SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Computational Storage – Driving Success, Driving Standards

A SNIA Webcast Discussion with
Bill Martin, Samsung; Jason Molgaard, Arm;
Oscar Pinto, Samsung; Scott Shadley, NGD Systems

Live October 26, 2021 at 10:00 am PDT

Today's Speakers

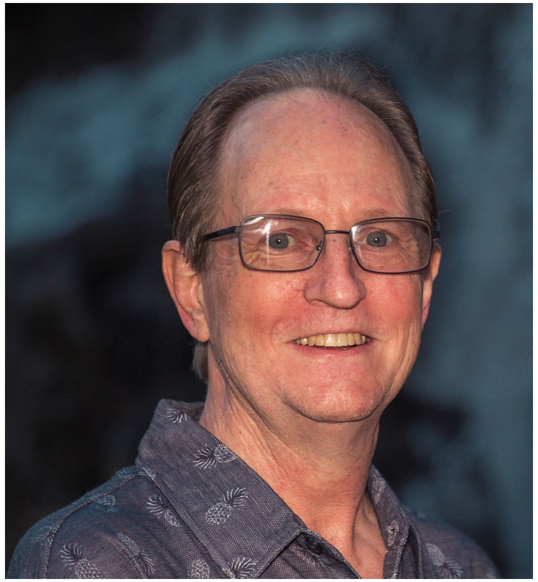

Bill Martin
Samsung Semiconductor Inc.
Editor, SNIA Computational Storage Architecture and Programming Model

Jason Molgaard
Arm
Co-Chair, CMSI Computational Storage Technical Work Group

Oscar Pinto
Samsung Semiconductor Inc.
Editor, SNIA Computational Storage API

Scott Shadley
NGD Systems
Co-Chair, CMSI Computational Storage Technical Work Group

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.

SNIA-at-a-Glance

185
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

SNIA. CMSI | COMPUTE, MEMORY, AND STORAGE

Computational Storage - A Quick History and Status

Common Language, Common Goals

- The challenge with driving new technology can be the convolution of data
  - The ability to say the same thing with different words
- Computational Storage had many names – back as far as 2010
  - Scale-In
  - In-Situ Processing
  - Compute to Data
  - In-Data Processing
- A change to the taxonomy model was needed and a SNIA Technical Work Group (TWG) was formed

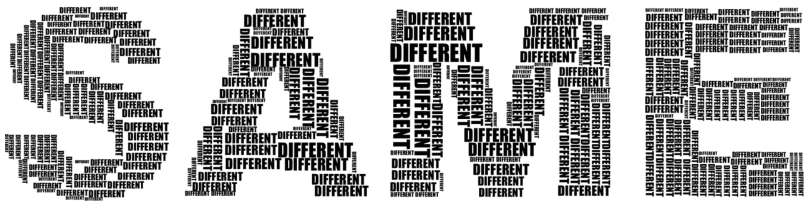

The Ongoing Work of SNIA to Define Standards

- TWG Working group is continuing to see growth
  - Member count is up, Users ‘following’ and ‘participating’
  - 51 companies, 261 individual members

51 Participating Companies - 261 Member Representatives

- Work in the Special Interest Group
  - CS SIG – Webinars, Blogs, Events
- Collaborating with other Groups
  - NVM Express – Computational Programs
  - Be sure to check out the [Storage Developer Conference session presented by the Co-Chairs on that work]

The Efforts to Get Information Out is Continuing

- ComputerWeekly.com
  - 13-part Series

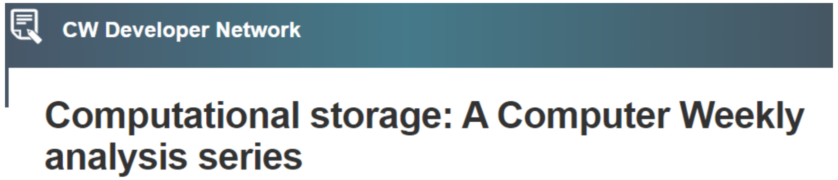

- Gartner Analysts
  - 2018 and 2021 'Cool Vendor'
  - Hype Cycle Entry

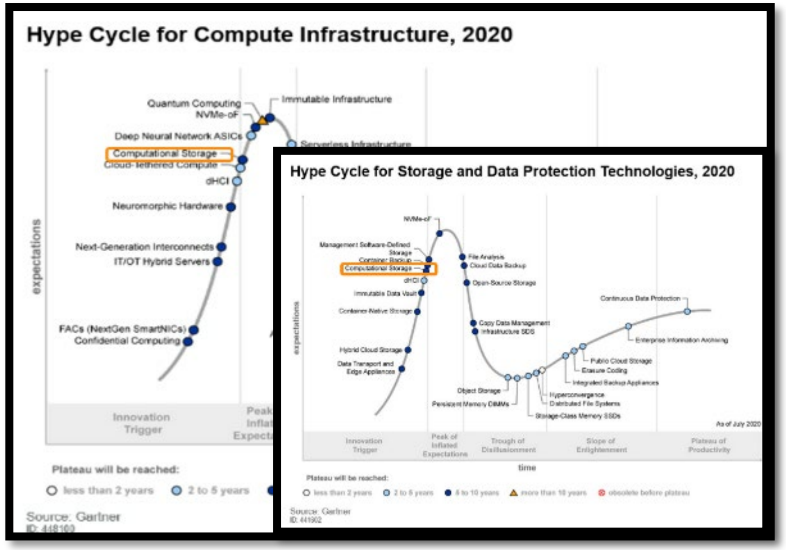

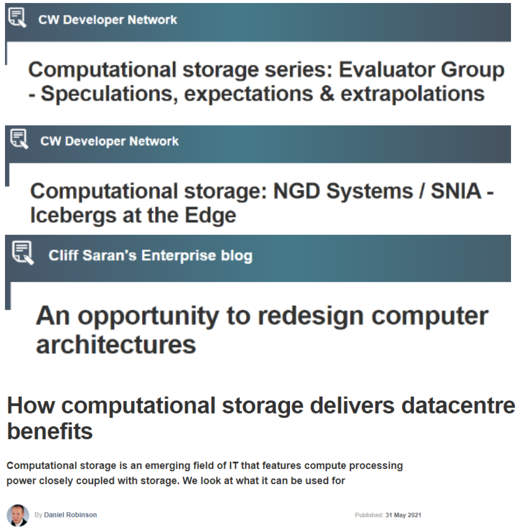

- Sponsored Efforts

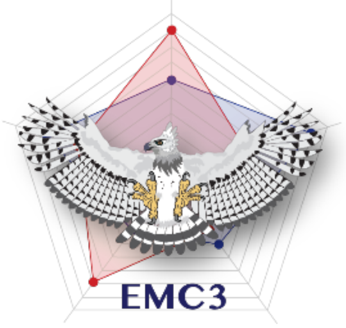

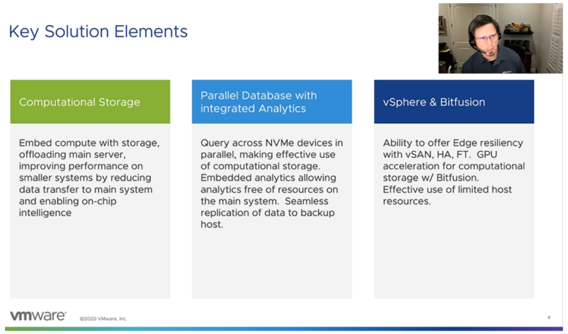

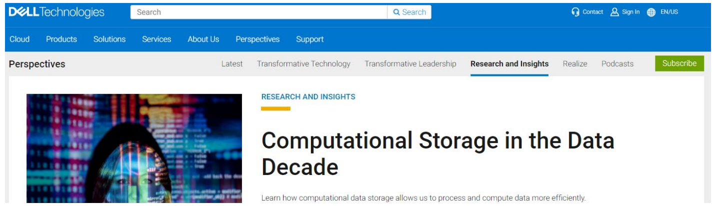

SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Computational Storage - The Work Being Done

A Potential Use Case for Computational Storage

- Generate Metadata database (e.g tags) over a large set of unstructured data locally stored on the drive, with an integrated AI inference engine
- Operation may be:
  - Triggered by a host processor
  - Done offline as a background task (batches)
- Metadata database may be then used by upper layer Big Data Analytics software for further processing
- Can work both on direct attached storage or on remote over the network storage
- Examples: Video search, Ad insertion, Voice call analysis, Images, Text scan, etc

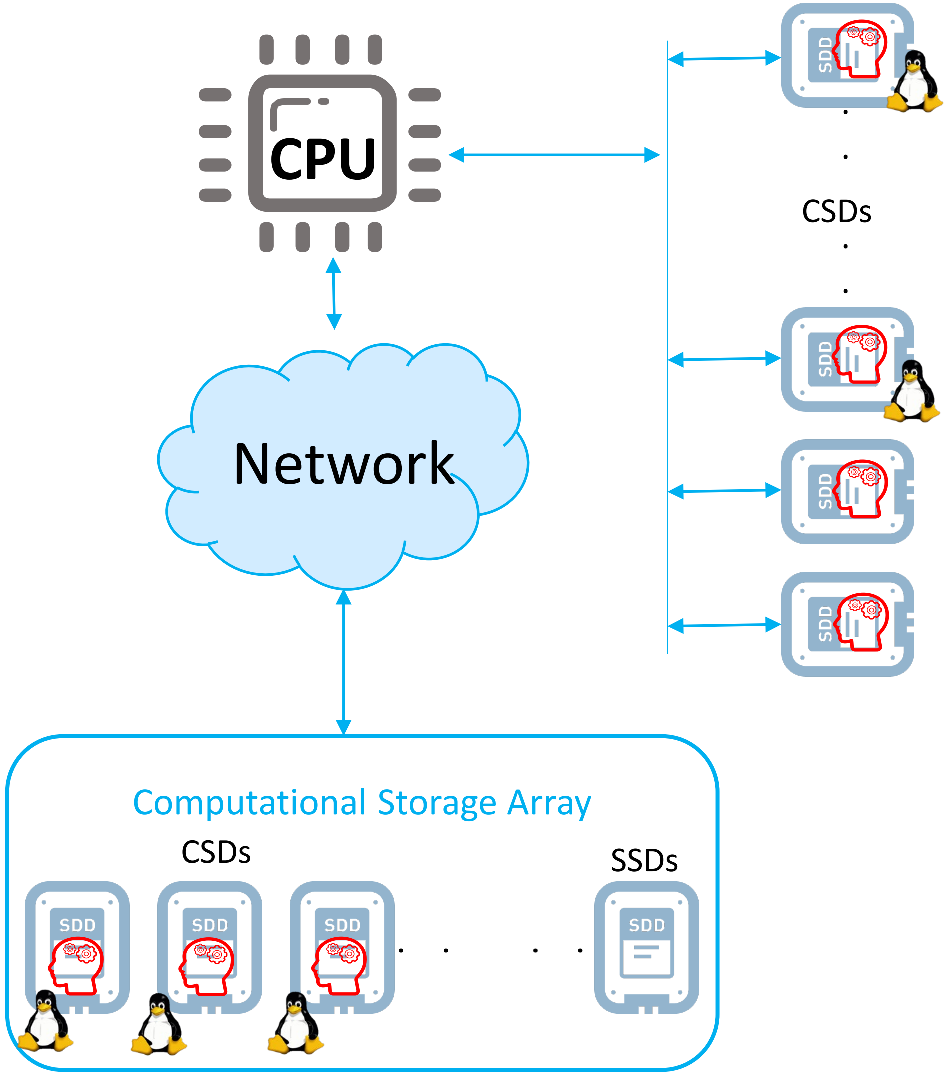

 |

Current Progress of TWG Output

- Architectural Document has been released
  - V0.8 is now in Public Review
  - Many updates from 0.5
- API Document has been released
  - V0.5 is now in Public Review
- Security now being reviewed
  - In collaboration with SNIA Security TWG
- Today we'll be speaking about Architecture and API

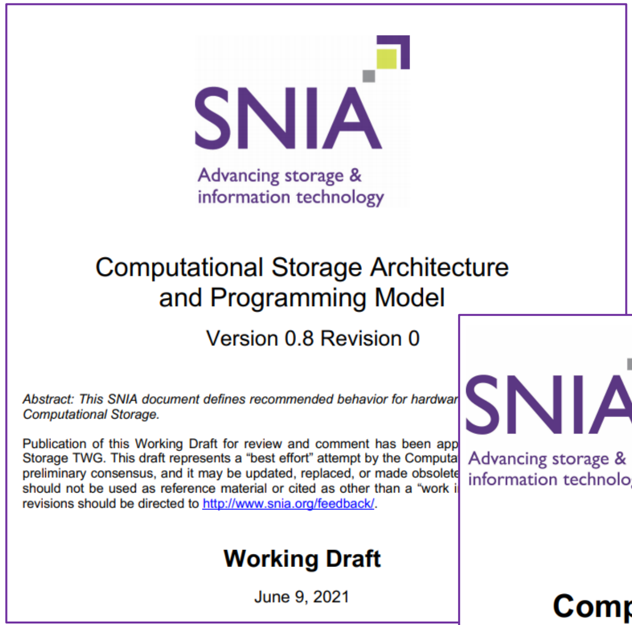

SNIA
Advancing storage & information technology

Computational Storage Architecture and Programming Model

Version 0.8 Revision 0

Abstract: This SNIA document defines recommended behavior for hardwar… Computational Storage.

Publication of this Working Draft for review and comment has been app… Storage TWG. This draft represents a "best effort" attempt by the Computa… preliminary consensus, and it may be updated, replaced, or made obsolete… should not be used as reference material or cited as other than a "work i… revisions should be directed to http://www.snia.org/feedback/.

Working Draft

June 9, 2021

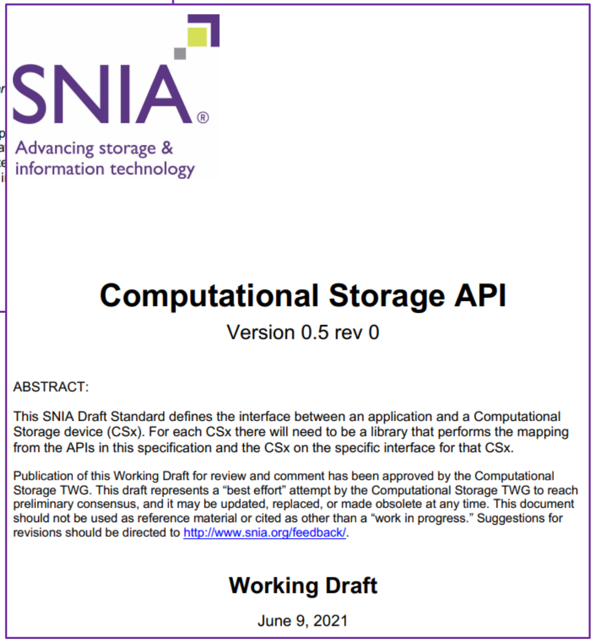

SNIA®
Advancing storage & information technology

Computational Storage API

Version 0.5 rev 0

ABSTRACT:

This SNIA Draft Standard defines the interface between an application and a Computational Storage device (CSx). For each CSx there will need to be a library that performs the mapping from the APIs in this specification and the CSx on the specific interface for that CSx.

Publication of this Working Draft for review and comment has been approved by the Computational Storage TWG. This draft represents a "best effort" attempt by the Computational Storage TWG to reach preliminary consensus, and it may be updated, replaced, or made obsolete at any time. This document should not be used as reference material or cited as other than a "work in progress." Suggestions for revisions should be directed to http://www.snia.org/feedback/.

Working Draft

June 9, 2021

The Taxonomy of Computational Storage

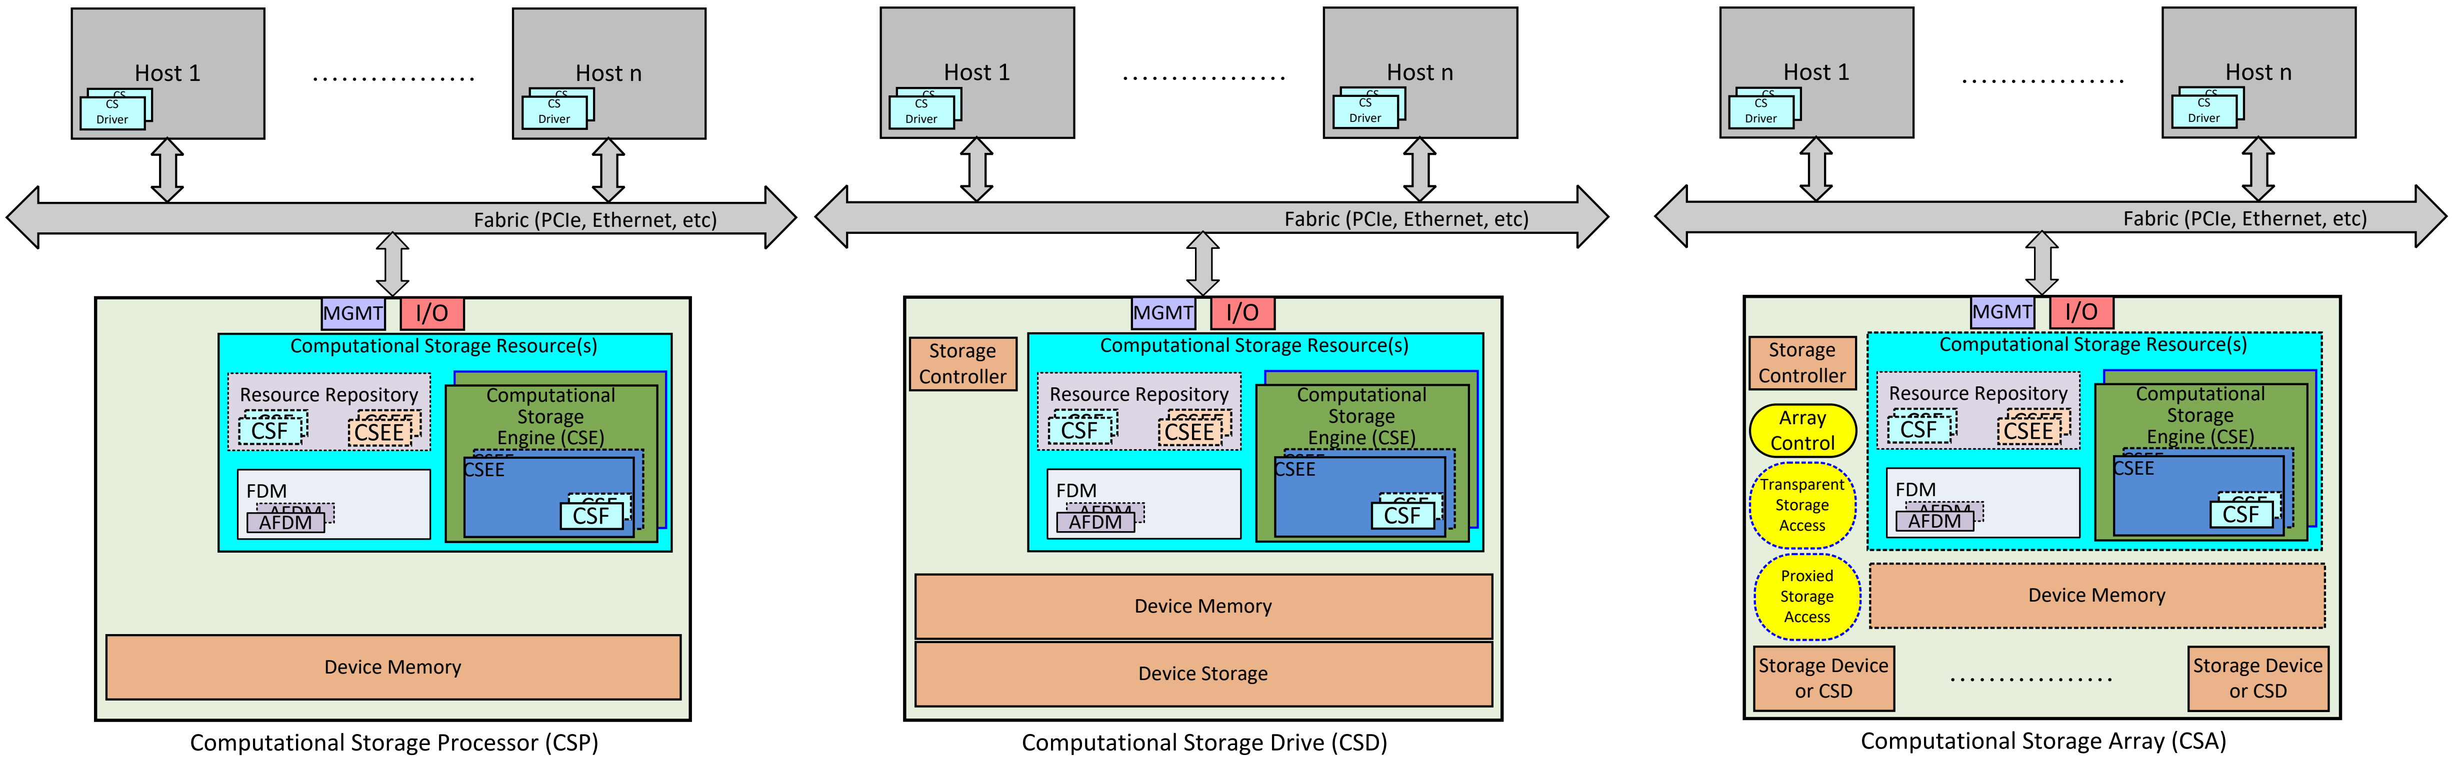

Computational Storage Architecture

Implementing The Taxonomy of Computational Storage

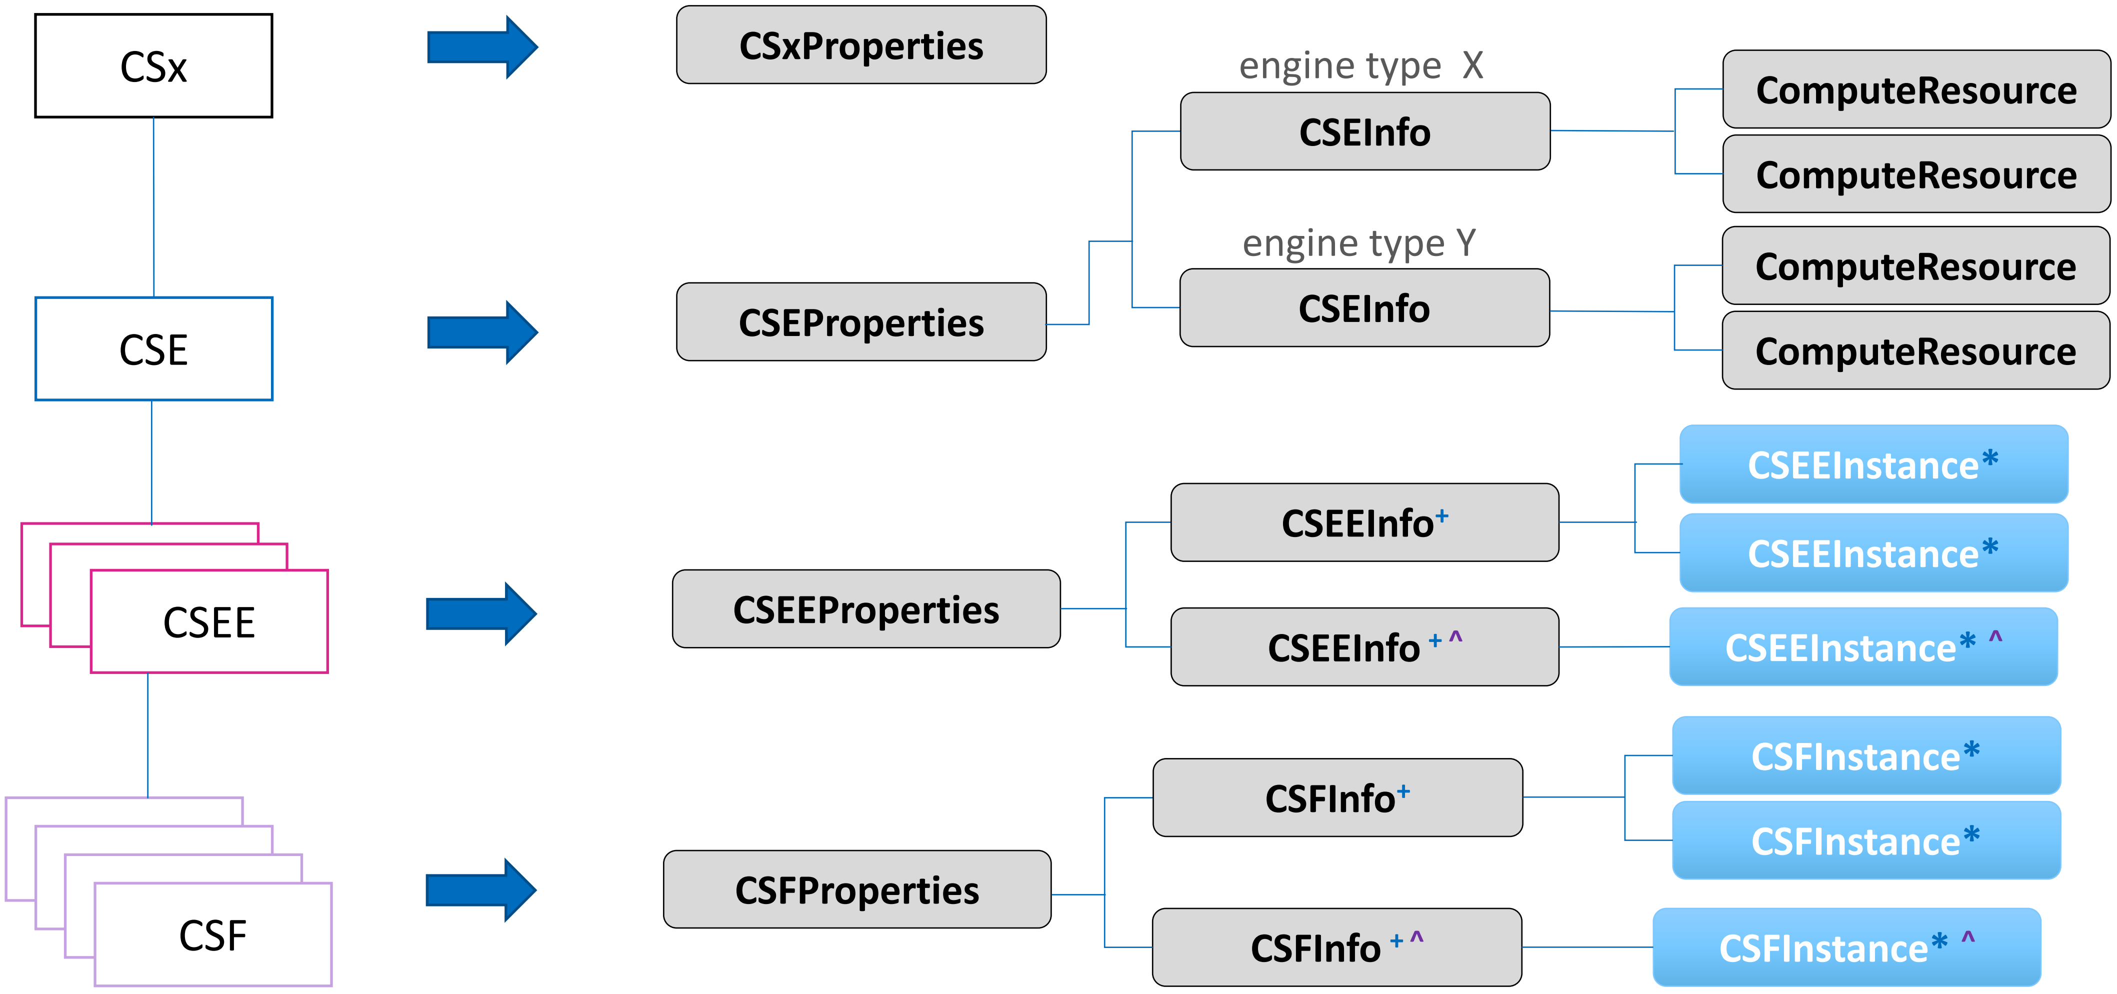

***CSFInstance, CSEEInstance** – activated for usage*
***+CSEEInfo, CSFInfo** – each in repository*

***^CSEEInfo, CSEEInstance** – hard-coded in CSE*
***^CSFInfo, CSFInstance** – hard-coded in CSEE*

Direct Computational Storage Implementation

- Assumption
  - Data on which computation is to be performed is placed in the FDM, prior to the request to the CSE, through some process that is not shown in this figure
  - Result data, if any, is returned to the host through some process that is not shown in this figure
- Process
  1. The host sends a command to the CS controller to invoke the CSF;
  2. The CSE performs the requested computation on data that is in FDM and places the result, if any, into FDM; and
  3. The CSE returns a response to the host.

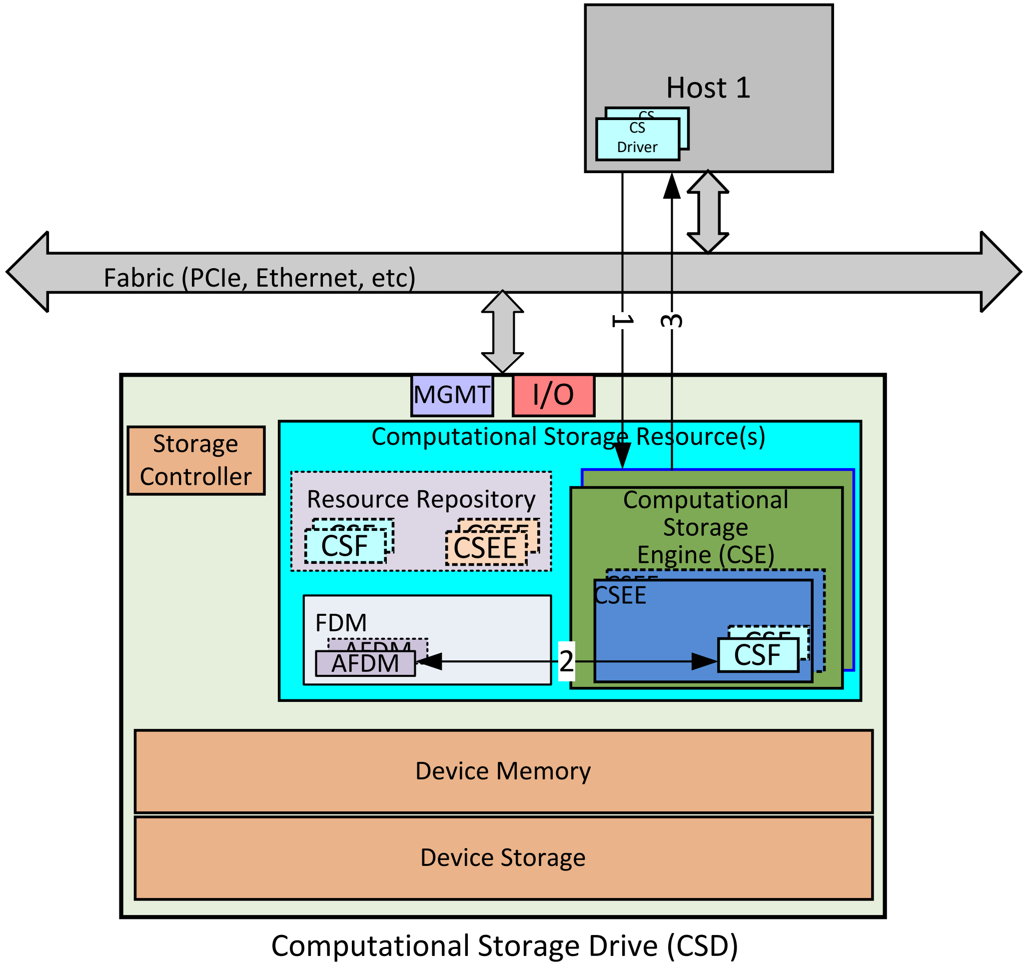

Computational Storage on Device Data

- Assumption
  - This example is for a computation on data that is in device storage
- Process
  1. The host sends a command to invoke the CSF;
     a. The command specifies the Device Storage location of the data;
  2. The CSE moves data from Device Storage to FDM;
  3. The CSE performs the requested computation on data that is in FDM and places the result, if any, into FDM; and
  4. The CSE returns a response to the host.

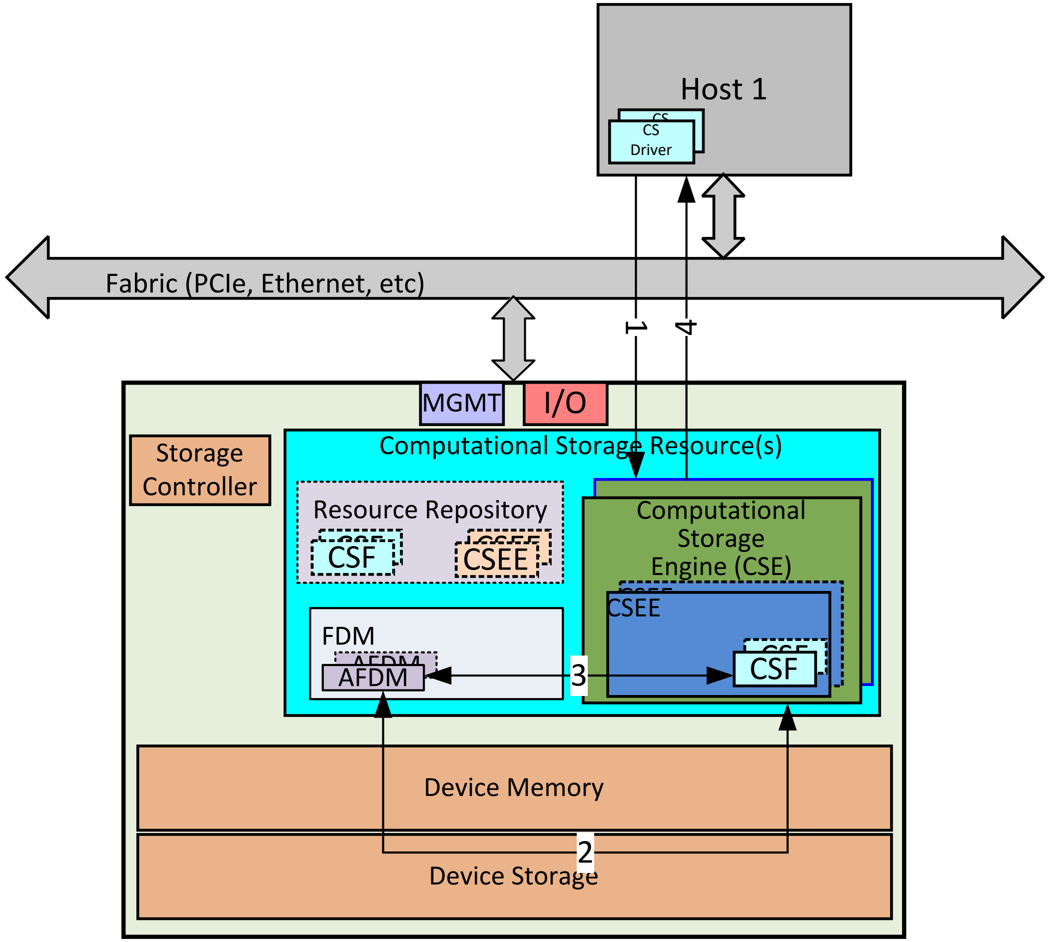

Computational Storage Drive (CSD)

SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Indirectly Using Computational Storage on Device Data

- Assumption:
  - This example is for a device to host operation
- Process
  1. The host sends a storage request to a Storage Controller where:
     a. that storage request is associated with a target CSF; and
     b. the storage controller determines what CSF is associated with the storage request;
  2. The Storage Controller moves data from storage into the FDM;
  3. The Storage Controller instructs the CSE to perform the indicated computation on the data in the FDM;
  4. The CSE performs the computation on the data and places the result, if any, into the FDM; and
  5. The Storage Controller returns the computation results, if any, from the FDM to the host.

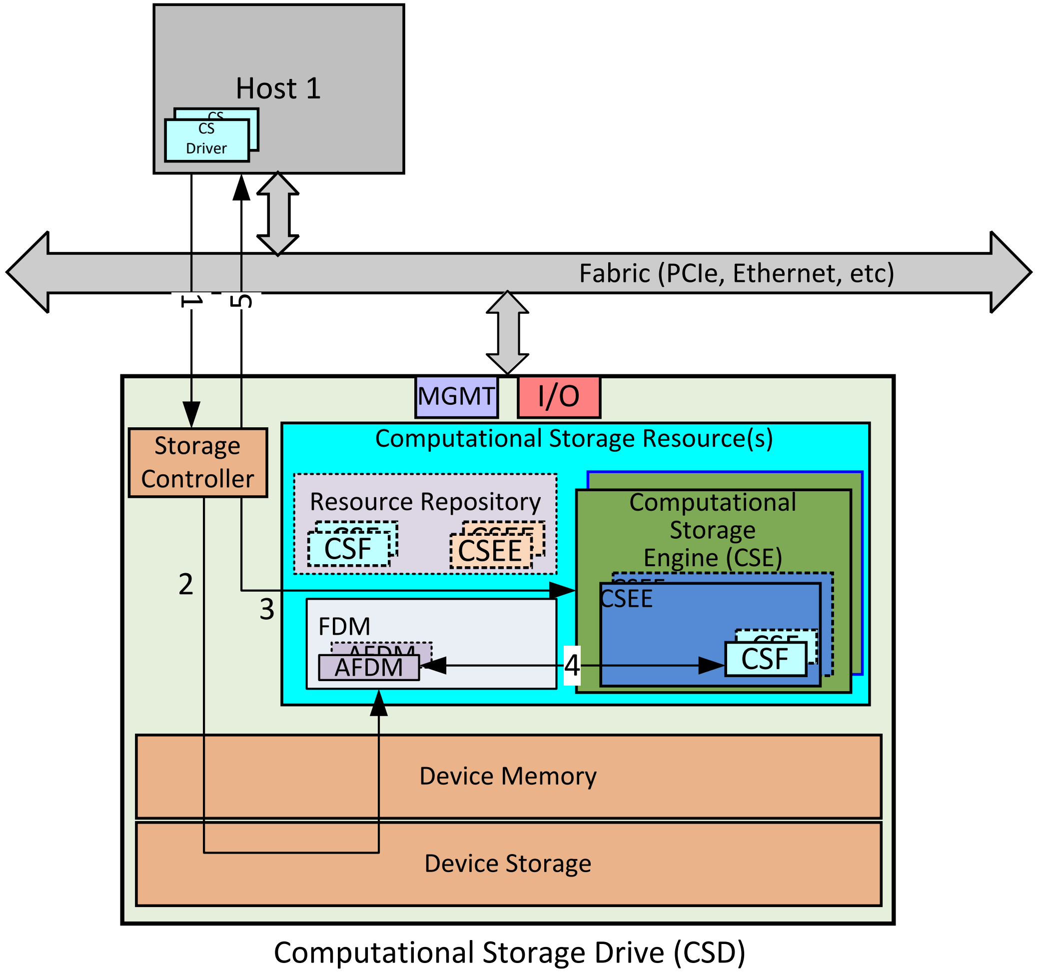

SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Computational Storage APIs

CS API Library Overview

- One Set of APIs across all CSx types
  - CSP, CSD, CSA
  - Common set of APIs for different CS devices
- One interface to different device and connectivity choices
  - Hardware ASIC, CPU, FPGA, etc
  - NVMe/NVMe-oF, PCIe, custom, etc
- Configurations may be local/remote attached
- Hides vendor specific implementation details below library
- Abstracts device specific details
- APIs to be OS agnostic

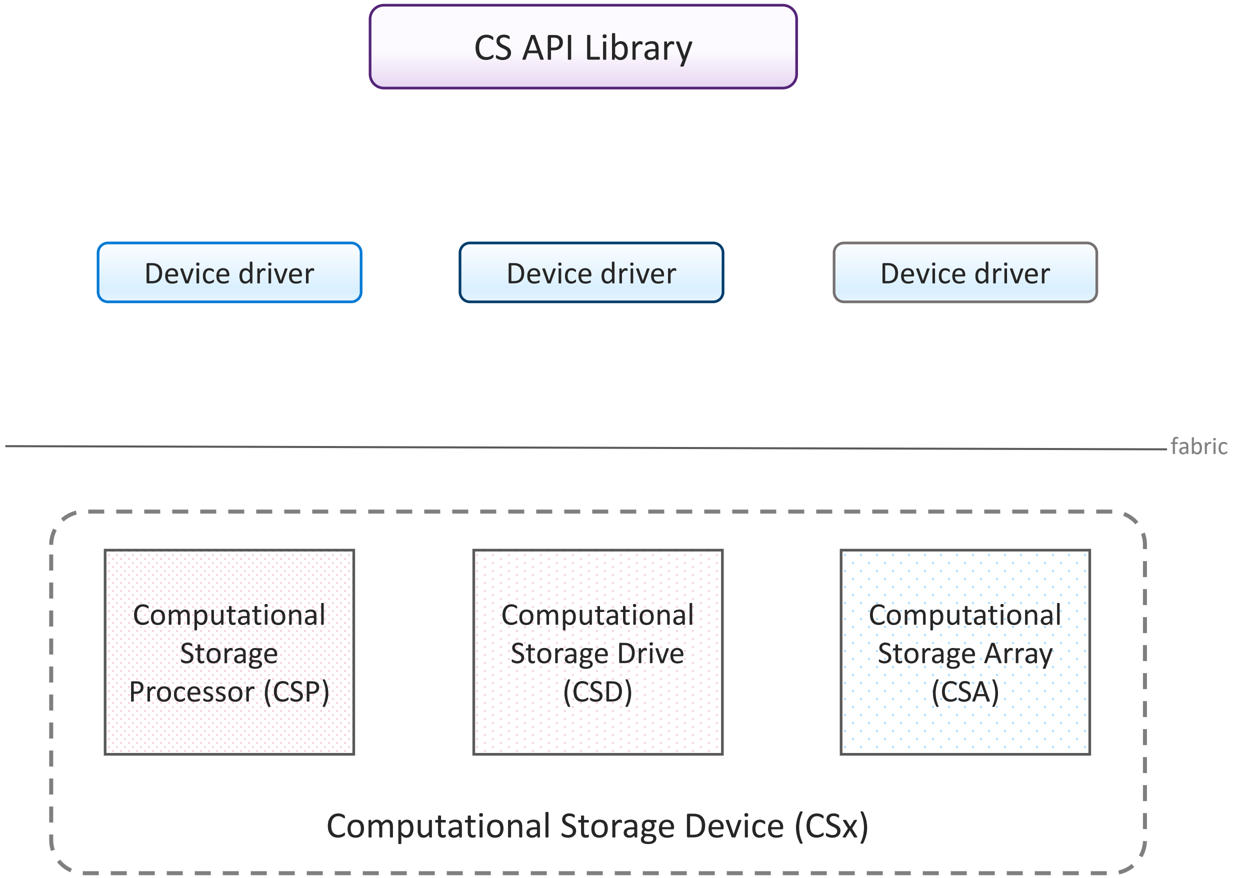

SNIA. CMSI | COMPUTE, MEMORY, AND STORAGE

About API Library

- Uniform interface for multiple configurations
  - APIs provided in common library
- Each CSx managed through its own device stack
  - Library may interface with additional plugins based on implementation requirements
  - Plugins help connect a CSx to abstracted CS interfaces
- Extensible Interface
- API Requirements
  - One interface across CS devices: CSP, CSD, CSA
  - Discovery
  - Device Access
  - Device Memory (mapped/unmapped) allocations
  - Near Storage Access
  - Copy Device Memory
  - Download Functions (CSFs)
  - Execute CSFs
  - Device Configuration & Management
  - Security

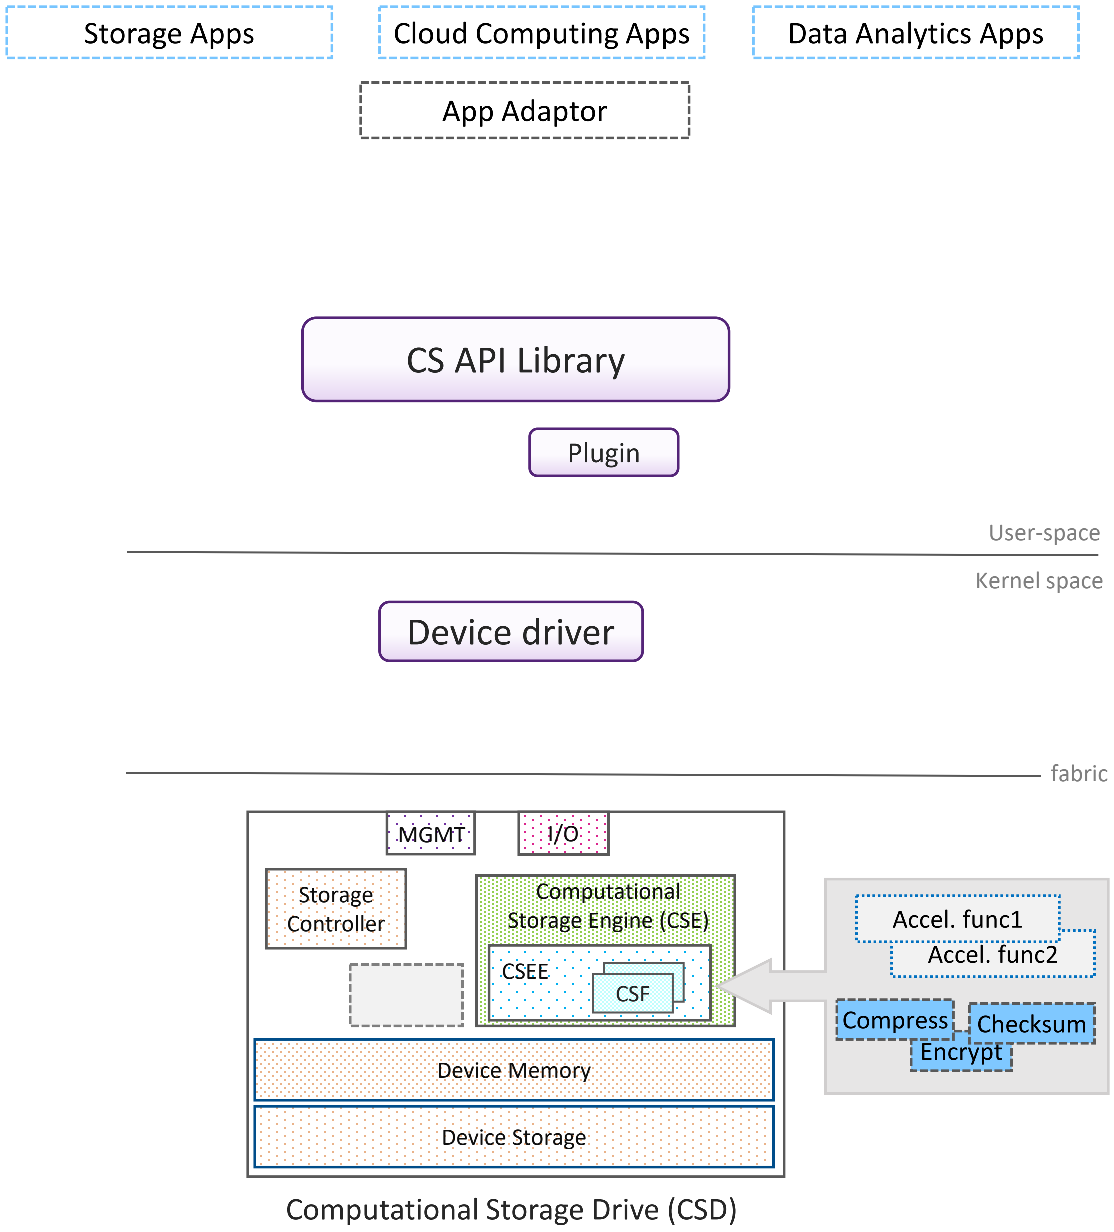

SNIA CMSI | COMPUTE, MEMORY, AND STORAGE

Applying Computational Storage

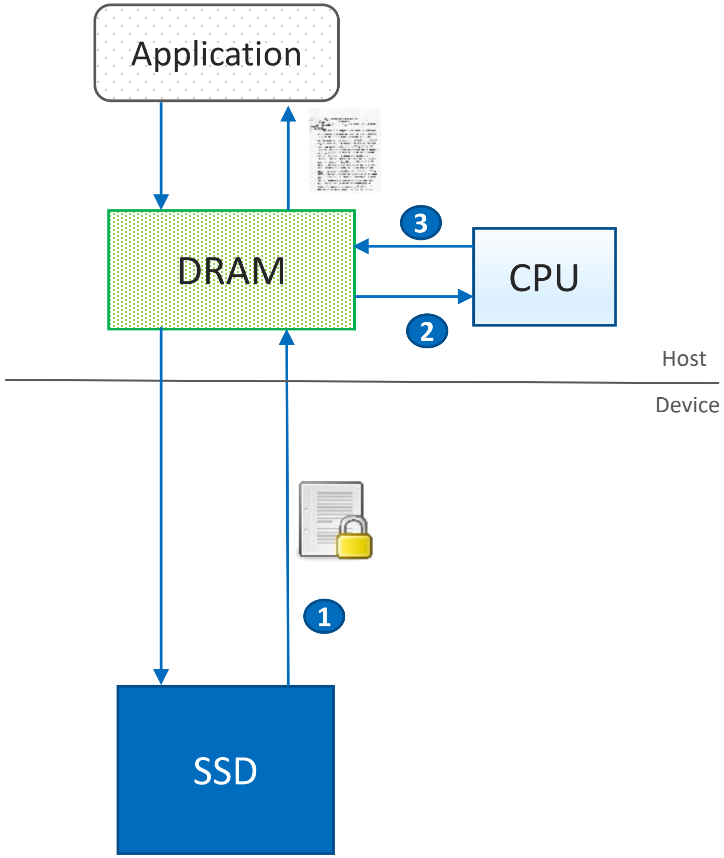

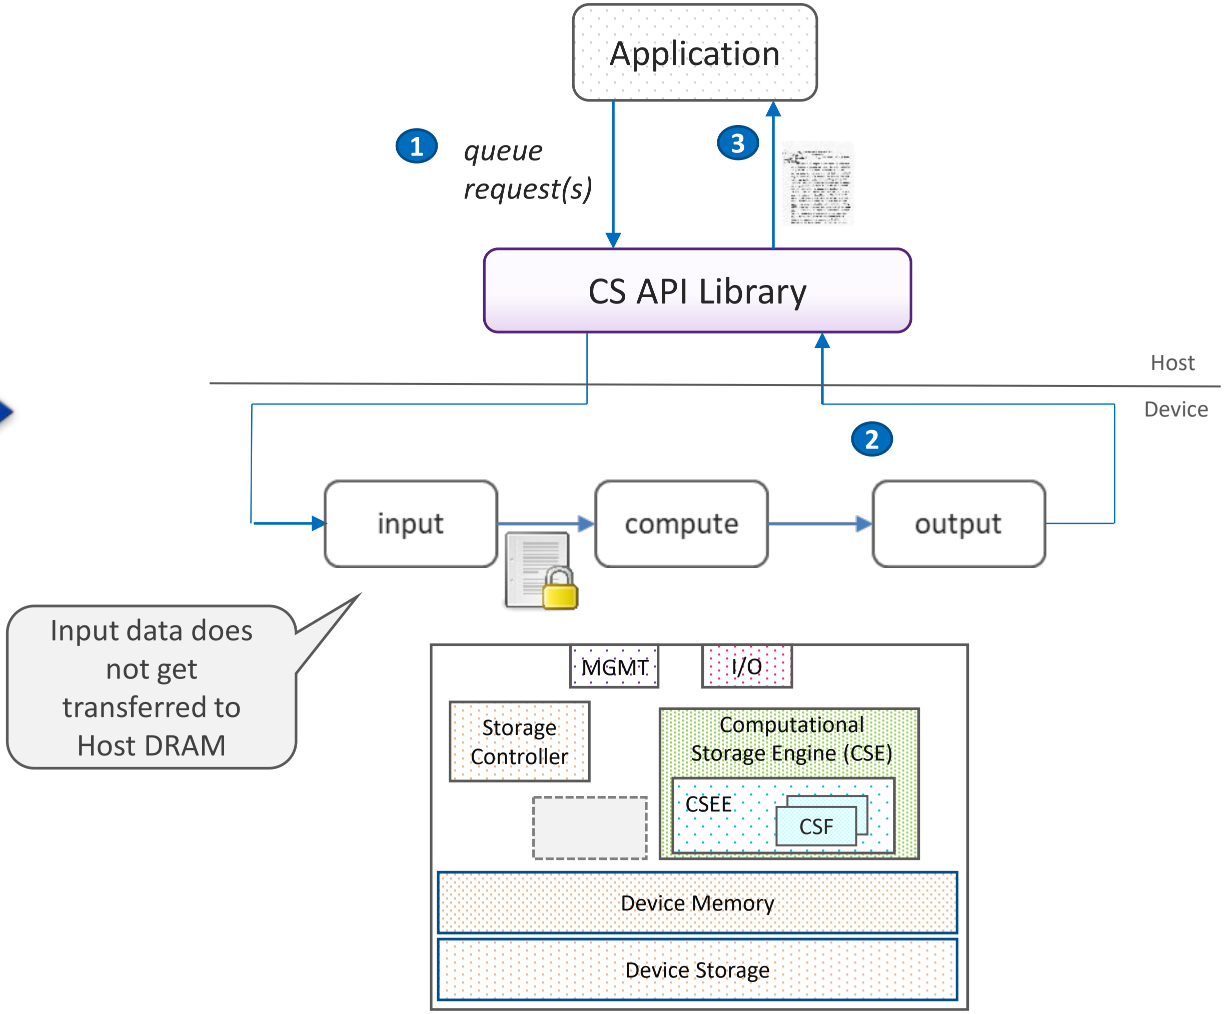

Computational Storage Drive (CSD)

 |

Example with CS APIs

1. Discover CSx & CSF
2. Allocate Device Memory
3. Queue Storage Request
4. Queue Compute Request
5. Queue Copy Memory Request

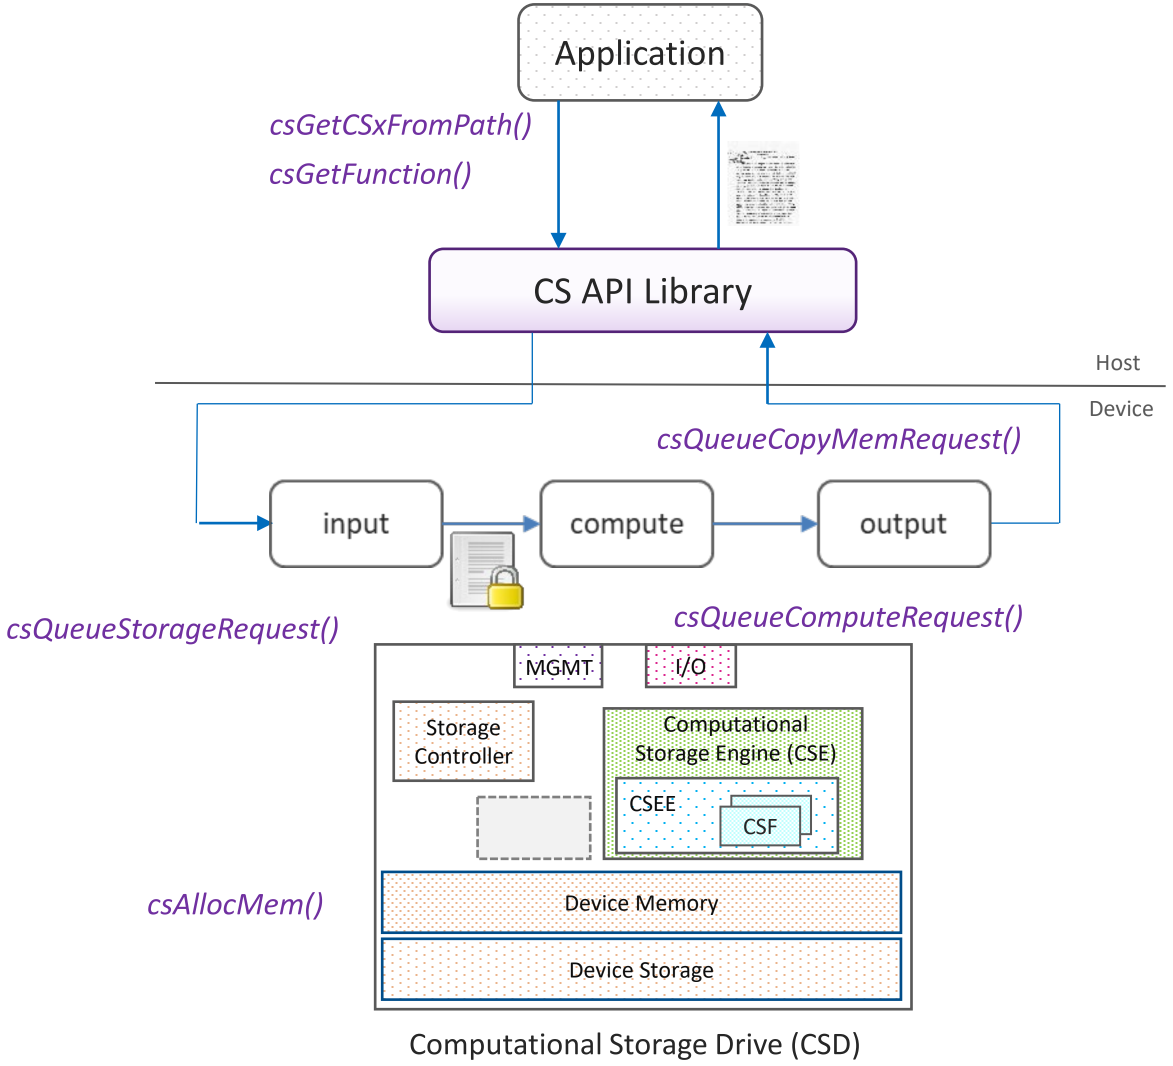

What Next?

Explore SNIA Computational Storage Activities

- SNIA Computational Storage Technical Work Group
  - Actively working on establishing hardware and software architectures to allow for compute to be more tightly coupled with storage at the system and drive level
- SNIA Computational Storage Special Interest Group
  - Fostering the acceptance and growth of computational storage in the marketplace
- SNIA Computational Storage Architecture and Programming Model v0.8 rev 0
  - Defines recommended behavior for hardware and software that supports Computational Storage
- SNIA Computational Storage API v0.5 rev 0
  - Defines the interface between an application and a Computational Storage device (CSx)

Is your company a SNIA member? ***Find out*** ***…….and*** ***contribute to the specifications***

Implementing computational storage? ***Provide a public feedback comment***

Thanks for Watching Our Webcast

- Please rate this webcast and provide us with feedback
- A link to this webcast and the PDF of the slides are posted to the SNIA Compute Memory and Storage Initiative website at https://www.snia.org/forums/cmsi/knowledge/articles-presentations
- You can also find this webcast and many other videos and presentations on today's topics in the SNIA Educational Library
- A Q&A from this webcast will be posted to the SNIA Compute, Memory, and Storage Blog

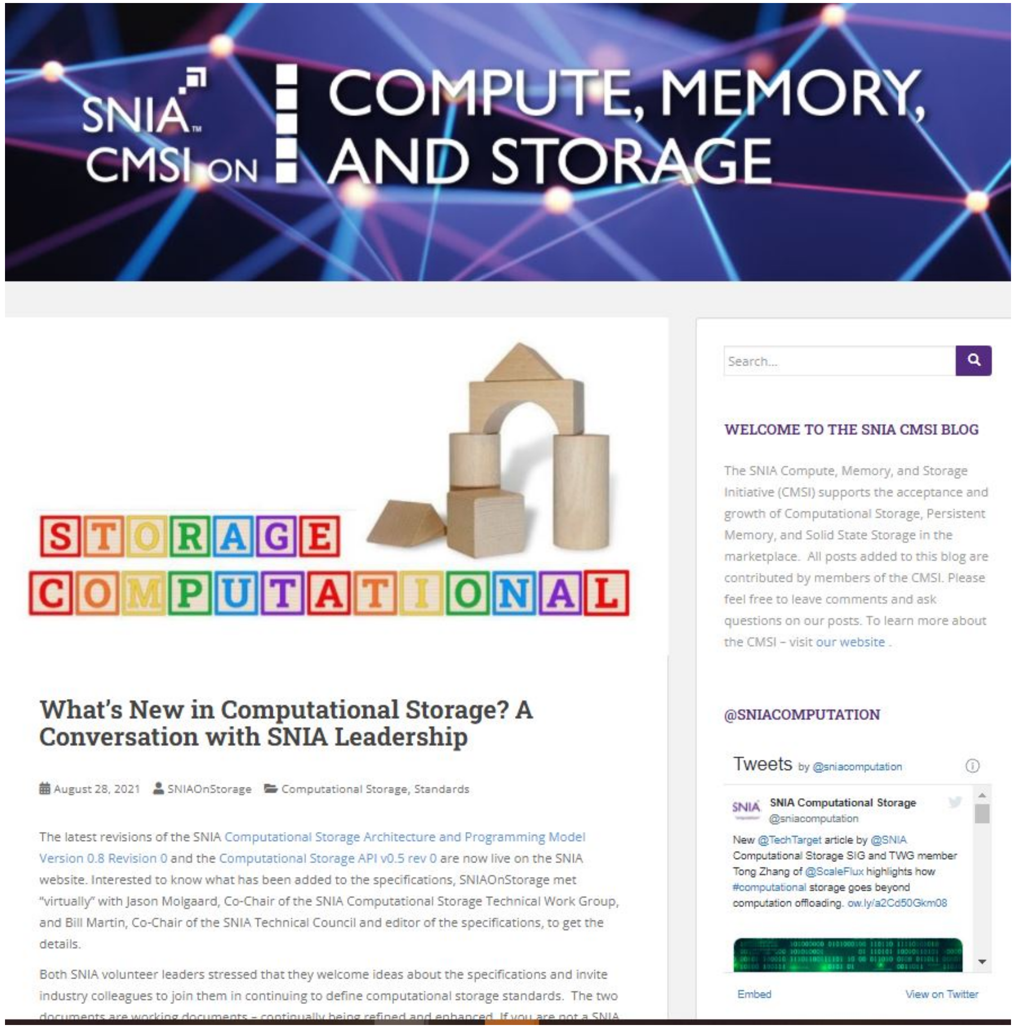

Check out our blog: sniacmsiblog.org

We welcome your questions

Thank you for watching
and
please rate the session

SNIA CMSI | COMPUTE, MEMORY, AND STORAGE