

How about we drill down on these a bit more.

On Mon, 2007-03-12 at 02:00 +0100, Herbert Poetzl wrote:

- > - shared mappings of 'shared' files (binaries
- > and libraries) to allow for reduced memory
- > footprint when N identical guests are running

So, it sounds like this can be phrased as a requirement like:

"Guests must be able to share pages."

Can you give us an idea why this is so? On a typical vserver system, how much memory would be lost if guests were not permitted to share pages like this? How much does this decrease the density of vservers?

- > - virtual 'physical' limit should not cause
- > swap out when there are still pages left on
- > the host system (but pages of over limit guests
- > can be preferred for swapping)

Is this a really hard requirement? It seems a bit fluffy to me. An added bonus if we can do it, but certainly not the most important requirement in the bunch.

What are the consequences if this isn't done? Doesn't a loaded system eventually have all of its pages used anyway, so won't this always be a temporary situation?

This also seems potentially harmful if we aren't able to get pages *back* that we've given to a guest. Tasks can pin pages in lots of creative ways.

- > - accounting and limits have to be consistent
- > and should roughly represent the actual used
- > memory/swap (modulo optimizations, I can go
- > into detail here, if necessary)

So, consistency is important, but is precision? If we, for instance, used one of the hashing schemes, we could have some imprecise decisions made but the system would stay consistent overall.

This requirement also doesn't seem to push us in the direction of having distinct page owners, or some sharing mechanism, because both would be consistent.

- > - OOM handling on a per guest basis, i.e. some
- > out of memory condition in guest A must not
- > affect guest B

I'll agree that this one is important and well stated as-is. Any disagreement on this one?

-- Dave

Containers mailing list
Containers@lists.osdl.org
<https://lists.osdl.org/mailman/listinfo/containers>
