SNIA STORAGE SECURITY SUMMIT
Wednesday, May 11, 2022 • Virtual

Persistent Data for Secured Containers

A Realizable Vision?

Nick Connolly, Chief Scientist, DataCore Software

A SNIA Event

SNIA Legal Notice

- The material contained in this tutorial 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 containin 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.

Agenda

- Container Environment
- Securing Data
  … at Rest
  … in Use
  … in Motion
- Long Term Research
- Conclusion

Container Environment

Application Environment

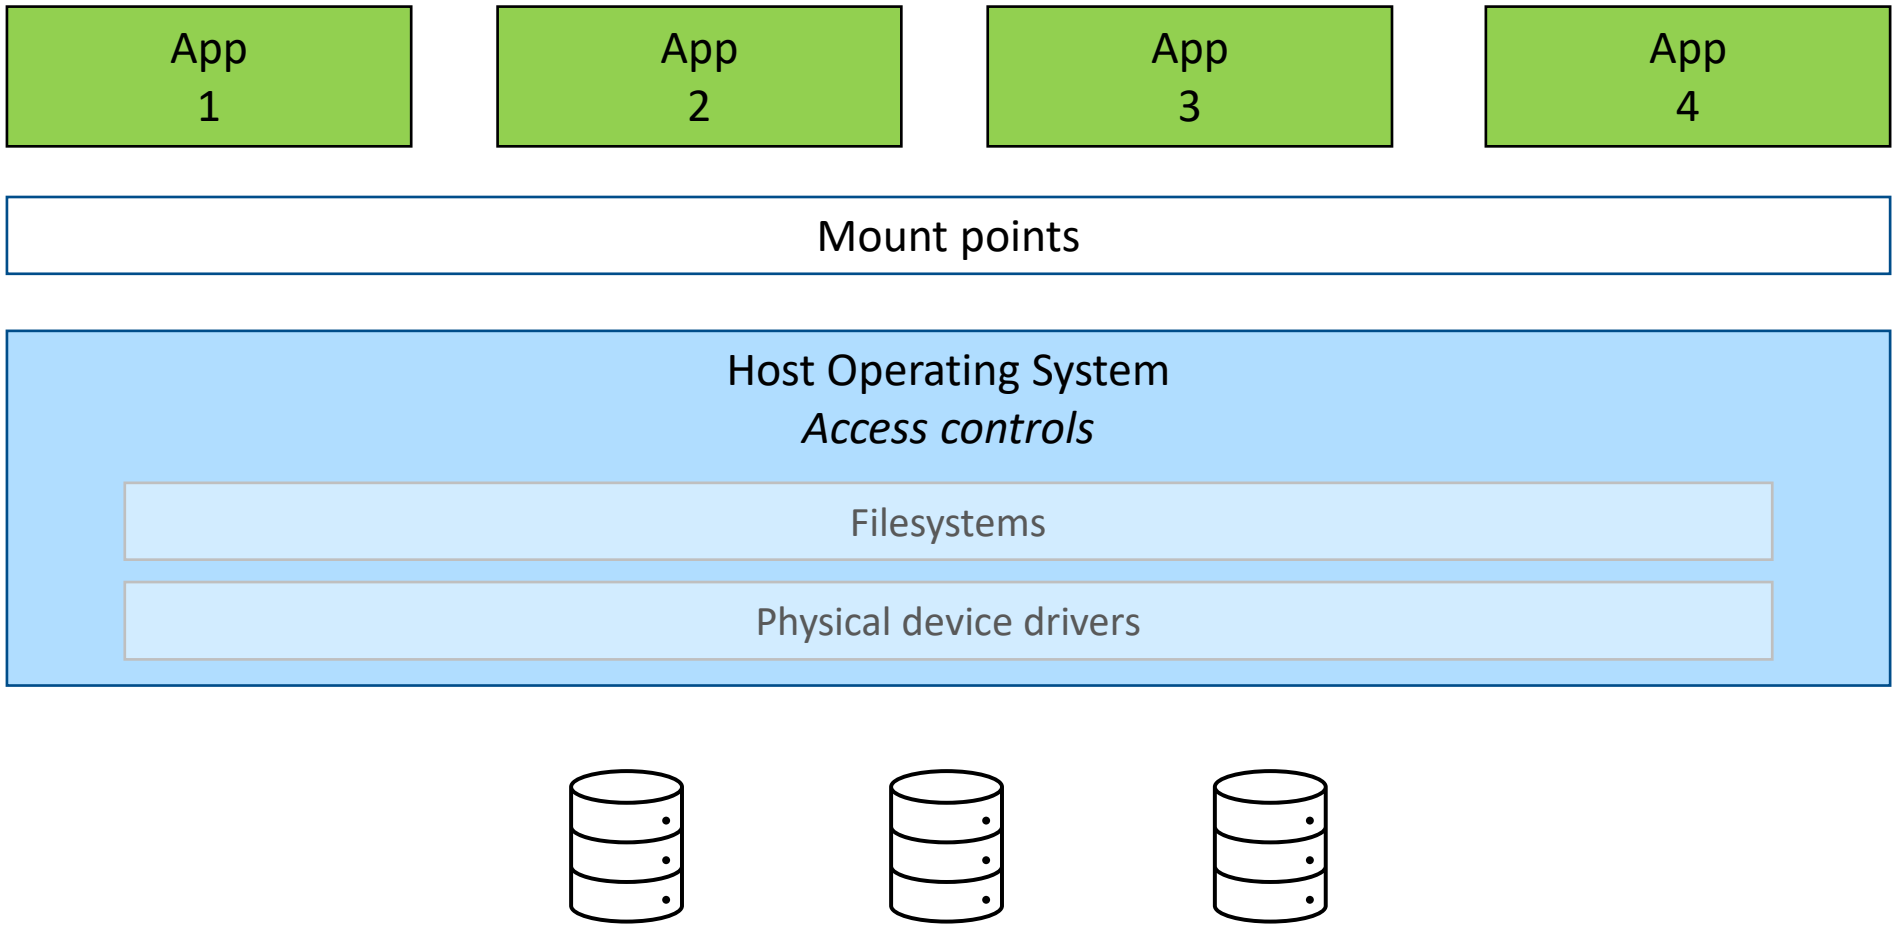

Virtualized Application Environment

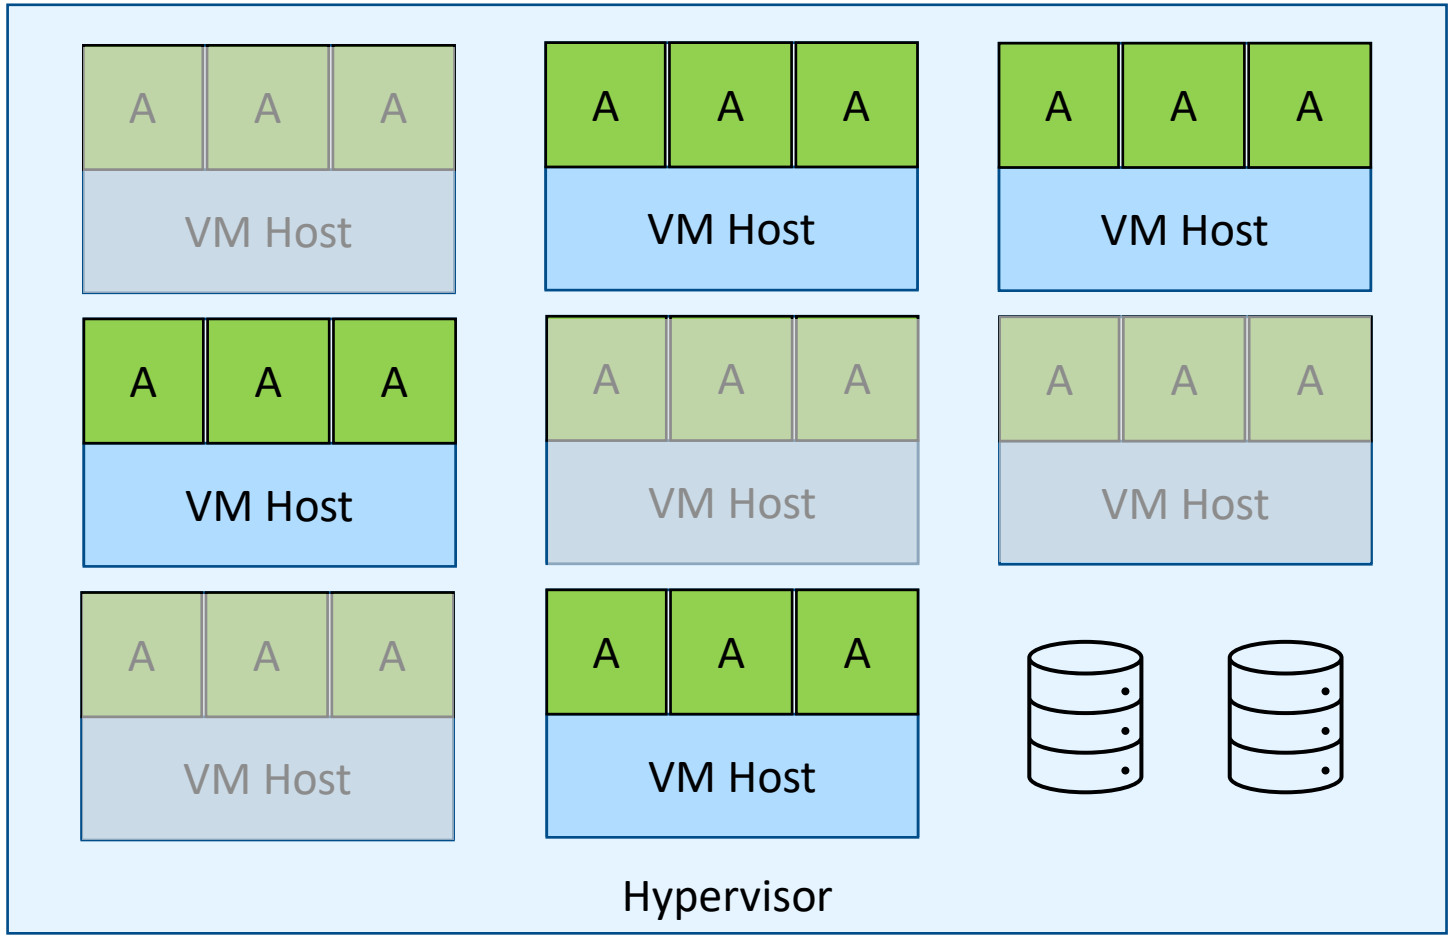

Containerized Environment

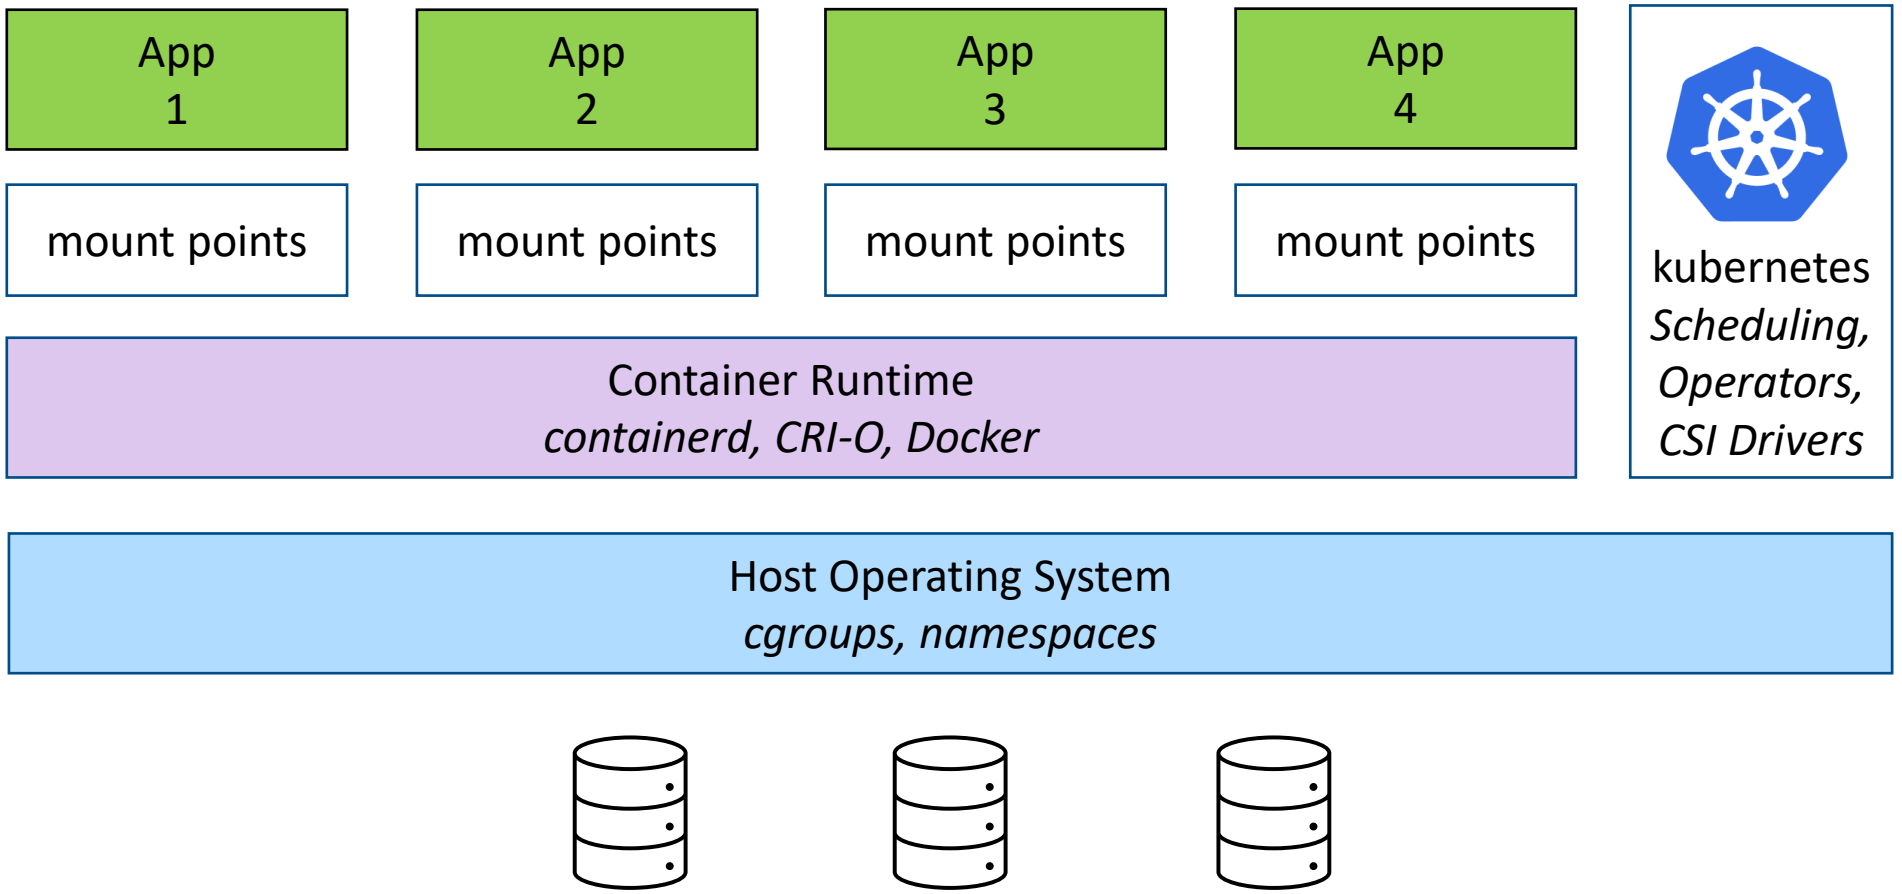

Containerized Storage Interface

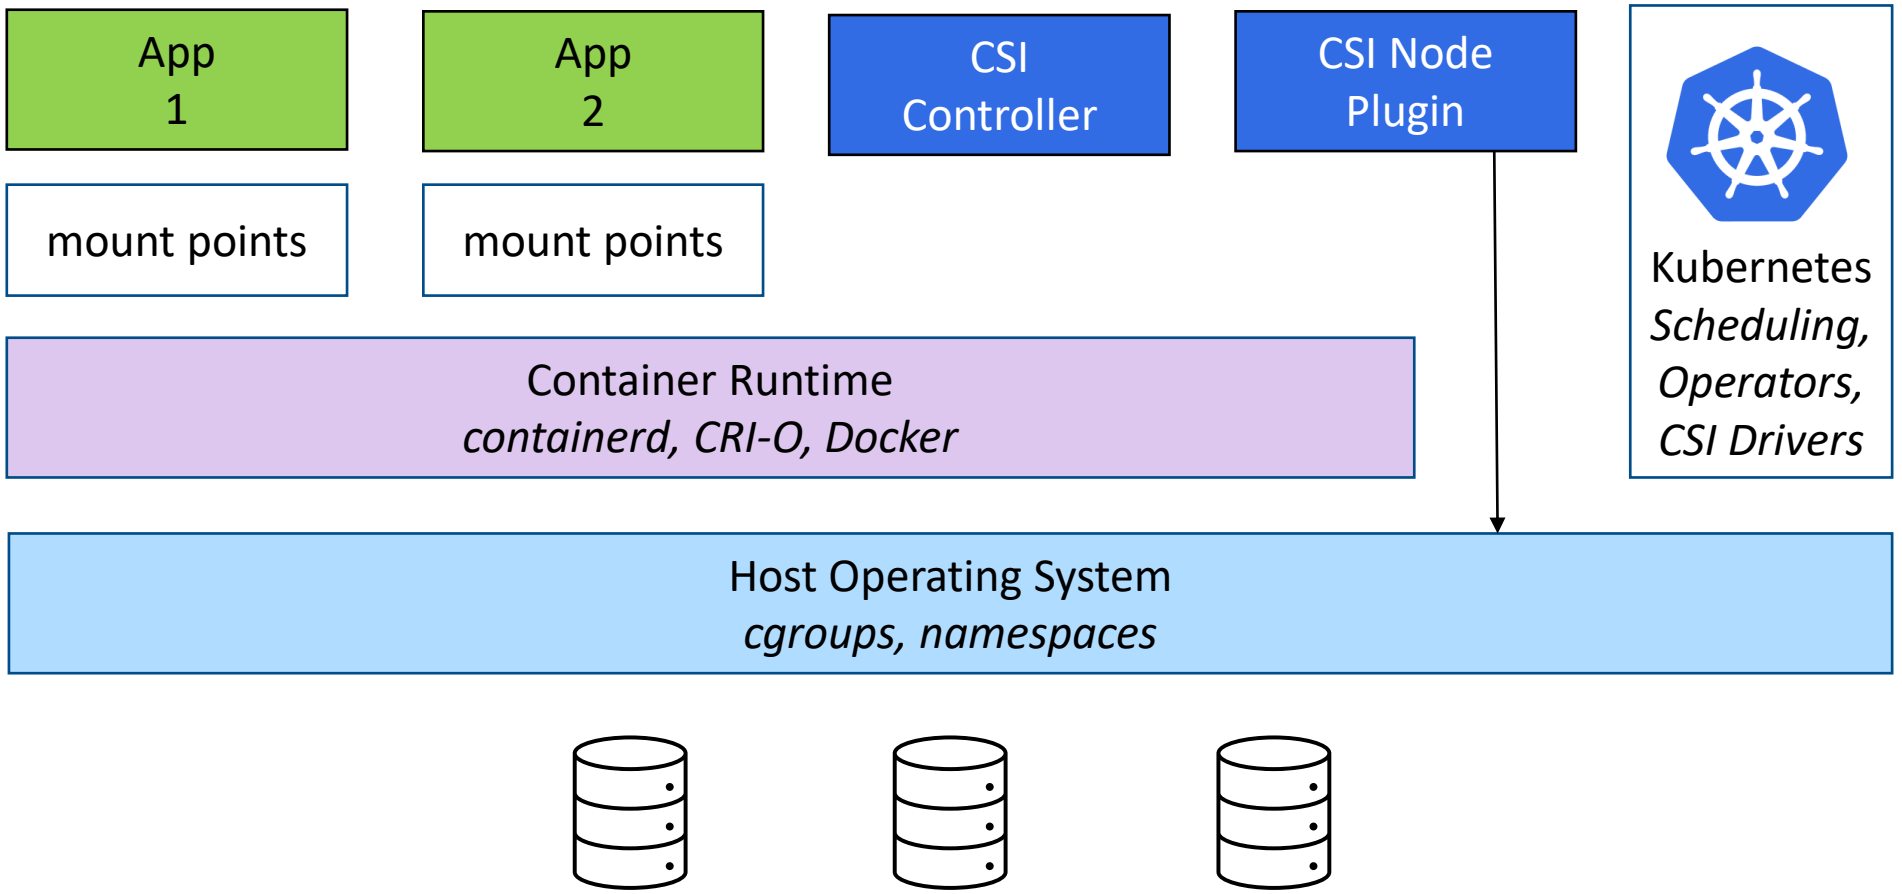

Container Native Storage

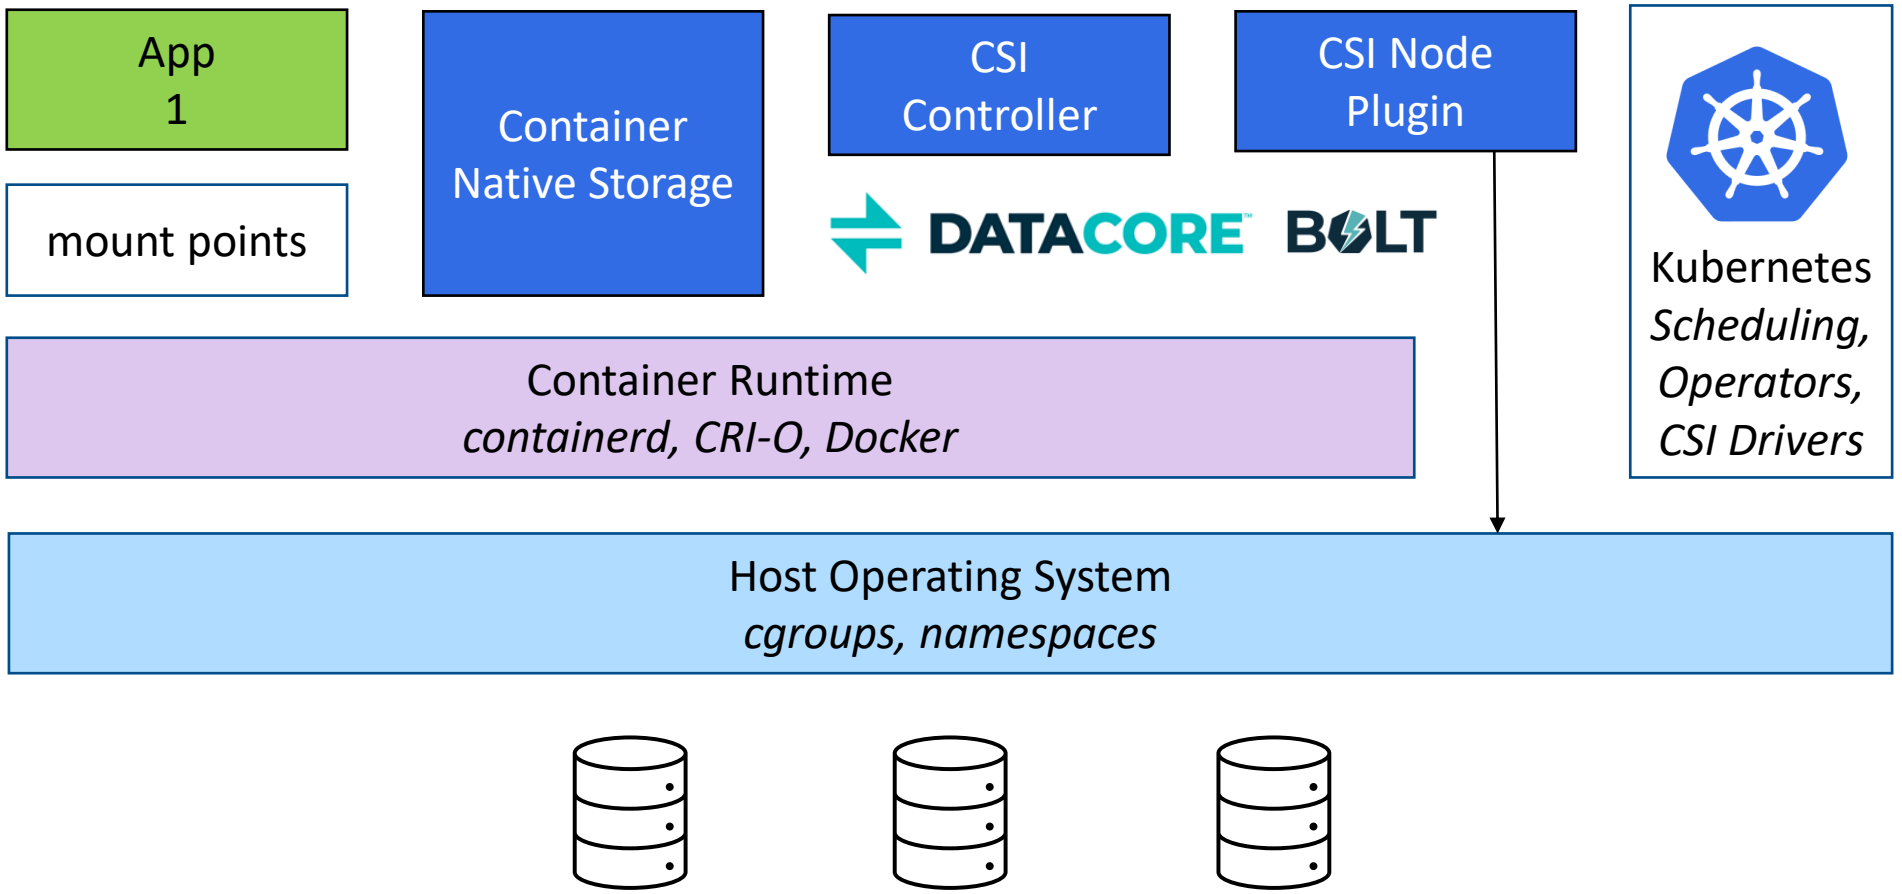

Managed Kubernetes

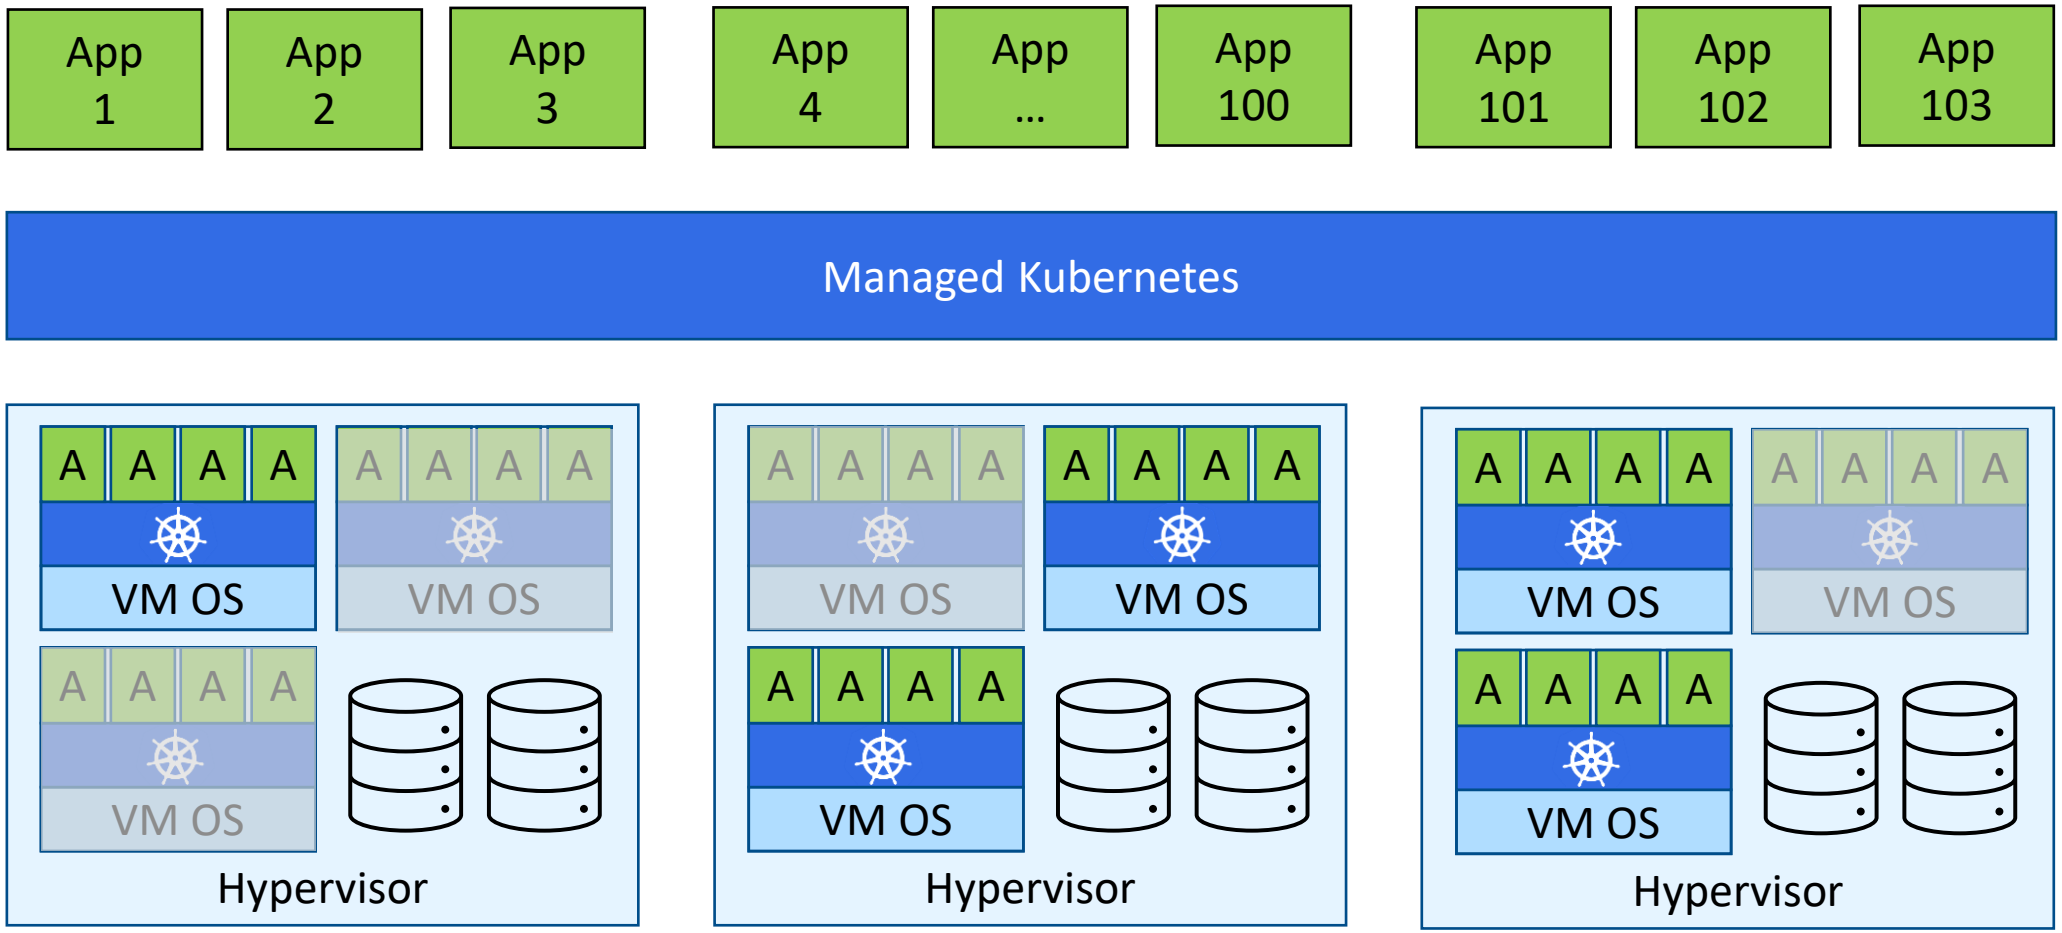

Securing Data at Rest

File Based Encryption

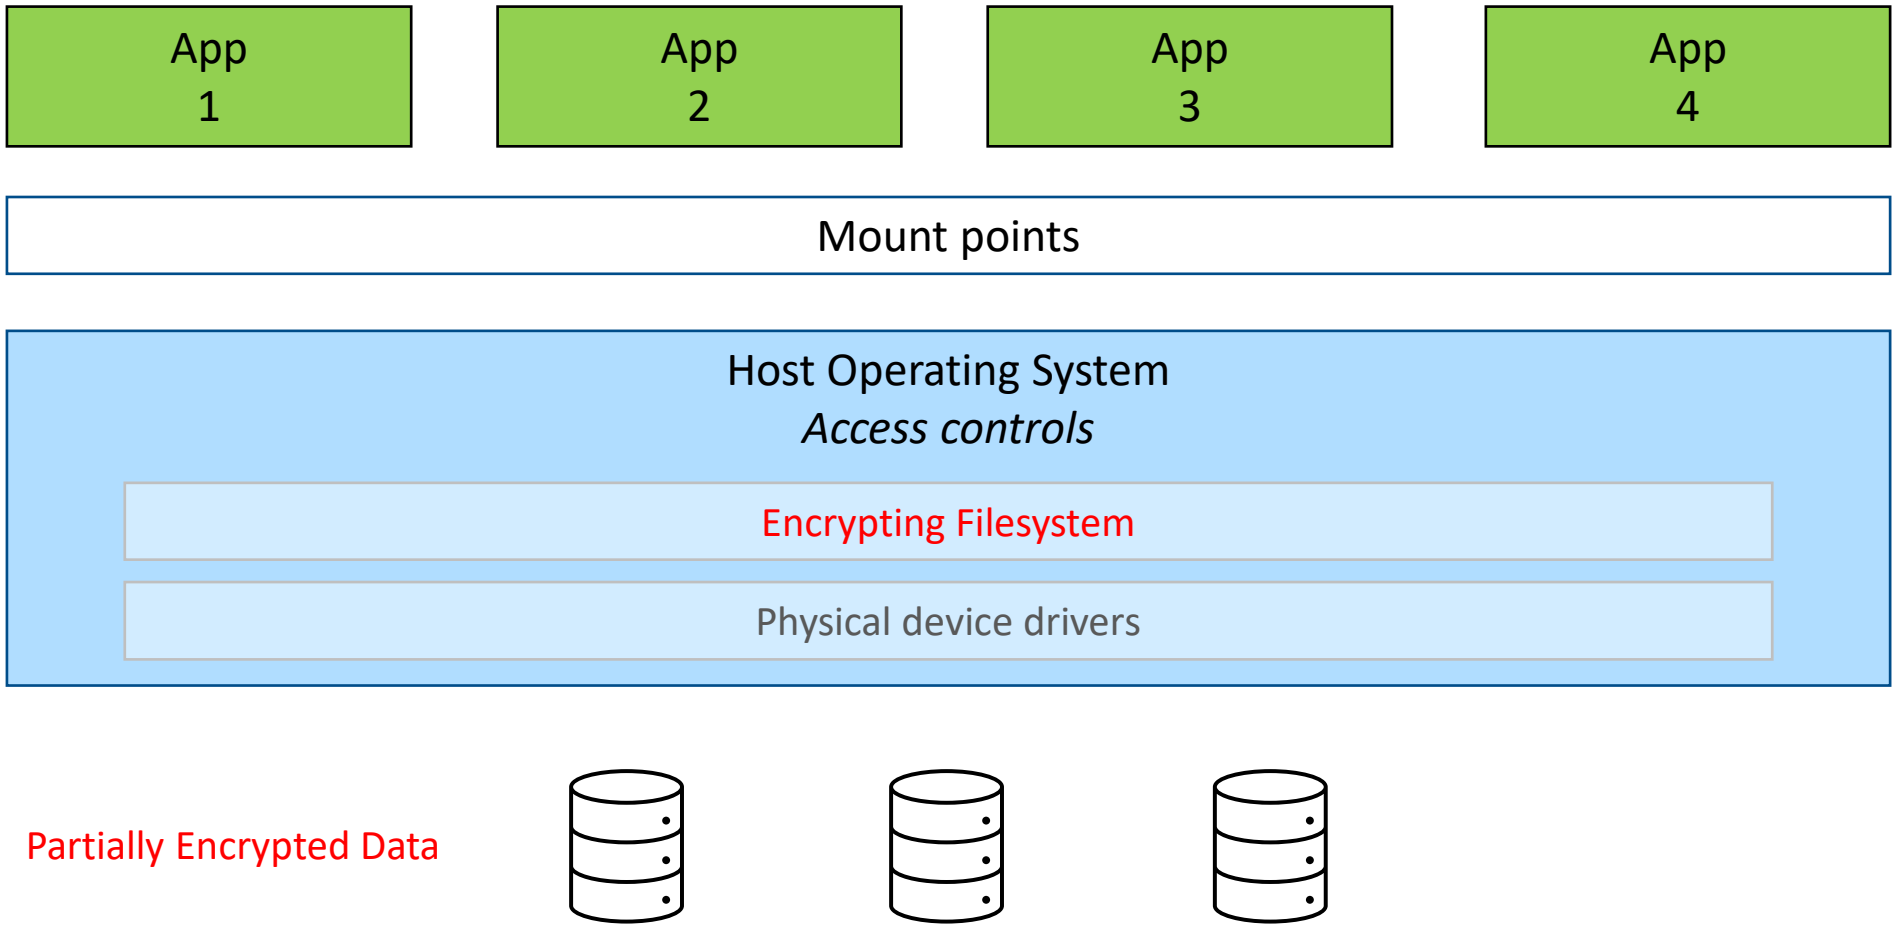

File Based Encryption

- Filenames may still be visible
  - Can disclose sensitive information
- Keys may be vulnerable
  - Protected by a weak user password
  - Administrative access is compromised
- Encryption overhead

Sources of Non-Encrypted Data

- Swap space
- Temporary files (/tmp, local working copies)
- Log files
- Free disk blocks with copies of data
- Overprovisioned space on SSDs

Full Disk Encryption

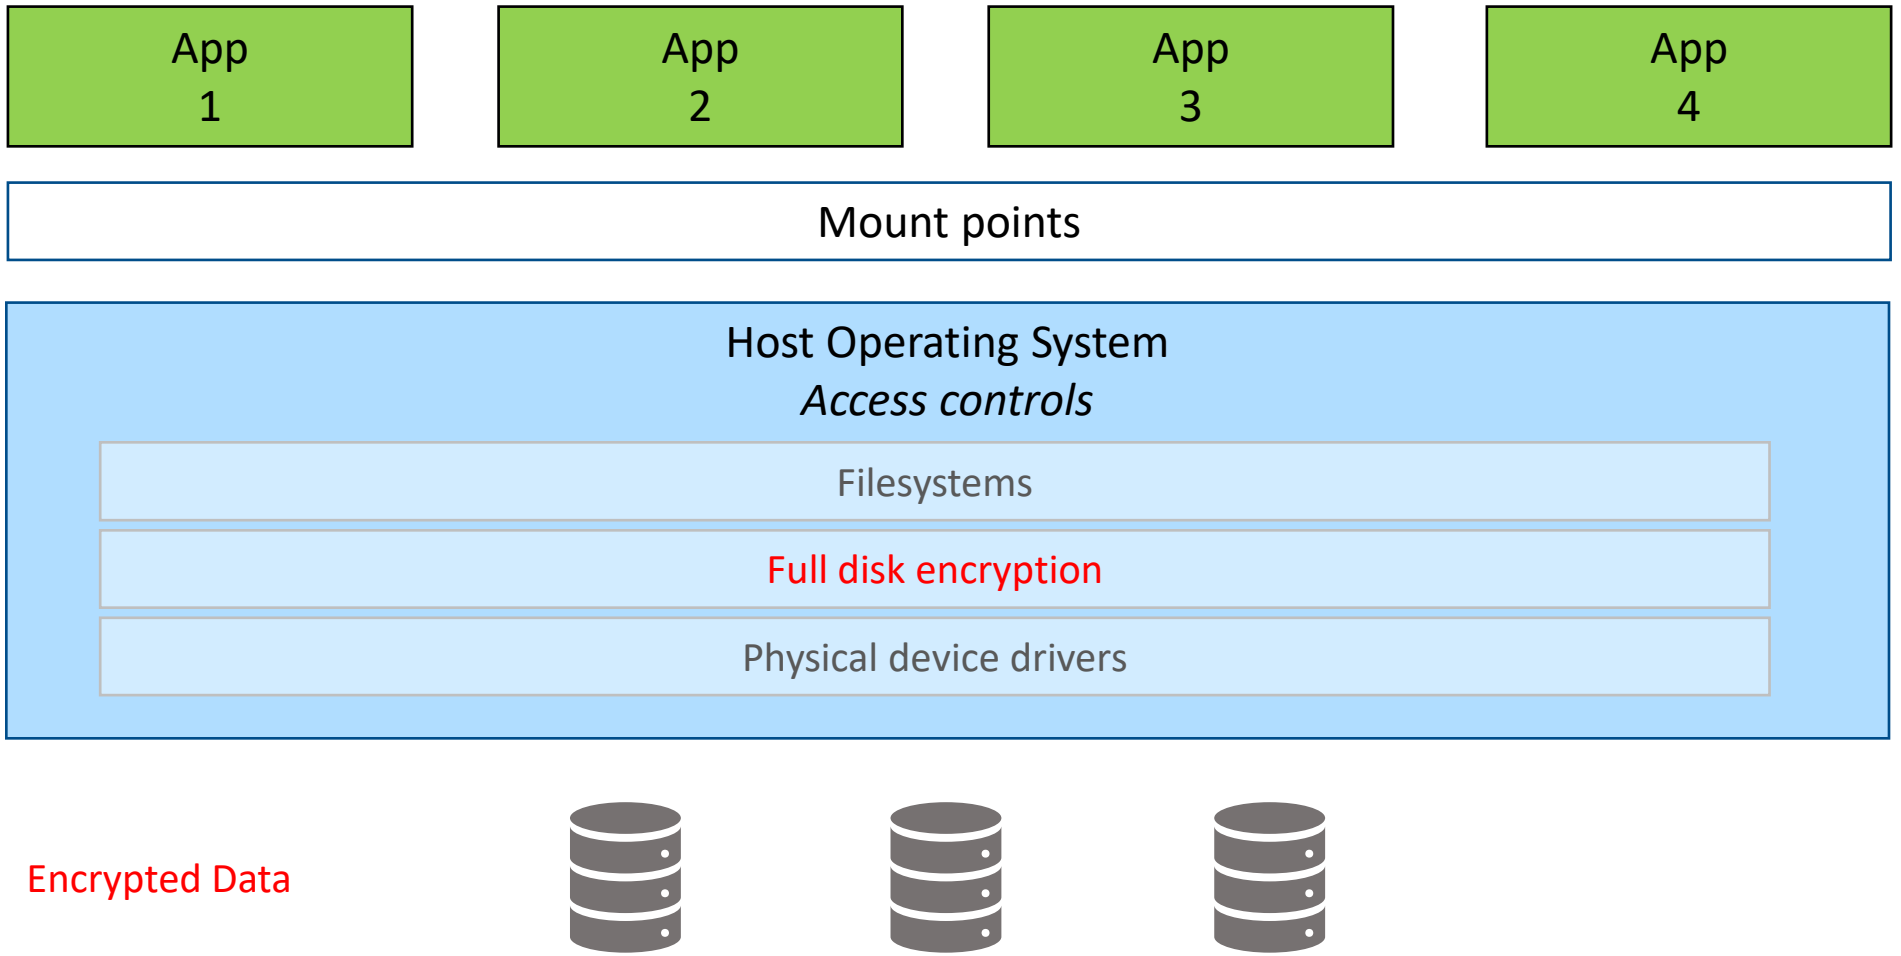

Self Encrypting Drives

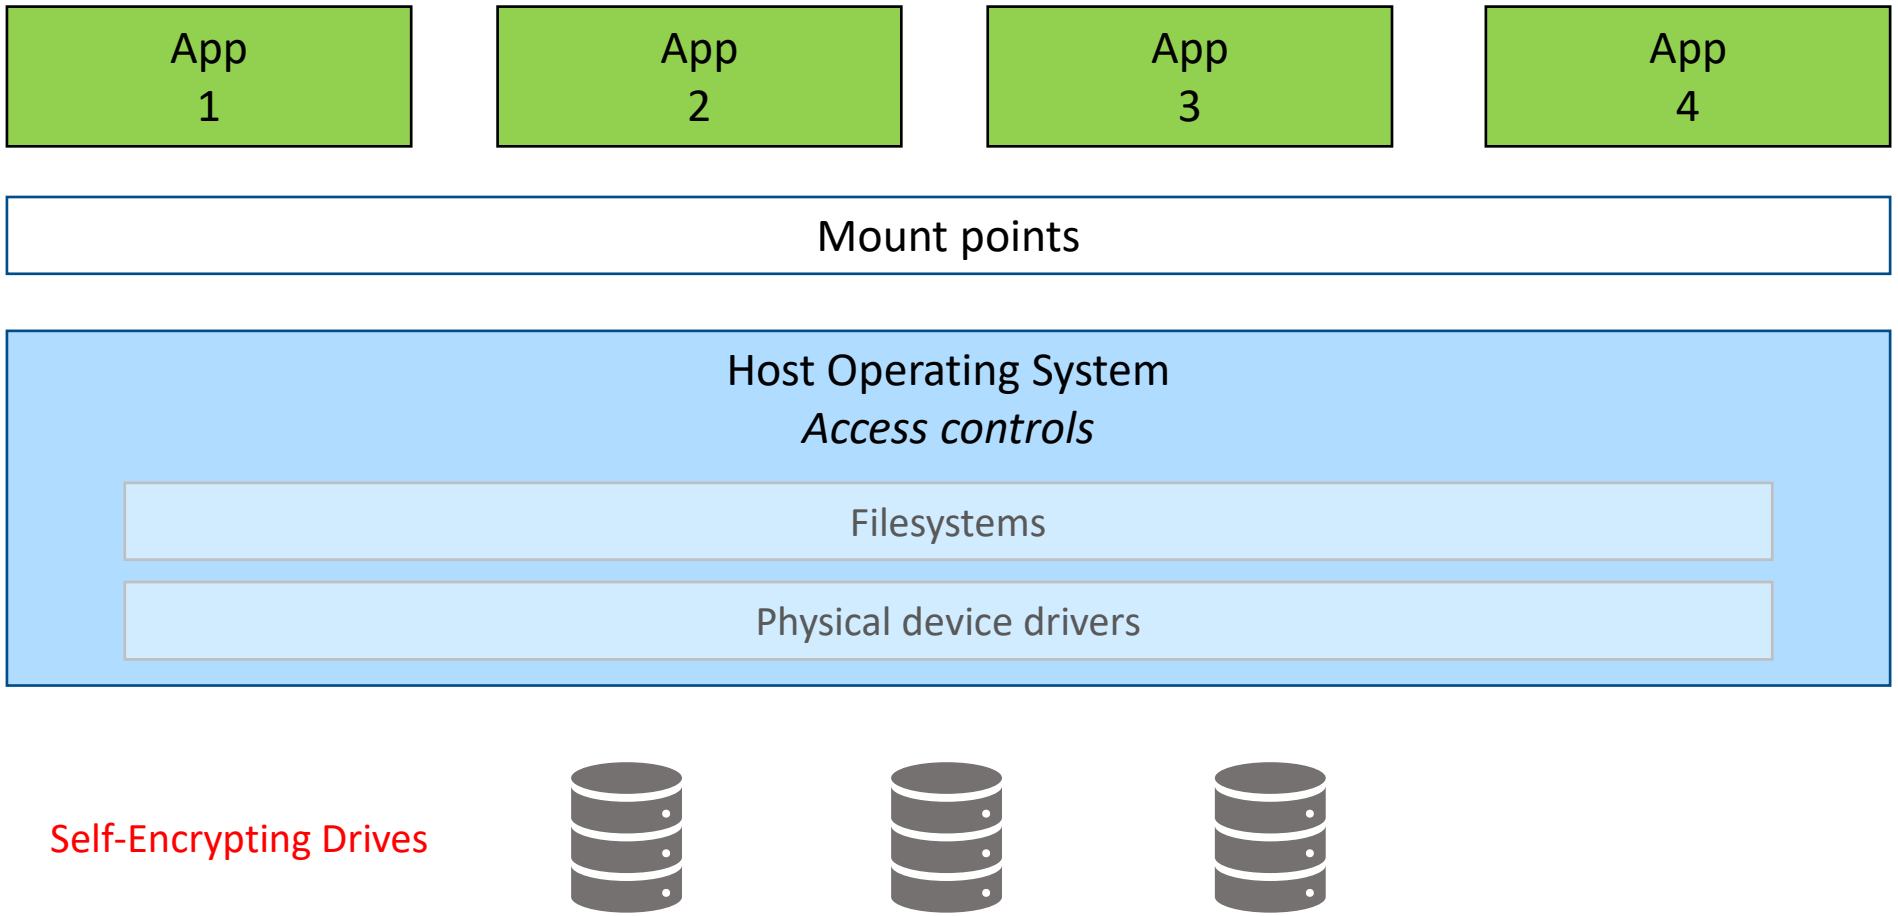

Storage Area Network

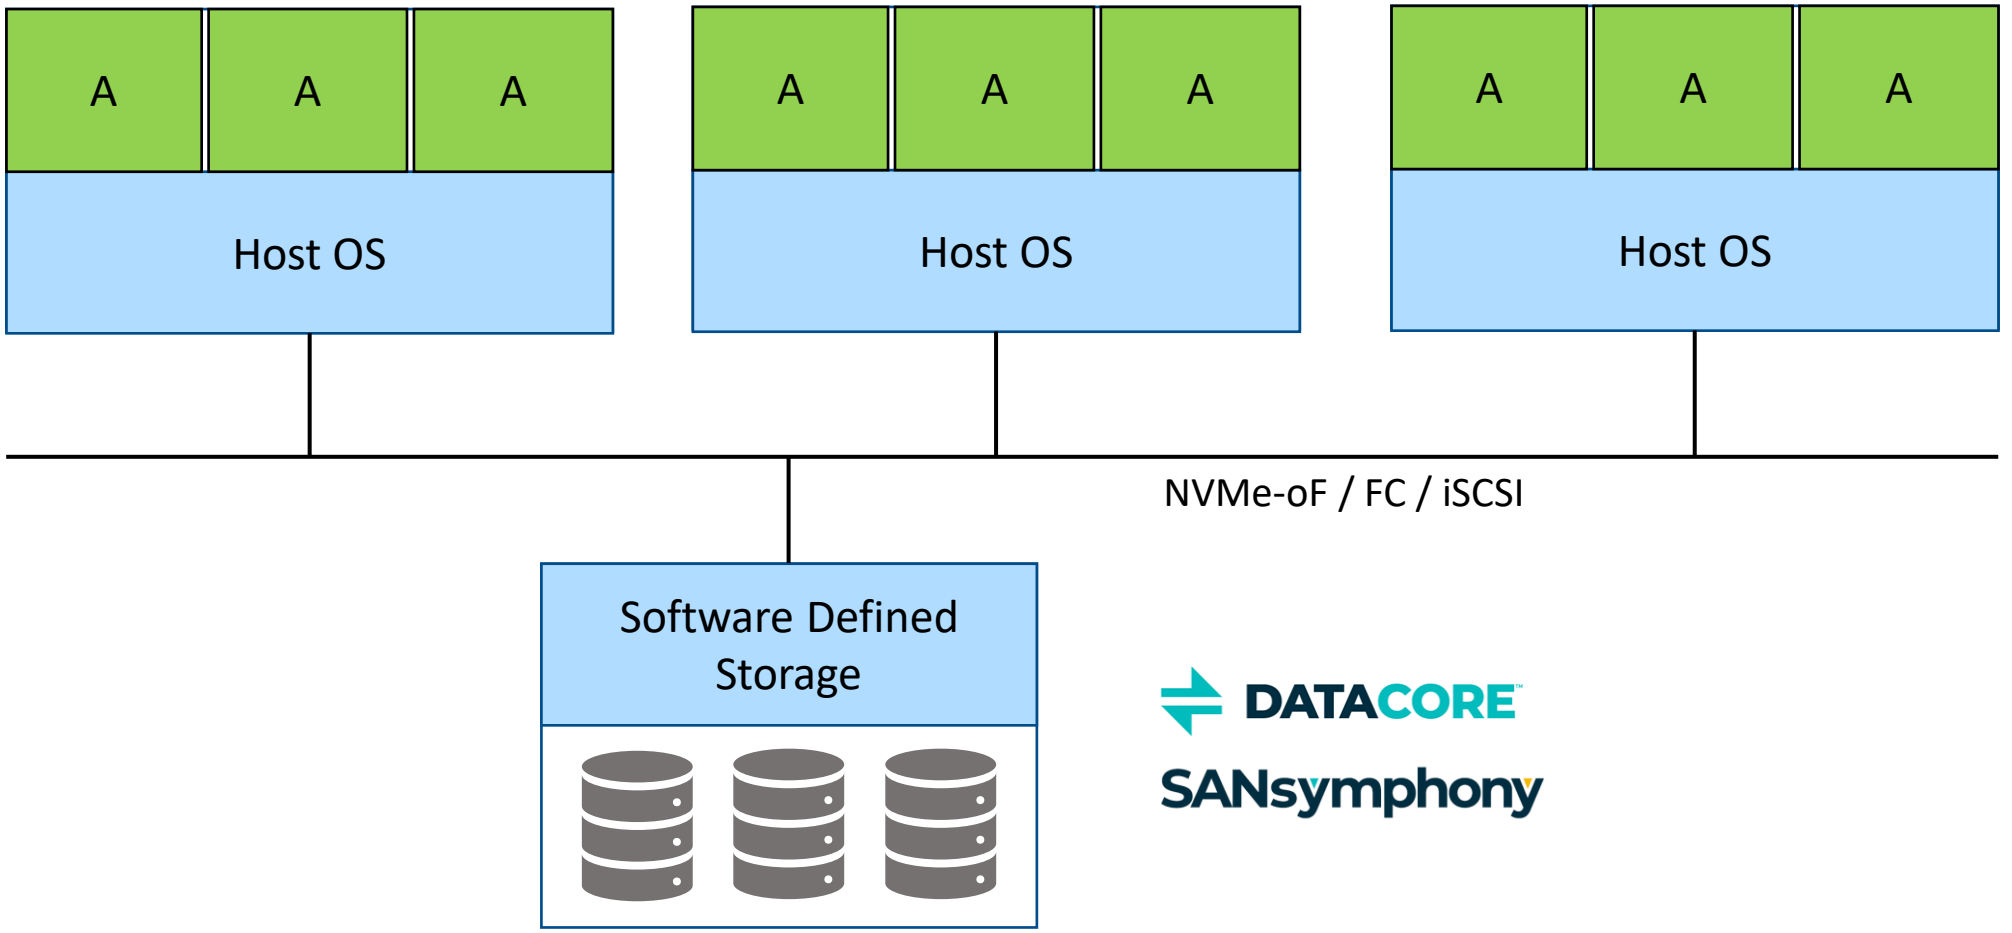

Disaggregated Storage

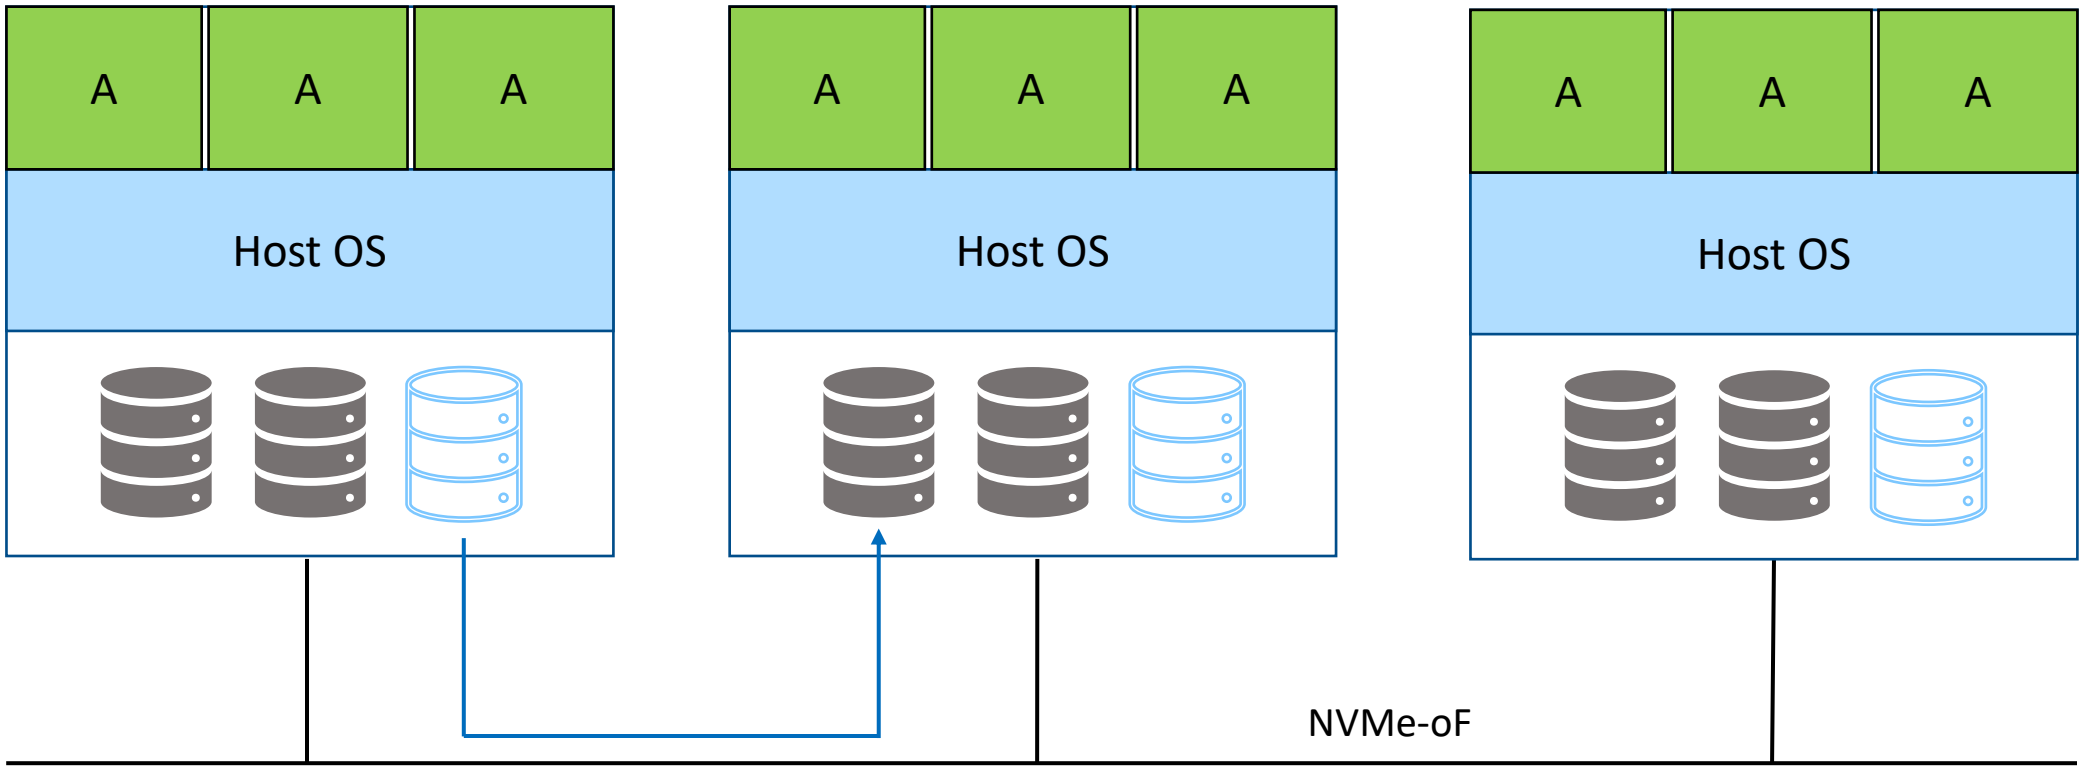

Kubernetes

- Enabled through Storage Class
- Defined with a yaml file
- The exact syntax is specific to each storage provisioner
- E.g., with EBS based storage:

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: slow
provisioner: kubernetes.io/aws-ebs
parameters:
  type: io1
  iopsPerGB: "10"
  fsType: ext4
  encrypted: "true"
```


Securing Data in Use

Secure Enclave

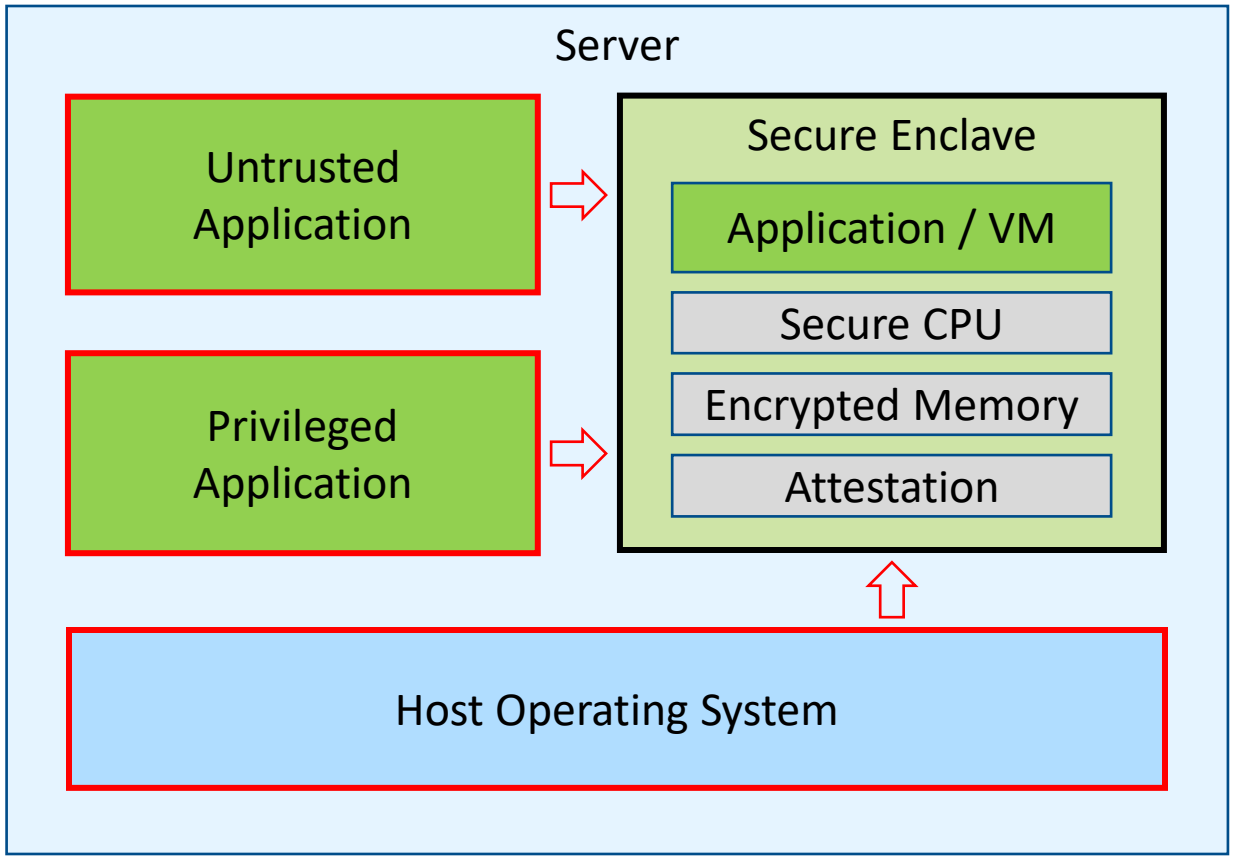

Confidential Computing

- Confidential Computing Consortium
  - Linux Foundation project to define and accelerate adoption
  - https://confidentialcomputing.io/
- Hardware support
  - Intel Software Guard Extension (SGX) - isolates applications
  - AMD Secure Encrypted Virtualization – isolates VMs
  - ARM TrustZone
- Offered by major cloud providers
- E.g., Microsoft Azure
  - Application Enclaves – securely run an application
  - Confidential VMs – run a virtual machine, with optional full disk encryption
  - Confidential Containers

Confidential Kubernetes Nodes

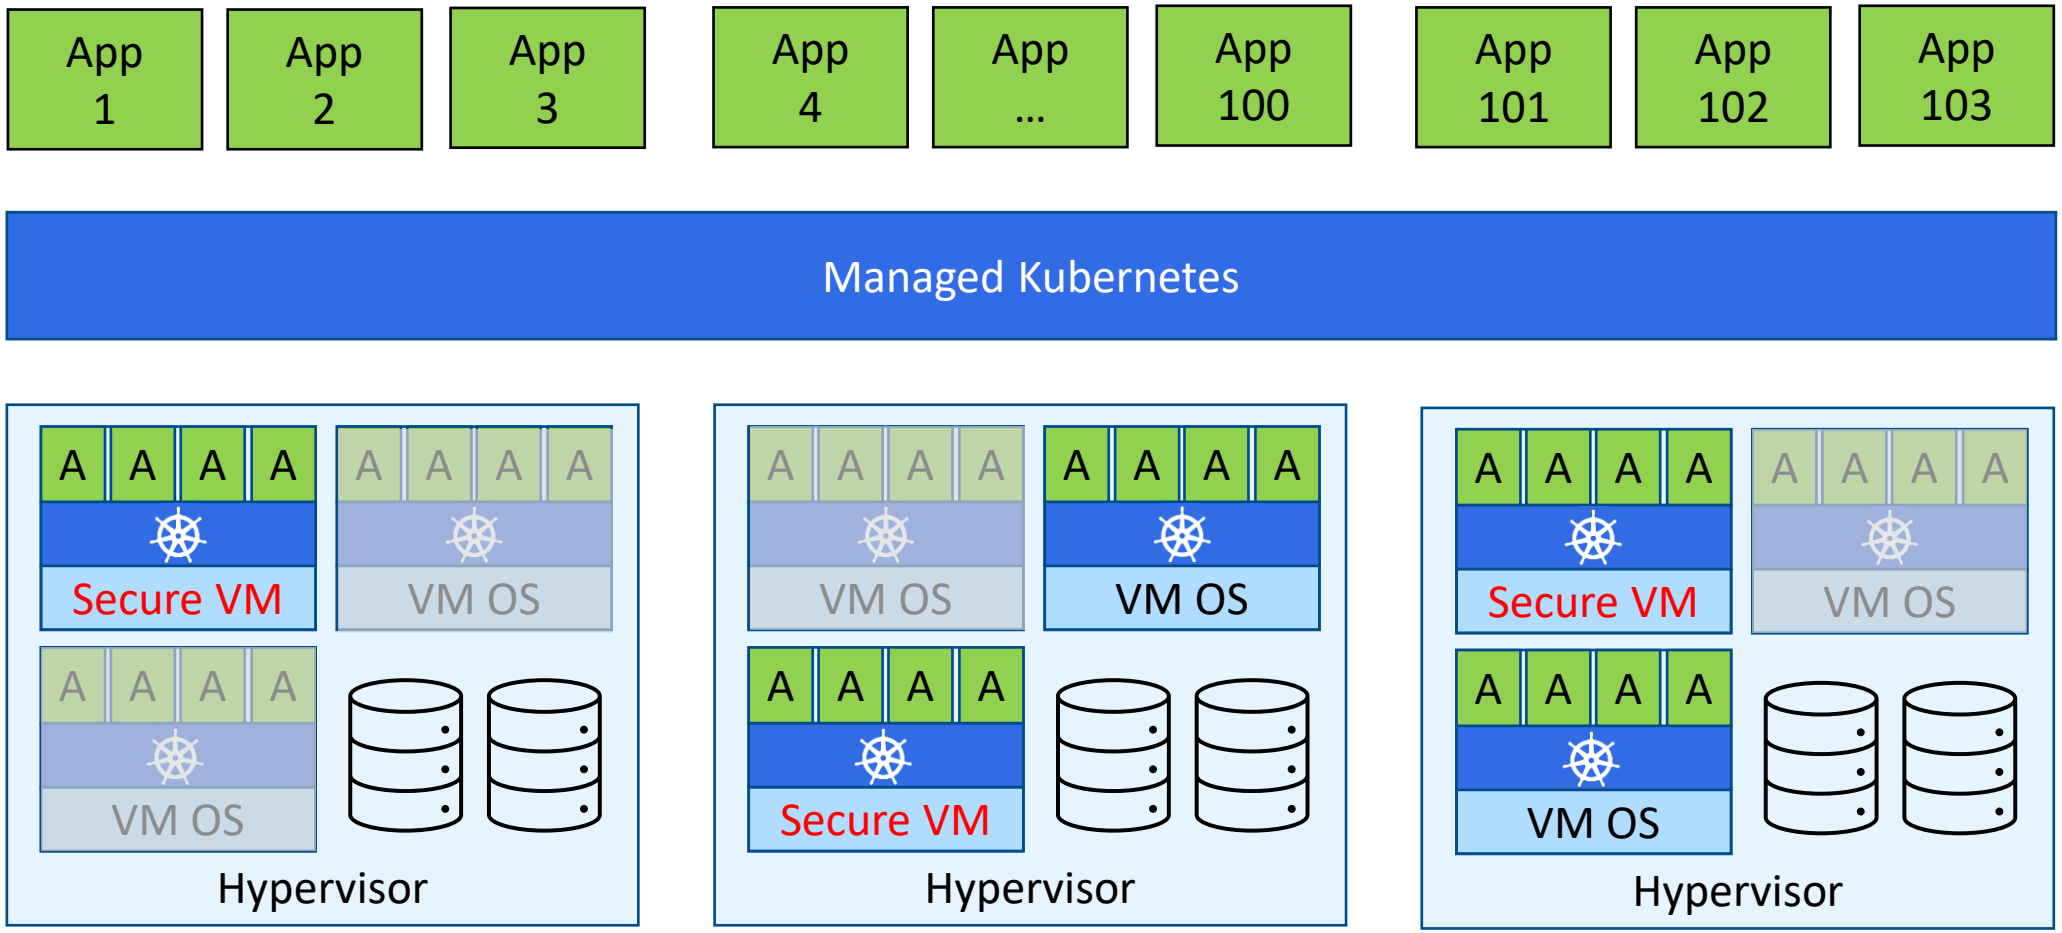

Confidential Containers

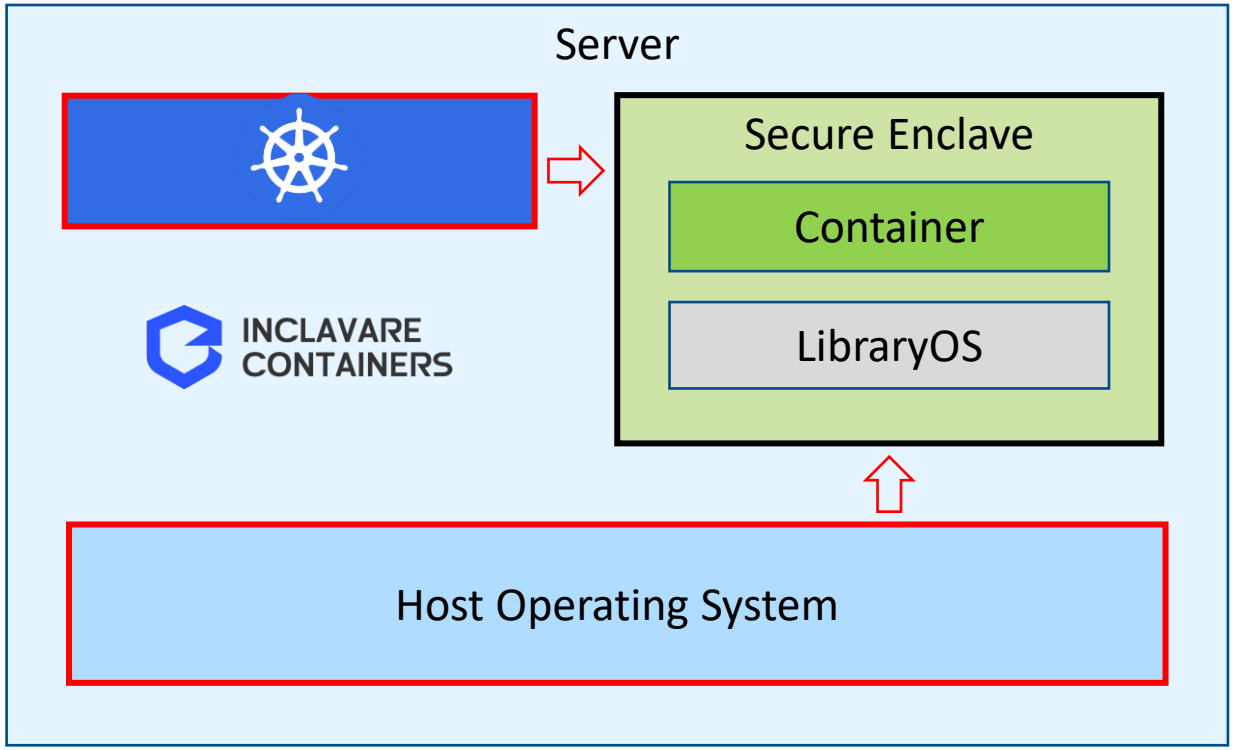

Library OS

- Occlum
  - Memory-safe, multi-process library OS for Intel SGX
  - Written in rust for memory safety
  - https://github.com/occlum/occlum
- Gramine
  - Lightweight library OS with Intel SGX support
  - Designed to run a single application
  - Minimal host requirements
  - https://github.com/gramineproject/gramine

Securing Data in Motion

S3 Object Store

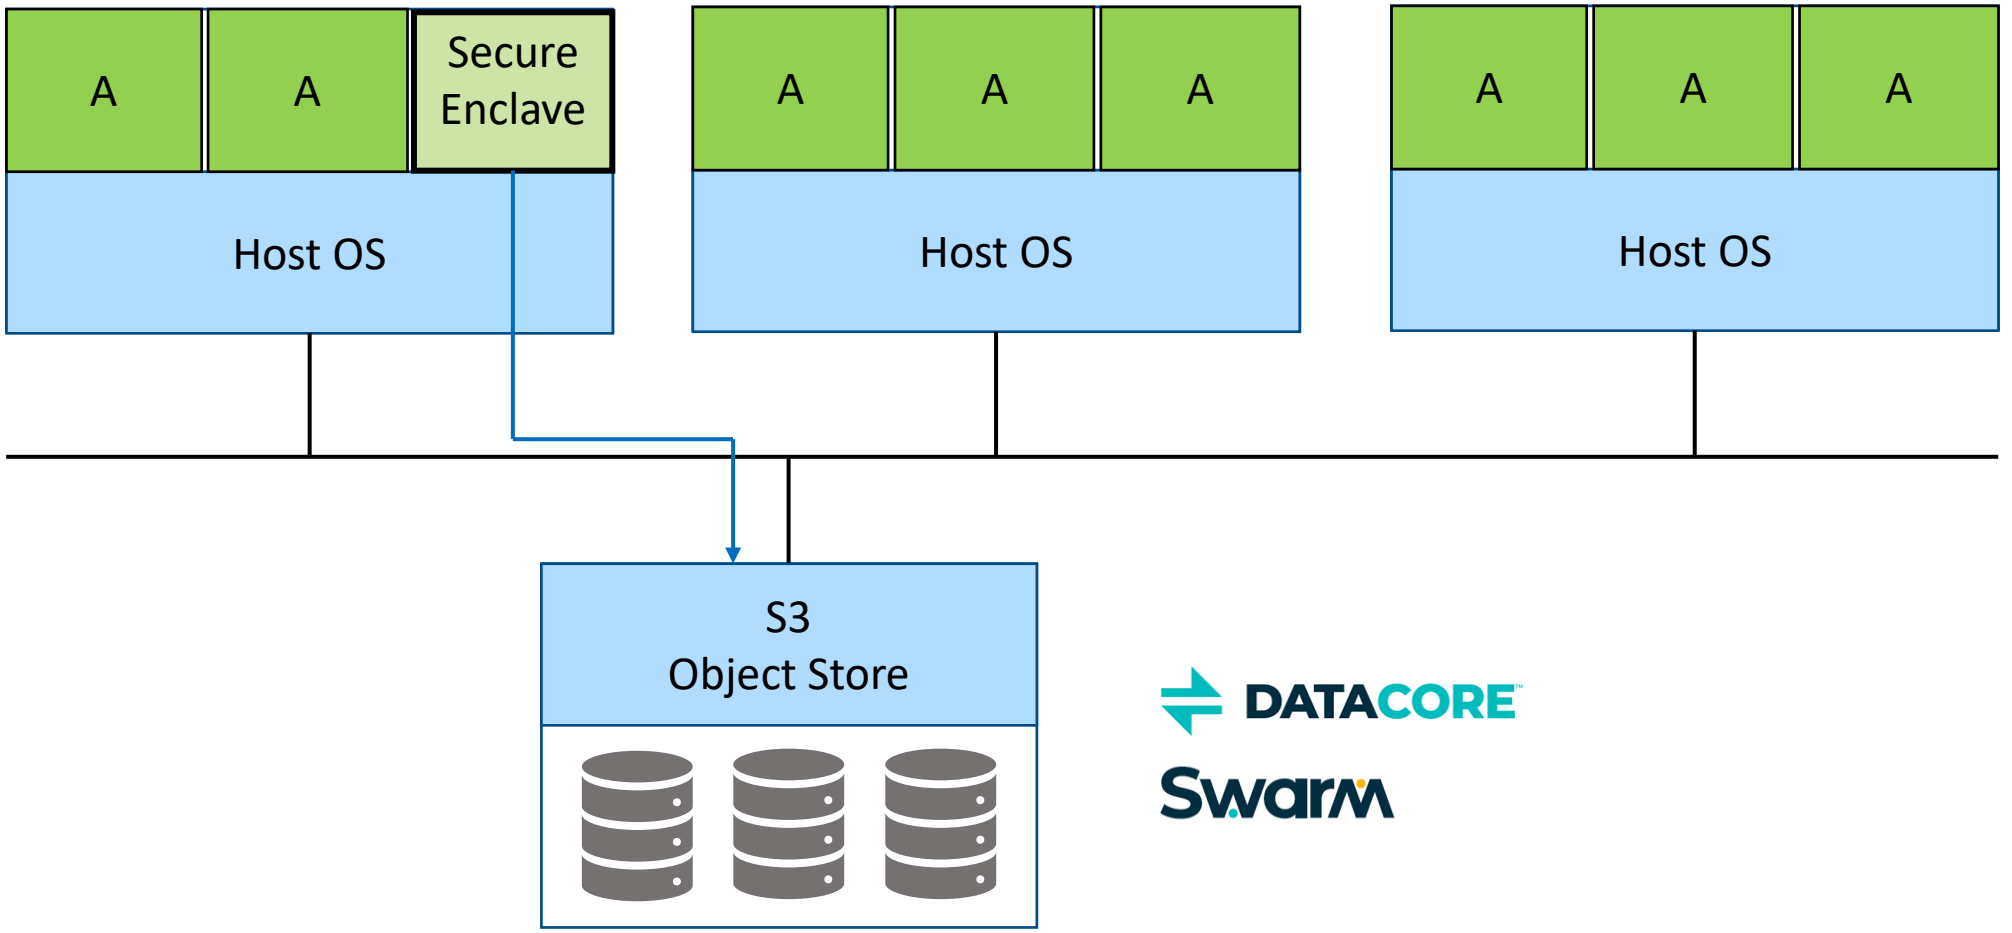

DATACORE

Swarm

File I/O

- Occlum supports multiple filesystems
- Read-only hashed filesystem
  - For integrity protection
- Writable encrypted filesystem
  - For confidentiality protection
- Untrusted host filesystem
  - For data exchange with the host OS

Storage Performance Development Kit (SPDK)

- Tools and libraries for writing:
  - High performance, scalable
  - User-mode storage applications
- Cutting Edge
  - Leverage the latest NVMe features
  - Poll-mode and event-loop for maximum performance
  - Lockless, thread-per-core design
- https://www.spdk.io

Direct Connect

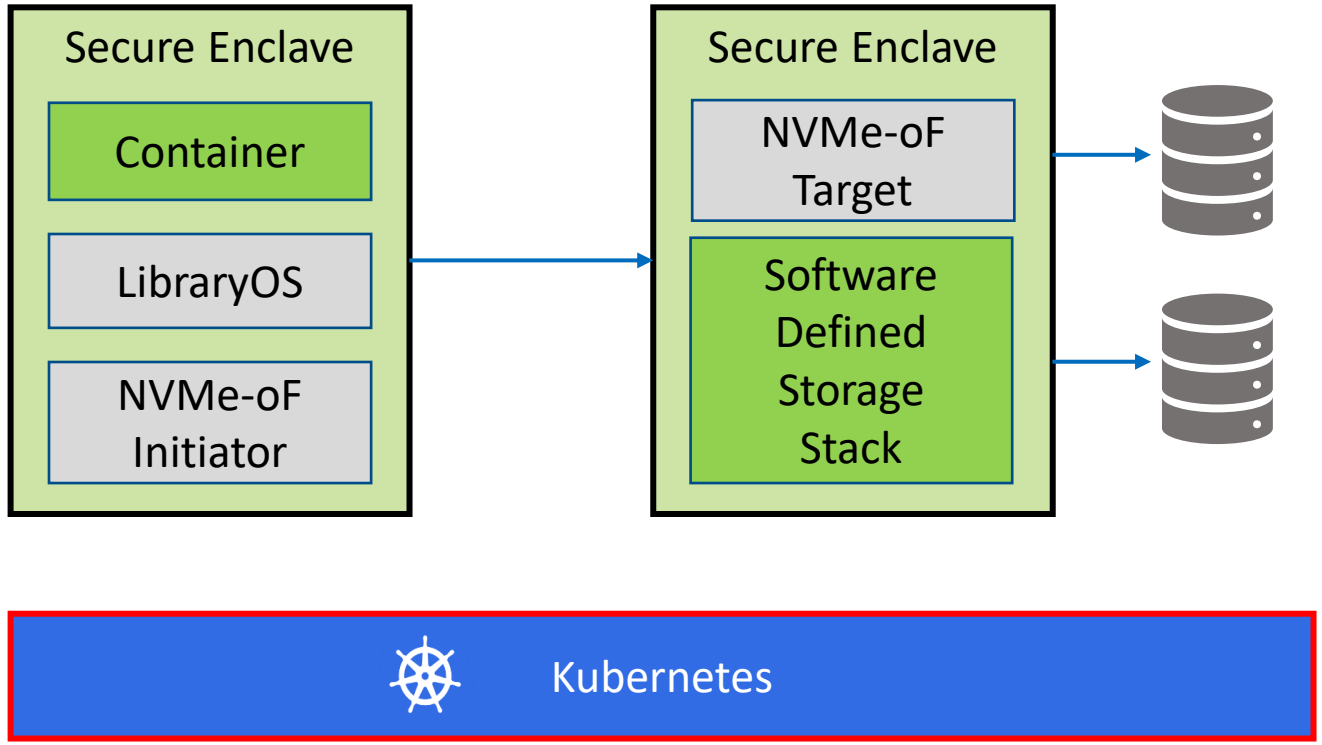

Long Term Research

CHERI (Capability Hardware Enhanced RISC Instructions)

- Joint research project between SRI International and the University of Cambridge
- Revisit fundamental design choices to dramatically improve system security
- Extends instruction set to enable fine-grained memory protection
  - Pointers have associated bounds and permissions
  - Invalid memory references throw an exception
- Long-term direction is towards software compartmentalization
  - Hardware ‘capabilities’ enforce software isolation
  - Granular and scalable data sharing
- https://www.cl.cam.ac.uk/research/security/ctsrd/cheri/

Arm Morello Program

- Five-year research program (launched in 2019)
- Defines a new prototype security architecture based on CHERI (Capability Hardware Enhanced RISC Instructions)
- Morello is a research and prototyping program
  - To determine if Morello prototype architecture is viable
  - To create more secure hardware architecture for processors of the future
- System on Chip (SoC) implementation of the architecture
  - Will provide a Digital Security by Design (DSdB) technology platform prototype
  - Enabling industry and academic partners to test real-world use cases and inform future development
- https://www.arm.com/morello

Conclusion

Conclusion

- Security is always a trade off between:
  - Sensitivity of the data
  - Risk
  - Practicality
  - Cost

- In a containerized environment, securing data:
  - At rest, is vendor specific, but straightforward
  - In use, is achievable but has higher costs and is in the ‘early adopters’ phase
  - In motion, is viable with an object store, but cutting edge for block storage

Questions?

- Please ask questions in Slack

Please take a moment to rate this session.

Your feedback is important to us.

