Building Shared Collections Using the Storage Resource Broker

Reagan W. Moore
moore@sdsc.edu
http://www.sdsc.edu/srb

Storage Resource Broker

- Data grid middleware
  - **Organize distributed data into shared collections.**
  - **Support access through**
    - **C library calls**
    - **Java class libraries and GridSphere portal**
    - **Python/Perl load libraries**
    - **Interactive browsers (Web, Perl, PHP, Windows)**
    - **Digital libraries (DSpace, Fedora).**
- Manage properties of the shared collection needed by
  - **Preservation environments**
  - **Digital libraries**
  - **Real-time sensor systems**
  - **Secure data management environments.**
- Used in production
  - **SDSC collections**
  - **Internationally shared collections**

UCSD

SDSC

Using a Data Grid – *in Abstract*

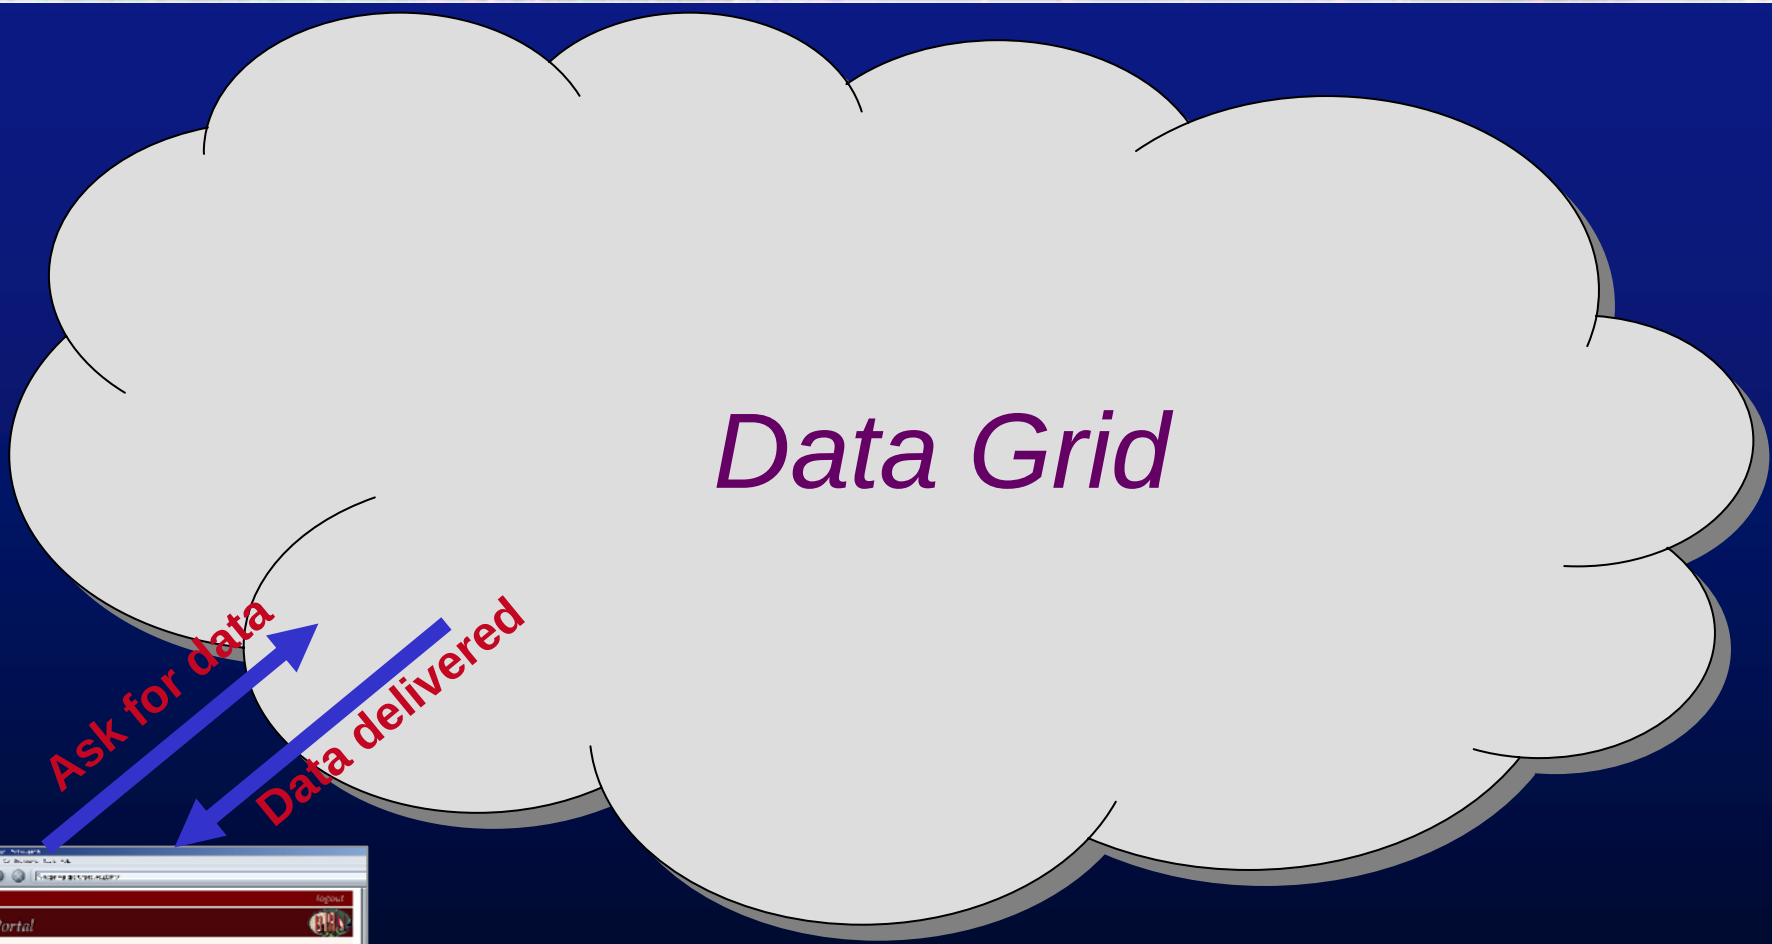

- User asks for data from the data grid
- The data is found and returned
  - Where & how details are hidden

UCSD

SDSC

Using a Data Grid - *Details*

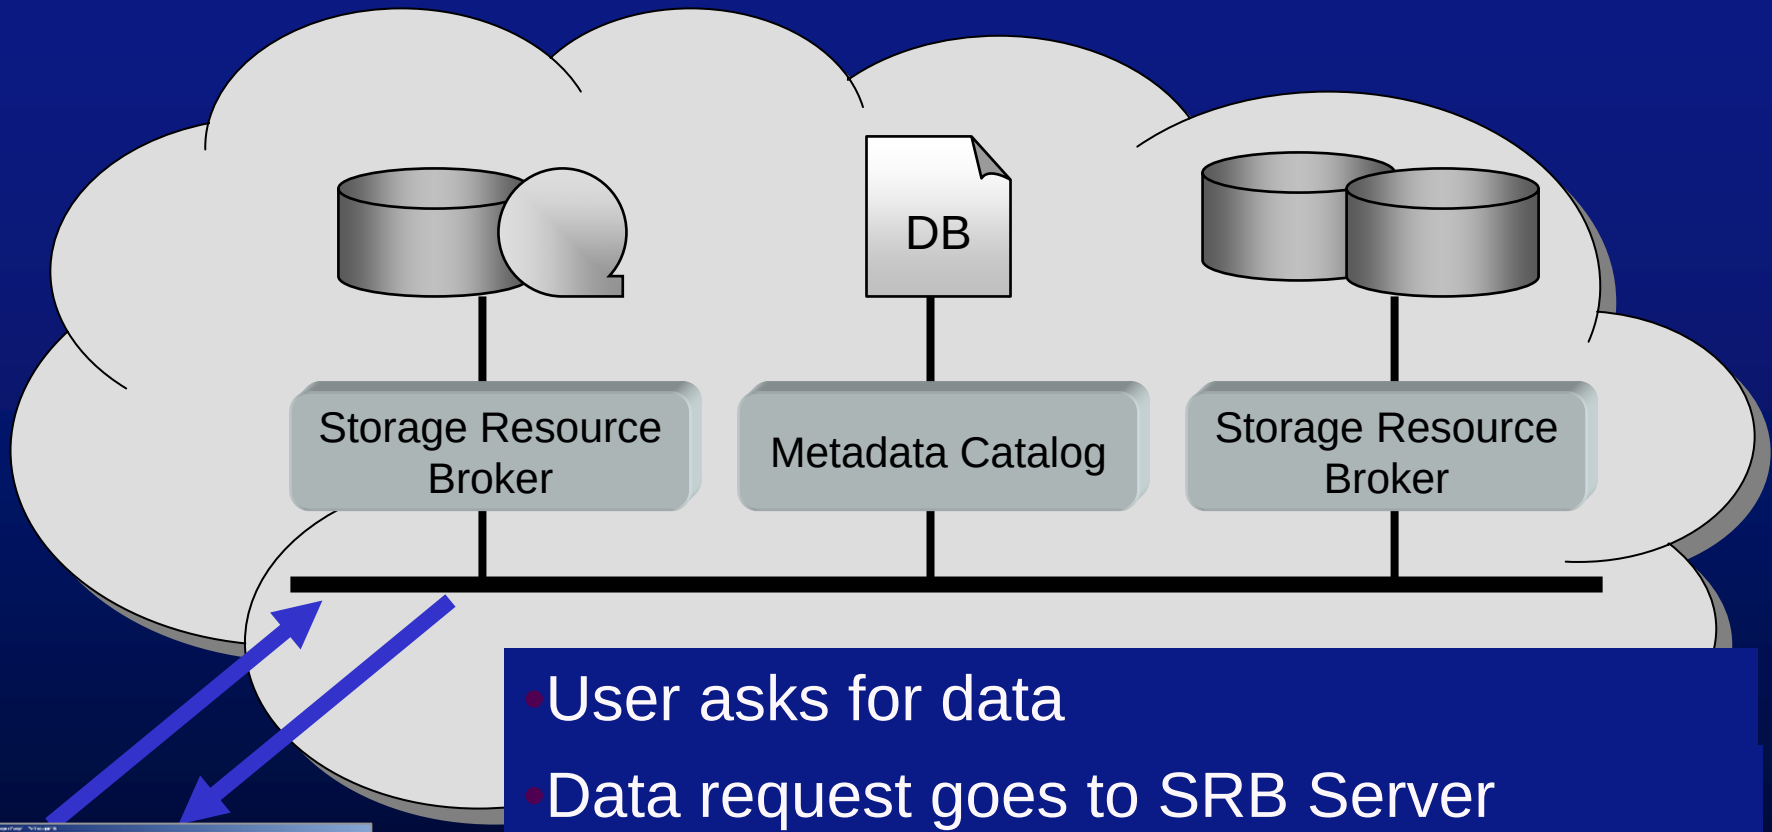

- User asks for data
- Data request goes to SRB Server
- Server looks up data in catalog
- Catalog tells which SRB server has data
- 1st server asks 2nd for data
- The data is found and returned

UCSD

SDSC

Using a Data Grid - *Details*

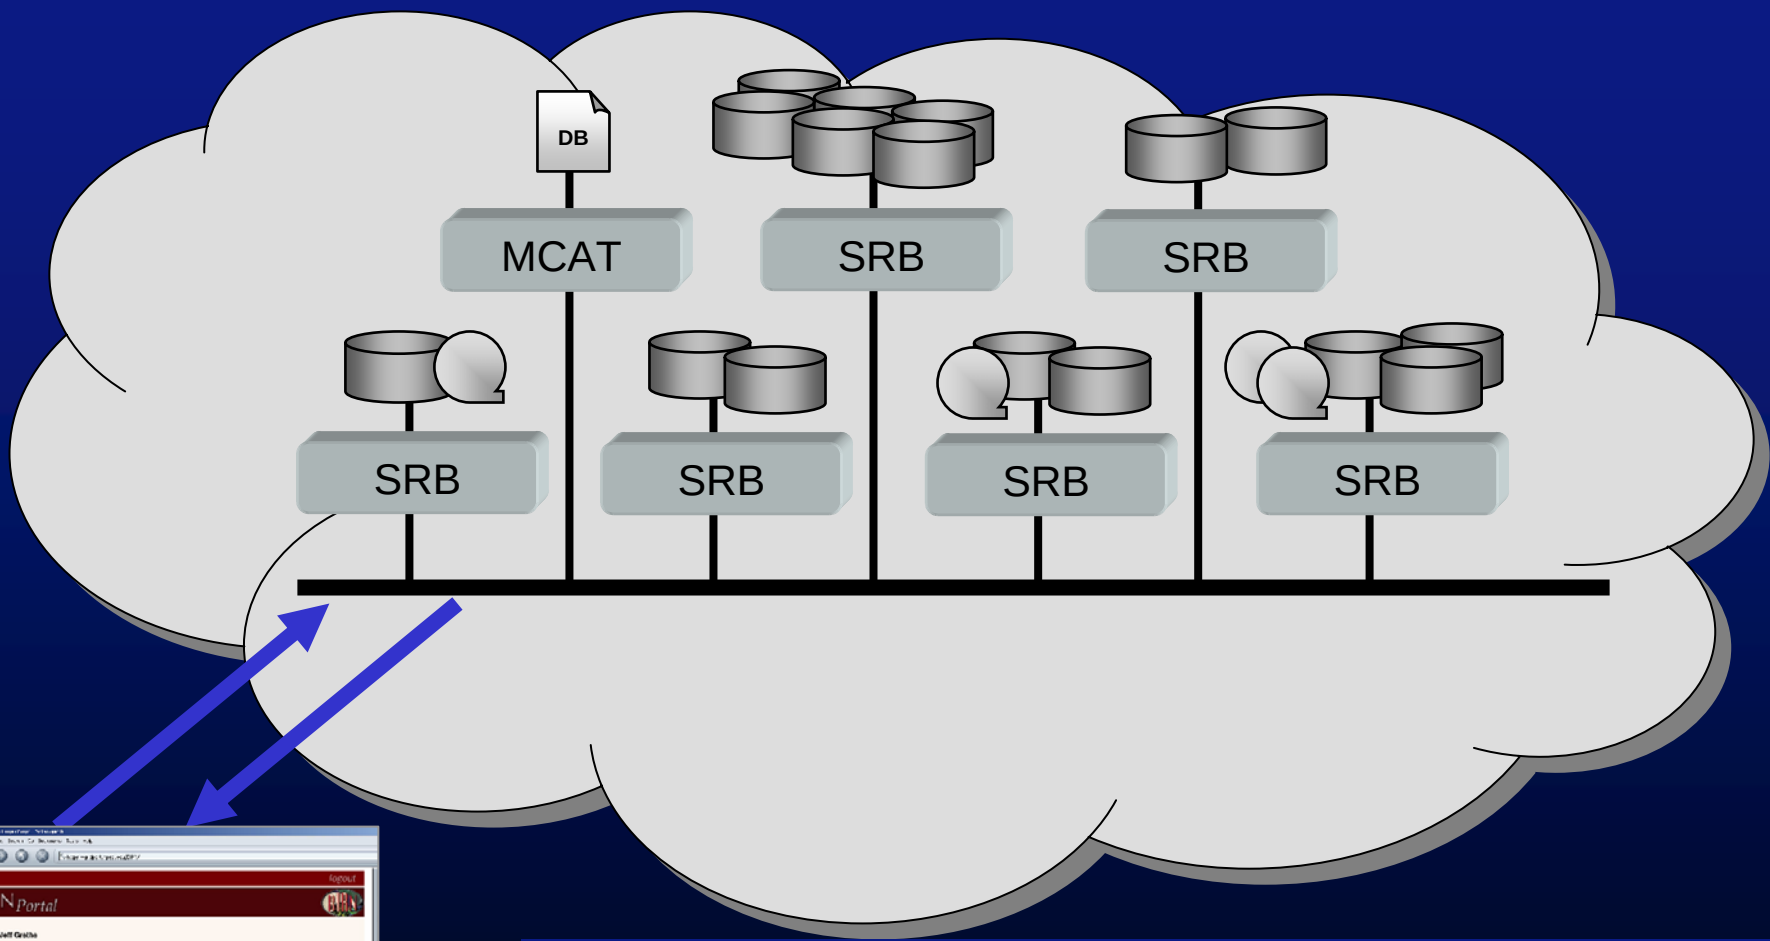

- Data Grid has arbitrary number of servers
- Complexity is hidden from users

UCSD

SDSC

Shared Collections

- **Purpose of SRB data grid is to enable the creation of a collection that is shared between academic institutions**
  - Register digital entity into the shared collection
  - Assign owner, access controls
  - Assign descriptive, provenance metadata
  - Manage state information
    - Audit trails, versions, replicas, backups, locks
    - Size, checksum, validation date, synchronization date, …
  - Manage interactions with storage systems
    - Unix file systems, Windows file systems, tape archives, …
  - Manage interactions with preferred access mechanisms
    - Web browser, Java, WSDL, C library, …

Shared Collections

- **Data grids** support the creation of shared collections that may be distributed across multiple institutions, sites, and storage systems.
- **Digital libraries** publish data, and provide services for discovery and display
- **Persistent archives** preserve data, managing the migration to new technology
- **Real-time sensor systems** federate name spaces across independent environments

Date	5/17/02		6/30/04			5/5/06		
Project	GBs of data stored	1000Õs o files	GBs of data stored	1000Õs o files	Users with ACLs	GBs of data stored	1000Õs o files	Users with ACLs
Data Grid								
NSF / NVO	17,800	5,139	51,380	8,690	80	100,990	13,217	100
NSF / NPACI	1,972	1,083	17,578	4,694	380	34,830	7,239	380
Hayden	6,800	41	7,201	113	178	8,013	161	227
Pzone	438	31	812	47	49	23,099	13,287	68
NSF / LDAS-SALK	239	1	4,562	16	66	115,178	146	67
NSF / SLAC-JCSG	514	77	4,317	563	47	17,095	1,775	55
NSF / TeraGrid			80,354	685	2,962	202,226	4,443	3,267
NIH / BIRN			5,416	3,366	148	16,288	15,306	361
Digital Library								
NSF / LTER	158	3	233	6	35	236	34	36
NSF / Portal	33	5	1,745	48	384	2,620	53	460
NIH / AfCS	27	4	462	49	21	733	94	21
NSF / SIO Explorer	19	1	1,734	601	27	2,605	1,121	27
NSF / SCEC			15,246	1,737	52	167,140	3,471	73
Persistent Archive								
NARA	7	2	63	81	58	2,916	2,004	58
NSF / NSDL			2,785	20,054	119	5,653	50,600	136
UCSD Libraries			127	202	29	190	208	29
NHPRC / PAT						1,338	519	28
TOTAL	28 TB	6 mil	194 TB	40 mil	4,635	701 TB	113 mil	5,393

Biomedical Informatics Research Network
BIRN Data Grid

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Mark Ellisman

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

National Science Digital Library

- **URLs for educational material for all grade levels registered into repository at Cornell**
- **SDSC crawls the URLs, registers the web pages into a SRB data grid, builds a persistent archive**
  - 750,000 URLs
  - 13 million web pages
  - About 3 TBs of data

UCSD

SDSC

View All Metadata

Data Object: **oai.nsdl.org.GROW.70.html**
Parent Collection: **/home/nsdl.sdsc/2003-06-10T13:53:37Z/oai:nsdl.org:GROW:70**
Owner: **nsdl@sdsc**

Control of water pollution from agriculture - FAO irrigation and drainage paper 55

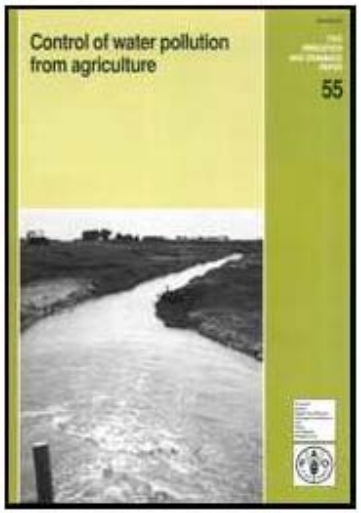

Table of Contents

Southern California Earthquake Center

SCEC Digital Library

- Intuitive User Interface
  - –Pull-Down Query Menus
  - –Graphical Selection of Source Model
  - –Clickable LA Basin Map (Olsen)
  - –Seismogram/History extraction (Olsen)
- Access SCEC Digital Library
  - –Data stored in a data grid
  - –Annotated by modelers
  - –Standard naming convention
  - –Automated extraction of selected data and metadata
  - –Management of visualizations

Select Receiver (Lat/Lon)

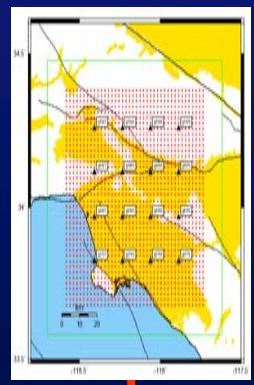

Select Scenario
Fault Model
Source Model

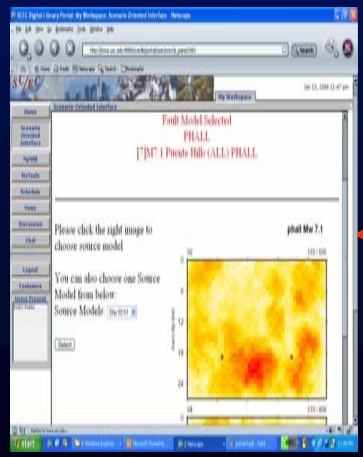

SCEC Community Library

Output
Time History
Seismograms

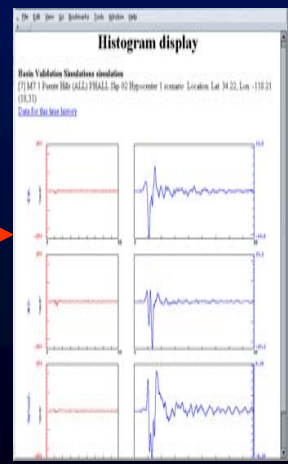

UCSD

SDSC

Terashake Data Handling

- **Simulate 7.7 magnitude earthquake on San Andreas fault**
  - 50 Terabytes in a simulation
  - Move 10 Terabytes per day
- **Post-Processing of wave field**
  - Movies of seismic wave propagation
  - Seismogram formatting for interactive on-line analysis
  - Velocity magnitude
  - Displacement vector field
  - Cumulative peak maps
  - Statistics used in visualizations
  - Register derived data products into SCEC digital library

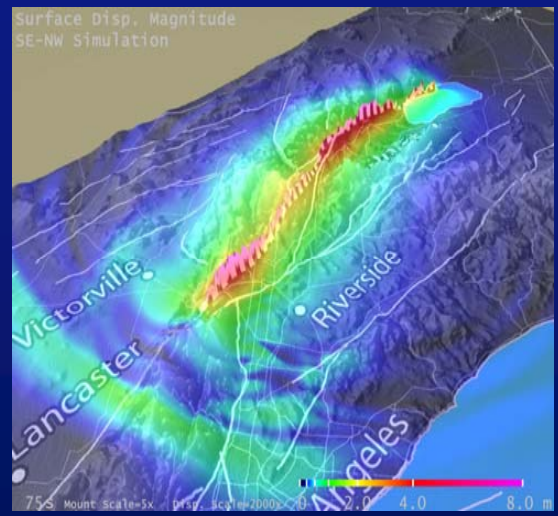

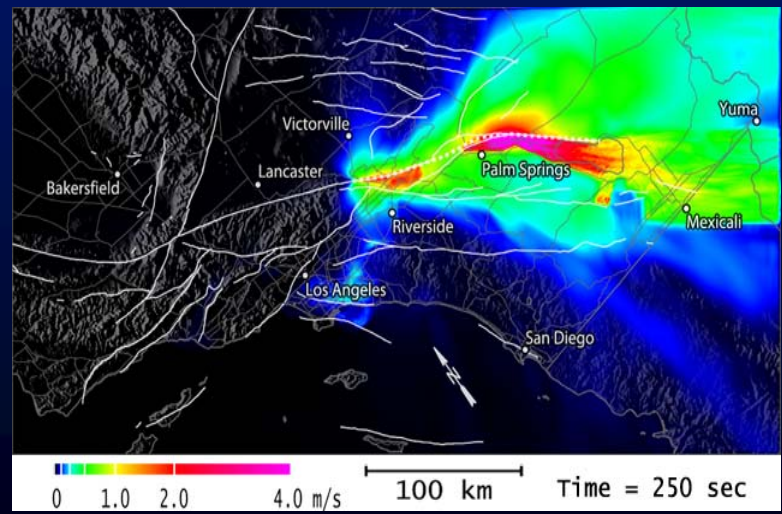

UCSD

SDSC

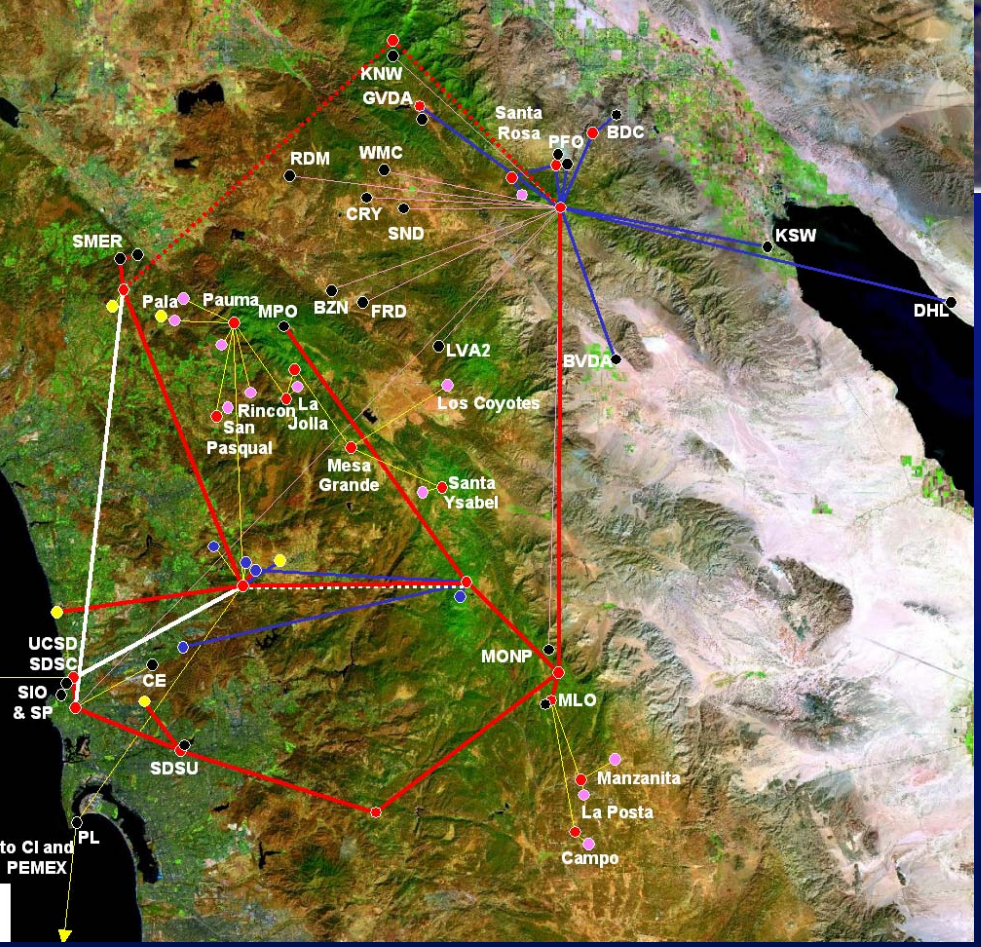

Humidity
Climate
Ecological
Wireless
Oceanography

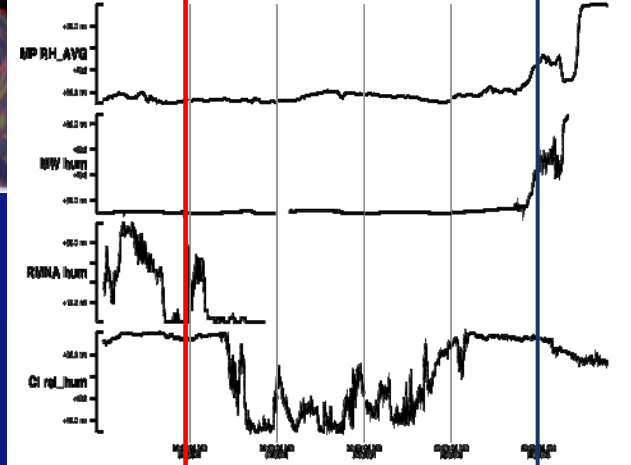

Wind Speed
Climate
Ecological
Wireless
Oceanography

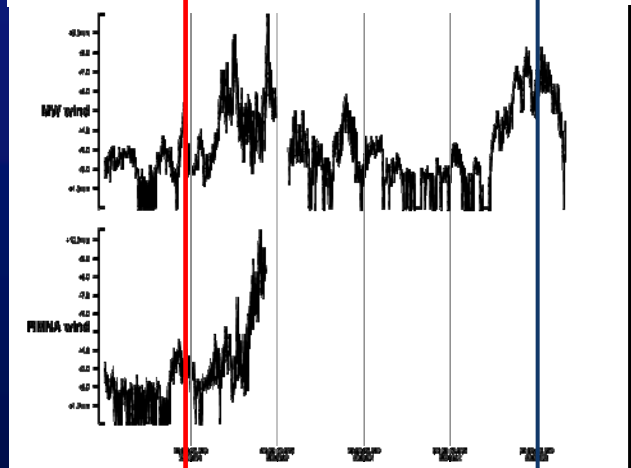

ROADNet Sensor Network Data Integration

Seismic
Geophysics

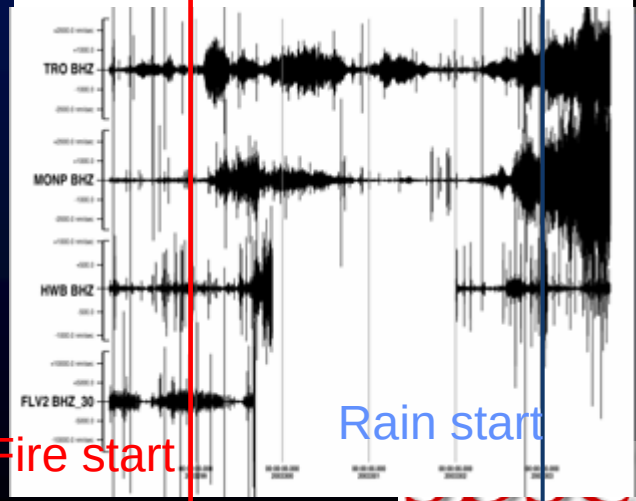

Frank Vernon - UCSD/SIO

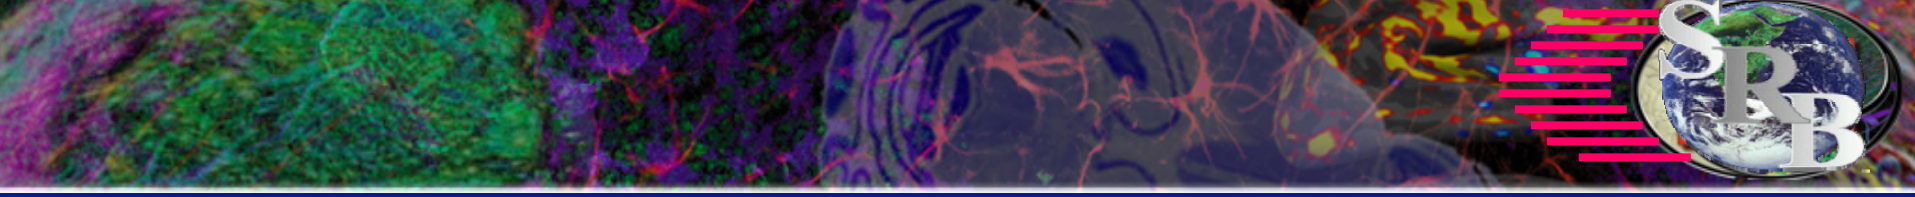

UCSD

SDSC

Logical Name Spaces

Data Access Methods (C library, Unix, Web Browser)

Data Collection

Storage Repository

- Storage location
- User name
- File name
- File context (creation date,…)
- Access constraints

Data Grid

- Logical resource name space
- Logical user name space
- Logical file name space
- Logical context (metadata)
- Control/consistency constraints

Data is organized as a shared collection

UCSD

SDSC

Federation Between Data Grids

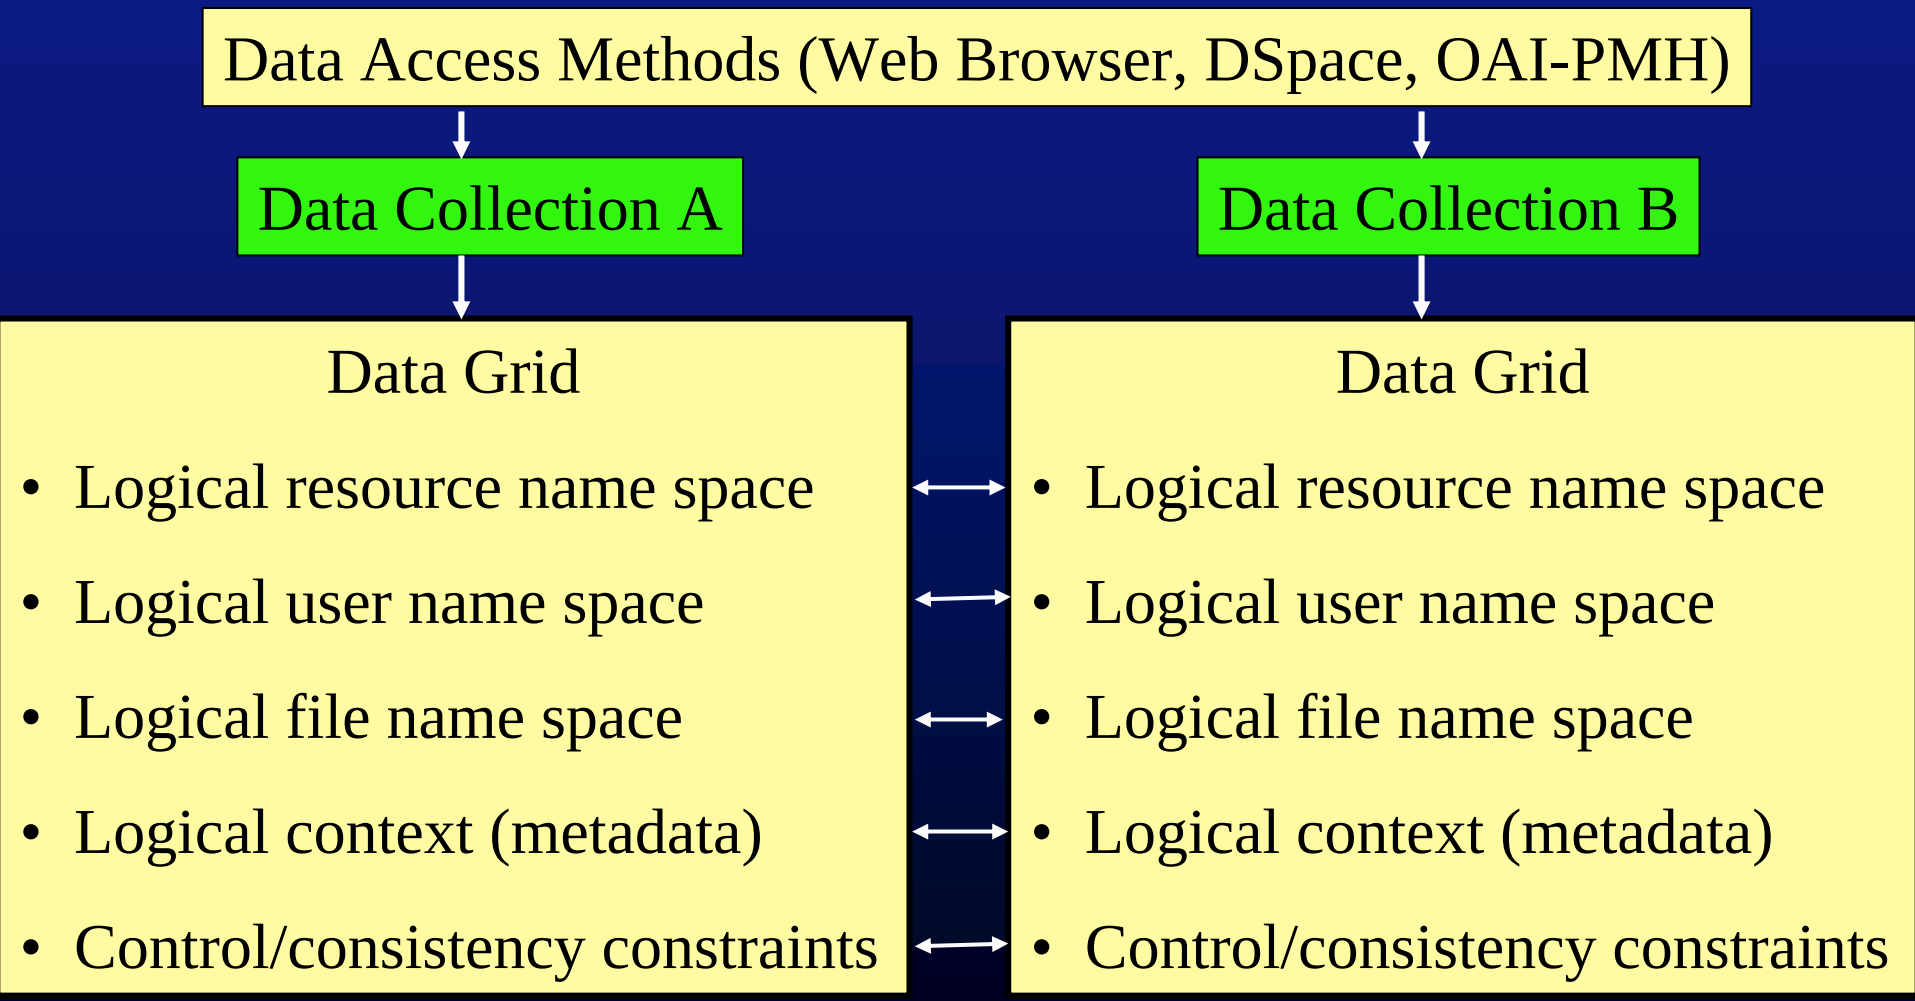

Access controls and consistency constraints
on cross registration of digital entities

NOAO Astronomy Data Grid

- Chile
- Tucson, Arizona
- NCSA, Illinois

- A functioning international Data Grid for Astronomy

Manchester-SDSC mirror

Moved over 400,000 images

UCSD

SDSC

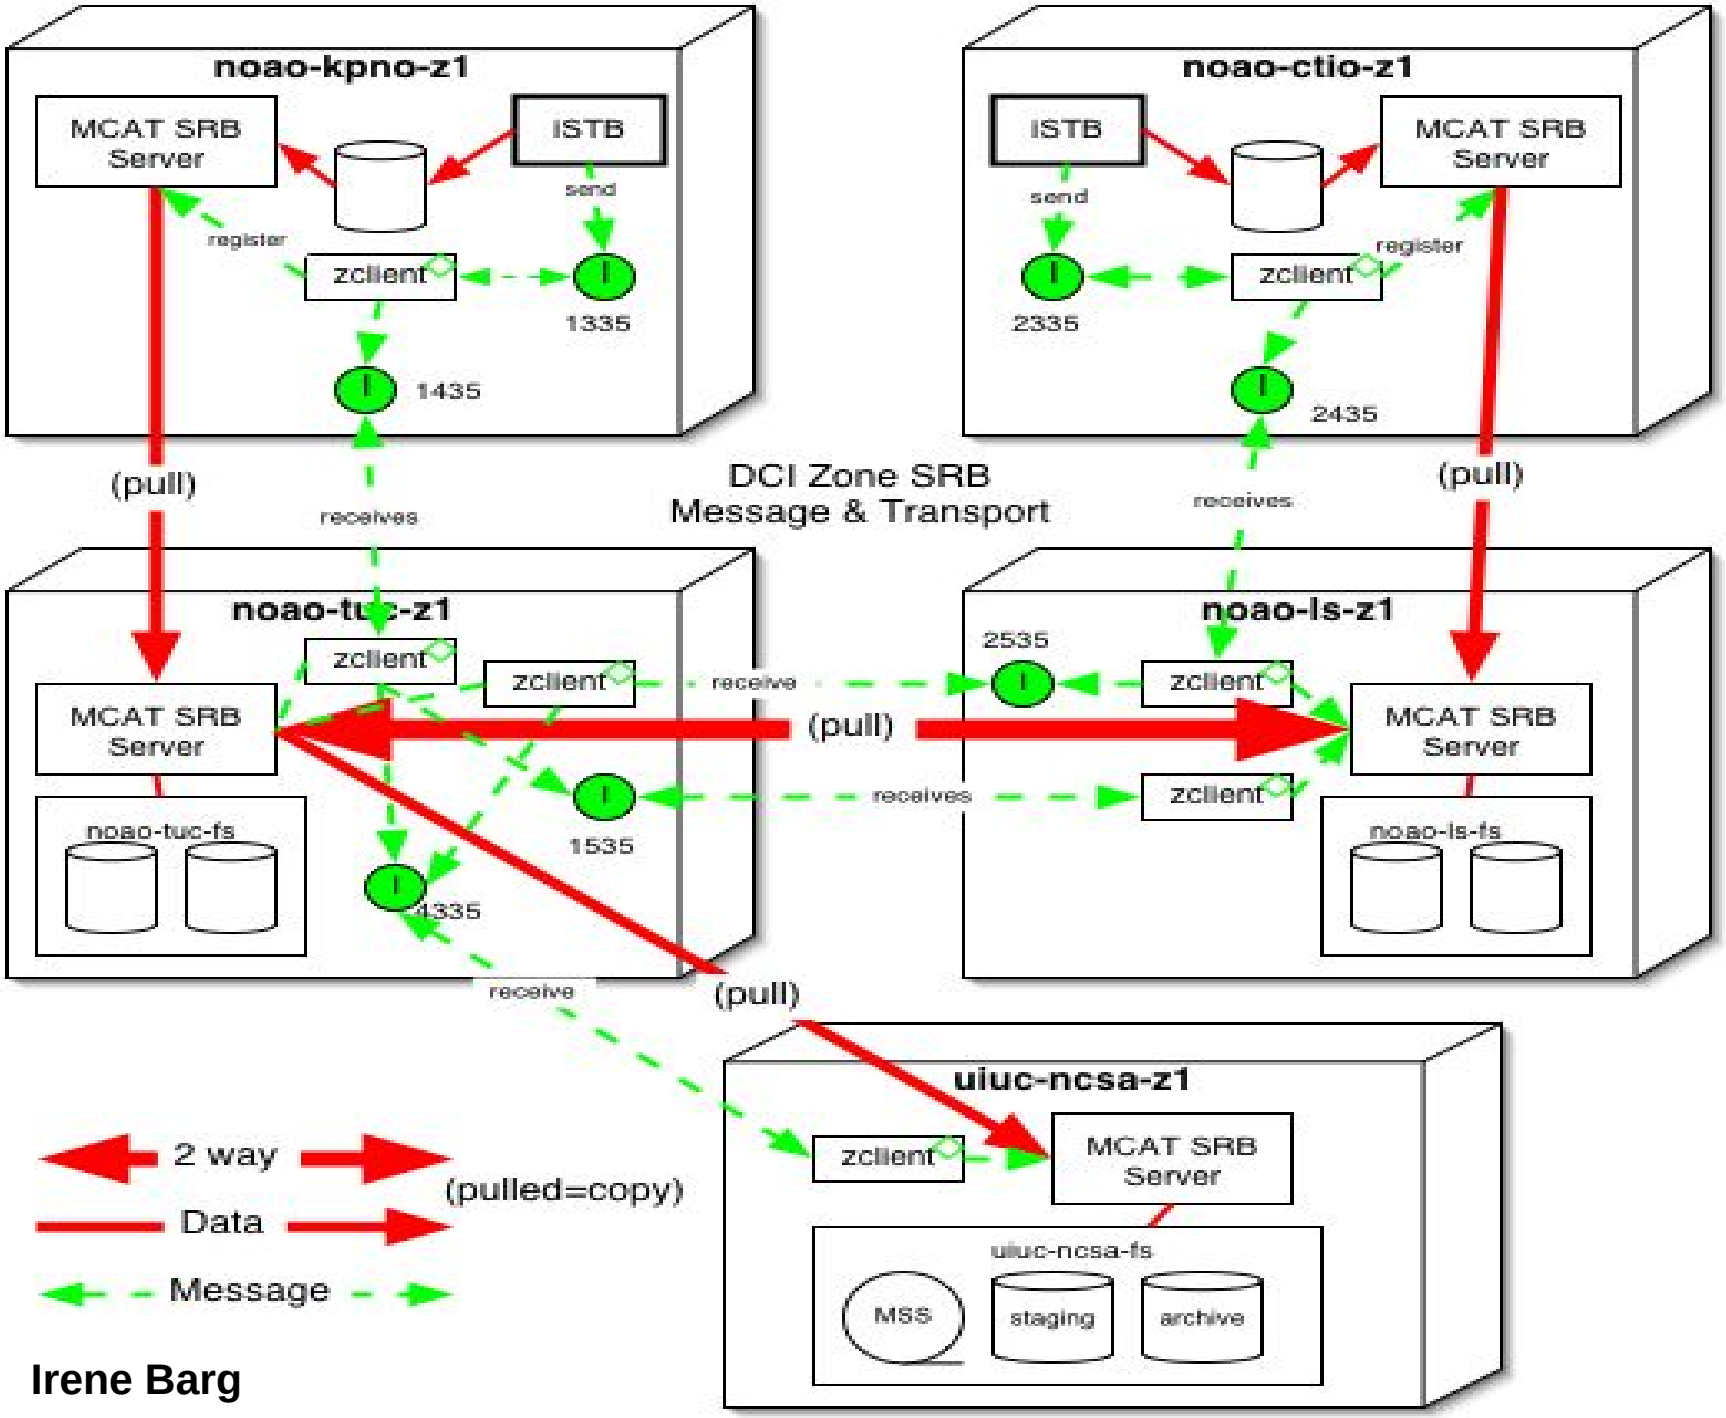
noao-kpno-z1
MCAT SRB Server
ISTB
send
register
zclient
1335
1435
noao-ctio-z1
ISTB
send
MCAT SRB Server
register
zclient
2335
2435
(pull)
receives
DCI Zone SRB
Message & Transport
(pull)
receives
noao-tuc-z1
zclient
zclient
receive
MCAT SRB Server
(pull)
noao-tuc-fs
1535
receives
4335
receive
(pull)
noao-ls-z1
2535
zclient
zclient
MCAT SRB Server
zclient
noao-ls-fs
uiuc-ncsa-z1
zclient
MCAT SRB Server
uiuc-ncsa-fs
MSS
staging
archive
2 way
(pulled=copy)
Data
Message

Worldwide University Network Data Grid

- SDSC
- Manchester
- Southampton
- White Rose
- NCSA
- U. Bergen

Manchester-SDSC mirror

- A functioning, general purpose international Data Grid for academic collaborations

UCSD

SDSC

WUNGrid Collections

- **BioSimGrid**
  - Molecular structure collaborations
- **White Rose Grid**
  - Distributed Aircraft Maintenance Environment
- **Medieval Studies**
- **Music Grid**
- **e-Print collections**
  - DSpace
- **Astronomy**

BaBar High-energy Physics

- Stanford Linear Accelerator
- Lyon, France
- Rome, Italy
- San Diego
- RAL, UK

Manchester-SDSC mirror

- A functioning international Data Grid for high-energy physics

Moved over 170 TBs of data

UCSD

SDSC

SRB Objectives

- **Automate all aspects of data discovery, access, management, analysis, preservation**
  - Security paramount
  - Distributed data
- **Provide distributed data support for**
  - Data sharing - data grids
  - Data publication - digital libraries
  - Data preservation - persistent archives
  - Data collections - Real time sensor data

Storage Resource Broker 3.3.1

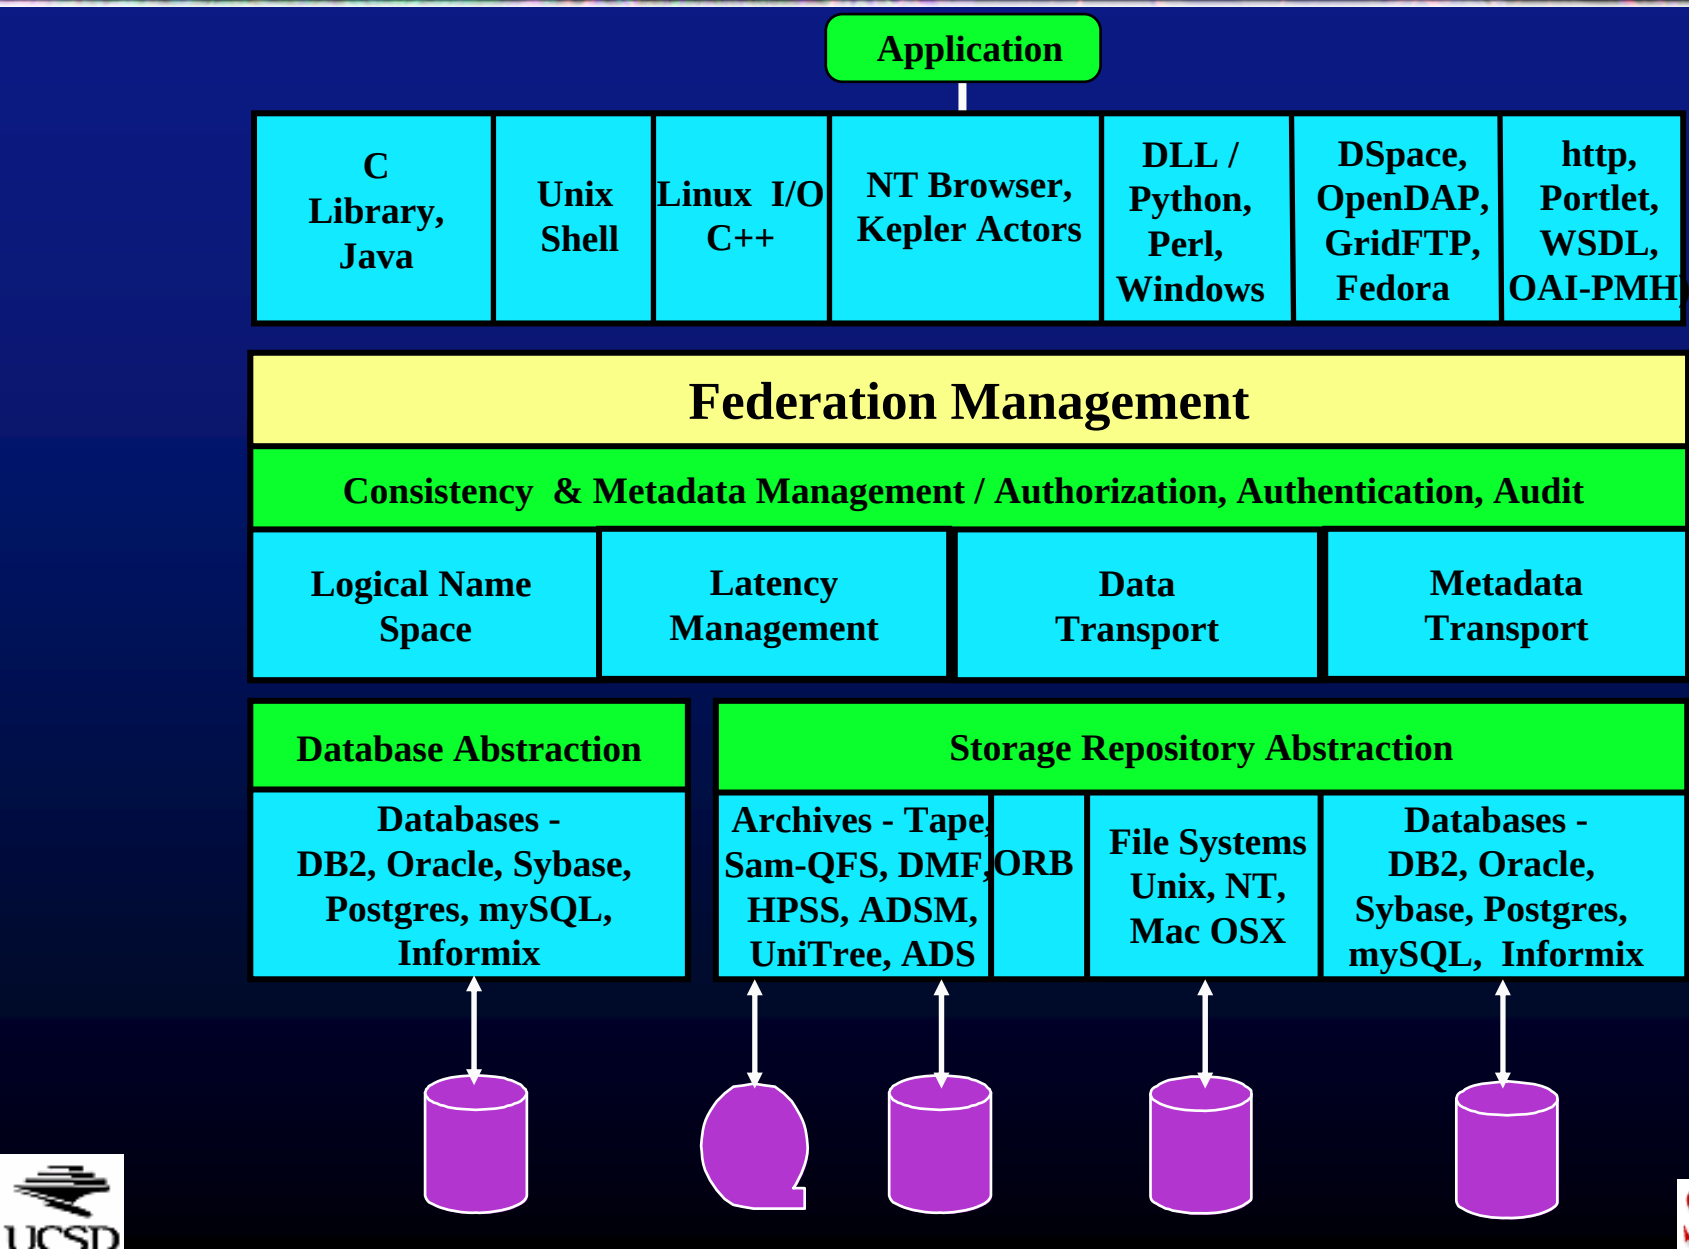

UCSD

SDSC

Data Grid Operations

- **File access**
  - Open, close, read, write, seek, stat, synch, …
  - Audit, versions, pinning, checksums, synchronize, …
  - Parallel I/O and firewall interactions
  - Versions, backups, replicas
- **Latency management**
  - Bulk operations
    - Register, load, unload, delete, …
  - Remote procedures
    - HDFv5, data filtering, file parsing, replicate, aggregate
- **Metadata management**
  - SQL generation, schema extension, XML import and export, browsing, queries,
- **GGF, "Operations for Access, Management, and Transport at Remote Sites"**

Types of Risk

- **Media failure**
  - Replicate data onto multiple media
- **Vendor specific systemic errors**
  - Replicate data onto multiple vendor products
- **Operational error**
  - Replicate data onto a second administrative domain
- **Natural disaster**
  - Replicate data to a geographically remote site
- **Malicious user**
  - Replicate data to a deep archive

How Many Replicas

- **Three sites minimize risk**
  - Primary site
    - Supports interactive user access to data
  - Secondary site
    - Supports interactive user access when first site is down
    - Provides 2nd media copy, located at a remote site, uses different vendor product, independent administrative procedures
  - Deep archive
    - Provides 3rd media copy, staging environment for data ingestion, no user access

Deep Archive

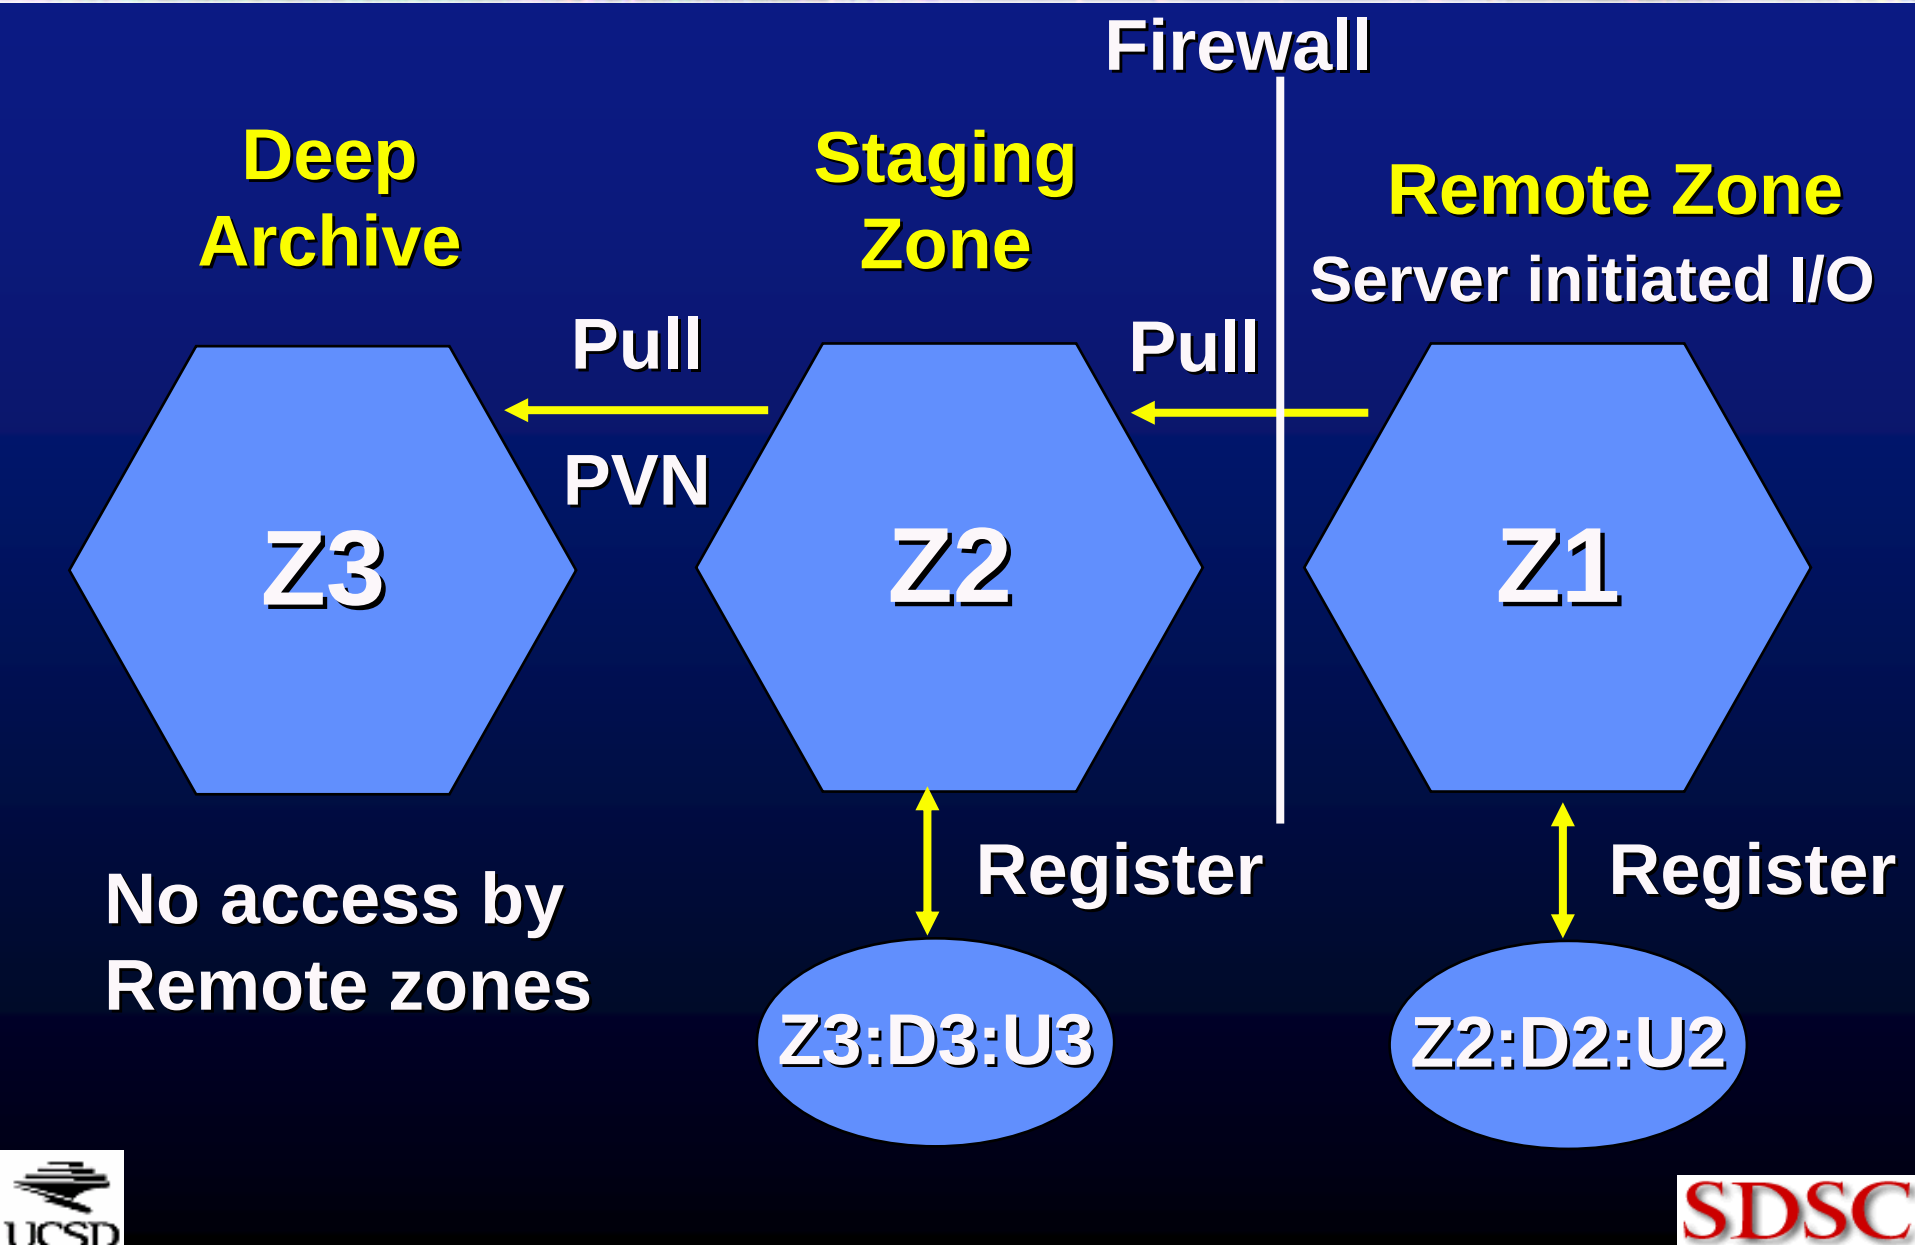

SRB Developers

Reagan Moore	- PI
Michael Wan	- SRB Architect
Arcot Rajasekar	- SRB Manager
Wayne Schroeder	- SRB Productization
Charlie Cowart	- inQ
Lucas Gilbert	- Jargon
Bing Zhu	- Perl, Python, Windows
Antoine de Torcy	- mySRB web browser
Sheau-Yen Chen	- SRB Administration
George Kremenek	- SRB Collections
Arun Jagatheesan	- Matrix workflow
Marcio Faerman	- SCEC Application
Sifang Lu	- ROADnet Application
Richard Marciano	- SALT persistent archives

Contributors from UK e-Science, Academia Sinica, Ohio State University, Aerospace Corporation, …

75 FTE-years of support
About 300,000 lines of C

UCSD

SDSC

Development

- **SRB 1.1.8 - December 15, 2000**
  - Basic distributed data management system
  - Metadata Catalog
- **SRB 2.0 - February 18, 2003**
  - Parallel I/O support
  - Bulk operations
- **SRB 3.0 - August 30, 2003**
  - Federation of data grids
- **SRB 3.4.1 - April 30, 2006**
  - Feature requests (quotas)

For More Information

Reagan W. Moore
San Diego Supercomputer Center
moore@sdsc.edu
http://www.sdsc.edu/srb/

UCSD

SDSC