



# **Building High Speed Erasure Coding Libraries for ARM and x86 Processors**

*Per Simonsen, CEO, MemoScale*

*May 2017*

# *Agenda*

- MemoScale company and team
- Erasure coding - brief intro
- MemoScale erasure codes
- Performance benchmarks
- Library and plugins

# *MemoScale - the Company*

- Spin off from Norwegian University of Science and Technology
- Result of a research project in algorithm development
- Raised capital spring 2017
- Main product: a high performance erasure coding library

## *MemoScale - the Team*

- 6 people – hard core tech team
- Core competencies:
  - algorithm development for storage, networks and cryptography
  - optimizations of SW/HW interactions
- Fastest SHA3 candidate: Blue Midnight Wish
- Fastest (to our knowledge) implementation of traditional type of Reed Solomon erasure coding

# *What is Erasure Coding?*

*Erasure coding is a cost effective method for protecting against data loss. The method can replace alternative data protection such as replication and traditional RAID.*

# *Replication: Copies of Data*



## 3-replication: Tolerates two *Losses*



# *Erasure Coding Step 1: Split Data Into Fragments*





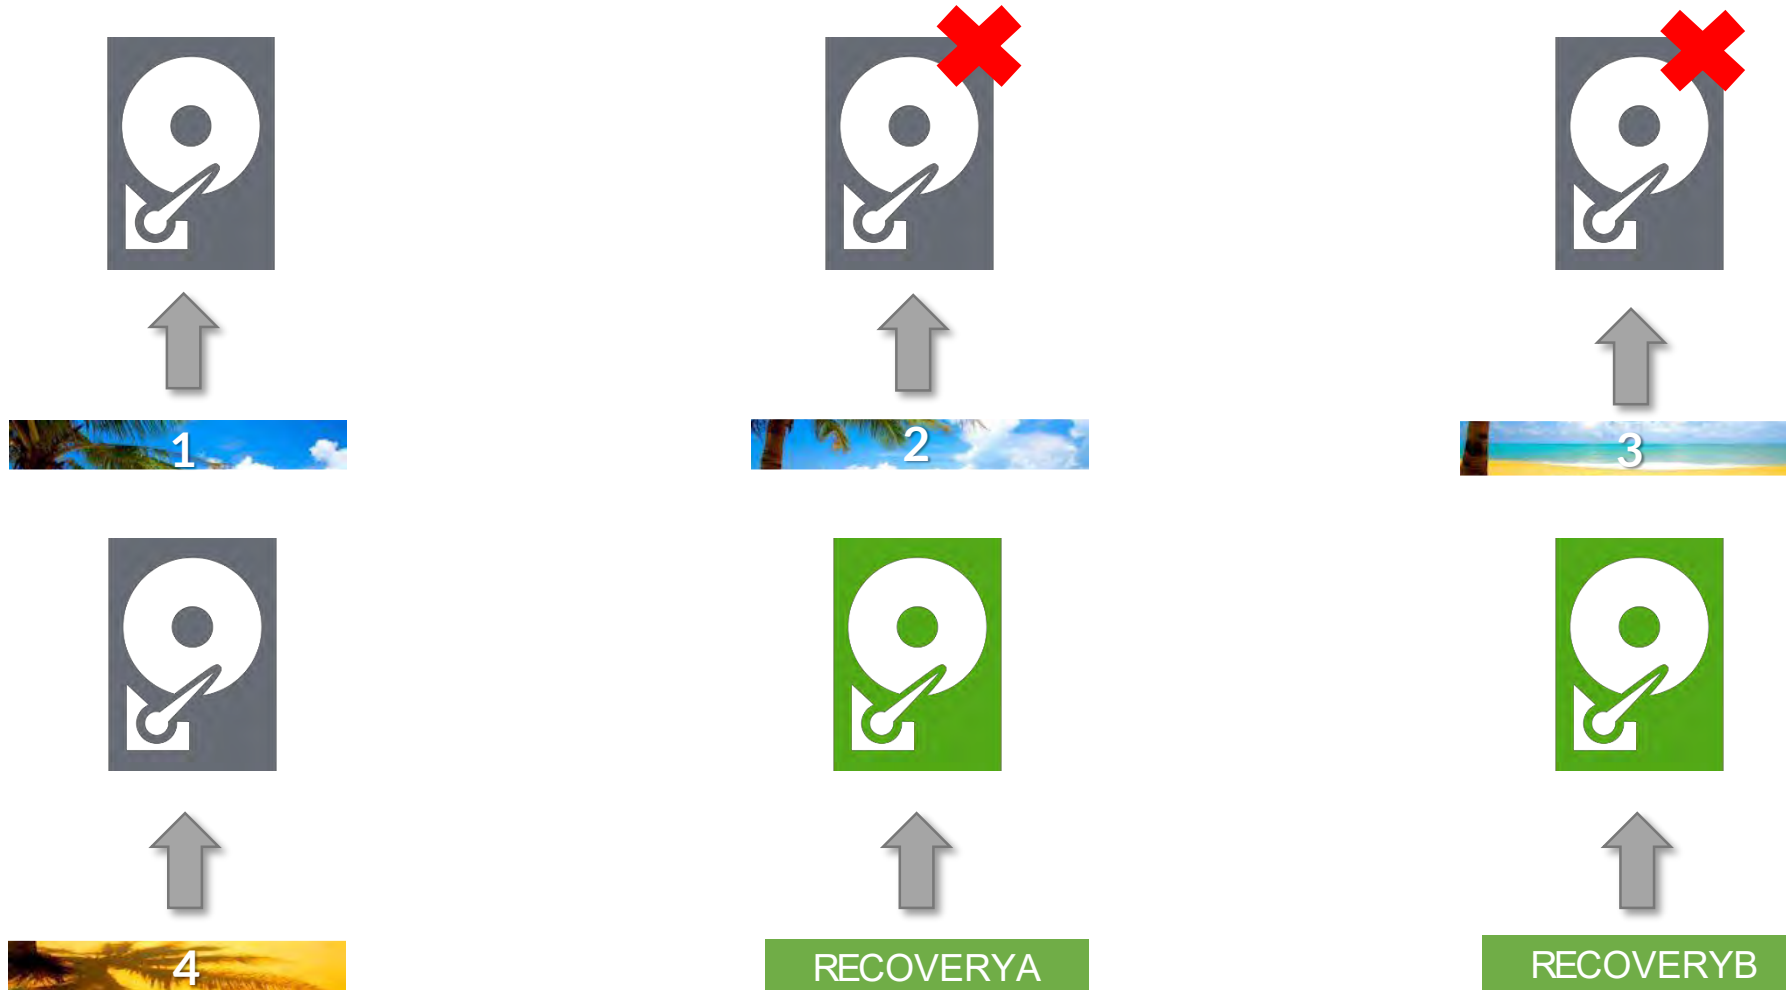
# *Erasure Coding Step 2: Add Recovery Fragments*




## Recovery fragment

- A "combination" of data fragments
- Can recover any lost fragment

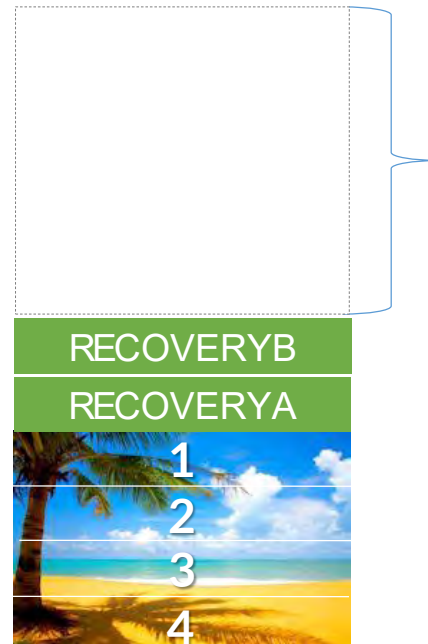
# Erasure Coding (4+2): Tolerates Two Losses



# Storage Space Required



**Replication**



**Erasure coding**

- 50 % reduction in storage space required
- 100 % increase in usable capacity

# *Challenges of Erasure Coding*

High network traffic

High I/O

Compute intensive

# *MemoScale Erasure Coding Library*

Reduced network traffic ✓

Reduced I/O ✓

Reduced compute ✓



1. Improved performance
2. Higher levels of cost savings/usable capacity
3. Improved data loss resiliency

# *MemoScale Erasure Coding Library*

Reduced network traffic ✓

Reduced I/O ✓

Reduced compute ✓




1. Improved performance
2. Higher levels of cost savings/usable capacity
3. Improved data loss resiliency

## *Why Reduce Recovery Traffic?*

- Faster recovery
- Better protection level
- More network capacity available
- Configurations with less overhead

# Recovery Traffic Increases With Erasure Coding

Replication



Erasure coding\*




\* Reed Solomon erasure coding: recovery traffic for a drive equals number of data fragments x drive size




# *MemoScale: ~1/3 of the Recovery Traffic\**

Replication



Erasure coding



\*Memoscale erasure code has approximately 1/3 of the recovery traffic of Reed Solomon erasure code for the same configurations of data blocks and redundancies. The code is MDS (no extra redundancy blocks required).



# *MemoScale Erasure Coding Library*

Reduced network traffic ✓

Reduced I/O ✓

Reduced compute ✓




1. Improved performance
2. Higher levels of cost savings/usable capacity
3. Improved data loss resiliency

# *MemoScale Compute Optimizations*


- Reed Solomon optimized for fast compute
- MemoScale erasure code optimized for fast compute
  - Has MDS property like Reed Solomon erasure code
- For challenging compute scenarios: extremely low latency, high throughput or low end processors

# *MemoScale Compute Optimizations*



**ARM**


**AMD** 



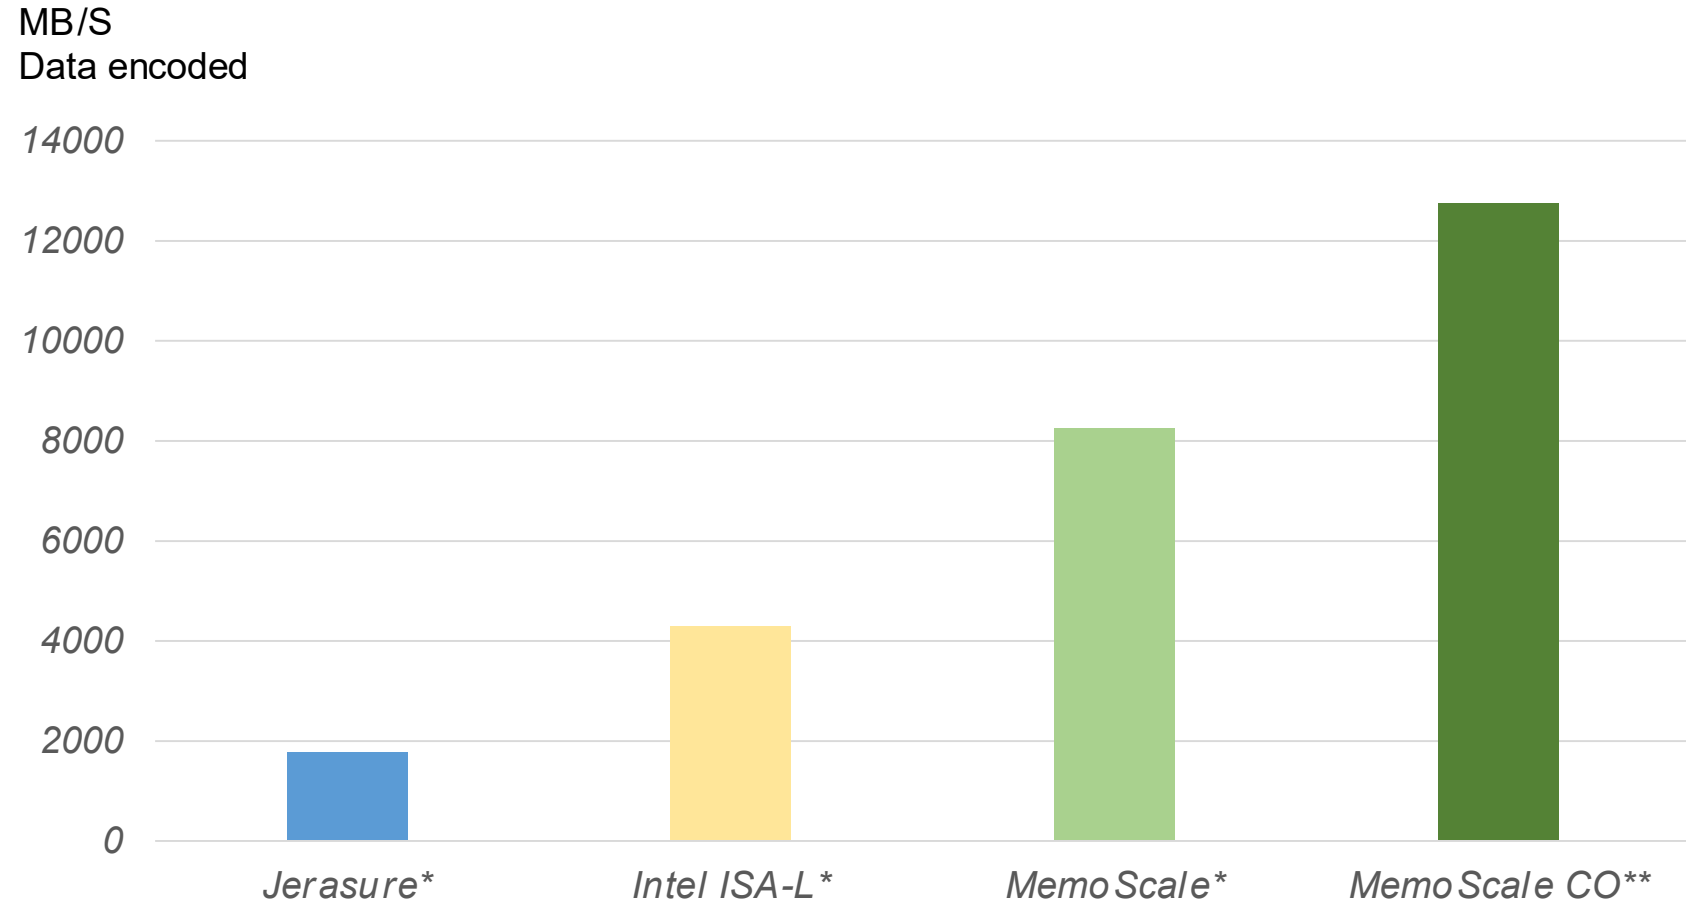
# Test Setup

<b>Data blocks</b>	10
<b>Redundancy blocks</b>	4
<b>Block size</b>	4 MB
<b>Number of cores/threads</b>	1
<b>Measurement criteria</b>	Data encoded excl. redundancies, MB/s
<b>Codes tested</b>	Reed Solomon and MemoScale compute optimized EC
<b>Libraries tested</b>	Jerasure, Intel ISA-L, MemoScale

# Encoding Speed Benchmark – Intel E5-2676 processor



Processor:  
Intel Xeon E5-  
2676 v3 2.4 Ghz



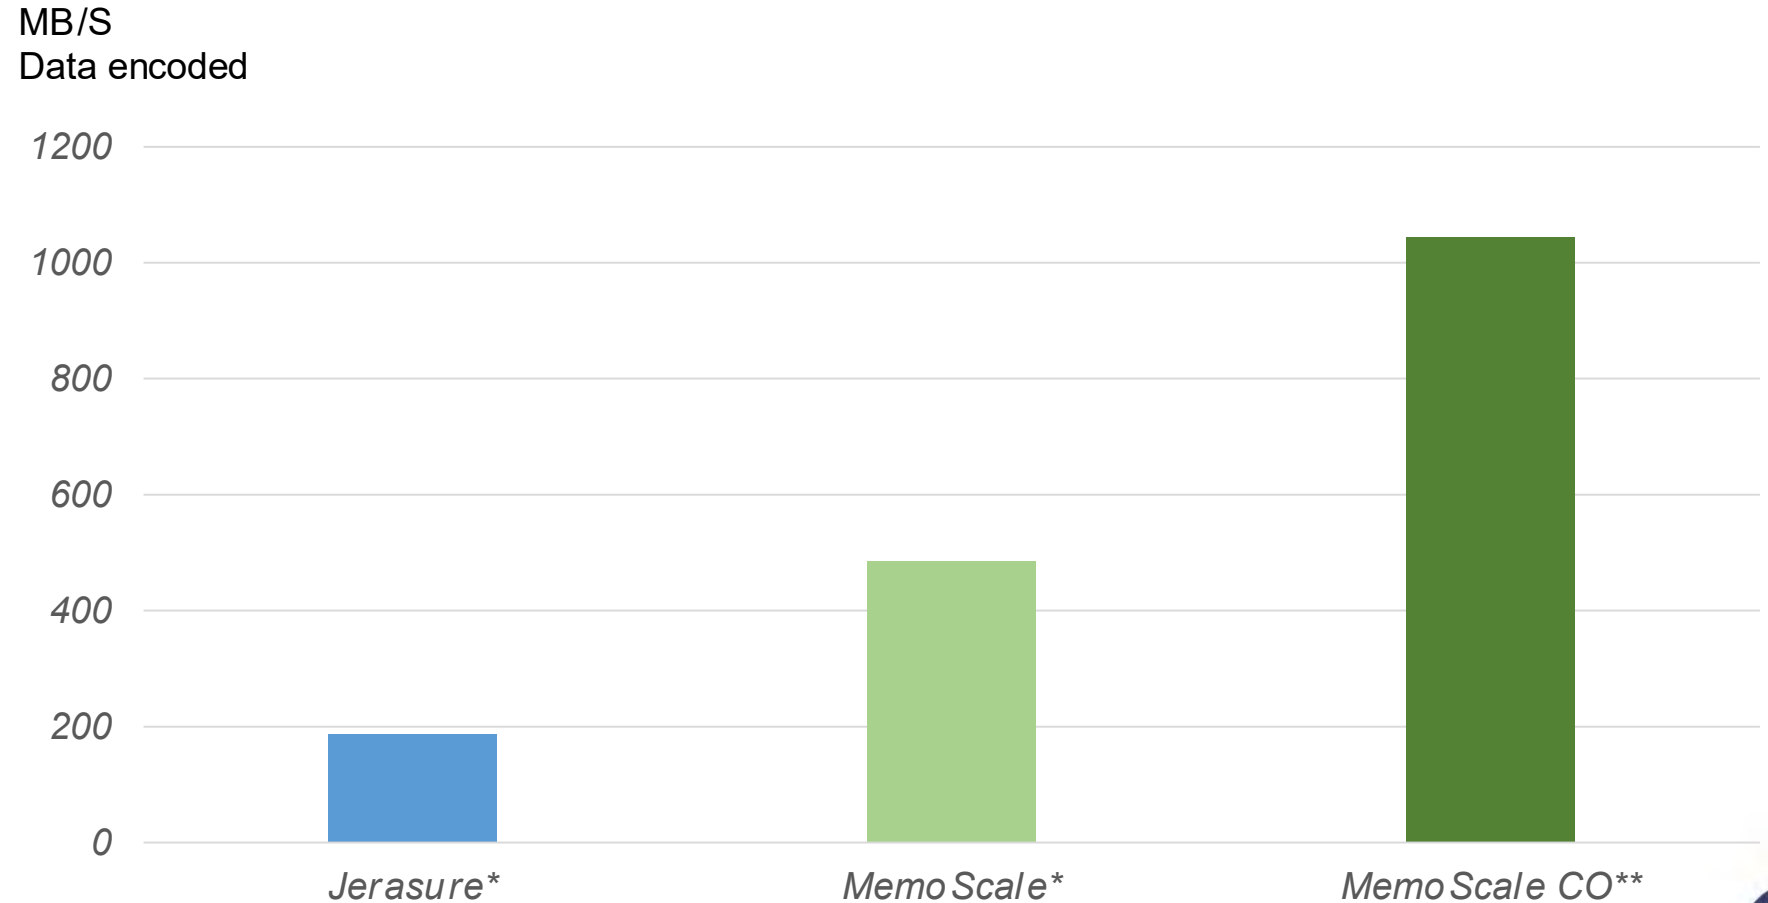
\* Erasure code: Reed Solomon    \*\* Erasure code: MemoScale compute optimized



# Encoding Speed Benchmark – ARM A53 Processor

# ARM

Processor:  
HiSilicon Kirin 620  
processor, ARM  
Cortex-A53 64-bit  
SoC 1.2ghz



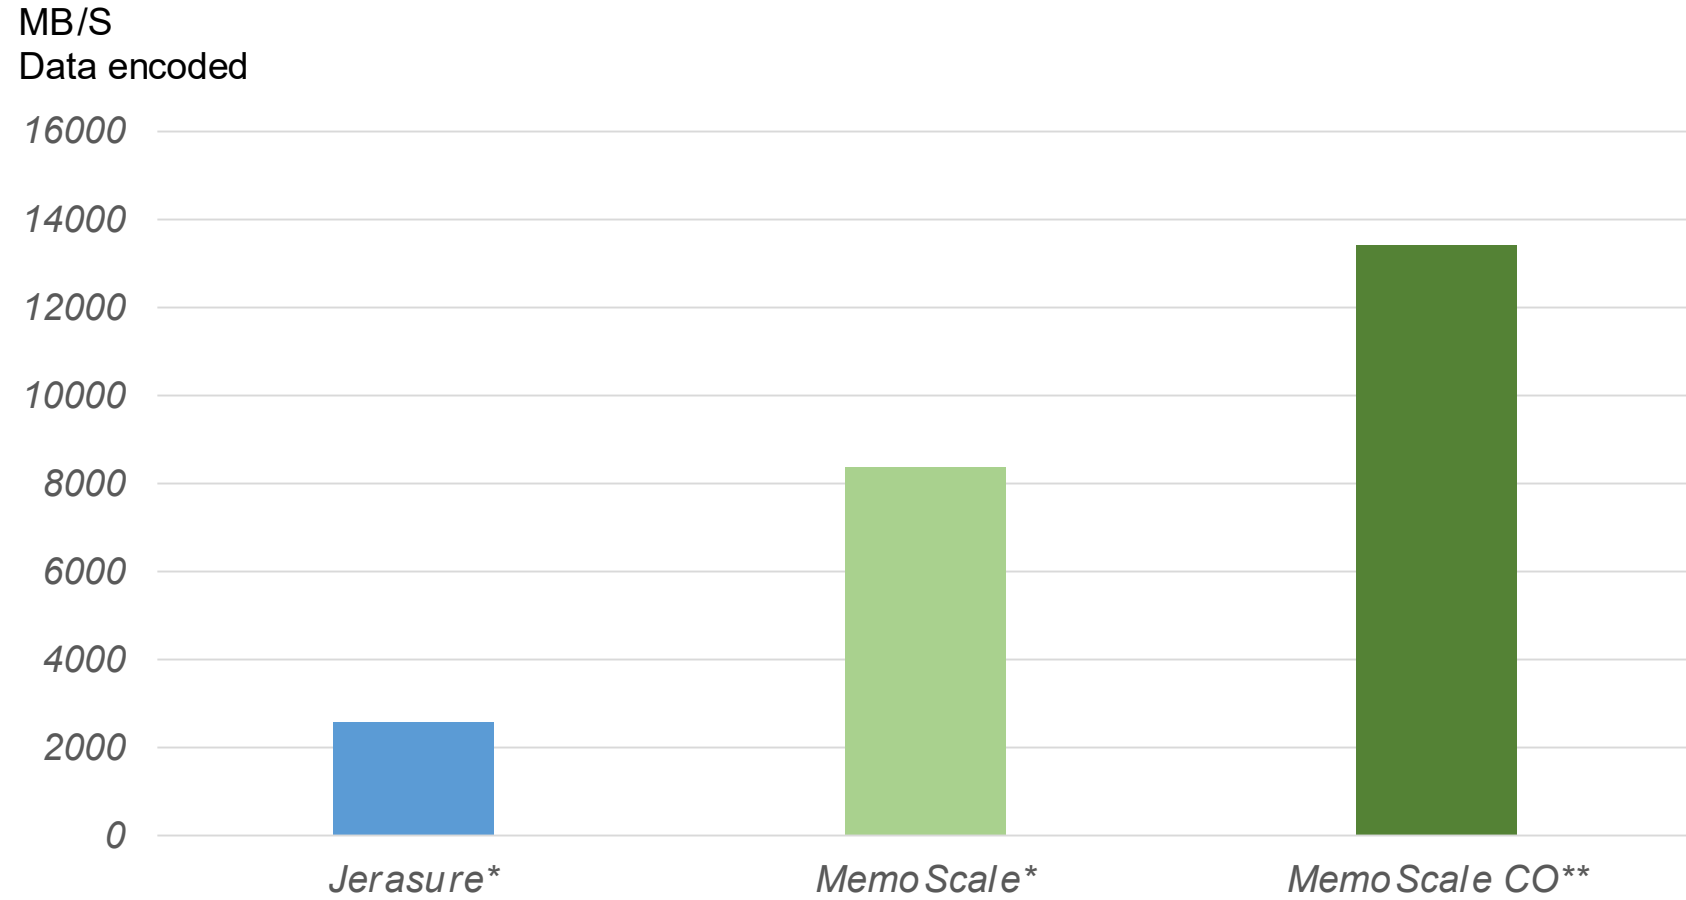
\* Erasure code: Reed Solomon \*\* Erasure code: MemoScale compute optimized



# Encoding Speed Benchmark – AMD Ryzen Processor



Processor: AMD  
RYZEN 7 1700  
8-Core 3.0 GHz  
(3.7 GHz Turbo)



\* Erasure code: Reed Solomon \*\* Erasure code: MemoScale compute optimized





# *MemoScale C-library*



- For integration in storage systems
- Erasure codes
  - Reed Solomon
  - Code with reduced recovery traffic
  - Compute optimized code
- Support and updates
- QA

# *MemoScale C-library and Plugins*



**C-library**



**SWIFT**



**HDFS**

Coming soon..



**Thank You  
for the Attention!**



*Contact:*  
***Per Simonsen - CEO***  
*MemoScale AS*

[\*per.simonsen@memoscale.com\*](mailto:per.simonsen@memoscale.com)

*Phone: +1 (888) 783-7278*

*www.memoscale.com*