Using Open-Source Software and Low-Cost Computers for Earth-Observation Data Processing

M. Önnestam
ESA Directorate of Applications Programmes, ESRIN, Frascati, Italy

W. Balzer
Deutsches Zentrum für Luft- und Raumfahrt (DLR), Oberpfaffenhofen, Germany

M. Curtis
National Remote Sensing Centre Ltd. (NRSC), Farnborough, United Kingdom

The introduction of open-source software compliant with industry standards has allowed the use of low-cost standard computers for a new set of Earth-observation applications. The Linux operating system is an example of this kind, which makes it possible to use ordinary PCs for demanding data-processing needs. ESA has benefited from this opportunity by replacing old and difficult to maintain systems with new ones providing both a significant increase in performance and a reduction in investment as well as maintenance costs.

Processing Earth-observation data

The Agency's Earth-observation projects involve space missions in which instruments carried by polar-orbiting satellites provide data to numerous users all over the World. This data arrives in a raw format requiring processing on the ground to convert it to data products that can be understood and exploited by users. This processing has until recently required very specific hardware and software, implying expensive electronic and computer equipment. The evolution of technology over the last few years, however, has now made it possible to use standard, high-performance computers even for these demanding tasks.

The Agency utilises a number of facilities in Europe for its Earth-observation data-processing tasks, where equipment, systems and resources have been assembled to form entities that provide the required services. The ERS mission is such an example, supported by so-called 'Processing and Archiving Facilities' (PAFs) where the data are archived and made available for processing. The PAFs have data-processing systems that read data from the archives and produce data products for the users.

System re-hosting

The current evolution of technology in general and computers and electronics in particular, means that systems that were fully state of the art in their design when built rapidly become obsolete as more modern solutions become available. This creates problems in maintaining equipment, as spare parts for and knowledge of the systems become scarcer after a number of years. It also becomes increasingly difficult to add or change functions of the original systems, as their performance is limited and difficult to increase. It can also be difficult to find new peripheral equipment compatible with the older technology.

The basic algorithms used for data processing meanwhile often need to be maintained in the same state throughout the lifetime of an Earth-observation mission, which can last for many years. The in-orbit instrumentation always remains the same, although it may degrade slightly. This means that the original software and algorithms need to be kept for several years, whilst the computer hardware needs regular replacement.

The process of changing the computer equipment while maintaining the software and algorithms is called 're-hosting'. It means that hardware is replaced by new equipment while the software remains, wherever possible, identical. The main problem in this process is that new hardware often requires a new intermediate layer, the Operating System, where basic input and output and most importantly mathematical functions are implemented in different ways. This leads to situations where the original software needs to

be updated although the intention is to keep it unchanged. It also means that thorough validation and testing is required to prove that the re-host generates output indistinguishable from that of the original system.

Open-source software

The recent evolution in networking using the Internet has enabled software developers around the World to co-operate in the so-called 'open-source' movement. This has been based on the GNU General Public license for software whereby anyone has the right to modify the code and re-distribute it as they see fit. The Linux operating system is one example of this international co-operation. It is technically a Unix kernel managing the computer hardware and networks that was first developed for the PC architecture. It is now also available for many other platforms.

It is possible for anyone, through the licence, to acquire the source code of Linux by downloading it from an Internet server. In practice, however, is it more convenient to buy a CD-ROM on which the source code is distributed together with support programs and applications.

The fact that the source code is open and freely available makes modification and tailoring for a specific application possible. The code can be changed and re-built for the specific environment by the teams building the application. This means also that the application when ready is more amenable to future re-hosting, if required.

The GOME Data Processor re-hosting

The GOME Data Processor (GDP) at the German Processing and Archiving Facility (D-PAF) in Oberpfaffenhofen generates data products for the monitoring of the Earth's atmosphere. This processor has been in operational use since the launch of ESA's ERS-2 Earth-observation mission in 1995. The output products derived from the Global Ozone Monitoring Experiment (GOME) instrument provide global information about the atmosphere's content of trace gases like ozone.

The GOME algorithms and processing software used by the GDP have evolved over the years since the start of operations, due both to increased knowledge of the instrument's capabilities and to innovations in processing techniques. This led to data acquired at different points in time being processed with different versions of the data processor, leading in turn to small but significant differences in the product output. However, the scientific users of the products looking for global trends and changes ideally require a coherent data set processed with the same software version covering the full mission. The requirement for re-processing capabilities was therefore acknowledged by the Agency during 1998.

The GDP at that time was coping successfully with the processing of new data, but could not simultaneously support extensive re-processing. ESA therefore asked the institution hosting the processing facility, namely the Deutsches Zentrum für Luft- und Raumfahrt (DLR), to examine the possibility of upgrading the processing capacity. It was soon apparent that there were two possibilities, to extend the system with the same brand of hardware or to replace the system with low-cost PCs running Linux. The second possibility offered a price/performance ratio an order of magnitude better than the more traditional first alternative.

The fact that the original processing system used hardware running a Unix operating system helped at this stage when the effort to re-host was analysed. Linux, being a flavour of Unix aiming for standard compliance, supported the application well without requiring major coding changes. The Agency therefore finally selected the PC-based solution whereby twenty identical units would be assembled into a 'PC-farm' to execute the most demanding parts of the data processing (Fig. 1).

The original software already supported an architecture in which the processing could be spread over a network of computers. The main advantage of this was that one unit could fail without disturbing the others and that the overall performance would only be slightly degraded.

The scope of the project at this stage was:

- Re-hosting of the most demanding part of the application, the data-processing engine, starting from calibrated instrument data and ending with user products containing geo-physical information.
- Maintenance of all existing external interfaces by keeping the system front-end as it was without changing the hardware.
- Implementation of memory handling to handle the differences in processor architecture correctly.

The GDP re-hosting project now consisted of two major activities related to the application software: the original software had to be re-built for the new Linux environment, and the differences in processor architecture had to be handled correctly. The software re-build demonstrated the fact that Linux and software

Figure 1. The Linux farm for GOME data processing at the German Processing and Archiving Facility (D-PAF)

designed for Linux follows accepted standards more strictly than other commercially available products, where machine-specific extensions are often introduced. The tool translating the code to machine instructions, the compiler, in this case showed this behaviour. The original code used many of these specific extensions, which the Linux compiler did not support. A complete inspection of the code was therefore needed to remove all such extensions. This also had the beneficial effect that several previously undetected errors were spotted and also corrected.

The second major task was the implementation of the memory handling, the so-called 'byte swapping'. The new PC architecture differed in its memory organisation from the existing hardware. The project therefore implemented a byte-swap routine at an early stage that corrected for this.

After these two activities had been completed, the project proceeded smoothly and the application software could be compiled and executed on the new PC hardware just 60 days after the kick-off. Comparisons between outputs from the new and the existing chains showed very small relative differences, all within the level of mathematical precision, but very few cases where the relative difference was high but the absolute difference was low, because the computed values were close to zero.

The D-PAF now uses the re-hosted GDP for its GOME data processing. The first re-processing campaign took place during autumn 1999. The processor generated user products for all ERS-2 GOME data over a number of months, achieving a speed whereby one year of data got processed in one month, despite several hardware problems. This can be compared with the previous performance of processing two months of data in one month with no operational failures.

The UK-PAF re-hosting

The United Kingdom Processing and Archiving Facility (UK-PAF) has provided data-processing and product-generation services for the user community since the first ESA Earth-observation mission, ERS-1, was launched in 1991. The products have been derived from the source data of the main ERS instruments, i.e. the Radar Altimeter (ALT), the Synthetic Aperture Radar (SAR) and the Along-Track Scanning Radiometer (ATSR). The architecture known as the Earth Observation Data Centre was first assembled to support the ERS-1 mission. It was subsequently extended during 1994 and 1995 to support ERS-2.

One of the main UK-PAF data products is the Radar Altimeter Waveform Product (ALT.WAP), which contains the processed RA telemetry as well as important annotations such as correction factors for atmospheric conditions. This is a global data product offering a unique data set covering oceans, ice- and land-masses, which no other Radar Altimeter data set is capable of delivering. The ALT.WAP processing chain is located only at the UK-PAF and its hardware has not been updated since the ERS-2 launch.

In 1998, the Agency began replacing obsolete hardware in its Earth-observation data-rocessing facilities, while also ensuring year-2000 system compliance. The UK-PAF host, the National Remote Sensing Centre Ltd. (NRSC), carried out a study of the implications for their facility. It was found that the architecture from 1995 utilised computer hardware that was obsolete such that it made maintenance very expensive and the replacement of peripheral equipment like magnetic tape devices almost impossible. Moreover, the operating system was not year-2000 compliant. It was also recognised that

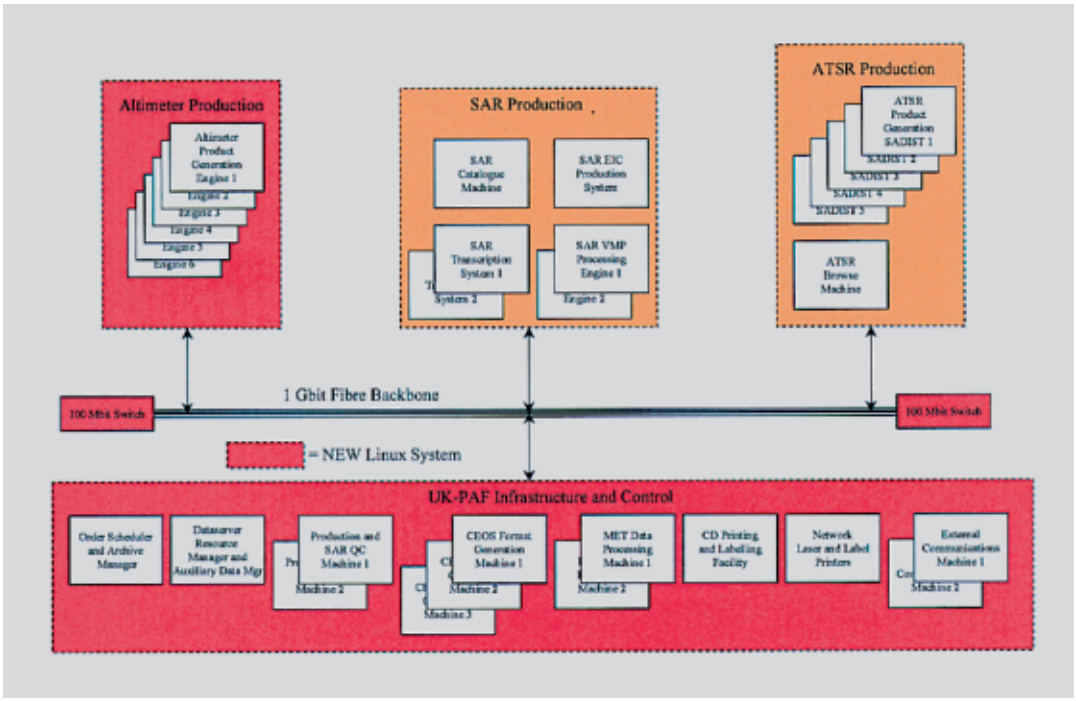

Figure 2. The system architecture of the UK Processing and Archiving Facility (UK-PAF)

critical subsystems did not have any redundancy, which would cause long periods of unavailability in case of failures.

All of these facts led to an NRSC study whereby the possibility of an upgrade was considered. The alternatives were soon recognised, as in the case of the D-PAF, as either replacing the hardware with a modern version from the same manufacturer, but still using the same basic software, or replacing the hardware with PCs running Linux, which would also imply system re-hosting. The fact that the operating system in use at the UK-PAF was an older version complicated the alternative of a simple hardware replacement, as the existing operating system would not execute on any new platform available. This meant that both alternatives would effectively lead to a re-hosting, and the study showed that the effort required would be of the same magnitude in both cases.

The success of the D-PAF re-hosting, together with the significantly better price/performance ratio of the PC-based solution combined with much lower maintenance costs, led to the selection of this alternative and the re-hosting project was started in late spring 1999. It included all external interfaces, the control of the SAR and ATSR processing chains, and the full ALT.WAP processing. The SAR and ATSR processing chains were not included as they are delivered by other entities for the UK-PAF's use.

The project elected to maintain the existing architecture, whereby discrete functions are statically distributed over a network of computers (Fig. 2). The lower hardware cost allowed generous redundancy and each function could thus reside on at least two computer hosts. This, combined with the use of RAID (Rapid Array of Independent Discs) technology allowing secure data storage and redundant local-area networks, removed all possibility of outages due to a single failure (Fig. 3).

The re-hosting project not only aimed at replacing obsolete hardware, but also had the goal of securing UK-PAF operations in the year 2000. This meant that the schedule had to be very ambitious, to allow startup before the end of 1999. The project was therefore divided into three phases, with the most critical functions to be addressed first, thus securing at least limited use. The first of the phases secured the

Figure 3. The scheduler subsystem of the UK-PAF. The black unit on the left is the RAID unit, providing full data redundancy

Figure 4. One of the UK-PAF data-processing sub-systems. The unit on the left is a robotic cassette-handling facility enabling continuous data processing

external interfaces of the UK-PAF to enable at least manually controlled data processing. The second integrated the functions of controlling the SAR and ATSR processing chains, thus allowing automated processing for these sensors. The last provided the fully automated ALT.WAP product processing.

The project experienced some initial delay due to problems with one of the Custom Off-the-Shelf Software (COTS) products being used. This item provided the management of the UK-PAF databases and had just been released as a native version for Linux. The project initially successfully used a preview version that was not officially supported, but encountered problems when trying to use a later supported version. These problems proved so serious that the preview version was actually recovered and used in the final operational system.

The first two phases, allowing SAR and ATSR product generation and supporting all external interfaces, were finished just in time for Christmas 1999 and the UK-PAF successfully completed the year 2000 roll-over. The final phase did, however, encounter problems that delayed the re-starting of ALT.WAP product generation.

The ALT.WAP re-hosting differed from the rest of the project in that it touched upon the data processing itself, and not just the processing control. This meant that the issues of memory organisation and mathematical precision also encountered in the D-PAF case became valid. The first of these caused some problems and delay, but were resolved quickly in a similar way to the D-PAF when a general byte-swapping strategy was implemented. The latter, on the other hand, created problems that were more difficult to understand. The values in the output products differed from those of the original chain in a way that could not be justified. Finally, a single instruction, namely a numerical division, appeared to be critical where the rounding systematically differed compared to the original system. This error was then propagated throughout the chain to the output results. The difference could thus be explained and accepted and the processor became ready to use.

The UK-PAF is now fully operational for all production chains ten months after project kick-off. The new systems perform according to expectations and well above the previous system levels (Fig. 4).

Re-hosting achievements and benefits

The two Linux re-hosting projects described in this article have both shown the excellent performance of this low-cost alternative. It is evident in particular that the possibilities of multiplying the hardware in a way that traditional alternatives do not permit raises the performance significantly. This means in the end that much higher throughputs can be achieved for a similar or even a lower cost.

The re-hosting to Linux on PCs requires specific attention to two issues: the differences in memory organisation between different hardware types, and the differences in mathematical precision. Re-hosting projects should therefore identify and establish strategies to handle these issues in the early stages, in order to avoid surprises when the new outputs are compared with those of the existing systems.

The two implementations have also demonstrated the reliability of the Linux operating system. There have been very few reported outages and all have been related to hardware or other components outside the control of the operating system. Projects implementing this technology should therefore select hardware with a good reputation for reliability in order to benefit fully from the good performance of the Linux operating system itself.

All things considered, Linux is proving a very good candidate for use as the operating system of choice in Earth-observation data-processing facilities. esa