Subject: Re: mem limit does not work... ?
Posted by disaster on Sun, 19 Feb 2006 10:26:10 GMT
View Forum Message <> Reply to Message

Quote:
1. can you please provide /proc/user_beancounters and a script?...

No Problem - here are the details:

user_beancounters before script:

cat /proc/user_beancounters
Version: 2.5

uid	resource	held	maxheld	barrier	limit	failcnt
4:	kmemsize	434432	477994	13107200	13107200	0
	lockedpages	0	0	32	32	0
	privvmpages	3915	5868	32768	32768	0
	shmpages	640	656	8192	8192	0
	dummy	0	0	0	0	0
	numproc	7	8	65	65	0
	physpages	945	945	32768	32768	0
	vmguarpages	0	0	32768	32768	0
	oomguarpages	945	945	32768	32768	0
	numtcpsock	2	2	80	80	0
	numflock	1	2	100	110	0
	numpty	1	1	16	16	0
	numsiginfo	0	1	256	256	0
	tcpsndbuf	2224	8896	319488	524288	0
	tcprcvbuf	0	4272	319488	524288	0
	othersockbuf	2224	7200	132096	336896	0
	dgramrcvbuf	0	4272	132096	132096	0
	numothersock	1	5	80	80	0
	dcachesize	91840	94136	1048576	1097728	0
	numfile	86	90	2048	2048	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	dummy	0	0	0	0	0
	numiptent	4	4	128	128	0

My Script - you don't like... and this time with top output

```
perl -e '$a="a";for(;;) {$a.=$a; print "allocated ".int(length($a)/1024/1024)."MB\n";}'
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
```

allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 0MB
allocated 1MB
allocated 2MB
allocated 4MB
allocated 8MB
allocated 16MB
allocated 32MB
allocated 64MB
allocated 128MB
allocated 256MB
allocated 512MB
allocated 1024MB
Out of memory!

Top Output

Mem: 2071904k total, 1067136k used, 1004768k free,   10272k buffers
Swap: 1437776k total,   0k used, 1437776k free,   46148k cached

  PID USER   PR NI VIRT RES SHR S %CPU %MEM   TIME+ Command

 4544 root   25   0 1027m 878m 2576 R 94.0 43.4   0:09.74 perl

As you can see, there is a total of 1GB RAM and perl eats 43% which is about 800MB the res 200MB are used from normal system.

user_beancounters after script start

cat /proc/user_beancounters
Version: 2.5
       uid resource   held   maxheld   barrier   limit   failcnt
         4: kmemsize   434432   1485574   13107200   13107200   0
            lockedpages   0   0   32   32   0

privvmpages	1621	263957	32768	32768	10
shmpages	640	656	8192	8192	0
dummy	0	0	0	0	0
numproc	7	8	65	65	0
physpages	946	263318	32768	32768	0
vmguarpages	0	0	32768	32768	0
oomguarpages	946	263318	32768	32768	0
numtcpsock	2	2	80	80	0
numflock	1	2	100	110	0
numpty	1	1	16	16	0
numsiginfo	0	1	256	256	0
tcpsndbuf	2224	8896	319488	524288	0
tcprcvbuf	0	4272	319488	524288	0
othersockbuf	2224	7200	132096	336896	0
dgramrcvbuf	0	4272	132096	132096	0
numothersock	1	5	80	80	0
dcachesize	91840	94136	1048576	1097728	0
numfile	86	90	2048	2048	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
dummy	0	0	0	0	0
numiptent	4	4	128	128	0

Quote:
2. failct=4 and messages like:
Fatal resource shortage: kmemsize, UB 4.
mean that you hit kmemsize limit 4 times.

Sorry but failcnt is at privvmpages not at kmemsize (at kmemsize only in my example before...).

Quote:
3. messages like:
Uncharging too much 1 h 0, res unused_privvmpages ub 4
mean that there is some discrepancy in privvmpages accounting.
This is why I would be happy to your script which reproduces this problem.

No problem - you have the small one line script above...

I would be happy if you can help me - if you need anything - write it down - i'll do that.

Thanks!