

# The $\beta$ -Grid: A National Infrastructure for Computer Systems Research

Ian Foster

Argonne National Laboratory

The University of Chicago

Randy Butler, Charlie Catlett

National Center for Supercomputer Applications

# Overview

- Views on the convergence of networking, storage, and computing
  - ◆ Particular focus on the “Grid” community
- The need for a national infrastructure to support research into future the services & applications that enable/depend on this convergence
- Discussion of approaches to building such a national infrastructure


# Basic Assumption: Convergence

- In today's Internet, the network is a dumb bitway that moves data to and from endsystems (computer and storage)
- In future networks, boundaries between computing, storage, communication will blur
  - ◆ Networks will incorporate substantial embedded storage and computing
  - ◆ Sophisticated middleware will exploit these resources to provide value-added services
- I.e., the network will evolve into a Grid

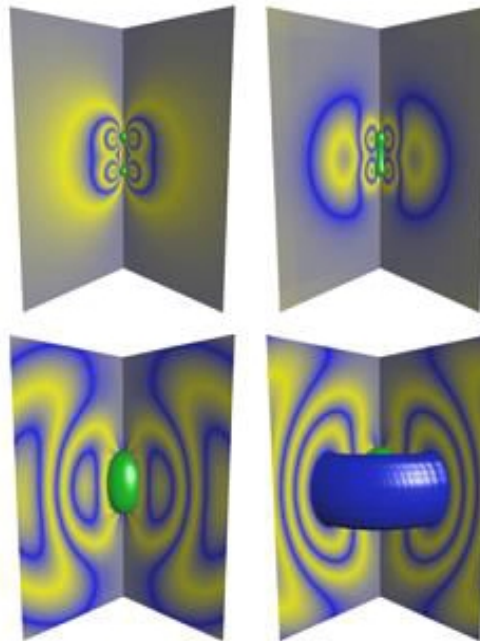
# Grid Applications



Teleimmersion/distance collaboration



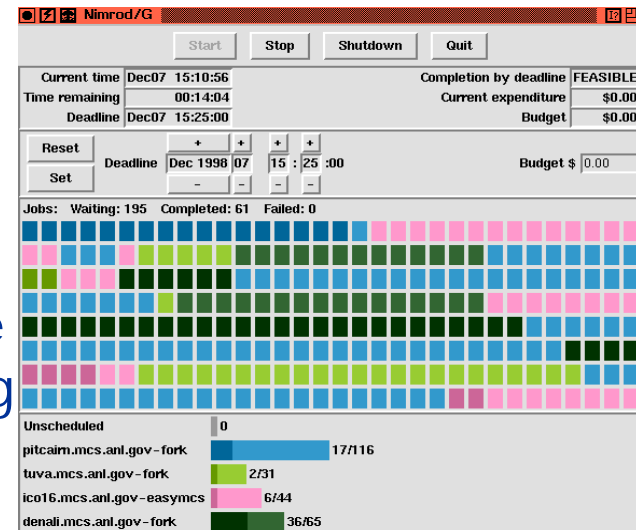
Record-setting distributed supercomputing



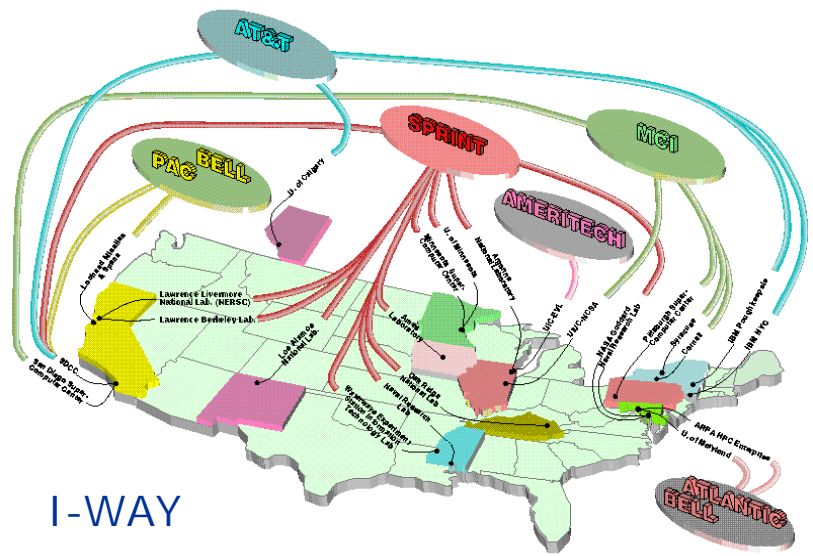
TransAtlantic remote visualization/steering



Online analysis of instrument data



Parameter studies with deadline scheduling



I-WAY




The Alliance National Technology Grid

## National and International Grid Testbeds



NASA's Information Power Grid

Ian Foster



High-performance networks from across the United States connect with international networks at the STAR TAP access point in Chicago, Illinois

# Evolution of Grid Concepts

- Initial goal was enabling remote access to unique end system resources
- Evolution to the creation of enhanced “network” services that rely on geographically distributed resources
- Next logical step is to (logically) migrate resources from end systems into networks
- Research challenges then lie in the definition of “Grid services” (middleware)

# Unifying Concept: An Integrated Grid Architecture

Applns

... a rich variety of applications ...

Appln  
Toolkits

Remote  
data  
toolkit

Remote  
comp.  
toolkit

Remote  
viz  
toolkit

Async.  
collab.  
toolkit

...

Remote  
sensors  
toolkit

Grid  
Services

Protocols, authentication, policy, resource  
management, instrumentation, discovery, etc., etc.

Grid  
Fabric

Archives, networks, computers, display devices, etc.;  
associated local services

# Computer Science Communities with an Interest in Convergence

- Logistical networking
- Web caching
- Grids/collaboratories/etc.
- Networking
- Distributed operating systems
- Distributed systems
- Distributed algorithms
- Agents
- And friends in various application domains




# These Communities Need Experimental Infrastructure!

- CAIRN
  - ◆ Network testbed, not endsystems
- The L-Bone (Logistical Backbone)
  - ◆ Storage focused
- D'Agents
  - ◆ Mobile agent testbed
- ACCESS (Tom Anderson)
  - ◆ Proposal, focused on low-level systems
- GUSTO
  - ◆ Volunteer contributions, focused on apps

# GUSTO Computational Grid Testbed

as of November 1998



# We Need New Experimental Infrastructure

- Simulation infrastructure
  - ◆ E.g., Chien's "Micro Grid" activity
- Experimental testbed(s)
  - ◆ dedicated to research
  - ◆ of a scale to permits realistic experimentation
  - ◆ with flexibility to enable wide range of experiments
  - ◆ highly instrumented as befits lab equipment
  - ◆ of a scale that encourages participation by adventurous applications groups

# Hypothesis/Proposal

- Substantial scale is desirable to enable interesting experiments
  - ◆ Substantial node capabilities: computing & storage
  - ◆ High-speed network access
  - ◆ Large number of nodes
  - ◆ Geographical distribution
- Feasible, and indeed useful, for multiple communities to share infrastructure

# What Scale is Needed?

- Individual nodes provide
  - ◆ OC48 (2.4 Gb/s=300 MB/s) network interfaces and paths to disk
  - ◆ Reasonable processing capabilities on that data: e.g., 50 I/B/s = 15 BIPS
  - ◆ Storage for say 2 hours of data = 2 TB
- 20-100 such nodes for good aggregate capability, scale, national-scale distribution
  - ◆ 40-200 TB, 300-1500 BIPS aggregate
- High-speed network links, national spread

# Software Requirements

- Support access by large user community
    - ◆ Resource management, security, accounting
  - Support modification by many developers
    - ◆ Automated software update, distribution, configuration, intrusion detection
  - Enable large-scale application evaluation
    - ◆ Support co-existence of (experimental) applications & experimental services
  - Support wide range of experiments
    - ◆ Good core services, easily extended
-

# Conducting Experiments

- Enquire as to current testbed state
  - ◆ E.g., via Grid Information Service
- Schedule testbed resources for experiment
  - ◆ CPU, disk, network interfaces
- Establish software for experiment
  - ◆ Specify required revision level for key software elements (automated update)
  - ◆ Install experimental software as required
- Monitor operation of experiment
  - ◆ Via instrumentation in network, kernel, etc.

# The $\beta$ -Grid Project

(“Broadband Experimental Terascale Access”:  
Foster, Butler, Catlett)

- NSF ANIR funding to U. Chicago & NCSA
- Supports
  - ◆ Planning activities: goals, tech, mgmt
  - ◆ Experimentation with node architecture
- Prototype  $\beta$ -Grid nodes deployed at Argonne and USC/ISI
- Prototyping  $\beta$ -Grid reference impln: Globus, NWS, DPSS, scheduler, etc., etc.



# $\beta$ -Grid Reference Implementation Builds on Grid Infrastructure

- Grid Security Infrastructure
  - ◆ Authentication, authorization, policy (?)
- Grid Information Service
  - ◆ LDAP-based
- Integrated instrumentation
  - ◆ Evolving: NWS is part of the package
- Resource mgmt: CPUs, storage, networks (?)
- Much more is needed: config mgmt,

# Missing Functionality


- User management
    - ◆ Establishment and management of trust relationships
  - Accounting, billing, policy
  - Configuration management
    - ◆ Automatic software update
  - Intrusion detection
  - Support for co-existence of multiple uses:  
network, middleware, application  
researchers
-

# Summary

<http://dsl.cs.uchicago.edu/beta>

- Future networks will feature considerable embedded storage and computing
  - Testbeds are needed now to enable the CS research that will define these architectures
  - $\beta$ -Grid is a sketch for what such a testbed might look like
  - Can support a wide range of experiments for both pure CS and application developers
  - Looking for input on feasibility, requirements, experiments, technology, participation, ...
-

# Access Grid



# Example: Climate Model Data

- Service = Delivery of derived data products to individual or collaboration
- Climate data distribution grid includes
  - ◆ High-speed networking
  - ◆ Distributed data caches
  - ◆ Transcoders/generators of various sorts
- Supported by security, accounting, resource management, replication, etc. services
- This grid created dynamically in support of a specific requirement: Virtual Private Grid