

Trove and PVFS2

PVFS Development Team

April 15, 2020

1 Motivation and Goals

The purpose of this document is to describe the use of Trove in PVFS2.

PVFS2 deals with four basic types of objects:

- directories
- metafiles
- datafiles
- symlinks

We will discuss how Trove is used to store these objects in the upcoming sections. Additionally we will cover how PVFS2 bootstraps from the Trove perspective; that is, how it obtains a “root handle” and so on.

2 Current Implementation (03/21/2003)

This section describes the implementation as of the date above. Eventually, when the implementation catches up with the long term plan, this entire section will probably disappear.

2.1 PVFS2 Objects

At this time the type of an object is stored in at least one place, the dataspace attributes. These can be retrieved using `job.trove_dspace_getattr`. Additionally, as we will see, directories and metafiles store metadata as keyvals.

Directories are of type `PVFS_TYPE_DIRECTORY`. They are actually stored as two dataspaces in the current implementation. The first dataspace is used solely to store the attributes of the directory (under the key `metadata`, as a `PVFS_object_attr` structure) and, if entries have been created, the handle of a second dataspace where the directory entries are stored (under the key `dir_ent`, as a `PVFS_handle` type). The second dataspace is marked as type `PVFS_TYPE_DIRDATA` to differentiate it. This one holds the directory entries, with keys being short names of

