From Multi-Cores to Clouds with ProActive Parallel Suite: UC in Biotech, IT, Finance, and Engineering

D. Caromel, et al.

Agenda

1. **Background: INRIA, ActiveEon**
2. **CLOUD Computing**
3. **ProActive Parallel Suite**
   Programming, Scheduling, Resourcing
4. **Use Cases & Demos**
5. **Conclusion: Cloud Revolution ?**

Cloud Computing Revolution ?

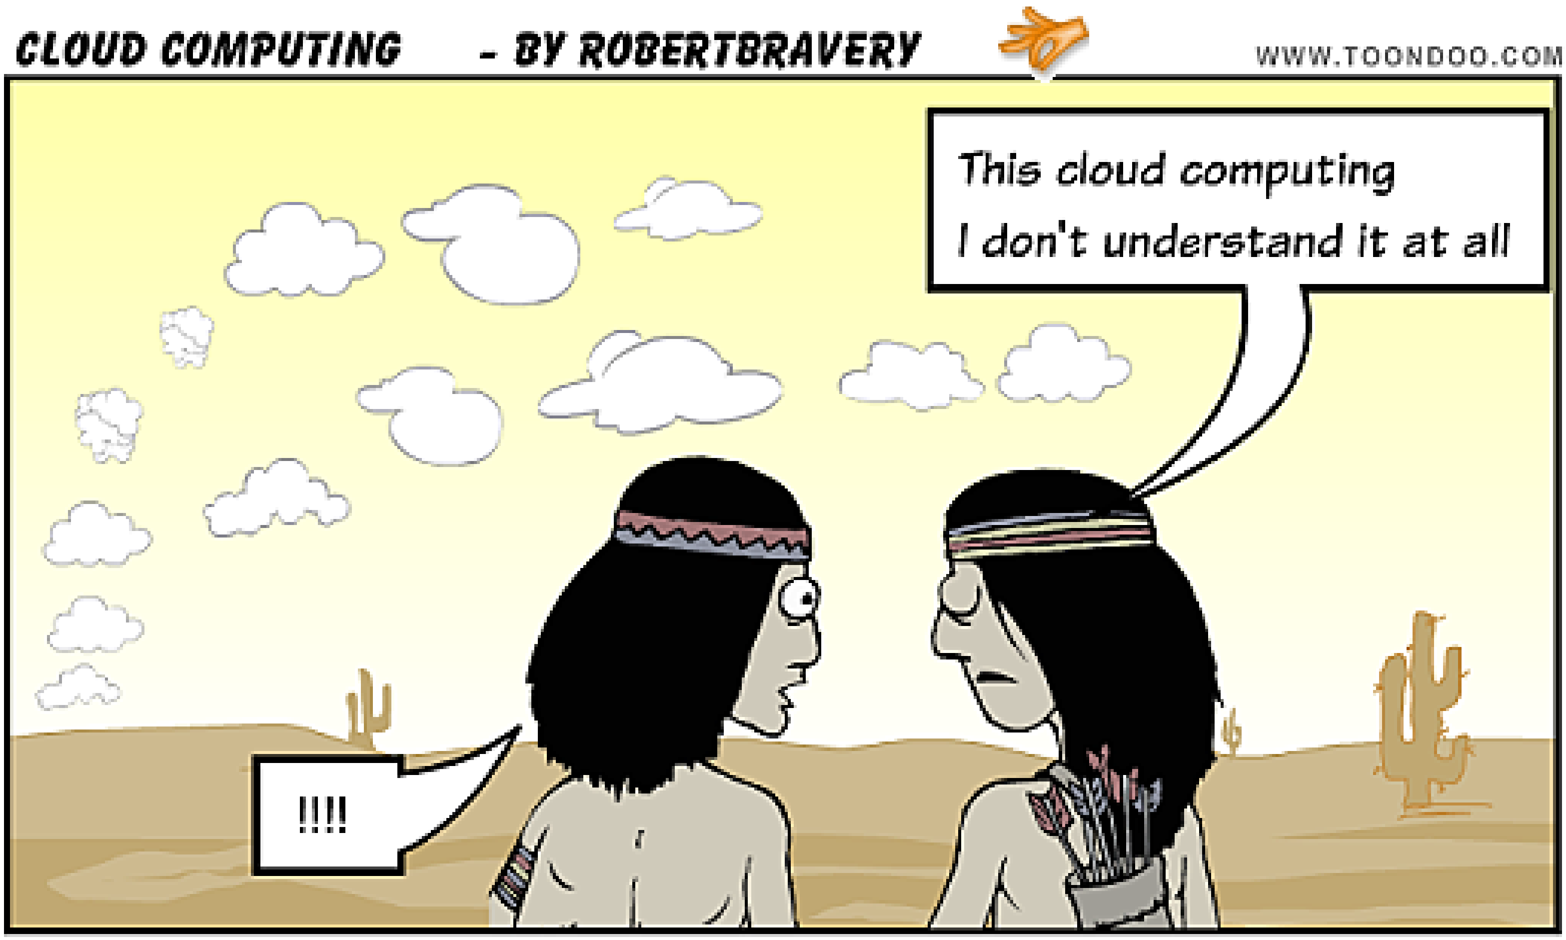

ProActive Parallel Suite

CLOUD Revolution ?

- 1990: PCs
- 2000: Internet for Companies
- 2010: Cloud for Companies

Concept: John McCarthy in 1961 originally coin the expression "Utility Computing" (Electricity, Water, Gas)

Today: How could we do without Internet and Google Search ?
In 2020: we will not imagine working without Clouds

Today: We buy Network, Hardware, Software, Services
Tomorrow: Cloud Services (hiding N, H, S)

ProActive Parallel Suite

1. Background

OASIS (HC: 35)

ProActive Parallel Suite

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 and PhD Scholarship Avail. !

ActiveEon SCALE BEYOND LIMITS

8 INRIA's Research Centres

3 800 HC, 217 M Euro

2 900 Scientists
- 1200 Researchers, Faculty members
- 1200 Doctoral students
- 500 Post-Doct & Visiting scientists

1 000 Engineers, Technicians and Staff

8 Research Centres in France
68 Associated Teams worldwide

4 000 Scientific Publications / year
230 Active patents

89 Innovative companies created

Startup Company Born of INRIA

Some Customers:

Some Partners:

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

2. Cloud Computing

The CLOUD Solution

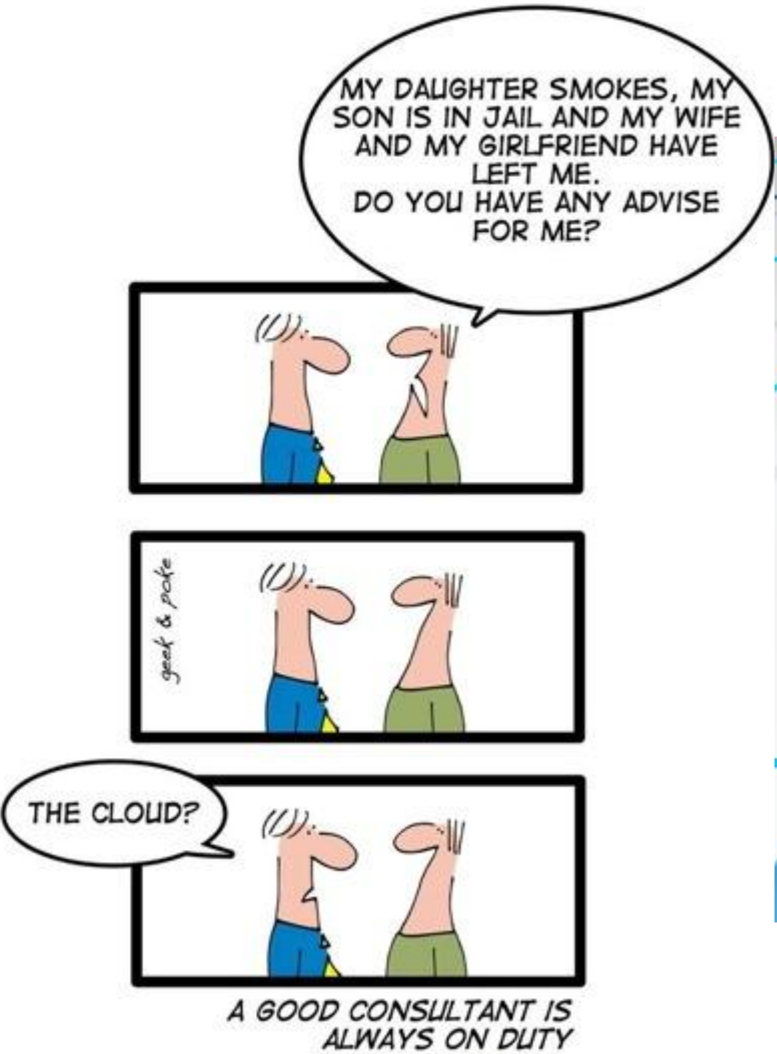

Source: ScienceDaily

Clouds: Basic Definition

- Dynamically scalable, often virtualized resources
- Provided as a service over the Internet
- Users need not have knowledge of, expertise in, or control over the technology infrastructure

XaaS: Anything as a Service

- Software as a service (SaaS), CRM, ERP
- Platform as a service (PaaS), Google App Engine
- Infrastructure as a service (IaaS), Amazon EC2

ProActive Parallel Suite

Clouds in Picture

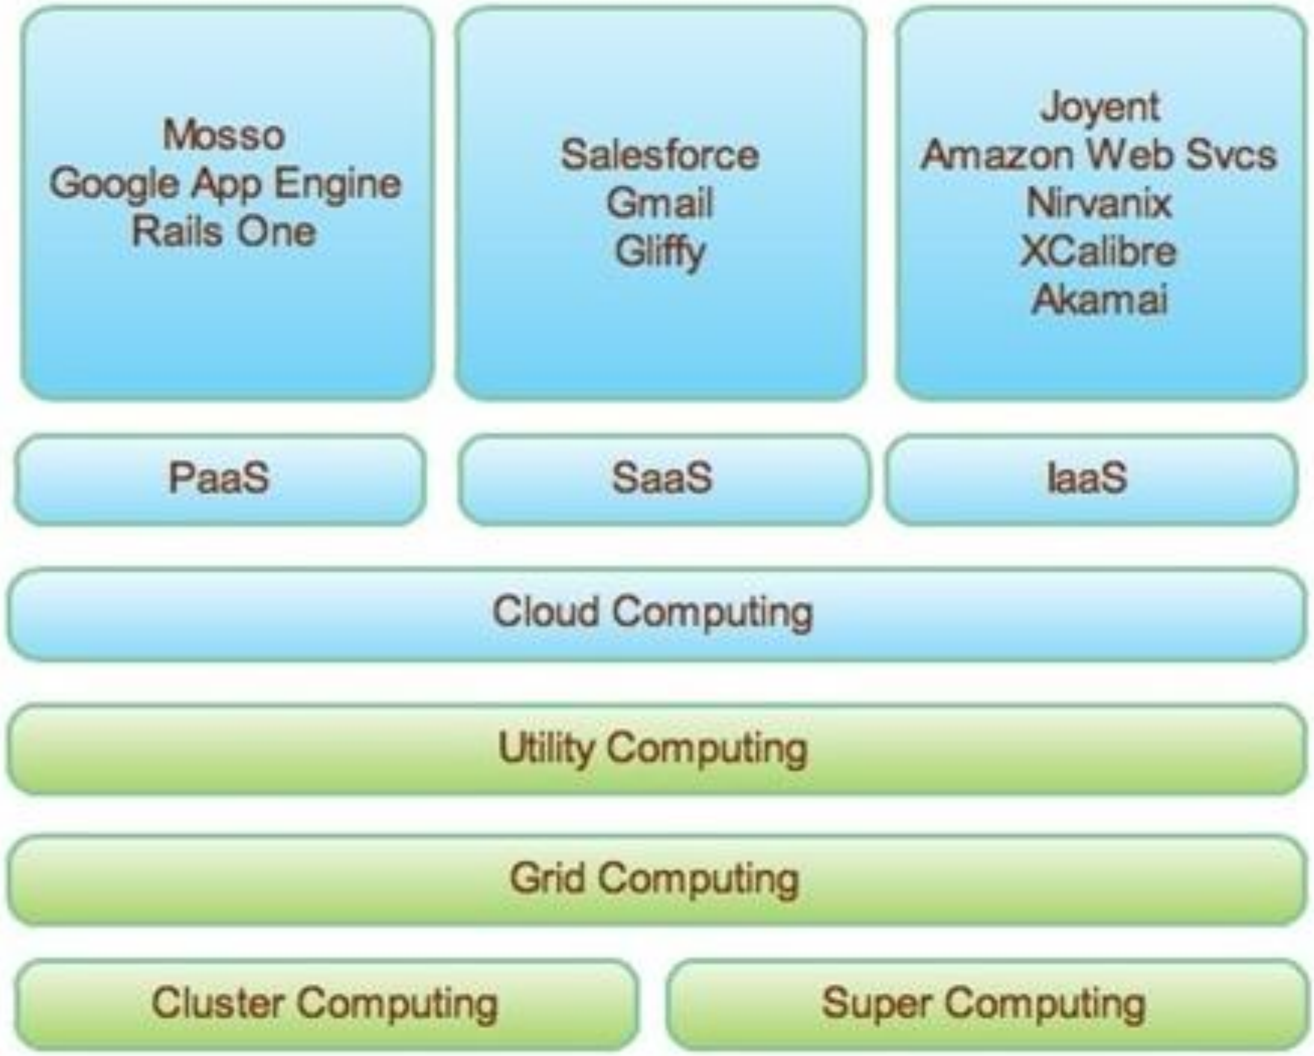

From [Joseph Kent Langley]

ProActive Parallel Suite

From Grids to Clouds

- Grid Computing
  - Several administrative Domains
  - Virtual Organizations
  - Trading not based on Currency

→(Too) Hard

- Cloud solves the issue:
  - Pay as you Go

Distributed, //, & Grid Technologies for Clouds

ProActive Parallel Suite

Multi-Core Push

Symetrical Multi-Core: 8-ways Niagara II

- 8 cores
- 4 Native threads per core
- Linux see 32 cores!

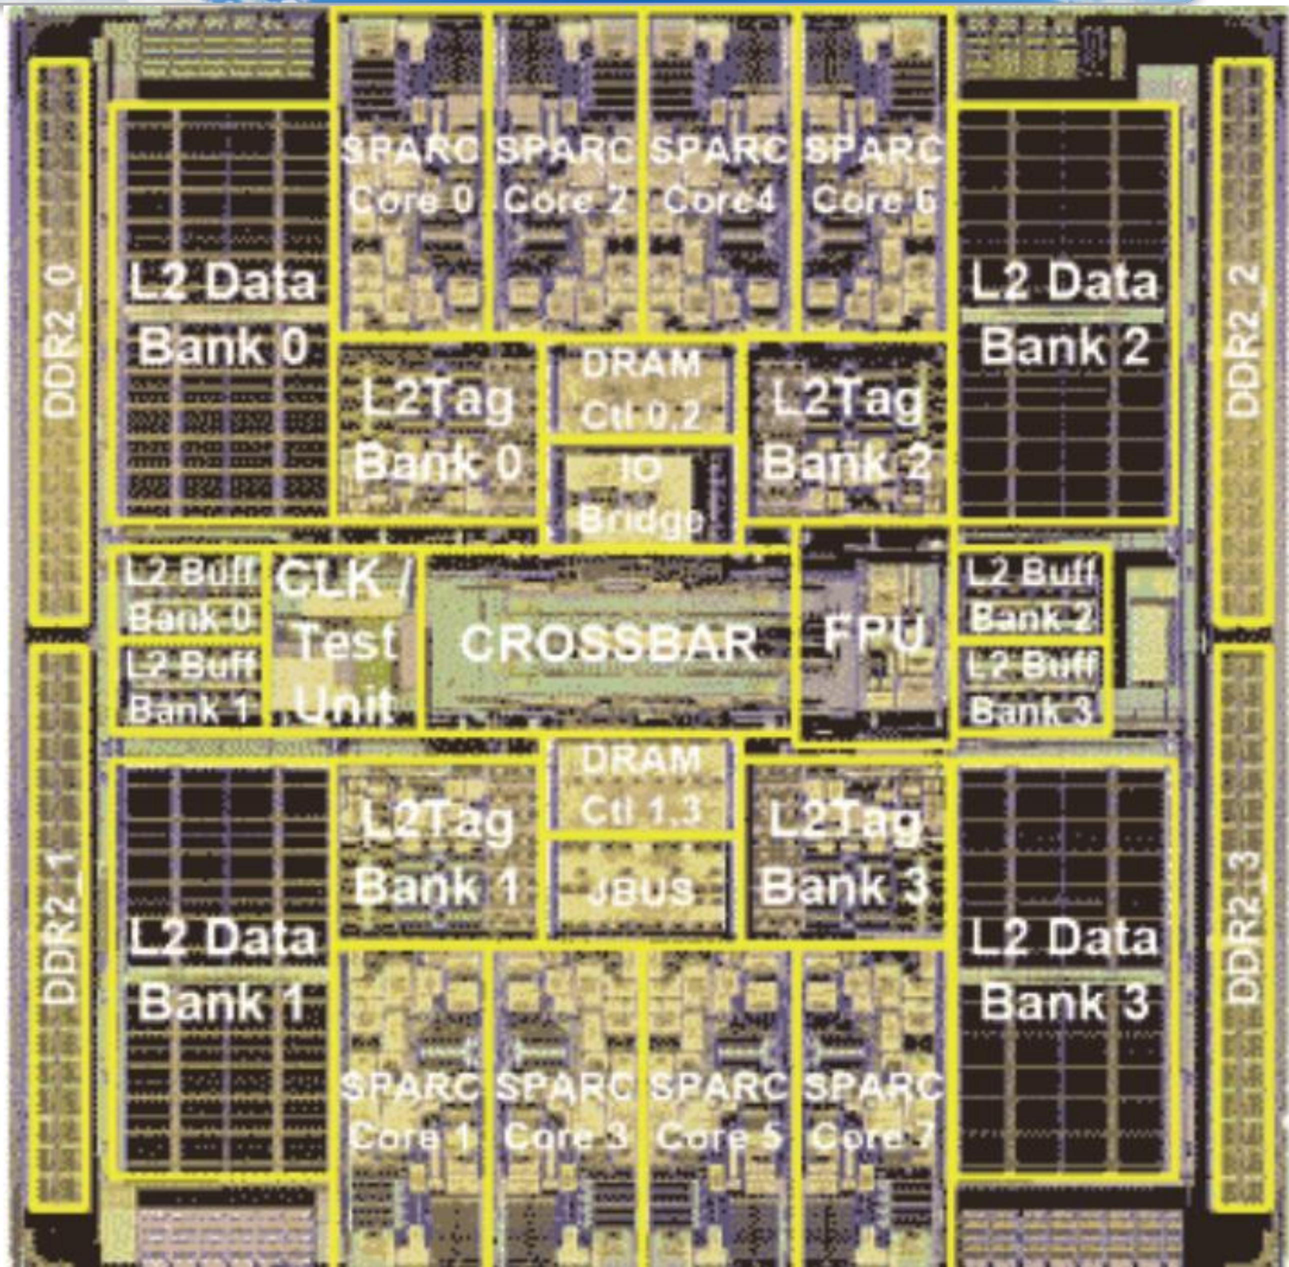

Today Off The Shelf Multi-Cores, 3 GHz

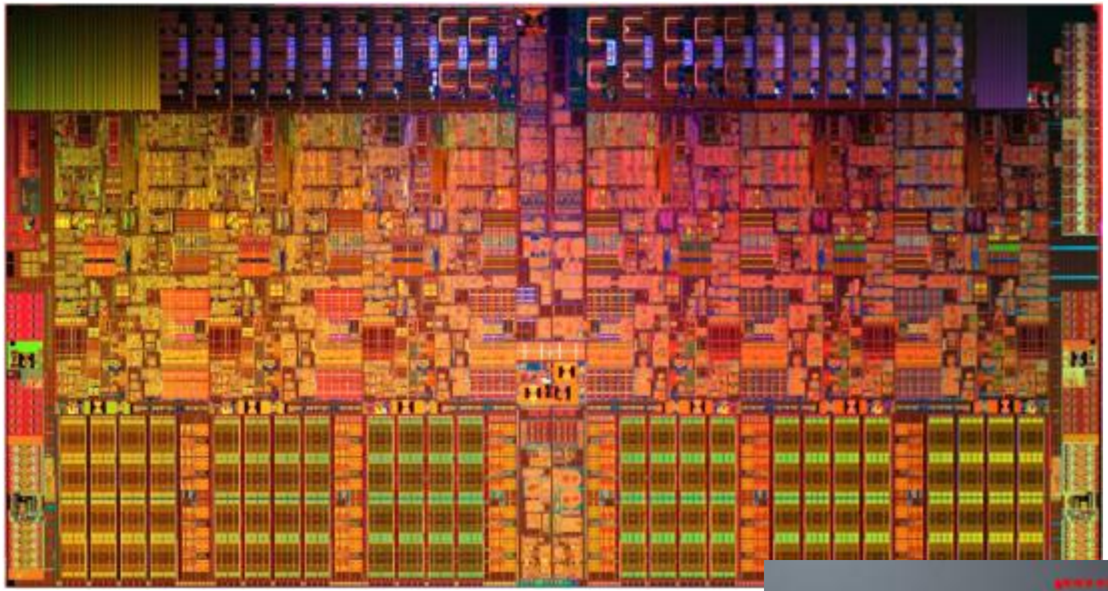

Intel Xeon 5670,
6 cores

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

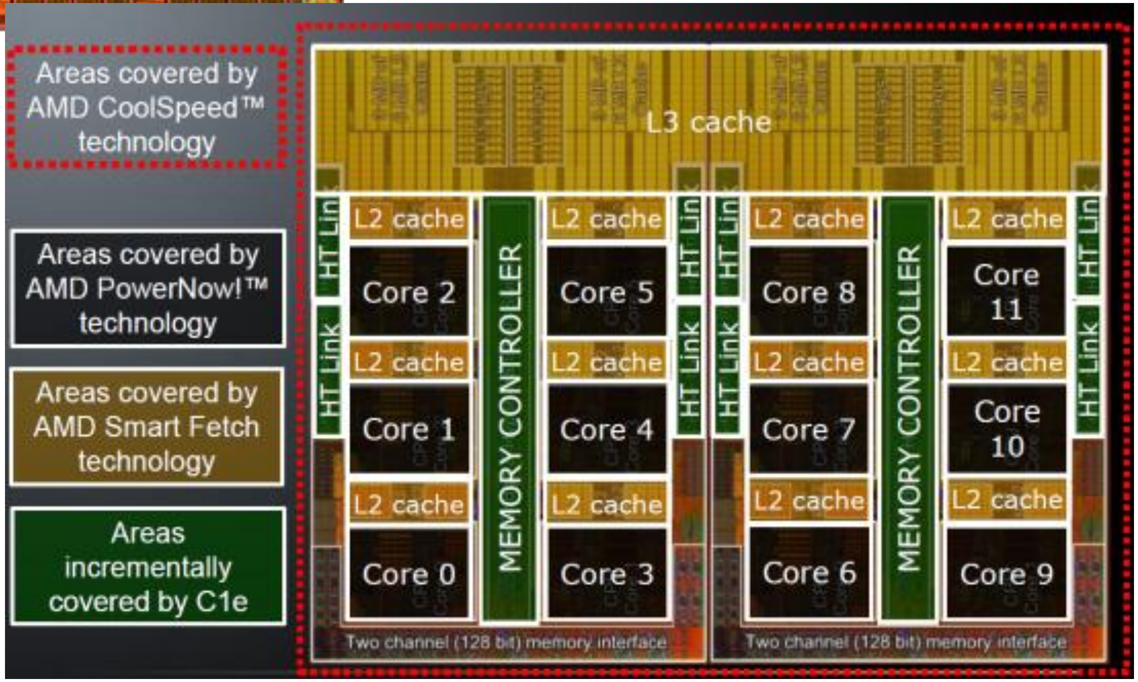

Multi-Cores: A Few Key Points

- Moore's Law rephrased:
  Nb. of Cores double every 18 to 24 months

- Key expected Milestones: Cores per Chips (OTS)
  - 2012: 32 to 64
  - 2014: 64 to 128

- 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

Virtualization

Virtualization

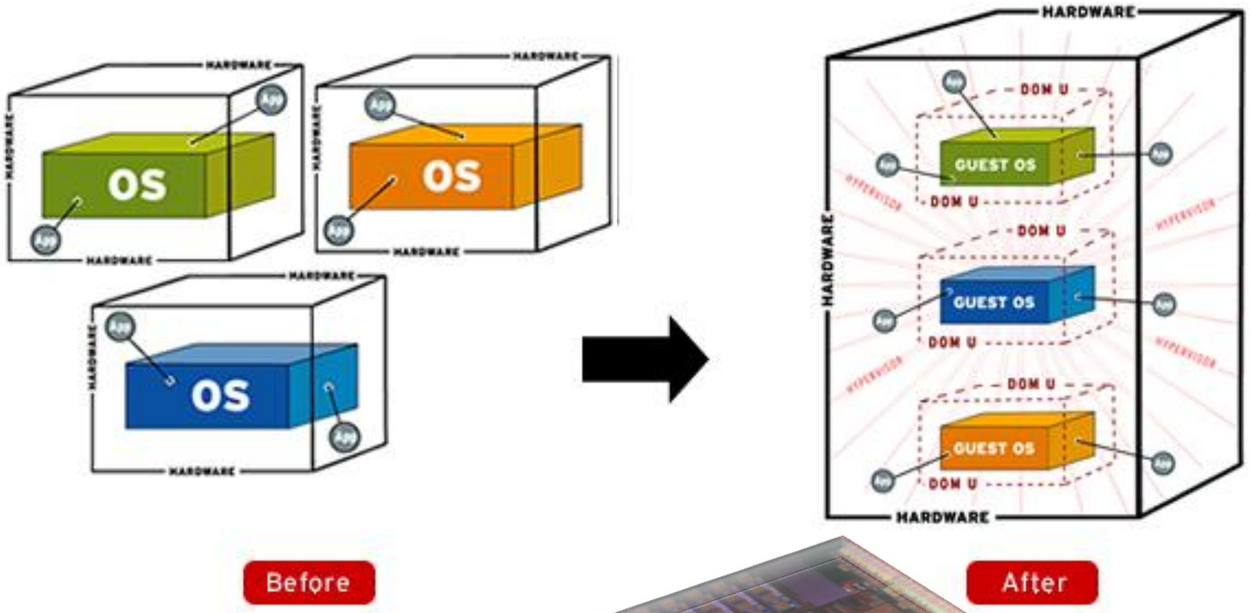

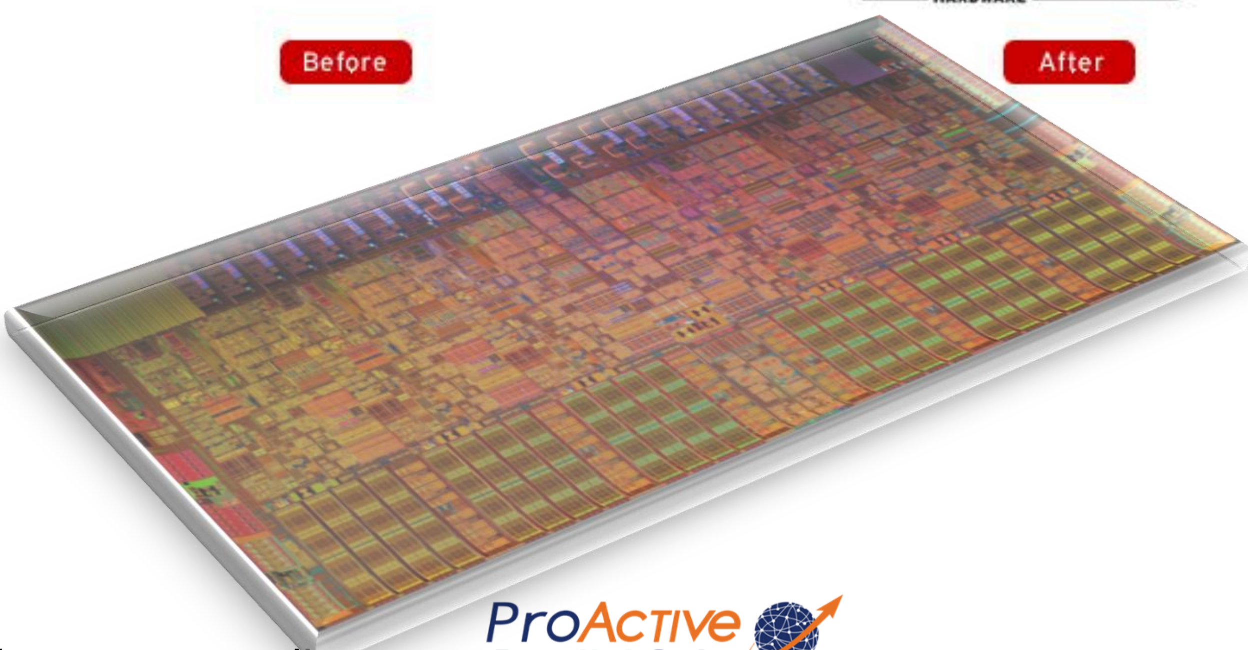

Source: http://www.apac.redhat.com

Virtualization

Sun, Blog [Marc Hamilton]

ProActive Parallel Suite

Virtualization

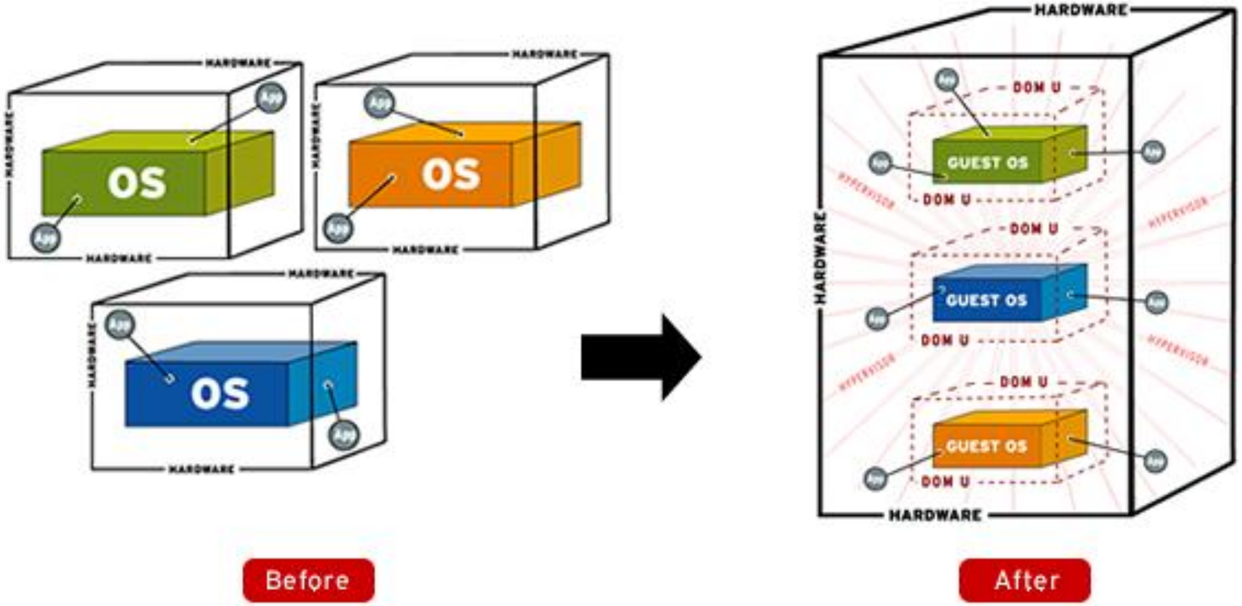

What we Used to do as Syst. Admin.

ProActive Parallel Suite

With Virtualization + Software Appliance

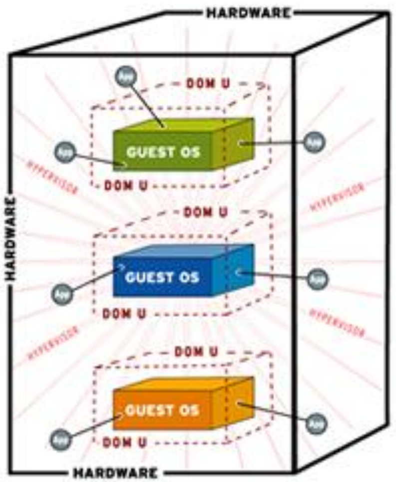

ProActive Parallel Suite

3. ProActive Parallel Suite

ProActive
Parallel Suite

Cloud Solution: ProActive Parallel Suite

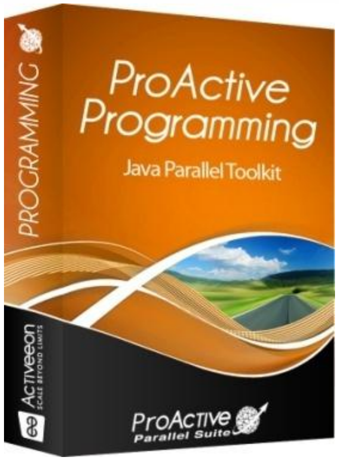

Java Parallel Toolkit

Multi-Platform Job Scheduler

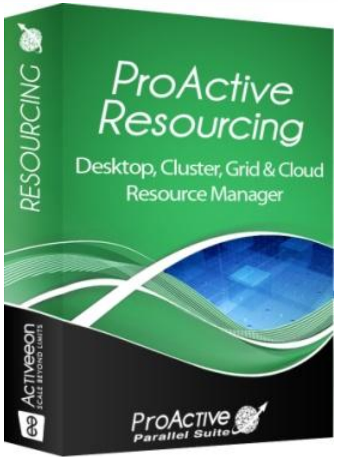

Resource Manager

Used in Production Today:
50 Cores → 300 Cores 2010

Strong Differentiation:

- ❑Java Parallel Programming + Integration +
- ❑Portability: Linux, Windows, Mac +
- ❑Versatility: Desktops, Cluster, Grid, Clouds = Perfect Flexibility

ProActive Programming: Active Objects

ProActive Programming View

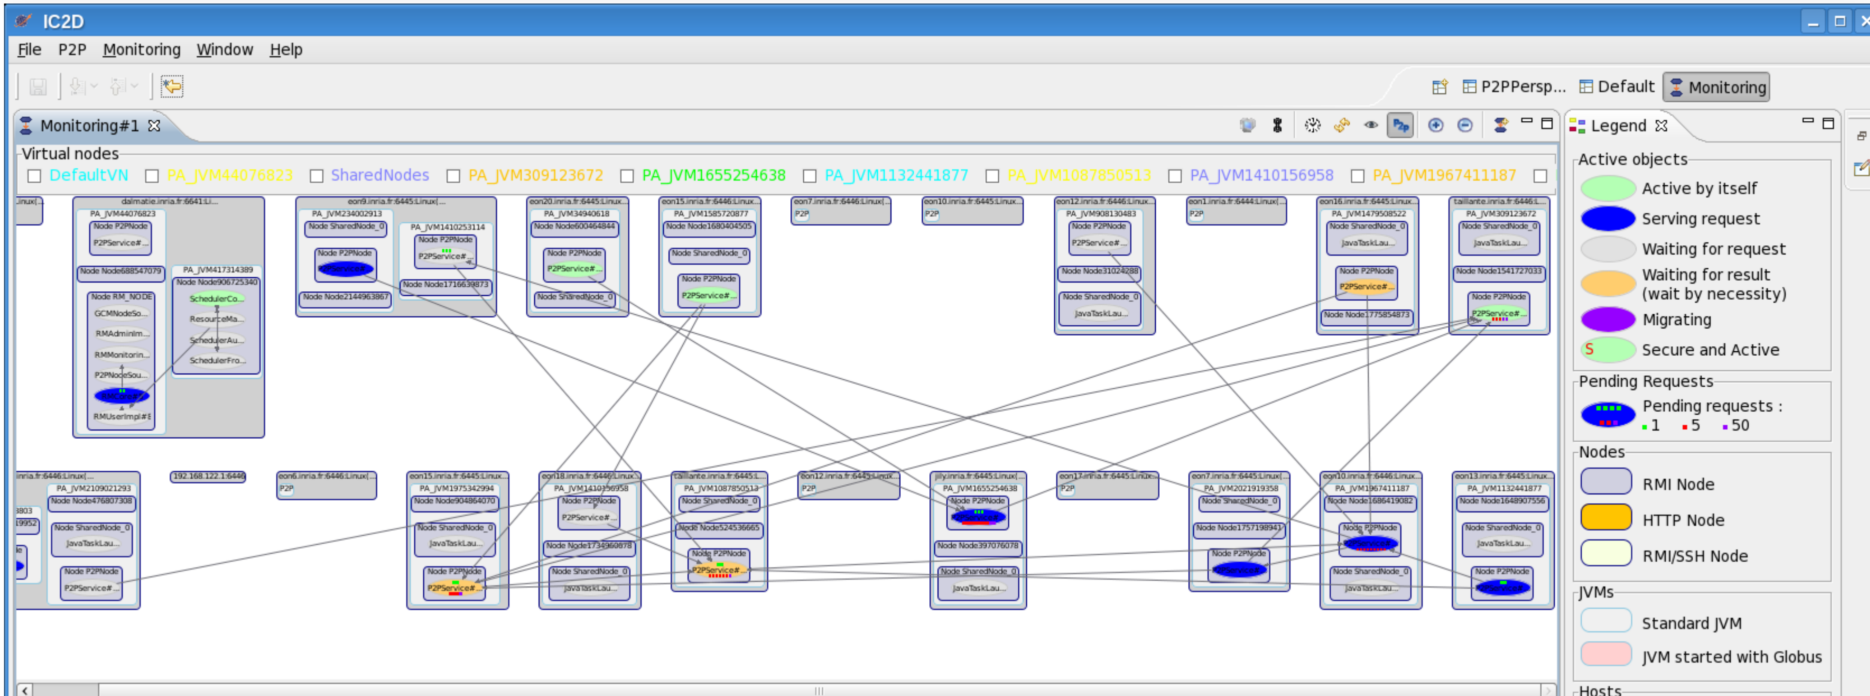

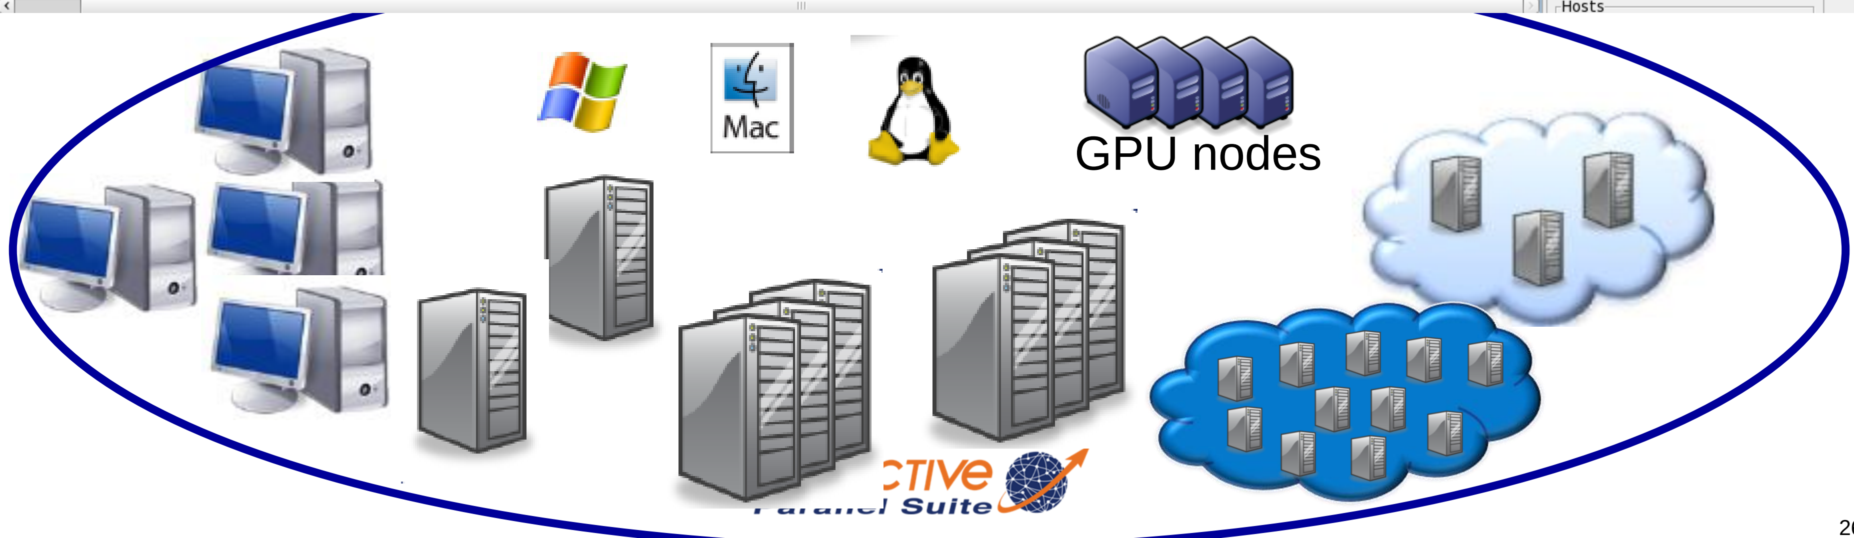

ProActive Programming View

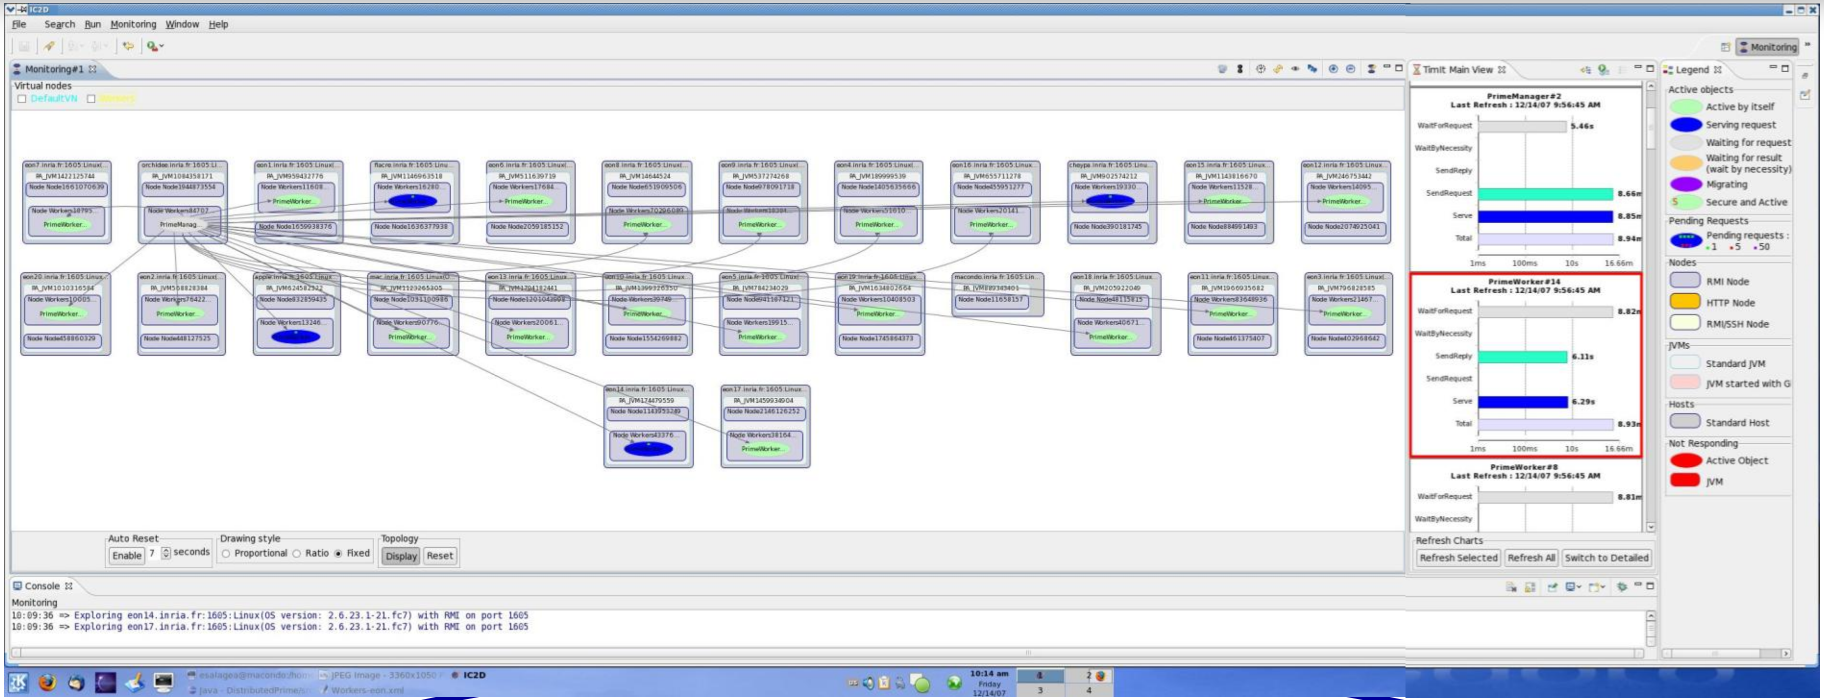

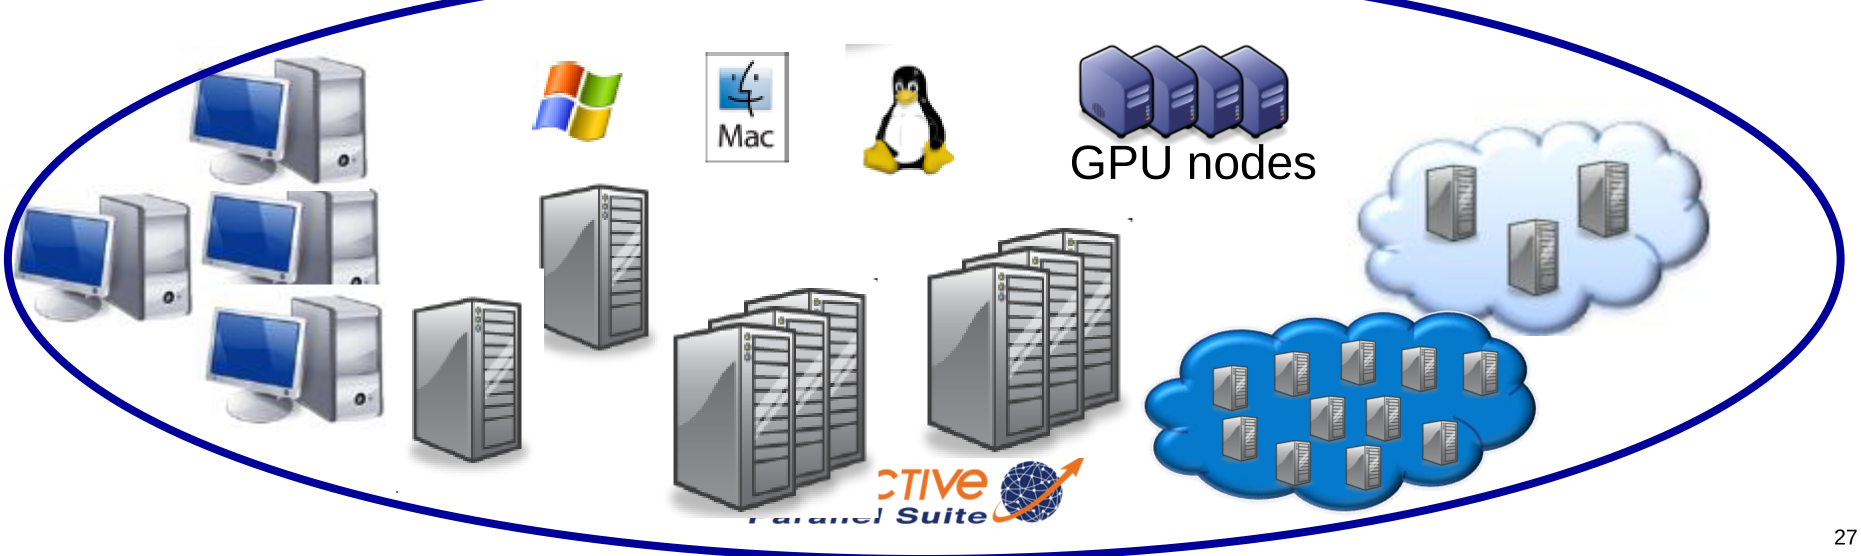

ProActive : Active objects

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

  `...`
- `v1.bar(); //Wait-By-Necessity`

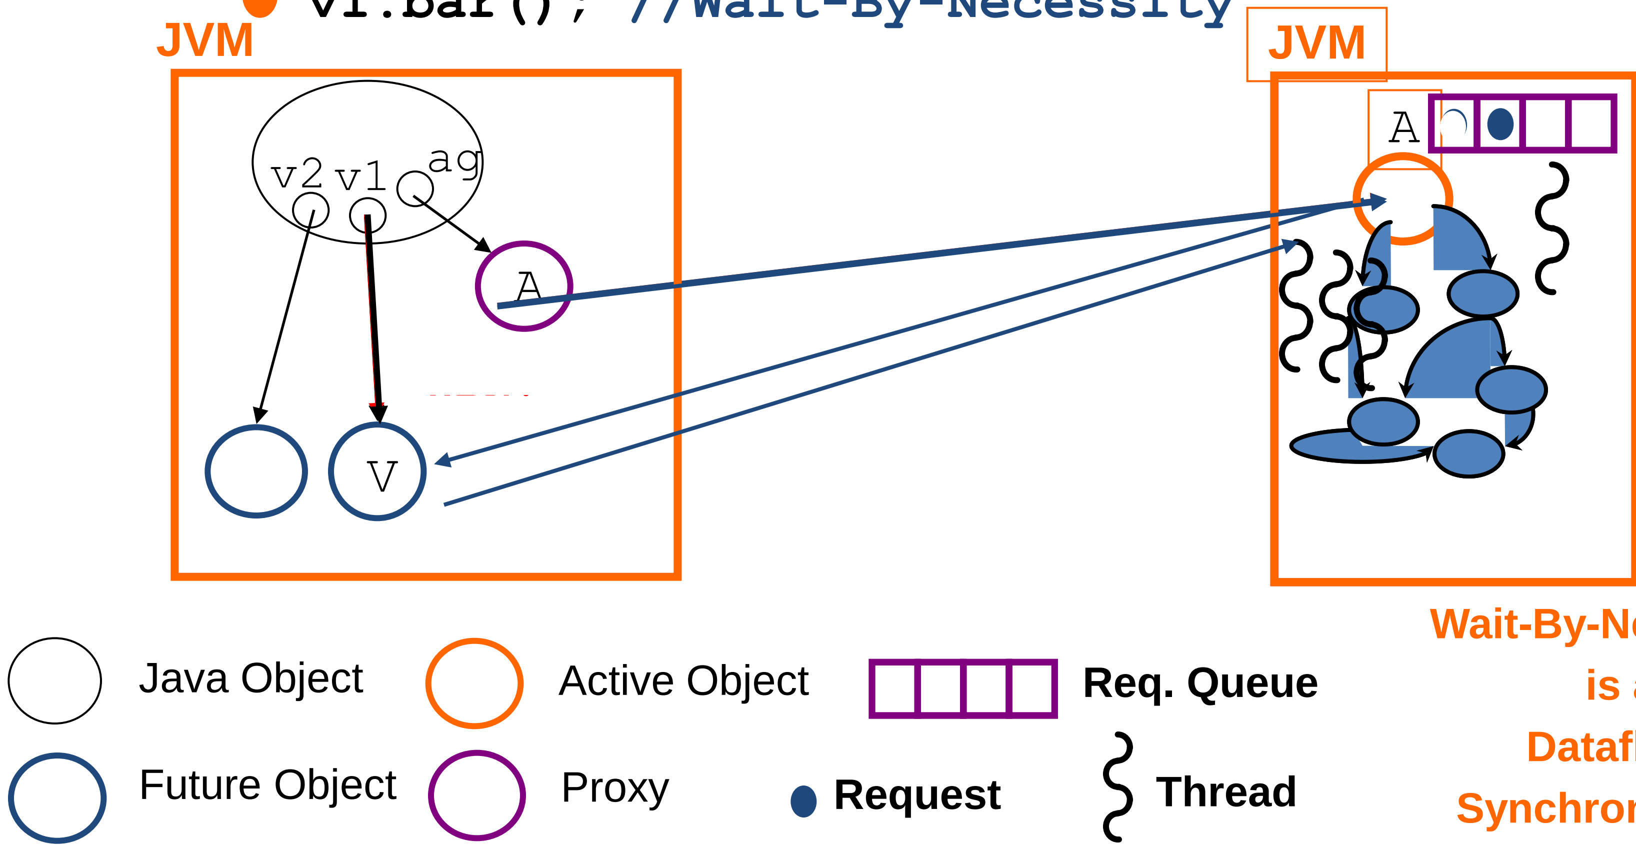

Wait-By-Necessity is a Dataflow Synchronization

ProActive Parallel Suite

Standard system at Runtime: No Sharing

NoC: Network On Chip

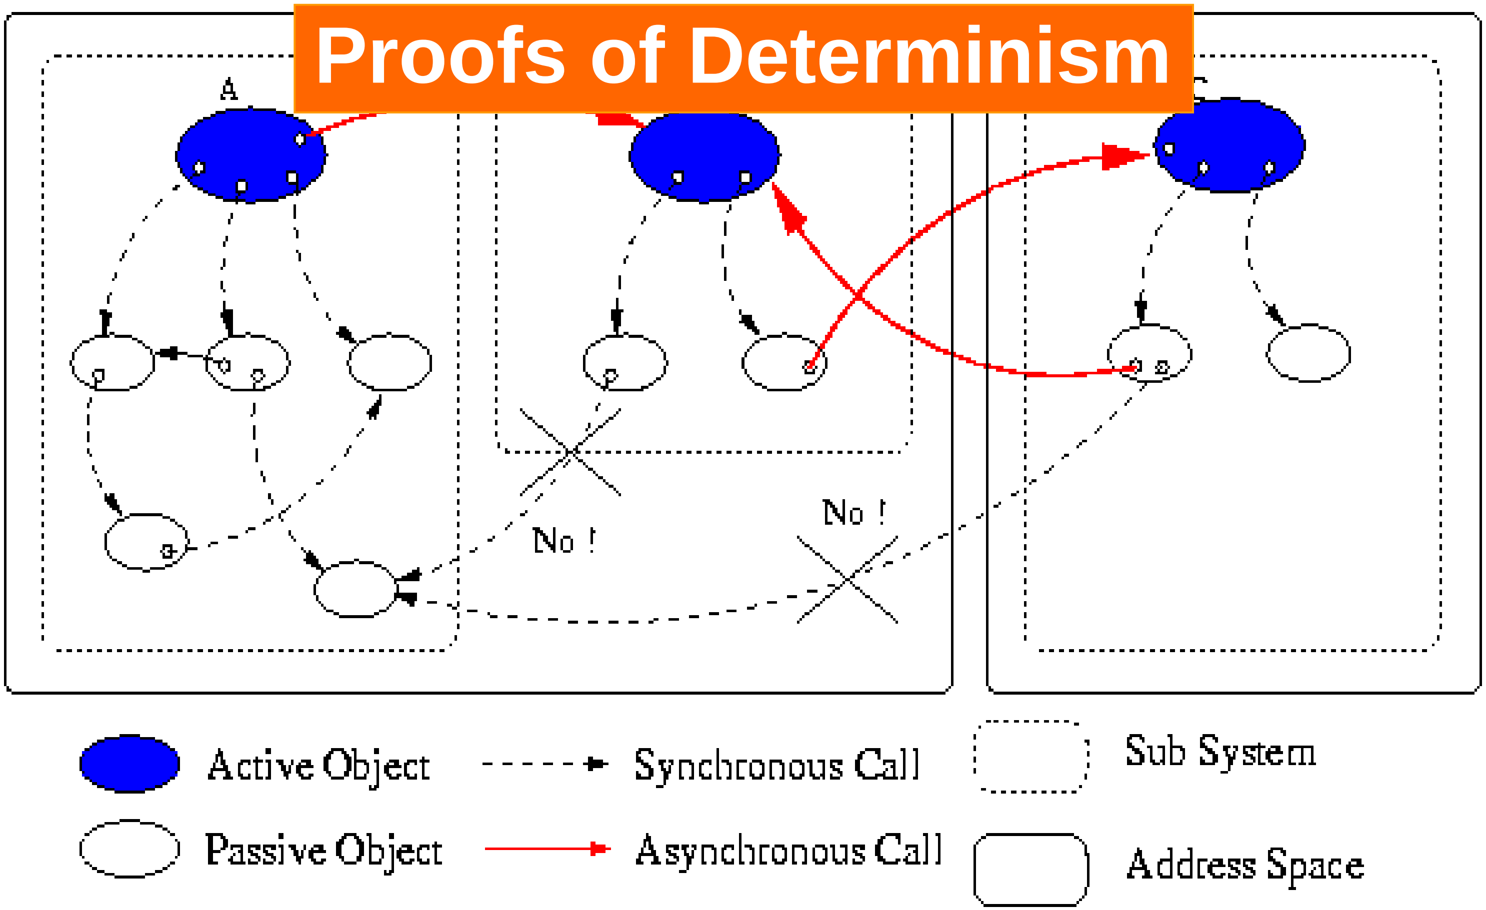

ProActive Parallel Suite

ASP: Asynchronous Sequential Processes

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

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

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

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

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

Sending a Reply

Standard system at Runtime: No Sharing

NoC: Network On Chip

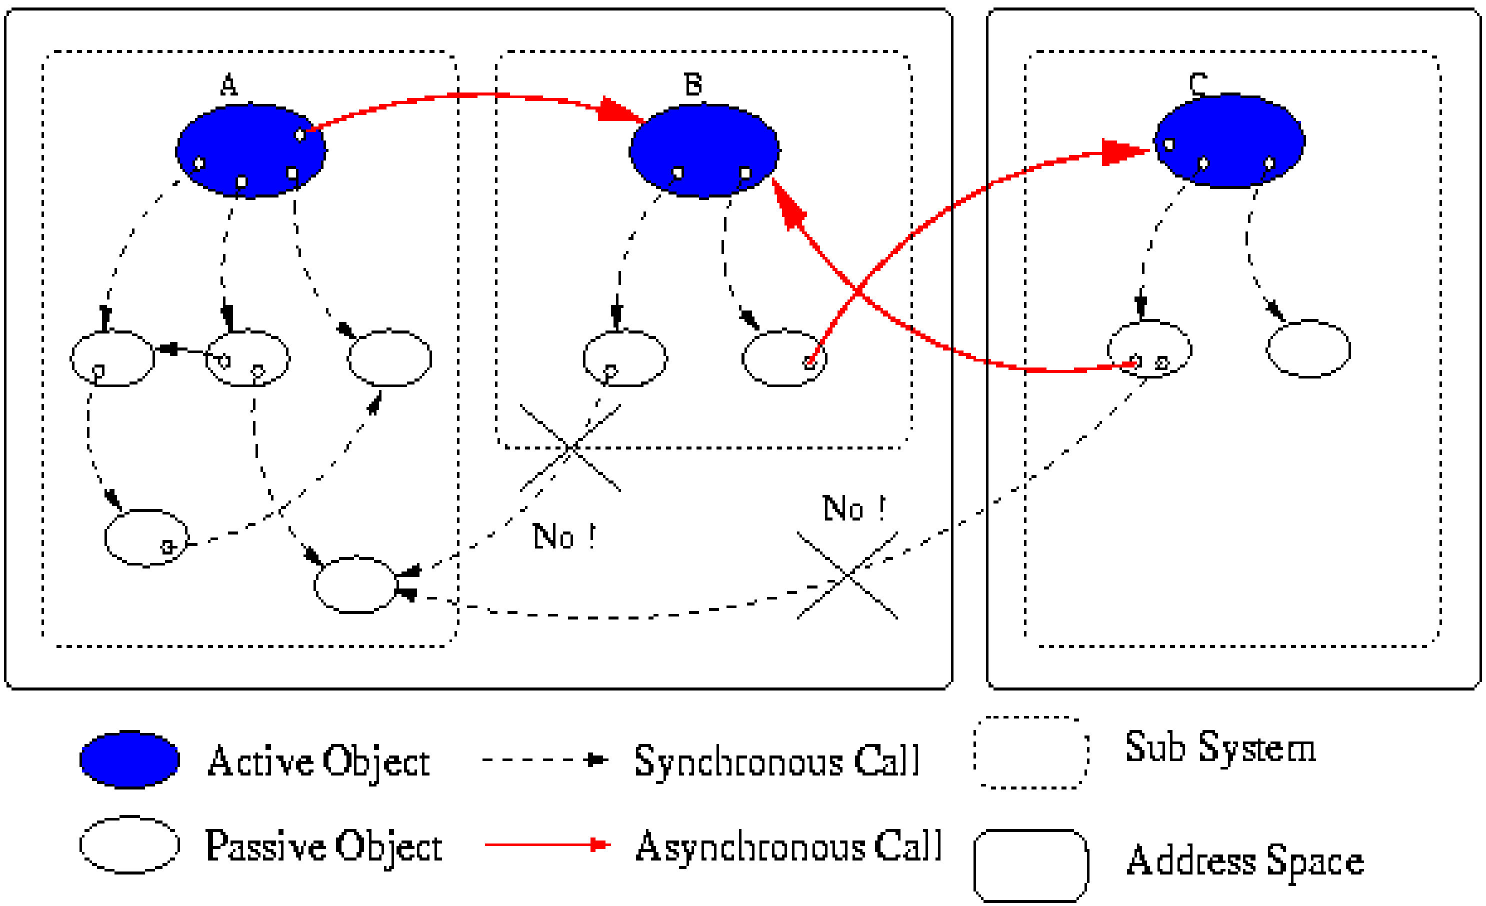

ProActive Parallel Suite

Distributed Objects On Chip

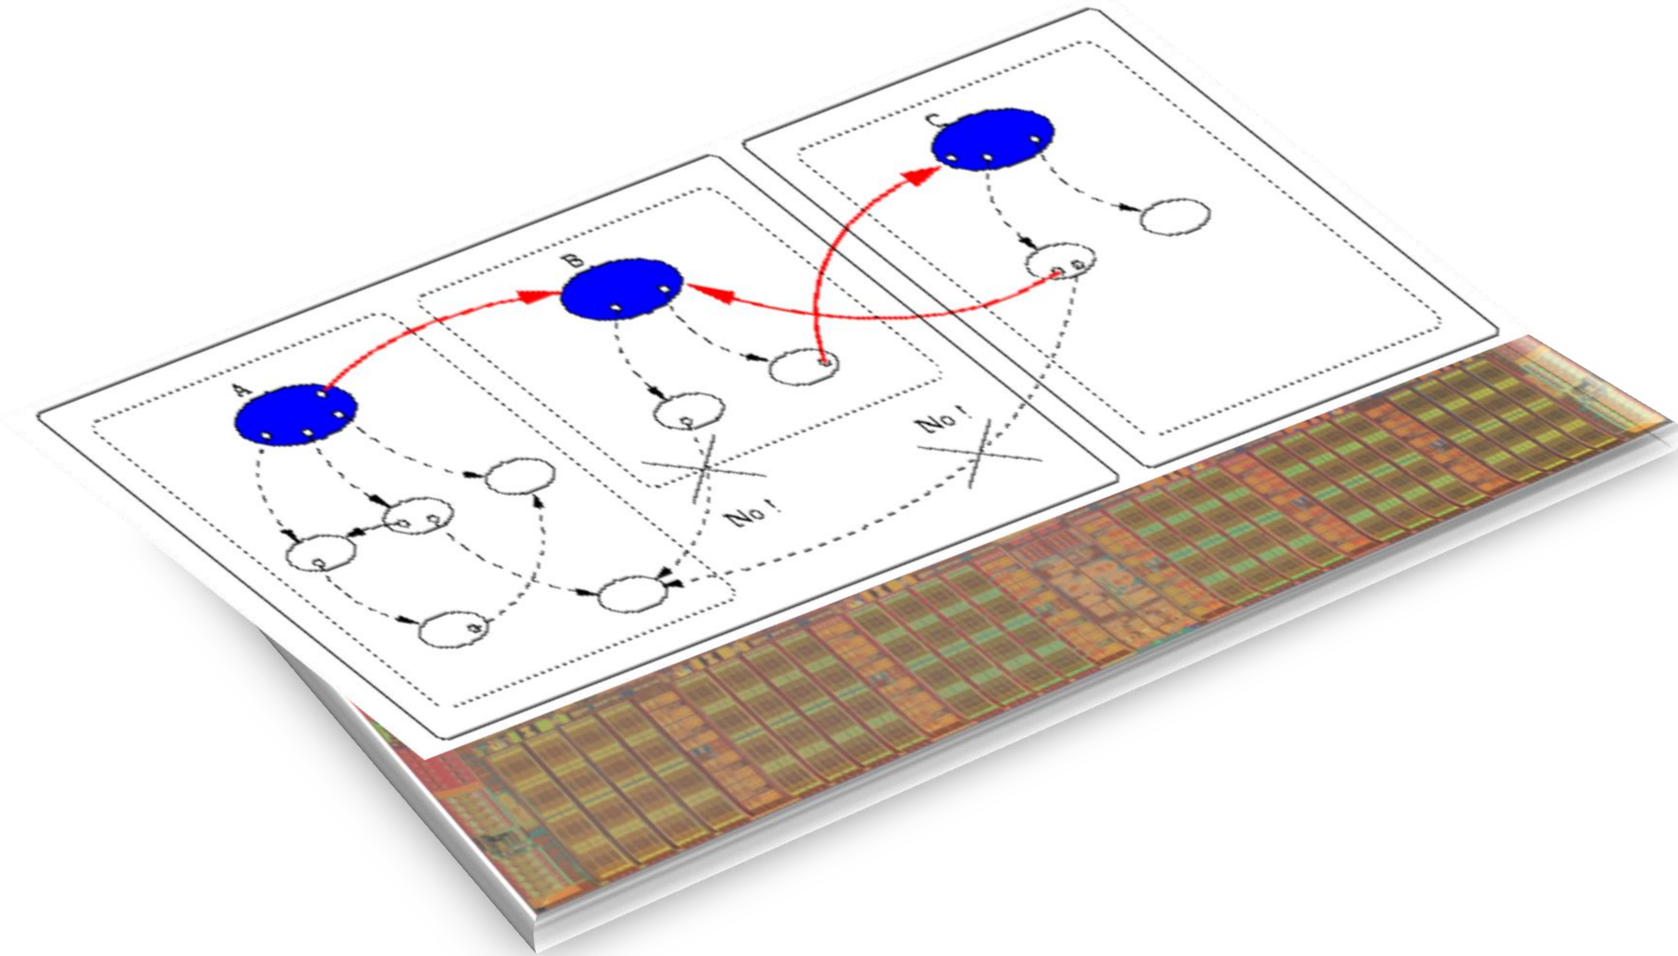

ProActive Parallel Suite

Distributed Objects On Chip, Boards, Clouds

Bi-Socket Board

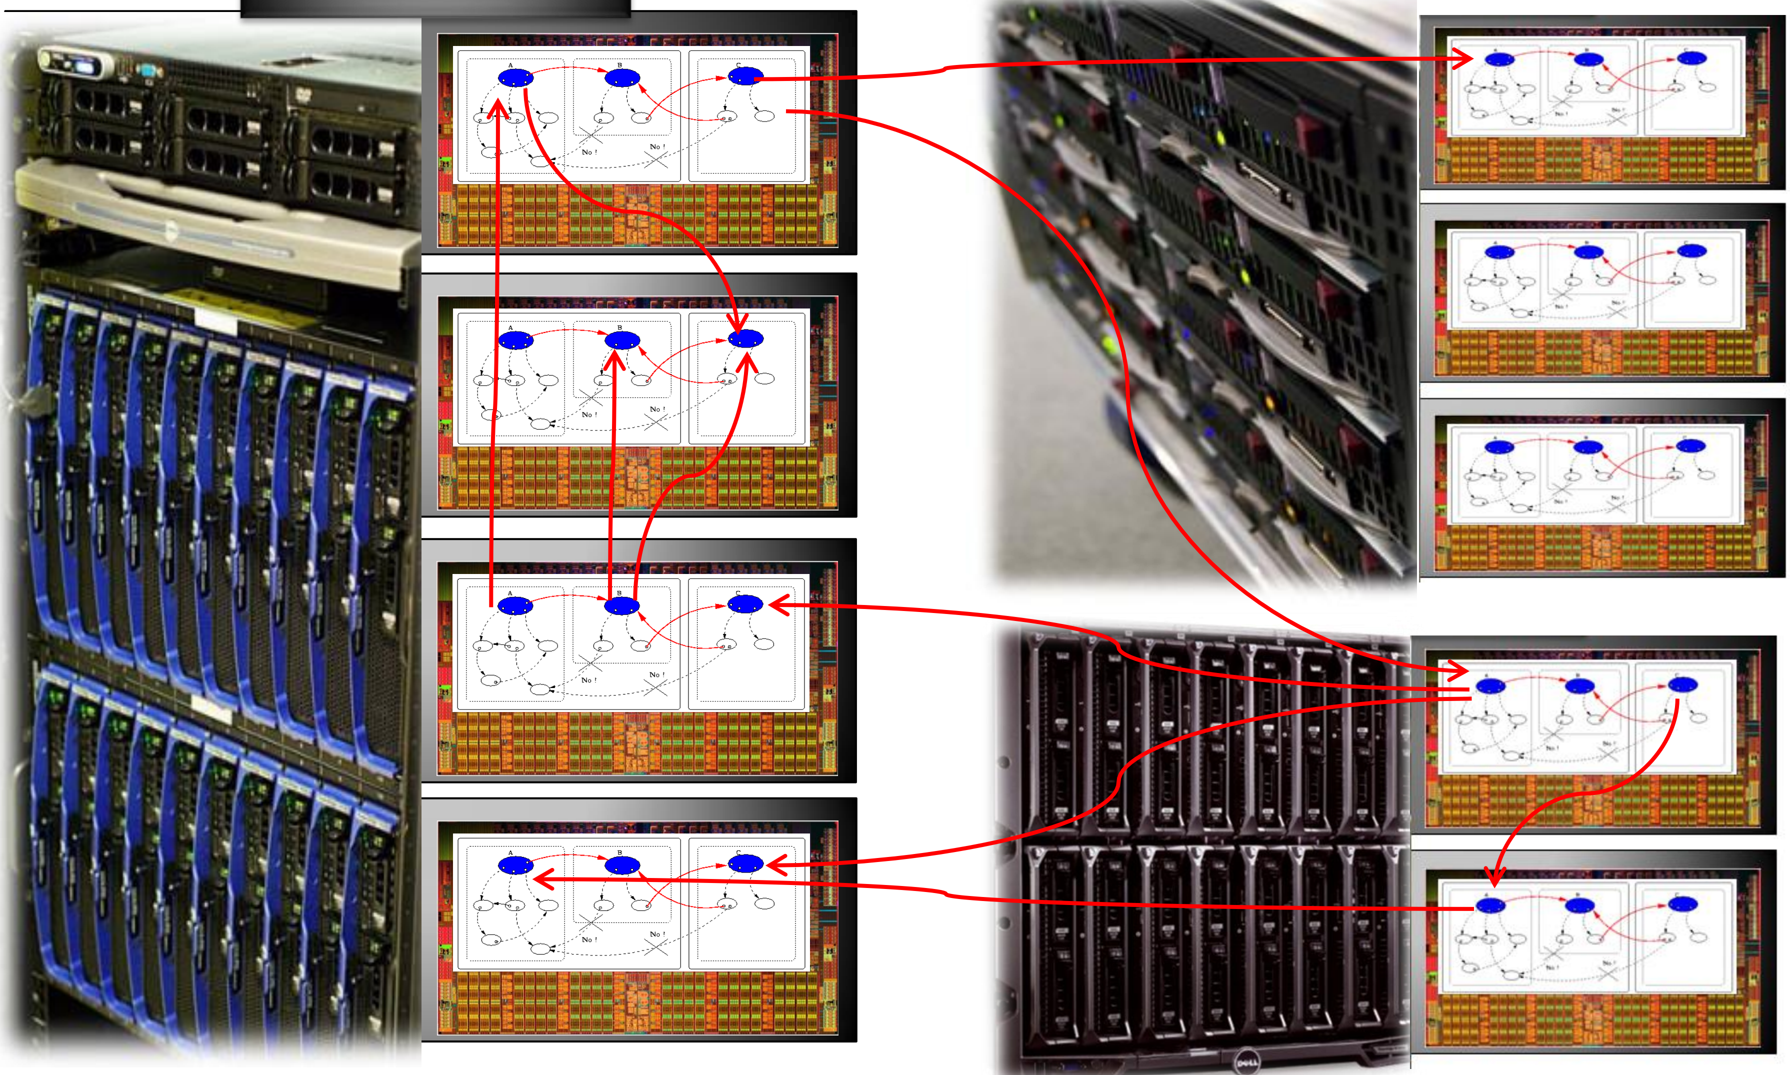

ProActive Parallel Suite

Broadcast and Scatter

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

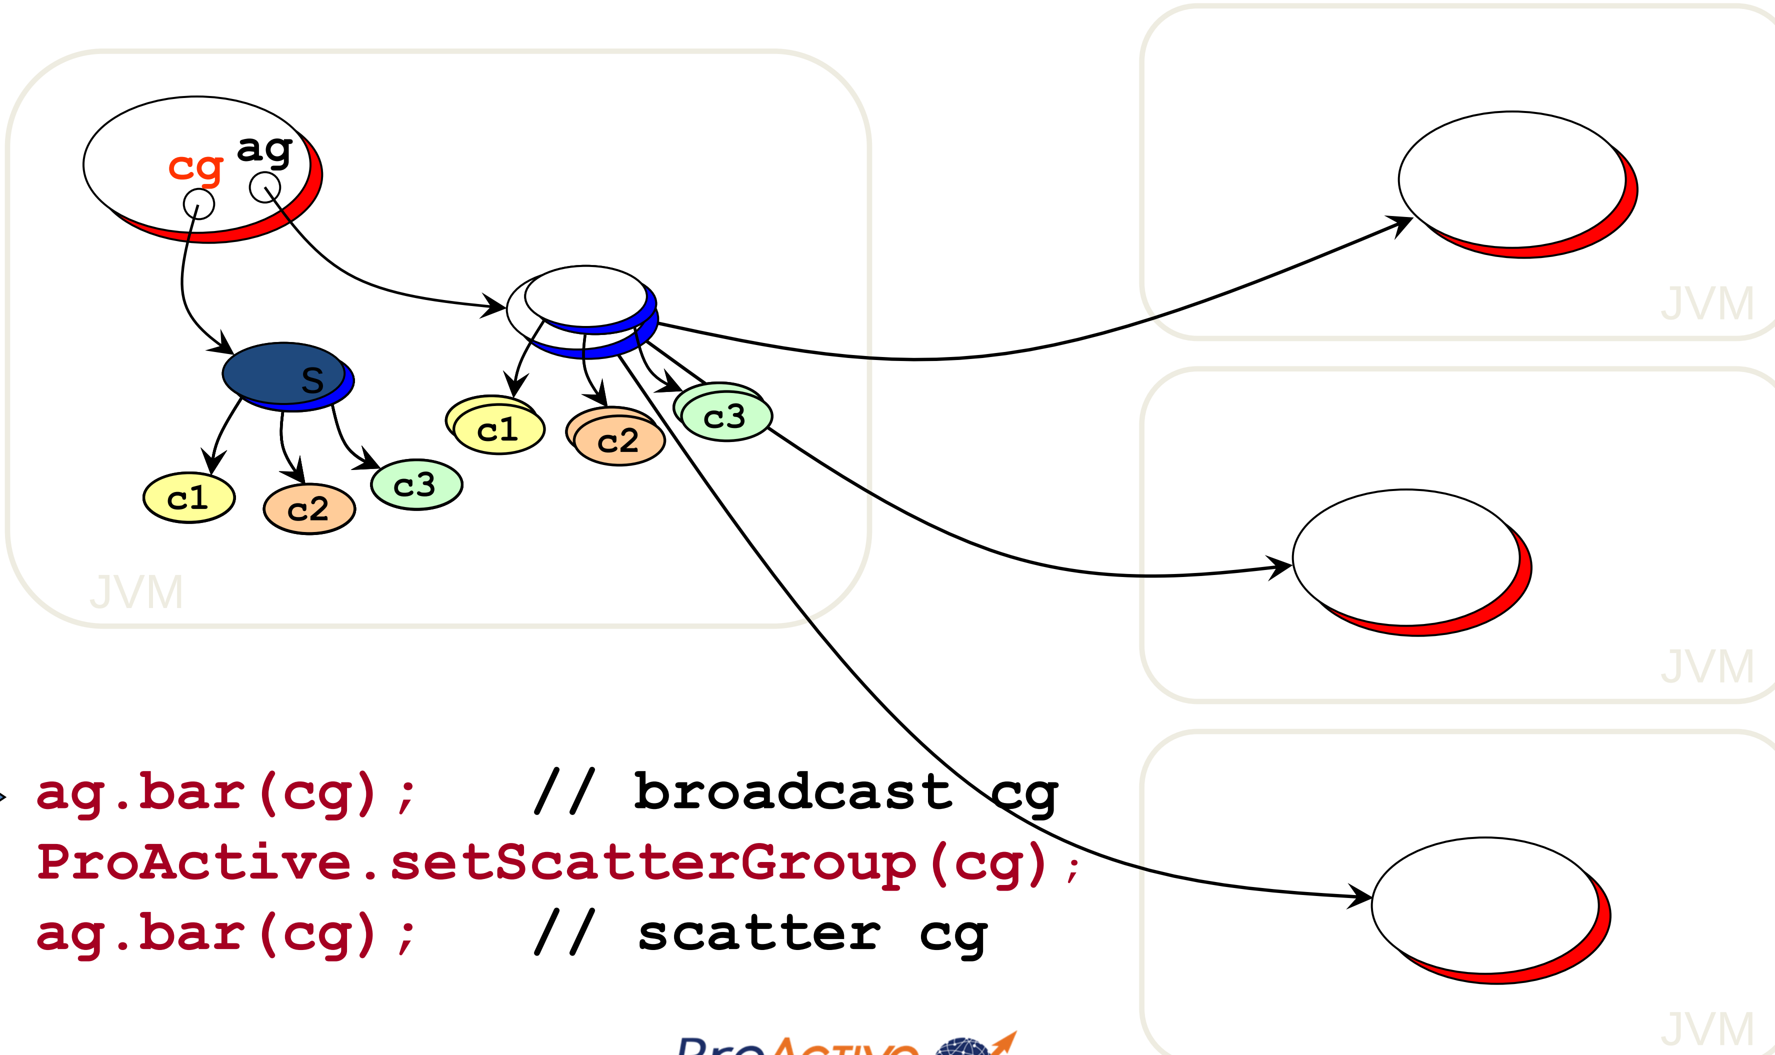


```
ag.bar(cg);    // broadcast cg
ProActive.setScatterGroup(cg);
ag.bar(cg);    // scatter cg
```

ProActive Parallel Suite

Dynamic Dispatch Group

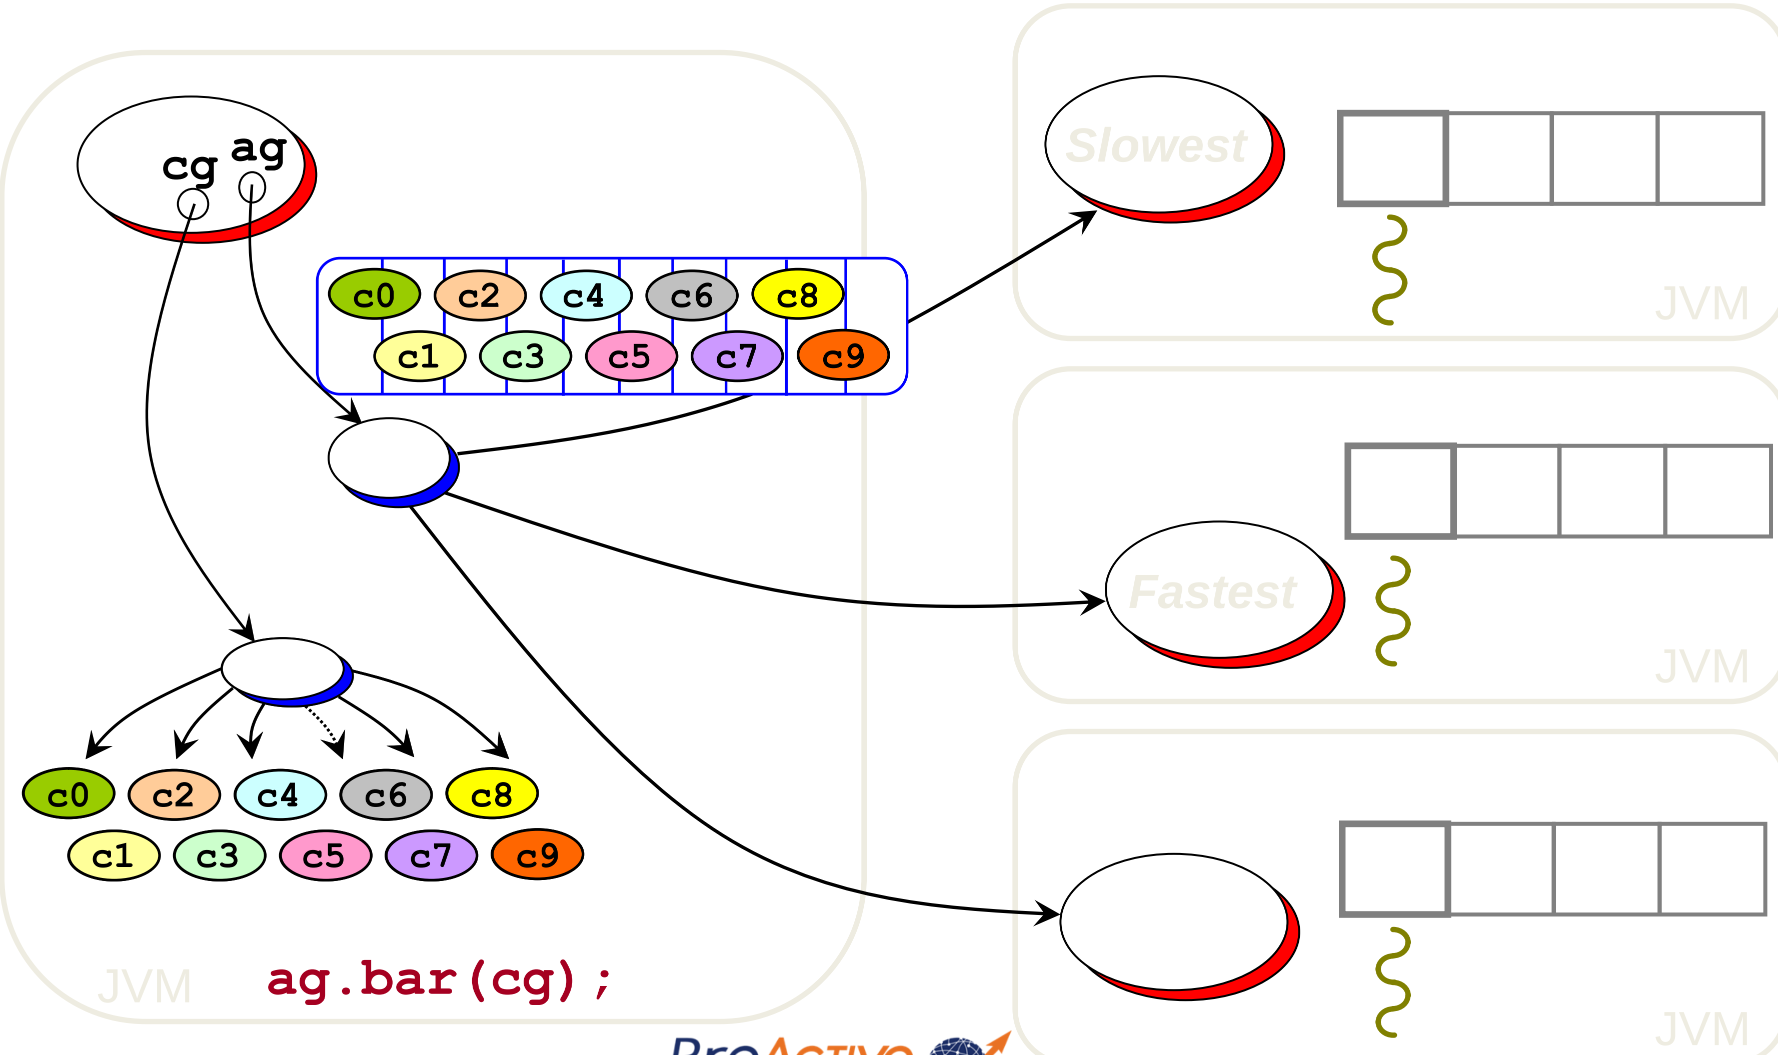

ProActive Parallel Suite

IC2D: Optimizing

IC2D

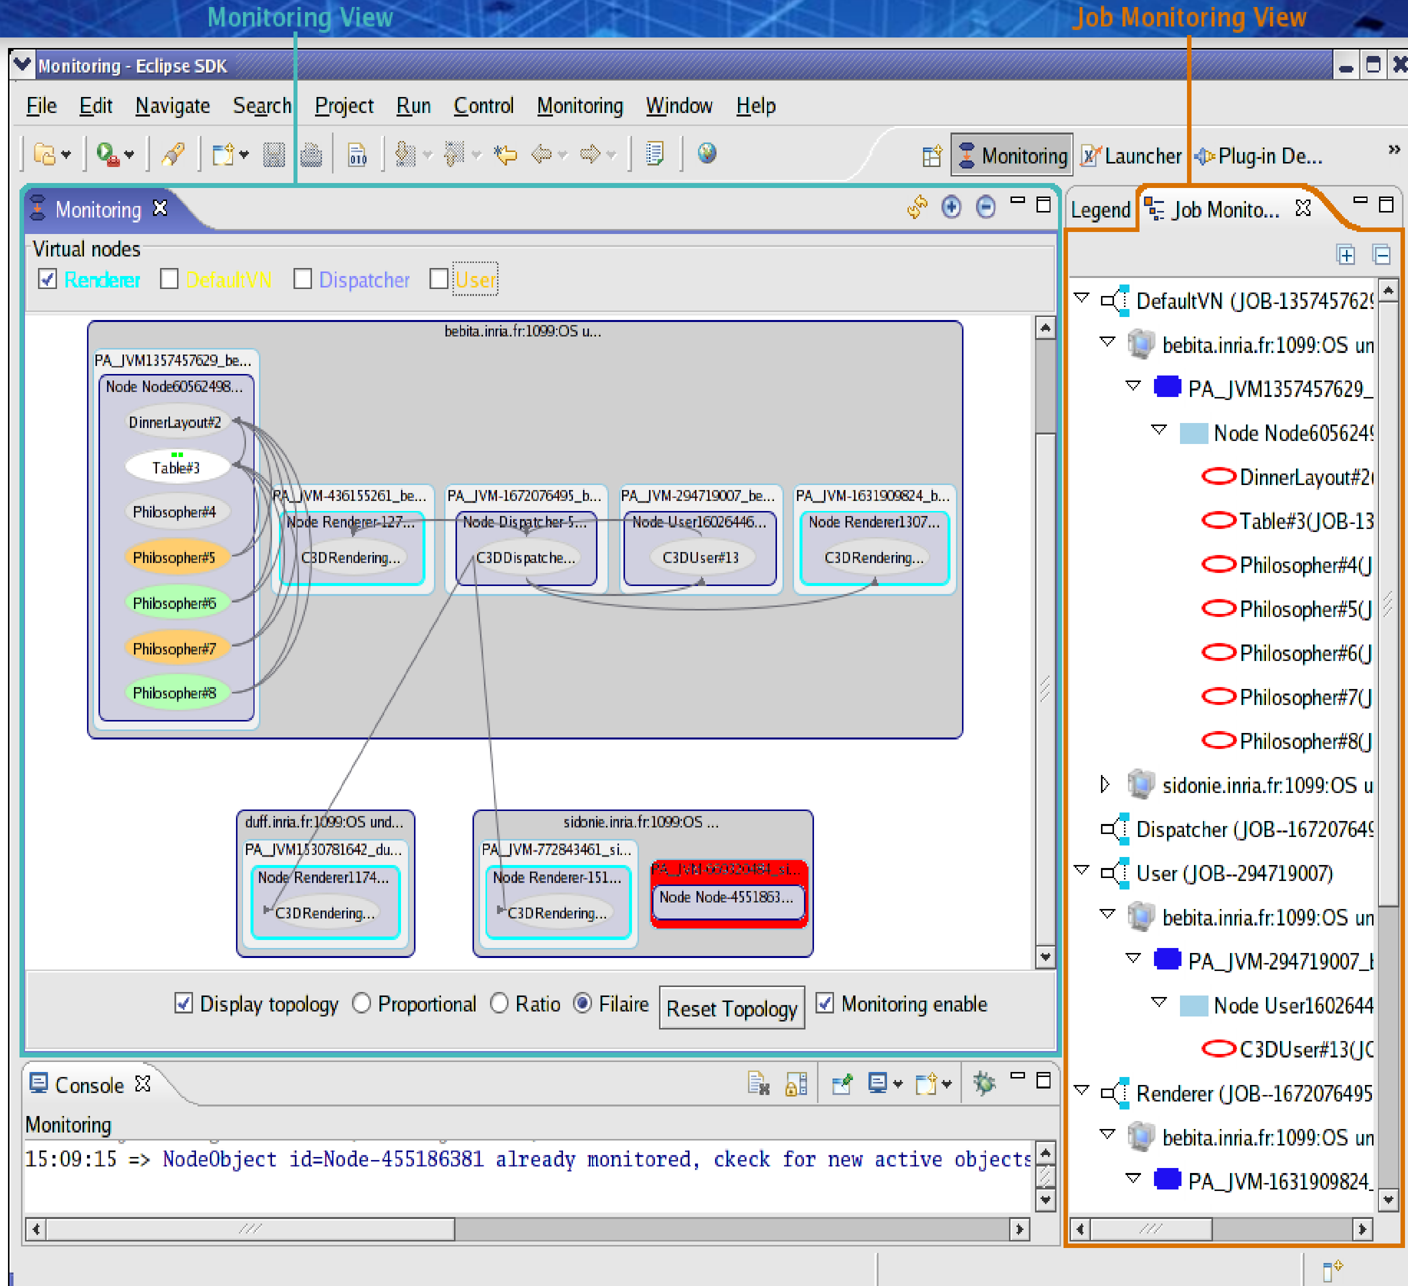

Monitoring View
Job Monitoring View
Monitoring - Eclipse SDK
File Edit Navigate Search Project Run Control Monitoring Window Help
Monitoring Launcher Plug-in De...
Monitoring
Virtual nodes
Renderer DefaultVN Dispatcher User
bebita.inria.fr:1099:OS u...
duff.inria.fr:1099:OS und...
sidonie.inria.fr:1099:OS ...
Display topology Proportional Ratio Filaire Reset Topology Monitoring enable
Console
Monitoring
15:09:15 => NodeObject id=Node-455186381 already monitored, ckeck for new active objects
Legend Job Monito...
DefaultVN (JOB-1357457629
bebita.inria.fr:1099:OS un
sidonie.inria.fr:1099:OS u
Dispatcher (JOB--167207649
User (JOB--294719007)
Renderer (JOB--1672076495

IC2D

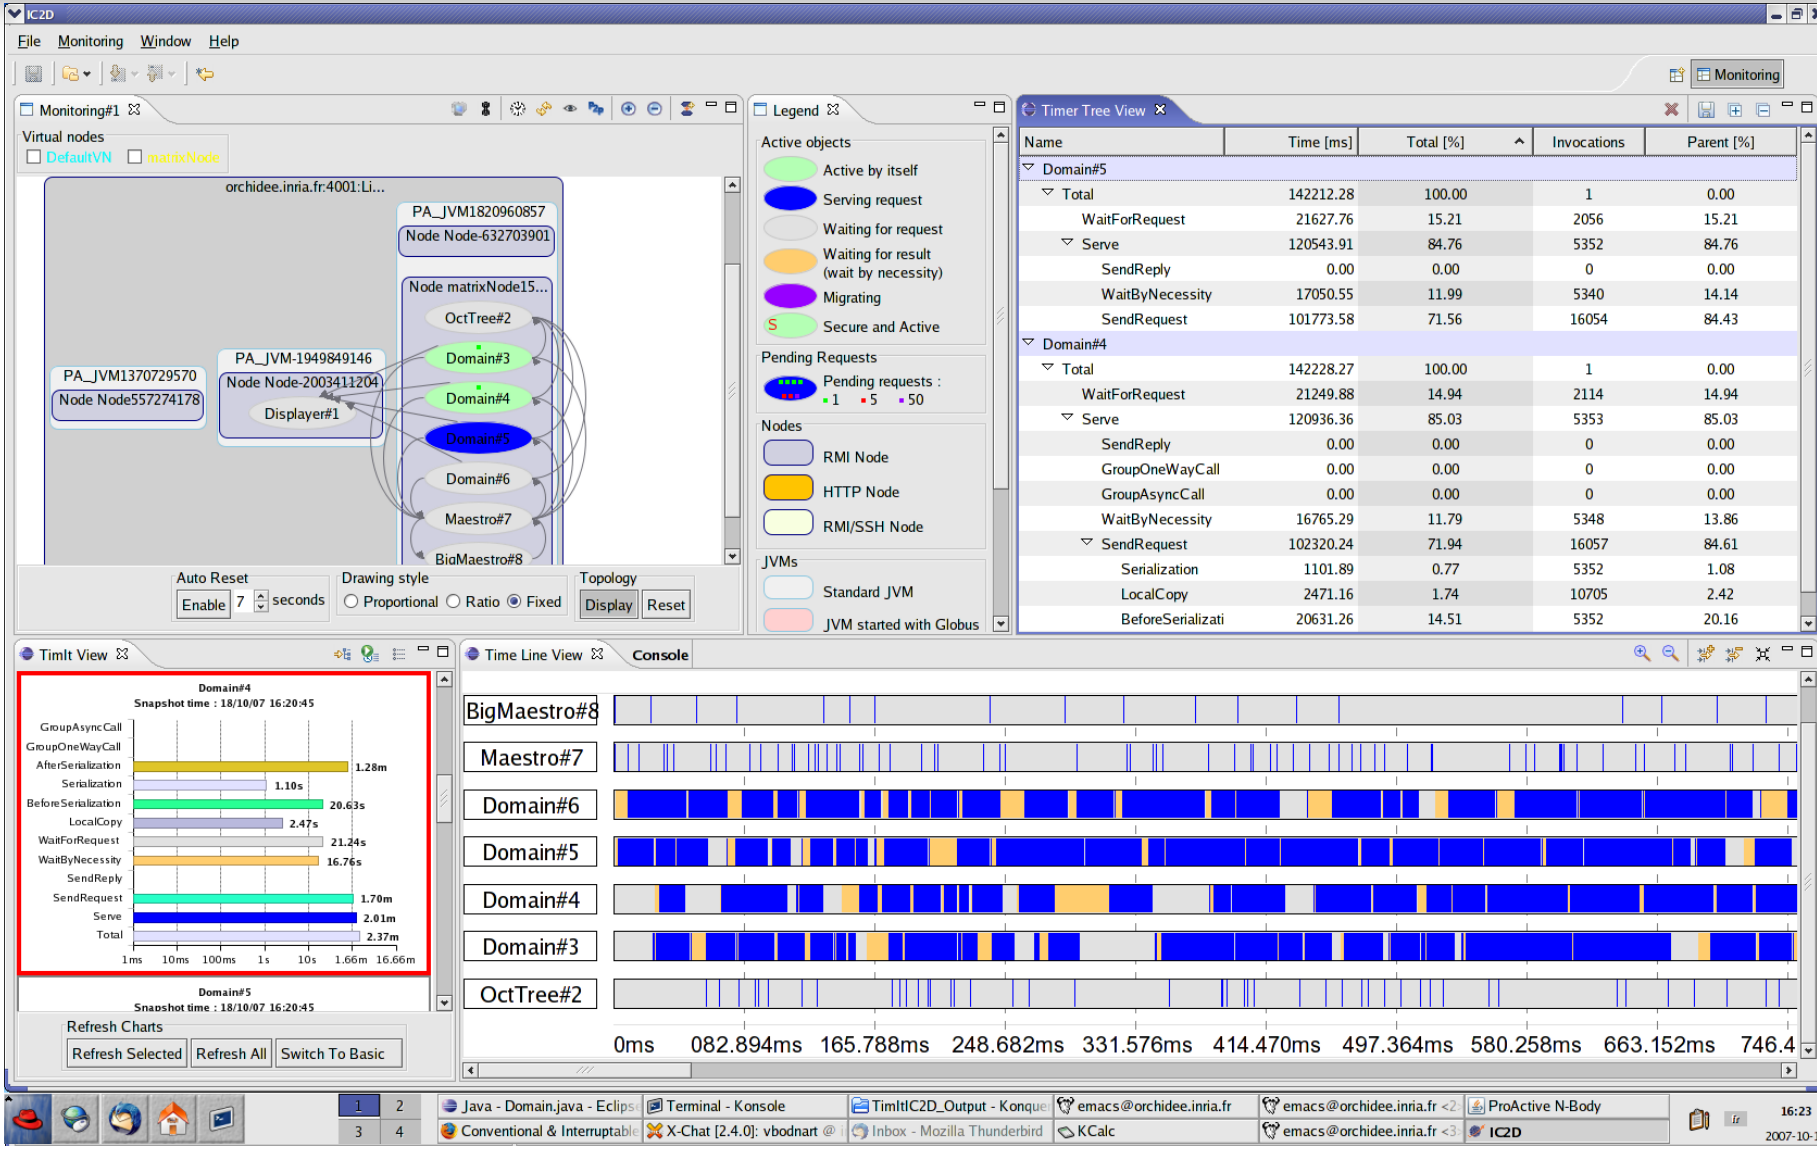

ProActive Parallel Suite

Video 1:
IC2D Optimizing
Monitoring, Debugging, Optimizing

ProActive
Parallel Suite

Scheduling & Resourcing

ProActive Scheduling

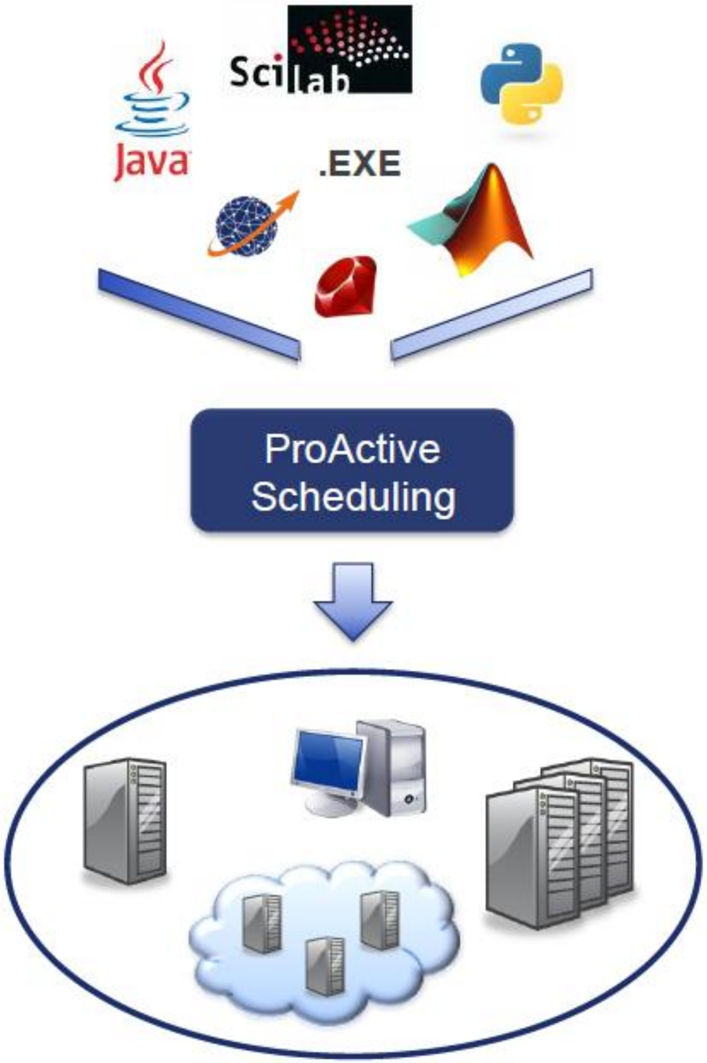

ProActive
Parallel Suite

ProActive Scheduling Big Picture

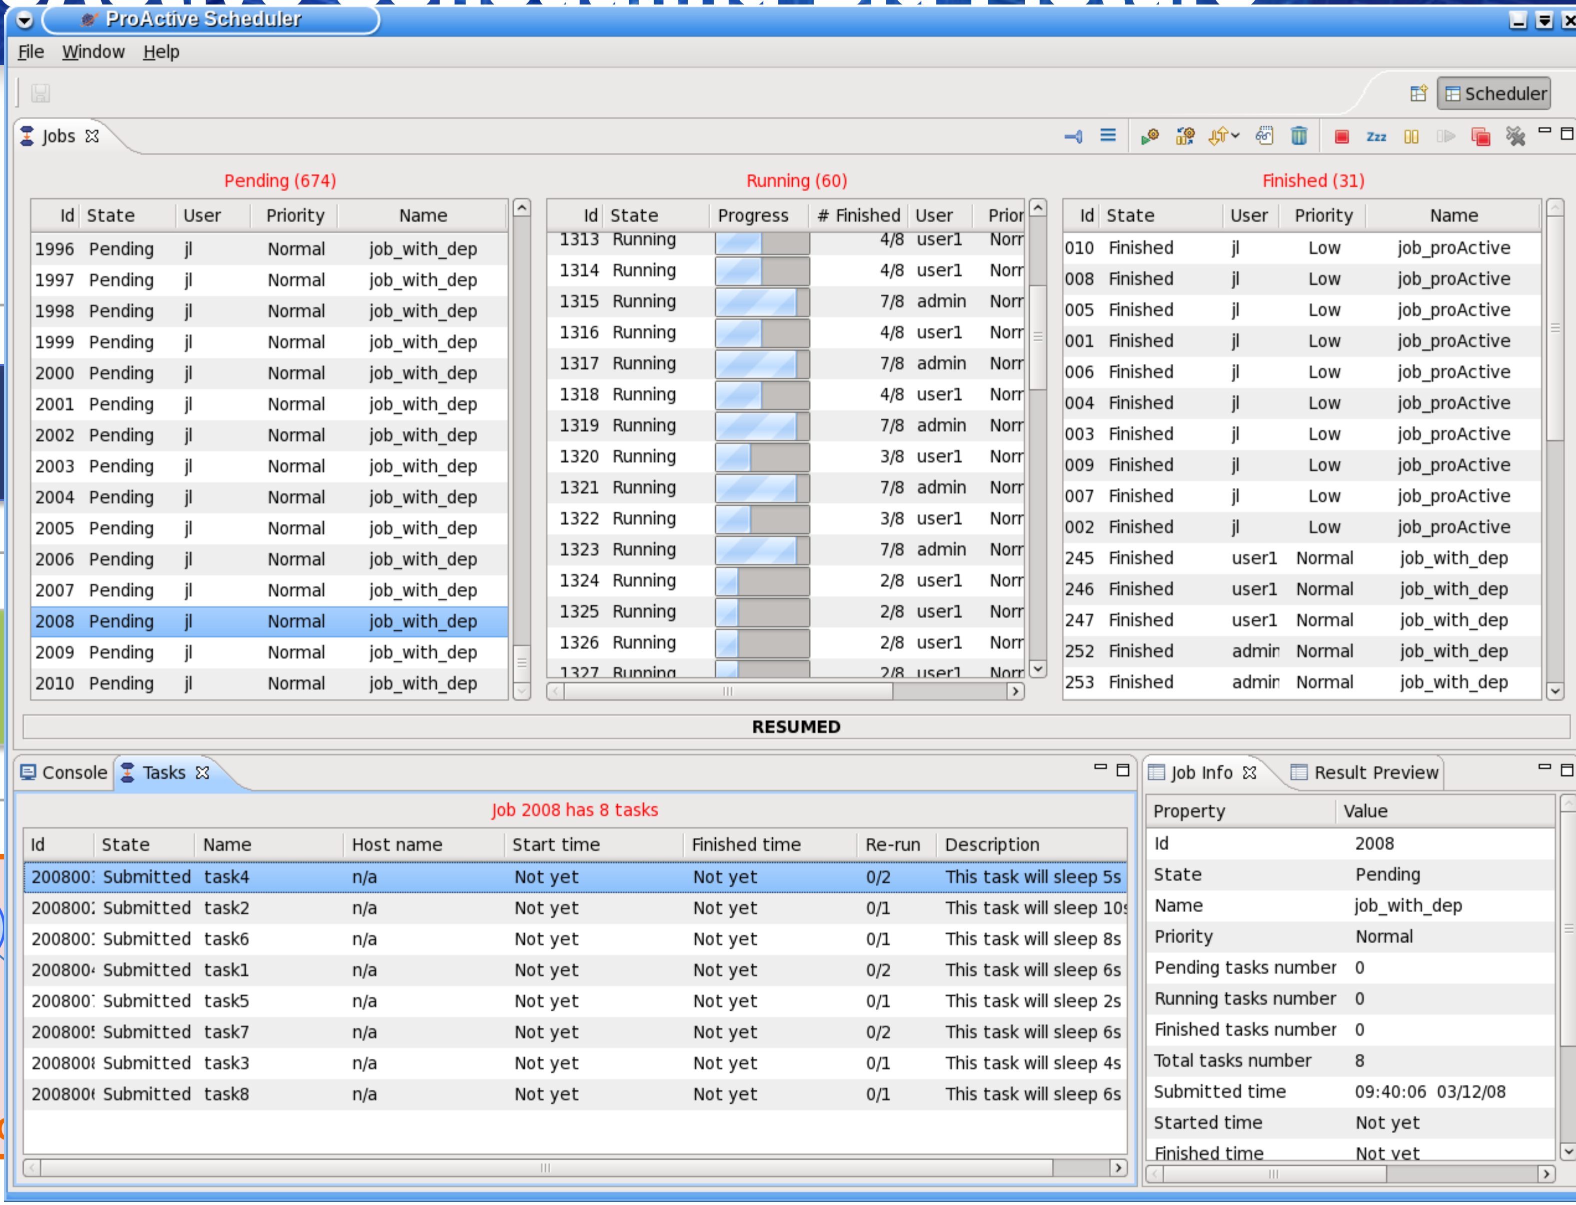

d

ce

RESC

Workflow Example : Picture Denoising

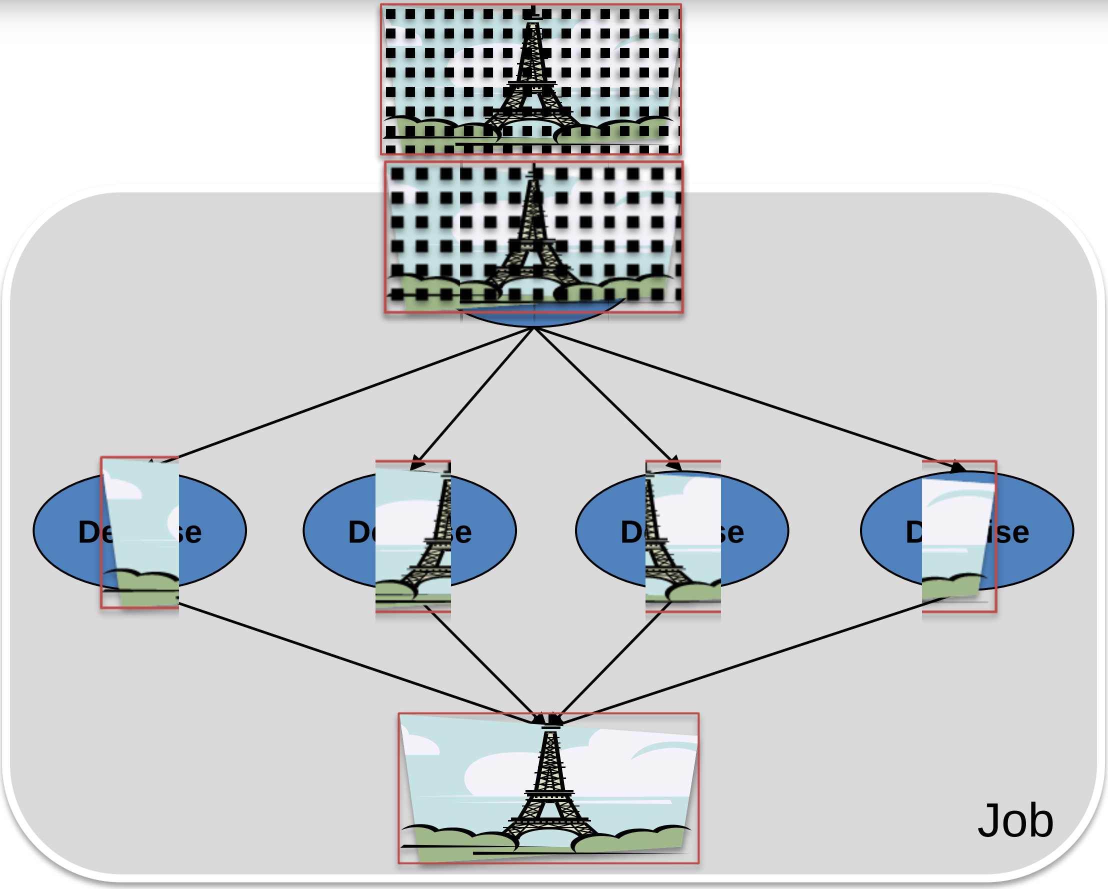

•with selection on native executable availability (ImageMagik, GREYstoration)
- Multi-platform selection and command generation

•with file transfer in pre/post scripts

ProActive Resourcing

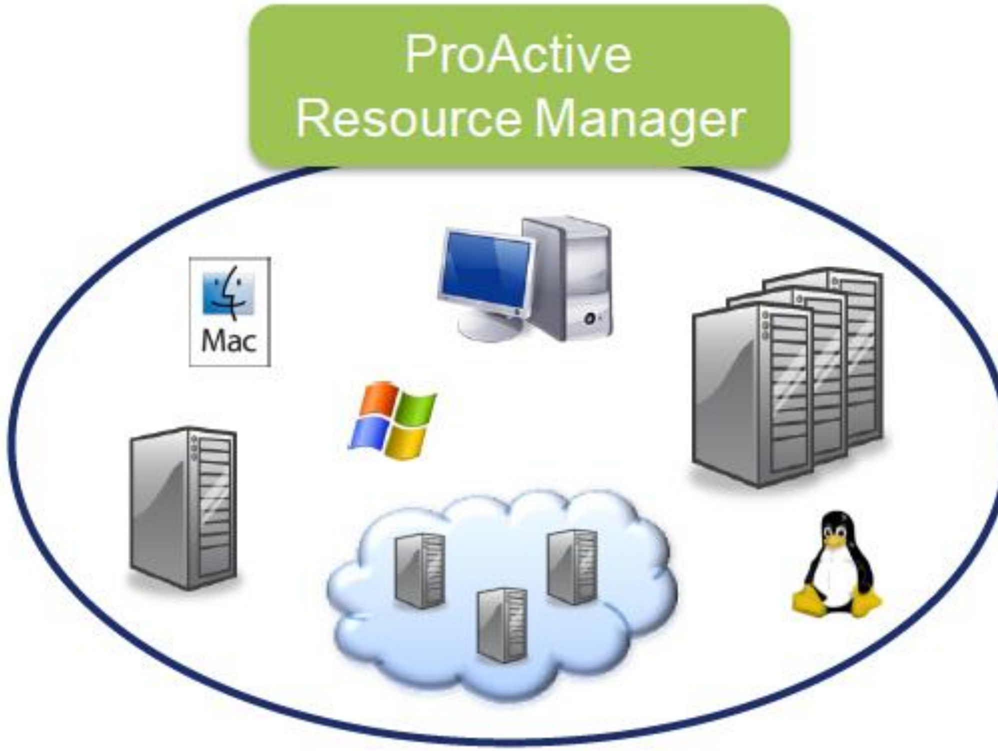

44

RESOURCING User Interface

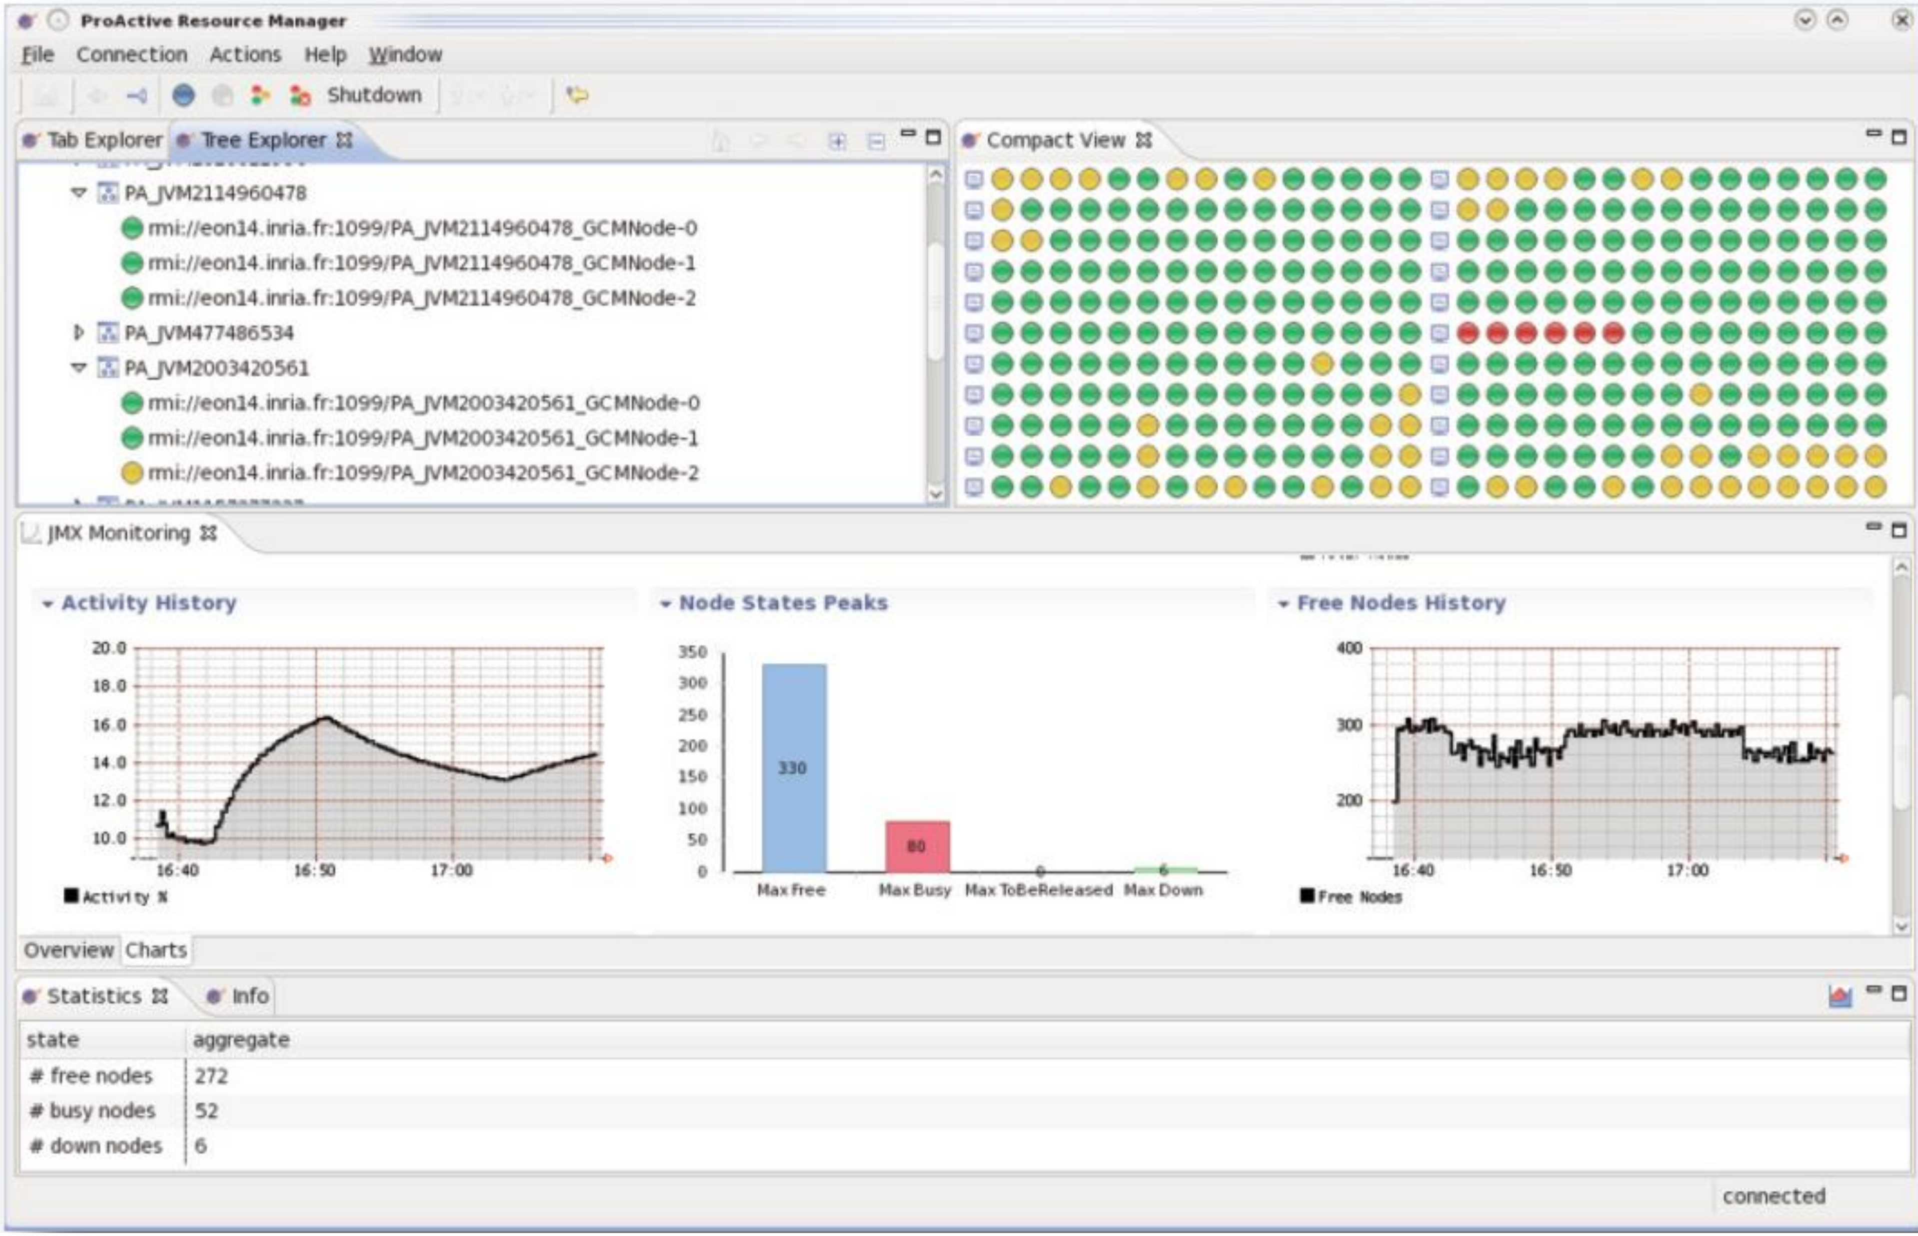

Clusters to Grids to Clouds: e.g. on Amazon EC2

ProActive Parallel Suite

Private, Public & Hybrid Clouds

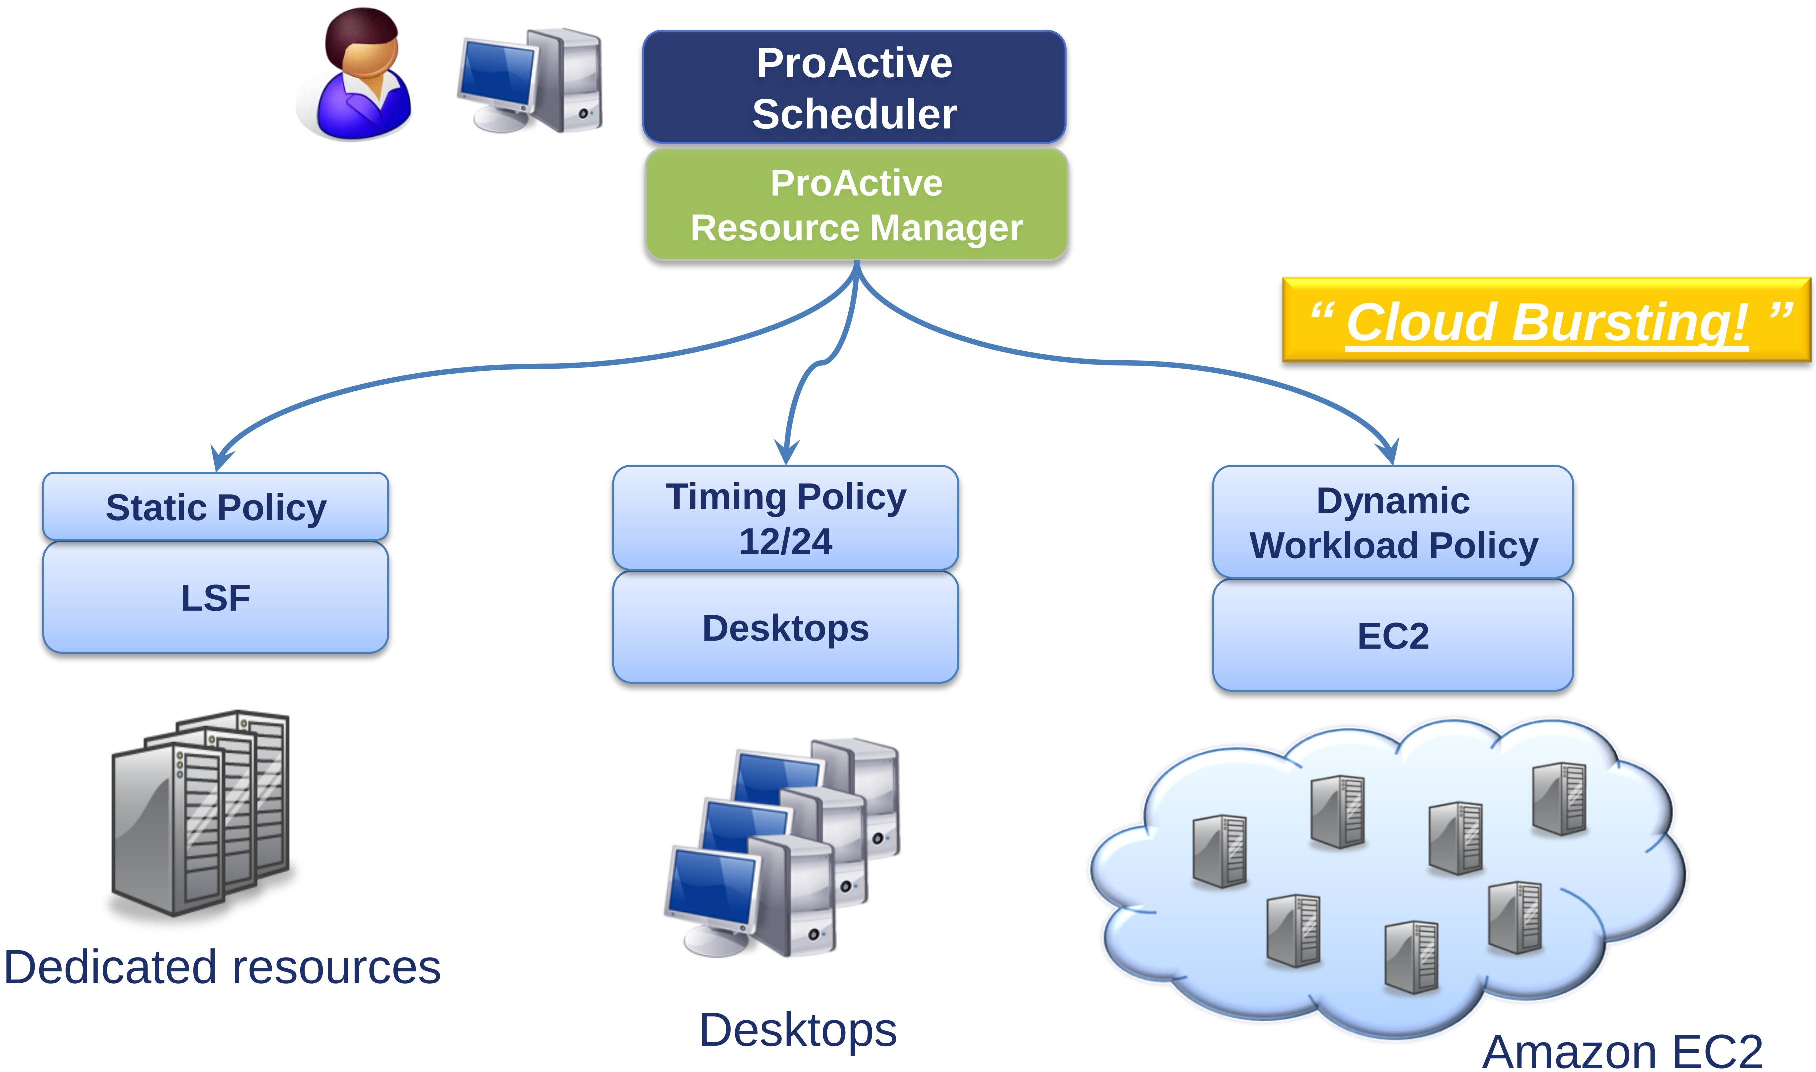

ProActive Parallel Suite

Cloud Seeding with ProActive

- Amazon EC2 Execution
- *Cloud Seeding* strategy to mix heterogeneous computing resources :
  - External GPU resources

" Cloud Seeding "

ProActive Parallel Suite

Cloud Seeding with ProActive

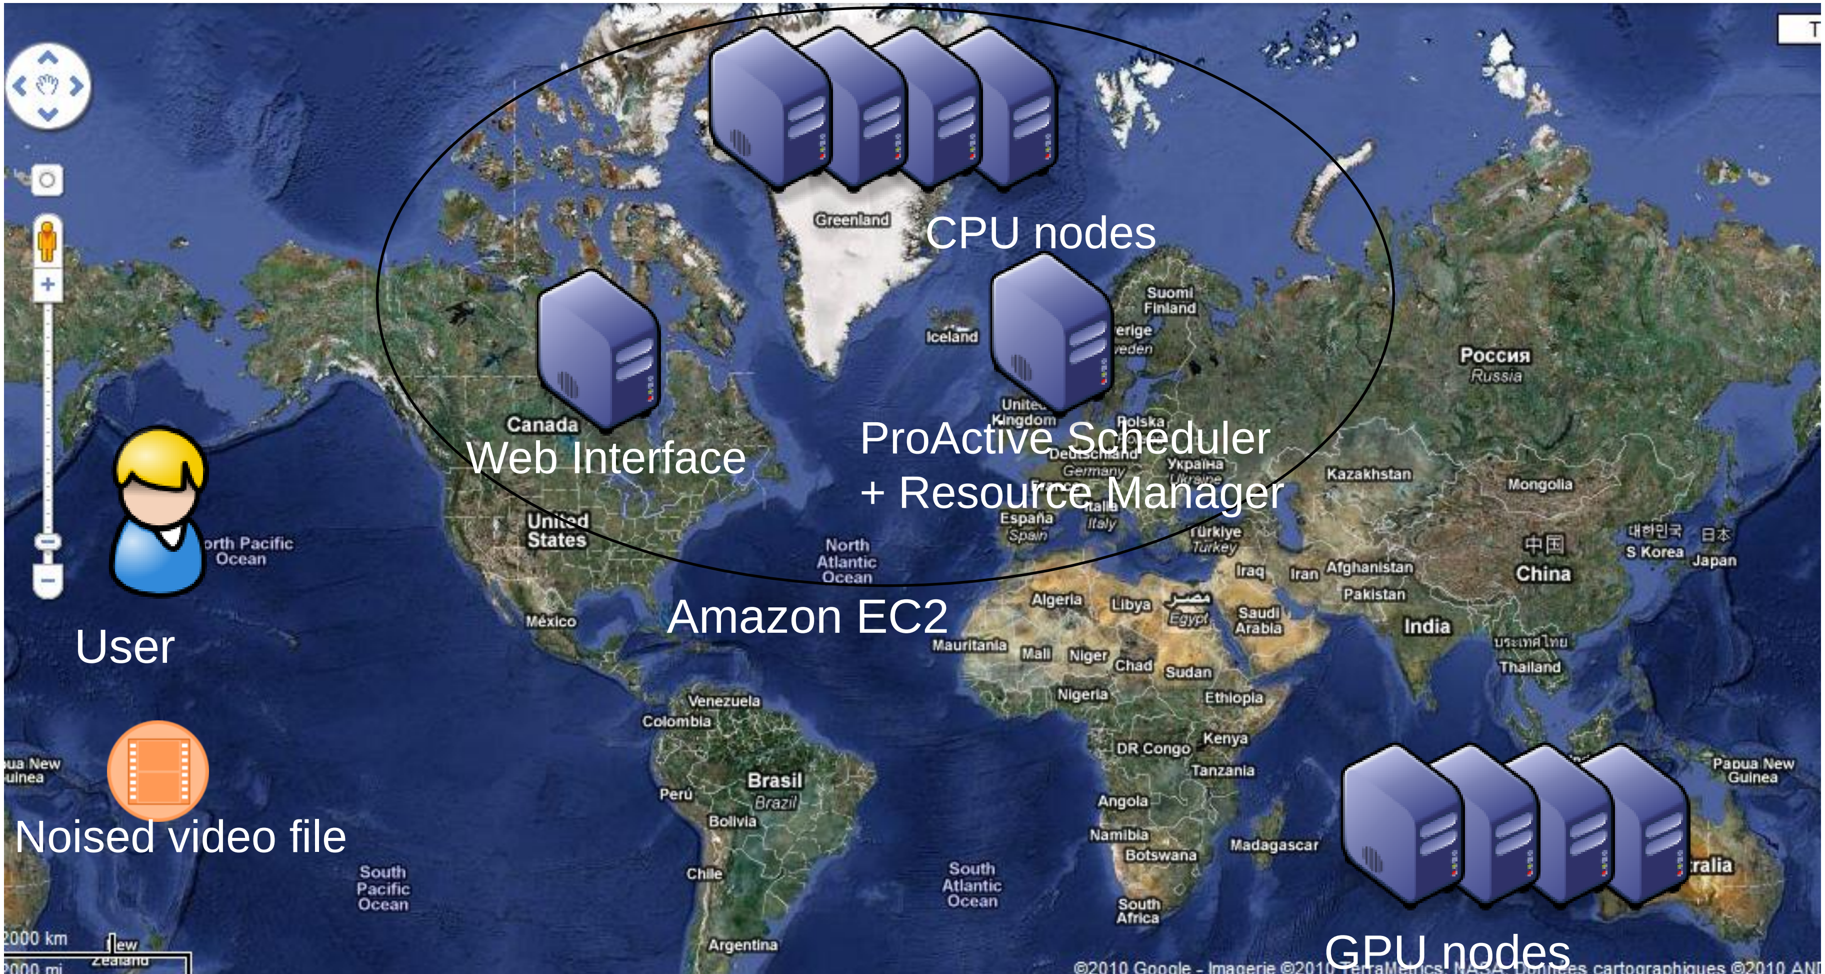

ProActive Parallel Suite

Cloud Seeding with ProActive

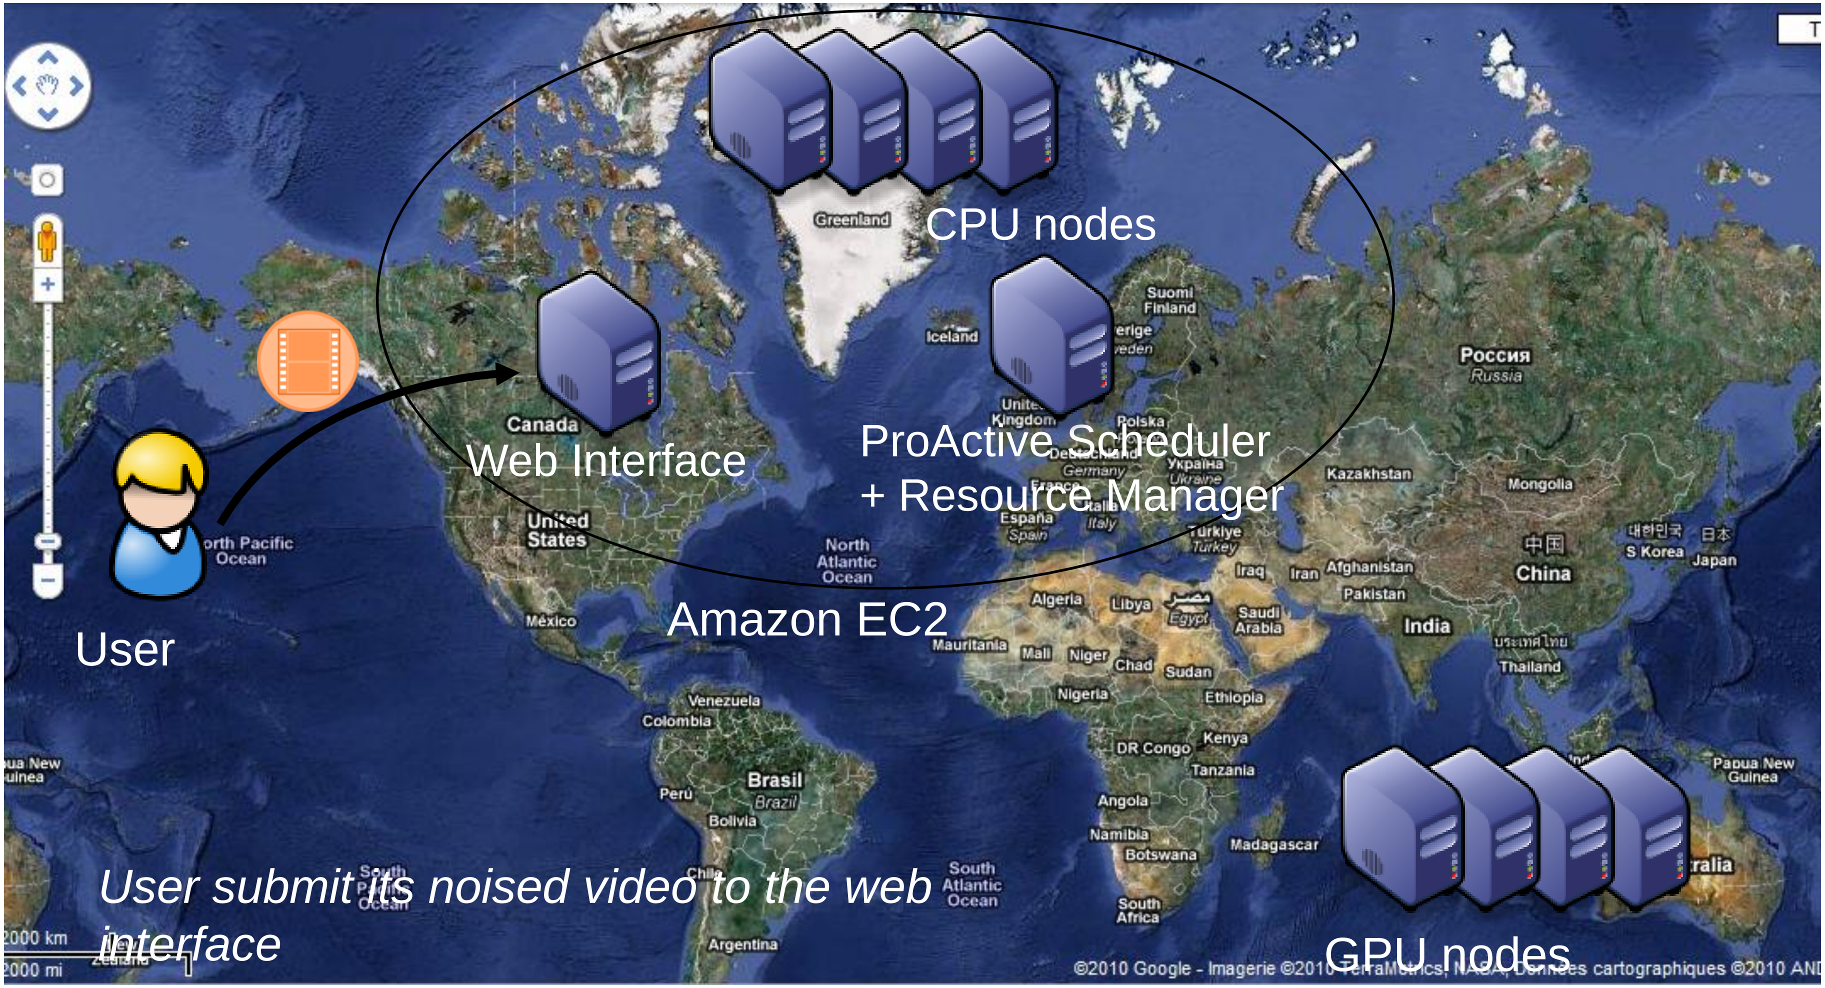

Cloud Seeding with ProActive

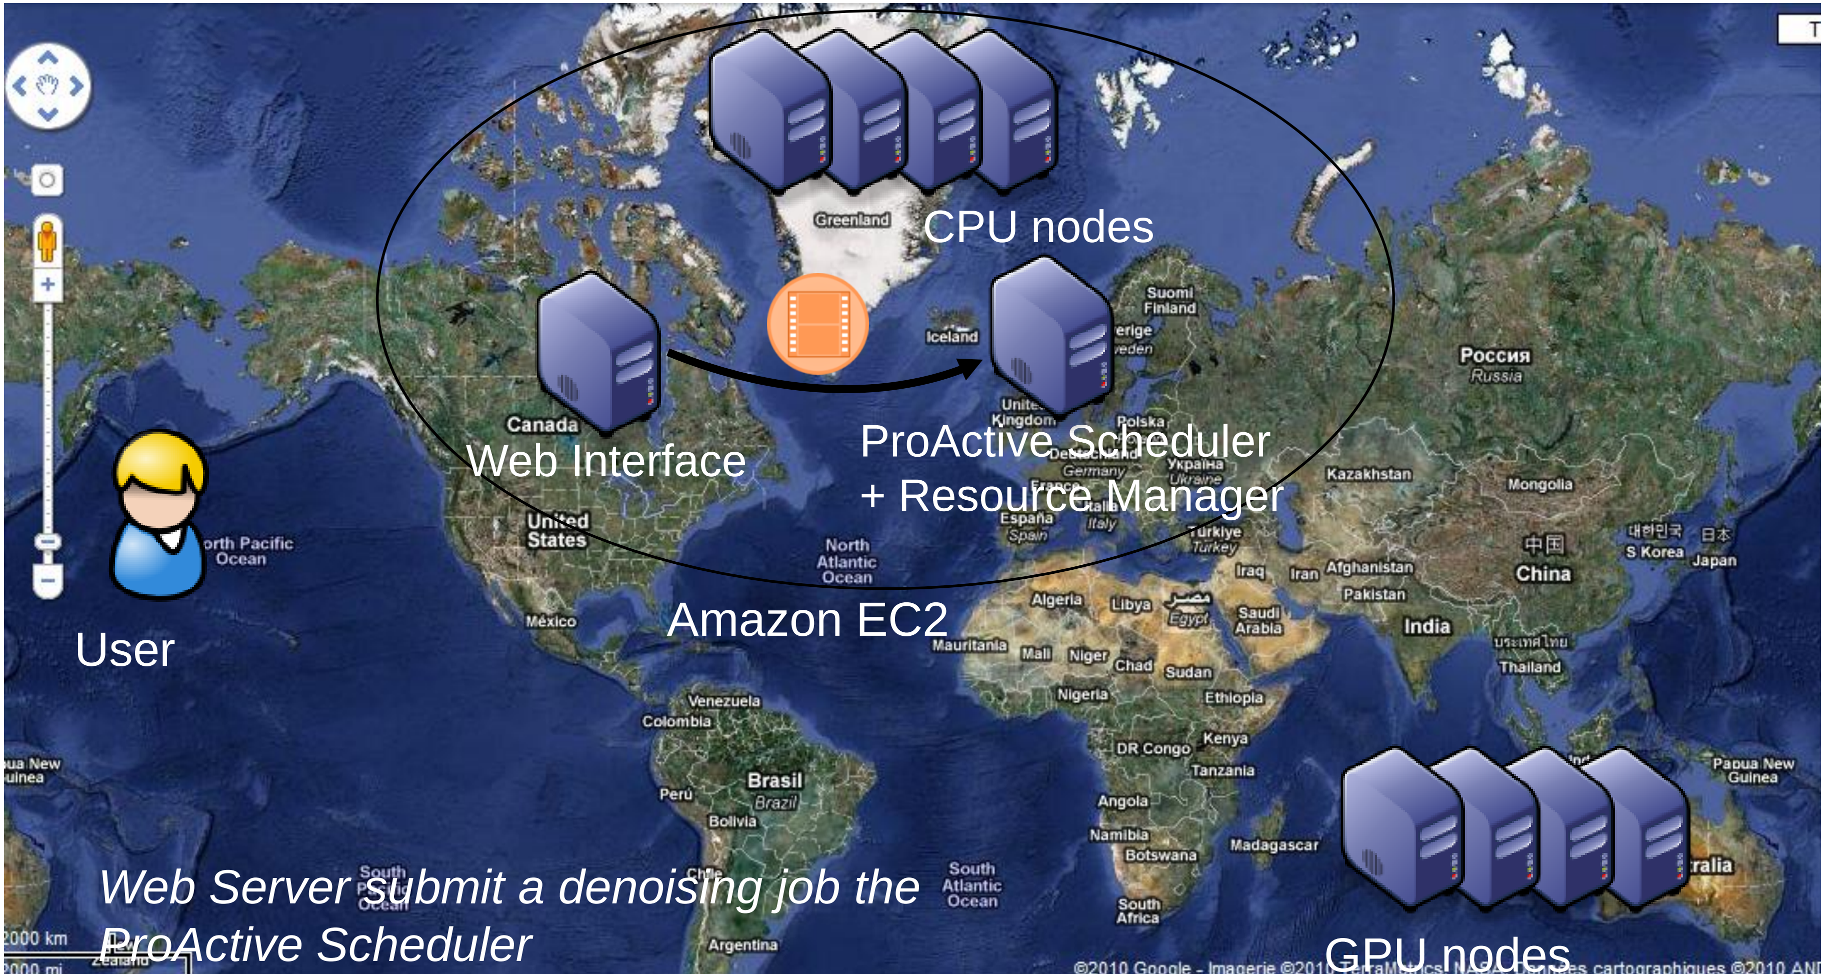

Web Server submit a denoising job the ProActive Scheduler

ProActive Parallel Suite

Cloud Seeding with ProActive

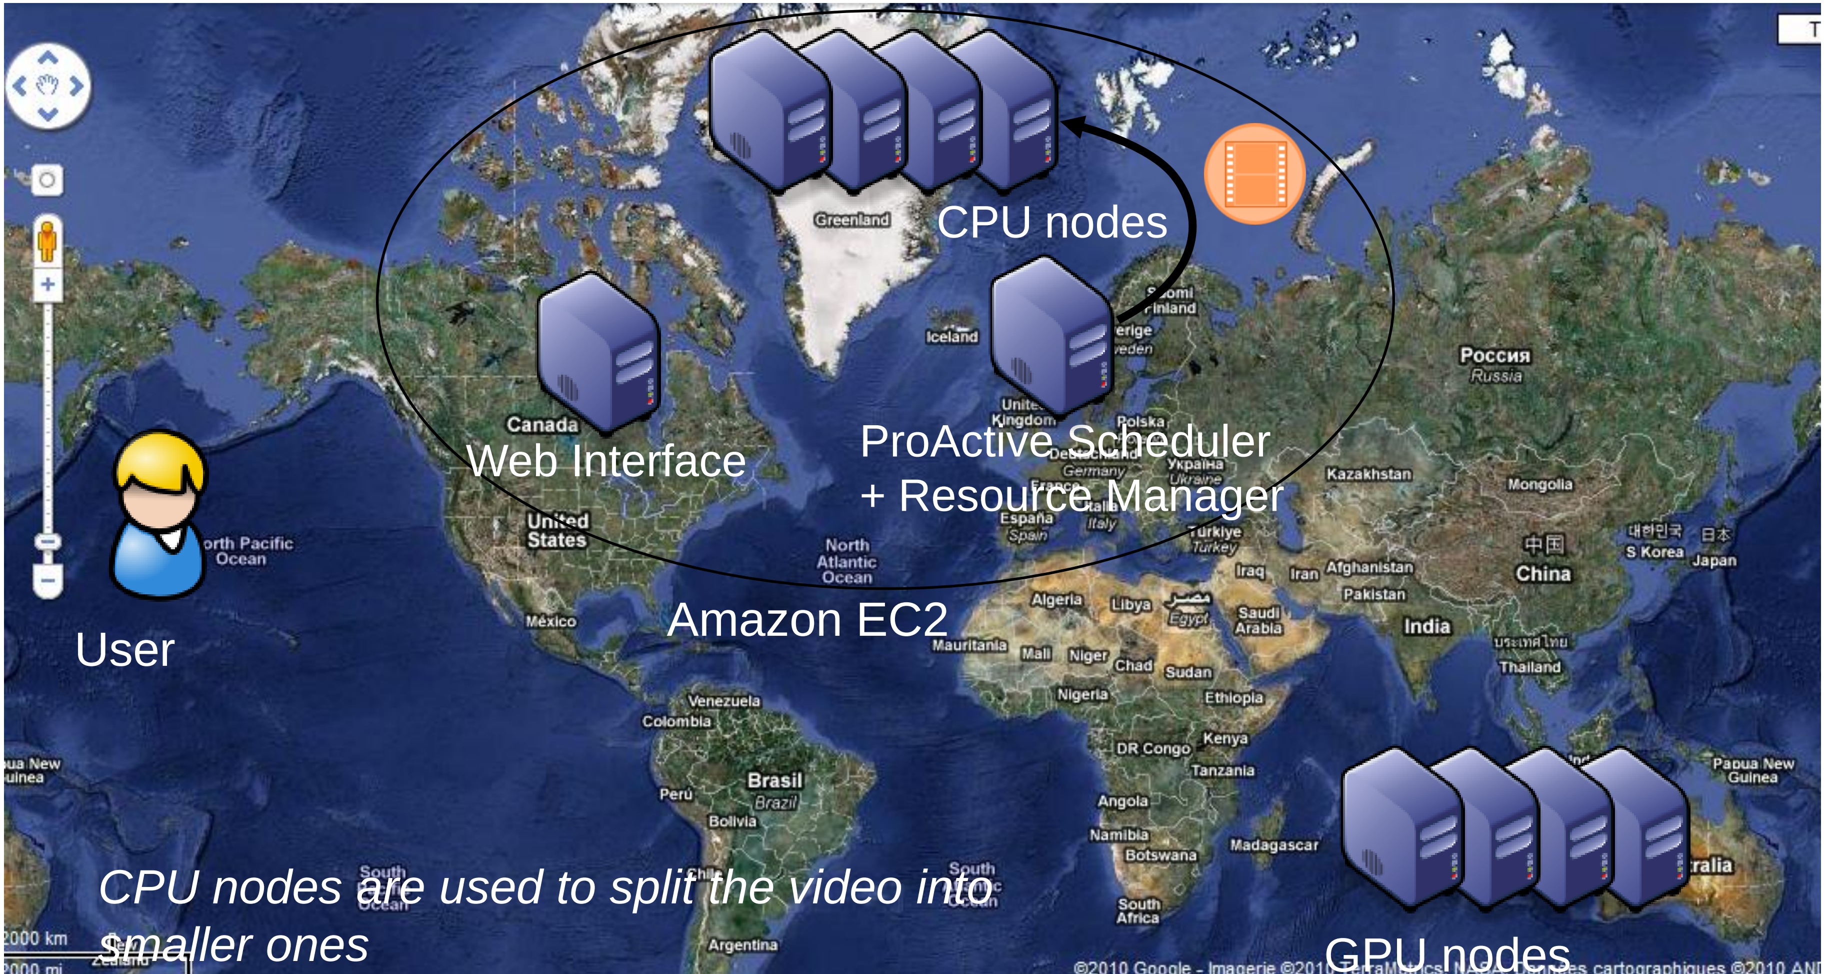

ProActive Parallel Suite

Cloud Seeding with ProActive

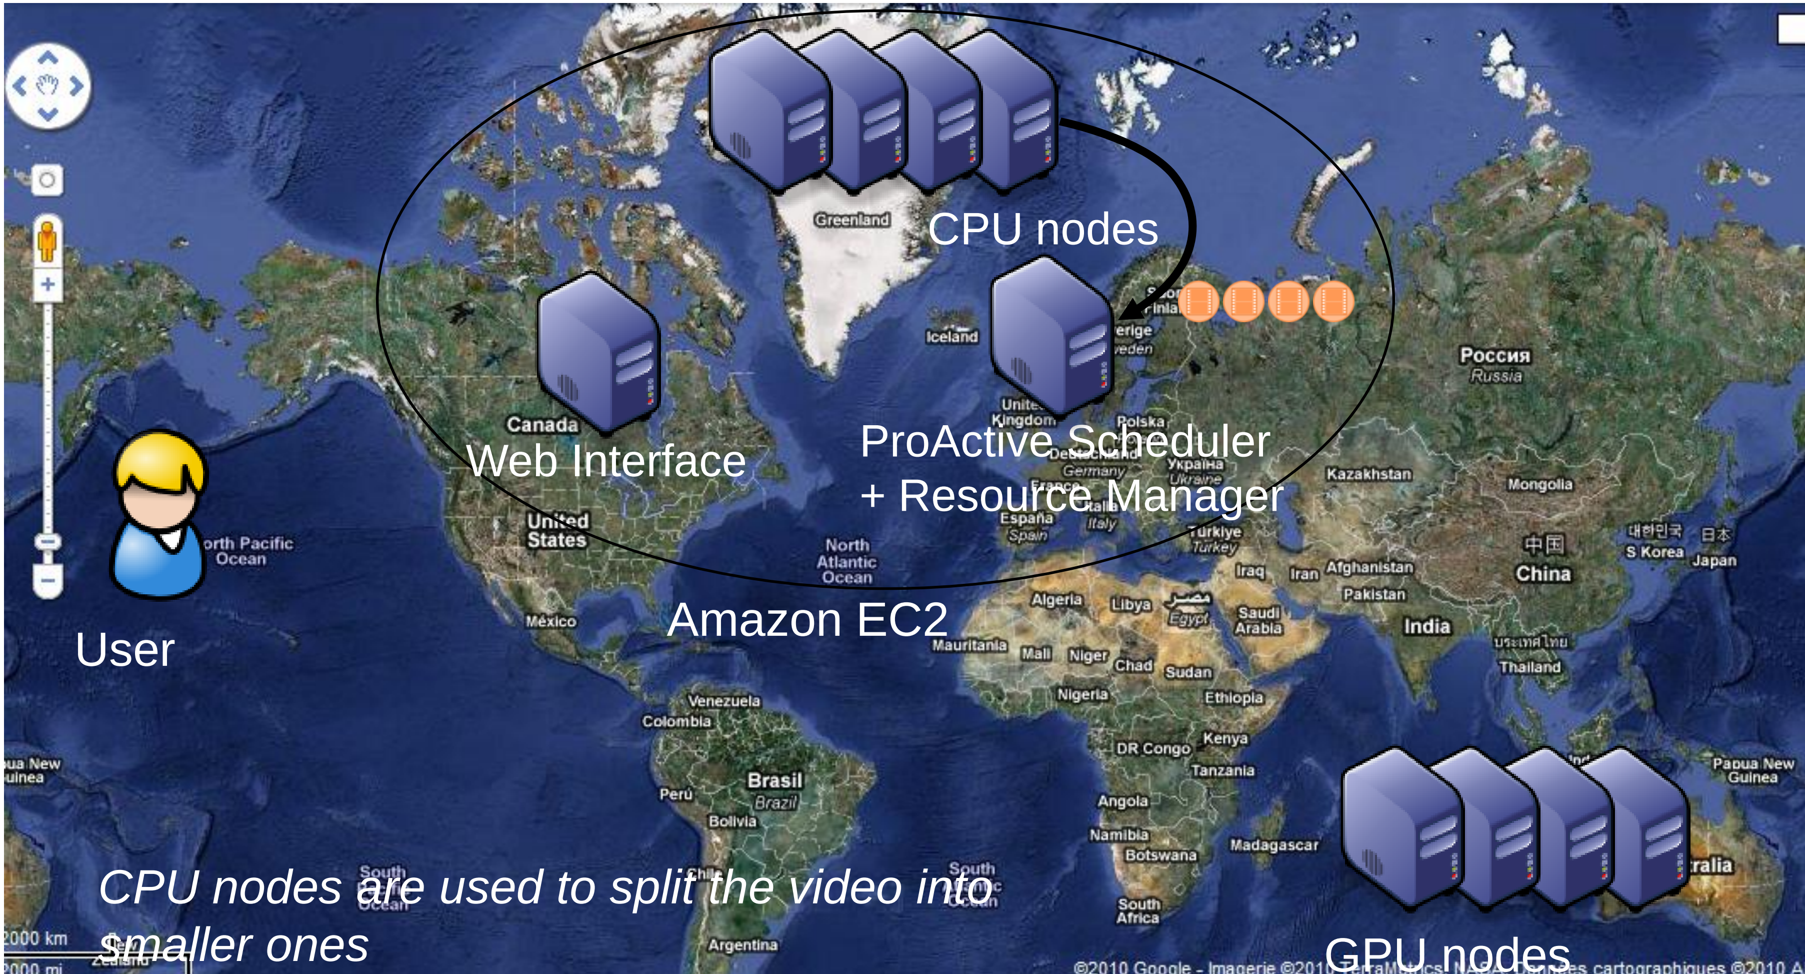

CPU nodes are used to split the video into smaller ones

Cloud Seeding with ProActive

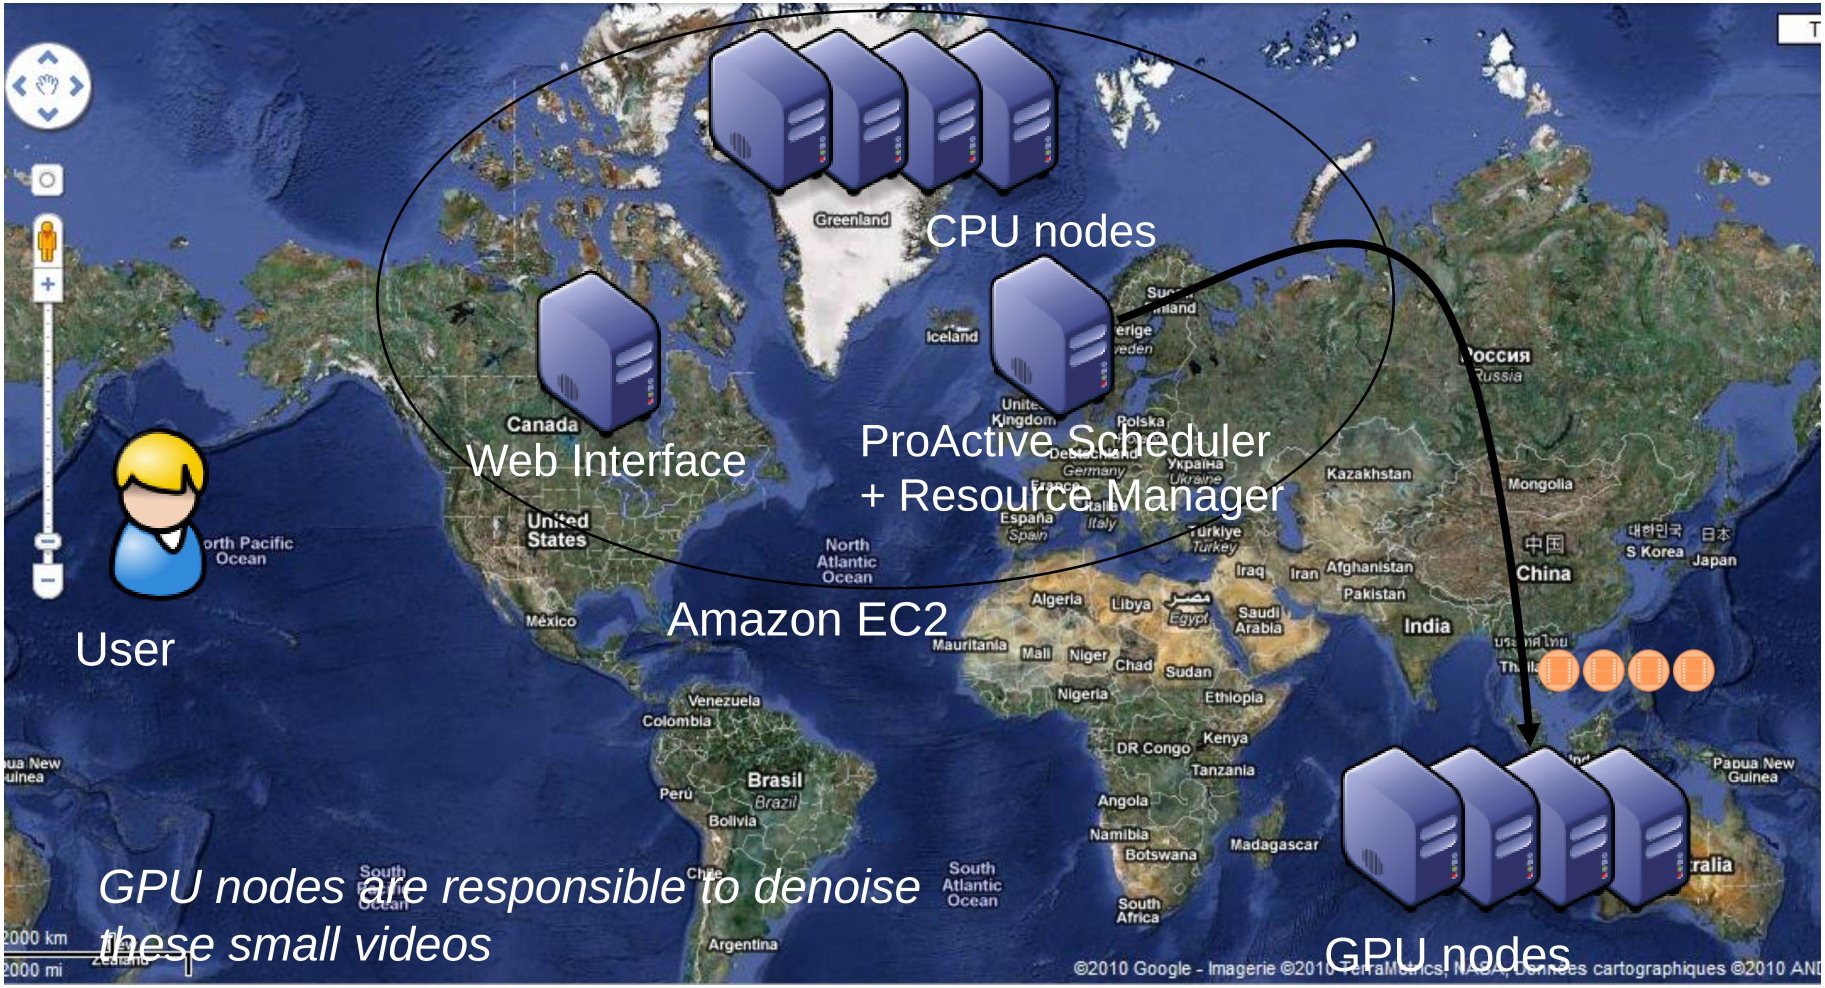

ProActive Parallel Suite

Cloud Seeding with ProActive

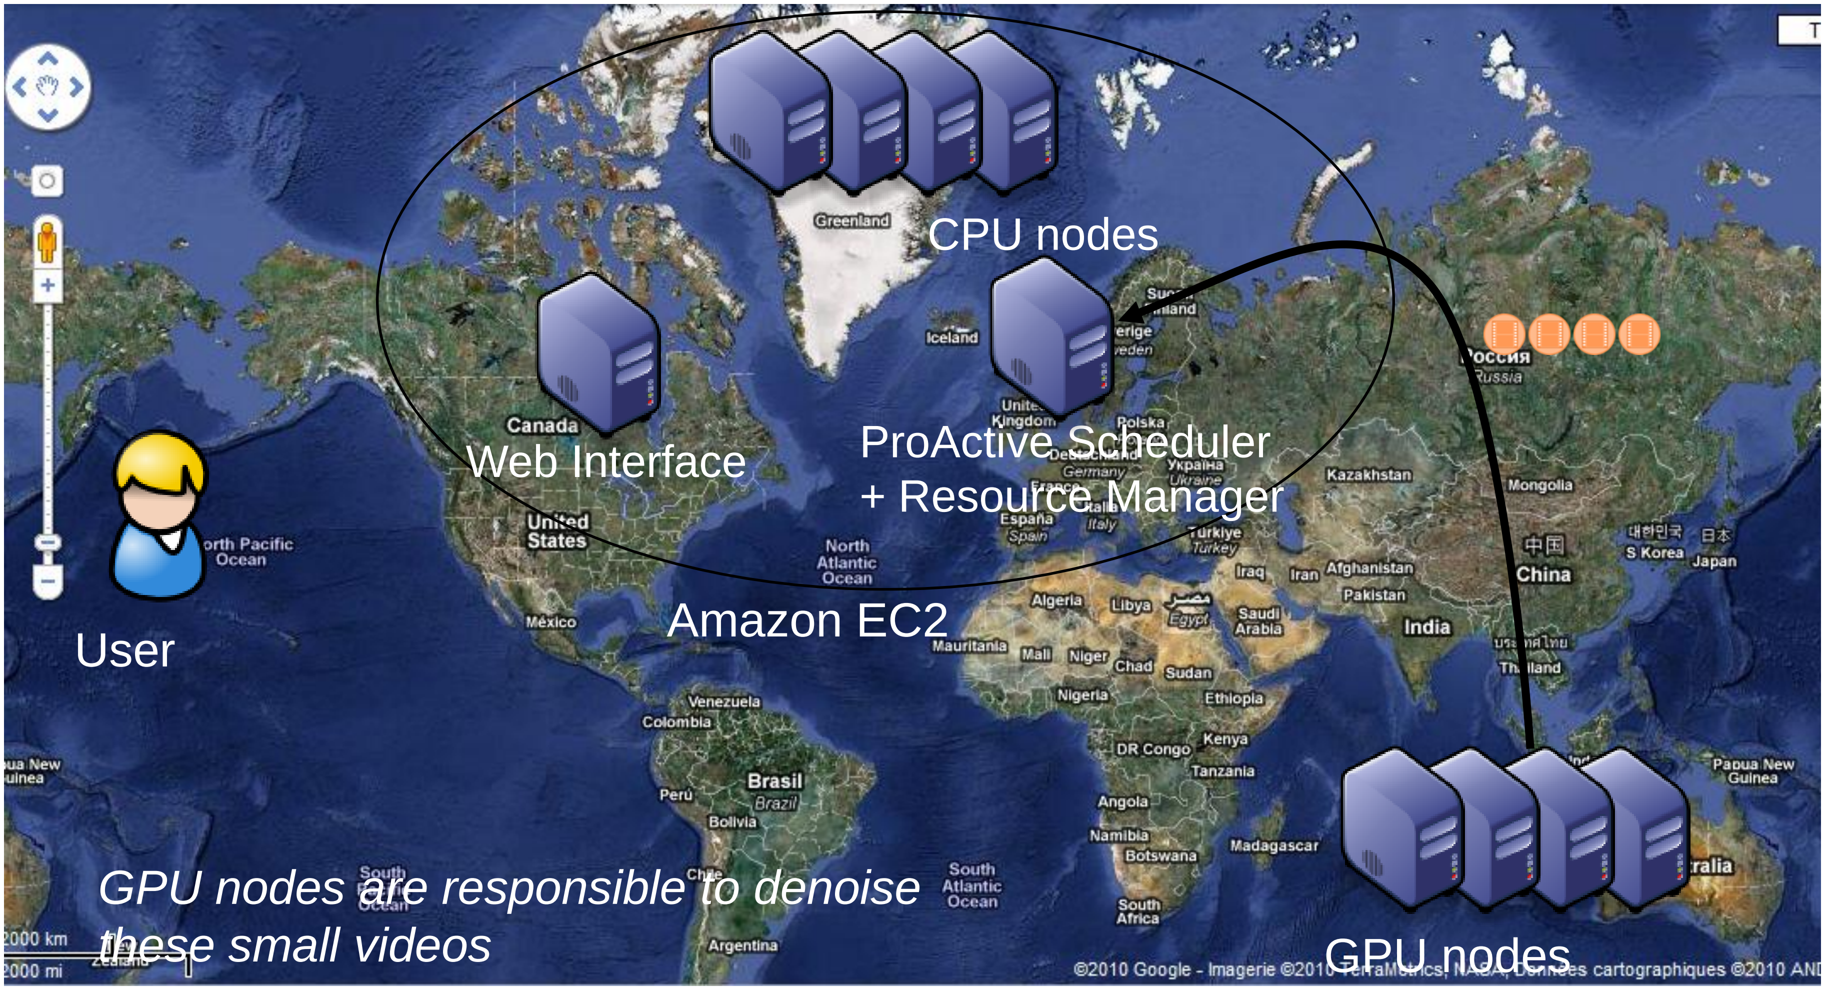

ProActive
Parallel Suite

Cloud Seeding with ProActive

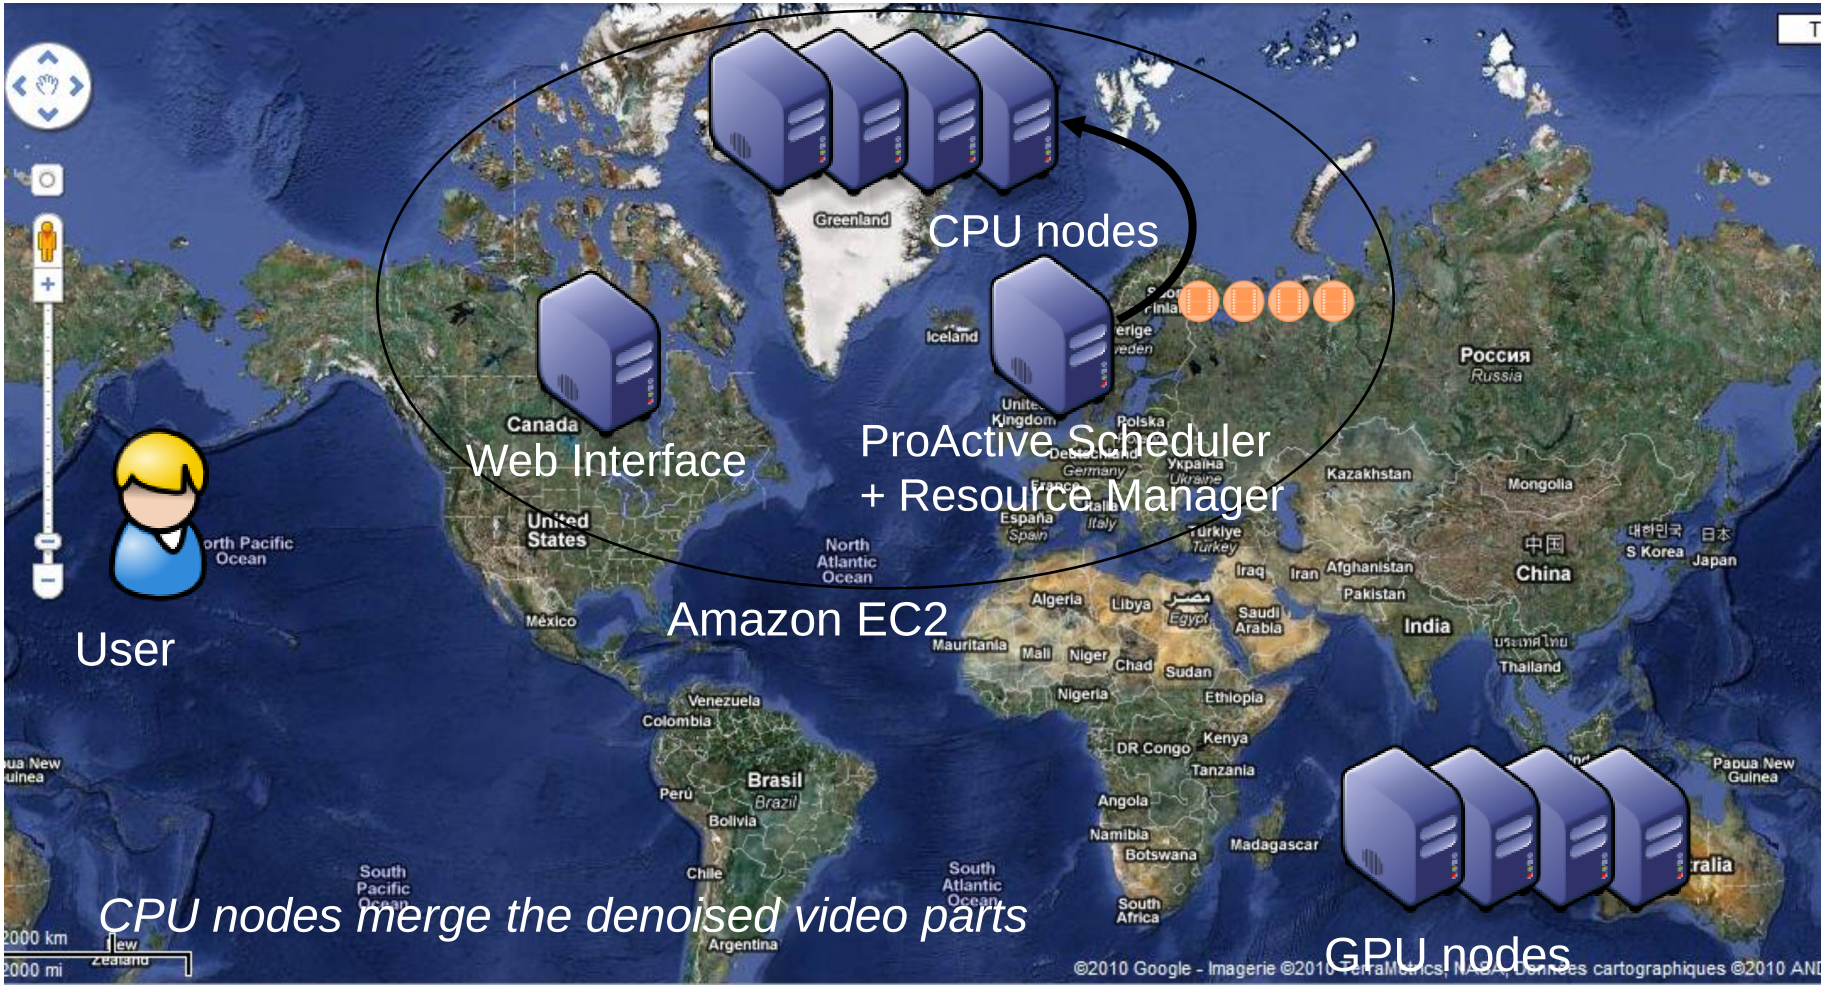

ProActive Parallel Suite

Cloud Seeding with ProActive

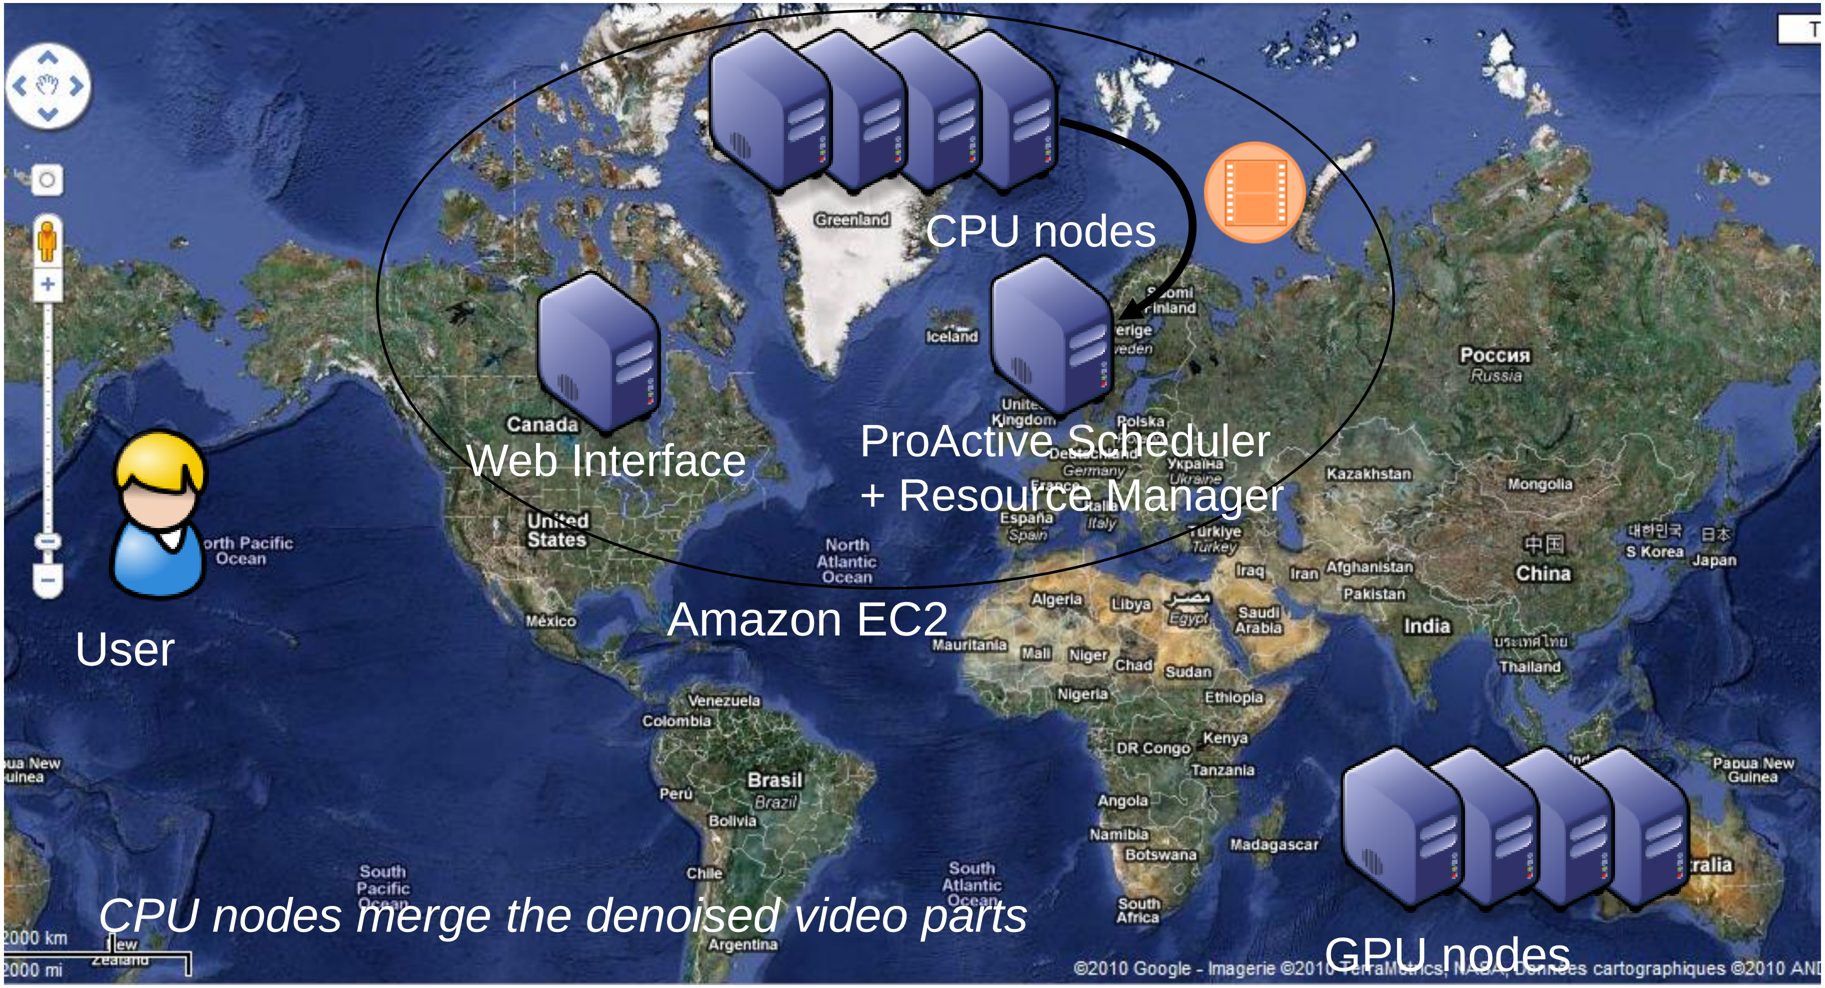

ProActive Parallel Suite

Cloud Seeding with ProActive

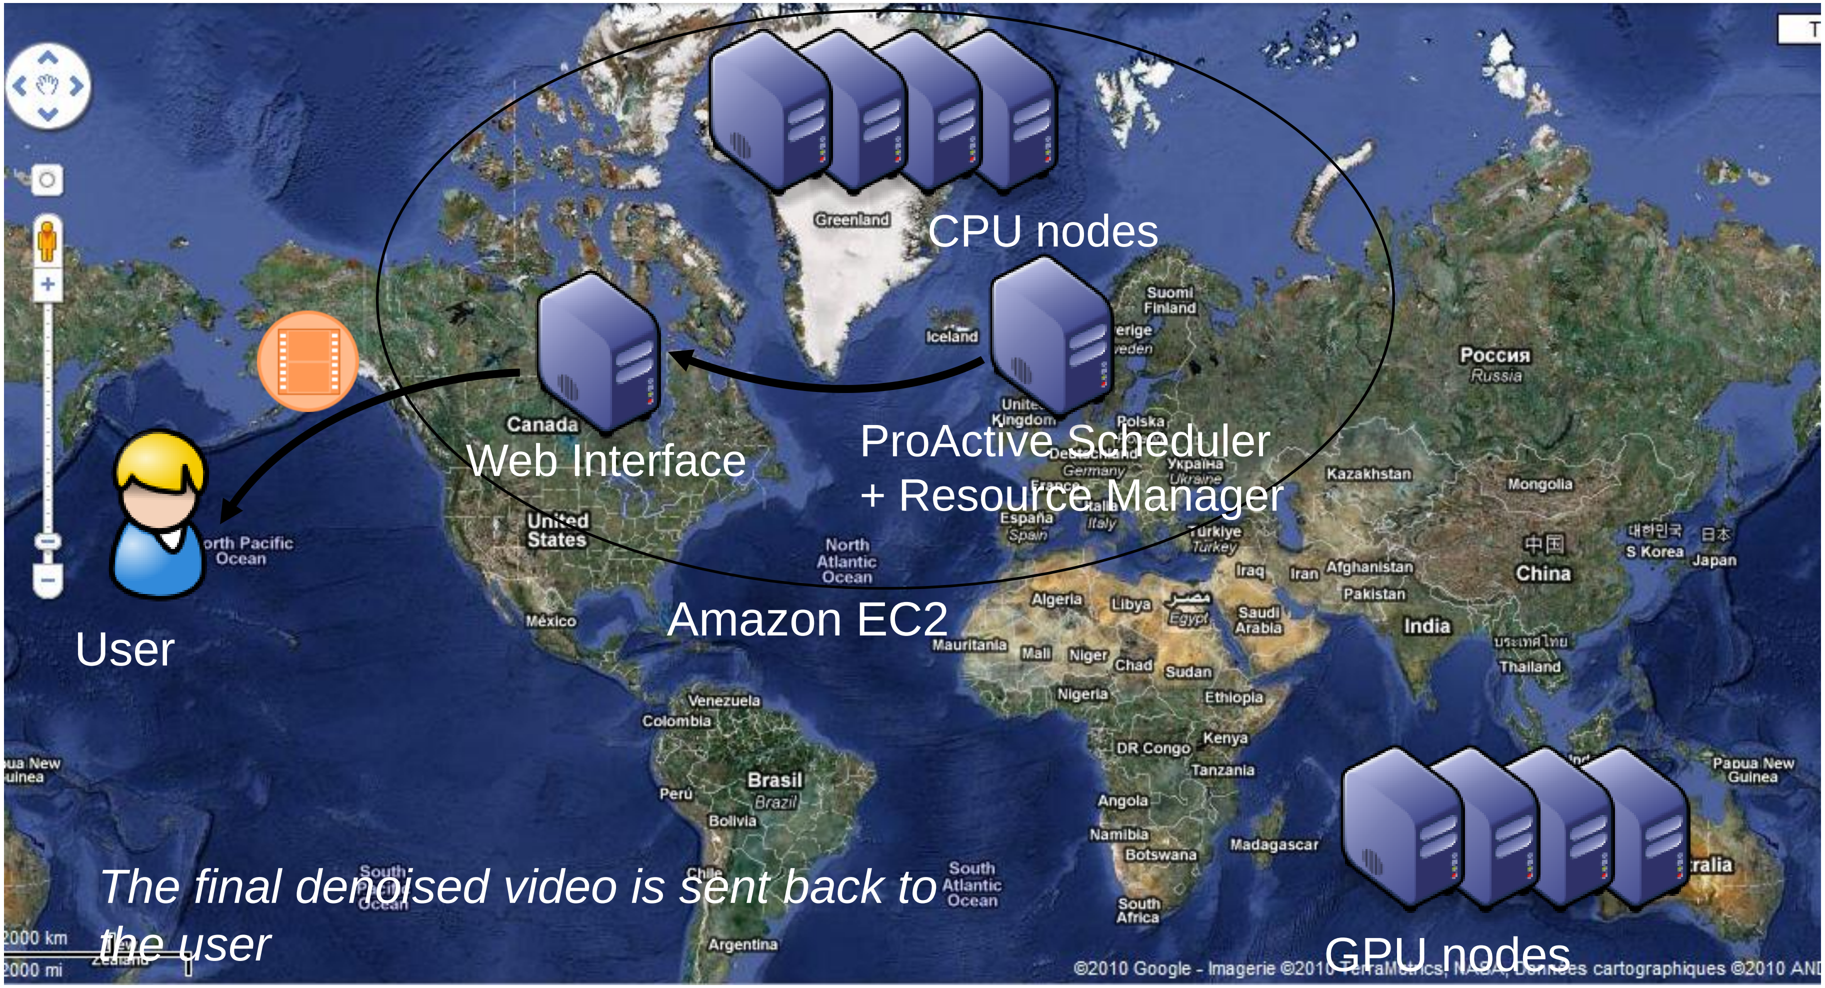

ProActive Parallel Suite

Real Live Demo: ProActive Scheduler & Resource Manager

ProActive
Parallel Suite

Use Case 1: BioTechs

IPMC Use Case and Collaboration

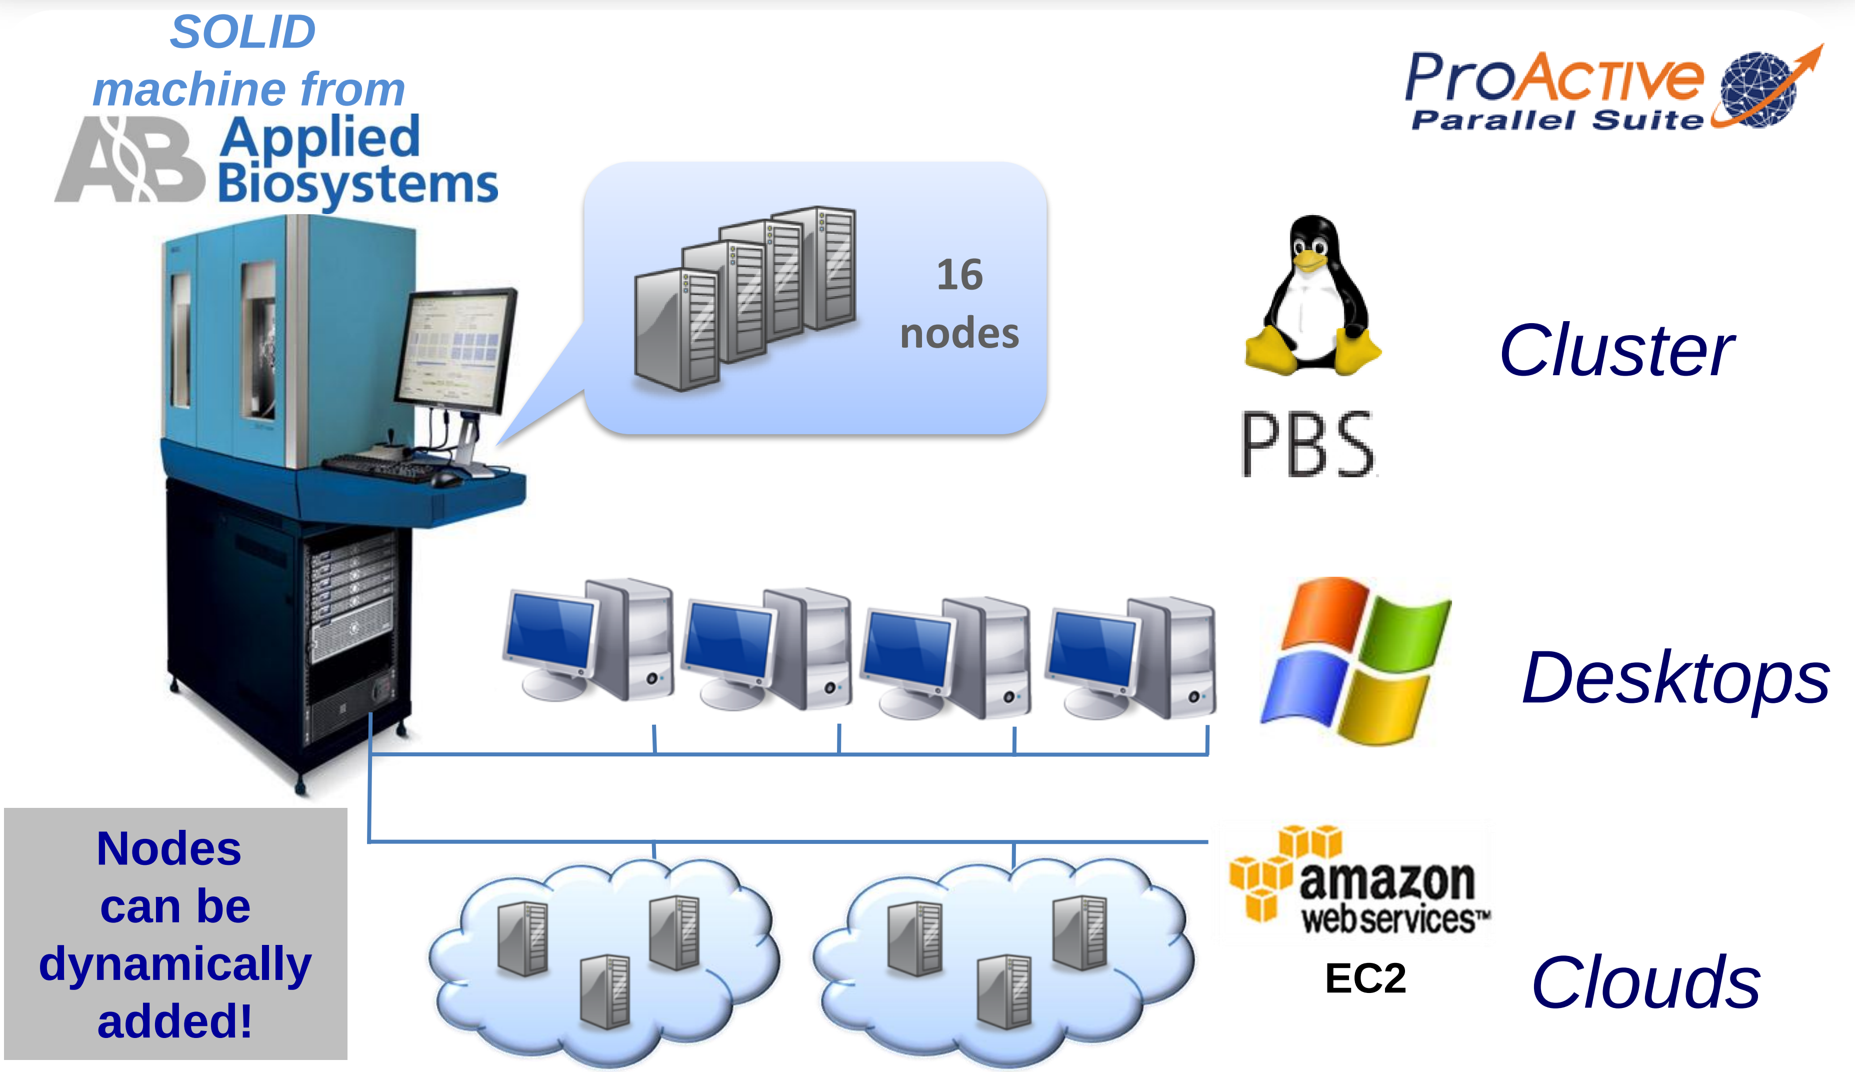

Benchmarks

- 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)

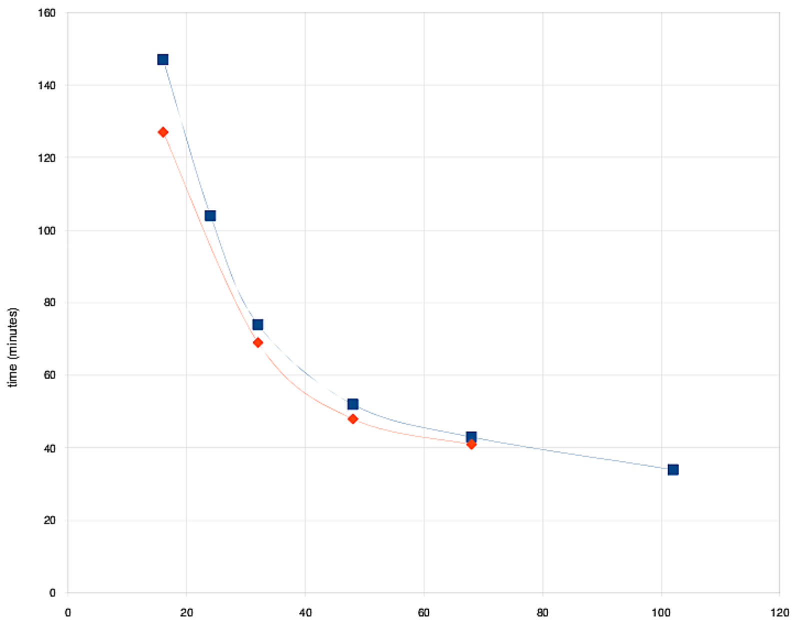

4 Time FASTER from 20 to 100
Speed Up of 80 / Th.
Sequential : 50 h ➔ 35 mn

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

UC 2: IT

SOA Analysis of Web Server Logs

hp invent

ORACLE

INRIA SOPHIA ANTIPOLIS

Activeeon SCALE BEYOND LIMITS

ProActive Parallel Suite

Parallel Services

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

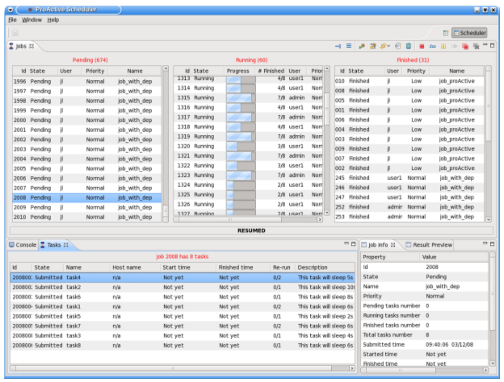

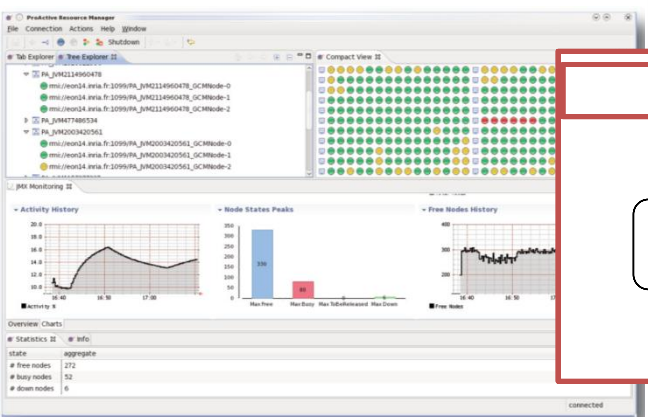

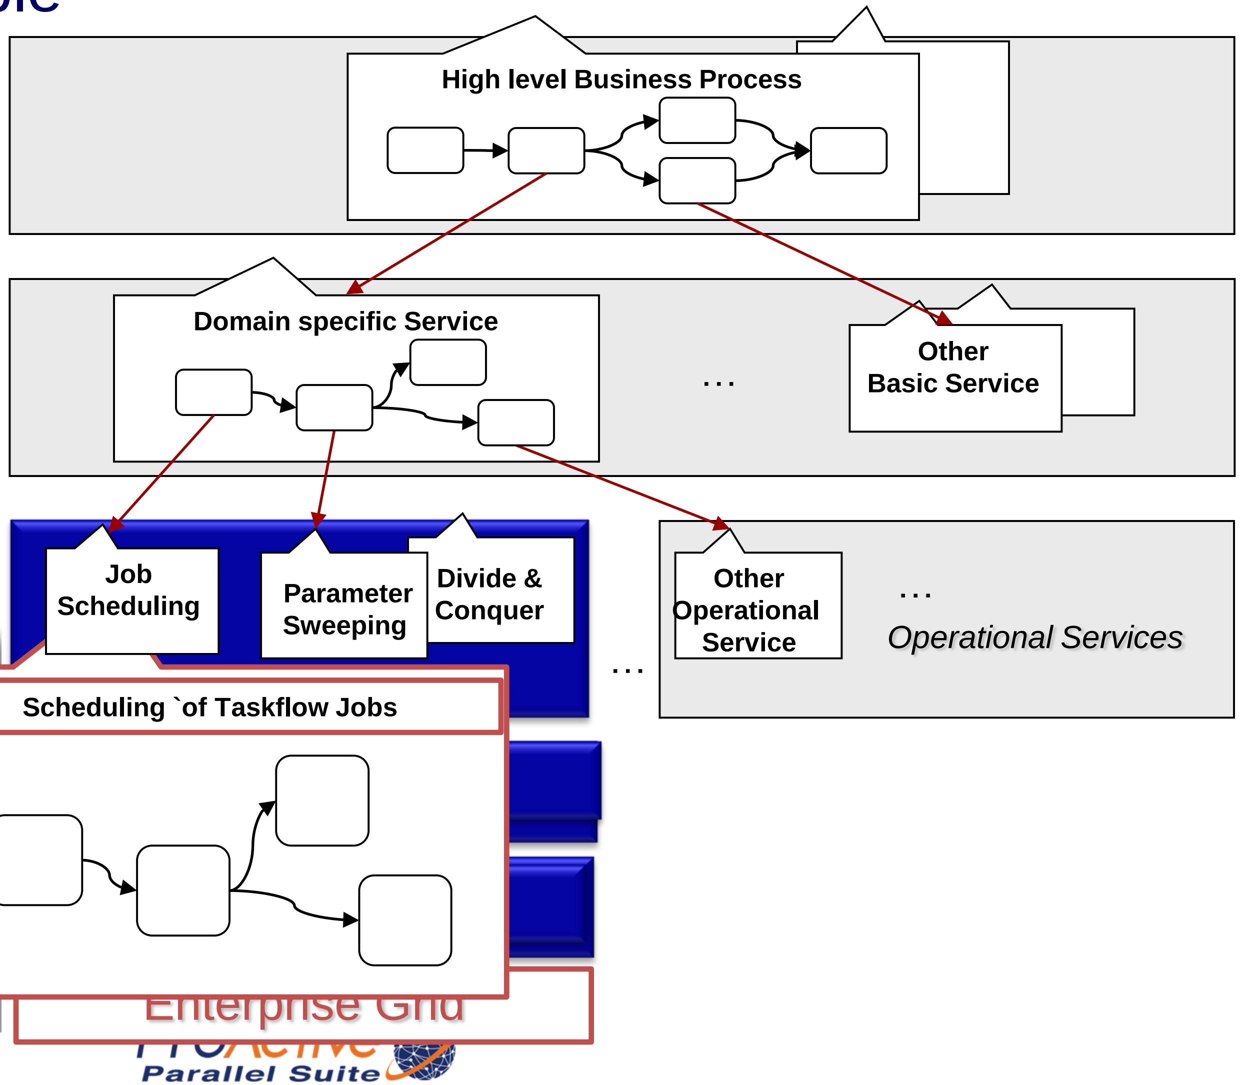

AGOS Platform Management

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

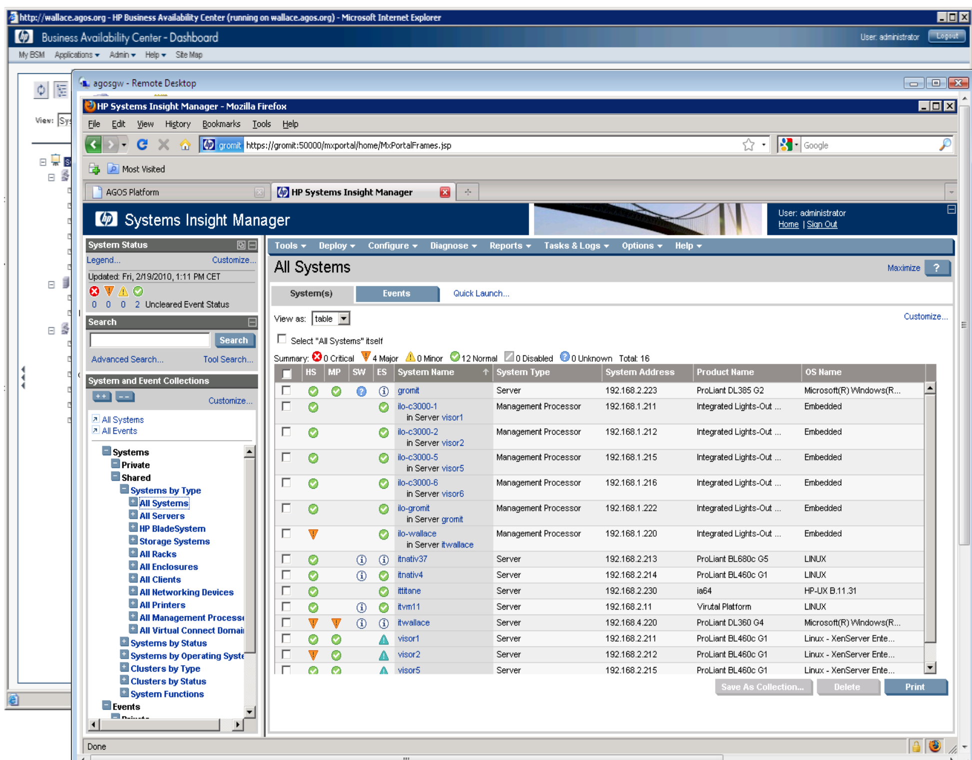

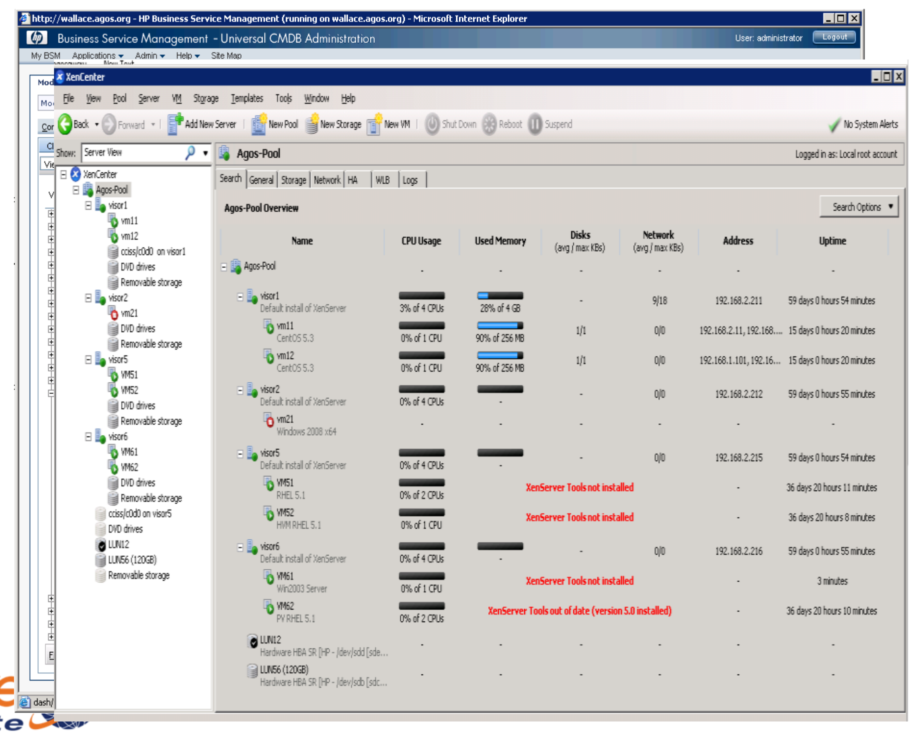

UC 3: Acceleration of Financial Valuations

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

How Does it Work? Price-it Computing Distribution

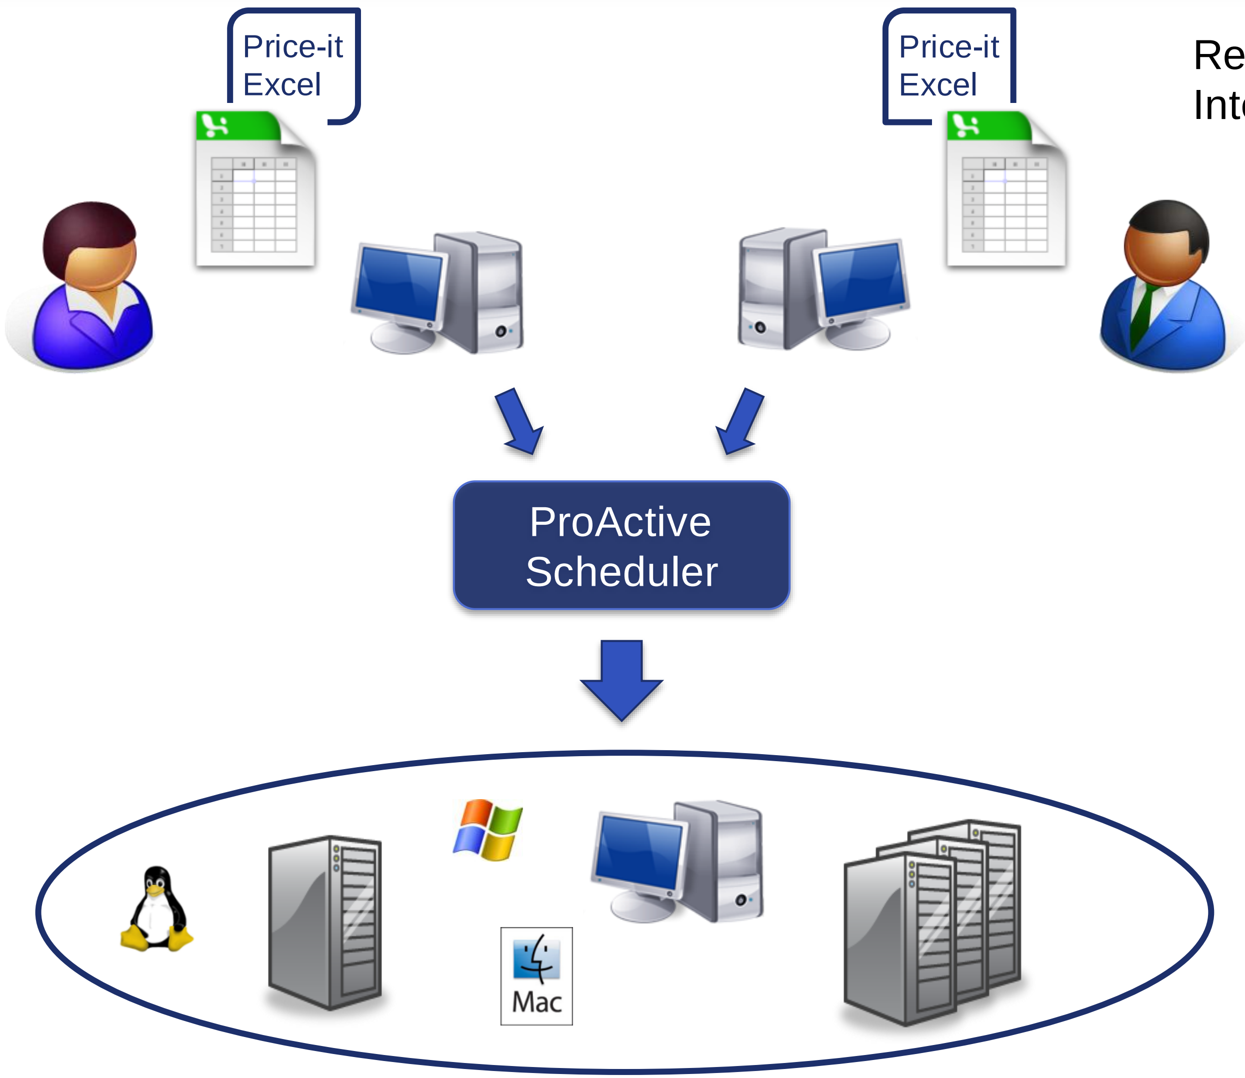

Regular Price-it Excel Interface

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!

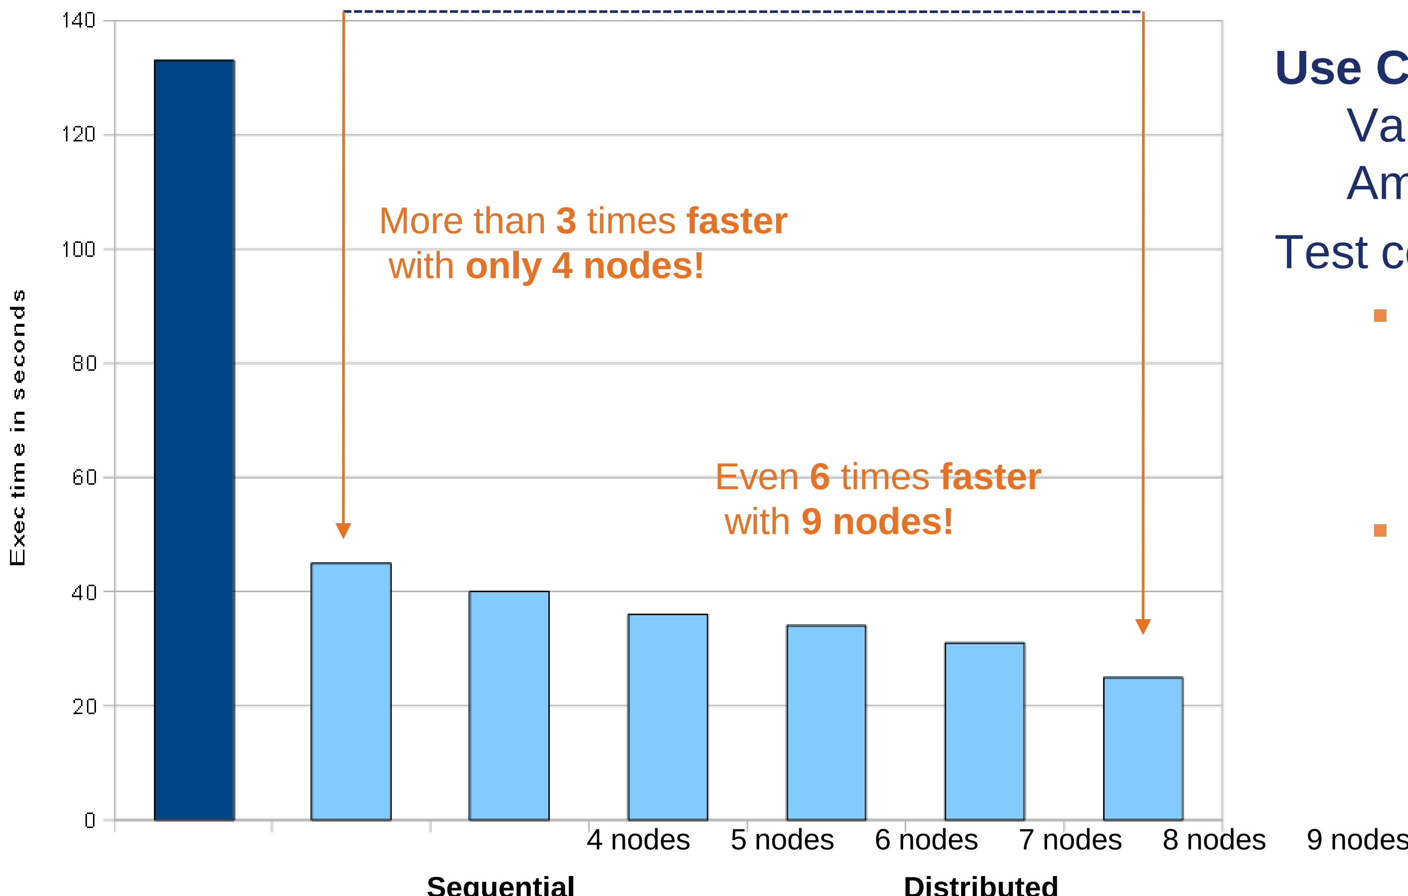

Use Case: Bermuda Vanilla, Model American MC

Test conditions:

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

Use Case 4: OMD2

Distributed Multi-Disciplinary Optimizations

ProActive Parallel Suite

Coupling Mechanics, Aerodynamics …

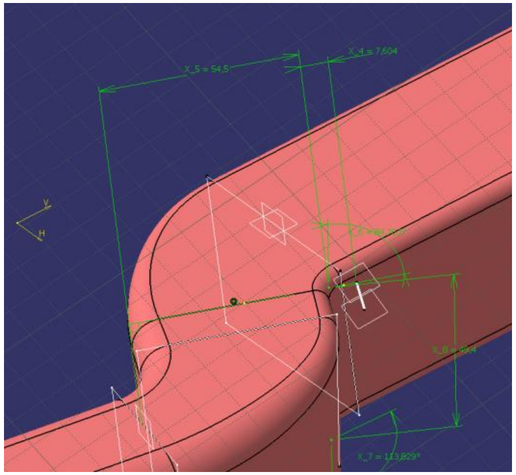

3D Air Conditionning

10min CPU

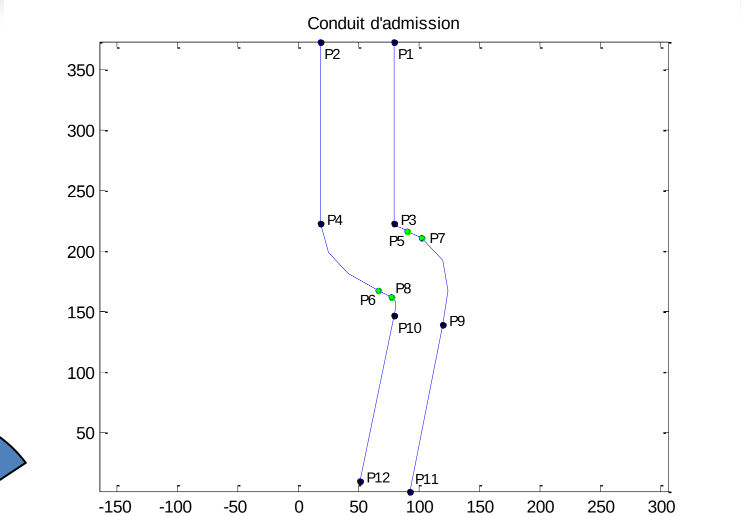

2D Air Conditionning

<1min CPU

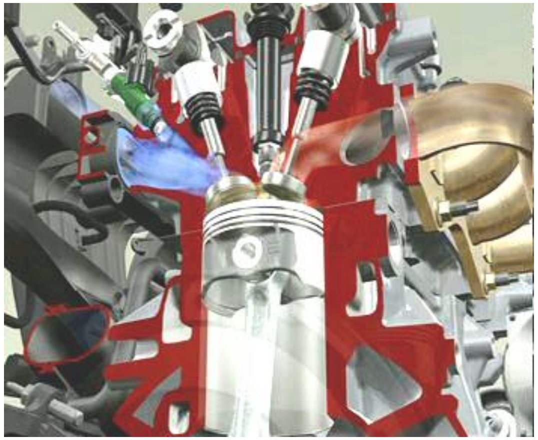

Cylinder Head

100h CPU

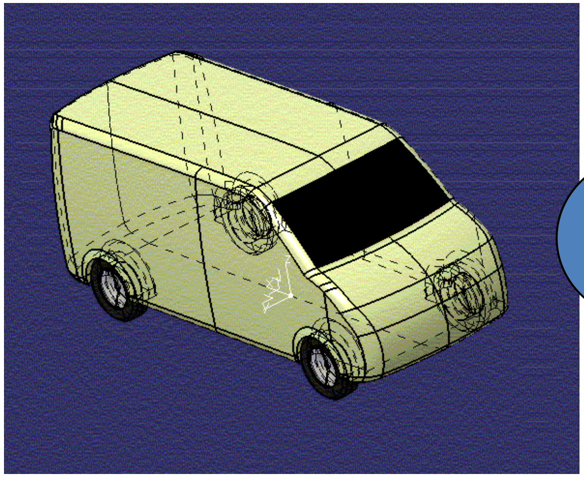

External Aerodynamic

1000h CPU

ProActive Parallel Suite

ProActive OMD2 Demo

1000 Cores
Production
Cloud Portal

ProActive Parallel Suite

Use Case 4-Bis: Hydrodynamic with K-Epsilon and FineMarine

K-EPSILON
HYDRODYNAMIC & AERODYNAMIC

Hydrodynamic Optimization: Workflow generated from a GUI

ProActive Studio → Graphical Workflow Editor

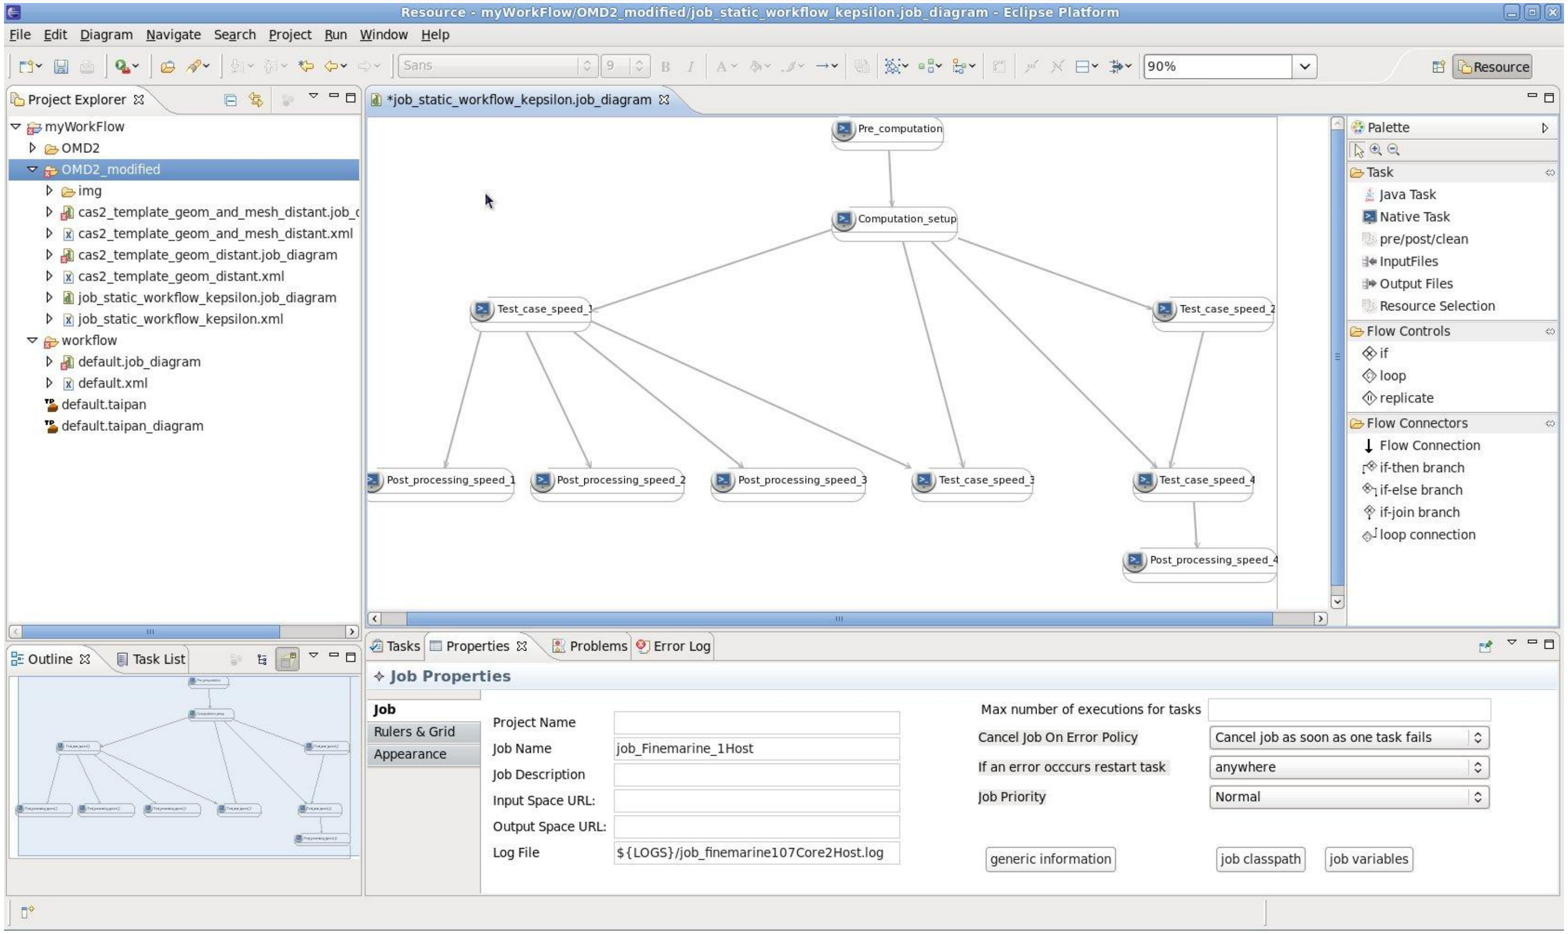

Real Live Demo:
CFD Distributed // Workflow

ProActive Parallel Suite

Hydrodynamic Optimization: Execution

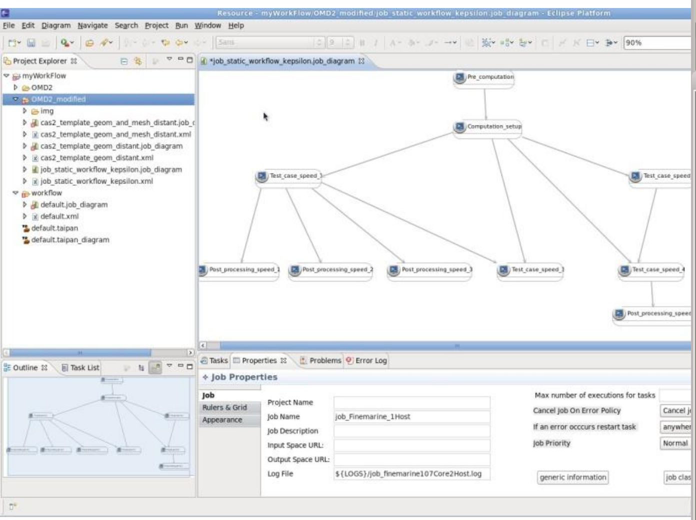

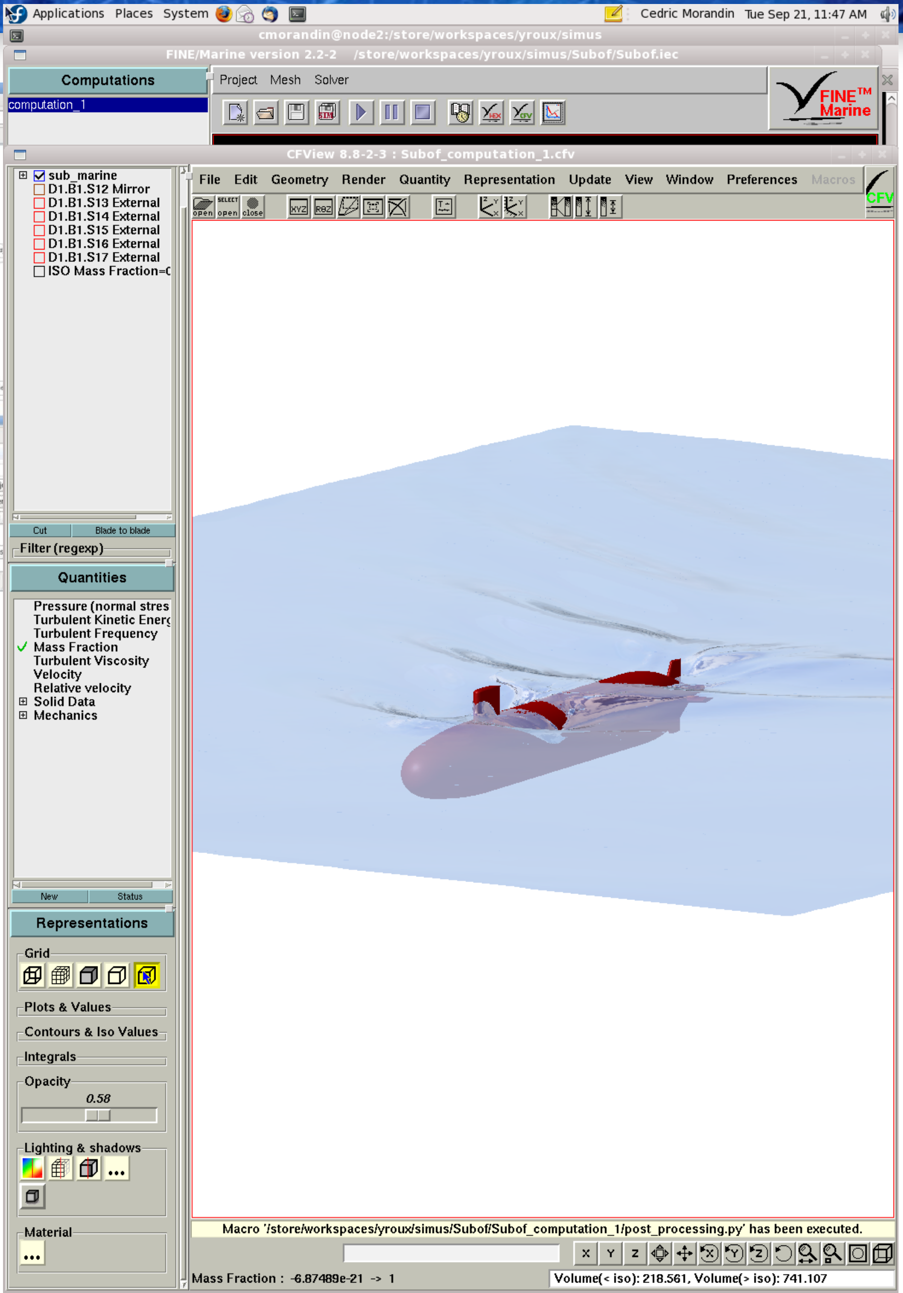

Hydrodynamic: Remote Steering during execution

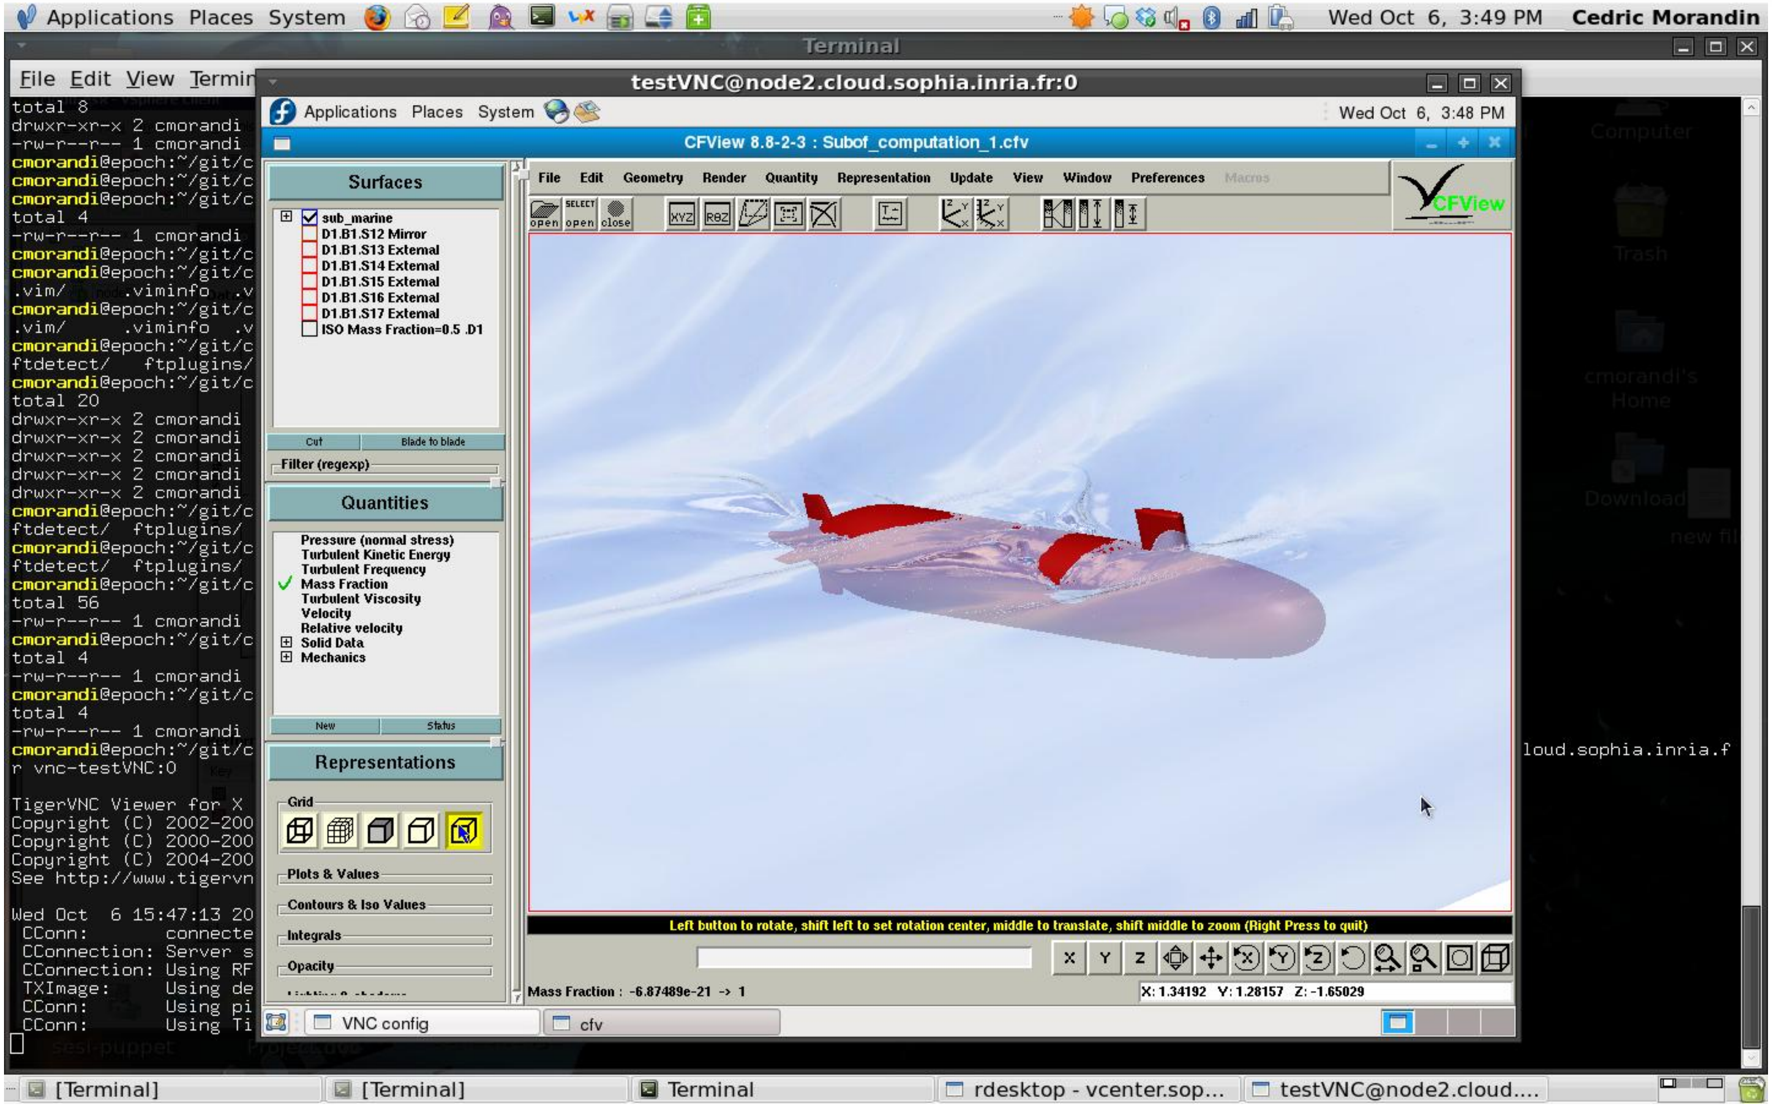

ProActive Parallel Suite

Conclusions

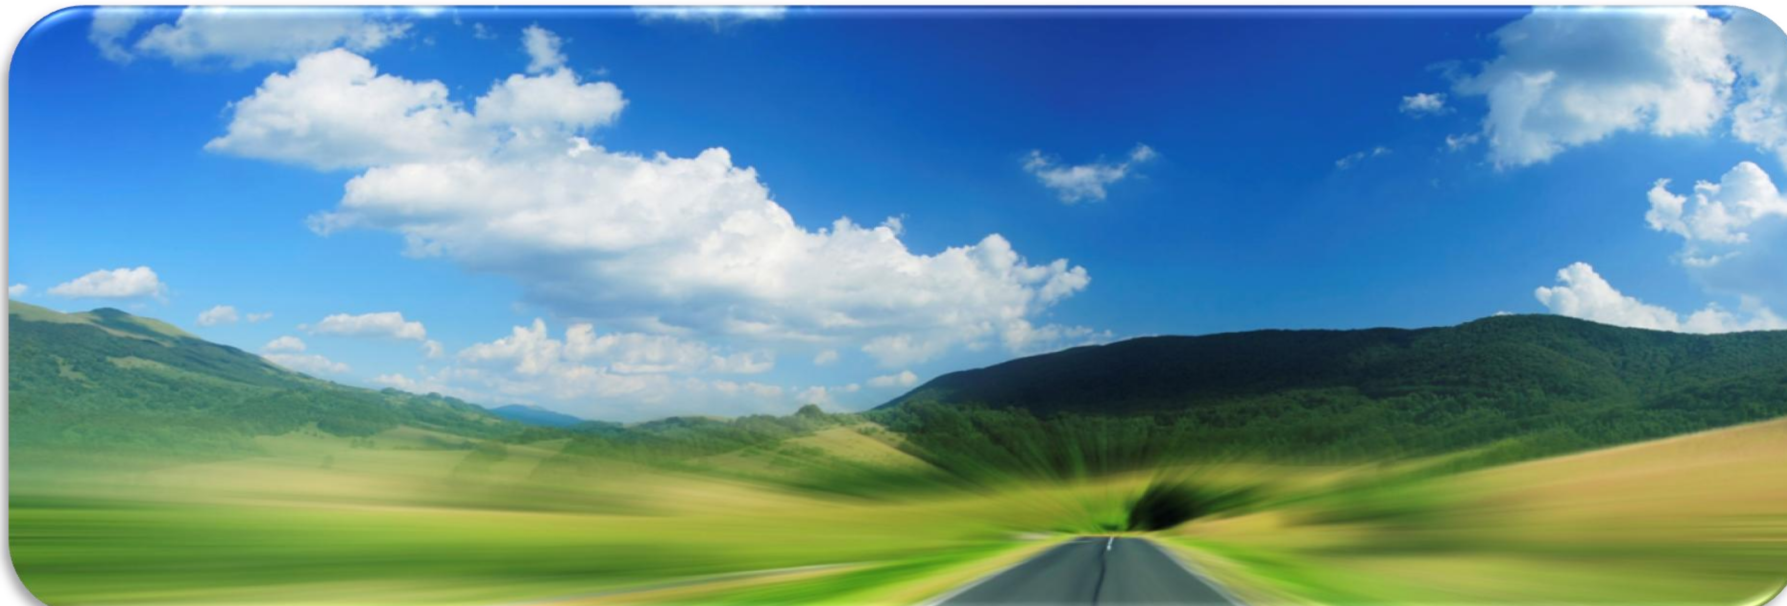

Conclusion: Technology Preview

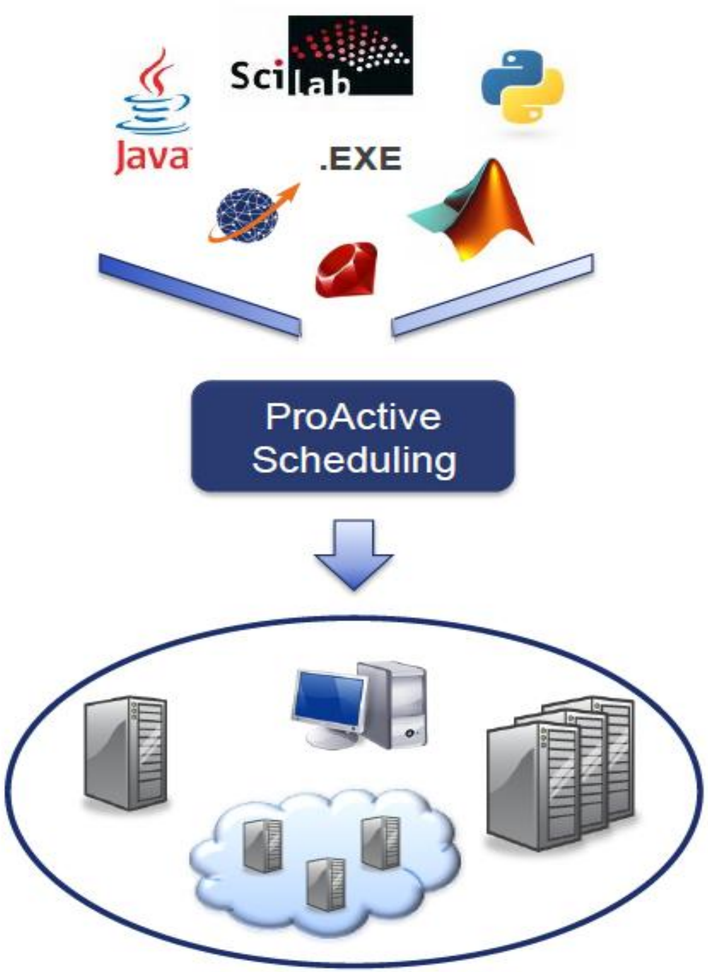

- ProActive Fine Grain CLOUD management:
  - ➔Pricing at the second (like GSM)
- Open Source Cloudware Initiative (OSCi)

➔ Elastic Clouds

http://ProActive.inria.fr

Industrial (1750) & Cloud Revolution Compared

	Industrial Revolution	Cloud Revolution
Concept	Mechanization and centralization of manufacturing activities	Computing as a Utility Centralization of Data Center
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


Thank you for your attention!

http://proactive.inria.fr

ProActive Parallel Suite