

Designing for Cost Economics While Driving Innovation

How tech leaders can build infrastructure that absorbs cost pressure without putting innovation on hold

Even before the latest market shifts, data was already outpacing budgets and storage costs were climbing. Then NAND flash prices more than doubled in five months, contracts shifted, and vendors that once offered stable pricing started moving the goalposts. Now tech leaders are defending infrastructure spend to finance teams while trying to keep AI initiatives, application modernization, and other strategies moving forward.

The instinct in this situation is to cut, freeze projects, and buy time. That instinct is understandable, but it carries a hidden cost. Every dollar saved by deferring infrastructure investment tends to show up again, with interest, when the business is ready to move and the environment isn't.

There is a better problem to solve. It's not how to spend less, but how to build an environment where cost efficiency and innovation reinforce each other rather than compete. That's not a philosophical point. It's an architectural one.

The real problem isn't pricing

NAND flash pricing has always been cyclical. AI infrastructure build-outs are consuming flash volumes at a historically uncommon pace, and the supply side hasn't kept up. Prices for NAND chips have more than doubled since mid-2025, and they continue to rise. That's a real constraint, and it isn't going away quickly.

But the organizations struggling most with this environment weren't undone by a price change. They were undone by environments with no room to absorb it: underutilized infrastructure, fragmented storage, workloads locked to specific systems, and data overduplicated across environments to keep IT moving. When something shifts in a system already stretched thin, there's nowhere to flex.

Pricing volatility surfaced the problem, but the problem was already there.

What looks like a procurement problem is actually about efficiency and predictability. The environments that handle volatility well aren't the ones with the best vendor relationships or the fastest procurement cycle. They're the ones built to absorb change rather than react to it.

The efficiency problem



Underutilized capacity

Capacity is overpurchased to stay safe, leaving resources locked in place and unable to serve new workloads without manual intervention.



Fragmented data

Data is copied across environments to keep AI and pipelines moving, with each copy increasing cost and adding governance complexity.



Manual operations

Every workload is different depending on who set it up. Without automation, costs scale with headcount rather than efficiency.

Cost efficiency and innovation are not a trade-off

The framing that pits innovation against cost control is worth questioning. It assumes that innovation requires spending more and that saving money means doing less. Neither is necessarily true.

The friction isn't between ambition and budget. It's between the architecture you have and the architecture you need. Most enterprises are running data infrastructure designed for a different era. Storage is siloed, cost structures are based on forecasts instead of actual usage, and upgrades are events rather than continuous processes. That design made sense when workloads were predictable and AI wasn't on the roadmap. It doesn't make sense now.

Efficiency as the foundation for scale

The organizations moving fastest on AI and modern applications aren't spending more, they're spending differently. Capacity that was tied to individual workloads is now pooled. Data that was duplicated across environments is now accessible without being recreated. Operations that required manual intervention now run automatically. The freed budget and freed attention both go toward innovation.

Efficiency isn't a holding pattern while you wait for the market to settle. It's the precondition for moving when the opportunity arises.

Predictability as the engine of investment

Finance teams don't kill innovation projects because they're against innovation. They kill them because the cost basis is unpredictable. When infrastructure spending swings with commodity markets and there's no clear line between committed costs and consumed capacity, every project is at risk in the next budget cycle.

Predictable infrastructure economics lets you make different commitments. When you know what storage will cost over three years rather than hoping the forecast holds, you can plan AI initiatives with real milestones, invest in application modernization on an actual schedule, and defend those investments to the board with numbers that won't require an asterisk.

What the architecture actually changes

The organizations making the most progress on cost and innovation aren't running separate programs. They've made five architectural shifts that compound:

1. **Pooled capacity:** Workloads run where resources are available instead of where data was originally stored.
2. **Consumption-aligned spend:** Cost tracks actual usage instead of forecast commitments.
3. **Duplication eliminated at the source:** Data is accessible without being recreated across every environment.
4. **Operations automated through policy:** Provisioning is consistent, and governance is a default rather than a follow-up.
5. **Infrastructure kept continuously current:** Upgrades stop being events that require timing, justification, and risk.

Each shift makes the others easier to sustain, but the reason they matter isn't cost reduction. It's what they free up. When the environment no longer requires constant intervention to remain operational, the attention and budget devoted to maintenance become available for something else. That's not a secondary benefit. It's the whole point.

The organizations that treat infrastructure efficiency as a cost exercise tend to get some incremental savings, but not enough to make a serious impact. The ones that treat it as an architectural decision tend to get a different kind of environment entirely, one that can absorb a market shock without putting the innovation agenda on hold.

Most organizations won't get here by optimizing around the edges. The limitation isn't awareness. It's the underlying model.

Where Everpure fits

Everpure™ was built on a premise that most storage vendors didn't accept 15 years ago: all-flash was the right direction, and the economics would follow. They did, and that orientation shapes how the platform works today.

The Everpure Platform is a single intelligent system spanning public cloud and on-premises environments, modern applications, and traditional workloads. It unifies storage, automates operations, and builds in efficiency and predictability not as a configuration, but as the foundation.

Evergreen//One

Evergreen//One™ is a storage-as-a-service model that aligns spend to actual consumption. There's no overprovisioning requirement, end-of-year scramble, or surprise costs when the market moves. Pricing is locked, and new capacity delivers in as little as four weeks. Lifecycle management is included, so upgrade cycles are no longer a project on someone's plan. For finance, that means a cost structure that the board can forecast. For engineering, it means having capacity available when the project needs it.

DirectFlash

Everpure DirectFlash® technology communicates directly with flash memory rather than routing through interfaces designed for spinning disk. The result is better performance, lower power consumption, and higher efficiency per terabyte, enabling more life out of every drive. In a market where flash costs have more than doubled, getting more out of what you have isn't a secondary benefit. It's the point.

Data reduction that doesn't cost performance

The Everpure data reduction engine delivers high reduction ratios without meaningful performance trade-offs. Less physical storage for the same data footprint means lower cost per effective terabyte and reduced exposure to commodity pricing swings.

Everpure Cloud

Everpure Cloud provides enterprises with consistent storage economics and a single operating environment across on-premises and cloud environments. Workloads move when it makes sense for the business, not based on where the data happens to be sitting. Cloud storage costs come down significantly, and data migration happens without lock-in.

Portworx

Portworx® brings modern application velocity to environments that require Kubernetes to run alongside traditional workloads. There's no second operating model or extra complexity as applications evolve. The infrastructure moves with the application estate. It doesn't make it wait.

What good infrastructure economics actually looks like

Budget freed from keeping the lights on becomes available for the next initiative not in theory, but in the next planning cycle. Costs the finance team can forecast mean investment decisions are made on business value, not hedged against infrastructure uncertainty. A vendor that delivers what it commits to across a contract term that doesn't require renegotiation every time a commodity market moves changes how the business thinks about infrastructure as an asset.

That's not a bold claim. It's what good infrastructure has always looked like. The volatility in the market has simply made it clearer how many environments weren't built that way.

The organizations that will move fastest when conditions stabilize are building that foundation now.

Starting the conversation

The practical first step isn't a purchasing decision. It's an honest assessment of where your current environment leaks cost and where it blocks agility. Most organizations find that the two problems share the same root causes: fragmentation, manual operations, and a procurement model that requires overbuying to stay safe.

These questions surface the gaps. The answers, or the absence of them, tend to point toward the same places.

Questions worth asking your team

- What percentage of our provisioned storage is actually utilized? Where is capacity sitting idle?
- How much of our storage cost is tied to duplicated data, and how many of those copies are essential?
- What does a workload deployment actually cost us in people-hours, from ticket to production?
- When did we last assess whether our commercial model matches how we actually consume storage?
- How much of our infrastructure team's time goes toward lifecycle management rather than new projects?

What the answers reveal

Vague answers, contested numbers, or uncomfortable silence aren't edge cases. They're diagnostic. Each one maps to a specific structural problem that won't resolve on its own.

If your team says...	It usually means...
"We think utilization is around 60%, but it's hard to measure."	Capacity is siloed. Each system looks at its own numbers, not the environment as a whole. Actual usable headroom is invisible.
"We have a lot of copies, but they're all for different teams."	Data architecture follows team boundaries, not data flows. Every boundary is a duplication point.
"Provisioning takes a few days, maybe a week. It depends."	Manual processes create variable outcomes. "It depends" is the sound of an environment running on coordination, not automation.
"We signed a three-year deal in 2022. We're locked in until next year."	The commercial model is misaligned with actual consumption. You're likely carrying unused capacity you're still paying for.
"Refresh planning takes up most of Q3 every other year."	The life cycle is still managed as a project. Every cycle is a risk: wrong timing, compressed lead times, and forced architecture decisions.

Most environments have at least two or three of these in play at the same time. That's not a failure of planning. It's the natural result of infrastructure built incrementally, vendor by vendor, cycle by cycle, without a model designed to hold it together.

The organizations that have changed that underlying model aren't spending more. They're building on a foundation where every subsequent decision costs less and moves faster.

Why Everpure

Everpure has spent 15 years building for the architecture that storage should be, not the one most vendors were comfortable selling. That means all-flash from the start, DirectFlash technology that squeezes more out of every drive, and a data reduction engine that cuts footprint without cutting performance.

Evergreen//One turns the commercial model into an advantage: spend is tied to actual consumption, pricing is locked against market swings, capacity is delivered in weeks, and lifecycle management is built in so the environment stays current without a refresh project on the calendar every other year.

The Everpure Platform connects on-premises and cloud environments under a single operating model, with consistent economics, unified governance, and full visibility across the environment. Everpure Cloud eliminates the cost premium of running in the public cloud. Portworx runs modern applications alongside traditional workloads without a second infrastructure layer. And Pure1* gives you a single view across everything, so the numbers are never hard to find.

The result is an environment that gets more capable over time, not more expensive to maintain.

Get Started With Your Enterprise Data Cloud Blueprint

[Visit Our Website](#)

[800.379.PURE](tel:800.379.PURE)

