

---

Subject: Re: [RFC][PATCH 2/7] RSS controller core  
Posted by [akpm](#) on Tue, 13 Mar 2007 06:04:39 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> On Mon, 12 Mar 2007 23:41:29 +0100 Herbert Poetzl <herbert@13thfloor.at> wrote:  
> On Mon, Mar 12, 2007 at 11:42:59AM -0700, Dave Hansen wrote:  
> > 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?  
>  
> sure, one reason for this is that guests tend to  
> be similar (or almost identical) which results  
> in quite a lot of 'shared' libraries and executables  
> which would otherwise get cached for each guest and  
> would also be mapped for each guest separately

nooooooooo. What you're saying there amounts to text replication. There is no proposal here to create duplicated copies of pagecache pages: the VM just doesn't support that (Nick has some protopatches which do this as a possible NUMA optimisation).

So these mmaped pages will continue to be shared across all guests. The problem boils down to "which guest(s) get charged for each shared page".

A simple and obvious and easy-to-implement answer is "the guest which paged it in". I think we should firstly explain why that is insufficient.

---

Containers mailing list  
[Containers@lists.osdl.org](mailto:Containers@lists.osdl.org)  
<https://lists.osdl.org/mailman/listinfo/containers>

---