STORAGE DEVELOPER CONFERENCE
SDC 21
BY Developers FOR Developers

Virtual Conference
September 28-29, 2021

A SNIA Event

CS TWG Update

What A Year It Was!

Presented by the Co-Chairs of the CS TWG

Jason Molgaard, Solutions Architect, ARM

Scott Shadley, VP Marketing, NGD Systems

Agenda

- Updates on the TWG Membership
- Updates on the TWG Work Efforts
- Status of the Architecture
- Status of the SW API
- What is Next?

The Growth and Evolution

TWG Membership

The Continued Growth of Experience

- 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 session presented by the Co-Chairs of that work

The Efforts to Get Information Out is Continuing

- ComputerWeekly.com
  - 13-part Series

CW Developer Network

Computational storage: A Computer Weekly analysis series

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

Gartner COOL VENDOR

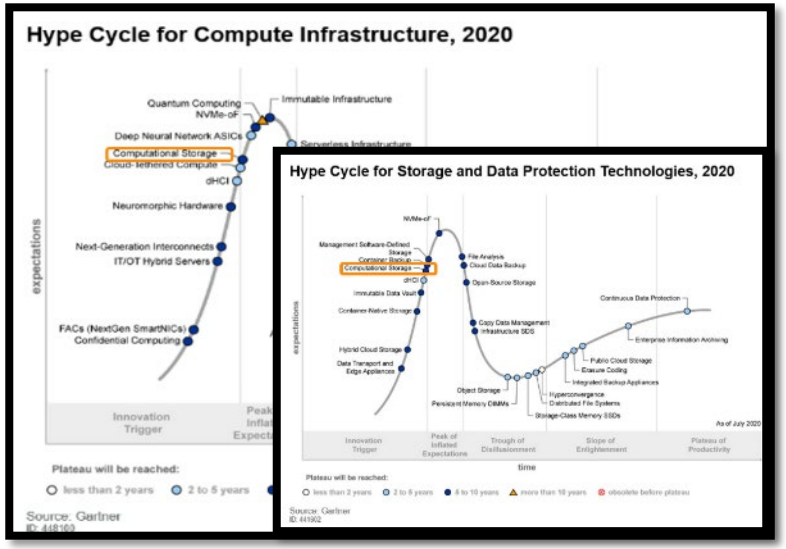

CW Developer Network

Computational storage series: Evaluator Group - Speculations, expectations & extrapolations

CW Developer Network

Computational storage: NGD Systems / SNIA - Icebergs at the Edge

Cliff Saran's Enterprise blog

An opportunity to redesign computer architectures

How computational storage delivers datacentre benefits

Computational storage is an emerging field of IT that features compute processing power closely coupled with storage. We look at what it can be used for

- Sponsored Efforts

EMC3

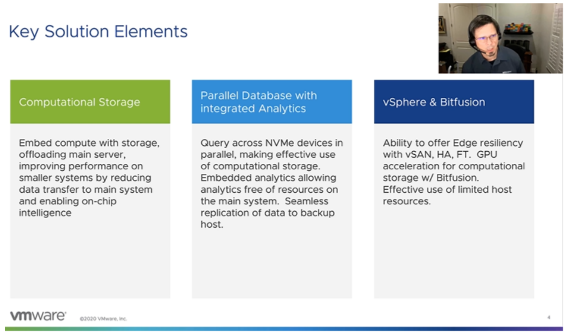

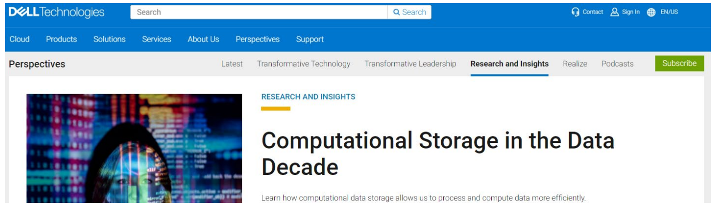

STORAGE DEVELOPER CONFERENCE SDC 21

The Output of the Work

Publicly Reviewable Documents

Current Progress of TWG Output

- Architectural Document has been Released
  - V0.8 is now in Public Review
  - Many updates from 0.5
- Initial release of API Document
  - First level support of customer interface
- Security now being reviewed
  - In Collaboration with Security TWG
- Sessions at this Event on API and Architecture!!
  - Bill Martin – Architecture work
  - Oscar Pinto – API work

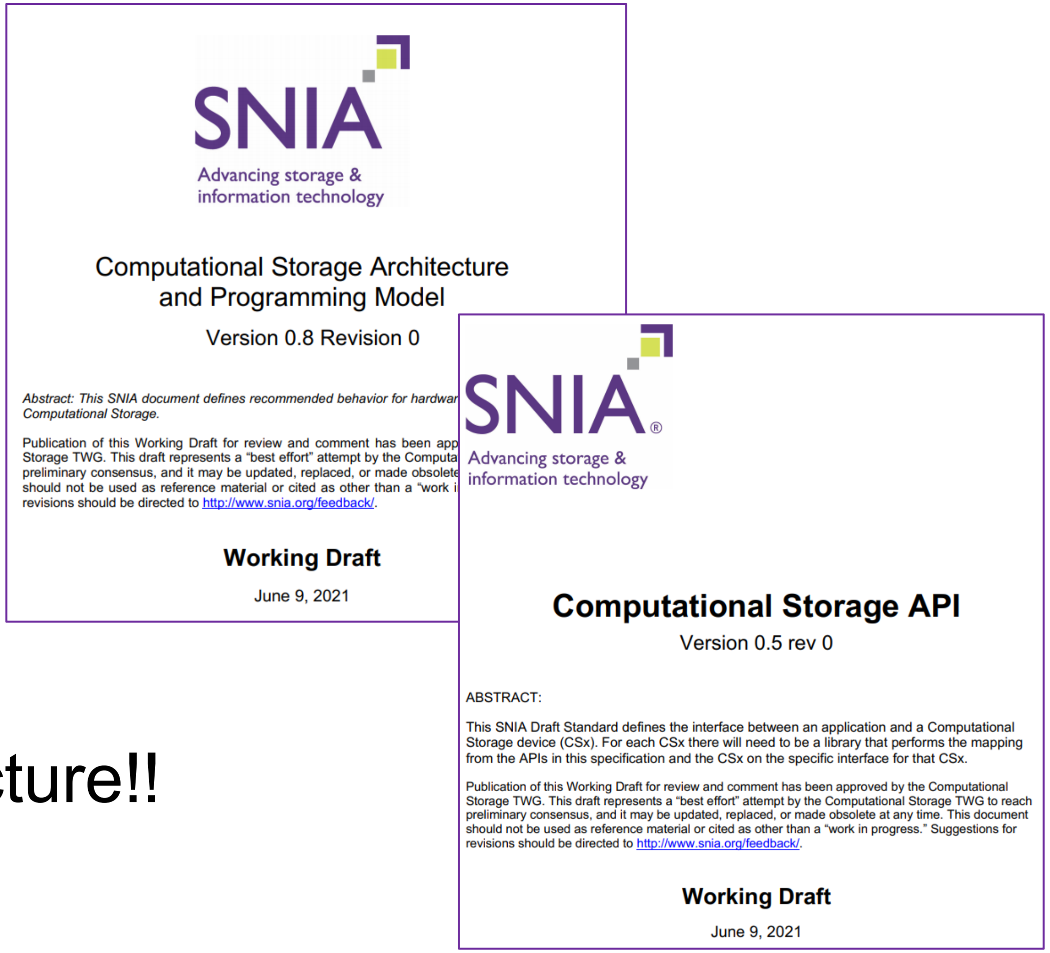
SNIA
Advancing storage & information technology

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 Use Cases – Annex – Review?

- Many Members supported actual use cases, Feedback!!

ANNEX B. (INFORMATIVE) ILLUSTRATIVE EXAMPLES 37

B.1 CSFs on a Large Multi-Device Dataset using Ceph ...

B.2 Using a Containerized Application within Linux (CSEE with included CSF)...

B.3 Data Deduplication CSF ...

B.4 A Computational Storage Function on a NVM Express and OpenCL based Computational Storage Drive Illustrative Example.. 60

B.5 Data Compression CSF Example ..

B.6 Data Filter CSF Example....

B.7 Scatter Gather CSF..

What is the CS Architecture?

Decoding CS, CSF, CSE, CSEE, CSR, CSP, CSD, CSA

What Has Changed? Architecture

- With More Members, there is More Understanding

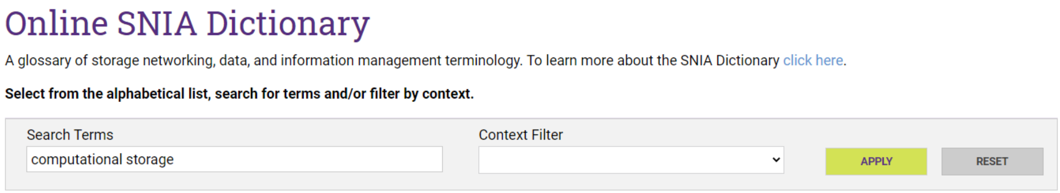

1) Re-naming ‘internal’ Computational Storage Processor –
   a. Component within a CSx is now called a Computational Storage Engine (CSE)

2) New architectural concept of a Computational Storage Engine Environment (CSEE)
   a. This would be attached to a specific CSE and defines the environment for CSF

3) New architectural element of a Resource Repository that contains:
   a. CSEEs that are available for activation on a CSE
   b. CSFs that are available for activation on a CSEE

4) Discovery and configuration flow documented

A Brief Rundown of the Dictionary Changes

REMOVAL OF - CSS, FCSS, PCSS to better align the Architectural use case of the Computational Storage Devices (CSx) - Addition of replacement terms below

Computational Storage

Computational Storage Resource (CSR)

Computational Storage Engine (CSE) ⟷ **Computational Storage Function (CSF)**

Computational Storage Devices (CSx)

- **Computational Storage Processor (CSP)**
- **Computational Storage Drive (CSD)**
- **Computational Storage Array (CSA)**

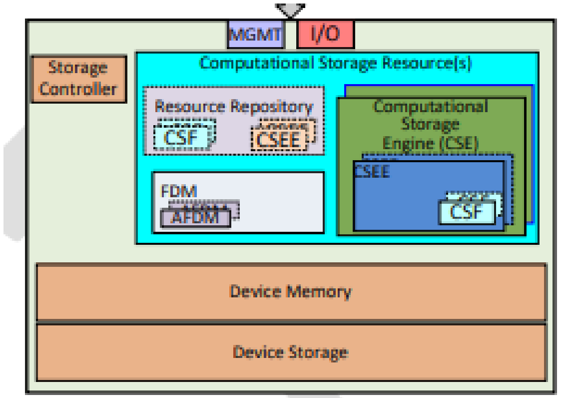

Computational Storage Drive (CSD)

Discovery Process Flow

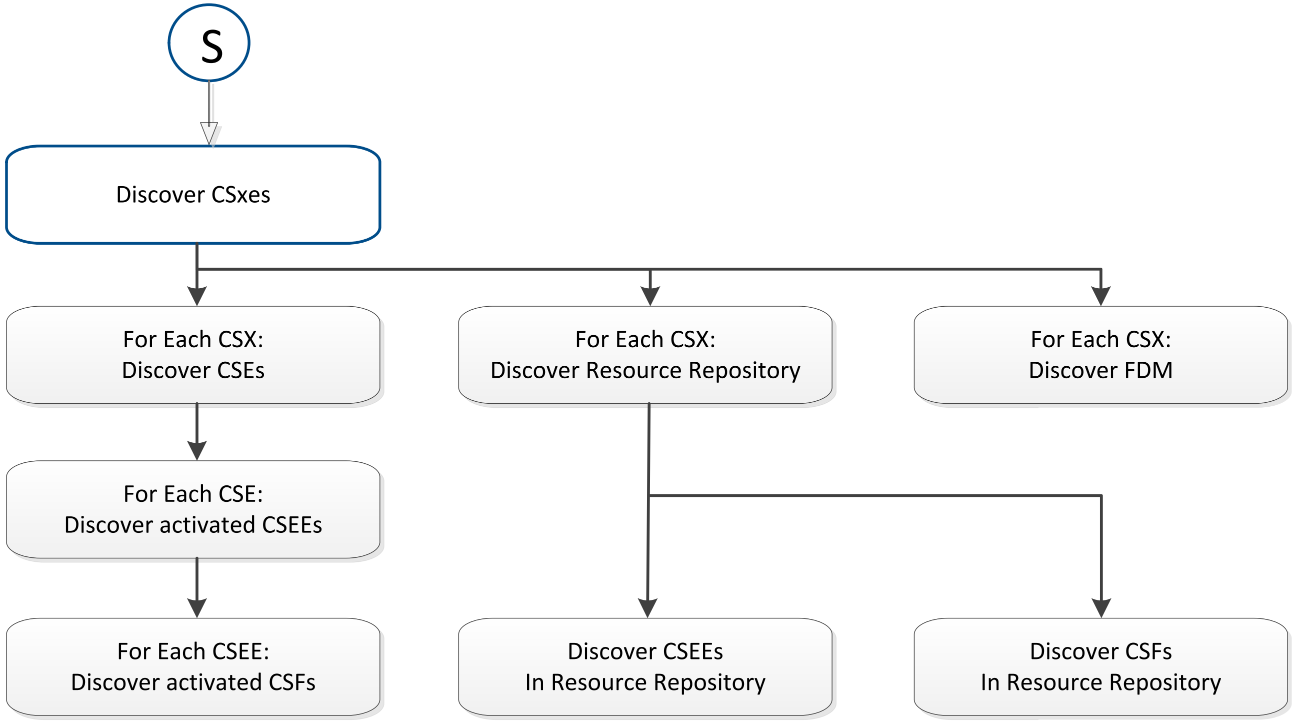

Session at this Event on this work

STORAGE DEVELOPER CONFERENCE SDC 21

What about the SW?

A brief look into the API

What Has Been Created? - API

1) Proposes an Application Programming Interface to Computational Storage devices

2) Allows a user application on a host to have a consistent interface to any vendor’s CS device

3) Vendor defines a library for their device that implements the API
    a. Mapping to wire protocol for the device is done by this library
    b. Functions that are not available on a specific CS device may be implemented in software

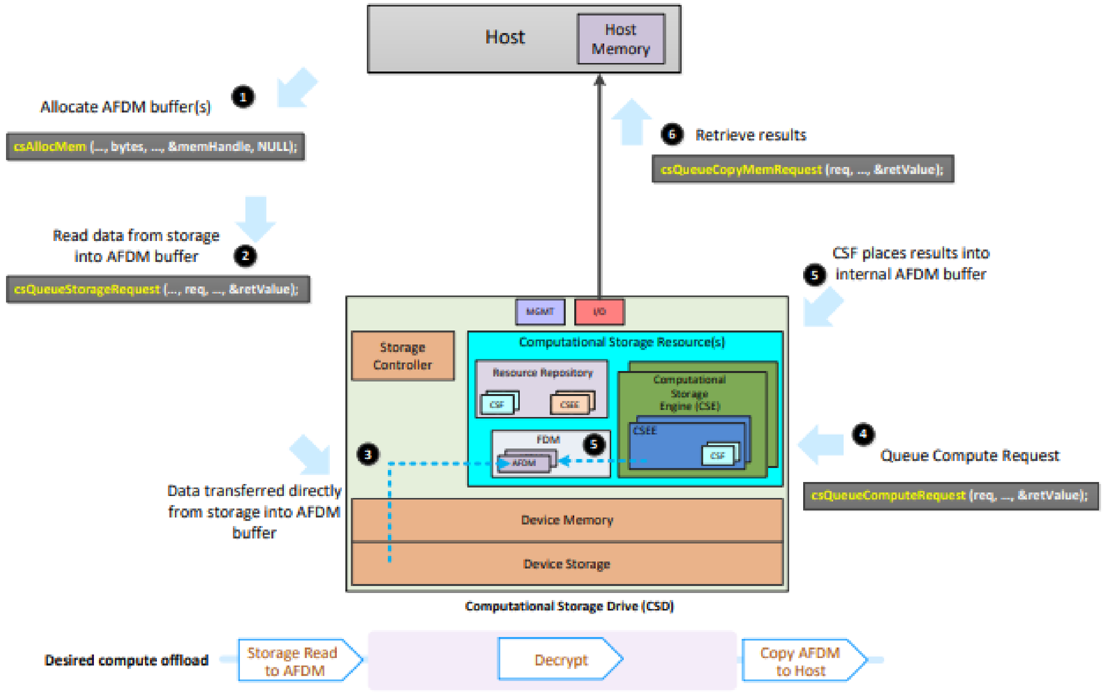

Figure 3: Example API flows

STORAGE DEVELOPER CONFERENCE SDC 21

What is Being Accomplished?

4 APIs Overview

Computational storage is possible with CSEs that are able to execute compute tasks typically run on a host CPU. These CSEs may use FDM that is different from the host memory and memory for storing CSFs. This memory is where computation programs run when they do. A mechanism is needed to transfer data to and from AFDM. These data transfers are required for inputs and outputs to the CSE compute functions. Data transfers to AFDM may be from host memory and/or storage. There are specific APIs that target these operations and interactions with the CSE. This section targets the usage of APIs and how they are able to be used with CSEs for computational storage.

This standard defines a base set of functions. Additional libraries are able to be built on this base set of functions. This version of the standard is tailored towards a host orchestrated interface. There are additional APIs required for a fully device managed interface. A future revision of this standard will also cover how the device may manage the FDM without host control.

As Computational Storage APIs provide mechanisms to allocate AFDM, there is a requirement that the case of computation overrunning the AFDM needs to be documented. This will be documented in a future revision of this standard.

If the device that this API interfaces to does not implement a particular function the API may return an error or implement an emulation of the function. The default is to return an error. The mechanism to control which of these options is performed is to be defined in a future revision of this standard.

COMPUTATIONAL STORAGE

Computational Storage APIs

Oscar Pinto, Principal Engineer, Samsung Semiconductor Inc

What Next?

STORAGE DEVELOPER CONFERENCE SDC 21

Moving Beyond Architecture

- Security and Computational Storage
  - Security TWG collaboration
- Illustrative Examples Growth
  - More and more ways to deploy
- Deployment Models and Feedback
  - Customer use cases, market growth

STORAGE DEVELOPER CONFERENCE SDC 21

SNIA

SEARCH LOGIN JOIN SNIA

ABOUT | STANDARDS | EDUCATION | TECHNOLOGY FOCUS AREAS | NEWS & EVENTS | RESOURCES | MEMBERSHIP

- Cloud Storage Technologies
- Computational Storage
  - Computational Storage Use Cases
  - Compute, Memory, and Storage Initiative
  - Computational Storage Technical Work Group

Home » Technology Focus Areas » Computational Storage

Computational Storage

Today, Computational Storage is transforming enterprises worldwide. The SNIA Computational Storage Technical Work Group (TWG) is 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. In addition, the SNIA Compute, Memory, and Storage Initiative (CMSI) is focused on fostering the acceptance and growth of computational storage in the marketplace with the activities of the Computational Storage Special Interest Group. To achieve those goals, the CMSI provides education, performs market outreach, and influences and promotes standards.

NVMe Computational Storage Task Group

The charter of Computational Storage Task Group is to develop features associated with the concept of Computational Storage on NVM Express devices. The scope of work encompasses how these features are discovered, configured and used inside an NVM Express framework. Examples of these features include general compute, compression, encryption, data filtering, image manipulation and database acceleration.

The target audience consists of the vendors and customers of NVMe Storage Devices that support computational features.

nvm EXPRESS®

Join NVM Express

Session at this Event on NVMe Work

Kim Malone, Stephen Bates – Co-Chairs

SDC 2021 Recap of Sessions

COMPUTATIONAL STORAGE

Computational Storage Architecture Simplification and Evolution

Jason Molgaard, Sr. Principal Storage Solutions Architect, Arm Inc.

COMPUTATIONAL STORAGE

Computational Storage APIs

Oscar Pinto, Principal Engineer, Samsung Semiconductor Inc

COMPUTATIONAL STORAGE

Scientific Data Powered by User-Defined Functions

Lucas Villa Real, Research Software Engineer, IBM Research

COMPUTATIONAL STORAGE

SkyhookDM: An Arrow-Native Storage System

Jayjeet Chakraborty, IRIS-HEP Fellow IRIS-HEP and CROSS, UC Santa Cruz

COMPUTATIONAL STORAGE

Computational Storage Deployment with Kubernetes and Containerized Applications

Scott Shadley, Co-Chair, Computational Storage TWG - VP Marketing, NGD Systems

COMPUTATIONAL STORAGE

Accelerating File Systems and Data Services with Computational Storage

Brad Settlemyer, Storage Researcher, Los Alamos National Laboratory

COMPUTATIONAL STORAGE

Stop Wasting 80% of Your Infrastructure Investment!

Tony Afshary, Sr. Director, Product Line Management, Pliops

COMPUTATIONAL STORAGE

Computational Storage

Jai Menon, Chief Scientist, Fungible

COMPUTATIONAL STORAGE

The Building Blocks to Design a Computational Storage Device

Jerome Gaysse, Senior Technology and Market Analyst, Silinnov Consulting

Please take a moment to rate this session.

Your feedback is important to us.

