Cloud on a Budget: Real-World Tactics

Steven Sklar

Who Am I?

Who Am I?

- 10+ Years of Startup Experience
  - I’ve seen things…
- Variety of Stages (and Budgets!)
  - Seed through Series E
  - 5 Engineers to 100s
- Economics and Finance Degree
  - Always thinking about the money
- Software Developer turned DevOps Engineer
  - Experienced all parts of the SDLC

AWS

Native AWS integration. One CloudFormation template to deploy.

Azure

Azure-native deployment with Lighthouse-managed auditable security.

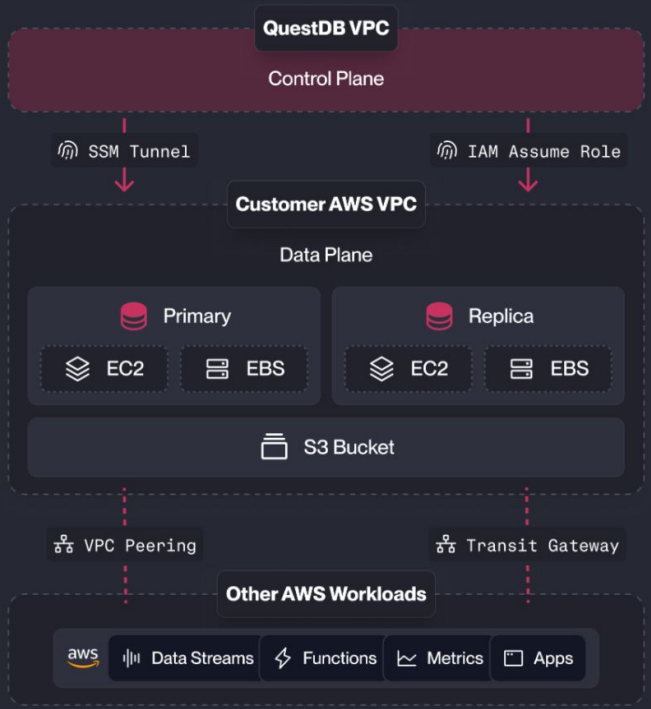

Your AWS, QuestDB's operation

Seamless AWS integration

One CloudFormation template grants us secure access to one of your accounts. We then provision the compute, networking, and storage resources for your cluster.

We support VPC peering and Transit Gateway connections to keep your traffic private as it flows to and from your other AWS workloads.

No Kubernetes required!

The Cloud is Expensive

Top cloud challenges for all respondents

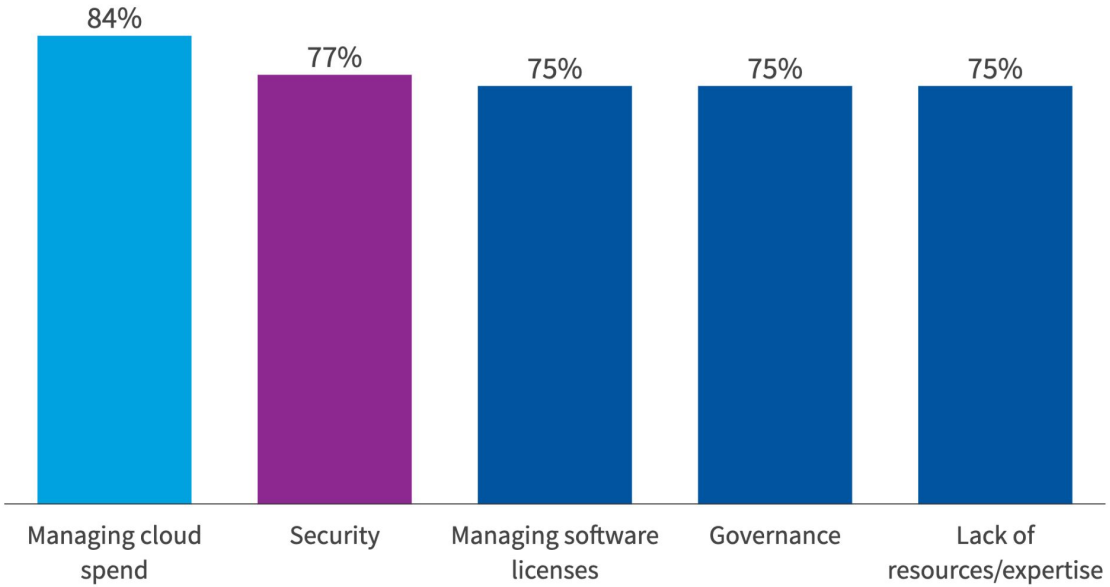

N=759

Source: Flexera 2025 State of the Cloud Report

flexera

$195 billion

$723 billion estimated cloud spend (Gartner 2024)

27% of cloud spend is “wasted” (Flexera 2025)

Google Cloud (over)Run: How a free trial experiment ended with a $72,000 bill overnight

reddit WEBDEV comments other discussions (3)

8736

Netlify just sent me a $104K bill for a simple static site

Question (self.webdev)

submitted 1 year ago * by liubanghoudai24

So I received an email from Netlify last weekend saying that I have a $104,500.00 bill overdue. At first I thought this is a joke or some scam email but after checking my dashboard it seems like I am truly owing them 104K dollars:

That's 190TB bandwidth in 4 days

So I was like 😅😅😅 and think okay maybe I got ddos attacked. Since Netlify charges 55$/100GB for the exceeding bandwidth, the peak day Feb 16 has 33385/55 * 100GB = 60.7TB bandwidth in a day. I mean, it's not impossible but why attack a simple static site like mine? This site has been on Netlify for 4 years and is always okay with the free tier. The monthly bandwidth never exceeded even 10GB, and has only ~200 daily visitors.

I contacted their billing support and they responded me that they looked into it and the bandwidth came from some user agents, meaning it is a ddos attack. Then they say such cases happen and they usually charge their customer 20% on this. And since my amount is too large, they offer to discount to 5%, which means I still need to pay 5 thousand dollars.

This feels more like a scam to me. Why do serverless platforms like Netlify and Vercel not have ddos protection, or at least a spend limit? They should have alerted me if the spending skyrocketed. I checked my inbox and spam folder and found nothing. The only email is "Extra usage package purchased for bandwidth". It feels like they deliberately not support these features so that they can cash grab in situations like this.

The ddos attack was focused on a file on my site. Yes it's partly my fault to put a 3.44MB size sound file on my site rather than using a third-party platform like SoundCloud. But still this doesn't invalidate the point of having protection against such attacks, and limit the spending.

I haven't paid that $5k yet and decided to post here to hear what others think first. And yes I have migrated my site to Cloudflare. Learned my lesson and will never use Netlify (or even Vercel) again.

UPDATE: Thank you all for the suggestions I have posted this on HackerNews.

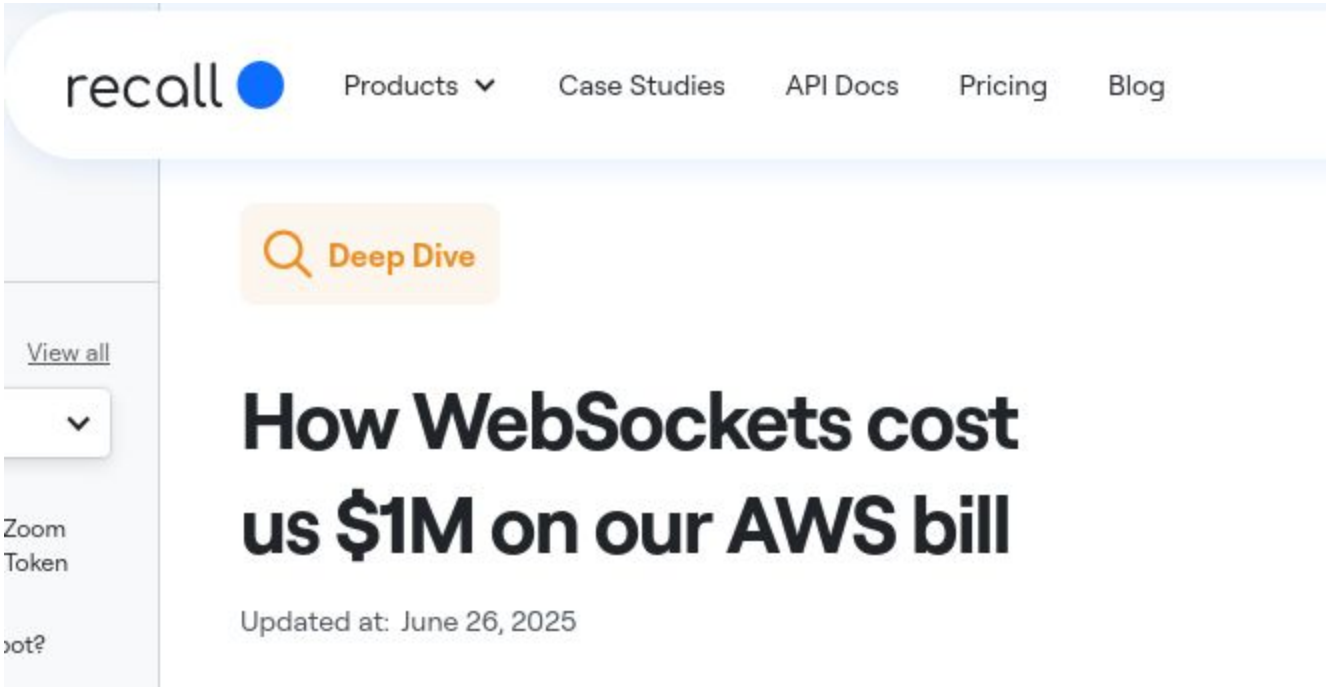
recall
Products
Case Studies
API Docs
Pricing
Blog
View all
Zoom
Token
oot?
Deep Dive
How WebSockets cost us $1M on our AWS bill
Updated at: June 26, 2025

Reliability, At What Cost?

Do We Really Need Five 9s?

Availability	Downtime / Year	Downtime / Month	Downtime / Week	Downtime / Day
99.999%	5.256 minutes	0.438 minutes	0.101 minutes	0.014 minutes
99.995%	26.28 minutes	2.19 minutes	0.505 minutes	0.072 minutes
99.990%	52.56 minutes	4.39 minutes	1.011 minutes	0.144 minutes
99.950%	4.38 hours	21.9 minutes	5.054 minutes	0.72 minutes
99.900%	8.76 hours	43.8 minutes	10.108 minutes	1.44 minutes
99.500%	43.8 hours	3.65 hours	50.538 minutes	7.2 minutes
99.250%	65.7 hours	5.475 hours	75.808 minutes	10.8 minutes
99.000%	87.6 hours	7.3 hours	101.077 minutes	14.4 minutes

SLAs Are One Thing…

Service	Stated Uptime
Google Workspace	99.9%
AWS RDS (Multi-AZ)	99.95%
AWS EC2 (Multi-AZ)	99.99%
Azure 2+ VMs across AZs	99.99%
Azure SQL Database	99.995%

Sources:

- https://workspace.google.com/terms/sla/
- https://aws.amazon.com/rds/sla/
- https://aws.amazon.com/compute/sla/
- https://www.microsoft.com/licensing/docs/view/Service-Level-Agreements-SLA-for-Online-Services
- https://azure.microsoft.com/en-us/blog/understanding-and-leveraging-azure-sql-database-sla/

But Outages Still Happen

TECH

Azure outage: Microsoft still working on fix, says recovery expected in several hours

PUBLISHED WED, OCT 29 2025•12:43 PM EDT | UPDATED WED, OCT 29 2025•7:39 PM EDT

Jonathan Vanian
@IN/JONATHAN-VANIAN-B704432/

SHARE

And Keep On Happening

Amazon outage 'resolved' as Snapchat and banks among sites impacted

20 October 2025

Share Save

Liv McMahon, technology reporter and Lily Jamali, North America technology correspondent

And Are Now Caused By LLMs

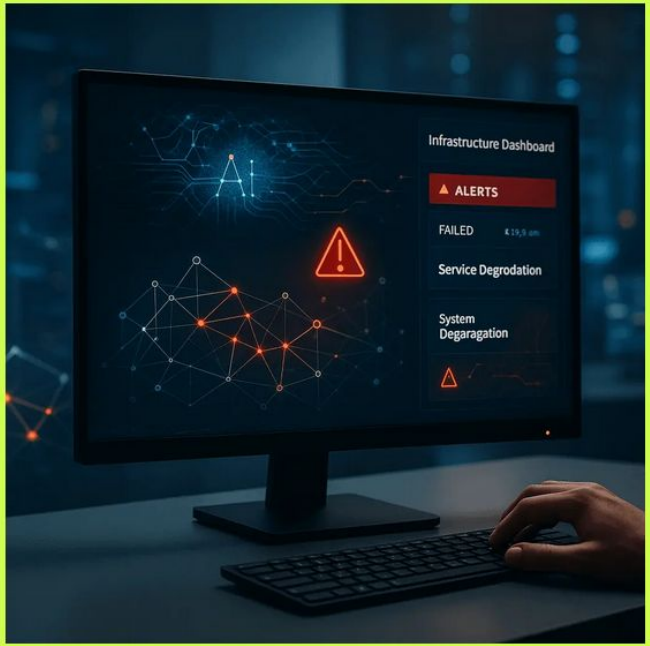

ENTERPRISE/SAAS / AWS

Amazon blames AI-assisted deployments for AWS outages

AWS infrastructure issues tied to AI production changes spark internal review

by **The Tech Buzz**

PUBLISHED: TUE, MAR 10, 2026, 4:25 PM UTC | UPDATED: WED, MAR 11, 2026, 1:29 PM UTC

Not Just Cloud Providers Are Having Problems

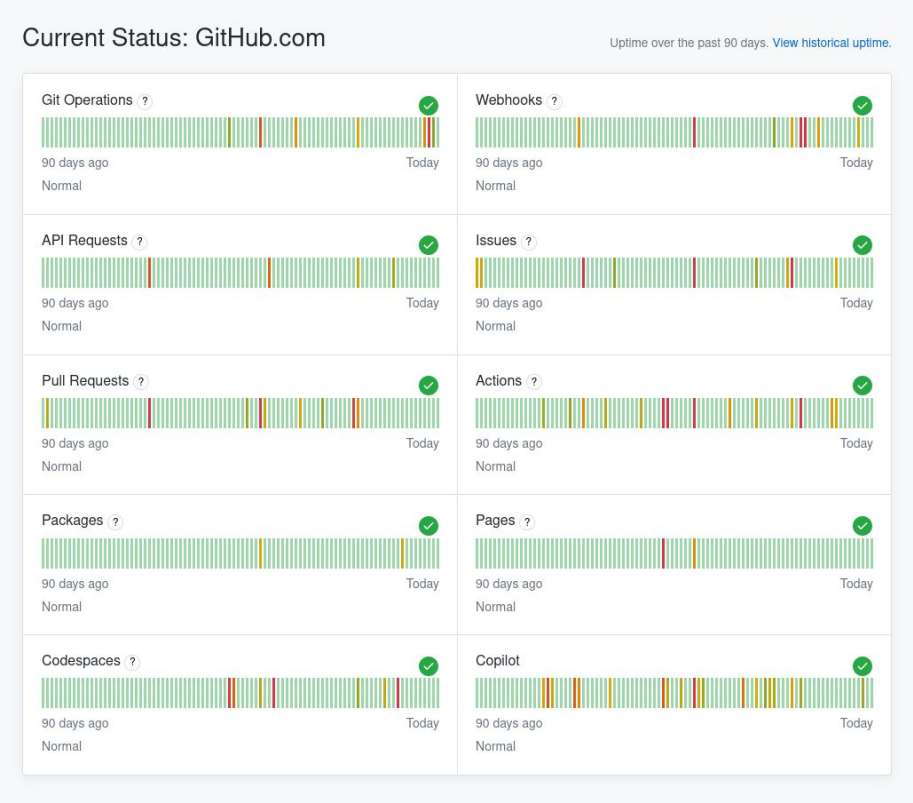

You’re Not Amazon

Probably…

Do I Really Need This Managed Service?

Databases: AWS RDS Entry-To-Mid Level

- Assumptions
  - Region: us-east-2
  - DB: Postgres
  - 1 Node, single AZ
  - Instance: db.m4.xlarge (4 vCPU/ 16GB)
  - Storage: 100GB gp3
  - Backups: 100GB

- Total Cost: $287.45/mo ($3,449.45/yr)
  - Compute: $266.45
  - Storage + Backups: $20.00

Databases: EC2 + Postgres

- Assumptions
  - Region: us-east-2
  - DB: Postgres
  - 1 Node, single AZ
  - Instance: m4.xlarge (4 vCPU/ 16GB)
  - Storage: 100GB gp3, 3k iops, 250 MiB throughput
  - Backups: 2x Daily, 1% changed per backup

- Total Cost: $165.50/mo ($1,986.00/yr)
  - Compute: $138.00
  - Storage + Backups: $27.50

Databases: AWS RDS with Addons

- Assumptions
  - Previous Configuration
  - Database Insights
  - RDS Proxy

- Total Cost: $367.75/mo ($4,413.00/yr)
  - Compute: $266.45
  - Storage + Backups: $20.00
  - Database Insights: $36.50
  - RDS Proxy: $43.80

Databases: AWS RDS Multi-AZ

- Assumptions
  - Previous Configuration
  - Primary in one AZ, Standby in another AZ

- Total Cost: $645.70/mo ($7,748.40/yr)
  - Compute: $532.90
  - Storage + Backups: $32.50
  - Database Insights: $36.50
  - RDS Proxy: $43.80

Databases: EC2 + Postgres

- Assumptions
  - Region: us-east-2
  - DB: Postgres
  - 2 Nodes, 2 AZs
  - Instance: m4.xlarge (4 vCPU/ 16GB)
  - Storage: 100GB gp3, 3k iops, 250 MiB throughput
  - Backups: 2x Daily, 1% changed per backup
  - Cross-AZ replication traffic: ?

- Total Cost: $331.0/mo ($3,972)
  - Compute: $276.00
  - Storage + Backups: $55.0

Messaging: SQS

- Assumptions
  - Region: us-east-2
  - 1 million messages per day (FIFO queue – single delivery guaranteed)
  - Internal traffic only

- Total Cost: $15.21/mo ($182.52/yr)

Storage: S3

- Assumptions
  - Region: us-east-1
  - Intelligent Tiering
  - 1 TB storage
  - 16MB avg object size
  - 100,000 PUT/COPY/POST/LIST
  - 10,000,000 GET/SELECT

- Total Cost: $28.22 ($338.64/yr)
  - **11 nines durability** (99.999999999%)
  - Strong read-after-write consistency
  - Automatic scaling, lifecycle policies, server-side encryption, access logging, and more

Managed Services: Kubernetes

- Assumptions
  - Region: us-east-2
  - Kubernetes Version: v1.36 (released March 2026)

- Total: $73.00/mo ($876/yr)
  - Don’t forget cross-AZ traffic (mandatory due to multi-AZ requirement)
  - And load balancing costs for network ingress/egress

Managed Services: Kubernetes

- Assumptions
  - Region: us-east-2
  - Kubernetes Version: v1.31 (released Aug 2024)

- Total: $438.00/mo ($5,256/yr)
  - $365.00/mo for Extended Support ($4,380/yr)
  - Eventually, you will be forced to upgrade

Wait, I’m Paying for What?!?

Volume Types

- Always use gp3

	3k IOPS	6k IOPS	10k IOPS	16k IOPs
gp2	$100	$200	$333.40	$533.0
disk size	1 TB	2 TB	3.34 TB	5.34 TB
gp3	$80	$95	$115	$145
disk size	1 TB	1 TB	1 TB	1 TB

Unused Volumes

- Assumptions
  - gp3
  - 500 GB
  - 10k IOPS
  - 500 MiB Throughput

- Total Cost: $90/mo ($1,080/yr)

Snapshots

- If you really need the data, take a snapshot!
- $0.05/GB-month per snapshot (stored incrementally)

- Full Disk Snapshots: $25/mo ($300/yr)
- Archived Snapshots: $6.25/mo ($75/yr) + $15 one-time restore

AMIs

- AMIs are backed by snapshots
- Deleting the AMI DOES NOT delete the underlying snapshot

CloudWatch Logs

- $0.50 / GB ingested
- $0.03 / GB stored per month

- 100 GB of logs stored for a month
  - $50.00/mo (ingest) + $3.00/mo (storage) = $53.00/mo
- Forget to rotate out old logs?
  - After 12 months…
  - $50.00/mo (ingest) + $36.00/mo (storage) = $86.00/mo
  - And increasing every month after…

S3 Tiering and Lifecycles

- Standard ($0.023 / GB-month)
  - Free retrieval
- Infrequent Access
  - Standard-IA ($0.0125 / GB-month)
  - One Zone-IA ($0.01 / GB-month)
- Glacier
  - Instant Retrieval (milliseconds)
  - Flexible Retrieval (minutes)
  - Deep Archive (hours)
- Don’t forget about retrieval costs!

S3 Tiering and Lifecycles

For 1 TB of data…

Tier	Storage Cost (per month)
S3 Standard	$23.00
S3 Standard-IA	$12.50
Glacier Flexible	$4.00
Glacier Deep Archive	$1.00

S3 Tiering and Lifecycles

For 1 TB of data (storage & retrieval costs)…

GB Retrieved	**Standard**	**Standard-IA**	**Glacier Flexible**	**Glacier Deep Archive**
Retrieval Time:	Milliseconds	Milliseconds	Minutes-Hours	Hours (9-12)
0 GB	$23.00	$12.50	$4.00	$1.00
100 GB	$23.00	$13.50	$5.00	$3.00
250 GB	$23.00	$15.00	$6.50	$6.00
500 GB	$23.00	$17.50	$9.00	$11.00
750 GB	$23.00	$20.00	$11.50	$16.00
1000 GB	$23.00	$22.50	$14.00	$21.00

S3 Tiering and Lifecycles

- S3 Intelligent Tiering
  - Moves data between tiers at no retrieval charge
  - $0.0025 per 1k objects monitored
- Use those lifecycle policies!

For 1 TB of data…

Objects	**Avg Size**	**Monthly Monitoring**	**Cheaper Than IA?**
1,000,000	1 MB	$2.50	Yes
10,000,000	100 KB	$25.00	Marginally
100,000,000	10 KB	$250.00	No

Networking

- Ingress
  - Free!
- Egress
  - Public Internet
    - Starts at $0.09/GB (first 10 TB/month)
  - Cross-AZ
    - $0.01/GB transferred
  - VPC Endpoints
    - Gateway Endpoints are Free!
      - S3 & DynamoDB
    - Private Link Endpoints are Not
      - $7.20/month per endpoint
      - $0.01/GB transferred

Networking

- NAT Tax
  - $32.40/month per NAT Gateway ($388.80/yr)
  - $0.045/GB data processed through NAT
- Fck-nat
  - https://fck-nat.dev/stable/

Bandwidth	Instance type	Price per Month
32Mbps	t4g.nano	$3.06
64Mbps	t3.micro	$7.59
1.6Gbps	c6gn.medium	$32.81
3.125Gbps	c7gn.medium	$48.25
5Gbps	c7gn.large	$132.20
25Gbps	r6in.8xlarge	$1074.56
50Gbps	c7gn.8xlarge	$1457.66

AWS Data Exports

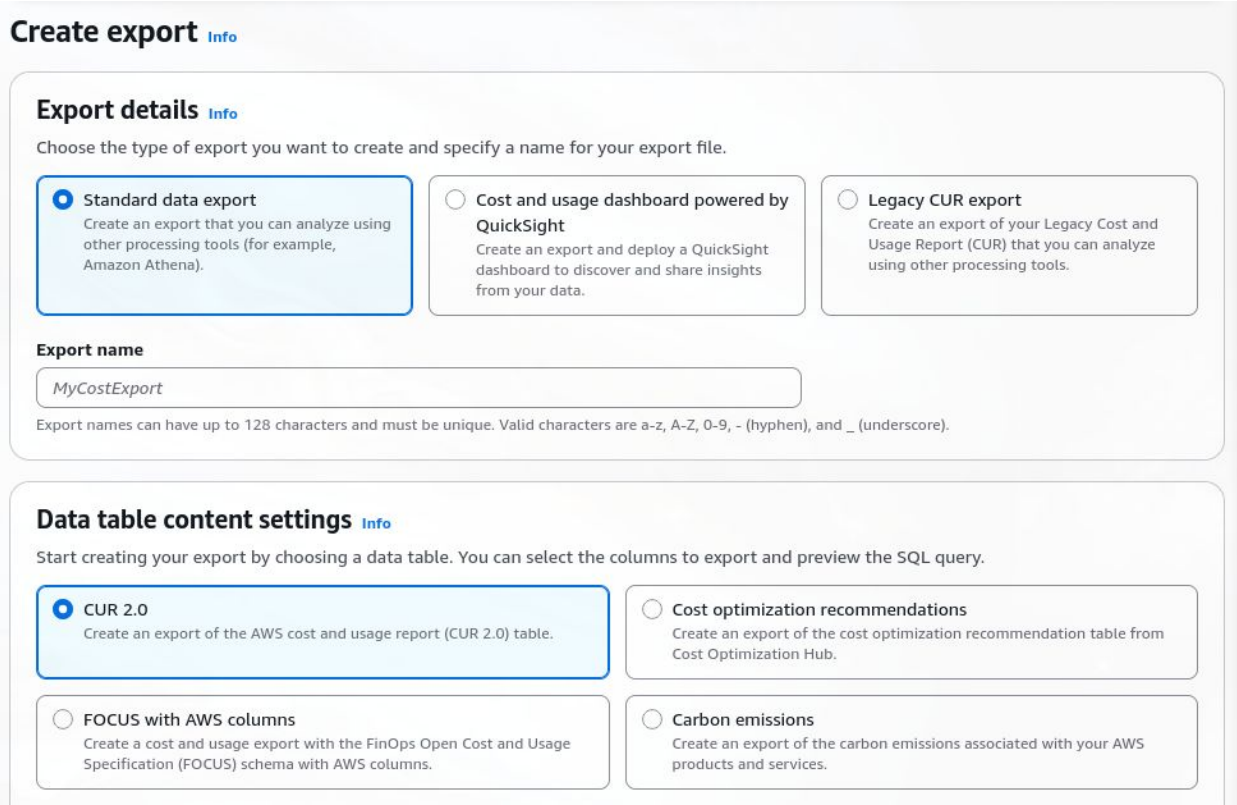

The Closest Thing to A Free Lunch

Commitment Pricing: AWS

- Compute Savings Plans
  - The most flexibility
  - Applies across any EC2 family, region, and Fargate
  - Up to 66% off for 3-year terms
- EC2 Instance Savings Plans
  - Scoped to a specific instance family and region
  - Offer deeper discounts
- Reserved Instances
  - Older model
  - Still useful for capacity reservations
  - Can apply to non EC2-services like RDS, ElastiCache, OpenSearch, Redshift
  - Can be bought and sold on the marketplace

Commitment Pricing: AWS

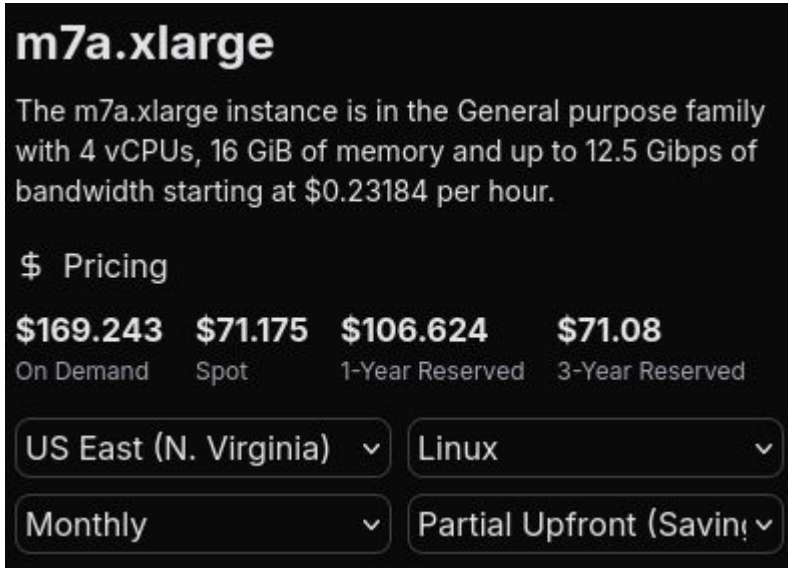

Spot	58%
3-Yr Reserved	58%
1-Yr Reserved	37%

Which provider discounts do you use?

AWS

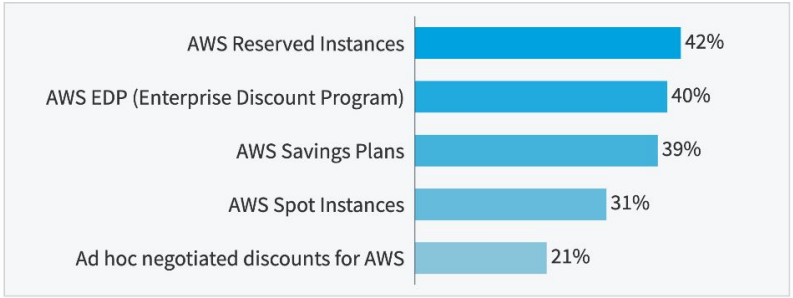

Microsoft Azure

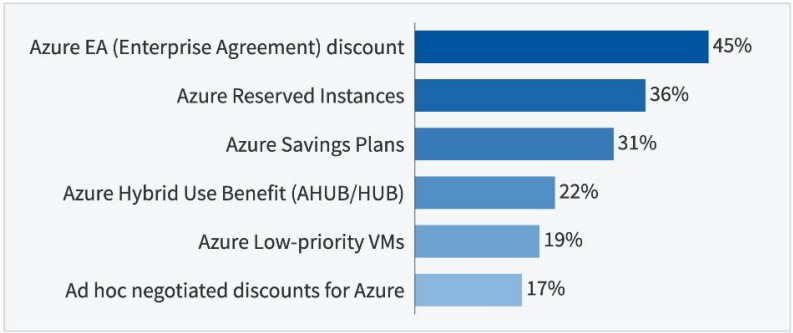

Google Cloud Platform

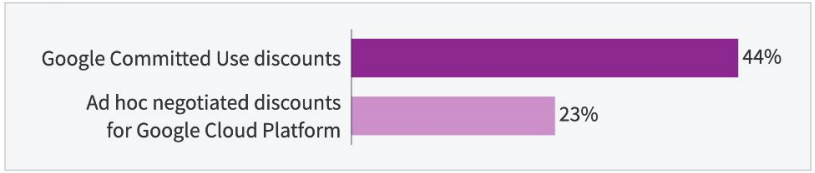

IBM Cloud

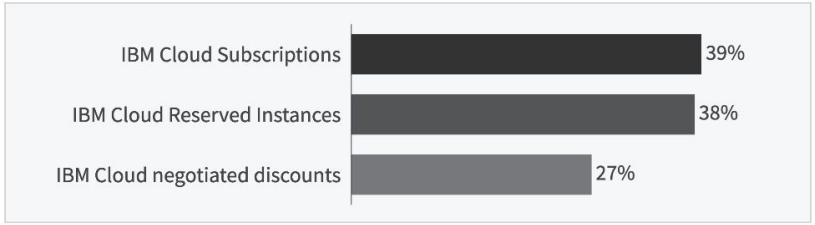

Oracle Cloud Infrastructure

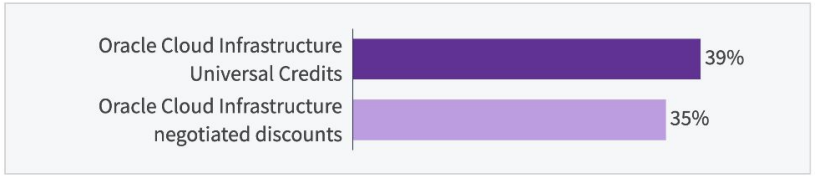

N=759

Source: Flexera 2025 State of the Cloud Report (Figure 28)

flexera

Commitment Pricing Drawbacks

- Lock-in
  - Make sure you buy the correct plans, some are non-refundable
  - How static are your workflows?
  - How good is your forecasting?

Do I Really Need AWS/Azure/GCP?

“Budget” Clouds

- DigitalOcean
- Hetzner Cloud
- OVHcloud
- Vultr
- Linode

DigitalOcean Pricing

CPU Options | **Regular** | Premium Intel | Premium AMD

Memory	vCPU	Transfer	SSD	$/hr	$/mo
512 MiB	1 vCPU	500 GiB	10 GiB	$0.00595	$4.00
1 GiB	1 vCPU	1,000 GiB	25 GiB	$0.00893	$6.00
2 GiB	1 vCPU	2,000 GiB	50 GiB	$0.01786	$12.00
2 GiB	2 vCPUs	3,000 GiB	60 GiB	$0.02679	$18.00
4 GiB	2 vCPUs	4,000 GiB	80 GiB	$0.03571	$24.00
8 GiB	4 vCPUs	5,000 GiB	160 GiB	$0.07143	$48.00
16 GiB	8 vCPUs	6,000 GiB	320 GiB	$0.14286	$96.00

Hetzner Cloud Pricing

	VCPU	RAM	SSD	Price	
CX23	2 Intel®/AMD	4 GB	40 GB	max. $ 5.59 /month	$ 0.0090 /hour
CAX11	2 Ampere®	4 GB	40 GB	max. $ 6.09 /month	$ 0.0098 /hour
CX33	4 Intel®/AMD	8 GB	80 GB	max. $ 8.59 /month	$ 0.0138 /hour
CAX21	4 Ampere®	8 GB	80 GB	max. $ 10.09 /month	$ 0.0162 /hour
CX43	8 Intel®/AMD	16 GB	160 GB	max. $ 14.59 /month	$ 0.0234 /hour
CAX31	8 Ampere®	16 GB	160 GB	max. $ 19.09 /month	$ 0.0306 /hour
CX53	16 Intel®/AMD	32 GB	320 GB	max. $ 27.09 /month	$ 0.0435 /hour
CAX41	16 Ampere®	32 GB	320 GB	max. $ 37.59 /month	$ 0.0603 /hour

excl. VAT.

Serverless Frameworks

- Fly.io
- Vercel
- Railway
- Cloudflare
- Netlify

DBaaS

- MongoDB Atlas
- Supabase
- Neon
- CockroachDB
- Aiven
- PlanetScale
- ScyllaDB
- Upstash

Dedicated Servers

- Hetzner
- OVH
- Leaseweb
- Vultr Bare Metal
- https://lowendtalk.com (for the brave)

The Top 5 Things To Do on Monday

The Top 5 Things To Do on Monday

- Delete unused Volumes/Snapshots/AMIs
- Make sure you are using VPC Gateway Endpoints for DynamoDB and S3
- Enable S3 lifecycle policies and Cloudwatch retention
- Buy a Savings Plan if you haven’t already
- Create an AWS Data Export and throw it into an LLM with python & pandas