STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

Why Today's Cloud and Hyperscale Topologies Use SAS:

A Look Into Meta's Grand Canyon

Rick Kutcipal, Product Planning – Broadcom
Board of Directors – SCSI Trade Association

Madhavan Ravi, Hardware Systems Engineer - Meta

Why Modern Hyperscale Architecture Use SAS

24G SAS Advancements for Hyperscale Environments

Investments in standards & products to enable hyperscale architectures

Capacity optimized features

Specialized performance solutions

Power considerations

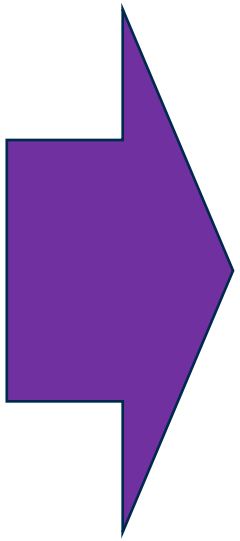

How Emerging Architectures are Leveraging 24G SAS

Scalability

Reliability

Flexibility

Cost

Hyperscale Innovations in SAS

Shingled Magnetic Recording (SMR) – Improvements in arial density (capacity)

Next Generation SMR – Flexibility in SKU management

Repurposing Depopulation (DePop) – Data Center management

Multi-Actuator – Addressing the performance per capacity issue

Command Duration Limits – Improving performance without sacrificing latency

Importance of Scalability in Hyperscale Architectures

Data Growth

Storage Management Agility

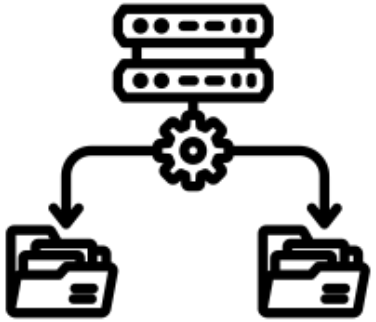

Service Level Agreements

Cost Optimizations

The Scalability of SAS

- SAS is a connection-based Storage protocol and is inherently scalable
- Supports practical topologies up to 2,000 devices
- Devices can be added or removed
  - Individually or in groups (enclosures)
  - Without interrupting IO traffic
- Supports active connections without extra components
  - Active Cu (10m – 20m)
  - Active optical (>300m)

Reliability

- SAS was originally developed to meet the needs of the Enterprise
  - Data availability
  - Data integrity
  - 24/7 operation
- Features like
  - Native support for redundant paths
  - Forward Error Correction (FEC)
  - End-to-end data protection
  - S.M.A.R.T

The Flexibility of SAS

- SAS protocol supports
  - SAS HDDs and SDDs
  - SATA HDDs and SDDs
- SATA NL HDDs have significant price advantage over comparable SSDs
- SATA Tunneling Protocol (STP) enables SATA end devices

CapEx – Media Cost

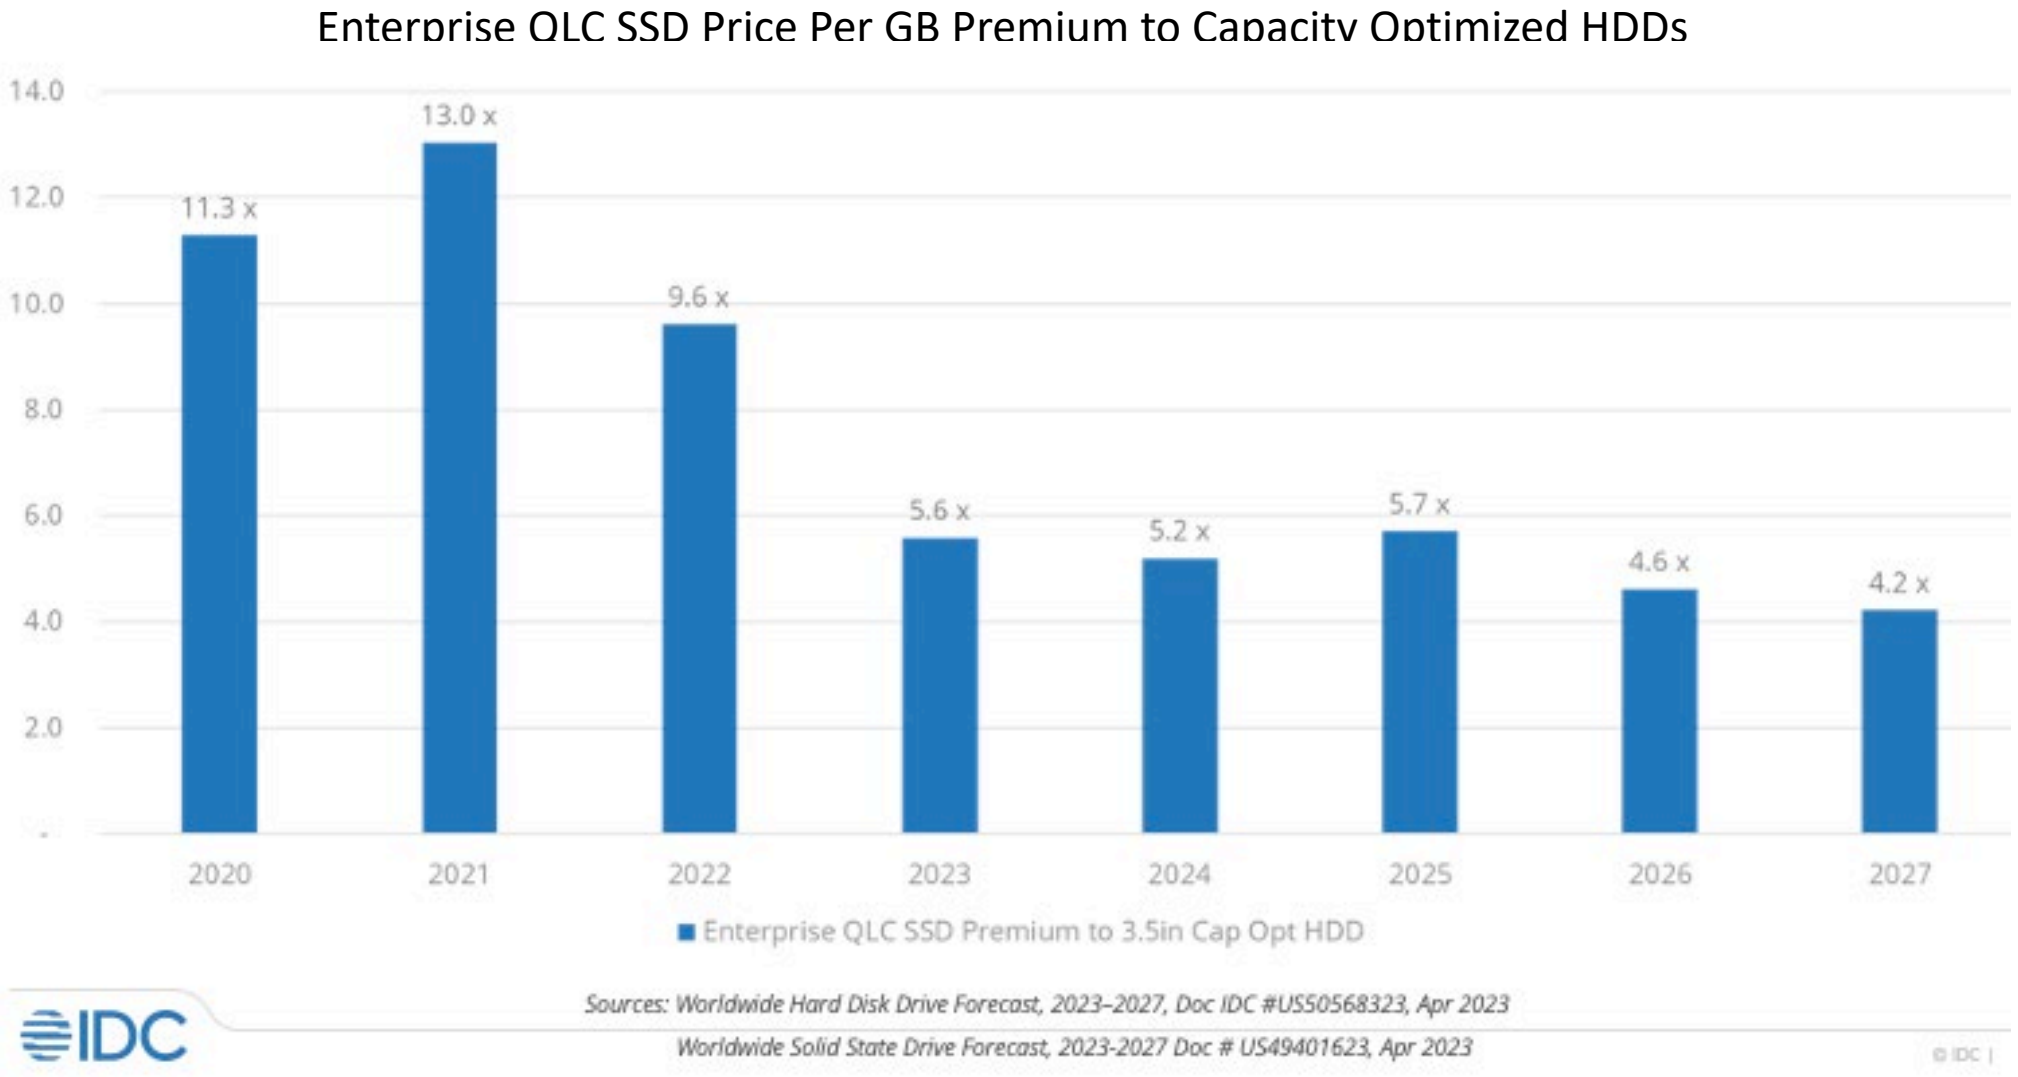

SDC 23

Enterprise Storage Capacity
by Technology, Exabytes

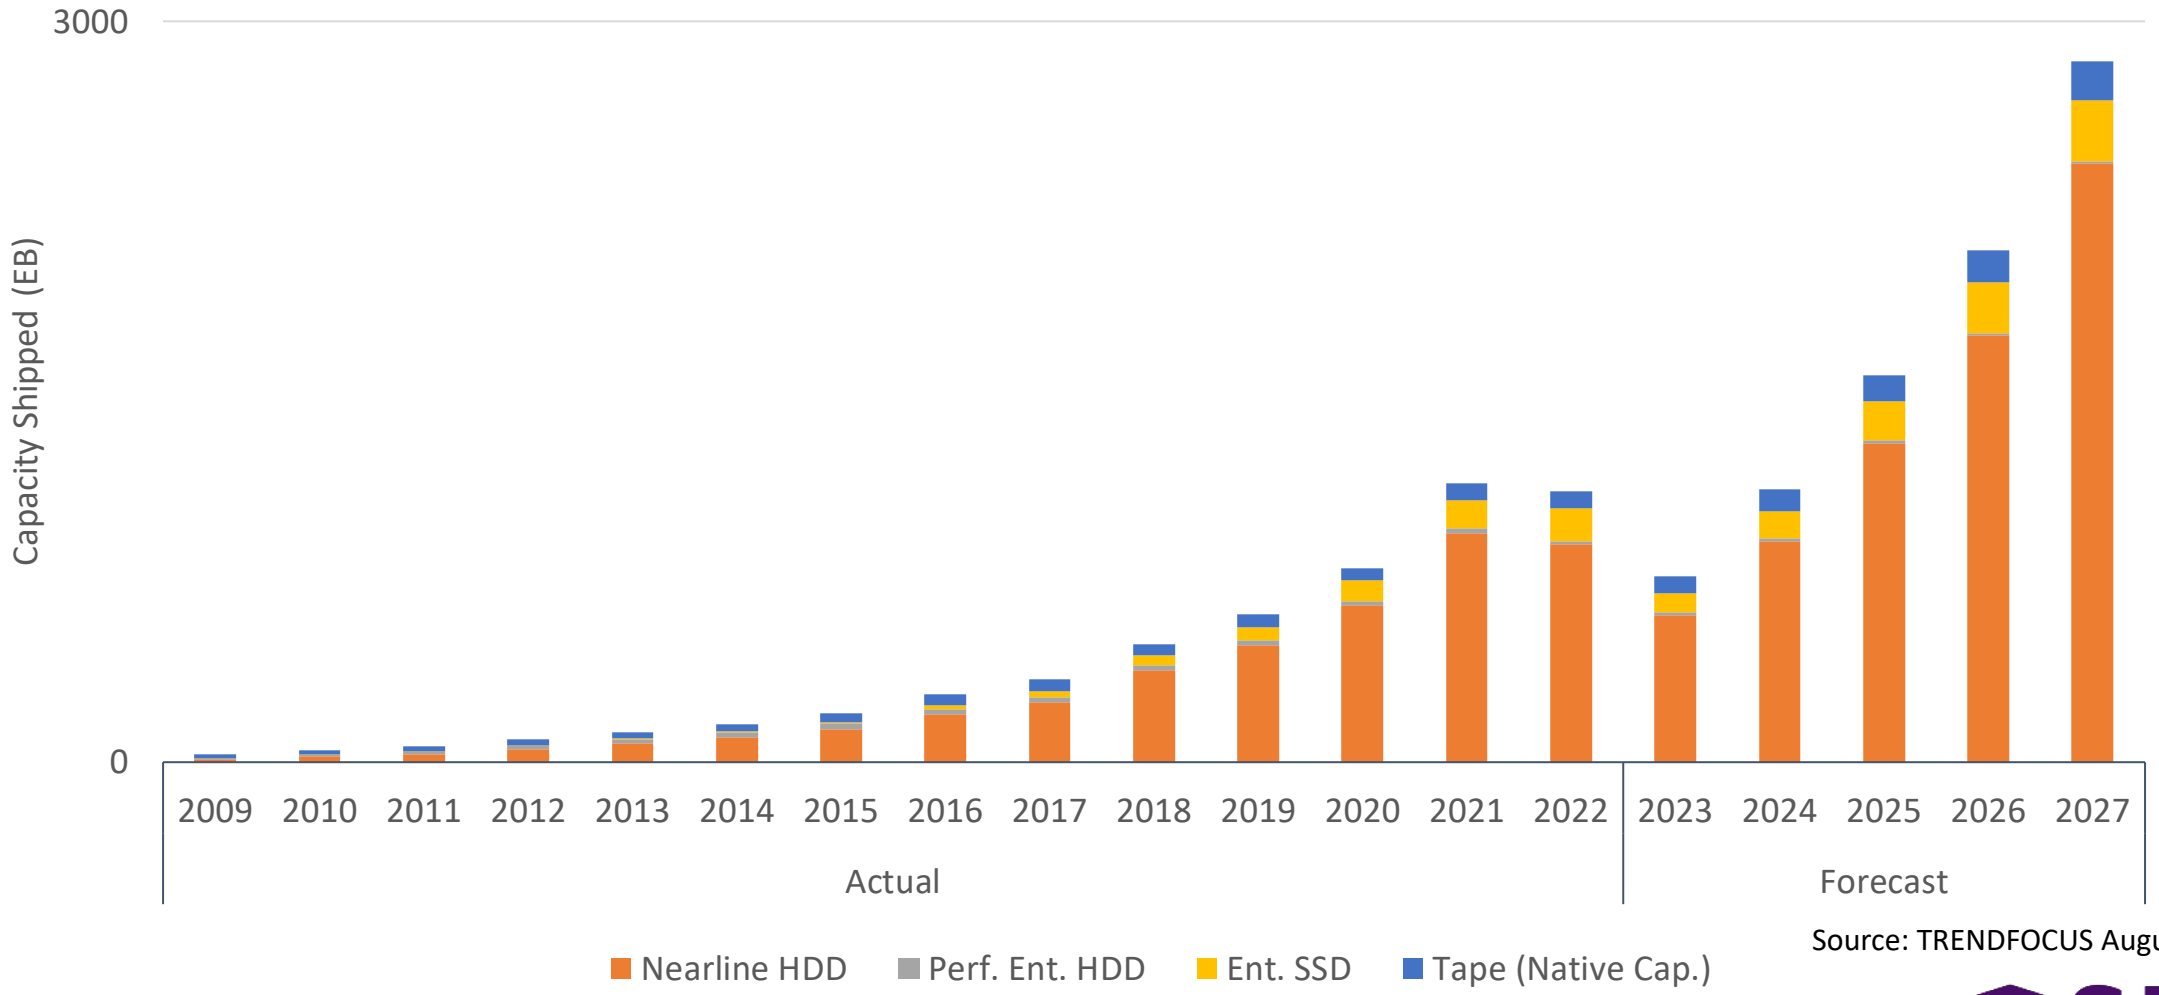

Source: TRENDFOCUS August 2023

OpEx

- The 2 largest contributors to OpEx is power costs and staffing

HDD / SSD Slot Level Power Comparison

	Micron 6500 – 30TB	Seagate Exos X22	Comments
Idle Power	5.0 W	5.7 W	14% SSD advantage[1]
Read Power	15.0 W	9.4 W	37% HDD advantage
Write Power	20.0 W	6.4 W	68% HDD advantage
Read Intensive Power	10.3 W	7.4 W	28% HDD advantage[2]
Write Intensive Power	12.3 W	6.2 W	50% HDD advantage[2]
Read Intensive TB/W	3.0 TB/W	3.0 TB/W	Equivalent[2]
Write Intensive TB/W	2.5 TB/W	3.5 TB/W	40% HDD advantage[2]

[1]Assumes no HDD power management i.e. the drives are not spun down
[2]Assumes 50% idle and a 45%/5% R/W W/R mix

How and Where Grand Canyon Uses 24G SAS

- Warmstorage usecase (aka Tectonic)
  - HDD fan-out per CPU
  - Flexibility to take SATA and SAS HDDs

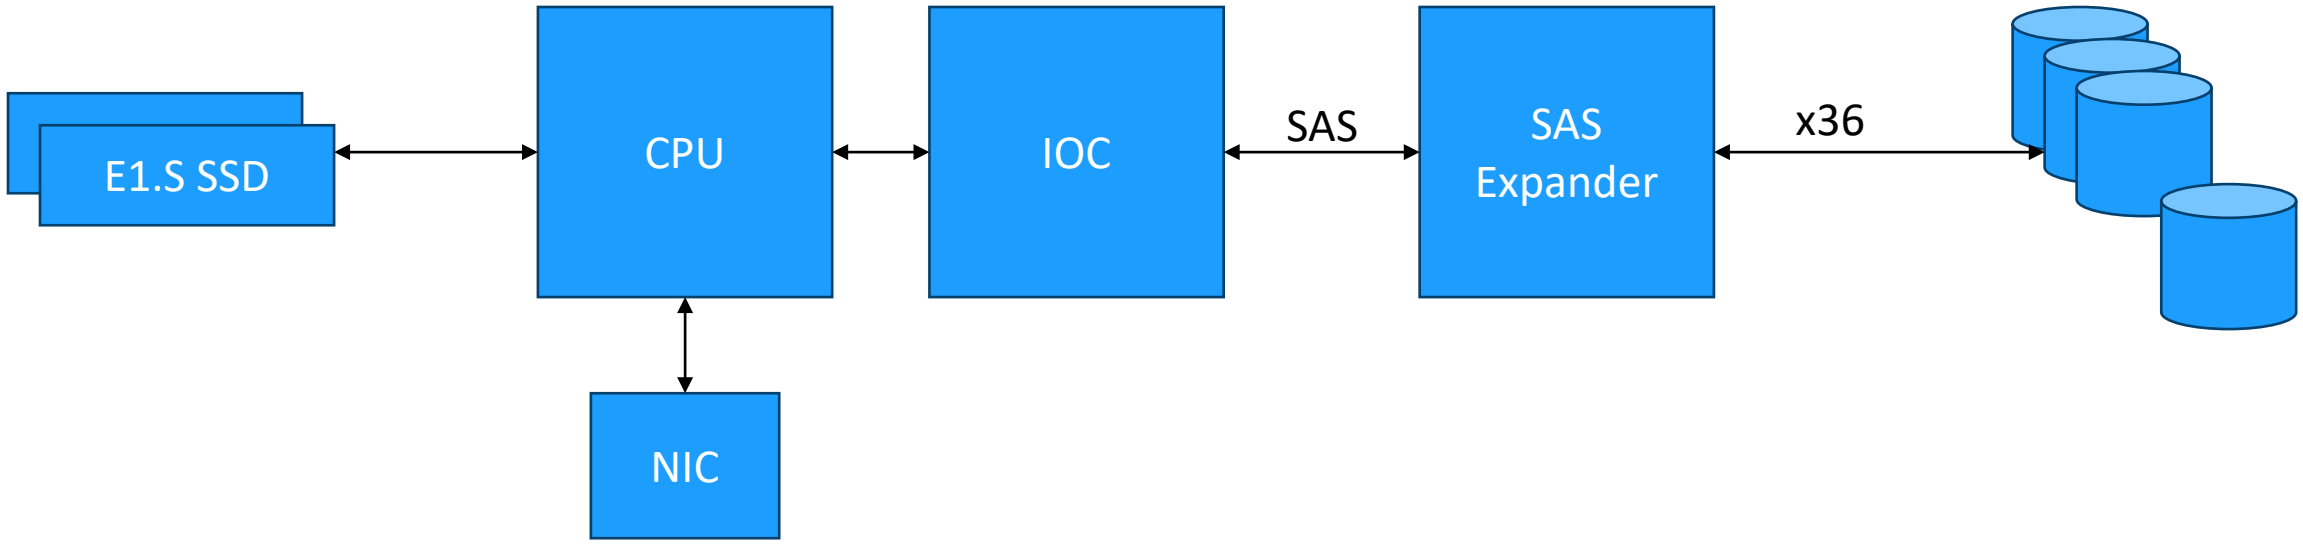

How and Where Grand Canyon Uses 24G SAS

- Coldstorage usecase
  - Even larger HDD fan-out
  - JBOD chassis daisy-chaining

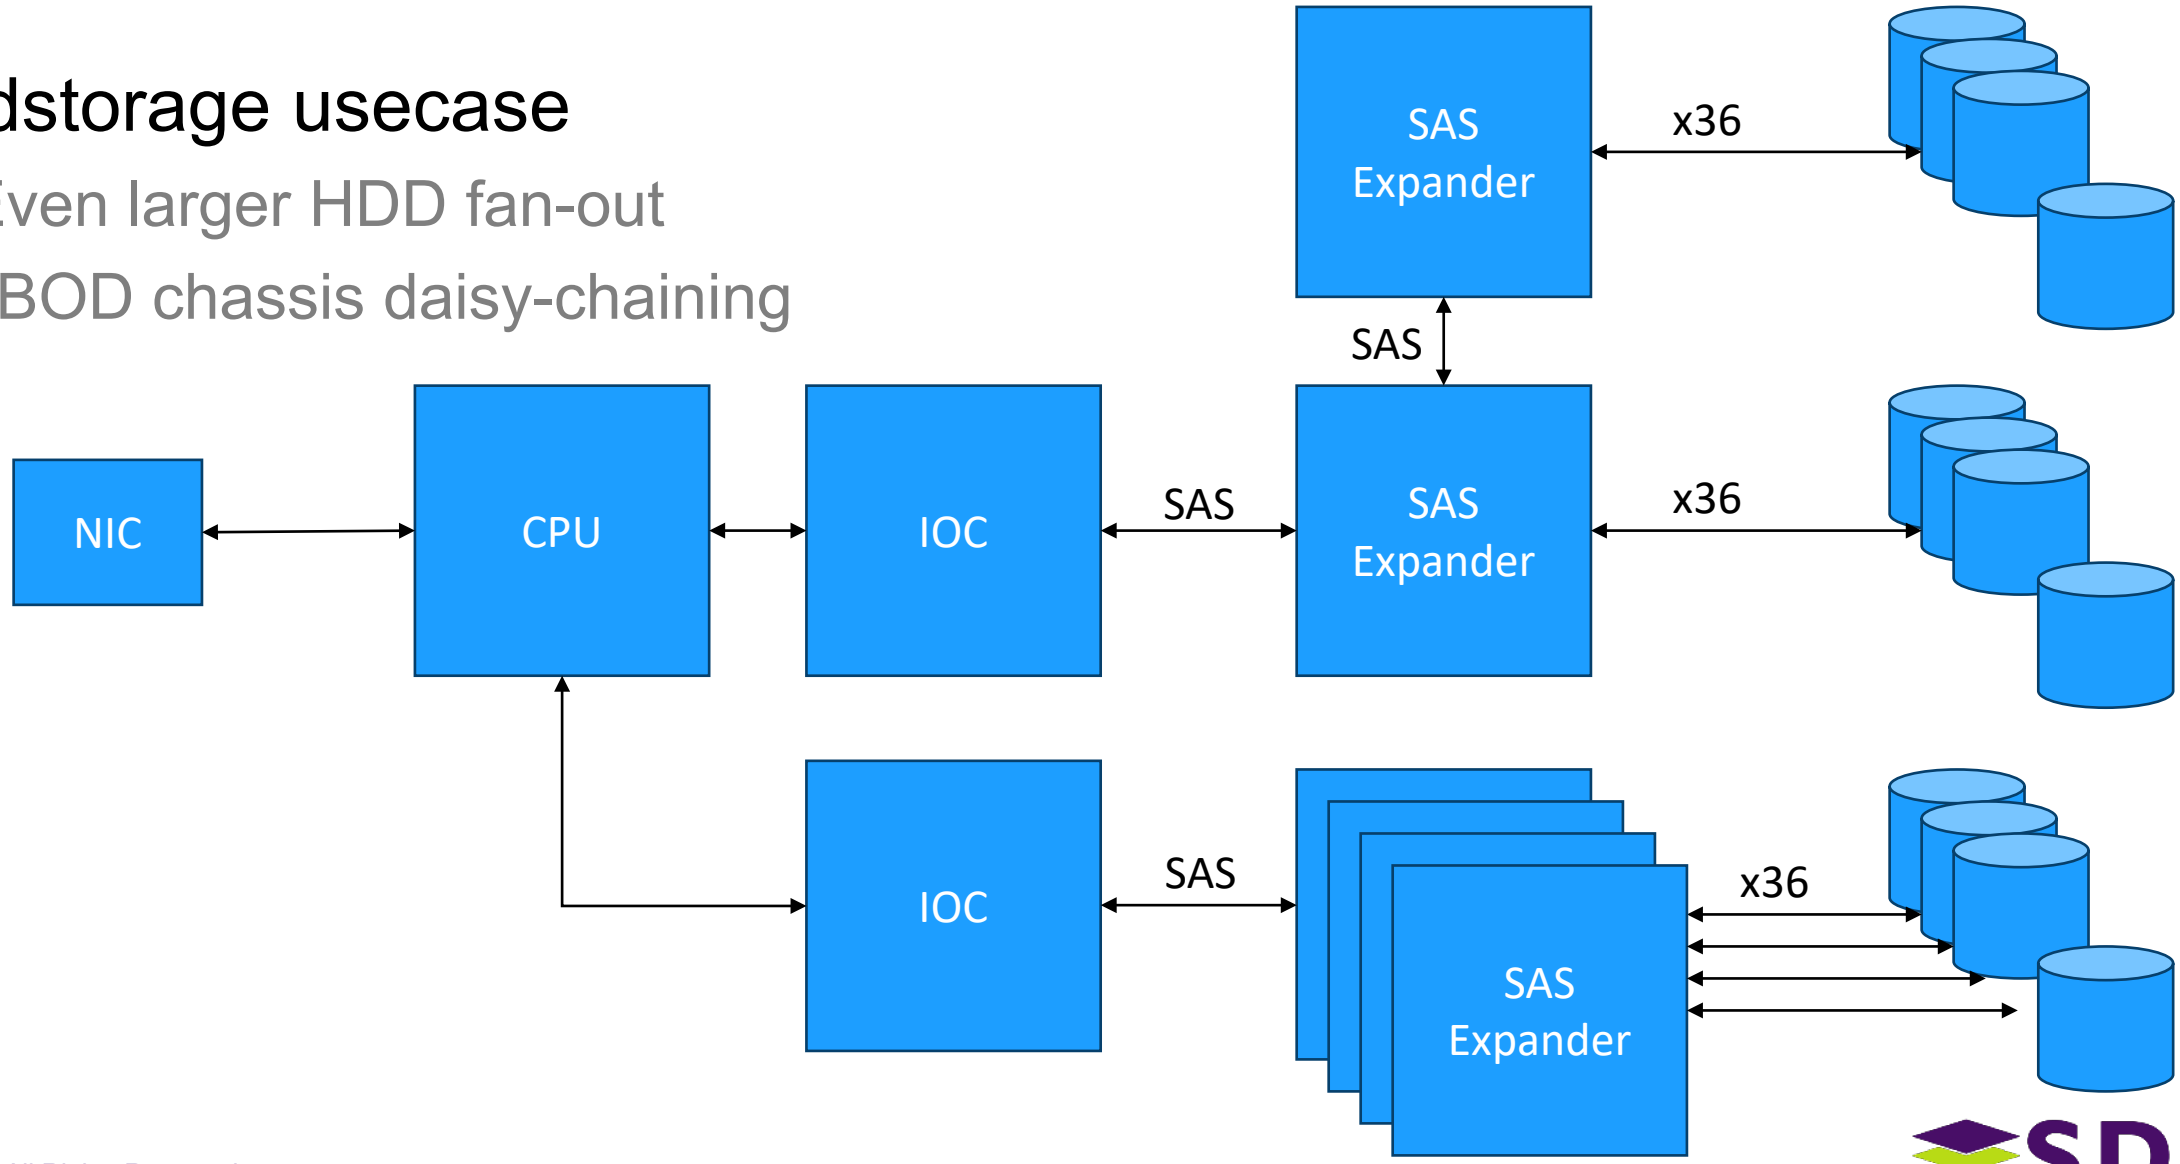

Why Grand Canyon Uses 24G SAS

- Scalability
  - Large HDD fan-out behind a CPU; 1:36 and 1:216 CPU:HDD ratio deployed in Meta
  - Beefy CPU in the future → even larger fan-out
- Flexibility
  - Enabled SAS to SATA switch for HDDs on the same platform
  - Support for SMR HDDs
  - Support for CDL on SAS as well as SATA drives
  - Enough interface perf to enable Dual Actuator adoption when necessary
- Reliability
  - SAS to SATA switch was not easy, but wasn't painful.
  - STP has come a long way. IOC+Expander+SATA HDD error handling has come a long way.

- Proven Interface to enable cost-efficient and performant HDD-based systems

Follow SNIA STA Forum

https://www.snia.org/sta-forum

Come chat with us tonight at the

STA BOF

8pm

Salons II & III

https://twitter.com/SAS_Storage

https://www.linkedin.com/company/scsi-trade-association

Serial Attached SCSI Playlist on SNIAVideo
https://www.youtube.com/@SNIAVideo

Please take a moment to rate this session.

Your feedback is important to us.

- Hyperscale data centers, which are large-scale facilities built to handle enormous amounts of data and support cloud computing services, often utilize a variety of infrastructure technologies to meet their unique requirements. While there isn't a strict rule that hyperscale data centers exclusively use SAS (Serial Attached SCSI) infrastructure, there are several reasons why SAS has been widely adopted in these environments:

1. Scalability: SAS infrastructure offers excellent scalability, allowing data centers to easily expand their storage capacity and meet growing demands. SAS supports daisy-chaining of devices, enabling the connection of multiple drives in a single chain, which simplifies cabling and reduces the number of host bus adapters required. This scalability is essential for hyperscale data centers, where storage needs can rapidly increase over time.
2. Reliability: SAS technology provides high levels of reliability and fault tolerance, crucial for hyperscale data centers that require maximum uptime and data availability. SAS drives often incorporate features like error detection and correction, redundancy, and data integrity checks to ensure data reliability and minimize the risk of data loss or corruption.
3. Flexibility: SAS infrastructure offers backward compatibility with SATA (Serial ATA) drives, allowing data centers to leverage existing investments in SATA drives while gradually migrating to SAS-based storage solutions. This flexibility is valuable for hyperscale data centers that may have a diverse range of storage devices and need to integrate new hardware seamlessly.

- It's worth noting that hyperscale data centers may employ a mix of storage technologies, including SAS, SATA, NVMe (Non-Volatile Memory Express), and other solutions, depending on their specific needs and workload requirements. The choice of infrastructure in a hyperscale data center is often based on a combination of performance, scalability, reliability, cost-effectiveness, and compatibility considerations.

