Subject: [RFC] [PATCH 0/2] namespace enter through hijack
Posted by serue on Wed, 29 Aug 2007 20:04:20 GMT
View Forum Message <> Reply to Message

This patchset implements namespace entering by forcing a task in the target namespace to clone itself. This has some advantages over just replacing a random task's namespace pointers to the target ones. In particular

1. If switching pid namespaces, the stack of upids is automatically correctly generated.
2. Security context is inherited from the target task. Assuming a security module which labels data based on the task security context, like selinux, this may prevent severe mislabeling of container data by an inadvertant host system administrator action. Whether that works or not will still depend on the policy and the task cloned.

While this version takes a pid of a process to clone (for convenience of prototyping) we may prefer to use a ns_container name and choose one of it's tasks, to prevent pid wraparound.

Tested and 'works for me', but at the moment I'm just sending this out for discussion.

Alternatives to this include bind_ns()+switch_ns() by Cedric and the ns_container namespace entering enhancements I've previously sent.

thanks,
-serge

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