



OCP SUMMIT

March 20-21
2018
San Jose, CA

OPEN. FOR BUSINESS.



SAS / SATA / NVMe Storage

Rackspace Barreleye G2 48V OpenPOWER Platform

Adi Gangidi

Sr. Systems Design Engineer, Rackspace

OPEN. FOR BUSINESS.




Agenda

- Introduction
- What is tri-mode ?
- Why tri-mode ?
- Implementation
 - Details & Alternatives
- Samples




Zaius & Barreleye G2

Capabilities



ZAIUS Motherboard

- 2 x POWER9 LaGrange
- 48V input
- Front IO & service access
- 80 Lanes of PCIe Gen4
- 32 Lanes of OpenCAPI / NVLink 2.0
- Open Source BMC & Host Firmware



BARRELEYE G2 Chassis


- Full-depth 48V open rack v2
- Hot swap fans and VGA access
- 2 OU chassis supports FHFL cards
- *High density & hot swap storage bay*
- *Tri-mode Support (SAS / SATA / NVMe)*




Storage Today & Past

- Separate connector pinout / backplane
- Separate mid-plane
- Separate controller




SAS



SATA





NVMe



Tri-Mode Storage:

- Interchangeable SAS / SATA / NVMe Support in the SAME slot
- Hot-swap between different devices (Where firmware support is available)
- 1 Backplane
- 1 Mid-plane / Expander board
- 1 Controllers

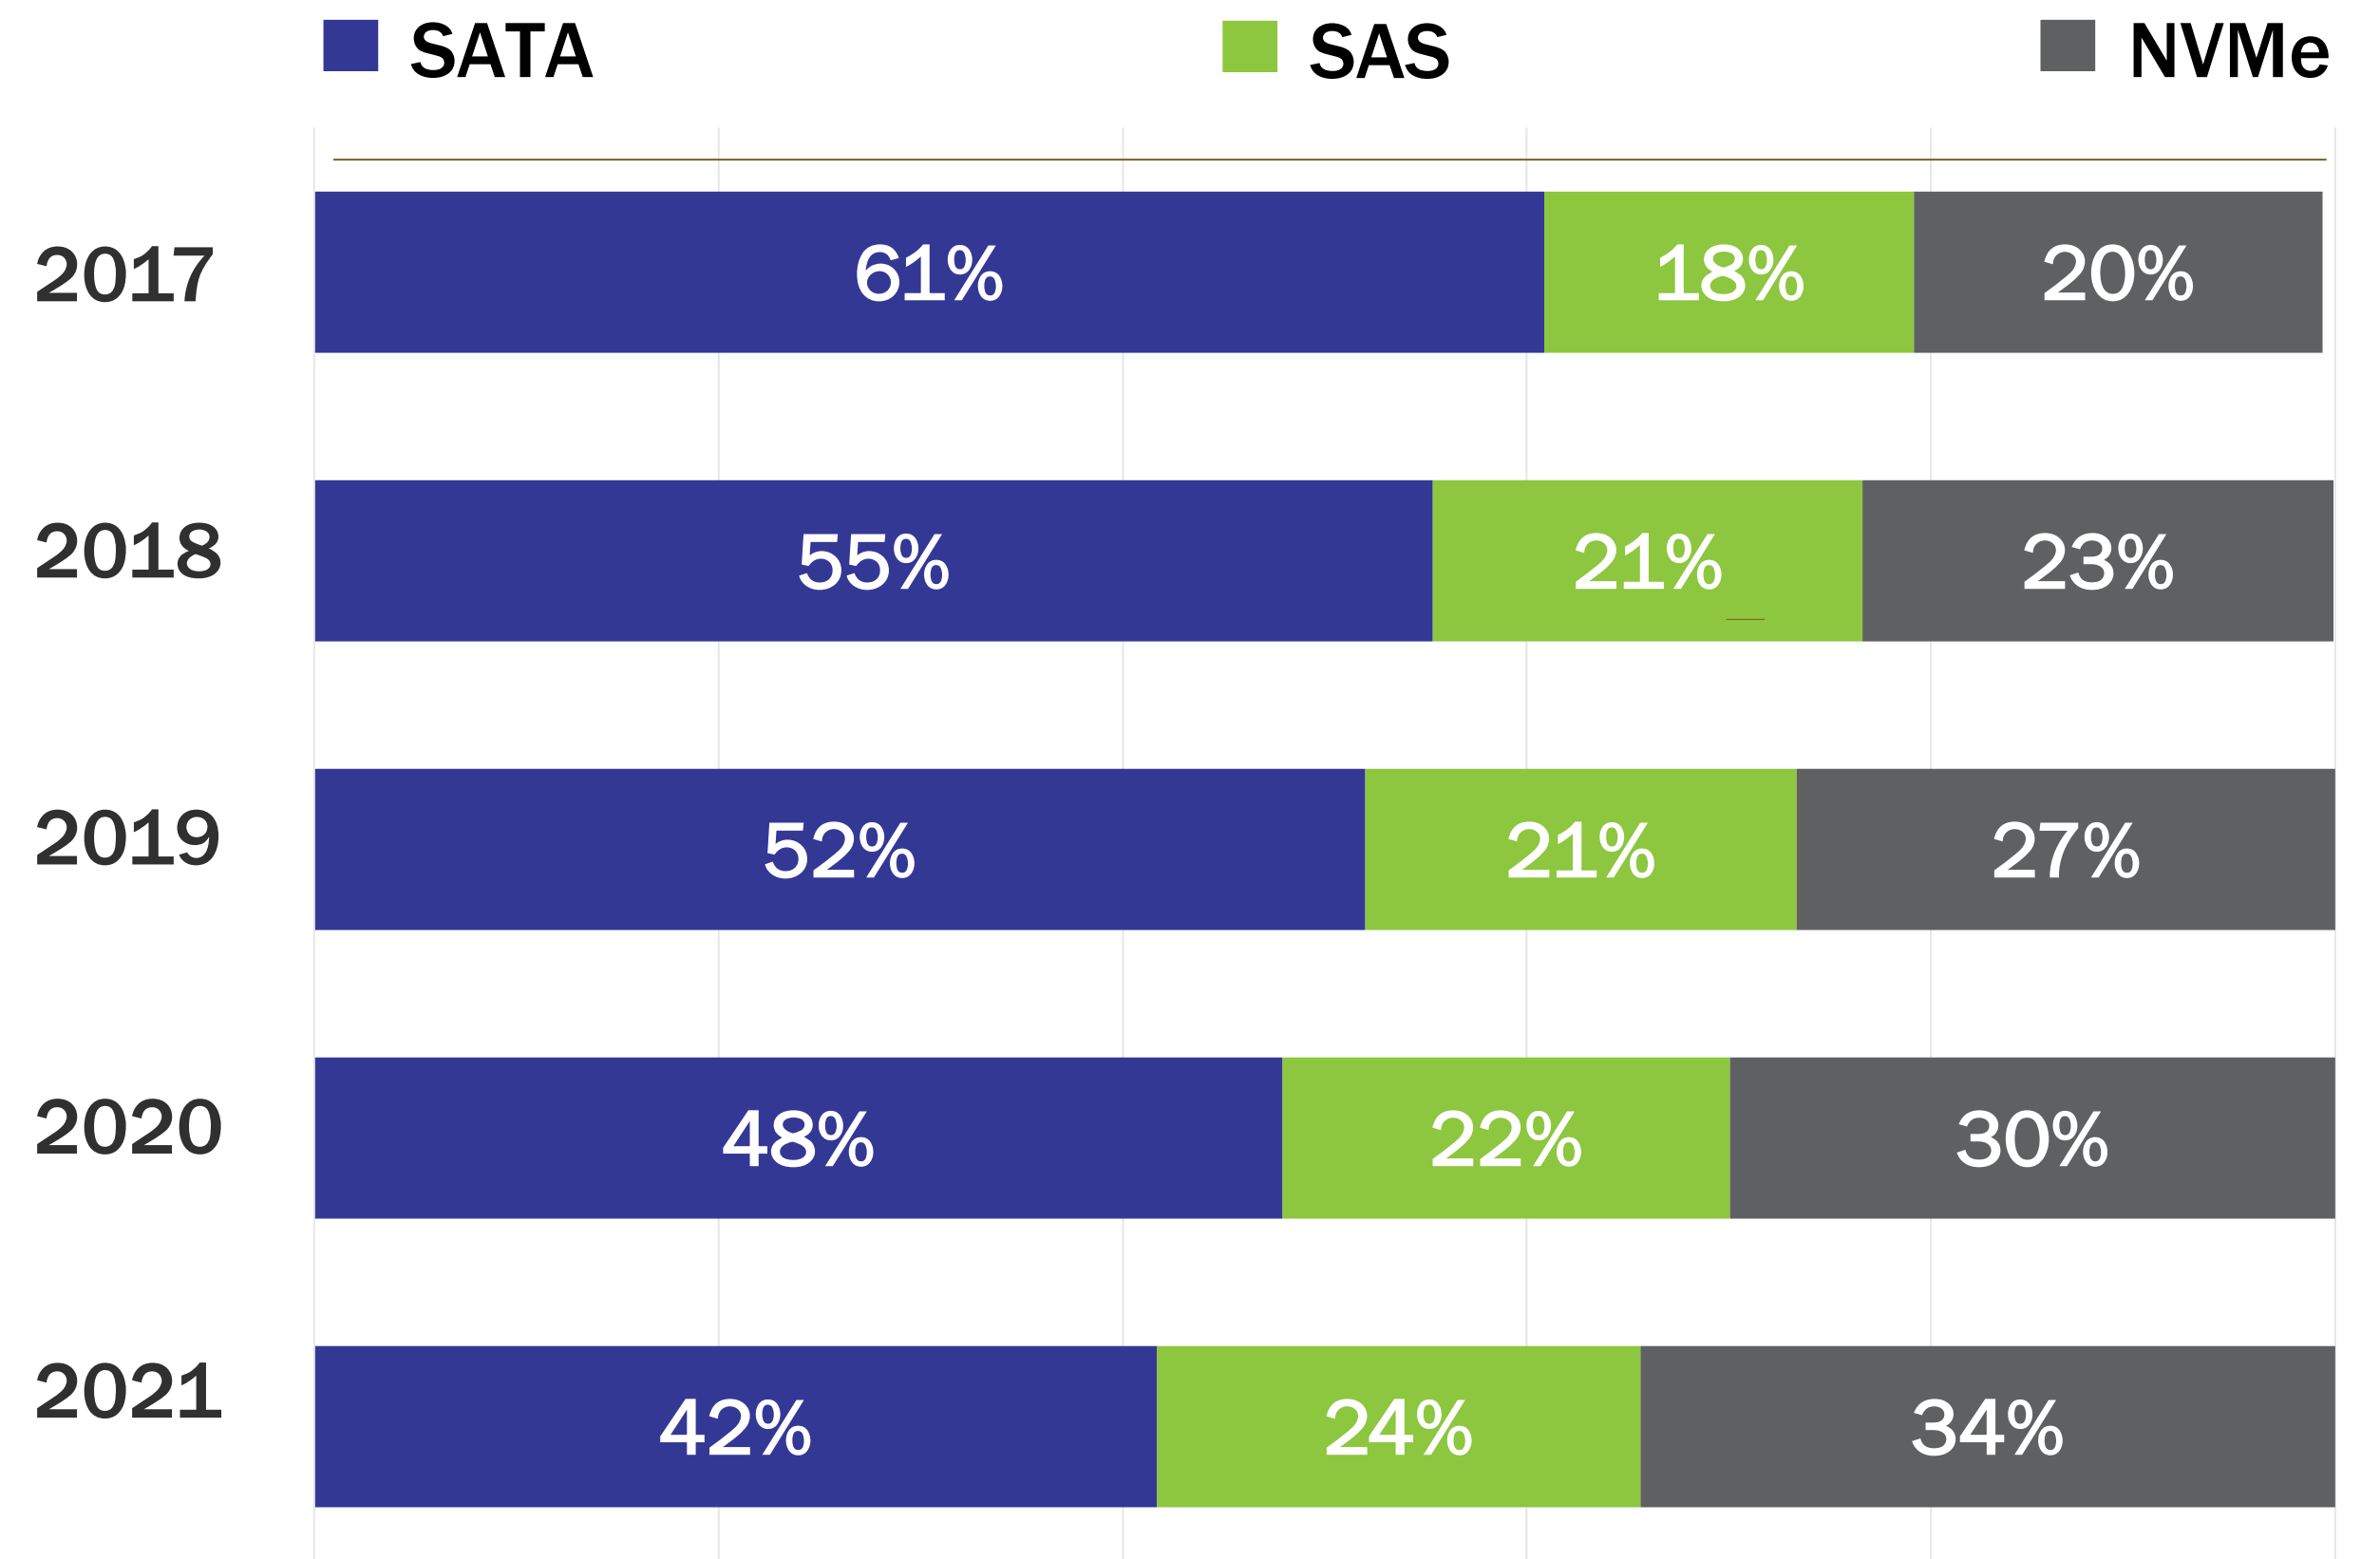
Tri-mode / Universal Backplane



Why Tri-mode?

- Mixed market forecast
- Bare metal as service
 - Avoid Reconfiguration
 - Hardware RAID
 - Hot-swap support
- Rackspace Perspective
 - Complex infrastructure
 - Increasingly mixed deployments (HCI)
 - Price Volatility

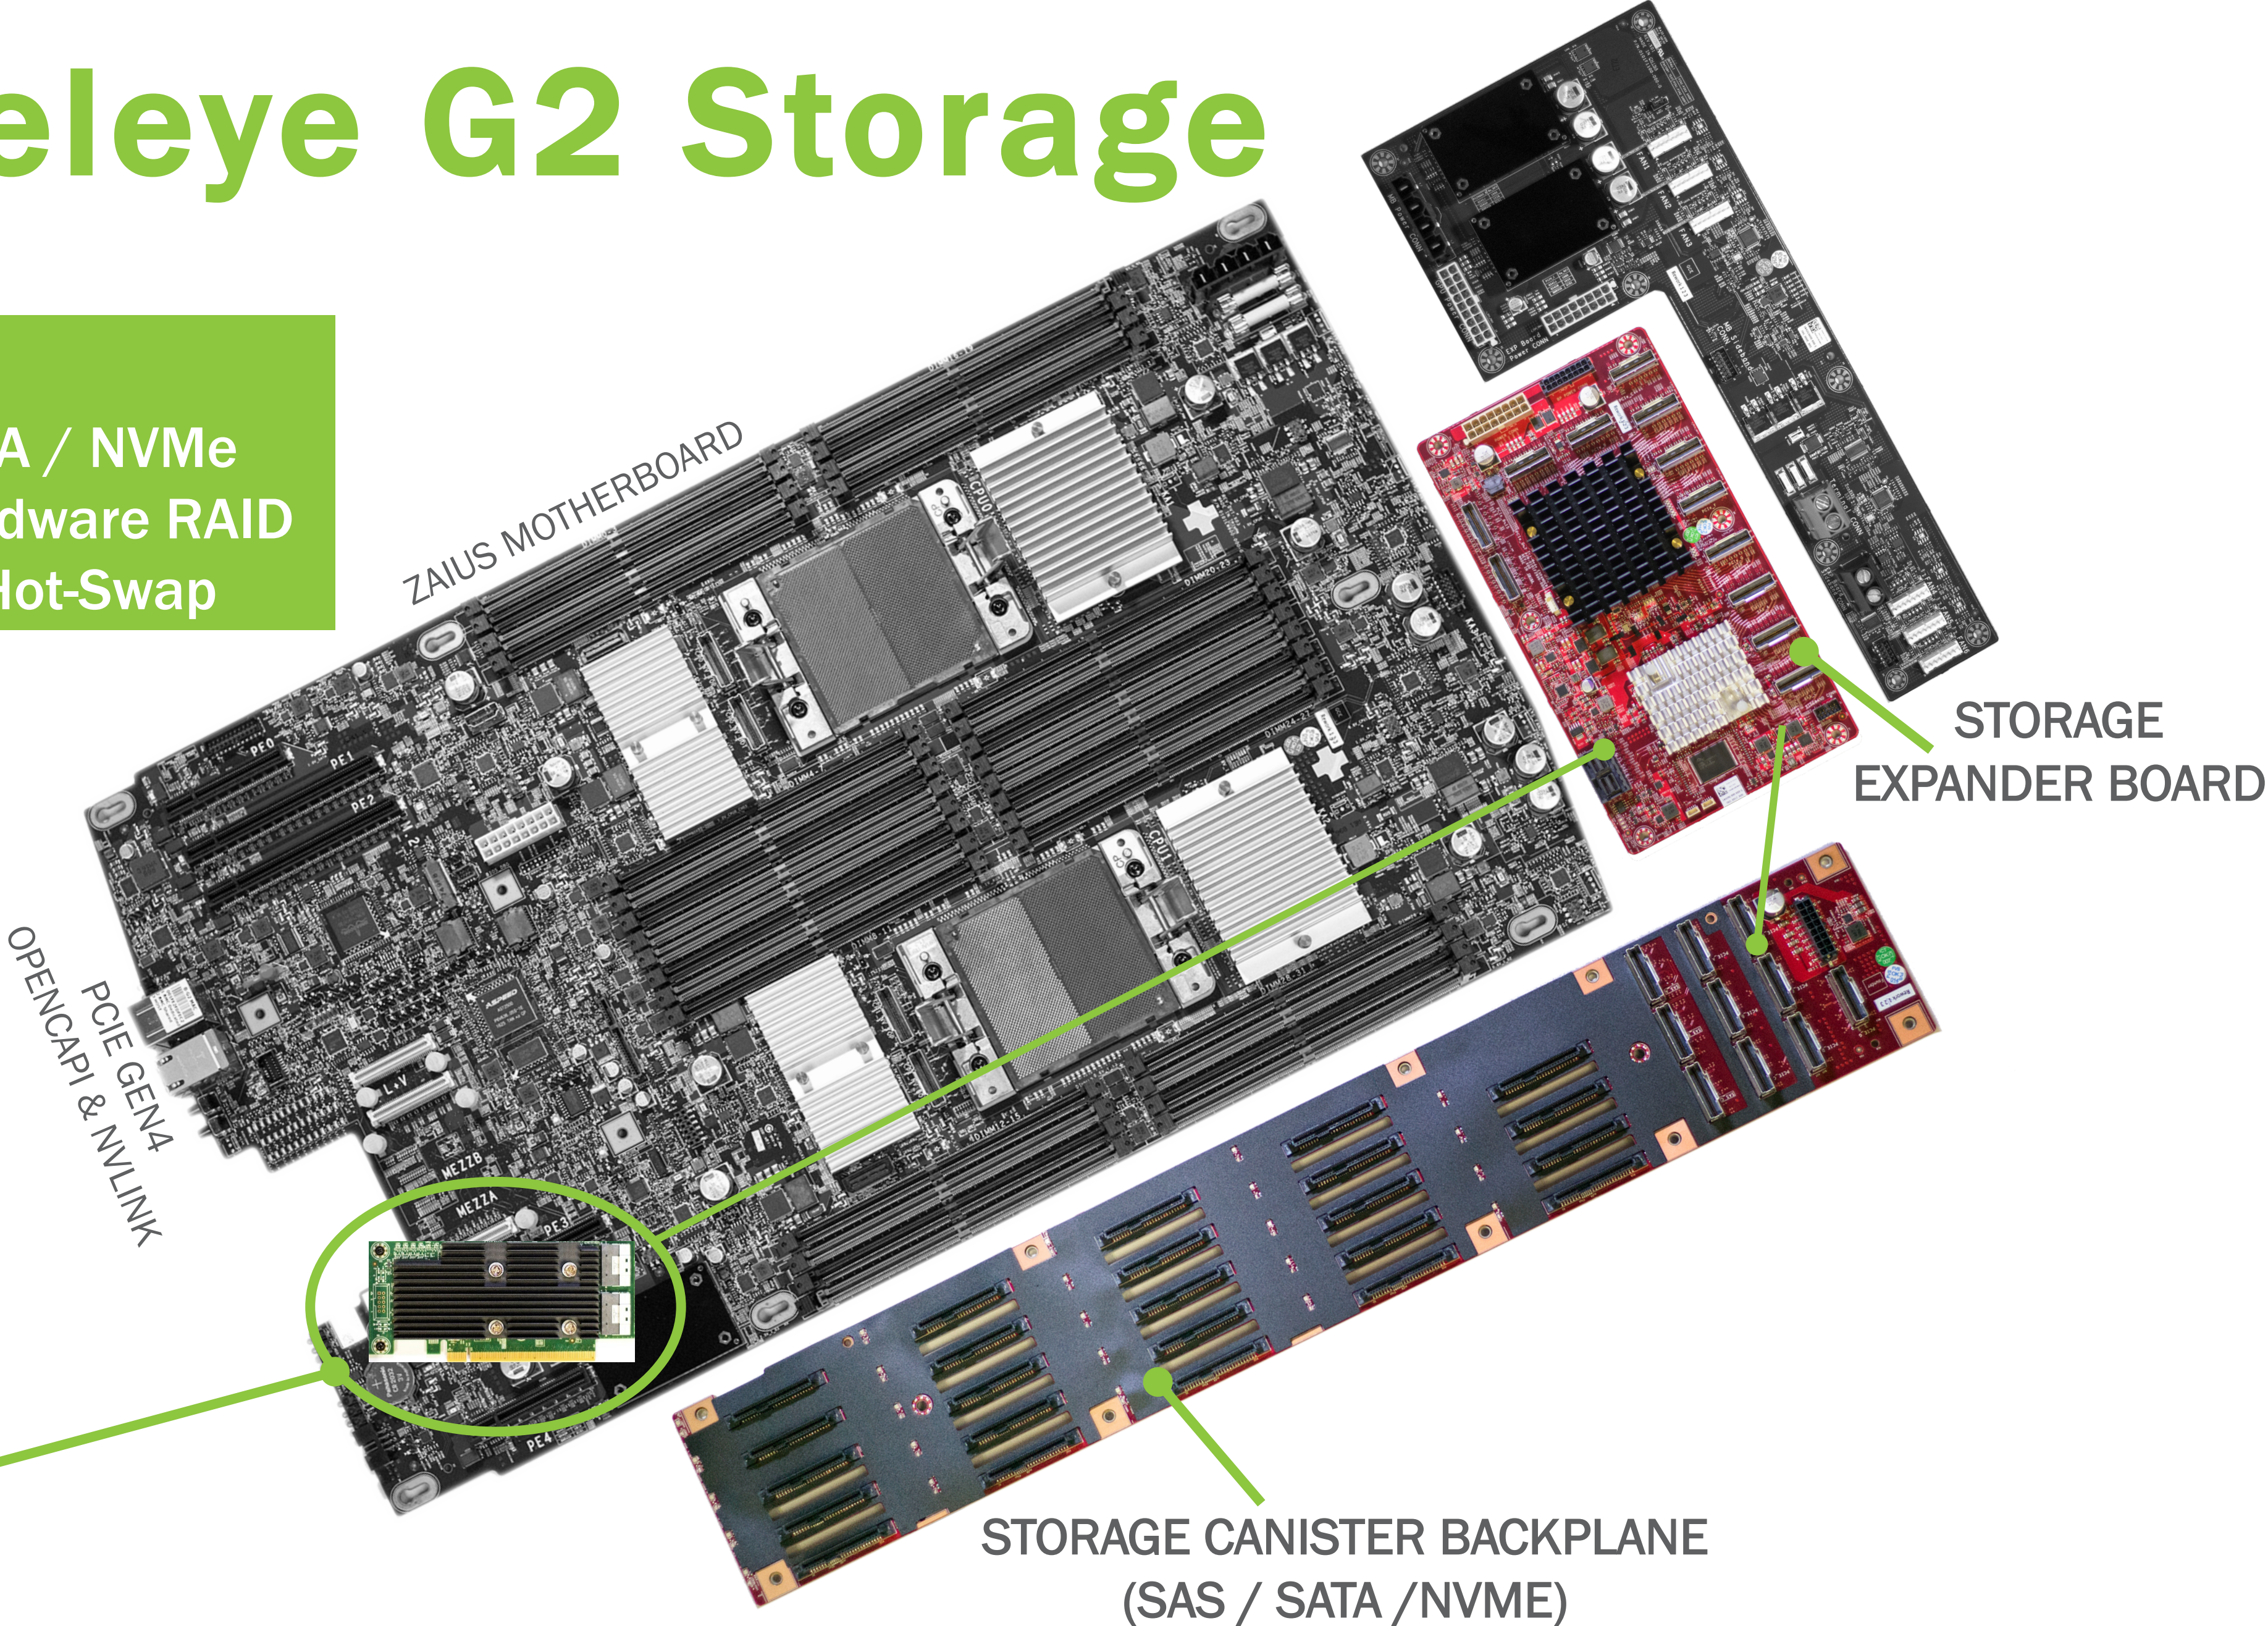
Percentage of Total Enterprise SSD consumption
Forecast based on number of units



Barreleye G2 Storage

Tri-mode


- 24 Drives
- SAS / SATA / NVMe
- NVMe Hardware RAID
- Tri-mode Hot-Swap



Tri-mode Storage

Demo 1

Hot-Swap SATA with NVMe



Tri-mode Storage

Demo 2

Building hardware RAID on NVMe with Broadcom 9460-16i



```
root@ubuntu:/home/fxn# ./storcli64 /c0 add vd type=raid0 drives=65:3,7
Controller = 0
Status = Success
Description = Add VD Succeeded
```

```
root@ubuntu:/home/fxn# ./storcli64 /c0 add vd type=raid0 drives=1,12
Controller = 0
Status = Success
Description = Add VD Succeeded
```

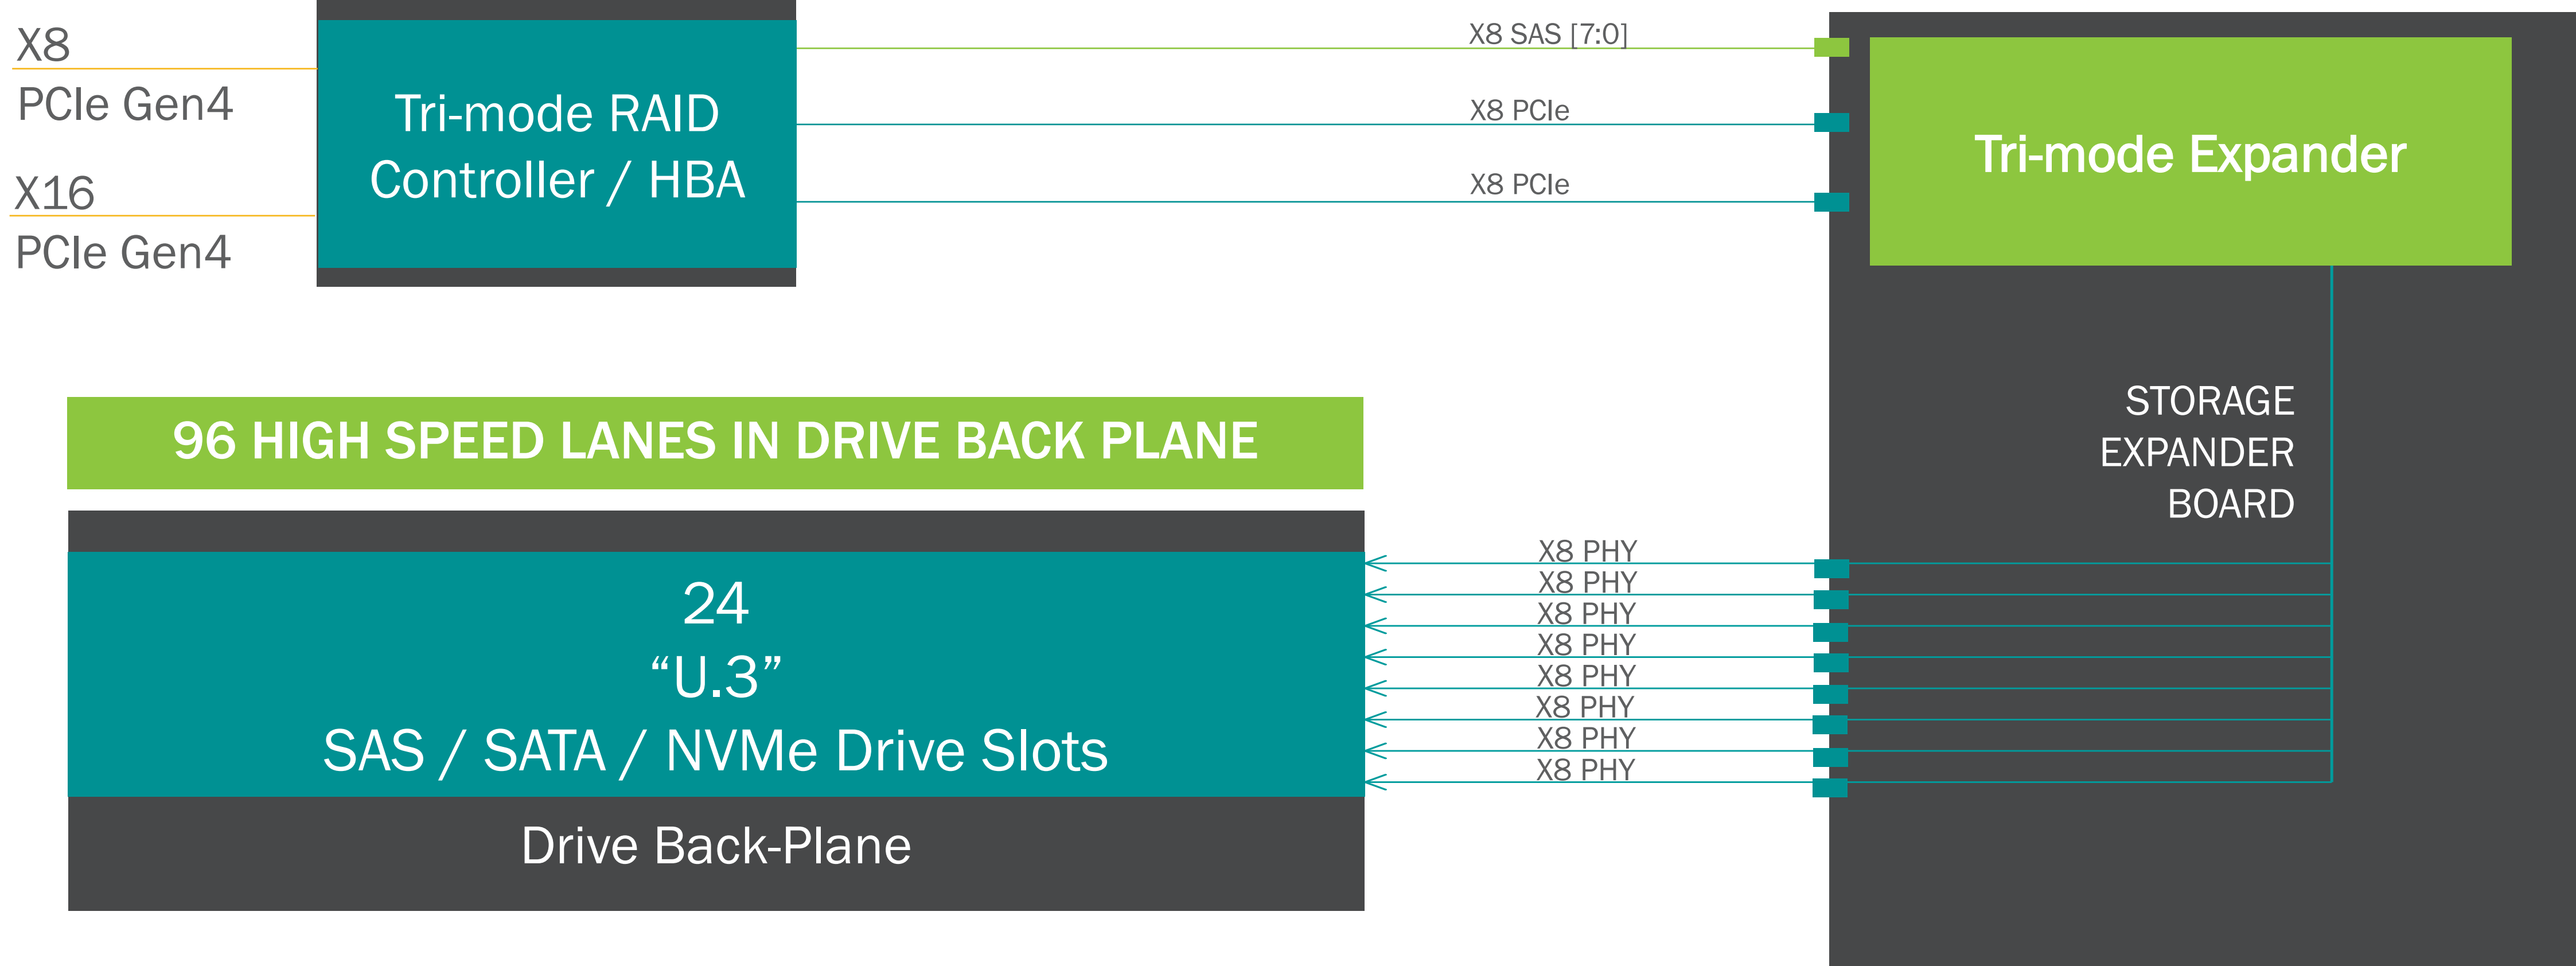
SATA HW
RAID

EID:Sl	DID	State	DG	Size	Intf	Med	SED	PI	SeSz	Model	Sp
65:3	14	Onln	1	237.968 GB	SATA	SSD	N	N	512B	TOSHIBA THNSNH256GCST	U
65:7	15	Onln	1	237.968 GB	SATA	SSD	N	N	512B	TOSHIBA THNSNH256GCST	U
65:9	13	UGood	-	278.875 GB	SAS	HDD	N	N	512B	ST9300653SS	U
:1	1	Onln	0	1.454 TB	PCIe	SSD	N	N	512B	Micron_9200_MTFD	U
:12	12	Onln	0	1.454 TB	PCIe	SSD	N	N	512B	Micron_9200_MTFD	U

NVMe
RAID

Ideal Tri-mode Storage

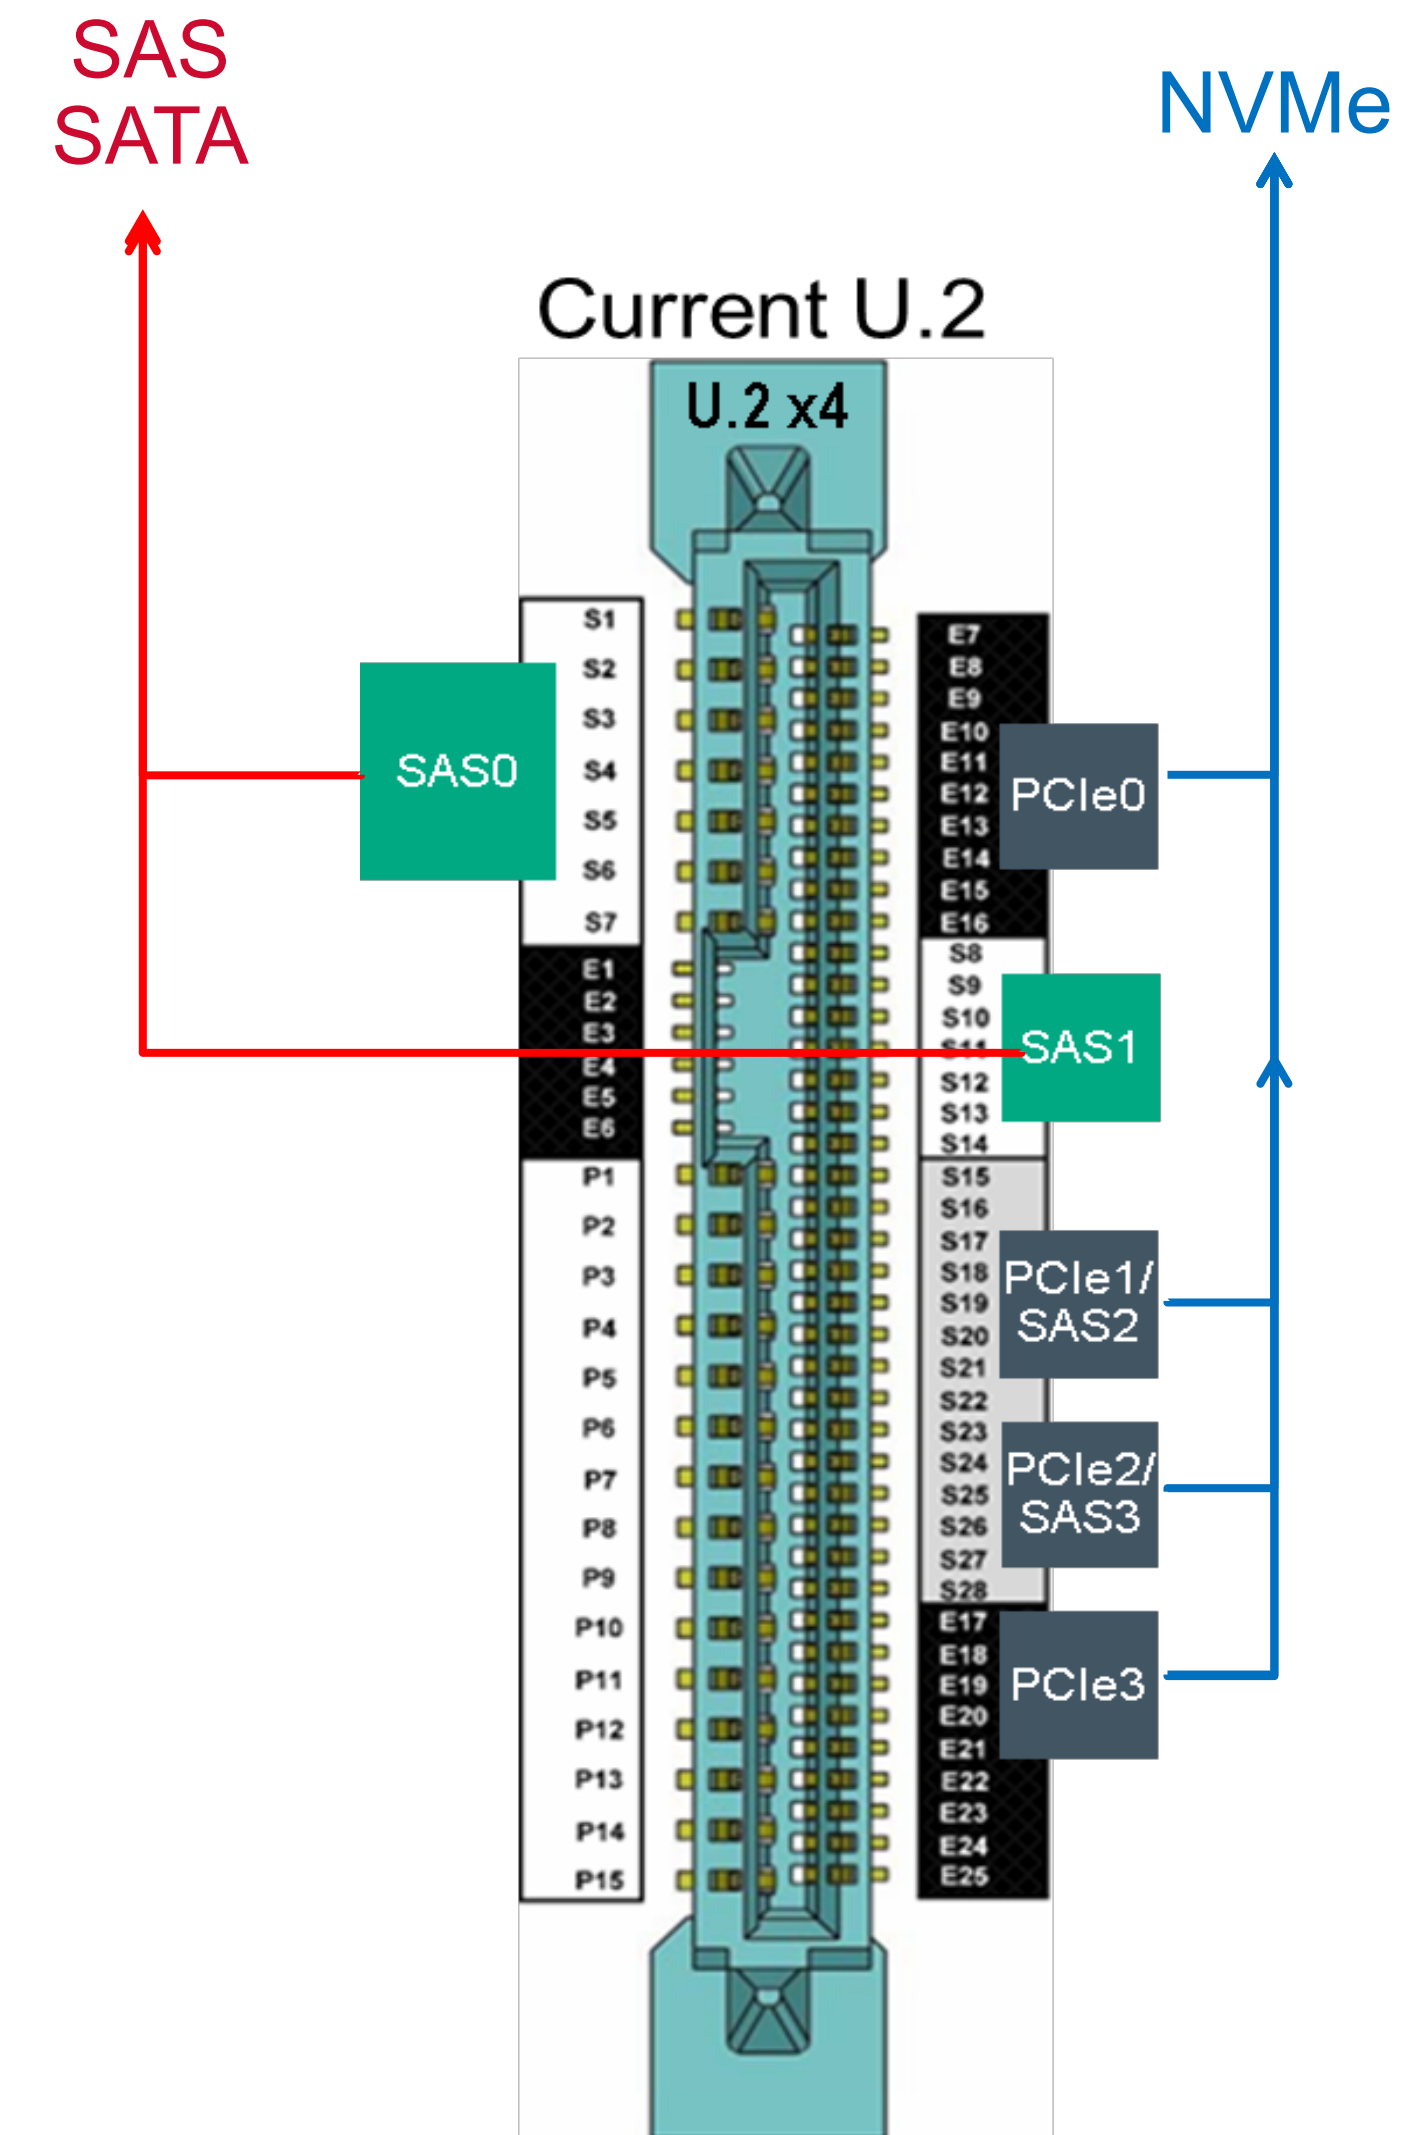
Reference Architecture



Tri-Mode Storage:

Requirements of Specification / Samples

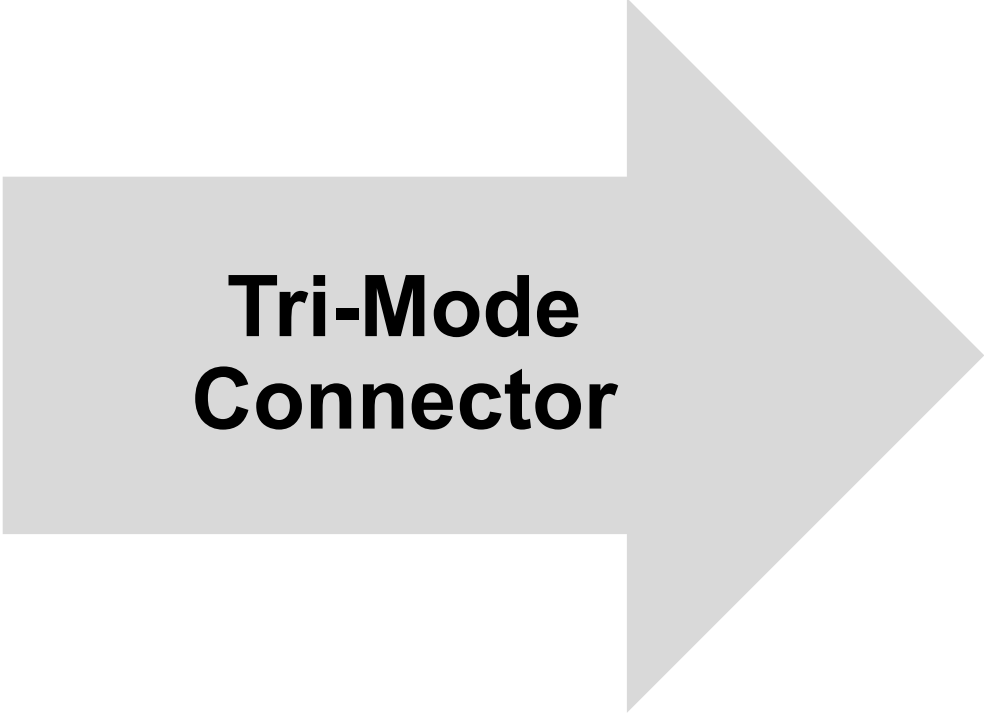
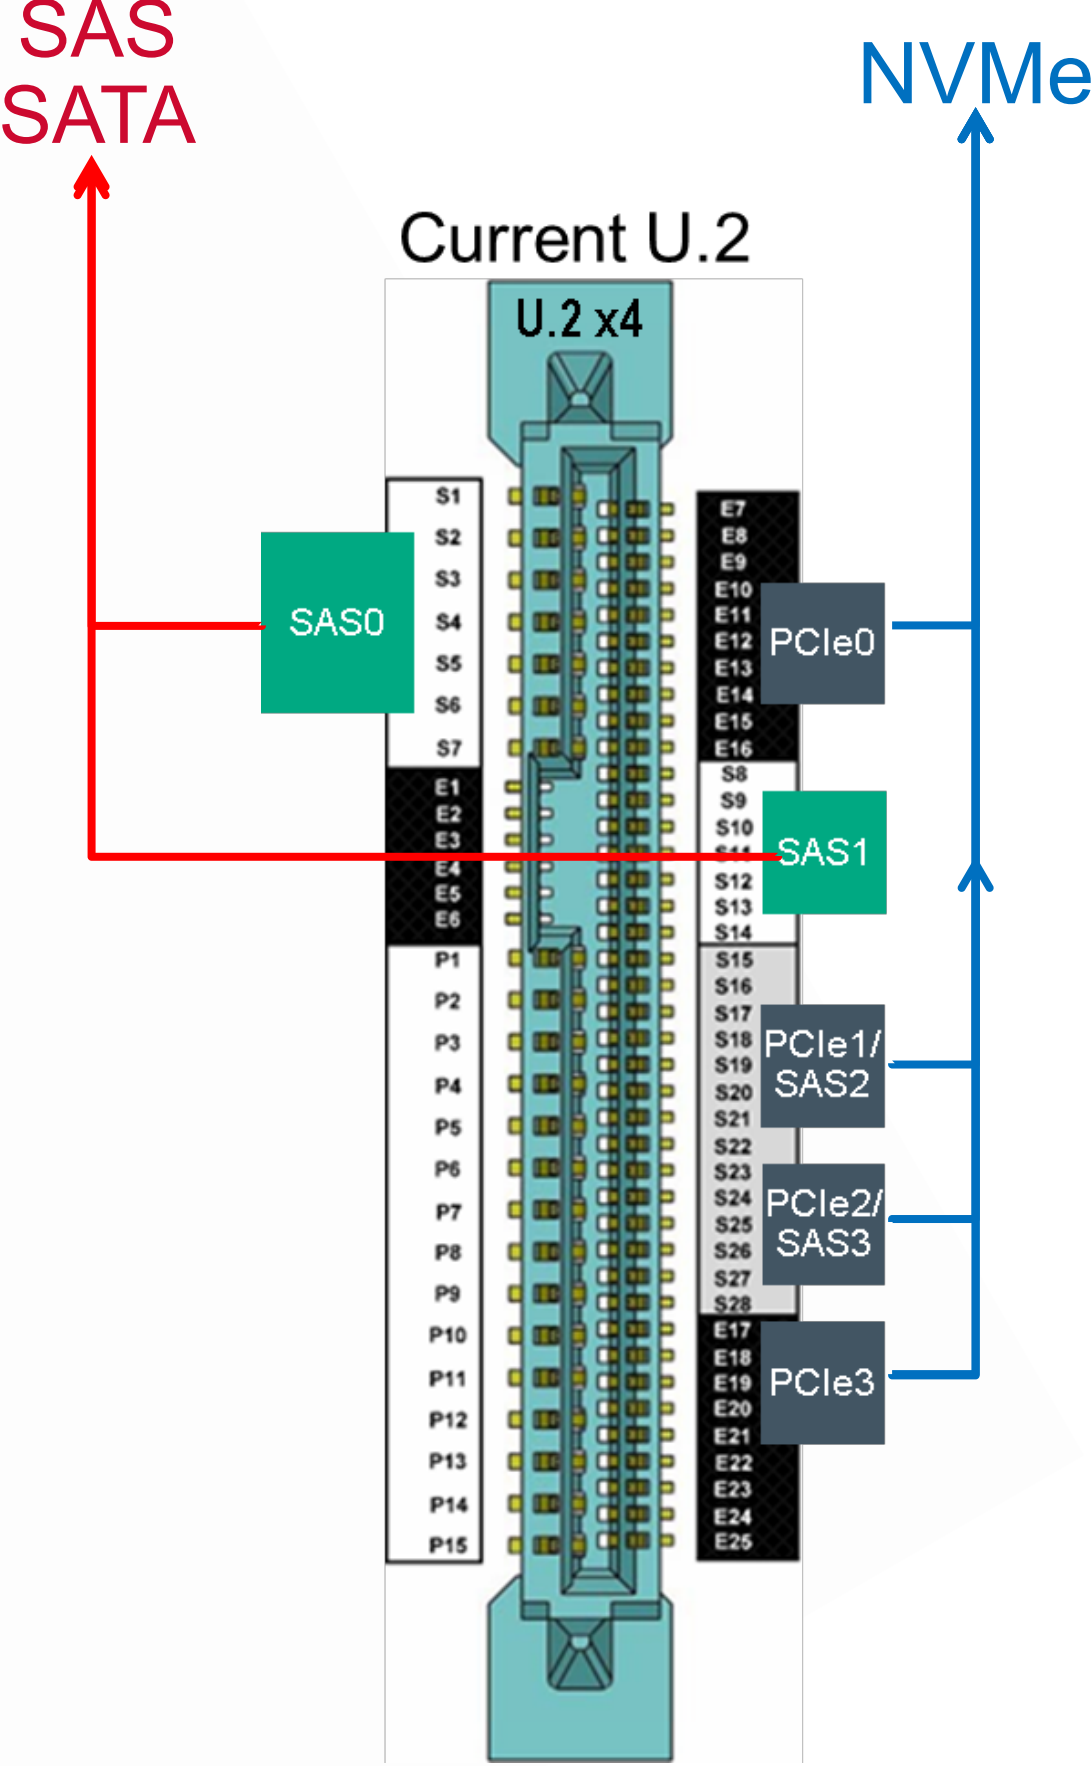
- One Backplane
 - One connector
 - Less high-speed lanes to backplane
- One Mid-plane
 - Tri-mode Expander
- One HBA / RAID Controller



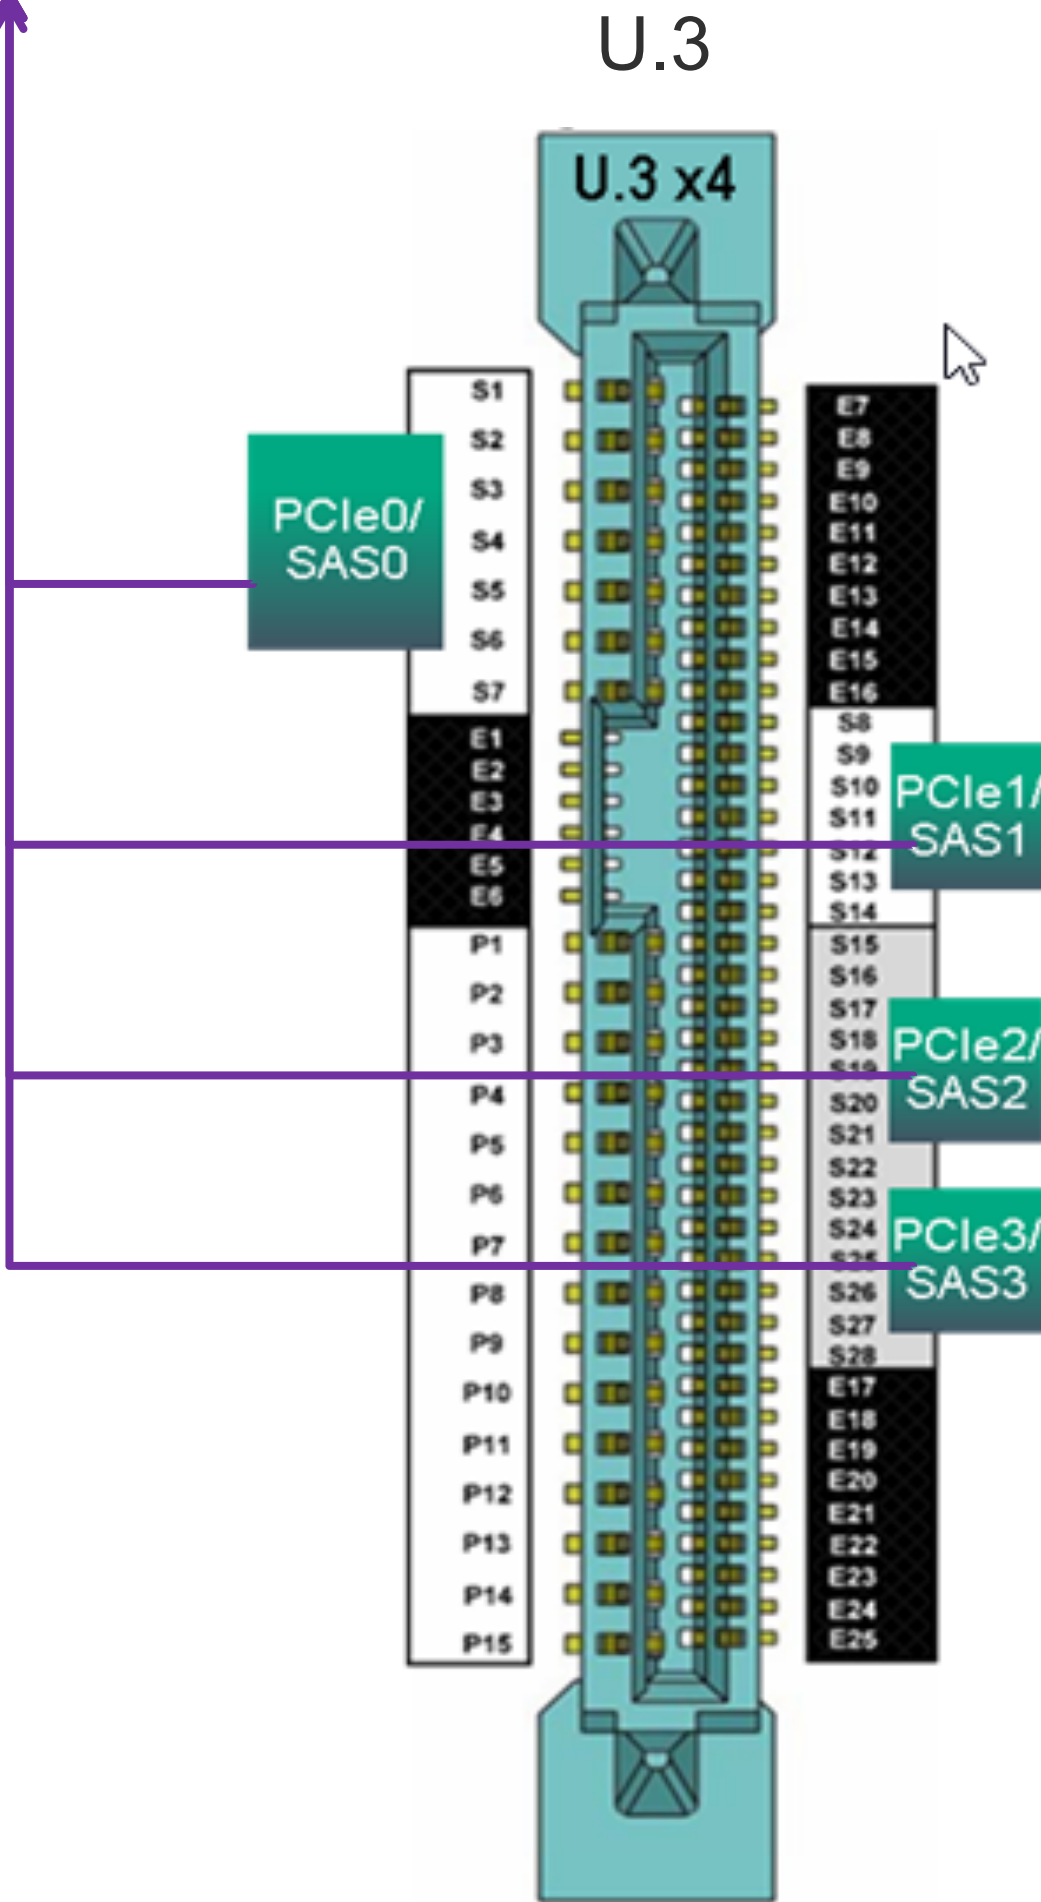
Tri-Mode Storage:

Requirements of Specification / Samples

- One Backplane
 - Single connector
 - Less high-speed lanes to backplane
 - Drives should change form-factor



SAS
SATA
NVMe

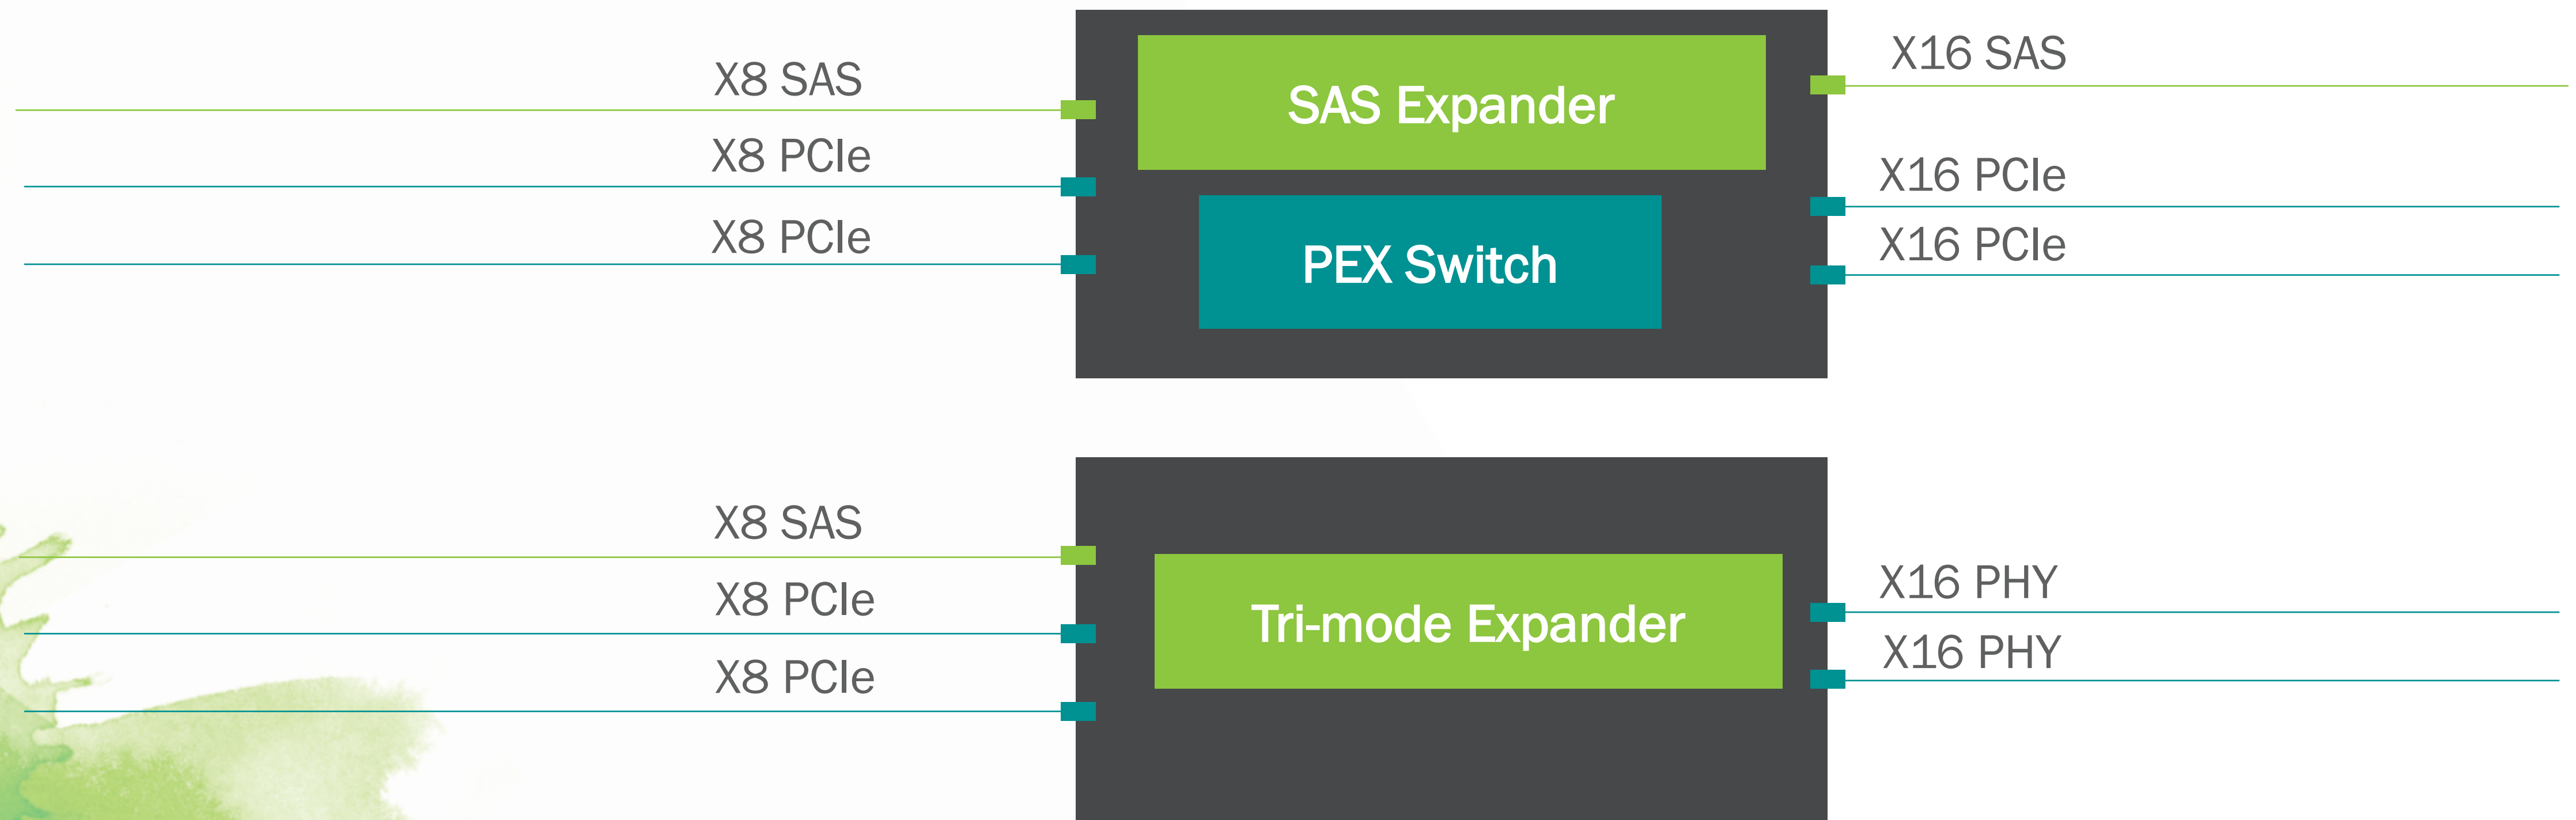


Tri-Mode Storage:

Requirements of Specification / Samples

One Mid-plane Solution



- Tri-mode Expander



Tri-Mode Storage:

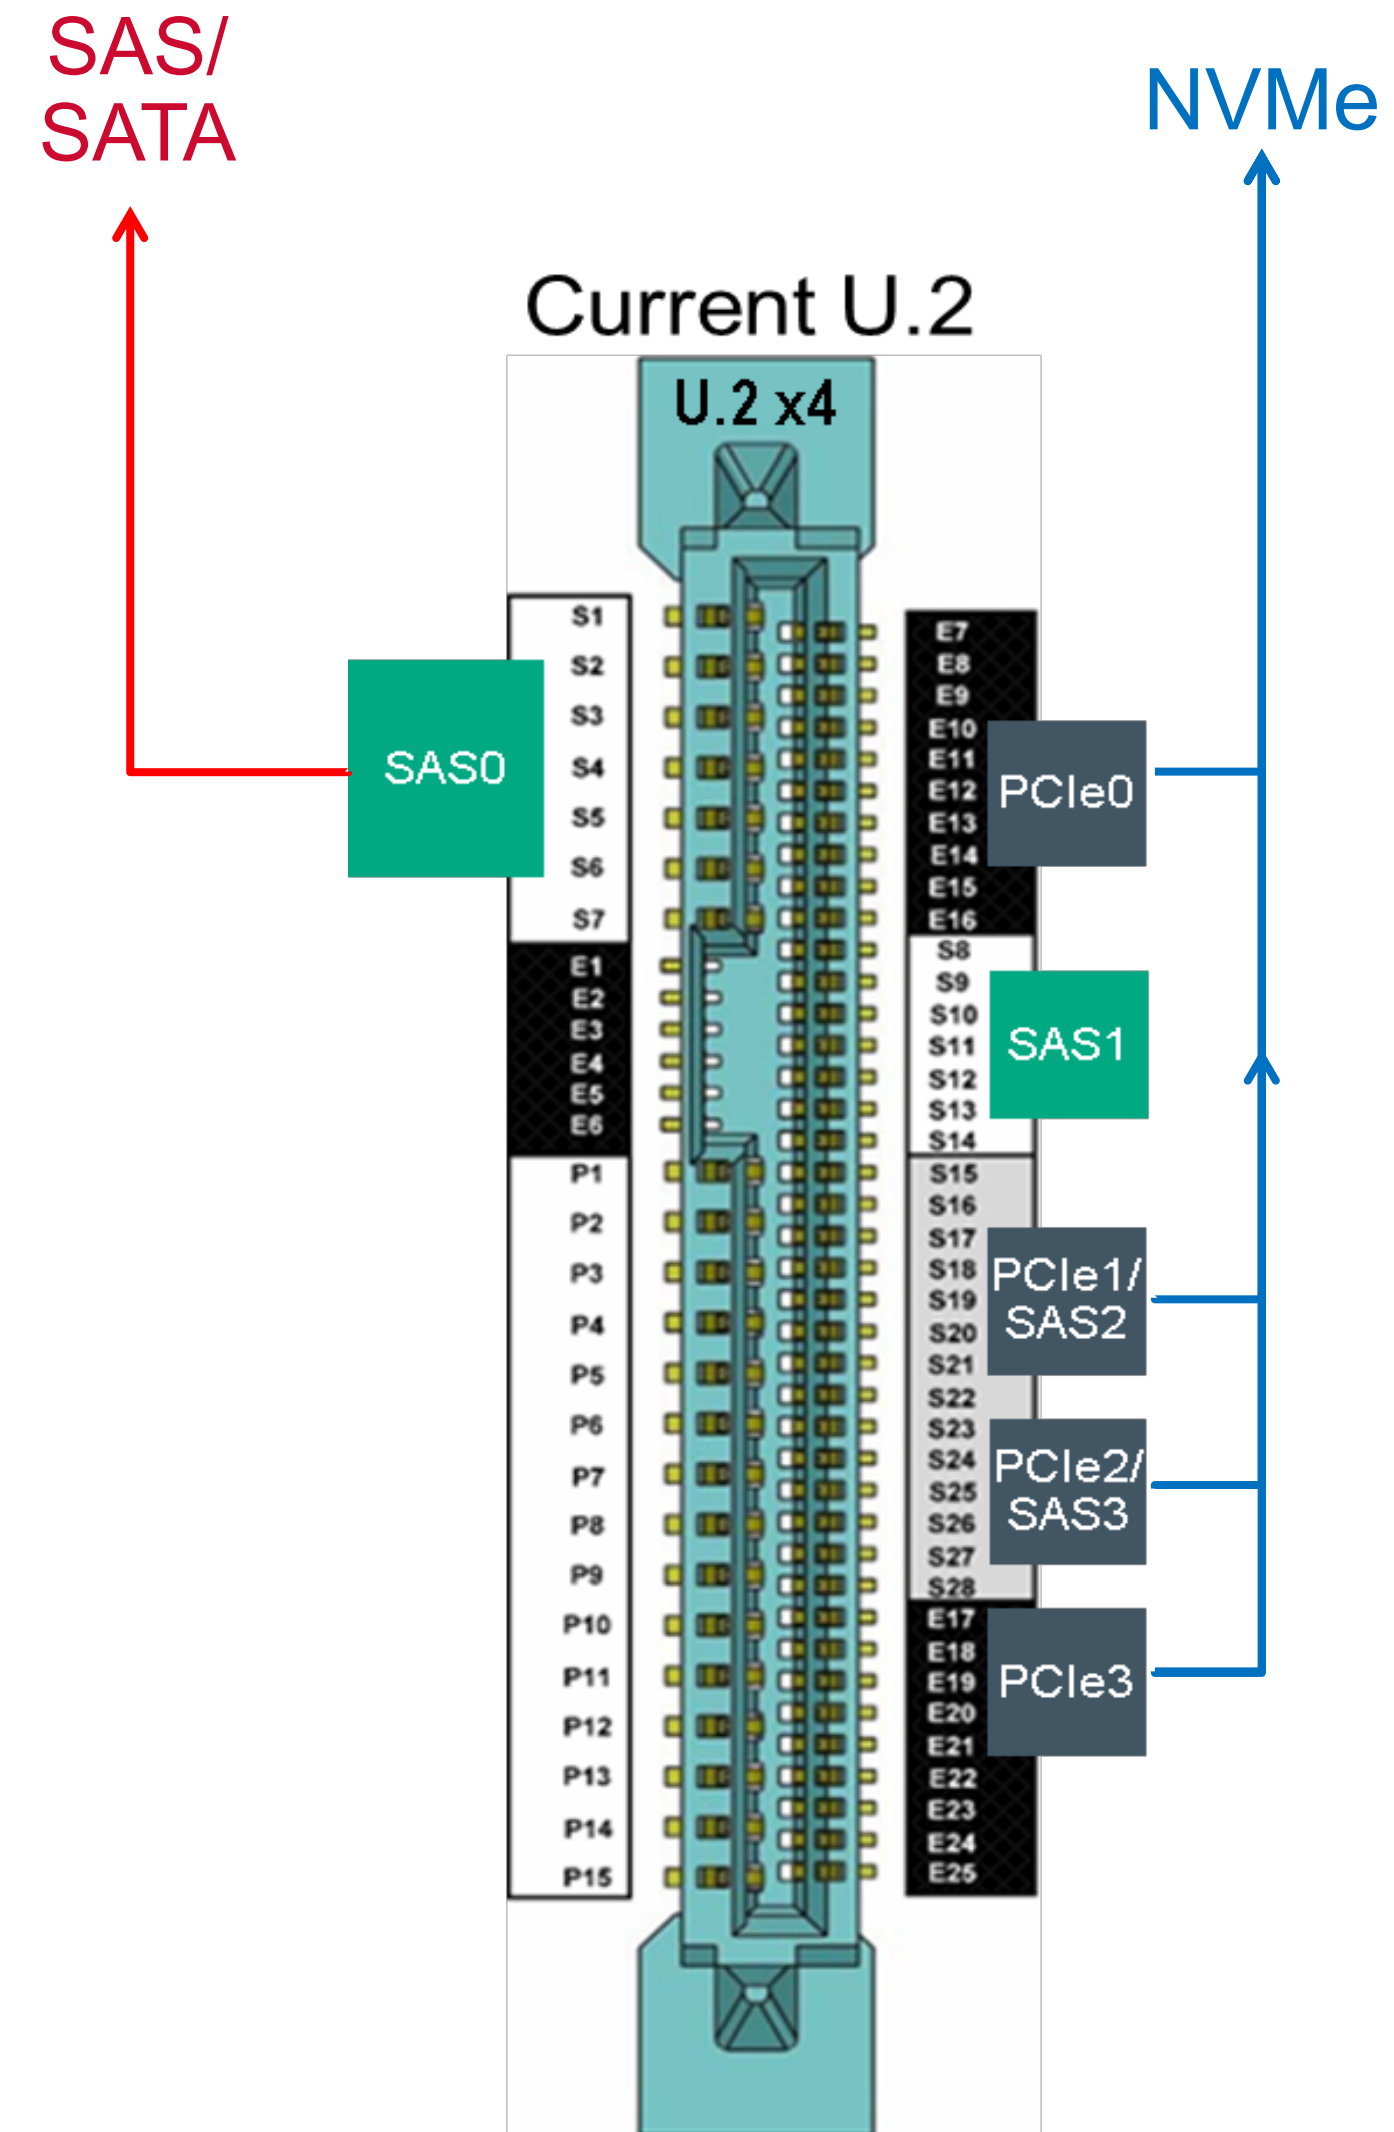
Requirements of Specification / Samples

- One HBA / RAID Controller



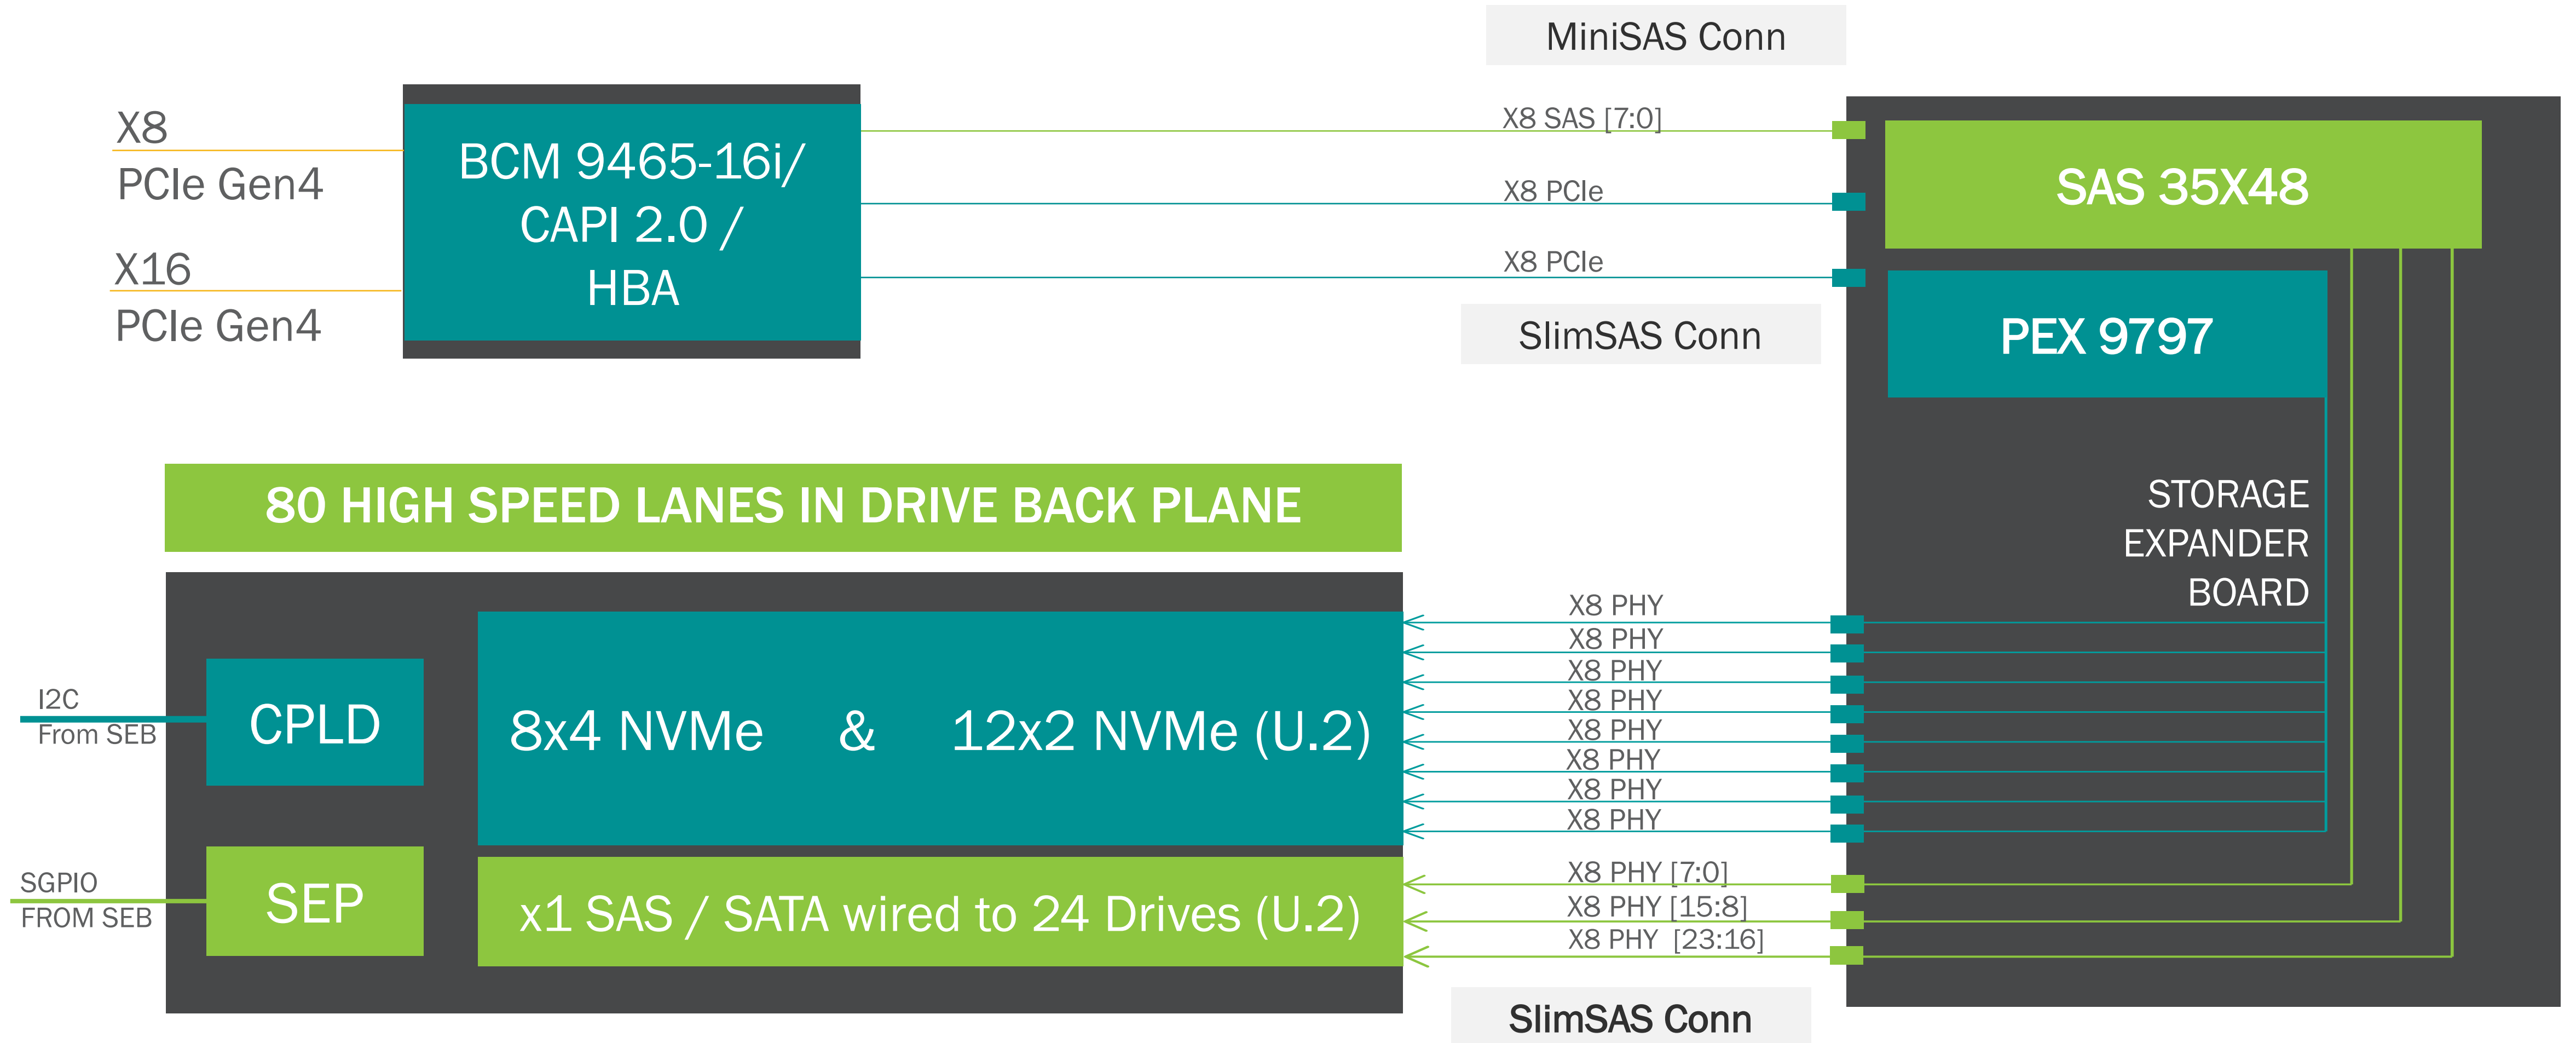
Double Plumb U.2

Single lane SAS / SATA , Dual lane NVMe





Barreleye G2 Tri-mode Storage

Actual Implementation






Barreleye G2 Tri-mode Storage

Actual Implementation: PCIe Gen3 U.2 Backplane




Barreleye G2 Tri-mode Storage

Using Separate Controllers for SAS / SATA and NVMe




Barreleye G2 Tri-mode Storage

Using Tri-mode Controller for SAS / SATA and NVMe



Barreleye G2 Tri-mode Storage

Using Separate Controllers for SAS / SATA and NVMe



Specification Update

Universal Backplane Implementation Guidance

- SFF-TA-1001 (U.3): Common drive bay definition
 - Specification ratified October 2017
 - Defines a universal backplane definition for SAS, SATA, and NVMe drives
 - Simplifies multi-protocol backplane design and lowers cost
 - U.3 Drives backwards compatible to U.2 bays
- SFF-TA-1005 (UBM): Universal Backplane Management
 - Backplane management framework
 - LEDs and much more
 - Unifies capabilities of various backplane management schemes
 - SGPIO (SFF-8485), 2Wire SES, and I2C PCA9555
 - Provides a method to manage and control SAS/SATA/NVMe backplanes
 - Resolves x2 and x1 NVMe drive support challenges



PCIe Gen4 Switch
Broadcom



NVMe NoLoad
Eideticom



Flash as Commodity
Burlywood



Flash Storage Accelerator
Molex



Tri-mode Storage Enhancements

PCIe Gen3 U.2 Universal Backplane

Possible Adoption & Upgrades

Adopting PCIe Gen3 U.2

Tri-mode Backplane

- SATA support by adding 1 or 2 lanes per U.2 NVMe connector

Upgrading to PCIe Gen4 U.3 Tri-mode Backplane

- Amphenol
- PCIe Gen4 NVMe (16 GT/s)
- SAS 4.0 (22.5GT/s)
- Tri-mode expander
- U.3 Connectors

SAMSUNG


Micron

BROADCOM

MEMORY/
STORAGE

POWER9

IBM



ACCELERATORS

molex
ALPHA DATA
EIDETICOM

ingrasys

MOTHERBOARD, PCBA, CHASSIS, INTEGRATION

Design Package

Available on OCP Server Wiki and GitHub

Enhancements Coming Throughout 2018

GitHub

<https://github.com/opencomputeproject/zaius-barreleye-g2> (don't forget to install Git LFS)

OCP

http://www.opencompute.org/wiki/Server/Working#Open_Rack



THANK
YOU



OCP SUMMIT

Archive Material



OPEN
Compute Project

U.3 vs U.2

PCIe Gen3 U.2 Universal Backplane Possible Adoption & Upgrades




FIGURE 3-1 PORT USAGE OVERVIEW