Recent developments and operational experience at IM of using Paipix for integrating meteorological systems

António Amorim [1] and Lígia Amorim [2]

[1] SIM/IDL and Fac. Ciências da Universidade de Lisboa, Ed. C8, Campo Grande
[2] IM – Instituto de Meteorologia, I. P. , Rua do Aeroporto, Lisboa

March 12, 2008

Abstract

The PAIPIX software project, that started as the remastering of live Linux, has been upgraded to provide full support for system installation, maintenace and upgrading of server systems. The new developments were aimed at maintaining operational systems and have been used in the new data server infrastructure at Instituto de Meteorologia. The services and applications provided by PAIPIX are discussed together with the main steps involved in the development of the software package repositories and creation of the full DVD images.

1 Introduction

The challenge to provide meteorologists with an integrated computer environment both for data processing, operational forecast and software development becomes difficult to meet for institutes that do not have separate teams allocated to the maintenance of workstations, of database servers, of data processing clusters and of high performance computing infrastructures. The number of different services that have to be supported tends to be constant and does not scale with the organization size. At Instituto de Meteorologia this problem is being addressed by the PAIPIX system based on the Debian distribution that integrates both the desktop environment, the data processing servers and the database infrastructure.

PAIPIX is a free software project that is being developed by SIM at the University of Lisboa that maintains a special restricted version, including all the required meteorological tools, to be used by meteorologists. The project has started as a remaster of live linux distributions that were common several years ago and has evolved into a distribution, with full support for hard disk installation and upgrading ,that is available both for Intel/AMD 32 and 64 bit architectures. The PAIPIX development work includes maintaining a special debian repository of scientific applications that is freely available and a specific repository for applications that are restricted to the meteorological institutes[1]. From these repositories and our debian application mirrors, a full list of ISO images is built nightly each including both a live compressed file system, a disk installer and a DVD based repository of all the required applications. This paper provides a description of the evolution of PAIPIX from the initial live system, towards a solid 32 and 64 bit Debian based installation and upgrade enabling tool. The new system is both aimed to be used at the individual workstations and at the institute main data processing servers. Section 2 describes the infracture and the services implemented at Instituto de Meteorologia and discusses the main supported applications. In section 3 the PAIPIX application repositories are discussed together with the packaging system and the debian infrastructure that supports maintenance and upgrading of the different applications. The DVD image building system and the automated testing mechanisms are described in Section 4.

2 Services and Infrastructure

The operational infrastructure that is running PAIPIX at IM consists of 12 servers each with two Intel Xeon Quad core processors, 8 GB of memory and about 10 TB of disk space. It is using hardware raid controllers over SATAII disks of 750GB each. This infrastructure, depicted in Fig. 1(left), makes available a total of 96 computing cores an about 120TB of disk space. The choice to avoid a central storage system was taken not only to minimize acquisition and maintenance costs but also to define redundant schemas using machines that are as much independent as possible. The system is becoming operational together with a IBM system p5, composed of a cluster of 10 nodes with 16 CPU's each, mainly intended for running mesoscale forecast models.

[1]Can only be used under license from ECMWF

All the hardware, based on Intel processors, is running the restrict distribution of PAIPIX 8 that includes a complete set of meteorological applications. The PAIPIX system was installed using a preconfigured and fully automated network installation departing from an instaler set up using BOOTP/FPTP and a PAIPIX image that is made available as a package repository from the principal node. The full installation of a single node, including disk partitioning and formating takes less than 1 hour and does not require any human intervation. This installation procedure was developed not only to support the initial setup of the 12 servers but also to minimize the dead time in the eventuality that a machine needs to be re-installed.

The machine configuration and monitoring applications, including the Open- Manager from Dell, were made available in PAIPIX and the whole hardware and software infrastructure is being continuously monitored [6](see Fig. 1 center and right).

The set of applications in PAIPIX is mostly selected form the stable Debian [3] distribution upgraded by using the "back-porting" of very recent applications to stable framework as is the case of the linux kernel that is set to version 2.6.22.

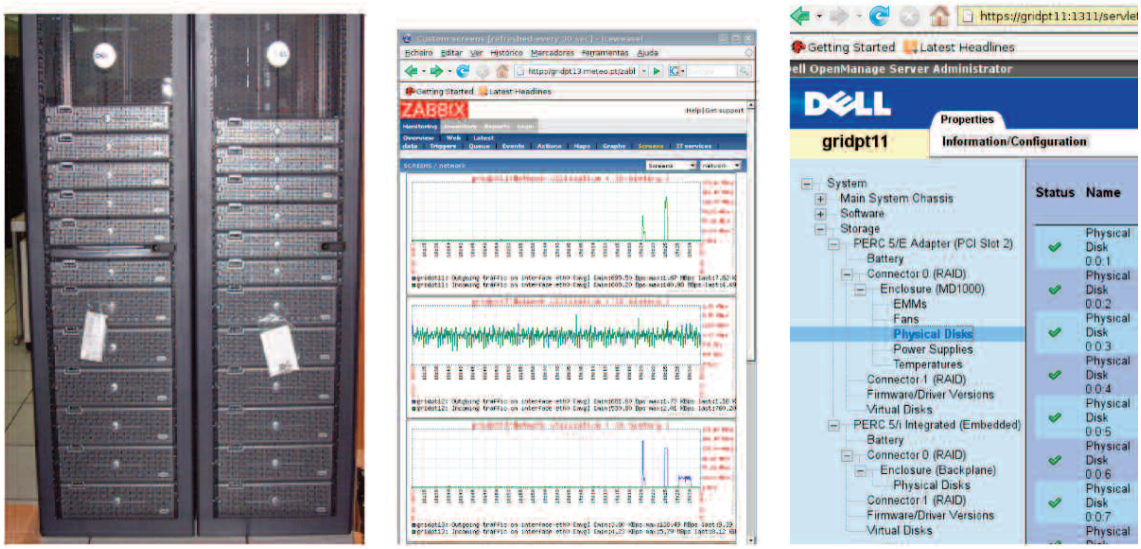

Fig. 1 The PAIPIX server infrastructure at IM, I.P. (left). The monitoring services and hardware administration interfaces (center and right).

2.1 Database servers and TIDB2

The TIDB2 scientific database interface developed with our group [4],[5] is used to manage the meteorological data. The TIDB2 framework includes both the main C++ programing interface and C and Fortran bindings. A set of tools to store and retrieve selected information from the database is also included as well as the ktidbexplorer interactive data browser. All these components are availably by default in the PAIPIX installation both for 32 and 64 bit architectures. The GRIB and BUFR plugins are compiled on top of the ECMWF gribapi [9] and bufr [10] libraries.

The TIDB2 database interface isolates the application form any specific database management solution. We have chosen to use the MySQL Open Source database servers that seam to provide better support for databases with a large number of intermediate size BLOBS. This is the case for the managing of GRIB and BUFR data in TIDB2. The MySQL database main and backup servers are running in two of the cluster dedicated machines with access to the large hardware RAID volumes that are interfaced trough local RAID controllers. The database server configuration was optimized for performance by increasing substantially the "key buffer", "table cache" and "query buffer" spaces.

To achieve convenient scaling of the access times to the databases with increasing amounts of data being stored, the data is automatically partitioned among different databases tables according the time period (year) and to the specific values of certain GRIB or BUFR parameters. The allocation of different objects to different tables is done using a special database table that automatically specifies the conditions associated with the objects to be managed by different tables.

2.2 *SMS and Subversion operations*

The SMS [7] server and xcdp, cpd client applications are installed in each machine with PAIPIX (meteorology version). The PAIPIX package automatically creates a local sms user and installs a system startup script that will invoke the sms user scripts for starting up and shutting down SMS. The SMS package includes a template suite that can bootstrap a subversion [8] repository copy for the sms user. It allows any user, with appropriate access to the subversion repository, to commit code or task scripts to sms. This suite defines tasks to checkout the repositories and to invoke configure, make and make install at the repository main directories.

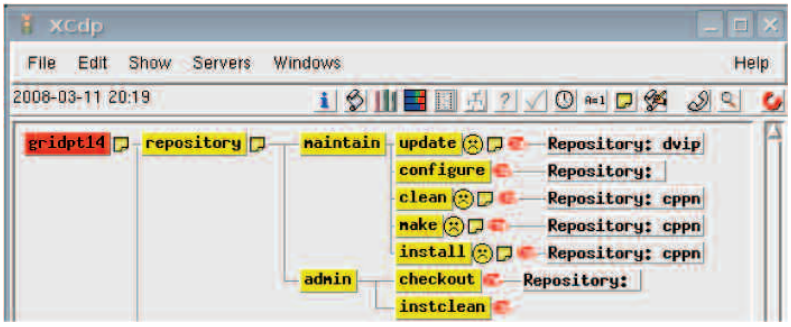

Fig. 2 The SMS suite that "bootstraps" the access to the Subversion repository.

The data server infrastructure at IM supports the operational SMS server that runs on a machine that will execute locally most of the operational procedures. This machine is not intended for user login thus defining an architecture where all operational applications must be submitted to the source code repository and compiled by the sms user on the dedicated machine. For testing purposes, the machines supporting user login also run local SMS servers.

2.3 *Visualization, data processing and software development*

The set of applications included in the public version of PAIPIX already cover most of the software development tools used in scientific applications. The gnu C++, C and gfortran compilers together with the Java Development Kit from SUN are already included in all live and installed systems. In parallel all application interfaces for meteorology data processing and visualization whose license has recently become free like Magics++, GRIB API, GRIBEX, BUFR, EMOSLIB, Vis5D and GRADS, are also standard in the public and special versions of PAIPIX.

There is however a larger set of very powerfull applications with restricted licenses that are commonly used in meteorological services and that have also been ported to PAIPIX. These include for visualization Metview and Magics and for software development, the Portland compilers. These applications can be accessed as PAIPIX/debian packages including all the necessary package dependencies in order to automatically install the extra required software.

Building both the free and non free packages is performed using the debian "pbuilder" [11] tool that compiles each package by first installing a minimal system as a separate "chroot" and then installing all the required development packages in other to provide each time a clean compilation environment. Building these packages required the development or porting of auxiliar packages like Open Motif that have then become part of the PAIPIX package repository.

In the case of the 64 bit systems, like the ones installed in the IM data server infrastructure, several applications can only be compiled using 32 bit mode (- m32 flag). The debian amd64 system used by PAIPIX is a pure 64 bit system requiring the 32 bit support libraries to be available as special compatibility packages like the debian ia32-libs or the PAIPIX ia32-libs-paipix packages.

3 Packages and Application Repositories

Users that would not like to bind themselves into the PAIPIX complete set of applications can still make use of the additional package repositories if they are using a Debian based system. The necessary steps involve:

a) Adding the PAIPIX repositories to the /etc/sources.list file by appending the lines:
   - deb http://mirror.sim.ul.pt/debian-paipix etch main contrib non-free # for the PAIPIX public repository
   - deb http://mirror.sim.ul.pt/debian-specific etch main contrib non-free # for the PAIPIX repository of restricted packages[2].

b) Install the applications by using the "apt-get update" and "apt-get install" commands.

[2]Only made available upon request to the ECMWF

4 Building and testing of the PAIPIX system

The new PAIPIX system is built using a modified version of the debian "livehelper" application. The final DVD "iso" files contain both the PAIPIX live system with all basic applications and an installer system that is tunned to automatically install in hard disk all packages existent in the live system. Many extra packages, like most of the services, can be installed after the system boots, either live or from hard disk, by using the PAIPIX install menu.

The PAIPIX DVD "iso" building infrastructure, based on live-helper, runs every night in order to produce updated images for all the PAIPIX available languages. In parallel, the meteorology restricted PAIPIX-IM systems are also built. This process takes about 12 hours in a two Xeon Quad core machine, from the SIM-University Lisboa cluster that is also running the PAIPIX system.

Following the upgrade work carried out on the PAIPIX development infrastructure, with the migration to the debian live-helper system, the amount of manual intervention needed to maintaining upgraded versions of the PAIPIX images was greatly reduced. However applying an exhaustive test suit to all the images both running as live or installed on hard disk remains a very time consuming activity that must precede the publication of a new release of PAIPIX. To minimize the work involved in testing, all live and install tests are now performed on a server machine by using the open source -Kernel-based Virtual Machine- kvm [12] emulation system.

The main task that was not yet made automatic consists of the set of procedures that must be used to test the GUI applications. The Java application vncrobot [13] is presently being evaluated carry these tests automatically.

5 Conclusions and Outlook

A new data server infrastructure running PAIPIX-IM 64 bits was installed at IM and is being used for data processing and development tasks. The installation and maintenance tasks are extremely simplified by maintaining all the required software in debian/PAIPIX packages. Exact copies of the same computing environment can be easily obtained in the developer workstations by installing locally the same PAIPIX system, allowing the full integration of the server and workstation environment.

The PAIPIX system, that was initially developed for personal workstation usage, has been upgraded over the last two years, to enable a data server infrastructure with full support for installation and system upgrade. With the new more sustainable software framework and the very powerfull hardware in place, the future developments will focus on introducing innovative applications and products as well as on establishing a permanent fully automatic testing facility. The Free PAIPIX will continue to incorporate all tools that will become available to support Universities, in particular in meteorological related studies.

All the work on the PAIPIX system has been very much concentrated in a single developer but the new system was built on top of an infrastructure that can easily support work done in parallel by many developers, not only in creating new debian science packages but also in defining the set of applications to be used and the testing procedures to be applied.

References

[1] http://www.paipix.org

[2] Making the ECMWF tools, including METVIEW, available in a restricted version of the PAIPIX scientific live Linux., António Amorim and Lígia Amorim, Tenth Workshop on Meteorological Operational Systems, European Center For Medium-Range Weather Forecast, Reading, Nov. 2005, ECMWF Proceedings Series, July 2006, page 111

[3] http://www.debian.org

[4] "Implementation and testing object extensions of open source RDBMS for meteorological data", João Simões, António Amorim and Maria Monteiro, Tenth Workshop on Meteorological Operational Systems, European Center For Medium-Range Weather Forecast, Reading, Nov. 2005, ECMWF Proceedings Series, July 2006, page 82

[5] Experience with the TIDB2 database interface in managing meteorological observation and forecast data, João Simões, António Amorim and Maria Monteiro, Eleventh Workshop on Meteorological Operational Systems, European Center For Medium-Range Weather Forecast, Reading, Nov 2007.

[6] http://www.zabbix.com/

[7] http://www.ecmwf.int/products/data/software/sms.html

[8] http://subversion.tigris.org/

[9] http://www.ecmwf.int/products/data/software/grib api.html

[10] http://www.ecmwf.int/products/data/software/bufr.html

[11] http://pbuilder.alioth.debian.org/

[12] http://packages.qa.debian.org/k/kvm.html

[13] http://www.vncrobot.com/