
Subject: Recommendations for setting overcommit_memory

Posted by [mustardman](#) on Thu, 30 Jan 2014 19:35:49 GMT

[View Forum Message](#) <> [Reply to Message](#)

Just discovered I have a problem with oom killer and overcommit_memory setting. I wasn't aware this could be a problem but now that I have run out of memory it apparently is.

I thought I had plenty of memory because although physical ram is all used up, Swap is not and there is no heavy swap usage. This is not an overselling situation trust me. No more VPS's are being added to this node.

However I seem to be right on the edge and I just realized that oom killer has been running around deciding who lives and dies. I thought by default Linux was smart enough to know it still has plenty of swap to not do that but apparently not.

Turns out overcommit_memory...at least on CE6/OVZ 2.6.32 defaults to 0 which takes a heuristic approach that seems to be getting it wrong.

Long story short, what is the recommended setting for overcommit_memory to prevent oom killer from shutting down processes when there is still plenty of swap? Right now I am testing setting it at 2. If I understand RH's explanation correctly, combined with the overcommit_ratio default of 50, it will not use oom killer until I get to SWAP + 50% RAM which is quite a bit higher than where I am at now with oom killer initiating. That should prevent this problem right?

Subject: Re: Recommendations for setting overcommit_memory

Posted by [mustardman](#) on Fri, 31 Jan 2014 23:51:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

Ok so I guess it wasn't the Node having a problem but a container on the node. I concluded that because the first oom-killer message is oom-killer in ub xxx. Where xxx is apparently the container ID and it's always the same one. Seems to happen regularly on the hour and the container isn't even close to hitting their memory limits. Not sure what is going on now.

I am seeing this on several nodes. Always just one container.
