

Mac OS X: a mature UNIX OS for (High Energy) Physics

Part 1: the UNIX workstation

O. Ciaffoni, S. Dell'Agnello, L. Di Cesare (INFN-LNF)
LNF, December 2, 2004

OUTLINE

- Mac OS X UNIX (v10.3, “Panther”)
- The Mac OS X workstation for HEP
- Developer tools
- OpenAFS
- Security
- Integration with MS Windows
- Mobility
- OS X customers in science
- OS X @Fermilab
- Where to go from here

Conceptual Scheme of an ideal Mac OS X Computer

You are here

User interface and desktop environment:

AQUA, by Apple (proprietary) AND/OR
GNOME 2, KDE 3 UNIX (open source)

Client/server graphic system:

QUARZ, by Apple (proprietary) AND/OR
X11, UNIX open source system by XFree86

Core Operating System

DARWIN: it is a **UNIX** “BSD” (Berkeley Software Distribution) customized by Apple; **open source**

Kernel: **MACH v3** by Carnegie-Mellon Univ.; **open source**

Wait: what version of Mac OS X are you using ?

- Stop using Mac OS X v10.0 (Cheetah) or v10.1 (Puma) !
- Use at least Mac OS X v10.2.8 (Jaguar) or (recommended) latest v10.3.6 Panther: they have way more functionalities, fewer bugs and are faster
 - Choose right upgrade option for you
 - Scratch&Install, Archive&Install, Upgrade
- Mac OS X 10.4 Tiger will ship in ~spring 2005
 - If your mac will be running reliably, wait for the upgrade from Panther to Tiger; remember that no UNIX upgrade by a major version number is trivial

The Mac OS X workstation for HEP

- Partnership with IBM
 - IBM **PowerPC CPU architecture** (G5, customized version of Power4)
 - KLOE just deployed IBM p640 servers, AIX OS, each with 4 Power4 CPUs
 - IBM ***xlf, xlc* compilers** (fortran and C/C++)
- Native open source (Linux) software
 - by fink.sourceforge.net
 - **X11**
 - XFree86 4.3. WM: Apple quartz or anything you like
 - emacs, gimp, gv, xv, xfig, tex/latex, imagemagick ...
- **CERNLIB04, Geant3 (fortran), Geant4 (OO), ROOT4**
- Apple PREVIEW: **free** PS to PDF distiller/writer

What OS for physics experiments ?

- Online data acquisition (DAQ)
- Offline data analysis
- Offline simulation software
- Mission-critical task (all the above) performed almost exclusively using **UNIX** (AIX, SunOS, LINUX, IRIX (SGI), BSD
 - Performance, reliability, flexibility, security, cost reasons
 - *BSD is very stable, holds record of longest up time (and it's free, like LINUX)
- **Windows** and classic **Mac OS** used almost only in
 - Control systems of hw apparatus (LabView, Intellution, ...)
 - Specialized applications (geodesy, astronautics)

Developer tools

- System compiler: **gcc 3.3** (opensource)
 - typical linux (CDF2 code can now run under OS X)
- Xcode
 - Apple suite for C, Objective-C, C++, Java ...
- CHUD: computer understanding developer tools
 - **SHARK: impressive OO graphic debugger**
 - Profiling tools ...
- Powerful graphical scripting suite
 - Both Applescript and UNIX SHELL scripting (bash, tcsh, ...)
- Opensource (by fink.sourceforge.net)
 - **g77**
 - Perl, ruby, python, tcl/tk ...
 - Almost anything LINUX (apache, mysql, mozilla, ...)

OpenAFS

- **AFS** = the HEP “golden” distributed File System
 - pioneered by CMU, developed as a product of Transarc (now IBM Pittsburgh Labs); now it’s free (www.openafs.org)
- Almost every major Univ/Lab/Institute in the world has one or more AFS domains (“cells”)
 - Can be mounted from anywhere there is a WAN
 - Flexible ACL, reasonable security
- Command line and GUI functionality

Security

- Regular, prompt and effective security updates
- Core security provided by MIT **KERBEROS V5** and built into Mac OS X
 - Mandatory in some US DOE lab; ex: FERMILAB can be accessed only via **Kerberized ssh (or telnet or ftp)**
- Mac OS X, AFS and Kerberos authentication can be combined into a single login password/screen (OS X's)
 - typical choice of US Universities
 - **klog (afs), kinit (kerberos)**
- VPN
 - both IPsec and Microsoft's
- Safe sharing defaults; easy start/stop of personal (built-in) firewall

Integration with Microsoft Windows (Part 1)

- MS **Office 2004** Pro (INFN has “campus” license)
 - Word, excel, powerpoint, entourage
 - Windows Media Player
 - PC emulator (Virtual PC 7); no need of LAN
 - Remote Desktop Connection ; needs LAN (see below)
- MS Internet Explorer (included w/OS X)
 - webssh, webmail: OK !
- MS **Remote Desktop Connection**
 - It’s a remote login; uses CPUs, disks, applications of the remote Windows server (NT, 2000, XP, 2003)
 - Not an emulator, it’s a real PC via LAN
 - Analogous opensource product is *rdesktop* from fink (uses rdp protocol)

Mobility

- Integrated wireless networking (“WiFi”)
 - 11 to 55 Mbit/sec, 50 meters range
- Integrated wireless peripheral connectivity (“Bluetooth”)
 - 1 Mbit/sec, 10 meter range
 - Mice, keyboards, cell phones, printers, cameras ...
 - Good candidate for short-range “control” systems in HEP
- For portable computers, it is MANDATORY to start and stop AFS on network connect and disconnect
 - *source start_afs.csh*
 - *source stop_afs.csh*

How to start/stop manually OpenAFS

- Move `/Library/StartupItems/OpenAFS` elsewhere, for example to your Desktop
- Create with an editor a file named `start_afs.csh` (assuming your shell is *tcsh*) with this content:

```
#!/bin/sh
sudo ~your_username/Desktop/OpenAFS/OpenAFS
#
```

- In a term/xterm type `source start_afs.csh` to start OpenAFS

- Create with an editor a file named `stop_afs.csh` with this content:

```
#!/bin/sh
sudo umount -f /afs
sudo /sbin/kextunload -b org.openafs.fileSystems.afs
#
```

- In Mac OS X 10.2 Jaguar: `sudo umount /afs` (without “-f”)
- In a term/xterm type `source stop_afs.csh` to stop OpenAFS

OS X customers in science

- Mac OS X being used as the other UNIX
 - SNO (**Sudbury Neutrino Observatory**) for DAQ !
 - Research in outer space: **NASA et al**
- Mac OS X in the US deployed by
 - **Bio-sciences** (Virginia Tech ~1100 G5's were 3rd supercomputer at the end of 2003; 11 TeraFlops, 5 M\$)
 - **Army** (Colsa Corporation's ~1600 Xserve G5's likely to be in the top 5 supercomputer; ~25 TeraFlops)
 - **Universities**
- HEP is less important than Bio-Sciences and Army
 - But Apple Higher Education at the Paris EXPO was interested in the ATLAS and CMS “corporations”...
- Anything based on OO programming and gcc can **in principle** run under OS X

OS X @Fermilab

- [FreeBSD cluster in the Accelerator Controls Department](#), used for web services, data acquisition (through java based applications), java development, etc.
- Most desktops though are in Windows 2000 and they communicate to the FreeBSD cluster through [Samba \(most Windows users don't even know it is a BSD disk they are mounting on their desktop\)](#). It is pretty much kerberized.
- Mac OS X is used as a desktop/portable through the lab. It is not officially supported by the Fermilab Computing Division, though the number of Macs is growing, especially among the user community from the universities.
- Matt Crawford is a Computing Professional in the Security part of the Computing Division and he uses a [Powerbook because it fits with everything he needs to do with Windows and Kerberos, plus has the Unix part under the hood](#).
- CDF is moving to using GCC as the compiler of choice, so the [CDF offline software will run on Macs](#).
- I have also used ROOT regularly and it is [officially supported by CERN on Mac OS X](#). I have used RDC to talk with Windows Terminal Servers, though with all the stuff in Panther I don't need much access to the Windows desktop

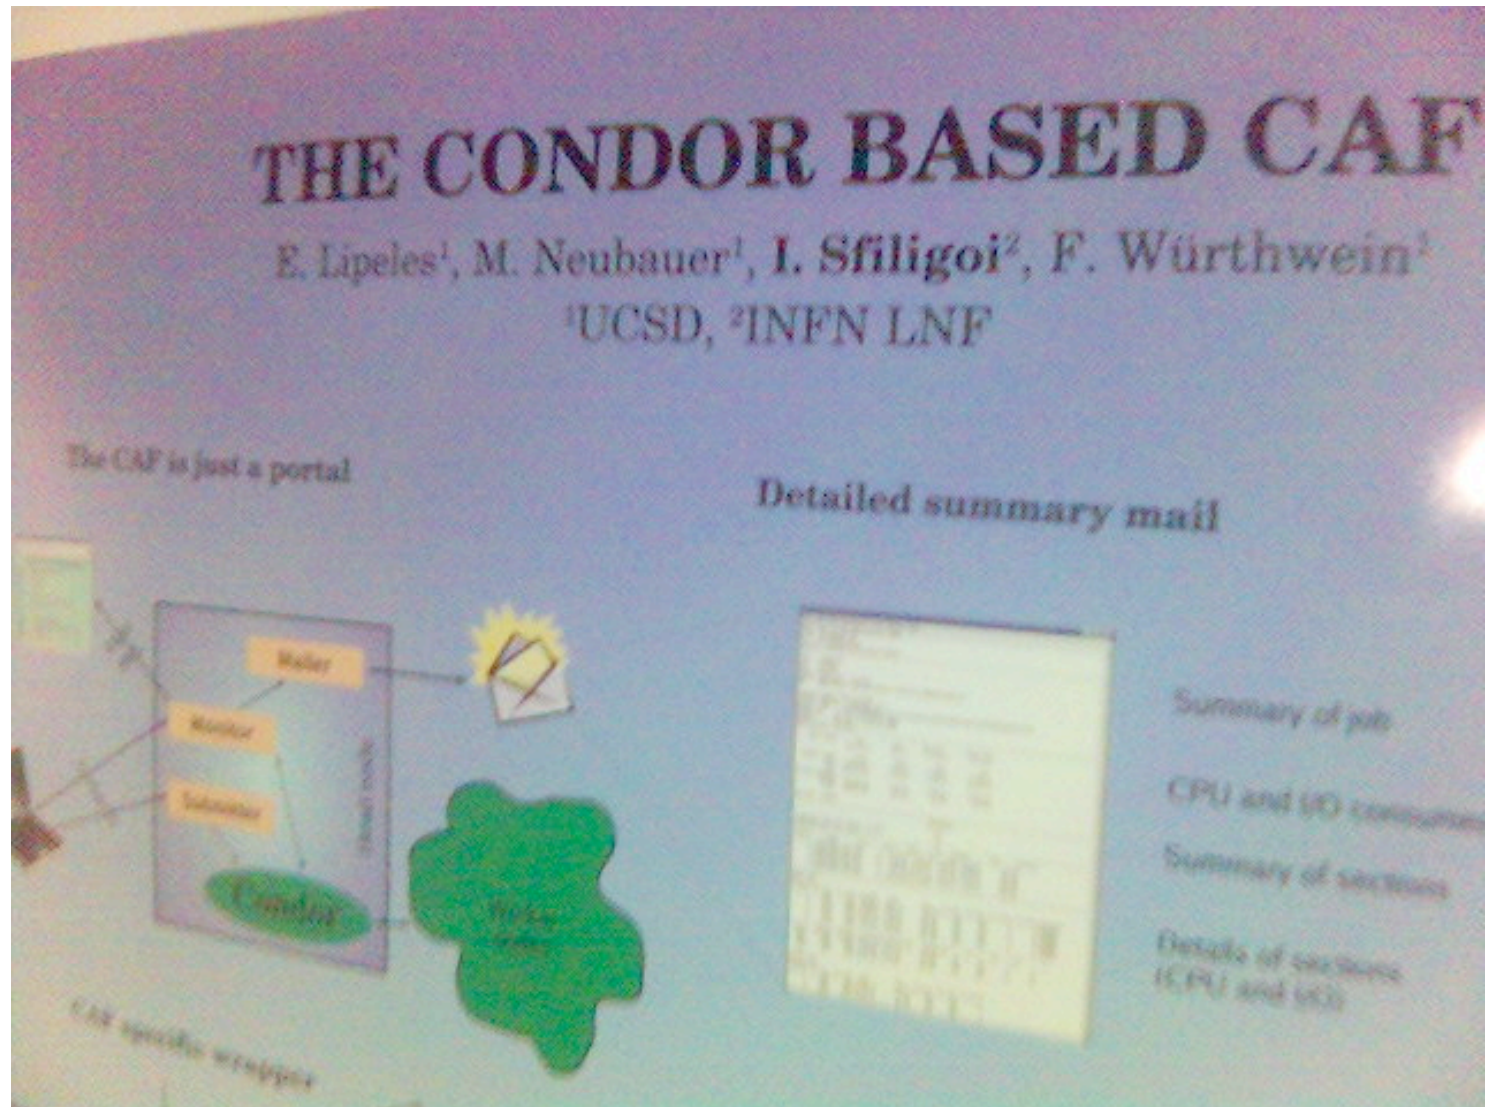
Where to go from here (Part 2)

- Server side of Mac OS X
- Workgroup solution feasible now at LNF --> Part 2
- Apple XGRID (an advanced implementation of **UNIX zeroconf**)
 - Tech. Preview 2 now, will ship next year
- Apple XSAN
 - **IBM SAN solution deployed by KLOE few weeks ago**
- Condor UNIX batch system
 - Industry standard
 - Heavily deployed under LINUX by CDF2 (I. Sfiligoi)
 - PIXAR (switching to Mac OS X/G5s)

CONDOR @ CDF2 (Fermilab)

Posters on
1st floor
of LNF AE

CAF =
CDF2
Analysis
Farm



Appendix

- In the following we append the presentation that we were invited to give at the meeting of Apple Higher Education Managers and Engineers with representatives of select Institutes, Labs and Universities from all over western europe at the Septempber 2004 Apple EXPO in Paris, France

Computing Tools, Applications used and needed in (High Energy) Physics

O. Ciaffoni (orlando.ciaffoni@Inf.infn.it),
S. Dell'Agnello (simone.dellagnello@Inf.infn.it)

INFN-LNF, ITALY

Paris, Apple EXPO 2004

OUTLINE

- A “physicist” perspective; select examples
- OS choice
- Desktops/Workstations and the Mac
- AFS, the “Golden” Distributed File System
- Present Computing Farms
 - Two examples: mid-size (Frascati/INFN, ITALY) and large-size (Fermilab/DOE, US) farms in experimental physics
- Future Farms
 - the LHC at CERN, Geneva, Swiss (towards “corporate experimental physics”)
 - Theoretical physics: LATTICE simulation, more Teraflops computing to come
- Plus, specialized “engineering” software (will not discuss it today)
 - Geodesy (LGC by CERN) and Astronautics (GEODYN, NASA)

OS used in Physics

- Online Data acquisition (DAQ) from physics experiments
- Offline Data Analysis
- Offline Simulation Software
- All the above performed mainly using **UNIX** OSes
 - AIX, SunOS, LINUX, IRIX (SGI)
- Windows and **Mac OS** up to now are used mainly in
 - Control systems of hw apparatus (LabView, Intellution, ...)
 - Specialized applications (Geodesy: metrology of physics installations)
- Mac OS X now starting to be used as the other UNIXes above: for example SNO (Sudbury Neutrino Observatory uses Mac OS X for DAQ !!!)

Desktop/Workstation: status and needs

- OS X is back on the desk of physicists (OS 9 was disappearing)
- What made the difference ? Availability of
 - **free and OSX-native X11** by Fink and now also by Apple
 - X11 versions of key sw packages (by fink.sourceforge.net):
 - **PAW** and **CERNLIB** (developed at CERN, Geneva, Suisse; Fortran based); past and present
 - **ROOT** (CERN, OO; Taligent origins); present and future
 - GIMP (=“Photoshop” for physicists)
 - OPENOFFICE (unfortunately there is MSOffice, but at least we don't dual boot our Macs like Linux users)
- **AFS** distributed file system (see next slide)
- MIT **Kerberos v5** (mandatory in DOE labs; ex: FERMILAB, near Chicago); **Kerberized ssh**
- Apple PREVIEW: free PDF viewer/distiller

The “Golden” Must-Have Distributed File System

- **AFS** = Andrew File System (not Apple File Services !). Pioneered by CMU, developed as a product of Transarc (now IBM Pittsburgh Labs); now it's free
- Almost every major Univ/Lab/Institute in the world has one or more AFS domains (“cells”)
 - Free versions: OPENAFS.org (more popular), ARLA.org; OPENAFS is now funded by Intel, Morgan Stanley, CMU, MIT, USENIX ...
 - Can be mounted from anywhere there is a WAN
 - Flexible ACL, reasonable security
 - OSX, AFS, Kerberos authentication can be combined (typical choice of US Universities)
- Current v1.2.11 not fully functional in Panther (up to **Kernel Panics**); command line worked better in Jaguar
 - Answer: Apple will try to support it *directly* in Tiger (2005)

Present Computing Farms

- Mid-size experiment and farms: **KLOE** at INFN-Frascati (Rome), ITALY; collaboration (100 people) active mainly in Italy
- Online farms
 - Database and data handling server: **1 IBM F50, AIX, 4 166-MHz PPC 603e, 2 GB RAM**
 - File server, offline disk pool, recall cache, tape library: **2 IBM H80s, AIX, each w/6 500-MHz RS64-III**
 - IBM Tape libraries for a total of **~ 1.1 PB (yrs 2001-5)**
- Farms for offline processing
 - **23 IBM B80, AIX, each w/4 375-MHz Power3 CPUs**
 - **10 SUN E450, SOLARIS, 4 UltraSPARC II CPUs**
 - Bought IBM p640 servers, AIX, each w/4 Power4 CPUs
(G5 ~ Power4, isn't)
- Total of ~ 200 PowerPC CPUs (like G3/G4/G5s), no Pentium PCs !

Present Computing Farms

- Large-size experiment and farm: **CDF II** at Fermilab/DOE, US; Collaboration (800 people) distributed across the whole world
- Farms use **Pentium/Athlon PCs** w/customized RedHat **Linux**: free OS and **gcc** compiler; **SGIs** still exist for interactive use and data staging
 - **Online**: ~100 single CPU PCs
 - **Offline**: ~1000 Dual CPU PCs at Fermilab (present), plus ~1000 Dual CPU PCs at Supercomputing Center in San Diego (in progress)
- **Dataset**: ~ **1.1 PB (2002-6)**; **ENSTORE** tape libs
- **Offline analysis model**
 - **FBSNG** (present) and **CONDOR** (present and future) batch managers (CONDOR available also for Mac OS X !)
 - Data are distributed using “**SAM**”, a small-scale prototype of GRID computing developed at Fermilab
 - Access to the data requires **KERBEROS V5** authentication

Future farm environments (I)

- New projects and/or upgraded facilities at Frascati (KLOE2) and Fermilab
- Experiments at the LHC of CERN, Geneva, Swiss
 - Likely largest physics project in the world for this decade
 - > 5000 physicists (“corporate” experimental physics) from the whole world; cost-to-completion ~ 2500 MEuro, then expected duration of ~20 yrs
 - World Wide GRID computing is a must ! Great deal of R&D all over the world; will have several large production centers (“tiers”; Frascati will probably be one of them)
 - Scale up by 10+ previous numbers relative to present farm/experiments

Future farm environments (II)

- Theoretical physics: **LATTICE** simulation
 - Solution of fundamental physics equations from first principles in a lattice of space-time points; it's like “solving the physics equations which govern NATURE by brute force”; APE project by INFN
 - Former precision was ~10%; **Teraflop** super-computing is needed to reach 1% by decreasing significantly the lattice pitch down to 10^{-13} cm
- GEANT3 Monte Carlo simulation (Fortran-based): heavily used in physics, space and medicine; the next version GEANT4, which is OO, is the next frontier
 - In physics typical samples of simulated data are about the same size of experimental data ==> prev. slides apply

What we would like from APPLE ?

- More visibility for the Mac OS X servers and XGrid
 - To be demonstrated to our community with well-chosen pilot projects
- UNIX-style support for running and maintenance of the potential Mac OS X servers and farms
 - None so far among italian resellers; 3rd party solutions ?
 - Where is the analogous of the 200+ apps for life science customers that Apple pre-installs on their **workgroup clusters**? See the slide on “Desktops and the Mac” for what physicists need
- Porting of specialized engineering sw packages with an agreement with HEP laboratories and Collaborations
- Direct support for OpenAFS