


# Increasing Performance and Scalability of dCache Storage Resource Manager

Timur Perelmutov<sup>1</sup> for the dCache team<sup>1,2,3</sup>

<sup>1</sup>Fermi National Accelerator Laboratory, <sup>2</sup>Nordic Data Grid Facility,

<sup>3</sup>Deutsches Elektronen-Synchrotron



## Scalability and Performance issues

Often requests served to clients that have just expired or are about to expire. Problems are aggravated by immediate retries

High CPU load due to GSI Authentication and Credential Delegation

High volume of blocking unscheduled SRM LS operations cause high CPU load on SRM and PnfsManager and starvation of the network connection slots

Large load due to the internal request timeouts and retries.

SRM as a single point of entry into the storage on the grid is a natural bottleneck

dCache SRM design may have additional obstacles to scaling

## Solutions

Require clients to specify the request lifetimes and use exponential back-off in case of server overloads (WLCG wide agreements)

Cache public and private key pairs used in GSI authentication  
GSI handshake CPU usage went down from 97% to 38 %

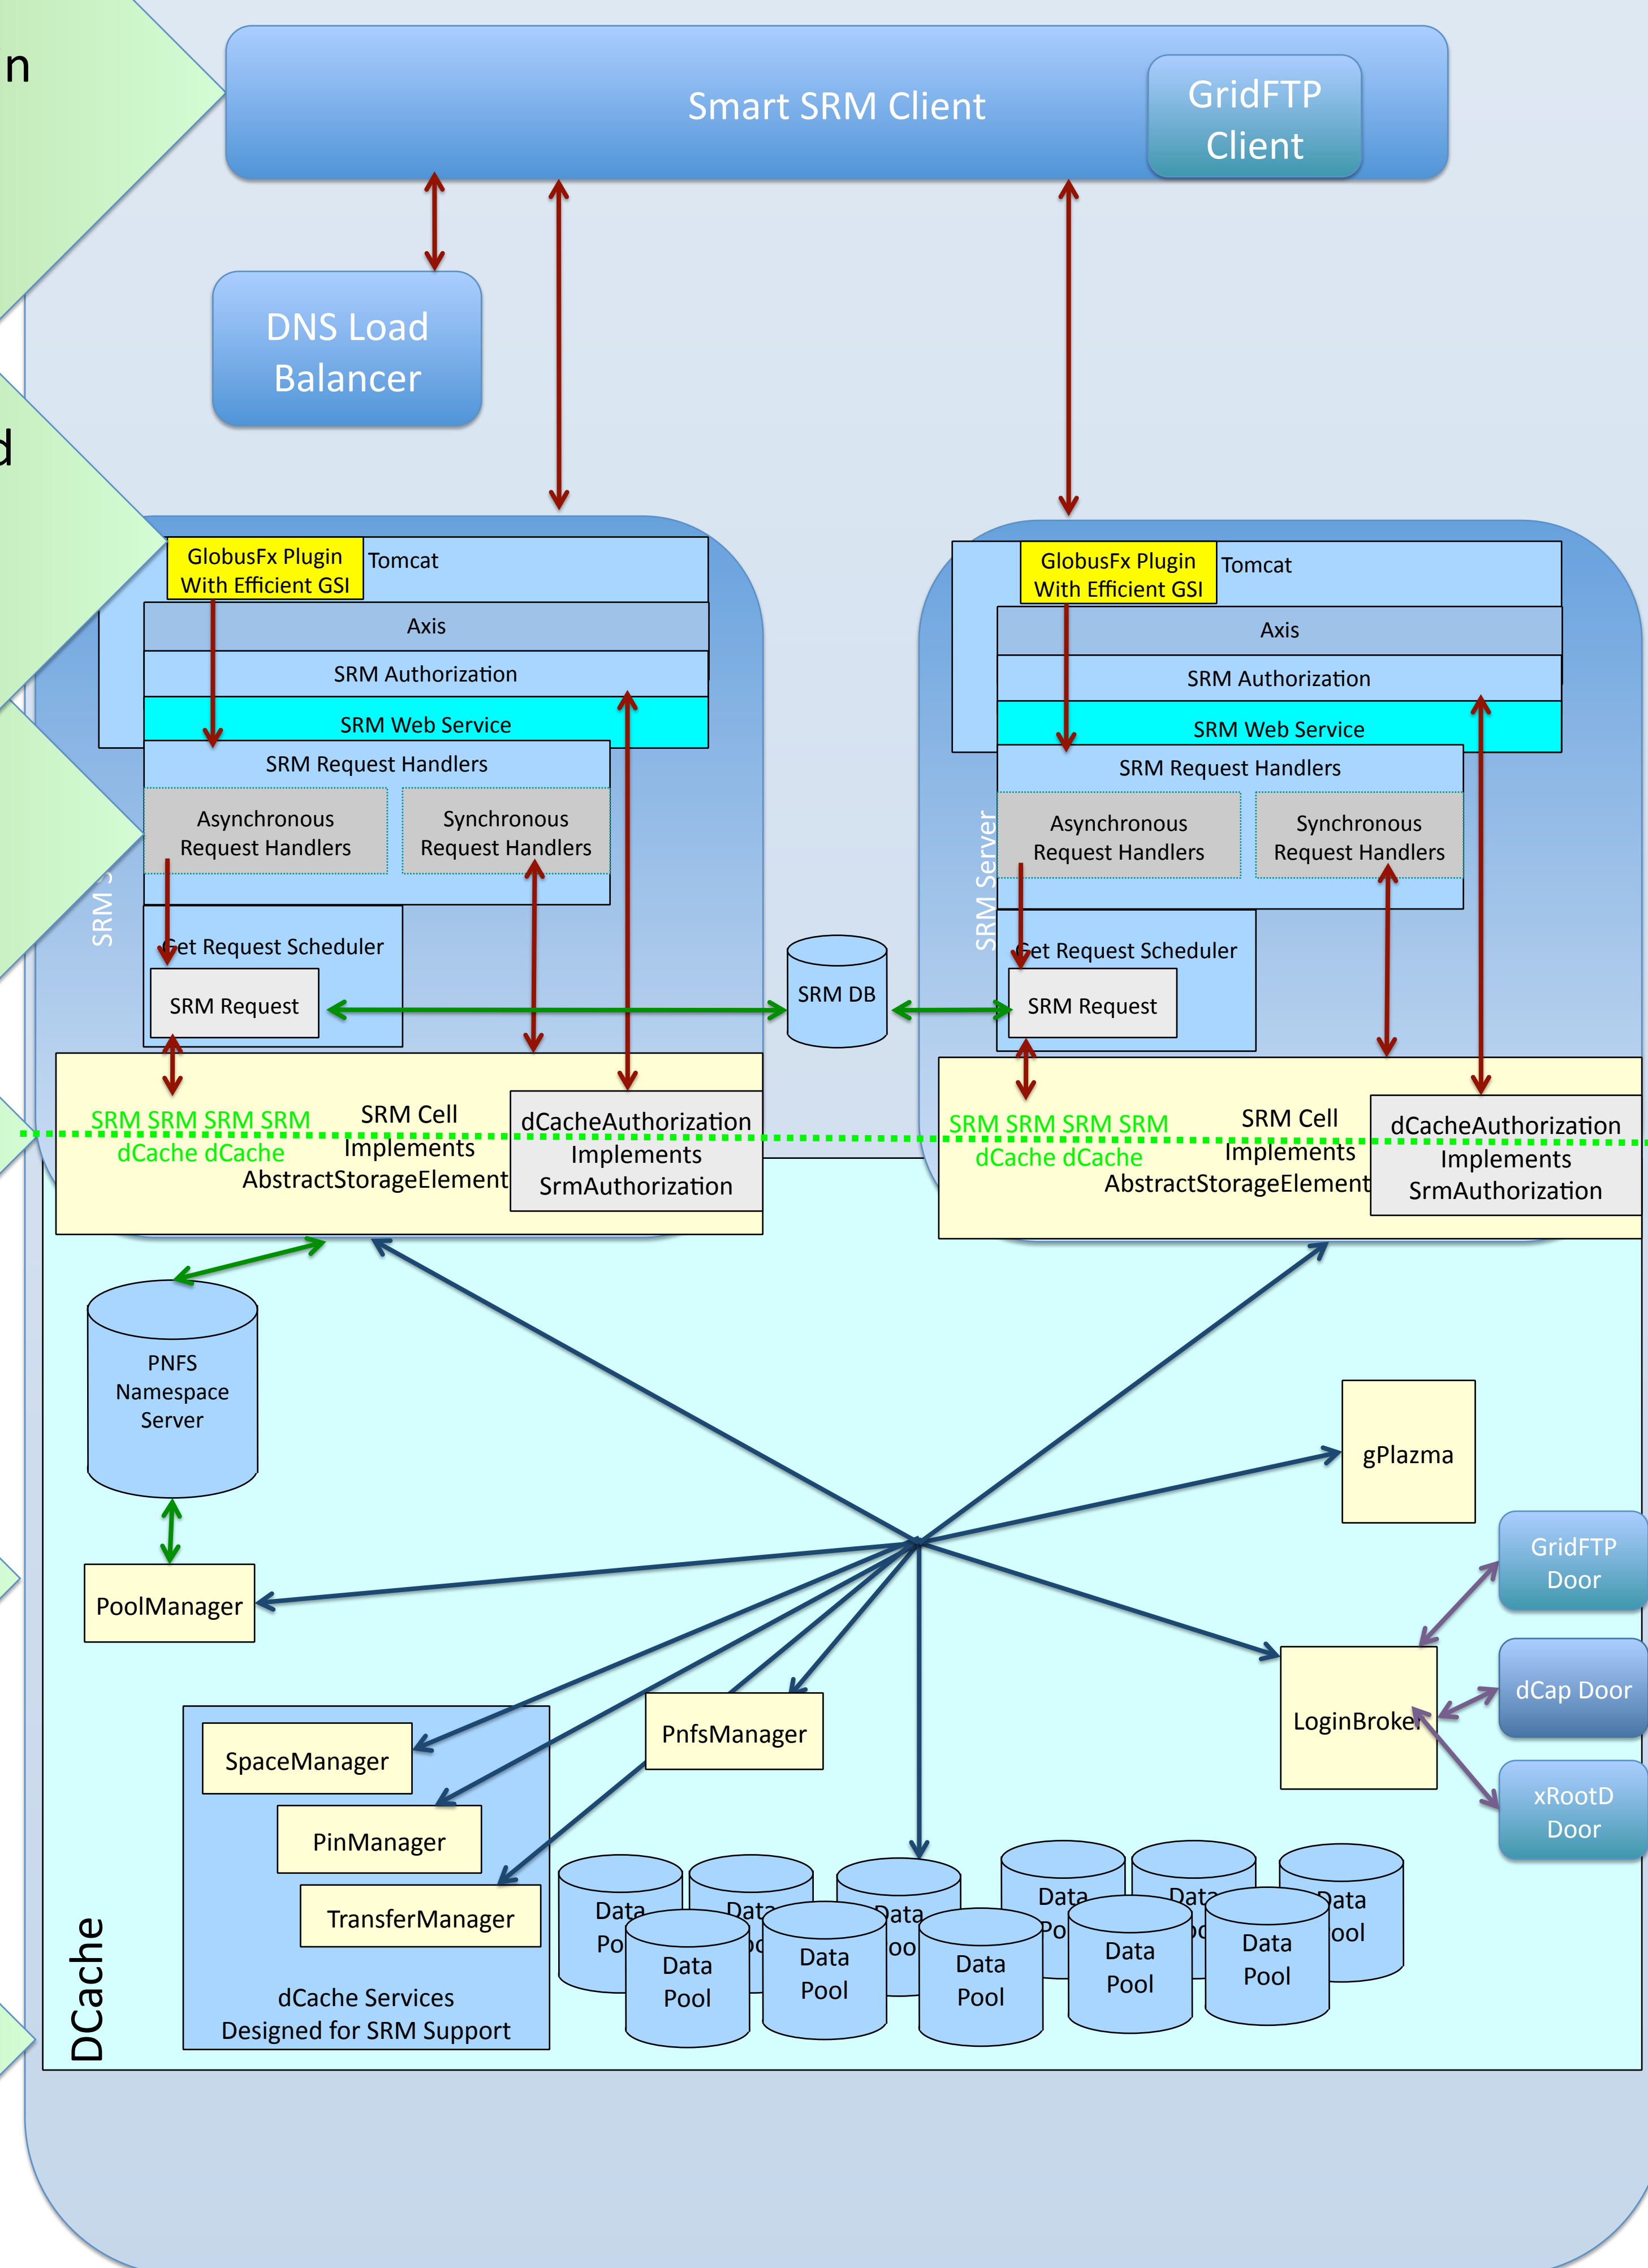
Convert SrmLs to a scheduled asynchronous operation

Use lifetime attributes in internal dCache messaging, aggressively remove expiring requests

Use DNS load balancer and a database as a centralized storage of requests data to replicate SRM Server functionality

Experts from BNL, NDGF and FNAL are reviewing design, recommendations by the end of March 2009

## Scalable dCache SRM Server



Computing in High Energy and Nuclear Physics,  
21 - 27 March 2009, Prague, Czech Republic

Work supported by the U.S. Department of Energy  
under contract No. DE-AC02-07CH11359