

GIGAWULF: Powering the Isaac Newton Group's Data Pipeline

Robert Greimel,¹ Nicholas A. Walton, Don Carlos Abrams

Isaac Newton Group, Apartado de Correos 321, 38700 Santa Cruz de La Palma, Canary Islands, Spain

Mike Irwin, James R. Lewis

Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge, UK

Abstract. The Isaac Newton Group of Telescopes (ING) is currently implementing a unified data flow and processing system to support all its astronomical instrumentation systems. The ING data infrastructure is tightly integrated, encompassing the initial data acquisition, pipeline processing and subsequent archiving and distribution of raw and processed data products.

In this paper we describe *Gigawulf*—a Beowulf cluster based on commodity PCs running Linux—which provides cost effective processing power for the ING's image data pipeline processing. The configuration and performance of Gigawulf is discussed. The operation of the data pipeline on the cluster together with the integration of the *Gigawulf* pipeline processor system with the DVD-R tower archiving system is detailed.

1. Introduction

The Isaac Newton Group (ING) operates three telescopes, including the 4.2-m William Herschel Telescope (WHT) and the 2.5-m Isaac Newton Telescope (INT). In the era of 8-m telescopes, there is increasing pressure to operate 4-m class telescopes in the most economic, efficient, and effective manner possible. The ING is addressing these three 'E's, in part with its implementation of an improved streamlined data flow system encompassing data acquisition to archiving of processed data products.

This paper describes the ING's implementation of the processing unit, 'Gigawulf',² a Beowulf class cluster, which ensures sufficient computational capability to handle the ING data flow.

¹Institut für Geophysik, Astrophysik und Meteorologie, Universität Graz, Austria

²<http://zwolle.ing-slo.iac.es/Beowulf>

2. The ING Data Flow System and Data Rates

In order to support new large array detectors and increase on-sky observing efficiencies, the ING has recently introduced a new Data Acquisition System (DAS, Rixon et al. 2000). A key consideration in its implementation was that it would provide the front-end to the ING data flow system (Walton et al. 1998). Thus attention was paid to items such as correct and sufficient FITS (see e.g., Walton & Rixon 2000) header information being made available for each science and calibration data file.

All raw data are archived on DVD-R media, mass on-line availability provided by a dual juke box system with a current capacity of ~ 6 TB. The data are accessible via a WWW based front end to a Sybase database (Lewis & Walton 1998).

The advent of large format CCD arrays and large infra-red detectors has led to an explosion in data volumes. For example, the current data rates at the ING are determined by the Wide Field Camera on the INT which typically generates ~ 8 GB/night and the IR camera on the WHT generating some ~ 4 GB/night. In total, data flows can amount to 15–20 GB/night from all telescopes.

The current day availability of affordable processing power and storage capability has opened the possibility to provide (semi)-processed data products at the point of data origin to the visiting astronomer. These reduced data products will form the core data resource of new ‘Virtual Observatories’ (see e.g., <http://www.astro.caltech.edu/nvoconf/>).

3. The ING Data Reduction System

Details of the ING’s data processing pipeline, as implemented for the reduction of imaging data, are described elsewhere (Irwin & Lewis 2001). The basic data reduction steps are: linearity correction, bias subtraction, flat fielding and the application of a basic astrometric solution. Additional steps in the pipeline are de-fringing, an accurate astrometric solution, object detection, classification and catalogue generation.

Currently, the pipeline (*quick look* and *full*) is implemented for the reduction of imaging data, optical and near infra-red, only. A pipeline will be introduced for Echelle and multi-fibre spectroscopic data by the end of 2001.

The *quick look* pipeline delivers a processed image to the observers within five minutes of image acquisition. This enables immediate assessment of the image quality and instrument performance. The *quick look* pipeline differs from the science pipeline in two areas: it applies calibration files from the most recent science pipeline run instead of the calibration frames from the current run, and it implements a subset of the full reduction, terminating with the de-fringing stage.

The science pipeline provides the observer with reduced data shortly after the end of the observing run. To provide the highest quality data a limited amount of human intervention is necessary, mainly in rejecting poor calibration frames. This pipeline has been running on a Sun UltraSparc system servicing ING Wide Field Survey Data since August 1998 (Lewis et al. 1999).




Figure 1. Schematic overview of the *Gigawulf* cluster.

4. Powering the Pipeline: *Gigawulf*

In order to provide an economic processing unit for the pipeline, it was decided to use commodity PC components. Recent advances in the clustering of PCs, making use of the Linux (see e.g., <http://www.linux.com>) operating system have made this feasible. Linux based PC farms have been found by a number of groups to offer a powerful and cost effective solution for large computational problems. Indeed, a small number of PC systems have been developed for use in astronomical data processing environments (see e.g., Gravitor³).

The ING data pipeline is a coarse grained parallel processing case. To a first approximation, each science data frame is processed in an identical fashion, with no cross reference to any other. Therefore a night's data can be equally distributed between the nodes. PC clusters are ideally suited for this case (see discussion by Brown 1999⁴).

4.1. The *Gigawulf* Structure

Gigawulf is a 'Beowulf' type cluster (see <http://www.beowulf.org> for a definition and related links) of eight high end PCs. Each node consists of an AMD Athlon 950 MHz processor with 256 MB of main memory. The seven slave nodes have 30 GB EIDE hard disks while one node, subsequently called the head node, has two 75 GB EIDE hard disks. The head node also has a DDS-3 DAT tape robot as well as a second network card which provides the connection to the telescopes and data archives. The network in the cluster is 100 Mbps apart from the head node, which has a gigabit connection. A schematic view of the system is shown in Figure 1.

³http://obswww.unige.ch/~pfennige/gravitor/gravitor_e.html

⁴http://www.phy.duke.edu/brahma/beowulf_advanced.ps

4.2. ‘Gigawulf’ Software Issues

To minimise the operational and maintenance overheads, the Scyld Beowulf extension (Scyld Computing Corporation^{TM5}) to Linux (currently based on RedHat’s⁶ 6.2 distribution) has been used as the operating system for *Gigawulf*.

Scyld Beowulf supports standard Linux interfaces and tools. It enhances the Linux kernel with features (provide by bproc⁷) that allow users to start, observe, and control processes on cluster nodes from the cluster’s head node (Hendriks 1999⁸). With this arrangement, software needs only to be configured on the head node. The result is that the cluster appears to be more like a traditional multi-processor computer to a user or developer. This reduces the cost of cluster application development, testing, training, and administration.

The existing pipeline software, has been ported to run on *Gigawulf* with only small modifications. Scripts have been developed to handle the input of data from the telescopes and output of reduced data to archiving media (DVD-R and DDS-3 tape).

5. Conclusion

In terms of performance, the *Gigawulf* cluster is currently ~ 10 times more cost effective than using alternative computing hardware, for instance UltraSparc computers. The cost of use has been minimised because only limited changes have been required in order to run the existing data reduction pipeline software on it. For the future, *Gigawulf* will be enlarged with the addition of more slave nodes.

References

- Irwin, M. J. & Lewis, J. R. 2001, NewAR. (in press)
- Lewis, J. R. Bunclark, P. S., & Walton, N. A. 1999, in ASP Conf. Ser., Vol. 172, Astronomical Data Analysis Software and Systems VIII, ed. David M. Mehringer, Raymond L. Plante, & Douglas A. Roberts (San Francisco: ASP), 179
- Lewis, J. R. & Walton, N. A. 1998, in SPIE Proc., Vol. 3349, 263
- Rixon, G. T., Walton, N. A., Armstrong, D. B., & Woodhouse, G. 2000, in SPIE Proc., Vol. 4009, 132
- Walton, N. A., Bunclark, P. S., Fisher, M. P., Jones, E. L., Ress, P. C. T., & Rixon, G. T. 1998, in SPIE Proc., Vol. 3351, 197
- Walton, N. A. & Rixon, G. T. 2000, ING Newsl., 3, 31

⁵<http://www.scyld.com>

⁶<http://www.redhat.com>

⁷<http://www.beowulf.org/software/bproc.html>

⁸<http://www.beowulf.org/software/bproc.html/bproc-LinuxExpo99.ps.gz>