Solution Profile

Supercharge Kafka: Accelerate Real-Time Innovation at Scale

With AI and machine learning (ML) revolutionizing industries faster than ever, staying competitive comes down to how well you harness and analyze data in real time.

Fuel Your Competitive Edge with Kafka

When you act on data faster, you make smarter decisions, keep customers engaged, and foster meaningful innovation.

Apache Kafka®, the backbone of real-time data streaming, gives you the power to process billions of events every day. Whether it's IoT device streams, social media feeds, customer transactions, predictive analytics or AI pipelines, Kafka helps you turn streams of raw data into actionable insights. It's how you spark innovation, make confident choices and deliver standout experiences that set your business apart.

Over 70% of Fortune 500 companies rely on Kafka to stay ahead. Retail leaders use it to streamline supply chains in real time, and financial institutions count on it to detect fraud the moment it happens. As data volumes skyrocket, your ability to process billions of events daily will determine how adaptable and impactful your business can be in our increasingly data-driven world.

What's Holding Your Kafka Strategy Back?

Despite Kafka's power, managing its demands can create roadblocks that slow your progress. These hurdles directly impact how efficiently you can operate and the outcomes your business delivers:

- **Scaling Bottlenecks:** Adding Kafka brokers often requires more hardware, which drives up costs and complexity.
- **Retention and Compliance Challenges:** Balancing real-time processing with long-term storage for analytics and governance can be resource intensive.
- **Fragmented Workflows:** Siloed systems force you to manage redundant storage and complex integrations, drain your time and resources, and make managing data harder than it needs to be.
- **Operational Strain:** Managing replication, rebalancing and recovery across your Kafka clusters can consume valuable resources and open the door to errors.

If these challenges remain unchecked, they can stall your operations, increase costs and make it harder to deliver the fast, meaningful insights Kafka is designed to provide.

Empowering Your Kafka Potential To Achieve More

Tackling the inherent complexities of managing Kafka environments requires a unified strategy that simplifies your operations, scales with your needs and helps you keep costs in check. Here's what you should prioritize to overcome these challenges and fully unlock Kafka's capabilities:

- **Scalability Without Complexity:** Decouple compute and storage so you can scale each independently to match your business needs — without over-provisioning or costly sprawl.
- **Effortless Retention and Compliance:** Use solutions that handle real-time processing and secure, long-term retention with minimal effort.
- **Seamless Workflows:** Gain unified access to real-time and historical data without redundant systems or complex migrations.
- **Simplified Operations:** Automate essential tasks like replication, recovery and rebalancing so you spend less time troubleshooting and more time driving innovation.

To maximize Kafka's full potential — and empower your business to innovate, act faster and stay ahead — you need a solution that scales with your growing demands, streamlines your daily workflows, reduces operational headaches and delivers the reliability you count on to meet critical deadlines.

Hitachi Vantara Virtual Storage Platform One: A Unified Solution for Kafka Excellence

Your Kafka environment powers real-time insights and drives innovation, but managing its high-throughput, metadata-intensive workloads can be complex. Hitachi Vantara Virtual Storage Platform One (VSP One) changes the game by integrating file, object and block storage on a unified data platform with hybrid cloud flexibility. Purpose-built for Kafka's unique demands, VSP One simplifies your workflows, boosts performance and helps you focus on delivering impactful results.

How VSP One Transforms Your Kafka Environment

When managing Kafka's demanding workloads, you need a solution that not only keeps up but also helps you get ahead. Here's how VSP One delivers a scalable, efficient and resilient data foundation tailored for Kafka's unique needs:

Accelerate Kafka Broker Performance

You need your Kafka brokers to focus on what they do best: managing partitions, processing messages and delivering high-throughput performance. VSP One makes that possible by offloading replication and storage tasks, significantly reducing CPU strain and eliminating East-West traffic. This optimization boosts broker efficiency and accelerates data streaming.

With VSP One's AI- and ML-powered 4:1 Data Reduction Guarantee, you manage Kafka's metadata-intensive workloads without over-provisioning, cutting infrastructure costs while maintaining exceptional performance across your clusters.
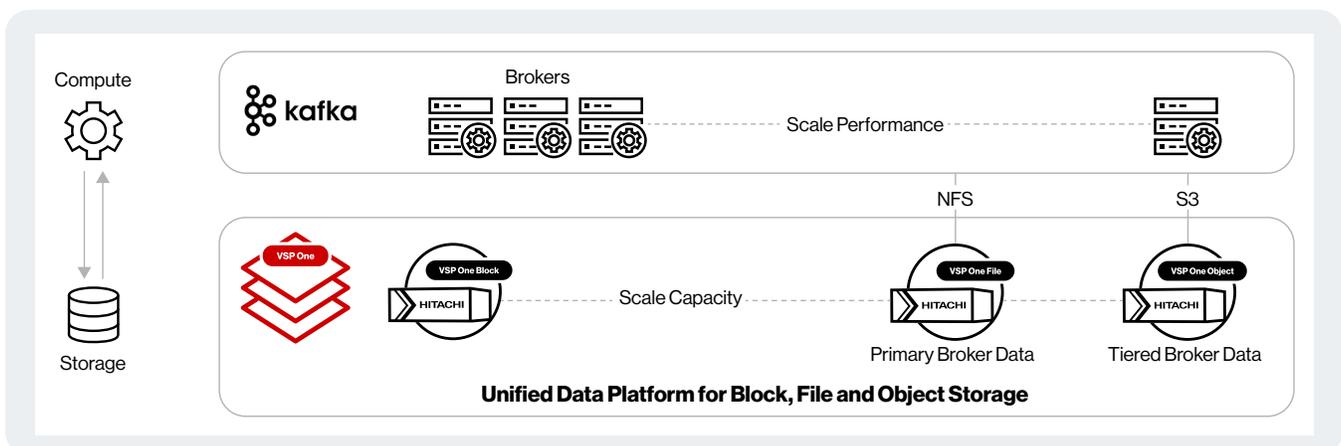
Maximize Real-Time Streaming and Historical Data Retention

Kafka's strength lies in handling active event streams and long-term historical data — VSP One enhances this dual functionality with integrated block, file and object storage on a unified data platform. You get seamless real-time performance for active data streams, while the Effective Capacity Guarantee ensures you store more data in less physical space.

By reducing costs and delivering high performance, VSP One keeps your data ready for analytics, compliance and advanced workloads like AI and ML. Your event streams run at their peak, and you stay agile for future demands.

Stay Resilient at Any Scale for Critical Kafka Workloads

Downtime isn't an option when your business relies on real-time data. With VSP One, your Kafka clusters stay online — even if a node or broker fails — ensuring your operations never miss a beat. By offloading resiliency functions to VSP One File, broker recovery is almost instantaneous and requires minimal intervention. When a broker fails, no data is lost, and

Compute — Brokers — Scale Performance — NFS — S3 — Storage — VSP One — VSP One Block — VSP One File — Primary Broker Data — VSP One Object — Tiered Broker Data — Scale Capacity — **Unified Data Platform for Block, File and Object Storage**

deploying a replacement broker allows your cluster to pick up right where it left off, returning to full operational status with virtually no impact on performance.

VSP One File eliminates the need for brokers to rebuild resiliency within the cluster. Instead of consuming CPU cycles and generating East-West traffic to recover, VSP One ensures that resiliency is maintained at the storage layer. This approach drastically reduces recovery times, keeps brokers focused on streaming data rather than rebuilding it, and minimizes administrative overhead. Backed by a 100% Data Availability Guarantee, VSP One ensures that you can confidently scale your Kafka environment, meet high-availability demands and keep your business running smoothly — no matter what.
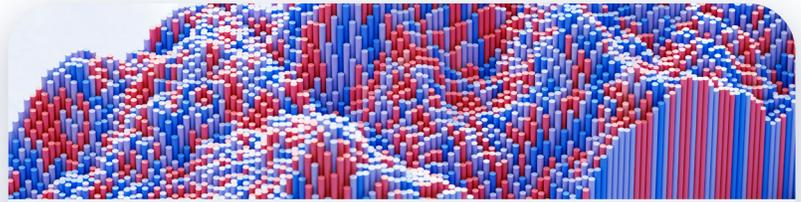
Simplify Kafka Cluster Management

Managing replication, rebalancing and recovery shouldn't slow you down. VSP One simplifies these critical tasks, freeing you to focus on driving growth and innovation. By decoupling compute and storage, you can scale each independently to handle sudden traffic spikes or support long-term data growth effortlessly.

With Modern Storage Assurance, your investment stays future-proof, delivering years of optimized storage performance that evolves with your needs. You'll spend less time firefighting and more time achieving meaningful outcomes.

Redefine What's Possible With Kafka and VSP One

VSP One unifies Kafka's operational and storage demands, so you can focus on what drives your business forward: delivering real-time insights, empowering your teams and exceeding customer expectations. With infrastructure engineered for Kafka's unique challenges, VSP One provides the scalable, reliable foundation you need to tackle today's demands and tomorrow's opportunities.

Let's redefine what Kafka can do — together.

Why VSP One Makes Kafka Work Better for You

To get the most out of Kafka, you need more than a solution that handles its complexities — you need one that simplifies your workflows, accelerates your insights and keeps your operations running smoothly. VSP One helps you:

- **Accelerate Your AI/ML Pipelines.** Process billions of events daily with a unified platform designed for Kafka's real-time and batch workloads. Backed by a 100% Data Availability Guarantee, you can trust your pipelines to run continuously, without interruptions, enabling faster insights and nonstop innovation.

- **Cut Costs Without Compromising Performance.** Decouple storage from compute to avoid broker sprawl and shrink infrastructure overhead. With VSP One's 4:1 Data Reduction Guarantee, you maximize storage efficiency, reduce costs and scale your business without over-provisioning.

- **Streamline Compliance and Governance Without the Stress.** Stay ahead of regulations with immutable snapshots, built-in data protection and integrated file and object storage that simplifies long-term retention. Powered by Modern Storage Assurance, VSP One ensures your storage adapts to evolving demands: It provides the performance and efficiency you need to meet compliance requirements without disrupting workflows.

- **Maximize Kafka Streaming and Analytics at Scale.** Handle Kafka's metadata-intensive workloads with FPGA-accelerated storage designed to optimize performance. Lower system CPU utilization, eliminate the need for 3x replicas and free your brokers to focus on managing partitions, processing messages and scaling operations.

Hitachi Vantara
A Hitachi Group Company