
Subject: [PATCH 0/8] RSS controller based on process containers (v2)

Posted by [xemul](#) on Mon, 09 Apr 2007 12:18:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

Adds RSS accounting and control within a container.

Major change: current scanner code reuse.

Tasks and files accounting is not included as these containers are simple enough to be implemented later.

Based on top of Paul Menage's container subsystem v8.

Note, that only first three patches from Paul's patchset should be used as BeanCounters, CPU accounting etc are not needed for RSS container.

RSS container includes the per-container RSS accounting and reclamation, and out-of-memory killer. The container behaves like a standalone machine - when it runs out of resources, it tries to reclaim some pages, and if it doesn't succeed, kills some task which `mm_struct` belongs to the container in question.

Changes from previous version:

- * reclamation code reuse. I.e. the scanner code used to reclaim pages on global memory shortage is used as-is in per-container page reclamation;
 - * active/inactive per-container lists (a bit reworked patch from Balbir);
 - * more elegant top container creation (thanks Paul for clarifying).
-