

ProActive Parallel Suite and OW2 OSCi: From Multi-Cores to Multi-Clouds

D. Caromel, et al.

Agenda

1. CLOUD Computing & Virtualization
2. OSCi 4 Domains
3. ProActive Parallel Suite
Programming, Scheduling, Resourcing
4. Use Cases & Demos
5. Conclusion: Cloud Revolution ?



Cloud Computing Revolution ?

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: Cloud taken for granted ?


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




CLOUD: the Multi-Core Push

Symmetrical Multi-Core: 8-ways Niagara II

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



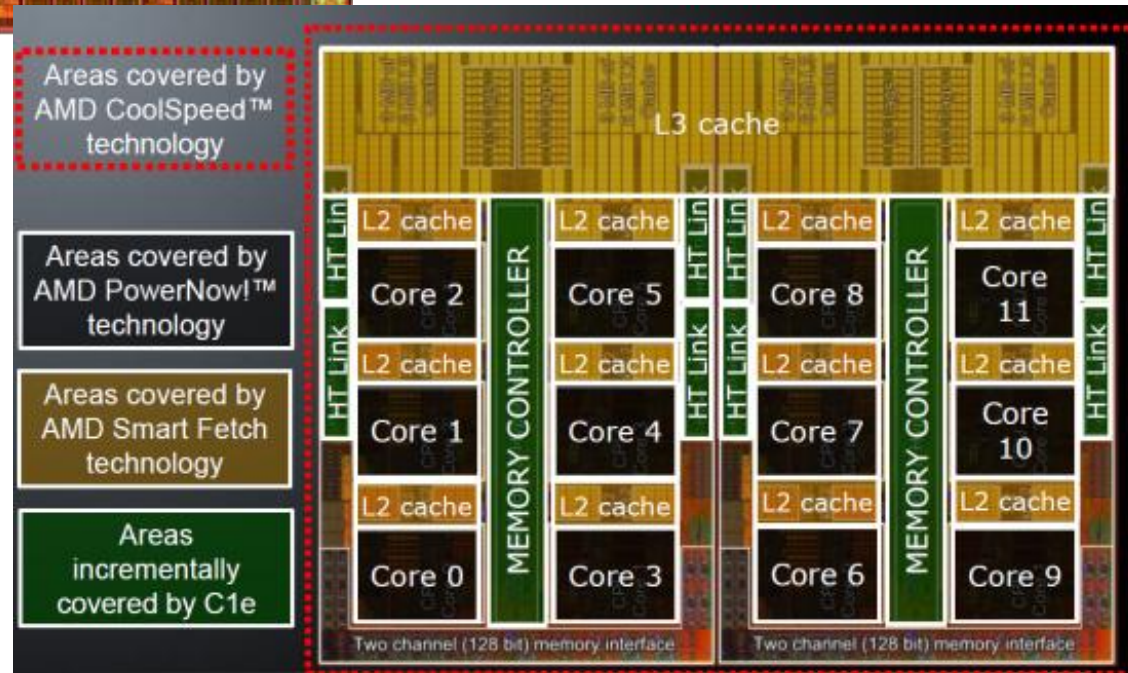
Today Off The Shelf Multi-Cores, 3 GHz



AMD's Opteron 6174,
"Magny-Cours",
12 cores

Intel Xeon 5670,
6 cores


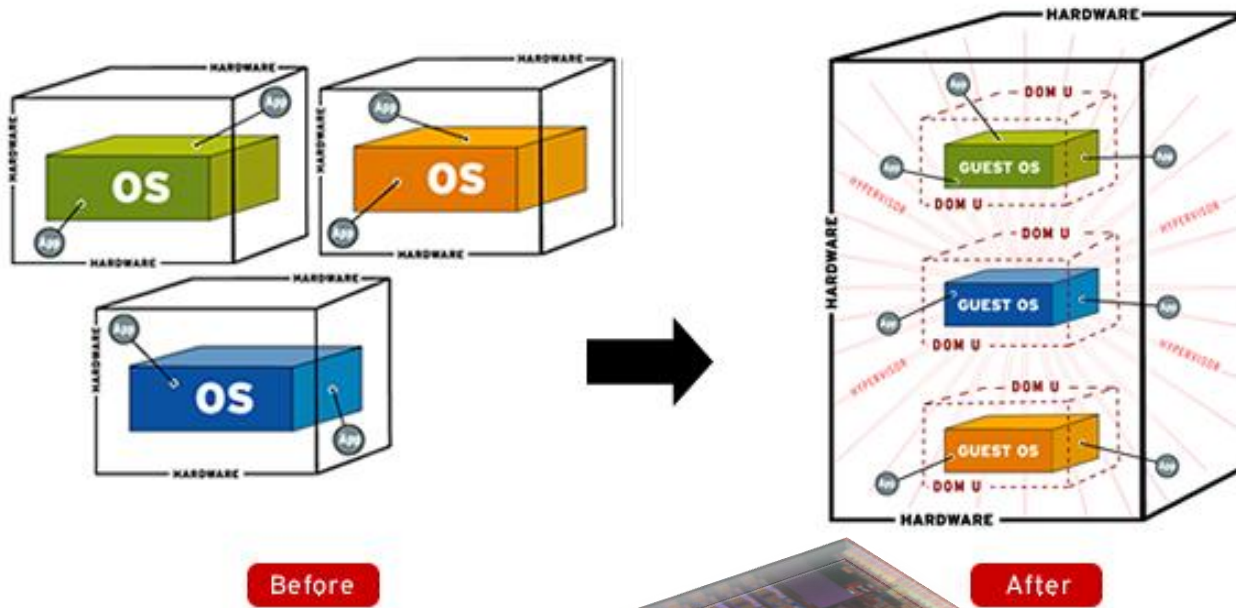
- ❑ **Moore's Law rephrased:**
Nb. of Cores double /18-24 months
- ❑ **Key expected Milestones:**
Cores per Chips (OTS)
 - 2012: 32 to 64
 - 2014: 64 to 128
- ❑ **1 Million Cores Parallel Machines in 2014**



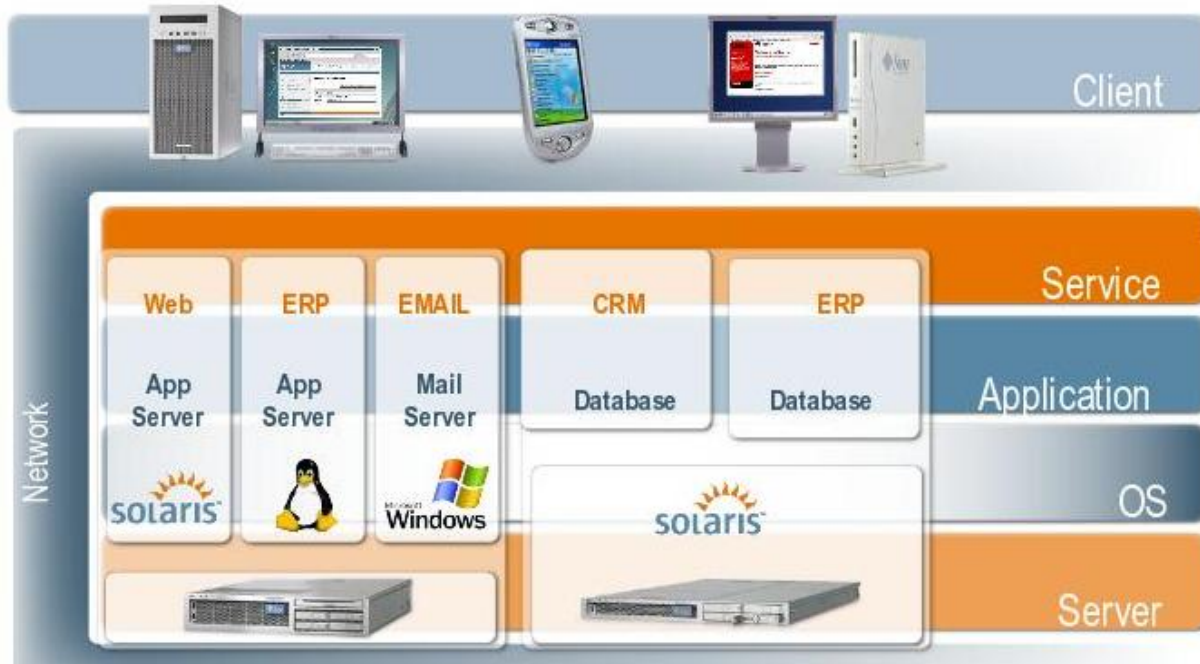


Virtualization

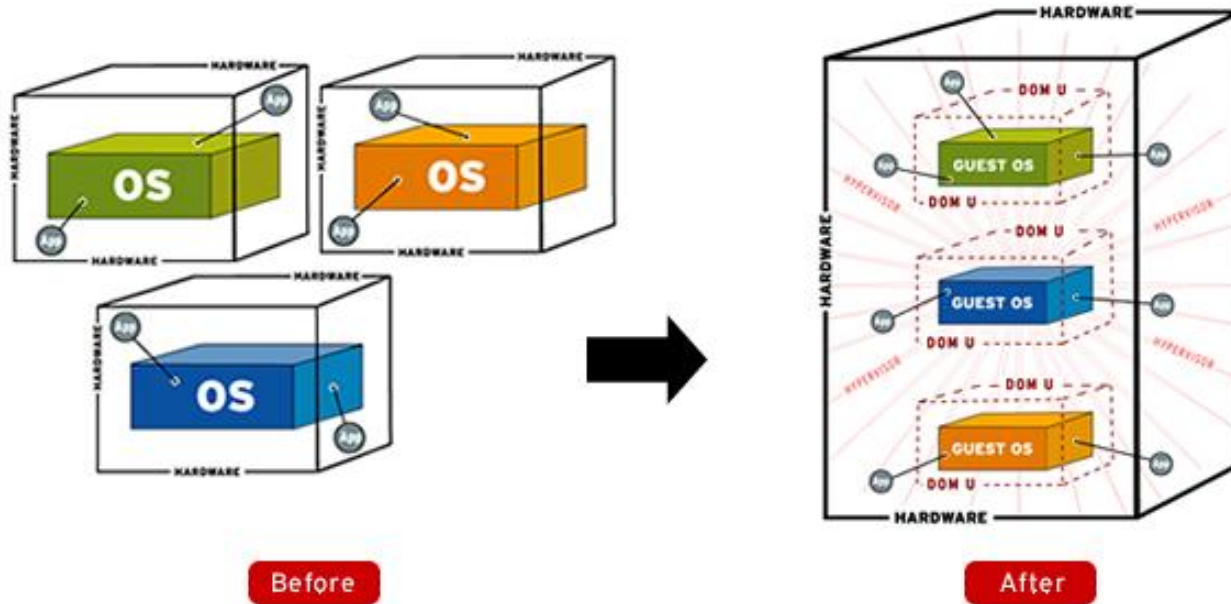
Virtualization



Virtualization




Virtualization



What we Used to do as Syst. Admin.



With Virtualization + Software Appliance



OW2 OSCi Open Source Cloud initiative

OW2 OSCi

Open Source Cloud initiative



4 Strategic Domains

- ➔ Self-sizing and green PaaS
- ➔ Massively Distributed Services
- ➔ BI 4 Cloud
- ➔ Massively Distributed clouds



Domain 1: Self-sizing and green PaaS

➔ Goal: improve the platform efficiency

- Performance
- Energy cost

➔ Challenges

- Cluster growth/shrink capabilities according to workload
- Data center multi-tenancy with limited capacities requires arbitration policies
- Data center placement for minimizing the energy consumption

➔ Use case: Orange Infrastructure

- Large-scale data centers (+40,000 x86 servers)
- JavaEE JOnAS middleware stack: +250 applications, +1,000 application server instances

➔ Technologies

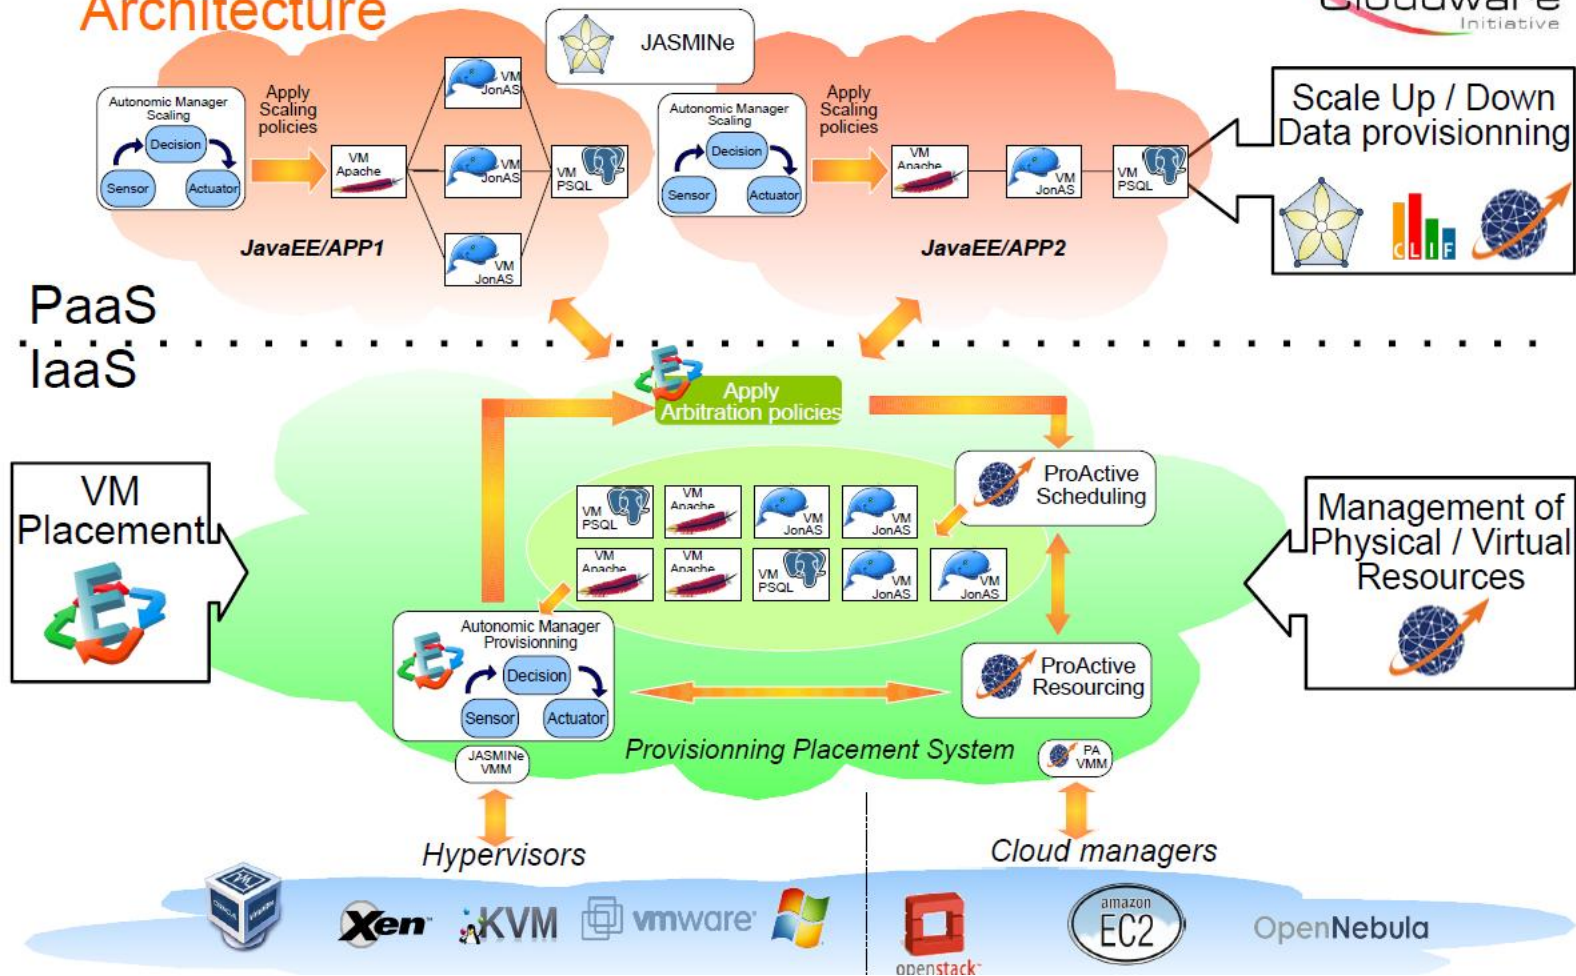
- OW2: JOnAS, Jasmine, ProActive, Clif, Entropy
- 3rd party: Xen, KVM + Collaboration in progress: OpenStack, OpenNebula, Ubuntu
- Technology to be developed:

➔ Partners

- Bull, Ecole Mines Nantes, Inria, Orange, ActiveEon
O-Engine, Iscas, Buaa, PKU, UCM, ...

SHORT TERM
Work
in Progress

Self-sizing and green PaaS Architecture





Domain 2: Massively Distributed Services

➔ Goal: Massive Heterogeneous Services with clouds

- Distributed Service Bus (DSB) made of Petals nodes organized in domains/clouds
- Highly distributed service architectures

➔ Challenges

- Large scale P2P distributed registry
- QoS policy management over heterogeneous domains
- Inter clouds interoperability


➔ Technologies

- OW2: Petals ESB, ProActive
- 3rd Party: Eucalyptus, OpenNebula, NiftyName, OpenStack, ...
- Technology to be developed:...

➔ Partners

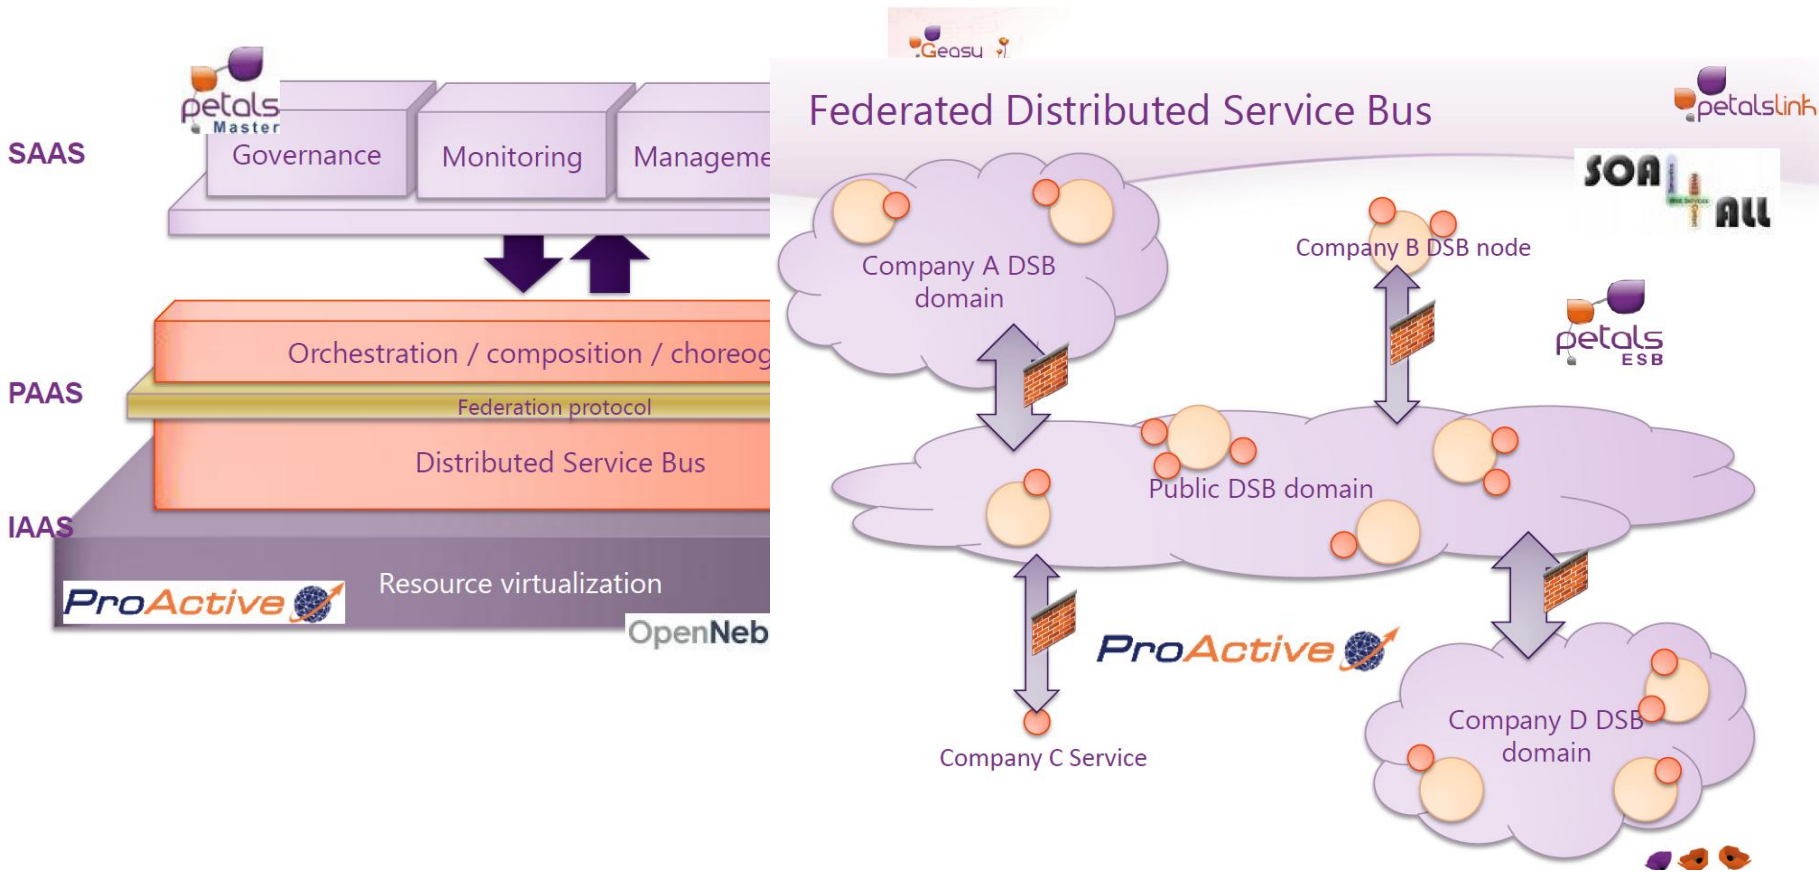
- Petals Link, ActiveEon, Inria, USP, Fokus


SHORT TERM



Domain 2: Massively Distributed Services

Architecture / Technologies





Domain 3: OW2 BI4Cloud initiative



BI4cloud, providing images pre-compiled with components for Business Intelligence from industry's leading vendors


OW2 members:

- ➔ Ingres
- ➔ Talend
- ➔ Jaspersoft
- ➔ Engineering/SpagoBI
- ➔ BonitaSoft

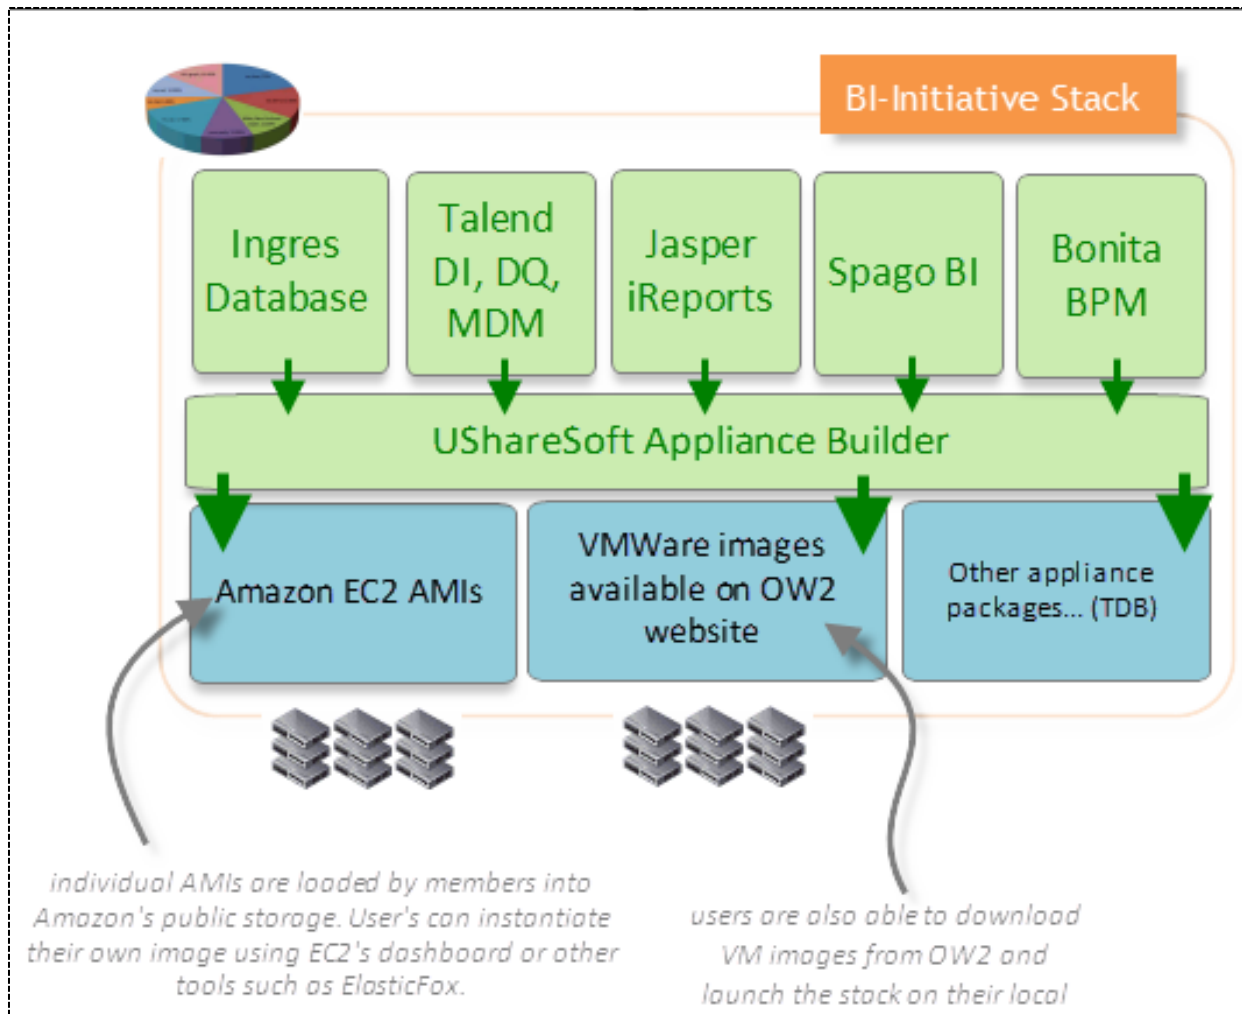
3rd Party/non OW2 partners:

- ➔ usharesoft





Domain 3: OW2 BI4Cloud initiative





Domain 4: Massively distributed clouds

➔ Goal: improve sovereignty and privacy

- non hierarchical and massively distributed cloud
- Next generation cloud architecture

➔ Challenges

- Large scale distributed clouds, super elasticity for massively decentralized systems
- Security, safety and privacy e.g. cryptography for the masses, personal control
- High speed public network

➔ Use Case: domestic cloud for citizens

➔ Technologies

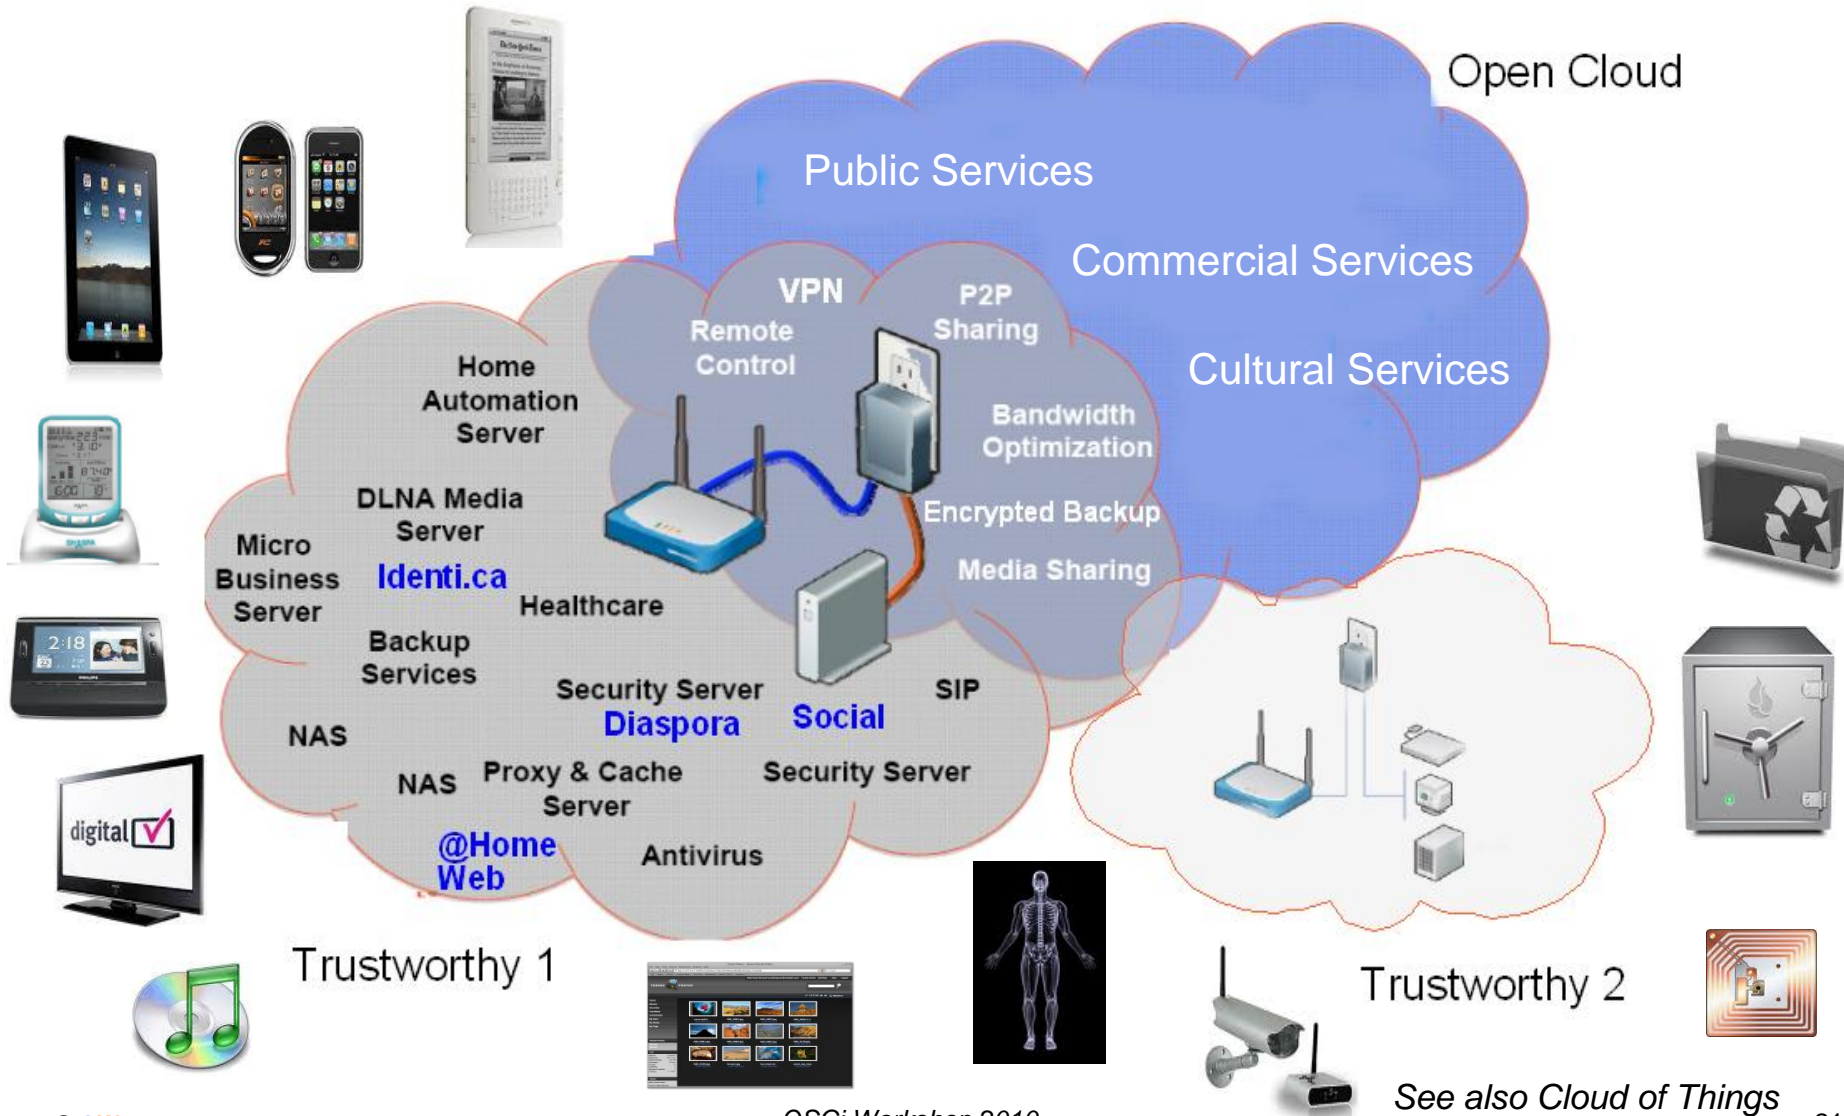
- Key R&D trend for next 5 years
- Technology to be developed:...

➔ Partners

- INRIA Sardes (JB Stefani), other partnership to be developed including 3rd party technology providers (e.g. Hedera Tech)


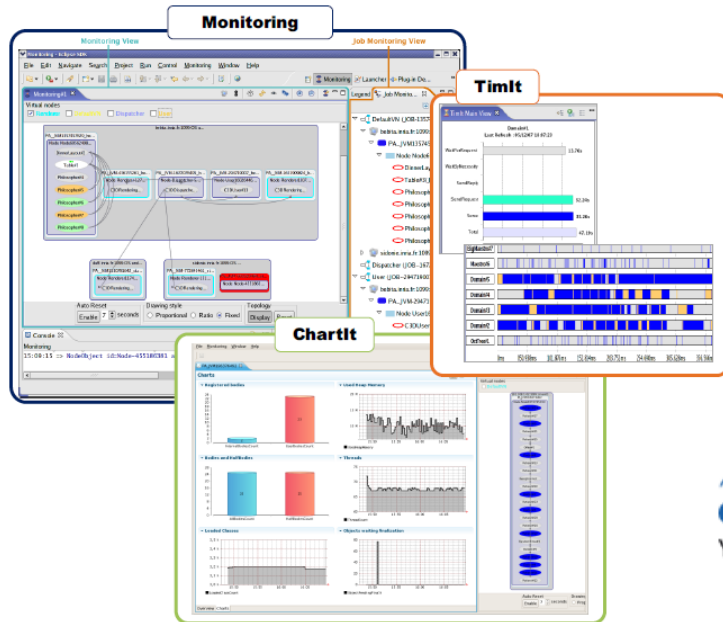
LONG TERM
Discussions
in Progress

Massively distributed clouds Architecture (as for 2015 - Draft)




3. ProActive Parallel Suite


Cloud Solution: ProActive Parallel Suite



Java Parallel Toolkit



Multi-Platform Job Scheduler



Resource Manager

amaDEUS
Your technology partner

Used in Production Today:
50 Cores → 300 Cores 2010


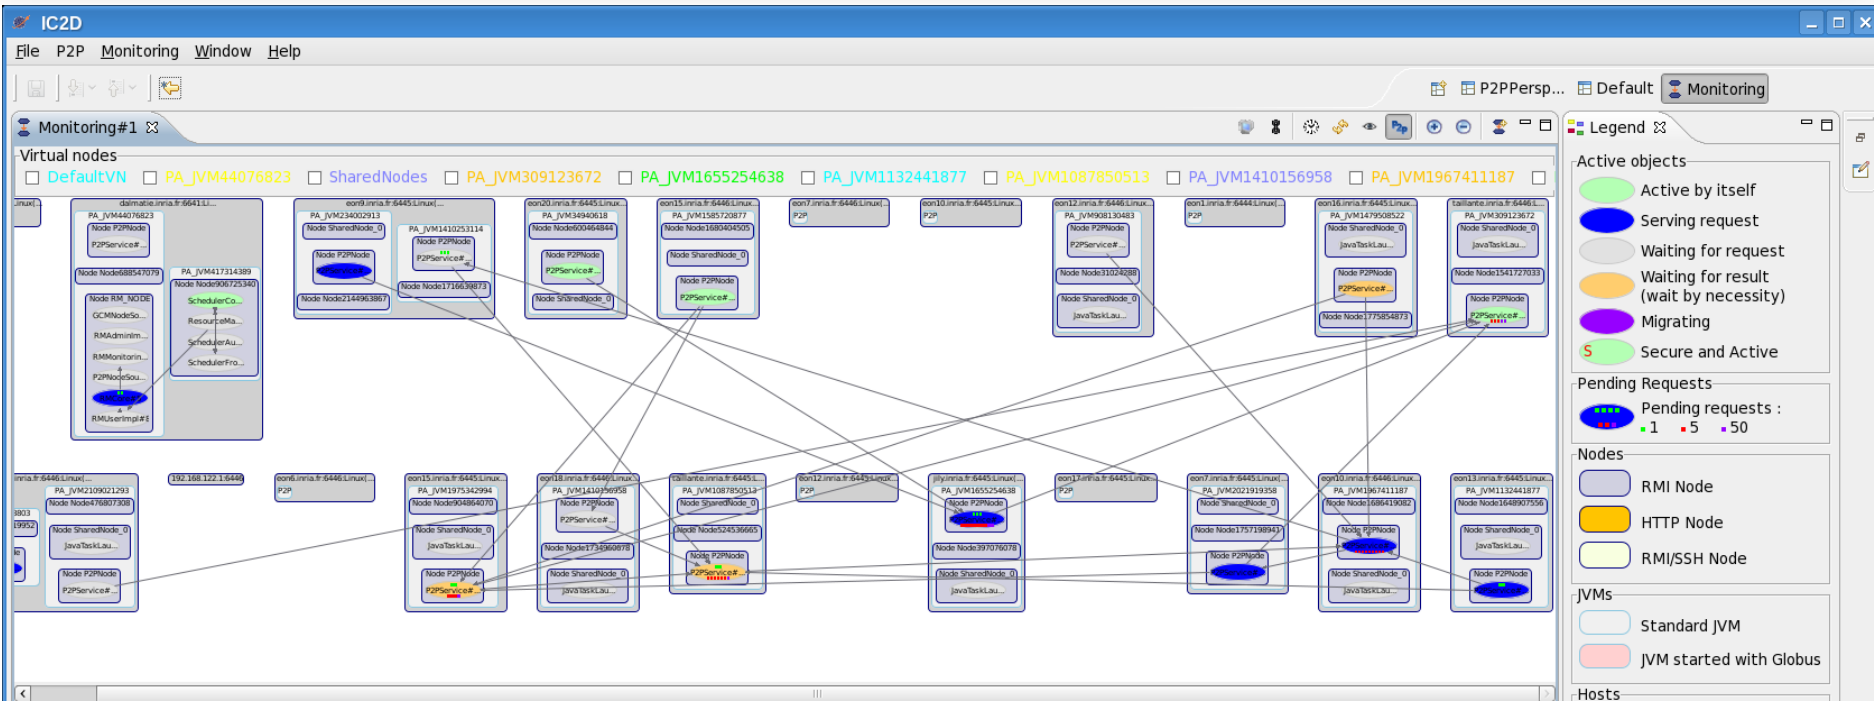
Strong Differentiation:

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


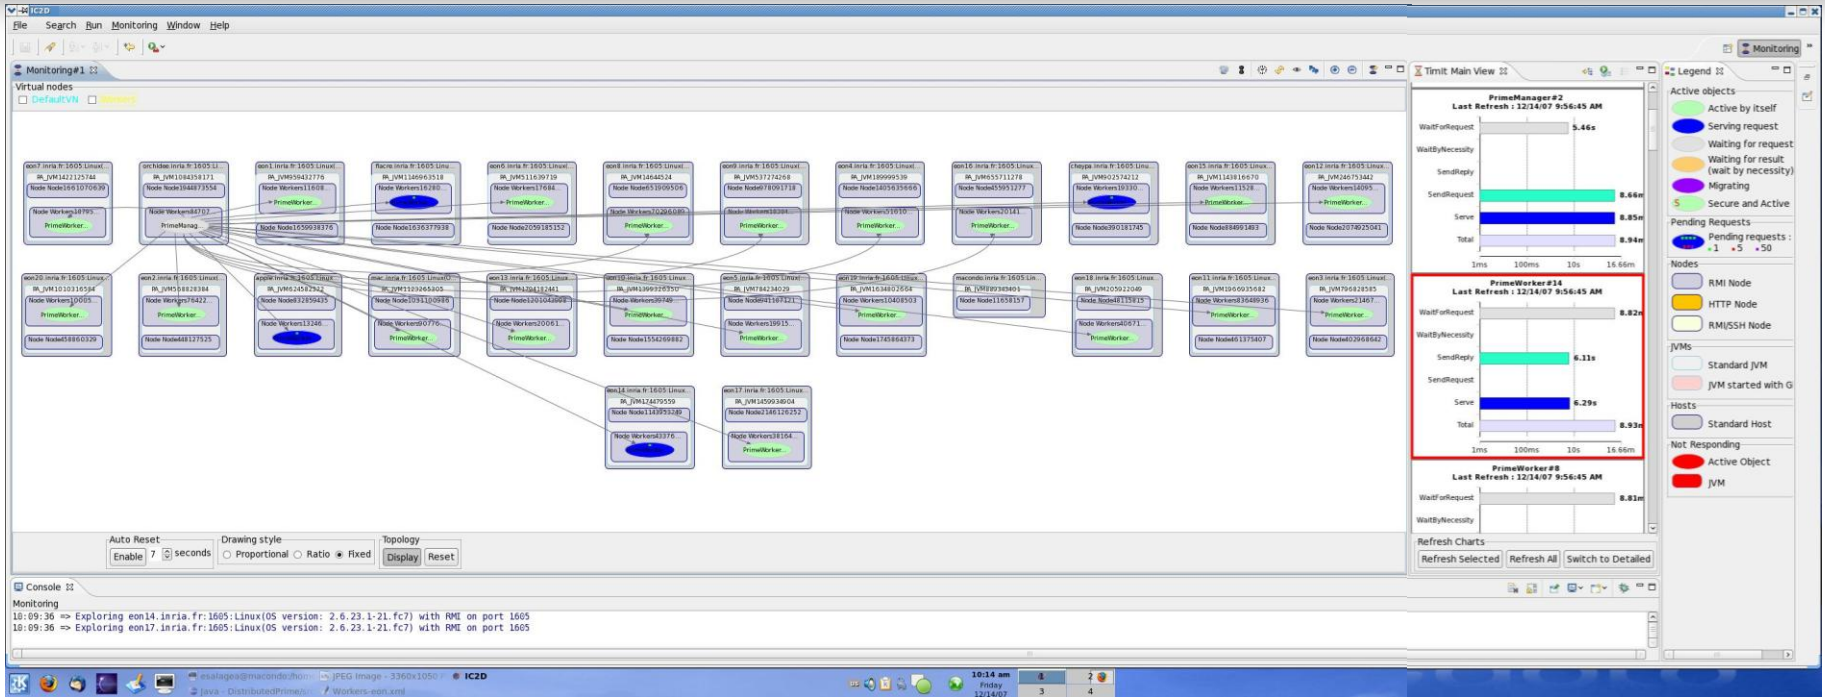


ProActive Programming: Active Objects

ProActive Programming View

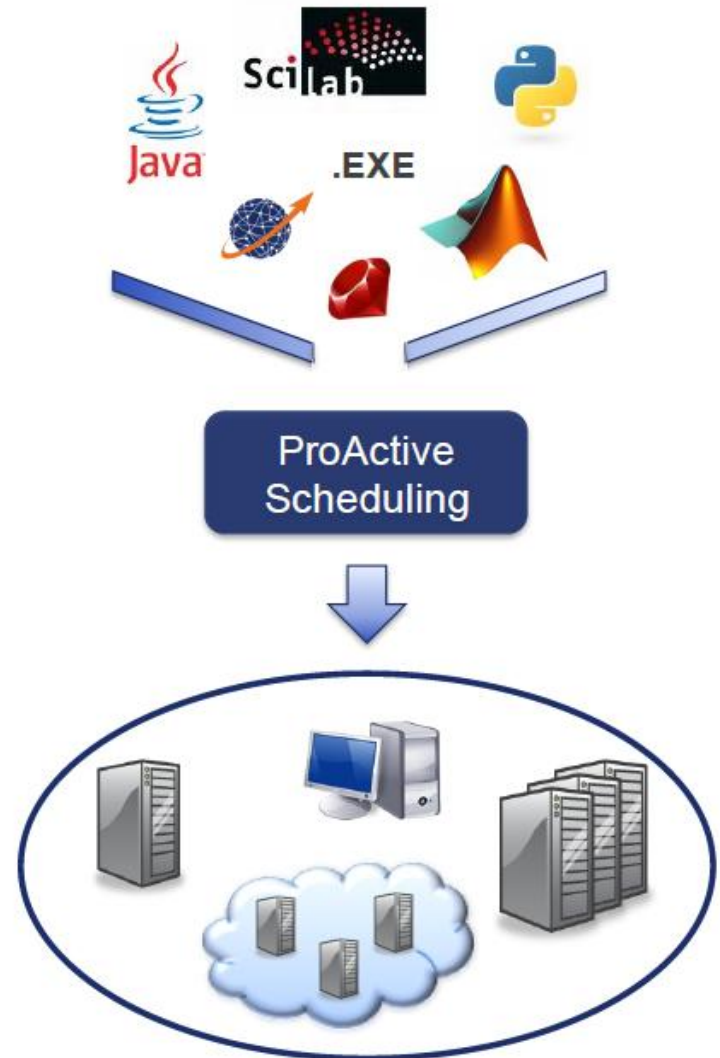


ProActive Programming View



Scheduling & Resourcing

ProActive Scheduling



ProActive Scheduling Big Picture

ProActive Scheduler

File Window Help

Scheduler

Jobs

Pending (674) Running (60) Finished (31)

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

Id	State	Progress	# Finished	User	Priority
1313	Running	<div style="width: 25%;"></div>	4/8	user1	Normal
1314	Running	<div style="width: 25%;"></div>	4/8	user1	Normal
1315	Running	<div style="width: 25%;"></div>	7/8	admin	Normal
1316	Running	<div style="width: 25%;"></div>	4/8	user1	Normal
1317	Running	<div style="width: 25%;"></div>	7/8	admin	Normal
1318	Running	<div style="width: 25%;"></div>	4/8	user1	Normal
1319	Running	<div style="width: 25%;"></div>	7/8	admin	Normal
1320	Running	<div style="width: 25%;"></div>	3/8	user1	Normal
1321	Running	<div style="width: 25%;"></div>	7/8	admin	Normal
1322	Running	<div style="width: 25%;"></div>	3/8	user1	Normal
1323	Running	<div style="width: 25%;"></div>	7/8	admin	Normal
1324	Running	<div style="width: 25%;"></div>	2/8	user1	Normal
1325	Running	<div style="width: 25%;"></div>	2/8	user1	Normal
1326	Running	<div style="width: 25%;"></div>	2/8	user1	Normal
1327	Running	<div style="width: 25%;"></div>	2/8	user1	Normal

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 Info Result Preview

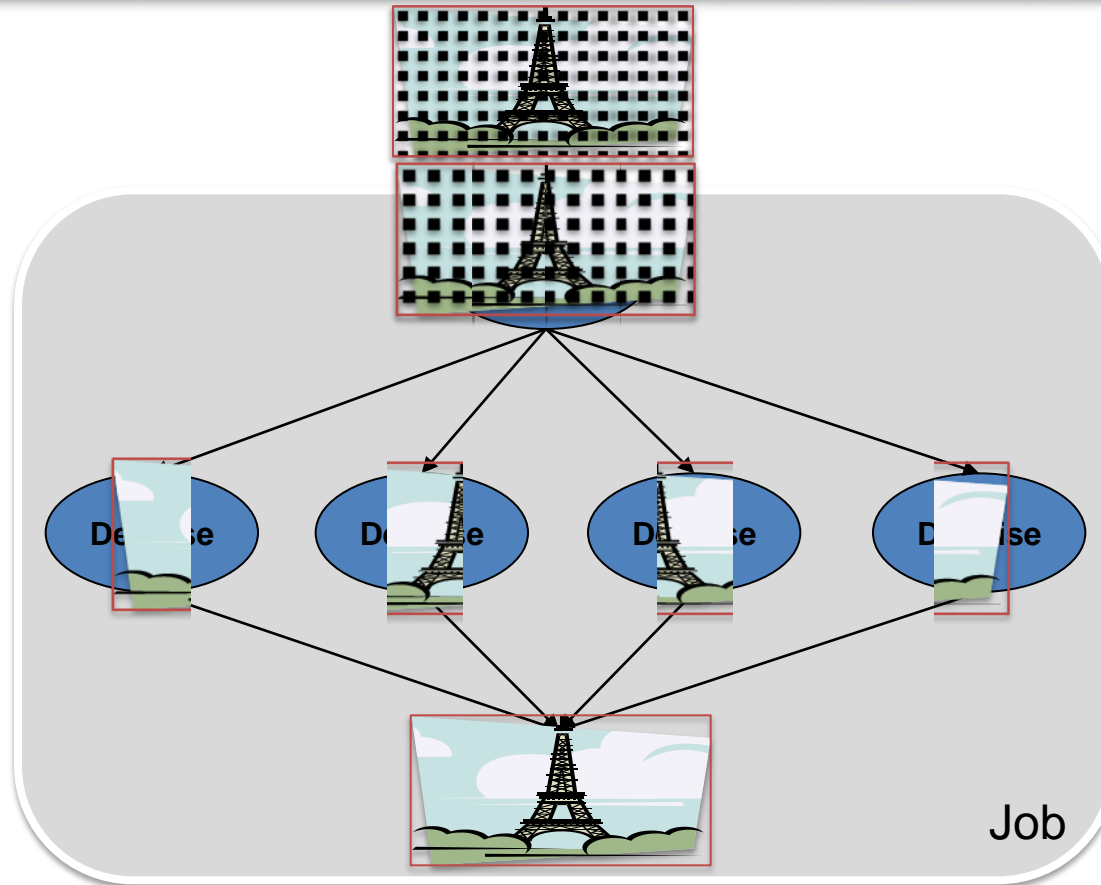
Job 2008 has 8 tasks

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

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

RESO

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



ProActive
Resource Manager



RESOURCING User Interface

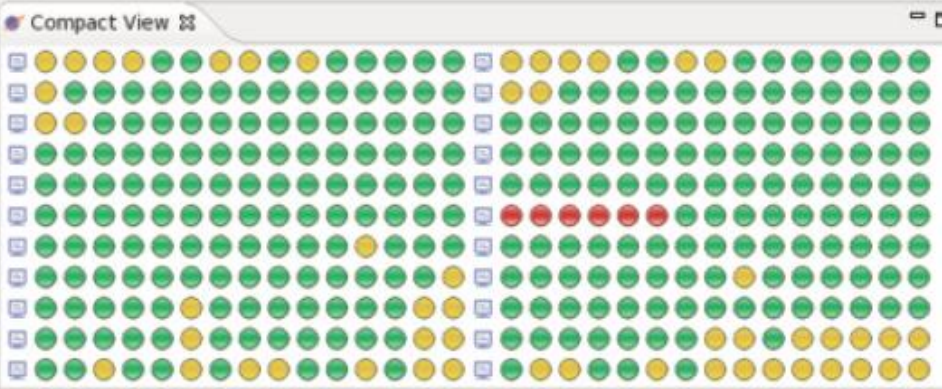
ProActive Resource Manager

File Connection Actions Help Window

Shutdown

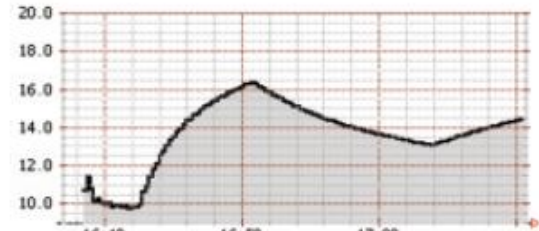
Tab Explorer Tree Explorer Compact View

- PA_JVM2114960478
 - rmi://eon14.inria.fr:1099/PA_JVM2114960478_GCMNode-0
 - rmi://eon14.inria.fr:1099/PA_JVM2114960478_GCMNode-1
 - rmi://eon14.inria.fr:1099/PA_JVM2114960478_GCMNode-2
- PA_JVM477486534
- PA_JVM2003420561
 - rmi://eon14.inria.fr:1099/PA_JVM2003420561_GCMNode-0
 - rmi://eon14.inria.fr:1099/PA_JVM2003420561_GCMNode-1
 - rmi://eon14.inria.fr:1099/PA_JVM2003420561_GCMNode-2

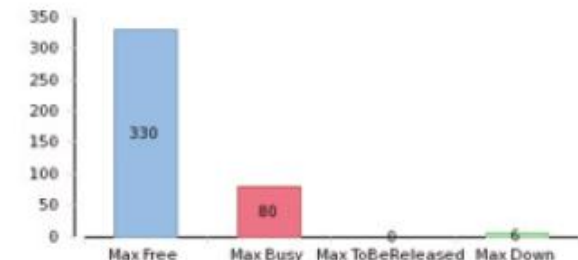


JMX Monitoring

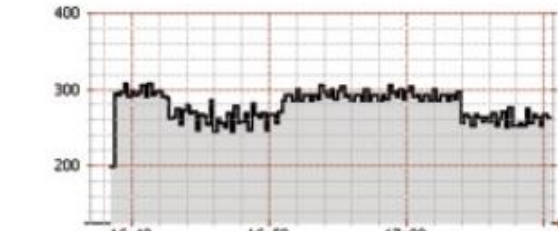
Activity History



Node States Peaks



Free Nodes History



Overview Charts

Statistics Info

state	aggregate
# free nodes	272
# busy nodes	52
# down nodes	6

connected

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

Private, Public & Hybrid Clouds



ProActive Scheduler

ProActive Resource Manager

“Cloud Bursting!”

Static Policy

LSF

Timing Policy
12/24

Desktops

Dynamic Workload Policy

EC2



Dedicated resources



Desktops



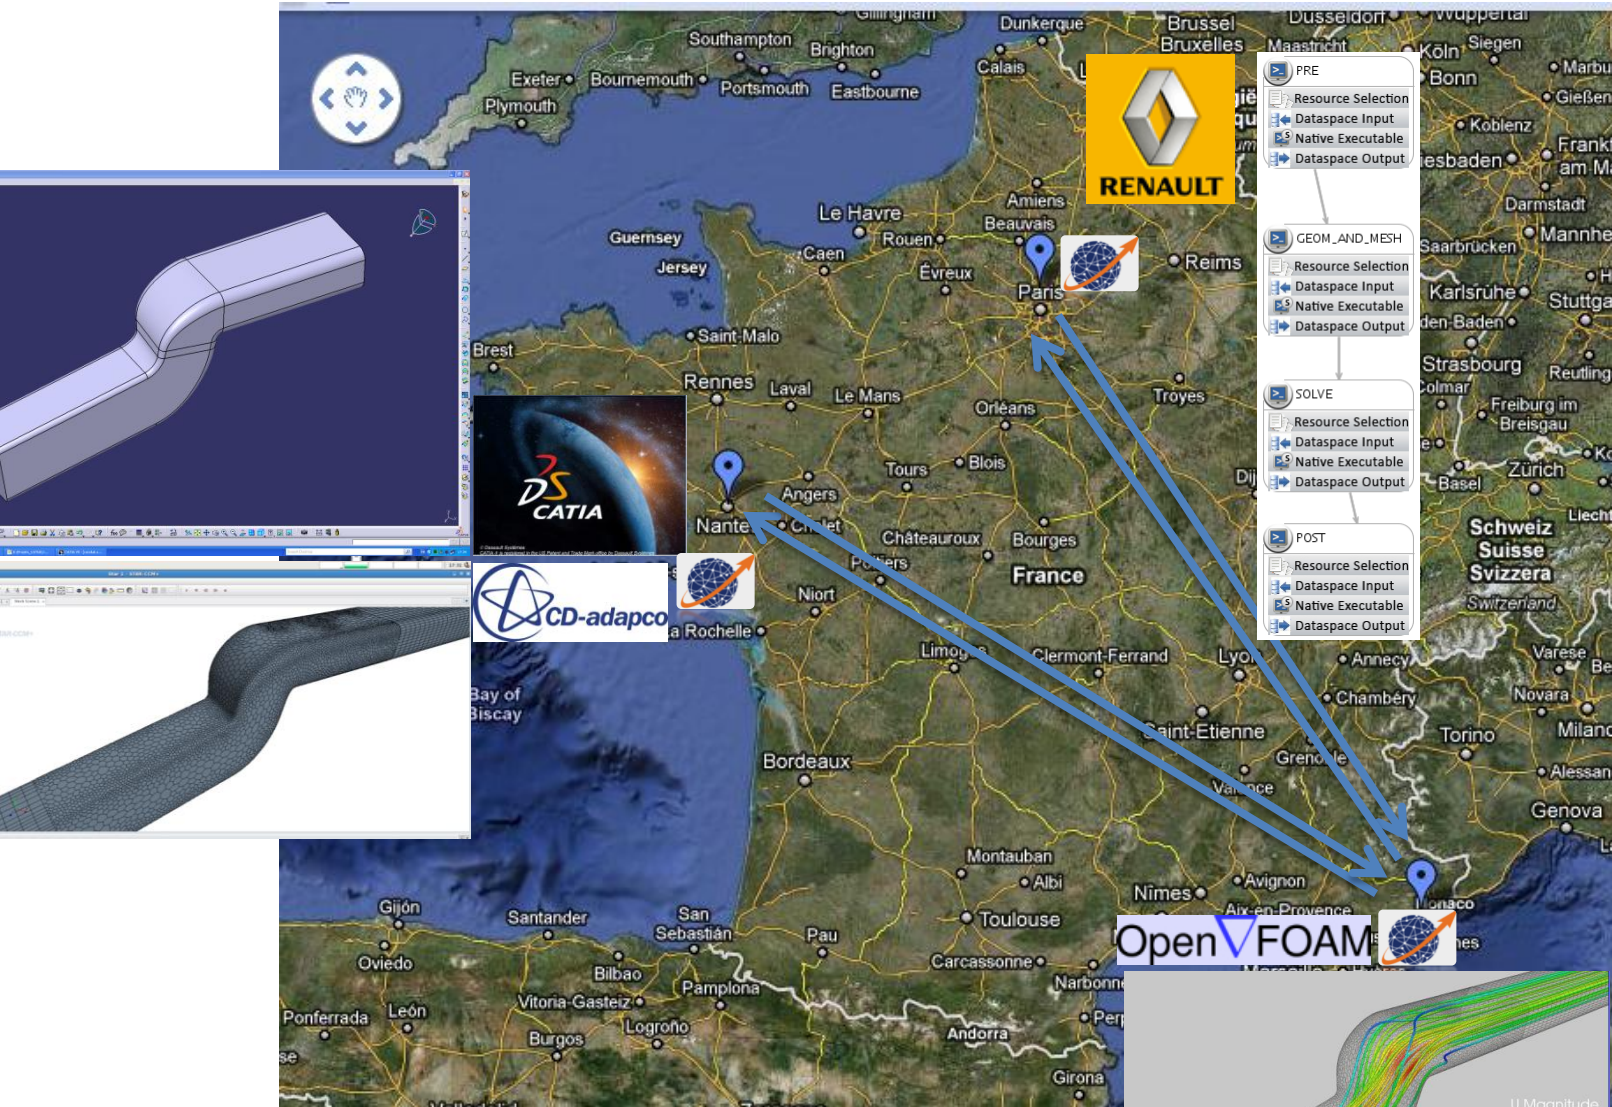
Amazon EC2

Use Case 1: OMD2

Distributed Multi-Disciplinary Optimizations

HPC

ProActive OMD2 Demo



1000 Cores
Production
Cloud Portal



ProActive
Parallel Suite

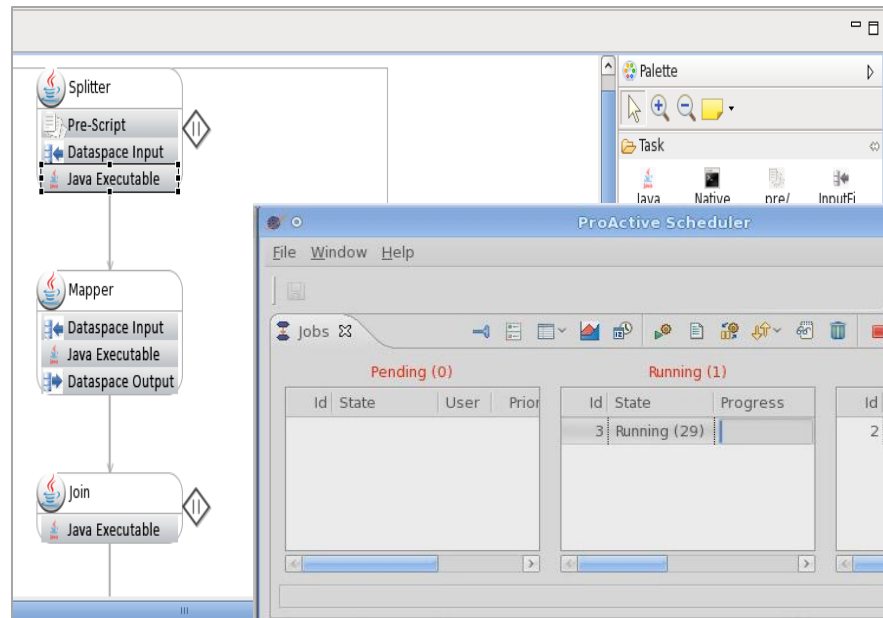
Use Case 2:

Bonita BPM to Clouds with ProActive Parallel Suite on an IT Map-Reduce Application

BonitaSoft + ProActive: BPM Map/Reduce in the Cloud



The screenshot shows the BonitaSoft BPM editor interface. At the top, there's a VLC media player window titled 'bonita_rest.avi'. Below it is the BonitaSoft application window with a menu bar (Process, Edit, Run, View, Extensions, Connectors, Simulation, Help) and a toolbar with icons for New, Open, Save, Print, Import, Export, Copy, and Paste. The main workspace displays a process diagram for 'MyProcessDiagram1 (1.0)'. The diagram starts with a 'Start' node, followed by a 'MapReduce' task, then a 'read results' task, and finally an 'End' node. A left-hand palette contains various BPM elements. At the bottom, there are 'Overview' and 'Details' panels. The 'Details' panel is expanded to show the configuration for the 'MapReduce' task, including a 'Connectors' section with 'PACConnector -- On Finish' selected.



The screenshot shows the ProActive Scheduler interface. The top menu includes 'File', 'Window', and 'Help'. Below the menu is a 'Scheduler' toolbar. The main area displays job status information:

- Pending (0)**: No pending jobs.
- Running (1)**: One job is running. A table shows Job 3 in a 'Running (29)' state.
- Finished (1)**: One job is finished. A table shows Job 2 in a 'Finished' state, completed by user 'demo'.

Below the job status, there are several panels:

- Console**: Shows 'Job 3 has 38 tasks'.
- Statistics**: A table showing system statistics:

Name	Value
# free nodes	6
# busy nodes	34
# down nodes	0
# total nodes	40
- Compact View**: A grid of colored circles representing nodes.

At the bottom right, it says 'connected'.

Video:
BonitaSoft + ProActive:
BPM Map/Reduce in the Cloud



Real Live Demo: ProActive Scheduler & Resource Manager



Conclusion

Industrial (1750) & Cloud Revolution Compared

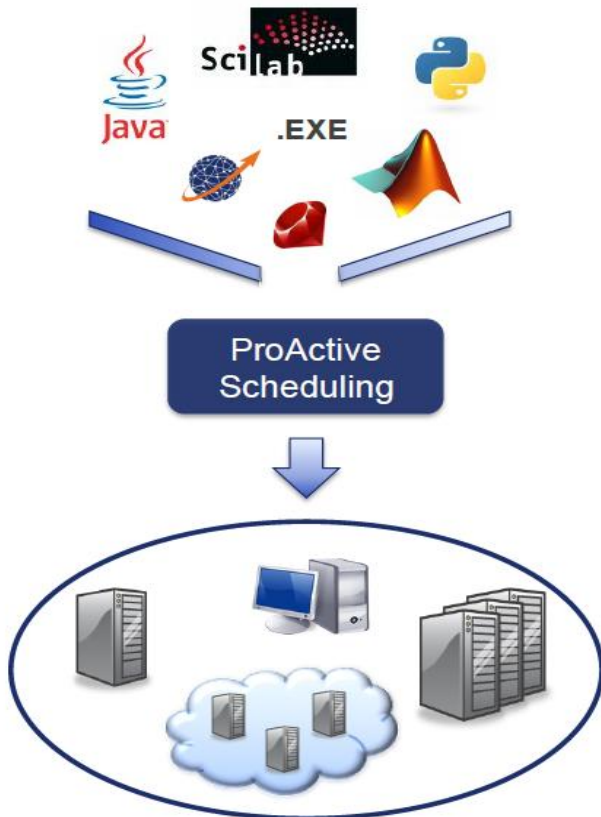
	Industrial Revolution	Cloud Revolution
Concept	Mechanization and centralization of manufacturing activities	Computing as a Utility Centralization of Data Center
Technology	Supporting new technos (Mechanic, Tool Machines, etc.)	Distributed Computing Virtualization Multi-Cores Network
Socio Economical 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

Conclusion: OW2 OSCi initiative

ProActive
Parallel Suite



OW2
Consortium



Cloud Elastic Mashup
PaaS Billing
VM Cloud-testing
Monitoring GreenITSLA
Integration Orchestration Contract
Workflow Server Scale-out
Models Interoperability
Management Autonomic
Distributed Optimization Mediation BPM
Placement SaaS Provisioning
Privacy Composition Data
Energy Application-Server
Auto-scaling Interoperable Grid
Open Manager Flexibility Abstraction
Scalability JavaEE Portal
SOA Application
Constraint Dynamically
IaaS Datacenter
Components Services

4 Strategic Domains

- ❑ Self-sizing and green PaaS
- ❑ Massively Distributed Services
- ❑ BI 4 Cloud
- ❑ Massively Distributed clouds

→ Fine Grain CLOUD management:
→ Pricing at the second (like GSM)
Elastic Clouds

→ Full-Fledged Cloud Solution in OS

ProActive
Parallel Suite



<http://proactive.inria.fr>



Thank you for your attention!