

---

Subject: Re: RSS controller v2 Test results (lmbench )  
Posted by [Balbir Singh](#) on Fri, 18 May 2007 03:46:46 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Andrew Morton wrote:

> On Thu, 17 May 2007 23:20:12 +0530

> Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>

>> A meaningful container size does not hamper performance. I am in the process  
>> of getting more results (with varying container sizes). Please let me know  
>> what you think of the results? Would you like to see different benchmarks/  
>> tests/configuration results?

>>

>> Any feedback, suggestions to move this work forward towards identifying  
>> and correcting bottlenecks or to help improve it is highly appreciated.

>

> <wakes up>

>

> Memory reclaim tends not to consume much CPU. Because in steady state it  
> tends to be the case that the memory reclaim rate (and hopefully the  
> scanning rate) is equal to the disk IO rate.

>

With the memory controller, I suspect memory reclaim will become a function  
of the memory the container tries to touch that lies outside its limit.

If a container requires 512 MB of memory and we configure the container  
size as 256 MB, then we might see aggressive memory reclaim. We do provide  
some statistics to help the user figure out if the reclaim is aggressive,  
we'll try and add more statistics.

> Often the most successful way to identify performance problems in there is  
> by careful code inspection followed by development of exploits.

>

> Is this RSS controller built on Paul's stuff, or is it standalone?

>

It's built on top of the containers infrastructure. Version 2 was posted  
on top of containers v8.

> Where do we stand on all of this now anyway? I was thinking of getting Paul's  
> changes into -mm soon, see what sort of calamities that brings about.

>

The RSS controller was posted by Pavel based on some initial patches by me,  
so we are in agreement w.r.t approach to memory control. Vaidy is working  
on a page cache controller, we are able to use the existing RSS infrastructure  
for writing the page cache controller (unmapped). All the stake holders are

on cc, I would request them to speak out on the issues and help build a way to take this forward.

I've been reviewing and testing Paul's containers v9 patches. As and when I find more issues, I plan to send out fixes. It'll be good to have the containers infrastructure in -mm, so that we can start posting controllers against them for review and acceptance.

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL

---