

NVMe mapping for SCSI Layout

Christoph Hellwig

pNFS scsi layout status

- ❑ Enhancement / replacement of the block layout RFC to take advantage of unique device identification and persistent reservations
- ❑ Explicitly tied into SCSI
- ❑ Released as RFC 8154 in May 2017

Multi-host NVMe devices

- ❑ Since version 1.1 NVMe allows access to the same namespace from multiple hosts
 - Using multiple “controllers”
 - Including a version of SCSI Persistent Reservations
- ❑ Multi host PCIe devices are a rather niche market
- ❑ In June 2016 the NVMe over Fabrics spec has been published
 - NVMe over RDMA, FC, etc

Using the SCSI layout with NVMe

- ❑ NVMe devices can implement all features needed for the SCSI layout
- ❑ The **NVM Express - SCSI Translation Reference** document translates the required features between NVMe and SCSI
- ❑ But unfortunately it's not a IETF-quality spec

Using the SCSI layout with NVMe

- *“Using the Parallel NFS (pNFS) SCSI Layout with NVMe”* (**draft-hellwig-nfsv4-scsi-layout-nvme**)
 - For now just points to the sections in the NVMe STLR document relevant for **draft-ietf-nfsv4-scsi-layout**
 - So far only published as informal
 - Multiple suggestions to make it a standards track document

Questions?