DOME: operational metrics under one roof

F. Primas[a], S. Marteau[a], L. E. Tacconi-Garman[a], V. Mainieri[a], M. Rejkuba[a], S. Mysore[a]
[a]User Support Department, European Southern Observatory, Karl-Schwarzschildstrasse 2, 85748 Garching, Germany

ABSTRACT

Thirteen VLT/I instruments plus some extra critical components like the block-scheduling of the Laser Guide Star Facility and VLTI baselines make for a rather complex machine that constantly challenges our operational efficiencies. DOME (Dashboard for Operational Metrics at ESO) is an ongoing project developed, implemented and maintained by the ESO User Support Department. It aims at providing an ESO-internal dashboard where key operational metrics are published and updated at regular intervals. Here, we will present the project and report on the indicators that have been looked at until now.

Keywords: Operations, metrics, efficiency, completion rates, observing modes

1. INTRODUCTION

DOME is an ongoing project, intended to provide snapshot and monitoring information on a variety of operational metrics to all ESO Operations-related groups, and eventually to the users' community.

The motivation for the project is at least twofold: i) to improve our understanding of operational efficiencies with a regular monitoring effort; ii) to respond in a professional and reliable way to the increasing demand of providing a variety of operational statistics. In recent years, the User Support Department received a large number of (usually urgent) requests, from different parties within the organisation. These requests always required the revision of already existing database queries and/or development of new ones, which sometimes revealed limitations of our current database structure and/or operational model. It is thus time to define and implement a consistent and reliable set of operational metrics that can help us to monitor our efficiency and to face last minute but important requests. As mentioned above, an important part of this project is also the documentation of caveats and/or limitations of our current operations schema, so that these can be properly addressed in future developments (e.g. the set-up of the E-ELT operational databases).

The first release of DOME-based metrics took place in early January 2012. We started from the first phase of an observing cycle (Phase 1) and have thus focused on "Proposals and Scheduling." Here, we present the main results of this first release and provide an update on what we are tackling at the moment.

2. THE FRAMEWORK

One of the starting points of DOME has been the list of all operational statistics the User Support Department has received during the last 10 years (more than 50 different requests), complemented by those requests that we serve on a regular basis (e.g. to the Observatory in Chile, to the Observing Programme Office). Next, we performed a critical revision of all the already available database (SQL[1]) queries and combined them according to their main operational area (Phase 1, Phase 2[2], post-execution, etc). While at this, we folded in our past experience within VLT operations and added other possibly interesting key operational metrics.

---

[1] SQL = Structured Query Language, a standard language for accessing and manipulating databases

[2] In ESO terminology, Phase 1 and Phase 2 refer to all tasks/actions related to the preparation and submission of respectively observing proposals (Phase 1) and of observing material to be executed in Service or Visitor Mode (Phase 2)

During this first part of the project, which was completed with the first release in January 2012, we defined the project workflow and carefully checked and optimised the corresponding database queries. At the same time, the DOME project team discussed the outcome of the project with other ESO operational groups/departments, taking note of their needs in terms of metrics they would like to be able to monitor. Finally, we critically evaluated the results and decided how to present them (in terms of interfaces, graphical solutions, etc.).

3. INITIAL RESULTS

As already mentioned in the Introduction, the first part of the DOME project focuses on 'Proposals and Scheduling', i.e. on metrics mostly related to what in ESO jargon we refer to as 'Phase 1.' Although these aspects do not seem to have much to do with "completion rates" and how efficient the Observatory is, it is very important to set the stage, i.e. to know what our users' community is asking for and what gets actually scheduled at the telescopes. The Very Large Telescope (plus VLTI) operates in a mixed Service and Visitor Mode (SM and VM for short) with Service Mode runs grouped according to their scientific priority in three rank classes A, B and C. Furthermore, runs can be of different types: Large, Chilean (when the PI is affiliated to a Chilean research institute), normal, Guaranteed Time Observation (GTO), Director Discretionary Time (DDT), Target of Opportunity and/or Rapid Response Mode (ToO and RRM respectively). A-class runs have the highest priority, whereas C-class are often defined as 'filler' runs, i.e. runs that can be executed under very relaxed atmospheric conditions and when no observations from higher ranked runs are available.

We identified three top metrics in this operational domain:

1. Time requested in Phase 1 proposals (Service vs. Visitor Mode, on which instrument, average time request, etc.)
2. Time allocated during scheduling (similarly to the point above)
3. Statistics about Principal Investigators of submitted and scheduled proposals

For each of these main points, we derived a range of graphs and trends: per period, per instrument, per rank class, etc. PI statistics aim at gathering and following up the distribution of ESO Principal Investigators in terms of how many ask for Service Mode, Visitor Mode, how many "new" PIs ESO is still able to attract every period and for which instrument.

3.1 Requested vs. allocated time and VLT/I oversubscription rates

Figure 1 shows the evolution of the time requested for all VLT telescopes and VLTI based on submitted proposals. The requested time is shown separately for Service and Visitor Mode and it is always counted in hours. For Visitor Mode proposals we applied the following conversion rule: 1 night = 10 hours for odd-numbered periods, 1 night = 8 hours for even-numbered periods, corresponding to Chilean winter and summer semesters, respectively. The relatively constant and higher than originally expected 50% fraction of SM time requested on VLT/VLTI is an indicator that the ESO astronomers' community has a high level of confidence in the Observatory and likes the Service Mode option.

Similarly, Figure 2 shows the allocated (i.e. scheduled) time both for Service and Visitor Mode on all four UTs and the VLTI since Period 63 (which marks the start of operations of the first VLT Unit Telescope). The numbers shown here refer only to "newly scheduled" runs within a given period, i.e. we did not count the amount of time assigned to carried-over runs (i.e. those A-class runs that could not be completed within their original period, hence they were granted carryover status). Numbers provided for early periods (P63 to P75) might be slightly uncertain due to a different schema for logging the amount of carryover time in the database.

Combining the number shown in Figures 1 and 2, one easily derives the overall ratio of requested/allocated time, or *pressure factor*, as a function of period (cf. Figure 3). All telescopes (UTs and VLTI) and instruments are combined, and the three lines in the graph represent Service Mode, Visitor Mode, and the mode-independent pressure factors. We emphasise that these graphs represent averages over all telescopes. However, there are quite substantial differences from one telescope/instrument to the other. For example, UT2 has reached pressure factors of 6, while other telescopes have remained more accessible.

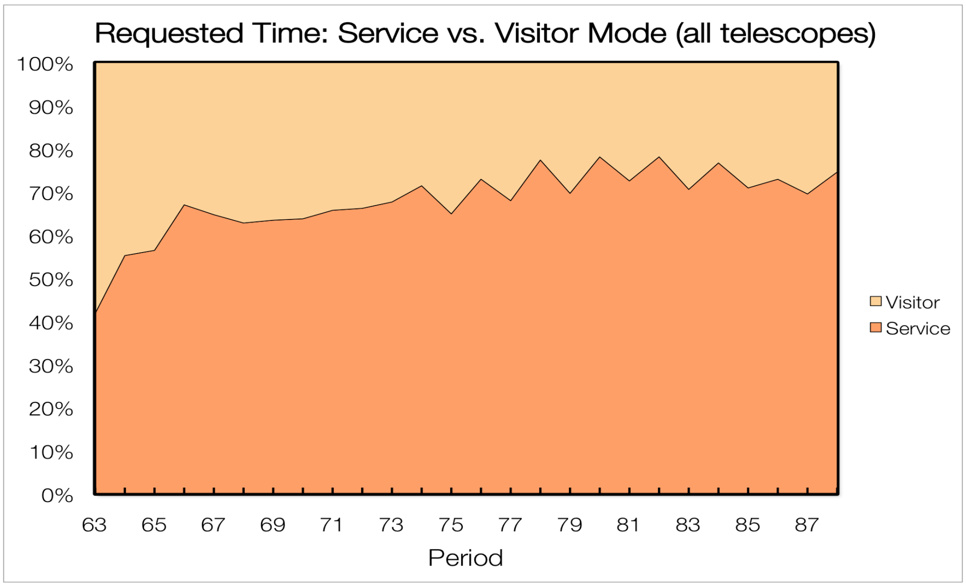

Figure 1. Relative percentage of Service vs. Visitor Mode hours requested on VLT/VLTI per ESO observing period.

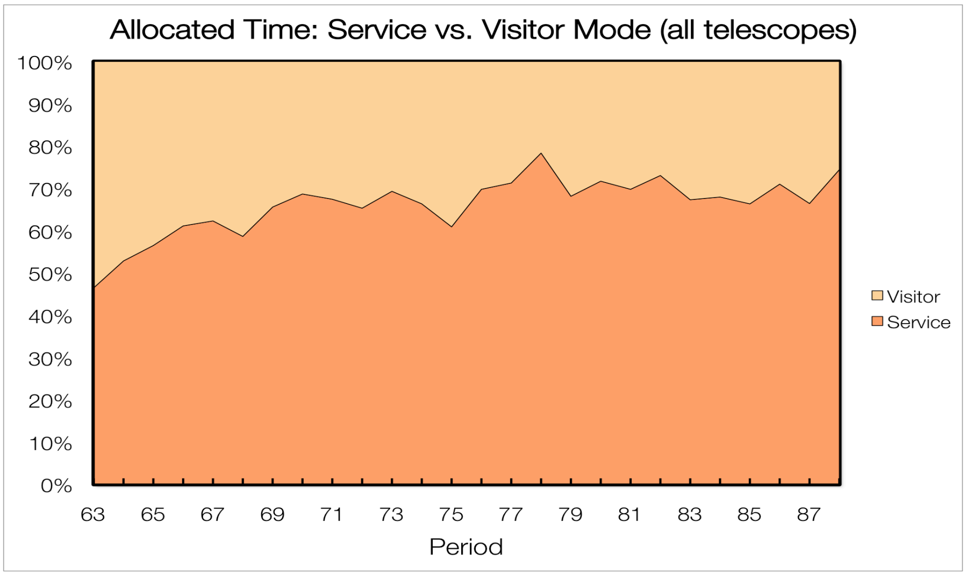

Figure 2. Percentage of time allocated on VLT/VLTI, per period, in both Service and Visitor Modes.

3.2 More Service and Visitor Mode figures

Another interesting aspect to look at is the average length of time requests and allocations.

Figure 4 shows how the number of newly allocated Service Mode runs and their corresponding number of hours has evolved with time, since the beginning of VLT/I operations (i.e. Period 63). We extracted both numbers — allocated runs and hours — as a whole (all instruments together), per instrument and per rank class. Here, we decided to show the evolution of the mean length of SM runs, grouped by rank class (cf. Figure 4). The larger mean for C-class runs is expected, since these are mostly ‘filler’ runs.

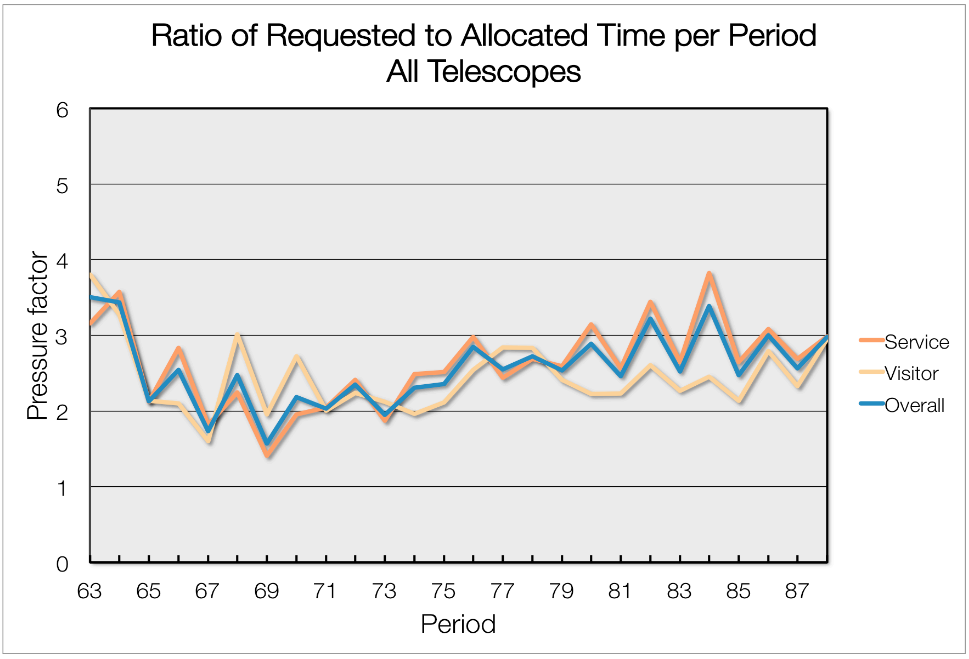

Figure 3. Pressure factors at the VLT. To ease the identification of the lines, over the P80-P85 range, one sees Service, Overall and Visitor displayed from top to bottom.

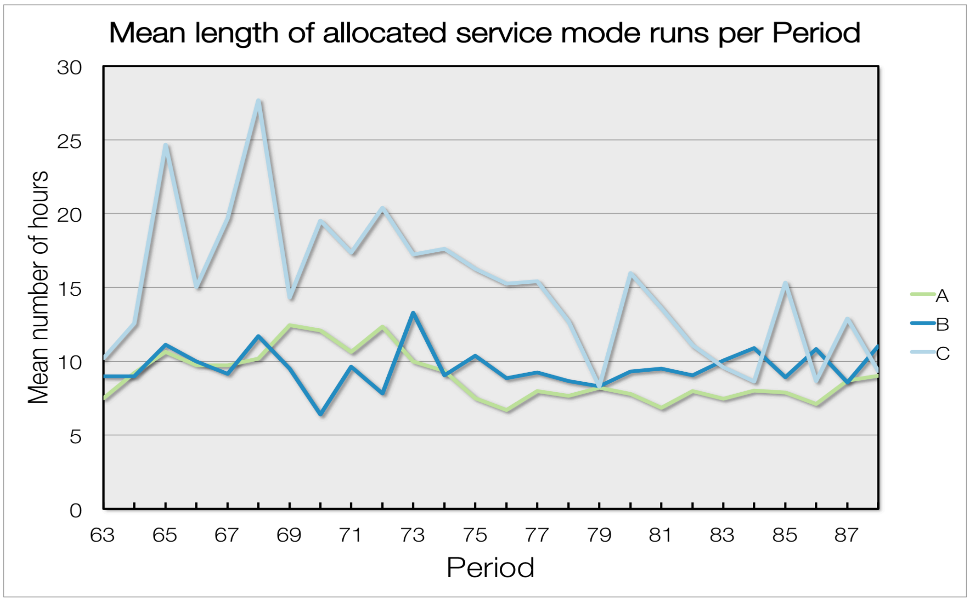

Figure 4. Mean number of hours allocated to Service Mode runs, per period and rank class. Over the P80-83 range, class-A, -B and –C appear from top to bottom.

A similar plot can be easily derived also for newly allocated Visitor Mode runs, the only difference being the unit of time, which in the case of Visitor Mode runs is 'nights' and not hours. By comparing Figures 4 and 5 one clearly notices that the mean length of SM runs has not evolved much with time, whereas the average length of Visitor Mode runs has steadily decreased (accompanied by an increase number of approved Visitor Mode runs and Visiting Astronomers travelling to Paranal Observatory).

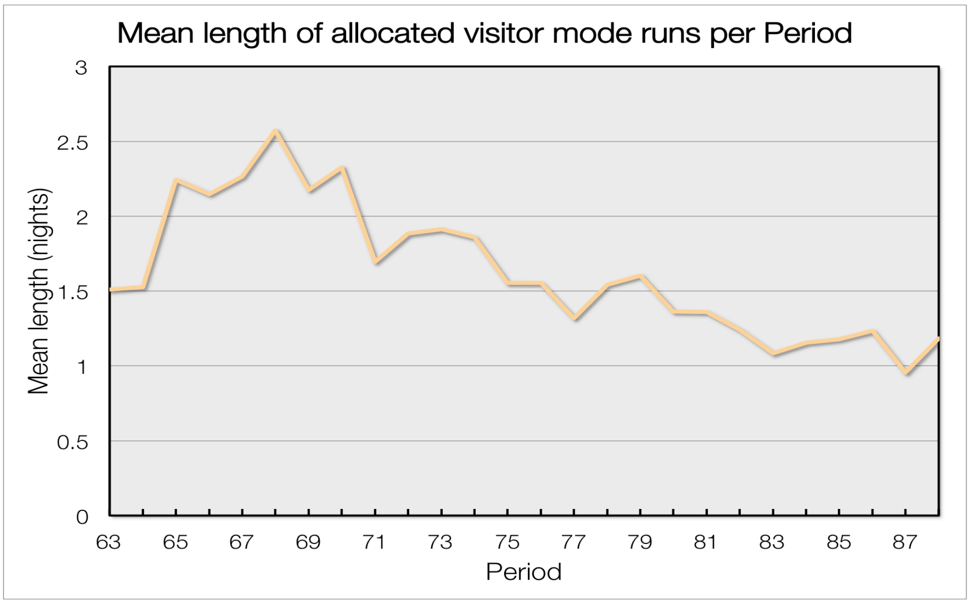

Figure 5. The evolution of the mean number of nights allocated to Visitor Mode runs since the start of VLT operations. In recent periods, the mean length of a VM run seems to have stabilized around 1night/run (or slightly above that).

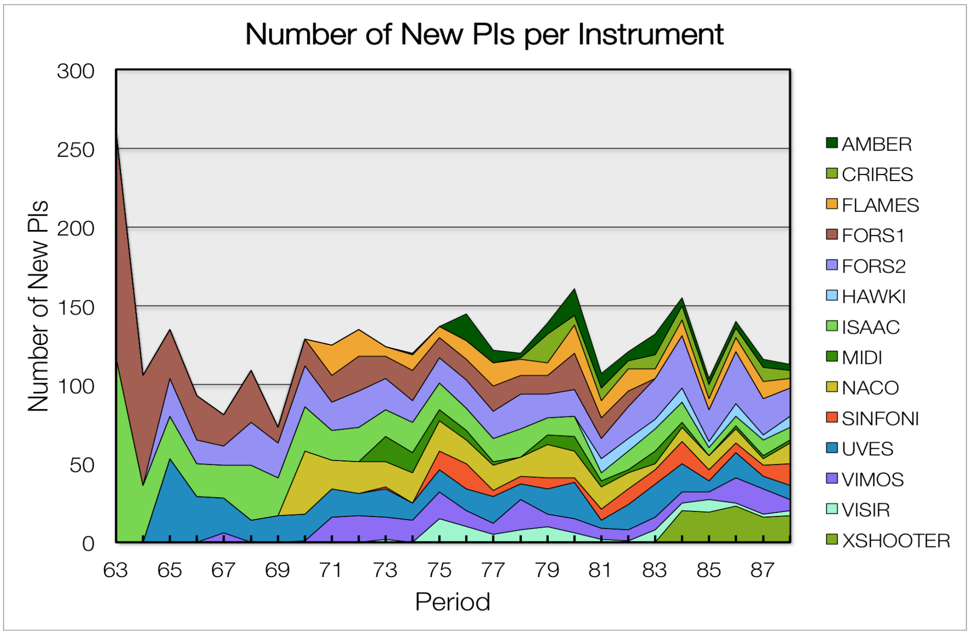

Figure 6. Which are the instruments that attract new Principal Investigators? Colour-coding follows instruments list, from top to bottom.

3.3 Principal Investigators of scheduled observing runs

We also looked at Principal Investigators of submitted and approved ESO proposals, both for Service and Visitor Mode for all VLT telescopes, plus VLTI. This may not seem very relevant, but it is interesting to check how many new PIs ESO facilities attract every period. Out of 300 individual PIs who get time in any given period, approximately one third of them are new Principal Investigators, who never applied before for ESO time (as PI). Figure 6 above shows the distribution of these new PIs per instrument: interestingly enough, the strongest 'attractors' are not only the newest instruments on the VLT (e.g. X-Shooter) but also the oldest one (e.g. FORS).

4. WHAT'S NEXT

We just started to tackle the core of our operational metrics, i.e. completion rates and related efficiencies. Our goal is to extract and monitor completion rates of A-, B- and C-class runs per period, per instrument, per run type. This seems a straightforward exercise, however there are several aspects/caveats to take into account. For instance, how do we define 'efficiency'? One can refer to Observatory efficiency at-large (shutter-open time) which is mostly a measure of how stable and well functioning a given telescope/instrument is. One could look at number of scientific publications, but this is strongly connected to the original science cases, to the team of proposers, etc; it is a nice add-on but we should not measure our efficiency just based on these figures. What we ultimately would like to know is how efficient we are in delivering the approved observations in a timely fashion. This may be somewhat inaccurate for Visitor Mode runs because of their larger dependence on actual atmospheric conditions, although information is logged in the so-called End-of-Mission reports that Visiting Astronomers are kindly asked to fill out at the end of their observing runs. On the other hand, it should be rather reliable for Service Mode runs. Here, possibly, the challenge is to add the time dimension to basic figures such as 'number of class-A, -B and –C completed', because when one derives the time it took us to complete a given run, one would also need some information about the start of the visibility window of the corresponding targets, if and when the observing material was delivered to us in time, etc.

At the time of writing, we have decided to start from deriving the simplest information possible and then develop more complex queries and explore a larger parameter space. We have thus started from deriving the number of class-A, -B and –C runs that are completed (based on 'run_status' entries in the database) and we plan to refine these figures in terms of timescale (what is the fraction of class-A runs that gets completed within the period they were initially scheduled for?) and of completion fraction (important especially for class-B and –C runs, i.e. to which percentage do they get observed?). It is our goal to accompany these metrics with monitoring the number of hours that are invested each period in repeating the same observations (due, e.g., to changing external condition). As well, we plan to standardize the number of runs (and eventually total amount of time) that is approved each period for carry-over status (i.e. not yet completed class-A runs). Clearly, the latter should be provided per telescope and per instrument, as to identify as soon as possible operational issues with specific (usually, most challenging) instruments and configurations.

ACKNOWLEDGMENTS

This is a project intended to serve all ESO Operations-related groups and departments. Therefore, constructive feedback is acknowledged from all main stakeholders.