FUJITSU

shaping tomorrow with you

Efficient Unstructured Data Compression for Block Storage Systems

Nov 12, 2018 - PDSW-DISCS 2018

Hiroki Ohtsuji, Yoshiyasu Doi
Fujitsu Laboratories Ltd.

Introduction

FUJITSU

- Data compression in storage systems
  - There is a trade-off between:
    - I/O performance
    - Compression ratio
- Compressing unstructured data blocks
  - Underlying storage system does not know the structure of data.
  - Storage systems have to compress them in an efficient manner without using structural information of application specific data formats.
- This presentation describes a method to improve the compression ratio with the minimal I/O overhead.

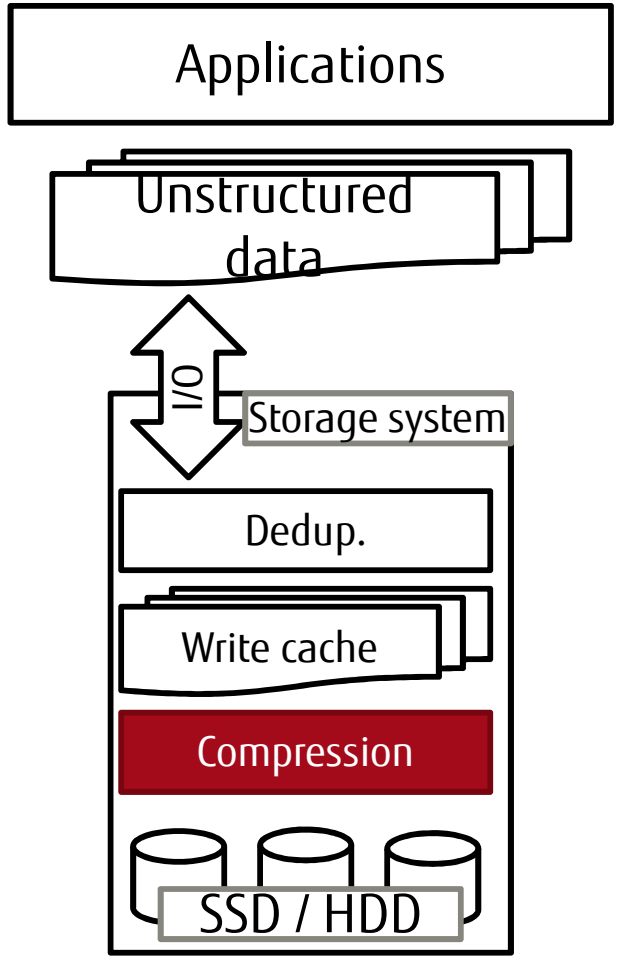

The effect of compression chunk size

FUJITSU

- Using larger chunk size helps dictionary-based compression algorithms to detect more redundant data.
  - We can expect better compression ratio.
  - 4 KB block: 26 %
  - -> 64 KB block: 9 %

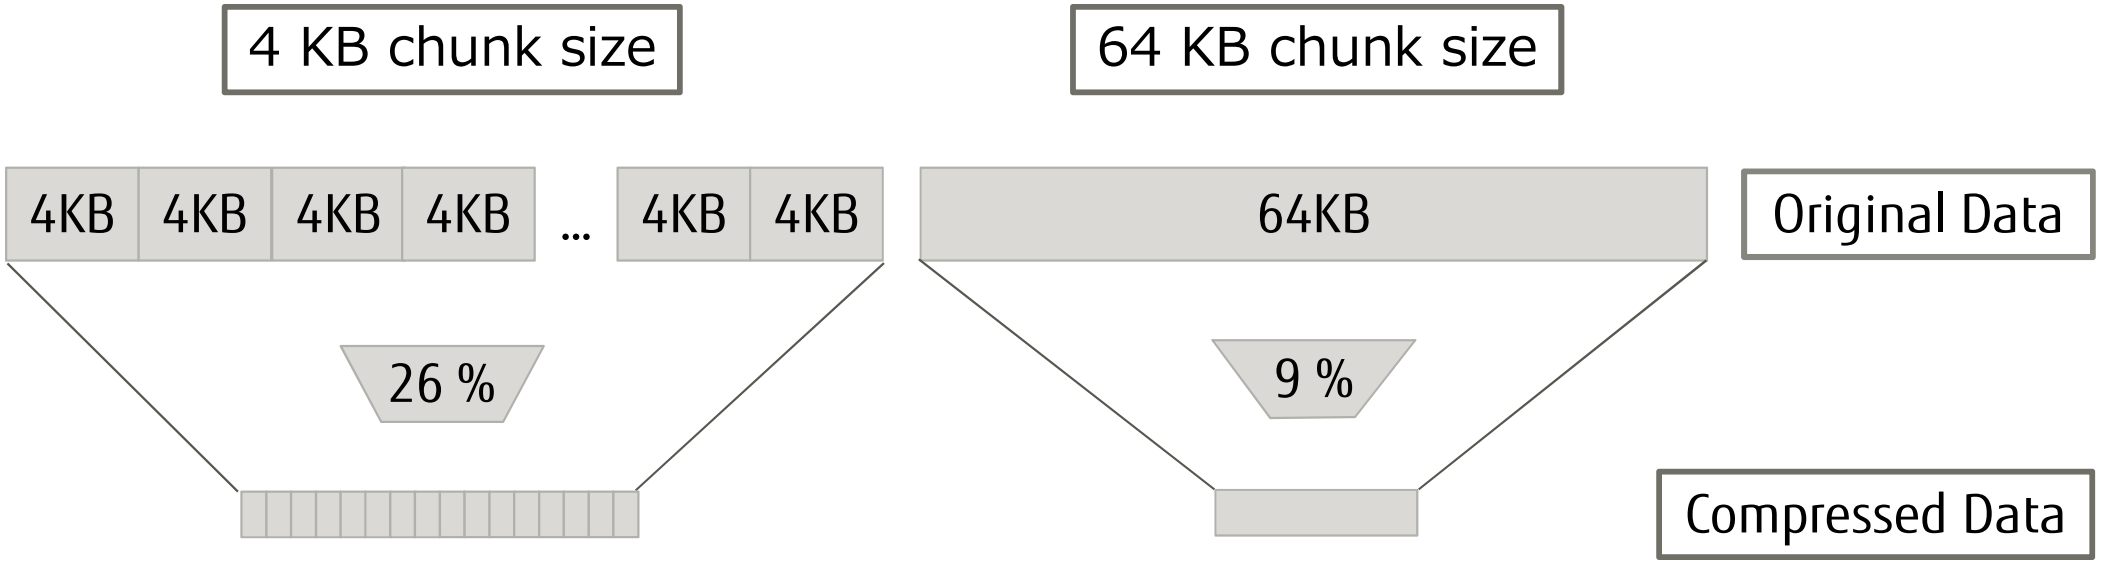

Performance bottleneck: Read amplification

FUJITSU

- A drawback of large compression chunk size
  - The storage systems have to read and decompress unnecessary blocks
    - -> Read amplification problem

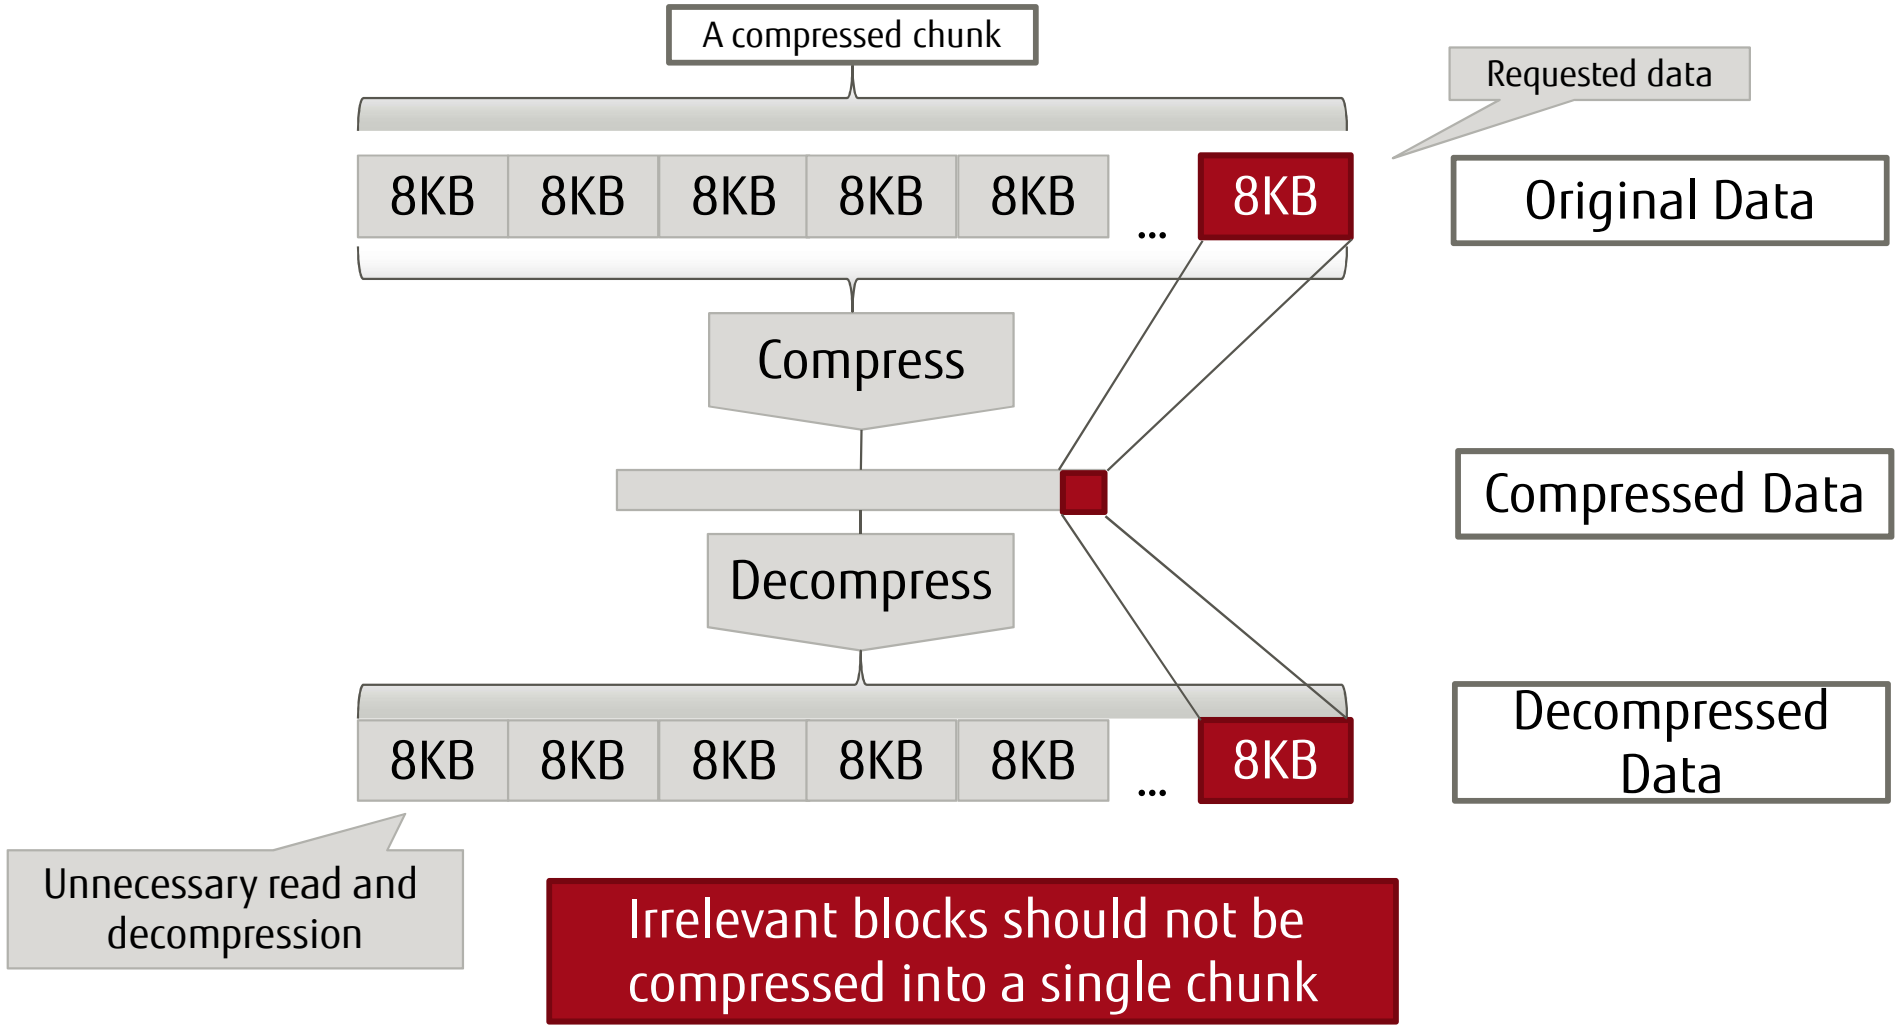

Dynamic chunking for compression

FUJITSU

- Changing the compression chunk size dynamically
  - In accordance with
    - The time stamp of write traces
    - The effective read ratio (read amplification ratio) of read I/O accesses

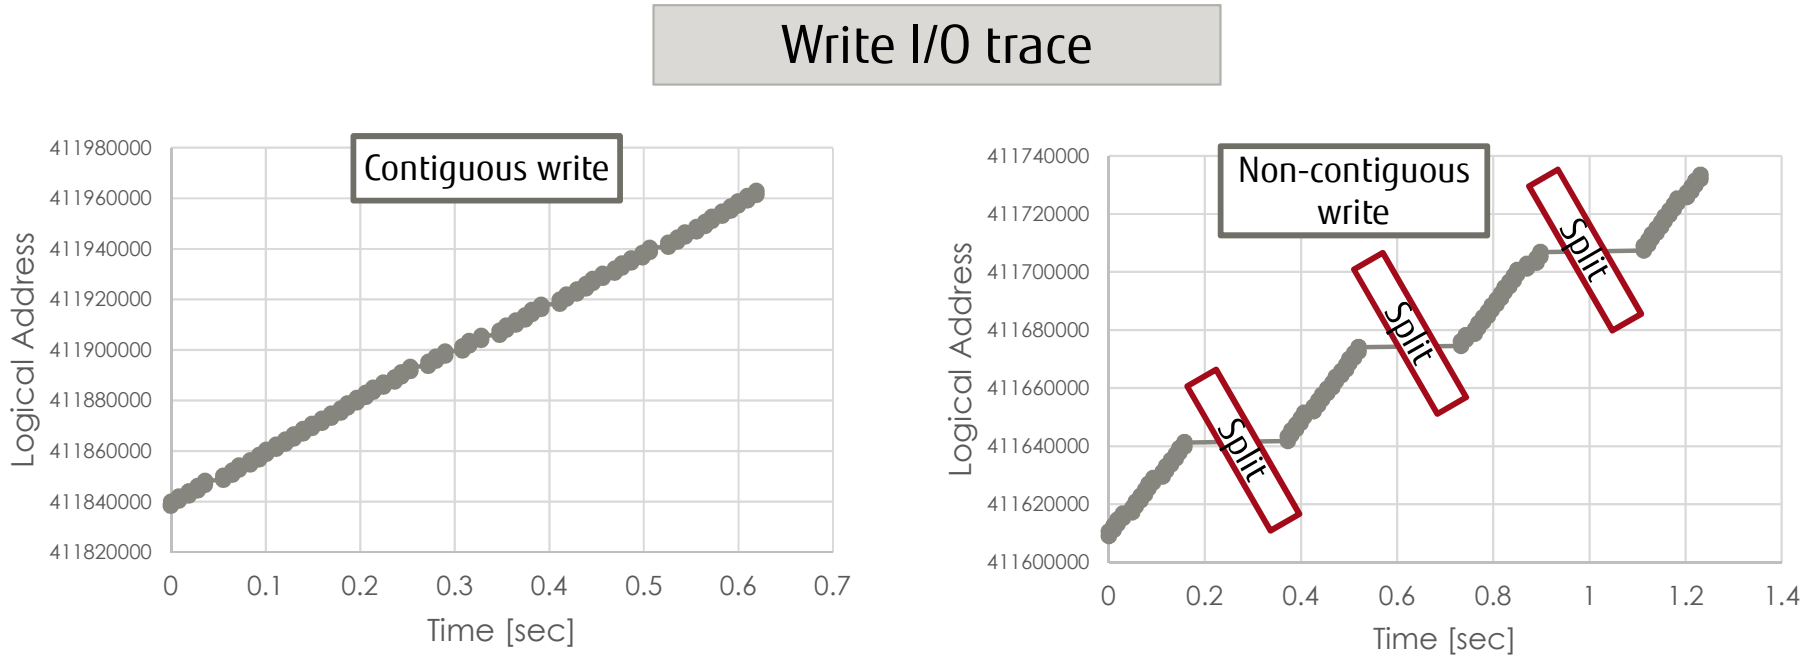

System Design: Dynamic chunking

FUJITSU

- Modified write cache page table
  - Write time stamp for the dynamic chunking method
- Monitoring read I/O operations (read cache)
  - Measuring actual read amplification ratio to change the maximum chunk size
    - -> In order to mitigate the worst case.

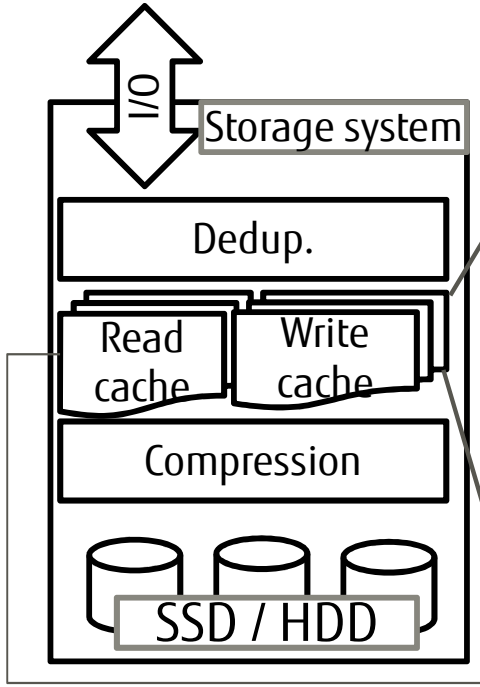

Write cache page table

Page ID	Block (Hash)	Reference Count (Dedup.)	Physical address	Offset in the chunk	Write time stamp	Logical Unit #
0	0x2345	2	0xabcd	0	10	1
1	0x3456	1	0xbcde	0	100	1
2	0x4567	1	0xbcde	1	101	1

An additional module measures the effective read ratio in order to adjust the maximum size of compressed chunks.

Compression ratio

- According to the preliminary evaluation of the compression ratio,
  - Expected improvement of the compression ratio
    - 26 % to 7% (-19 points)

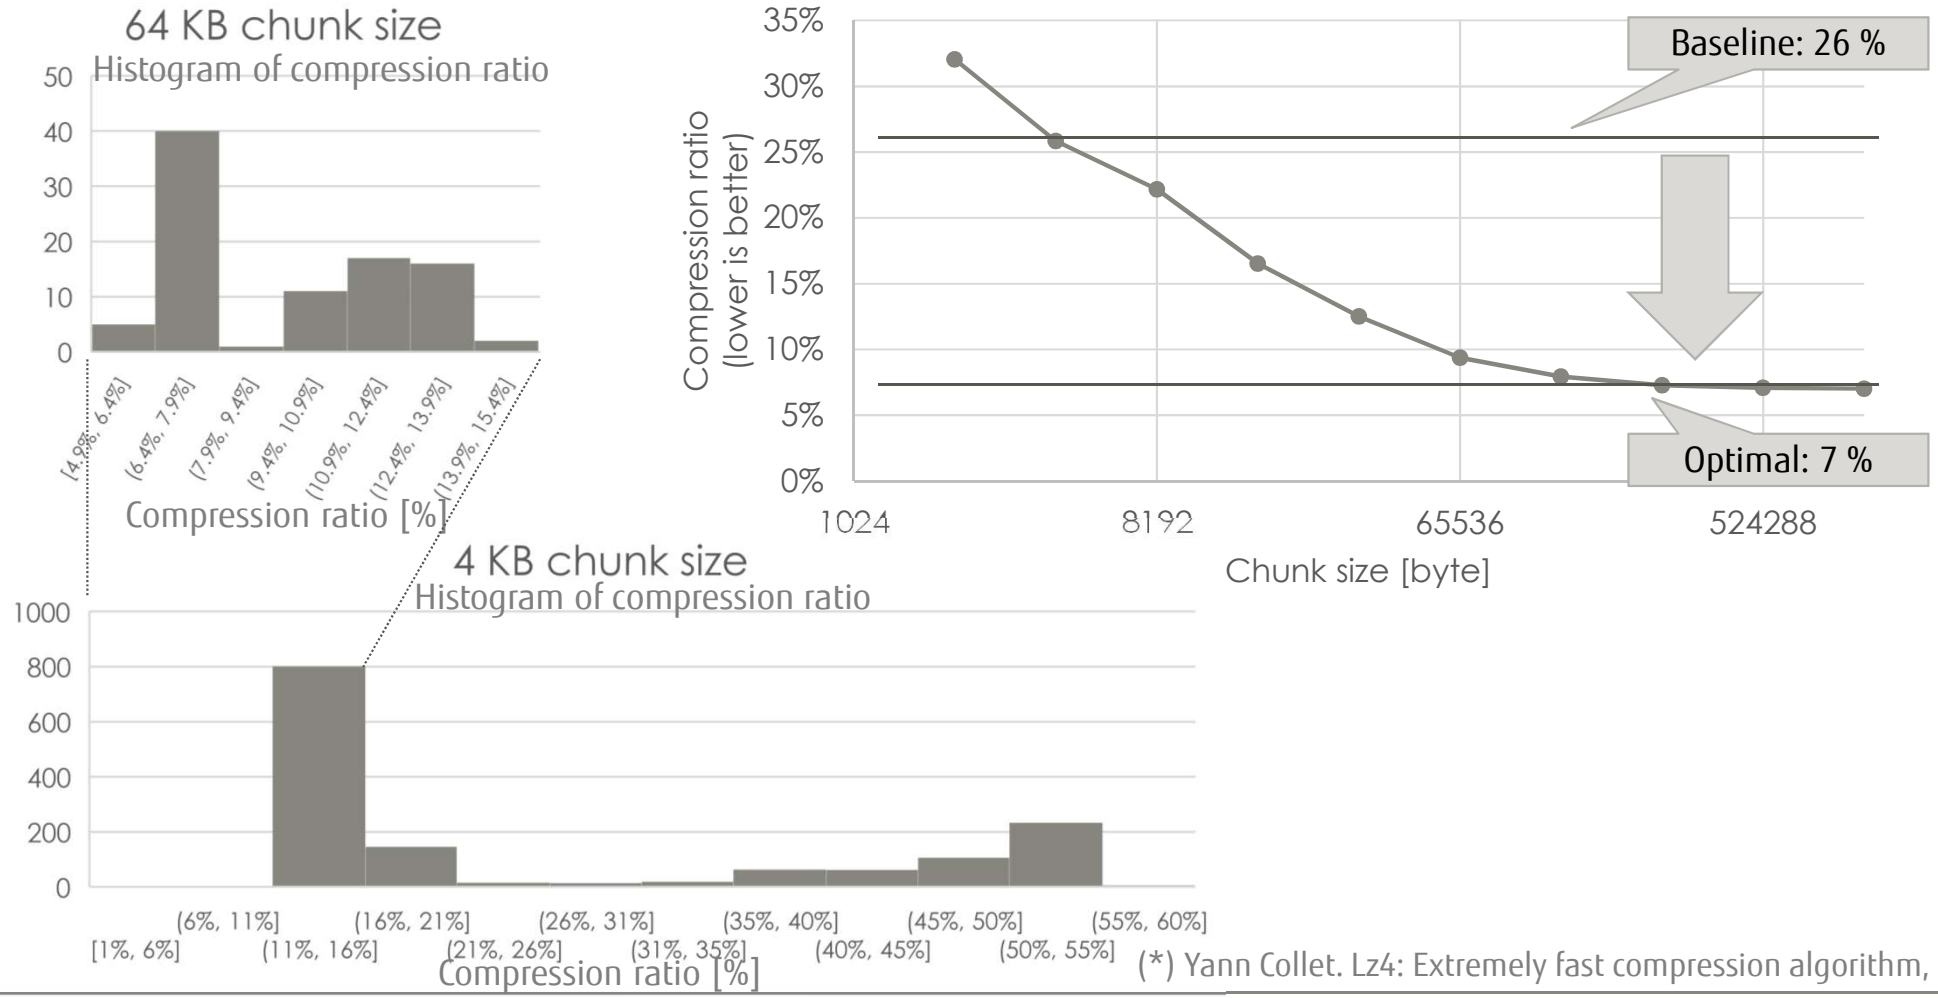

(*) Yann Collet. Lz4: Extremely fast compression algorithm, 2013.

Summary and Future work

FUJITSU

- Compressing unstructured data in storage systems
  - Requires a dynamic mechanism to improve the compression ratio
- Solving the read amplification problem
  - Reducing unnecessary read accesses by using the dynamic compression chunking method
    - Using write I/O traces
    - Monitoring read I/O operations
- Expected improvement of compression ratio
  - -19 points, compared with the baseline
- Future work
  - Evaluating the method with a complete implementation of the system with real data sets and complicated I/O patterns.

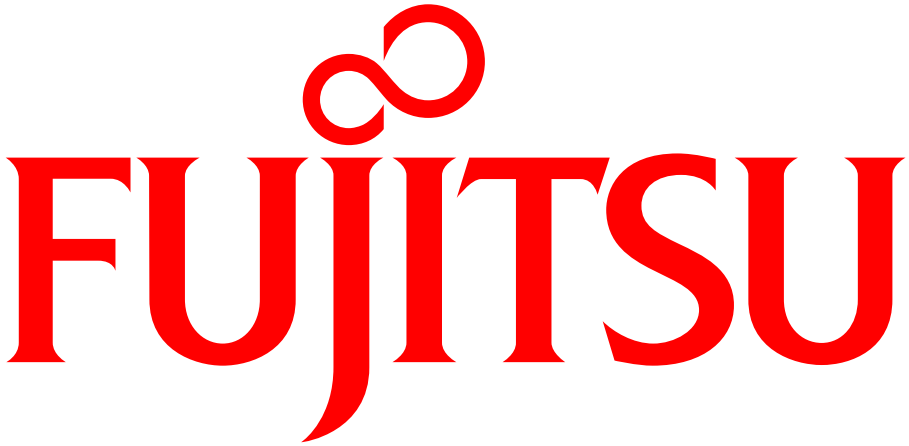

shaping tomorrow with you