



Host Managed Live Migration

Sponsored by NVM Express organization, the owner of NVMe® specifications

Speakers



Mike Allison
Sr. Director NAND Product Planning
Standards

SAMSUNG

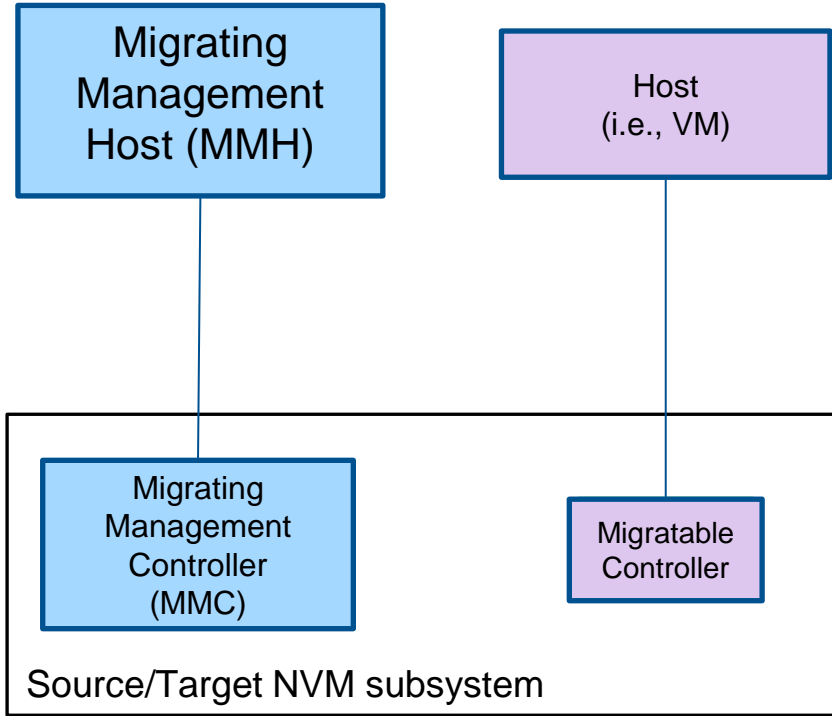
Agenda

Review the NVM Express ratified TP4159 PCIe[®] Infrastructure for Live Migration

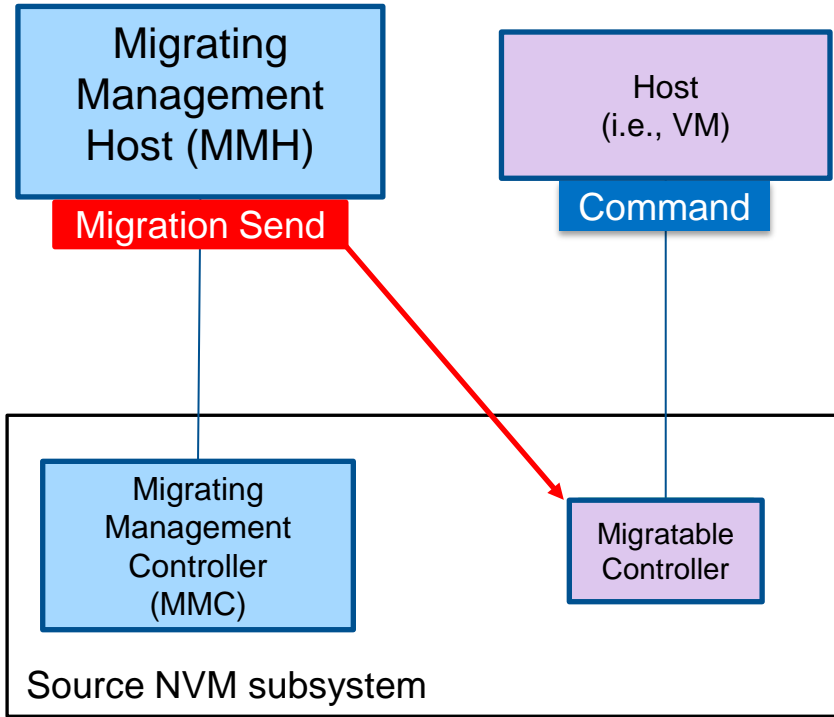
- Migrating a controller
 - Suspend/Resume a controller
 - Get/Set Controller State
- Migrating namespaces attached to a migrating controller
 - Logging Logical Block Address (LBA) modifications to a LBA Migration Queue
- Track host memory modifications due to migrating controller processing commands



Migrating a Controller



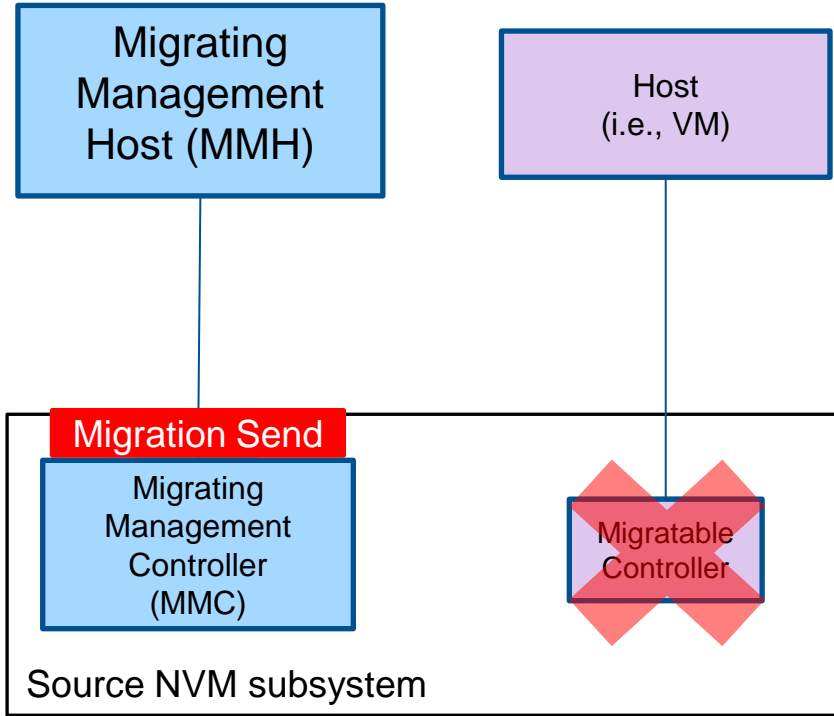
Migrating a Controller – Suspending a Controller



Migrating Management Host (MMH) suspends the controller to be migrated

- Submit a **Migration Send** command to a Migrating Management Controller (MMC) specifying:
 - Controller Identifier of the Migratable Controller
 - Suspend Management Operation
- Migratable Controller:
 - Halts fetching commands
 - Completes the processing of all previously fetched commands

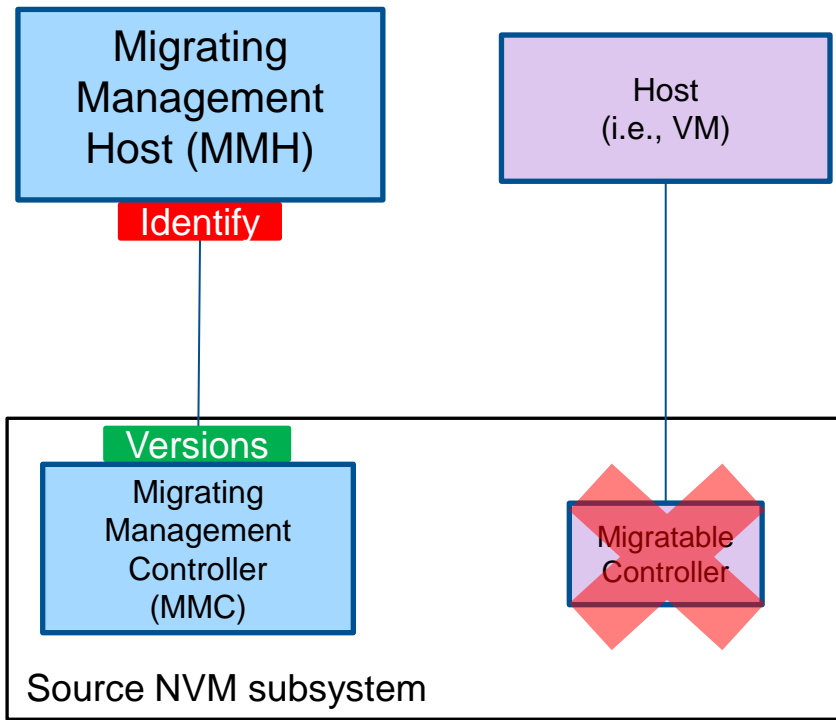
Migrating a Controller – Suspending a Controller



Migrating Management Host (MMH) suspends the controller to be migrated

- Submit a **Migration Send** command to a Migrating Management Controller (MMC) specifying:
 - Controller Identifier of the Migratable Controller
 - Suspend Management Operation
- Migratable Controller:
 - Halts fetching commands
 - Completes the processing of all previously fetched commands

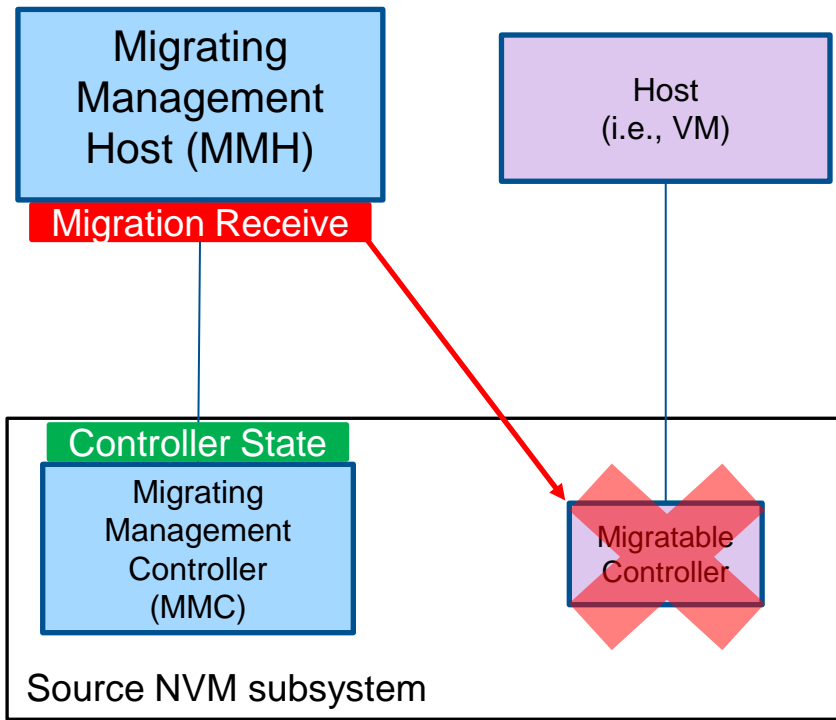
Migrating a Controller – Get the Controller State



MMH obtains the controller state of the suspended controller being migrated

- Submit an Identify command specifying a **Controller Namespace Structure (CNS) value 20h**
- MMC returns supported versions:
 - NVMe® technology defined state (#)
 - Vendor specific state (UUID)

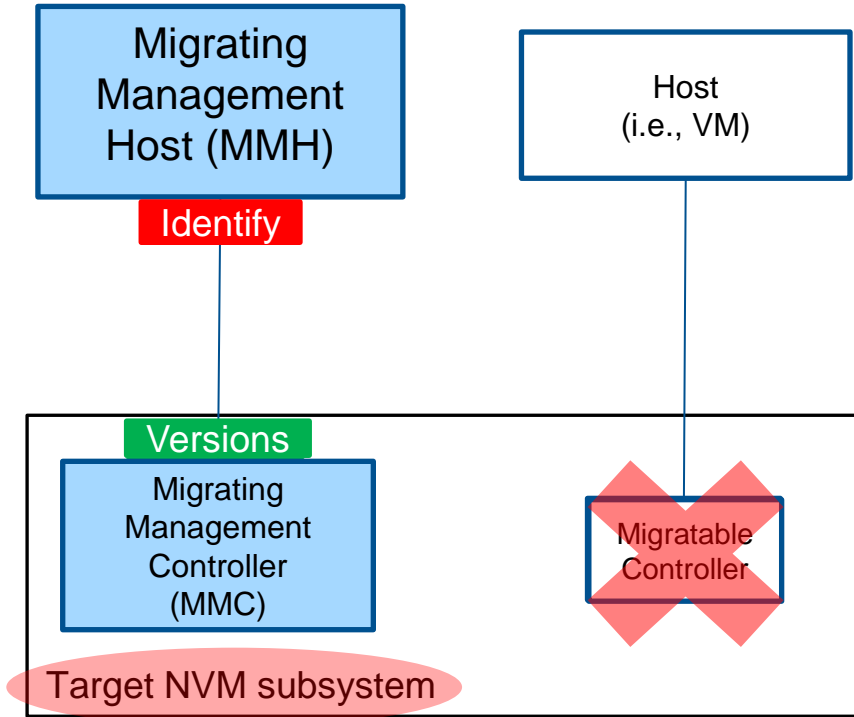
Migrating a Controller – Get the Controller State



MMH obtains the controller state of the suspended controller being migrated

- Submit an Identify command specifying a **Controller Namespace Structure (CNS) value 20h**
- MMC returns supported versions:
 - NVMe® technology defined state (#)
 - Vendor specific state (UUID)
- Submits one or more **Migration Receive Send** commands to a MMC specifying:
 - Controller Identifier of the Migratable Controller
 - Specific controller state versions
 - Offset
 - Size
- MMC returns the controller state

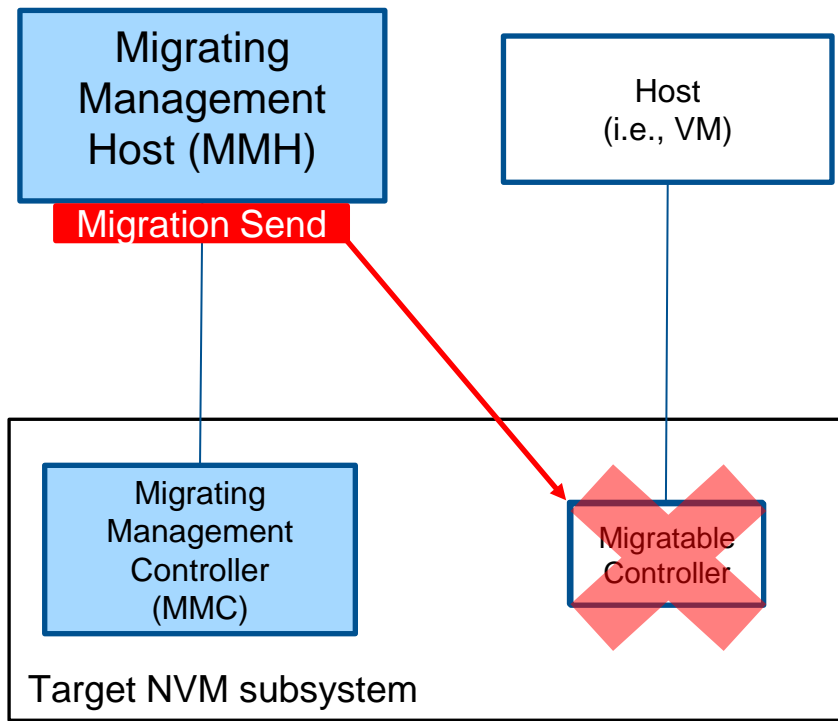
Migrating a Controller – Set the Controller State



MMH puts the controller state of the suspended controller being migrated into a different controller

- Submit an Identify command specifying a **CNS value 20h**
- MMC returns controller state versions:
 - NVMe[®] technology defined state
 - Vendor specific state

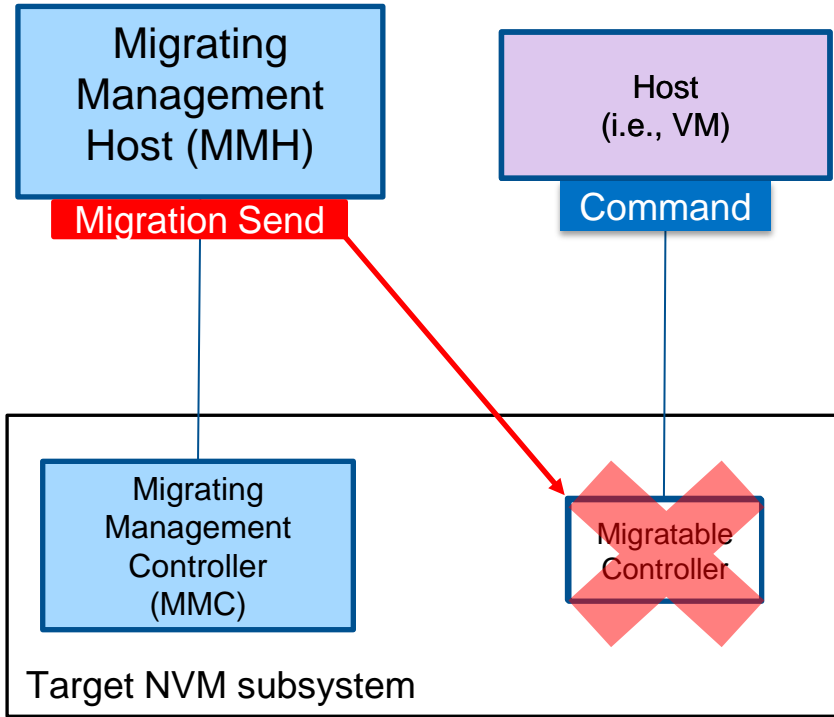
Migrating a Controller – Set the Controller State



MMH puts the controller state of the suspended controller being migrated into a different controller

- Submit an Identify command specifying a **CNS value 20h**
- MMC returns controller state versions:
 - NVMe® technology defined state
 - Vendor specific state
- Submits one or more **Migration Send** commands to a MMC specifying:
 - Controller Identifier of the different controller
 - Get Controller State Management Operation
 - Specific controller state versions
 - Order of commands
 - First, middle, last

Migrating a Controller – Resume the Controller

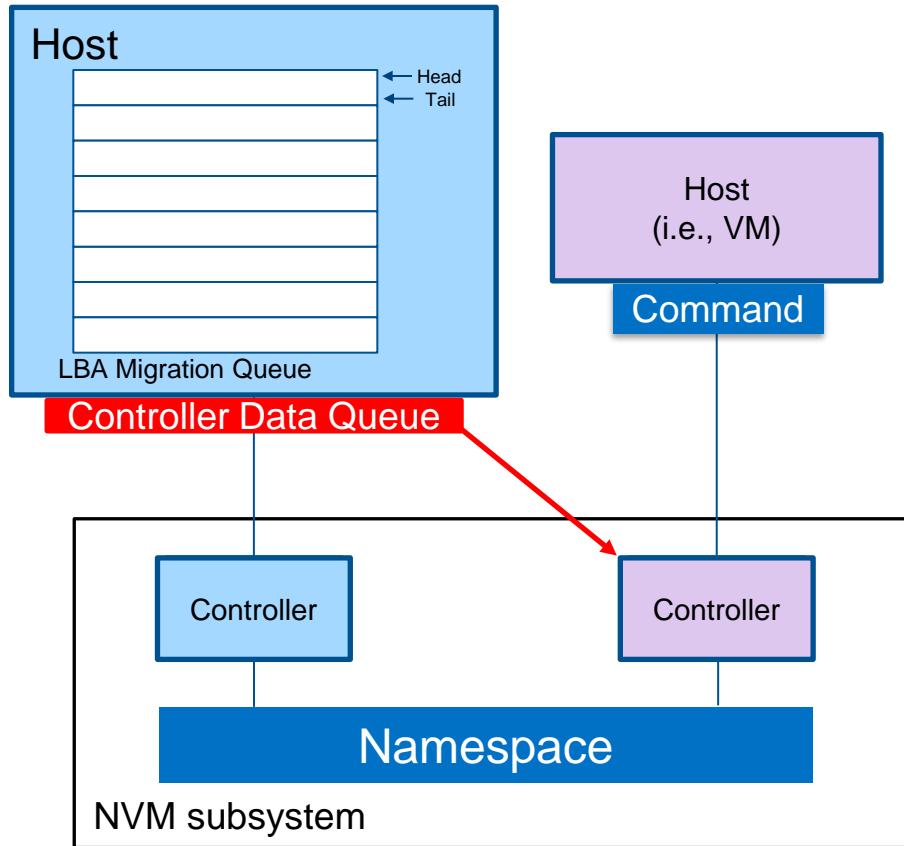


MMH responsibility to migrate the VM

MMH resumes the controller that has migrated

- Submit a **Migration Send** command to a Migrating Management Controller (MMC) specifying:
 - Controller Identifier of the Migratable Controller
 - Resume Management Operation
- Migratable Controller:
 - Starts fetching commands

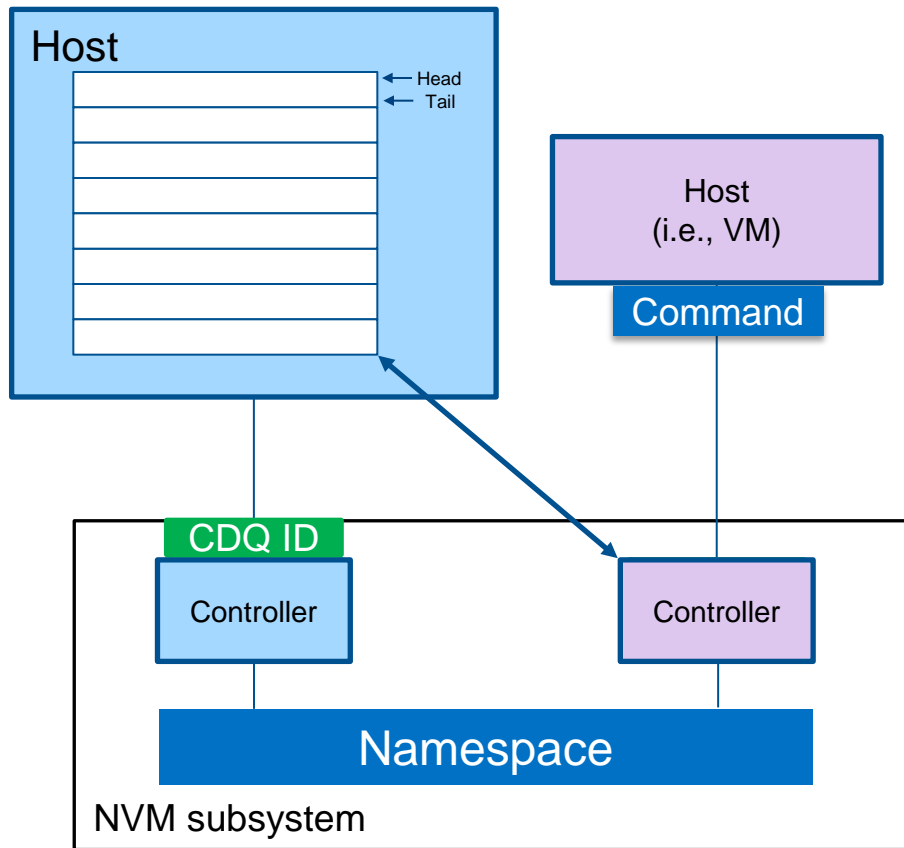
Live Migration of Attached Namespaces



Create a LBA Migration Queue

- Host allocates memory for a LBA Migration Queue
- Host submits a **Controller Data Queue** command to a controller specifying:
 - Controller Identifier of a controller to track LBA modifications by that controller
 - Log User Data Changes Management Operation
 - Host Memory Address
 - Dword Size

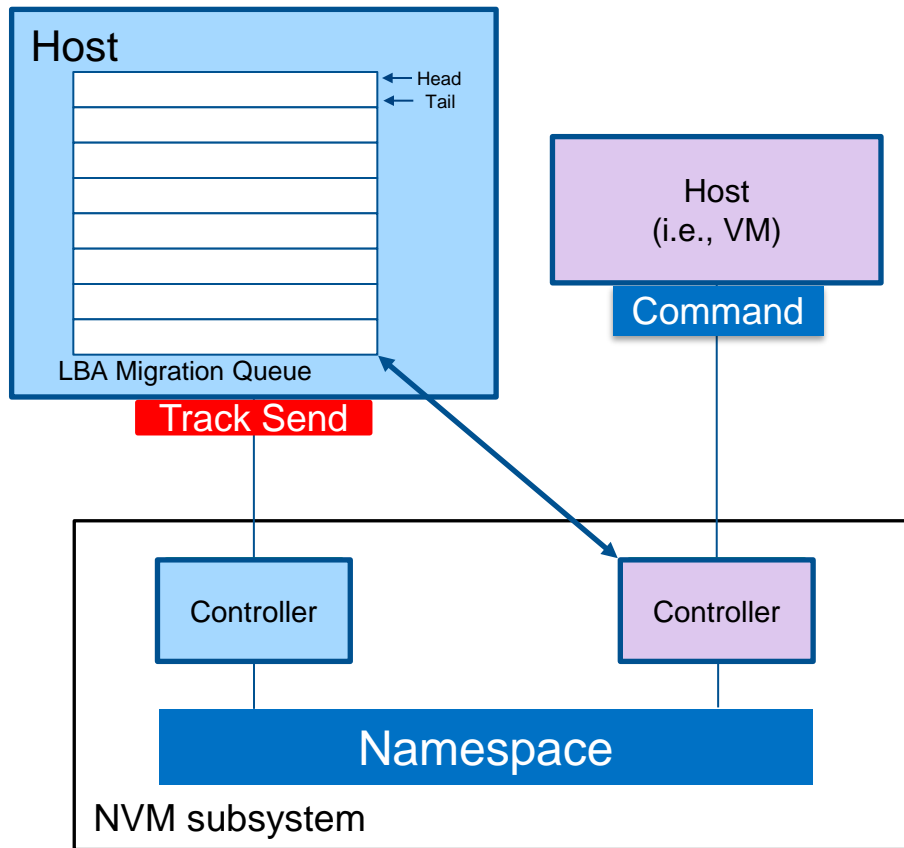
Live Migration of Attached Namespaces



Create a LBA Migration Queue

- Host allocates memory for a LBA Migration Queue
- Host submits a **Controller Data Queue** command to a controller specifying:
 - Controller Identifier of a controller to track LBA modifications by that controller
 - Log User Data Changes Management Operation
 - Host Memory Address
 - Dword Size
- Controller returns a Controller Data Queue Identifier

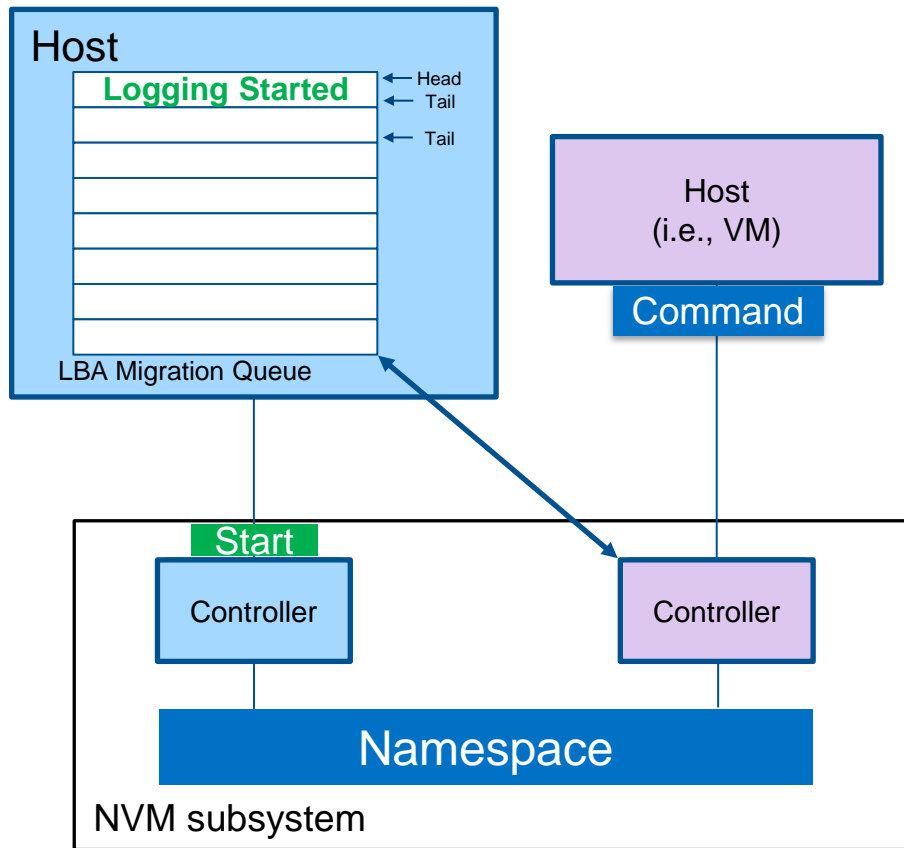
Live Migration of Attached Namespaces



Start logging LBA modification into the LBA Migration Queue

- Host submits a **Track Send** command to a controller specifying:
 - Log User Data Management Operation
 - Controller Data Queue Identifier of the LBA Migration Queue
 - Start Logging

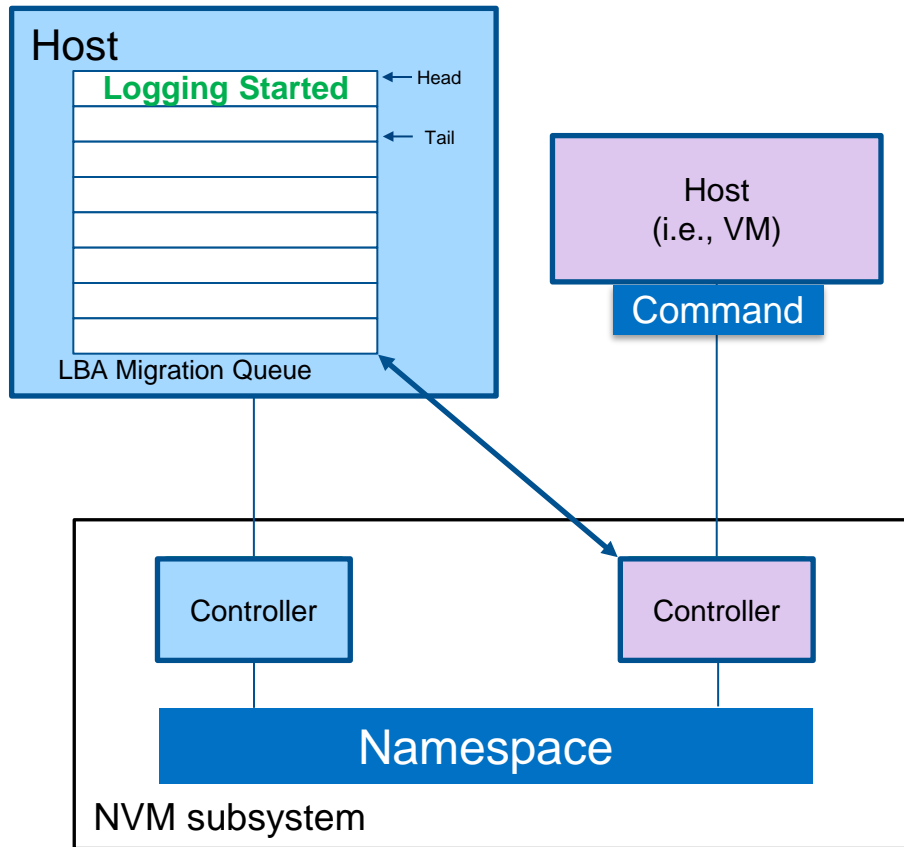
Live Migration of Attached Namespaces



Start logging LBA modification into the LBA Migration Queue

- Host submits a **Track Send** command to a controller specifying:
 - Log User Data Management Operation
 - Controller Data Queue Identifier of the LBA Migration Queue
 - Start Logging
- The controller that created the LBA Migration Queue to post entries specifying:
 - Changes logging
 - Start, stop, full, controller suspended
 - LBA modifications by the controller associated to the LBA migration queue

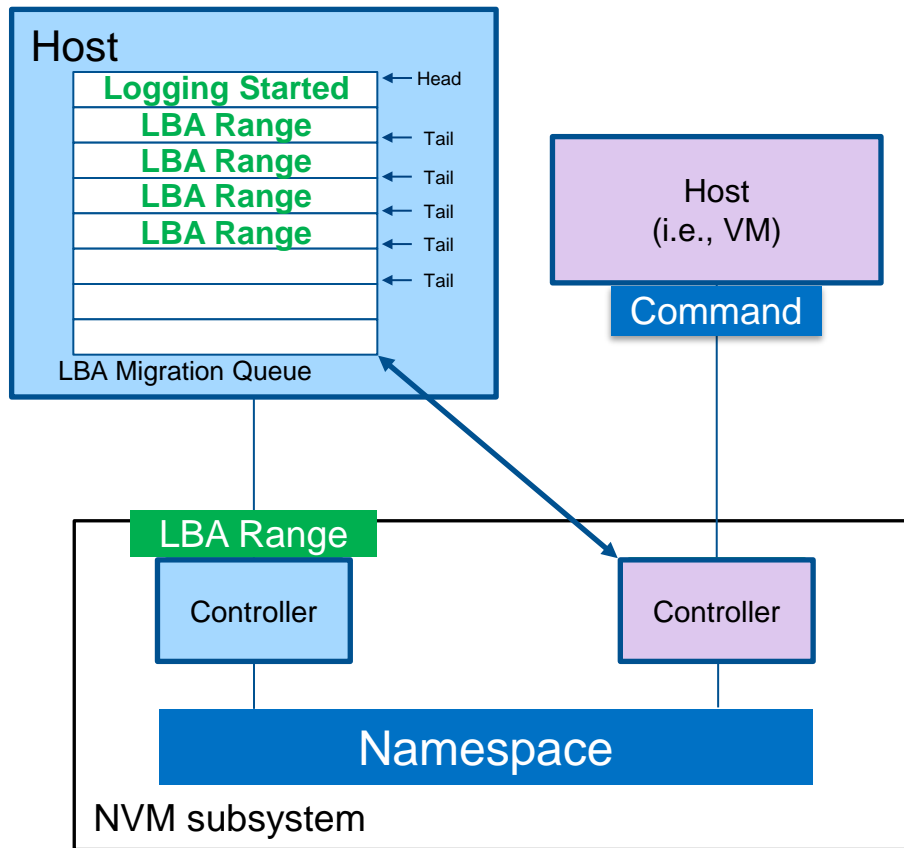
Live Migration of Attached Namespaces



Host copies namespace to another NVM Subsystem while changes are being logged

- Host may use the LBA Status command to copy the allocated logical blocks
- Ratified TP4165 Tracking LBA Allocation with Granularity

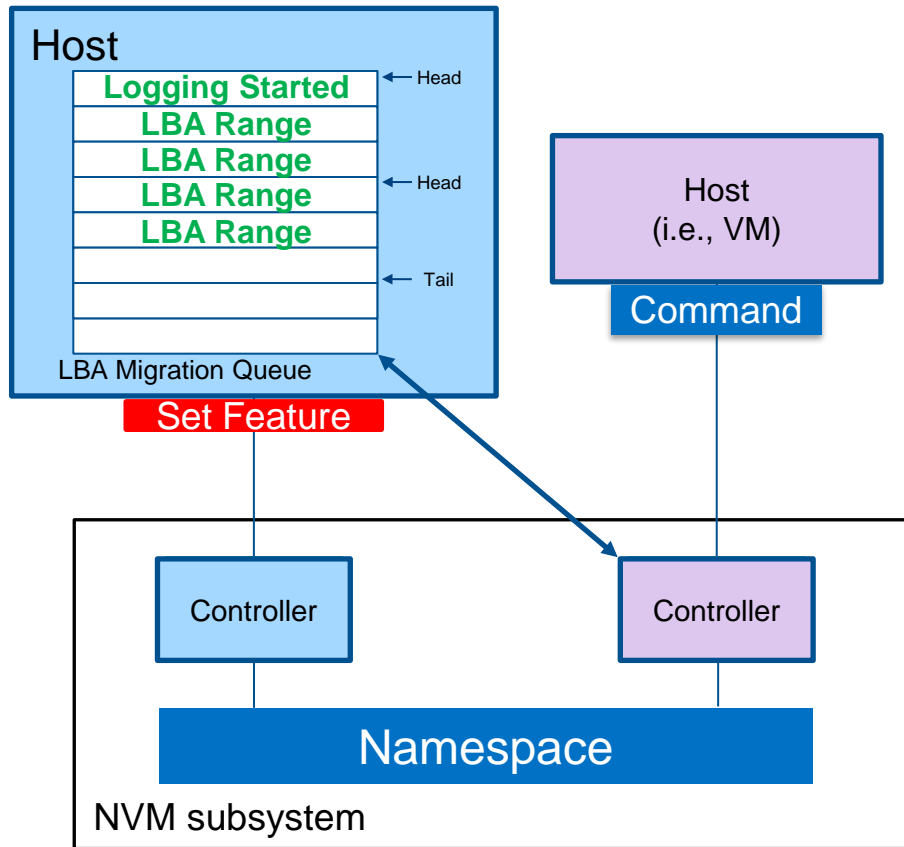
Live Migration of Attached Namespaces



Start logging LBA modification into the LBA Migration Queue

- Host submits a **Track Send** command to a controller specifying:
 - Log User Data Management Operation
 - Controller Data Queue Identifier of the LBA Migration Queue
 - Start Logging
- The controller that created the LBA Migration Queue to post entries specifying:
 - Changes logging
 - Start, stop, full, controller suspended
 - LBA modifications by the controller associated to the LBA migration queue

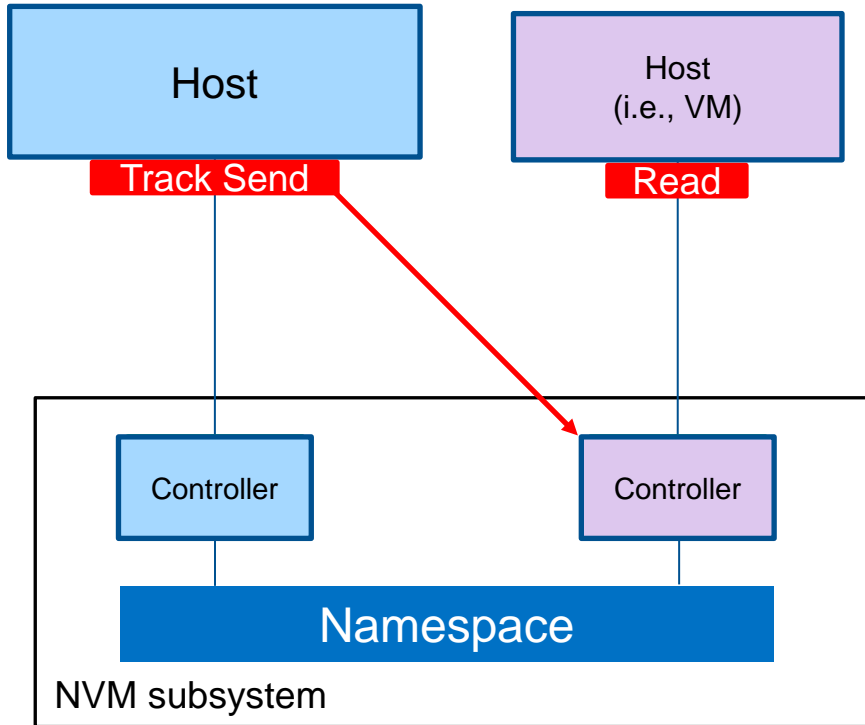
Live Migration of Attached Namespaces



Host processes the entries in the LBA Migration Queue

- Host submits a Set Features command specifying:
 - **Controller Data Queue** feature
 - Controller Data Queue Identifier of the LBA Migration Queue
 - Head Pointer Location
 - Optional Tail Pointer trigger to cause a **One-shot Controller Data Queue Tail Pointer** event


Tracking Host Memory Modification



Start tracking the host memory modifications by another controller due to processing NVMe[®] commands

- Host submits a **Track Send** command to a controller specifying:
 - Controller Identifier of a controller to track host memory modifications by that controller
 - Memory ranges


Tracking Host Memory Modification



Start tracking the host memory modifications by another controller due to processing NVMe[®] commands

- Host submits a **Track Send** command to a controller specifying:
 - Controller Identifier of a controller to track host memory modifications by that controller
 - Memory ranges
- As controller processes command
 - Remembers all host memory modifications within those Memory ranges

Tracking Host Memory Modification



Obtain the host memory medications

- Host submits a **Track Receive** command to a controller specifying:
 - Tracked Memory Ranges Management Operation
 - Size of the returned data
 - Controller Identifier of a controller to track host memory modifications by that controller
- the controller returns:
 - Suspended state of the specified controller
 - If more data is ready to be reported
 - Zero or more modified host memory ranges

TP4159 PCIe[®] Infrastructure for Live Migration

- Added a new Migration Send command and Migration Receive command:
 - Enables suspending and resuming a controller from processing NVMe[®] commands
 - Enables obtaining the state of a controller and setting the state of a controller
- Added a new Track Send command and a Track Receive command:
 - Enables the live migration of namespaces by identifying the logical blocks modified by a controller
 - Enables a host to obtain host memory modifications by a controller processing NVMe commands
- Added new Controller Data Queues to create a LBA Migration Queue
- Added a new Controller Data Queue feature
 - Allows a host to specify Head pointer updates and send a new One-shot event
- Added a new CNS value to the Identify command to report the supported controller state versions



Questions?



