

OW2 ProActive Parallel Suite: Building Flexible Enterprise CLOUDs

by Denis Caromel, Cédric Dalmaso, Christian Delbe, Fabrice Fontenoy and Oleg Smirnov

The ProActive Parallel Suite features Infrastructure as a Service (IaaS) capabilities together with an innovative parallel programming model and a distributed workflow environment. It involves the OASIS team from INRIA Sophia Antipolis, which initiated the development early 2000 and ActiveEon, an INRIA spin-off created in 2007, which together co-develop ProActive and provide users with professional services.

Federating large sets of distributed resources is an important issue for companies. Scientists and engineers in many fields including finance, engineering, entertainment, and energy need increasing amounts of computational power. As a consequence, companies and laboratories are placing increasing demands on their existing infrastructure. Due to peak workloads, companies also want to gain flexibility. Finally, green IT with improved control of infrastructure usage is another reason for desiring more precise control with optimization of the overall workload. In order to address these strong industrial and scientific needs INRIA and ActiveEon provide IT manager with a simple way to aggregate native or virtualized resources.

Resource Provisioning

ProActive Resourcing is the first building block used to provide heterogeneous resource management. ProActive Resourcing provides an open source intelligent and adaptive application deployment engine to virtualize hardware resources and monitor and control all computing resources. With ProActive, resource management is easier and highly configurable. It leverages the existing infrastructure of an organization, from dedicated clusters to heterogeneous distributed resources, building a Private Cloud with the capacity to manage virtualization and software appliances.

We introduce the concept of Node Source which associates the method of acquisition of computing resources with the policy determining when these resources have to deploy or to be released. Node sources allow a company to have business driven management of their computing power. New resources can be automatically acquired at any given time, and acquisition can be automatically triggered according to the load

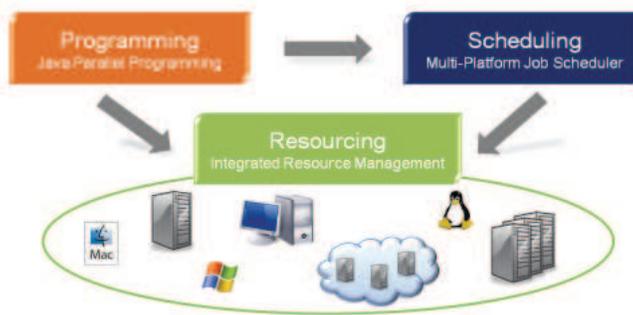


Figure 1: ProActive Parallel Suite.

of the current infrastructure. These new resources can come from another business unit in the same company or from outside the company like a Data Center or a public cloud.

The solution, developed in Java, is highly portable and can be deployed on Unix-like, Windows, and Mac operating systems. Resources can be virtualized or not. The following virtualization environments are supported: VMware, KVM, Xen, Xen Server, QMU, and Microsoft Hyper-V. In case of native resources directly accessed through well known protocols, RSH and SSH can be used, as well as third party schedulers like PBS, LSF, SGE, IBM Load Leveler, Oar, Prun.

Job Scheduling and Workload Management

ProActive Scheduling is an open source multi-platform job scheduler managing the distribution of workflows and application execution over the available computing resources. ProActive Scheduling ensures more work is done with fewer resources: maximum utilization and optimal allocation of existing IT infrastructure, reducing administration costs and future hardware expenditures.

A command line interface (CLI) as well as a graphical user interface (GUI)

based on Eclipse RCP provides users and administrators with all the tools needed to easily submit, monitor and retrieve results, and administrate enterprise cloud. Moreover, workflows can be built using various methods such as XML file, flat file, Web Service and programming API in Java and C/C++. The ability to use a range of languages facilitates integration with any kind of application when part of a workload needs to be delegated to other resources.

Operating a production Platform: ProActive PACA Grid

The ProActive PACA Grid is a Computing Cloud operated by INRIA, University of Nice and CNRS-I3S laboratory; it is funded by the PACA Lander and the European Commission. The Cloud platform makes available a set of machines to laboratories and SMEs. The resources are accessible via graphical interactive interfaces launchable from the PACA Grid website, in a portal mode.

The machines are currently deployed within INRIA Sophia Antipolis networks. The Cloud aggregates dedicated machines, both Linux and Windows, GPU processors, and spare desktop machines, dynamically added during nights and week-ends. Infrastructure and workload are managed using ProActive Resourcing Scheduling. It is

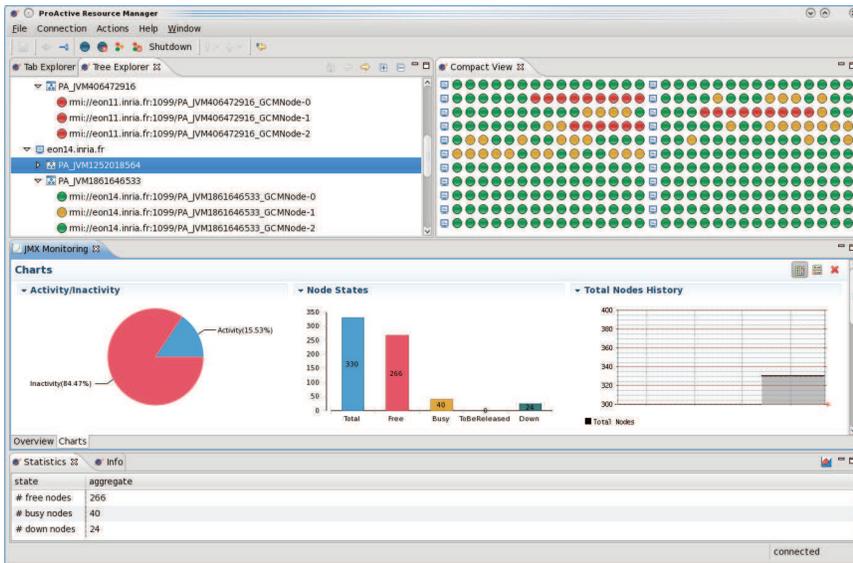


Figure 2: ProActive Resource Manager.

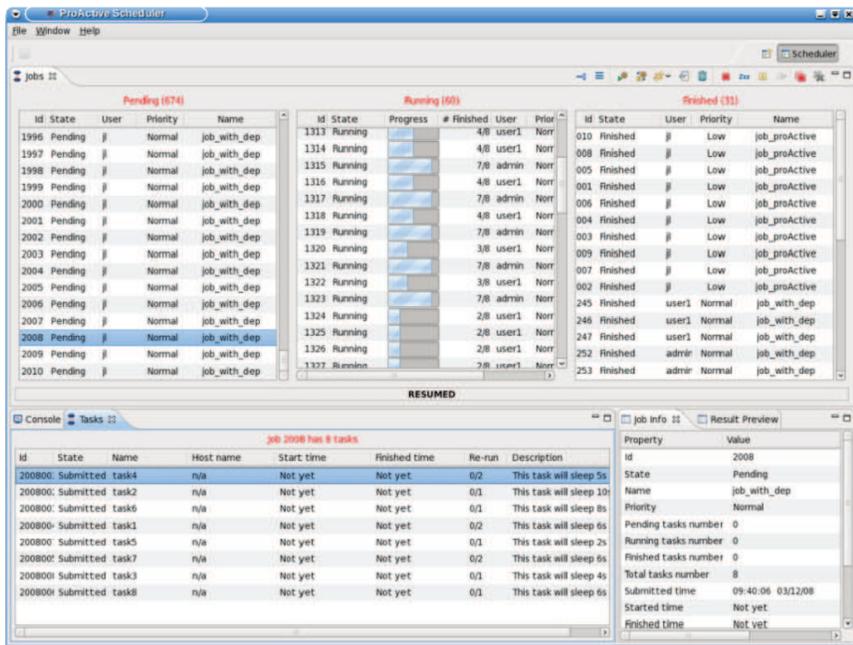


Figure 3: ProActive Scheduler.

integrated with the infrastructure through JMX/Nagios for the monitoring and LDAP for the authentication. The use of PACA Grid is simplified by the DataSpaces feature which automatically transfers input file parameters and brings home output results. Today, ProActive PACA Grid features in production over 1000 CPU Cores, 4 TByte of RAM, 15 TByte of storage, and 480 CUDA Cores for about 2 TFlops.

Conclusion and Perspectives

The comprehensive ProActive Parallel Suite toolkit will be soon enriched with a graphical editor for Workflows allowing design and monitoring of jobs of tasks, and a Web Portal to benefit

from thin client interfaces. In addition, ProActive is one of the three building blocks of the Open Source OW2 Cloud Initiative (OSCI) recently launched by the consortium.

The OASIS has collaborated with many EU partners such as University of Pisa, IBM, Atos Origin, Thales, Microsoft, HP, Oracle and Telefónica, including co-operation with other ERCIM members including CNR and Fraunhofer, and EU projects GridCOMP, CoreGRID, SOA4ALL, and TEFIS. The ProActive team has also built close relationships with international partners like University of Adelaide, Tsinghua University, University of Chile, STIC in Shanghai.

Links:

- ActiveEon SAS:
<http://www.activeeon.com/>
- ProActive Parallel Suite
<http://proactive.inria.fr/>
- PACA Grid:
<http://proactive.inria.fr/pacagrid/>
- OW2 OSCI:
<http://www.ow2.org/view/Cloud/>

Please contact:

Denis Caromel,
INRIA Sophia Antipolis, France
Tel: +33 4 92 38 76 31
E-mail: Denis.Caromel@inria.fr