CERN
European Organization for Nuclear Research
Organisation Européenne pour la Recherche Nucléaire

Grid and Cloud Computing
an Overview

Tony Cass
Leader, Database Services Group
Information Technology Department

6th September 2010

Disclaimer and Approach

- This **is not** a rigorous, scholarly discourse on the two approaches to computing service provision.
- This **is** a pragmatic and biased description from someone with long experience as a provider of CPU and data services to a specific community (HEP…).
- Questions asked
  - What problem is being solved?
  - Who owns the resources?
  - What is the enabling technology?

Outline

- Grids
- Clouds
- Analogies
- A Digression
- Problems
- Synthesis?
- Summary

Grids

The skeptic can be forgiven for wondering if there is more to the Grid than … a ‘funding concept’

Ian Foster, “What is the Grid? A Three Point Checklist”

Concept Crystallised

"A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive and inexpensive access to high-end computational capabilities"

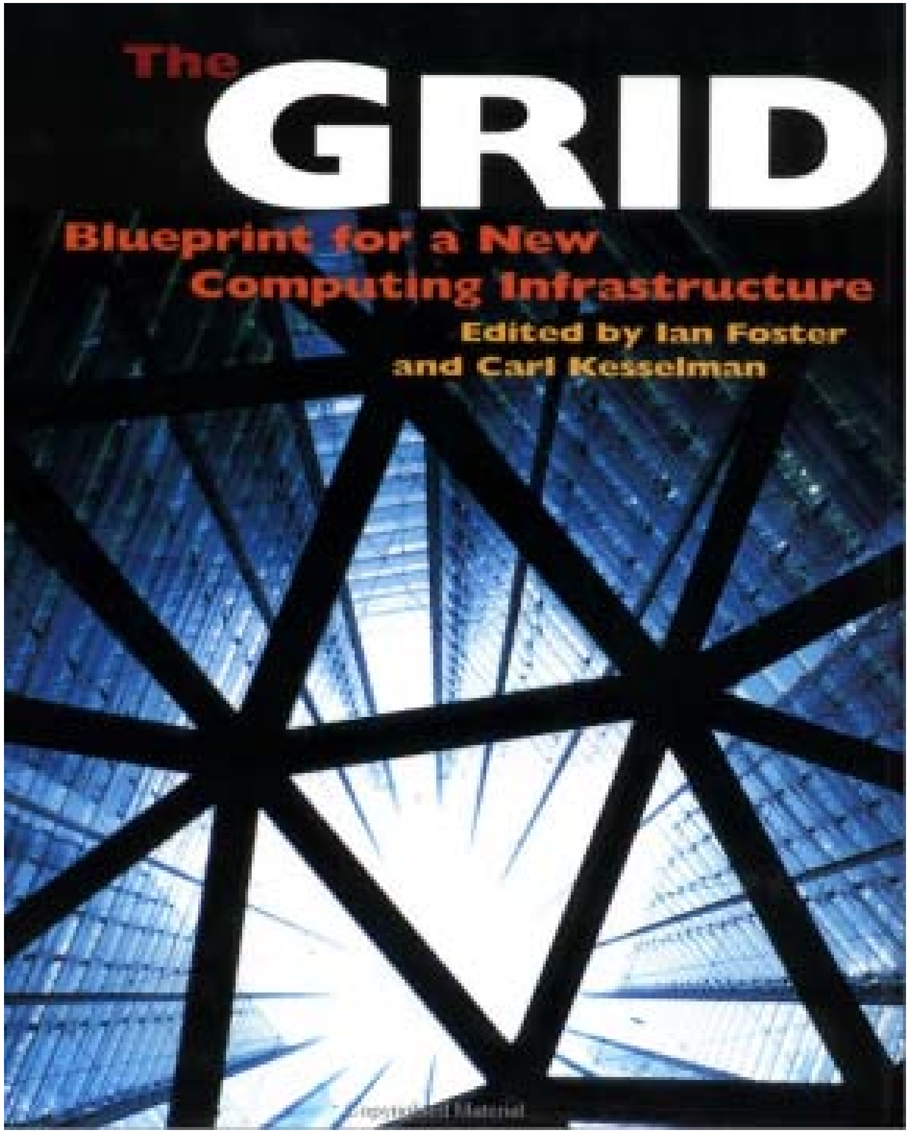

What problem is being solved?

- Clearly: "**access to high-end computational capabilities**"
  - In 1990s, many supercomputing and national computing centres existed, but these were disconnected from researchers. Users needed a dedicated account at each centre they wanted to use and had to connect to a front-end system.
  - The Grid model foresaw integration of such facilities with the researcher's local environment: the end-user's PC plugged into electricity on tap via the power socket and computing on tap via the network socket.

Who owns the resources?

- **The end-user**
  - Perhaps not directly, but certainly indirectly: end-users have a right to use the facilities that exist either because they have been given a grant or because they are part of a community that has purchased the resources.
  - Interestingly, “The Grid” also presents proposals for integration of data-collection systems (e.g. electron microscopes) into the grid as well as data processing (and data storage) systems.

Enabling technology?

- **Affordable high-speed networking.**
  - From “The Grid”: “Furthermore we are starting to see the construction of large-scale, high-performance national- and international-scale networks.” and “Existing high-speed networks ... have enabled extensive experiments with ... the coupling, *usually by heroic efforts*, of geographically distributed supercomputers.”
    My emphasis.
  - Grid middleware was designed and introduced to remove the need for heroic efforts and place high-end computational capabilities within easy reach for mere mortals.

The Grid in Pictures

Interconnected Computing Centres

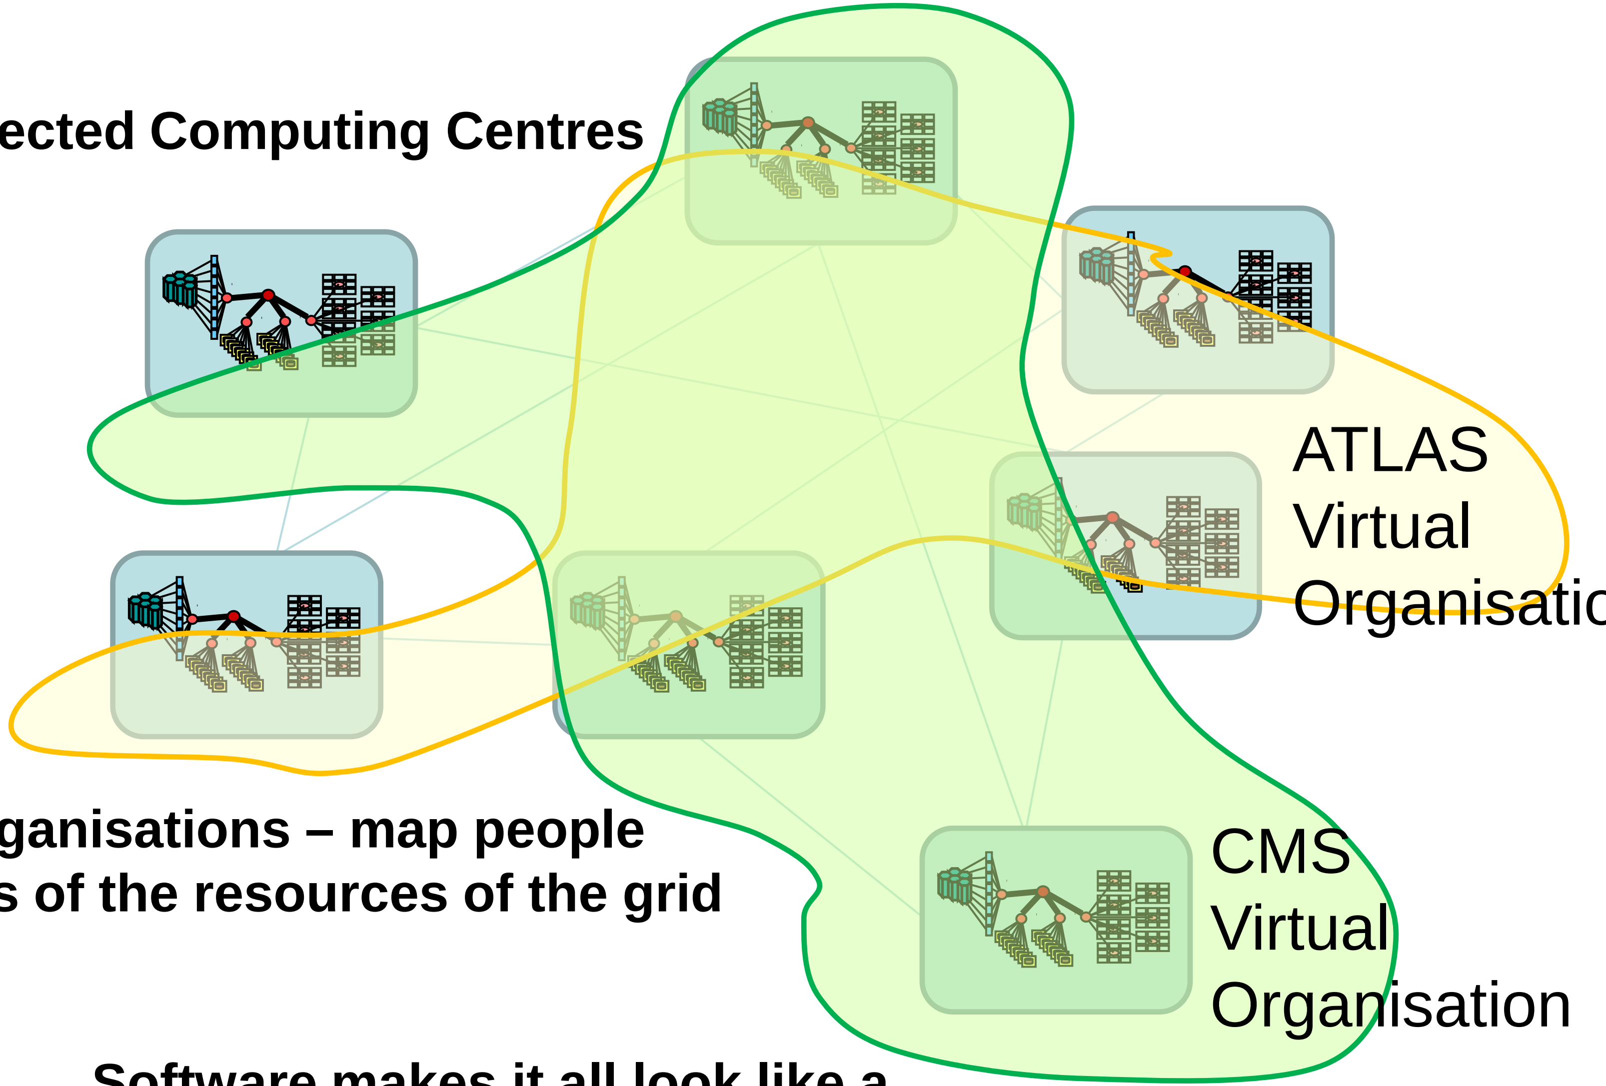

Virtual Organisations – map people to subsets of the resources of the grid

Software makes it all look like a single computing service for each Virtual Organisation

The Grid in Action

Graphics and animation courtesy of André-Pierre Olivier

Grid Components

- Middleware
  - For the various components of the previous slide
    - Compute & Storage Elements
    - Information System
    - Resource Broker / Workload Management System
    - User Interface
    - Logging
  - Many different implementations developed, although mostly with reliance on the Globus toolkit
    - gLite, ARC, iVDGL, OSG, UNICORE
    - Different implementations do interoperate (c.f. wLCG), but not 100% seamlessly.
- TRUST

International Grid Trust Federation

- The international community is deploying large scale distributed computing grids on a production scale, across organisations, across countries, and across continents, for the advancement of science and engineering. In shaping this common grid infrastructure, many of these grids are relying on common practices, policies and procedures to reliably identify grid subscribers and resources.
- The International Grid Trust Federation (IGTF) is a body to establish common policies and guidelines between its Policy Management Authorities (PMAs) members and to ensure compliance to this Federation Document amongst the participating PMAs. The IGTF does not provide identity assertions but instead ensures that within the scope of the IGTF charter the assertions issued by accredited authorities of any of its member PMAs meet or exceed an authentication profile relevant to the accredited authority.

IGTF International Grid Trust Federation AP|EU|TAG

APGrid PMA Asia-Pacific Grid Policy Management Authority

eu gridpma

TAGPMA

Ian Foster's Checklist

A Grid is a system that

1. coordinates resources that are not subject to centralized control ...
2. ... using standard, open, general purpose protocols and interfaces ...
3. ... to deliver nontrivial qualities of service.

Worldwide resources

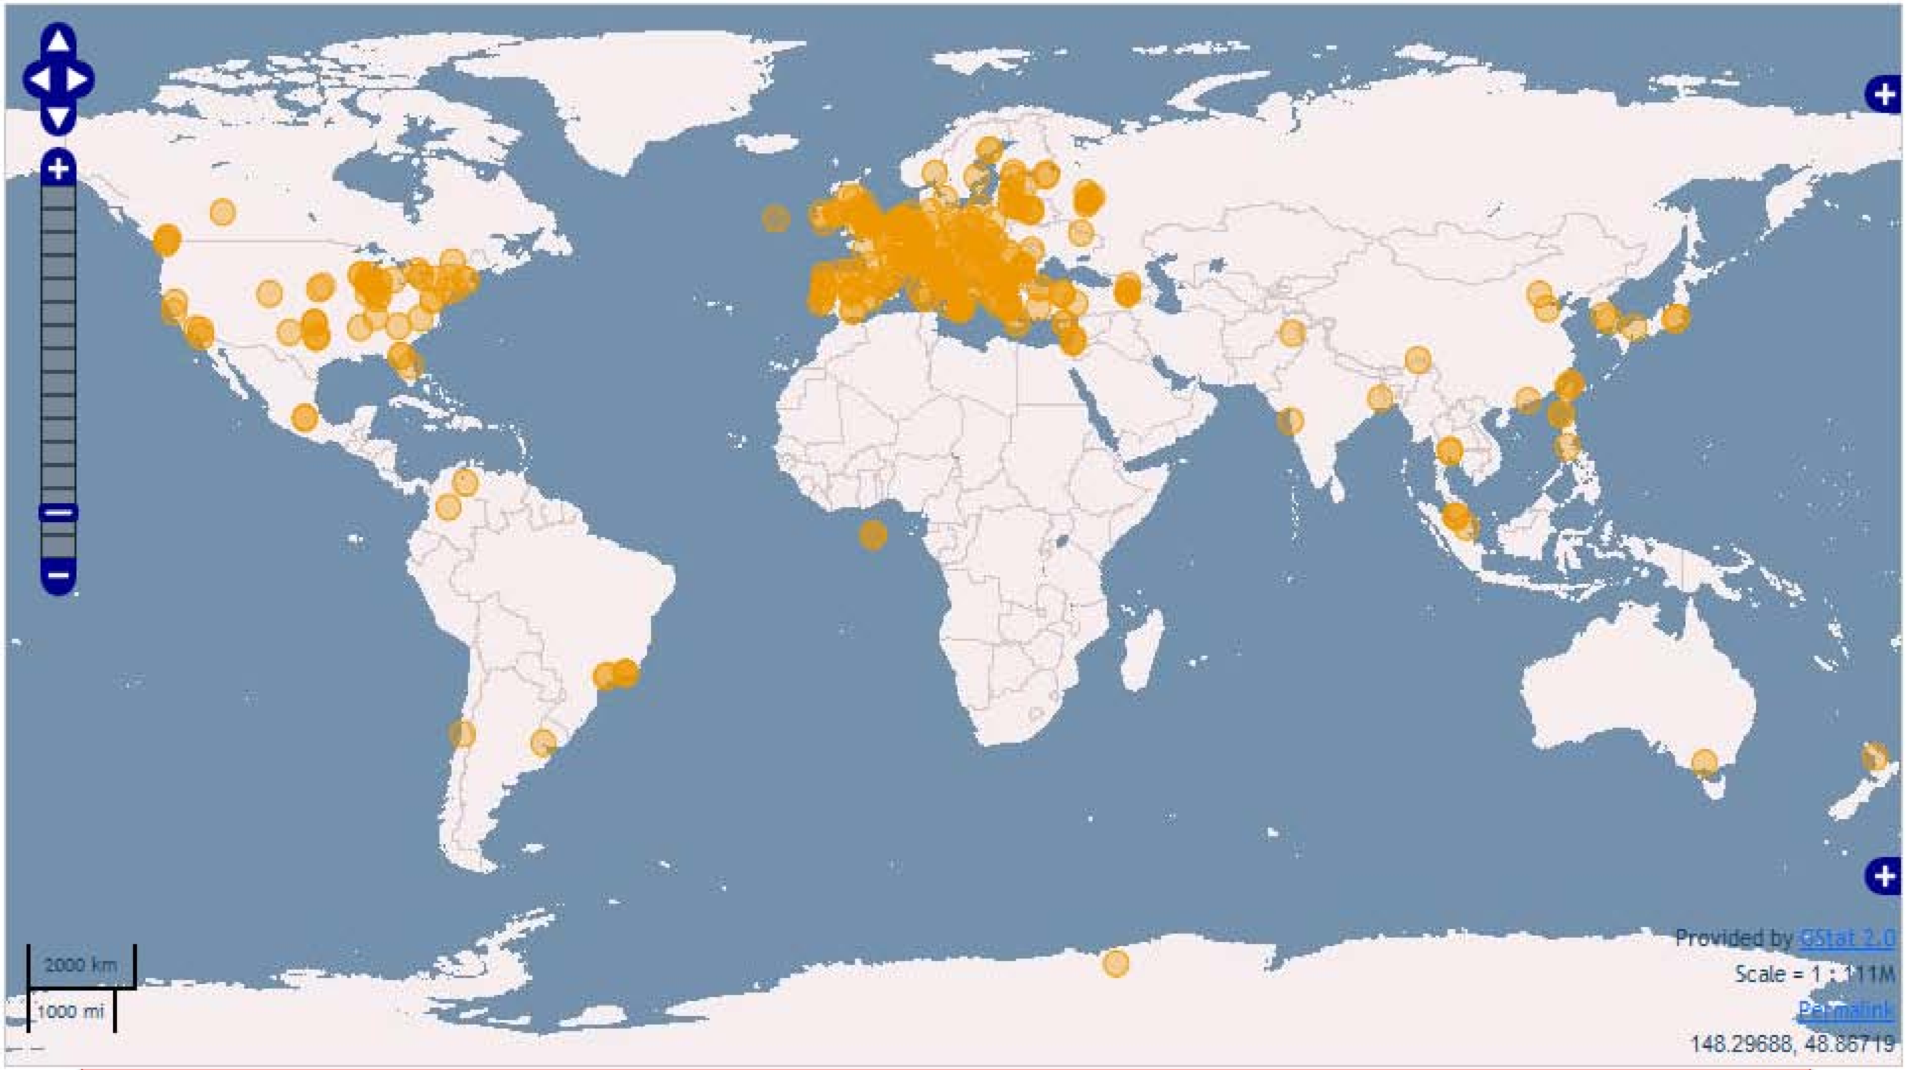

Use of CPU …

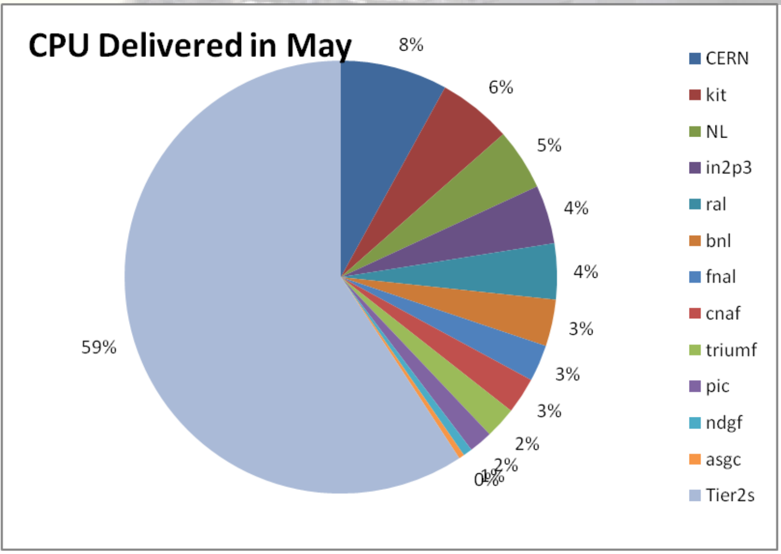

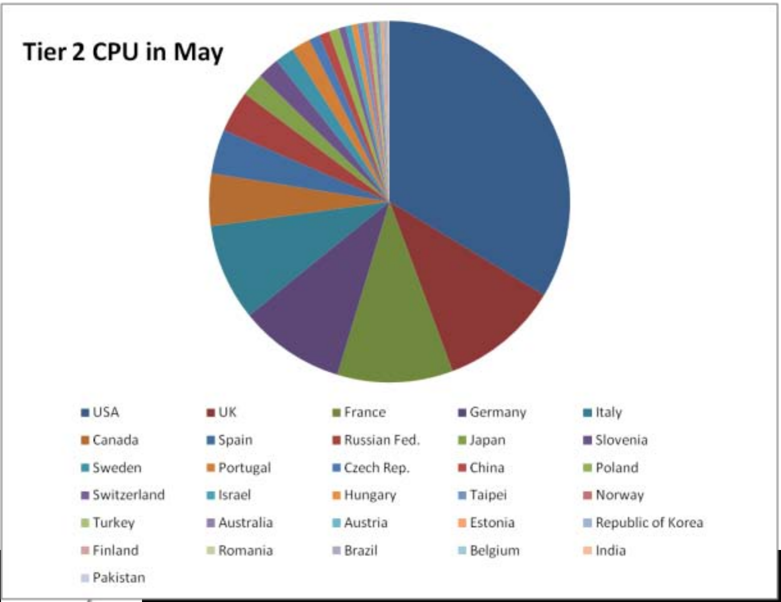

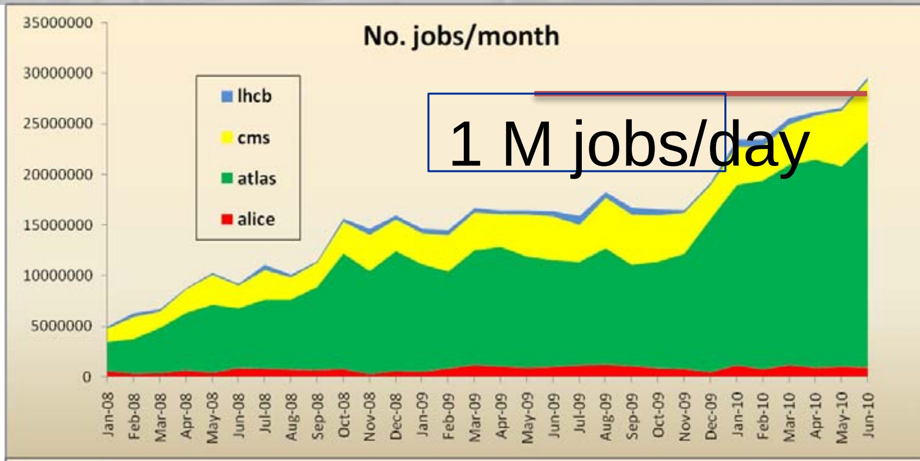

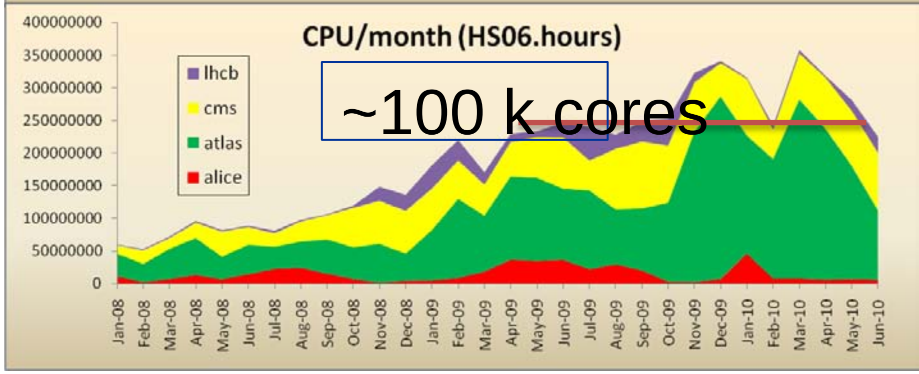

- Peaks of 1M jobs/day now
- Use ~100k cores equivalent
- Tier 2s heavily used wrt Tier 1s

In terms of data transfers ...

Final readiness test (STEP'09)

Preparation for LHC startup

LHC physics data

2009: STEP09 + preparation for data

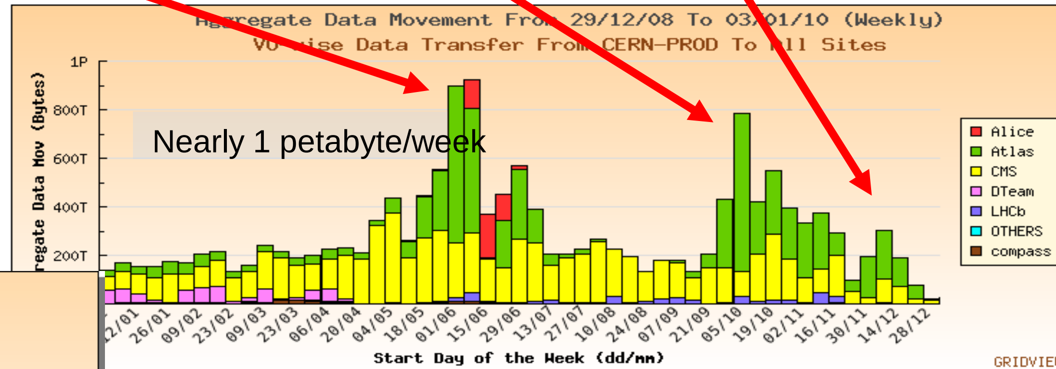

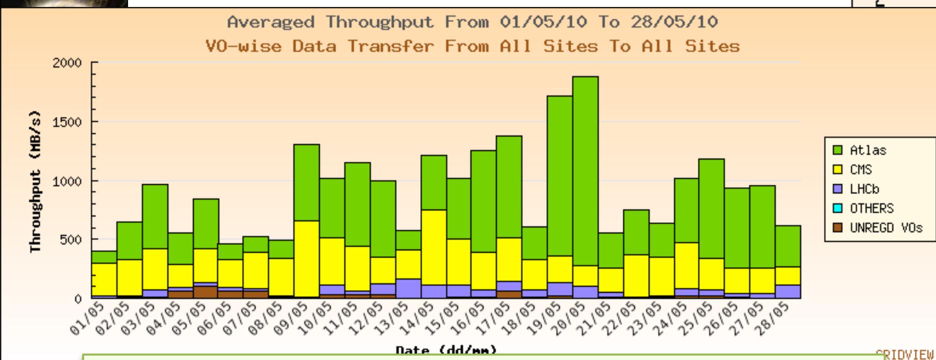

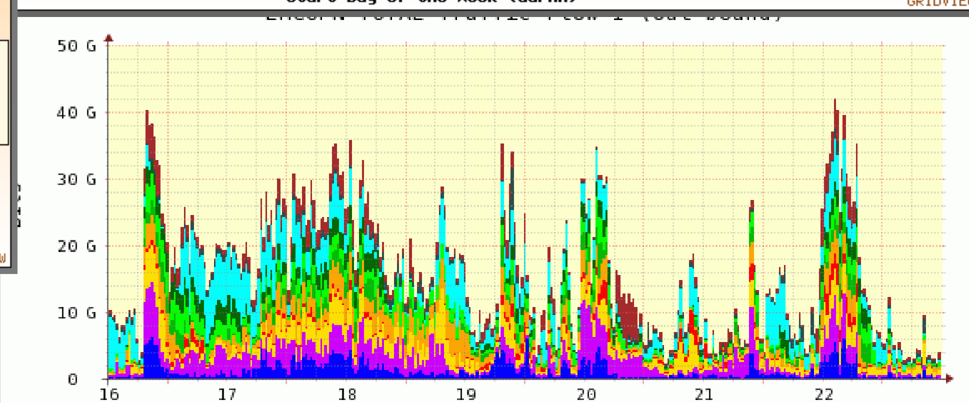

Data export during data taking:

- According to expectations on average

Tier 0 traffic:

> 4 GB/s input

> 13 GB/s served

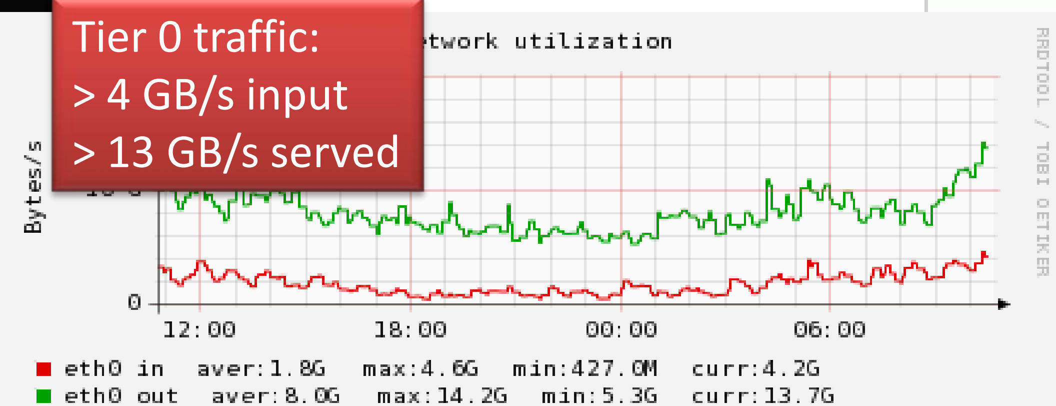

Traffic on OPN up to 70 Gb/s!

- ATLAS reprocessing campaigns

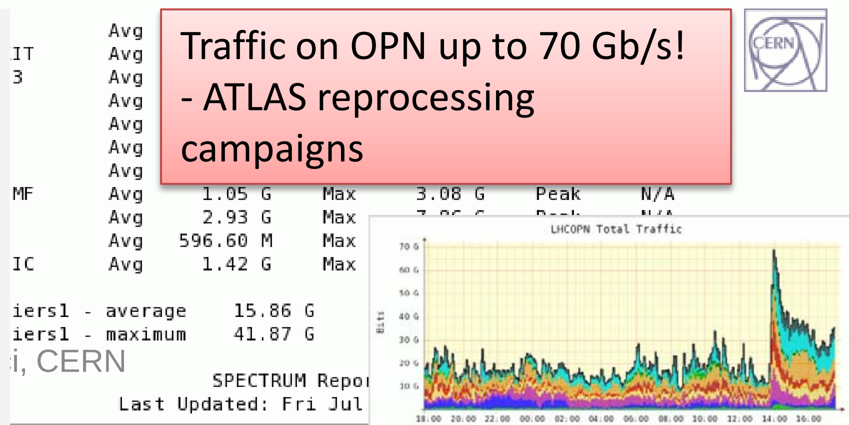

Site availability and readiness

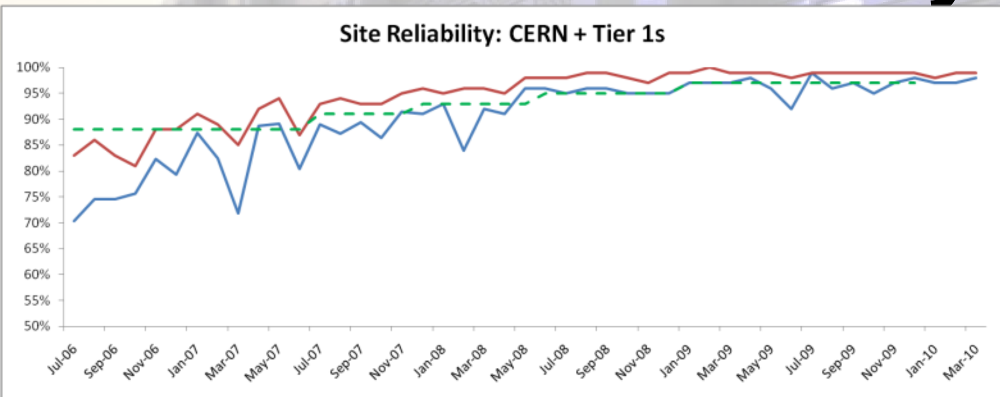

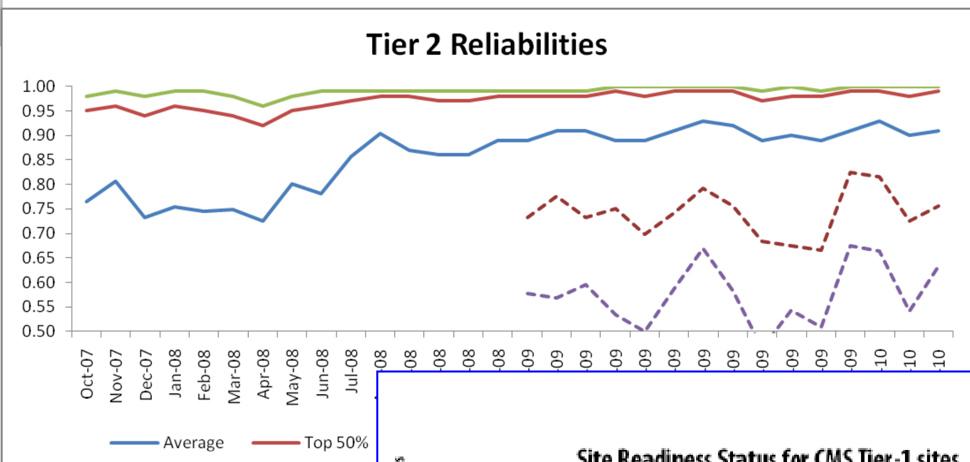

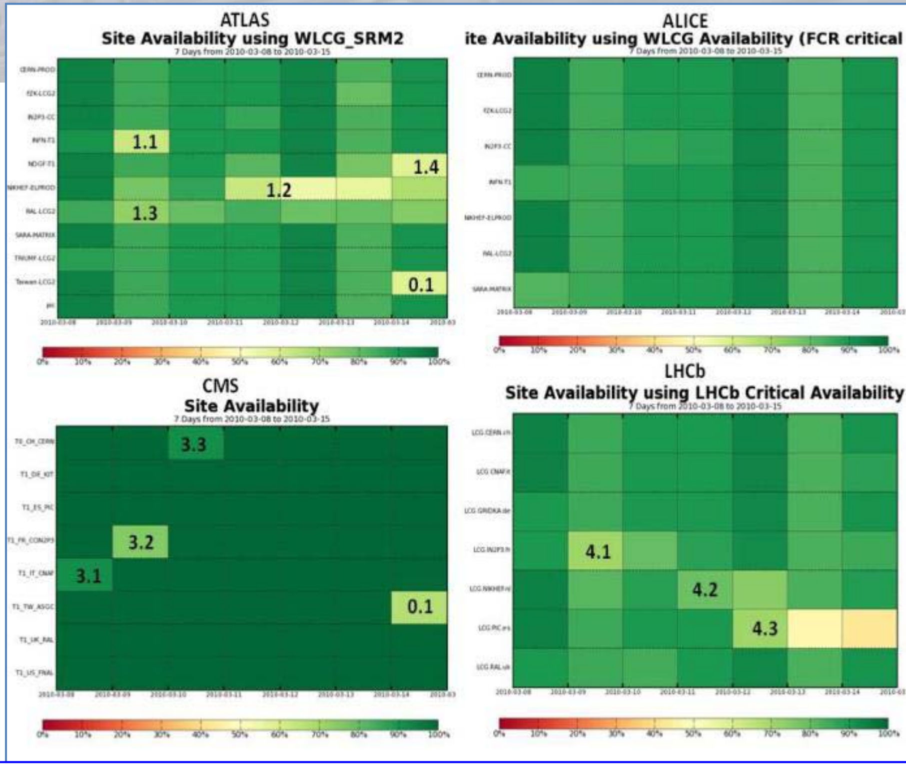

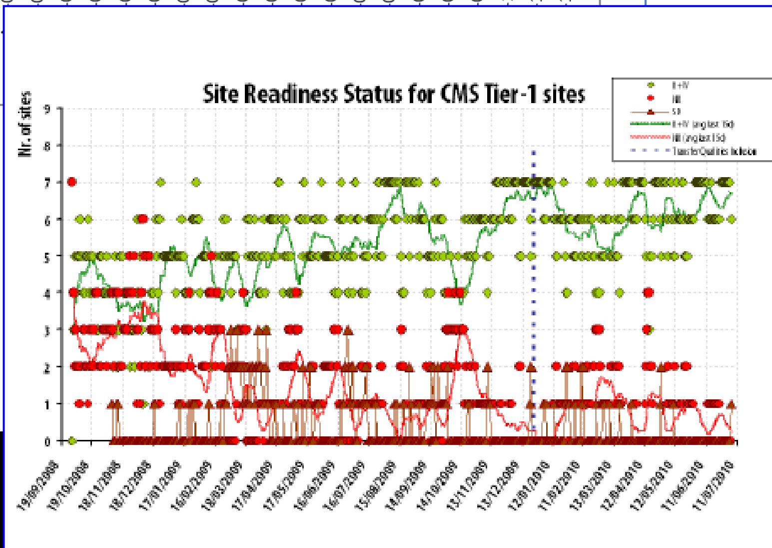

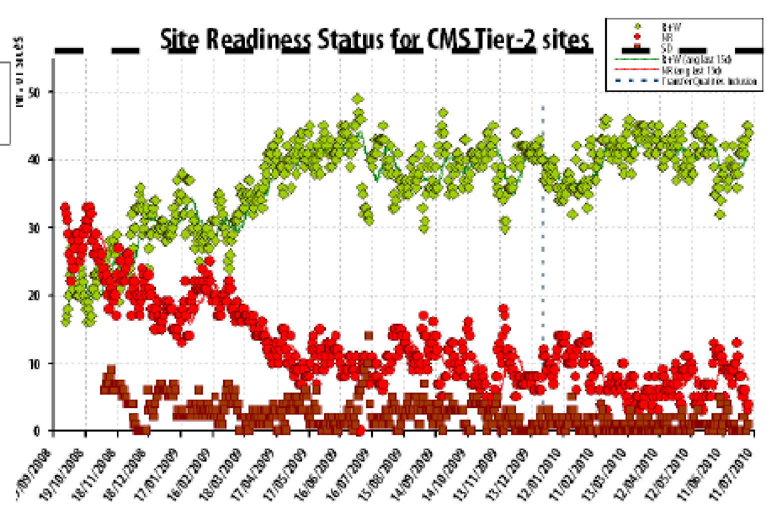

Progress …

- Overall is clear – physics output in very short time
- Huge effort: Combination of experiment sw & computing and grid infrastructures
- And a lot of testing !

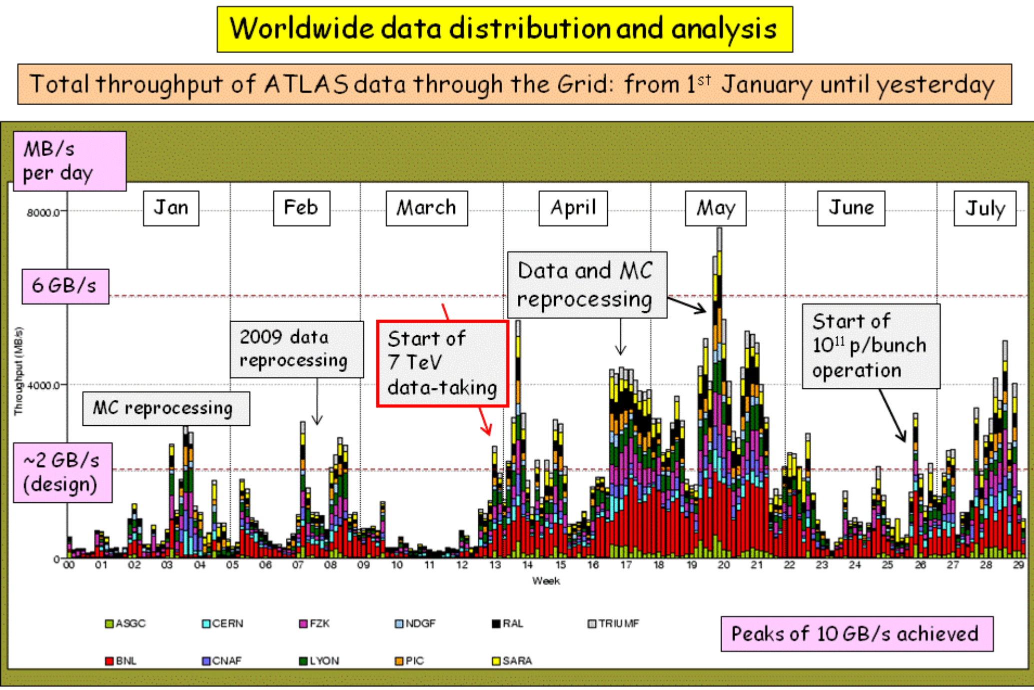

GRID-based analysis in June-July 2010:
> 1000 different users, ~ 11 million analysis jobs processed

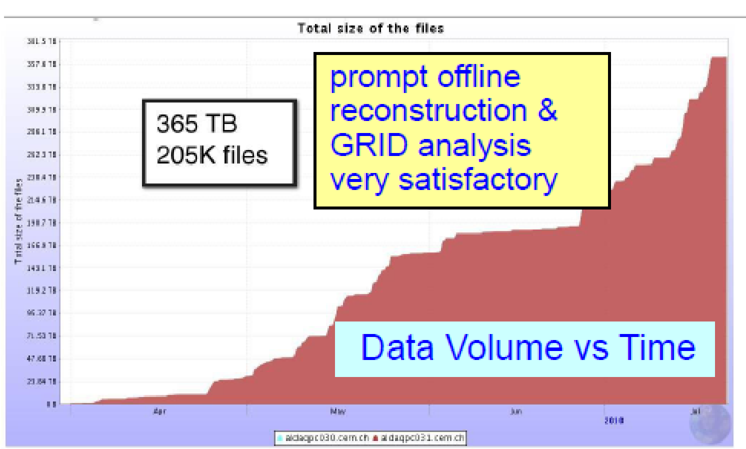

Data Processing, Transfer and Analysis Activities

Excellent experience so far: the whole offline and computing organization + GRID infrastructure performing very well.

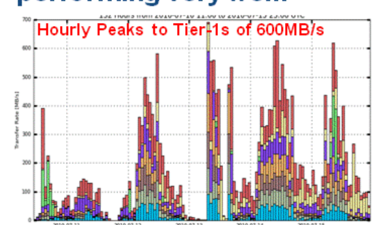

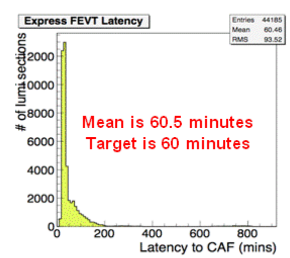

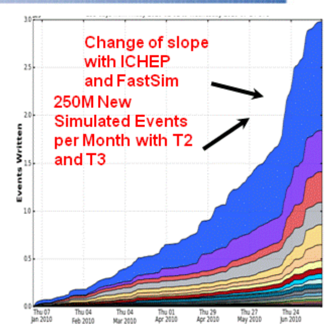

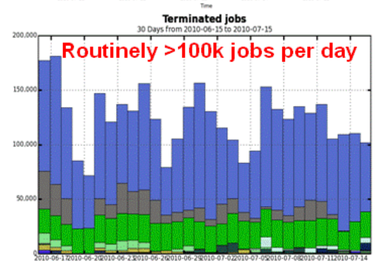

>500 individuals submitting jobs

Being less parochial…

NEWS from: HealthTech Wire date: 23 Aug 2010

eHealth in India: Advancing on all levels

HYDERABAD, INDIA - (HealthTech Wire / Perspective - by Armin Scheuer) - Hyderabads airport is a ...

HealthGrid©

A Human Grid Initiative

Who are we? Projects Dissemination Q.A.R.M. Training R&D Tools Partners Jobs

You are here: Who are We? :: Initiative

Search

The initiative

Concepts :: The Vision :: Missions :: Tools

Concepts

The HealthGrid initiative, supported by the HealthGrid association since 2003, was created to bring a long-term continuity, to reinforce and to promote awareness of the possibilities and advantages linked to the deployment of Grid technologies in health at the planet level. The HealthGrid community consists of over 1000 members from about 50 countries around the world.

The association works indeed for the initiative and the projects in which the association is involved and makes the initiative grow up.

Make use of the new information technologies in order to provide to the patients a less expensive, faster and more efficient individualised healthcare is HealthGrid's main goal.

Support the initiative

Breaking News

2010-08-04

IEEE eScience Workshop on High-Performance Computing in [...]

The purpose of this workshop is to provide the opportunity for participants to discuss and share the latest research in parallel and distributed high performance computing systems ... [More]

> All News

Event in the spotlight

2010-08-16

Clouds

The skeptic can be forgiven for wondering if there is more to Clouds than a ‘revenue concept’

Tony Cass, “Grid and Cloud Computing; an Overview”

Why “Cloud”?

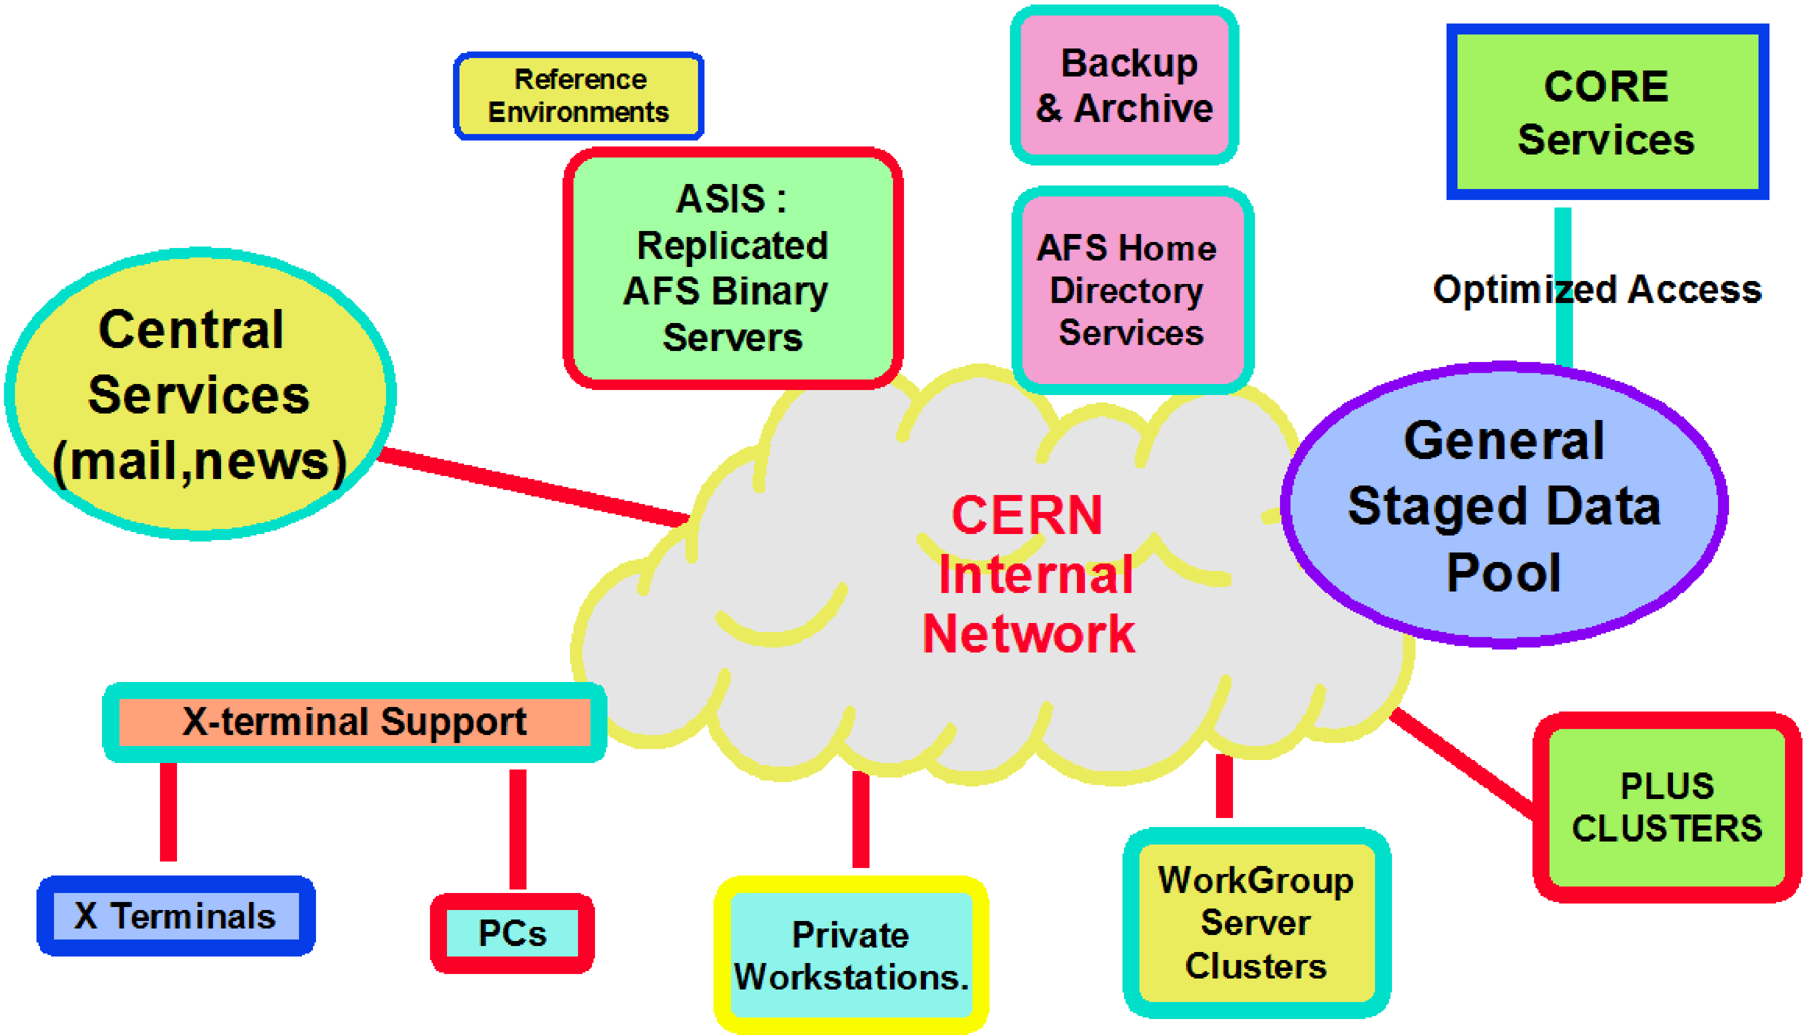

Reference Environments
Backup & Archive
CORE Services
ASIS : Replicated AFS Binary Servers
AFS Home Directory Services
Optimized Access
Central Services (mail,news)
CERN Internal Network
General Staged Data Pool
X-terminal Support
X Terminals
PCs
Private Workstations.
WorkGroup Server Clusters
PLUS CLUSTERS

XaaS

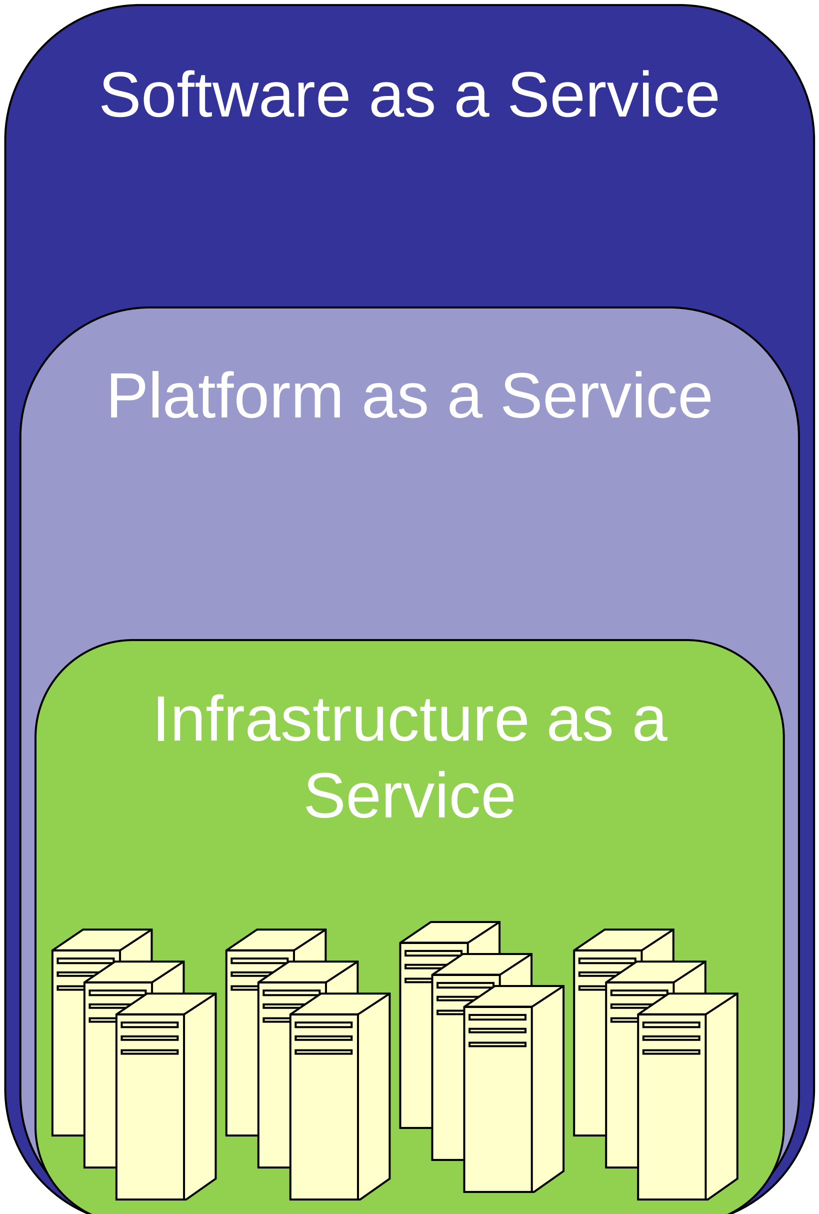

On demand access to applications

skype

Gmail by Google

Service-now.com
On Demand IT Service Management

Google Docs

Platform for building & delivering Web applications

force.com platform as a service

Azure

Utility computing

amazon web services™ EC2 S3

ElasticHosts
Flexible servers in the cloud

flexiscale™

GOGRID

greenqloud

What problem is being solved?

- **How do I monetise my idle resources?**
  - Services we would recognise today as "Cloud services" (notably remote data storage) existed in the early 2000s but were not sold as such.
  - The hype around "Cloud Computing" really started with the introduction of Amazon's EC2, a service that could be introduced because Amazon had free capacity since their computing facilities were sized to meet peek demand.
- This, however, is the cynic's (or sceptic's) view.

What problem is being solved?

- **User view**
  - I have a computing need but cannot afford to build the necessary infrastructure myself—but I can afford to "pay as I go".
- **System View**
  - Applications are isolated from the underlying hardware and operating system.
    - Unlike the Grid, there is no need for the CPU consumer to have any knowledge of the environment at the supplier.
  - Suppliers can be isolated from user complexity
    - Schedule only virtual machine images; users choose the image contents according to need.
      - compare CERN batch system where we have many subclusters for different VO needs, e.g. production vs analysis.

Who owns the resources?

- No question: the service provider!
  - Users pay according to usage...

Enabling technology?

- **(x86) Virtualisation**
  - Easily available, low-overhead virtualisation enables users to provide their own computing environment running on an unknown (and irrelevant) substrate
    - Virtualisation overhead is
      - low for CPU and network I/O, likely to be needed by users
      - high for I/O to local disk, but this is less likely to be an issue
- A view slanted towards IaaS...
  - ... but virtualisation doubtless also helps SaaS and PaaS providers, removing the link between the physical hardware and the logical service as seen by the users
    - c.f. comments in Ulrich Schwickerath's presentation.

Cloud Models

Model	Definition	Examples
Public	Infrastructure is owned by a single organisation and made available to other organisations	•**Commercial Cloud Providers** •**Community Public Clouds** by ICT service centres to enable scientific and educational projects to experiment with cloud computing •**Special Purpose Clouds** with dedicated capabilities (Science Clouds, HPC Clouds) •**Regional Clouds** to address regulatory or latency issues
Private	Infrastructure is owned by a single organisation and made available only to the organisation	•Optimise and support **internal operation** •**Saas/PaaS support** •IT consolidation within **large organisations** (Government Clouds, University Clouds)
Hybrid	Infrastructure is a composition of two or more clouds	•**Cloudbursting** to address peak demands •**Cloud Federation** to share infrastructure with partners •**Cloud Aggregation** to provide a larger resource infrastructure

With apologies to Ignacio M. Llorente

Gartner's View

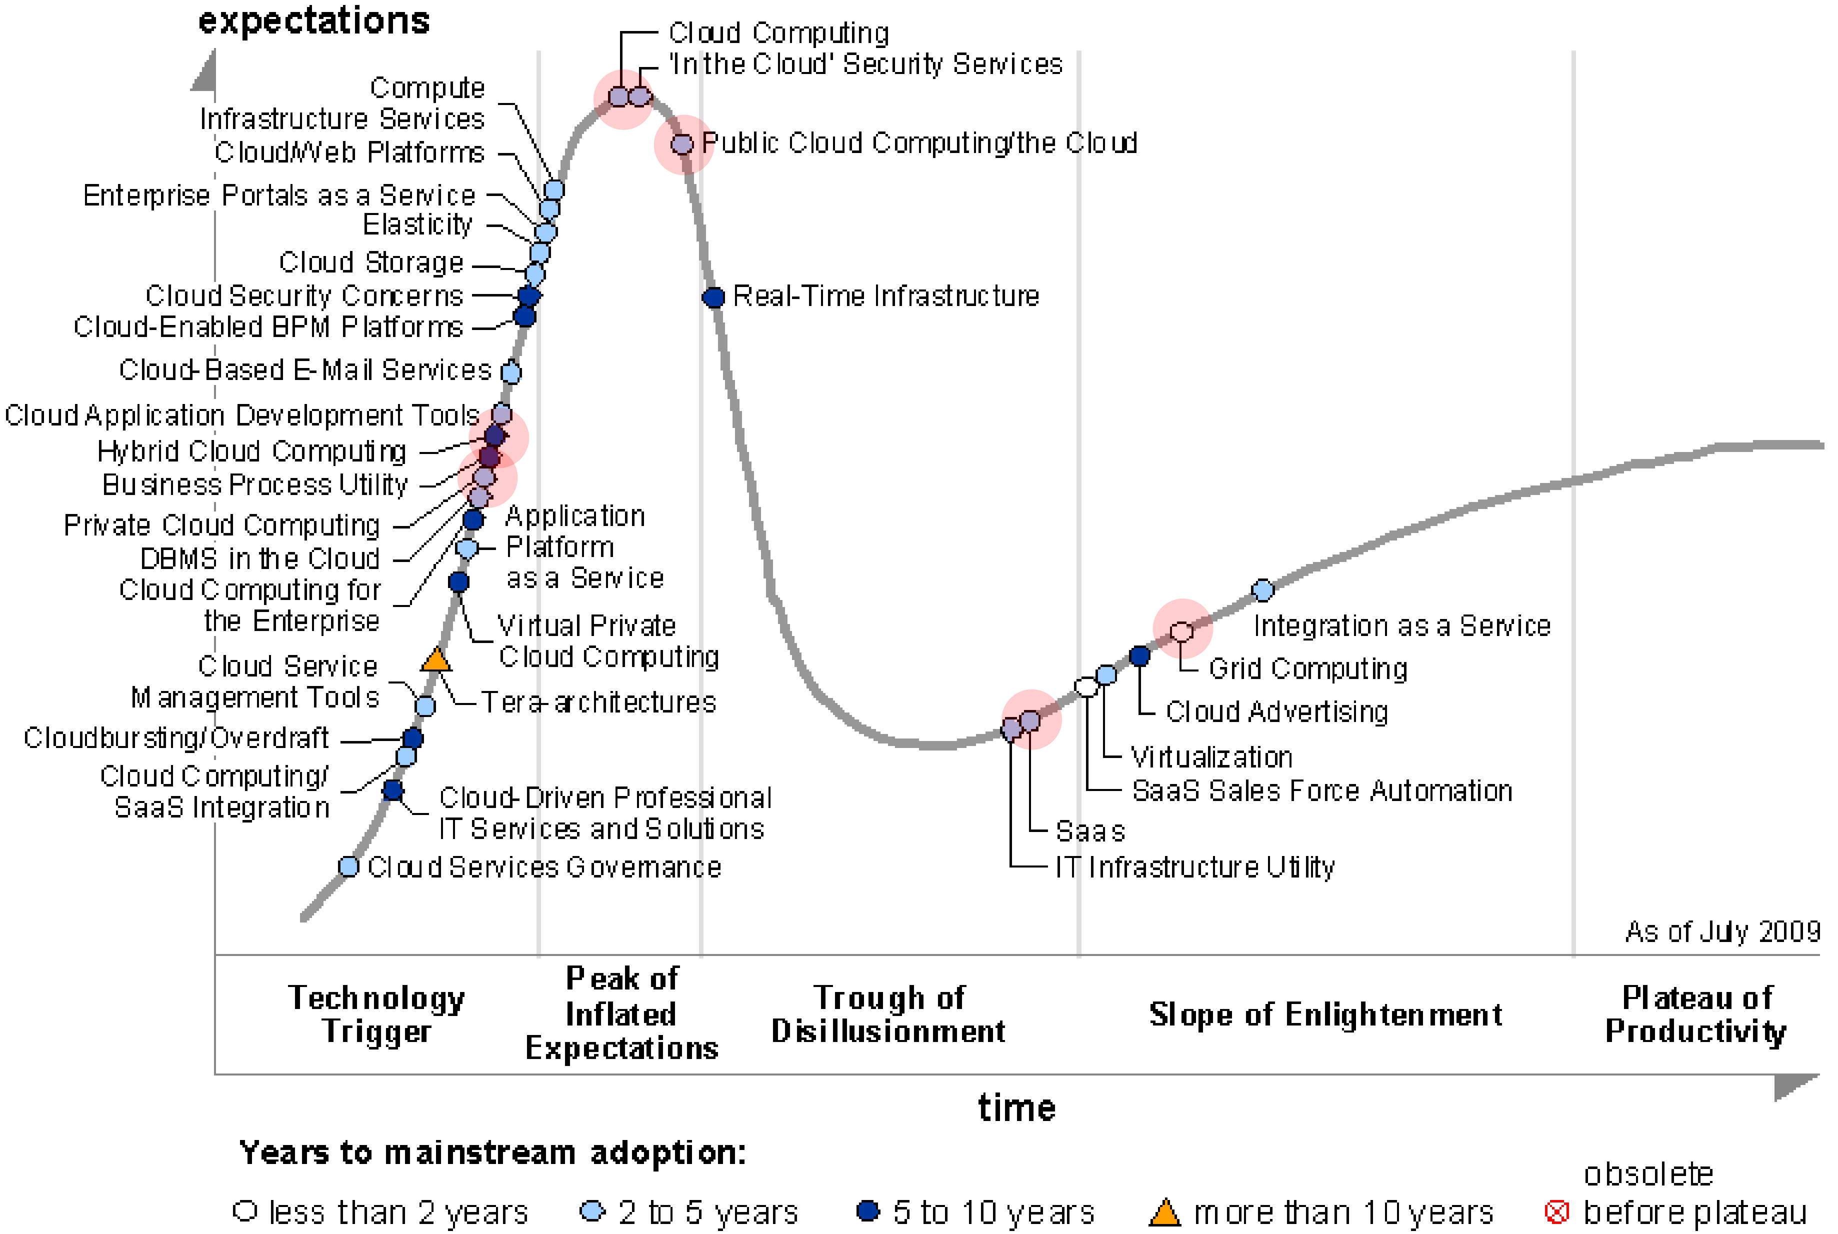

Analogies

Mountains & Theme Parks

- Mountains

- Theme Parks

Mountains and Theme Parks

Other Projects

Management

THE MOUNTAINS OF ENGLAND AND WALES

The Mountains of England and Wales are those summits in England and Wales which are at least 2000 feet (610 metres) high, with a minimum of 50 feet (15 metres) of ascent on all sides. All the mountains were visited personally by us and we checked the marginal tops using on-the-spot surveys. Full details of the tops, and how to ascend them in a series of interesting walks is given in: England & Wales

Other Users

Sys Admins

Information

News

Meetings

Wednesday 25th August

Session-5: The Scaling of the *Peaks*

13:01-18:00	A-Team - Fell Running (Chair: Pete Clarke/Andrew Sansum)
	B-Team - Hiking - Loughrigg Fell (Chair: Steve Lloyd)
	C-Team - Strolling (Chair: Roger Jones)
	D-Team - Sedentary (Chair: Chair)
	W-Team - Watery (Chair: Neasan O'Neill)
19:00-	Collaboration Dinner at The Low Wood Hotel. Sponsored by Boston

Mountains & Theme Parks

- Mountains
  - Owned by?
  - Organised for?

- Theme Parks
  - Owned by?
  - Organised for?

Mountains & Theme Parks

- Mountains
  - Owned by? Users
  - Organised for? Users
  - Open Interface? Yes

Grids

- Theme Parks
  - Owned by? Providers
  - Organised for? Providers
  - Open Interface? No

Clouds

A Digression

Volunteer Computing

- Where do SETI@Home & BOINC fit?
  - What problem is being solved?
    - Lack of resources.
  - Who owns the resources?
    - Providers. But there is no one provider and no SLA.
  - Enabling technology?
    - Cheap CPU and home networks.
- Not “cloud”, but more “cloud” than “grid”
- Actually, who are the users and who the providers?
  - Credit given by CPU consumers is a valuable commodity...

Virtualisation and VC

- Lack of control over the execution environment has limited the range of applications that can exploit volunteer computing.
  - LHC@HOME required a significant effort to fully understand IEEE floating point issues and use of a specialist library to ensure consistent results.
- Running a virtual image as the BOINC work unit provides a guaranteed environment and fully isolates users from the workload.
  - Some interesting work is being done at CERN to integrate CernVM and BOINC with the aim of providing HEP work for the volunteer community.

Problems

Cloud Issues

- Cost
  - Service providers must make a profit as well as amortising their costs. Large scale consumption is thus unlikely to be cost effective compared to a private service. But
    - this is not an issue when discussing marginal capacity
    - CapEx vs OpEx considerations may favour IaaS solutions.
- Opacity
  - Not an issue for CPU-bound work, but data-heavy applications can't know either where data is stored wrt the compute nodes or the connecting bandwidth. An issue as
    - charges are usually levied for both data storage and data transfer out of the storage unit, and
    - Low-bandwidth data connections reduce CPU usage efficiency, increasing cost.
  - The problem is acknowledged by Amazon at least but, to date, no service offerings specifically address data-processing requirements.

Grid Issues - I

- Early vs. Late Job Binding
  - Grid sites generally want to maintain a high average CPU utilisation.
    - Easiest to do this if there is a local queue of work to select from when another job ends.
  - Users are generally interested in turnround times as well as job throughput.
    - Turnround is reduced if jobs are held centrally until a processing slot is known to be free at a target site.
  - Grid workload management systems (which look at expected wait times at sites) can't guarantee rapid job execution.
    - Clouds can…
      - ... even if this may be at the expense of overall throughput...

Grid Issues - II

- Software (in)consistency
  - HEP users have large libraries of code which need updating frequently.
    - Guaranteeing consistency across thousands of processors at hundreds of sites is difficult, if not impossible...
    - One inconsistent site or node can cause significant problems as users frequently need many jobs to be completed successfully before they can proceed.
  - Not such a concern for many other communities as they frequently use standard applications (c.f. the computational chemistry grid) and vary the input data.
- Software/OS/hardware compatibility
  - A related issue: Grid users need to be aware of the OS (and perhaps hardware) at different sites.
    - Or, at least, specify their requirements in the JDL and possibly restrict the set of potential execution targets.

Pilot Job Frameworks

- HEP users soon developed Pilot Job Frameworks to address both Grid issues.
  - These frameworks submit "pilot jobs" to the grid. When a pilot job is scheduled at a site it
    - verifies local software installation, installing the correct software if necessary, then
    - connects to a central, VO-specific workload queue to download a "joblet" which will carry out the useful work
      - in many cases, pilot jobs will continue to download joblets until the alloted cpu time has been reached.

Synthesis?

Synthesis?

- Two aspects of the Cloud/IaaS model are highly attractive for service providers:
  - Applications are isolated from the underlying hardware.
    - Unlike the Grid, there is no need for the CPU consumer to have any knowledge of the real hardware at the supplier.
  - Suppliers can be isolated from user complexity
    - Schedule only virtual machine images; users choose the image contents according to need.
      - compare CERN batch system where we have many subclusters for different VO needs, e.g. production vs analysis.
- Can we integrate virtual machine images in a Grid world to provide these advantages?

Synthesis?

- **Yes!**
  - VOs provide virtual machine images and these are scheduled by sites either
    - on demand, or
    - dynamically according to observed utilisation.
  - Preferably, these VM images contain the VO pilot job client
    - and the CernVM filesystem provides dynamic access to the VO software libraries.
  - Grid infrastructure provides information about resources available (CPU, data, ...)
- **Issues**
  - Trust (again!)
  - VO concerns over virtualisation overheads
    - Need to balance single job performance against overall system utilisation.

Cloudburst?

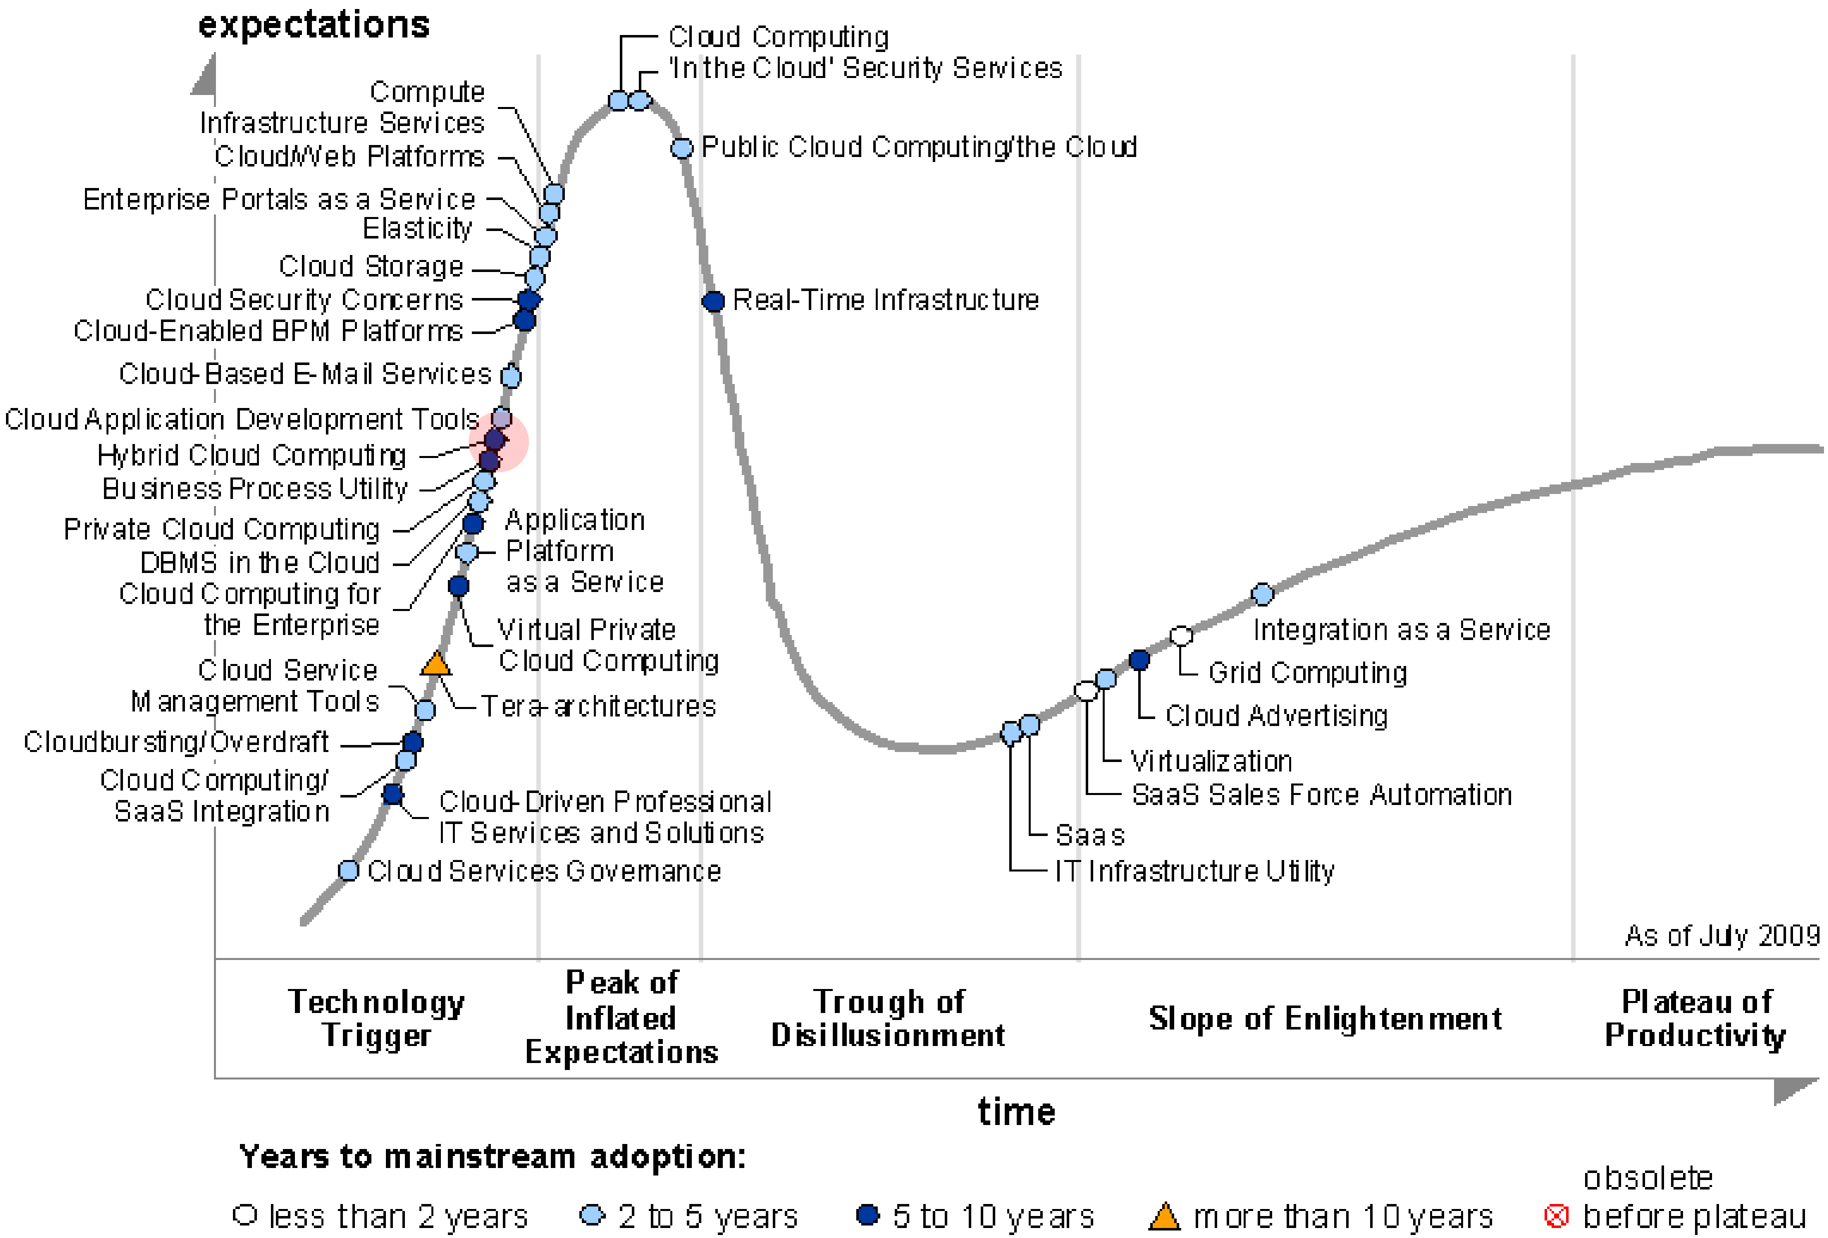
expectations
Cloud Computing
'In the Cloud' Security Services
Compute
Infrastructure Services
Cloud/Web Platforms
Public Cloud Computing/the Cloud
Enterprise Portals as a Service
Elasticity
Cloud Storage
Cloud Security Concerns
Real-Time Infrastructure
Cloud-Enabled BPM Platforms
Cloud-Based E-Mail Services
Cloud Application Development Tools
Hybrid Cloud Computing
Business Process Utility
Application Platform as a Service
Private Cloud Computing
DBMS in the Cloud
Cloud Computing for the Enterprise
Virtual Private Cloud Computing
Integration as a Service
Grid Computing
Cloud Service Management Tools
Tera-architectures
Cloud Advertising
Cloudbursting/Overdraft
Virtualization
Cloud Computing/SaaS Integration
Cloud-Driven Professional IT Services and Solutions
SaaS Sales Force Automation
SaaS
Cloud Services Governance
IT Infrastructure Utility
As of July 2009
Technology Trigger
Peak of Inflated Expectations
Trough of Disillusionment
Slope of Enlightenment
Plateau of Productivity
time
Years to mainstream adoption:
less than 2 years
2 to 5 years
5 to 10 years
more than 10 years
obsolete before plateau

Summary

Summary

- Grids and Clouds address different areas of the distributed computing problem.
  - Grids provide an easy interface to resources which users have a right to use.
  - Clouds provide “pay-as-you-go” access to services set up by a provider
    - IaaS can be exploited by users with CPU or data storage needs who understand the cost implications
      - likely to be most beneficial for addressing peak processing needs with limited data transfer.
    - SaaS and PaaS are not as directly relevant
      - but can be interesting, e.g. database services (and other applications for sites as opposed to users—e.g. gmail)
- Neither can replace (or is better than) the other
  - but there is certainly scope to exploit virtualisation and Cloud interfaces in a Grid environment to reduce application/hardware and user/administrator coupling.