NVM Express® Host-Managed Live Migration

Presented by Mike Allison

Sr. Director
NAND Product Planning – Standards
Samsung

SNIA COMPUTE, MEMORY, AND STORAGE SUMMIT

Solutions, Architectures, and Community

VIRTUAL EVENT, MAY 21-22, 2024

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload

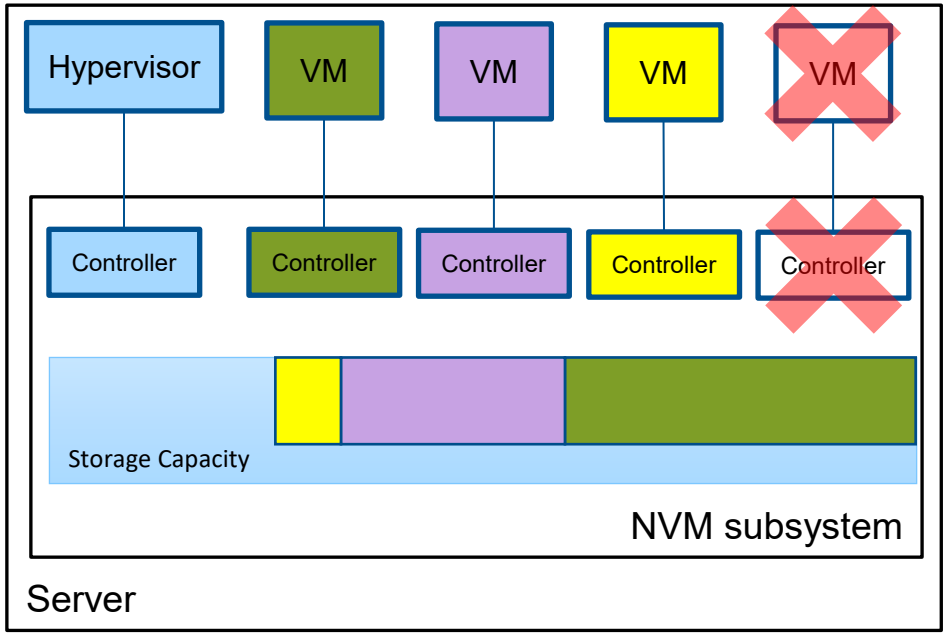

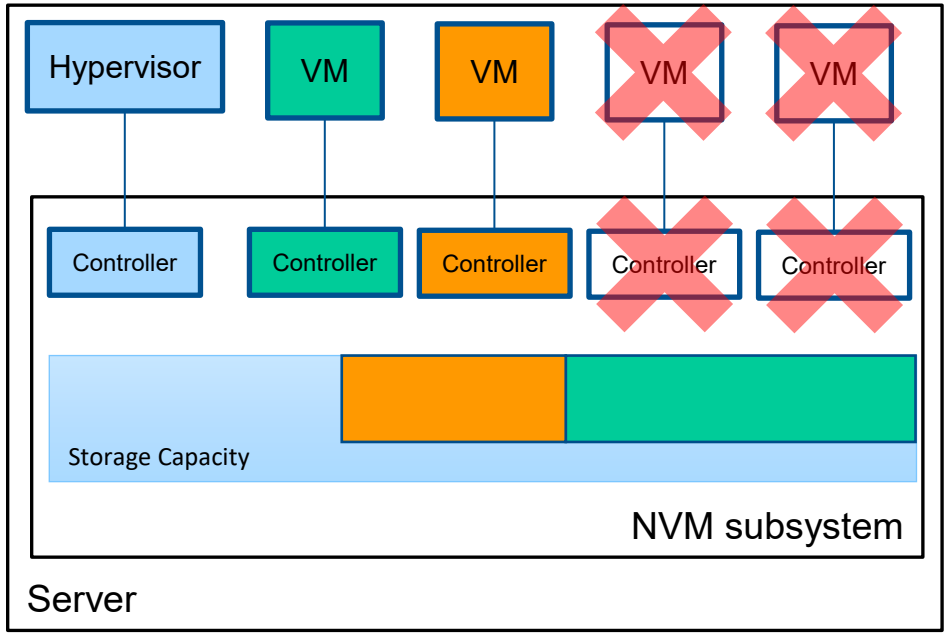

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload

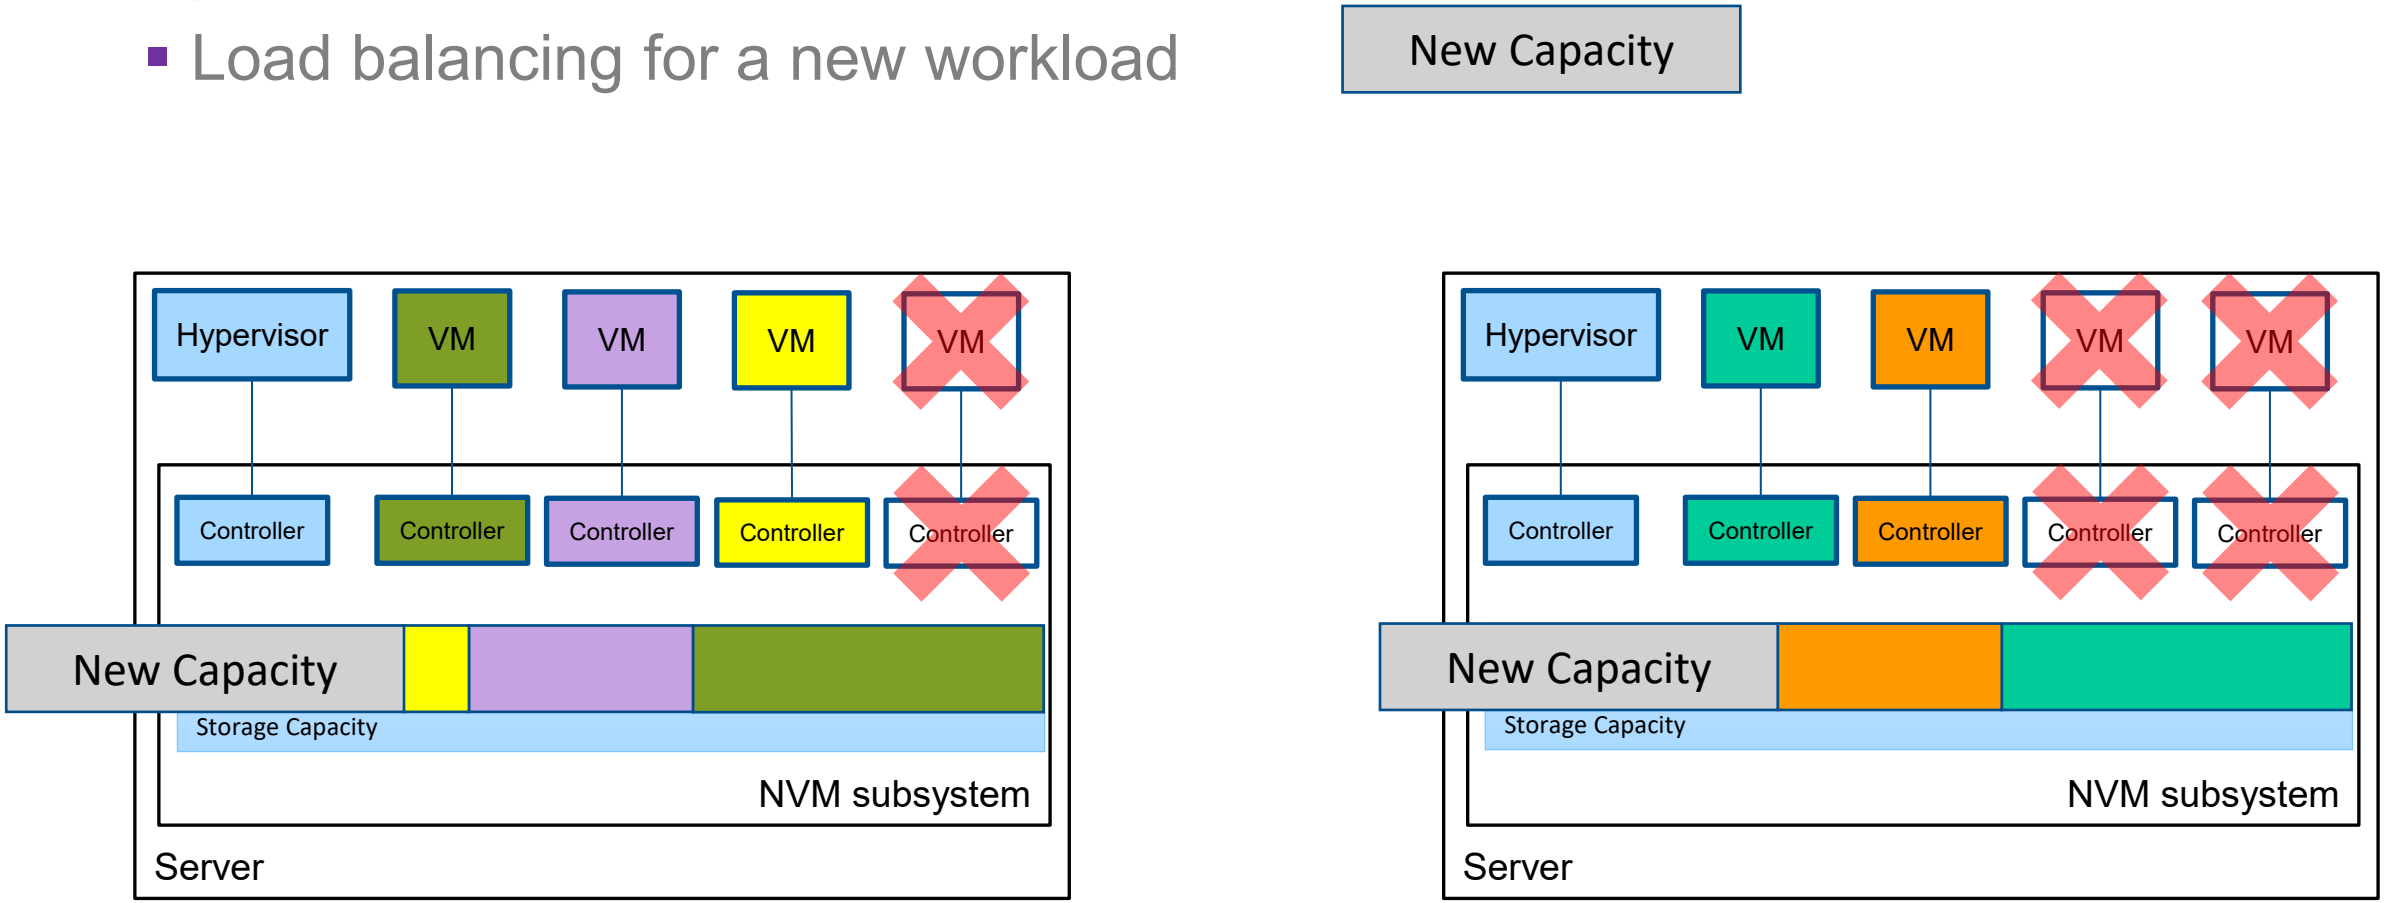

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload

New Capacity

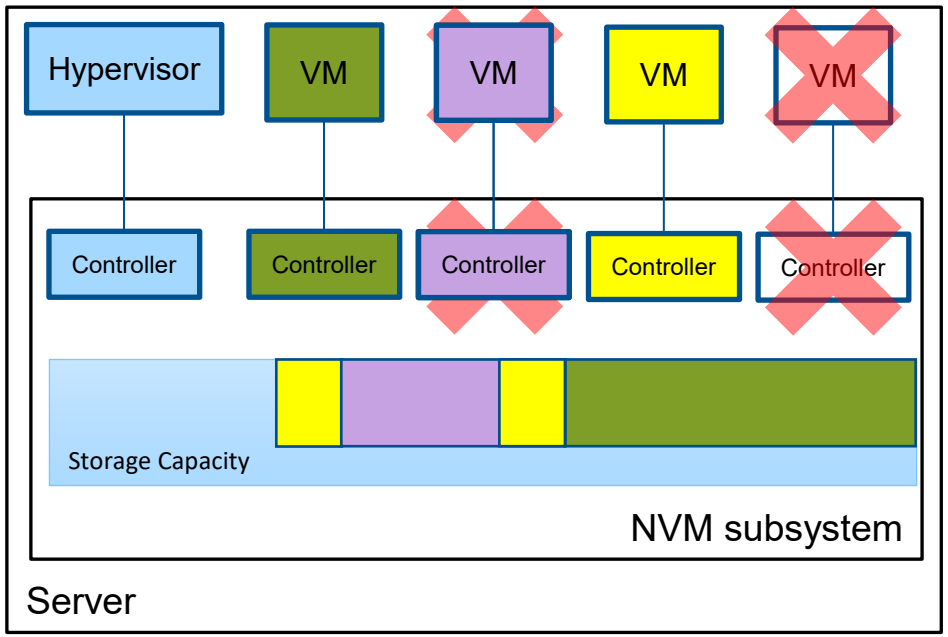

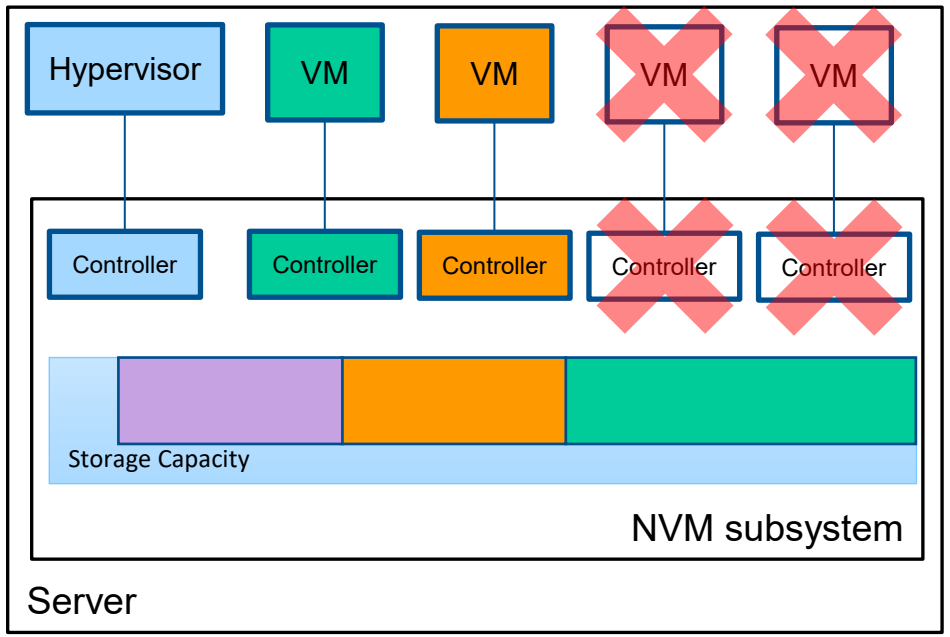

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload

New Capacity

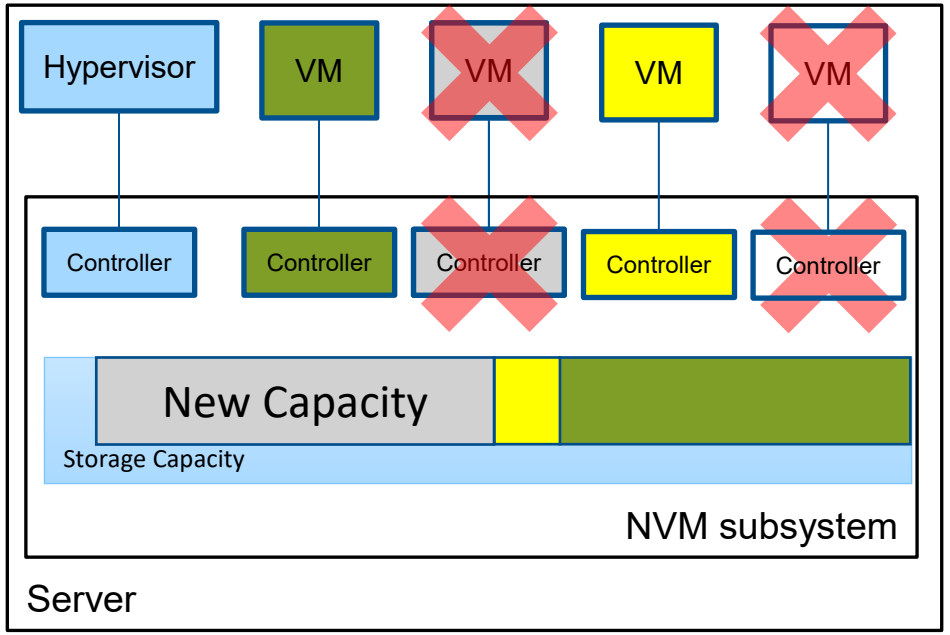

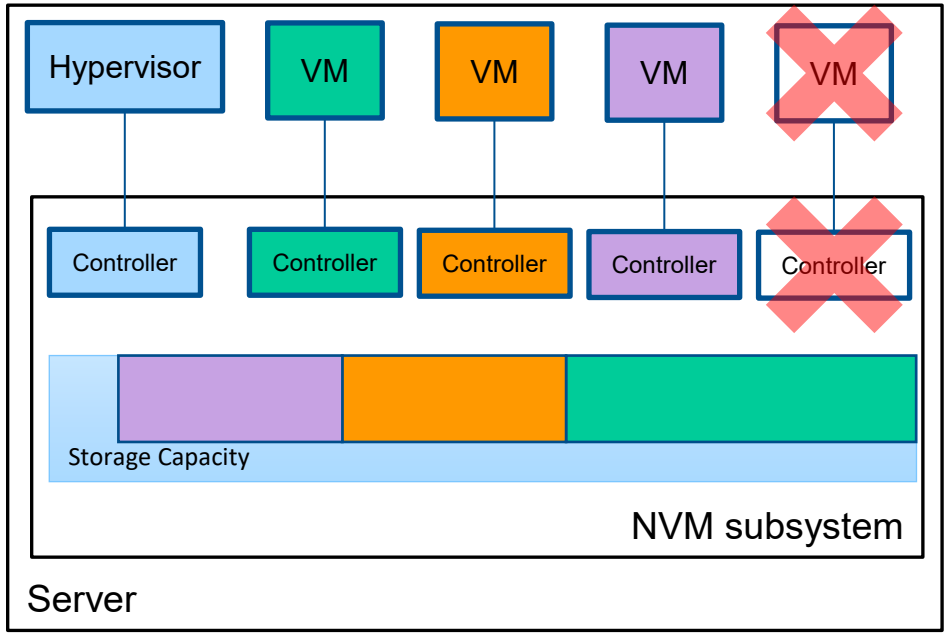

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload

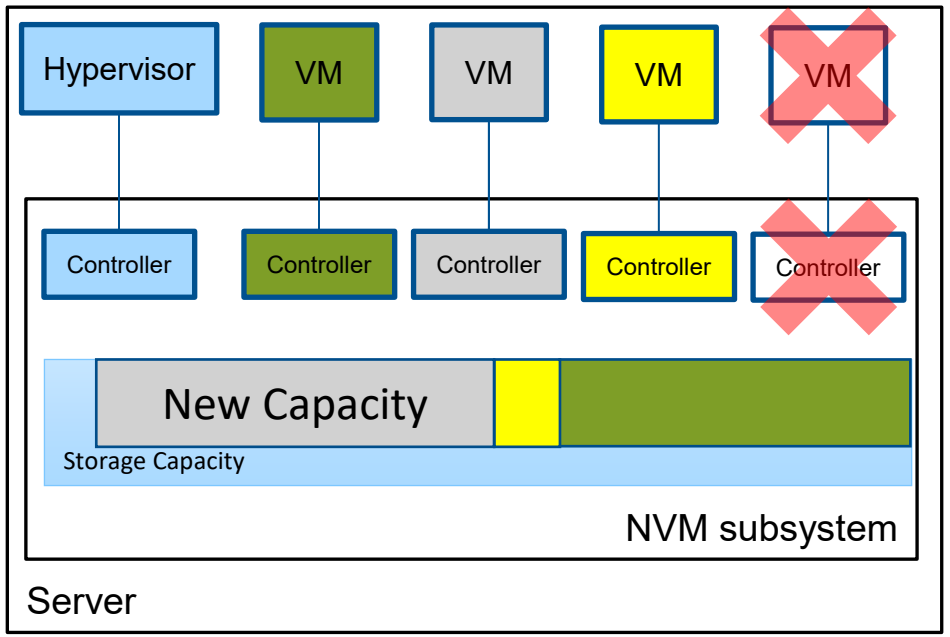

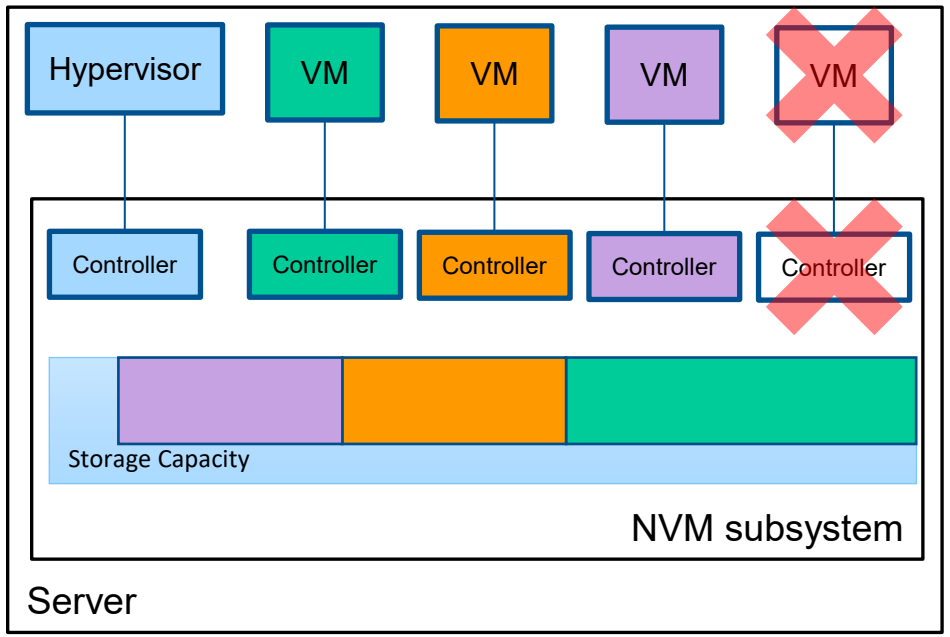

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload
  - Hardware maintenance

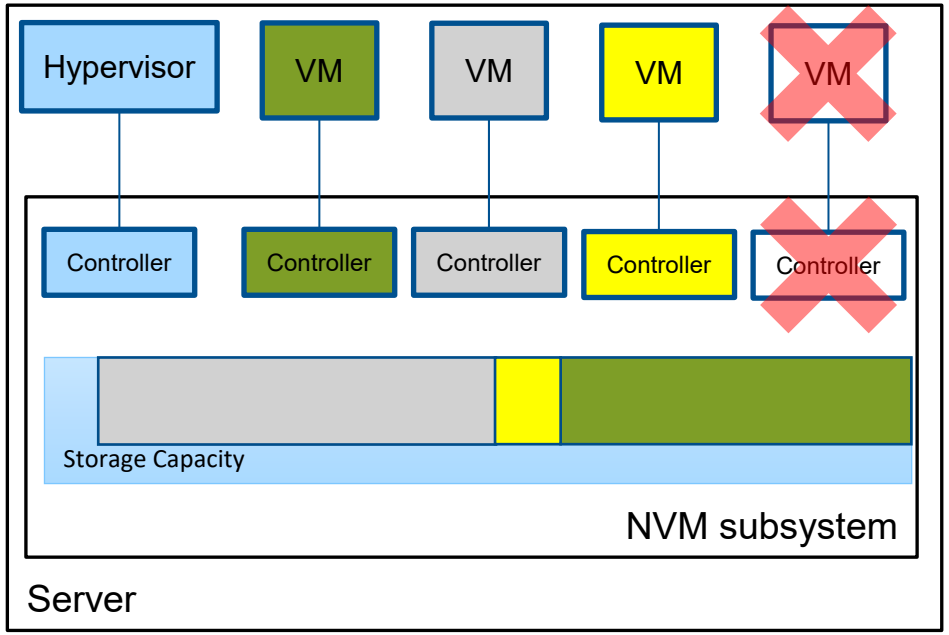

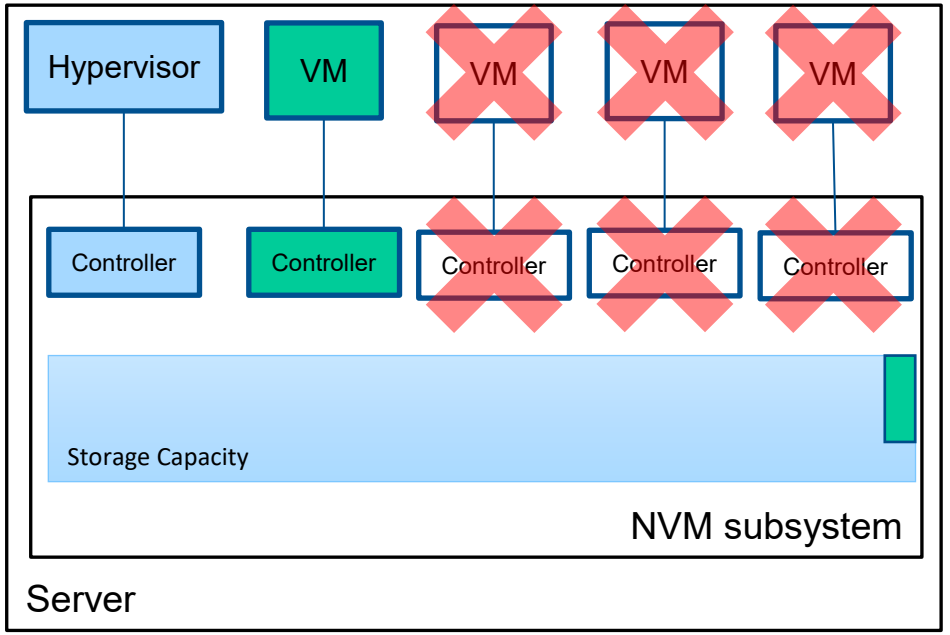

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload
  - Hardware maintenance

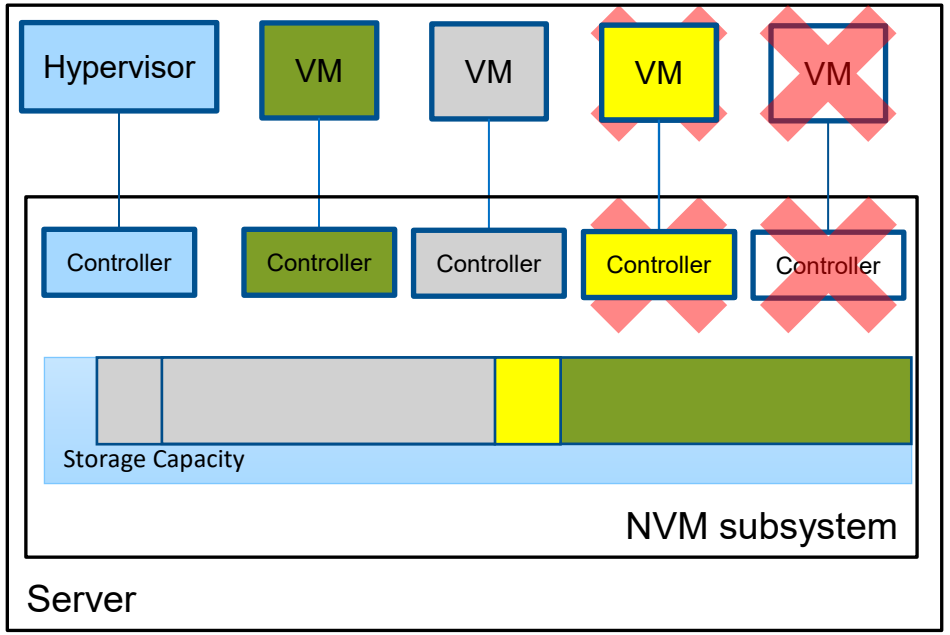

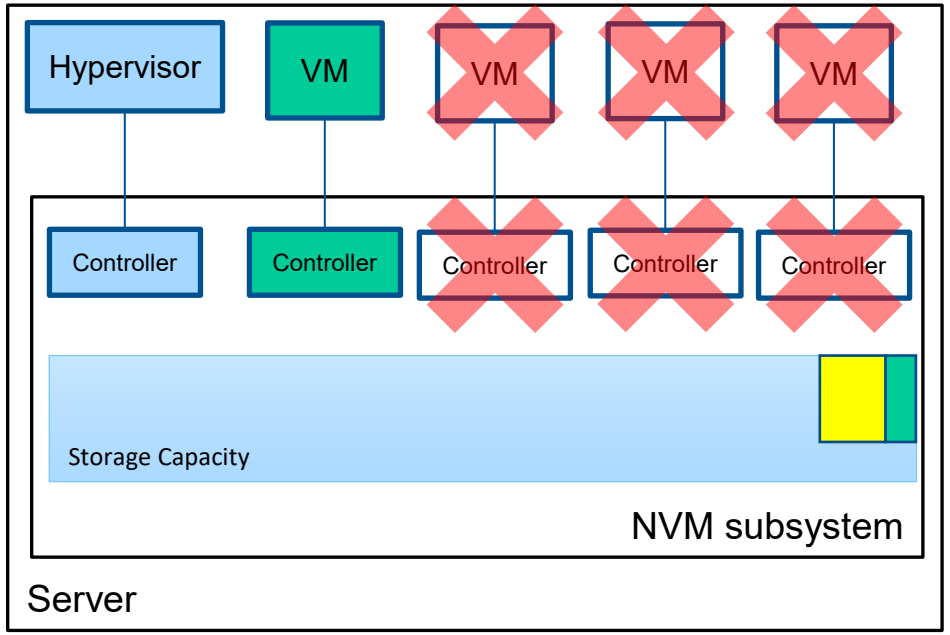

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload
  - Hardware maintenance

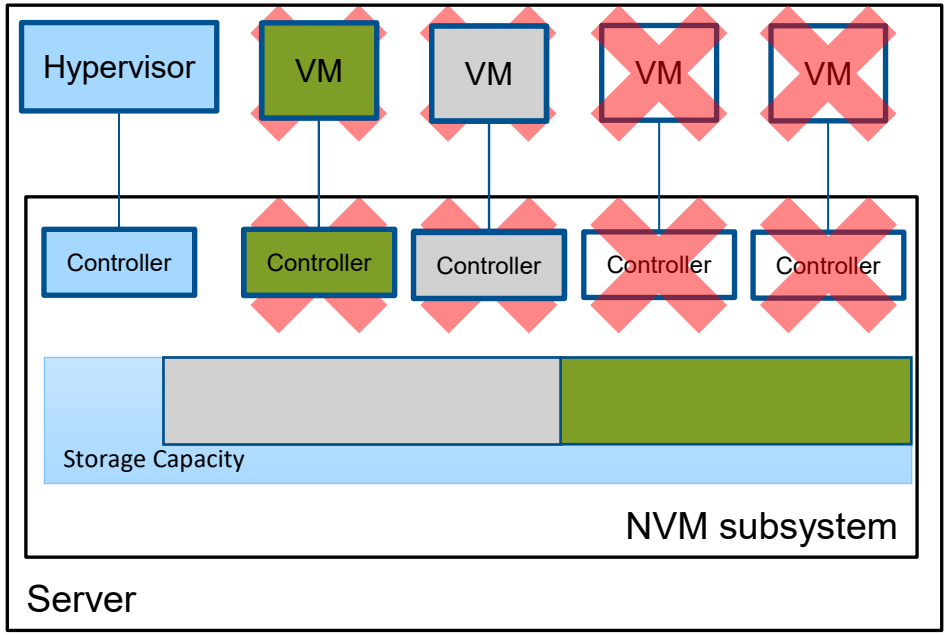

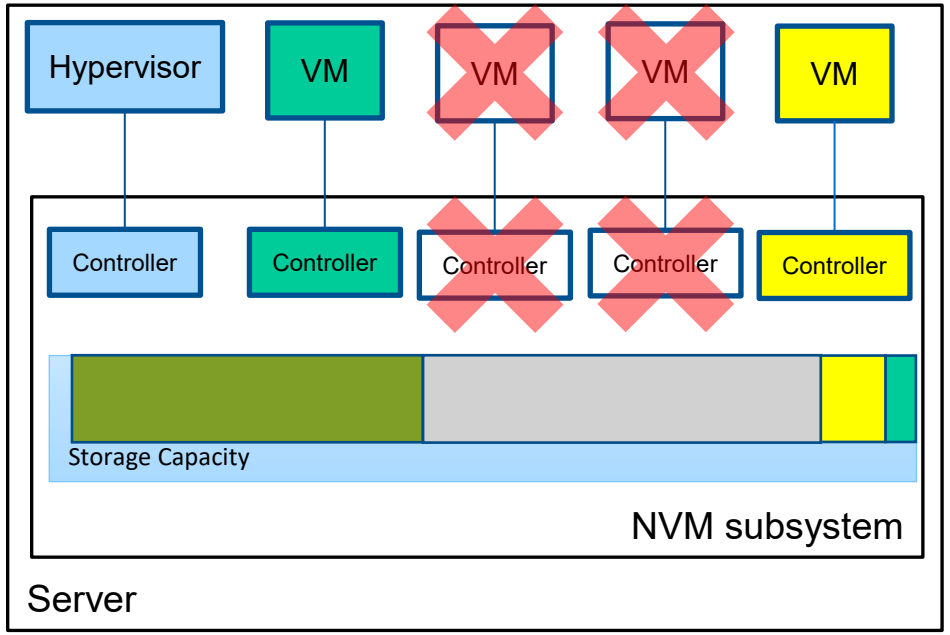

Why Live Migration in NVMe™?

- Key Use Cases:
  - Load balancing for a new workload
  - Hardware maintenance

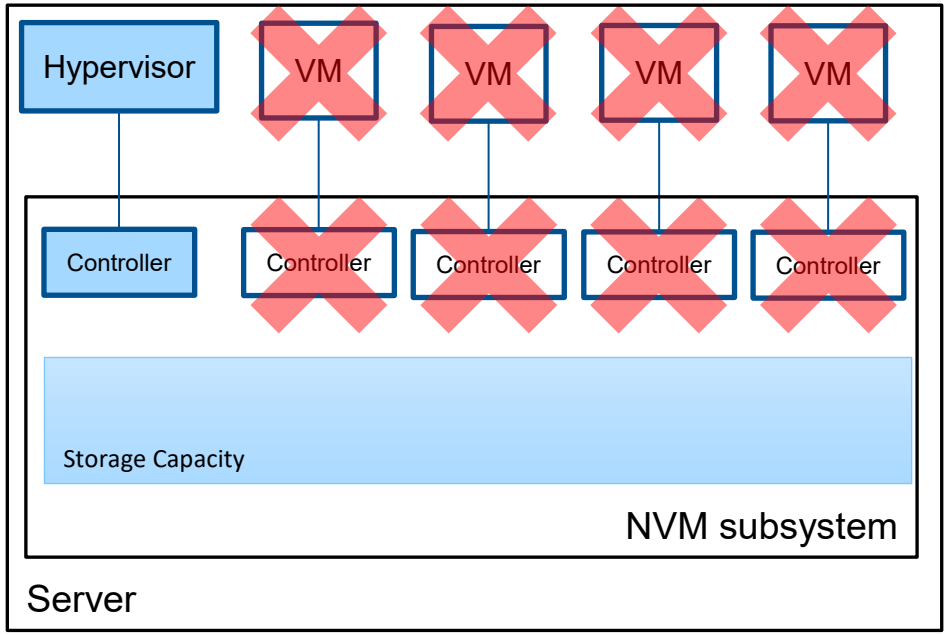

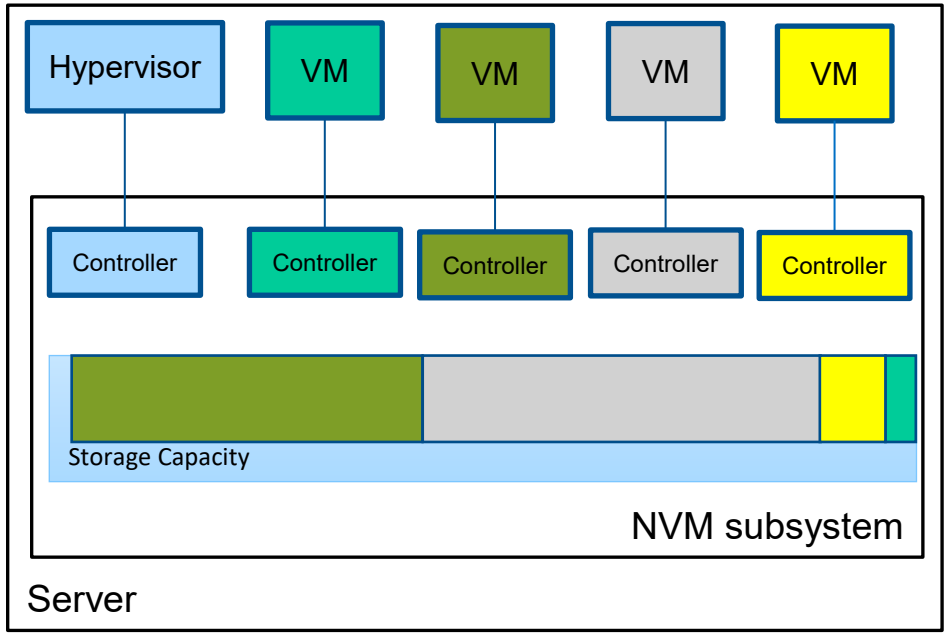

NVMe™ Host Managed Live Migration

- NVMe Host Managed Live Migration provides services to a Host to:
  - Get the NVMe state from a source controller
  - Put that NVMe state into the target controller and resume operation

- Source NVM Subsystem

NVMe™ Host Managed Live Migration

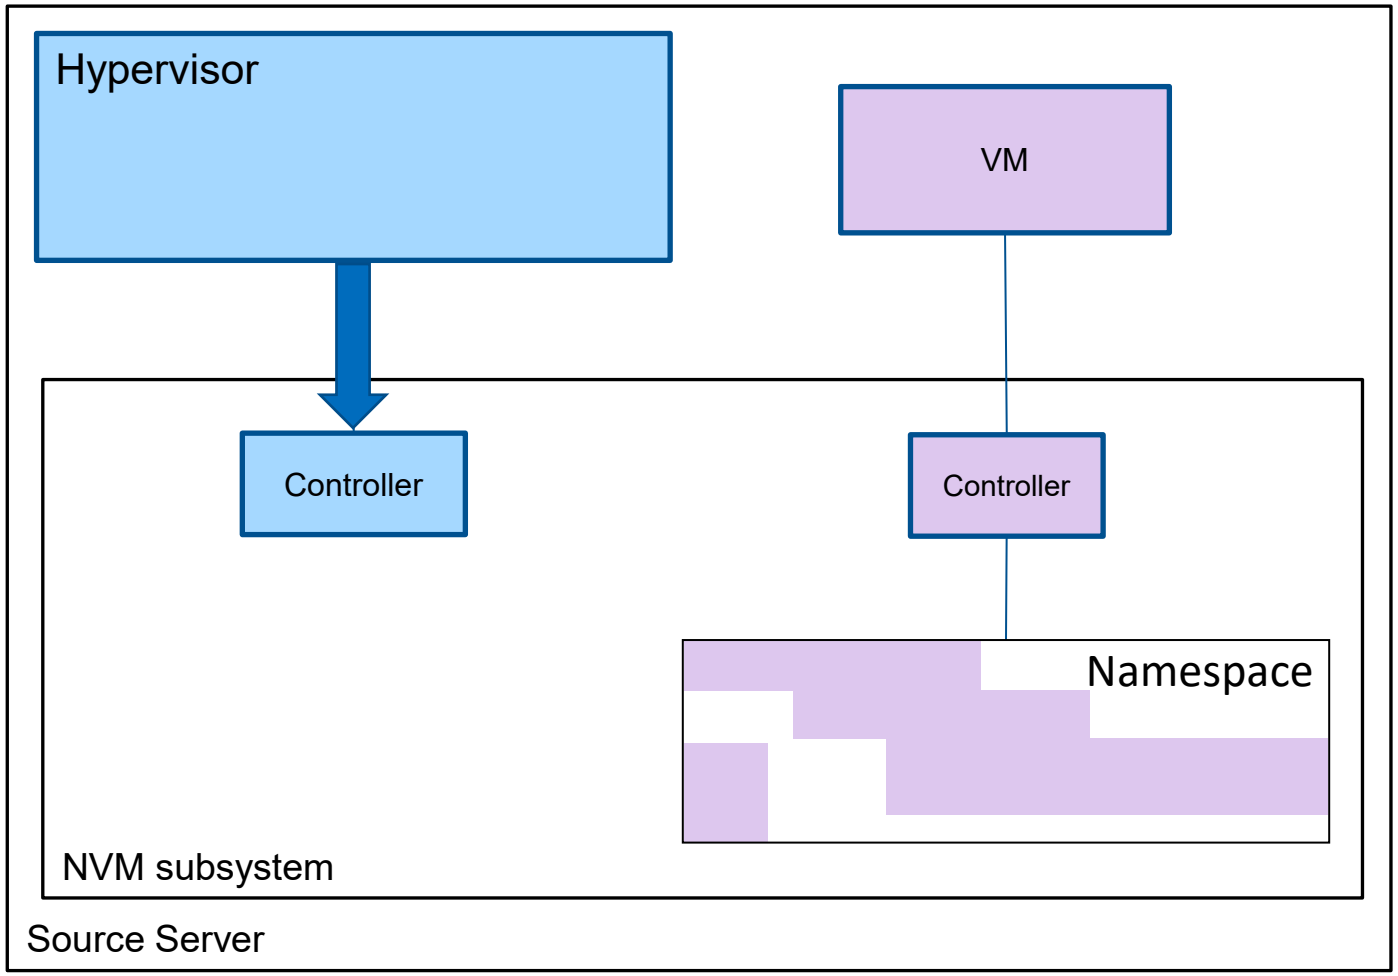

- Source NVM Subsystem
  - Read allocated logical blocks
    - Ratified TP4165 Tracking LBA Allocation with Granularity
    - Allows the host to request allocated LBA reporting

NVMe™ Host Managed Live Migration

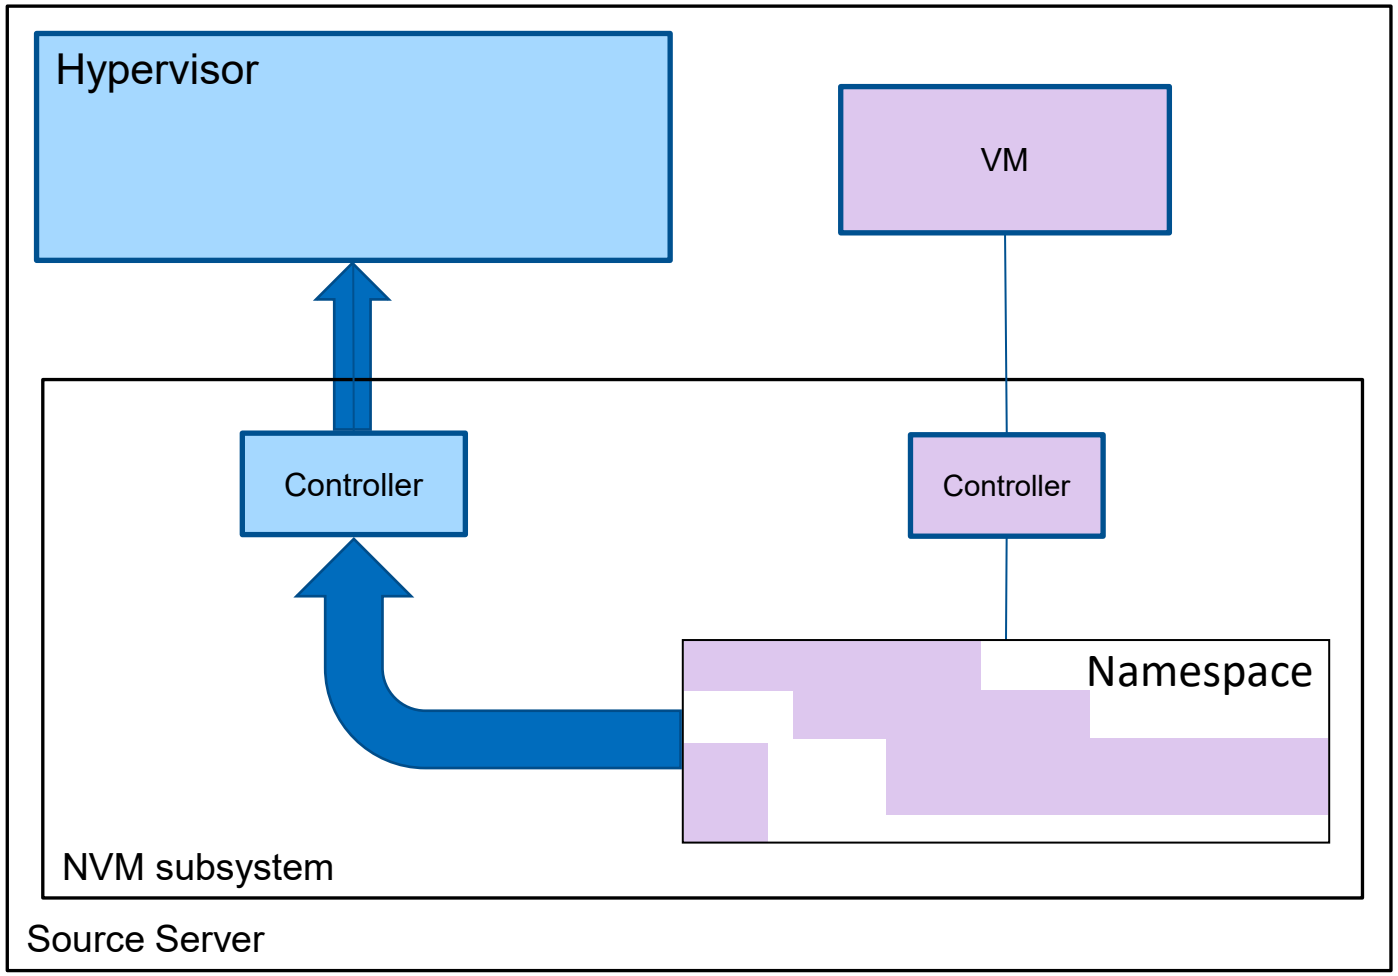

- Source NVM Subsystem
  - Read allocated logical blocks
    - Ratified TP4165 Tracking LBA Allocation with Granularity
    - Allows the host to request allocated LBA reporting
    - Controller reports sequential allocated LBAs on a reported granularity
    - Can be used for other use cases like Snapshots

NVMe™ Host Managed Live Migration

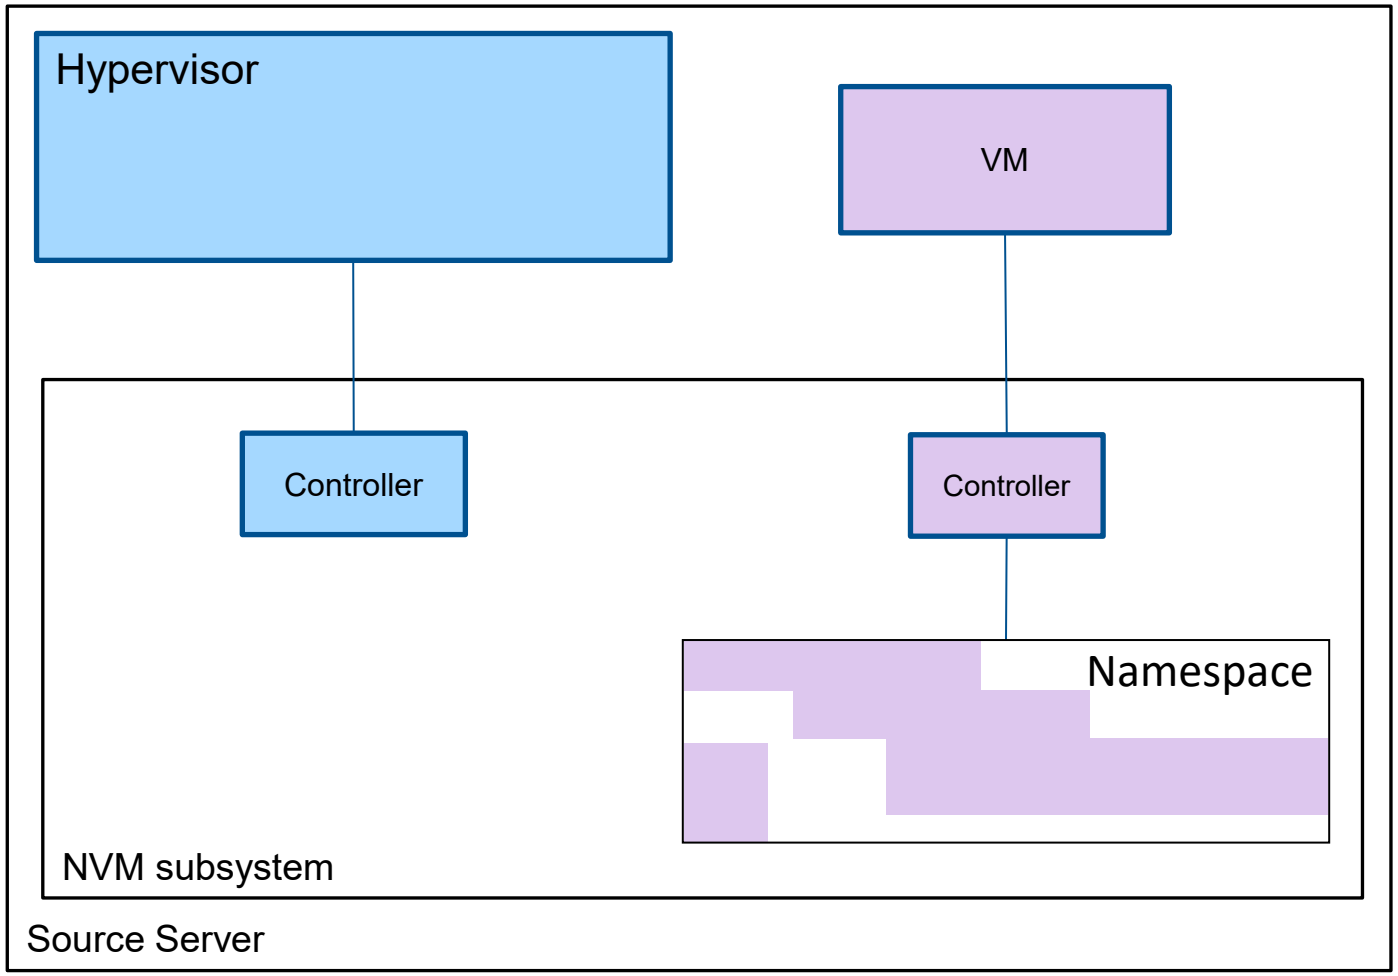

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
      - For the use case of migrating PCIe SSD namespaces to another PCIe SSDs

NVMe™ Host Managed Live Migration

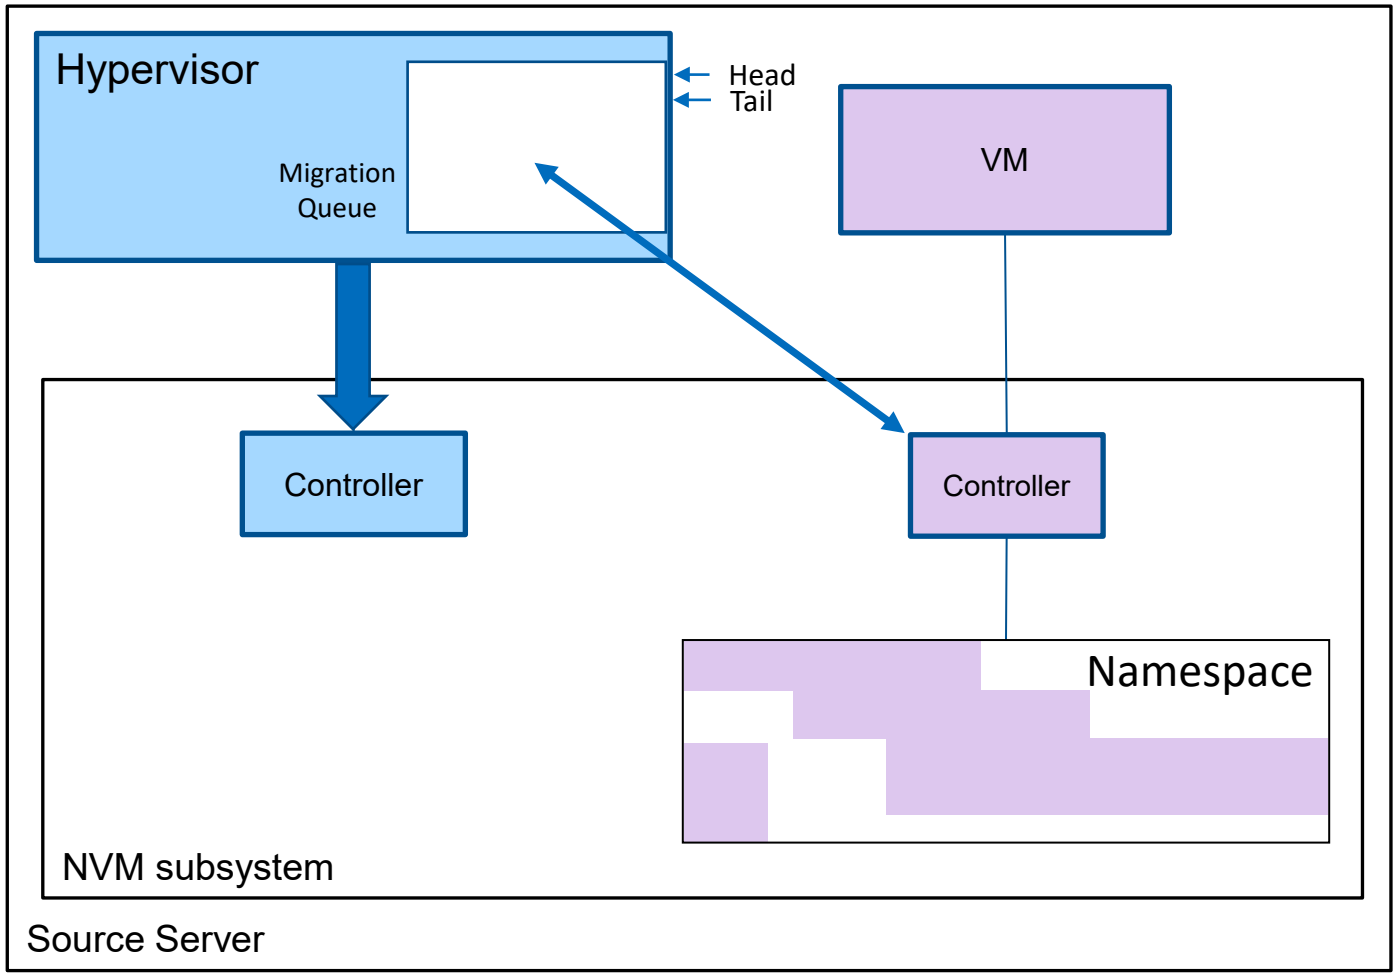

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
      - Create a Migration Queue
      - Start Logging

NVMe™ Host Managed Live Migration

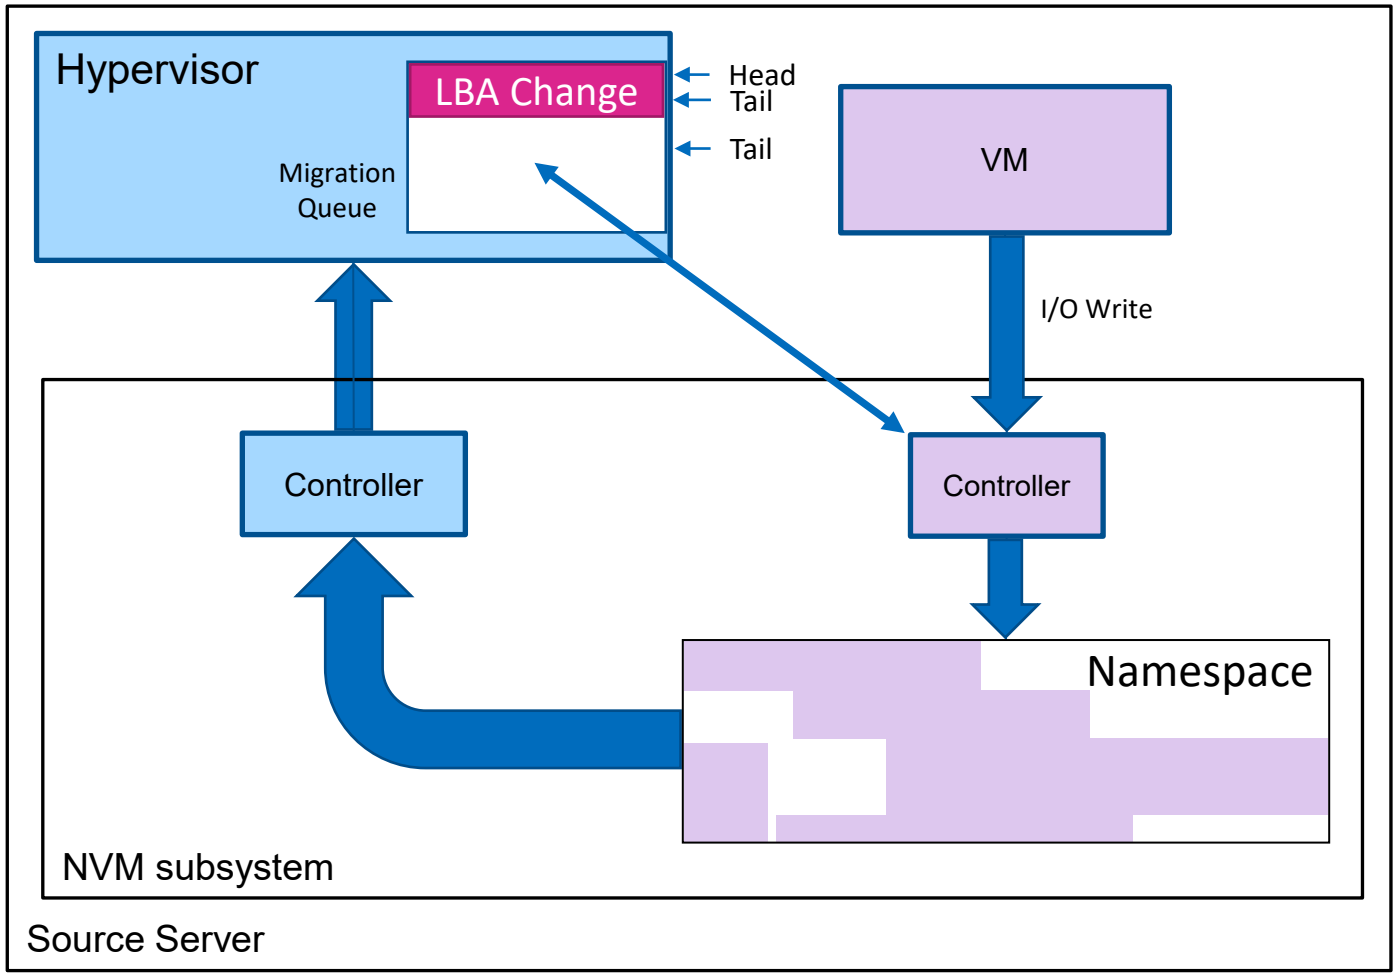

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
      - Create a Migration Queue
      - Start Logging
      - Post logical block changes

NVMe™ Host Managed Live Migration

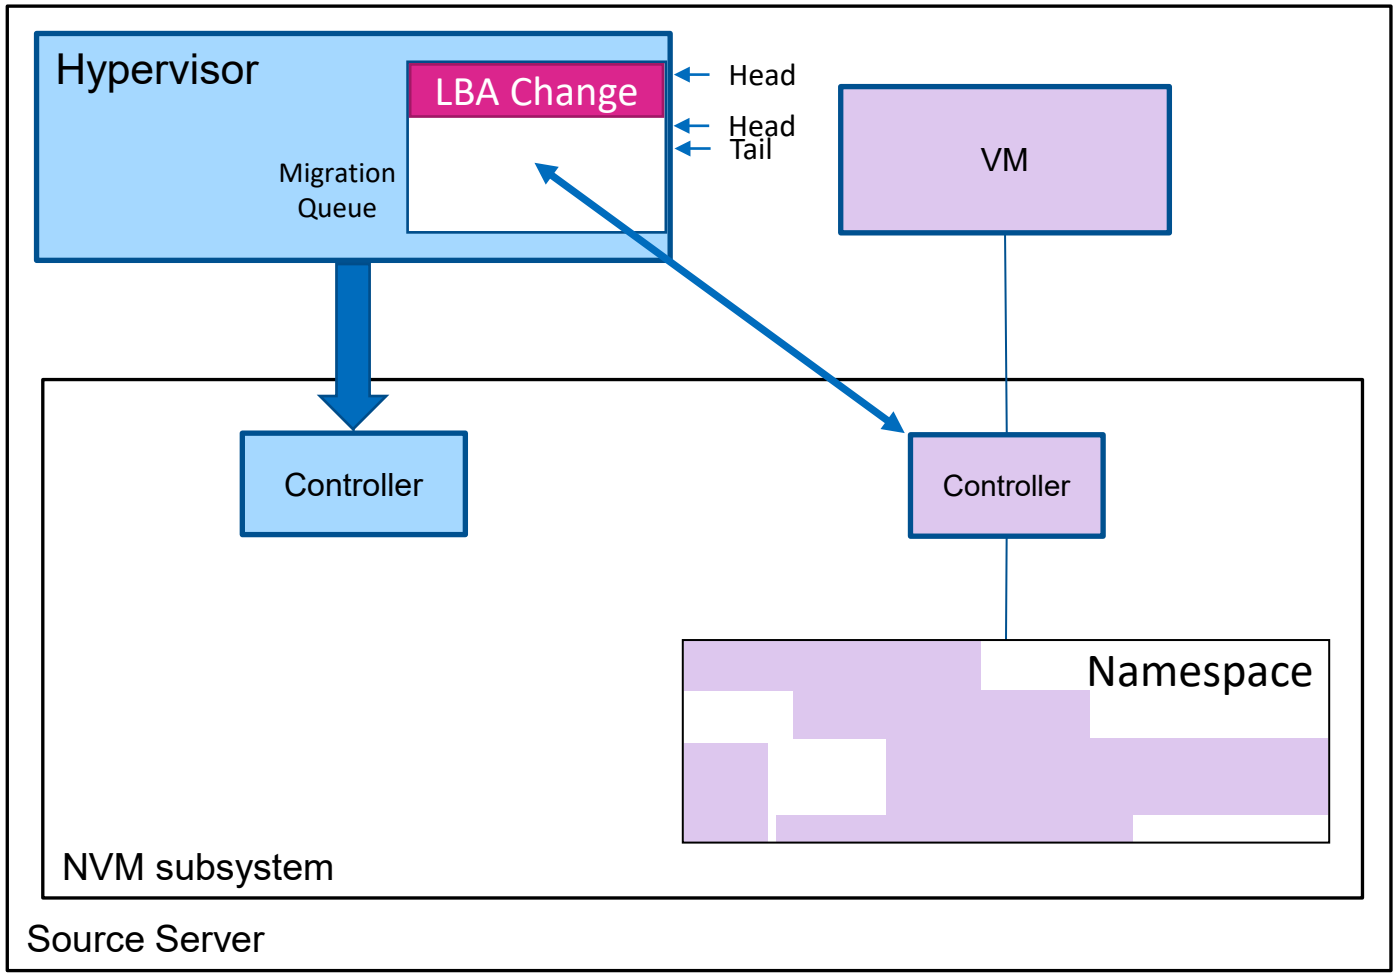

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
      - Create a Migration Queue
      - Start Logging
      - Post logical block changes
      - Update the Head pointer

NVMe™ Host Managed Live Migration

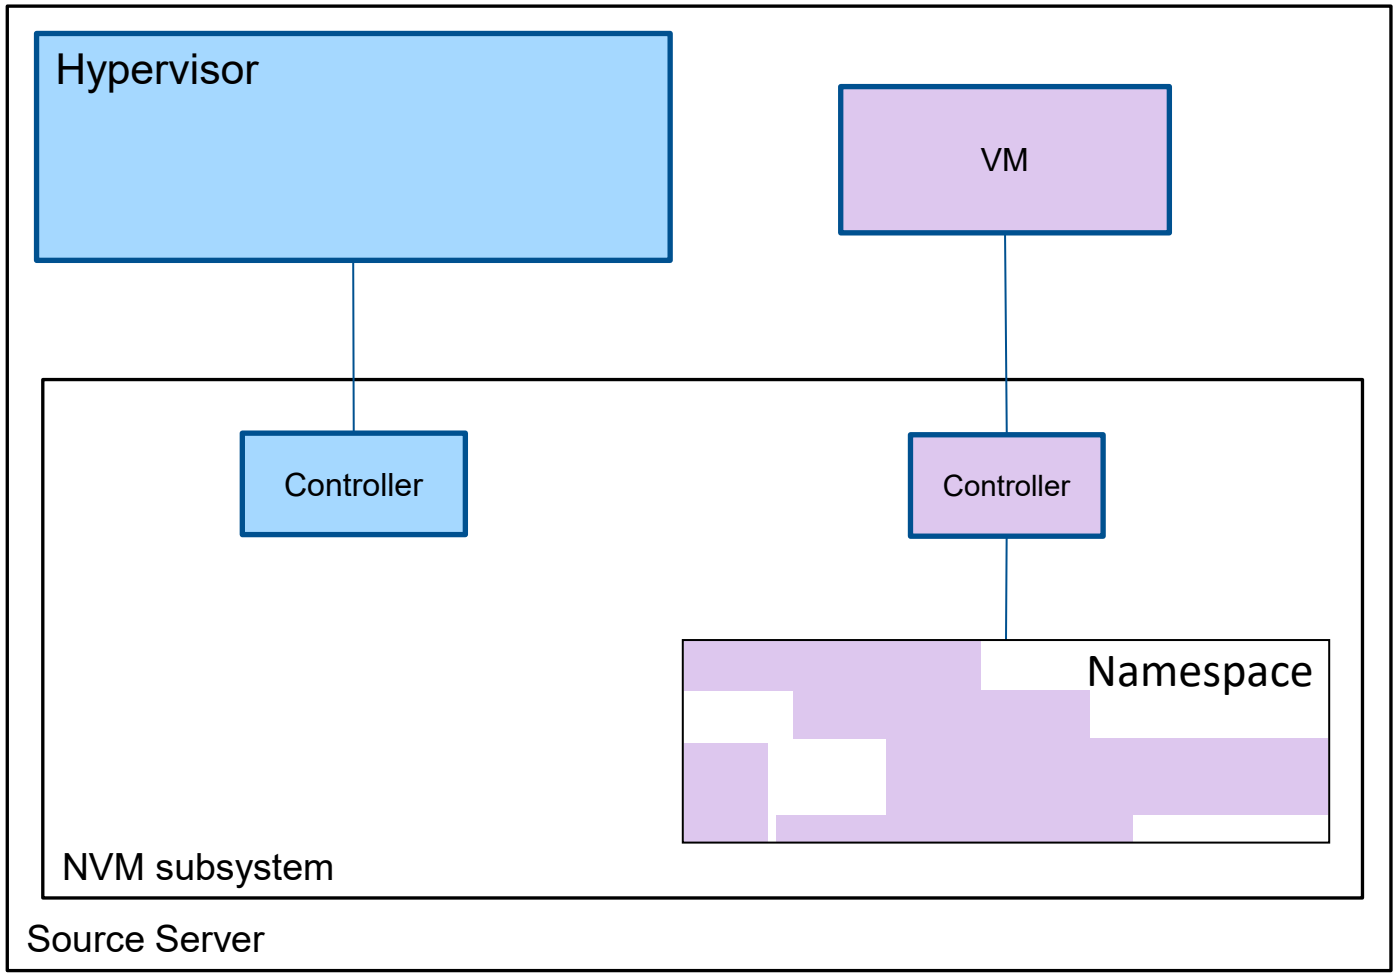

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
      - Use case for migrating DPU controllers using NVMe-oF

NVMe™ Host Managed Live Migration

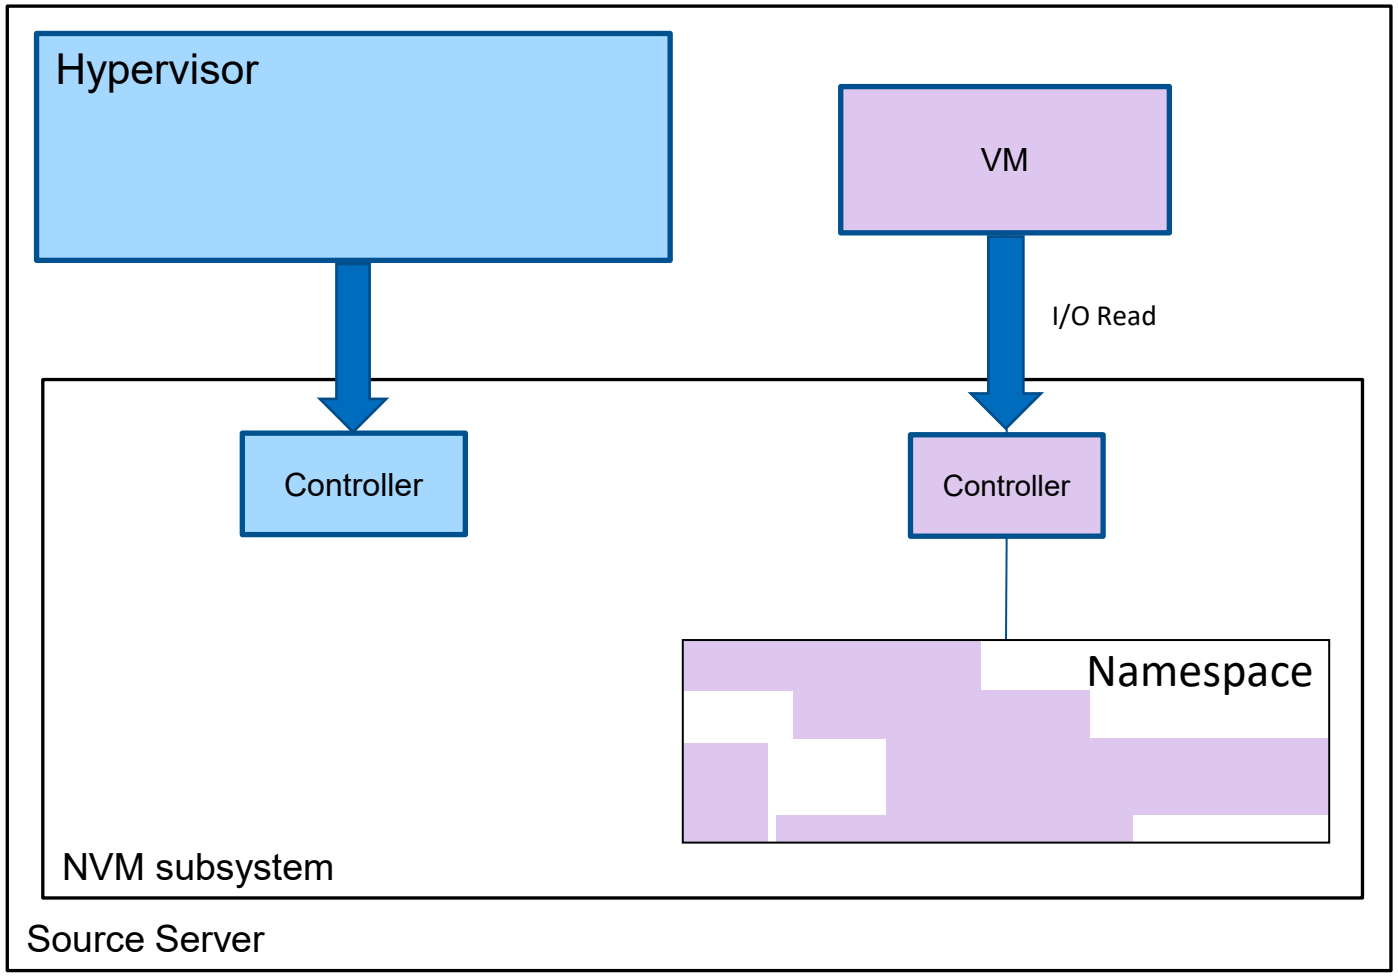

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
      - Start logging and aggregating changes

NVMe™ Host Managed Live Migration

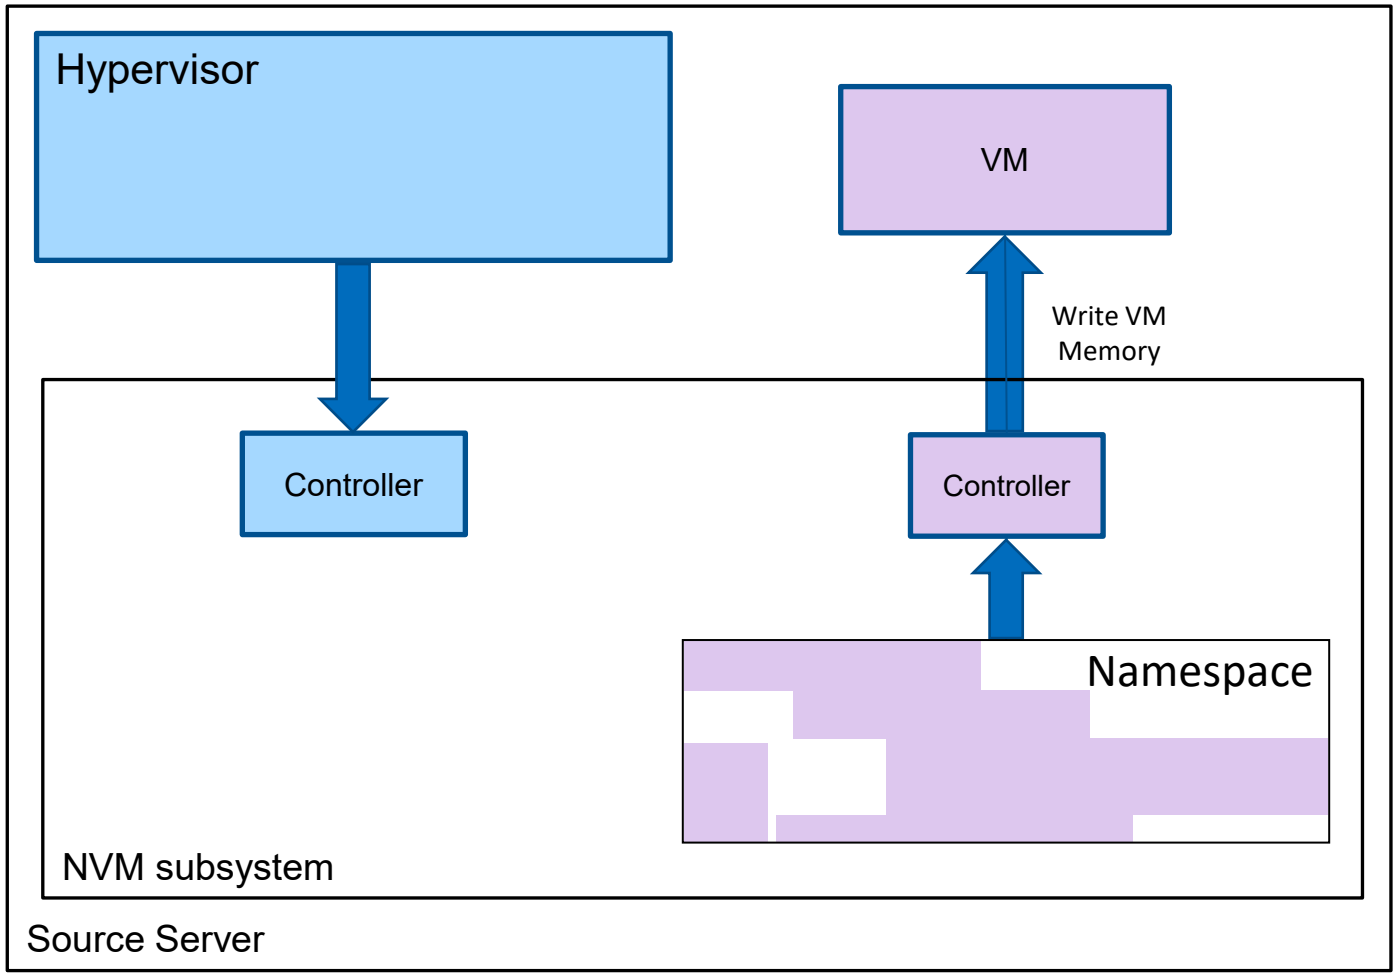

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
      - Start logging and aggregating changes
      - Query aggregated changes

NVMe™ Host Managed Live Migration

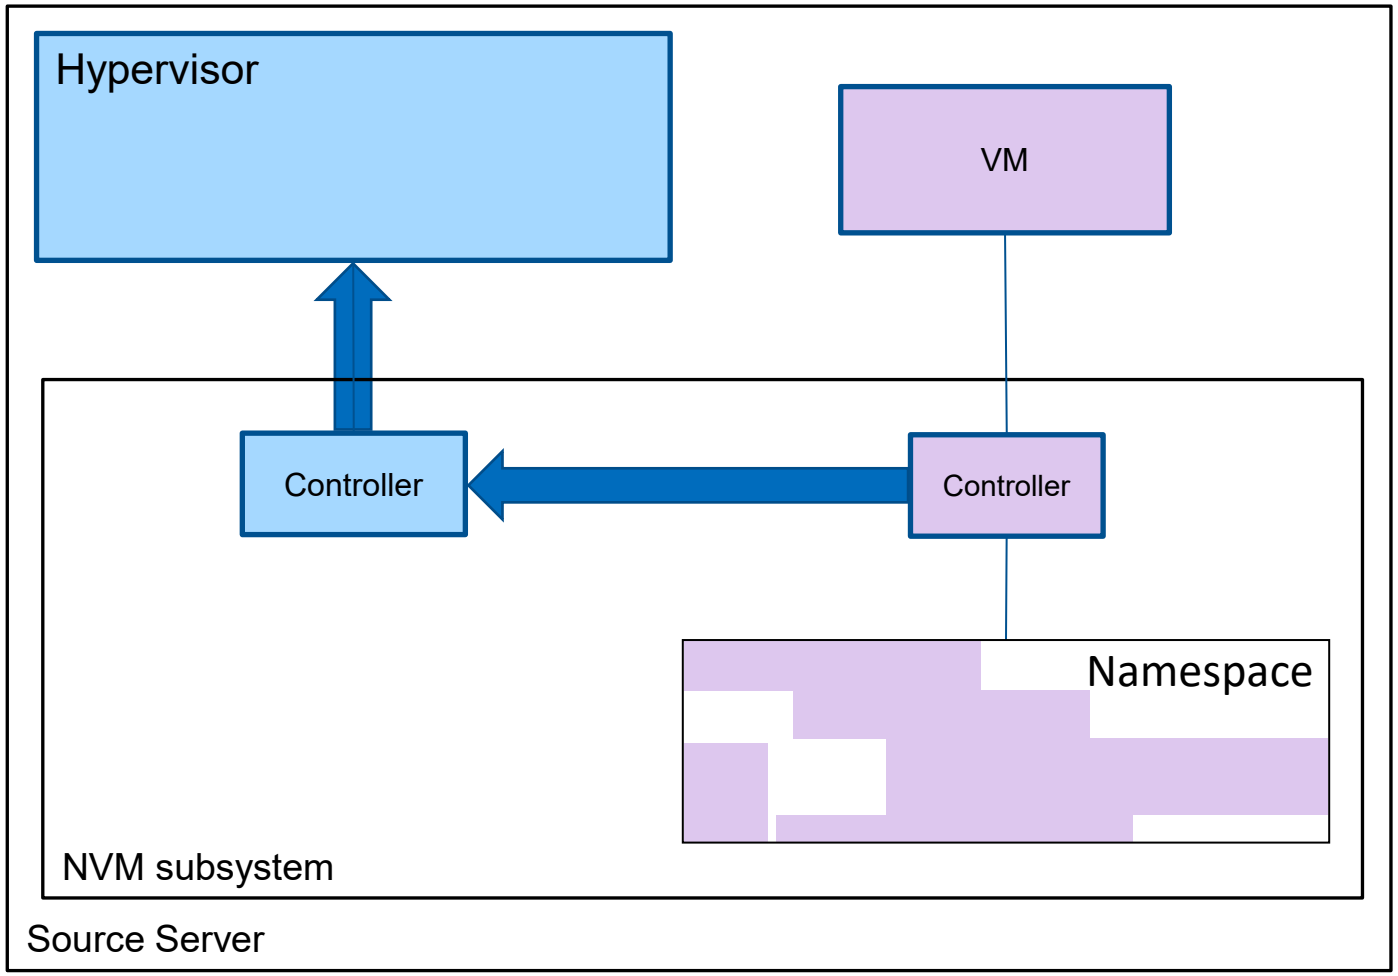

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
      - Start logging and aggregating changes
      - Query aggregated changes

NVMe™ Host Managed Live Migration

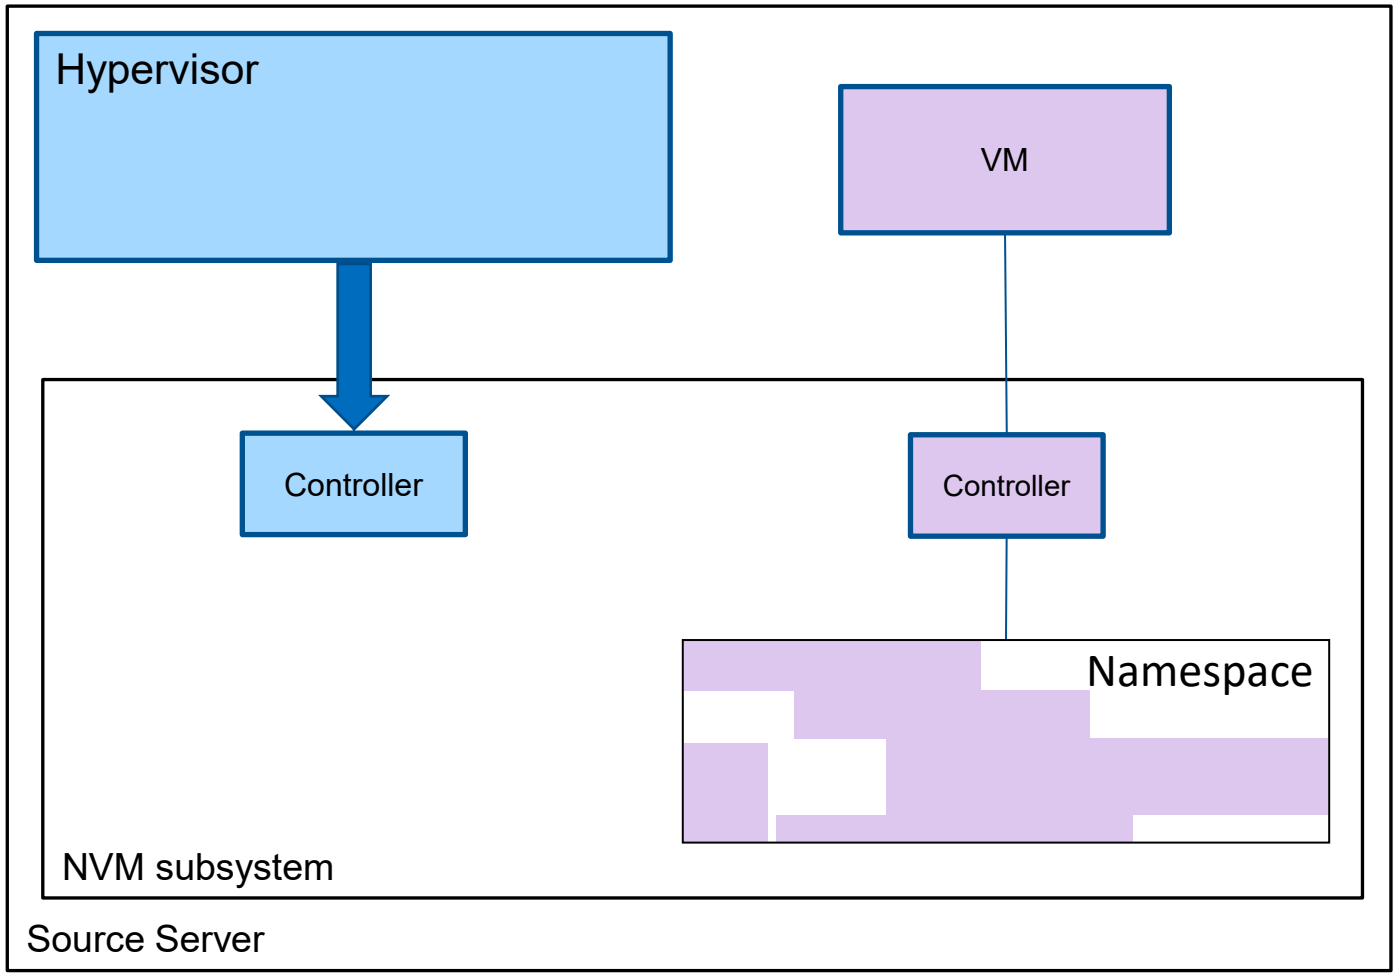

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
  - Read controller state

NVMe™ Host Managed Live Migration

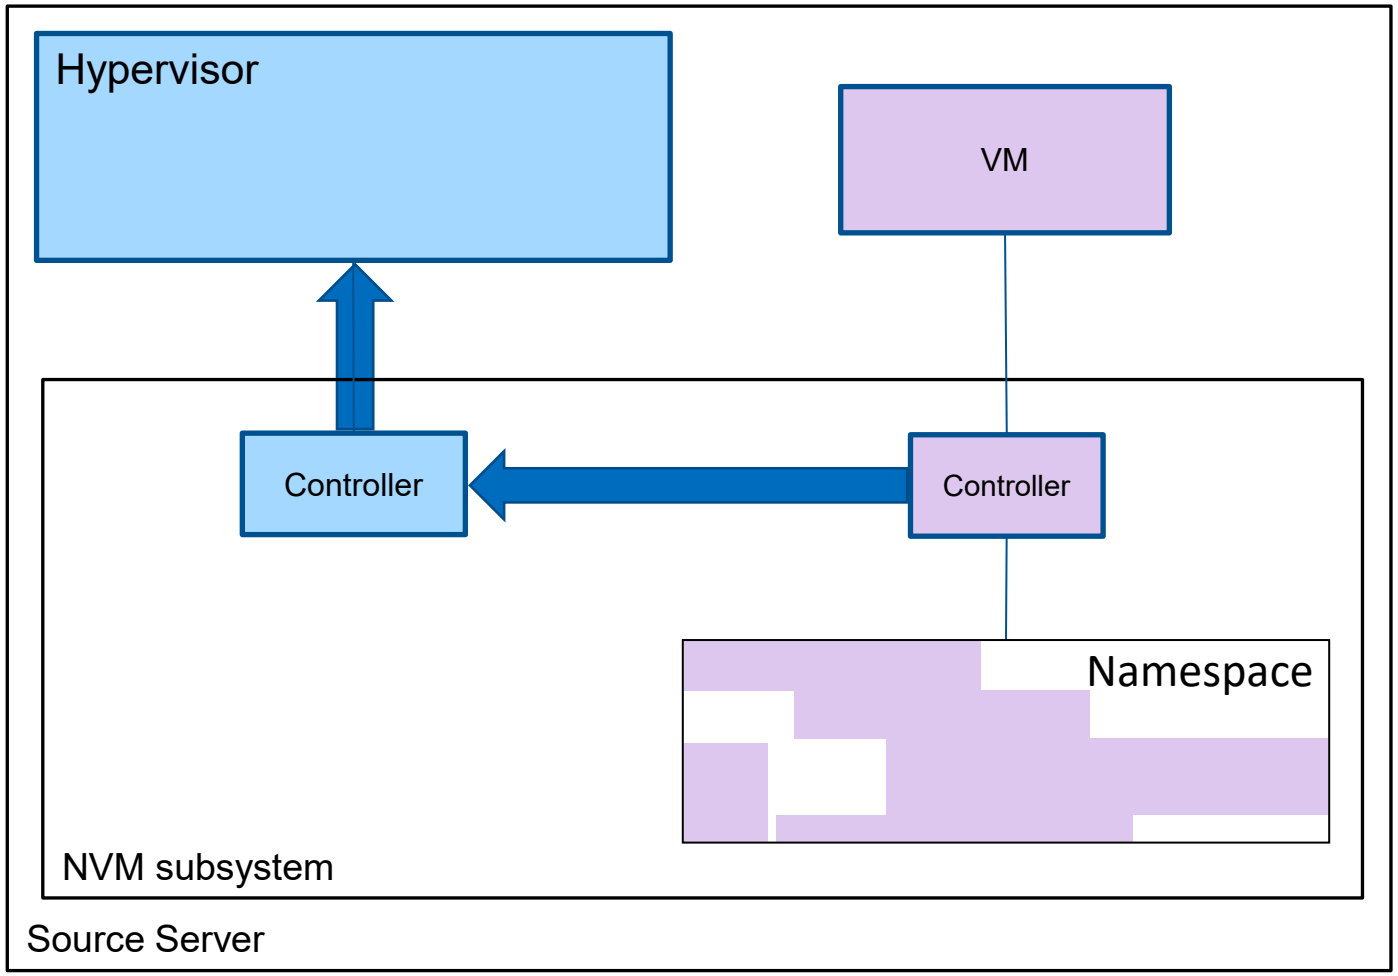

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
  - Read controller state
    - NVMe state
    - Vendor specific state

NVMe™ Host Managed Live Migration

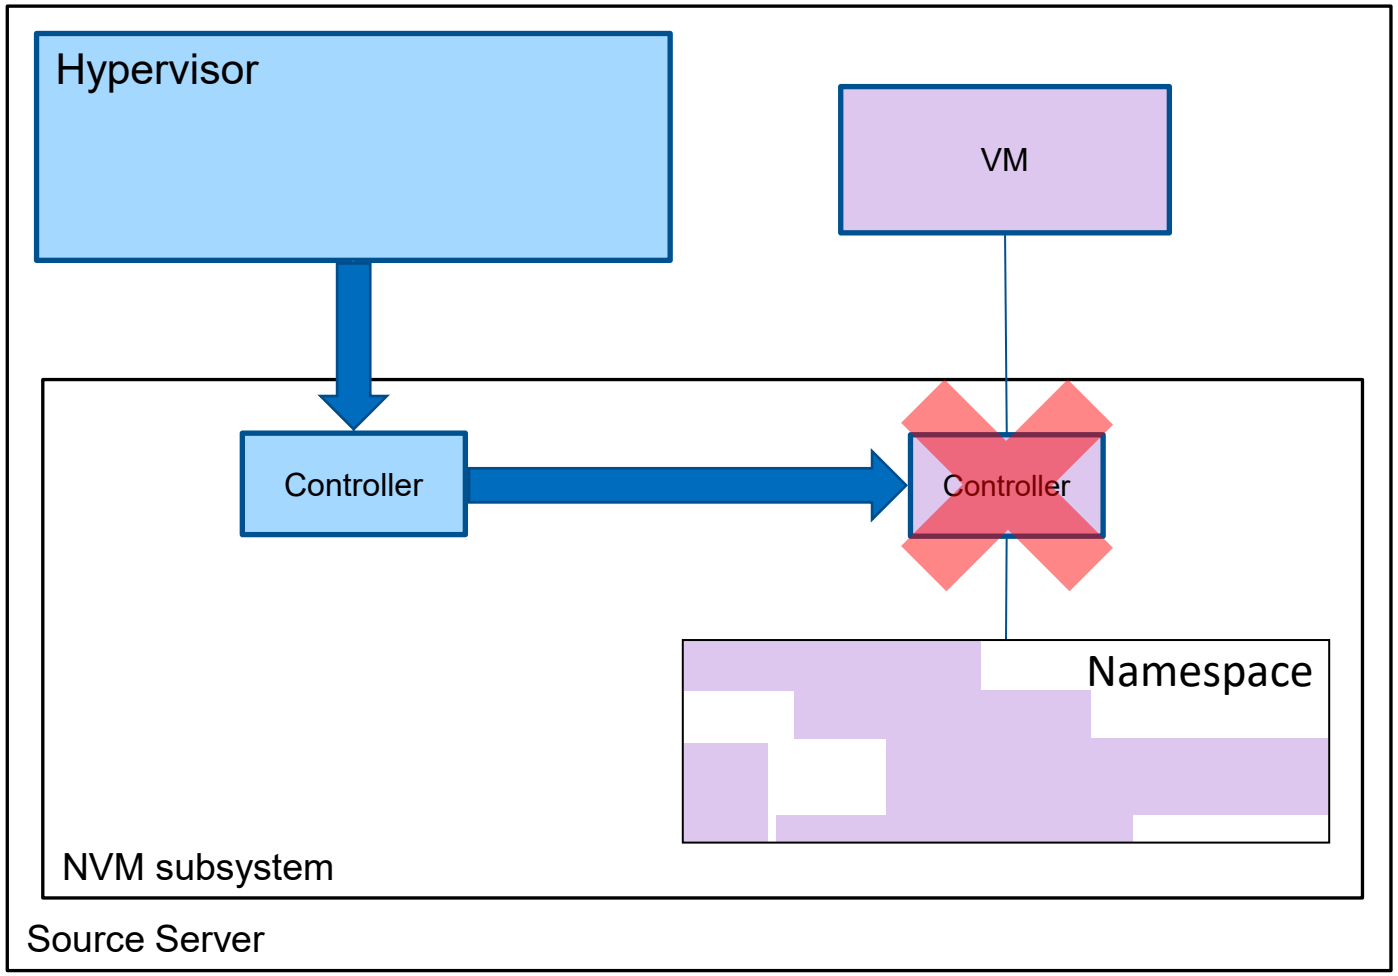

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
  - Read controller state
  - Pause controller

NVMe™ Host Managed Live Migration

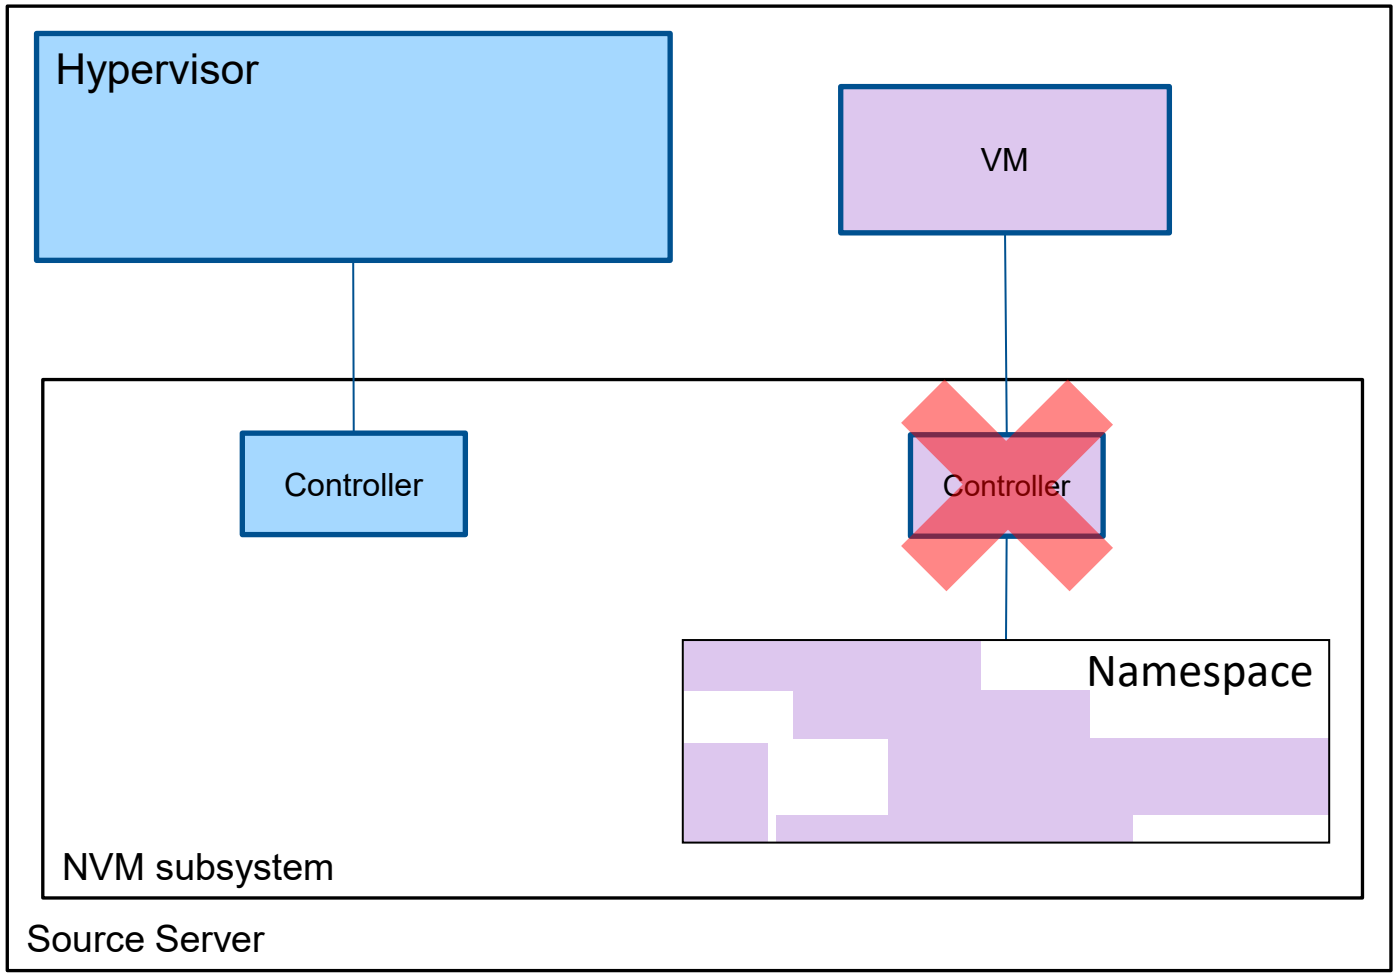

- Source NVM Subsystem
  - Read allocated logical blocks
  - Request logging changes
    - Logical Blocks
    - VM memory
  - Read controller state
  - Pause controller

NVMe™ Host Managed Live Migration

- NVMe Host Managed Live Migration provides services to a Host to:
  - Get the NVMe state from a source controller
  - Put that NVMe state into the target controller and resume operation
- Target NVM Subsystem

NVMe™ Host Managed Live Migration

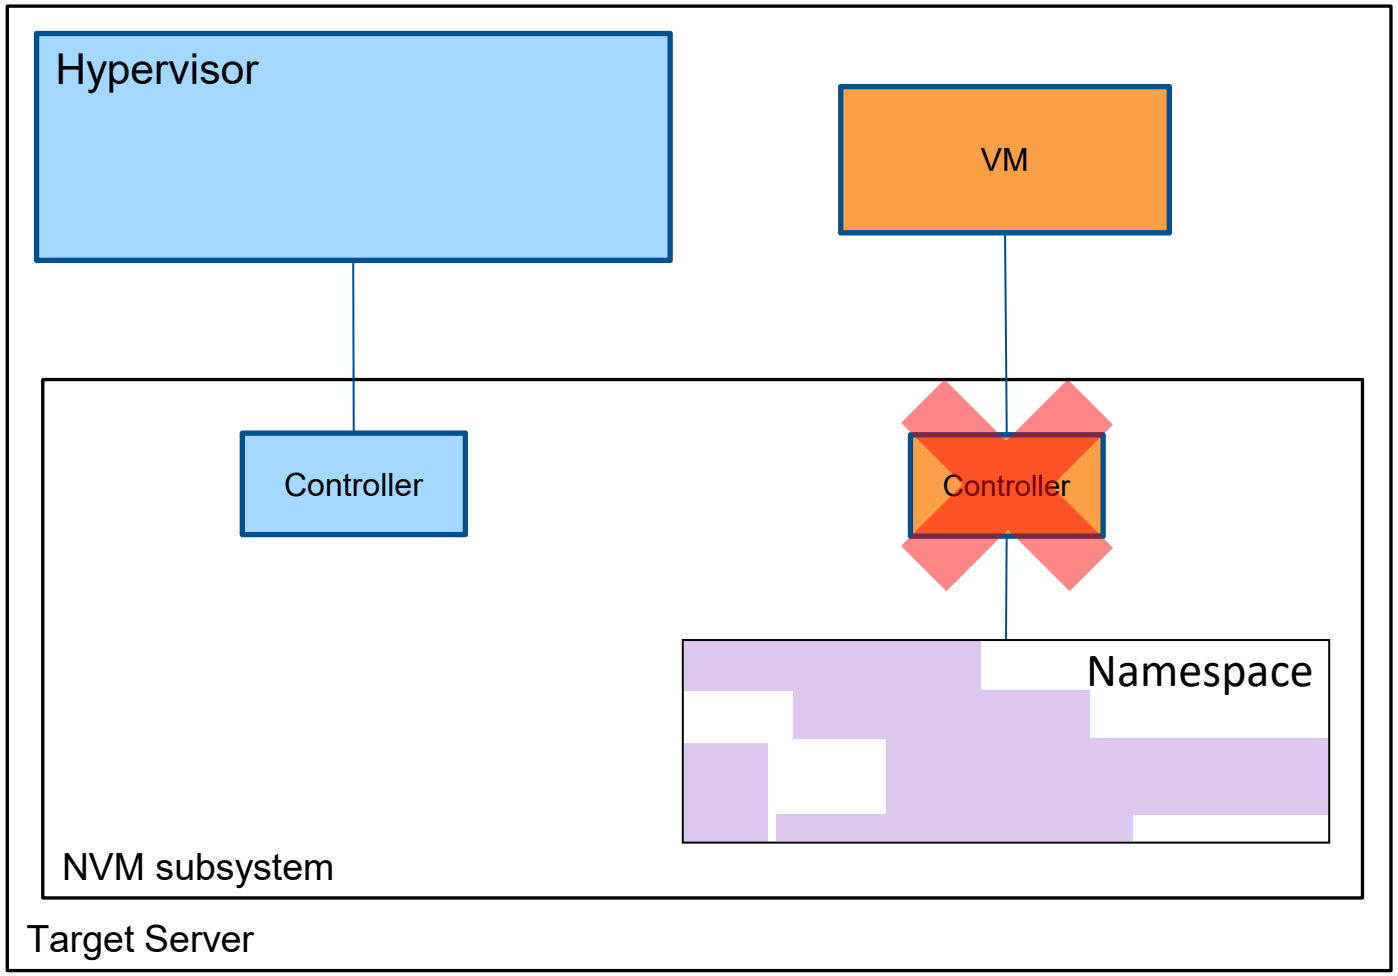

- Target NVM Subsystem
  - Set the Controller State

NVMe™ Host Managed Live Migration

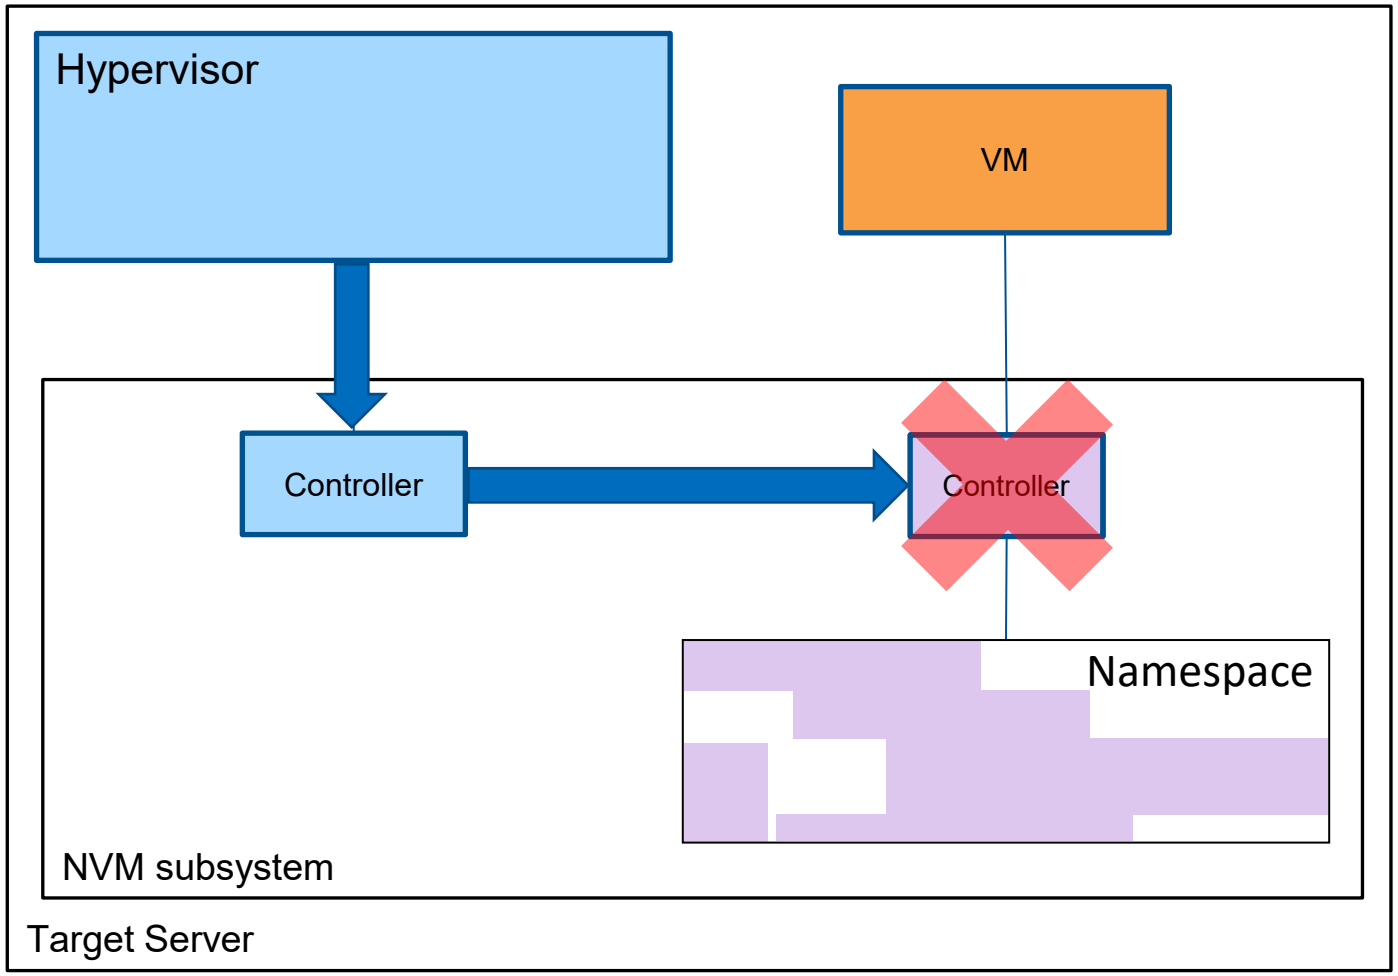

- Target NVM Subsystem
  - Set the Controller State

NVMe™ Host Managed Live Migration

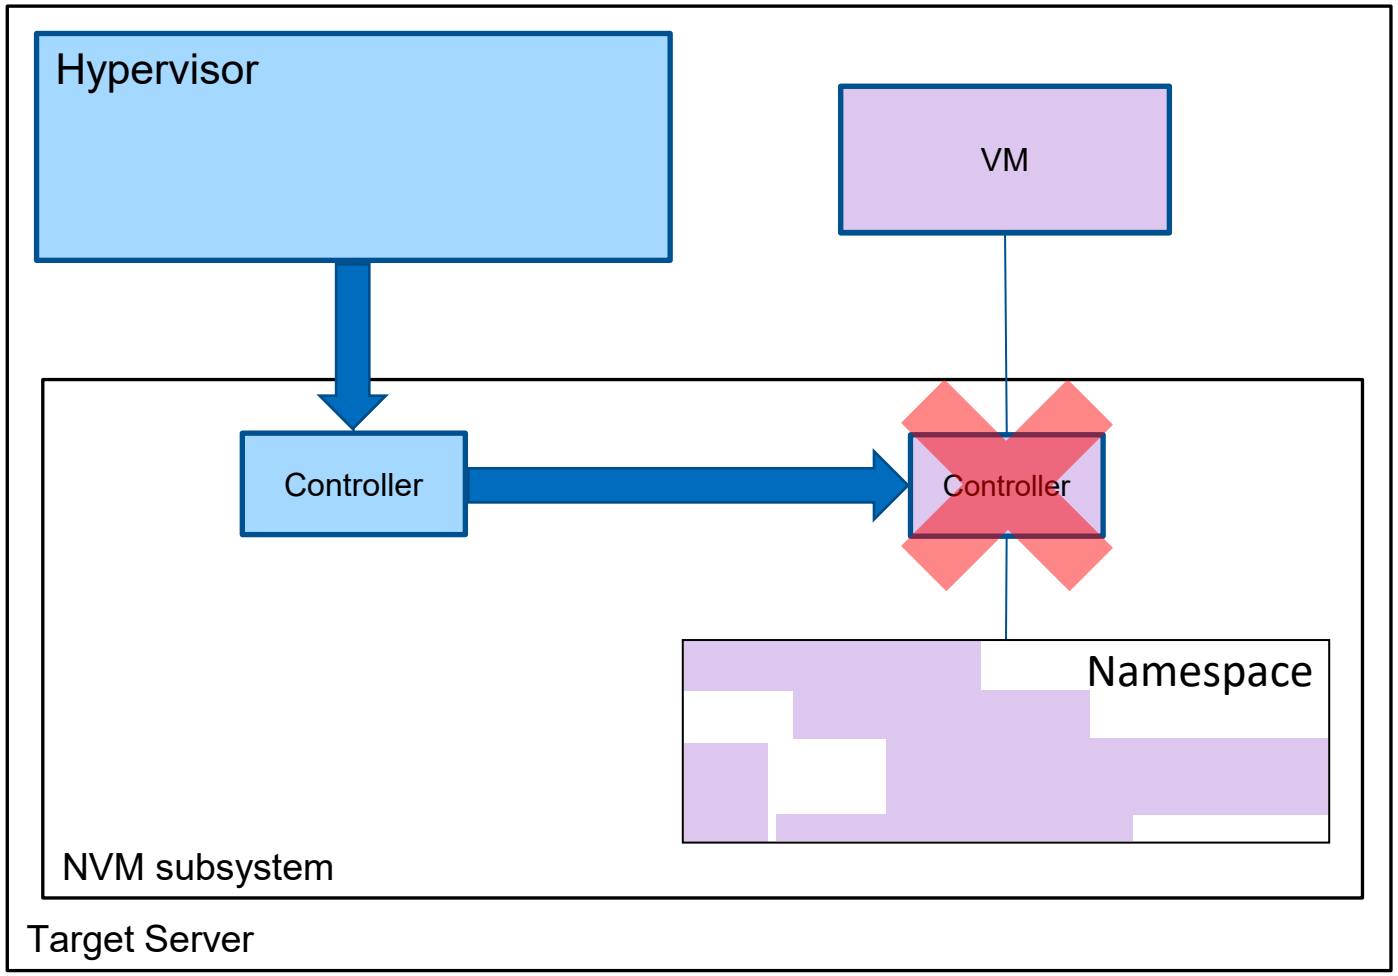

- Target NVM Subsystem
  - Set the Controller State
  - Resume operations on the migrated controller

NVMe™ Host Managed Live Migration

- The reading and writing of namespaces is using existing read and write I/O commands.
- The transfer of namespaces and controller state between servers is outside the scope of the NVMe specifications.

SNIA Developers Conference 2024

- I am planning on submitting a presentation for SNIA SDC 2024 to present the ratified version of the specification detailing the full NVMe™ protocol supporting Host Managed Live Migration.

Please take a moment to rate this session.

Your feedback is important to us.