

Frequently Asked Questions about PVFS2

PVFS2 Development Team

April 17, 2006

Contents

1	Basics	1
1.1	What is PVFS2?	1
1.2	What does the “V” in PVFS2 stand for?	1
1.3	Is PVFS2 an attempt to parallelize the *NIX VFS?	1
1.4	What are the components of PVFS2 that I should know about?	1
2	Supported Architectures and Hardware	1
2.1	Does PVFS2 require any particular hardware?	2
2.2	What architectures does PVFS2 support?	2
2.3	Does PVFS2 work across heterogeneous architectures?	2
2.4	Does running PVFS2 require a particular kernel or kernel version?	2
2.5	What specific hardware architectures are supported by the PVFS2 kernel module?	2
2.6	Does the PVFS2 client require a patched Linux kernel?	2
2.7	Can I build the PVFS2 kernel code directly into the kernel, rather than as a module?	2
2.8	Is there a MacOS X/Cygwin/Windows client for PVFS2?	3
3	Installation	3
3.1	How do I install PVFS2?	3
3.2	How can I store PVFS2 data on multiple disks on a single node?	3
3.3	How can I run multiple PVFS2 servers on the same node?	3
3.4	Can I use multiple metadata servers in PVFS2?	3
3.5	Does using multiple metadata servers reduce the chance of file system corruption during hardware failures?	4
3.6	How many servers should I run?	4
3.7	Can PVFS2 servers listen on two network interfaces simultaneously (i.e. multihome)?	4
3.8	How can I automount PVFS2 volumes?	4
3.9	Can I mount more than one PVFS2 file system on the same client?	4
3.10	How can I upgrade from PVFS1 to PVFS2?	5
4	Reporting Problems	5
4.1	Where can I find documentation?	5
4.2	What should I do if I have a problem?	5
4.3	How do I report a problem with PVFS2?	5

5	Problems and Solutions	6
5.1	When I try to mount, I get 'wrong fs type, bad option, bad superblock...'	6
5.2	PVFS2 server consumes 100% of the CPU	6
5.3	PVFS2 write performance slows down dramatically	7
5.4	I get "error while loading shared libraries" when starting PVFS2 programs	7
5.5	Make kmod24 fails with "structure has no member..." errors	7
5.6	When i try to mount a pvfs2 file system, something goes wrong.	8
5.7	I did all three of the above steps and I still can't mount pvfs2	8
5.8	I'm running Redhat and the pvfs2-server can't be killed! What's wrong?	8
5.9	Why do you single out Redhat users? What's so different about Redhat than other distributions?	8
5.10	Where is the kernel source on a Fedora system?	9
5.11	What are extended attributes? How do I use them with PVFS2?	9
5.12	What are Access Control Lists? How do I enable Access Control Lists on PVFS2?	9
5.13	On SLES 9, 'make kmod' complains about mmgrab and flush_icache_range being undefined	10
5.14	Can we run the Apache webserver to serve files off a PVFS2 volume?	10
6	Performance	10
6.1	I ran Bonnie and/or IOzone and the performance is terrible. Why? Is there anything I can do?	10
6.2	Why is program XXX so slow?	11
6.3	NFS outperforms PVFS2 for application XXX. Why?	11
6.4	Can the underlying local file system affect PVFS2 performance?	11
6.5	My app still runs more slowly than I would like. What can I do?	12
7	Fault Tolerance	12
7.1	Does PVFS2 support some form of fault tolerance?	12
7.2	Can PVFS2 tolerate client failures?	12
7.3	Can PVFS2 tolerate disk failures?	12
7.4	Can PVFS2 tolerate network failures?	13
7.5	Can PVFS2 tolerate server failures?	13
8	File System Interfaces	13
8.1	How do I get MPI-IO for PVFS2?	13
8.2	Can I directly manipulate PVFS2 files on the PVFS2 servers without going through some client interface?	13
9	Management	13
9.1	How can I back up my PVFS2 file system?	13
9.2	Can I add, remove, or change the order of the PVFS2 servers on an existing PVFS2 file system?	14
9.3	Are there tools for migrating data between servers?	14
9.4	Why does df show less free space than I think it should? What can I do about that?	15
9.5	Does PVFS2 have a maximum file system size? If so, what is it?	15