



SDC 18


September 24-27, 2018
Santa Clara, CA

www.storagedeveloper.org


Recent Developments in The Linux I/O Stack

Martin K. Petersen
Oracle


Linux I/O Development Activity



Subway Map Development Activity




IDE & SATA



IDE & SATA

- ❑ Legacy IDE subsystem
 - Quirky controllers on quirky platforms
 - Recent activity mostly kernel housekeeping
- ❑ Libata for modern ATA devices
 - New controller support, mostly embedded systems
 - Otherwise mostly controller or SSD quirks
 - Plumbing for ZAC

SCSI Initiator & Target




SCSI Initiator & Target

- ❑ SCSI Initiator
 - blk-mq enablement
 - ZBC enhancements
 - UFS
 - NCR5380!
 - Storage array quirks
 - FC-NVMe

SCSI Initiator & Target

- SCSI Target
 - tcmu
 - Bug fixes


Block Layer




Block Layer

- ❑ Zoned device enhancements and multiqueue support
- ❑ Kyber
- ❑ BFQ
- ❑ Planned obsolescence for legacy I/O path


NVM Express



NVM Express + Block Layer




NVM Express + Block Layer + SCSI



NVM Express

- ❑ FC-NVMe
- ❑ Multipathing, ANA
- ❑ Trace points
- ❑ PCIe SGL support
- ❑ Support for files as backing store for target
- ❑ Controller quirks


Persistent Memory



Persistent Memory

- ❑ NV-DIMM vs. ACPI
- ❑ MAP_SYNC
- ❑ Poison handling

Development Activity Trends



What's in the pipeline?

- ❑ NVMe multipathing refinement
- ❑ NVMe over TCP/IP
- ❑ Peer-2-peer DMA
- ❑ Zoned devices not just for disks anymore
- ❑ Multi-actuator disk drives