Quantifying SSD Power Efficiency in the Data Center

Sujit Somandepalli
Principal Storage Solutions Engineer
Micron Technology, Inc.

Steven Wells
Fellow, Architect Storage Systems
Micron Technology, Inc.

Agenda

- Motivation
- Methods for measuring power consumption
- Metrics captured
- Reporting Options

Circa 2008 Question – which one gives better battery life?

3W HDD?

✓ 6W SSD?

Circa 2008 Question – which one gives better battery life?

3W HDD?

✓ 6W SSD?

Why?

How?

Data Centers consume a lot of power!

- AI is poised to drive 160% increase in data center power demand[1]
- SSDs can have a meaningful impact to platform power
- Measuring power consumption is the first step to optimization

- Important to correlate drive performance in real world applications to drive power consumption

1https:/www.goldmansachs.com/intelligence/pages/AI-poised-to-drive-160-increase-in-power-demand.html

Measuring Drive Power

- Digital Multimeter measuring voltages and current to the SSD
  - Pros: Very Accurate
  - Cons: Scalability and cost for data center deployments
- QUARCH Programmable Power Module
  - Pros: Very popular for measuring power, simulating power failures and more
  - Cons: Scalability, size for data center scale
- Custom hardware between SSD and power source
  - Pros: Scalable for data center scale
  - Cons: Multiple drives in a chassis
- Self-reported Power Measurements from the SSD
  - Pros: Scalable, measures per drive power
  - Cons: Not standard today, requires additional firmware / software

Self reported power standard options, what makes sense?

Visit us at the Micron Booth!

Dashboard like “Fuel Efficiency”?

Power Exceeded Notification?

Long term Power Histogram?

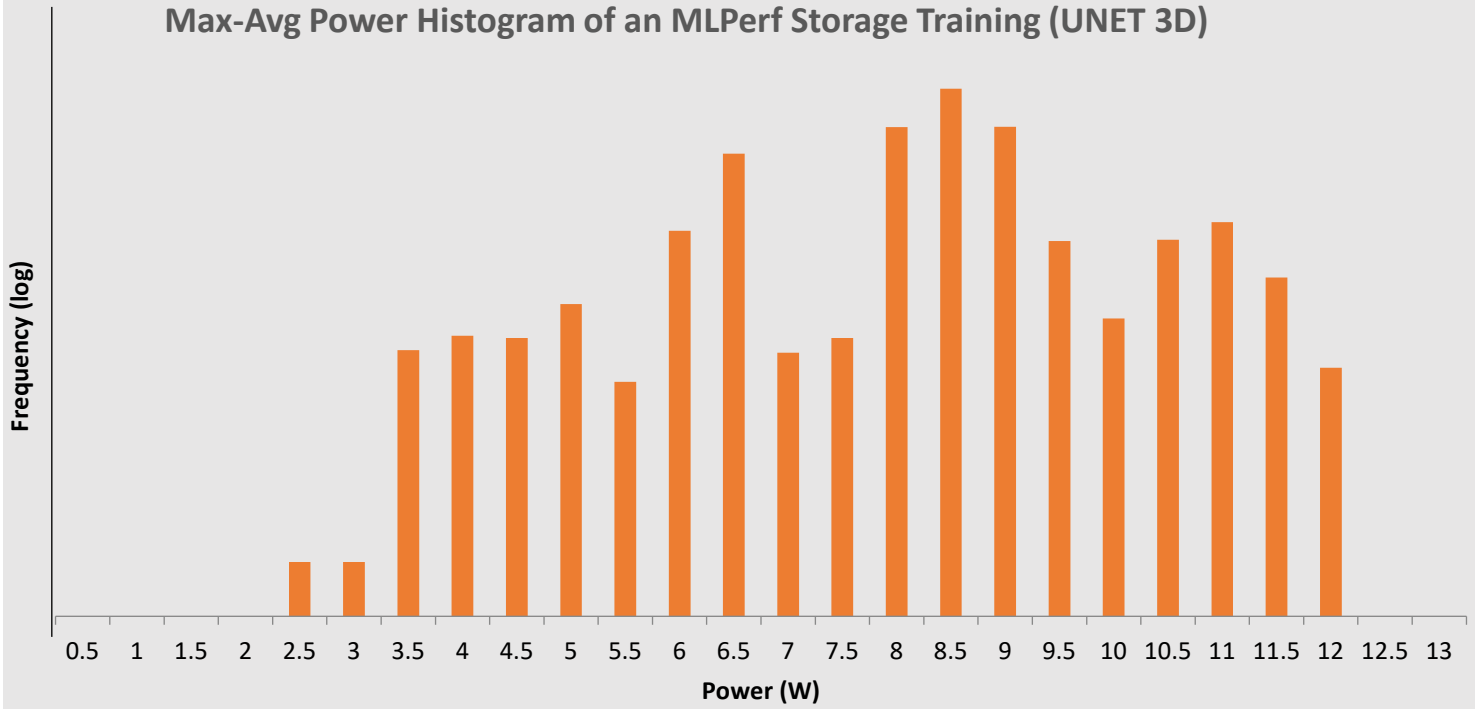

Best in class AI performance and energy efficiency

NVIDIA H100 GPU platform GNN workload results with BaM and GIDS[1,2]

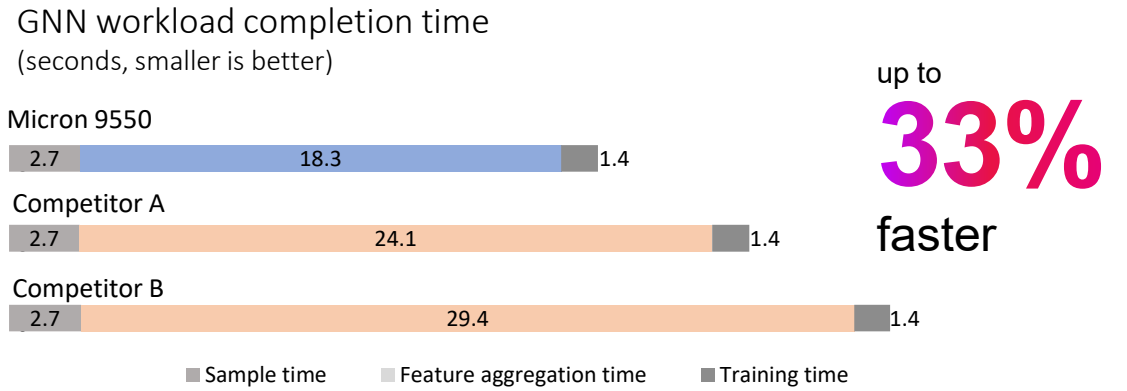

up to **33%** faster

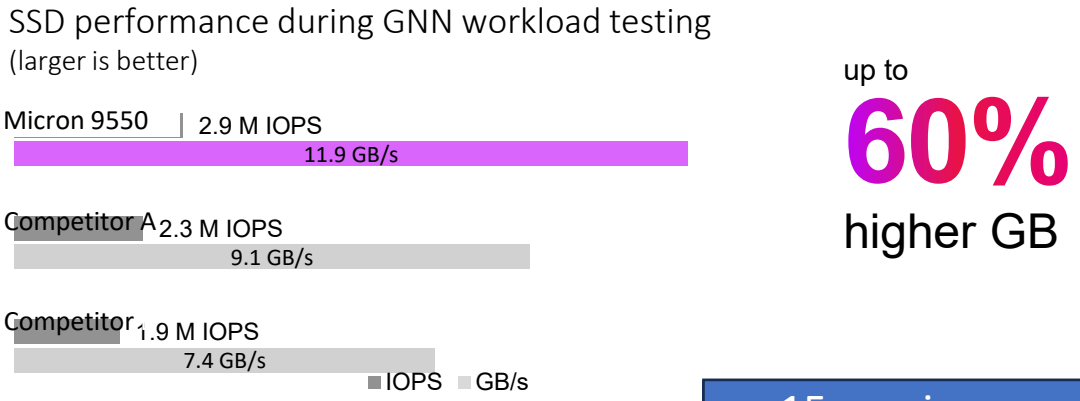

up to **60%** higher GB

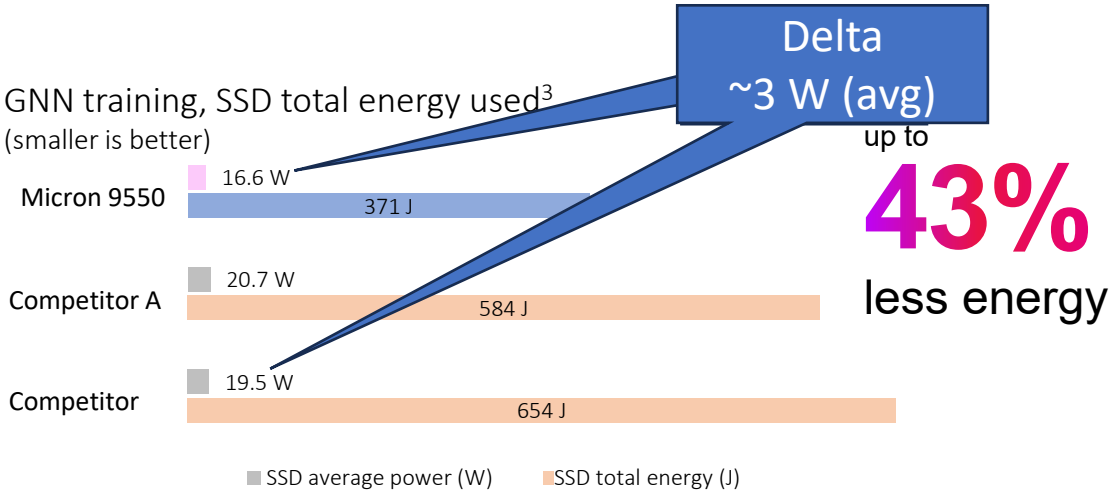

up to **43%** less energy

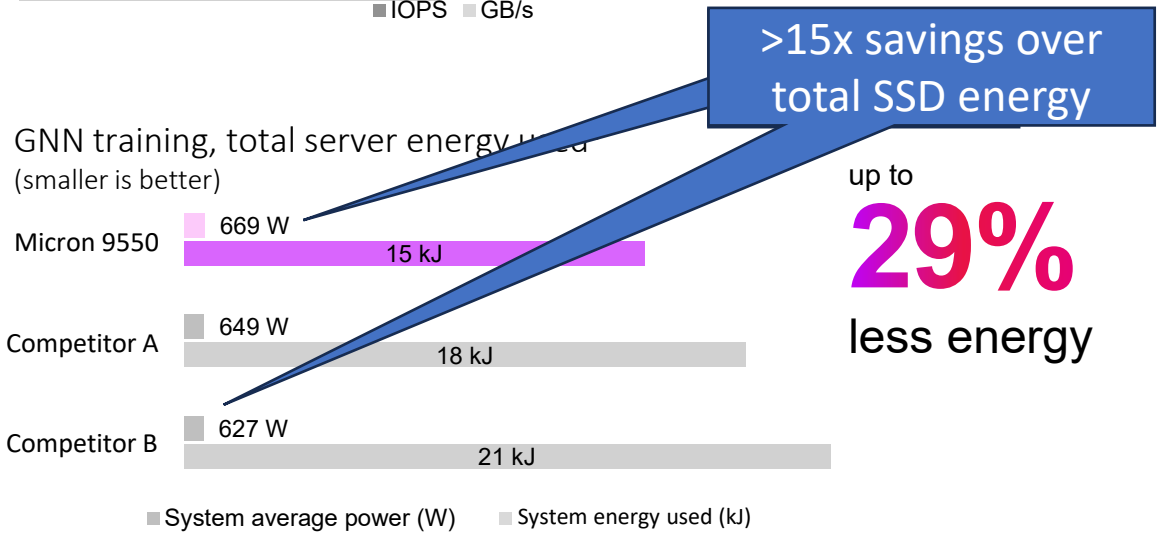

up to **29%** less energy

1. AI performance results using Big accelerator Memory (BaM) and GPU Initiated Direct Storage (GIDS) for Graph Neural Network (GNN) training workloads.
2. System under test: 2x Intel 8568Y+, 48-core, 1.5TB DDR5, NVIDIA H100NVL-96GB GPU PCIe Gen5x16, Ubuntu 20.04 LTS (5.4.0-182), NVIDIA Driver 535.161.08, CUDA 12.4 versus top competition as tested in Micron's labs.
3. Power efficiency units: J=Joules, W=Watts, kJ=Kilojoules

Aerospike DB – SSD energy correlates to System Energy

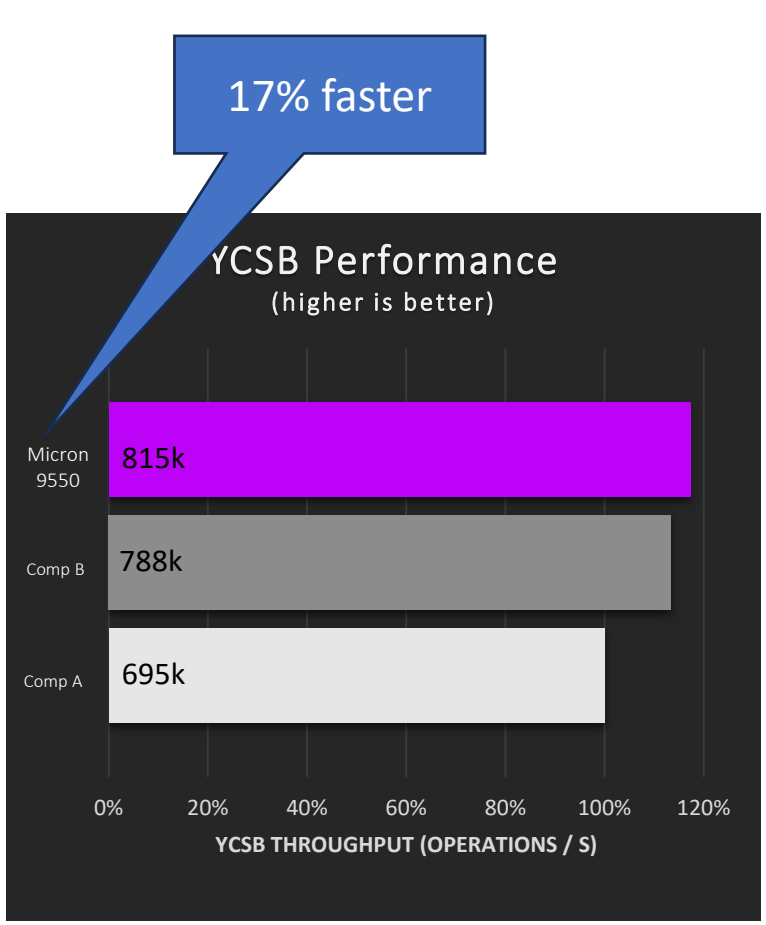

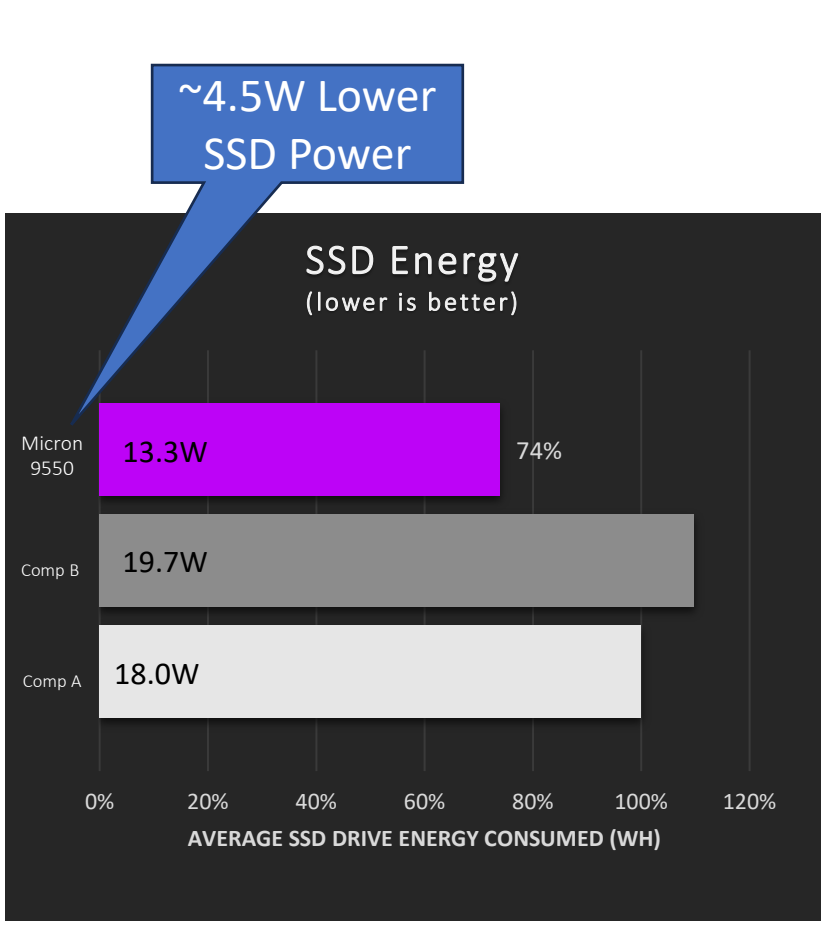

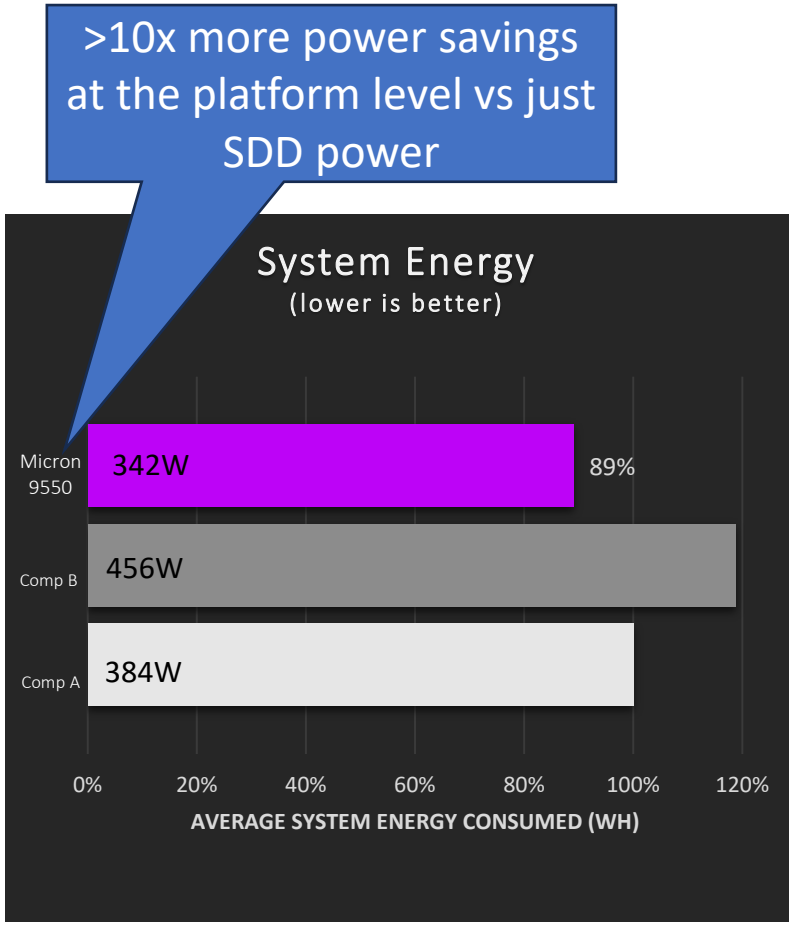

Tests were run on a Supermicro AS-1115CS-TNR
4th Generation EPYC 9564 Processor
12x 64GB Micron DDR5 RDIMMs
Micron 9550 Pro 7.68TB PCIe Gen5 SSD

Drive Power Consumption in real-world applications

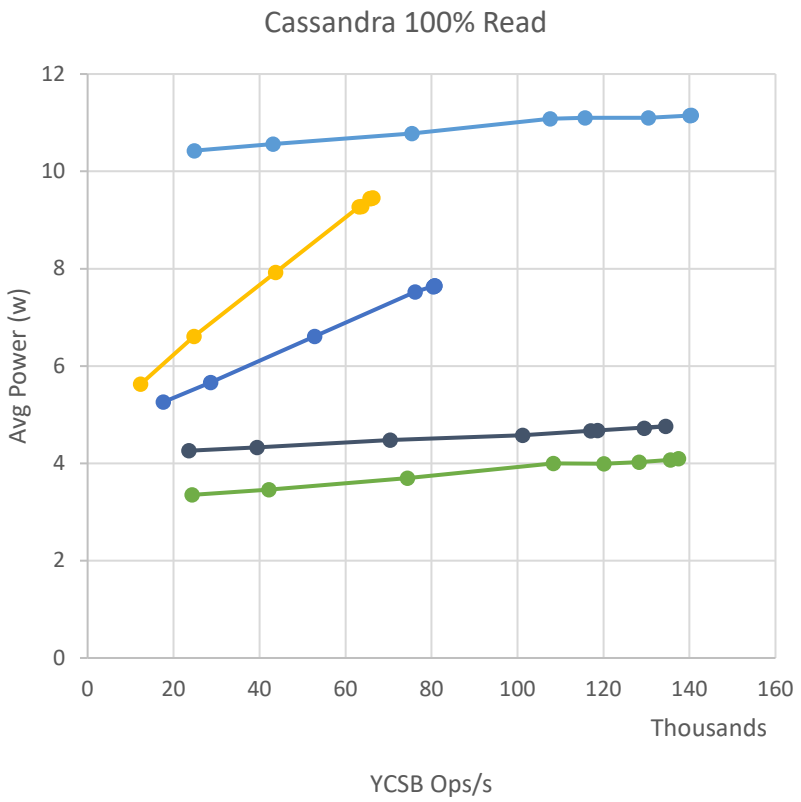

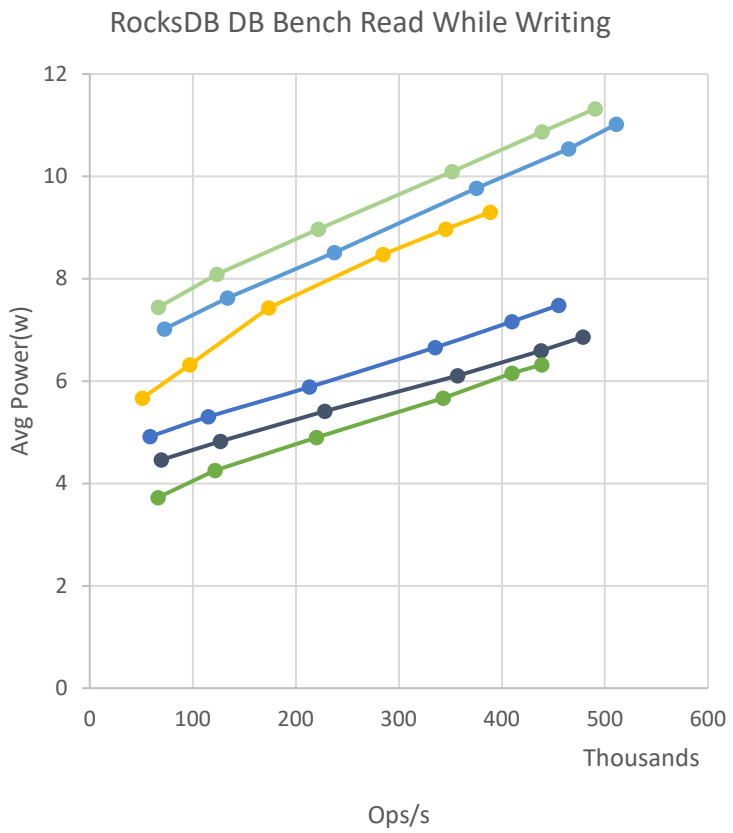

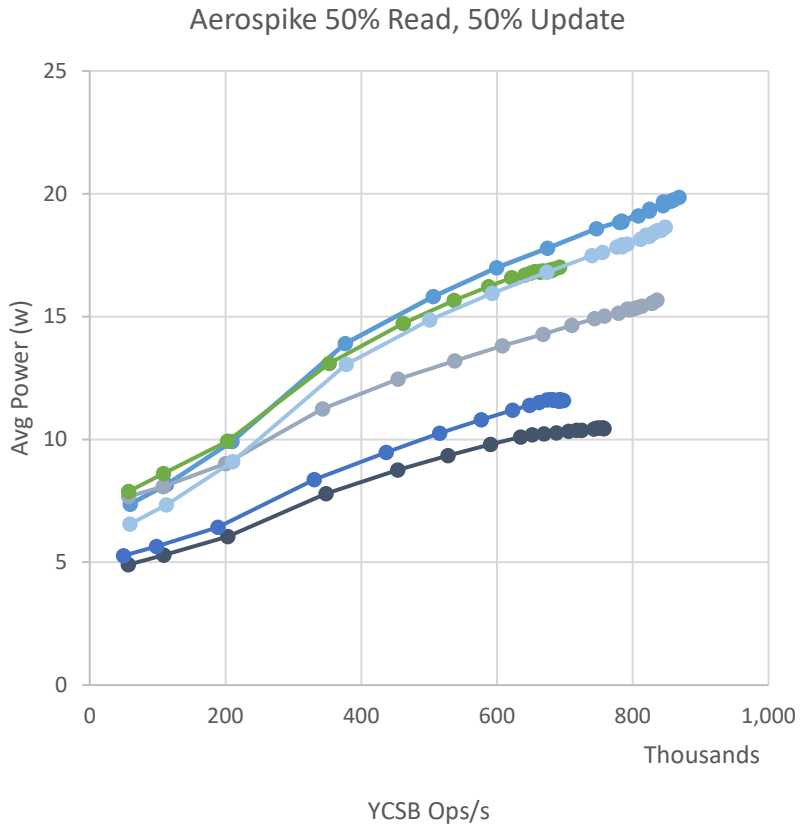

Thank You

For more power efficiency data in NoSQL databases, checkout

DCTR-304-1 – Enterprise Storage Part 2
Thursday August 8, 2024 @ 1:25pm
by: Ryan Meredith & Sayali Shirode

FMS the Future of Memory and Storage

Micron