
Subject: Re: [RFC][PATCH] UBC: user resource beancounters
Posted by [Rohit Seth](#) on Wed, 16 Aug 2006 18:53:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Wed, 2006-08-16 at 19:24 +0400, Kirill Korotaev wrote:

- > The following patch set presents base of
- > User Resource Beancounters (UBC).
- > UBC allows to account and control consumption
- > of kernel resources used by group of processes.
- >
- > The full UBC patch set allows to control:
- > - kernel memory. All the kernel objects allocatable
- > on user demand should be accounted and limited
- > for DoS protection.
- > E.g. page tables, task structs, vmas etc.
- >

Good.

- > - virtual memory pages. UBC allows to
- > limit a container to some amount of memory and
- > introduces 2-level OOM killer taking into account
- > container's consumption.
- > pages shared between containers are correctly
- > charged as fractions (tunable).
- >

I wouldn't be too worried about doing fractions. Make it unfair and charge it to either the container who first instantiated the file or the container who faulted on that page first.

Though the part that seems important is to be able to define a directory in fs and say all pages belonging to files underneath that directory are going to be put in specific container. Just like you are having resource beans associated with sockets, have address_space or inode also associated with resource beans. (And it should be possible to have a container/resource bean without any active process but set of address_space mappings with its own limits and current usage).

-rohit
