ProActive Hybrid Workflows with CPUs and GPUs and various Use Cases

D. Caromel, et al.

Agenda

1. Background: INRIA, ActiveEon
2. Multi-Core and Virtualization
3. ProActive Parallel Suite
   Programming, Scheduling, Resourcing
4. Use Cases & Demos (PACA Grid)
5. Conclusion

Cloud Computing Revolution ?

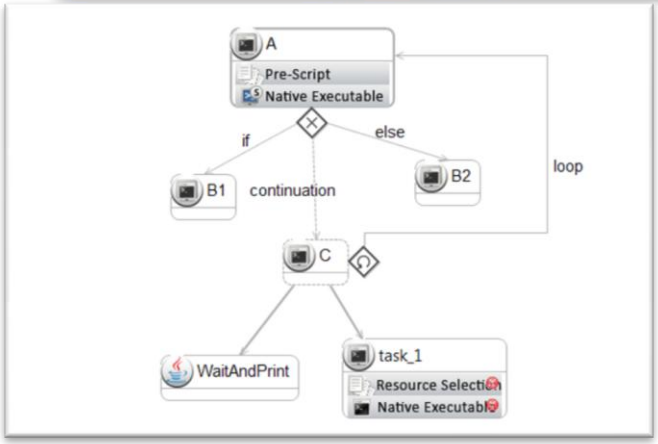

Workflow Execution
Studio Editor and Visualization
Parallel Programming in Java

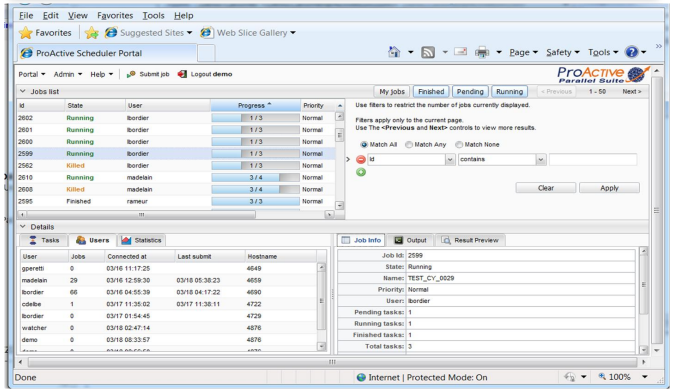

Portal, Multi-Application & Multi-Tenant Enterprise Orchestration

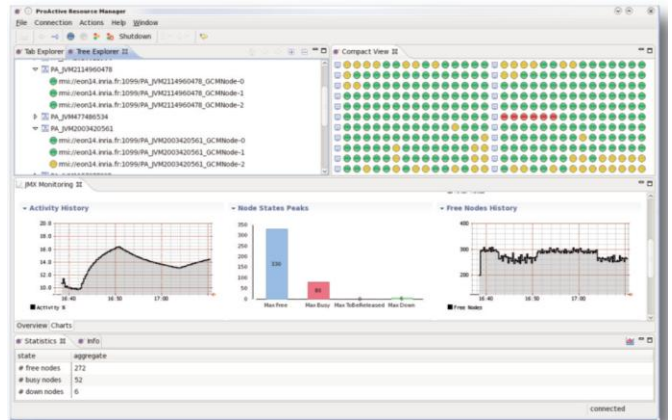

Physical and Virtual Machines Management

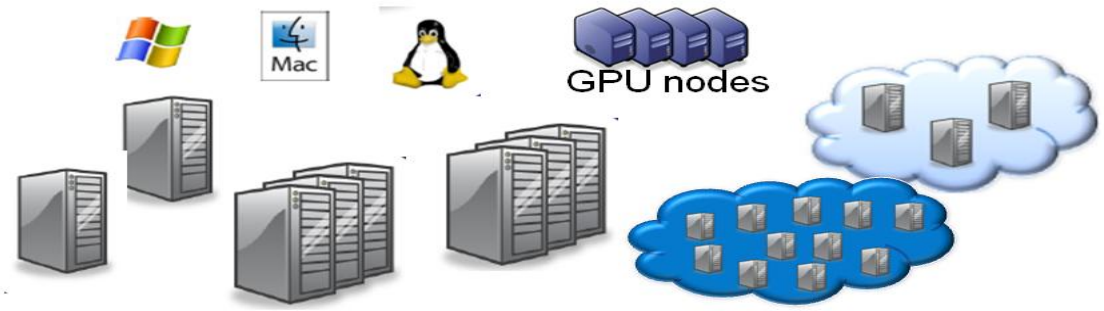

INRIA OASIS Team Composition (35)

- Researchers (5):
  - D. Caromel (UNSA, Det. INRIA
  - E. Madelaine (INRIA)
  - F. Baude (UNSA)
  - F. Huet (UNSA)
  - L. Henrio (CNRS)
- PhDs (11):
  - Antonio Cansado (INRIA, Coni
  - Brian Amedro (SCS-Agos)
  - Cristian Ruz (INRIA, Conicyt)
  - Elton Mathias (INRIA-Cordi)
  - Imen Filali (SCS-Agos / FP7 S
  - Marcela Rivera (INRIA, Conicy
  - Muhammad Khan (STIC-Asia)
  - Paul Naoumenko (INRIA/Régio
  - Viet Dung Doan (FP6 Bionets)
  - Virginie Contes (SOA4ALL)
  - Guilherme Pezzi (AGOS, CIFR
- + Visitors + Interns

Located in Sophia Antipolis, between Nice and Cannes,
Visitors Welcome!

OASIS Team & INRIA

- A joint team, about 35 persons
- 2004: First ProActive User Group
- 2011, July: ProActive 5.1, Distributed & Parallel: From Multi-cores to Enterprise GRIDs & Clouds

ProActive Parallel Suite

Startup Company Born of INRIA

- Co-developing, Support for ProActive Parallel Suite
- Worldwide Customers: Fr, UK, Boston USA

ActiveEon Overview

- **ActiveEon**, a software company **born of INRIA**, founded in 2007. HQ in the French scientific park Sophia Antipolis
- **Co developing** with INRIA ***ProActive Parallel Suite®***, a Professional Open Source middleware for parallel, distributed, multi-core computing
  30 peoples in total
- Core **mission**: Scale Beyond Limits
- Providing **a full range of services** for ProActive Parallel Suite
- **Worldwide** customers and production users:

RENAULT, hp, ifp Energies nouvelles, orange, EDF, GMO, DEXIA, THALES, sirehna a DCNS company, Tsinghua University, cea, amadeus Your technology partner, EMBL-EBI

ProActive Parallel Suite


Multi-Core Push

Symetrical Multi-Core: 8-ways Niagara II

- 8 cores
- 4 Native threads per core

- Linux see 32 cores!

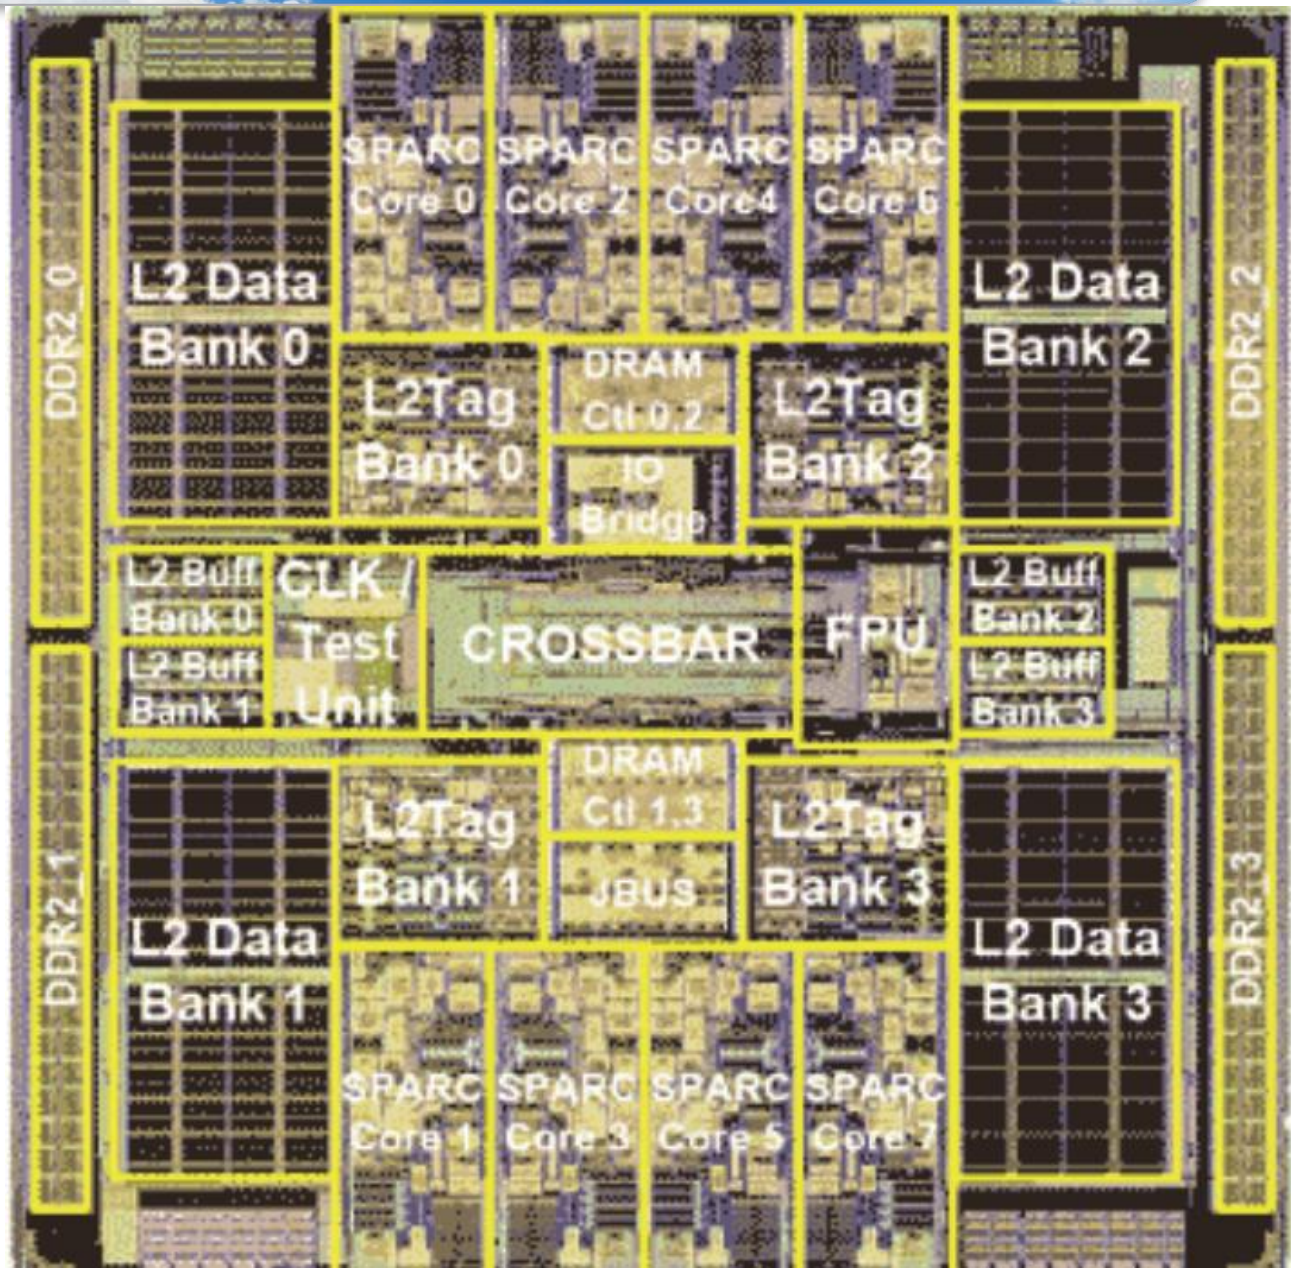

Off The Shelf Multi-Cores, 3 GHz

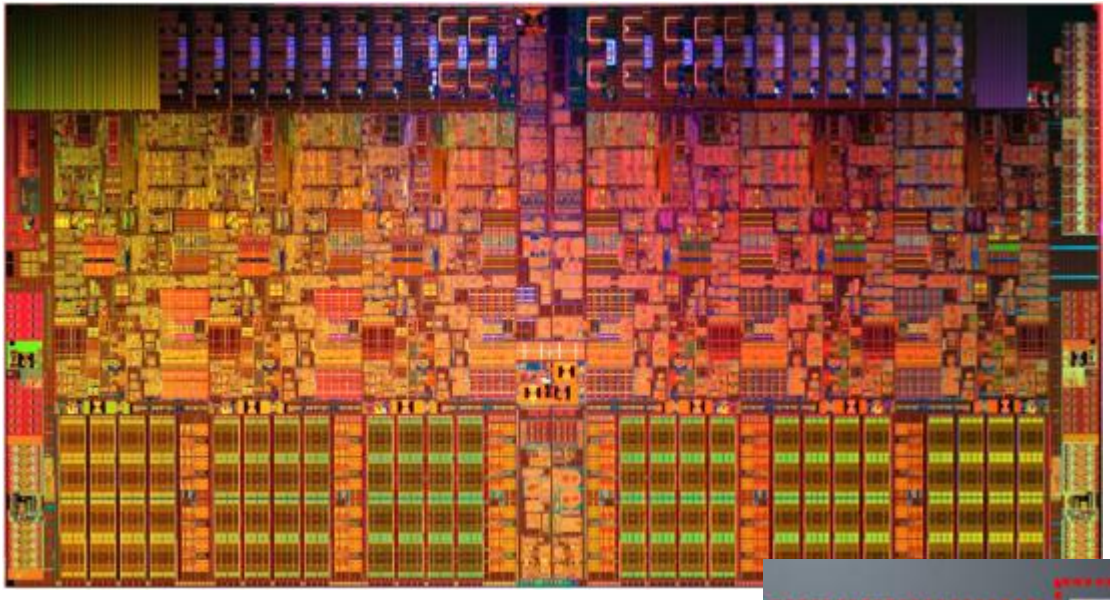

Intel Xeon 5670,
6 cores, 10, … 12

AMD's Opteron 6174,
“Magny-Cours”,
12 cores

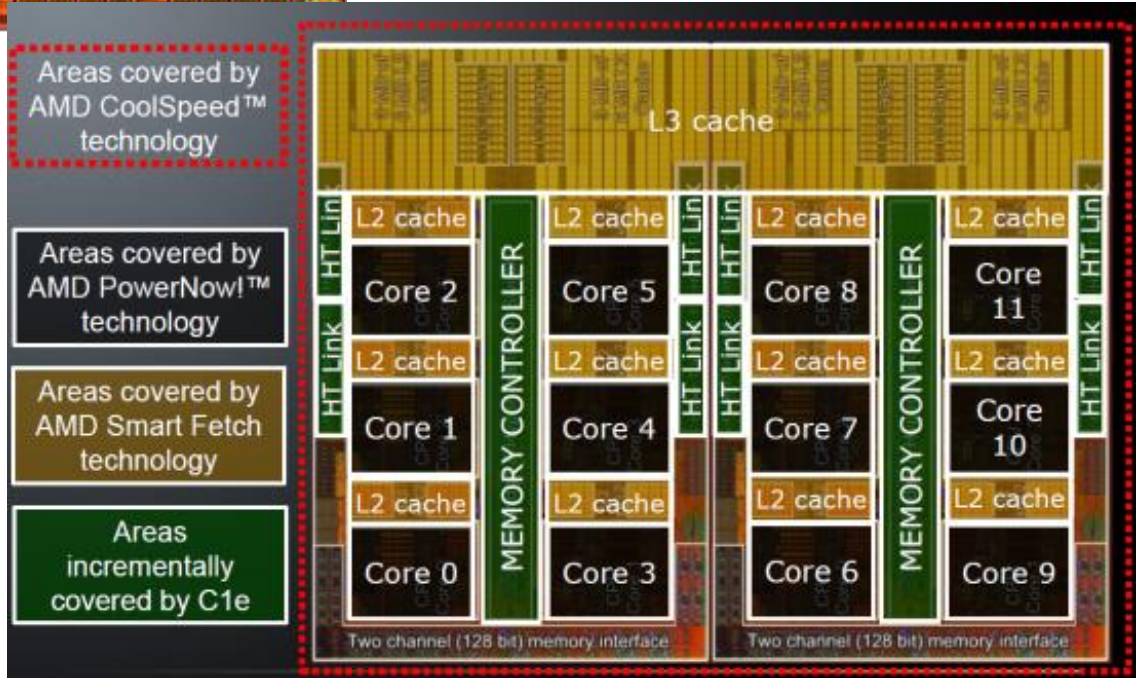

Multi-Cores
A Few Key Points

- Not Shared Memory (NUMA)
- Moore's Law rephrased:

  Nb. of Cores double every 18 to 24 months
- Key expected Milestones: Cores per Chips (OTS)
  - 2012: 32 to 64
  - 2013: 64 to 128
  - 2015: 128 to 256

  1 Million Cores Parallel Machines in 2014

  100 M cores coming in 2020
- Multi-Cores are NUMA, and turning Heterogeneous (GPU)

  They are turning into SoC with NoC: NOT SMP!

Key Point: Locality will more than ever be Fundamental

- Let the programmer control it
- No global shared memory

*At user choice **PGAS:** Partitioned Global Address Space*

ProActive Parallel Suite

Virtualization

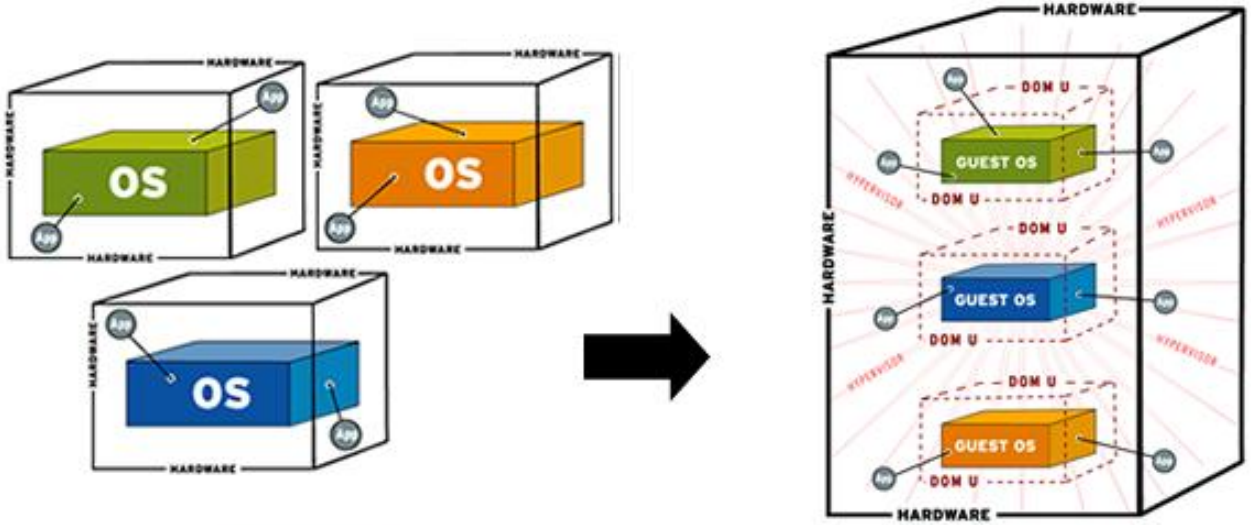

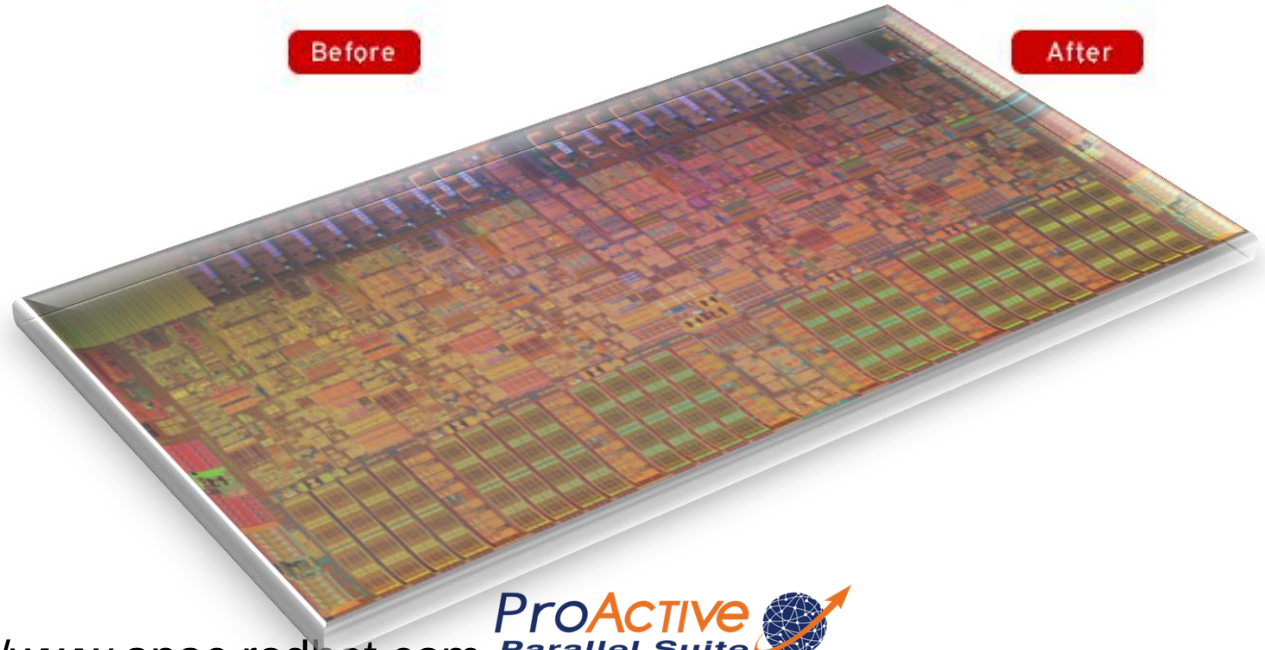

Source: http://www.apac.redhat.com

ProActive Parallel Suite

Virtualization

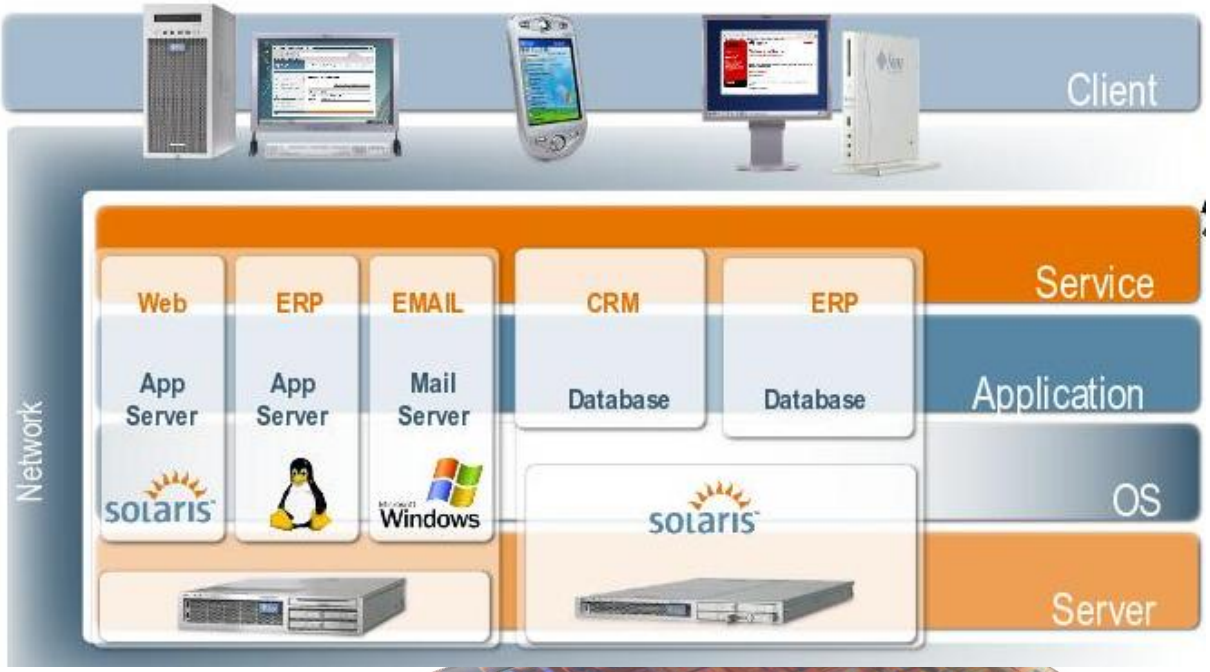

Sun, Blog [Marc Hamilton]

ProActive Parallel Suite

Virtualization

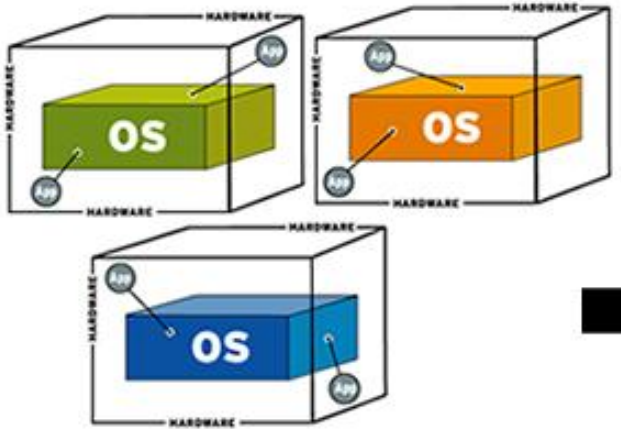

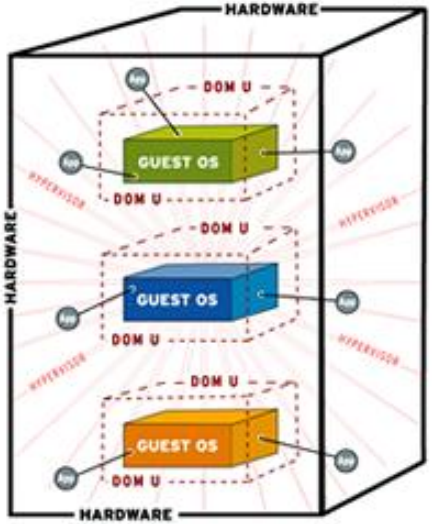

Before

After

What we Used to do as Syst. Admin.

ProActive Parallel Suite

With Virtualization + Software Appliance

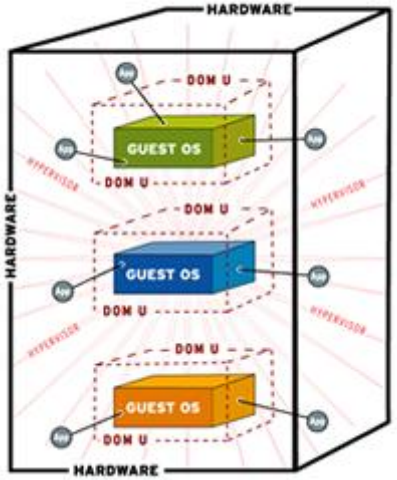

ProActive Parallel Suite

ProActive Parallel Suite

ProActive Parallel Suite

- Professional Open Source middleware for parallel, distributed, multi-core, Grid and Cloud computing

ProActive Parallel Suite

- Composed of three modules:

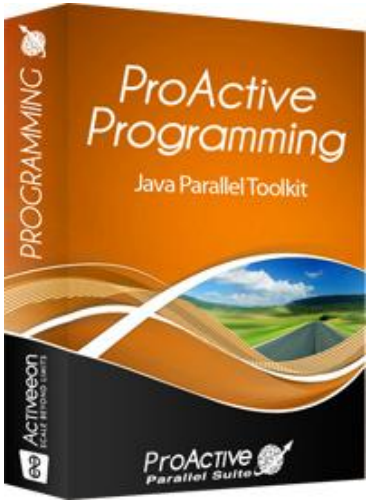

Java
Programming Library

Multiplatform
Job scheduler

Global
resource manager

Activeeon
SCALE BEYOND LIMITS

ProActive Parallel Suite

ProActive Parallel Suite

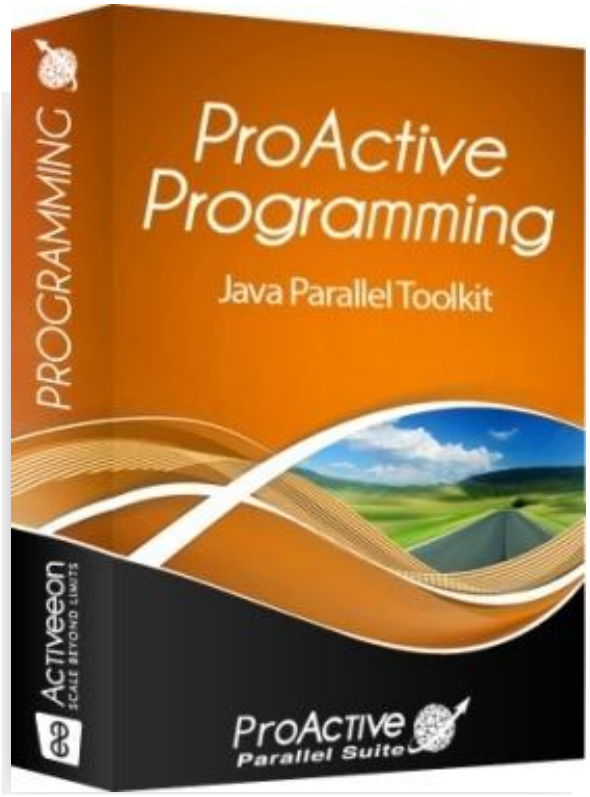

- Workflows in Java
- Master/Workers
- SPMD
- Components
- …

Core API
Active Objects
Asynchrony
Futures
Groups
Mobile Agents
MOP / AOP

ProActive Parallel Suite

ProActive Programming: Active Objects

ProActive Parallel Suite

ProActive Programming View

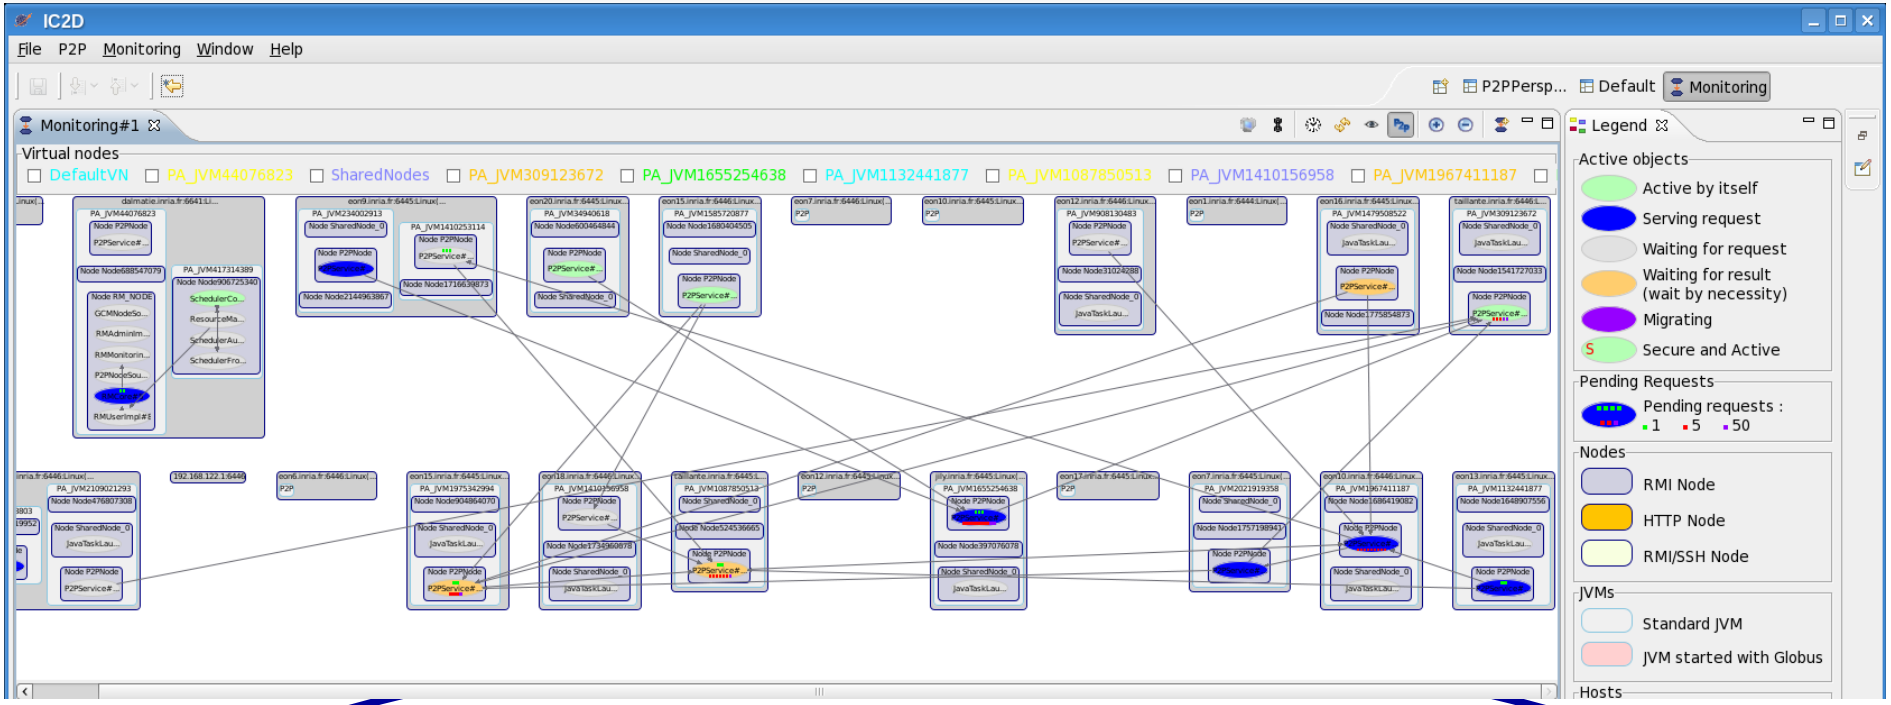

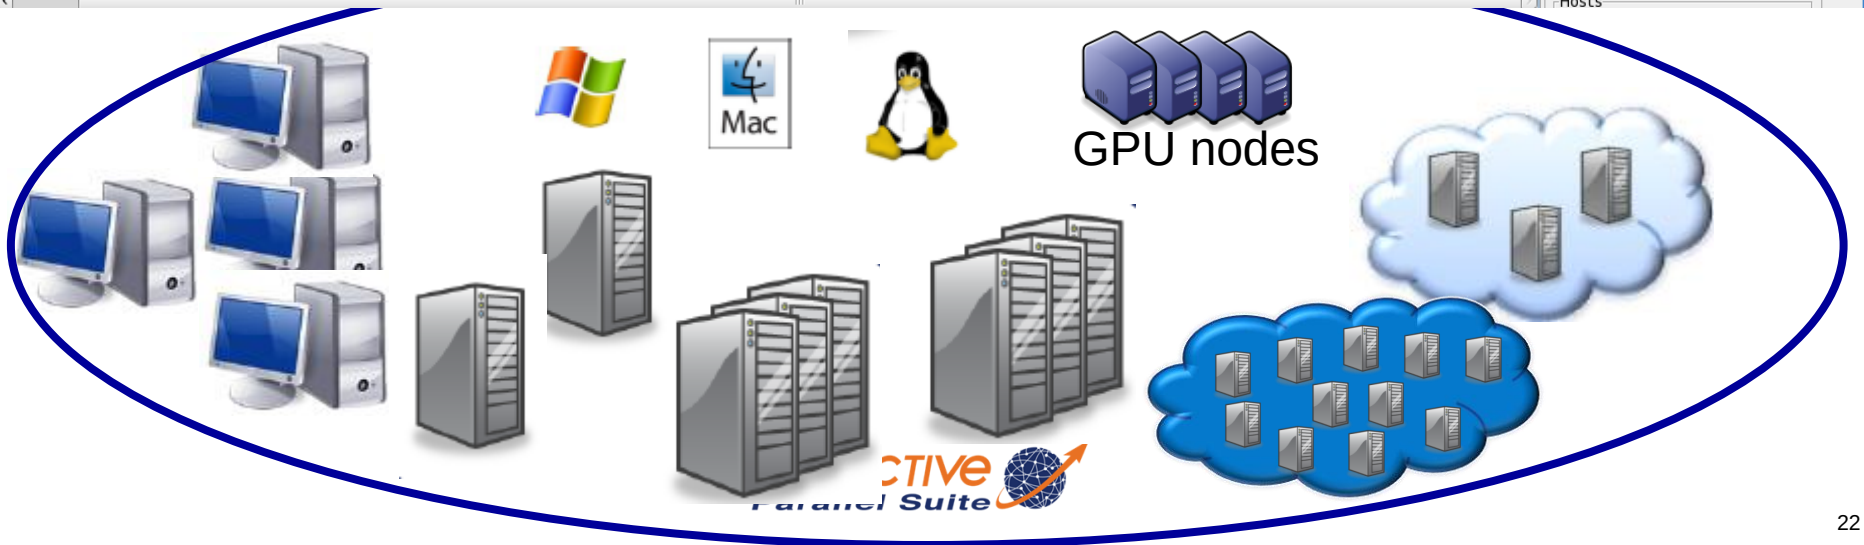

ProActive Programming View

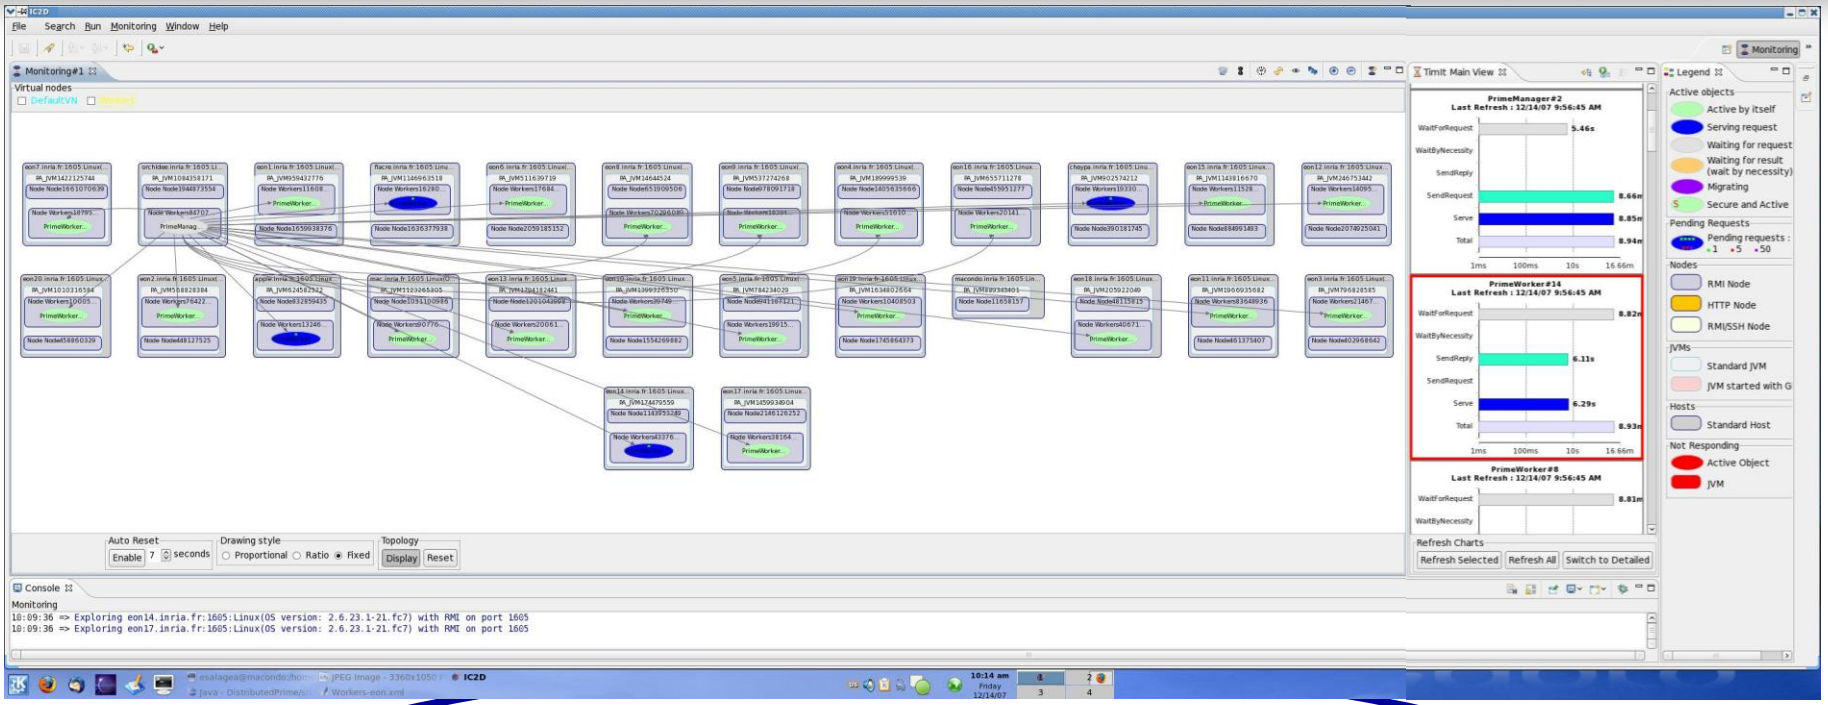

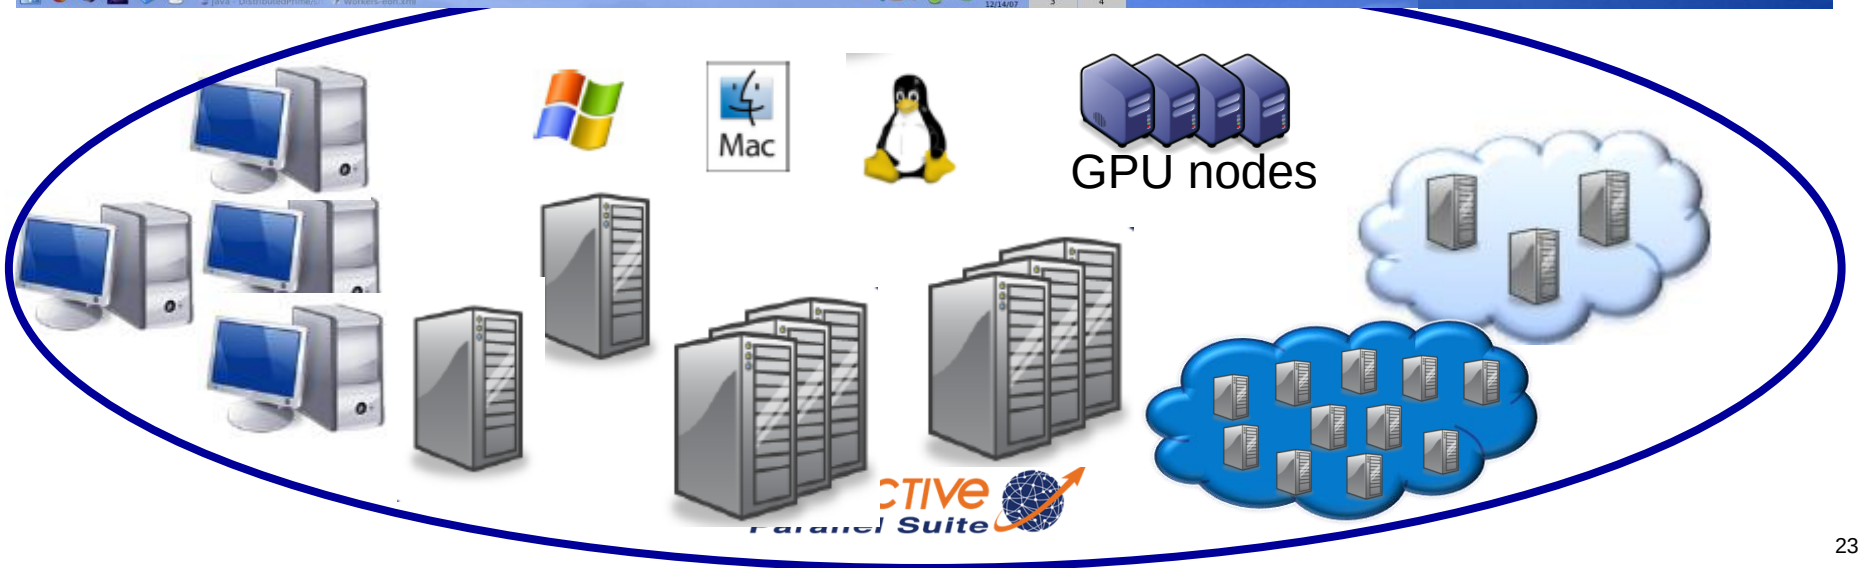

ProActive : Active objects

```
A ag = newActive ("A", […], VirtualNode)
V v1 = ag.foo (param);
V v2 = ag.bar (param);
...
v1.bar(); //Wait-By-Necessity
```

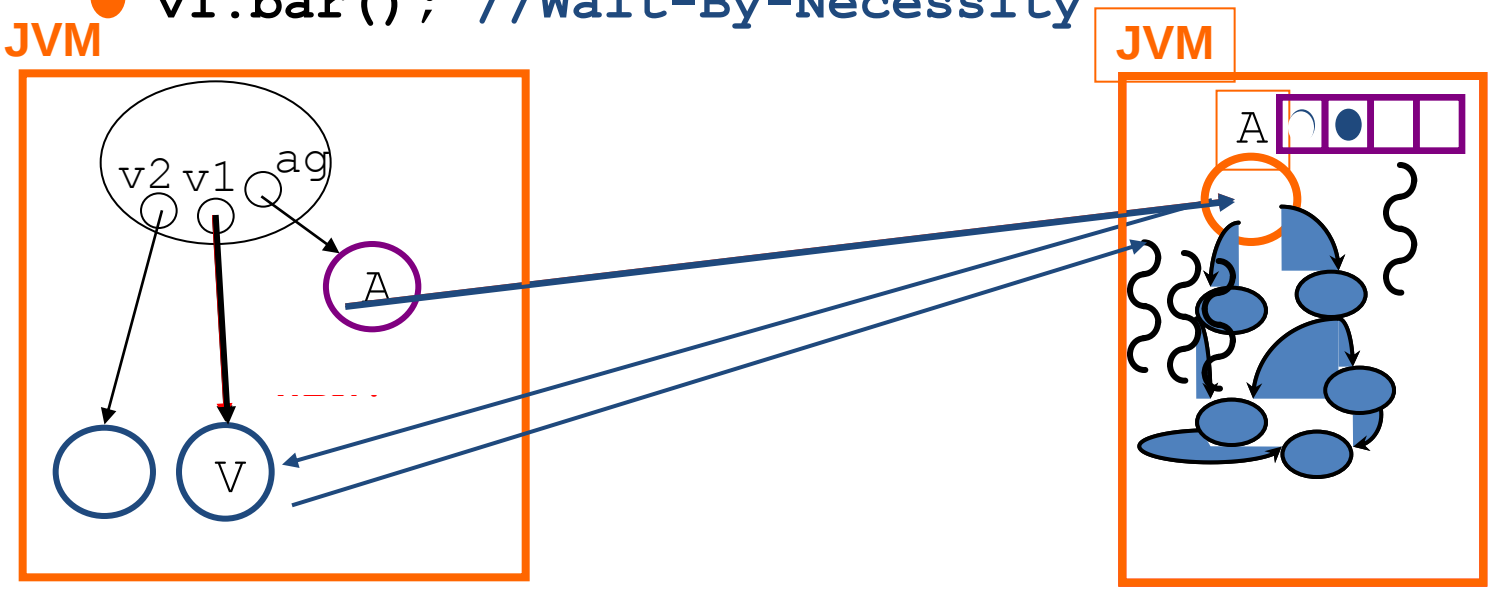

Java Object

Active Object

Req. Queue

Future Object

Proxy

Request

Thread

Wait-By-Necessity is a Dataflow Synchronization

ProActive Parallel Suite

24

Standard system at Runtime: No Sharing

NoC: Network On Chip

Proofs of Determinism

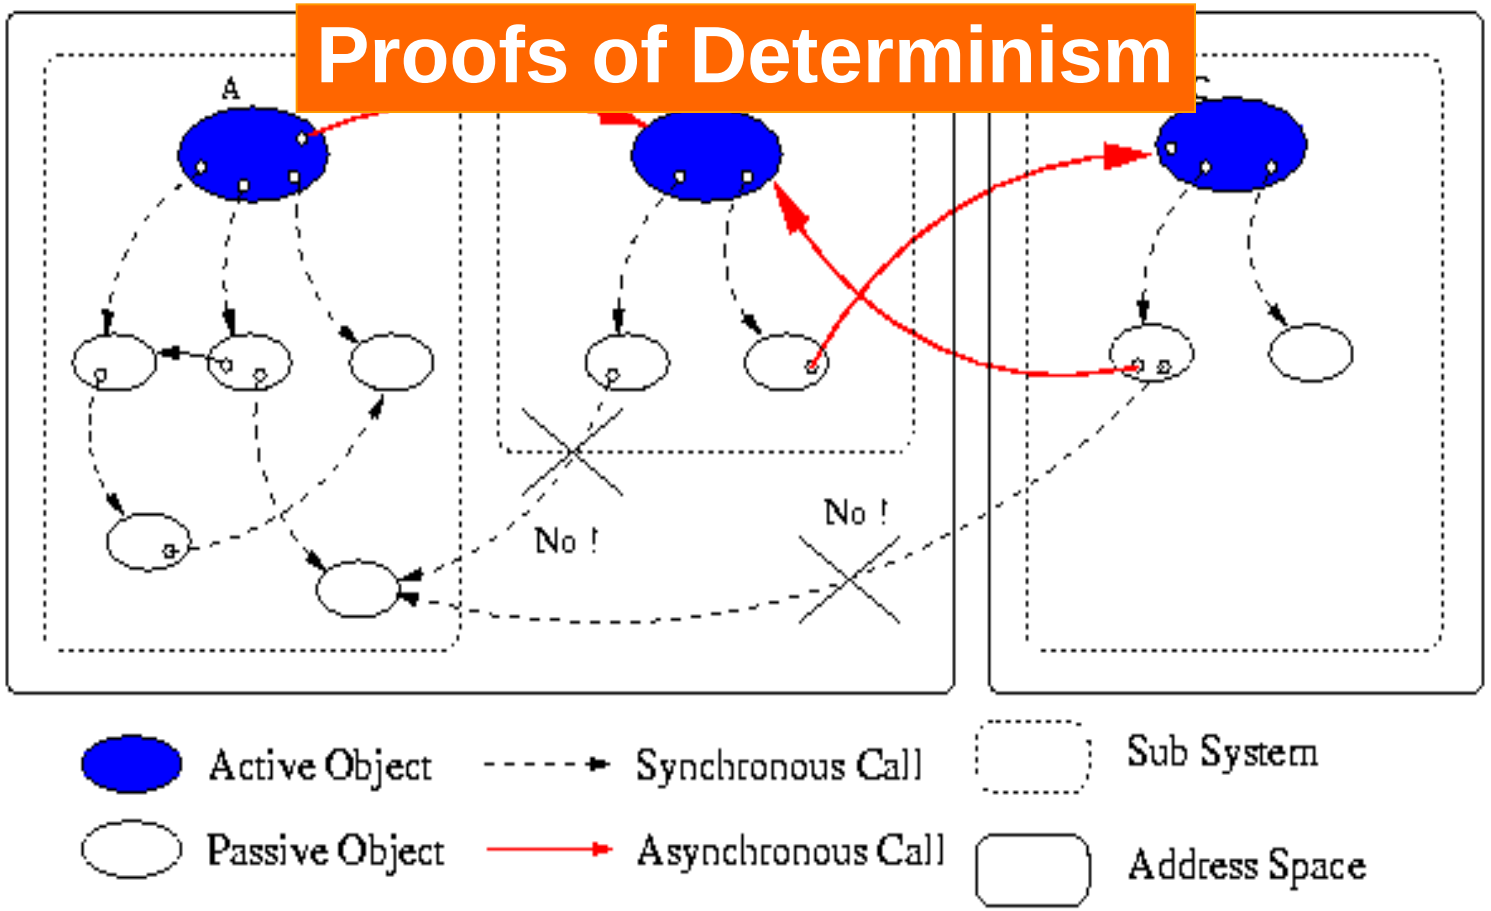

ProActive Parallel Suite

25

ASP: Asynchronous Sequential Processes

$$\frac{(a,\sigma) \rightarrow_S (a',\sigma')}{\alpha[a;\sigma;\iota;F;R;f] \parallel P \longrightarrow \alpha[a';\sigma';\iota;F;R;f] \parallel P}\ (\text{LOCAL})$$

Local

$$\frac{\begin{array}{c}\gamma \text{ fresh activity} \qquad \iota' \notin dom(\sigma) \qquad \sigma' = \{\iota' \mapsto AO(\gamma)\} :: \sigma \\ \sigma_\gamma = copy(\iota'',\sigma) \qquad Service = (\text{ if } m_j = \emptyset \text{ then } FifoService \text{ else } \iota''.m_j())\end{array}}{\begin{array}{c}\alpha[\mathcal{R}[Active(\iota'',m_j)];\sigma;\iota;F;R;f] \parallel P \\ \longrightarrow \alpha[\mathcal{R}[\iota'];\sigma';\iota;F;R;f] \parallel \gamma[Service;\sigma_\gamma;\iota'';\emptyset;\emptyset;\emptyset] \parallel P\end{array}}\ (\text{NEWACT})$$

Creating an Activity

$$\frac{\begin{array}{c}\sigma_\alpha(\iota) = AO(\beta) \qquad \iota'' \notin dom(\sigma_\beta) \qquad f_i^{\alpha\rightarrow\beta} \text{ new future} \qquad \iota_f \notin dom(\sigma_\alpha) \\ \sigma'_\beta = Copy\&Merge(\sigma_\alpha,\iota' \ ; \ \sigma_\beta,\iota'') \qquad \sigma'_\alpha = \{\iota_f \mapsto fut(f_i^{\alpha\rightarrow\beta})\} :: \sigma_\alpha\end{array}}{\begin{array}{c}\alpha[\mathcal{R}[\iota.m_j(\iota')];\sigma_\alpha;\iota_\alpha;F_\alpha;R_\alpha;f_\alpha] \parallel \beta[a_\beta;\sigma_\beta;\iota_\beta;F_\beta;R_\beta;f_\beta] \parallel P \longrightarrow \\ \alpha[\mathcal{R}[\iota_f];\sigma'_\alpha;\iota_\alpha;F_\alpha;R_\alpha;f_\alpha] \parallel \beta[a_\beta;\sigma'_\beta;\iota_\beta;F_\beta;R_\beta :: [m_j;\iota'';f_i^{\alpha\rightarrow\beta}];f_\beta] \parallel P\end{array}}\ (\text{REQUEST})$$

Sending a Request

$$\frac{R = R' :: [m_j;\iota_r;f'] :: R'' \qquad m_j \in M \qquad \forall m \in M, m \notin R'}{\alpha[\mathcal{R}[Serve(M)];\sigma;\iota;F;R;f] \parallel P \longrightarrow \alpha[\iota.m_j(\iota_r) \Uparrow f,\mathcal{R}[[]];\sigma;\iota;F;R' :: R'';f'] \parallel P}\ (\text{SERVE})$$

$$\frac{\iota' \notin dom(\sigma) \qquad F' = F :: \{f \mapsto \iota'\} \qquad \sigma' = Copy\&Merge(\sigma,\iota \ ; \ \sigma,\iota')}{\alpha[\iota \Uparrow (f',a);\sigma;\iota;F;R;f] \parallel P \longrightarrow \alpha[a;\sigma';\iota;F';R;f'] \parallel P}\ (\text{ENDSERVICE})$$

Service

$$\frac{\sigma_\alpha(\iota) = fut(f_i^{\gamma\rightarrow\beta}) \qquad F_\beta(f_i^{\gamma\rightarrow\beta}) = \iota_f \qquad \sigma'_\alpha = Copy\&Merge(\sigma_\beta,\iota_f \ ; \ \sigma_\alpha,\iota)}{\begin{array}{c}\alpha[a_\alpha;\sigma_\alpha;\iota_\alpha;F_\alpha;R_\alpha;f_\alpha] \parallel \beta[a_\beta;\sigma_\beta;\iota_\beta;F_\beta;R_\beta;f_\beta] \parallel P \longrightarrow \\ \alpha[a_\alpha;\sigma'_\alpha;\iota_\alpha;F_\alpha;R_\alpha;f_\alpha] \parallel \beta[a_\beta;\sigma_\beta;\iota_\beta;F_\beta;R_\beta;f_\beta] \parallel P\end{array}}\ (\text{REPLY})$$

Sending a Reply

Distributed Objects On Chip

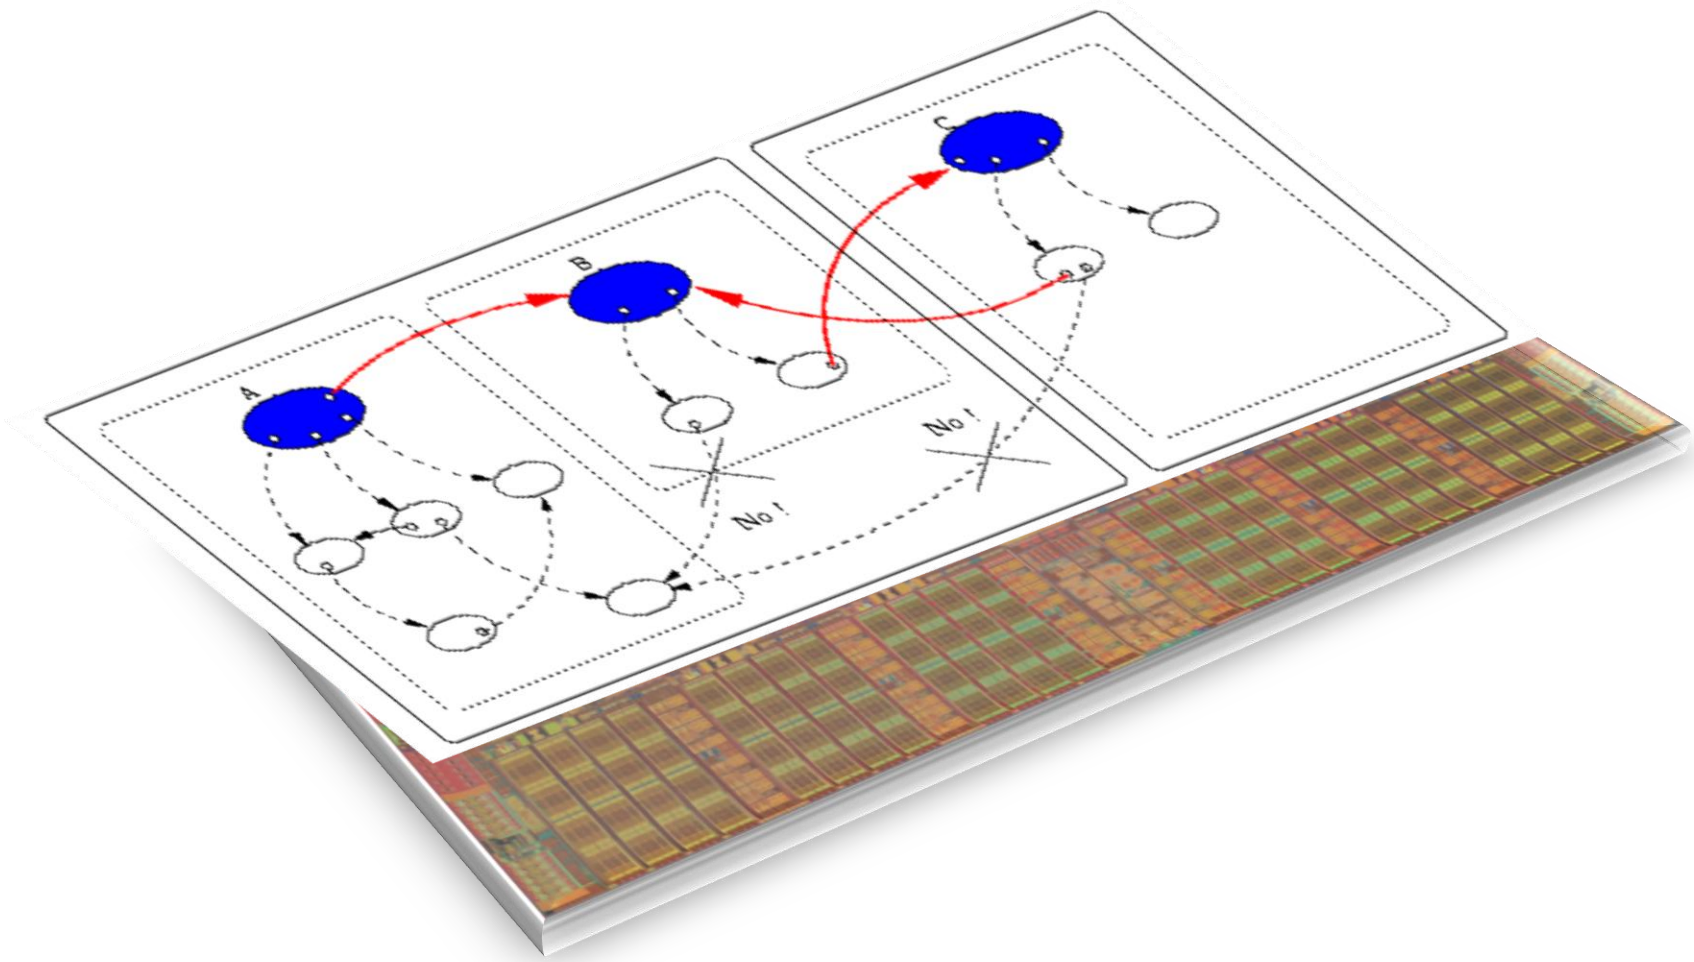

ProActive Parallel Suite

Distributed Objects On Chip, Boards, Clouds

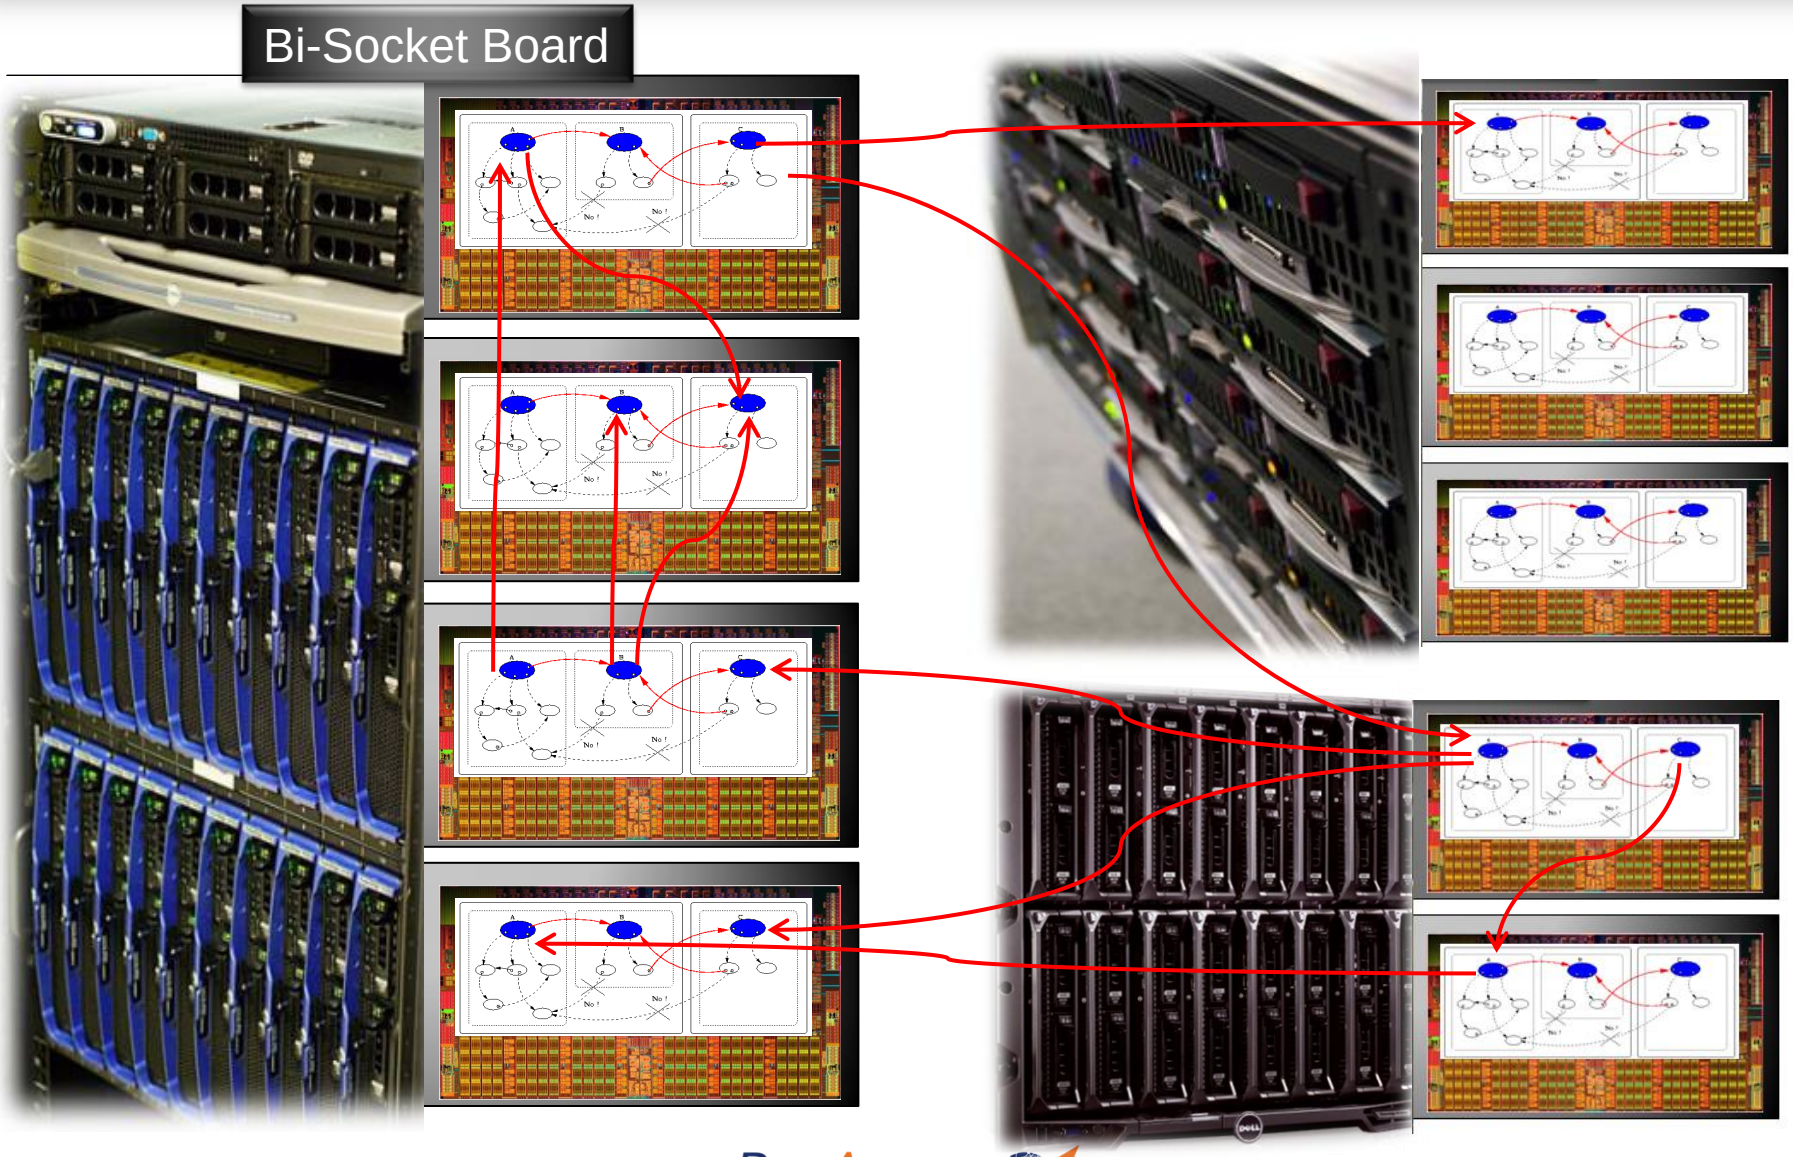

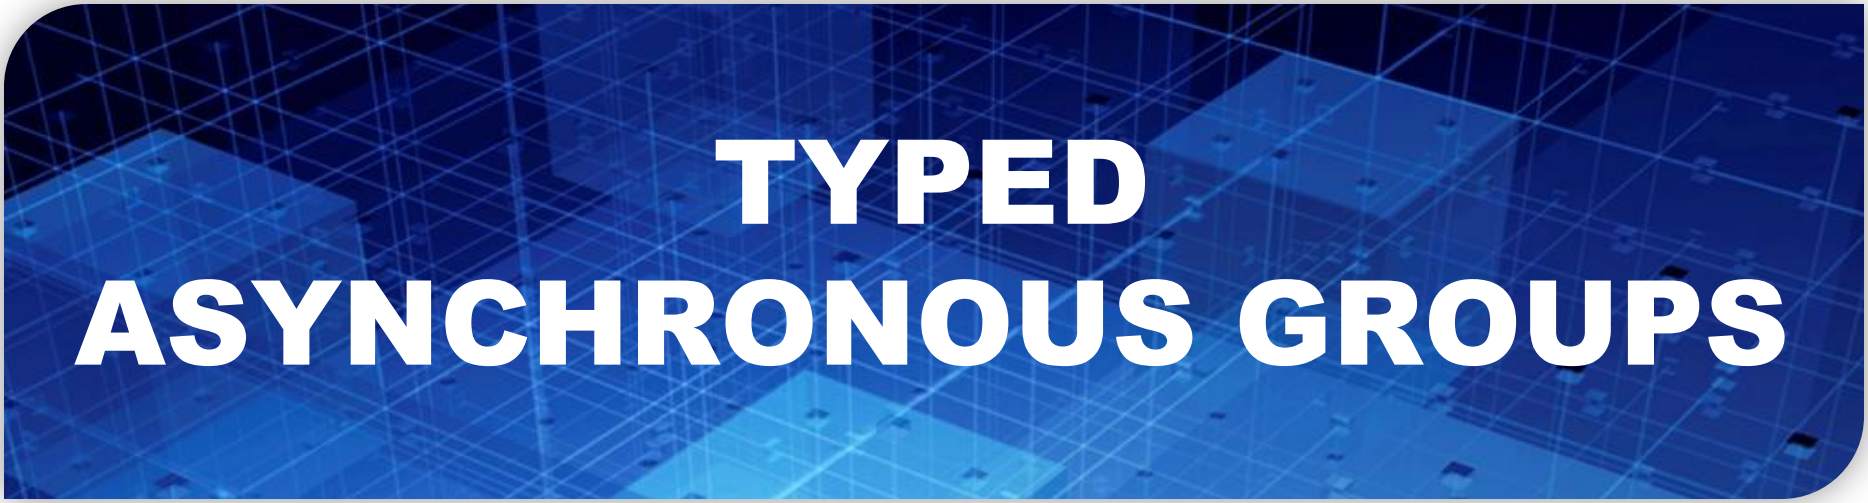
TYPED
ASYNCHRONOUS GROUPS

Broadcast and Scatter

Broadcast is the default behavior
Use a group as parameter, Scattered depends on rankings

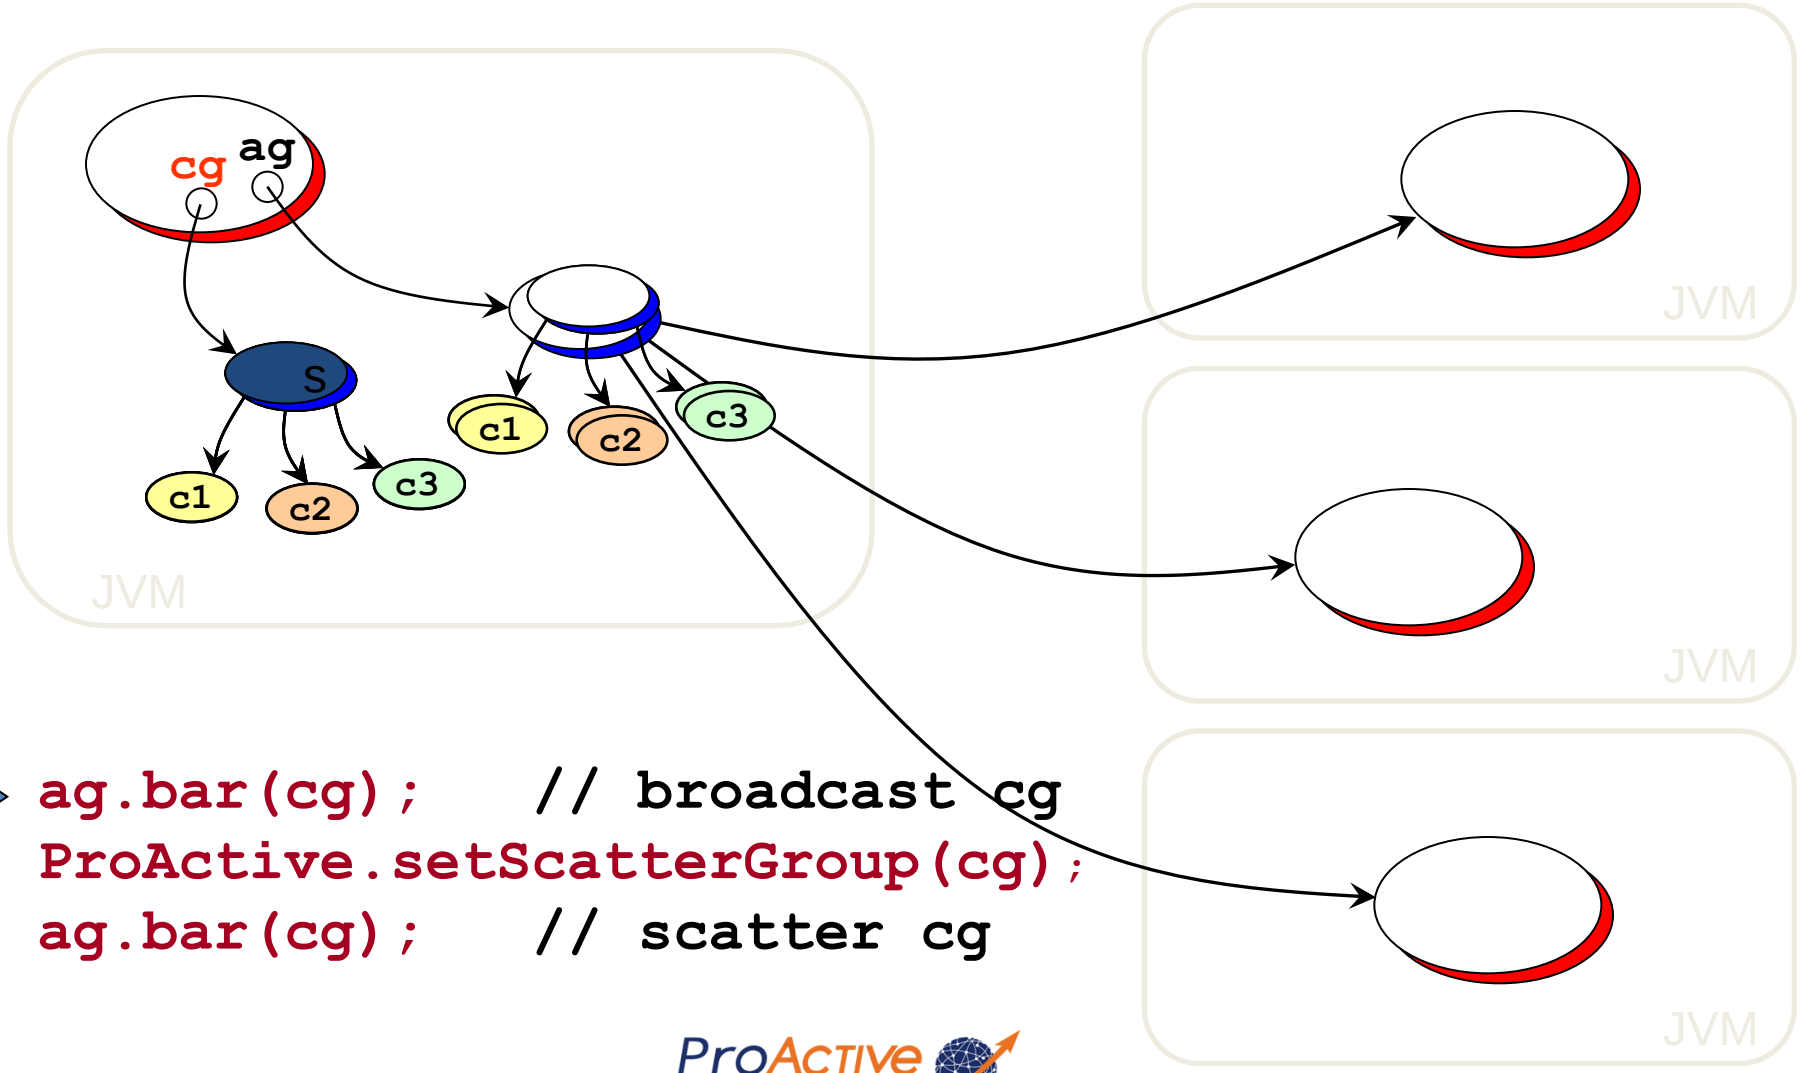


```
ag.bar(cg);    // broadcast cg
ProActive.setScatterGroup(cg);
ag.bar(cg);    // scatter cg
```

ProActive Parallel Suite

30

Dynamic Dispatch Group

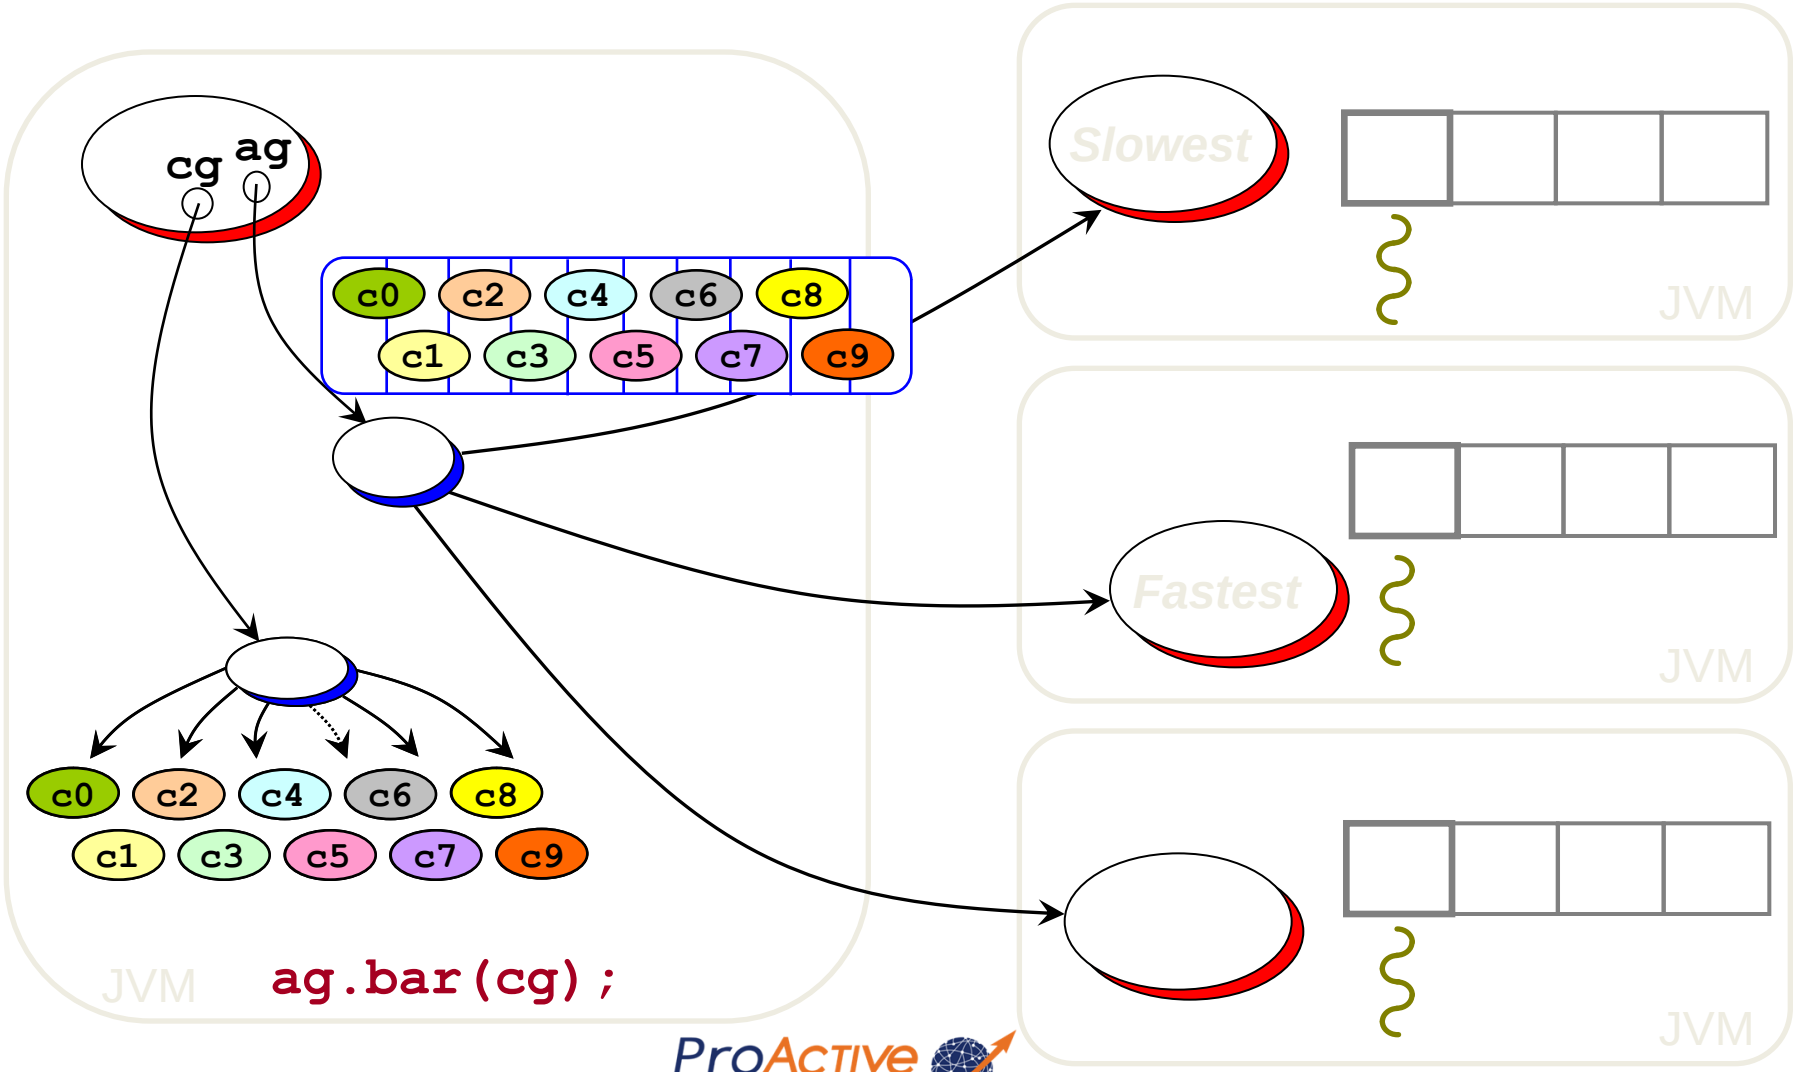

ProActive Parallel Suite

31

Object-Oriented SPMD

Key Point

- **"MPI and programming languages from the 60's will not make it"**
  - **Jack Dongarra,** 2/13/2009,
  - Wake Forest University talk
- **"It is time to get ride of MPI"**
  - **Alan Edelman, MIT,** 06/16/2010,
  - ScilabTec'10 Users' Day

OO SPMD: Object-Oriented SPMD

```
A ag = newSPMDGroup ("A", […], VirtualNode)
        // In each member
        myGroup.barrier ("2D"); // Global Barrier
        myGroup.barrier ("vertical"); // Any Barrier
        myGroup.barrier ("north","south","east","west");
```

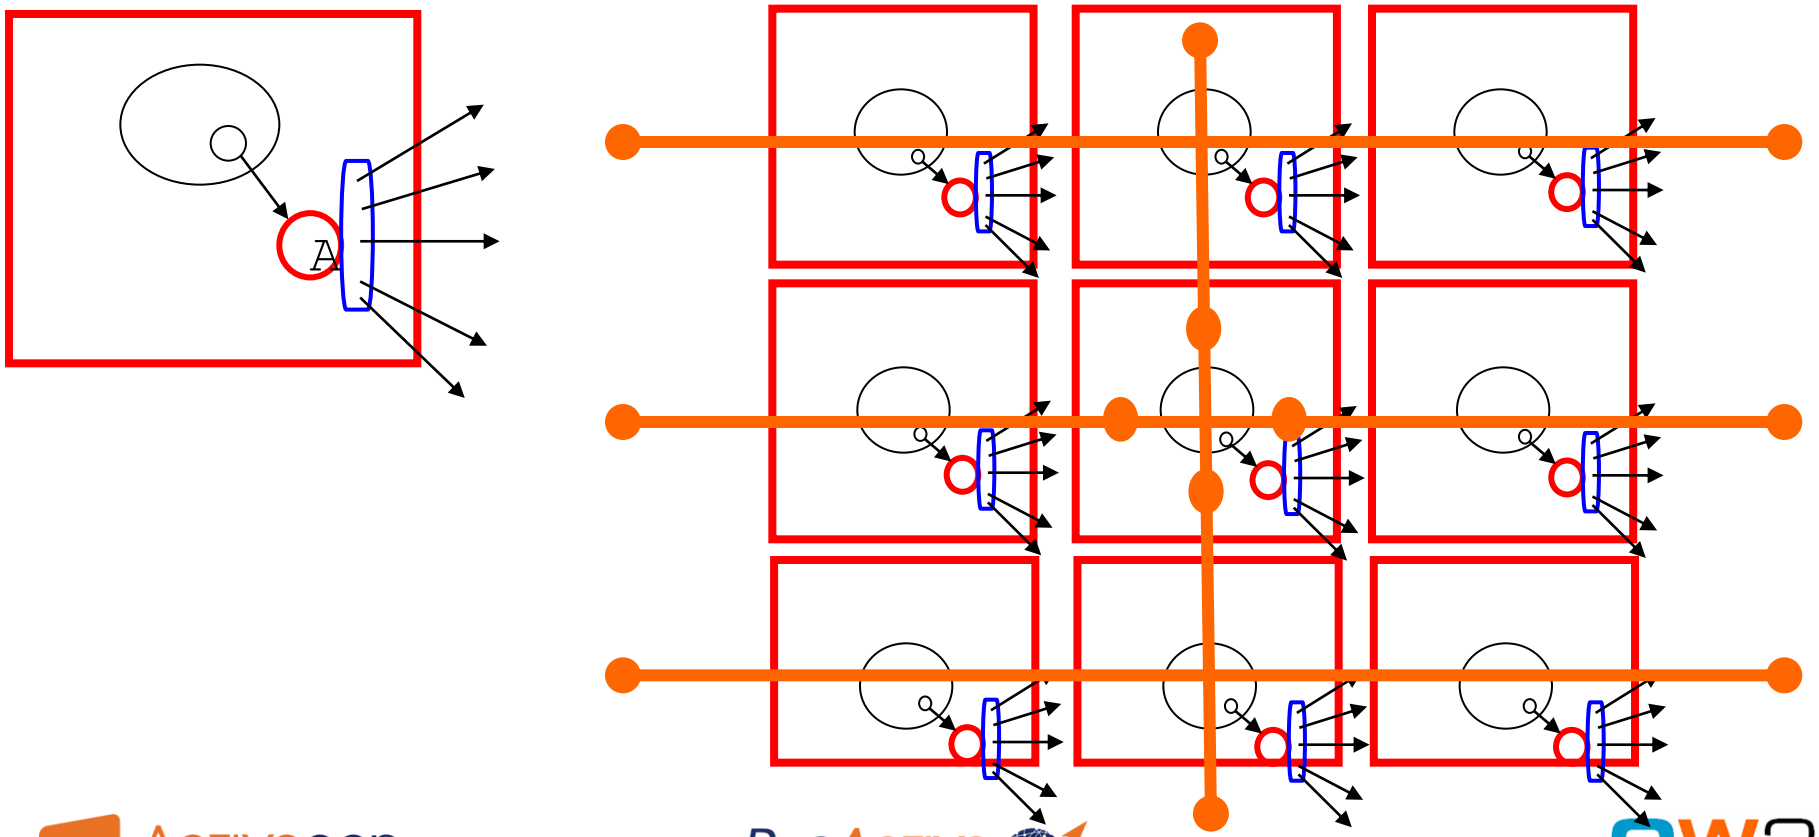

NAS Parallel Benchmarks

- Experimented on 3D ElectroMagnetism, and Nasa Benchmarks

- **Designed by NASA to evaluate benefits of high performance systems**
- **Strongly based on CFD**
- **5 benchmarks (kernels) to test different aspects of a system**
- **2 categories or focus variations:**
  - **communication intensive and computation intensive**

Communication Intensive
CG Kernel (Conjugate Gradient)

- Floating point operations
- Eigen value computation
- High number of unstructured communications

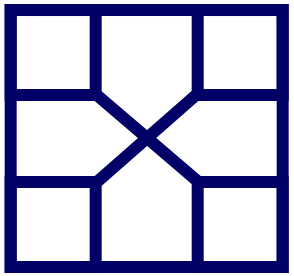

- 12000 calls/node
- 570 MB sent/node
- 1 min 32
- 65 % comms/WT

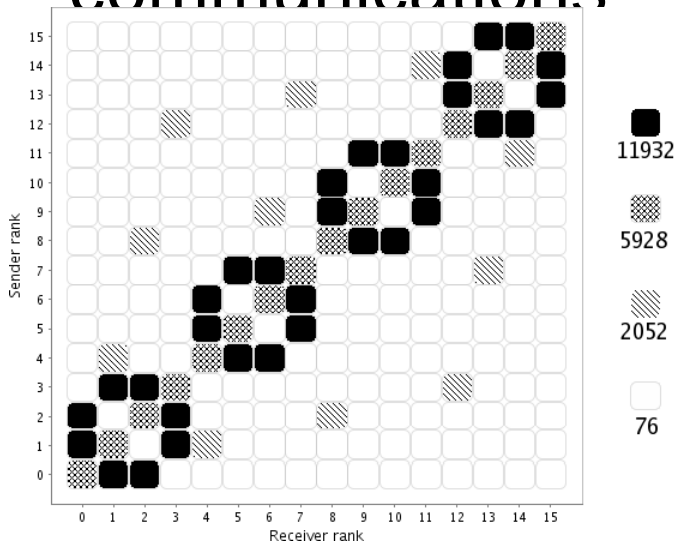

Message density distribution

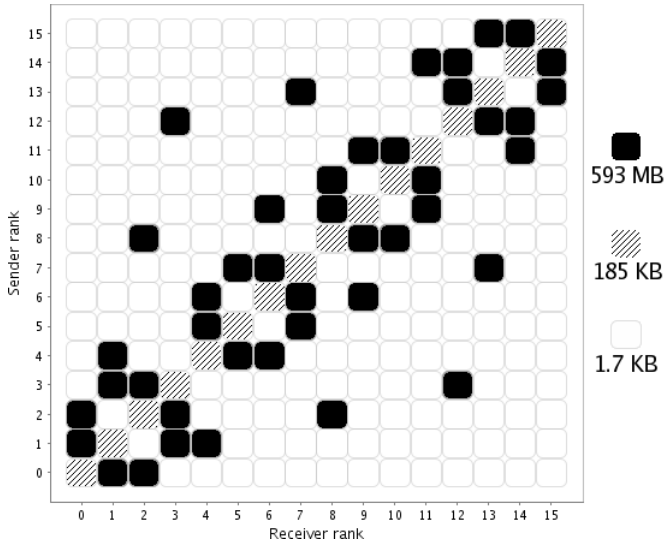

Data density distribution

Communication Intensive
CG Kernel (Conjugate Gradient)

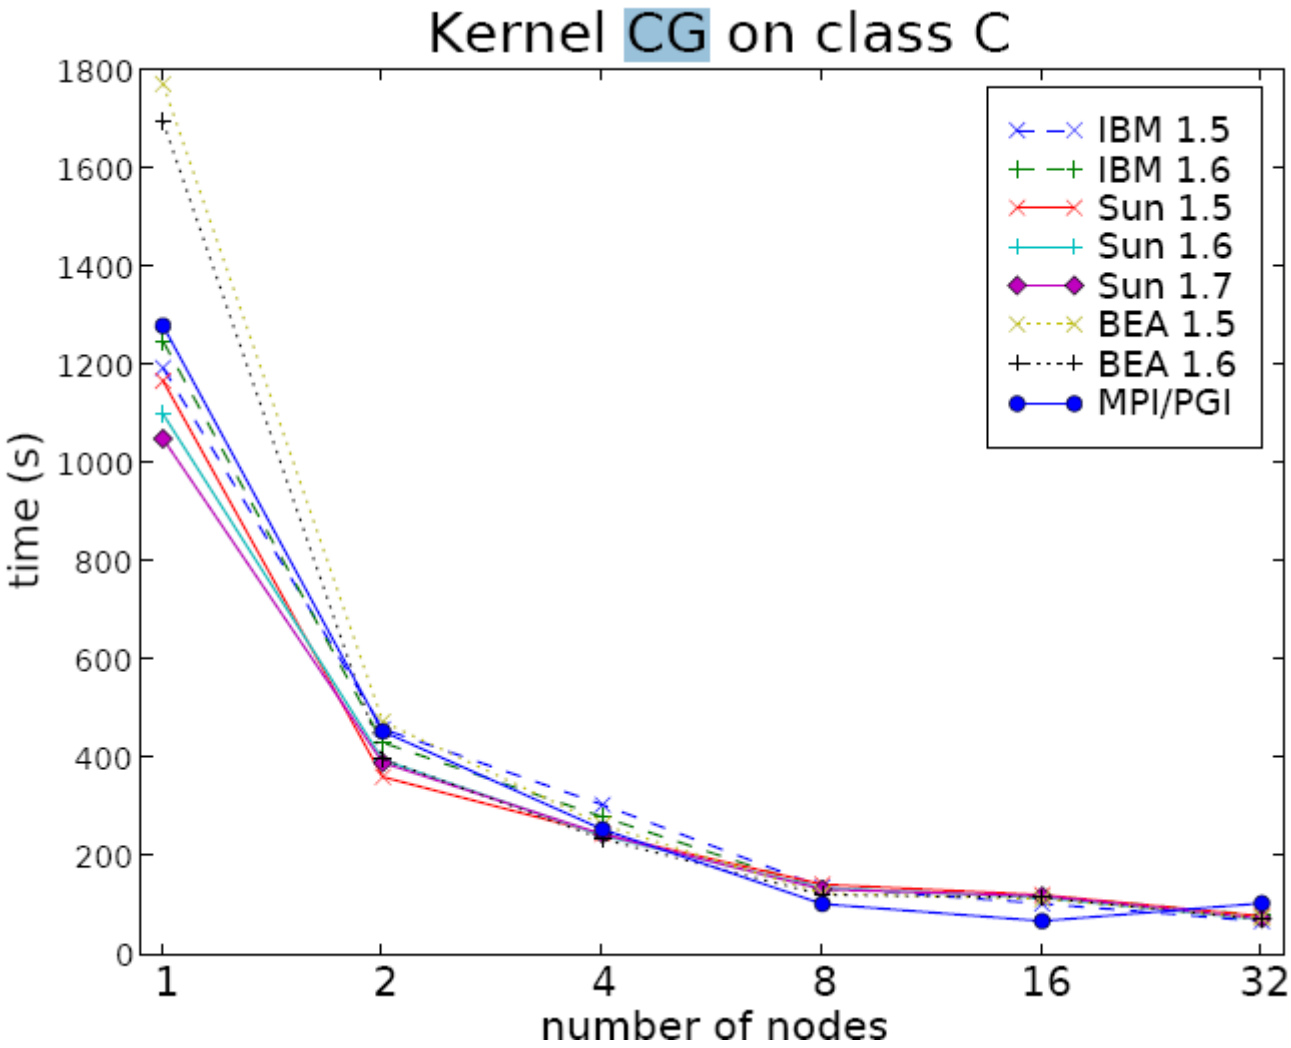

Comparable Performances

Scheduling and Workflow Orchestration

ProActive Scheduling

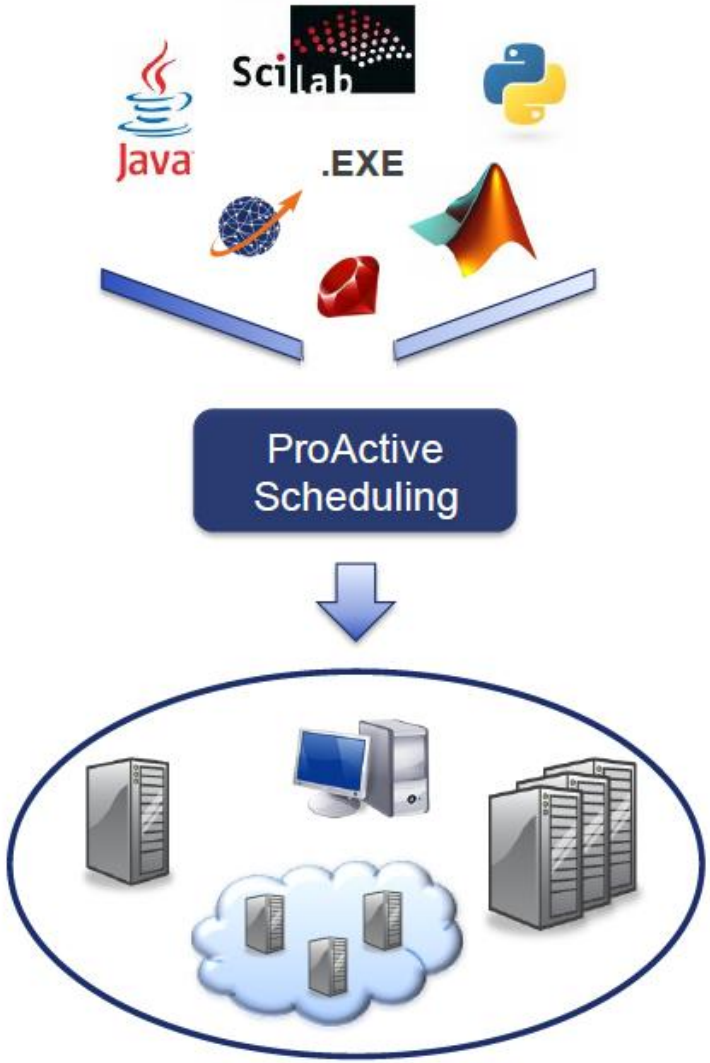

ProActive Parallel Suite

39

ProActive Scheduling

- Job Scheduling
- Optimize:
  - workload distribution
  - computing resources usage
  - software licences
- Straightforward batch, jobs and workflow construction
- Any kind of jobs supported:
  - Java, Applications and Scripts
- Data management with automatic transfer

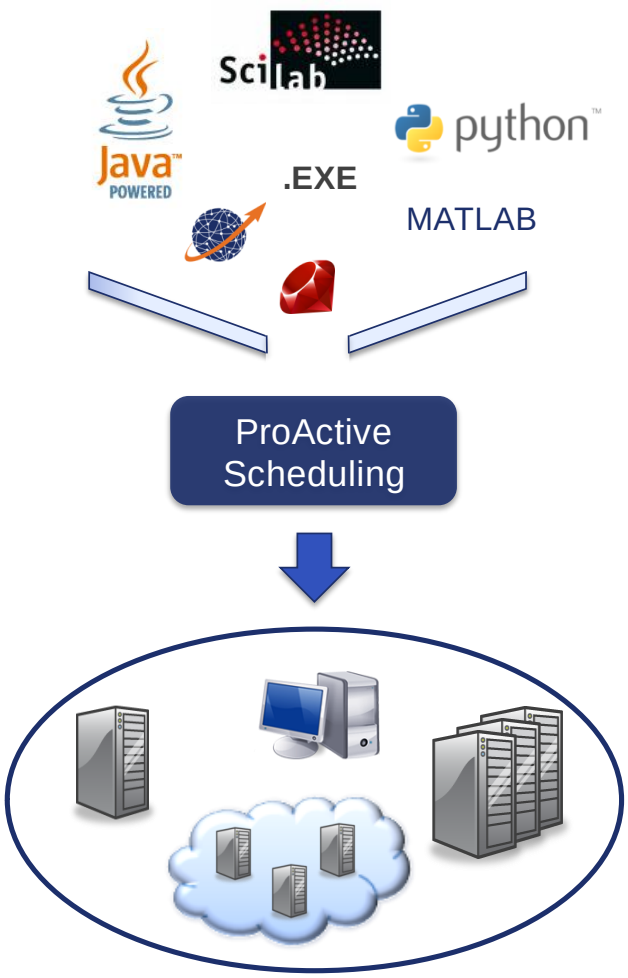

ProActive Scheduling Big Picture

ProActive Scheduler

File Window Help

Scheduler

Jobs

Pending (674)

Id	State	User	Priority	Name
1996	Pending	jl	Normal	job_with_dep
1997	Pending	jl	Normal	job_with_dep
1998	Pending	jl	Normal	job_with_dep
1999	Pending	jl	Normal	job_with_dep
2000	Pending	jl	Normal	job_with_dep
2001	Pending	jl	Normal	job_with_dep
2002	Pending	jl	Normal	job_with_dep
2003	Pending	jl	Normal	job_with_dep
2004	Pending	jl	Normal	job_with_dep
2005	Pending	jl	Normal	job_with_dep
2006	Pending	jl	Normal	job_with_dep
2007	Pending	jl	Normal	job_with_dep
2008	Pending	jl	Normal	job_with_dep
2009	Pending	jl	Normal	job_with_dep
2010	Pending	jl	Normal	job_with_dep

Running (60)

Id	State	Progress	# Finished	User	Prior
1313	Running		4/8	user1	Norr
1314	Running		4/8	user1	Norr
1315	Running		7/8	admin	Norr
1316	Running		4/8	user1	Norr
1317	Running		7/8	admin	Norr
1318	Running		4/8	user1	Norr
1319	Running		7/8	admin	Norr
1320	Running		3/8	user1	Norr
1321	Running		7/8	admin	Norr
1322	Running		3/8	user1	Norr
1323	Running		7/8	admin	Norr
1324	Running		2/8	user1	Norr
1325	Running		2/8	user1	Norr
1326	Running		2/8	user1	Norr
1327	Running		2/8	user1	Norr

Finished (31)

Id	State	User	Priority	Name
010	Finished	jl	Low	job_proActive
008	Finished	jl	Low	job_proActive
005	Finished	jl	Low	job_proActive
001	Finished	jl	Low	job_proActive
006	Finished	jl	Low	job_proActive
004	Finished	jl	Low	job_proActive
003	Finished	jl	Low	job_proActive
009	Finished	jl	Low	job_proActive
007	Finished	jl	Low	job_proActive
002	Finished	jl	Low	job_proActive
245	Finished	user1	Normal	job_with_dep
246	Finished	user1	Normal	job_with_dep
247	Finished	user1	Normal	job_with_dep
252	Finished	admin	Normal	job_with_dep
253	Finished	admin	Normal	job_with_dep

RESUMED

Console Tasks

Job 2008 has 8 tasks

Id	State	Name	Host name	Start time	Finished time	Re-run	Description
200800	Submitted	task4	n/a	Not yet	Not yet	0/2	This task will sleep 5s
200800	Submitted	task2	n/a	Not yet	Not yet	0/1	This task will sleep 10s
200800	Submitted	task6	n/a	Not yet	Not yet	0/1	This task will sleep 8s
200800	Submitted	task1	n/a	Not yet	Not yet	0/2	This task will sleep 6s
200800	Submitted	task5	n/a	Not yet	Not yet	0/1	This task will sleep 2s
200800	Submitted	task7	n/a	Not yet	Not yet	0/2	This task will sleep 6s
200800	Submitted	task3	n/a	Not yet	Not yet	0/1	This task will sleep 4s
200800	Submitted	task8	n/a	Not yet	Not yet	0/1	This task will sleep 6s

Job Info Result Preview

Property	Value
Id	2008
State	Pending
Name	job_with_dep
Priority	Normal
Pending tasks number	0
Running tasks number	0
Finished tasks number	0
Total tasks number	8
Submitted time	09:40:06 03/12/08
Started time	Not yet
Finished time	Not yet

RESOU

va

ion

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

ProActive Orchestration Portal

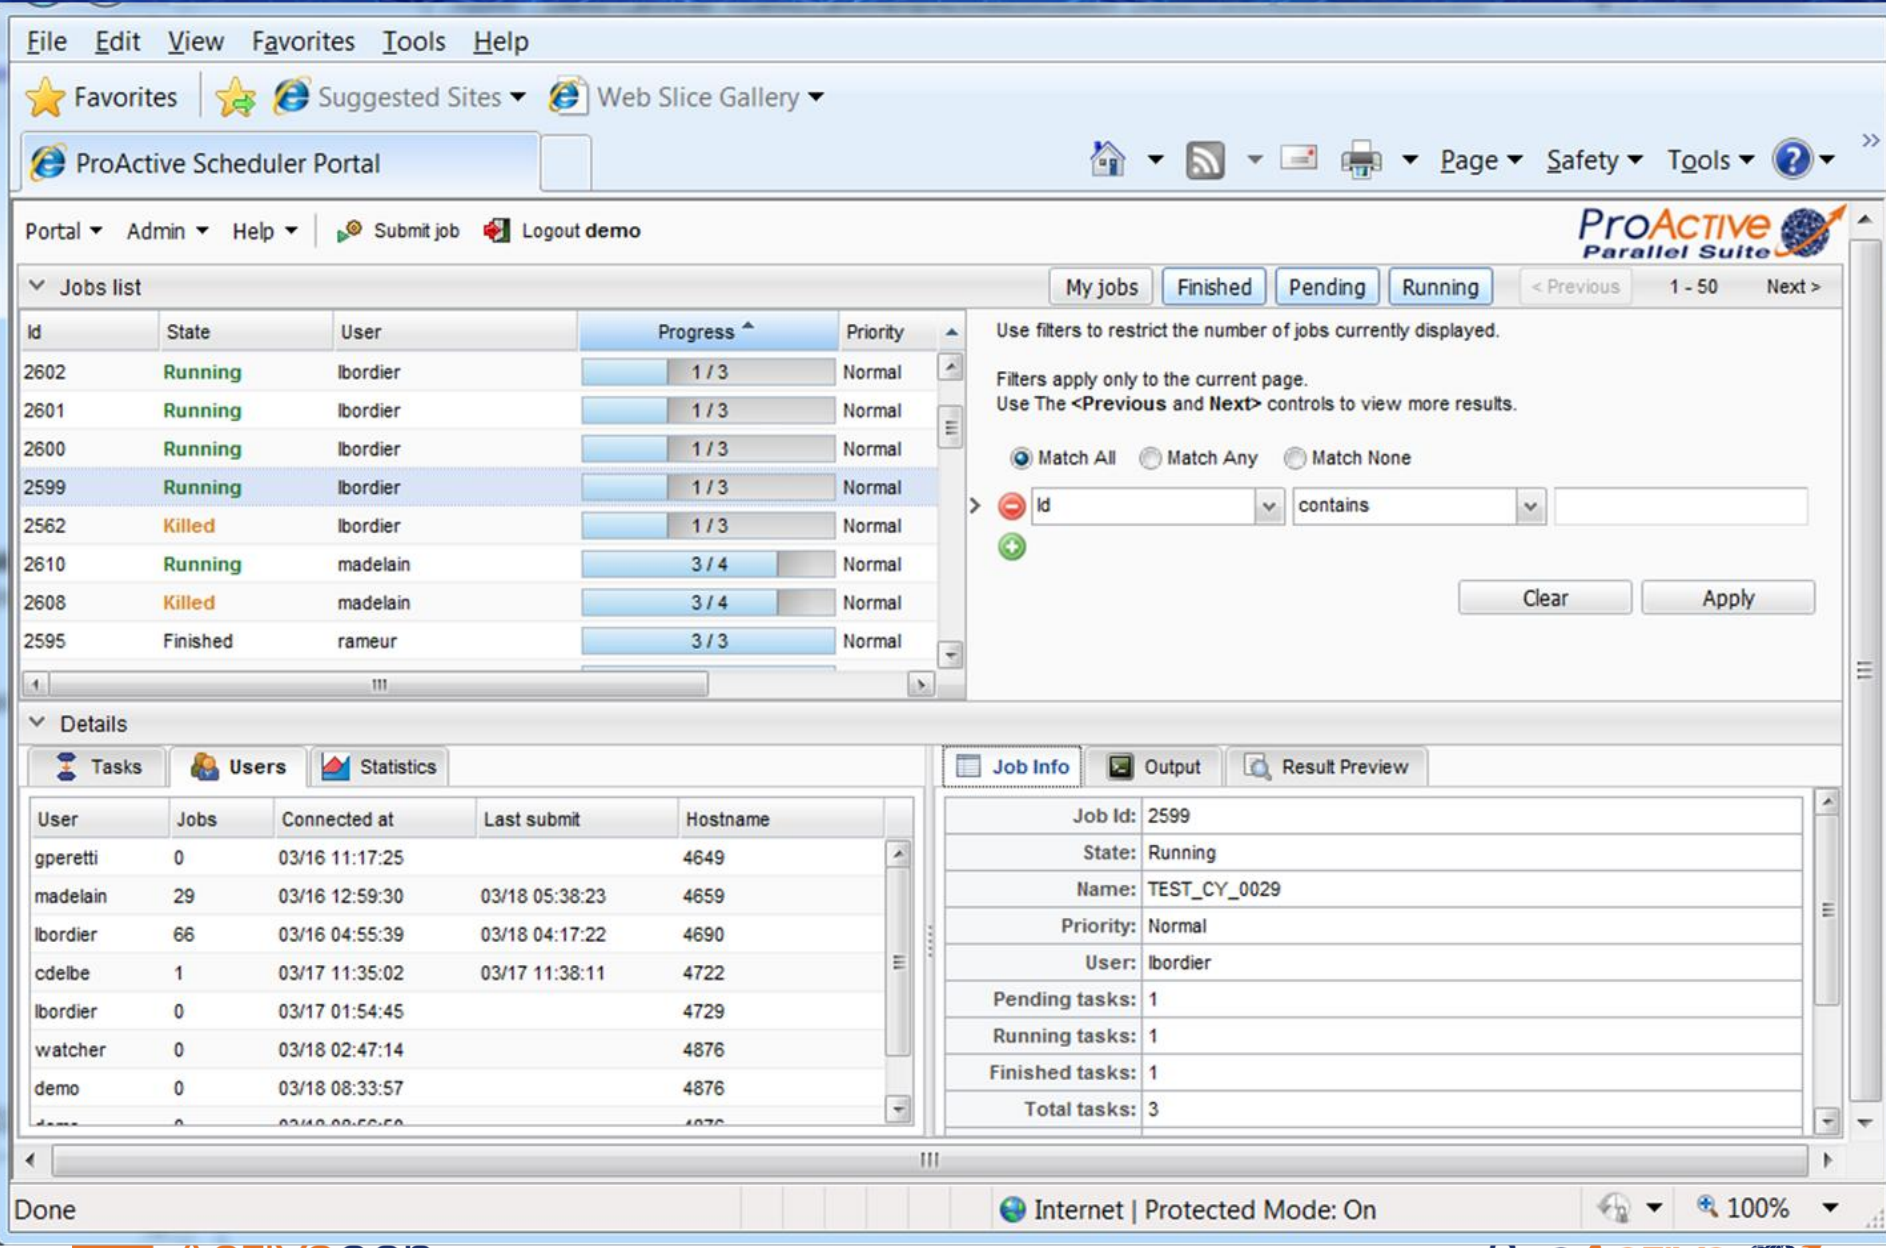

ProActive Scheduling & Orchestration

- Provides highly configurable scheduling policies
- Self-healing with automatic restart from latest valid point
- Hardware and software fault tolerant mechanism for task execution
- Graphical interface and command line client
- Seamless integration with third-party application
  - Java, Web Service and C/C++
- Accounting per user

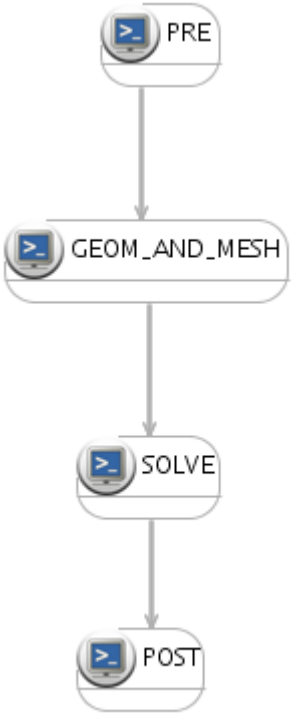

Workflow Example : Picture Denoising

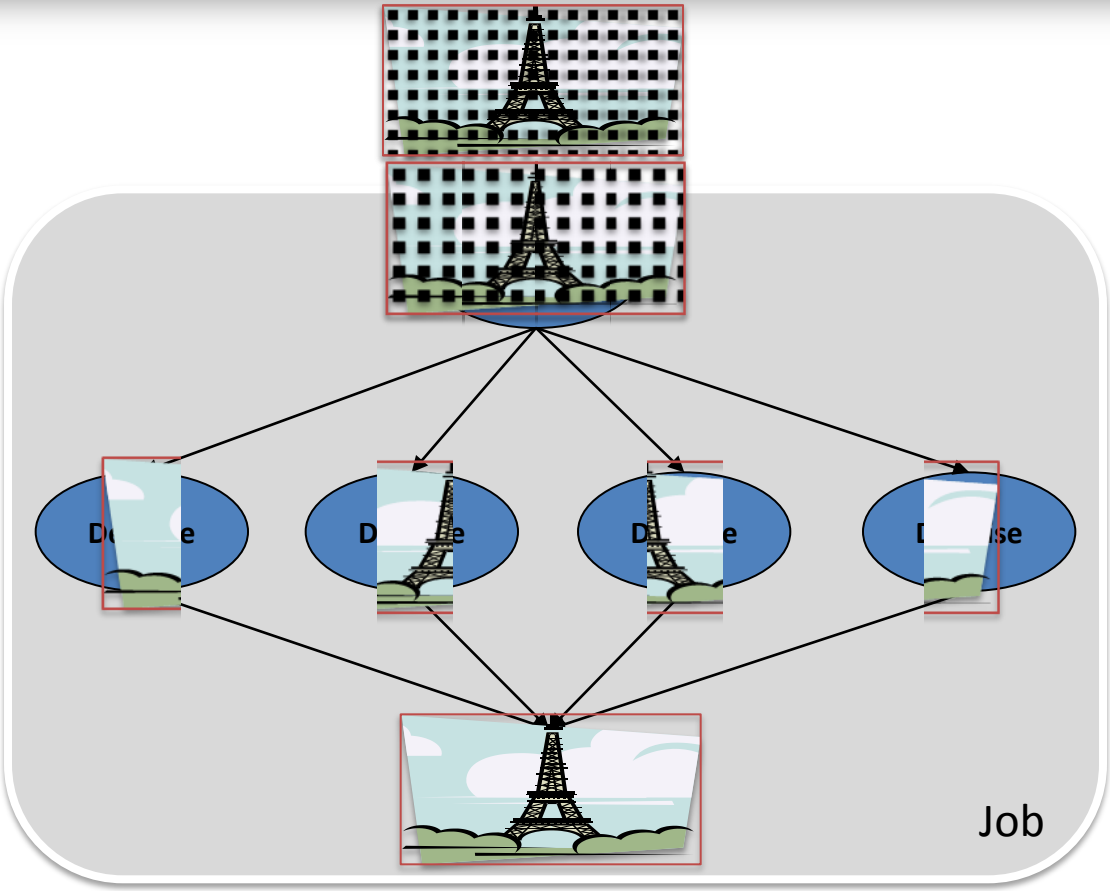

- with selection on native executable availability (ImageMagik, GREYstoration)
  - Multi-platform selection and command generation
- with file transfer in pre/post scripts

Workflow Studio

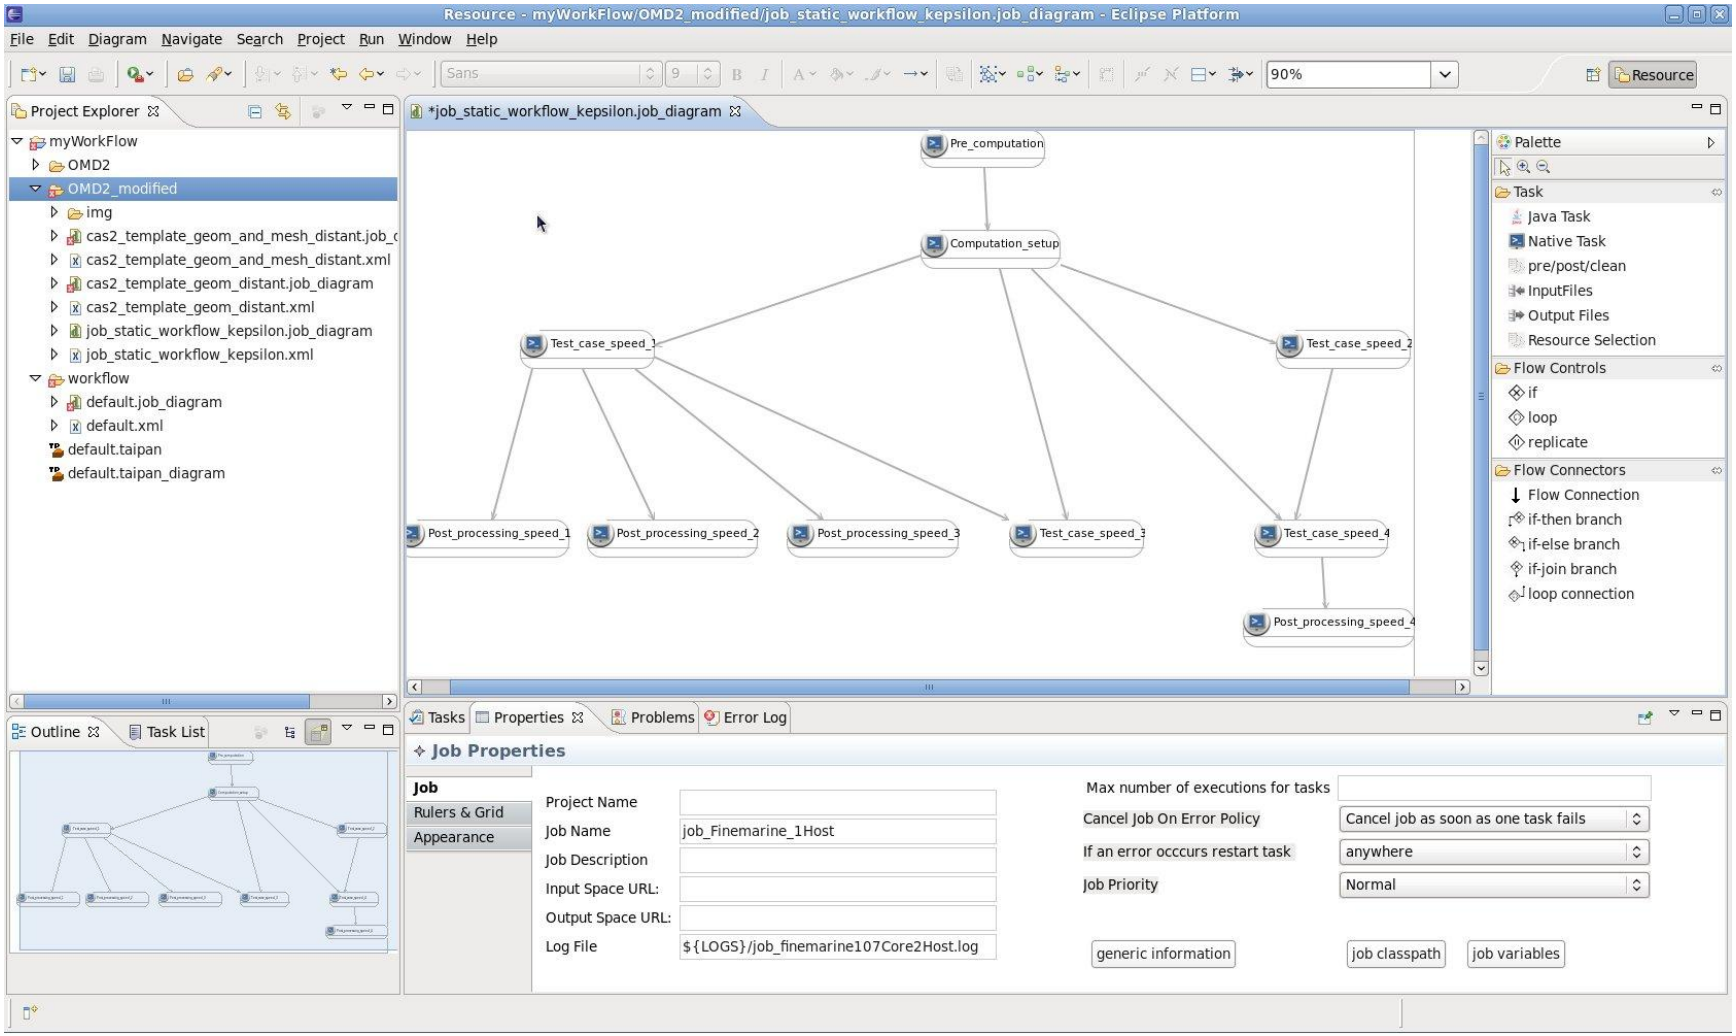

Activeeon
SCALE BEYOND LIMITS

Studio Demonstration

ProActive
Parallel Suite

Heterogeneous Resource Management

ProActive Resourcing

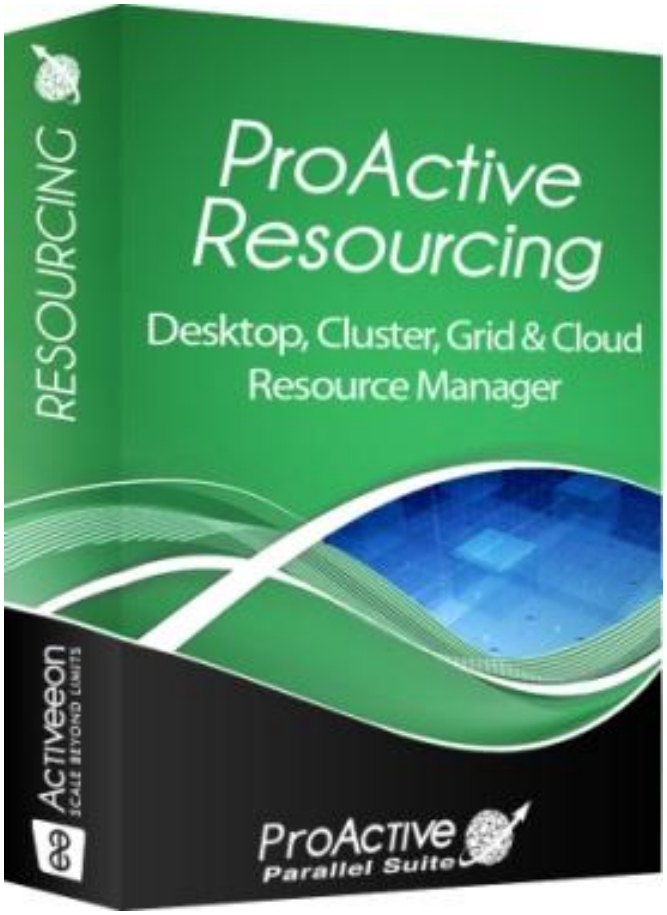

ProActive Parallel Suite

48

ProActive Resourcing

- Virtualizes organizations' existing infrastructure for a heightened computing power
- Enables business driven computing resources acquisitions
  - Elastic computing platform
- Reaps the benefits from Clouds, e.g. Amazon EC2, and latest hardware architecture
- Aggregates and leverages any type of resources
- Manage your virtual machines
  - VMWare, Xen, KVM, Hyper-V, ...
- Accounting per resource provider

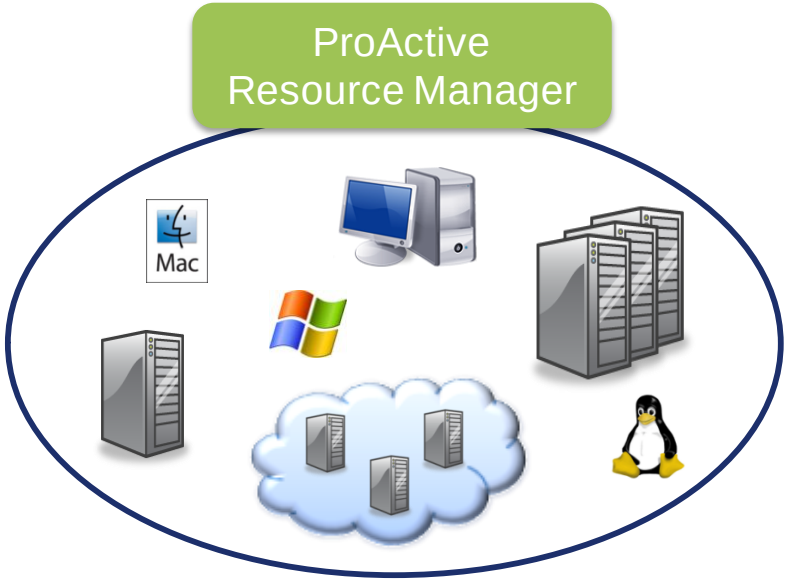

Activeeon
SCALE BEYOND LIMITS

ProActive
Parallel Suite

Resource Manager GUI

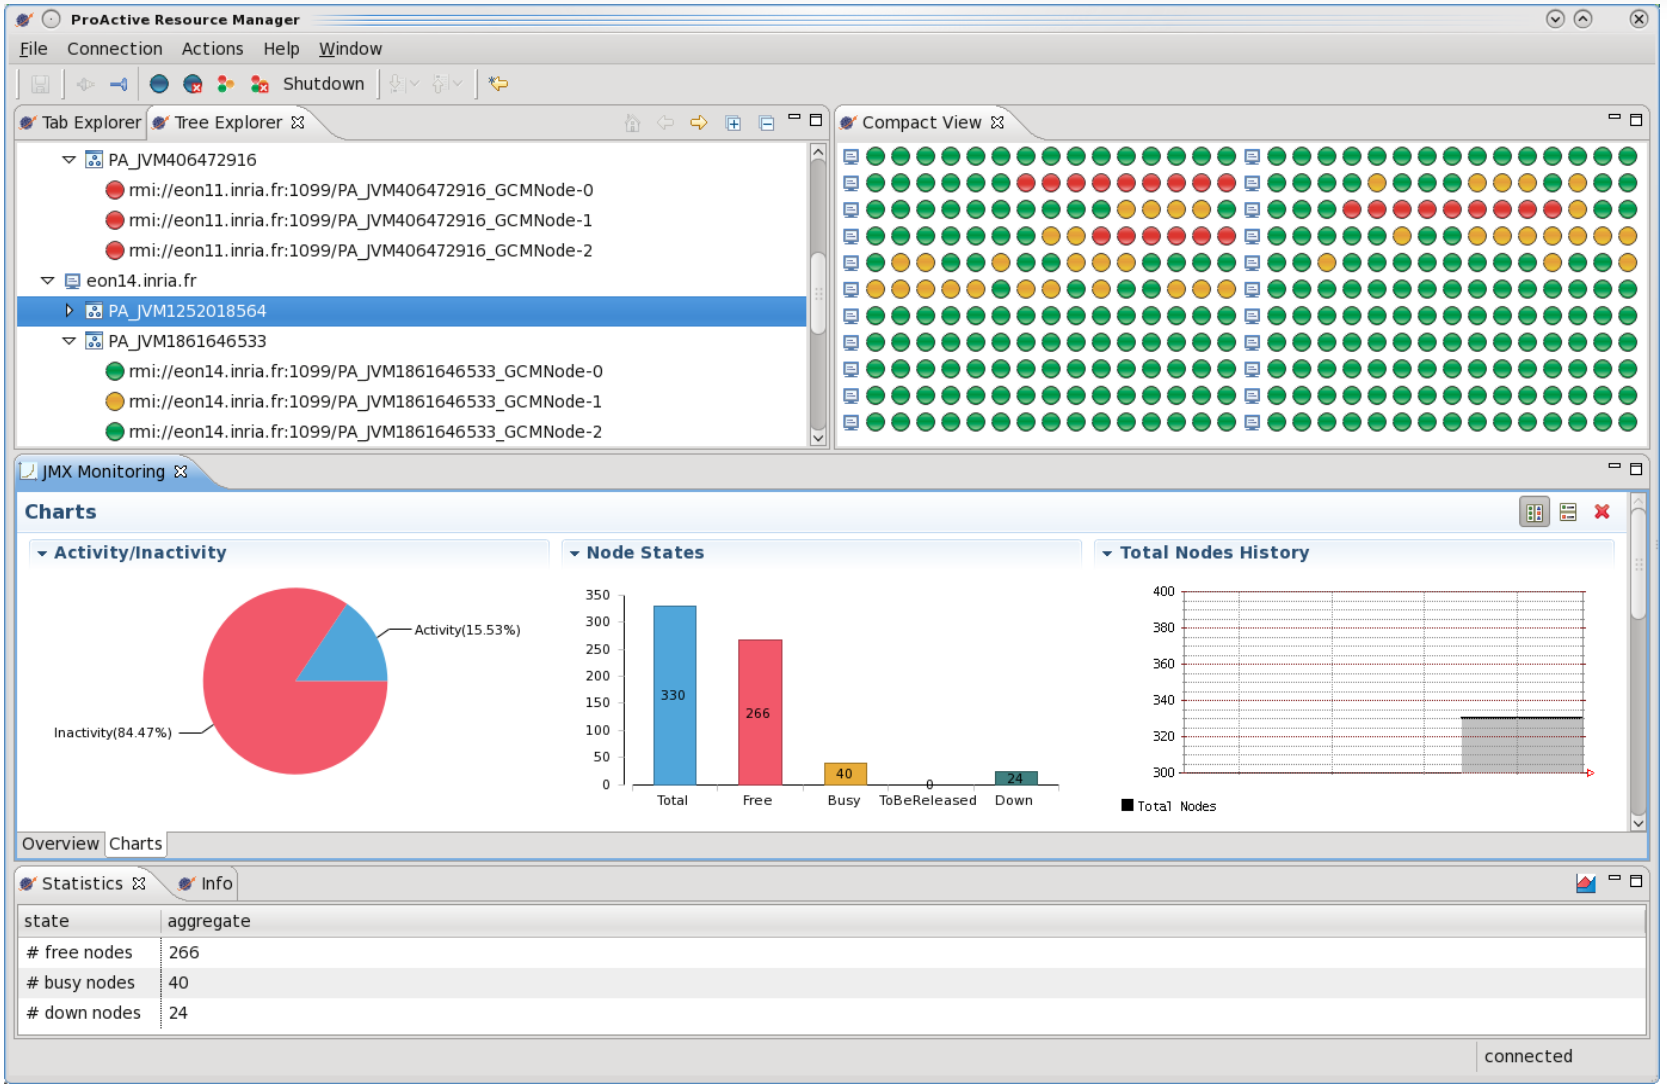

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

Topology with the ProActive Resourcing

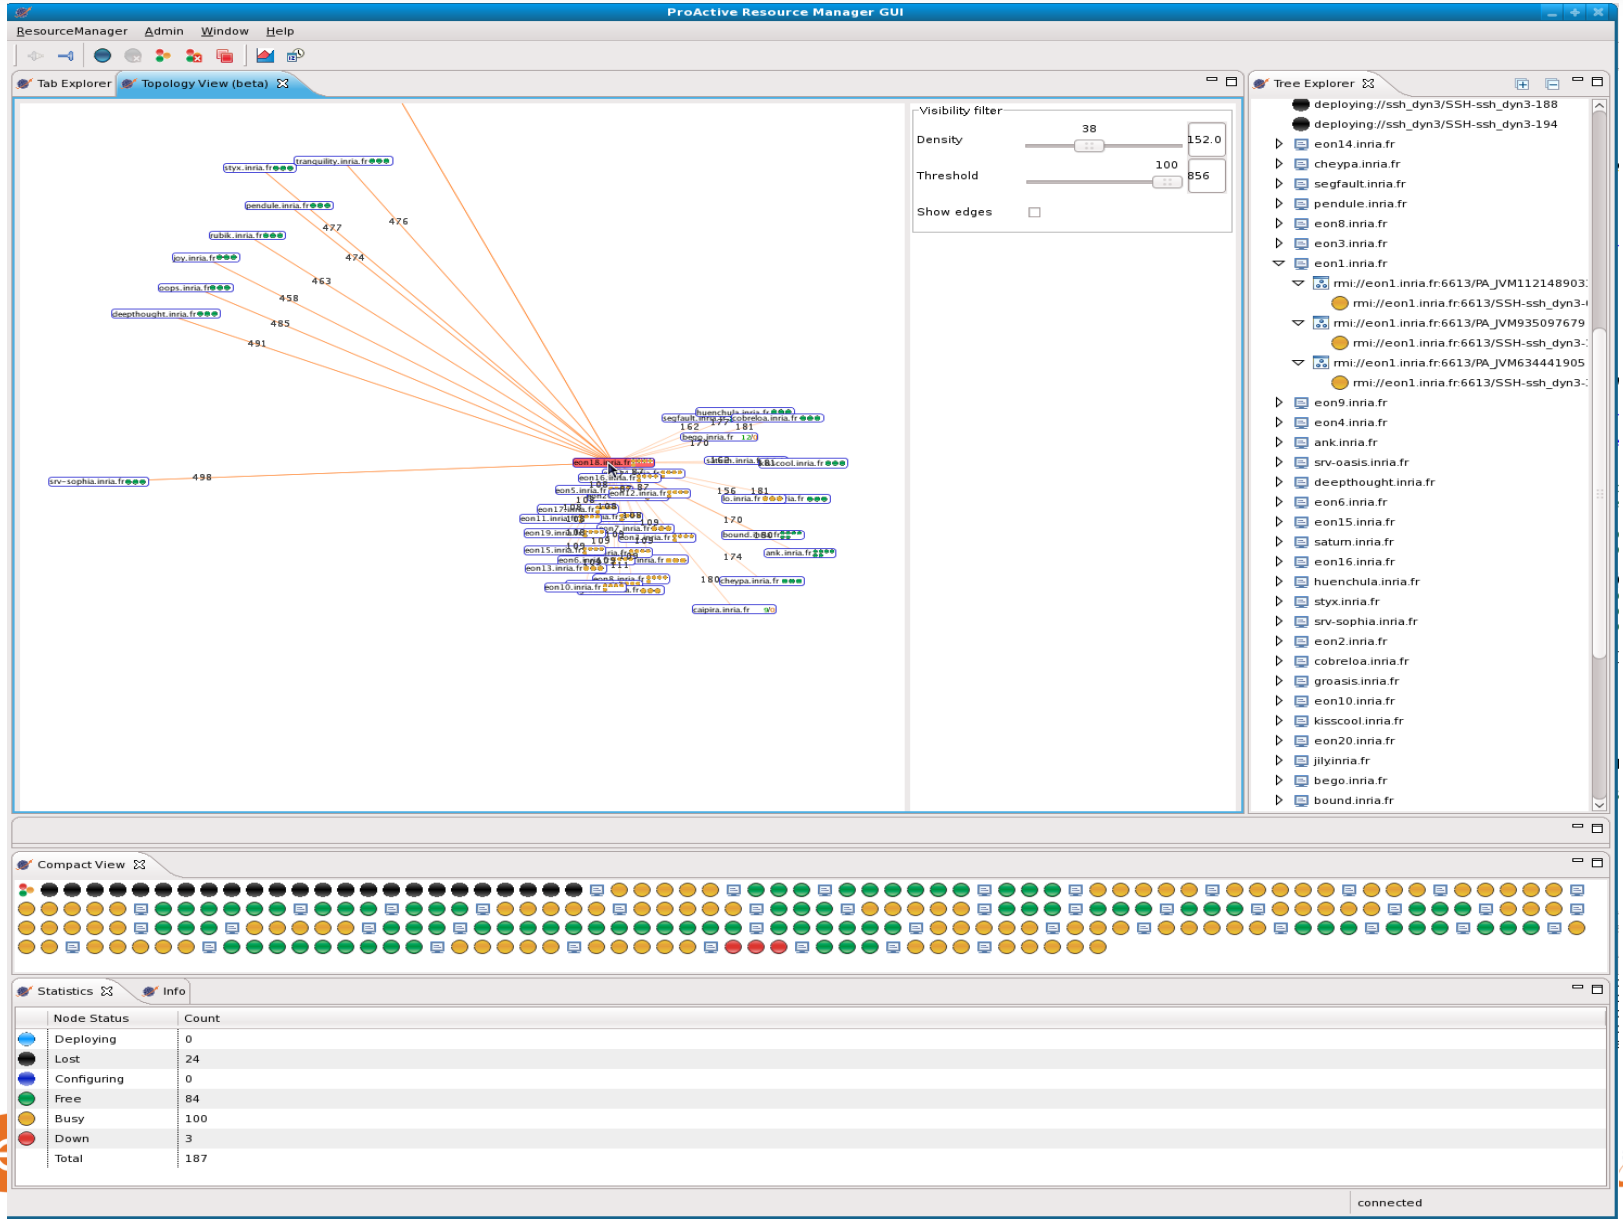

Private, Public & Hybrid Clouds

Enterprise
Private
Hybrid
Cloud

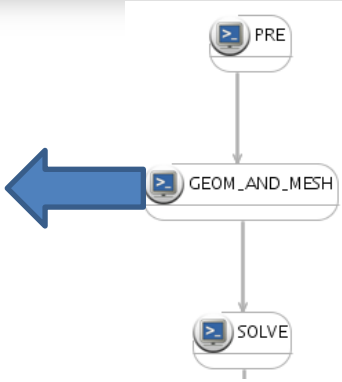

Application
and
Workflow
Acceleration

Static Policy

LSF

Timing Policy
12/24

Desktops

Dynamic Workload Policy

On Burst

Dedicated resources

On premise

Desktops

External Data Center, Amazon EC2, Azure

Activeeon
SCALE BEYOND LIMITS

ProActive
Parallel Suite

ProActive: Security for your Enterprise Cloud platform

- User Authentication (LDAP or Files)
- Authentication and encryptions of network communications using SSH, SSL and PKI
- Secure communication routing through Firewall and NAT configuration of LAN network

Use Cases and Demonstration on a Production Platform

ProActive Parallel Suite

The ProActive PACA Grid Platform (4)

Total:

- 1 368 Cores
- 480 CUDA Cores
- 30TB Storage

Publically Available Today for Production

Activeeon
SCALE BEYOND LIMITS

Use Case 1:
CPU + GPU Workflow

ProActive Parallel Suite

Workflow ProActive for CPU and GPU

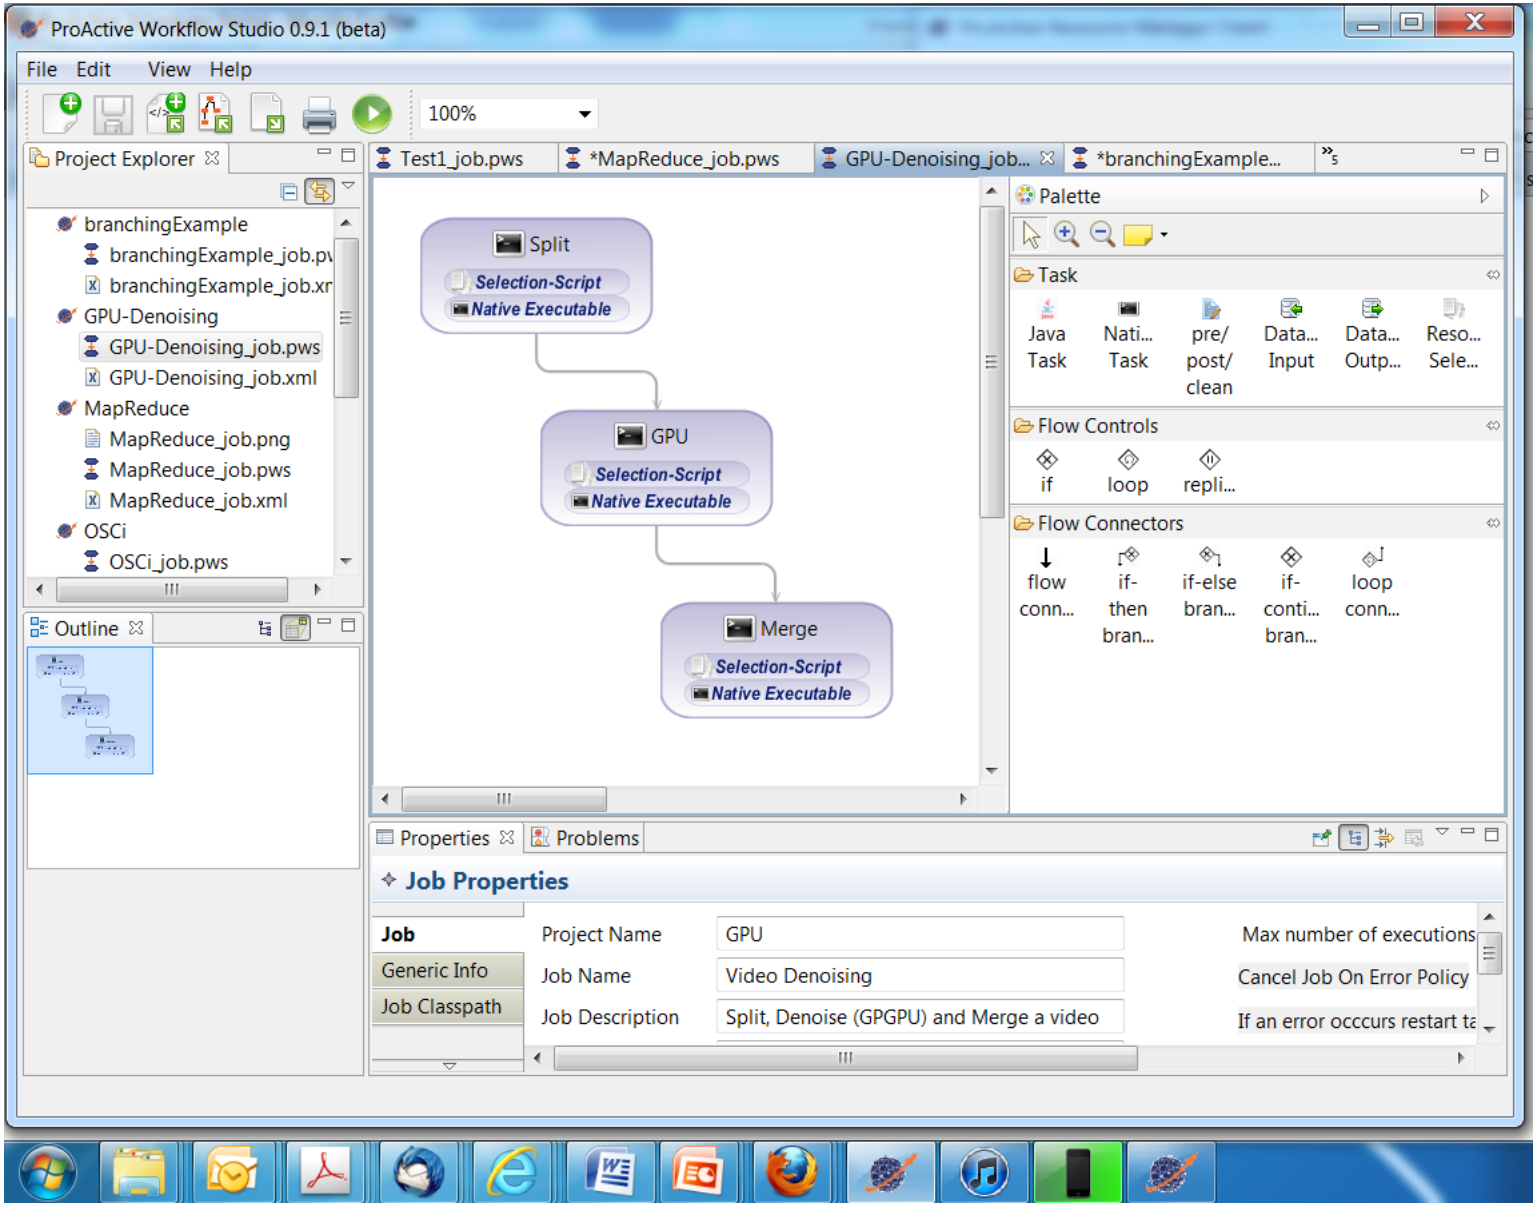

ProActive Parallel Suite

CPU + GPU Workflow

Live Demo

ProActive Parallel Suite

58

CPU + GPU ProActive Workflows

- Resource selection for each Task of a ProActive Workflow
- Selection of Host with GPU capacity
- Data Transfer to the GPU Host
- Configuration of GPU Capacity at the level of Admin (Number of GPU Nodes, size)
- Freedom to request one or several GPU capacities for one GPU program
- Global Scheduling (Multi-Tenant, Multi-Application) of GPU Tasks

ProActive Parallel Suite

Use Case 2: OMD2

Distributed Multi-Disciplinary Optimizations with Remote Visualization

ProActive Parallel Suite

OMD2

Open Source Interfaces

For

Distribued Multi-Disciplinary Optimisations

OMD2 OPTIMISATION MULTI DISCIPLINAIRE DISTRIBUÉE

ProActive Parallel Suite

OMD2 :

Open Source Interfaces
For
Distribued Multi-Disciplinaires
Optimisations

OMD2 OPTIMISATION MULTI DISCIPLINAIRE DISTRIBUÉE

RENAULT

Laboratoire Roberval
Unité de recherche en mécanique

CENTRALE PARIS

Ecole Nationale Supérieure des Mines
SAINT-ETIENNE

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
INRIA

CD-adapco

digiteo
Research in information sciences and technologies

SIREHNA
a DCNS company

ACTIVEEON
SCALE BEYOND LIMITS

ProActive
Parallel Suite

Distributed Workflow

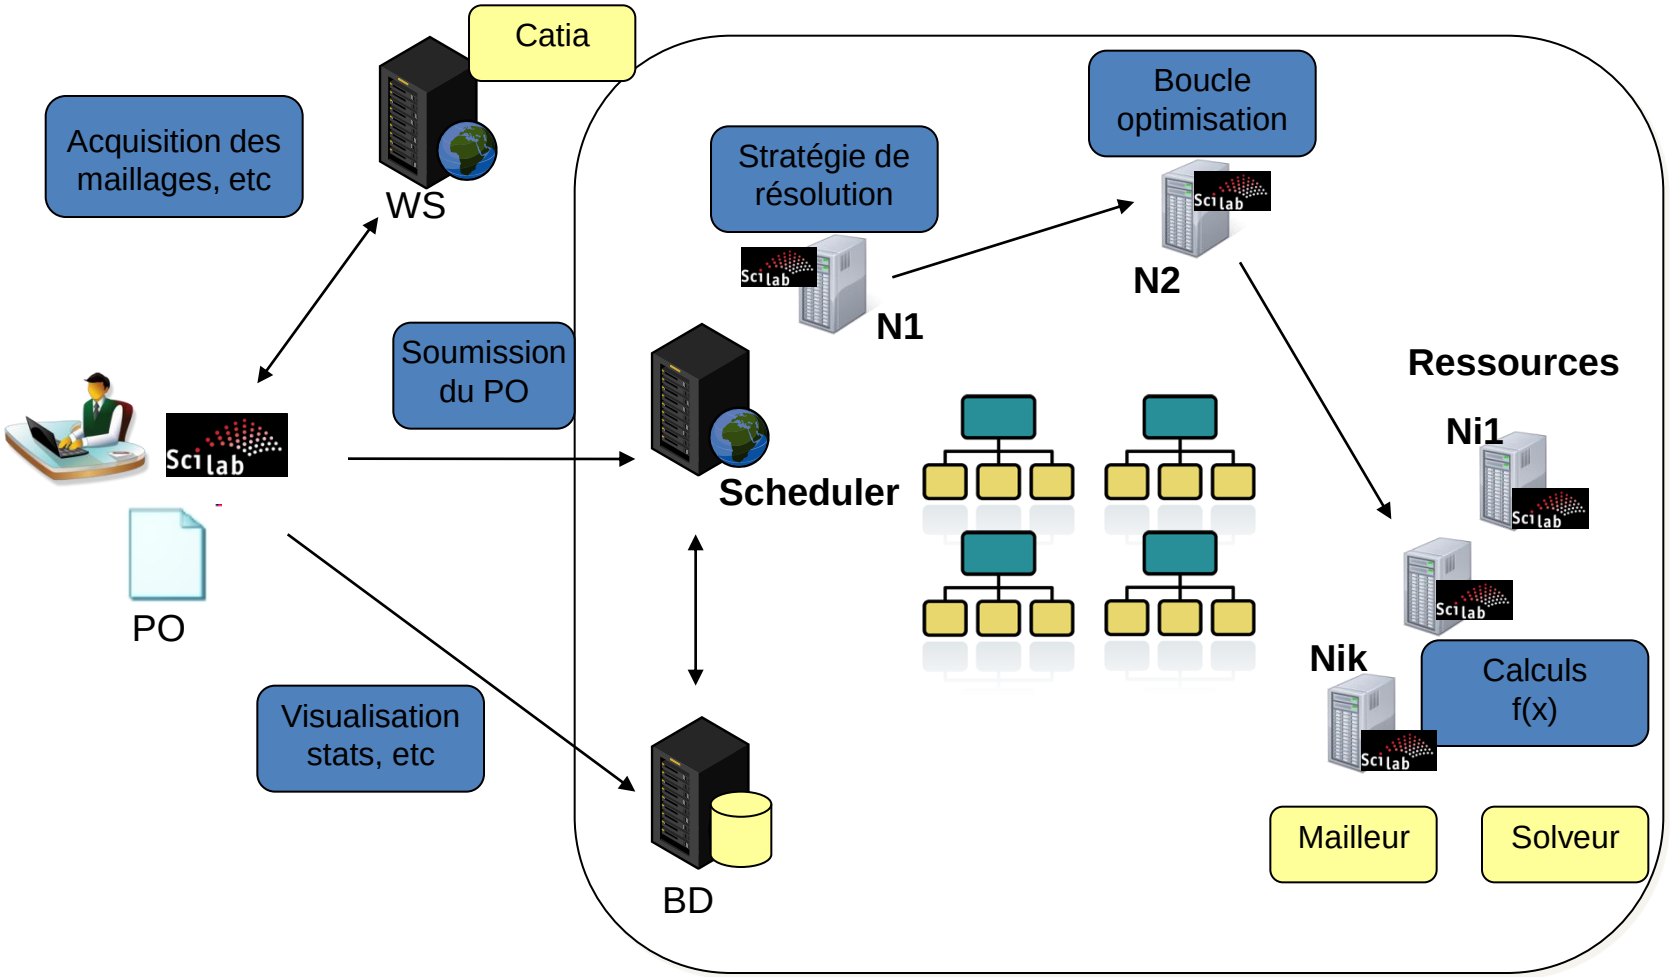

OMD2 OPTIMISATION MULTI DISCIPLINAIRE DISTRIBUÉE

Coupling Mechanics, Aerodynamics ...

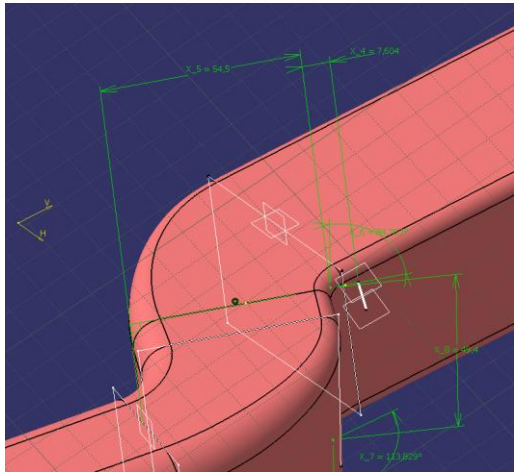

10min CPU

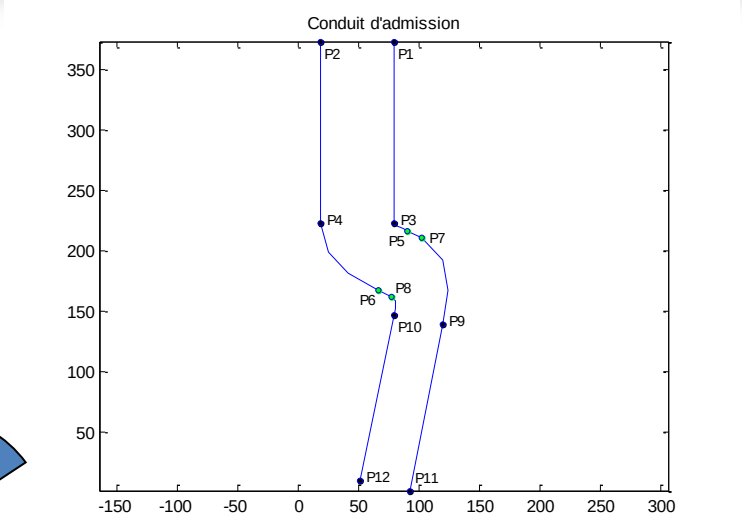

<1min CPU

3D Air Conditionning

2D Air Conditionning

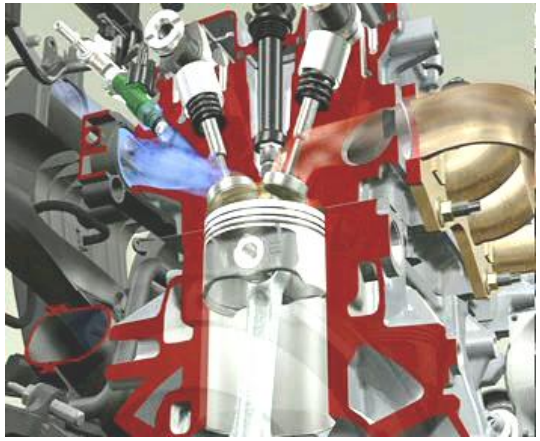

100h CPU

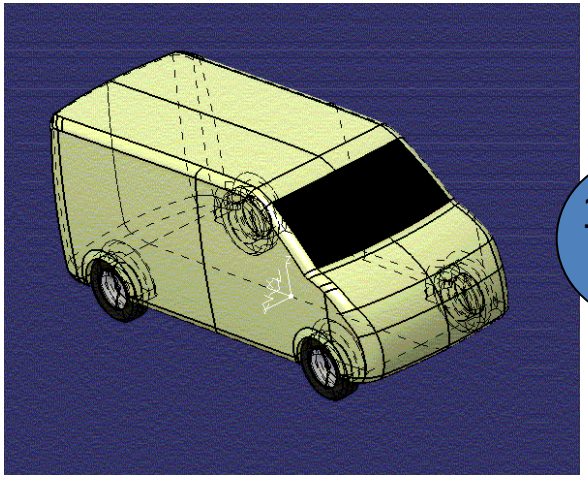

1000h CPU

Cylinder Head

External Aerodynamic

ProActive Parallel Suite

ProActive OMD2 Demo

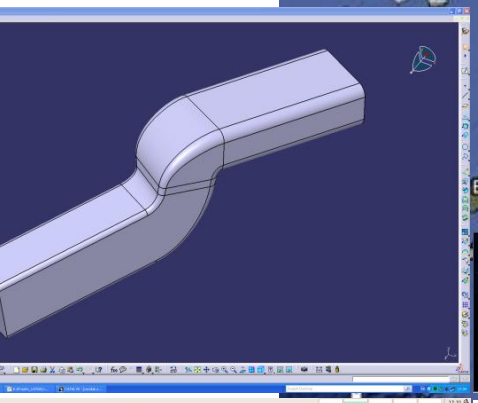

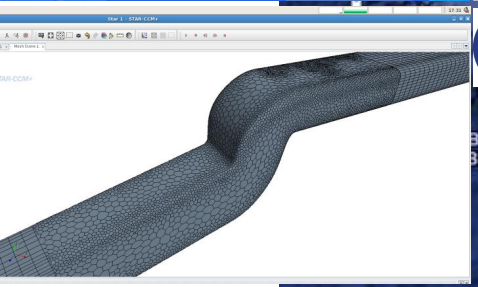

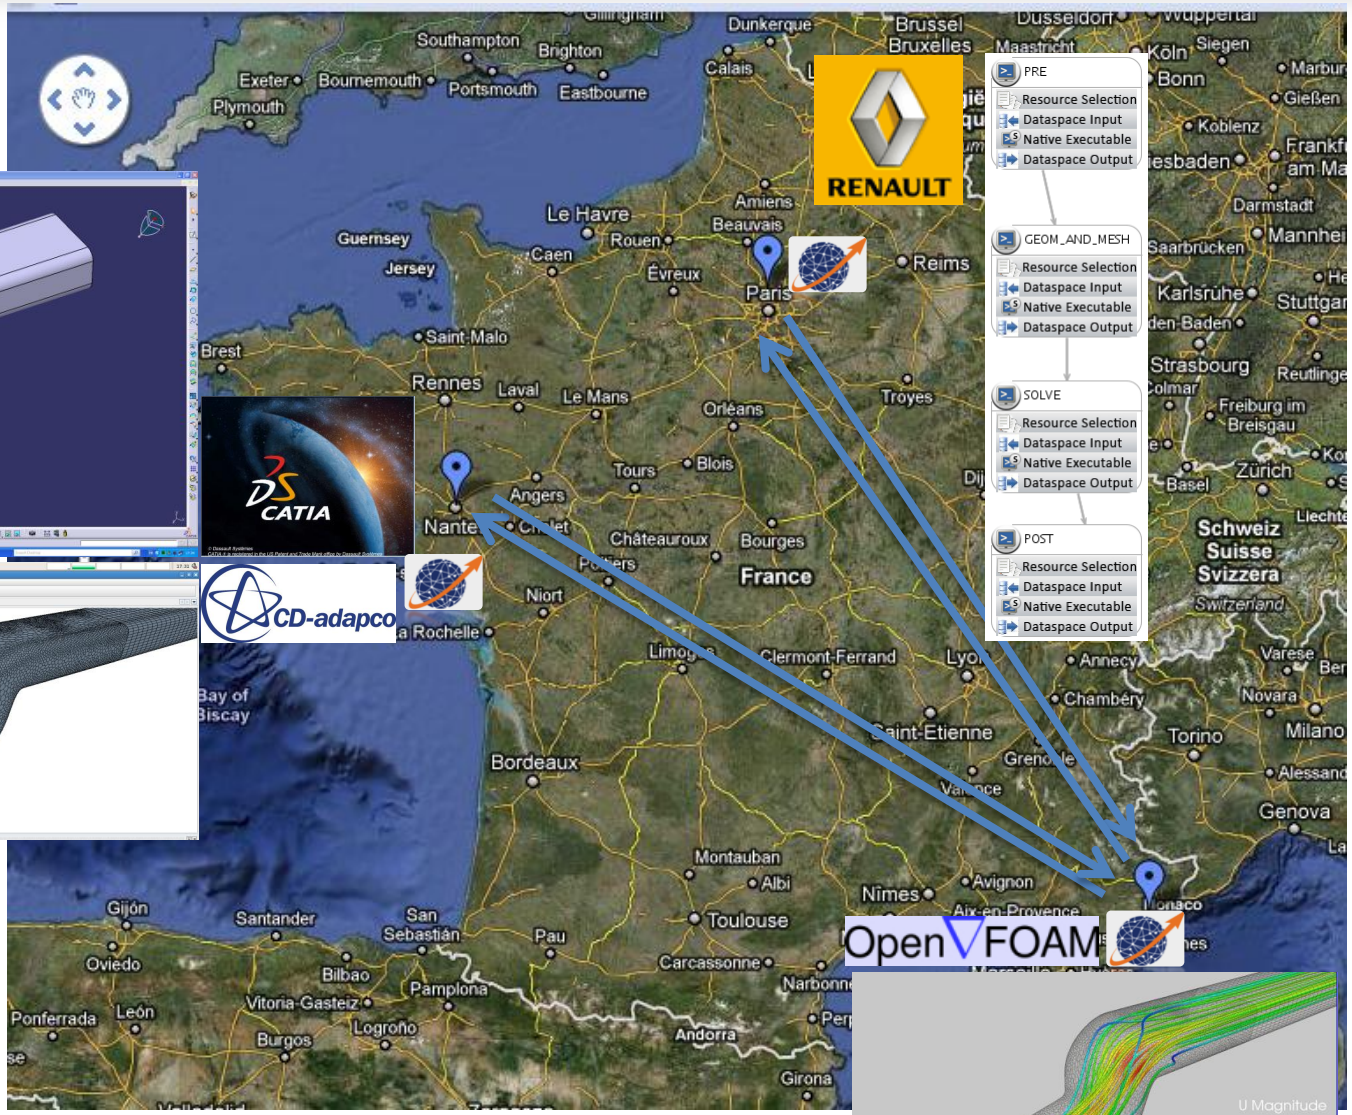

1000 Cores Production Cloud Portal

Demonstration

Engineering Optimizations: Renault UC

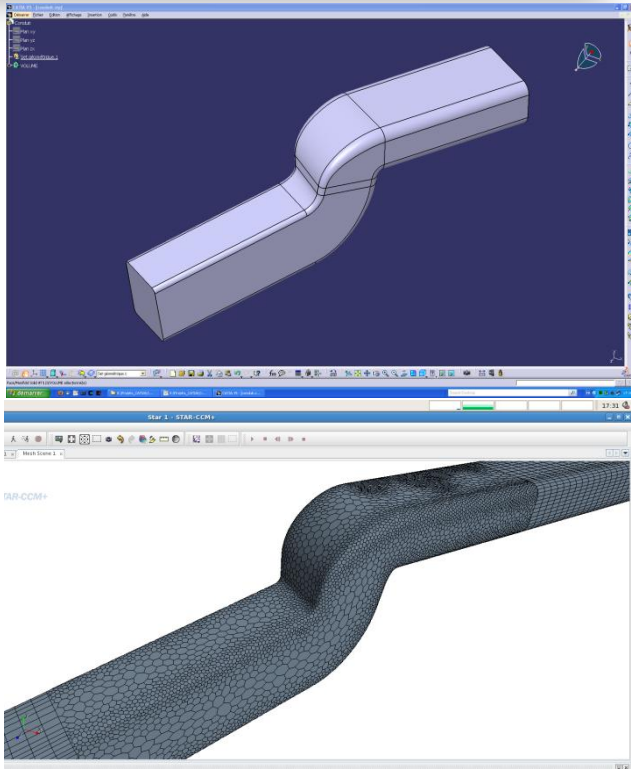

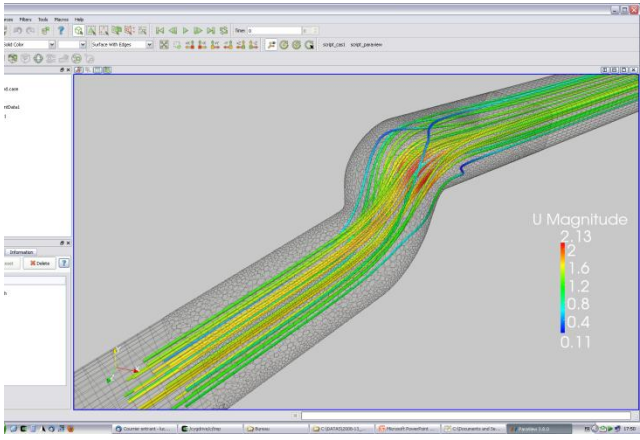

Use Case 3: Map Reduce

ProActive Parallel Suite

ProActive MapReduce (CO, SP2, Task 2.1)

- Same APIs as Hadoop
  (Easy switch from Hadoop to ProActive)
- Does not requires an HDFS File System
- Runs on general purpose, Multi-tenant, Multi-Applications Grids and Clouds
- Available as PaaS in Java

ProActive Parallel Suite

Workflow ProActive MapReduce

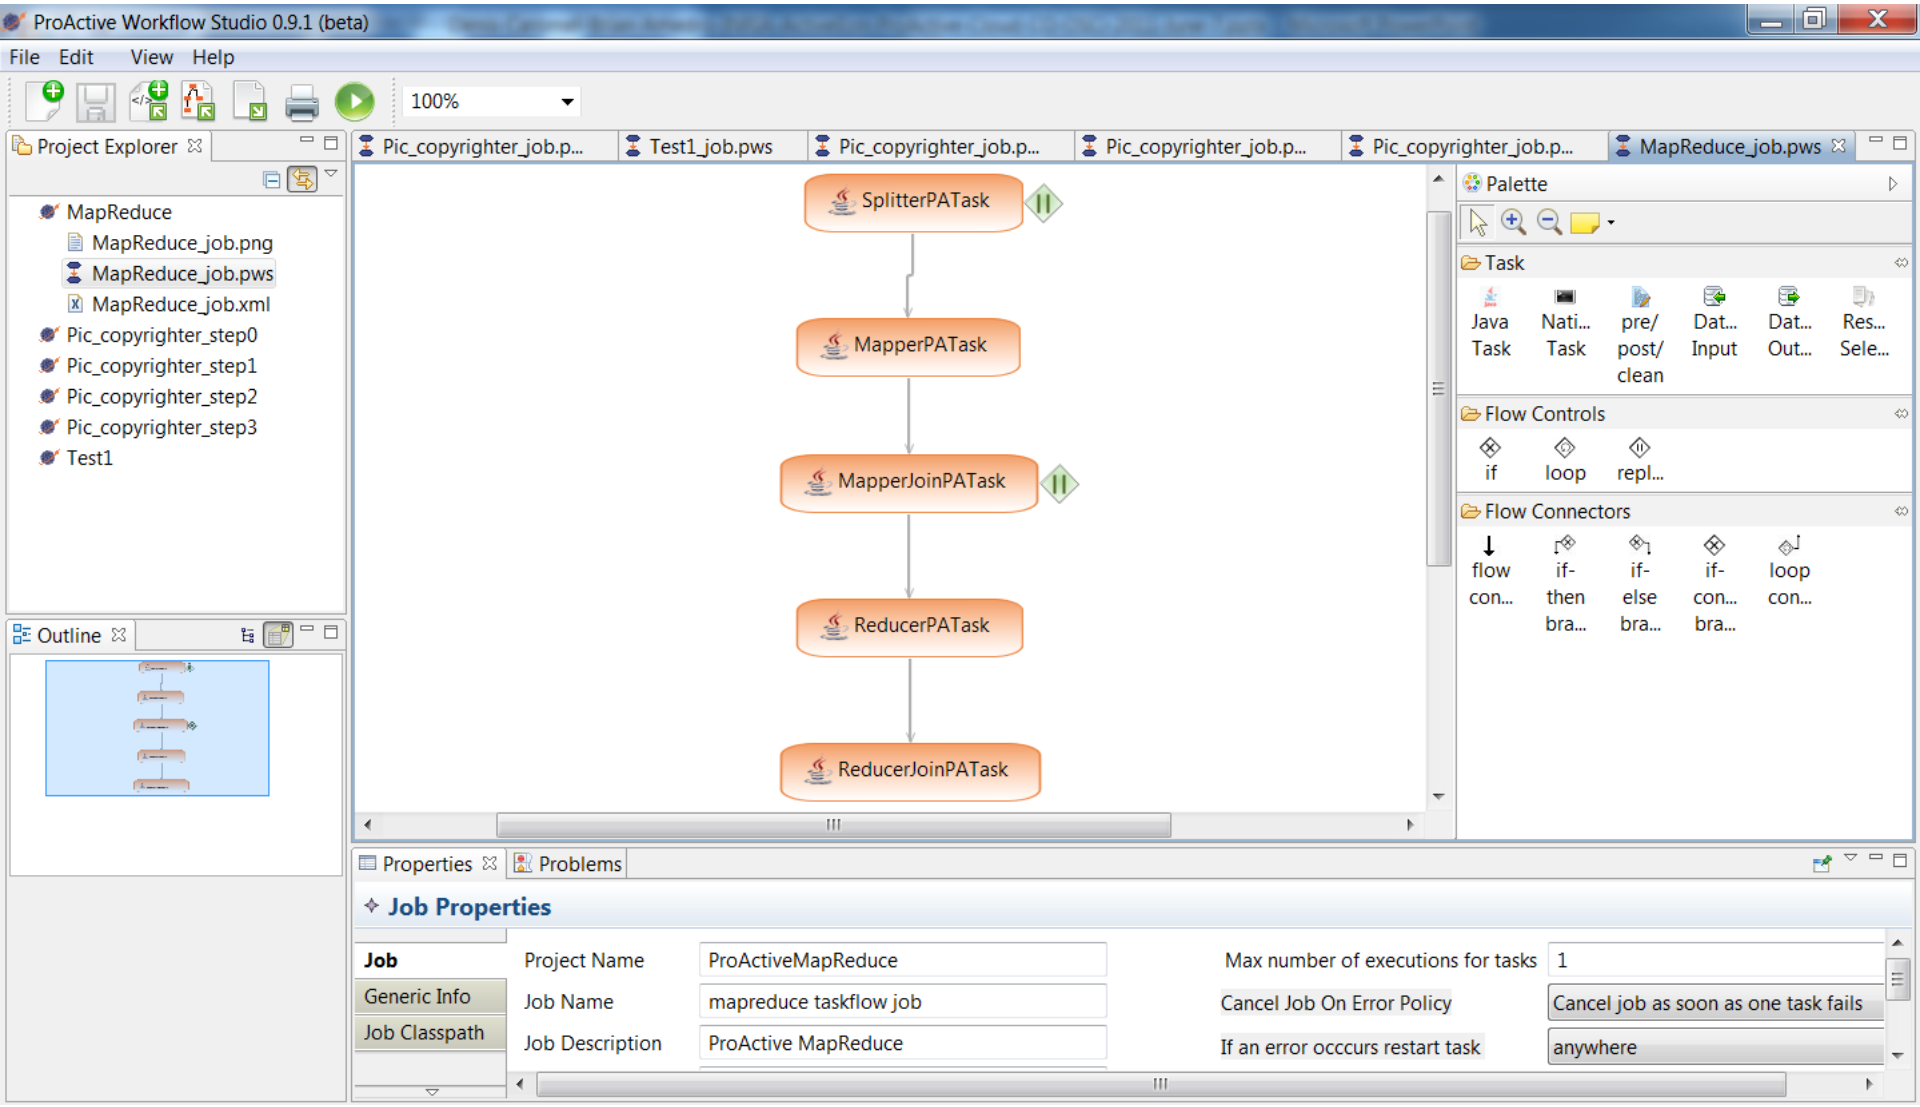

ProActive Parallel Suite

Map Reduce Demo

ProActive Parallel Suite

ProActive MapReduce vs. Hadoop+HDFS

File Size	Sequential	Hadoop	PA MapReduce	Speedup
0.7 GB	5m 04s	1m 17s	1m 05s	4.6
4.3 GB	25m 31s	2m 30s	2m 20s	10.9
7.3 GB	46m 00s	3m 31s	3m 30s	13.1
20 GB	2h 07m 00s	8m 30s	7m 09s	17.8
50 GB	5h 19m 00s	21m 05s	25m 11s	12.7
100 GB	10h 38m 00s	43m 23s	58m 42s	10.9

- Data available in a NAS (General purpose storage)
- Transfer to HDFS for Hadoop
- Used directly without copy for ProActive
- Use Case of Map/Reduce on fresh data
- Different ProActive Map/Reduce configuration for recurrent MR on in place Data (e.g. ProActive HDFS interface)

ProActive Parallel Suite

Use Case 4: BioTechs

ProActive Parallel Suite

IPMC Use Case and Collaboration

ProActive Parallel Suite

PBS

Cluster

Desktops

amazon web services™

EC2

Clouds

Nodes can be dynamically added!

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

OW2 Consortium

Benchmarks

- The distributed version with ProActive of Mapreads has been tested on the INRIA cluster with two settings: the Reads file is split in either 30 or 10 slices
- Use Case: Matching 31 millions Sequences with the Human Genome (M=2, L=25)

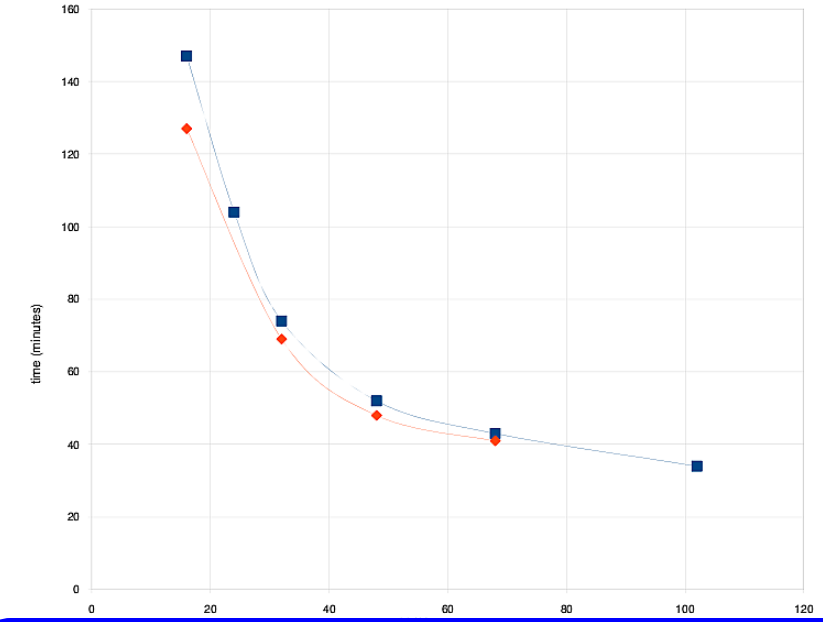

4 Time FASTER from 20 to 100
Speed Up of 80 / Th.
Sequential : 50 h ➔ 35 mn

EC2 only test: nearly the same performances as the local SOLiD cluster (+10%)

For only $3,2/hour, EC2 has nearly the same perf. as the local SOLiD cluster (16 cores, for 2H30)

ProActive Parallel Suite

Use Case 5: Hydrodynamic with K-Epsilon and FineMarine

Hydrodynamic Optimization: Workflow generated from a GUI

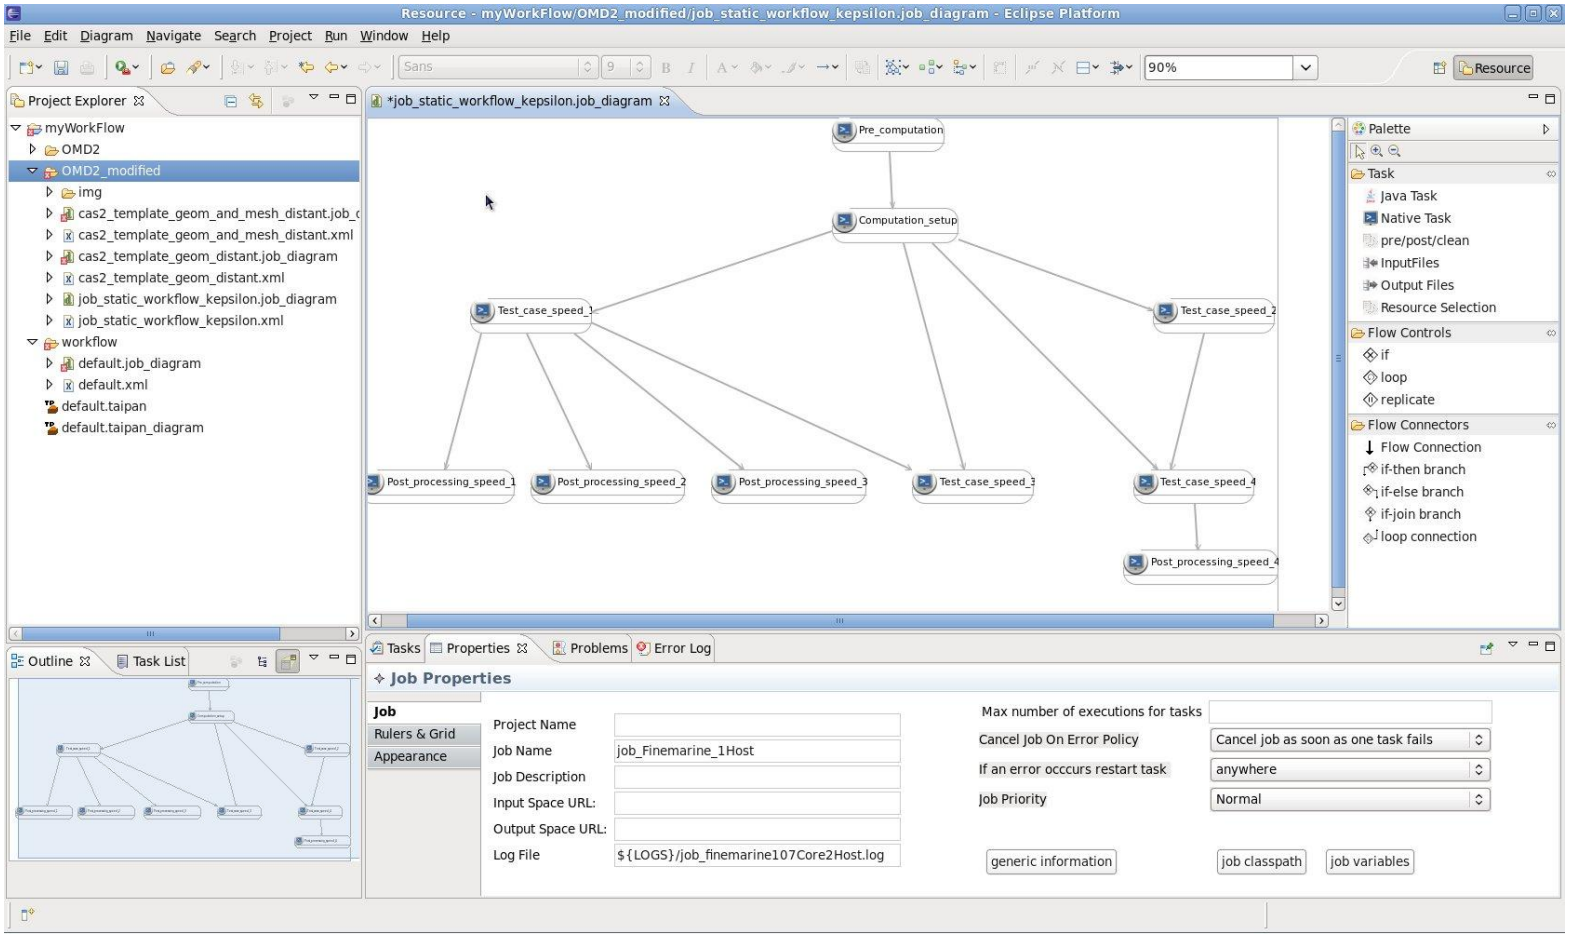

Hydrodynamic Optimization: Execution

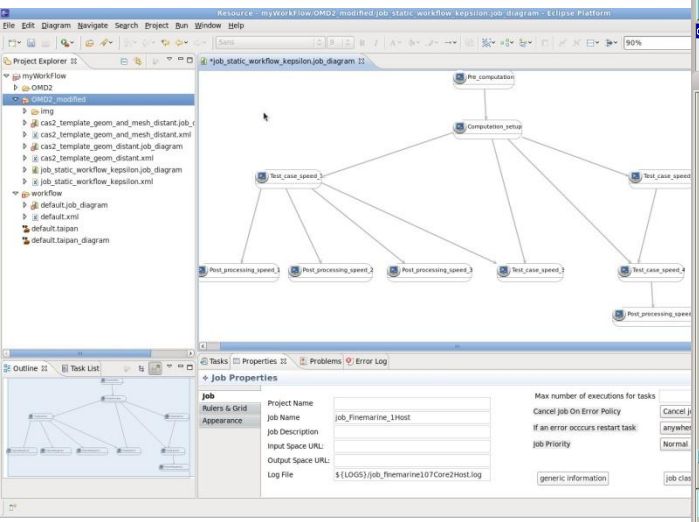

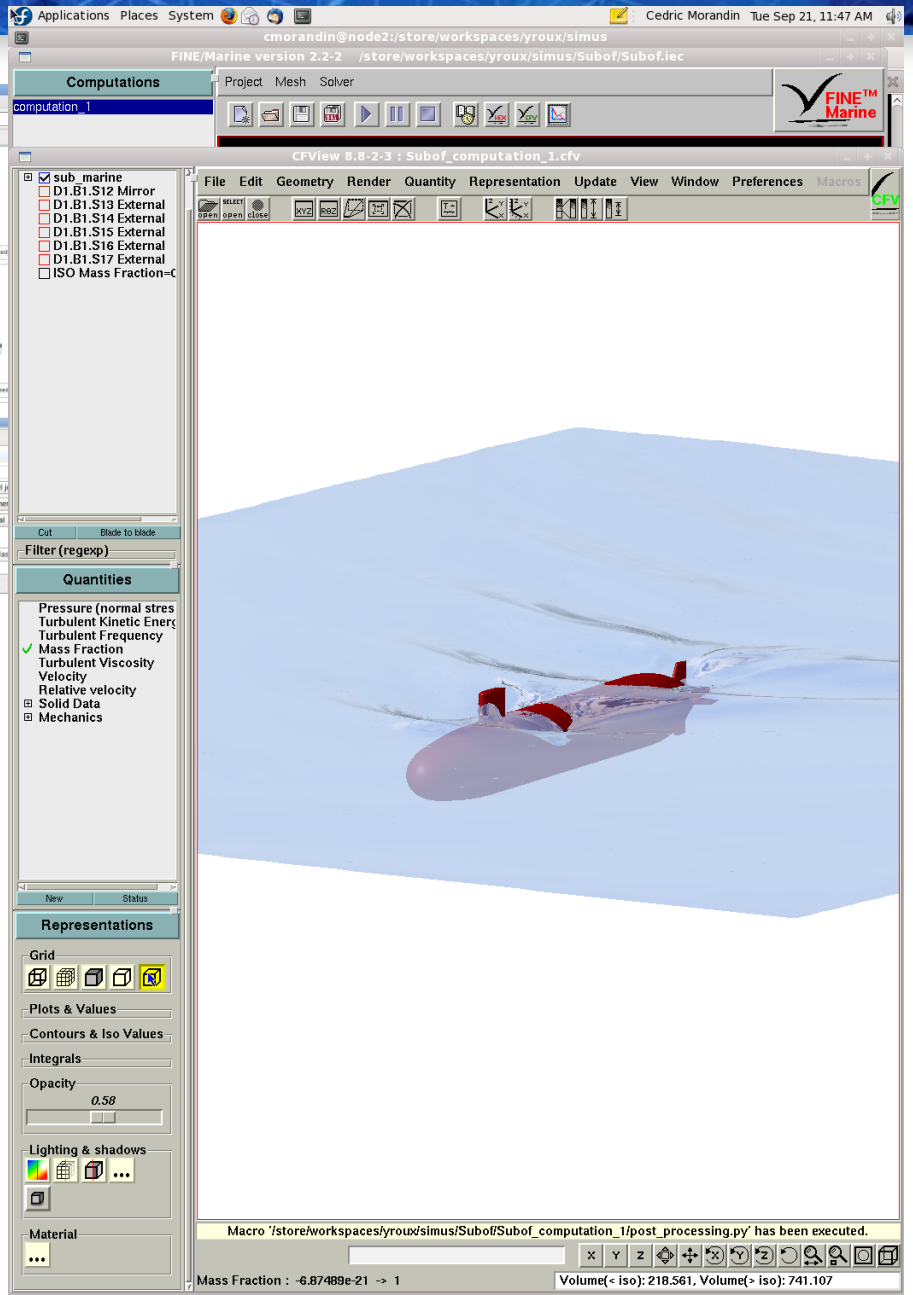

Activeeon
SCALE BEYOND LIMITS

OW2 Consortium

Hydrodynamic:
Remote Steering during execution

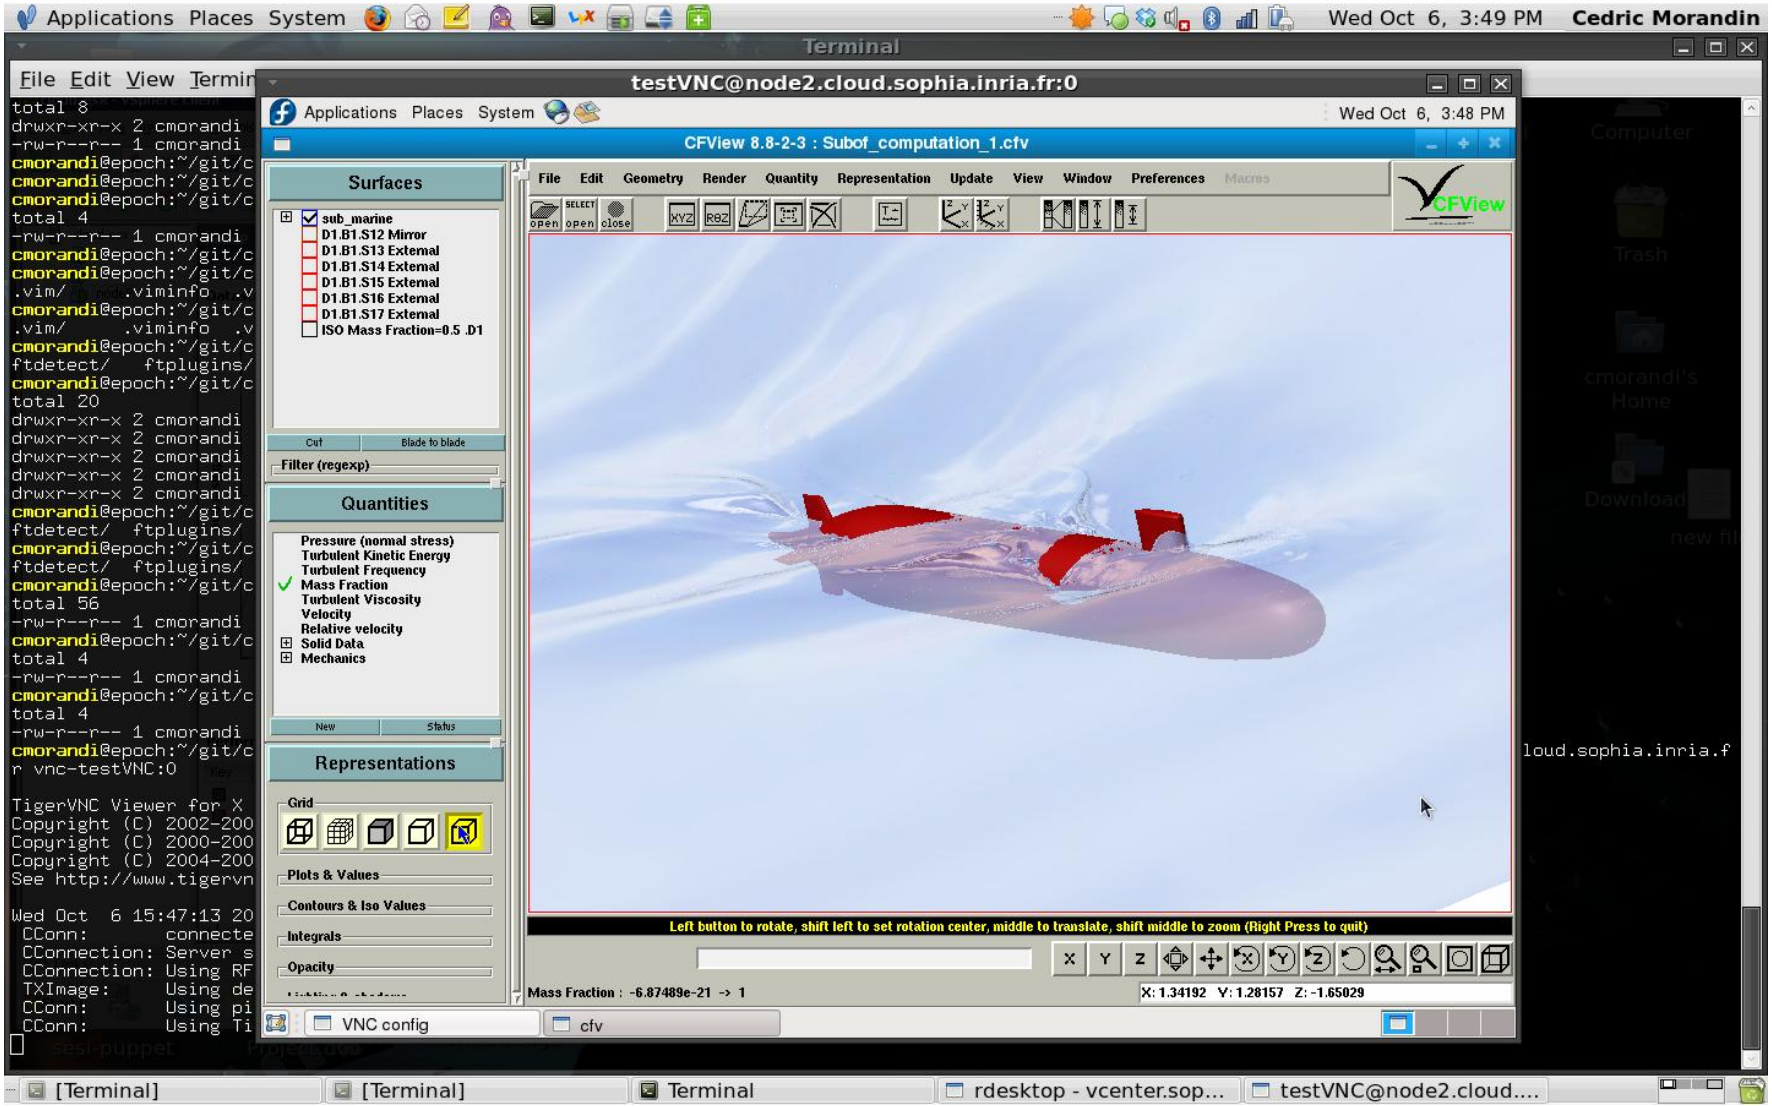

ProActive Parallel Suite

PRICING PARTNERS

UC 6: Acceleration of Financial Valuations

C++ library developed by Pricing Partners
Pricing solution dedicated to highly complex derivatives,
Greek computation

ProActive Parallel Suite

How Does it Work?
Price-it Computing Distribution

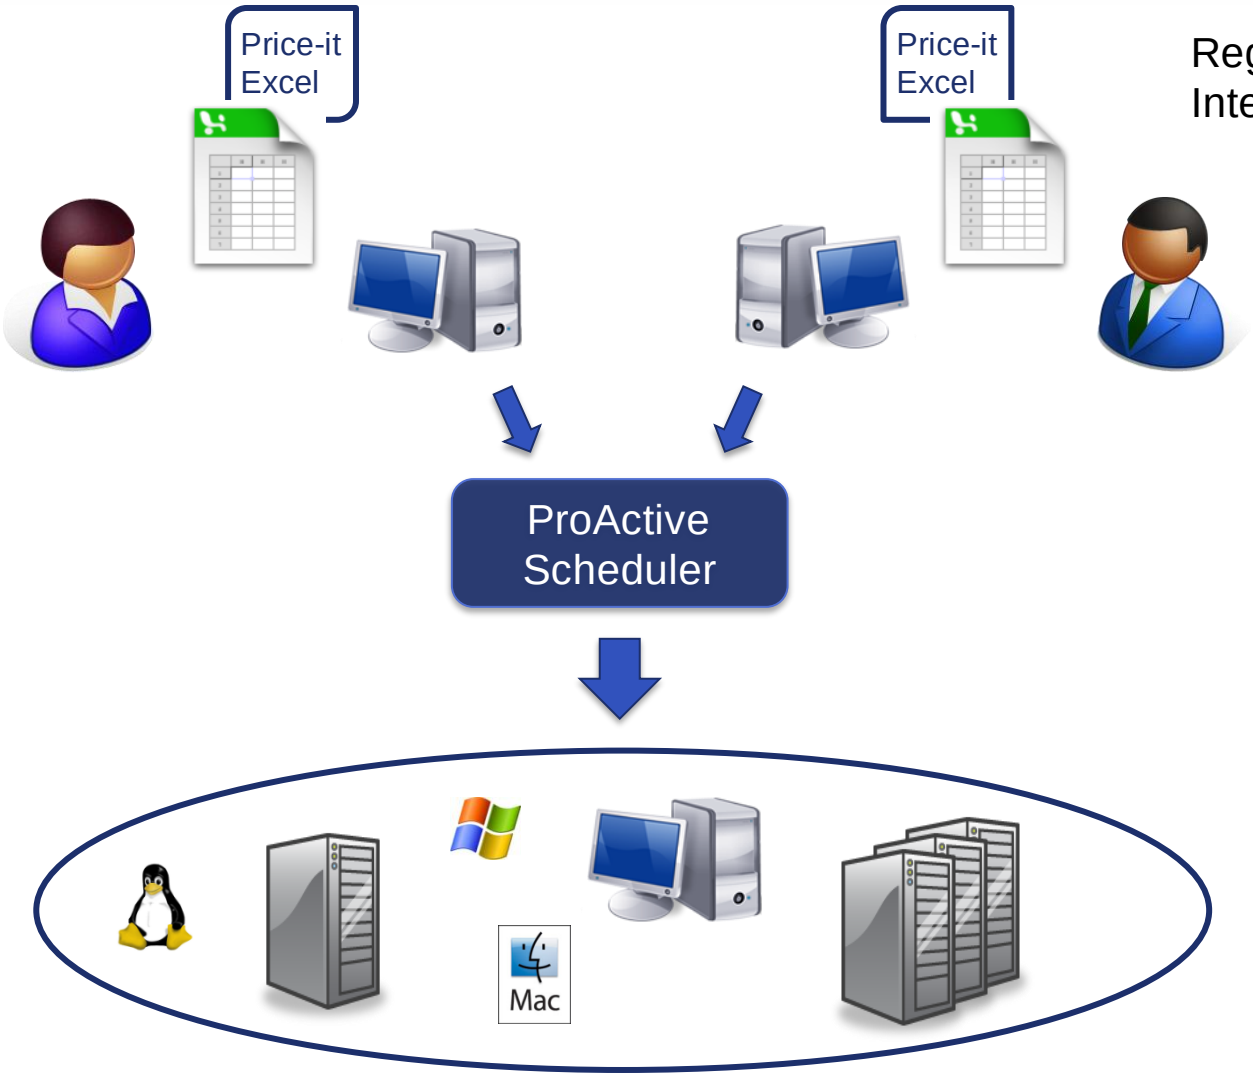

Regular Price-it Excel Interface

Automatic execution via job scheduler

Pool of shared resources

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

Accelerated Price-it Performances

- **Increased Productivity**: Reduces Price-it Execution Time by 6 or more!

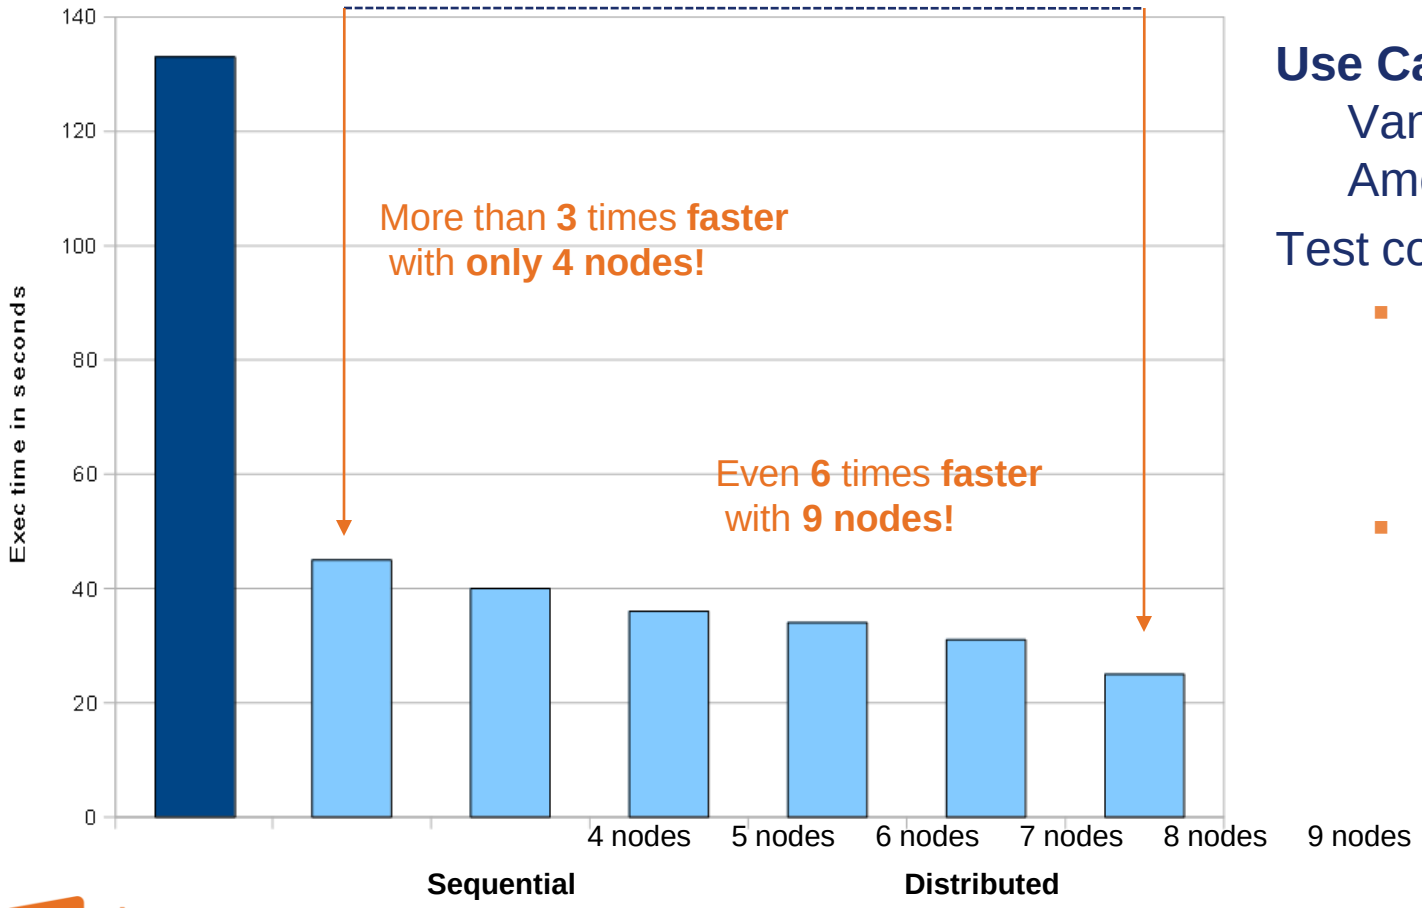

Use Case: Bermuda Vanilla, Model American MC

Test conditions:

- One computation is split in 130 tasks that are distributed
- Each task uses 300ko

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

UC 7: IT

SOA Analysis of Web Server Logs

Parallel Services

- Separation: BPEL – Parallel Serv. – Task Flow
- Standards et Portable
- Flexibility

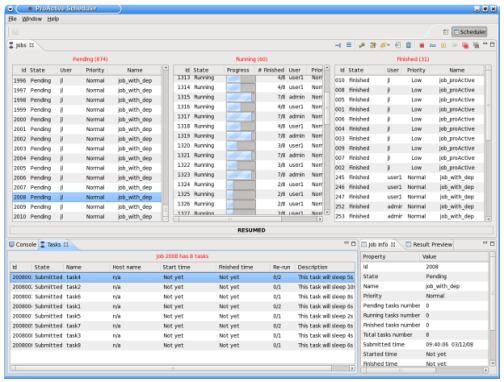

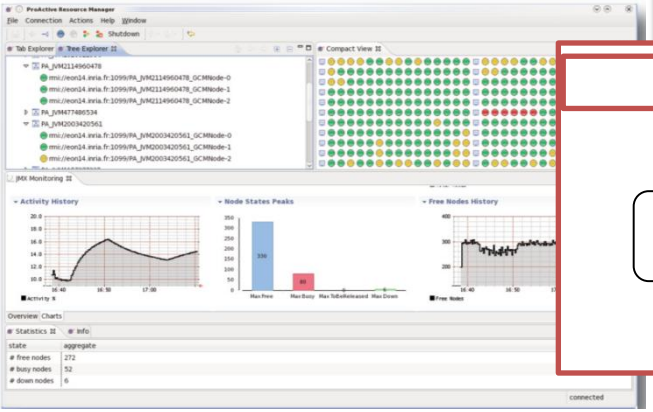

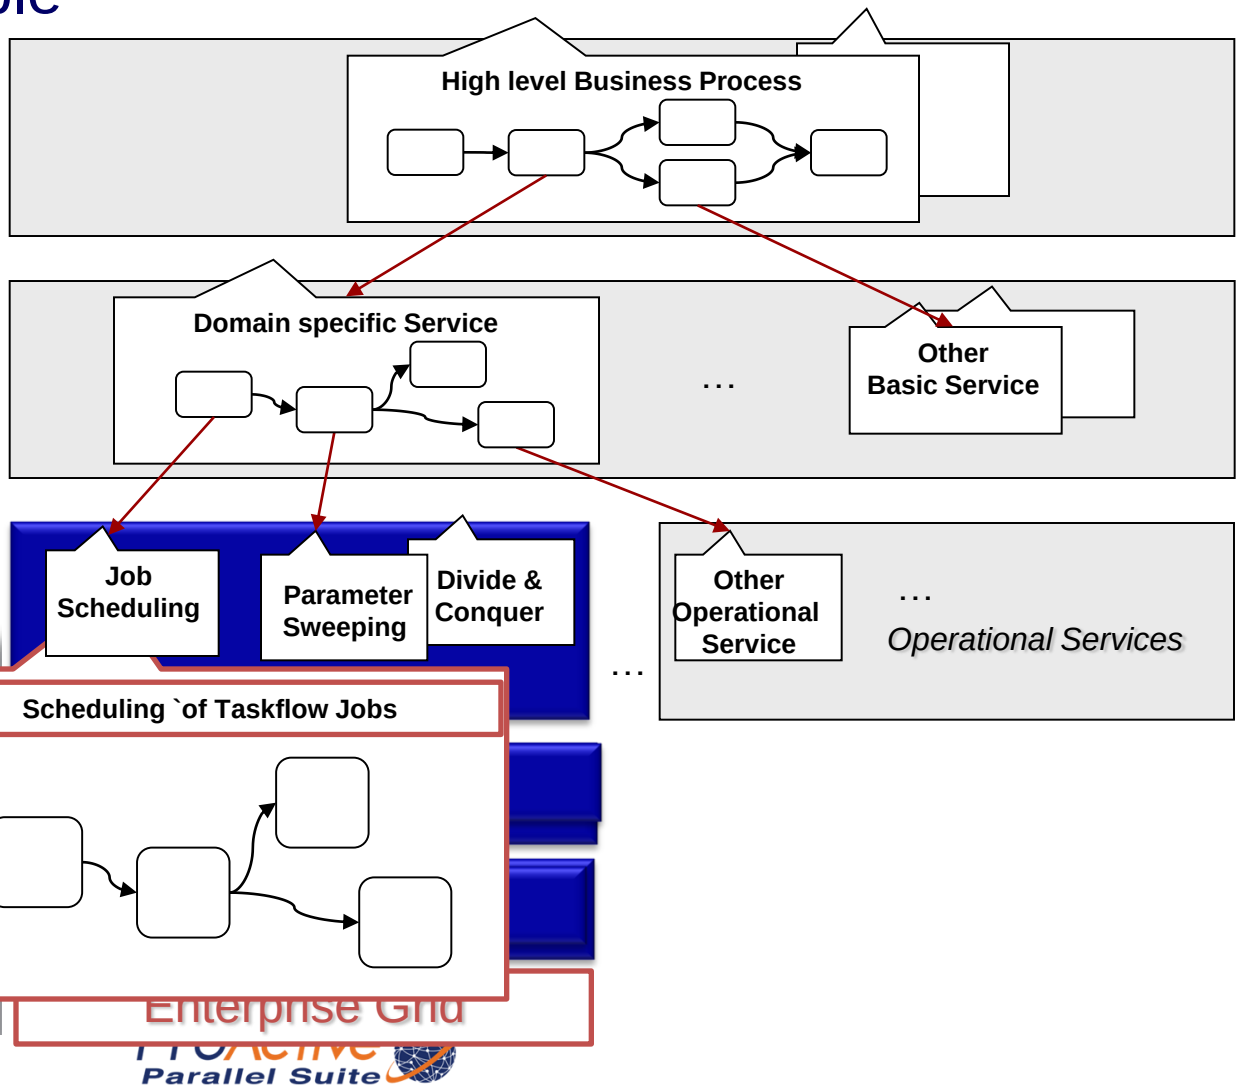

84

AGOS Platform Management

HP- Business Availability Center (HP-BAC)

- Monitoring of entire platform
- Cover all layers in the scope
- Monitoring dashboard and reports

Tasks Scheduler & Resources manager

- Integration with grid
- Indicator on running jobs
- Hypervisor & VM management

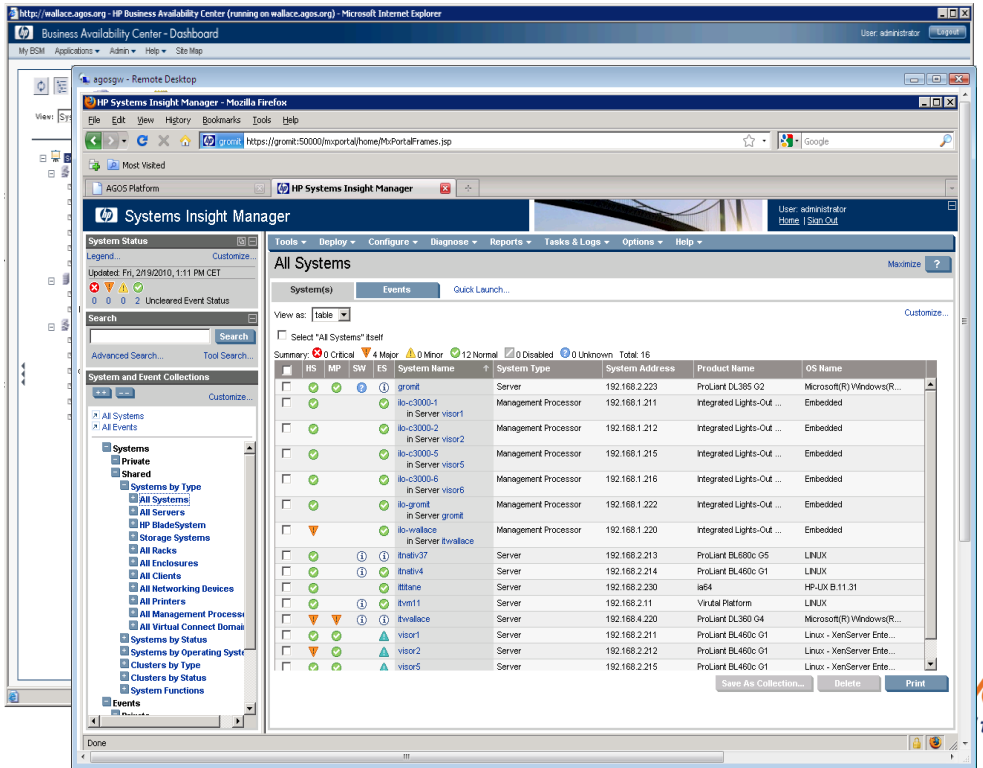

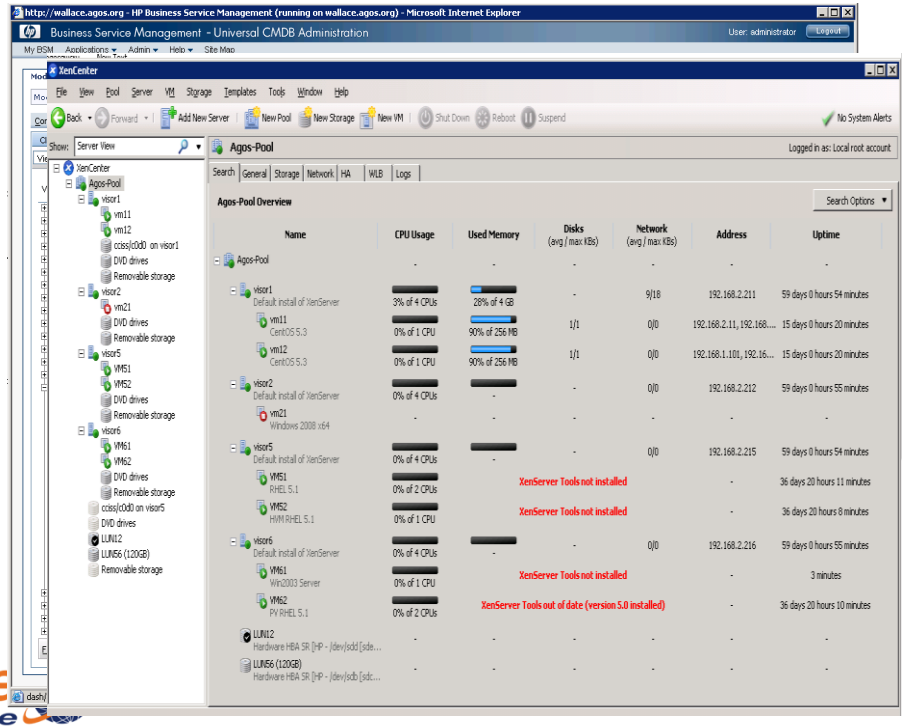

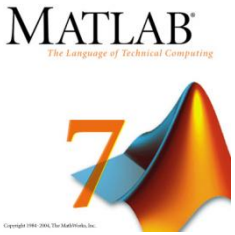

Integration with Scilab and Matlab

ProActive Parallel Suite

Integration with Scilab and Matlab

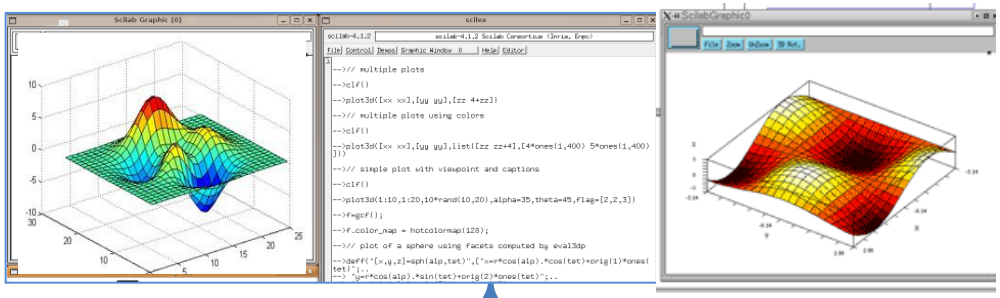

Scilab

MATLAB
The Language of Technical Computing

Static Policy

LSF

Timing Policy 12/24

Desktops

Dynamic Workload Policy

EC2

Dedicated resources

Desktops

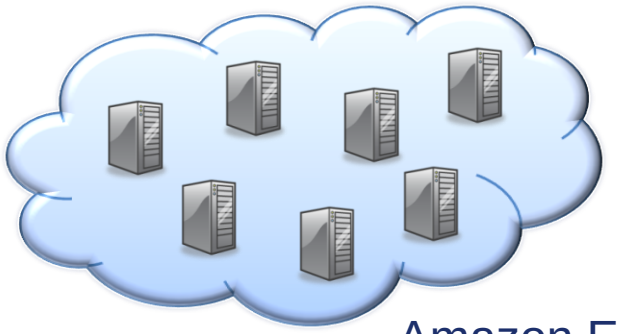

Amazon EC2

ProActive Parallel Suite

Interface ProActive ⇔ Scilab

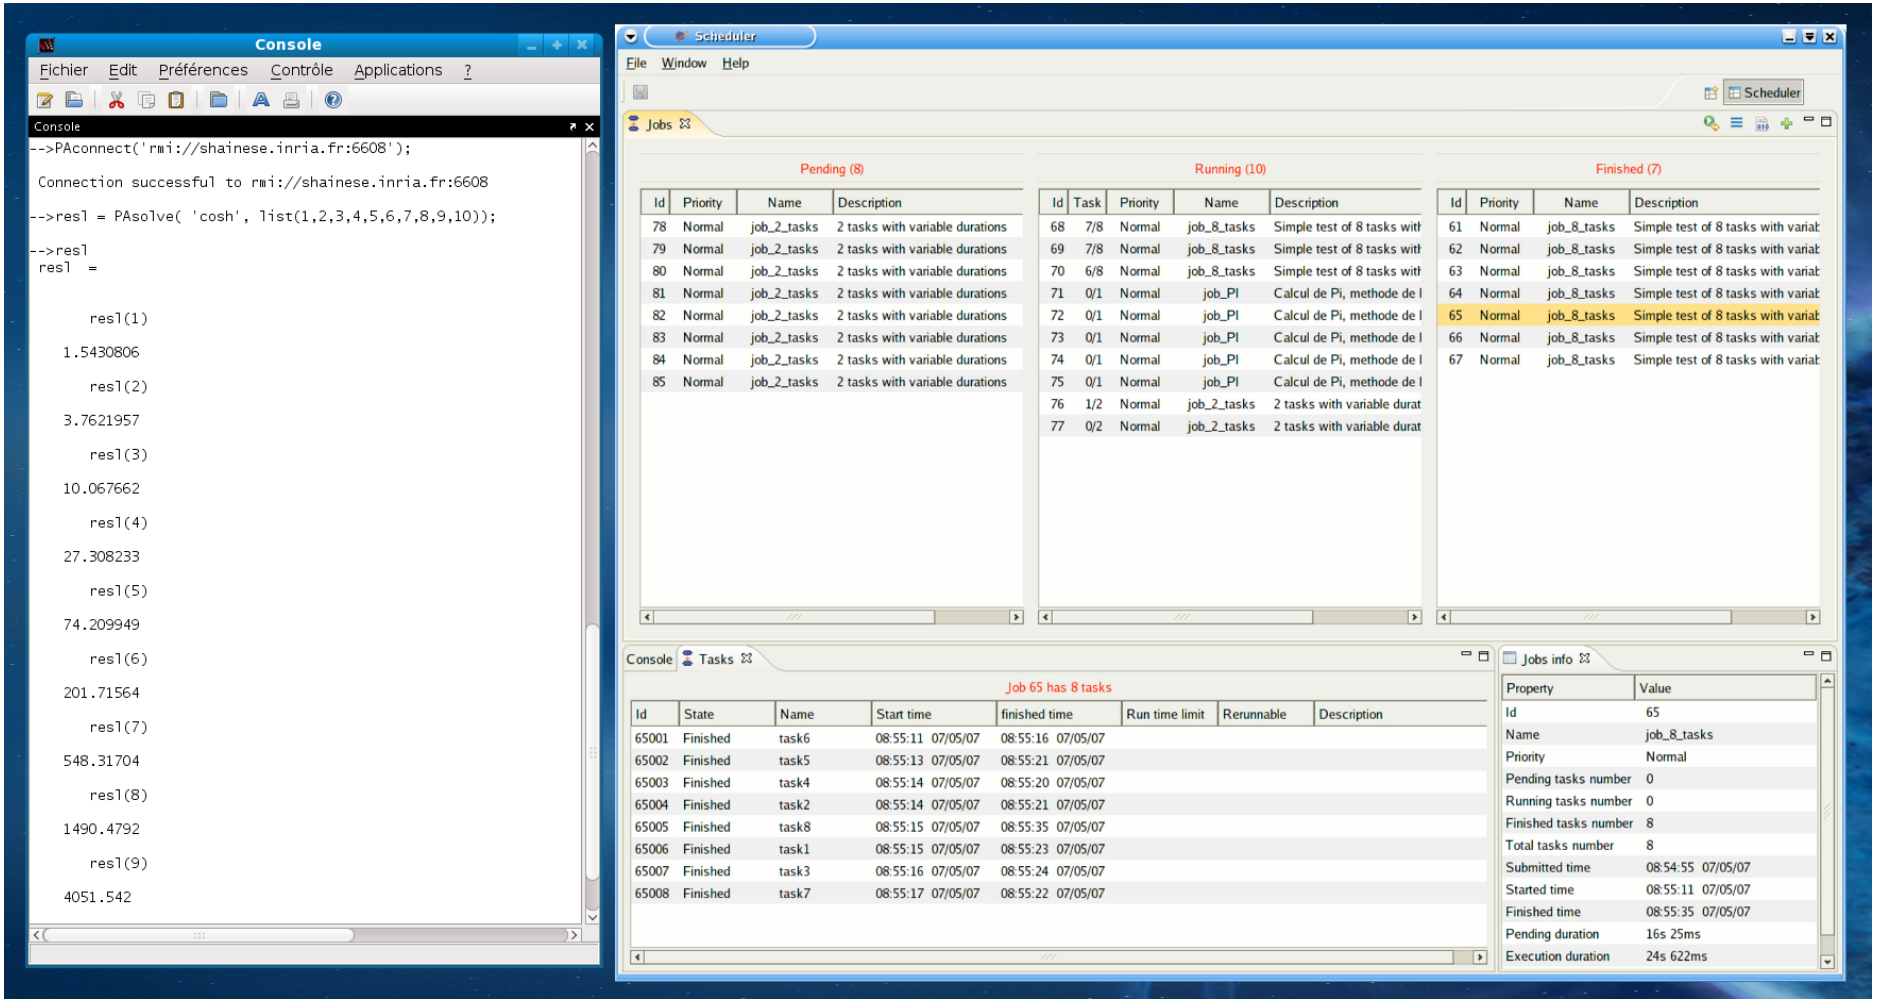

ProActive Parallel Suite

Interface ProActive ⇔ Matlab

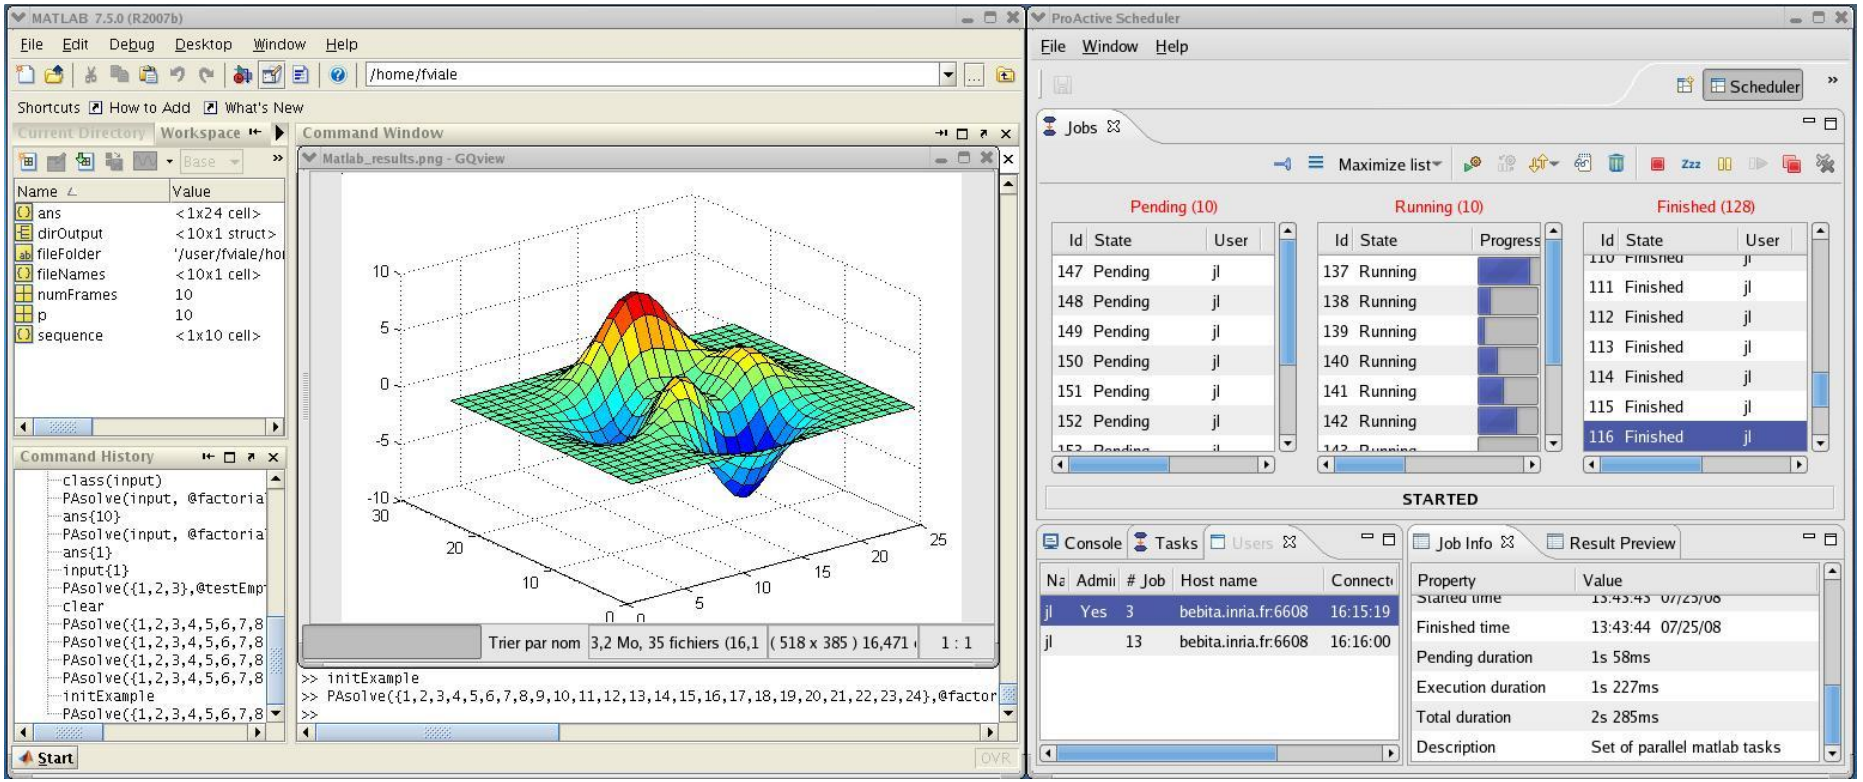

ProActive Parallel Suite

Integration with Scilab and Matlab

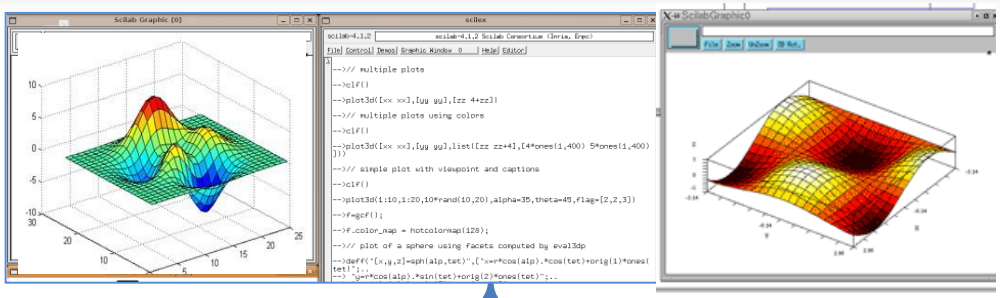

Scilab

MATLAB
The Language of Technical Computing

Static Policy

LSF

Timing Policy 12/24

Desktops

Dynamic Workload Policy

EC2

Dedicated resources

Desktops

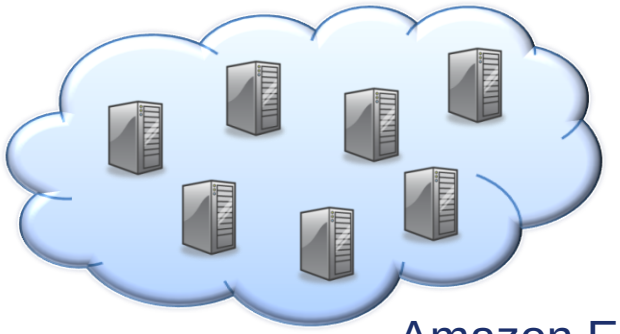

Amazon EC2

ProActive Parallel Suite

Integration with Applications

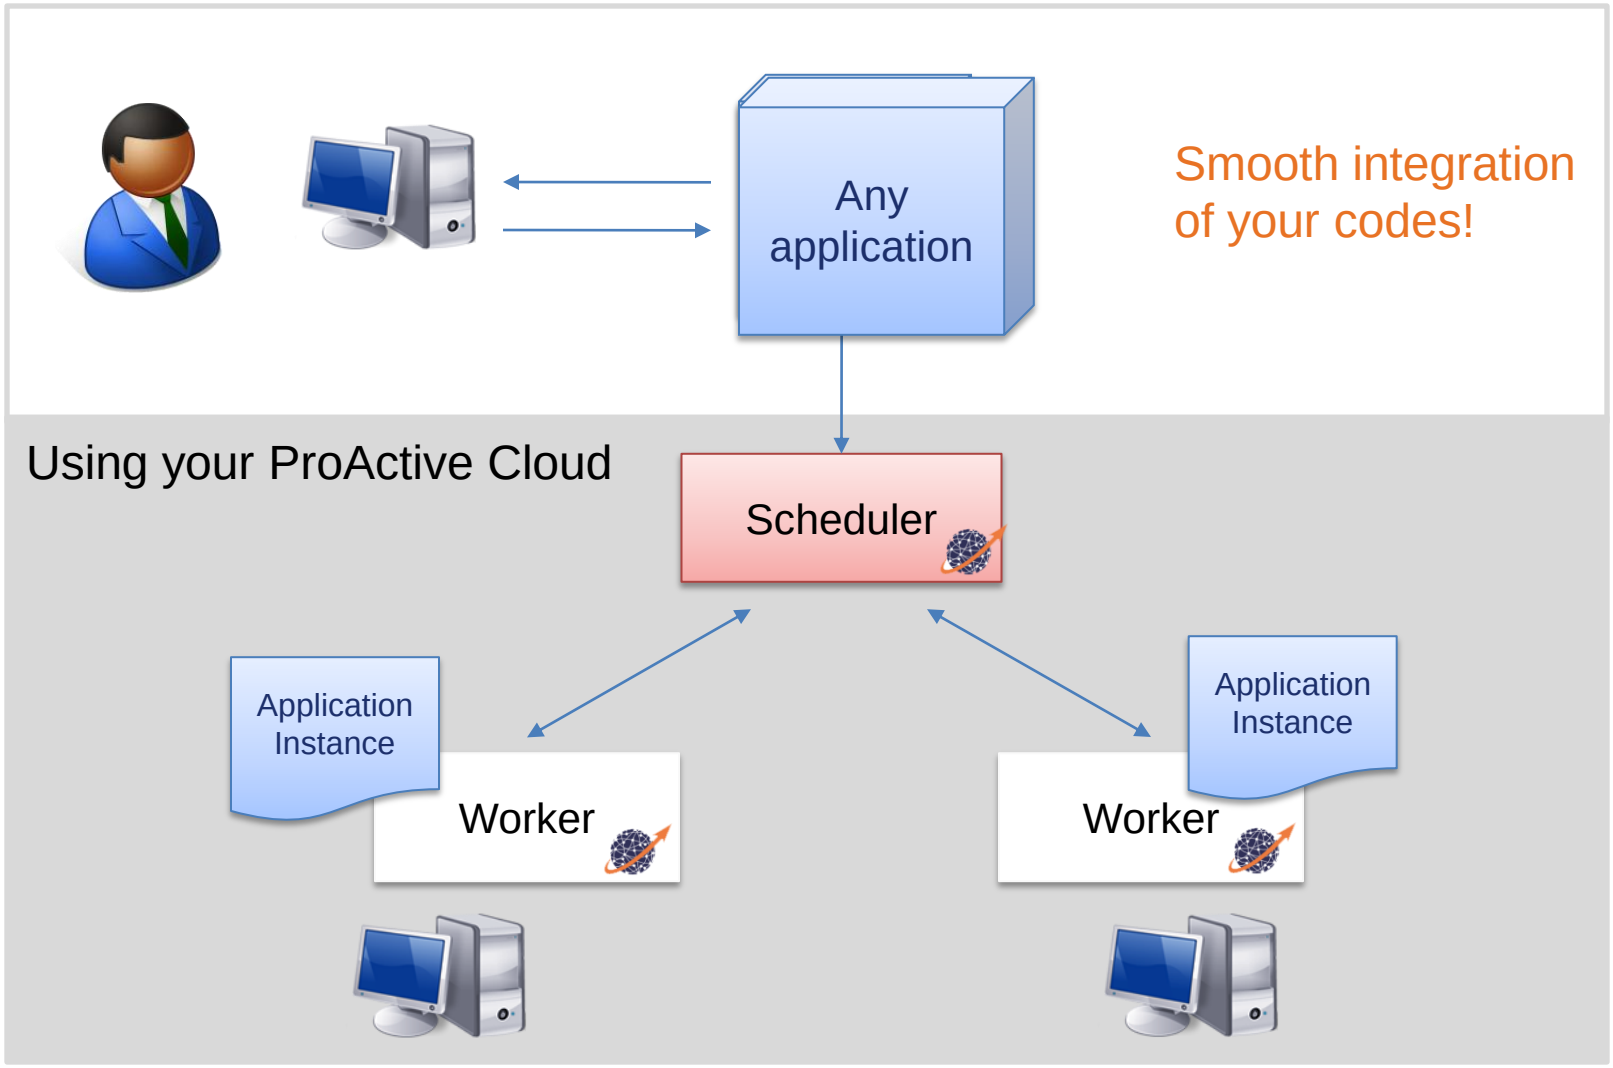

ProActive Parallel Suite

Conclusion

Conclusion

ProActive Parallel Suite

- **Flexibility**
  Clutch Power
- **Portability:**
  Windows, Linux, Mac
- **Versatility:**
  Desktops, Grids, Clouds

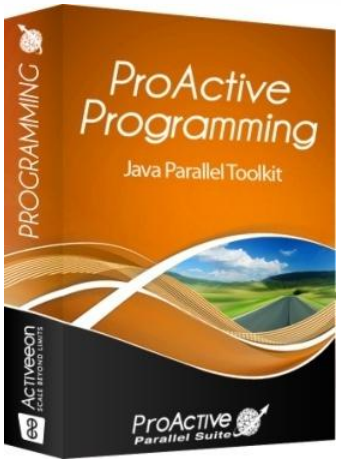

Java Parallel Toolkit

Multi-Platform Job Scheduler

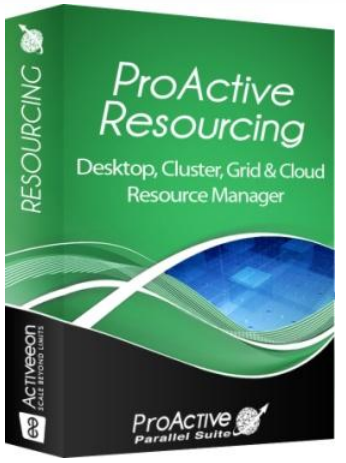

Resource Manager

Free Professional
Open Source Software

ProActive.inria.fr

Multi-Core: No sharing Parallel Programming Model
Cloud: Smooth transition needed:
Consolidation + Interoperability

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

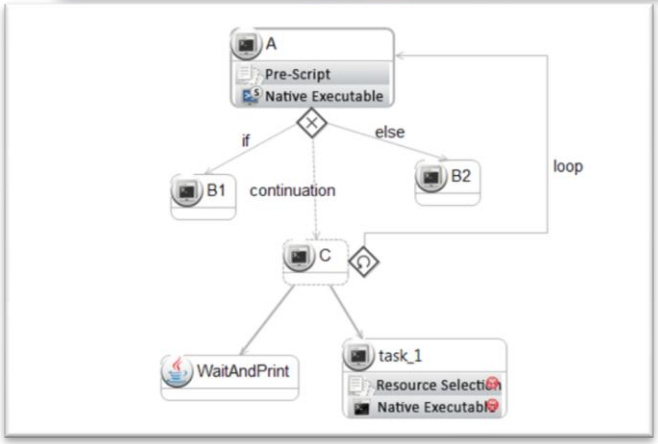

Workflow Execution
Studio Editor and Visualization
Parallel Programming in Java

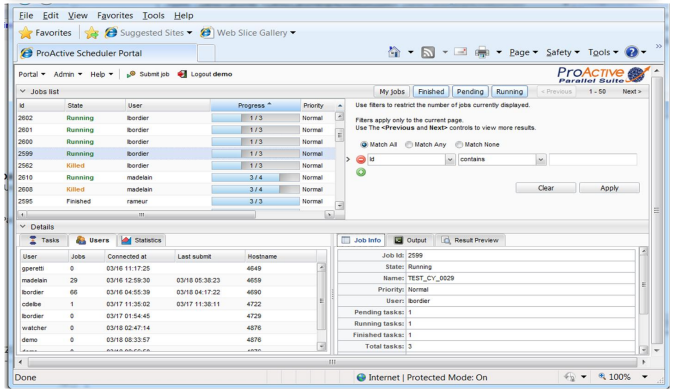

Portal, Multi-Application & Multi-Tenant Enterprise Orchestration

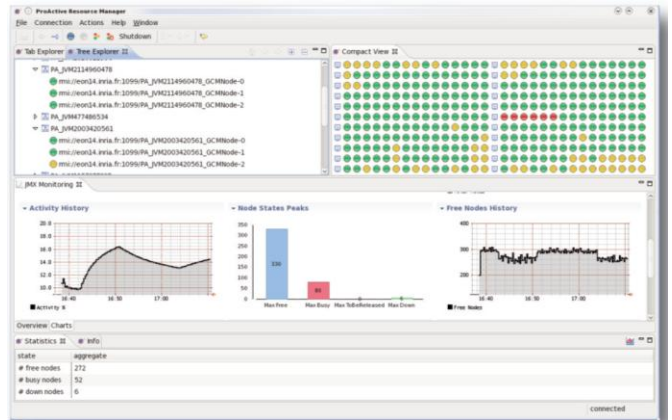

Physical and Virtual Machines Management

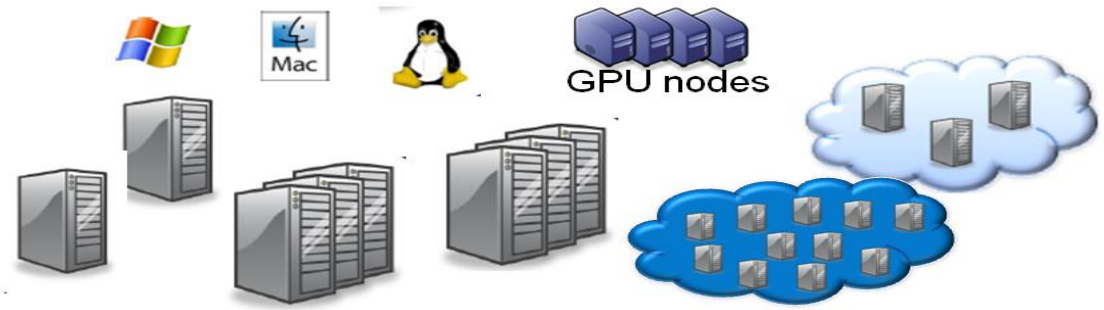

Conclusion

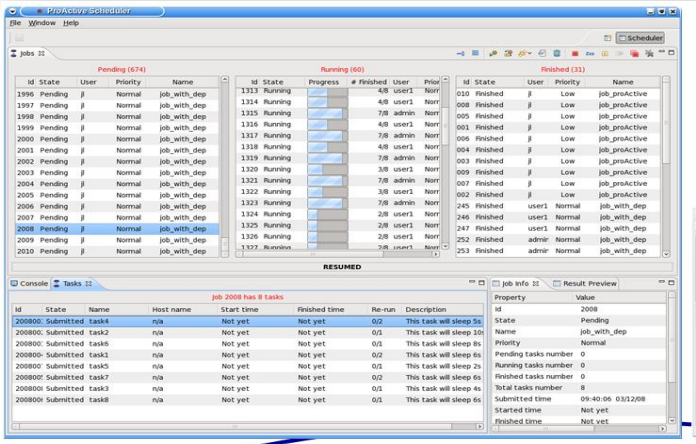

ProActive Parallel Suite

Computing Portal for Heterogeneous Resources

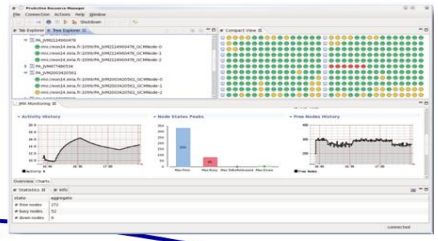

Free Professional
Open Source Software

OW2 Consortium

- **Portability:**
  Windows, Linux, Mac
- **Versatility:**
  Desktops, Grids, Clouds
- **Infrastructure Management**
  Dynamic (Local, Remote, Public Cloud)
- **Application Acceleration:**
  APIs: Java, C++, Script, WS REST full

ProActive Parallel Suite

Industrial (1750) & Cloud Revolution Compared

	Industrial Revolution	Cloud Revolution
Concept	Mechanization and centralization of manufacturing activities	Computing as a Utility Centralization of Data Center, Automation
Technology	Supporting new technos (Mechanic, Tool Machines, etc.)	Distributed Computing Virtualization Multi-Cores Network
Socio Economical Factors	Large new demand was ready to use the new offer. (A change in business attitude & organization)	IT Cost Reduction Pressure CIO Nightmare CEO Out-of-DataCenter CapEx

➔All elements converge for a strong Cloud Revolution

Sources & Inspiration: Simon Wardley (CSC) Scott Stewart

ProActive Parallel Suite

http://proactive.inria.fr

ProActive Parallel Suite


Thank you for your attention!