
Subject: [patch -rss] Make RSS accounting display more user friendly

Posted by [Balbir Singh](#) on Wed, 20 Jun 2007 15:46:55 GMT

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

Display the current usage and limit in a more user friendly manner. Number of pages can be confusing if the page size is different. Some systems can choose a page size of 64KB.

NOTE: Values shown in MB and GB could be rounded off.

Sample display

```
:~ # cat /container/rss_limit
17179869183 (GB), 9223372036854775807 pages
:~ # cat /container/rss_usage
36 (MB), 9220 pages
:~ #
```

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

```
include/linux/res_counter.h | 6 ++++++
kernel/res_counter.c        | 27 ++++++-----
2 files changed, 29 insertions(+), 4 deletions(-)
```

```
diff -puN kernel/res_counter.c~rss-ui-display-kdb kernel/res_counter.c
--- linux-2.6.22-rc24-mm2/kernel/res_counter.c~rss-ui-display-kdb 2007-06-20
20:15:46.000000000 +0530
+++ linux-2.6.22-rc24-mm2-balbir/kernel/res_counter.c 2007-06-20 21:12:29.000000000 +0530
@@ -61,7 +61,8 @@ void res_counter_uncharge(struct res_cou
}
```

```
-static inline unsigned long *res_counter_member(struct res_counter *cnt, int member)
+static inline unsigned long *res_counter_member(struct res_counter *cnt,
+ int member, int *format_mem)
{
    switch (member) {
        case RES_USAGE:
@@ -69,6 +70,7 @@ static inline unsigned long *res_counter
        case RES_LIMIT:
            return &cnt->limit;
        case RES_FAILCNT:
+ *format_mem = 0;
            return &cnt->failcnt;
    };
};
```

```
@@ -81,10 +83,26 @@ ssize_t res_counter_read(struct res_coun
```

```

{
    unsigned long *val;
    char buf[64], *s;
+ int format_mem = 1;
+ unsigned long val_kb = 0, val_mb = 0, val_gb = 0;

    s = buf;
- val = res_counter_member(cnt, member);
- s += sprintf(s, "%lu\n", *val);
+ val = res_counter_member(cnt, member, &format_mem);
+ /*
+  * NOTE: GB and MB values might be rounded off
+  */
+ if (format_mem) {
+ val_gb = (*val << PAGE_SHIFT) >> GB_SHIFT;
+ val_mb = (*val << PAGE_SHIFT) >> MB_SHIFT;
+ val_kb = (*val << PAGE_SHIFT) >> KB_SHIFT;
+ if (val_gb)
+ s += sprintf(s, "%lu (GB), %lu pages\n", val_gb, *val);
+ else if (val_mb)
+ s += sprintf(s, "%lu (MB), %lu pages\n", val_mb, *val);
+ else
+ s += sprintf(s, "%lu (KB), %lu pages\n", val_kb, *val);
+ } else
+ s += sprintf(s, "%lu\n", *val);
    return simple_read_from_buffer((void __user *)userbuf, nbytes,
        pos, buf, s - buf);
}
@@ -95,6 +113,7 @@ ssize_t res_counter_write(struct res_cou
    int ret;
    char *buf, *end;
    unsigned long tmp, *val;
+ int nop;

    buf = kmalloc(nbytes + 1, GFP_KERNEL);
    ret = -ENOMEM;
@@ -111,7 +130,7 @@ ssize_t res_counter_write(struct res_cou
    if (*end != '\0')
        goto out_free;

- val = res_counter_member(cnt, member);
+ val = res_counter_member(cnt, member, &nop);
    *val = tmp;
    ret = nbytes;
out_free:
diff -puN include/linux/res_counter.h~rss-ui-display-kdb include/linux/res_counter.h
--- linux-2.6.22-rc24-mm2/include/linux/res_counter.h~rss-ui-display-kdb 2007-06-20
20:15:46.000000000 +0530

```

+++ linux-2.6.22-rc24-mm2-balbir/include/linux/res_counter.h 2007-06-20 21:01:57.000000000
+0530
@@ -13,6 +13,12 @@

#include <linux/container.h>

```
+#ifndef KB_SHIFT
+#define KB_SHIFT 10
+#define MB_SHIFT 20
+#define GB_SHIFT 30
+#endif
+
+/*
+ * the core object. the container that wishes to account for some
+ * resource may include this counter into its structures and use
```

—

--

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: [patch -rss] Make RSS accounting display more user friendly
Posted by [Paul Menage](#) on Thu, 21 Jun 2007 19:17:27 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/20/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>

>

> Display the current usage and limit in a more user friendly manner. Number
> of pages can be confusing if the page size is different. Some systems
> can choose a page size of 64KB.

I'm not sure that's such a great idea. "Human-friendly"
representations would make programmatic parsing harder.

What's wrong with just showing counts in bytes in the control files?

Paul

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Pavel Emelianov](#) on Fri, 22 Jun 2007 02:09:19 GMT

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

Paul Menage wrote:

> On 6/20/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>>

>>

>> Display the current usage and limit in a more user friendly manner.

>> Number

>> of pages can be confusing if the page size is different. Some systems

>> can choose a page size of 64KB.

>

> I'm not sure that's such a great idea. "Human-friendly"

> representations would make programmatic parsing harder.

>

> What's wrong with just showing counts in bytes in the control files?

Nothing wrong, but currently they are shown in "natural" points, i.e. in those that the controller accounts them in. For RSS controller the natural point is "page", but auto-converting them from pages to bytes is wrong, as not all the controllers account in pages. I think that the better way is to show the value in the internal units and specify (somehow) what these units are. The userspace then can read them and convert accordingly.

> Paul

Thanks,
Pavel

Containers mailing list

Containers@lists.linux-foundation.org

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Paul Menage](#) on Fri, 22 Jun 2007 15:39:50 GMT

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

On 6/21/07, Pavel Emelianov <xemul@openvz.org> wrote:

>

> Nothing wrong, but currently they are shown in "natural" points, i.e. in

> those that the controller accounts them in. For RSS controller the natural

> point is "page", but auto-converting them from pages to bytes is wrong, as

> not all the controllers account in pages.

This exposes more implementation detail than I think is good. Something like a memory controller should use a more abstract interface like bytes, and do the conversion to/from its internal "natural" units like pages internally. E.g. the example cpu accounting controller that I included with my patch set reports usage in milliseconds, even though it counts internally in jiffies.

Paul

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Balbir Singh](#) on Sat, 23 Jun 2007 03:48:07 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/22/07, Paul Menage <menage@google.com> wrote:
> On 6/21/07, Pavel Emelianov <xemul@openvz.org> wrote:
> >
> > Nothing wrong, but currently they are shown in "natural" points, i.e. in
> > those that the controller accounts them in. For RSS controller the natural
> > point is "page", but auto-converting them from pages to bytes is wrong, as
> > not all the controllers account in pages.
>
> This exposes more implementation detail than I think is good.
> Something like a memory controller should use a more abstract
> interface like bytes, and do the conversion to/from its internal
> "natural" units like pages internally. E.g. the example cpu accounting
> controller that I included with my patch set reports usage in
> milliseconds, even though it counts internally in jiffies.
>
Hi, Paul,

Bytes are too fine grained. Do we also accept the input in bytes? kB are probably the best, but then with the limit set to LONG_MAX, kB wrap around (if the input is in units of pages) on 64 bit machines (the same applies to bytes as well).

The problem with input in bytes is that the user will have to ensure that the input is
a multiple of page size, which implies that she would need to use the calculator every time.

Containers mailing list

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Paul Menage](#) on Mon, 25 Jun 2007 07:19:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/22/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>
> The problem with input in bytes is that the user will have to ensure
> that the input is
> a multiple of page size, which implies that she would need to use the
> calculator every time.
>

Having input in bytes seems pretty natural to me. Why not just have the RSS controller round the input to the nearest page (or whatever granularity of memory the controller is able to limit at)?

Paul

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Balbir Singh](#) on Mon, 25 Jun 2007 12:47:39 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On 6/22/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
>>
>> The problem with input in bytes is that the user will have to ensure
>> that the input is
>> a multiple of page size, which implies that she would need to use the
>> calculator every time.
>>
>
> Having input in bytes seems pretty natural to me. Why not just have
> the RSS controller round the input to the nearest page (or whatever
> granularity of memory the controller is able to limit at)?
>
> Paul

Hi, Paul,

I am not a CLUI expert, but rounding off bytes will something that the administrators will probably complain about. Since we manage the controller memory in pages, it might be the easiest unit to use. The output is totally different matter.

Having said that, I am not opposed to your suggestion, I'll see if I can find good CLUI guidelines.

--

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: [patch -rss] Make RSS accounting display more user friendly
Posted by [dev](#) on Tue, 26 Jun 2007 12:39:38 GMT
[View Forum Message](#) <> [Reply to Message](#)

Paul Menage wrote:

> On 6/22/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>

>>The problem with input in bytes is that the user will have to ensure
>>that the input is

>>a multiple of page size, which implies that she would need to use the
>>calculator every time.

>>

>

>

> Having input in bytes seems pretty natural to me. Why not just have
> the RSS controller round the input to the nearest page (or whatever
> granularity of memory the controller is able to limit at)?

totally agree with Paul.

Kirill

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Balbir Singh](#) on Tue, 26 Jun 2007 12:54:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev wrote:

> Paul Menage wrote:

>> On 6/22/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:

>>

>>> The problem with input in bytes is that the user will have to ensure

>>> that the input is

>>> a multiple of page size, which implies that she would need to use the

>>> calculator every time.

>>>

>>

>> Having input in bytes seems pretty natural to me. Why not just have

>> the RSS controller round the input to the nearest page (or whatever

>> granularity of memory the controller is able to limit at)?

>

> totally agree with Paul.

>

> Kirill

>

Kirill

If someone assigns a rss_limit of 1 byte and sees a usage of 1 page,
won't that be confusing. But having said that it's not a big
change, it should be easy to accommodate.

--

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: [patch -rss] Make RSS accounting display more user friendly
Posted by [dev](#) on Tue, 26 Jun 2007 14:05:53 GMT
[View Forum Message](#) <> [Reply to Message](#)

Balbir Singh wrote:

> Kirill Korotaev wrote:

>

>> Paul Menage wrote:

>>

>>>On 6/22/07, Balbir Singh <balbir@linux.vnet.ibm.com> wrote:
>>>
>>>
>>>>The problem with input in bytes is that the user will have to ensure
>>>>that the input is
>>>>a multiple of page size, which implies that she would need to use the
>>>>calculator every time.
>>>>
>>>
>>>Having input in bytes seems pretty natural to me. Why not just have
>>>the RSS controller round the input to the nearest page (or whatever
>>>granularity of memory the controller is able to limit at)?
>>
>>totally agree with Paul.
>>
>>Kirill
>>
>
> Kirill
>
> If someone assigns a rss_limit of 1 byte and sees a usage of 1 page,
> won't that be confusing. But having said that it's not a big
> change, it should be easy to accommodate.

Well, from my experience pages are hardly understandable by people.
So bytes are always better and more convenient for non-programmers.
Rounding is not that big issue, since people still get the result
they expect (unlike to the case when they mess up with the page size).

Thanks,
Kirill

P.S. 1 byte limit is not that common :)

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Peter Zijlstra](#) on Mon, 09 Jul 2007 11:38:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Sun, 2028-02-27 at 02:39 -0500, Balbir Singh wrote:

> I am not a CLUI expert, but rounding off bytes will something that
> the administrators will probably complain about. Since we manage

> the controller memory in pages, it might be the easiest unit to use.
> The output is totally different matter.
>
> Having said that, I am not opposed to your suggestion, I'll see if
> I can find good CLUI guidelines.

Pages are generally considered a bad unit for user-space exposed parameters because a page can have a wide spectrum of sizes on some machines.

The typical unit used in its stead is KiB. Although I could imagine MiB being more useful in this case :-)

Perhaps a new proc parser that takes postfix [KMG] units would be handy..

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Balbir Singh](#) on Tue, 10 Jul 2007 04:44:31 GMT
[View Forum Message](#) <> [Reply to Message](#)

Peter Zijlstra wrote:

> On Sun, 2028-02-27 at 02:39 -0500, Balbir Singh wrote:
>
>> I am not a CLUI expert, but rounding off bytes will something that
>> the administrators will probably complain about. Since we manage
>> the controller memory in pages, it might be the easiest unit to use.
>> The output is totally different matter.
>>
>> Having said that, I am not opposed to your suggestion, I'll see if
>> I can find good CLUI guidelines.
>
> Pages are generally considered a bad unit for user-space exposed
> parameters because a page can have a wide spectrum of sizes on some
> machines.
>

Exactly!

> The typical unit used in its stead is KiB. Although I could imagine MiB
> being more useful in this case :-)
>

I think a routine that can handle either KiB or MiB would probably be the best

> Perhaps a new proc parser that takes postfix [KMG] units would be
> handy..
>

Hmm.. yes.. a library routine would be nice!

--

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: [patch -rss] Make RSS accounting display more user friendly
Posted by [rdunlap](#) on Tue, 10 Jul 2007 15:29:04 GMT
[View Forum Message](#) <> [Reply to Message](#)

On Mon, 09 Jul 2007 21:44:31 -0700 Balbir Singh wrote:

> Peter Zijlstra wrote:
> > On Sun, 2028-02-27 at 02:39 -0500, Balbir Singh wrote:
> >
> >> I am not a CLUI expert, but rounding off bytes will something that
> >> the administrators will probably complain about. Since we manage
> >> the controller memory in pages, it might be the easiest unit to use.
> >> The output is totally different matter.
> >>
> >> Having said that, I am not opposed to your suggestion, I'll see if
> >> I can find good CLUI guidelines.
> >
> > Pages are generally considered a bad unit for user-space exposed
> > parameters because a page can have a wide spectrum of sizes on some
> > machines.
> >
>
> Exactly!
>
> > The typical unit used in its stead is KiB. Although I could imagine MiB
> > being more useful in this case :-)
> >
>
> I think a routine that can handle either KiB or MiB would probably be the best

>
> > Perhaps a new proc parser that takes postfix [KMG] units would be
> > handy..
> >
>
> Hmm.. yes.. a library routine would be nice!

you mean something like lib/cmdline.c::memparse() ?

~Randy

*** Remember to use Documentation/SubmitChecklist when testing your code ***

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

Subject: Re: [patch -rss] Make RSS accounting display more user friendly
Posted by [Balbir Singh](#) on Tue, 10 Jul 2007 15:40:47 GMT
[View Forum Message](#) <> [Reply to Message](#)

Randy Dunlap wrote:

>> Hmm.. yes.. a library routine would be nice!
>
> you mean something like lib/cmdline.c::memparse() ?
>

Excellent, one of the reasons to love Linux (there's already
code) :-)

--

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>
