

# Using System Commander in Patagonia - A Multi Use Cluster of PCs for Computation, Collaboration and Education

White Paper

by Felix Rauch, Christian Kurmann and Thomas M. Stricker

Laboratory for Computer Systems  
ETH - Swiss Institute of Technology  
CH-8092 Zuerich, Switzerland

{rauch,kurmann}@inf.ethz.ch, tomstr@acm.org

## Introduction

Clusters of PCs are emerging rapidly as the best low cost platform for several applications that require true supercomputing performance regarding the amount of computation involved or regarding the amount of data to be communicated between the PCs of a cluster.

We can identify three modes of usage for such compute clusters in practice. A first kind of PC cluster is used for scientific research, development in engineering and some selected business applications. Those clusters are built from up to hundreds of rack mounted PCs, typically stored in a cooled machine room and interconnected with a high speed system area network. As a second mode of usage, many corporations dispatch hundreds of PCs to their employees desks to give them personal computing support and access to all important information needed to do business. Such PC installations are rarely viewed as a cluster, but future applications like multimedia support for collaboration (e.g. teleconferencing) or data mining will soon require the computation-, storage- and communication resources of supercomputers. A third kind of clusters is found in all training and education setups, where PCs are installed in classrooms and used for education and corporate training.

In this white paper we report about the use of [System Commander](#) to configure and operate such clusters as well as about a few interesting technological problems we solved to make our [Patagonia cluster](#) at the Swiss Institute of Technology, ETH Zürich (see Figure 1) a viable research vehicle in all three mentioned areas of cluster computing. We indicate that it is impossible to accommodate all needs with a single operating system installation and demonstrate that flexibility can be achieved with a multi-boot setup, featuring multiple independent operating system installation on a single computer. The OSES include our own ETH Oberon System, Microsoft Windows NT and Linux.



Figure 1: The Patagonia cluster.

## Different Requirements for Education and Research

In the Patagonia cluster, we plan for research users in computer sciences and scientific computing as well as for a broad spectrum of educational users taking courses with entirely different requirements in terms of system software, application software and account management. Our highly selective university offers mainly combined undergraduate/graduate degrees and therefore many highly specialized courses require elaborate software installations (e.g. Microsoft SQL Server, Oracle or large Circuit Design Environments).

The different operating systems and application sets for education and research require software installations with different administrative setups and different protection levels.

## Operating Systems

Researchers in computer sciences and scientific computing need the freedom to try different parameters and improvements to their application code including fine-tuning and customizing the OS; therefore an open and flexible operating system with a readily available full source code is often required.

Educational users on the other hand, are mostly unsophisticated simple users, so a complex and unhandy operating system or difficult booting procedure is unacceptable. Even for a pure educational cluster the option to boot different operating systems is a big advantage as different lectures might have different needs (e.g. our non-standard operating system Oberon is the programming environment of choice for several courses here at ETH).

In an educational environment with a large user base (over 1000 students for CS majors alone) it is also important that users can individually configure their accounts and that all customization state is associated with central storage space on the server and not with the machines.

## Different Requirements for Security and Maintenance

Users in scientific computing research projects typically work in the same group and usually trust each other. They are also fairly competent users who know what they are doing. In this mode of operation, security and fool-proofness of the installation is not very important. Security can be achieved at a very coarse grain level (e.g. by one single System Commander password to boot an installation). But once a cluster is also used by students, some of them will inevitably try to stretch its limits of allowed use, some will try to break into installed applications or crack the entire system security. Also some unsophisticated users might not understand the system enough to realize when they are doing something harmful or cause permanent damage to the installation. An educational or dual-use system must therefore protect itself from unwanted changes of state in the operating system images. The different operating systems need to be properly sandboxed (protected from each other) by maintaining a reasonable level of overall operating system security. Maintenance of the computational installation is mostly done by the computer scientists themselves, but the educational environment will need designated, professional system administrators. Both groups should be able to maintain their own parts without much interference.

## System Software in the Patagonia Cluster at ETH Zürich

For the educational mode of operation, two types of operating systems are currently used to support different courses in computer science: Microsoft Windows NT and ETH System Oberon.

Windows NT is a well-known commodity OS that is most suitable for introduction courses involving the use of word processors, spreadsheets, databases and global information systems such as WWW and e-mail. Both, a standard English and a localized German user interfaces is required to cover the education of a broad range of students in computer science and other departments. Computer scientists, that work on several platforms simultaneously prefer a consistent English user interface across all platforms, while most other students learn faster if the course materials written in German are consistent with the user interface of the operating system and the application software. Until all common operating systems and common applications support switching between multiple languages on the fly, separate installations must be maintained. System Commander helps to run a German and English version selectively on the same Machine (as a major technical university we have generous site licensing agreements, so the doubled software cost is not an issue). Also the original and nationalized versions of most application programs can not be installed concurrently and therefore the Patagonia cluster features two fully isolated Windows NT installations for education in German and in English. These images are provided in addition to the Windows NT image with experimental drivers and development tools for research. Again the OS sand boxing techniques created for the multi purpose clusters offer new possibilities for the educational mode of operation. All we need is large disk and a tool like system commander to boot the different installation.

Besides the two Windows NT partitions for education in German and English, the disks of the Patagonia cluster also

host the [System Oberon](#). Oberon is a programming language, a run-time- and operating system with an integrated development environment for object oriented or structured programming. The Oberon system is kept lean, simple and easily fits into a very small disk partition or main memory. All data is stored on the local disk, which is mounted read only. The modules and objects of the system are only copied to a RAM-disk when required. As there are no user accounts in Oberon and no home directories, no server is needed. Network access is used only for common Internet services like e-mail, WWW and printing. The programming work of students can be kept on a simple floppy disk or on a ZIP drive.

Server installations over Gigabit Networks are planned but for the time being, a working set consisting of the most common software is installed on the local disk drives for faster access, ease of installation and to take as much load as possible from the central Windows NT server. The thousands of student home directories are stored on a central server running Solaris. To make the system as transparent as possible the boot-partition as well as all other local partitions are hidden and protected from the users access. The partition of the active operating system is remapped and appears to the users of any Windows NT Education image as if there were only one single partition on the C: drive.

For research in scientific computing, research in systems or electronic collaboration systems, we chose to install Linux and Windows NT as operating systems. Linux is an *Open Source* OS which offers great flexibility to researchers. Some software and drivers for our advanced networking hardware are unfortunately only available for Windows NT, therefore this OS is also installed for selected applications and performance testing in the research mode of our cluster. For maximum flexibility, the Windows NT used in research is installed in a separate partition that is completely isolated from educational use and is password protected as an entire image. Additionally this partition is also maintained by the scientists and file protection is not enabled since it is not needed.

## Details of the Security Setup

The most important goals of a successful security setup are not to inconvenience the users with unknown and therefore distracting operational procedures of unknown OSes. The purpose of a security setup is to protect the integrity of the system installation from corruption and the different users from each other. In a multi-boot setup multiple levels of security are needed. System Commander features a simple but complete security system that is well engineered and strongly protected.

- The first level of this security setup handles the booting procedure. The *System Commander* utility (see Figure 2) offers menu controlled selection of the operating system of choice and protects research systems and administration setup with a sophisticated password protection scheme, which allows user groups to be defined. This protects educational users from booting unsupported OSes with which they would experience unexpected authentication failures. Further, it protects from booting from external attached devices and floppy disks.

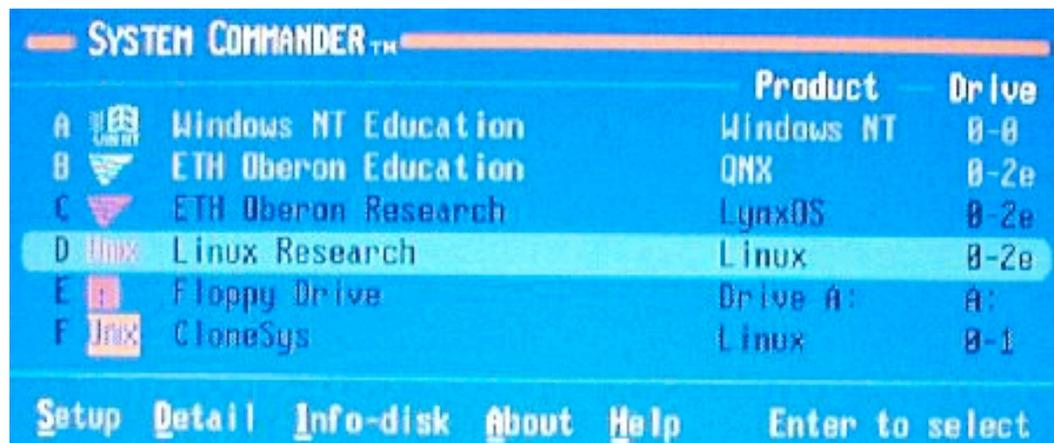


Figure 2: System Commander boot screen as used in Patagonia-Cluster at ETH Zürich.

- The second level of system security deals with the visibility (or better non-visibility) of other non-active operating systems, their partitions including the boot partition with its administrative tools. It should remain impossible to access or modify data from alternate partitions at least not for a normal, mortal user. In order to dissuade educational users from cracking system security, a [shareware tool](#), called *DeviceLock* hides unused partitions and denies access to them from Windows NT. The same task is accomplished in UNIX systems by configuring the corresponding mounting tables. Mounting and unmounting volumes are protected

operations that require administrative privileges or are accomplished automatically at boot time.

- For the purpose of allowing the same application links (and shortcuts stored in the user profiles) across all Windows NT installations on campus, we enlist the device name remapping mechanism of Windows NT and remap the partitions in a way that the partition of the active OS image is always the C: drive. With this setup it is even possible to switch between two different Windows NT systems transparently for an English or German version, while maintaining the same user profiles and application links. All these tricks and treats would not be necessary under a well designed OS like Linux that offers mount points in the file system and keeps mount a privileged operation.

As the cluster is also used by many students, the system must make a reasonable effort to protect itself from being modified or damaged. Within an executing OS image some protection on the file level can be accomplished by invoking the appropriate security setup mechanisms offered by UNIX and Windows NT. In the education environment a Windows NT Domain Server authenticates users and controls access to the local and remote files in the cluster - for the research environment a Linux server handles all user authentication employing the NIS protocol. The Oberon system copies itself onto a ram disk upon startup and the access to the disks is restricted to read-only access at the driver level.

## CloneSys - The Maintenance Operating System

With such a large number of Operating System spread over several dozens of cluster machines, the maintenance could easily become a nightmare. Sophisticated setups with System Commander and remapped partitions often prevent automatic install scripts from proper operations. We therefore devised a operating system independent maintenance strategy by replicating raw partition images.

This system was built on top of a small LINUX installation, which we call CloneSys due to the support for partition cloning. Again System Commander protects the system from accidentally booting into maintenance mode. Thanks to our high speed networking facilities and a broad cast like multi-drop distribution utility we manage to install an entire cluster with 5 operating system on 24 machines within 40-50 Minutes. A single operating system (e.g. a 2 GB Windows NT installation) is properly replaced or restored on all machines in less than 10 minutes.

## Conclusions

In our work on the Patagonia multi use cluster technologies we proved that it is indeed possible to user a single infrastructure for high performance computation, educational use and research in electronic collaboration. The system was in use as a prototype at ETH Zürich since spring 1999 and went in production with 50 education and 25 research machines by late 1999. In the white paper we presented some details of the installation, giving a special focus on details of the protection of the different supported installations. For achieving this protection, the System Commander as well as the Device Lock utilities proved to be very useful and stable tools for such a specialized setup, while offering an easy to use and secure user interface even for non-technical users of our cluster.