STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

NVM Express® TP4146a Flexible Data Placement

Overview

Presented by

Mike Allison (Samsung) and John Rudelic (Solidigm)

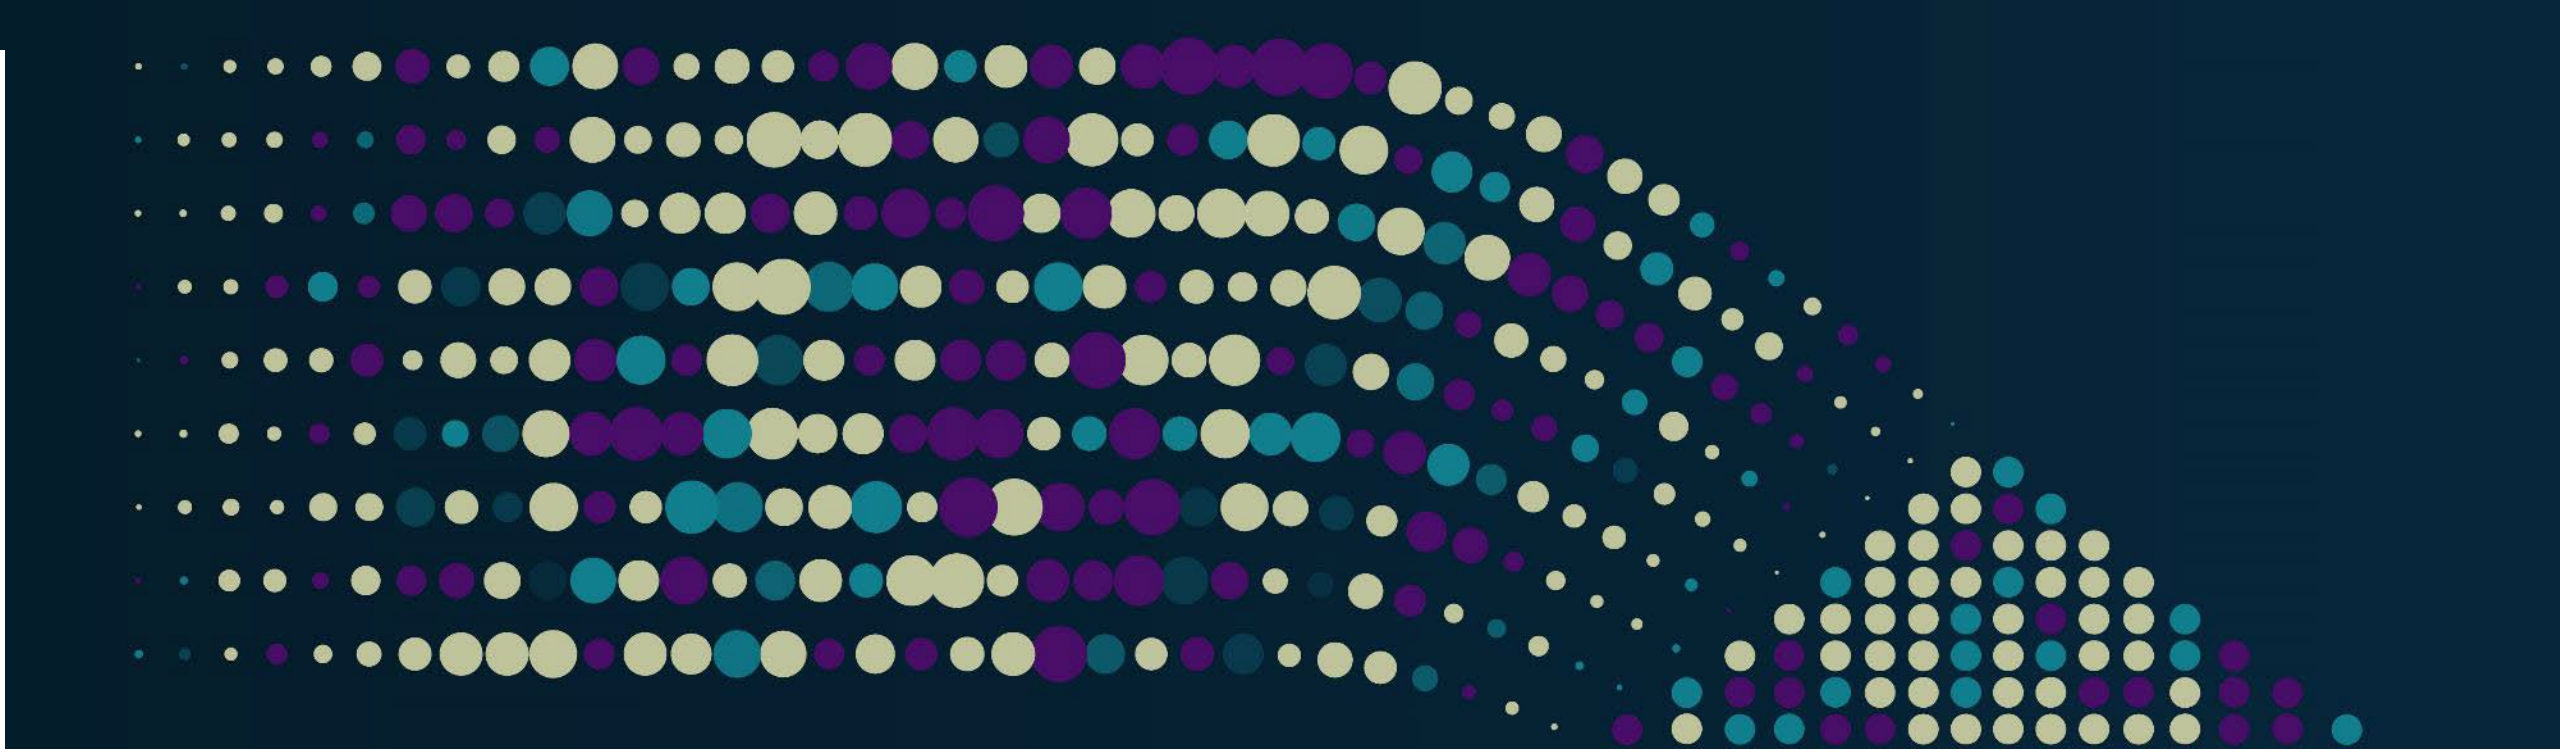

FDP Architecture

SAMSUNG SOLIDIGM SDC 23

NVMe™ Storage Entities

- An FDP configuration consists of:
  - One or more Reclaim Units (RUs)
  - One or more Reclaim Groups (RGs)
  - One or more Reclaim Unit Handles (RUHs) that reference to a Reclaim Unit in each RG
- An Endurance Group may supports one or more FDP configurations
- Write commands specify which RU to write the logical blocks by specifying:
  - An RUH
  - An RG

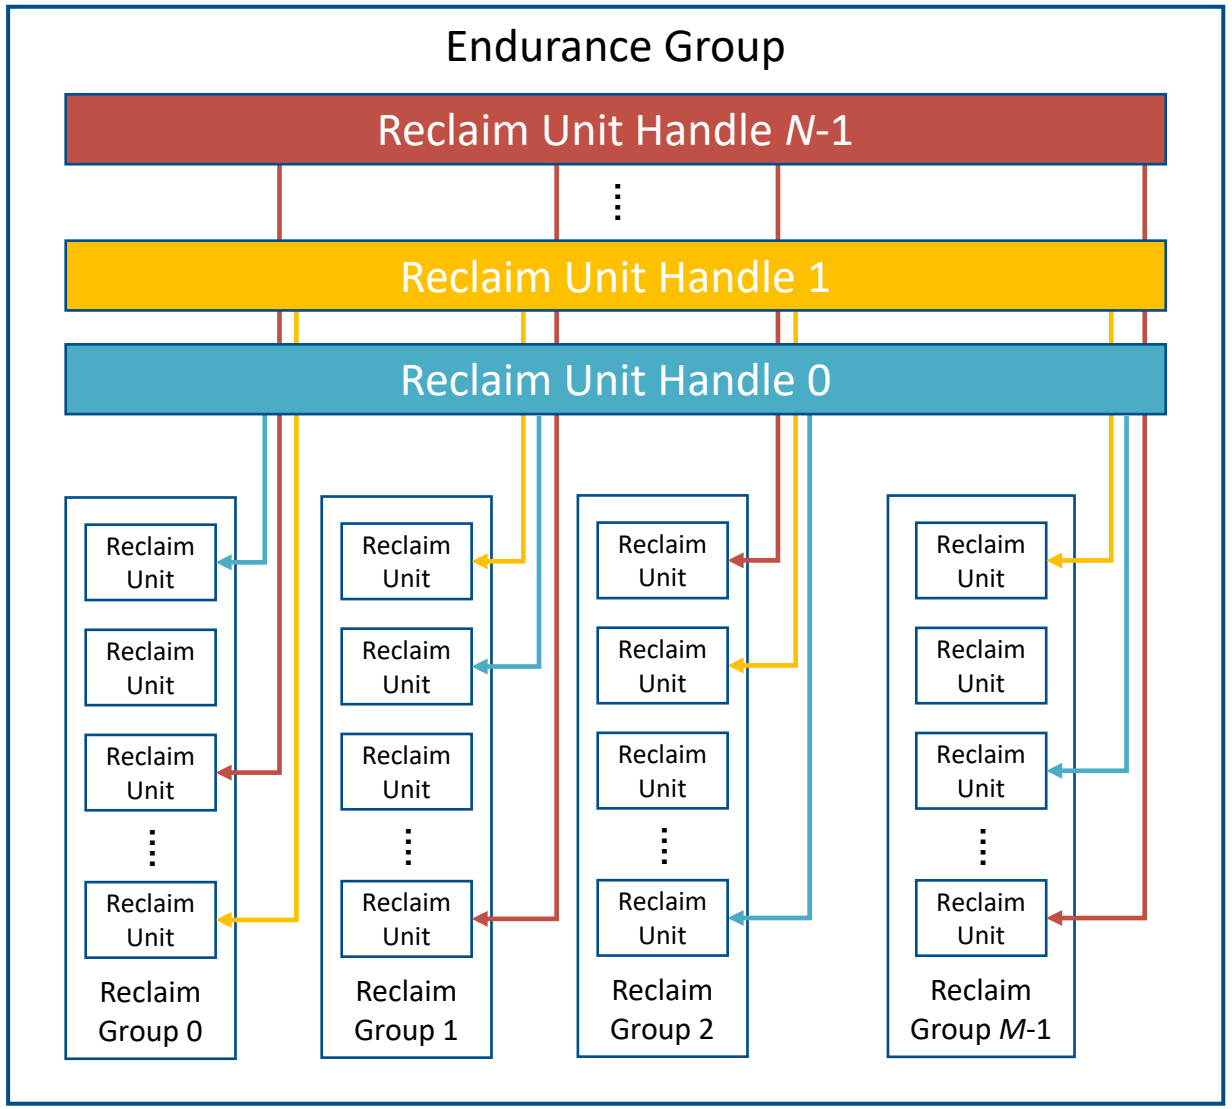

SAMSUNG SOLIDIGM SDC 23

Namespace Creation

- Namespace creation using the Namespace Management command
  - Specify the Endurance Group with FDP enabled
  - Host may provide a Placement Handle List

Placement Handle	Reclaim Unit Handle Identifier
0	1
1	0
2	*N*-1

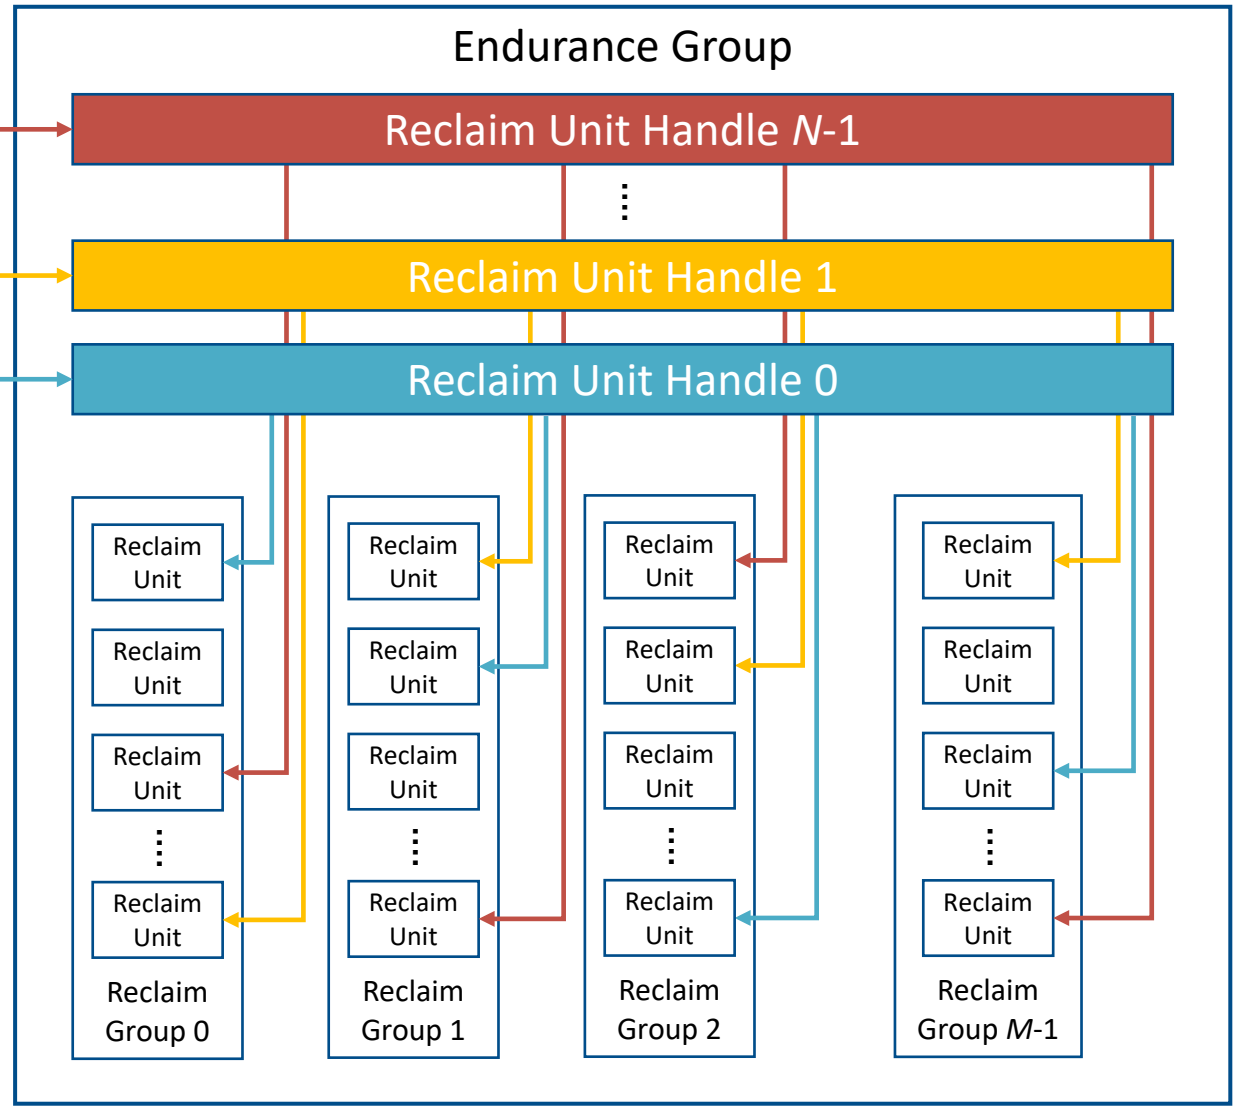

SAMSUNG SOLIDIGM SDC 23

Log Pages

FDP Log Pages
FDP Configurations
Reclaim Unit Handle Usage
FDP Statistics
FDP Events

SAMSUNG SOLIDIGM SDC 23

Log Pages – FDP Configurations

FDP Configuration Log Page
#FDP Configurations (N)
Version
Size
FDP Config Descriptor 0 (FDPCD0)
FDP Config Descriptor 1 (FDPCD1)
…
FDP Config Descriptor N (FDPCDN)

FDP Configuration Descriptor
Descriptor Size
FDP Attributes
Vendor Specific Size (VSS)
RGs (NRG)
RU Handles (NRUH)
MAX Placement IDSs (MAXPIDS)
Namespaces supported
Reclaim Unit Nominal Size (RUNS)
Estimated RU Time Limit (ERUTL)
RUH Descriptor List (RUHD0)
RUH Descriptor List (RUHD1)
RUH Descriptor List (RUHDN)
Vendor specific

FDP Attributes
FDP Config Valid
FDP Volatile Write Cache (FDPVWC)
RG ID Format (RGIF)

RUH Descriptor
Handle type (initial or persistent)
Vendor specific

SAMSUNG SOLIDIGM SDC 23

Log Pages – FDP Reclaim Unit Handle Usage

FDP Reclaim Unit Handle Usage
#Reclaim Unit Handles (NRUH)
RUH Usage Descriptor 0 (RUHUD0)
RUH Usage Descriptor 1 (RUHUD1)
…
RUH Usage Descriptor N (RUHUDN)

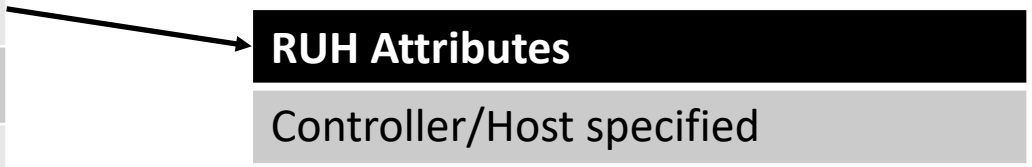

RUH Attributes
Controller/Host specified

SAMSUNG SOLIDIGM SDC 23

Log Pages – FDP Statistics

FDP Statistics Log Page
Host Bytes with Metadata Written (HBMW)
Media Bytes with Metadata Written (MBMW)
Media Bytes Erased (MBE)

SAMSUNG SOLIDIGM SDC 23

Log Pages – FDP Events

FDP Events Log Page
#FDP Events (N)
FDP Event 1
FDP Event 2
…
FDP Event N

FDP Event
FDP Event Type
FDP Event Flags (FDPEF)
Placement Identifier (PID)
Event Timestamp
Namespace Identifier (NSID)
Event Type Specific
Reclaim Group Identifier
Reclaim Unit Handle Identifier
Vendor Specific Info

FDP Event Type
HOST EVENTS
RU not fully written to capacity
RU active3 time limit exceeded
Controller level reset – modified RUH
Invalid PID
Vendor specific
CONTROLLER EVENTS
Media reallocated
Implicitly modified RUH
Vendor Specific

FDP Event Flags (FDPEF)
Location Valid (LV)
NSID Valid (NSIDV)
Placement ID (PIV)

SAMSUNG SOLIDIGM SDC 23

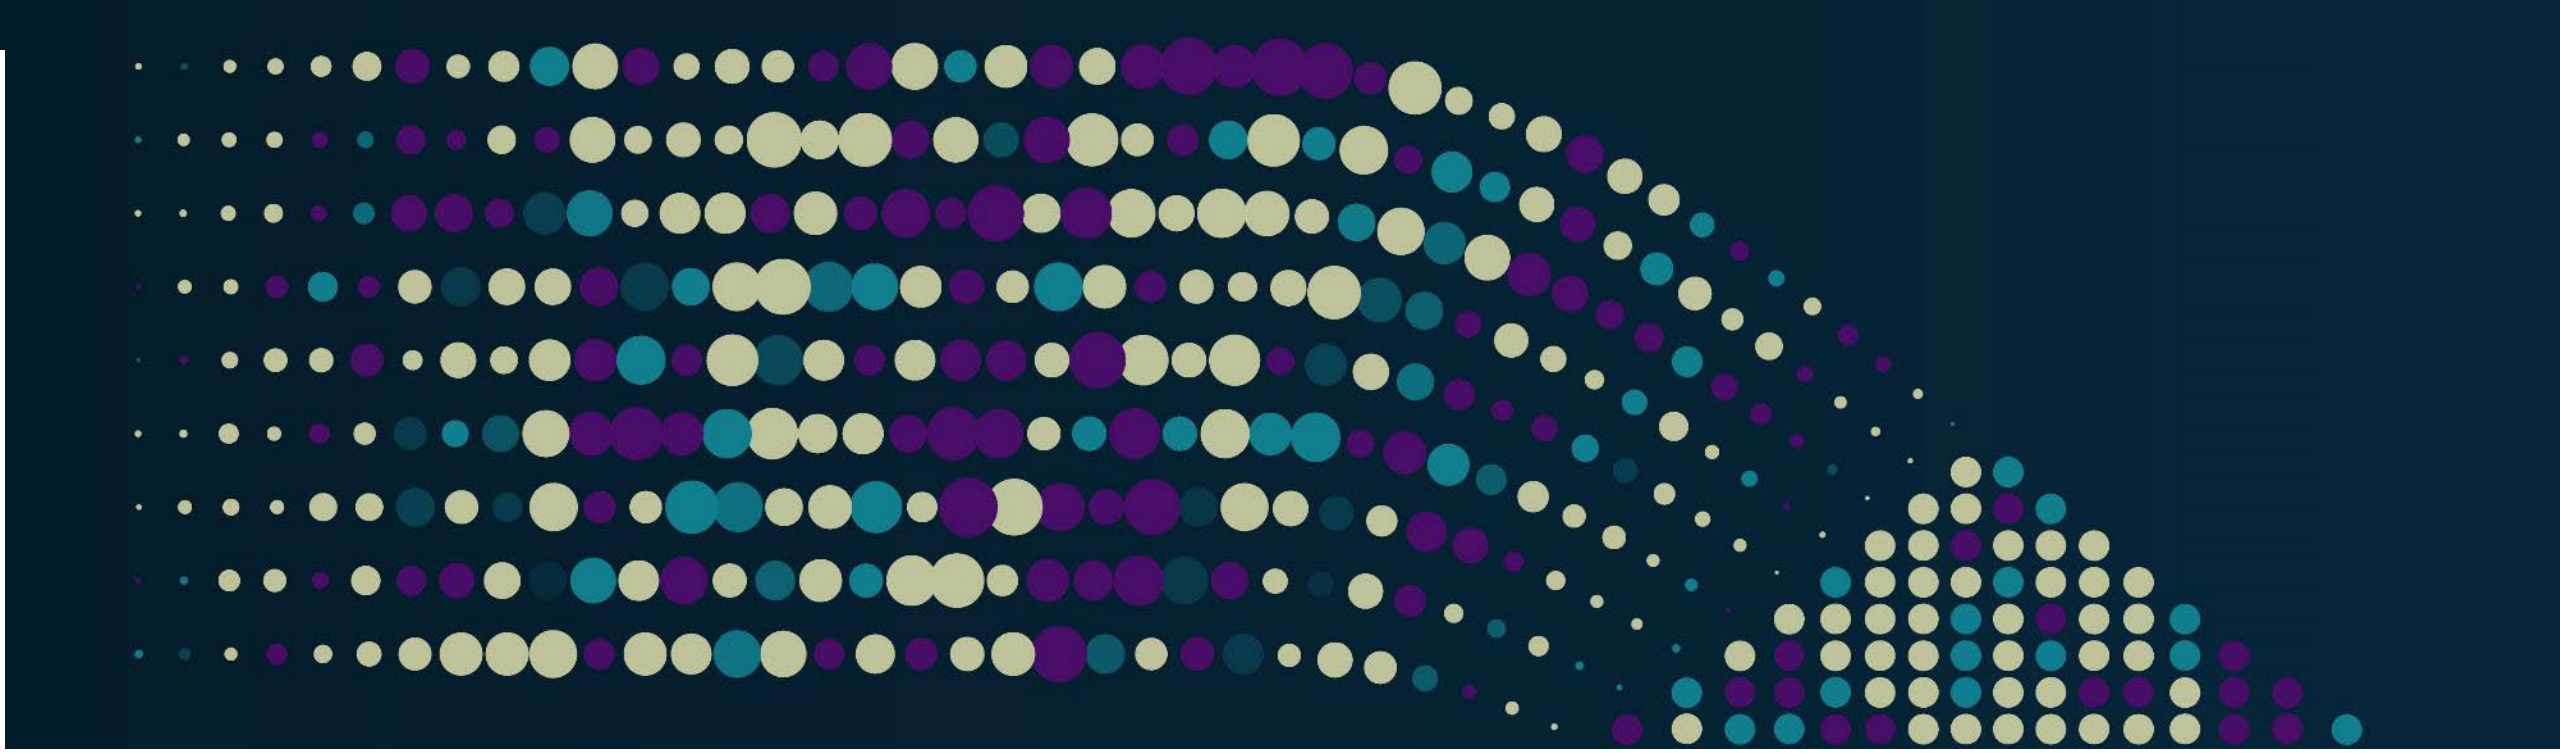

Configuring & Monitoring

SAMSUNG SOLIDIGM SDC 23

Configuring FDP - Pseudocode

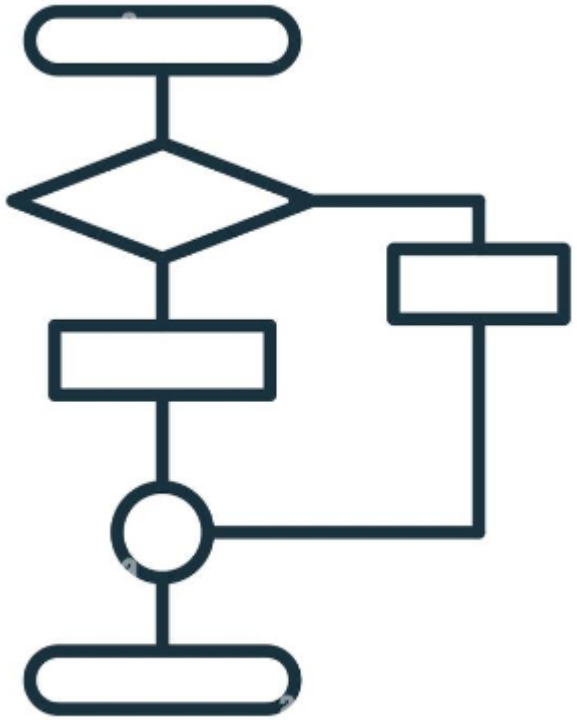

```
#Check NVMe Identify Controller Data Structure
if (FDP Support Bit = 1)  #FDP Supported
        #Configure FDP
        #Read FDP Log Page for FDP configurations
        #Set feature – enable FDP configuration
        #(Optional) Configure Placement handle list
        #(Optional) Enable FDP events
        #Set feature – enable FDP set FDPE=1
else #FDP Not supported
```


SAMSUNG SOLIDIGM SDC 23

Monitoring FDP - Pseudocode

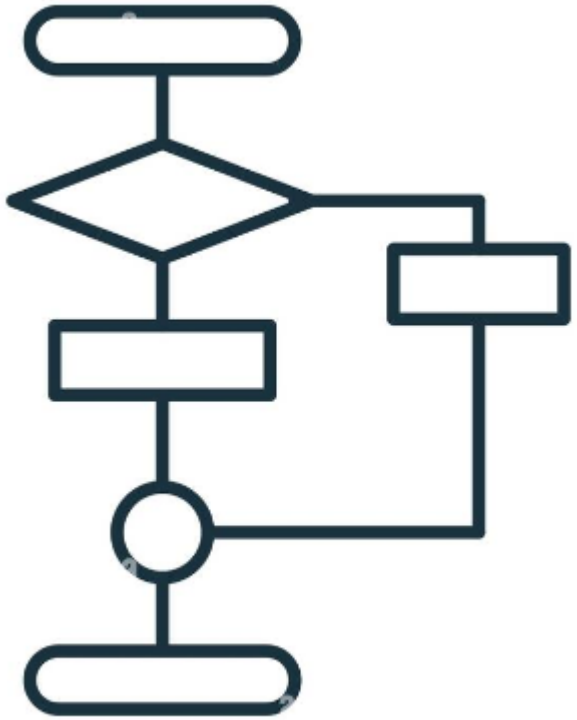

```
#FDP Monitoring Loop
#Check FDP Events
      #respond to event
#Check FDP RUH
      #respond to RUH
#Check FDP Statistics
      #respond to Statistics
```


SAMSUNG SOLIDIGM SDC 23

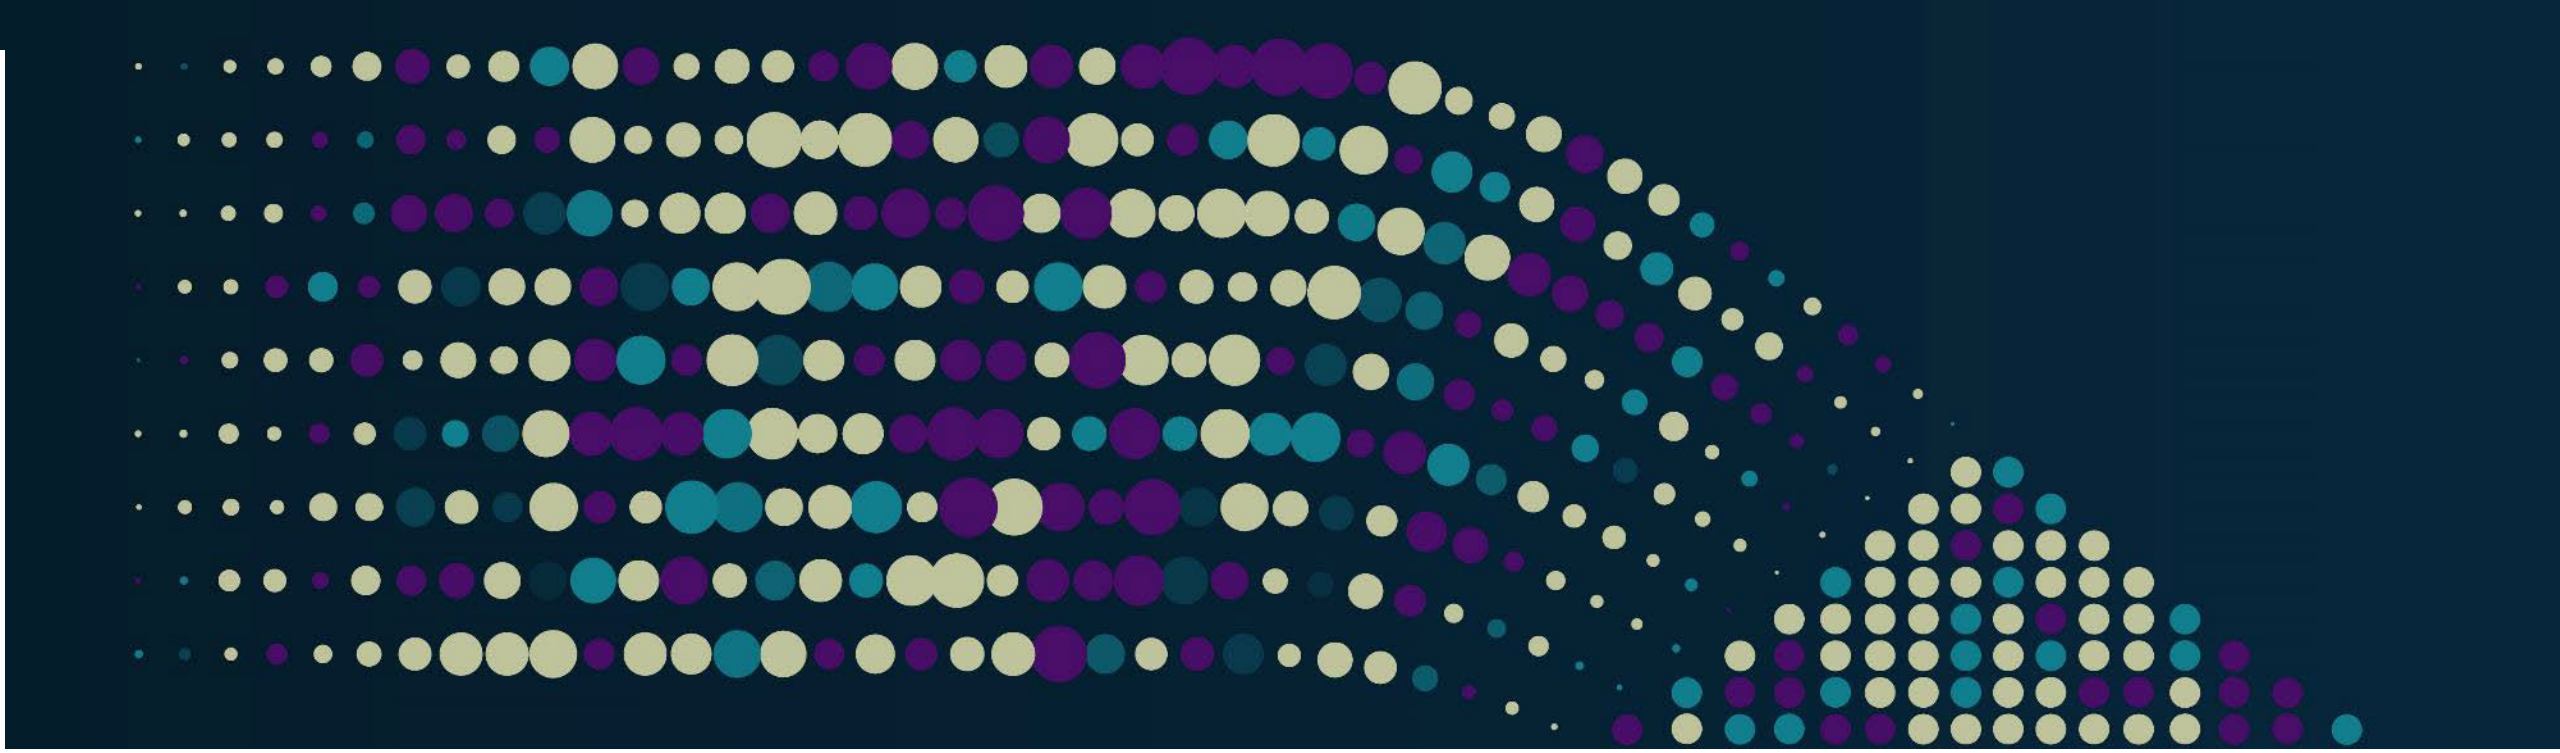

FDP Writing Command Examples

SAMSUNG SOLIDIGM SDC 23

Write to a Reclaim Unit

Host issues a Write command to Namespace A specifying Placement Handle 1 and Reclaim Group 0

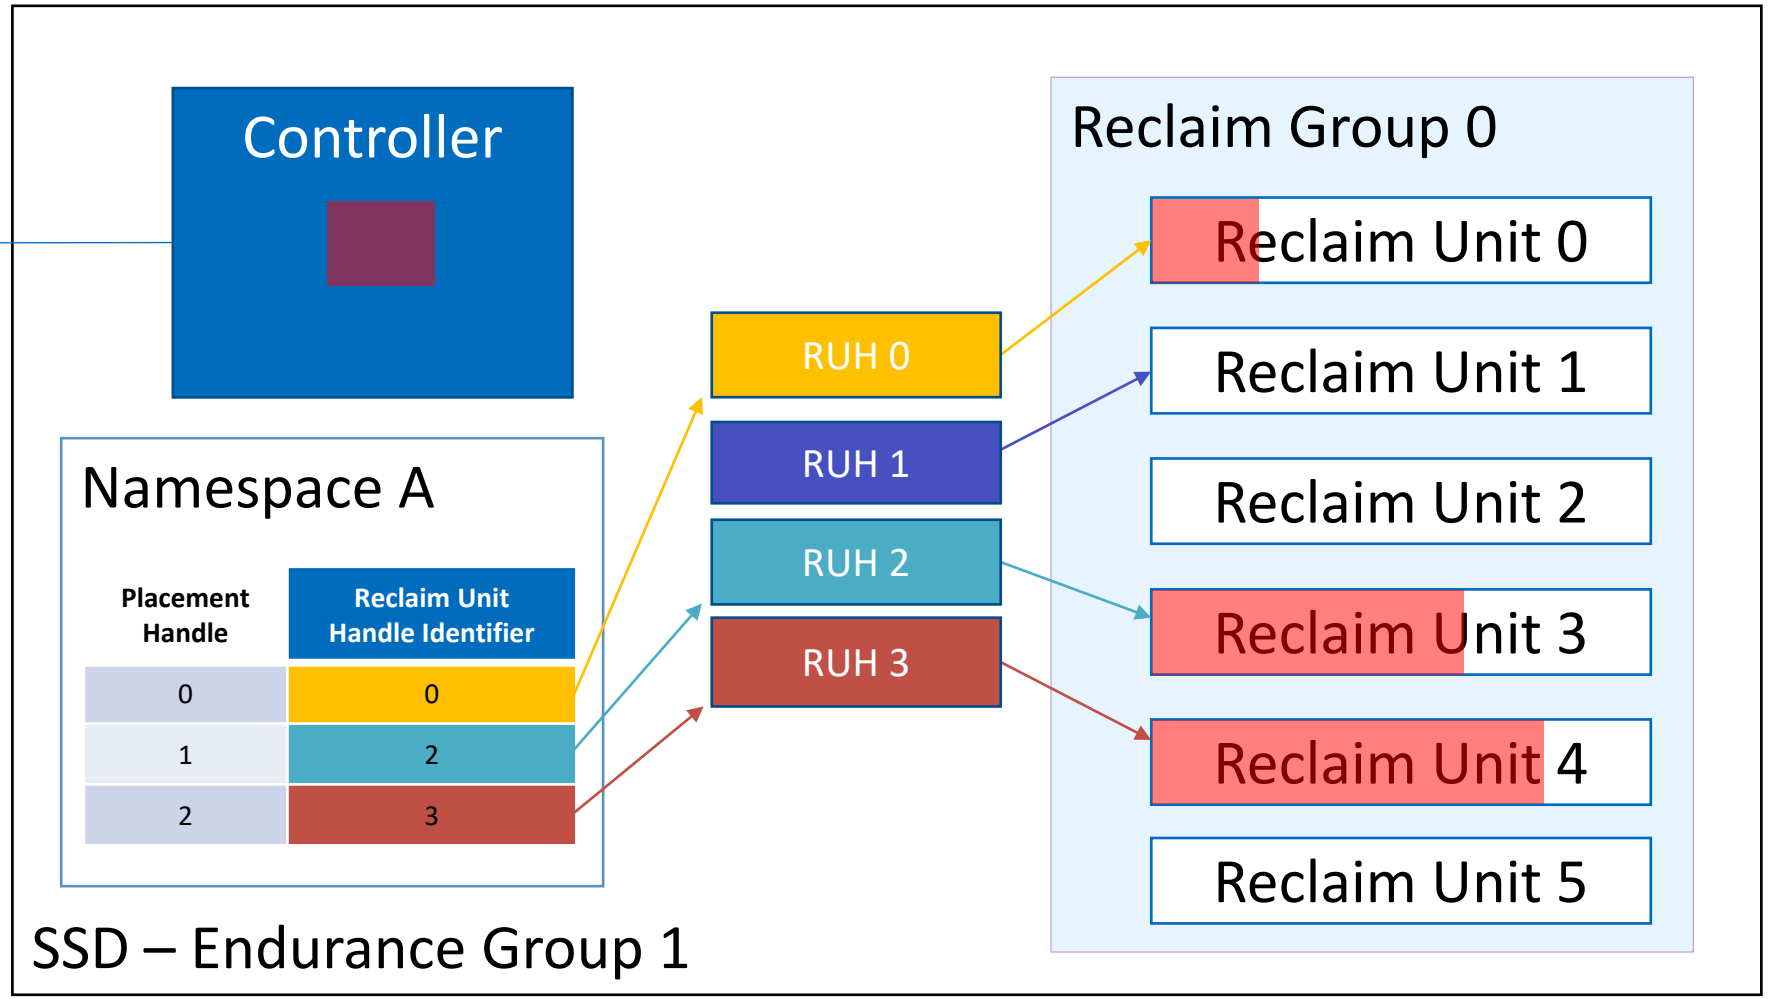

SAMSUNG SOLIDIGM SDC 23

Write to a Reclaim Unit

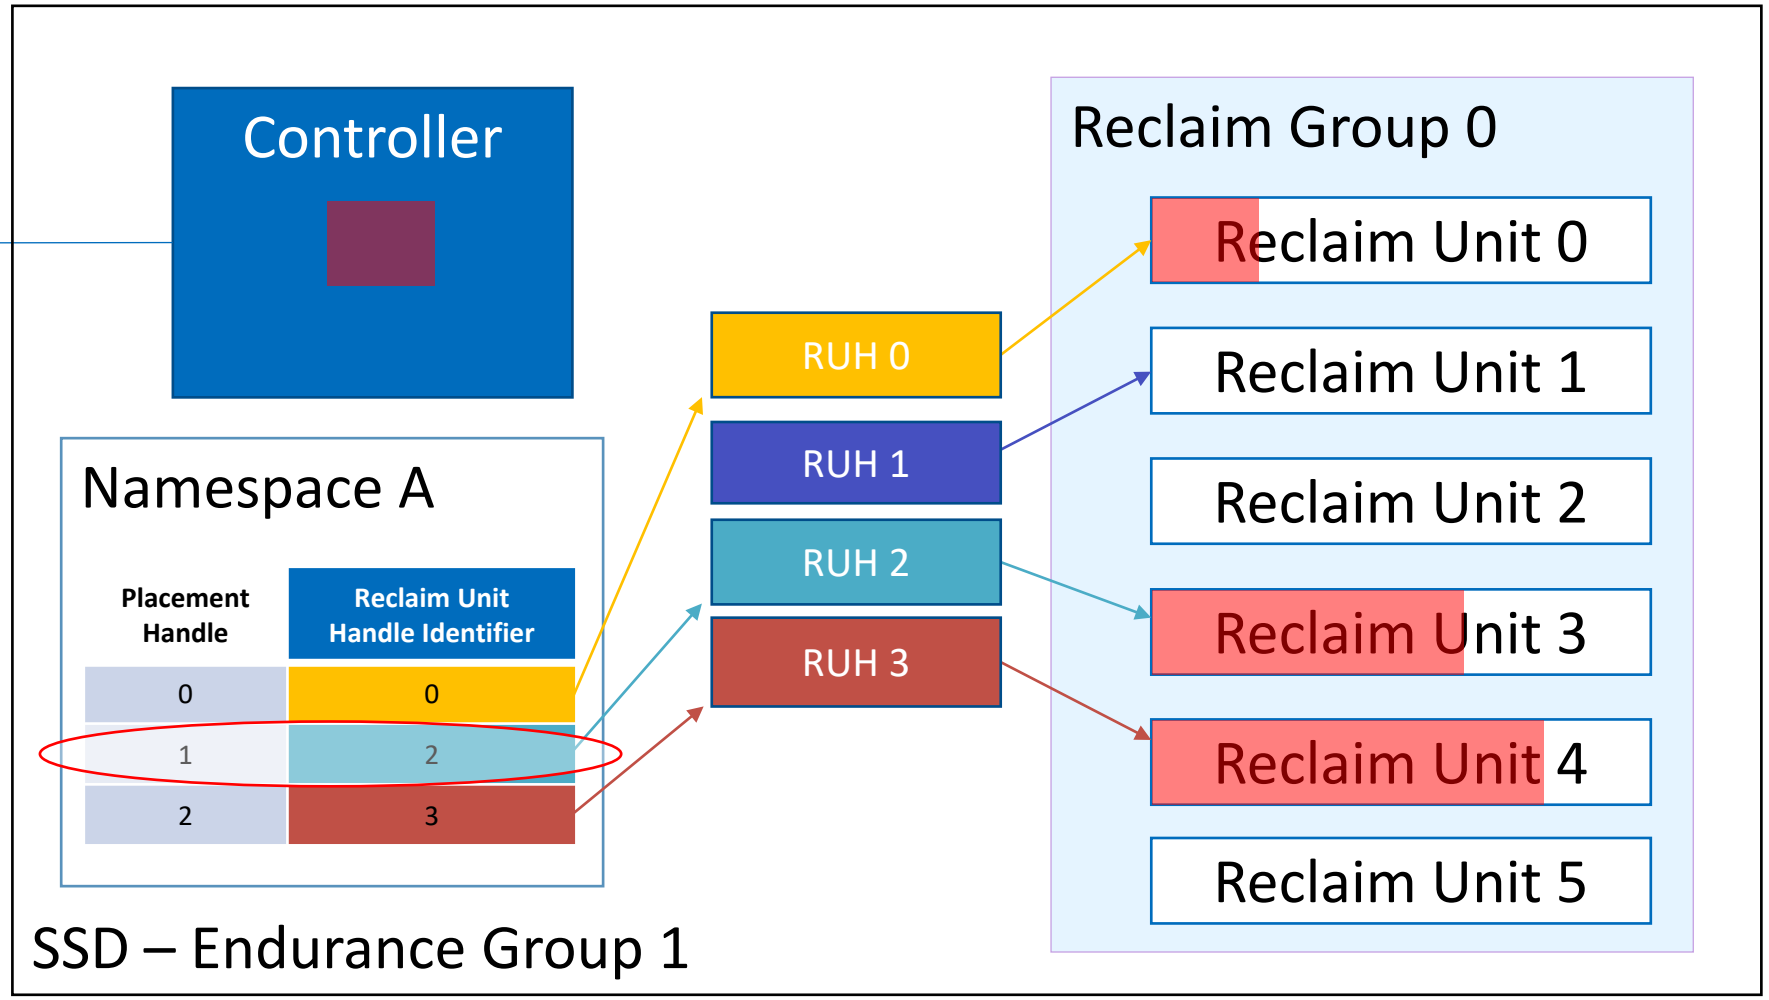

Controller looks up Placement Handle 1 in Namespace A and determines to place the data using Reclaim Unit Handle 2 on Reclaim Group 0

SAMSUNG SOLIDIGM SDC 23

Write to a Reclaim Unit

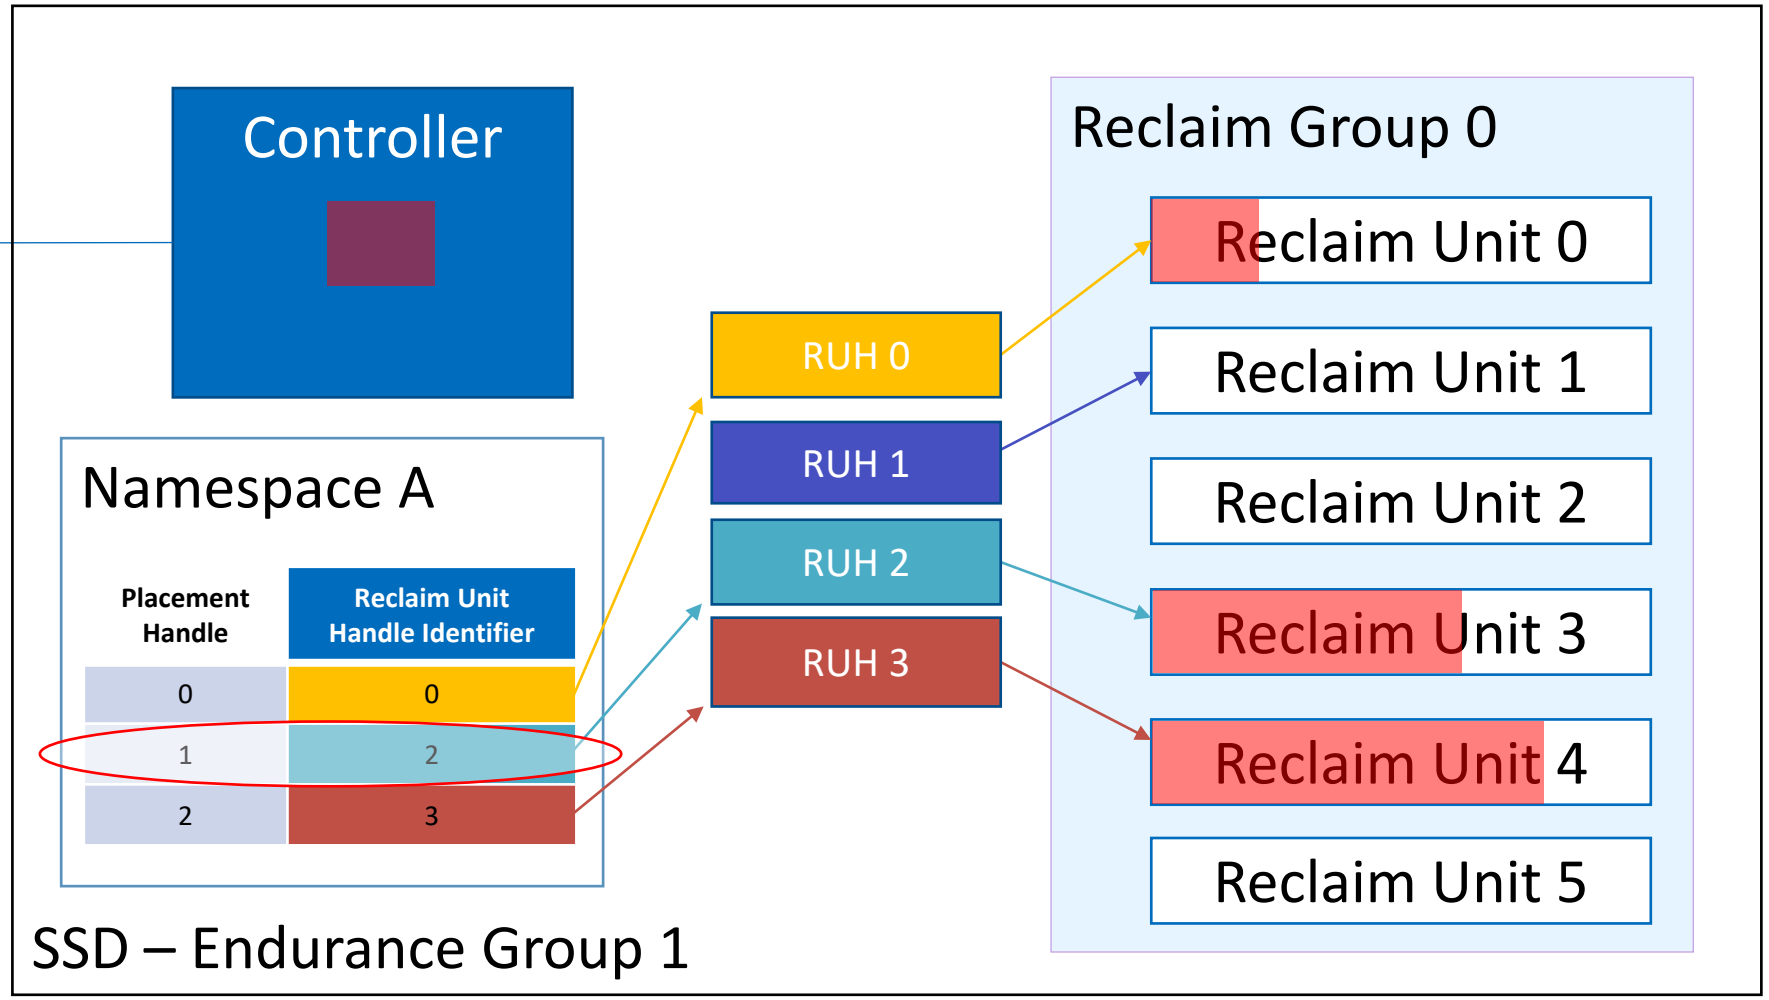

Controller looks up Placement Handle 1 in Namespace A and determines to place the data using Reclaim Unit Handle 2 on Reclaim Group 0

SAMSUNG SOLIDIGM SDC 23

Write to a Reclaim Unit

Controller uses Reclaim Unit Handle 2 to place the data in Reclaim Unit 3 of Reclaim Group 0

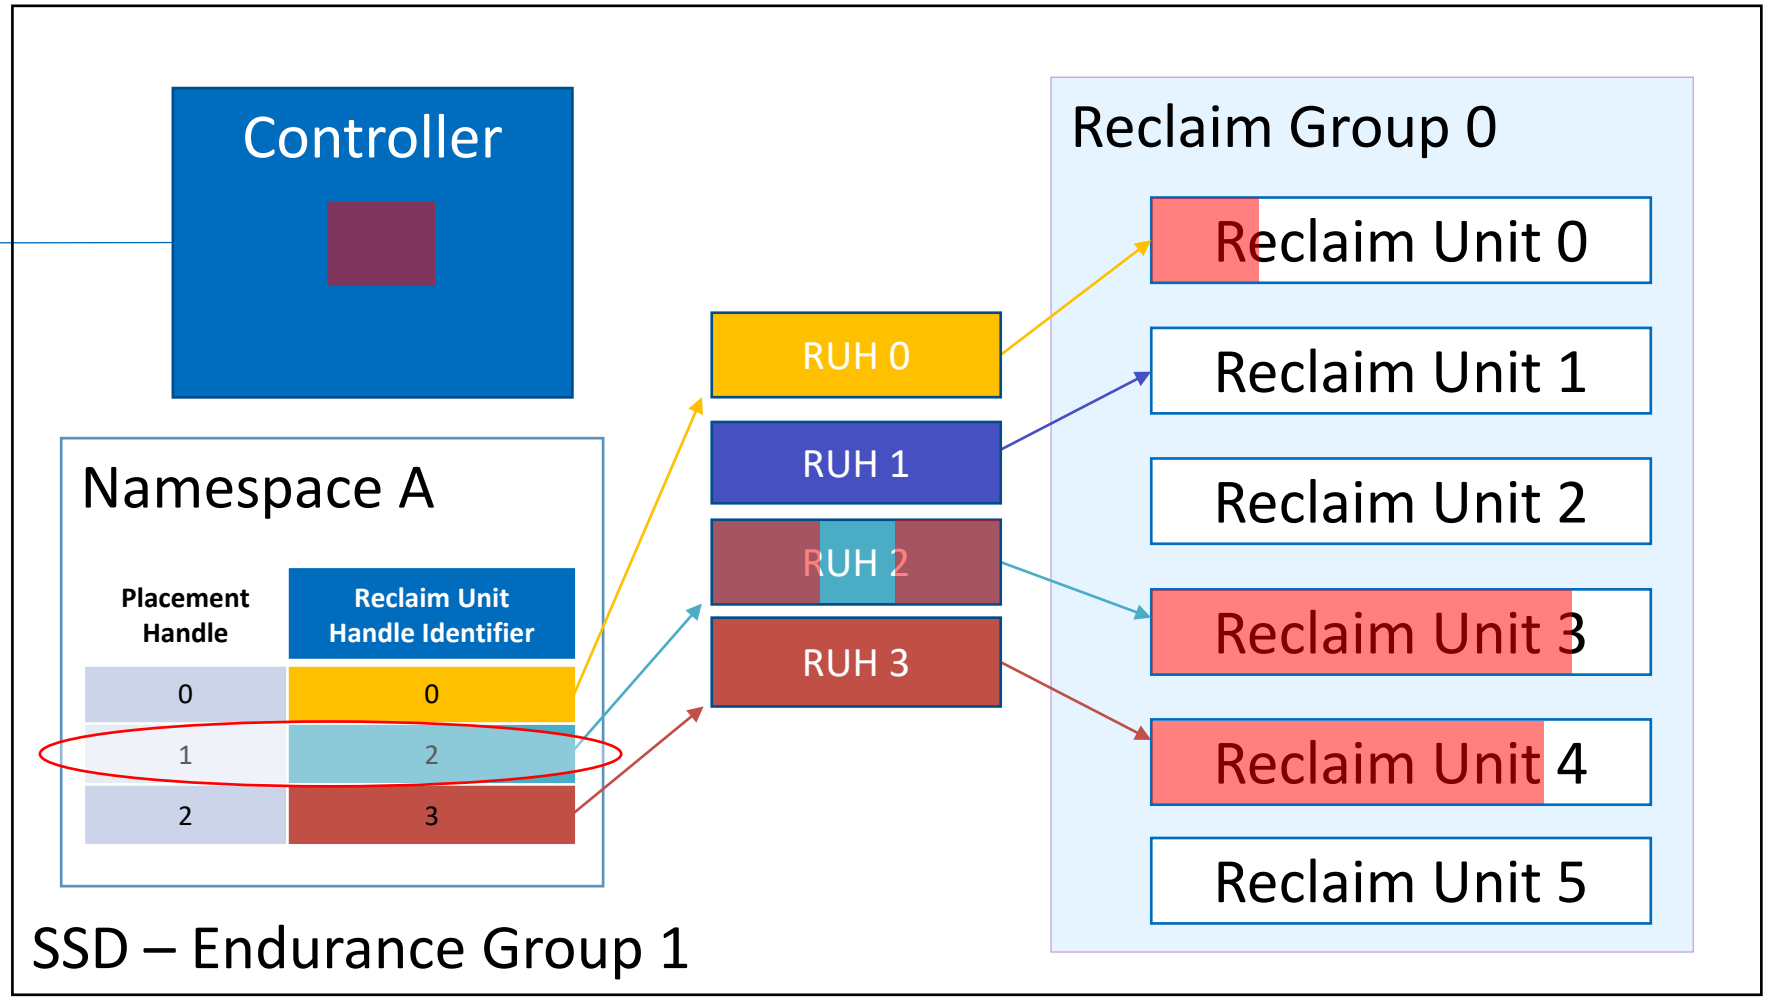

Write to fill the capacity of a Reclaim Unit

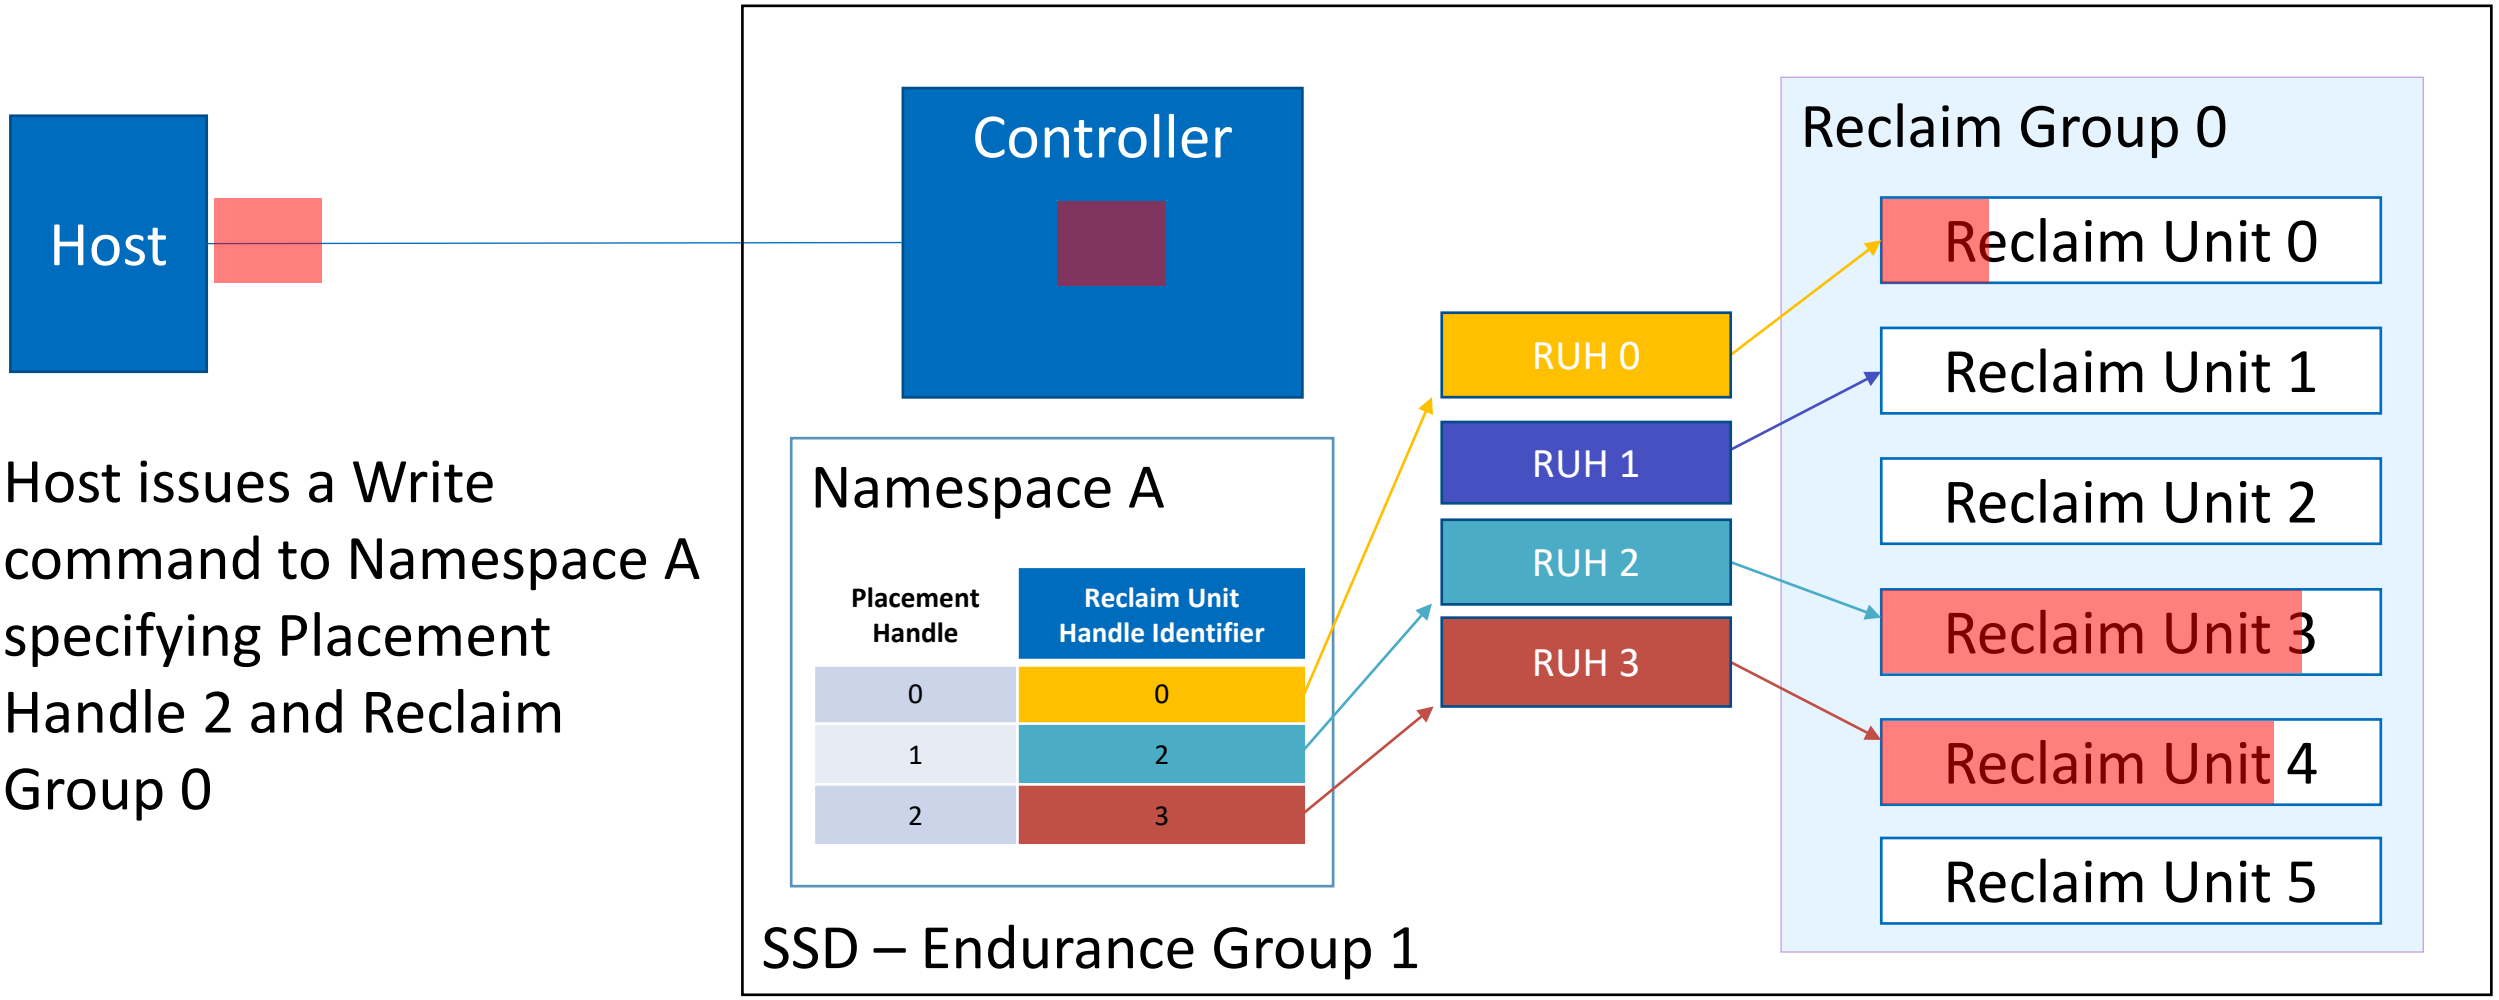

Host issues a Write command to Namespace A specifying Placement Handle 2 and Reclaim Group 0

SAMSUNG SOLIDIGM SDC 23

Write to fill the capacity of a Reclaim Unit

Controller looks up Placement Handle 2 in Namespace A and determines to place the data using Reclaim Unit Handle 3 on Reclaim Group 0

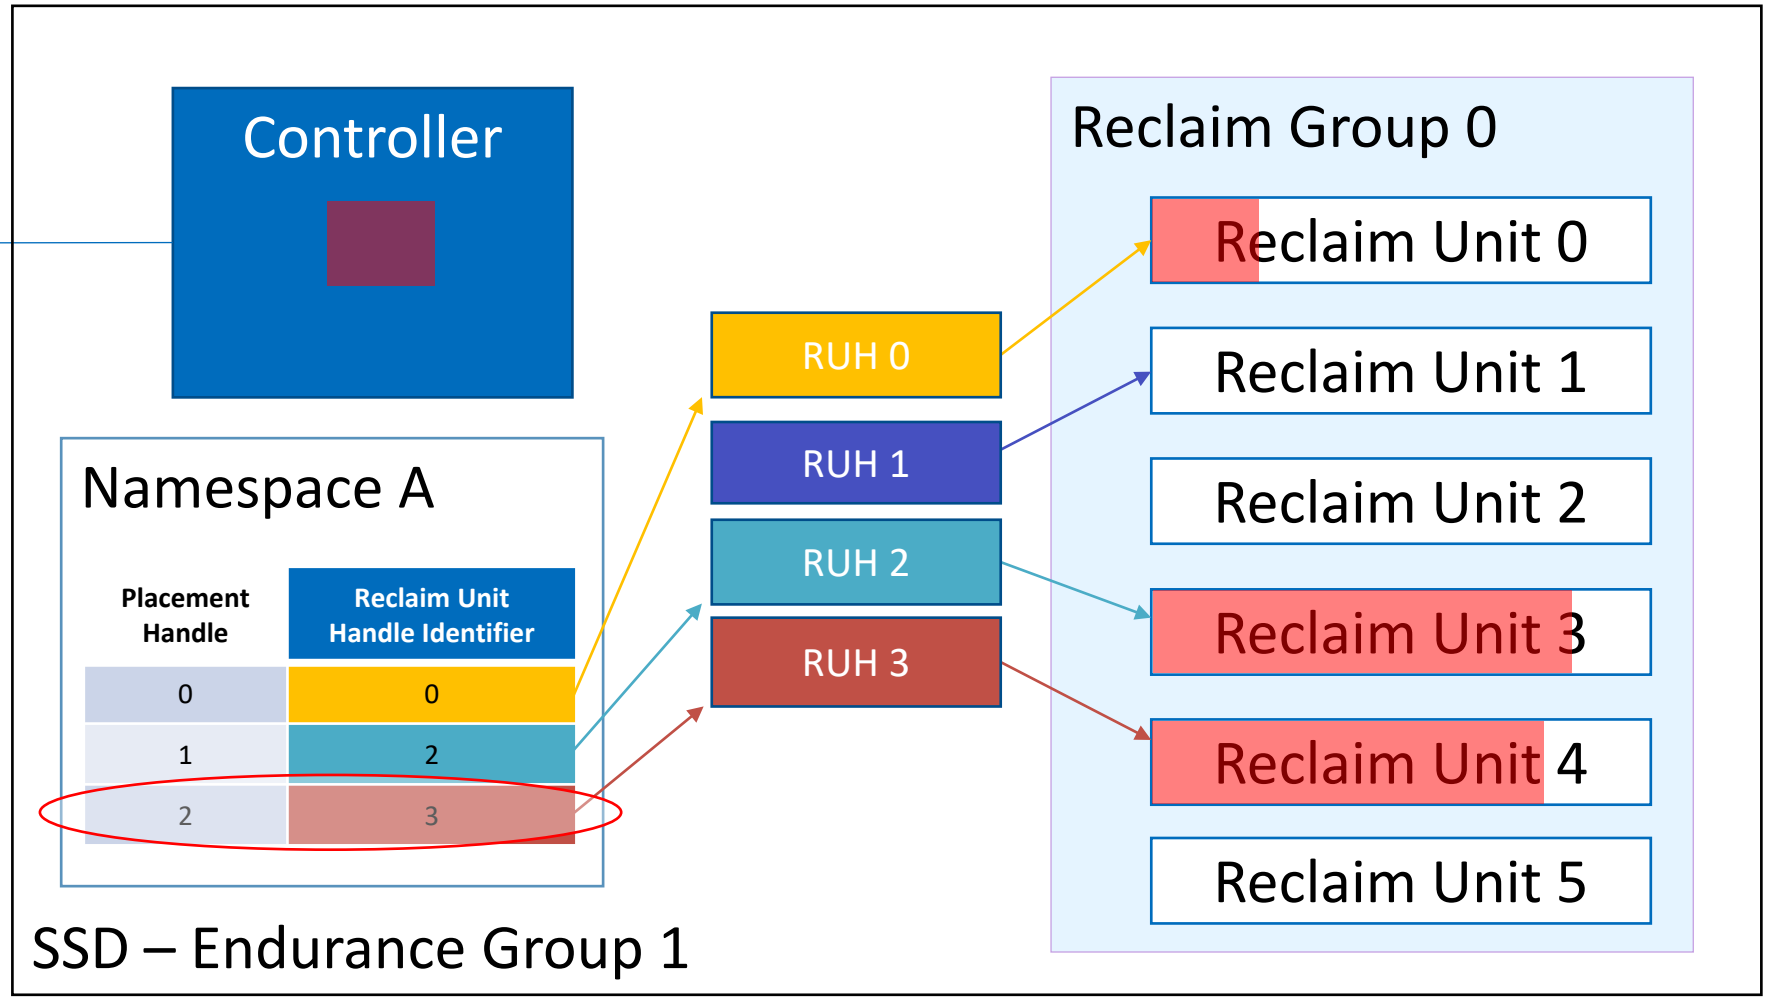

Write to fill the capacity of a Reclaim Unit

Controller uses Reclaim Unit Handle 3 to place the data in Reclaim Unit 4 in Reclaim Group 0

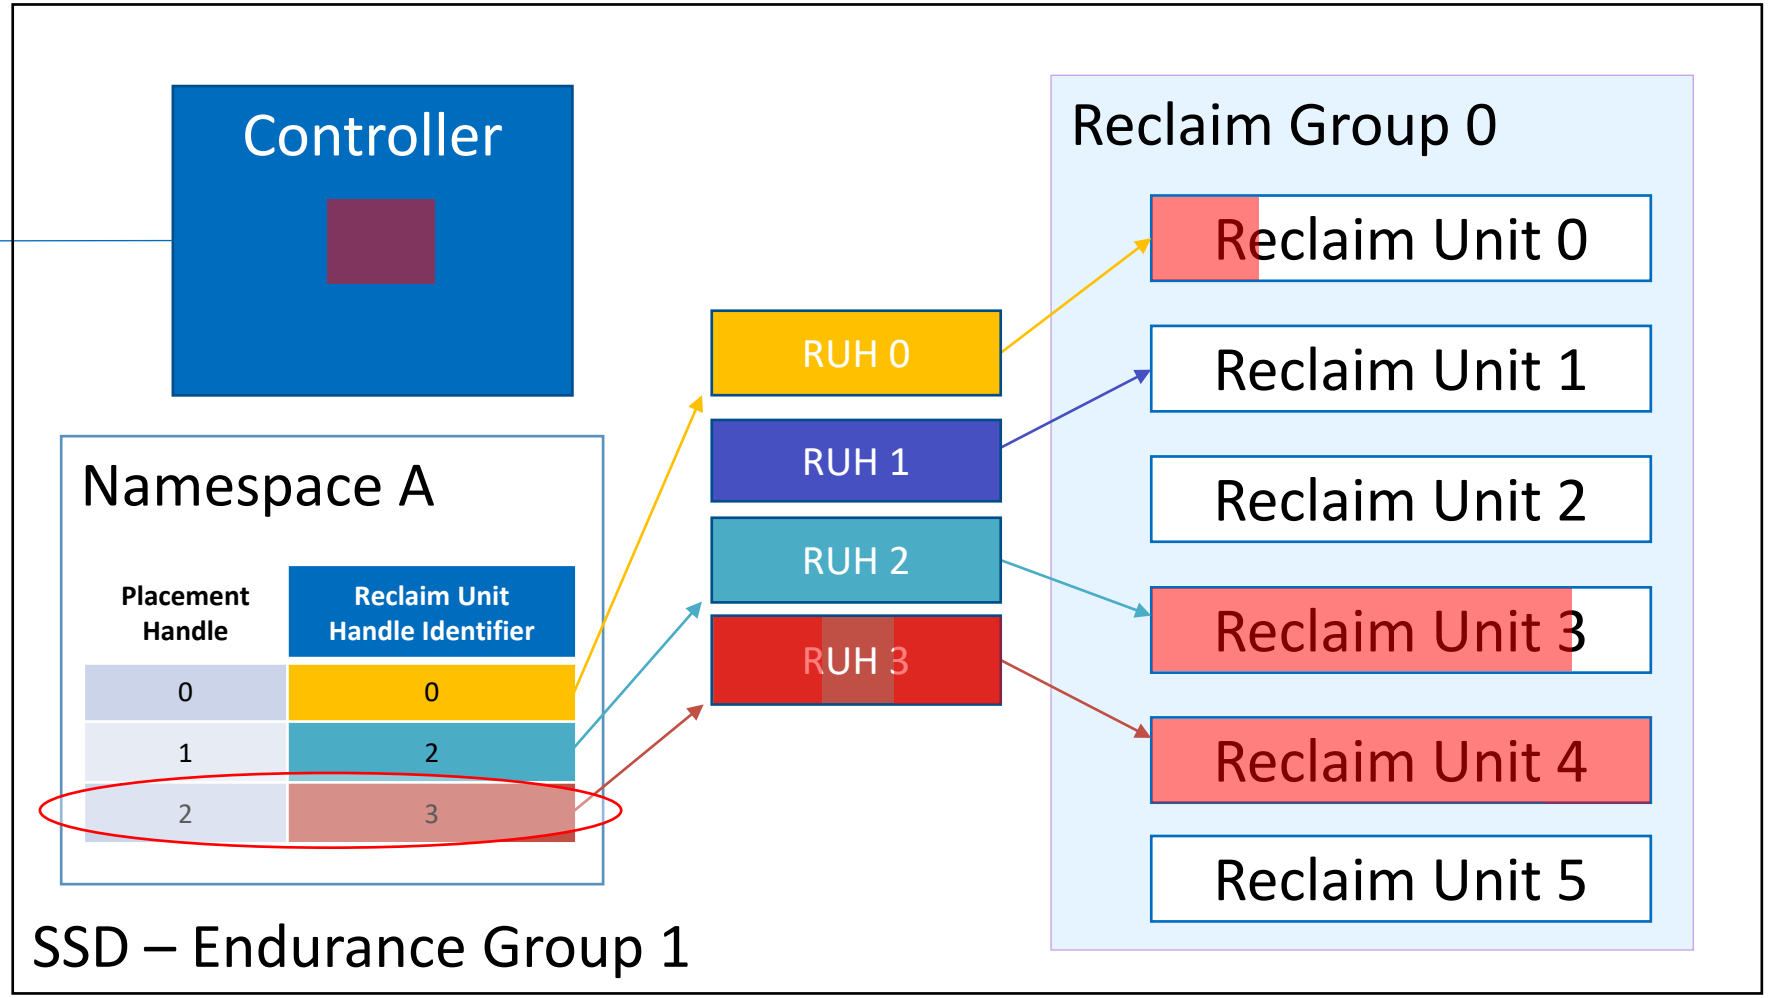

Write to fill the capacity of a Reclaim Unit

Controller detects Reclaim Unit 4 is written to capacity and modifies Reclaim Unit Handle 3 to reference the empty Reclaim Unit 5

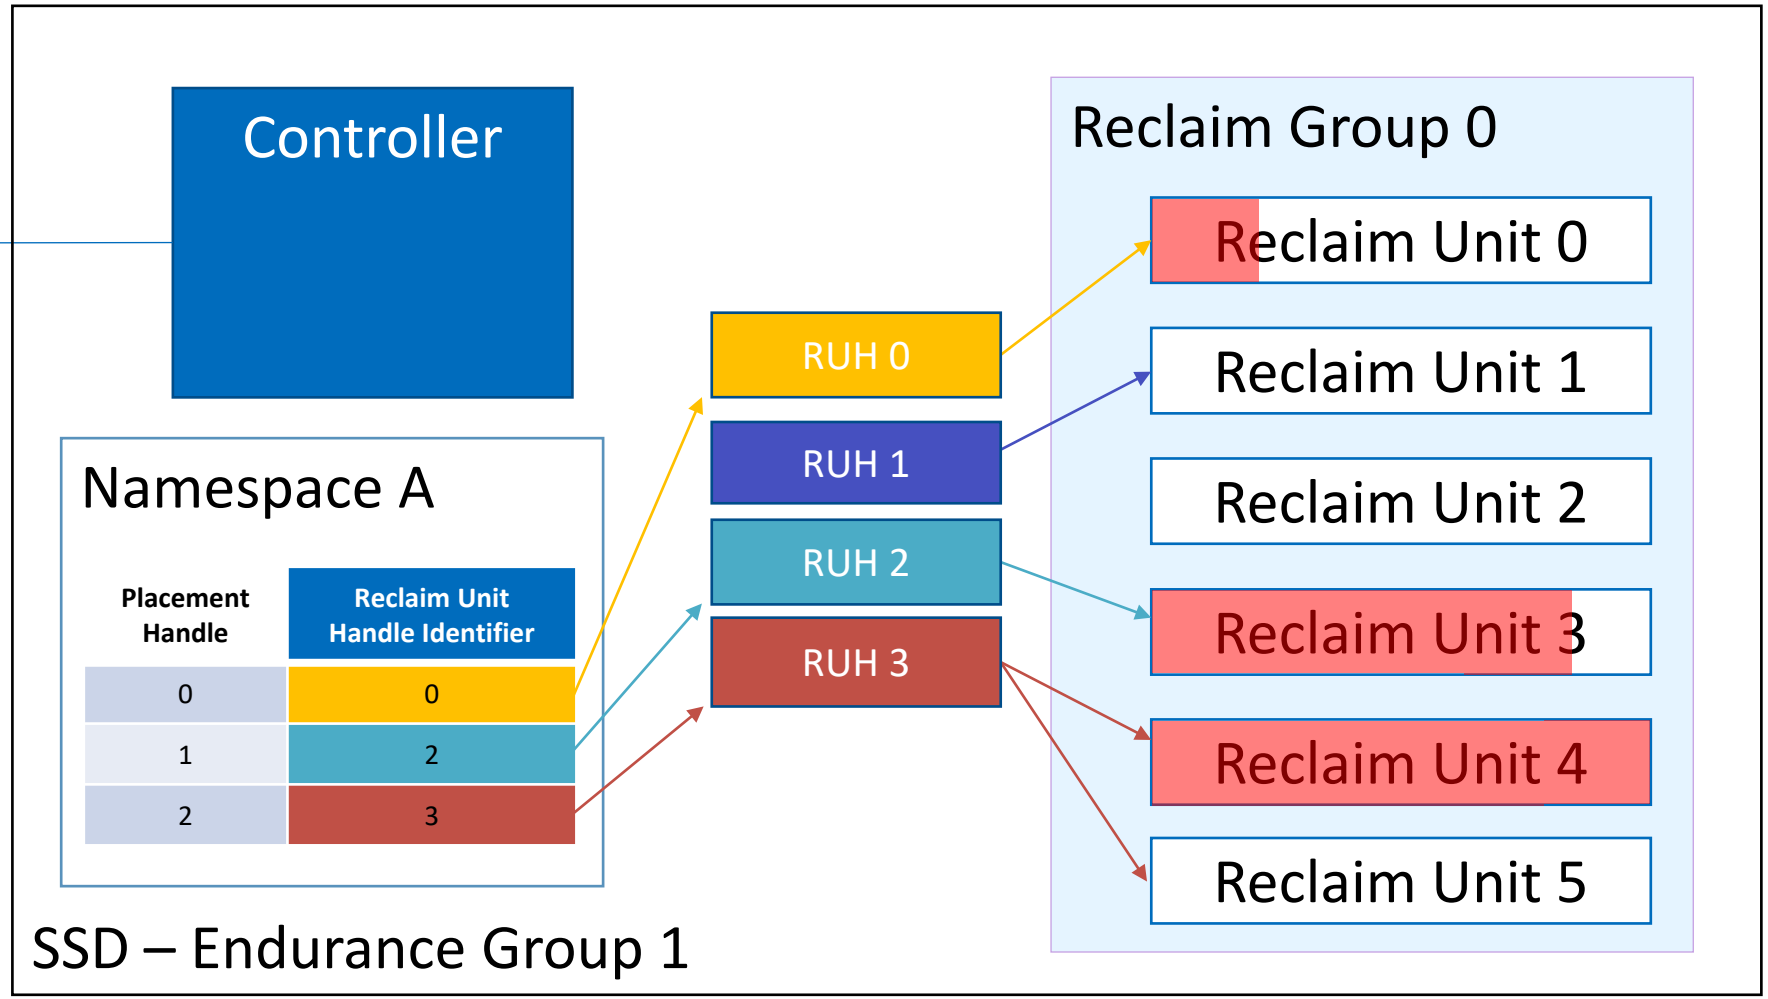

Write to fill the capacity of a Reclaim Unit

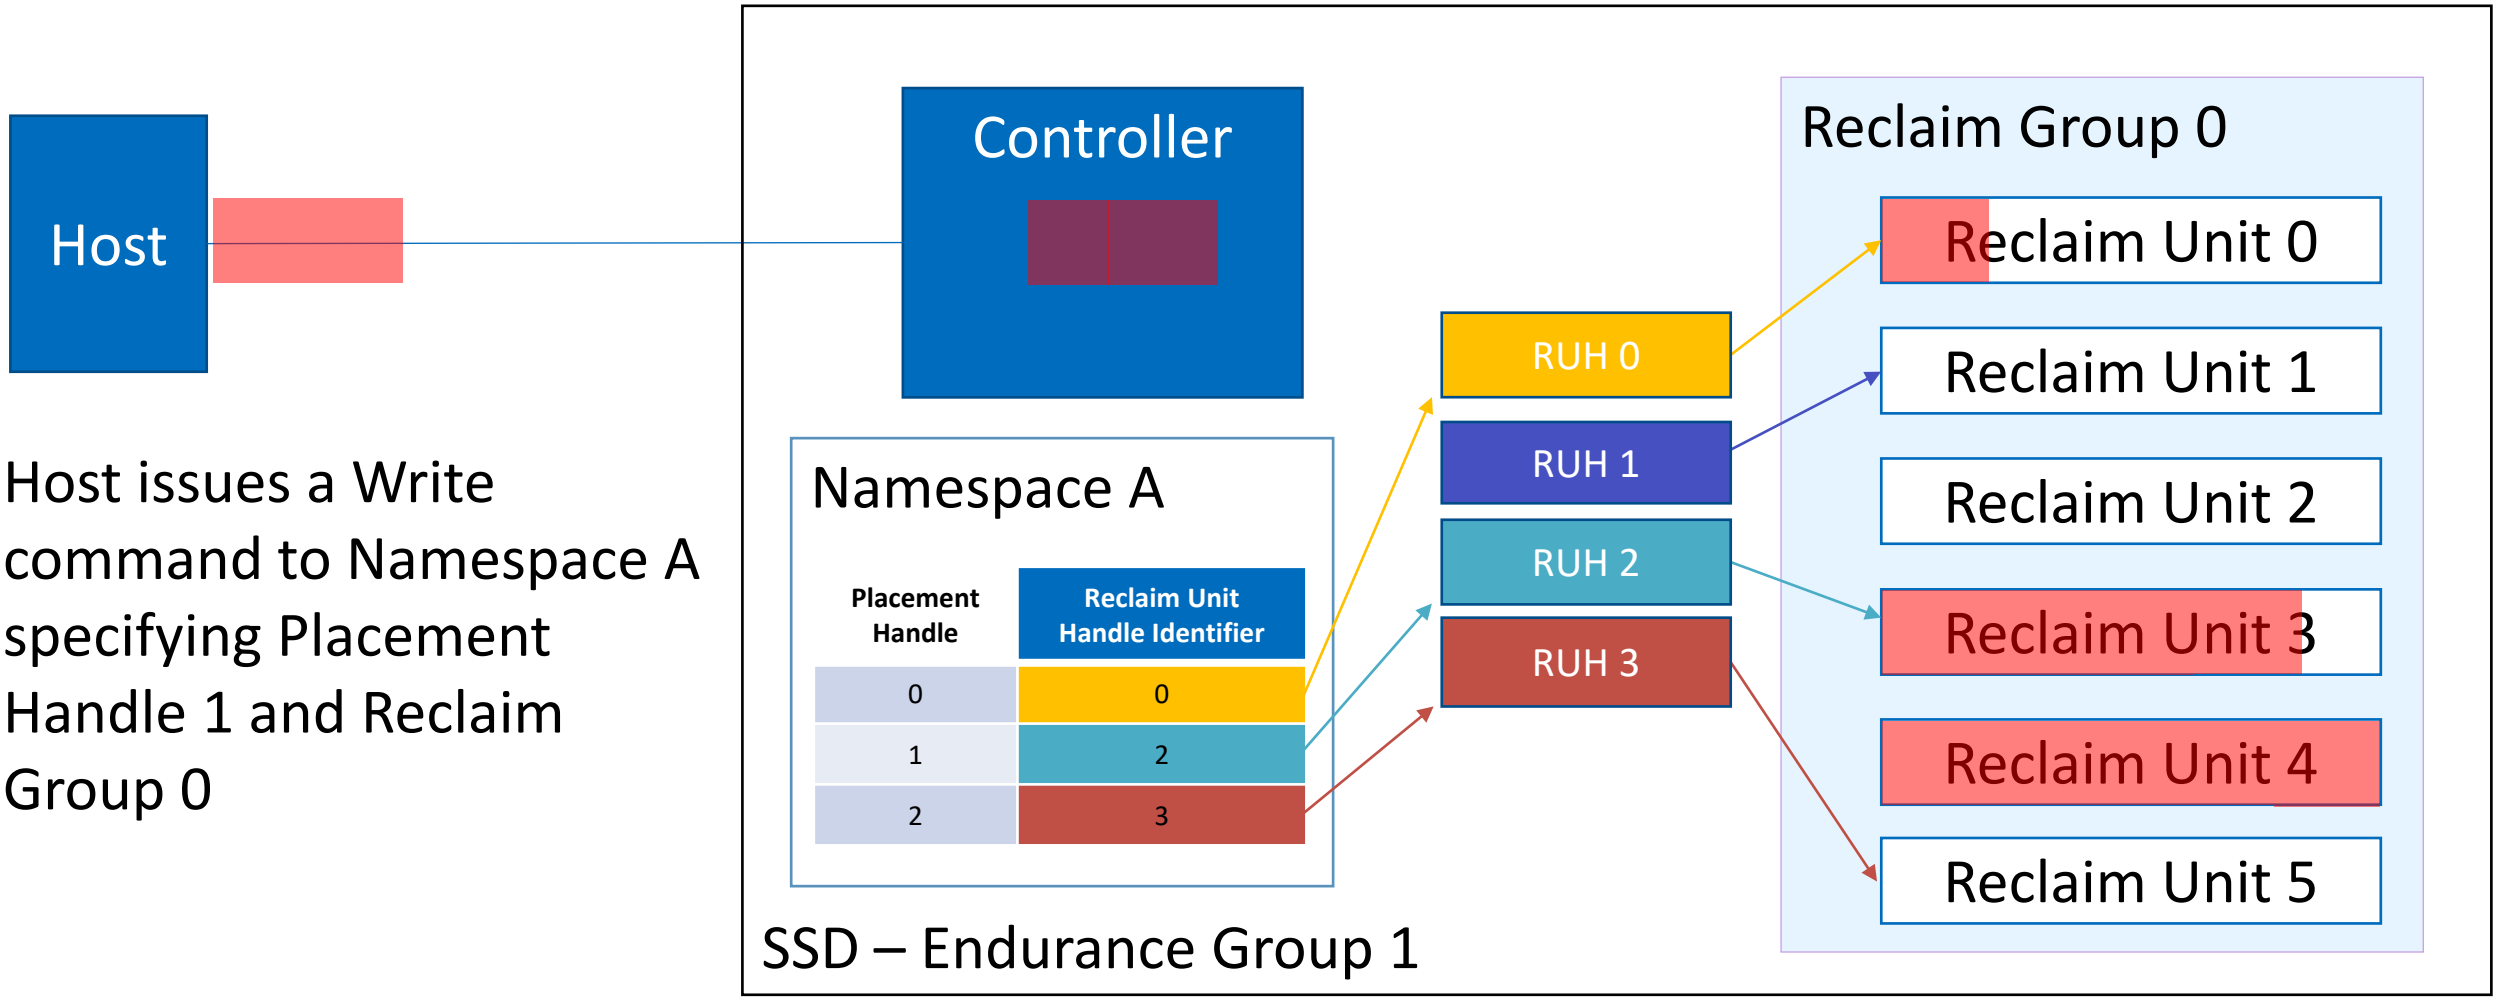

Host issues a Write command to Namespace A specifying Placement Handle 1 and Reclaim Group 0

SAMSUNG SOLIDIGM SDC 23

Write to fill the capacity of a Reclaim Unit

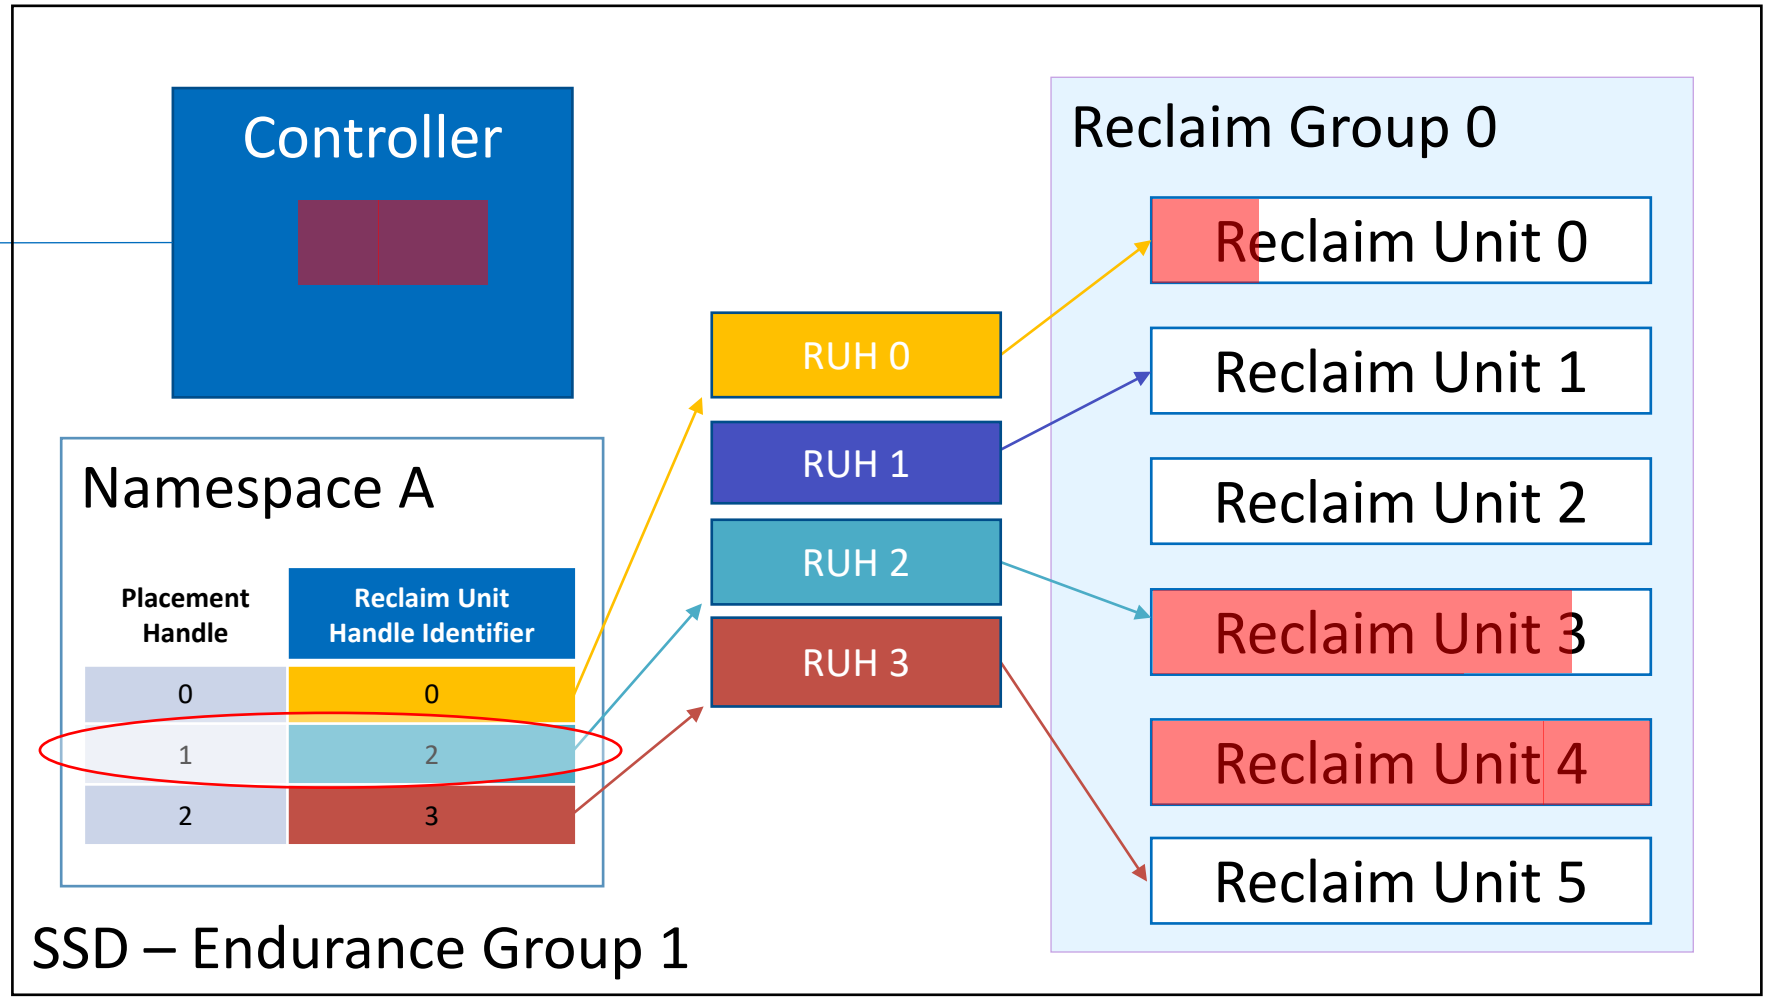

Controller looks up Placement Handle 1 in Namespace A and determines to place the data using Reclaim Unit Handle 2 on Reclaim Group 0

Write to fill the capacity of a Reclaim Unit

Controller uses Reclaim Unit Handle 2 to places a portion of the data in Reclaim Unit 3 in Reclaim Group 0 filling Reclaim Unit 3 to capacity

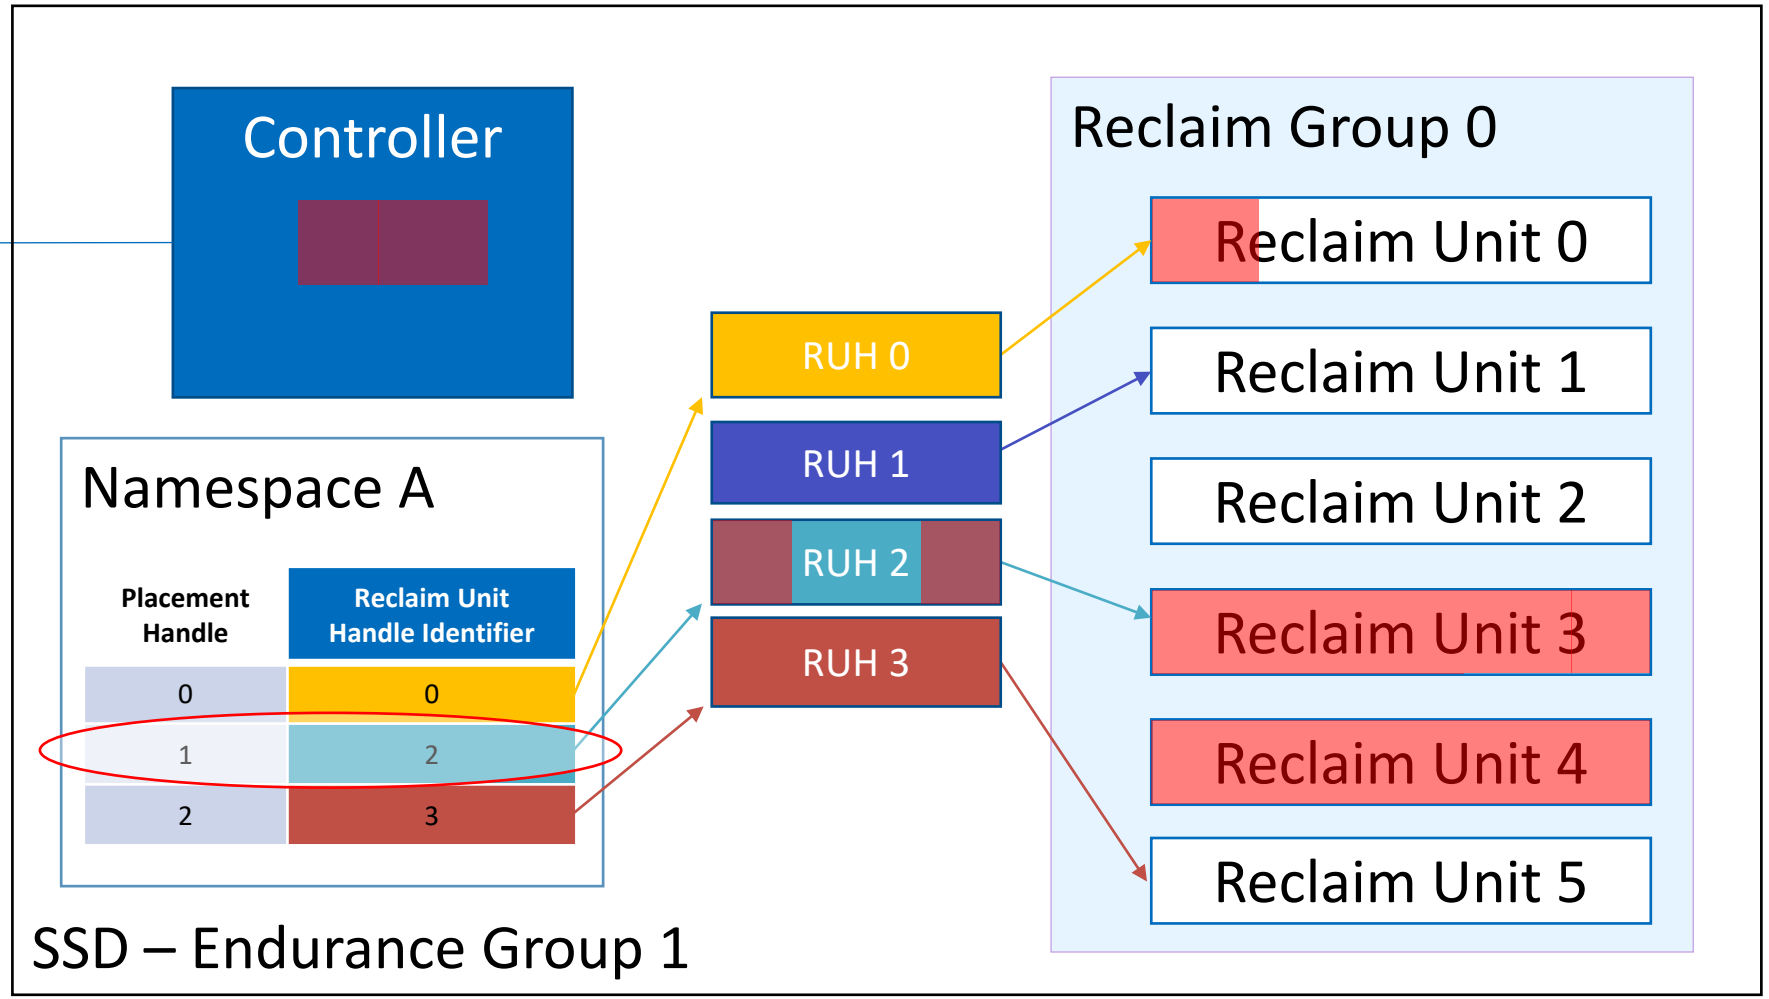

Write to multiple Reclaim Units

Controller Reclaim Unit Handle 2 to reference the empty Reclaim Unit 2

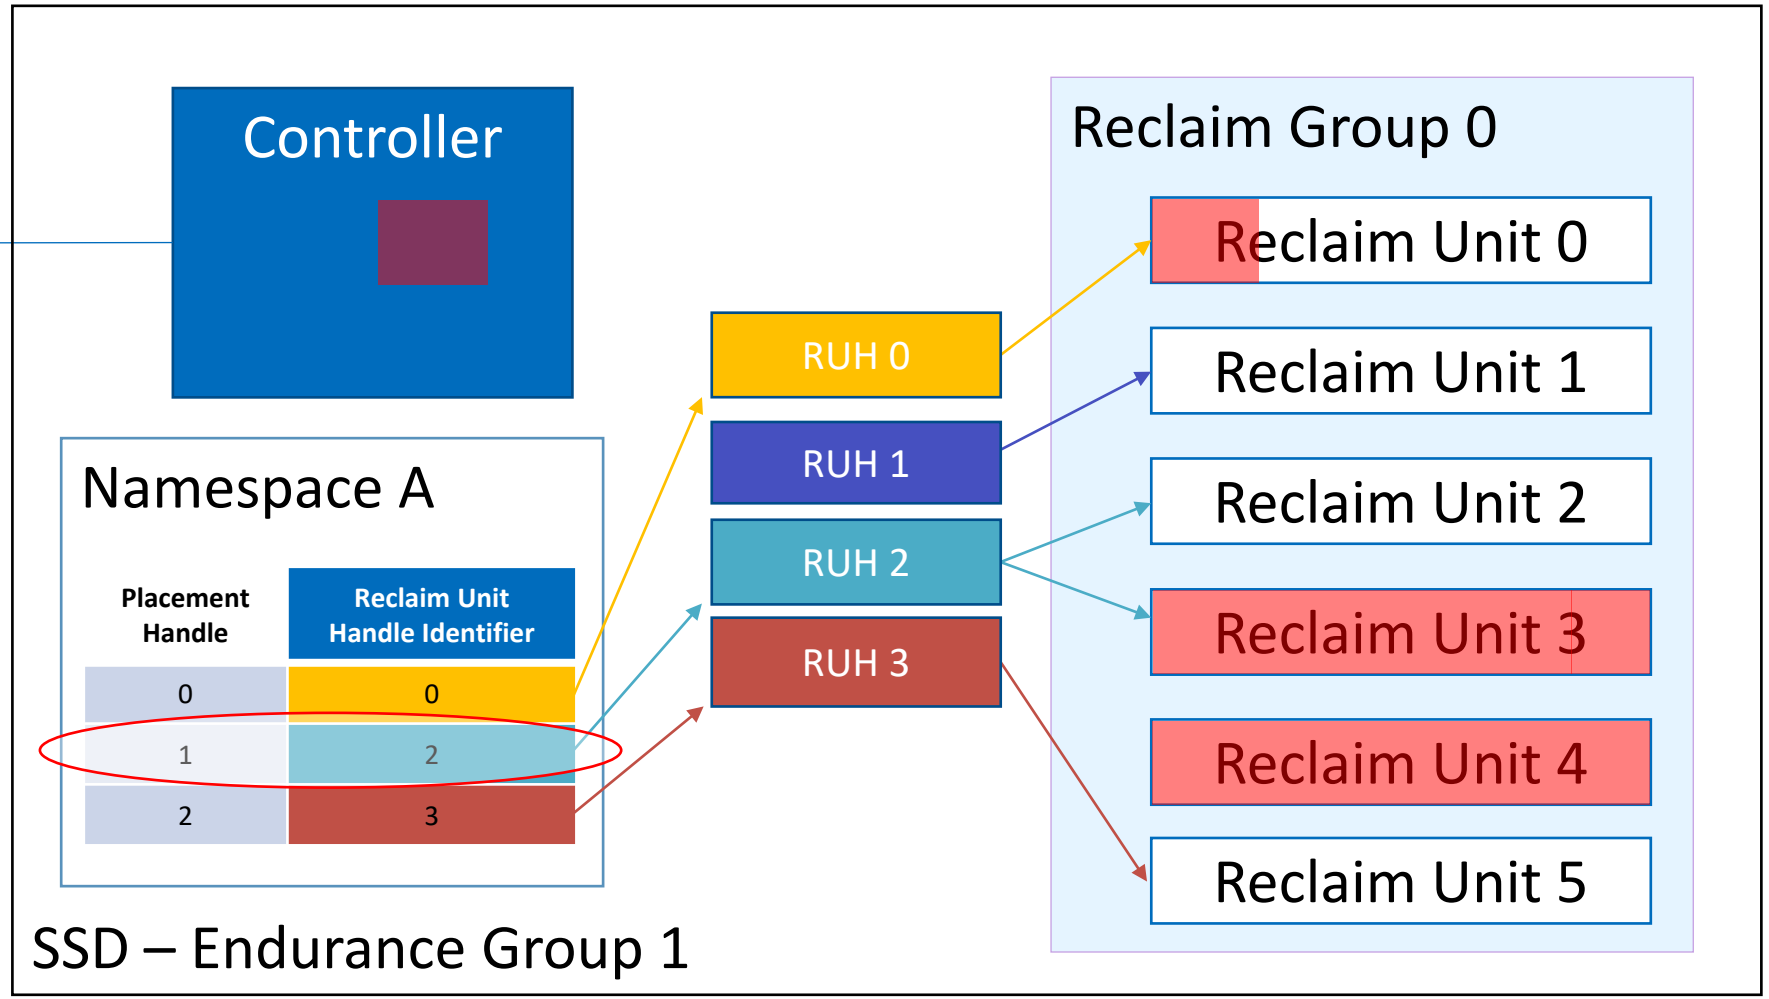

Write to multiple Reclaim Units

Controller uses Reclaim Unit Handle 2 to place the remaining portion of the data in Reclaim Group 0

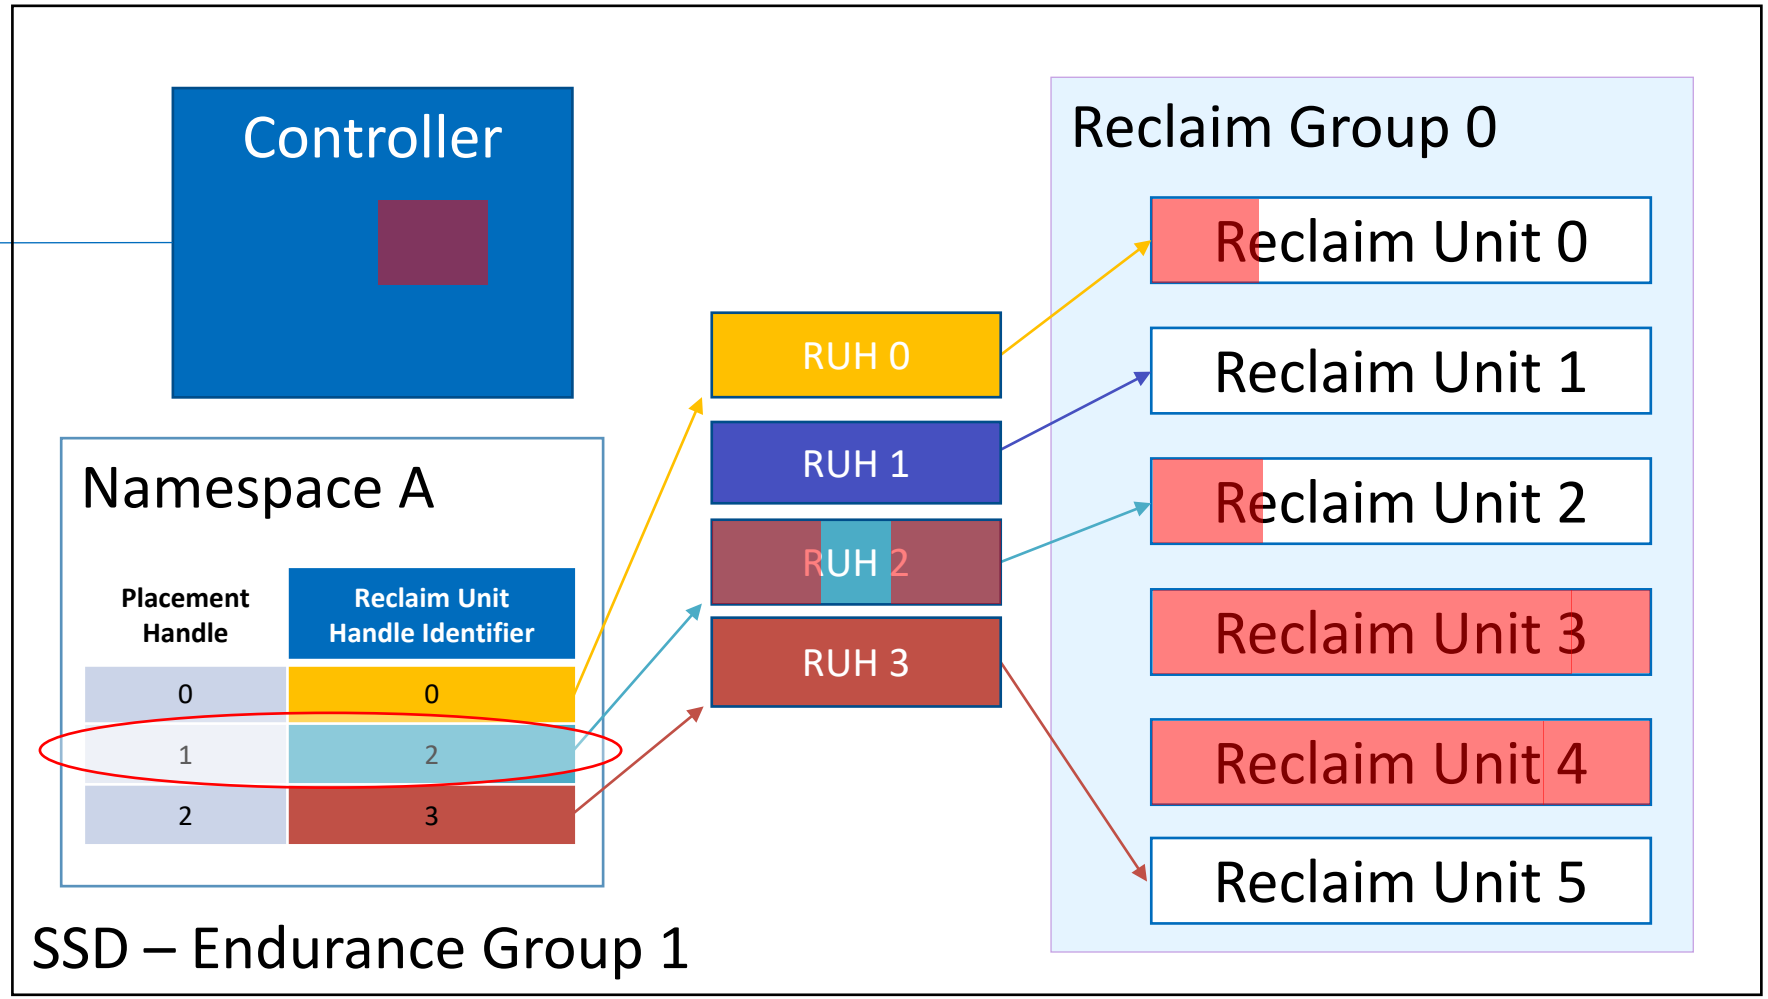

Backwards Compatible Writes

Host issues a Write command to Namespace A without specifying a Placement Handle or an Reclaim Group

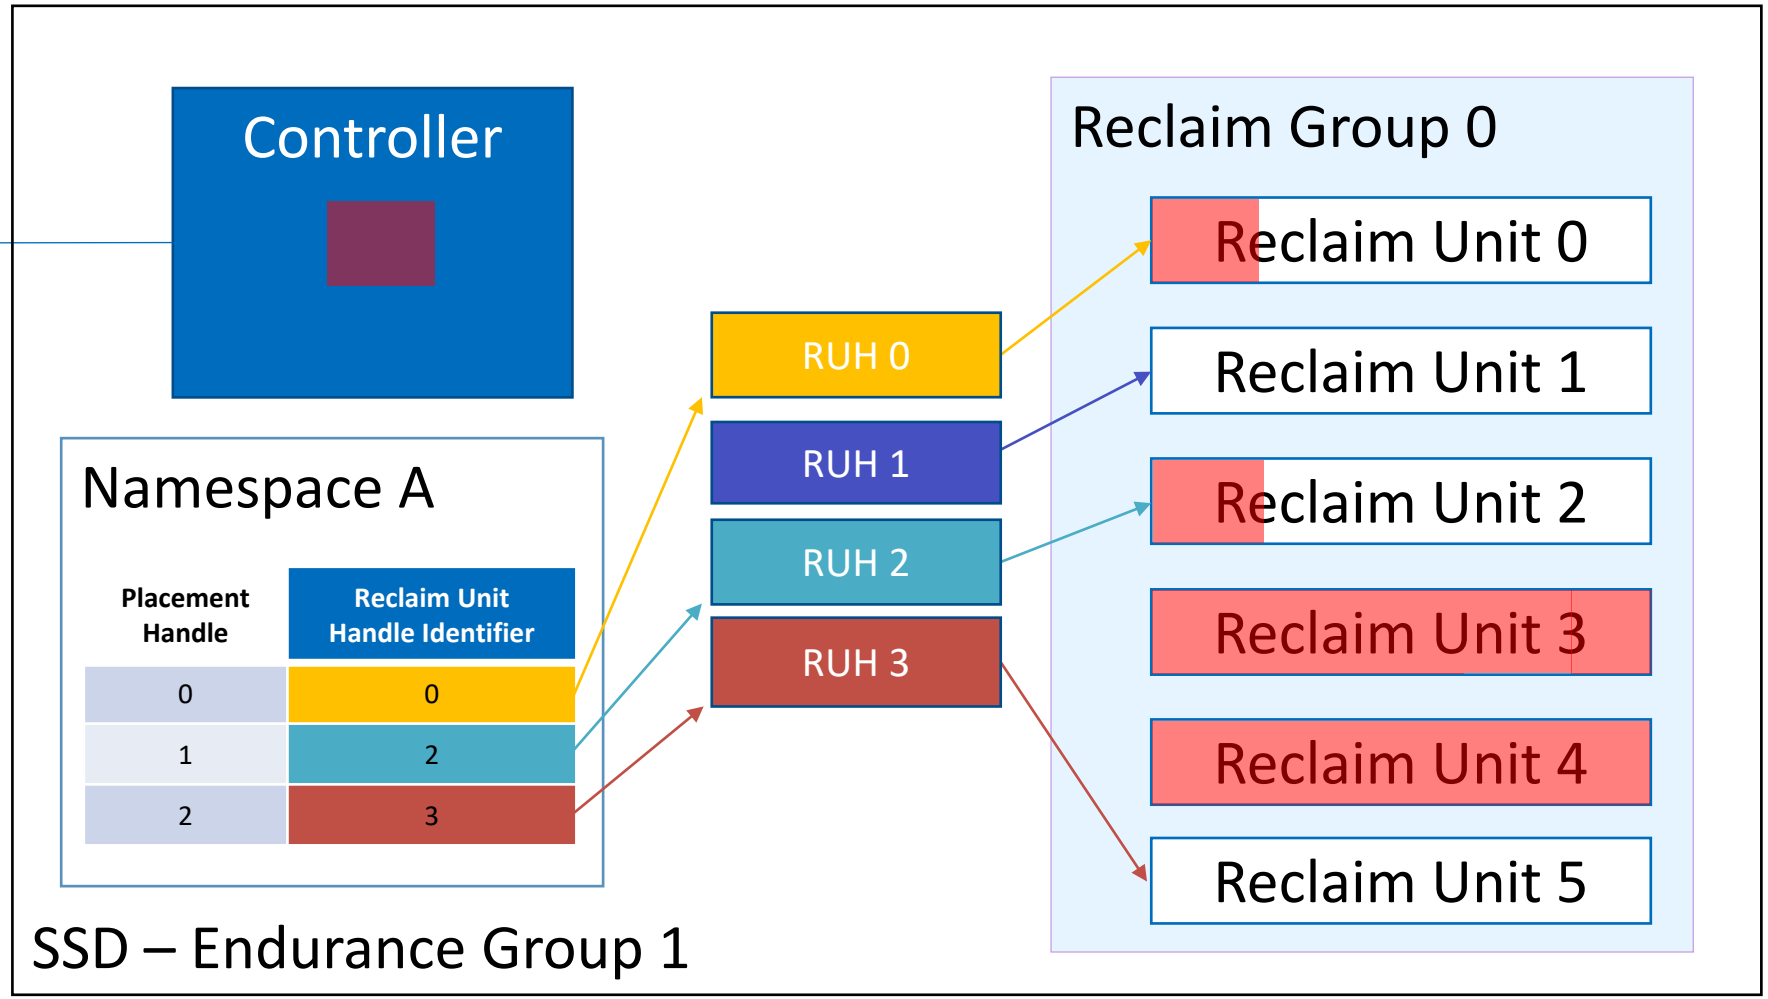

SAMSUNG SOLIDIGM SDC 23

Backwards Compatible Writes

Controller looks up Placement Handle 0 in Namespace A and determines Reclaim Unit Handle 0 is used for the write and selects Reclaim Group 0

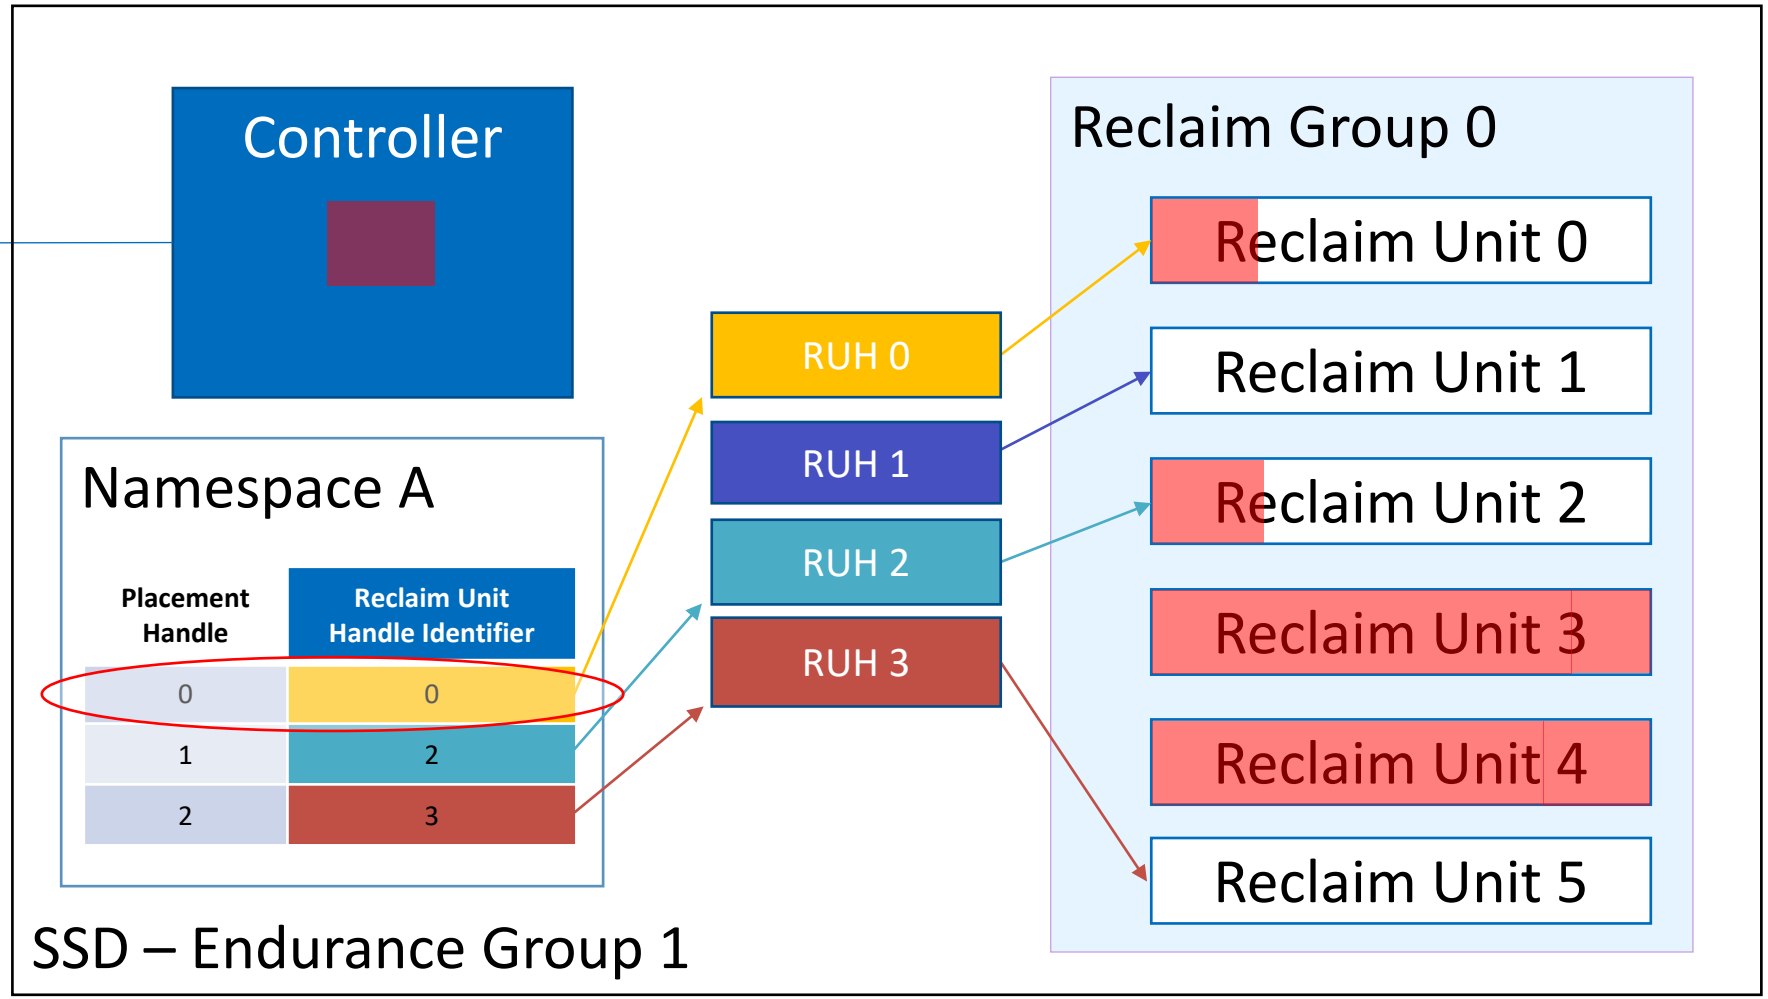

Backwards Compatible Writes

Controller uses Reclaim Unit Handle 0 to place the data in Reclaim Unit 0 in Reclaim Group 0

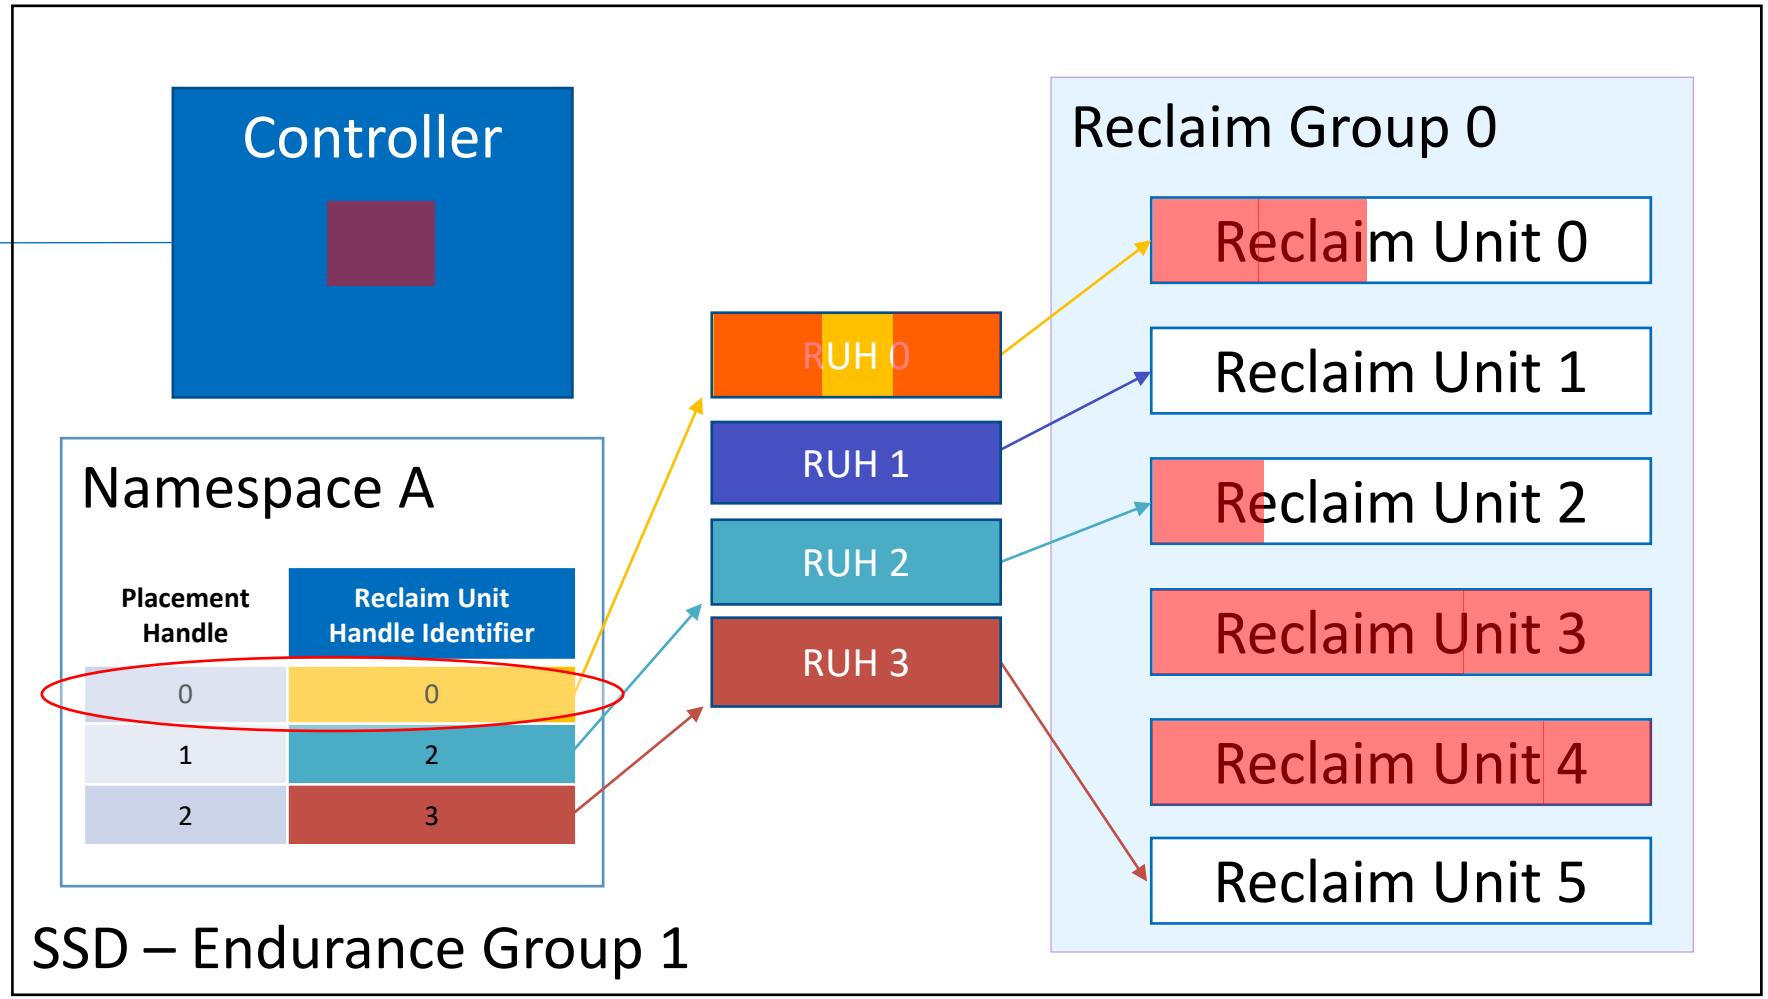

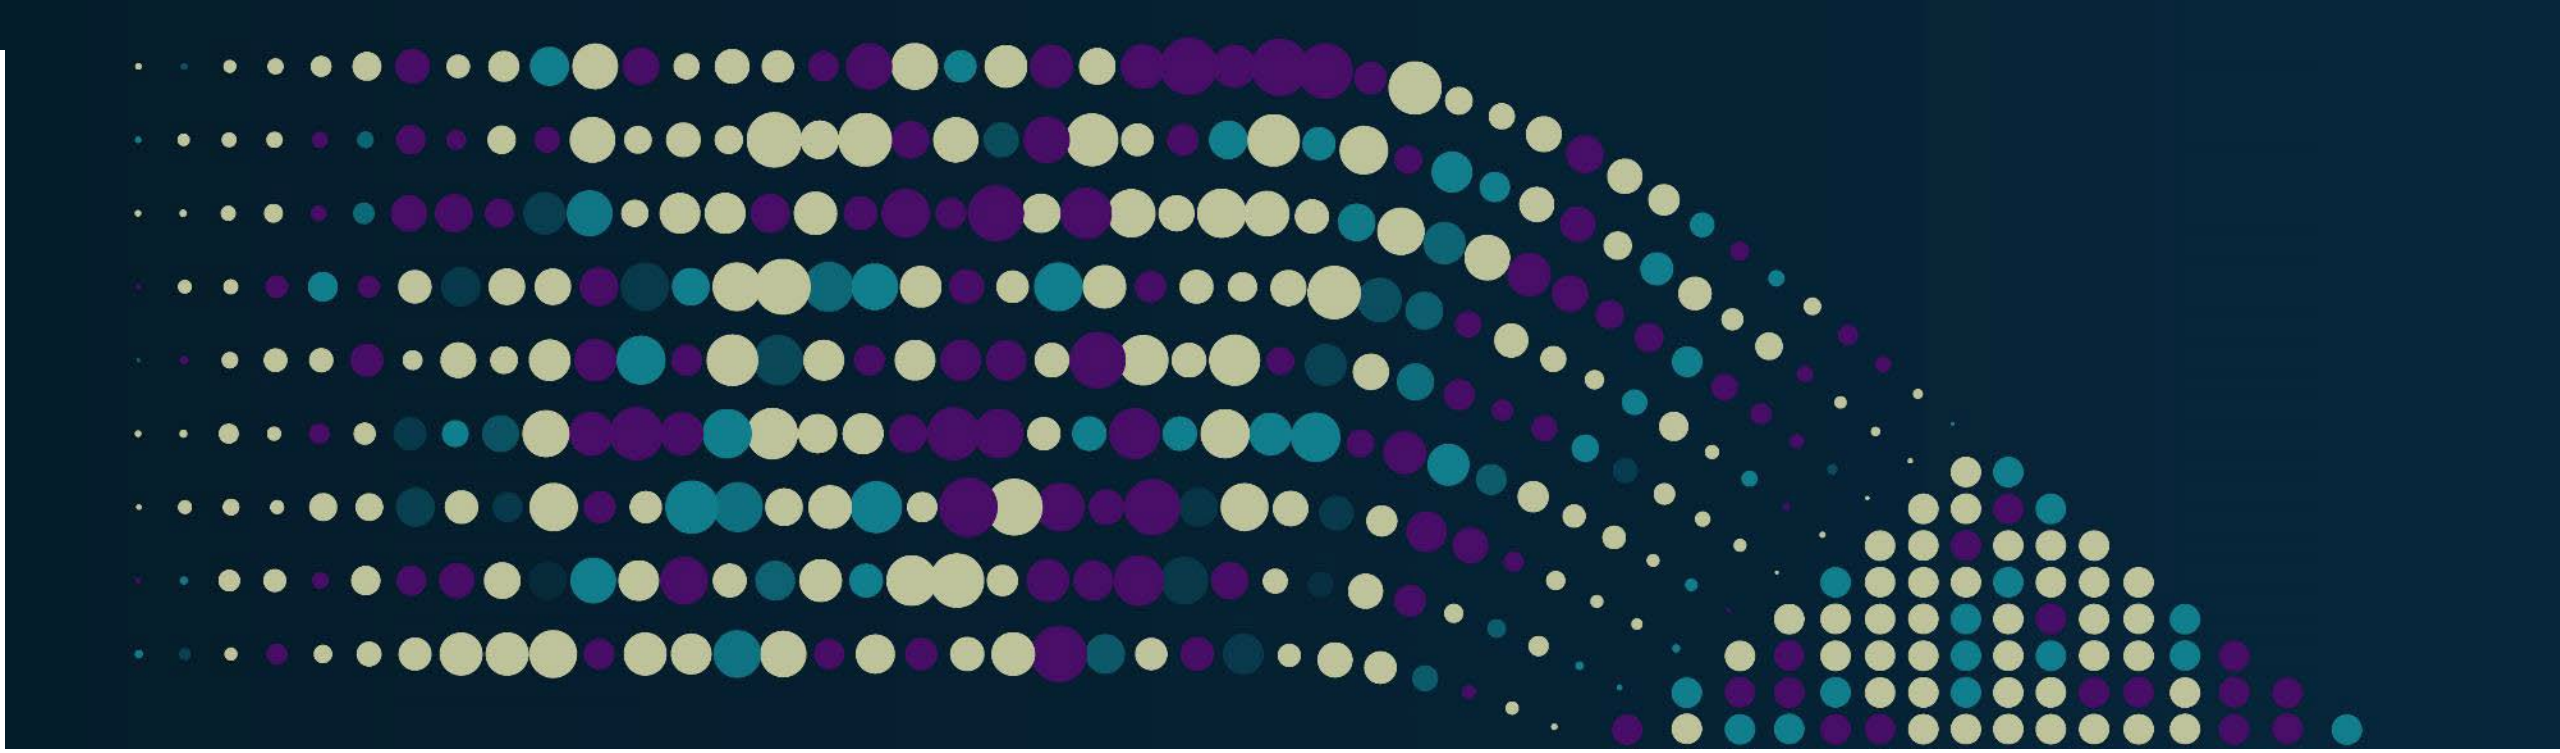

Comparing NVMe® Data Placements

Streams

Flexible Data Placement (FDP)

Zoned Namespaces (ZNS)

SAMSUNG SOLIDIGM SDC 23

Streams	FDP	ZNS
Non-conforming writes are not logged	Non-conforming writes are logged	Error on Write
Known alignment only after format	Commands available to host to stay aligned	Always aligned by interface rules
WAF = 1 achievable without feedback	WAF = 1 achievable with feedback	WAF = 1 guaranteed
Backwards compatible	Backwards compatible	Not Backwards compatible
No information that controller moved user data	Post logging that controller moved user data	Notification for Host to move user data
Placement identifier not tied to LBA	Placement identifier not tied to LBA	Placement identifier is the LBA
Stream Granularity Size (SGS)	Reclaim Units	Zones
SGS capacity = SGS Size	Reclaim Unit capacity = Reclaim Unit size	Zone capacity <= Zone size
No Host metadata per SGS	No Host metadata per Reclaim Unit	Host metadata per Zone
Namespace capacity defines # SGS	Endurance Group capacity defines # Reclaim Units	Namespace Capacity defines # zones
Sequential, Random, and Over Write	Sequential, Random, and Over Write	Sequential Write
Writes allowed to cross Boundaries	Writes allowed to cross Boundaries	Writes not allowed across Boundaries
QD > 1: LBA known at Write Submission	QD > 1: LBA known at Write Submission	QD > 1: LBA known at Write Completion Zone Append command
Stream written by a single namespace	Reclaim Unit written by one or more namespaces	Zone written by a single namespace
API is Stateless	API is Stateless	API is Stateful
Requires Full FTL Table	Requires Full FTL Table	Full FTL Table not required
Dynamic write resource allocation	Static write resource allocation	Dynamic write resource allocation

SAMSUNG SOLIDIGM SDC 23

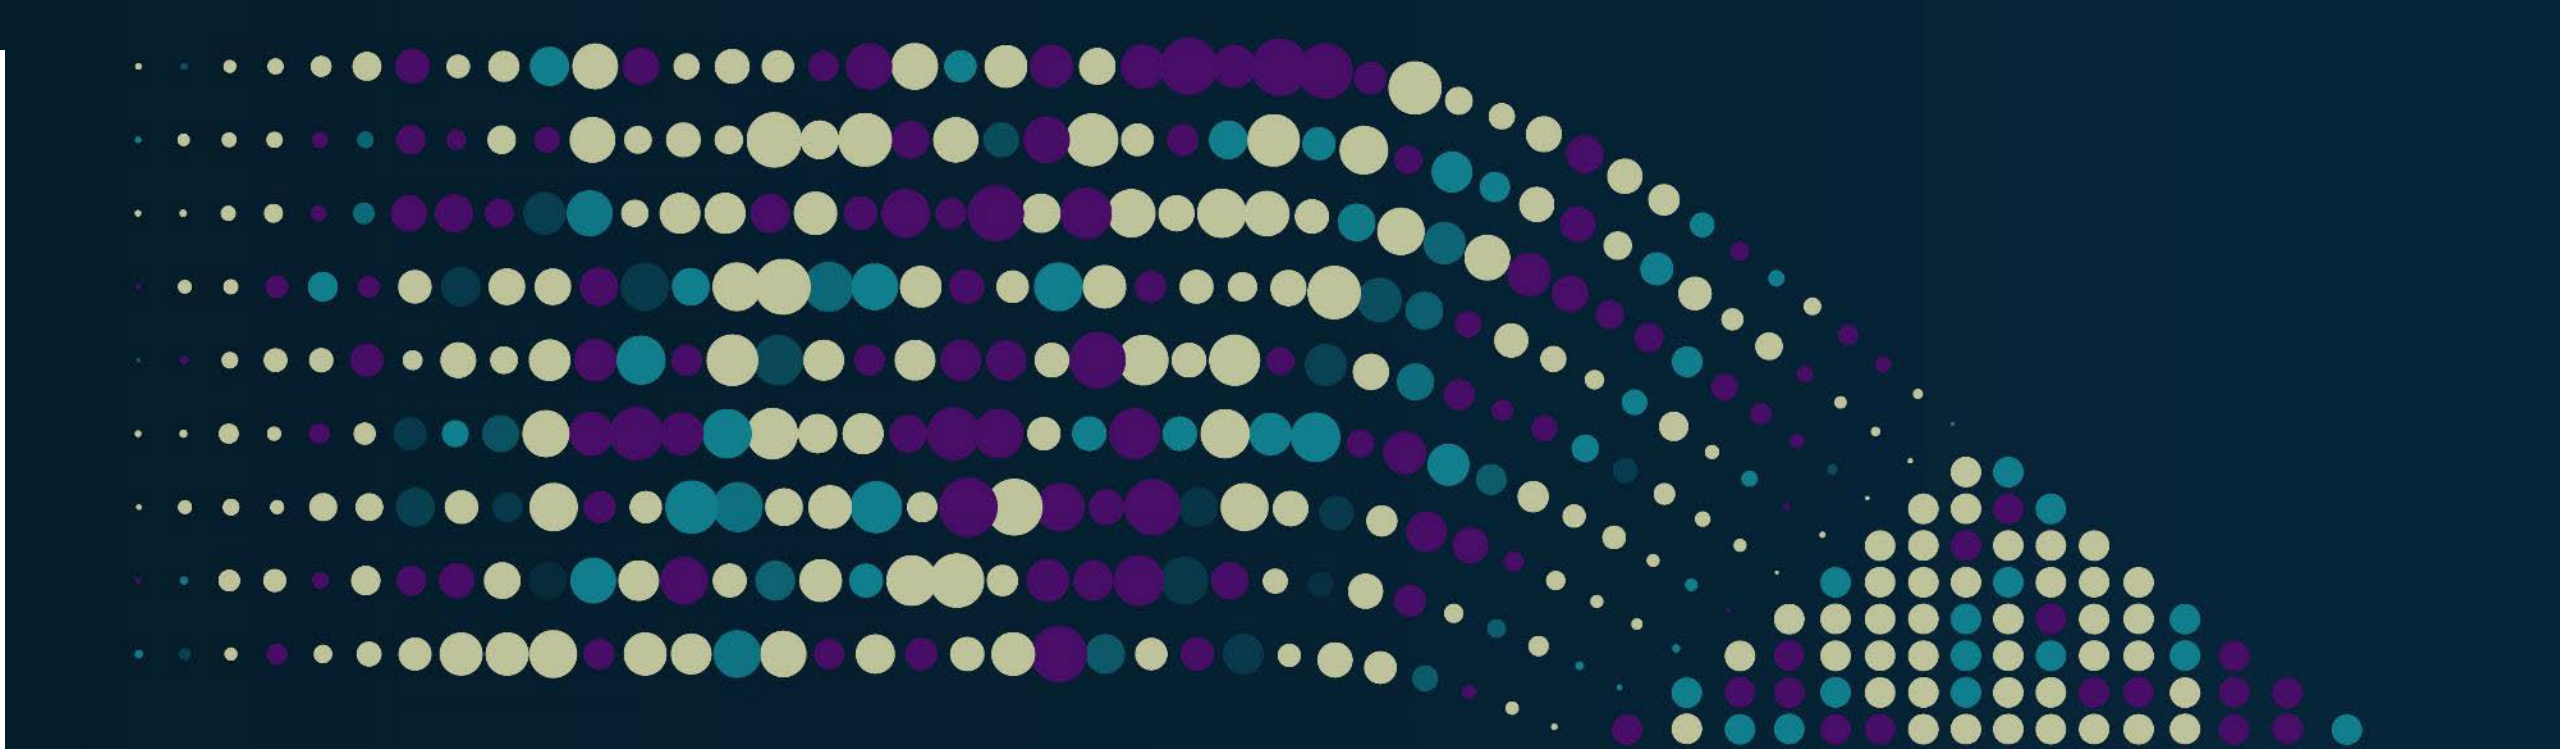

Please take a moment to rate this session.

Your feedback is important to us.

SAMSUNG SOLIDIGM SDC 23