
Subject: Re: Mainstream progress?

Posted by [Carl-Daniel Hailfinger](#) on Wed, 04 Jun 2008 13:32:20 GMT

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

On 13.03.2008 14:20, Pavel Emelyanov wrote:

> Kir Kolyshkin wrote:

>

>> Pavel,

>>

>> Please correct/enlarge me if there's anything missing/incorrect.

>>

>> Carl-Daniel Hailfinger wrote:

>>

>>> Hi,

>>>

>>> the ongoing merge of OpenVZ-based lightweight virtualization into

>>> mainline Linux seems to be progressing nicely and I have a few questions

>>> about it:

>>> - How usable will net namespaces be in 2.6.25? I'm especially interested

>>> in assigning a real hardware ethernet device to a VE.

>>>

>>>

>> I guess some network namespace stuff is NOT completed in upcoming 2.6.25

>> (a lot of stuff done, but not finished)

>>

>> For 2.6.26, it will be more-or-less working, with just a few exceptions

>> such as conntracks and IPv6 -- maybe.

>>

Do you have an overview of the current 2.6.26 state?

>>> - Will there be an OpenVZ patch or against 2.6.25?

>>>

>>>

>> As of now we do not plan to have one. Porting to each mainstream version

>> is too much work to do -- instead we concentrate on

>> (1) merging stuff to mainstream

>> (2) supporting our stable kernels

>> (3) stabilizing 2.6.24 (for now)

>>

>

> 2.6.26 will be the next kernel we will port on.

>

Yes, I saw the new tree on git.openvz.org. I hope to use a 2.6.26-based OpenVZ kernel in the near future when 2.6.26 is officially released.

>>> - Do you expect net namespace/virtualization work to be completely
>>> merged in 2.6.26?
>>>
>>>
>> As said before, we guess that network namespaces will be more or less
>> useable in 2.6.26.
>>

It would be cool if you can give some pointers on how to use the merged
part of network namespaces (with OpenVZ tools or any other tool).

>>> - Which OpenVZ features will not make it to mainline in the 2.6.26
>>> timeframe?
>>>
>>>
>> A lot. Network namespaces is just one piece of a pie.
>>
>> To paint in big strokes: resource management and live migration are
>> still on our todo list.
>>
>> Resource management: some stuff is in -mm kernel, need to add much more.
>> Will work on that after finishing with net namespaces.
>>
>
> Memory management will change with 2.6.26. API will be kept, but
> we'll have a per-ve swapout and (probably) drop privvmpages accounting.
>
>
>> Live migration: this is much more far perspective, we haven't even
>> started to discuss that seriously on containers@ list.
>>

That's unfortunate, but I can understand concentrating on a few features
to get them merged completely.

Regards,
Carl-Daniel
