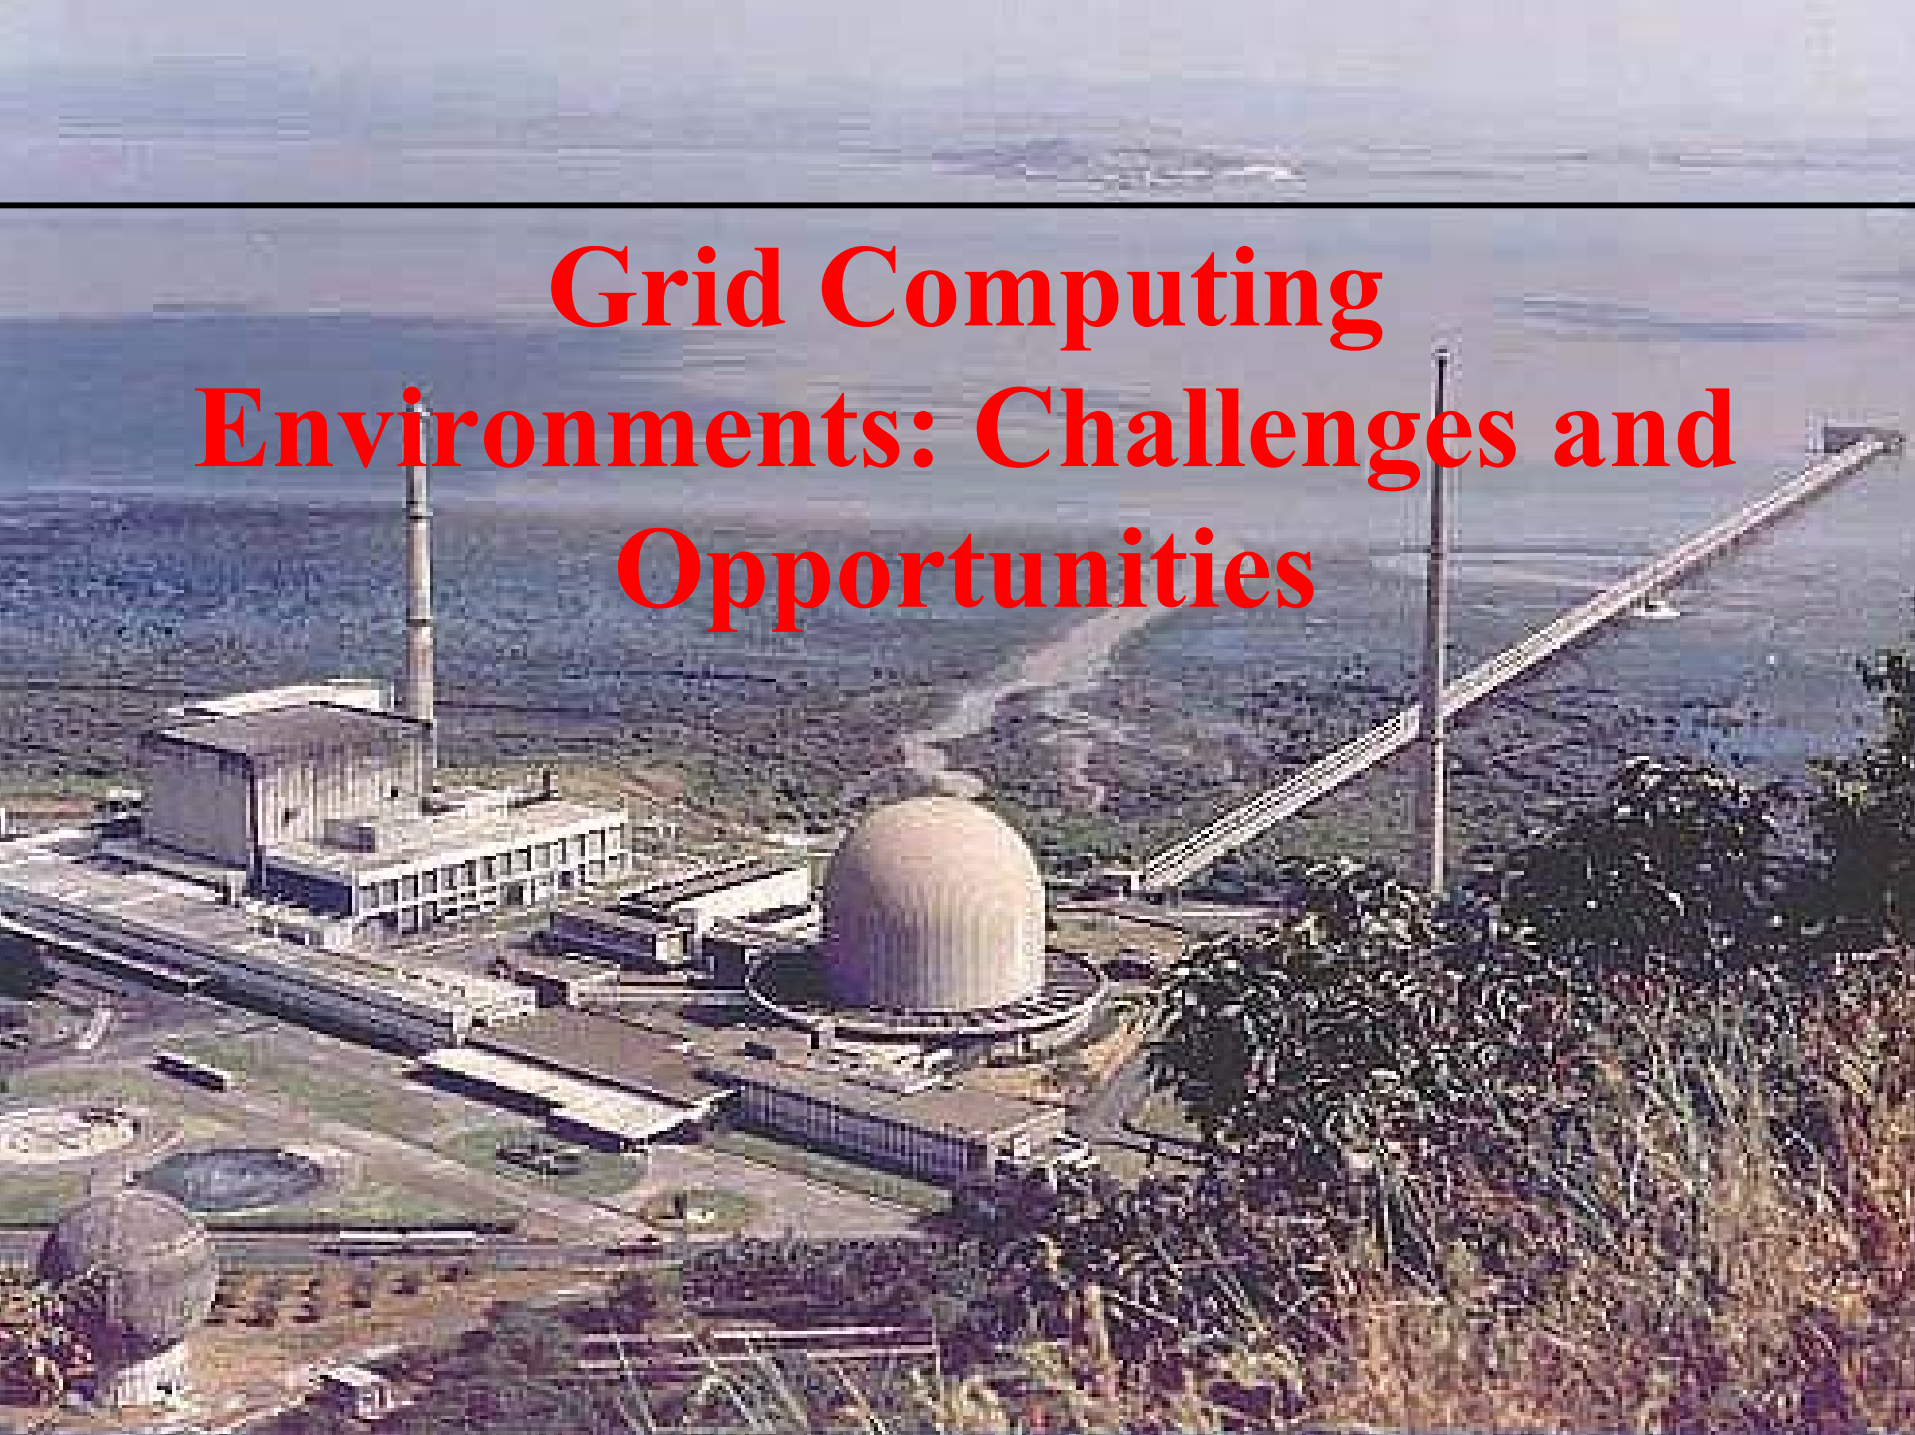

Grid Computing Environments: Challenges and Opportunities

Computing Environments

- Computing & Communication Trends
- Cluster Computing
- High Performance Computing (HPC)
- Peer to Peer Computing (P2P)
- Internet Computing (Web Services)
- Collaboratories
- Grid or Meta Computing
- Conclusions

Computing Requirements

- Today's science is based on computations, data analysis, data visualization & collaborations
- Computer Simulation &Modeling is more cost effective than experimental methods
- Scientific and Engineering problems are becoming more complex & users need more accurate, precise solutions to their problems in shortest possible time
- Data Visualization is also very important

High-end Computing Applications

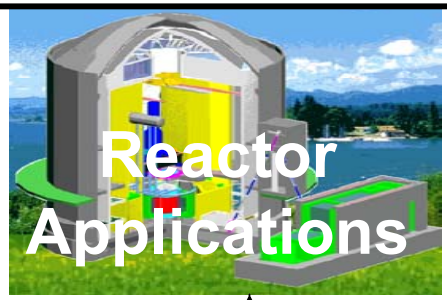

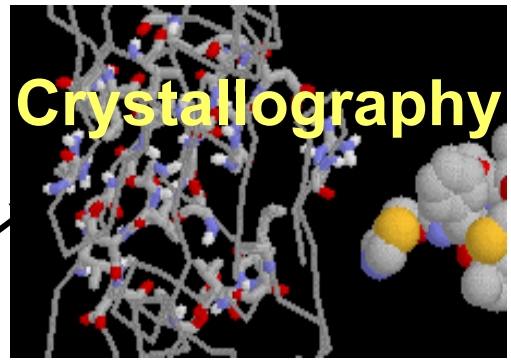

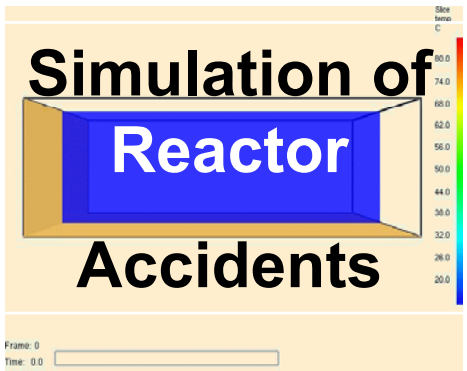

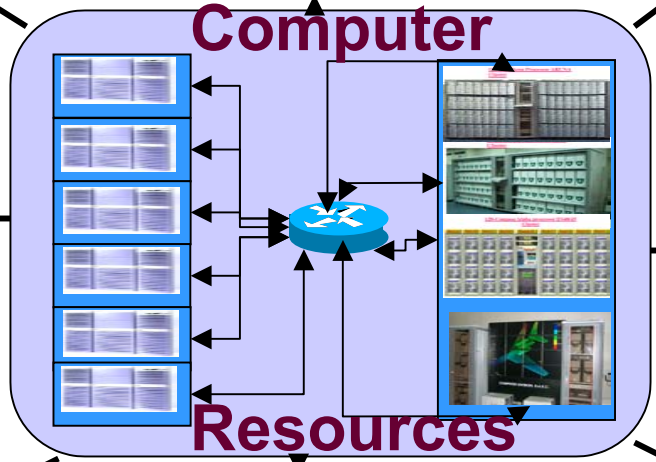

Physics Applications

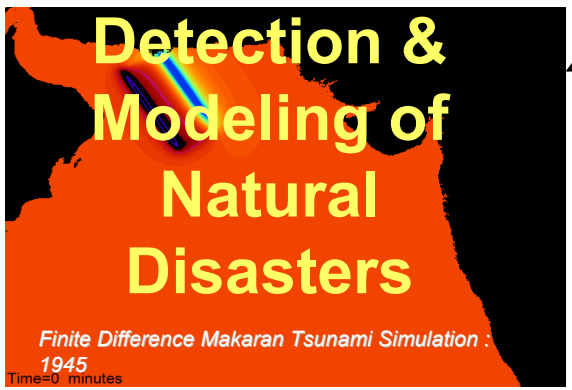

Material Characterization

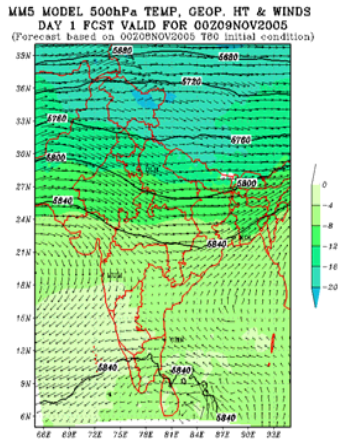

Weather Forecasting

Computational Science Pipeline

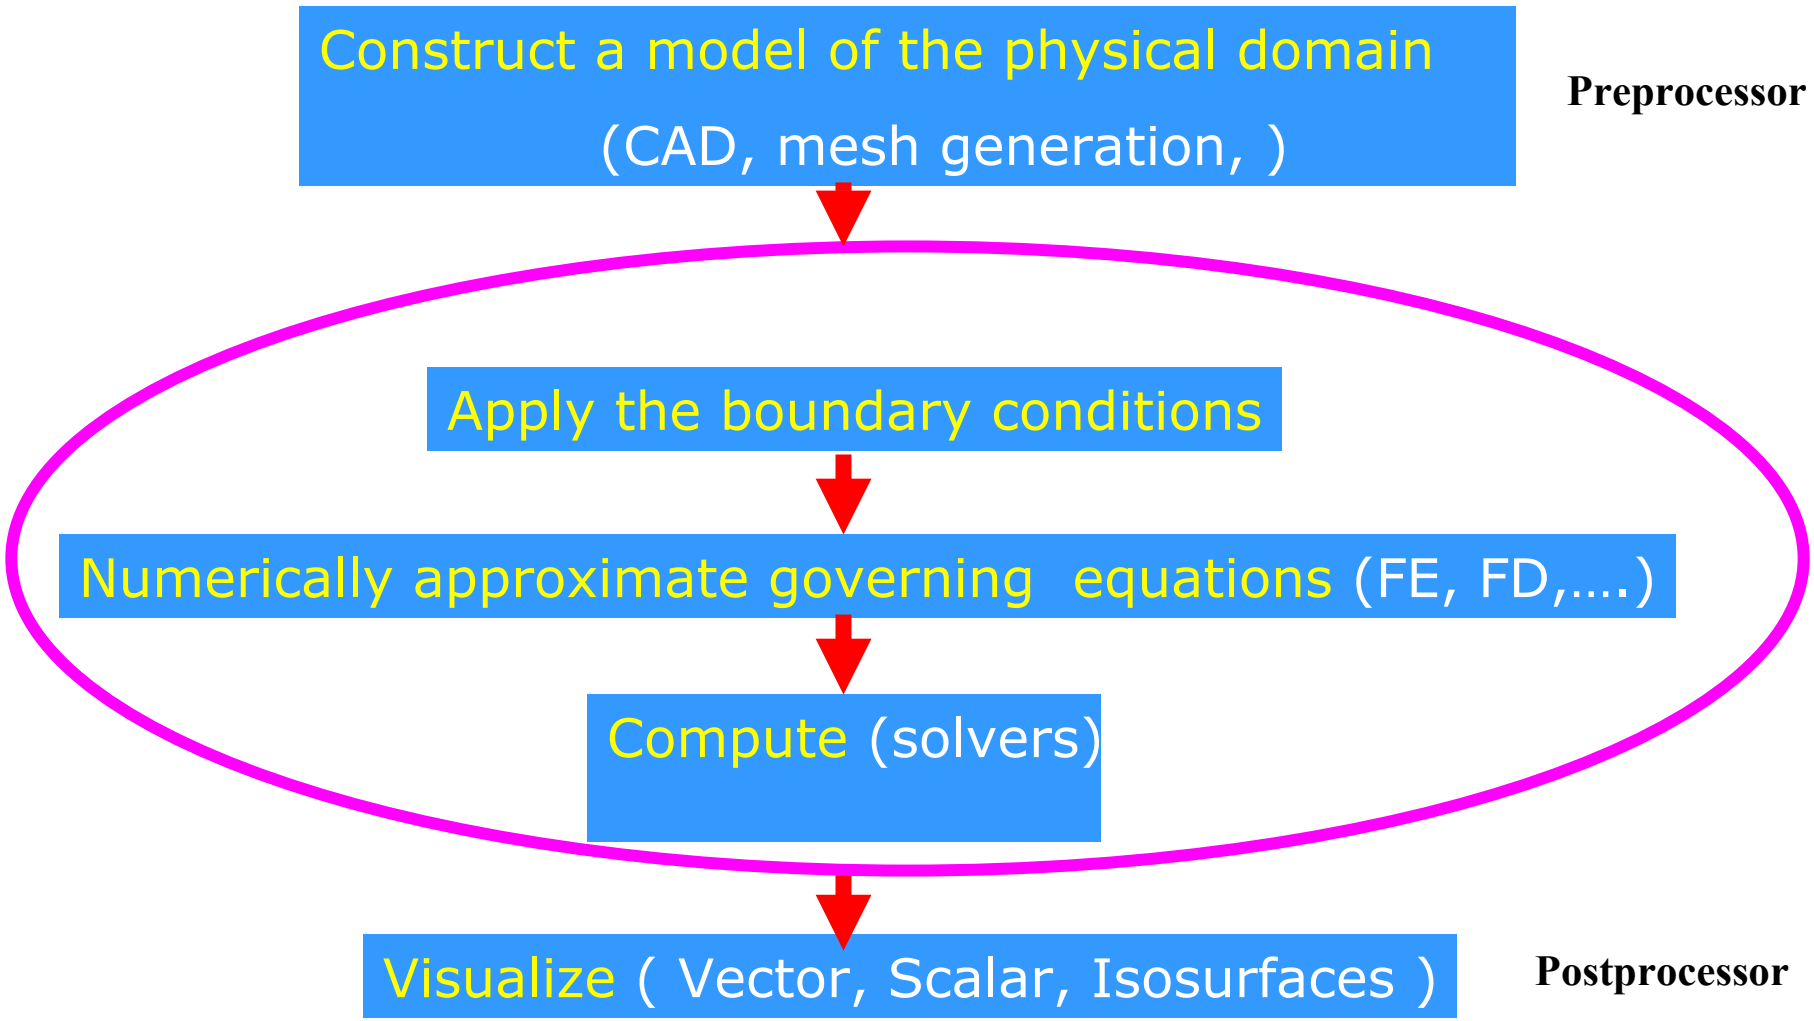

Supercomputer

Preprocessor → **Solver** → **Postprocessor**

Front-end

Vector and Multiprocessor

Cray Supercomputer

SGI Origin 2000 – 3000

Multiple Graphics HW

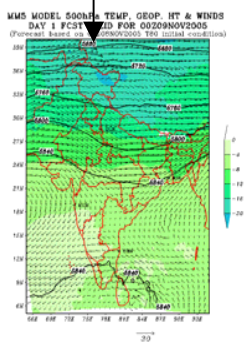

Costs millions of $ and not available to India

Growth of Computing power in BARC

Parallel Cluster Computers

Use locally available commodity hardware

Open source software

Inhouse monitoring management and development tools

Cheaper alternative for supercomputing

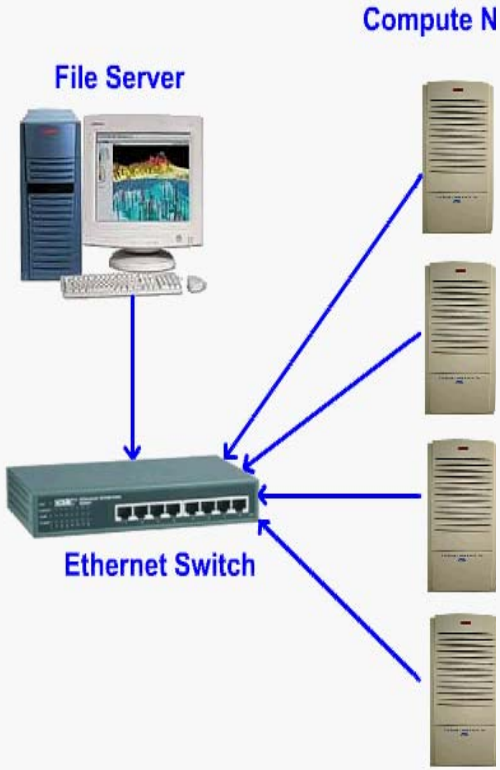

ANUPAM Series of BARC Supercomputers

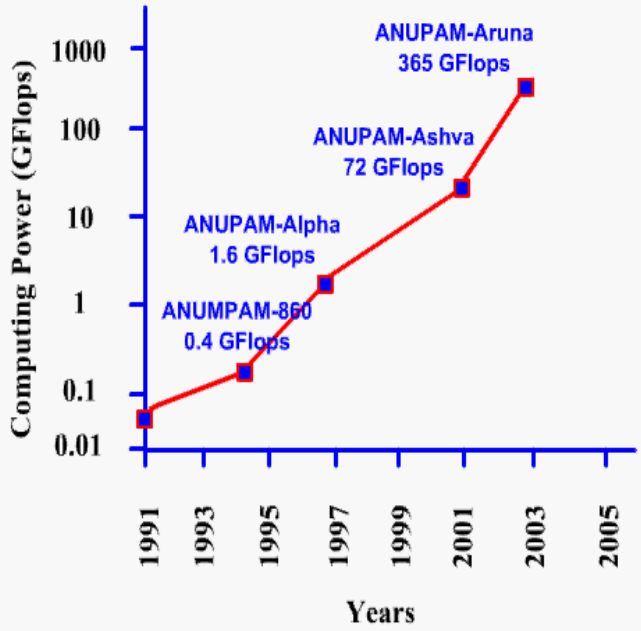

Computing Power increased by order of 1000 from ANUPAM-860 to ANUPAM-ARUNA

ANUPAM-Ashva 72 GFlops

ANUPAM-Aruna 365 GFlops

Supercomputer with Extremely low cost

Parallel Computer?

- Parallel Computing: the use of multiple computers or processors working together on a common task
  - each processor works on its section of the problem
  - processors are allowed to exchange information with other processors
- Two big advantages of parallel computers:
  1. total performance
  2. total memory
- Parallel computers enable us to solve problems that:
  - benefit from, or require, fast solution
  - require large amounts of memory

Performance: Parallel Computers

- Earth Simulator, Japan: 35.86 TFLOPS
- ASCI Q, LANL, USA: 7.22 TFLOPS
- TeraScale Cluster, Virginia Tech University: 11 TFLOPS
- IMSc, Chennai: KABRU 1 TF
- C-DAC Param: 535 GFLOPS
- BARC Anupam: > 2 TFLOPS
- HRI Cluster 480 GFlops

- **All Top 500 Supercomputers nowadays are parallel computers**

Cluster based Visualization Why?

- **Current high end visualization hardware is expensive and not flexible enough**
  - **No commodity building block**
  - **Need redesign in order to keep pace with technology**
  - **Limited resolution, low rendering speed**
- **Solution is Parallel Visualization/Rendering**

20-MegaPixel Tiled Display System

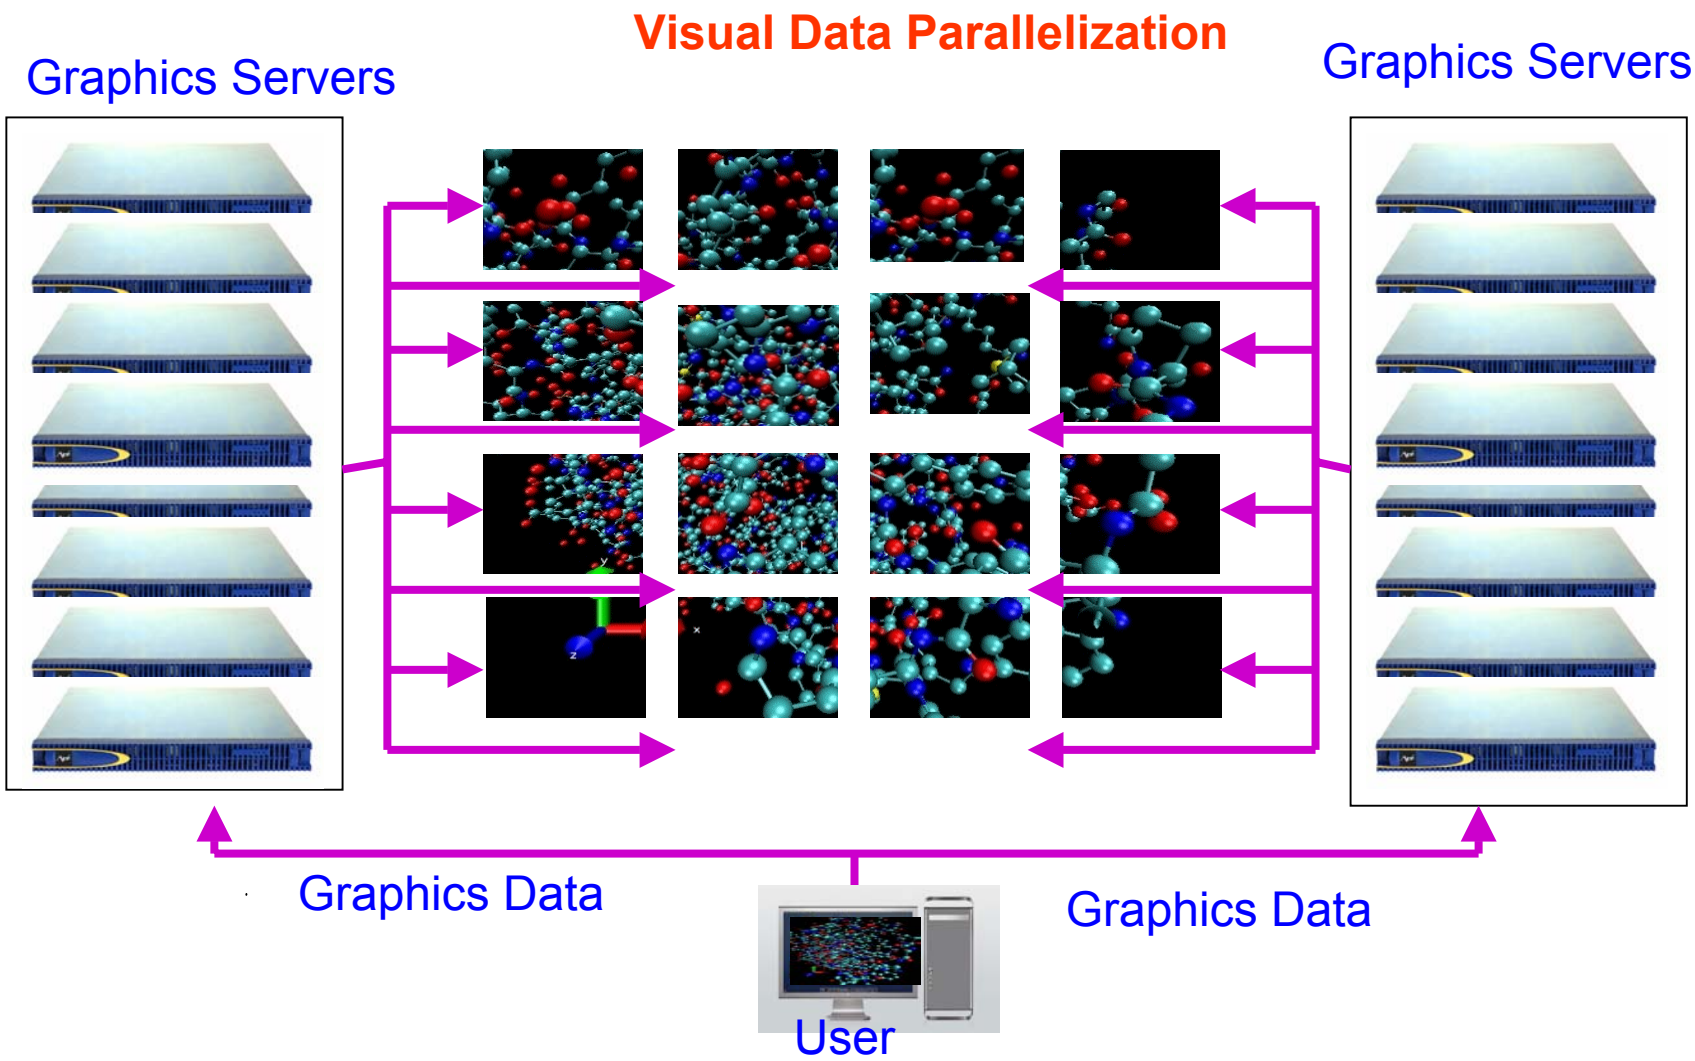

Data Visualization Software

- **RedHat Linux 7.2**
- **Chromium**
- **DMX**

Application Areas

- **Deep & Rich Scientific Visualization**
- **Large Scale 3D Simulations**
- **High Resolution 2D Content**

CollabCAD (NIC 3 D Viewer)

Tiled Image Viewer

Cost less than $ 60, 000

Tiled MPEG/AVI Player

OpenVRML

VTK

Graphical Control Panel for Tiled Display

High Performance Visualization Cluster

First of its kind in the country

Tiled display giving very high resolution (20Mpixel), high-speed rendering needed for scientific visualization

- Tiled 4x4 LCD Panels
- 1 Master Client, 16 Graphics Servers
- 5120 x 4096 total resolution 1280x1024 per LCD 20 Million Pixels
- 16 Times rendering speed

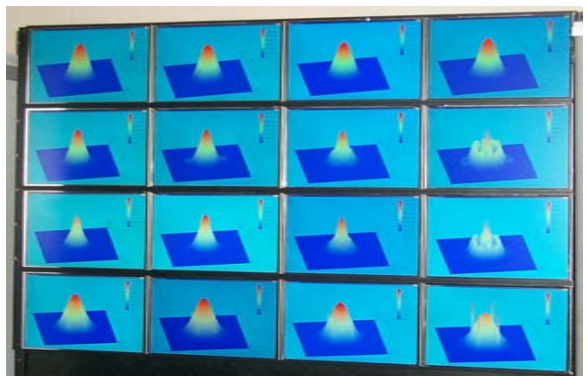

Exploiting parallelism for

- Processing (parallelization of computations)
- I/O (parallel file system for large storage)
- Visualization (parallelized graphic data pipeline/Tile Display Unit)

Has given high compute and high throughput Computational system

Limitations of Parallel Computing

Programming so many nodes concurrently remains a major barrier for most applications

- Source code should be known & parallisable
- Scalable algorithm development is not easy
- User has to worry about message passing, synchronization and scheduling of his job
- Fortunately lot of free MPI codes and even parallel solvers are now available

However all resources are allotted for a single job and idle when there is no parallel job

- 15% users only require these solutions, rest can manage with normal PCs

Computations of future?

- **PC of 2003 is as fast as supercomputer of 1990's**

  Still not adequate for many and users still continue to buy new machines

- **PC now comes with over 80 Gigabyte of disk as much as entire 1990's computer centre (Storage Capacity doubles every 12 months)**

  By 2006 several physics projects like CERN or Astrophysics, Biology, HEP will produce multiple of Petabytes of data per year

- **To work with a colleague even across campus LAN on Petabyte scale data set, requires ultra fast network**

  Even though cpu power, disc storage & communication continues to improve, computing resources are failing to satisfy users demand, they are difficult to use. No where near user expectations : software is the issue !!!

Compute-Communication Cross over

"When the network is as fast as the computer's internal links, the machine disintegrates across the net into a set of special purpose appliances" (George Gilder)

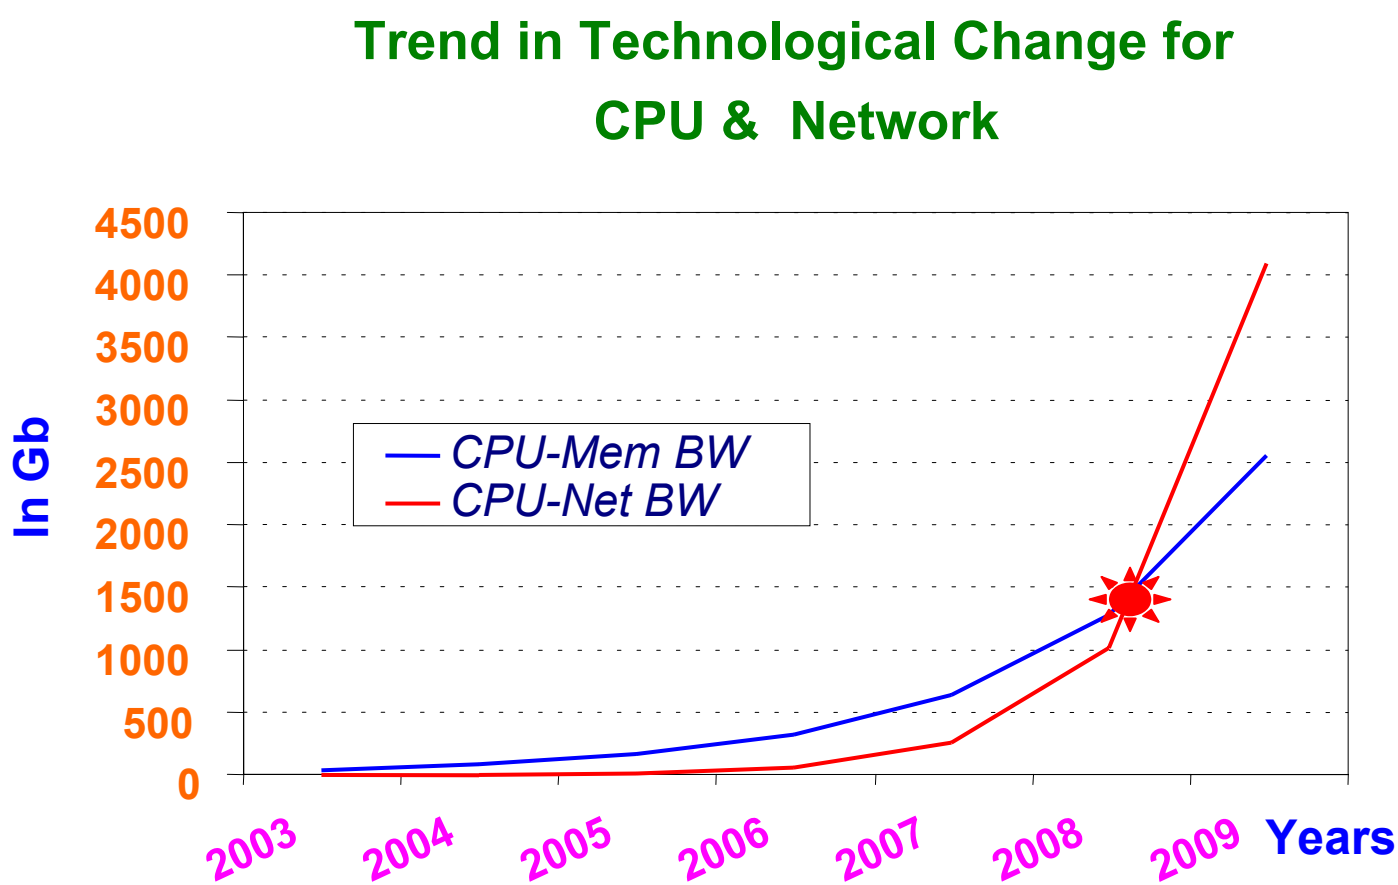

PC Clusters:Multiple PCs

- Over the last few years, computing power of Intel PCs have gone up considerably (from 100 MHz to 3.6 GHz in 8 years) with fast, cheap network & disk (in built)
- PCs are freely available from several vendors
- Emergence of free Linux as a robust, efficient OS with plenty of applications
- Linux clusters (use of multiple PCs) are now rapidly gaining popularity in academic/research institutions because of low cost, high performance and availability of source code

OSS:New Direction in Cost Saving

- The ideas behind "open source" and "free software" are not new
  - Lots of software has been free from the beginning
- What is new is the emergence of business and/or political movements to define, advocate, and institutionalize these notions
- These movements have been galvanized by the success of Linux (or more properly GNU/Linux systems)
  - *Much* of the software in Linux system is GNU
  - Allows speedy yet low cost deployment
- Software evolves slowly but will have to catch up to satisfy user's expectations and OSS is the way

Clusters : Primary IT infrastructure

Clustering is replacing all traditional Computing platforms and can be configured depending on the method and applied areas

- **LB Cluster -** *Network load distribution and LB*
- **HA Cluster -** *Increase the Availability of systems*
- **HPC Cluster (Scientific Cluster) -** *Computation-intensive*
- **Web farms -** *Increase HTTP/SEC*
- **Rendering Cluster –** *Increase Graphics speed*

HPC : High Performance Computing HA : High Availability LB : Load Balancing

Issues in building large clusters

- Scalability of interconnection network
- Scalability of software components
  - Communication Libraries
  - I/O Subsystem
  - Cluster Management Tools
  - Applications
- Auto-Installation and Management Procedures
- Troubleshooting Procedures

Other Issues in operating large clusters

- Space Management
  - Node form factor (Desktop, rack mounted 1U/2U)
  - Layout of the nodes and other issues like noise
  - Cable routing and weight
  - Cooling arrangements
- Power Management
- Centralized Infrastructure Management Software
- Cost of ownership is not really low
- Performance/Price/Power Consumption

The P2P Computing

- Computing based on P2P architecture allows to share distributed resources with each other with or without the support from a server.
- How do you manage under utilized resources?

  - It is seen that utilization of desktop PC is typically <10 %, and this percentage is decreasing even further as PCs are becoming more powerful

  - Large organizations must be having more than thousand PCs, each delivering > 20 MFlops and this power is growing with every passing day ……Trick is to use Cycle Stealing mode

  - Each PC now has about 120Gbyte disc capacity

  120Gb X 1000 = 120 Terabyte storage space is available ; Very large File storage

  – How do you harness power of so many PCs in a large organization? …….. Issue of "Ownership" hurdle , to be resolved

  – Latency & bandwidth of LAN environment is quite adequate for P2P computing……Space management no problem; use PCs wherever they are!!

INTERNET COMPUTING

- Today you can't run your jobs on the Internet
- Internet Computing using idle PC's, is becoming an important computing platform (Seti@home,Napster,Gnutella,Freenet, KaZak)
  - www is now a promising candidate for core component of wide area distributed computing environment.
  - Efficient Client/server models & protocols
  - Transparent networking, navigation & GUI with multimedia access & dissemination for data visualization
  - Mechanism for distributed computing such as CGI.Java
- With improved performance (price/performance) & the availability of Linux, Web Services (SOAP, WSDL, UDDI,WSFL), COM technology it is easy to develop loosely coupled distributed applications

Collaboratory: Working Together Apart

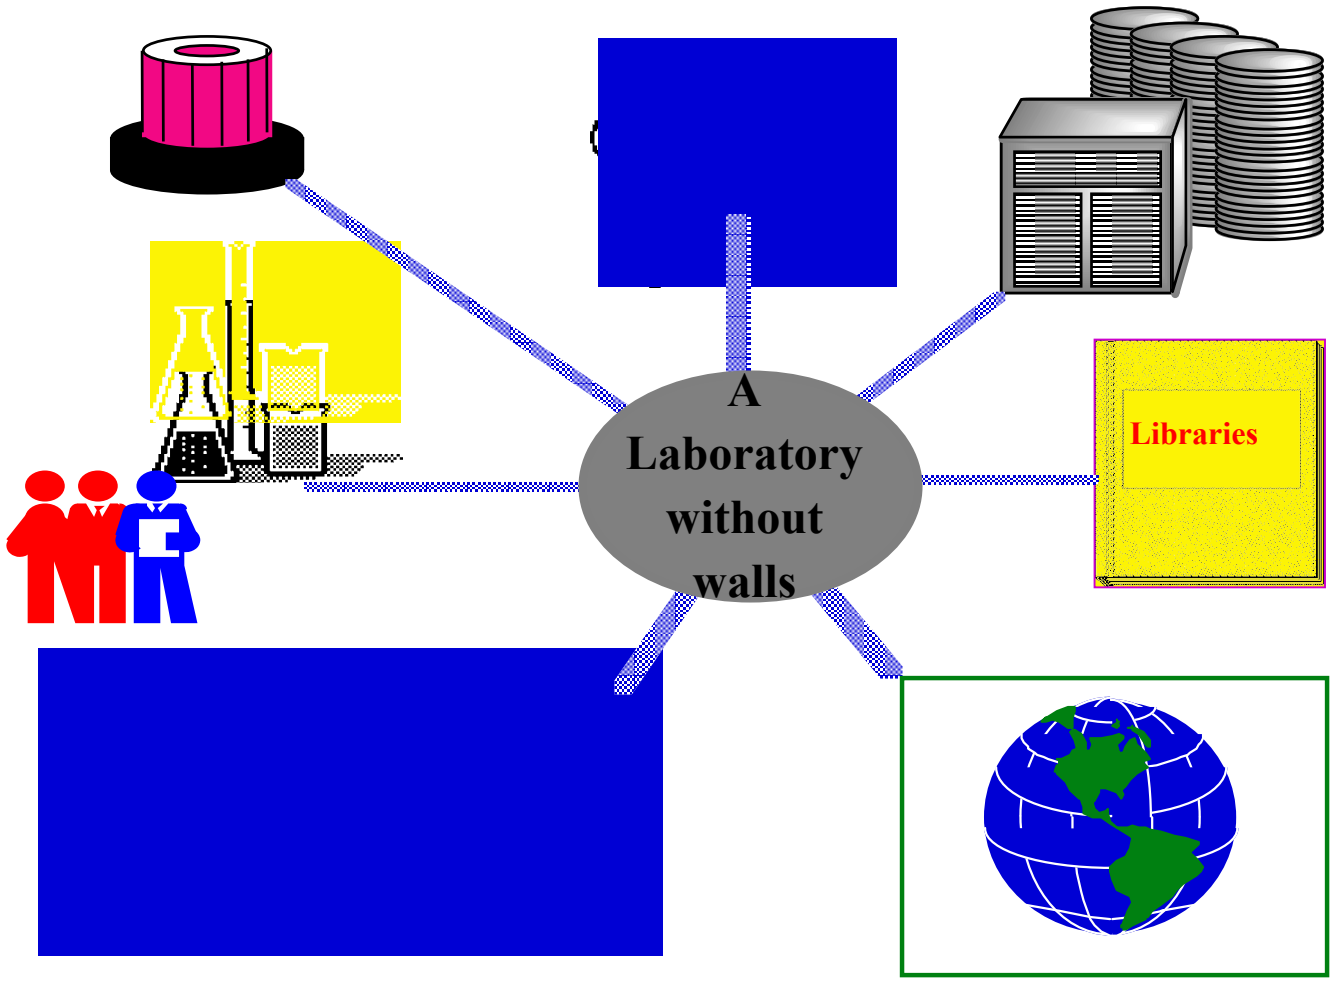

Tools

- Chat
- Email
- Video Conferencing
- VR
- White Boards
-- Web Portal

Web based remote monitoring & Control

- Main Computer Hall
- PC-Instruments Lab
- Mod Lab 2-294S Hall
- Hall # 7
- North Gate
- Central Stores Unit
- Home

Main Computer Hall :

Parameters

13 Dec 99 23:24:55

Light [tp1]	Normal
Light [tp2]	Normal
Temp [tp1]	20.46 C
Temp [tp2]	20.35 C
Humidity [tp1]	64.75
Humidity [tp2]	65.53
Reserve	N.C
Reserve	N.C
Power -5V DC	OFF
Power +5V DC	OFF
Power -12V DC	OFF
Power +12V DC	OFF

Password :

Relay 1
Relay 2
Relay 3
Relay 4

Back Home

TELESCIENCE for NRL

ANUNET with max of 512 Kbps Bandwidth

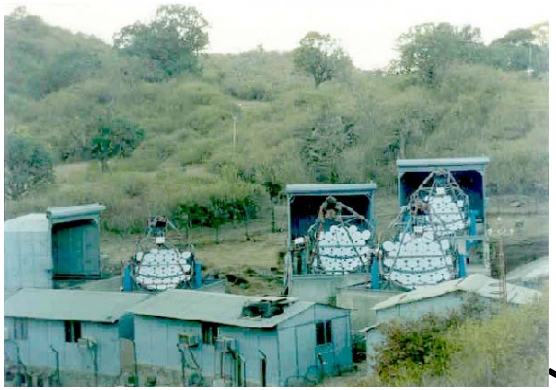

NRL Laboratory

Simulations
On ANUPAM at
BARC, Mumbai

Instrumentation

Remote Control & Mon

National Collaborations

- MOU with NCMRWF, New Delhi for developing parallel version of weather models on ANUPAM
- MOU with NIC & VSSC for Collab-CAD software development
- MOU with VJTI, Mumbai (Solve CFD related computational problems)

Projects through BRNS

- IIT, Mumbai (Parallel Data Base, Parallel File System)
- IIT, Kanpur (Parallel CFD Solver)
- IISC, Banglore (Wormhole Routing, Parallel Compiler)
- SNDT, Mumbai (Parallelization of Algorithms)

Consortium approach for Software

Collaborative Software Architecture based on Free & Open Source Software components (FOSS)

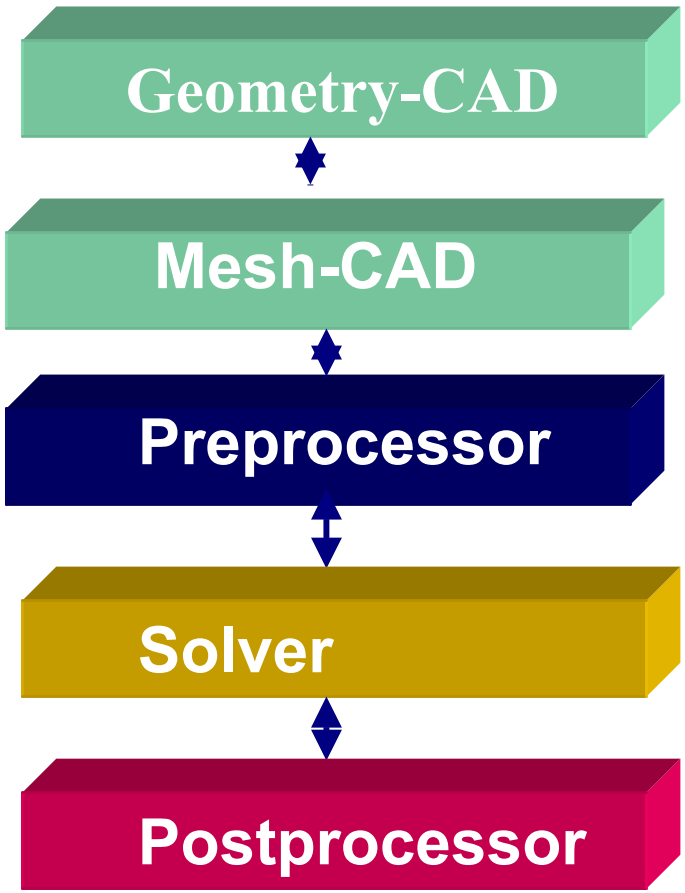

Collab-CAD, NIC, DIT

IITB

IITK

CD, BARC

Use Compatible Interfaces

- CGNS
- IGES

Snapshots of CollabCAD, 3D Java Viewer

Scientist also now need Pre & Post for scientific applications

Analysis – a very general model
Can we tie all components tightly by software?

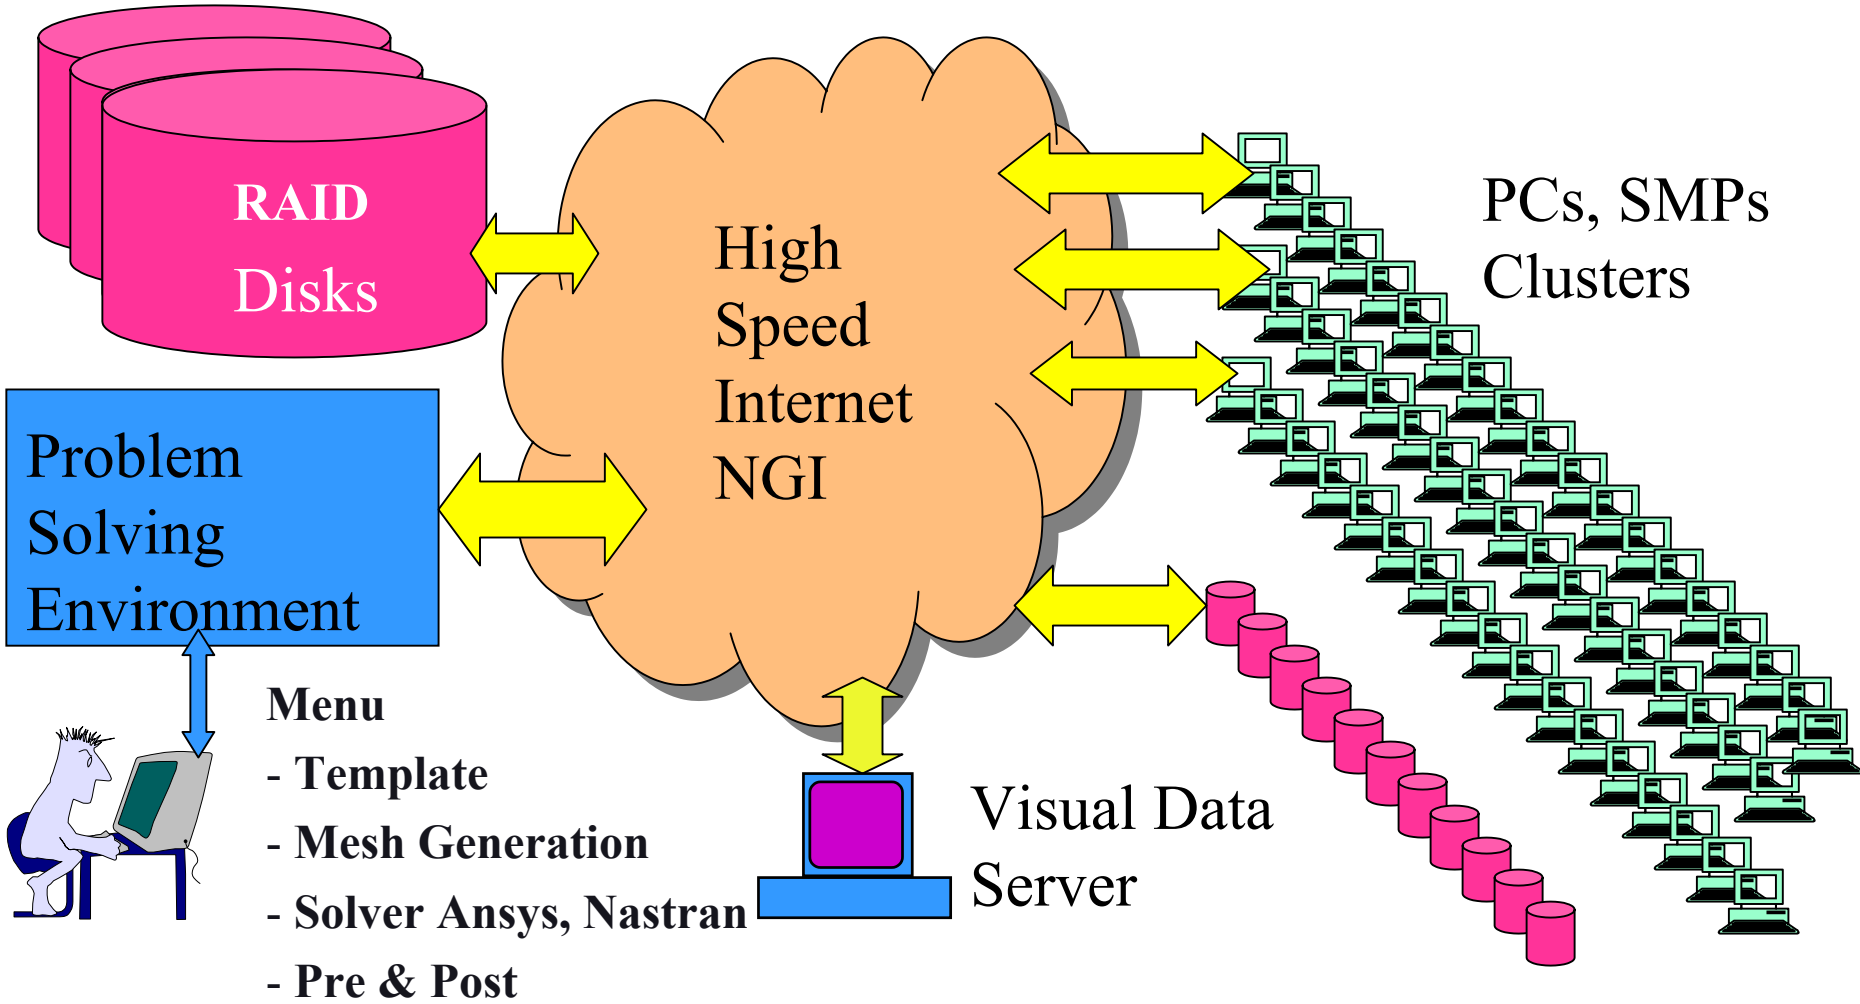

Computer Assisted Science & Engineering CASE

Virtual Organization : Grids

Resource sharing and coordinated problem solving in dynamic, multiple R&D units : Millions of users, Thousand Organizations, Many Countries

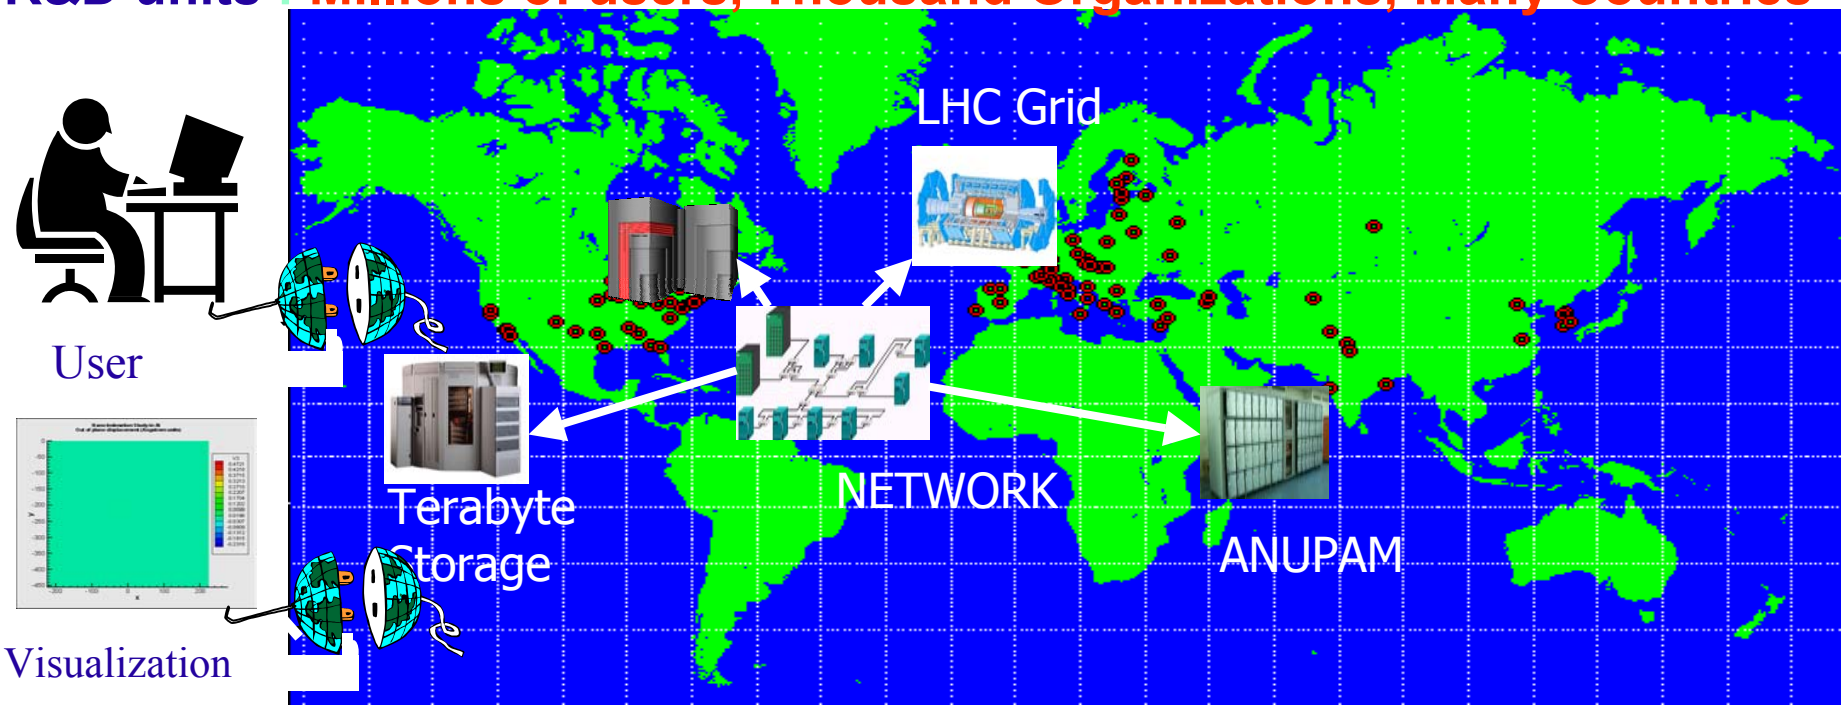

Making Information Technology (IT) as easy to use as plugging into electrical or TV socket

Simple philosophy:
Seamless integration of distributed environments to aggregate Computing power, Storage Capacity, Software and Visual Capability

Are Grids a Solution?

"Grid Computing" means different things to different people.

Goals of Grid Computing	Technology Issues
Reduce computing costs	Clusters
Increase computing resources	Internet infrastructure
Reduce job turnaround time	MPP solver adoption
Enable parametric analyses	Administration of desktop
Reduce Complexity to Users	Use middleware to automate
Increase Productivity	Virtual Computing Centre

"Dependable, consistent, pervasive access to resources"

What is needed?

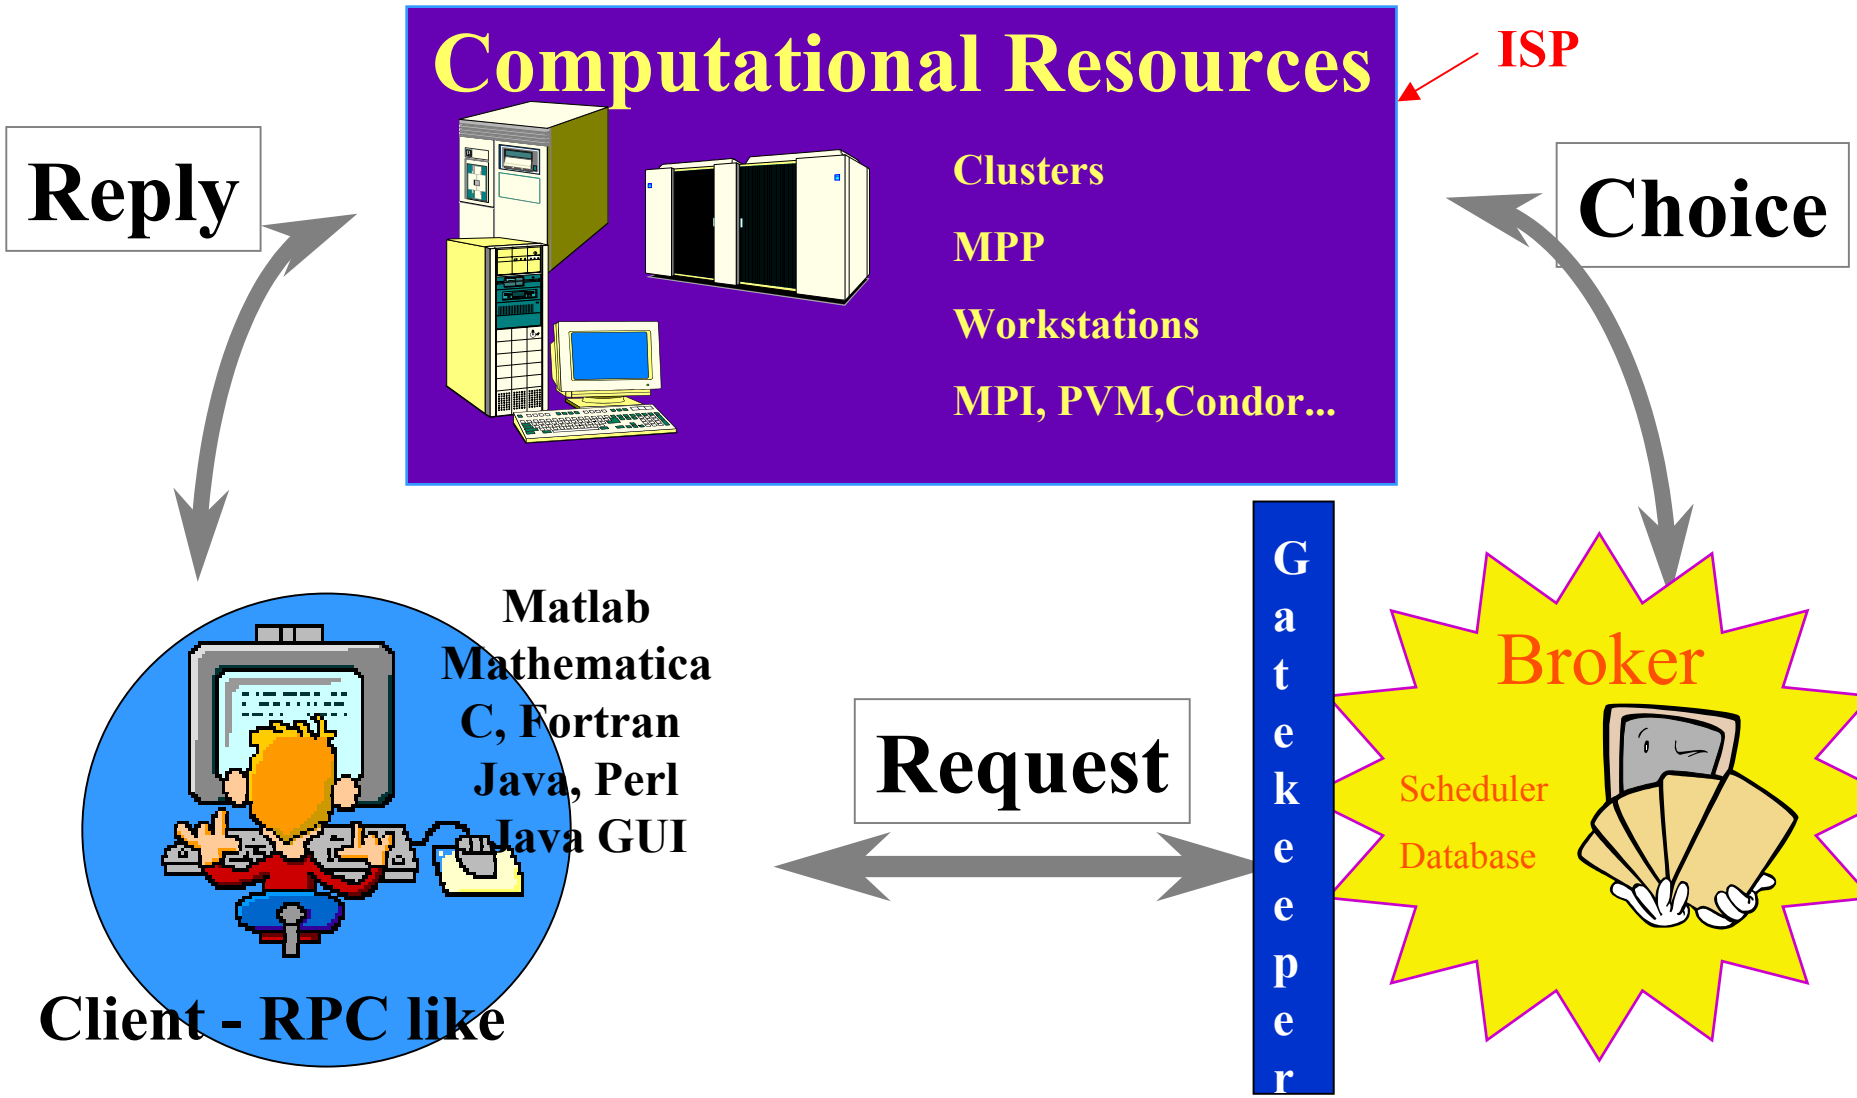

GRID CONCEPT

User Access Point

Result

Resource Broker

Grid Resources

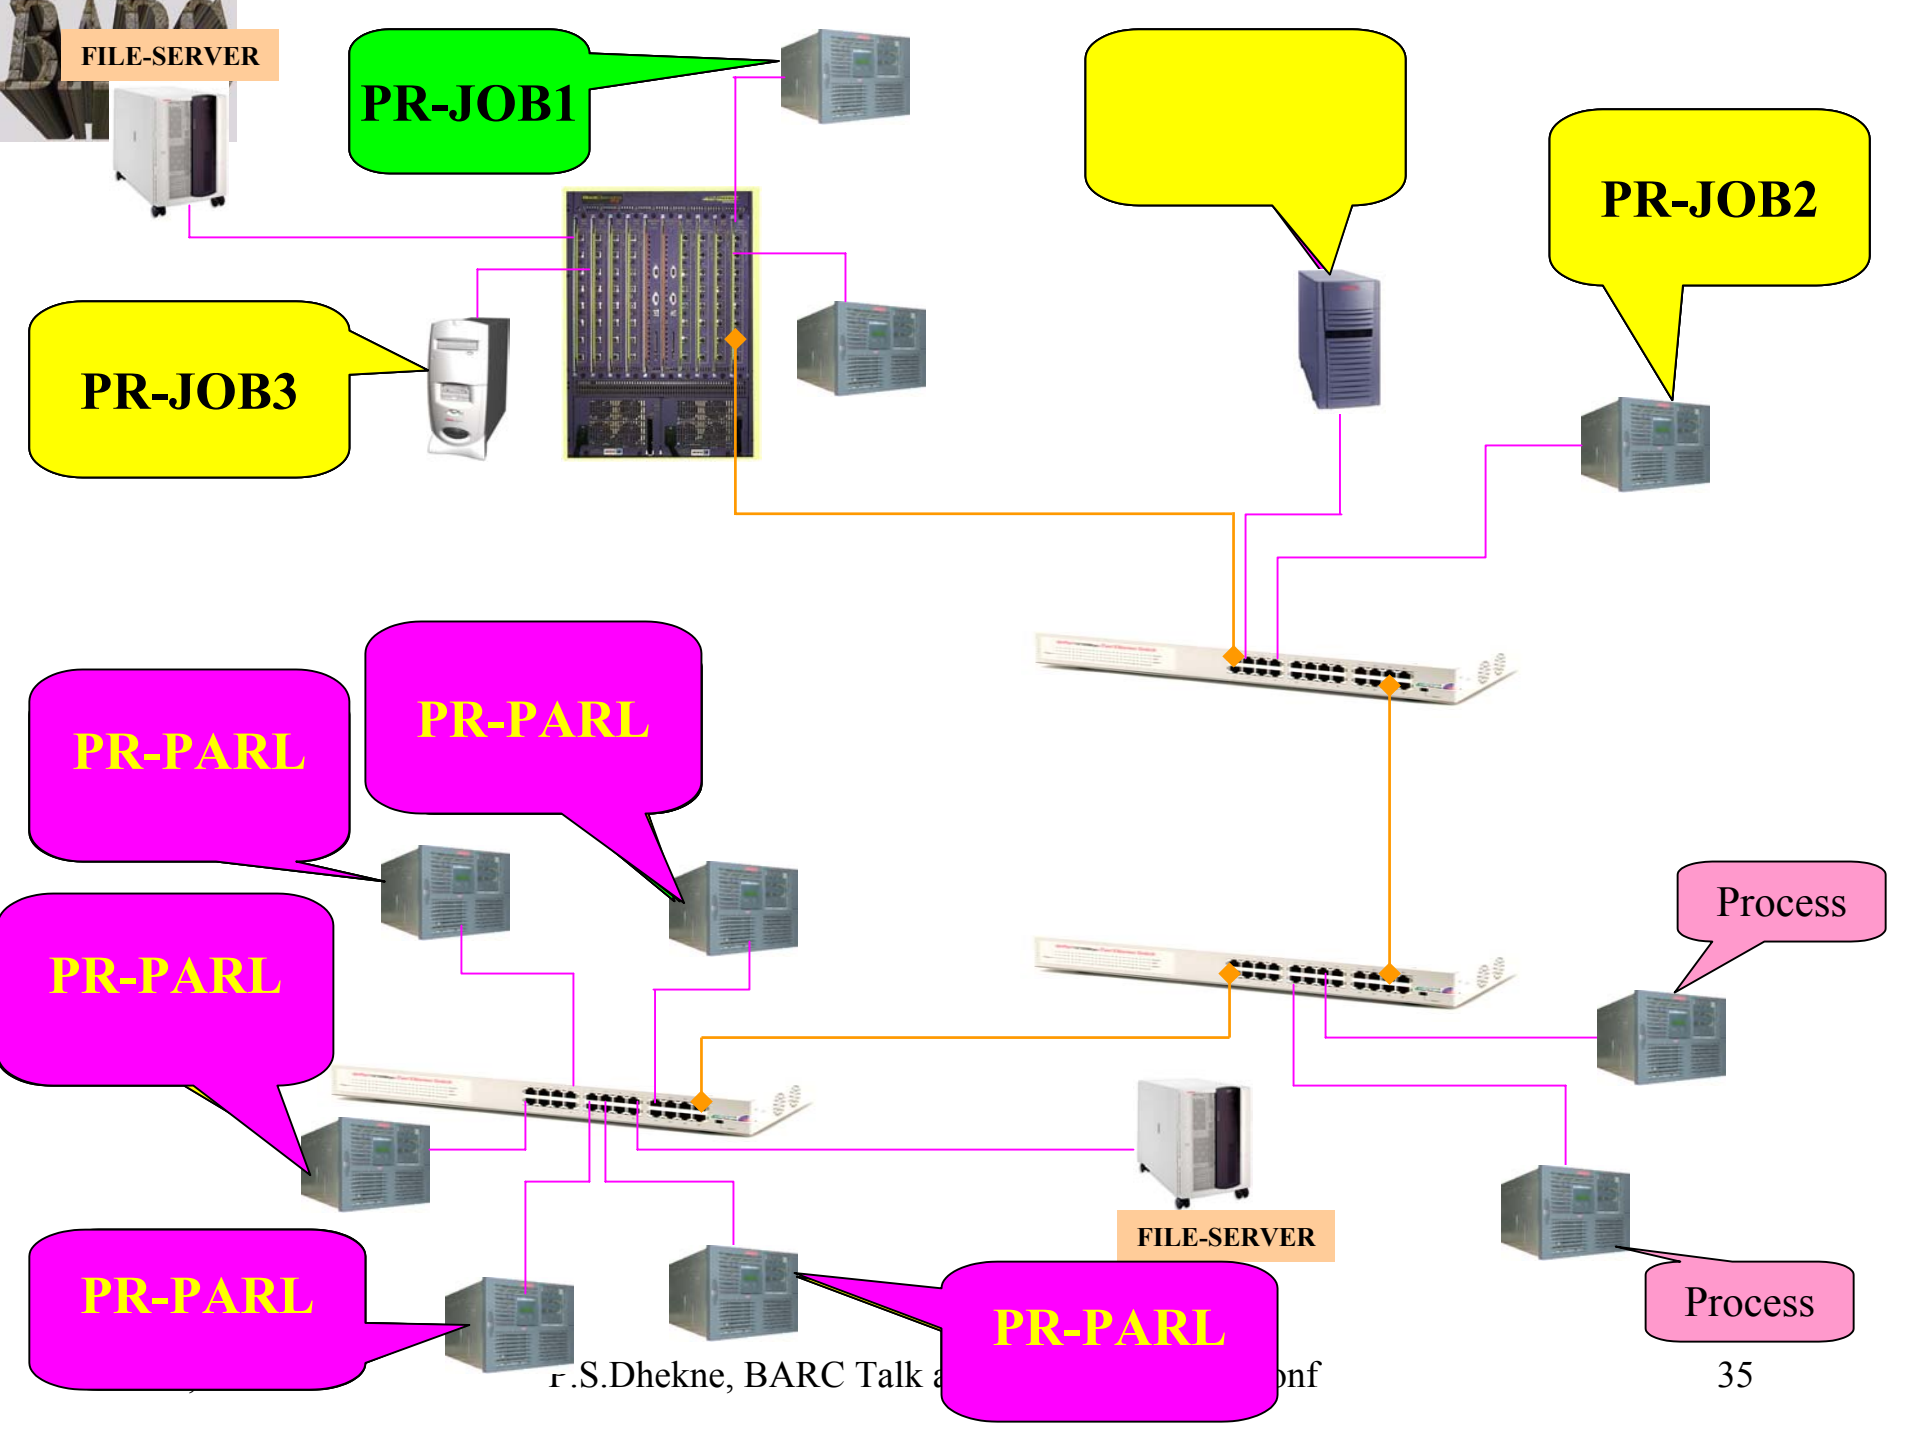
FILE-SERVER
PR-JOB1
PR-JOB2
PR-JOB3
PR-PARL
PR-PARL
PR-PARL
PR-PARL
PR-PARL
Process
Process
FILE-SERVER

What does the Grid do for you?

- You submit your work
- And the Grid
  - Finds convenient places for it to be run
  - Organises efficient access to your data
    - Caching, migration, replication
  - Deals with authentication to the different sites that you will be using
  - Interfaces to local site resource allocation mechanisms, policies
  - Runs your jobs, Monitors progress, Recovers from problems,Tells you when your work is complete
- If there is scope for parallelism, it can also decompose your work into convenient execution units based on the available resources, data distribution

Computational Science Pipeline

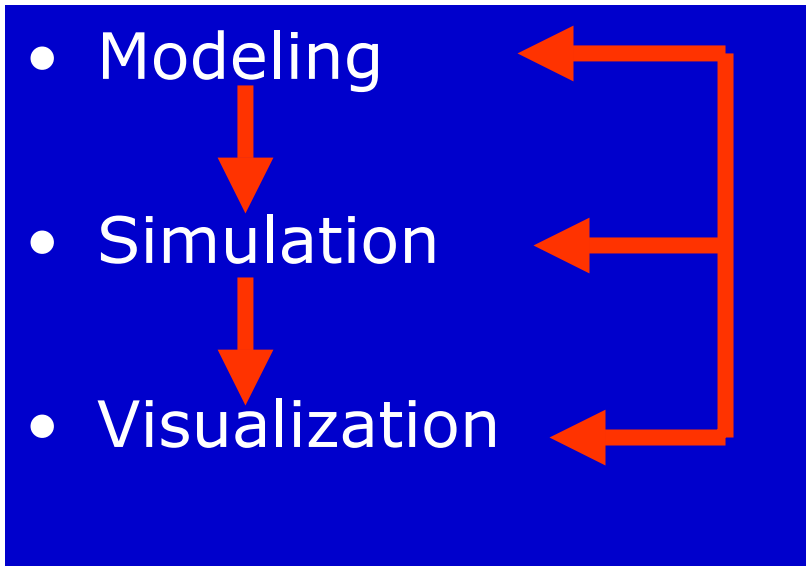

Integration

Three different compute systems, different application software for each function

But now user wants integrated solution

Why was it not done before?

- Use of Internet as infrastructure has become attractive now
  - Increasing bandwidth, advanced web services on Internet II and NGI
- Highly user friendly & mature Web Tech. : no training needed
- Provide access via "The Grid" to scarce remote computing resources of millions of PCs from anywhere to anybody
- Computing Problems are getting bigger (2x each year); need MPP
- Harnessing the power of Internet to aggregate and share resources spread across the globe is very challenging and highly cost effective and can give you unlimited computing capability
- Multiple ways are possible to implement Grid applications

Grid Services ("Middleware")

- **Standard services that**
  - **Provide uniform, high-level access to a wide range of resources (including networks)**
  - **Address inter-domain issues of security, policy, etc.**
  - **Permit application-level management and monitoring of end-to-end performance**
- **By integration of Distributed Technologies from Web Services, Dbase, Clustering, P2P, Collaboratories, High End Graphics to PKI based Security, you can provide these services**
- **Middleware-level and higher-level APIs and tools targeted at application programmers**
  - **Map between application and Grid**

Grid Computing Layers

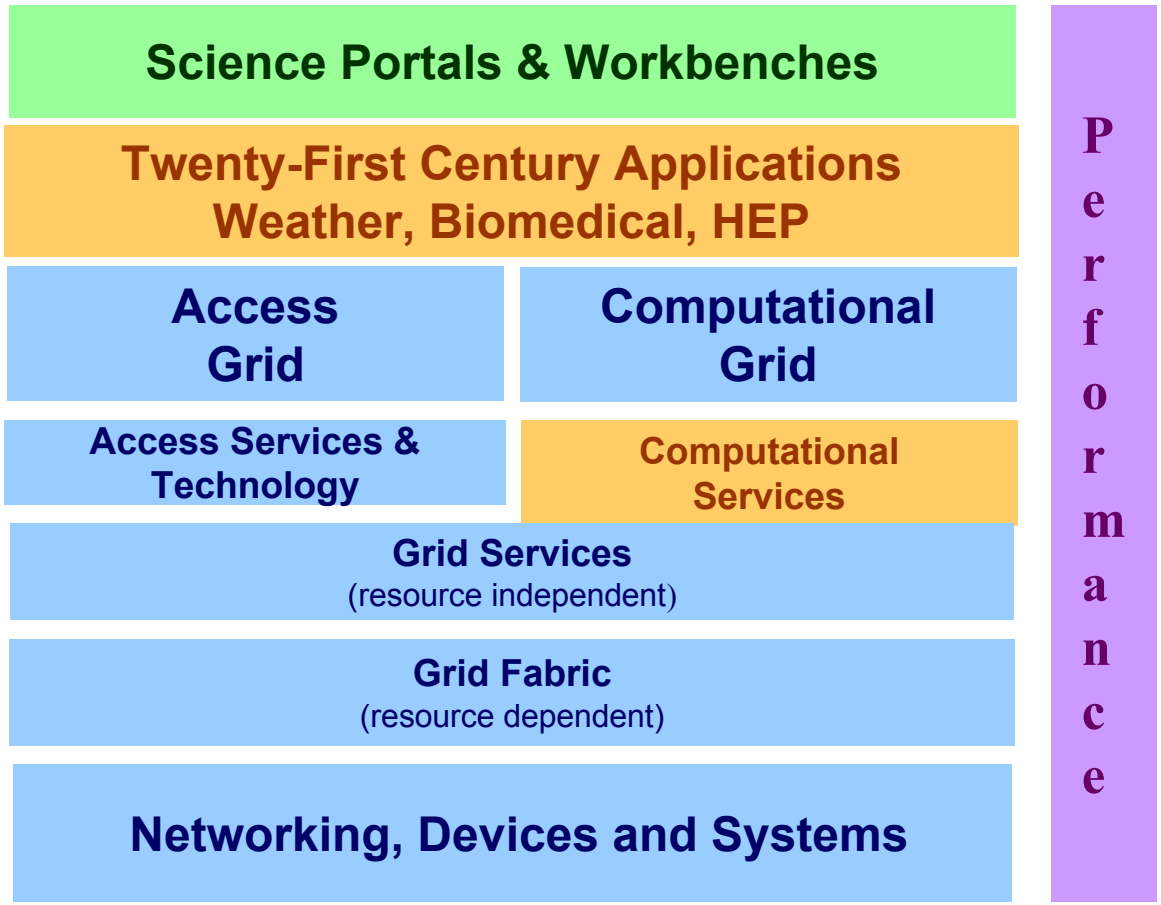

The Internet and burgeoning wired and wireless provide universal connectivity

On Going Projects

- **Globus - Sum of services architecture**
- **Legion - Object based & Parallel Env.**
- **Condor Group – Cluster Computing Env.**
- **Others**
  - **CERN, 3 Tier Data Grid**
  - **Harness Geist**
  - **Dongarra**
  - **Grid Forum**
  - **Particle Physics Data Grid ..... & many more**

Major Data Grid Projects

- **Earth System Grid (DOE Office of Science)**
  - **DG technologies, climate applications**
- **European Data Grid (EU)**
  - **DG technologies & deployment in EU**
- **GriPhyN (NSF ITR)**
  - **Investigation of "Virtual Data" concept**
- **Particle Physics Data Grid (DOE Science)**
  - **DG applications for HENP experiments**

LHC Computing

- **LHC (Large Hadron Collider) will begin taking data in 2006-2007 at CERN, Geneva.**
- **Data rates per experiment of >100 Mbytes/sec.**
- **>1 Pbytes/year of storage for raw data per experiment.**
- **World-wide collaborations and analysis.**
  - **Desirable to share computing and analysis throughout the world**
  - **Computing requirement is so huge that it can't be met by a single Computing Centre**

Data Grids for HEP

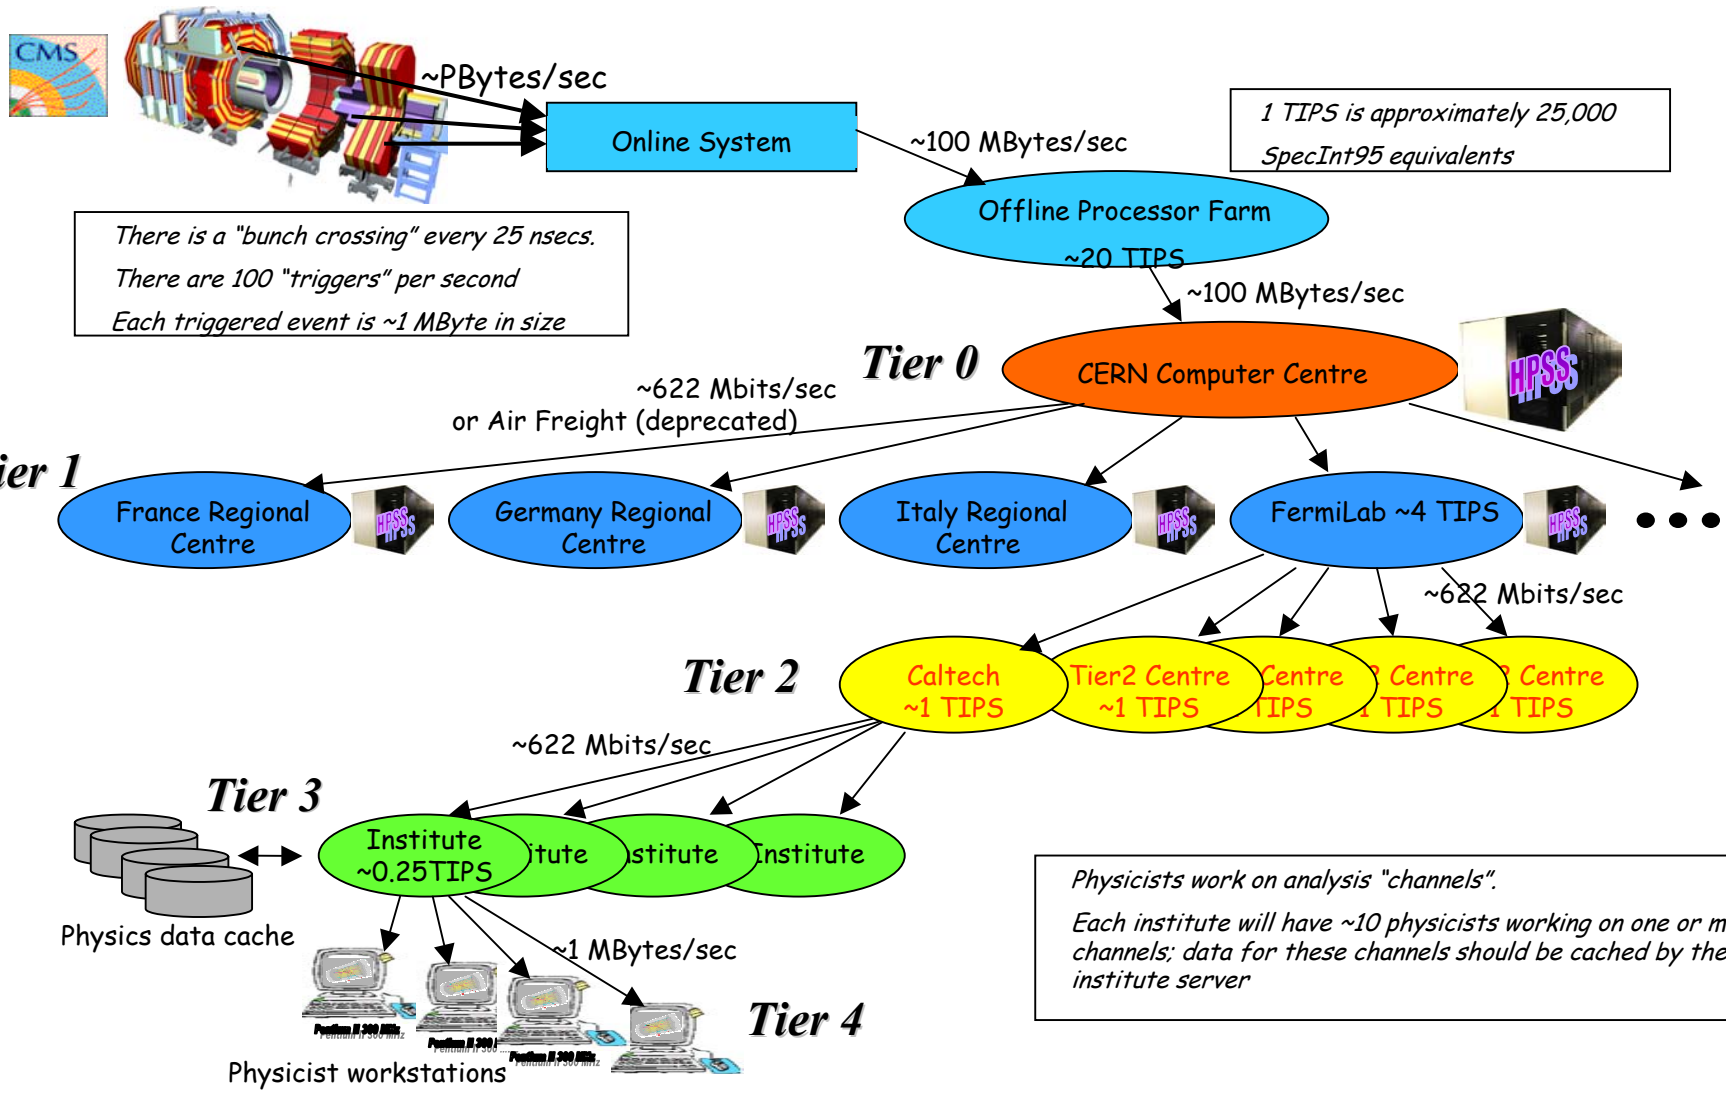

Image courtesy Harvey Newman, Caltech

International Collaboration

- India became a CERN observer state in 2002
- Large Hadron Collider (LHC) Grid Software Development, DAE-CERN Protocol agreement on computing for LHC data analysis, a DATA Grid called LCG

  ~10 people working in India for 5 years amounting to 7.5 MSWF
- BARC developed software is deployed at LCG, CERN

  - Co-relation Engine, Fabric management

  - Problem Tracking System (SHIVA)

  - Grid Operations (GRID VIEW)

  - Quattor a system administration toolkit

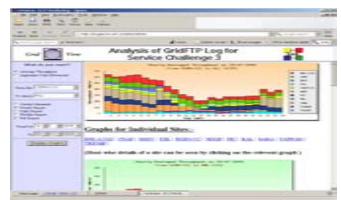

LHC: CERN-DAE COLLABORATION

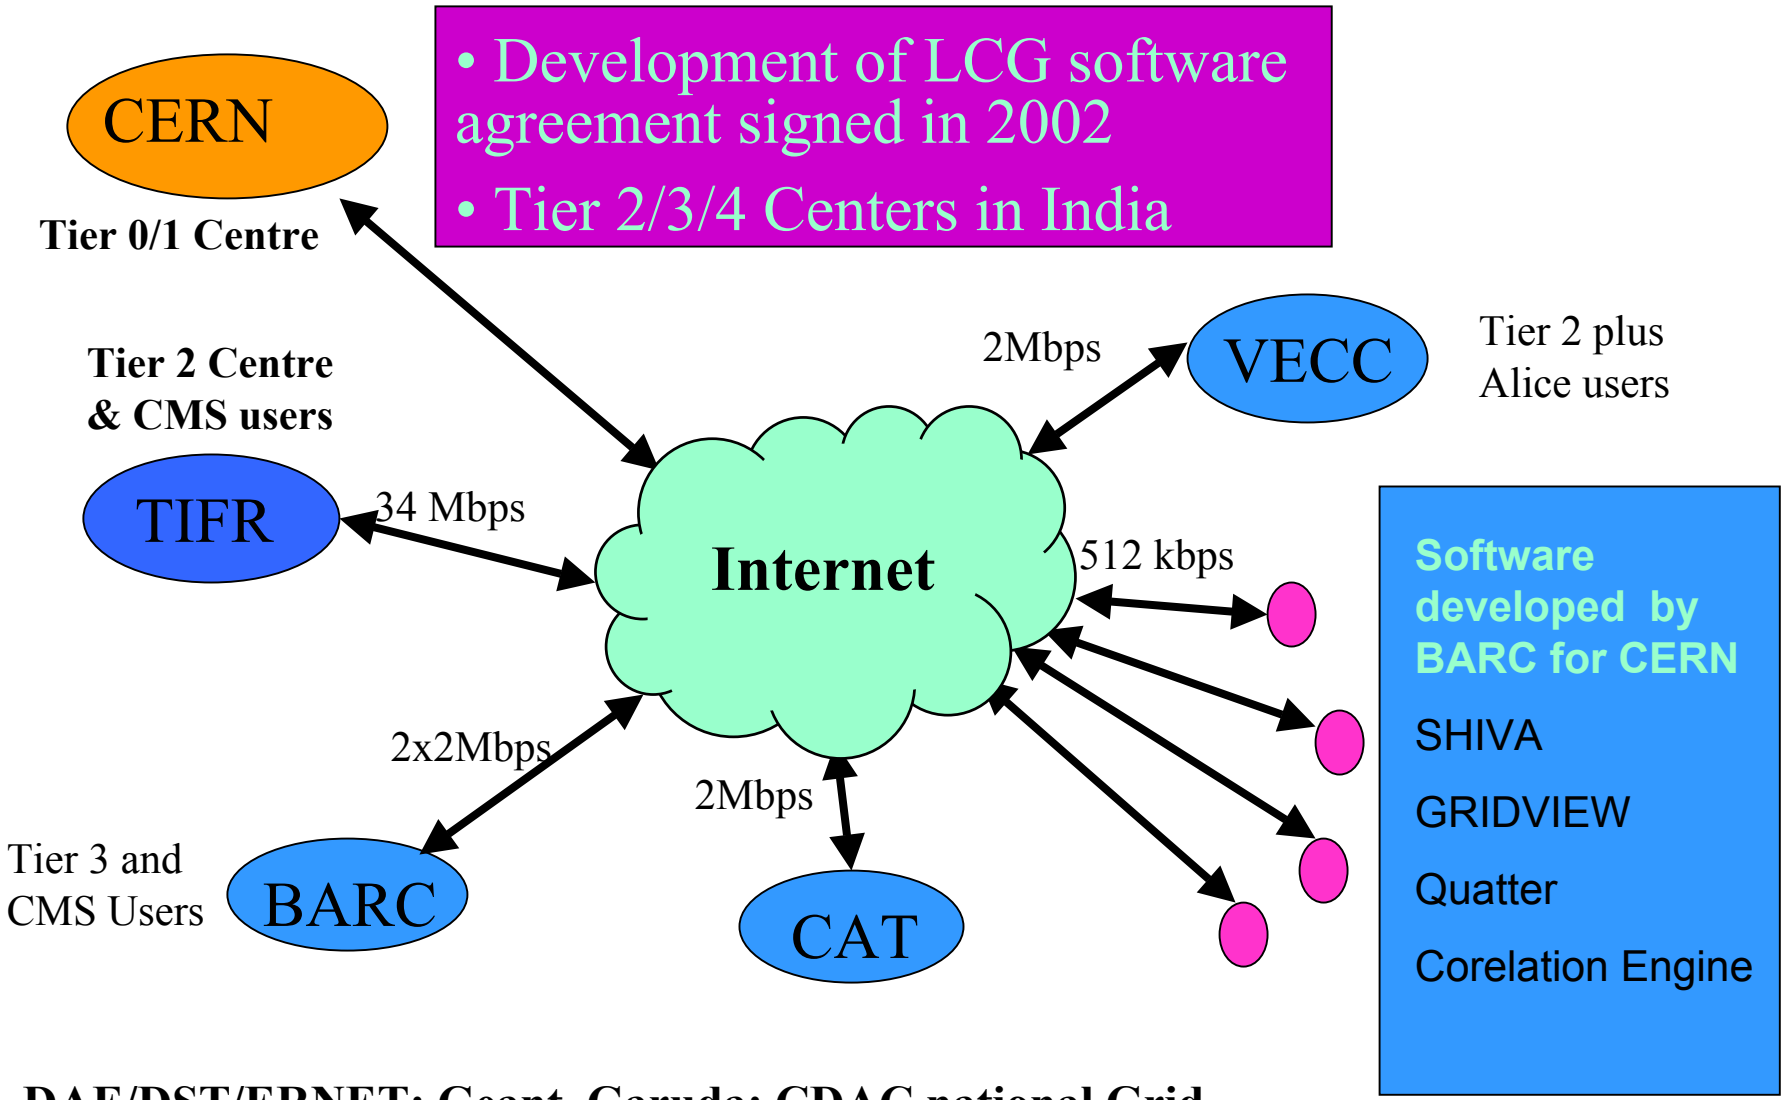

DAE/DST/ERNET: Geant, Garuda: CDAC national Grid

Campus Grid at BARC

Computing Grid system has been set up as a Test-Bed using existing Grid Technology Components

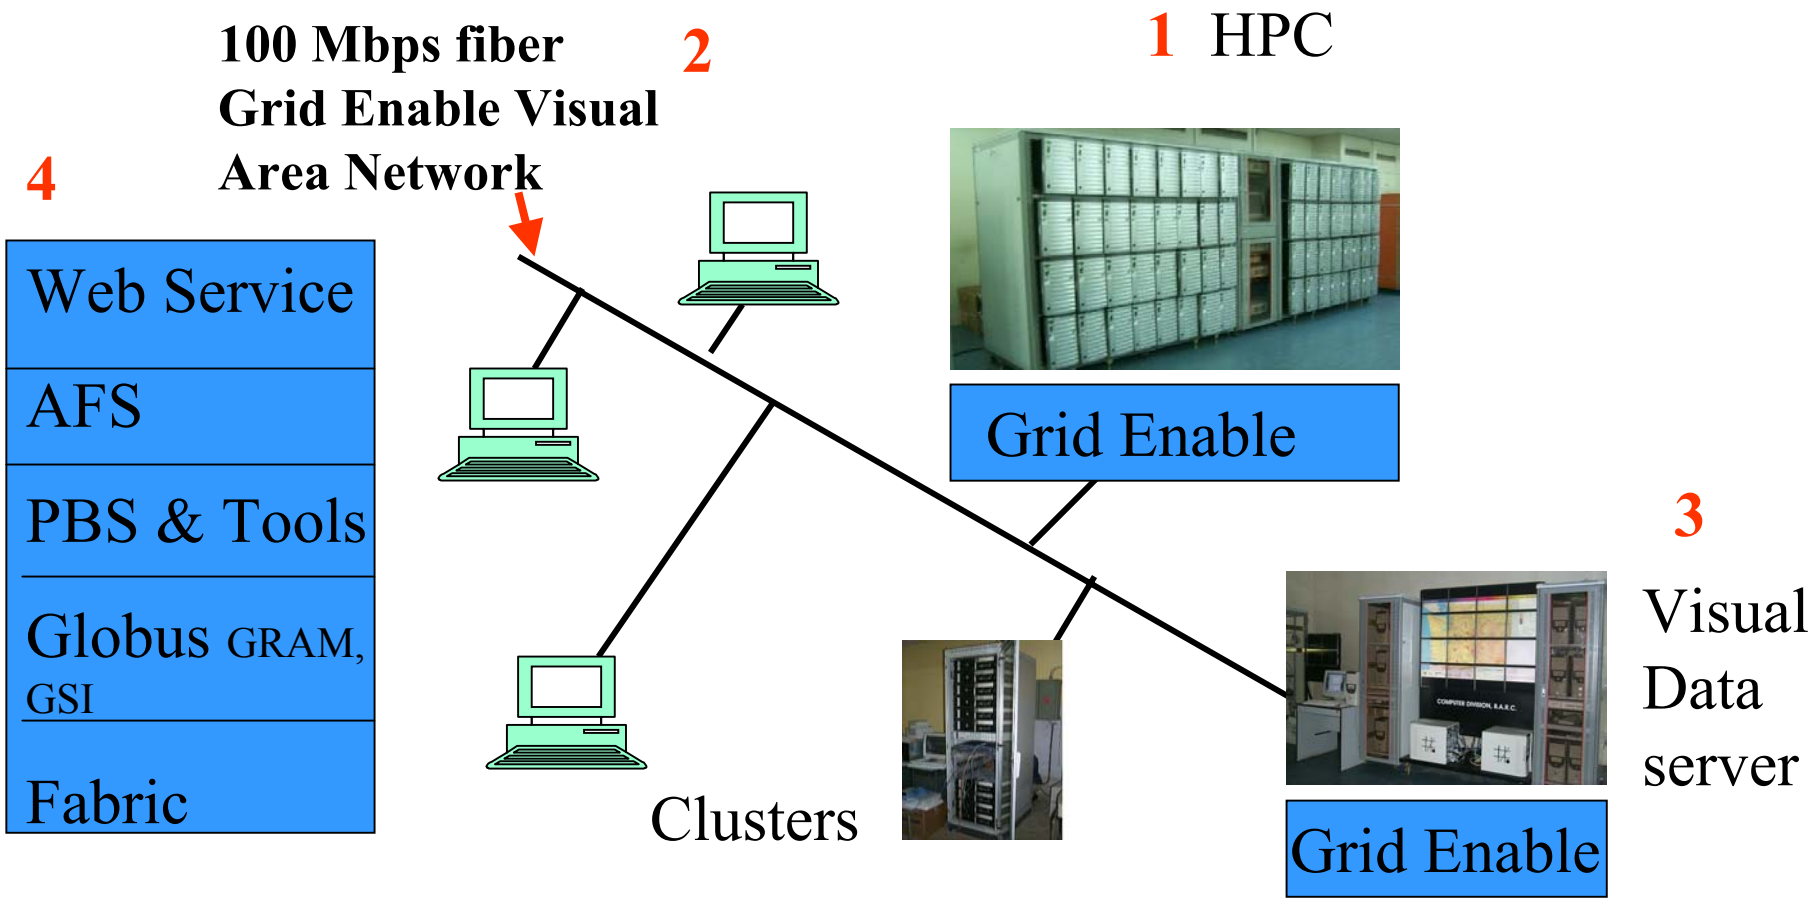

GRID PORTAL ARCHITECTURE

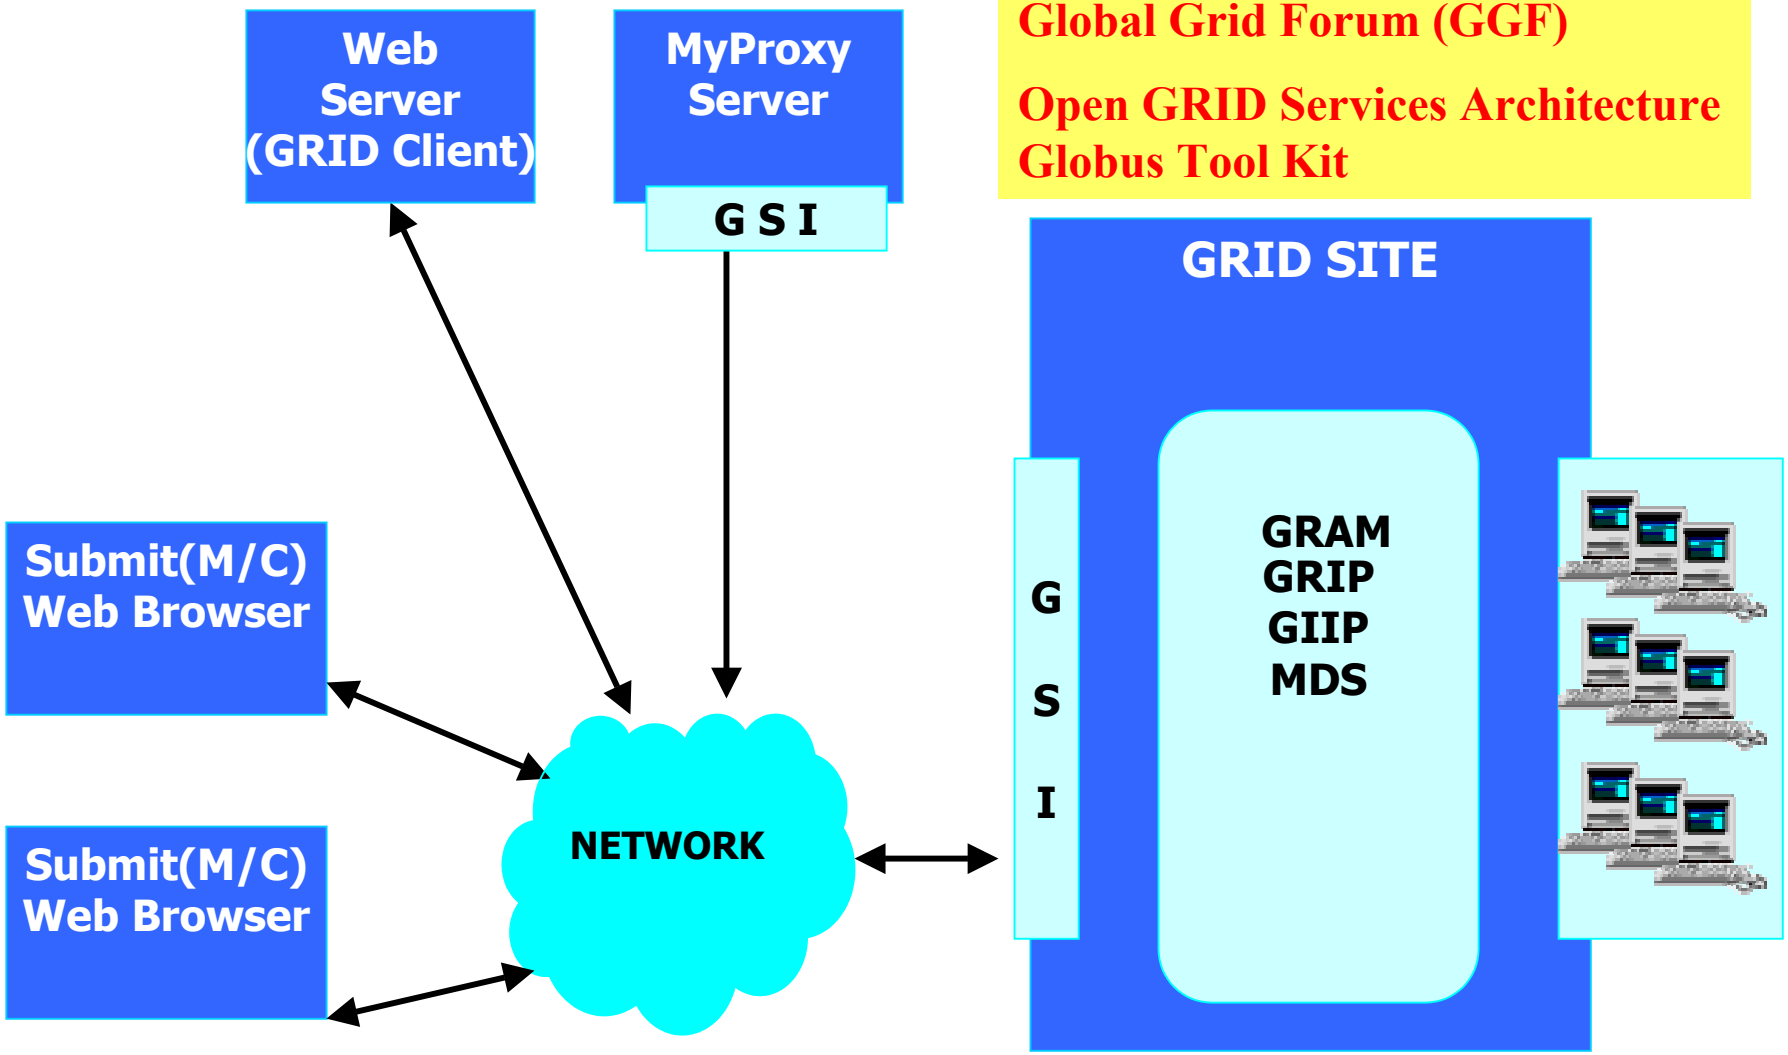

Snap (Showing Job Location)

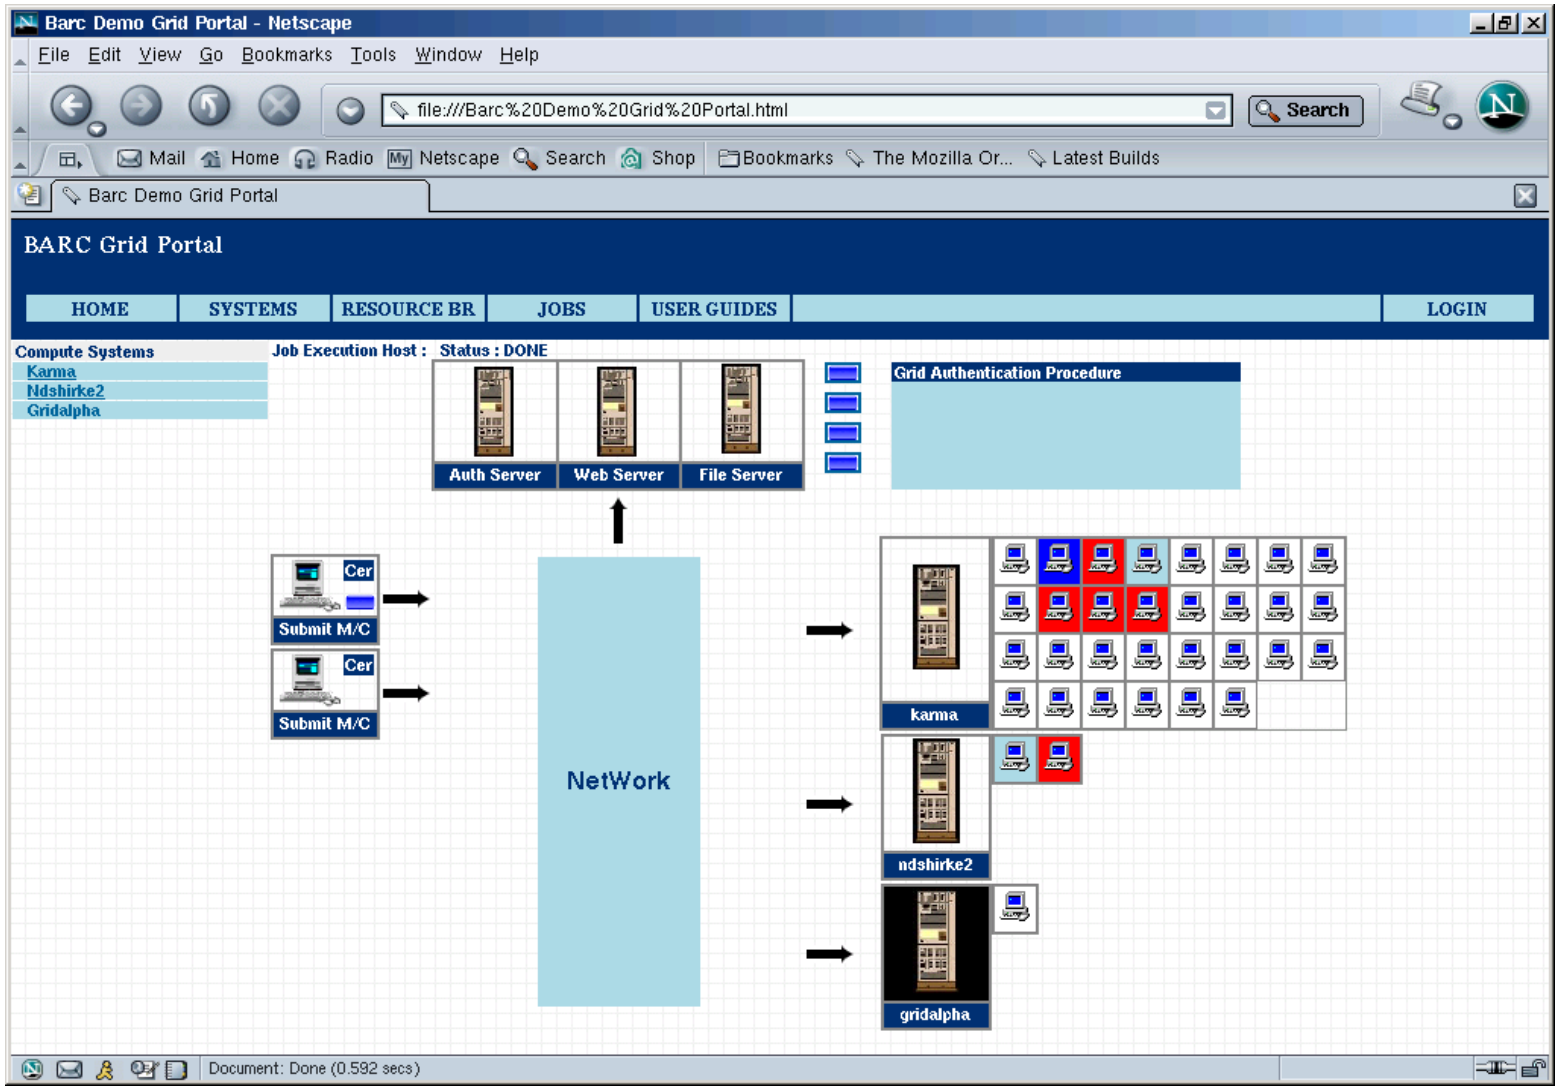

DAE Grid

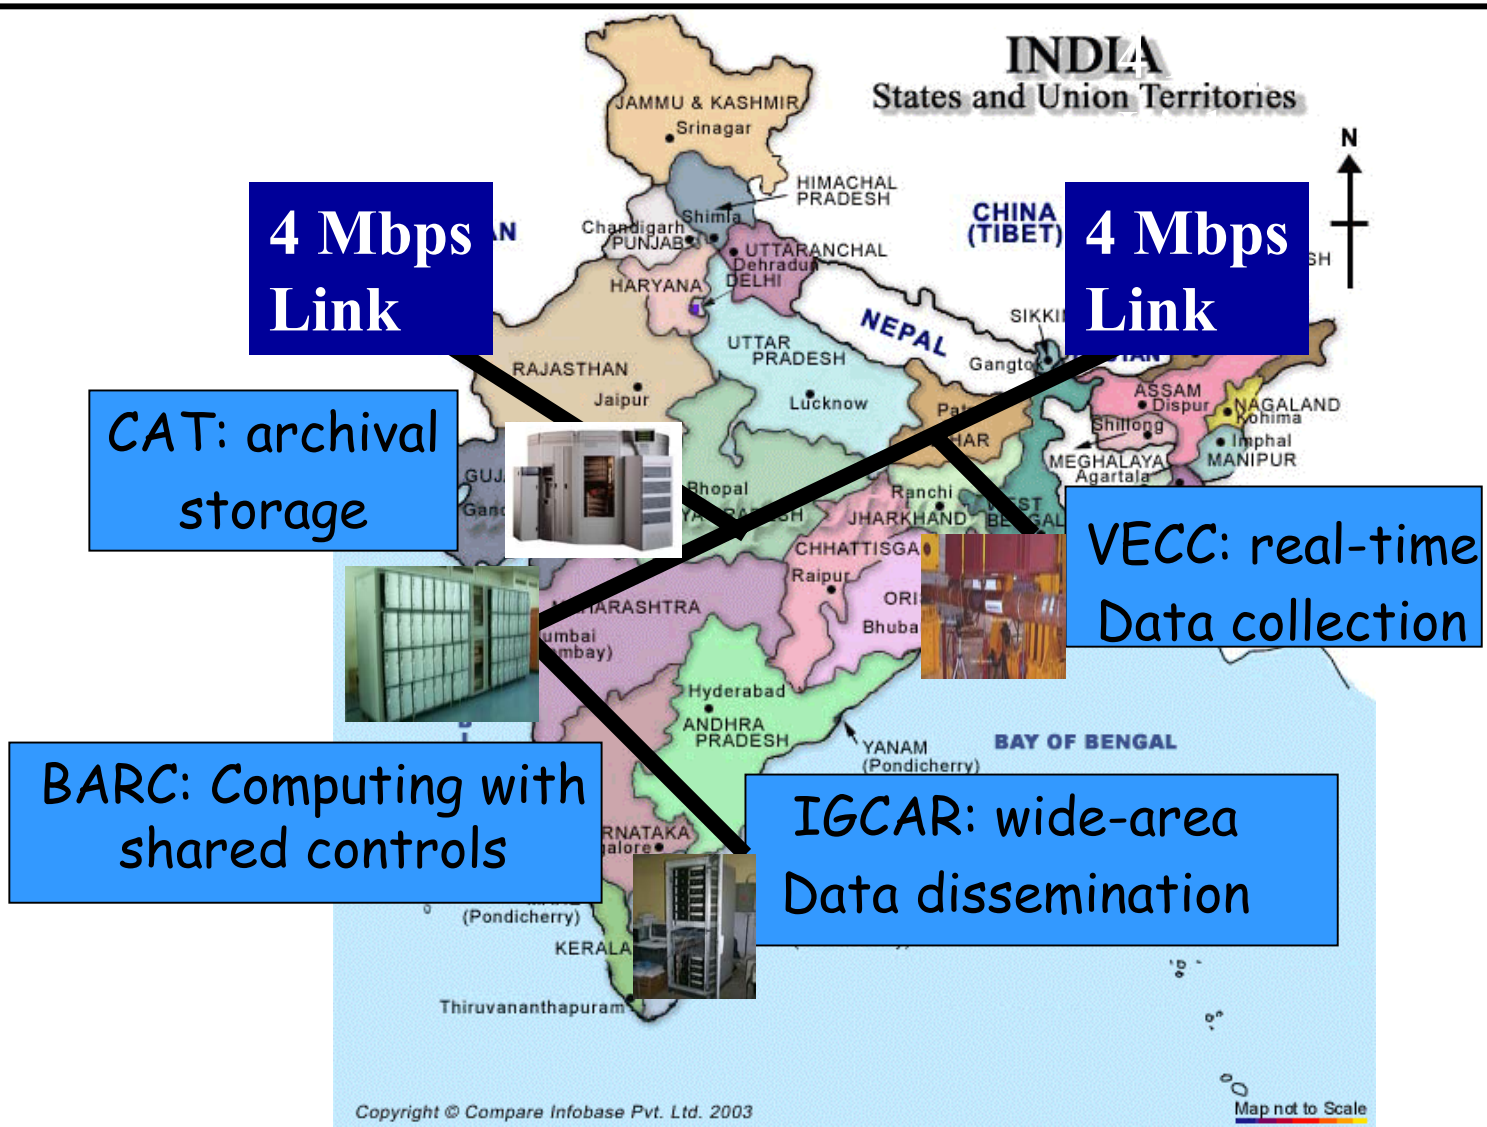

4 Mbps Link

4 Mbps Link

CAT: archival storage

VECC: real-time Data collection

BARC: Computing with shared controls

IGCAR: wide-area Data dissemination

Resource sharing and coordinated problem solving in dynamic, multiple R&D units

Real Challenge

- Future HEP experiments require massive amounts of computing, including data collection and storage, data access, database access, computing cycles, etc.
- How to provide seamless and transparent access to Grid Services - without compromising on security & convenience
- Conservation of network bandwidth (use on demand basis)
- How to do it inexpensive way ?
- How to meter use of resources ?

Issues Involved?

- Relatively New & Immature Technology
- Few standards & absence of good Tools
- Lack of expertise of users, developers & system support people
- Portability & Scalability to be resolved
- Added complexity due to WAN connectivity

Take-way Message !!!

- Users can expect UNLIMITED computing capacity in the future
- More efforts are required in developing application software which can use this enormous computing power and you can leverage on OSS for speedy development
- Though a research area today, Grid & Collaboratory concepts are becoming increasingly popular & will be in place sooner than you think.
- Time has come now, where you can't excel even in your own field unless you collaborate with others.


Thank You