STORAGE DEVELOPER CONFERENCE
SDC 21
BY Developers FOR Developers
Virtual Conference
September 28-29, 2021
A SNIA Event

Managing Open Fabrics with a Standards-based Interface: Bringing Gen-Z, Redfish, and Swordfish Together

Erich Hanke
Principal Engineer of Storage and Memory Products
IntelliProp

Agenda

- OFA and DMTF Redfish and SNIA Swordfish Collaboration Overview
- OFA Open Fabric Management Framework Overview
- Gen-Z “Zephyr” Fabric Manager Introduction
- OFA and Gen-Z Consortium Proof of Concept

The Ecosystem for Fabrics in the Datacenter is Changing

The Fabric Landscape is Changing

- Rapidly increasing types of fabric interconnects
- Each fabric has its strengths, features, and management tools
- Each fabric has its own configuration mechanisms and interfaces

The Workload and Resource Ecosystem is Changing

- New compute, storage, and accelerator resources are becoming available
- HPC Clusters and Cloud Computing Environments:
  - Running increasingly diverse and dynamic workloads
  - Incorporating both distributed computing capabilities and heterogeneous hardware solutions

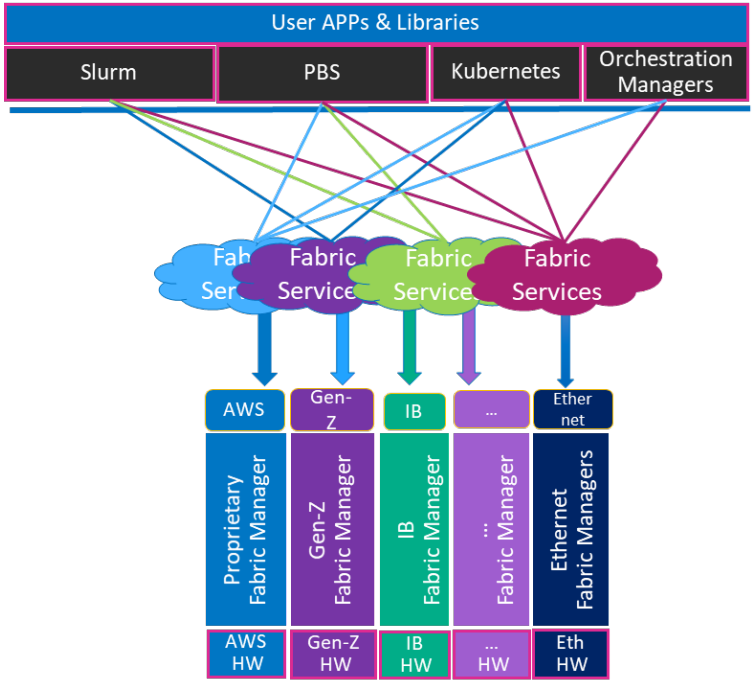

This Creates Problems in Manageability

Administrative Management Challenges

- No common fabric manager (FM) interface or fabric model available to link applications with remote resource
- Workload management and optimization is different for each type of fabric
- Administrators are being asked to manage and increasingly heterogeneous fabric infrastructure, each with its own management standard and model

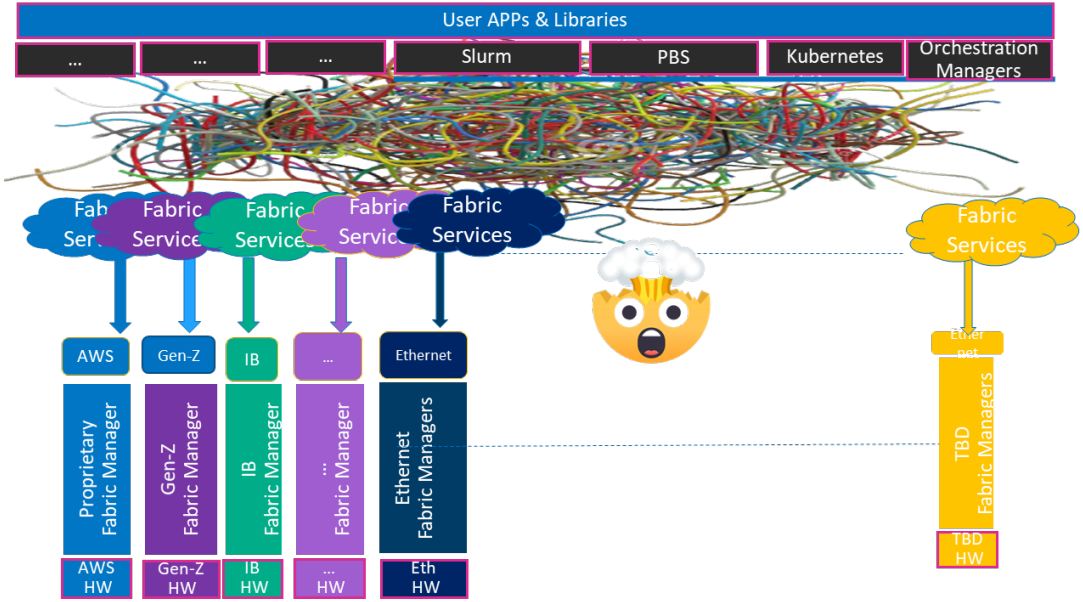

We Can Fix These Management Problems

- **Need:**
  - **Interoperability**
  - **Standards based Ecosystem Management**
- **Keep fabric specific management where required**
- **Create Open and General Management Interface**

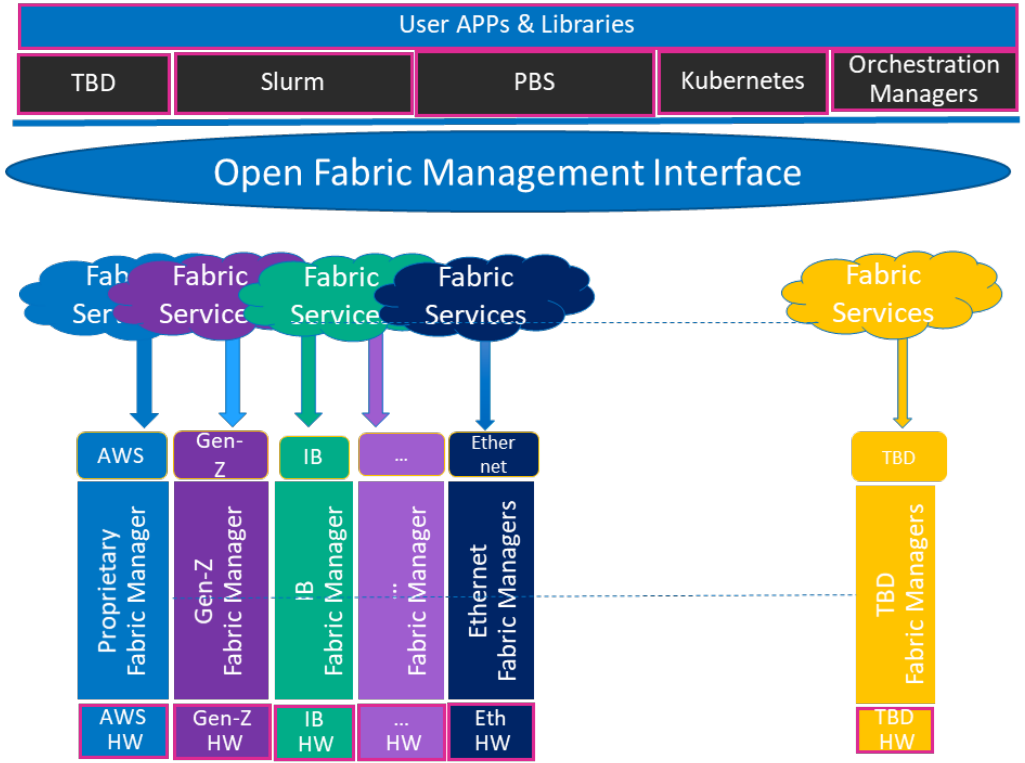

Generic Fabric Management

- **Common APIs**
  - Control Services
    - Discovery and Inventory
  - Communication Services
    - Connection Management
    - Address Vectors
  - Partition Services
    - Zones
    - Connections
  - Messaging Services
    - Queues and Contexts
    - Events and Errors
    - Atomics and other Sync
  - Security
    - Encryption
    - Authentication
    - Isolation

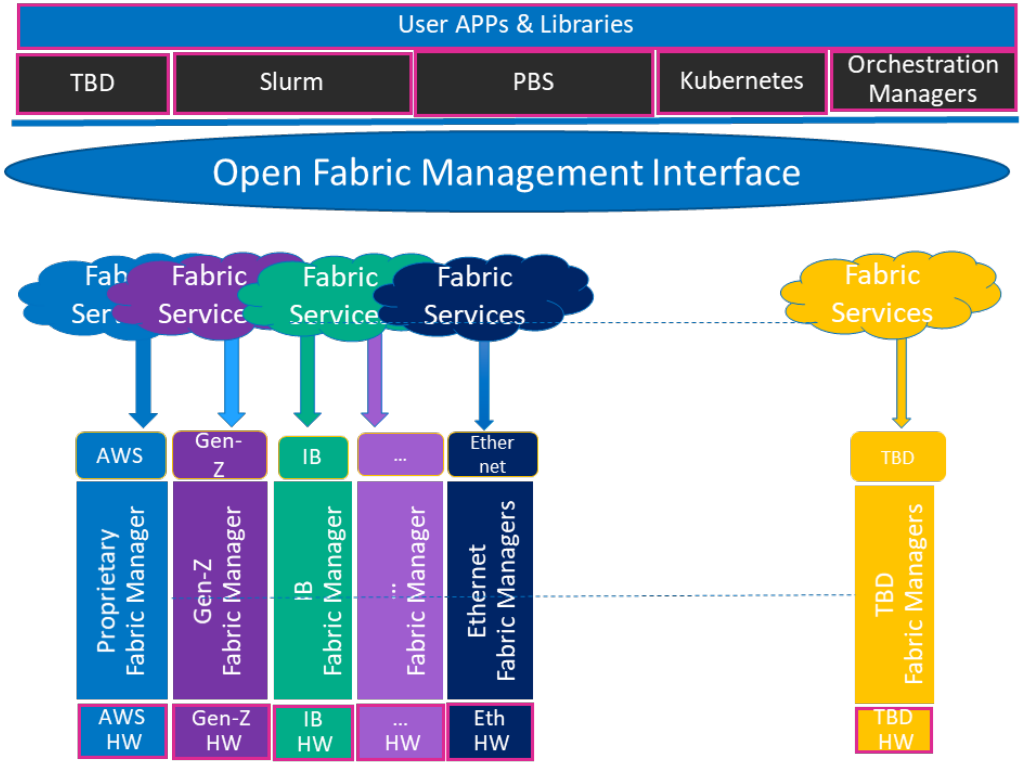

Open Fabric Management Framework Architecture

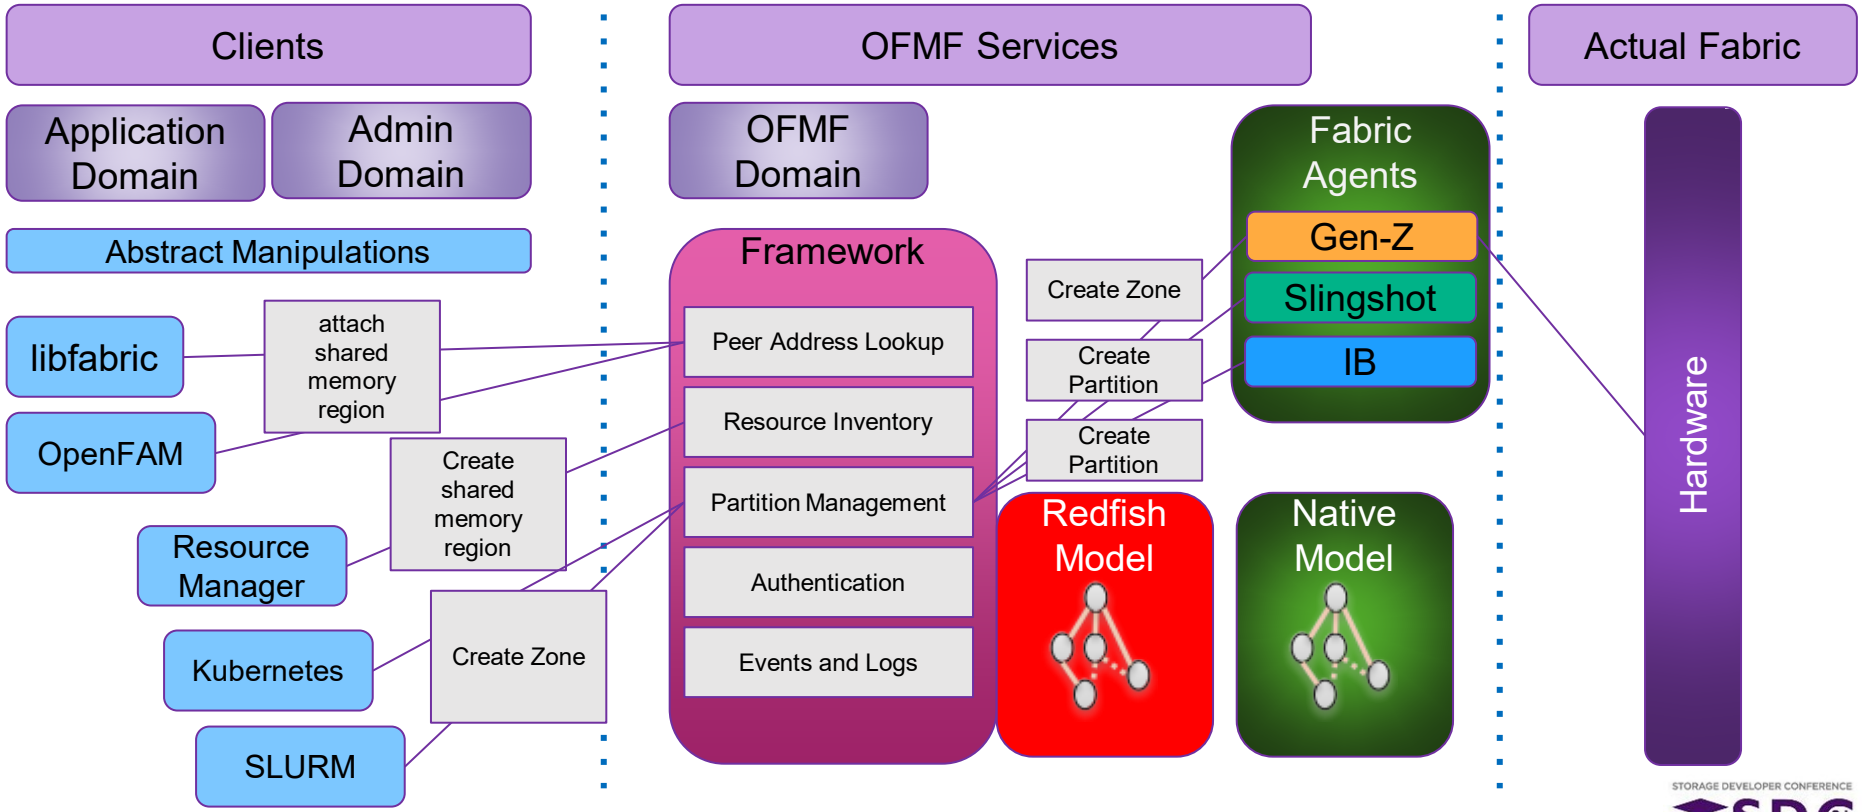

OFA and DMTF Redfish and SNIA Swordfish Collaboration

- Open Fabrics Alliance / Gen-Z Consortium MOU
  - Standardization of Open Source Fabric Management Software
- DMTF (formerly known as: Distributed Management Task Force)
  - DMTF Creates open manageability standards spanning cloud, virt, network, servers, and storage
- SNIA Swordfish
  - Unified approach for the management of storage and servers in hyperscale and cloud infrastructure environments
  - Extension of DMTF Redfish specification

Fabric Specific Agent

- One Agent per vendor-specific fabric implementation
- Provides a connection from OFMF to VS FM
- Represents the underlying Fabric object to OFMF
  - Listens to subnet-manager
  - Translation layer for fabric-specific taxonomy to Redfish fabric schema
  - Translates logical connection information to physical routes
  - Communicates the updates to OFMF
    - Underlying hardware element information
    - Connection information, links, ports, and paths

OFMF Planned Work Items

- Gather more client-driven use-cases
- Map together redfish/swordfish management interface with OFA Open Fabric Manager functionality
- Ensure wide fabric management coverage:
  - Gen-Z
  - Slingshot
  - InfiniBand
  - OmniPath
  - RoCE
  - iWARP
  - Ethernet
  - FiberChannel
  - Future Fabrics...

Proof of Concept

- Build on Emulation Environments to showcase functionality
- Utilize a set of PoC hardware and PoC FM that is under development
- Put a stake in the ground for an industry Convention for first PoC

- Gen-Z Consortium Proof of Concept Working Group (PoCWG)
  - Prototype Hardware
    - Host Bridges / Fabric Adaptor
    - ARM Hosts
    - x86 Hosts
    - Discrete Switches and Integrated Switches
    - Gen-Z Memory Module (ZMM) for Fabric Attached Memory (FAM)

Gen-Z Zephyr Fabric Manager

Linux Based Gen-Z Subsystem, Bridge driver and Fabric management software

- Gen-Z Subsystem, Fabric Manager (Zephyr), and Linux Local Management Service (LLaMaS) are open source and available on Github:
  - https://github.com/linux-genz
- IntelliProp Vendor Bridge driver for Orthus/Sphinx Fabric Adaptor Bridges
- Gen-Z Utilities and debug tools available on Github:
  - https://github.com/linux-genz/genz-utils

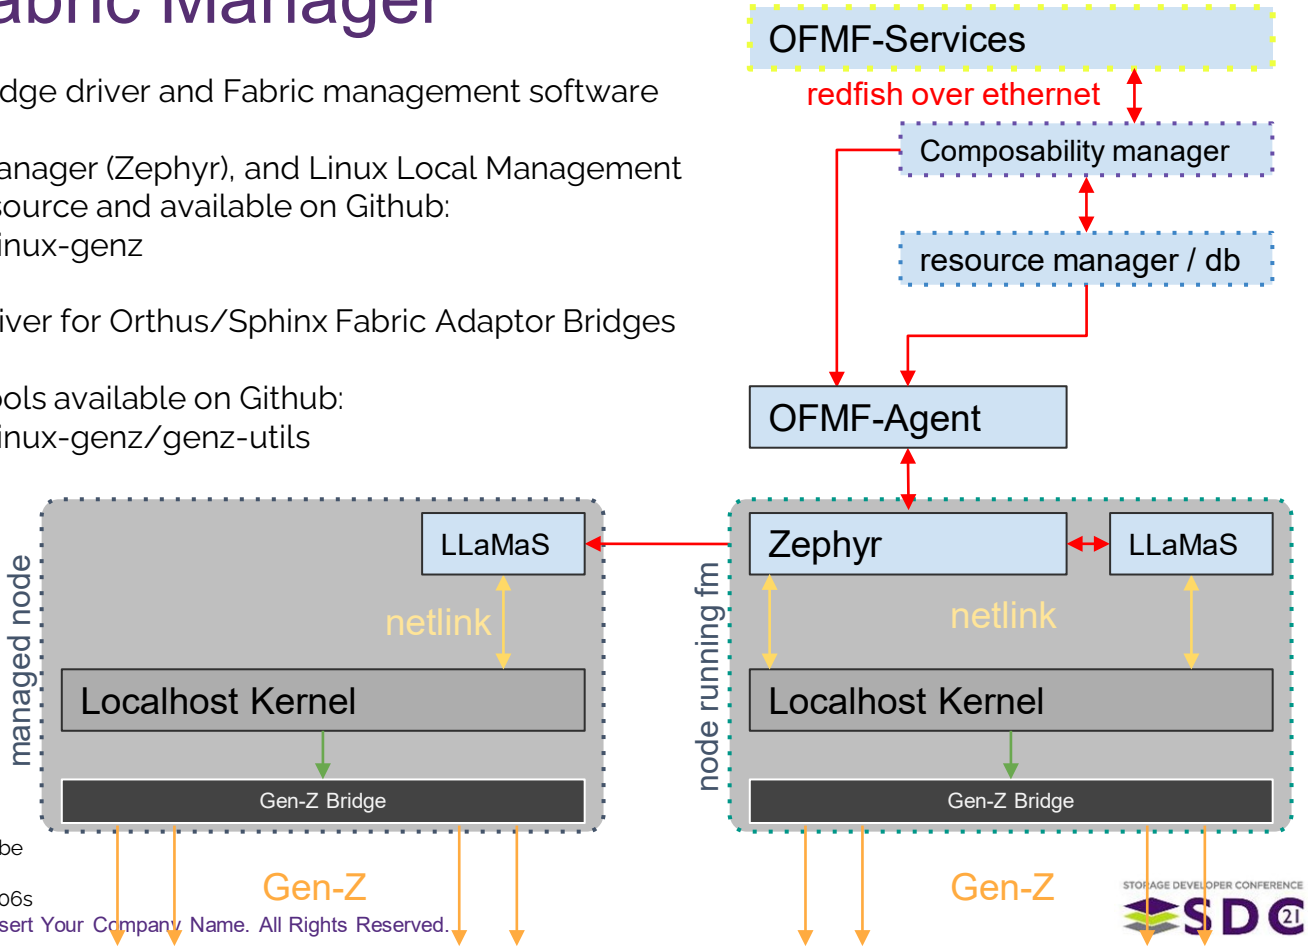

*See the full OFA Presentation by Jim Hull on OFA YouTube channel!
https://www.youtube.com/watch?v=6XWgGYyhg6s&t=606s

Sysfs tree

/sys/devices/genz/bridge0

OFA and Gen-Z Consortium PoC

Gen-Z Zephyr FM

- Performs recursive walk of fabric
- Configures components
- Makes routing connections between components based on requests
- Represents fabric topology and object descriptions to Agent
- Utilizes a python module, networkX, to create topology graphs

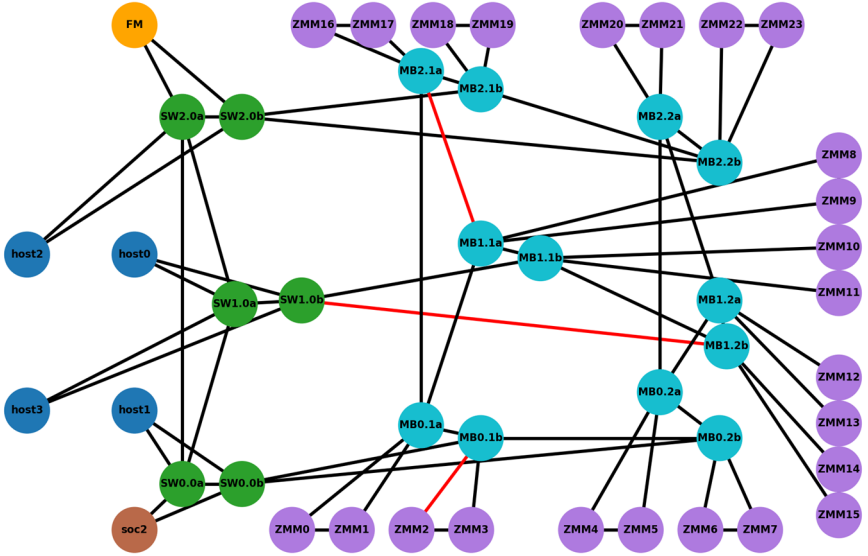

Demonstration Hardware

BittWare XUP-P3R:

Functions:

- CXL/PCIe Bridges (GZB):
  a. Typhon / Sphinx

Gen-Z FAM ZMM:

Functions:

1. Gen-Z FAM
2. Gen-Z CFAM

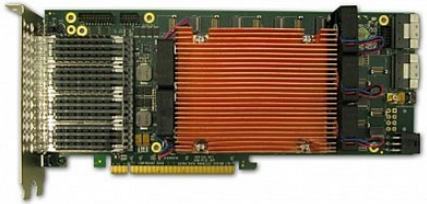

Alpha Data ADM-PCIE-9H7:

Functions:

- Gen-Z 12P Switch

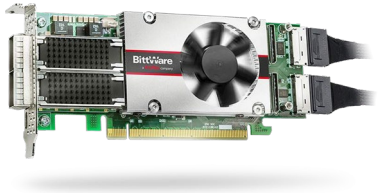

BittWare 250-SoC

Functions:

- Orthus ARM Host

SC21 Demo of PoC

PoC Demo:

- HyperX Fabric Topology
- Management running OFMF Services and PoC Fabric Agent connection to Zephyr
- Multipath Fabric Demo showcasing memory fabric resiliency

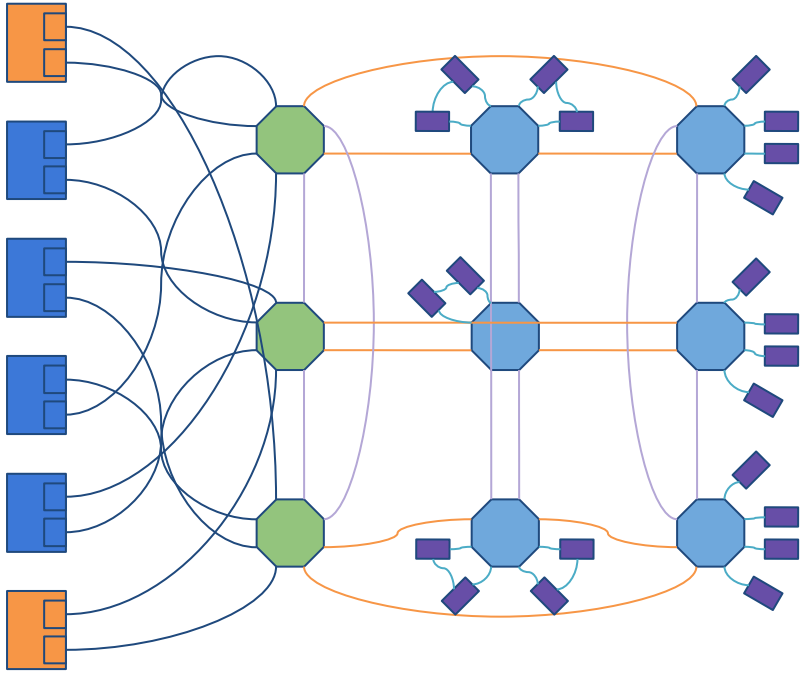

Key:

- New coherent host bridge / fabric adaptor
- Gen-Z Discrete Switch Boxes
- Gen-Z Media Boxes with ZMM
- ARM Based Gen-Z Host
- Gen-Z Fabric Attached Memory (ZMM)

Super Computing 2021 (SC21)

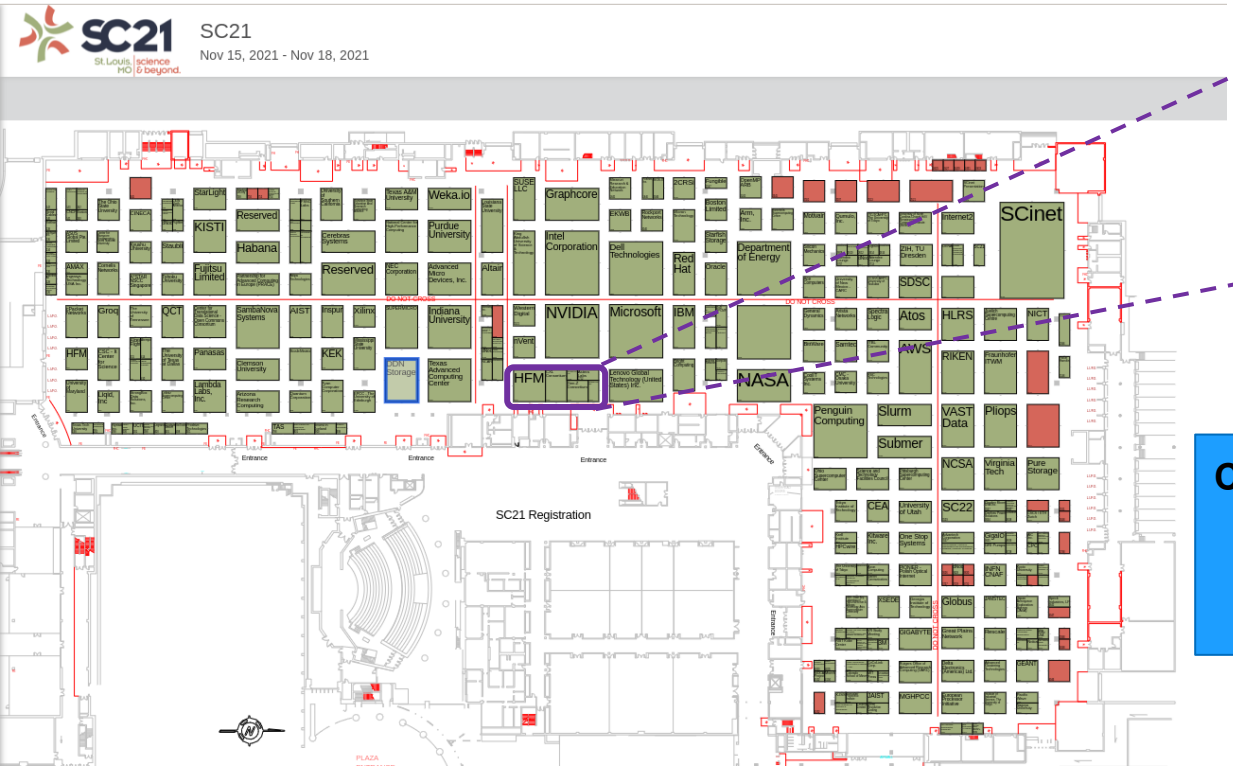

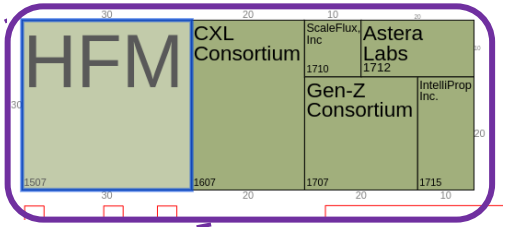

Come see a live (or virtual) demo!

- Gen-Z Consortium [1707]
- Open Standards Pavilion [1507]
- IntelliProp [1715]

How to Learn More or Contribute

SNIA Swordfish™

- Swordfish Standards
  - Schemas, Specs, Mockups, User and Practical Guides, … 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

How to Learn More or Contribute

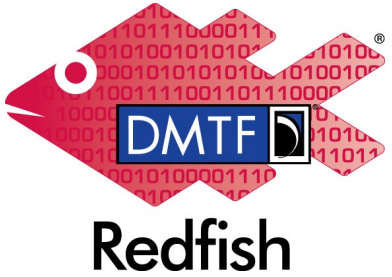

DMTF Redfish™

- Redfish Standards
  - Specifications, whitepapers, guides,… https://www.dmtf.org/standards/redfish
- OpenFabrics Alliance: OFMF
- 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 OpenFabrics Alliance
  - https://www.openfabrics.org/membership-how-to-join/

How to Learn More or Contribute

Gen-Z Consortium

- Gen-Z Specification
  - Specifications, whitepapers, guides,… https://genzconsortium.org
  - How to Join: https://genzconsortium.org/about-us/membership/become-a-member

- Gen-Z Linux Subsystem
  - Github
    - Linux Subsystem: https://github.com/linux-genz
    - MicroDevelopment Kit (uDK): https://github.com/linux-genz/udk
    - Gen-Z Utilities: https://github.com/linux-genz/genz-utils

- GLSS mailing list subscription
  - https://groups.google.com/g/genz-linux

Thank you!