



EXECUTIVE TECHNOLOGY BRIEFING

Performance Without Compromise:

GateStor Virtual Storage Architecture™ and the ISO 20022 Imperative for Financial Institutions

Prepared by GateStor Data Systems Corporation

40–60%

of enterprise IT spend is storage
infrastructure

10×

faster node interconnect vs. fastest
network alternatives

15 EB

maximum cluster capacity (1024
nodes × 15 PB)

Executive Summary

ISO 20022 — the global financial messaging standard now mandated across SWIFT, TARGET2, CHIPS, Fedwire, and major card networks — places unprecedented demands on the storage infrastructure underpinning payment systems, securities settlement, and core banking operations. High message volumes, strict latency requirements, zero-tolerance for data corruption, and continuous availability obligations are not aspirational targets; they are contractual and regulatory commitments. This briefing explains how GateStor's Virtual Storage Architecture™ (VSA), delivered through the SolidPOWER7000® and EPOCH® Power Cluster platforms, directly satisfies the performance, availability, and reliability requirements embedded in the ISO 20022 framework — and why the architectural choices that distinguish VSA from conventional RAID-based storage are not incremental improvements but fundamental prerequisites for compliance.

ISO 20022 and the Storage Infrastructure Imperative

ISO 20022 is far more than a messaging standard. It is the architectural backbone of the modern financial system's data exchange layer, governing how payment instructions, securities transactions, trade finance messages, and regulatory reports flow between institutions, clearing houses, central banks, and end clients. Its adoption is now mandatory across the world's largest payment infrastructures, with full migration deadlines enforced by SWIFT for cross-border transactions, the European Central Bank for TARGET2, and the U.S. Federal Reserve for Fedwire Funds Service.



What ISO 20022 introduces that previous standards did not is a dramatic increase in message richness, complexity, and volume. Where legacy MT messages carried compact, abbreviated fields, ISO 20022 MX messages carry structured XML payloads with extensive data elements — purpose codes, legal entity identifiers, extended remittance information, and structured address fields. This richer data structure means that each message is larger, each transaction generates more I/O operations, and the storage systems that process and archive these messages must deliver substantially higher throughput without sacrificing the sub-millisecond response times that settlement finality requires.

Why This Matters for Your CTO and CIO

ISO 20022 compliance is not achieved at the application layer alone. Every message that your institution sends or receives must be processed, validated, stored, and retrievable in real time. The storage infrastructure beneath your payment engines, core banking platforms, and compliance archives is either a strategic enabler of ISO 20022 compliance — or its most significant bottleneck. Institutions that treat storage as a commodity procurement decision rather than a compliance-critical architecture choice are accepting systemic risk.

The Limitations of Conventional RAID Storage in an ISO 20022 World

For decades, RAID (Redundant Array of Independent Disks) technology has served as the dominant approach to enterprise storage. Its core value proposition — combining multiple disks into a single logical unit with built-in redundancy — was well suited to the storage demands of legacy banking architectures. However, the ISO 20022 migration fundamentally changes the performance envelope within which financial institutions must operate, exposing architectural trade-offs in RAID that were previously manageable but are now disqualifying.

The Write Penalty Problem

RAID levels 5 and 6, the most widely deployed configurations in enterprise banking storage, impose a structural write penalty. Every write operation requires a read of existing data, a calculation of updated parity information, and then two separate write operations — one for the data and one for the parity block. This Read-Modify-Write cycle introduces latency that is architecturally unavoidable in conventional RAID implementations. Under the high-frequency, high-volume write loads that ISO 20022 message processing generates, this penalty compounds across millions of daily transactions, producing throughput degradation that directly impacts settlement processing windows.



The Capacity Versus Protection Trade-Off

RAID-1 mirroring, the other common enterprise protection strategy, eliminates write penalties but sacrifices 50% of raw storage capacity to maintain the duplicate copy. At the petabyte-scale storage footprints required for ISO 20022 message archives, transaction logs, and audit trails — which regulators require to be retained for multi-year periods — this 1:1 capacity cost becomes economically prohibitive and operationally inefficient.

Cluster Architecture Fragmentation

Network-based storage clusters, the traditional scaling path for financial institutions, introduce metadata directory overhead, node-specific data segmentation, and inter-node communication latency that creates non-deterministic performance under load. When data spills from one node's storage allocation to an adjacent node — an inevitable occurrence in high-growth environments — the resulting data movement operations cause severe, unpredictable performance degradation at precisely the moments of peak transaction volume when reliability is most critical.

The Virtual Storage Architecture™: A Purpose-Built Response

GateStor's Virtual Storage Architecture™ was designed from first principles to eliminate the trade-offs described above. Rather than incrementally improving RAID technology, VSA replaces the fundamental control model of storage management with an approach centered on intelligent, real-time coordination of every storage operation at the data level. Understanding what this means in practice requires examining three core architectural innovations.

1. Virtual Memory Mapping: Eliminating Unnecessary I/O

Conventional storage arrays manage data through caching — holding frequently accessed blocks in high-speed memory to reduce disk access latency. This approach works well until cache saturation occurs, at which point performance degrades sharply and becomes unpredictable. In a financial messaging environment processing millions of ISO 20022 MX transactions per day, cache saturation is not an edge case; it is a routine operating condition.

VSA replaces caching with virtual memory mapping, treating all storage — from system memory through to spinning disk and NVMe flash — as a single, continuous logical pool. The system maintains a real-time map of the location and status of every data block across this unified space. This enables the architecture to make intelligent, context-aware decisions: deferring writes to disk until a complete stripe is available, discarding transient data that will be superseded before it ever needs to be committed to permanent storage, and servicing read requests from memory without any disk access when data is already present in the memory tier.



The result is performance that remains consistent and predictable regardless of load — a critical property for payment infrastructure where SLA commitments require deterministic sub-millisecond response times throughout the trading day.

2. Eliminating the RAID Write Penalty

VSA eliminates the Read-Modify-Write cycle that penalizes RAID 5 and 6 implementations by aligning the system's memory page size with its RAID stripe size. Because every write operation delivers a complete stripe to the storage layer, parity is calculated on the fly against the incoming data without any need to first read existing blocks. Partial stripe parity calculations — the root cause of the traditional RAID write penalty — are architecturally eliminated rather than mitigated.

For ISO 20022 payment processing environments where write-intensive workloads are constant and sequential, this is not a marginal performance improvement. It is the difference between storage that scales linearly with transaction volume and storage that becomes a system-wide bottleneck as message throughput increases.

3. The EPOCH® Unified Bus Cluster: Hardware-Level Scalability

The most architecturally distinctive element of the GateStor platform is its cluster interconnect model. Where conventional storage clusters connect nodes via network protocols — introducing Ethernet, InfiniBand, or Fibre Channel overhead and all the complexity that network-based clustering requires — the EPOCH® Power Cluster connects SolidPOWER7000® nodes through a PCIe-based unified hardware bus.

This distinction has profound implications. At the hardware bus level, every node in the cluster can directly address the memory and storage resources of every other node, without routing through network protocol stacks, without maintaining metadata directories that map data locations, and without the data segmentation constraints that require complex redistribution operations when capacity thresholds are crossed. Transfer speeds across the unified bus are five times faster than even the fastest network-based alternatives currently available.

Each SolidPOWER7000® node (active/active) supports up to 120 122.8TB NVMe drives and 15 petabytes of raw capacity. A fully configured EPOCH® cluster of 1024 nodes delivers up to 15 exabytes of seamlessly addressable storage, with performance that scales linearly as nodes are added — no rebalancing, no metadata reorganization, no performance interruption.



Mapping VSA Capabilities to ISO 20022 Requirements

The ISO 20022 standard, and the operational frameworks that financial market infrastructures layer on top of it, embed specific technical requirements for the systems that process, store, and transmit conforming messages. The following analysis maps the key infrastructure requirements across three dimensions — Performance, Availability, and Reliability — to the VSA capabilities that directly satisfy them.

Performance Requirements

ISO 20022 Requirement	VSA Capability	Operational Benefit
High-throughput message ingestion and processing without latency accumulation under peak load	Virtual memory mapping with complete-stripe writes eliminates Read-Modify-Write penalty; data broadcast to all cluster nodes simultaneously on write	Deterministic sub-millisecond I/O response maintained regardless of transaction volume or concurrent host connections
Concurrent support for heterogeneous workloads (payments, securities, FX, AML screening) on shared infrastructure	Selectable per-dataset RAID levels (0,1,3,5,6,10,550,560) applied concurrently; logical volumes independently configured per workload	Each application receives the protection/performance profile appropriate to its criticality without infrastructure silos or dedicated arrays
Scalable throughput as message volumes grow (SWIFT projects 50%+ volume growth from ISO 20022 adoption)	EPOCH® Unified Bus Cluster scales linearly to 1024 nodes; performance grows proportionally with capacity additions	Infrastructure investment protects forward as message volumes grow; no architectural re-platforming required
Low-latency access to historical message data for straight-through processing, reconciliation and AML	Pre-fetching and memory-resident data servicing; NVMe flash tier delivers 5M+ IOPS; SSDs extended as memory per node	Compliance archive queries and real-time enrichment lookups execute without impacting live transaction processing

Availability Requirements

ISO 20022 Requirement	VSA Capability	Operational Benefit
Continuous operation during hardware failures (payment system RTOs measured in seconds, not hours)	Six concurrent mirroring levels including intra-mirroring, cross-host and remote mirroring; hot recovery at all levels; unified bus enables surviving nodes to access any failed node's storage instantly	Node or drive failure produces zero interruption to payment processing; no manual failover intervention required
Non-disruptive capacity expansion to accommodate growing message archives and transaction logs	Physical storage added at any time without system downtime; logical volumes re-allocated dynamically	Regulatory archive growth managed without maintenance windows that would impact payment SLAs



Multi-site resilience for systemically important financial institutions (SIFIs) under DORA, FFIEC, and BCBS requirements	Remote and Reflective Mirroring provide peer-to-peer cross-site replication; Replicated Mirroring distributes critical data to multiple locations simultaneously	Recovery Point Objective (RPO) near zero; geographically distributed data copies satisfy regulatory resilience requirements for critical infrastructure
Support for heterogeneous host environments across merged/acquired entities	Host-agnostic architecture supports InfiniBand, 100GbE, Fibre Channel, OmniBUS®, SATA and SAS concurrently; legacy Fibre Channel and SCSI integrated without disruption	M&A integration and multi-platform environments serviced from a single storage fabric without re-platforming legacy systems

Reliability & Data Integrity Requirements

ISO 20022 Requirement	VSA Capability	Operational Benefit
Zero data loss for in-flight transactions (settlement finality is legally binding under ISO 20022 frameworks)	Non-volatile memory retains all modified data in transit until securely committed to disk; UPS/battery-backed memory is a platform requirement; write-back acknowledgment mode with NV-RAM protection	Settlement instructions and payment confirmations cannot be lost due to power failure or system interruption after write acknowledgment
Data integrity verification across multi-petabyte archives (regulatory audit and litigation requirements)	ECC (Error Checking and Correcting) memory enforced across all nodes; multi-dimensional RAID levels 530 and 630 provide double-drive-fail tolerance beyond standard RAID 6	Bit-level data integrity maintained across entire storage fabric; catastrophic drive failure scenarios that defeat standard RAID protection are covered
Immutable audit trails for regulatory reporting (MiFID II, EMIR, Dodd-Frank, BCBS 239)	WORM snapshot integration through VaultSAFE® architecture; data protection levels changeable without physical reconfiguration ensures audit metadata remains intact	Regulatory audit archives are demonstrably tamper-evident and retrievable on demand without impacting production system performance
Predictable performance behavior for SLA compliance reporting to regulators and correspondent banks	Memory mapping delivers consistent, non-erratic performance as load increases (unlike cache-based systems that degrade at saturation); SNMP monitoring from any management station	Performance SLAs can be measured, reported, and defended to regulators with confidence; non-deterministic performance events eliminated architecturally



Strategic and Competitive Implications for Banking Leadership

Regulatory Risk Reduction

Financial regulators, including the ECB, Federal Reserve, Bank of England, and MAS, are increasingly scrutinizing the technology infrastructure of systemically important institutions as part of operational resilience assessments. The ability to demonstrate that storage infrastructure is architecturally designed for the performance and availability requirements of ISO 20022 — rather than bolted together from commodity components — is becoming a material factor in regulatory examinations. VSA's deterministic performance model, its non-volatile data protection, and its continuous availability architecture provide defensible evidence of infrastructure fitness that commodity RAID solutions cannot offer.

Total Cost of Ownership Over the ISO 20022 Migration Lifecycle

The economics of ISO 20022 infrastructure must be evaluated over a five-to-ten-year horizon that accounts for message volume growth, regulatory archive expansion, and the technology refresh cycles that conventional storage architectures require. VSA's linear scaling model means that capacity and performance investments made today extend seamlessly to accommodate future volumes without architectural re-platforming. The elimination of the RAID write penalty reduces the compute overhead allocated to storage operations, freeing processing capacity for message enrichment and compliance workloads. And the ability to apply different data protection levels per dataset means that costly high-redundancy configurations can be reserved for mission-critical data sets while less critical data is stored more efficiently.

Operational Resilience as a Competitive Differentiator

In an ISO 20022 world, the reliability of your payment infrastructure is directly visible to your correspondent banking partners, corporate clients, and market counterparties. Settlement failures, processing delays, and reconciliation errors — however caused — carry reputational consequences that compound over time. Institutions that can demonstrate architecturally grounded operational resilience, supported by storage infrastructure with provable continuous availability characteristics, are better positioned to attract and retain the large corporate and institutional clients for whom payment reliability is a non-negotiable requirement.

Technical Architecture Summary for ISO 20022 Environments

For the institution's technology leadership team, the following summarizes the specific VSA capabilities and their mapping to ISO 20022 infrastructure requirements:



- Virtual Memory Mapping replaces caching with a continuous, predictable memory management model that eliminates performance degradation under peak ISO 20022 message loads.
- Complete-stripe write architecture removes the RAID 5/6 write penalty, delivering consistent write performance for the high-throughput message ingestion that ISO 20022 demands.
- EPOCH® Unified Bus Cluster provides hardware-level node interconnect at 10× the speed of network alternatives, with linear scaling to one exabyte and no data segmentation constraints.
- Concurrent multi-level RAID (0,1,3,5,6,10,550,560) per dataset enables ISO 20022 message archives, live transaction data, and compliance logs to coexist on shared infrastructure with individually optimized protection profiles.
- Six-level concurrent mirroring including remote and reflective configurations satisfies DORA, FFIEC, and BCBS operational resilience requirements for geographically distributed data protection.
- Non-volatile memory with write-acknowledgment guarantees zero data loss for in-flight settlement instructions, satisfying the finality requirements of ISO 20022 payment frameworks.
- Dynamic, non-disruptive reconfiguration allows protection levels, capacity allocations, and host assignments to be modified without system downtime or interruption to payment processing.
- Host-technology agnosticism supports the heterogeneous infrastructure reality of most financial institutions, including legacy Fibre Channel and SCSI environments alongside modern NVMe and InfiniBand platforms.
- SNMP-based enterprise management from any station provides the centralized monitoring and reporting capability required for SLA compliance documentation and regulatory oversight.

Conclusion

ISO 20022 is not simply a messaging format upgrade. It is a fundamental transformation of the data demands placed on financial institutions' technology infrastructure, and its implications reach from the application layer down to the physical storage systems that underpin every transaction processed, every message archived, and every regulatory report generated.

GateStor's Virtual Storage Architecture™, embodied in the SolidPOWER7000® and EPOCH® Power Cluster platforms, was designed precisely for the conditions that ISO 20022 creates: deterministic high-throughput performance, continuous availability, multi-layered data integrity, and seamless scalability across the transaction volumes that financial institutions will process in the years ahead.

For banking executives responsible for infrastructure investment decisions, the relevant question is not whether conventional storage can be made to work in an ISO 20022 environment



with sufficient workarounds and overprovisioning. It is whether an architecture purpose-built for this challenge — one that eliminates write penalties by design, scales linearly by hardware, and protects data through non-volatile memory at the hardware level — represents a sounder foundation for the institution's obligations to its clients, its counterparties, and its regulators.

To arrange an architecture briefing or proof-of-concept evaluation:

GateStor Data Systems Corporation | 22 Cotton Rd., Building E, Nashua, NH 03063

Tel: 603.879.0216 | www.gatestor.com