

# NVME ≠ NVME-OF

## Turn Risk Into Opportunity

Racing to reap the benefits of your HPC investment? Whether the project is AI/ML, eCommerce, Scientific Research, or Financial Risk Analysis, transforming your business can make or break your career. Making better decisions faster is the prime directive. As a CIO/CTO, the decisions on infrastructure spending and the results can make or break your career.

Organizational and team leaders for infrastructure (server, storage, network) have a new bottleneck – NVMe Storage.

## What?

That's the fastest and best stuff on the market.

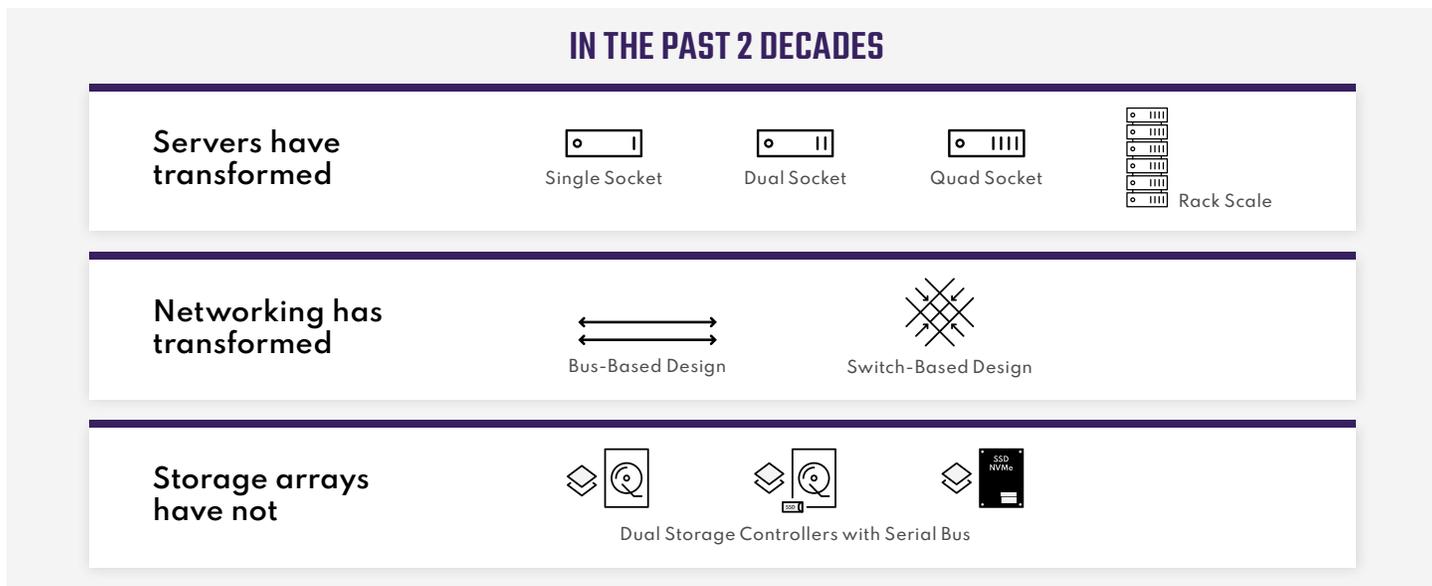
## Why?

Aging storage array architectures.

In the last decade, CPU and GPU laws boosted server performance tenfold, 100Gbps server networks are widely available.

Storage systems and storage networking have lagged. Storage arrays are still using internal serial bus networks? After 20 years. Come on, man!

Your dad's storage arrays were designed to support hard disk drives or SSDs using SATA or SAS protocols.



The SATA/SAS interface is an order of magnitude slower than that of an NVMe SSD.

	SATA/SAS SSD	NVMe SSD	Difference
Read BW (MB/sec)	500	3,300	6.6X
4K Read IOPS	64K	800K	12.5X
Write BW (MB/sec)	475	2,100	4.4X
4K Read IOPS	15K	200K	13.3X
Latency	Milliseconds	Microseconds	100X

**Highly parallel, low-latency, industry standards inherent in today's NVMe and NVMe-oF are reshaping storage systems design.**

Legacy AFA or Appliances with NVMe	NVMe-oF Storage Arrays
Bandwidth	7X Higher
Read/Write IOPS	10X Greater
Latency	4X Lower

In a few short years, NVMe has gone from being the new kid on the block to the dominant SSD interface. But adding NVMe SSDs to legacy storage systems is like putting new chrome wheels to an old car. It looks better but performs the same. Legacy storage systems were designed to support a hard disk drive using SATA or SAS protocols, not NVMe.

You want modern technology, but some storage vendors have left you hanging. Big Storage has a roadmap, an evergreen clause, and great golf outings, but incrementalism will not propel your organization.

**What about alternatives?**

When your data grows, application performance needs to be increased just to maintain the status quo. Adding more compute and storage resources will solve the problem, but it is costly and minimally better.

SATA or SAS interfaces for flash don't access the data fast enough, so you look to high-speed, high NVMe SSDs.

**What happens?**

Adding NVMe SSDs inside your servers to boost application performance comes with its own set of issues. You must tune the application and invest in more SSDs and servers for things

like RAID, backup, and other copy operations. Having a single shared namespace is almost impossible. And you must justify these expenses to the broader organization.

Software-defined storage (or SDS) does create a shared namespace. But it runs on every server node, patches and upgrades at scale are a nightmare. And SDS takes compute resources away from HPC applications and increases their cost.

When you look at NVMe-based storage arrays, you will find that many vendors do not customize their hardware. Instead, they provide a generic server installed with SDS software that uses NVMe SSDs. These are, at best, low-end storage appliances. Others do customize their hardware but treat NVMe flash as a "bolt-on" and add it to their legacy AFAs. Both types of vendors say their solution boosts application performance, and they have numbers to show you that it flies. It looks good...but

On your HPC filesystems and applications, AFAs fall flat. You face a variety of issues, including no improvement in user experience, upgradeability challenges, support issues, and more. Your teams are now questioning your plan to purchase several of these storage systems for limited incremental improvements.

## NVMe ≠ NVMe-oF

You must turn risk into opportunity and find a better approach to accelerating actionable results. The business (and your career) depends on it. You need a better way

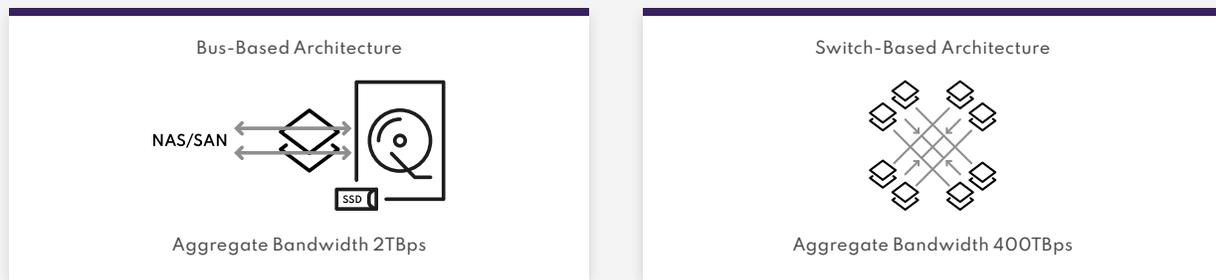
It's called NVMe-Over-Fabrics (NVMe-oF), and it uses a network fabric (such as Ethernet and InfiniBand) that you already have. It maximizes the parallelism and performance of flash, boosting application performance.

What if storage was designed like a network switch with the performance and parallelism

of multiple NVMe SSDs accurately aligned with compute, memory, and network?

An array designed from the ground up for NVMe and NVMe-oF is simply a better way. You can run RDMA or TCP in the same system at the same time. You can have a single shared namespace that scales in a linear fashion. NVMe is ubiquitous and NVMe-oF is ready today.

### STORAGE SYSTEMS DESIGN BUS VS. SWITCH



## DON'T BE LEFT BEHIND

Deploy an NVMe-oF array now, and be a data center hero

To how one NVMe-oF array vendor is redefining HPC performance visit

[www.pavilion.io/solutions](http://www.pavilion.io/solutions)