EPAC 2002 Computing Post-Mortem

Michel Jouvin – jouvin@lal.in2p3.fr
LAL – Bat. 200 – B.P. 34 – 91898 Orsay Cedex – France

EPAC 2002 was held in La Villette Congress Centre in Paris. It was attended by 900 people. LAL, one of the IN2P3 laboratory in charge of the organization of this EPAC edition, was asked to take responsibility for the computing infrastructure deployment and support. This has been done by people of the computing division at LAL and this note summarizes the choices made, the infrastructure deployed, the problem we faced and the lessons we learnt.

Expressed needs

We had to address the following needs :

- proceedings processing during the conference : a team of 20 people working during the conference, dedicated to this processings.
- a Web Café for attendees, allowing access to the internet and last minute edition of attendees proceedings and talks.

Proceedings Office configuration

Hardware

Proceedings Office was equipped with 20 PCs and 2 Macintoshes. The configuration was :

- PC : Pentium III 1Ghz, 256 MB of memory, 20 GB of disk, LCD 18" display (with integrated sound), running Windows 2000.
- Macintosh : G4 1Ghz, 256 MB of memory, 20 GB of disk, 19" display, running MacOS 9

In addition to the desktop computers, 4 printers were dedicated to the Proceedings Office :

- 2 Lexmark T620 : 20 ppm black and white laser printers, with duplex capabilities
- 1 Ricoh Afficio 3800C : a high end colour laser printer, able to sustain 20ppm for color printing, with duplex capabilities and a low cost per page.

- 1 Xerox/Tektronix Phaser 850 : a very high quality colour laser printer, with duplex capabilities but slower than Ricoh Afficio (around 5 ppm sustained rate).

All these hardware were rented, except the colour printers that were loaned by IN2P3 laboratories. The PCs were Compaq Deskpro but we issued no requirement on PC branding.

Proceedings Office users were satisfied by the configuration. The total number of desktop was enough for both people doing the real processing and people in charge of the welcome desk of the Proceedings Office. The ratio between Macintosh and PC was right too. One minor mistake made in Macintosh configuration was the lack of floppy drive. During the conference a few attendees had only a Macintosh formatted floppy to submit their proceedings and we had some difficulties to download them.

About printers, Ricoh Afficio 3800C was very much appreciated and there was no waiting queue for printed papers. Tektronix Phaser 850 was not really used, as expected, but was here for backup in case of a problem with the Afficio.

Software

Software deployed on each PC in the Proceedings Office were :

- Office XP
- Internet Explorer 6
- Netscape 4.7
- Acrobat 5
- Pitstop Pro for Acrobat 5 (evaluation version)
- Srip32
- MikTeX 2.1
- PowerArchiver (WinZip equivalent)
- SSH client 3.1 (from ssh.com)
- LeechFTP
- F-Secure (anti-virus)

Configuration on the Macintosh was pretty the same (with software equivalent for ssh, ftp, …).

Web Café Configuration

Main goals of the Web Café were :

- Access to internet for attendes
- Ability for attendees to modify their presentation or proceeding and submit them through the submission server.

For cost reasons, we chose not to offer colour printing in the Web Café.

Hardware

Based on previous editions of EPAC, we deployed 30 PCs and 2 Macintoshes in the Web Café. In addition, we offered 24 twisted-pairs Ethernet connections for laptop and wireless Ethernet coverage. Desktop configuration in the Web Café was :

- PC : Pentium III 733 Mhz, 128 MB of memory, 10 GB of disk, LCD 15" display, running Windows 2000.
- Macintosh : IMac, 128MB of memory, 10 GB of disk, running MacOS 10.

One black and white printer was deployed in the Web Café (Lexmark T620, 20ppm, duplex capability). Like for Proceedings Office, the PCs and printer have been rented. PCs were Compaq Deskpro. IMac have been loaned by LAL.

The hardware configuration proved to be adequate for the number of attendees. The only minor problem was the lack of floppy in Macintosh, requested by a few attendees. Web Café was almost full during the breaks but we observed no queue for accessing it. The same is true for the laptop connection through wired Ethernet. Wireless was not heavily used as this is still a new technology. But this service was appreciated by the people who used it.

Software

Software deployed on each PC in the Web Café were :

- Office XP
- Internet Explorer 6
- Netscape 4.7
- Acrobat Reader 5
- PowerArchiver (WinZip equivalent)

- SSH client 3.1 (from ssh.com)
- LeechFTP
- F-Secure (anti-virus)

Windows infrastructure design

We decided to deploy the PCs using an ActiveDirectory domain, as LAL has a lot of experience with ActiveDirectory. The main reasons to do that were :

- Ability to easily manage software deployment through ActiveDirectory Group Policies. This includes initial deployment but also the ability to reconfigure and redeploy with a very low manpower (in fact just a machine reboot).
- Ability to do initial system installation in an efficient and totally unattended way, using Windows 2000 Remote Installation Services. This basically allow to boot a generic floppy that launch the Windows 2000 installation without any input required.
- Ability to share files between all the PCs in a secured environment. This was a strong requirement for the proceedings office : they wanted to be able to share files between them and to be sure that other people (in particular Web Café users) will have no access to these shared resources.

This choice was very successful. One PC installation from scratch, including required applications (with some large ones like Office and Acrobat), was completed in 20 minutes. And we have been able to install 30 PCs in parallel in less than 1 hour.

The main difficulty we faced was the need to repackage almost all applications except MS Office as the standard installer provided is not compatible with ActiveDirectory Group Policy technology (MSI installers). LAL had a tool to do this repackaging (Wise for Windows Installer) as we are using it at LAL and had some experience that helped a lot. This repackaging leads to some problems at the beginning, described later.

Two dedicated PCs were acting both as ActiveDirectory domain controllers and application/file servers, for redundancy. Main application and file services deployed were :

- File service for shared files, including Appletalk Appleshare access
- Print server, including Appletalk access
- DHCP
- Remote Installation Services

- Web server for proceedings submission

These 2 servers had a very similar configuration : Pentium III (1Ghz, dual 866Mhz), 2 disks (9 and 36 GB), 512 MB of memory, 100Mb/s Ethernet connection. One of the PC was the main server, the only one accessed by users. The other one was a backup server. Main user disk areas were synchronized between both servers using Microsoft DFS. Thus it would have been possible to restart applications on the backup server in case of major failure of the main one in one hour (including reconfiguration of IIS, DHCP, RIS…).

Backup was done on both servers independently, on disk. The backup policiy was an incremental backup every hour for the 'PAPER UPLOAD' (submitted proceedings) and 'PAPER PROCESSED' (processed proceedings) areas. That meant 2 copies per server and was thought enough. Thus we decided not to use any tape based backup.

Disk size was appropriate. At the end of the conference, documents areas total size was about 6 GB and backup about 15 GB (over 36 GB available for the document disk).

Proceedings Submission Server

On the main server, we ran a Web server (IIS 5.0) used for proceedings submission. This Web server was started 2 months before the conference and was shut down one month after the conference.

This Web server has been running at LAL before and after the conference and was transferred at La Villette during the conference. This was eased by having a DNS domain name specific to this conference edition (epac02.org, see below). But, during the conference, it added the requirement of an fixed IP address for this server (to remain reachable from the Internet).

With this Web server, attendees were able to submit their proceedings. The proceeding itself was stored on the Web server but metadata information about the proceeding was registered in an Oracle database located at CERN. For many reason, it was not reasonable to try to locate this database on the Web server itself (in particular it was impossible to provide the level of Oracle management provided by CERN). This characteristic put a strong requirement

on network connection availability as both proceeding submission and processing were impossible without an active connection between the Web server and the Oracle server.

Network infrastructure

All the desktops and servers deployed during the conference were interconnected by the La Villette 100Mb/s Ethernet infrastructure available inside the congress centre. This was a great help. We just provided one Ethernet switch in each room to connect all the equipments deployed in a room.

The connection to Internet was provided through an ADSL link (1Mb/s download, 256 Mb/s upload) to Wanadoo (France Telecom) ISP. We had to subscribed the fixed IP option because of the proceedings submission server. We used an entry level ADSL router (Netgear MR 314) to provide Network Address translation between internet and the computers deployed inside de congress centre. All the equipments connected to the conference LAN were sharing one non routed class C IP network (255 addresses) and the same ADSL link for outside access.

Bandwidth was well adapted to the needs of this conference, both for Web Café and Proceedings Office. We experimented no network bottleneck. The main problem was the lack of backup solution. We suffered an 8 hours outage due to a Wanadoo problem and during this period, Proceedings Office was unable to work as they needed access to the Oracle database located at CERN.

We chose to have one DNS domain (epac02.org) dedicated to the conference, used by both servers, desktops and printers. All the conferences being in a non routed network, we had in fact one DNS server used by machines on the conference LAN (with dynamic update enabled), advertising non routed address for epac02.org, and one DNS server (located at LAL) serving epac02.org for access from outside. In this external server was registered only the EPAC 2002 proceedings submission Web server.

All the equipments deployed at the conference were using DHCP to get their address. The DHCP server was running on one of the Windows 2000 domain controller. The ADSL router and all the wireless access point had DHCP server disabled. The main reason for that was the Windows Remote Installation Services : it is easier to have the DHCP server collocated with the RIS server. DHCP was configured with a pool of 250 addresses, shared by

equipment we deployed and attendees laptops, and a very short address lifetime (2 hours, with automatic expansion). We experimented no depletion of addresses.

Main problems

We experimented four major problems with computing during this conference :

- Lack of validation of Proceedings Office desktops before they started to work
- DNS problem
- DFS replication inconsistencies
- Network outage

Desktop configuration validation

Proceedings Office members arrived 3 days before the conference start. They began to work at LAL where we deployed 10 PCs and 2 Macintosh for them. These PCs where configured with the same Windows infrastructure and configuration described before : in fact the servers, the network router and these PCs were configured at LAL and just moved to the conference.

During the first day we had several problems with the repackaged applications we used to deploy them through ActiveDirectory. This could have been avoided if one or two of the proceedings processing experts would have been able to come at LAL one day before and test everything. We got the required software list in advance but we didn't have the expertise to validate that specific software, like Acrobat Plugins, were working properly.

DNS configuration

We experimented an erratic problem (not permanent, not reproducible) when trying to connect to outside locations. It took some time to identify that this was related to DNS resolution that was not working for some domains (especially for CERN). We were not able to identify clearly the problem but it was probably related to our NAT configuration, leading to our DNS server being unable to forward name resolution request for some domains.

We found a workaround by defining DNS redirectors on the epac02.org DNS servers, forcing these servers to forward their request for non epac02.org names to a specific DNS server that was doing the forwarding properly.

DFS replication inconsistencies

As described in the Windows Infrastructure part, we were relying on DFS to replicate data between both servers. We experimented a few problems at the beginning because one of the server was improperly installed without the required service packs, leading to DFS database corruptions. After clearing this issue, we had no other problem.

When we discovered the problem we chose not to advertise DFS to users and require them to access files on only one server. DFS was used only to provide a transparent replication, without any risk of impact on users in case of problem.

Network outage

The most severe problem that affected the computing during the conference was a network outage of 8 hours the third day, because of a major problem affecting ADSL service of our ISP (Wanadoo). This had consequences on the work of the Proceedings Office as they were relying on connection to CERN to access the proceedings database.

We had no backup solution for the ADSL connection and it was a mistake. For example, we could have planned an ISDN backup to allow a minimal service for Proceedings Office only, on a much slower connection.

Conclusions

We think that the general design of EPAC 2002 computing infrastructure was right. In particular, the choice to deploy an ActiveDirectory domain allowed a very easy and smooth deployment, an easy document sharing between Proceedings Office members and a configuration tuning on site with a very low manpower.

Excepted the network outage, the configuration ran smoothly during the conference and we were able to support the users on site with one person per day. If Internet connection is required for Proceedings Office work and not only for Web Café, a backup solution is clearly a requirement. Or proceedings processing should be revisited to allow a ‘disconnected mode’ (run the database locally or maintain a local copy of the database for example).

For the Proceedings Office, it is necessary to have a configuration validation before the work starts, with one “expert” testing the configuration of all required tools. Without our ActiveDirectory infrastructure, the need to reconfigure some of the tools because of the lack of validation would have been a nightmare.