



SNIA™ | COMPUTE, MEMORY,
CMSI | AND STORAGE



SNIA™ | NETWORKING
NSF | STORAGE

Storage Life on the Edge: Edge Use Cases

A CMSI-NSF Webcast

March 22, 2022



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.

Today's Presenters



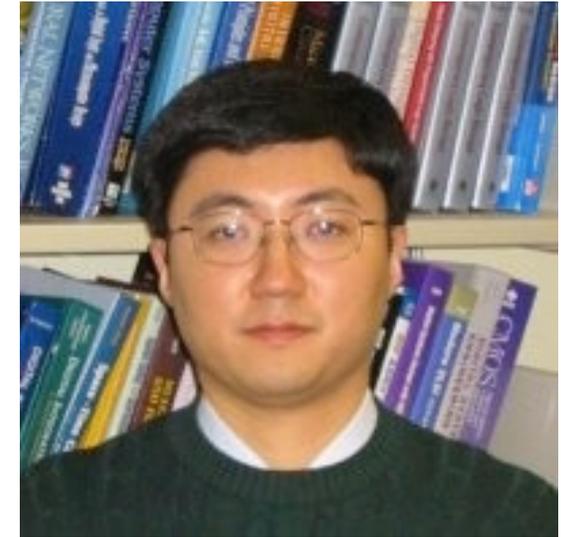
Moderator:
Bill Martin
Chair, SNIA Compute,
Memory, and Storage
Initiative



Mayank Saxena
Sr. Director Engineering
Samsung



Stephen Bates
CTO
Eideticom



Tong Zhang
Chief Scientist
ScaleFlux

SNIA-at-a-Glance



180
industry leading
organizations



2,500
active contributing
members



50,000
IT end users & storage
pros worldwide

Learn more: snia.org/technical

 **@SNIA**

**Ethernet, Fibre Channel,
Infiniband**

iSCSI, NVMe-oF, NFS, SMB

**Virtualized, HCI,
Software-Defined Storage**

**Storage Protocols
(block, file, object)**

Securing Data

Technologies We Cover

**Computational Storage
Standards**

**Persistent Memory
Programming**

**Smart Data Accelerator
Interface**

**Solid State Drives,
Solid State Performance**

**SSD Form Factors,
NVMe SSD Classification,
TCO Model for Storage**

Agenda

- Storage Life on the Edge – A Recap
- Edge Use Cases
- Panel Discussion
- Next Up.....



Storage Life on the Edge: Managing Data from the Edge to the Cloud and Back

- On January 26, 2022 we covered:
 - Data and compute pressure points: aggregation, near & far edge
 - Supporting edge workloads
 - Analytics and AI considerations
 - Understanding data lifecycle to generate insights
 - Governance, security & privacy overview
 - Optimizing for CapEx & OpEx
- Available on-demand at <https://bit.ly/StorageLifeEdge>

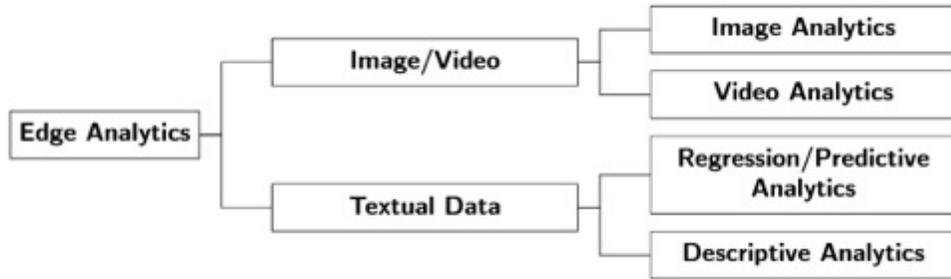




Mayank Saxena

Samsung

Edge Data Needs



Type of Edge Analytics

Parameters	Cloud Analytics	Edge Analytics
Deployed	Datacenter	Edge Device (e.g., Routers, Sensors)
Age	Older	Younger
Development	Matured	Fragmented
Size of data	Big data	Small
Application	Advanced and sophisticated	Not advanced and unsophisticated
Storage	Distributed System	Small memory of Edge device
QoS dependence	None	Cloud
Performance	Varied	Near real time

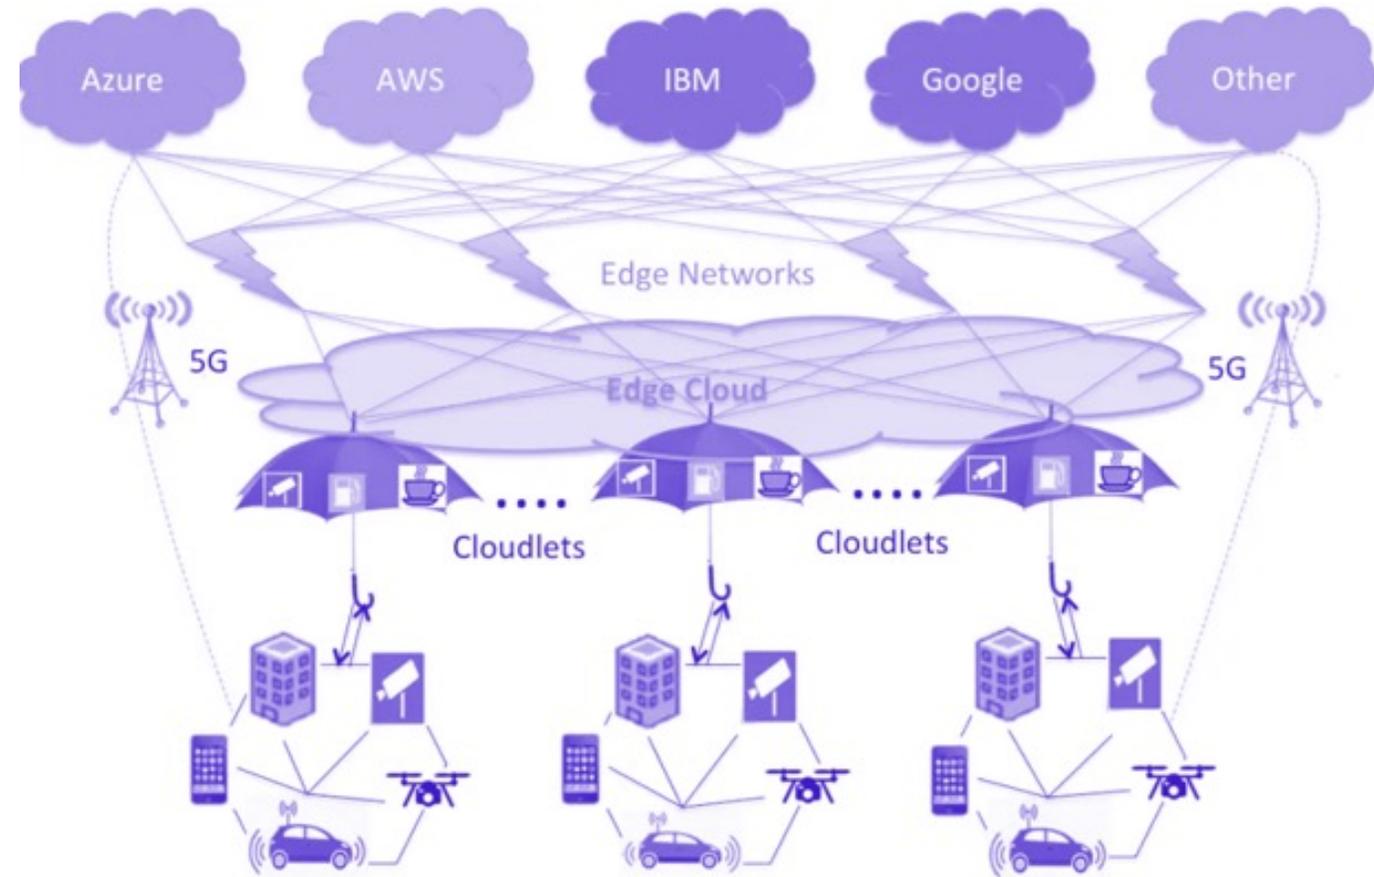
Cloud to Edge: Need for Speed

■ Use Cases

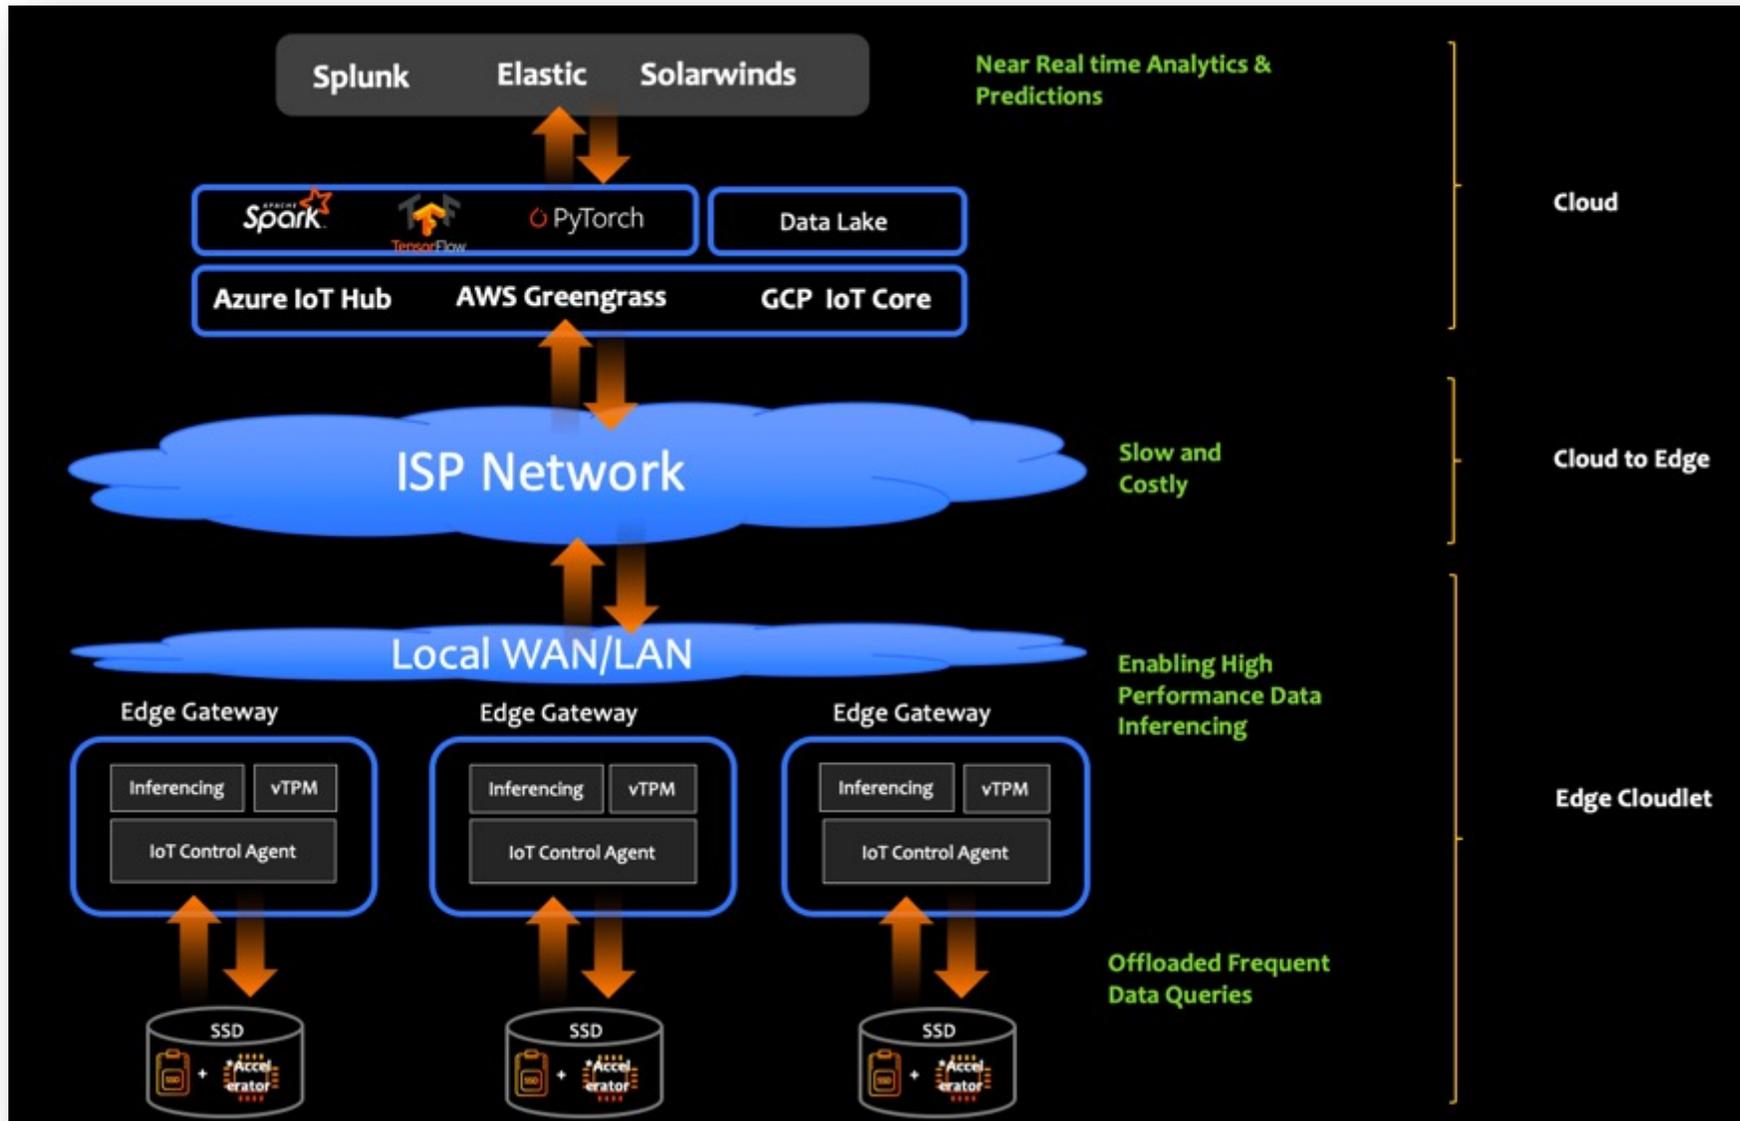
- Distributed Analytics at Edge Cloud
- Federated AI
- Near Real Time Anomaly detections

■ Challenges with Edge Cloud

- Cloudlets – cluster of edge gateways with storage and complete constraints
- Slow mesh network for Cloudlets



High Performance Edge Data





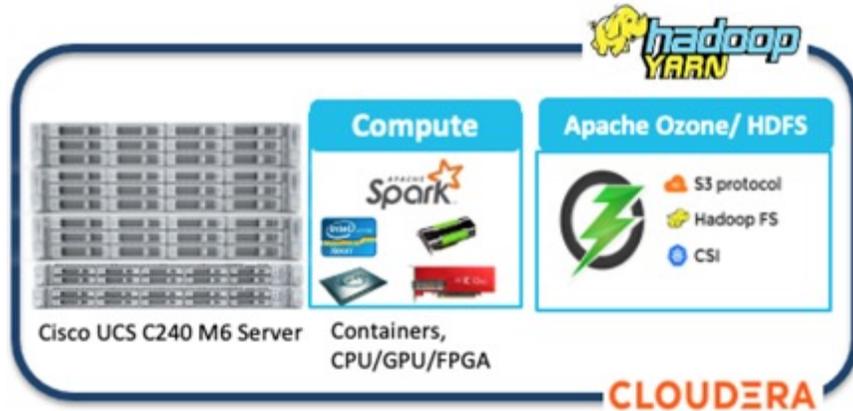
Stephen Bates

Eideticom

NoLoad Transparent Compression - Spark Acceleration

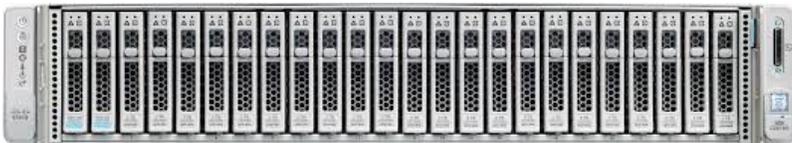


Cisco Data Intelligence Platform



NoLoad Benefits:

- Zero application changes
- **3x** increased capacity
- **Deployment flexibility** in compute, storage and cloud

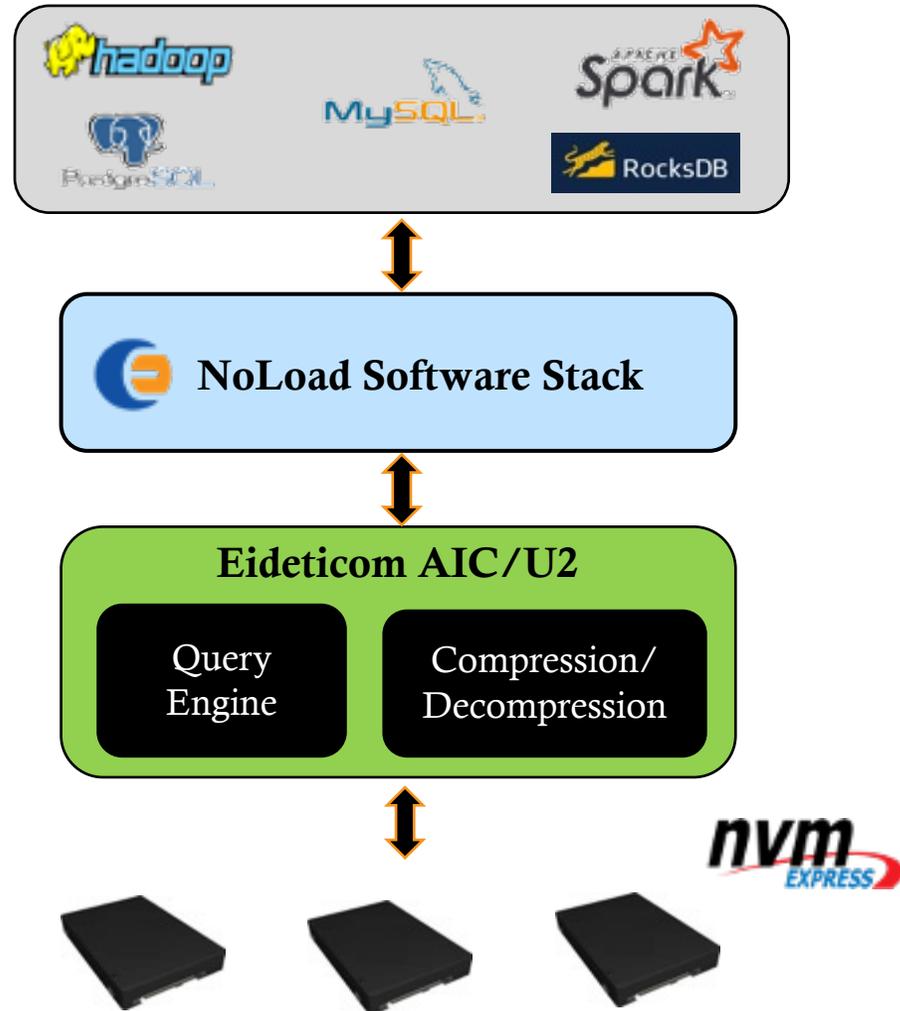


1x Capacity



3x Capacity

NoLoad Query Analytics – Database Acceleration



Highlights:

- NoLoad Query Engine **accelerates** data queries
- **Implements** CSV/JSON/Parquet parsing and query execution
- **70-80%** Improved CPU Offload with **5-10x** Increased Performance
- NVMe driver achieves **low latency high throughput** data transfers
- **Deployment flexibility** in compute, storage and cloud
- **NoLoad Query Engine** can be paired with NoLoad compression and decompression engines



Tong Zhang

ScaleFlux

Far Edge: Retail Chain Example

Essential applications with tight latency requirements

Running in 100s or 1000s of locations

With minimal or no on-site IT Staff

Challenges



- ❑ No true on-site data center
 - ❑ Limited environmental controls & floorspace
 - ❑ Low-power Servers to help with thermals
 - ❑ Low cost per Server to scale to 1000s of locations
- ❑ Software suite for centralized management

Solutions



- Computational Storage Drives
- Low-overhead Edge Computing SW



Results



- ✓ CSDs expand Server functionality
 - ✓ **Low-power & low-cost capacity expansion** via CSD hardware compute engines (HCE)
 - ✓ **Improved thermal management** from distributed computing
 - ✓ **No added physical footprint** since the HCE is integrated directly into the SSD
- ✓ Optimized Edge SW addresses system constraints & management
 - ✓ Brings **HCI** capabilities to the Edge
 - ✓ **Remote management** of 1000s of sites
 - ✓ **Reduced CPU power** to meet SLAs
 - ✓ **Reduced Server costs** via lower-grade CPUs, less DRAM, and smaller form factors

Far Edge / Near Edge: Data Security Example

Companies must protect against mis-use of data

Generated in 10s, 100s, or 1000s of Edge locations

And transferred to the Core or Cloud for processing

Challenges



- ❑ Each movement of data creates an exposure opportunity
 - ❑ Man-in-the-middle intercept
 - ❑ Unintended inferencing / model training

Solutions



- Computational Storage Drives
- Data Transformation Application



Results



- ✓ Original data never leaves the drive, **only a representation** of the data is transferred
- ✓ CSDs running Data Transformation enhance data security
 - ✓ **Original data protected** from man-in-the-middle malicious actors
 - ✓ **Transformations minimize the potential for mis-use** of data for unintended inference/ model training while enabling the intended task to be performed



Panel Discussion and Q&A

Storage Life on the Edge is a Series! Next Up:

“Storage Life on the Edge: Security Challenges”

April 27, 2022

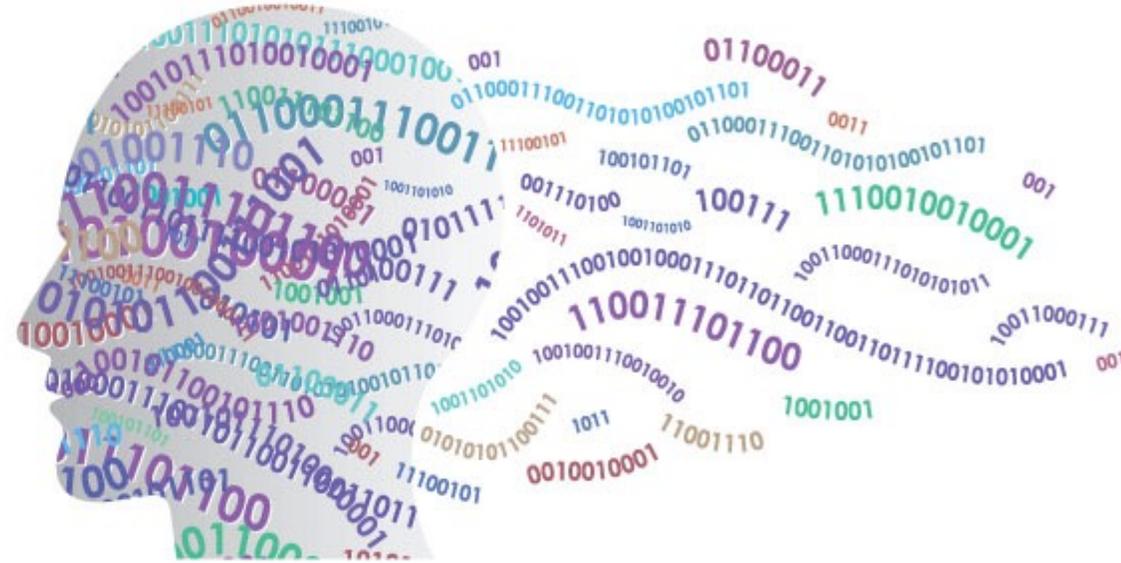
10:00 am PT / 1:00 pm ET

Register here: <https://bit.ly/SNIAEdgeLifeSecurity>



SNIA PERSISTENT MEMORY
+ **SUMMIT 2022**
COMPUTATIONAL STORAGE

VIRTUAL EVENT • MAY 25-26, 2022



- The 10th annual PM+CS Summit is a virtual event this year with live keynotes and panels, breakout sessions, and demos
- Complimentary registration is now open – save your seat today!

<https://www.snia.org/events/persistent-memory-summit/pm-cs-summit-2022-registration>

Thanks for Watching Our Webcast

- Please rate this webcast and provide us with feedback
- You can 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 On Storage Blog at <https://sniablog.org/>
- Follow us on Twitter: [@SNIA](#)
- Learn more about compute, memory, and storage at www.snia.org/CMSI
- Learn more about networked storage at www.snia.org/NSF



The screenshot shows the SNIA On Storage website. At the top, there is a dark blue header with the text "SNIA ON STORAGE" in white. Below this is a light gray section with a large image of a person sitting on a rock, with the text "STORAGE LIFE ON THE EDGE" in blue. Below the image is a blog post titled "Our Storage Life on the Edge Webcast Series Continues...." dated March 21, 2022. The post text begins: "The second webcast in our Storage Life on the Edge series is coming up on March 22, 2022 at 10:00 am Pacific time. This panel, moderated by Bill Martin, SNIA Compute, Memory, and Storage Initiative Chair, takes a deeper dive to focus on edge use cases in the computational storage space. Our panelists Mayank Saxena from Samsung, Stephen Bates from Eideticon, and Tong Zhang from ScaleFlux will discuss edge to cloud use cases where storage and compute resources need to be deployed in practical topologies that deliver the very best in application performance. They'll examine high performance edge data needs, database acceleration solutions, meeting retail chain..."

On the right side of the page, there is a search bar and a "RECENT POSTS" section listing several articles. At the bottom right, there is a "CONNECT WITH US" section with icons for RSS, Facebook, Twitter, YouTube, and LinkedIn.

Thank You for Watching!