

Storage Industry Update



Ross Stenfort, Meta

Form Factors Today:

E1.S Capacity Scaling:

Assume: 8 Placements & 32 Die stacks

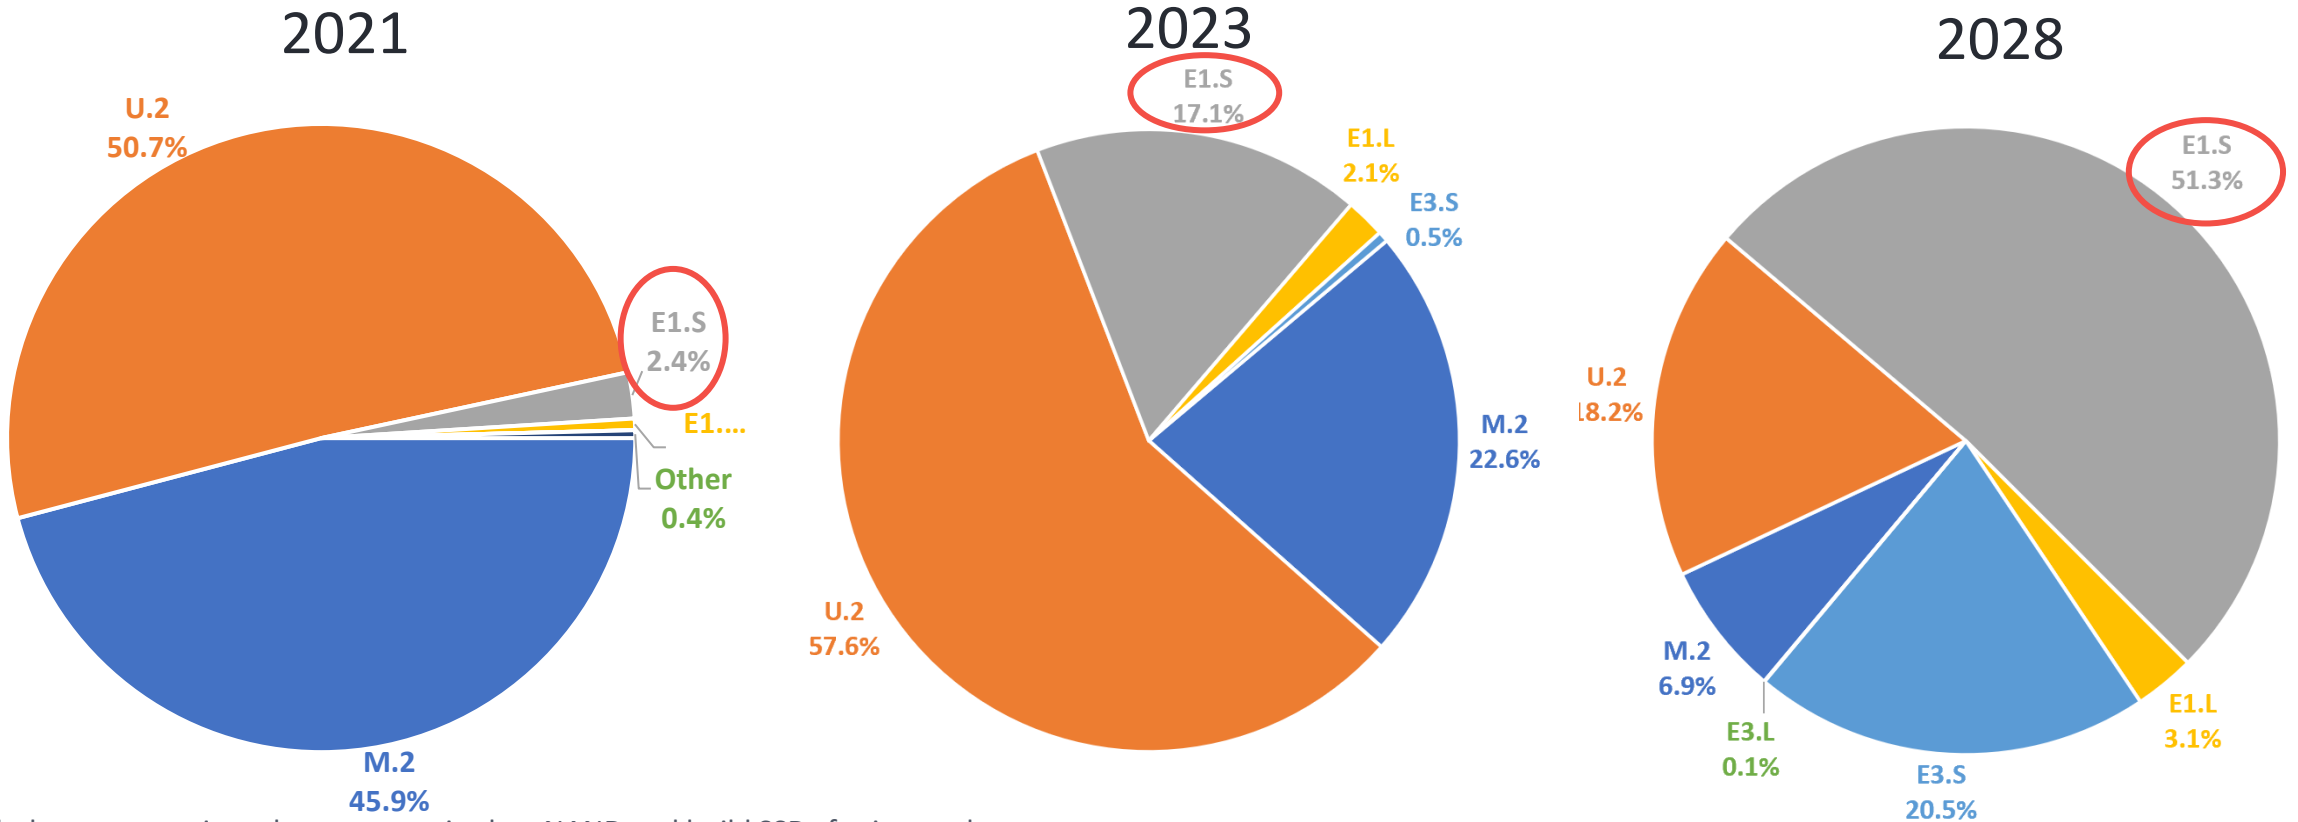
- TLC (1 Tb Die): 32 TB
- QLC (2 Tb Die): 64 TB
- QLC (4 Tb Die): 128TB

**Key Take
Always:**

- **E1.S Market Growth**
 - 2021: 2%
 - 2023: 17%
 - 2028: 51%
- **Capacity**
 - E1.S scales to 128 TB
- **E1.S Growth due to:**
 - Density, Power, Performance capabilities

Industry Adoption Trend (Units*)

TRENDFOCUS



*Data excludes consumption where companies buy NAND and build SSDs for internal use.

Flexible Data Placement (FDP) Update

❖ *Benefits*

- Improves Endurance, Performance, Quality of Service through data placement

❖ *Key Adoption Drivers:*

- Low development effort with high benefits
- Backwards compatibility
 - Applications are not required to understand FDP to get benefits
 - Applications which understand FDP have increased benefits
- Open-Source Support

❖ *Open-Source Support:*

- **Linux Kernel:**
 - Full support through I/O Passthru (**Complete**)
 - Lifetime Hints (**In progress**)
- **xNVMe:** Full Support (**Complete**)
- **QEMU:** FDP Emulation (**Complete**)
- **Fio:** Full Support (**Complete**)
- **nvme-cli:** Full Support (**Complete**)
- **Cachelib:** Full Support (**Complete**)

OCP Datacenter NVMe SSD Specification



OPEN
Compute Project

Datacenter NVMe® SSD Specification

Version 2.5 (10302023)

Author: Ross Stenfort, Meta
Author: Lee Prewitt, Microsoft
Author: Paul Kaler, HPE
Author: David Black, Austin Bolen, Dell Technologies
Author: Chris Sabol, Charles Kunzman, Google

❖ Latest Released Specification

- [OCP Datacenter NVMe SSD Specification V2.5](#)
- **Key Updates:**
 - Debuggability – Human Readable Telemetry with open-source tooling (OCP NVMe CLI)
 - Security
 - Flexible Data Placement (FDP)
 - Issues found in deploiments, clarifications, bug fixes

❖ Next Version 2.6

- In Progress – Targeting 2024 Release
- **Key Updates**
 - Additional Telemetry Improvements
 - Additional Health Statistics
 - Open-Source Tooling Support Required
 - OCP NVMe CLI
 - Clarifications, bug fixes

Key Features for Managing at Scale

- ❖ OCP Health Information Log
- ❖ OCP Latency Monitoring Feature
- ❖ OCP Formatted Telemetry for Human Readable Logs
- ❖ Open-Source OCP NVMe CLI

Improved Debug Method Results:

- Faster Debug & Qualification Cycles
- Effective Debug at Scale

*Improvements Based
On Deployment
@Scale*

OCP Storage Project Update

2023

NVMe HDD rev. 1.0	[spec]	Microsoft, Seagate, Western Digital
Datacenter SAS-SATA Device rev. 1.0	[spec]	HPE, Meta, Microsoft
Datacenter NVMe SSD rev. 2.5	[spec]	Dell, Google, HPE, Meta, Microsoft
NVMe Telemetry Scripts	[scripts]	Samsung
2 nd OCP Storage Teck Talk	[virtual]	May 16 th

2024

Datacenter SAS-SATA Device rev. 1.5	Coming	HPE, Dell, Meta, Microsoft
Datacenter NVMe SSD rev. 2.6	Coming	Dell, Google, HPE, Meta, Microsoft
EDSFF SSD Reference Design	Coming	ScaleFlux
OCP Marketplace	Growing Fast	8 SSDs (including OCP V2.5 support) and Multiple OCP SSD Testing Services
Meta Open-Source SSD Qualification Tests	Complete	OCP Test Framework OCP Storage Test Cases
3 rd OCP Storage Teck Talk	[virtual]	May 15 th

- ❖ [OCP Storage Project Link](#)
- ❖ Meetings are 2nd Thursday of the month
 - If you have ideas to present in monthly storage project meetings reach out

Thank You