Subject: Re: [RFC] Virtualization steps
Posted by dev on Wed, 29 Mar 2006 01:39:00 GMT
View Forum Message <> Reply to Message

Nick,

>> 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.
>
> Are any future hardware solutions likely to improve these problems?
Probably you are aware of VT-i/VT-x technologies and planned virtualized
MMU and I/O MMU from Intel and AMD.
These features should improve the performance somehow, but there is
still a limit for decreasing the overhead, since at least disk, network,
video and such devices should be emulated.

>> OS kernel virtualization
>> ~~~~~~~~~~~~~~~~~~~~~~~~
>
> Is this considered secure enough that multiple untrusted VEs are run
> on production systems?
it is secure enough. What makes it secure? In general:
- virtualization, which makes resources private
- resource control, which makes VE to be limited with its usages
In more technical details virtualization projects make user access (and
capabilities) checks stricter. Moreover, OpenVZ is using "denied by
default" approach to make sure it is secure and VE users are not allowed
something else.

Also, about 2-3 month ago we had a security review of OpenVZ project
made by Solar Designer. So, in general such virtualization approach
should be not less secure than VM-like one. VM core code is bigger and
there is enough chances for bugs there.

> What kind of users want this, who can't use alternatives like real
> VMs?
Many companies, just can't share their names. But in general no
enterprise and hosting companies need to run different OSes on the same
machine. For them it is quite natural to use N machines for Linux and M
for Windows. And since VEs are much more lightweight and easier to work
with, they like it very much.

Just for example, OpenVZ core is running more than 300,000 VEs worldwide.

Thanks,
Kirill