Subject: Re: ploop and trim/ discard support Posted by Kirill Korotaev on Mon, 17 Sep 2012 16:46:51 GMT

View Forum Message <> Reply to Message

On Sep 17, 2012, at 19:53, Corin Langosch < corin.langosch@netskin.com> wrote:

>

> On 13.09.2012 at 09:22 +0200, Kirill Korotaev <dev@parallels.com> wrote:

>>

- >> No, AFAIR we should use TRIM on ext4 and it simply reports unused space.
- >> Balloon is used for "resize" via allocating some space and hiding it
- >> from user, but for compacting it's a bit bad since can cause ENOSPC
- >> while it's really not...

>>

- > So this whole ballooning is only a work around as trim/ discard support
- > for ext4 is only available in kernel >= 2.6.33? Once openvz is rebased
- > to a newer kernel (3.2.x?) it can/ will be dropped? :)

>

You've mixed 2 different scenarios:

1. vzctl set --diskpace

When you resize CT to smaller sizes we do not want to resize live file system and move data around causing I/O.

So we use balloon to reserve some space in CT and "pretend" that CT was made smaller.

TRIM has nothing to do with this scenario, cause it wouldn't prevent file system from allocating its free space.

2. compacting

When CT has used some space and then files were removed image requires "compaction" to free this space back to host.

This is where both ballooning and TRIMing can help. But ballooning reserves disk space, so it can lead to ENOSPC inside CT and thus is better to avoid (remember, it reserves space!).

TRIM on the other hand is a standard way to cause file system to report it's unused space and this is what we use.

TRIM support present in our RHEL6 kernels, so switching to >= 2.6.33 is not required and won't result in any benefits in this area.

Thanks, Kirill

File Attachments

1) smime.p7s, downloaded 1635 times