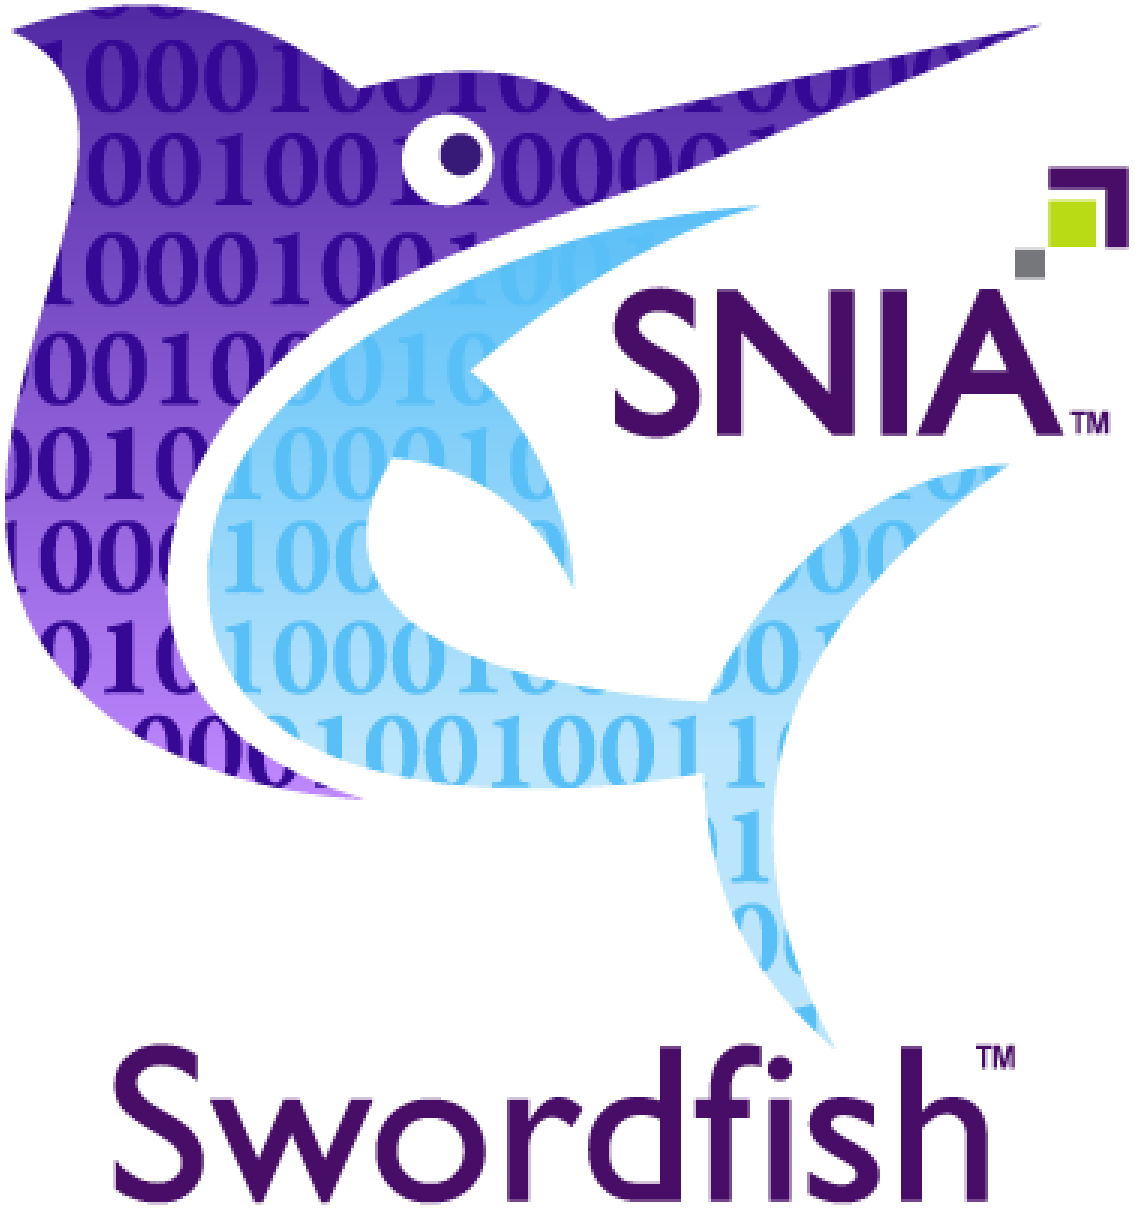

Flash Memory Summit

Catch the Wave – Managing NVMe-oF™ in the Enterprise

Richelle Ahlvers, Intel

Curtis Ballard, HPE

About the Presenter

Richelle Ahlvers
Storage Technology
Enablement Architect, Intel

Richelle Ahlvers is a Storage Technology Enablement Architect at Intel, where she promotes and drives enablement of new technologies and standards strategies. Richelle has spent over 25 years in Enterprise R&D teams in a variety of technical roles, leading the architecture, design and development of storage array software, storage management software user experience projects including mobility, developing new storage industry categories including SAN management, storage grid and cloud, and storage technology portfolio solutions.

Richelle has been engaged with industry standards initiatives for many years and is actively engaged with many groups supporting manageability including SNIA, DMTF, NVMe, OFA and UCIe. She is Vice-Chair of the SNIA Board of Directors, Chair of the Storage Management Initiative, leads the SSM Technical Work Group developing the Swordfish Scalable Storage Management API, and has also served as the SNIA Technical Council Chair and been engaged across a breadth of technologies ranging from storage management, to solid state storage, to cloud, to green storage. She also serves on the DMTF Board of Directors as the VP of Finance and Treasurer.

About the Presenter

Curtis Ballard

Strategist, Emerging Storage Technology, Hewlett Packard Enterprise

Curtis Ballard is a Distinguished Technologist with Hewlett Packard Enterprise in the HPE Storage organization. He works on emerging storage technologies and industry engagement. Curtis is on the NVM Express Board of Directors, the SNIA Technical Council, and is the vice-chair of the INCITS SCSI Storage Interfaces Technical Committee. Curtis has worked on storage products across the industry including host drivers, interface ASICS, standalone drives, and complex storage systems.

Flash Memory Summit

Abstract

- The enterprise storage market is rapidly expanding to include NVMe and NVMe-oF products pervasively. This provides a challenge: how do you manage these as part of your enterprise datacenter?
  - As the NVM Express family of specifications continue to develop, the corresponding Swordfish management capabilities are also evolving: the SNIA Swordfish specification has expanded to include full NVMe and NVMe-oF enablement and alignment across DMTF, NVMe, and SNIA for NVMe and NVMe-oF use cases.

    In conjunction with Redfish management of servers, Swordfish's capabilities to manage NVMe and NVMe-oF devices in the enterprise provide a seamless management ecosystem. Dive in and catch up on the latest developments of the SNIA Swordfish specification:
- This presentation will provide an introduction to managing NVMe and NVMe-oF with Swordfish, using an example of this functionality introducing the complexity of discovery controllers with the simplified model presented to Swordfish clients.

NVMe Subsystem Model

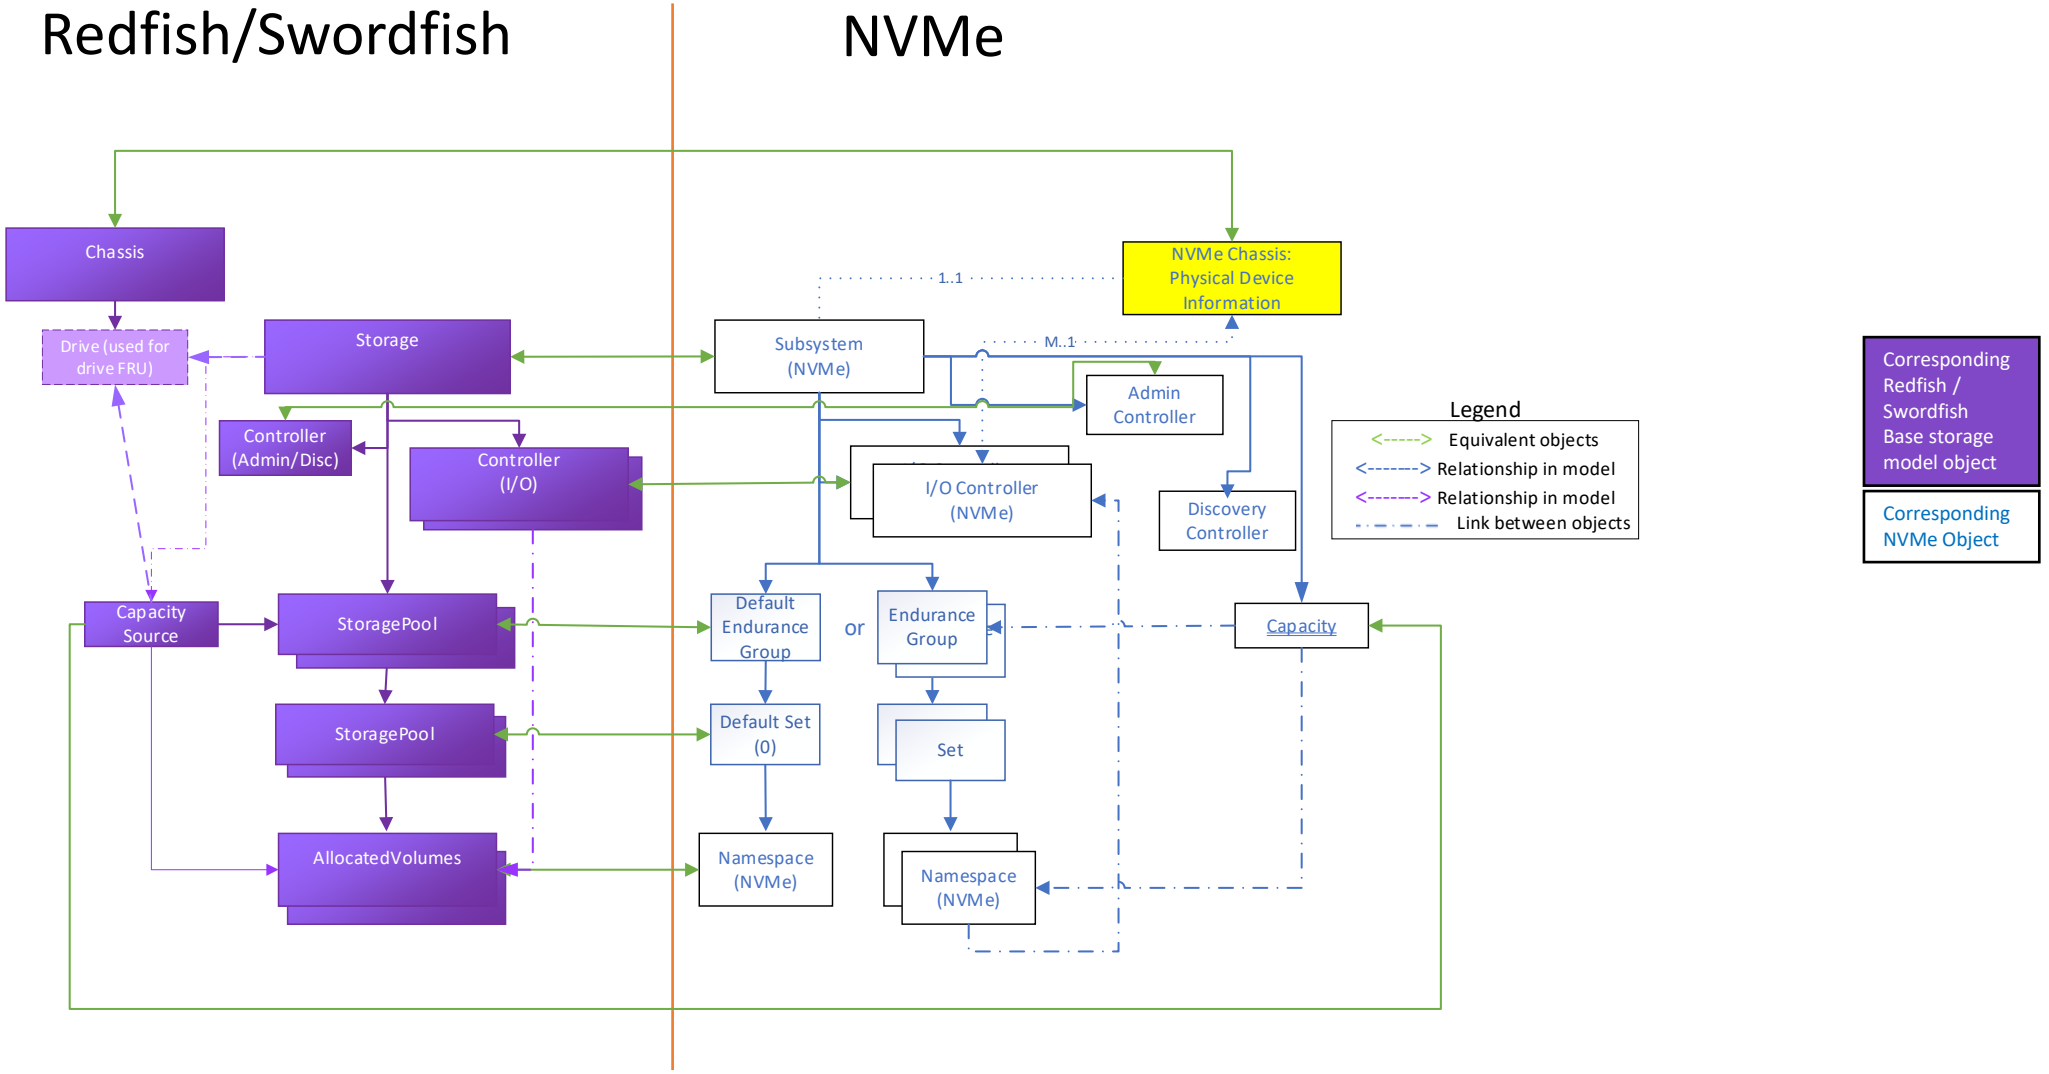

Swordfish and NVMe: Basic Functionality

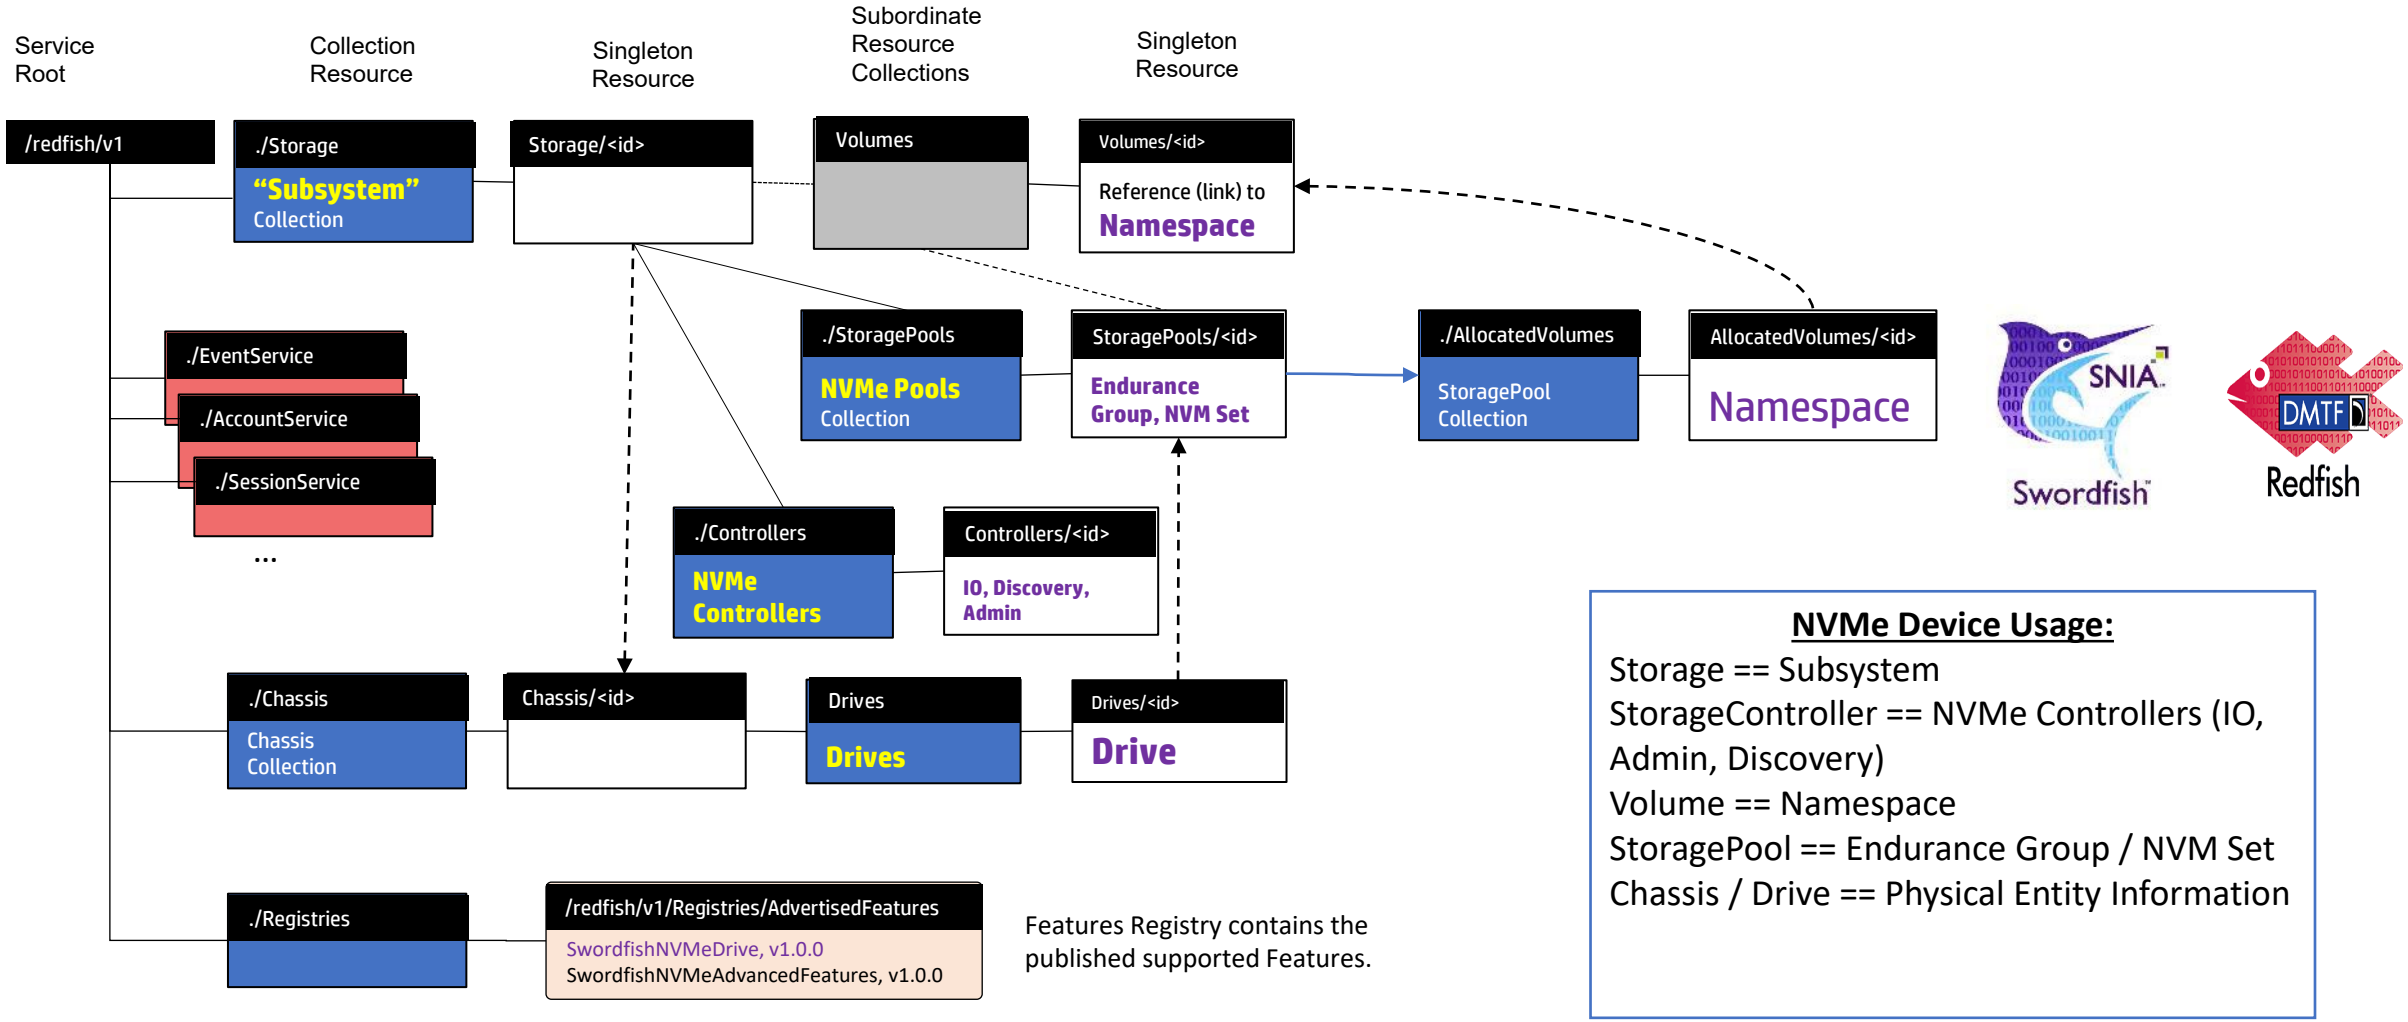

NVMe Device Usage:

Storage == Subsystem
StorageController == NVMe Controllers (IO, Admin, Discovery)
Volume == Namespace
StoragePool == Endurance Group / NVM Set
Chassis / Drive == Physical Entity Information

Adding Network and fabric…

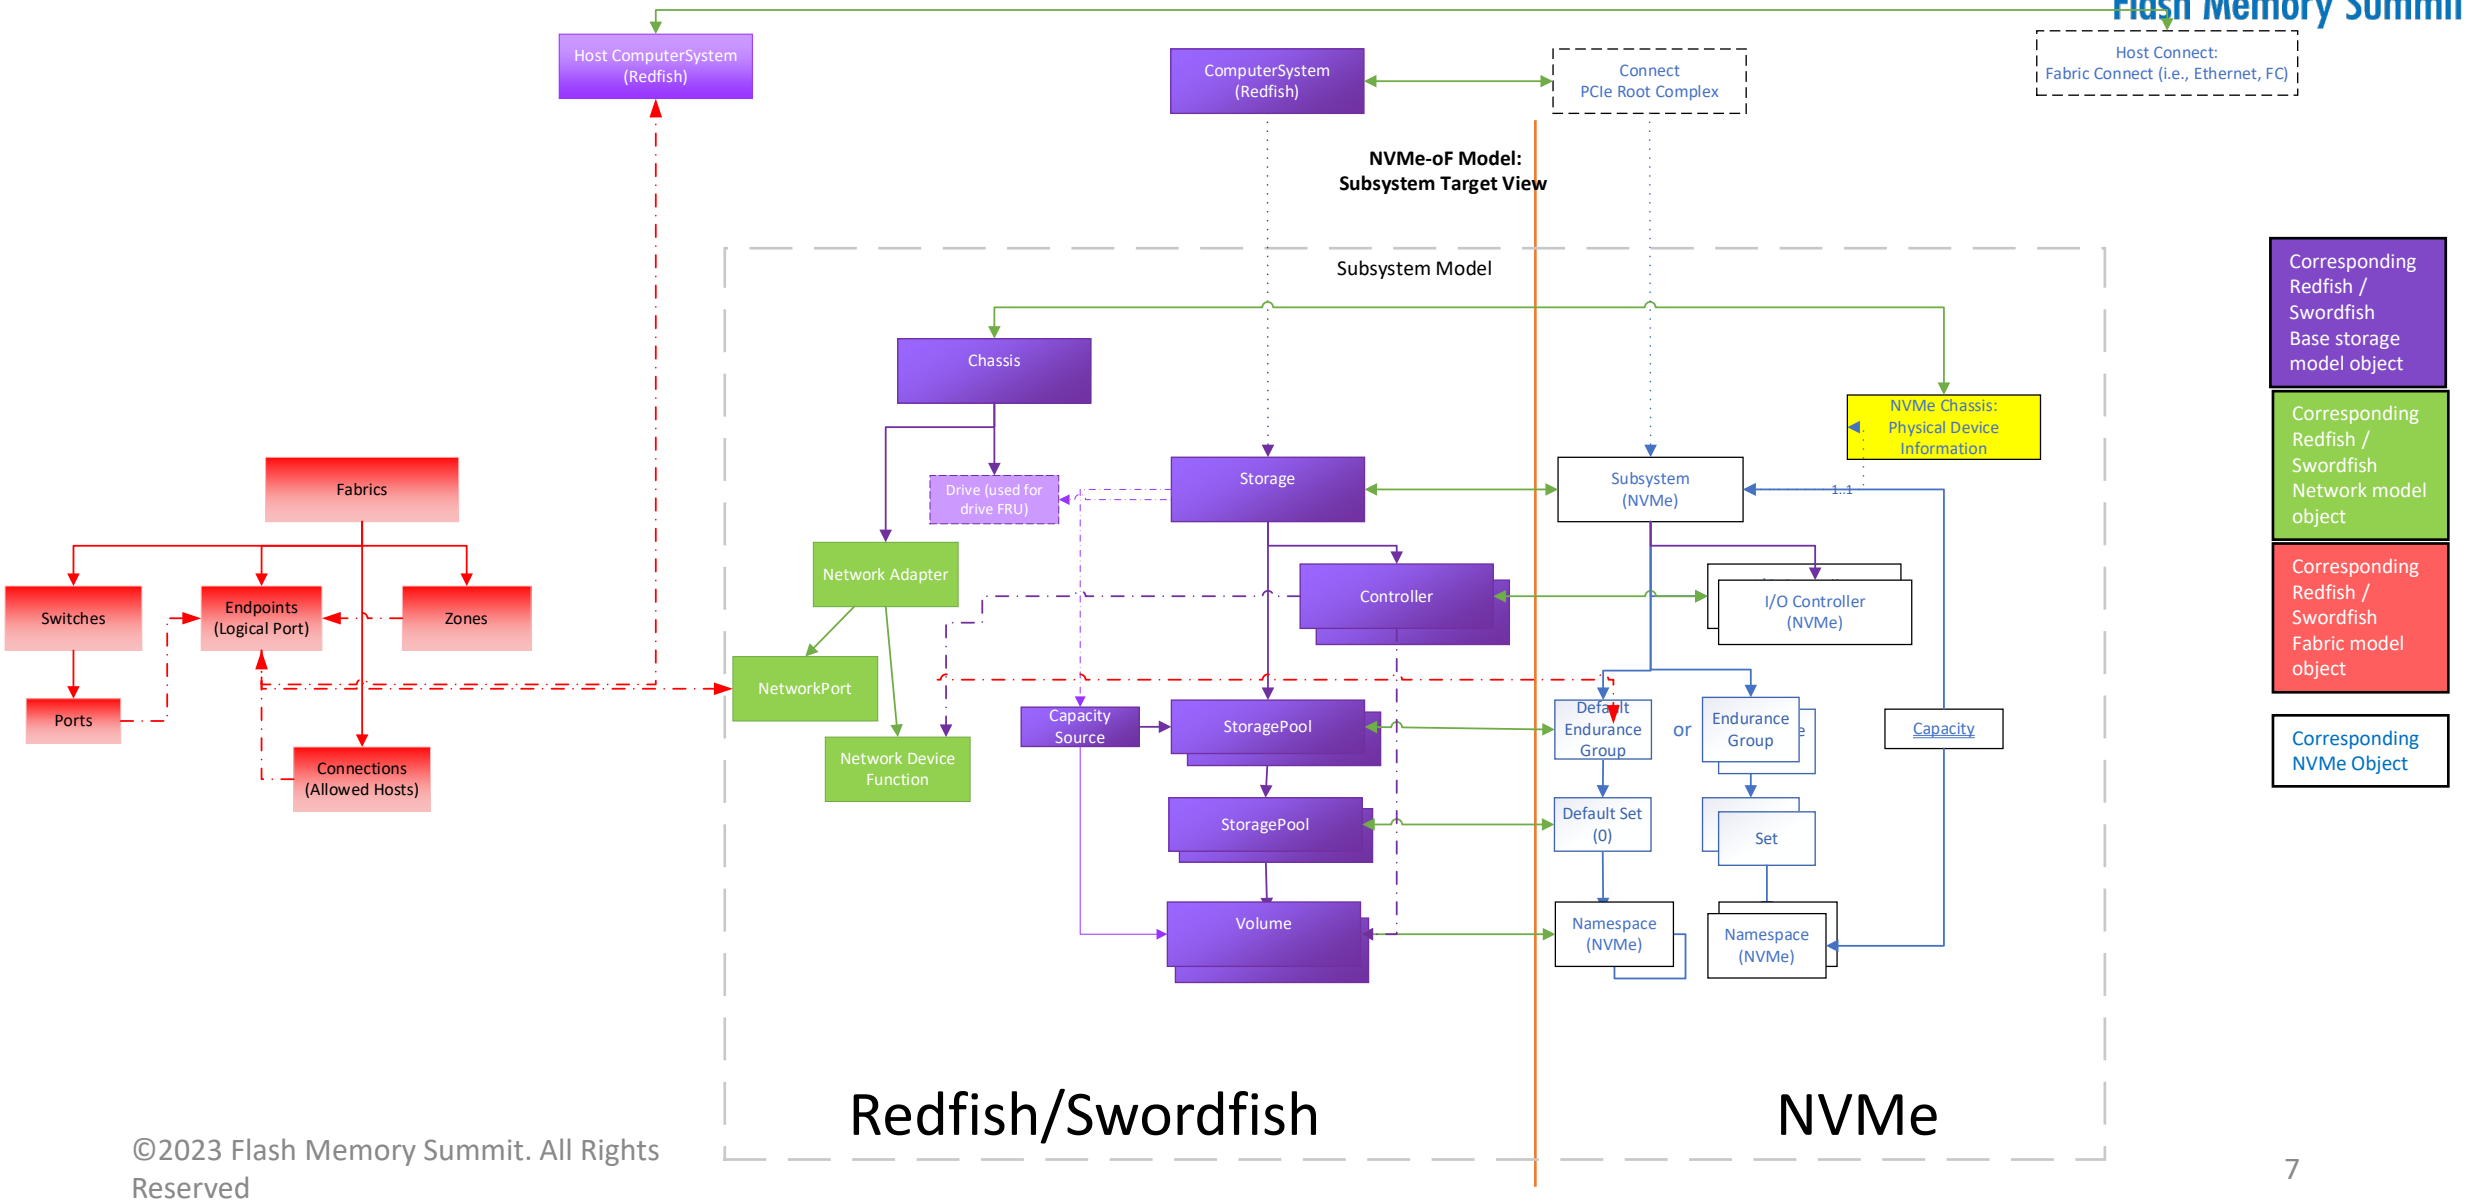

Redfish/Swordfish Hierarchy: Managing Extended Connectivity

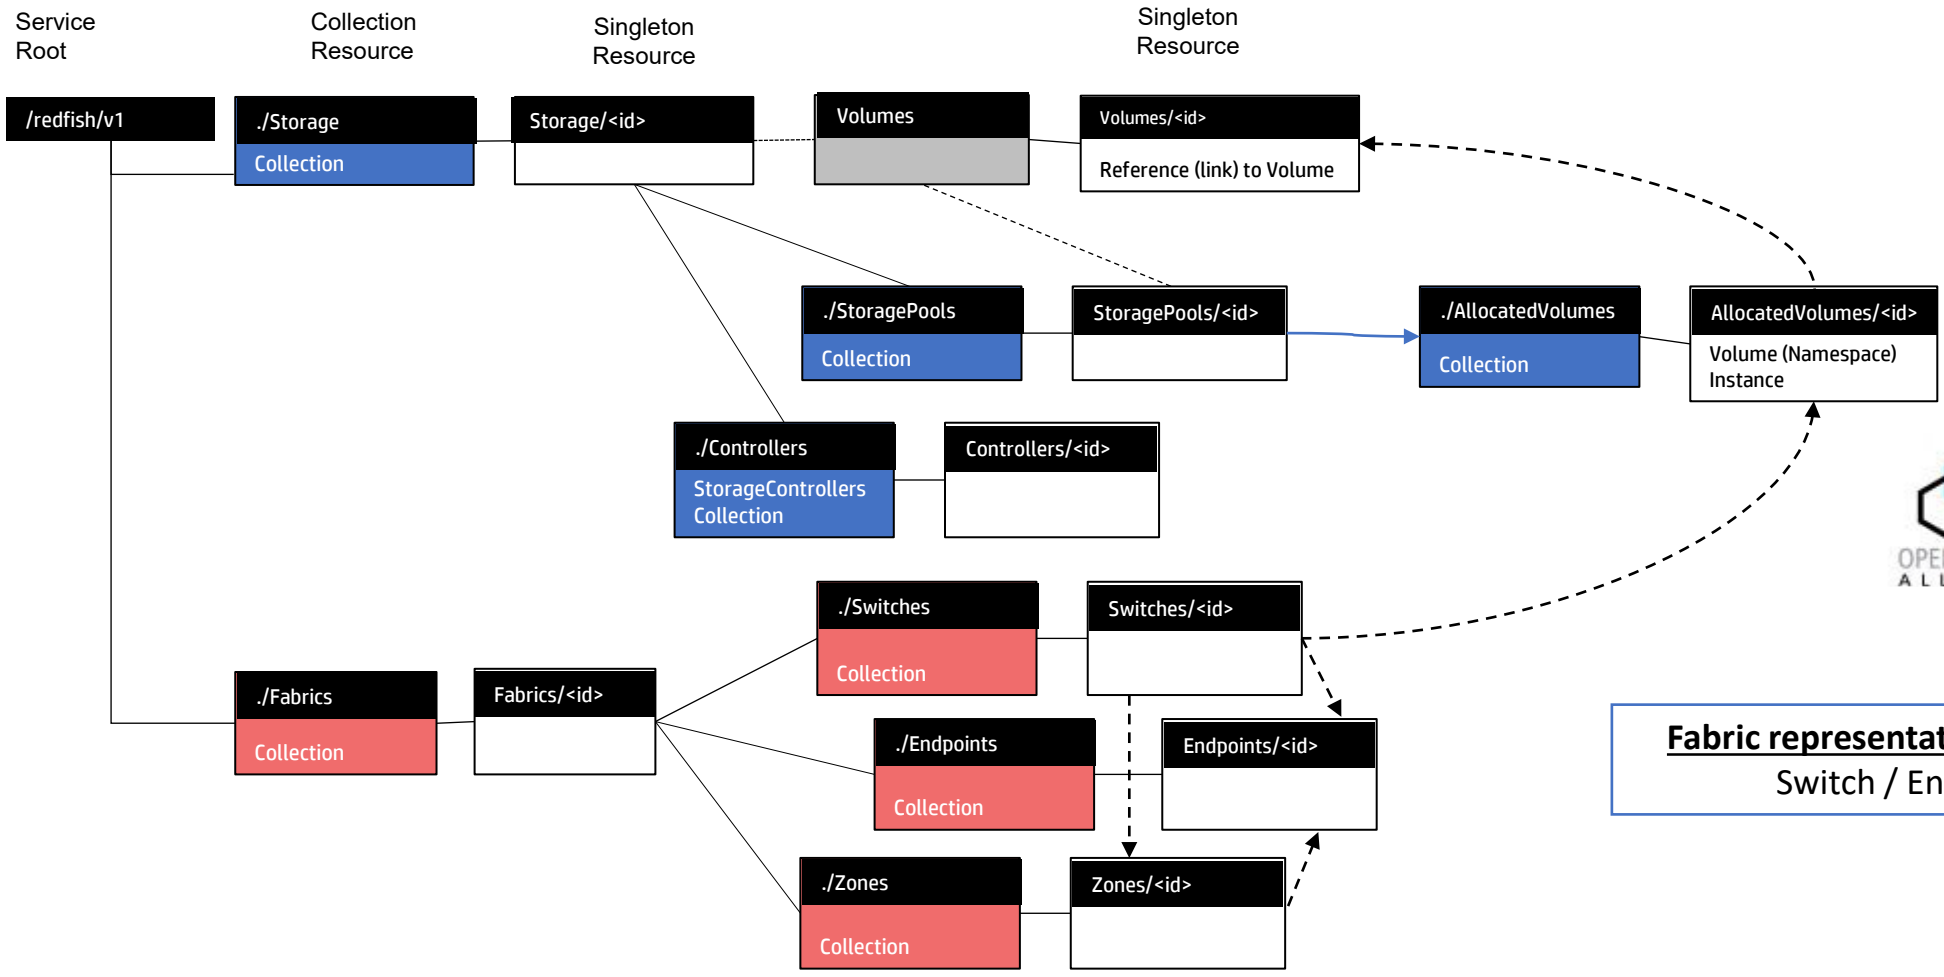

Redfish/Swordfish Hierarchy: Adding Multi-System Access Management

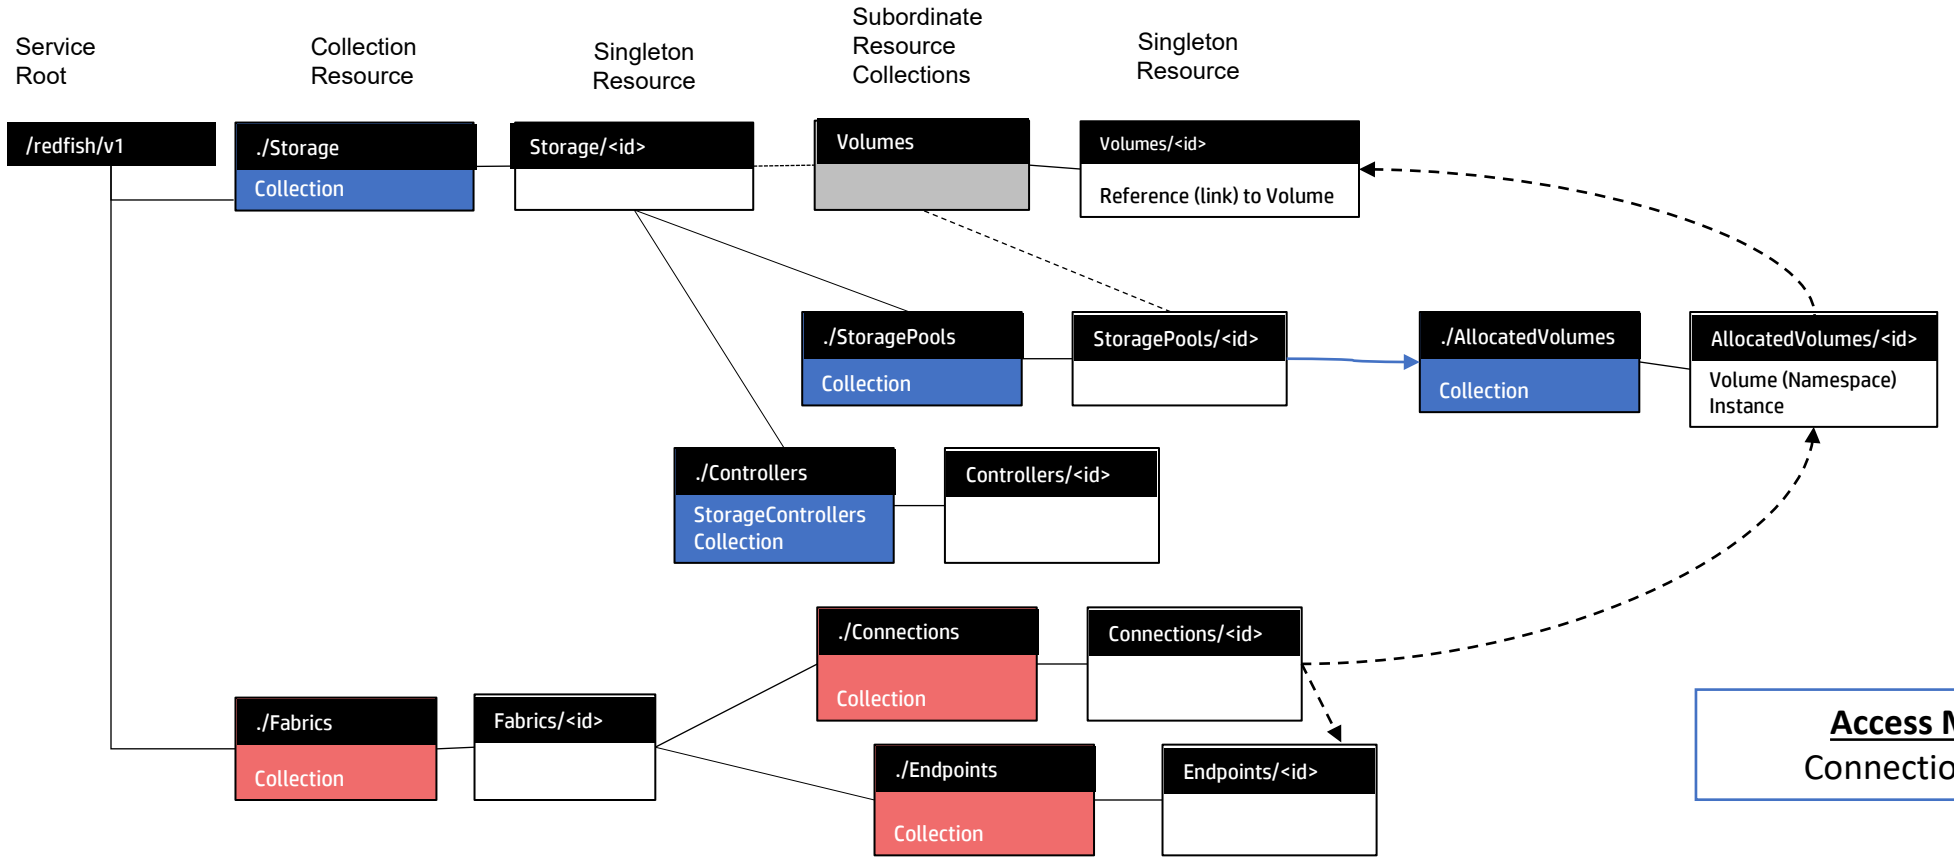

Redfish

Swordfish

Example

- Let’s look at an example of a complex NVMe-oF concept, and how it is modeled in Swordfish and presented to end users and clients.

NVMe-oF™ Storage Device Discovery

- NVMe-oF storage device discovery uses ***Discovery Controllers***
- Two types
  - Direct Discovery Controller
  - Centralized Discovery Controller
    (Next slide)

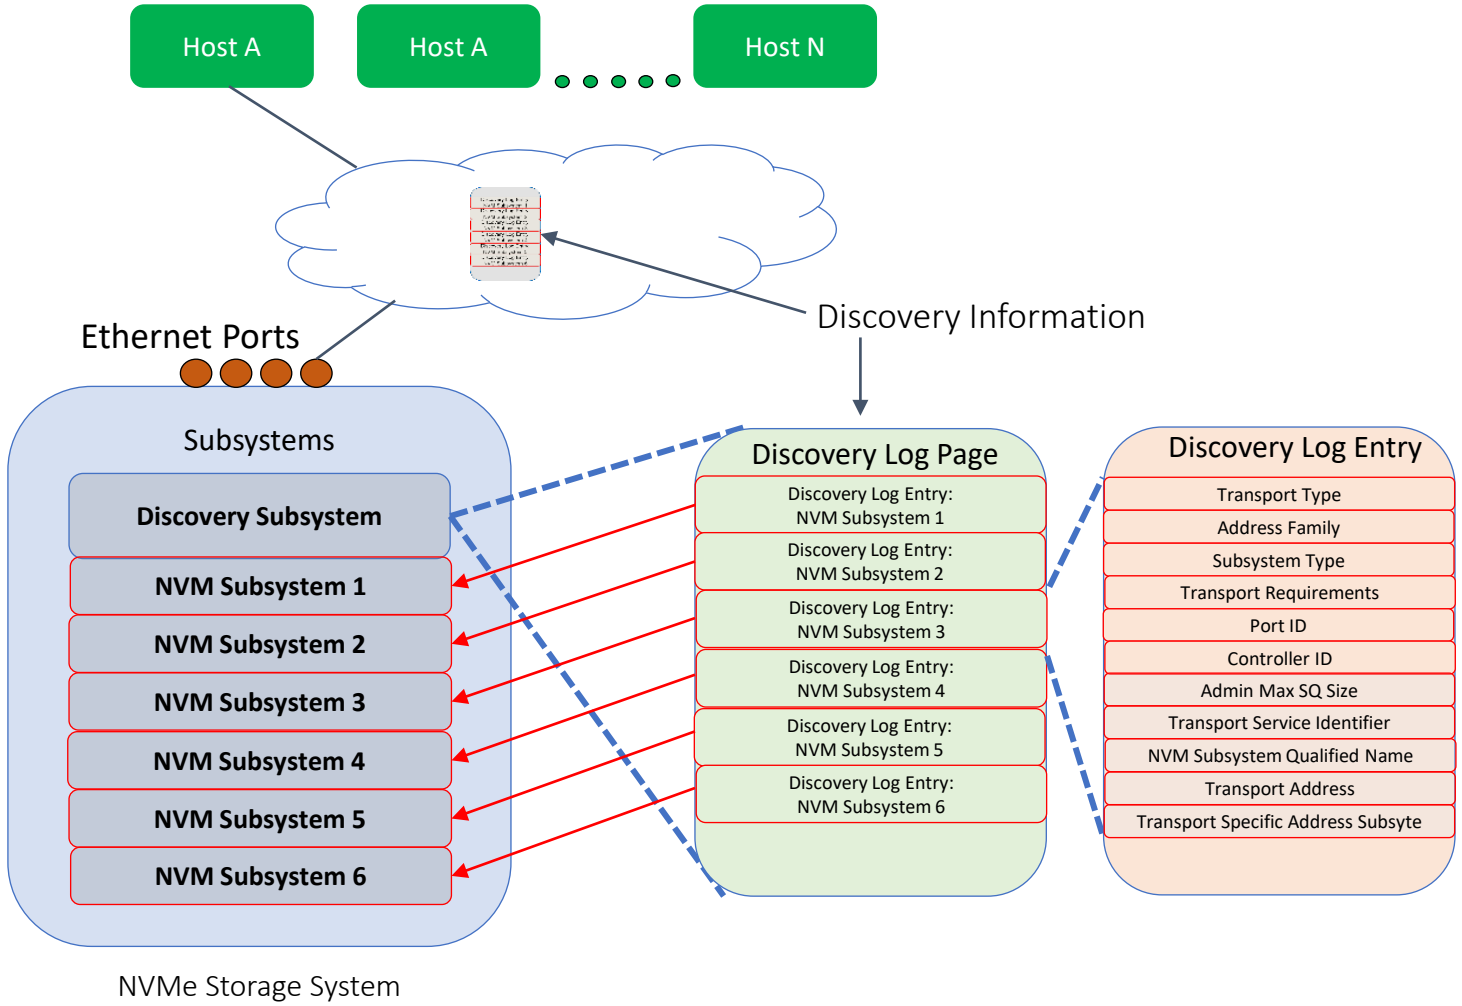

Centralized Discovery Controller

- A Centralized Discovery Controller aggregates discovery information from several NVMe™ storage systems to report discovery information for the full fabric.

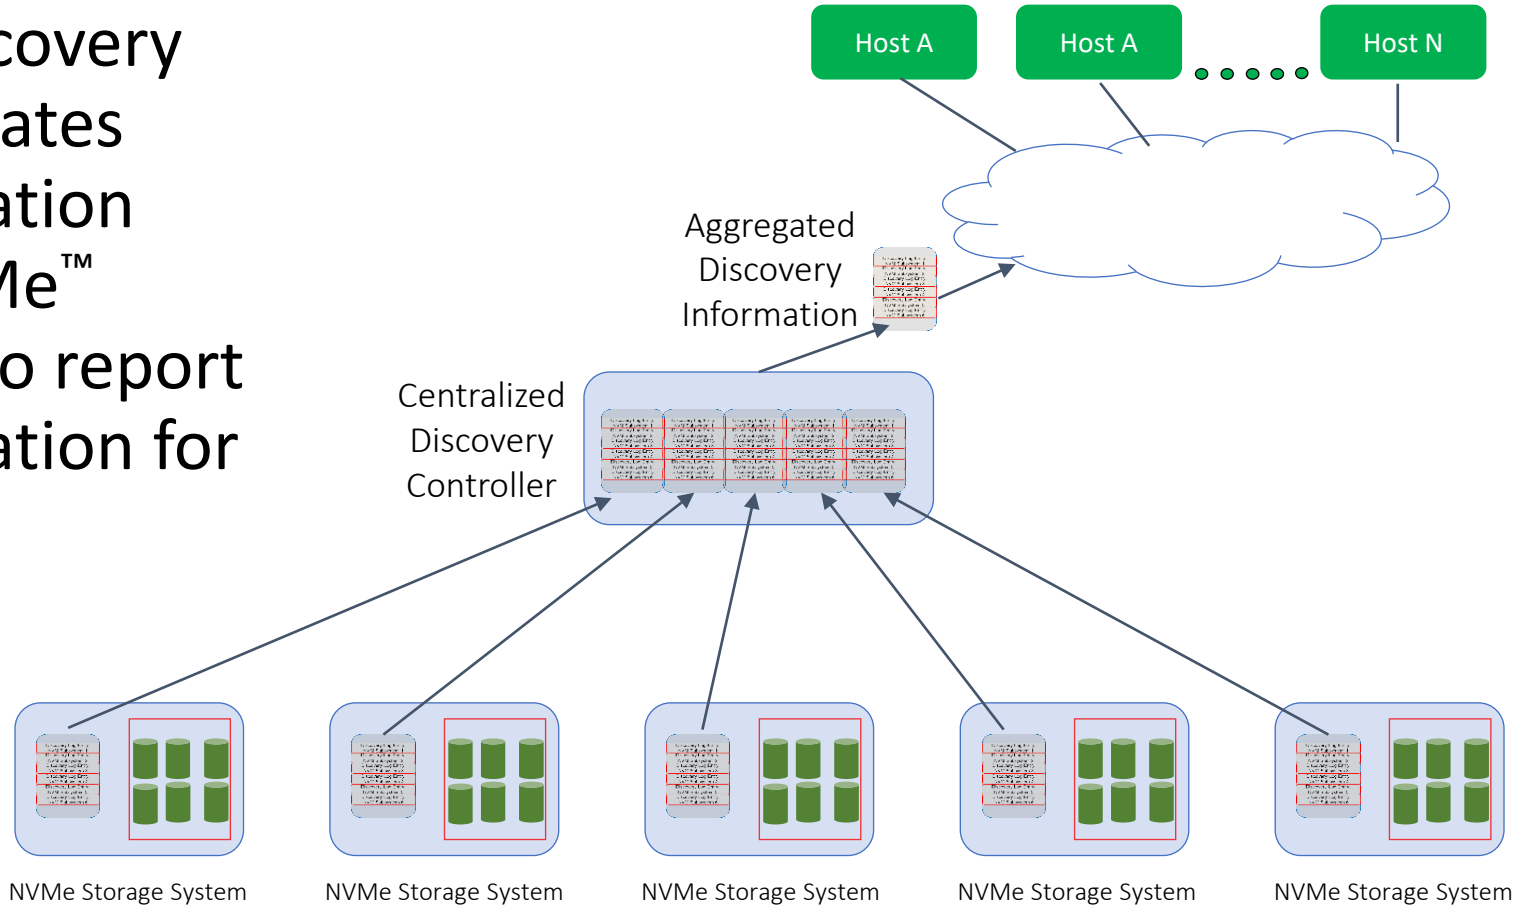

Swordfish Representation of Discovery Controllers

- Discovery Controllers require no configuration by the end user / client.
- So, we have created an extremely simplified, read-only model with information in two places:
- 1. Subsystems.
  - Subsystems have pointers to subsystems which contain discovery controllers
- 2. Discovery Controllers.
  - Discovery controllers have pointers to the subsystems they have discovered

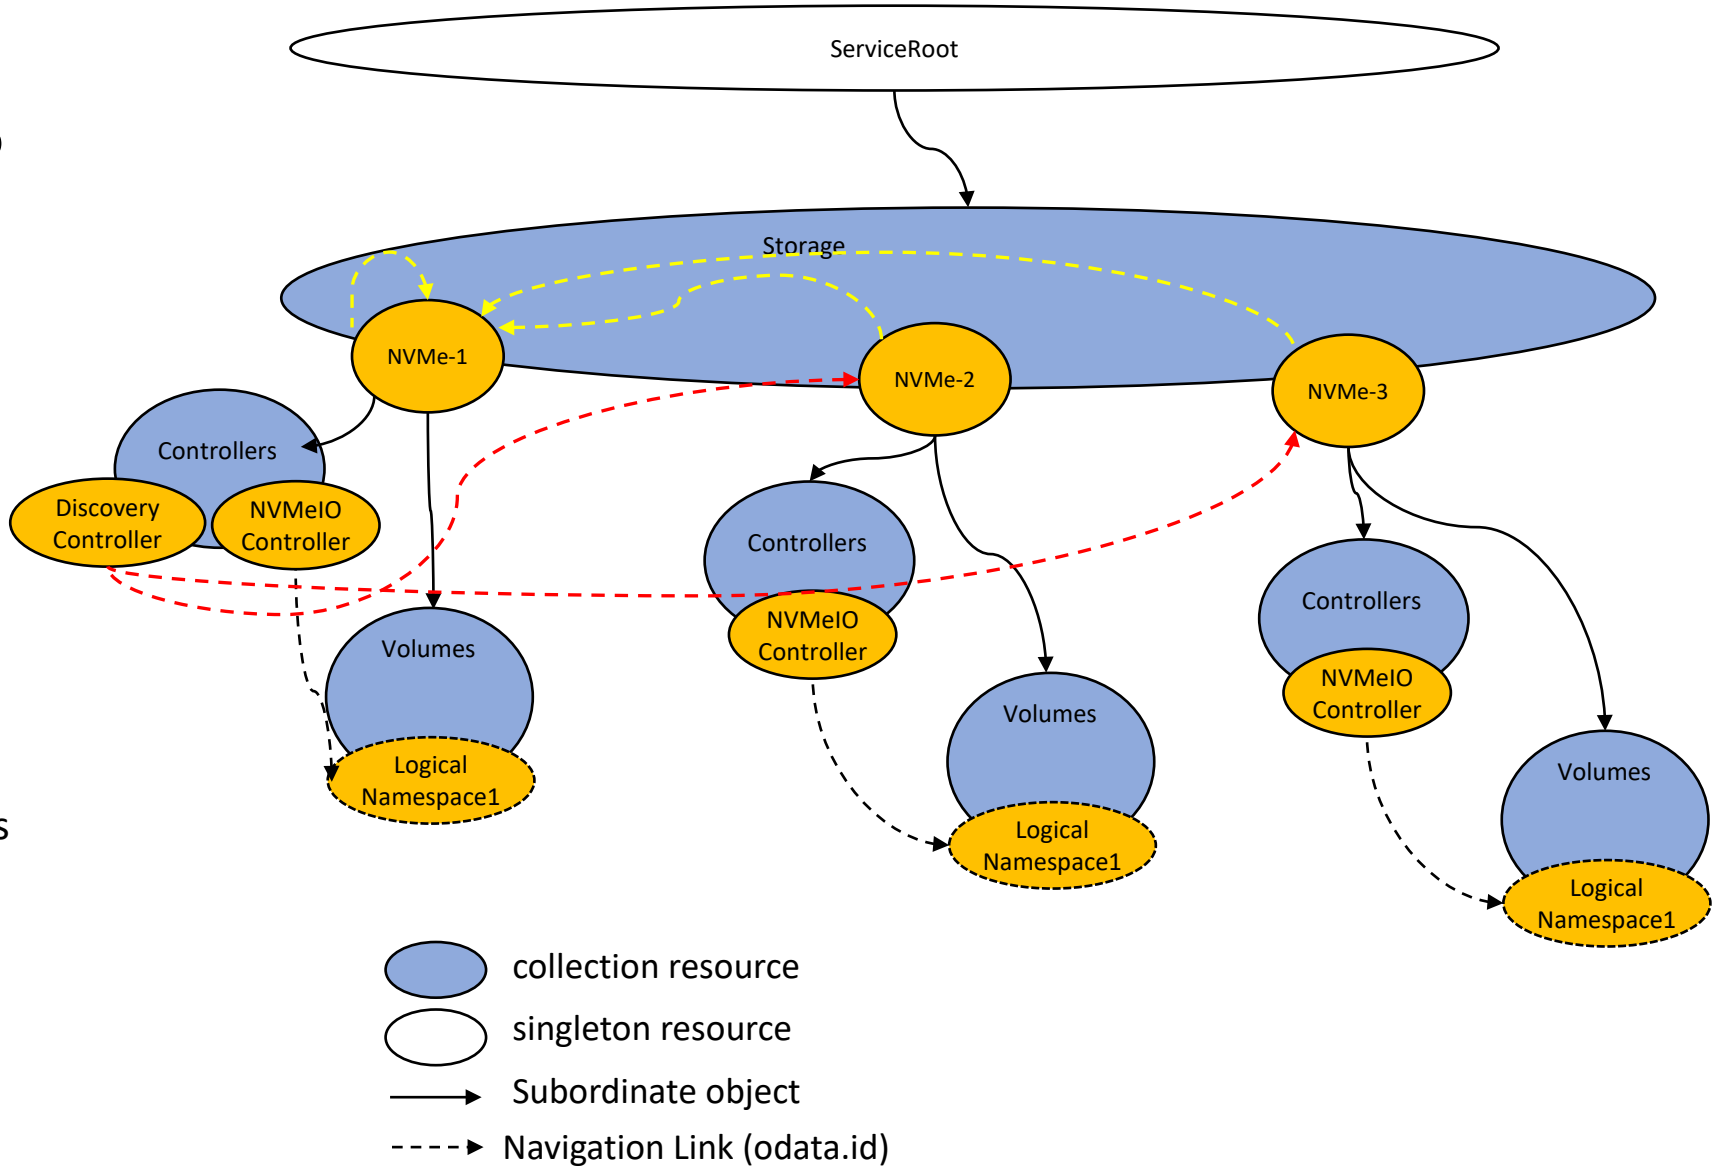

Mockup of Subsystem

Mockup of Discovery Controller

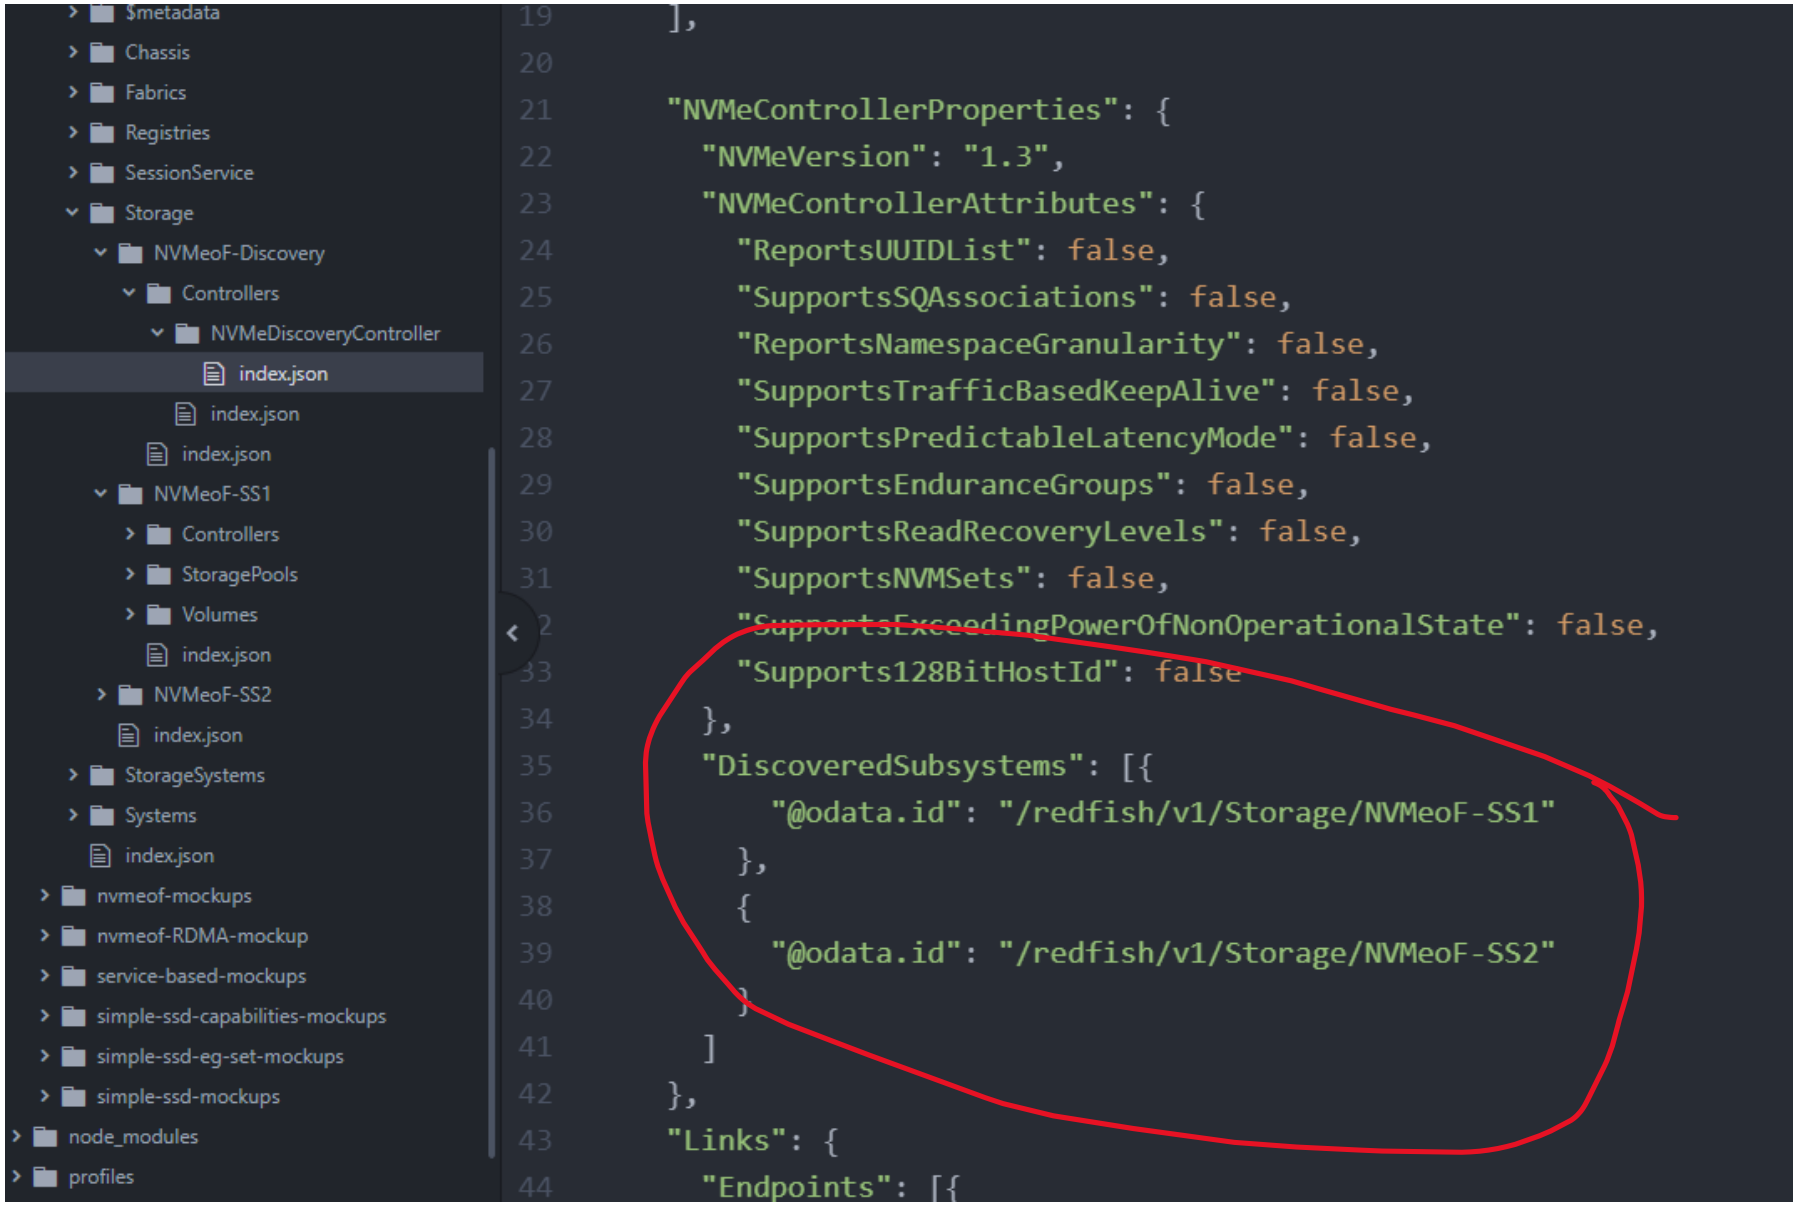

Where to find more info…

SNIA Swordfish™

- Swordfish Standards
  - Schemas, Specs, Mockups, User and Practical Guide`s, … https://www.snia.org/swordfish
- Swordfish Specification Forum
  - Ask and answer questions about Swordfish
  - http://swordfishforum.com/
- Scalable Storage Management (SSM) TWG
  - Technical Work Group that defines Swordfish
  - Influence the next generation of the Swordfish standard
  - Join SNIA & participate: https://www.snia.org/member_com/join-SNIA
- Join the SNIA Storage Management Initiative
  - Unifies the storage industry to develop and standardize interoperable storage management technologies
  - https://www.snia.org/forums/smi/about/join

DMTF Redfish™

- Redfish Standards
  - Specifications, whitepapers, guides,… https://www.dmtf.org/standards/redfish

Open Fabric Management Framework

- OFMF Working Group (OFMFWG)
  - Description & Links https://www.openfabrics.org/working-groups/
- OFMFWG mailing list subscription
  - https://lists.openfabrics.org/mailman/listinfo/ofmfwg
- Join the Open Fabrics Alliance
  - https://www.openfabrics.org/membership-how-to-join/

NVM Express

- Specifications https://nvmexpress.org/developers/
- Join: https://nvmexpress.org/join-nvme/