



# An Introduction to Grid Computing

Mark Servilla

Network Information System

LTER Network Office

University of New Mexico





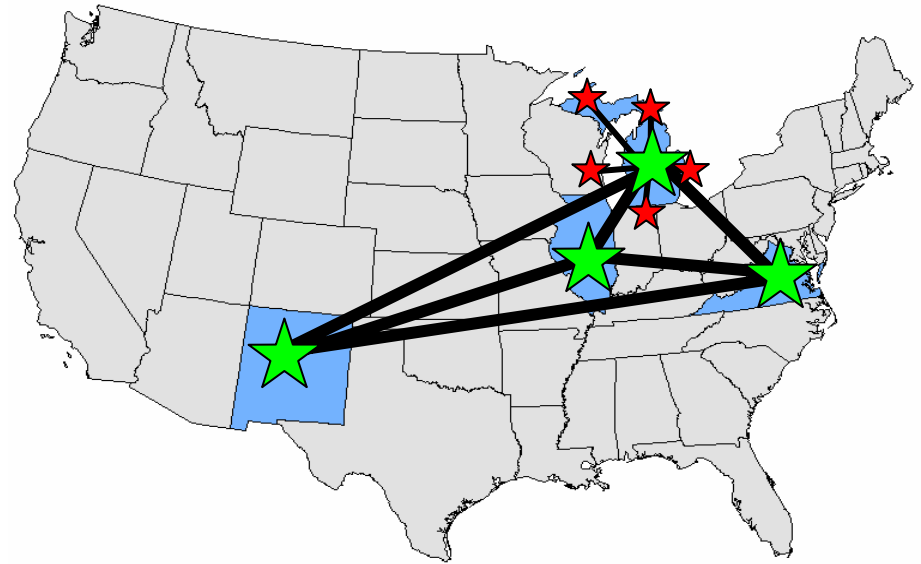
# Agenda

- Motivation
- Definitions
- Evolution of the Grid
- Characteristics
- Computing Model
- Examples
- References



# Scenario

- Data sensors geographically distributed
- Storage centers geographically distributed
- Analytical resources geographically distributed
- Colleagues at different institutions or research sites
- Seamless integration of all resources



What infrastructure would you need to make this possible?

**A Grid!**



# Definitions

- *"... uses the resources of many separate computers connected by a network (usually the internet) to solve large-scale computation problems" – Wikipedia*
- *"... the illusion of a simple yet large and powerful self-managing virtual computer out of a large collection of connected heterogeneous systems sharing various combinations of resources" – IBM Redbooks*
- *"Grid Computing enables virtual organizations to share geographically distributed resources as they pursue common goals, assuming the absence of central location, central control, omniscience, and an existing trust relationship." – Globus Alliance*



# The Evolution of Grid Technology

- High-Performance Computing
- Cluster Computing
- Peer-to-Peer Computing
- Internet Computing



# High-Performance Computing

- Traditionally known as super-computing
- Specialized for parallel processing algorithms
- Shared equally among academia, research, and commercial sectors

Industry Area	Sample Companies
Telecommunication	Sprint, Duetsche Telekom
Finance	Charles Schwab
Automotive	BMW, GM
Database	State Farm, Starbucks
Transportation	Oy Saimaa Lines
Electronics	Cisco, Motorola
Geophysics	Aramco, Shell
Aerospace	Dassault Avation
Energy	Centrica
Worldwide Web	Newsky, Amazon
Information Services	EDS
Chemistry	Bayer
Manufacturing	Alcoa
Mechanics	Hitachi
Pharmaceutics	Aventis Pharma

Source: University of Mannhiem



# Cluster Computing

- Originated 1994 – Beowulf cluster NASA
- Classification
  - High-availability – fault tolerant
  - Load-balancing – simultaneous user base
  - High-performance – computational use
- Massively-parallel (2 to 1000's of nodes)
- Commodity hardware (Intel, AMD, PowerPC)
- Low-cost software (Linux, FreeBSD, MacOS)
- Interconnected via high-speed private networks
- Shared storage SAN/NAS
  
- System X at Virginia Tech – 12.25 Tflops, 7<sup>th</sup> fastest high-performance computer in the world (Nov 2004, TOP500)





# Peer-to-Peer Computing


- Primarily used for distributed storage and file-sharing
- Early models (rcp, scp, ftp)
  - Restricted to LANs, or
  - Limited to known peers
- Internet-based models
  - Centralized (Napster, Kazaa\*)
  - Decentralized (Gnutella)

\*100,000,000 downloads by 2004; 2-million new downloads a week






# Centralized





# Decentralized





# Internet Computing

- Volunteer or philanthropic computing; utilizes personal desktop computers connected to the Internet
- Desktop computers idle approximately 95% of the their lifespan
- Divide and Conqueror approach
  - Tasks broken into smaller subtasks
  - Desktop executes subtasks during idle time
  - Desktop sends data back to central server, which aggregates results

Area of Interest	Projects
Science	SETI@Home Analytical Spectroscopy Research Group Evolutionary Research eOn Climateprediction.com Distributed Particle Accelerator Design
Life Sciences	Folderol Folding@Home Genome@Home FightAIDS@Home Übero Drug Design Optimization Lab The Virtual Laboratory Project Distributed Folding Community TSC Find-a-Drug
Cryptography	Distributed.net ECCp-109
Mathematics	Great Internet Mersenne Prime Search Proth Prime Search ECMNET n!+1 and n!-1 Prime Search Minimal Equal Sums of Like Powers GRISK MM61 Project 3x + 1 Problem Pi(x) Project Distributed Search for Fermat Number Divisors PCP@Home Generalized Woodall Numbers Generalized Fermat Prime Search ZetaGrid Strong Pseudoprime Search Wilson Prime Search Largest Proth Prime Search Search for Primes of the Form $k \cdot 2^n - 1$ Weferich Prime Number Search Seventeen or Bust Factorizations of Cyclotomic Numbers

Source: Grid Technology Partners



# Synthesis entrée Grid

- High-performance computing
  - pioneered the use of “parallel” algorithms
- Cluster computing
  - demonstrated the nature of shared computing and storage
  - load balancing protocols
- Peer-to-peer computing
  - distributed storage resource with no central authority
- Internet computing
  - geographically distributed virtual organization
  - fabric of the project vanishes with completion of the task



# Characteristics of a Grid

- Resources that
  - are connected via a network
  - are geographically distributed
  - heterogeneous hardware and/or software
  - managed transparently for performance and fault tolerance
- Illusion of a virtual organizations without
  - a central authority
  - a central control
- Explicit trust relationships between users and resources
- A system that scales in space and time




# Types of Resources

- Computation
  - utilization of computing cycles found on processors of the machines on the grid
- Storage
  - to increase capacity, performance, sharing, and reliability of data
- Communication
  - to increase capacity, performance, and reliability of data communication
- Collaboration tools
  - to facilitate collaboration through conferencing, visualization, and data sharing
- Software and Licenses
  - to share site-specific software and/or licenses
- Special equipment, capacities, architectures, and policies
  - printers, imaging, sensors, or other local specialty resources



# Grid Ingredients





# Grid Topologies

- Departmental Grids
  - localized to a specific group of people
  - generally, same hardware and software
  - designed for high throughput and high performance over a dedicated network
- Enterprise Grids
  - service to numerous groups within a single company or campus
  - resource heterogeneity increases
  - company-wide local area network
- Extraprise Grids
  - service to multiple companies, partners, and customers within a particular domain
  - domain based private network
- Global Grids
  - established over the public-Internet





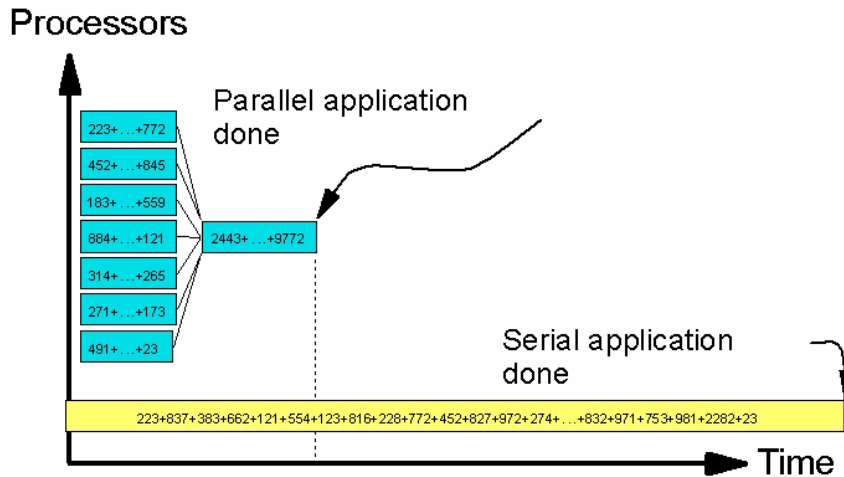
# Resource-based Grids

- Compute Grids
  - desktop nodes
  - server nodes
  - high-performance computing clusters
- Data Grids
  - performance-based distributed storage
  - replication for fault-tolerance
- Collaboration Grids
  - support for video-conferencing, visualization and data sharing
- Utility Grids
  - maintained and managed by a commercial service provider
  - compute resources acquired on a per-need basis
  - application resources that are purchased on a per-use or per-minute basis




# Application Characteristics

## Optimized for parallel execution



- **Perfect Parallelism** – computations run autonomously (Monte Carlo Simulations)
- **Data Parallelism** – operations performed on data simultaneously (db searches)
- **Functional Parallelism** – multiple operations are performed simultaneously

## Not capable of parallel computation



**Fibonacci Series (1, 1, 2, 3, 5, 8, 13, 21,...)**  
 $F(k+2) = F(k+1) + F(k)$



## Questions to ask? When thinking Grid

- Identity and Authentication
  - Is this user who he says he is?
  - Is this program the right program?
- Authorization and Policy
  - What can the user do on the grid?
  - What can the application do on the grid?
  - What resources are the user and or application allowed to access?
- Resource Discovery
  - Where are the resources?
- Resource Characterization
  - What types of resources are available?



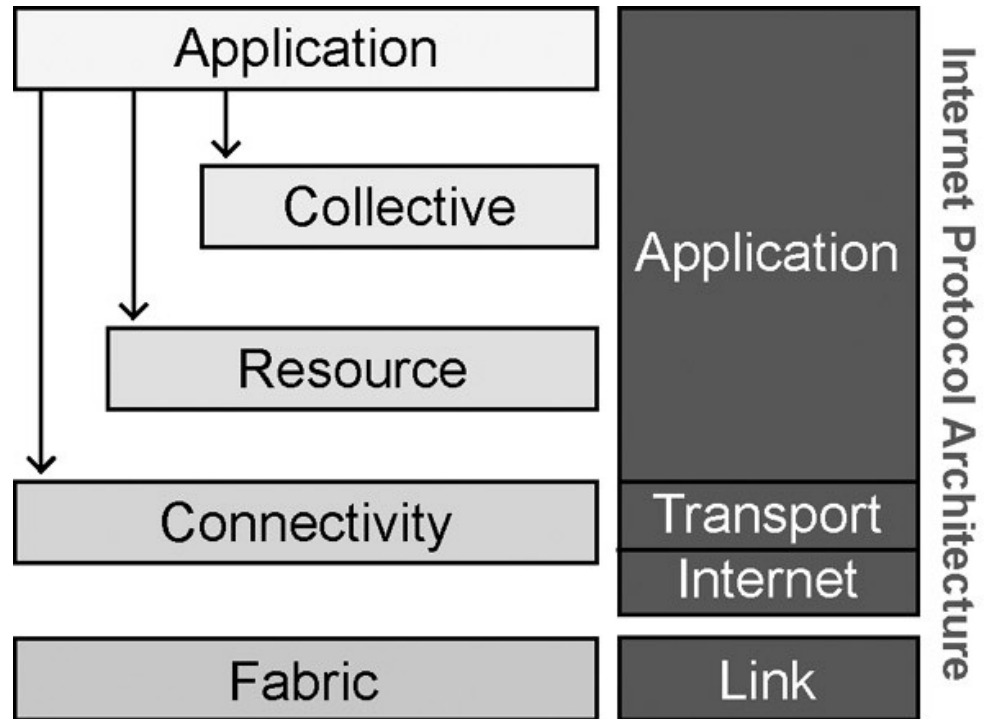
## Questions to ask? When thinking Grid

- Resource Management
  - Which resource can be used at what time and for what purpose?
- Accounting/Billing/Service Level Agreement
  - How much of the resources is being used?
  - What is the rating schedule?
  - What is the SLA?
- Security
  - How do I make sure that this is done securely?
  - How do we know if we have been compromised?
  - What steps are taken once a security breach is detected?



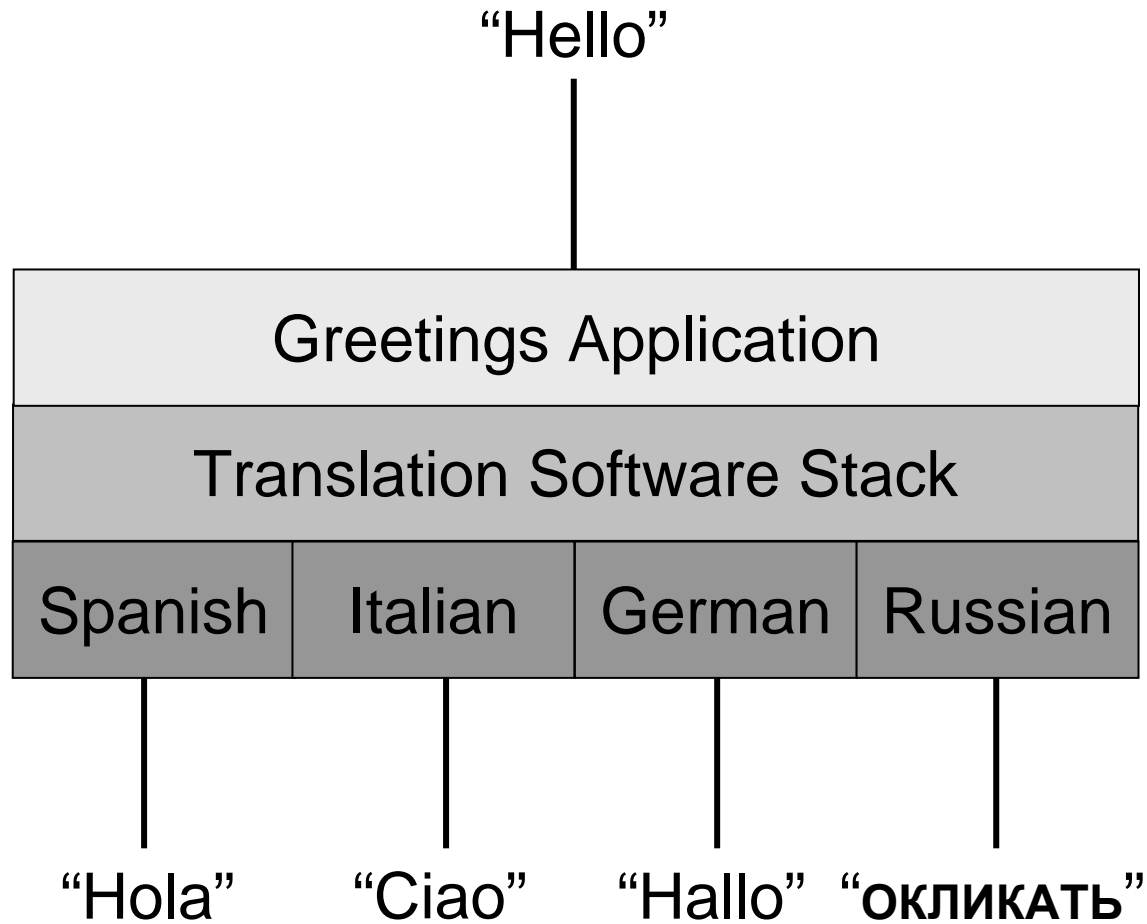
# Grid Computing Model (the Globus view)

- Software stack consisting of
  - Standards
  - Protocols
  - APIs and SDKs
- Loosely based on the Internet model





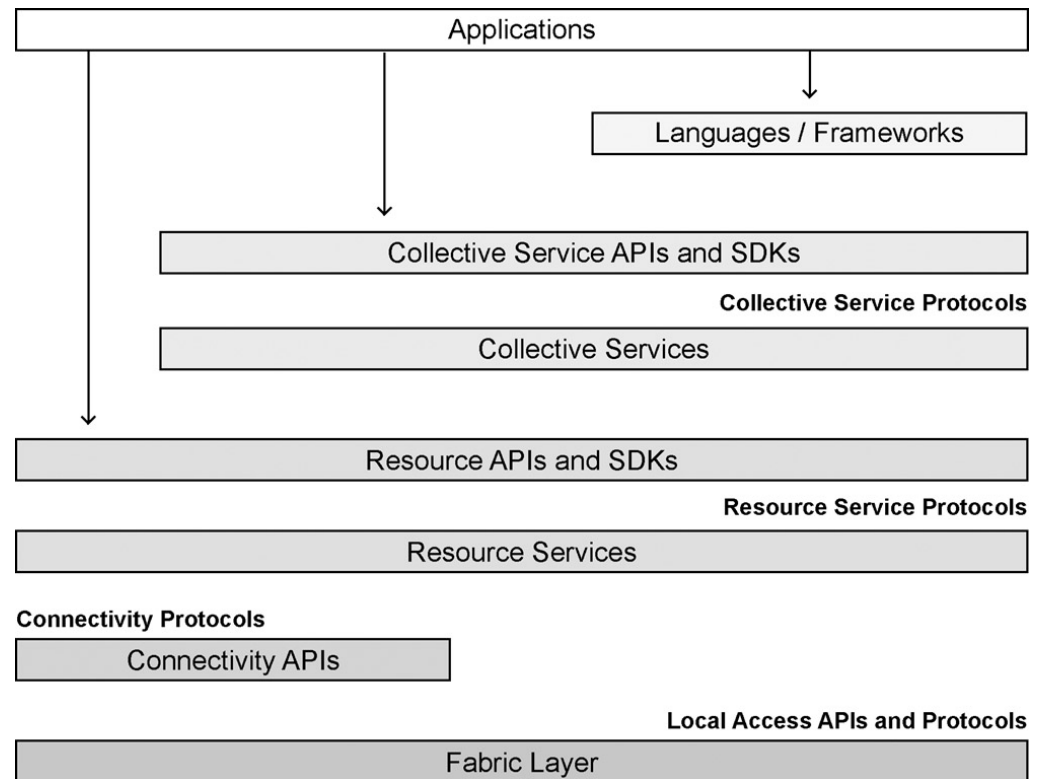
# Analogy





## A detailed view...

- **Application** – protocols targeted at a specific application or class of applications
- **Collective** – protocols for system-wide deployment (versus local)
- **Resource** – protocols to initiate and control sharing of local resources (GRAM, GridFTP, GRIS)
- **Connectivity** – protocols for grid-specific network transactions (IP, DNS, WSDL); Security implementation (GSI)
- **Fabric** – protocols and interfaces to resource being shared

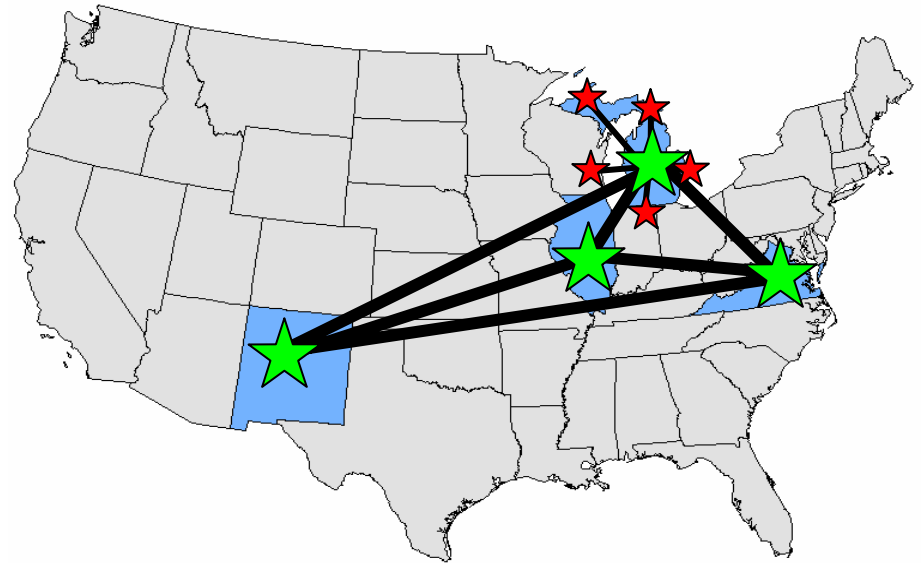




# Biophony Grid Portal

## Environmental Acoustic Pattern Matching

- Data sensors geographically distributed
- Storage centers geographically distributed
- Analytical resources geographically distributed
- Colleagues at different institutions or research sites
- Seamless integration of all resources







# Biophony Grid Portal



- Environmental Acoustic Monitoring
  - Embedded microphone
  - Sampling every 30 minutes for 30 seconds
  - Conversion to \*.wav file format
  - Raw data sent to MSU data store
  - Metadata sent to LNO Metacat database





# Biophony Grid Portal

- Ornithologic (bird) pattern matching
  - Analyze signal using DSP/Matlab workbench
  - Match unknown dataset to a known signature file

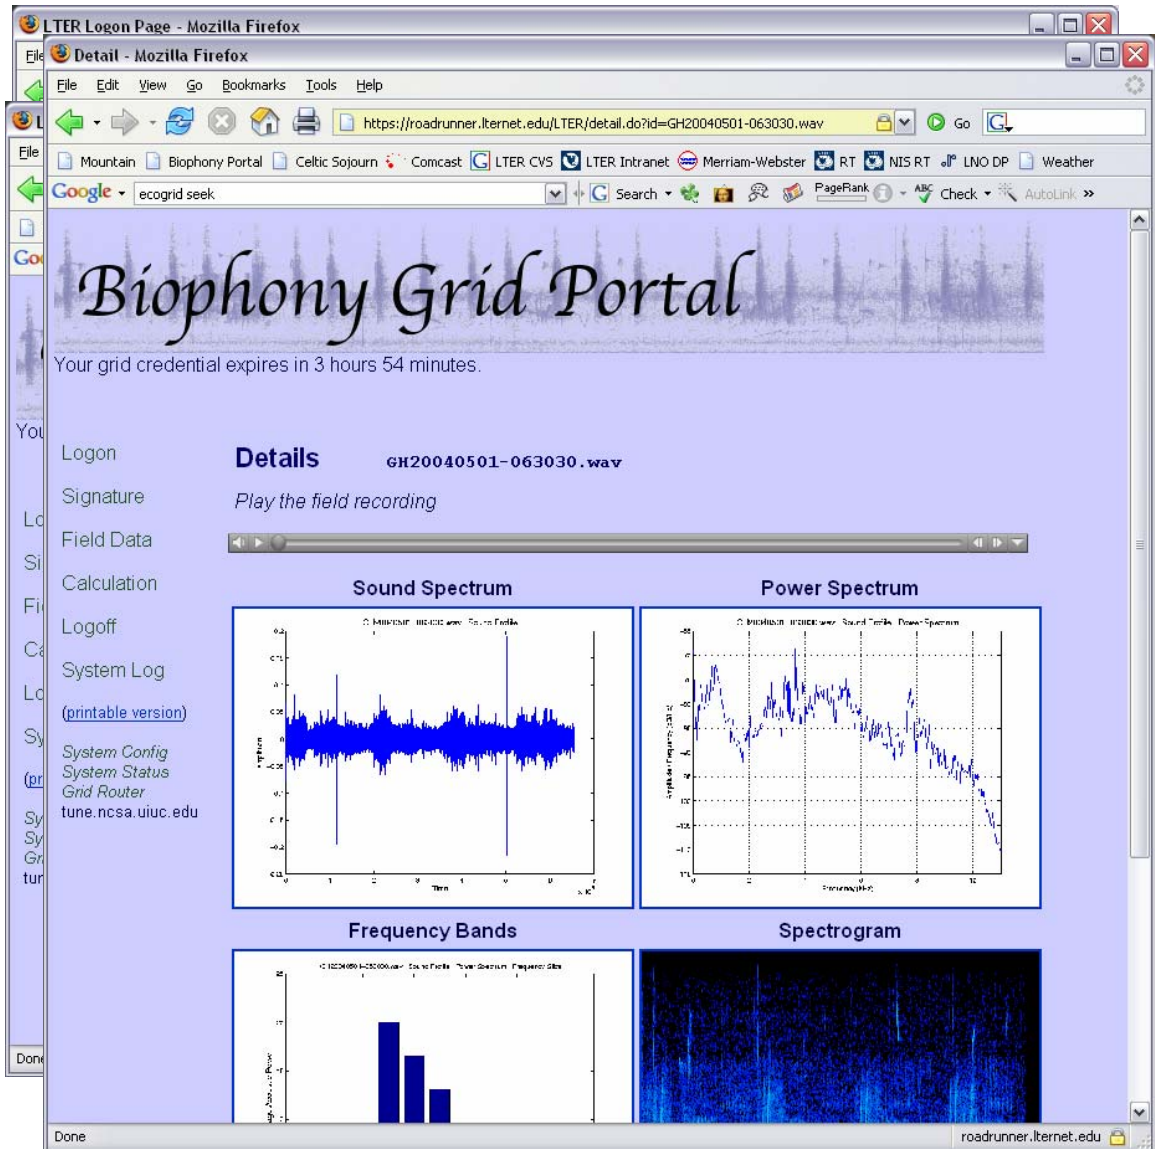




# Biophony Grid Portal

## User Scenario

- User logs onto portal (single sign-on)
- Selects known signature file(s)
- Performs data search via Metacat for unknown acoustic file(s)
- Performs acoustic matching on HPC; match probability identified to user
- Performs detailed analysis





## System Audits & Provenance

# Biophony Grid Portal

Biophony Grid Portal

Your grid credential expires in 3 hours 44 minutes.

LTER Log Page - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

https://roadrunner.lternet.edu/LTER/log.do

Mountain Biophony Portal Celtic Sojourn Comcast LTER CVS LTER Intranet Merriam-Webster RT NIS RT LNO DP Weather

Google ecogrid seek Search PageRank Check AutoLink Options ecogrid seek

### System Log / Data Provenance Log

Ligon

Signature

Field Data

Calculation

Logoff

System Log

(printable version)

System Config

System Status

Grid Router

tune.ncsa.uiuc.edu

Number of Records: 20 Update

ID	Hostname	Priority	Date/Time	Message
477258	grid.matrix.msu.edu	notice	2006-01-05 22:03:06	GRID-FTP: DATE=20060105220305.942041 HOST=grid.matrix.msu.edu PROG=globus-gridftp-server NL EVNT=FTP_INFO START=20060105220305.926224 USER=servilla FILE=/data/Signatures/Images/011.bmp BUFFER=0 BLOCK=262144 NBYTES=7990 VOLUME=/ STREAMS=1 STRIPES=1 DEST=[141.142.58.2] TYPE=ERET CODE=226
477248	grid.matrix.msu.edu	notice	2006-01-05 22:02:50	GRID-FTP: DATE=20060105220250.292130 HOST=grid.matrix.msu.edu PROG=globus-gridftp-server NL EVNT=FTP_INFO START=20060105220250.289264 USER=servilla FILE=/home/servilla/lter_data_transfer.sh BUFFER=0 BLOCK=262144 NBYTES=365 VOLUME=/ STREAMS=1 STRIPES=1 DEST=[129.24.124.180] TYPE=STOR CODE=226
476954	roadrunner.lternet.edu	notice	2006-01-05 22:00:08	METACAT: MetaCat: user [GSI DN: "/C=US/O=National Center for Supercomputing Applications/CN=Mark Servilla", context present] try to login
476953	roadrunner.lternet.edu	notice	2006-01-05 22:00:08	METACAT: MetaCat: --- Action: login
476889	roadrunner.lternet.edu	notice	2006-01-05 15:00:07	myproxy-server: <5197> Delegating credentials for /C=US/O=National Center for Supercomputing Applications/CN=Mark Servilla lifetime=14400
476871	roadrunner.lternet.edu	notice	2006-01-05 21:58:36	METACAT: MetaCat: user [GSI DN: "/C=US/O=National Center for Supercomputing Applications/CN=Mark Servilla", context present] try to login
476870	roadrunner.lternet.edu	notice	2006-01-05 21:58:36	METACAT: MetaCat: --- Action: login
476869	roadrunner.lternet.edu	notice	2006-01-05 21:58:36	METACAT: MetaCat: --- Action: login

Done

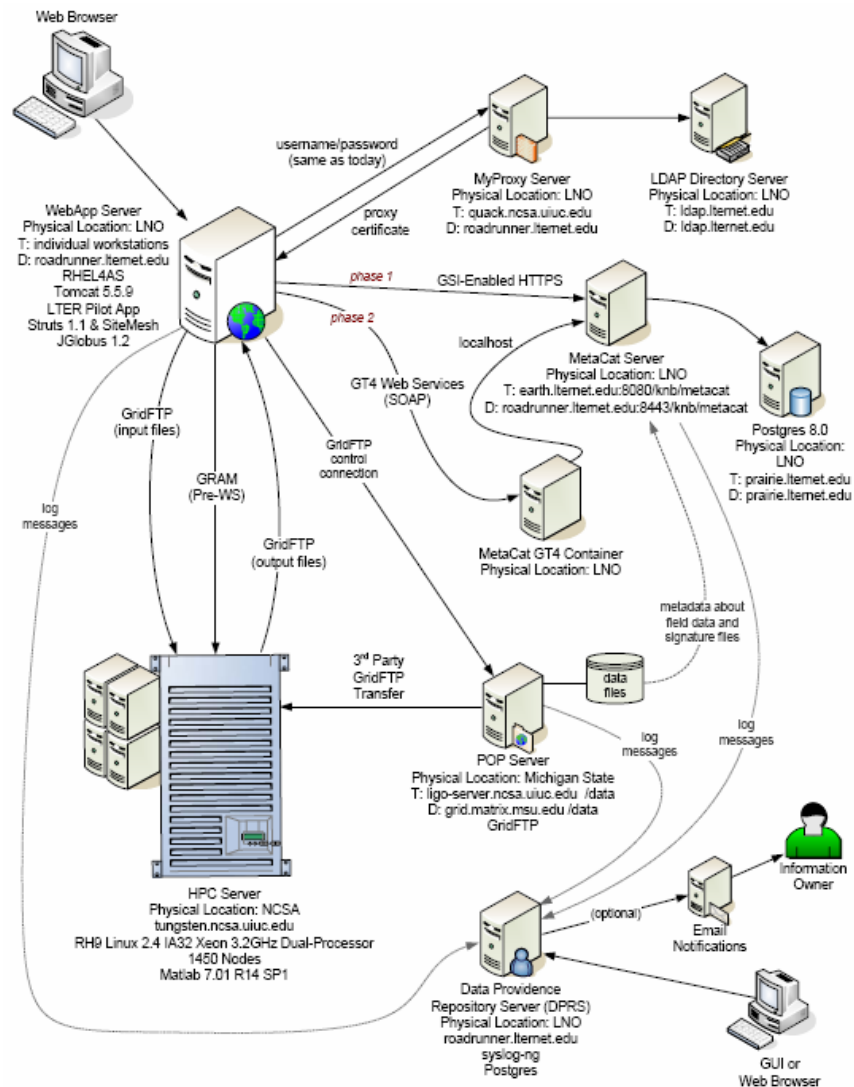
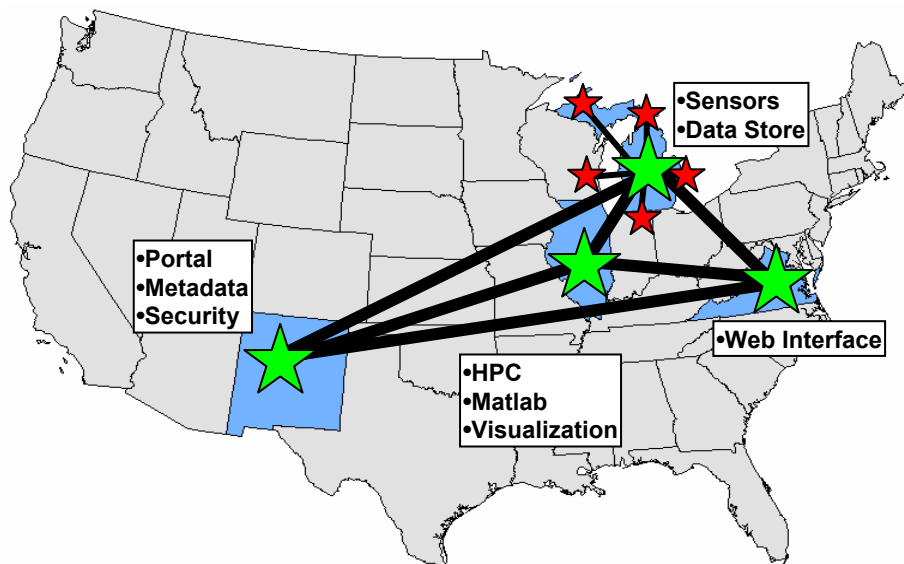
roadrunner.lternet.edu





# Biophony Grid Portal

## System Architecture & Topology





# Other Grid Projects

- TeraGrid




- Network for Earthquake Engineering and Simulation



- Biomedical Informatics Research Network



- EcoGrid





# Grid Organizations

- Globus Alliance
  - *Globus Toolkit™* – Reference implementation of the grid architecture and grid protocols
  - <http://www.globus.org>
- NSF Middleware Initiative (NMI)
  - Supports the design, development, testing, and deployment of middleware for HPC
  - <http://www.nsf-middleware.org>
- GRIDS Center
  - Grid Research Integration Deployment and Support Center – part of NMI
  - <http://www.grids-center.org>
- Global Grid Forum
  - Main standards body governing the world-wide grid community
  - <http://www.globalgridforum.org>





# Recommended Texts

- ***Grid Computing: A Practical Guide to Technology and Applications***
  - Ahmar Abbas
  - Charles River Media © 2004
- ***Introduction to Grid Computing with Globus***
  - Luis Ferreira et al.
  - IBM Redbooks © 2004
- ***Enabling Applications for Grid Computing with Globus***
  - Bart Jacob et al.
  - IBM Redbooks © 2003
- ***Grid Services Programming and Application Enablement***
  - Luis Ferreira et al.
  - IBM Redbooks © 2004





# Acknowledgements

- This material is based upon work supported by:
- The National Science Foundation under Grant Numbers: 0129792 ,9980154, 0225676 and 0072909.
- Collaborators: University of New Mexico (Long Term Ecological Research Network Office), NCEAS (UC Santa Barbara), San Diego Supercomputer Center, University of Kansas (Center for Biodiversity Research), University of Virginia, University of California Berkeley (Hastings Biological Station), University of Wisconsin
- The Andrew W. Mellon Foundation