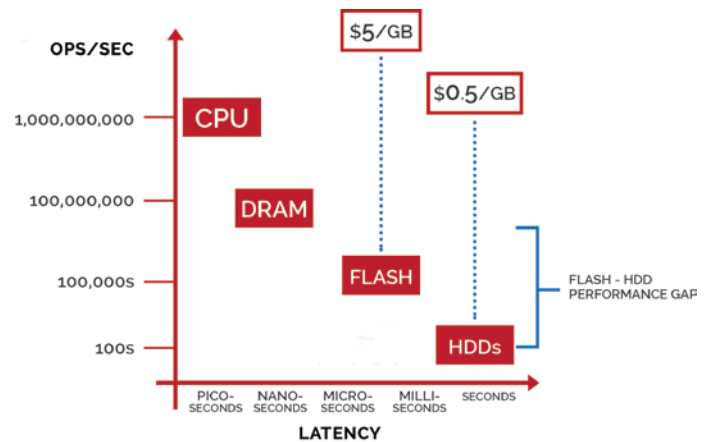


Storage Fusion Xcelerator™ (SFX)

The Only Flash Caching Technology Designed for Big Data

Now, realize the performance promise of flash at a price-point that's closer to lower cost, high density hard disk drives.

Flash cache file systems are purpose-built to bridge the growing performance gap between CPUs and hard disk drives (HDDs). In today's market, flash memory commands premium pricing, as it offers orders of magnitude improvements in reducing latency and increasing IO operations per second (IOPS). DataDirect Networks' SFX technology provides a unique, cost-effective approach by highly optimizing flash utilization and delivering on the the performance promise of flash at HDD prices.



SERVER STORAGE PERFORMANCE GAP CHALLENGE

Unlock the Performance Promise of Flash for Big Data

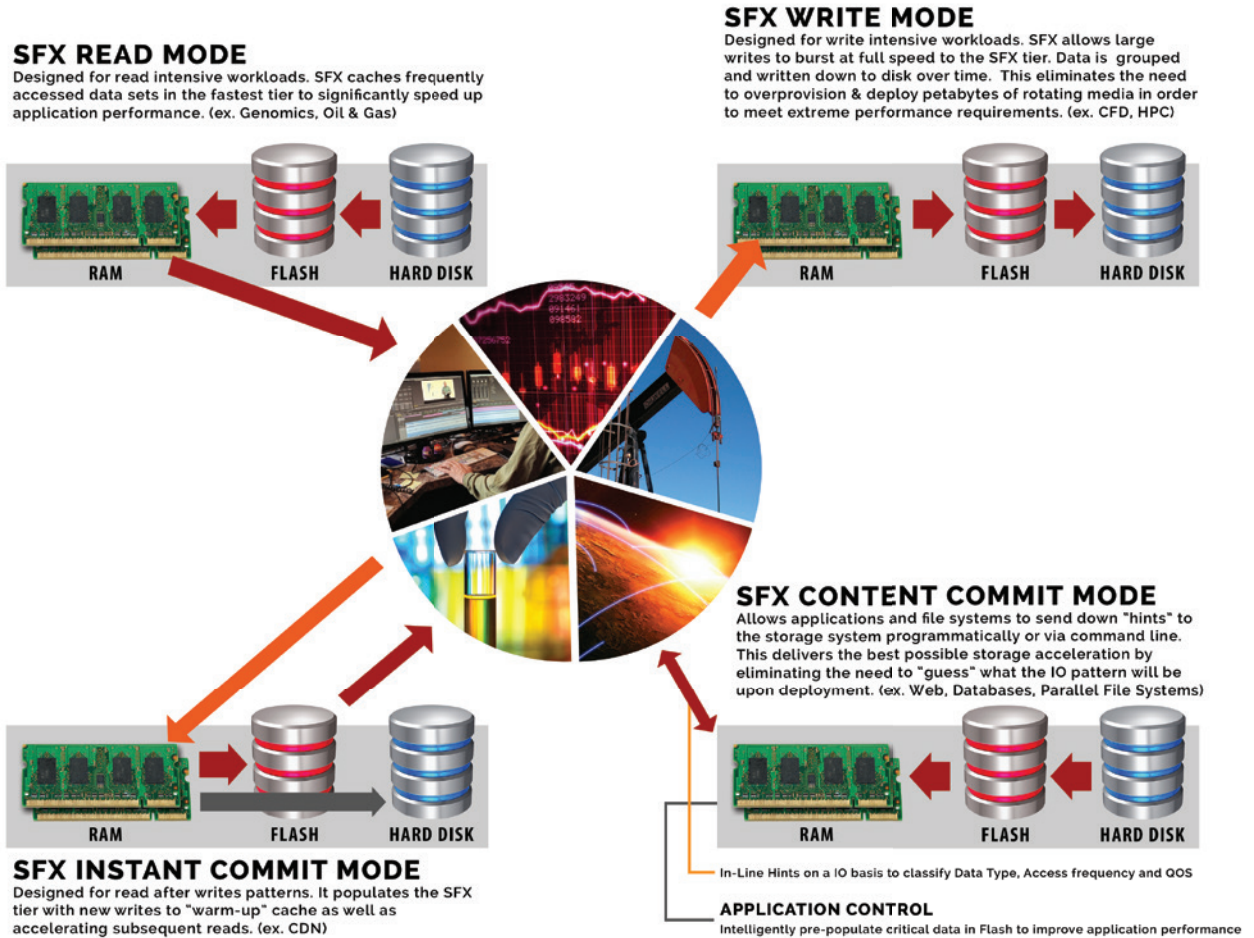
Big data has demanding data access patterns that can turn otherwise fast storage systems into bottlenecks that limit productivity. Buying additional memory and faster storage systems are not the only options for speeding-up big data analysis today. Recently, flash caching solutions have been introduced and positioned as big data analysis acceleration solutions. While general purpose flash cache solutions are able to accelerate some big data workflows, they are most effective with small block IO for databases like Oracle® and MySQL®, and applications like Exchange®. Effectively solving big data problems requires understanding various data types and the specific analysis workflow to speed results. For example, long running parallel applications create checkpoints to avoid restarting calculations if a system failure occurs before the job completes. Checkpointing can take upwards of 15% of job execution time and requires extremely fast write capability to reduce checkpointing duration and to better optimize CPU utilization. Transcoding movies involves ingesting large amounts of data, while immediately accessing that same data in order to convert it to a new encoding format. Achieving optimal data workflows for these and other workflows using flash memory requires giving the applications control of the physical placement of data within the storage system. SFX offers four modes of data control which enables the system to load the right data into flash, at the right time. This maximizes cache hit rates and delivers quicker time to results.

SFX BENEFITS

- ▶ Designed for Big Data e.g. Life Sciences, Rich Media, Financial Services, Oil & Gas.
- ▶ Analyzes and adapts to IOPS and streaming workloads differently.
- ▶ Comprehensive - works in 4 modes to handle any workload and application.
- ▶ Higher cache utilization – application & file system hinting promotes/demotes data.
- ▶ Quicker response times and IOPS for same capacity at lower cost and power consumption.
- ▶ Integrated with DDN's GRIDScaler™, EXAScaler™ and hScaler™ parallel file systems.

Controlling Flash Content to Accelerate Time to Results

DDN designed SFX from the ground-up to optimize cache content for any workload and application. SFX provides four modes of operation to accelerate the different access patterns of data:



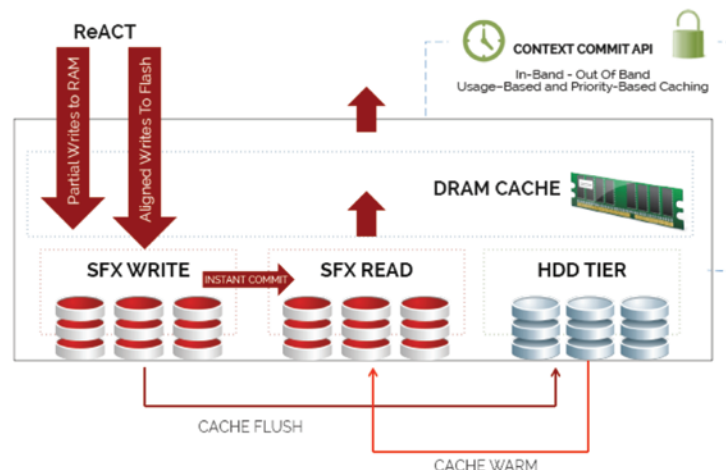
SFX ReACT™s to Mixed Data Streams

The four SFX data control modes work together with DDN's innovative Real-time Adaptive Caching Technology (ReACT) to further optimize cache utilization. ReACT analyzes incoming data streams by directing large aligned block writes to stream directly to rotating media, while caching small block IO to improve performance. This eliminates the need to repopulate the cache because of these large writes. ReACT, combined with the SFX mode that best fits a big data workflow, significantly increases performance and reduces the number of high capacity HDDs required in any solution. This lowers acquisition costs, saves power and delivers unmatched TCO.

SFX & ReACT

THE MANY DIMENSIONS OF SFX ACCELERATION FOR BIG DATA

- ReACT works with SFX to bypass DRAM for fully aligned writes
- Small IOs get cached while large IOs are identified and bypassed
- Helps accelerate Big Data workloads with a combination of highly streaming as well as transactions IOs



Accelerating DDN's GRIDScaler and EXAScaler Parallel File Systems with SFX

Parallel file systems are designed by high performance computing (HPC) engineers for research computing. Their job is to balance data access across large numbers of storage resources to enable reading and writing from many servers in parallel. They scale to 1000s of clients and are ideal for concurrent file and directory access. Parallel file systems are a natural fit for streaming and large file access, but are not well known for small file performance. DDN combines GRIDScaler and EXAScaler with SFX to dramatically improve metadata performance, small IO operations from large data files and small file throughput. Applications direct GRIDScaler or EXAScaler to load frequently accessed data into SFX cache either programmatically or through hints via a command line interface. Pre-warming cache (or pinning critical data in cache), ensures the most critical data is readily available for the most demanding HPC applications in a way not previously possible.

Delivering TCO Savings with SFX

Finally, SFX eliminates the need to purchase more HDDs than needed for capacity, in order to reach a performance requirement. Now, you will free-up budget for other critical IT initiatives. SFX provides the flexibility to optimize big data analysis solutions across performance, capacity, power consumption and cost. The right combination of SFX, solid state disk drives (SSDs) and HDDs enables the most efficient solution, delivering cost savings and increased productivity. These benefits are only possible with an SFX-powered solution, which is designed for today's big data.

Specifications	SFA7700™	SFA12KX-20E™	SFA12KX-40™
SFX Flash Cache Capacities	Read Min: 800GB Write Min: 1TB Max: 6.4TB	Read Min: 800GB Write Min: 1TB Max: 6.4TB	Read Min: 800GB Write Min: 1TB Max: 12.8TB
Max. Capacity	1.5PB	6.7PB	6.7PB
SSD Drives	1.6TB, 800GB, 400GB read intensive 6Gb/s SAS (read/write ratio 90/10), <3 full writes/day 800GB, 400GB, 200GB mixed use 6Gb/s SAS (read/write ratio 70/30) <10 full writes/day 400GB, 200GB write intensive 6Gb/s SAS (read/write ratio 50/50) <30 full writes/day		
File System Solutions	GRIDScaler, EXAScaler		

Learn more at ddn.com

DataDirect™
NETWORKS

9351 Deering Avenue
Chatsworth, CA 91311

1.800.837.2298

www.ddn.com