Subject: Re: [RFC][PATCH 7/7] CGroup API: Update cpusets to use cgroup structured file API
Posted by Balbir Singh on Mon, 18 Feb 2008 11:13:33 GMT
View Forum Message <> Reply to Message

* Balbir Singh <balbir@linux.vnet.ibm.com> [2008-02-18 16:42:05]:

> Li Zefan wrote:
> > Paul Jackson wrote:
> >> Ok ... this would (I suspect, just from code reading, no bytes were
> >> harmed in actual testing of this) have a minor change to how white
> >> space is handled writing integer flags to cpuset files, and a minor
> >> inconstency.
> >>
> >> 1) Existing cpuset code lets you set a flag (e.g. cpu_exclusive) by doing:
> >> echo '1 rumplestiltskin' > cpu_exclusive   # same as: echo 1 > cpu_exclusive
> >>   With this patch, that probably fails, EINVAL.
> >>
> >> 2) With this patch, one can write "1" or "1\n" to cpuset integer files, but one
> >>   cannot successfully write "1\r\n" or "1 " or "1 \n". However, for the cpuset
> >>   control files that take strings, not single integers, one -can- have any mix
> >>   of trailing white space.
> >>
> >> So far as I know, I have no requirement to write rumplestiltskin to cpuset files ;).
> >> So I'm content to let the minor change in (1) pass without further comment.
> >>
> >> I'd like to recommend consideration of the following patch, to address the
> >> minor inconsistency of (2), and to save a few bytes of kernel text space.
> >>
> >
> > For memory controller, we have to do this:
> > echo -n 4m > memory.limit_in_bytes
> > '-n' is necessary. This is another inconsistency..
>
Hi. Li,

I have a similar patch that fixes the inconsistency.
It's attached below. Andrew, could we please consider this patch for -mm

The memory controller has a requirement that while writing values, we need
to use echo -n. This patch fixes the problem and makes the UI more consistent.

Signed-off-by: Balbir Singh <balbir@linux.vnet.ibm.com>

 Documentation/controllers/memory.txt | 8 ++++----
 kernel/res_counter.c | 1 +

```
 2 files changed, 5 insertions(+), 4 deletions(-)

diff -puN mm/memcontrol.c~memory-controller-fix-crlf-echo-issue mm/memcontrol.c
diff -puN kernel/res_counter.c~memory-controller-fix-crlf-echo-issue kernel/res_counter.c
--- linux-2.6.25-rc2/kernel/res_counter.c~memory-controller-fix-crlf-echo-issue 2008-02-18
16:15:02.000000000 +0530
+++ linux-2.6.25-rc2-balbir/kernel/res_counter.c 2008-02-18 16:16:16.000000000 +0530
@@ -113,6 +113,7 @@ ssize_t res_counter_write(struct res_cou

  ret = -EINVAL;

+ strstrip(buf);
  if (write_strategy) {
   if (write_strategy(buf, &tmp)) {
    goto out_free;
diff -puN Documentation/controllers/memory.txt~memory-controller-fix-crlf-echo-issue
Documentation/controllers/memory.txt
---
linux-2.6.25-rc2/Documentation/controllers/memory.txt~memory-controller-fix-crlf-echo-issue 2008
-02-18 16:18:26.000000000 +0530
+++ linux-2.6.25-rc2-balbir/Documentation/controllers/memory.txt 2008-02-18
16:18:44.000000000 +0530
@@ -164,7 +164,7 @@ c. Enable CONFIG_CGROUP_MEM_CONT

 Since now we're in the 0 cgroup,
 We can alter the memory limit:
-# echo -n 4M > /cgroups/0/memory.limit_in_bytes
+# echo 4M > /cgroups/0/memory.limit_in_bytes

 NOTE: We can use a suffix (k, K, m, M, g or G) to indicate values in kilo,
 mega or gigabytes.
@@ -185,7 +185,7 @@ number of factors, such as rounding up t
 availability of memory on the system.  The user is required to re-read
 this file after a write to guarantee the value committed by the kernel.

-# echo -n 1 > memory.limit_in_bytes
+# echo 1 > memory.limit_in_bytes
 # cat memory.limit_in_bytes
 4096 Bytes

@@ -197,7 +197,7 @@ caches, RSS and Active pages/Inactive pa

 The memory.force_empty gives an interface to drop *all* charges by force.

-# echo -n 1 > memory.force_empty
+# echo 1 > memory.force_empty

 will drop all charges in cgroup. Currently, this is maintained for test.
```

@@ -238,7 +238,7 @@ rmdir() if there are no tasks.
 The type of memory accounted by the cgroup can be limited to just
 mapped pages by writing "1" to memory.control_type field

-echo -n 1 > memory.control_type
+echo > memory.control_type

 5. TODO

_

--
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