Subject: VE memory misconfigured / misreported
Posted by csigman on Wed, 20 Aug 2008 15:38:43 GMT
View Forum Message <> Reply to Message

I've just created a VE running debian with apache, tomcat, mysql, and a custom java server app. On a regular, non-virtual server, the whole thing uses about a half gig of ram at most, however when I've got it running in the VE it is at least reported that it's using a whole lot more than that (about 2 gigs), even when idling, and thus I've tweaked my config to give it a lot more ram (at least inside the VE). Part of my issue stems from a gap in understanding as to what all of the memory configuration options actually do, but none the less, I'd really like to get a config, if possible, that shows a real and reliable amount of memory usage. Below is my current config:

KMEMSIZE="119516908:179249920"
LOCKEDPAGES="256:256"
PRIVVMPAGES="14939613:22406240"
SHMPAGES="31504:31504"
NUMPROC="240:240"
PHYSPAGES="0:2147483647"
VMGUARPAGES="415932:2147483647"
OOMGUARPAGES="831864:2147483647"
NUMTCPSOCK="360:360"
NUMFLOCK="188:206"
NUMPTY="16:16"
NUMSIGINFO="256:256"
TCPSNDBUF="1720320:2703360"
TCPRCVBUF="1720320:2703360"
OTHERSOCKBUF="1126080:2097152"
DGRAMRCVBUF="262144:262144"
NUMOTHERSOCK="360:360"
DCACHESIZE="3409920:3624960"
NUMFILE="9312:9312"
AVNUMPROC="180:180"
NUMIPTENT="128:128"

There's a few changes that are still in there from when I was trying to figure it all out, so just be aware if you see something crazy, that it might not be because I'm retarded with the config (even though I kind-of am).

Subject: Re: VE memory misconfigured / misreported
Posted by piavlo on Wed, 20 Aug 2008 22:44:50 GMT
View Forum Message <> Reply to Message

AFAIK you should be looking at oomguarpages
(oomguarpages parameter accounts the total amount of memory and swap space used by the processes of a particular container. The barrier of the oomguarpages parameter is the out-of-memory guarantee) in /proc/user_beancounters at "held" column to get initial estimation.
More details at http://wiki.openvz.org/UBC_systemwide_configuration

In particular at http://wiki.openvz.org/UBC_systemwide_configuration#Utilizat ion_2 if you are concerned with RAM utilization

---

Subject: Re: VE memory misconfigured / misreported
Posted by csigman on Thu, 21 Aug 2008 14:26:03 GMT
View Forum Message <> Reply to Message

So, I did a cat on user_beancounters, and I don't think that oomguarpages is really doing me much. Here's a before and after I adjusted using information from there:

Before:

uid	resource	held	maxheld	barrier	limit	failcnt
109:	kmemsize	5761486	7302244	119516908	179249920	0
	lockedpages	0	0	256	256	0
	privvmpages	455772	693191	504140	554554	51
	shmpages	1422	1438	31504	31504	0
	dummy	0	0	0	0	0
	numproc	122	130	240	240	0
	physpages	48396	57609	0	2147483647	0
	vmguarpages	0	0	415932	2147483647	0
	oomguarpages	48396	57609	831864	2147483647	0
	numtcpsock	24	38	360	360	0
	numflock	12	62	188	206	0
	numpty	2	2	16	16	0
	numsiginfo	0	10	256	256	0
	tcpsndbuf	213792	411356	1720320	2703360	0
	tcprcvbuf	206684	799604	1720320	2703360	0
	othersockbuf	19980	30136	1126080	2097152	0
	dgramrcvbuf	0	8364	262144	262144	0
	numothersock	14	20	360	360	0
	dcachesize	0	0	3409920	3624960	0
	numfile	2503	2625	9312	9312	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	10	10	128	128	0

After:

uid	resource	held	maxheld	barrier	limit	failcnt
109:	kmemsize	4761437	5333291	9776841	119516908	0
	lockedpages	0	0	256	256	0
	privvmpages	487310	756408	919491	14939613	0
	shmpages	142	142	31504	31504	0
	dummy	0	0	0	0	0
	numproc	108	126	240	240	0
	physpages	42568	55944	0	2147483647	0
	vmguarpages	0	0	415932	2147483647	0

oomguarpages	42568	55944	2128630	2147483647	0
numtcpsock	16	20	360	360	0
numflock	12	13	188	206	0
numpty	1	2	16	16	0
numsiginfo	0	6	256	256	0
tcpsndbuf	142080	0	1720320	2703360	0
tcprcvbuf	214188	215384	1720320	2703360	0
othersockbuf	17760	40548	1126080	2097152	0
dgramrcvbuf	0	8364	262144	262144	0
numothersock	12	23	360	360	0
dcachesize	0	0	3409920	3624960	0
numfile	2214	2525	9312	9312	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	10	10	128	128	0

I feel that the maxheld are misleading though, as htop shows max memory usage of almost 3gigs of the 3.5 allocated. The other very confusing thing is the resource usage on the host machine is definitely not nearly that much (1.5 gigs of ram used with 6 other ve's running)

Subject: Re: VE memory misconfigured / misreported
Posted by piavlo on Thu, 21 Aug 2008 15:25:36 GMT
View Forum Message <> Reply to Message

So your VE current RAM usage is 171M since
(physpages*4096+kmemsize+allsocketbuf)/1024/1024=
=42568*4096+4761437+(142080+214188+17760+0)/1024/1024=
=179493993/1024/1024=171
Also since oomguarpages equals physpages (since no VE pages are swaped out by the kernel) the RAM+SWAP=RAM in your case.

Now total current allocated “virtual” (but not used) memory for VE is 1908MB(privvmpages*4096+kmemsize+allsocketbuf)/1024/1024=
=487310*4096+4761437+(142080+214188+17760+0)/1024/1024=
2001157225/1024/1024=1908MB

This huge difference is since your VE apps ask for too much memory but don't actually use it. For example i took a look at my VE with nginx webserver and php-cgi and it's RAM usage is 55M while “virtual” memory is 92M. So look in tool like htop to find which of the apps is asking for too much memory but not using it.

It's all in http://wiki.openvz.org/UBC_systemwide_configuration

---

Subject: Re: VE memory misconfigured / misreported

Posted by csigman on Thu, 21 Aug 2008 16:11:37 GMT
View Forum Message <> Reply to Message

So, I guess the problem is that the swap and actual ram are shared from the ve's perspective, and java apps might allocate a huge amount of ram, but generally this is mostly in swap. I've toned down the memory usage of some of the java parts and that makes it so that it's using a lot less, but still not the most ideal setup. Luckily this is development, and production isn't in a VE.

Subject: Re: VE memory misconfigured / misreported
Posted by piavlo on Thu, 21 Aug 2008 16:36:59 GMT
View Forum Message <> Reply to Message

csigman wrote on Thu, 21 August 2008 19:11So, I guess the problem is that the swap and actual ram are shared from the ve's perspective, and java apps might allocate a huge amount of ram, but generally this is mostly in swap.
 Apps allocate vitrual memory which is neither in ram nor is swap. Only the real used memory is in ram or swap, and VE does not know where it is (it only sees "virtual" ram), but this does not matter.
Quote:I've toned down the memory usage of some of the java parts and that makes it so that it's using a lot less, but still not the most ideal setup. Luckily this is development, and production isn't in a VE. I think you are missing the point, the app virtual memory allocation does not depend if you are using OpenVZ or Vanilla kernel. Currently your VE is using 171M from RAM+SWAP on HN or equally 171M of "virtual" ram on VE, so it is still safe for production.

Subject: Re: VE memory misconfigured / misreported
Posted by csigman on Thu, 21 Aug 2008 22:02:53 GMT
View Forum Message <> Reply to Message

Thanks for the explanation, that helps a lot!

Subject: Re: VE memory misconfigured / misreported
Posted by mifritscher on Thu, 02 Oct 2008 11:42:31 GMT
View Forum Message <> Reply to Message

The problem seems to be the shared memory: Each process using it has a big virtual memory, but this memory is only alloc one time. On openvz, this is counted shared_memory_size*accessing_processes.

This happens particulary with cachers like xcache,