UNVARNISHED REVIEWS

Amazon Aurora vs. Google Cloud Spanner vs. Azure SQL Database: The HA Multiplier

Market Position: Big Picture

Amazon Aurora, Google Cloud Spanner, and Azure SQL Database are the three most-compared managed relational database services from the major cloud hyperscalers, the default choice for organizations already committed to one cloud provider and evaluating that provider's flagship database offering. This report covers a narrower, more specific question than this site's broader AWS vs. Azure vs. Google Cloud report, which models total multi-service cloud TCO; this report is specifically about the managed relational database layer, and the real cost differences between these three that a multi-service comparison doesn't surface.

All three are credible, well-reviewed services: Aurora and Azure SQL Database both lead with MySQL/PostgreSQL or SQL Server compatibility and instance-based scaling, while Spanner leads with horizontal, globally-distributed strong consistency at the architectural cost of requiring more infrastructure for the same redundancy guarantee. The finding that should change how any of the three gets budgeted: a fair comparison isn't about per-unit rate, it's about what each platform's architecture requires for production-grade high availability, and Spanner's requirement is structurally larger than the other two's.

Platform Ratings at a Glance

Platform	G2 Rating	Reviews (G2)	Note
Amazon Aurora	4.5 / 5	~491	Largest review base of the three
Azure SQL Database	4.5 / 5	~241	Enterprise-skewed reviewer base (39.6%-43.2%)
Google Cloud Spanner	4.3 / 5	~44	Smallest review base; enterprise-skewed (43.2% of reviews)

Aurora and Azure SQL Database rate identically in user ratings, with Spanner trailing slightly behind and carrying the smallest review sample of the three. Note on sourcing for this comparison specifically: unlike most reports on this site, no independent transaction-benchmark data (Vendr, CostBench, or similar) exists for these three platforms, since all three bill as consumption within an organization's existing AWS, Azure, or Google Cloud account rather than as a separately negotiated SaaS contract. Pricing in this report is sourced from each vendor's published rate cards.

Amazon Aurora: Instance-Based, With a Serverless Option

Aurora prices two ways. Provisioned instances are billed hourly by instance class, a representative mid-tier instance (db.r6i.large) runs approximately $0.29/hour. Aurora Serverless v2 bills per Aurora Capacity Unit (ACU) at roughly $0.06/ACU-hour, scaling automatically with load. Storage is billed separately at approximately $0.10/GB-month, and I/O requests at roughly $0.20 per million requests, a real, usage-scaling cost dimension that a flat per-GB storage rate doesn't capture.

Aurora's standard high-availability pattern uses two instances, a primary and a replica, across Availability Zones within a single region. This is Aurora's architectural baseline for production redundancy, and it's the comparison point that matters most against Spanner's very different minimum footprint.

Azure SQL Database: vCore-Based, With the Widest Range of Purchasing Models

Azure SQL Database prices per vCore, provisioned rates run approximately $0.50/hour for 2 vCores, scaling to roughly $2.52/hour at 10 vCores in representative published rate tables. A serverless compute tier is also available, billing around $0.378/vCore-hour with automatic pause-and-resume for intermittent workloads, and Azure SQL uniquely among the three offers a free tier for small workloads.

Reviewers consistently flag two things about Azure SQL Database: strong ease-of-administration and setup scores relative to competitors, and frequent complaints about cost and pricing complexity once real workloads scale beyond the entry tier, a pattern independent review aggregation specifically calls out as one of the most-mentioned negatives across verified reviews.

Google Cloud Spanner: Globally Consistent, With a Structurally Larger Minimum Footprint

Spanner prices by Processing Units (PUs) and storage: 100 PUs run approximately $0.10/hour, and 1,000 PUs (equivalent to one full node) run approximately $1/hour, roughly $720/month per node running continuously. Storage is billed at approximately $0.30/GB-month, notably higher than Aurora's or Azure SQL's per-GB storage rate.

The architectural difference that matters most for budgeting: Spanner's horizontally-distributed, globally-consistent design means production-grade redundancy commonly requires three or more nodes, not the two-instance pattern Aurora and Azure SQL use for HA. A minimum viable production Spanner deployment at three nodes runs approximately $2,160/month, before storage or network costs, a real, structural consequence of Spanner's architecture, not a pricing markup. This tradeoff buys genuinely different capability, Spanner's multi-region strong consistency has no equivalent in Aurora or Azure SQL without significant additional application-level engineering, but it means Spanner is the wrong platform to evaluate on a simple per-instance price comparison.

Pricing (July 2026)

Platform	Compute Basis	Representative Rate	Storage
Amazon Aurora	Instance-hourly (provisioned) or ACU-hourly (Serverless)	~$0.29/hour (db.r6i.large); ~$0.06/ACU-hour (Serverless)	~$0.10/GB-month
Azure SQL Database	vCore-hourly (provisioned or Serverless)	~$0.50/hour (2 vCore) to ~$2.52/hour (10 vCore)	Included in tier; varies by service level
Google Cloud Spanner	Processing Unit-hourly (node-based)	~$1/hour per 1,000 PU (1 node, ~$720/month)	~$0.30/GB-month

TCO Comparison: Minimum Production HA Deployment, 1 Month

This scenario models each platform at its own architectural minimum for high availability, since the three don't require the same infrastructure footprint for comparable redundancy. Spanner is shown at two configurations, since its per-node rate itself depends on regional scope, a distinction independent competitive-pricing analysis (Vendr, via Cockroach Labs' own published comparison data) makes explicit but that a single headline Spanner rate can obscure.

Platform	Minimum HA Topology	Monthly Compute Cost
Azure SQL Database	2 vCore provisioned tier	~$368
Amazon Aurora	2 instances (primary + replica), db.r6i.large	~$423
Google Cloud Spanner (regional)	3 nodes, single-region, ~$0.90-$1.80/node-hour	~$2,160

Platform	Minimum HA Topology	Monthly Compute Cost
Google Cloud Spanner (multi-region)	3 nodes, multi-region, ~$3.00-$6.00/node-hour	~$9,855

Spanner's regional minimum HA deployment runs roughly 5.1x Aurora's and 5.9x Azure SQL's comparable setup. But regional deployment isn't Spanner's differentiating use case, multi-region, globally-consistent deployment is, and at multi-region node rates the same 3-node minimum runs closer to 23x-27x Aurora's or Azure SQL's cost. Independent transaction-benchmark data provides a useful real-world anchor at larger scale: at a matched 32-vCPU, 1TB, multi-region configuration, Vendr's competitive analysis puts annual costs at roughly $40,000-$80,000 for Aurora (with Reserved Instances), $80,000-$120,000 for CockroachDB (negotiated), and $100,000-$150,000 for Spanner (with committed-use discounts), a real, verified data point showing Spanner's multi-region premium narrows substantially at scale and with negotiated commitments compared to the node-hour list-rate math alone. This is not evidence that Spanner is overpriced at either configuration, it reflects a genuinely different architecture built for a different problem than Aurora or Azure SQL solve for. The finding is that a buyer comparing Spanner's regional rate to Aurora's or Azure SQL's, when the actual requirement is multi-region consistency, will significantly underestimate Spanner's real cost, and a buyer who doesn't need Spanner's specific consistency guarantees is very likely paying for capability their workload doesn't require.

Costs the Pricing Table Misses

- **Aurora's I/O charges scale with usage, not just data size.** $0.20 per million requests sounds trivial until a high-transaction workload is modeled explicitly; this is a real, variable cost line that a flat per-GB storage estimate misses entirely.
- **Spanner's per-GB storage rate is roughly 3x Aurora's.** $0.30/GB-month against $0.10/GB-month; at meaningful data volumes, this compounds the node-count multiplier into an even larger total gap.
- **Azure SQL's cost-complexity is a documented, repeated complaint, not an isolated one.** Independent review aggregation specifically flags pricing complexity as one of the most-mentioned negatives across verified Azure SQL Database reviews; get a detailed, itemized quote rather than budgeting off the headline per-vCore rate.
- **None of the three publish negotiated enterprise discount data independently of their parent cloud commitment.** Unlike most platforms this site covers, discount leverage here runs through the broader AWS, Azure, or Google Cloud enterprise agreement, not a standalone database-specific negotiation.
- **Cross-region replication and egress are billed separately from the base compute rate at all three.** Multi-region deployments, especially relevant for Spanner's core use case, should model network costs explicitly rather than assuming the node or instance rate is the full story.

The Decision Framework

Choose Amazon Aurora if: You're already on AWS, need MySQL or PostgreSQL compatibility, and want the lowest-cost path to single-region high availability of the three.

Choose Azure SQL Database if: You're Microsoft-stack-first, want the widest range of purchasing models (provisioned, serverless, free tier for small workloads), and can budget for real cost growth as usage scales past the entry tier.

Choose Google Cloud Spanner if: Your application genuinely requires global, horizontally-scaled strong consistency, a real technical requirement, not a nice-to-have, and you can budget for a minimum three-node production footprint from the outset rather than being surprised by it after evaluating on a per-node rate alone.

Everyone: Model your actual minimum production-grade topology, not just the headline hourly rate, before comparing these three. The gap between Aurora and Azure SQL is modest at comparable HA configurations; the gap to Spanner is real and architectural, and only makes sense to pay if Spanner's specific consistency guarantees are a real workload requirement.

The Bottom Line

Amazon Aurora, Google Cloud Spanner, and Azure SQL Database are all credible, well-reviewed managed relational database services, each representing its cloud provider's flagship offering. The real finding this comparison surfaces isn't which platform is cheapest per hour, it's that Spanner's architecture requires meaningfully more infrastructure than Aurora or Azure SQL for comparable production redundancy, roughly 5x the monthly cost at regional scale, and considerably more once multi-region distribution, Spanner's actual differentiating use case, is factored in. That premium buys real, differentiated capability for workloads that need it, and is a real, avoidable overspend for workloads that don't.