

Asine SSD Architecture: Toward Deterministic Flash Storage

A Position Paper on Predictable Storage for Space, Defense, and Mission-Critical Systems

ASINE Group | Hod HaSharon, Israel

Abstract

Modern SSDs have transformed throughput and median latency. They remain, however, fundamentally non-deterministic: garbage-collection storms, thermal throttling, wear-leveling side effects, and opaque firmware policies produce tail latencies that are orders of magnitude above the median. For space and defense systems — where a missed downlink window, a stalled onboard AI inference cycle, or a delayed sensor log can carry mission consequences — this unpredictability is not a performance footnote. It is an architectural disqualifier.

The Asine SSD Architecture addresses this directly. Built on ASINE's ruggedized flash platforms already deployed in defense and aerospace applications, it introduces seven coordinated mechanisms — deadline-aware flash translation, reserved garbage-collection epochs, workload isolation domains, host-negotiated latency classes, admission-token flow control, predictive thermal budgeting, and cryptographically signed latency attestation — that together allow p99.999 latency to be engineered and verified rather than hoped for.

The central claim is simple: flash storage for mission-critical infrastructure must be designed around worst-case guarantees, not benchmark averages. Asine is that design.

1. The Problem: Fast Is Not the Same as Predictable

An NVMe SSD rated at 7 GB/s sequential throughput may return a write in 50 μ s one moment and stall for 200 ms the next. In consumer storage, that variance is tolerated. In the systems this paper addresses, it is not.

Consider three concrete failure modes:

- **SmallSat downlink window:** A satellite has a 90-second ground contact. Onboard data staging to flash must complete within a hard deadline. A GC storm during staging means data that will not be recovered for days — or at all.
- **Onboard AI inference:** Edge inference for autonomous systems requires bounded model-swap and result-logging latency. An unbounded storage pause breaks the control loop.
- **Distributed training checkpoint:** A storage stall during checkpoint blocks hundreds of GPUs simultaneously. The cost is proportional to cluster size.

The underlying architecture is unchanged across all three cases: SSDs are engineered for statistical performance — optimized for throughput and average latency — while mission-critical infrastructure requires temporal contracts. The Asine architecture closes that gap.

2. Why Conventional SSDs Cannot Be Made Deterministic

Four structural causes produce non-determinism in commodity and even enterprise-grade SSDs:

- **FTL entropy.** NAND cannot overwrite in place. Logical-to-physical remapping accumulates relocation debt that must be paid via garbage collection. GC timing is workload-dependent; identical host writes do not imply identical response times.
- **Internal resource contention.** DRAM metadata caches, NAND channels and dies, ECC engines, and wear-leveling logic all share internal bandwidth. Contention is neither exposed nor controllable at the host interface.
- **Opaque firmware policies.** Hosts cannot observe free-block pressure, GC urgency, wear variance, thermal headroom, or internal queue state. Cross-layer scheduling is impossible when the host is blind to the conditions that produce latency spikes.
- **Benchmark culture.** A drive where 99.99 % of requests complete in 70 μ s and 0.01 % complete in 120 ms scores well in most review dashboards. For deterministic systems, that drive is broken. What matters is the shape and stability of the tail — not the median.

3. Determinism as a Design Objective

Asine defines a deterministic storage class C_i as one satisfying:

$$\Pr(L_i \leq B_i) \geq 1 - \epsilon$$

where L_i is observed latency, B_i is a declared bound, and ϵ is a violation probability target on the order of 10^{-9} . The architecture targets strong determinism: bounds must hold under sustained mixed workloads, near-full capacity, thermal stress, multi-tenant contention, post-power-loss recovery, and aged NAND conditions. Idle-state determinism is not the engineering challenge. The challenge is guaranteeing bounds when the system is doing the hard work it was built for.

4. The Seven Mechanisms

The Asine architecture achieves determinism through seven coordinated mechanisms. Each addresses a specific failure mode identified in Section 2.

#	Mechanism	What it guarantees
1	Deadline-aware FTL	Every I/O carries a deadline and priority class; mapping decisions optimize schedulability, not average write amplification.
2	Reserved relocation epochs	GC runs in pre-declared windows (e.g., 2 ms every 500 ms). The host knows the schedule; surprise storms are eliminated.
3	Capacity reservation zones	~70 % deterministic / ~20 % relocation reserve / ~10 % recovery pool. Users trade capacity for tighter bounds explicitly.
4	Isolation domains	Namespaces map to physical channel groups. A noisy tenant cannot contaminate another tenant's latency profile.
5	Thermal-budget scheduling	Predictive rate-shaping replaces reactive throttling; thermal budget is a schedulable resource, not a last-minute cut.
6	State telemetry export	Relocation debt, thermal headroom, wear variance, and bound-confidence scores are exported in real time for cluster schedulers.
7	Cryptographic attestation	Signed telemetry provides verifiable p99.999 distributions, GC schedule integrity, and throttle-event logs.

5. Host–Device Contract Layer

5.1 Latency Classes

Applications declare a latency class at I/O submission time. The device enforces it at the controller — not simulated in host software. For space and defense platforms, L0 is the relevant class: it guarantees read-dominated paths within 50 μs , covering onboard inference data access and downlink-window staging triggers.

Class	Bound	Intended use
L0	$\leq 50 \mu\text{s}$	Read-dominated, time-critical paths (e.g., onboard AI inference, downlink-window data staging)
L1	$\leq 200 \mu\text{s}$	Mixed I/O — command handling, telemetry logging
L2	Soft / throughput	Bulk transfer, image downlink queues
L3	Best effort	Archival, background housekeeping

5.2 Admission Tokens

The device issues admission tokens representing bounded write capacity under current conditions. When tokens are exhausted, the host delays I/O at a visible orchestration layer rather than allowing unbounded queues to accumulate inside the drive. This makes backpressure explicit and manageable — the failure mode is a visible queue, not a silent stall.

6. Scheduling Model

I/O classes are modeled as sporadic real-time tasks. The system is schedulable if total utilization — including relocation, ECC, and metadata overhead — remains within a statically budgeted bound:

$$U_{\text{total}} = U_{\text{io}} + U_{\text{d}^c} + U_{\text{e}^c} + U_{\text{meta}}$$

All four components are statically budgeted. Deterministic I/O classes are never scheduled onto an over-committed substrate. The controller microarchitecture enforces lane separation between real-time, standard, and maintenance workloads in silicon — not through firmware heuristics that can be overridden under load.

7. Reliability Without Latency Collapse

The assumption that strong reliability requires unpredictable pauses is false. Reliability mechanisms are redesigned to be incremental and schedulable:

- Incremental metadata journaling, not large blocking checkpoints
- Parallel parity refresh spread across low-utilization intervals
- Opportunistic scrubbing tuned to thermal and utilization envelopes
- Reserved recovery pools allocated in advance, not on demand

Post-power-loss recovery targets a defined bound (< 250 ms). Post-boot behavior is constrained to avoid the minutes of latency chaos that follow power events in conventional devices — a failure mode with direct mission consequences in space and defense deployments.

8. Deployment Contexts

The mechanisms described above are not context-neutral. The following table maps deployment environments to their specific determinism requirements:

Deployment context	Storage determinism requirement
Space / SmallSat	Downlink window data staging, onboard AI inference timing, mission-critical command logging
Defense / Edge	Real-time sensor fusion, UAV control loops, field-deployed historian logging
AI Training Clusters	Checkpoint determinism — prevent GPU idle cascades during multi-node training runs
Financial Systems	Bounded persistence windows for ordering guarantees and regulatory auditability
Industrial Automation	Historian logging and control loops operating without storage-induced pauses

9. Benchmarking Deterministic Storage

Standard benchmarks — fio averages, peak MB/s, headline IOPS — are insufficient for evaluating deterministic storage. The relevant metrics are:

- Latency percentiles: p99, p99.9, p99.99, p99.999
- Maximum observed latency over multi-hour windows
- Burst recovery time after load spikes
- Bound violation count per billion I/Os
- Aged-drive behavior after multi-PB write endurance

The Determinism Score (DS) summarizes tail behavior:

$$DS = 1 / (1 + \sigma_{tail} + V + R)$$

where σ_{tail} is tail-latency variance, V is normalized violation rate, and R is recovery instability after perturbation. Higher scores correspond to tighter, more stable tails. A high-throughput drive with poor tail stability scores low. Asine is optimized for DS, not for the benchmark averages that have historically defined the field.

10. Future Work

- Deterministic QLC via multi-tier write staging — extending guarantees to higher-density NAND
- Formal verification of FTL scheduler correctness and latency-contract satisfaction
- Cross-rack deterministic storage fabrics for distributed mission-critical infrastructure

Conclusion

The question that matters

The field has long asked: "How fast is your drive?" For space, defense, and mission-critical infrastructure, the relevant question is different: "Can I trust it at the exact moment I need it?" The Asine SSD Architecture is designed to answer yes — not as a benchmark claim, but as a verifiable, cryptographically attestable property of the device.

By exposing internal state, budgeting maintenance work in scheduled epochs, isolating workloads to physical channel groups, shaping thermals predictively, and enforcing deadline-aware scheduling at the controller, Asine transforms NAND flash from probabilistic accelerator into time-bounded storage infrastructure.

The transition from fast to certain is not a firmware parameter adjustment. It is an architectural commitment — one that ASINE's ruggedized flash platforms are uniquely positioned to deliver.

Executive thesis: Asine SSD Architecture is the transition of flash storage from high-speed best-effort media into mathematically bounded infrastructure.