

The cloud you can't escape

User lock-in, value erosion, and the hyperscaler playbook

Introduction: From digital paradise to financial dependence

Ever get the feeling you're paying more but getting less?

Maybe it's your streaming provider: once, you had access to every movie in the world; now, you're paying for five different subscriptions, and waiting for a new episode to drop each week. Or maybe it's a ridesharing app, which used to be fast and fun, but is now slower and more expensive than a regular taxi. Or it's that food delivery service, packed with hidden fees and last-minute tips that double the price of a meal.

Back in 2017, Netflix told customers, "love is sharing a password." But apparently you can put a price on love, because in 2023, password sharing cost \$7.99. The company hooked users with amazing price and value and then, systematically, started chipping that value away. These things don't happen by accident. They're part of a predictable playbook that vendors use to lure you in, lock you down, and extract maximum profits while destroying the services you loved. Author Cory Doctorow calls this Enshittification: A strategy that starts out great for customers and ends up feeling like being stuck in an expensive, declining ecosystem.

It's a universal pattern with three parts:

- **Stage 1: The Hook.** To attract a massive user base, the product or service is incredibly good to its users. They offer you low prices, generous features, and a great experience to build market share and create a "sticky" ecosystem.
- **Stage 2: The Squeeze.** Once users are locked in, the platform begins to abuse them to benefit its business customers (like developers, advertisers, or creators). Features are put behind paywalls, and organic reach is throttled to encourage paid promotion.
- **Stage 3: The Harvest.** Finally, having locked in both users and businesses, the platform turns on everyone. Sometimes known as rent-taking, this stage exploits customers' inability to move by funneling all surplus to its shareholders until the platform becomes a frustrating, barely-usable shell of its former self.

Once you see it, you recognize it everywhere. And if you're in enterprise IT, it probably sounds familiar. That's because the largest cloud computing providers, known as hyperscalers, have been running this playbook for the last decade. It's not an accident—it's their business model. It's why CIOs cite cost overruns, slower support response times, and complex billing as growing concerns—and why migrating away from proprietary platforms can become an expensive, complex project.

An example: Financial engineering at the customer's expense

Here's a simple example of financial engineering masquerading as operational improvement. Starting in 2020, hyperscalers started quietly extending the lifespan of servers on paper. They told Wall Street this was possible because of operational excellence—but the reality was, critical workloads were being hosted on increasingly aged infrastructure.

- **Amazon Web Services (AWS)** systematically extended its server depreciation schedules from 3 years to 4 years (2020), then to 5 years (2022), and finally to 6 years (2024), boosting its operating income by \$2.3 billion in 2020 alone.
- **Microsoft Azure** followed the same playbook, extending its server depreciation schedule from 4 years to 6 years in 2023, adding \$3.7 billion to its fiscal year profits.
- **Google Cloud** also extended its schedules from 3 years to 6 years, saving \$3.4 billion in 2023.

AWS CFO Brian Olsavsky claimed that the company had, “refined our software to run more efficiently on the hardware” which “lowers stress on the hardware and extends the useful life.” But academic and industry data paint a starkly different picture: Hardware fails more often as it ages. Server failure rates climb from 5% in year one to 11% by year four and 18% by year seven. The cost of maintaining that hardware increases by 148% by year five and 300% by year seven—costs that get passed on to customers.

So what did AWS do when performance actually mattered to them? In 2025, the company reversed course for its AI hardware, shortening the depreciation schedule from 6 years back to 5. It was an admission that the accounting fiction was unsustainable for high-performance workloads, resulting in a \$700 million hit to operating income and a \$920 million charge for retired equipment.

This practice, known in accounting as “sweating the assets”, is just one example of the value degradation playbook in action. While you paid the same or higher prices, the underlying infrastructure was becoming older and less efficient by design. Hyperscalers can only do this because they've locked you into their ecosystem, making it prohibitively hard for you to take your business elsewhere. The cloud promise of portable, elastic workloads turned out to be a long-term dependency.

When lock-in becomes dependence

Once a platform has attracted enough users, its next priority is to make it incredibly difficult for them to leave. Hyperscalers have perfected this art through two primary mechanisms: high exit costs and deep technical integration.

The digital tollbooth: Egress fees

Data egress fees are, simply put, a penalty for moving your own data out of a cloud provider's network. They're a departure tax—one that's been declared illegal in some jurisdictions. While it costs next to nothing to upload data, downloading it or transferring it elsewhere comes with a hefty price tag. A comparison reveals just how non-competitive these fees are:

Provider	Cost per Gigabyte (GB)
Google Cloud	\$0.12
AWS	\$0.09
Azure	\$0.0087
Alternative Providers	\$0.01

The real-world impact is staggering. These fees are 8 to 12 times higher than those of more affordable competitors and can constitute 10-15% of a company's total cloud bill. For an enterprise moving 50TB of data per month, that amounts to \$51,000 annually just for the privilege of accessing their own information. Hyperscalers claim they offset their networking costs—but once regulators stepped in, they agreed to waive those fees for departing customers. In practice, these charges discourage customers from pursuing more cost-efficient options.

The velvet cage: Proprietary services

The second layer of dependence involves moving customers from portable services (Infrastructure-as-a-Service, or IaaS) to deeply integrated, proprietary platforms (Platform-as-a-Service, or PaaS).

AWS Lambda is a classic example. It offers tremendous value by allowing developers to run code without managing servers (“serverless”). However, using Lambda means embedding an application deep within the AWS ecosystem, relying on AWS-specific tools for permissions (IAM), logging (CloudWatch), and dozens of other integrations. This strategy is not limited to one service. DynamoDB, AWS’s database service, creates even deeper lock-in through its proprietary data model, with features like single-table design and specific partition keys that have no direct equivalents, making migration a fundamental application rewrite. The other hyperscalers have similar proprietary platforms that are impossible to leave.

This lock-in is by design. The cost and effort to escape become prohibitively high, with estimates showing that migrating off proprietary PaaS services can take 3-12 months and incur overhead costs of 40-200%. The convenience that once felt liberating becomes a form of dependency—an elegant constraint.

Value extraction laid bare: Broadcom and VMware

The acquisition of VMware by Broadcom in 2023 became one of the clearest examples of value extraction. Almost as soon as Broadcom acquired the virtualization software company, it began one

of the clearest examples of value extraction at scale. Almost immediately after the acquisition, Broadcom launched a series of aggressive pricing and licensing changes that rippled through the industry.

The price hikes were extraordinary:

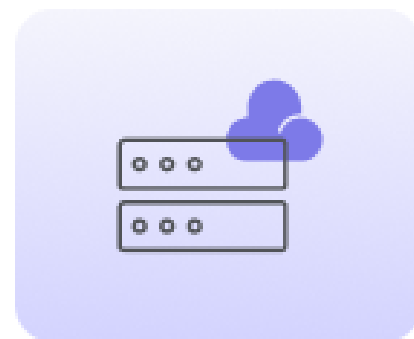
- **AT&T** faced a quote for a 5-year commitment that represented a 1,050% price increase.
- **Computershare** received a new quote that was 10 to 15 times their previous costs.
- **European cloud providers** reported increases ranging from 800% to 1,500%.

Susan A. Johnson, EVP at AT&T, didn’t mince words in 2024: “This proposed annual increase of 1,050 percent in one year is extreme and certainly not how we expect strategic partners to engage in doing business with AT&T.”

The shifts didn’t stop there. Broadcom discontinued all perpetual software licenses overnight, forcing customers into expensive subscriptions. It changed licensing from per-socket to per-core and raised the minimum core count per subscription from 16 to 72, forcing customers to pay for phantom capacity.

Broadcom didn’t stop at its customers, though: it went after its partner ecosystem, terminating all existing programs and replacing them with an invitation-only model requiring \$500,000 to \$1M in annual revenue. The company then pursued legal action against former partners and customers alike.

This was a wake-up call for technology leaders. If a market leader could so rapidly alter its pricing and licensing model, what’s to prevent others from following? CIOs across industries began asking the same question: how dependent have we become on a handful of providers, and how much control have we ceded in the process?



This playbook isn't unique; its patterns are echoing across the digital services you use every day. Hyperscaler support quality plummeted: Microsoft unilaterally removed free support from its enterprise agreements, forcing customer to pay more for the same service that was once standard.

It's not just you: Recognizing the pattern

This structural pattern is not a series of isolated business decisions; it is the universal playbook for dominant platforms. Once you recognize the cycle of an early, user-friendly “hook” followed by a systematic degradation of value, and an onerous fee, you see it everywhere: In record clubs, social media platforms, retail marketplaces, airlines, loyalty programs, and even dating platforms

Inevitably, however, customers wake up. The bills pile up, service quality declines, and the frustration grows.

The escape plan

Faced with escalating bills and diminishing returns, technology executives are no longer passively accepting the costs of the cloud. A growing movement known as “repatriation”—moving workloads out of the public cloud and back to on-premises infrastructure or open platforms—is gaining momentum. Moving significant workloads away from hyperscalers to alternate hosting, compute, and hardware platforms is not only plausible—it's possible.

The Repatriation Math

For companies operating at scale, the financial case for leaving the cloud can be overwhelming.

- **Dropbox:** Saved \$74.6 million over two years by moving its data from AWS to custom infrastructure, with an initial investment payback period of just 8 months.
- **37signals (Basecamp):** Cut cloud spending by 60%, achieving over \$2 million in annual cost reductions after purchasing its own servers.
- **GEICO:** After its annual cloud bill ballooned to over \$300 million—2.5 times higher than projected—the insurance giant is moving away from the public cloud to regain control of its costs.

AI is further exacerbating such costs. You can buy a state-of-the-art GPU server for what it costs to rent the equivalent from AWS for a couple of months. For any AI job running continuously, buying your own hardware is 50% cheaper in less than a year. These price discrepancies are forcing CIOs to rethink their entire compute strategy.

GEICO, 37signals, and Dropbox had immediate pain, and the technical wherewithal to assess their entire computing posture so they could adjust. It's a luxury many other firms lack, which is why they turn to trusted cloud advisors. The first step is an assessment that offers the basic facts the executive team needs to put a repatriation strategy in place: How vulnerable are you? How egregious are your overages? How dependent are you on proprietary platforms and tools? Which workloads are ready to be moved, and which must be disentangled? And, importantly, which workloads make sense for hyperscalers? It's important to acknowledge that **elastic, large-scale cloud platforms play a vital role in any modern computing strategy—just not the only role.**

The Smart Path Forward: Strategic Hybrid Cloud

The solution for most isn't to abandon hyperscalers, but to be strategic. The best companies run stable, predictable workloads on cost-effective on-premises or open-platform infrastructure (for example, migrating Databricks workloads to Apache Spark hosted on a more affordable cloud platform), while using hyperscalers for what they do best: handling unpredictable spikes in demand, bursty and transient workloads, and prototyping.



This isn't a fringe movement. Recent data shows that 86% of CIOs plan to repatriate at least some of their workloads, and 73% of organizations already employ a hybrid cloud strategy. Once you understand the enshittification playbook, you can chart a path back to control of your compute strategy.

Conclusion: You're not mismanaging—it's the model

If your cloud bills are spiraling out of control while service gets worse, you're not alone. If costs were stable, you'd expect budgets to resemble estimates—but 69% of enterprises exceeded their cloud budgets last year. Soaring costs aren't a sign of your failure to manage resources; they're symptoms of a deliberate platform strategy. You're not mismanaging, you're experiencing structural exploitation.

Andreessen Horowitz famously summarized the economic reality of the cloud for growing businesses: "You're crazy if you don't start in the cloud; you're crazy if you stay on it."

Like every other business running this playbook, hyperscalers offered an incredible launchpad—but once you achieved scale, the incentives shifted. Recognize the patterns and find the right partner to help you make strategic decisions—whether through repatriation, hybrid cloud, or switching providers.

The entire hyperscaler business model is built on one central assumption: that you believe it is too difficult and too expensive to leave. They've done all of their math to maximize their profit. **The only way to take back control is to do yours.**



READY TO LEARN MORE?

Own Your Destiny in the Cloud with Aptum.

Visit aptum.com to
start the conversation →