Emerald Tiers: Focusing on SSD+MAID Through a Green Lens

Zhaokang Ke [1] *Jim Diehl* [1] *Ya-Shu Chen* [2] *David H.C. Du* [1]

[1]*University of Minnesota, Twin Cities*

[2]*National Taiwan University of Science and Technology*

University of Minnesota
Driven to Discover®

Data Growth & Carbon Concerns

- Enterprise storage demand is increasing
- Data center electricity use is rapidly increasing
- **Sustainability** becomes crucial as storage scales up
- Storage systems emit carbon in two ways
  - **Operational Carbon Emissions**
  - **Embodied Carbon Emissions**

Figure 1: Enterprise Data Shipments and the Active Installed Base, 2020-2035

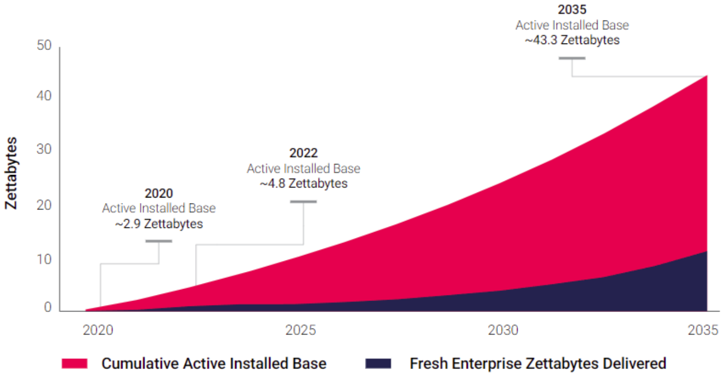

Source: Furthur Market Research (January 2024)

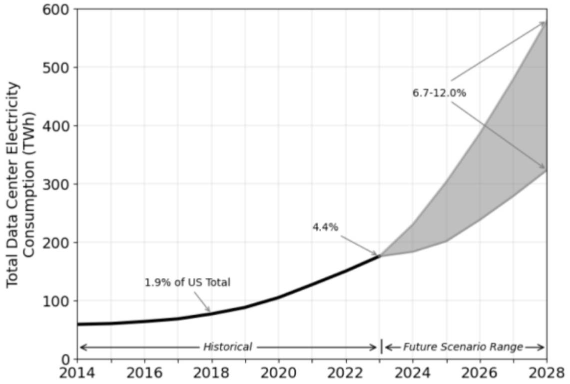

Figure ES-1. Total U.S. data center electricity use from 2014 through 2028.

Fig. 1 source: *The Sustainable Preservation of Enterprise Data*
https://furthurdata.com/wp-content/uploads/2024/02/Sustainable-Preservation-of-Enterprise-Data_V13.pdf
Fig. ES-1 source: *2024 United States Data Center Energy Usage Report* https://doi.org/10.71468/P1WC7Q

Operational Carbon Emissions

- Calculated as: *Grid Carbon Intensity (kgCO_2 / kWh) × Energy Use (kWh)*
- Reduced by using cleaner energy and/or reducing energy use
- **HDDs（higher operational carbon per I/O）**
  - Major source: Mechanical motor spinning platters 24/7
  - How to reduce: Change to a low-power mode or shut down
- **SSDs（lower operational carbon per I/O）**
  - When active, newer SSDs can use more power per device than HDDs but offer far higher performance
  - Can quickly transition into and out of lower-power idle modes if not accessed

Embodied Carbon Emissions

- Emissions from device's production including material extraction, fabrication, packaging, transport, and disposal
- Reduced by **extending device lifetime**
- **HDDs（lower embodied carbon emissions）**
  - Major source: Metals, motors, platters, etc.
  - How to reduce: Minimize spin-up/down cycles and/or reduce run time
- **SSDs（higher embodied carbon emissions）**
  - Major source: NAND chip production
  - How to reduce: Prolong device life via effective garbage collection, wear leveling, and write amplification control

Embodied Carbon Emissions

- Emissions from device's production including material extraction, fabrication, packaging, transport, and disposal
- Reduced by **extending device lifetime**
- **HDDs（lower embodied carbon emissions）**

Challenge—balancing **performance**, **cost**, and **both operational** and **embodied carbon**

ime

- **SSDs（higher embodied carbon emissions）**
  - Major source: NAND chip production
  - How to reduce: Prolong device life via effective garbage collection, wear leveling, and write amplification control

Contributions

- We propose a **tiered** **SSD+MAID** architecture that balances performance, cost, and carbon emissions.
- We develop a comprehensive **carbon emission model** capturing both **operational** and **embodied** emissions.
- We conduct extensive evaluations showing **SSD+MAID** can offer superior trade-offs over **all-SSD**, **all-HDD**, and **SSD+HDD** systems.

MAID: A Historical Solution

- Introduced 2002: Massive Arrays of Idle Disks
- Idea: Most data is cold → keep most HDDs spun down
- Goal: Cut **operational power** by keeping most HDDs spun down
- Design: Always-on disk cache for frequent data, spun-down disks for cold data
- Why MAID failed: High latency because of slow spin-up
- Can we revisit MAID—retain its strengths, fix the weaknesses?

Our Goal: ⚖️ Balance performance, cost, carbon emissions

Carbon-Aware Tiered Design: SSD+MAID

- Ensure older data remains readily accessible for a period before archiving it to long-term storage (e.g., tape, silica, DNA)
- Proposed tiers:
  - **SSDs**: Fast tier for hot data
  - **Active-on HDDs**: For moderately-accessed data
  - **Powered-off HDDs**: Low-power for rare access

SSD+MAID Data Flow

MAID

- Writes: SSDs → active-on HDDs → powered-off HDDs
- Reads: Served from SSDs, active-on HDDs, or powered-off HDDs (with spin-up delay)
- Migration: Batch every 1-2 days to minimize spin-ups

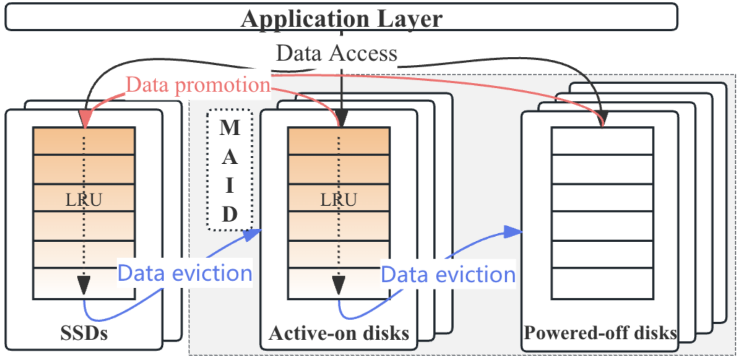

Carbon Emission Model: Embodied + Operational

- Total Carbon = Embodied + Operational

$$C_{\text{total}} = C_{\text{emb}} + C_{\text{oper}}$$

- Embodied: SSD (TBW, aging); HDD (hours, spin-up cycles)

$$C_{\text{emb}} = N_{\text{SSD}} \cdot \max\left(\frac{w \cdot 365}{\text{TBW}_{\text{SSD}}}, \frac{1}{L_{\text{SSD}}}\right) + N_{\text{HDD}} \cdot 365 \cdot \max\left(\frac{T_{\text{HDD,active}}}{L_{\text{HDD}}}, \frac{N_{\text{spin-up}}}{\text{Cycles}}\right)$$

- Operational: SSD (active, idle); HDD (active, spin-up energy)

$$C_{\text{oper}} = 365 \cdot [E_{\text{SSD}}^{(\text{day})} + E_{\text{HDD}}^{(\text{day})}] \cdot \text{CI}$$

- We assume data access follows the power law and data exits our system for deep archive after one year

Experimental Setup

- **Workload-aware**: Tested with **IBM Cloud Object Storage workloads** (read/write-heavy, mixed)
- We choose the highest-capacity representative **enterprise-class** devices for which **sustainability reports** are available

Table 1: Workload Characteristics

Characteristic	Read-heavy	Mixed	Write-heavy
Trace number	IBM_011	IBM_074	IBM_020
Pow-law α	2.58	2.25	2.52
Write (GB/day)	498.96	857.34	3102.39
Read (GB/day)	2174.44	857.23	181.36
Write req. (ops/day)	41228	53,708	128,203
Read req. (ops/day)	135,661	53,290	25,822
Avg. read size (MB)	16.41	16.47	7.19
Supported servers	6	4	1
Storage capacity (TB)	1200	1200	1200

Table 2: Device Specs for Carbon Emission Model

Dev.	Cap.	Emb. C	Power (opr/idle)	Lifetime	R/W (MB/s)
SSD	15.36 TB	38.08 kg	8.5 / 4.7 W	5 yr / 28,000 TBW	2100 / 1000
HDD	22.00 TB	28.92 kg	6.0 / 5.7 W	43,800 h / 600,000 cycles	285 / 285

Evaluation 1: Impact of Active-on/Powered-off Disk Ratio on Carbon Emissions

- SSDs + only active-on HDDs emits the most carbon (4,52,0)
- Too few active-on disks cause **frequent spin-ups**, increasing emissions (4,0,52)
- A **balanced ratio** SSD+MAID lowers both operational and embodied carbon (4,5,47)

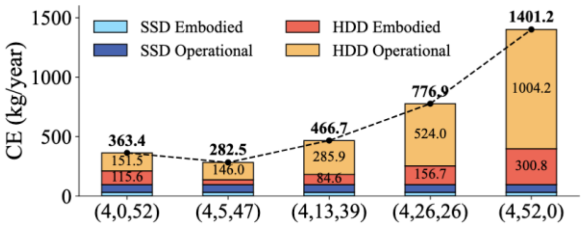

Figure 2: Read-heavy, Fixed SSD Size

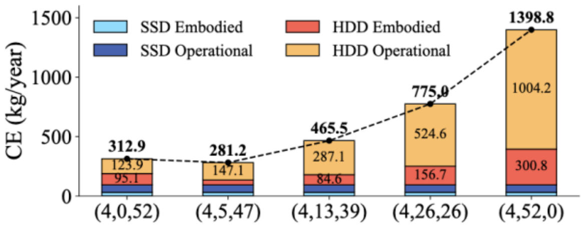

Figure 3: Mixed, Fixed SSD Size

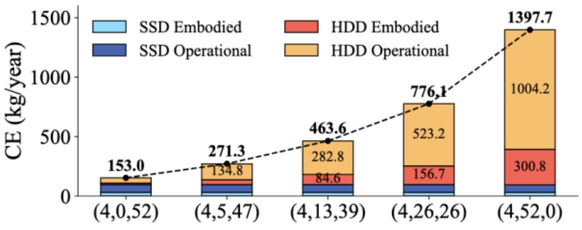

Figure 4: Write-heavy, Fixed SSD Size

(x,y,z) in the figures is (#SSDs, #active-on HDDs, #powered-off HDDs)

Evaluation 2: SSD Capacity vs. Carbon Emissions

- More SSDs → higher carbon due to per-TB cost
- SSD+MAID hybrids beat all-SSD and pure active-on HDD config
- MAID-only config (0,14,41) has the lowest emissions

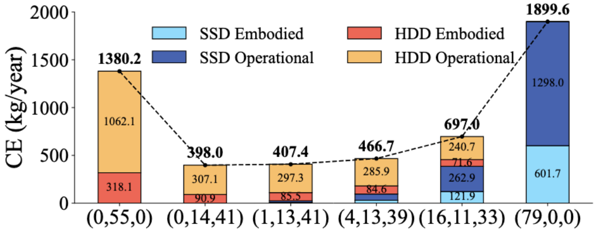

Figure 5: Read-heavy, Var. SSD Size

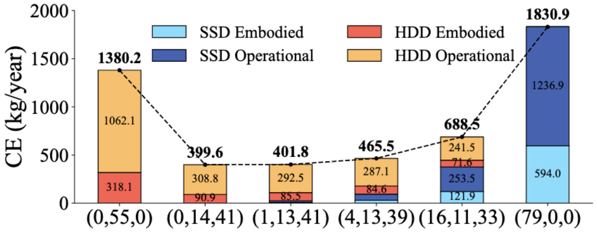

Figure 6: Mixed, Variable SSD Size

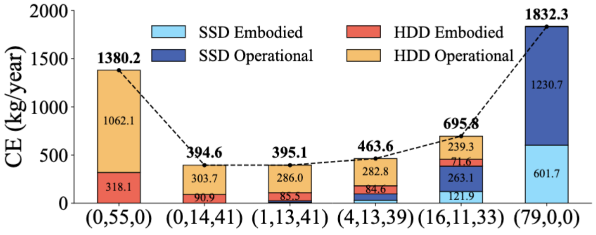

Figure 7: Write-heavy, Var. SSD Size

(x,y,z) in the figures is (#SSDs, #active-on HDDs, #powered-off HDDs)

Evaluation 3: System Latency

- All-SSD has **lowest latency**; adding SSDs helps reduce delays
- Latency worsens when more requests hit powered-off HDDs
- Higher skew (larger α) → fewer spin-ups → lower latency
- Larger request size **dilutes spin-up cost**, reducing latency penalty multiple of tiered configurations

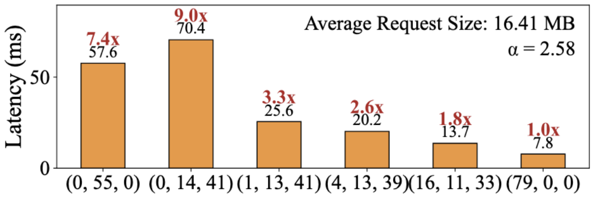

Figure 8: Read-heavy Latency

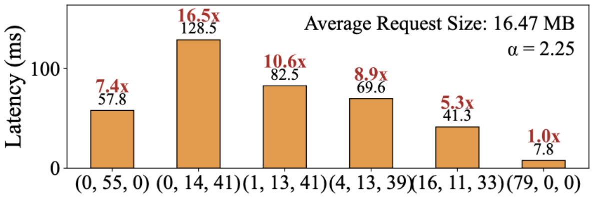

Figure 9: Mixed Latency

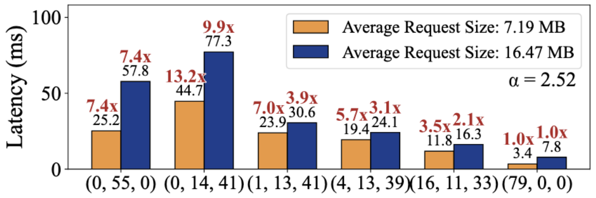

Figure 10: Write-heavy Latency

(x,y,z) in the figures is (#SSDs, #active-on HDDs, #powered-off HDDs)

Evaluation 4: Cost-Benefit Analysis

- **All-SSD** is **14× more expensive** than MAID; hybrid SSD+MAID (4,13,39) is a **cost-effective middle ground**
- Hybrid config (4,13,39) achieves:
  - **1.17× carbon** of MAID, but only ¼ **of all-SSD**
  - **1.7× cost** of MAID, but **~9× cheaper** than all-SSD
  - **3.5× faster** than MAID, **2.6× slower** than all-SSD

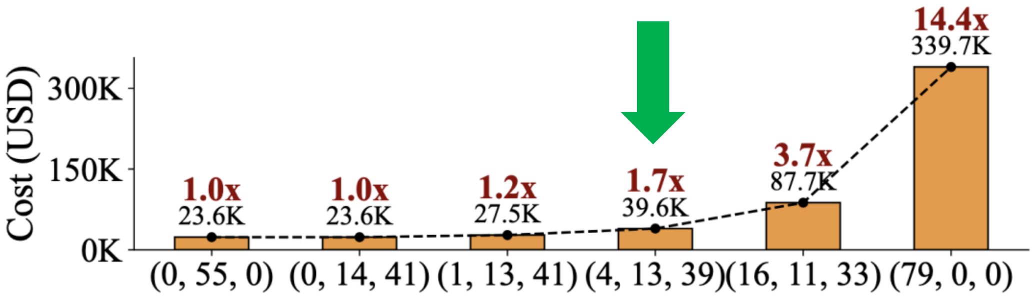

Conclusion

- Proposed a **tiered SSD+MAID architecture** balancing performance, cost, and carbon
- Developed a **workload-driven carbon model** capturing:
  - Access Model (power-law behavior) + Spin-Up Model (cold disk activation)
  - Operational + Embodied emissions, performance, and financial cost
- Enables **informed design trade-offs** for sustainable storage
- Call to action: **Vendors, please report embodied carbon** for modern devices

Thanks!
Questions?

University of Minnesota
Driven to Discover®

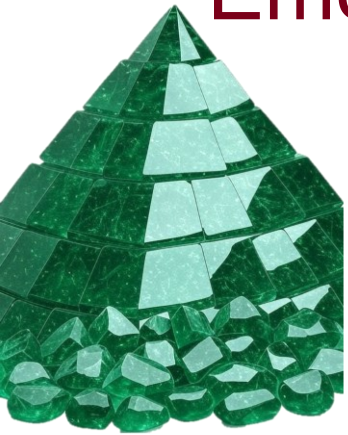

Emerald Tiers: Focusing on SSD+MAID Through a Green Lens

Zhaokang Ke [1] *Jim Diehl* [1] *Ya-Shu Chen* [2] *David H.C. Du* [1]

[1]*University of Minnesota, Twin Cities*

[2]*National Taiwan University of Science and Technology*

UNIVERSITY OF MINNESOTA Driven to Discover®

TAIWAN TECH NATIONAL TAIWAN UNIVERSITY OF SCIENCE AND TECHNOLOGY

UNIVERSITY OF MINNESOTA Driven to Discover®