Benefits of FDP on device power efficiency

August 6th, 2025

John Mazzie
MTS, Systems Performance Engineer

micron Intelligence Accelerated™

FMS the Future of Memory and Storage

Overview

- Why power is important
- Multi workload test results
- Uniform workload test results

Why power is important

- Power usage going up
  - GPUs, DPUs, CPU
  - Datacenters Power Usage
    - 4.4% of US Power in 2023, 176TWh
    - Projected 6.7 to 12% of US Power by 2028, 325 to 580TWh
- Saving power enables more efficient usage of other components
  - More performance in same power footprint
  - Same performance in lower power footprint
- Goal: Show power and relative benefits of FDP

Multiple Workloads

Separate microservices sharing a device

Multiple Distinct Workloads

- Workload Description
  - 8 Separate Jobs
  - Sequential Workloads ranging 4KiB to 128KiB
  - Non-FDP Workload
    - Distinct LBA ranges used for each job.
    - All jobs are writing to the same namespace in parallel
  - FDP Workload
    - 8 equal namespaces (1 reclaim group per namespace)
    - Each job is writing to different namespaces at the same time

Distinct Workloads – Write Amplification

- Measured at 1-minute intervals for 10 hours
- Workload is all overwrites, no discards sent.
- GC Handled by SSD
- WAF
  - Non-FDP: ~2.3 Average at Steady Date
  - FDP: ~1.05 at Steady State

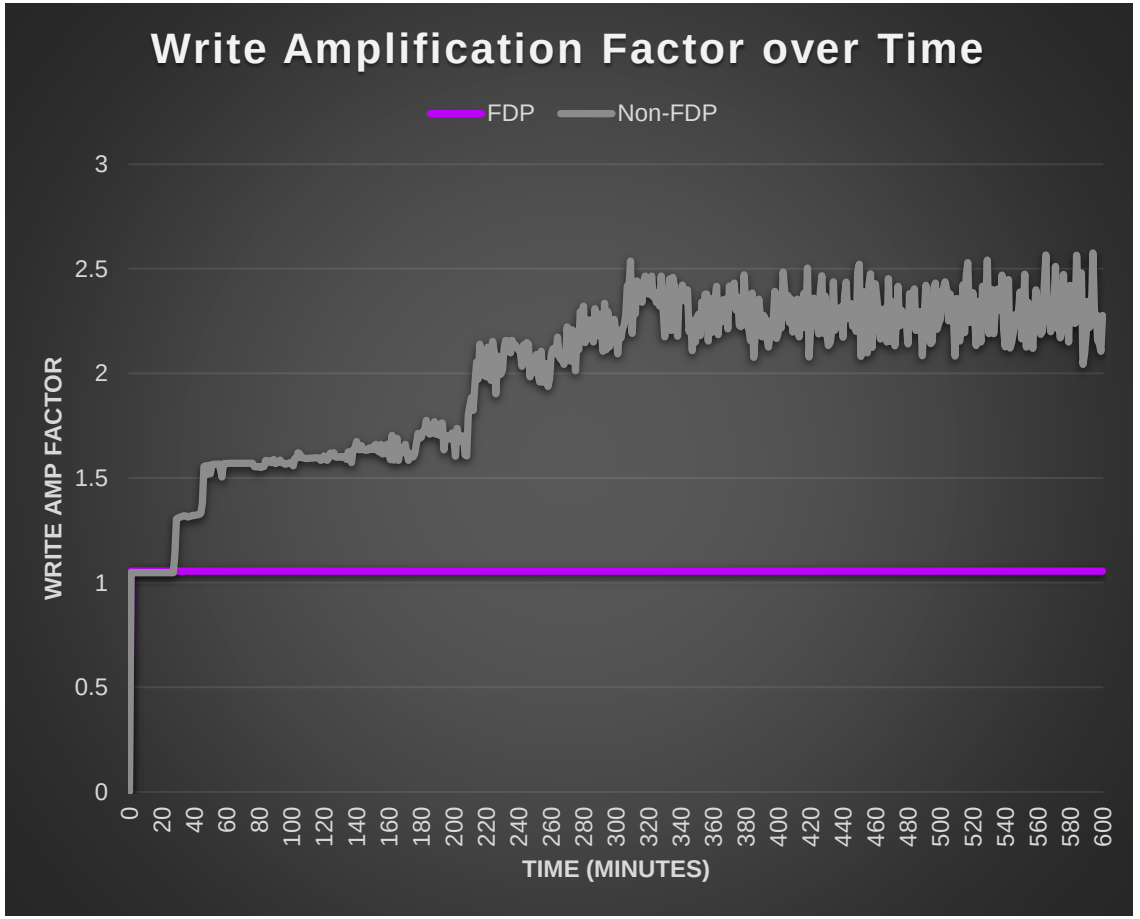

Distinct- Throughput and Power Results

- Throughput
  - Non-FDP: ~2000MiB/s
  - FDP: ~5000MiB/s
- Power
  - Non-FDP: 13.09W Average
  - FDP: 12.36W Average
- Throughput Efficiency
  - Non-FDP: ~155MiB/s per Watt
  - FDP: ~408MiB/s per Watt

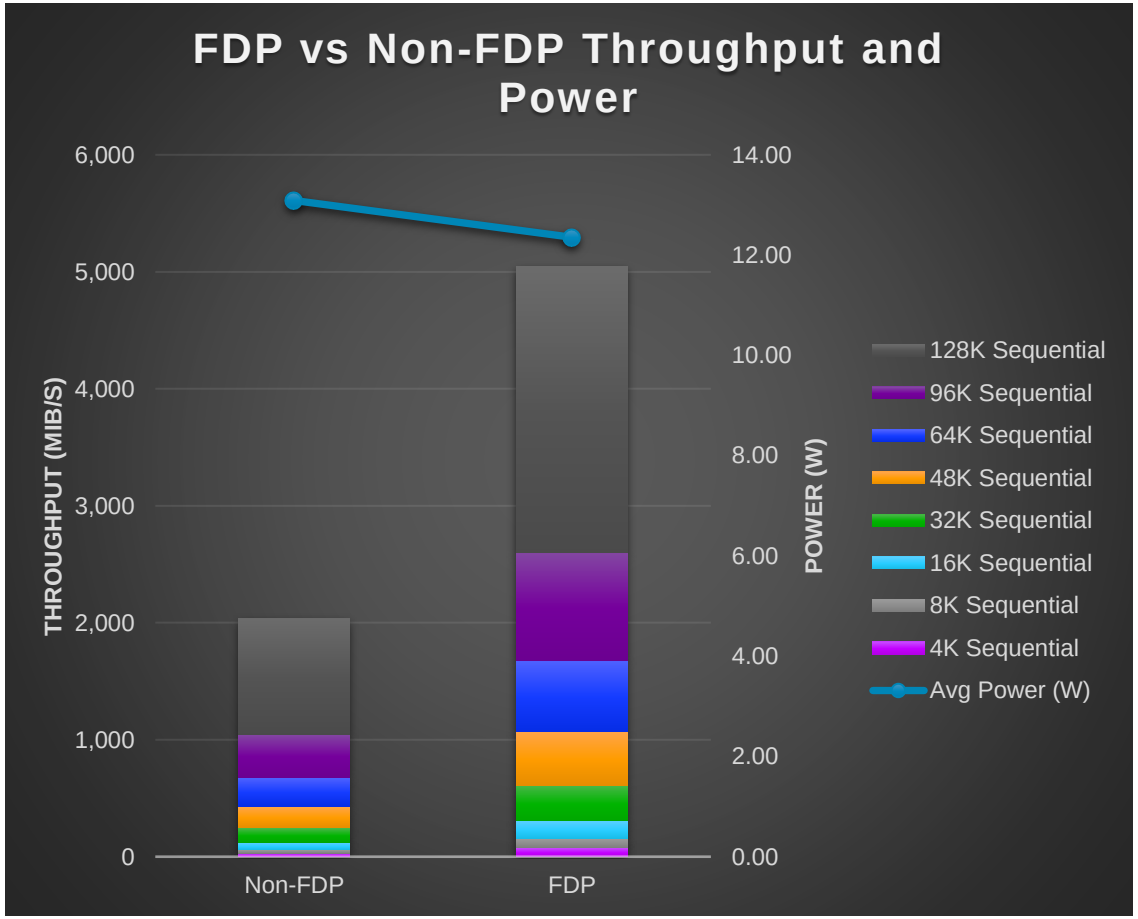

Uniform Workloads

GPU Checkpointing

Uniform Workloads

- Workload Description
  - 8 Separate Jobs
  - Sequential Workloads - 128KiB Block Size
  - Non-FDP Workload
    - Distinct LBA ranges used for each job.
    - All jobs are writing to the same namespace in parallel
  - FDP Workload
    - 8 equal namespaces (1 reclaim group per namespace)
    - Each job is writing to different namespaces at the same time

Uniform Workloads– Write Amplification

- Measured at 1-minute intervals for 10 hours
- WAF
  - Non-FDP: ~2.0 Average. Not consistent over 1 minute
  - FDP: ~1.05 at Steady State

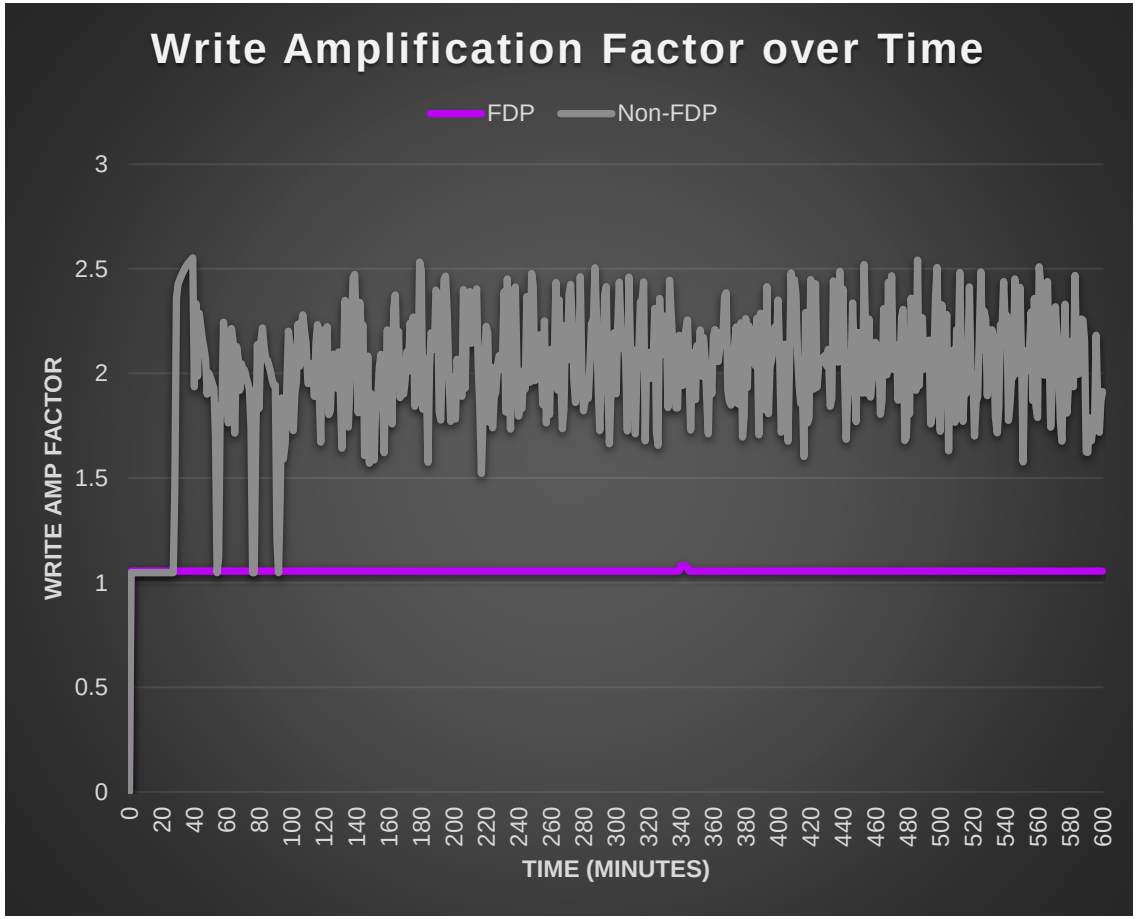

Uniform Workloads - Throughput and Power Results

- Throughput
  - Non-FDP: ~2000MiB/s
  - FDP: ~5000MiB/s
- Power
  - Non-FDP: 12.95W Average
  - FDP: 12.28W Average
- Throughput Efficiency
  - Non-FDP: ~157MiB/s per Watt
  - FDP: ~409MiB/s per Watt

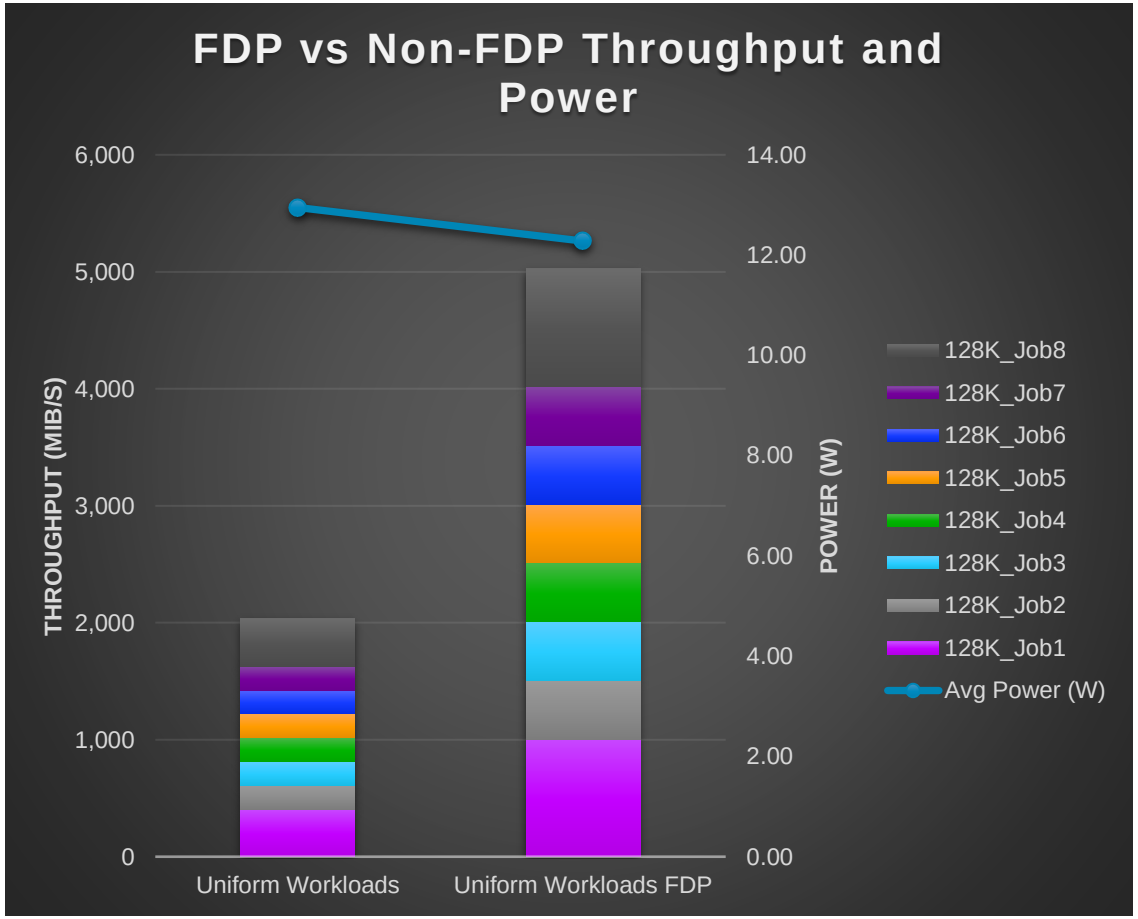

Power State Restriction

- Uniform Workload (Sequential 128K – 8 Jobs)
- Power Efficiency
  - Non FDP – 16W Power State
    - 16W Power State
      - ~2000 MiB/s, 12.94 W
  - FDP
    - 16W Power State
      - ~5000 MiB/s, 12.36 W
    - 10W Power State
      - ~1937 MiB/s, 6.47 W

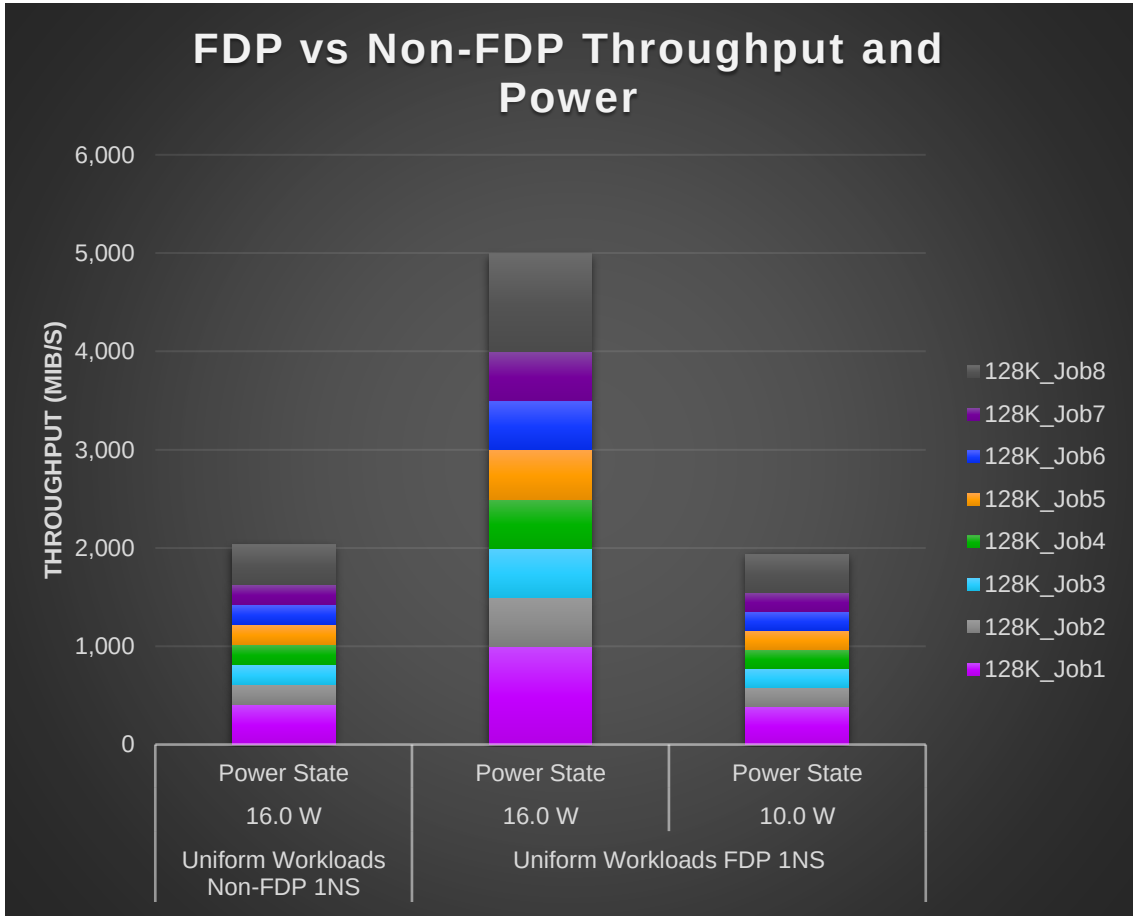

micron Intelligence Accelerated™

FMS the Future of Memory and Storage