



2020 OFA Virtual Workshop

# **ENHANCING NVME AND NVME-OF CONFIGURATION AND MANAGEABILITY WITH SNIA SWORDFISH TO ENABLE SCALABLE INFRASTRUCTURES**

**Rajalaxmi Angadi (Intel)**

**Phil Cayton (Intel)**

**Richelle Ahlvers (Broadcom)**

**Intel Corporation/Broadcom Inc**

# SNIA

## SNIA-at-a-Glance



**185**  
industry leading  
organizations



**2,000**  
active contributing  
members



**50,000**  
IT end users & storage  
pros worldwide

Learn more: [snia.org/technical](https://snia.org/technical)



# DISCLAIMER

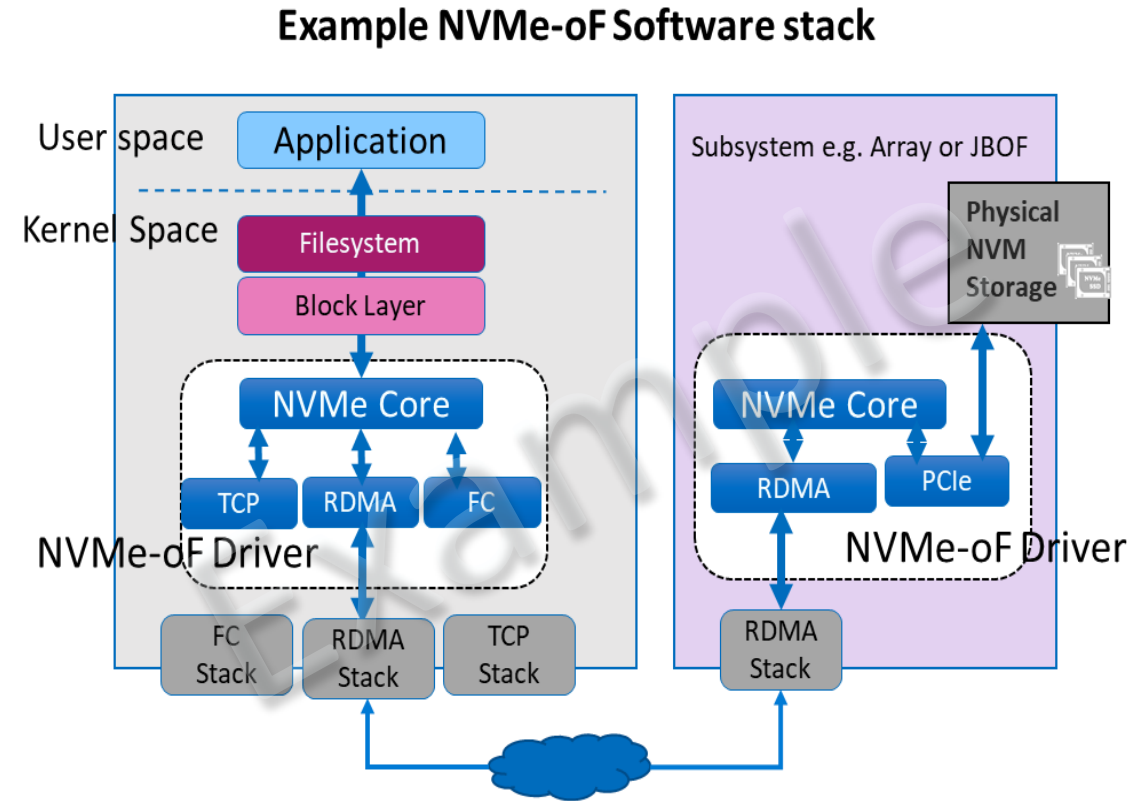
- **The information in this presentation:**
  - Represents a snapshot of the work in progress within SNIA
  - It is evolving
- **For additional information:**
  - SNIA website: [www.snia.org/swordfish](http://www.snia.org/swordfish)
  - NVMe Consortium website: [nvmexpress.org](http://nvmexpress.org)

# AGENDA

- **NVMe-oF Refresher**
- **Managing NVMe and NVMe-oF**
- **Introducing DMTF Redfish + SNIA Swordfish**
- **Managing NVMe and NVMe-oF in SNIA Swordfish**
- **Summary**

# NVM EXPRESS OVER FABRICS (NVME-OF) REFRESHER

- Extends the efficiency of NVMe over a variety of fabrics
- Builds on base NVMe architecture with thin encapsulation of base NVMe across fabrics
- Enables low-latency/high IOPS access to remote NVMe Storage
- Same architecture regardless of transport
- End-to-end mechanism to transfer NVMe commands and data structures
- Presents an abstraction for exporting NVM Subsystems over fabrics
- Restricted/Unrestricted NVMe-oF Subsystems access



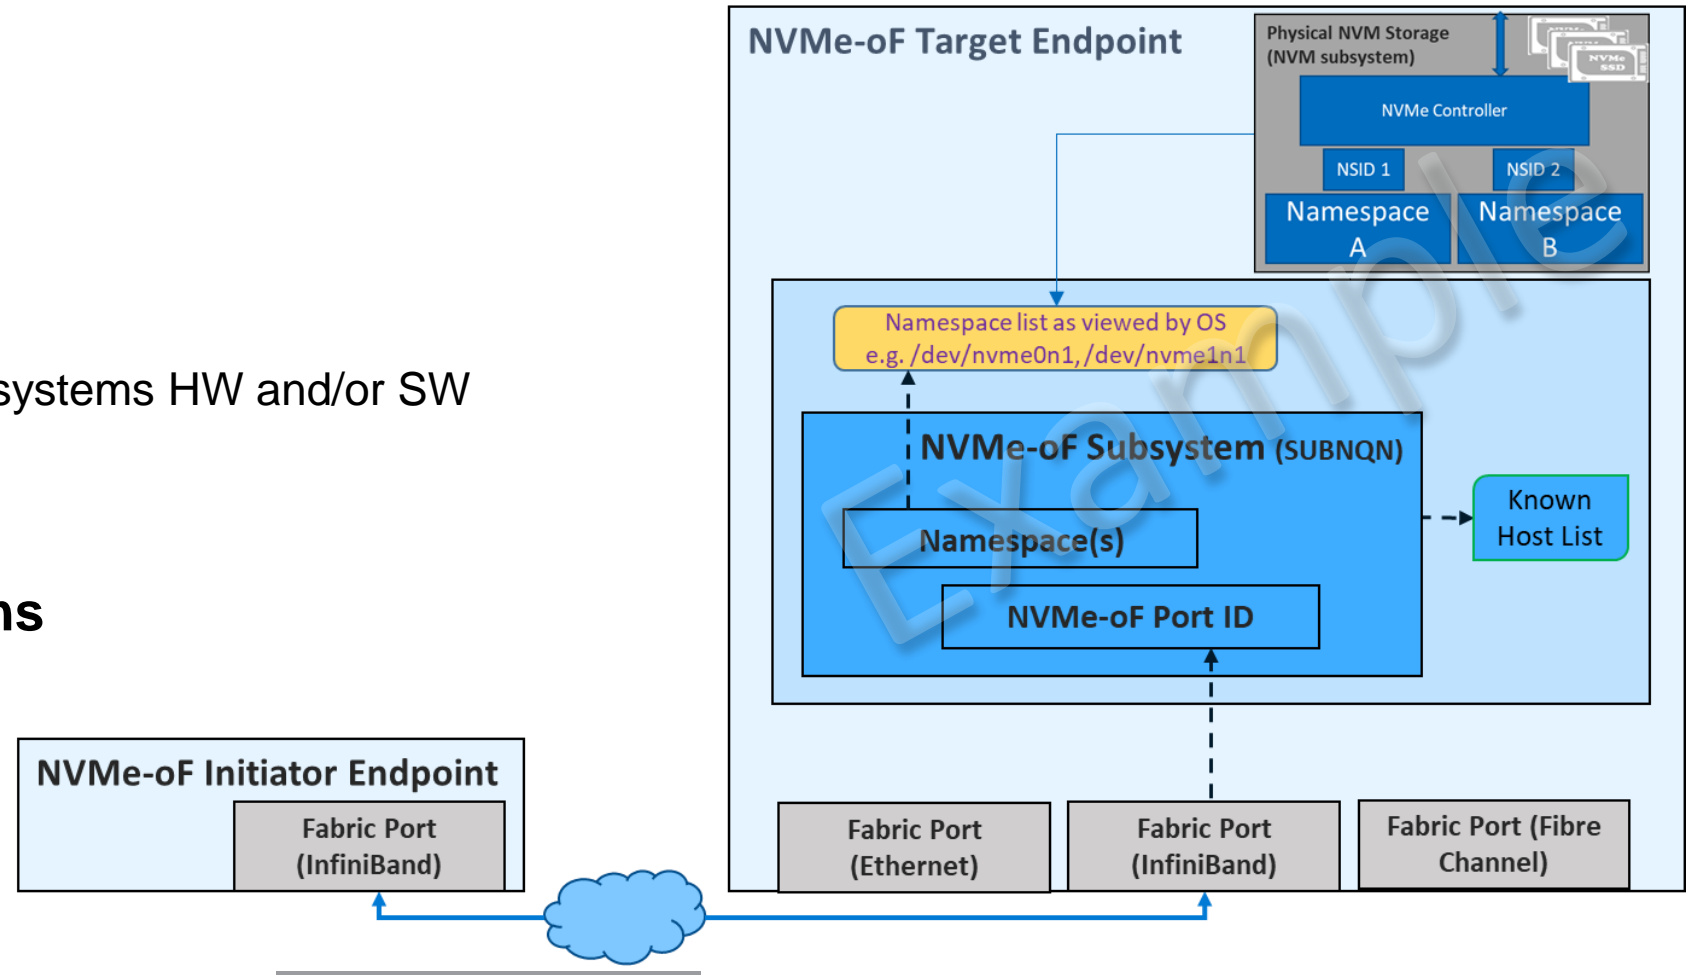
# MANAGING NVME/NVME-OF

## ■ Management points

- NVMe/NVMe-oF
  - Subsystem(s)
  - Controller(s)
  - Namespace(s)
- NVMe-oF
  - Initiator (host) Endpoint systems HW and/or SW
  - Fabrics

## ■ Management Operations

- Set configuration
- Get configuration
- Health status
- Reset



# INTRODUCING DMTF REDFISH & SNIA SWORDFISH

## ■ DMTF Redfish® - Infrastructure Management Standard

- IPMI Successor for extended Management Scope
- Focus: management of scale-out commodity servers
- Design: RESTful API, OData, HTTP operations(GET/ PUT/POST)
- Three main categories for server management
  - Systems – server, CPU, memory, devices, etc.
  - Managers – BMC, Enclosure Manager, etc.
  - Chassis – racks, enclosures, blades, etc.
- Expanding to cover data center infrastructure, fabrics, network management



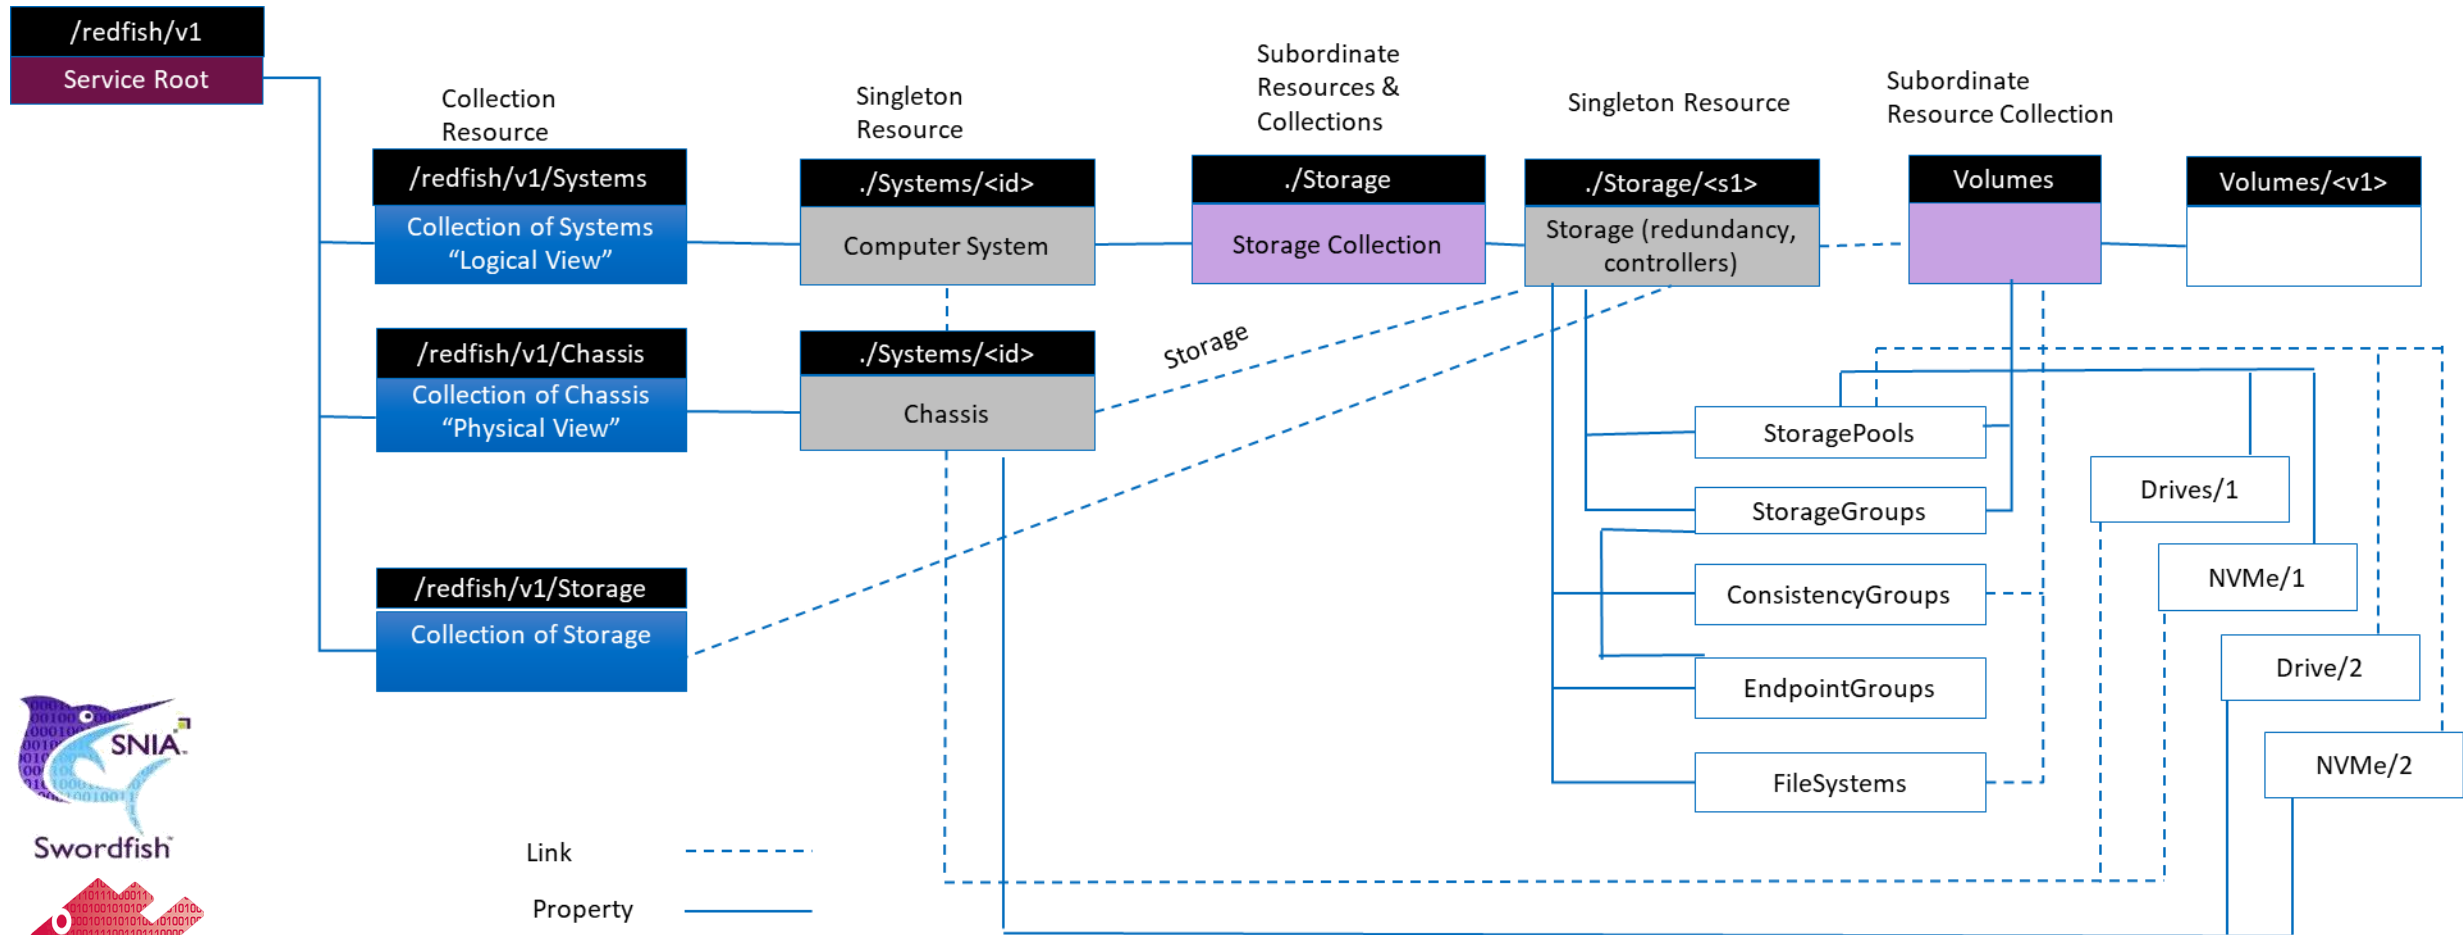
## ■ SNIA Swordfish® Storage Management and Ecosystem Standard

- Uses and Extends Redfish Schemas
- Focus: Storage Management
  - Logical Storage (Block, Object, File)



- Expanding to encompass NVMe and NVMe-oF

# SWORDFISH STORAGE




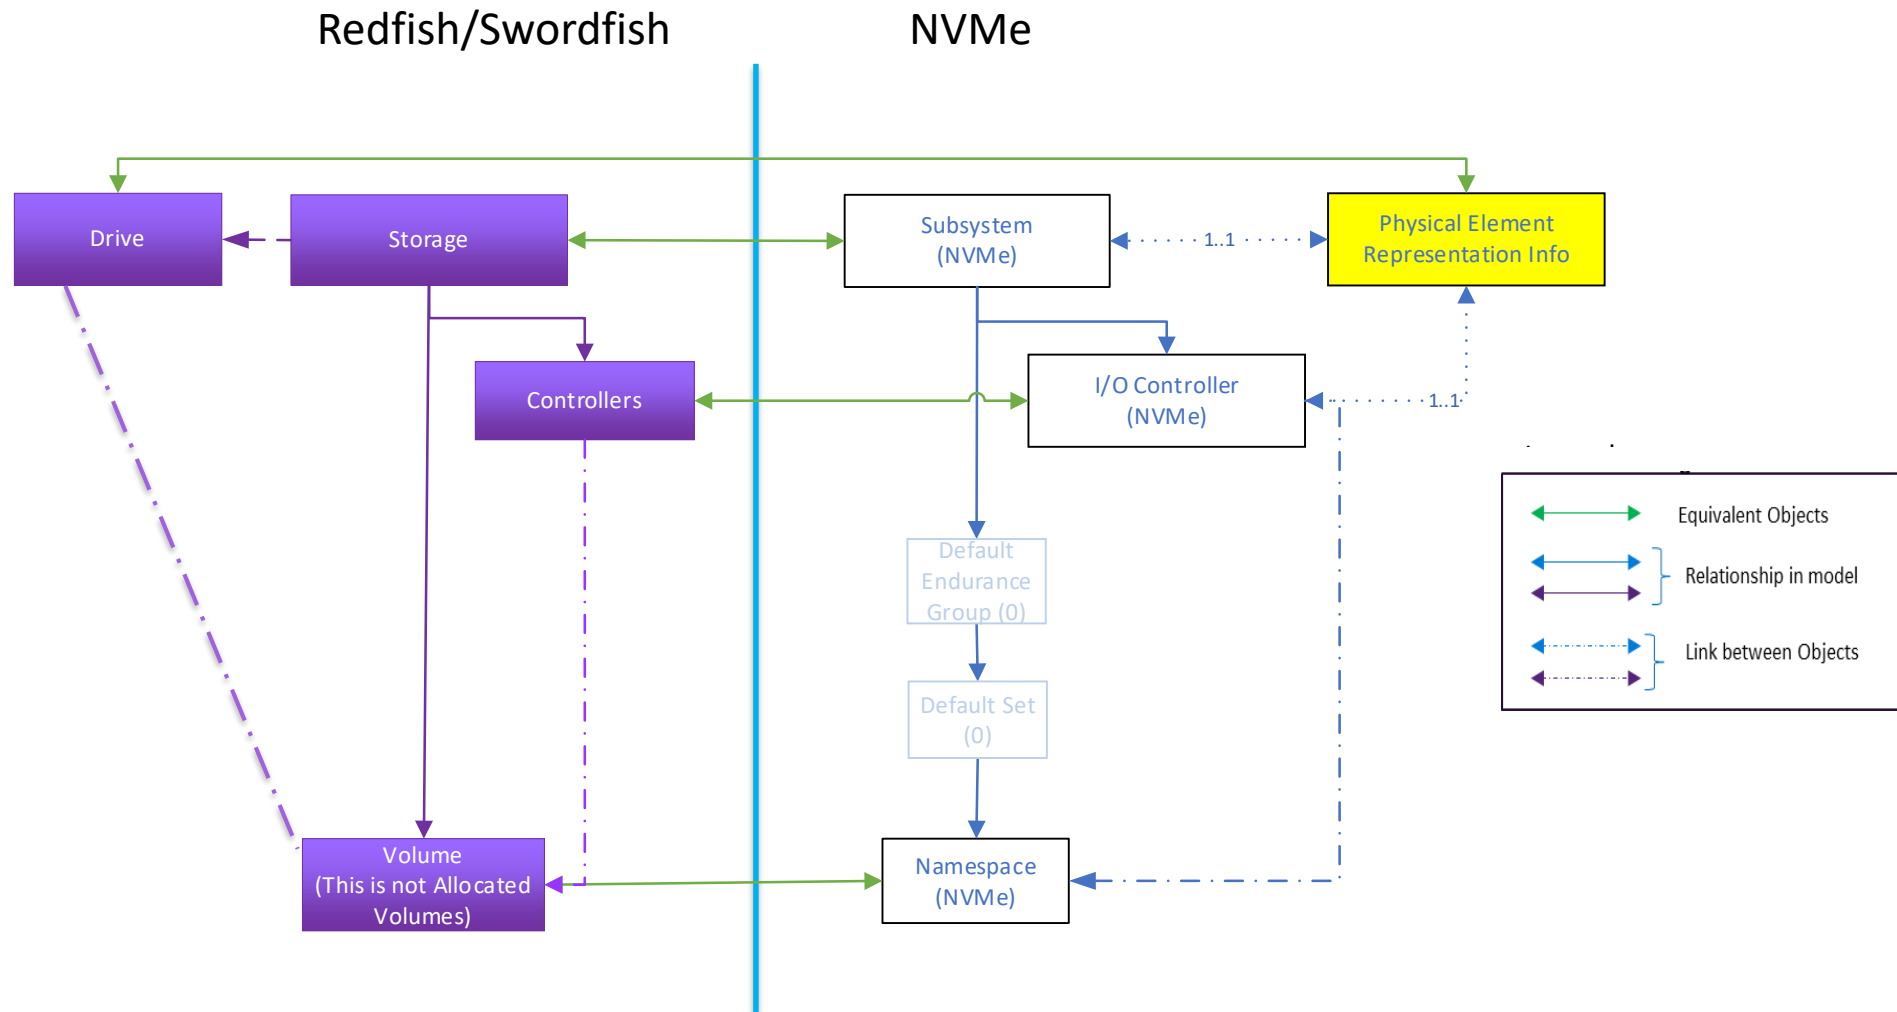
Swordfish



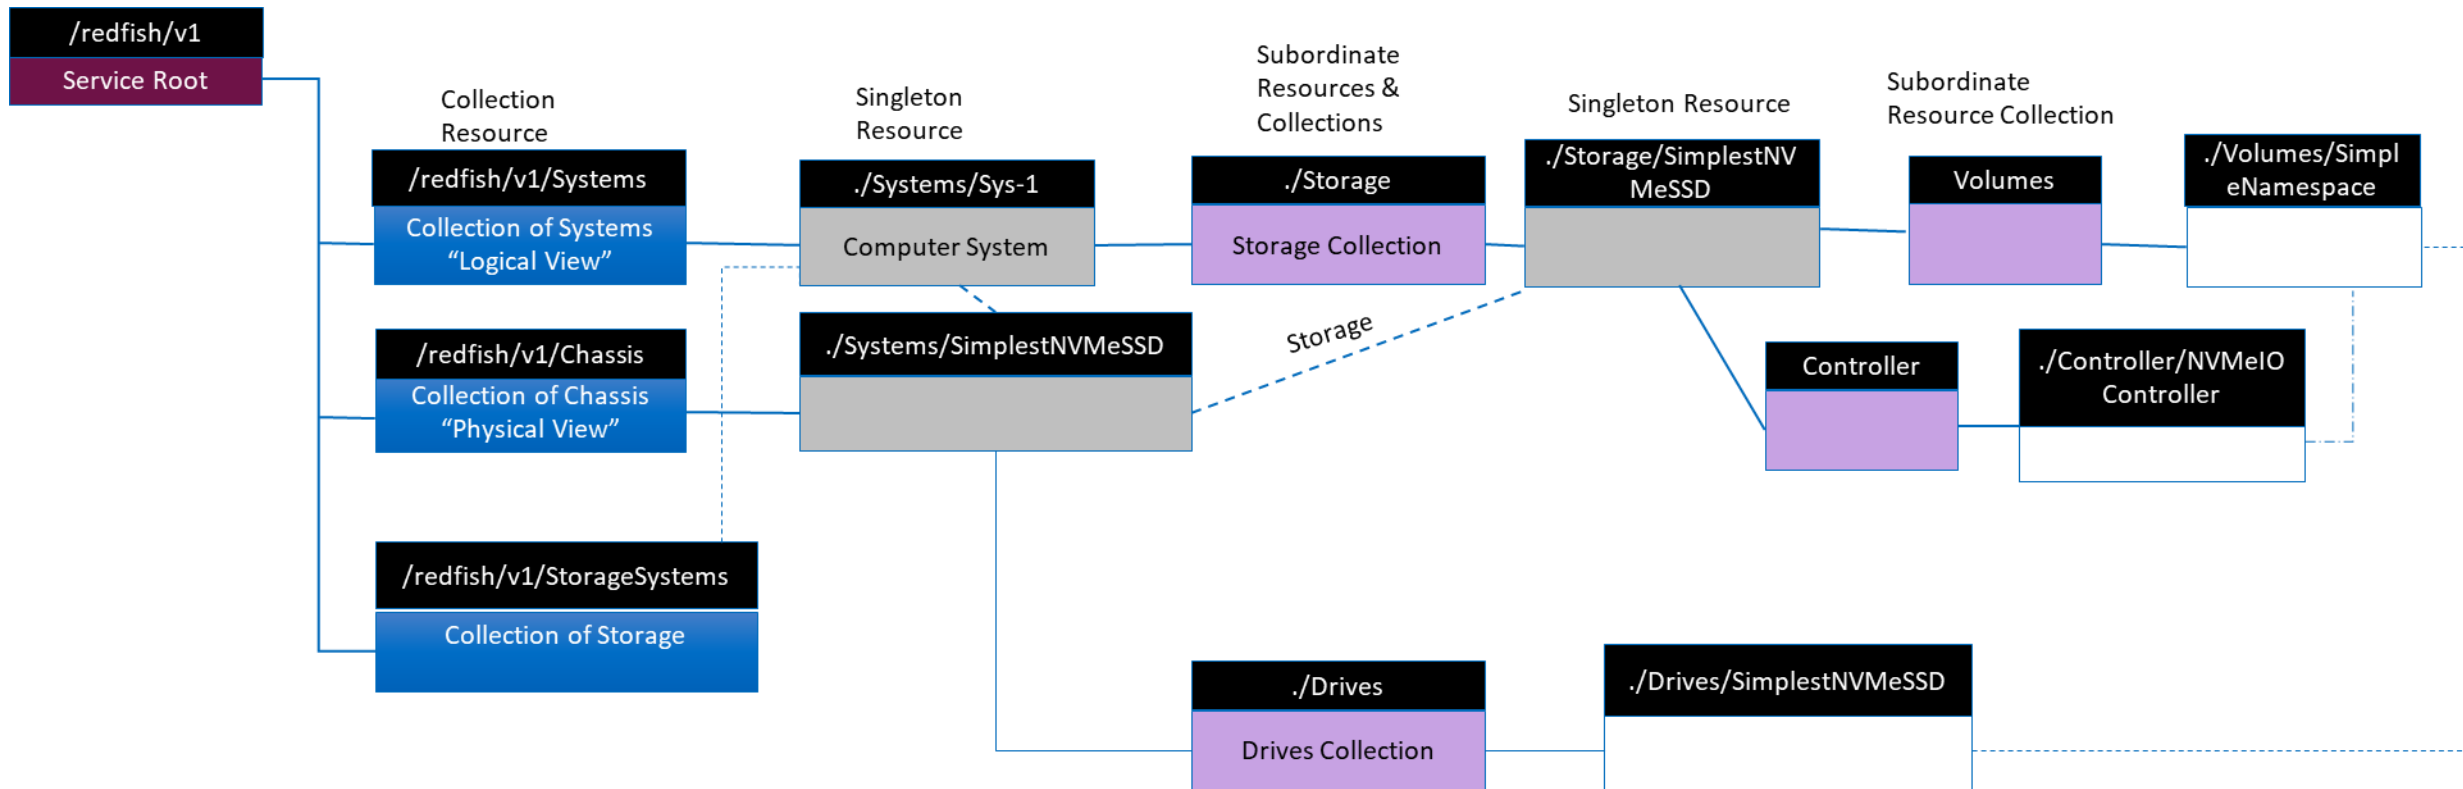
Redfish



# MANAGING NVME USING SWORDFISH



# NVME WITH SWORDFISH



Link  
Property

# SWORDFISH CONFIGURATIONS: NVME (SYSTEMS/SYS-1)

- SNIA SSM Technical Working Group
- Mockups at <http://swordfishmockups.com>

*Note: Mockups are representations of implementations*



```
{
  "@odata.type": "#ComputerSystem.v1_8_0.ComputerSystem",
  "Id": "Sys-1",
  "Name": "WebFrontEnd483",
  "SystemType": "Physical",
  "AssetTag": "Chicago-45Z-2381",
  "Manufacturer": "Contoso",
  "Model": "3500RX",
  "SKU": "8675309",
  "SerialNumber": "Sys-1",
  "PartNumber": "224071-J23",
  "Description": "Web Front End node",
  "UUID": "38947555-7742-3448-3784-823347823834",
  "HostName": "web483",
  "Status": {
    "State": "Enabled",
    "Health": "OK",
    "HealthRollup": "OK"
  },
  "HostingRoles": [
    "StorageServer"
  ],
  "Storage": {
    "@odata.id": "/redfish/v1/Systems/Sys-1/Storage"
  },
  "Links": {
    "Chassis": [
      {
        "@odata.id": "/redfish/v1/Chassis/SimplestNVMeSSD"
      }
    ]
  },
  "@odata.id": "/redfish/v1/Systems/Sys-1",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```

# SWORDFISH CONFIGURATIONS: NVME (STORAGE COLLECTION)

/redfish/v1/Systems/Sys-1/Storage

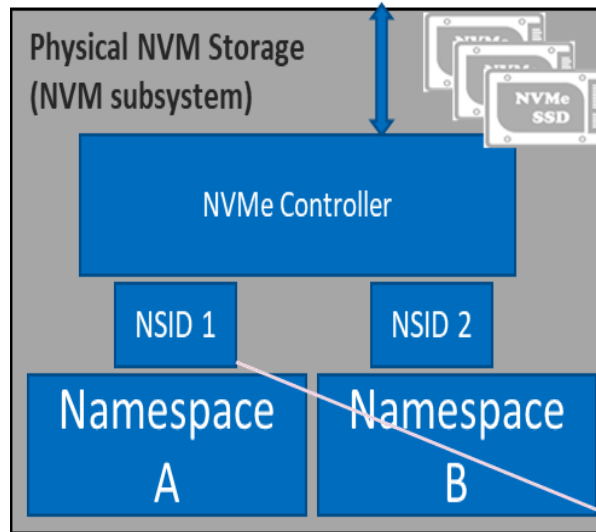
```
{
  "@odata.type": "#StorageCollection.StorageCollection",
  "Name": "Storage Collection",
  "Members@odata.count": 1,
  "Members": [
    {
      "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD"
    }
  ],
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```

/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD

```
{
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved.",
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD",
  "@odata.type": "#Storage.v1_9_0.Storage",
  "Id": "1",
  "Name": "NVMe Simplest Configuration",
  "Description": "Mockup of simplest NVMe simple config with 1 Subsystem, 1 I/O Controller and 1 Namespace.",
  "Status": {
    "State": "Enabled",
    "Health": "OK",
    "HealthRollup": "OK"
  },
  "Identifiers": [{
    "DurableNameFormat": "NQN",
    "DurableName": "nqn.2014-08.org.nvmexpress:uuid:6c5fe566-10e6-4fb6-aad4-8b4159f50245"
  }],
  "Controllers": {
    "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers"
  },
  "Volumes": {
    "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes"
  }
}
```



# SWORDFISH CONFIGURATIONS: NVME (VOLUME/NAMESPACE)

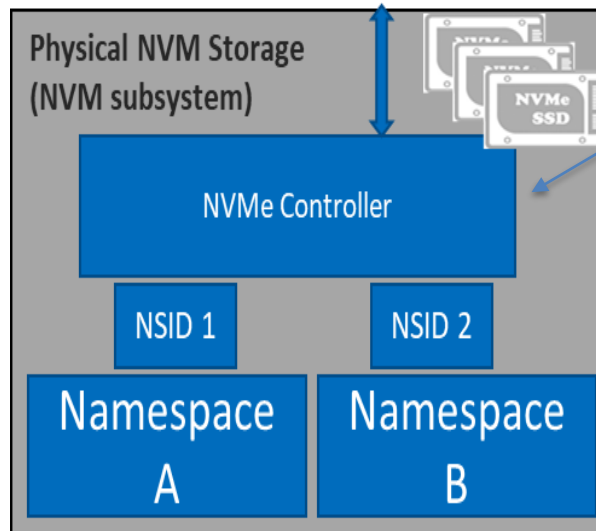


```
{
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved.",
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimpleNVMeSSD/Volumes/SimpleNamespace",
  "@odata.type": "#Volume.v1_5_0.Volume",
  "Id": "1",
  "Name": "Namespace 1",
  "LogicalUnitNumber": 1,
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "Capacity": {
    "Data": {
      "ConsumedBytes": 0,
      "AllocatedBytes": 10737418240
    },
    "Metadata": {
      "AllocatedBytes": 536870912
    }
  },
  "NVMeNamespaceProperties": {
    "NamespaceFeatures": {
      "SupportsThinProvisioning": false,
      "SupportsAtomicTransactionSize": false,
      "SupportsDeallocatedOrUnwrittenLBAError": false,
      "SupportsNGUIDFieldValue": false,
      "SupportsIOPerformanceHints": false
    },
    "NumberLBAFormats": 0,
    "FormattedLBASize": "LBAFormat0Support",
    "MetadataTransferredAtEndOfDataLBA": false,
    "NVMeVersion": "1.4"
  }
}
```



GET /redfish/v1/Systems/Sys-1/Storage/SimpleNVMeSSD/Volumes/SimpleNamespace/1 HTTP/1.1

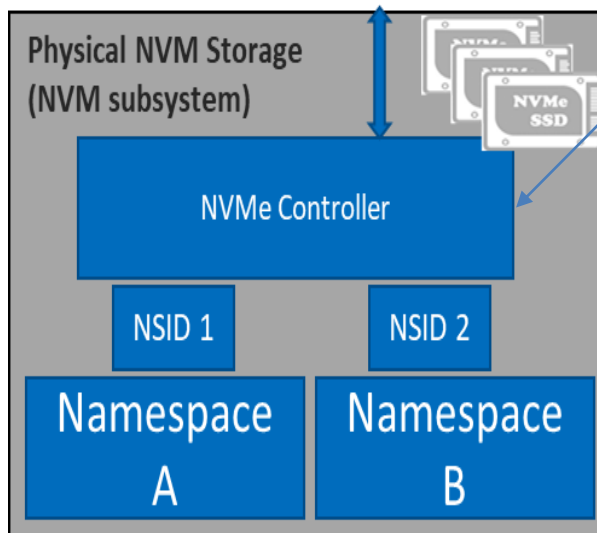
# SWORDFISH CONFIGURATIONS: NVME (CONTROLLER)



/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers/NVMeIOController


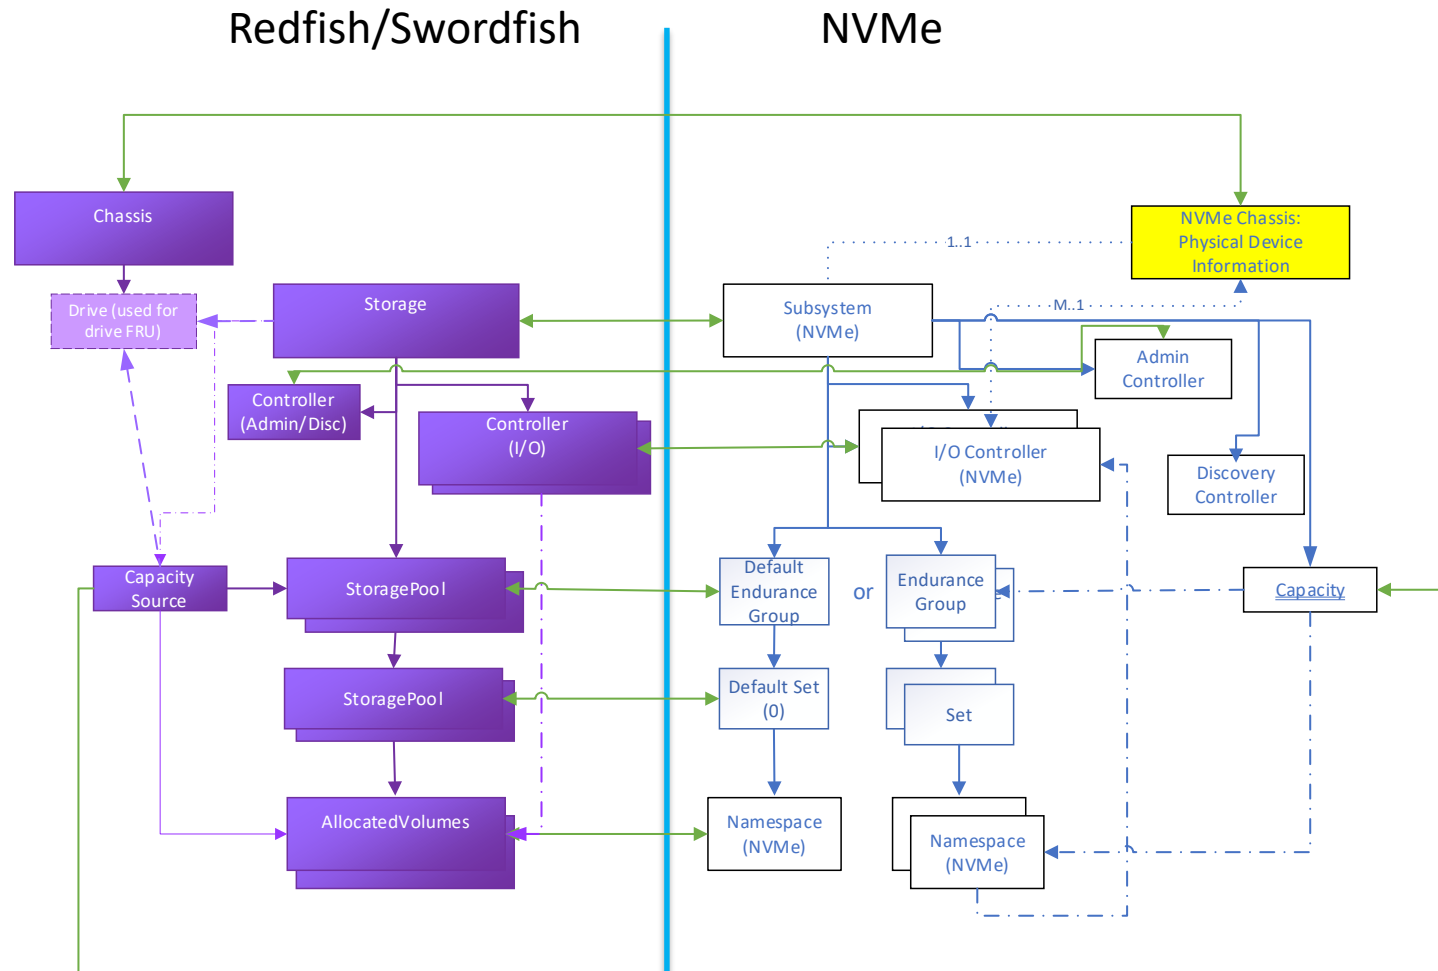
```
{
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved.",
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers/NVMeIOController",
  "@odata.type": "#StorageController.v1_0_0.StorageController",
  "Id": "1",
  "Name": "NVMe I/O Controller",
  "Description": "Single NVMe I/O Controller presented to host.",
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "Id": "NVMeIOController",
  "Manufacturer": "Best NVMe Vendor",
  "Model": "Simple NVMe Device",
  "SerialNumber": "NVME123456",
  "PartNumber": "NVM44",
  "FirmwareVersion": "1.0.0",
  "SupportedControllerProtocols": [
    "PCIe"
  ],
}
```

# SWORDFISH CONFIGURATIONS: NVME (CONTROLLER...)



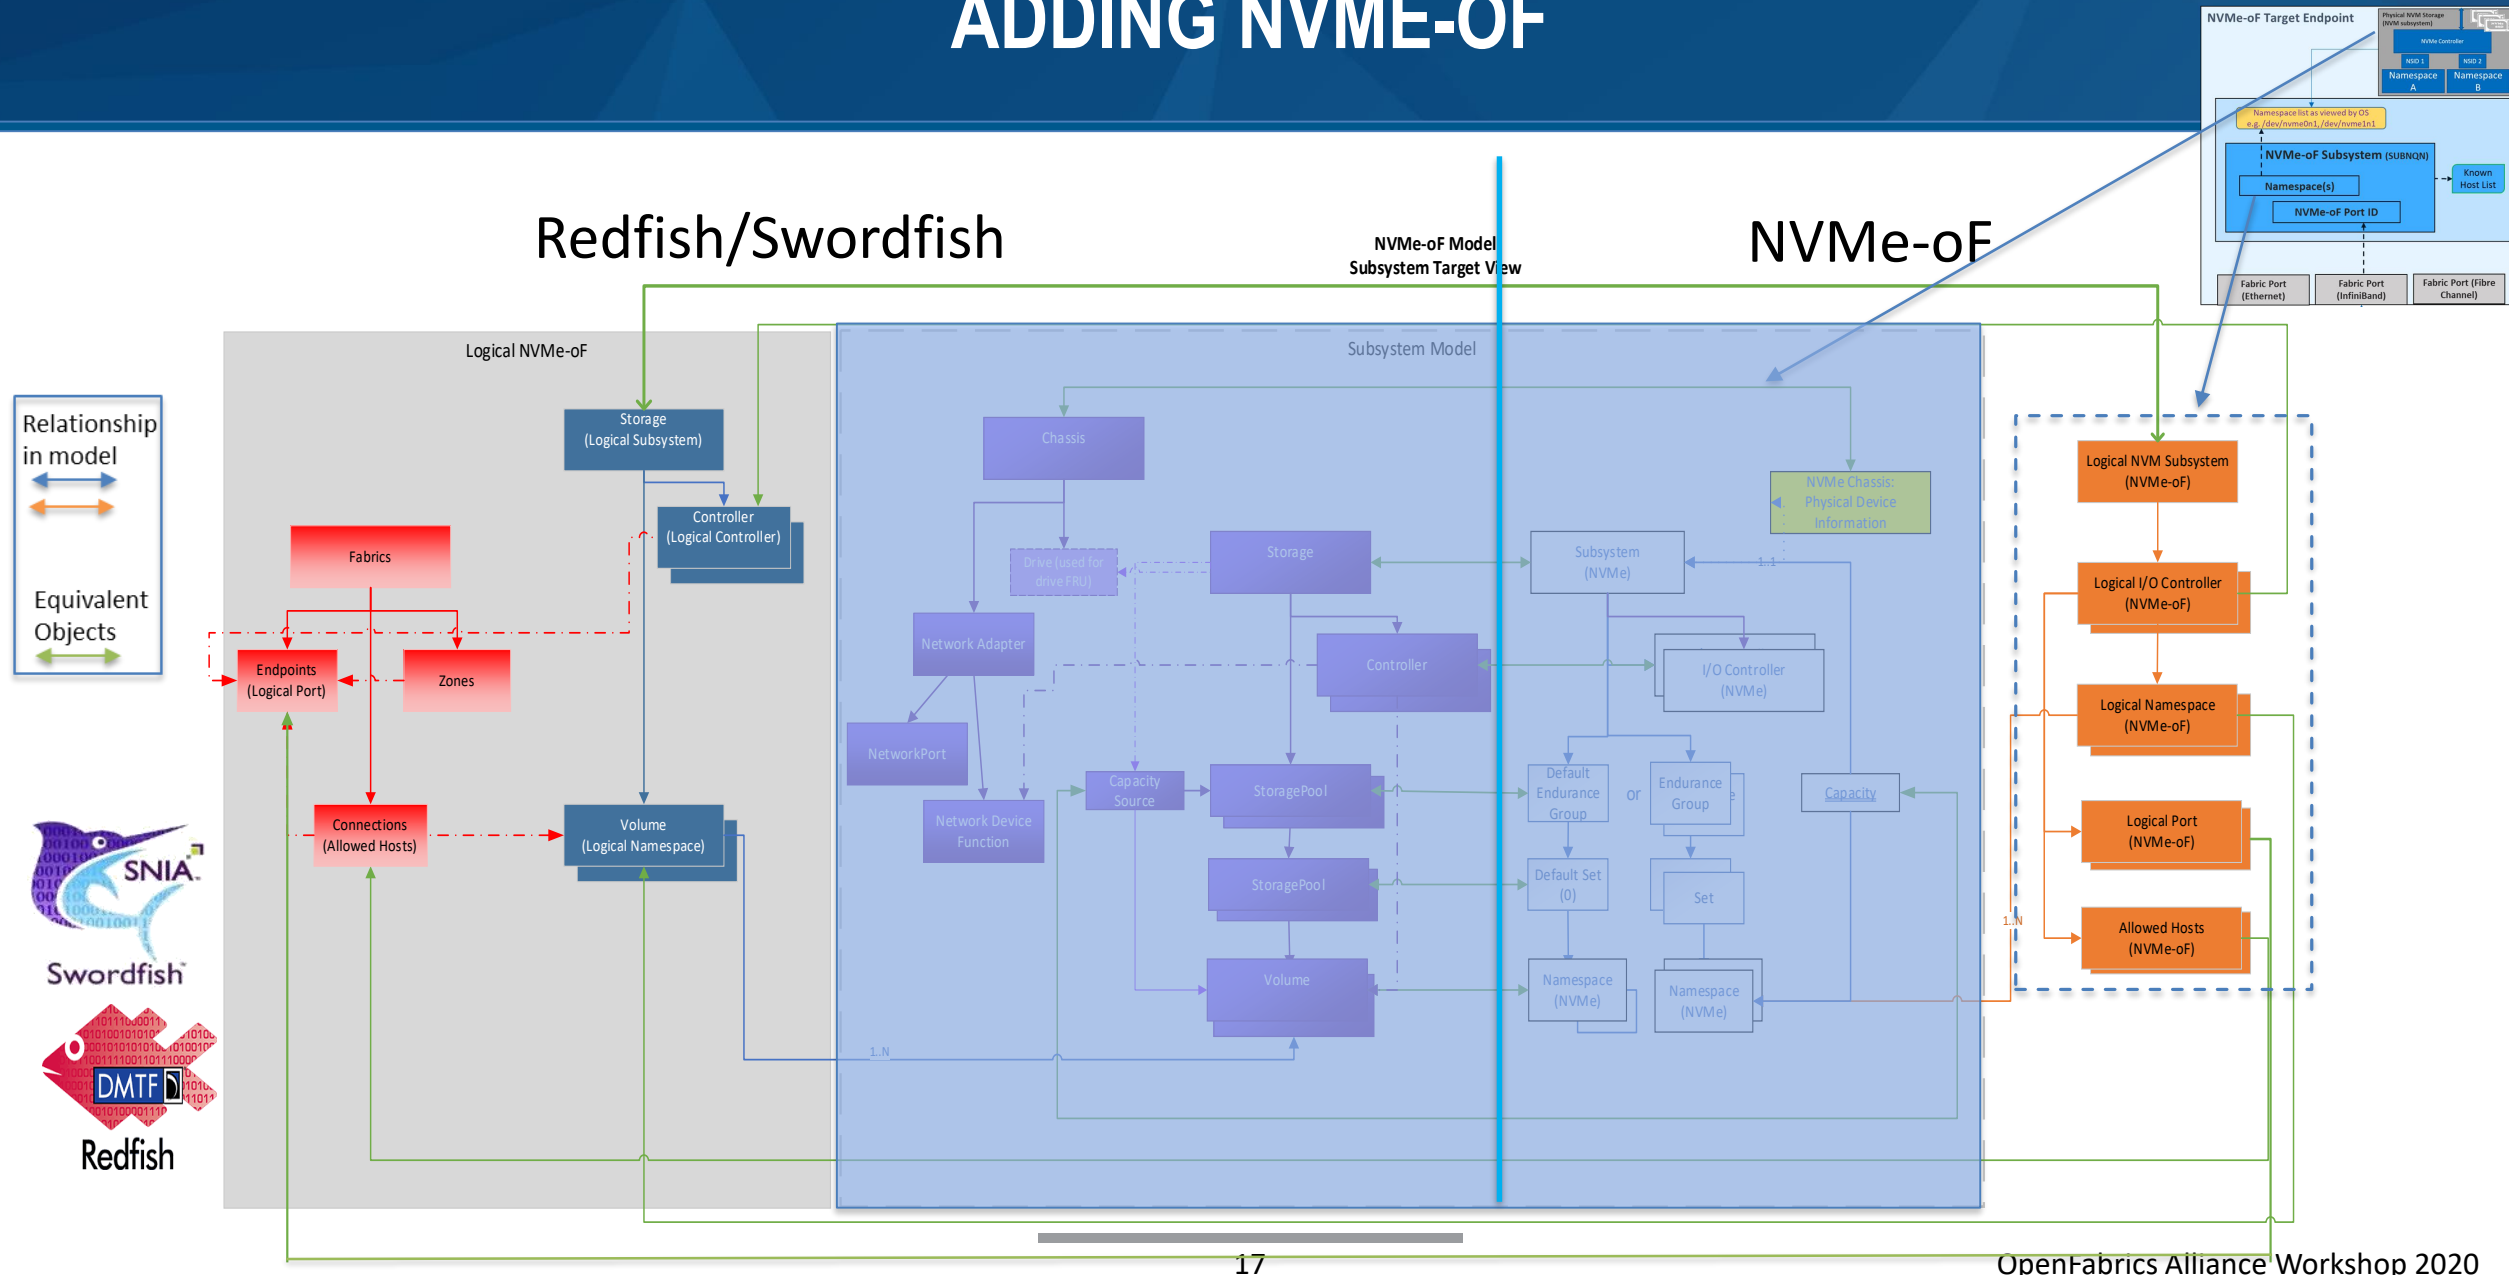
```
"NVMeControllerProperties": {  
  "NVMeVersion": "1.3",  
  "NVMeControllerAttributes": {  
    "ReportsUUIDList": false,  
    "SupportsSQAssociations": false,  
    "ReportsNamespaceGranularity": false,  
    "SupportsTrafficBasedKeepAlive": false,  
    "SupportsPredictableLatencyMode": false,  
    "SupportsEnduranceGroups": false,  
    "SupportsReadRecoveryLevels": false,  
    "SupportsNVMSets": false,  
    "SupportsExceedingPowerOfNonOperationalState": false,  
    "Supports128BitHostId": false  
  }  
},  
"Links": {  
  "AttachedVolumes": [{  
    "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace"  
  }]  
}
```

# NVME WITH ENDURANCE GROUP AND SET

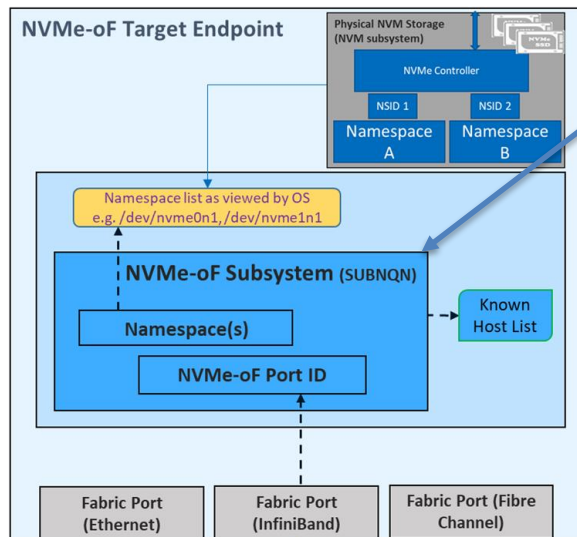




# ADDING NVME-OF

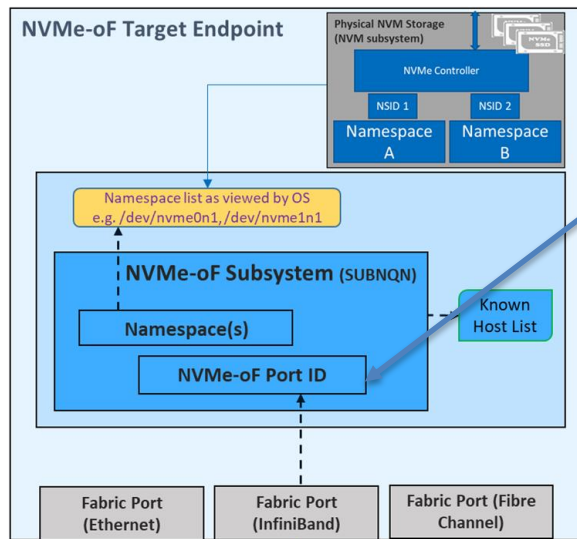


# SWORDFISH CONFIGURATIONS: NVME-OF (SUBSYSTEM)



```
{
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved.",
  "@odata.id": "/redfish/v1/Storage/NVMe-oF-Subsystem",
  "@odata.type": "#Storage.v1_9_0.Storage",
  "Id": "1",
  "Name": "NVMe-oF Logical NVM Fabric System",
  "Description": "Mockup of NVMe-oF Logical NVM Fabric System with 1 Logical Subsystem,
    Logical I/O Controller and 1 Logical port and 1 allowed host.",
  "Status": {
    "State": "Enabled",
    "Health": "OK",
    "HealthRollup": "OK"
  },
  "Identifiers": [{
    "DurableNameFormat": "NQN",
    "DurableName": "nqn.2014-08.org.nvmexpress:uuid:6c5fe566-10e6-4fb6-aad4-8b4159f50245"
  }],
  "Controllers": {
    "@odata.id": "/redfish/v1/Storage/NVMe-oF-Subsystem/Controllers"
  },
  "Volumes": {
    "@odata.id": "/redfish/v1/Storage/NVMe-oF-Subsystem/Volumes/LogicalNamespace1"
  }
}
```

# SWORDFISH CONFIGURATIONS: NVME-OF (ENDPOINT)



```
{
  "@odata.type": "#Endpoint.v1_4_0.Endpoint",
  "Id": "1",
  "Name": "NVMeEndpoint",
  "Description": "Endpoint connected Logical Namespace (NVMe-oF)",
  "EndpointProtocol": "NVMeOverFabrics",

  "ConnectedEntities": [{
    "EntityType": "Volume",
    "EntityRole": "Target",
    "Identifiers": [{
      "DurableNameFormat": "NGUID",
      "DurableName": "FDECBA9876543210h"
    }]
  }],

  "IPTransportDetails": [{
    "TransportProtocol": "RDMA",
    "IPv4Address": {
      "Address": "192.168.155.22"
    },
    "Port": 4420
  }],

  "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Endpoints/NVMeEndpoint",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```

# SUMMARY AND WRAP-UP

- **Schema modeling is a work in progress**
  - Mockups – Simple NVMe, NVMe with Endurance Group and Set, JBOF, Fabrics attached array
  - Advanced NVMe features like NVMe Power profiles properties will be added in the future
  - Redfish and swordfish initial draft incorporating NVMe and NVMe-oF management will be release in June 2020 for public review
- **Review the currently released [Swordfish Mockups](#)**
  - Ensure the schema is defined sufficiently to represent your desired implementation
  - Mockup your use case & submit it to the Swordfish forum

**Join SNIA and the Scalable Storage Management TWG & help define Schema**

<https://members.snia.org/wg/ssmtwg/dashboard>



2020 OFA Virtual Workshop

**THANK YOU**

Rajalaxmi Angadi

Intel Corporation