

ATA

Improving Performance with Data Compression Accelerators

Juan Deaton, Ph.D.
Research Scientist and Engineer



AHA Introduction

- **Moscow, ID**
- **Relevant History**
 - 1988 NASA Startup
 - 2006 First GZIP Compression IC
 - 2014 80Gbps PCIe GZIP Accelerator
 - **Fastest GZIP Accelerator**
- **Coding Technology Experts**
 - Error Correction
 - Data Compression
 - Encryption





Are you prepared for the future?

- **Flash is the Future**


- Lower Cost. High Performance

- **Biggest Data: Industrial IoT**

- \$225 billion by 2020



Visions

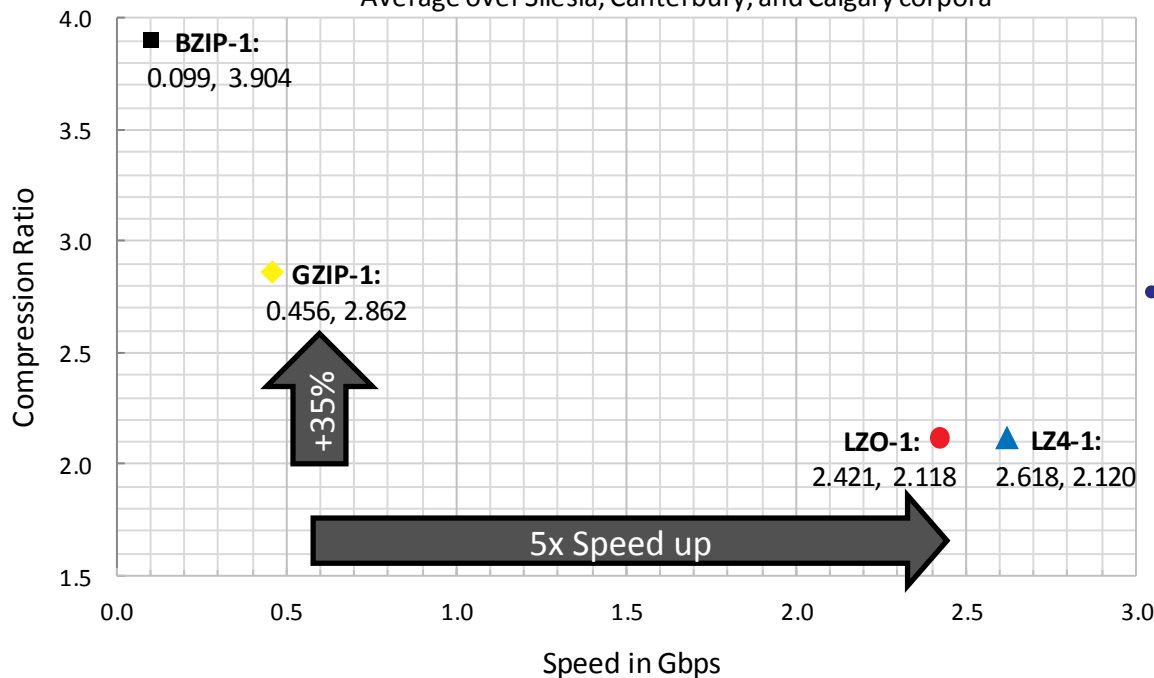


Source: © Wikibon 2015. 4-Year Cost/TB M

Data Compression Challenge

Single Core Performance

Average over Silesia, Canterbury, and Calgary corpora



- **Top Speed w/ LZ4 & 20 cores**

- 20 Cores X 2.618 Gbps = 52.4 Gbps


100% of server processing
data compression

- **Options**


- Customer buys more servers
 - High performance, lower cost?

AHA37X Family

- **Interface**
 - PCIe interface
- **Algorithms**
 - GZIP/ZLIB
- **Board by Compression Speed**
 - AHA371 – 10 Gbps
 - AHA372 – 20 Gbps
 - **AHA374 – 40 Gbps**
 - AHA378 – 80 Gbps



AHA371, AHA372



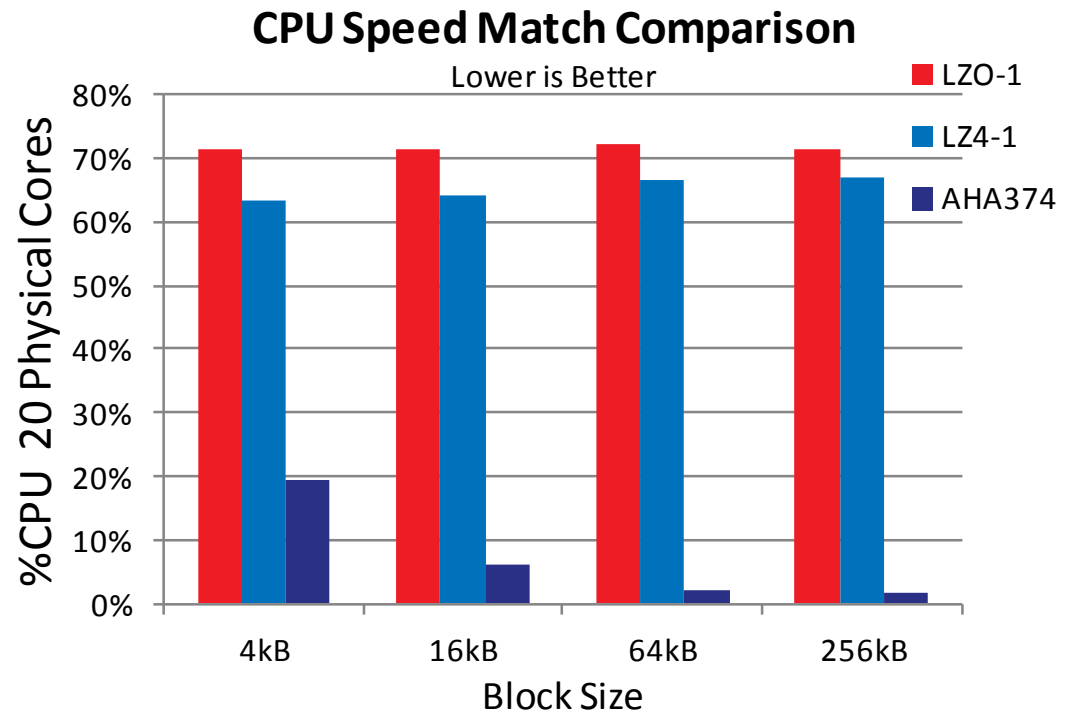
AHA374, AHA378

Experiments Summary

- **Compression Techniques Examined**
 - LZ0, LZ4, AHA374, GZIP-1
- **Compared 40Gbps/8GBps Speed Match**
 - CPU Utilization
 - Power Consumption
 - Compression Ratio
- **Block Sizes**
 - 4kB, 16kB, 64kB, 256kB
- **Hardware**
 - Server: HP Proliant DL380 Gen8
 - CPU: E5-2660v2 @ 2.20GHz
 - Dual 10 Core

Reduced CPU Utilization

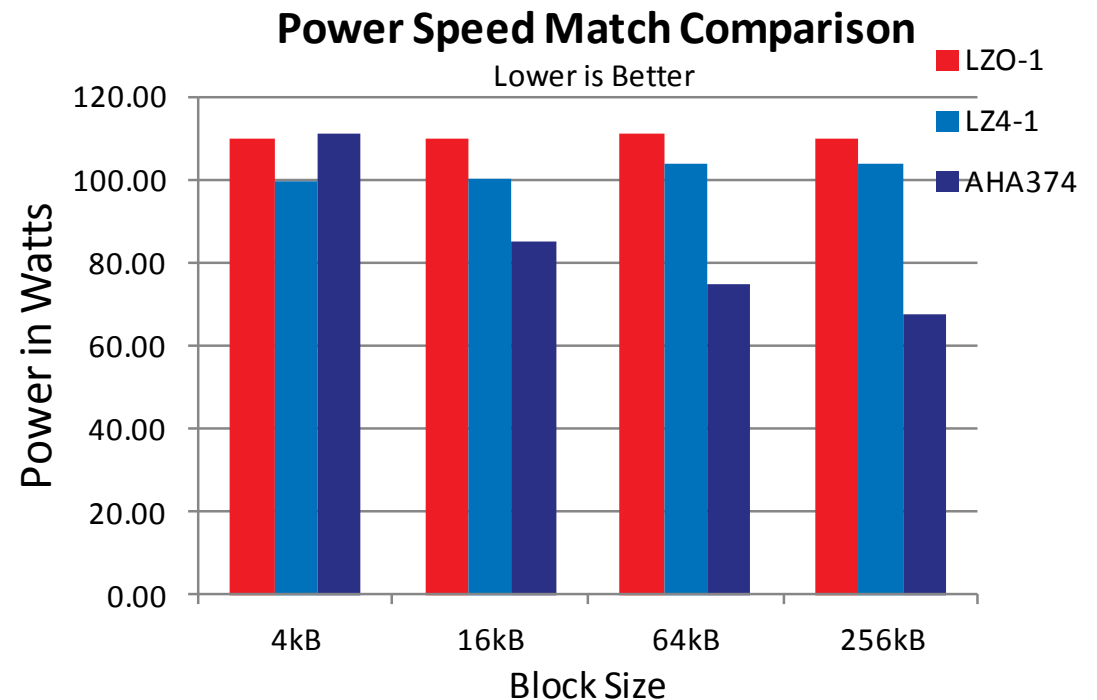
- **40 Gbps (8GBps) Speed Match**
 - Normalized for Comparison
- **%Utilization of 20 Cores**
 - 40% = 8 cores
 - 50% = 10 cores
 - 70% = 14 cores
- **GZIP/BZIP CPU Not Shown**
 - GZIP 410% - 309%
 - BZIP 4247% - 2098%



AHA 374 offloads CPU resources

Lower Average Power

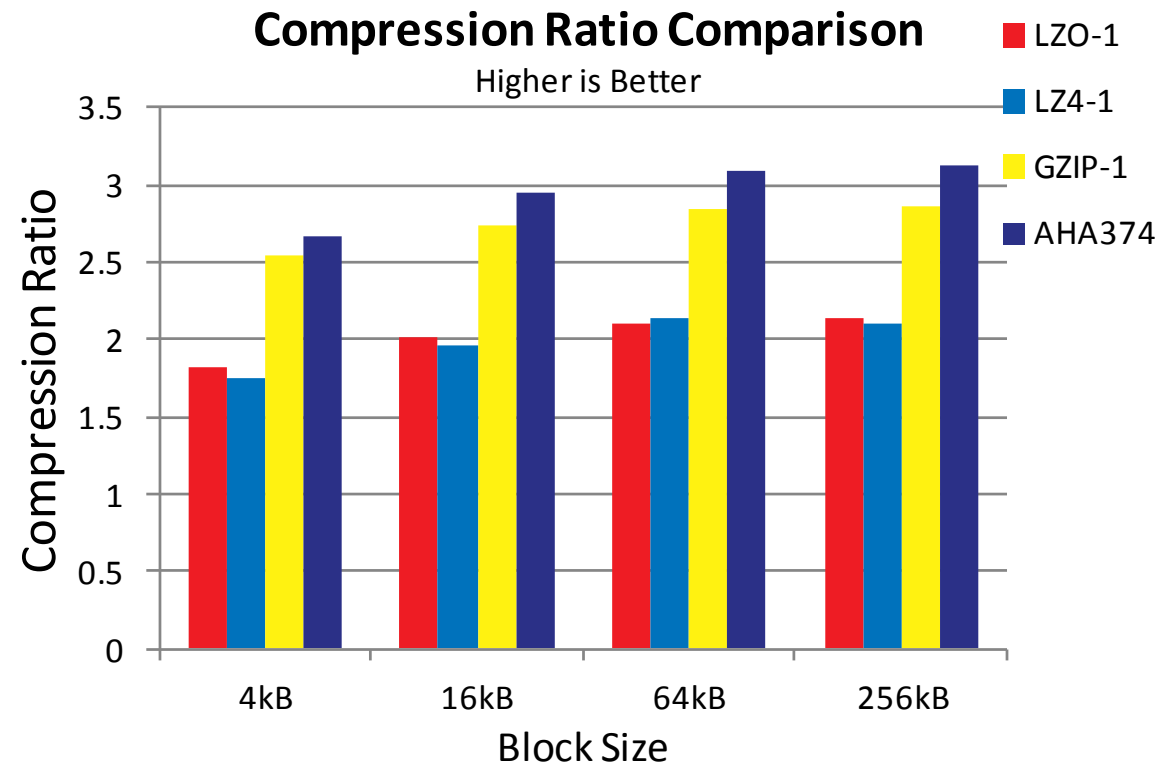
- **40 Gbps (8GBps) Speed Match**
 - Normalized for Comparison
- **Power Calculations**
 - Subtracted Idle Power
- **LZO/LZ4 Power**
 - Modeled from experiments
- **AHA374 Power**
 - Direct Measured



At >4kB blocks AHA374 uses less power

Increased Capacity

- **Average over Corpora**
 - Silesia, Canterbury, Calgary
- **4kB BZIP 2.68:1**
 - AHA374 2.66, 99% of BZIP
- **AHA374 ~GZIP-5**
 - GZIP-1 lowest GZIP CR
- **Longer Cycle Life**
 - Higher Compression Ratio



Summary

- **Industrial IoT**

- Biggest Ephemeral Data
- High Data Rates

- **AHA374 GZIP Accelerator**

- High Compression Speed
- Frees CPU Resources
- Increases Capacity and Cycle Life



Sales and Contact Information

- Website
 - www.aha.com
- Sales Contact
 - sales@aha.com
- Booth #826
 - Free Evaluations

