
Subject: Some observations from ploop testing
Posted by [jjs - mainphrame](#) on Sat, 24 Mar 2012 17:35:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

I've been creating simfs and ploop based containers and exercising them in different ways. While the ploop-based containers are basically working, in my testing a ploop-based CT seems to require more resources than an equivalent simfs-based CT. On my modest 32 bit test rig with 1 GB RAM, I've been running dbench on simfs based CTs and looking at performance with new kernel versions. But when running dbench tests on a ploop based CT with the same resources, it has not been able to finish because the machine runs out of resources, performance slows to a crawl and even host processes are killed off.

I'll try to get some more memory for this machine for further testing.

Regards,

Joe

Subject: Re: Some observations from ploop testing
Posted by [Kirill Korotaev](#) on Sat, 24 Mar 2012 18:40:02 GMT
[View Forum Message](#) <> [Reply to Message](#)

Can you please report slabtop output? We've just fixed obe memory leak. Thanks!

Sent from my iPhonespam SPAMSPAM

On 24.03.2012, at 21:57, "jjs - mainphrame" <jjs@mainphrame.com> wrote:

> I've been creating simfs and ploop based containers and exercising them in different ways. While the ploop-based containers are basically working, in my testing a ploop-based CT seems to require more resources than an equivalent simfs-based CT. On my modest 32 bit test rig with 1 GB RAM, I've been running dbench on simfs based CTs and looking at performance with new kernel versions. But when running dbench tests on a ploop based CT with the same resources, it has not been able to finish because the machine runs out of resources, performance slows to a crawl and even host processes are killed off.

>
> I'll try to get some more memory for this machine for further testing.
>
> Regards,
>
> Joe
> <ATT00001.c>

Subject: Re: Some observations from ploop testing
Posted by [jjs - mainphrame](#) on Sat, 24 Mar 2012 19:57:35 GMT
[View Forum Message](#) <> [Reply to Message](#)

I'm running slabtop every 30 seconds during a dbench run and the thing that is growing the fastest and taking the lion's share is biovec-256 - you can see it growing at 30 second intervals.

OBJS	ACTIVE	USE	OBJ SIZE	SLABS	OBJ/SLAB CACHE SIZE	NAME
88	58	65%	3.00K	44	2	352K biovec-2
134	96	71%	3.00K	67	2	536K biovec-2
152	100	65%	3.00K	76	2	608K biovec-2
112	74	66%	3.00K	56	2	448K biovec-2
140	94	67%	3.00K	70	2	560K biovec-2
74	56	75%	3.00K	37	2	296K biovec-2
144	102	70%	3.00K	72	2	576K biovec-2
114	82	71%	3.00K	57	2	456K biovec-2
154	116	75%	3.00K	77	2	616K biovec-2
80	60	75%	3.00K	40	2	320K biovec-2
164	122	74%	3.00K	82	2	656K biovec-2
152	114	75%	3.00K	76	2	608K biovec-2
70	46	65%	3.00K	35	2	280K biovec-2
1004	1004	100%	3.00K	502	2	4016K biovec-2
1952	1952	100%	3.00K	976	2	7808K biovec-2
2946	2946	100%	3.00K	1473	2	11784K biovec-2
3876	3876	100%	3.00K	1938	2	15504K biovec-2
4858	4858	100%	3.00K	2429	2	19432K biovec-2
5844	5844	100%	3.00K	2922	2	23376K biovec-2
6782	6782	100%	3.00K	3391	2	27128K biovec-2
7766	7766	100%	3.00K	3883	2	31064K biovec-2
8774	8774	100%	3.00K	4387	2	35096K biovec-2
9774	9774	100%	3.00K	4887	2	39096K biovec-2
10750	10750	100%	3.00K	5375	2	43000K biovec-2
11696	11696	100%	3.00K	5848	2	46784K biovec-2
12700	12700	100%	3.00K	6350	2	50800K biovec-2
13676	13676	100%	3.00K	6838	2	54704K biovec-2
14644	14644	100%	3.00K	7322	2	58576K biovec-2
15620	15620	100%	3.00K	7810	2	62480K biovec-2
16568	16568	100%	3.00K	8284	2	66272K biovec-2
17582	17582	100%	3.00K	8791	2	70328K biovec-2
18562	18562	100%	3.00K	9281	2	74248K biovec-2
19558	19558	100%	3.00K	9779	2	78232K biovec-2
20500	20500	100%	3.00K	10250	2	82000K biovec-2
21424	21424	100%	3.00K	10712	2	85696K biovec-2
22414	22414	100%	3.00K	11207	2	89656K biovec-2
23404	23404	100%	3.00K	11702	2	93616K biovec-2
25252	25252	100%	3.00K	12626	2	101008K biovec-2
27192	27192	100%	3.00K	13596	2	108768K biovec-2
29172	29172	100%	3.00K	14586	2	116688K biovec-2

31112	31112	100%	3.00K	15556	2	124448K	biovec-2
33006	33006	100%	3.00K	16503	2	132024K	biovec-2
34998	34926	99%	3.00K	17499	2	139992K	biovec-2
36820	36820	100%	3.00K	18410	2	147280K	biovec-2
38750	38750	100%	3.00K	19375	2	155000K	biovec-2
40480	40480	100%	3.00K	20240	2	161920K	biovec-2
42362	42362	100%	3.00K	21181	2	169448K	biovec-2
44264	44264	100%	3.00K	22132	2	177056K	biovec-2
46182	46182	100%	3.00K	23091	2	184728K	biovec-2
48058	48058	100%	3.00K	24029	2	192232K	biovec-2
49982	49974	99%	3.00K	24991	2	199928K	biovec-2
51894	51894	100%	3.00K	25947	2	207576K	biovec-2
53828	53808	99%	3.00K	26914	2	215312K	biovec-2
55596	55596	100%	3.00K	27798	2	222384K	biovec-2
57484	57484	100%	3.00K	28742	2	229936K	biovec-2
59352	59352	100%	3.00K	29676	2	237408K	biovec-2
61304	61286	99%	3.00K	30652	2	245216K	biovec-2

Joe

On Sat, Mar 24, 2012 at 11:40 AM, Kirill Korotaev <dev@parallels.com> wrote:

> Can you please report slabtop output? We've just fixed obe memory leak.
> Thanks!
>
> Sent from my iPhonespam SPAMSPAM
>
> On 24.03.2012, at 21:57, "jjs - mainphrame" <jjs@mainphrame.com> wrote:
>
>> I've been creating simfs and ploop based containers and exercising them
> in different ways. While the ploop-based containers are basically working,
> in my testing a ploop-based CT seems to require more resources than an
> equivalent simfs-based CT. On my modest 32 bit test rig with 1 GB RAM, I've
> been running dbench on simfs based CTs and looking at performance with new
> kernel versions. But when running dbench tests on a ploop based CT with the
> same resources, it has not been able to finish because the machine runs out
> of resources, performance slows to a crawl and even host processes are
> killed off.
>>
>> I'll try to get some more memory for this machine for further testing.
>>
>> Regards,
>>
>> Joe
>> <ATT00001.c>
>

Subject: Re: Some observations from ploop testing
Posted by [Kirill Korotaev](#) on Sun, 25 Mar 2012 06:22:25 GMT
[View Forum Message](#) <> [Reply to Message](#)

Exactly. It's fixed in next kernel version. Sorry for that.

Btw, when comparing performance be aware that disk performance depends a lot on placement.
Disk beginning is typically around 2x faster.

Sent from my iPhonespam SPAMSPAM

On 25.03.2012, at 0:21, "jjs - mainphrame"
<jjs@mainphrame.com<<mailto:jjs@mainphrame.com>>> wrote:

I'm running slabtop every 30 seconds during a dbench run and the thing that is growing the fastest and taking the lion's share is biovec-256 - you can see it growing at 30 second intervals.

OBJS	ACTIVE	USE	OBJ SIZE	SLABS	OBJ/SLAB	CACHE SIZE	NAME
88	58	65%	3.00K	44	2	352K	biovec-2
134	96	71%	3.00K	67	2	536K	biovec-2
152	100	65%	3.00K	76	2	608K	biovec-2
112	74	66%	3.00K	56	2	448K	biovec-2
140	94	67%	3.00K	70	2	560K	biovec-2
74	56	75%	3.00K	37	2	296K	biovec-2
144	102	70%	3.00K	72	2	576K	biovec-2
114	82	71%	3.00K	57	2	456K	biovec-2
154	116	75%	3.00K	77	2	616K	biovec-2
80	60	75%	3.00K	40	2	320K	biovec-2
164	122	74%	3.00K	82	2	656K	biovec-2
152	114	75%	3.00K	76	2	608K	biovec-2
70	46	65%	3.00K	35	2	280K	biovec-2
1004	1004	100%	3.00K	502	2	4016K	biovec-2
1952	1952	100%	3.00K	976	2	7808K	biovec-2
2946	2946	100%	3.00K	1473	2	11784K	biovec-2
3876	3876	100%	3.00K	1938	2	15504K	biovec-2
4858	4858	100%	3.00K	2429	2	19432K	biovec-2
5844	5844	100%	3.00K	2922	2	23376K	biovec-2
6782	6782	100%	3.00K	3391	2	27128K	biovec-2
7766	7766	100%	3.00K	3883	2	31064K	biovec-2
8774	8774	100%	3.00K	4387	2	35096K	biovec-2
9774	9774	100%	3.00K	4887	2	39096K	biovec-2
10750	10750	100%	3.00K	5375	2	43000K	biovec-2
11696	11696	100%	3.00K	5848	2	46784K	biovec-2
12700	12700	100%	3.00K	6350	2	50800K	biovec-2
13676	13676	100%	3.00K	6838	2	54704K	biovec-2
14644	14644	100%	3.00K	7322	2	58576K	biovec-2
15620	15620	100%	3.00K	7810	2	62480K	biovec-2
16568	16568	100%	3.00K	8284	2	66272K	biovec-2
17582	17582	100%	3.00K	8791	2	70328K	biovec-2

18562	18562	100%	3.00K	9281	2	74248K	biovec-2
19558	19558	100%	3.00K	9779	2	78232K	biovec-2
20500	20500	100%	3.00K	10250	2	82000K	biovec-2
21424	21424	100%	3.00K	10712	2	85696K	biovec-2
22414	22414	100%	3.00K	11207	2	89656K	biovec-2
23404	23404	100%	3.00K	11702	2	93616K	biovec-2
25252	25252	100%	3.00K	12626	2	101008K	biovec-2
27192	27192	100%	3.00K	13596	2	108768K	biovec-2
29172	29172	100%	3.00K	14586	2	116688K	biovec-2
31112	31112	100%	3.00K	15556	2	124448K	biovec-2
33006	33006	100%	3.00K	16503	2	132024K	biovec-2
34998	34926	99%	3.00K	17499	2	139992K	biovec-2
36820	36820	100%	3.00K	18410	2	147280K	biovec-2
38750	38750	100%	3.00K	19375	2	155000K	biovec-2
40480	40480	100%	3.00K	20240	2	161920K	biovec-2
42362	42362	100%	3.00K	21181	2	169448K	biovec-2
44264	44264	100%	3.00K	22132	2	177056K	biovec-2
46182	46182	100%	3.00K	23091	2	184728K	biovec-2
48058	48058	100%	3.00K	24029	2	192232K	biovec-2
49982	49974	99%	3.00K	24991	2	199928K	biovec-2
51894	51894	100%	3.00K	25947	2	207576K	biovec-2
53828	53808	99%	3.00K	26914	2	215312K	biovec-2
55596	55596	100%	3.00K	27798	2	222384K	biovec-2
57484	57484	100%	3.00K	28742	2	229936K	biovec-2
59352	59352	100%	3.00K	29676	2	237408K	biovec-2
61304	61286	99%	3.00K	30652	2	245216K	biovec-2

Joe

On Sat, Mar 24, 2012 at 11:40 AM, Kirill Korotaev <dev@parallels.com<<mailto:dev@parallels.com>>> wrote:
 Can you please report slabtop output? We've just fixed obe memory leak. Thanks!

Sent from my iPhonespam SPAMSPAM

On 24.03.2012, at 21:57, "jjs - mainphrame" <jjs@mainphrame.com<<mailto:jjs@mainphrame.com>>> wrote:

> I've been creating simfs and ploop based containers and exercising them in different ways.
 While the ploop-based containers are basically working, in my testing a ploop-based CT seems to require more resources than an equivalent simfs-based CT. On my modest 32 bit test rig with 1 GB RAM, I've been running dbench on simfs based CTs and looking at performance with new kernel versions. But when running dbench tests on a ploop based CT with the same resources, it has not been able to finish because the machine runs out of resources, performance slows to a crawl and even host processes are killed off.

>
 > I'll try to get some more memory for this machine for further testing.
 >

> Regards,
>
> Joe
> <ATT00001.c>

Users mailing list
Users@openvz.org<mailto:Users@openvz.org>
<https://openvz.org/mailman/listinfo/users>

<ATT00001.c>

Subject: Re: Some observations from ploop testing
Posted by [jjs - mainphrame](#) on Sun, 25 Mar 2012 07:16:36 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for the update and the supplementary information!

Joe

On Sat, Mar 24, 2012 at 11:22 PM, Kirill Korotaev <dev@parallels.com> wrote:

> Exactly. It's fixed in next kernel version. Sorry for that.
>
> Btw, when comparing performance be aware that disk performance depends a
> lot on placement. Disk beginning is typically around 2x faster.
>
> Sent from my iPhonespam SPAMSPAM
>
> On 25.03.2012, at 0:21, "jjs - mainphrame" <jjs@mainphrame.com> wrote:
>
> I'm running slabtop every 30 seconds during a dbench run and the thing
> that is growing the fastest and taking the lion's share is biovec-256 - you
> can see it growing at 30 second intervals.
>
> OBJS ACTIVE USE OBJ SIZE SLABS OBJ/SLAB CACHE SIZE NAME
> 88 58 65% 3.00K 44 2 352K biovec-2
> 134 96 71% 3.00K 67 2 536K biovec-2
> 152 100 65% 3.00K 76 2 608K biovec-2
> 112 74 66% 3.00K 56 2 448K biovec-2
> 140 94 67% 3.00K 70 2 560K biovec-2
> 74 56 75% 3.00K 37 2 296K biovec-2
> 144 102 70% 3.00K 72 2 576K biovec-2
> 114 82 71% 3.00K 57 2 456K biovec-2
> 154 116 75% 3.00K 77 2 616K biovec-2
> 80 60 75% 3.00K 40 2 320K biovec-2
> 164 122 74% 3.00K 82 2 656K biovec-2

```

> 152 114 75% 3.00K 76 2 608K biovec-2
> 70 46 65% 3.00K 35 2 280K biovec-2
> 1004 1004 100% 3.00K 502 2 4016K biovec-2
> 1952 1952 100% 3.00K 976 2 7808K biovec-2
> 2946 2946 100% 3.00K 1473 2 11784K biovec-2
> 3876 3876 100% 3.00K 1938 2 15504K biovec-2
> 4858 4858 100% 3.00K 2429 2 19432K biovec-2
> 5844 5844 100% 3.00K 2922 2 23376K biovec-2
> 6782 6782 100% 3.00K 3391 2 27128K biovec-2
> 7766 7766 100% 3.00K 3883 2 31064K biovec-2
> 8774 8774 100% 3.00K 4387 2 35096K biovec-2
> 9774 9774 100% 3.00K 4887 2 39096K biovec-2
> 10750 10750 100% 3.00K 5375 2 43000K biovec-2
> 11696 11696 100% 3.00K 5848 2 46784K biovec-2
> 12700 12700 100% 3.00K 6350 2 50800K biovec-2
> 13676 13676 100% 3.00K 6838 2 54704K biovec-2
> 14644 14644 100% 3.00K 7322 2 58576K biovec-2
> 15620 15620 100% 3.00K 7810 2 62480K biovec-2
> 16568 16568 100% 3.00K 8284 2 66272K biovec-2
> 17582 17582 100% 3.00K 8791 2 70328K biovec-2
> 18562 18562 100% 3.00K 9281 2 74248K biovec-2
> 19558 19558 100% 3.00K 9779 2 78232K biovec-2
> 20500 20500 100% 3.00K 10250 2 82000K biovec-2
> 21424 21424 100% 3.00K 10712 2 85696K biovec-2
> 22414 22414 100% 3.00K 11207 2 89656K biovec-2
> 23404 23404 100% 3.00K 11702 2 93616K biovec-2
> 25252 25252 100% 3.00K 12626 2 101008K biovec-2
> 27192 27192 100% 3.00K 13596 2 108768K biovec-2
> 29172 29172 100% 3.00K 14586 2 116688K biovec-2
> 31112 31112 100% 3.00K 15556 2 124448K biovec-2
> 33006 33006 100% 3.00K 16503 2 132024K biovec-2
> 34998 34926 99% 3.00K 17499 2 139992K biovec-2
> 36820 36820 100% 3.00K 18410 2 147280K biovec-2
> 38750 38750 100% 3.00K 19375 2 155000K biovec-2
> 40480 40480 100% 3.00K 20240 2 161920K biovec-2
> 42362 42362 100% 3.00K 21181 2 169448K biovec-2
> 44264 44264 100% 3.00K 22132 2 177056K biovec-2
> 46182 46182 100% 3.00K 23091 2 184728K biovec-2
> 48058 48058 100% 3.00K 24029 2 192232K biovec-2
> 49982 49974 99% 3.00K 24991 2 199928K biovec-2
> 51894 51894 100% 3.00K 25947 2 207576K biovec-2
> 53828 53808 99% 3.00K 26914 2 215312K biovec-2
> 55596 55596 100% 3.00K 27798 2 222384K biovec-2
> 57484 57484 100% 3.00K 28742 2 229936K biovec-2
> 59352 59352 100% 3.00K 29676 2 237408K biovec-2
> 61304 61286 99% 3.00K 30652 2 245216K biovec-2
>
> Joe

```

>
> On Sat, Mar 24, 2012 at 11:40 AM, Kirill Korotaev <dev@parallels.com>wrote:
>
>> Can you please report slabtop output? We've just fixed obe memory leak.
>> Thanks!
>>
>> Sent from my iPhonespam SPAMSPAM
>>
>> On 24.03.2012, at 21:57, "jjs - mainphrame" <jjs@mainphrame.com> wrote:
>>
>>> I've been creating simfs and ploop based containers and exercising them
>> in different ways. While the ploop-based containers are basically working,
>> in my testing a ploop-based CT seems to require more resources than an
>> equivalent simfs-based CT. On my modest 32 bit test rig with 1 GB RAM, I've
>> been running dbench on simfs based CTs and looking at performance with new
>> kernel versions. But when running dbench tests on a ploop based CT with the
>> same resources, it has not been able to finish because the machine runs out
>> of resources, performance slows to a crawl and even host processes are
>> killed off.
>>>
>>> I'll try to get some more memory for this machine for further testing.
>>>
>>> Regards,
>>>
>>> Joe
>>> <ATT00001.c>
>>
> <ATT00001.c>
>
>
