On Fri, 2006-08-18 at 12:08 +0400, Andrey Savochkin wrote:
>
> A) Have separate memory management for each container,
>    with separate buddy allocator, lru lists, page replacement mechanism.
>    That implies a considerable overhead, and the main challenge there
>    is sharing of pages between these separate memory managers.

Hold on here for just a sec...

It is quite possible to do memory management aimed at one container
while that container's memory still participates in the main VM.

There is overhead here, as the LRU scanning mechanisms get less
efficient, but I'd rather pay a penalty at LRU scanning time than divide
up the VM, or coarsely start failing allocations.

-- Dave