

I know this is has probably been discussed ad nauseum, but I haven't found what I'm looking for yet, so I thought I would ask here.

I have been running OpenVZ for a few years, but in the last couple of weeks, I have noticed over the past couple of weeks that several VMs were getting out of spec settings, mainly dcachesize growing too large.

These VMs started on a Debian openvz box, and as my virtual infrastructure grew, I started using a pair of proxmox-ve machines (which is Debian-lenny based as well), which are clustered.

I have 8 VMs that were created over time, some on 32-bit host machines, some on 64-bit. Thus, some have /proc/user_beancounters that look like:

1:	kmemsize	13775736		
15028224	48811846	51254098		63446
	lockedpages	0		
447	393216	393216	0	
	privvmpages	15152		
105895	426752	439252		0
	shmpages	648		
1304	21504	21504	0	
	dummy	0		
0	0	0	0	
	numproc	47		
72	240	240	0	
	physpages	166345		
425143	0	2147483647		0
	vmguarpages	0		
0	426752	2147483647		0
	oomguarpages	6374		
97683	426752	2147483647		0
	numtcpsock	44		
48	360	360	0	
	numflock	1		
7	188	206	0	
	numpty	0		
2	16	16	0	
	numsiginfo	1		
27	256	256	0	
	tcpsndbuf	525744		
1026352	4212558	6014798		0
	tcprrcvbuf	524552		
3052984	4212558	6014798		0

	othersockbuf	46240		
65808	1126080	2097152		0
	dgramrcvbuf	0		
101600	262144	262144		0
	numothersock	75		
82	360	360		0
	dcachesize	9997638		
10000000	8000000	10000000		0
	numfile	508		
695	9312	9312		0
	dummy	0		
0	0	0		0
	dummy	0		
0	0	0		0
	dummy	0		
0	0	0		0
	numiptent	20		
20	128	128		0

While others have effectively unlimited barrier and limit settings:

	7: kmemsize	93292551	107253760	
9223372036854775807	9223372036854775807			0
	lockedpages	0		
16	393216	393216		0
	privvmpages	299033		
413214	524288	536788		0
	shmpages	68	724	
9223372036854775807	9223372036854775807			0
	dummy	0		
0	0	0		0
	numproc	86		
108	1024	1024		0
	physpages	321589		
496217	0	9223372036854775807		0
	vmguarpages	0		
0	524288	9223372036854775807		0
	oomguarpages	155305		
180405	524288	9223372036854775807		0
	numtcpsock	13	17	
9223372036854775807	9223372036854775807			0
	numflock	3	9	
9223372036854775807	9223372036854775807			0
	numpty	0		
2	255	255		0
	numsiginfo	1		
15	1024	1024		0
	tcpsndbuf	226720	329312	

9223372036854775807	9223372036854775807	0
tcprcvbuf	277072	5662864
9223372036854775807	9223372036854775807	0
othersockbuf	43928	66680
9223372036854775807	9223372036854775807	0
dgramrcvbuf	0	5648
9223372036854775807	9223372036854775807	0
numothersock	63	69
9223372036854775807	9223372036854775807	0
dcachesize	88045648	101016538
9223372036854775807	9223372036854775807	0
numfile	360	605
9223372036854775807	9223372036854775807	0
dummy	0	
0	0	0
dummy	0	
0	0	0
dummy	0	
0	0	0
numiptent	20	20
9223372036854775807	9223372036854775807	0

I have three questions. First, I know that leaving everything unlimited is a quick path to running out of resources on the host machine. That said, I've been having troubles recently with the VMs with "normal" settings. It started out with dcachesize going out of spec, which, when I adjusted it, within an hour, I started getting out of memory errors, requiring me to up the kmemsize...This then caused problems on another "normal" VM, and so forth.

As I said, I know setting everything to unlimited is probably not recommended, so what is the recommended way to set the proper values for user_beancounters? Every time I change values in user_beancounters, something else comes unglued, except for the ones that have unlimited kmemsize and dcachesize.

Is there a tool to set up the values based on the use of the particular VM?
Is there any more information I need to provide?

Thanks,
--b

Subject: Re: Optimizing resources from /proc/user_beancounters
Posted by [JR Richardson](#) on Fri, 14 Oct 2011 21:35:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

> I know this is has probably been discussed ad nauseum, but I haven't found

> what I'm looking for yet, so I thought I would ask here.
>
> I have been running OpenVZ for a few years, but in the last couple of weeks,
> I have noticed over the past couple of weeks that several VMs were getting
> out of spec settings, mainly dcache size growing too large.
>
> These VMs started on a Debian openvz box, and as my virtual infrastructure
> grew, I started using a pair of proxmox-ve machines (which is Debian-lenny
> based as well), which are clustered.
>
> I have 8 VMs that were created over time, some on 32-bit host machines, some
> on 64-bit. Thus, some have /proc/user_beancounters that look like:

```

> 1: kmemsize          13775736
> 15028224          48811846          51254098          63446
>   lockedpages              0
> 447          393216          393216          0
>   privvmpages          15152
> 105895          426752          439252          0
>   shmpages              648
> 1304          21504          21504          0
>   dummy              0
> 0          0          0          0
>   numproc              47
> 72          240          240          0
>   physpages          166345
> 425143          0          2147483647          0
>   vmguarpages              0
> 0          426752          2147483647          0
>   oomguarpages          6374
> 97683          426752          2147483647          0
>   numtcpsock          44
> 48          360          360          0
>   numflock              1
> 7          188          206          0
>   numpty              0
> 2          16          16          0
>   numsiginfo          1
> 27          256          256          0
>   tcpsndbuf          525744
> 1026352          4212558          6014798          0
>   tcprcvbuf          524552
> 3052984          4212558          6014798          0
>   othersockbuf          46240
> 65808          1126080          2097152          0
>   dgramrcvbuf          0
> 101600          262144          262144          0
>   numothersock          75

```

```

> 82          360          360          0
>          dcache size      9997638
> 10000000      8000000      10000000          0
>          numfile          508
> 695          9312          9312          0
>          dummy          0
> 0          0          0          0
>          dummy          0
> 0          0          0          0
>          dummy          0
> 0          0          0          0
>          numiptent        20
> 20          128          128          0
>
> While others have effectively unlimited barrier and limit settings:
>
> 7: kmemsize      93292551      107253760
> 9223372036854775807 9223372036854775807          0
>          lockedpages      0
> 16          393216      393216          0
>          privvmpages      299033
> 413214          524288      536788          0
>          shmpages          68          724
> 9223372036854775807 9223372036854775807          0
>          dummy          0
> 0          0          0          0
>          numproc          86
> 108          1024          1024          0
>          physpages      321589
> 496217          0 9223372036854775807          0
>          vmguarpages      0
> 0          524288 9223372036854775807          0
>          oomguarpages      155305
> 180405          524288 9223372036854775807          0
>          numtcpsock      13          17
> 9223372036854775807 9223372036854775807          0
>          numflock          3          9
> 9223372036854775807 9223372036854775807          0
>          numpty          0
> 2          255          255          0
>          numsiginfo      1
> 15          1024          1024          0
>          tcpsndbuf      226720      329312
> 9223372036854775807 9223372036854775807          0
>          tcprcvbuf      277072      5662864
> 9223372036854775807 9223372036854775807          0
>          othersockbuf      43928      66680
> 9223372036854775807 9223372036854775807          0

```

```

> dgramrcvbuf 0 5648
> 9223372036854775807 9223372036854775807 0
> numothersock 63 69
> 9223372036854775807 9223372036854775807 0
> dcachesize 88045648 101016538
> 9223372036854775807 9223372036854775807 0
> numfile 360 605
> 9223372036854775807 9223372036854775807 0
> dummy 0
> 0 0 0 0
> dummy 0
> 0 0 0 0
> dummy 0
> 0 0 0 0
> numiptent 20 20
> 9223372036854775807 9223372036854775807 0
>
> I have three questions. First, I know that leaving everything unlimited is a
> quick path to running out of resources on the host machine. That said, I've
> been having troubles recently with the VMs with "normal" settings. It
> started out with dcachesize going out of spec, which, when I adjusted it,
> within an hour, I started getting out of memory errors, requiring me to up
> the kmemsize...This then caused problems on another "normal" VM, and so
> forth.
>
> As I said, I know setting everything to unlimited is probably not
> recommended, so what is the recommended way to set the proper values for
> user_beancounters? Every time I change values in user_beancounters,
> something else comes unglued, except for the ones that have unlimited
> kmemsize and dcachesize.
>
> Is there a tool to set up the values based on the use of the particular VM?
> Is there any more information I need to provide?
>

```

> Thanks,
 > --b

Try using vzsplint to segment your VE's equally, start there and increase/decrease resources per the demand of each VE. Once you adjust your config conf files, use vzcfgvalidate to ensure your beans are adjusted properly.

Good luck.

JR

--

JR Richardson
 Engineering for the Masses

Subject: Re: Re: Optimizing resources from /proc/user_beancounters
Posted by [Brad Alexander](#) on Sat, 15 Oct 2011 14:11:28 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thank you, JR. As it turns out, I was *severely* starving my VMs. vzsplrit worked like a charm...

--b

On Fri, Oct 14, 2011 at 5:35 PM, JR Richardson <jmr.richardson@gmail.com>wrote:

> > I know this is has probably been discussed ad nauseum, but I haven't
> found
> > what I'm looking for yet, so I thought I would ask here.
> >
> > I have been running OpenVZ for a few years, but in the last couple of
> weeks,
> > I have noticed over the past couple of weeks that several VMs were
> getting
> > out of spec settings, mainly dcachesize growing too large.
> >
> > These VMs started on a Debian openvz box, and as my virtual
> infrastructure
> > grew, I started using a pair of proxmox-ve machines (which is
> Debian-lenny
> > based as well), which are clustered.
> >
> > I have 8 VMs that were created over time, some on 32-bit host machines,
> some
> > on 64-bit. Thus, some have /proc/user_beancounters that look like:

```
> >
> > 1: kmemsize          13775736
> > 15028224          48811846          51254098          63446
> >   lockedpages              0
> > 447          393216          393216              0
> >   privvmpages          15152
> > 105895          426752          439252              0
> >   shmpages              648
> > 1304          21504          21504              0
> >   dummy              0
> > 0              0              0              0
> >   numproc              47
> > 72          240          240              0
> >   physpages          166345
> > 425143          0          2147483647              0
> >   vmguarpages              0
> > 0          426752          2147483647              0
> >   oomguarpages          6374
> > 97683          426752          2147483647              0
```

```

>> numtcpsock 44
>> 48 360 360 0
>> numflock 1
>> 7 188 206 0
>> numpty 0
>> 2 16 16 0
>> numsiginfo 1
>> 27 256 256 0
>> tcpsndbuf 525744
>> 1026352 4212558 6014798 0
>> tcprcvbuf 524552
>> 3052984 4212558 6014798 0
>> othersockbuf 46240
>> 65808 1126080 2097152 0
>> dgramrcvbuf 0
>> 101600 262144 262144 0
>> numothersock 75
>> 82 360 360 0
>> dcachesize 9997638
>> 10000000 8000000 10000000 0
>> numfile 508
>> 695 9312 9312 0
>> dummy 0
>> 0 0 0 0
>> dummy 0
>> 0 0 0 0
>> dummy 0
>> 0 0 0 0
>> numiptent 20
>> 20 128 128 0
>>
>> While others have effectively unlimited barrier and limit settings:
>>
>> 7: kmemsize 93292551 107253760
>> 9223372036854775807 9223372036854775807 0
>> lockedpages 0
>> 16 393216 393216 0
>> privvmpages 299033
>> 413214 524288 536788 0
>> shmpages 68 724
>> 9223372036854775807 9223372036854775807 0
>> dummy 0
>> 0 0 0 0
>> numproc 86
>> 108 1024 1024 0
>> physpages 321589
>> 496217 0 9223372036854775807 0
>> vmguarpages 0

```



```

>> 0          524288 9223372036854775807          0
>>          oomguarpages          155305
>> 180405          524288 9223372036854775807          0
>>          numtcpsock          13          17
>> 9223372036854775807 9223372036854775807          0
>>          numflock          3          9
>> 9223372036854775807 9223372036854775807          0
>>          numpty          0
>> 2          255          255          0
>>          numsiginfo          1
>> 15          1024          1024          0
>>          tcpsndbuf          226720          329312
>> 9223372036854775807 9223372036854775807          0
>>          tcprcvbuf          277072          5662864
>> 9223372036854775807 9223372036854775807          0
>>          othersockbuf          43928          66680
>> 9223372036854775807 9223372036854775807          0
>>          dgramrcvbuf          0          5648
>> 9223372036854775807 9223372036854775807          0
>>          numothersock          63          69
>> 9223372036854775807 9223372036854775807          0
>>          dcachesize          88045648          101016538
>> 9223372036854775807 9223372036854775807          0
>>          numfile          360          605
>> 9223372036854775807 9223372036854775807          0
>>          dummy          0
>> 0          0          0          0
>>          dummy          0
>> 0          0          0          0
>>          dummy          0
>> 0          0          0          0
>>          numiptent          20          20
>> 9223372036854775807 9223372036854775807          0
>>
>> I have three questions. First, I know that leaving everything unlimited
> is a
>> quick path to running out of resources on the host machine. That said,
> I've
>> been having troubles recently with the VMs with "normal" settings. It
>> started out with dcachesize going out of spec, which, when I adjusted it,
>> within an hour, I started getting out of memory errors, requiring me to
> up
>> the kmemsize...This then caused problems on another "normal" VM, and so
>> forth.
>>
>> As I said, I know setting everything to unlimited is probably not
>> recommended, so what is the recommended way to set the proper values for
>> user_beancounters? Every time I change values in user_beancounters,

```

> > something else comes unglued, except for the ones that have unlimited
> > kmemsize and dcachesize.
> >
> > Is there a tool to set up the values based on the use of the particular
> VM?
> > Is there any more information I need to provide?
> >
> > Thanks,
> > --b
> Try using vzsplitt to segment your VE's equally, start there and
> increase/decrease resources per the demand of each VE. Once you
> adjust your config conf files, use vzcfgvalidate to ensure your beans
> are adjusted properly.
>
> Good luck.
>
> JR
> --
> JR Richardson
> Engineering for the Masses
>
