

DATA FLOW SYSTEM OF THE EUROPEAN SOUTHERN OBSERVATORY

2005 COMPUTERWORLD HONORS CASE STUDY

SCIENCE

ESO HAS REVOLUTIONIZED THE OPERATIONS OF GROUND-BASED ASTRONOMICAL OBSERVATORIES WITH A NEW END-TO-END DATA FLOW SYSTEM, DESIGNED TO IMPROVE THE TRANSMISSION AND MANAGEMENT OF ASTRONOMICAL OBSERVATIONS AND DATA OVER TRANSCONTINENTAL DISTANCES. [20055322]



SUMMARY

For the efficient management of astronomical observations and data over transcontinental distances, the European Southern Observatory has implemented an end-to-end data flow system. Using this system, ESO has revolutionized the operations of ground-based astronomical observatories.

APPLICATION

Research in astronomy provides several societal benefits. Easily observed astronomical events have formed the basis for time keeping, navigation, and the debunking of myths in cultures around the world. The beauty of the night sky and its rhythms are at once stunning and compelling. The boldness of our collective efforts to comprehend the universe inspires us, while the dimensions of space and time humble us. Astronomy encompasses the full range of natural phenomena from the physics of invisible elementary particles, to the nature of space and time, to biology, thus providing a powerful framework for illustrating the unity of natural phenomena and the evolution of scientific paradigms to explain them.

The European Southern Observatory (<http://www.eso.org>) is an intergovernmental research organization consisting of 11 European countries created to construct and operate astronomical observatories in the Republic of Chile. The coupled goals of the efficiently using telescope time and building a scientifically useful data archive led to the development of the ESO Data Flow System (DFS). The DFS is a collection of tools and processes for managing the end-to-end information flow involved in planning and executing astronomical observations and the storage of the observational data in the ESO science archive. At its lowest layer, the DFS relies heavily on enterprise-class relational database management system (RDMS) technology as databases are used as a persistent state for most of the objects and information flowing through the system. At its higher layers, it uses a combination of Web-based tools and client-server applications implemented in Java and C++. In particular, robust and instantaneous database replication at transcontinental distances is critical for reacting rapidly to changing conditions or to address operational problems at the observatories.

The Chilean observatories are located in the Atacama Desert, one of the clearest and driest places on the Earth, making it ideal for astronomical observations. Nevertheless, atmospheric conditions vary on an hour-by-hour basis in an unpredictable manner. When physically present at one of the ESO telescopes, astronomers are often confronted with the problem of rapid weather variations or weather that is simply not good enough to execute their observations. These phenomena are often called the weather lottery – sometimes you win, sometimes you lose. Losing this lottery is both frustrating and inefficient - when observations cannot be completed, astronomers must often wait for a year or more before they can attempt their observations again. There is an additional travel overhead: it typically takes two days each way to travel back and forth to the Chilean observatories from Europe. Losing the weather lottery means that those four days have been wasted.

To mitigate this problem, ESO has adopted an approach known as service observing. This approach has been used since 1997 at most ESO telescopes, including the Very Large Telescope (VLT), a cluster of four 8.2m telescopes on Cerro Paranal which is the world's largest optical astronomy observatory. With this approach, astronomers use software tools to completely design their observing programs at their home institutions. These program descriptions are then submitted over the Internet to ESO Headquarters, located in a suburb of Munich, Germany. After a quality control check, these descriptions are transmitted to Chile and stored there in local databases. When weather conditions are right, the observations are then executed. The data generated are then sent back to ESO Headquarters. After

Robert Carrigan,
Chairman of the Chairmen's Committee

Ron Milton,
Vice-Chairman of the Chairmen's
Committee

Dan Morrow,
Chief Historian

suitable processing and a quality assurance check, the data are forwarded to the originating astronomer, completing the end-to-end service observing cycle. In this system, there are no weather lottery losers – and the four travel days per user are eliminated as well. The European astronomical community has been extraordinarily enthusiastic about this new approach – in a typical year, ESO executes the scientific programs for more than 1000 users. Following the lead of ESO, other ground-based observatories are now implementing service-observing systems, all of which rely upon on accurate data flow.

The savings are staggering. ESO supports about 1,000 user sessions per year. At the moment about 500 are in service mode (astronomers stay at home). This translates to about 1,000 round trip airfares (plus ground transport plus accommodation and supplies) being saved. That's about 2 million Euro or US \$2.6 million. Furthermore, by asking for service mode observations it is possible to perform science programs that demand rare conditions. For example, a project may need conditions that are in the top 10% of the clear nights. This means if an astronomer travels to the telescope for 10 nights of observing in a year, he/she only gets one night per year of the conditions required. At that rate, the astronomer's scientific program will require 10 years to complete. With service observing, the project will run only on those nights in the top 10% of conditions. Since there are about 30 nights per year in the top 10% of conditions, the astronomer can finish the work in less than one year. This automated scheduling of required conditions both improves the outcomes of the scientific endeavors and speeds the project completion timelines.

ESO is equally concerned about improving the productivity of the astronomers of tomorrow. At most ground-based observatories, astronomical data are only accessible by the original investigator. Since the early 1990s, ESO has been archiving all astronomical and relevant engineering data collected at its facilities. After a one-year waiting period to allow the original investigators enough time to complete their own research, archive data is available to any qualified researcher in the ESO community. In 2005, the ESO archive will be opened to the entire worldwide research community, increasing the value of the archive. By facilitating data re-use, using DFS, the value of the original data is multiplied. The ESO archive currently contains approximately 30 Terabytes, but will grow into a Petabyte class archive by the end of the decade. Indeed, the annual data intake is expected to exceed 0.5 TB per day by 2010. Accessing and using such large datasets in an efficient manner will require new approaches and new technology, much of which will be based on DFS.

BENEFITS

For the ESO user community (ESO's primary customers), the benefits are clear. Observational programs can be designed and submitted in the comfort of a user's home institution, without traveling to the observatory. This leads to fewer operational mistakes and therefore more efficient data acquisition. Observations are executed under the most appropriate weather conditions – the weather lottery is eliminated and substantial financial savings are realized. This also means that demanding science programs can also make use of large amounts of relatively rare and excellent atmospheric conditions. The stress and inefficiency of transcontinental travel is eliminated, and millions of dollars are saved. The user receives processed data that has met well-defined quality assurance criteria and are ready for scientific analysis. Finally, users are supported by a team of ESO astronomers who are experts in all aspects of DFS operations. The result has been a significant increase in the scientific productivity of the ESO user community. As measured by the number of papers in peer-reviewed journals, ESO has moved ahead of its rival facilities and is now one of the leading astronomical facilities in the world, second only to the Hubble Space Telescope. Coupled with cutting edge optical telescopes and astronomical instruments at the Chile sites, the DFS has contributed to this success by providing the fundamental IT infrastructure for observation and data management.

These benefits have allowed ESO users to explore some of the key astronomical questions in the most efficient manner possible – including the formation and origin of planets around other stars, the origin of fundamental atomic elements in the early universe, and the mysterious dark matter that seems to fill and bind the entire universe. Researchers were relying heavily on the DFS when they discovered that the expansion of the universe seems to have accelerated in the early universe driven by the equally mysterious dark energy.

For ESO itself, the development of the DFS has also had tangible benefits including improved user satisfaction, improved scientific return on capital and operational investment, and improved proficiency with cutting-edge IT technologies. In the area of science archives, for example, ESO is now the world-leader in ground-based astronomy and has positioned itself to play a major leadership role in the development of the Virtual Observatory (VO). The VO is a worldwide collaboration to connect major

astronomical science archives through the use of new interoperability protocols and Grid technology. Proto-VO demonstration projects have already resulted in the detection of heretofore-unknown examples of exotic astronomical objects.

IMPORTANCE

Without the use of an IT-based strategy within the DFS, it would have been impossible to overcome the weather lottery problem and promote data re-use without unacceptably high personnel costs for ESO and tremendous paperwork burdens for the ESO user community. Although many of the DFS paradigms had already been explored and deployed by space observatories such as the Hubble Space Telescope, ESO was the first ground-based observatory to rely heavily on an IT strategy for its science operations. Furthermore, ESO has had to extend previously developed paradigms to cope with weather related problems in science program scheduling and execution.

Fundamentally, the ESO DFS relies on three key IT archetypes: Web-based information management tools, client-server architectures, and enterprise-class relational database management systems. In many ways, ESO is no different from any other round-the-clock service provider dealing with a geographically distributed customer base.

Web-based interfaces are used to provide information and tools to users, including information on how to submit observing programs, status of current observing programs, and tools for observing program planning. Web-based interfaces are also used to gather information from users, including user contact information and observing program proposals.

The main user tool for developing and submitting observing programs is a classic Java-based client running on the astronomer's desktop and exchanging information with a server located at the ESO Headquarters in Germany. Once completed, user programs are managed and executed by ESO operations staff using another Java-based tool. Not only does that tool provide database-browsing functionality to select the most appropriate program for execution under the current weather conditions, but it also exchanges with the observatory real-time system the information necessary to execute the observation.

Automatic on-line processing checks the quality of the observations and monitors the health of the instruments. This software subsystem consists of rule-based application for automatic data organization, an execution framework supporting parallelism and C components implemented as plug-ins that manage number crunching. The Quality Control information resulting from this processing is stored in the database and is used as a health indicator of the system.

An enterprise-class RDMS is used to operate and synchronize the operational databases in Germany and Chile. In particular, replication technology is absolutely critical to managing the constant round-the-clock, two-way information flow across multiple time zones at trans-continental distances. Database administration and content management support is provided at all ESO operations sites.

With the rapid growth in data volume, ESO is also moving rapidly to cluster computing technology. By 2010, ESO must have the capacity to store and process close to 1 TB of science data per day from a heterogeneous set of astronomical facilities. Cluster-based solutions for both data storage and data processing are being developed in-house.

ORIGINALITY

Traditionally, ground based astronomical observatories have been used as facilities where scientists apply for observing time, eventually travel to the remote sites where telescopes are located, carry out their observations by themselves and finally take their data back to their home institutes to do the final scientific analysis. As observatories become more complex and located in ever more remote locations (to reduce light pollution), this operational concept (coupled with the weather lottery effect) becomes less and less effective. In particular, the lack of data re-use has been increasingly seen as scientifically unproductive. Such thoughts, together with the experience made with space-based observatories such as the Hubble Space Telescope (HST), guided the design and implementation of the ESO Data Flow System. The DFS allows both traditional on-site observing as well as service observing, where data is collected by observatory staff on behalf of the ESO user community based on user submitted descriptions and requirements. In either case, the data is captured by DFS and saved in the ESO science archive. After a one-year proprietary period during which the original investigators have private access to their data, suitably qualified researchers can access the data for their own use. ESO was the first

ground-based observatory to implement these operational concepts and tools with a complete system. It was also the first ground-based observatory to build and maintain such an extensive science archive that does not only contain observational data, but also auxiliary information describing the operation process. In both areas, ESO remains the world-leader in end-to-end observatory operations on the ground.

Over time, the DFS has evolved to meet new operational scenarios, the requirements of new instruments and to profit from new technologies.

The dataflow system supports the following services:

- Observing proposal preparation by the ESO user community
- Observing proposal peer review and scheduling
- Observation description creation and submission for successful proposals
- Observation scheduling based on user requirements and weather conditions
- Post-observation execution data capture and on-line archiving
- Data processing pipelines to support data quality assurance and to create science-ready data products
- Long-term on-line data archiving, with safe storage combined with rapid access
- Data retrieval and delivery to primary and archive researchers

These functions support the overall cradle-to-grave data flow and operations process in place at the ESO observatories.

SUCCESS

Since 1999, ESO has used the DFS to capture almost 4 million datasets totaling more than 30TB in size. All these data are available on-line with a maximum access time of about 30 seconds to any randomly selected dataset. In a collaboration with the Space Telescope European Coordinating Facility (ST-ECF) (funded by the European Space Agency) and the Space Telescope Science Institute (STScI) (funded by NASA), the ESO archive also hosts all data collected by the Hubble Space Telescope (HST). The ESO archive provides a homogeneous way to query and retrieve ESO and HST data through a combination of generic and specialized query mechanisms providing novice and expert users with access to all the parameters describing the archived data. Based on these interfaces, archive users request about three times more data in the course of a year than the amount added by new observations in the same period – in other words, the data re-use multiplier is a factor of three on average, and is much higher for certain high-value datasets.

The whole system has proven to be reliable and robust as no observational data has been lost since its implementation. The newly introduced observing concepts as well as system tools are well accepted by the astronomers. The tools and therefore the telescopes/instruments are understandable and easy to use, and do not require a deep knowledge of the different hardware and software. There are very few bug reports and the telescope downtime due to DFS software problems has so far been minimal.

DIFFICULTY

Putting together a data flow system for a wide variety of complex scientific instruments is one issue. Implementing such a system in a way that it can be used in a 24/7 operational environment with the highest possible availability and efficiency is a much larger challenge. Another difficulty comes from distributed nature of ESO science operations – with four sites (Cerro La Silla, Cerro Paranal, ESO/Santiago, ESO/Garching) on two continents (South America and Europe). Some data must flow as quickly as possible from one site to another in order to monitor science data quality and to react to special situations. This requires special implementation and maintenance of data flow paths in general and database replication in particular. Because of cost constraints, a full network-based data transfer cannot be implemented. Therefore, the bulk of the data is sent on either DVDs or magnetic disks from Chile to Germany where the main ESO archive is located.

The final challenge comes from the rapid growth in data complexity and volume. In 1997, the initial release of the DFS had to deal with two relative simple instruments at one site in Chile producing several gigabytes per night. At present, the DFS must handle 21 instruments with a broad range of data complexity at two sites in Chile producing several 10s of GB per night. Over the next three to five years, the system will manage more complex data with volumes of several 100s of GB per night. Constantly increasing data complexity requires constantly improving tools for observing planning, data processing,

and data quality assurance. Constantly increasing data volume means tracking the development of mass storage and processing capability, with an emphasis on Linux-based computing clusters to address both issues.