Section 1 – introduction to grid computing

- Introduction :
  - Who am I
  - Life as a physicist
  - The "3rd way" of science and e-Science
- Distributed Computing and e-Infrastructure
  - Trends in high-performance computing
  - E-Infrastructure and grids
- Why grids ?
- Virtual Organisations and Grid Middleware
- Overview of gLite

meraka INSTITUTE
Bruce Becker, for SAGrid JRU and Ops | BBecker@csir.co.za | www.sagrid.ac.za

Section 2 – overview of gLite middleware

- General overview
- Security System
- Information Service
- Workload management service
- Data management on the grid

Section 3 : The grid in (South) Africa

- Outline of SAGrid and integration with global grid activities
- What is available ?
- How to get started ?
- How to run my application ?
- How to get help ?

Bruce Becker : ex physicist,
Coordinator, South African National Compute Grid.

UCT CERN

A Large Ion Collider Experiment

European Organisation for Nuclear Research

4

Life as a physicist

- As research questions have become more and more fundamental, we have had to build ever-more complex tools to study them
- This observation holds both for the physical tools
    - Experiment detectors
    - Accelerators
    - Acquisition systems
    - Data repositories
- As well as the theoretical tools (which I won't pretend to know about)
- All of this involves an impressive investment in terms of time, money and intellect...
- Due to the nature of the topics probed by frontier physics, many new methodologies and tools have been first tested and explored by them :
    - e.g. internet, cryogenics, data processing
- If you become a physicist, you will inevitably need to deal with computing...

The “3rd Way” of science

- Science has traditionally been conducted in essentially two ways :
    - Observation : do an experiment and publish the result, in a carefully controlled way, to make some statement about nature.
    - Theory : construct a logical and self-consistent description or model of some phenomenon and present it for falsification by experiment (or logic...)
- However, we have lived through a great revolution : we live in the information age !
- The rapid advances in data processing and capacity have given us a way to study previously intractable problems :
    - No mathematical solution
    - No explicit theory
    - No way to study global data sets
    - Limited by small samples
    - Etc....
- A third way of science has emerged over the last 20-30 years which allows us to make discoveries and generate knowledge **in the computer, via :**
    - Monte-Carlo simulation
    - Massive data set analysis and model/parameter scanning
    - etc...

We're not alone anymore !

It is practically impossible nowadays to perform scientific research without computers

- almost all scientists are facing increasingly complicated problems:
  - Simulation of molecular drug candidates and their interactions with specific proteins
  - High energy physics experiments : produce about 10 PB of data per year, and thousand of physicist will want to analyze this data
  - Earth science : keeps track of ozone level with satellite observations, producing about 100 GB of images per day - O(1PB/year) – which has to be available to public and government bodies worldwide
- As a matter of fact, facing such problems, scientists simply “hit the wall” : It became very difficult, expensive and sometimes simply impossible achieve their goals using traditional HPC paradigms
- So scientists started dreaming... What could we do if only :
  - **We had infinite storage space**
  - **We had infinite computing power**
  - **If these resources were always available**
  - **If we could share them safely and securely with our collaborators and others ?**
- **What needs to be done to achieve this ?**

Trends in computing

The price for the same amount of computing power has dropped consistently over the last 20 years

	1991	1998	2005
Sample System	**Cray Y-MP C916**	**Sun HPC10000**	**Shuttle@ NewEgg.com**
Performance	**~10 GFlops**	**~10 GFlops**	**~10 GFlops**
Price	**$40,000,000**	**$1,000,000 (40x drop)**	**$4,000 (250x drop)**
Customers	**Government Labs**	**Large Enterprises**	**Every Engineer and Scientist**

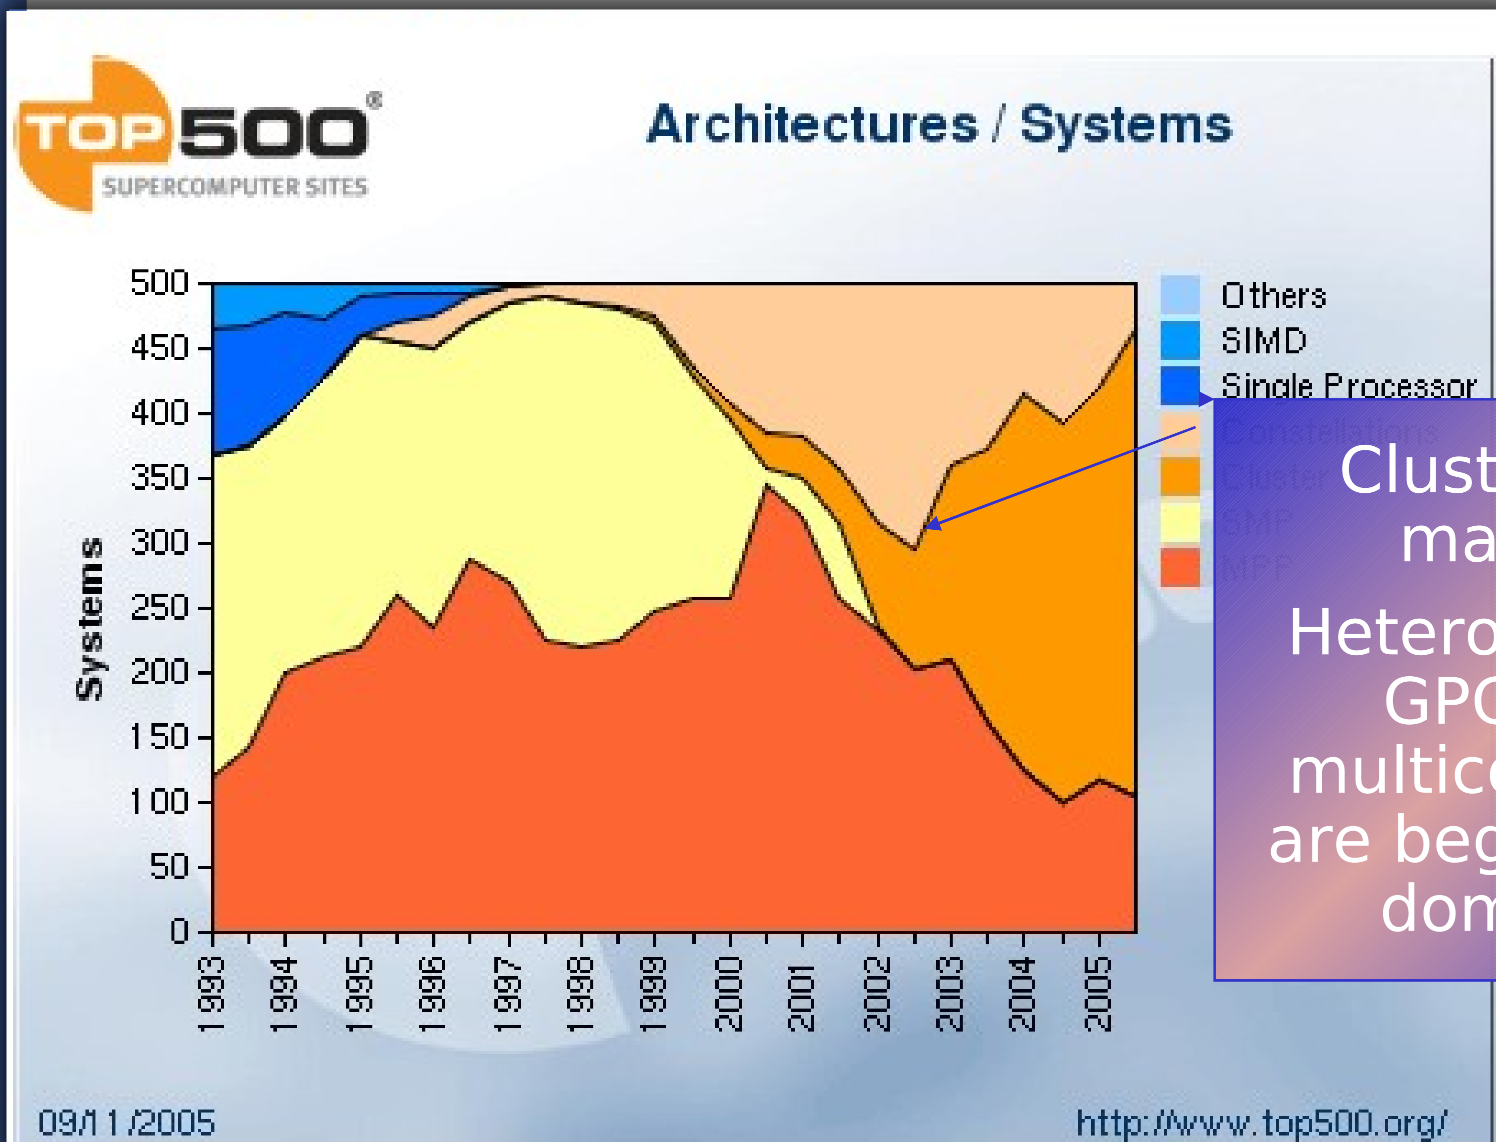

Clusters are majority

Heterogeneous GPGPU + multicore CPUs are beginning to dominate

Computing clusters have "commoditized" supercomputing

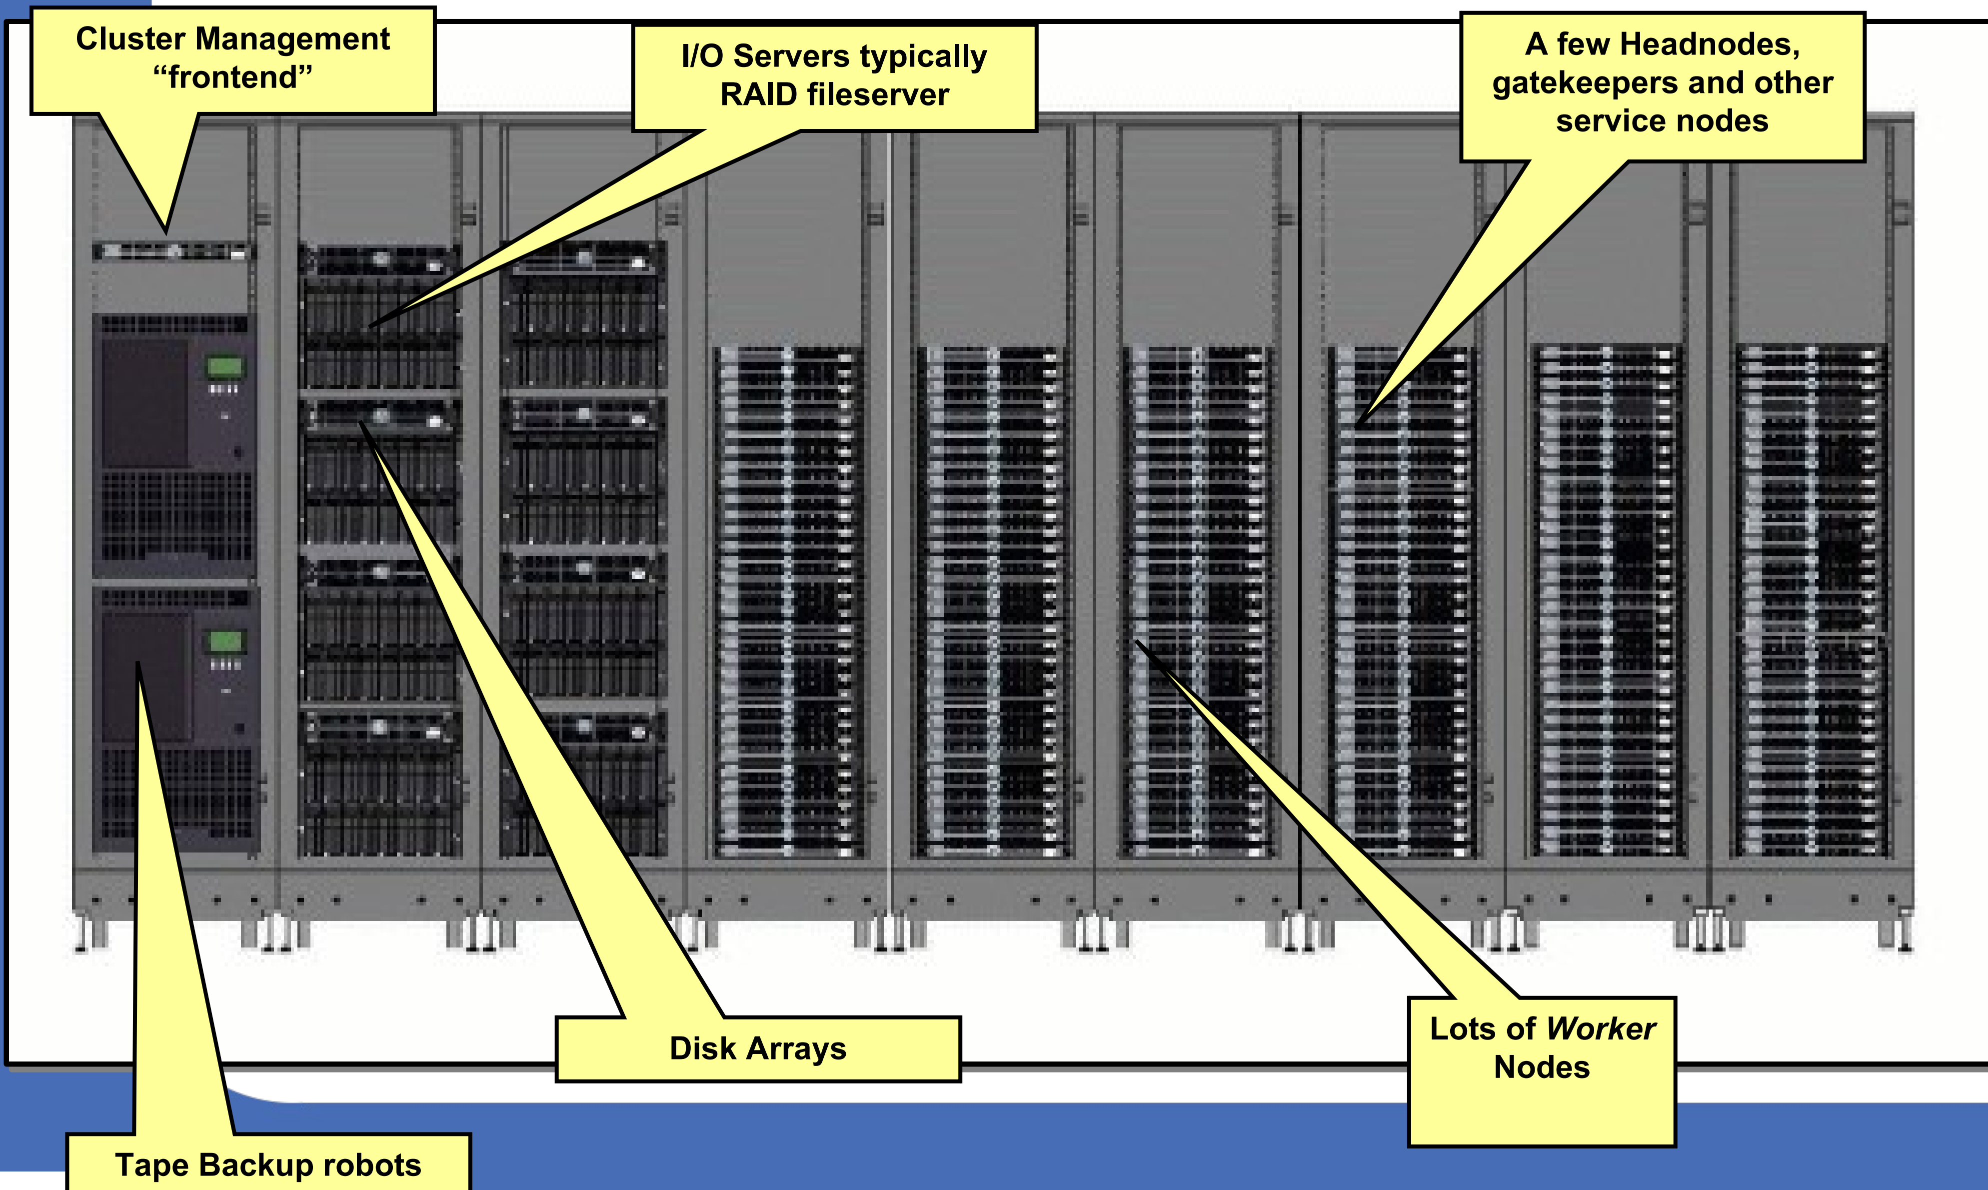

“The network is the computer”

Improvement of Network Technologies

- Network vs. Computer perfomances at constant cost:
  - Computer speed doubles every 18 months
  - Data storage density doubles every 12 months
  - **Network speed doubles every 9 months**
- 1986 → 2000
  - Computers: x500
  - **Networks: x340,000**
- 2001 → 2010
  - Computers: x60
  - **Networks: x4,000**

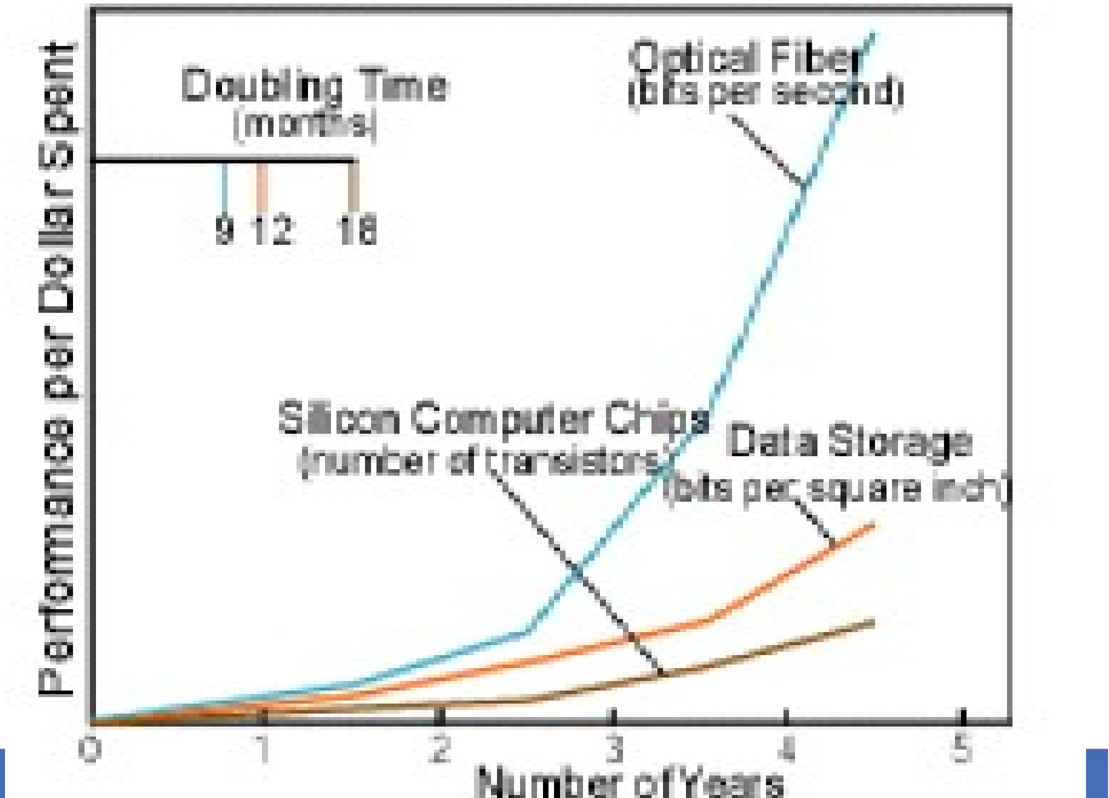

Scientific American (Jan 2001) by Cleo Vilett; source: VinedKhoslan, Kleiner, Caufieldand Perkins

The “global” network coverage means high-speed access... to computing resources, data, instruments – and other researchers !

See http://www.feast-project.org/ ; soon AfricaConnect project (Dante)

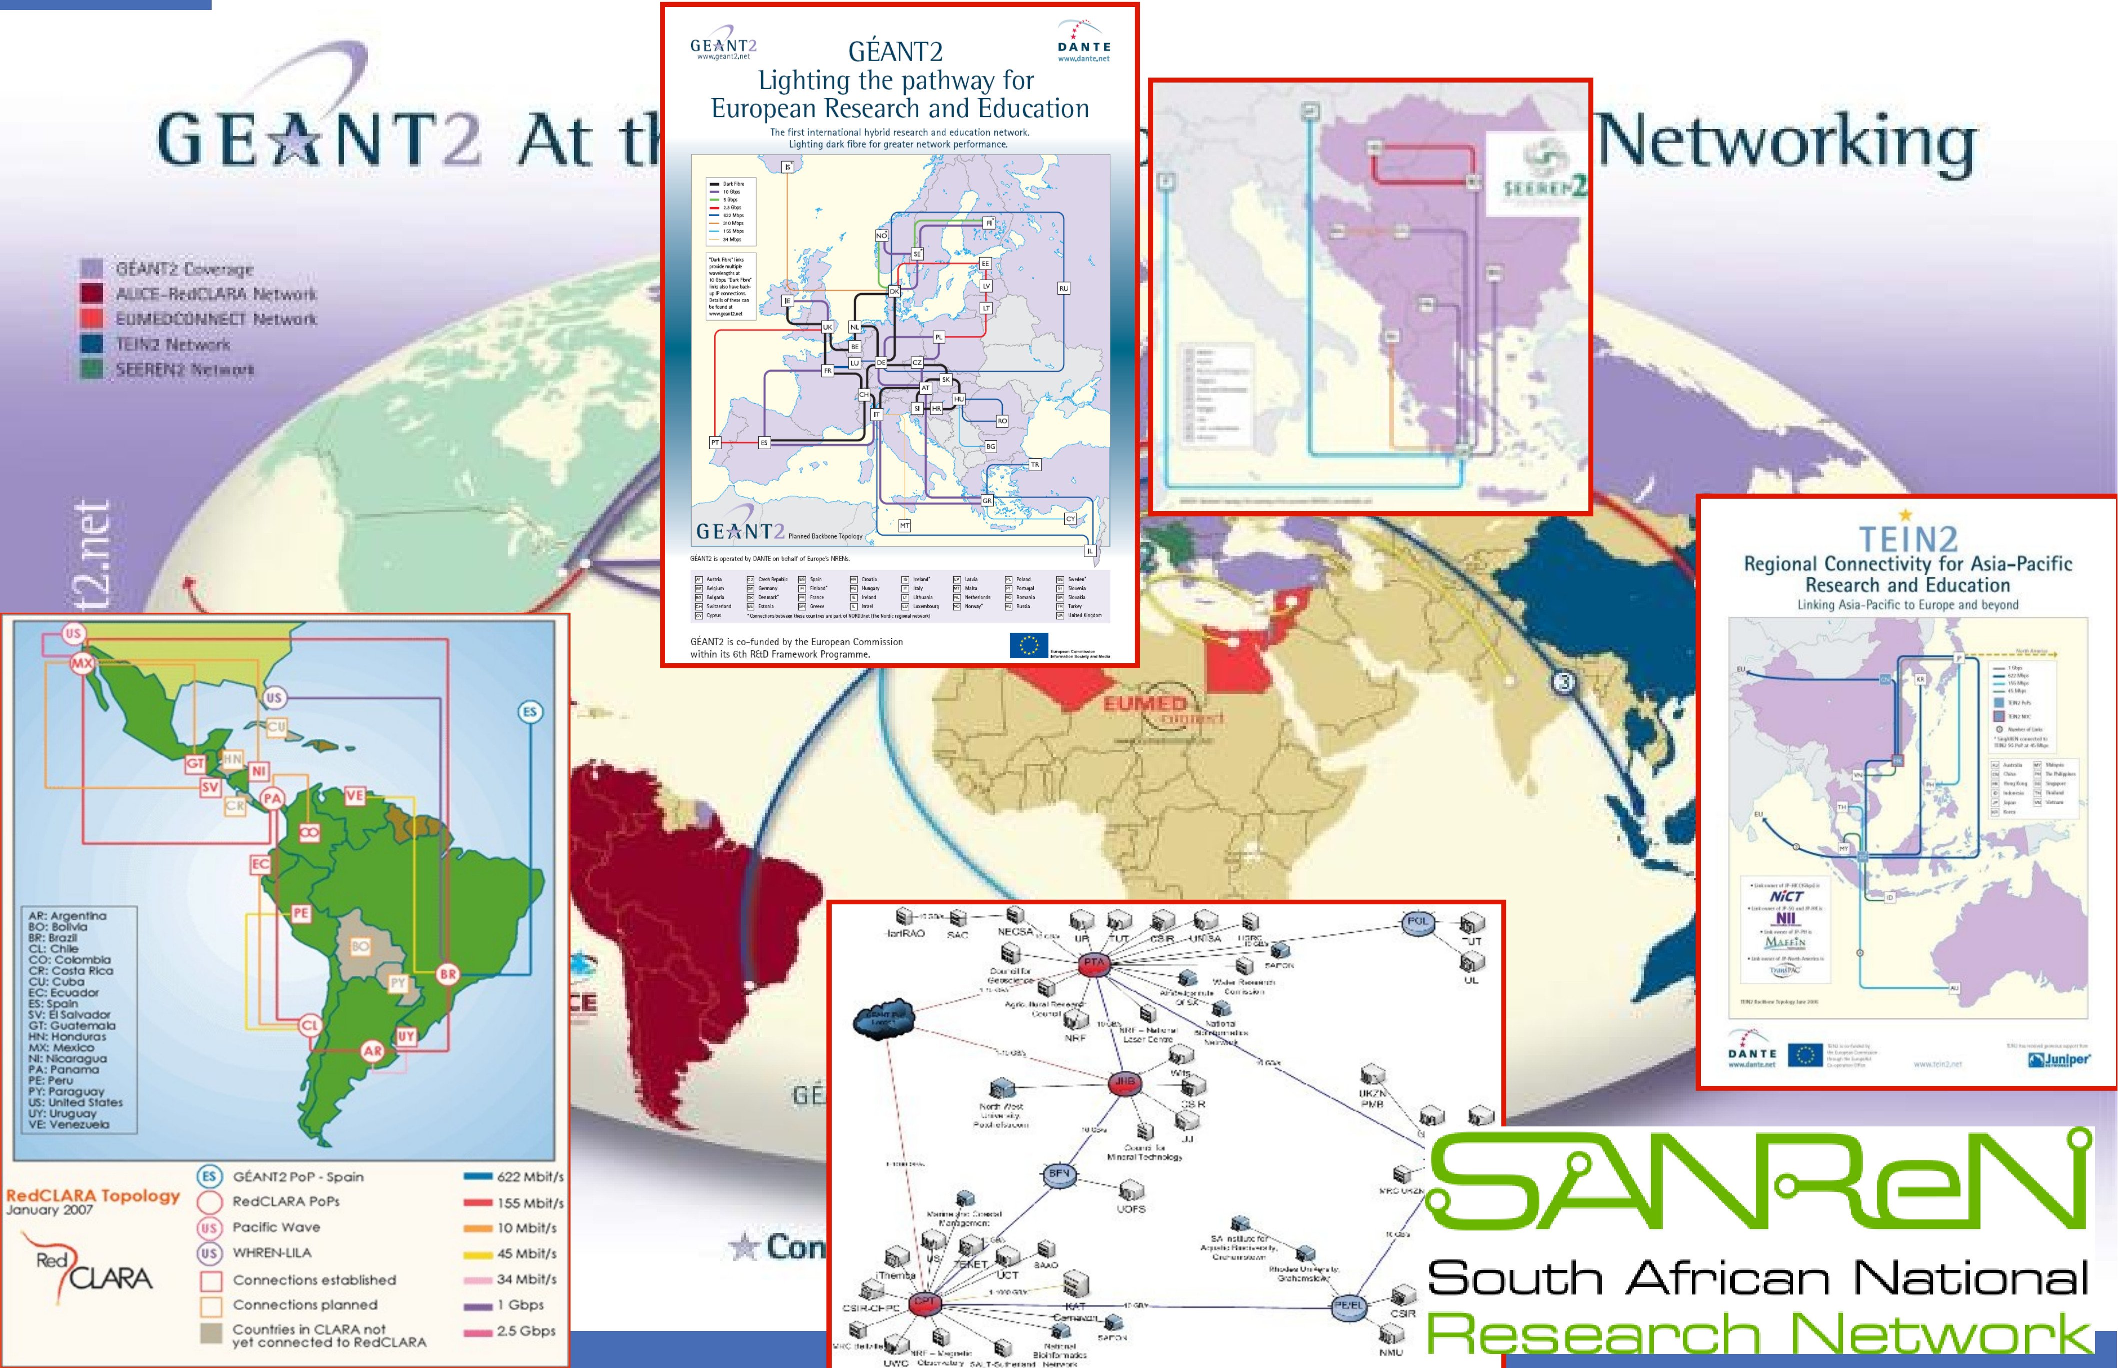

Evolution of Computing and Communication Technologies: 1960-2010

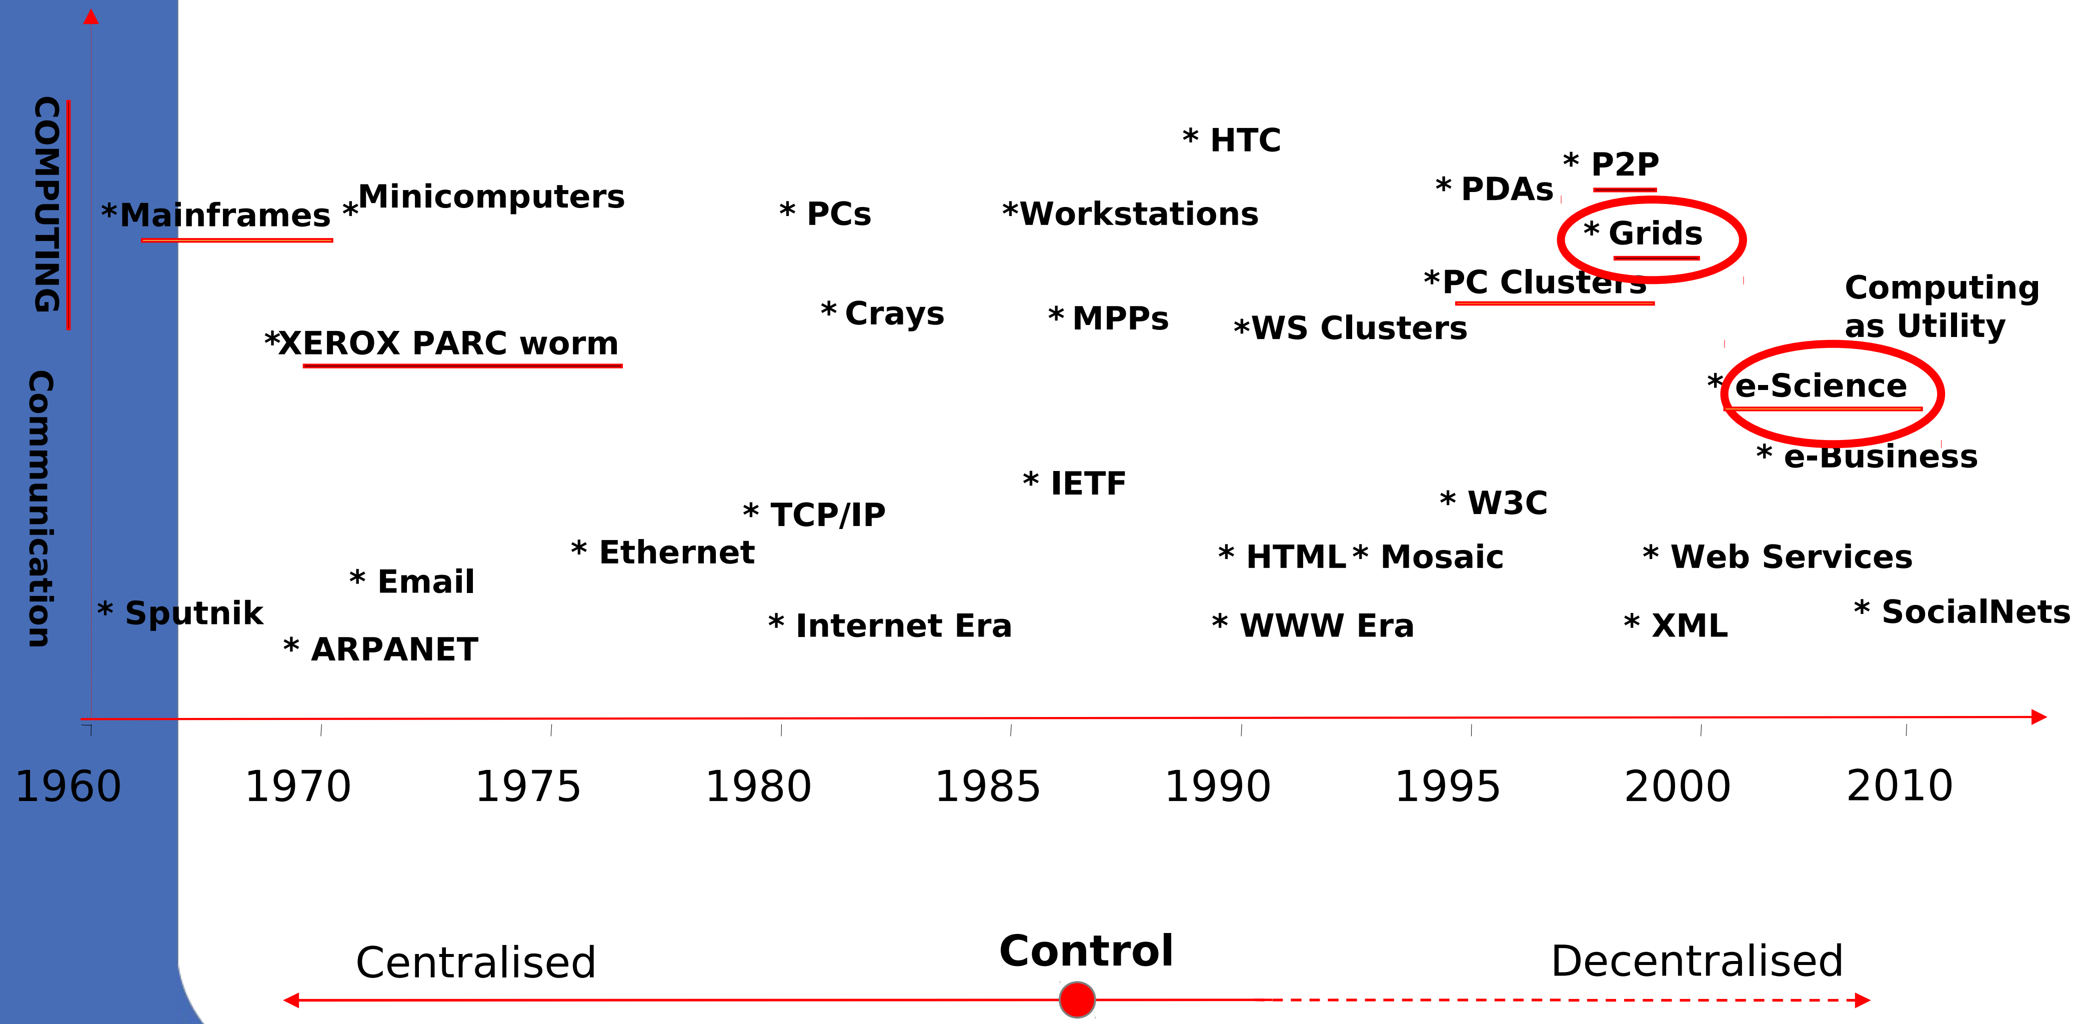

The current e-Science paradigm :

Researchers are close to their dream :

- Computing capacity is almost unlimited : commodity systems are cheap and easily networked.
- Data storage capacity is growing (although data repositories are the biggest problem facing today's infrastructure)
- The world has become very small and your access to these resources depends not on where you are, but who you are :
  - Do you belong to a group or experiment ?
  - Do you have some role to play (software developer, analysis manager, monte-carlo data manager, etc) ?
- Sociological issues come into play...
  - Distinction between
    - those who maintain, deploy and administer the resources, and
    - those who use, extend and rely on the resources
  - Applications have have become so complex that they are largely used as instances of generic frameworks to do specific things :
    - Provided almost **as a service**

A scenario : this scientist walks into a bar...

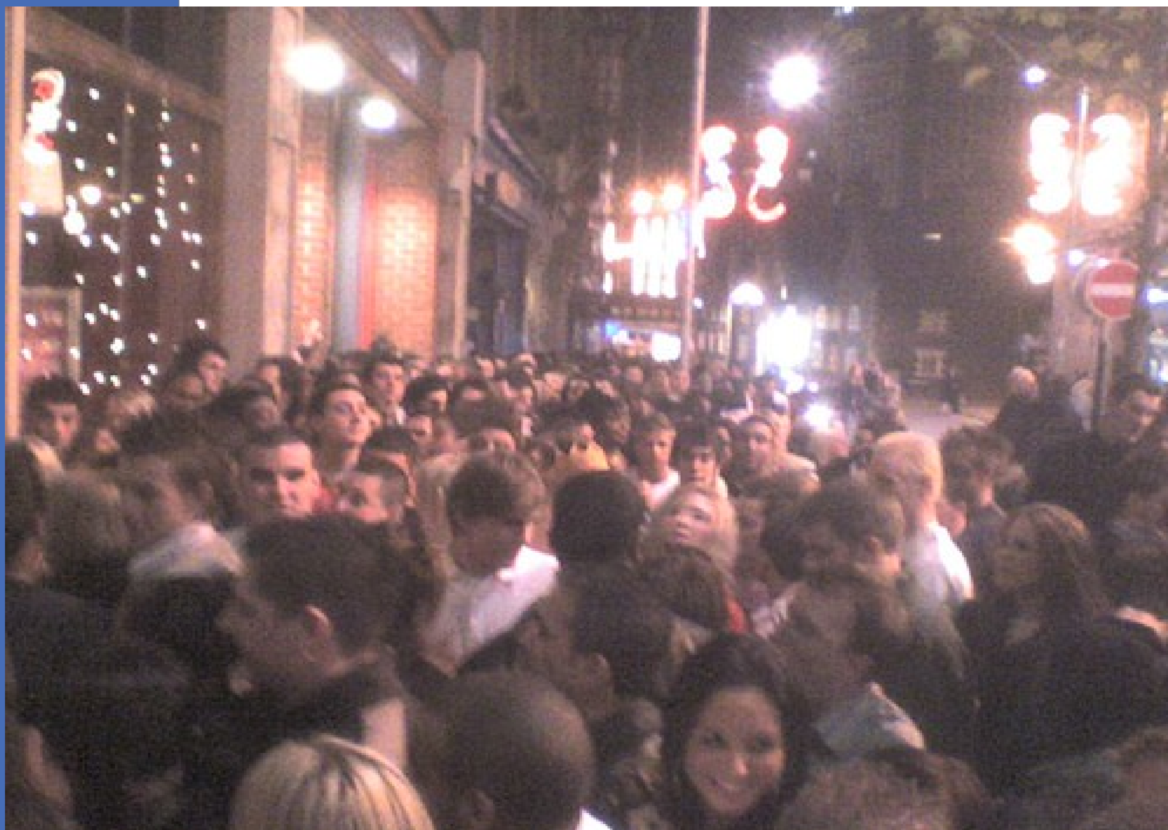

meraka INSTITUTE

What users want – best of all worlds for no extra effort

Some users want highly varied, modular application selection

Amazing infrastructure

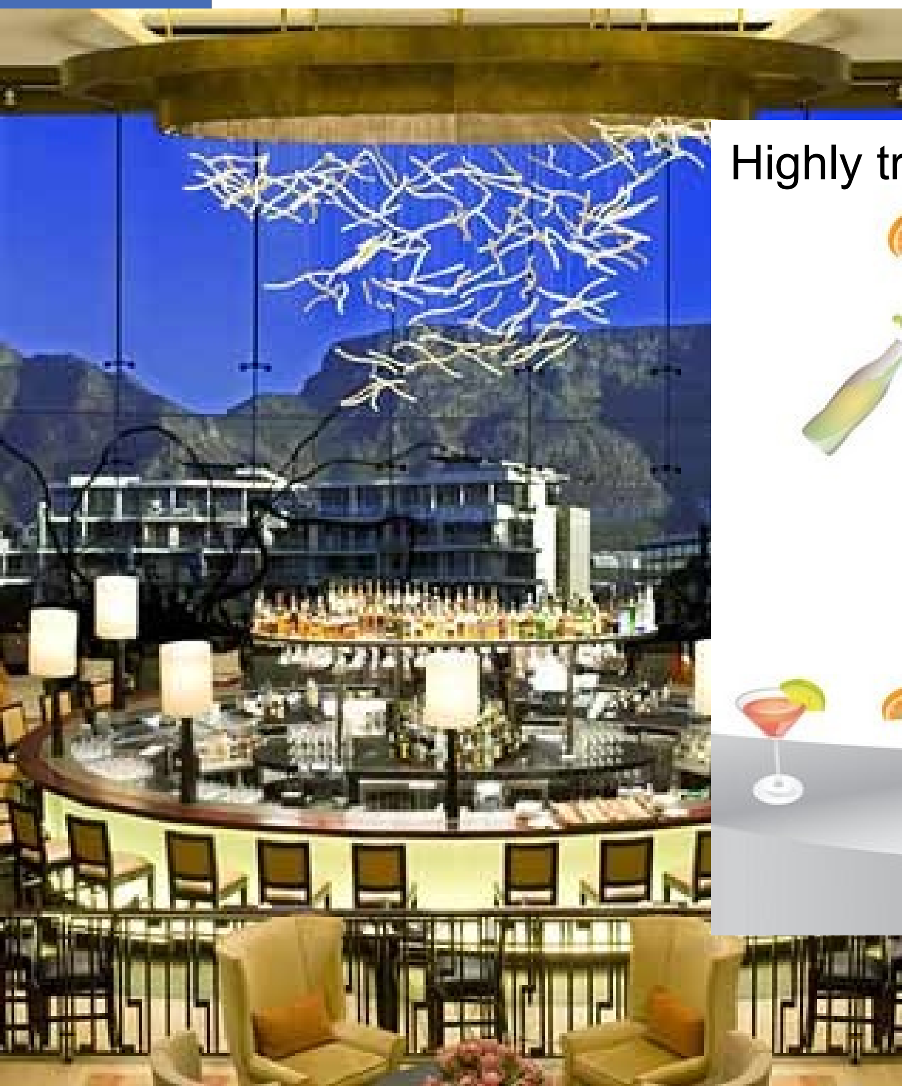

Highly trained support

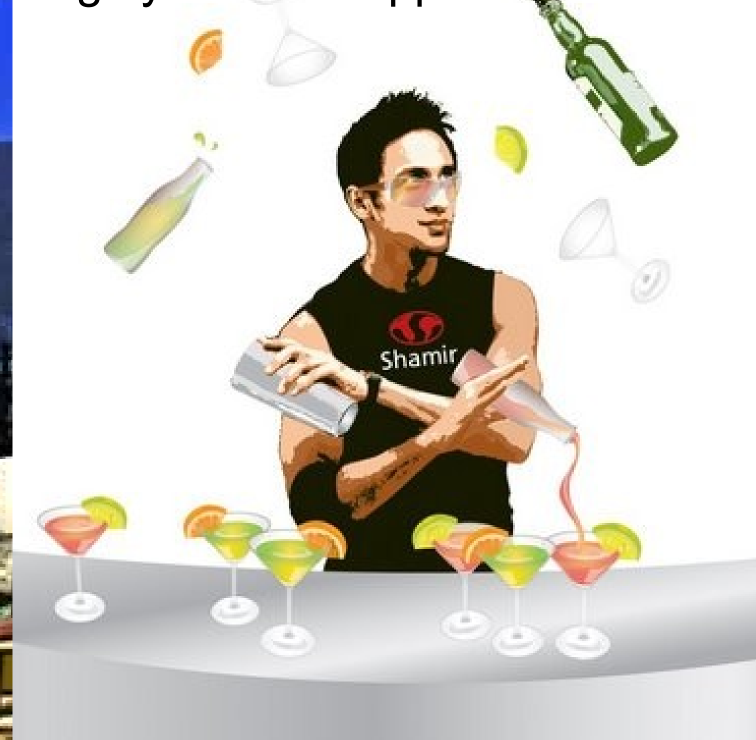

Vertically integrated Highly specialised applications

Best laid plans

we want to avoid...

We have to standardise on something...
not just anything : good, functional, palatable, interoperable
not necessarily perfect...

To talk about infrastructure, we have to standardise on something...

- Technologies
- **Methodologies**

What is e-Science ? The most widely used quote…

'e-Science is about global collaboration in key areas of science, and the next generation of infrastructure that will enable it.'

John Taylor

Director General of Research Councils

Office of Science and Technology

2000

Bruce Becker, for SAGrid JRU and Ops | BBecker@csir.co.za | www.sagrid.ac.za

e-Infrastructure

'e-Science is about global collaboration in key areas of science, and the next generation of infrastructure that will enable it.'

e-Infrastructure = Networks + Grids ..

\+ **Operations, Support, Training…**

\+ **Data centres, archives, instruments…**

- Networks connect resources
- Grids enable flexible use of networked resources: "virtual computing" across administrative domains

Grid concepts

“A computational grid is a hardware and software infrastructure that provides dependable consistent, pervasive and inexpensive access to high-end computational capabilities” , Ian Foster and Carl Kesselman, 1998

Foster's three-point checklist :

1. Coordinate resources that are not subject to a centralized control
2. Using standard, open, general-purpose protocols and interface
3. Deliver non trivial quality of services

Bruce Becker, for SAGrid JRU and Ops | BBecker@csir.co.za | www.sagrid.ac.za

What is a grid, what is not ?

A local cluster system, through PBS or Platform's LSF can't be considered a Grid

– it's centrally managed

The web is not (yet) a Grid : open, standard protocols, no central management, but no guaranteed nontrivial level of service.

Multi site schedulers, managed by Condor or Platform's MultiCluster, although using application-specialized protocols to integrate distributed resources and deliver non-trivial quality of services – They can be considered first-generation Grids

The three point checklist applies very well to the large-scale grids being built by scientific communities :

– EGEE, GridPP, TeraGrid, OSG, NAREGI, etc.

– They are often based on open, general purpose protocols (Globus Toolkit and others) to manage sharing and address services as security, reliability and performance.

Why Grids? - Resource Sharing across organisations

- Resource= computer, data, instrument, storage, ...
  - Share investment
    - many small computers = one high throughput architecture
  - Meet occasional peak demands with large number of resources
  - Enable access to specific significant resources
    - E.g. from national investment in data or compute resources
  - Allows communities with small budgets means to access significant resources

Why Grids ? Orchestrate multiple resources

- Enabling whole-system approach to complex problems

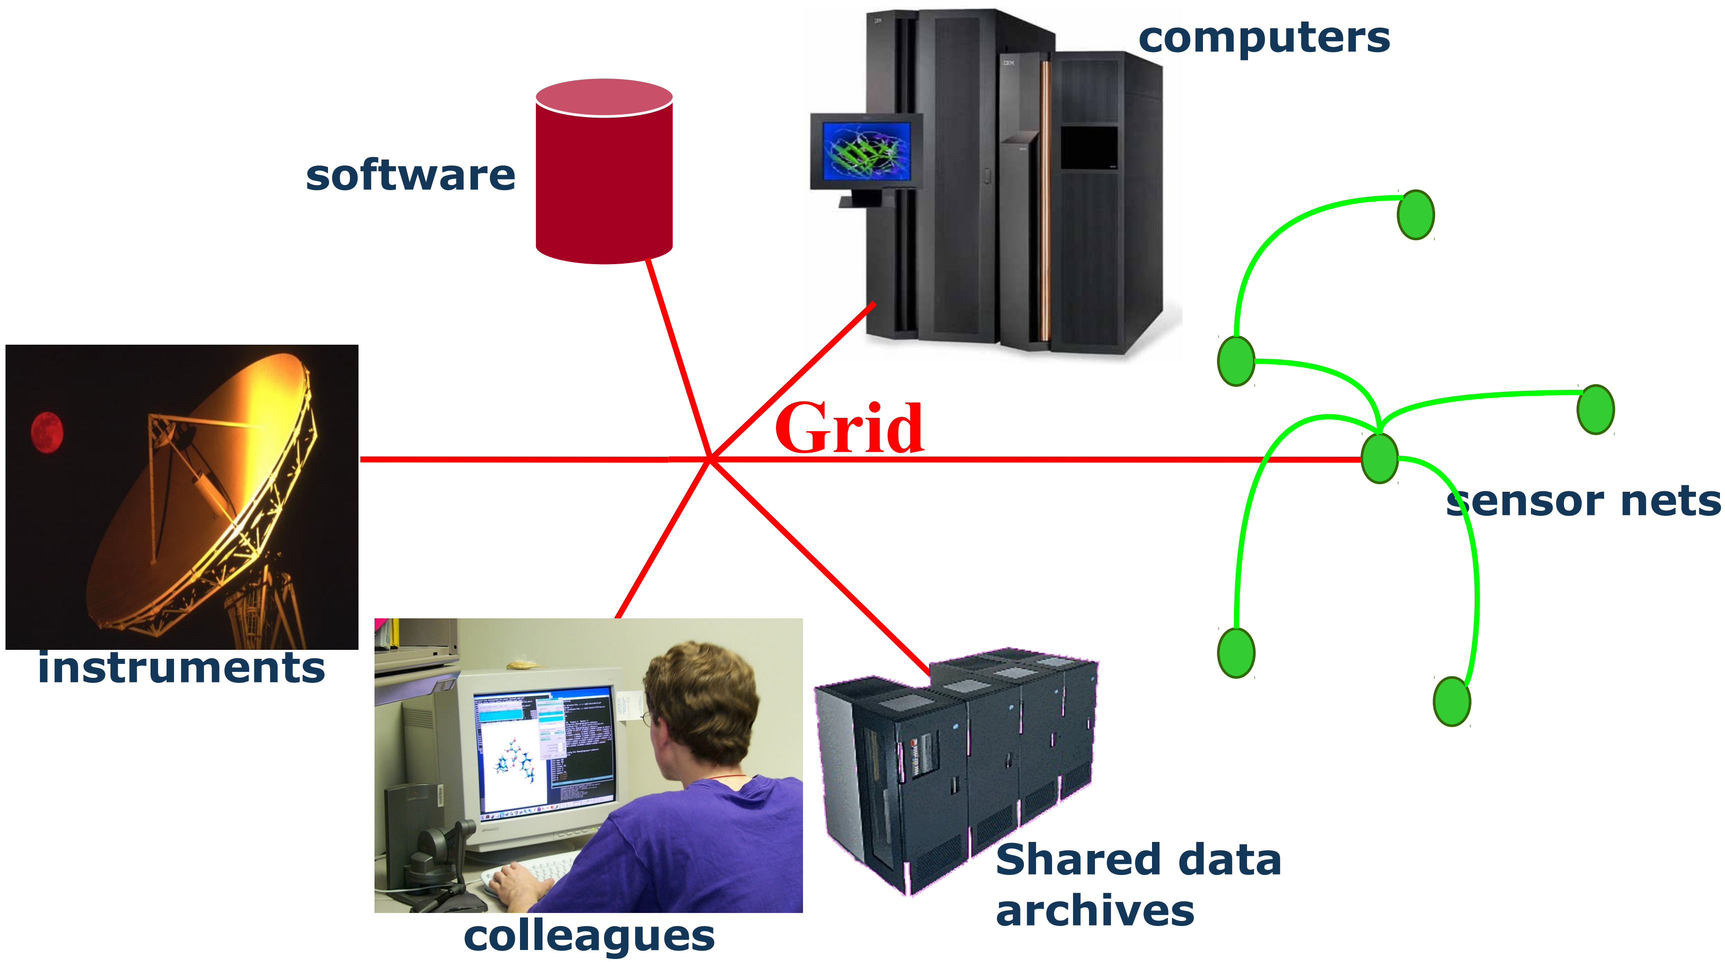

Why Grids ?
3 – Enable effective collaboration

3. Enable more effective collaboration across institutes

Collaborative "virtual computing"

Sharing data, computers, software
Enabled by Grids

Improvised cooperation

Email
File exchange
ssh access to run programs
Enabled by networks:
national, regional and International

People with shared goals

Why Grids?
4 : Move towards “service orientation” in research

Grid becomes sustained infrastructure

- As rail and roads stimulated orchestration of services for commerce, so….
- Networks and grids stimulate information-based commerce and research

- Negotiating access is faster than buying new machine!
  - E.g. for bird-flu investigations with EGEE infrastructure
- But note – need maintain balance between contributing resources to grid and their use

Effect

- Researchers do research
- Systems people run the systems
- Quicker startup to new research

Bruce Becker, for SAGrid JRU and Ops | BBecker@csir.co.za | www.sagrid.ac.za

The five big ideas behind grid computing

- Resource sharing: This is the very essence of the grid
- Secure access : There must be an high level of trust between resource providers and use, which often will never know who each other are
- Resource use : Access to resource needs to be balanced, such that computers are used more efficiently
- The death of distance : Distance should no longer matter; a calculation on the other side of the globe instead of next door should not result in any significant reduction of speed or efficiency
- Open standards : the only way to assure everyone can use the grid and contribute to its development.

Virtual organisations and grids

- What is a Virtual Organisation?
  - People in different organisations seeking to cooperate and share resources across their organisational boundaries
  - E.g. A research collaboration
- Each grid is an infrastructure enabling one or more "virtual organisations" to share and access resources
- Each resource is exposed to the grid through an abstraction that masks heterogeneity, e.g.
  - Multiple diverse computational platforms
  - Multiple data resources
- Resources are usually owned by VO. Negotiations lead to VOs sharing resources so each can benefit at times of peak demand

The many scales of grids

Physical resource		Associated e-Infrastructure
International instruments,..		**International grid (EGEE)**
		Regional – e.g. SE Europe EUMedGrid, EUIndiaGrid
National datacentres, HPC, instruments		**National grids (e.g. SAGrid)**
Institutes' data; Condor pools, clusters		**Campus grids**
	Desktop	

Wider collaboration greater resources

Little interoperability across these scales of grids – yet.

Grid Middleware

- When using a PC or workstation you
  - Login with a username and password ("Authentication")
  - Use rights given to you ("Authorisation")
  - Run jobs
  - Manage files: create them, read/write, list directories
- Components are linked by a bus
- Operating system
- One admin. domain

- When using a Grid you
  - Login with digital credentials – single sign-on ("Authentication")
  - Use rights given you ("Authorisation")
  - Run jobs
  - Manage files: create them, read/write, list directories
- Services are linked by the Internet
- Middleware
- Many admin. domains

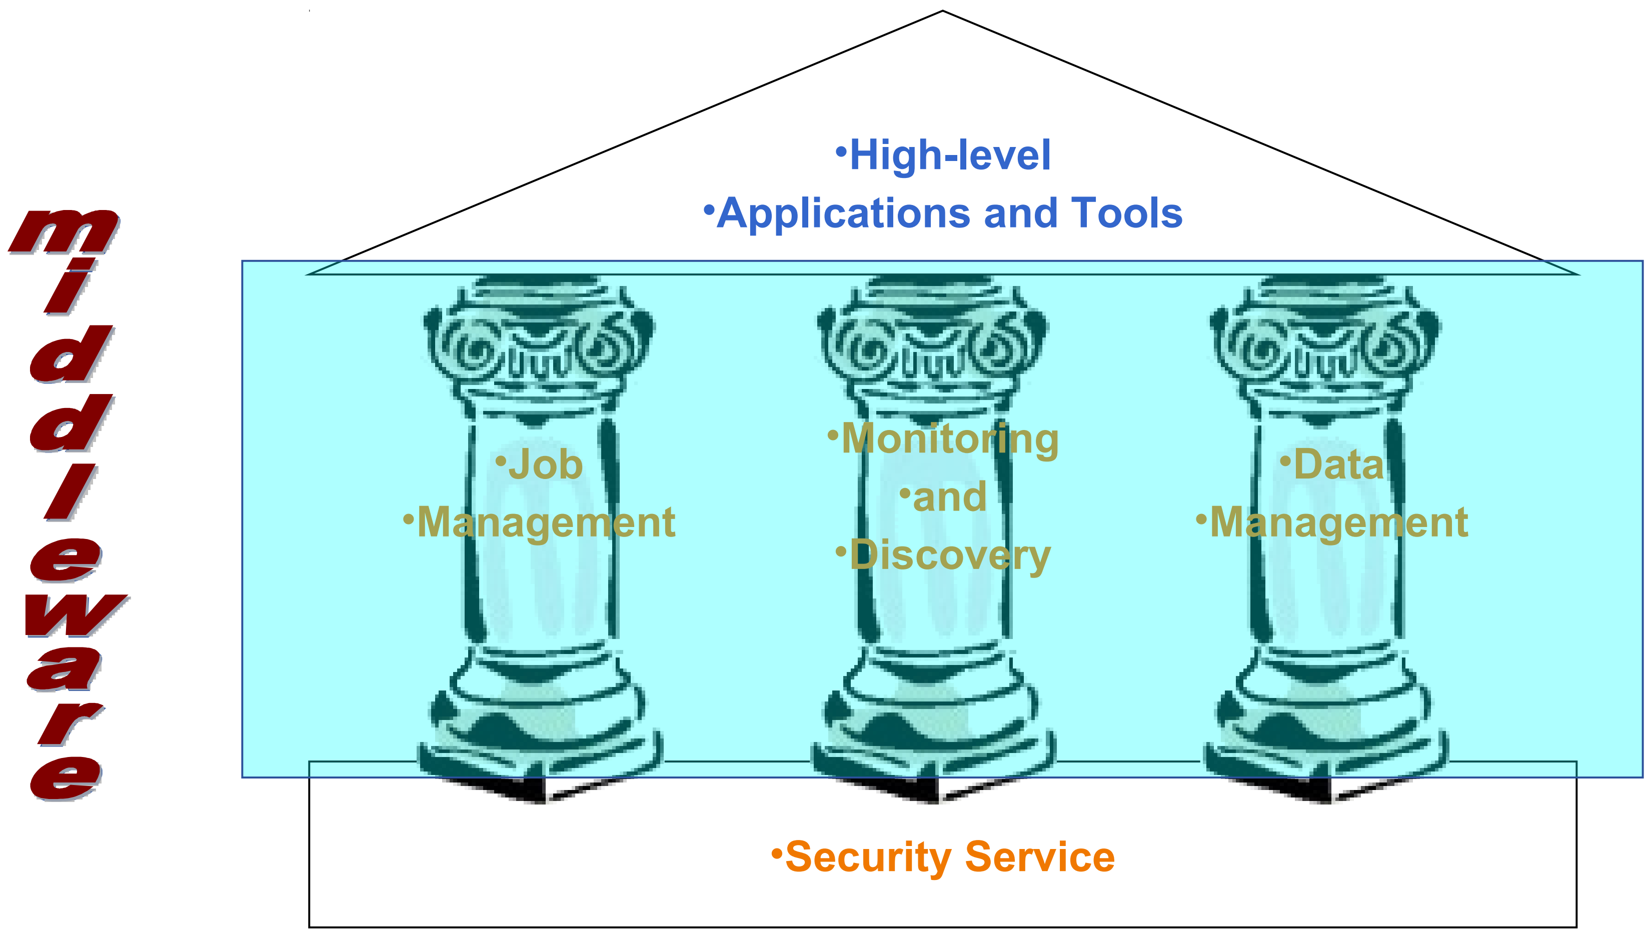
middleware
•High-level
•Applications and Tools
•Job
•Management
•Monitoring
•and
•Discovery
•Data
•Management
•Security Service

Middleware and Software Example – Biomedical applications

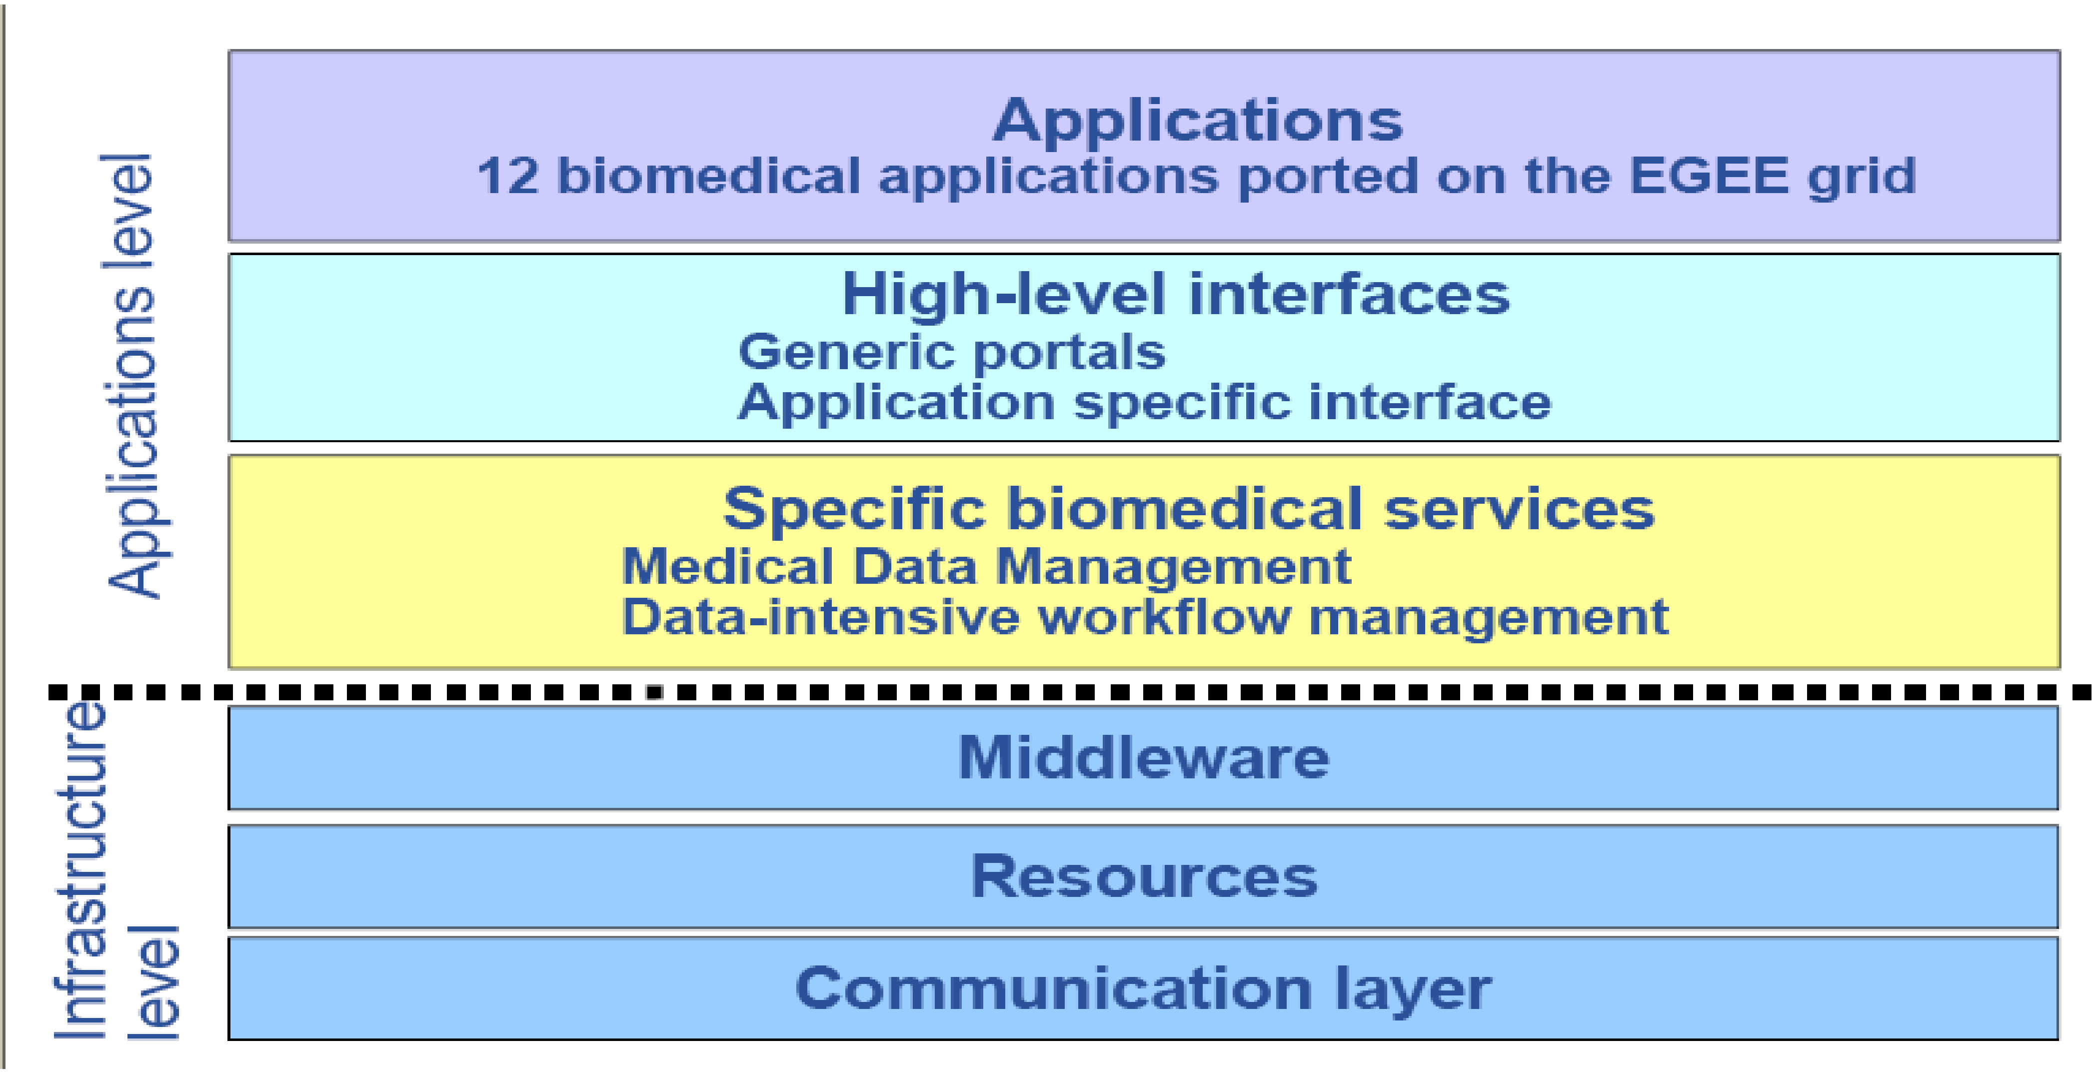

Biomedical community and the Grid, EGEE User Forum, March 1st 2006, I. Magnin

Grid Topology and Services : Make life easy for the users !