

Devpts namespace patchset

In continuation of the implementation of containers in mainline, we need to support multiple PTY namespaces so that the PTY index (ie the tty names) in one container is independent of the PTY indices of other containers. For instance this would allow each container to have a '/dev/pts/0' PTY and refer to different terminals.

[PATCH 1/7]: Propagate error code from devpts_pty_new
[PATCH 2/7]: Factor out PTY index allocation
[PATCH 3/7]: Enable multiple mounts of /dev/pts
[PATCH 4/7]: Allow mknod of ptmx and tty in devpts
[PATCH 5/7]: Implement get_pts_ns() and put_pts_ns()
[PATCH 6/7]: Determine pts_ns from a pty's inode
[PATCH 7/7]: Enable cloning PTY namespaces

Todo:

- This patchset depends on availability of additional clone flags. and relies on Cedric's clone64 patchset. See

<http://marc.info/?l=linux-kernel&m=120272411925609&w=2>

- Needs some cleanup and more testing
- Ensure patchset is bisect-safe

Changelogs from earlier posts to Containers@.

Changelog[v2]:

(Patches 4 and 6 differ significantly from [v1]. Others are mostly the same)

- [Alexey Dobriyan, Pavel Emelyanov] Removed the hack to check for user-space mount.
- [Serge Hallyn] Added rcu locking around access to sb->s_fs_info.
- [Serge Hallyn] Allow creation of /dev/pts/ptmx and /dev/pts/tty devices to simplify the process of finding the 'owning' pts-ns of the device (specially when accessed from parent-pts-ns) See patches 4 and 6 for details.

Changelog[v1]:

- Fixed circular reference by not caching the pts_ns in sb->s_fs_info (without incrementing reference count) and clearing the sb->s_fs_info when destroying the pts_ns
- To allow access to a child container's ptys from parent container, determine the 'pts_ns' of a 'pty' from its inode.
- Added a check (hack) to ensure user-space mount of /dev/pts is done before creating PTYs in a new pts-ns.
- Reorganized the patchset and removed redundant changes.
- Ported to work with Cedric Le Goater's clone64() system call now that we are out of clone_flags.

Changelog[v0]:

This patchset is based on earlier versions developed by Serge Hallyn and Matt Helsley.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
