



## CS561 Spring 2022 - Research Project

**Title:** Developing SPDK application for modern NVMe Storage Devices

**Background:** Over the past couple of decades, SSDs has gained tremendous popularity as storage device because of their lower access time. Most off-the-shelf SSDs use the traditional SATA interface, while, various high- end devices use the PCIe interface offering a new improved protocol called Non-Volatile Memory Express (NVMe) [1]. NVMe SSDs offer low latency, almost as fast as DRAM while persisting data. Among these technologies, 3D XPoint is the most mature one which is already on the market via Intel's Optane series. These devices can have access latency as low as 8us. For such extremely fast devices, the software stack (i.e., file system) can be a bottleneck.

Intel's Storage Performance Development Kit (SPDK) [2] allows user to bypass the file system and analyze the device's raw performance. SPDK provides a set of libraries for writing high-performant, scalable, user-mode storage applications. It does so by directly accessing raw memory via a userspace polled-mode driver. To run an SPDK application (i) the device must not have an active file system, (ii) hugepages must be allocated to facilitate its driver (2GB by default), and (iii) the device must be bound to a Virtual Function IO (VFIO) kernel driver rather than the native kernel drivers to allow direct device access to userspace. Analysis of such devices reveal that for these fast devices, the file system can be a bottleneck [3]. We envision to explore the uncharted realm of *filesystem-less* storage application.

**Objective:** The objective of the project is to develop common database related applications (bufferpool, LSM-tree, B+ tree, etc.) using SPDK on NVMe devices. By developing such applications, we can perform a comparison study with existing systems, thus, identifying the impact of the device without the filesystem.

**Responsible mentor:** *Tarikul Islam Papon*

### References

- [1] NVM Express. 2015. What is NVMeTM? <https://nvmexpress.org> (2015).
- [2] SPDK. 2016. Storage Performance Development Kit (SPDK). <https://spdk.io> (2016).
- [3] Tarikul Islam Papon and Manos Athanassoulis. 2021. A Parametric I/O Model for Modern Storage Devices. In Proceedings of the International Workshop on Data Management on New Hardware (DAMON)



**CAS CS 561: Data Systems Architectures**  
**Data-intensive Systems and Computing Lab**  
Department of Computer Science  
College of Arts and Sciences, Boston University  
<http://bu-disc.github.io/CS561/>

