Visible & Infrared Survey Telescope for Astronomy

Data Flow System

Document Title: **VISTA UK Pipeline Operations**

Document Number: **VIS-TRE-IOA-20000-0017**

Issue: **1.0 pre 3**

Date: **2006-09-22**

Document Prepared by:	**Peter Bunclark** (CASU)	**Signature And date:**	

Change Record

Issue	Date	Sections Affected	Description of Change/Change Request Reference/Remarks
0.1	2006-09-20	All	New document

Notification List

The following people should be notified by email that a new version of this document has been issued:

IoA	W Sutherland
RAL	G Dalton
QMUL	J Emerson
ATC	Malcolm Stewart

Contents

List of Figures

1 Introduction

The Visible and Infrared Survey Telescope for Astronomy (VISTA) is a new 4-metre telescope designed specifically for imaging survey work at visible and near-infrared wavelengths. The pawprint of the VISTA infrared camera (VIRCAM) will cover 0.59 degree^2 in ZYJHK_s passbands, using a 4×4 array of 2k×2k non-buttable chips with ~0.34″ pixels resulting in a 269MB FITS file per exposure. The camera is expected to produce a median of 210GB a night with a realistic high volume night producing 630GB.

1.1 Scope of this Document

This document discusses the practical considerations of producing science-data products in the UK from the flow of raw VIRCAM data, identifying processing power and person power required. It is assumed that the CASU group in Cambridge will be successful in funded to perform the Pipeline Operations. Where appropriate, requirements from the UK VISTA User Requirements [AD2] are referenced.

1.2 Which Pipeline?

VIRCAM data will pass through at least three pipelines from the beginning of operations:

- Paranal: The “QC1” (Quality-Control) pipeline will run at the telescope in near real time. It will use purpose-written CPL recipes to compute a range of QC parameters which will enable early assessment of the quality of the data. Calibration data (such as flat fields) will be drawn from a library which will be updated only infrequently. Because of the near-time aspect, each recipe can only use data taken up to that point.
- Garching: The same software will be run again at ESO headquarters in Europe. The difference will be a slightly more flexible approach to input data, e.g. twilight flats taken later than the science frames could be used.
- UK: The UK pipeline will be responsible for producing much of the science-grade reductions of VIRCAM data. Many of the algorithms used in the QC pipeline will be re-employed, the difference being that great care will be taken in the production of calibration data on input, and similarly only science-grade products will be released. In addition, some extra procedures not needed for the QC pipelines will be provided.

1.3 Applicable Documents

[AD1] *Data Flow for the VLT/VLTI Instruments Deliverables Specification,* VLT-SPE-ESO-19000-1618, issue 2.0, 2004-05-22.
[AD2] *UK VISTA User Requrements,* VDF-SPE-IOA-00009-0001, Issue 3.0, 2005-07-05.
[AD3] *\VISTA Infra Red Camera DFS System Impact,* VIS-SPE-IOA-20000-0001, Issue 1.3, 2006
[AD4] *VISTA Infra Red Camera DFS Calibration Plan*, VIS-SPE-IOA-20000-00002, issue 1.3, 2005-12-25.
[AD5] *VISTA Data Reduction Library Design,* VIS-SPE-IOA-20000-0010, Issue 1.4, 2006
[AD6] *Jimbo's UK pipeline doco.*

1.4 Reference Documents

[RD1] *VISTA Infra Red Camera DFS Data-Reduction Specifications,* VIS-SPE-IOA-20000-00003, issue 1.0, 2004-12-15.
[RD2] *VISTA Science Requirements Document*, VIS-SPE-VSC-00000-0001, issue 2.0, 2000-10-26
[RD3] *Estimate of the Number of CPU Nodes to Exectute the Vircam Pipeline in Parnal*, 2006-07-28.

1.5 Abbreviations and Acronyms

ATC	Astronomical Technology Centre (Edinburgh)
DFS	Data Flow System
PSF	Point Spread Function
SRD	Science Requirements Document
VDFS	VISTA Data Flow System
VIRCAM	VISTA Infrared Camera
VISTA	Visible and Infrared Survey Telescope for Astronomy
WFAU	Wide Field Astronomy Unit (ATC)
WFCAM	Wide Field Camera (on UKIRT)

2 Data Transport

2.1 Paranal to Cambridge

Transmission of around a third to a half terabyte a day from Paranal to Europe is currently not possible using networking. Data is written to hard disks using a Linux operating system and the ext3 file system. The disks are physically located in a Chenbro chassis and are permanently mated to the corresponding mounting bracket. Currently, a capacity of 250GB is standard. Arrangements are made to ship these disks periodically to Garching where they are inserted into an identical Chenbro chassis and the data read off.

Subsequently, because the network connection to ESO headquarters is also insufficient to support bulk data transfers, disks will be shipped to Cambridge for insertion into another similar chassis and read again for use in the UK pipeline. It is expected that there will be a three-week cycle before a particular disk is returned for re-use.

2.2 Cambridge to Edinburgh

The situation here is much simpler. With gigabit connectivity either over JANET or the reserved “UKLight” connection, reduced science products will be flagged as ready-to-transfer and copied up to the WFAU archive automatically using high-performance networking protocols.

3 Data Ingestion

Data ingestion is one of the most manually-intensive tasks in the pipeline, taking non-trivial operator effort. This is not merely the routine cycle, but substantial time in recovering lost data:

3.1 Copy ESO Disks

The data ingestion portal will consist of a relatively modest Linux system connected to the ESO-standard disk chassis. Data will be copied over to raw data storage via gigabit network using NFS protocol. A standalone ingestion system is deemed desirable because the constant physical insertion/removal of disks will inevitably result in frequent system reboots. From a systems point of view, there may be merit in running this machine as a near clone of an ESO workstation.

3.2 Verify Raw Data

Following copying, the new raw data collection will be verified by several means, including checking the FITS MD5 checksum, running “fitsverify” and cataloguing the files to compare with the expected list of observations. Verification failure will trigger a recovery process which would vary depending on the scale of the data dropout; a single file might be transferred over the network, a whole disk failure would require a resend from Garching or Paranal as appropriate.

3.3 Backup Raw Data

The opportunity exists at this stage to backup the raw data onto archival media, for example LTO-n tapes in a tape robot. This step may be deemed unnecessary since the data will exist at Paranal, Garching and on disk at Cambridge.

3.4 Recycle Disks

Shortly after the integrity of the new batch of raw data is affirmed, the transport disks are erased (probably at low level with “`mkfs`”), are replaced in their shipping container and sent back to Garching.

4 Triage

By this stage the data is safely on disk as an exact copy of that which left Paranal. However there will certainly be a fraction which is unusable for example underexposed or saturated twilight flats. The procedures detailed here constitute part of the fulfilment of *requirement 5.7*.

4.1 Bad Frames by QC

The QC pipeline would generate parameters indicating problems, and if these are routinely available it may be possible to pre-identify bad frames. At any rate, the science pipeline will be generating exactly the same set of QC parameters and so after reprocessing frames producing outlying QC parameters will be examined and flagged as unusable.

4.2 Bad Frames by Inspection

In particular, it is essential that frames contributing to calibration data, i.e. darks and flats are high quality. Unlike the near-time QC pipeline, the science pipeline has the advantage of being able to use data from both before and after a particular science observation but great care is necessary in building the calibration products and skilled human intervention is a necessary part of the process.

5 Computing Requirements

The requirements for processing power can be based on experience with WFCAM and from benchmarking results from early versions of the VIRCAM reduction recipes. Storage capacities are based on the likely data-acquisition curve for VISTA. The DFS Impact Document [AD3] claims a likely annual data rate of 95Tb. The UK User Requirements [AD2], *requirement 5.4*, states that the pipeline should operate at Tb/day for 10 days. Also, sufficient headroom must be available to allow scheduled and unscheduled downtime (hardware and software upgrades, system failure recovery). Given the well-know projection of future price/performance in IT, even at this stage the exact specification of equipment cannot, and should not, be made. Final choices of actual hardware will be based on a total systems approach, taking into account not only simple processing per pound, but physical footprint (real estate), power requirements (including cooling), reliability and serviceability.

5.1 Processing Power

Enough processing power must be available to meet the peak sustained data rate of 1Tb/day stated above.

WFCAM

The ESO DFS group have benchmarked the QC pipeline [RD3] and derived a requirement of 16 cpus to cope with near-time data analysis. The base model tested is a 2.4GHz Opteron with 4Gb memory. While that system must be able to run at peak data rates, and the UK one only at mean data rates, the UK pipeline will be doing extra processing compared with the QC pipeline. If these two effects balance out, then 16 cpus of this specification can be taken as a good independent recommendation.

5.2 Disk Storage

Enough disk storage must be available to hold the likely mean data acquisition stated above. A reasonable purchasing model might be to acquire enough disks to hold a quarter's worth of raw and reduced data at once. Given a factor of 2 expansion in reduced frames we need 3 times the raw data but luckily it will compress by a factor of 3 and so the first tranche of operational disk storage would be around 30Tb. At the time of writing this might be 3 10Tb SATA raid arrays costing £1k/Tb but as usual the cost is expected to reduce.

5.3 Off-line archive

It is desirable to backup the reduced products. Possibly the most effective way to do this is to simply provide duplicate disk storage off site.

6 People Power

The pipeline will operate nearly continuously and although it is massively automated, at any one time one person will need to monitor progress and perform interactive quality control. The data ingestion cycle takes about 3 weeks, within which in a given week there is about a day of disk handling. Data-transfer disk problem-shooting might take a day a month. Hardware commissioning and maintenance will require around a tenth FTE and operating system maintenance and upgrades a further 0.1 FTE. Upgrading and testing of new releases of VDFS software is estimated at 0.1 FTE.

Finally of course allowance must be made for management and administrative overheads. Operations staff members are expected to participate in relevant external meetings and conferences, and exploit training and development opportunities. To cover these each FTE estimate is multiplied by 1.2.

Activity	FTE
Running Pipeline	1.20
Data Ingestion	0.24
Data Ingestion follow up	0.06
Hardware	0.12
System Management	0.12
VDFS Software management	0.12
Operations group manager	0.50
Bodge factor	0.66

Figure 6-1 Conceptual Ingestion Server and Blade Pipeline Engine

__oOo__