

# CS514: Intermediate Course in Computer Systems

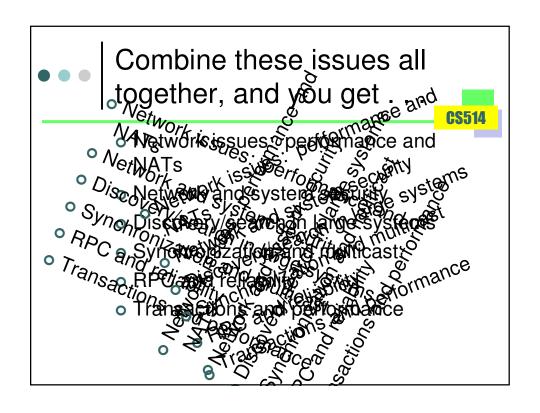
Lecture 39: April 25, 2003 "Grid Computing Systems"

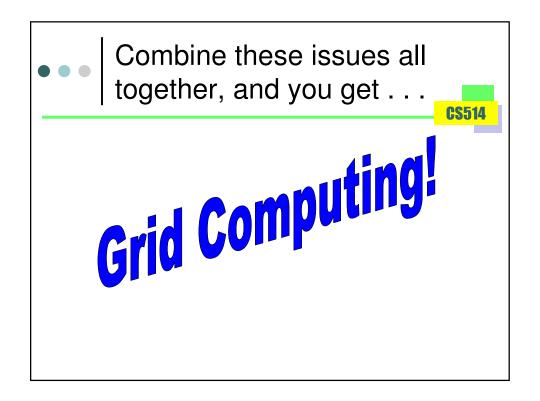


# We've looked at some hard issues . . .



- Network issues: performance and NATs
- Network and system security
- o Discovery/search in large systems
- Synchronization and multicast
- RPC and reliability
- Transactions and performance







# What is grid computing?



- Means lots of things to different people
- "Resource sharing and coordinated problem solving in dynamic, multiinstitutional virtual organizations"
- "Coordinated resource sharing in a distributed dynamic heterogeneous environment to enable efficient largescale problem solving"

## • • •

# What is grid computing?

**CS514** 

o "Enable communities ("virtual organizations") to share geographically distributed resources as they pursue common goals—in the absence of central control, omniscience, trust relationships"



# Common elements in definitions:

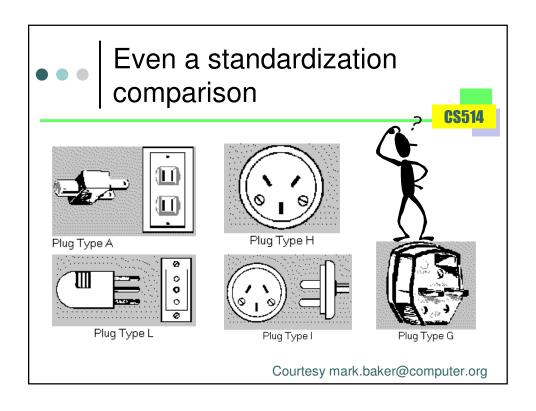


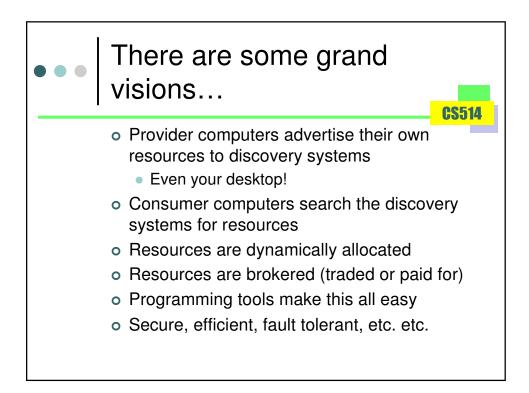
- o Dynamic
- o Resource sharing
- o Across organizations
- Where resources can be:
  - Raw materials of computing
    - CPU, storage
  - Execution of some specific function or service
    - A library routine, specific data, etc.



# Sometimes compared to power grid

- This is where the "grid" in "grid computing" comes from
- Plug into the computational grid as easily as we plug into power grid today
- Power grid has many sources and consumers of power, gets traded, moved around, etc.
- Even so, I'm not sure it is a good analogy







# Kind-of like web services on steroids!

**CS514** 

- And in fact, some web services standards are becoming the basis for grid computing standards...
  - WSDL has grid service description tag
    - Web Services Definition Language
  - SOAP for transport
  - UDDI registry and WSIL for service location
    - · IBM, Microsoft, HP, SAP registries
    - Web Services Inspection Language
    - Universal Description, Discovery, and Integration

## • • •

## Some grid computing history

**CS514** 

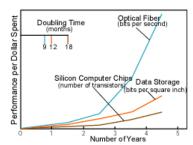
- Early 90s
  - Gigabit testbeds, metacomputing
- Mid to late 90s
  - Early experiments (e.g., I-WAY), academic software projects, application experiments
- o 2002
  - Dozens of application communities & projects in scientific and technical computing
  - Major infrastructure deployments
  - De facto standard technology: Globus Toolkit<sup>TM</sup>
  - Growing industrial interest
  - Global Grid Forum: ~1000 people, 30+ countries

from Globus/IBM tutorial

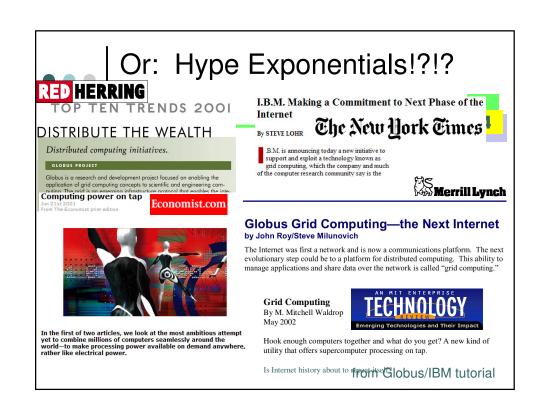
# Why now? Network Exponentials



- o Network vs. computer performance
  - Computer speed doubles every 18 months
  - Network speed doubles every 9 months
  - Difference = order of magnitude per 5 years
- o 1986 to 2000
  - Computers: x 500
  - Networks: x 340,000
- o 2001 to 2010
  - Computers: x 60
  - Networks: x 4000



Moore's Law vs. storage improvements vs. optical improvements. Graph from Scientific American (Jan-2001) by Cleo Vilett, source Vined Khoslan, Kleiner, Caufield and Perkins. from Globus/IBM tutorial





### What we'll look at

**CS514** 

- Two very different current successes
  - seti@home (and similar efforts)
  - Cornell etc. Adaptive Software Project
- And a number of different grid computing middleware projects
  - NetSolve
  - Globus/IBM

### • • •

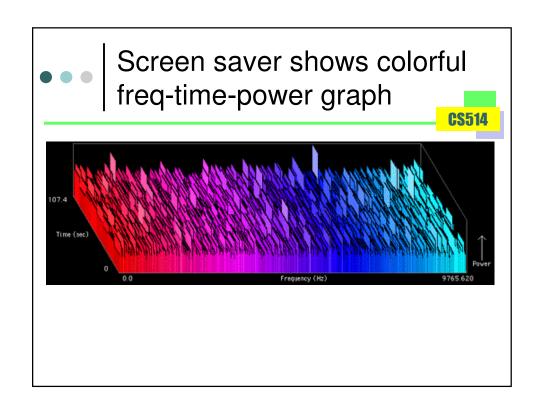
# seti@home goals

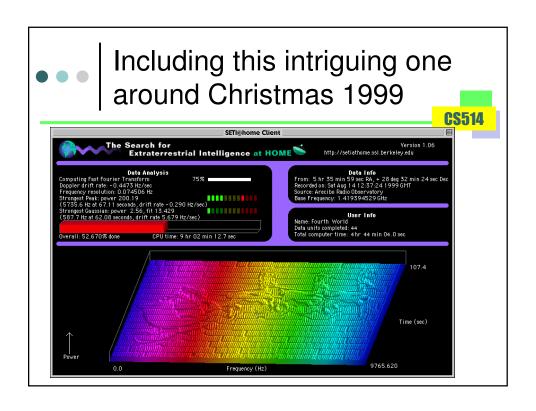
- Look for signs of extra-terrestrial life
- By harnessing the spare power of hundreds of thousands of Internetconnected computers
- To do signal processing on space signals
  - Piggybacked on data collected from Arecibo radio telescope, Puerto Rico

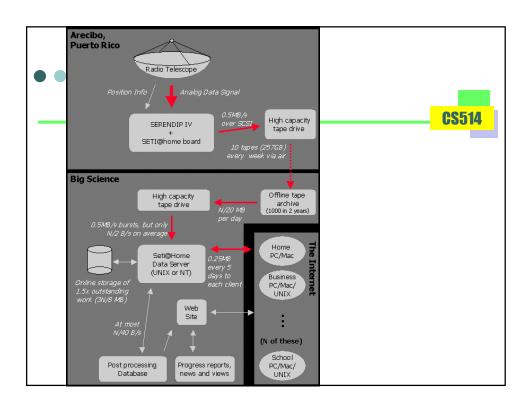


# Brilliant concept for user incentive

- "Each participant will have the slight but captivating possibility that his or her computer will detect the faint murmur of a civilization beyond Earth."
- o Cool screen saver!
- o User rankings on web site
  - Most hours logged









### Impressive results

**CS514** 

- Over 4 million users logging around 1.5 million CPU hours
- Why does this work?
- o seti@home is "embarrassingly parallel"
- Data processing can be broken out into individual tasks
  - no inter-processor communications
  - no fault-tolerance issues



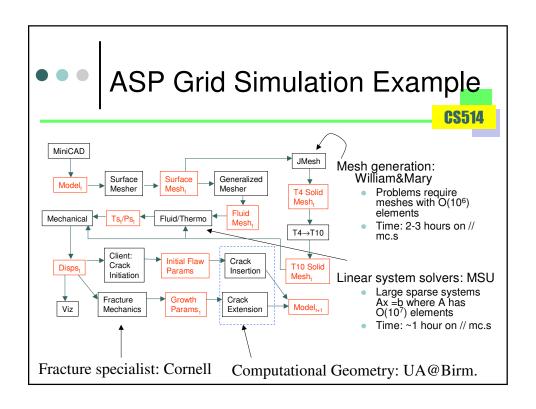
# Common class of scientific computing problem

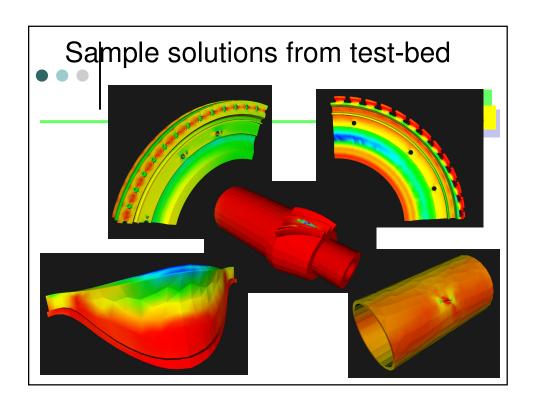
- o Globus folks call this a "Data Grid"
- Examples include
  - High-energy physics
  - Biotech: screen molecules for new drugs
    - "patriotgrid" in the case of anthrax treatment (United Devices)
  - Math problems, code breaking

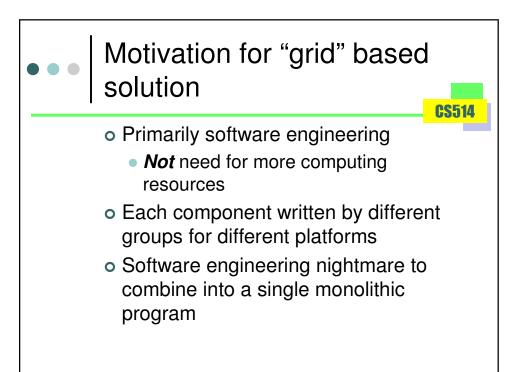
# • • Adaptive Software Project

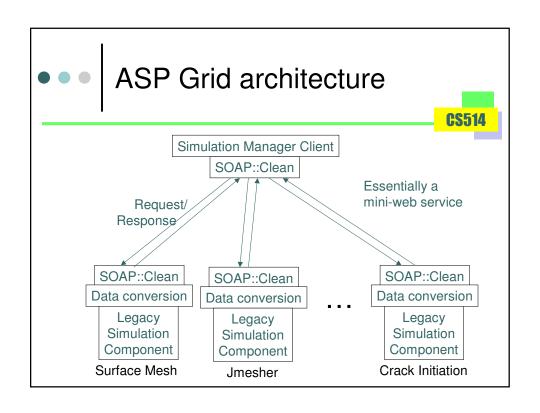
- **CS514**
- Cornell, Mississippi State, Ohio State,
   William and Mary, Clark Atlanta
- Inter-disciplinary team to simulate fracture mechanics in an internal combustion engine

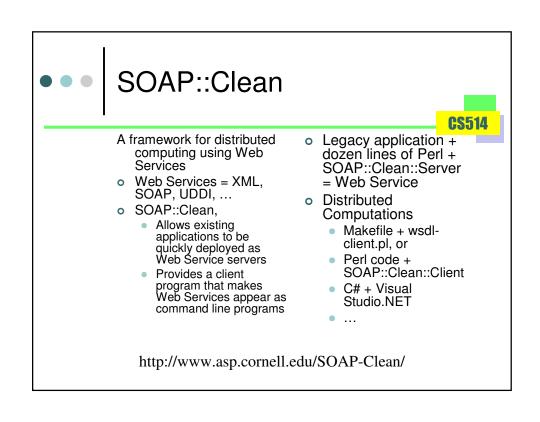
Slides courtesy of Paul Stodghill













# Performance using SOAP::Clean



Component	Wght	Request size	Response size	Base Running	Remote Running	Overhead	
		(bytes)	(bytes)	(seconds)	(seconds)	(seconds)	(%)
Surface Mesh	2	135,074	377,992	63.69	74.85	11.16	17.52%
Jmesher	2	512,386	1,334,431	107.26	116.51	9.25	8.62%
T4→T10	2	1,468,764	4,454,170	70.01	81.32	11.31	16.15%
Generalized Mesh	2	513,015	2,824,138	39.17	30.63	-8.54	-21.80%
Fluid/Thermal	2	7,416,514	1,038,255	1584.36	1603.39	19.03	1.20%
Crack Growth	1	135,017	141,467	0.42	7.72	7.30	1738.10%
Fracture Mechanic	1	6,621,903	3,201	36.43	56.59	20.16	55.34%
Crack Initiation	1	141,299	180,465	0.31	8.65	8.34	2690.32%
Weighted Total		26,989,725	20,383,105	3766.14	3886.36	120.22	3.19%

### Preliminary results

- Small problem size
- Room for additional optimizations

Base Running = time to directly run components, without the o Overhead will decrease as use of Web Services

Remote Running = client at UA@Birm., server at CUCS

### Conclusions

- o 3.19% overhead is fantastic for an interstate computation.
- the problem size increases.
  - Data is *O*(*n*).
  - Running times are > O(n).



# Advantages of grid-based simulation



- Simplifies project management
  - no need for everyone to agree on a common implementation language or hardware platform
  - need agree only on data exchange format (XML/SOAP)
- Avoids software maintenance problem
  - each project site maintains its own code but makes it available to other partners as a web service
- In future
  - computations scheduled for execution wherever there are free resources
  - computations may even migrate during execution where more resources become available



# Implications for software



- Software needs to be adaptive
  - adaptation for efficiency
    - application must be optimized dynamically when computation starts on or migrates to a new platform
  - adaptation for survival
    - adapt gracefully to processor and link failures: self-healing software
- → Software must become more intelligent



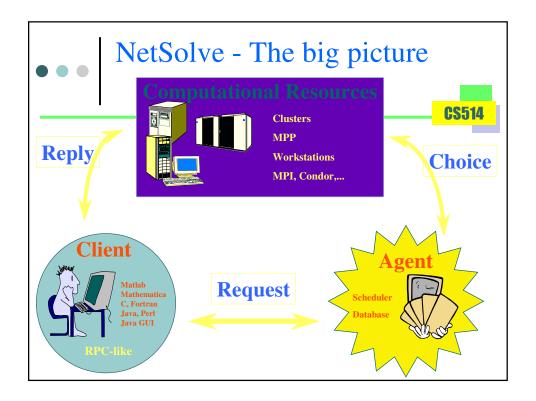
# Two examples of working grid computation

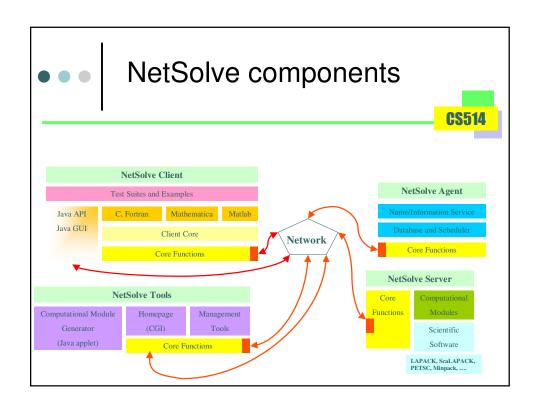
- One works because "embarrassingly" easy to parallelize
- One works because component teams work closely together to define interactions
- Neither really satisfies definition of "dynamic resource sharing across organizations"

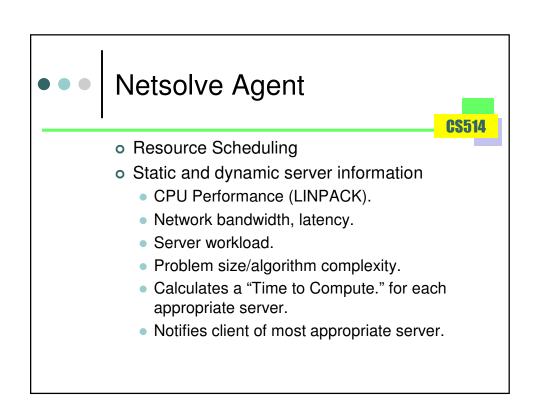
# • • Netsolve

- Jack Dongarra et.al. (many et.al.)
  - Univ of Tennessee
- Motivated by Problem Solving Environments (PSE)
  - MatLab, Mathematica, etc.
- Goal more to find resources specific to a problem than to find massive computing resources
  - Certain computation libraries, computations that require large intermediate storage

Some slides courtesy Casanova and Dongarra









### **Netsolve Client**

**CS514** 

- Function Based Interface.
- Client program embeds call from NetSolve's API to access additional resources.
- Interface available to C, Fortran, Matlab, Mathematica, and Java.
- Opaque networking interactions.
- NetSolve can be invoked using a variety of methods: blocking, non-blocking, task farms (for simulation runs), etc.

## • • •

# Netsolve programming

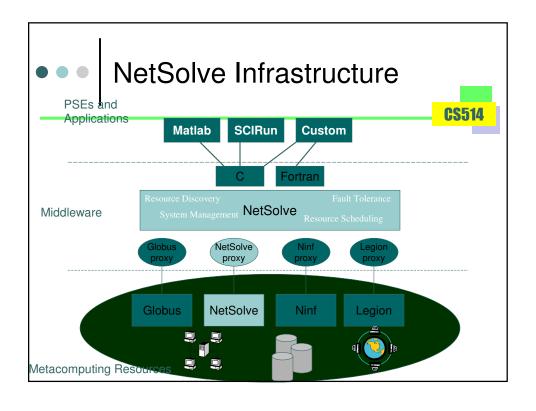
**CS514** 

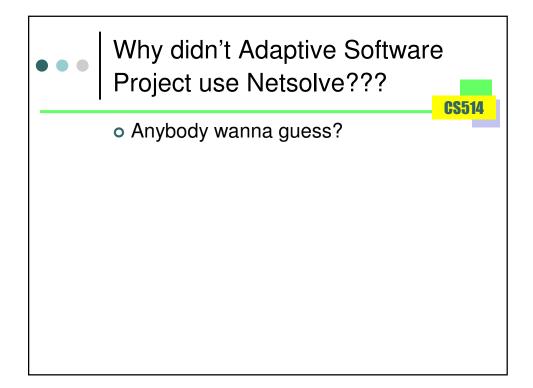
- o Intuitive and easy to use.
- o Matlab Matrix multiply e.g.:

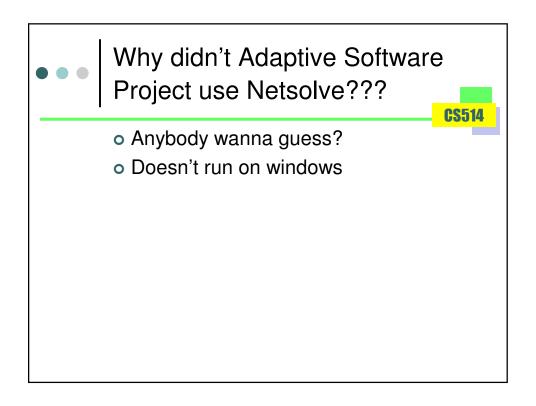
```
A = matmul(B, C);A = netsolve('matmul', B, C);
```

• Request farming:

# Not meant to scale to massive infrastructures Uses Kerberos for security Plugs in to other grid middleware



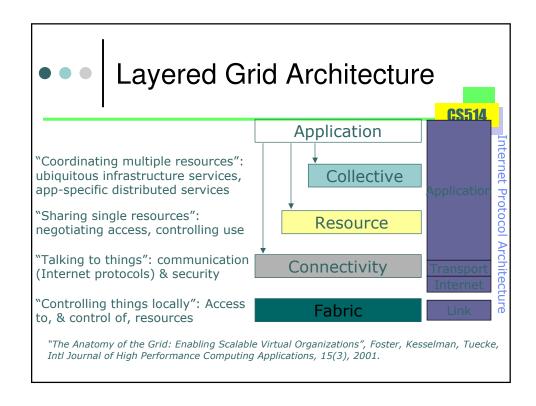






# The Globus Project (www.globus.org)

- o Ian Foster, Carl Kesselman
  - and about 60 people!
- Argonne National Lab, Univ of Chicago, USC-ISI
- Globus Toolkit v2 (GT2)
- Many projects (a few dozen)
- Picked up by IBM
  - Which is cramming Globus into a web services framework
- Has won mindshare, now it just needs to work! Slides courtesy "Introduction to Grid Computing and Globus Toolkit", only 220 pages long...





## **Key Protocols**



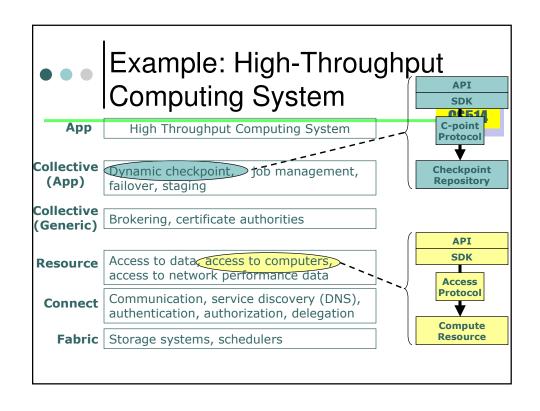
- The Globus Toolkit<sup>™</sup> v2 (GT2) centers around four key protocols
  - Connectivity layer:
    - Security: Grid Security Infrastructure (GSI)
  - Resource layer:
    - Resource Management: Grid Resource Allocation Management (GRAM)
    - Information: Grid Resource Information Protocol (GRIP/LDAP)
    - Data Transfer: Grid File Transfer Protocol (GridFTP)
- Also key collective layer protocols
  - Monitoring & Discovery, Replication, etc.

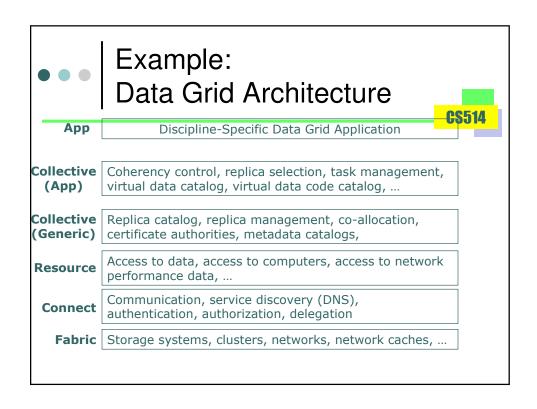


# GT-Based Grid Tools & Applications

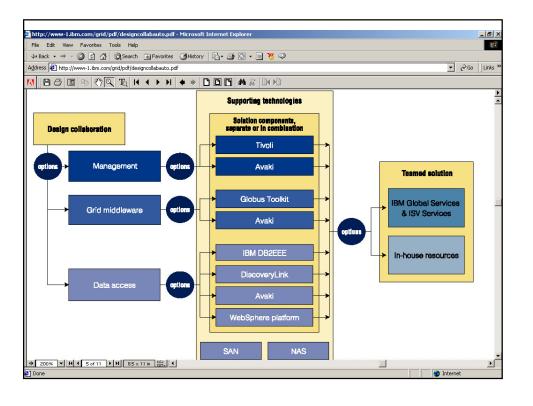


- Data Grids
  - Distributed management of large quantities of data physics, astronomy, engineering
- High-throughput computing
  - Coordinated use of many computers
- Collaborative environments
  - Authentication, resource discovery, and resource access
- Portals
  - Thin client access to remote resources & services
- And combinations of the above





# Target specific industries with compute intensive tasks Aerospace Automotive Life sciences Financial Sell them on idea of utilizing wasted computing resources in the organization In coordination with other web services, data sharing, and workflow products...





# A hard sell

- Requires buying, customizing, and managing new software
- Are the problems associated with that outweighed by increased utilization of resources?