
Subject: [RFC][PATCH 0/7] Resource controllers based on process containers
Posted by [xemul](#) on Tue, 06 Mar 2007 14:42:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

This patchset adds RSS, accounting and control and limiting the number of tasks and files within container.

Based on top of Paul Menage's container subsystem v7

RSS controller includes per-container RSS accouter, reclamation and OOM killer. It behaves like standalone machine - when container runs out of resources it tries to reclaim some pages and if it doesn't succeed in it kills some task which mm_struct belongs to container in question.

Num tasks and files containers are very simple and self-descriptive from code.

As discussed before when a task moves from one container to another no resources follow it - they keep holding the container they were allocated in.

The difficulties met during using of Pauls' containers were:

1. Container fork hook is placed before new task changes. This makes impossible of handling fork properly. I.e. new mm_struct should have pointer to RSS container, but we don't have one at that early time.
2. Extended containers may register themselves too late. Kernel threads/helpers start forking, opening files and touching pages much earlier. This patchset workarounds this in not-so-cute manner and I'm waiting for Paul's comments on this issue.