

# NFS-RDMA Status



OPENFABRICS  
ALLIANCE

**James Lentini**  
**Network Appliance**

[jlentini@netapp.com](mailto:jlentini@netapp.com)

11/17/2006

# NFS-RDMA



## ➤ **NFS-RDMA** is

- An RPC-layer protocol that allows NFS to use RDMA networks (such as Infiniband and iWARP)
- A transparent solution for applications, NFS protocol features, and NFS users
- A significant performance boost to clients
  - Reduces client CPU overhead
  - Utilizes high-bandwidth, low-latency fabrics
- A single-wire host cluster solution

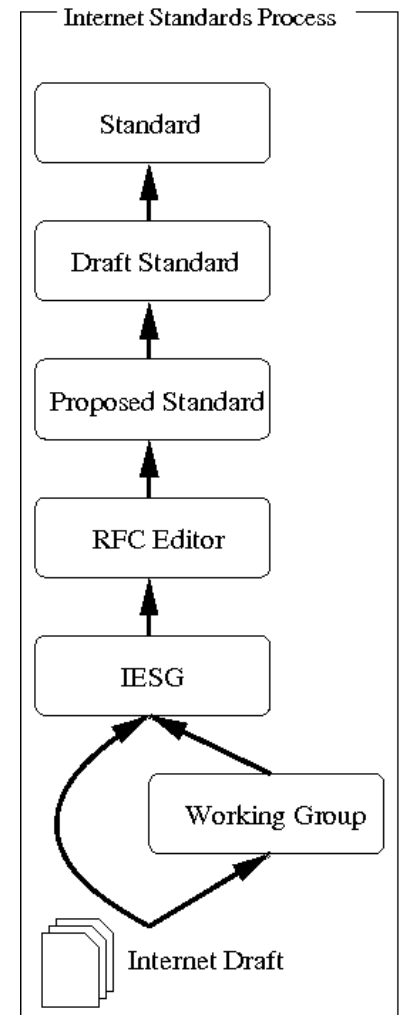
# Protocol Standardization



Three documents in the IETF NFSv4 Working Group:  
<http://www.ietf.org/html.charters/nfsv4-charter.html>

- NFS RDMA Problem Statement
- RDMA Transport for ONC RPC
  - Describes the RPC-RDMA protocol for sending RPC messages on an RDMA transport (IB, iWARP)
- NFS Direct Data Placement
  - Describes the NFSv2/v3/v4 mapping to RPC-RDMA operations

All documents are in working group last call.



# Open Source Implementations



- Linux client and server

<http://www.sourceforge.net/projects/nfs-rdma>

- OpenSolaris client and server

<http://www.opensolaris.org/os/project/nfsrdma>

- Excellent performance results obtained

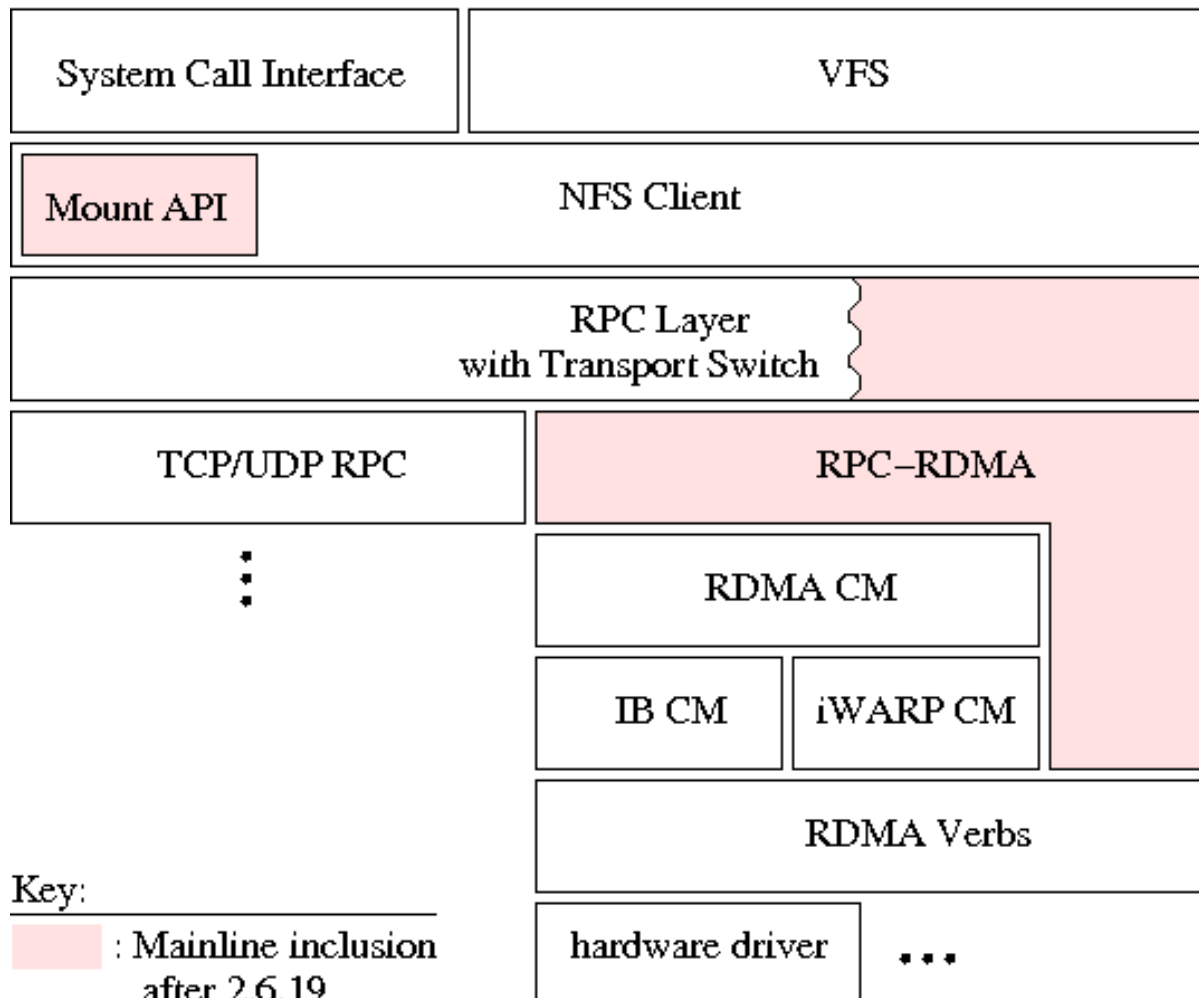
- 900+MB/s direct i/o at <10% client CPU

# Linux Client and Server Status

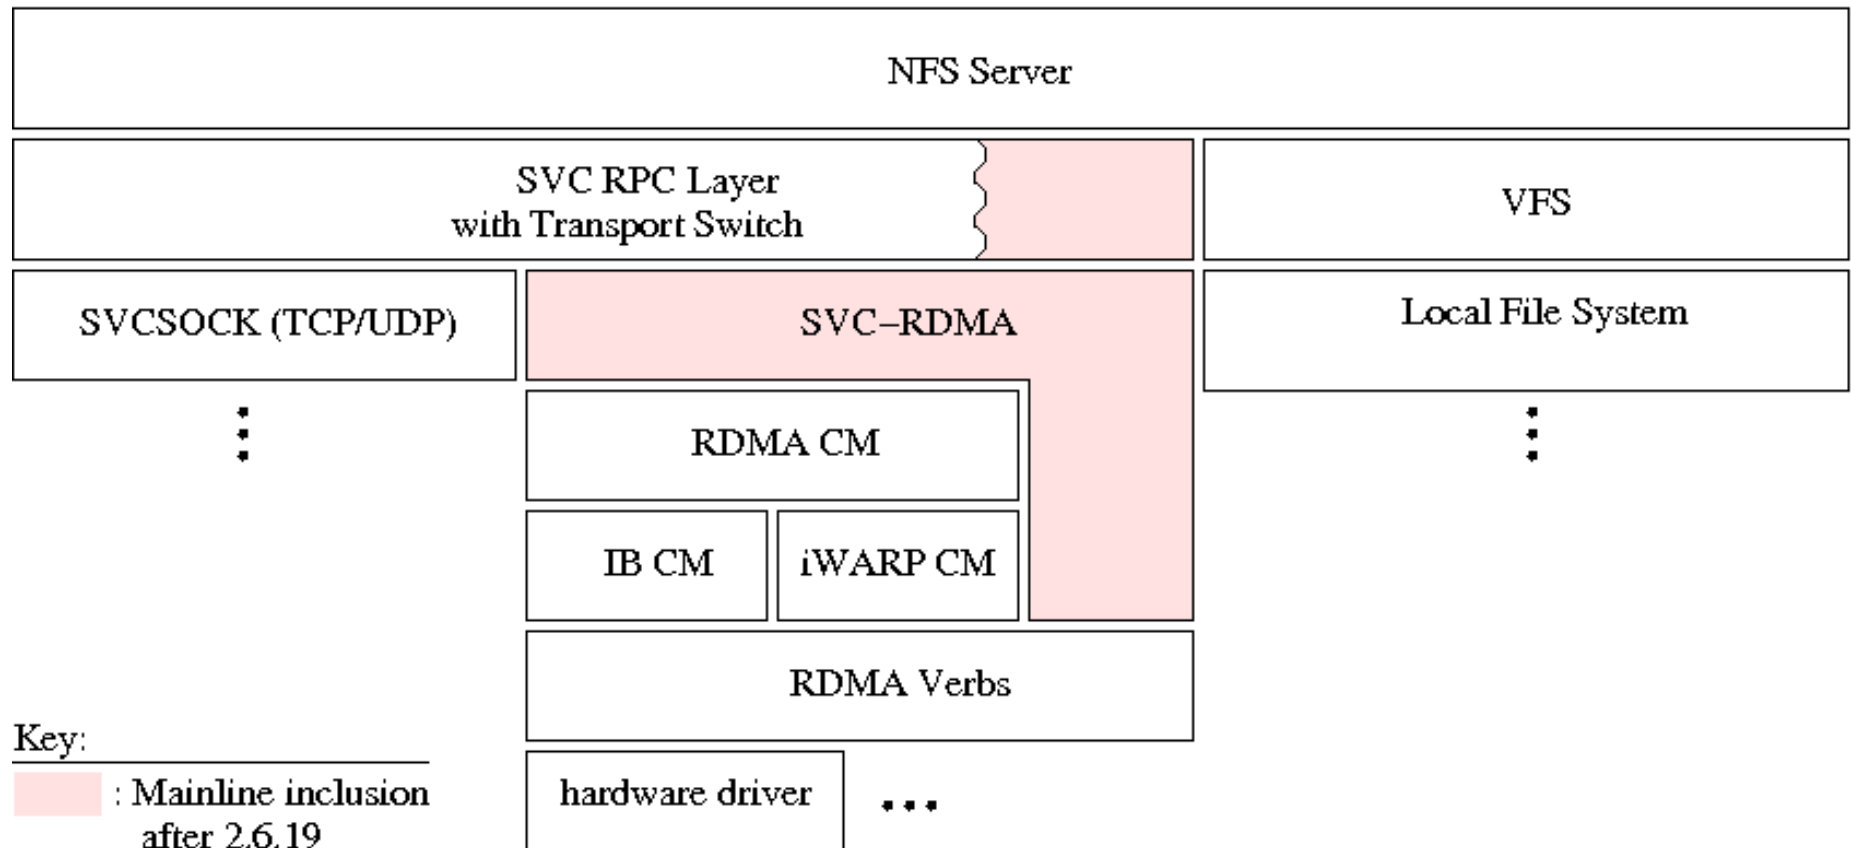


- Stable code (passes NFS Connectathon tests, iozone, etc)
- Available under a dual BSD/GPL license
- Full implementation supporting NFSv2, v3, v4
- Data transfer via send/recv, RDMA read, and RDMA write

# Linux Client Architecture



# Linux Server Architecture



# Future Plans



- Submission of Linux client and server to kernel.org
  - NFS-RDMA depends on partially merged RPC Transport Switch
- OpenSolaris project has approved inclusion of NFS-RDMA