

Nominee: ScaleFlux Computational Storage

Nomination title: **Unleashing Data Driven Applications at Scale**

ScaleFlux is the leader in production deployment of Computational Storage (introduced mid-2017) which provides low-latency, high capacity NAND flash storage (up to 6.4TB per standard PCIe card or U.2 drive) and, in addition, uniquely unlocks computationally intense application level bottlenecks with compute acceleration engines. This subsystem solution provides unparalleled application-level performance (database transactions per second and job run time) and is easy to integrate into standard Linux server environments, similar to a NVMe SSD (Flash solid-state drive).

The architecture of Computational Storage is the first application proven solution that brings compute acceleration to the data as opposed of traditional architectures where all data is moved from Flash storage to a processing subsystem (i.e. x86 host CPU). This is very different than standard NVMe SSD competitive offers from the likes of Samsung, Intel, WD, Toshiba and Micron as those solution only provide flash storage I/O with other any compute acceleration forcing the host-CPU to handle all intense processing functions.

By uniquely and simply unlocking both storage I/O and compute bottlenecks cohesively, mainstream data driven database, big data and data warehousing applications like Aerospike, Hadoop, HBase, MySQL, PostgreSQL, Spark and many others get the most responsive performance, storage capacity utilization, and infrastructure scaling efficiency with Computational Storage over standard NVMe SSDs. For example, full benchmark information of validated applications are located at <http://www.scaleflux.com/applicationvalue.html> and include:

- 2.7x better transaction rate with Aerospike
- 3.5x better transaction rate with PostgreSQL
- 3.6x better flash capacity utilization with MySQL
- 2.3x faster queries with Vitesse Data Deepgreen Db
- 2.2x better KV-Store write throughput with RocksDB
- 2.4x better transaction rate with Apache HBase
- 1.7x more jobs completed (vs. HDD) with Hadoop
- 1.7x more jobs completed with Apache Spark

Another major innovation from ScaleFlux is that fact that its Computational Storage leverages programmable hardware in the form of FPGAs (Field Programmable Gate Arrays). The end benefit is that new computational functions can be dropped into the system after initial product deployment. This allows customers to evolve their infrastructure platforms as new applications



and computational challenges emerge over time effectively providing the ability to treat their hardware like software. Data centers are now enabled with new levels of agility never seen before in a flash storage solution, which is quickly becoming one of the largest IT spends.

ScaleFlux has been deployed at several end users who have reported a dramatic improvement in the time-to-insight when analysing the extremely valuable transactional data. Examples include credit card (financial) companies, online travel ticketing, financial risk management analysis, host cloud services and more. ScaleFlux Computational Storage not only saves infrastructure cost by deploying less servers and storage to meet a given workload need because each server with ScaleFlux gets more jobs completed or more transactions processed, it also enables customers to generate new revenue streams based on faster data insight. For example, our credit card customer can offer faster group targeted marketing after users purchase the same type of products.

ScaleFlux will continue to innovate from both the Flash storage I/O perspective by leveraging state-of-the-art 3D NAND Flash that continues to get more affordable over time. In addition, ScaleFlux will offer and add more compute engines to its product portfolio that can be used and integrated into applications for the ultimate in server and storage infrastructure performance, efficiency, and agility. The solution will remain as easy to install as a standard NVMe, yet provide much more value – all in the same economic footprint.

Why nominee should win

- Simple, evolutionary integration (easy to use and install) of high capacity flash storage (like an SSD) with revolutionary impact to true application performance
- Delivers unprecedented performance in terms of data-driven application transactions per second and shorter run time of jobs
- Enables server and storage infrastructure to scale performance across servers and storage efficiently and cost-effectively by balancing the system with additional compute resources while Terabytes of Flash storage are added to the system
- Provides unprecedented agility in the data center by enabling hardware to be re-programmed in-system (like software) to evolve infrastructure as new applications are developed and deployed