
Subject: Slow Live Migration (prlctl migrate)
Posted by [alenco](#) on Thu, 14 Mar 2019 19:37:45 GMT
[View Forum Message](#) <> [Reply to Message](#)

Between 2 OpenVZ 7 servers got 1000mbit / 1Gbit connections, which performs fine for data transfer, and offline migration. But when i do a live migration it goes back to ~180bmit when the "Live migration stage started". Before that, in the state " Copying static data ", full 1000mbit is being used.

Is there something wrong, or is this normal behavior. Don't see any issues with CPU or OI being to bottleneck here. No high loads, and both SSD servers.

Would like to figure this out, because network is being upgraded to 10Gbit, so would like to use this speed offcourse.

Thanks in advance.

Subject: Re: Slow Live Migration (prlctl migrate)
Posted by [khorenko](#) on Fri, 15 Mar 2019 10:30:23 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi,

well, some decrease in network bandwidth utilization is expected:

- on "Copying static data" stage we send a lot of data - disk, memory, big portions => full bandwidth utilization

- on "Live migration stage" we scan disk and memory dirty masks to get data which has been changed while we were coping data on the previous stage.

Once "dirty" data is found, it's sent to the destination => so generally it is lot of small sends and thus bandwidth utilization is less.

Note1: there are many iterations on this stage.

Note2: processes are still running during this stage, so this time is not included into "downtime".

Is it ok if the speed of second stage is only ~180bmit?

Well, i'm not ready to answer this without deep investigation, but as this time is not in the "downtime", hope it's not that important.

And what's your actual whole migration time and downtime? (just approximately, downtime may be measured via "ping CT").

And what size of CT disk and RAM, it does matter surely.

Subject: Re: Slow Live Migration (prlctl migrate)
Posted by [alenco](#) on Wed, 20 Mar 2019 19:19:18 GMT
[View Forum Message](#) <> [Reply to Message](#)

Thanks for your reply, really much appreciated.

Well for example, i just tried to migrate a 25GB disk CT, with 8GB RAM. The copying static data stops after 5,5GB. Then starts the Live Migration stage. Always at 18MegaByte per second. Not slower, not faster. Like its a cap or something.

Looking at the server, there is at least 20GB of disk space in use (90% is static content like images). You would expect to get more data trough the first stage. Looking at the private folder on source node, its 27GB.

So if the slow speed happens when working with dirty data, i have no issue with it, but when i have to move 1TB of CT than it takes a very long time.

Is there a way to do export a debug log on this, that i can read and check myself for logical explanations.

This happens for all servers that i try on whatever node, if its small or big. All nodes are upgraded to latest versions of Virtuozzo 7 and kernel.

Even when i do this on 2 nodes, that does not have any load on it at all, it does not make a difference. Both nodes, when at the live migration stage, are not using any io / memory or cpu at all. So no bottleneck in there.

Thanks in advance.