



dCache, the Update

Patrick Fuhrmann

for the dCache people





Responsibility, dCache

Patrick Fuhrmann Rob Kennedy

Responsibility, SRM

Timur Perelmutov

Core Team (Desy and Fermi)

Jon Bakken

Mathias de Riese

Micheal Ernst

Alex Kulyavtsev

Birgit Lewendel

Dmitri Litvintsev

Tigran Mrktchyan

Martin Radicke

Neha Sharma

Vladimir Podstavkov

External

Development

Nicolo Fioretti, BARI

Abhishek Singh Rana, SDSC

Support and Help

Maarten Lithmaath, CERN

Owen Synge, RAL





dCache, the Idea

The story so far ...

what's new , what's coming next

Feature Track

SRM V2

Replica Manager

Information Provider

File Hopping

Smart Tape System flushing

Generic Tertiary Storage Interface

dCache partitioning

VOMS integration

Performance Track

medium term strategy

Multiple I/O queues

Multiple Pnfs (file systems)

Chimera





Distributed Peta Byte Disk Store with single rooted filesystem providing posix like and wide area access protocols.

Distributed cache system to optimize access to Tertiary Storage Systems

Grid Storage Element coming with standard data access protocols, Information Provider Protocols and Storage Resource Manager.





The story so far ...





What's new ?

What's coming next ?





Feature Track





Please refer to talk by Alexander Kulyavtsev

Resilient dCache

- Controls number of copies for each dataset in dCache
- Makes sure $n < \text{copies} < m$
- Adjusts replica count on pool failures
- Adjusts replica count on scheduled pool maintenance
- Makes use of local disk space when running on farm nodes
- **doesn't work with HSM back-end yet**





Please refer to talk by Timur Perelmutov

Storage Resource Manager V2

- Directory methods/function available
- Explicit Space Management is in development phase





Please refer to talk by Abhishek Singh Rana

gPLAZMA (grid-aware PLuggable AuthoriZation MAnagement): Introducing RBAC (Role Based Access Control) Security in dCache

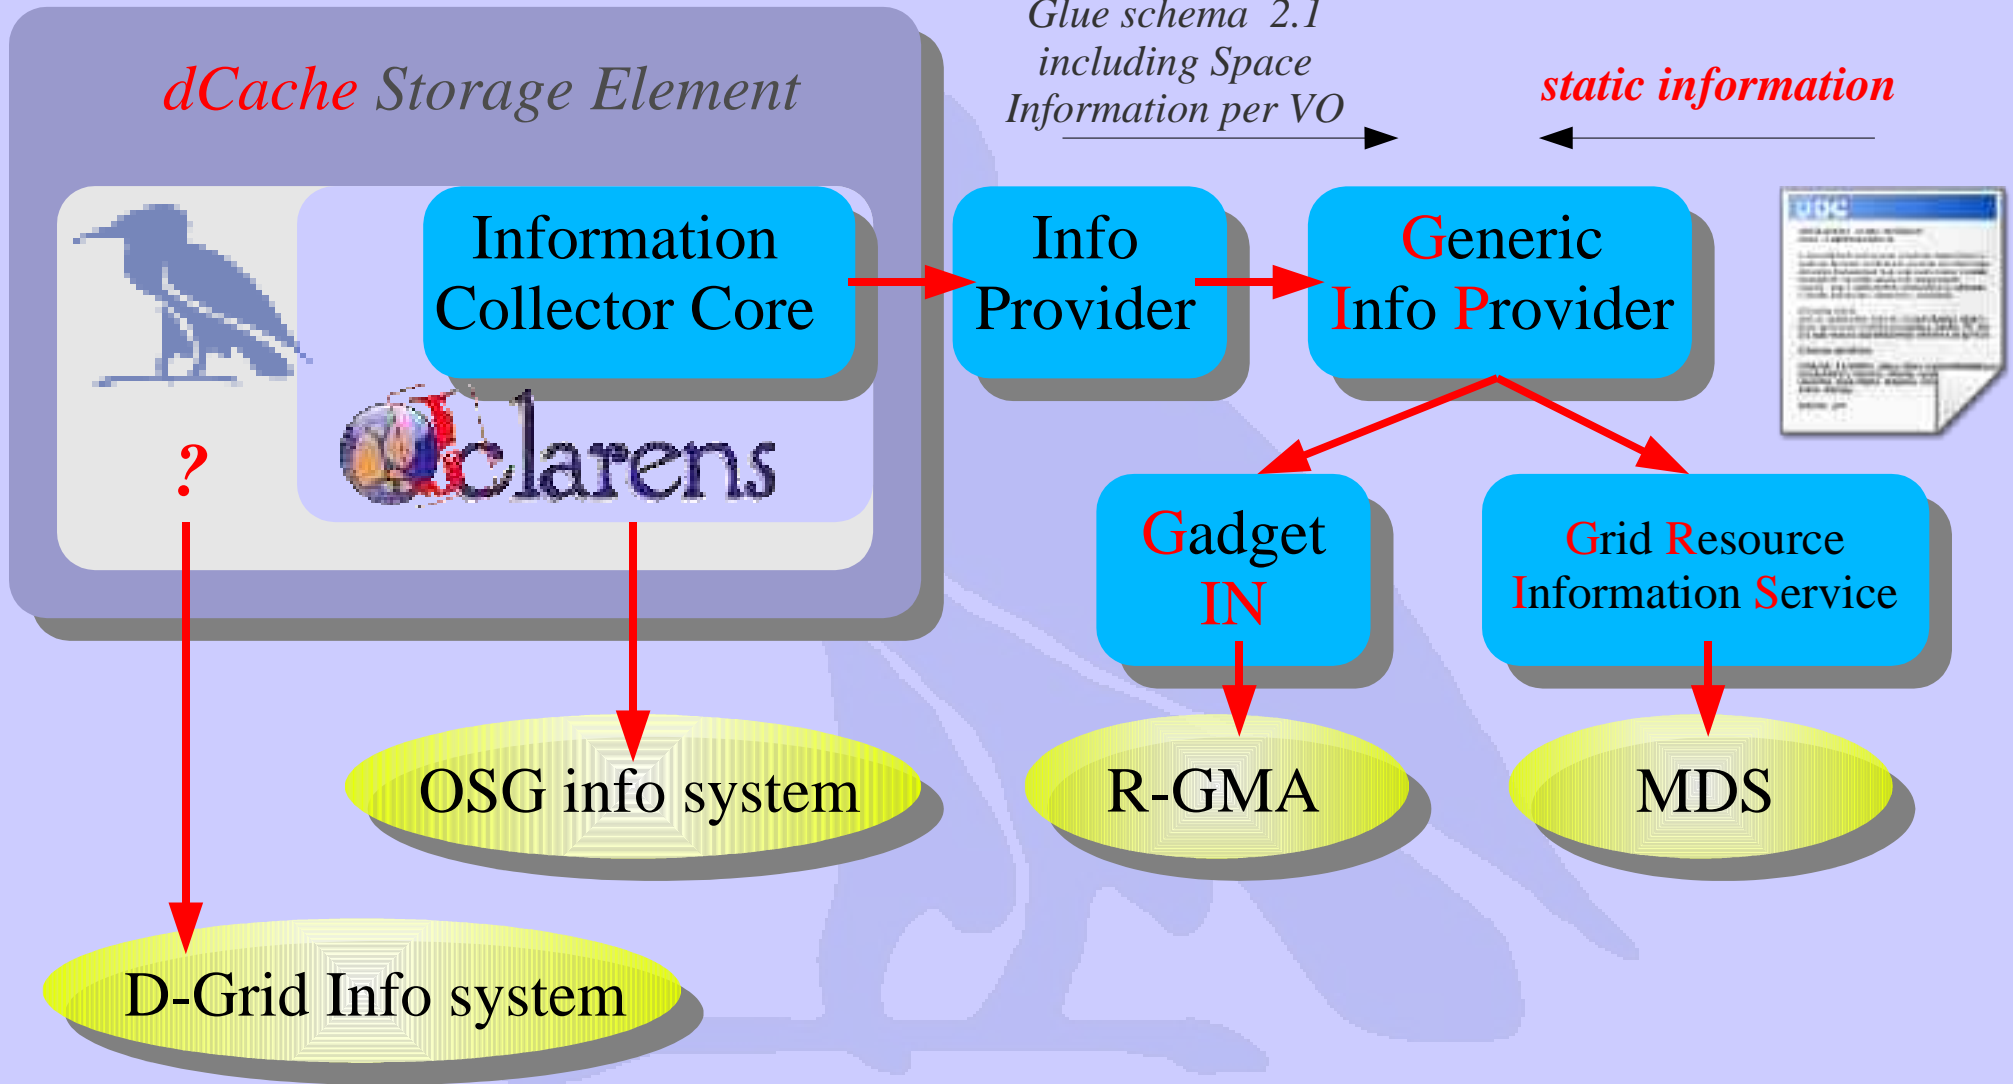
Virtual Organization Membership Service Integration

- Pluggable authorization
- Backward compatible to current system.
- Storage Authorization Call-outs for SRM and GridFtp





by courtesy of Nicolò Fioretti

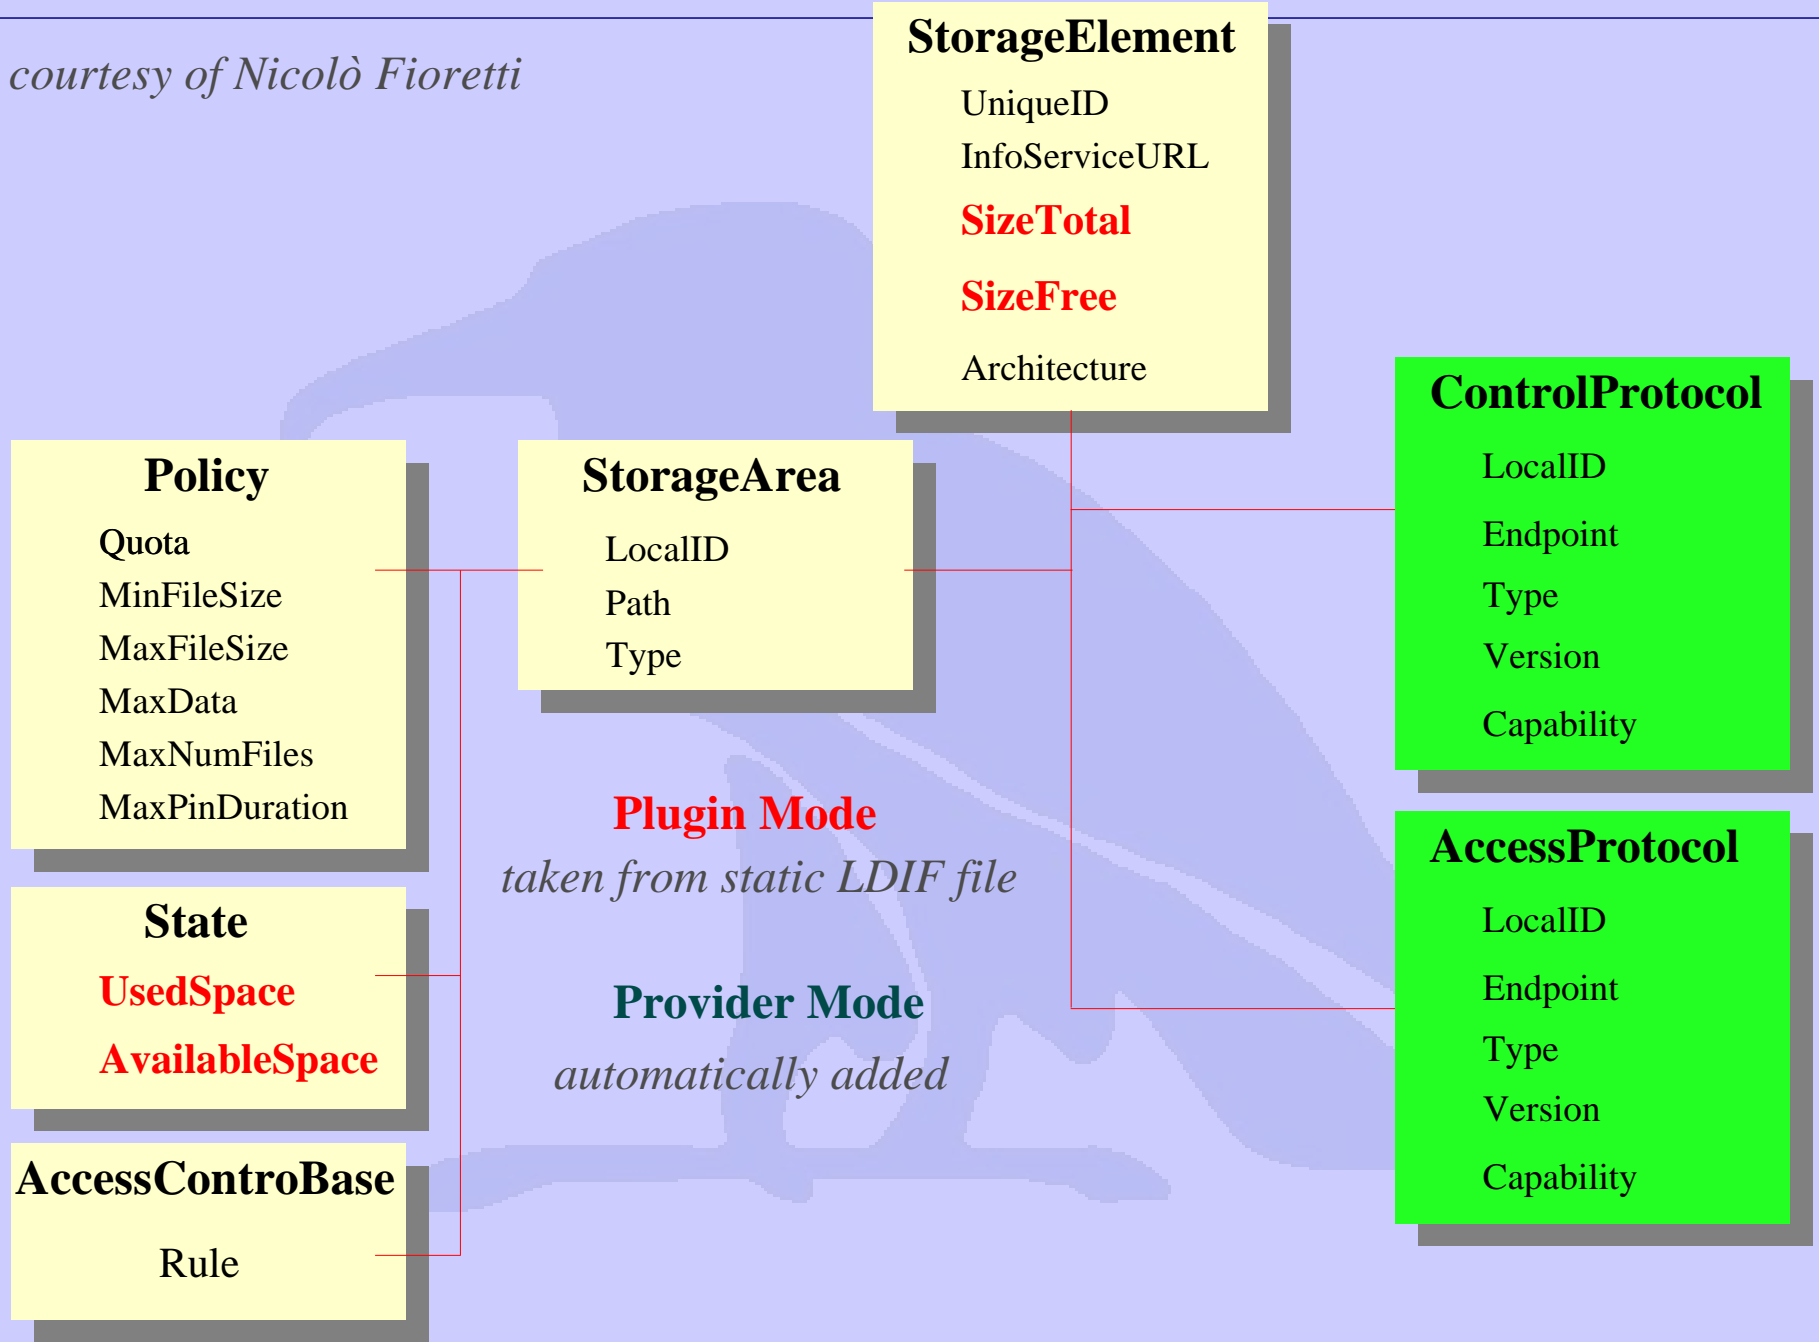


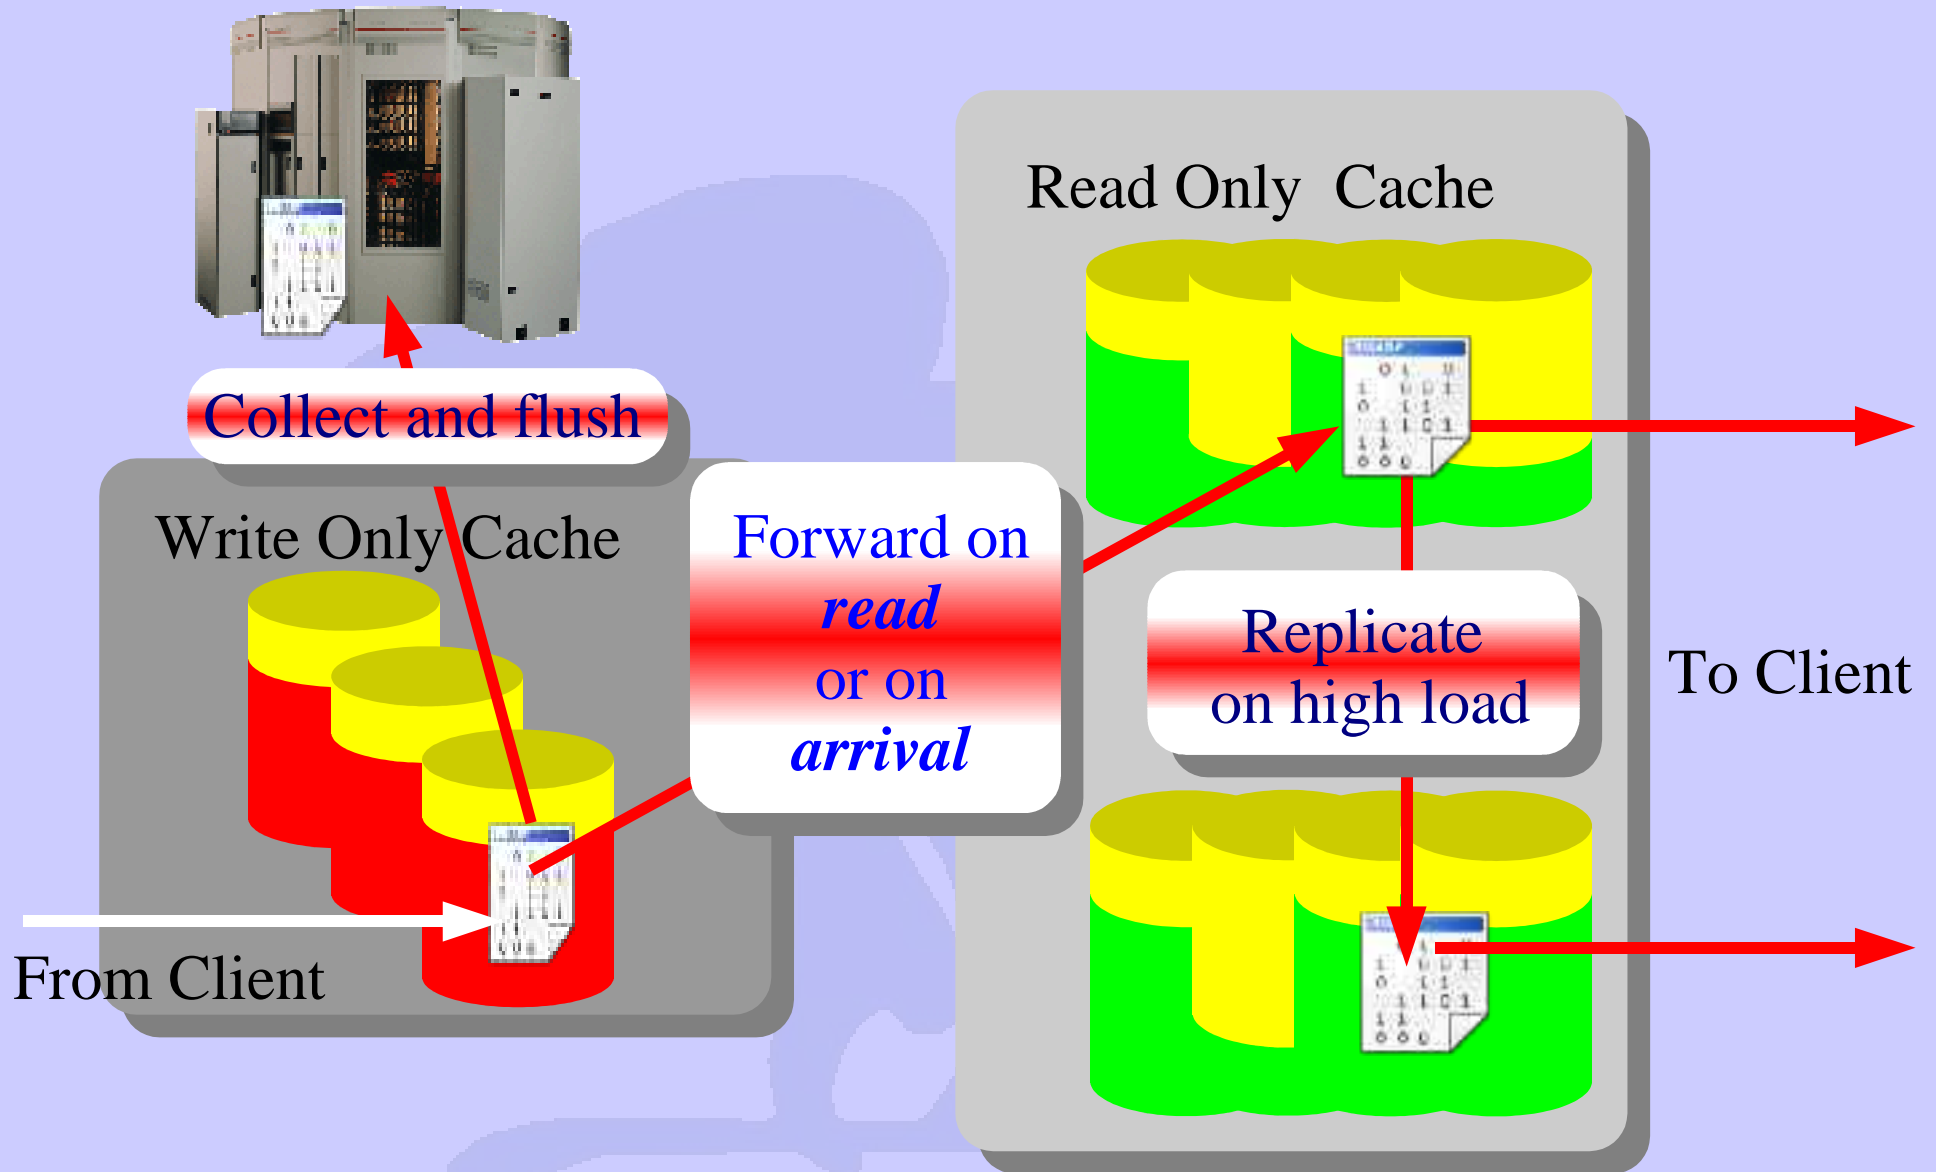
with many thanks to Jean-Philippe Baud and Lawrence Fields





by courtesy of Nicolò Fioretti





even more dynamic when using API





One dCache instance, but pool-groups with different characteristics

*prefer balanced
movers*

*prefer filling empty pools
first*

- *small file set*
- *high throughput*
- *all files on all pools*

Resilient Pool Group,
*file multiplicity centrally
managed,
no HSM back end*





Centrally controlled flushing of precious datasets to back-end tape system(s)

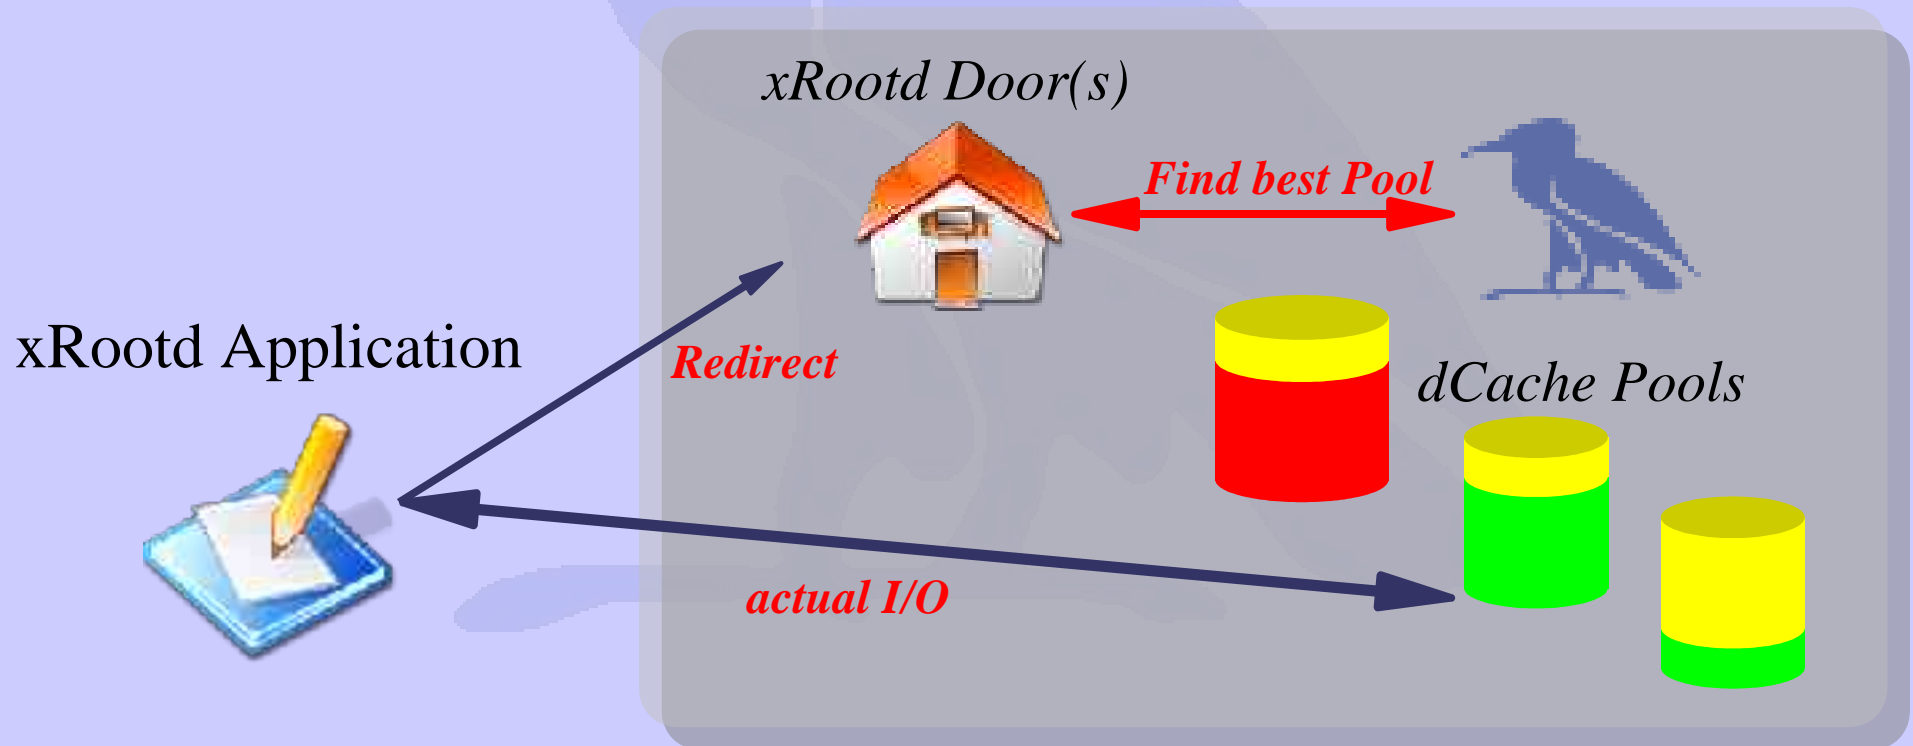
- determines number of concurrent streams to HSM system.
- never allow client data flowing onto pools, while same pool is flushing.
- never allow pool flushing, while data coming in from clients.
- event driven API for fine grained customization of HSM flush.





xrootd implemented :


- *regular I/O ok, including forwarding*
- *no asynchronous mode yet*
- *no security yet*
- *we need help from experts, specs sometimes not clear enough*





Performance Track







Mover Queue(s)

dCap

gridFtp

xRootd

Application

dccp

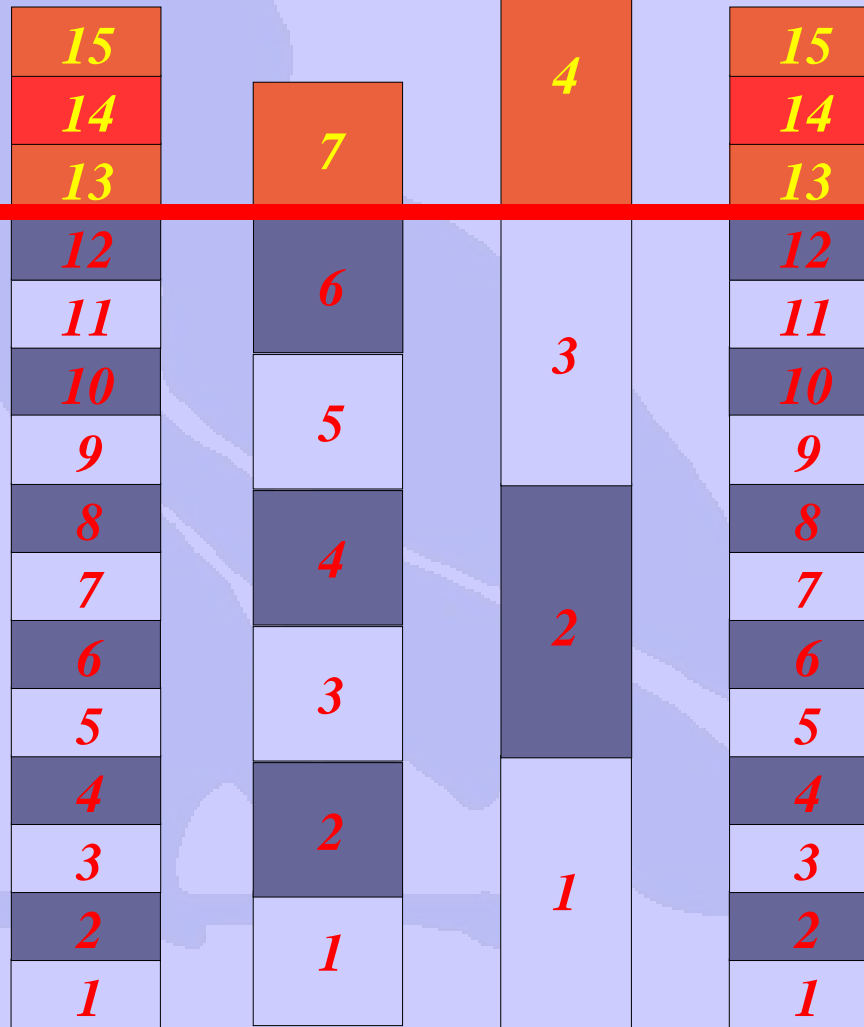
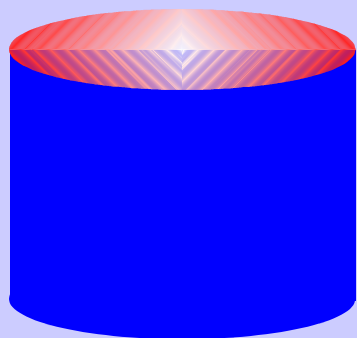
Application

Cost = 1

Waiting

Active

Pool





Separate Pnfs Server

