Building Preservation Environments with Data Grid Technology
(NARA Research Prototype Persistent Archive)

Reagan Moore
Richard Marciano
Arcot Rajasekar
Michael Wan
Wayne Schroeder
Antoine de Torcy
Sheau-Yen Chen

http://www.sdsc.edu/NARA/
http://www.sdsc.edu/srb/

Topics

- Moore, R., "Building Preservation Environments with Data Grid Technology", American Archivist, vol. 69, no. 1, pp. 139-158, July 2006.
- Identify relevant preservation concepts
- Prove concepts in NARA Research Prototype Persistent Archive
- Identify future technology development goals

Preservation Concepts

- Paper records
  - Authenticity
  - Integrity
- Digital records
  - Authenticity
  - Integrity
  - Infrastructure independence
  - Scalability

Traditional Preservation Concepts

- Authenticity - the inextricable linking of identity metadata to a record
  - Date record is made
  - Date record is transmitted
  - Date record is received
  - Date record is set aside (i.e.,, filed)
  - Name of author (person or organization issuing the record)
  - Name of addressee (person or organization for whom the record is intended)
  - Name of writer (person or organization responsible for the articulation of the record's content)
  - Name of originator (electronic address from which the record is sent)
  - Name of recipient(s) (person or organization to whom the record is sent)
  - Name of creator (person or organization in whose archival fonds the record exists)
  - Name of action or matter (the activity in the course of which the record is created)
  - Name of documentary form (e.g., e-mail, report, memo)
  - Identification of digital components
  - Identification of attachments (e.g.,, digital signature)
  - Archival bond (e.g.,, classification code)
  - Assertions about the creation of the record
- Assertions made by the archivist about the creator of the record and the creation process

Traditional Preservation Concept

- Integrity - the management of record correctness and the chain of custody
  - Name(s) of the handling office/officer
  - Name of the office of primary responsibility for keeping the record
  - Indication of annotations or comments
  - Indication of actions carried out on the record (e.g., audit trail)
  - Indication of technical modifications due to transformative migration
  - Integrity signature
  - Validation date for last integrity check
- Assertions made by the archivist about the management of the records

Preservation of Digital Records

- Extract a digital record from the environment in which it was created,
- Import the digital record into the preservation environment
- Given that the preservation environment will evolve, the process must be repeated with each new generation of storage technology
  - Extract from the old technology and import onto the new storage technology
- Two more preservation concepts:
  - **Infrastructure independence**
  - **Scalability**

Extraction of Electronic Records

Data Access Method (Web Browser)

↕

Storage Repository

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

Extract

Digital record (file)

Identifiers (names used to manage the file)

Provenance metadata (creator, time stamps)

Integrity metadata (digital signature)

Encoding format

To Import into a Preservation Environment

- Archivist needs **persistent** identifiers:
  - Names used to describe archivists, files, and storage systems
  - Names for authenticity and integrity metadata attributes
  - Storage locations for electronic records
- Archivist needs to control the properties of the electronic record
  - Access controls for processing the electronic records
  - Locations of replicas
  - Audit trails on actions performed on electronic record
  - Checksums for validating integrity

Digital Preservation Concept

- Infrastructure Independence
  - The ability to manage all of the properties of the electronic records independently of the choice of storage systems
  - Provides persistent names to identify persons, files, storage systems as well as manage authentication and authorization
    - Data virtualization
    - Trust virtualization

Data Grids

- Data grids implement data virtualization
  - Ability to manage properties of shared collection independently of the choice of storage system
  - Ability to access records stored in all types of storage systems
  - Ability to support multiple types of access mechanisms independently of choice of type of storage system
- Data grids implement trust virtualization
  - Ability to authenticate users independently of the local administrative domain
  - Ability to manage access controls independently of the local file system or archive
- Data grids are used to manage shared collections that are distributed across multiple sites and multiple storage systems

Astronomy Data Grid

- National Optical Astronomy Observatory
  - Chile
  - Tucson, Arizona
  - Kitt Peak
  - NCSA, Illinois
- Replicate images taken by a telescope in Chile to an archive at NCSA

- A functioning international Data Grid for Astronomy

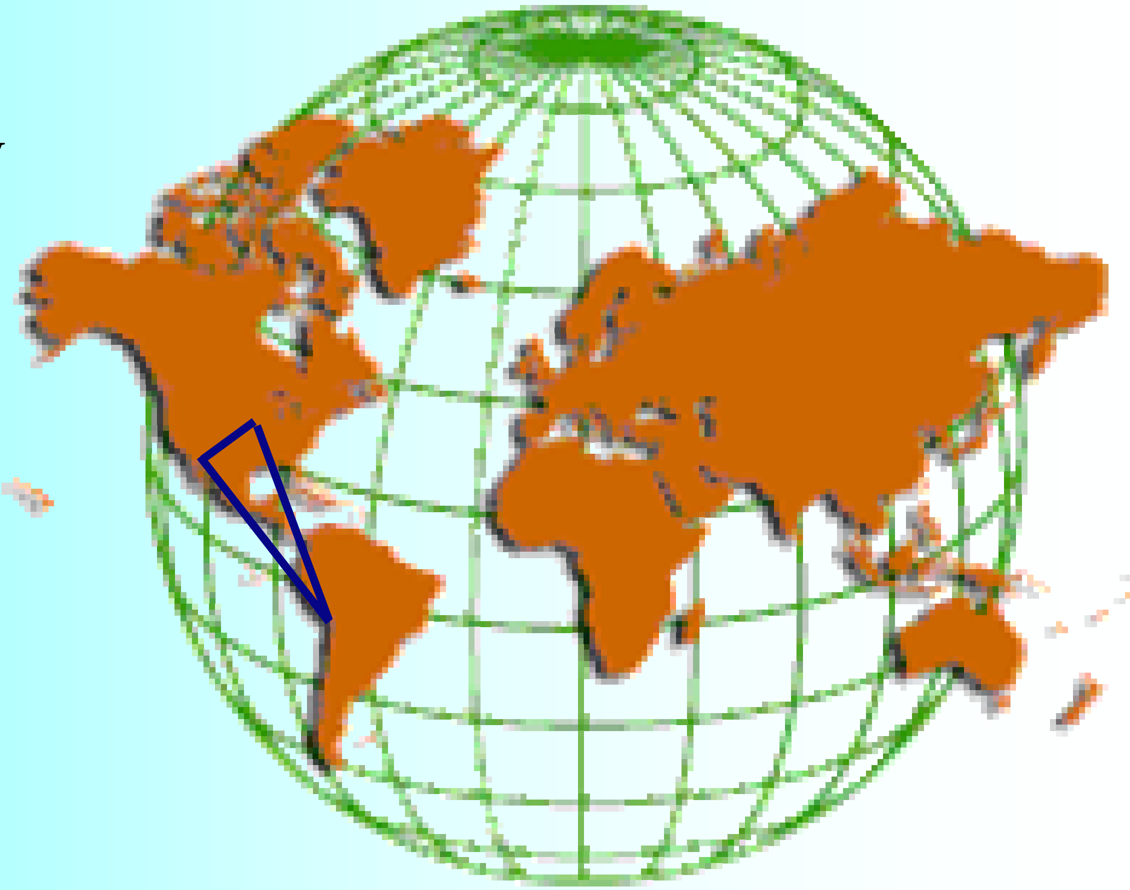

Manages over 400,000 images

Generic Distributed Data Management

- Data grids support
  - Shared collections distributed across international organizations
  - Digital libraries for the publication of records
  - Real-time sensor data systems
  - Persistent archives that manage technology evolution
- Data grids manage
  - Small collections with a few thousand files and a few Gigabytes of data
  - Large collections with 800,000 Gigabytes (800 Terabytes) of data and 100 million files
  - Collections stored within a single computer
  - Collections distributed across multiple international sites

Using a Data Grid – *in Abstract*

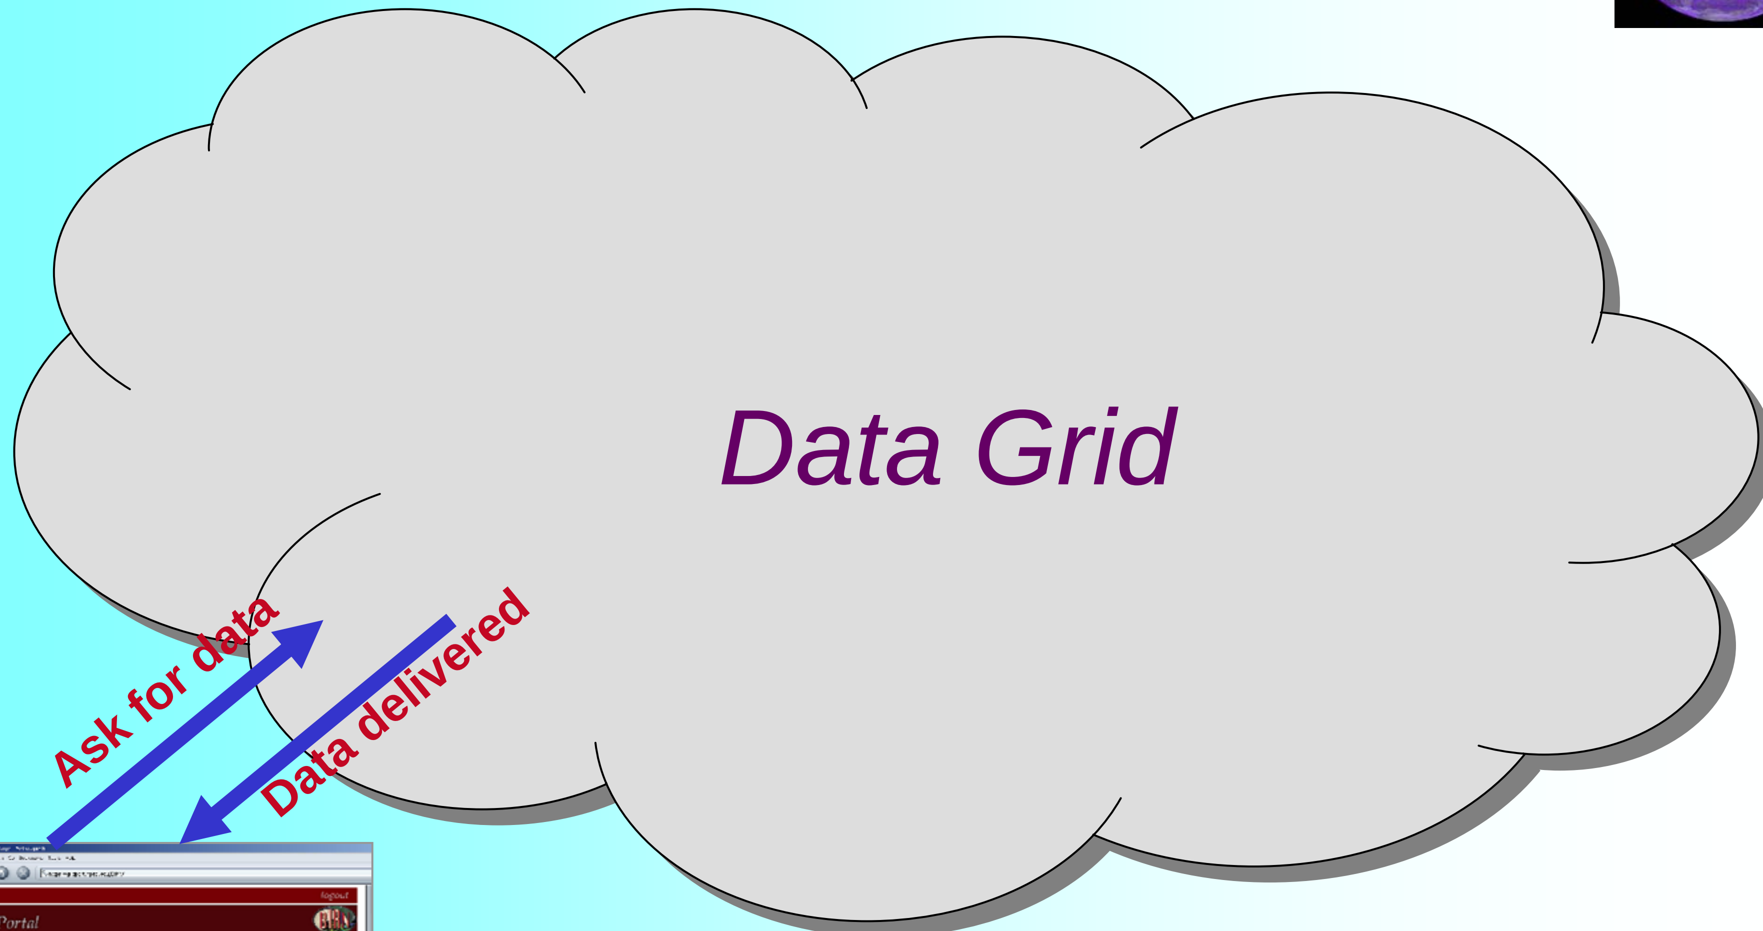

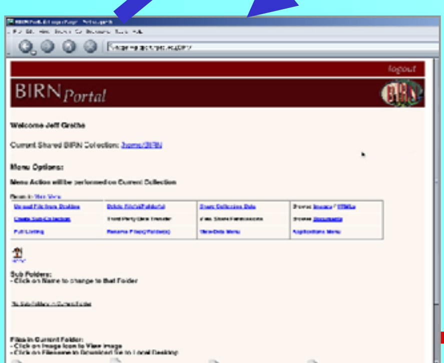

- User asks for data from the data grid
- The data is found and returned
  - Where & how details are hidden

Using a Data Grid - *Details*

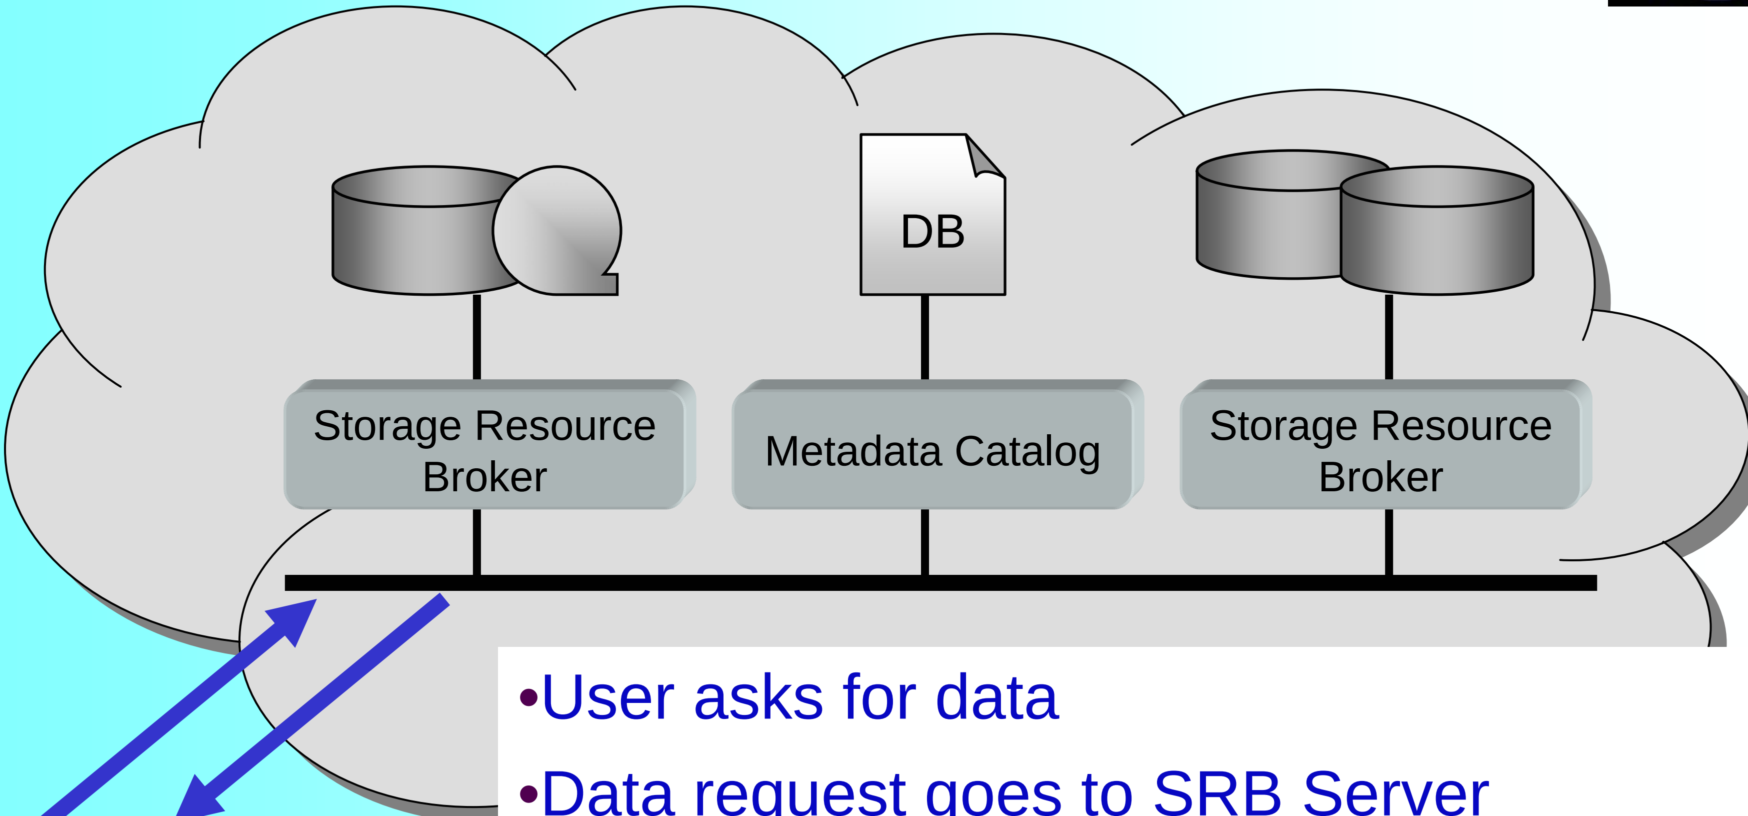

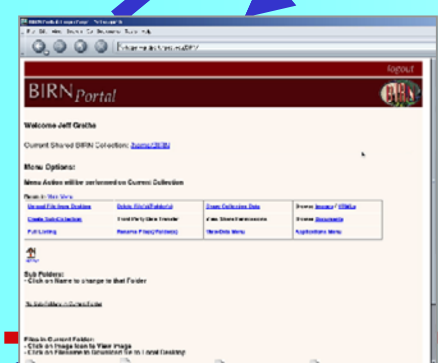

- User asks for data
- Data request goes to SRB Server
- Server looks up data in catalog
- Catalog tells which SRB server has data
- 1st server asks 2nd server for data
- The data is found and returned

Using a Data Grid - *Details*

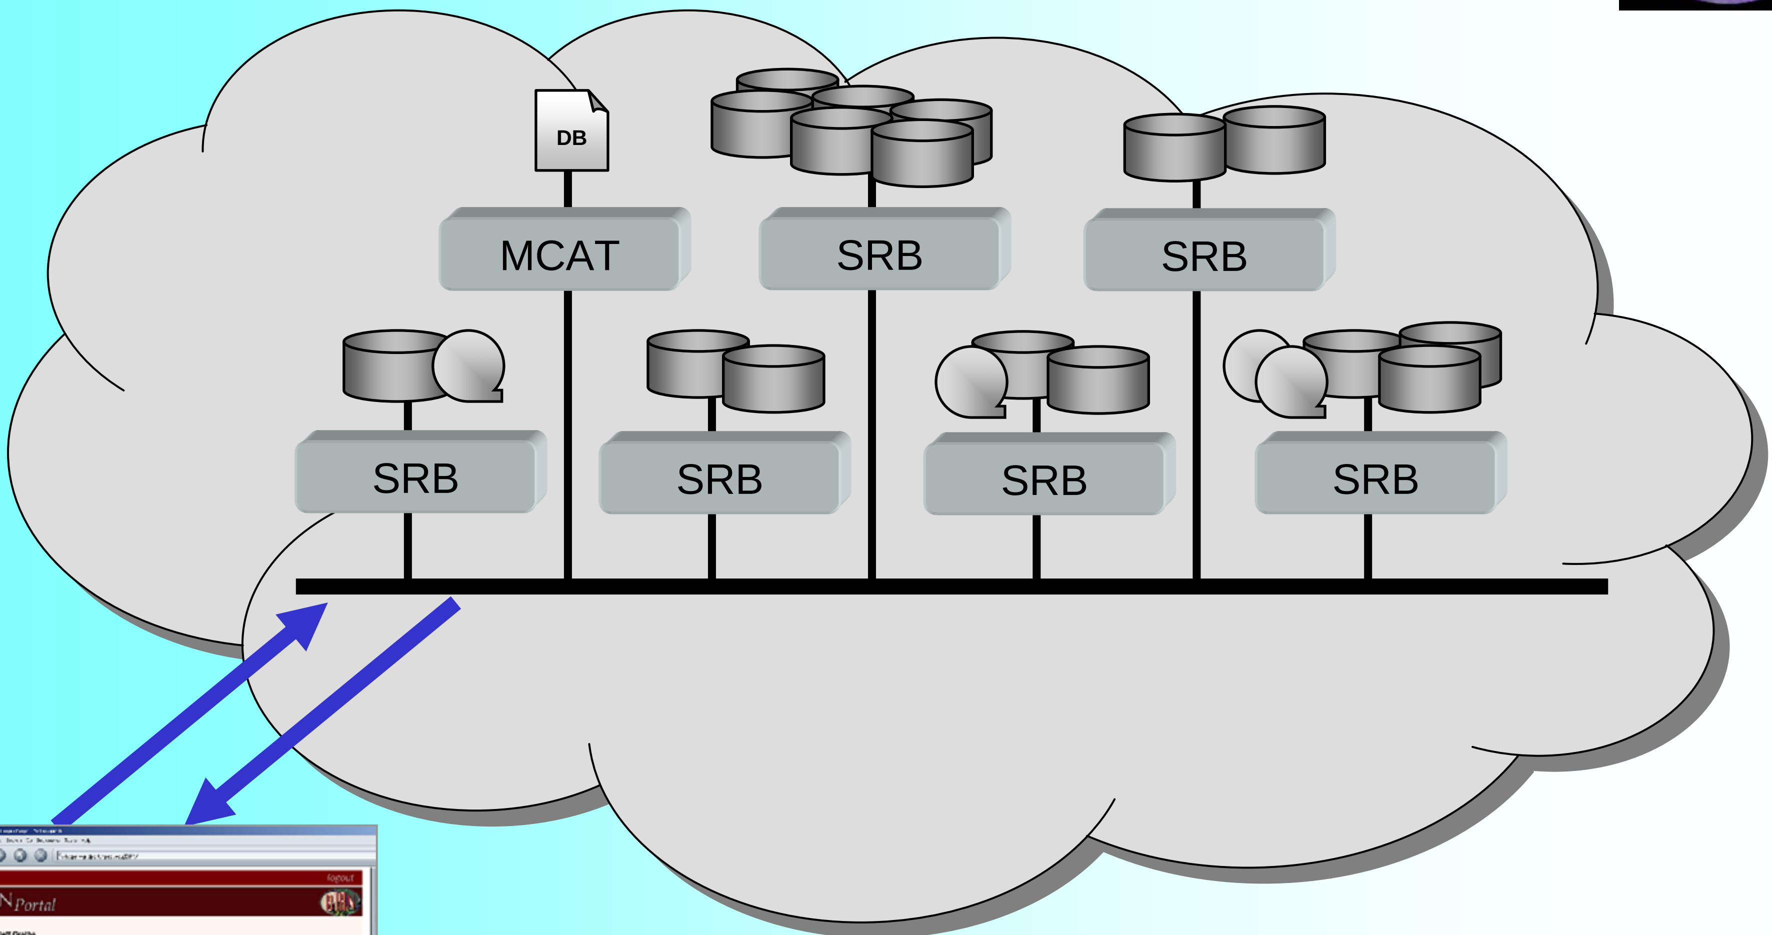

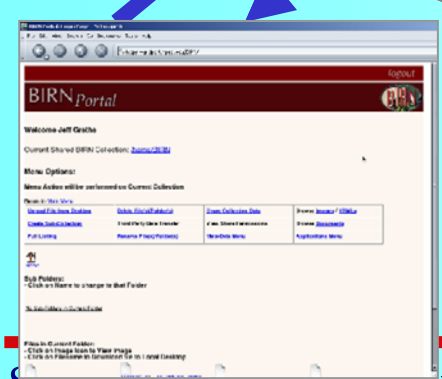

- Data Grid has arbitrary number of servers
- Heterogeneity is hidden from users

Infrastructure Independence

- Manage migration to new choices of storage systems or access protocols or security technology
  - At the point in time when the electronic records are migrated from old technology to new storage systems, both systems are present
  - The ability of data grids to interoperate with multiple types of storage systems means they can be used to manage data migration to new types of storage systems
- Data grids provide the fundamental mechanisms needed to implement infrastructure independence
  - Storage system protocol converters (SRB drivers)
  - Application interface protocol converters
  - Management of record authenticity and integrity

Import into Preservation Environment

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 Software

- 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

Logical Arrangement of Digital Records

- Persistent identifier for the record is the logical file name
- Arrangement hierarchy imposed on the logical file names as collection hierarchy
  - Record group / Record series / File-unit / Item / Object
  - Associate Life Cycle Data Requirements Guide attributes with each level of the collection hierarchy
  - Separate extensible schema associated with each logical file name
- Information about all operations performed upon digital record are mapped to the logical file name
  - Integrity attributes
  - State information describing the result of each operation
  - **Logical file name is the link between authenticity information and the record**

NARA Research Prototype Persistent Archive

- Implemented using the SDSC Storage Resource Broker (SRB) Data Grid
  - Demonstrated migration to new types of storage systems
    - Added commodity-based disk file storage systems
  - Demonstrated evolution of access methods
    - Added interfaces to web browsers, workflow systems
  - Demonstrated migration to new transport mechanisms
    - Added mechanisms to support interaction with network firewalls and support bulk load of records
  - Demonstrated replication of records across multiple systems
  - Demonstrated automated loading of authenticity metadata
  - Demonstrated mechanisms to implement a deep archive

National Archives and Records Administration Research Prototype Persistent Archive

Federation of Five Independent Data Grids

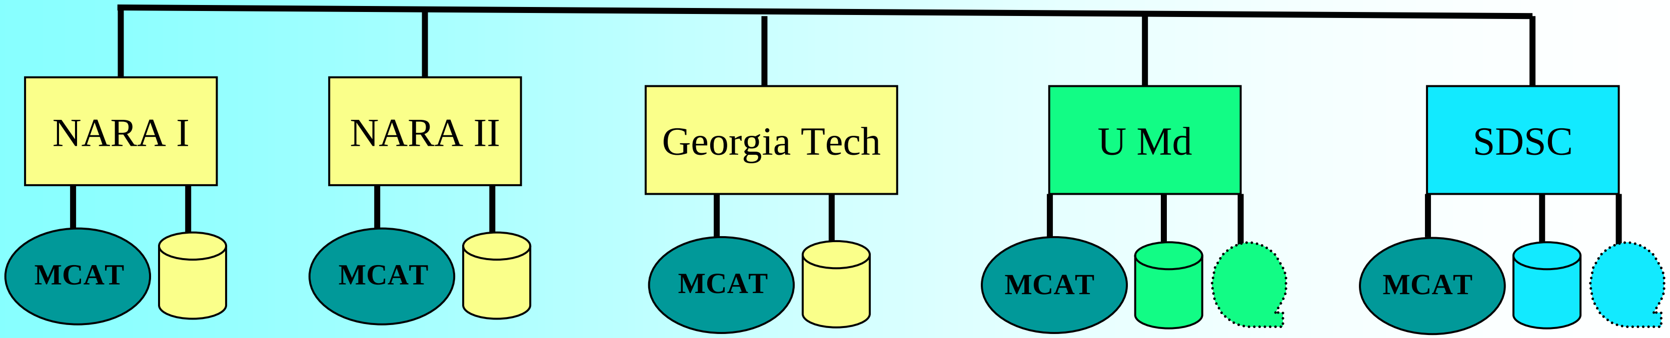

Extensible Environment, can federate with additional research and education sites

Federation Between Data Grids

Data Access Methods (Web Browser, DSpace, OAI-PMH)

Data Collection A

Data Collection B

Data Grid (A)

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

Data Grid (B)

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

Access controls and consistency constraints on cross registration of digital entities

NARA Research Prototype Persistent Archive

Powerful platform for demonstrating preservation

- Interoperability across diverse platforms
  - Multiple access mechanisms
  - Multiple types of storage systems
- Extensible schema to support authenticity and integrity attributes
  - Archivist defined metadata attributes
  - Audit trails
  - Checksums
- Mitigation of risk of data loss
  - Replication of data
  - Federation of catalogs
- Synchronization across zones
- Deep archive
- Collaboration between SDSC, University of Maryland, Stanford Linear Accelerator, Georgia Institute of Technology

Scalability: Automation of Preservation Processes

- Generic operations for managing interactions with files
  - Open, close, read, write, seek, stat, synch, …
  - Authentication and authorization
- Latency management operations
  - Aggregation of files in containers
  - Bulk load, unload, registration
  - Remote procedures for metadata extraction, file filtering
- Database interaction operations
  - Registration of SQL command strings
  - Extensible metadata schema
  - Table import and export
- Integrity mechanisms
  - Replication, checksum validation, synchronization
  - Audit trails
  - Federation

NARA Research Prototype Persistent Archive

- Prototype persistent archive design based on:
  - Data virtualization - ability to manage collection properties independently of storage system
  - Trust virtualization - ability to manage authentication and authorization independently of administrative domains
  - Latency management - scalable operations
  - Collection management - impose logical arrangement such as LDCRG hierarchy
  - Federation - ability to create preservation environments that span multiple data grids
- Automation of operations across a million records

Risk Mitigation Against Data Loss

- How many replicas are enough?
- Media corruption
  - Maintain a copy on a second set of media
- Systemic vendor error
  - Maintain a copy on a separate vendor’s product
- Operational error
  - Maintain a copy in a separate administrative domain
- Natural disaster
  - Maintain a copy at a geographically remote site
- Malicious users
  - Maintain a copy in a deep archive under archivist control

Deep Archive

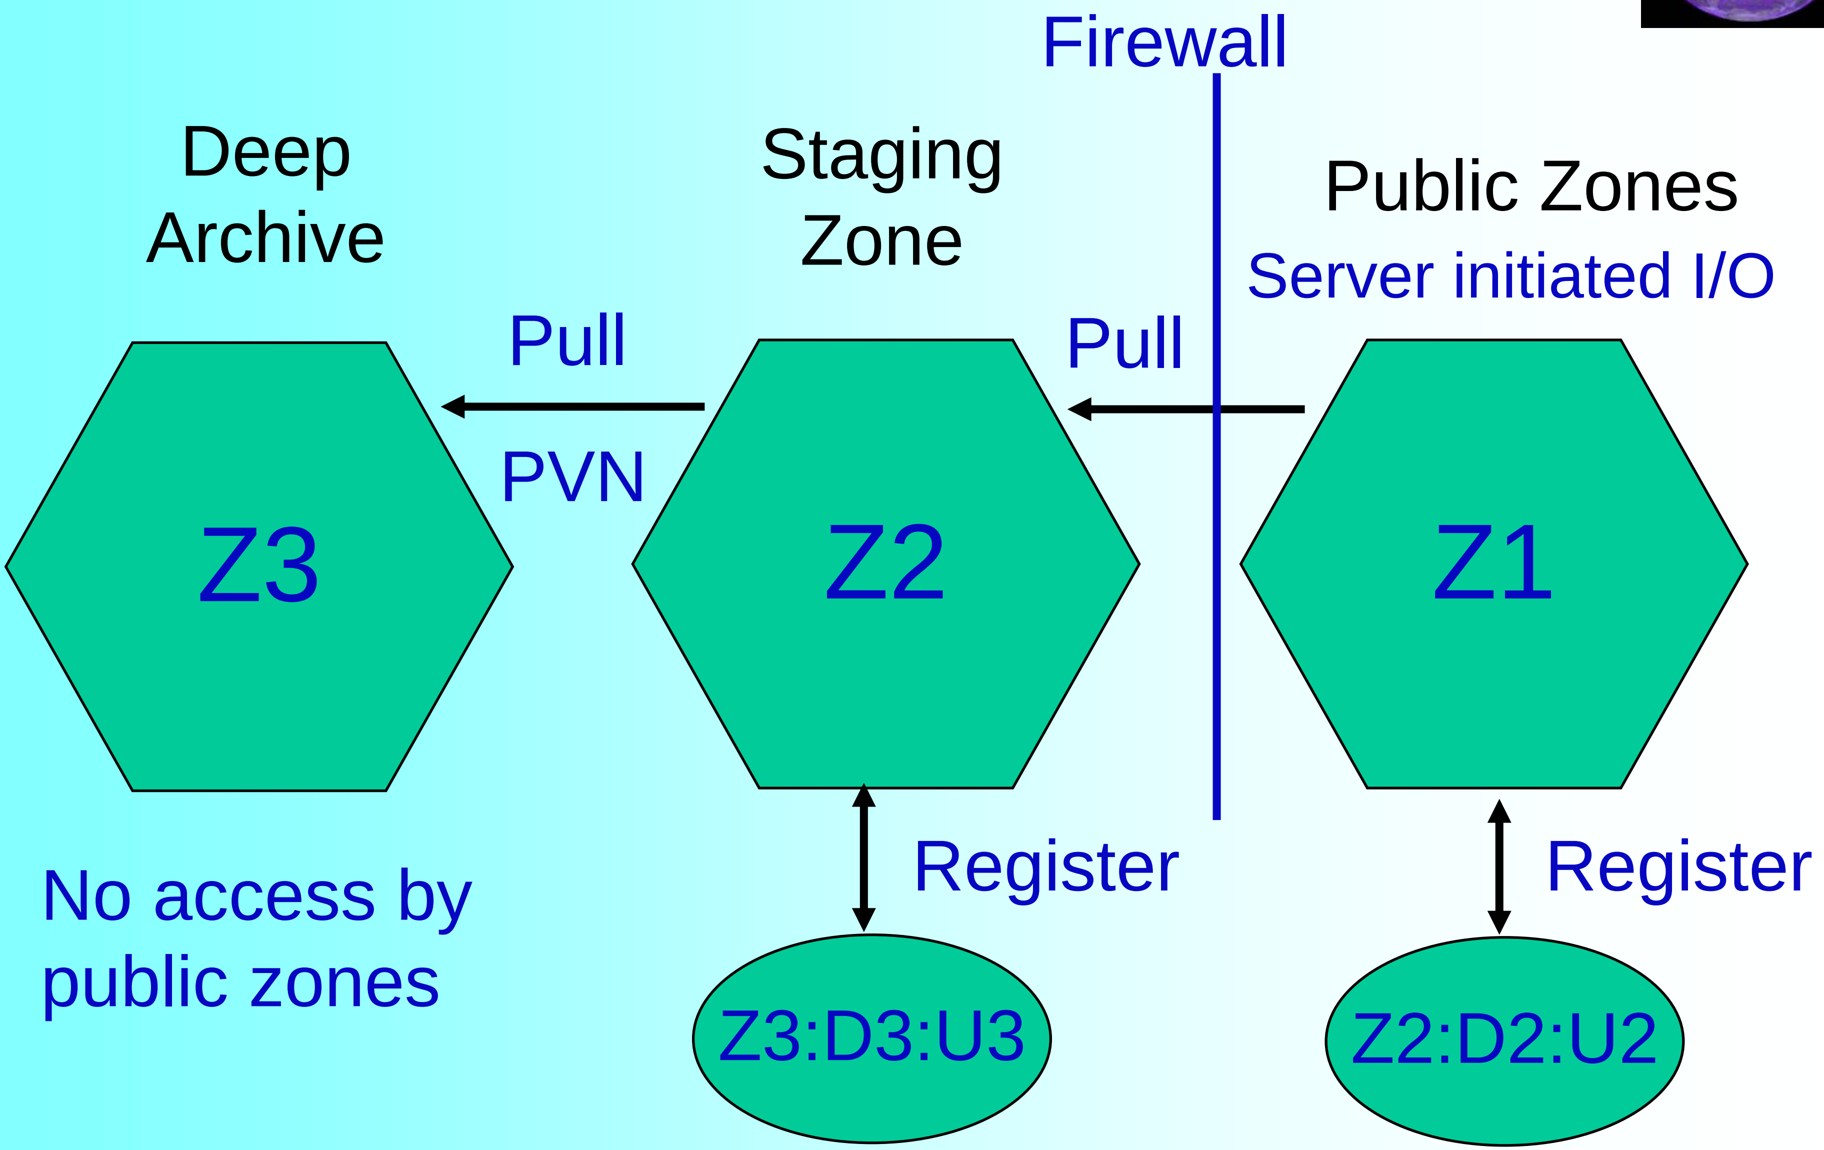

NARA Leadership

- Technology demonstrated in the NARA Research Prototype Persistent Archive is now being applied in multiple national and international collaborations
  - NSF National Science Digital Library persistent archive
  - NOAO preservation data grid for astronomy images
  - California Digital Library - Digital Preservation Repository
  - Taiwan Preservation Data Grid
  - European Union CASPAR preservation environment
- NARA support has been essential in the continued development of data grid technology for building preservation environments
  - Software distributed to 174 institutions in 2004-2005
  - Half of the sites are international

SDSC Research Objectives

- Understand principles underlying digital preservation
  - Authenticity - assertions made about creator
  - Integrity - assertions made by archivist about management
  - Infrastructure independence - ability to migrate to new or alternate technology
  - Scalability - automation of preservation policies
- Map from preservation principles to Information Technology concepts
  - **Data virtualization**
  - **Trust virtualization**
  - **Latency management**
  - **Shared collection management**
  - **Federation**
  - **Policy virtualization**

SRB Developers

Reagan Moore	- PI
Richard Marciano	- Sustainable Archives and Library Technology
Michael Wan	- SRB Architect
Arcot Rajasekar	- SRB Manager, Information Architect
Wayne Schroeder	- SRB Productization, Security
Charlie Cowart	- inQ, NSDL application
Lucas Gilbert	- Jargon, DSpace/Fedora integration
Bing Zhu	- Perl, Python, Windows load libraries
Antoine de Torcy	- mySRB web browser, NARA collections
Sheau-Yen Chen	- SRB Administration
George Kremenek	- SRB Collections
Arun Jagatheesan	- Matrix workflow
Leesa Brieger	- NVO Application
Sifang Lu	- ROADnet Application
Students	

75 FTE-years of development and application

About 300,000 lines of C

Storage Resource Broker 3.4.2

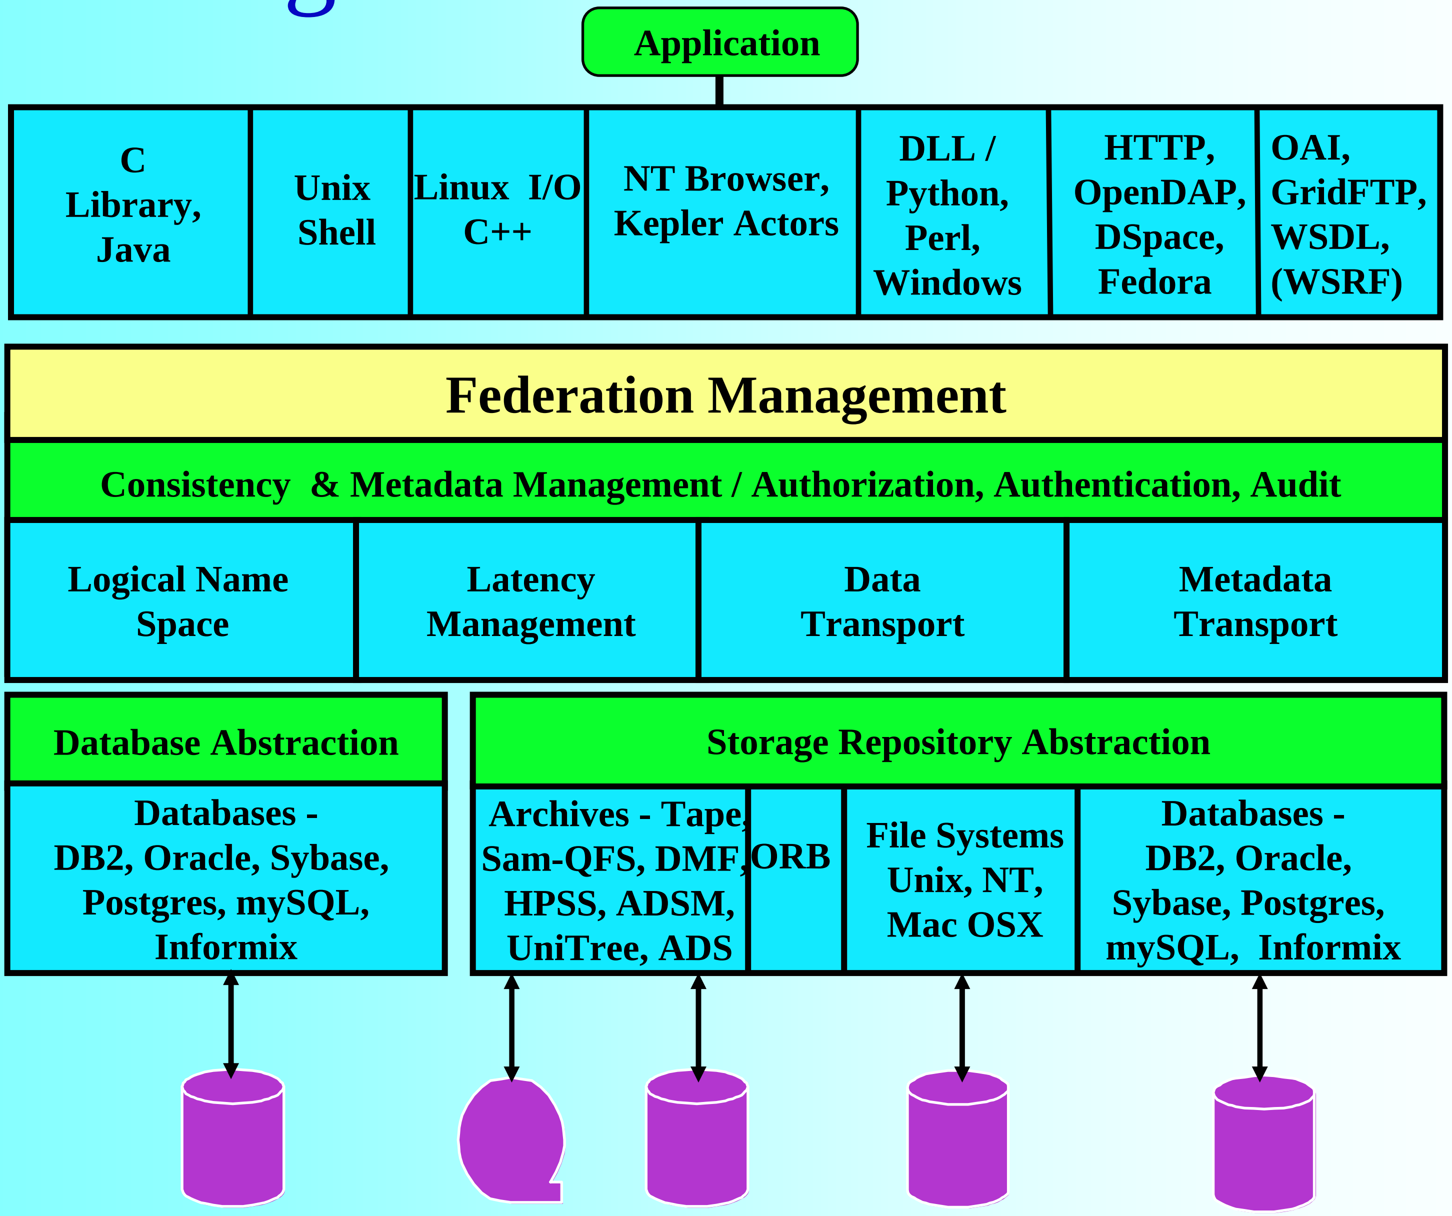

integrated Rule-Oriented Data System

- Traditional shared collection
  - Metadata catalog manages persistent state information
- Add rule engine
  - Choose rule set to apply to a given record series
  - Allow dynamic rule changes
    - Track version of rule, date version was applied and the level of granularity (item, sub-collection)
    - Manage temporary state information needed for rule execution
    - Manage persistent state information resulting from rule application
- Apply rules that control assertions about the collection
  - Data distribution, replication, access constraints
  - Integrity validation
  - Metadata consistency

NSF and NARA Funded Research

- Automation of the application of management policies
  - Identify, characterize, and manage rules for:
    - Clients (user access, allowed operations, views)
    - Item state information (update, consistency, validation)
    - Collection properties (global state, integrity, replication)
  - Dynamically apply the constraints to collection subsets
    - Modify rules without rewriting code
    - Re-apply rules to enforce new policies
  - Create modular system which will allow components of the data management environment to be updated independently
- Demonstrate use of dynamic rules to implement and manage preservation policies
  - Automation of the evaluation of the RLG/NARA assessment criteria for trusted digital repositories

iRODS - integrated Rule-Oriented Data System

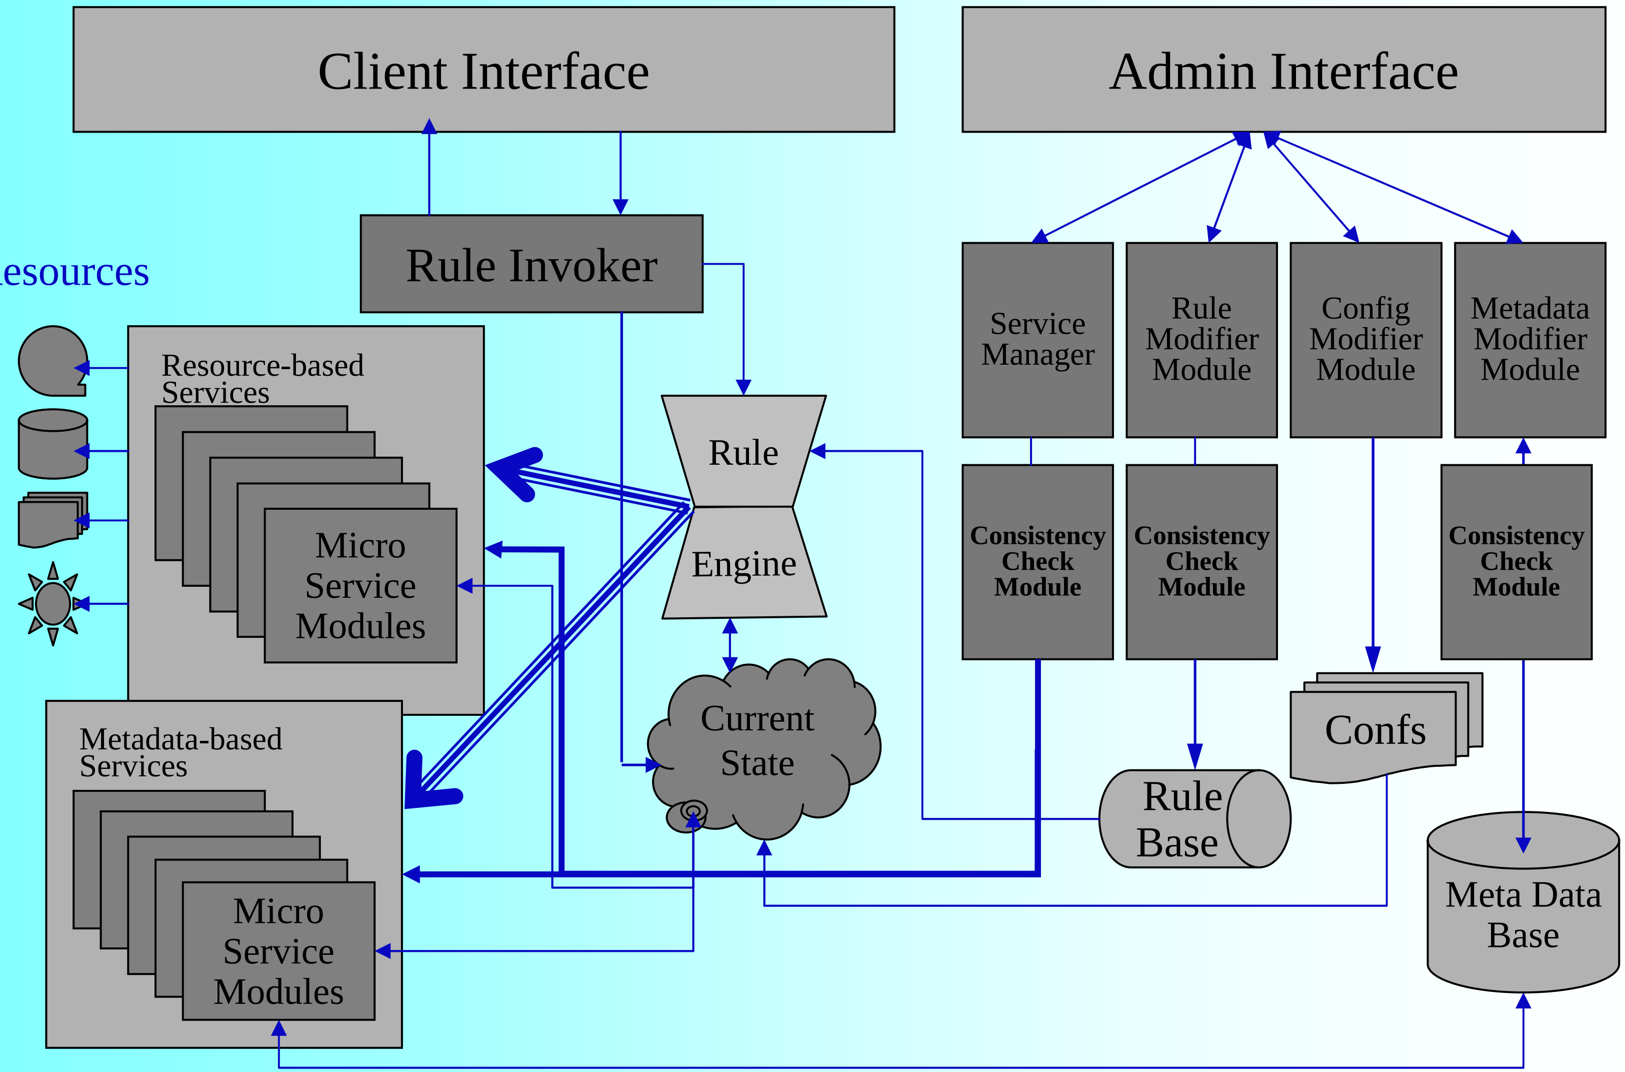

For More Information

Reagan W. Moore

moore@sdsc.edu

http://www.sdsc.edu/NARA/

http://www.sdsc.edu/srb/