

# The Paraná Digital Project

Marcos Castilho  
Center for Scientific Computing and Free Software  
C3SL/UFPR

September, 25th 2008

Managing a Grid of Computer Laboratories  
for Educational Purposes

# Abstract

- ▶ We present a grid-based model for managing hundreds of distributed educational computing laboratories.
- ▶ The aim is to dispense the need of specialized staff in each place whilst maximizing the laboratories uptime
- ▶ The model was implemented in order to manage more than two thousand laboratories spread across the entire public school network of Paraná State (Brazil)

# Motivation

- ▶ Computers and the Internet are important for educational purposes
- ▶ Cost reduction of computer hardware led to expansion of digital inclusion policies
- ▶ The aim is to provide laboratories for every public school

# First problems

- ▶ The lack of specialized workforce to manage and maintain these laboratories frequently turns them useless
- ▶ Computers get obsolete, the software doesn't work properly, computers get infected, among other reasons
- ▶ In public schools this situation is worsened by the lack of funds for frequent updates

## Problems that seem to have no solution

- ▶ It's impossible to have an expert to manage every school
- ▶ Educators and students should only worry about the educational use of the system, and not about its technical details
- ▶ This demands an effort in order to define a new administration model

# Problems do have a solution!

## The Grid-based model

- ▶ Management tasks are automated
- ▶ Self-monitoring systems
- ▶ Maintenance costs are minimized
- ▶ Dispenses specialized staff in every laboratory
- ▶ Maximises system uptime

## The model was implemented in Brazil

- ▶ 2.100 schools
- ▶ 1.500.000 students
- ▶ 57.000 teachers
- ▶ 42.000 working points
- ▶ 399 cities
- ▶ 199,314km<sup>2</sup>
- ▶ GNU/Linux software, under Free Software licence (GPL)
- ▶ Up and running since June, 2006

# The Paraná Digital Project

- ▶ It is a partnership between
  - ▶ State Secretary of Education
  - ▶ State Computer Company of Paraná State
  - ▶ State Secretary of Science and Technology
  - ▶ Inter-American Development Bank (IDB)
  - ▶ United Nations Development Programme (UNDP)
  - ▶ Electric Company of Electricity (COPEL)
  - ▶ Universidade Federal do Paraná (C3SL)



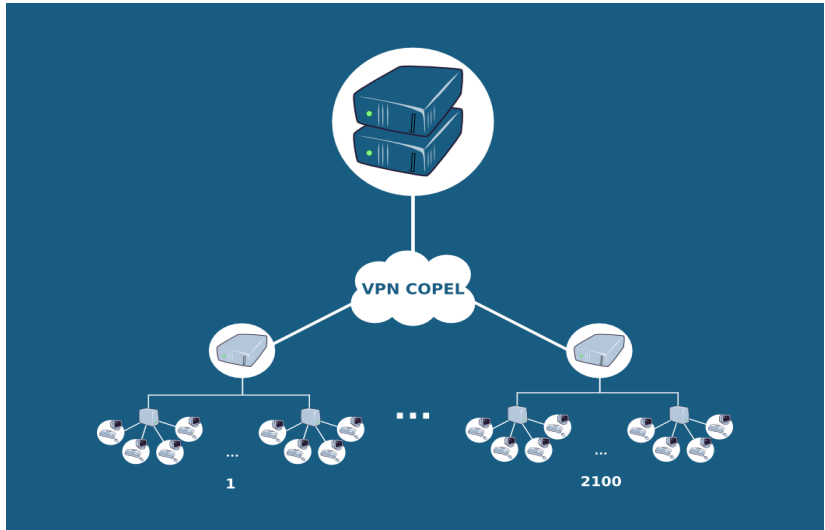
# View of Paraná State in Brazil



# View of Paraná State



# PRD's architecture



## PRD's architecture

- ▶ In each school:
  - ▶ a laboratory composed of 20 X-terminals
  - ▶ one processing server called *the school server*
- ▶ The school server acts simultaneously as
  - ▶ processing and storage unit, gateway to the network, firewall and access point to the Core
  - ▶ it runs a Debian-based GNU/Linux distribution,
  - ▶ all servers have the same software packages installed.
- ▶ At the Core, a proxy-controlled connection to the Internet is provided

# The grid-based model

- ▶ The computer laboratories are interconnected forming a grid
- ▶ The hardware, operating system and main applications are fairly homogeneous
- ▶ Global tasks: are executed to guarantee the expected services, with maximum performance
- ▶ Local tasks: are related to specific aspects of each laboratory and those demanding physical access to the hardware
- ▶ Main idea: To implement self-management concepts

# The local managers

- ▶ They are not computer experts
- ▶ They take high level and simple decisions
- ▶ They use a user-friendly interface to perform local tasks:
  - ▶ User accounts, disk quotas, etc.
  - ▶ Initial installation or re-installation procedures
  - ▶ Contact the call center in case of problems

## Global managers: the Core

- ▶ A small group of specialized system managers
- ▶ They are able to manage hundreds of laboratories
- ▶ They specify what are the system's features, software to be installed, what should be restricted or allowed
- ▶ Their decisions are propagated through self-management
  - ▶ Self-configuration
  - ▶ Self-optimization
  - ▶ Self-recuperation
  - ▶ Self-protection

# Self-configuration

- ▶ Concerns the adaptation of new components or new execution environments without a significant human intervention
- ▶ Is a continuous process aiming to keep the system configured
- ▶ The configuration policies are defined for the entire grid by the Core



# Implementing self-configuration

- ▶ Users have access to the server through the X-terminals
- ▶ The server is the only machine to be configured and maintained
- ▶ The software and configuration of the X-terminals are determined at the server
- ▶ The servers' self-configuration is based on software package management and unified mirrors

## Initial installation

- ▶ A CD-ROM containing a standard system image
- ▶ The server checks periodically for updates at the central mirror
- ▶ X-terminals must automatically recognize and configure hardware
- ▶ It allows automated installation and configuration by the local manager
- ▶ This automated procedure minimizes service disruption and data loss
- ▶ A regular user must have an account created by the local manager.
- ▶ The local manager does not have root powers
- ▶ This ensures that critical tasks are globally defined on the grid.

# Self-optimization

- ▶ Allows the configuration of parameters that impact on the overall performance
- ▶ Monitoring systems are fundamental in a self-optimized system
- ▶ Collected data can be used to optimize system parameters
- ▶ New configurations can be automatically applied globally in the grid
- ▶ The identification of possible failure points can be used to avoid or to fix bugs

# Self-optimization

- ▶ The analysis of historical data can also be used to optimize the system
- ▶ This helps to scale the hardware in each laboratory or to plan a hardware upgrade
- ▶ There is a web page with strategic information from the whole grid

# Self-recuperation and Self-protection

- ▶ Failures in a system can cost many weeks of work
- ▶ It's possible to automatically detect, diagnose, treat and prevent many software and hardware problems
- ▶ The downtime can be reduced if some aspects of the hardware and the operating system are tracked
- ▶ Examples: Contact the Core depending on memory, hard disk and temperature sensors

## Self-recuperation and Self-protection

- ▶ Redundant Array of Inexpensive Disks (RAID) prevent loss of data
- ▶ Downtime can be minimized by verifying the filesystem's integrity
- ▶ The system must automatically recover damaged files
- ▶ There must be a recuperation process to reinstall the system without user data loss
- ▶ If everything fails, enter in remote assistance mode (ssh from the Core)

# System Recovery

- ▶ Automatic recovery: it formats only the root filesystem, preserving user data and local information. A collection of scripts saves important files, restoring them after system re-installation.
- ▶ The remote help: Live-CD allows the Core to connect to the school server and try to find the problem or to do a backup.

# System Upgrade

- ▶ Frequent system upgrades are necessary to provide new functionalities, address security problems and propagate new software, tools, or policies from the Core
  - ▶ automatic daily upgrades and
  - ▶ triggered upgrades



# The daily automatic upgrade

- ▶ Based on Debian's *apt-get* tools
- ▶ Every night, each school server looks for new software packages
- ▶ The single mirror ensures that all servers will install exactly the same software
- ▶ It is quite simple to propagate a new tool or configuration over the entire network

# The triggered upgrade

- ▶ The Core can force all school servers to upgrade
- ▶ This happens as soon as the network link turns on
- ▶ Use cases:
  - ▶ the Kernel exploit that allowed an ordinary user to become root
  - ▶ the upgrade from *sarge* to *etch* (in February 2008)

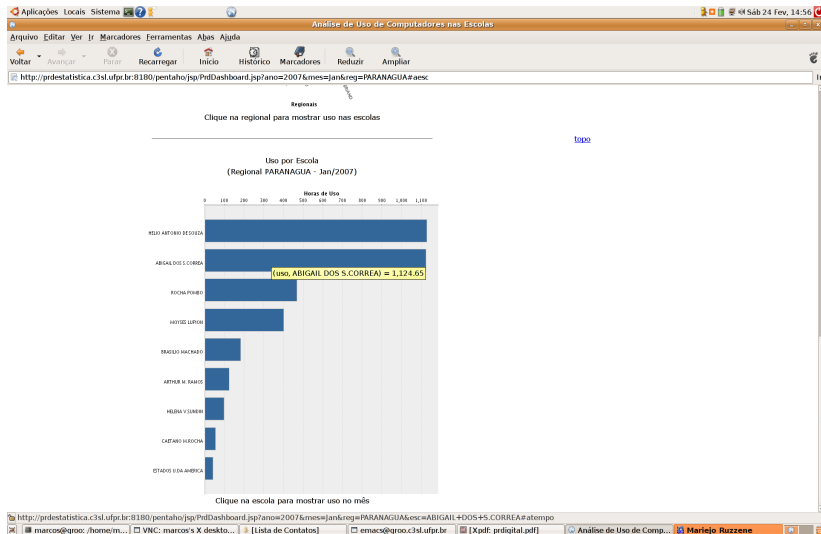
# System Monitoring

- ▶ Is an essential feature of autonomic systems, and provides information to allow the system's self-optimization and self-recuperation
- ▶ It is at the heart of the PRD network, revealing the real state of the whole grid
- ▶ In the PRD model, there are two different monitoring systems:
  - ▶ the statistics center
  - ▶ diagnosis system

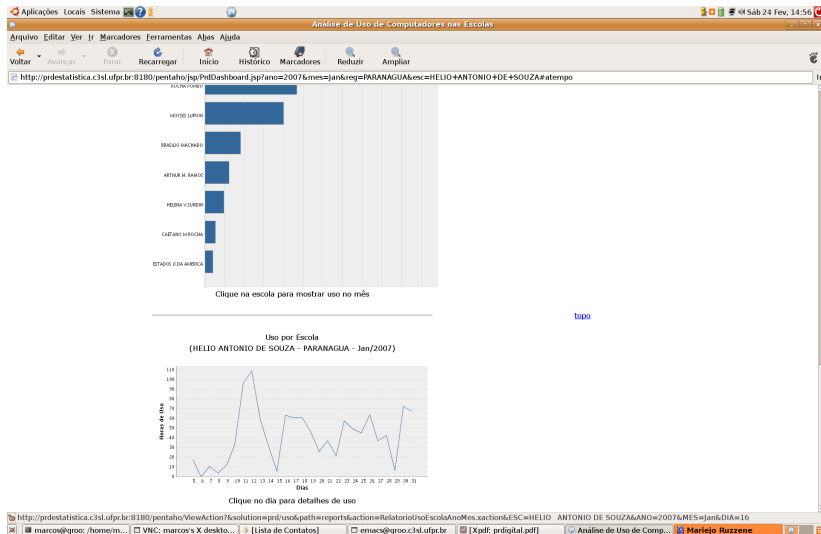
# The Statistics Center

- ▶ A web site with strategic information
- ▶ It allows an overview of the network's growth and provide data concerning the laboratories usage
- ▶ This is automatically stored in the central database
- ▶ This provides strategic information for decision support


# Snapshot of the statistics center web page



# Snapshot of the statistics center web page



# Examples of collected data giving decision support



# The Instant Diagnosis System

- ▶ Conceived to provide real time diagnosis of the network's status
- ▶ Detects possible faults in the laboratories and triggers alarms at the Core
- ▶ Capable to detect whether the schools are online, identify old software, inconsistencies in a filesystem, checksum errors, hard disk problems and so on
- ▶ Allows prevention of critical failures, either by acting remotely or by timely providing hardware replacement



# Manual System Inspection

- ▶ Manual interventions are undesired
- ▶ The core can remotely log in a server via SSH
- ▶ But just to inspect the server's behavior
- ▶ All changes should happen on a global manner in the grid
- ▶ This means: global management!

# Main contributions

- ▶ Software
  - ▶ 40 million dollars saved due to free software
  - ▶ Software is free and has high quality
  - ▶ More than 100 thousand packages are made available
- ▶ Hardware
  - ▶ Carefully dimensioned hardware configuration
  - ▶ “Obsolete” hardware can be reused
- ▶ 15 million dollars saved due to the multiseat
- ▶ Maximization of hardware lifetime

# Conclusion

- ▶ This model allows the administration of thousands of computing laboratories with minimum human intervention
- ▶ A large scale implementation of this model is provided, in which a management team composed of 12 highly trained Unix managers is able to control all software-related issues from the entire network

## Final considerations

- ▶ The system works pretty well . . .
- ▶ . . . but a working system does not mean it will be used!

## Contact and links

- ▶ [castilho@ufpr.br](mailto:castilho@ufpr.br)
- ▶ [jefferson.s@lactec.org.br](mailto:jefferson.s@lactec.org.br)
- ▶ <http://www.c3sl.ufpr.br/prd>
- ▶ <http://yoda.c3sl.ufpr.br/SDI>
- ▶ <http://www.prdestatistica.seed.pr.gov.br>