Flash Memory Summit

SOLIDIGM

Prevent Silent Data Corruption with PI from NVMe *An Experts Take*

Niels Reimers

Solidigm Senior Strategic Planner

Silent Data Corruption (SDC)

Flash Memory Summit

What is Silent Data Corruption?

- Data is read with ***no error***, but data is ***not what was last written***
  - Frequency is low and impact is high

The proportion of outages costing over $100,000 has soared in recent years
Uptime Institute
June 8, 2022

A bank mistakenly put $120,000 into a couple's account. They spent it, police say
By Theresa Waldrop, CNN
September 9, 2019

SDCs are expensive, let's take a brief run through:

- *SDC sources*
- *SDC detection needs*
- *Protection Information as an SDC mitigation*

Silent Data Corruption Sources

What are the causes of Silent Data Corruption?

- Data Misplacement Errors
  - Data is written to or read from the wrong location
- Data Content Errors
  - Data content changes, either at rest or in flight
- Lost I/O Errors
  - A write is signaled as complete, but data has not been written

Where can Silent Data Corruption occur?

- Software and hardware, anywhere data is operated on, moved or stored
  - Storage stack, DMA engines, memory, switches, NICs, DPUs, HBAs, HDDs, SSDs

When can Silent Data Corruption occur?

- Any time in the life of the data
  - Normal operating conditions, power fails, data rebuilds/recoveries, snapshots, tiering, backups, protocol conversions, compression, deduplication, encryption

Silent Data Corruption Detection

Flash Memory Summit

Solution

- Ensure correctness of data through the life of the data
- “End to end” across data center(s)
  - Belt and suspenders, additive to all other data checks

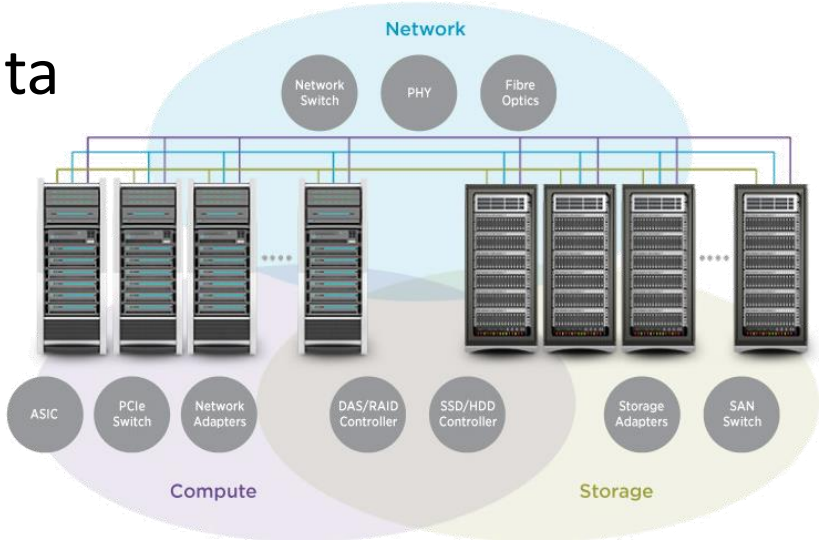

Implementation

- Is the data correct?
  - CRC or Checksum
- Is it the correct data?
  - In space : Carry LBA with data to compare with address in command
  - In time : Carry write counter with data to catch misplaced data and lost writes
  - In application : Carry tags to denote application connection

Silent Data Corruption Detection

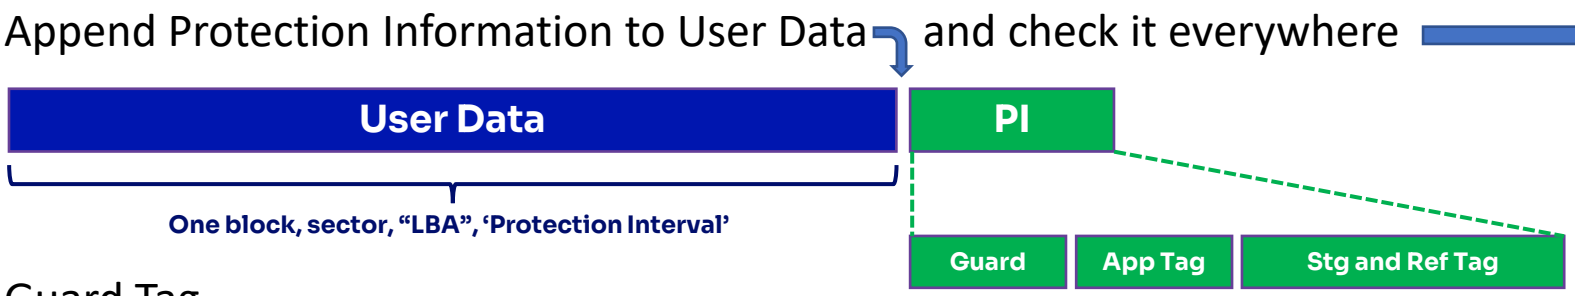

Guard Tag

- CRC (or Checksum in some applications)

Storage and Reference Tag

- High order bits used by storage devices
- Low order bits contain LBA

Application Tag

- Bit vector used by host applications

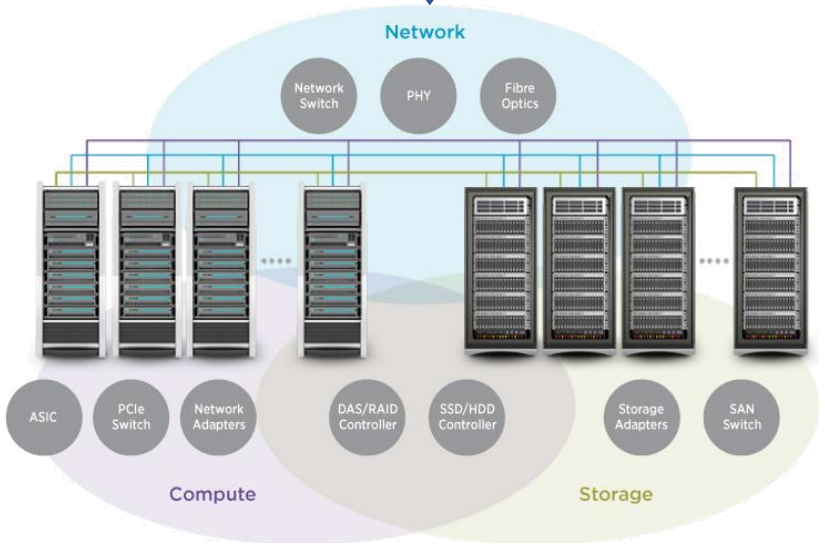

Operations : Insert, forward, remove

Control bits and masks to select checks and operations

Wide industry support including Intel DSA, Linux OSes, NICs, ROCs, HBAs, DPUs, HDDs, SSDs and storage arrays

Standardized in T10 and NVMe


SOLIDIGM™