Access Control for a Replica Management Database

J. M. Wozniak, P. Brenner
D. Thain, A. Striegel, J. A. Izaguirre

SSS 2006

Supported by: NSF DBI-0450067

UNIVERSITY OF NOTRE DAME

Ad Hoc Access Control

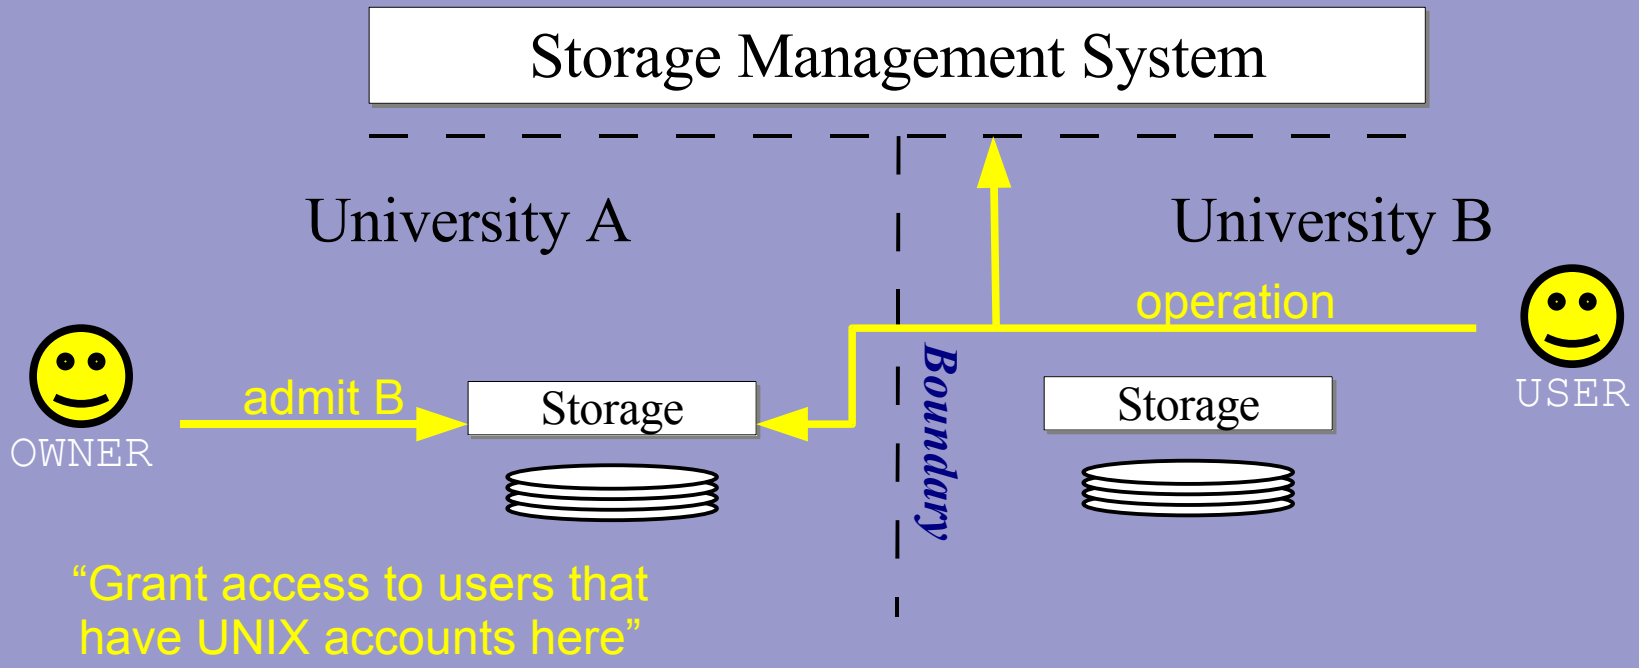

- Storage owners can grant ad hoc access for remote users
- Such usage may force a change in an overarching system
- How do we make authenticate these operations - without centralized administration?

Storage Fabric

- Cooperative storage clusters in collaborative research environments
- Highly volatile connections
- Multiple-use machines
- Varying free disk space, CPU load, etc.
- *No master list of storage resources...*

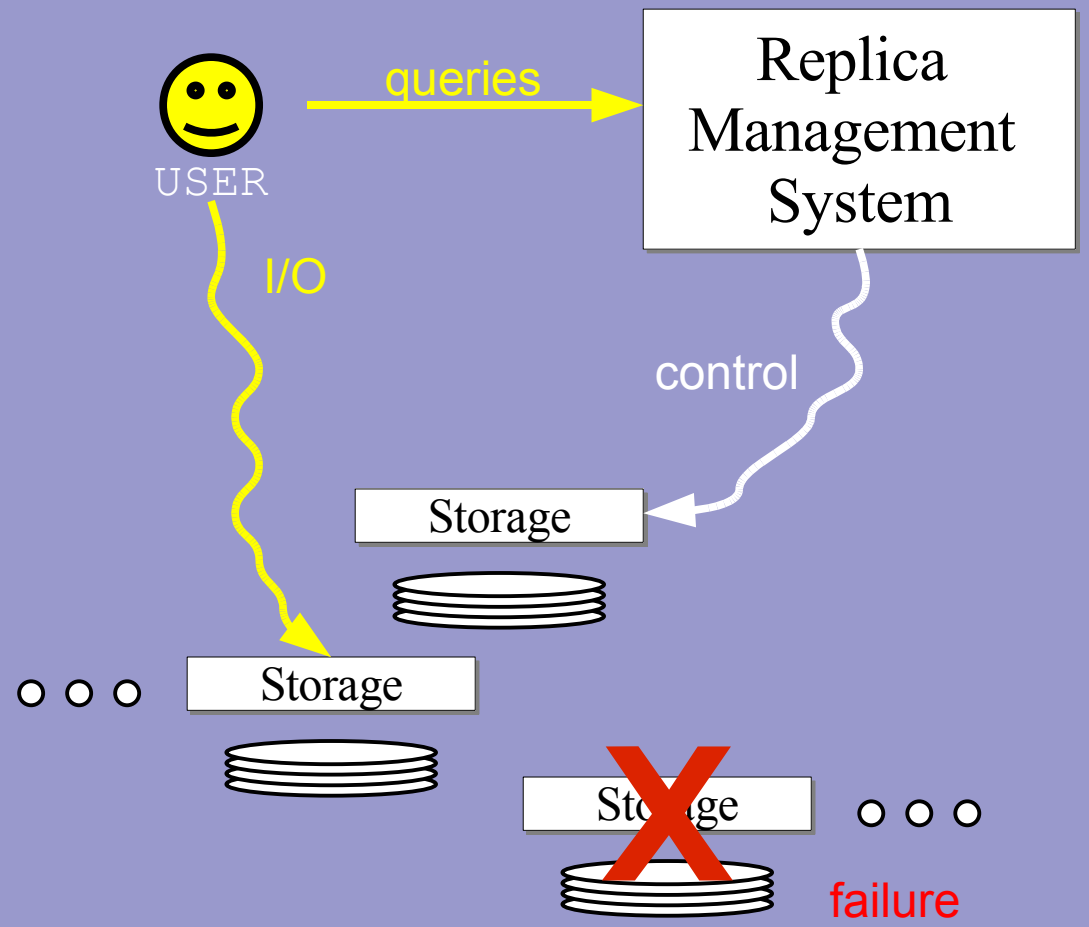

Plethora of Access Modes

- Cooperative storage users in collaborative research environments
- Highly volatile user groups
- Multiple-identity users
- Varying access to subsets of available systems
- *No master list of users...*

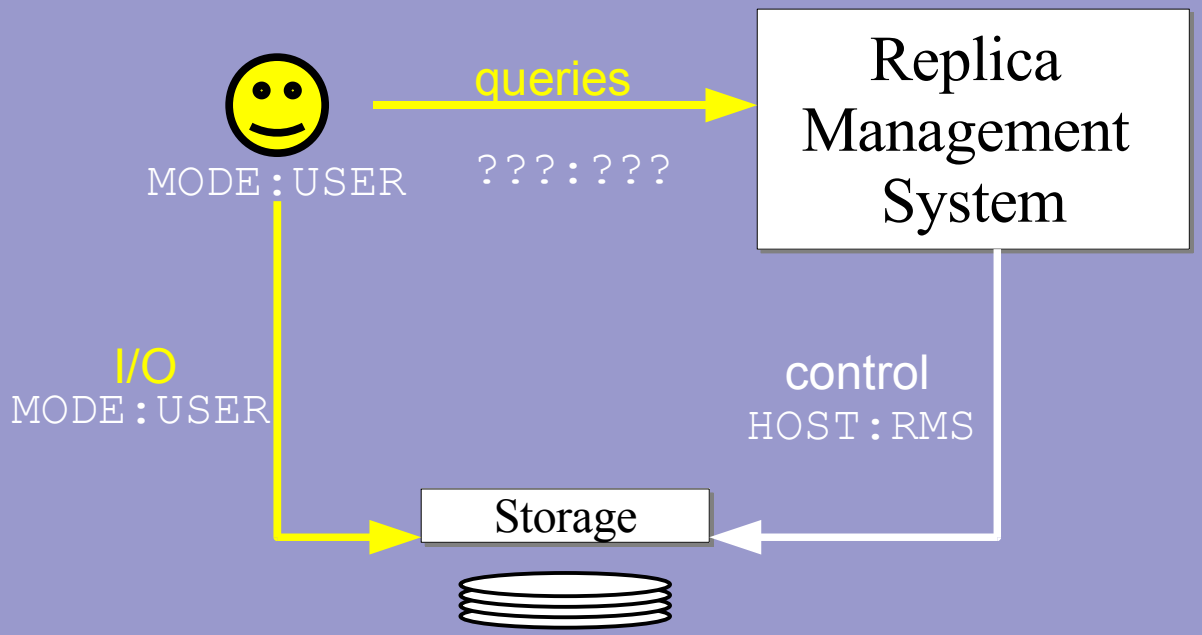

Ad Hoc Storage Networks

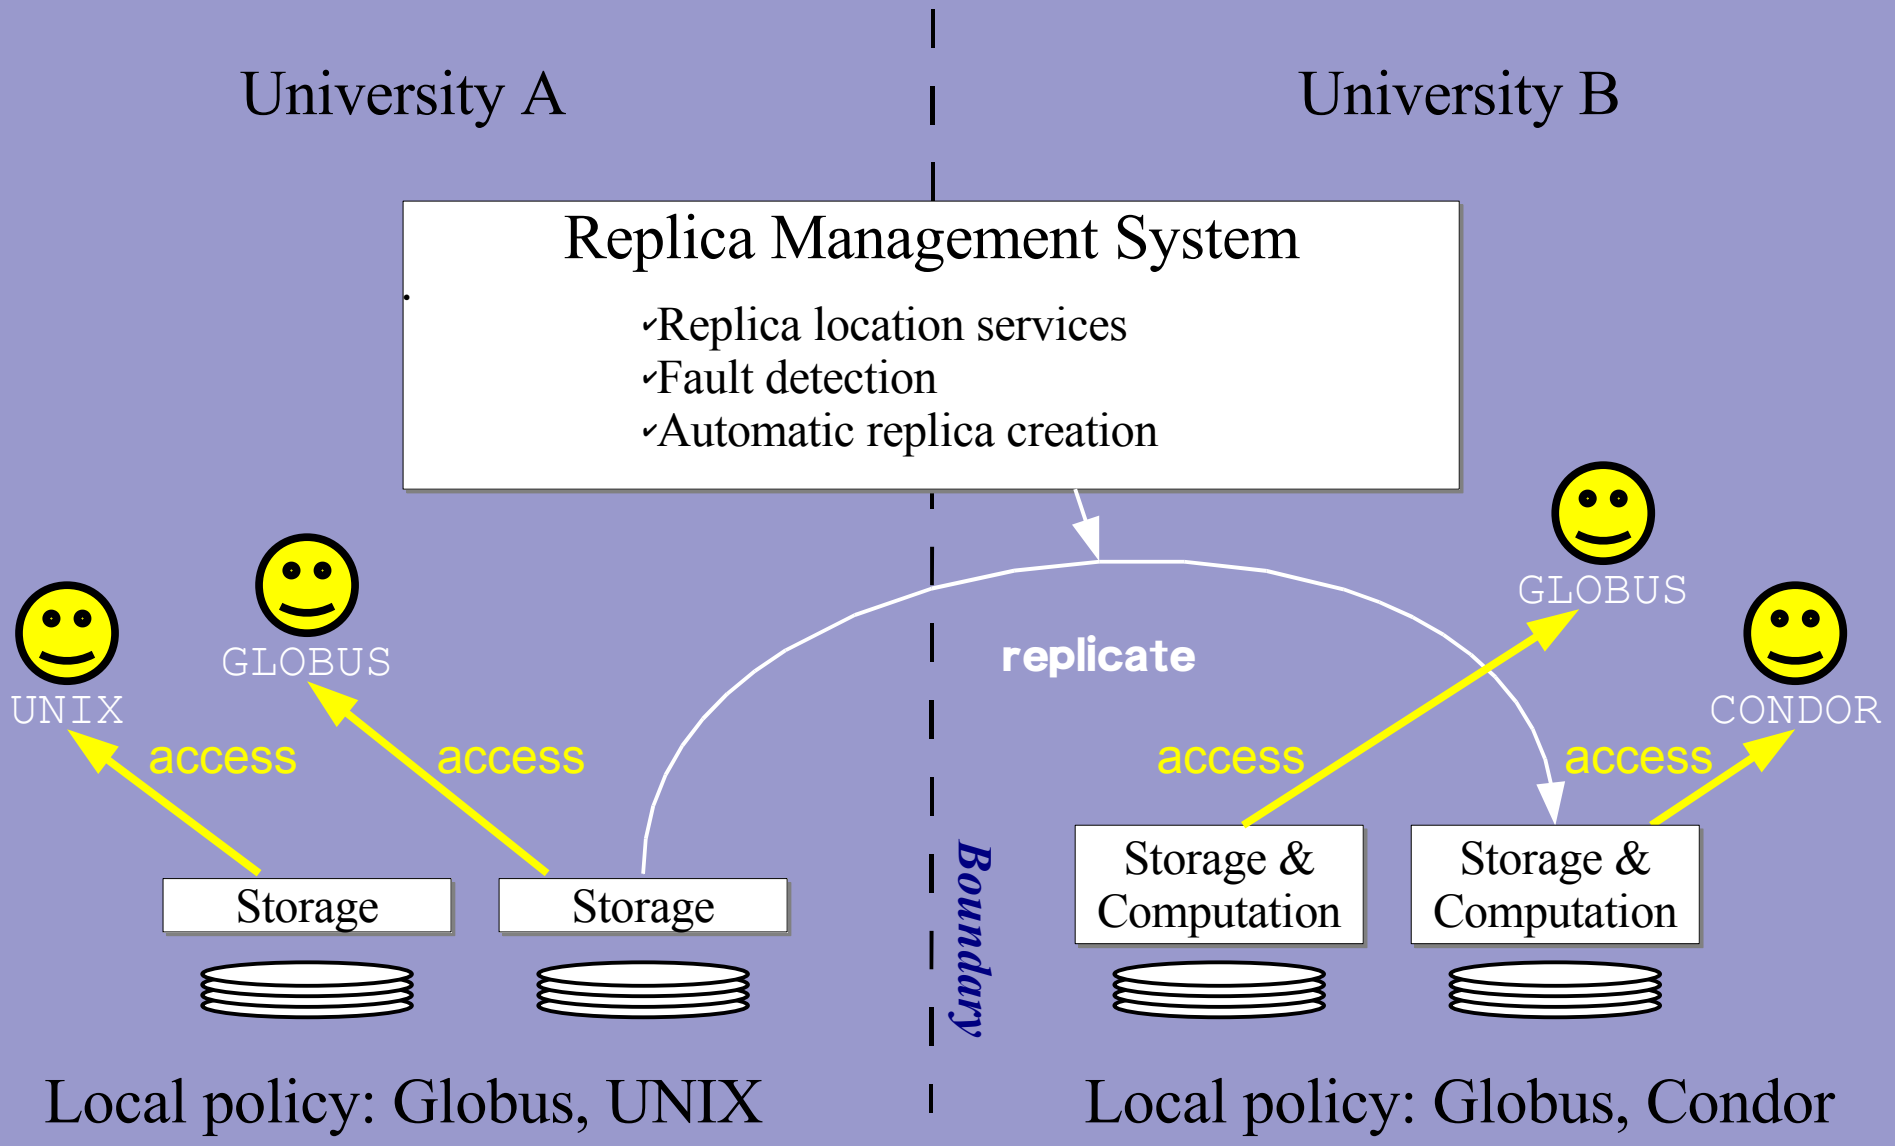

Outline

- Data description
- The GEMS architecture
- Protocol explanation
- Usage scenarios
- Conclusion

Grid Enabled Molecular Simulation

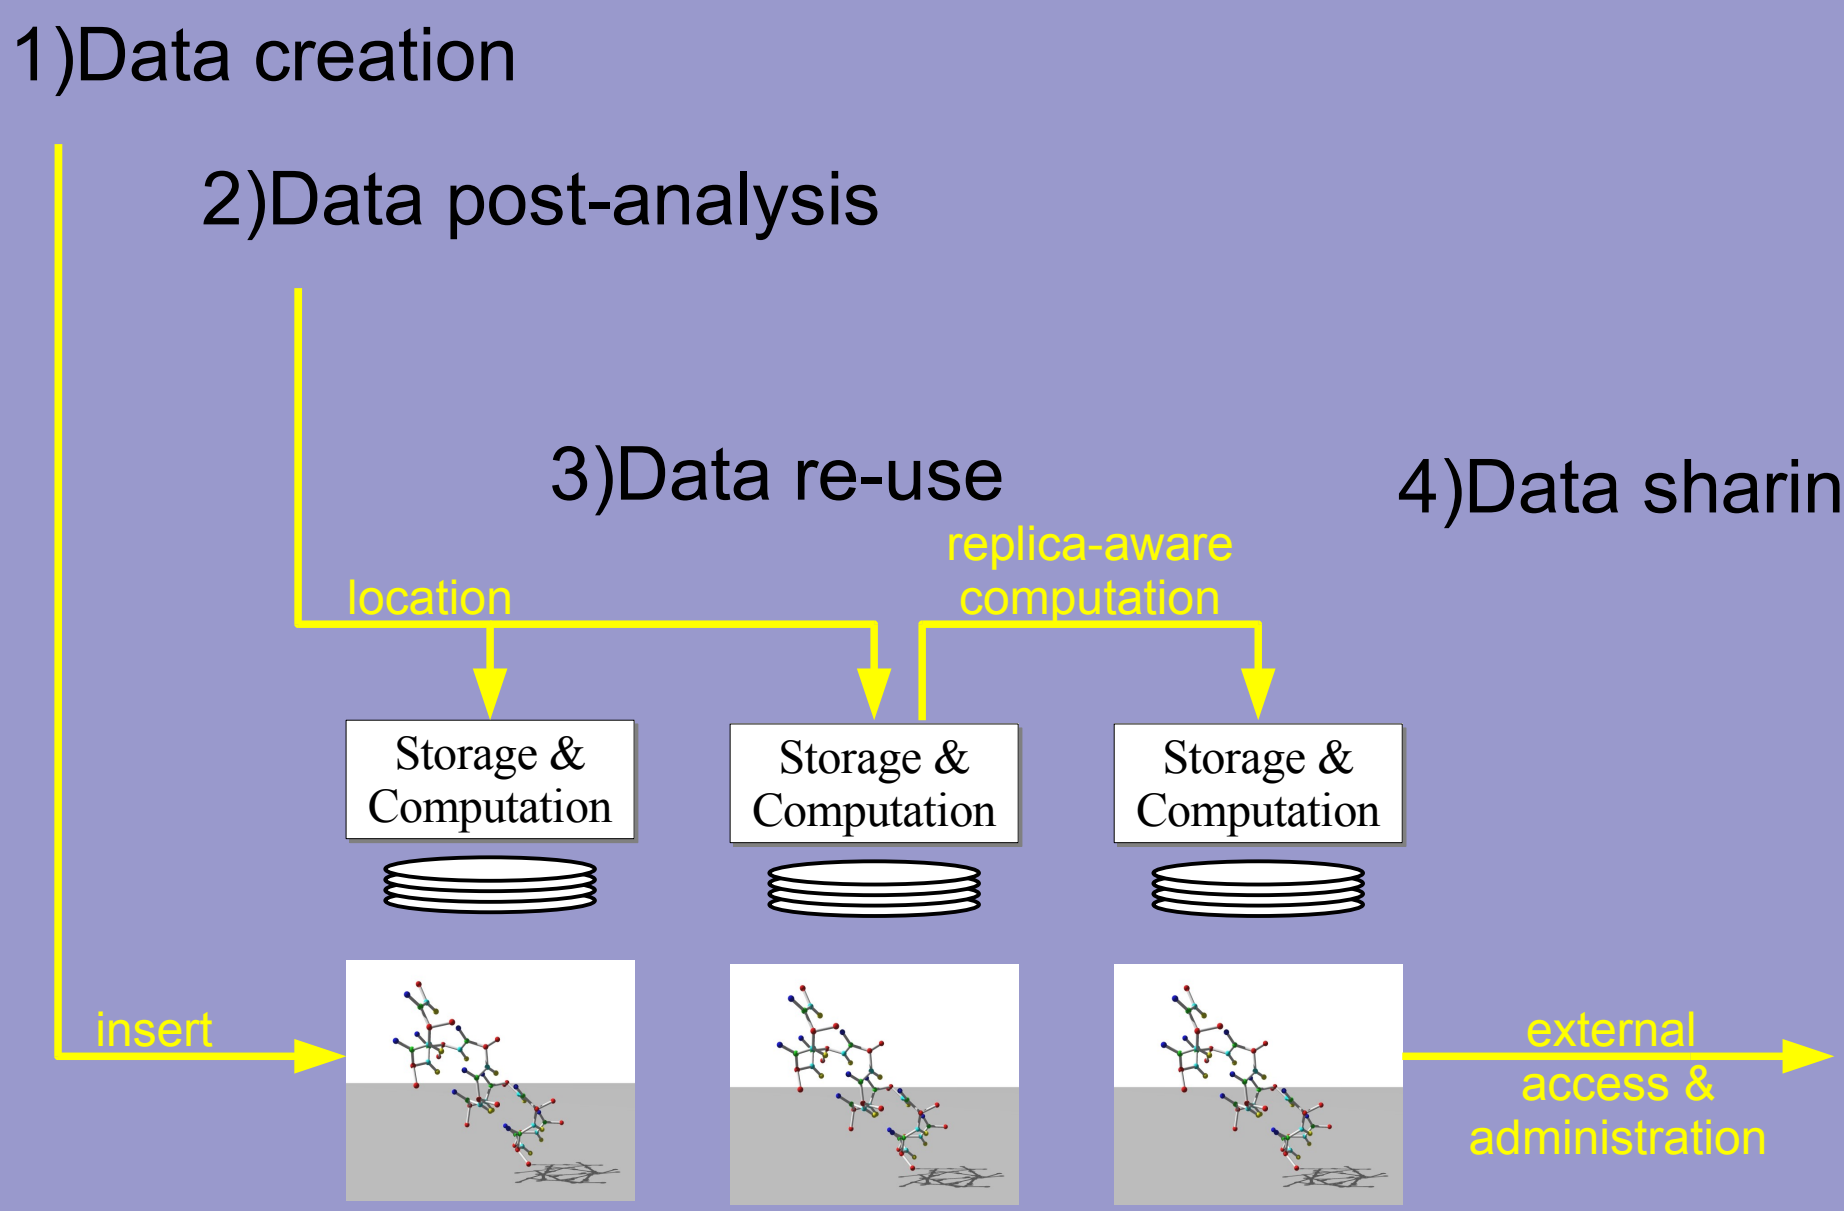

Data Properties

- Start with large data sets in directory trees
- Management aware of data properties
- Records define where they may be stored, who may access them
- Fields specify a management plan

Replica Management Database

For config C_i:

METADATA: $k_1 = v_1$, ...

FILES: $\{f_1[3] @ \{host_1, host_2, ...\} ... \}$

MAP: $\{cluster_1$ *is* $\{host_1\},$
$cluster_2$ *is* $\{host_5, host_7\}...\}$

ACL: {`HOST:RMS` *has* **RA**,
`UNIX:USER1` *has* **RA**,
`UNIX:USER2` *has* **R**}

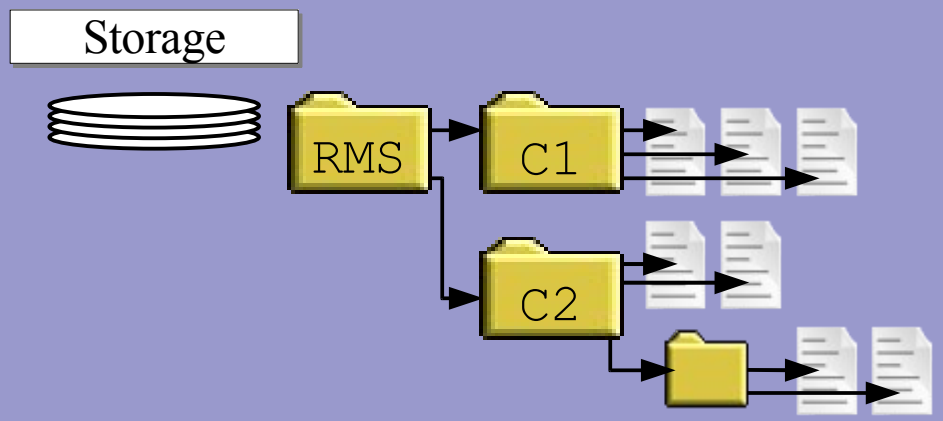

Browsing Simulations in GEMS

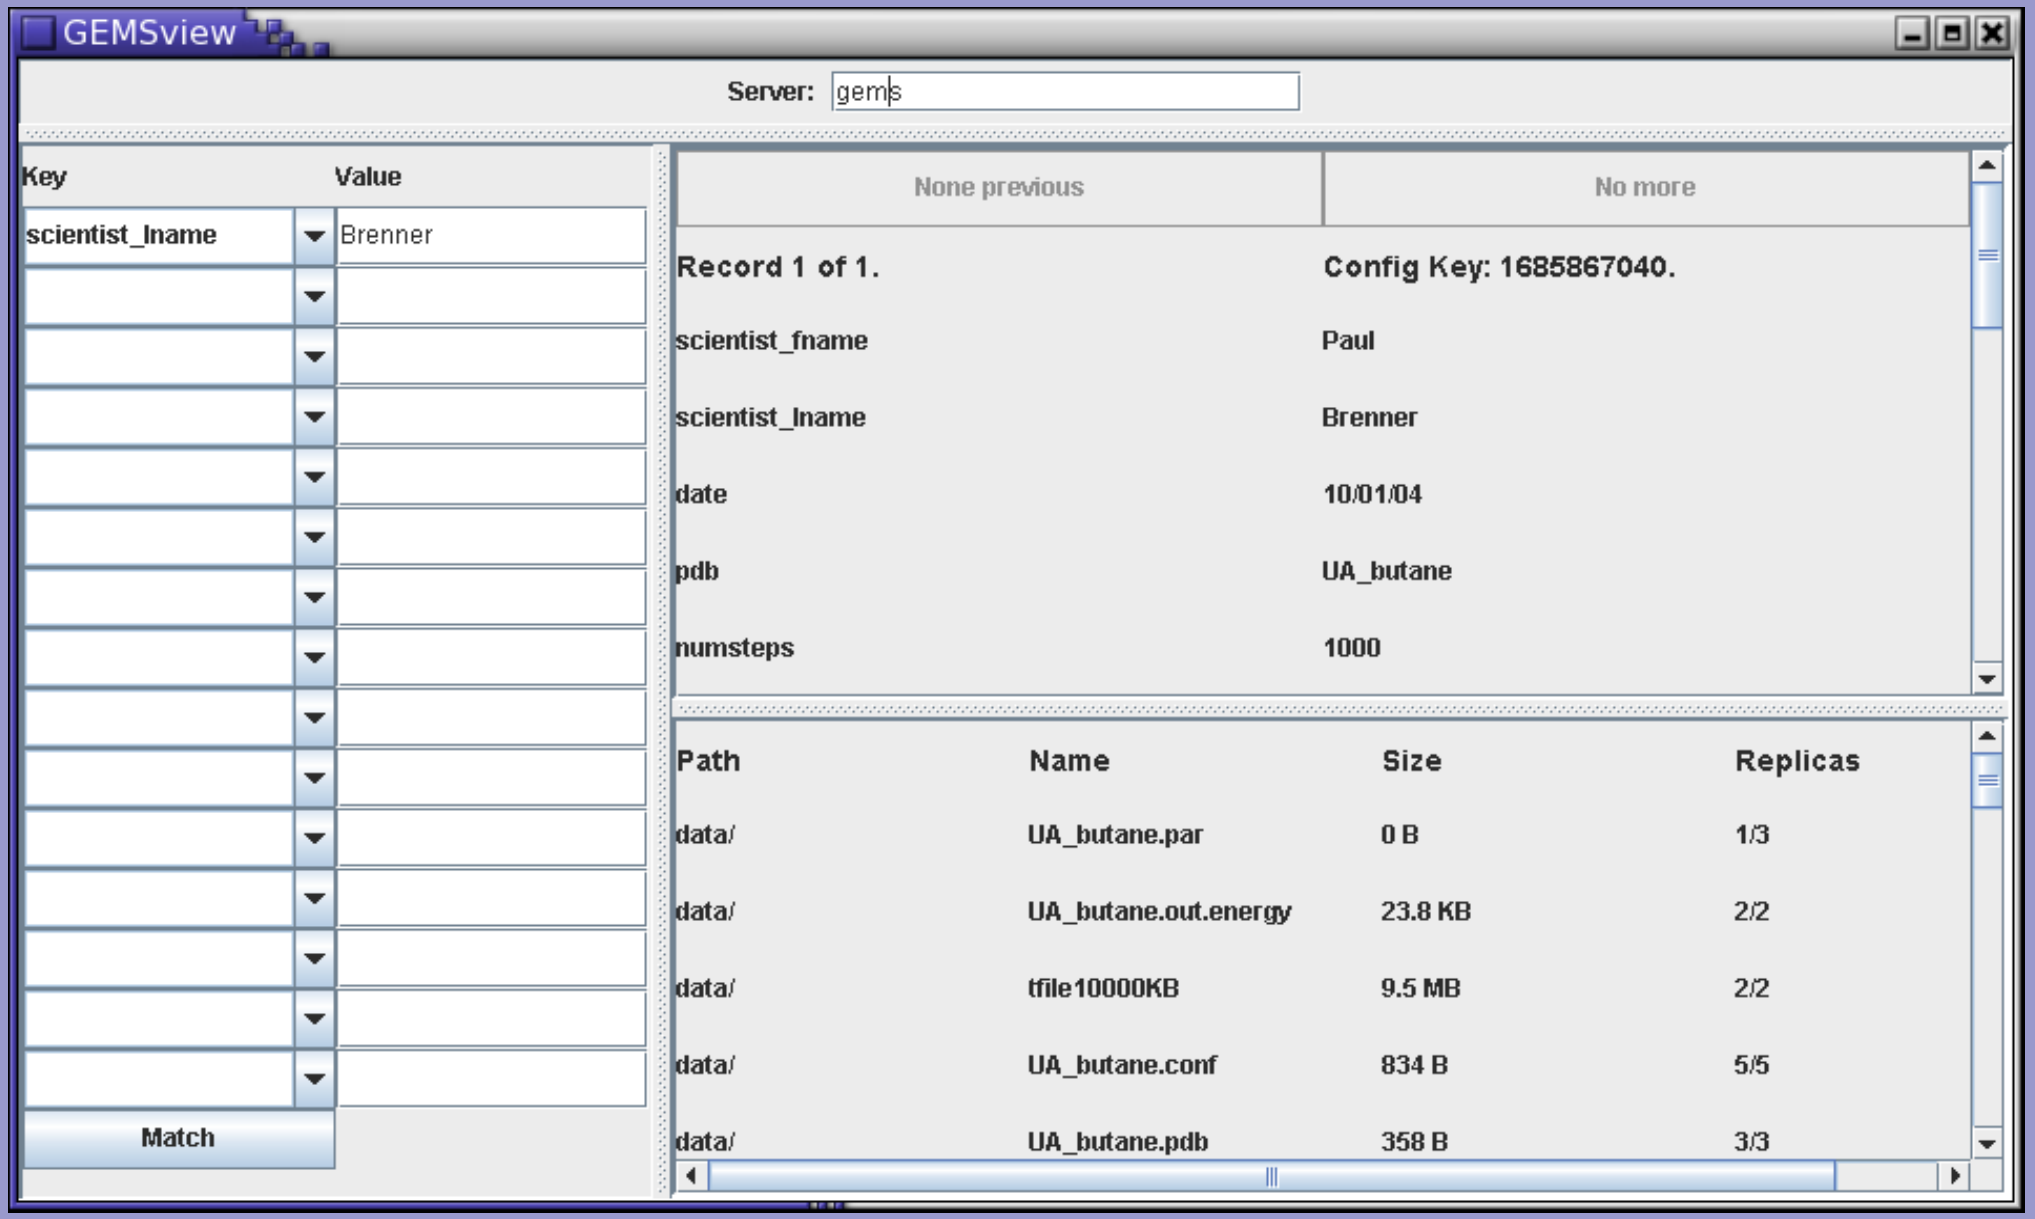

Architecture

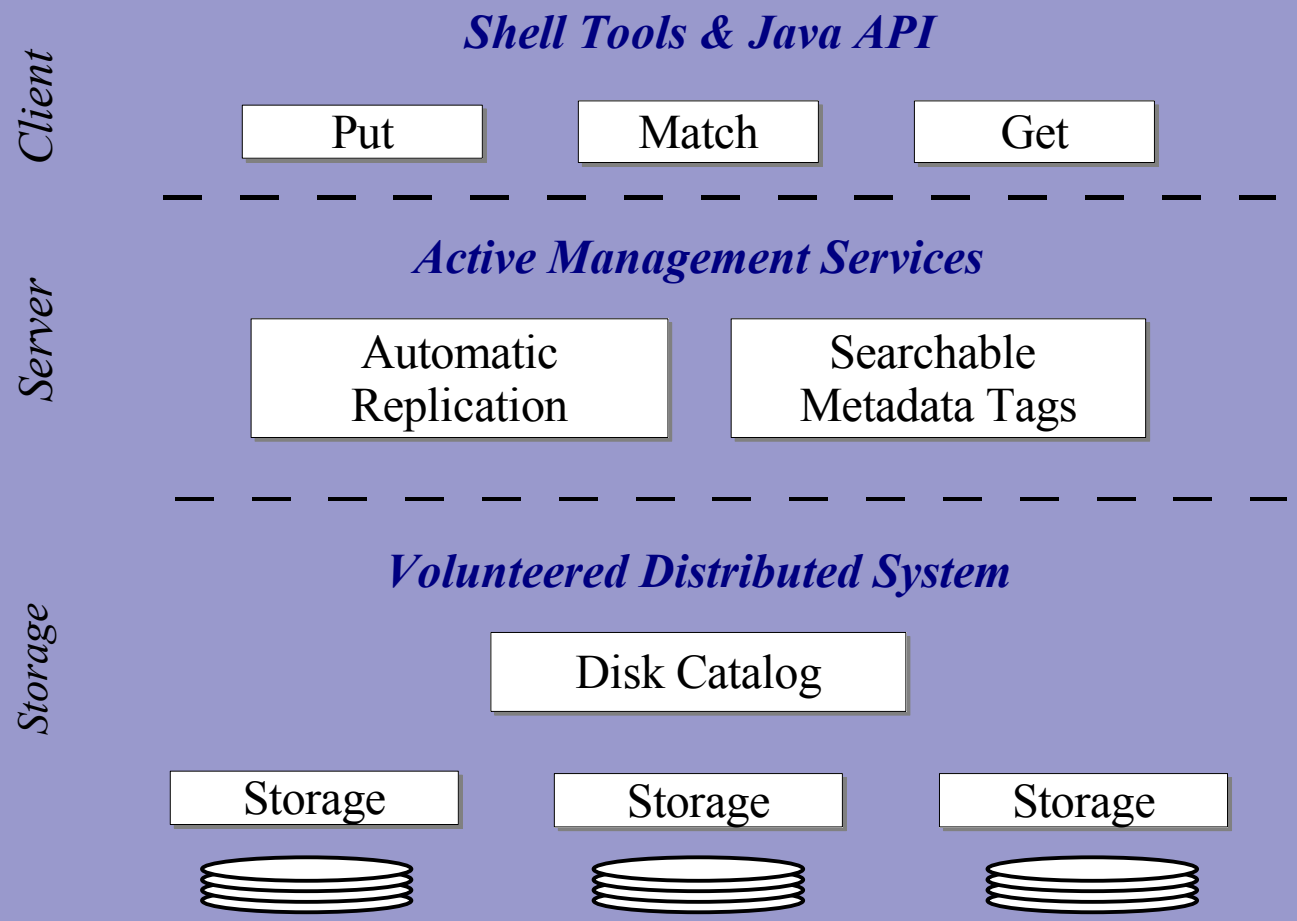

Record-wise Cluster Topology

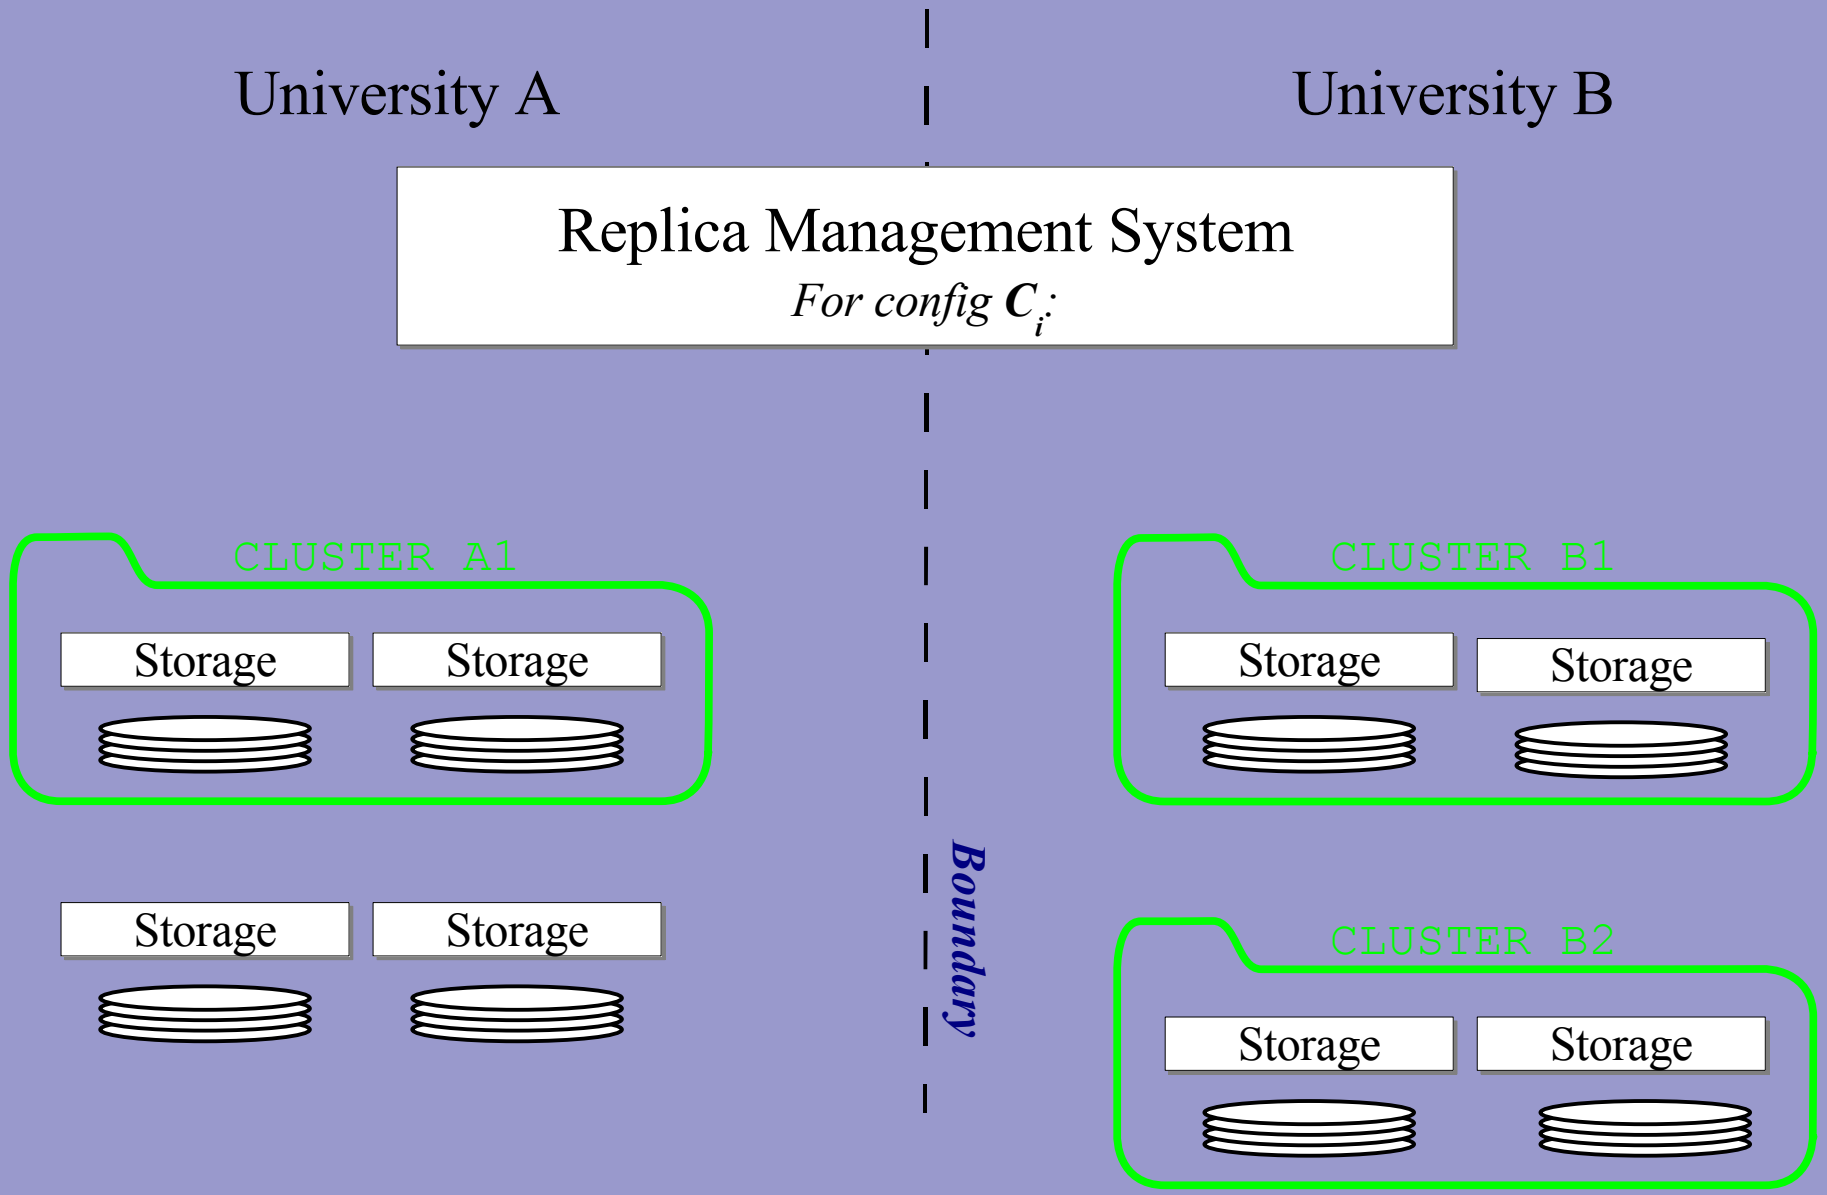

Replica Location and Access

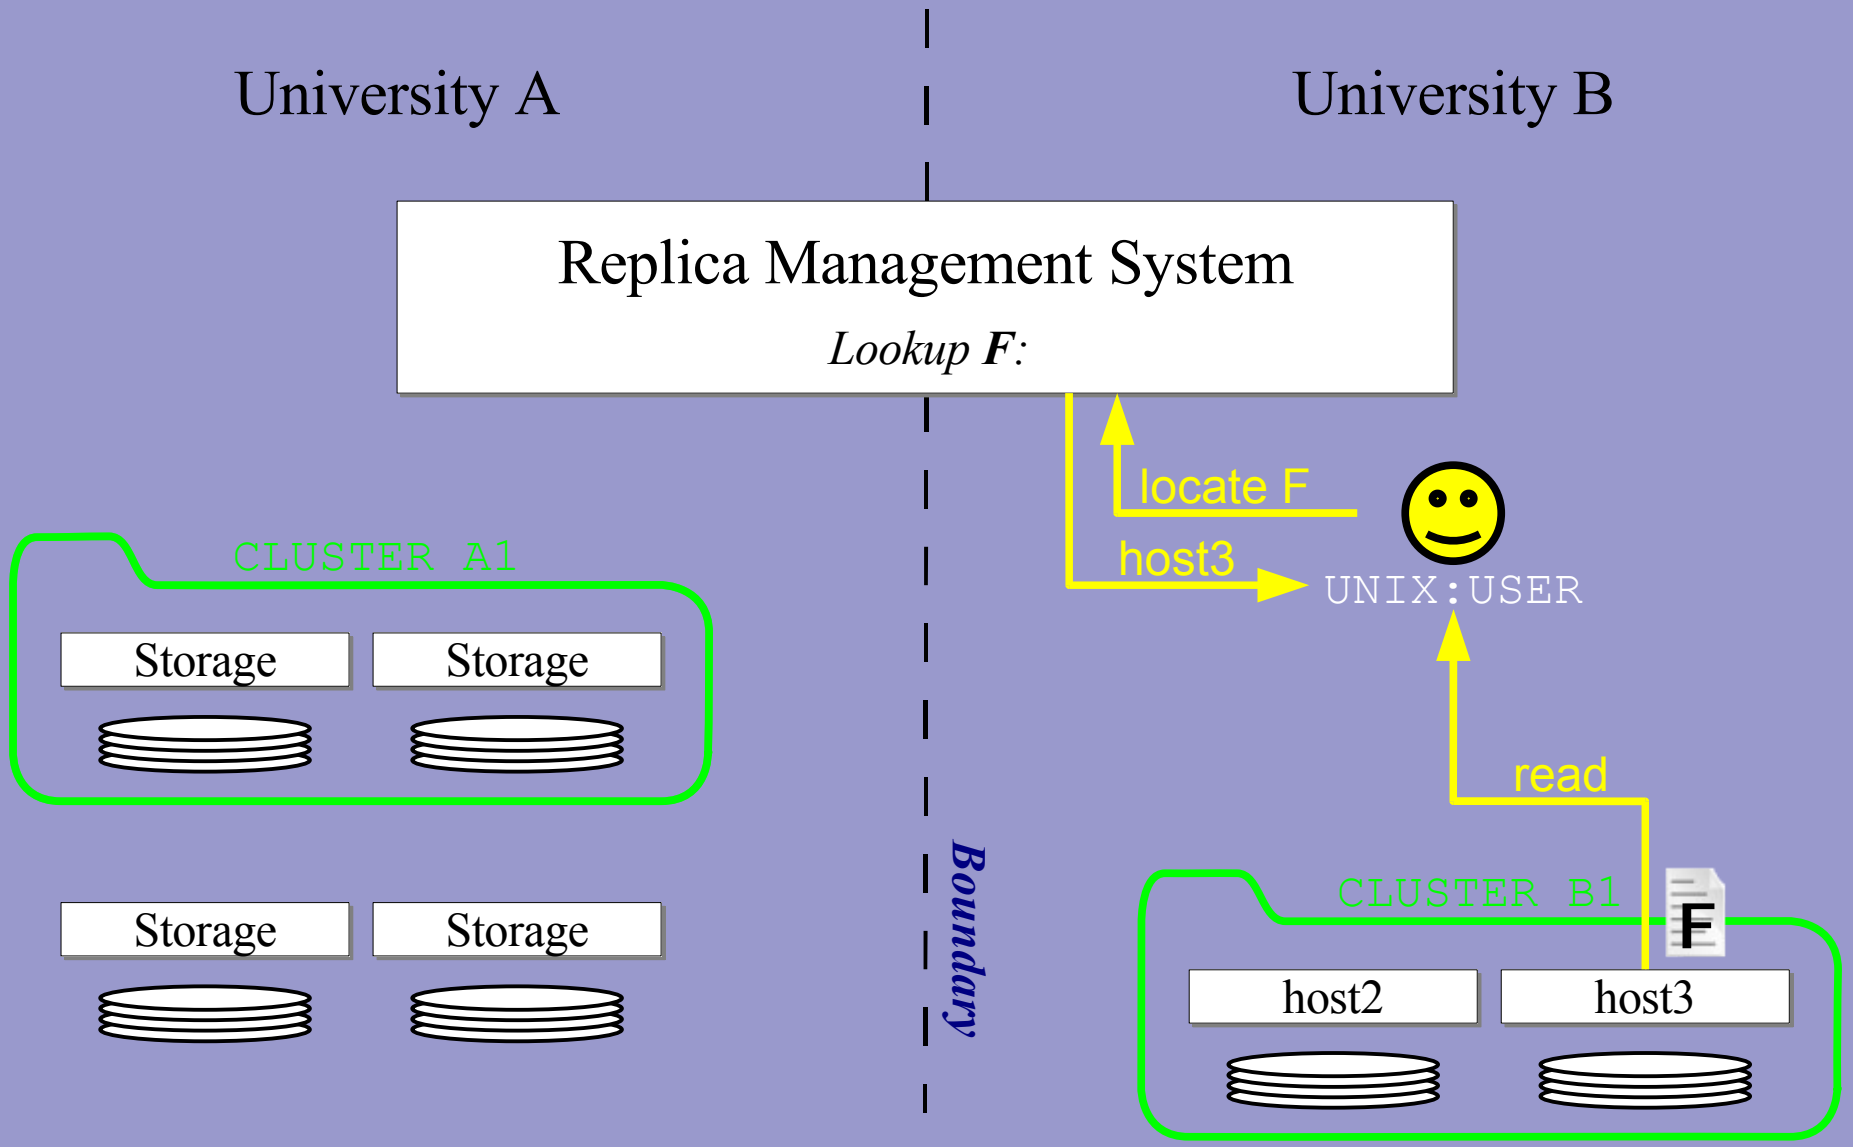

- Topology considered when locating replicas

Database Access

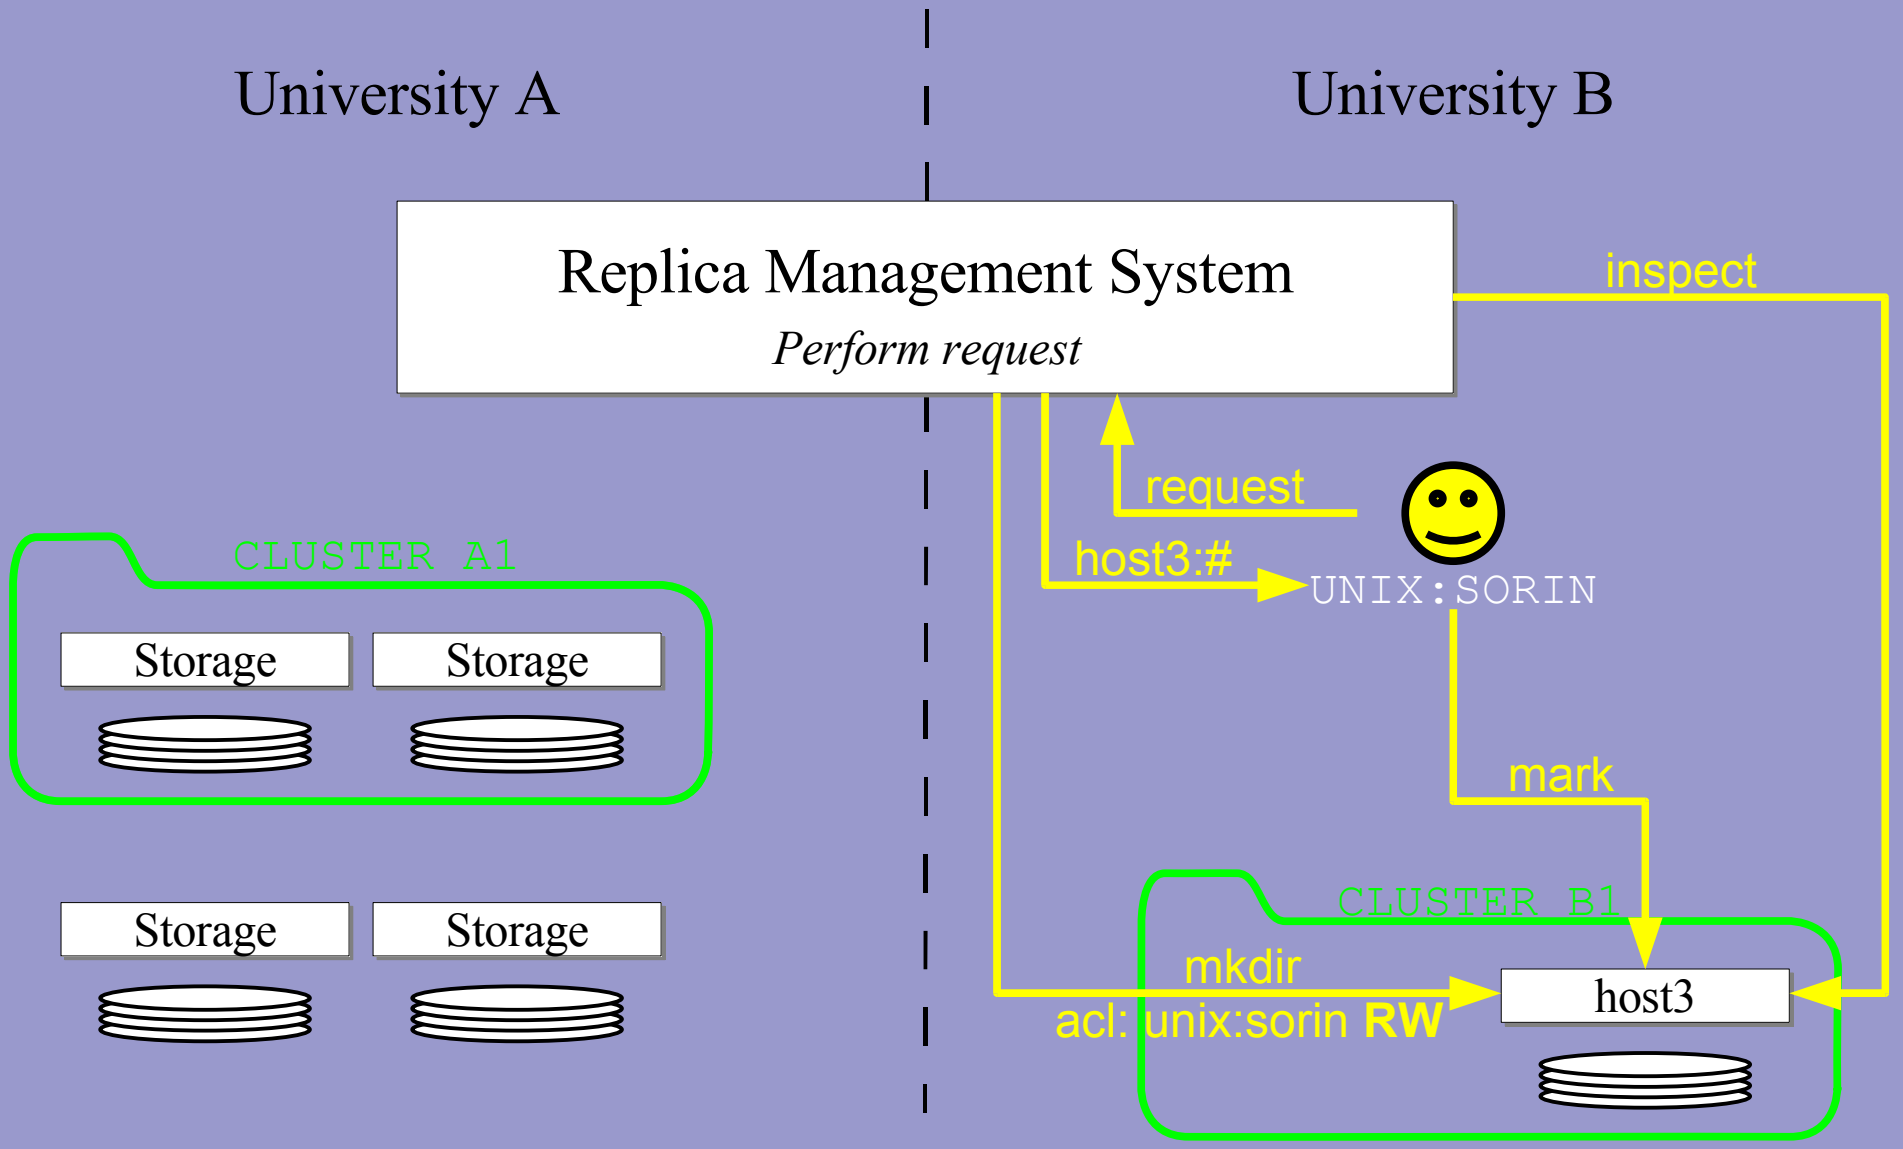

- Database affected by config deletion, modification

Channel Authentication

- RMS must determine the ability of a channel to affect a given record
- Record is dependent on storage servers
- RMS may rely on storage servers to “speak for” users
- Users specify permitted storage servers via per-record storage map

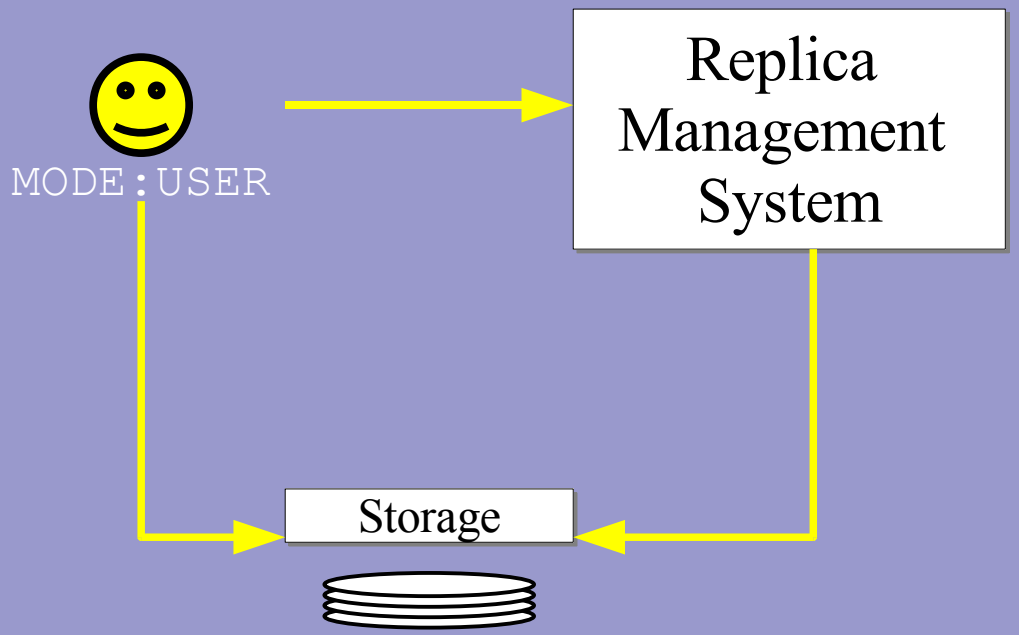

- Mode, user are opaque and are treated as symbols

Rendition Protocol

- Client requests challenge for config *C* on channel *N*
- RMS creates challenge using eligible host and ACL for *C*
- Client meets challenge
- RMS checks challenge
- Channel is authenticated for config *C*

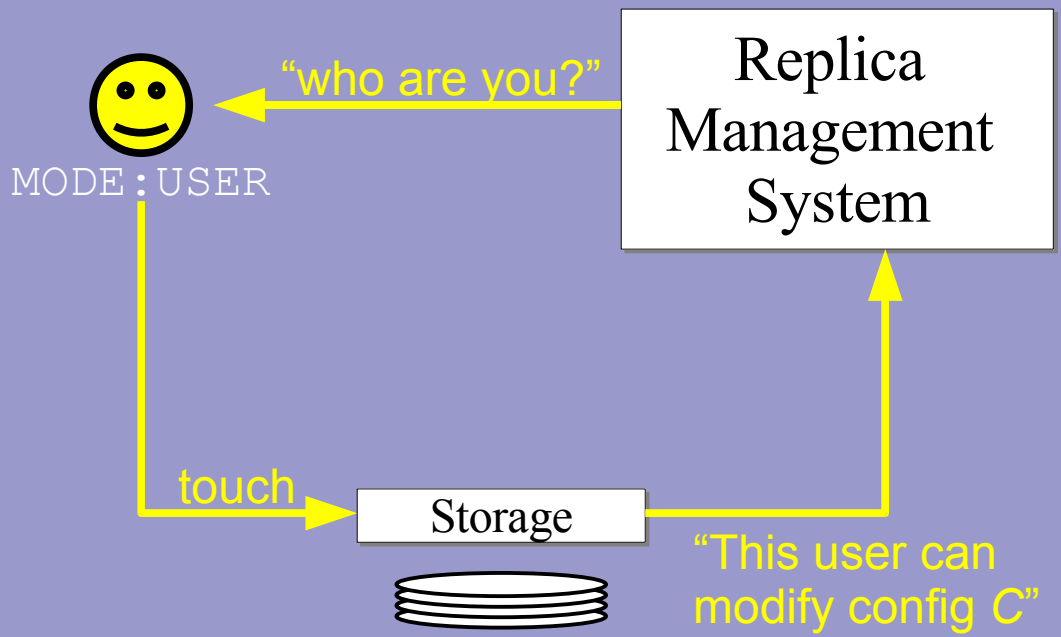

Cross-boundary Scenarios

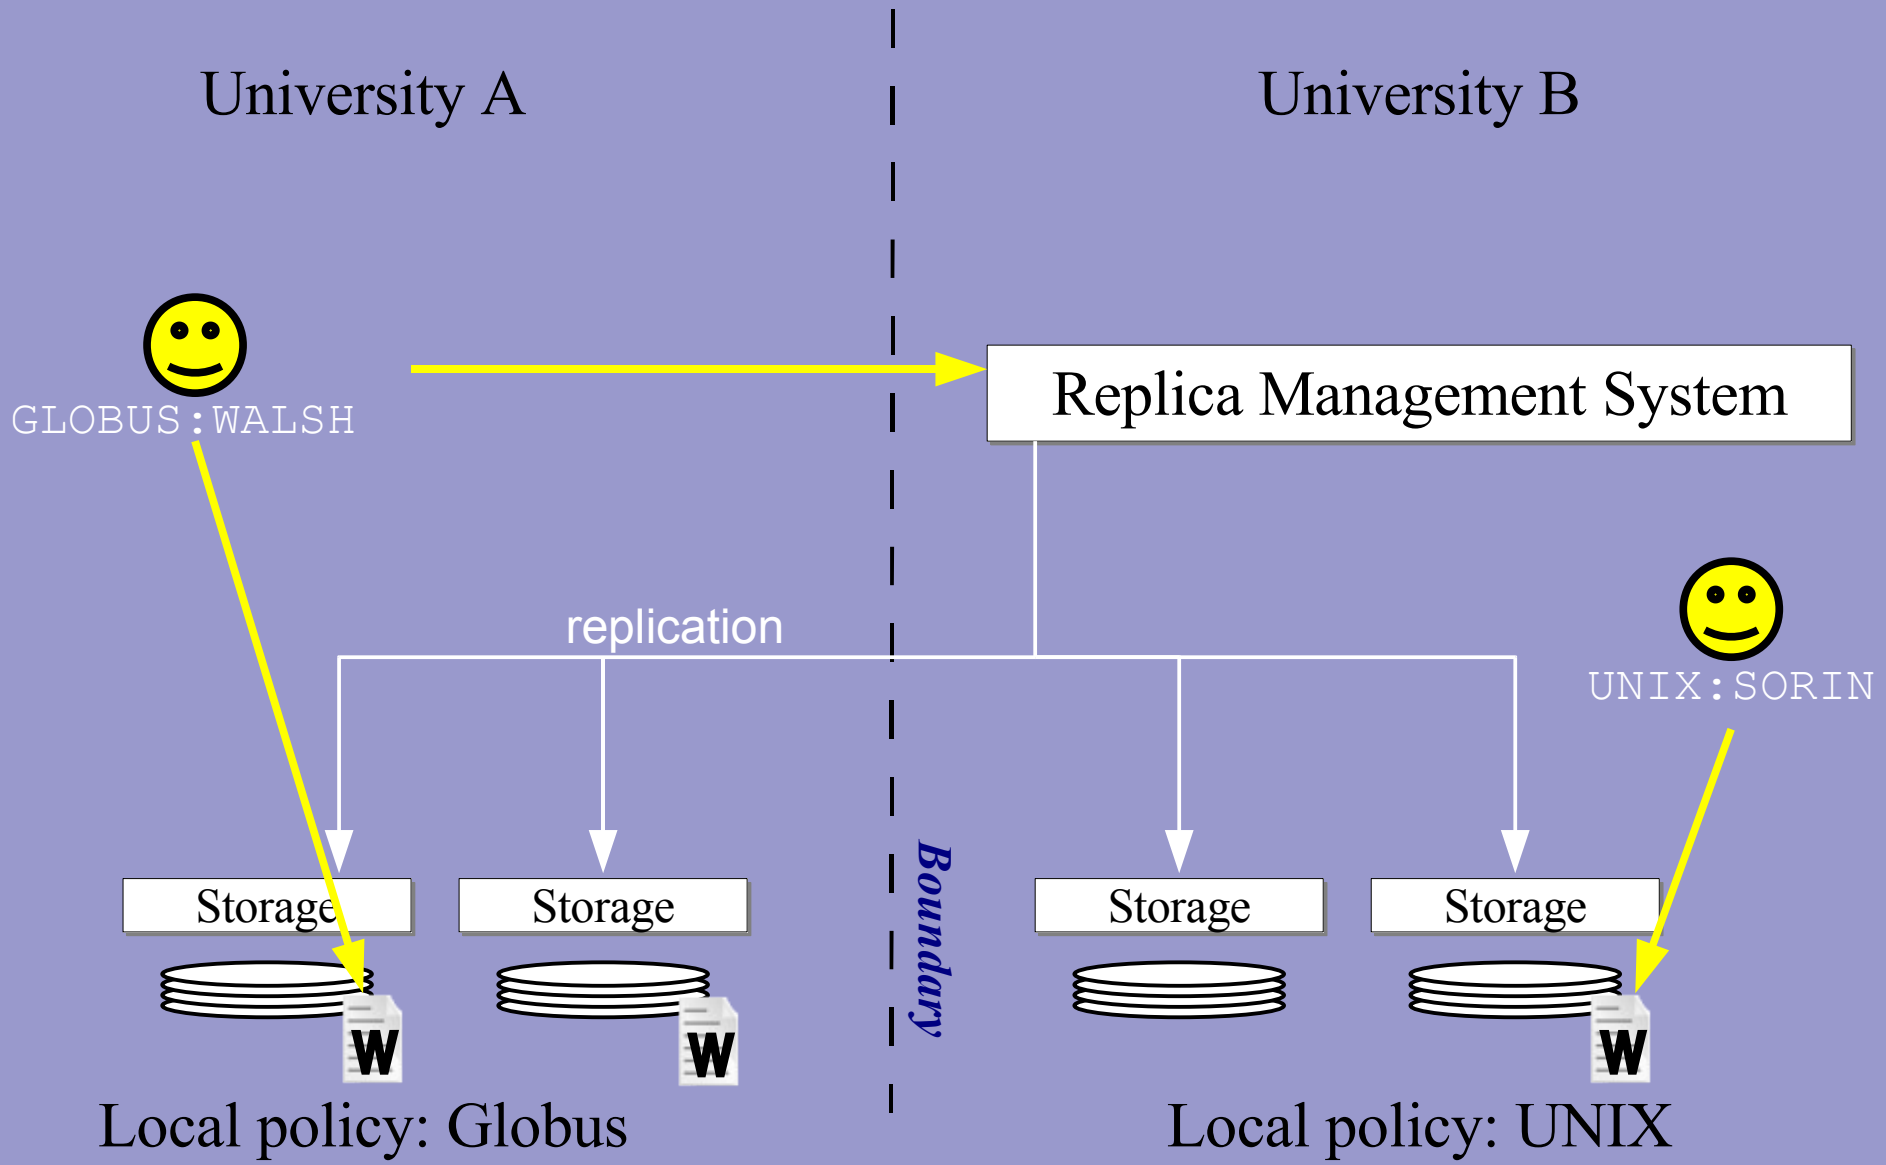

Highlights

- RMS treats users as symbols
- Users may administer data on machines to which they cannot directly authenticate
- Remote replicas are very useful: running remote jobs, collaboration
- Data storage system is essentially as secure as the data creation system or storage servers
- Secure replica placement becomes crucial

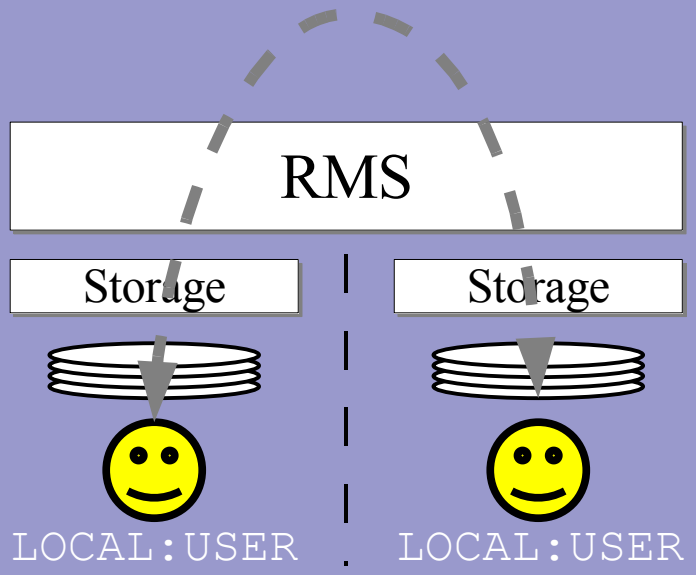

- New collaboration technique built around a replica system

Conclusion

- Heavy emphasis on widely distributed, uncontrolled nature of storage fabric.
- Able to authenticate users indirectly via rendition protocol.

- Generosity and Gluttony in GEMS: Grid Enabled Molecular Simulations. Justin Wozniak, Paul Brenner, Douglas Thain, Aaron Striegel, Jesus Izaguirre In Proceedings of IEEE High Performance Distributed Computing, July 2005.
- *GEMS Home:* `http://gipse.cse.nd.edu/GEMS`

UNIVERSITY OF NOTRE DAME