



# **Ceph Benchmarking: An update on the vision of CBT**

## **Ceph Day London 2025**

Lee Sanders  
ljsanders@uk.ibm.com

# Vision for CBT



- Single Click -Standardized Test Methodology – Block/Object/File- comparable results, with comparison tools.  
**Demonstrate improvements, easily spot regressions. – Integration with Ceph-CI/Teuthology**
- Whole ceph community to use this as the standard way of evaluating performance.
- 30+ Hockey stick curves with a defined test order, wide variety of customer like workloads.
- Regression test a **build/fix** in around 12 hours. Or a **3 hour** quick regression test.
- Still maintain developer focus - optional enablement of graphs with CPU utilization, code profiling features.
- **Upload reports to a Github repo for everyone to see. - Useful for customer sizings**
- **Crowd funding performance data from the community with wide variability in configs**

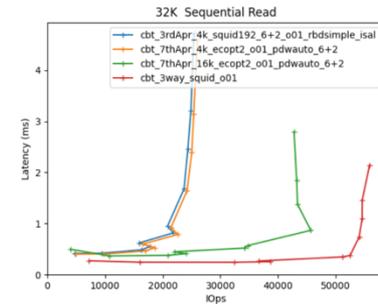
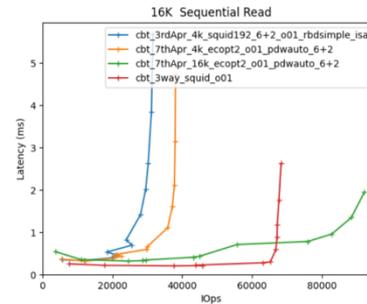
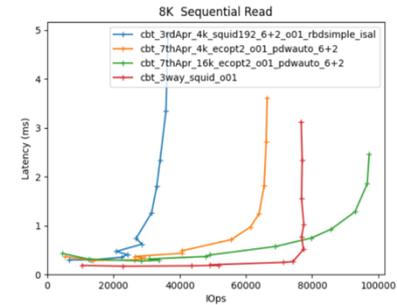
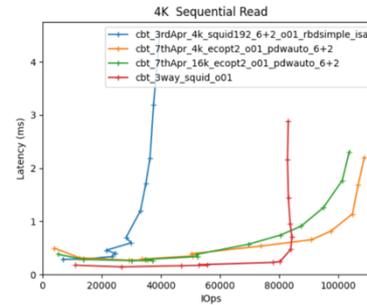


# Improvement: CBT Report Generation



## Response Curves

### Sequential Read



- Individual results reports
- Comparison reports of multiple results
- Intermediate format between IO exerciser (eg. FIO) and report generation. Avoids needing to re-write report generator for each exerciser
- Reports generated in PDF and Markdown format



# Improvement: CBT Report Generation



- N-way table comparison of best throughput in curve, included in report

Sequential Write	cbt_10thMar_4k_main_8vol_cacheon_rbdoff_isal	cbt_7thMar_4k_ls26_8vol_cacheon_rbdoff_isal	%change
<a href="#">4K</a>	<a href="#">3433@149.7ms</a>	7503@136.5	119%
<a href="#">8K</a>	<a href="#">2008@512.0ms</a>	5297@242.2	164%
<a href="#">16K</a>	<a href="#">5099@100.3ms</a>	5406@94.7	6%
<a href="#">32K</a>	<a href="#">9593@80.0ms</a>	9773@78.5	2%
<a href="#">64K</a>	<a href="#">215@156.5ms</a>	233@144.3	8%
<a href="#">256K</a>	<a href="#">400@62.8ms</a>	456@73.5	14%
<a href="#">512K</a>	<a href="#">292@229.8ms</a>	305@220.4	4%
<a href="#">1024K</a>	<a href="#">469@357.8ms</a>	484@347.3	3%

Random Read	cbt_10thMar_4k_main_8vol_cacheon_rbdoff_isal	cbt_7thMar_4k_ls26_8vol_cacheon_rbdoff_isal	%change
<a href="#">4K</a>	<a href="#">69659@5.5ms</a>	69552@5.5	-0%
<a href="#">8K</a>	<a href="#">47136@8.1ms</a>	46343@8.3	-2%
<a href="#">16K</a>	<a href="#">29788@12.9ms</a>	29079@13.2	-2%
<a href="#">32K</a>	<a href="#">26545@9.6ms</a>	26213@9.8	-1%
<a href="#">64K</a>	<a href="#">1502@11.2ms</a>	1420@5.9	-5%
<a href="#">256K</a>	<a href="#">1708@19.6ms</a>	1708@9.8	0%
<a href="#">512K</a>	<a href="#">1775@14.2ms</a>	1775@14.2	0%
<a href="#">1024K</a>	<a href="#">1771@16.6ms</a>	1772@16.6	0%



# Improvement: Workloads section of YAML



- **Workloads** provides ability to create a standardized/deterministic test methodology.
- Inclusion of **workloads** type from librbdfio to all benchmark types.
- iodepth (per volume) and total\_iodepth
- New “script” syntax to execute script between each test. Useful for measuring empty vs populated capacity.

```
workloads:
  precondition:
    jobname: 'precond1rw'
    mode: 'randwrite'
    time: 600
    op_size: 65536
    numjobs: [ 1 ]
    total_iodepth: [ 16 ]
    monitor: False # whether to run the monitors along the test

  seq32kwrite:
    jobname: 'seqwrite'
    mode: 'write'
    op_size: 32768
    numjobs: [ 1 ]
    total_iodepth: [ 2, 4, 8, 16, 32, 64, 128, 256, 512, 768 ]

  seq32kread:
    jobname: 'seqread'
    mode: 'read'
    op_size: 32768
    numjobs: [ 1 ]
    total_iodepth: [ 2, 4, 8, 12, 16, 24, 32, 64, 96, 128, 192 ]
```



# In Progress: Cephadm config generator



- Generate configuration details for inclusion into report.
- Details of configuration, ceph commit version, number of OSDs, volume and size, Erasure Coding config/Replica's, drive sizes etc.





- Vision: Single click. Currently multi step:
  - Create configuration
  - Performance run `cbt.py`
  - Process results `fio_common_output_wrapper.py`  
`generate_performance_report.py`  
`generate_comparison_performance_report.py`
- Focused on Block. **Next:** Will be doing NVME GW (Raw) and Object
- Interested to hear from the community on opinions of Object exerciser tools.  
(eg. Warp vs Elbencho vs HSBench)



# Futures – Common repository



- Report in PDF and Markdown format - ideal for github
- Example here:  
<https://github.com/lee-j-sanders/perfresults>





- Performance (throughput/latency) over time graphs to highlight variance min/max
- Histograms
- Include CPU utilization/memory utilizations in the reports to see how system utilization is affected by increase in throughput.





Interested in hearing your feedback

If you'd like to try out CBT and get involved  
please contact us!

Slack: #cbt in [ceph-storage.slack.com](https://ceph-storage.slack.com)

Github: <https://github.com/ceph/cbt>

Ceph Performance Weekly @ 8am PST - 4pm London Time

<https://meet.google.com/uhb-cysu-nvg>

Lee Sanders

[lsanders@uk.ibm.com](mailto:lsanders@uk.ibm.com)

Jose Juan Palacios Perez

[perezjos@uk.ibm.com](mailto:perezjos@uk.ibm.com)

Chris Harris

[harrisr@uk.ibm.com](mailto:harrisr@uk.ibm.com)