
Subject: Re: [0/1] [patch -mm] Add containerstats (v3)

Posted by [akpm](#) on Fri, 08 Jun 2007 19:48:16 GMT

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

On Fri, 08 Jun 2007 23:43:46 +0530

Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

> This patch implements per container statistics infrastructure and re-uses
> code from the taskstats interface.

boggle.

Symbol: CONTAINERS [=y]

Selected by: CONTAINER_DEBUG || CPUSETS && SMP || CONTAINER_CPUACCT

Paul, that's just bizarre. How come it was done this way?

<struggles for a while, works out how to make CONFIG_CONTAINERS go away>

OK, so taskstats.c still compiles, but I'm surprised. Shouldn't all that newly-added container stuff in taskstats.c be inside CONFIG_CONTAINERS?

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)

Posted by [Balbir Singh](#) on Sat, 09 Jun 2007 02:52:43 GMT

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

Andrew Morton wrote:

> On Fri, 08 Jun 2007 23:43:46 +0530

> Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>

>> This patch implements per container statistics infrastructure and re-uses
>> code from the taskstats interface.

>

> boggle.

>

> Symbol: CONTAINERS [=y]

> Selected by: CONTAINER_DEBUG || CPUSETS && SMP || CONTAINER_CPUACCT

>

> Paul, that's just bizarre. How come it was done this way?

>

> <struggles for a while, works out how to make CONFIG_CONTAINERS go away>

>

> OK, so taskstats.c still compiles, but I'm surprised. Shouldn't all that

> newly-added container stuff in taskstats.c be inside CONFIG_CONTAINERS?

Hi, Andrew,

I've added a definition of containerstats_build() even when CONFIG_CONTAINERS is turned off, it simply returns -EINVAL. That's why taskstats.c compiles, I could add an #ifdef and move the containers commands under CONFIG_CONTAINERS, but I felt taskstats.c could be clean without any #ifdef's hanging around (as much as possible).

In the next iteration, I could move out all the code to containerstats.c and make containerstats.c depend on CONFIG_CONTAINERS

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)
Posted by [Paul Menage](#) on Sat, 09 Jun 2007 07:09:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/8/07, Andrew Morton <akpm@linux-foundation.org> wrote:

> On Fri, 08 Jun 2007 23:43:46 +0530

> Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>

> > This patch implements per container statistics infrastructure and re-uses

> > code from the taskstats interface.

>

> boggle.

>

> Symbol: CONTAINERS [=y]

> Selected by: CONTAINER_DEBUG || CPUSETS && SMP || CONTAINER_CPUACCT

>

> Paul, that's just bizarre. How come it was done this way?

Containers on their own without any subsystems aren't hugely useful. So the plan was that any container subsystem would "select CONTAINERS" to cause the container framework to be built. I'm happy to change it to invert the dependencies if you'd prefer.

Paul

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)
Posted by [akpm](#) on Sat, 09 Jun 2007 08:02:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sat, 9 Jun 2007 00:09:55 -0700 "Paul Menage" <menage@google.com> wrote:

> On 6/8/07, Andrew Morton <akpm@linux-foundation.org> wrote:
> > On Fri, 08 Jun 2007 23:43:46 +0530
> > Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
> >
> > > This patch implements per container statistics infrastructure and re-uses
> > > code from the taskstats interface.
> >
> > boggle.
> >
> > Symbol: CONTAINERS [=y]
> > Selected by: CONTAINER_DEBUG || CPUSETS && SMP || CONTAINER_CPUACCT
> >
> > Paul, that's just bizarre. How come it was done this way?
>
> Containers on their own without any subsystems aren't hugely useful.
> So the plan was that any container subsystem would "select CONTAINERS"
> to cause the container framework to be built. I'm happy to change it
> to invert the dependencies if you'd prefer.
>

- CONTAINER_DEBUG should depend on CONTAINERS
- the CPUSETS && SMP is weird and should be deleted, unless I'm missing something
- CONTAINERS should depend on CPUSETS
- That leaves CONTAINER_CPUACCT. One _could_ `select' CONTAINERS if CONTAINER_CPUACCT but I suspect that'll screw up the above (select does odd things). I think it's reasonable to have a CONTAINERS menu and under that, all the clients of it: CONTAINER_CPUACCT, CONTAINER_MEMORY_CONTROL, etc.

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)
Posted by [Paul Menage](#) on Sat, 09 Jun 2007 08:07:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/9/07, Andrew Morton <akpm@linux-foundation.org> wrote:

>
> - CONTAINER_DEBUG should depend on CONTAINERS

CONTAINER_DEBUG is actually a container subsystem whose sole purpose is to provide debugging information about any hierarchy that it's mounted as a part of. So in some senses it's in the same boat as something like cpusets or the RSS controller. CONFIG_CONTAINER_DEBUG doesn't affect any of the container framework code.

>
> - the CPUSETS && SMP is weird and should be deleted, unless I'm missing
> something

Cpusets depends on SMP in the vanilla tree, so that's not anything new that I added.

>
> - CONTAINERS should depend on CPUSETS

You mean the other way around?

>
> - That leaves CONTAINER_CPUACCT.

Really, CONTAINER_CPUACCT should have the same relationship to CONTAINERS as CPUSETS does.

Paul

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)
Posted by [akpm](#) on Sat, 09 Jun 2007 08:17:40 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sat, 9 Jun 2007 01:07:53 -0700 "Paul Menage" <menage@google.com> wrote:

> On 6/9/07, Andrew Morton <akpm@linux-foundation.org> wrote:

> >

> > - CONTAINER_DEBUG should depend on CONTAINERS

>

> CONTAINER_DEBUG is actually a container subsystem whose sole purpose

> is to provide debugging information about any hierarchy that it's

> mounted as a part of. So in some senses it's in the same boat as

> something like cpusets or the RSS controller. CONFIG_CONTAINER_DEBUG

> doesn't affect any of the container framework code.

Oh, that's right.

But it still should depend on CONTAINERS

> >

> > - the CPUSETS && SMP is weird and should be deleted, unless I'm missing

> > something

>

> Cpusets depends on SMP in the vanilla tree, so that's not anything new

> that I added.

Oh, OK, so CPUSETS is not a client of CONTAINERS: so it depends on CONTAINERS

> >

> > - CONTAINERS should depend on CPUSETS

>

> You mean the other way around?

yup

> >

> > - That leaves CONTAINER_CPUACCT.

>

> Really, CONTAINER_CPUACCT should have the same relationship to

> CONTAINERS as CPUSETS does.

>

Would it not be simplest to have CONTAINERS as the top-level user-configurable item and to then have everything else depend on it?

select is a nasty thing - we repeatedly have problems when using it.

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

Subject: Re: [0/1] [patch -mm] Add containerstats (v3)
Posted by [Paul Menage](#) on Sat, 09 Jun 2007 08:20:55 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/9/07, Andrew Morton <akpm@linux-foundation.org> wrote:

>

> Would it not be simplest to have CONTAINERS as the top-level
> user-configurable item and to then have everything else depend on it?

>

Yes, OK - it can go that way around too. I guess my thought was that people would be more interested in enabling/disabling the subsystems themselves (cpusets, RSS, etc) than the underlying framework. But if select has problems then it's no trouble to invert it.

Paul

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
