

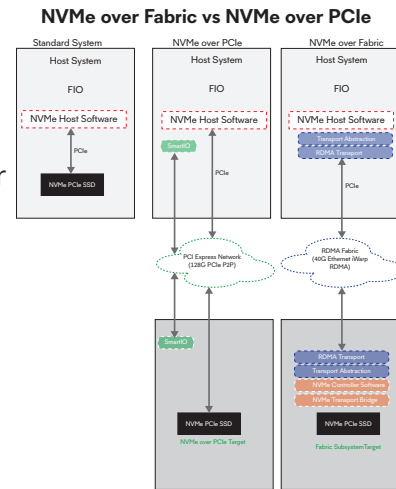
# NVMe over PCIe

## PCIe Device lending

NVM Express (NVMe) over Fabrics defines a common architecture that supports a range of networking hardware (e.g. Infiniband, RoCE, iWARP) for a NVMe block storage protocol over a networking fabric.

NVMe devices have a direct PCIe interface. NVMe over Fabrics defines a software stack to implement transport abstraction layers at both sides of the fabric interfaces. These layers translate native PCIe transactions and disk operations over the fabric.

A PCIe fabric is different. Native PCIe transactions (TLPs) are forwarded automatically over the fabric with no protocol conversion. Standard PCIe NT technology is used to route the PCIe traffic from the host computer to the NVMe device. Device Interrupts are also automatically routed through the PCIe fabric. This image illustrates the difference between using a PCIe Fabric and other Fabrics. PCIe Fabrics eliminate the transport abstraction, thus providing a much lower latency. It also still supports features such as RDMA.



## Benefits

- Share SRIOV NVMe drives between systems
- Fast access to remote drivers
- Use native NVMe driver
- Near native performance with no translation

## Features

No modifications to Linux or device drivers. Works with all PCIe based storage solutions

Send and receive native NVMe commands directly to and from the fabric.

Optimized for low latency. Latency as low as 540 ns data transfers.

Supports direct memory region access for applications.

Scales to 100s of devices or more.

Multiple ports for simultaneous communication.

Self-throttling, guaranteed delivery at hardware level, no dropped frames or packets due to congestion.

Support for hot plugging of cables, full error containment and transparent recovery.