Scientific workstations using Debian

GNOME.Asia conference

April 2010

LEADING THE ENERGY CHANGE

EDF, a world energy leader

37,9 millions
of **clients** worldwide

169 139
collaborators worldwide

66,3 G€
revenue
49 % outside of France

618,5 TWh
produced worldwide

117,1 g
CO_2 per produced kWh

- World leader in nuclear power operation, Europe leader in hydro power
- **Strong Europe implantation**: France, UK, Germany, Italy…
- Europe leader in distribution, transport and sales
- Industrial operator in **Asia**: China, Laos
- Natural gas : a major player

Key figures

Revenue 2009

%

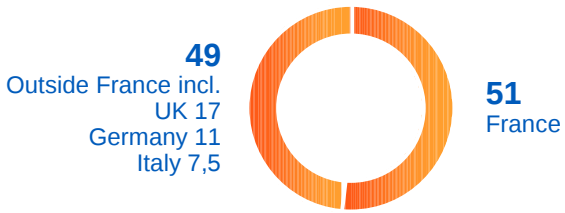

EBITDA 2009

%

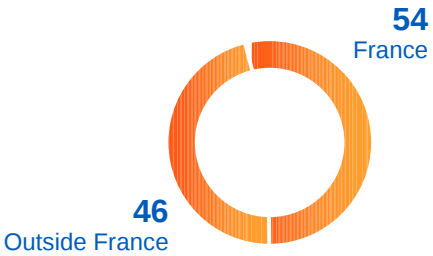

Production mix 2009

%

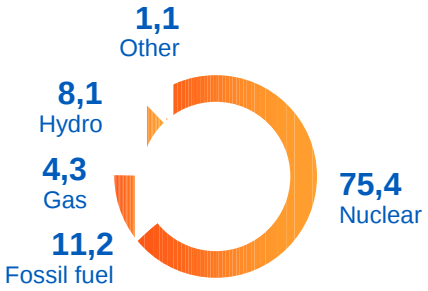

Total : 618,5 TWh

1. Why scientific workstations

What a scientific information system looks like

The role of scientific workstations in the company

Scientific computing at a glance

Modeling

- Approximate reality with a model
- Often need for a *modeler* to translate a specific case into machine-readable data

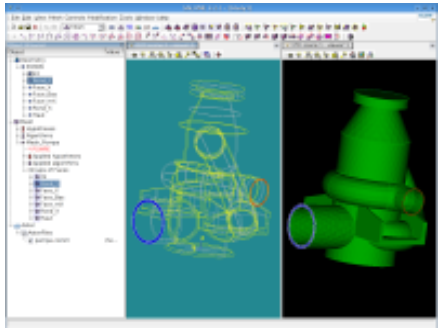

Simulating

- Execution of a numerical *code* computing the behavior of the model system
- A whole area of software development
- Need for the fastest hardware to work on *large arrays of floats*

Visualizing

- Results exploration and analysis
- Need for the best graphics hardware and displays

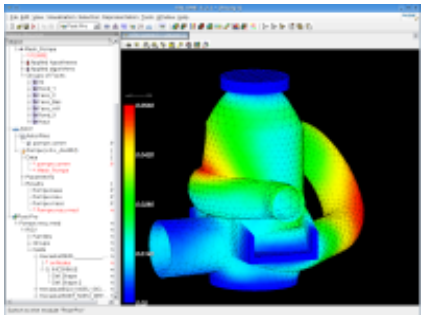

Why EDF needs scientific computing

R&D

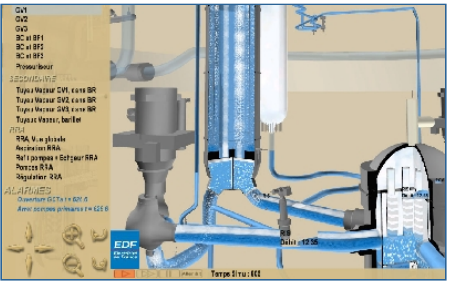

Conception

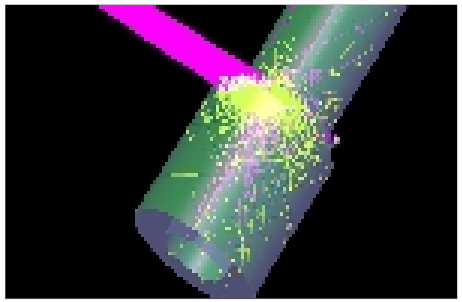

Information technology

Renewable energies

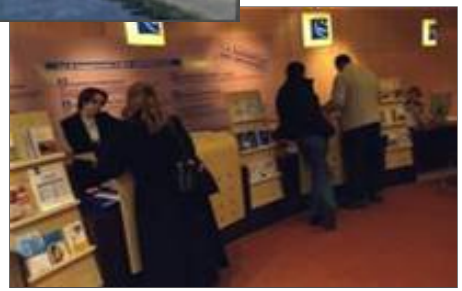

Sales

Engineering

Energy management

Electrical networks

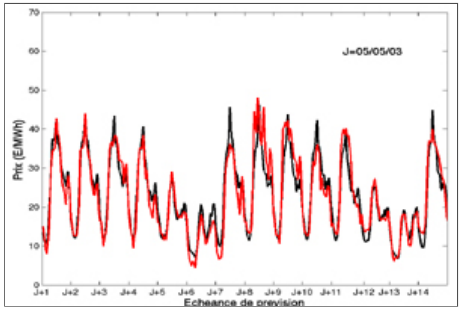

The role of a scientific workstation

- Running business applications
  - 1/3 in-house applications
  - 1/3 existing free software
  - 1/3 off-the-shelf non-free software
- Accessing computing resources
  - Preparing simulations
  - Running them remotely
  - Visualization and analysis
- Day-to-day engineering / research job
  - Productivity applications
  - Integration with the Information System

2. Our scientific workstations in 2011

~1000 high-end workstations and laptops
- **2/3 R&D**
- **1/3 Engineering**

The workstation offer

Hardware offer

- Workstation: dual quad-core + 12 GiB
- ~~Laptop~~ mobile workstation: dual-core, high-res screen
- High-end graphics

One single system: CALIBRE

- Rebranded Debian + a few applications
- Calibre 5 = Debian *etch*
- Calibre 6 = *lenny*
- Calibre 7 = *squeeze*
- **Same versions for clusters**

A comprehensive included software offer

- Standard desktop: GNOME
- Typical scientific applications for visualization and publishing
- Extensive on-demand installation for other software (incl. KDE)

Access to central scientific resources

- 8 clusters totaling ~400 Tflops
- Graphical cluster: 6×3 image wall, 24 Mpixels

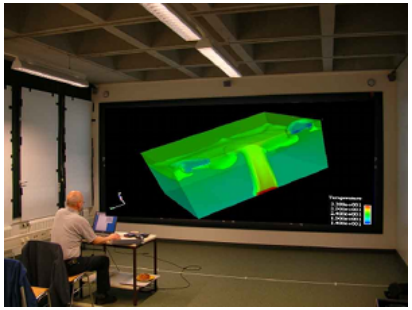

In-house development organization

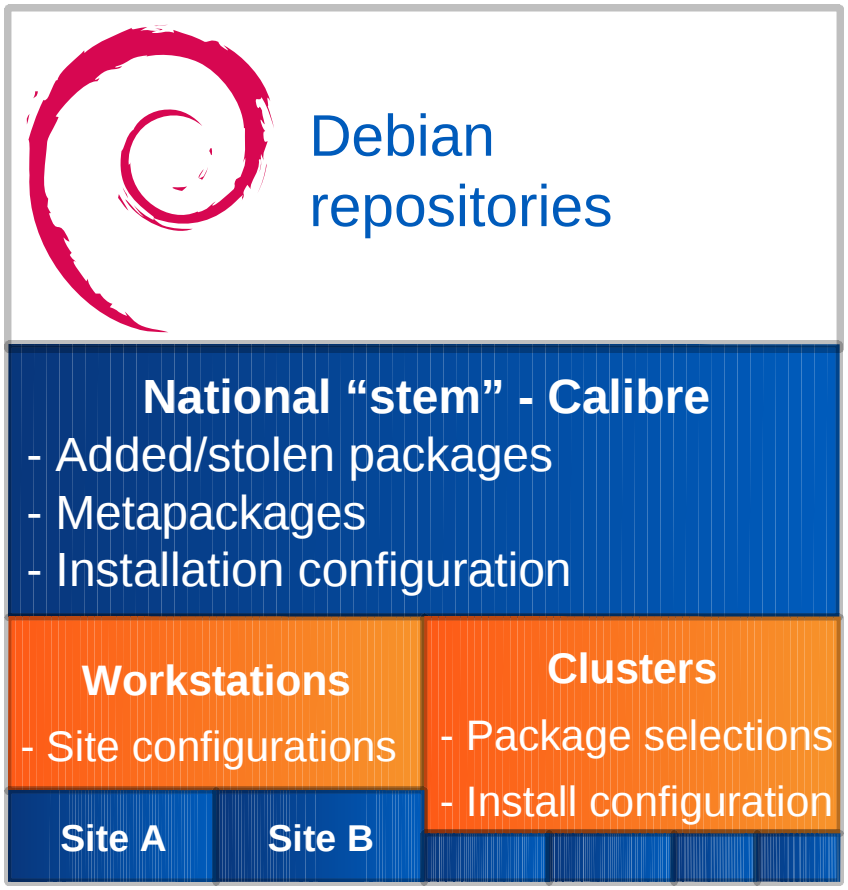

- One repository to rule them all
- Added packages:
  - Metapackages: only way to maintain consistency across upgrades
  - Backports / additional software
  - Configuration packages: violate policy hard
- Installation classes:
  - Hardware
  - Site
  - Basic package selections
- Site-specific infrastructure:
  - Authentication
  - File servers
  - Print servers
  - …

Support organization

- Dedicated scientific hot-line
  - Redirection from the centralized hot-line
  - Level 1 support
  - Help with basic usage / debugging shell scripts / …
- Level 2 support
  - Site integration
  - Workstation specifics
  - Engineering of upgrades
- Level 3 support
  - OS integration and upstream interaction
- Exploitation tools
  - Installation: FAI (Fully Automated Installer)
  - OCS Inventory NG / FusionInventory
  - Managed upgrades: PULSE 2
  - Many hacks in custom Debian packages
  - ***To be replaced by Puppet?***

3. Challenges

What happens when you introduce Linux desktops in a large company's network

Integrating with the information system

The good

- Network infrastructure
- Sharing data
- Printing

The bad

- MS Office documents & macros
  - OOo on Windows experimentation
- Corporate websites for IE6 only

The ugly

- Proprietary VPNs
- Bluecoat proxy
- Lotus Notes
  - Even MS Exchange would be easier
- Adobe Flash

Current and prospective solutions

- VMware player
  - Full Windows installation on each machine
  - Heavy and costly
  - Version 3 is extremely buggy
- Remote Windows access (RDP/ICA)
- Remote Linux access (NX/VNC)
  - No decent 3D support

Distribution life cycle and hardware support

- Lifespans:
  - A workstation: 3 years
  - An engineering project: 10 years
  - A nuclear reactor: **30-50 years**
- We need long release cycles
- Long-term security support
  - Currently done in-house: time-consuming and imperfect
  - Very hard to impossible for some desktop components

- Hardware qualification
  - Same hardware for 6-12 months
  - Manufacturers change specifications
  - Issues shared with Windows world
- Call for bids mechanism issues
  - “Of course it works on Linux”
  - Many manufacturers don’t answer
- Kernel obsolescence
  - Operating modern hardware with etch can be hard
  - High hopes for 2.6.32 long-term support
- Graphics drivers
  - So far, only nVidia
  - KMS will change that
  - Again, what about LTS?

Change resistance

- From other IT members
  - Only 1% of total workstations
  - A slow but strong tendency towards open standards
- From geek users / developers
  - They don’t like managed solutions
  - They like short release cycles
- From old-timers
  - Started with SunOS long ago
- From management?
  - They like managed and clean solutions
  - Scientific computing is strategic

Any questions?