

SAMSUNG

Configuring FDP

Guidelines for Ensuring Consistency and Compatibility

Robert Moss

Principal SSD Architect, TERCAT
Samsung Semiconductor, Inc.



the Future of Memory and Storage

Overview

FDP Provides a High degree of flexibility

- **FDP protocol provides a host with the capability to configure and specify:**
 - How data should be initially isolated
 - How data should remain isolated during garbage collection
- **Backwards compatibility with non-FDP aware hosts**

Overview

SNIA Recommendations for FDP Configuration

- SNIA Storage Data Placement TWG has created a white paper on FDP which includes guidelines for using FDP
- By constraining FDP SSD configurations, interoperability and consistent performance may be achieved between host systems and SSD vendors
- These configuration recommendations help maintain performance for non-FDP aware hosts as well as FDP-aware systems with workloads which may not be fully optimized for FDP
- The following slides will cover the guidelines and why they are being recommended

SNIA FDP Configuration Recommendations

We recommend a single RG for the entire SSD storage capacity.

- **Background:**

- Each RUH becomes unique within an RG resulting in a multiplication of RUH resources
- Controller must isolate garbage collection within an RG

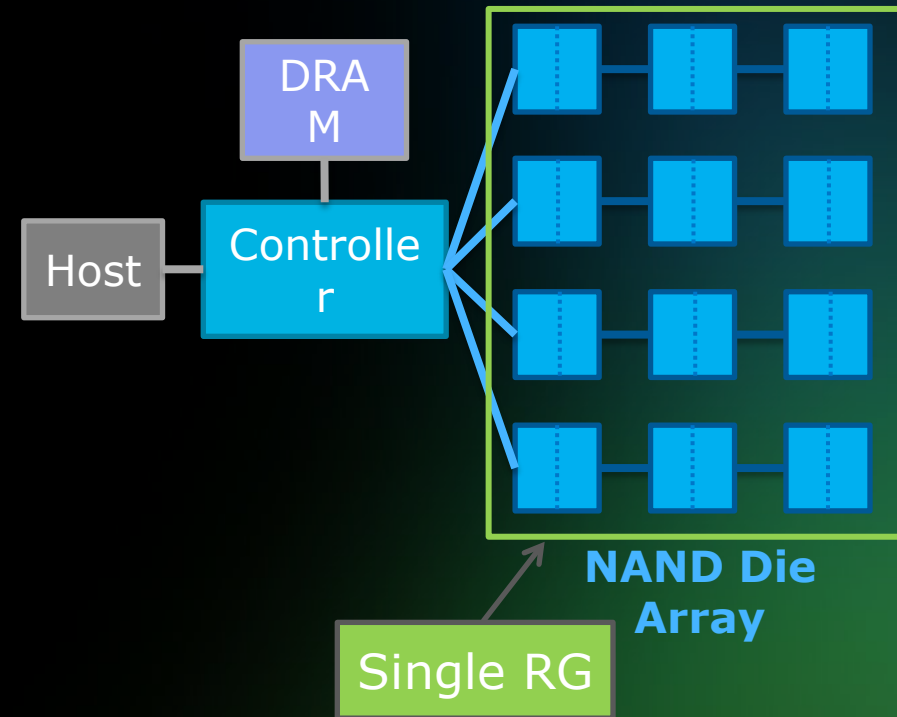
- **Reasoning:**

- Leverages NAND management algorithms that address performance and endurance
- Persistently isolated RUH can provide similar benefits to additional RGs
- Supporting multiple RGs reduces the number of RUHs that can be supported

SNIA FDP Configuration Recommendations

We recommend a single RG for the entire SSD storage capacity.

- One RG will include all of the SSD's storage capacity
- Leverages existing NAND management algorithms
- SSD can decide precise data placement within the RU that an RUH is filling
- No increased host complexity required for good system behavior



SNIA FDP Configuration Recommendations

We recommend setting the Reclaim Unit (RU) Size to match the superblock size of the SSD

- **Background:**

- NVMe defines an RU as a set of data that will be reclaimed together as a unit
- A Superblock in an SSD is a set of blocks across planes of a die where those die's are written and garbage collected together before being erased and reused
- Commonly a superblock may be composed of one erase block from every plane of every NAND die. This allows for parallelism for maximum write performance from an SSD.

- **Reasoning:**

- RU size is directly proportional to write performance for a given RUH
- Small RU size requires hosts to manage parallelism across multiple RUHs

SNIA FDP Configuration Recommendations


We recommend setting the Reclaim Unit (RU) Size to match the superblock size of the SSD

- SNIA FDP TWG does not provide any recommendation for superblock size, however note that there are impacts to system WAF and write performance in setting superblock size
- Small superblocks require the host to manage parallelism to achieve write performance across the die array
- Small superblocks restrict the capability of an SSD to control parallelism
- Appropriate selection of superblock size can reduce system WAF due to the host extent not being aligned to RU size

SNIA FDP Configuration Recommendations

We recommend setting the Reclaim Unit (RU) Size to match the superblock size of the SSD

- File Systems are built around Host Extents
- Host Extent is unlikely to match SSD RU size
- Reasons Host Extent may not match SSD RU
 - Vendor-to-Vendor mismatch
 - Generation over Generation SSD RU changes



SNIA FDP Configuration Recommendations

We recommend limiting the number of RUHs

- **Background:**

- Limited by controller resources (e.g., isolated stream support, open superblocks, open superblock time, and write buffering capacity)
- Different NAND technologies require different amounts of write buffer for programming requirements

- **Reasoning:**

- Current products support a maximum of 16-32 RUHs; with no significant trends towards higher numbers
- Smaller capacity, lower performance, and QLC SSDs may support only up to 8 RUH

SNIA FDP Configuration Recommendations

We recommend that the SSD and not the host control physical placement

Background :

- NVMe specification does not define mechanisms for associating an RG or an RUH with specific die
- NVMe specification only defines Namespaces as having a capacity

• Reasoning :

- Allowing the SSD to select physical placement maximizes legacy host SW infrastructure
- Permits the SSD to optimize performance & endurance decisions, rather than the host managing write & sequential read performance, channel & die conflicts, and endurance
- The lack of standardization for a host writing more to an RG than a set of die's physical capacity will result in inconsistent behavior between vendors

Summary

SNIA FDP Configuration Recommendations

- **Single RG for entire SSD storage capacity**
- **RU size equal to superblock size**
- **Limit number of RUHs to no more than 8-32**
- **No assumptions on physical placement constraints**

Thank You

SAMSUNG