

APACHE SPARK SOLUTIONS FOR ANALYTICS



Supercharging Spark with Vexata

- Performance at Scale**
 Vexata offers a resilient and high-performance enterprise storage solution, with linear scaling of performance as capacity scales.
- Breakthrough Economics**
 5X more performance at the same capacity at half the price and with lower complexity than leading allflash arrays (AFAs).
- Simplified Operations**
 Large predictive analysis and autonomous decision-making deployments run smoothly with lower operational risk.
- Next Generation Architecture**
 NVMe and NVMe-oF storage system that scales from TBs to PBs in a compact 6U chassis and is a SAN and NAS, gives Spark five times more performance without sacrificing capacity.

Enterprise data growth, especially the amount of active data that must be transacted and analyzed in real-time, is doubling every few years. Businesses extract value from real-time information by analyzing and discovering trends and patterns across petabytes of active and archived data. This analysis helps to achieve business objectives, driving revenue and creating competitive advantage.

When analyzing this data, a business needs servers and shared storage infrastructure designed for speed and scale. This must be achieved without compromising the enterprise data management features required to run Spark for ML/AI, real-time data analytics, and other large-scale data analysis.

What Spark users need is cost-effective, high-performance storage system that fully utilizes the performance of NVMe and modern, compute cores to get the most from Apache Spark analytics. Without performance delivered at scale, Spark applications must be deployed across multiple storage arrays and many servers to avoid performance bottlenecks, increasing cost and complexity.





Vexata VX-100 Scalable NVMe Systems

The VX-100 family of Scalable NVMe Flash Arrays are purpose built to support large scale block and file deployments, ideal for high performance Spark applications. Vexata unleashes the performance of NVMe storage to provide Spark users with a storage infrastructure that can be deployed across Gigabit Ethernet or Fibre Channel networks. Vexata delivers cost/performance ratios far superior to any AFA on the market. Vexata's next-generation architecture enables a **dramatic increase in performance at substantially less cost while greatly simplifying the data infrastructure.**



"Through our collaboration, Vexata not only provides a very fast storage layer for accessing the data, but it is so fast that, combined with Levyx's persistent data frame layer, they deliver an ideal alternative to memory for hosting live datasets in a bulletproof enterprise-grade system."

Reza Sadri, CEO of Levyx

			
PERFORMANCE AT SCALE	BREAKTHROUGH ECONOMICS	SIMPLIFIED OPERATIONS	NEXT GENERATION ARCHITECTURE
<p>Linear scaling of performance as capacity increases</p> <ul style="list-style-type: none"> 10X lower latency than AFAs across 100's of terabytes 5X higher DB query/update performance per server Ingest at 1.5TB/m Throughput of 3TB/m management reduces management cost Millions of IOPS Record-setting STAC-A3 testing using Spark¹ 	<p>Reduce application spend</p> <ul style="list-style-type: none"> Half cost of AFAs >4X reduction in servers Dramatic reduction in software licenses Multi-chassis management reduces OPEX 	<p>Consolidate multiple applications with unified monitoring and management</p> <ul style="list-style-type: none"> Consolidate storage services for applications Monitor & manage storage services across the hall or across the world Multi-chassis management Supports physical and virtualized application 	<p>Multiple simultaneous SAN/NAS applications</p> <ul style="list-style-type: none"> Ultra-high-throughput, ultra-low latency, and millions of IOPS Scales from TBs to PBs in a 6U chassis Supports NVMe attached Fibre Channel (FC-SCSI) and NVMe over fabric (NVMe-oF) using FC or Ethernet Supports Spark and VMs running Spark on vSphere, Hyper-V, and XenServer hypervisors Protocols FC-SCSI, S3, GPFS, SMB, HDFS, S3, NFS

¹The complete STAC Report™ is available at www.STACresearch.com/LEVX180608.

APACHE SPARK SOLUTIONS FOR ANALYTICS



Scalable Performance with Breakthrough Economics

The Vexata VX-100 scalable NVMe flash array performance is ideal for Spark analytic workloads, with over 70GB/s of mixed workload bandwidth, 7M IOPS, less than 200µs of sustained application latency, and no DRAM caching. With Vexata, linear scaling of these performance specifications from TBs to PBs is implemented through a bladed architecture that adds performance and capacity. Vexata provides this performance in dense, powerful system that, replaces racks of legacy storage equipment with a single 6U chassis.

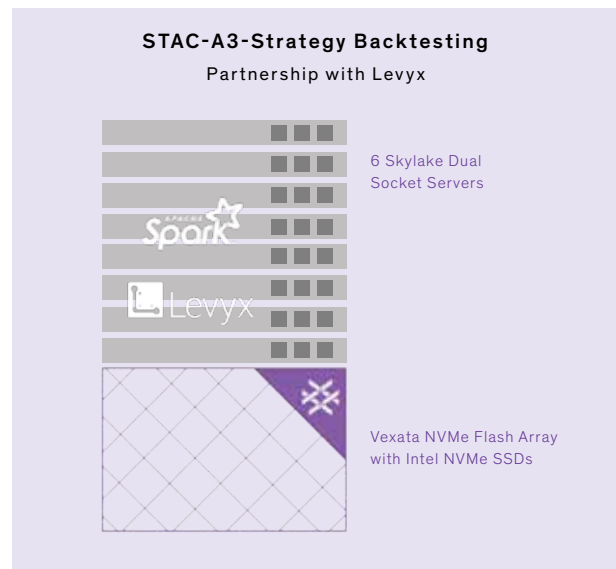
By delivering high-performance data infrastructure to accelerate data ingest and analysis across Spark analytics, data architects can support real-time analytics and deliver faster business outcomes. The Vexata VX-100 family is core to providing Apache Spark with the performance of NVMe SSDs at scale with enterprise-class services.

Spark use cases

- Strategy Backtesting
- Machine Learning/Deep Learning/AI
- Fraud Detection
- IoT, Big Data, and Machine Data

In recent tests submitted to STAC®, the Securities Technology Analysis Center, the Vexata VX-100F used Spark and beat existing performance benchmarks for backtesting of algorithmic trading systems.

- 69.5 Instrument sims/sec with only 6 nodes
- 69% faster on SWEEP.SPEED2.BASE
- 1/3rd the cost of previous solutions



Getting Started with Vexata

Vexata has improved application performance by 3-10X for some of the world's largest corporations on their existing infrastructure. Its patented architecture for accelerated storage ensures Spark applications are not waiting on I/O. The VX-100 delivers Tier 0 performance across large datasets at a fraction of the cost of existing storage offerings, providing Spark applications with solid-state NVMe storage performance at scale with enterprise-class services.

Vexata will provide a free assessment of your Spark applications. We will discuss workloads with your application/analytics team, work with your operations group to design an integration plan and provide your IT team with a budgetary TCO analysis.

ABOUT VEXATA: Founded on the premise that every business is challenged to deliver cognitive, data-intensive applications, Vexata delivers 10X performance AND efficiency improvements at a fraction of the cost of existing all-flash storage solutions.

Learn more at www.vexata.com. Contact us at 1.408.931.6334 or info@vexata.com.