aptum

Time to Grow Up

What happens after the blitz?

Every successful startup inevitably goes through several stages: An idea becomes a product, finds product-market fit, then scales rapidly to capture marketshare before incumbents can respond or copycats can compete. In the words of Reid Hoffman and Chris Yeh:

> "When a startup matures to the point where it has a killer product, a clear and sizable market, and a robust distribution channel, it has the opportunity to become a "scaleup"—a world-changing company that touches millions or even billions of lives. How do startups evolve into scaleups? In a word, by blitzscaling."

The whole point is to achieve "massive scale at incredible speed." Growth is the only job; it is, to quote Hoffman, "prioritizing speed over efficiency in the face of uncertainty."

But what happens when you've achieved market dominance? When you're a Dropbox, Spotify, Uber, or Salesforce? Now, you're a public company, facing scrutiny from activist investors. You're the market leader, so everyone's challenging you. You've saturated the market, putting pressure on margins. You need to do more with less. Efficiency and certainty take the place of experimentation and risk.

This transition is often accompanied by a change in management. The board brings in "Adult leadership" to oversee operations—after all, the skills needed to defy norms in the early stages of growth are seldom those needed to operate a large organization masterfully. Governance and procurement become more important, inviting scrutiny from regulators. With thousands or even millions of customers relying on you, you have to make a business case for changes and enhancements. Contracts limit what you can do.

It's not just time to change leadership—it's time to change your computing strategy.

Cloud computing was a boon to the startup world. Millions of new businesses no longer needed to dilute their ownership in order to buy dedicated hardware; instead, they could scale their compute resources as they grew. It was easy to get started: the "big three"—Amazon Web Services, Microsoft Azure, and Google Cloud Platform—showered aspiring founders with credits and set up startup programs.

But there was a catch. While uploading data was free, the compute power needed to do something useful with it was expensive. In the same way a taxi is convenient, but expensive if you use it all the time, "elastic" computing cost much more than a dedicated machine once your loads were predictable. "Own the base, rent the spike" was the wisdom of the day.

To keep customers hooked, hyperscalers layered convenient platforms (known as Platform-as-a-Service, or PaaS) atop their on-demand infrastructure (known as Infrastructure-as-a-Service, or IaaS). Why install and run your own instance of a free, open source database like MySQL when you can instead use AWS' turnkey platform, DynamoDB? Why build a data pipeline when you can just use AWS Kinesis?

Fast-growing startups loved this, because they could focus on building products and markets without worrying about the high costs.

The ease of use had a price, of course: once you built atop a proprietary platform, it was hard to leave. And it's not just the hyperscalers playing these tricks. For example, Databricks (the creators of the open-source Apache Spark project) offers a data platform that's quick to try—but over time, far more expensive than running your own version of Spark.

The numbers don't lie

So now you're beyond scale-stage growth, and trying to bring some maturity to your established business, and you're locked in. Cloud bills constantly exceed predictions. You're reliant on more and more services. You're ready to grow up, but you're trapped. You've blitzscaled your way straight into rent-taking.

- According to Harness' **FinOps in Focus 2025** report, an roughly 21% of the money enterprises spend on cloud computing ($44.5 billion globally) is wasted on under-utilized resources. That's enough to fund your startup from scratch many times over.
- Flexera's **State of the Cloud** reports consistently show cloud waste around 30% of total spend, with more than half of the respondents to a study by the FinOps foundation saying that workload optimization and waste reduction is their top priority.
- Cloud spend keeps exceeding budget. Flexera's 2025 report found that budgets ran over by 17% on average—with only 30% of respondents actually **knowing** where their cloud budget was spent.

These aren't hypotheticals. Some high-visibility tech darlings have made headlines by replatforming and repatriating workloads. 37signals (makers of Basecamp and HEY) got a $3.2M cloud bill in 2022, prompting them to rethink their strategy. They bought $700,000 worth of servers, lowering their annual costs by $2M; then copied 10 petabytes from AWS S3 to Pure Storage, cutting much of their remaining spend. The payoff took only 18 months, and they didn't add any team members.

As companies mature and start examining their cloud bills with adult scrutiny, it becomes increasingly clear that hyperscalers aren't partners—they're landlords. And like all landlords who sense a captive tenant, they're squeezing harder than ever: Extending hardware depreciation periods, ghosting customers for support, adding onerous fees for things like data egress and IP addresses, and worse.

At the same time, these companies are pouring billions of dollars into AI infrastructure, treating traditional compute as an afterthought. Your boring, predictable, revenue-generating workload is subsidizing their AI expansion.

You’re Viewing The Preview

Click here to get instant access to the complete whitepaper

You’re Viewing The Preview

Click here to get instant access to the complete whitepaper

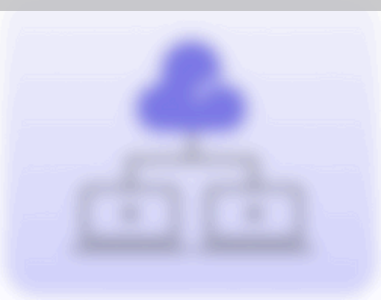

Ready to learn more?

Own Your Destiny in the Cloud with Aptum.