

# TS7700 R 5.0.1 Compression and Performance

Don Denning  
Joe Swingler

May 12, 2020

# Overview

- ❖ TS7700 R 5.0 major offerings:
  - ❖ Additional flexibility in capacity options.
  - ❖ A more organic disk subsystem providing additional synergy with IBM offerings delivering more integrated service and product lifecycle to our users.
  - ❖ IBM POWER9 server technology to help support performance, greater compression capabilities and grid network traffic improvements

# Background

- ❖ With release R 5.0, TS7700 IBM is encouraging customers to utilize software based compression – LZ4 and zSTD
- ❖ TS7700 has supported a form of ALDC compression in the FICON adapters since its first release
- ❖ The FICON adapter compression is an older algorithm that produces lower than average compression results
- ❖ FICON adapter compression is still supported

# Performance Benchmark Data Patterns

- ❖ We utilize two data patterns for TS7700 performance benchmark workloads:
  - RECS4 (near average compression results) – for general testing.
  - A00 -- for 1:1 cache-stressed testing in standalone configuration
- For customer workloads that don't compress, it's best to disable compaction via the z/OS DATACLASS so that the overhead of attempting compression can be avoided.
  - Previously compressed data
  - Encrypted data
  - LZ4 and ZSTD will detect expansion and store data at 1:1, but at a cost.

# TS7700 Performance Workloads and Measurements

## TS770 Performance Workloads and Measurements

- ❖ Unless otherwise stated, all runs were made with 128 concurrent jobs. Each job wrote/read a volume (2GB after compression at the TS7700), using 32768 byte block size, QSAM BUFFNO=20, and eight 16Gb FICON channels from a z13 LPAR.
- ❖ Clusters are located at zero or nearly zero distance to each other in laboratory setup.
- ❖ All runs were made using tuning values: DCT=125, PMPIOR=5600, PMTHLVL=6000, ICOPYT=ENABLED, Reclaim=disabled, LINKSPEED=1000. Number of premigration drives per pool=10.

# TS770 Performance Metrics


- ❖ **Peak Write Rate/Peak Channel:** host rate at which data can be written into the DASD cache with no required physical tape drive activity.
- ❖ **Sustained Write Rate/Sustained Channel:** rate at which data can be written into the DASD cache with equivalent concurrent destaging to physical tape drives.
- ❖ **PremigRate During Sustained:** Average rate at which data is written to physical tape drives while the host is actively writing to the TS7770 and the PMPRIOR threshold has been crossed.
- ❖ **Peak Premig Post Host Workload:** Average rate at which data is written to the physical tape drives after the host activity has been stopped. This shows the maximum rate at which data can be offloaded to the tape drives in the TS7770.
- ❖ **Grid Deferred Copy Mode:** copy to other Hydra does not have to complete prior to job end (DD copy mode).
- ❖ **Grid Immediate Copy Mode:** copy to other Hydra must complete prior to job end (RR copy mode).
- ❖ **Grid Sync Copy Mode:** tape synchronization up to SYNC level granularity across two clusters within a grid (SS copy mode).
- ❖ **Grid Copy Rate:** rate at which data is copied to other Hydra.

# Standalone TS7700 Performance


(Using A00 and RECS4 Data Patterns)




Standalone TS7770T Performance  
 10 Drawer Config  
 ZSTD (1:1 and 5.35:1)




Standalone TS7770T Performance  
 10 Drawer Config  
 LZ4 (1:1 and 3.43:1)




Standalone TS7770T<sub>cp1</sub> Performance vs Drawer Count  
ZSTD-RECS4 / 5.35:1



Standalone TS7770T<sub>cp1</sub> Performance vs Drawer Count  
LZ4-RECS4 / 3.44:1



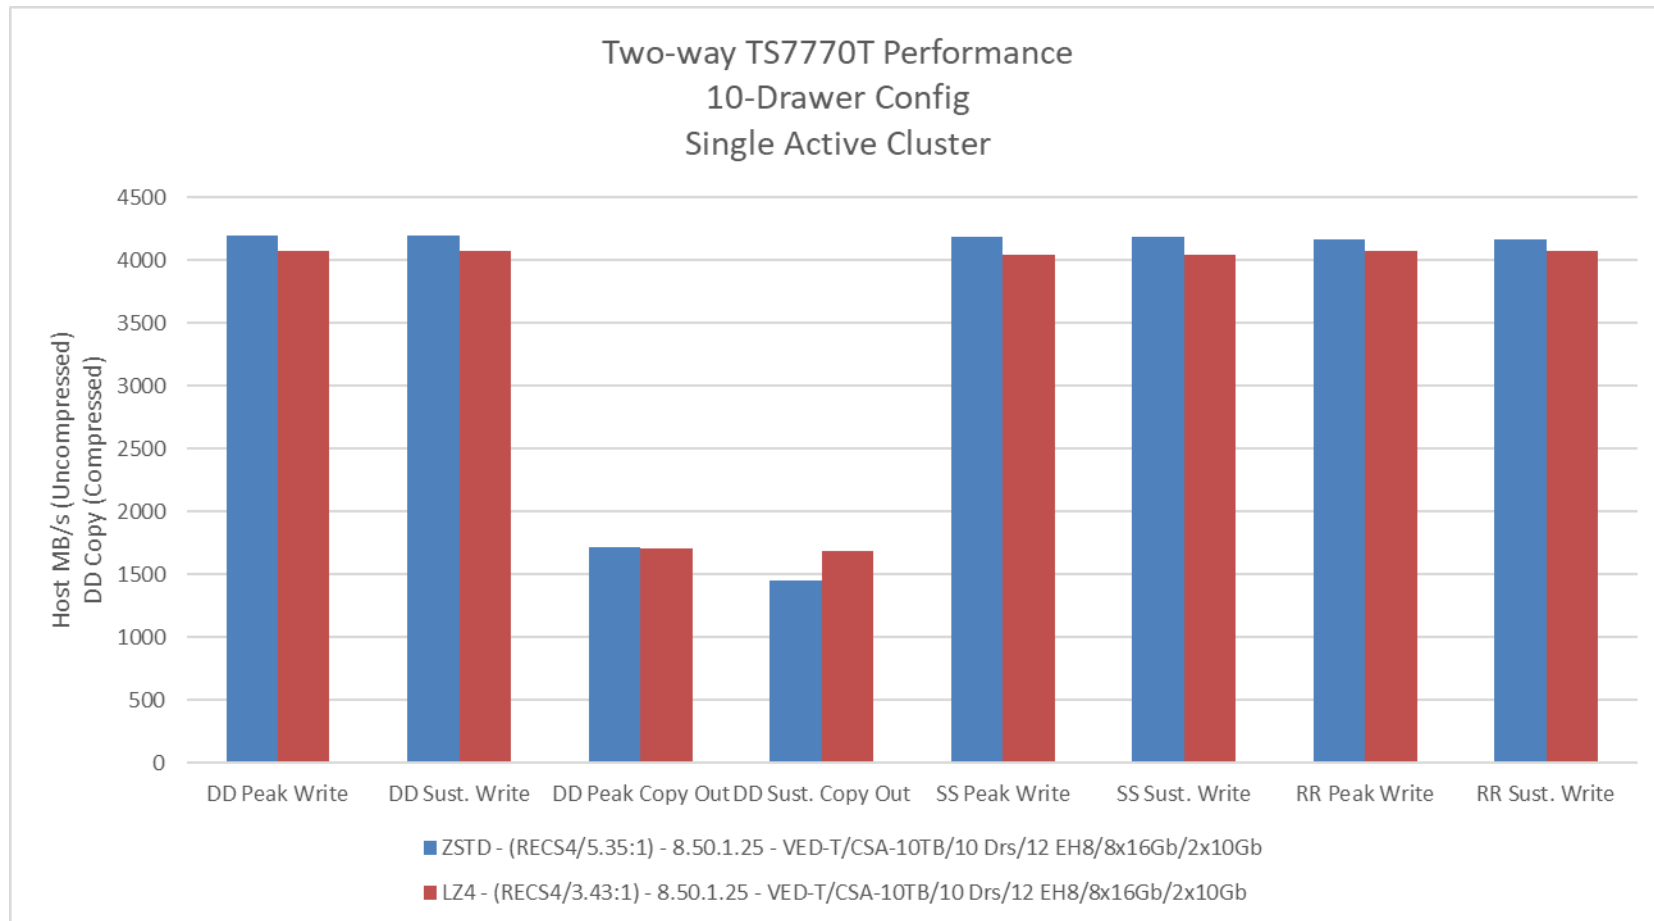
Standalone TS7770<sub>cp1</sub> Performance vs Drawer Count  
LZ4-A00 / 1:1




# TS7700 Grid Performance

## (8 x 16Gb FICON)

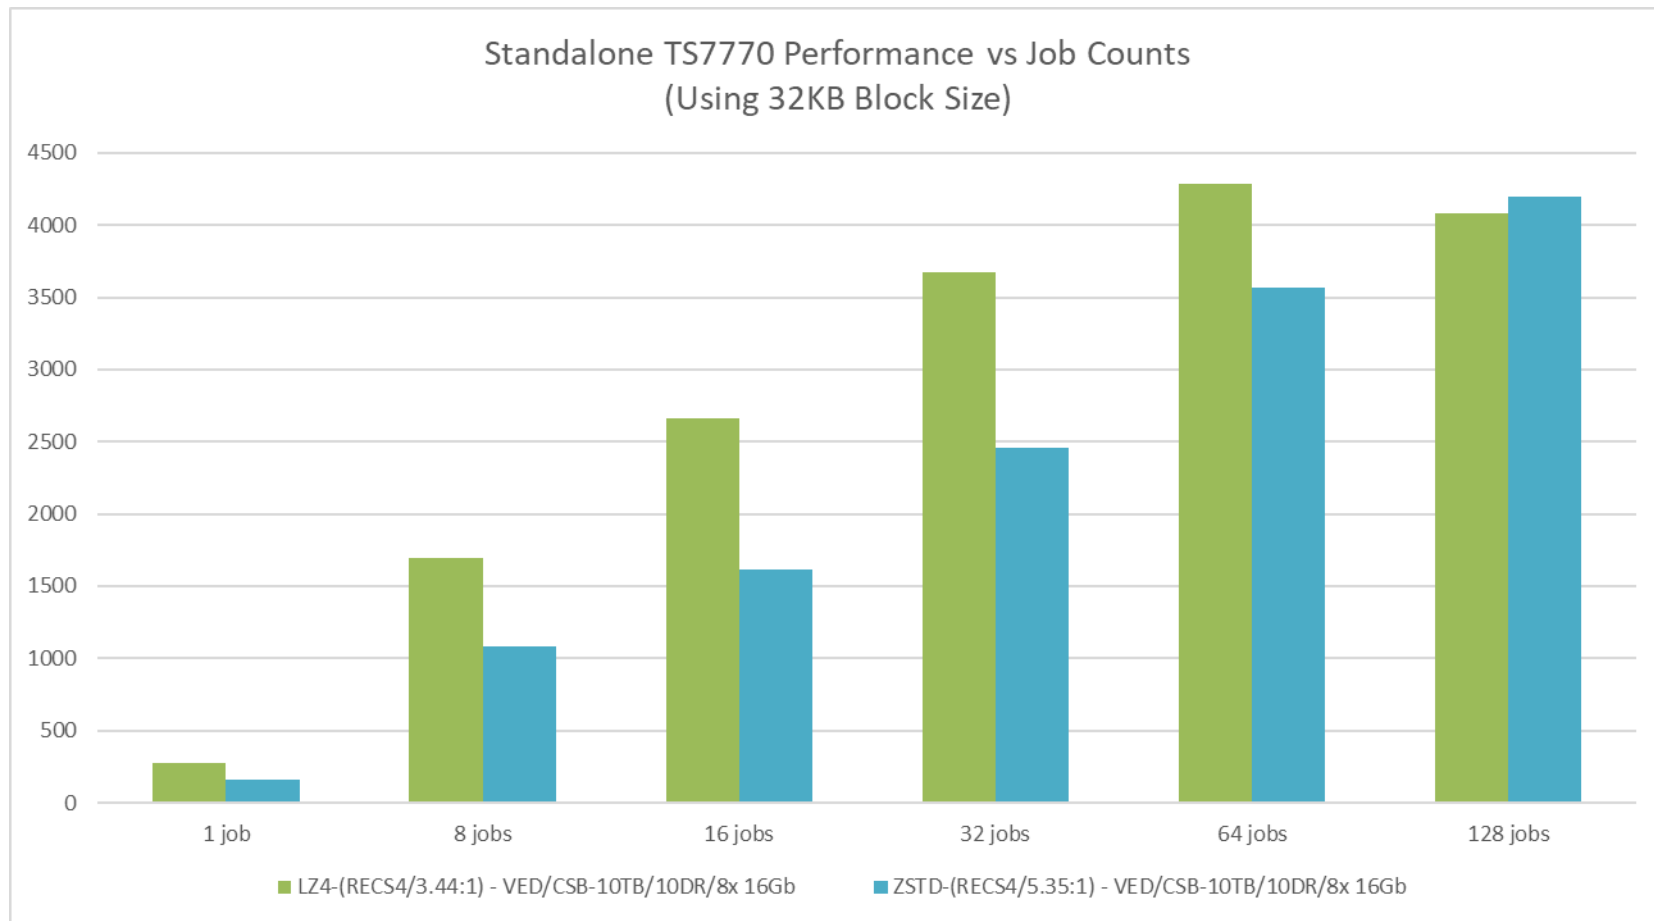
(Using RECS4 Data Patterns)








Standalone TS7770 Write Performance vs. Concurrent Job Counts  
(Using RECS4 Data Patterns)



Standalone TS7770  
Performance vs. Concurrent Job Counts and Block Sizes  
(Using RECS4 Data Patterns)



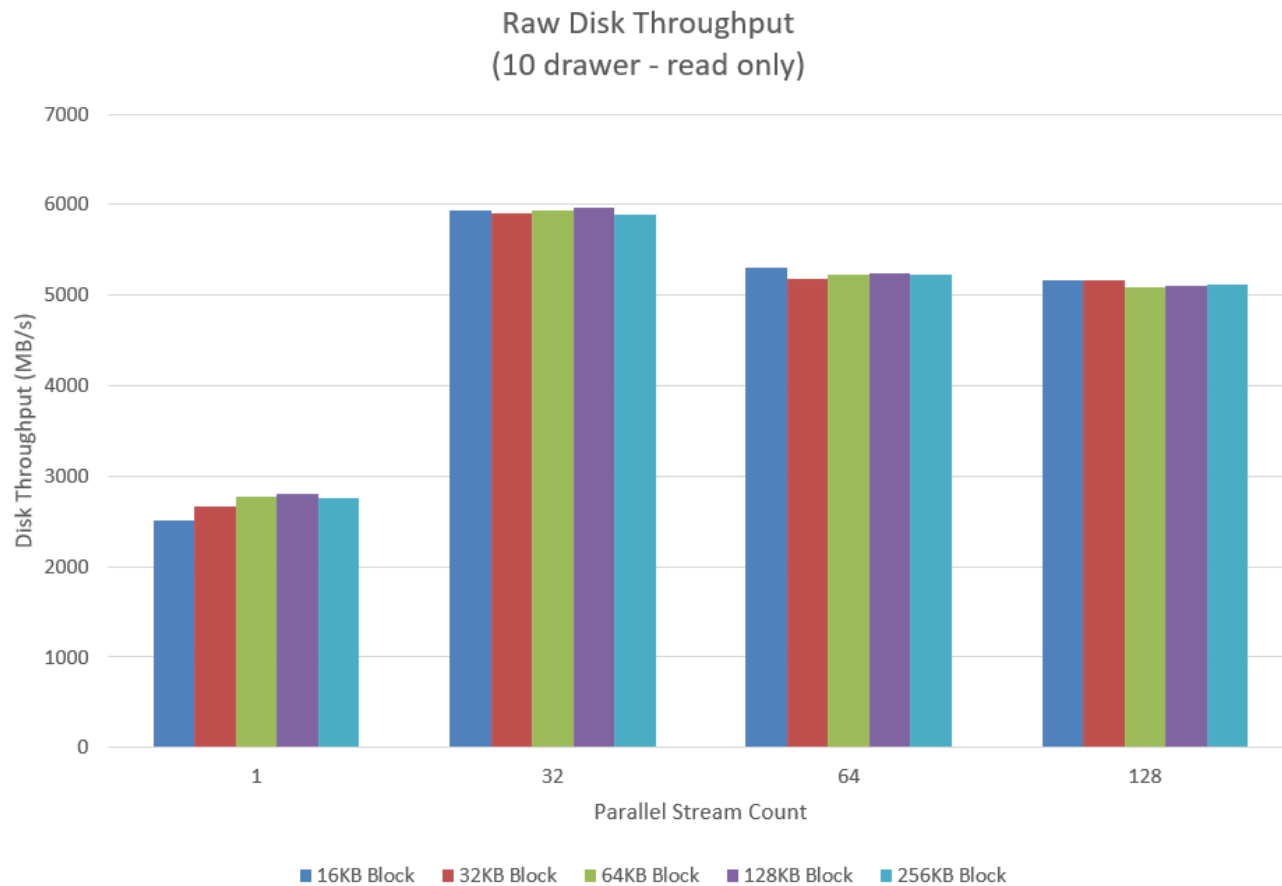
## Raw Disk Throughput Metrics (New for 5.0pga1)

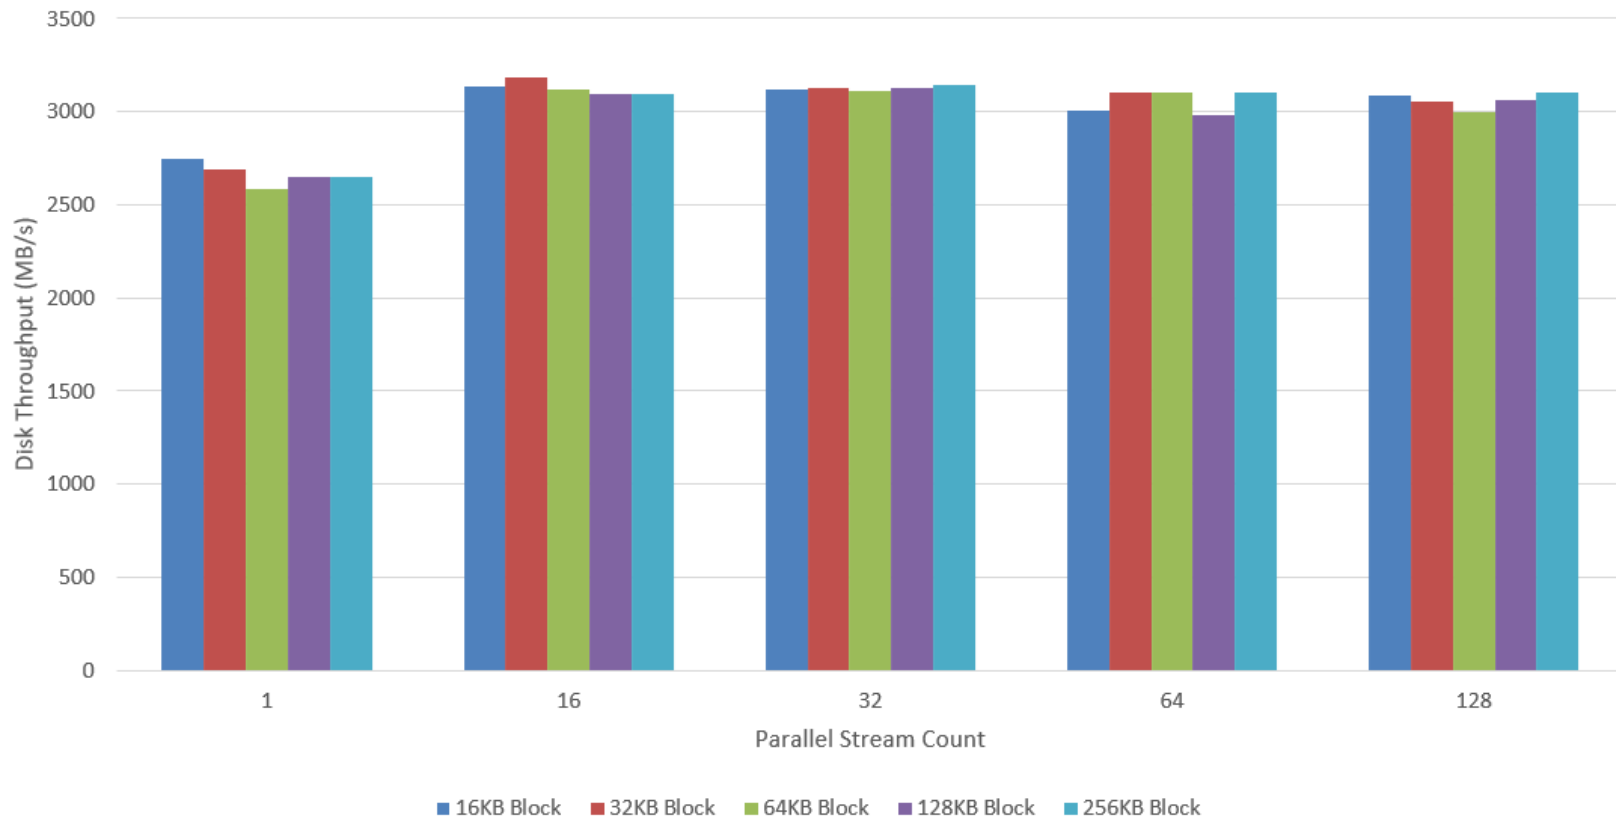
This section provides additional information of low level disk read and write throughput metrics.

Disk cache I/O was driven using an internal driver while no other host or TS7700 activity was occurring.


Test cases consist of reads, writes and a mixture of reads and writes with different block sizes and stream counts.

Results are provided for reference only and not a representation of resulting host rates, copy rates or premigration rates.

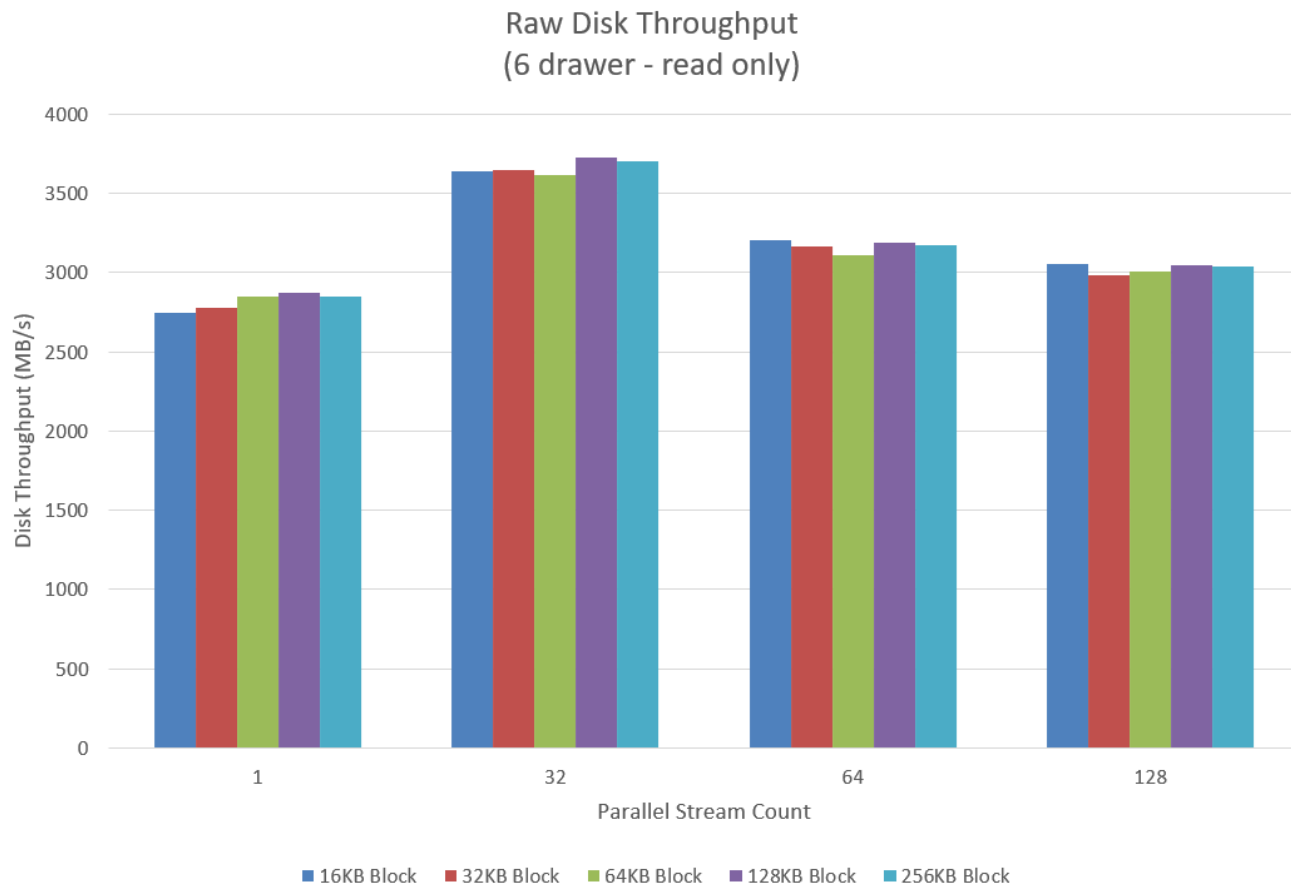



Raw Disk Throughput  
(10 drawer - write only)

## Raw Disk Throughput (10 drawer - read + write combinations at different block sizes)




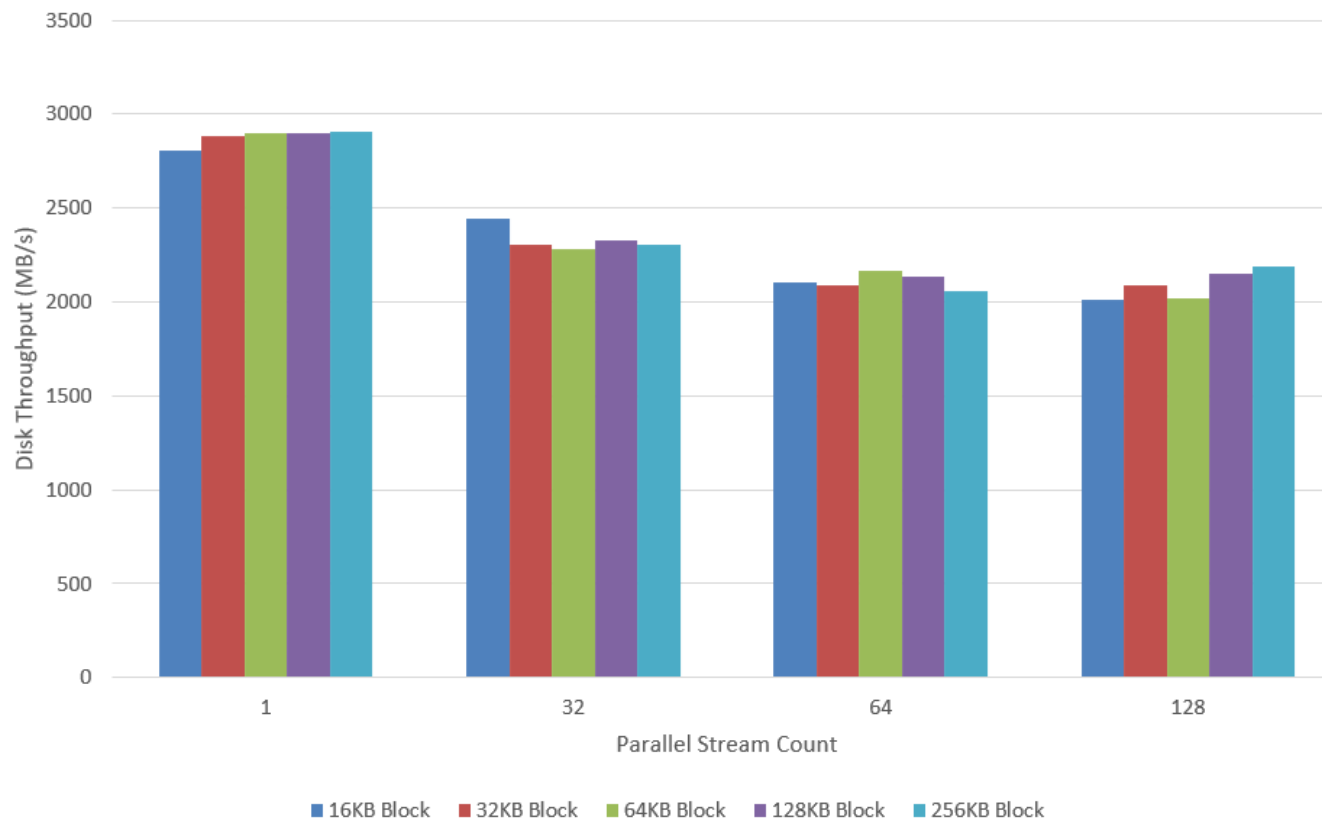





Raw Disk Throughput  
(6 drawer - write only)

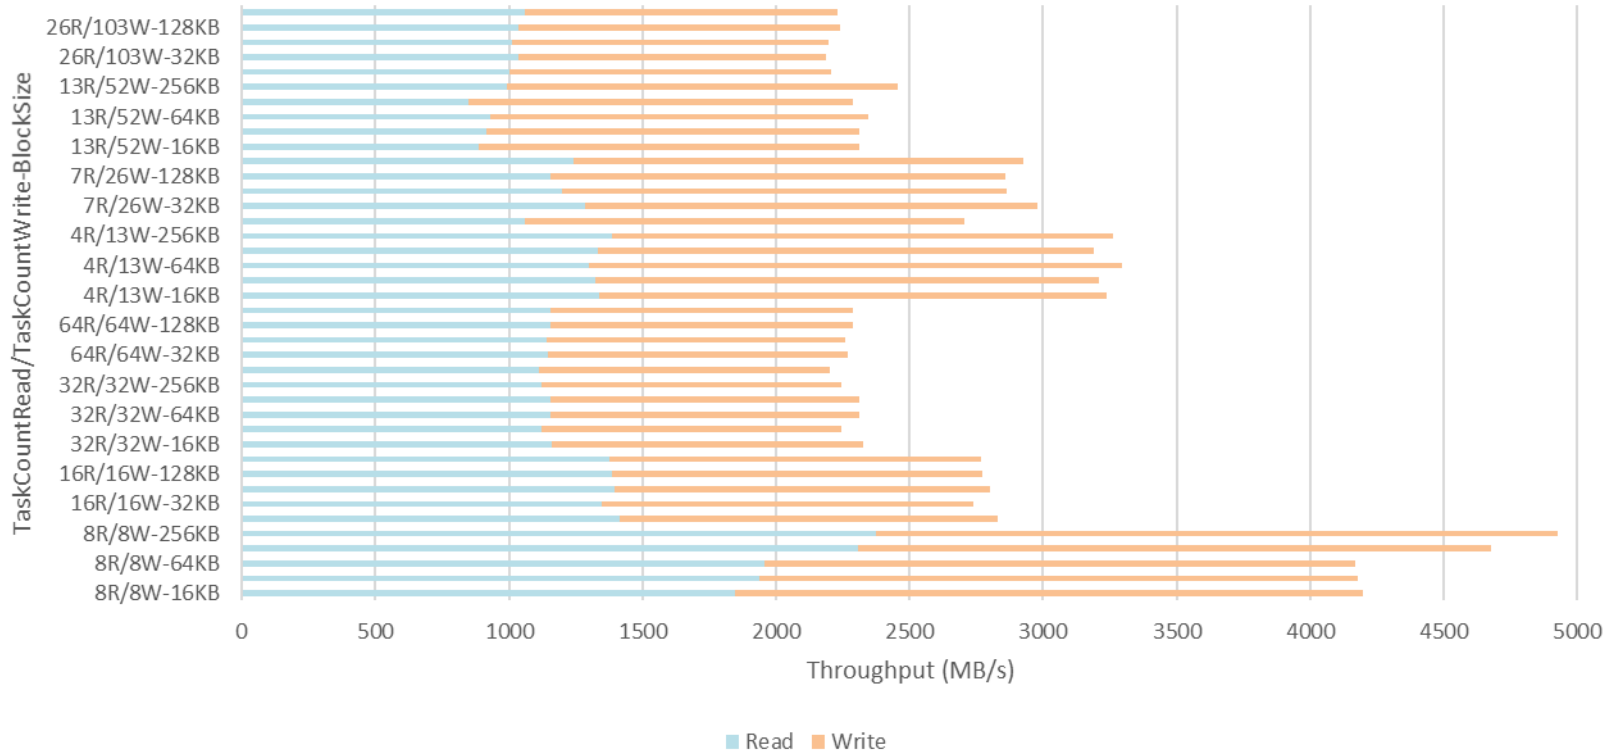
## Raw Disk Throughput (6 drawer - read + write combinations at different block sizes)





Raw Disk Throughput  
(4 drawer - read only)

Raw Disk Throughput  
(4 drawer - write only)

### Raw Disk Throughput (4 drawer - read + write combinations at different block sizes)




Raw Disk Throughput  
(2 drawer - read only)

Raw Disk Throughput  
(2 drawer - write only)



## Raw Disk Throughput (2 drawer - read + write combinations at different block sizes)



**Thank You.**  
**IBM Storage & SDI**

# Best Practices: Which one should I use?

- ❖ If compression is the goal, ZSTD would be the go-to algorithm. But performance is also a variable that customers need to take into account.

## A00 data pattern

Completely random data pattern aiming to stress the TS7700 cache.

- FICON compression shows 0.89:1 (expansion)
- LZ4 compression shows 1:1 (expansion automatically avoided)
- ZSTD compression shows 1:1 (expansion automatically avoided)