

---

Subject: Re: Re: [RFC] Virtualization steps

Posted by [Kirill Korotaev](#) on Wed, 29 Mar 2006 00:55:01 GMT

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

---

> Kirill Korotaev wrote:

>>> Oh, after you come to an agreement and start posting patches, can you  
>>> also outline why we want this in the kernel (what it does that low  
>>> level virtualization doesn't, etc, etc), and how and why you've agreed  
>>> to implement it. Basically, some background and a summary of your  
>>> discussions for those who can't follow everything. Or is that a faq  
>>> item?

>> Nick, will be glad to shed some light on it.

>>

>> First of all, what it does which low level virtualization can't:

>> - it allows to run 100 containers on 1GB RAM

>> (it is called containers, VE - Virtual Environments,

>> VPS - Virtual Private Servers).

>> - it has no much overhead (<1-2%), which is unavoidable with hardware

>> virtualization. For example, Xen has >20% overhead on disk I/O.

>

> I think the Xen guys would disagree with you on this. Xen claims <3%

> overhead on the XenSource site.

>

> Where did you get these figures from? What Xen version did you test?

> What was your configuration? Did you have kernel debugging enabled? You

> can't just post numbers without the data to back it up, especially when

> it conflicts greatly with the Xen developers statements. AFAIK Xen is

> well on it's way to inclusion into the mainstream kernel.

I have no exact numbers in the hands as I'm in another country right now.

But! We tested Xen not long ago with iozone test suite and it gave

~20-30% disk I/O overhead. Recently we were testing CPU scheduler and

EDF scheduler gave me 33% overhead on some very simple loads with almost

busy loops inside VMs. It also was not providing any good fairness on

2CPU SMP system to my surprise. You can object to me, but better simply

retest it if interested yourself. There were other tests as well, which

reported very different overheads on Xen 3. I suppose Xen guys do such

measurements themselves, no?

And I'm sure, they are constantly improving it, they are doing a good  
work on it.

Thanks,

Kirill

---