

POHMELFS high performance network filesystem release.

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-04/msg10508.html>

- *From:* Evgeniy Polyakov <johnpol@xxxxxxxxxxxx>
 - *Date:* Fri, 25 Apr 2008 19:07:56 +0400
-

Hi.

I'm please to announce POHMEL high performance network filesystem.
POHMELFS stands for Parallel Optimized Host Message Exchange Layered File System.

Development status can be tracked in filesystem section [1].

This is a high performance network filesystem with local coherent cache of data and metadata. Its main goal is distributed parallel processing of data. Network filesystem is a client transport. POHMELFS protocol was proven to be superior to NFS in lots (if not all, then it is in a roadmap) operations.

Basic POHMELFS features:

- * Local coherent (notes [5]) cache for data and metadata.
- * Completely async processing of all events (hard and symlinks are the only exceptions) including object creation and data reading.
- * Flexible object architecture optimized for network processing. Ability to create long paths to object and remove arbitrary huge directoris in single network command.
- * High performance is one of the main design goals.
- * Very fast and scalable multithreaded userspace server. Being in userspace it works with any underlying filesystem and still is much faster than async ni-kernel NFS one.

Roadmap includes:

- * Server extension to allow storing data on multiple devices (like creating mirroring), first by saving data in several local directories (think about server, which mounted remote dirs over POHMELFS or NFS, and local dirs).
- * Client/server extension to report lookup and readdir requests not only for local destination, but also to different addresses, so that reading/writing could be done from different nodes in parallel.
- * Strong authentication and possible data encryption in network channel.
- * Extend client to be able to switch between different servers (if one goes down, client automatically reconnects to second and so on).
- * Async writing of the data from receiving kernel thread into userspace pages via `copy_to_user()` (check development tracking blog for results).

POHMELFS high performance network filesystem release.

One can grab sources from archive [2] or check homepage [3].
Benchmark section can be found in th blog [4].

It is work-in-progress, and network protocol is not stable yet.

Thank you.

1. POHMELFS development status.

<http://tservice.net.ru/~s0mbre/blog/devel/fs/index.html>

2. Source archive.

<http://tservice.net.ru/~s0mbre/archive/pohmelfs/>

3. POHMELFS homepage.