FDP: Panel Introduction

Javier Gonzalez | Principal Engineer Samsung Semiconductor

FDP Status

Industry Adoption

- FDP is the de-facto data placement technology in NVMe SSDs
- Single SKU with backwards compatibility
  - Incremental adoption at deploy-time
- Simple host integration for legacy and NVMe-native applications on xPUs
  - Fast 80% WAF benefit in real deployments
  - Incremental engineering effort
- Benefits in TLC and QLC technologies
- Customer driven use-cases in DC and Entrp.
- Multi-vendor support

Ecosystem

Kernel v6.16

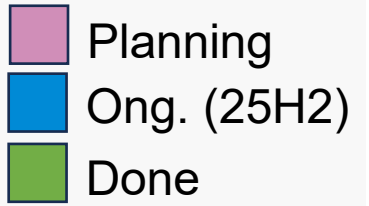

FDP Experimental Benefits

RocksDB

- Tuned SST hint. Upstreaming L4 to RocksDB mainline
- Same results in pre-upstream FS with FDP support

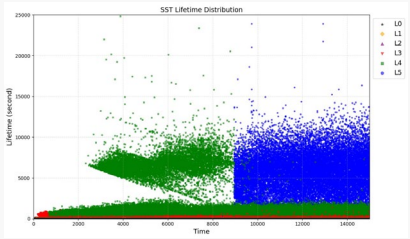

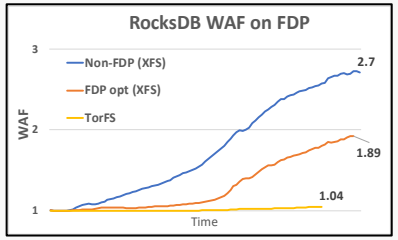

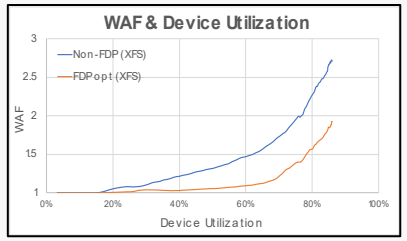

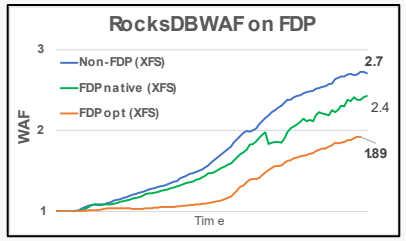

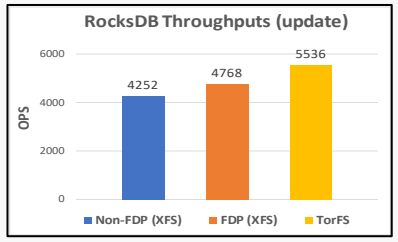

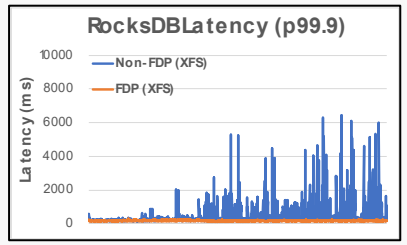

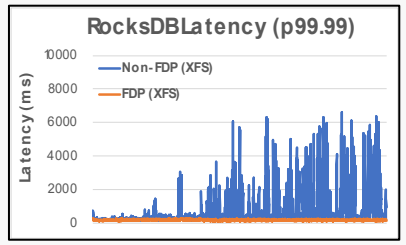

https://github.com/SamsungDS/TorFS

AI Training/Inference

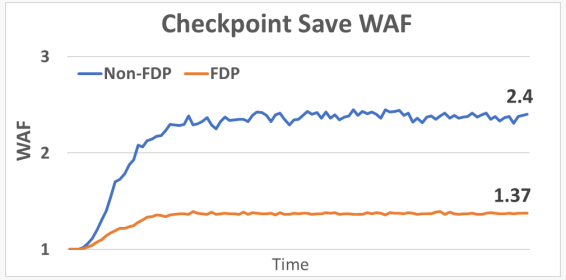

- O(TB) checkpoints
- 3 lifetimes

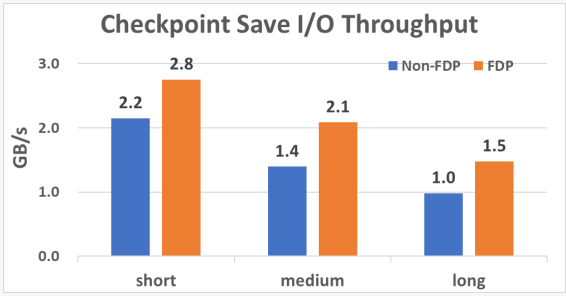

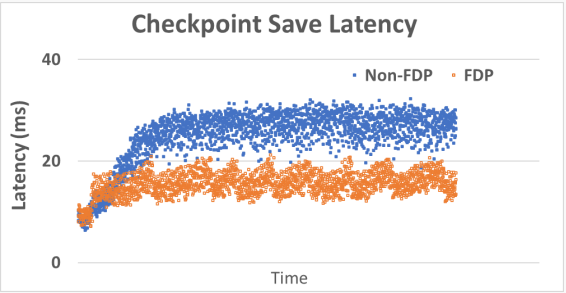

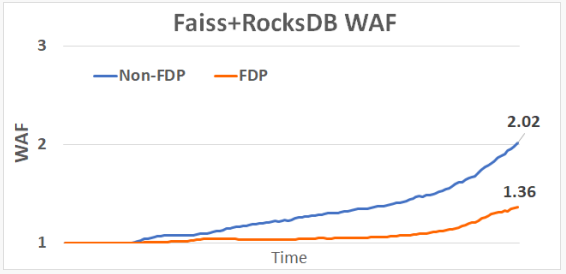

- RocksDB as backend storage to store vectorDB (Faiss) data

FDP for AI Workloads

FDP enables high-capacity SSD adoption for AI workloads

LLM Checkpointing	• Large dumps O(TB) • Different lifetimes simultaneously (e.g., fine tunning, recovery, versioning) • Transient optimizations vs stable checkpoints
Intermediate Result Store (e.g., Agent Memory	• Segregate ephemeral data (sessions) with semi-permanent (knowledge) • Trend to store popular queries
Data Labeling	• Segregate temporary staging data before moving to final datasets
Content Creation Pipelines	• Segregate intermediate and versions • Varying sizes on audio and video
KVCache memory expansion to SSD	• Embedding / RAG KVCache (e.g., LMCache) • Vector DBs with SSD caches (e.g., Milvus, Qdrant)

FDP: Panel Introduction

Javier Gonzalez | Principal Engineer Samsung Semiconductor