

Introduction to Grid Computing

Balázs Kónya

(balazs.konya@quark.lu.se)

What is the Grid?


- Grid is a **technology** to **share** and **seamlessly access resources** of the world:
 - computing cycles
 - datasets,
 - software,
 - special instruments
- the Holy Grail of **distributed** computing
- **Middleware**: a bag of software which implements Grid Standards & protocols
- World Wide Web access to information

- **World Wide Grid**  access to computing capacity and ...



What is the Grid?


- The future **infrastructure** of computing and data management
- The **Computational Power Grid**
- a very ambitious attempt to create a new **utility**, next to the already existing water, heating, electricity, ...
- the present hype in IT



source: IBM


History

- Grand Scientific Challenges of the 80's
 - parallel computation
 - high performance & high throughput computing
- Early "Testbeds" in the USA connected supercomputing centers at the late 90's
- Ian Foster, Carl Kesselman, July, 1998:
Blueprint for a new Computing Infrastructure



History cont.


- The Computational Grid <-> Power Grid analogy was suggested
- The birth of the "ancient" middleware solutions
 - Globus, Legion, Condor, NWS, SRB, NetSolve, AppLes, Unicore
 - "demonstration quality Testbeds: Gusto" no real users
 - loose of interest in Grids
- **2000+: The Grid revives and gets "Global"**
 - The High Energy Physics community picks up the nearly forgotten Grid idea
 - The appearance of the Global Grid Forum
 - de facto standard middleware: **Globus**
- the "Grid phenomena" or hype is started
 - Grid Projects are launched everywhere, governments & research agencies rush to support Grid project



History cont.

- Huge commercial interest: startup companies & the Big Names try to sell the Grid
 - IBM wants to Grid-enable the company's entire product portfolio
 - commercial Grid software (IBM, Platform Computing, SUN)
 - commercial support, consulting, training
 - serious research projects (mainly biology) among the customers
 - last Global Grid Forum meeting in Edinburgh July, 2002:
 - over 850 participants
 - Key speakers involved: IBM, Nec, Hewlet Packard, Microsoft, SUN
 - Daily Grid magazines: www.thegridreport.com, www.gridtoday.com, www.gridcomputingplanet.com
- GRID** today *The Grid Report*
the state of grid computing
- Everything is Grid, the word "Grid" is inflated to a marketing term
 - the divergence of Grid Toolkits and solutions

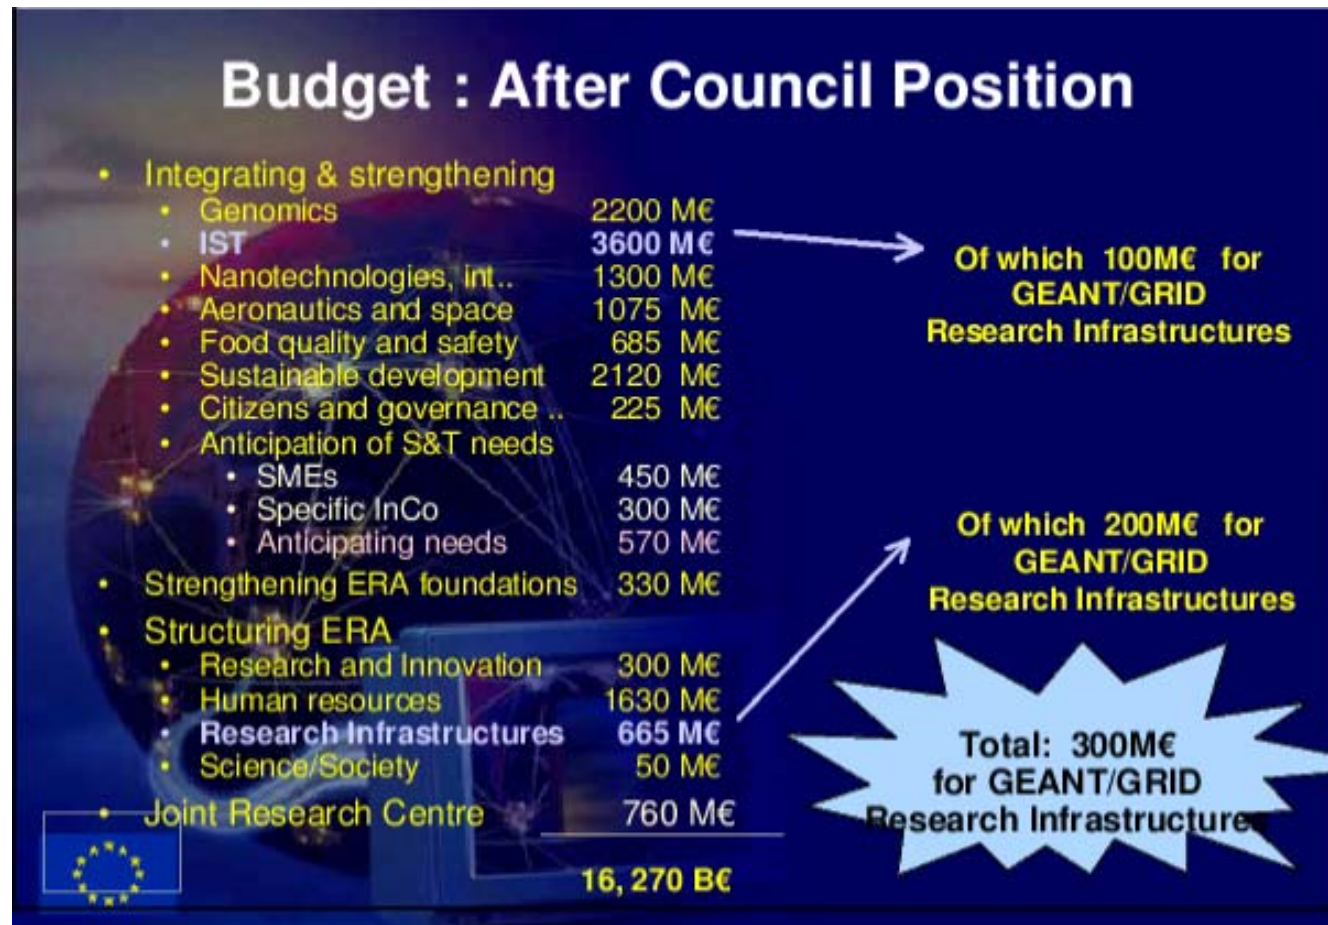
European projects




Current EU founded projects

GRIDSTART





USA projects

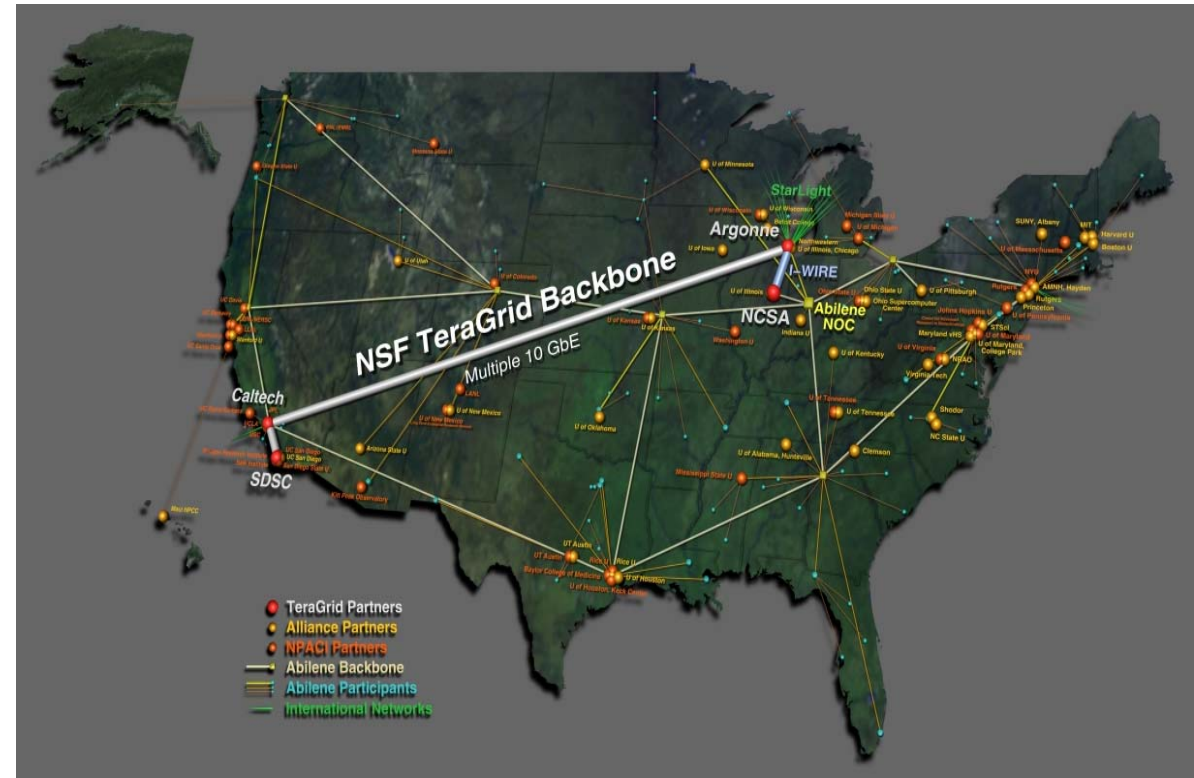


DISCOM
SinRG
APGrid
IPG ...

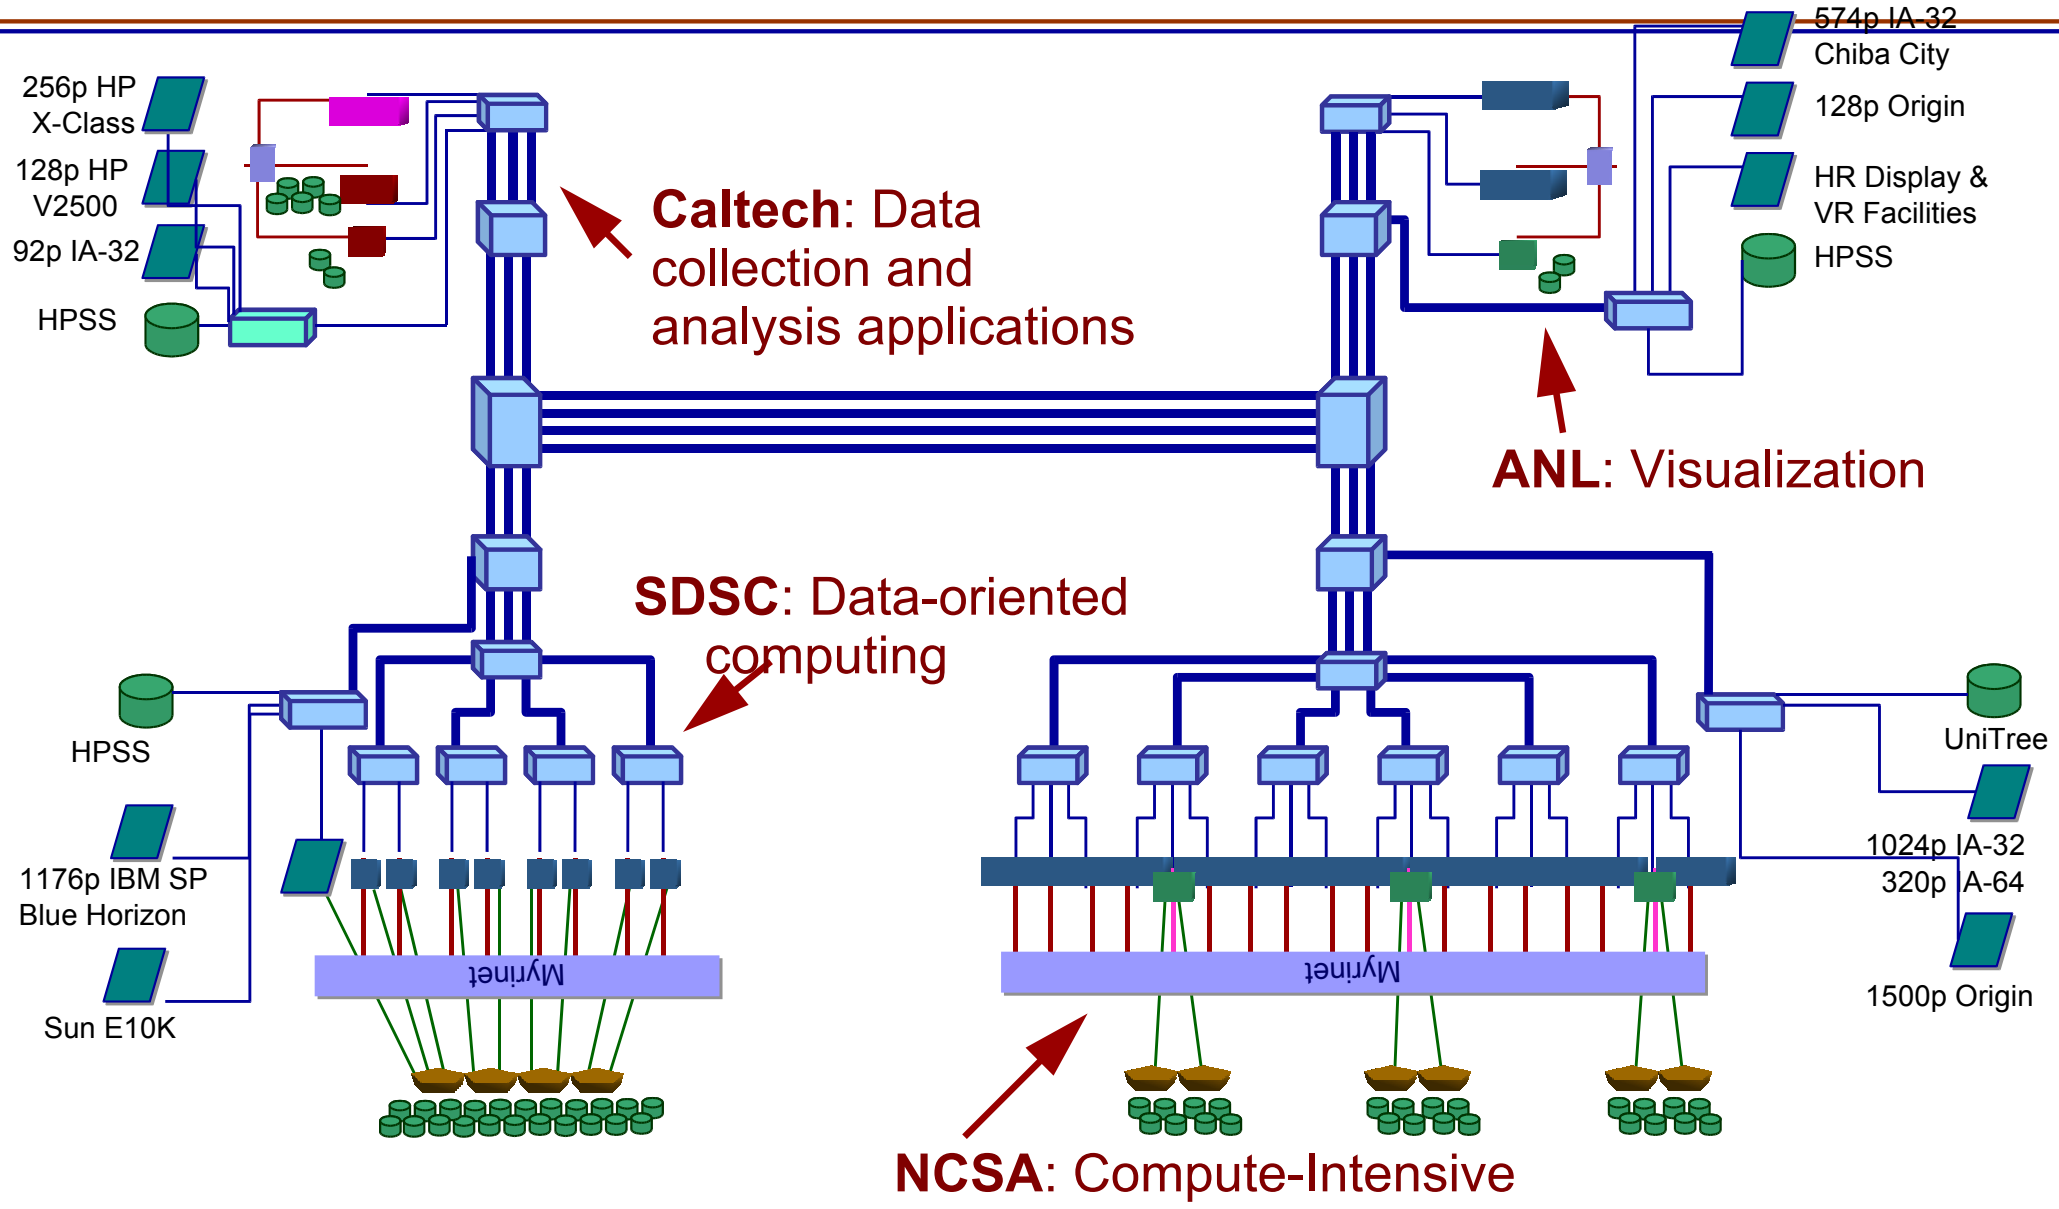


TeraGrid

- ➔ 53 million from the NSF
- ➔ 13.6 teraflops of Linux clusters
- ➔ 450 terabytes of data storage
- ➔ 4 sites
- ➔ 40 Gbits/sec (later 50-80) network connections
- ➔ Globus based Grid toolkits
- ➔ Visualisation environment



TeraGrid



Asia Pacific Projects



Japan: AIST GTRC
China: SDG
Korean Grid
Thailand: ThaiGrid
Australia: GRIDSLab



Grid in the NEWS

BusinessWeek online

BW MAGAZINE DAILY BRIEFING INVESTING GLOBAL BUSINESS TECHNOLOGY SMALL BUSINESS B-SCHOOL

JUNE 3, 2002

BW MAGAZINE
U.S. EDITION
[Full Table of Contents](#)
[Cover Story](#)
[Up Front](#)
[Readers Report](#)

INFORMATION TECHNOLOGY

Who Needs Supercomputers?

Grid software lets companies tap other machines in their network [gridcomputingplanet.com](#) : [News](#)

GRID COMPUTING PLANET.COM EARTHWEB®
Management Networking & Communications Web Development Hardware & Systems Software Development
internet.com™

FREE TECH NEWSLETTERS YOUR EMAIL [Sign Up](#)
Grid Computing Planet Text

PHYSICS TODAY.org

Gradschool

Shred Engine
STRATEGIES 2002
TRENDS & EXPO

Researchers Achieve Production Grid Breakthrough
By [Paul Shread](#)

Search

Welcome! [logoff](#)

FEATURE ARTICLES PHYSICS UPDATE LETTERS SEARCH & DISCOVERY ISSUES & EVENTS BOOKS WE HEAR THAT OBITUARIES

TABLE OF CONTENTS

PAST CONTENTS

LINKS TO PHYSICS TODAY ADVERTISERS

PLACE AN AD

BUYER'S GUIDE

ABOUT US

CONTACT US

Feature Article

The Grid: A New Infrastructure for 21st Century Science

As computer networks
of science and engineering

[Ian Foster](#)

White Papers IT Jobs Events Research Premium Services Training & Certification Media Kit
GRID COMPUTING PLANET.COM EARTHWEB®
IT Management Networking & Communications Web Development Hardware & Systems Software Development
internet.com™

[www.gridcomputingplanet.com](#) : [News](#)

FREE TECH NEWSLETTERS YOUR EMAIL [Sign Up](#)
Grid Computing Planet Text



[News Archives](#)

TeraGrid Receives \$35 Million From National Science Foundation
By [Paul Shread](#)

Grid in the NEWS

ZDNet Tech Update
Business & Technology Today

Home Applications Networking Security eBusiness Hardware Platforms IT Products

Bridges/Routers | Hubs/Switches | PCs/Notebooks | Servers | Storage | more...

Search: Tech Update

Grid computing rates a 6.5 on the hype meter

By **Dan Farber**
May 15, 2002

UNPLUGGED: WLADAWKSY-BERGER
SPECIAL REPORT
READ MORE

TalkBack
E-mail this
Print this

IN THIS STORY

1. Grid computing rates a 6.5 on the hype meter
2. Grid projects underway

SPECIAL COVERAGE

- ▶ What to do with Tablet PCs?
- ▶ Gartner's guide to strategic outsourcing
- ▶ Hot spots for wireless ISPs
- ▶ United States of Linux
- ▶ TechXNY unplugged
- ▶ Bluetooth bandwagon ready to roll
- ▶ Face to face with Matthew Szulik

Grid computing has been hailed by some as the new infrastructure for the 21st century, or as the next major computing revolution. Major companies including IBM,

NetworkMagazine.com

Find all of your Networking and IT needs on eBay



[click here](#)

Visit these other CommWeb channels

- [Buyer's Guide](#)
- [Lab Tests](#)
- [Case Studies](#)
- [Events](#)
- [Product Reviews](#)
- [Tutorials](#)
- [Roundtables](#)
- [Subscriptions](#)

Insanely Overhyped Technology of the Month- Grid Computing

Distributing executable processing capacity across thousands or millions of CPUs isn't the right approach for most of the problems that computers solve.

by Steve Steinke, Editor-in-Chief

Network Magazine
05/06/02, 12:52 p.m. ET

Utilities

BUILDING GRIDS: HYPE MEETS REALITY

The senior executives from Compaq (now part of the new HP), Platform Computing and renowned analysts from IDC invite you to attend an Exclusive Executive event on Grid Computing.

Vision...

- Cohen Communication Group:
 - Grid computing will be the driving force behind the 150% annual internet traffic expansion in 2005
 - in contrast to the 60% predicted growth rate accounted mainly for video streaming and video file transfer forecasted by McKinsey - JP Morgan
- IBM:
 - Grid is the next utility in the line of the water, drainage, gas and electricity systems
 - people will pay their “computing bills”

Oversold?

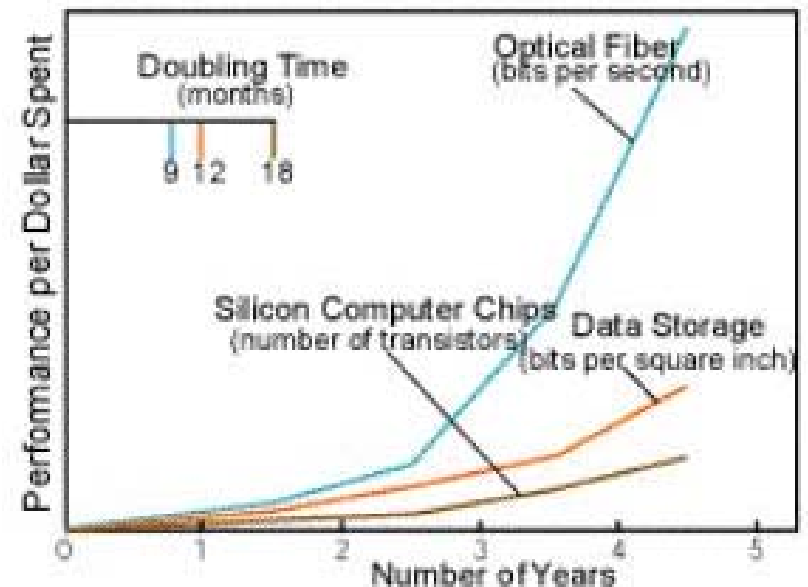
- The promise of the **Grid** has been **not oversold** but the **difficulty of developing** the necessary Grid infrastructure has been **underestimated**

Ian Foster:

- People used to **overestimate** the short term impact of change but **underestimate** the long-term effect


what is behind?

- Powerful PCs are everywhere
- Clusters are commodity
- Network & Storage & Computing **exponentials**:
 - Networking speed grows faster than computing power
 - Even data storage outperforms the CPUs



source: Scientific American, Jan 2001

The physicist's real challenge:



there are already ...

Walmart Inventory Control

- ➔ Satellite technology used to track every item
- ➔ Bar code information sent to remote data
- ➔ Inventory adjusted in real time to avoid shortages and predict demand
- ➔ Data management, prediction, real-time, wide-area synchronization



SETI@HOME

- ➔ 3.8M users in 226 countries
- ➔ 1200 CPU years/day
- ➔ 1.7 ZETAflop over last 3 years (10^{21})
- ➔ 38 TF sustained performance
(Japanese Earth Simulator is 40 TF peak)
- ➔ Highly heterogeneous: >77 different processor types



... distributed applications

Everquest

- ➔ 45 communal “world servers” (26 high-end PCs per server) supporting 430,000 players
- ➔ Real-time interaction, individualized database management, back channel communication between players



Napster, Gnutella, Kazaa, etc...

- ➔ file sharing
- ➔ ask the music industry :)



Google

- ➔ database, search engine
- ➔ more than 150 million searches per day, 2 billion indexed pages, more than 10.000 linux servers



there should be a Grid ...

- Existing real world examples demonstrate that it is **technically, commercially, and economically viable** to deploy robust, large-scale distributed applications
- The Grid
 - will extend those distributed applications
 - should accelerate the progress of distributed applications
 - will use common interfaces
 - will be based upon well-defined protocols & standards
 - will offer scalable Grid services for applications

where we are now?

- lots of theoretical papers
 - *The anatomy of the Grid:Enabling Scalable virtual organizations, I.Foster et.al.*
 - *The Physiology of the Grid: An Open Grid Services Architecture for Distributed System Integration, I.Foster,C.Kesselman, et. al.*
 - *The patology of the Grids, ???*
- non-existing (very few) TestBeds:
 - they are incompatible
 - Difficult to get access to them
 - very expensive to maintain
- non existing standards (GGF has not produced anything yet)
- “de facto standard” middleware is rather limited in functionality
- diverging solutions, huge amount of (overlapping) projects
- we are living in the Grid hype era

not even (hardly) addressed:

- political issues
- heterogeneity
- Grid-based authorization
- Grid scheduling
- Program development environments
- Debugging, compiling, performance tuning
- Fault tolerance
- Modeling of dynamic, unpredictable environments
- Grid market economy (allocation, accounting, cost models)

Definition

- **Ian Foster** (www.gridtoday.com/02/0722/100136.html):
 - coordinates resources that are not subject to a centralized control
 - using standard, open, general-purpose protocols and interfaces
 - delivers nontrivial qualities of service
- **Rajkumar Buyya:**
 - A type of parallel and distributed system that enables the sharing, selection, & aggregation of resources distributed in administrative domains depending on their availability, capability, performance, cost, and users quality of service requirements.

Simple Model of the Grid



Data Management



Resource & Job
Management


+ security




Information System

another model (basic elements)

Security



Computational Economy




Uniform Access



System Management


Resource Discovery




Resource Allocation
& Scheduling



Data locality



Application Development Tools



Network Management

source:Rajkumar Buyya

Lund, 24/01/2003

The layers of the Grid:

Grid Applications

science, engineering, commercial apps, web portals

Grid Programming environment

languages, interfaces, libraries, compilers, griddifying tools

User level Middleware

resource management and scheduling services

Low level Middleware

job submission, storage access, info service, accounting

Security Infrastructure


single log-on, authentication, authorization, secure communication

Grid Fabric

clusters, networks, batch systems, devices, databases

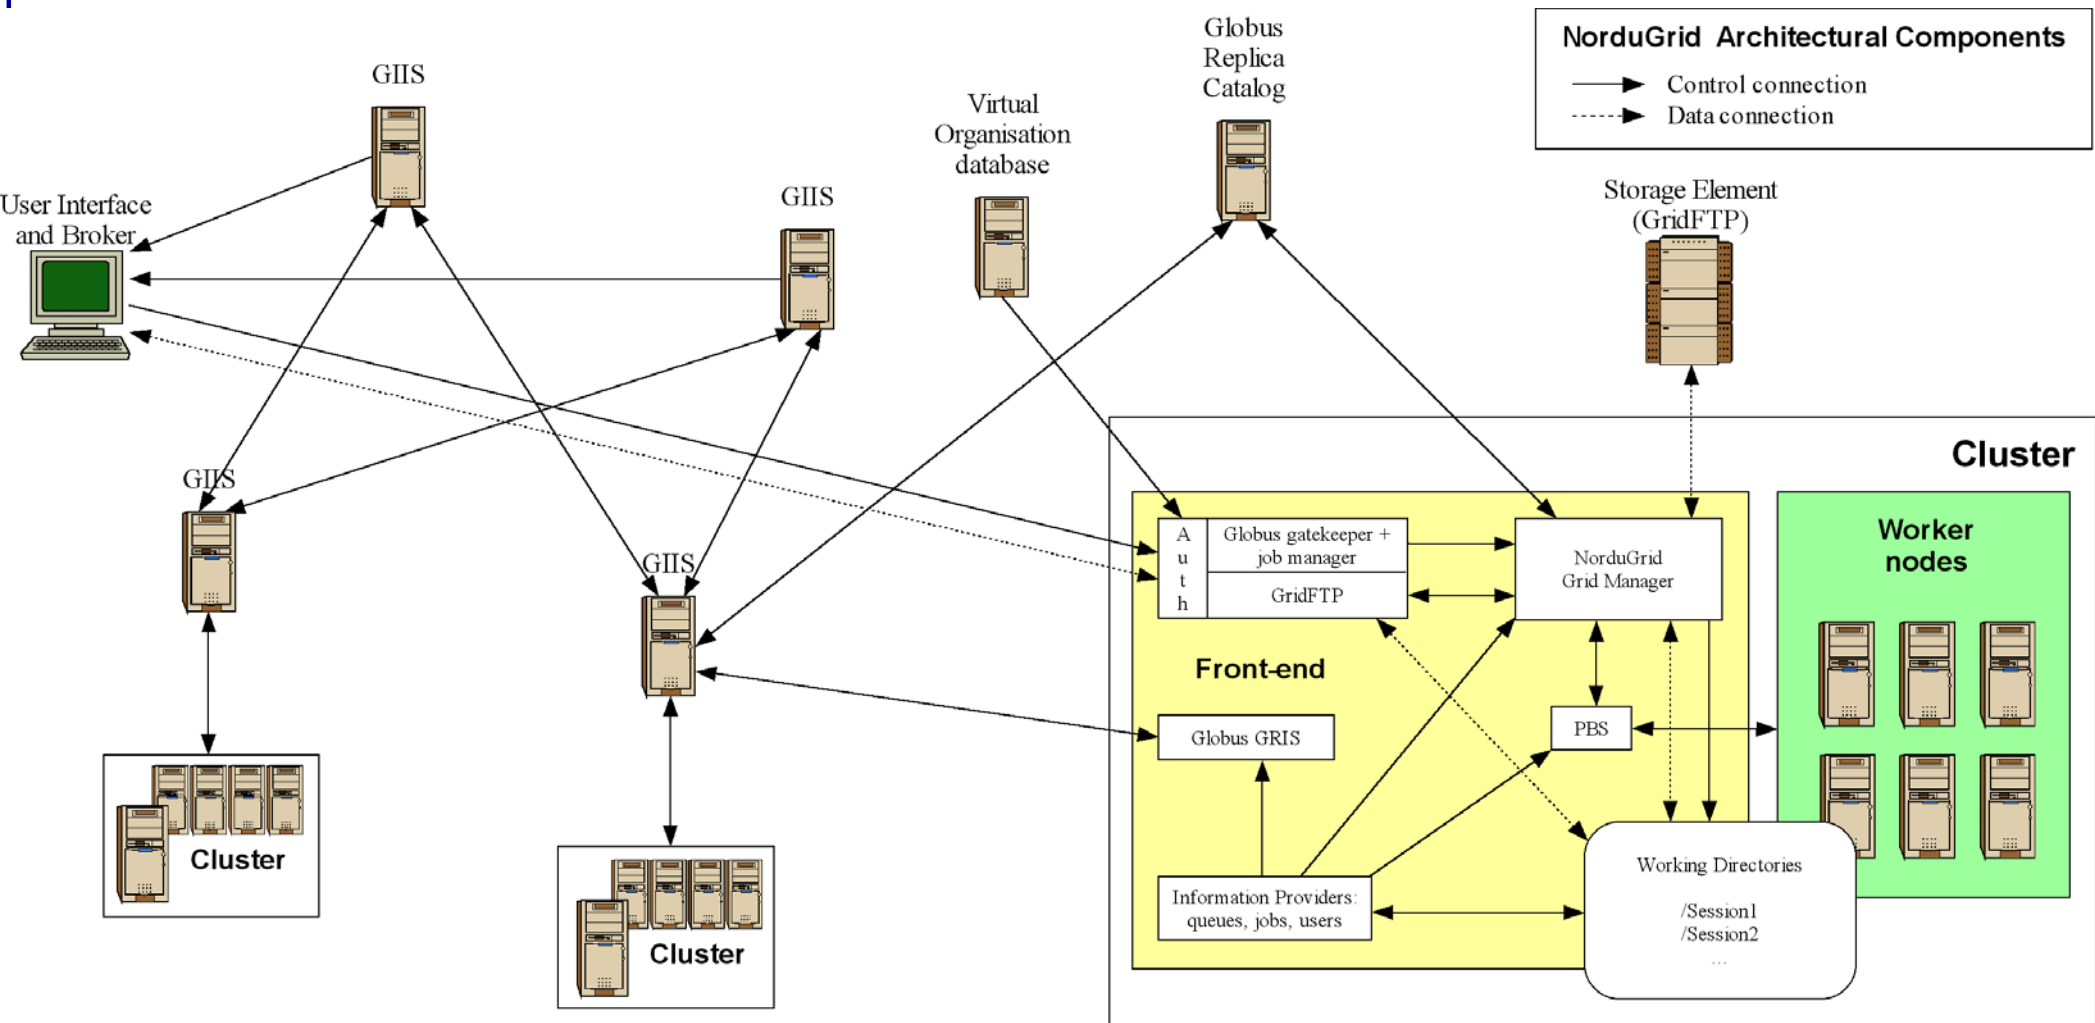
TeraGrid model of the Grid

- Linux Operating Environment
- Basic and Core Globus Services
 - ➔ GSI (Grid Security Infrastructure)
 - ➔ GSI-enabled SSH and GSIFTP
 - ➔ GRAM (Grid Resource Allocation & Management)
 - ➔ GridFTP
 - ➔ Information Service
 - ➔ Distributed accounting
 - ➔ MPICH-G2
 - ➔ Science Portals




- ➔ Advanced and Data Services
 - ➔ Replica Management Tools
 - ➔ GRAM-2 (GRAM extensions)
 - ➔ CAS (Community Authorization Service)
 - ➔ Condor-G (as brokering “super scheduler”)
 - ➔ SDSC SRB (Storage Resource Broker)
 - ➔ APST user middleware, etc.

The NorduGrid Architecture



Grid & Supercomputers

- The present day Supercomputers are the PC clusters
- Grid will provide a uniform access to all the resources
- The Supercomputing centers will be the power plants of the Grid



clusters, P2P, Grid

Cluster:

- single administrative domain
- centralized resource management, full control over resources
- suitable for strongly-coupled applications
- limited capacity

Grid:

- a layers on top of clusters, bunch of services on top of clusters
- “borrows” resources from clusters, capacity will be able to be reserved
- multiple administrative domains

Peer-to-Peer:

- millions of uncoordinated, unorganized desktops (screensavers)
- parallel application pools
- capacity varies and mostly unpredictable

Applications

- Applications are key to the success of Grid
- Application developers will only pick up the Grid IF:
 - Grid services will have a well-defined interface
 - Grid infrastructure some day be as natural part of the picture as the OS
- We are still very far from “throwing any application onto the Grid”
- Considerable porting effort in “Griddifications” of problems

targeted application areas

Genetics

Computer Graphics

Bioinformatics

VLSI Design

Data Mining

Molecular Dynamics

Network Simulation

Bioinformatics

Fluid Dynamics

Astrophysics

Drug Design

Protein Structure

Meteorology

Earth Observation

Civil Engineering

Financial Risk Analysis

“best” applications for the Grid

Decoupled applications (minimal communication)

- embarrassingly parallel apps, parameter sweeps

Staged/linked applications

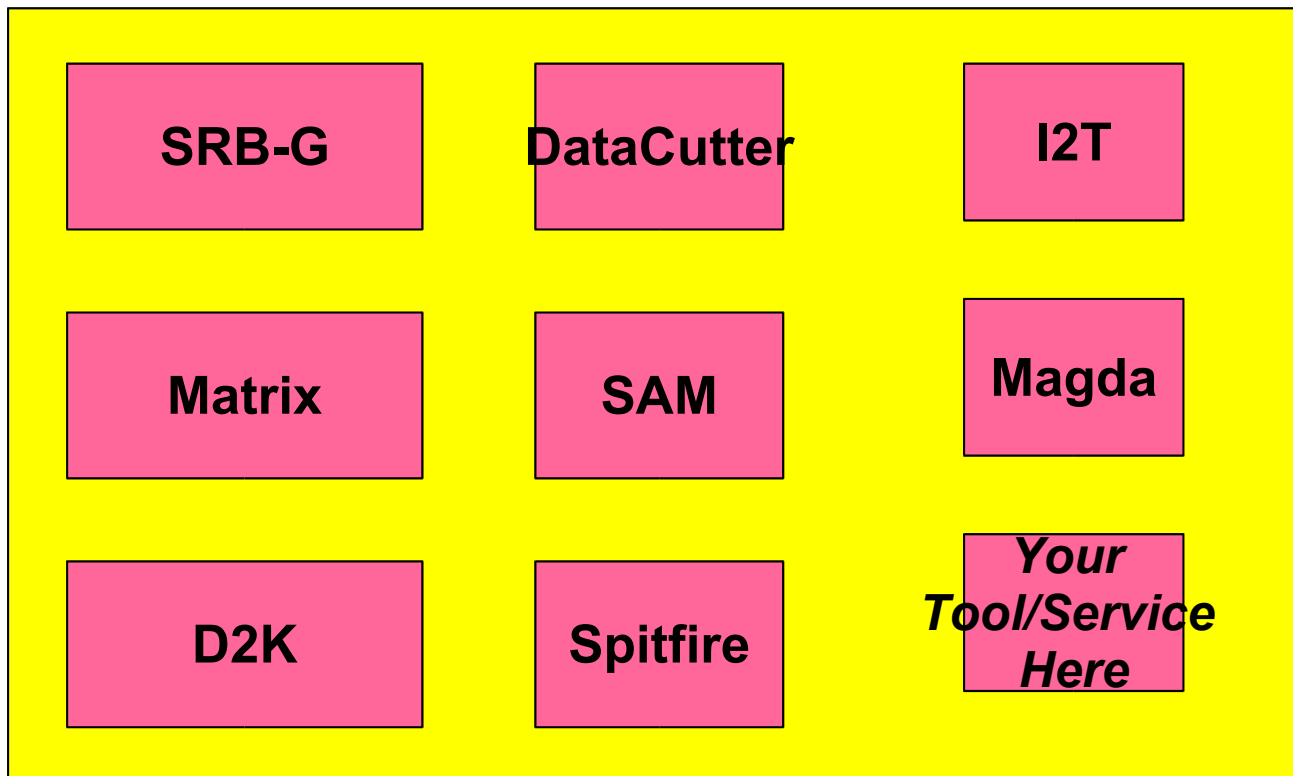
- (complete part A then do part B)
- Includes remote instrument applications
(get input from instrument at site A, compute/analyze data at site B)

Access to resources (mainly data)

- get “something” from/do “something” at site A
- dataGrids
- data & controlled/shared access to data is the critical issue of the future

“Data is the killer app”

there are many similar but incompatible solutions :



Alessandro Volta in Paris in 1801 inside French National Institute shows the battery while in the presence of Napoleon I



Fresco by N. Cianfanelli (1841)

(Zoological Section "La Specula" of National History Museum of Florence University)

Oh, mon Dieu !

What !?!?
This is a mad man...

....and in the future,
I imagine a
worldwide
Power (Electrical)
Grid



NorduGrid Project


www.nordugrid.org




- Create a Grid infrastructure in Nordic countries
- Operate a production quality Testbed
- Expose the infrastructure to end-users of different scientific communities
- Survey current Grid technologies
- Pursue basic research on Grid Computing
- Develop Middleware Solutions

brochure:www.nordugrid.org/documents/brochure.pdf

Participants



Copenhagen University: Niels Bohr
Institute, Research Center COM, DIKU



Oslo University, Bergen University



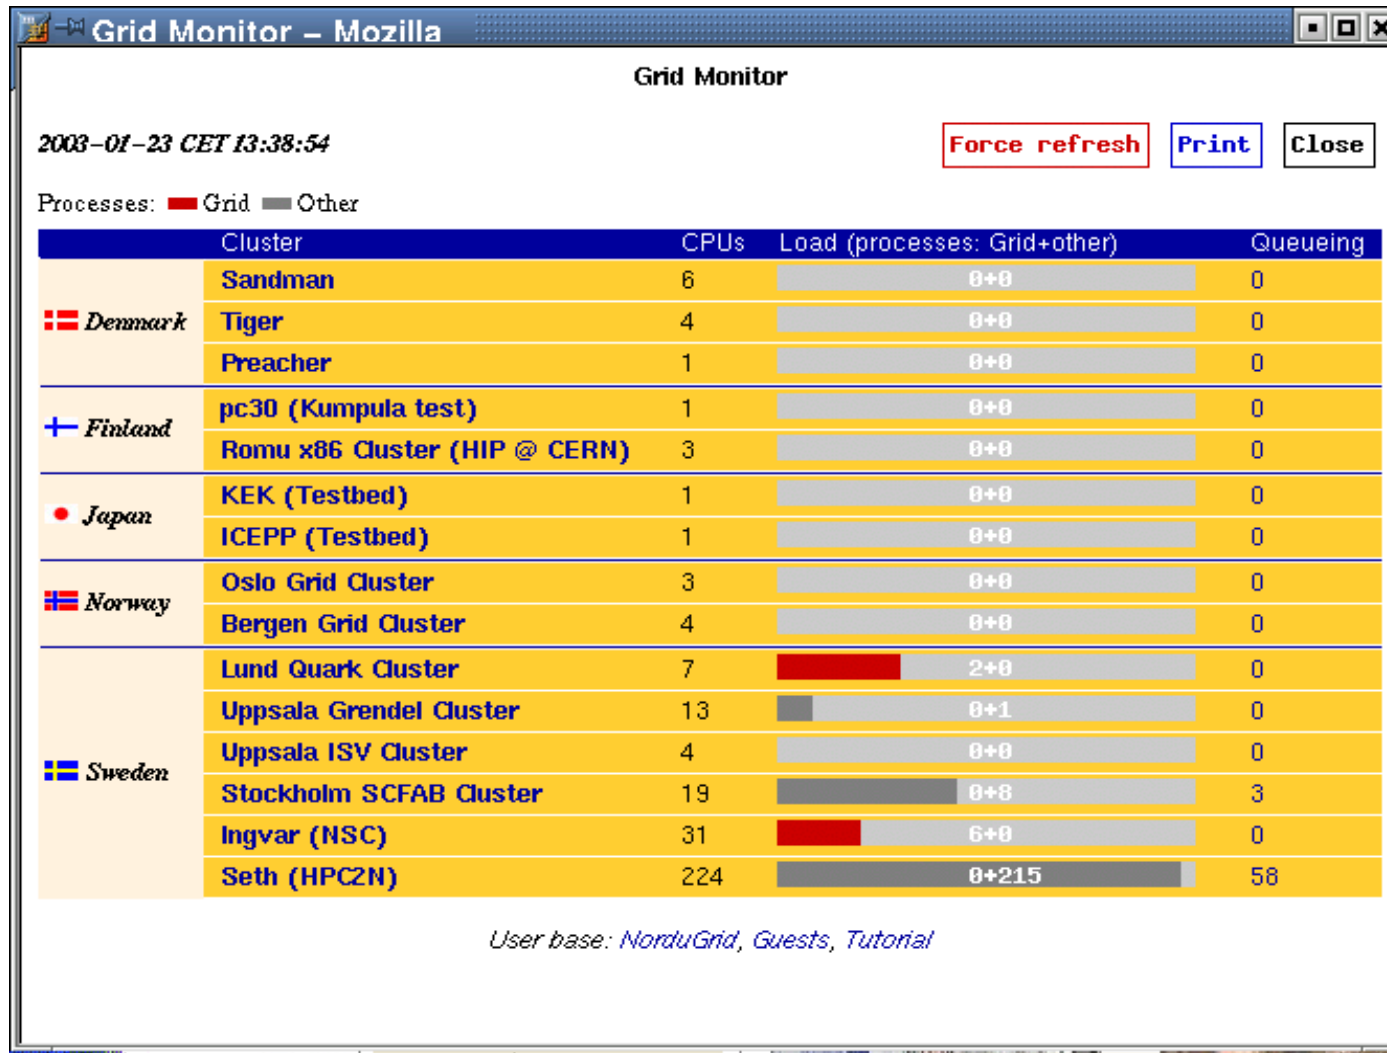
Lund University, Uppsala University,
Stockholm University, KTH



Helsinki Institute of Physics

resources:

www.nordugrid.org, and click on the GridMonitor



contact us

Dept. of Elementary Particle Physics, Physics Building A427,
balazs.konya@quark.lu.se

- ➔ The Grid needs you & your ideas :)
 - ➔ computing & physics diploma works within the NorduGrid
 - ➔ PhD topics
 - ➔ try to run your application on the Grid
 - ➔ or just look at the www.nordugrid.org

acknowledgement

- ➔ while I was preparing for this introductory talk I “borrowed” slides, ideas, pictures from general Grid-talks. I would like to thank all the authors of those talks. Especially to Rajkumar Buyya & Fran Berman*

*GGF5 Plenary Keynote: TeraGrid "State of the Grid 2002"