



GANESHA, a multi-usage with large cache NFSv4 server

Philippe Deniel
Thomas Leibovici
Jacques-Charles Lafoucrière

GANESHA and NFSv4



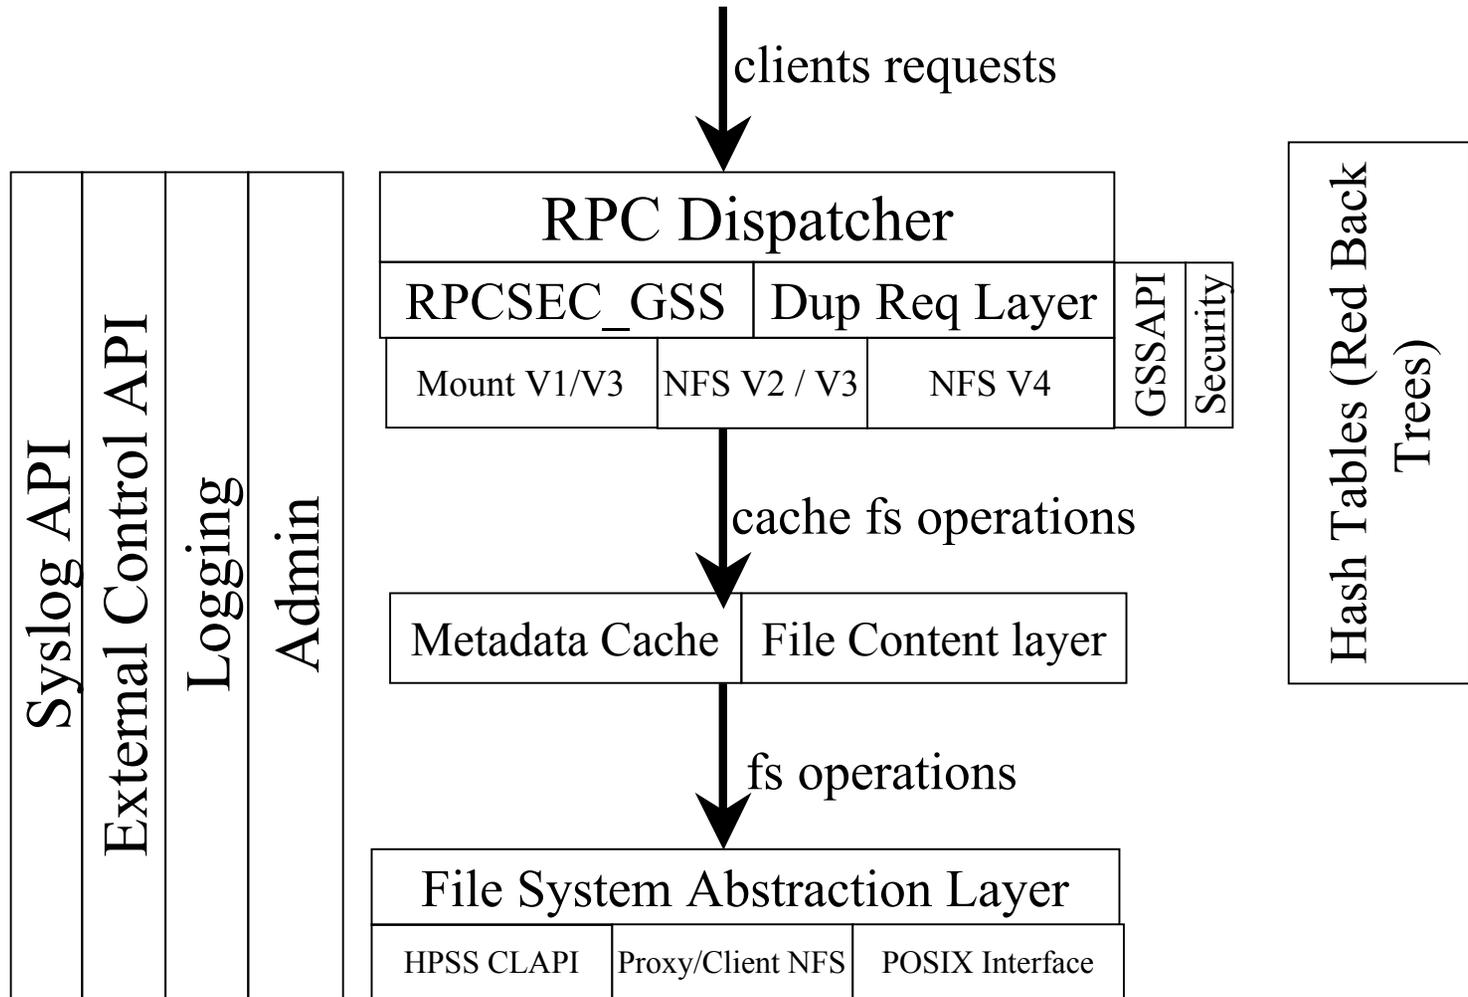
- NFSv4 is a complex and generic protocol. It can be used to access many different file systems
- GANESHA is a NFSv4 Server that focuses on the following points:
 - Managing huge metadata and data caches (up to millions of entries)
 - Being able to provide access to different sets of data

Why run in User Space ?



- Less constraints than in Kernel Space
- It's easier to manage huge pieces of memory (several Gigabytes) for building big caches
- Many mechanisms have native access in User Space (no need to use “rpc_pipefs”)
 - Kerberos (used for RPCSEC_GSS), LDAP, SNMP
- A code designed to run in User Space is more portable.
 - It now runs on Linux
 - Port to AIX and Solaris is a work in progress

Layered architecture



The File System Abstraction Layer modules



- We currently have two available FSALs
 - FSAL/HPSS: The HPSS Name Space via the HPSS CLAPI (HPSS is a HSM developed by the DOE and IBM Gov)
 - FSAL/POSIX: Any File System through « POSIX calls » (a Database is used to « reverse lookup » paths from file handles)
- The following FSALs are under development
 - FSAL/Proxy Client: with a FSAL which is a NFSv4 client, GANESHA will become a NFSv4 proxy
 - FSAL/LDAP: LDAP information is organized as trees. They can be browsed like the /proc fs. This access will be provided by this FSAL
 - FSAL/SNMP: this module will make it possible to browse MIBs information like /proc (SNMP OID are nice to build file handles with them).
- A FSAL will be developed on top of LUSTRE API to provide access to LUSTRE file system for computers with no LUSTRE client but with a NFS client.