NVMe ≠ NVMe-oF

TURN RISK INTO OPPORTUNITY

LEGACY AFAs		HYPERPARALLEL FLASH ARRAYS
CONSTRAINED PERFORMANCE Legacy AFAs use slow, hard -drive era protocols like SAS and SATA. Even NVMe AFAs use an old architecture and just add NVMe drives, throttling performance.	<<<< 10X >>>>	**FASTER PERFORMANCE** A modern hyperparallel system, designed for flash, combines NVMe and NVMe-oF to increase performance for both bandwidth and IOPS up to 10x.
HIGHER LATENCY Legacy AFAs are built around a fundamentally slow, serialized architecture. Parallelism is low and latency is high, which can cause you to miss SLAs.	<<<< 100X >>>>	**LOWER LATENCY** Leveraging NVMe-oF, a hyperparallel system can take advantage of the parallelism to deliver latencies that are 100x better than legacy designs.
VENDOR LOCK-IN Proprietary drivers limit your choice and can be a challenge to deploy and maintain, locking you in to your vendor's update cycle.	<<<< >>>>	**OPEN CHOICE** NVMe and NVMe-oF are industry standards with native driver support in all major OSs. No proprietary vendor drivers needed.
SLOWER AND EXPENSIVE Combining high latency and high costs, outdated AFAs give low performance per dollar, generating budgetary risks while leaving application performance on the table.	<<<< 25X >>>>	**FASTER AND AFFORDABLE** Replacing legacy AFAs with modern NVMe-oF empowers HPC with a low risk, scalable architecture that can deliver up to 25 times better IOPS per dollar.
WASTES SPACE Data centers often have to deploy dozens of AFAs to meet performance and capacity requirements increasing the risk of outgrowing your space.	<<<< 1/10TH >>>>	**SAVES SPACE** Hyperparallel NVMe-oF systems may use only 1/10th the rack space of legacy AFAs, while delivering the same or better performance, replacing racks of storage in just a few RUs.