e-Science and Cyberinfrastructure: A Middleware Perspective

Tony Hey
Corporate VP for Technical Computing
Microsoft Corporation

Licklider’s Vision

“Lick had this concept – all of the stuff linked together throughout the world, that you can use a remote computer, get data from a remote computer, or use lots of computers in your job”

Larry Roberts – Principal Architect of the ARPANET

The e-Science Vision

- e-Science is about multidisciplinary science and the technologies to support such distributed, collaborative scientific research
  - Many areas of science are now being overwhelmed by a ‘data deluge’ from new high-throughput devices, sensor networks, satellite surveys …
  - Areas such as bioinformatics, genomics, drug design, engineering and healthcare require collaboration between different domain experts
- ‘e-Science’ is a shorthand for a set of technologies to support collaborative networked science

Vision For Scientific Workflow

Example: Project NEPTUNE

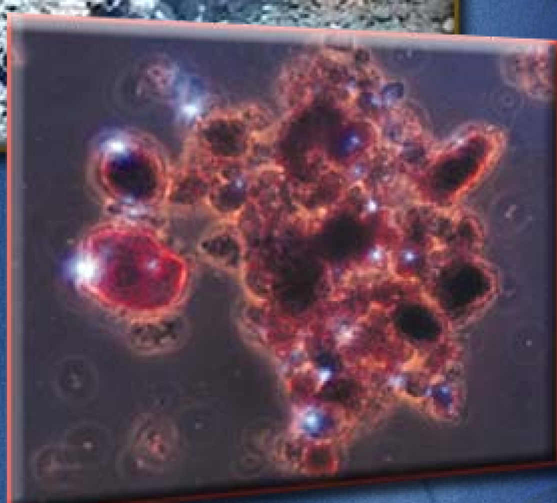

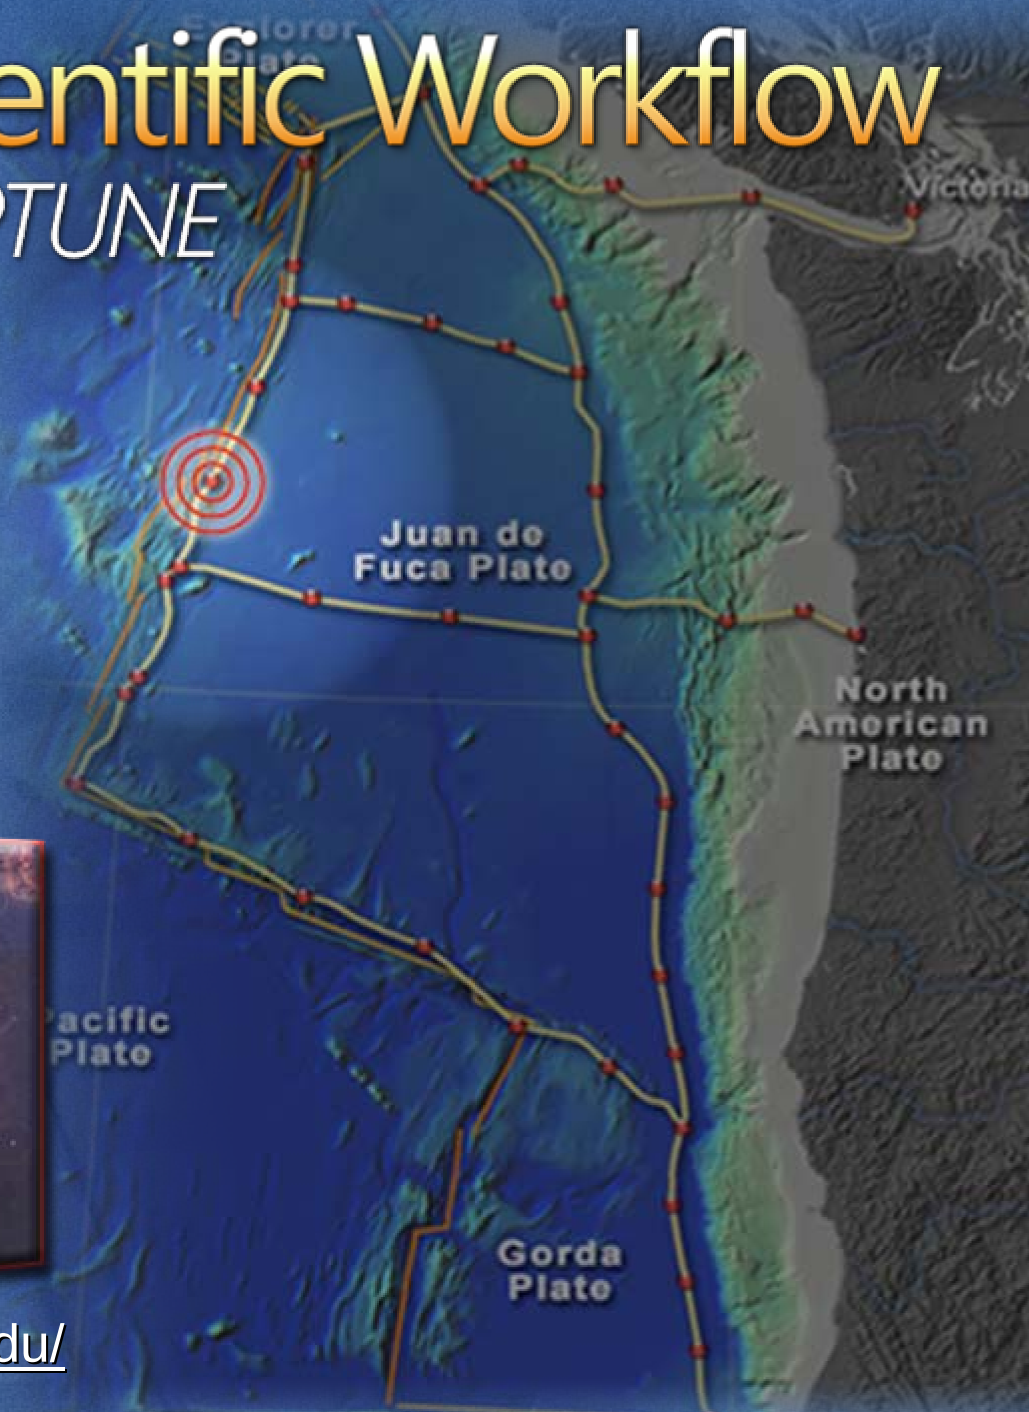

http://www.neptune.washington.edu/

Programmable Sensors & Remote Instruments

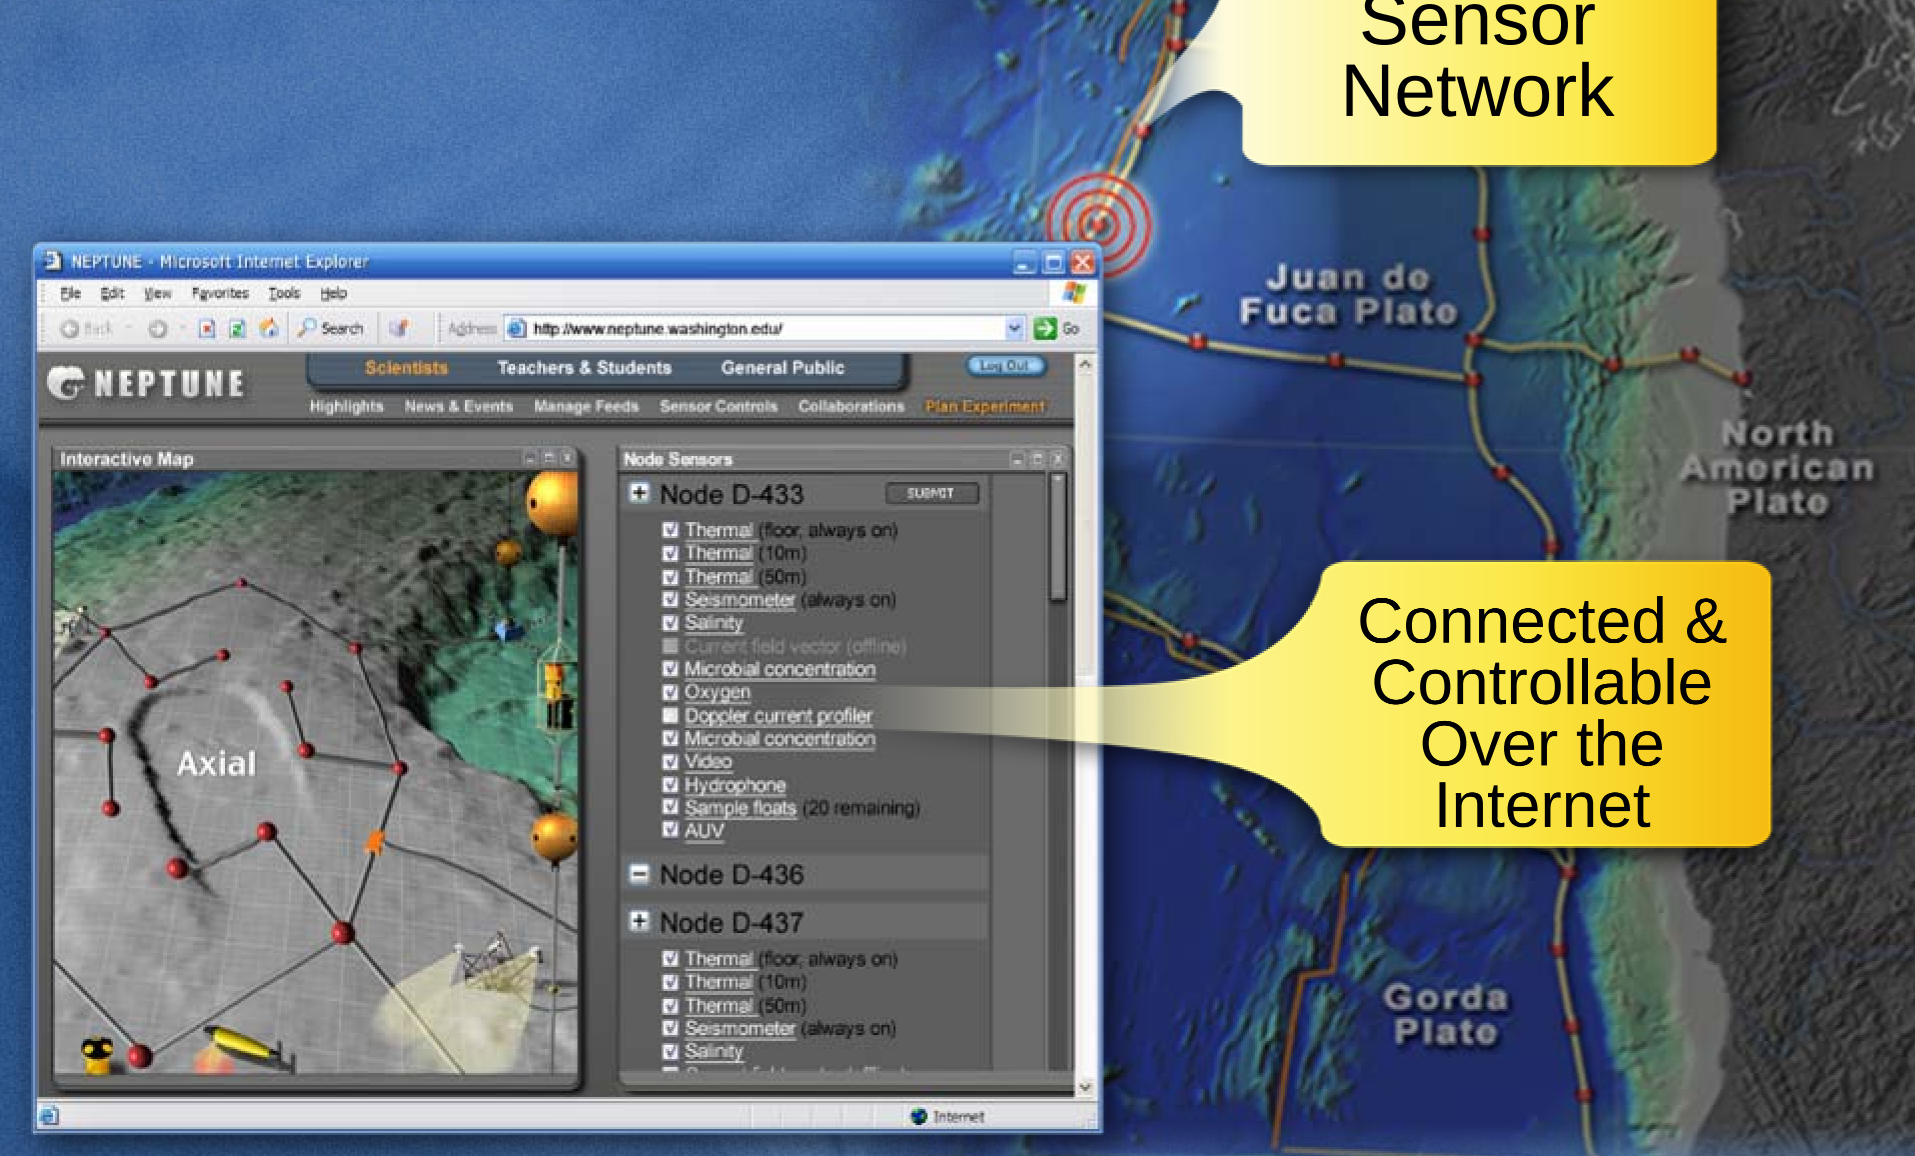

Data Workbench

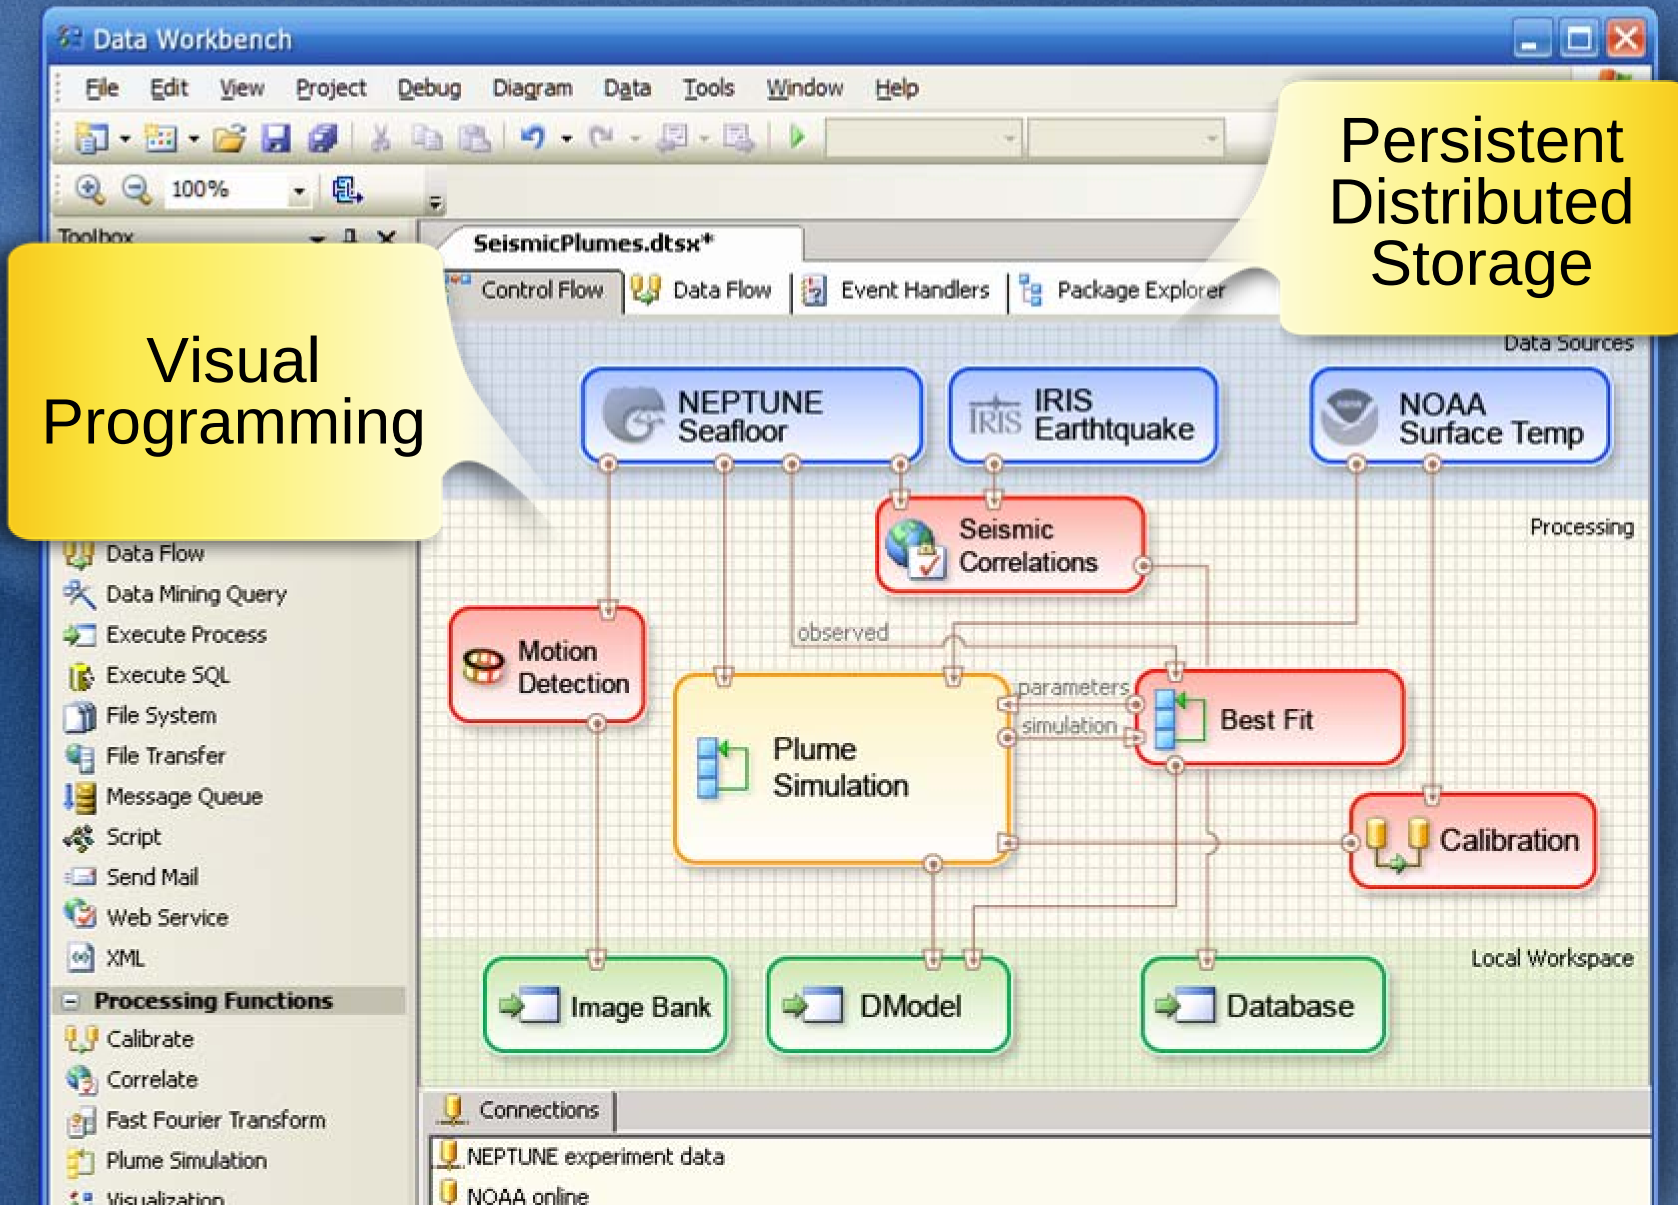

Data Workbench

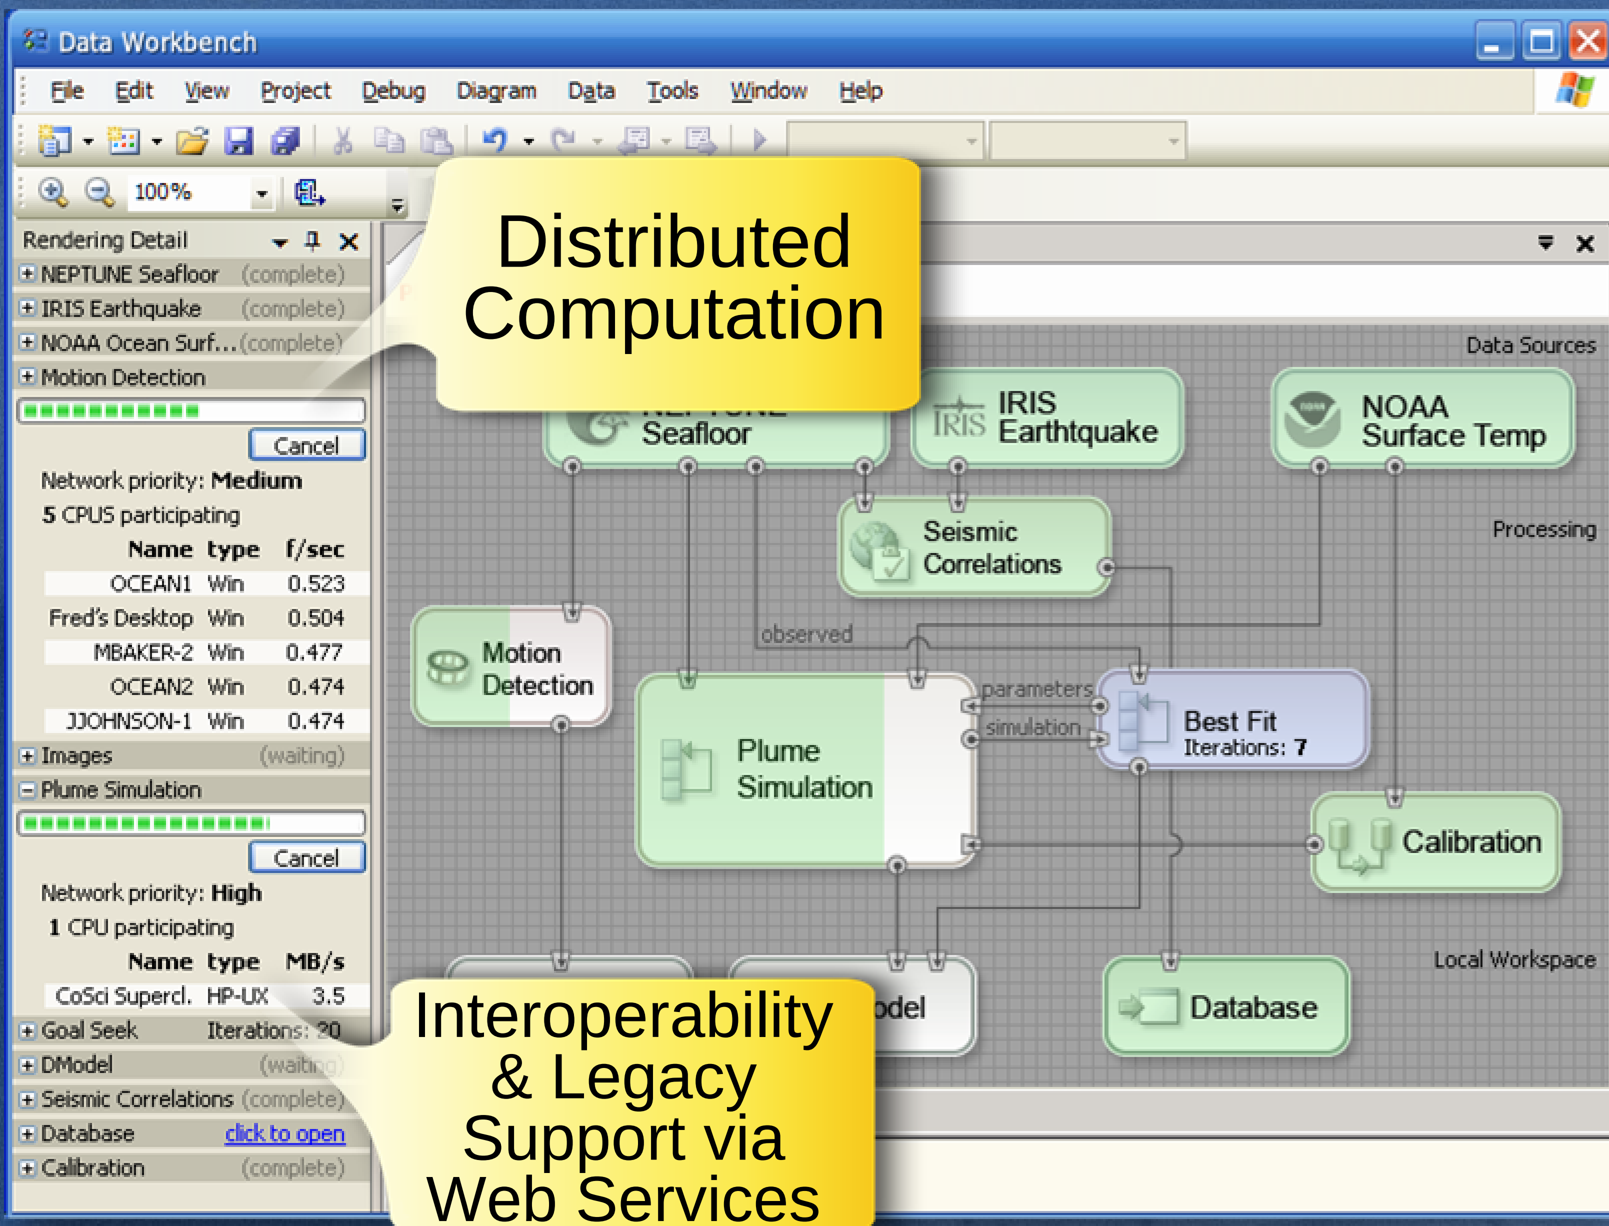

Research

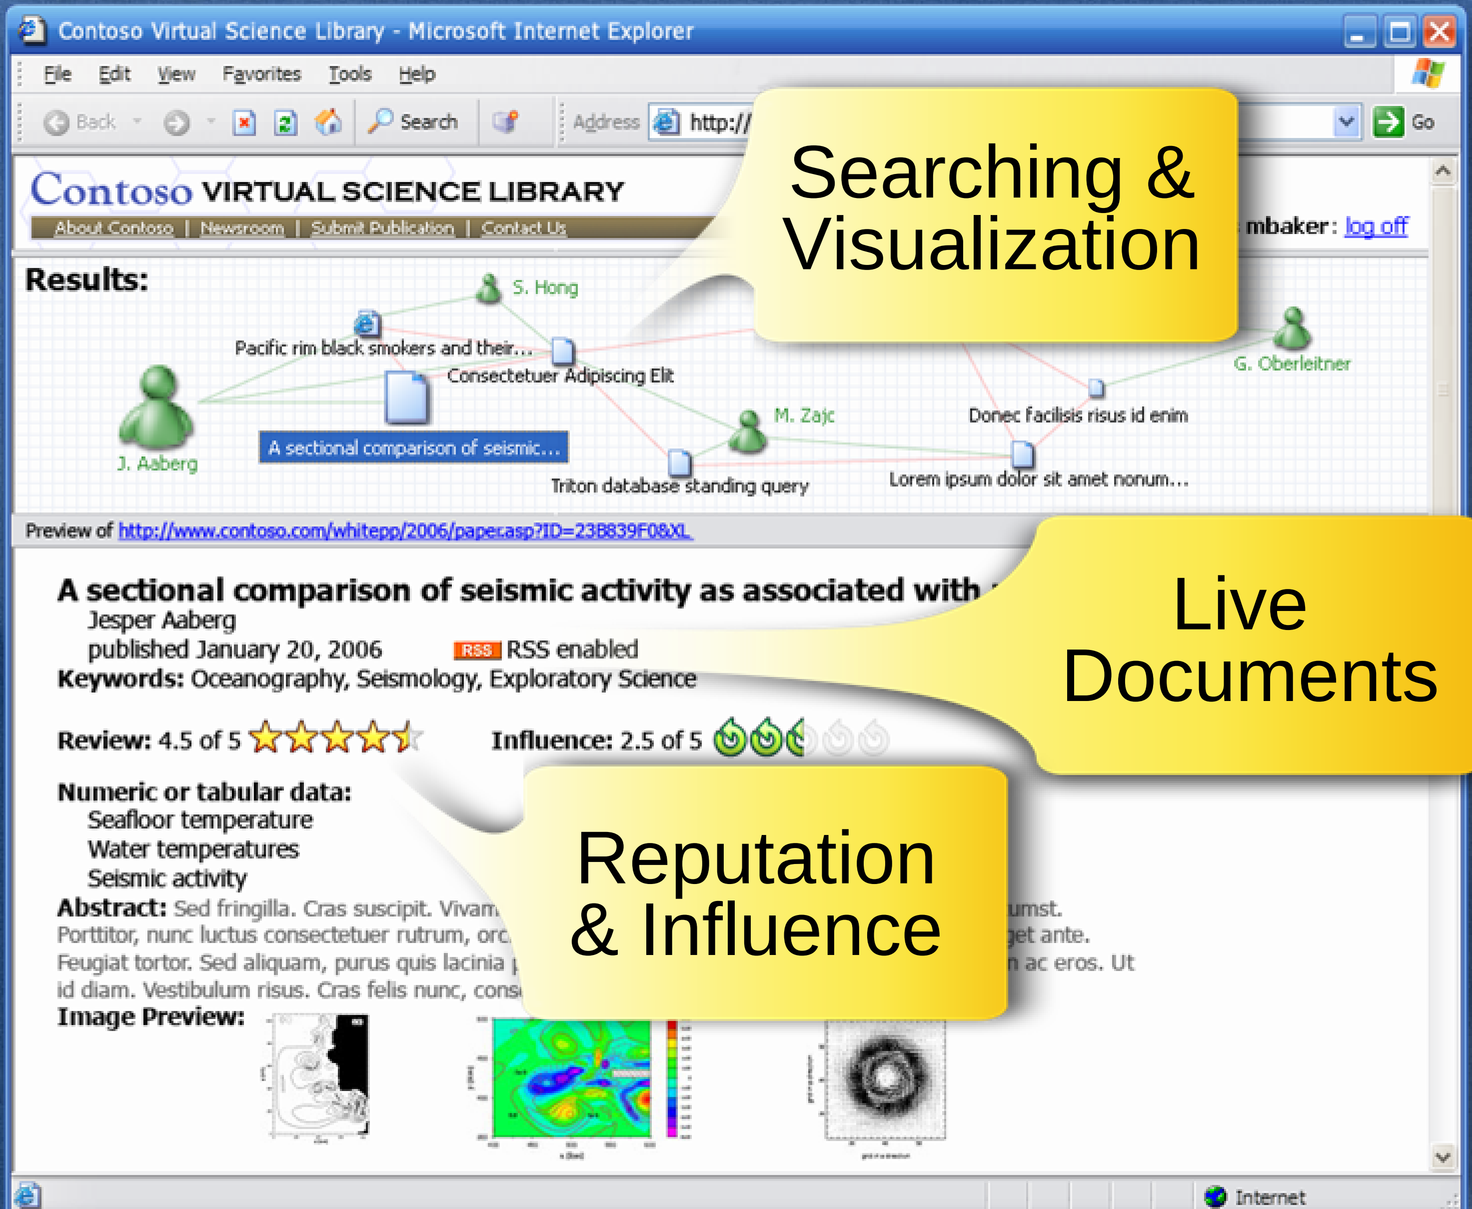
Contoso Virtual Science Library - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Address http:// Go
Contoso VIRTUAL SCIENCE LIBRARY
About Contoso | Newsroom | Submit Publication | Contact Us
mbaker: log off
Results:
S. Hong
Pacific rim black smokers and their...
Consectetuer Adipiscing Elit
G. Oberleitner
M. Zajc
Donec facilisis risus id enim
J. Aaberg
A sectional comparison of seismic...
Triton database standing query
Lorem ipsum dolor sit amet nonum...
Preview of http://www.contoso.com/whitepp/2006/paper.asp?ID=23B839F08XL
A sectional comparison of seismic activity as associated with
Jesper Aaberg
published January 20, 2006
RSS RSS enabled
Keywords: Oceanography, Seismology, Exploratory Science
Review: 4.5 of 5
Influence: 2.5 of 5
Numeric or tabular data:
Seafloor temperature
Water temperatures
Seismic activity
Abstract: Sed fringilla. Cras suscipit.
Image Preview:
Internet
Searching & Visualization
Live Documents
Reputation & Influence

Cyberinfrastructure

- Cyberinfrastructure and e-Infrastructure
  - In the US, Europe and Asia there is a common vision for the 'cyberinfrastructure' required to support the e-Science revolution
  - Set of Middleware Services supported on top of high bandwidth academic research networks
  - Software, hardware and organizations to support e-Science
- Similar to vision of the Grid as a set of services that allows scientists – and industry – to routinely set up 'Virtual Organizations' for their research – or business
  - The 'Microsoft Grid' vision is as much about integrating and managing data and information than about compute cycles

Technical Computing at Microsoft

- Advanced Computing for Science and Engineering
  - Application of new algorithms, tools and technologies to scientific and engineering problems
- High Performance Computing
  - Application of high performance clusters and database technologies to industrial and scientific applications
- Radical Computing
  - Research in potential breakthrough technologies

Fighting HIV with Computer Science

Nebojsa Jojic and David Heckerman

- A major problem: Over 40 million infected
  - Drug treatments are effective but are an expensive life commitment
- Vaccine needed for third world countries
  - Effective vaccine could eradicate disease
- Methods from computer science are helping with the design of vaccine
  - Machine learning: Finding biological patterns that may stimulate the immune system to fight the HIV virus
  - Optimization methods: Compressing these patterns into a small, effective vaccine

Developed Set of Specialist Tools

- Chromatogram deconvolution
- Pathway analysis/association/causal models
- Clustering/Trees (phylo, haplotypes etc.)
- Protein binding and folding
- Sequence diversity models (epitomes)
- Image analysis/classification
- Evolution modeling and inference
- Epitope prediction

HIV: The diabolical virus

The train-and-kill mechanism doesn’t work for HIV – the virus adapts through rapid mutation. As soon as the killer cells get the upper hand, the epitopes start changing.

Strategy:

- Find peptides or epitopes that occur commonly across a *population* of HIV viruses
- Compact the known or potential immune targets into a small vaccine

International Virtual Observatory

- Data has no commercial value
  - No privacy concerns
  - Can freely share results with others
  - Great for experimenting with algorithms
- Data is real and well documented
  - High-dimensional data
  - Spatial data
  - Temporal data
- Data from many different instruments, places and times
  - Federation is a key goal
  - There is a lot of data (petabytes)

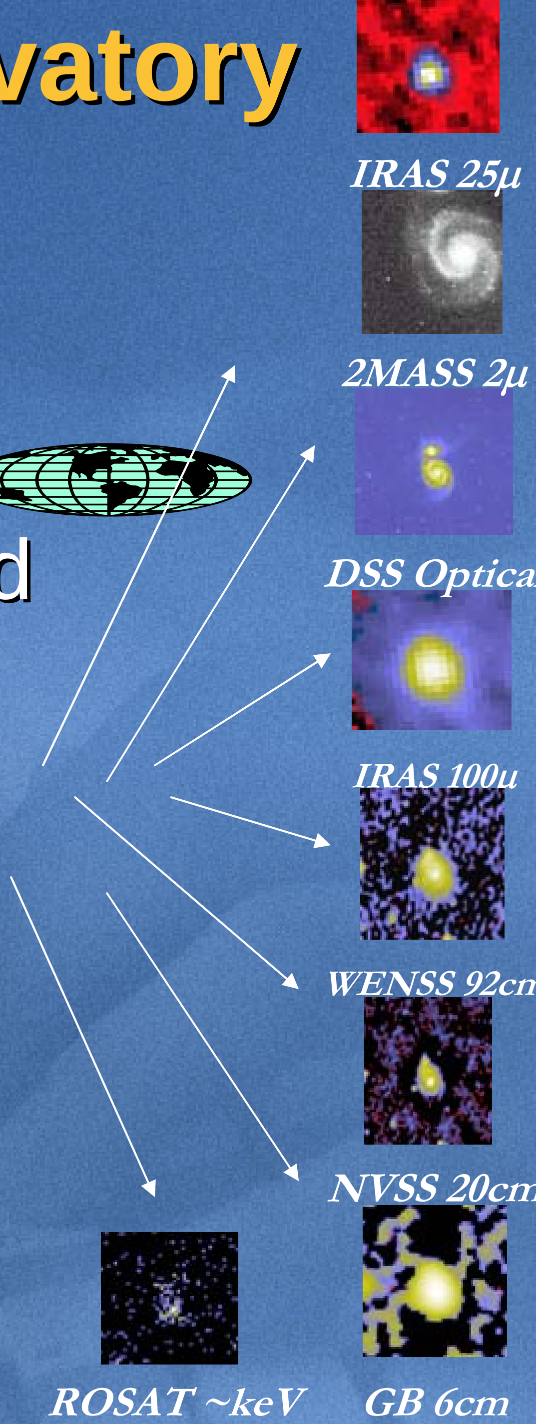

With thanks to Jim Gray

The Multiwavelength Crab Nebulae

Crab star
1053 AD

X-ray,
optical,
infrared, and
radio
views of the nearby
Crab Nebula, which is
now in a state of chaotic
expansion after a
supernova explosion
first sighted in 1054
A.D. by Chinese
Astronomers.

Slide courtesy of Robert Brunner @ CalTech.

SkyServer (http://cas.sdss.org)

- A modern archive
  - Access to Sloan Digital Sky Survey Spectroscopic and Optical surveys
  - Raw Pixel data lives in file servers
  - Catalog data (derived objects) lives in Database
  - Online query to any and all
- Interesting things
  - Spatial data search
  - Query interface via Java Applet
  - Query from Emacs, Python, ….
  - Template design cloned by other surveys
  - Web Services are core of it

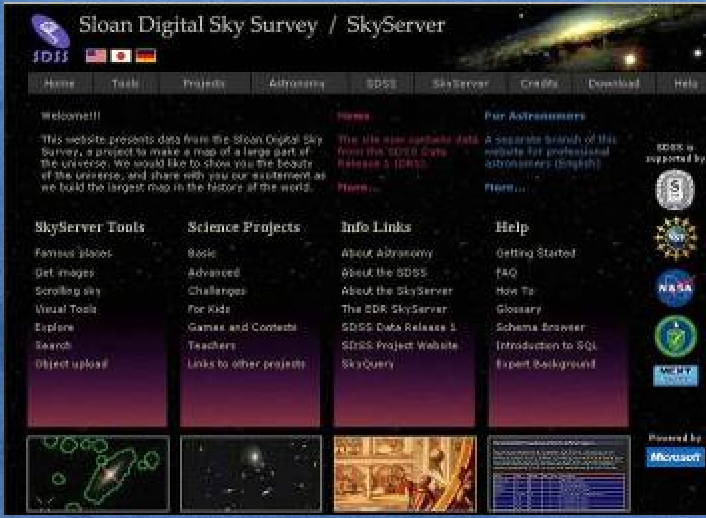

SkyQuery (http://skyquery.net/)

- Distributed Query tool using a set of Web Services
- Federates many astronomy archives from Pasadena, Chicago, Baltimore, Cambridge UK
- Grown from 4 to 15 archives,becoming international standard
- WebService 'Poster Child'
- Allows queries like:

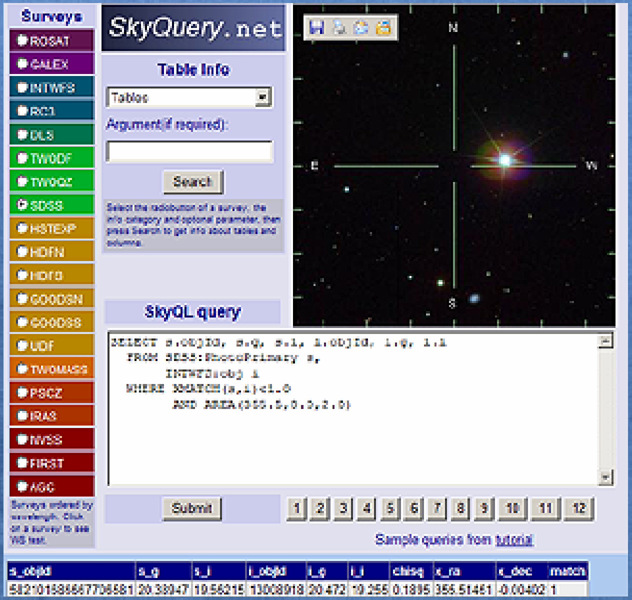


```
SELECT o.objId, o.r, o.type, t.objId
  FROM SDSS:PhotoPrimary o,
       TWOMASS:PhotoPrimary t
  WHERE XMATCH(o,t)<3.5
       AND AREA(181.3,-0.76,6.5)
        AND o.type=3 and (o.I - t.m_j)>2
```

IVO: An Astronomy Data Grid

- Working to build world-wide telescope
  - All astronomy data and literature
  - online and cross indexed
  - Tools to analyze it
- Built SkyServer.SDSS.org
- Built Analysis system
  - MyDB
  - CasJobs (batch job)
- OpenSkyQuery
  Federation of ~20 observatories.
- Results:
  - It works and is used every day
  - Spatial extensions in SQL 2005
  - A good example of Data Grid
  - A good example of Web Services

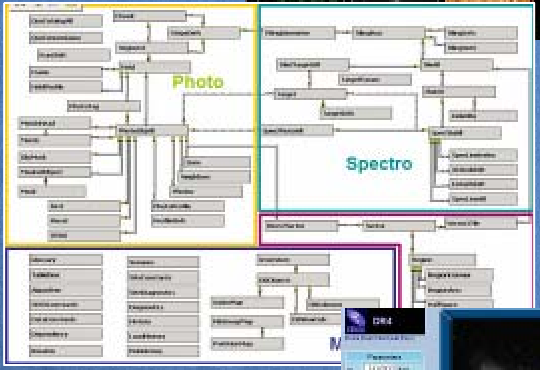

HPC: Top 500 Trends

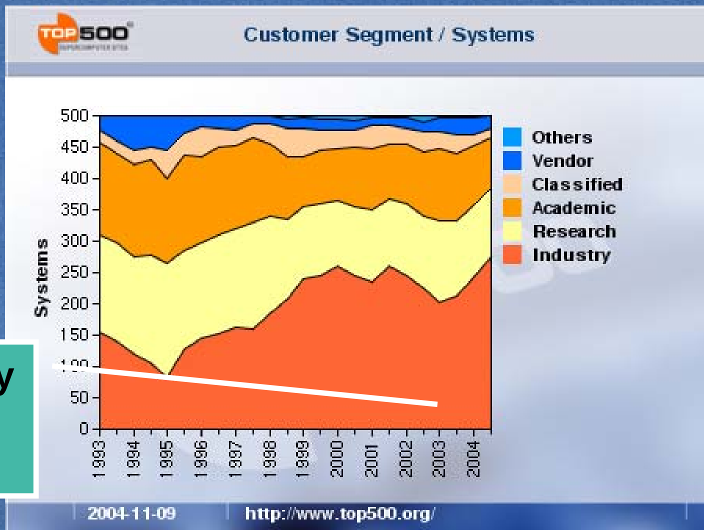

Industry usage rising

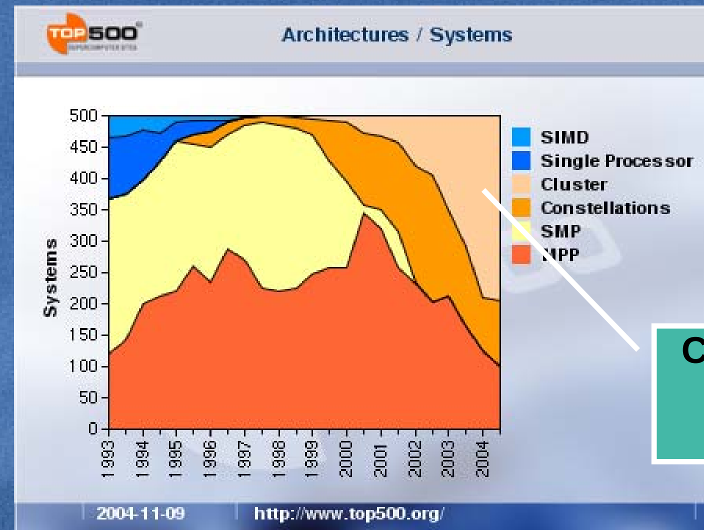

Clusters over 50%

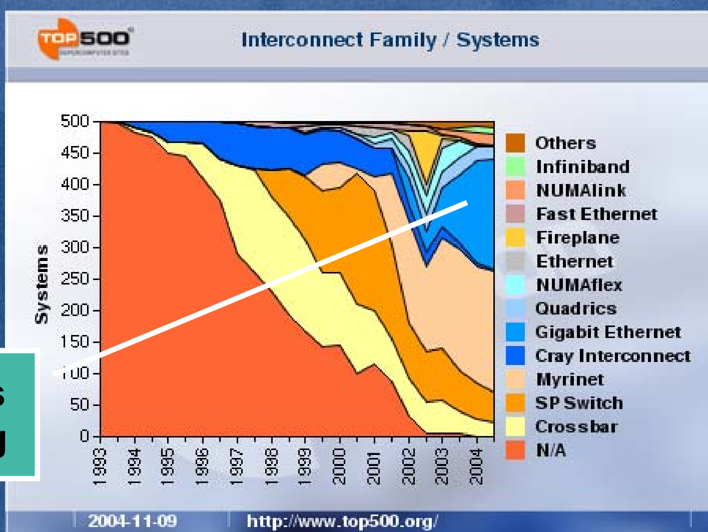

GigE is gaining

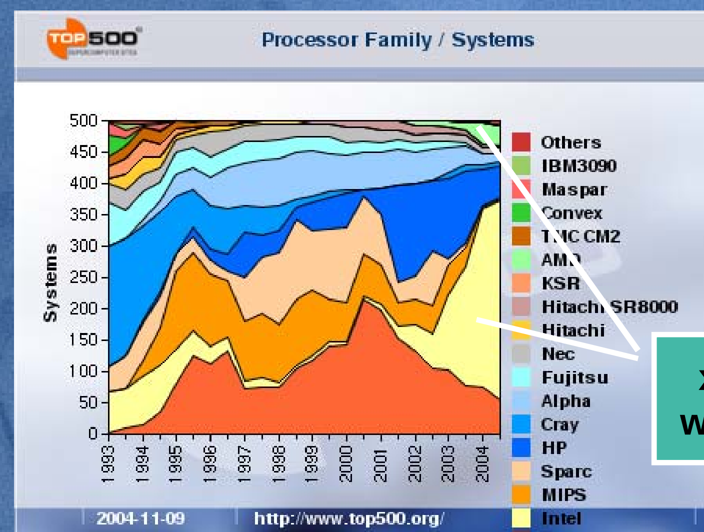

x86 is winning

HPC: Market Trends

2004 Systems

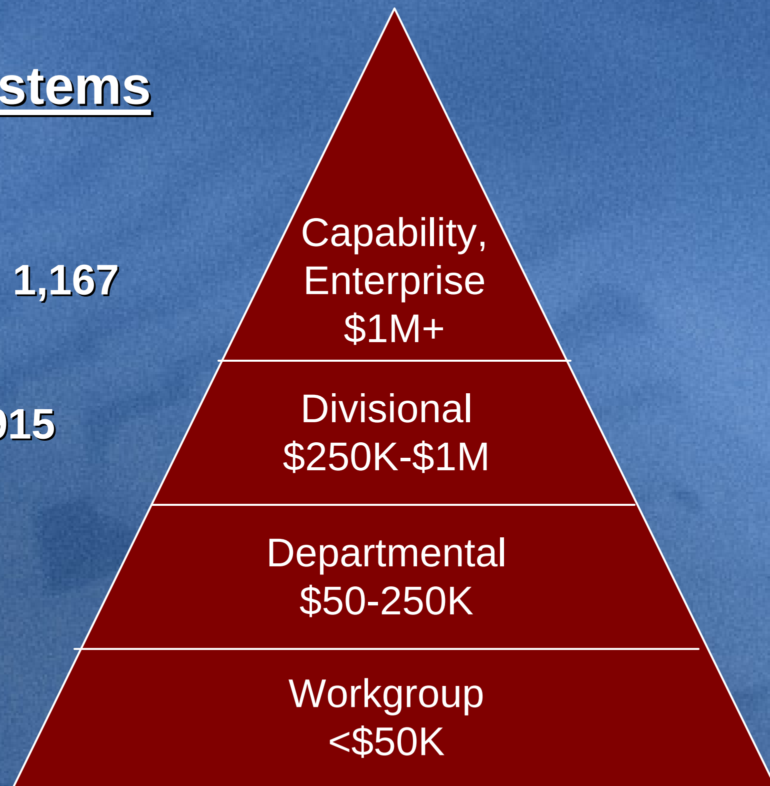

1,167

3,915

22,712

127,802

<$250K – 97% of systems, 52% of revenue

In 2004 clusters grew 96% to 37% by revenue
Average cluster size 10-16 nodes

Source: IDC, 2005

Continuing Trend Towards Decentralized, Networked Resources

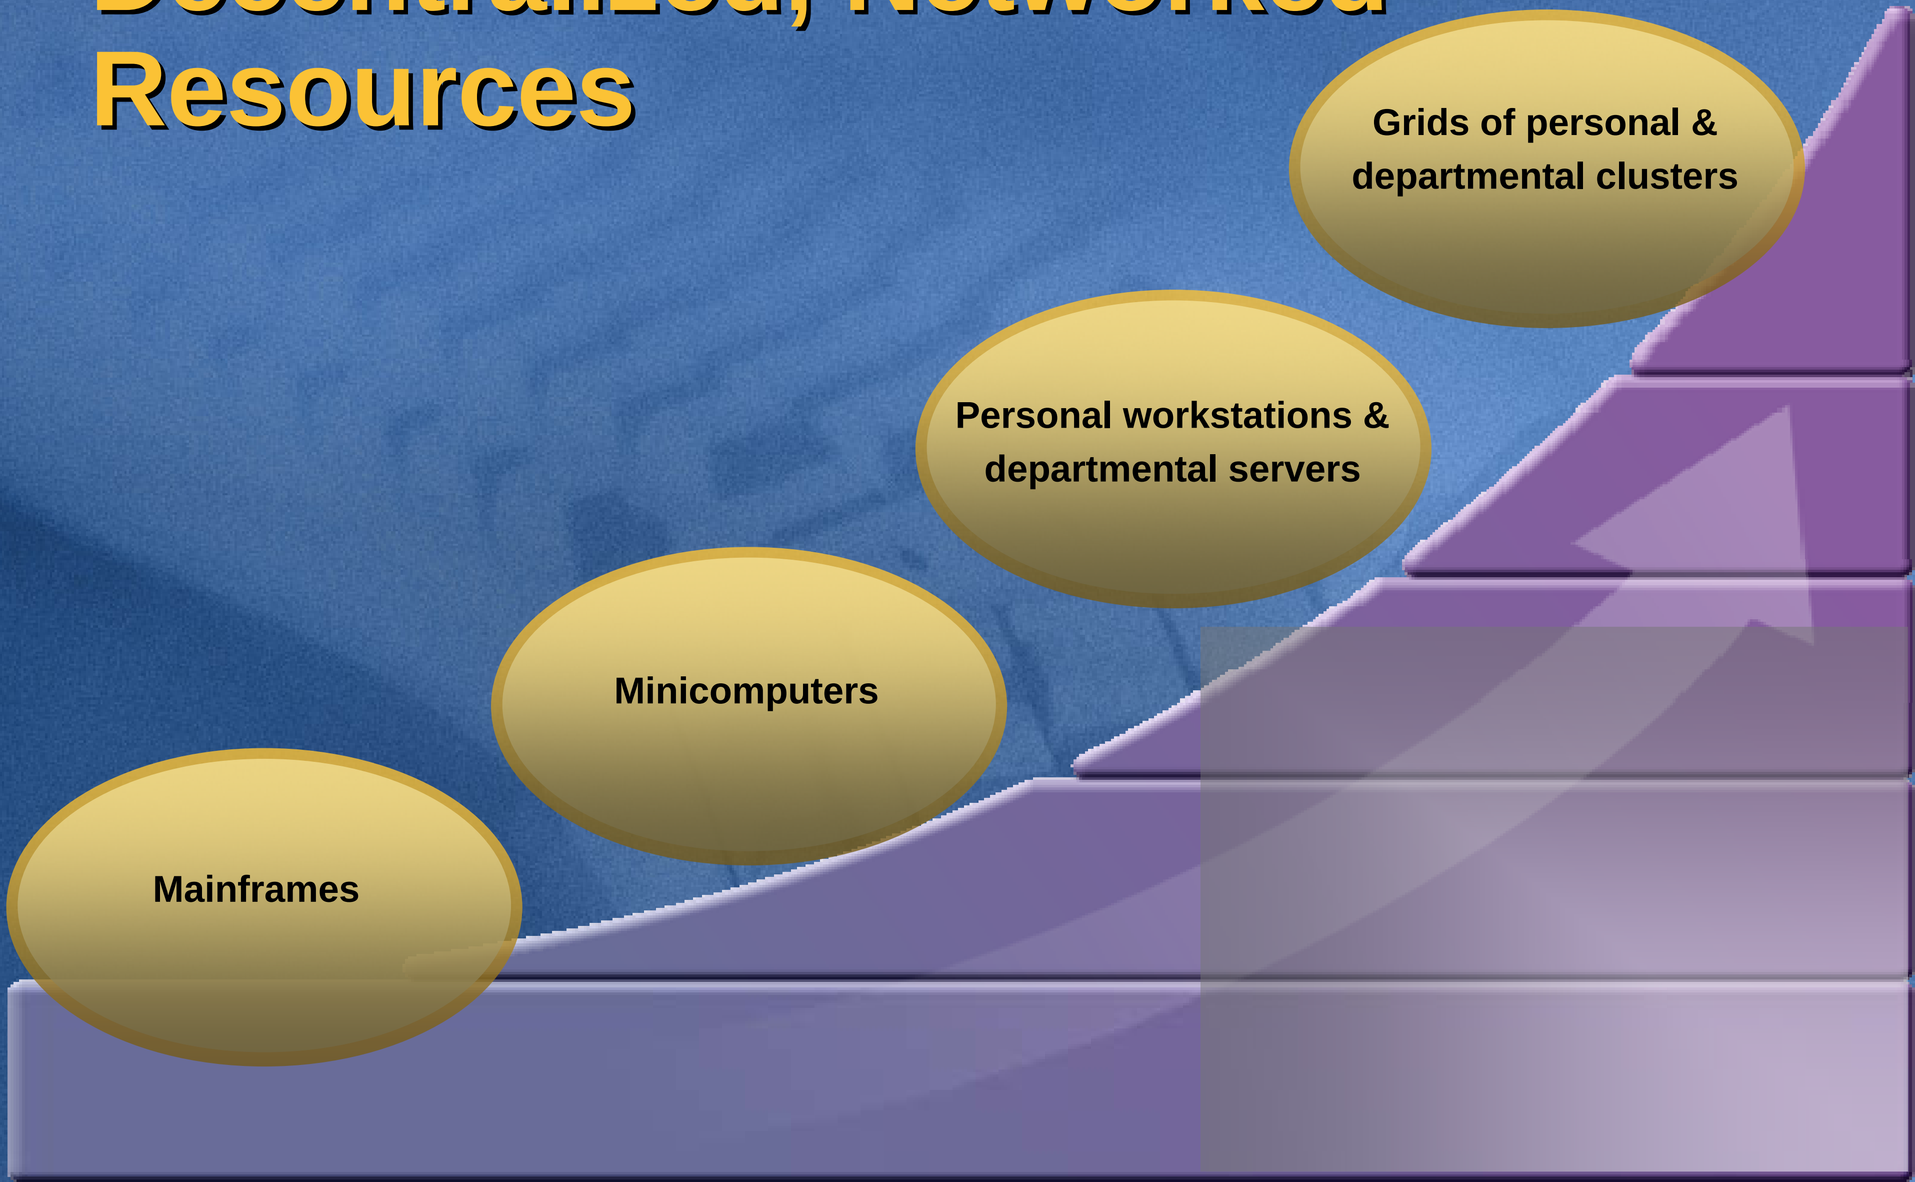

Microsoft Strategy for HPC

- Reduce barriers to adoption for HPC clusters
  - **Easy to deploy, manage and use**
- Provide application support in key HPC verticals
  - **Engagement with the top HPC ISVs**
- Leverage a breadth of standard tools
  - **Web Services, SQL, Sharepoint, Infopath, Excel**
- High Volume Market
  - **Enable broad HPC adoption**

Today's CPU Architecture
Heat becoming an unmanageable problem

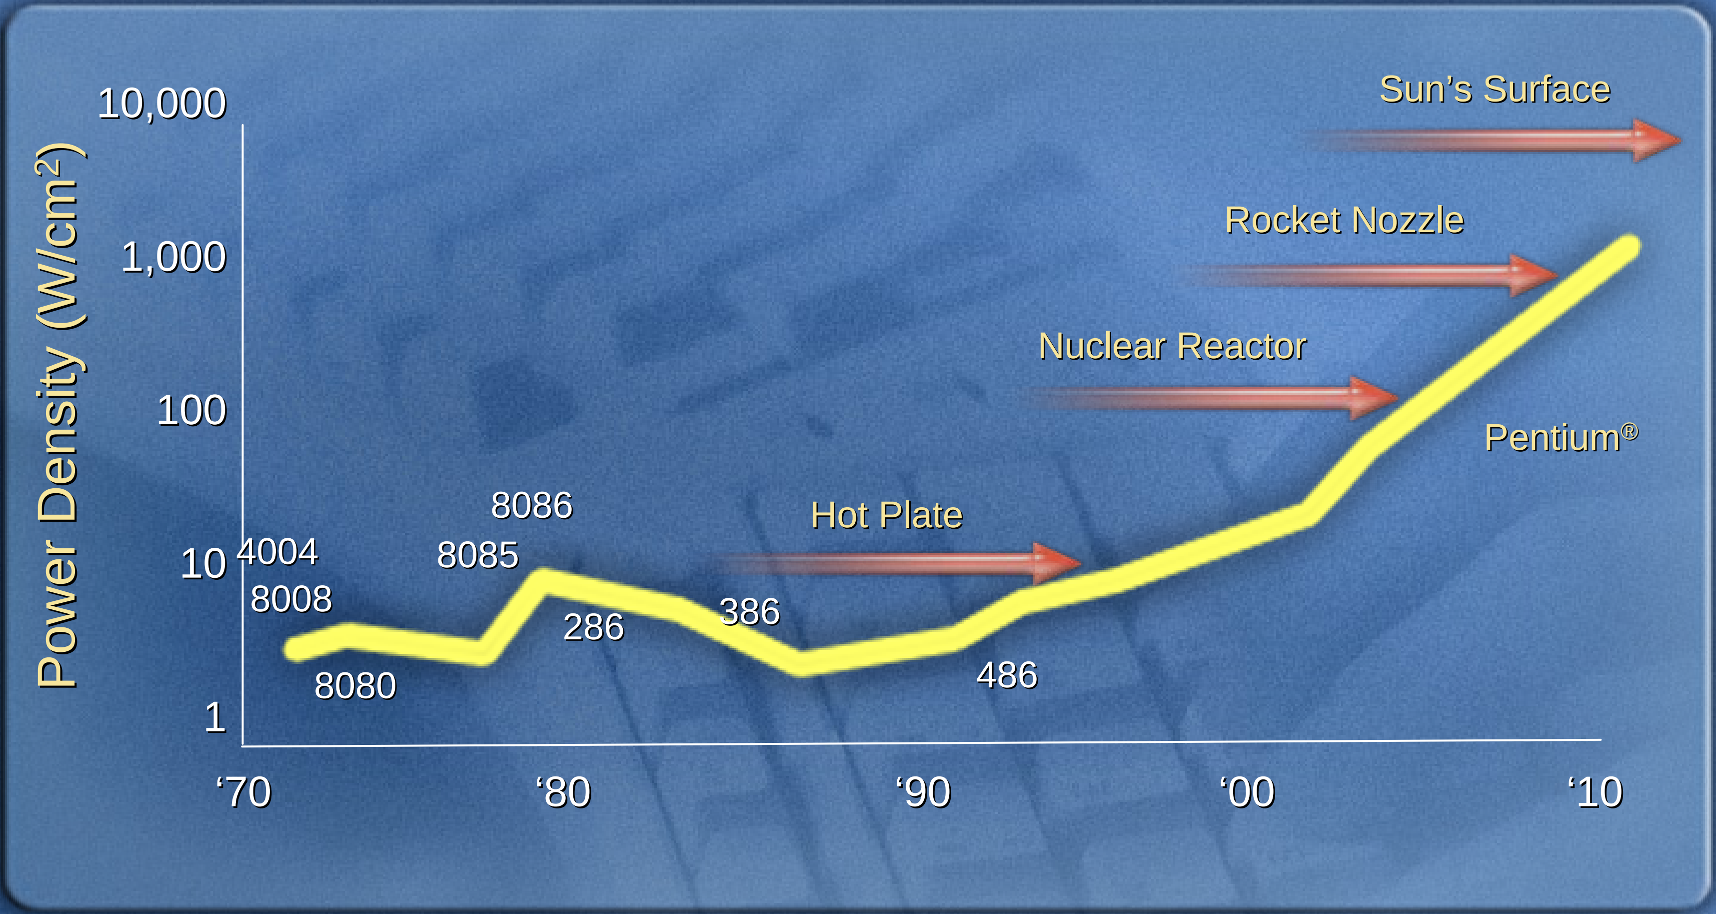

Radical Computing

- **The end of Moore's Law as we know it**
  - Number of transistors on a chip will continue to increase
  - No significant increase in clock speed
- **Future of silicon chips**
  - "100's of cores on a chip in 2015" (Justin Rattner, Intel)
  - "4 cores"/Tflop => 25 Tflops/chip
- **Challenge for IT industry and Computer Science community**
  - Can we make parallel computing on a chip easier than message-passing?

Service-Orientation for building Distributed Systems

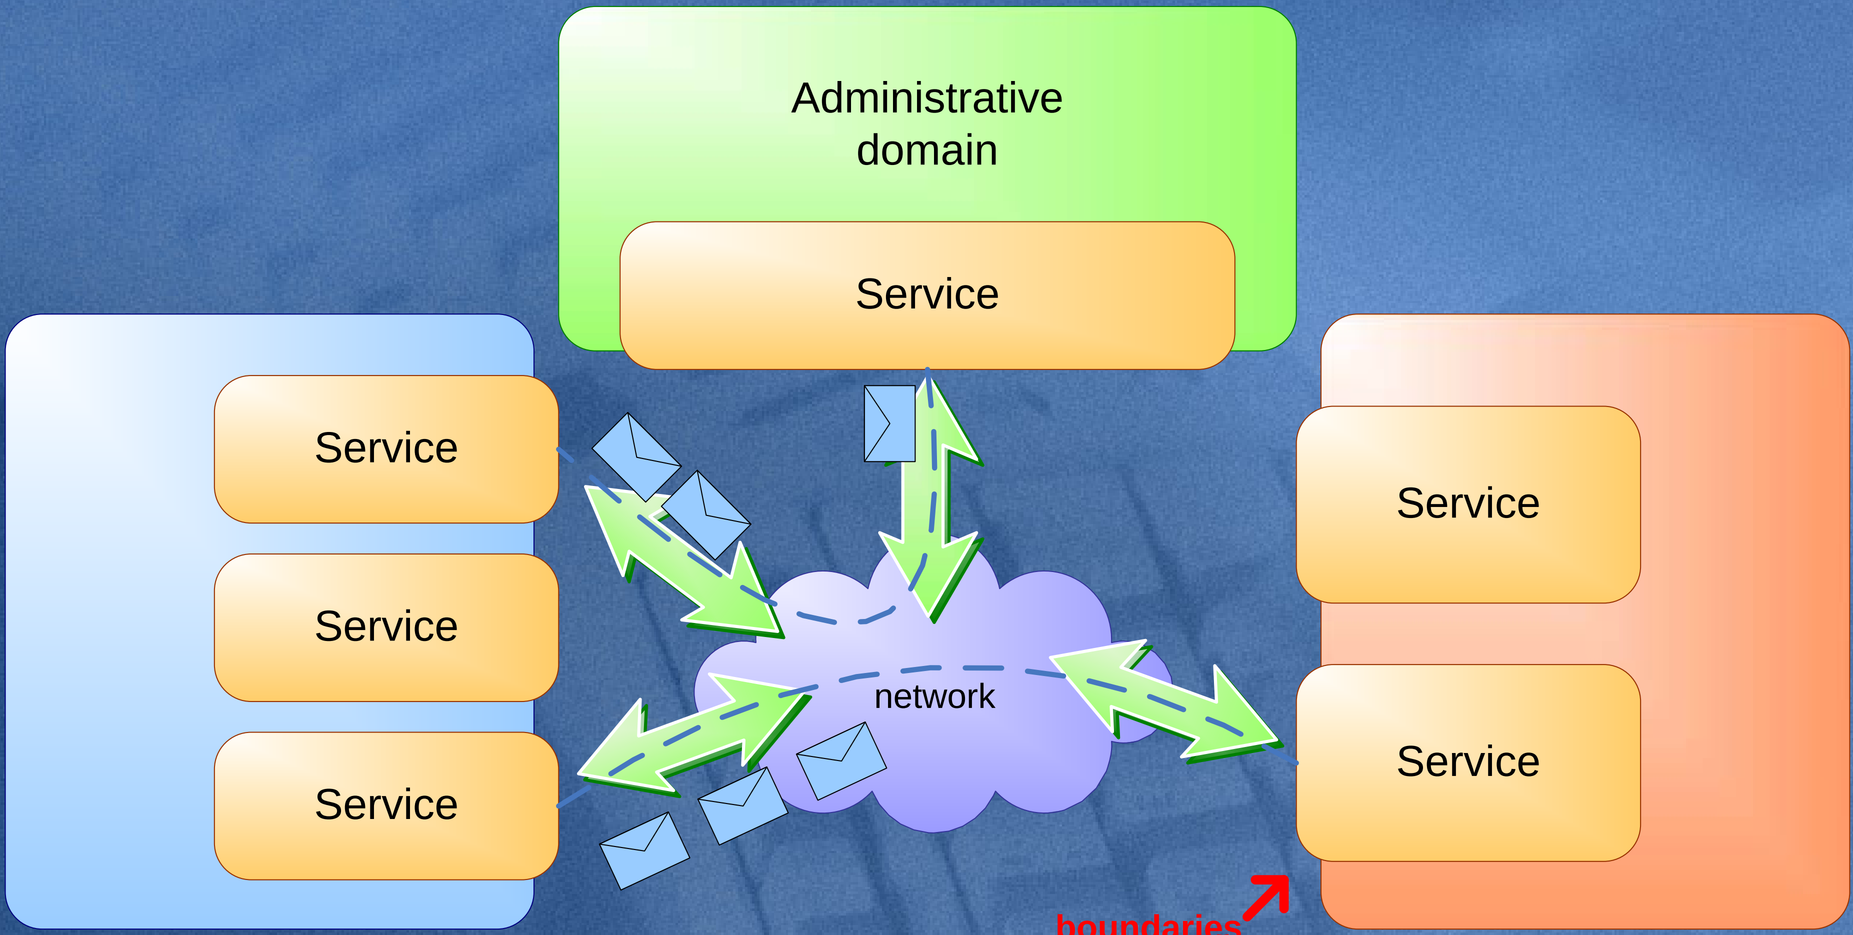

The Web Services ‘Magic Bullet’

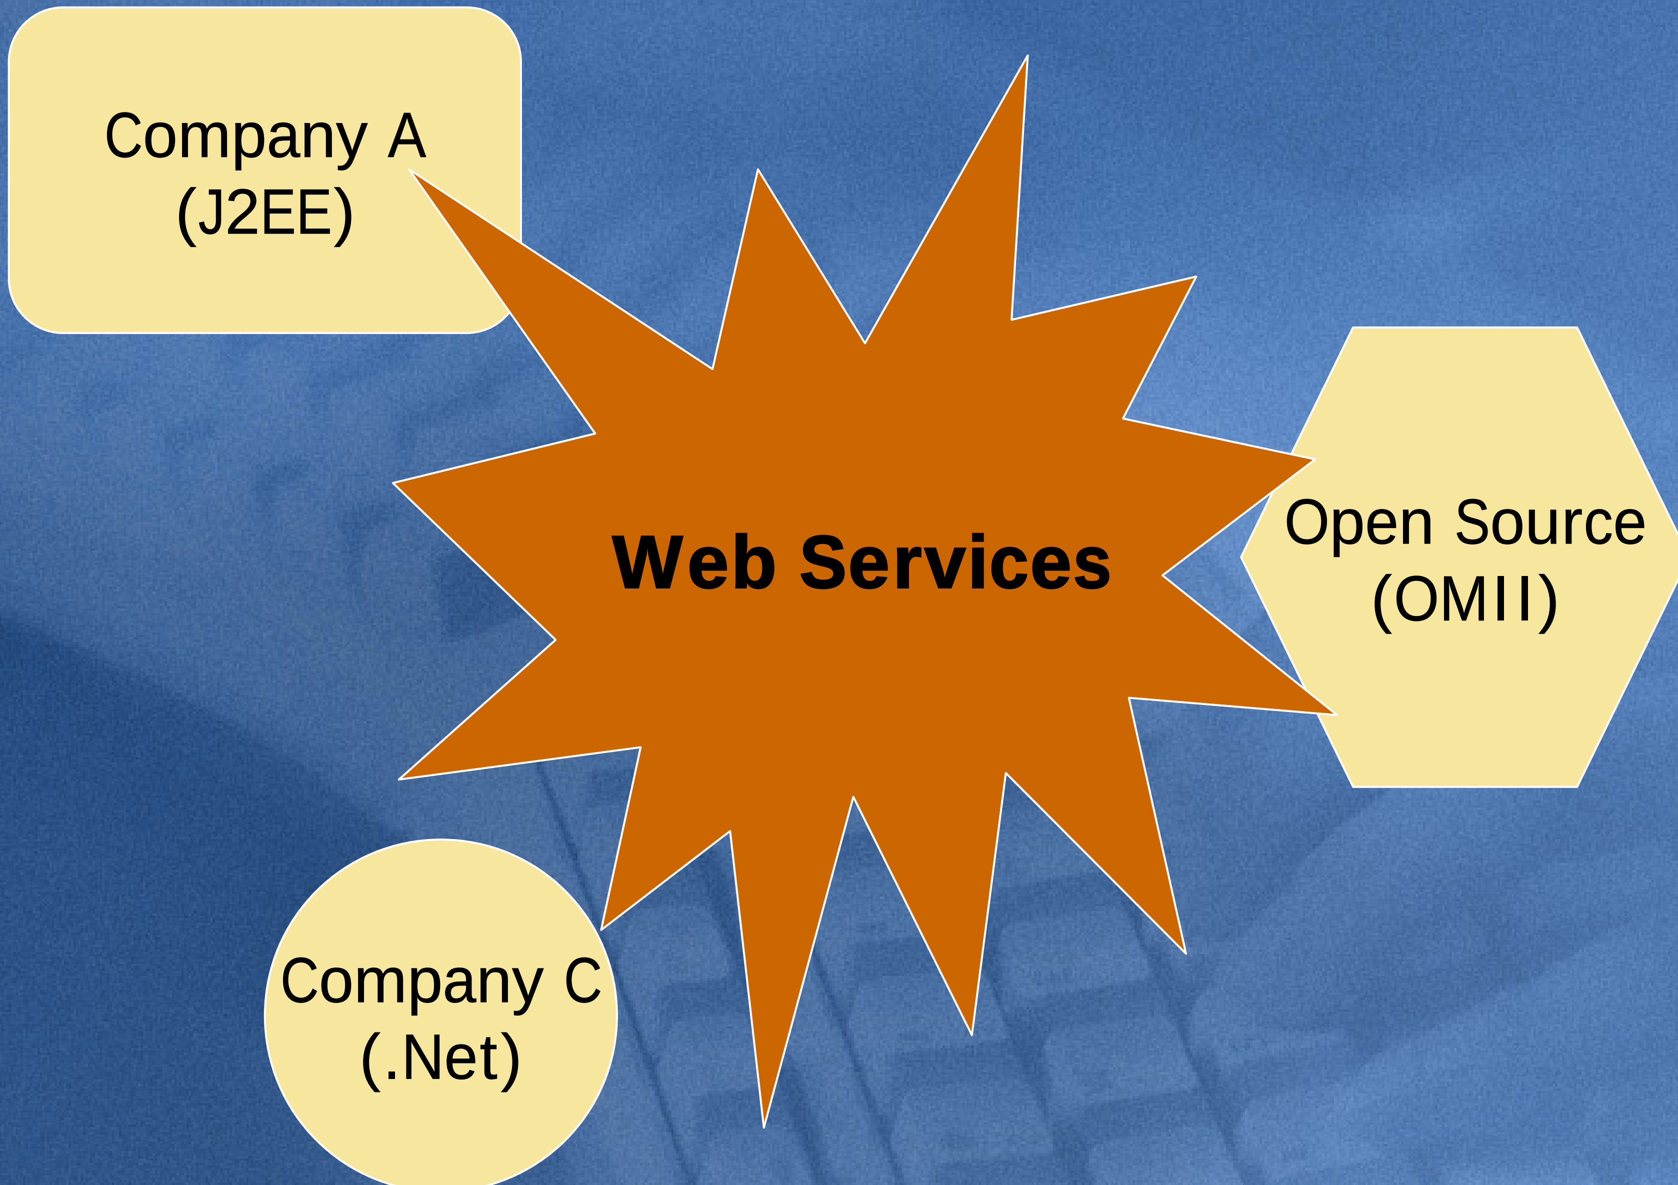

Convergence in Web Services Systems Management

- Different approaches lead to confusion and uncertainty
  - WS-DM and WS-Management
  - WS-RF and WS-Transfer
  - WS-Notification and WS-Eventing
- Microsoft, IBM, HP, and Intel agreed to a convergence roadmap
  - No specific timeline yet announced

The Web Services Ecosystem

Specifications that have/will enter a standardisation process
but are not stable and are still experimental

stable
profile

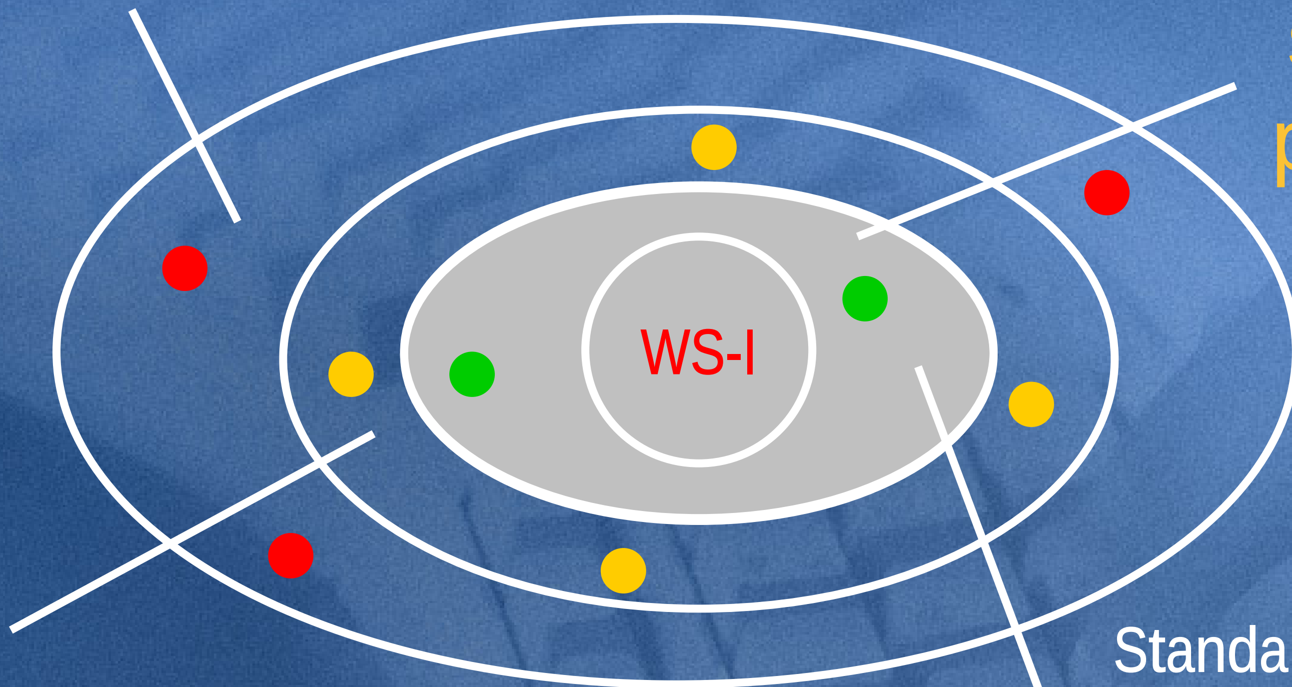

Specifications that are emerging
from standardisation process
and are recognised as being ‘useful’

Standards that have
broad industry support
and multiple interoperable
implementations

Web Services and the Grid

A Complicated Story:

- Basic Web Service specifications
  - WS-I (SOAP, WSDL) from 2001 onwards
- Web Service Grids
  - G-WSDL and OGSI (2001 – 2003)
  - WS-RF, WS-N and WS-DM (2004 - ?)
- Lesson:

  Build Web Service Grids incrementally only on stable, mature and widely-accepted WS foundations

Grids for Virtual Organizations

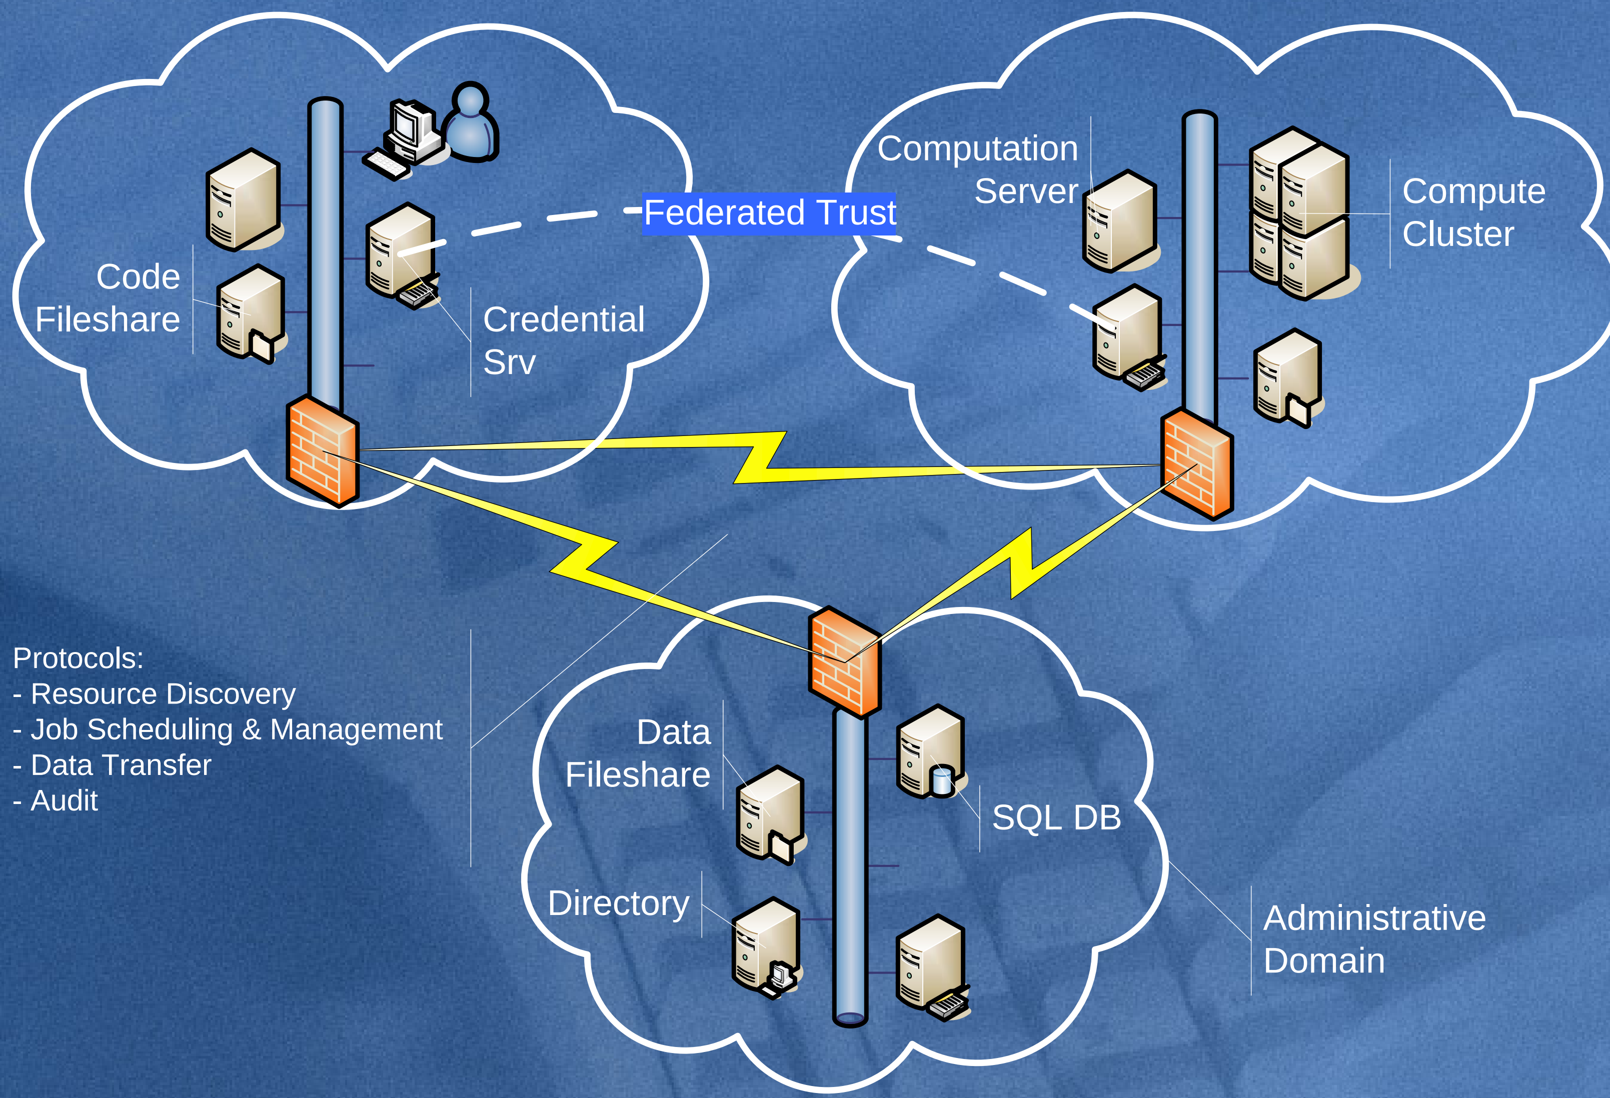

Grids for Virtual Organizations

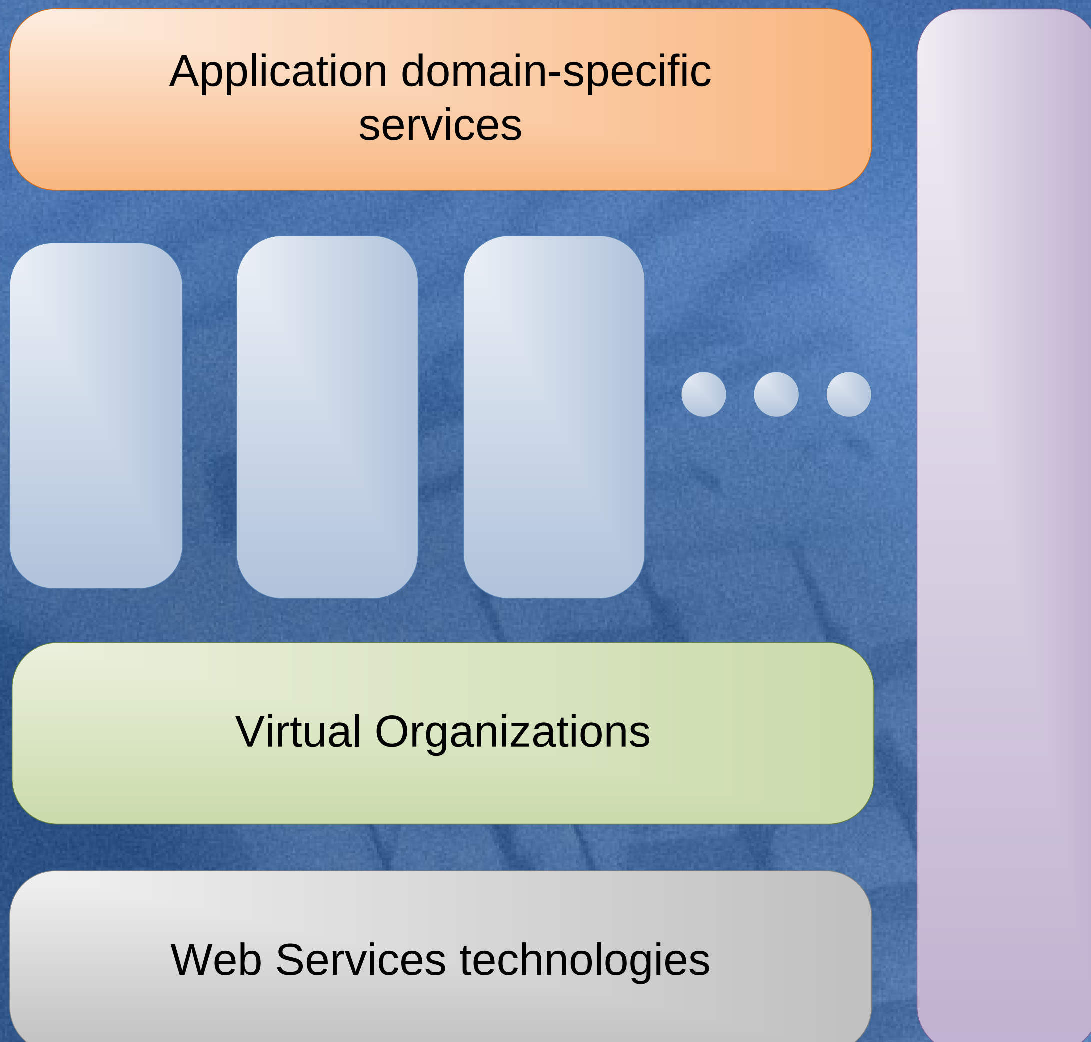

Premise: The Grid and Web communities could soon deliver some useful specifications for Web Service Grids

- By focusing on simple Grid services built on accepted Web Services we can reach agreement quickly
- Look at three key areas for Grids for Virtual Organizations
  - Security
  - HPC Services
  - Data Services

Virtual Organization Security

- Not yet routine and seamless: many technologies and standards exist in the security space
- Interoperability only works if proposed solutions are widely accepted by both industry and academia
- Larger problem than just for the GGF community
- IT industry will provide high quality, well documented tooling and services to construct secure Virtual Organizations

The OGSA HPC Profile

- Defines a minimalist base interface plus optional extensions
  - Small base interface enables simple interoperability widely and quickly
  - Common use cases covered by extensions
  - Extension model enables principled experimentation and evolution
- Defines minimal set of composable, extensible services
  - Job Submission
  - Data Staging

An OGSA Data Profile?

Guiding principles:

- Keep profile as simple as possible
  - Example of Amazon S3
- DAIS Working Group specifications
  - WS-DAI
  - WS-DAIR and WS-DAIX
- Build on only widely accepted Web Services
  - WS-I + ….

New Science Paradigms

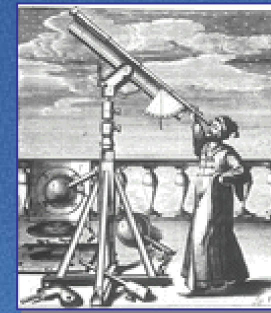

- **Thousand years ago:**
  Experimental Science
  - description of natural phenomena
- **Last few hundred years:**
  Theoretical Science
  - Newton's Laws, Maxwell's Equations …

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{4\pi G\rho}{3} - K\frac{c^2}{a^2}$$

- **Last few decades:**
  Computational Science
  - simulation of complex phenomena

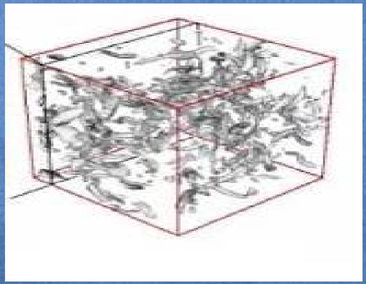

- **Today:**
  e-Science or Data-centric Science
  - unify theory, experiment, and simulation
  - using data exploration and data mining
  - Data captured by instruments
  - Data generated by simulations
  - Processed by software
  - Scientist analyzes databases/files

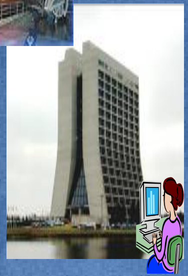

(With thanks to Jim Gray)

Key Data Issues for e-Science

- Networks
  - Lambda technology
- The Data Life Cycle
  - From Acquisition to Preservation
- Scholarly Communication
  - Open Access to Data and Publications

An International e-Infrastructure

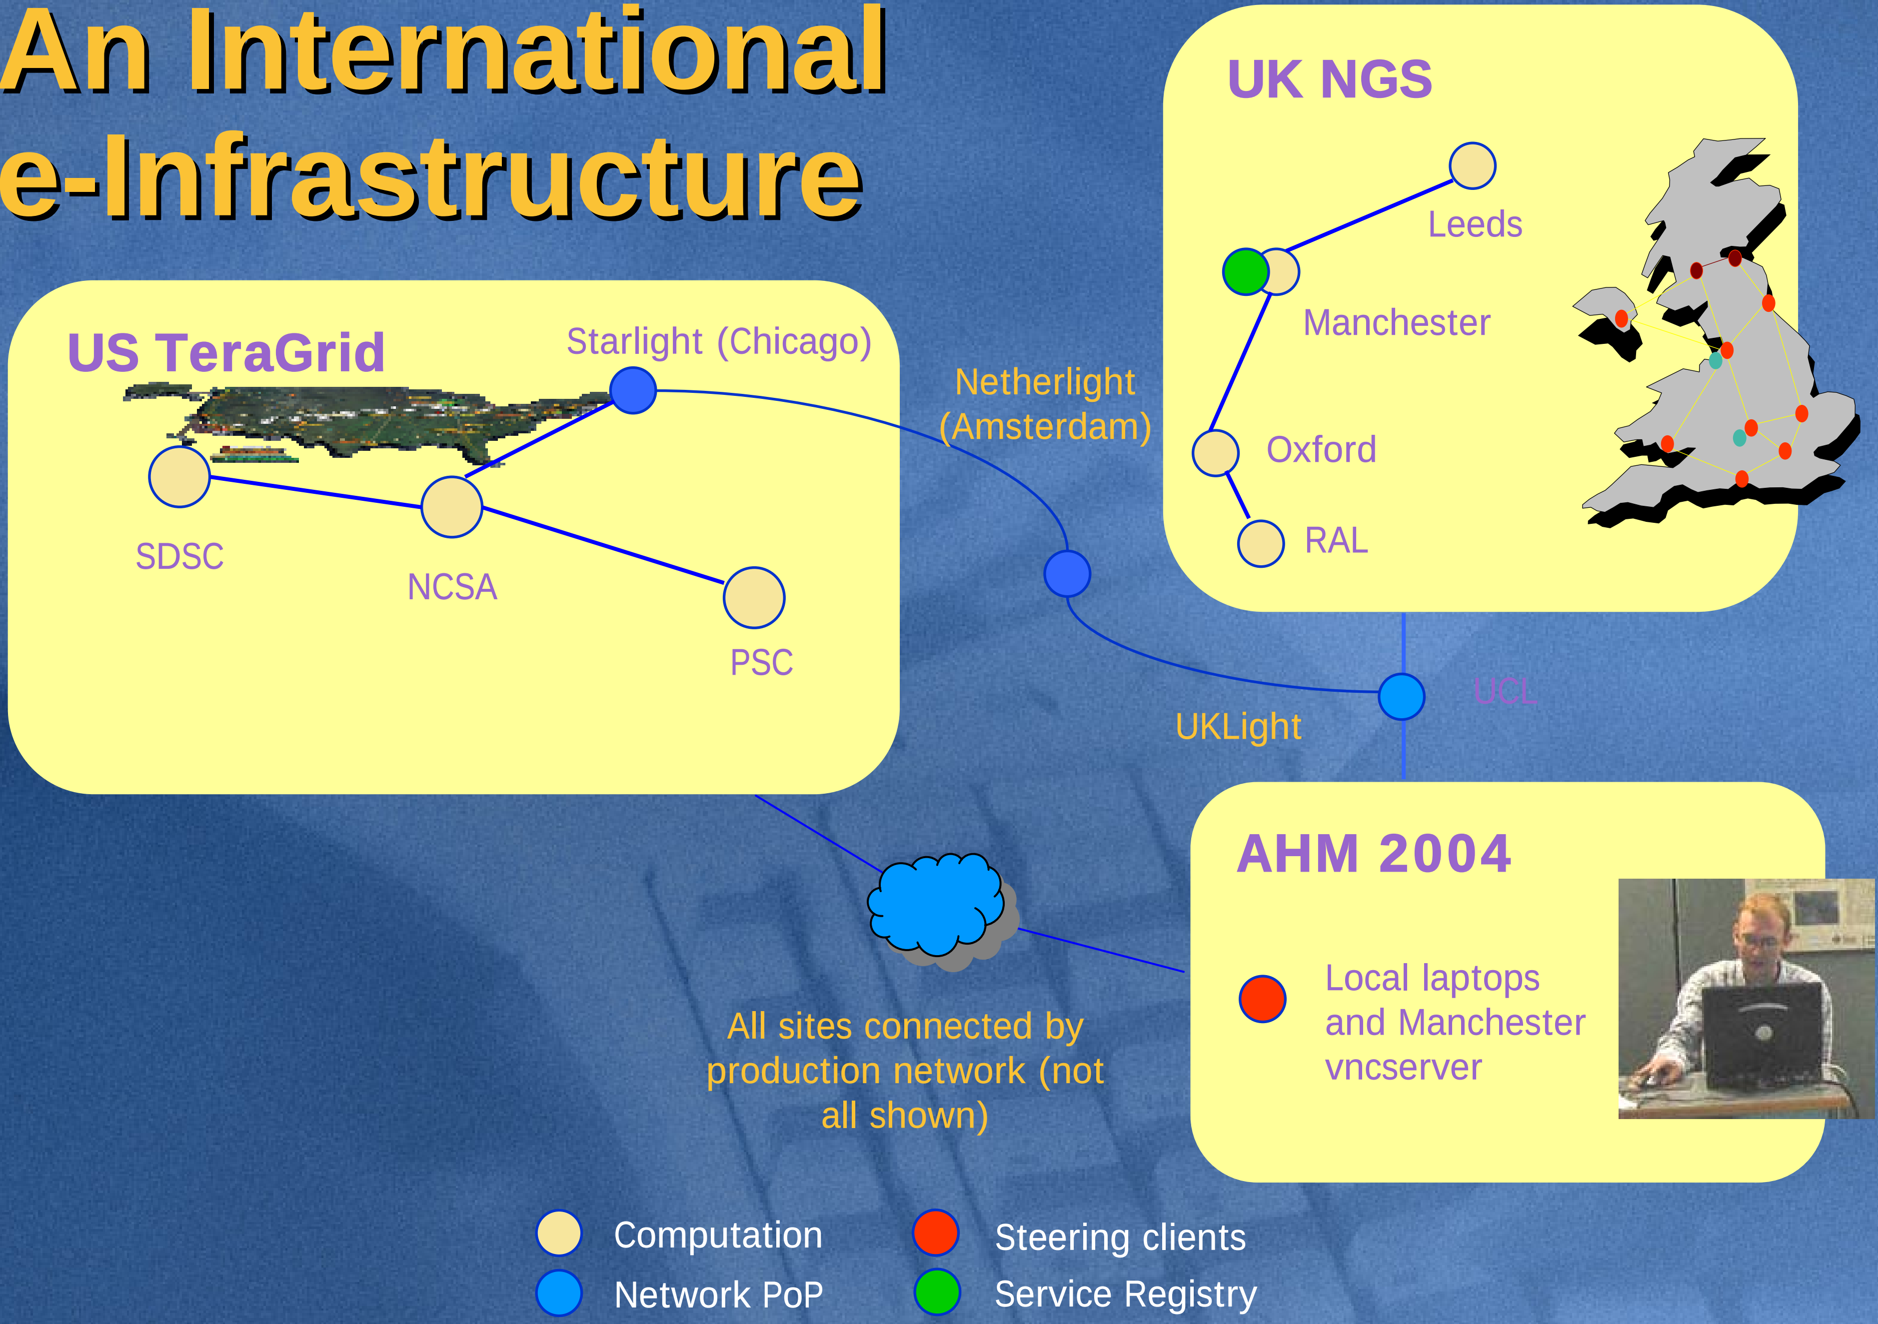

The Problem for the e-Scientist

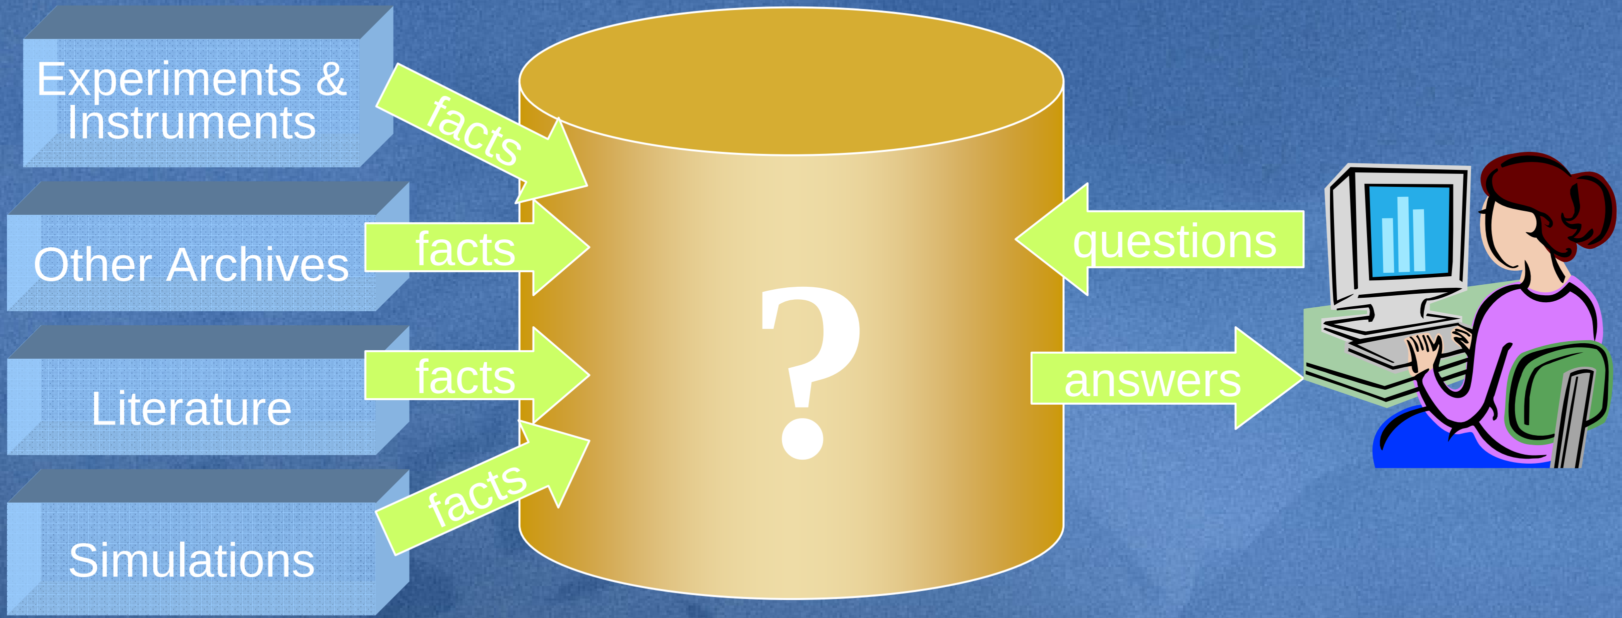

- Data ingest
- Managing a petabyte
- Common schema
- How to organize it?
- How to *re*organize it?
- How to coexist & cooperate with others?
- Data Query and Visualization tools
- Support/training
- Performance
  - Execute queries in a minute
  - Batch (big) query scheduling

The e-Science Data Life Cycle

- Data Acquisition
- Data Ingest
- Metadata
- Annotation
- Provenance
- Data Storage
- Data Cleansing
- Data Mining
- Curation
- Preservation

Scholarly Communication

- Global Movement towards permitting ‘Open Access’ to scholarly publications
  - Libraries can no longer afford publisher subscriptions
  - Principle that results of publicly funded research should be available to all
  - First World/Third World issue
- Open Archive Initiative (OAI)
  - Creation of ‘Subject Repositories’ such as arXiv for physics, astronomy and computer science, and PubMedCentral for Bio-Medical area
  - Global network of ‘Institutional Repositories’ being established using software such as MIT’s DSpace, Southampton’s EPrints and others

NSF ‘Atkins’ Report on Cyberinfrastructure

- ‘the primary access to the latest findings in a growing number of fields is through the Web, then through classic preprints and conferences, and lastly through refereed archival papers’
- ‘archives containing hundreds or thousands of terabytes of data will be affordable and necessary for archiving scientific and engineering information’

The Service Revolution

- Web 2.0
  - Social networks, tagging for sharing e.g. Flikr, Del.icio.us, MySpace, …
  - Wikis, Blogs, RSS …
- Software delivered as a service
  - Live services
    - Microsoft Office Live
    - XboxLive
    - AcademicLive
  - Mashups
    - Craigslist + GoogleMap
    - http://mashupcamp.com

An e-Science Mashup

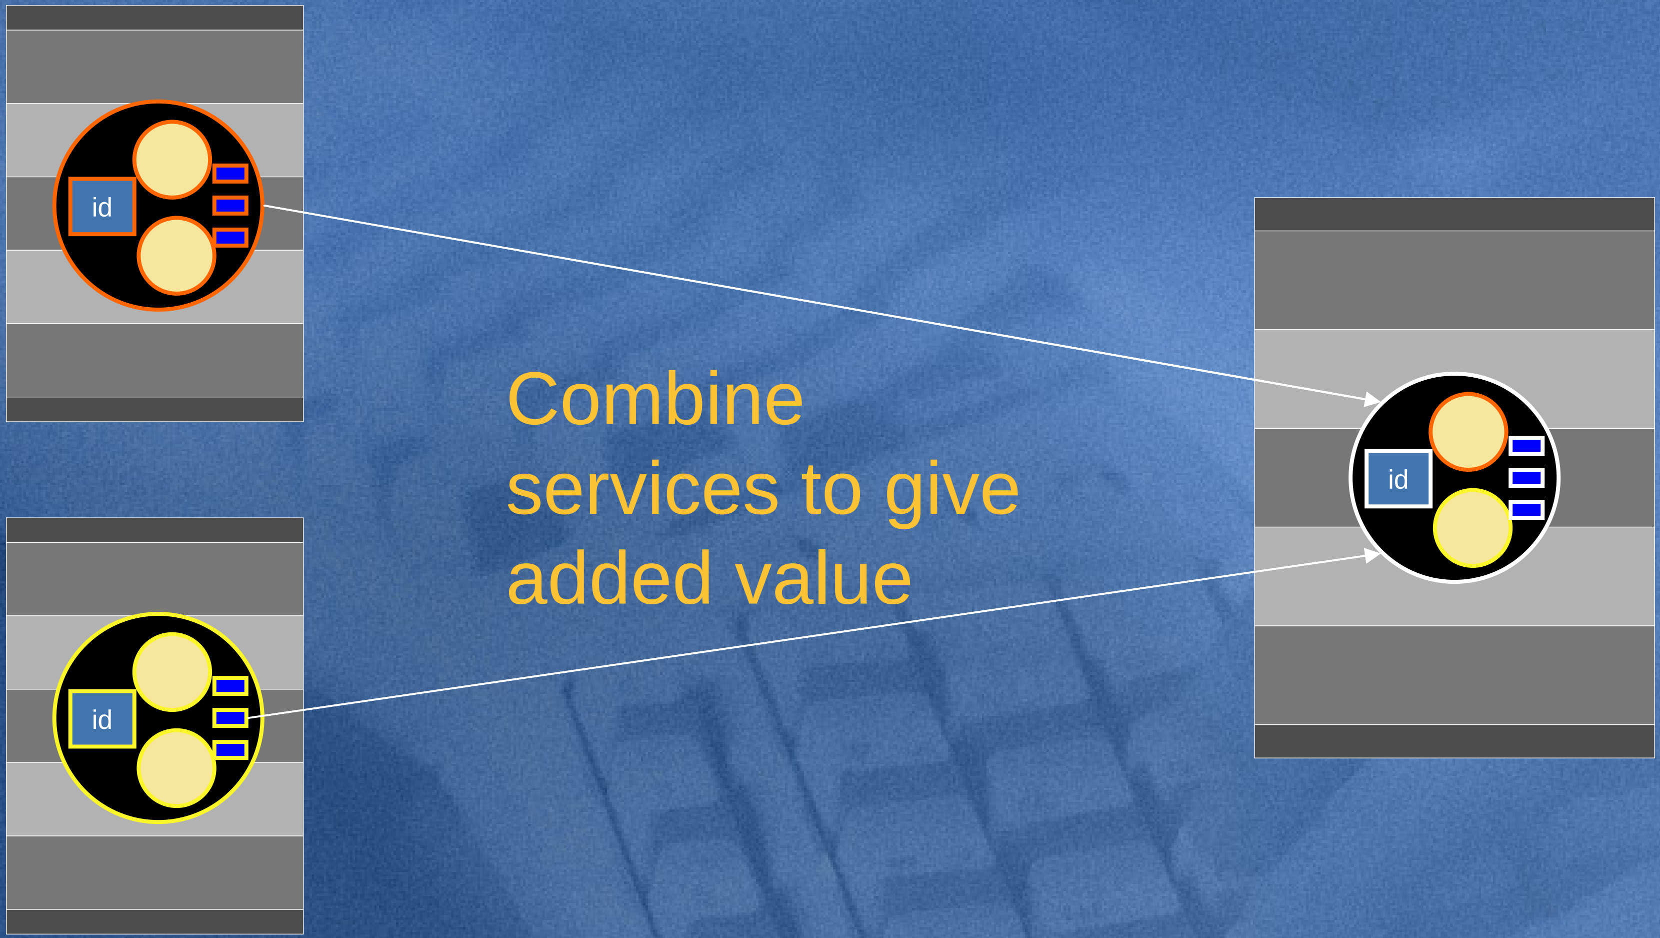

Combine services to give added value

The Semantic Grid

- In 2001, De Roure, Jennings and Shadbolt introduced the notion of the Semantic Grid
  - Advocated 'the application of Semantic Web technologies both *on* and *in* the Grid'
- Argued that users now required interoperability across time as well as space
  - Would allow both anticipated and unanticipated reuse of services, information and knowledge
- In 2005, experience with UK e-Science Projects led them to enumerate requirements for a Semantic Grid

The Semantic Grid and Web Science

- De Roure, Jennings and Shadbolt identified 5 key technologies for building a Semantic Grid:
  - 1) Web Services
  - 2) Software Agents
  - 3) Metadata
  - 4) Ontologies and Reasoning
  - 5) Semantic Web Services
- Web and Grid communities coming together in a common vision for high level semantic services connecting distributed data resources

Summary

Microsoft wishes to work with the Web, Grid and HPC communities:

- to utilize open standards and develop interoperable high-level services, work flows, tools and data services
- to accelerate progress in a small number of societally important scientific applications
- to assist in the development of interoperable repositories and new models of scholarly publishing
- to explore radical new directions in computing and ways and applications to exploit on-chip parallelism

Acknowledgements

With special thanks to Malcolm Atkinson, Neil Chuehong, Geoffrey Fox, Jim Gray, Marty Humphrey, Steven Newhouse, Stuart Ozer, Savas Parastatidis, Norman Paton and Paul Watson