Subject: vzctl: race condition at open("/sbin/init")
Posted by Vasily Kulikov on Wed, 25 Jul 2012 19:07:50 GMT
View Forum Message <> Reply to Message

Hi,

stat()+open() is not atomic in the code below, so there is a race condition. A container root may change /sbin/init between these calls to e.g. FIFO and then make the vzctl's process hang up on read().

I'd add O_NOCTTY to open's flags and change stat() before open() to fstat() just after open().

vzctl-3.3/src/lib/readelf.c:

```
int get_arch_from_elf(const char *file)
{
...
 if (stat(file, &st)) <<<<<
  return -1;
 if (!S_ISREG(st.st_mode))
  return -1;
 fd = open(file, O_RDONLY); <<<<<
 if (fd < 0)
  return -1;
 nbytes = read(fd, (void *) &elf_hdr, sizeof(elf_hdr));
...
}
```

Thanks,

--
Vasily

Subject: Re: vzctl: race condition at open("/sbin/init")
Posted by kir on Tue, 18 Sep 2012 15:09:53 GMT
View Forum Message <> Reply to Message

On 07/25/2012 11:07 PM, Vasily Kulikov wrote:
> Hi,
>
> stat()+open() is not atomic in the code below, so there is a race
> condition. A container root may change /sbin/init between these calls
> to e.g. FIFO and then make the vzctl's process hang up on read().
>
> I'd add O_NOCTTY to open's flags and change stat() before open() to

> fstat() just after open().

Thanks a lot for reporting!

Does this patch seems sufficient to you?
http://git.openvz.org/?p=vzctl;a=commitdiff;h=7c47a7953

>
>
> vzctl-3.3/src/lib/readelf.c:
>
> int get_arch_from_elf(const char *file)
> {
> ...
> if (stat(file, &st)) <<<<<
>   return -1;
> if (!S_ISREG(st.st_mode))
>   return -1;
> fd = open(file, O_RDONLY); <<<<<
> if (fd < 0)
>   return -1;
> nbytes = read(fd, (void *) &elf_hdr, sizeof(elf_hdr));
> ...
> }
>
> Thanks,
>

Subject: Re: vzctl: race condition at open("/sbin/init")
Posted by Vasily Kulikov on Tue, 18 Sep 2012 20:12:44 GMT
View Forum Message <> Reply to Message

On Tue, Sep 18, 2012 at 19:09 +0400, Kir Kolyshkin wrote:
> On 07/25/2012 11:07 PM, Vasily Kulikov wrote:
> >Hi,
> >
> >stat()+open() is not atomic in the code below, so there is a race
> >condition. A container root may change /sbin/init between these calls
> >to e.g. FIFO and then make the vzctl's process hang up on read().
> >
> >I'd add O_NOCTTY to open's flags and change stat() before open() to
> >fstat() just after open().
>
> Thanks a lot for reporting!
>
> Does this patch seems sufficient to you?
> http://git.openvz.org/?p=vzctl;a=commitdiff;h=7c47a7953

Yes, look good.

Thanks!

--
Vasiliy Kulikov
http://www.openwall.com - bringing security into open computing environments