



Flash Memory Summit

Applying AMD ACAP to Computational Storage

Chris Stillo, AMD Storage Solutions Architect



Agenda


- Intro
- ACAP Overview
 - Platform
 - Host Interface Options
 - Processor Subsystem
- Computational Storage (CS) on ACAP Example Application
 - HW Block Diagram
 - SW Stack
 - Design Metrics
- Other CS Applications
- CS Standardization
- Conclusion

Intro: Computational Storage Motivation


- Data is
 - Valuable
 - Big today
 - Growing fast / faster than compute
- Moving Data to Compute is
 - Expensive
 - Power Hungry
 - Best to minimize transactions
- Better to use compute in/near storage where data resides
 - Bulk of data crunching happens in storage
 - Results passed up to the CPU/network




Intro: Computational Storage Architectures



Computational Storage Drive



Computational Storage Processor



Computational Storage Array

AMD ACAP – Adaptive Compute Accel. Platform



Adaptable Engines

- 4X compute density¹
- Faster timing closure



Scalar Engines

- Platform Control
- Embedded Edge Compute
- A72 and R5F



PCIe Gen 4/5 & CCIX/CXL

- 2X PCIe & DMA bandwidth
- Cache-coherency



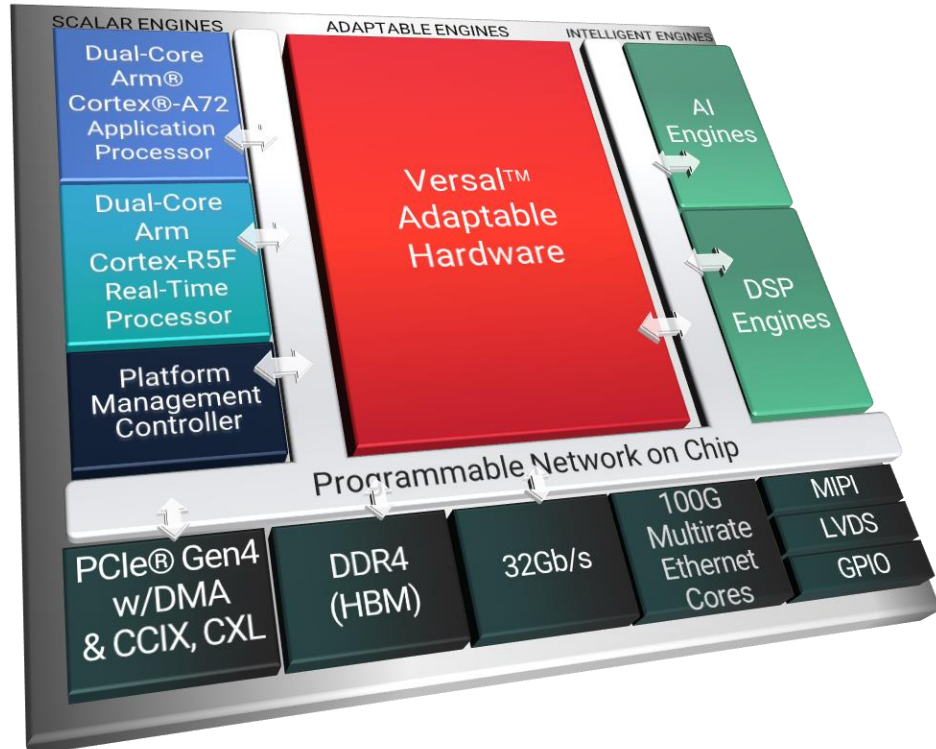
DDR4 & HBM Memory

- 3200-DDR4, 4266-LPDDR4
- 2X bandwidth/pin



Transceiver Leadership

- Broad range, 1G → 32G
- Power-optimized



Intelligent Engines

- AI Compute
- Diverse DSP Workloads



Programmable NoC

- Guaranteed Bandwidth
- Enables SW Programmability



Multi-rate Ethernet

- 10G/25G/40G/50G/100G
- Flexible, multi-standard



Programmable I/O

- Any interface or sensor
- Includes 3.2Gb/s MIPI

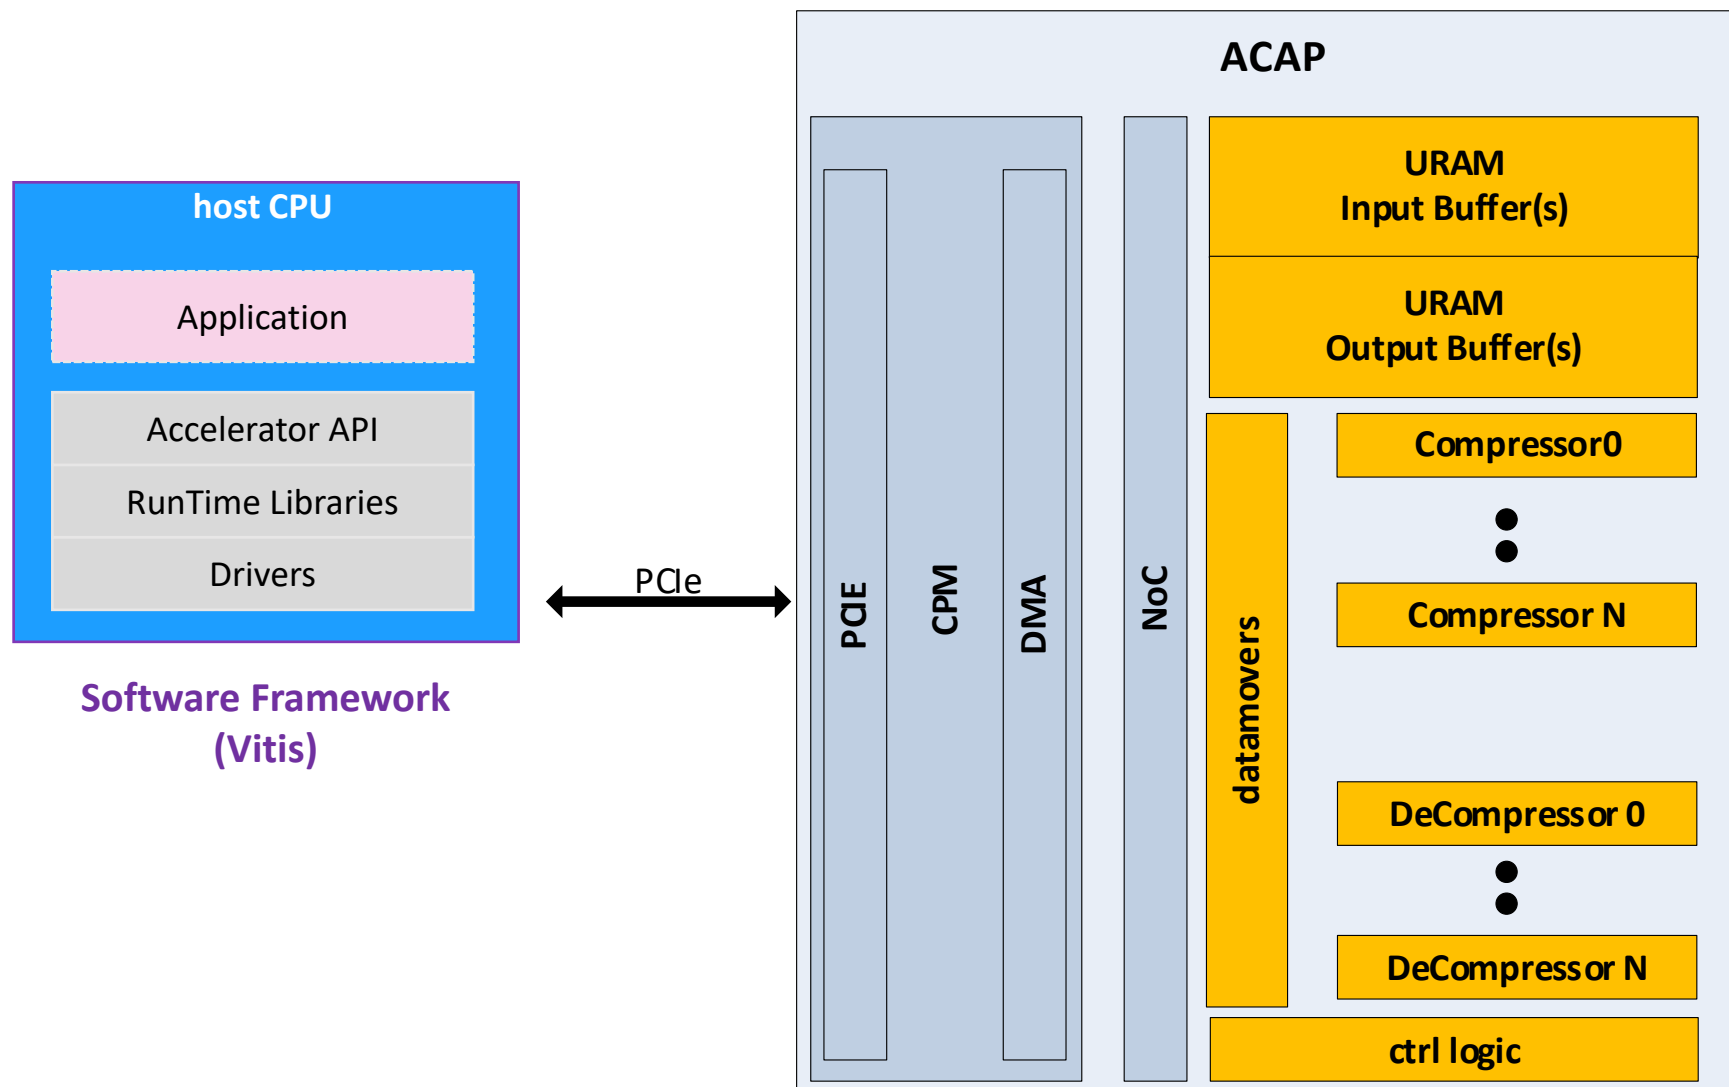


ACAP Host Interface

- PCIe EP hard logic
 - Quick time to EP bring-up... load host driver and go...
- Available NVMe Controller IP (NVMe TC)
 - Soft-logic wrapper associated with QDMA/CPM/PCIe hard-macros
 - NVMe TC IP upgrade path to support NVMe TP4091 – ROADMAP
- Thus, Varied Options for Endpoint:
 - PCIe EP
 - PCIe+NVMe EP
 - PCIe+NVMe (TP4091) EP



CS on ACAP Compr. Example: Block Diag.





AMD Vitis Data Compression IP Library (HLS)

- Vitis is the unified SW platform for AMD ACAP including drivers, APIs, etc.
- Vitis includes an open-source C/C++ library of acceleration IP, compatible with Vitis High-Level Synthesis
- Available compression IP shown at right →
- Decompression IP also available

Compression

Tables below showcases throughput details of compression for various Alveo accelerated data compression algorithms.

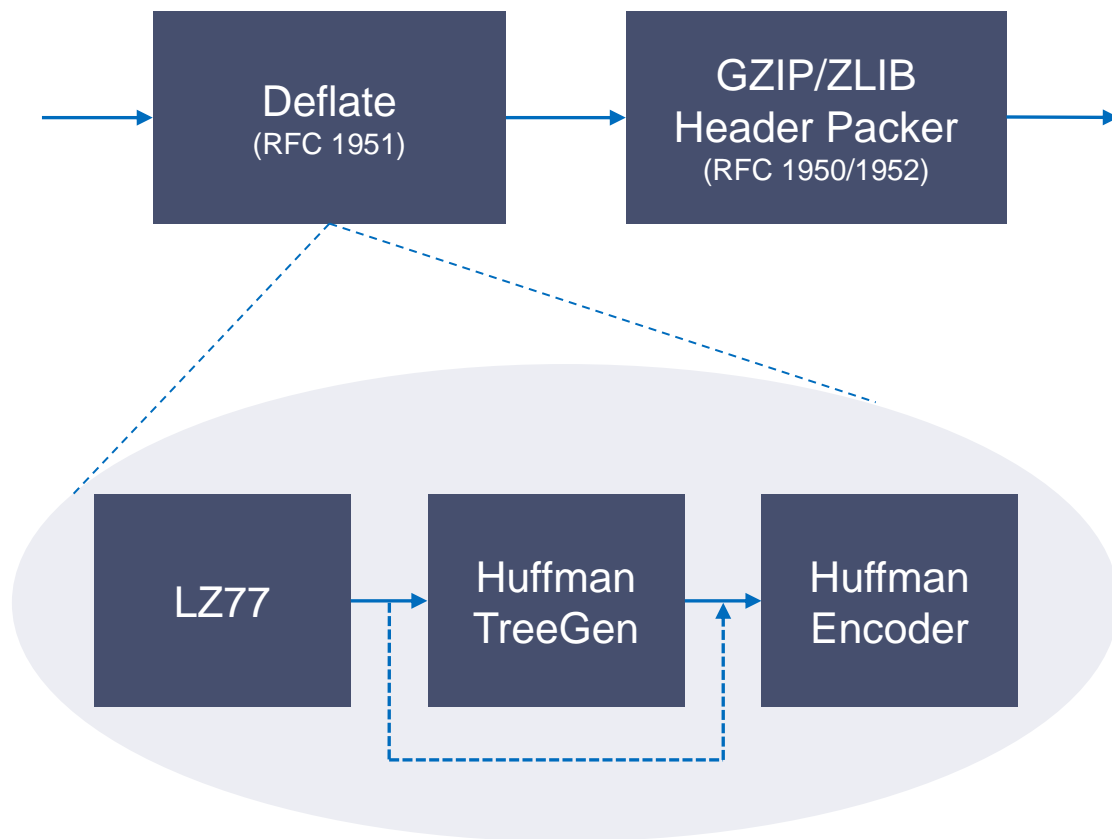
Architecture	Compression Ratio	Throughput	FMax	LUT	BRAM	URAM
LZ4 Streaming	2.13	290 MB/s	300MHz	3K	5	6
Snappy Streaming	2.13	290 MB/s	300MHz	3K	4	6
GZip/Zlib 32KB Memory Mapped	2.70	2 GB/s	300MHz	57K	135	64
GZip 32KB Compress Stream	2.70	2 GB/s	300MHz	54K	141	64
GZip 16KB Compress Stream	2.62	2 GB/s	282MHz	58K	164	48
GZip 8KB Compress Stream	2.50	2 GB/s	300MHz	57.5K	100	48
GZip Fixed 32KB Compress Stream	2.31	2 GB/s	300MHz	35K	45	64
Zlib 32KB Compress Stream	2.70	2 GB/s	300MHz	54K	128	64
Zlib 16KB Compress Stream	2.62	2 GB/s	300MHz	58K	160	48
Zlib 8KB Compress Stream	2.50	2 GB/s	300MHz	57.4K	96	48
Zlib Fixed 32KB Compress Stream	2.31	2 GB/s	300MHz	35.7K	39	64
Zstd Compress Quad Core	2.68	1.17 GB/s	284MHz	40K	79	37

- GZip/Zlib Memory Mapped and GZip/Zlib Compress Stream: Supports Dynamic Huffman

<https://www.xilinx.com/products/design-tools/vitis/vitis-libraries/vitis-data-compression.html>



Xilinx Compression IP: Deployment Options

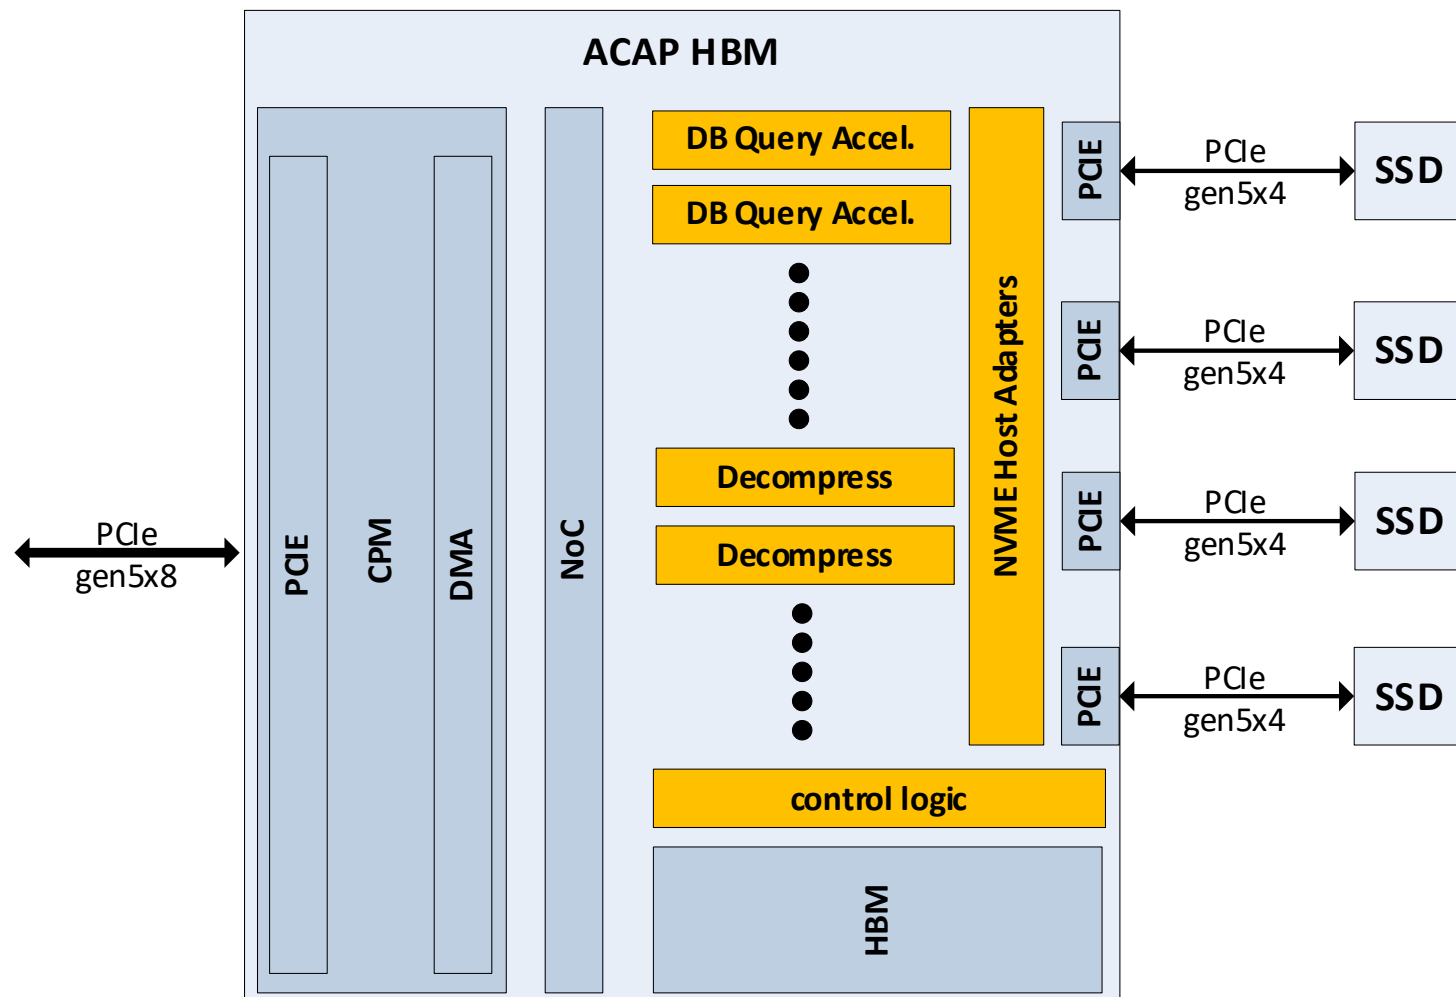


Implementing Hardware libraries and solutions to enable Compression IP targeted at the following:

- Optimizing available for:
 - “Transparent” (Low latency, high throughput)
 - Archival (High Compression ratio)
 - Target peak throughput
- Standards-based – Interoperable, software offload models.
 - LZ4
 - Zlib
 - ZSTD
- Closed system – Semi custom.
- Configurable mix of compressor/decompressor



ACAP HBM CSA Database Accelerator Concept



NOTES

- ACAP can support direct connection for up to 8 SSDs
- HBM can support up to 819 GB/s throughput and up to 32GB density
- DB Query Acceleration kernels can perform functions such as parse, scan, filter, aggregate or others
- On-board Query Accelerators reduce the amount of data transferred back to host
- Example dataflow (pipelined)
 - 4x SSDs read, write into HBM. Aggr. 48 GB/s
 - Decompress (2:1) HBM read 96 GB/s
 - Decompress output HBM write 96 GB/s
 - DB Query Acceleration read HBM 96 GB/s
 - DB Query Acceleration write buffer 100MB/s
 - DMA buffer to host memory 100MB/s

Other Potential ACAP CS Applications

- Encryption
 - https://github.com/Xilinx/Vitis_Libraries/tree/master/security
- Hashing
 - https://github.com/Xilinx/Vitis_Libraries/tree/master/security
- RegEx Search
 - Partner IP
- Erasure Coding Offload
- Transparent Compression (with partner Eideticom)
 - Demo at FMS 2022 SNIA Booth #725



Conclusion

- ACAP provides hardened logic for host interfaces, memory interfaces and interconnect, leaving programmable logic for deployment on acceleration functions in Computational Storage applications
- AMD Vitis Acceleration Libraries feature a rich variety of open-source IP, targetable to ACAP, and are open-source
- ACAP flexibility enables quickly tracking and deploying evolving Computational Storage standards

AMD Disclaimer and Attribution

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

© Copyright 2021 Advanced Micro Devices, Inc. All rights reserved. Xilinx, the Xilinx logo, AMD, the AMD Arrow logo, Alveo, Artix, Kintex, Kria, Spartan, Versal, Vitis, Virtex, Vivado, Zynq, and other designated brands included herein are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

