SNIA NSF | NETWORKING STORAGE

Compression: Putting the Squeeze on Storage

Live Webcast

September 2, 2020

11:00 am PT

Today's Presenters

Ilker Cebeli
Moderator
Samsung

John Kim
Chair, SNIA Networking Storage Forum
NVIDIA

Brian Will
Intel® QuickAssist Technology
Software Architect
Intel

SNIA-At-A-Glance

SNIA-at-a-Glance

185
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

Learn more: **snia.org/technical** **@SNIA**

Technologies We Cover

- Ethernet
- iSCSI
- NVMe-oF
- InfiniBand
- Fibre Channel, FCoE
- Hyperconverged (HCI)
- Storage protocols (block, file, object)
- Virtualized storage
- Software-defined storage

SNIA NSF | NETWORKING STORAGE

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Agenda

- Why, Where and What to Compress
- Data Compression Techniques

Why, Where and What to Compress

John Kim

Why Compress Data?

Yes, Compress!

- Save space
- Reduce network bandwidth
- Doesn’t lose data (usually)

No, Work Less!

- More time to write/read files
- Consumes processing power
- Long-term accessibility risk

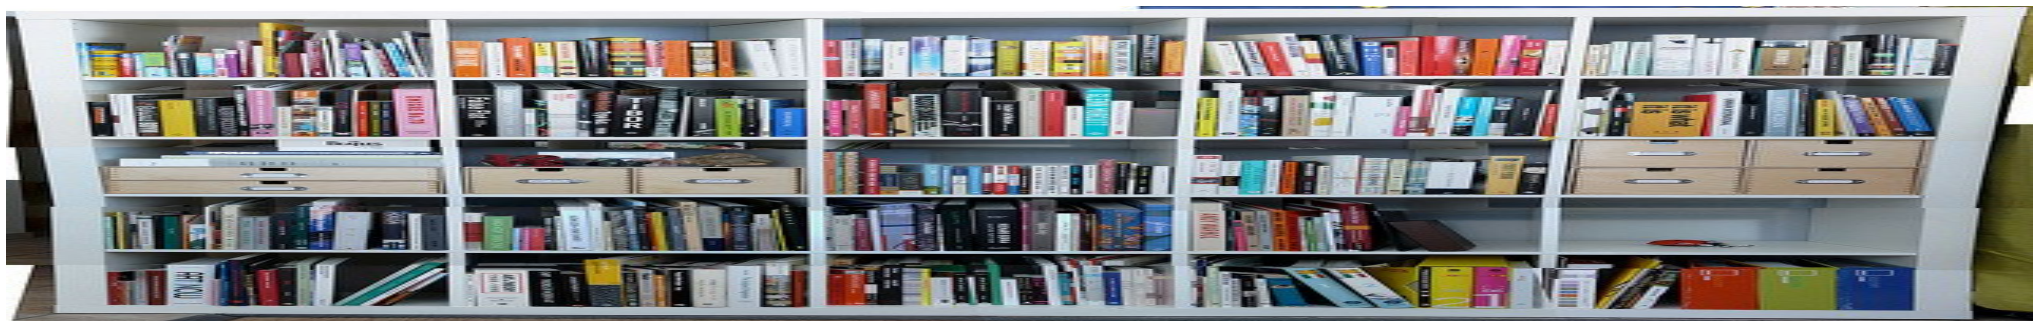

Where to Compress the Data

- In application – e.g. Photos, videos
- At client – e.g. Zip files
- On network – e.g. WAN optimization appliance
- In storage – Storage controller, SmartNIC or computational storage
- During backup – Backup software or hardware

- Can compress in more than one place
  - Especially if data is decompressed whenever it's received/read

Where Data Can be Compressed

- Client, server, storage, and/or backup compress data
- Can accelerate with software library, FPGA, SmartNIC or DPU

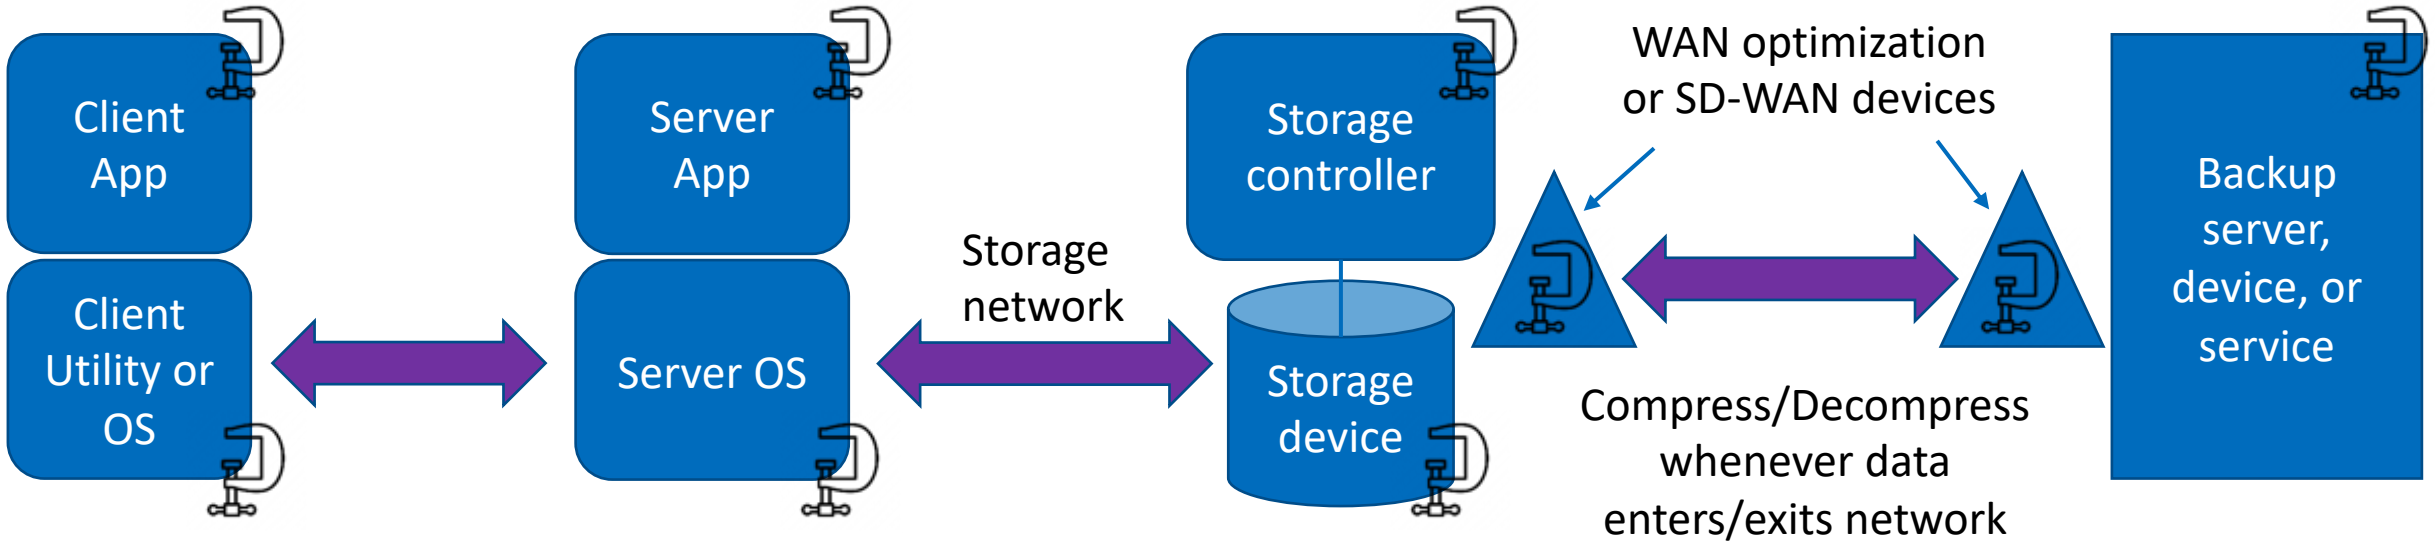

What Types of Data Should be Compressed

Yes, Compress!

- Contains empty space
- Local repetitive sequences
  - At bit/byte level
  - Within each file/object
  - Locally-repeated blocks

No, Don't Stress

- Already compressed
- Non-local repetition
  - At file/object level
  - Identical files/objects
  - Dispersed block repetition

Good and Bad Compression Candidates

Easy to compress File

Many repeated characters, large blocks of empty space. Many repeated characters. Many repeated characters.
ABCD ABCD ABCD ABCD
00000000000000000000000000000000
5678 5678 5678 5678 5678 5678 5678
My dog has fleas My dog has fleas
My dog has fleas

Hard to compress files

My dog has fleas	My cat has NVMe SSDs	My dog has fleas
My cat has NVMe SSDs	The skunk smells nice	My dog has fleas
The skunk smells nice	My dog has fleas	My cat has NVMe SSDs

Data Compression Techniques

Brian Will

What is Data Compression

- **Definition**: “The process of encoding data to reduce its size.”*
- What is lossless vs lossy?
  - Lossless: “compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed exactly”*
  - Lossy: “compression using a technique in which a portion of the original information is lost.”*

*SNIA: https://www.snia.org/education/online-dictionary/term/data-compression

LZ Lossless Compression Introduction

Dictionary Coding (example LZ77)

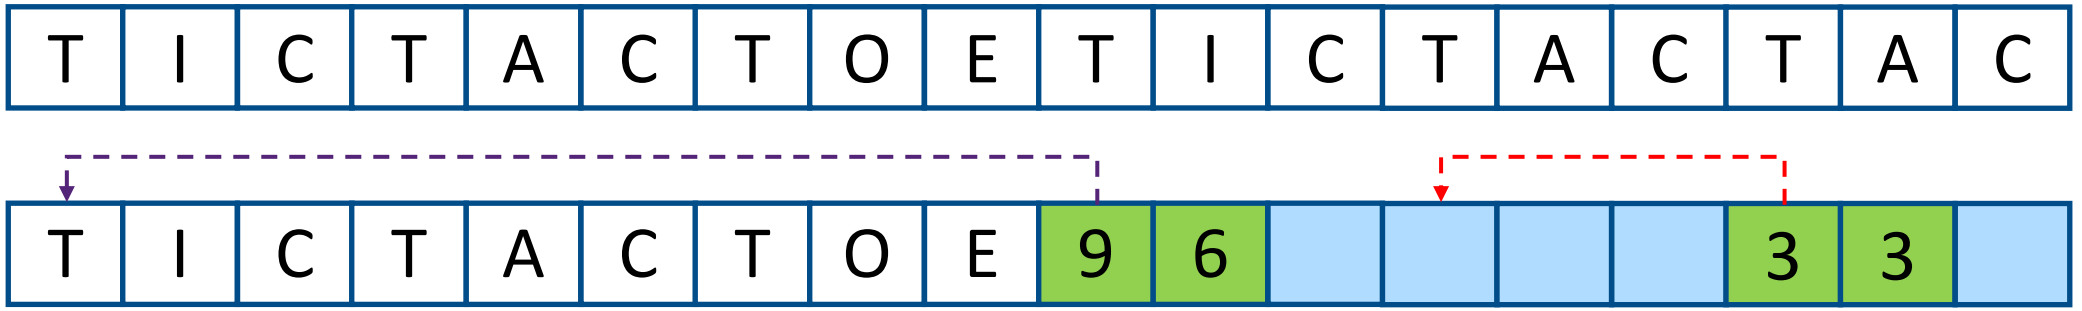

Min Match: 3
Window: > 6

Distance-Length Pair

Entropy Encoding (examples Huffman and Arithmetic coding etc.)

Char	Freq	Code (bits)
T	6	00
C	5	01
A	3	100
I	2	101
E	1	110
O	1	111

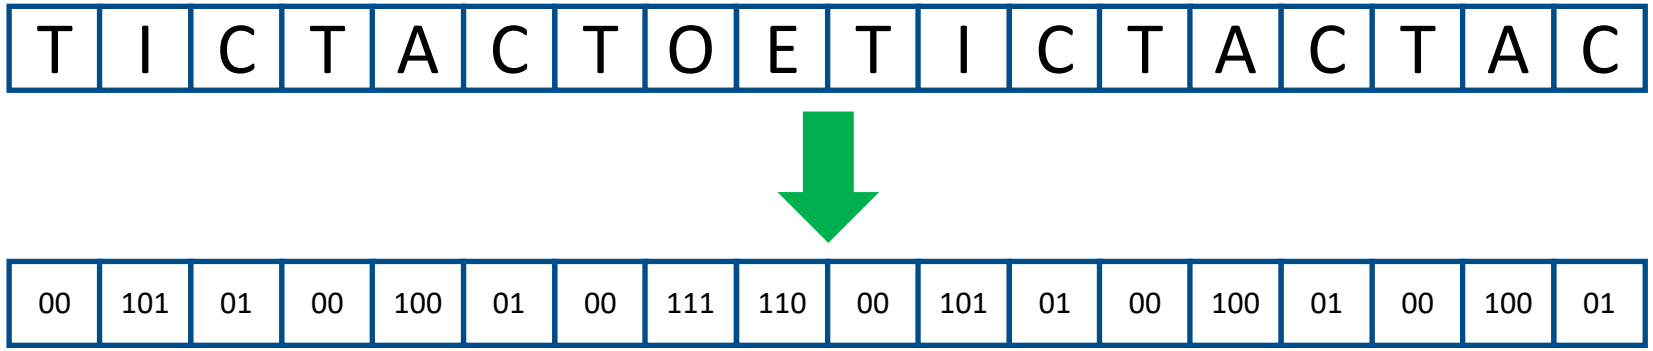

Different Compression Algorithms

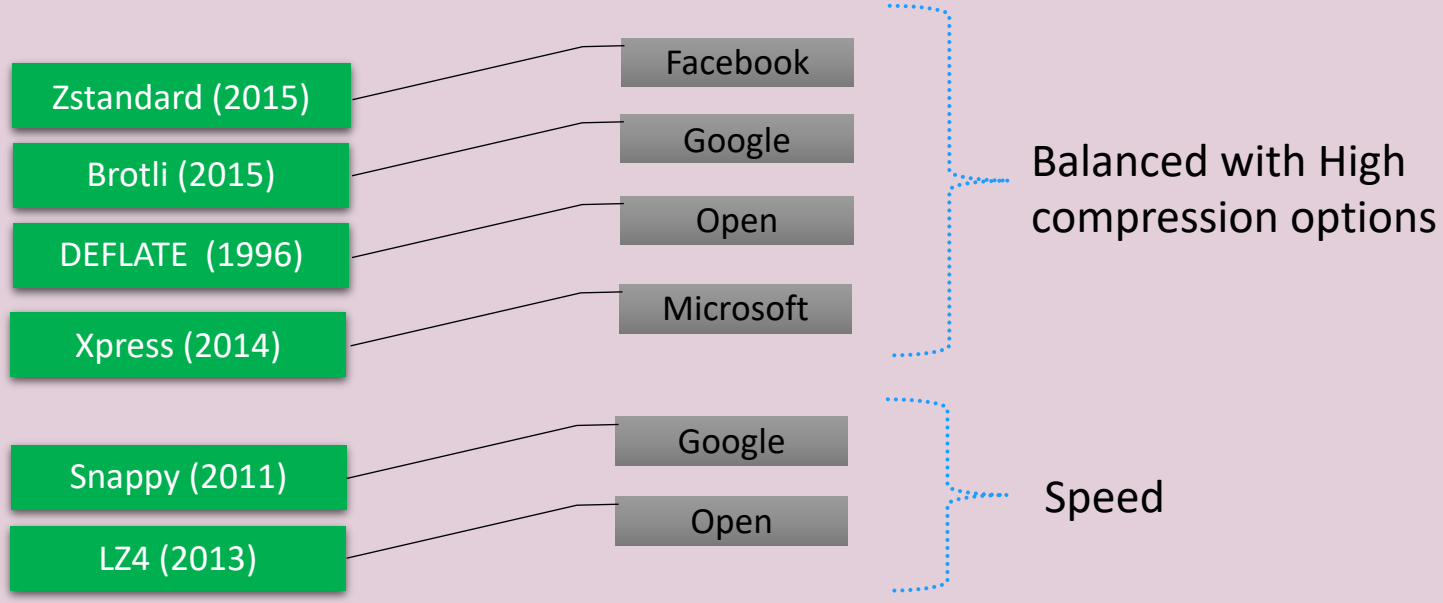

- To the left are some of the predominant lossless compression algorithms utilized today.
- They are integrated into many workloads/software stacks:
  - Linux Kernel Crypto Framework(LKCF)
  - ZFS, BTRFS, SquashFS
  - MongoDB, RocksDB, Hadoop, MySQL
  - WinZip, gzip
  - HTTP encoding (NGINX/HAProxy/Apache Traffic Server etc.)

Techniques employed:

- LZ77/LZ78
- Huffman Encoding
- Finite State Entropy Encoding
- Burrows-Wheeler transform
- Pre-built dictionaries
- etc.

This is only a small subset of the entire landscape..

lzf	lzo	lzma	Blosclz	zopfli	crush	lzg
lzham	lizard	rle	bzip2	lzw	lzss	ans
lzr	lzj	Lzc	DEFLATE64	PAQ	lzmw	lzrw

***Other names and brands may be claimed as the property of others.**

How Compression Affects Performance

Compression performance vs ratio*

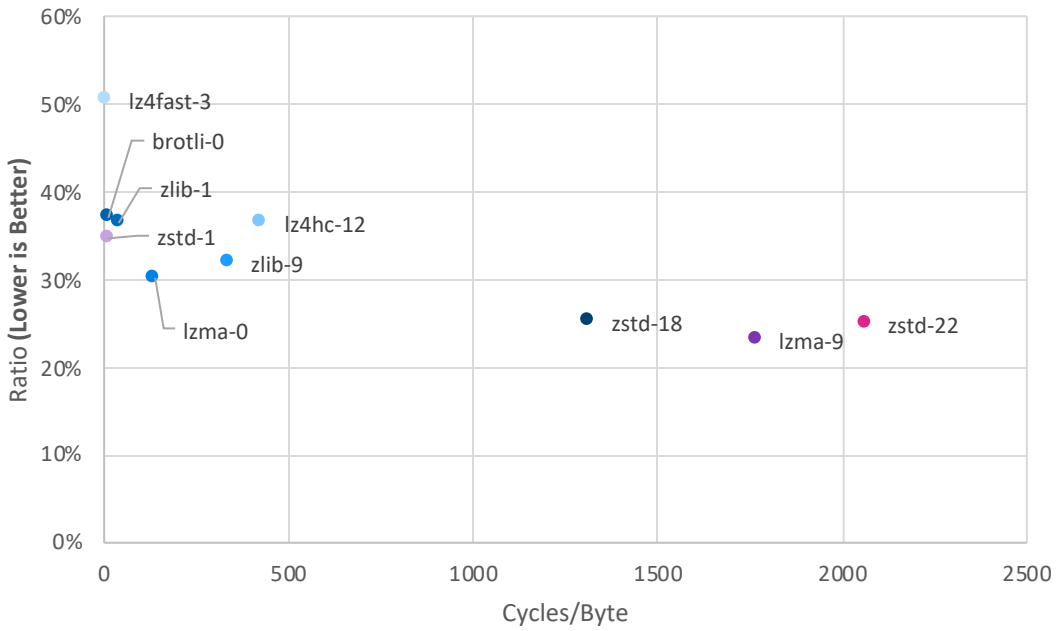

Decompression performance vs ratio*

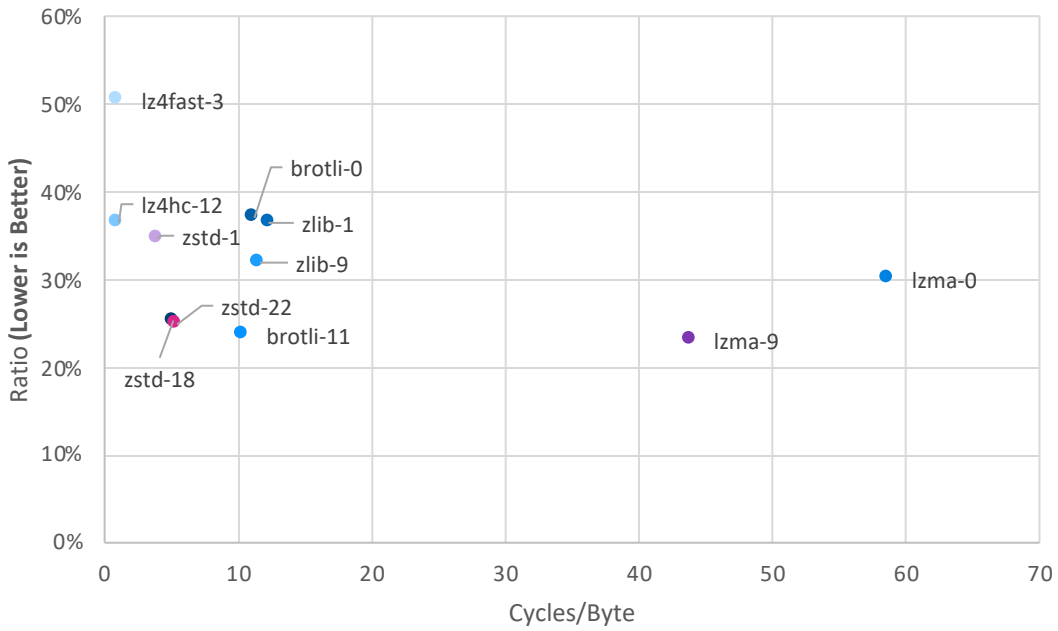

- Compression performance varies significantly based on the ratio level targeted.
  - 1-2% at higher ratios can have a significant cycle cost
- Decompression performance is reasonably clustered based on format rather then compression ratio
  - Much lower cost then compression.

* Data for these graphs taken directly from lzbench: https://github.com/inikep/lzbench

When to Compress: real-time vs post processing

- The choice between real-time and post processing comes down to one of performance (compression ratio & throughput) vs cost
  - Real Time:
    - File Systems reads (decompression)
    - HTTP encoding/decoding
  - Post-Processing/Offline:
    - Log files
    - Cold Storage Tiers

Summary

- Compression saves space, bandwidth, and improves total cost of ownership
- Some types of data can be compressed more than others
- Data can be compressed while transmitted, stored, and backed up
- There are several propriety and open source compression algorithms
  - Compression performance varies with compression ratio
- Choosing the right compression algorithm for workload is critical for performance and cost

Related Webcasts in this Series

- Everything You Wanted to Know About Storage But Were Too Proud to Ask: Data Reduction
  - BrightTALK on-demand: https://www.brighttalk.com/webcast/663/428420
  - SNIAVideo YouTube channel: https://youtu.be/N_y_1frsYtQ
- Not Again! Data Deduplication for Storage Systems – October 2020
  - Follow us for date and time @SNIANSF

After this Webcast

- Please rate this webcast and provide us with your feedback
- This webcast and a copy of the slides will be available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast, including answers to questions we couldn't get to today, will be posted on our blog at https://sniansfblog.org/
- Follow us on Twitter @SNIANSF

Thank You

