
Subject: openvz and qemu

Posted by [Dietmar Maurer](#) on Fri, 01 Feb 2008 06:05:14 GMT

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

Hi all,

i want to run a few qemu VMs inside VE0. Just wonder how to set
fairsched properties
for a single process?

- howto create a new fairsched group (for the qemu process)?
- howto assign weight/rate

or is that not possible at all?

- Dietmar

Subject: AW: openvz and qemu

Posted by [Dietmar Maurer](#) on Fri, 01 Feb 2008 06:23:51 GMT

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

OK, basically found the syscall interface in kernel/fairsched.c

But how do I find out the syscall numbers, i.e how do I
call sys_fairsched_mknod() with userspace syscall()

- Dietmar

> Von: users-bounces@openvz.org
> [mailto:users-bounces@openvz.org] Im Auftrag von Dietmar Maurer
> Gesendet: Freitag, 01. Februar 2008 07:05
> An: users@openvz.org
> Betreff: [Users] openvz and qemu
>
> Hi all,
>
> i want to run a few qemu VMs inside VE0. Just wonder how to
> set fairsched properties for a single process?
>
> - howto create a new fairsched group (for the qemu process)?
> - howto assign weight/rate
>
> or is that not possible at all?
>

Subject: Re: AW: openvz and qemu
Posted by [dev](#) on Fri, 01 Feb 2008 07:13:41 GMT
[View Forum Message](#) <> [Reply to Message](#)

check vzctl sources. it calls other OVZ-specific syscalls by their numbers, like sys_setluid, sys_setublimit etc.

For creating fairsched node you'll need to call fairsched_mknod first, then move the process to this node using fairsched_mvpr

Dietmar Maurer wrote:

> OK, basically found the syscall interface in kernel/fairsched.c

>

> But how do I find out the syscall numbers, i.e how do I

> call sys_fairsched_mknod() with userspace syscall()

>

> - Dietmar

>

>> Von: users-bounces@openvz.org

>> [mailto:users-bounces@openvz.org] Im Auftrag von Dietmar Maurer

>> Gesendet: Freitag, 01. Februar 2008 07:05

>> An: users@openvz.org

>> Betreff: [Users] openvz and qemu

>>

>> Hi all,

>>

>> i want to run a few qemu VMs inside VE0. Just wonder how to

>> set fairsched properties for a single process?

>>

>> - howto create a new fairsched group (for the qemu process)?

>> - howto assign weight/rate

>>

>> or is that not possible at all?

>>

>

>

Subject: AW: AW: openvz and qemu
Posted by [Dietmar Maurer](#) on Fri, 08 Feb 2008 17:07:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

> check vzctl sources. it calls other OVZ-specific syscalls by
> their numbers, like sys_setluid, sys_setublimit etc.

>

> For creating fairsched node you'll need to call

> fairsched_mknod first, then move the process to this node
> using fairsched_mvpr

OK, I wrote a small (perl) script to test those syscalls, and everything work as expected, for example:

```
..  
fairsched_mknod (0, 500, 400);  
fairsched_vcpus (400, 1);  
set_cpulimit (400, 25);  
fairsched_mvpr ($$, 400);  
fork_and_do_something(); # runs with 25% cpu  
wait  
fairsched_rmnod (400);  
..
```

work perfectly.

But when I call fairsched_rmnod(400) from another process while there is still a process running inside 400 the system freezes. I get a kernel panic somewhere inside move_task_off_dead_cpu.

any idea how to fix that?

- Dietmar

Subject: Re: AW: AW: openvz and qemu
Posted by [Kirill Korotaev](#) on Sat, 09 Feb 2008 11:07:11 GMT
[View Forum Message](#) <> [Reply to Message](#)

node should be deleted normally only when no processes left.
it's possible to fix syscall to return an error when node is non-empty...

Dietmar Maurer wrote:

```
>  
>> check vzctl sources. it calls other OVZ-specific syscalls by  
>> their numbers, like sys_setluid, sys_setublimit etc.  
>>  
>> For creating fairsched node you'll need to call  
>> fairsched_mknod first, then move the process to this node  
>> using fairsched_mvpr  
>  
> OK, I wrote a small (perl) script to test those syscalls, and everything  
> work as expected, for example:  
>  
> ..
```

> fairsched_mknod (0, 500, 400);
> fairsched_vcpus (400, 1);
> set_cpulimit (400, 25);
> fairsched_mvpr (\$\$, 400);
> fork_and_do_something(); # runs with 25% cpu
> wait
> fairsched_rmnod (400);
> ..
>
> work perfectly.
>
> But when I call fairsched_rmnod(400) from another process while there is
> still a process running inside 400 the system freezes. I get a kernel
> panic somewhere inside move_task_off_dead_cpu.
>
> any idea how to fix that?
>
> - Dietmar
>
>
>
>
>
>

Subject: AW: AW: AW: openvz and qemu
Posted by [Dietmar Maurer](#) on Sun, 10 Feb 2008 07:15:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

> node should be deleted normally only when no processes left.
> it's possible to fix syscall to return an error when node is
> non-empty...

Ok, will try make sure that no processes left. But I thing the syscall
should be fixed too.

- Dietmar

Subject: Re: AW: AW: AW: openvz and qemu
Posted by [Kirill Korotaev](#) on Mon, 11 Feb 2008 12:47:06 GMT
[View Forum Message](#) <> [Reply to Message](#)

care to create a patch?

Dietmar Maurer wrote:

>> node should be deleted normally only when no processes left.
>> it's possible to fix syscall to return an error when node is
>> non-empty...
>
> Ok, will try make sure that no processes left. But I thing the syscall
> should be fixed too.
>
> - Dietmar
>
>

Subject: AW: AW: AW: openvz and qemu
Posted by [Dietmar Maurer](#) on Mon, 11 Feb 2008 15:58:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

> care to create a patch?

Please can you give me a hint where i can find the list of
processes associated with a fairsched node?

- Dietmar

Subject: Re: AW: AW: AW: openvz and qemu
Posted by [Kirill Korotaev](#) on Mon, 11 Feb 2008 16:26:37 GMT
[View Forum Message](#) <> [Reply to Message](#)

AFAIR the only way is to take tasklist_lock
and check for tasks having vsched == your-node-vsched.

Kirill

Dietmar Maurer wrote:
>> care to create a patch?
>
> Please can you give me a hint where i can find the list of
> processes associated with a fairsched node?
>
> - Dietmar
>
>
