DOI: 10.18727/0722-6691/5062

End-to-End Operations in the ELT Era

Olivier R. Hainaut[1]
Thomas Bierwirth[1]
Stéphane Brillant[1]
Steffen Mieske[1]
Ferdinando Patat[1]
Marina Rejkuba[1]
Martino Romaniello[1]
Michael Sterzik[1]

[1] ESO

The Data Flow System is the infrastructure on which Very Large Telescope (VLT) observations are performed at the Observatory, before and after the observations themselves take place. Since its original conception in the late 1990s, it has evolved to accommodate new observing modes and new instruments on La Silla and Paranal. Several updates and upgrades are needed to overcome its obsolescence and to integrate requirements from the new instruments from the community and, of course, from ESO's Extremely Large Telescope (ELT), which will be integrated into Paranal's operations. We describe the end-to-end operations and the resulting roadmap guiding their further development.

The origins

At the time of the construction of the VLT in the late 1990s, it was already becoming clear that the "classical observing cycle of trekking to the telescope and returning home with a tape of unprocessed data", as Quinn (1996) described it, was no longer the most efficient way to deal with the coming generation of 8- to 10-metre telescopes. The first paradigm change introduced by ESO was Service Mode (SM), in which pre-defined observations are executed when the observing conditions match those needed for the specific science case. The second was to guarantee the calibration of the instruments to a pre-defined level of accuracy and to allow archive science, implying a well-defined calibration plan and continuous quality control of the calibration process. Quinn also noted that a direct consequence of this was that the "flow of data must be managed from start to end of the observation process", resulting in the development of the Data Flow System (DFS; Quinn 1996). It was designed between 1994 and 1995, and deployed on the NTT between 1996 and 1997 during the "NTT Big Bang", when that telescope was overhauled to become a prototype of the new VLT standard.

The main building blocks of the DFS were already implemented at that time as a series of interconnected but standalone processes and tools: proposal handling; observation handling; maintaining the science archive; and data processing pipelines. One of the key concepts at the core of the original DFS was the Observation Block (OB), the quantum unit that defines all the information required to execute an independent set of observations on a specific target. Over time, the system has become more complex as new types of observing programmes have been added, including new modes and new concepts on top of the individual OBs (for example, containers in the ESO Public Surveys). Many tools underwent substantial changes and upgrades to cope with these new concepts and technical requirements, while others grew more organically.

The DFS gradually evolved to reach the current implementation (Figure 1). The core concepts of the system are robust (for example, the quantum unit of the OB to define observations) and a set of interfaces ties the tools together. Nevertheless, in many areas the tools have grown beyond a level at which they can be efficiently maintained, are based on aging technologies, and/or don't take advantage of new technical capabilities. Furthermore, the original concept of data flow has been broadened to include not only the data themselves, but also the whole operation process, from the original proposal for observing time to the distribution of science-ready data packages — i.e., an "end-to-end operations" process. Last, but not least, ESO is building the ELT, and one of its top-level requirements is that its operations must be fully integrated with Paranal.

The time has come for an in-depth review of ESO's Data Flow System, and a (r)evolution.

Evaluation of the end-to-end process

In addition to being the backbone of the operations of about 20 instruments on 12 telescopes over two sites, the DFS provides a rather homogenous interface for users (including both the community and internal users). The system must be sufficiently versatile and flexible to cope with this complexity and to incorporate new and continuously evolving instruments. In particular, the DFS will have to efficiently support ELT operations. Furthermore, sub-processes and the corresponding tools must pass information from one stage to the next one effectively, avoiding duplication and averting inconsistencies.

Besides these top-level requirements, the evaluation of the DFS incorporates the ELT top-level requirements, and the requirements derived from the future instruments (in particular from the Echelle SPectrograph for Rocky Exoplanet and Stable Spectroscopic Observations [ESPRESSO], the 4-metre Multi-Object Spectroscopic Telescope [4MOST] and the Multi Object Optical and Near-infrared Spectrograph [MOONS]). It also includes the results of a review of the current system, summarised in a series of Messenger articles (Primas et al., 2014 for SM; Romaniello et al., 2015 for the Archive, Arnaboldi et al., 2014 for Phase 3; Sterzik et al., 2015 and Patat et al., 2017 for scientific return). The feedback from the community has been assimilated through the ESO2020 workshop and poll (Primas et al., 2015) and the recommendations of the ESO2020 Time Allocation and Science Data Management working groups, as well as input from ESO's advisory bodies: the Scientific Technical Committee and the Users Committee.

The tools and facilities offered in other observatories were also examined to estimate the evolving astronomical environment and the possibility for their re-use in ESO's tools. Finally, the need to ensure that the DFS could be maintained over an expected lifetime on the order of 20 years guided any technological choices.

All the new requirements were assessed; while some are indeed novel, they are all compatible with the current fundamental concepts — even if not with the current

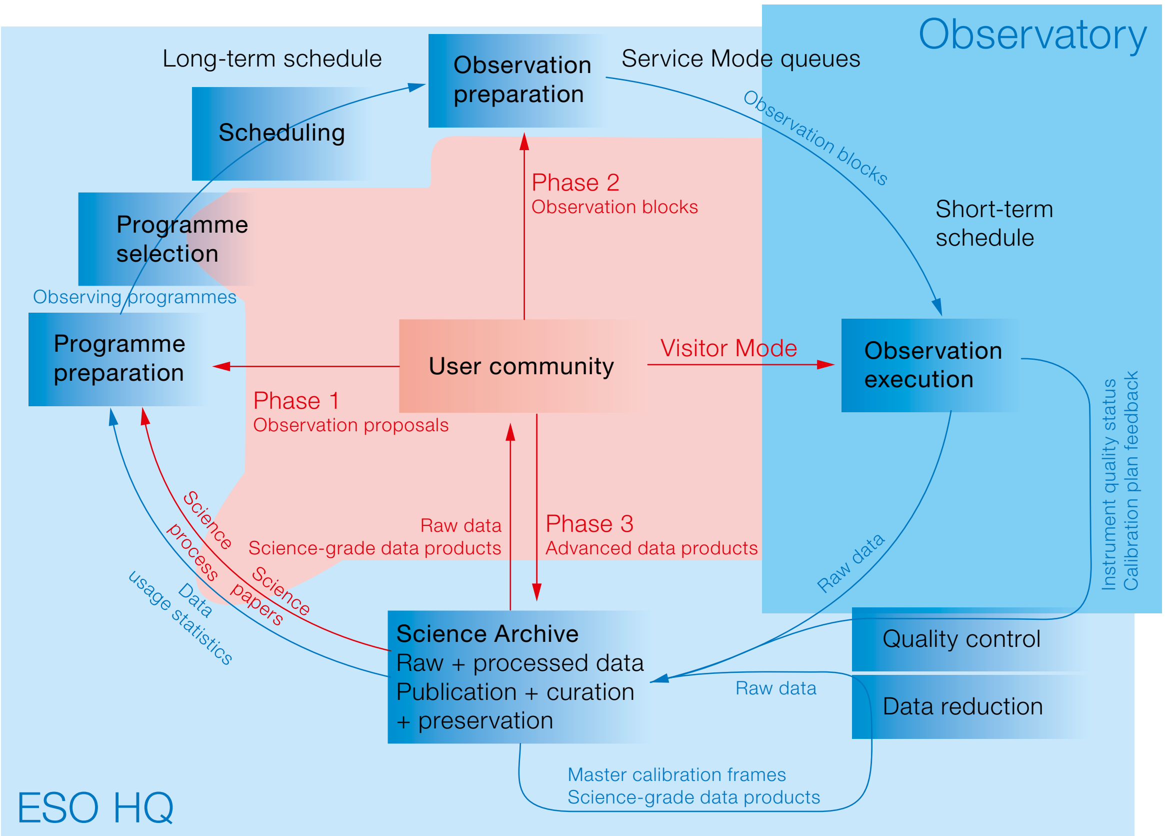

Figure 1. End-to-end operation process: the main sub-processes are indicated in boxes; the background colour indicates the main actors involved in each step. The arrows are labelled with the information flowing over them. This "DFS Wheel" is an adaptation of the original one presented by Quinn (1996). While the overall process will not change much over the deployment of the upgrades described in this roadmap, many of the sub-processes and the tools implementing them will undergo in-depth, and often complete, re-implementation.

tools — implying that the overall end-to-end operations, which has been in use since the start of VLT operations, is still sound. What is needed is an in-depth "clean-up" and update of the system rather than a complete re-engineering. As no exotic or extreme new requirement was identified, we did not investigate different operation schemes in detail. The new end-to-end operations model will be a modernisation of the current one, building on the many years of experience operating the VLT and, before that, La Silla.

All sub-processes were assessed, and areas requiring work were identified. An overall high-level plan was designed and broken down into individual projects. The dependencies between these projects were accounted for, together with specific milestones (for example, the first light of a given instrument) as well as the savings that they will generate in terms of maintenance and operations. The resources in time, equipment and staffing were evaluated for each project. Some of these projects were ongoing at the time of this exercise, so they were already collecting requirements and evaluating costs. Others will take place in several years' time; these have been estimated from the complexity of the existing system and the additional requirements. The preliminary estimates will be refined when the projects actually start — in the meantime they are sufficient for planning the roadmap presented here.

In the following section, we offer an overview of these projects and, by way of illustration, delve into some of them in more detail.

Roadmap towards a new data flow system

Platform

Breaking with a long tradition of requiring users to install and deploy programs on their own computers, new tools are being developed using application programming interfaces (APIs) to services hosted at ESO. Web-based user interfaces and user- or science case-specific scripts and even desktop applications can be built against these APIs. The strict separation of user interface and business logic has enormous advantages in terms of maintainability. Rather than asking users to download and install new software versions ESO can seamlessly fix bugs and roll out new features by deploying new software to an ESO server for immediate availability to all users.

The paradigm chosen is the representational state transfer (REST), an industry standard that the contracted software engineers are familiar with. Additionally, there are now powerful web frameworks that enable fast development of powerful, large-scale user interfaces. The framework chosen is Google's Angular 5. User interfaces tend to age faster than the underlying business logic. Thanks to their separation, interfaces can be updated without requiring a re-implementation of the business logic. Also, as the APIs are exposed (and documented), power users can access them directly using scripts to execute series of commands that would be cumbersome from a web interface; indeed, users can even develop their own custom interfaces.

Before the observations

Phase 1 includes the release of details of what is being offered in a particular observing cycle through a Call for Proposals, the preparation of proposals for observing time by the community, the handling of these proposals, and the organisation of and support to the Observing Programme Committee (OPC), which reviews the observing proposals submitted. The Principal Investigators (PIs) and co-investigators (Co-Is) of successful proposals then prepare the detailed specifications of their observations — the OBs — during Phase 2. Both Phase 1 and Phase 2 activities are supported by exposure time calculators (ETCs).

Following the plan, all tools related to Phase 1 and 2 and the ETCs will be modernised through the implementation of REST APIs and web interfaces. These projects have started and are at different stages of implementation.

Phase 2

For operational reasons, the Phase 2 systems were the first to be upgraded. They had already been updated to support survey operations on the Visible and Infrared Survey Telescope for Astronomy (VISTA) and the VLT Survey Telescope (VST) in 2010 and 2011, with the addition of scheduling containers. However, it was already recognised in 2014 that the current system would not be able to respond to the requirements from ESPRESSO's planet-hunting strategy and massive spectroscopic surveys with 4MOST. The Java desktop application called P2PP (Phase 2 Proposal Preparation) was introduced in 1997, replacing the original Tcl/Tk application. It is now being replaced by the new API-based P2[1], which reproduces its functionalities while adding new ones. With P2, the user creates OBs and stores them directly on the Garching repository. A bi-directional database replication ensures that the OB's content is the same in both the Garching and the Paranal repositories.

The P2 project was started in 2016; the P2 web tool and the matching visitor Observing Tool (vOT version 4) have been deployed for Visitor Mode support during 2017. The deployment for Service Mode will be done in two phases in 2018, first on Unit Telescope 2 (UT2) and on the survey telescopes for Period 101, followed by implementation on all Paranal telescopes for Service Mode for P102. A second wave of developments that will include specific VLTI and Adaptive Optics operational requirements, enable definition of complex observing strategies through nested scheduling containers, and also enable more streamlined connections with the Phase 1 ETCs will take place in 2019. The inclusion of the La Silla operations in the new P2 system will take place in 2018-2019, in time for the new instruments NIRPS and SOX. Once this is done, we will be able to decommission the legacy P2PP and underlying servers.

Phase 1

The Phase 1 support tools include the long-lived ESOFORM, a LaTeX package that has been used to define proposals for observing time since Period 60 (we are now in Period 100). However, the ESOFORM is only the tip of the iceberg of Phase 1 tools. A series of tools process the LaTeX files of the observing proposals and store some of the information in a database (unfortunately the LaTeX format means that some important information cannot be stored in such a way that can be easily retrieved programmatically), deal with the administration of the OPC and the OPC reviewer rankings and recommendations, interface with the scheduling system, and finally handle the communication to the proposers.

The maintenance of these tools has become cumbersome, and integrating new requirements (for example, the recommendations from the ESO2020 Time Allocation working group) would be impractical. These tools and the underlying databases have grown organically. In terms of technology and architecture, they require a complete redesign and implementation. Because of the tight integration of all of the Phase 1-related subsystems, and because of the fundamental changes in the underlying infrastructure, the change from the old to the new system can only take place when the whole system is ready.

Currently, the first two modules — which define what is offered in an observing cycle and how to define and submit proposals — are under development and are expected to be delivered in the first half of 2018. The next modules (OPC administration, handling OPC rankings, interactions with the submitters) will be developed in 2018 and 2019. The overall integration will take place in 2020 and we expect to have the new system in place by the end of 2020. While it would be nice to deploy the new proposal submission system earlier, its deeply rooted interconnections with the rest of the Phase 1 infrastructure make that impossible. From the user point of view, the proposal submission will be fully web-based, supporting collaborative work from the PI and Co-Is. Figure 2 shows screenshots of the P1 and P2 web user interfaces.

Exposure time calculators (ETCs)

The ETCs are already web-based, but implemented in a way that precludes easy interaction via scripts (or from other tools like Phase 1 and Phase 2 preparation) and that uses specific models for the instruments. This makes their maintenance cumbersome and their integration and use in batch mode difficult.
The ETC upgrade project was initiated in 2017 and will likely be completed by 2020. Some features of the user interfaces at other observatories, such as the possibility to save and share ETC sessions, are included in the requirements. Thanks to the modular structure of the ETCs, the benefits of the new system will appear gradually on the ETC web pages.

The new P1, P2 and ETC systems share exactly the same authoritative instrument definition model as that used for the instrument itself (i.e., the Instrument Packages), ensuring consistency between the tools. The tools will also be able to interact and exchange information using their APIs. This will make a tight integration possible between the ETCs and the P1 and P2 systems; all the calculations performed to estimate the telescope time required for a Phase 1 proposal are preserved in the proposal itself, and will be available as the basis on which to optimise the observing strategies at Phase 2. The ETC API will also make simulations available for Quality Control processes (see the After the observations

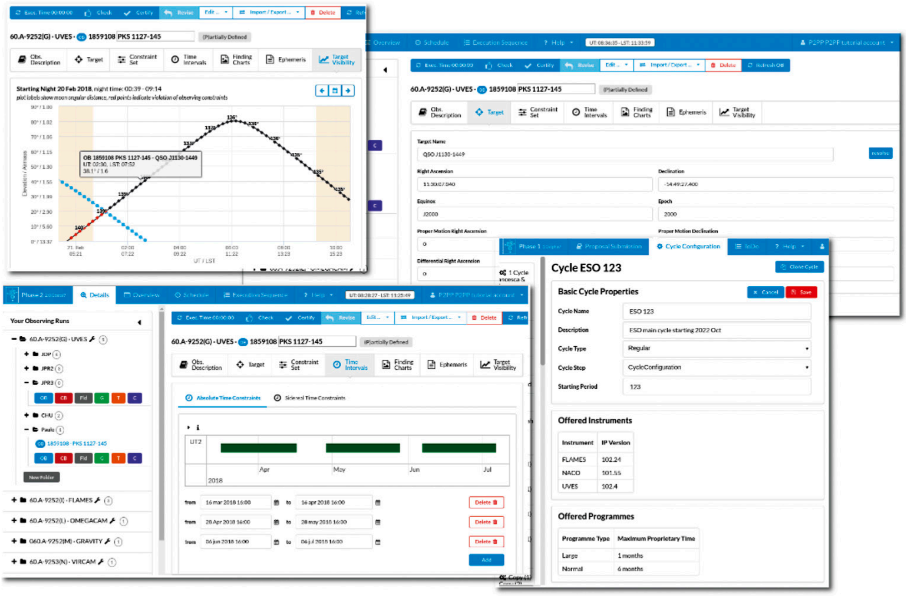

Figure 2. Screenshot of the P1 and P2 web tools developed to prepare proposals and observations.

section), and to investigators who want to explore a broad range of parameters or apply a simulation to a catalogue of objects.

Towards a unified GuideCam tool

OBs for several instruments must include very specific details on the instrument configuration (for example, fibre positions on a multi-object spectrograph). These are currently defined using a suite of auxiliary preparation tools (for example, FORS Instrument Mask Simulator [FIMS], KMOS ARM Allocator [KARMA]), many of which are based on the increasingly outdated Tcl/Tk technology. As these tools are too complex for implementation in current web technologies, we developed a platform and implemented it as a plug-in to the Aladin sky atlas tool[2]. This becomes the basis for all future observation preparation tools. This new tool, GuideCam[3], is already available for various instruments (Figure 3). Existing preparation tools for other instruments that will not be decommissioned in the foreseeable future will be retrofitted to GuideCam over the coming years. In 2018, GuideCam will also include a generic system to produce finding charts, to replace the aged SkyCat tool; this will make it much simpler to prepare Service-Mode-compliant finding charts. Furthermore, thanks to the integration between GuideCam and the P2 system, information will be embedded in finding charts and will be available to the staff reviewing and executing the observations.

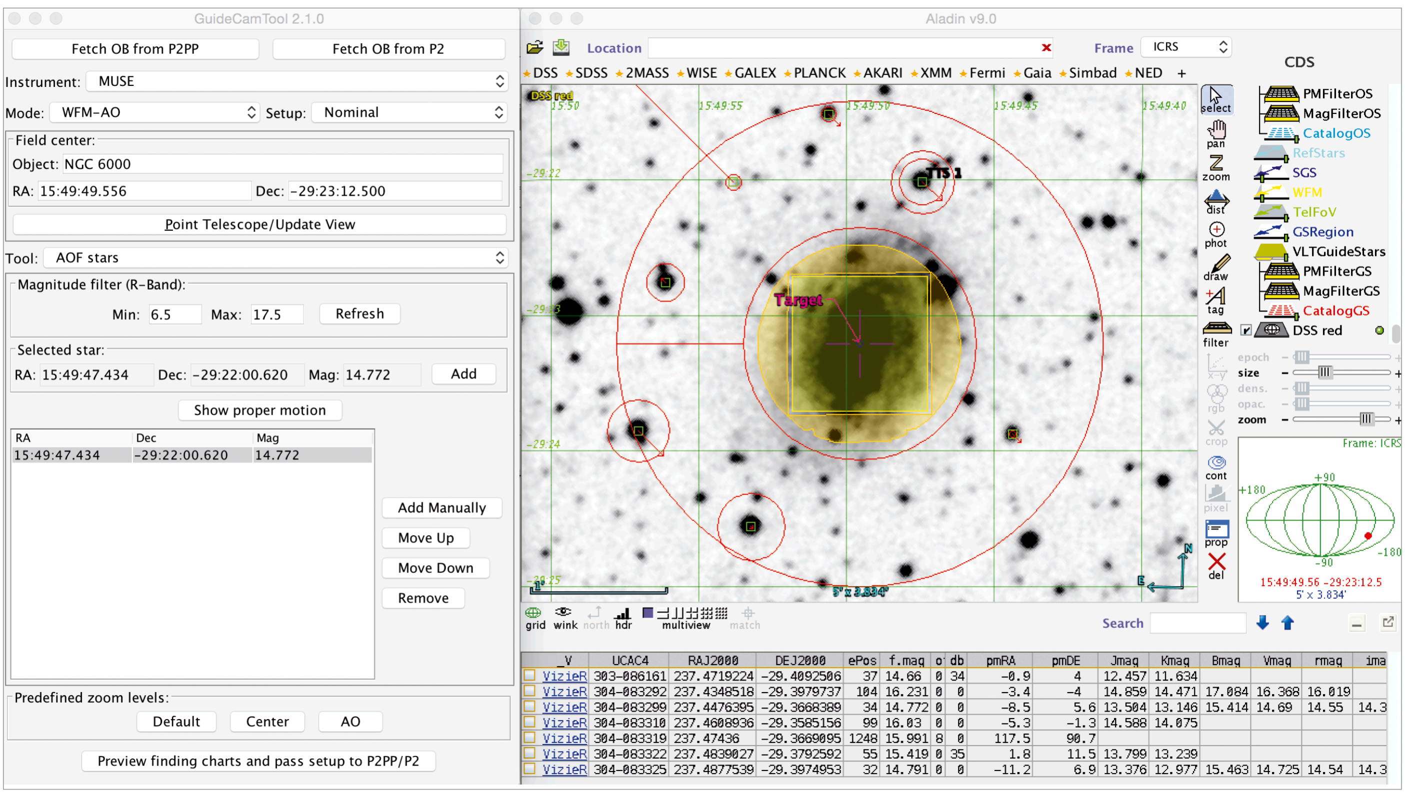

Figure 3. Screenshot of the Unified GuideCam Tool.

Scheduling the observations

Long-term telescope schedules are assembled by accounting for a number of factors, including the ranking of Phase 1 proposals by the OPC, ESO high-level science operation policies[4] and any constraints specified in the proposals themselves — both deterministic (for example, target coordinates and the phase of the Moon) and statistical (such as image quality). The output of the new P1 system will enable a much finer granular accounting of the constraints and geometry than the current system, ensuring that all of the information relevant for scheduling and observations execution is fully captured.

The scheduling tool that is currently deployed treats VLT units separately and the only case requiring coordinated allocations on multiple Unit Telescopes (i.e., the VLTI) has been managed manually, within pre-allocated slots. The start of ESPRESSO operations, with its capability of operating in Service Mode at any UT during the night, poses a new challenge, which can only be addressed by a parallel approach to VLT scheduling. This requires a radical change to the scheduling paradigm.

Therefore, we also plan to update the current Telescope allocation Tool (TaToo). The new scheduling tool will enable complex strategies, such as simultaneous or sequential scheduling on multiple telescopes, therefore addressing operational requirements for ESPRESSO. This project is expected to start in early 2019, once the later phases of the P1 project are underway.

Observations

Observations on the mountain are either performed in Visitor Mode, for which the observer travels to La Silla or Paranal to execute and fine-tune their observations in real time, or in Service Mode. For a Service Mode run, the observer selects the observation to be performed from the pool of Phase 2 OBs, executes it and evaluates its success based on fulfilment of predefined observing parameters. The selection of the "best" OB for execution is assisted by the Short-Term Scheduler (STS), also called the OB Ranking Engine (ORANG), which accounts for policies and user-specified priorities, as well as for the prevailing conditions. ORANG also takes into account additional constraints such as time criticality and links between OBs in OB containers. However, ORANG has no capabilities to account for the evolution of the observing conditions. We are working with the Paranal astro-weather group to integrate information provided by the Astronomical Site Monitor that goes beyond established weather parameters such as cloud transparency and seeing, including water vapour pressure, turbulence profile, etc.

Furthermore, ESO has embarked on several investigations with weather forecasting institutions (both academic and commercial) to refine the existing forecast to suit astronomical needs. The goal is to be able to forecast the main parameters that are relevant to the observations — and, even more importantly, how they change — with a precision that can assist in the selection of future OBs. We will likely review and upgrade the STS to account for astro-weather forecasts, with the ultimate goal of forecasting the seeing conditions and the turbulence profile, as these are particularly critical for the operation of the AO-assisted instruments — especially on the ELT. In parallel with this functional upgrade, a complete re-implementation of the underlying Observation Tool (OT) is required by its outdated code base. While the preliminary studies are already ongoing, the project itself will be launched only in 2020; we hope to have the astro-weather forecast module in place in 2022, and the full OT/STS deployed in 2023.

In addition to these visible changes, a series of behind-the-scenes projects need to take place. For instance, the Data Handling System (which moves data between workstations at the Observatory before bringing them to the Science Archive) has been maintained since its creation, but has never been re-evaluated at the overall process level. The integration of the requirements from the ELT, as well as from the Quality Control project (see below) is a good opportunity. This will take place in parallel with the implementation of the interfaces to the ELT between 2021 and 2023, including both "front end" (from the DFS to the ELT control system) and "back end" (from the ELT back to the DFS).

After the observations

Science Archive

The Science Archive System has always been considered the final repository of VLT data, thus preserving their legacy value. Over time, its role has become more and more central in overall operations. Thanks to advances in the internet, its role now also includes the distribution of data to both the original PI as well as to other independent archive researchers. Teams leading Public Surveys and Large Programmes deliver large, consistent, science-ready datasets called the External Data Products to the Archive via the Phase 3 process (Arnaboldi et al., 2014). In parallel, the development of science-grade pipelines has enabled ESO to systematically process the data from a (growing) number of instruments, creating the Internal Data Products (Romaniello et al., 2016). Both raw data and science-grade data products are very popular with the community, with the number of users and the request rate steadily increasing (Romaniello et al., 2016).

However, the core archive services — the ensemble of web interfaces and underlying system — have not evolved much since their deployment in the late 1990s, while web and database technologies have flourished and the Virtual Observatory (VO) protocols have matured. A project to re-implement the Archive Services is underway. Its goal is to let the users search for and discover ESO data, taking advantage of the detailed and consistent metadata describing them, either through an interactive web interface, or through a VO-compliant API for more complex queries. The API will also open the archive to the available VO tools (for example, Aladin, TopCat, etc). Finally, the Archive will offer previews of its assets, allowing users to quickly evaluate the suitability of a file for their purposes. The previews include progressive multi-scale images (c.f. Google Earth), which will be broadcast via the HiPS network[5]. Thanks to this, services like ESASky[6], which gives access to a highly curated subset of multi-wavelength/multi-observatory data, will also be able to access, display and retrieve suitable ESO data – for instance, data products from large surveys. The first release of this new Archive Service, which will include most of the data products, is scheduled around the time of publication of the present article and will be described in detail in the next issue of the Messenger.

Quality control (QC)

One of the key concepts of the original DFS was the introduction of a formal calibration plan ensuring that instruments can be calibrated to a pre-defined accuracy; the operations plan ensures that all the required ancillary frames are acquired and the quality control (QC) system verifies their validity and suitability through instrument health and trending parameters. A QC infrastructure has been developed over the years to monitor the instruments and detect deviations from their specifications before the instrument performance significantly degrades. The QC also produces certified calibration frames, which are stored in the Archive and can be used by the pipelines for the science data processing.

This infrastructure is robust and flexible, but uses a technology that has issues in terms of maintainability. Furthermore, we are now in a situation where the scope of QC can be expanded beyond the calibration and instrument stability. A QC infrastructure can be deployed directly at the telescope, together with the online pipelines, to provide a systematic real-time assessment of many parameters of the data as it is acquired. This will help to evaluate observations from new complex instruments, for which raw data are so entangled that an inspection of raw frames would not be sufficient to judge their quality (for example, integral field units and interferometric instruments). Furthermore, the pipeline generation of science-grade processed data implies that the data quality is also evaluated in a systematic way. Finally, the traditional instrument calibration QC produces parameters that can be compared directly with the output of the instrument simulators, closing the loop between the ETCs and the actual instrument. Thanks to the ETC API, deviations between the simulation and the measurement can result in flagging a problem, or in updating the ETC parameters. This major evolution of the QC processes constitutes a significant effort and the project is due to start in mid-2018, with new QC systems deployed gradually until 2022.

Data processing and pipelines infrastructure

The QC system, the online pipelines (at Paranal) and offline pipelines (on users' machines) rely on a series of data organisation tools (for example, to select and associate suitable calibration files) and data processing infrastructure that calls and manages pipeline recipes (for instance, ESOrex, and the workflow system Reflex). The pipeline recipes themselves are implemented using the low-level Common Pipeline Library, and more advanced algorithms from the High-level Data Reduction Library. The suitability of these tools will be reviewed in the coming years accounting for new internal and community requirements (for example, interfacing with Python) and new capabilities (for example, cloud computing). The outcome of this review will lead to an evolution of pipeline systems and their infrastructure. Whilst the scope and nature of this project are not defined yet, resources have been earmarked between 2020 and 2021.

Infrastructure

The DFS relies on a series of infrastructures, which must be maintained and allowed to evolve. A couple of examples are given below.

The first is the Next Generation Archive System (NGAS), the storage technology developed for the Science Archive. NGAS has evolved over the past decade, and is now, either directly or indirectly, used at ESO, ALMA and other institutions. As storage technologies evolve rapidly, NGAS must also follow.

The DFS also depends on many databases; while our demands are fairly modest by modern standards, we have original requirements (such as our multi-site architecture, or the spherical geometry of the celestial sphere), which are not fully supported by out-of-the-box products. Furthermore, the landscape of available database systems is evolving in terms of capabilities, support and cost.

Conclusions

An ambitious roadmap has been developed to overhaul the dataflow system supporting the VLT observations, the ultimate aim being to accommodate the requirements of ELT and new VLT instruments, and to ensure the system's future maintainability. This plan has been developed and endorsed internally at ESO. An external review is taking place to ensure its completeness, review its soundness and evaluate synergies with similar developments at other observatories. The developments are being staged to maximise the use of finite resources, while meeting deadlines. One of the main goals is to integrate all of the sub-processes and make the overall end-to-end operations seamless, both for the users and for the operators. While many new tools and systems will be deployed as soon as they become available, some require major infrastructure changes and must therefore be completed before they can enter operations. The ultimate deadline for this series of projects is ELT first light.

References

Arnaboldi, M. et al. 2014, The Messenger, 156, 24
de Zeeuw, T. 2016, The Messenger, 166, 2
Patat, F. et al. 2017, The Messenger, 170, 51
Primas, F. et al. 2014, The Messenger, 158, 8
Primas, F. et al. 2015, The Messenger, 161, 6
Quinn, P. 1996, The Messenger, 84, 30
Romaniello, M. et al. 2016, The Messenger, 163, 5
Sterzik, M. et al. 2015, The Messenger, 162, 2

Links

[1] Phase 2 demo environment: https://www.eso.org/p2demo
[2] Aladin sky atlas: aladin.u-strasbg.fr
[3] The unified GUideCam Tool: www.eso.org/sci/observing/phase2/SMGuidelines/GUCT.generic.html
[4] VLT/VLTI Science Operations Policy: eso.org/sci/observing/policies/Cou996-rev.pdf
[5] Hierarchical Progressive Surveys (HiPS): aladin.u-strasbg.fr/hips
[6] ESASky: sky.esa.int