Subject: Re: Stale mtab entries following stop Posted by laotse on Thu, 27 Aug 2009 12:29:38 GMT

View Forum Message <> Reply to Message

maratrus wrote on Wed, 19 August 2009 15:39Most of VEs contain /etc/mtab as a symlink to /proc/mounts. You can manually do it for this particular VE.

The mtab of the VE is sane and indeed is the link you suggested. The issue happens with the mtab of the node!

Quote: A VE's mount script which is supposed to mount a partition put an appropriate information to /etc/mtab.

The mount script does not fiddle with the mtab of the VE for mtab is a link to /proc/mounts. The mount shown in the original post is listed as:

/dev/mapper/r5-mail /srv/mail xfs rw,attr2,nobarrier,noquota 0 0

Quote: When a VE is going to be stopped, one of the init-scripts have to umount all mounted filesystems gracefully. Perhaps it umounted all filesystems indeed but somehow missed to clear /etc/mtab.

Since the init script runs inside the VE, it cannot access the mtab of the node. Neither could an init script during startup access the block device for mounting. Mounting and unmounting has to be performed by the node.

But, if there is something like the VE trying to unmount the mount performed by the node, then we probably have some start for investigating what is actually happening.

Quote: or you can mount a partition with a "-n" option.

This is probably a work-around for this issue to avoid flooding mtab. But I'd really like to know how mount/unmount all running in the node's context can perform such a misbehaviour in the first place.

Thanks for the help,

- lars.