
Subject: Re: Linux Memory Overcommit in OpenVZ
Posted by [charlesl](#) on Sat, 12 Jul 2008 13:59:50 GMT
[View Forum Message](#) <> [Reply to Message](#)

It has been nearly a month, almost 1000 views and not a single reply so I guess this has everyone else stumped too (or its so obvious it didn't deserve an answer) ...

I wrote a couple of programs to test and I'll share my results. The two programs were:

memalloc - Allocates memory in 100MB/second increments

memfill - Allocates and fills memory at rate 100MB/second

I then ran these in a VE and watched its user_beancounters from the HN, and the HN's /proc/meminfo too. Here are my findings:

memalloc

The privvmpages accounting increases until it hits the privvmpages barrier, at which failcnt increases and the program begins to fail.

But /proc/meminfo does not change - the kernel is not counting any allocations made via malloc.

memfill

The privvmpages accounting and physpages/oomguarpages accounting all increase. Since this is a test system not under load, the privvmpages barrier is reached first, its failcnt increases and the program fails.

In this case, /proc/meminfo does change - the kernel is counting used pages.

Conclusion

Note: naturally, "allocated" memory includes "used" memory. An allocation limit is implicitly a usable limit too. I use `_bar` for barrier and `_cur` for current (held) values.

The VE has a maximum allocation limit of `privvmpages_bar`.

The VE has a guaranteed allocation limit of `guarvmpages_bar`.

The kernel only cares about used memory - the sum of all VE's `physpages_cur/oomguarpages_cur` (+ HN overhead).

VEs can allocate all the memory they like up to `privvmpages_bar` under situations where the total used memory is low.

VEs using memory above their guarantees may fail above the guarantees under high overall memory consumption.

My only one remaining question is: clearly VEs cannot use arbitrary amounts of memory under high load conditions, but can VEs still allocate up to their `privvmpages_bar` even under high memory conditions, since the kernel doesn't care about allocations?
