Welcome to AG Retreat 2002

Rick Stevens

Argonne National Laboratory

University of Chicago

Access Grid Retreat 2002

- Welcome
- AG Progress Report
- The TeraGrid and the future of the Grid
- Updated Vision for the Access Grid
- Standards and the Role[s] of the Private Sector
- New Communities
- My Goals for the Next Two Days

Welcome and Few Announcements

- AG Tutorial – date confirmed 30 May – schedule to be announced but most likely to include Beginners and Advanced Practitioner classes
- Monday evening – Baja Mexican Fest – 6.30pm – Executive Ballroom tickets given at Registration
- AG Community Survey – please complete this if you haven't so we can present the results to the group on Tuesday
- Thanks to Microsoft for hosting the lunches
- Thanks to Insors for sponsoring dinner tonight

AG is a Tool Critical for the Grid's Success

"If I had eight hours to chop down a tree,
I'd spend six sharpening my axe."

-- Abraham Lincoln

AG Progress Report

- Deployment
- Documentation [AGDP]
- Progress towards release 2.0
- OS ecumenicalism [XP, Linux, ...]
- Minimum node spec [AG in one box]
- Virtual Venue services development

Near Term Grid Futures

- Global Grid Forum
- NSF TeraGrid Project
- iVDGL and NSF NMI Grid Center
- DOE Earth Systems Grid
- National Virtual Observatory
- National Earthquake Engineering Simulation Grid
- DOE Science Grid

The Grid Problem

Resource sharing + coordinated problem solving in dynamic, multi-institutional virtual organizations

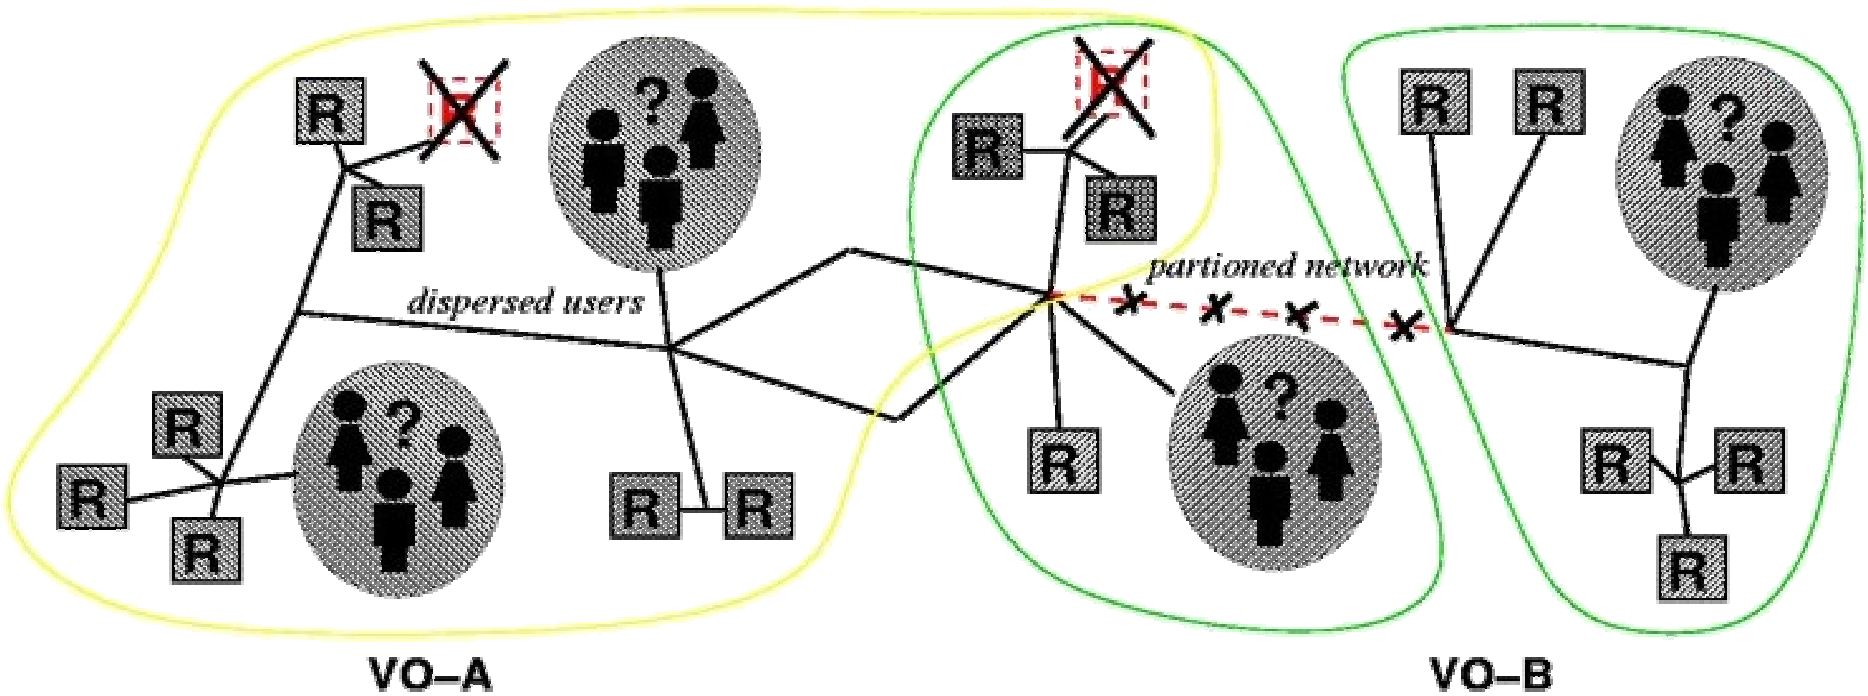

Elements of the Problem

- Resource sharing
  - Computers, storage, sensors, networks, ...
  - Sharing always conditional: issues of trust, policy, negotiation, payment, ...
- Coordinated problem solving
  - Beyond client-server: distributed data analysis, computation, collaboration, ...
- Dynamic, multi-institutional virtual orgs
  - Community overlays on classic org structures
  - Large or small, static or dynamic

Grid Communities & Applications: Data Grids for High Energy Physics

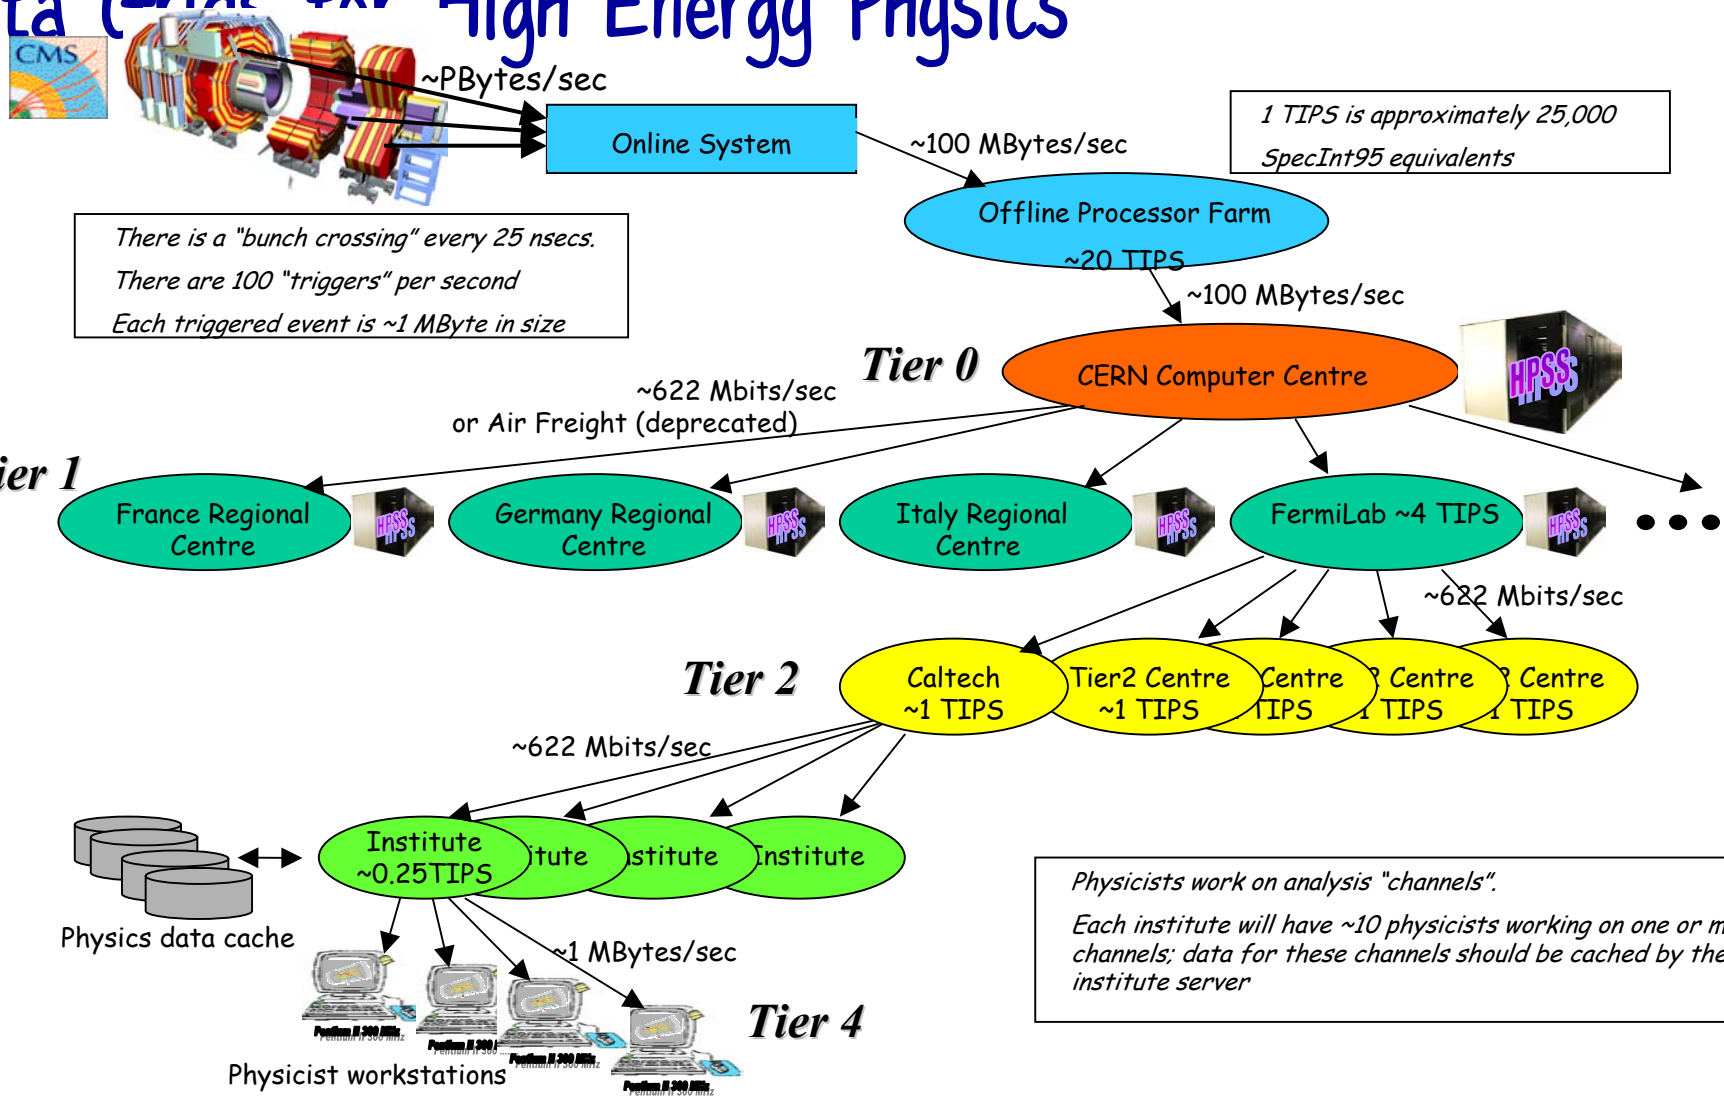

Image courtesy Harvey Newman, Caltech

The I-Way @ SC95 Conference

- First large-scale "modern" Grid experiment
  - provided the basis for modern Grid infrastructure efforts
- I-Way included
  - A Grid of 17 sites connected by vBNS
  - 60+ application groups
  - OC-3 backbone
  - Large-scale use of immersive displays
    - CAVE and I-Desk
- I-Soft programming environment
  - Pioneered security, scheduling ideas
  - Scheduling done with a "human-in-the-loop" [Warren Smith!]

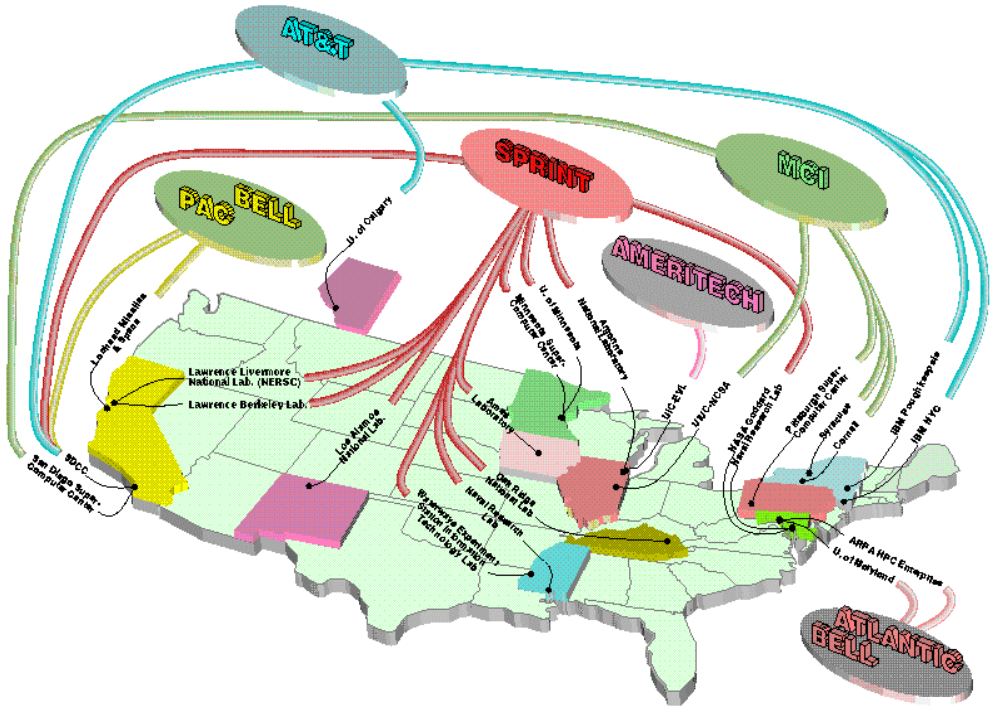

An Emerging Grid Community

1995–2000

- "Grid book" gave a comprehensive view of the state of the art
- Important infrastructure and middleware efforts initiated
  - Globus
  - Legion
  - Condor
  - NetSolve, Ninf
  - Storage Resource Broker
  - Network Weather Service
  - AppLeS, ...

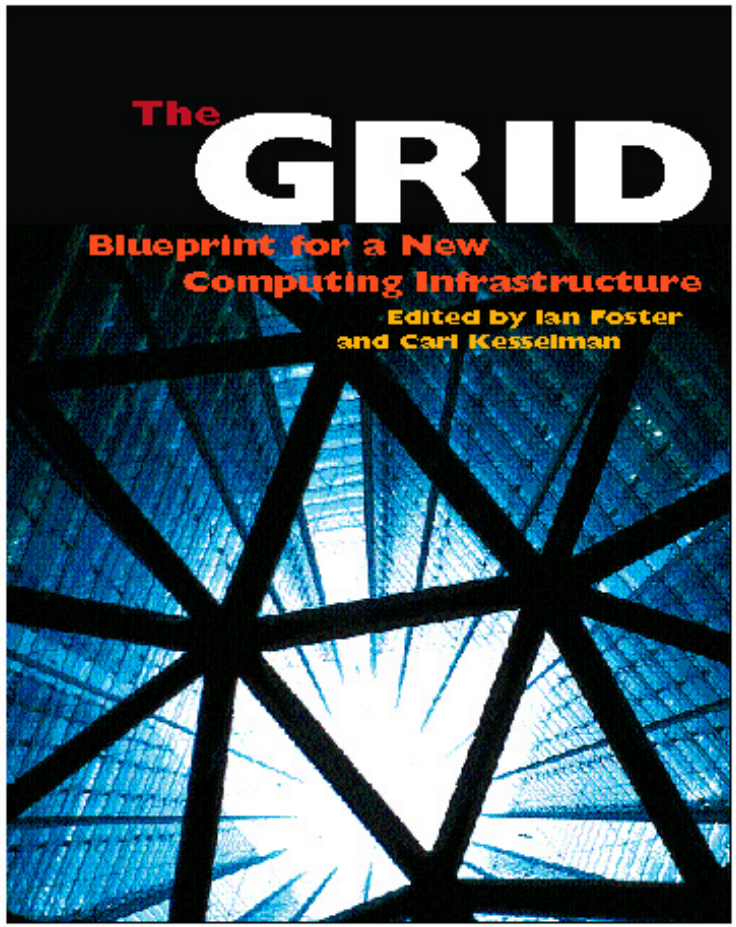

The Globus Project –a Layered Set of Services

- Leads: Ian Foster and Carl Kesselman
- Globus model focuses on providing key Grid services
  - Resource access and management
  - Grid FTP
  - Information Service
  - Security services (authentication, authorization, policy, delegation)
  - Network reservation, monitoring, control

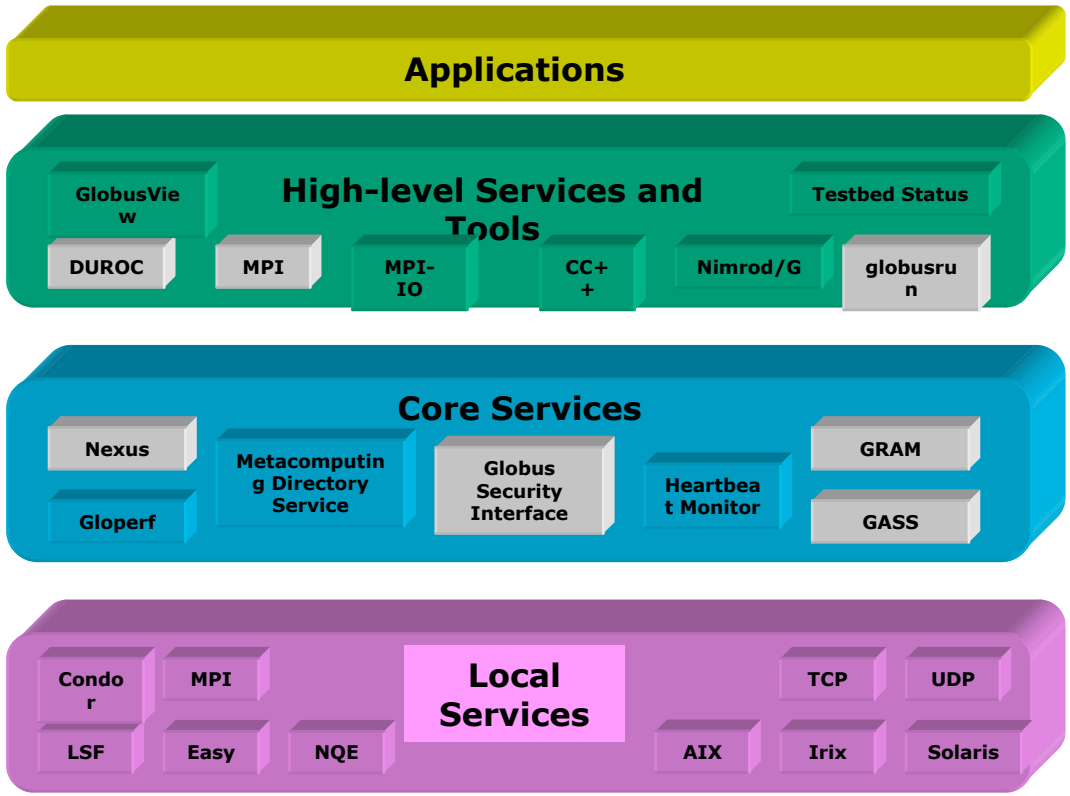

Grid Computing Today

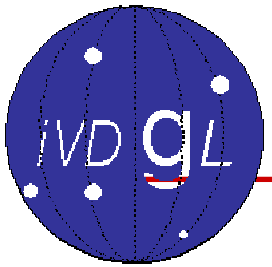

DISCOM

SinRG

APGrid

IPG …

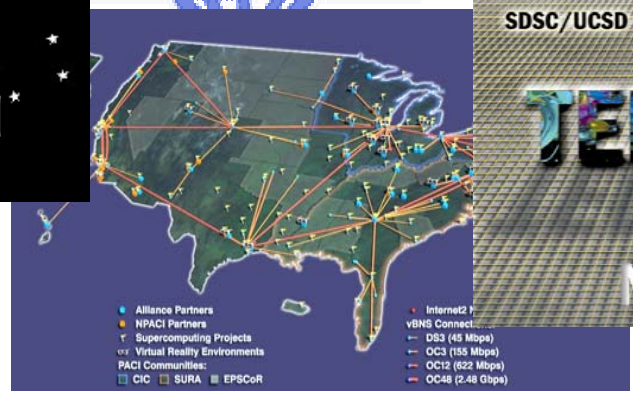

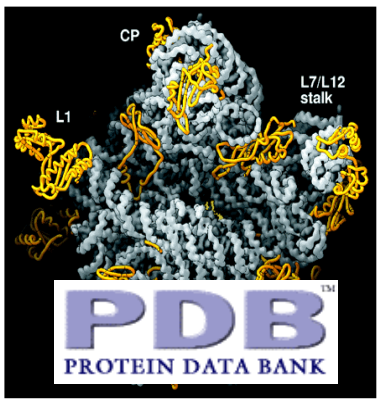

NSF DTF Solicitation

"NSF seeks to open a pathway to future computing, communications and information environments by creating a very large-scale system that is part of the rapidly expanding computational grid."

- RFP specifications
  - multi-site distributed facility
    - single site and Grid-enabled capabilities
    - at least one site with 5+ TF peak performance
    - at least one additional site coupled with the first
    - capabilities for both simulation and data exploration
  - networked system optimized to support the use of data stored at one site by a major computational resource at a geographically distant site
    - ultra high-speed networking
    - sophisticated data handling
  - visualization
  - newer generation processors and HPC equipment
  - fully coordinated with PACI resources and activities
  - production quality service

What We Proposed

- Distributed, multisite facility
  - single site and "Grid enabled" capabilities
    - uniform compute node selection and interconnect networks at 4 sites
    - central "Grid Operations Center"
  - at least one 5+ teraflop site and newer generation processors
    - SDSC at 4+ TF, NCSA at 6.1–8 TF with McKinley processors
  - at least one additional site coupled with the first
    - four core sites: SDSC, SDSC, ANL/UC, and Caltech
- Ultra high-speed networks
  - multiple gigabits/second
    - modular 40 Gb/s backbone (4 x 10 GbE)
- Remote visualization
  - data from one site visualized at another
    - high-performance commodity rendering and visualization system
    - Argonne hardware visualization support
    - data serving facilities and visualization displays

TeraGrid Cyberinfrastructure

- Not just a supercomputer, but rather
  - an integrated hardware and software environment
    - terascale computing, and data management
    - visualization and ultra high-speed networking
- Features
  - distributed instruments and Grid computing at terascale
  - large-scale data management and mining an integral component
  - commodity hardware for broad scalability and replication
  - open source software for community development
- Newer generation microprocessors
  - second generation McKinley IA-64
  - price/performance and open source software
- Leverages expertise and PACI partners
  - software development will be required [as noted by the solicitation]

PACI 13.6 TF Linux TeraGrid

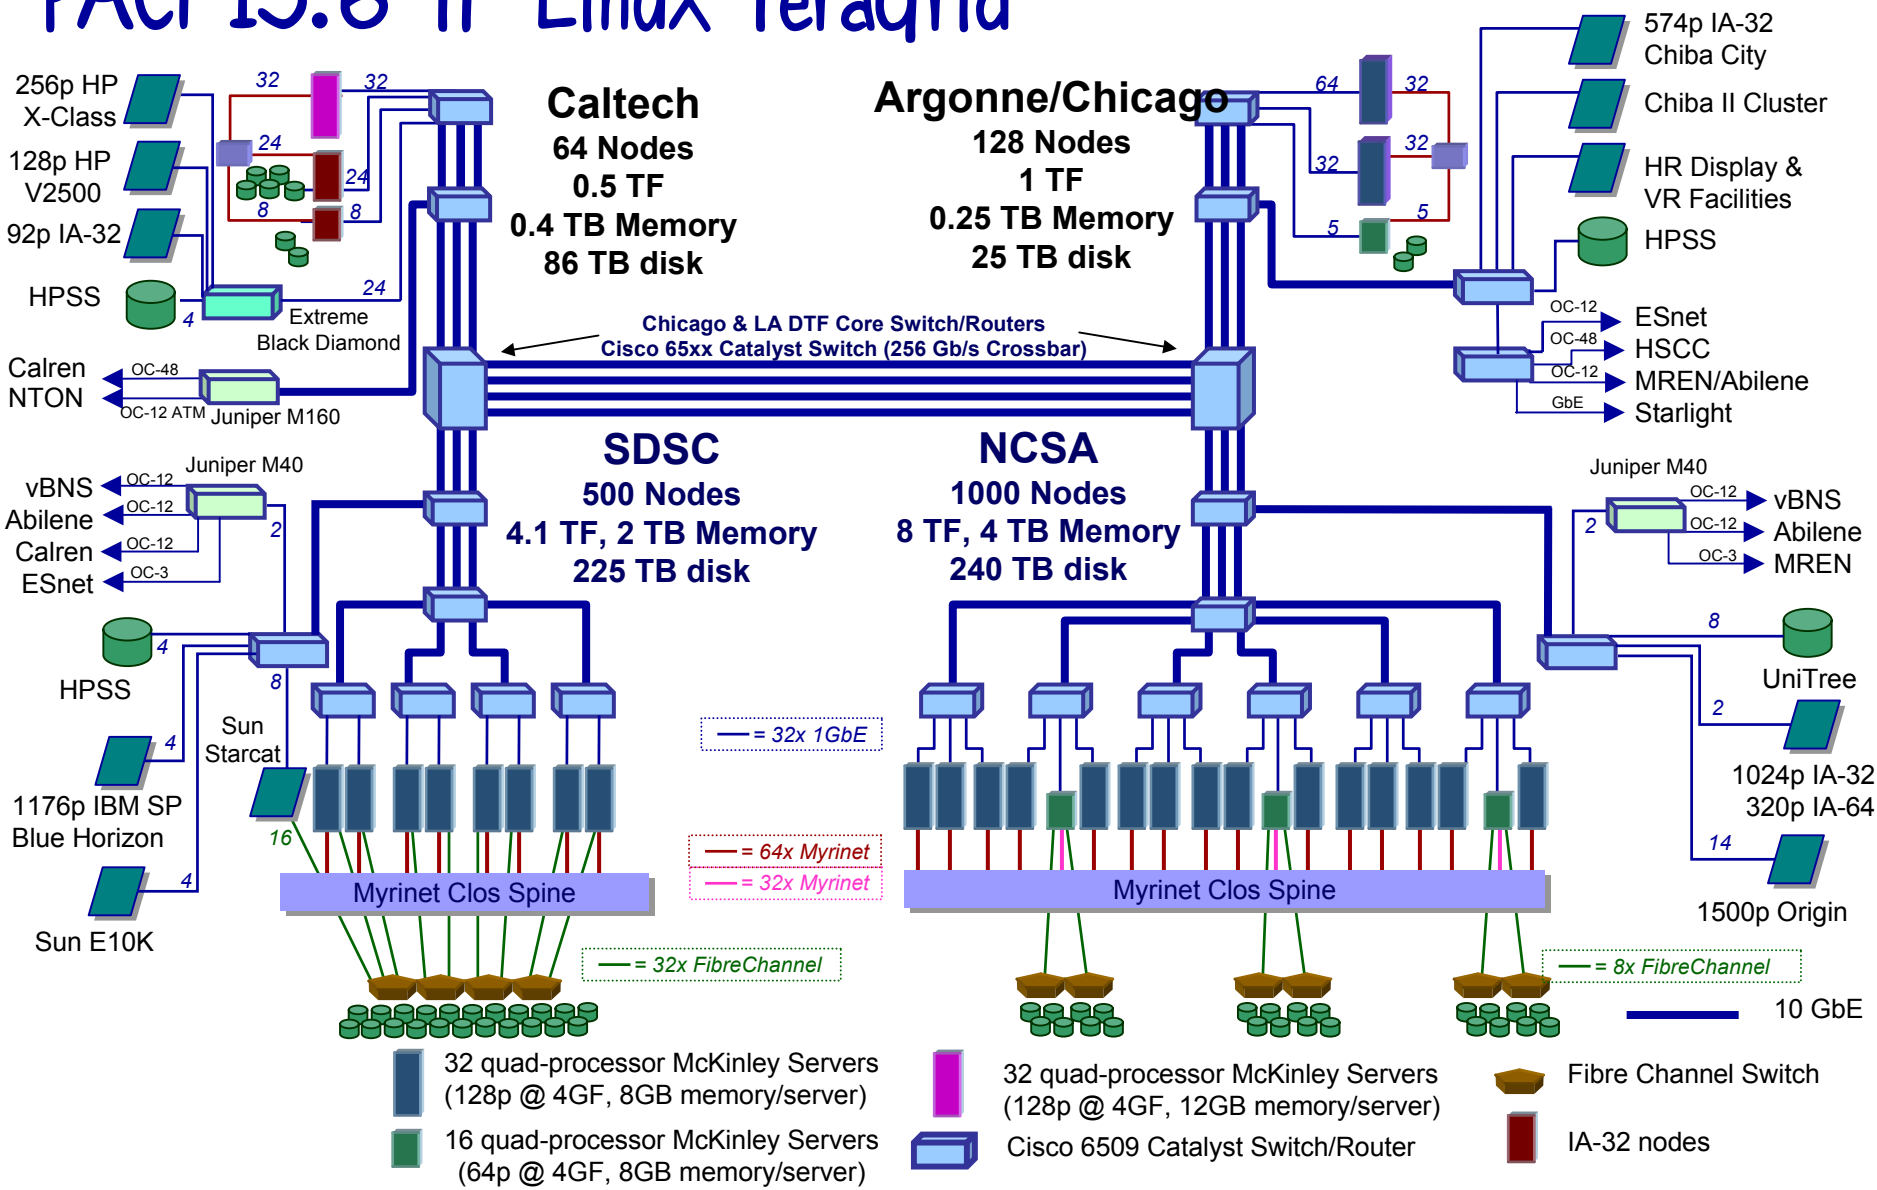

TeraGrid Wide Area Network

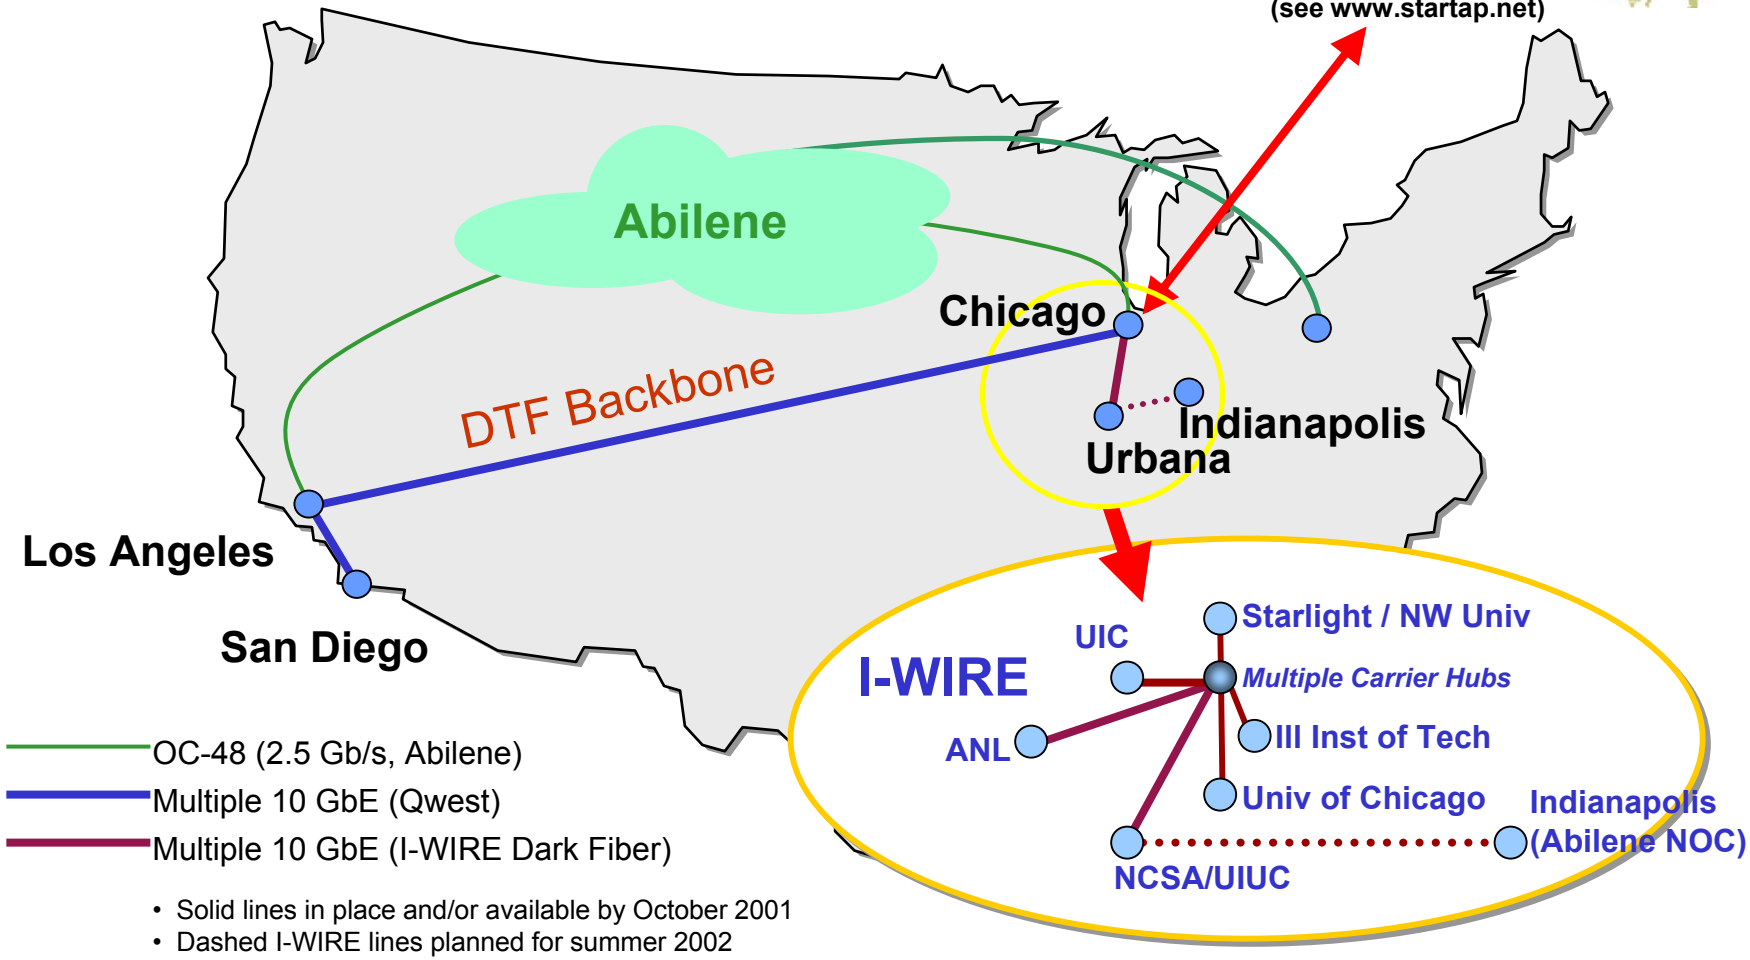

Accelerated Visualization

- 128-256 augmented nodes on ANL cluster
  - nvidia Geforce4 graphics processor
- Remote visualization server software
  - networked parallel multipass volume render
    - remote path planning and job monitoring
    - out of core support for TB class datasets
  - CorridorOne: high-performance remote visualization
    - streaming output to remote tiled displays
  - VisBench: web portal for remote visualization

Application Community Investment

- Reducing application development risk
  - NSF ITR GriPhyN [Grid Physics Network]
    - CMS and ATLAS Large Hadron Collider experiments
    - Sloan Digital Sky Survey [SDSS]
    - LIGO gravity wave detector
  - NSF NEESgrid
    - National Earthquake Engineering Simulation grid
  - NASA Information Power Grid [IPG]
    - distributed computation and resource access
  - UK E-science Grid
    - computational/data grids, middleware and hardware
  - EU Data Grid
    - computation and distributed terabyte/petabyte data bases
- And there are more coming every month
  - NSF MRE projects
  - DOE SciDAC projects

TeraGrid Scientific Impact

- Multiple classes of user support
  - each with differing implementation complexity
    - minimal change from current practice
    - new models, software, and applications
- Benefit to three user communities
  - existing supercomputer users
    - new capability [FLOPS, memory, and storage]
  - data-intensive and remote instrument users
    - linked archives, instruments, visualization and computation
    - several communities already embracing this approach
      - GriPhyN, BIRN, Sloan DSS/NVO, BIMA, ...
  - future users of MRE and similar facilities
    - DTF is a prototype for ALMA, NEESGrid, LIGO, and others

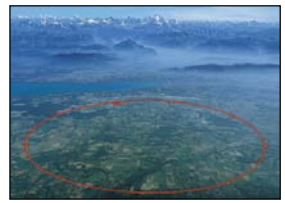

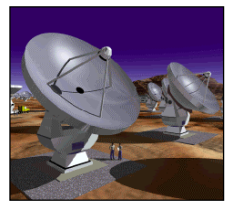

NEON
National Ecological Observatory Network

iVDGL

- International Virtual-Data Grid Laboratory
  - A place to conduct Data Grid tests at scale
  - Concrete manifestation of world-wide grid activity
  - Continuing activity that will drive Grid awareness
- Scale of effort
  - For national, intl scale Data Grid tests, operations
  - Computation & data intensive computing
- Who
  - Initially US-UK-Italy-EU; Japan, Australia
  - & Russia, China, Pakistan, India, South America?
  - StarLight and other international networks vital

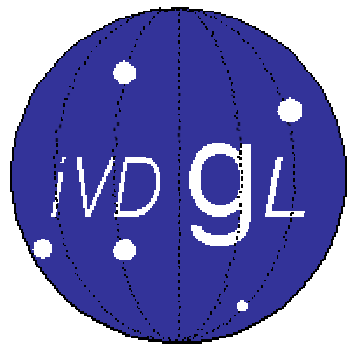

U.S. Co-PIs: Avery, Foster, Gardner, Newman, Szalay

iVDGL Map Circa 2002-2003

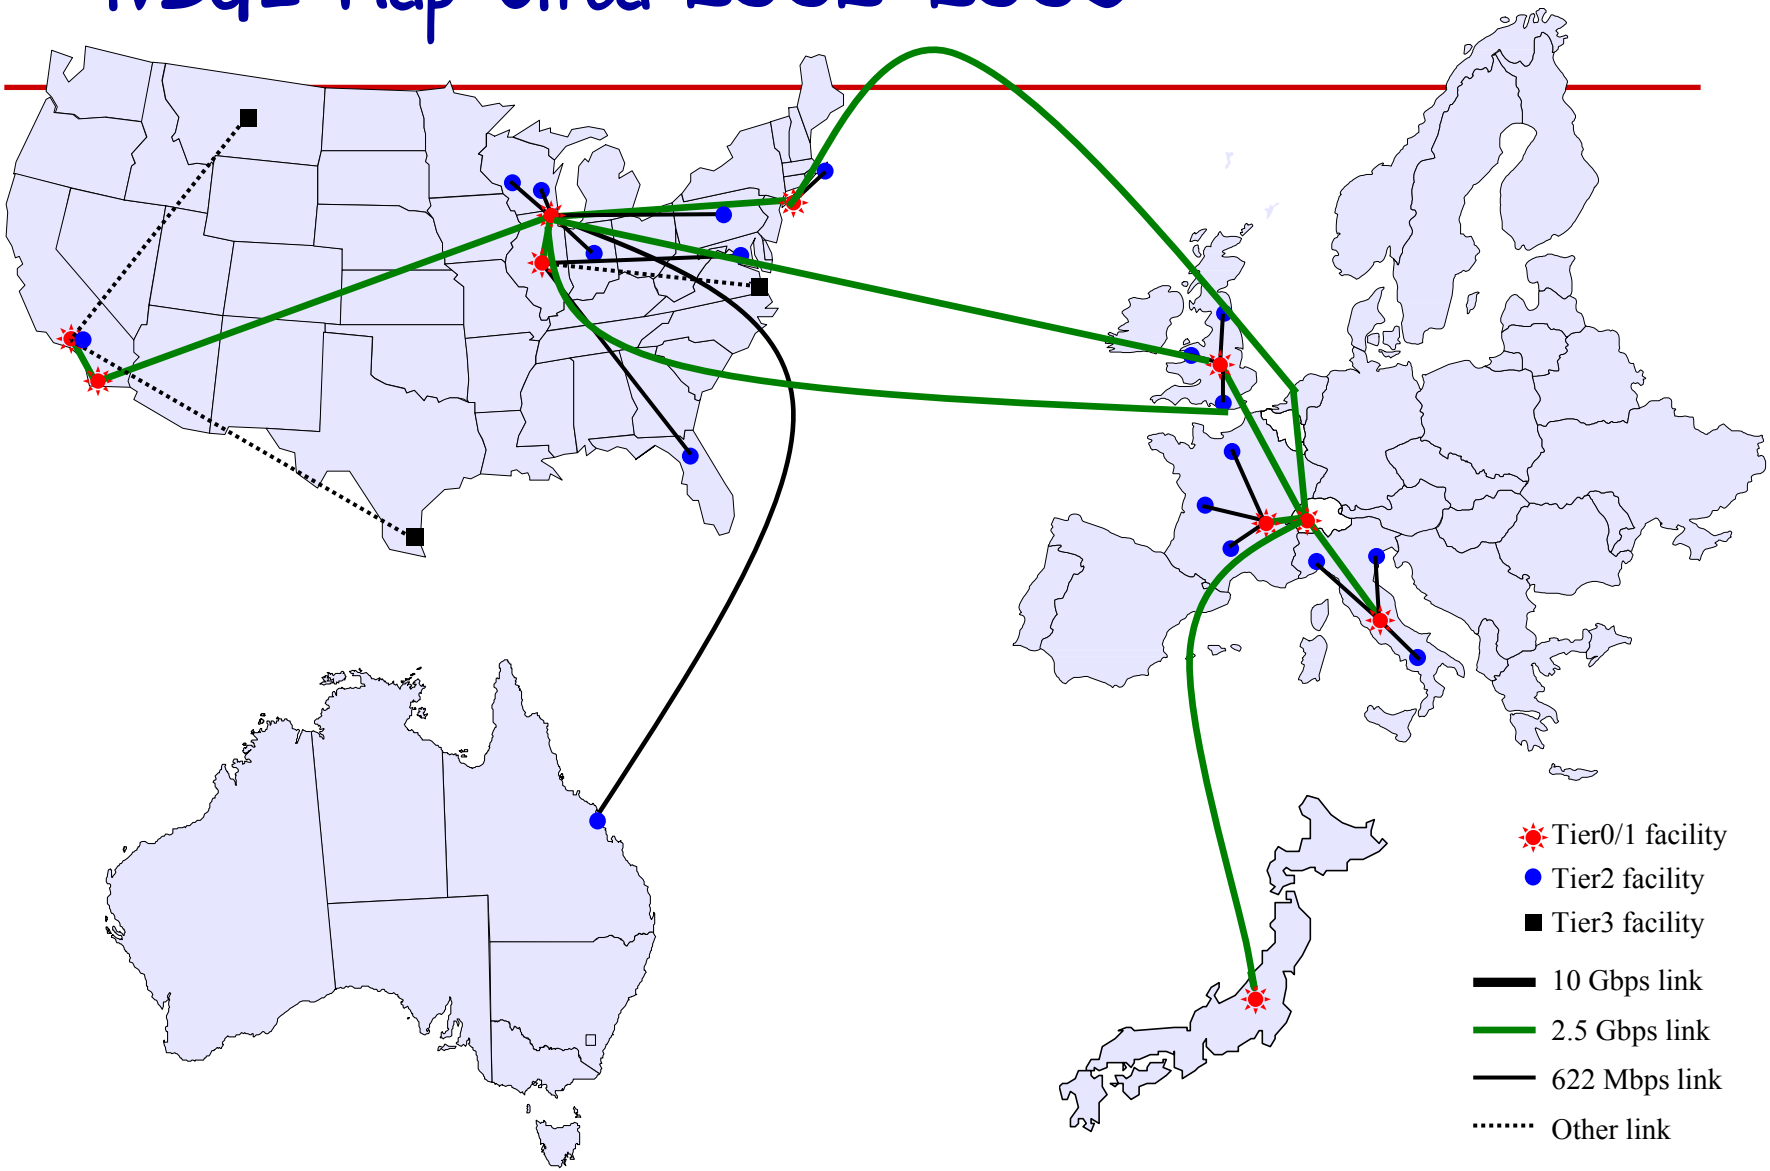

R. Stevens

Argonne National Laboratory + University of Chicago

An International Systems Biology Grid

- A Data, Experiment and Simulation Grid Linking:
  - People [biologists, computer scientists, mathematicians, etc.]
  - Experimental systems [arrays, detectors, MS, MRI, EM, etc.]
  - Databases [data centers, curators, analysis servers]
  - Simulation Resources [supercomputers, visualization, desktops]
  - Discovery Resources [search servers perhaps optimized]
  - Education and Teaching Resources [classrooms, labs, etc.]
- Different than and more fine grain than current Grid Projects
  - More laboratory integration [need small laboratory software interfaces]
  - Most of the participants will be experimentalists [team workflow, visualization]
  - More diversity of data sources and databases [integration, federation]
  - More portals to simulation environments [ASP models]

The Grid will have Broad and Global Impact

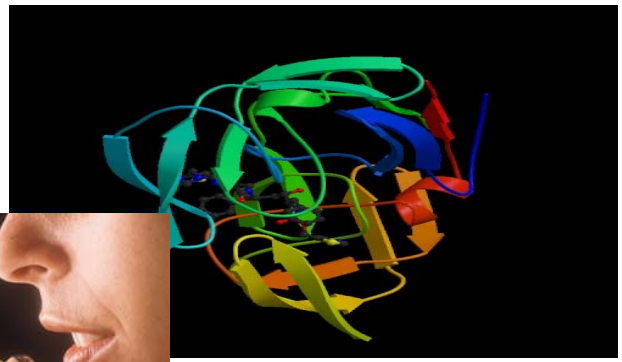

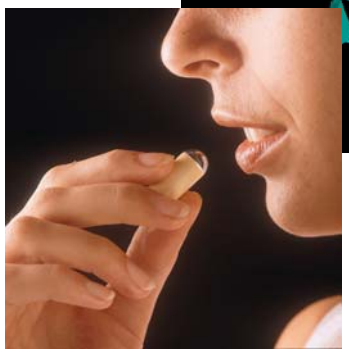

Personalized Medicine

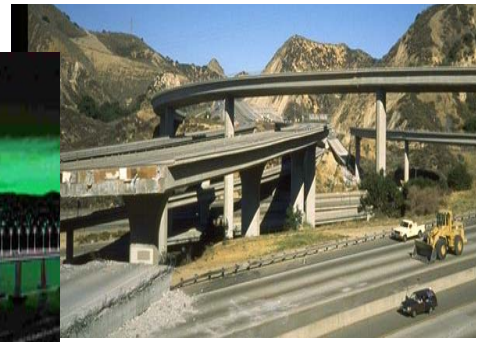

Sensors

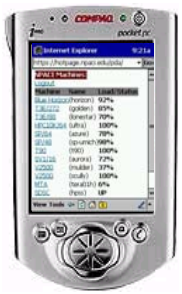

Wireless networks

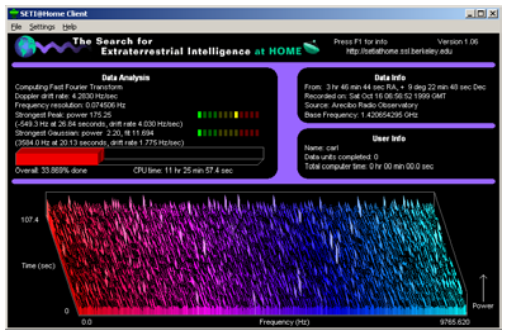

Knowledge from Data

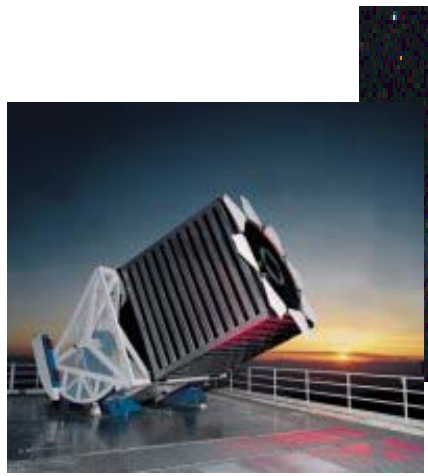

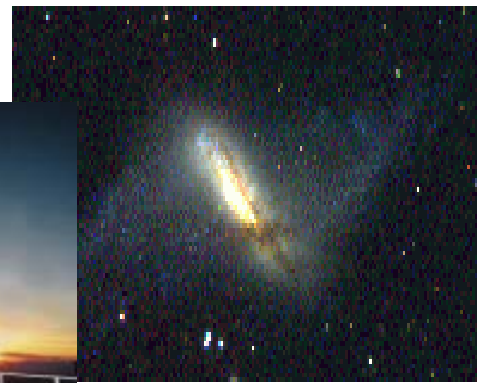

Instruments

Future Vision for the AG

- Integration with Grid services
  - Globus and integrated web services
- Drive the high-end but remain affordable
  - video and audio
- Interoperability with desktop based systems
  - via gateways at the network level
- Maintain focus on small group-to-group work environments
- Explore issues relating to advanced displays and high-throughput networking

Thinking about Future Venues

- Transitioning from thinking of Venues as scarce resources to ubiquitous and persistent resources
  - User, Group, Node, Site
  - Scheduling vs Reservation vs Navigation
  - Personal
  - Persistent
- 1995 slides on virtually free travel
  - Place to store things, help in finding folks and resources
  - Security and automatic (cross platform?) usability

Standards and the Role(s) of the Private Sector

- Standards via the Global Grid Forum
  - AG services APIs and Protocols (AG 2.0)
  - Interfaces to Grid services needed for collaboration
- PS can effectively address needs of the commercial sector for support and services
- PS provided expanded development resources focused on ease of use and maintainability
- PS is critical for large-scale deployments and network services
- Long term growth of the AG requires PS investment

New Research Communities

- Internet2 universities
- DOE Office of Science labs
- International Grid Projects
- Biology and Bioinformatics Centers
- Nanoscience and Nanotechnology Centers
- Medical Research and Clinical Centers
- Earth Science and Oceanography
- Paleontology Laboratories

Goals for the Next Two Days

- Current and future plans for the Access Grid
- AG related technologies and architectures
- Strategies for AG development and growth
- Obtain feedback from the AG community
- Engage community in applications discussions
- Discuss roles of standards in AG development