A step back from the fray

Grids and Clouds

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

GridKa School,
05.09.2011

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Distributed Computing

(must at least leave the building!)

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Analogy : Computing ➔Making Coffee
Borrowed from N. Drost,
"Real-World Distributed Supercomputing"

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

“Supercomputing”

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

“Cluster”

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

“Grid Computing”

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

“Cloud Computing”

Coffee Machine Rental

rent your espresso machine here

Home About Coffee Machine Leasing Coffee Shop Discussion News Shipping Rates

Commercial Espresso Machines

We rent out commercial espresso coffee machine packages on flexible terms -no deposit required.

Join our Coffee Rental FORUM to ask questions about renting or contribute your views!
Subscribe for latest news:

Name:

E-Mail Address:

Latest News

Lets keep in touch

Name:

E-Mail Address:

Subscribe

Email Marketing by WP Autoresponder

Searches

- La Marzocco uk price list
- hire large coffee machines
- la marzocco rental
- Coffee Machine Rental
- renting coffee
- free coffee machine for shop
- 2 group espresso machine
- espresso machine rentals in philippines

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Grid Computing in Reality

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Grid Computing: "More Than One"

- More than one machine
- More than one user
- More than one research community
- More than one administrative domain
- More than one geographical location

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

General (our) case: more than one of each!!!

Presentation from 2006

"The information system *IS* the grid!"

Grids are about *federation*.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

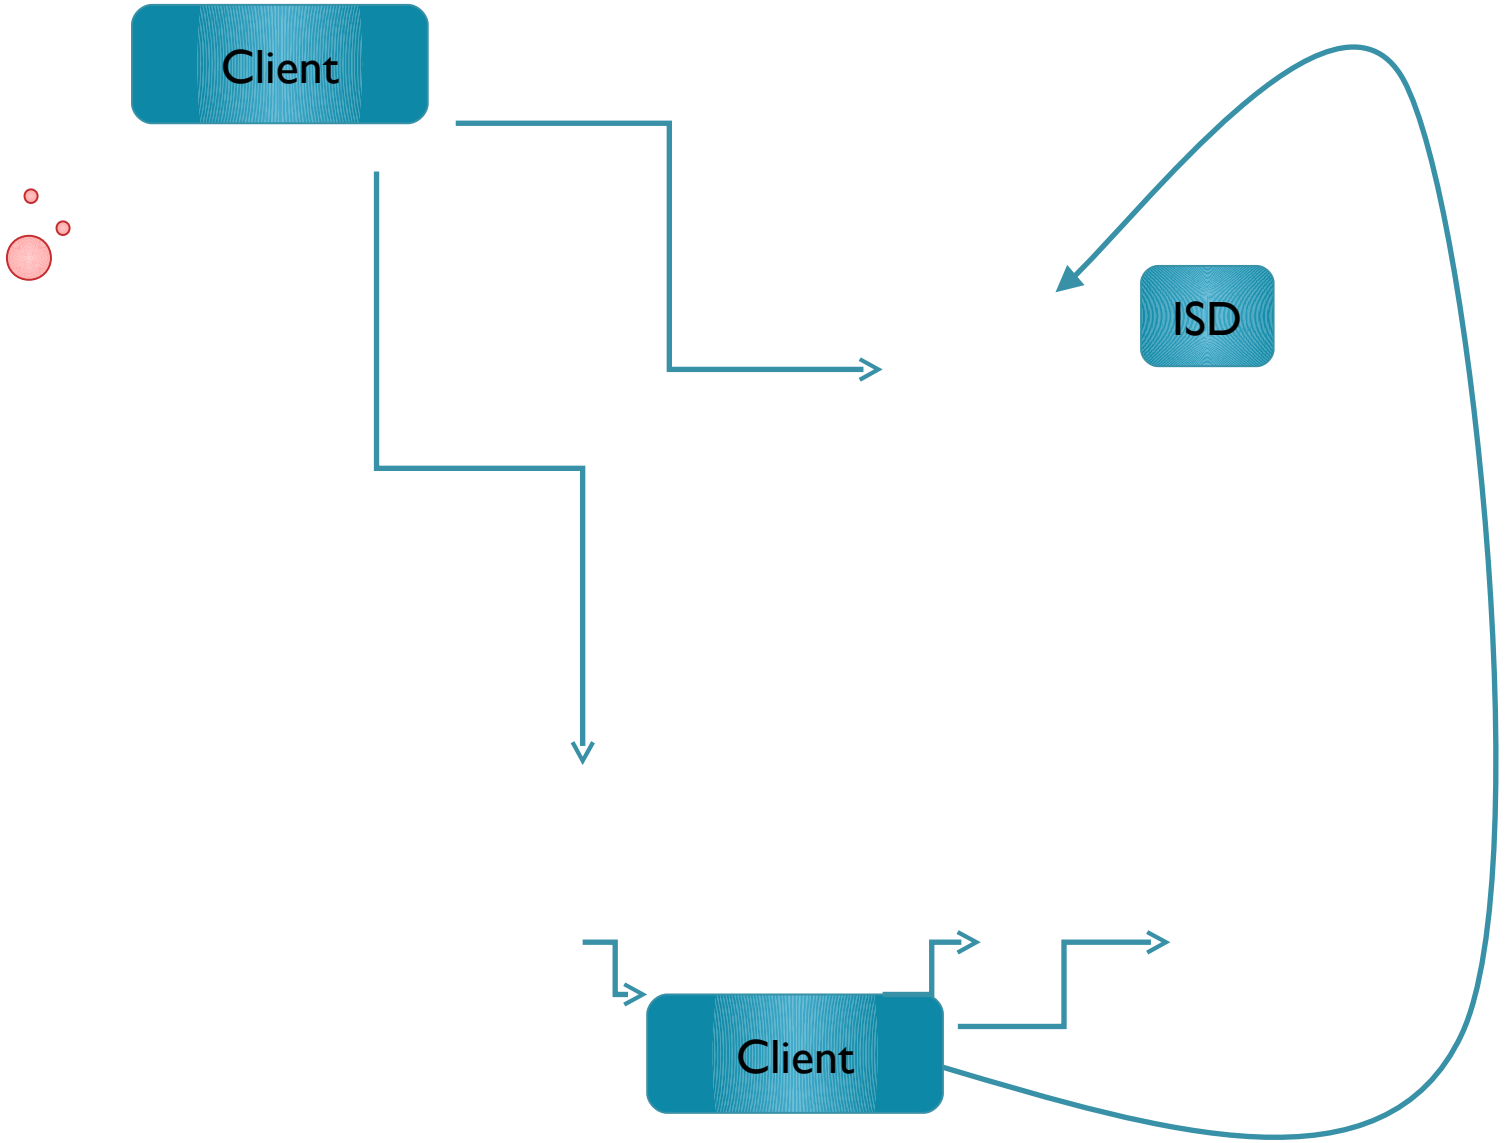

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

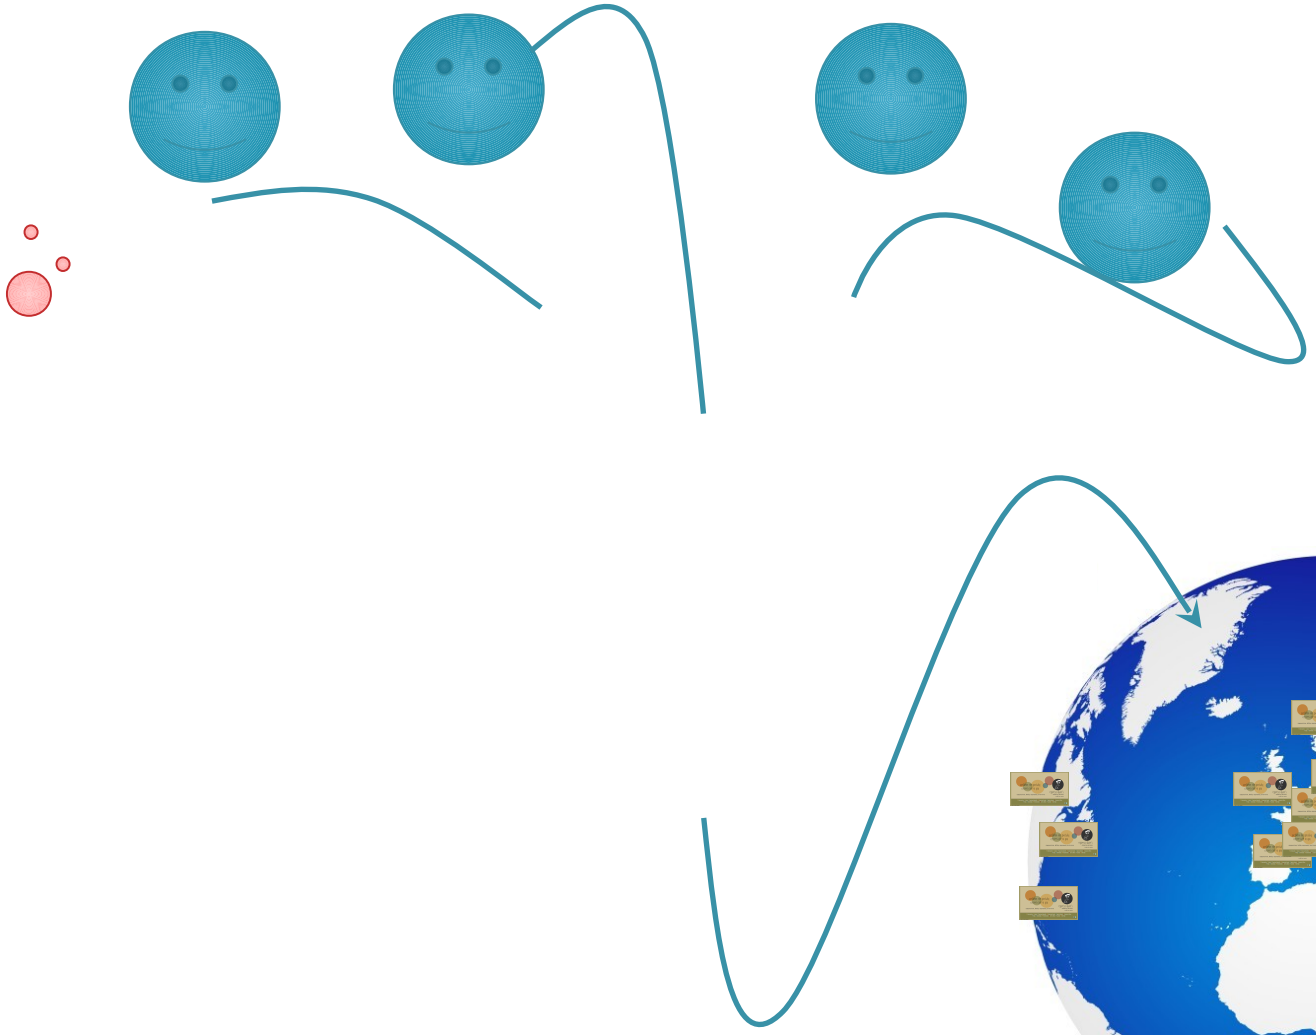

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

"WMS Model"

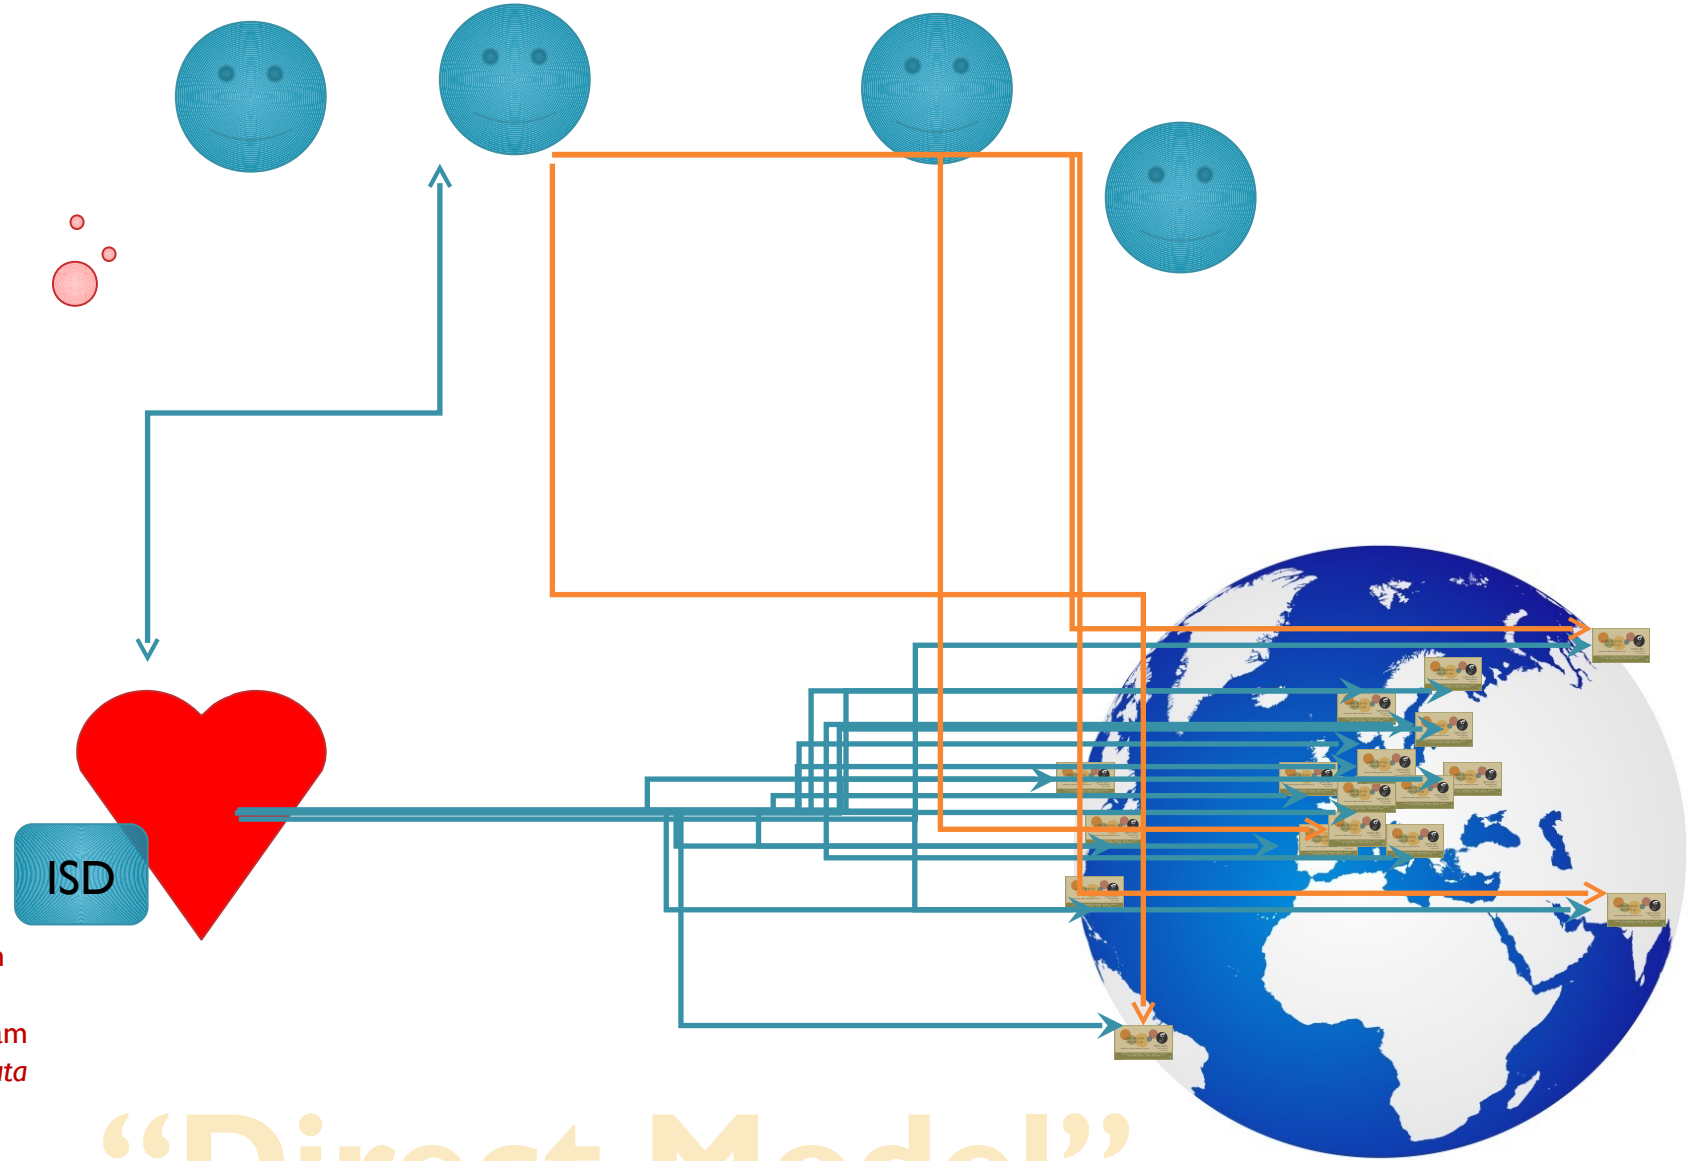

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

"Direct Model"

Data Models

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

WLCG Usage

Stolen From Ian Bird

- Use remains consistently high
  - 1 M jobs/day; >>100k CPU-days/day
  - Actually much more inside pilot jobs

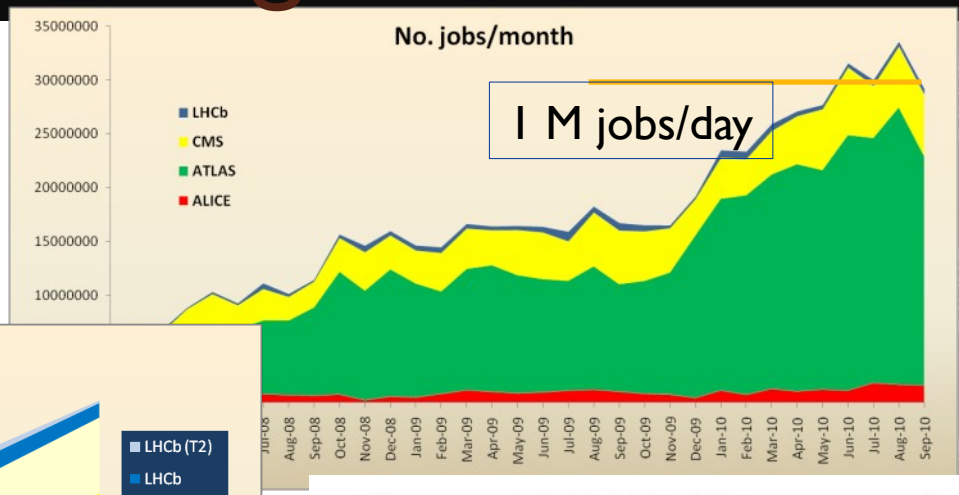

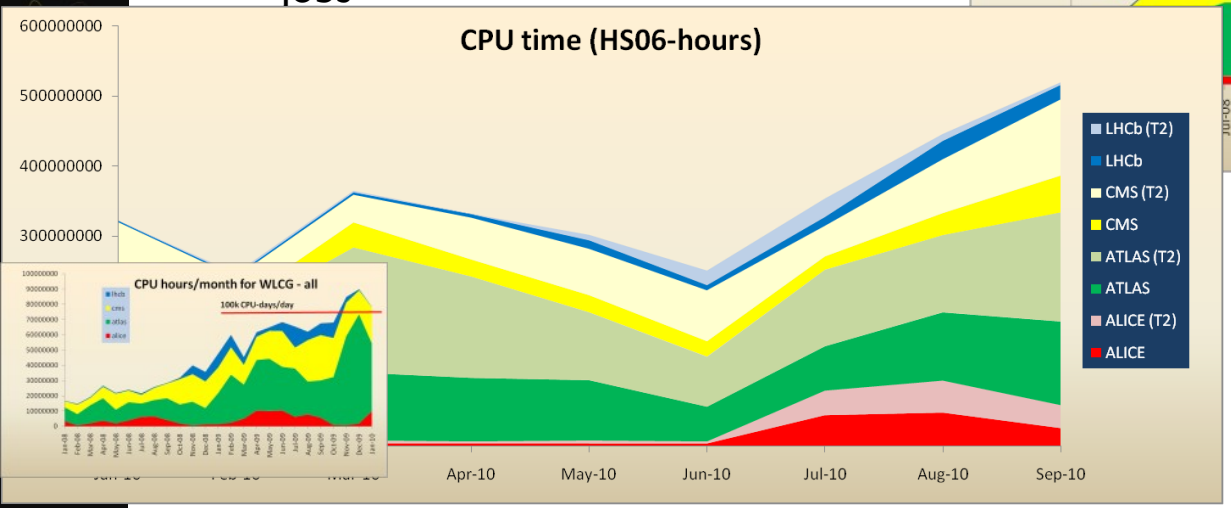

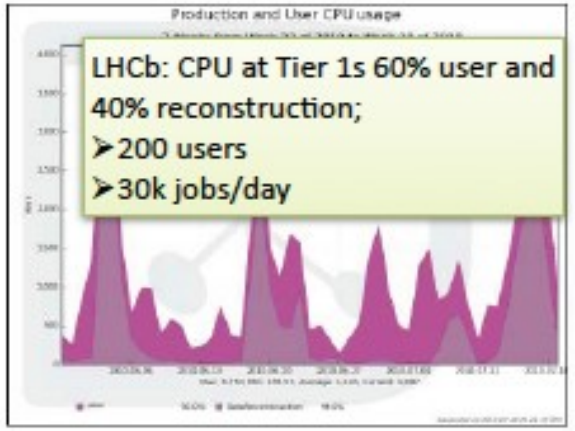

As well as LHC data, large simulation productions ongoing

CMS

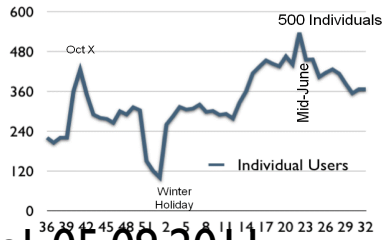

Large numbers of analysis users

CMS ~800, ATLAS ~1000, LHCb/ALICE ~200

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Multiple Communities is the toughest part of "our" grid.

Separation of data & access
Separation of users
Tracking of "whodunit"
Scheduling between groups (esp. during peak demand)
Agreement on OS, sw versions, etc.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Lots of domain-specific "grids"
Why share if you've got enough cash?
Pressure to consolidate …

See R. Pordes talk in Vilnius …

“Cloud Computing”

amazon web services™

Toodledo

Google

mobile me

“public clouds” few field reports
Amazon “renting”, science cloud do we really rent?
What’s different?

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Differences (so far)

- No federation. Cloud "site" is standalone
- Interfaces appear to be *much* simpler
  - Grid : specify everything
  - Clouds : as little as possible (too little?)
- More standardization (although they may still blow it.)

Cloud is about provisioning
Not so interesting?

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Virtualization & Clouds

Vs.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

In words

- Current most common case : user gets "core" on phys machine. Machine (mem, disk, bus, network) shared with other users from other communities.
- Virtualization : user gets an entire "machine" locked to a single core.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Advantages

- Decouple users from each other
- Decouple OS of users from each other and from site!
- Decouple OS space : no attack vectors between communities. Security focused now at site boundary.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Can virtualize other things

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Virtual Networks

- Network properties tagged to user
- Users isolated from each other
- IP address tagged to grid cert : *completely* changes security game.
- Example : virus research in cloud!

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Virtual architecture cont...

From Floris Sluiter, SARA

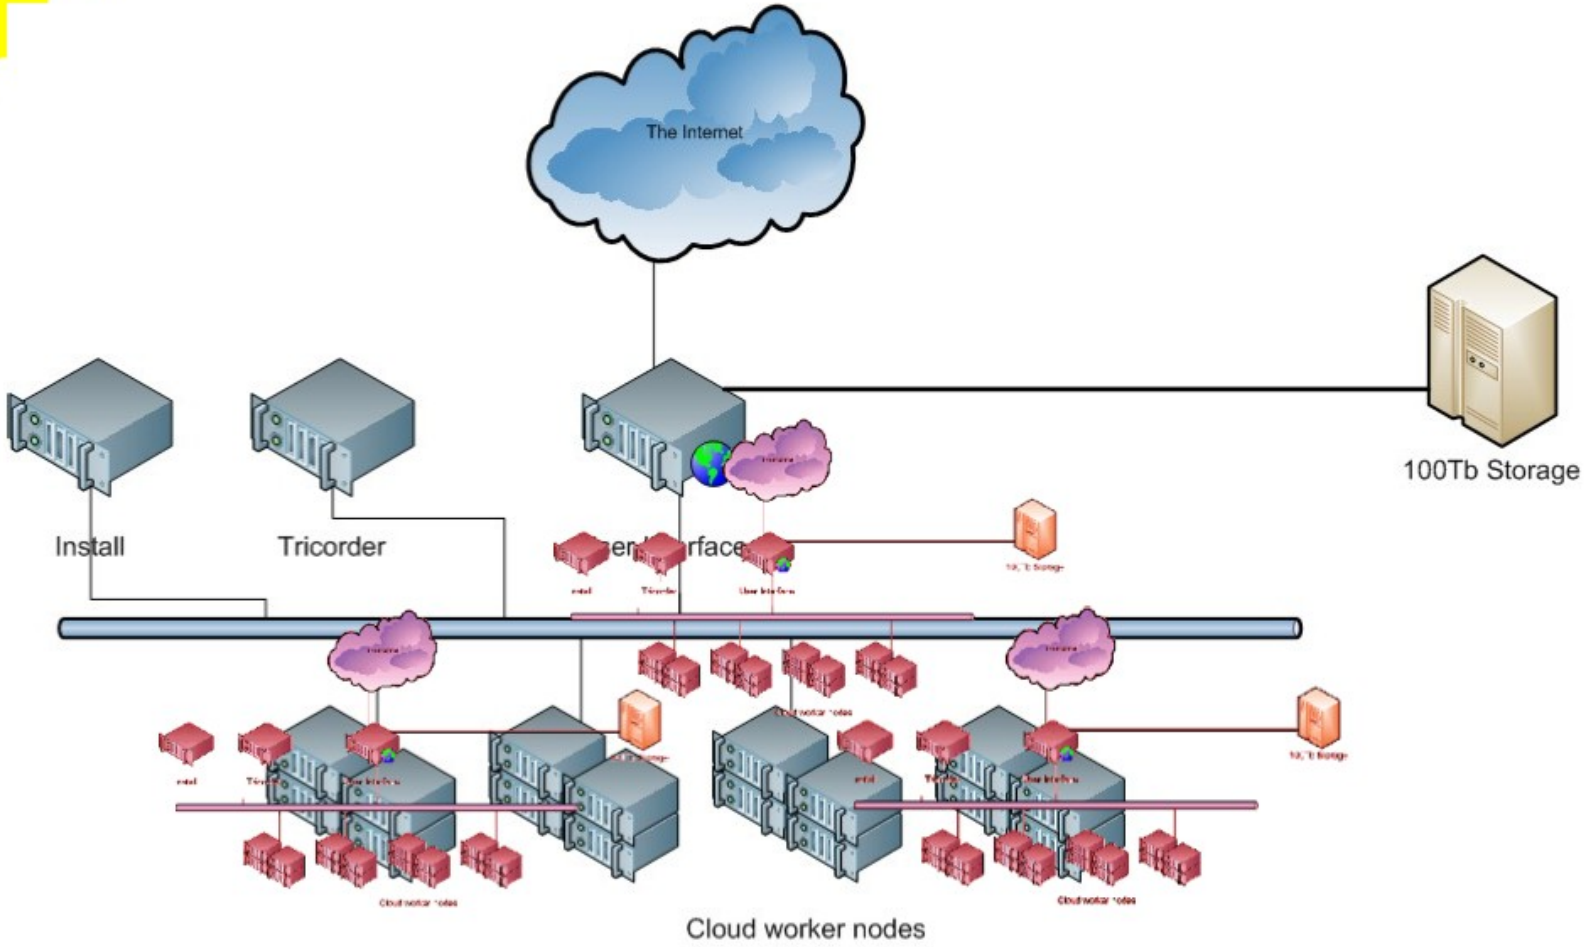

Vision

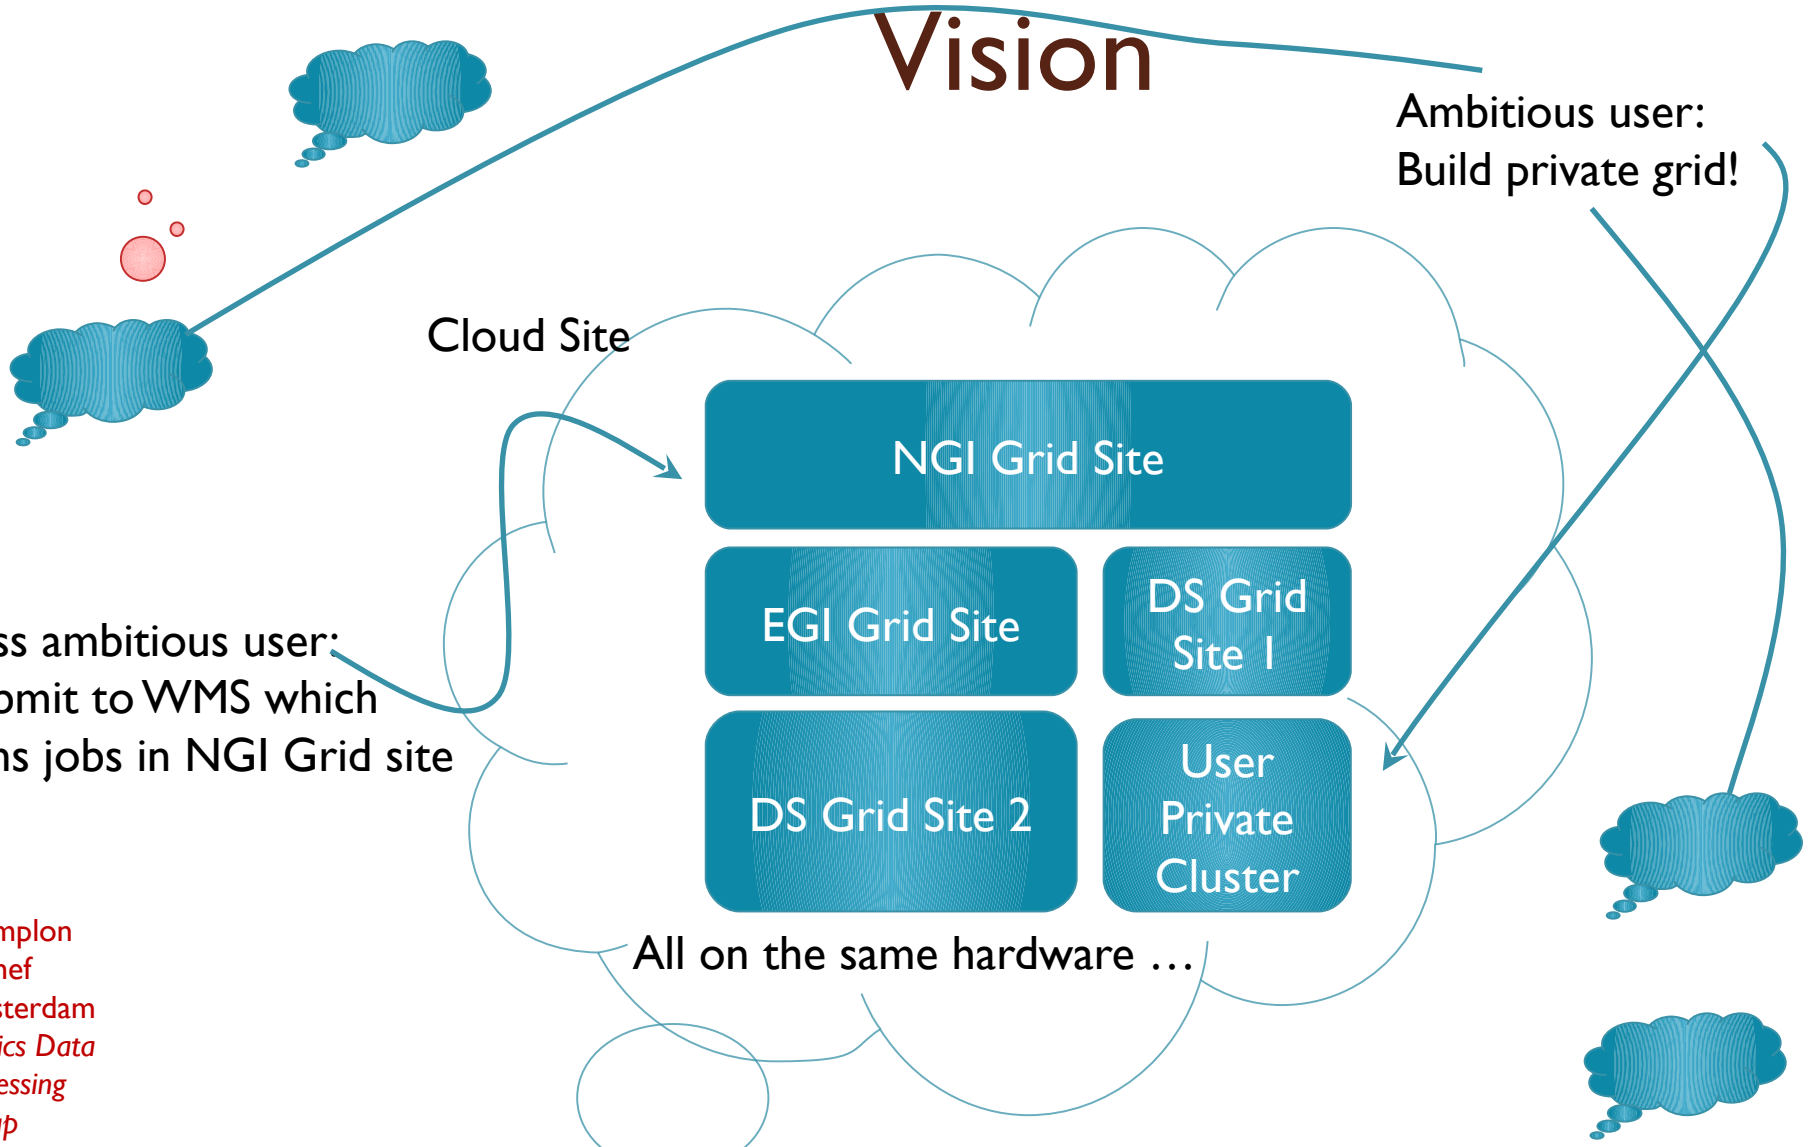

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Talk by Shane Canon @ ScienceCloud 2011

Scheduling!

Enterprise IT versus HPC

	Traditional Enterprise IT	HPC Centers
Typical Load Average	30% *	90%
Computational Needs	Bounded computing requirements – Sufficient to meet customer demand or transaction rates. (i.e. If you gave a typical business free computing, would they suddenly be able to take advantage of it?)	Virtually unbounded requirements – Scientist always have larger, more complicated problems to simulate or analyze.
Scaling Approach	Scale-in. Emphasis on consolidating in a node using virtualization	Scale-Out Applications run in parallel across multiple nodes.

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Cloud Security

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Grid needs Cloud to prosper; Cloud needs Grid to scale.

-- Ian Foster

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Data

- No good models
- Problems
  - Catalogs at various levels : sync
  - Dark Data
  - ‘awkward’ access

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Distributed Storage and Processing Cluster Architecture

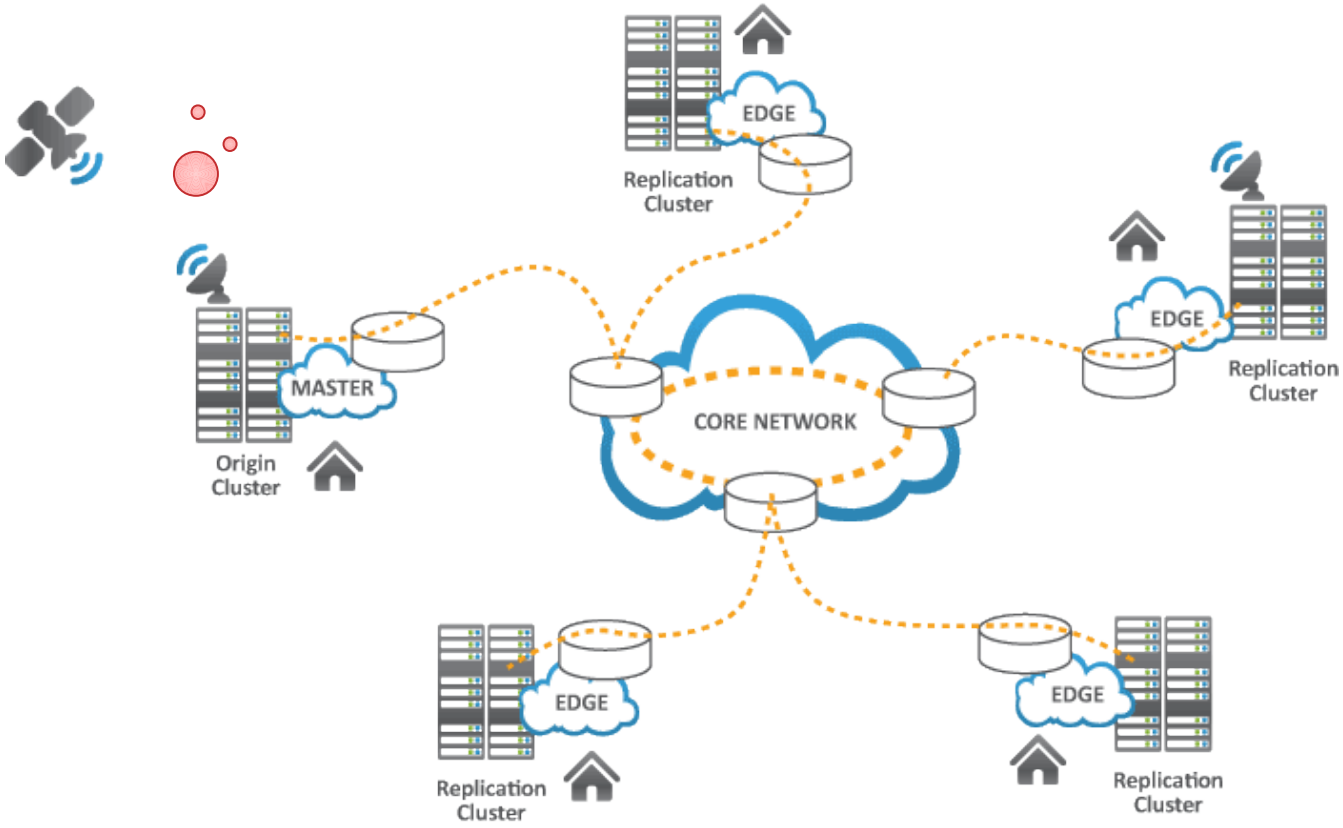

- CVMFS
- Squids
- Xroot/rd
- WebDAV/rd

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Scales

- Overall : probably ok (at least for HEP)
- More cores per box : uh oh
- Multicore apps? Multicore scheduling? Multicore VMs?
- Lots of ways to win in own programs
  - See S. Jarp talk @ CHEP 2010

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Even if we have highly favorable architectures (expensive systems with lots of bandwidth, good resiliency, etc.) I think the community and most, if not all, of the applications are still years away from having algorithms and software implementations that can exploit that scale of computing efficiently. – Andrew Jones, NAG

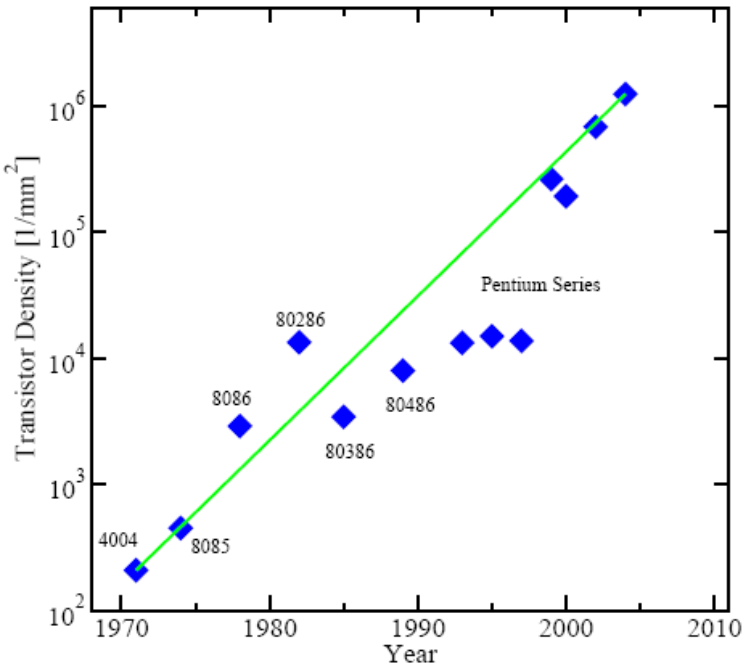

(a) Transistor integration density per die

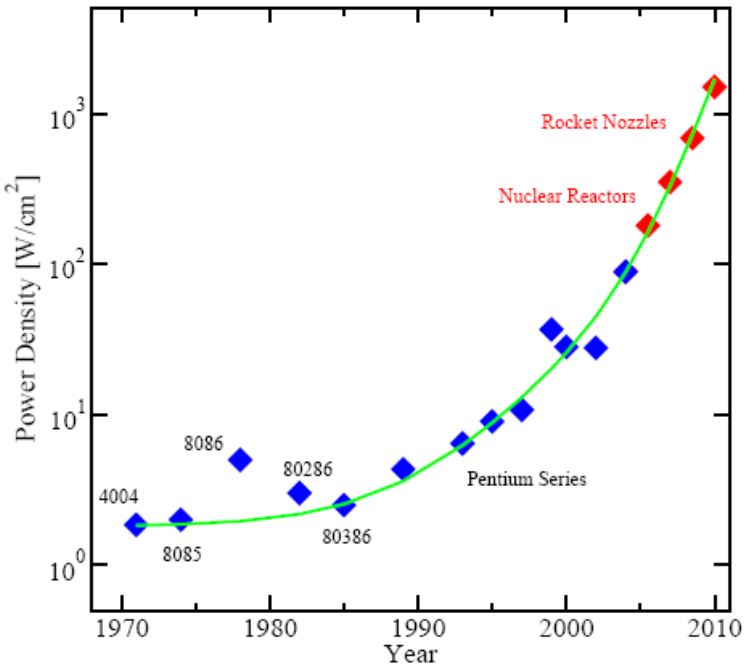

(a) Power loss density per die

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

Shift in R&D : low-power chips for mobile devices!

See B. Panzer talk @ DESY WLCG workshop.

Distributed Computing Dreams

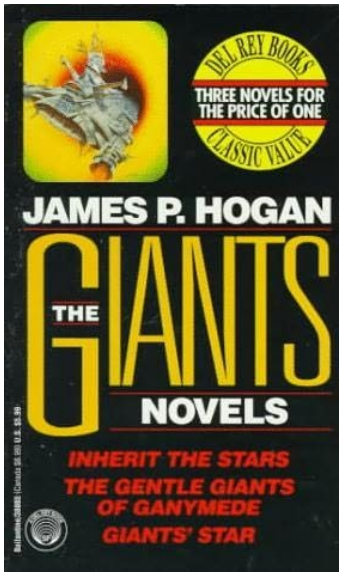

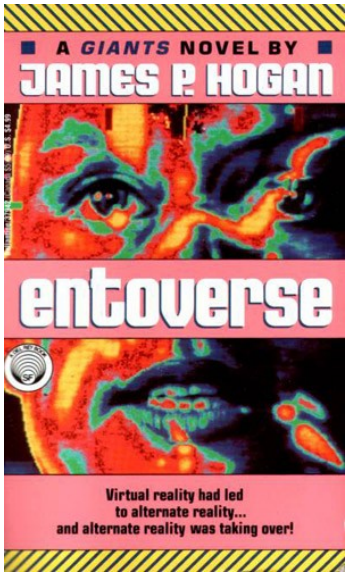

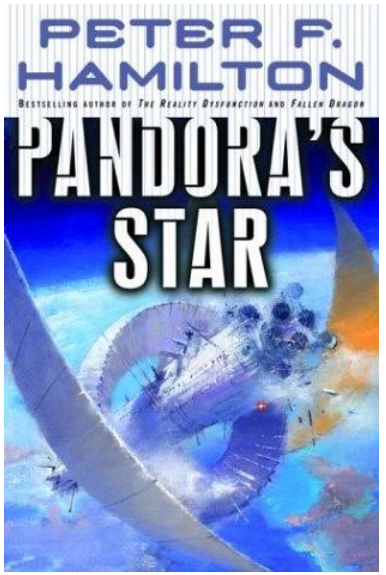

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group

J. Templon
Nikhef
Amsterdam
Physics Data Processing Group