Subject: Slow reading speed of ploops
Posted by Juves on Sun, 27 Aug 2017 23:55:15 GMT
View Forum Message <> Reply to Message

In the ploop the reading speed dramatically slow down.
In 2-4 times, depended with vhd space.
Containers don't have IO limitations.
How to resolve it?

Subject: Re: Slow reading speed of ploops
Posted by khorenko on Mon, 28 Aug 2017 08:29:22 GMT
View Forum Message <> Reply to Message

Too many undefined variables:
- what is the ploop state before experiment?
  * it may be fresh - so you read the data that was never written - and it requires additional actions to allocate blocks (ploop - is an expandable format!)
  * it may be old - then it may be significantly fragmented and random io is slower than the sequential one
- do you measure on the same node?
- what is the system? Kernel version?

Just ad hoc ideas.

Anyway - reading data which was not written by you (test) in advance - is a bad idea for testing.

Suggest following sequence:
xfs_io -c "pwrite 0 1G -S 0xa -b 1M" -df filename
xfs_io -c "pread 0 1G -b 1M" -df filename

Subject: Re: Slow reading speed of ploops
Posted by Juves on Mon, 28 Aug 2017 18:01:03 GMT
View Forum Message <> Reply to Message

ploops are near month old.
Used latest OpenVZ7 OS v7.0.5. Ext4 by default. All placed on one SSD NVMe, same server, node, etc. but different sizes.

Is any same tool for Ext4 for tests?

Subject: Re: Slow reading speed of ploops
Posted by khorenko on Mon, 28 Aug 2017 20:03:01 GMT
View Forum Message <> Reply to Message

xfs_io is a generic tool, it can be used for any fs checking.