Hyperscale Boot SSDs for future...

Karthik Shivaram, Meta
Charles Kunzman, Google

Connect. Collaborate. Accelerate.

Problem Statement

Flash Capacity Trend

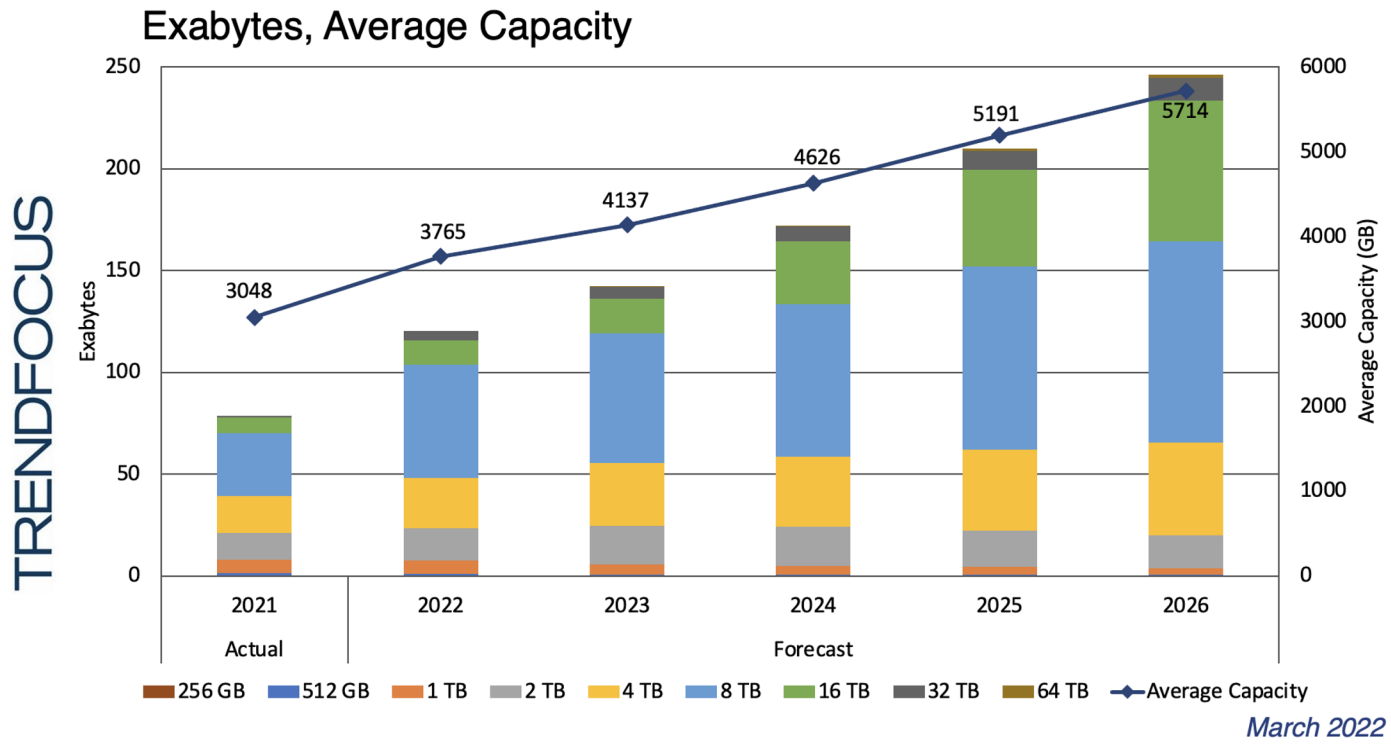

- **Data Drive capacity keeps increasing**
- **Boot Drive capacity trends are also increasing**
- **Increasing capacity = Increasing expense**

Client vs Hyperscale

Metric	Client	Hyperscale
Device Idle Time	More	Less
Power Saving Features	Required	Not Required
Performance Metric	Fresh out-of-box	Sustained
Monitoring Capabilities	Not important	Important
Endurance Requirements	Low	High

Client and Hyperscale Boot SSDs have different requirements

Hyperscale Boot I/O Profile Example

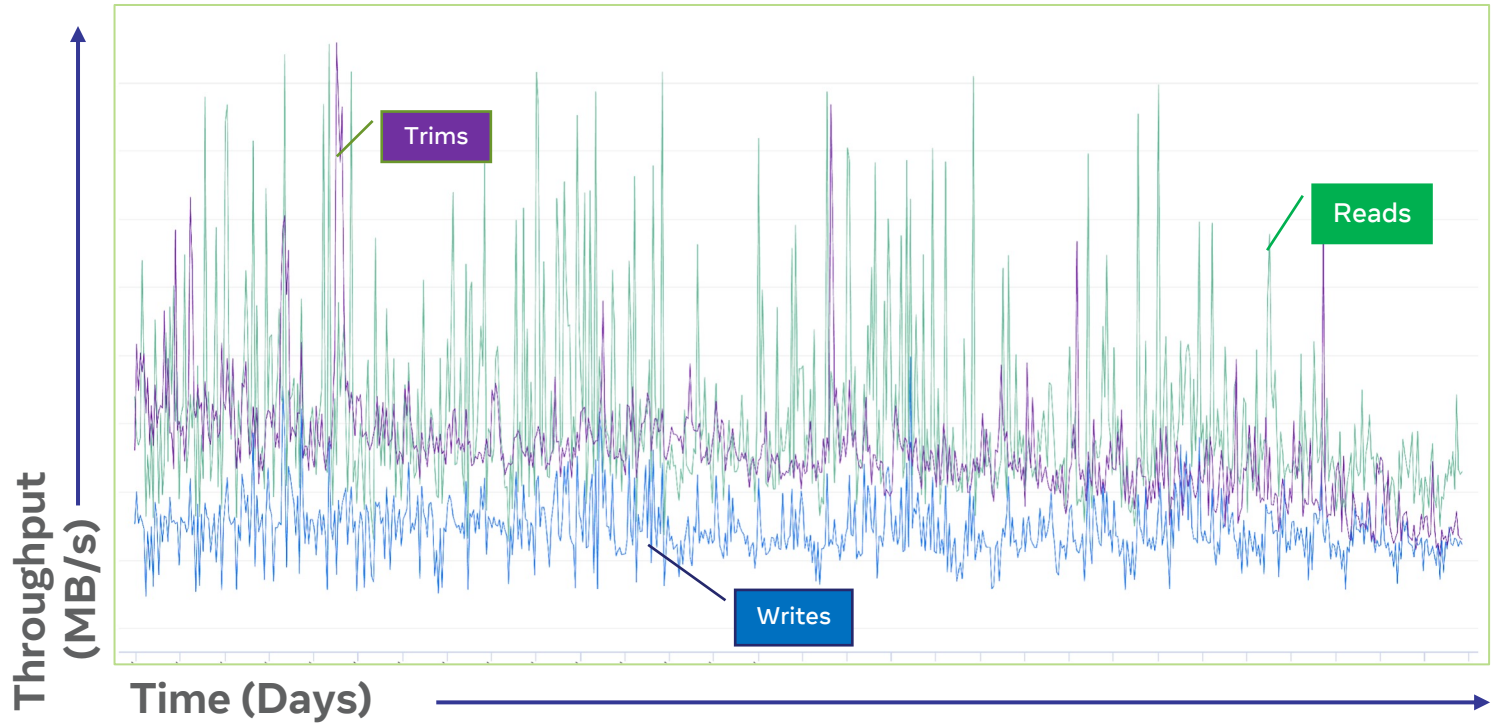

- Boot Drive experiences constant traffic with no idle time
- TRIM rate on Boot Drive is very high
  - Latency stalls due to TRIM are not desirable

Boot Drive I/O Traffic Breakdown

- Majority of the traffic is random in nature

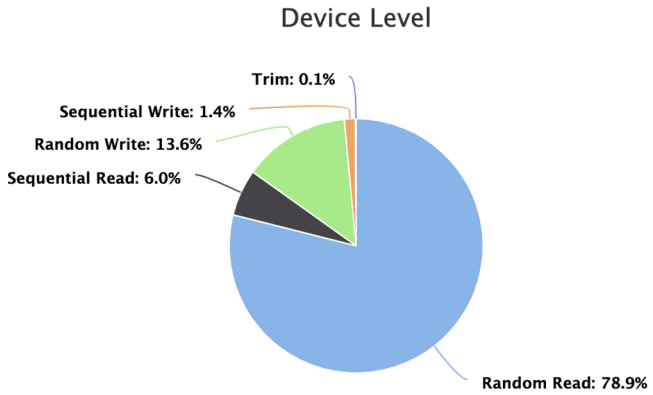

- Majority of the traffic has low queue-depth

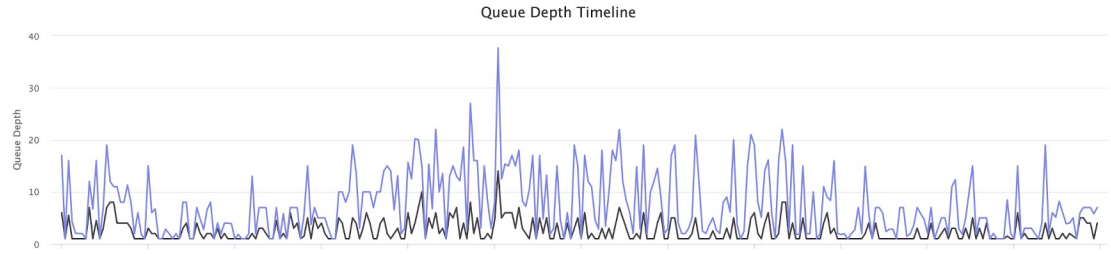

- **Majority of traffic is random in nature**
- **Workloads have low queue depth**
- **User experience is sensitive to latency**

Boot Drive Performance @Scale

- Performance methodology for Hyperscale Boot Drives is not clear
- No public minimum bar (or performance target) defined
- Leads to huge drive-to-drive performance variation

Hyper-Scalers struggle with huge variation in drive performance due to lack of public performance targets

Hyperscale Endurance & Monitoring Requirements

- Monitoring at scale is important
  - Boot Drives are deployed all over the world
- Monitoring helps predict & detect failing drives
- Boot SSDs need higher endurance to prevent early wear out
  - Reliability is extremely important as repair at-scale is extremely challenging

Hyperscale Boot SSDs require high endurance and enhanced monitoring at-scale

Security

- Hyperscale customers care and value Security and Privacy
- Industry is fragmented with many standards in this space
- Many “optional” features in Security Standards

Hyperscale Boot SSD require higher level of Security due to scale of deployment and sensitivity in nature of data!

Summary of Boot SSD Challenges

- Capacity of SSDs are increasing
  - Boot Drive capacity needs remain small
- Client SSDs are designed with a focus on Client use-cases
- Hyper-Scalers require higher endurance, robust security features and enhanced monitoring compared to Client SSDs
- Hyper-Scalers have confidential Boot SSD specifications which doesn't encourage industry collaboration
- Hyperscale Boot SSD test requirements are not clear

The solution…

Connect. Collaborate. Accelerate.

Meta & Google are collaborating to combine requirements to create a OCP Hyperscale Boot SSD Specification.

Benefits Open Boot Drive Spec

- Meta & Google have merged their SSD boot drive requirements into a single document enabling the following benefits:
  - Allows the wider market to:
    - Understand features Hyper-Scalers need to manage a boot SSD at-scale
    - Enable standardization of boot SSD across general market
  - Reduces SSD market fragmentation
  - Promotes industry alignment on SSD boot drive adoption
  - Enables use of open-source tools to manage & monitor boot SSDs at-scale
  - Allows 3rd parties to create test-suites to simplify the drive qualification process

Open requirements increase industry collaboration and enable Boot SSD standardization

Key Spec Focus Areas

- Specifies requirements for Hyperscale Boot SSD
- This includes requirements around:
  - NVM Express
  - PCI Express
  - SMART Logs
  - Reliability
  - Power
  - Performance
  - Security
  - Side-Band/SM-Bus
  - Sustainability
  - Thermal
  - Monitoring

Everything needed to build a Hyperscale Boot SSD!

Solution

Today

Lack of Industry Standards for Hyperscale Boot Drives

- SSD Boot Drives are customized but there is no Industry Standard to capture all the requirements.

Future

OCP Hyperscale Boot Drive Specification (Target Mid 2022)

- Benefits system makers and SSD providers.
- Enables additional collaboration between Hyper-Scalers and industry.