Processing Requirements for DMO's Data Processing and Quality Control group

Reinhard Hanuschik & John Pritchard 2010-03-26

This document describes the processing requirements of DMO's Data Processing and Quality Control group (QCG) for the next 3-4 years (2010-2013). It is intended as input for the OTS department to execute the project of replacing the current processing hardware.

In this document we describe in section 1 important working patterns of QC. In chapter 2 we provide statistics about current and future key metrics like data volume and processing jobs. Finally in chapter 3 we expand on an evaluation of the current system, to identify properties that should remain, and issues that should be improved with the new system.

1 HOW DOES QC WORK? AN OVERVIEW

1.1 Processing use cases at QC

Two main QC processing use cases exist:

- Day-to-day processing of the current data stream (*full QC processing*) **[UC1]**
- Reprocessing of a defined data set **[UC2]**

In the following we describe the workflow for UC1. The workflow for UC2 is a small subset of UC1 and is much more homogeneous and automated.

The intention of this section is to provide insight into automatic and interactive steps that are nested in the QC workflow, and why good interactivity is as important as computational performance.

Typical components of the QC workflow are a mixture of

- automatic procedures (cronjobs) triggered once per hour (→ each must execute, under all reasonable circumstances, faster than in one hour);
- interactive jobs (launched on the command line as *step A*, then *step B*, then *step C*);
- job execution in the background, often performed over night.

1.2 Automatic procedures

These procedures are mainly required to

- find, retrieve, and process new data incrementally;
- update QC information on the web

They are typically triggered once per hour. By the nature of their frequent launch pattern, they are usually not very demanding in terms of data volume, but none the less require good performance in order to be finished before the next instance starts, and in order to optimize feedback times.

1.3 Interactive workflows

These interactive workflow steps are a fundamental characteristic for the QC scientist's work. They are necessarily interactive since decisions need to be taken (certify or reject; analyze). Here, good interactivity (short response times, snappiness) is the main issue. Waiting for results for too long a time, or interrupts in machine response, or poor performance in general are generally considered as not acceptable. Although it's hard to quantify, the resulting requirement is that *the user should ideally always experience the snappiness of an idle system,* and that background number crunching or data downloads or disk I/O have no impact on interactive response.

1.4 Background jobs

Longer jobs that do not require interactivity are typically sent to the background, to be executed over night. They require computing performance, but neither good interactivity, nor optimization for disk I/O.

1.5 A bit more about daily workflow steps

Figure 1 displays a screenshot of the DFO monitor for the daily workflow. It is essentially organized in rows per date. Within a row, the workflow goes from left (processing of calibration data) to right (packing and finishing). A finished date disappears from the monitor, while a new date is automatically added at bottom.

The purpose of this figure is to demonstrate the complexity of the daily workflow and its main components: there are automatic parts (1), interactive parts (2,4,6,7), and background jobs (3,5). Number crunching occurs within steps 1 and 5. Find the full explanation of this figure in the Appendix 1.

This mixture of workflow and performance patterns is typical of the QC scientist's day-to-day work (UC1). It repeats for each supported instrument, where one QC scientist on average has two instruments to support. This is why there is the need for the compute platform to provide compute performance, high throughput, and good interactivity at the same time.

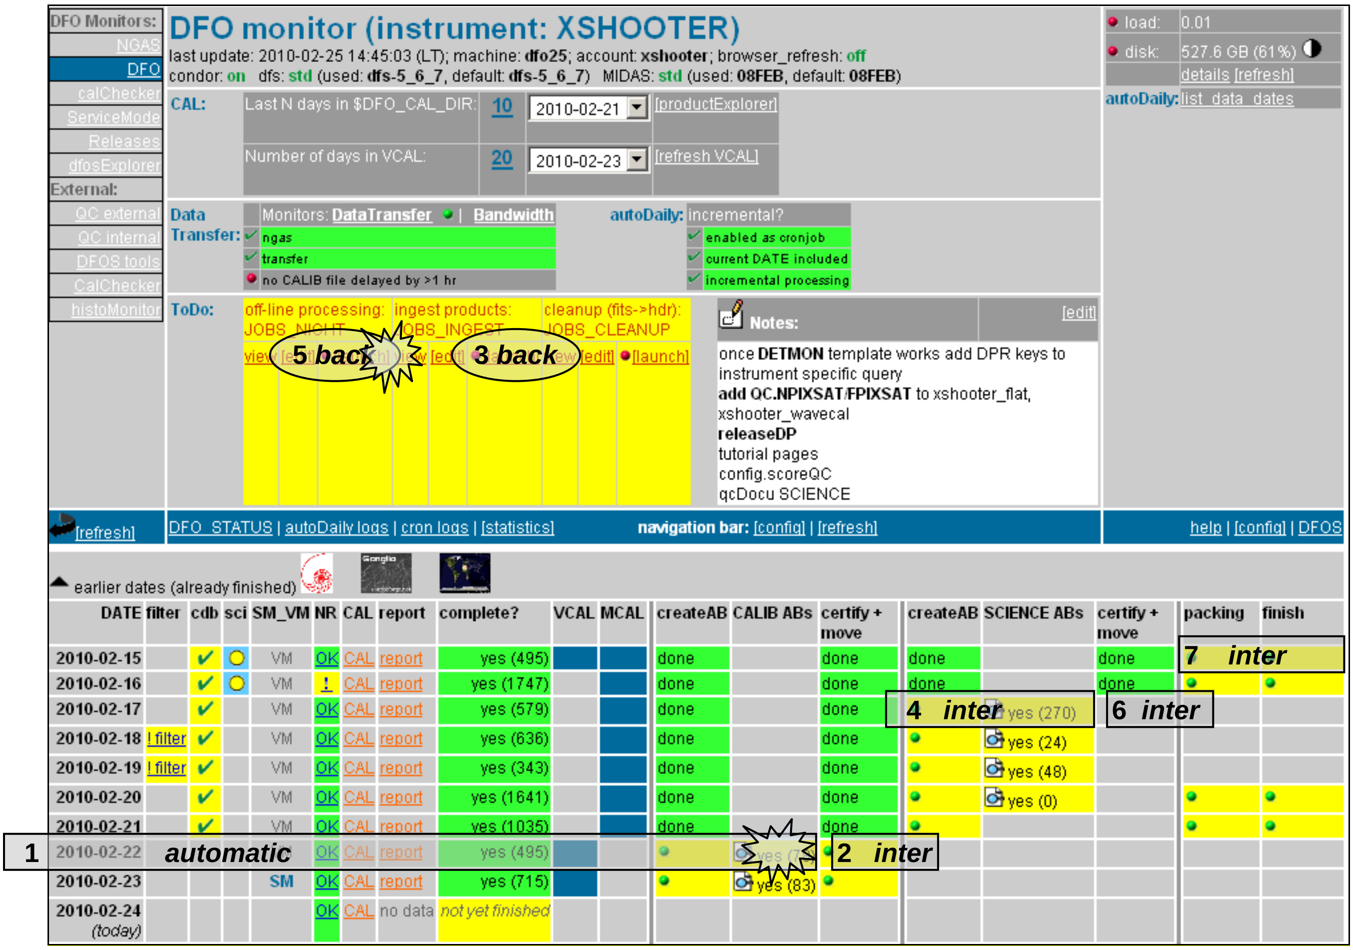

Figure 1: Screenshot of the DFO workflow monitor for daily processing. Steps 1 and 5, as marked by the "explosion cloud", require high compute power and network throughput. "Back" jobs execute in the background, "inter" jobs require interactivity.

1.6 Burst handling

There is the need to offer *resources on demand* to handle processing peaks (e.g. process 1000 ABs from a burst mode night where 50 would be normal; see Figure 2). Since it is unlikely that two such bursts from two different instruments occur at the same time, the resources can be found and assigned by harvesting idle resources if they exist.

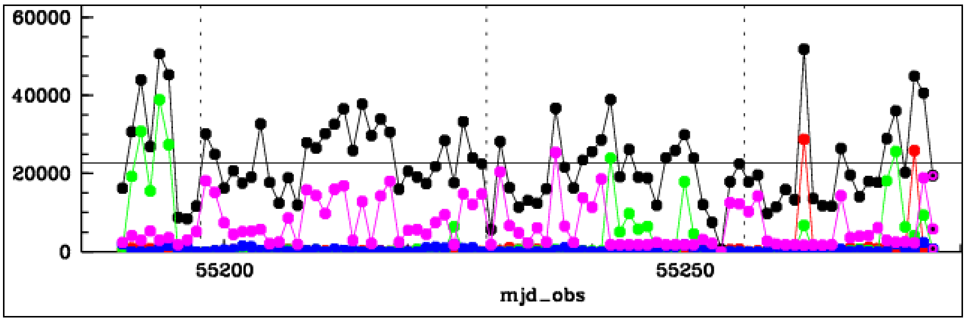

Figure 2: Bursts of NACO (green) and ISAAC (red) data acquisition. These are daily rates in (compressed) MB, for the period 2010 Jan-March. Average daily data production for these two instruments is in the range of 500 MB each. The observed bursts correspond to an increase by a factor of 80 (NACO) and 50 (ISAAC), resp. (purple is AMBER+MIDI; black is the sum of all VLT and VLTI instruments).

2 KEY PARAMETERS AND REQUIREMENTS FOR QC PROCESSING

2.1 UC1: Day-to-day processing of the current data stream

2.1.1 Data streams

UC1, daily processing of the current data stream, is the first important QC use case. There is always[1] one data flow per instrument, that is completely separate from the flows for all other instruments. There are currently 11 VLT instruments to support, 2 VLTI instruments and one survey instrument (VIRCAM)[2]. Within a year, there will be a third VLTI instrument, PRIMA, and a second survey instrument, OCAM. For the rest of the document it will be assumed that both PRIMA and OCAM are already operational.

Within the next 3 years, there will be three more so-called second-generation instruments going into operations (KMOS, SPHERE, MUSE). Each of these will replace a currently existing instrument. Hence the number of instruments for QC to support at any given time in the projected time range will be 16.

The data streams from these instruments fall into two distinct categories:

- Moderate data volume and moderate number of detectors (*N*, “multiplex factors”): 14 instruments (12 as of 2012) (“**moderate DV**”)
- High data volume and/or high number of detectors: 2 instruments (4 as of 2012). These are: VIRCAM (*N*=16, with *N*=number of detectors); OCAM (*N*=32); SPHERE (not high *N* but high DV); MUSE (*N*=24, to come in 2012) (“**high DV**”)

---

1 Here, “always” does not mean that every day there is data from all instruments.

2 Plus two separate processes, La Silla package preparation and pre-imaging processing, but these are essentially subsets of the processes described here and will be ignored for the rest of this document.

2.1.2 Class of moderate data volume

Total raw data volume to be processed for moderate DV levels out at about 1 TB per month[3]. There are about 20,000 processing jobs (Association Blocks, or ABs) per month. Products are created corresponding to about 0.6 TB per month. These numbers are expected to be stable for the next few years, although there is always a certain risk that the operational mode of an instrument evolves into a regime with significantly higher data volume.[4]

2.1.3 Class of high data volume

The high DV class currently consists of VIRCAM only. The current processing scenario foresees QC processing of 100% calibrations, and 10% science (for spot checks[5]). By projection based on the first 4 months of operations, we assume the following numbers: 4 TB per month to be processed; 40,000 processing jobs; 0.3 TB per month as products.

Lacking any measured numbers so far, we assume the same processing scenario and data volumes for OCAM.

For MUSE and SPHERE, QC will process 100% of all data. For MUSE, the consortium estimates about 1 TB per month of raw data on average[6]. The product data volume will be 3-5 times higher, up to 5 TB. Number of processing jobs will be in the 10,000 regime.

SPHERE will produce about 0.6 TB per month of raw data on average[7]. We expect roughly 10,000 processing jobs per month. Product data volume will be negligible.

The accumulated key numbers for all high-DV instruments then are: 10 TB per month of raw data to process, 100,000 processing jobs, 5 TB of products (mainly driven by MUSE).

2.1.4 Key performance numbers

Important features for **UC1** are:

1. This processing mode is under operational time pressure (the calibration data stream should be processed on near-real time, i.e. on timescales of hours rather than days), due to the requirement of *closing the loop for near-real time QC feedback to the mountain.* However this pressure exists only for a subset of the data (calibrations) that needs to be processed to retrieve quality information about the instruments
2. There is also some time pressure requirement (though relaxed in comparison: timescale of days) to provide data packages to the users. This applies to the complementary part of the data stream (science data).

Because of the operational pressure, *daily peaks drive the requirements more than monthly averages*. Daily numbers show a much larger scatter than the monthly numbers that can be more reliably estimated. Daily rates can show peaks of up to 10-50 times the average rates (see Figure 2).

---

[3] The following numbers are based on the WISQ information system which records key parameters of the QC workflow: http://www.eso.org/observing/dfo/quality/WISQ/overview.html

[4] In particular, ISAAC, NACO and VISIR can potentially produce in 10-50x higher data rates than currently, by running routinely in burst mode.

[5] The spot checks are done to provide some baseline for later verification of data quality when external scientists submit their complete set of processed data to the ESO archive. The 10% fraction could be increased or reduced by decision of the QC scientist. All VIRCAM and OMEGACAM numbers given here relate to the assumption of 10% science processing.

[6] Based on a daily average data rate of 30 GB of science data, the same in calibrations, and 15 days of operations per month. Daily rates may potentially have 10-50 times (!) higher peaks.

[7] 40 GB in an average night, or 600 GB in a 15-day month. Product volume is negligible.

Since peaks are usually followed by periods with low volume, we assume for the requirements a global contingency factor of 3 (rather than 10-50 for individual instruments and days)[8]. Then, the daily key performance numbers projected for 2012 become:

Regime	Number of instruments	Daily processing volume*	Daily processing jobs*	Daily product volume*
Moderate DV	14	100 GB (30 GB*3)	2000	60 GB
High DV	4	1 TB	10,000	1 TB**
Sum	(18)***	1.1 TB	12,000	1 TB

*Based on average rates *3
** Dominated by MUSE, therefore we cannot relax to a factor of 3 but follow the numbers given by the consortium
*** Sum is never higher than 16 but some moderate-DV instruments are replaced by high-DV ones.

2.1.5 Data transfer and bandwidth

Downloads. The daily processing volume puts a requirement on the bandwidth for downloading data into the compute platform. For operations, “bandwidth” must be interpreted as an end-to-end property. For downloads this is:

- Archive access
- Internal bandwidth.

With the current intranet, a typical transfer rate of 30 MB/s has been measured, which in theory can support up to 2.5 TB/day. The bottleneck is archive access: we currently measure an NGAS data access rate of about 4 MB/s corresponding to 0.3 TB/day which falls a factor 3-4 short of the above estimate for the daily processing volume (1 TB).

Uploads. The daily product volume determines the required bandwidth for uploads. As for the downloads, the current network bandwidth is ok while archive access (for the ingestion process) is the bottleneck. The current situation is that the ingestion time has two components, a measured flow rate of 4 MB/s and a fixed overhead of 3 sec per file, so that effectively the bottleneck is even worse than for archive downloads.

Both archive downloads and archive ingestion processes will need an improved bandwidth in order to first feed, and then deplete, the processing system in phase with the needed processing throughput.

2.1.6 Data access pattern

Once data are downloaded into the system, they need to be accessed by multiple processes in parallel. This is true for both regimes:

- Instruments with moderate DV need to share master calibrations among the processing jobs
- Instruments with large DV need, in addition, to split raw and product data and join new products.

This brings up the issue of disk I/O, which becomes serious if many (more than say 10 or 20) processes try to access the same file at the same time, which is a realistic scenario for calibration products (“master calibrations”). The maximum number of simultaneous accesses to the same file (or the same part of a file) corresponds to the maximum number of currently active processing jobs. For current VIRCAM processing on the QC blade cluster, this number is 34.

Several disk access patterns exist and it is probably worth elaborating on these in order to help make the right decisions with respect to storage and disk I/O options. For instance, it would be fully acceptable to have disk I/O management (provided by e.g. condor or the operating system) such that the maximum number of simultaneous reads or writes to the same file is limited.

---

[8] Partly because of statistics (two peaks happening at the same time are very unlikely), partly due to the option of a relaxed processing mode: processing of a factor 10 peak can be spread over 2-3 days (but not over 10 days).

2.1.7 Storage

Disk space is primarily required as temporary disk space for the certification area (where products are delivered by the pipelines and stored until review by the QC scientist). There is no requirement to store larger data volumes permanently. Raw data are deleted after processing. Product data are deleted after archive ingestion.

The temporary disk space must be big enough to accommodate data delivery peaks, processing peaks, and longer periods of data accumulation (during travel or holiday). For each instrument in the moderate DV regime, 1 TB of disk space is good enough (in total 14 TB). For high DV, we estimate that 20 TB in total is sufficient.

2.1.8 Number of cores

Experience with the current compute platforms has shown that for each moderate-DV instrument two cores combined with a local disk gives a reasonable balance between required interactivity and number crunching power. Effectively one core is then available for on-line commands, while the other core may be busy with scheduled jobs or with pipeline processing[9]. For background jobs (overnight processing), both cores are available.

For high-DV instruments parallel processing capabilities are required where, as a rule of thumb, the number of simultaneous processes corresponds to the number of detectors that can be processed simultaneously. Then, to first order, OMEGACAM processing (N=32) on such a platform could be performed in the same time as e.g. FORS2 processing (N=2) on a single dual-core blade.[10]

As a speculative guideline, this "naive" scenario, with parallel threads are mapped into cores, would assign 2 cores per moderate DV instrument as minimum, and *N* cores for high DV instruments, *N* being the total number of detectors. This scenario yields 100 cores[11] (14x2 + 16 + 32 + 24) but would still not deliver the required processing power for bursts.

A smarter approach would be to assume that the required load is always somehow distributed, while on the other hand there should be resources available on demand for the moderate DV instruments to support bursts. This dynamic resource allocation is in fact very important for the new processing platform.[12]

A dynamic platform would establish a minimum number of cores for background and interactive jobs, for each instrument, and offer dynamically additional cores for number crunching. Ideally as required, realistically limited by the total number of currently available cores. With such architecture, a total number of 60-80 cores would be sufficient and could provide support for bursts.

2.1.9 Execution times

As the WISQ system[13] shows, QC currently executes about 50,000 processing jobs per month, this takes about 2200 hours[14] processing time yielding an average pipeline recipe execution time of 2.6 min. The actual

---

[9] Even pipeline processing using both cores is possible during daytime. Although it can have a noticeable effect on interactive response, the incremental nature of the calibration processing during daytime generally means these periods are relatively short and the QC scientist can coordinate the need for interactive response fairly easily.

[10] Another dual-core processor would be reserved to interactive processes that do not scale by number of detectors but by number of files. This additional processor is neglected in the following since it is in the noise of the assumptions.

[11] This does not necessarily specify a multicore machine, in principle it could also be a cluster or a grid

[12] Recent experiments with XSHOOTER jobs for burst nights have been exported from a dfo blade to the QC cluster, thus establishing a "QC grid". The tremendous compute power of the QC cluster has provided a very reasonable execution time (5 hours on the QC cluster with 30 compute nodes instead of a whole weekend on a dual-core dfo blade), and required at the same time very small changes to the overall QC workflow since all steps before and after processing took place on the home host.

[13] http://www.eso.org/qc/WISQ/FULL/trend_report_AB_ALL_HC.html

number per processing job can of course vary greatly between instruments and recipes, execution times may take between a couple of seconds and an hour or more. The current number includes a certain mixture of simple tasks (like image subtraction, image stacking), where data volume is the issue, and complex, compute-intensive tasks (like spectrum extraction, PSF fitting, data cube construction), where processing complexity is the challenge. For all current instrument modes it is true that they either fall into categories "volume" or "complexity" but not in both at the same time.

This is also true for the expected new instruments, with the exception of MUSE. MUSE will start producing data in 2012. With a multiplex factor of 24, MUSE data can be processed like VIRCAM or OMEGACAM data on a multicore platform until it comes, as the last step in the processing cascade, to a combined processing of the data from all 24 detectors in one go (gridding to a common reference system). Here we face the new challenge of combining volume and complexity. In other words we expect **many and long** MUSE processing jobs[15]. We assume that the MUSE jobs will shift the *average execution time* per recipe to higher values. With a factor of 2 higher times then (which is very conservative), we end up with an average execution time of 5 minutes.

Another way of looking at execution times is to say that data processing under UC1 should always occur in near-real time. On average the incoming data must be treated in less time than it takes to acquire them. If not, data processing will get further and further behind data acquisition. This requirement must be fulfilled for each instrument individually, and is particularly challenging for MUSE:

We must be able to process a full night worth of MUSE data (calibration and science) in a few hours.

2.1.10 MUSE requires 64bit OS

For the last step in the MUSE processing cascade mentioned in the section before, QC will need about 10 GB of RAM for processing one data product, and we may need to combine many of those into a final data product which may require several 100 GB RAM. This amount goes beyond the memory limit of a 32bit OS.

For all current instruments including VIRCAM and OMEGACAM there is no such requirement, although it would certainly not hurt all other instruments to be able to go beyond the 32bit/~3GB per process memory limit. There are examples when already now processing becomes RAM-limited[16].

2.2 UC2: Reprocessing projects

Occasionally there are reprocessing projects that require partial or complete processing of the entire data set for a given instrument. A related case is the need for reprocessing because an error was discovered, or a significant improvement of the pipeline becomes available.

The first reference example for this UC2 was the UVES reprocessing project. Executed in 2007, it included 59,000 processing jobs, required a total execution time of 88 days (distributed over 6 cores) and produced 370,000 product files. [17] To have this project executed in one month of processing time, 18 cores would have been needed. Employing 66 cores (the current total number available to QC), the project would have needed just 8 days.

This use case needs massively parallel processing to provide maximum efficiency and reasonable process timescales. There is no natural upper limit to the number of streams (unlike in UC1). Typical for UC2 is that it

---

[14] Summing up parameters AB_texec and QC_texec.

[15] Recent numbers for the final, recombination step in MUSE are 1 hour execution time per detector on a single-core machine. Remember there are 24 detectors.

[16] Occasionally UVES visitors take FLATs with more than the usual 5 RAW frames. When this number is above 10, the pipeline recipe fails due to lack of memory. Similarly HAWK-I FLAT templates typically comprise 20-30 RAW files and occasionally 30-50, but the final product is derived from only the first 10 as a direct result of the 32bit memory limitations.

[17] See more under http://www.eso.org/~qc/monitor/REPROC/reproc.html

consists of science jobs only[18] that have no mutual dependencies and can therefore always execute at the saturation limit of the compute platform (be it caused by the number of available cores, or by the available bandwidth for data downloads or uploads).

The UC2 has no real-time requirement, but it should be efficient enough to be finished in a reasonable time (days, weeks or months depending on the project scale).

Key properties for UC2:

- No operational time pressure
- High multiplex factors for efficiency

It is likely that UC2 can be covered by the architecture for UC1 (large DV case) if that platform is properly dimensioned.

3 HIGH-LEVEL REQUIREMENTS FOR QCG

3.1 Summary of requirements

Performance:

- Because of the operational time pressure in UC1, end-to-end performance of the system is important: the whole chain, data download – processing – data upload, needs to be optimized. An overall throughput of 1 TB (inflow and outflow) in less than 24 hours should be supported day by day. We will need faster archive access for this.
- For the performance of the core system (processing platform), there is the requirement that average execution times for pipeline recipes should be the same or shorter than on the current systems. Because of the UC2, and in particular because of the requirement that repeated processing may be needed, the system must be able to process data actually even faster than they can be downloaded or uploaded. In numbers, 12,000 jobs should be processed in much less than 24 hours.
- With 12,000 jobs taking 5 minutes each (section 2.1.9), we would need at least 40 parallel processes to finish within 24 hours. To do this in 8 hours, we will need 120 parallel processes.
- We must be able to process a full night worth of MUSE data (calibration and science) in a few hours.
- Non-scalable components (like network connections) need to be high performing and peak oriented rather than average oriented.

Storage:

- Total disk space: 1 TB for each moderate-DV instrument
- 20+ TB for high-DV instruments in total[19]

Interactive response:

- Very good interactivity/response for interactive part (meaning no noticeable delay by concurrent processes in the background)
- No interference with other users (other than trivial things like the common queue is already full); e.g., poor interactivity due to disk I/O by other users is not acceptable.

18 However there could also be cases where the pipeline improvement affects also calibrations.

19 About 6 TB is needed to store the data for one month of VIRCAM processing.

Flexibility:

- Architecture should provide flexibility and scalability since key numbers for new instruments like SPHERE and MUSE are still very uncertain.
- There should be the chance to scale up hardware in two years from now, otherwise we need to buy a system now which might then turn out to be oversized.
- For handling data bursts and occasional reprocessing, there should be compute resources, memory, disk space available on demand (dynamic resource allocation).
- Timescale for dynamic resource allocation is *not* very short (no need for on-the-fly and automatic allocation): it is likely that allocation on demand, e.g. by configuration, is sufficient.
- Support inhomogeneous workload (we have different regimes: moderate data volume and high data volume).

Stability:

- Minimize changes to the current compute model (which has a reserved workspace per instrument, no limitation by other users, essentially no limitation by disk space); the only resource that is currently sometimes scarce is compute power.
- The stability requirement is not just about maintenance but about the architecture. Observed achievements with the dfo and QC cluster blades are one or two failures (unscheduled downtimes) per year per blade, this is acceptable. The (now historical) failure pattern of the QC cluster (single node failure propagates to other nodes) would be unacceptable. A failed component must be brought back without minimal interruptions of the whole system.

Simplicity:

- It is preferable having a simple system that is very stable, easy to maintain, and performing (this order), rather than one that has cutting edge technology and is ultra-performing, but experimental (not understood) or unstable. As experience with the QC cluster and the fastcache has shown, the latter system would require high manpower beyond proportion for support, both on the SOS and on the QC side.

Others:

- It is worth considering to have reserved nodes for frequent automatic services (like *trendPlotter*).
- 64bit platform: this is currently not an issue but is required for MUSE (lack of memory otherwise).

Appendix 1: The QC daily workflow expanded

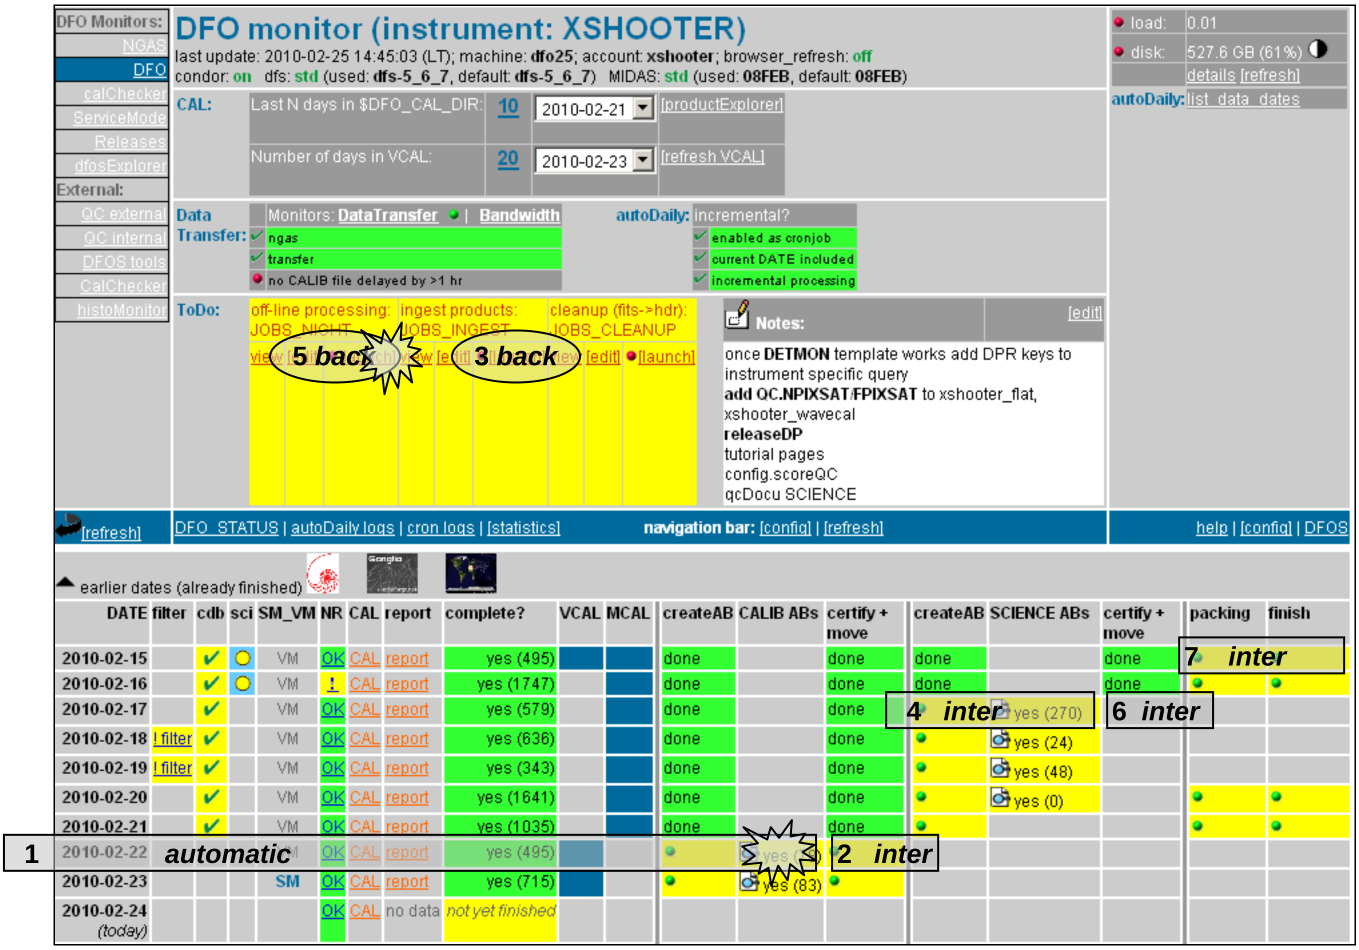

Figure 1 (repeated): Screenshot of DFO workflow monitor for daily processing. Steps 1 and 5, as marked by the “explosion cloud”, require high compute power and network throughput. “Back” jobs execute in the background, “inter” jobs require interactivity.

All steps for the first part of the workflow are actually automatically provided by a workflow tool, up to the certification of CALIB products (box “1”, date 2010-02-22). These steps are all provided by cronjobs, called in an incremental way once per hour. Pipeline processing is required as part of step 1, this is one of two steps in the daily workflow with high demands on compute power and network throughput. Once that date is finished the QC scientist pushes the green button to start an interactive session for the review of calibration products (“2”). This session requires good connectivity and fast response. Once all products are reviewed, the accepted ones are ingested into the archive. This step is launched as a job in the background (labeled “3”), but requires good network throughput in order not to take forever. Next, a processing job for science data is interactively created (they depend on CALIB data), for an earlier date than 2010-02-22 (box 4). This job, and other jobs, could be launched overnight (step 5). This is the second step requiring high compute performance and network bandwidth, due to the download of raw files. Step 6 then is again review of products, plus ingestion (skipped here). Finally there is step 7, updating all packages for a given date, and finishing. This is again fully interactive.

APPENDIX 2: CURRENT SYSTEM IN COMPARISON

Overview

The current system consists of 13 independent dual-core blades (dfo blades), plus 20 clusterized dual-core blades (QC cluster nodes), 66 nodes in total. The 13 dfo blades and the 20 QC cluster nodes have a similar, but not identical basic structure for the storage provided to each node. On all nodes, the root and swap partitions are provided by the first internal disk (sda, 250 GB) while home directories and common static software[20] are provided via NFS from a central filer. However the majority of local disk I/O is carried out on the so-called *data-disk* (the blue line in Figure 3, in comparison to the green line), which in the case of dfo blades is provided by a second internal disk (sdb, 1 TB). For the QC cluster nodes it is provided by the so-called fastcache which is a shared file system seen by all 20 nodes mounted over GFS based on a ~9 TB fibre channel raid array plus expansion.

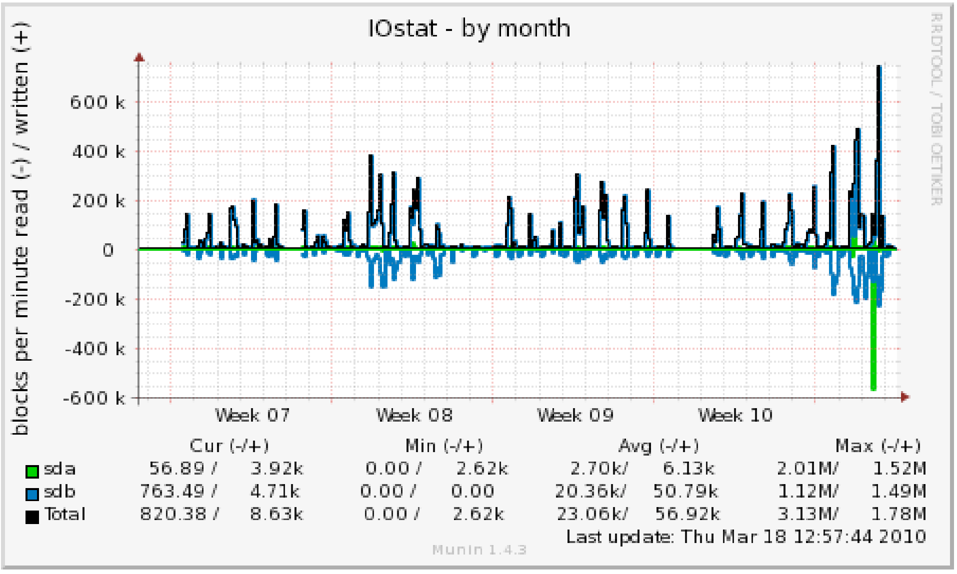

Figure 3: IO statistics for dfo blade *dfo21* for 2010 week 06 to week 11; first internal disk: sda (250 GB), second internal disk (“data disk”): sdb (1 TB). Reads display negative, writes positive.

Limitations

In general, the dfo blades and the QC cluster blades are performing well enough, with the exception of some accounts reaching performance limits (HAWKI; peaks/bursts of X-Shooter and UVES).

From experience, the following properties provide an acceptable work environment:

- Local disks 1 TB instead of 0.25 TB (no time wasted for disk space management)
- Shared disk for home accounts, pipeline installations etc.
- Total disk space: 1 TB for each moderate-DV; 20+ TB for high-DV instruments in total

Known issues are:

20 Such as /vlt, /opt, /opsw and /scisoft.

- Interactive response affected by disk I/O: no interactive work possible on cluster nodes if heavy processing (i.e. all QC Cluster compute nodes processing VIRCAM CALIBs) takes place, see next section
- Frequent pattern of cronjobs, competing for resources with interactive steps
- Data download pattern into fast cache too slow
- Move of product files from certification area to final area to archive: can be improved

Interactive response

Interactive response is a critical property of the system as this has a direct impact on the productivity of the QC scientist, whose time is the most expensive component of the overall system. Examples of degraded interactive response are:

- Seconds, rather than fractions of seconds, to open text files in vi
- Tens of seconds, rather than ~1 seconds, to open graphic files in xv
- Minutes, rather than ~10 seconds, to update html web pages

This is a very difficult property of the system to measure, as it is rather subjective. It is even more difficult to estimate for any prospective system. For this reason, any potential replacement should ideally be tested under realistic operational conditions.

dfo blade nodes

Degraded interactive response is most noticeable on the dfo blade nodes when pipeline processing is underway as pipeline processing is managed via condor, which makes use of both cores of the dfo blade node, meaning that CPU for the interactive process(es) must be shared with that of the pipeline processing. In general, since there are at any time only two pipeline processes running, disk I/O limitations are less of an issue, though there are likely particular cases of instruments or specific recipes for a given instrument which are more or less I/O intensive and thus I/O limitations no doubt also affect interactive response to a certain degree, especially if two such I/O intensive recipes happen to be running simultaneously. We thus believe the dominating limitation affecting interactive response for the dfo blade nodes is CPU.

It is important to note that each dfo blade node is completely isolated from all others in terms of CPU and data disk I/O, so to a certain extent this effect is limited for each QC scientist to his/her own actions, and it is (relatively) easy for a QC scientist to coordinate their own activities to allow for these periods of degraded interactive response.

QC cluster nodes

Currently two instrument accounts reside on the QC cluster: HAWKI and VIRCAM. Because each instrument operated on the QC cluster has a dedicated “operational node” where no pipeline processing is performed, interactive sessions have better access to CPU. On the other hand because VIRCAM pipeline processing is massively parallel with up to 34 processes all accessing (read and write) the fastcache shared file system simultaneously, disk I/O limitations have a greater impact.

Recent studies show that the QC cluster is currently disk I/O limited, a result that should be considered in the architecture of a new system.

VIRCAM data is (currently) received from Paranal once per week via USB disk, thus once per week there is a burst of VIRCAM processing of 7 days worth of calibrations. Typically this takes of the order of 36 hours to process. During this time interactive response is affected.[21]

[21] This effect was not noticed before VIRCAM operations began, i.e. when only HAWK-I operations were active on the QC cluster. Typically HAWK-I has only at most 10 processes active simultaneously, so this effect apparently only becomes noticeable for a multiplicity somewhere between 10 and 34.

It should also be noted that the HAWK-I and VIRCAM data access patterns are individually quite different. For HAWK-I each of the 10 processes that might be active simultaneously all depend on different input raw files, i.e. the 10 processes are in general accessing different files from the shared file system. For VIRCAM, 16 processes for the 16 detectors of a single fileset are launched simultaneously, i.e. 16 different processes spread over up to 16 different nodes are all trying to access the same files.[22]

Figure 4 and Figure 5 show IOStat and CPU performance trending plots for one of the QC Cluster compute nodes on a 7-day scale. The first point to notice in Figure 4 is the burst nature, this node is more or less idle for a large fraction of the week until the VIRCAM data arrives at which point it is processed during a little under 24hrs in this case, followed by a few, short-lived bursts a bit later in the week. Figure 5 shows the CPU usage for the corresponding period. What is striking in Figure 5 is that despite the fact that the cluster was "working flat out" total CPU usage (system+user+nice) never surpassed 140% (total available is 2x100% since two cores per node). Presumably due to the not insignificant IOwait occurring at the same time, demonstrating that the current QC cluster system is I/O limited, the fastcache is not fast enough to be able to make full usage of the available CPU power, at least not for VIRCAM.

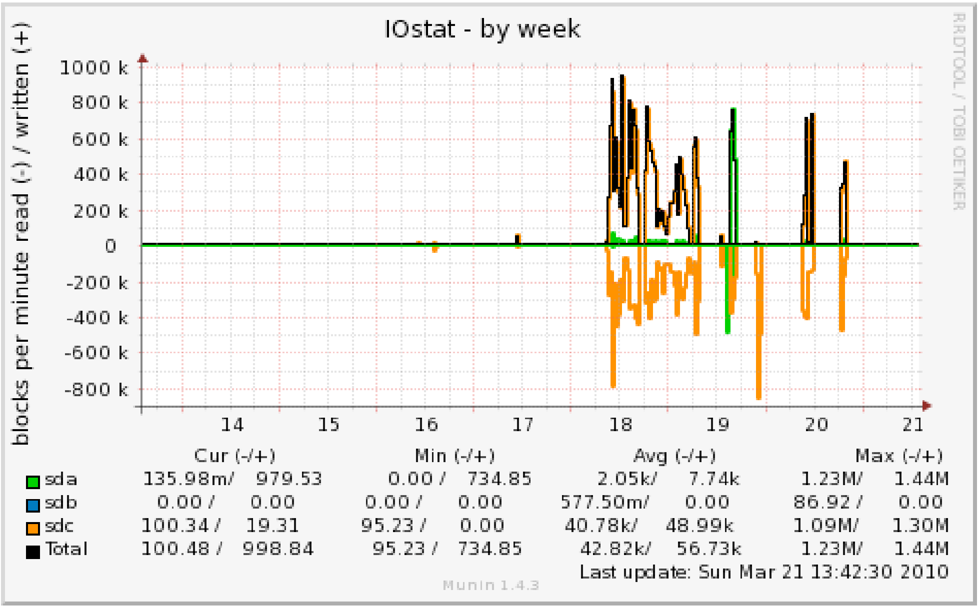

Figure 4: IO statistics for the QC cluster node *qc03* for week 2010-03-13 to 2010-03-21

22 In the best case, the 16 processes execute on 8 dual-core nodes. Then each file is read from the shared file system once by each of the eight blades and then cached by the blade in its RAM to be used by the two processes running on that node, i.e. 8 reads of the shared files system. In the worst case 16 different nodes are employed, and each of those 16 nodes must read each file from the shared file system once to then be cached in RAM for use by the single process active on each node, i.e. 16 sequential reads of the shared file system.

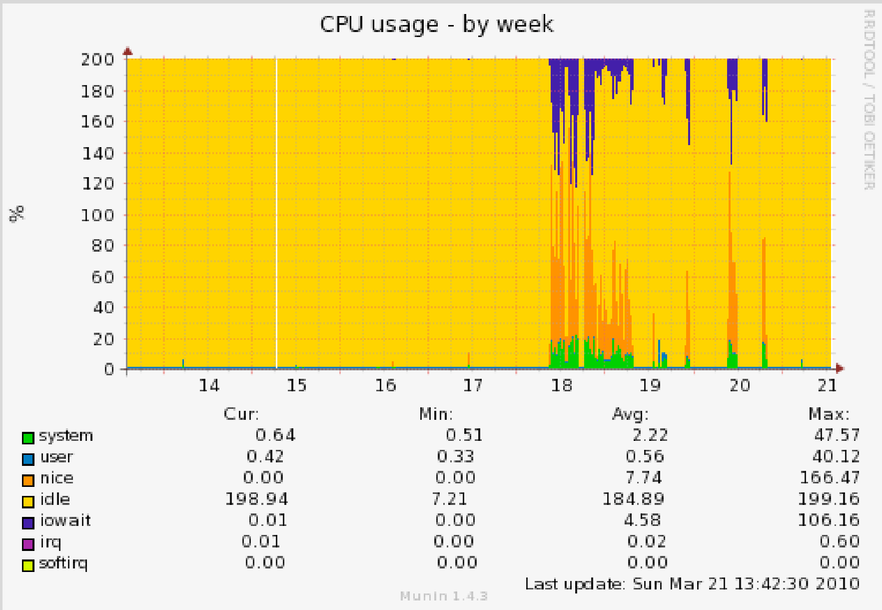

Figure 5: CPU usage for the QC cluster node *qc03* for week 2010-03-13 to 2010-03-21