
Subject: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Daniel Pittman](#) on Sat, 07 Apr 2007 07:46:52 GMT
[View Forum Message](#) <> [Reply to Message](#)

G'day. I have been playing with the 'upstart' daemon that replaces the traditional /sbin/init process in Ubuntu Edgy and Feisty.

I didn't have a huge degree of luck with this -- once I allocated a tty for the console I ended up with the system failing during the bootstrap process.

Before investing more time in this I was wondering if anyone else had been working on getting upstart to cooperate in the OpenVZ boot process?

Regards,
Daniel

--
Digital Infrastructure Solutions -- making IT simple, stable and secure
Phone: 0401 155 707 email: contact@digital-infrastructure.com.au
<http://digital-infrastructure.com.au/>

Subject: Re: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Daniel Pittman](#) on Sat, 07 Apr 2007 13:20:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

Daniel Pittman <daniel@rimspace.net> writes:

First, sorry for the cross-post. Since this is an issue with the interaction between the two software packages -- and seems to be of general interest on the OpenVZ users list where I first posted -- I felt it was best to copy all three locations up front.

> G'day. I have been playing with the 'upstart' daemon that replaces
> the traditional /sbin/init process in Ubuntu Edgy and Feisty.
>
> I didn't have a huge degree of luck with this -- once I allocated a
> tty for the console I ended up with the system failing during the
> bootstrap process.
>
> Before investing more time in this I was wondering if anyone else had
> been working on getting upstart to cooperate in the OpenVZ boot
> process?

I spent some more time tracking this down and seem to be a conflict between the expectations of the Upstart process and the environment provided by OpenVZ when starting init in a new VE.

All the testing was with vzctl 3.0.16-1dso2 (from Debian/testing) and Upstart 0.3.8 from Feisty.

At issue are the expectations about file description layout that are made in init/main.c of upstart; on line 150 the 'control_open' method is called.

That, eventually, binds a Unix domain socket for the Upstart control connection, used to support 'telinit' among other options.

Later, on line 209, Upstart cleans up the environment:

```
/* Close any file descriptors we inherited, and open the console
 * device instead. Normally we reset the console, unless we're
 * inheriting one from another init process.
 */
for (int i = 0; i < 3; i++)
    close (i);
```

During a normal boot sequence those three descriptors are bound to STDIN, OUT and ERR just as expected. The result is that Upstart correctly rebinds the console and everything "just works."

Under OpenVZ startup, however, things don't quite look the same. OpenVZ will exec the init process with **no** open file descriptors.

This causes the control_open method to bind file description 0 to the Unix domain socket -- and the later loop to carefully close the socket again.

I was able to verify my belief that this was causing the problems by replacing init with a shell script that simply did:

```
exec /sbin/init.real "$@" </dev/null >/dev/null 2>/dev/null
```

Once that was in place Upstart kept the control socket open and proceeded considerably further through the boot sequence.

It isn't quite clear to me who is at fault here -- I presume that in the real world the kernel and linuxrc processes never start upstart without the first three file descriptors open, but I don't know that is assured anywhere.

I am happy to help with further testing or to explain how to configure an OpenVZ environment to reproduce this, if desired.

Regards,
Daniel

--

Digital Infrastructure Solutions -- making IT simple, stable and secure
Phone: 0401 155 707 email: contact@digital-infrastructure.com.au
<http://digital-infrastructure.com.au/>

Subject: [Devel] Re: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Matthias Urlichs](#) on Sat, 07 Apr 2007 13:51:22 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

Daniel Pittman:

> Under OpenVZ startup, however, things don't quite look the same.
> OpenVZ will exec the init process with *no* open file descriptors.
>

Upstart/init is not the only program that breaks when stdin/out/err are closed when you call it.

On the other hand, it's not too much to expect upstart to be a bit more lenient in what it expects, given that there may or may not be a way to have these file descriptors open when launching it.

So in this case I'd like to apply the "when in doubt, do both" guideline and fix upstart *and* openVZ.

--

Matthias Urlichs | {M:U} IT Design @ m-u-it.de | smurf@smurf.noris.de
Disclaimer: The quote was selected randomly. Really. | <http://smurf.noris.de>

--

To be thrown on one's own resources is to be cast in the very lap of fortune; for our faculties undergo a development, and display an energy, of which they were previously unsusceptible.

-- Benjamin Franklin

Devel mailing list
Devel@openvz.org
<https://openvz.org/mailman/listinfo/devel>

Subject: [Devel] Re: OpenVZ and the Ubuntu Upstart init daemon

On Sat, 2007-04-07 at 23:20 +1000, Daniel Pittman wrote:

> > G'day. I have been playing with the 'upstart' daemon that replaces
> > the traditional /sbin/init process in Ubuntu Edgy and Feisty.
> >
> > I didn't have a huge degree of luck with this -- once I allocated a
> > tty for the console I ended up with the system failing during the
> > bootstrap process.
> >
> > Before investing more time in this I was wondering if anyone else had
> > been working on getting upstart to cooperate in the OpenVZ boot
> > process?
>
> I spent some more time tracking this down and seem to be a conflict
> between the expectations of the Upstart process and the environment
> provided by OpenVZ when starting init in a new VE.
>
I appreciate your debugging efforts; I've been asking OpenVZ users to
help with the problem for some time with little avail.

> All the testing was with vzctl 3.0.16-1dso2 (from Debian/testing) and
> Upstart 0.3.8 from Feisty.
>
> At issue are the expectations about file description layout that are
> made in init/main.c of upstart; on line 150 the 'control_open' method is
> called.
[...]
> This causes the control_open method to bind file description 0 to the
> Unix domain socket -- and the later loop to carefully close the socket
> again.

>
This was a bug introduced relatively recently while clearing out the
main() function, and has already been fixed in bzip trunk:

http://codebrowse.launchpad.net/~keybuk/upstart/main/revision/scott%40netsplit.com-20070313191319-gztu8c0r0sjla0hp?start_revid=scott%40netsplit.com-20070316171800-scmrd6w9r22uf4me

I've had reports of OpenVZ failing for some time before this though; can
you confirm that with this patch, Upstart works correctly?

Scott

--

Scott James Remnant
Ubuntu Development Manager
scott@ubuntu.com

Subject: [Devel] Re: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Daniel Pittman](#) on Sun, 08 Apr 2007 03:41:54 GMT
[View Forum Message](#) <> [Reply to Message](#)

Scott James Remnant <scott@ubuntu.com> writes:
> On Sat, 2007-04-07 at 23:20 +1000, Daniel Pittman wrote:

Please respect the reply-to -- the issues, if any, now seem to be located in the Upstart area rather than OpenVZ specific.

>> > G'day. I have been playing with the 'upstart' daemon that replaces
>> > the traditional /sbin/init process in Ubuntu Edgy and Feisty.

[...]

>> I spent some more time tracking this down and seem to be a conflict
>> between the expectations of the Upstart process and the environment
>> provided by OpenVZ when starting init in a new VE.
>
> I appreciate your debugging efforts; I've been asking OpenVZ users to
> help with the problem for some time with little avail.

Well, hopefully I should be able to dig through the issues and get to the bottom of the problems.

[...]

>> This causes the control_open method to bind file description 0 to the
>> Unix domain socket -- and the later loop to carefully close the socket
>> again.
>
> This was a bug introduced relatively recently while clearing out the
> main() function, and has already been fixed in bzd trunk:

Great. I have worked around that for the moment with my local shell script and am working on getting the rest of the boot process up and running.

Is it likely that this issue will be fixed before the Feisty release?
I know the freeze is in effect at the moment but this is definitely a big incompatibility with the OpenVZ system.

> I've had reports of OpenVZ failing for some time before this though;
> can you confirm that with this patch, Upstart works correctly?

No, but it certainly gets a lot further. I am still working on debugging the early "boot" process to find out where it fails.

I will follow up on this (on the Ubuntu list) once I have more results. Hopefully this should lead to full compatibility between the two.

Once I have final results I will also copy of OpenVZ list, but I don't think they need to know all the messy details of the debugging process.

Regards,
Daniel

--

Digital Infrastructure Solutions -- making IT simple, stable and secure
Phone: 0401 155 707 email: contact@digital-infrastructure.com.au
<http://digital-infrastructure.com.au/>

Devel mailing list
Devel@openvz.org
<https://openvz.org/mailman/listinfo/devel>

Subject: RE: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Stephen Fletcher](#) on Sun, 08 Apr 2007 05:15:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

Just set up OpenVZ on a new machine and followed how I did it on my first but main difference is I am using a different kernel build.

```
root@stephen2:/usr/src/openvz# vzctl start 100
UB parameter kmemsize not set
UB parameter lockedpages not set
UB parameter privvmpages not set
UB parameter shmpages not set
UB parameter numproc not set
UB parameter physpages not set
UB parameter vmguarpages not set
UB parameter oomguarpages not set
UB parameter numtcpsock not set
UB parameter numflock not set
UB parameter numpty not set
UB parameter numsiginfo not set
UB parameter tcpsndbuf not set
```

UB parameter tcpvbuf not set
UB parameter othersockbuf not set
UB parameter dgramrcvbuf not set
UB parameter numothersock not set
UB parameter numfile not set
UB parameter dcachesize not set
UB parameter numiptent not set

Wtf is its problem? :)

Thanks,
Stephen

Subject: Re: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Daniel Pittman](#) on Sun, 08 Apr 2007 05:25:19 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Stephen Fletcher" <fletch@cisco.com> writes:

G'day Stephen.

First of all: since this question isn't even vaguely related to problems with Ubuntu Upstart and OpenVZ it would have been polite to start a new thread.

Even if you had bothered to change the subject line there are a bunch of internal headers your mail client inserts that tell other mail software what you replied to.

As a result of this your message was specifically highlighted to me as part of a response to my post -- and being irritated because it has nothing to do with the subject at hand doesn't incline people to help you very much. :/

Regards,
Daniel

--

Digital Infrastructure Solutions -- making IT simple, stable and secure
Phone: 0401 155 707 email: contact@digital-infrastructure.com.au
<http://digital-infrastructure.com.au/>

Subject: RE: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Stephen Fletcher](#) on Sun, 08 Apr 2007 05:48:21 GMT
[View Forum Message](#) <> [Reply to Message](#)

Ill try and make up for my lameness :)

Bugzilla Ubuntu Edgy Upstart bugid
http://bugzilla.openvz.org/show_bug.cgi?id=436

Something to try based on someone elses previous post (didn't get it working for me)

Getting Edgy to work: patch /etc/init.d/rc, add comments so QUIET is always yes:

```
#if grep -w -q quiet /proc/cmdline 2>/dev/null; then
    QUIET=yes
#else
#   QUIET=no
#fi
```

Subject: Re: OpenVZ and the Ubuntu Upstart init daemon
Posted by [Daniel Pittman](#) on Sun, 08 Apr 2007 06:45:26 GMT
[View Forum Message](#) <> [Reply to Message](#)

"Stephen Fletcher" <fletch@cisco.com> writes:

> Ill try and make up for my lameness :)

Actually, the bug has been tracked down -- when upstart is run without any file descriptors open (as it is under OpenVZ but not during normal Ubuntu booting) it would accidentally close the control socket.

I am waiting to hear about the one remaining issue, a signal killing upstart in the VE where it doesn't kill the default init, before I can declare Feisty fully working ... but I do have a patched upstart and init scripts set that let me run up a full VE now. :)

Thanks for the comment though. It looks like that effect is achieved by OpenVZ -- at least in my version -- which claims the kernel command line had **only** the word 'quiet' present.

Regards,
Daniel

--

Digital Infrastructure Solutions -- making IT simple, stable and secure
Phone: 0401 155 707 email: contact@digital-infrastructure.com.au
<http://digital-infrastructure.com.au/>
