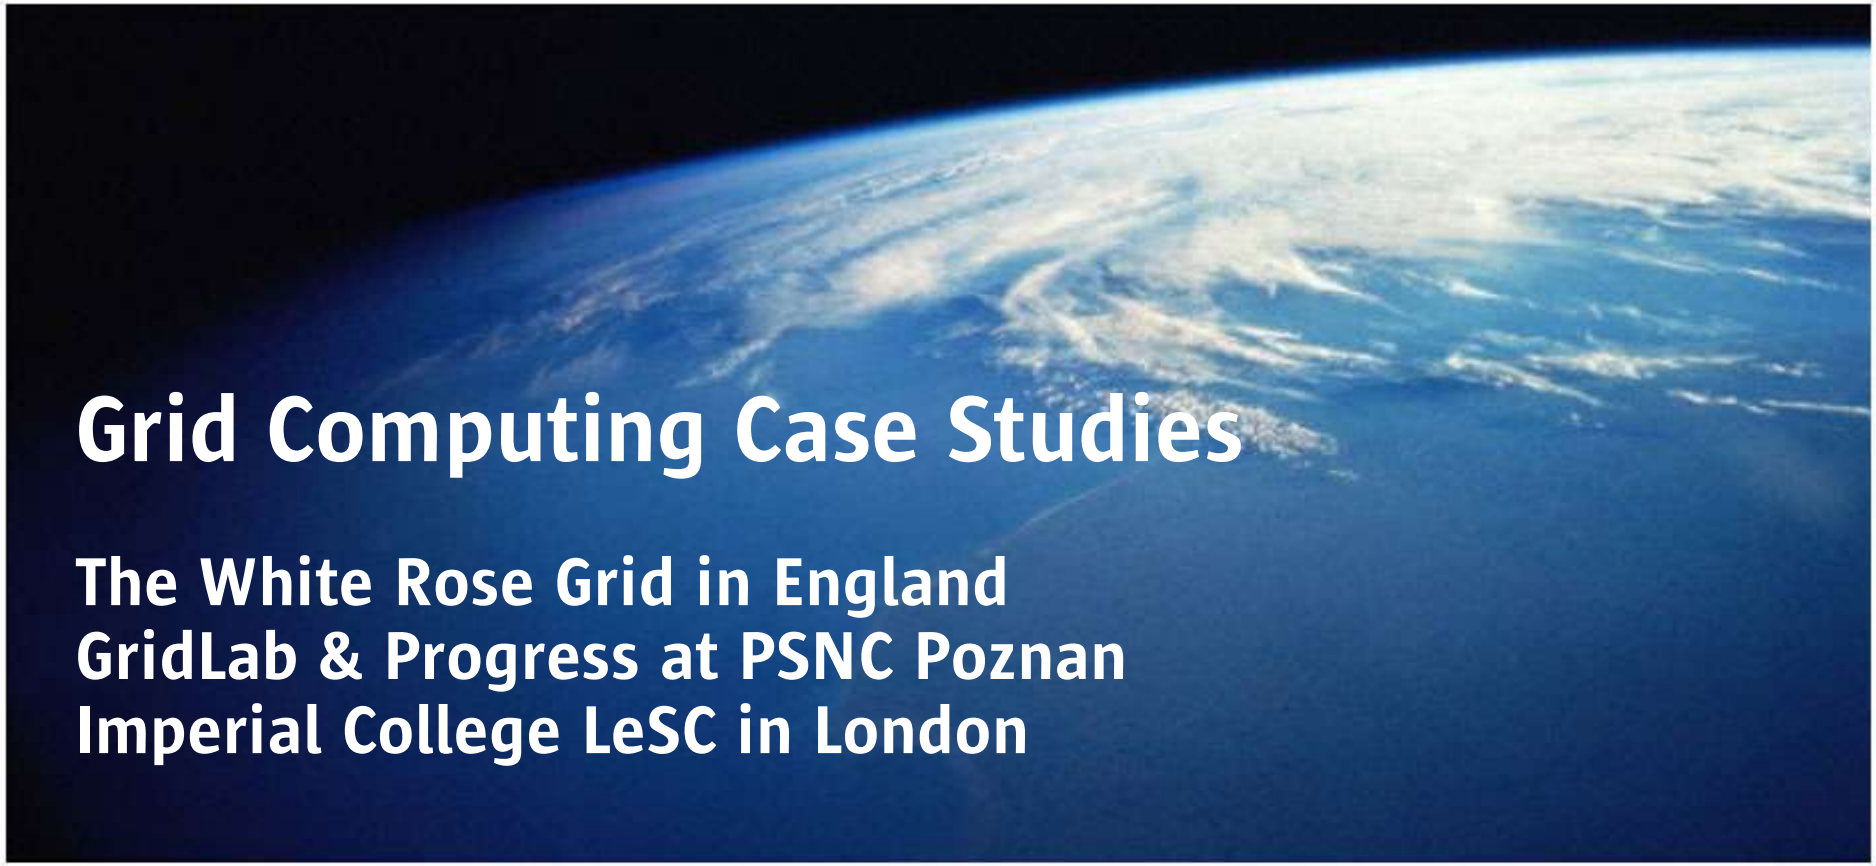

Grid Computing Case Studies

The White Rose Grid in England
GridLab & Progress at PSNC Poznan
Imperial College LeSC in London

Wolfgang Gentzsch, Director Grid Computing

Sun Microsystems

Sun microsystems
We make the net work.

Sun's Evolutionary Grid Strategy

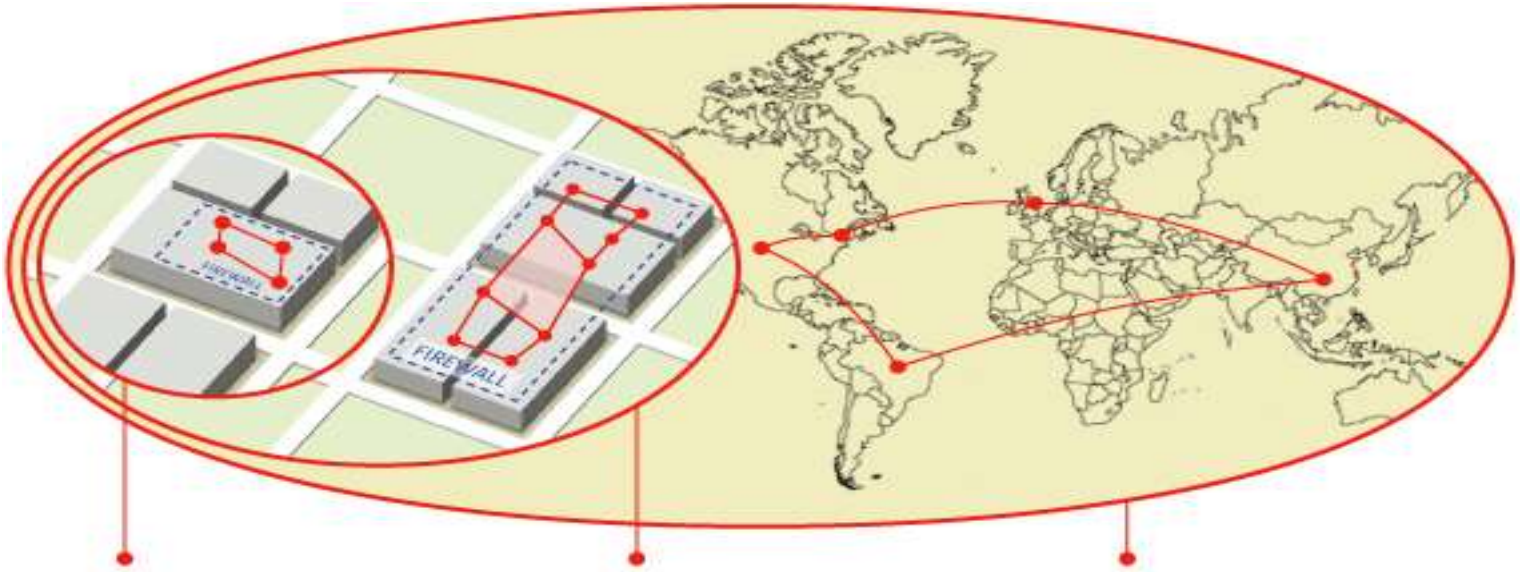

Department Grid
Departmental Computing

Enterprise Grid
Enterprise Computing

Global Grid
Internet Computing

From Department Grids, to Enterprise Grids, to Global Grids, to THE GRID

Critical Customer Requirements

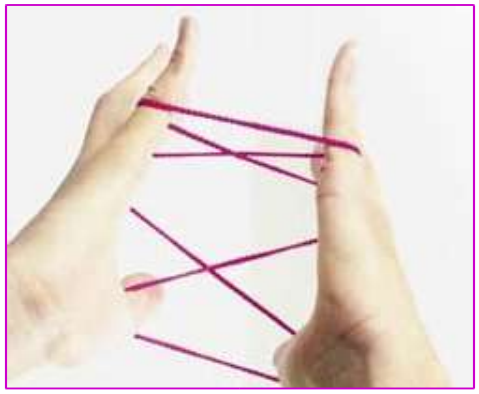

- ✓ More Compute Power
- ✓ Improved Resource Access
- ✓ Increased Productivity
- ✓ Better Utilization of Existing Resources
- ✓ Reduced Costs
- ✓ Reduced Management Complexity

Sun Grid Technology meets these needs *today!*

Grid Computing Benefits

- Access: transparent, remote, secure
- Virtualization: access services, not servers
- On Demand: resources at your finger tip
- Sharing: collaboration over the net
- Failover: migrate/restart applications
- Heterogeneity: platforms, OS's, software
- Utilization: increase from 20% to 80+%
- Productivity: more work in shorter time

Grid Computing & Web Services Environment

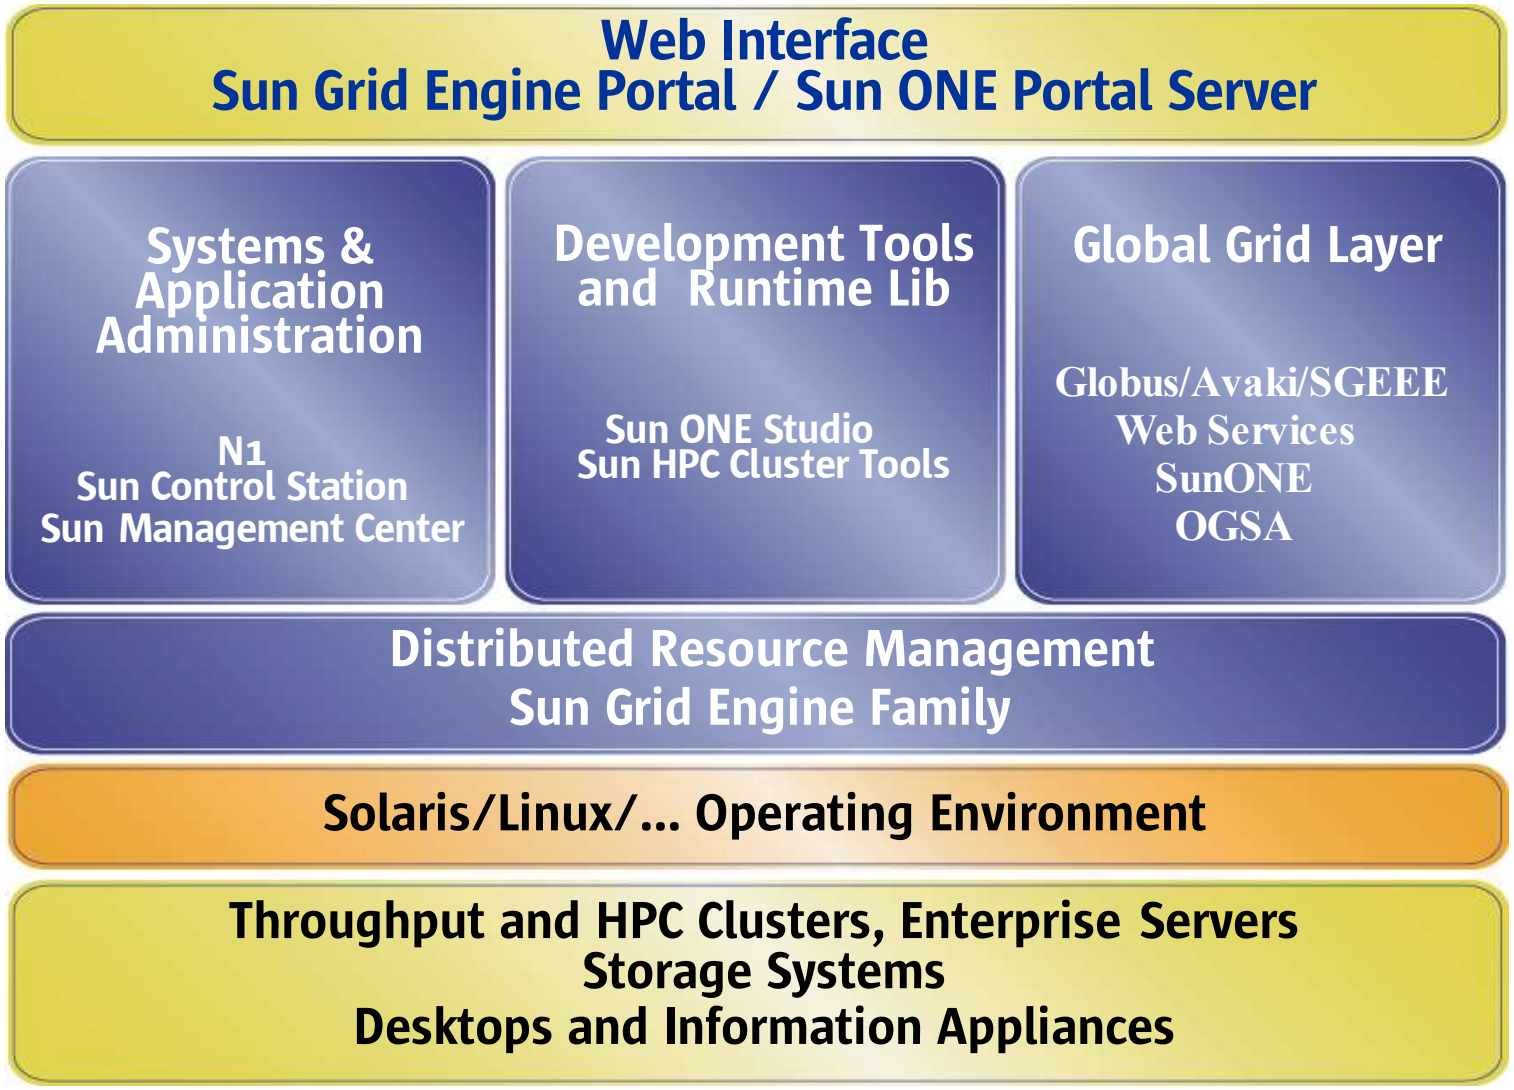

Sun Grid Products Roadmap

Sun Grid Engine 6.0

Scalability, analysis, monitoring, accounting, ease of installation, administration, scheduler, standards

Department & Enterprise

ClusterGrid SW Stack:
Grid appliance, preload, software product.
Avaki/SGE/Portal

Global Grid

Grid standards (OGSA), Globus GT3, research collaborations

H1FY04 | H2FY04 | H1FY05 | H2FY05

The Changing Grid Landscape

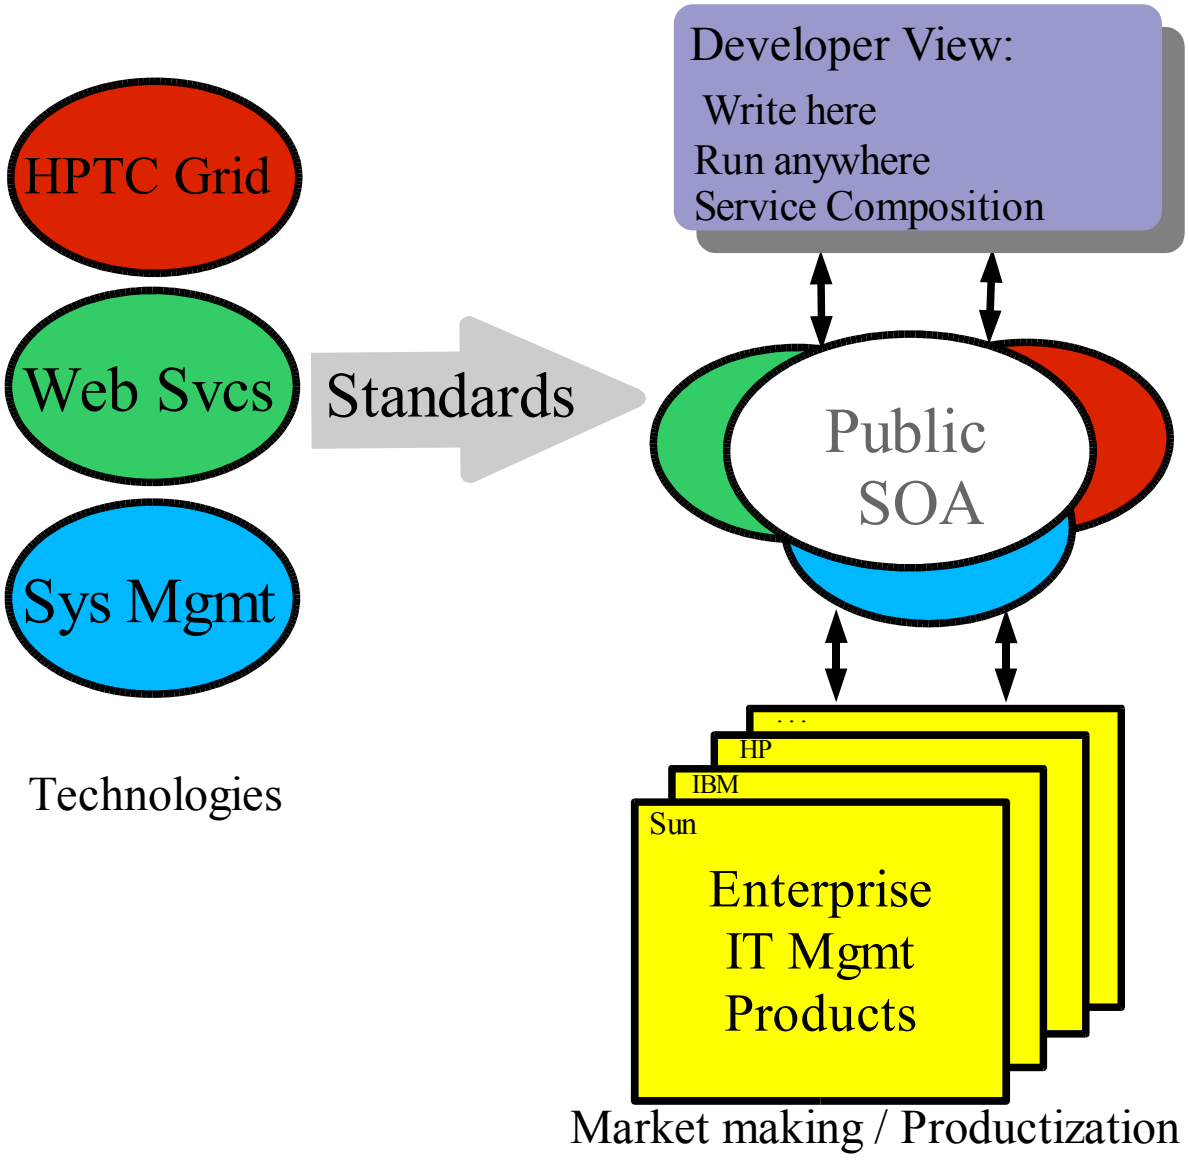

Grid Components: Compute Grid Racks "Grid-Ready"

An engineered, tested, integrated and supported Compute Grid *Solution*

- 1U-form-factor servers
- Gbit-E switch(es)
- Terminal server
- KVM
- Management Node with monitoring and management software
- Delivered in a rack

Part of Sun's new "Grid Infrastructure Solutions" Program

Sun Grid Manager

- SGEEE as dynamic job load manager
- SCS as integrated automatic grid install and config
- Software, image, inventory, and lights-out management
- Health and performance monitoring
- Automatic resource discovery via SCS/Jxta
- Single control/mgment grid interface (remote access)
- System setup/configuration in minutes (instead of days)
- Immediately utilize additional grid resources

Department Grid

Browser to CGM
(Remote Server Setup & Configuration)

Internet

Auto Download of Modules

Compute Grid Manager

SGEEE/SCS

Workload & System Mgmnt:
- Auto OS Deployment
- Grid Installation/Mgmt
- Central Server Mgmt
- Loadbalancing
- Resource Matching

Solaris Servers

Linux Servers

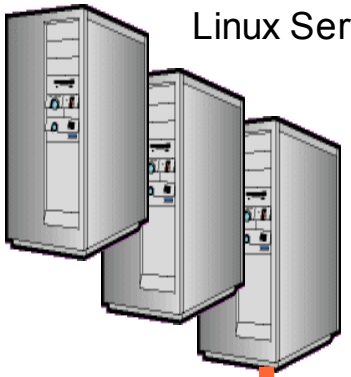

Workstations
(Linux or Solaris)

Corporate Firewall

Enterprise Grids

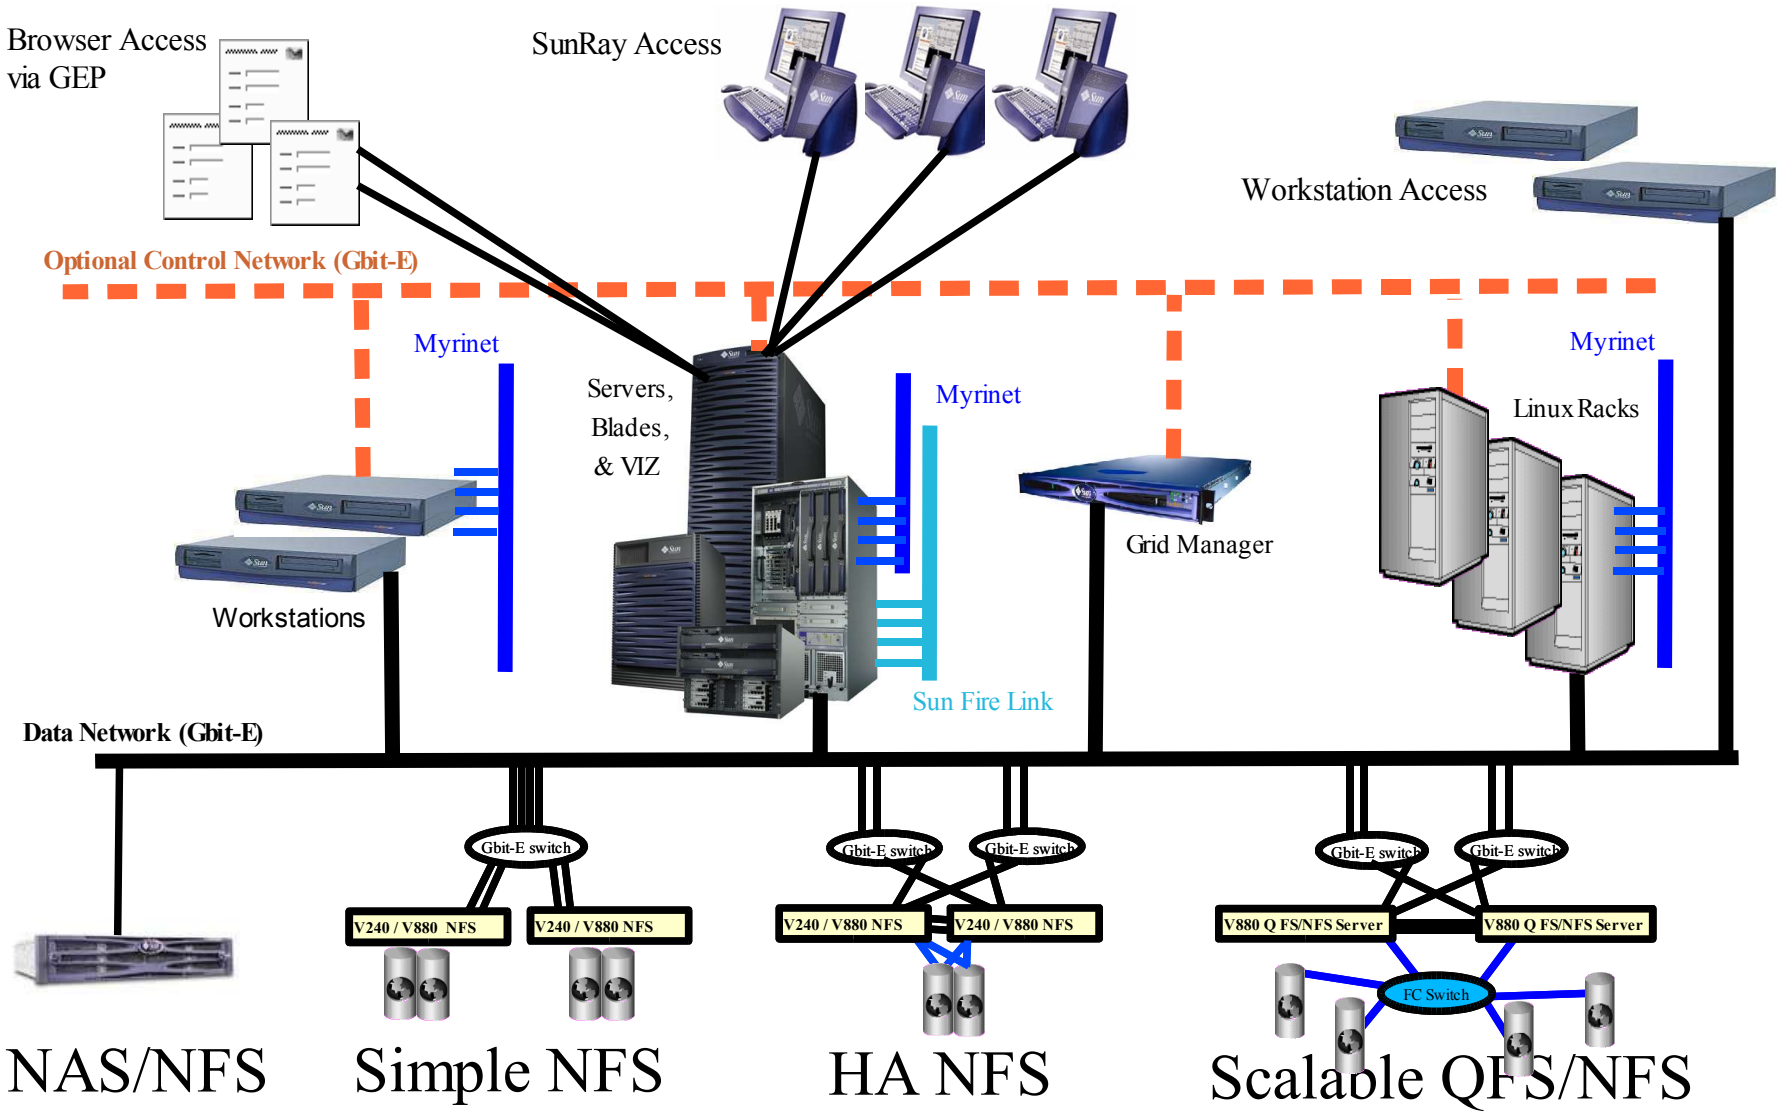

Sun Global Grid Partner Projects

- **ICENI**, Imperial College e-Science Netw. Infrastructure, London
- **GRIDS**, Grid Computing & Distributed Systems Lab, Melbourne
- **EZ-Grid**, Sun Center of Excellence for Grid Computing, Houston
- **White Rose Grid**, Universities of Leads, Sheffield, York, UK
- **NCSV**, Nanyang Center for Supercomp.& Visualization, Singapore
- **EPCC** Edinburgh Sun Data & Compute Grid Project
- **HPCVL** Canada, Secure innovative HPC/Grid environment
- **GridLab** European Project for Grid Application Infrastructure
- **myGrid** Infrastructure for an e-Biologist Workbench, Manchester
- **OSC Grid**, Sun Center of Excellence for BioInformatics, Ohio
- **AIST** Advanced Industrial Science & Technology Institute, Tokyo
- and more...

THE UNIVERSITIES OF LEEDS, SHEFFIELD AND YORK

THE WHITE ROSE GRID

Welcome

Sun microsystems

Global Grid Example: White Rose in England

- Leeds, York + Sheffield Universities
- Deliver stable, well-managed HPC resources supporting multi-disiplinary research
- Deliver a Metropolitan Grid across the Universities
- Globus, SGEEE, GPDK, MyProxy

Other Examples: AIST Japan, Nanyang Singapore, Houston, HPCVL, Poznan, GlobeXplorer, Imperial, ...

- Slides courtesy of Tom Jackson, DAME project manager, York

UNIVERSITY OF LEEDS

UNIVERSITY OF SHEFFIELD

UNIVERSITY OF YORK

- The White Rose Grid - Grid Infrastructure developed between the Universities of York, Sheffield and Leeds. Investment of £3.0m, linked to YHMAN (£1.3m)
- Production architecture is currently based on Globus Toolkit 2.4.3 with mix of Grid Services architecture from GT3.0. Test-grid runs in parallel with full GT3.0 implementation.
- Sun Grid Engine Enterprise Edition for Dynamic Resource Management
- Access portals are developed using jet stream/struts technology.

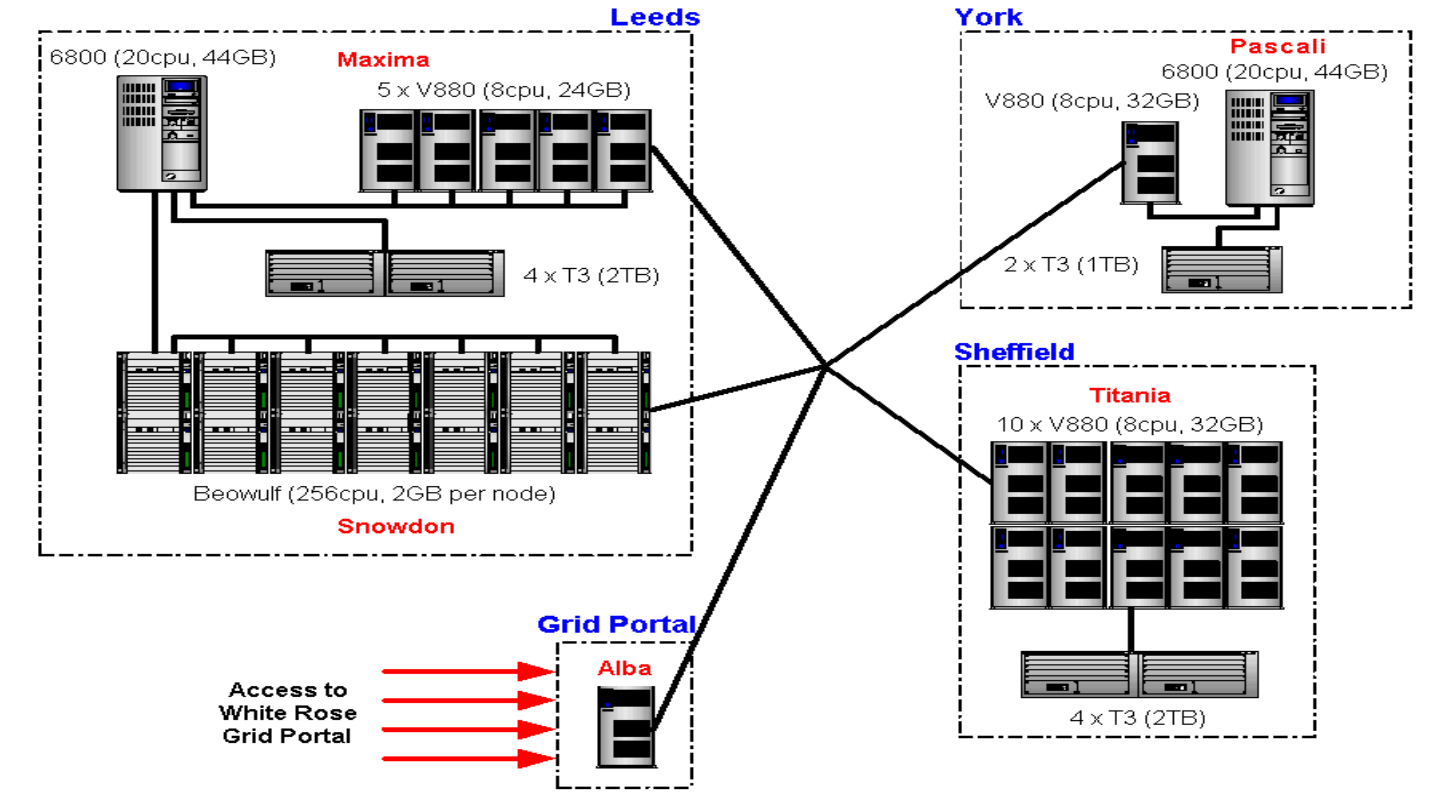
Leeds
6800 (20cpu, 44GB)
Maxima
5 x V880 (8cpu, 24GB)
4 x T3 (2TB)
Beowulf (256cpu, 2GB per node)
Snowdon
York
Pascali
6800 (20cpu, 44GB)
V880 (8cpu, 32GB)
2 x T3 (1TB)
Sheffield
Titania
10 x V880 (8cpu, 32GB)
4 x T3 (2TB)
Grid Portal
Alba
Access to
White Rose
Grid Portal

WRG Key Components

Globus Toolkit 2.0

Secure means for inter-campus actions

Grid Engine Enterprise Edition

Manages campus grid compute resources

Grid Portal Development Kit

Provides a portal interface into Globus Toolkit

MyProxy

Store and retrieve delegated X.509 credentials via the Grid Security Infrastructure (GSI)

Sun Grid Engine Enterprise Edition

Share Policies

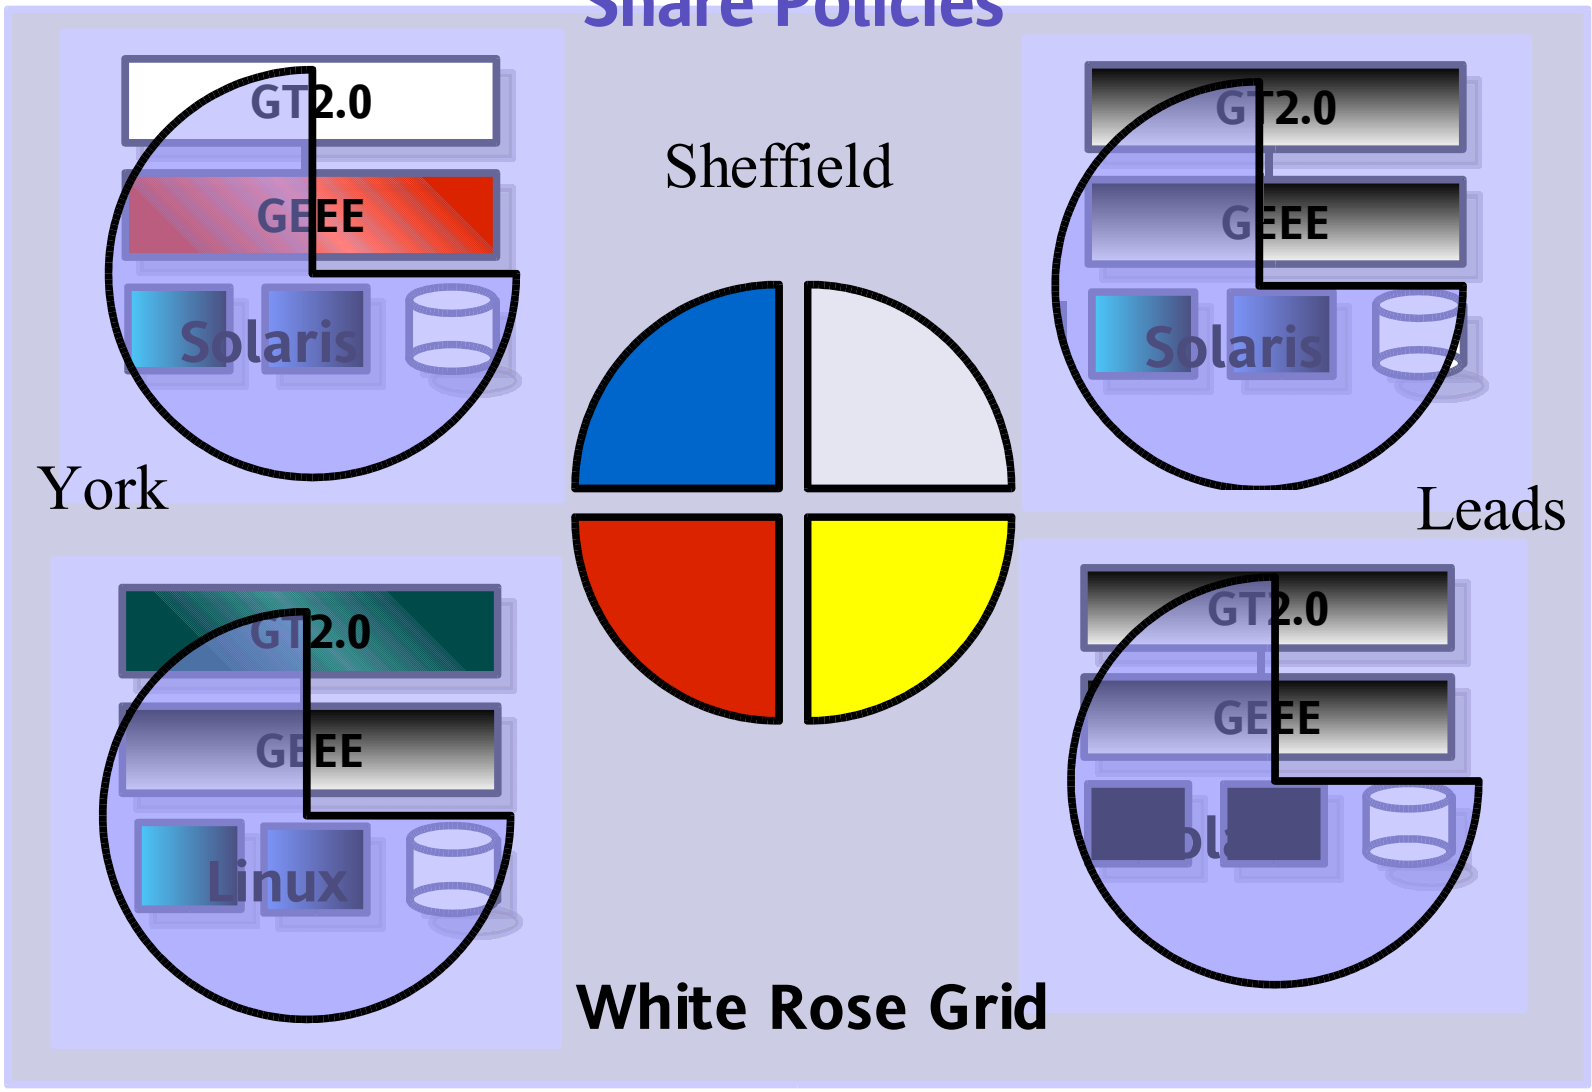

- Universities
  - Departmental use
    - Biology, Chemistry, Computing; Environment; Food Science; Maths; Mechanical Engineering; Physics and Materials Engineering - most fall into category of HPC users.
  - Collaborative Research and Development
    - Computation fluid dynamics; bone modelling; skin disease; digital preservation; elastohydrodynamic lubrication; visualisation.

- Commercial users
  - Research Labs
    - Central Science Laboratory
  - Small Medium Enterprises
    - Glass manufacture; digital media enterprise; software developers; IT solutions providers for libraries; computational fluid dynamics.
  - Large Enterprises
    - Aircraft maintenance; agrichemical; automotive; tool cutters; oil industry.

- Experiences
  - Implementations through several versions of Globus Toolkit have caused regular problems;
  - software portability – local variations in host environment mean that portability of code is not guaranteed;
  - Firewalls and security policies have been a barrier to progress;
  - End users do not want to interact with GT, hence a lot of effort required to build access portals for application hosting;
  - Scalability – unwieldy process for addition of new users and very limited scope to dynamically add new compute resource;
  - Software licensing issues are complex.

- Grid Wish List
  - End Users need far better support. GT too low level, and require desktop support for managing Grid applications;
  - Improved methods for handling heterogeneity of Grid infrastructures; to facilitate portability of code/applications;
  - Rapid development of security mechanisms and policies for Grid infrastructures.

GridLab and Progress Testbeds Working Together

Slides courtesy of Jarek Nabrzyski, PSNC

GridLab vs. Progress Architectures

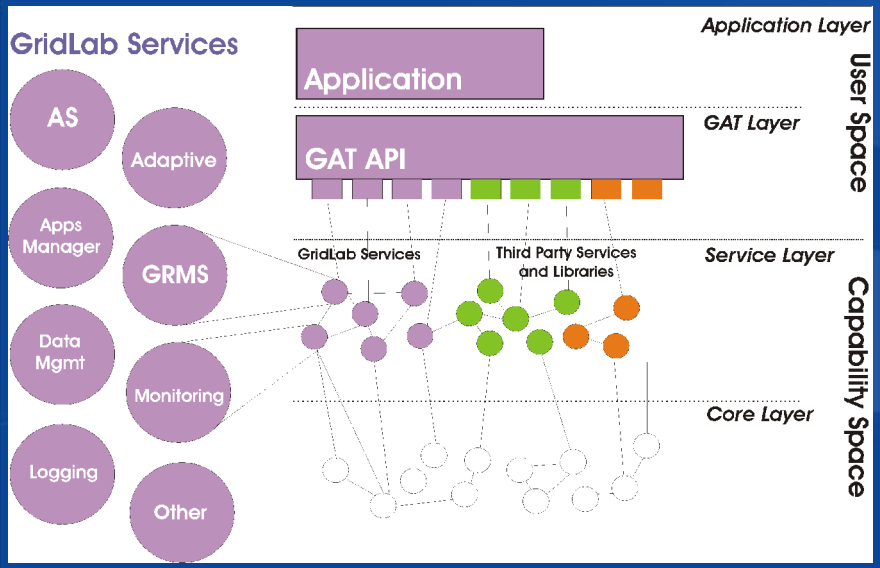

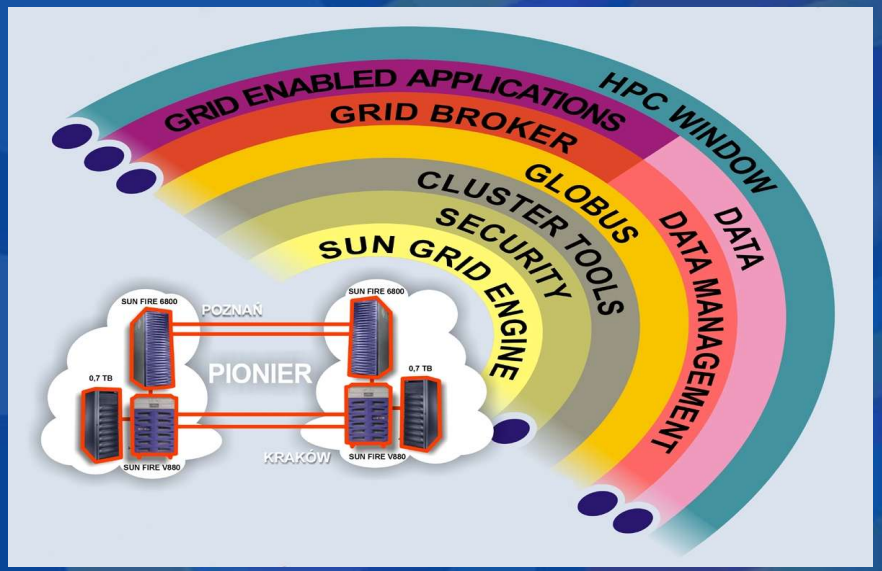

- Heterogeneous
- Europe-wide
- From clusters to supercomputers
- Globus 2.2.4

- homogeneous
- Poland-wide
- Clusters
- Globus 2.2.4

Applications

- GridLab
  1. Cactus & Triana
  2. Dynamic scenarios (migration, dynamic staging, look ahead, spawn off independent tasks)
  3. Enabling applications on the grid through Grid Application Toolkit
- Progress
  1. Bioinformatics
  2. Other domains coming
  3. Workflow jobs
  4. Enabling applications on the grid through Application Management Services

Moving towards Web Services World

- GridLab
  1. currently all the services are Web Services based
  2. current focus is on the GAT
  3. OGSA will be adressed in the final year of GridLab
  4. GAT allows to use whatever services
- Progress
  - Currently all the services are Web Services based
  - OGSA is coming in 2004
  - Grid Service Provider allows to use Web Services, OGSA interface is to be added

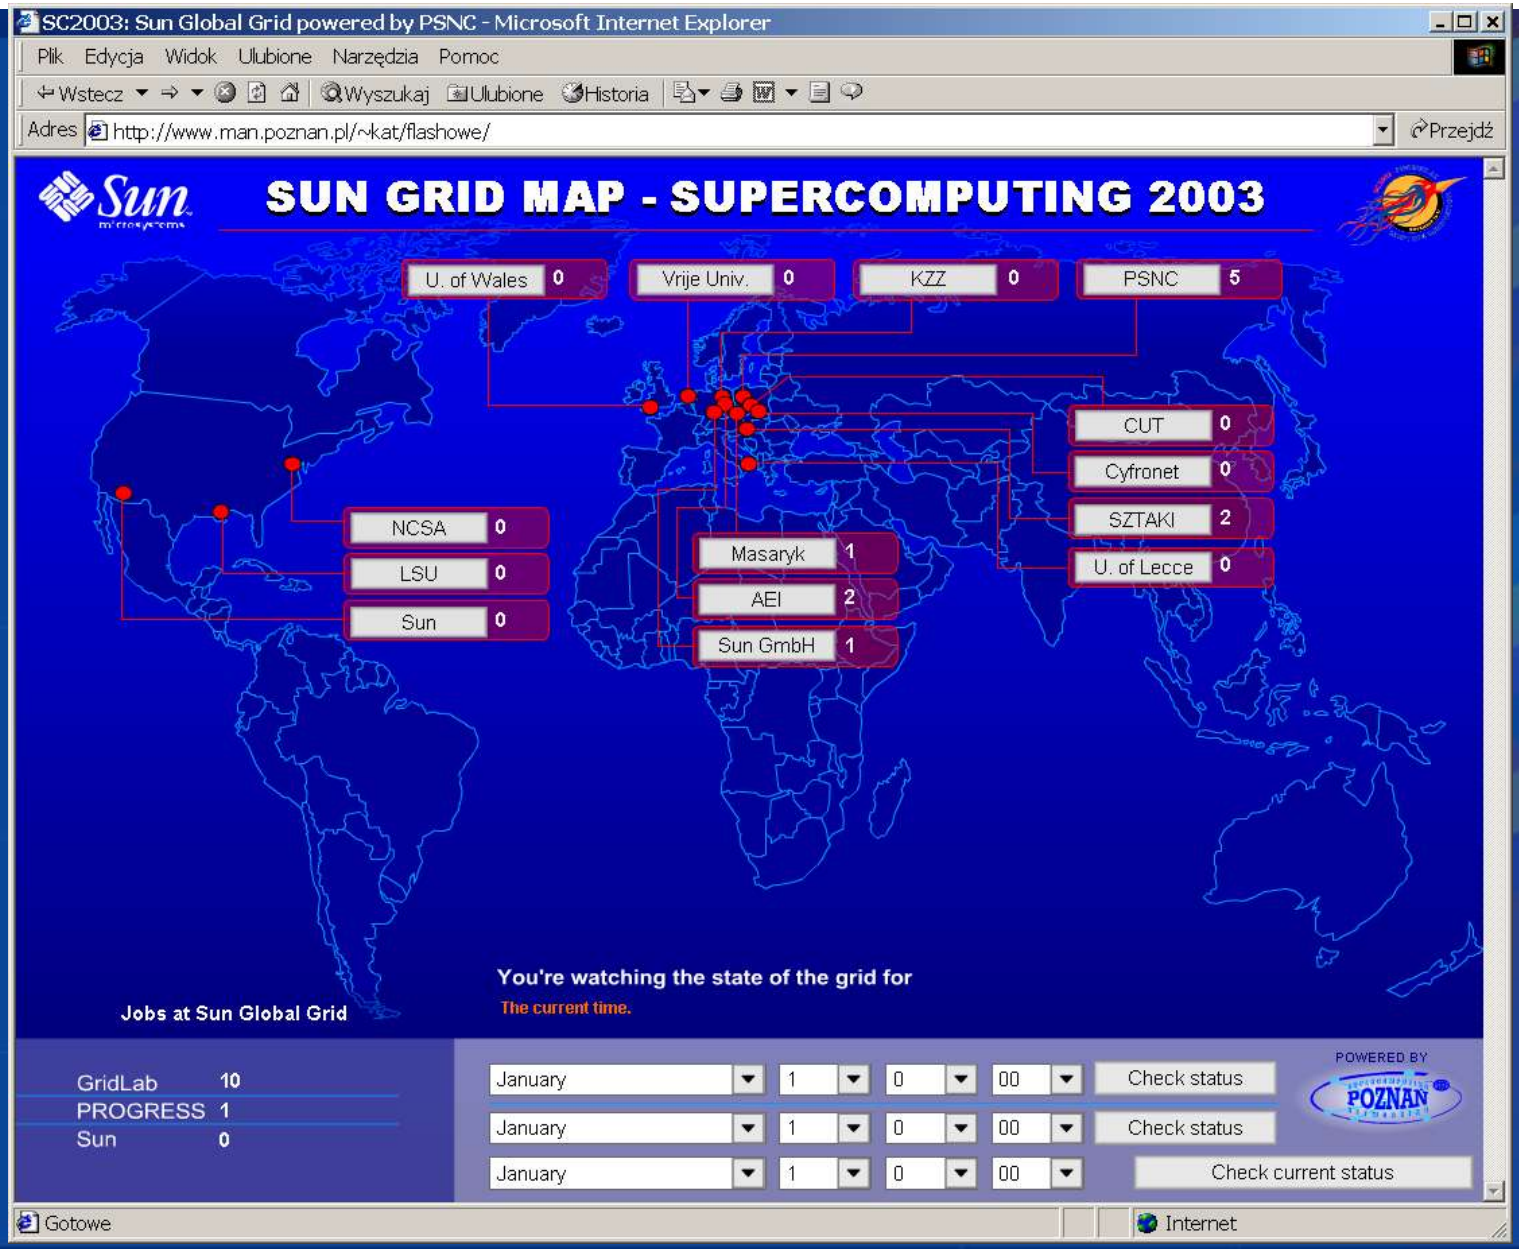
SC2003: Sun Global Grid powered by PSNC - Microsoft Internet Explorer
Plik Edycja Widok Ulubione Narzędzia Pomoc
Wstecz Wyszukaj Ulubione Historia
Adres http://www.man.poznan.pl/~kat/flashowe/
Przejdź
SUN GRID MAP - SUPERCOMPUTING 2003
U. of Wales 0
Vrije Univ. 0
KZZ 0
PSNC 5
CUT 0
Cyfronet 0
SZTAKI 2
U. of Lecce 0
NCSA 0
LSU 0
Sun 0
Masaryk 1
AEI 2
Sun GmbH 1
You're watching the state of the grid for
The current time.
Jobs at Sun Global Grid
GridLab 10
PROGRESS 1
Sun 0
January 1 0 00 Check status
January 1 0 00 Check status
January 1 0 00 Check current status
POWERED BY
POZNAN
Gotowe
Internet

How might OGSI/WS-agreement affect the next steps?

- It would allow to build a real global grid, with negotiations on all levels (costs, dynamic users, scheduling criteria, user preferences, resource usage policies etc.)
- Would allow to have even more decentralized architecture, with more autonomies for local centers
- Security is still an issue and is still behind.../ underestimated

Grid wish list, what next, what needs to be done?

- Security
- Grid user accounting
- Application support technologies (Application Toolkits, grid programming environments)
- Knowledge-based grids
- Fault tolerance, performance, production environments

Globus Grids in Production

Slides courtesy of Steven Newhouse, Technical Director

London e-Science Centre

London e-Science Centre

'Enabling the e-Scientist'

- Established applied multi-disciplinary research
- Industrial Collaborations:
  - Sun Centre of Excellence in e-Science
  - Intel Virtual European Centre of Grid Computing
- Cross-campus collaborations:
  - Bioinformatics
  - High Energy Physics
  - Computational Engineering
- Specialisation: Next Generation Grid Middleware
- http://www.lesc.imperial.ac.uk/

LeSC Grid Resources

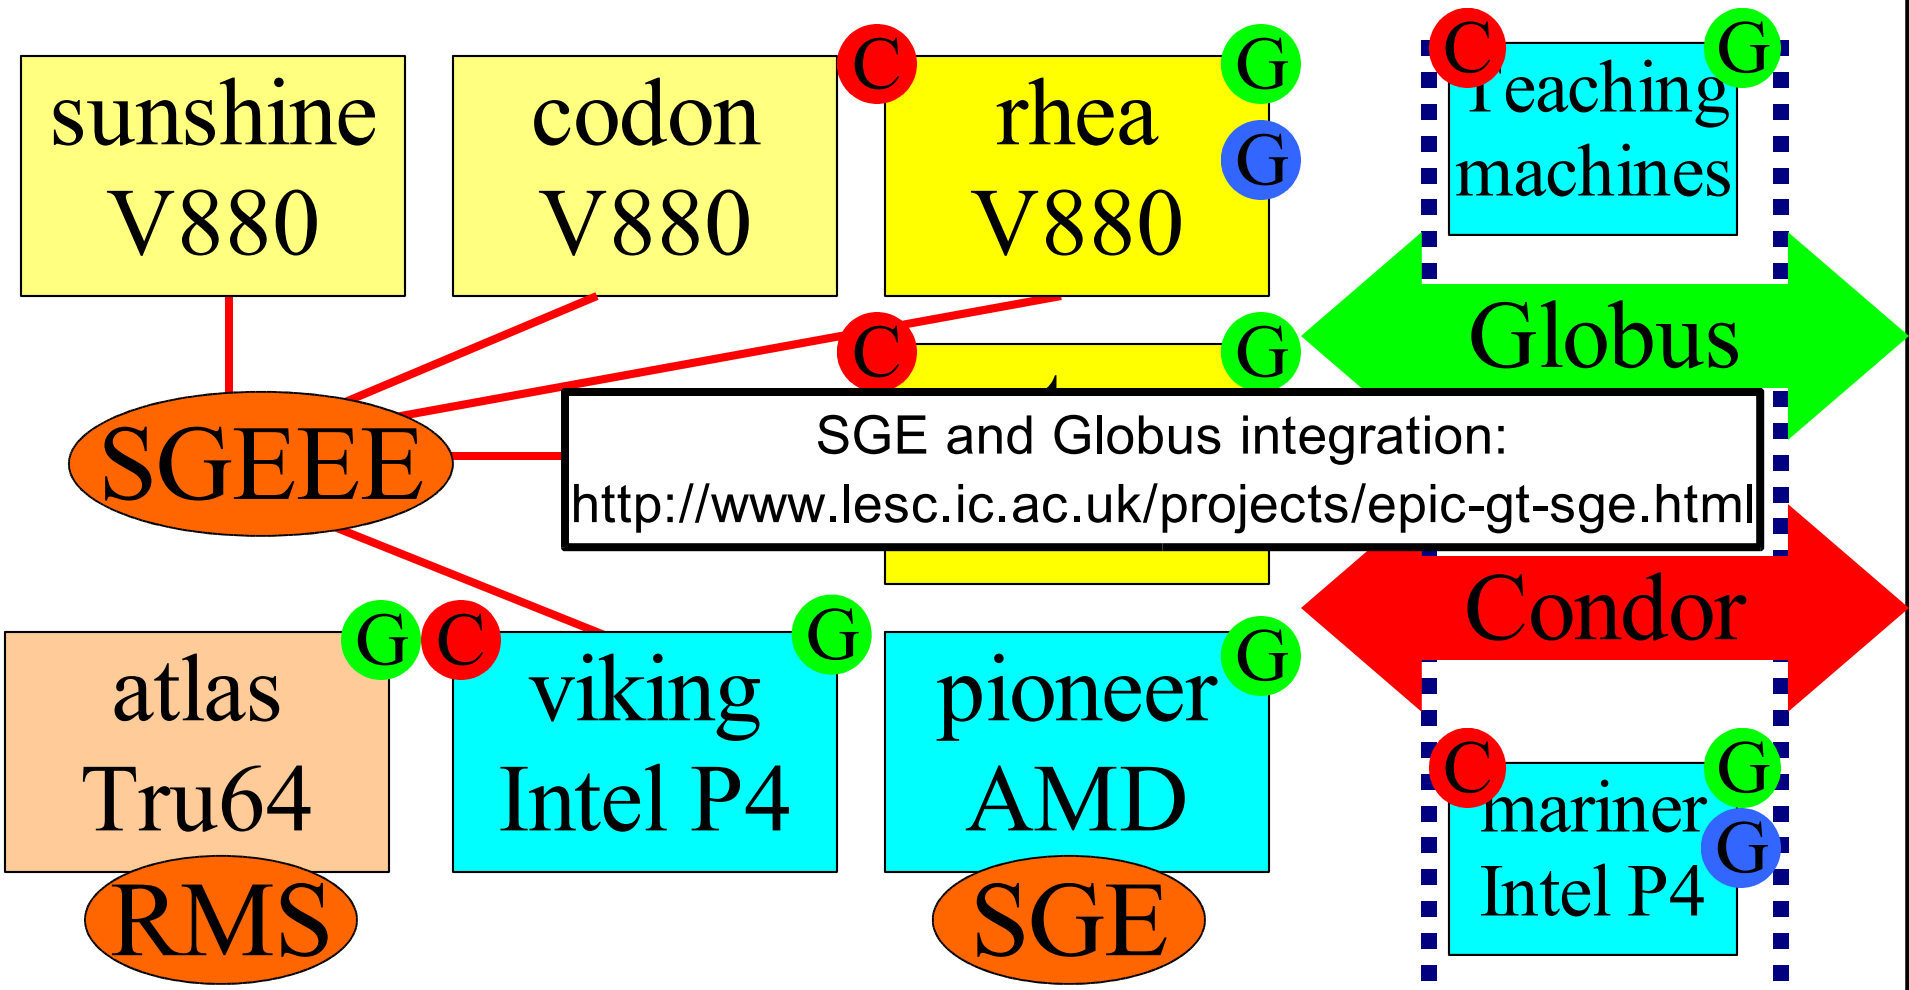

The Problem: Thermohaline circulation

- Ocean transports heat through the "global conveyor belt."
- Heat transport controls global climate.
- Wish to investigate strength of model ocean circulation as a function of two external parameters.
- Use GENIE-Trainer.

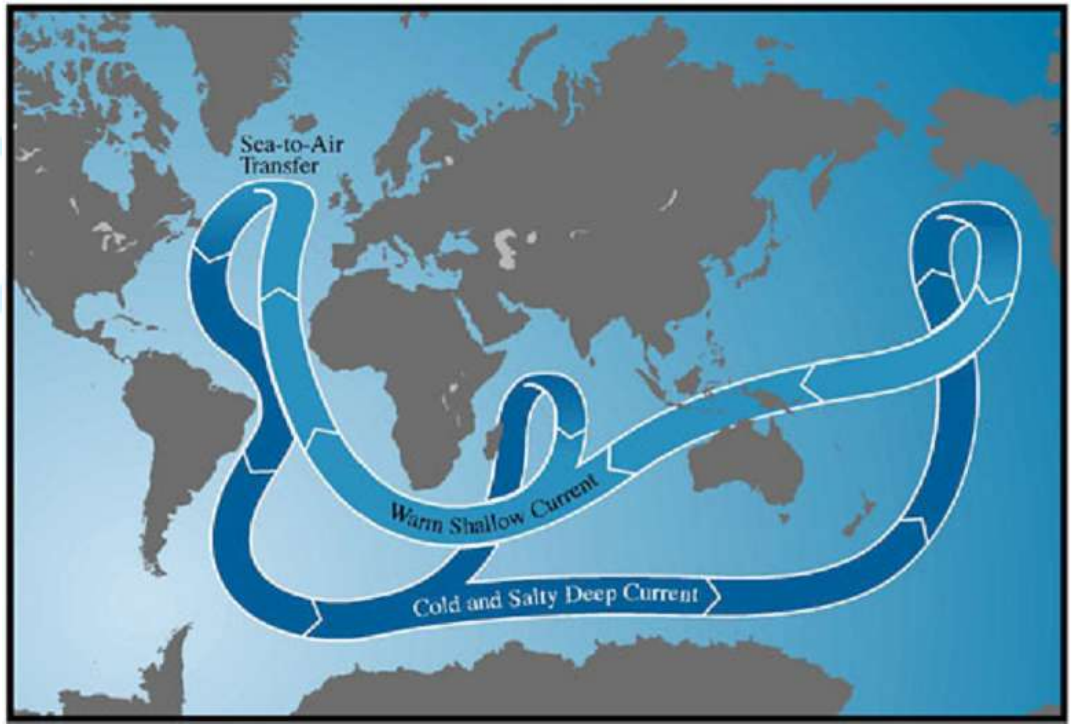

- Wish to perform 31×31 = 961 individual simulations.
- Each simulation takes ~4 hours to execute on typical Intel P3/1GHz, 256MB RAM, machine ⇒

time taken for 961 sequential runs ≈ 163 days!!!

The Results: Scientific Achievements

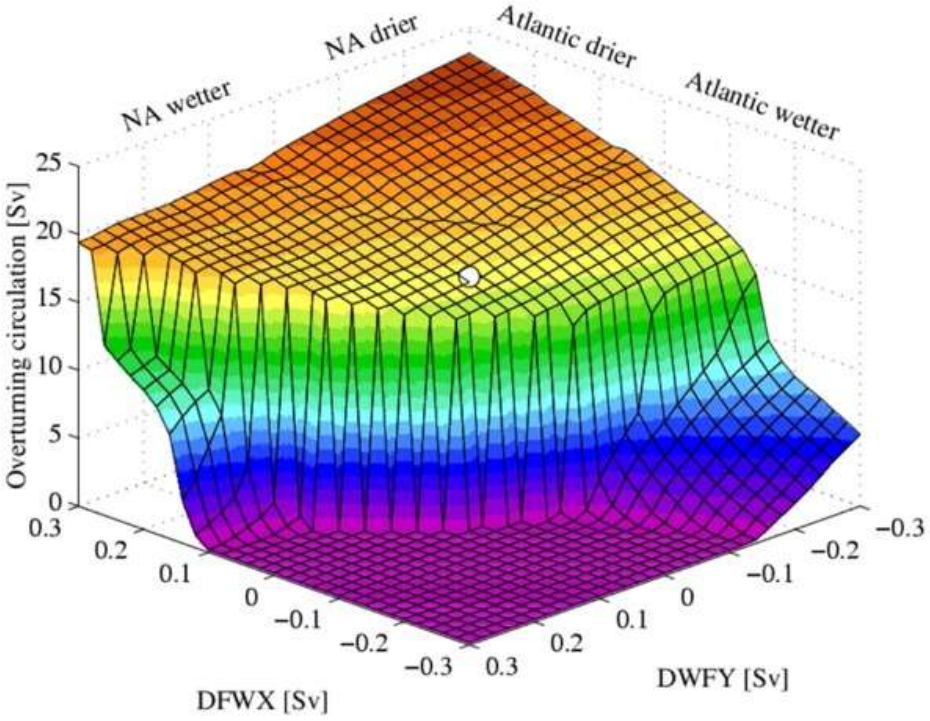

Intensity of the thermohaline circulation as a function of freshwater flux between Atlantic and Pacific oceans (DFWX), and mid-Atlantic and North Atlantic (DFWY).

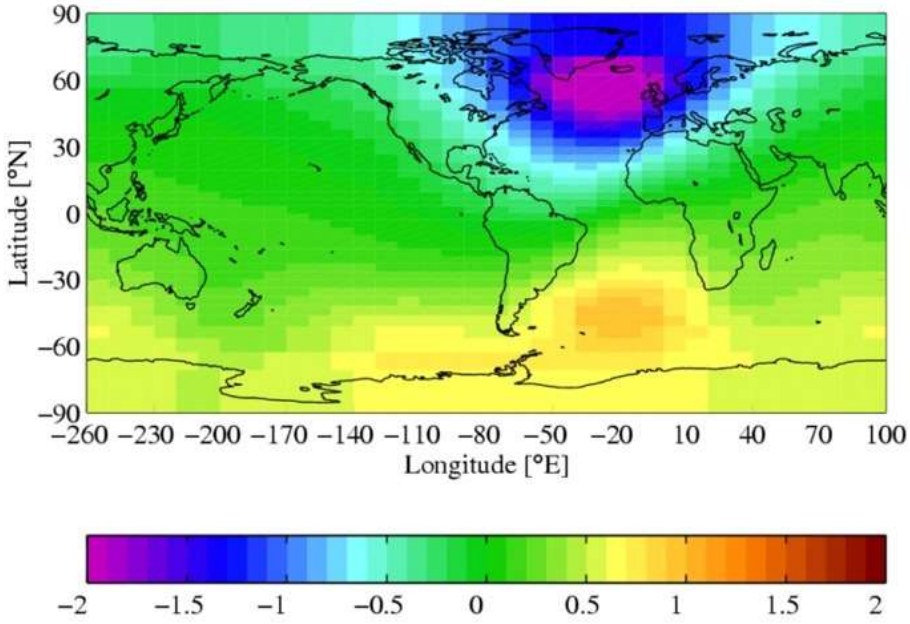

Surface air temperature difference between extreme states (off - on) of the thermohaline circulation.

North Atlantic 2°C colder when the circulation is off.

time taken for 961 runs over ~200 machines ≈ 3 days

ICENI:
IC e-Science Networked Infrastructure

The Iceni, under Queen Boudicca, united the tribes of South-East England in a revolt against the occupying Roman forces in AD60.

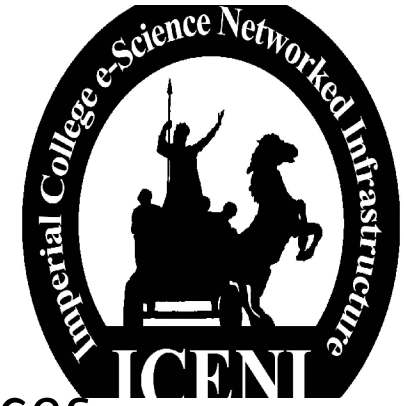

- Developed by LeSC Grid Middleware Group
- Use to define and develop higher-level services
- Collect and provide relevant Grid meta-data
- Interaction with other frameworks: OGSI, Jxta etc.
- Available under extended open source SISSL:
  - http://www.lesc.ic.ac.uk/iceni

The Future

- Highly stable production quality fabric
  - Intel+SPARC/Linux+Solaris/Sun Grid Engine
  - SGE and GT2 & GT3 integration

    http://www.lesc.ic.ac.uk/projects/epic-gt-sge.html
- Highly unstable grid layer
  - Is GT3 just a reference implementation?
  - What focus should be placed on production quality?
- Link WSA to elements of the Grid fabric
  - SGE 6.0

Thank You !

The New Time Machine

The Grid Engine

The Combustion Engine

The Steam Engine

wolfgang.gentzsch@sun.com
http://www.sun.com/grid