Subject: Re: [RFC][PATCH 2/7] RSS controller core
Posted by xemul on Mon, 12 Mar 2007 09:02:01 GMT
View Forum Message <> Reply to Message

>>> Maybe you have some ideas how we can decide on this?
>> We need to work out what the requirements are before we can
>> settle on an implementation.
>
> Linux-VServer (and probably OpenVZ):
>
> - shared mappings of 'shared' files (binaries
>   and libraries) to allow for reduced memory
>   footprint when N identical guests are running

This is done in current patches.

> - 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)

So what to do when virtual physical limit is hit?
OOM-kill current task?

> - 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)

This is true for current implementation for
booth - this patchset ang OpenVZ beancounters.

If you sum up the physpages values for all containers
you'll get the exact number of RAM pages used.

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

This is done in current patches.

Herbert, did you look at the patches before
sending this mail or do you just want to
'take part' in conversation w/o understanding
of hat is going on?

> HTC,
> Herbert

>
>> Sigh. Who is running this show?   Anyone?
>>
>> You can actually do a form of overcommittment by allowing multiple
>> containers to share one or more of the zones. Whether that is
>> sufficient or suitable I don't know. That depends on the requirements,
>> and we haven't even discussed those, let alone agreed to them.
>>
>> _____
>> Containers mailing list
>> Containers@lists.osdl.org
>> https://lists.osdl.org/mailman/listinfo/containers
>