
Subject: Re: [Ecryptfs-devel] [PATCH 1/2] eCryptfs: Introduce device handle for userspace daemon communication

Posted by [Michael Halcrow](#) on Tue, 15 Apr 2008 22:47:50 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Tue, Apr 15, 2008 at 02:04:53PM -0700, Andrew Morton wrote:

> On Tue, 15 Apr 2008 15:23:13 -0500

> Michael Halcrow <mhalcrow@us.ibm.com> wrote:

>

> > Functions to facilitate reading and writing to the eCryptfs
> > miscellaneous device handle. This will replace the netlink interface
> > as the preferred mechanism for communicating with the userspace
> > eCryptfs daemon.

> >

> > Each user has his own daemon, which registers itself by opening the
> > eCryptfs device handle. Only one daemon per euid may be registered at
> > any given time. The eCryptfs module sends a message to a daemon by
> > adding its message to the daemon's outgoing message queue. The daemon
> > reads the device handle to get the oldest message off the queue.

> >

> > Incoming messages from the userspace daemon are immediately
> > handled. If the message is a response, then the corresponding process
> > that is blocked waiting for the response is awakened.

>

> This is a drastic change, but the changelog doesn't tell us why it
> is being made!

This resurrects the question from 2006:

On Thu, Aug 24, 2006 at 08:54:19PM -0700, Andrew Morton wrote:

> - _why_ does it use netlink?

I responded:

> Netlink provides the transport mechanism that would minimize the
> complexity of the implementation, given that we can have multiple
> daemons (one per user).

A regular device file was my real preference from the get-go, but I went with netlink at the time because I thought it would be less complex for managing send queues (i.e., just do a unicast and move on). It turns out that we do not really get that much complexity reduction with netlink, and netlink is more heavyweight than a device handle.

In addition, the netlink interface to eCryptfs has been broken since 2.6.24. I am assuming this is a bug in how eCryptfs uses netlink, since the other in-kernel users of netlink do not seem to be having

any problems. I have had one report of a user successfully using eCryptfs with netlink on 2.6.24, but for my own systems, when starting the userspace daemon, the initial hello message sent to the eCryptfs kernel module results in an oops right off the bat. I spent some time looking at it, but I have not yet found the cause. The netlink interface breaking gave me the motivation to just finish my patch to migrate to a regular device handle. If I cannot find out soon why the netlink interface in eCryptfs broke, I am likely to just send a patch to disable it in 2.6.24 and 2.6.25. I would like the device handle to be the preferred means of communicating with the userspace daemon from 2.6.26 on forward.

Mike

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
