Subject: [PATCH 0/9] Core pid namespace enhancements
Posted by ebiederm on Wed, 12 Dec 2007 12:38:11 GMT

The following patchset updates the pid namespace infrastructure
so we don't constantly have to worry if we have been called
before or after exit_task_namespaces, by using the pid_namespace
obtained from a processes pid, handles the general case of setting
si_pid in struct sig_info, changes where we drop signals sent to init,
and enhances that changes to also work with the per namespace init.

Thus resolving most of the big gotchas with the current pid namespace
implementation.

Eric

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

Subject: Re: [PATCH 0/9] Core pid namespace enhancements
Posted by Sukadev Bhattiprolu on Tue, 18 Dec 2007 00:52:41 GMT

Eric W. Biederman [ebiederm@xmission.com] wrote:
|
| The following patchset updates the pid namespace infrastructure
| so we don't constantly have to worry if we have been called
| before or after exit_task_namespaces, by using the pid_namespace
| obtained from a processes pid, handles the general case of setting
| si_pid in struct sig_info, changes where we drop signals sent to init,
| and enhances that changes to also work with the per namespace init.
|
| Thus resolving most of the big gotchas with the current pid namespace
| implementation.
|
| Eric

The patchset looks good to me. My only testcase from the previous set
that broke was the case that Oleg pointed out

 - container-init sets a handler for SIGUSR1
 - container-init blocks SIGUSR1
 - a descendant of container-init posts SIGUSR1 to container-init
 - container-init sets SIGUSR1 to SIG_DFL and unblocks and takes the
   fatal signal.

While that discussion can continue...

<Acked-by>: Sukadev Bhattiprolu <sukadev@us.ibm.com>

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