Towards cloud computing, opportunities and challenges for e-Science

Dr. Fabrizio Gagliardi

Director for Europe, Middle East, Africa and Latin America

Technical Computing

Microsoft Research

Outline

- The experience of the Grid
- Examples beyond e-Science
- Views: To distribute or not to distribute
- Issues and new trends:
  - Virtualisation: Pros and Cons
  - Green Grid and Cloud Computing
- Conclusions

The experience of the Grid 1/3

- Grids for e-Science:
  a success story so far?
  - Several mature Grid Middleware stacks
  - Many HPC applications using the Grid
    - Some of them (HEP, Bio) in production use
    - Some of them still in testing phase: more effort still required to make the Grid their day-to-day workhorse
  - e-Health applications also part of the Grid
  - Some industrial applications:
    - CGG Earth Sciences

The experience of the Grid 2/3

- Grids beyond e-Science
  - Slower adoption: prefer different environments, tools and have different TCOs
    - Intra grids, internal dedicated clusters , cloud computing
  - e-Business applications
    - Finance, Enterprise Resource Planning, SMEs
  - Industrial applications
    - Automotive, Aerospace, Pharmaceutical industry, Telecom
  - e-Government applications
    - Earth Observation, Civil protection:
    - e.g. The Cyclops project

The experience of the Grid 3/3

- Industry also demonstrated interest in becoming an HPC infrastructure provider:
  - On-demand infrastructures:
    - Cloud and Elastic computing, pay as you go…
    - Data centers: Data getting more and more attention
  - Service hosting: outsourced integrated services
- EU "Pre-commercial procurement" for fostering technology transfer: Research⇔Industry
  - Example: PRACE building a PetaFlop Supercomputing Centre in Europe

Examples beyond e-Science

EU BEinGRID: Computational Fluid Dynamics

BEinGRID

Demonstrator 1: CFD

European Commission
Information Society Technologies

Computational Fluid Dynamics & Computer Aid Design

The *actors*:

- ICON ………………...as the end-user
- ICCS / NTUA ……….as the pilot coordinator and grid expert
- OpenCFD …………...as the solution provider

The *plot*: emphasize the business benefits of performing Computational Fluid Dynamics (CFD) simulation of fluid flow for aircrafts or automotive vehicles within a grid environment (Airbus / Audi)

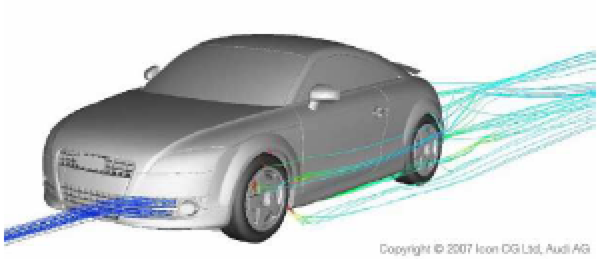

Business Experiments in GRID

Examples beyond e-Science

CYCLOPS: Forest Fire propagation

Selected use case : Large forest fire propagation

- Propagation models used on the field are simplistic due to computational limitations
- More sophisticated models require heavy computations, large number of field and meteo parameters
- French research project PAREFEU, link with GMES service PREVIEW

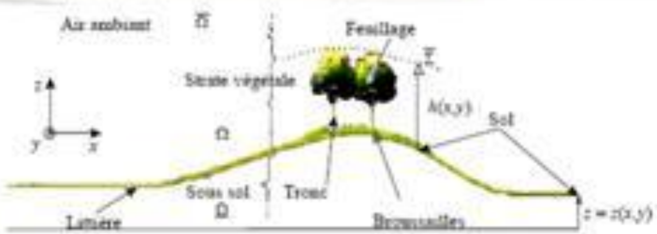

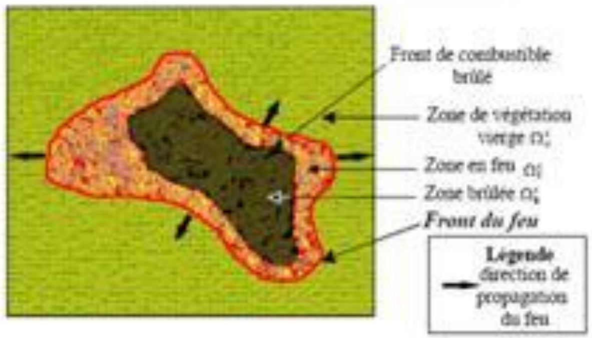

After Sero-Guillaume et al.

nativi@imaa.cnr.it

Information Society and Media

11

Examples beyond e-Science

EGEODE VO : Seismic processing based on Geocluster application by CGG company (France)

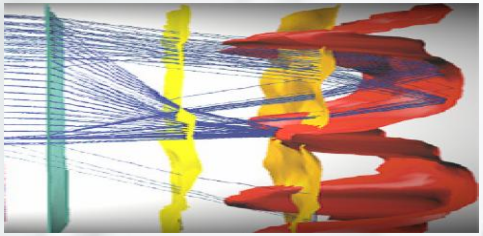

degree

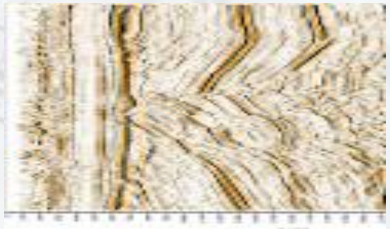

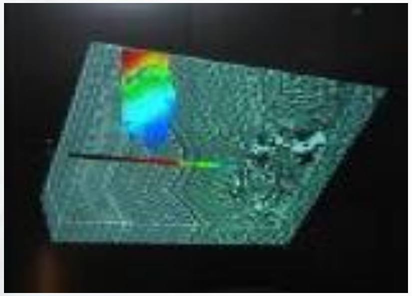

egee
Enabling Grids for E-sciencE

CGGVERITAS

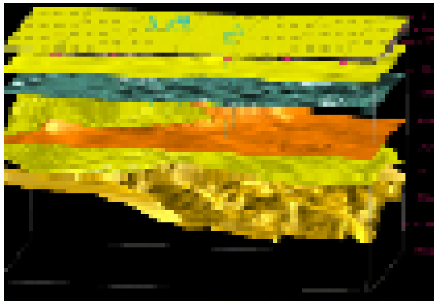

In summary…

- Grid computing has delivered an affordable and high performance computing infrastructure to scientists all over the world to solve intense computing and storage problems within constrained research budget

- This has also been effectively used by industry to increase the usage of their computing infrastructure and reduce Total Cost of Ownership (TCO)

- Grid is not only aggregating computing resources but also leveraging international research networks to deliver an effective and irreplaceable channel for international collaboration

The flip side…

- Major issues with wide adoption of Grid computing in eScience, e-Business, industry etc. have to do with:
  - Cost of operations and management complexity
  - Not a solution for all problems (latency, fine grain parallelism are difficult)
  - Difficult to use for the average scientist
  - Security and reliability
- Power consumption and heat dissipation are becoming a limiting factor to consumer based distributed systems
- We are observing the limits of Moore's law

Switching Gears:

“To Distribute or Not To Distribute”

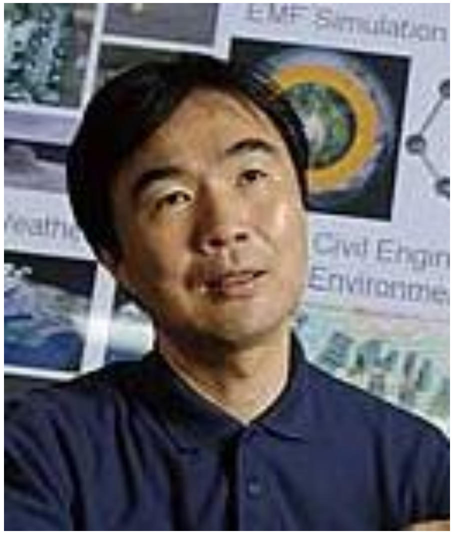

- Prof. Satoshi Matsuoka, TITech
- Keynote at Mardi Gras Conference, Baton Rouge, Jan.31, 2008
  - **In the late 90s, petaflops were considered very hard and at least 20 years off …**
  - **… while grids were supposed to happen right away**
  - **After 10 years (around now) petaflops are “real close” but there's still no "global grid”**
  - ***What happened?***

What Happened?

- It was easier to put together massive clusters than to get people to agree about how to share their resources
- For tightly coupled HPC applications, tightly coupled machines are still necessary
- Grids are inherently suited for loosely coupled apps (e.g., Monte Carlo, Parameter Sweep), or enabling access to machines and data, and the integration of the two
  - **With Gilder's Law, bandwidth to the compute resources will promote thin client approach**
  - **Example: *Tsubame* machine in Tokyo**

Virtualization and Cloud Computing

- *“There’s Grid in them than clouds!”*
  - I. Foster’s blog, ANL & UC, Jan. 8, 2008
- Clouds have a very simple user API effectively hiding all the complexity of an *ad hoc grid* on the back-end
  - e.g., Amazon’s EC2 & S3, IBM’s Blue Cloud and others …

- ***If so, will this enable mass-market grids?***
  - **Users don’t have to be aware of using “a grid”**
- ***If so, what does “cloud interoperability” require?***
  - **Is *virtualization* a means of achieving this?**
- ***Major opportunity for synergy***

What is Virtualization?

Virtualization is the isolation of one computing resource from the others:

or

Virtual Applications
Any application on any computer on-demand

Virtual Presentation
Presentation layer separate from process

Virtual Machine
OS can be assigned to any desktop or server

Virtual Storage
Storage and backup over the network

Virtual Network
Localizing dispersed resources

Slide Courtesy of Neil Sanderson, UK Product Manager, Virtualisation and Management, Microsoft Ltd.

Virtualisation: Pros

Control costs

- Optimize utilisation of servers
- Reduce costs: hardware, power, space
- Reduce overall system complexity
- Reduce application compatibility conflicts

Improve Availability

- Disaster recovery capabilities
- Streamline server maintenance, and isolate risk
- Enable access to any desktop application anywhere

Drive Agility

- Dynamic resource allocation
- Faster provisioning of services / workloads to support business growth and meet SLAs

Slide Courtesy of Neil Sanderson

Virtualisation: Cons?

- New technology
- Management overhead
- Single Point of Failure
- Software Licensing?
- Performance?
- Policy issues
- Security Concerns

Virtualisation market: The game is only starting (1/2)
Only 5% of servers are virtualised

- Computerworld
  - “Although virtualization has been the buzz among technology providers, only 6% of enterprises have actually deployed virtualization on their networks, said Levine, citing a TWP Research report. **That makes the other 94% a wide-open market**.”
  - “We calculate that roughly 6% of new servers sold last year were virtualized and project that 7% of those sold this year will be virtualized and **believe that less than 4% of the X86 server installed base has been virtualized to date.**
- Pat Gelsinger, Intel VP Sept. 2007
  - **“Only 5% of servers are virtualized.”**

World Wide Virtualization Adoption

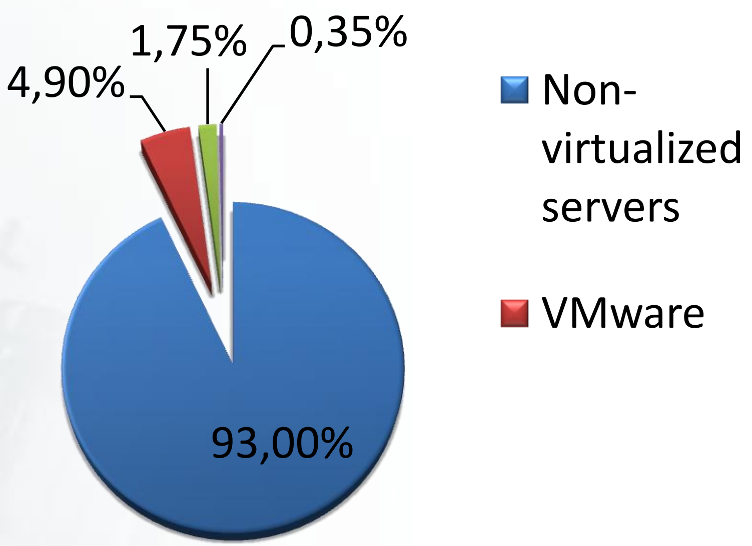

Slide Courtesy of Neil Sanderson

Virtualisation market: The game is only starting (2/2)
Multiple solutions – New entrants

SWsoft tries to virtualize China before VMware speaks

Market Scan
The Competition's Catching

FTN Midwest Securities analyst Trip Chowdhry said that the problem is that investors wrongly believed that VMware itself was the trend. Now, he says, they are realizing that virtualization is the trend.

lets Leopard virtualize
Tired of thinking different
1 Nov 18:33

XenSource preaches the joy of platform virtualization
Interop New management, new credo, same old enemy
By Paula Rooney • Wednesday 24 Oct 2007 18:21

Slide Courtesy of Neil Sanderson

Virtualisation at a glance

- Virtualisation: A promising technology
- It's not only about virtual machines…
- Virtualisation not much exploited in e-Science so far
- However exploited in Cloud computing and hosted services
  - Which emerge as the next incarnation of distributed computing
- Virtualisation is a potentially better "Green" technology
  - Less space, less energy consumption
- Microsoft: a player in the field
  - HyperVisor

Emerging new trends: Green Grid and Pay per … CPU/GB

- The Green Grid, IBM Big Green and other IT industry initiatives try to address current HPC limits in energy and environmental impact requirements
- Computer and data centers in energy and environmental favorable locations are becoming important
- Elastic computing, Computing on the Cloud, Data Centers and Service Hosting are becoming the new emerging solutions for HPC applications
- Many-multi-core and CPU accelerators are promising potential breakthroughs

Cloud computing and storage on demand (1/2)

Cloud Computing: http://en.wikipedia.org/wiki/Cloud_computing

Amazon, IBM, Google, Microsoft, Sun, Yahoo, major 'Cloud Platform' potential providers

- **Operating compute and storage facilities around the world**
- **Have developed middleware technologies for resource sharing**
- **First services already operational - Examples:**
  Amazon Elastic Computing Cloud (EC2) -Simple Storage Service (S3)

Cloud computing and storage on demand (2/2)

Sun's Grid Utility Expands Beyond the United States

Published: May 3, 2007

by Timothy Prickett Morgan

Nearly two months ago, Sun Microsystems and a raft of software partners put a jukebox of sorts out onto its network.com Sun Grid compute ar... since end users don't have to get utility licens... problem that Sun had, however, is that only c... capacity.

Starting today, however, the Sun Grid utility, ... by default, is now available in 24 additional co... is now allowing customers to log in and buy c... Australia, Austria, Belgium, Canada, China, C... Greece, Hungary, India, Ireland, Italy, Japan, N... Sweden, and the United Kingdom. To help sp... is giving away 200 CPU hours of compute time... individual--who starts up an account. This is ...

Google, IBM Partner on Utility Computing Cloud

Published: October 8, 2007

by Timothy Prickett Morgan

Search engine and Web advertising goliath Google and server maker and supercomputing giant ...orking with some prominent American universities to set up ...sist those universities in teaching and developing ...esent the future of corporate computing some day.

...ernet-based applications, Google perhaps knows more about ...re utilized on the fly as workloads change) or cloud ...eems to imply utility-style infrastructure as well as ... with the size of the computing resources dedicated to ...ld encompass many different kinds of utilities, with varying ...tics, and a workload would make use of the computing ...her way to refer to a network of machines connected by ...load progressed. For instance, a Linux-X64 cluster might be ..., but a vector parallel computer might be used to run ...data and then a separate virtualization system might be used ...imilarly, large-scale parallel systems can be used to search ...es well), perhaps looking for spam lurking in wiki and forum ...d in business transactions.

COMPUTERWORLD Hardware

SEARCH

...enter ...sts | Five Technologies Simplifying Infrastructure Management | Powering Change in the Data Center | Time Warner Cable of New York, Case Study

Supercomputing at less than the cost of a gallon gas

New service allows users to offload computational work to supercomputer

By Patrick Thibodeau   Comments 3   Recommended 45   Share

April 1, 2008 (Computerworld) A supercomputing, pay-per-use service that would make supercomputing affordable to small and midsize businesses, as well individuals who need added computation power, was announced today.

Interactive Supercomputing Inc., a Waltham, Mass.-based company said its Star-P On-Demand lets users tap into 168 processor cores on Intel ...

This is not a software-as-a-service offering that gives users access to the application... the compute ability... users would take...

http://www.itjungle.com/bns/bns100807-story02.html

http://www.itjungle.com/tug/tug050307-story05

http://www.computerworld.com/action/article.do?command=viewArticleBasic&taxonomyName=mainframes_and_supercomputers&articleId=9073758&taxonomyId=67&intsrc=kc_top

Amazon EC2 and S3

- EC2 Beta Service: Web-Services based
  http://www.amazon.com/gp/browse.html?node=201590011
  - $0.10 per hour - Small Instance (Default)
    - 1.7 GB of memory, 1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit), 160 GB of instance storage, 32-bit platform
    - EC2 Compute Unit (ECU) - One EC2 Compute Unit (ECU) provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor
- S3 storage services: WS-based (REST and SOAP)
  http://www.amazon.com/S3-AWS-home-page-Money/b/ref=sc_fe_l_2?ie=UTF8&node=16427261&no=3440661&me=A36L942TSJ2AJA
  - *Storage: $0.15 per GB-Month of storage used*
  - *Data Transfer: $0.10 per GB - all data transfer IN*
    $0.18 per GB - first 10 TB / month data transfer OUT
    $0.16 per GB - next 40 TB / month data transfer OUT
    $0.13 per GB - data transfer out / month over 50 TB

Services may be given below actual cost for various reasons

Virtualised

EGEE cost estimation (1/2)

Capital Expenditures (CAPEX):

a. Hardware costs: 55.000 CPUs – 25 PB storage ~ in the order of 100M € (60-140M)
Depreciating the infrastructure in 5 years: 25 M € year (10-15M to 40-45M)
b. Cooling and power installations (supposing existing housing facilities available) 25% of H/W costs: 25 M €, depreciated over 5 years: 5 M € (2-8 M €)

Total: ~ 30 M € / year (15 M € - 45 M €)

Slide Courtesy of Fotis Karayannis

EGEE cost estimation (2/2)

Operational Expenditures (OPEX):

a. 20 M € per year for all EGEE costs (including site administration, operations, middleware etc.

b. Electricity ~10% of h/w costs: 10 M € per year (other calculations lead to similar results)

c. Internet connectivity: Supposing no connectivity costs (existing over-provisioned NREN connectivity)

Total 30 M € / year

CAPEX+OPEX= 60 M € per year (45-75 M)

Slide Courtesy of Fotis Karayannis

EGEE if performed with Amazon EC2 and S3

Similar order of magnitude (~47M €), probably more cost effective

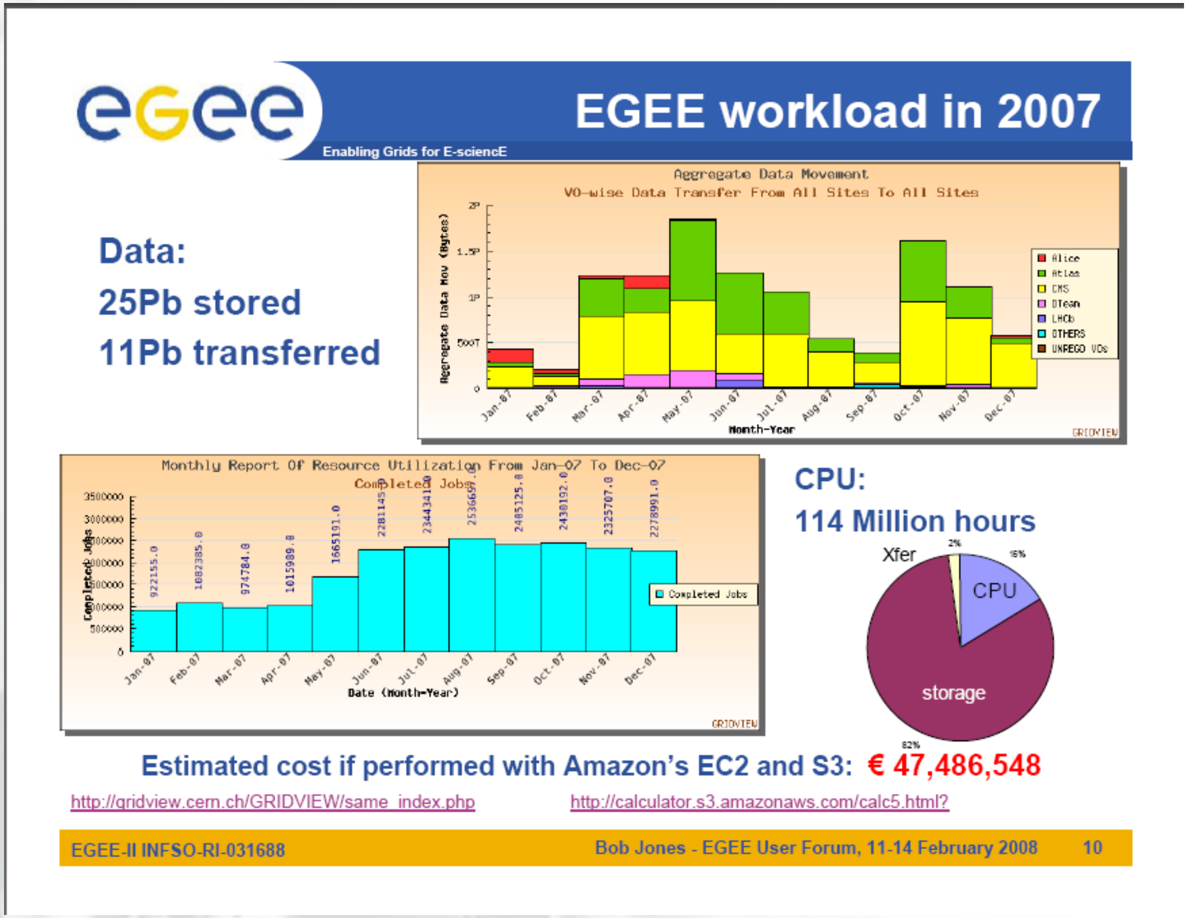

Slide Courtesy of Bob Jones

The Goal… More Time For Science

Highly Skilled Scientist Spending to Much Time Doing Non-scientific Work- Past and Present Approach are Manually Intensive

Integrated Information Management – Contextual, Collaborative and Rich Content

Today

Non-Scientific Activities

Tomorrow

More Time on Real Science

Not Enough Science

Conclusion (1/2)

- We are at a flex point in the evolution of distributed computing (nothing new under the sun…)
- Grid remains a good solution for a reduced number of communities (and often for social/political reasons)
- Cloud computing and hosted services are emerging as the next incarnation of distributed computing with some obvious additional advantages (think of data centres located in Iceland or Siberia for cheap cooling and electricity…)

Conclusion (2/2)

- The emphasis should move in making computing easier for the "normal scientist"
- We should critically re-think and avoid over engineered solutions (learn from the past experience)
- If we will be successful we will be able to enable major new scientific discoveries while industry and commerce will follow as it has always happenedin computing…

Thanks

Thanks to the organizers for the kind invitation and to all of you for your attention

fabrig@microsoft.com

•Resources

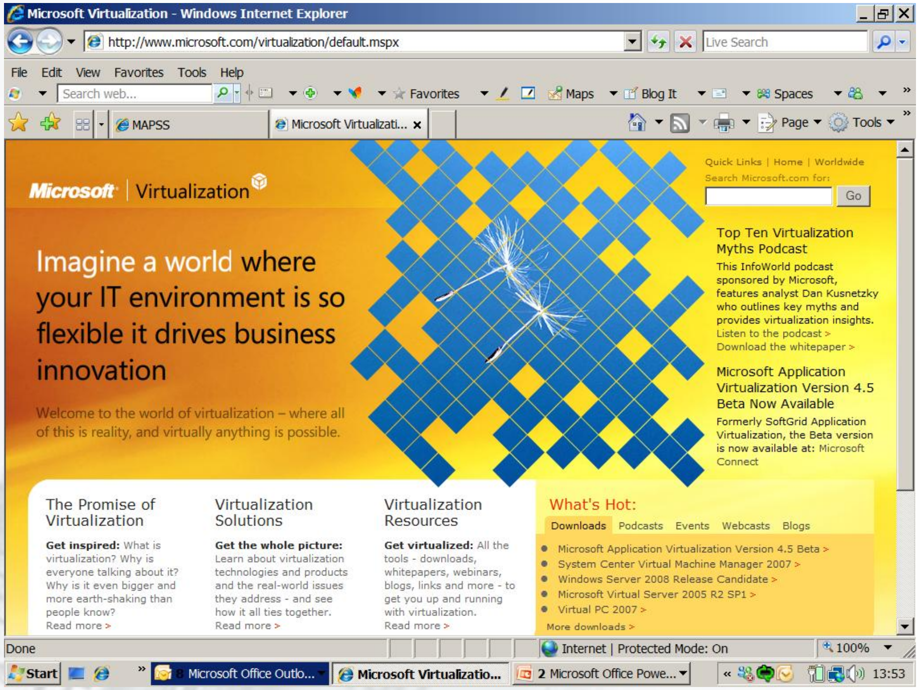

http://www.microsoft.com/virtualisation