Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by Paul Menage on Thu, 08 Mar 2007 00:58:43 GMT

View Forum Message <> Reply to Message

On 3/7/07, Sam Vilain <sam@vilain.net> wrote:

- > But "namespace" has well-established historical semantics too a way
- > of changing the mappings of local * to global objects. This
- > accurately describes things liek resource controllers, cpusets, resource
- > monitoring, etc.

Sorry, I think this statement is wrong, by the generally established meaning of the term namespace in computer science.

>

- > Trying to extend the well-known term namespace to refer to things that
- > are semantically equivalent namespaces is a useful approach, IMHO.

>

Yes, that would be true. But the kinds of groupings that we're talking about are supersets of namespaces, not semantically equivalent to them. To use Eric's "shoe" analogy from earlier, it's like insisting that we use the term "sneaker" to refer to all footware, including ski boots and birkenstocks ...

Paul

Containers mailing list Containers@lists.osdl.org https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by ebiederm on Thu, 08 Mar 2007 01:32:10 GMT

View Forum Message <> Reply to Message

"Paul Menage" <menage@google.com> writes:

- > On 3/7/07, Sam Vilain <sam@vilain.net> wrote:
- >> But "namespace" has well-established historical semantics too a way
- >> of changing the mappings of local * to global objects. This
- >> accurately describes things liek resource controllers, cpusets, resource
- >> monitoring, etc.

>

- > Sorry, I think this statement is wrong, by the generally established
- > meaning of the term namespace in computer science.

>

>>

>> Trying to extend the well-known term namespace to refer to things that

>> are semantically equivalent namespaces is a useful approach, IMHO.

>>

>

- > Yes, that would be true. But the kinds of groupings that we're talking
- > about are supersets of namespaces, not semantically equivalent to
- > them. To use Eric's "shoe" analogy from earlier, it's like insisting
- > that we use the term "sneaker" to refer to all footware, including ski
- > boots and birkenstocks ...

Pretty much. For most of the other cases I think we are safe referring to them as resource controls or resource limits. I know that roughly covers what cpusets and beancounters and ckrm currently do.

The real trick is that I believe these groupings are designed to be something you can setup on login and then not be able to switch out of. Which means we can't use sessions and process groups as the grouping entities as those have different semantics.

Eric

Containers mailing list Containers@lists.osdl.org https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by Sam Vilain on Thu, 08 Mar 2007 02:47:24 GMT

View Forum Message <> Reply to Message

Paul Menage wrote:

- > Sorry, I think this statement is wrong, by the generally established
- > meaning of the term namespace in computer science.

>

Sorry, I didn't realise I was talking with somebody qualified enough to speak on behalf of the Generally Established Principles of Computer Science.

- >> Trying to extend the well-known term namespace to refer to thingsthat
- >> are semantically equivalent namespaces is a useful approach, IMHO.

>>

>>

- > Yes, that would be true. But the kinds of groupings that we're talking
- > about are supersets of namespaces, not semantically equivalent to
- > them. To use Eric's "shoe" analogy from earlier, it's like insisting
- > that we use the term "sneaker" to refer to all footware, including ski

> boots and birkenstocks ...

>

I see it more like insisting that we use the term "clothing" to also refer to "weapons" because for both of them you tell your body to "wear" them in some game.

This is the classic terminology problem between substance and function. ie, some things share characteristics but does that mean they are the same thing?

Look, I already agreed in the earlier thread that the term "namespace" was being stretched beyond belief, yet instead of trying to be useful about this you still insist on calling this sub-system specific stuff the "container", and then go screaming that I am wrong and you are right on terminology.

I've normally recognised[1] these three things as the primary feature groups of vserver:

- isolation
- resource limiting
- resource sharing

So I've got no problem with using "clothing" remaining for isolation and "weapons" for resource sharing and limiting. Or some other suitable terms.

Sam.

1. eg, http://utsl.gen.nz/talks/vserver/slide4c.html

Containers mailing list Containers@lists.osdl.org https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by Herbert Poetzl on Fri, 09 Mar 2007 00:53:57 GMT

View Forum Message <> Reply to Message

On Wed, Mar 07, 2007 at 06:32:10PM -0700, Eric W. Biederman wrote:

> "Paul Menage" <menage@google.com> writes:

>

- >> On 3/7/07, Sam Vilain <sam@vilain.net> wrote:
- >>> But "namespace" has well-established historical semantics too a way
- >>> of changing the mappings of local * to global objects. This
- >>> accurately describes things liek resource controllers, cpusets, resource

>>> monitoring, etc. >> >> Sorry, I think this statement is wrong, by the generally established >> meaning of the term namespace in computer science. >> >>> Trying to extend the well-known term namespace to refer to things >>> that are semantically equivalent namespaces is a useful approach. >>> IMHQ. >> >> Yes, that would be true. But the kinds of groupings that we're talking >> about are supersets of namespaces, not semantically equivalent to >> them. To use Eric's "shoe" analogy from earlier, it's like insisting >> that we use the term "sneaker" to refer to all footware, including ski >> boots and birkenstocks ... > Pretty much. For most of the other cases I think we are safe referring > to them as resource controls or resource limits. > I know that roughly covers what cpusets and beancounters and ckrm > currently do. let me tell you, it also covers what Linux-VServer does :) > The real trick is that I believe these groupings are designed to > be something you can setup on login and then not be able to switch > out of. Which means we can't use sessions and process groups as the > grouping entities as those have different semantics. precisely, once you are inside a resource container, you must not have the ability to modify its limits, and to some degree, you should not know about the actual available resources, but only about the artificial limits HTC. Herbert > Eric > Containers mailing list > Containers@lists.osdl.org > https://lists.osdl.org/mailman/listinfo/containers

Containers mailing list Containers@lists.osdl.org https://lists.osdl.org/mailman/listinfo/containers Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by Paul Jackson on Fri, 09 Mar 2007 04:30:41 GMT

View Forum Message <> Reply to Message

- > The real trick is that I believe these groupings are designed to be something
- > you can setup on login and then not be able to switch out of. Which means
- > we can't use sessions and process groups as the grouping entities as those
- > have different semantics.

Not always on login. For big administered systems, we use batch schedulers to manage the placement of multiple jobs, submitted to a run queue by users, onto the available compute resources.

But I agree with your conclusion - the existing task grouping mechanisms, while useful for some purposes, don't meet the need here.

--

I won't rest till it's the best ...
Programmer, Linux Scalability
Paul Jackson <pj@sgi.com> 1.925.600.0401

Containers mailing list
Containers@lists.osdl.org
https://lists.osdl.org/mailman/listinfo/containers

Subject: Re: [ckrm-tech] [PATCH 0/2] resource control file system - aka containers on top of nsproxy!

Posted by Srivatsa Vaddagiri on Fri, 09 Mar 2007 18:19:08 GMT View Forum Message <> Reply to Message

On Fri, Mar 09, 2007 at 01:53:57AM +0100, Herbert Poetzl wrote:

- > > The real trick is that I believe these groupings are designed to
- > > be something you can setup on login and then not be able to switch
- >> out of. Which means we can't use sessions and process groups as the
- > > grouping entities as those have different semantics.

>

- > precisely, once you are inside a resource container, you
- > must not have the ability to modify its limits, and to
- > some degree, you should not know about the actual available
- > resources, but only about the artificial limits
- >From non-container workload management perspective, we do desire dynamic manipulation of limits associated with a group and also the ability to move tasks across resource-classes/groups.

--

Regards, vatsa

Containers mailing list Containers@lists.osdl.org https://lists.osdl.org/mailman/listinfo/containers