
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
```
