

/ ENTERPRISE FILE INFRASTRUCTURE / WHITE PAPER

Beyond the Refresh: Why the NAS Model is Breaking for Good

Why hardware disruption, storage inflation, and AI competition are forcing a permanent rethink of enterprise file storage.

May 2026



Contents

/ 00 EXECUTIVE SUMMARY	3
/ 01 THE MARKET	4
Force 1: The AI Hardware Reallocation	4
Force 2: Price Inflation Across the NAS Stack	4
Force 3: Exploding Unstructured Data	5
/ 02 THE FINANCIALS	6
Hard Costs: The Visible Tier	6
Soft Costs: The Hidden Multiplier	6
The 5-Year TCO Reality	7
/ 03 MOVING WITHOUT THE RISK	8
How Migration Actually Works	8
Where to Start	8
No Sequencing Problem	11
/ 04 THE NASUNI FILE DATA PLATFORM	9
UniFS®: One Global File System	9
Built-In Data Protection: Backup and DR	9
Operational Simplicity: The Management Dividend	10
AI Readiness: The Strategic Foundation	10
/ 05 THE OUTCOME: MANAGE, PROTECT, ACTIVATE	11
What Leaders Should Do Next	11
About Nasuni	12



/ EXECUTIVE SUMMARY

Why predictability broke — and why the old model can't come back.

Enterprise storage infrastructure has entered a period of structural disruption. This white paper examines why the forces reshaping the NAS market are permanent, not cyclical, and makes the financial and operational case for a new model.

63%

**5-YEAR TCO REDUCTION
VS. ON-PREMISES**

75%

**REDUCTION IN IT
ADMINISTRATION TIME**

<15M

**DISASTER RECOVERY
VS. HOURS ON LEGACY**

\$7.4M

**SAVINGS FOR 500TB
ACROSS 24 LOCATIONS**

/ EXECUTIVE SUMMARY

For decades, enterprise file storage followed a reassuringly predictable rhythm: a hardware refresh every three to five years, a budget cycle that accommodated it, and procurement timelines that could be planned. That model is now under structural strain from multiple directions simultaneously. This is not a cyclical blip that patient organizations can wait out. The forces reshaping the NAS market today are architectural and enduring.

AI hyperscalers are consuming storage supply at unprecedented scale. NAND flash capacity is already allocated well into 2027 by major suppliers. Enterprise hard drive prices rose an average of 46% in just four months leading into early 2026. (Source: Tom's Hardware, January 2026) Quote validity windows have collapsed from weeks to days. And unstructured data is growing at 55 to 65% annually, compounding the capacity pressure regardless of market conditions.

At the same time, IT leaders are being asked to do more with less. AI initiatives are competing directly for the same capital budgets earmarked for infrastructure refresh. The result is a convergence of pressures that has elevated what used to be a routine procurement decision to board-level strategic attention.

This white paper makes the case that the response to this moment should not be another hardware purchase. It should be a permanent rethinking of how enterprise file data is managed, protected, and made available, one that eliminates refresh dependency entirely and builds a foundation capable of supporting whatever the organization needs next.

The perfect storm: how four forces converged to break enterprise storage planning.

/ FORCE 1: THE AI HARDWARE REALLOCATION

The root cause of today's storage supply disruption is structural, not cyclical. The explosive growth of AI infrastructure, particularly the demand for high-bandwidth memory to serve GPU clusters, has fundamentally redirected the semiconductor manufacturing capacity that enterprise NAS hardware has relied on for decades. High-bandwidth memory requires approximately four times the wafer capacity per gigabyte compared to conventional DRAM. As AI hyperscalers lock in multi-year supply agreements, the economics for memory vendors are unambiguous: prioritize the higher-margin products.

NAND suppliers have reported their entire 2026 production capacity is already sold out, with analysts signaling limited relief through 2027. (Source: Phison CEO Khein-Seng Pua / Tom's Hardware, November 2025) Enterprise SSD revenues grew 28% quarter-over-quarter in Q3 2025 alone, not because supply expanded, but because hyperscalers were scrambling to stockpile inventory. (Source: TrendForce, December 2025) Phison reported average lead times of 224 days through Q3 2025, with the CEO confirming all 2026 NAND capacity was sold out. (Source: Tom's Hardware, November 2025) HDD lead times at major suppliers extended beyond two years as a result. (Source: Digitimes / Tom's Hardware, November 2025)

"I believe there's a bit of a structural change to this. It's not really a cycle you can just wait out. This is an inflection point in the whole storage refresh cycle."

JERRY CARTER
Chief Technology Officer,
Nasuni

/ FORCE 2: PRICE INFLATION ACROSS THE NAS STACK

Enterprise hard drive prices rose an average of 46% in the four months preceding January 2026, with individual models showing increases of 23% to 66%. (Source: Tom's Hardware, January 2026) Seagate's enterprise lines showed 12 to 18% increases in Q1 2026 alone. These are not transient adjustments. Cloud providers committed more than \$405 billion in capital expenditure in 2025 alone (a 62% year-over-year increase) with storage infrastructure accounting for a significant and growing share, effectively outbidding enterprise buyers for priority allocation. (Source: IO Fund / Big Tech earnings, 2025) When hardware estimates arrive at 125% to 200% of expected figures, the ripple effect across backup, DR, replication, and migration budgets is substantial.



The perfect storm: how four forces converged to break enterprise storage planning.

/ FORCE 3: EXPLODING UNSTRUCTURED DATA

Unstructured data now accounts for 80 to 90% of all new enterprise data and is growing at 55 to 65% annually. (Source: Gartner) By 2026, 74% of enterprises are storing more than 5 petabytes of unstructured data, a 57% increase over 2024. (Source: Komprise, 2026 State of Unstructured Data Management) This growth does not pause during hardware procurement difficulties. The gap between available hardware and data growth is widening in both directions simultaneously.

/ FORCE 4: AI BUDGET COMPETITION

AI initiatives are competing directly for the same capital budgets earmarked for infrastructure refresh. IT teams that are already constrained are being asked to both maintain legacy infrastructure and build the data foundation for AI-driven workflows. The irony: the file data NAS serves is precisely what AI workflows need. Organizations that cannot stabilize their file foundation discover they cannot advance their AI strategy either.

“What used to be pretty routine planning around infrastructure cycles has now created specific planning risk and gotten all the way up to board-level executive attention.”

NICK BURLING

Chief Product Officer, Nasuni



The hidden cost of doing nothing: why the refresh trap costs more than the hardware.

/ HARD COSTS: THE VISIBLE TIER

Legacy file infrastructure is a composite system, not a single purchase. Every organization maintaining on-premises NAS is simultaneously maintaining several interrelated capital investments:

- Production file storage hardware — NAS systems and file servers refreshed every 3 to 5 years per location, with multiple capital events per decade for larger organizations.
- Backup infrastructure — separate software licensing, media servers, storage media, and operational overhead to manage schedules, test restores, and coordinate retention policies.
- Disaster recovery infrastructure — dedicated DR sites, standby NAS hardware, ongoing testing cadences, and coordination overhead across changing application landscapes.
- Multi-site data transfer — MPLS circuits, WAN acceleration appliances, file transfer software, and replication tooling for every location.
- Data center facilities — electricity, cooling, physical security, and space that scales with every hardware addition.
- Maintenance and support escalation — vendor support contracts on aging NAS hardware typically increase 8 to 15% annually, and out-of-cycle support surcharges are increasingly common on supply-constrained components. Each refresh resets this clock; each year without a refresh compounds it.
- Refresh cycle compounding — each successive hardware refresh is more expensive than the last, because data volumes grow between cycles and the infrastructure required to hold them expands accordingly. Mid-cycle capacity additions are frequently the worst case: they often trigger broader upgrades to controllers, networking, or software licensing that were not in the original budget.

/ SOFT COSTS: THE HIDDEN MULTIPLIER

IT administration overhead is where the 75% reduction Nasuni customers report actually comes from. It is not a rounding-up of vague efficiency gains but the elimination of specific recurring work: monitoring utilization across locations and raising capacity alerts before systems degrade; scheduling, verifying, and troubleshooting backup jobs; managing tape rotation and offsite retention; maintaining DR site readiness and running failover tests; coordinating vendors across NAS, backup software, replication tools, and WAN acceleration; executing data migrations during every hardware refresh. On Nasuni, backup is continuous and automatic. DR is a console action. Capacity is unlimited. There are no tapes, no replication jobs, no DR site to staff, and no migration to plan.

“I don’t have to buy large storage devices anymore because our data is in cloud storage. I’ll add capacity annually, depending on our needs, but I’m not going out spending half a million dollars on a new device every couple of years.”

– Stephen Held, Vice President and Chief Information Officer, Leo A Daly

Business productivity losses represent a second category easier to quantify than most teams realize. Take annual revenue and divide by 250 working days. That daily figure is your exposure for any event that meaningfully interrupts file access: a ransomware incident, a DR failover, a site that goes down because NAS hardware failed during a supply delay. Legacy infrastructure creates this exposure structurally: recovery is measured in hours or days, cross-site transfers degrade under load, and systems approaching 90% utilization throttle before IT can respond. Nasuni eliminates each of those conditions architecturally, not by working around them.

People risk is the third category that rarely appears in a formal budget line but shows up in every engineering org that has spent years running legacy NAS. Senior infrastructure engineers who spend the majority of their cycles on storage operations (capacity planning, vendor escalations, refresh execution, tape management) are not doing architecture work, security hardening, or AI enablement. That is an opportunity cost measured in the projects that do not get started. It is also a retention problem: the engineers capable enough to lead infrastructure transformation are often capable enough to find roles where they are building, not maintaining.

/ THE 5-YEAR TCO REALITY

The following comparison is drawn from Nasuni’s Business Value Assessment methodology, based on actual customer data. The example reflects a large, globally distributed organization, consistent with the scale and profile of Nasuni’s enterprise customer base.

Cost Category	Legacy On-Premises	Nasuni Cloud File Storage
Primary file storage hardware	Full refresh every 3–5 years per location; multiple capital events per decade	Eliminated — replaced by cloud object storage subscription
Backup infrastructure	Separate licensing, media servers, storage media, operational overhead	Included — Continuous File Versioning built in at no additional cost
Disaster recovery	Dedicated DR site, standby NAS hardware, ongoing test cycles	Included — edge appliance provisioned in under 15 minutes
Multi-site data transfer	MPLS, WAN acceleration, replication tooling per site	Eliminated — unified global namespace, no WAN infrastructure needed
IT administration	Full allocation: backup, DR, multi-site, patch, vendor management	~75% reduction via Nasuni Management Console
5-year TCO (500TB, 24 locations)	\$11.7 million	\$4.3 million — 63% reduction, \$7.4M saved

“Our financial people were particularly happy to see us move from a large cyclical CapEx model to a predictable month-to-month operational model. It was one of the few things I’ve done in thirty-four years that actually made the CFO smile.”

– Glen Ridnour, Chief Technology Officer, Huitt-Zollars

Moving without the risk — why the blocker isn't what you think.

/ HOW MIGRATION ACTUALLY WORKS

The reason most organizations delay moving off legacy NAS is not the financial case. The TCO math is clear once the full cost picture is assembled. The blocker is almost always a version of the same fear: that migration is a high-risk event requiring a cutover window, a period of downtime, and a moment when the organization is exposed. That fear is understandable. It is also the single most expensive misconception in enterprise file storage, because it causes organizations to absorb another refresh cycle, with all its cost, disruption, and risk, rather than take a path that carries none of those consequences. Nasuni's architecture was built to make the forklift assumption false by design:

- Nasuni Edge Appliances are lightweight virtual machines that deploy in any on-premises location, branch office, or cloud data center, typically in hours, without disrupting existing infrastructure.
- Active files are cached locally at LAN speed. More than 98% of all I/O requests are served from local cache, providing the performance that engineering and design workflows require.
- The authoritative "gold" copy of all file data is stored in cloud object storage (Azure Blob, Amazon S3, or Google Cloud Storage) with continuous synchronization via the UniFS® global file system.
- File paths, permissions, Active Directory integrations, and NTFS/SMB access controls are preserved throughout. Users see no change in how they access files.

/ WHERE TO START

The practical answer to where to start is that the starting workload matters less than organizations assume. The transition mechanics are the same regardless: Nasuni deploys alongside existing infrastructure, data moves to cloud object storage in the background, and users access files through the same paths and permissions they use today. They do not experience a migration. Organizations with an active pain point (a DR deployment that has never been tested, a branch running out of capacity, a backup process requiring manual intervention) solve that problem first. Those that prefer a lower-profile start displace backup or archive storage first. In either case, no hardware purchase, change freeze, or migration weekend is required. The organization waiting to complete a hardware refresh before evaluating alternatives is choosing the most disruptive and expensive option on the table.

/ NO SEQUENCING PROBLEM

A common concern in cloud migration is sequencing: move data first, and applications left on-premises suffer latency reaching back across the WAN to access it. Move applications first, and they are disconnected from their data. Nasuni's hybrid architecture makes the question irrelevant. Data is already in cloud object storage and cached locally on-premises. Applications in the on-premises environment read from local cache at LAN speed, while cloud-based workloads access the same authoritative data directly. There is no sequencing problem, no migration window, and no period where applications are reaching across a narrow link to data that hasn't caught up yet. The architecture resolves both access patterns simultaneously from day one. Nasuni's implementation methodology includes structured user acceptance testing and engagement planning, so the transition is operationally supported end-to-end, not just architecturally non-disruptive but actively managed.

"We know that we can scale and grow without investing additional on-prem storage. At the end of the day, the users did not notice any issue when transitioning from the traditional Windows file shares."

DAVE DEPILLIS

Director of IT, Cole,
Scott & Kissane

The Nasuni File Data Platform: one platform replacing three infrastructure layers.

Nasuni is not a cloud storage bucket, a data lake, or a backup appliance. It is an enterprise file data platform: a unified system that replaces NAS hardware, backup infrastructure, and disaster recovery through a single cloud-native architecture, while delivering the performance, governance, and operational characteristics enterprise file access requires.

/ UNIFS®: ONE GLOBAL FILE SYSTEM

The foundation of the Nasuni platform is UniFS®, a patented global file system that stores a single authoritative copy of all file data in cloud object storage and presents it to users through Nasuni Edge Appliances as a familiar SMB/NFS file share. UniFS provides capabilities that distributed NAS architectures fundamentally cannot: a single authoritative file version accessible from any location simultaneously; a global namespace spanning every location at LAN speed; infinite scalability on demand; and built-in deduplication and compression that reduce cloud storage consumption by an average of 40% (Source: Nasuni Business Value Assessment).

/ BUILT-IN DATA PROTECTION: BACKUP AND DR

Nasuni Continuous File Versioning® creates immutable snapshots of all file data at one-minute granularity, stored automatically in cloud object storage. This is a core platform capability at no additional licensing cost. It eliminates separate backup software, media servers, and storage media entirely.

The immutability of these snapshots is critical for ransomware resilience. Nasuni's cloud-stored snapshots cannot be encrypted or deleted by ransomware even if the on-premises environment is compromised. A Nasuni Edge Appliance can be provisioned and made accessible in an alternative location in under 15 minutes, eliminating dedicated DR sites entirely.

"We eliminated 500 file servers and all the backup licensing that went with them. It's something we've been wanting to do for twenty years."

CHRIS MOORHEAD, Director of IT Systems and Engineering, Tetra Tech

/ OPERATIONAL SIMPLICITY: THE MANAGEMENT DIVIDEND

The Nasuni Management Console provides centralized visibility and control across every edge appliance, volume, share, and user, regardless of how many locations the organization operates. This single interface is what enables the 75% reduction in file storage administration time Nasuni customers consistently report. Instead of managing dozens of individual NAS systems each with their own backup schedules, capacity thresholds, and vendor contacts, IT teams manage one platform.

/ AI READINESS: THE STRATEGIC FOUNDATION

The file data that NAS serves (documents, designs, CAD files, project records, engineering data) is precisely the unstructured data that AI applications need to deliver business value. Nasuni's architecture satisfies the foundational requirements AI workflows depend on: a single authoritative file version with full revision history, permissions enforced as a global file system property rather than a per-location configuration, and the same governance model that controls human access to file data.

As Nick Burling, Nasuni's Chief Product Officer, has framed it: **"How do we help you turn data from being a cost to be managed to an asset to be leveraged?"** That question is at the heart of what a cloud-native file foundation enables, and the framing that resonates most with the executive stakeholders now involved in these decisions.

In practice, most organizations discover that their file data is not AI-ready. Not because the data doesn't exist, but because of how it is stored and managed. AI workflows require data that is current, consistently governed, and grounded in a single source of truth. Legacy NAS environments typically produce the opposite: multiple copies of the same files across locations, no authoritative version, and permissions that vary by site. The infrastructure problems most organizations have to solve before AI workflows become viable (fragmented data, inconsistent governance, duplicated datasets) are resolved by the act of moving to Nasuni, not by additional work after the fact. UniFS maintains a single authoritative version with full revision history. Permissions are a global property, not a per-location configuration. The same governance model that controls human access governs how the data is exposed to any downstream system.

For organizations that are planning AI initiatives (retrieval-augmented generation, document intelligence, engineering data search, or any workflow that requires accessing enterprise file content), the file infrastructure decision and the AI readiness decision are the same decision. Organizations that modernize their file foundation as a cost and resilience play discover, without additional work, that they have also built the data layer their AI strategy requires. This same architectural foundation is what enables Nasuni's emerging AI Activate capability, using Model Context Protocol to extend permission-aware access to AI agents and large language models within the same governance model that controls human access.



"How do we help you turn data from being a cost to be managed to an asset to be leveraged?"

NICK BURLING,
Chief Product Officer, Nasuni

The strategic payoff: Manage, Protect, Activate.

What does day-to-day operations actually look like for an infrastructure team that has made this transition? Not in abstract terms, but in terms of what changes on a normal Tuesday, and what changes when something goes wrong. On a normal day: there is no storage capacity alert to triage, because capacity does not have a ceiling. There is no backup job to verify, because Continuous File Versioning is running automatically at one-minute granularity across every location. There is no replication lag to investigate, because every location reads from the same global namespace.

The Nasuni Management Console shows the health of every edge appliance, every volume, and every share in a single interface. An engineer who would have spent the first two hours of the morning on storage operations spends those two hours on something else: security work, automation, architecture, AI enablement. During a ransomware incident: the response is not a multi-day recovery operation coordinating backup vendors, tape retrieval, and system rebuilds. It is a console action. The most recent clean snapshot is identified, revert is initiated, and operations resume. The immutable cloud-stored snapshots are untouched because ransomware on the on-premises environment cannot reach object storage. For organizations on legacy infrastructure, this same scenario means days of downtime, potential data loss back to the last successful backup, and a recovery process that assumes all the infrastructure is still intact, an assumption ransomware operators specifically target. When a site goes offline: a replacement edge appliance is provisioned and operational in under 15 minutes.

On legacy NAS, the equivalent event means sourcing hardware in a supply environment where lead times are measured in months, rebuilding configurations, and restoring data from backup: a process measured in days to weeks. The difference is not incremental. It is architectural. At budget time: infrastructure spend is a predictable monthly subscription line, not a capital event that competes with AI projects, security investments, and digital transformation initiatives. The outcomes below describe this shift at a strategic level:

Strategic Pillar	What It Eliminates	What It Delivers
Manage	Hardware refresh cycles, CapEx volatility, capacity anxiety, multi-vendor complexity, fragmented visibility	Predictable OpEx, centralized management, unlimited scale, up to 60% total cost reduction, IT time reclaimed
Protect	Ransomware exposure, DR hardware dependency, backup software sprawl, supply-chain-dependent resilience, recovery measured in days	Immutable snapshots, built-in ransomware recovery, <15-min DR, permission governance, 1-minute RPO
Activate	Siloed data AI can't reach, duplicated datasets, broken pipelines during refresh events, governance drift	AI agents access the same file data humans access, with the same governance, permissions, and version history, and no separate pipelines, duplicated datasets, or additional infrastructure required

/ WHAT LEADERS SHOULD DO NEXT

If you have read this far, the conceptual case is probably not what is holding the decision back. What holds technical buyers at this stage is typically one of three things: uncertainty about what a first engagement actually looks like, not having the business case structured in terms that move internal stakeholders, or not having heard from a peer who has been through it. All three are solvable, and the path forward on each is direct.

- **Request a no-cost Business Value Assessment.** Nasuni's BVA is a structured analysis, not a sales call. It takes your actual environment (number of locations, total data under management, current infrastructure stack, existing support and maintenance contracts) and produces a specific, defensible 5-year TCO model comparing your current trajectory against a Nasuni deployment. The output is the document that moves internal budget conversations, because it answers the questions CFOs and CIOs actually ask: what does this cost today, what does it cost on Nasuni, and what is the risk exposure of doing nothing. Available at no cost at nasuni.com/contact.
- **Frame the internal conversation around operational risk, not storage specs.** The stakeholders who need to approve this decision are not evaluating architecture. They are evaluating exposure. The questions that move those conversations: What is the financial exposure if a hardware order is delayed three months during a capacity crisis? What does a ransomware recovery look like today versus on immutable snapshots? What AI or digital transformation initiative is blocked because the data infrastructure is absorbing the budget and the IT team's attention? These are not rhetorical. They have specific answers, and the BVA provides the numbers to back them.
- **Talk to a customer who has been through it.** Customers like Glen Ridnour at Huitt-Zollars, Stephen Held at Leo A Daly, Chris Moorhead at Tetra Tech, and Dave DePillis at Cole, Scott & Kissane represent a profile common across Nasuni's customer base in engineering, AEC, manufacturing, energy, and media organizations managing distributed infrastructure across dozens of locations. There are infrastructure leads at organizations like yours who made this decision and can describe exactly what happened: what the migration felt like, what changed operationally, and what the environment looks like now. Nasuni can connect you directly. That conversation typically accomplishes more than any analysis.

"Don't be afraid to step back and reevaluate. Define your success in terms of business goals. Talk about the reduced risk of downtime or data loss, and the periodic disruption you're not going to have anymore because you're just going to add another increment of data and move on down the road."

GLEN RIDNOUR, Chief Technology Officer, Huitt-Zollars

/ ABOUT NASUNI

Nasuni is the unstructured data foundation for enterprise teams and the AI that supports them. We manage, protect, and activate unstructured data so enterprise teams and the AI that supports them can do their best work. Built on the patented UniFS® global file system, the Nasuni File Data Platform replaces legacy NAS hardware, backup infrastructure, and disaster recovery with a single cloud-native solution delivering LAN-speed performance, infinite scalability, built-in ransomware protection, and sub-15-minute disaster recovery, across any number of locations, on any major cloud provider. Nasuni serves organizations across more than 70 countries. Headquarters: Boston, Massachusetts.

Learn more: nasuni.com | Request a no-cost Business Value Assessment: nasuni.com/contact

Nasuni, UniFS, and Continuous File Versioning are registered trademarks of Nasuni Corporation. All statistics and market data are sourced from publicly available analyst reports, industry surveys, and Nasuni's proprietary Business Value Assessment data. Customer quotes are reproduced from a publicly available Nasuni webinar. © 2026 Nasuni Corporation. All rights reserved.

Let's talk

See how much your organization could save with Nasuni. Generate a personalized savings report in minutes, or connect with our experts for a custom assessment tailored to your environment, your data profile, and your goals.

[LEARN MORE](#)



Nasuni is the unstructured data foundation for enterprise teams and the AI that supports them. Built on the patented UniFS® global file system, the Nasuni File Data Platform delivers LAN-speed performance, infinite scalability, built-in ransomware protection, and sub-15-minute disaster recovery across any number of locations on any major cloud provider. Nasuni delivers services to organizations in more than 70 countries.