SOLUTION BRIEF

STEELDOME AND SUPERMICRO OFFER NVME STORAGE AT SCALE

Optimizing Time-To-Market for Your AI Infrastructure

TABLE OF CONTENTS

Executive Summary

SteelDome delivers StratiSYSTEM, an infrastructure operating system that delivers clustered storage (StratiSTOR), virtualization and containers (StratiSERV), or an integrated hyperconverged deployment model (HyperSERV). The result is a single, cohesive foundation for running modern workloads with higher resilience, simpler operations, and predictable performance—without stitching together separate products or managing multiple silos.

Paired with Supermicro BigTwin®, SteelDome becomes a validated, high-density solution built for hyperconverged and storage-intensive environments. BigTwin's architecture—four hot-pluggable nodes in 2U, all-NVMe front bays, AIOM (OCP 3.0) networking, redundant Titanium-level power, and optional M.2 RAID1 NVMe boot—aligns with the realities of cluster operations: node-level fault isolation, fast serviceability, consistent performance, and clean scale-out growth.

Expected Outcomes

- Consolidate storage, compute, and orchestration into a single platform—eliminating separate silos and reducing operational overhead
- Improve availability through clustered services, including storage resiliency, virtualization high availability, and live migration

- Accelerate time-to-value with validated Supermicro bundles and deployment patterns optimized for BigTwin density and serviceability

- Deliver performance at scale with high-throughput, low-latency data paths for modern workloads (e.g., NVMe, RDMA, GPUDirect) where appropriate

Solution Overview

SteelDome's architecture consists of multiple components, each delivering distinct capabilities. The diagram below presents a high-level view of the integrated SteelDome and Supermicro architecture.

- StratiSTOR: Massively scalable, software-defined storage providing unified block, file, and object access
- StratiSERV: Virtualization and container platform for orchestrating workloads with enterprise capabilities, including migration, high availability, monitoring, and automation
- HyperSERV: An integrated hyperconverged deployment model that combines StratiSTOR + StratiSERV on the same nodes for unified compute and storage

The platform is designed to be flexible and extensible—supporting broad integration options, modern networking, and multiple deployment models—so customers can select the right architecture without being forced into rigid patterns.

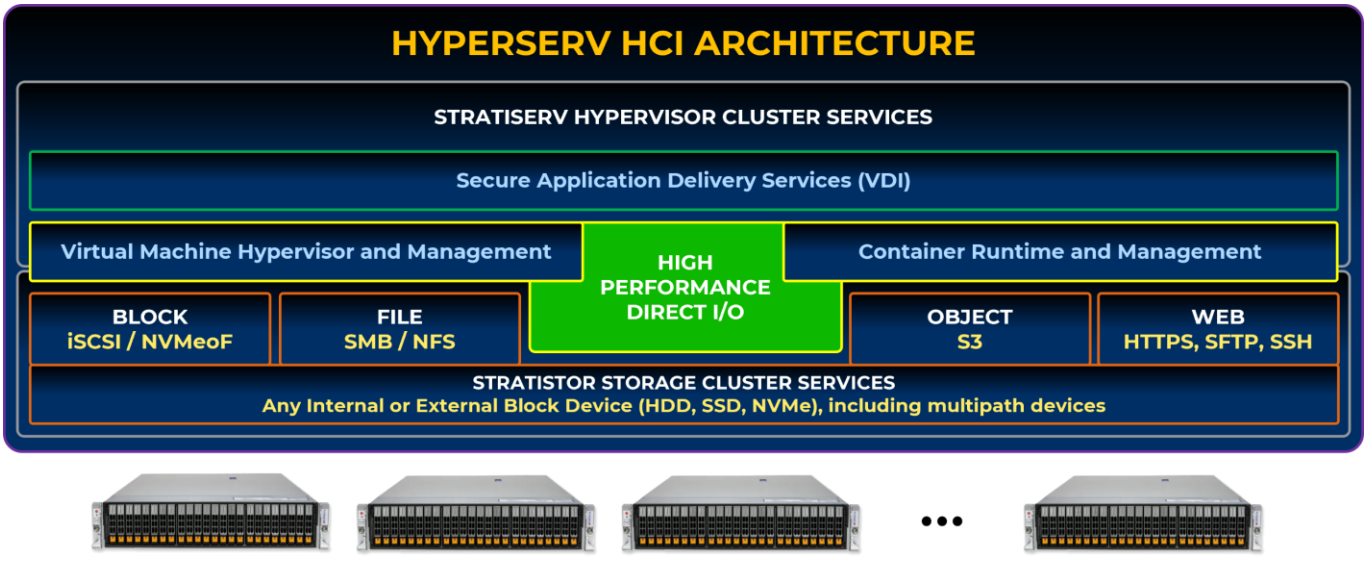

Figure 1 - SteelDome and Supermicro Integrated Architecture

Supermicro BigTwin systems (SYS-221BT-HNR / SYS-222BT-HNR) deliver a high-density, cluster-friendly design:

- Four hot-pluggable nodes in a 2U chassis for density and serviceability
- 6× hot-swap 2.5" NVMe bays per node for an all-flash NVMe design
- AIOM (OCP 3.0) networking for standardized, flexible high-speed NIC options
- Optional M.2 NVMe RAID1 boot for clean separation of boot and data
- Redundant Titanium-level power supplies and a shared cooling design

SteelDome's Supermicro pairing is engineered for validated compatibility and a streamlined operational experience—so customers can focus on outcomes rather than integration and interoperability work. For more information about these two system models, see Supported/Validated System Configurations later in this document.

Customer Challenges and Use Cases

In today's IT environments, customers commonly face the following challenges:

Infrastructure silos and complexity:

Separate SAN/NAS systems, hypervisor stacks, and container platforms multiply operational overhead, increase compatibility risk, and complicate change management—especially for lean teams or distributed sites.

Non-disruptive growth without migrations:

As legacy designs reach their limits, organizations are often forced into forklift upgrades or disruptive migration projects—adding risk, downtime, and cost precisely when stability matters most.

Performance variability under real workloads:

Mixed I/O patterns and multi-tenant environments can introduce unpredictable latency, "noisy neighbor" behavior, and inconsistent throughput—particularly during peak events such as patch cycles, reporting windows, or onboarding surges.

High availability constraints at small scale:

Smaller clusters often want high availability but run into minimum-size requirements, expensive dependencies, or complex architectures that don't match operational reality.

Hardware lifecycle friction:

Refresh cycles become painful when software is tightly coupled to a single vendor, a single architecture, or a rigid bill of materials—reducing flexibility and increasing risk over time.

High-Value Use Cases

The Supermicro BigTwin + SteelDome solution is purpose-built for the following high-value use cases:

Hyperconverged modernization (VMs + containers + unified storage):
HyperSERV enables a "compute-over-data" model—running virtual machines and containers on the same nodes that provide clustered storage. This reduces cost and complexity while improving performance by keeping workloads close to data. BigTwin supports this model exceptionally well, with 4 nodes in 2U, enabling a minimum viable cluster quickly and straightforwardly with linear scale-out.

Private cloud foundations (VM orchestration + storage services):

StratiSERV provides live migration, storage live migration, high availability, and automation features that make private cloud operations practical and maintainable. BigTwin increases compute density per rack while preserving serviceability through hot-pluggable nodes and front-access NVMe, making it ideal for environments that need high utilization without operational fragility.

High-performance file and block storage for enterprise applications:

SteelDome supports broad protocol access (SMB/NFS/iSCSI/NVMeoF/S3) and advanced networking options (RDMA, SR-IOV, VLAN/VxLAN, and more). BigTwin's all-NVMe design aligns naturally with storage-intensive roles such as high-performance file services, consolidated block storage, and mixed enterprise application I/O.

AI/ML and HPC-adjacent pipelines (data paths + GPU integration):

SteelDome supports GPU-aware data movement and acceleration capabilities where applicable, including GPUDirect. BigTwin nodes provide PCIe Gen5 expansion and flexible AIOM networking options, and can be configured for higher-TDP profiles (including liquid-cooling options) to support performance-oriented deployments.

Edge and remote sites (start small, expand safely):

SteelDome supports single-node, two-node (with quorum), and 3+ node configurations—enabling right-sized deployments that scale without redesign. BigTwin density is especially valuable where rack space and power are constrained, delivering multi-node cluster capability in a compact footprint.

Key Features and Benefits

StratiSTOR (clustered storage foundation)

- Unified storage foundation across heterogeneous nodes and configurations
- Built-in data integrity and resilience (checksum-protected structures and self-healing behavior)
- Performance acceleration via SDCache and modern I/O paths
- Broad protocol support for block, file, and object access

Why BigTwin helps here:

BigTwin's per-node NVMe layout enables consistent per-node performance symmetry—critical in clustered designs where predictability is a requirement, not a best-effort outcome. Hot-swap NVMe and node-level serviceability reduce operational friction during expansions, replacements, and maintenance.

StratiSERV (virtualization + containers)

- Live Migration and Storage Live Migration for non-disruptive maintenance and workload balancing
- High-availability orchestration to reduce downtime during node failures
- Hardware passthrough for near-native performance (GPUs, NICs) when required
- Virtual TPM, vGPU support, and confidential computing support (Intel TDX / AMD SEV)

With PCIe Gen5 expansion per node and standardized AIOM NIC options, BigTwin supports dense virtualization while keeping network and expansion strategy consistent, repeatable, and easy to scale across nodes and deployments.

HyperSERV (integrated HCI)

HyperSERV combines StratiSTOR + StratiSERV into a unified architecture for VM, container, and storage workloads on the same nodes—simplifying design, reducing footprint, and improving data locality.

SDCache Acceleration (RAM + Disk) — Performance for Every Configuration

SDCache is SteelDome's built-in acceleration layer designed to deliver exceptionally high performance across a wide range of deployment styles—capacity-first, balanced, or performance-first. It can use RAM, dedicated cache drives, or both to boost I/O performance across virtually any underlying media. This is especially valuable in cost-effective, capacity-forward configurations (often HDD-heavy), where latency and random I/O are the limiting factors. It also improves responsiveness and burst behavior in SSD and NVMe deployments by absorbing hotspots and smoothing demand.

Why SDCache matters in real deployments:

Most environments don't run a single, predictable workload. They run a mix: virtual machines, databases, file services, containers, analytics jobs, and sometimes AI pipelines. That mix is what creates performance volatility—brief spikes, contention, and small-block randomness that can dominate the user experience.
SDCache addresses that reality by providing a tunable acceleration layer that:

- Reduces latency for hot working sets and repeated access patterns
- Improves consistency during bursts (boot storms, patch cycles, batch runs)
- Increases perceived performance without requiring every drive to be premium media
- Helps maintain predictable behavior under mixed workloads and concurrency

In practical terms, SDCache can make a cost-effective disk layout behave like a much more expensive tier—without changing the cluster design or forcing a redesign.

Workloads that benefit immediately:

- VM-heavy clusters: boot storms, patch cycles, and mixed random I/O that typically stress storage latency
- Databases and metadata-heavy applications: small-block random reads/writes and high-concurrency patterns
- File services: mixed operations, latency-sensitive directory work, and multi-user concurrency
- AI and analytics pipelines: bursty staging, repeated reads, and hotspot datasets

Two acceleration paths (use one or both)

1) RAM-based SDCache (ultra-low latency)

RAM-based caching allocates a portion of system memory as an extremely fast cache tier. This is ideal for reducing latency and accelerating repeated access to hot blocks—especially metadata and small random reads that often drive perceived responsiveness.

Typical outcomes:

- Faster “first response” behavior for interactive workloads
- Improved consistency for mixed workloads with frequent reuse
- Strong gains where latency matters more than raw throughput

2) Disk-based SDCache (SSD/NVMe as a cache tier)

Disk-based caching uses designated SSD or NVMe devices to accelerate slower media, especially in HDD-heavy or mixed-media designs. It delivers high impact when capacity is the priority, but performance still needs to feel premium.

Typical outcomes:

- Dramatically improved random I/O behavior versus HDD-only designs
- Better multi-tenant behavior under concurrency
- More predictable performance during peak periods

Operational integration (how it’s enabled and managed)

SDCache is integrated into the StratiSYSTEM workflow rather than bolted on afterward. It is configured through guided tooling and validated activation steps designed to prevent misconfiguration and ensure the cache layer aligns with the physical devices present in the system.

A common approach is to:

- Define cache-capable devices and intent (RAM, disk, or both)
- Generate SDCache configuration using built-in tools
- Apply configuration with validation to ensure consistency
- Update or regenerate the configuration as devices are added over time

Net result: high performance without exotic complexity

SDCache is a major reason SteelDome can deliver consistent performance across diverse Supermicro BigTwin configurations. It provides a practical performance lever whether you are:

- Building an all-NVMe hyperconverged cluster for maximum performance
- Building a mixed-media design for balanced cost and speed
- Building a capacity-forward configuration while still expecting strong responsiveness

The key is that SDCache scales with the right combination of RAM allocation, cache device selection, and operating mode aligned to workload intent—without requiring customers' intervention.

Key Value Propositions and Differentiation

1) Cluster-first design that scales without disruption

SteelDome supports growth from single-node to multi-node deployments, with clear guidance for two-node + quorum and 3+ node HA configurations. BigTwin makes scale-out practical by enabling customers to add nodes (or chassis) while preserving consistent fault domains and per-node hardware behavior.

2) StratiSTOR clustered storage foundation for resilience and flexibility

SteelDome uses StratiSTOR as its clustered storage foundation, providing mature distributed behavior that aligns well with mixed workloads. In practical terms, this delivers resilient storage services, strong operational durability, and broad compatibility with modern infrastructure patterns—without locking customers into rigid architectural constraints.

3) Performance-focused architecture (NVMe, GPUDirect, advanced networking)

SteelDome supports high-throughput, low-latency architectures through modern I/O paths and advanced networking. BigTwin adds the hardware runway: all-NVMe front bays per node, PCIe Gen5 expansion, and AIOM networking to keep bandwidth scalable as workload demands grow.

4) Operational simplicity through integrated capabilities

SteelDome is designed to reduce day-2 operational burden through integrated capabilities spanning resilience, automation, and non-disruptive maintenance. BigTwin complements this by simplifying physical operations: hot-pluggable nodes and hot-swap NVMe reduce maintenance windows and make dense deployments serviceable without drama.

5) Validated compatibility with freedom of configuration

SteelDome and Supermicro together provide a validated path that still preserves customer choice. Organizations get vendor-grade hardware and a pre-integrated software foundation—without being locked into a single rigid configuration template.

Supported / Validated System Configurations

Supermicro SYS-221BT-HNR (BigTwin, 2U 4-node, all-NVMe)

Platform highlights (per node):

- Dual Socket E (LGA 4677) supporting 5th/4th Gen Intel Xeon Scalable processors
- Up to 16 DIMMs and up to 4 TB DDR5-5600 (1DPC)
- 6× front hot-swap 2.5" PCIe Gen5 NVMe bays per node
- AIOM networking, PCIe Gen5 expansion, and optional M.2 NVMe RAID1 boot

Supermicro SYS-222BT-HNR (BigTwin, 2U 4-node, all-NVMe; newer CPU platform)

Platform highlights (per node):

- Dual Socket E2 (LGA 4710) for Intel Xeon 6700/6500 series variants
- Up to 16 DIMMs and up to 4 TB DDR5-6400 (1DPC)
- Up to 6 front hot-swap 2.5" NVMe bays per node with flexible AIOM networking options
- Strong alignment with hyperconverged, high-performance file, and container-centric deployment styles

Recommended / Optional Configuration Guidance (Practical Sizing Patterns)

SteelDome sizing is workload-driven, with a few guiding principles that translate well to BigTwin deployments:

- Boot design: Use the optional M.2 NVMe RAID1 boot configuration to keep the front NVMe bays focused on data and performance tiers
- Network design: Plan for at least two networks (front-end/quorum and back-end replication), and select AIOM NICs that match bandwidth and feature requirements
- Storage layout: For metadata-heavy or latency-sensitive workloads, consider dedicating NVMe resources to isolate high-sensitivity activity
- Cooling/TDP headroom: Consider higher-performance cooling options when targeting sustained CPU utilization or performance-focused profiles

Deployment Overview

The deployment strategy for this solution is straightforward: start small and scale safely.

Supported cluster sizes

SteelDome supports multiple cluster footprints:

- Single node: Supported. The fault domain is disk-level; a typical baseline includes ample memory and dual 10Gb networking.
- Two nodes: Supported with an external quorum node requirement, enabling a minimum HA footprint.
- Three+ nodes: A fully functional HA clustered platform with a node-level fault domain.

BigTwin mapping

A single BigTwin chassis provides four nodes, making a 3+ node HA design achievable in one compact footprint—while keeping an additional node available for capacity headroom, failover margin, or dedicated roles.

Cluster Concepts and Redundancy (Operational Behavior)

- Node-level resiliency: Individual nodes can fail while services continue through clustered storage and compute HA mechanisms.
- Network separation: A front-end/quorum network and a back-end replication network provide predictable performance and clean fault isolation.
- Hyperconverged simplicity: Enabling StratiSERV alongside the storage cluster yields HyperSERV, reducing complexity and improving locality for many workloads.

A practical BigTwin deployment flow (typical)

1. Design + sizing workshop: Map workload mix (VM density, I/O profiles, protocol needs) to node counts, NVMe population, and NIC bandwidth.
2. Hardware staging (Supermicro BigTwin): Configure boot devices, select AIOM NICs, and populate NVMe per node. Symmetry across nodes is strongly recommended for predictable cluster behavior.
3. Cluster bring-up: Establish front-end/quorum and back-end replication networks, then form the cluster based on the chosen topology.
4. Enable virtualization/containers where needed: Activate StratiSERV to support VM and container workloads, using HA and migration features for maintenance and workload balancing.
5. Validation + baseline performance: Run a short validation using platform performance tooling and representative workloads to confirm throughput and latency targets.

Conclusion and Summary

SteelDome on Supermicro BigTwin provides a validated, high-density path to modern infrastructure—unifying storage, virtualization, and orchestration in a single platform. Designed to scale without disruptive migrations and built to support performance-intensive workloads with strong resilience, the combination of cluster-first software and cluster-friendly BigTwin hardware enables customers to deploy quickly, operate, and scale confidently.

Suggested next actions:

1. Select a starting topology (single node, 2-node + quorum, or 3+ nodes) based on availability goals and footprint constraints.
2. Choose the BigTwin model (SYS-221BT-HNR vs. SYS-222BT-HNR) aligned to CPU generation, memory targets, and expansion headroom.
3. Define a reference configuration (NVMe population strategy, AIOM NIC bandwidth/features, boot layout) and run a short proof-of-value validation.
4. Plan phased expansion by adding nodes or chassis while maintaining symmetry and predictable fault domains.

By adopting the integrated SteelDome + Supermicro solution, organizations can reduce TCO, improve performance, and achieve cost-effective outcomes across their infrastructure.

For More Information

Supermicro: www.supermicro.com
Supermicro BigTwin Systems: https://www.supermicro.com/en/products/bigtwin
SteelDome: www.steeldome.com

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements. Visit www.supermicro.com

STEELDOME

SteelDome is built on a software-first philosophy that treats infrastructure as a controlled, resilient platform rather than a collection of tools. By unifying storage, virtualization, hyperconverged systems, and data protection, SteelDome enables organizations to operate critical workloads with precision, stability, and architectural freedom.

Visit www.steeldome.com