SEAGATE

SSD with Compression: Implementation, Interface and Use Case

Erich F. Haratsch

Seagate Technology

SSD with Compression: Introduction

- Data reduction techniques such as compression and deduplication have been employed in some storage systems, but are not widely available inside SSDs yet
- Confusion exists about benefits, use cases and data integrity when SSDs implement compression

Compressibility of Data

- Data bases, OS files, application data are typically highly compressible
- Image and video files may have some small compressibility left

	Typical Average Compressibility per Workload			
Compression Algorithm	MySQL	Oracle	Win8	Linux VM
gzip	60%	70%	50%	60%

- Compressibility =

(1 – OutBytes / InBytes) * 100%

SSD with Compression

- Compression algorithm needs to be lossless
- Compression needs to run inline at full data rate: low impact to write and read latencies
- Compression needs to be done before encryption and ECC encoding
- Compression reduces data written to media
- Write original data if data is incompressible

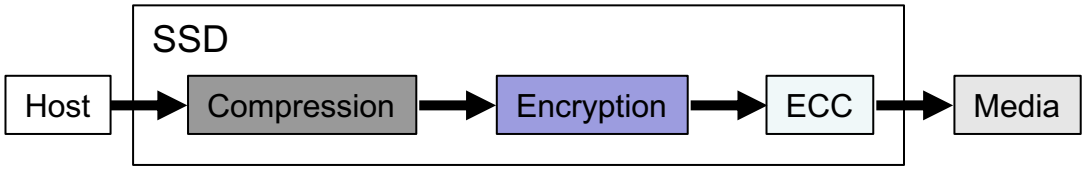

SSD Flash Translation Layer with Compression

- Traditional FTL writes data chunks with equal physical size that fit into a flash page
- With compression, FTL needs to have ability to manage physical data units of variable size

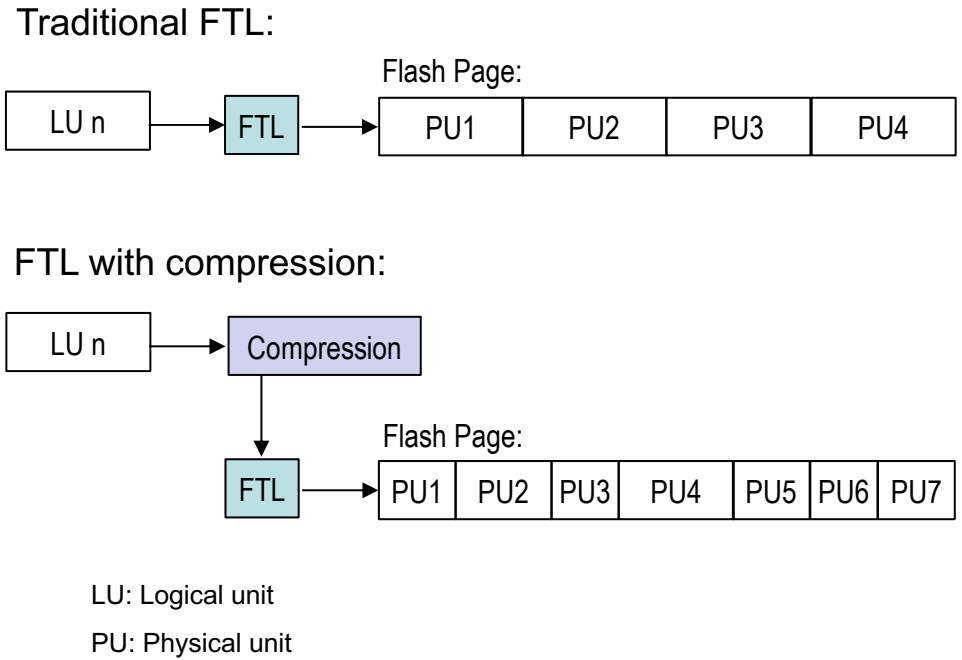

Use Compression to Increase Effective Overprovisioning

- Logical capacity does not change
- Reduces write amplification
- Increases random write and mixed read/write performance
- Increases endurance

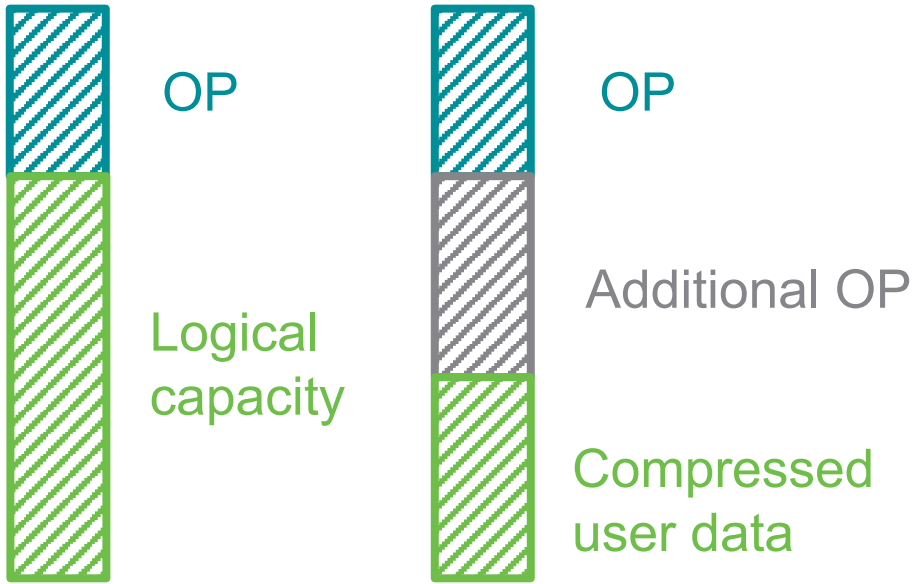

Flash Memory Summit

Write Amp, OP and Compression

- For random write workloads, write amplification increases as OP decreases
- Compression increases available OP
- Compression reduces WA and therefore increases endurance and performance

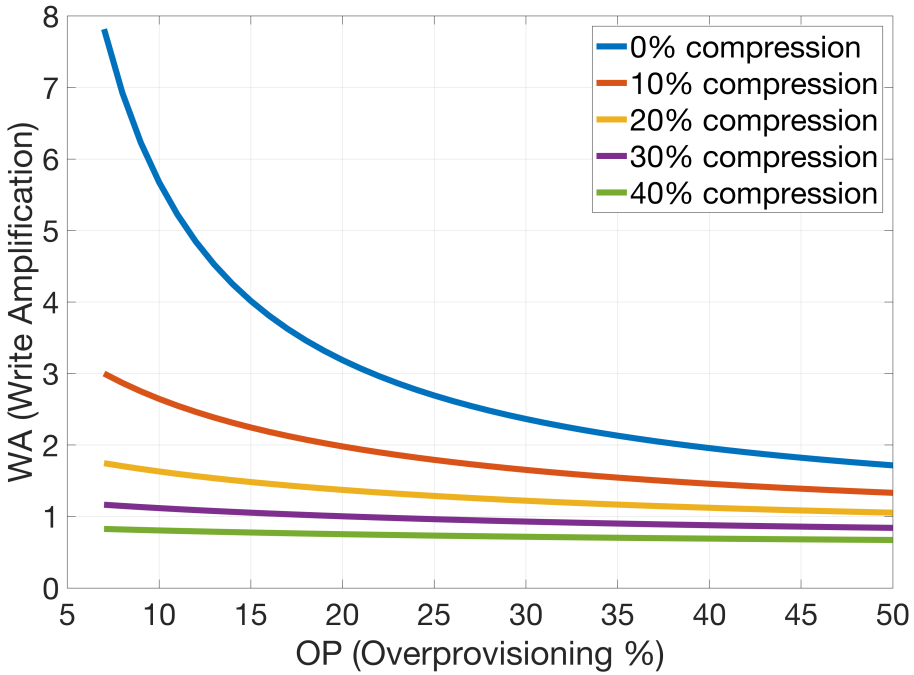

Use Compression to Increase Logical Capacity

- Report higher logical capacity to host
- Actual logical capacity depends on data entropy
- Host needs to monitor free physical space

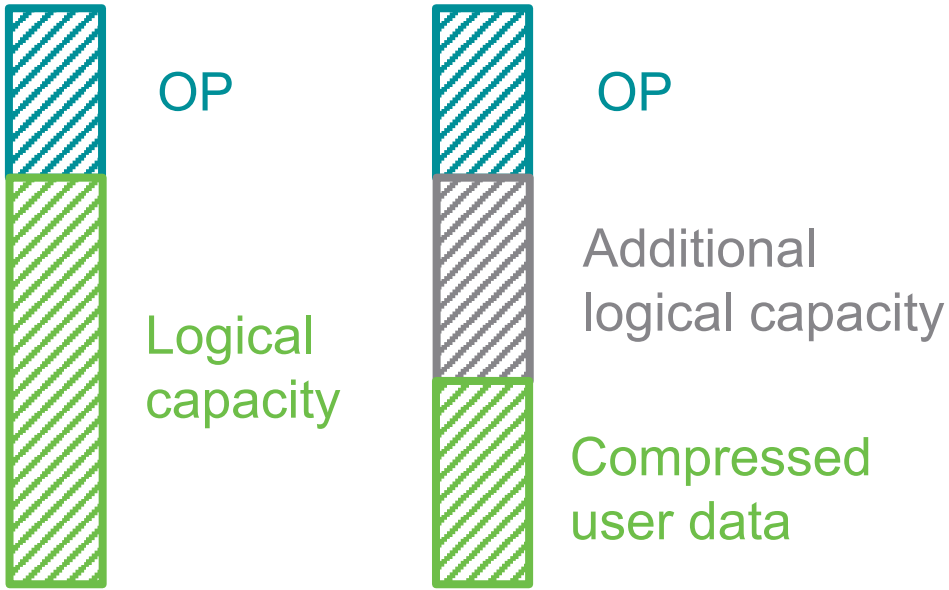

QLC SSD with Compression

- QLC NAND media has typically low endurance and performance characteristics
- Compression can make QLC SSDs more attractive by increasing
  - endurance
  - performance
  - user capacity

Flash Memory Summit

SSD Product with Compression

- Nytro® 1000 SSD series
- Seagate DuraWrite™ lossless data reduction technology is designed to increase performance and deliver high-power efficiency
- Tunable capacity for performance- or capacity-optimized SSD solutions
- Seagate Secure technology with secure supply chain, SD&D, Seagate Instant Secure Erase, and SED options
- Easy deployment in legacy storage infrastructures with SATA 6Gb/s interface
- Consistent IOPS performance with low latency for faster random access
- Won Best of Show award at Flash Memory Summit 2018

Conclusion

- Compression inside the SSD
  - Increases effective overprovisioning
  - Reduces write amplification
  - Increases endurance and performance
  - Increases logical capacity
  - Can make QLC SSDs more attractive for the data center