Subject: [PATCH] Allow signalling container-init
Posted by Sukadev Bhattiprolu on Wed, 08 Aug 2007 23:47:37 GMT
View Forum Message <> Reply to Message

Pavel,

Should we include this in the patchset ?

Sukadev

From: Sukadev Bhattiprolu <sukadev@us.ibm.com>
Subject: [PATCH] Allow signalling container-init

Only the global-init process must be special - any other container-init
process must be killable to prevent run-away processes in the system.

TODO: Ideally we should allow killing the container-init only from
ancestor containers and prevent it being killed from that or
descendant containers. But that is a more complex change and
will be addressed by a follow-on patch. For now allow the
container-init to be terminated by any process with sufficient
privileges.

Signed-off-by: Sukadev Bhattiprolu <sukadev@us.ibm.com>
---
kernel/signal.c | 6 ++----
1 file changed, 2 insertions(+), 4 deletions(-)

Index: lx26-23-rc1-mm1/kernel/signal.c
==========
--- lx26-23-rc1-mm1.orig/kernel/signal.c 2007-08-07 13:52:12.000000000 -0700
+++ lx26-23-rc1-mm1/kernel/signal.c 2007-08-08 15:09:27.000000000 -0700
@@ -1861,11 +1861,9 @@ relock:
continue;

/*
- * Init of a pid space gets no signals it doesn't want from
- * within that pid space. It can of course get signals from
- * its parent pid space.
+ * Global init gets no signals it doesn't want.
*/
- if (current == task_child_reaper(current))
+ if (is_global_init(current->group_leader))
continue;

if (sig_kernel_stop(signr)) {

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