
Subject: Re: [PATCH 4/4] The control group itself
Posted by [Paul Menage](#) on Thu, 21 Feb 2008 20:47:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Thu, Feb 7, 2008 at 5:01 AM, Pavel Emelyanov <xemul@openvz.org> wrote:

```
> +
> +[cb] <major>:(<minor>|*) [r-][w-]
> + ^      ^      ^
> + |      |      |
> + |      |      +--- access rights (1)
> + |      |
> + |      +-- device major and minor numbers (2)
> + |
> + +-- device type (character / block)
> ...
> +When reading from it, one may see something like
> +
> +   c 1:5 rw
> +   b 8:* r-
> +
```

In the interest of avoiding proliferating cgroup control file formats, I'm wondering if we can abstract out the general form of the data being presented here and maybe simplify it in such a way that we can hopefully reuse the format for other control files in the future?

For example, one way to represent this would be a map from device strings such c:1:5 to permission strings such as rw. Or maybe numerical device ids to numerical permission values.

The alternative might be to accept that there are two kinds of control files - those which are likely to be programmatically read (e.g. resource usage values), and those that are likely to be programmatically written but only actually read by humans for debugging purposes (like this one) and make it clear up-front when a control file is added which type they're considered to be. We could then ignore the API consistency requirements for the debugging-readable files.

On a separate note, can you document the recommended way to completely overwrite the set of device permissions for a cgroup? Does this involves writing a "--" permission for every device that you don't want in the cgroup?

Paul

Containers mailing list
Containers@lists.linux-foundation.org

