

SNIA DEVELOPER CONFERENCE



BY Developers FOR Developers

September 16-18, 2024
Santa Clara, CA

Distributed-NVMe

a high performance distributed block storage system

Peng Yu

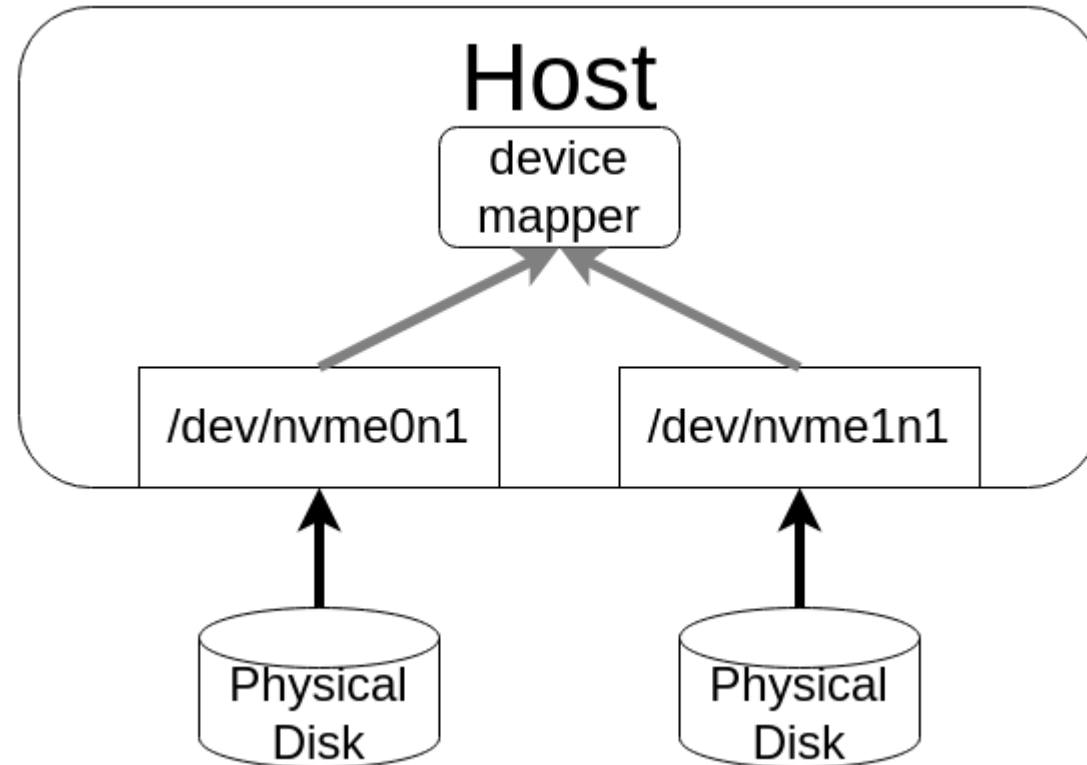
Introduction

- A distributed block storage system.
- Standard NVMe-oF interface.
- High performance, more than 200M IOPS for a single block device
- NVMe multipath for HA.
- Using raid internally for data redundancy.
- A large granularity thin provision, typically 1G – 100G.
- Open source

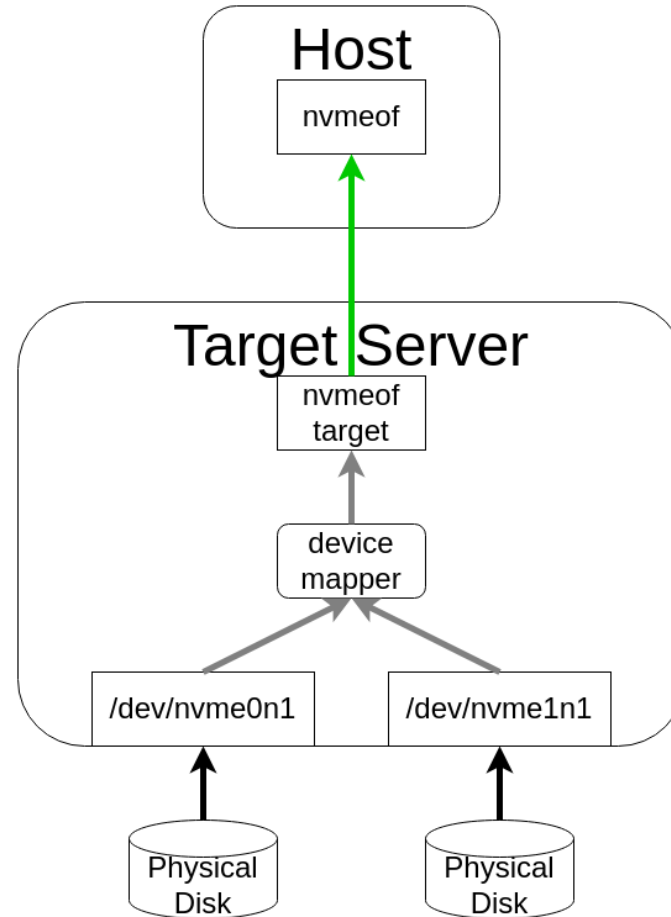


Architecture

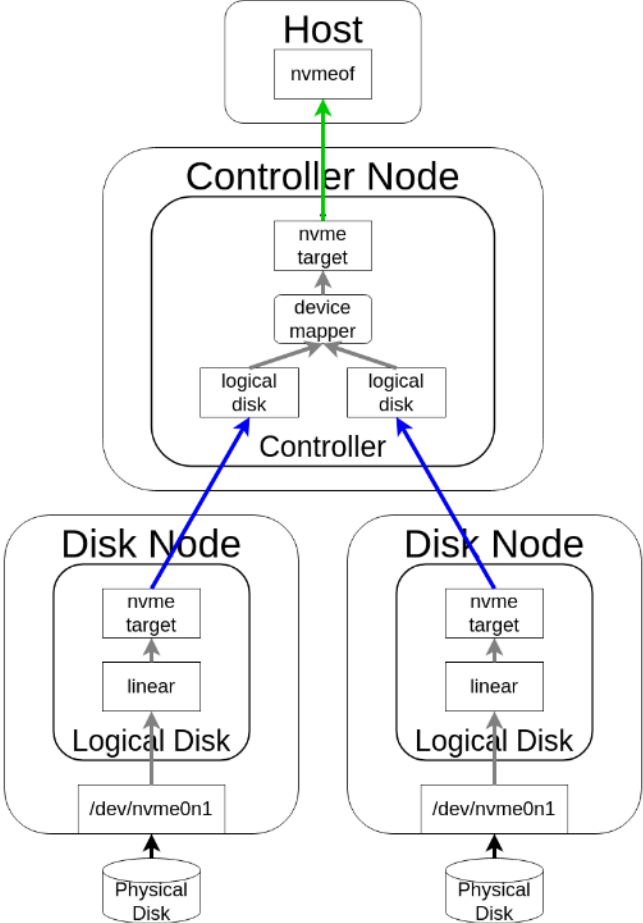
Single host



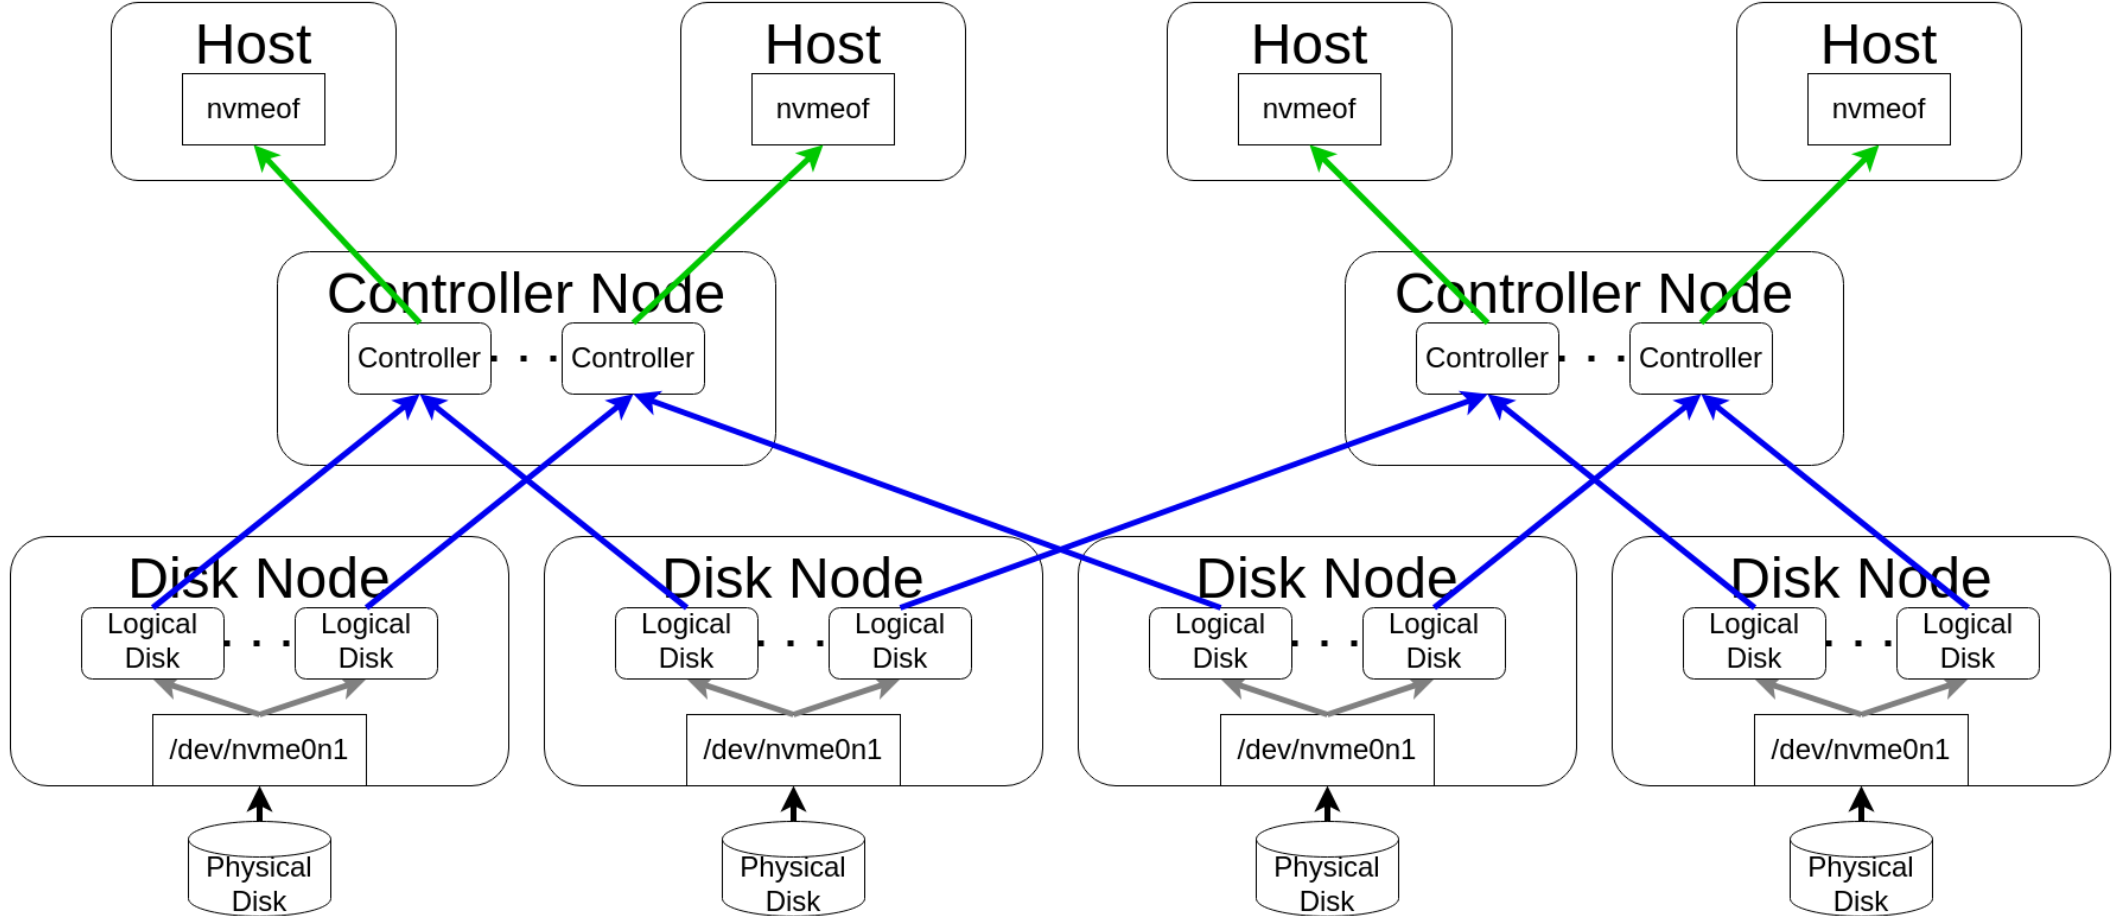
Host and target



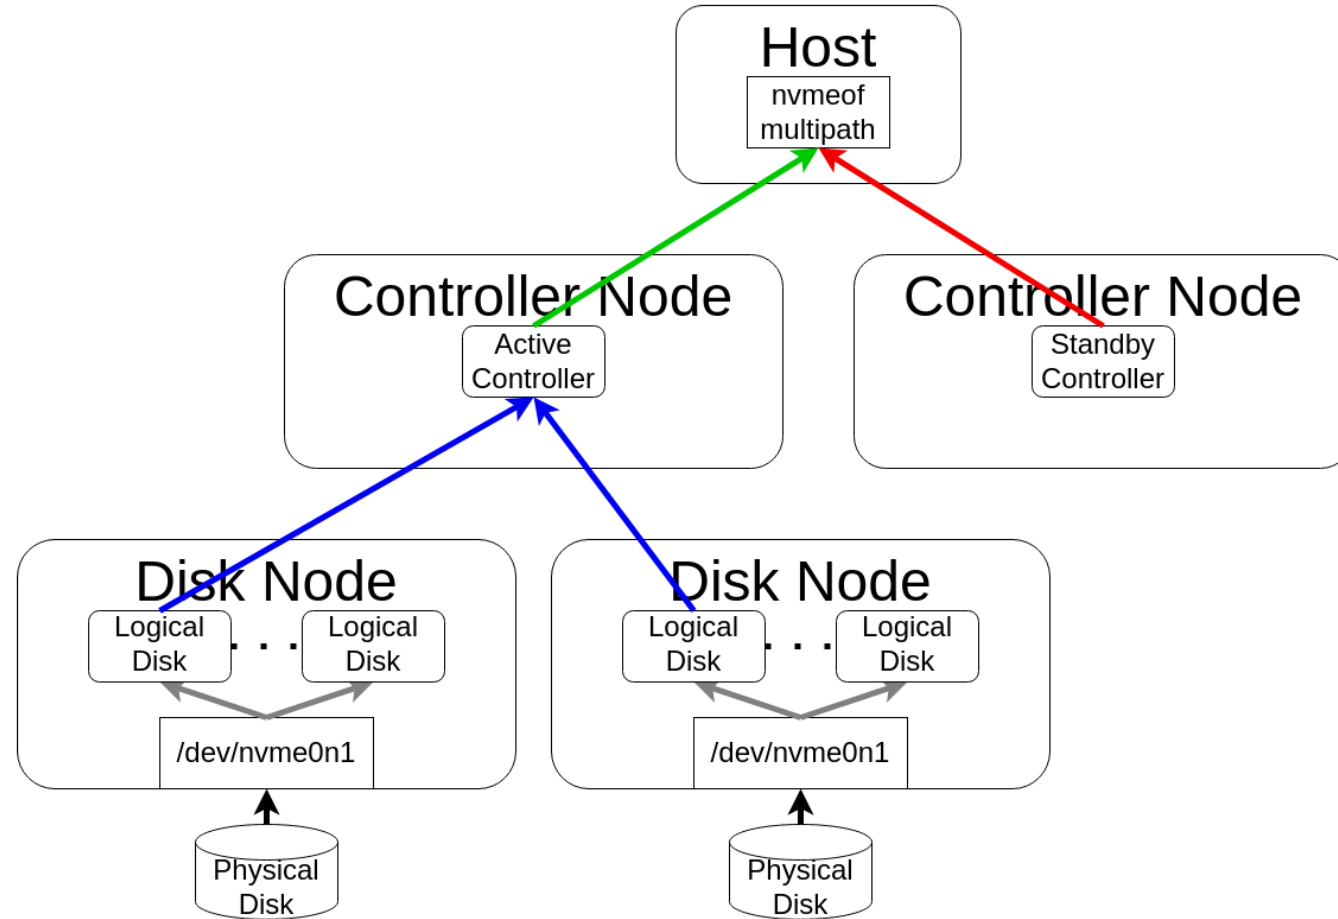
Controller node and disk node




Multiple controller nodes and disk nodes




Active and standby controller




Dispatcher layer



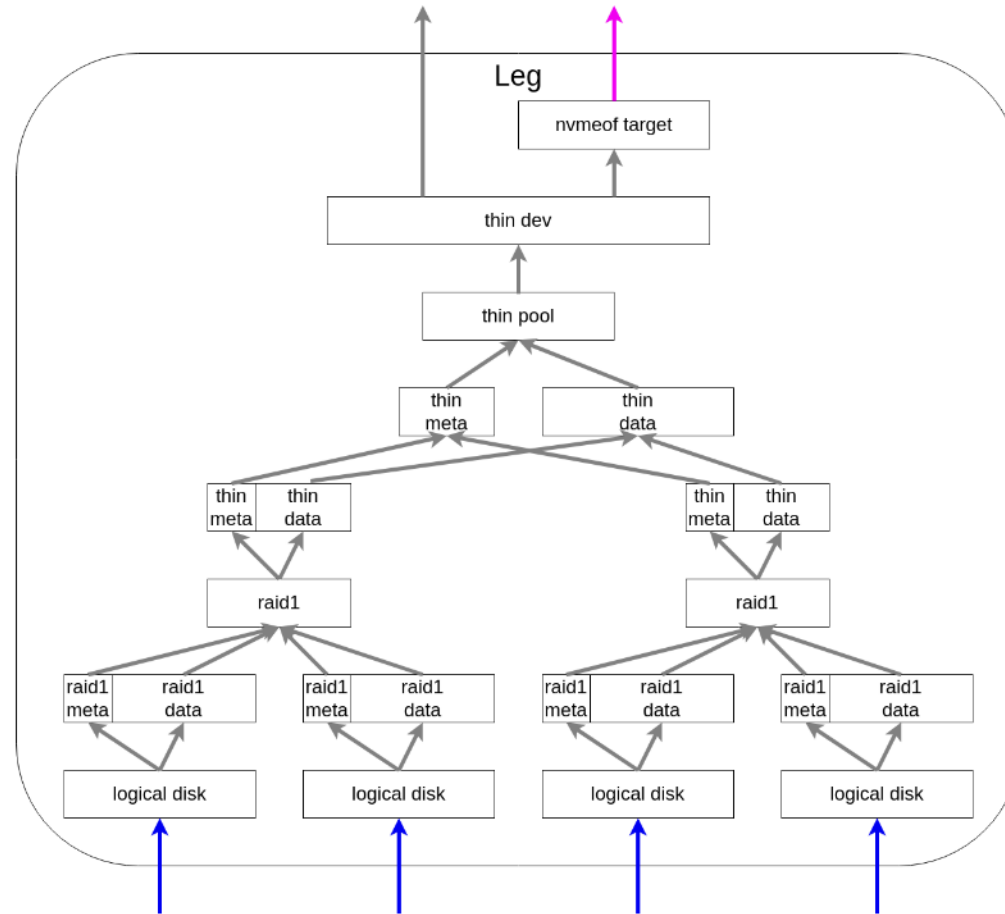
Compact mode



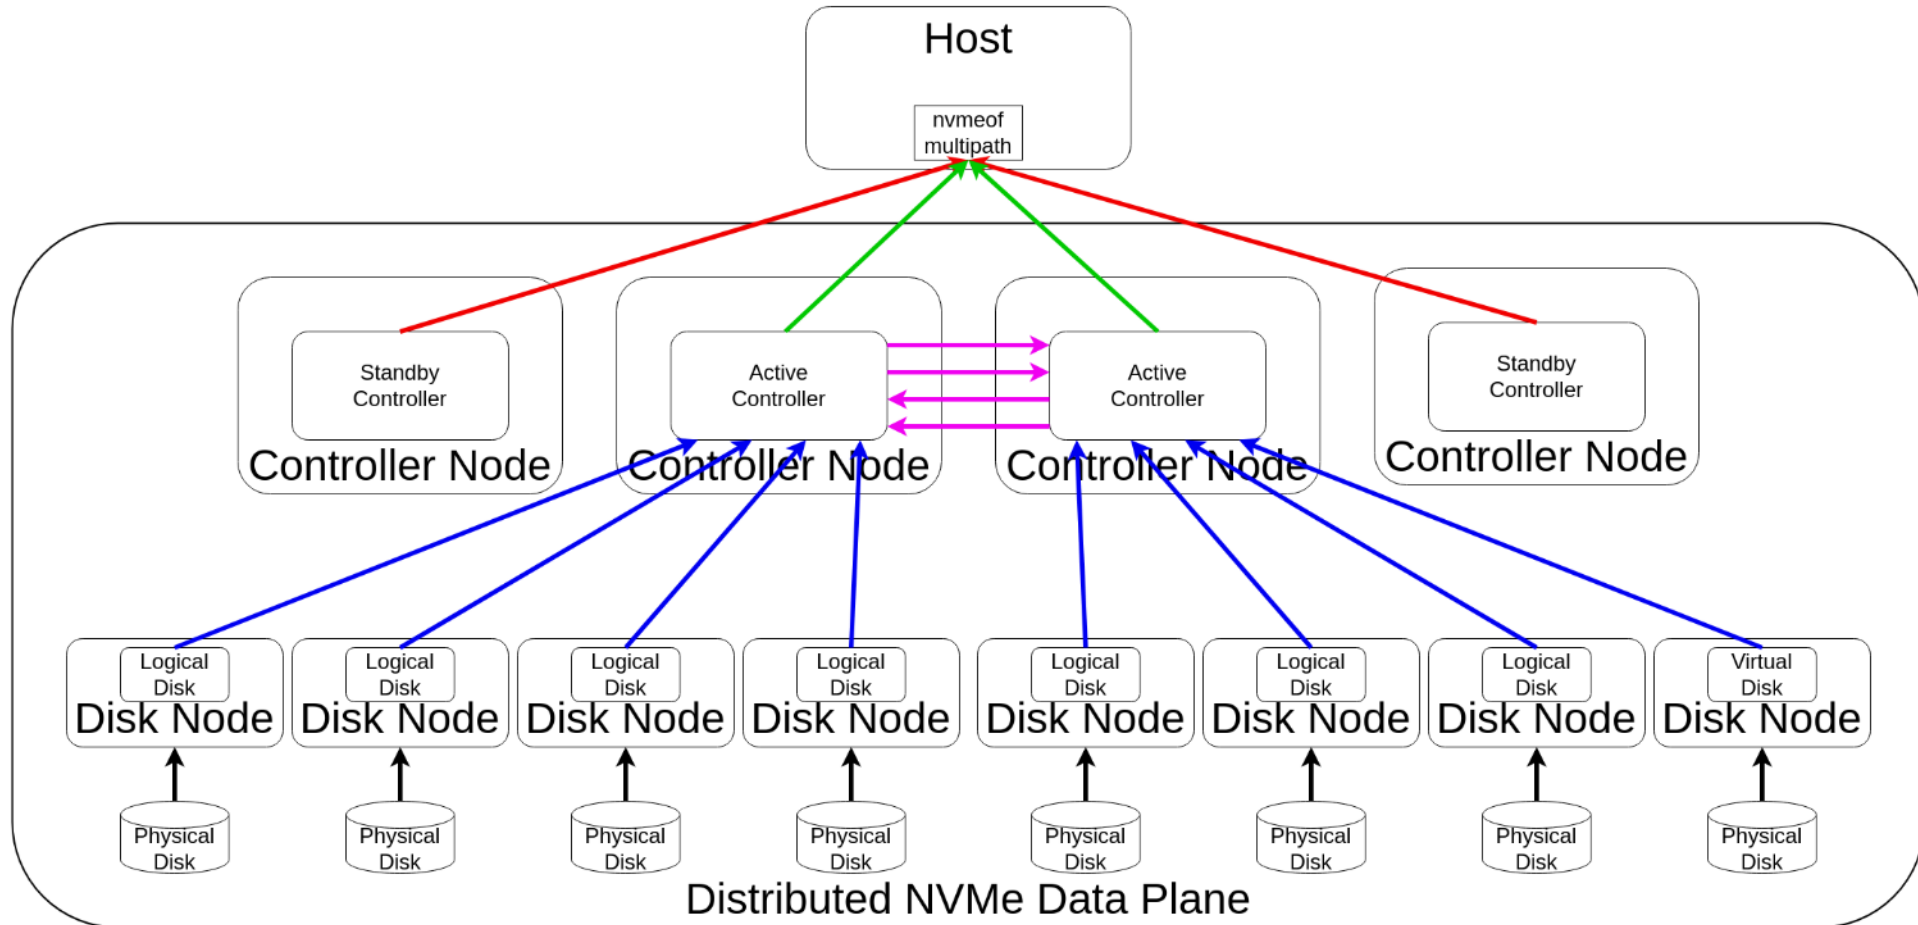
Compact mode with standby controller



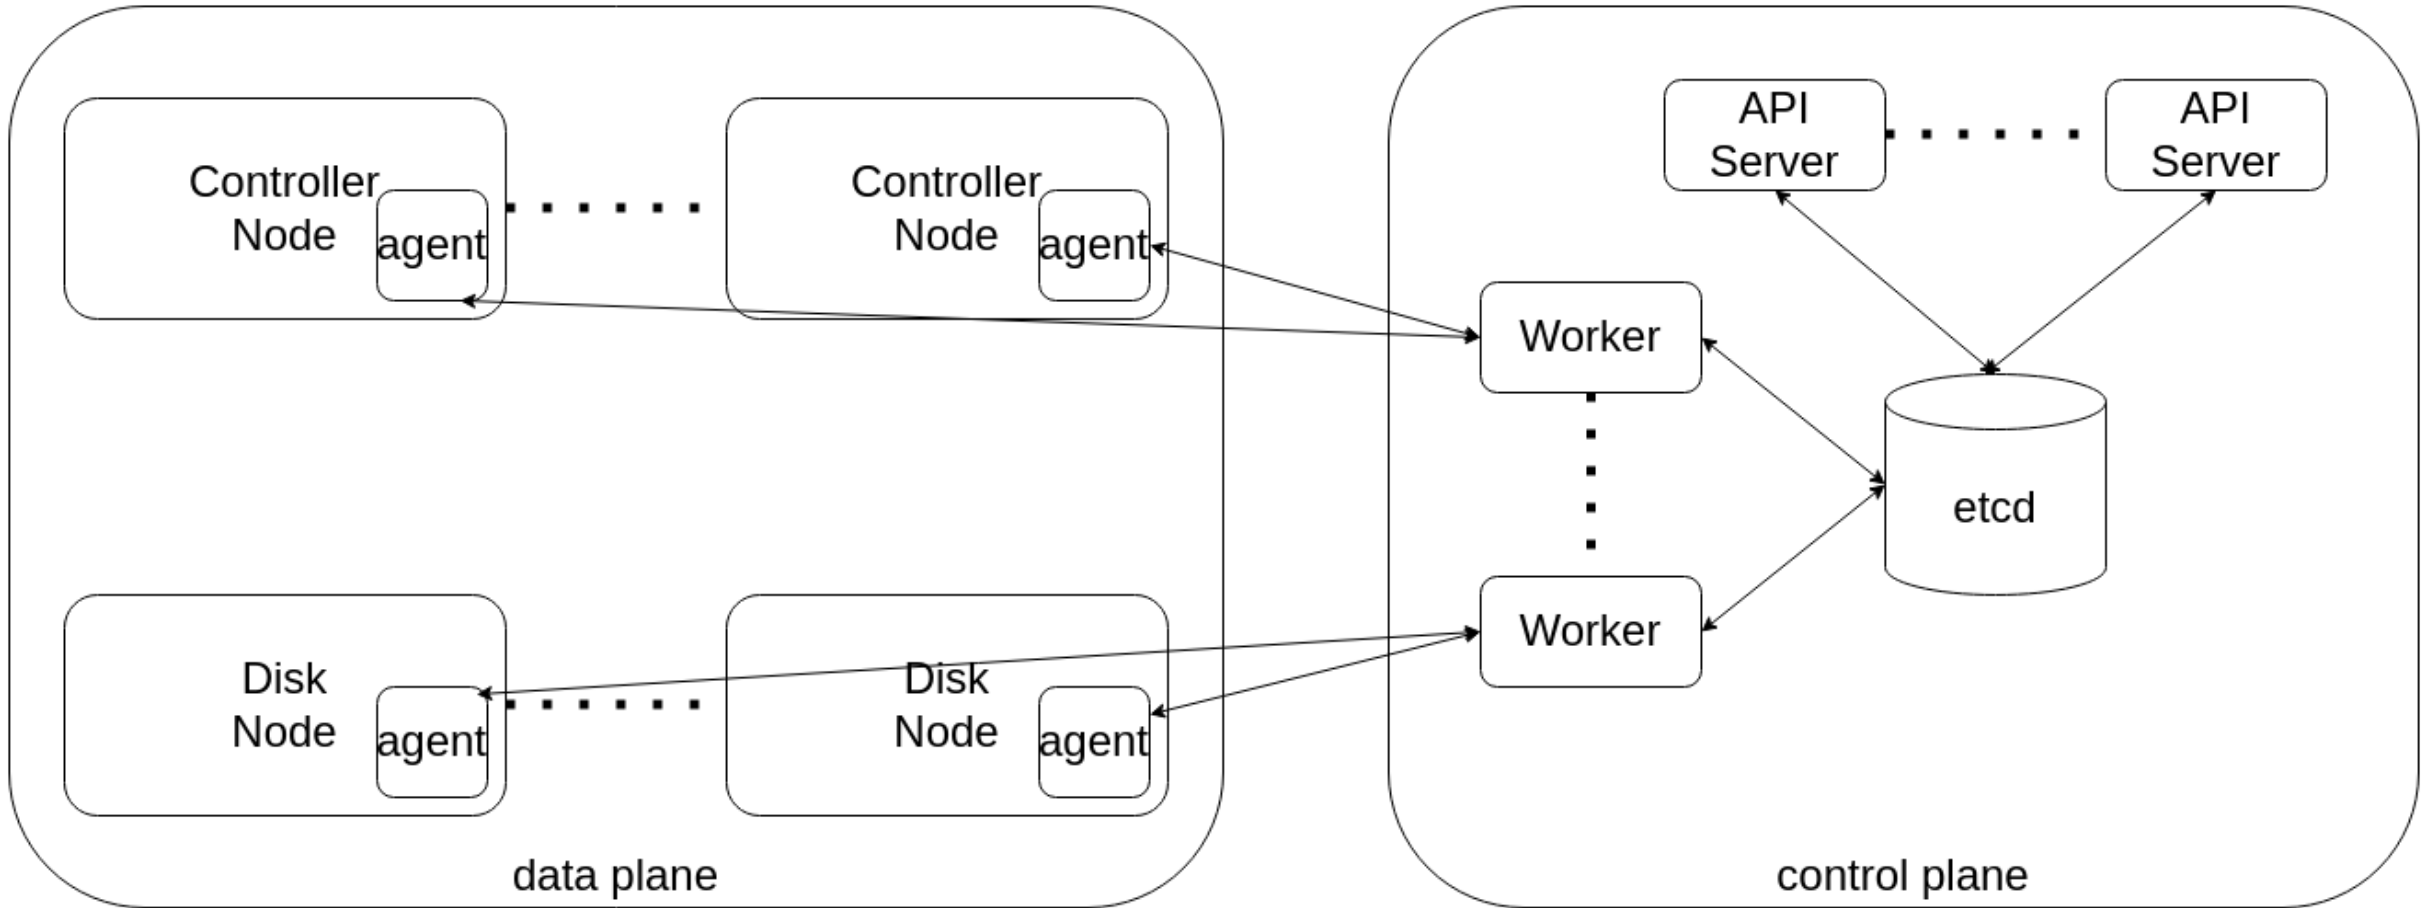
Leg



Host and virtual block device



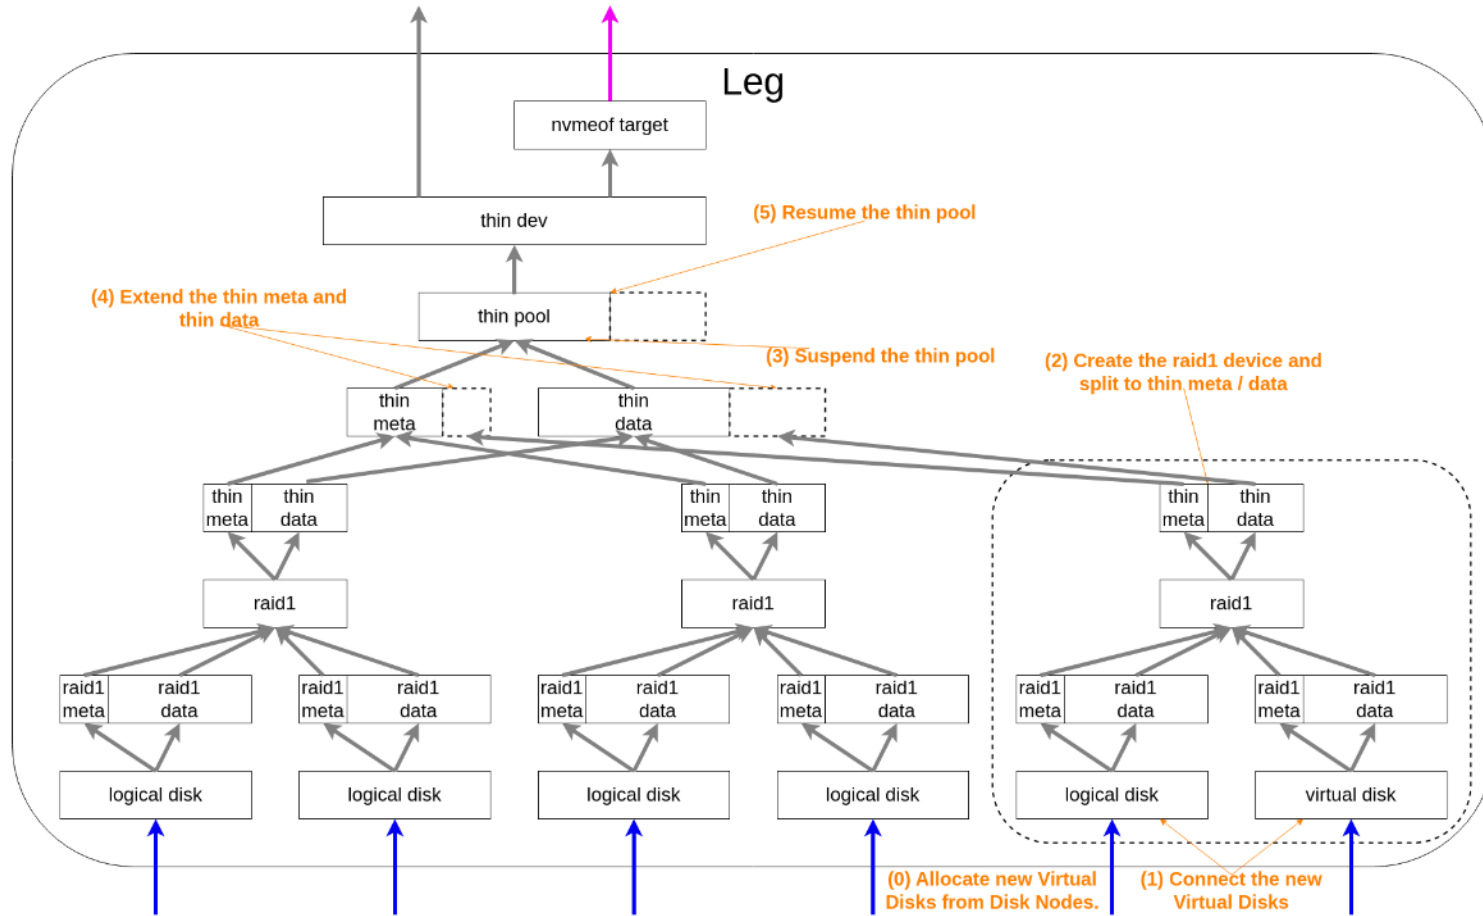
Control plane and data plane



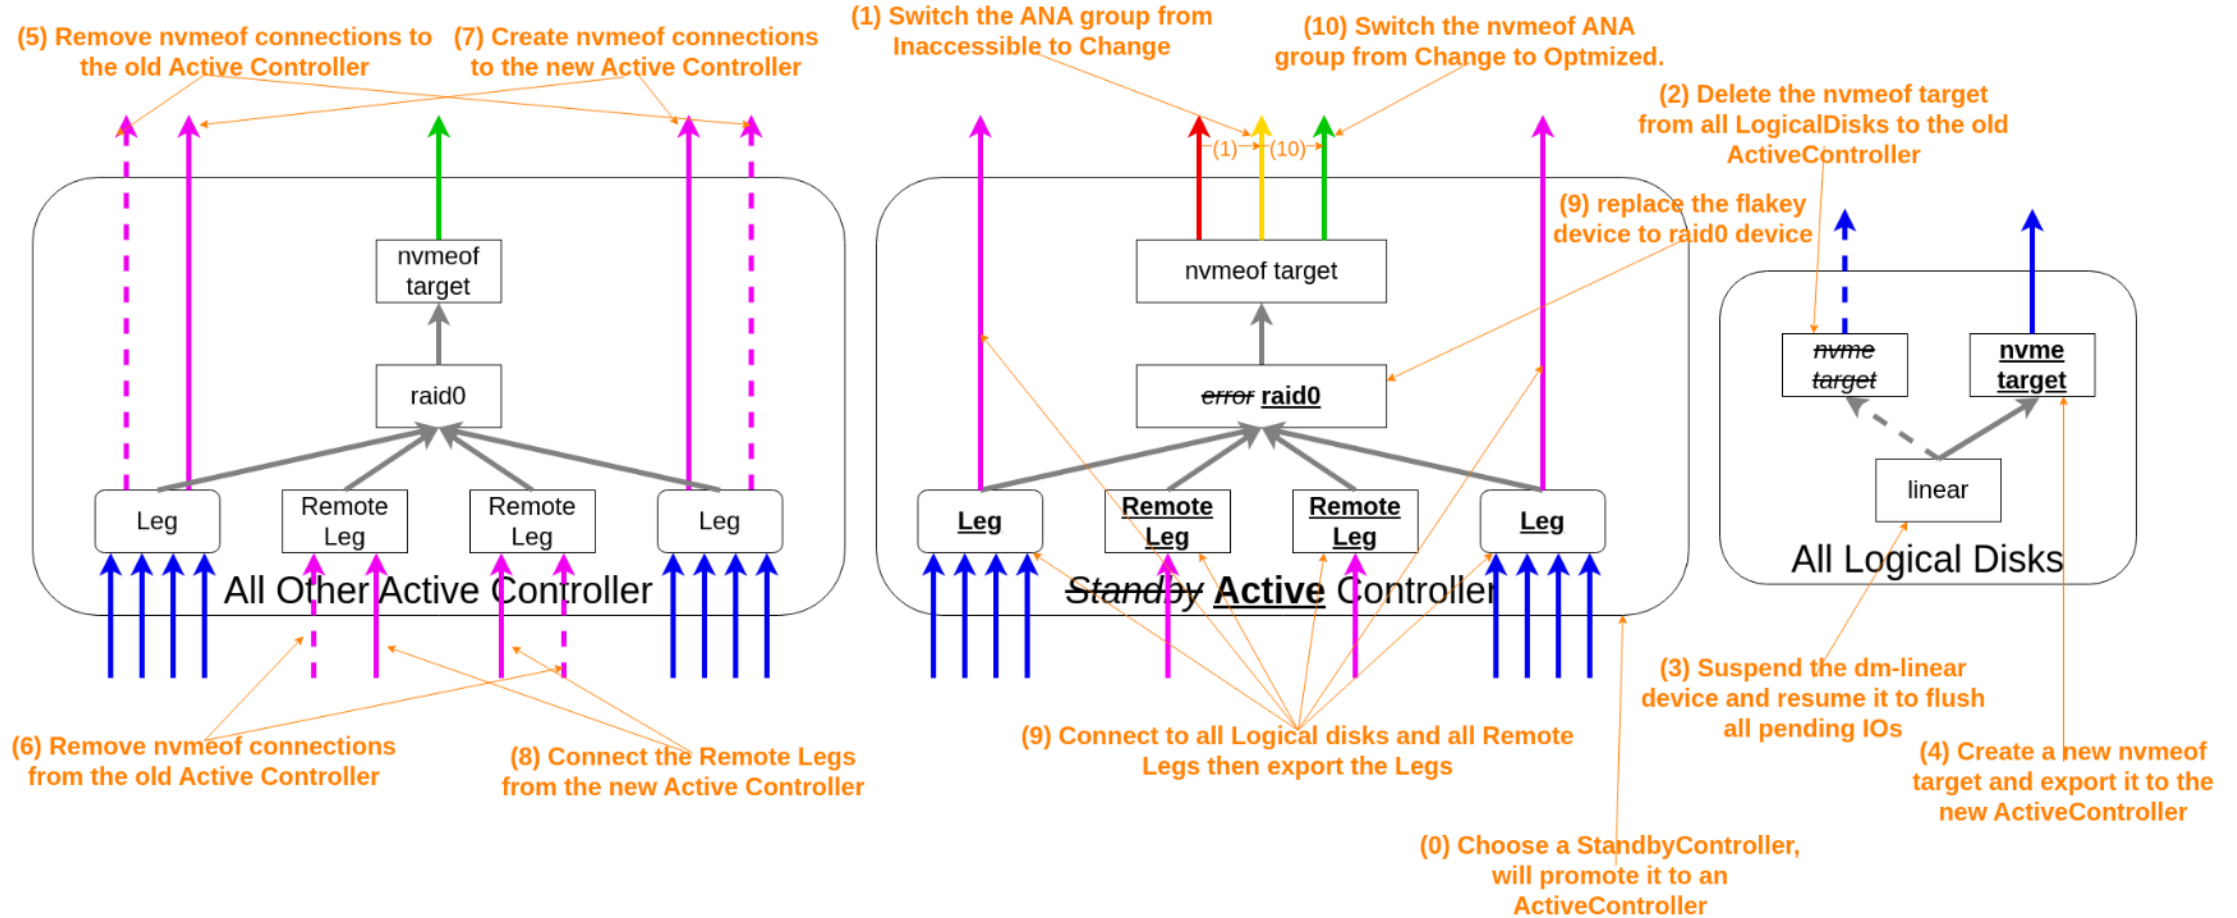


Operations

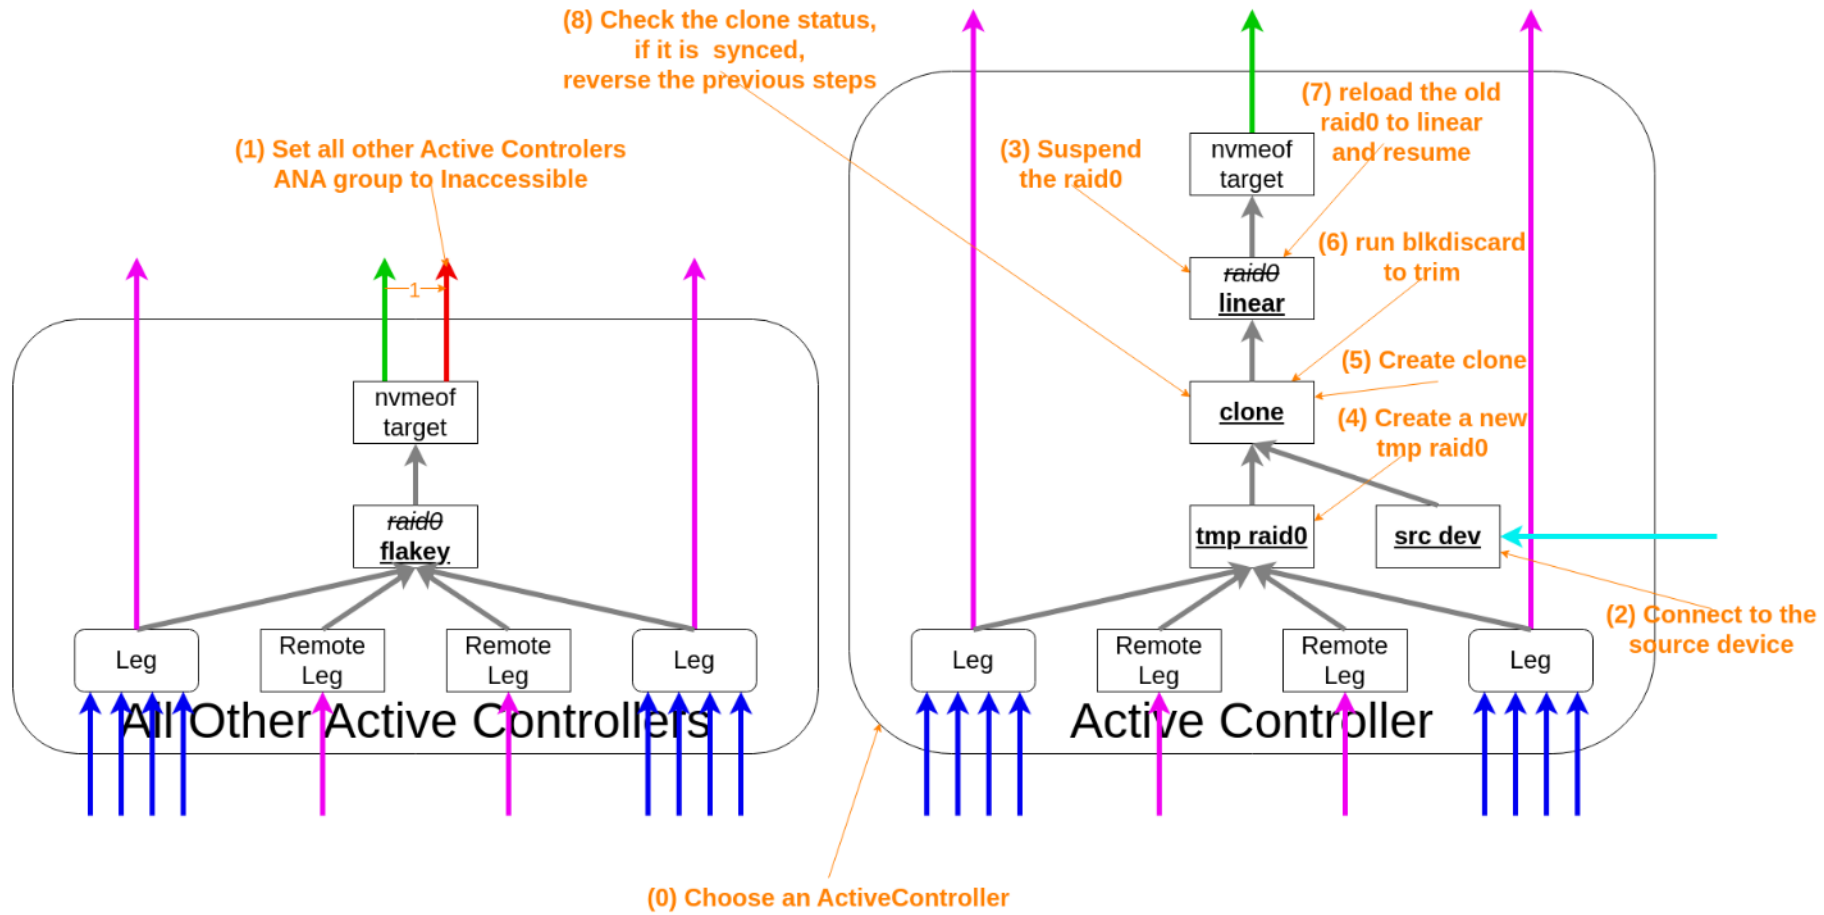
Extend



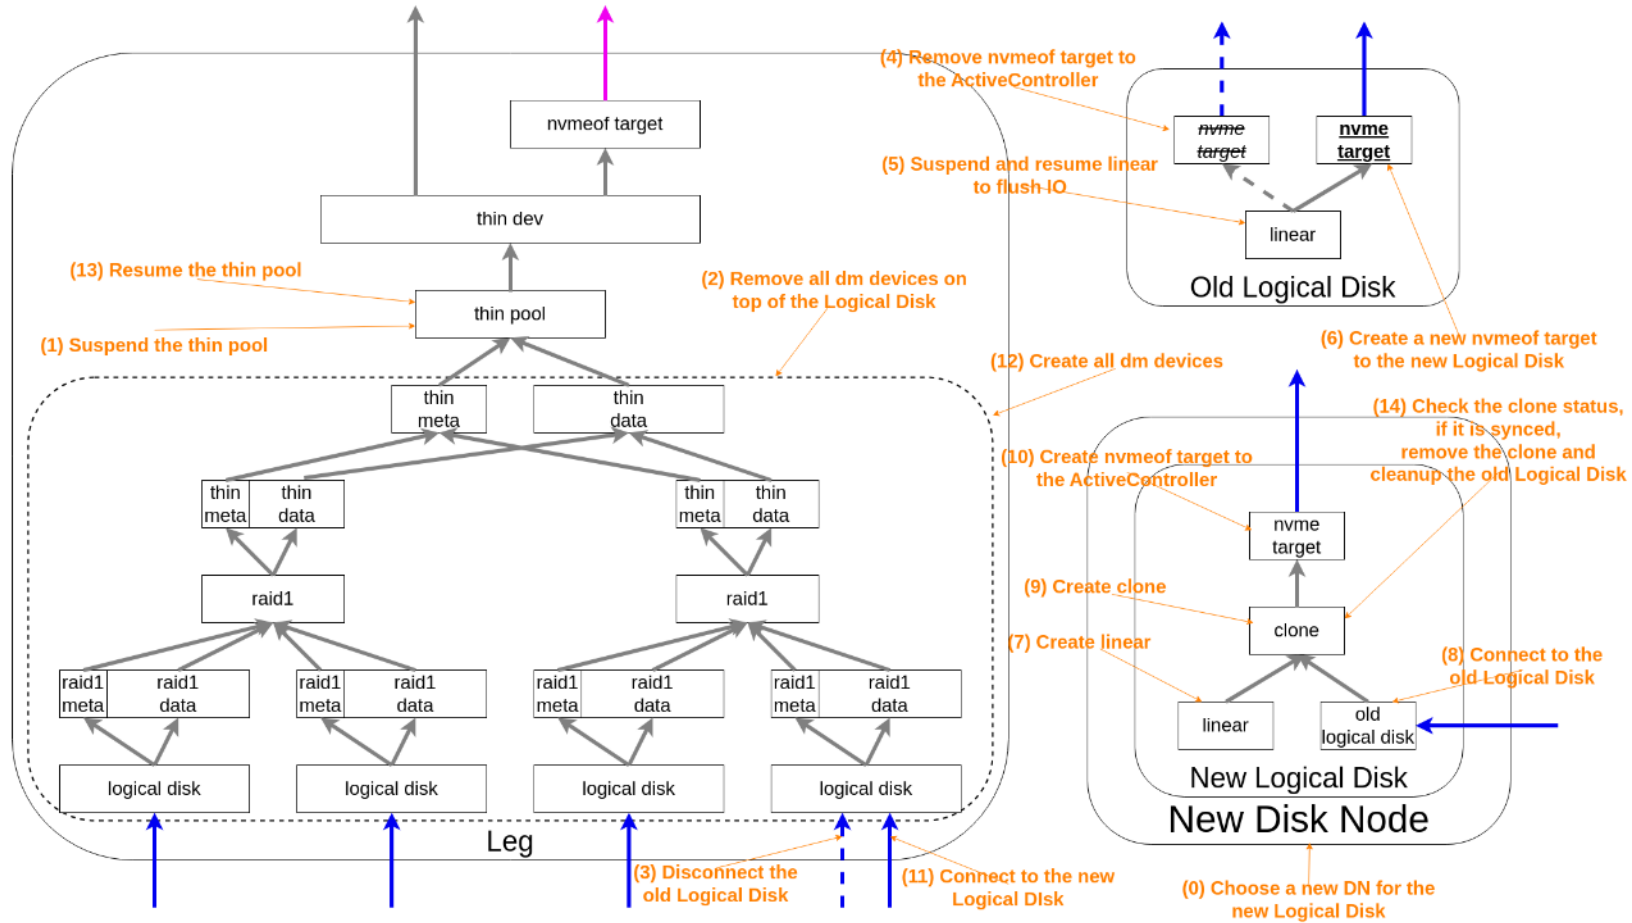
Failover



Clone

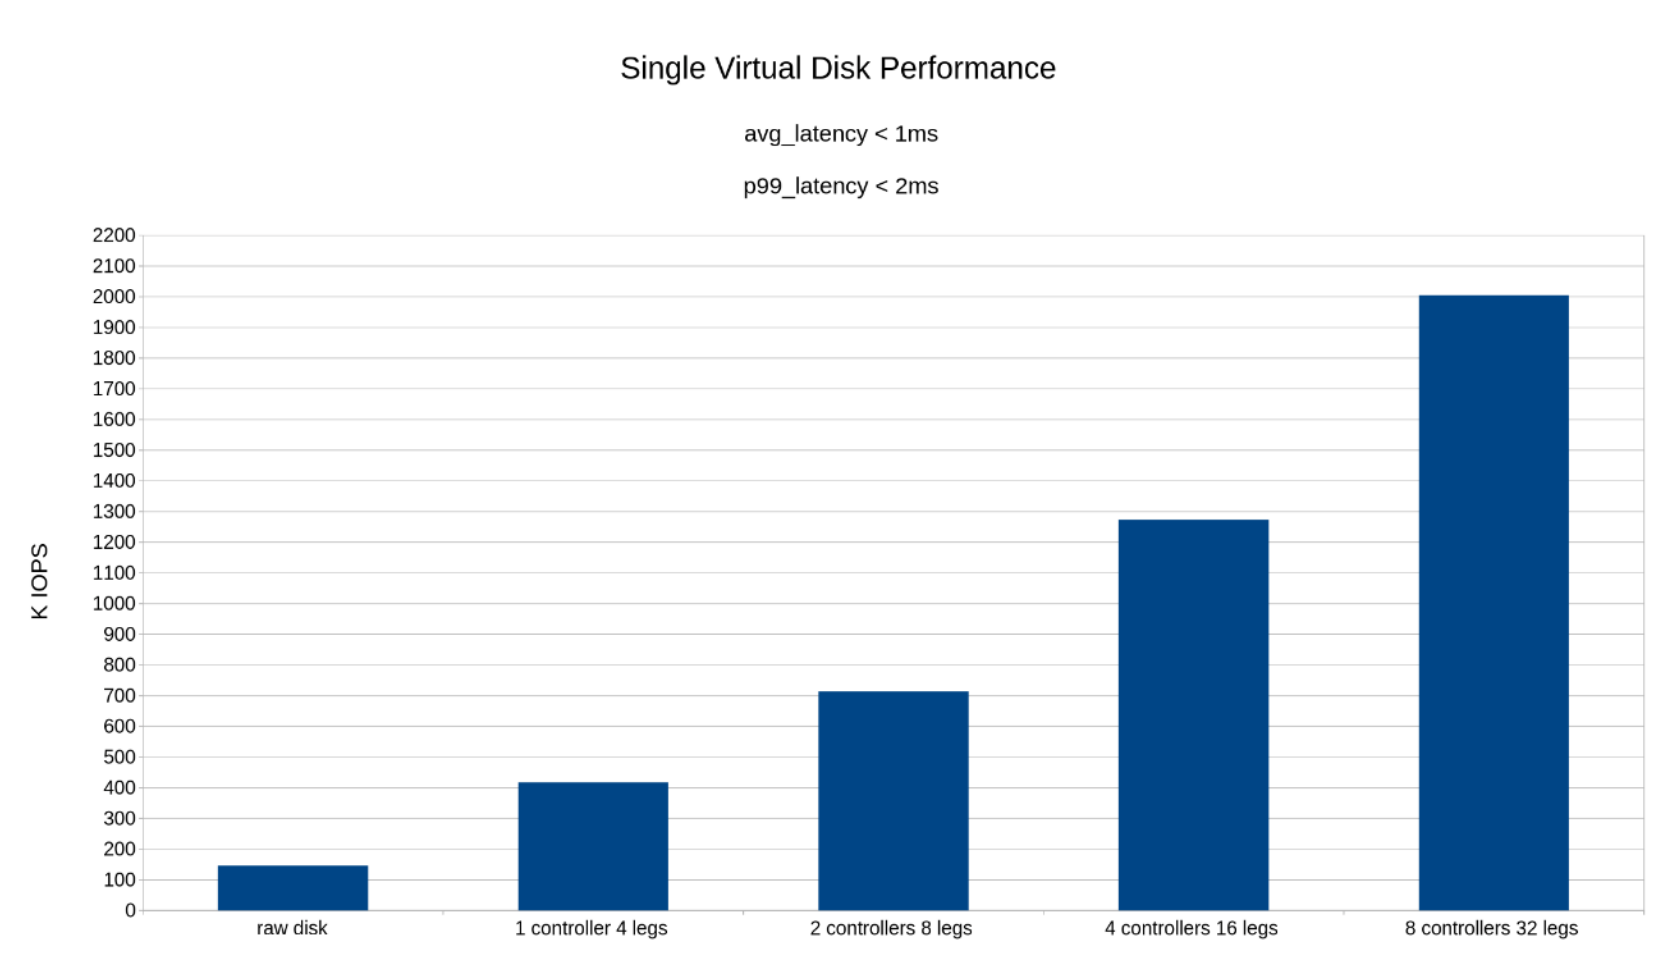


Move data





Performance



Future plan

- Finish the functions in this presentation
- Implement CSI driver for k8s
- Implement CDC (Centralized Discovery Controller)



Please take a moment to rate this session.

Your feedback is important to us.