

Current PVFS2 status

PVFS2 Development Team

Last Updated: September 2003

Contents

| | | |
|----------|---|----------|
| 1 | Introduction | 2 |
| 2 | Known limitations and missing features | 2 |
| 3 | Experimental features | 2 |
| 4 | Placeholder / deprecated code | 2 |
| 5 | Open issues | 2 |
| 6 | Good examples | 3 |

1 Introduction

This document describes the current status of PVFS2 development. This document only includes issues related to functionality or correctness. No performance optimizations are listed for now.

2 Known limitations and missing features

This section lists file system limitations for which we have a known solution or plan.

- efficient conversion of MPI datatypes to PVFS2 datatypes in ROMIO
- hooks for tuning consistency semantics
- hooks for controlling distribution and distribution parameters
- standardizing error code format
- integration of user level buffer cache
- eliminating memory leaks
- consistent error handling in client and server state machines
- simple failover plan

3 Experimental features

These are features that are implemented but have not been thoroughly tested.

- GM network support
- IB network support

4 Placeholder / deprecated code

These parts of the code have a working implementation, but we intend to replace them as time permits.

- “contig” request encoder implementation
- pvfs2-client implementation

5 Open issues

The items on this list are known problems that have not been resolved.

- access control / security
- how to manage client side configuration (fstab information)
- support for 2.4 series kernels

- how to add file systems to an existing system interface run time instance (proper vfs bootstrapping)
- managing server configuration files
- nonblocking I/O at system interface
- how to handle I/O failures, unposting, etc.
- TCP module scalability
- extended attributes
- redundancy

6 Good examples

This section points out specific areas of the code that demonstrate best practice for PVFS2 development.

- ?