# From Chile to Europe in minutes: Handling the data stream from ESO's Paranal observatory

Martino Romaniello[a], Stefano Zampieri[a], Cecilia Cerón[b], Andrew Wright[b], Reinhard Hanuschik[a], Cédric Ledoux[b], Fernando Comerón[a]

[a]European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching bei München, Germany;

[b]European Southern Observatory, Alonso de Córdova 3107, Vitacura - Santiago, Chile

**ABSTRACT**

The ESO telescopes in Chile are operated in a geographically distributed scheme, in which some of the essential steps in the end-to-end observing chain take place in Europe. Most notably, the health status of the instruments as derived from the data themselves is monitored in Europe and the results fed back to the observatory within the hour. The flexibility of this scheme strongly depends on the speed with which the data stream produced by the telescopes can be sent to Europe for analysis and storage. The main challenge to achieve a fast intercontinental data transfer is the data volume itself, which currently reaches an average 25 GB/night (compressed) for the four VLT Unit Telescopes. Since mid 2008, this stream has been entirely transferred through the internet via a 4.56 Mbit/s bandwidth assured via a Quality of Service policy, which sufficed to transfer an average night of data within a few hours. A very recent enlargement of this capacity to 9.12 Mbit/s will soon allow the addition of the calibration data for VISTA, the new infrared survey telescope on Paranal, to the data stream transferred through the internet. Ultimately, the average data volume produced on Paranal once the visible VLT Survey Telescope (VST) and the full complement of second-generation VLT instruments becomes available is expected to exceed 200 GB/night (compressed). Transferring it over the internet will require a new fiber-based infrastructure currently under construction, as well as the use of additional high bandwidth channels. This infrastructure, provided by the European Union co-funded project EVALSO, should provide a data transfer capacity exceeding 1 Gbit/s that will allow the transfer to Europe of the entire Paranal data stream, as well as that of the nearby Observatory of Cerro Armazones and of the future European Extremely Large Telescope, with a delay of minutes at most since the data were taken.

**Keywords:** Information and Communications Technology, Quality Control, Observatory Operations

## 1. INTRODUCTION

The operation of world-class ground-based astronomical facilities is nowadays a global enterprise in several respects. Besides their observing time being made normally available to research teams throughout the world, operations processes such as time allocation, observation preparation, data processing and data archival can take place at locations thousands of kilometers away from the telescope. The centers where such processes take place thus become a vital part of the facility. Underpinning the capability of managing operations in such geographically distributed manner is the ability to transfer the vast amounts of data generated by current astronomical instrumentation across the operations sites in a reasonably short time, so that the outcome of the various processes composing the operations scheme of the facility can feed back on each other in a timely manner, with the ultimate goal of making available to researchers high quality datasets ready for analysis on the shortest possible timescale.

The remote location of those facilities usually implies that no advanced communications infrastructure is available in the proximities of the chosen site previous to their construction. Data transfer must either take place via physical storage media transferred among the sites by conventional means such as express mail, or via dedicated network infrastructure built for this purpose connecting the facility with a high bandwidth channel and all the way to the remote operations centers.

The ESO observatory of Cerro Paranal (Chile), hosting the Very Large Telescope (VLT), the VLT Interferometer (VLTI), the Visible and Infrared Survey Telescope for Astronomy (VISTA) and, in the future, the VLT Survey

Telescope (VST) has been no exception to these demands. The operations facilities on Paranal are complemented in the ESO Headquarters in Garching near Munich (Germany), where user support, data processing and permanent storage takes place. During the initial 10 years of operations, limitations of the available bandwidth coupled with the growing overall data rates coming from the increasing number of telescopes and instruments on the mountain forced ESO to transfer the data from Chile to Europe using physical media, with the subsequent time lag of up to two weeks between the generation of the data and their detailed quality analysis. Only recently it has become possible to use the Internet to transfer the data stream produced by the VLT within hours or less of it being produced. Recent upgrades have made possible both to increase the data transfer speed and, in the near future, to add the VISTA survey telescope (currently the facility with the highest data rate production on Paranal) to the data stream transferred through this channel. However, qualitative leaps lie ahead with the arrival to Paranal in the period 2009-2013 of the VLT 2$^{nd}$ generation instruments with multiple, large-format detectors, and of the VST. The data production rates of these new instruments are expected to exceed the volume of data that can be transferred over a 24 hours period beyond the capabilities of the current communications infrastructure. Furthermore, the communications infrastructure on Paranal will be used in the future to transfer the data obtained with the European Extremely Large Telescope (E-ELT), to be built in the nearby Cerro Armazones and expected to enter operations around 2019. In this paper we present the current and upcoming data volume panorama and the perspectives to deal with it up to the E-ELT era.

## 2. VLT END-TO-END OPERATIONS AND ITS COMMUNICATIONS NEEDS

Since the design of its most important elements in the mid 1990s[1], VLT operations were conceived as an integrated system[2] encompassing the whole process from proposal preparation and submission to data archival and reuse, a model strongly inspired by the operation of space-borne facilities. An important rationale for this model is the need to flexibly allocate a large fraction of the available time in the so-called Service Mode. Operations in this mode have become in practice the baseline operations mode of the VLT[3], with between 60% and 70% of the available time for scientific observations being scheduled in this way. In its current implementation by ESO, Service Mode users submit fully predefined observation specifications to the observatory, and execution takes place when the external constraints indicated by the users are met. Underlying this model is a system with tools to manage the flow of data and information, heavily relying on databases that are replicated between the operations sites. A central element in this scheme is the ESO archive, located at the ESO Headquarters in Garching, where all observations collected at the observatory are stored and made available to operations groups in charge of quality control, instrument trending and health check and data package preparation and distribution[4]. Of course, they are also made available to external users, both for download of newly acquired data and for reuse of legacy data obtained for other programmes. The essential role played in operations by the archive thus makes the operations workflow critically dependent on the speed with which the data obtained at the observatory can be ingested in it.

The end-to-end model has been validated by its successful implementation for over 10 years[5], and is now applied with minor variations to the operation of the La Silla facilities and to the other facilities on Cerro Paranal. Many elements of this model have been adopted by the planning of operations of the ALMA millimeter and submillimeter array[6]. Conveniently evolved in the coming years, it will also form the basis of the operation of the E-ELT[7], allowing its integration in a common scheme with the VLT.

## 3. CURRENT DEMANDS

The transfer of data through the Internet[8] from Paranal to Garching was declared operational in late 2009, following approximately one year of tests and optimization of the new system. At present, the transfer of data from Paranal (and also from the ESO La Silla observatory near La Serena, Chile) to the ESO Headquarters in Garching takes place through two channels:

- *Online transfer:* data are staged on a dedicated area at the observatory, queued, and transferred through the Internet to a staging area in Garching, from where they are permanently copied to the ESO Science Archive Facility (SAF). The SAF hardware consists of a primary archive based on spinning disks, and a secondary archive that uses a tape library[4]. The queuing system is based on priorities, as set in configuration files under the control of the Data Handling Administration groups on La Silla and Paranal. By tuning the configuration, the most critical files for operations, such as the calibrations needed to measure the health status of the instruments, are assigned a higher transfer priority. Under normal circumstances, at present 50% of the data assigned to the online transfer reach Garching within 10 minutes from acquisition at the telescope, and 90% within 30 minutes. The transfer time becomes occasionally longer when programs having very high data acquisition rates are executed.

- *Physical disk transfer:* data are copied from the staging areas at the observatories onto USB disks, which are then shipped to Garching on a weekly basis. The disks are shipped from the mountains to the ESO central offices in Vitacura, Santiago de Chile, on Thursday, and from there they carry on to Garching, where they are received and processed typically on Monday or Tuesday. They are then archived in the Primary and Secondary Archive and made available for processing. Therefore, once a shipment reaches Garching, it contains data that are between about 4 and 10 days old. Files, or category of files, are assigned to the disk transfer by means of the same configuration files used to regulate the online transfer. The double channel is designed to ensure the prompt online transfer of high priority data without saturating the available bandwidth, in which case data can be made to overflow to the disk channel. Most notably, VISTA and VST data are currently foreseen to be transferred solely via the disk channel.

### 3.1 Why fast data transfer?

Having the data available in the ESO archive in Garching within a short time from acquisition has positively impacted both the user experience and the observatory operations:

- *Principal Investigators* of programs executed in Service Mode can access their proprietary data in almost real time through their accounts in the ESO User Portal[9] to speed up their scientific exploitation. Although they are not present at the facility when their observations are executed, most Service Mode users are eager to gain access to their data as soon as they have been obtained to be able to start the scientific analysis. It is thus no surprise that this service has proven to be rather popular, with more than 100 dataset downloads a month on the average. PIs can access their data, both raw and pipeline products, upon authentication with their User Portal credentials.
- *Data handling operations* are simplified on both sides of the Atlantic by eliminating most of the logistics associated with the transfer and handling of physical media.
- *Science Operations on Paranal* is the place where the fast data transfer has had its biggest impact. Thanks to it, the dedicated Data Processing and Quality Control group in Garching can perform in-depth quality control[10] in quasi real time using the most recent files. In particular, the quality control loop involving the crucial calibration needed to measure and certify the health status of instruments is typically closed within one hour from when the suitable frames are acquired. This removes most of the needs for similar activities on the mountain.

Arguably, excluding specific cases like, e.g. Target of Opportunity and Rapid Response Mode observations, an almost immediate availability of the data is in principle a useful convenience for users, which is appreciated and even expected nowadays from a leading observatory, but not really a compelling necessity. However, given the strong competition in many of the scientific disciplines routinely addressed by VLT observations, a delay of weeks between the execution of the observations and their delivery to users may have serious and undesirable consequences in delaying the publication of the results. On the other hand, for operations the fast data transfer opens up an entirely different operational scheme, where the mission critical tasks needed to assess the suitability of instruments to be used are not carried out at the observatory, but rather in Garching, in application of the principle ‘bring data to people, not people to data’. Only the results are then fed back to the observatory itself, enabling the decision-making required for operations. This is one the cornerstone concepts of an evolution of the operational model for the VLT that is being developed and implemented, known as “SciOps 2.0”. At its very basis lies the ability to transfer at least the most critical data for quality control on timescale much shorter than the interval between when they are acquired and when decisions on possible corrective actions on the instruments have to be applied: in practice, the few hours between when the daytime calibrations are taken after the end of an observing night and the beginning of the next.

### 3.2 Current data volumes and data transfer capabilities

The current average production rate of VLT-VLTI instruments is 20-25 GB of data (all data volumes given in this paper are after compression) per 24-hour period. Until very recently, the nominal capacity of the Internet connection between

Paranal and the ESO headquarters in Garching has been 6.4 Mbits/s, whereas the total bandwidth from Vitacura to Garching is currently 50 Mbits/s. Out of these 6.4 Mbits/s, a nominal 4.5 Mbits/s was reserved to the Scientific Data Transfer, corresponding to about 50 GB over a 24-hour period. The data transfer software[8] is designed to maximize the transfer throughput and it tries to occupy all of the available bandwidth. This has made it necessary to set a ceiling of 80% (now 70%) of the total capacity to the bandwidth that can be used by the data transfer, so as to protect the performance for other Internet traffic, e.g. e-mails or browsing.

Apart from a couple of instances, in May and July 2009, the bandwidth available until recently from Paranal could accommodate the entire stream from VLT-VLTI, even in case of episodic very high nightly data production like the burst observing modes. The indication from this first year of operations of the Scientific Data Transfer is that the nominal available bandwidth should be over dimensioned by *at least* a factor of two with respect to the average data rate to be transferred.

An important upgrade has taken place in May 2010 leading to an approximate doubling of the capacity, not only providing a larger margin to deal with data rate peaks but also opening the possibility to add to the data stream transferred via the Internet the calibration files being produced by the VISTA survey telescope. The bandwidth currently reserved for data transfer is currently 9.12 Mbit/s, out of a total available bandwidth of 11.2 Mbit/s.

# 4. VLT/VLTI SECOND GENERATION, SURVEY TELESCOPES, AND BEYOND

## 4.1 VLT second generation instruments and survey telescopes

With the new telescopes and instruments scheduled to start operating at Paranal in the next few years[11], the amount of data that will be produced, and that will have to be transported to Garching, is about to increase significantly. The most significant contributors will be VISTA (entering in operations in 2010, with an expected average production of 75 GB/night of compressed data, including all calibrations), VST (2010, 35 GB/night compressed), SPHERE (2011, 40 GB/night) and MUSE (2012, 30 GB/night). The other Paranal instruments will keep contributing with about 25 GB/night (compressed). All in all, the average data production of Paranal in year 2012 will be of about 205 GB/night (compressed). Transferring it to Garching within 24 hours via the Internet corresponds to a sustained bandwidth of about 19 Mbits/s.

## 4.2 VLTI 2$^{nd}$ generation instrumentation

VLTI 2nd generation instruments[11] are expected to be operational somewhat later than their VLT counterparts, in the year 2014+ timeframe, and to yield somewhat smaller data volumes: MATISSE (2014, 10 GB/night; these data do not compress very well, so the figure here is for uncompressed data), GRAVITY (2014, 3 GB/night) and VSI (2014+, 3 GB/night). The total average data production of Paranal in the years beyond 2014 will thus be of about 225 GB/night (compressed). Transferring it to Garching within 24 hours via the Internet corresponds to a sustained bandwidth of about 21 Mbits/s. The conclusion of this cursory analysis is that the current bandwidth from Paranal to Garching will fall short by a factor ~2 from that needed to transfer the data produced on Paranal when the 2nd generation VLT-VLTI instruments will come online over the next few years. Furthermore, the requirement of transferring the entire data volume in 24 hours is a minimum that guarantees that no backlog will be produced, but still implies transfer times of up to several hours for some categories of data.

## 4.3 The E-ELT

As a part of the detailed design phase in which the E-ELT is at the time of this writing, a number of instrumentation Phase A studies[12] have been concluded. The nightly data volumes to be produced by E-ELT instrumentation is obviously difficult to predict nearly one decade in advance and with the actual selection of instruments to be built still pending, but the instrument concepts that have been developed have provided already a crude estimate of the average nightly data production to be set at the level of 1-2 TB per night, with large variations depending on the instruments and modes actually used on a given night. Assuming compression factors similar to those of the current VLT instrumentation, a rough estimate of the data transfer requirements of the Paranal observatory, including the E-ELT on Armazones, is 0.5 TB to be reached around year 2020, a quantity that may be further increased through the upgrade of planned Paranal instruments with larger format detectors.

## 5. DEALING WITH FUTURE DEMANDS: EVALSO

The daily transfer of the Paranal and Armazones data production over the internet will require a bandwidth increase of roughly one order of magnitude with respect to the capabilities of the current setup over the next decade. While this may be considered as feasible given the ever increasing availability of bandwidth with time, it must be recalled that the baseline requirement is to be able to transfer the complete data stream over a period of 24 hours. However, some of the advantages of the data transfer over the internet materialize only if the transfer time becomes much shorter and closer to real time. For instance, the implementation of a remote interaction mode involving the transfer of data just acquired from the observatory to a user located on another continent, which is being considered to be offered at the E-ELT in the future[6], would require Gbit/s-level capabilities over the whole path. Guaranteeing the closure of the quality control loop within less than one hour for all the future facilities on Paranal and Armazones also requires significantly higher bandwidth.

The EVALSO project (for 'Enabling Virtual Acces to Latin South American Observatories)[13], funded by the European Union under its Framework Program 7 and currently under construction, is expected to provide such capabilities for Paranal and Armazones already in the very near future.

The main goal of EVALSO is to create the missing parts of the physical infrastructure to connect the Paranal and Cerro Armazones observatories to Europe with a high capacity link[13]. To this end, a consortium was formed in 2007 by seven European institutions (the GARR consortium, the University of Trieste, and the Astronomical Observatory of Trieste in Italy, Queen Mary University of London, NOVA in the Netherlands, the Astronomical Observatory of the Ruhr University of Bochum, and ESO) directly interested in the exploitation of possibilities offered by the availability of a high capacity link with the Paranal facilities and the Observatory of Cerro Armazones, plus the REUNA and RedCLARA networks in Chile. The project plans to use the ALICE/ALICE2 research network infrastructure created with EU support since 2003 for high-capacity interconnectivity within South America and transatlantic connection to European National Research Networks (NREN) via GEANT2. In this collaboration, ESO is procuring the infrastructure needed to connect Paranal to the existing networks linking Santiago with Europe, and in particular a 75 km-long fiber link between Paranal and the access point to the Chilean backbone. Its completion is expected to take place in the third quarter of 2010.

In the EVALSO implementation currently being put in place, the capacity of the path between Paranal and Armazones is limited by that of the transatlantic link, which with the planned ALICE2 upgrade is expected to exceed 1 Gbit/s. Preliminary tests on this segment using the current (pre-upgrade) ALICE infrastructure have already achieved a sustained transfer rate exceeding 50Mbit/s between Santiago and Garching over the REUNA network. The tests were conducted using the *bbcp* file transfer tool, which is capable of transferring files at approaching line speeds in the WAN. The most critical bbcp configurations for achieving a good transfer speed are the TCP window size and the number of parallel streams, for which a single set of values based on initial tests were chosen. Considering that all existing links have higher nominal capacity and that the planned upgrade within ALICE and the trans-Atlantic link will even increase such limits, it is envisaged that it may be possible to achieve faster transfer rates by further optimizing the system and network parameters.

Real-time performance (low jitter, no packet drop) is not guaranteed by the transatlantic link via RedCLARA, which makes parts of the EVALSO infrastructure possibly unsuitable as a general-purpose communications infrastructure all the way to Europe. However, such high standard of real-time performance is not required by data transfer, and even with such caveat the new infrastructure should thus provide sufficient data transfer capacity for at least the next decade on Paranal and Armazones, enabling data files produced by the scientific instruments to be stored in the Garching archive within seconds from being produced and making data processing to start in a practically instantaneous manner. The high capacity of already existing NREN would make it possible the extension of fast data transfer to many other locations in Europe, which may be used in the future to enable remote interaction with the facility as described above[14].

## REFERENCES


[1] Quinn, P.J., Albrecht, M.A., Ballester, P., Banse, K., Chavan, A.M., Grosbol, P., Peron, M., Silva, D.R., 1998, SPIE 3349, 2

[2] Peron, M., 2008, in "The 2007 ESO Instrument Calibration Workshop", Springer-Verlag.

[3] Comerón, F., 2004, Organizations and Strategies in Astronomy, vol. 5, ed. A. Heck, ASSL Vol. 31, Kluwer Acad. Publ., Dordrecht

[4] Eglitis, P., Suchar, D., 2010, in "Ensuring long-term preservation and adding value to scientific and technical data", ESA, in press

[5] Kaufer, A., Comerón, F., 2010, SPIE 7737, in press (this conference)

[6] Andreani, P., Zwaan, M., 2008, SPIE 7016, 31

[7] Spyromilio, J., Comerón, F., D'Odorico, S., Kissler-Patig, M., Gilmozzi, R., 2008, The Messenger, 133, 2

[8] Zampieri, S., Forchi, V., Gebbinck, M.K., Moins, C., Padovan, M., 2009, ASP Conf. Ser. 411, 540

[9] Chavan, A.M., Tacconi-Garman, L.E., Peron, M., Sogni, F., Dorigo, D., Nass, P., Fourniol, N., Sforna, D., Haggouchi, K., Dolensky, M., 2008, SPIE 7016, 30

[10] Hanuschik, R., 2007, ASP Conf. Ser. 376, 383

[11] Moorwood, A., 2009, in "Science with the VLT in the ELT era", Springer-Verlag

[12] D'Odorico, S., Casali, M., Mainieri, V., 2009, in "Science with the VLT in the E-ELT era", Springer-Verlag

[13] Filippi, G., Jaque, S., Liello, F., Chini, R., Utreras, F., Wright, A., Lemke, R., 2010, SPIE 7740, in press (this conference)

[14] Comerón, F., Filippi, G., Emerson. J.P., 2008, SPIE 7016, 59