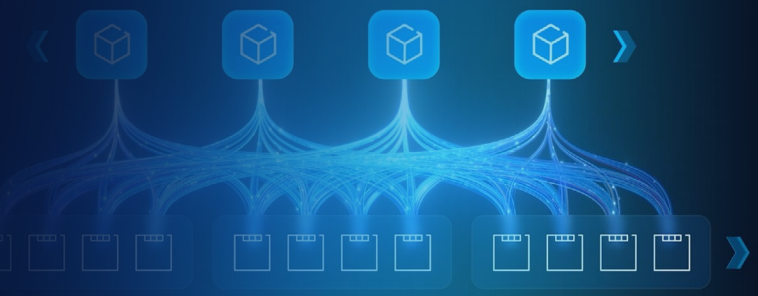
Solution Brief

VAST DASE Architecture

A New Foundation for AI Infrastructure

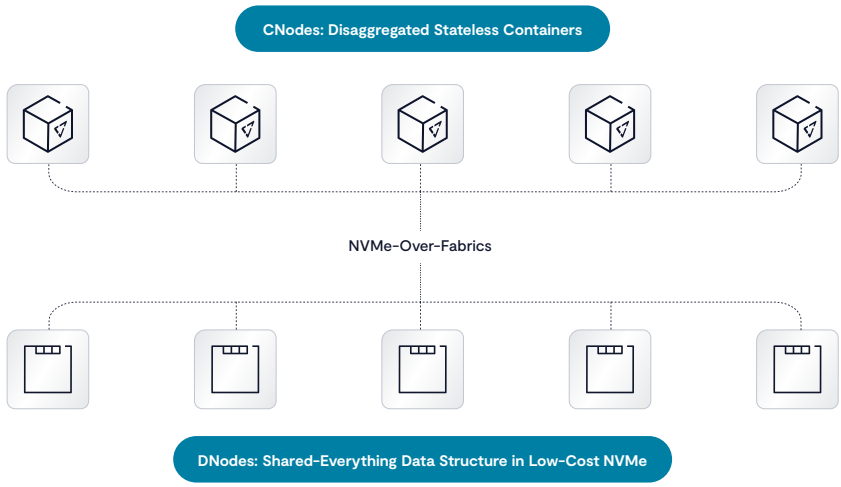


For more than two decades, enterprise data infrastructure has been built on shared-nothing architectures that require constant coordination between independent servers. While effective for early cloud-scale workloads, these systems force organizations to choose between simplicity, scalability, performance, resilience, and economics. AI, analytics, and real-time applications now demand all five simultaneously, exposing the limitations of architectures built around distributed coordination. VAST's Disaggregated Shared-Everything (DASE) architecture was designed from a clean slate to eliminate these limitations, creating a unified AI Operating System that breaks the architectural compromises of the past.

Designed To Eliminate Yesterday's Infrastructure Tradeoffs

Shared-Everything Architecture Replaces Legacy Shared-Nothing Designs

Traditional distributed systems depend on tightly coordinated storage nodes that constantly exchange metadata, locking information, and cache updates as clusters grow. This east-west communication becomes increasingly difficult to scale while introducing operational complexity and recovery challenges. DASE fundamentally rethinks this model



Why VAST

Shared-Everything Architecture

Eliminate distributed system bottlenecks with globally shared transactional state.

Independent Scaling

Scale compute and storage independently to optimize performance and infrastructure investment.

Infrastructure Consolidation

Replace fragmented storage, databases, streaming, and AI infrastructure with one unified architecture.

by separating stateless compute (CNodes) from persistent shared system state and introducing globally shared transactional data structures (DNodes). Every stateless CNode can process any workload against the same shared data on all-flash DNodes without communicating directly with other processors, dramatically simplifying cluster operations while delivering predictable performance from petabytes to exabytes.

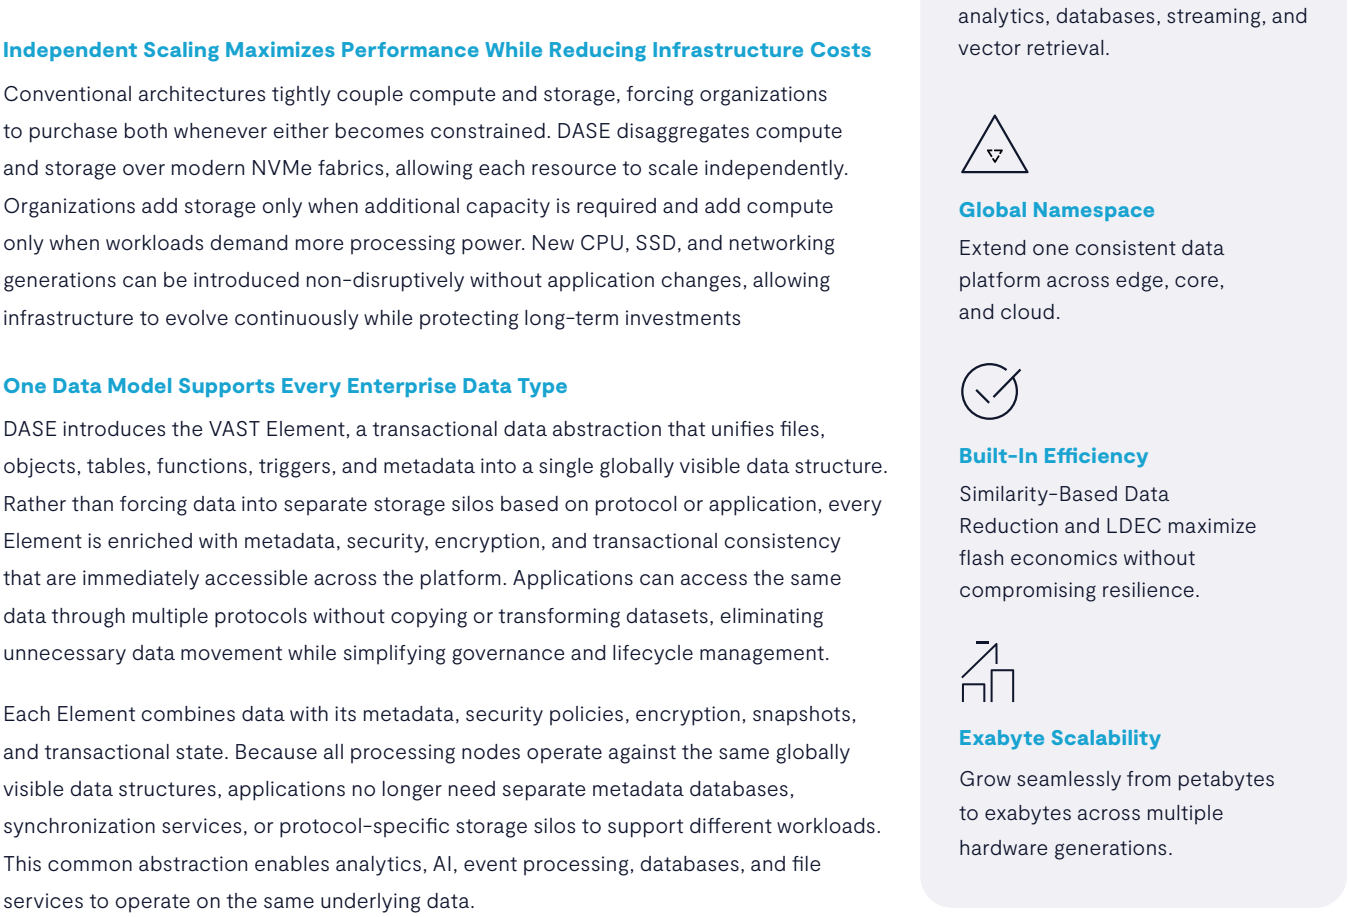
Independent Scaling Maximizes Performance While Reducing Infrastructure Costs

Conventional architectures tightly couple compute and storage, forcing organizations to purchase both whenever either becomes constrained. DASE disaggregates compute and storage over modern NVMe fabrics, allowing each resource to scale independently. Organizations add storage only when additional capacity is required and add compute only when workloads demand more processing power. New CPU, SSD, and networking generations can be introduced non-disruptively without application changes, allowing infrastructure to evolve continuously while protecting long-term investments

One Data Model Supports Every Enterprise Data Type

DASE introduces the VAST Element, a transactional data abstraction that unifies files, objects, tables, functions, triggers, and metadata into a single globally visible data structure. Rather than forcing data into separate storage silos based on protocol or application, every Element is enriched with metadata, security, encryption, and transactional consistency that are immediately accessible across the platform. Applications can access the same data through multiple protocols without copying or transforming datasets, eliminating unnecessary data movement while simplifying governance and lifecycle management.

Each Element combines data with its metadata, security policies, encryption, snapshots, and transactional state. Because all processing nodes operate against the same globally visible data structures, applications no longer need separate metadata databases, synchronization services, or protocol-specific storage silos to support different workloads. This common abstraction enables analytics, AI, event processing, databases, and file services to operate on the same underlying data.



AI Operating System

One platform for AI pipelines, analytics, databases, streaming, and vector retrieval.

Global Namespace

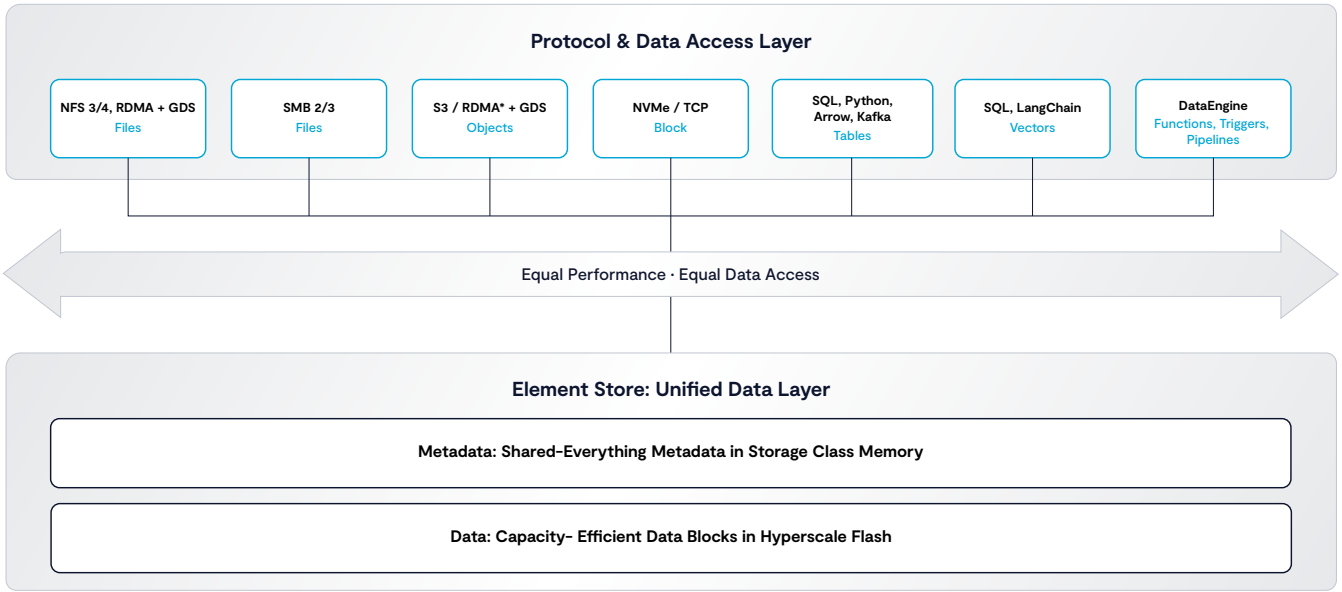
Extend one consistent data platform across edge, core, and cloud.

Built-In Efficiency

Similarity-Based Data Reduction and LDEC maximize flash economics without compromising resilience.

Exabyte Scalability

Grow seamlessly from petabytes to exabytes across multiple hardware generations.

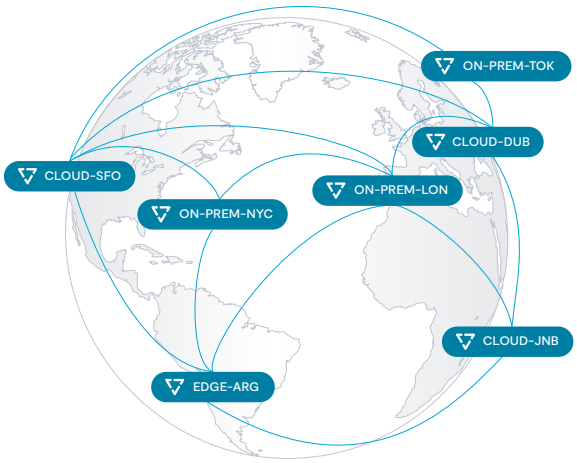


Built-In Resilience Simplifies Data Protection At Any Scale

DASE incorporates several architectural innovations that improve efficiency without sacrificing resilience. Byte-granular snapshots enable nearly unlimited recovery points with minimal overhead, while globally visible transactional metadata supports consistent time travel and fine-grained replication. Unlike conventional erasure coding, which reconstructs entire protection groups during failures, Locally Decodable Erasure Coding (LDEC) rebuilds only the affected data, dramatically accelerating recovery while maintaining extremely low storage overhead. Reduction reduces storage requirements across even compressed and encrypted datasets. Together, these capabilities make all-flash infrastructure economically practical while delivering enterprise-class data protection.

One Global Platform From Edge To Cloud

DASE extends beyond a single cluster through the VAST DataSpace, creating a globally federated data platform with a shared namespace that spans data centers, edge locations, and cloud environments. Policy-based data sharing allows organizations to present a unified view of distributed data while maintaining local performance and global consistency. By eliminating the traditional tradeoffs between locality and scalability, DASE enables enterprises to build globally distributed AI agents, analytics, databases, streaming, and HPC workloads on a single architectural foundation.





DASE: A New Architectural Foundation

Rather than optimizing one layer of the infrastructure stack, DASE rethinks the underlying architecture itself. By eliminating distributed coordination, unifying data under a single transactional model, and disaggregating compute from persistent state, VAST provides a foundation capable of supporting enterprise storage, databases, streaming, analytics, and AI on one continuously scalable platform.

About VAST Data

VAST Data is the AI Operating System company – powering the next generation of intelligent systems with a unified software infrastructure stack that was purpose-built to unlock the full potential of AI. The VAST AI OS consolidates foundational data and compute services and agentic execution into one scalable platform, enabling organizations to deploy and facilitate communication between AI agents, reason over real-time data, and automate complex workflows at global scale. Built on VAST's breakthrough DASE architecture – the world's first true parallel distributed system architecture that eliminates tradeoffs between performance, scale, simplicity, and resilience – VAST has transformed its modern infrastructure into a global fabric for reasoning AI. Learn more at vastdata.com and follow VAST Data on LinkedIn, YouTube and X.

©2026 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.
07/26 | 12917