XtreemOS: a Grid OS providing native Virtual Organization support

Christine Morin
XtreemOS scientific coordinator
Yvon Jégou
Adrien Lèbre
Thierry Priol
Oscar Sanchez
INRIA Rennes - Bretagne Atlantique research centre (France)

Haiyan Yu
CAS/ICT (China)

Erica Yang, Brian Matthews
STFC (UK)

Luis Pablo Prieto
TID (Spain)

1- Objectives

The emergence of Grids enables the sharing of a wide range of resources for solving large-scale computational and data intensive problems in science, engineering and business. While much has been done to build Grid middleware on top of existing operating systems, little has been done to extend the underlying operating systems for enabling and facilitating Grid computing, for example by embedding important functionalities directly into the operating system. XtreemOS project aims at investigating and proposing new services that should be added to current operating systems to build a Grid infrastructure in a simple way.

This approach can be seen to have some advantages over conventional Grid middleware toolkits, which may have different programming interfaces and lack of a unifying model. A common interface can be provided to simplify the task of the application developer on the Grid by making the Grid support native to the operating system, and also by removing layers of abstraction, leading to higher dependability of services.

Therefore, the goals of XtreemOS project are to design, implement, evaluate and distribute an open source Grid operating system which supports Grid applications, and capable of running on a wide range of underlying platforms, from clusters to mobiles. Installed on each participating machine (PC, cluster of workstations, mobile device), the XtreemOS system will provide for the Grid what a traditional operating system offers for a single computer: abstraction from the hardware and secure resource sharing between different users. It will thus considerably ease the work of users belonging to virtual organisations by

giving them the illusion of using a traditional computer, and releasing them from dealing with the complex resource management issues of a typical Grid environment. The approach being investigated is to base XtreemOS on the existing well-accepted open source Linux OS. The underlying Linux OS will be extended as needed to support virtual organisations spanning across many machines and to provide appropriate interfaces to Grid OS services. A set of system services will provide users with all the Grid capabilities associated with current Grid middleware, but fully integrated into the OS. By integrating Grid capabilities into the Linux operating system, XtreemOS will also provide a more robust, secure and easier way to manage a Grid infrastructure for system administrators. This will be experimentally demonstrated with a set of real applications, provided by industrial partners that cover a large spectrum of application fields.

2- Overall Architecture

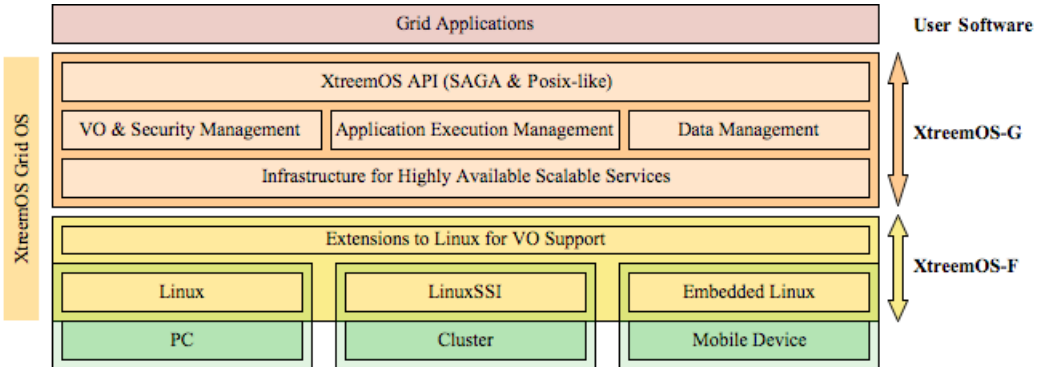

Figure 1

The overall XtreemOS architecture is depicted in Figure 1. The XtreemOS architecture is divided logically into two layers: XtreemOS Foundation layer, XtreemOS-F; and XtreemOS Grid-support layer, XtreemOS-G.

XtreemOS-F provides a modified Linux kernel embedding native support for Virtual Organizations, in three major variants: one for PCs and workstations; one for cluster computing (based on the existing Kerrighed Single-System Image OS); and one for small mobile devices such as PDAs and smartphones.

XtreemOS-G provides several Grid OS distributed services to securely manage computation and data resources: Application Execution Management (AEM), that allows starting, monitoring and controlling applications and selecting and allocating resources to them; a Grid file system (XtreemFS), that provides reliable and efficient data management; and VO and Security Management (VOM), that guarantees a secure operational environment for the Grid, covering all common requirements for information security as well as those intrinsic to the Grid. Additionally, a highly-available and scalable infrastructure is set up, in order to deploy these services over a large number of nodes. This infrastructure supports publish/subscribe services, directory service, node/resource discovery, node virtualization and distributed servers. XtreemOS provides a common API to Grid applications, based on the emerging Simple API for Grid Applications (SAGA) standard under development within the OGF.

3- Virtual Organizations and Security Management

XtreemOS aims to provide native support for the management of VOs in a secure and scalable way, without compromising on flexibility and performance. VO Management (VOM) covers all the infrastructural services that are needed to manage the entities involved in a VO and ensure a consistent and coherent exploitation of the resources, capabilities, and information inside the VO under the governance of the VO policies. In XtreemOS we are aiming to integrate VOM as part of the OS. More specifically, VOM can be implemented as a service that can be integrated directly with existing authentication infrastructure. First, this approach reduces the management and performance overheads introduced by the layers of controls. Second, the hassle of accessing VO resources can be reduced. The

policies specified by a VO, such as security, resource limitations, scheduling priorities and rules on how shared resources could be used by VO members, will be finally checked and ensured at resource nodes. In order to adapt to different VO models and reduce kernel code changes, XtreemOS uses PAM (Pluggable Authentication Modules). This way, a system administrator can add (possibly VO-specific) authentication methods just by installing new PAM modules. Local user accounts in XtreemOS are allocated dynamically on each resource to match the actual global users exploiting that resource. The XtreemOS PAM plugins are in charge of implementing (or interfacing to) a local service allocating fresh local UID/GID couples upon request. The dynamic allocation of user accounts ensures XtreemOS scalability and reduces the complexity of VO management: no need to configure resources when users are added or removed from VOs. Dynamic management of local UID/GID also provides some level of isolation between Grid users: they do not share access to local files, and it is possible to hide the real identity of a user in the local name space.

4- Concluding Remarks

XtreemOS is 4-year project funded by the European Commission, started in June 2006.

XtreemOS Grid OS, currently under implementation, will be released as open source software. A first public release of XtreemOS basic version is planned in Spring 2008.

References

http://www.xtreemos.eu/