




*BY Developers FOR Developers*

**Storage Developer Conference**  
**September 22-23, 2020**

# **NVMe and NVMe-oF Configuration and Manageability with Swordfish and Redfish**

**Rajalaxmi Angadi - Intel Corporation**  
**Krishnakumar Gowrawaram – Cisco Systems**



# Agenda

- Resources in NVMe/NVMe-oF
- NVMe Resource Management
- Management with NVMe-MI
- NVMe Management with Swordfish
- Swordfish Models
- Summary and Wrap up

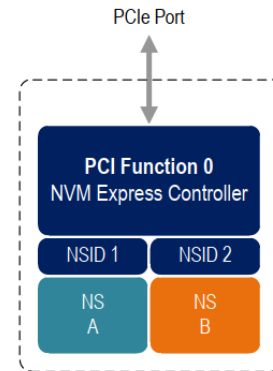
# Disclaimer

- The information in this presentation represents a snapshot of work in progress within SNIA
- This information is subject to change without notice.
- For additional information, see the SNIA website: [www.snia.org/swordfish](http://www.snia.org/swordfish)




# Resources to manage in NVMe/NVMe-oF

- Namespace
- Controllers
- Subsystem
- Endurance Groups
- Sets
- Domains



Example: NVM Subsystem with 1 controller, 2 Namespaces



Example: Endurance group and sets

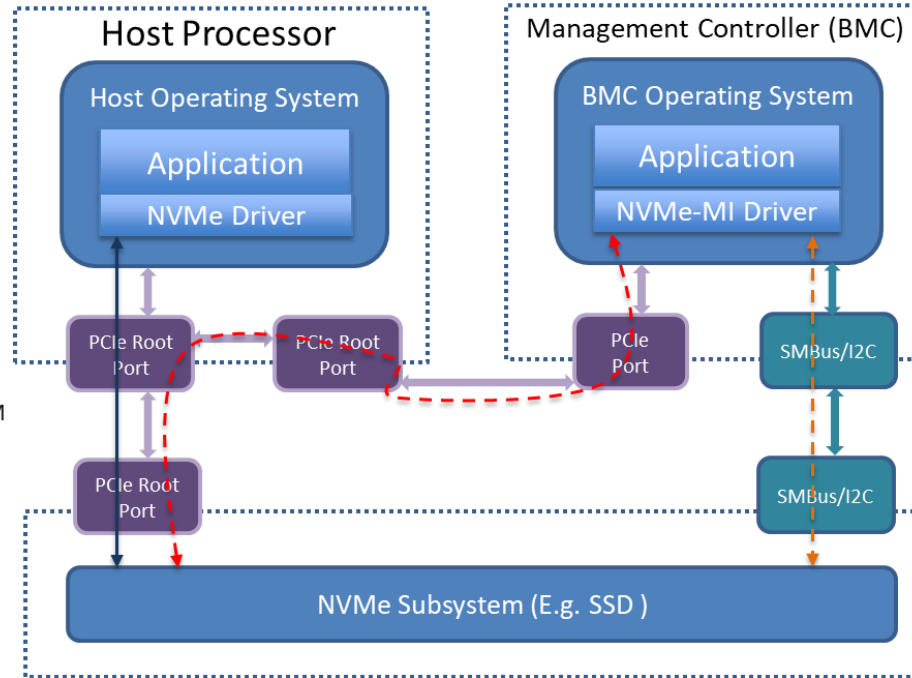
# NVMe Resource Management

- **Out-of-Band Management** – Management of hardware resources and components independent of host operating system
  - NVMe™ Out-of-Band Management Interfaces
    - SMBus/I2C
    - PCIe Vendor Defined Messages (VDM)
- **In-Band Management** – allows application to tunnel NVMe-MI commands through NVMe™ driver with NVMe Admin commands NVMe-MI Send/Receive commands. This is through the host operating system

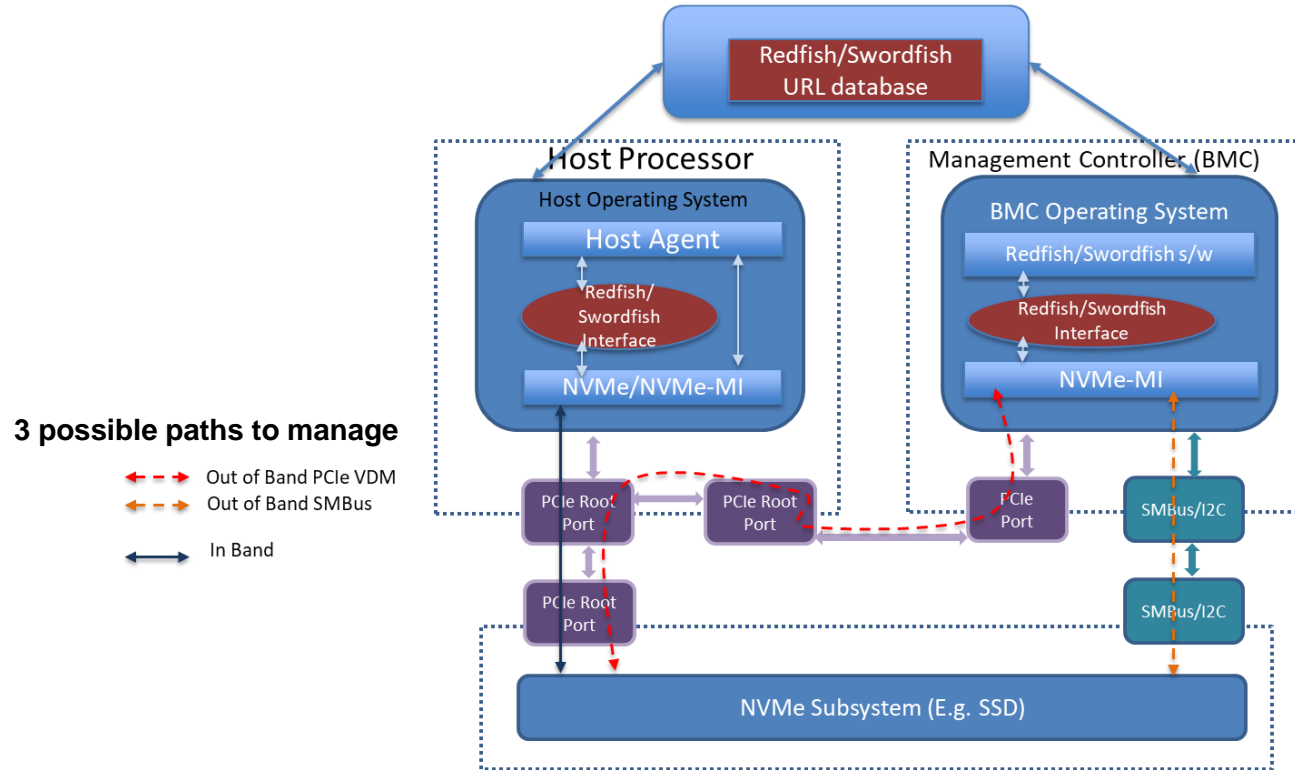
# NVMe Management with NVMe-MI

## 3 possible paths to manage

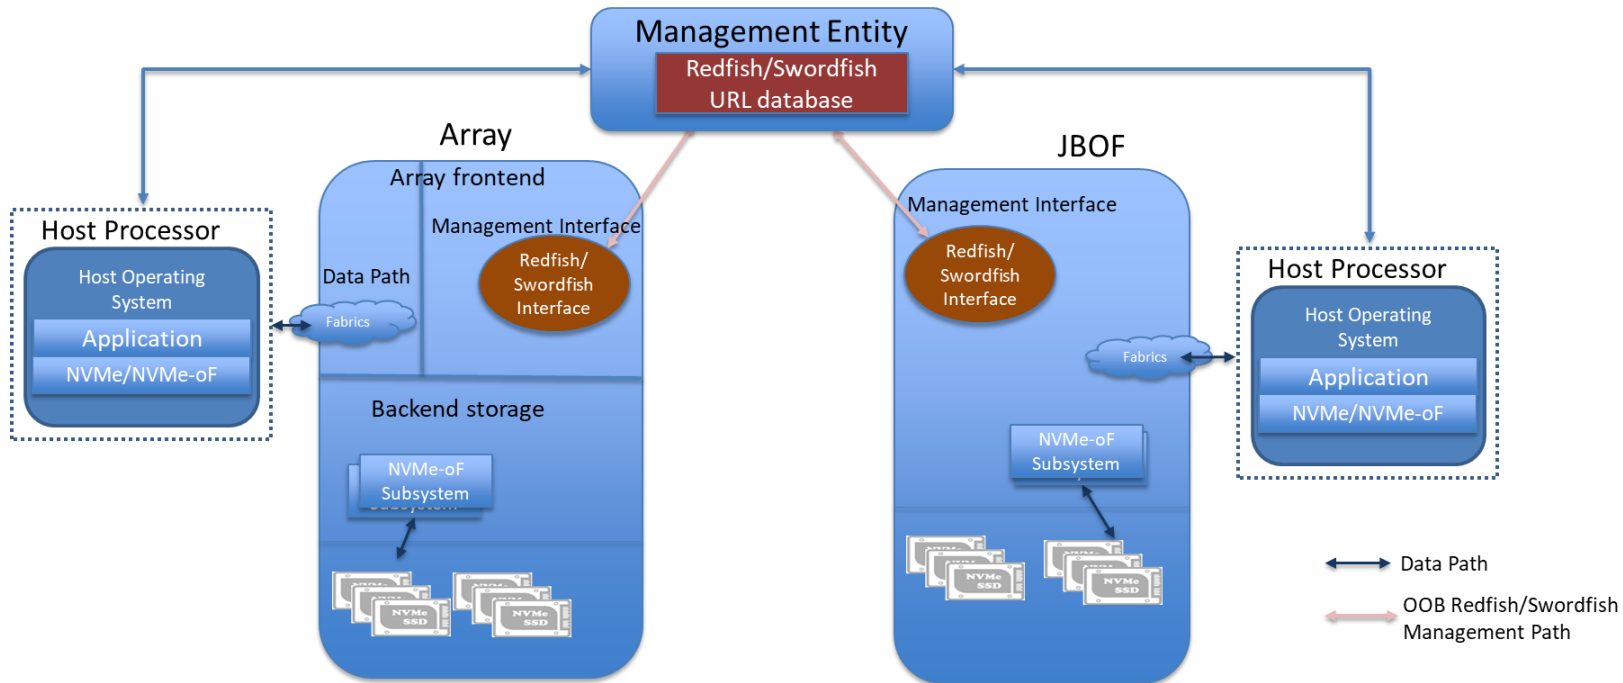
- Out of Band PCIe VDM
- Out of Band SMBus
- In Band



# NVMe Management with Swordfish: PCIe attached SSD



# NVMe Management with Swordfish: Array/JBOF



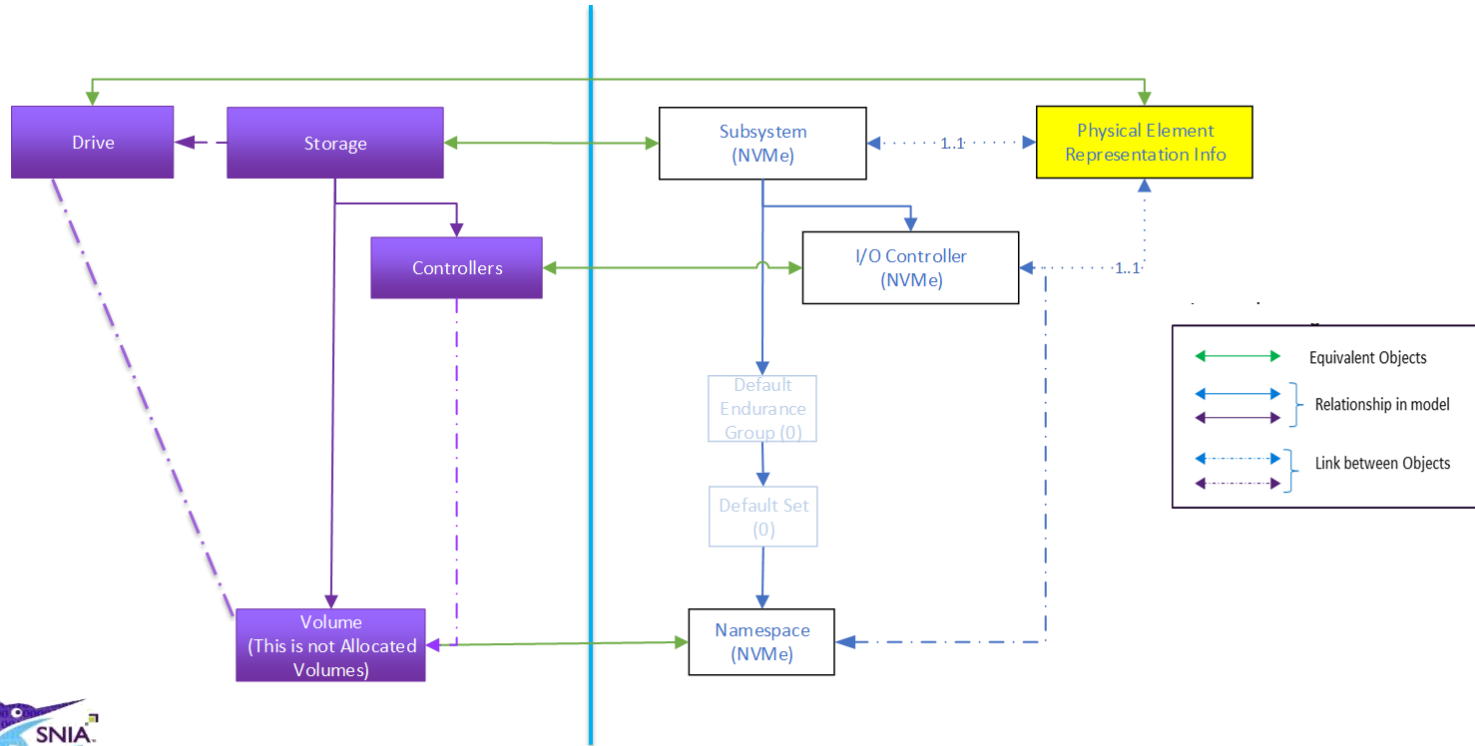
\*One of the possible Models - implementation layer choice



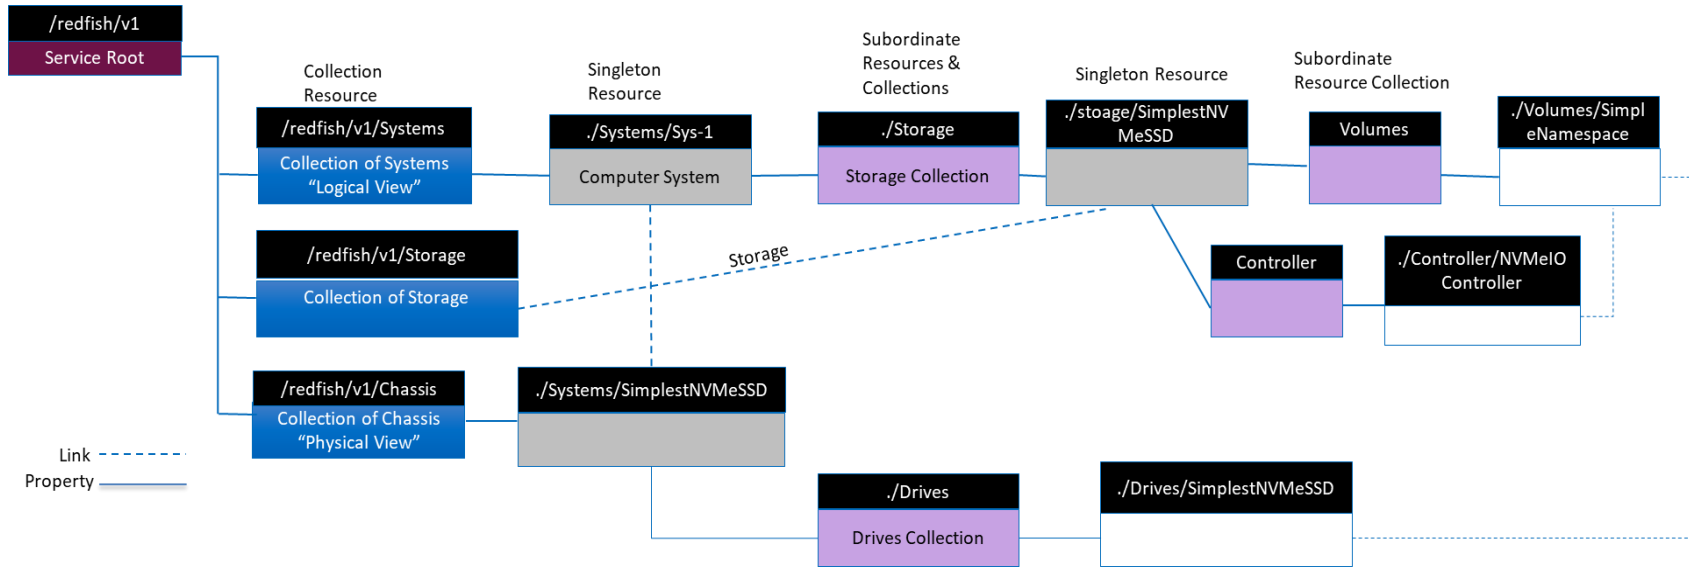


# NVMe Swordfish Models

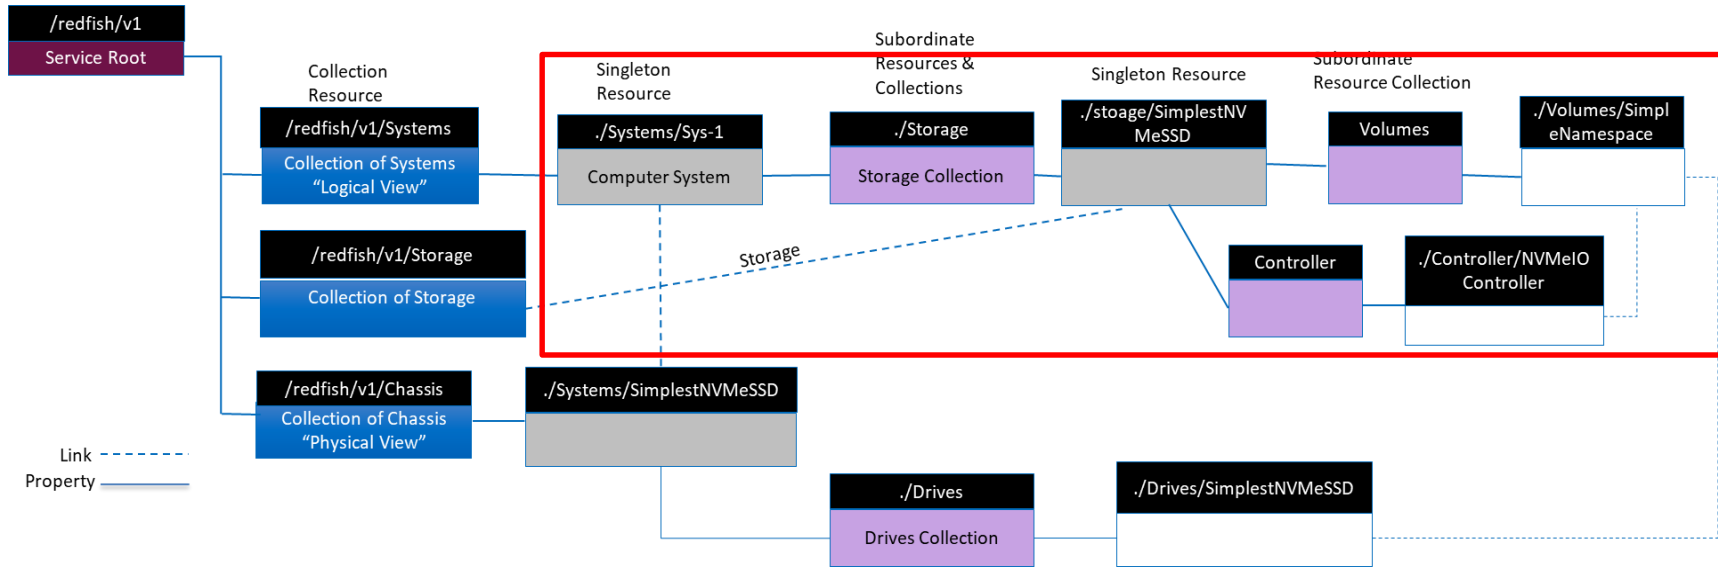
# NVMe mapping in Swordfish



# Swordfish NVMe Model



# Swordfish NVMe Model



# Swordfish configurations: NVMe (systems/sys-1)

/redfish/v1/Systems/Sys-1



- 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": "224071-323",
  "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."
}
```

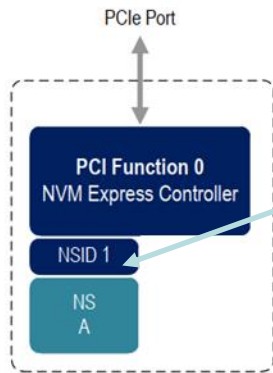
/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD

```
{
  "@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/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace



```
{
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved.",
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace",
  "@odata.type": "#Volume.v1_5_0.Volume",
  "Id": "1",
  "Name": "Namespace 1",
  "LogicalUnitNumber": 1,
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "Identifiers": [{
    "DurableNameFormat": "NQN",
    "DurableName": "nqn.2014-08.org.nvmexpress:uuid:6c5fe566-10e6-4fb6-aad4-8b4159029384"
  }],
  "Capacity": {
    "Data": {
      "ConsumedBytes": 0,
      "AllocatedBytes": 10737418240
    }
  },
  "NVMeNamespaceProperties": {
    "NamespaceId": "0x22F",
    "NamespaceFeatures": {
      "SupportsThinProvisioning": false,
      "SupportsAtomicTransactionSize": false,
      "SupportsDeallocatedOrUnwrittenLBAError": false,
      "SupportsNGUIDReuse": false,
      "SupportsIOPerformanceHints": false
    },
    "NumberLBAFormats": 0,
    "FormattedLBASize": "LBAFormat0Support",
    "MetadataTransferredAtEndOfDataLBA": false,
    "NVMeVersion": "1.4"
  },
  "Links": {}
}
```



# Example: Getting Namespace info

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

```
{
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace",
  "@odata.type": "#Volume.v1_5_0.Volume ",
  "@odata.context": "/redfish/v1/$metadata#Drive.Drive",
  "Id": "1",
  "Name": "Namespace 1",
  "CapacityBytes": 10737418240,
  "Oem": {
  },
  "Status": {
    "Health": "OK",
    "State": "Enabled"
  },
  "PhysicalLocation": {
    "PartLocation": {
      "LocationType": "Slot",
      "ServiceLabel": "NVME-1"
    }
  }
}
```



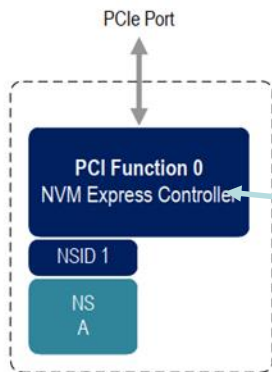
# Example: Updating Namespace State

```
PATCH -data '{"status": {"state": "Disabled"}}'  
/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace HTTP/1.1
```

```
{  
  "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Volumes/SimpleNamespace",  
  "@odata.type": "#Volume.v1_5_0.Volume",  
  "@odata.context": "/redfish/v1/$metadata#Drive.Drive",  
  "Id": "1",  
  "Name": "Namespace 1",  
  "CapacityBytes": 10737418240,  
  "Oem": {  
  },  
  "Status": {  
    "Health": "OK",  
    "State": "Disabled"  
  },  
  "PhysicalLocation": {  
    "PartLocation": {  
      "LocationType": "Slot",  
      "ServiceLabel": "NVME-1"  
    }  
  }  
}
```

# Swordfish configurations: NVMe (Controller)

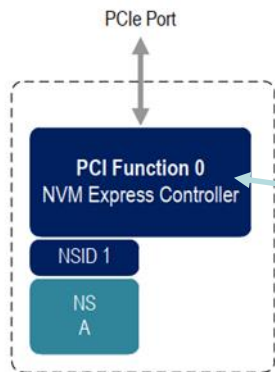
/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers/NVMeIOController



```
{
  "@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"
  ],
  "SupportedDeviceProtocols": [
    "NVMe"
  ],
}
```

# Swordfish configurations: NVMe (Controller...)

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



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

# Example: Getting Controller info

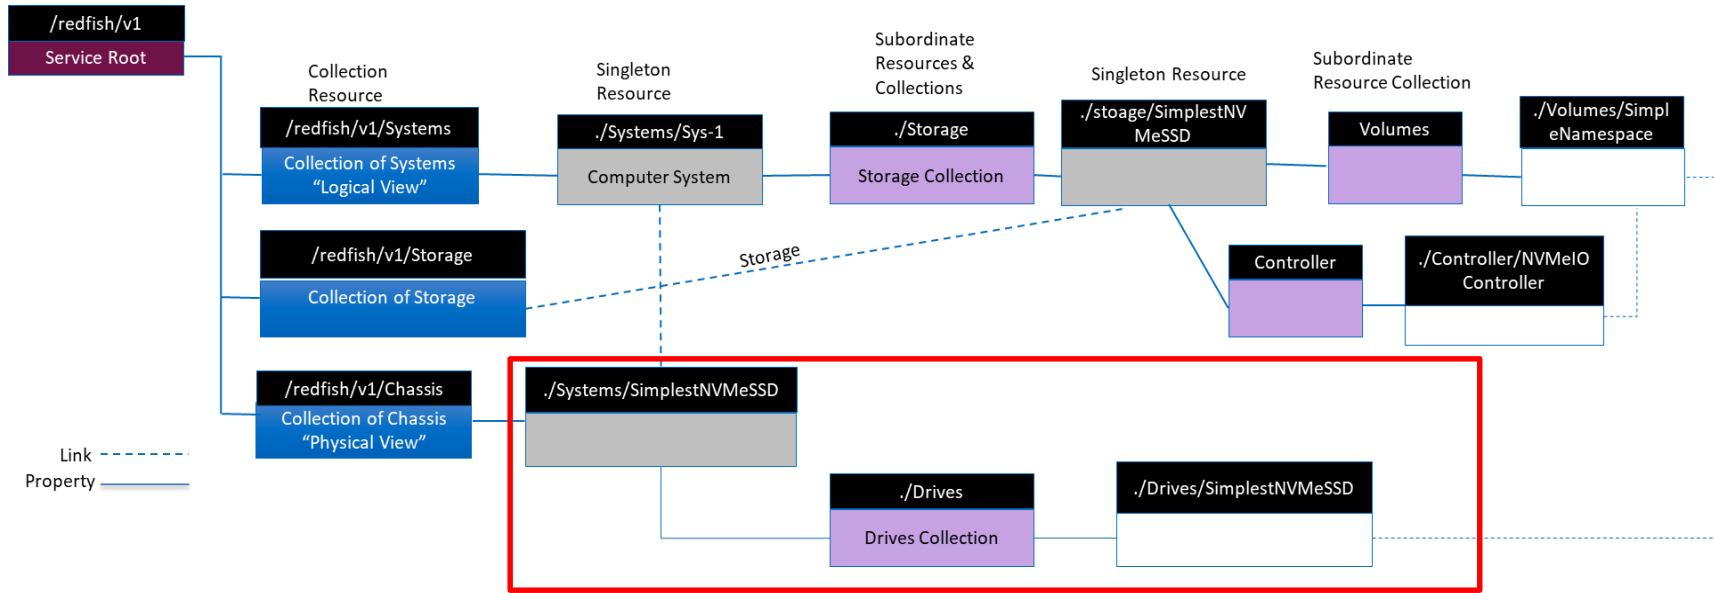
GET /redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers/NVMeIOController HTTP/1.1

```
{
  "@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 ",
  "Manufacturer": "Best NVMe Vendor",
  "Model": "Simple NVMe Device"
  "SerialNumber": "NVME123456 ",
  "FirmwareVersion": "1.0.0",
  ....
  "ReportsNamespaceGranularity": "false",
  ....
}
```

PATCH -data '{"NVMeControllerProperties": {"ReportsNamespaceGranularity": "True"}}'  
/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD/Controllers/NVMeIOController HTTP/1.1

```
{
  "@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 ",
  "Manufacturer": "Best NVMe Vendor",
  "Model": "Simple NVMe Device"
  "SerialNumber": "NVME123456 ",
  "FirmwareVersion": "1.1.0",
  ....
  "ReportsNamespaceGranularity": "True",
  ....
}
```

# Swordfish NVMe Model



# Swordfish configurations: NVMe (Drive)

/redfish/v1/Chassis/SimplestNVMeSSD


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

```
{
  "@odata.type": "#Chassis.v1_14_0.Chassis",
  "Id": "SimplestNVMeSSD",
  "Name": "SimplestNVMeSSD",
  "ChassisType": "Module",
  "Manufacturer": "NVMeDriveVendorFoo",
  "Model": "NVMeMODEL",
  "SKU": "6914260",
  "SerialNumber": "529QB9450R6",
  "PartNumber": "166480-S23",

  "PowerState": "On",
  "IndicatorLED": "Off",
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "Thermal": {
    "@odata.id": "/redfish/v1/Chassis/SimplestNVMeSSD/Thermal"
  },
  "Drives": {
    "@odata.id": "/redfish/v1/Chassis/SimplestNVMeSSD/Drives"
  },
  "Links": {
    "Storage": [
      {
        "@odata.id": "/redfish/v1/Systems/Sys-1/Storage/SimplestNVMeSSD"
      }
    ]
  },
  "@odata.id": "/redfish/v1/Chassis/SimplestNVMeSSD",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```



# Swordfish configurations: NVMe (Drive..)

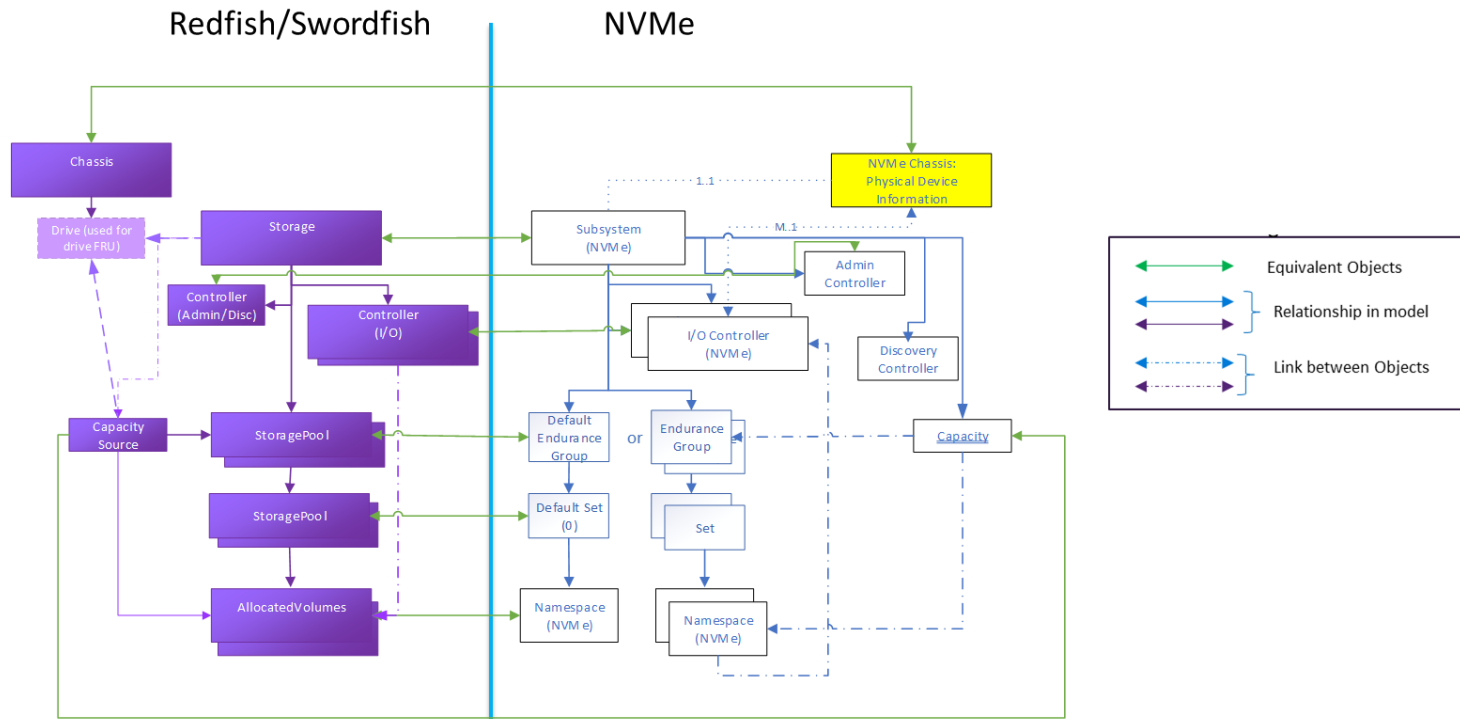


SSD Drive

/redfish/v1/Chassis/SimplestNVMeSSD/Drives/SimplestNVMeSSD

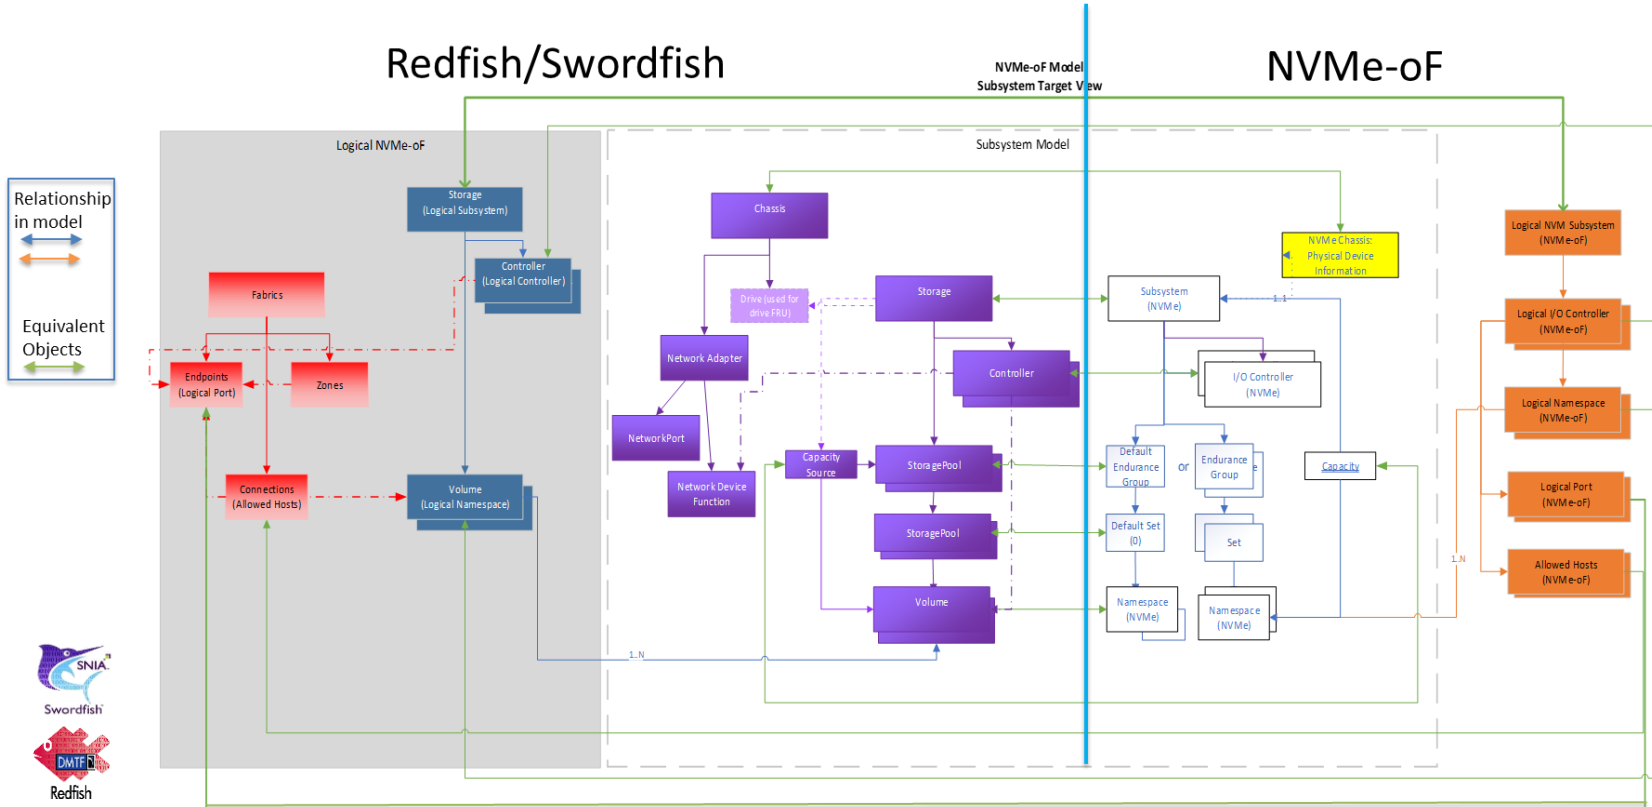
```
{
  "@odata.id": "/redfish/v1/Chassis/NVMeDrive1/Drives/NVMeDrive1",
  "@odata.type": "#Drive.v1_9_0.Drive",
  "IndicatorLED": "Lit",
  "Model": "ST9146802SS",
  "Revision": "S20A",
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "CapacityBytes": 899527000000,
  "FailurePredicted": false,
  "Protocol": "NVMe",
  "MediaType": "SSD",
  "Manufacturer": "Contoso",
  "SerialNumber": "72D0A037FRD26",
  "PartNumber": "SG0GP8811253178M02GJA00",
  "Identifiers": [{
    "DurableNameFormat": "NAA",
    "DurableName": "500003942810D13A"
  }],
  "CapableSpeedGbs": 12,
  "NegotiatedSpeedGbs": 12,
  "Actions": {
    "#Drive.Reset": {
      "target": "/redfish/v1/Chassis/NVMeDrive1/Drives/NVMeDrive1/Actions/Drive.Reset"
    }
  }
}
```

# NVMe with Endurance Groups and Sets

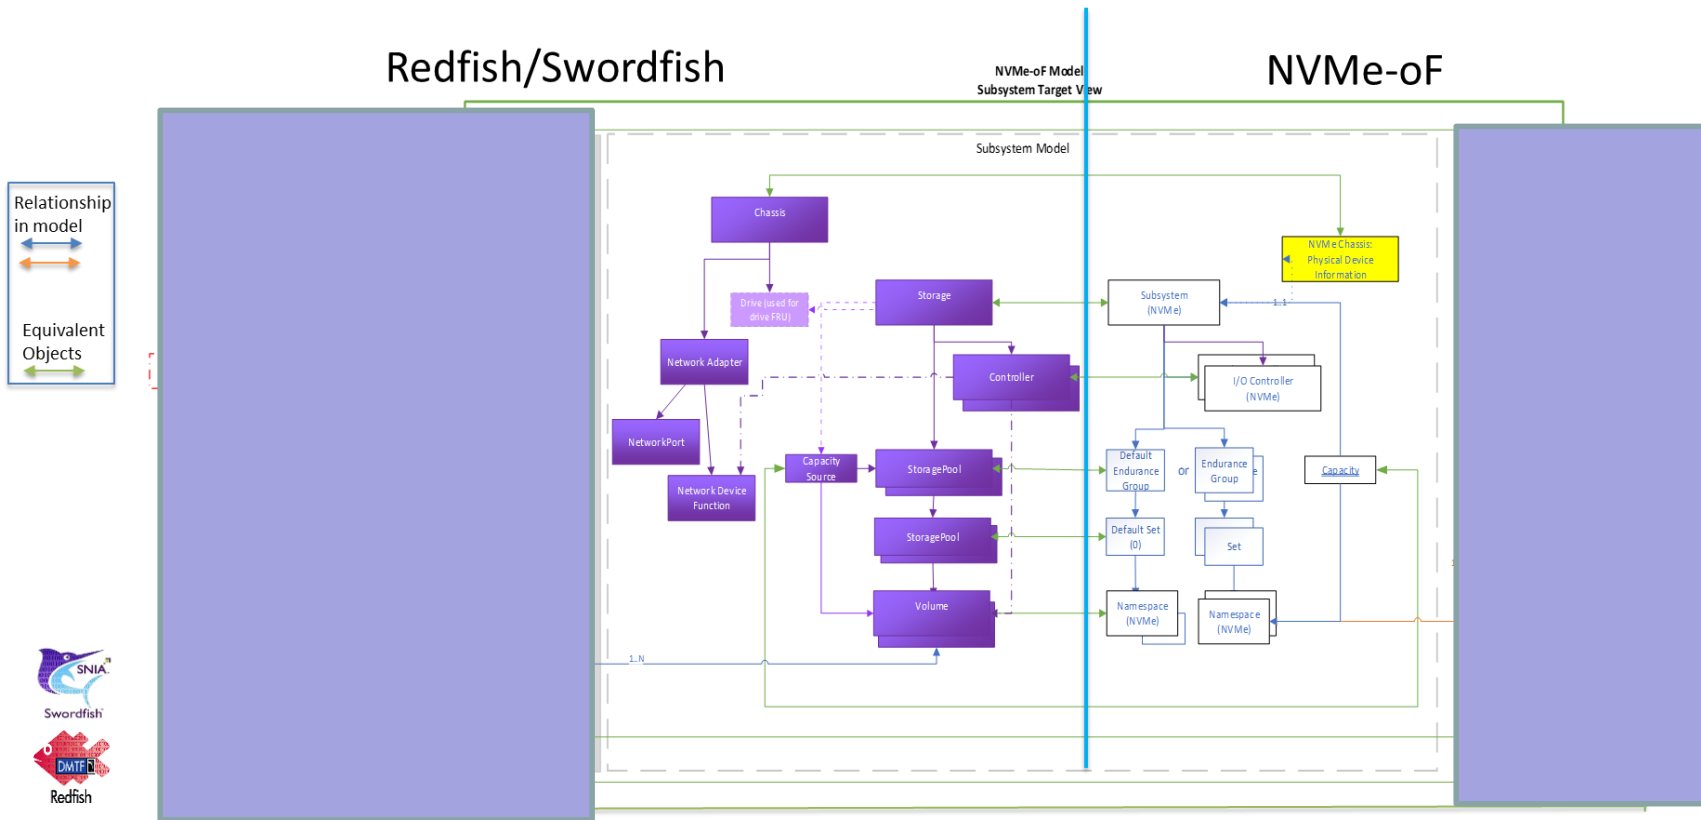




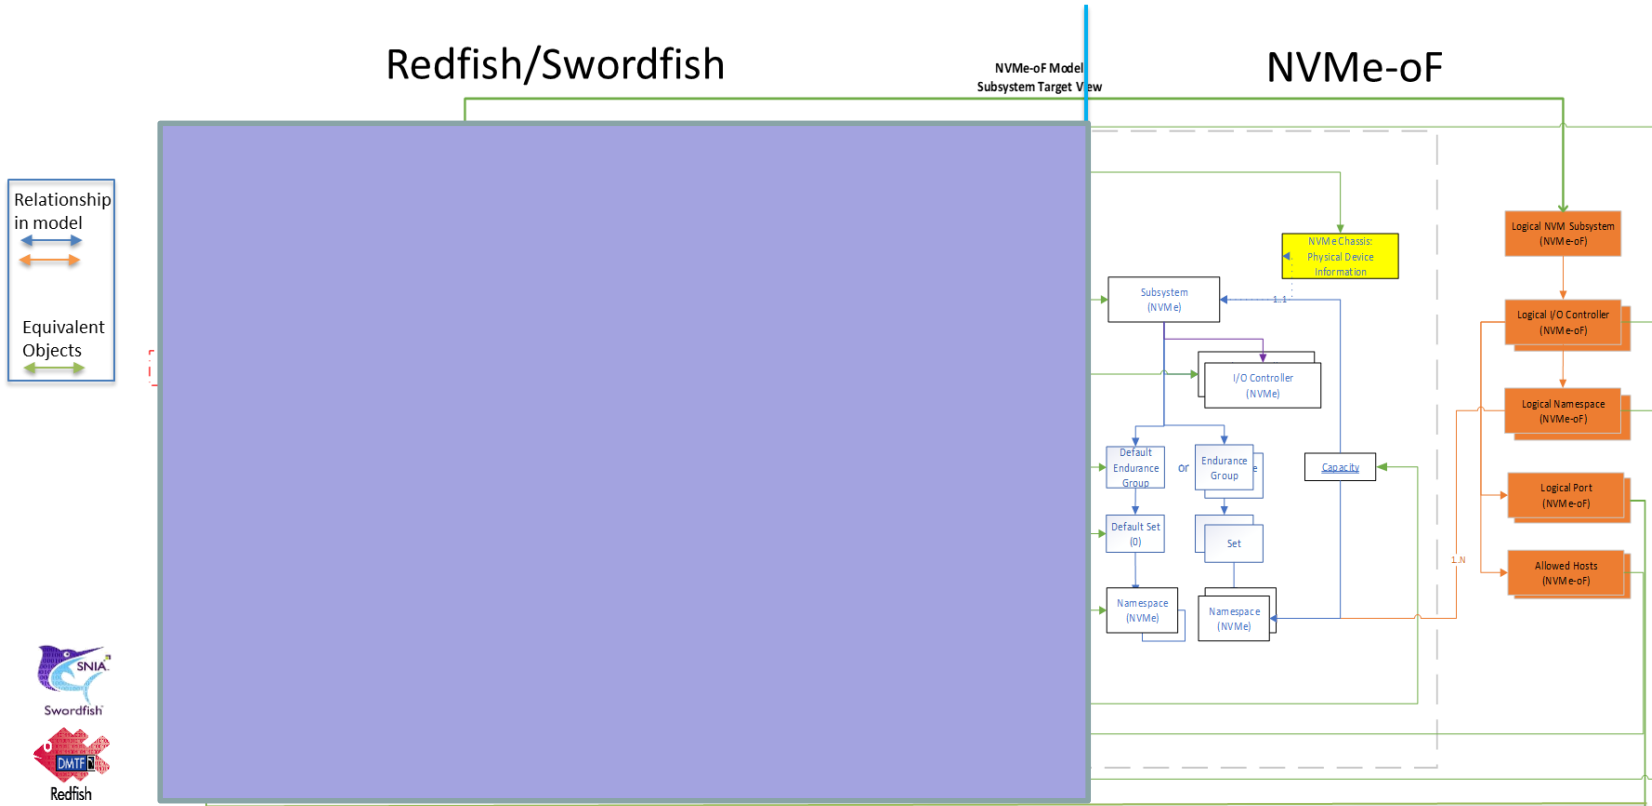
# NVMe-oF Mapping in Swordfish



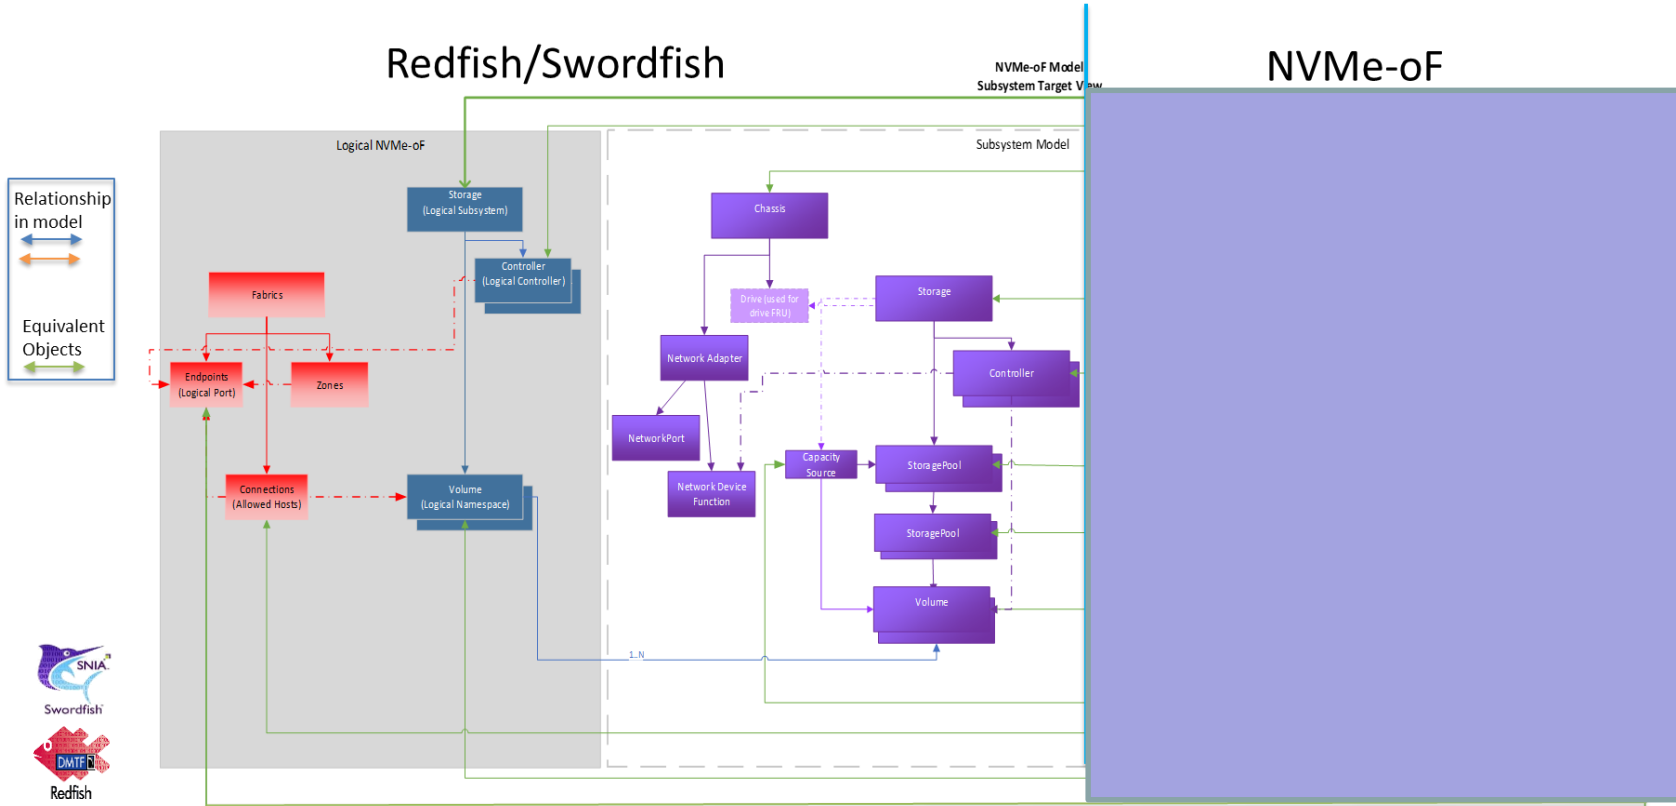
# NVMe-oF Mapping in Swordfish



# NVMe-oF Mapping in Swordfish



# NVMe-oF Mapping in Swordfish



# Redfish/Swordfish configurations: NVMe-oF (Fabrics)

/redfish/v1/Fabrics

```
{
  "@odata.type": "#FabricCollection.FabricCollection",
  "Name": "Fabric Collection",
  "Members@odata.count": 1,
  "Members": [
    {
      "@odata.id": "/redfish/v1/Fabrics/NVMe-oF"
    }
  ],
  "@odata.id": "/redfish/v1/Fabrics",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```

```
{
  "@odata.type": "#Fabric.v1_2_0.Fabric",
  "Id": "NVMe-oF",
  "Name": "NVMe-oF Fabric",
  "FabricType": "NVMeOverFabrics",
  "Status": {
    "State": "Enabled",
    "Health": "OK"
  },
  "Zones": {
    "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Zones"
  },
  "Endpoints": {
    "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Endpoints"
  },
  "Connections": {
    "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Connections"
  },
  "@odata.id": "/redfish/v1/Fabrics/NVMe-oF",
  "@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```



Swordfish



Redfish

# Redfish/Swordfish configurations: NVMe-oF

## (Endpoints and Connections)

/redfish/v1/Fabrics/NVMe-oF/Endpoints/NVMeEndpoint

```
{
  "@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": "F0ECBA9876543210h"
    }]
  }],

  "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."
}
```

/redfish/v1/Fabrics/NVMe-oF/Connections/Host1

```
{
  "@odata.type": "#Connection.v1_0_0.Connection",
  "Id": "Host1",
  "Name": "Connection info for host 1",
  "ConnectionType": "Storage",
  "VolumeInfo": [{
    "AccessCapabilities": ["Read", "Write"],
    "Volume": {
      "@odata.id": "/redfish/v1/Storage/NVMe-oF-Subsystem/Volumes/LogicalNamespace1"
    }
  }],
  "Links": {
    "Endpoints": [{
      "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Endpoints/Host"
    },
    {
      "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Endpoints/NVMeEndpoint"
    }
  ],
  "Zones": [{
    "@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Zones/1"
  }]
},
"@odata.id": "/redfish/v1/Fabrics/NVMe-oF/Connections/Host1",
"@Redfish.Copyright": "Copyright 2014-2020 SNIA. All rights reserved."
}
```



Swordfish



Redfish

# Summary & Wrapup

- Mockup for various models  
PCIe attached NVMe, JBOF, NVMe-oF, ...  
Mockups @ <http://swordfishmockups.com>
- NVMe-Swordfish Mapping Guide  
“Swordfish NVMe Model Overview and Mapping Guide”
- Latest Swordfish bundle released in Aug 2020  
1.2.1 version

# Thank you for watching

- **SNIA Swordfish™ Standards**
  - Schemas, Specs, Mockups, Users Guide, Practical Guide & more  
<https://www.snia.org/swordfish>
- **Redfish / Swordfish Specification Forum**
  - This is where you can ask and answer questions about Redfish and 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 and participate: [https://www.snia.org/member\\_com/join-SNIA](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>







# Thank you



**Please take a moment  
to rate this session.**

**Your feedback matters to us.**