

Technical Leader's Guide to Evaluating Data Platforms

Five criteria that separate storage vendors from true data platforms, and how Everpure answers each one

Who this guide is for

This guide is designed for technical leaders evaluating storage and data management platforms for enterprise environments. It is structured around the five questions that matter most and that vendors are least prepared to answer with specificity.

Use each criterion as a structured line of inquiry in your vendor conversations. The “What to measure” sections give you concrete questions to bring into every evaluation. The “Everpure” response” sections show you what a genuine answer looks like so you can judge the quality of what you hear from others.

1 Total cost predictability

Why it matters

For most organizations, storage is one of the top three IT budget lines and one of the least predictable. Traditional infrastructure models are built around a refresh cycle that forces capital expenditure every three to five years, often with significant cost overruns driven by unplanned capacity growth, emergency performance upgrades, and migration projects that consume months of engineering time.

Beyond hardware, the hidden costs compound quickly: software licensing tied to capacity tiers, support contracts that escalate annually, and the operational overhead of managing increasingly complex, multivendor environments. Leaders who evaluate platforms only on day-one acquisition cost routinely find themselves with a budget problem by year two.

The evaluation question isn't what the platform costs to buy. It's what it costs to operate and whether that number is knowable in advance.

What to measure

Build a five-year TCO model that includes hardware acquisition or subscription cost, software and licensing fees, support and maintenance contracts, staff time allocated to platform management, cloud egress and data movement costs, and unplanned refresh or upgrade costs from prior cycles.

- What is the total five-year cost, including support, licensing, and upgrades?
- Is performance guaranteed at a fixed price, or does it require additional investment to maintain?
- What happens to the cost when capacity grows by two times? By five times?
- What migration or transition costs are embedded in your current vendor relationship?
- Can the vendor provide a fixed, all-in number for a defined performance tier or only a price list?

Everpure response

Most vendors sell you infrastructure and then bill you for the privilege of keeping it running. Refresh cycles, capacity upgrades, and support escalations are revenue events for them and cost surprises for you. The model is designed to create dependency, not partnership.

Everpure flips this. Evergreen//One™ is a subscription that covers hardware, software, support, and guaranteed performance SLAs in a single, predictable payment. There are no forced refresh cycles, surprise upgrade fees, or penalty for growth. When Everpure introduces new technology, existing Evergreen//One customers receive it as part of their subscription, not as an additional line item.

The financial accountability goes further. Everpure backs its availability guarantee with SLA credits—financial consequences for Everpure if commitments aren't met. That alignment of incentives is rare in this market, and it's a meaningful signal about how Everpure views the customer relationship.

2 Performance at scale

Why it matters

Performance benchmarks are easy to produce under controlled conditions. What matters is how a platform performs under the messy reality of production workloads—mixed read/write patterns, concurrent access from dozens of applications, AI pipelines running alongside transactional databases, and backup jobs competing for the same resources.

The more consequential question is how performance changes as capacity fills. Many platforms are architected to deliver strong performance at low utilization and degrade as you approach capacity. For organizations running AI, analytics, or real-time data products, that degradation isn't an abstraction. It shows up as missed SLAs, slower model training, and frustrated business stakeholders.

Technical leaders evaluating platforms need to move beyond spec sheets and test under conditions that reflect their actual environment, not a vendor's demo environment.

What to measure

Design a benchmark that reflects your real workload mix. Test at 50%, 80%, and 95% capacity. Measure latency, IOPS, and throughput at each level under concurrent workload pressure.

- What is the guaranteed latency SLA, and under what conditions does it apply?
- How does the architecture handle workload contention between high-priority and background tasks?
- What is the performance profile at 90%+ capacity utilization?
- How long does it take to provision new capacity, and is there a performance impact during provisioning?
- Can the vendor demonstrate performance on your workloads, not a vendor-supplied benchmark?

Everpure response

Legacy storage architectures make a trade-off: they are optimized for cost at the expense of consistent performance. Hybrid arrays, tiered storage, and spinning disk all introduce latency variability that is tolerable for archive workloads and catastrophic for AI pipelines. Most vendors manage this by routing your most sensitive workloads to the most expensive tier and charging accordingly.

Everpure is built on Everpure all-flash architecture, which delivers submillisecond latency consistently not just at low utilization, but under full production load across mixed workload types. There is no tiering tax or performance cliff at 90% capacity. Quality-of-service controls protect critical applications from resource contention without requiring manual tuning or a separate infrastructure.

Everpure also offers the Everpure//EX evaluation program, a structured proof of concept that runs Everpure in your environment against your workloads before you make any commitment. This is the opposite of a benchmark; it is evidence from your own data center. We offer it because we are confident in what you will find.

3 Data intelligence and context

Why it matters

The gap between a storage platform and a data platform comes down to one question: Does the system know what the data means? Legacy storage treats every byte as equivalent—a backup file, customer record, and model training dataset are all just blocks on a volume. That design made sense when storage was a utility. It no longer makes sense when data is a strategic asset.

Modern data environments require platforms that understand data context natively: classification by type, sensitivity, and regulatory category; lineage that tracks data from source to consumption; workflow awareness that knows which pipelines depend on which datasets; and semantic access that allows teams to query across sources without brittle extract, transform, and load (ETL) pipelines or manual tagging.

Without this intelligence layer, data governance is manual, AI initiatives are bottlenecked by data preparation, and compliance is a perpetual fire drill. The platform that understands your data is the platform that accelerates everything built on top of it.

What to measure

Ask vendors to demonstrate classification, lineage, and workflow orchestration as native capabilities, not add-on modules or third-party integrations. The depth of integration matters as much as the feature list.

- Can the platform automatically classify data based on content, metadata, and policy without manual tagging?
- Does it track data lineage end to end, from ingestion through transformation to consumption?
- How does it surface workflow dependencies?
Can it detect that an upstream issue blocks a downstream pipeline?
- What is the query model for accessing data across heterogeneous sources?
- How does the platform enforce data policy?
At the storage layer, the access layer, or both?

Everpure response

Every other platform on your shortlist manages storage. Everpure manages data. That distinction is not marketing language—it is an architectural difference with measurable consequences for your AI programs, your compliance posture, and the productivity of every team that depends on data infrastructure.

The Everpure semantics layer provides native data intelligence: classification, lineage, and semantic context that is embedded in the platform, not layered on through third-party integrations. This means your governance policies travel with your data. It means your AI teams spend less time on data preparation and more time on model development. And it means your compliance team can answer audit questions in hours, not weeks.

The operational context this creates is what separates Everpure from every legacy storage vendor. When your platform understands what data is, not just where it lives, you stop managing infrastructure and start accelerating the business that depends on it.

4 Resilience and recovery

Why it matters

The threat landscape has changed permanently. Ransomware is no longer an edge case; it is a mainstream operational risk, and sophisticated attackers now specifically target backup infrastructure before launching an encryption attack. Traditional recovery models were not designed for this scenario.

At the same time, business tolerance for downtime has collapsed. Applications that were once noncritical are now woven into revenue-generating workflows. A recovery time objective (RTO) measured in hours may trigger SLA penalties, customer churn, or regulatory consequences that dwarf the cost of the infrastructure that failed.

Resilience evaluation must go beyond RTO and recovery point objective (RPO) specifications. Leaders need to understand how a platform behaves during an attack and whether recovery is a tested, repeatable workflow or an undocumented, manual scramble.

What to measure

Run a tabletop exercise with your shortlisted vendors. Walk through a realistic ransomware scenario and ask them to demonstrate, not describe, the recovery workflow. The gap between what vendors claim and what they can demonstrate is often significant.

- Are snapshots immutable and air-gapped by default, or is immutability an optional add-on?
- What is the blast radius of a ransomware attack? Which systems are isolated, and how quickly?
- What is the actual tested RTO for a full environment recovery, not just a single volume?
- How does the platform detect anomalous behavior indicative of an active attack?
- Is the availability guarantee backed by financial SLA credits, or is it a marketing claim?

Everpure response

Most vendors approach resilience as a backup problem. They sell you a recovery product to protect you from failure in their primary product. Everpure approaches resilience as a platform requirement, meaning the capabilities that protect your data are native to the architecture, not attached to it.

SafeMode™ Snapshots provide immutable, administrator-locked copies of data that cannot be deleted or encrypted by ransomware, even if an attacker gains administrative credentials. This operates at the storage layer, below the reach of most attack vectors, and is on by default, not provided as a premium feature. Combined with anomaly detection that surfaces behavioral indicators of an active attack, Everpure gives your security team visibility and control before the encryption event, not after.

For organizations that need to demonstrate recovery capability to their board or auditors, Everpure offers ransomware recovery SLAs, a contractual commitment to restore operations within a defined timeframe that is supported by the Everpure engineering team. This is a meaningful accountability structure that separates Everpure from every vendor that hands you a runbook and wishes you luck.

5 Operational simplicity

Why it matters

Complexity has costs that rarely appear on a procurement spreadsheet. Every additional management interface, manual tuning requirement, and vendor-specific certification adds to the burden on your team. These accumulate into a hidden tax on engineering capacity that compounds over time. Organizations that have grown through acquisition or vendor sprawl often find that their infrastructure team spends more time managing tools than delivering outcomes.

The shift to AI and real-time data products has made this worse. Platform teams are being asked to support more workloads faster, while the infrastructure beneath them grows more complex. A platform that requires expert-level administration to maintain constrains your team's ability to scale.

Operational simplicity is not a soft requirement. It is a force multiplier. The right platform lets a smaller team manage more infrastructure, respond faster to requests, and redirect energy toward work that creates business value.

What to measure

Ask vendors for reference customers with a similar environment complexity and team size. The metric you want is admin hours per petabyte managed and how that number has changed over time as the environment grew.

- How many management interfaces are required to operate the platform fully?
- What percentage of routine operations—provisioning, tiering, and replication—are automated by default?
- What is the learning curve for new administrators, and what certifications are required?
- How does the platform surface problems before they become incidents?
- What do day-two operations look like at three times your current scale?

Everpure response

Legacy platforms are complex by design. Complexity creates switching costs, drives professional services revenue, and keeps customers dependent on vendor expertise for routine operations. The incumbent vendors in this market have spent decades building environments that require them to maintain those environments.

Everpure is built on the opposite premise. Pure1®, the management service of Everpure, provides a single interface for monitoring, managing, and optimizing the entire Everpure environment across on-premises, public cloud, and edge environments. There's no separate console for each product line, no manual correlation across siloed data, and no tuning required to maintain the performance levels you bought.

Pure1 AI-driven analytics surface recommendations proactively, identifying patterns that indicate emerging issues and recommending corrective action before impact occurs. The result is infrastructure operations that shift from reactive firefighting to proactive management. Everpure customers consistently manage significantly more capacity per administrator than they did on previous platforms not because their teams improved, but because the platform no longer requires heroics to operate.

Ready to run your evaluation?

Get a customized TCO analysis, a live performance demonstration against your workloads, and a reference introduction to customers in your industry.

Get Started With Your Enterprise Data Cloud Blueprint

Visit Our Website

800.379.PURE

