Beyond monitoring

How we built an open-source, self-healing Postgres agent

Alex Francoeur

October 1, 2025

Beyond monitoring

~~How we built~~ and why we're building an open-source, self-healing Postgres agent

Alex Francoeur

October 1, 2025

Me, on a slide

- Product @ Xata
- Greater Boston
- Happily married with 2 kids
- Proud father of 6 chickens and 1 dog
- Soccer / football player, runner, snowboarder and year-round-coach
- Putzer-arounder, tinkerer, hacker (well, probably more of a vibe coder now)

Professional me, on a slide

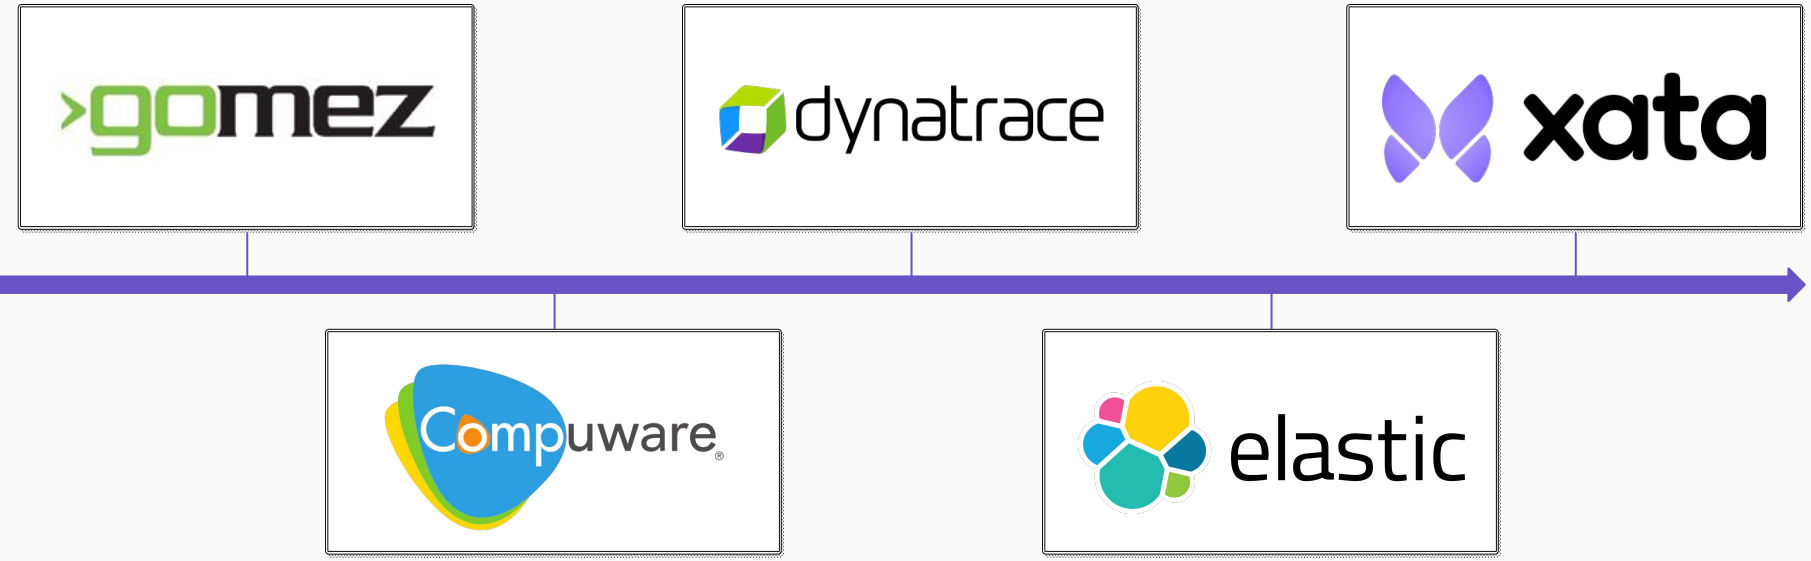

Postgres at scale

“Postgres at scale” for Xata means more than just handling large data or compute. It’s about scaling team productivity and operations.

xataio/pgroll

xataio/pgstream

xataio/agent

xata

Foundations
1960s-1980s

Monitoring relied on console lights, SMF records, and simple SNMP alerts.

DB administrators focused on mainframe RDBMS, backups, and console-based operations.

Centralization

1990s-Mid 2000s

System management platforms (HP OpenView, IBM Tivoli, BMC, CA) centralized infrastructure monitoring.

DBAs managed client-server databases with replication, clustering, and tuning.

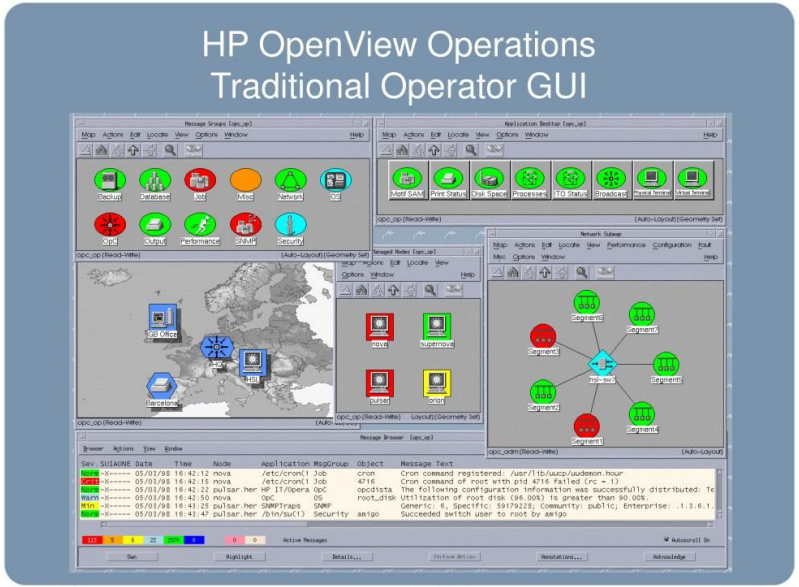

Logs & Apps

2000s

Splunk popularized log search; early APM tools tied performance to end-user experience

Ops DBAs supported web-scale apps, ensuring uptime and managing complex HA setups.

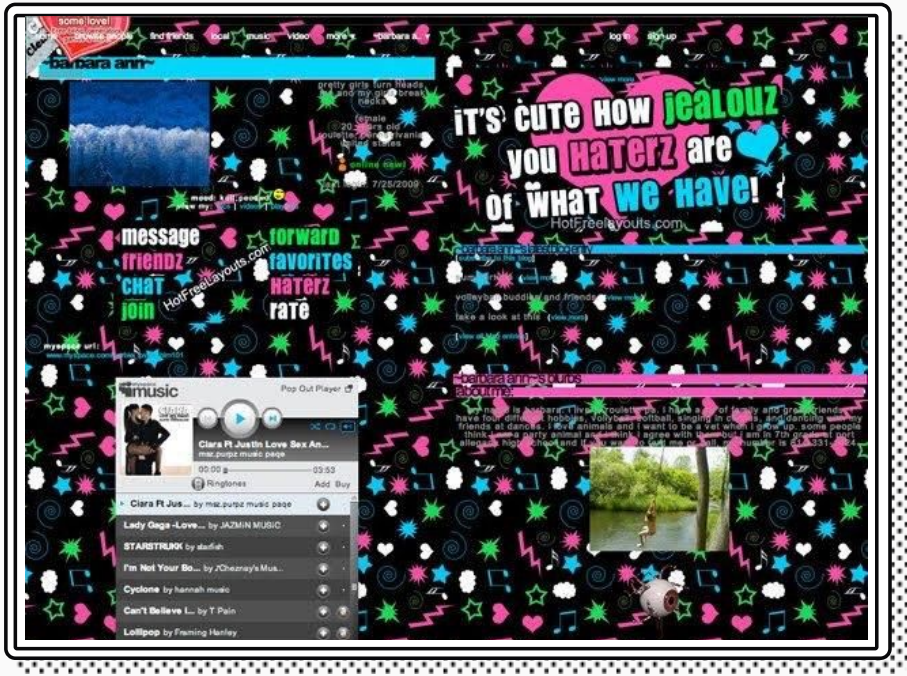

Cloud and DevOps

2010s

ELK, Prometheus, and SaaS platforms like Datadog enabled developer-centric, cloud-native monitoring.

Traditional DBA roles blurred into SRE and DevOps, with automation and CI/CD at the core.

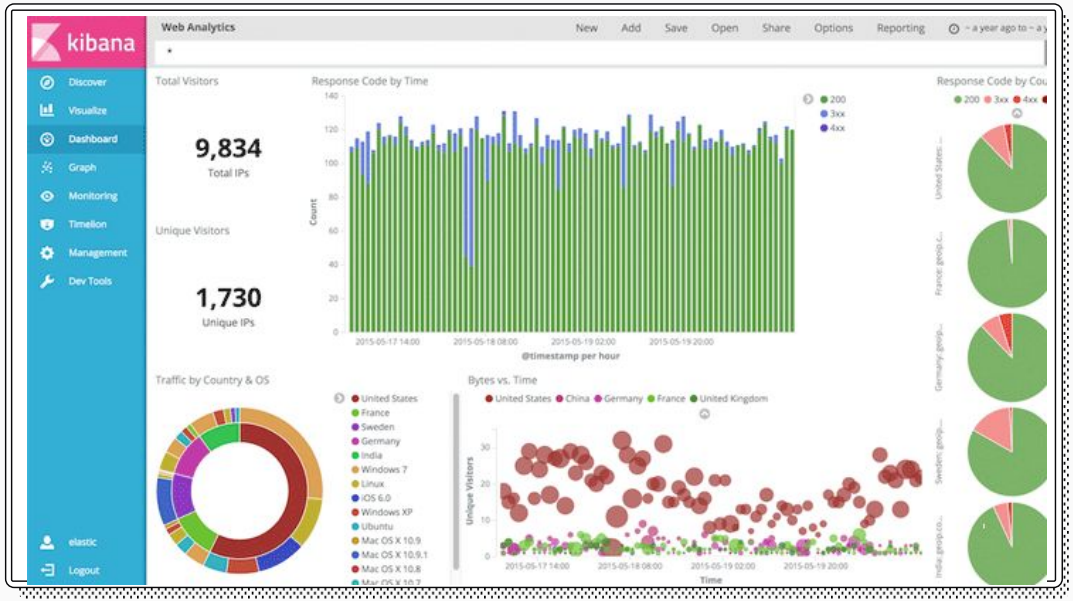

Scale & Specialization 2018+

Loki and Tempo introduced efficient storage for logs and traces to manage telemetry growth.

Platform and DataOps roles emerged, integrating database management with pipelines and analytics.

Intelligence Era

2020+

OpenTelemetry, eBPF, and ML brought predictive insights and auto-remediation into observability

DevSecOps practices embedded security and compliance into database workflows.


WAIT! NOT THIS GUY!
HE ALWAYS SAID THANK YOO!
10 YEARS AGO
ChatGPT
Hashem
Great answer. Thank you!
Grok

Monitor

Detect

Notify

Investigate

Remediate

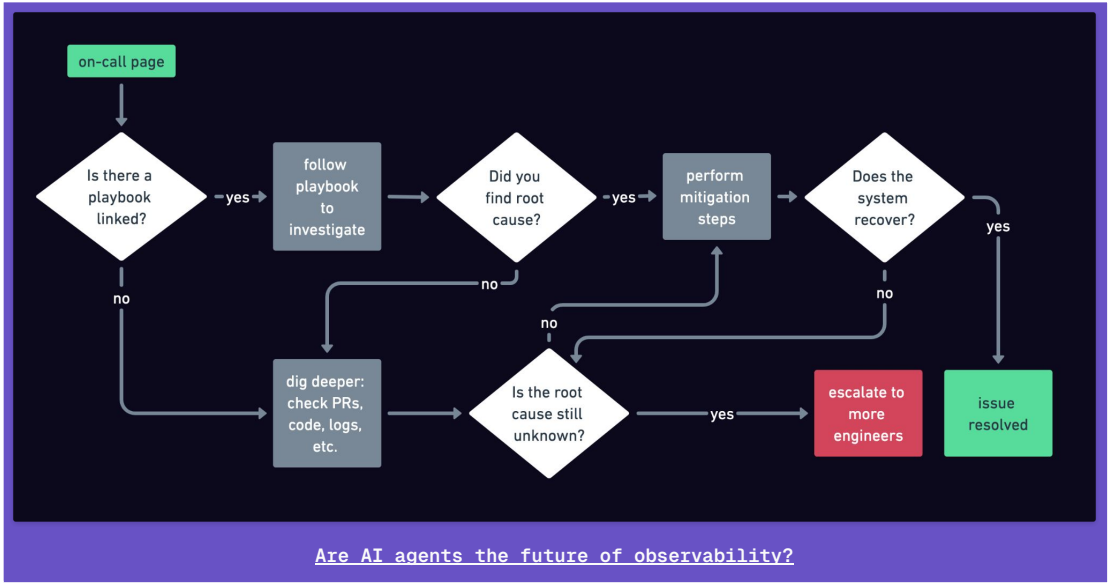

Are AI agents the future of observability?

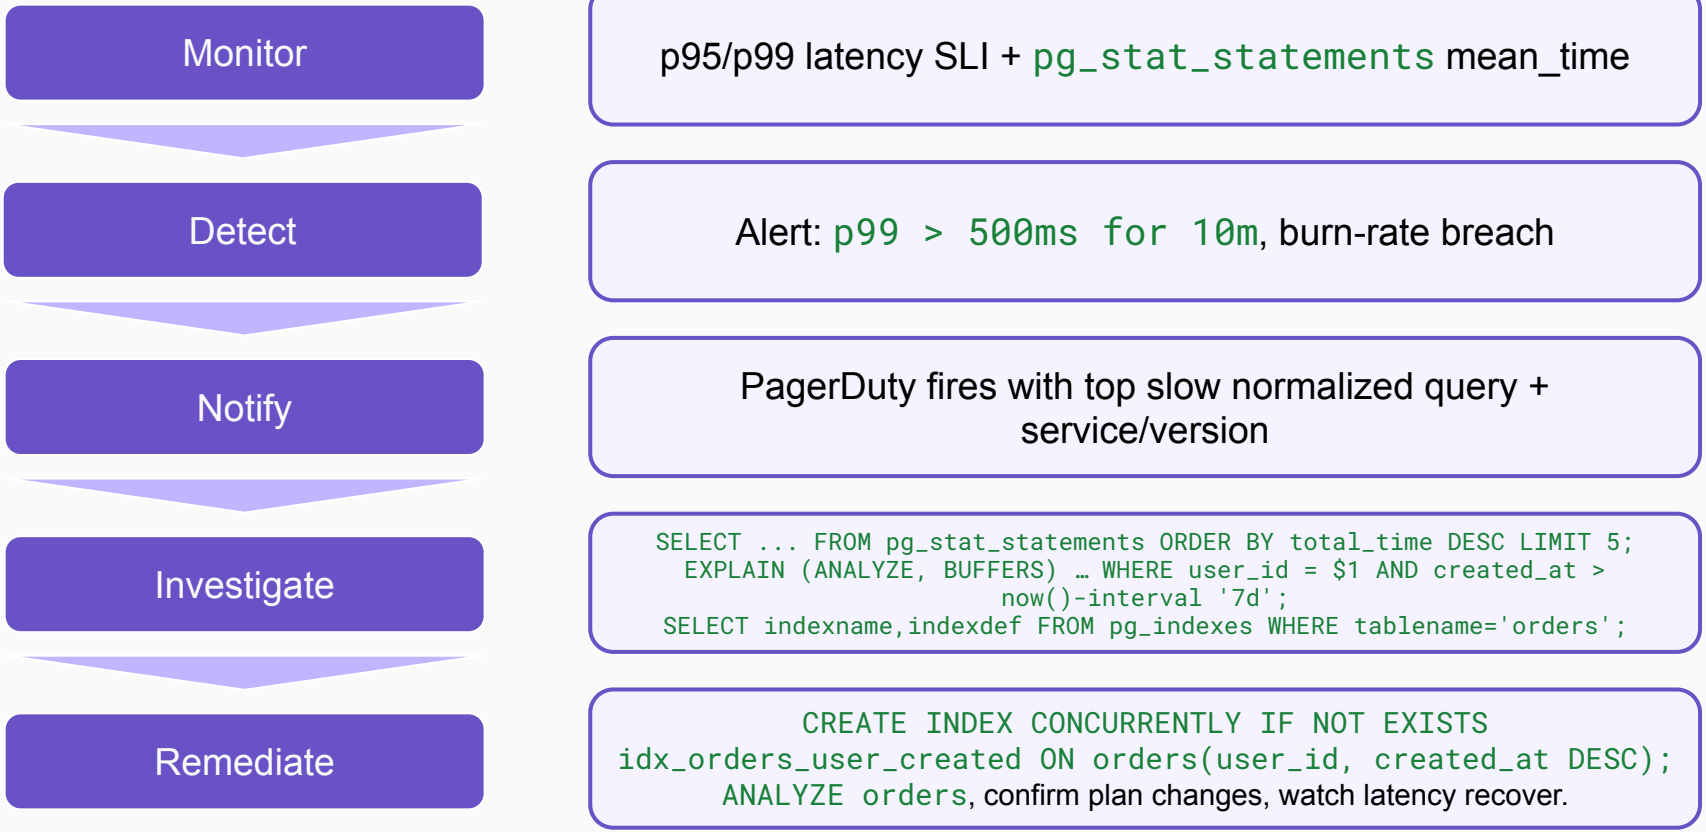
Monitor
p95/p99 latency SLI + pg_stat_statements mean_time
Detect
Alert: p99 > 500ms for 10m, burn-rate breach
Notify
PagerDuty fires with top slow normalized query + service/version
Investigate
SELECT ... FROM pg_stat_statements ORDER BY total_time DESC LIMIT 5;
EXPLAIN (ANALYZE, BUFFERS) … WHERE user_id = $1 AND created_at > now()-interval '7d';
SELECT indexname,indexdef FROM pg_indexes WHERE tablename='orders';
Remediate
CREATE INDEX CONCURRENTLY IF NOT EXISTS idx_orders_user_created ON orders(user_id, created_at DESC);
ANALYZE orders, confirm plan changes, watch latency recover.

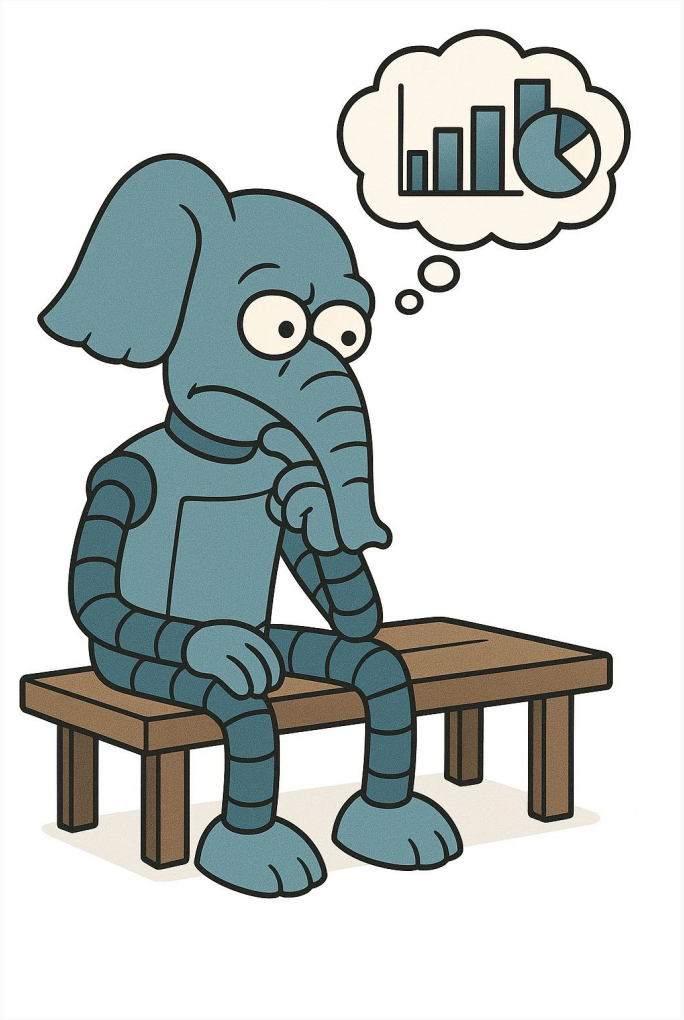

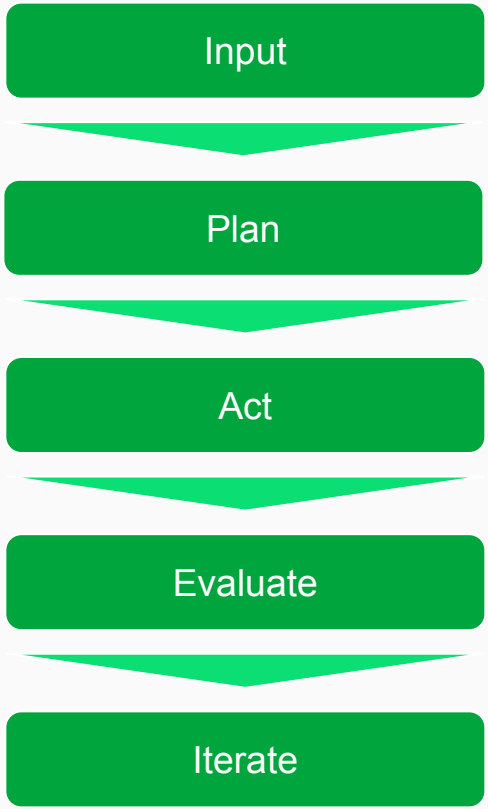

SLI breach + top query trace

Hypothesis “missing composite index”

Run `EXPLAIN`, create index concurrently, kick `ANALYZE`

Compare p95 and `mean_time` before/after, ensure error budget stabilizes

Add CI EXPLAIN gate for critical queries, tighten alert context (include `rows`, `heap_blks_read`, `idx_scan` deltas)

Microsoft

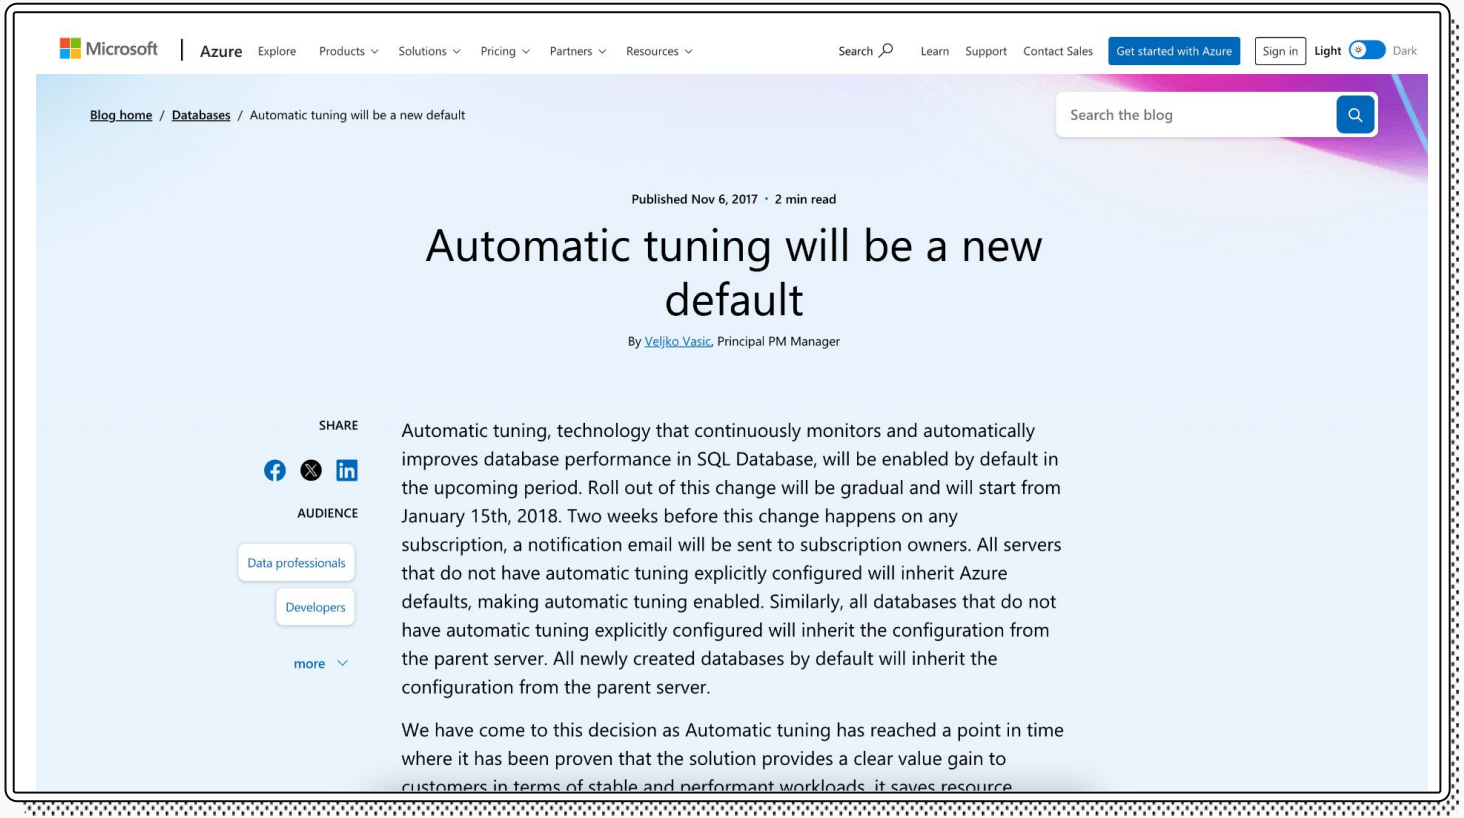

https://azure.microsoft.com/en-us/blog/automatic-tuning-will-be-a-new-default

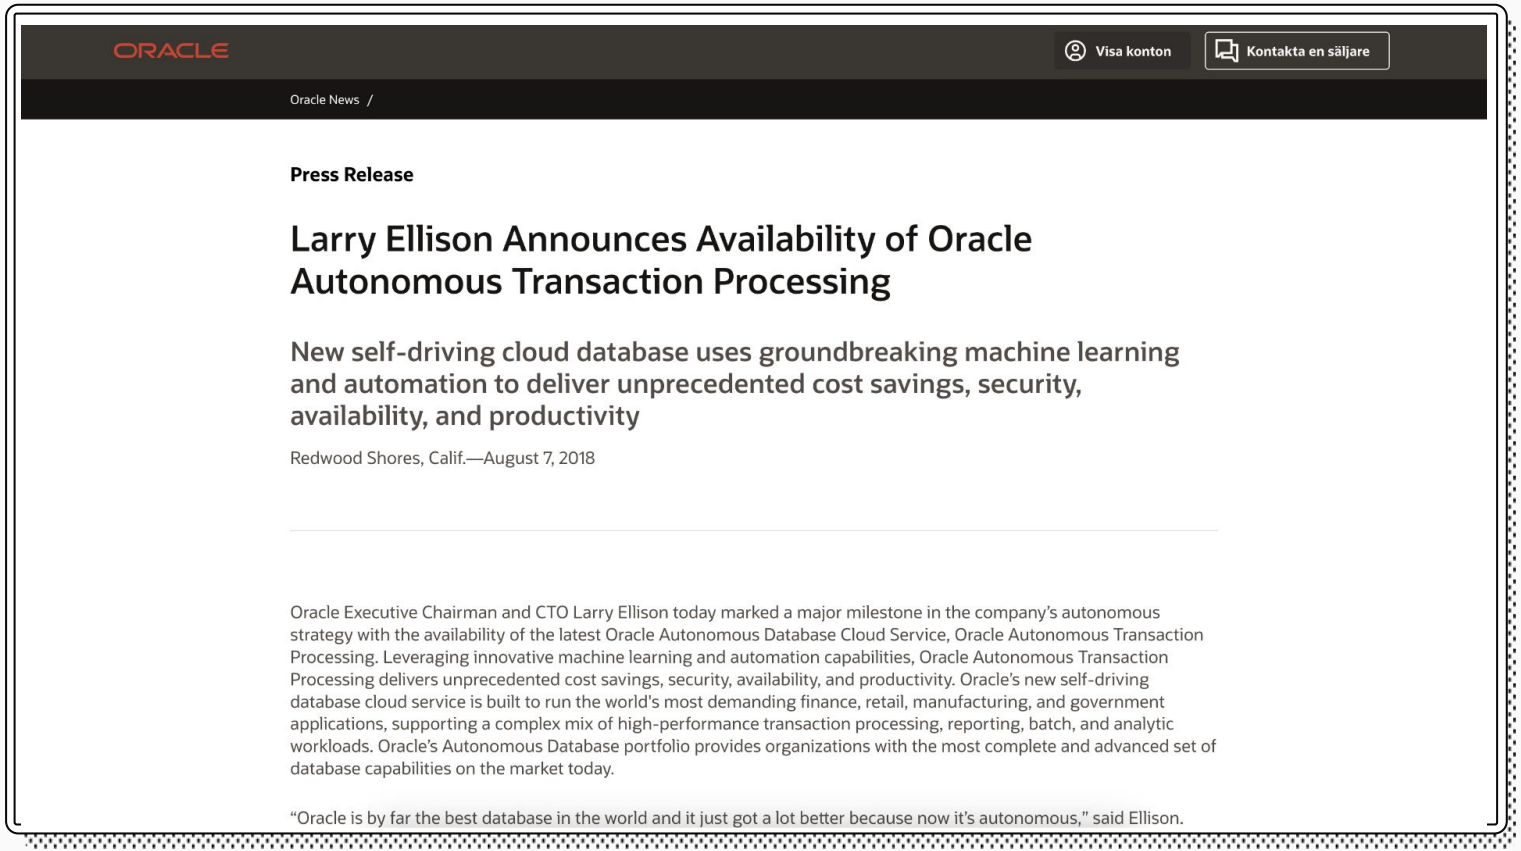

https://www.oracle.com/se/corporate/pressrelease/oracle-autonomous-transaction-processing-2018-08-07.html

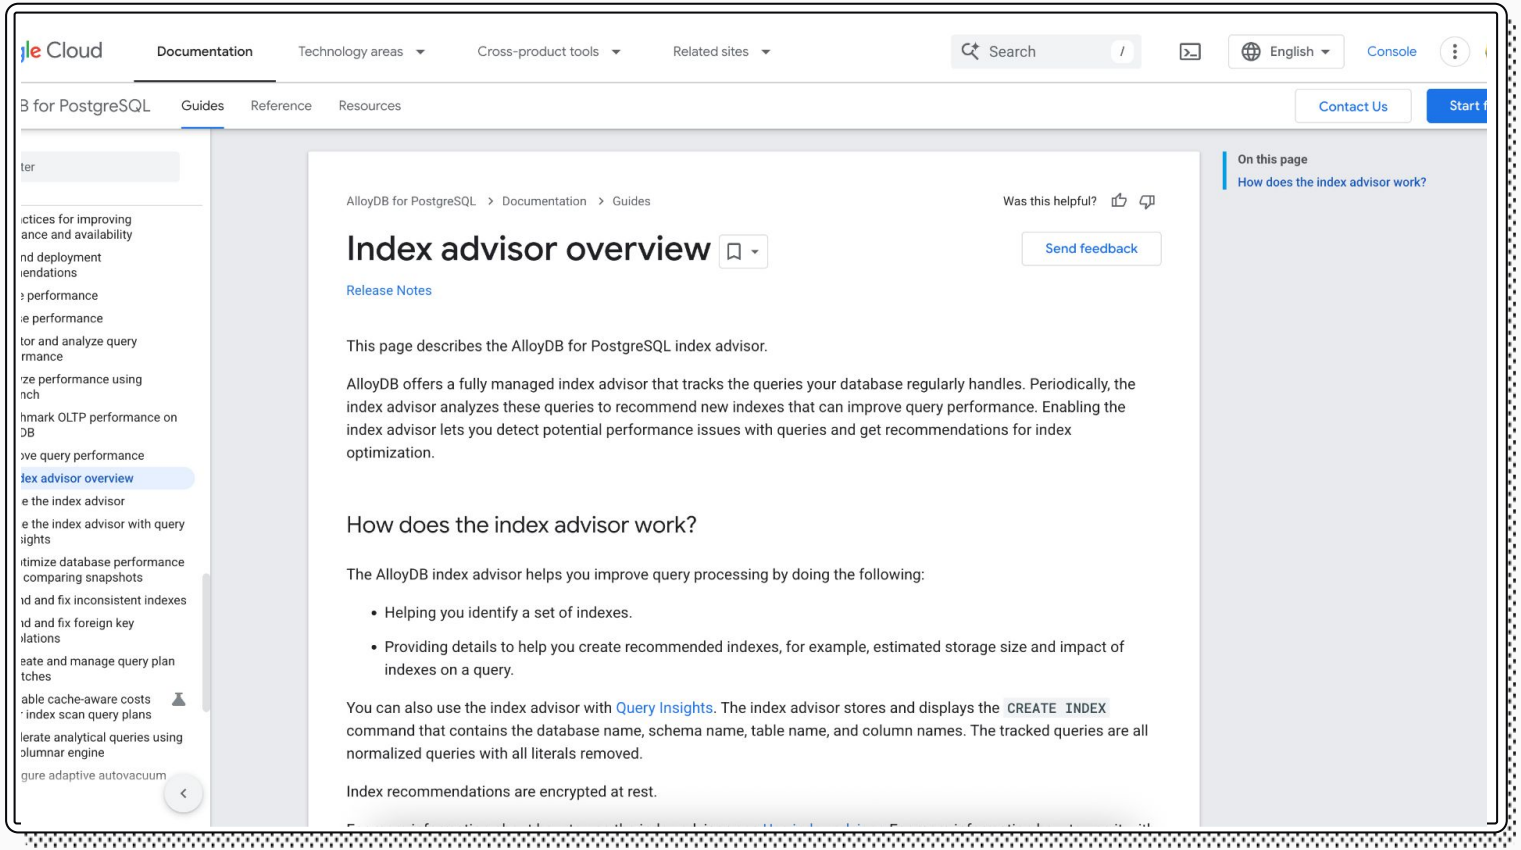

https://cloud.google.com/alloydb/docs/index-advisor-overview

supabase

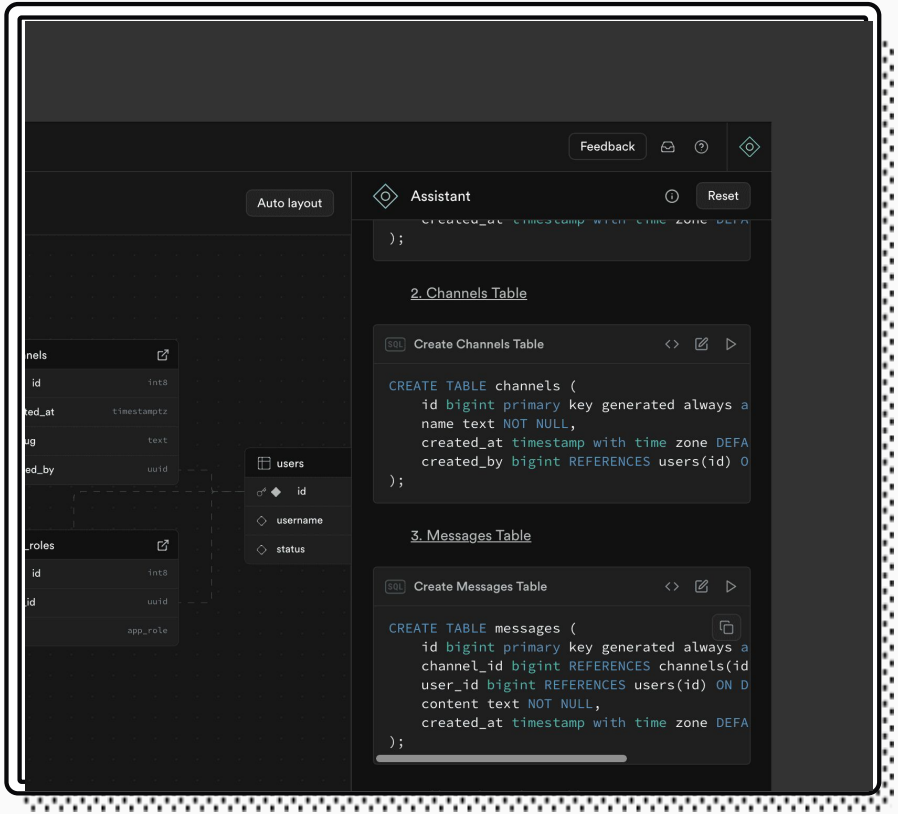

https://supabase.com/blog/supabase-ai-assistant-v2

ClickHouse

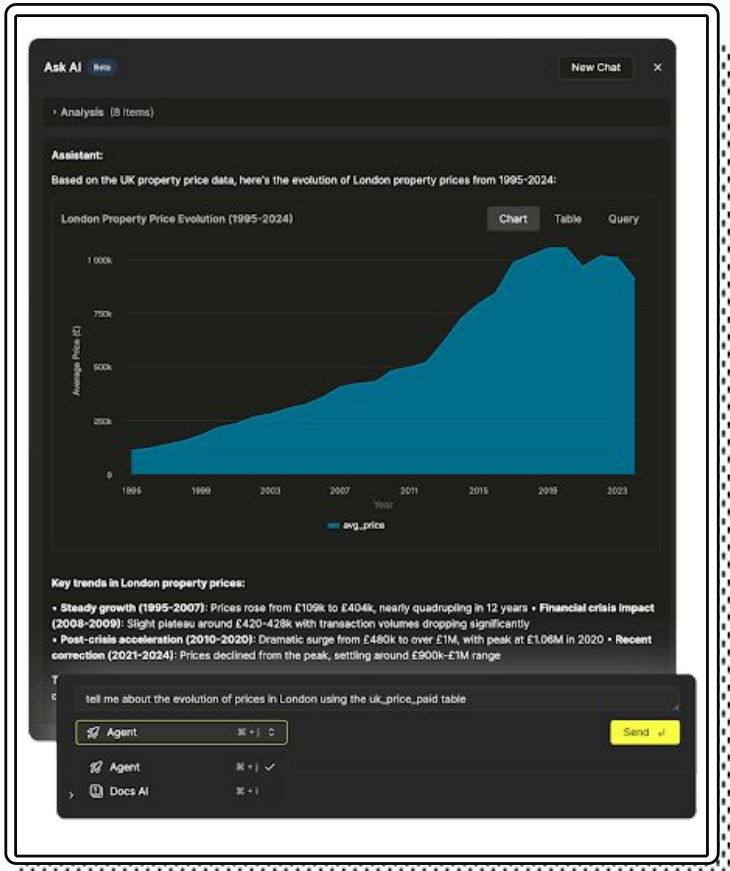

https://clickhouse.com/blog/agentic-analytics-ask-ai-agent-and-remote-mcp-server-beta-launch

PlanetScale

6. Learning, optimizing: Intelligent

The Future Database will be intelligent. It will learn and optimize processes for its developers. Imagine a database that routes queries automatically based on connection locations. Or a database that shards based on geography and will even suggest which tables to shard and where.

A future engineer will be able to optimize queries, improve database performance and even make adjustments intelligently based on recommendations from a database that learns behavior. A database designed to always do the right thing for the developer is coming.

https://planetscale.com/principles

3. **AI-powered database automation (DBA)** – There's also the idea of AI as your DBA. Most databases need manual tuning, index management, query optimization. What if an AI handled that? We're actively thinking about this—turning database management into an AI-driven system that automates maintenance, scaling, and performance tuning.

These shifts we're seeing indicate databases need to adapt. AI Agents need persistent memory—which means databases need to be faster, more elastic, and easier to integrate into automated workflows. Neon is perfectly positioned for this—serverless, autoscaling, and deeply integrated into modern dev environments.

The industry is shifting toward automation, self-healing databases, AI-driven infrastructure, and AI-assisted development. Neon is at the center of it all.

https://m12.vc/news/founders-feature-neon/

Integrations

Postgres

Connect PostgreSQL databases to Tembo for performance monitoring and optimization.

Features

- Monitor slow queries and missing indexes
- Detect unused indexes

https://postgres.ai/blog/20250725-self-driving-postgres

PostgresAI

Nikolay Samokhvalov
CEO & Founder

I'm excited to announce that Postgres AI has started work on a new project – open-source Self-Driving Postgres (SDP).

In the AI era, Postgres is the natural choice for AI builders. With fast-growing database clusters, the highest level of automation is essential. AI-driven growth demands efficient, proactive, and intelligent database management. Our goal is to reduce manual interventions as much as possible to achieve the highest level of operational efficiency and reliability.

How can we define levels of automation?

For self-driving cars, there is a widely used approach – SAE J3016 Automation Levels. We can apply a similar methodology to each area of database operations:

Automation levels (SAE J3016-inspired, simplified)

Level	Name	Description
0	No Automation	Fully manual management
1	DBA Assistance	Recommendations provided, DBA action
2	Partial Automation	Basic tasks automated, DBA oversight
3	Conditional Automation	Autonomous within boundaries
4	High Automation	Predictive and proactive optimization
5	Full Automation	Complete autonomy

https://postgres.ai/blog/20250725-self-driving-postgres

From DBA to DB Agent

https://xata.io/blog/dba-to-db-agent

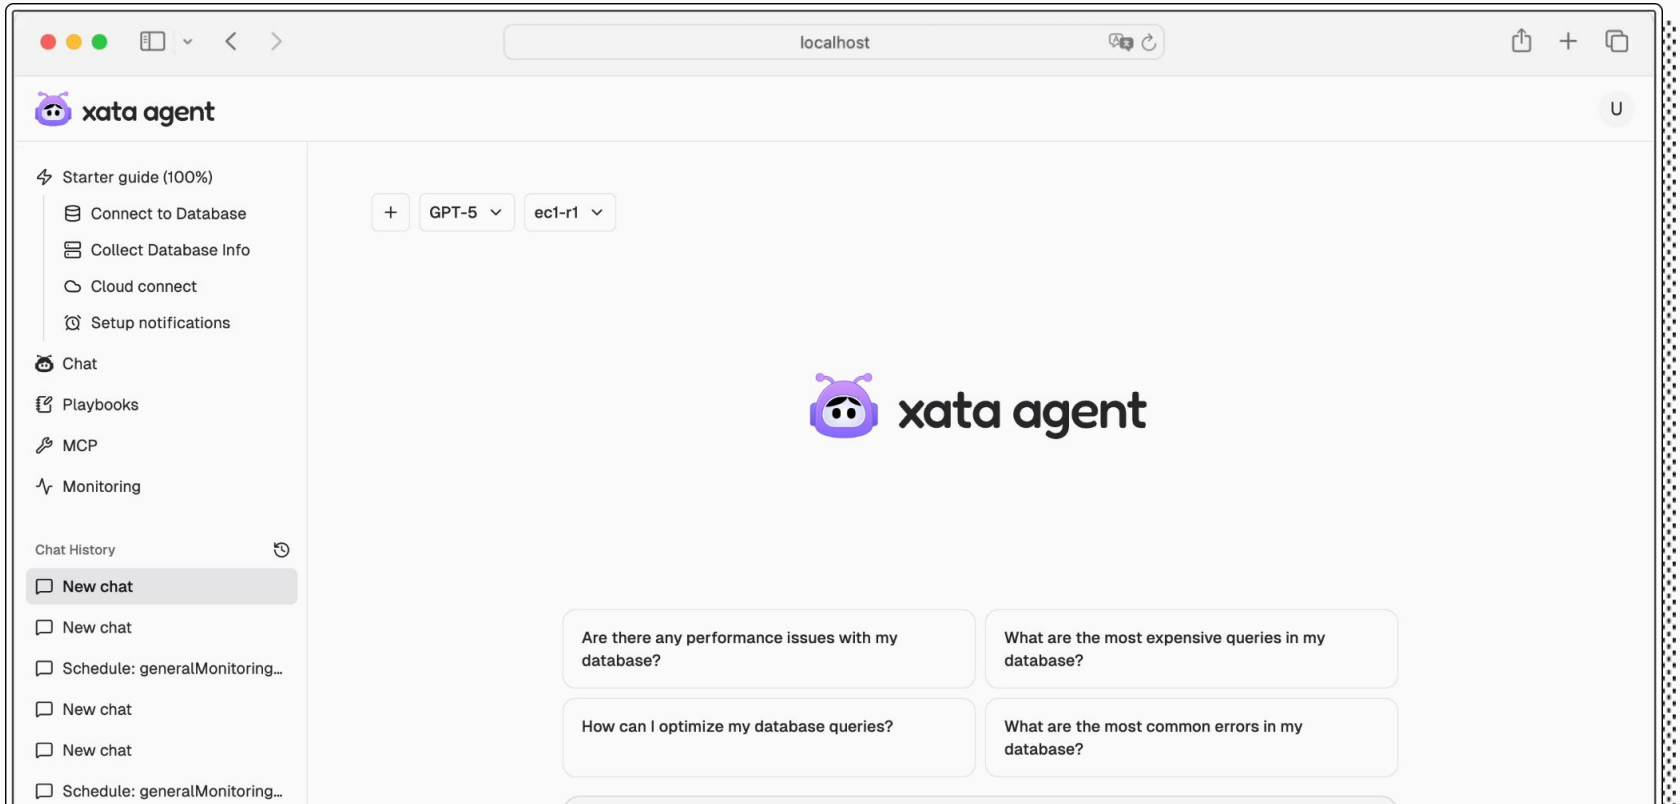

Tier 1 Triage

- CPU utilization
- Memory consumption
- Performance issues
- Query optimization
- Resource utilization

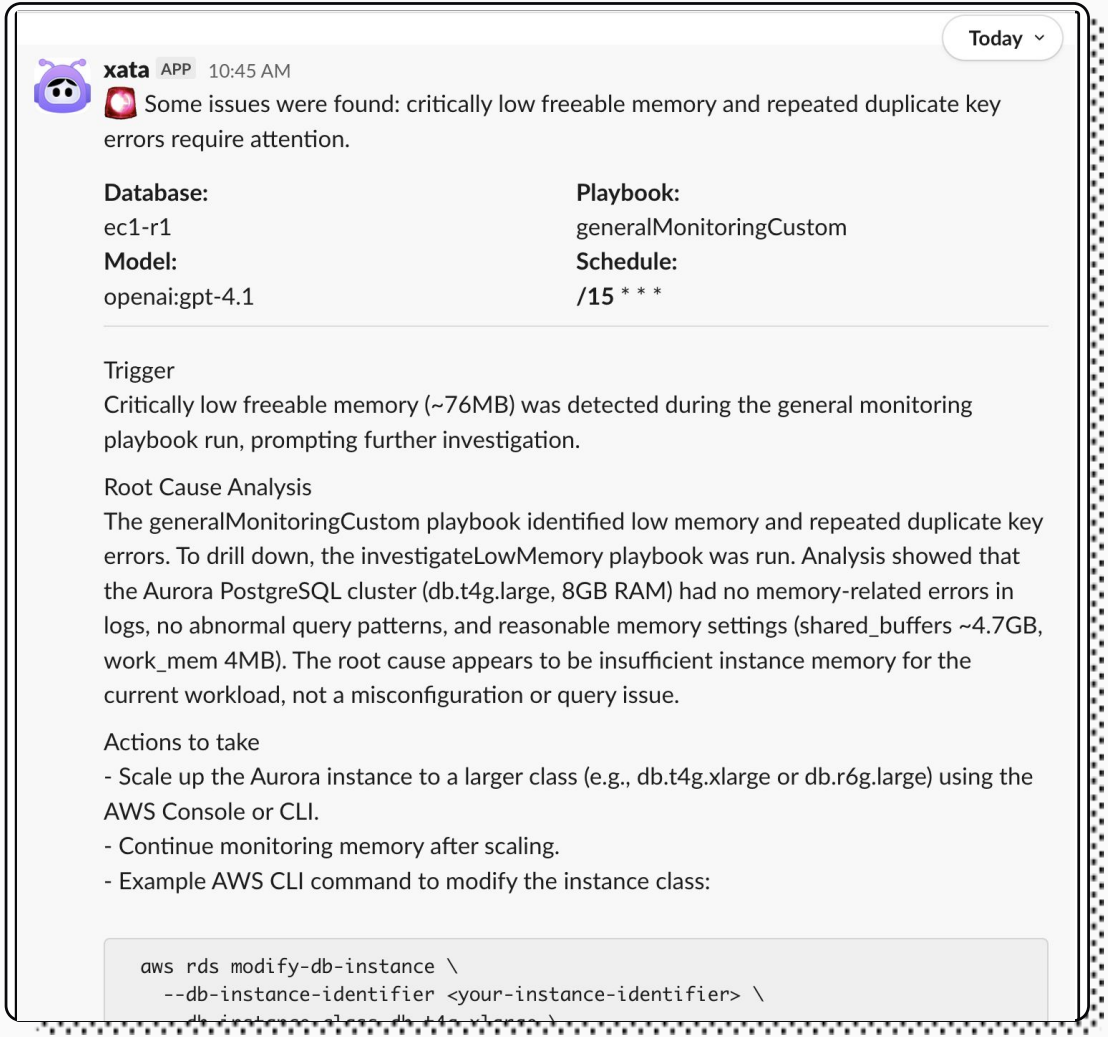

xata

Walk through

xata

Building the Agent

Xata Agent

xataio/agent

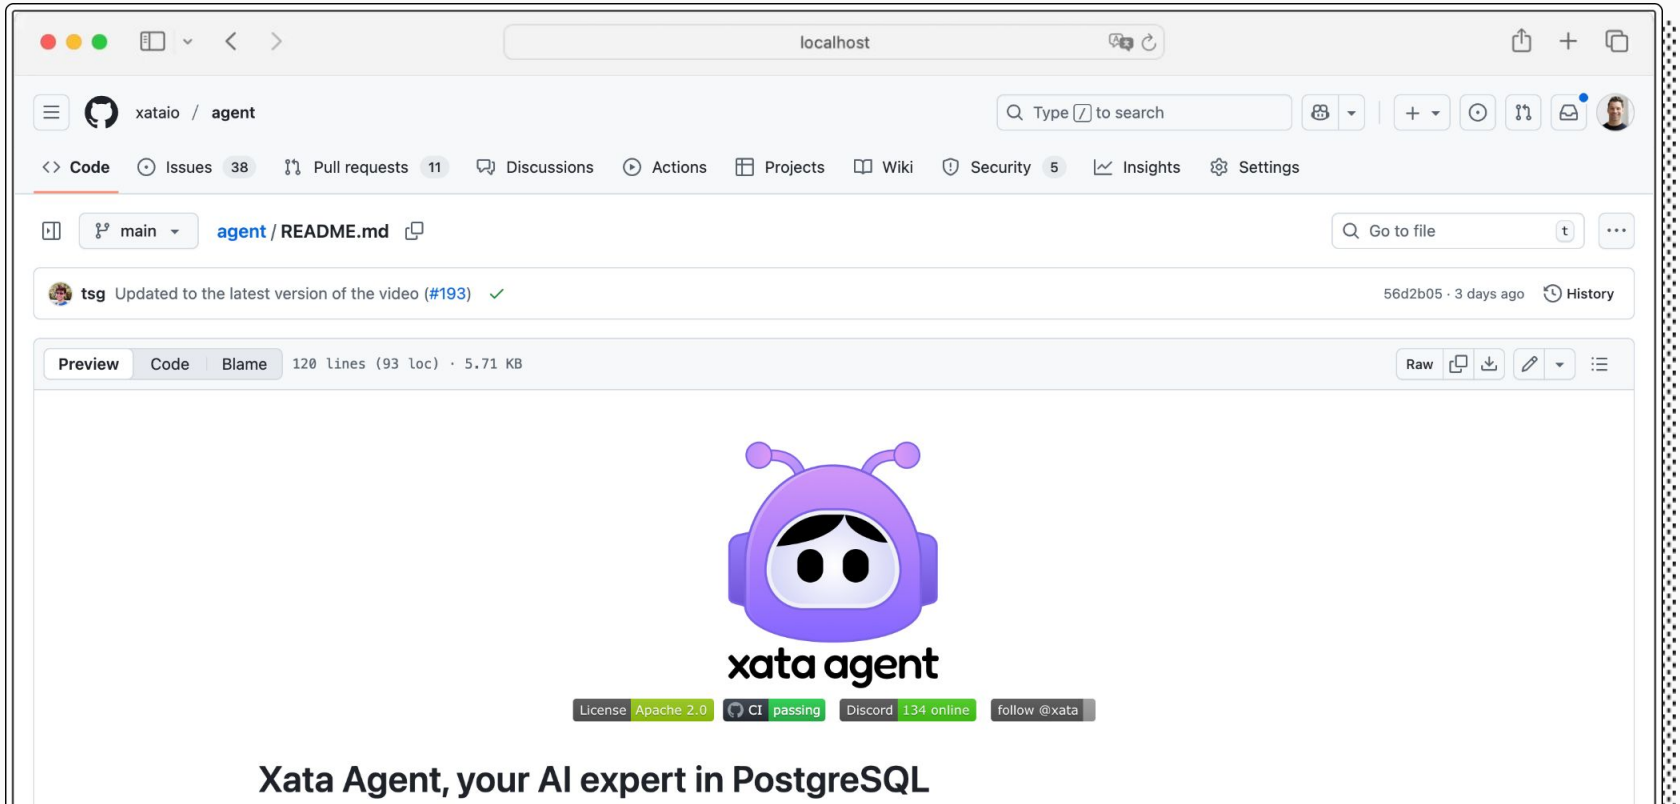

Tech Stack

- Next.js
- Postgres
- TypeScript
- Vercel AI SDK
- MCP TypeScript SDK
- Self-host via docker-compose
- Mastra TypeScript AI Framework
- Ollama for local models
- LangFuse for observability
- LiteLLM as AI API gateway

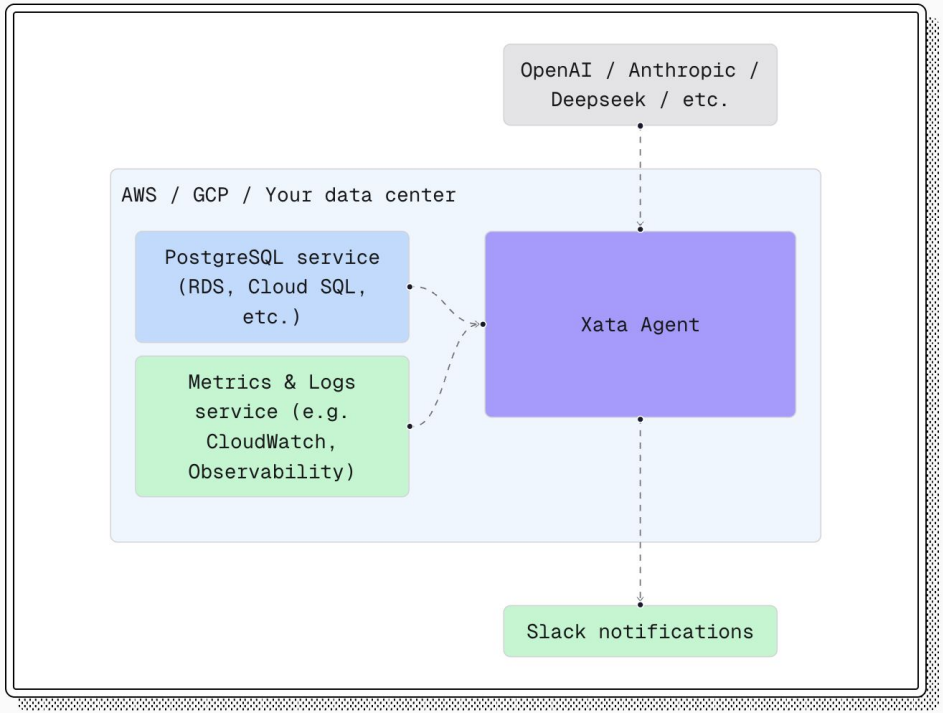

xata

Schedule

Run playbooks at scheduled or agent defined intervals

Playbooks

Sequence of steps that the agent can follow to detect, diagnose, and resolve issues

Tools

Functions that can be called by the agent to gather more context and act

Approve

Workflows to approve recommendations from the agent

xata

Schedule

Input

Schedules / Monitors

- Execute playbooks periodically
- Allows chaining of playbooks:
  - Based on the findings of a playbook, “drill down” using another playbook.
- Keep history of runs and outcomes
  - This will become the agent “memory”

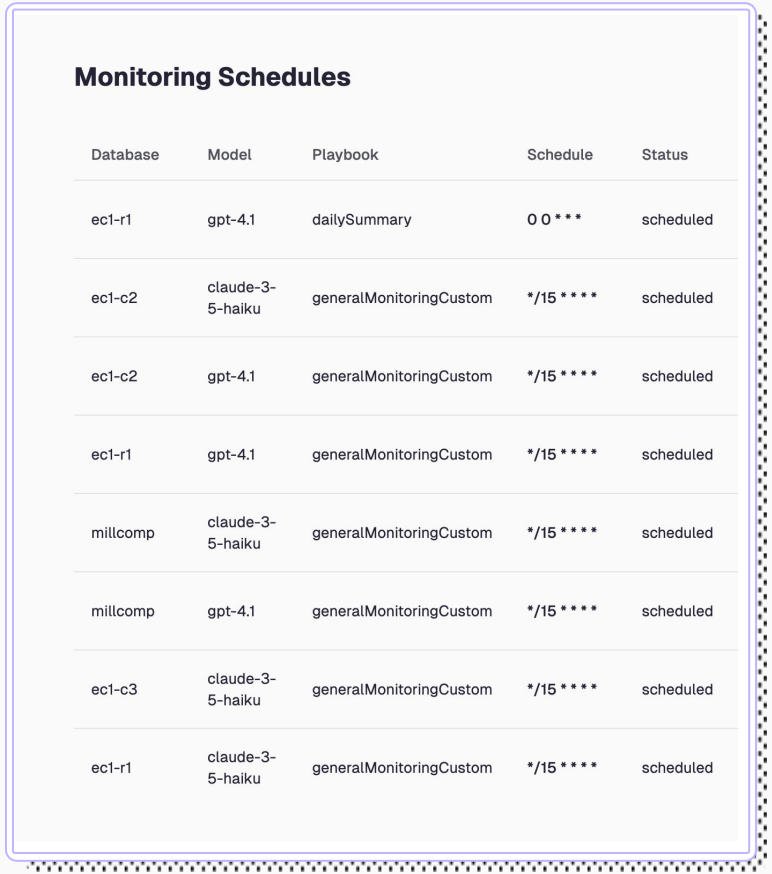

Monitoring Schedules

Database	Model	Playbook	Schedule	Status
ec1-r1	gpt-4.1	dailySummary	0 0 * * *	scheduled
ec1-c2	claude-3-5-haiku	generalMonitoringCustom	*/15 * * * *	scheduled
ec1-c2	gpt-4.1	generalMonitoringCustom	*/15 * * * *	scheduled
ec1-r1	gpt-4.1	generalMonitoringCustom	*/15 * * * *	scheduled
millcomp	claude-3-5-haiku	generalMonitoringCustom	*/15 * * * *	scheduled
millcomp	gpt-4.1	generalMonitoringCustom	*/15 * * * *	scheduled
ec1-c3	claude-3-5-haiku	generalMonitoringCustom	*/15 * * * *	scheduled
ec1-r1	claude-3-5-haiku	generalMonitoringCustom	*/15 * * * *	scheduled

xata

Playbooks

Plan

Playbooks

- Playbook examples:
  - generalMonitoring
  - investigateHighCPU
  - investigateLowMemory
  - investigateHighConnectionCount
  - tuneSettings
  - dailySummary

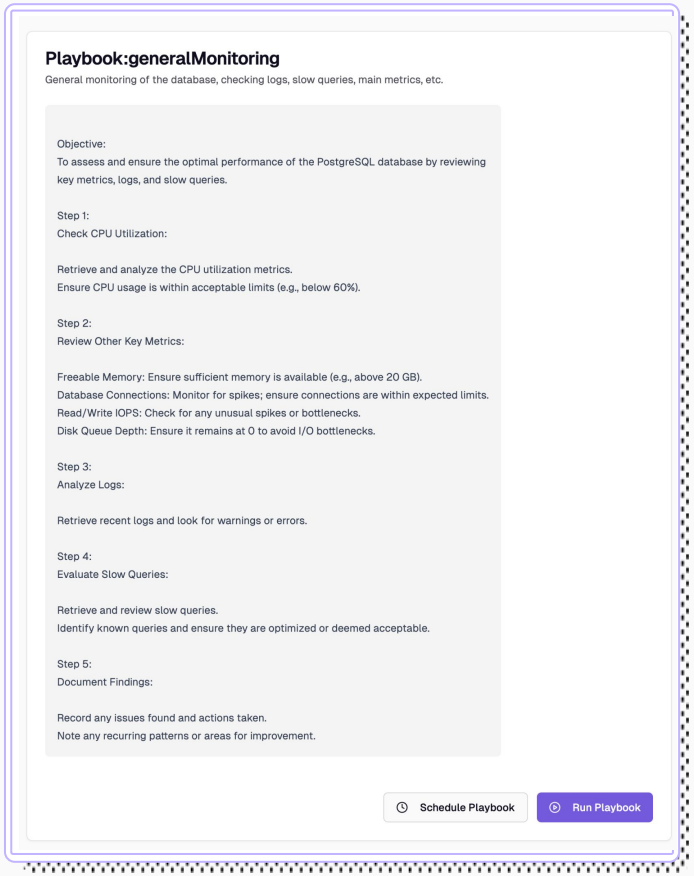

Playbooks are editable

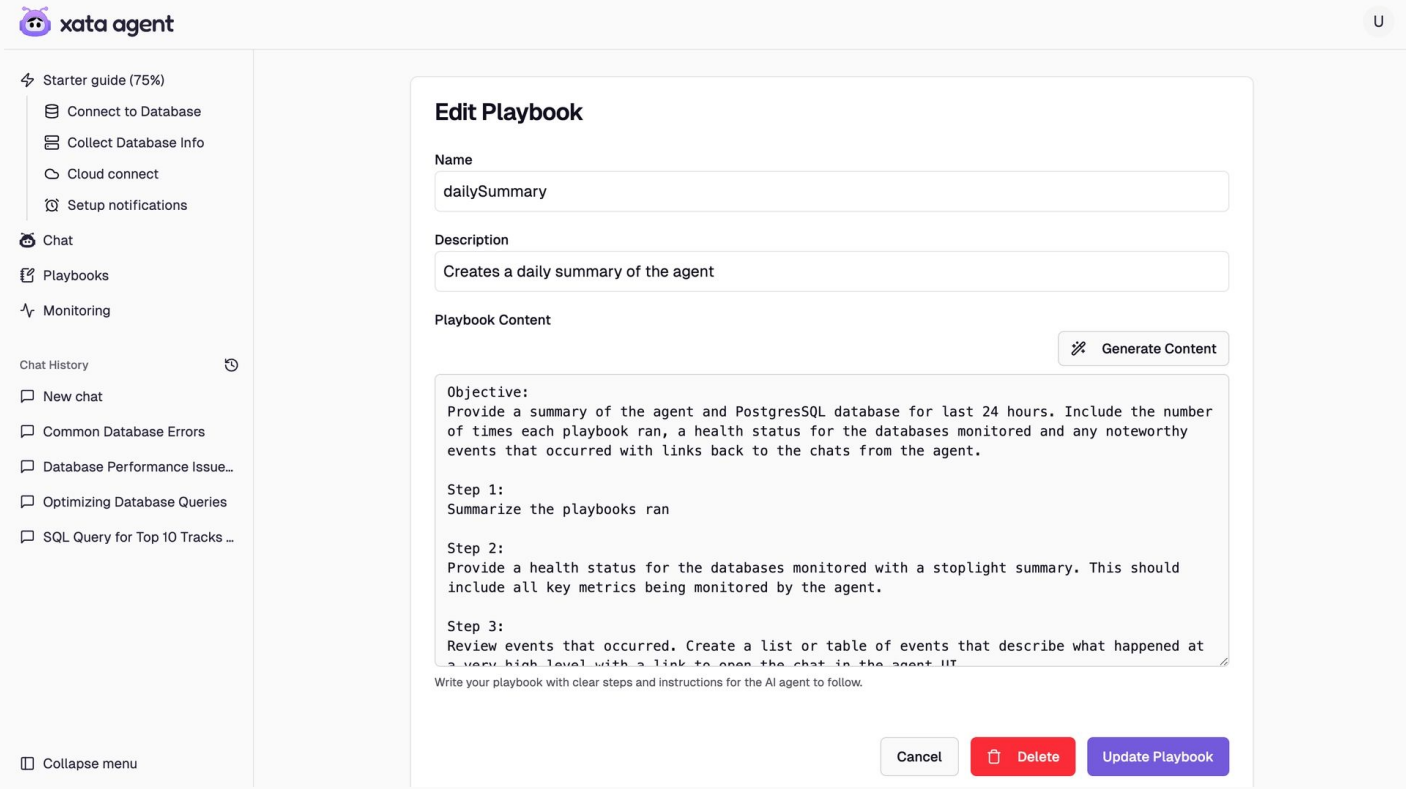

xata

Tools

Act

Tools

- A mix of simple and complex tools for the agent to act
- Tool examples:
  - getInstanceMetric()
  - getInstanceLogs()
  - getSlowQueries()
  - explainQuery()
  - getSettings()
  - getTablesInfo()
  - getPostgresExtensions()

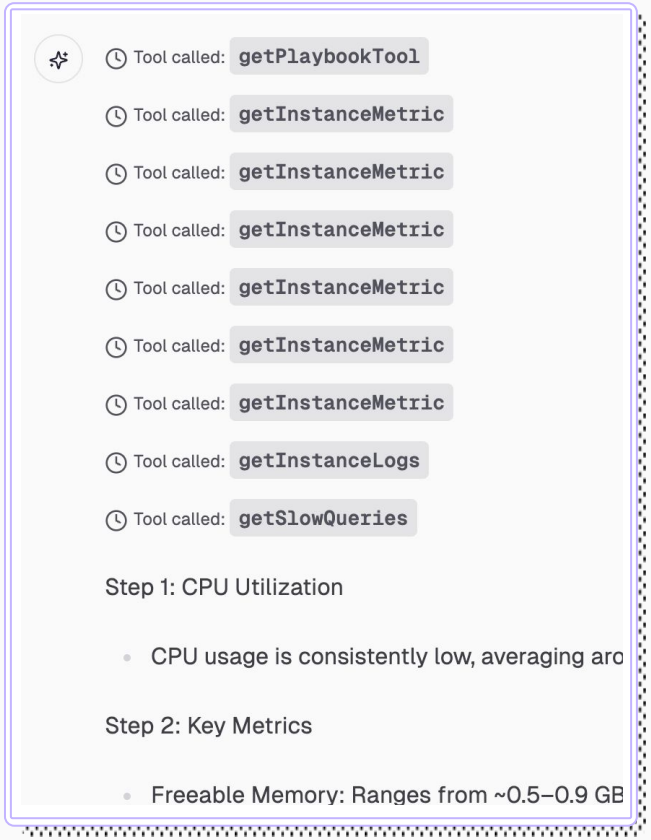

xata

Example tool implementation: getSlowQueries

```
export async function getSlowQueries(client: ClientBase, thresholdMs: number):
Promise<SlowQuery[]> {
  const query = `
    SELECT
      calls,
      round(max_exec_time/1000) max_exec_secs,
      round(mean_exec_time/1000) mean_exec_secs,
      round(total_exec_time/1000) total_exec_secs,
      query,
      queryid::text as queryid
    FROM pg_stat_statements
    WHERE max_exec_time > $1
    ORDER BY total_exec_time DESC
    LIMIT 10;
  `;
  const result = await client.query(query, [thresholdMs]);
  return result.rows;
}
```

Model Context Protocol

- Proposed by Anthropic but now the industry standard
- An MCP server exposes
  - Resources
  - Tools
  - Prompts
- Many companies are adding MCP servers for their services
- Range from API wrapper to use case oriented

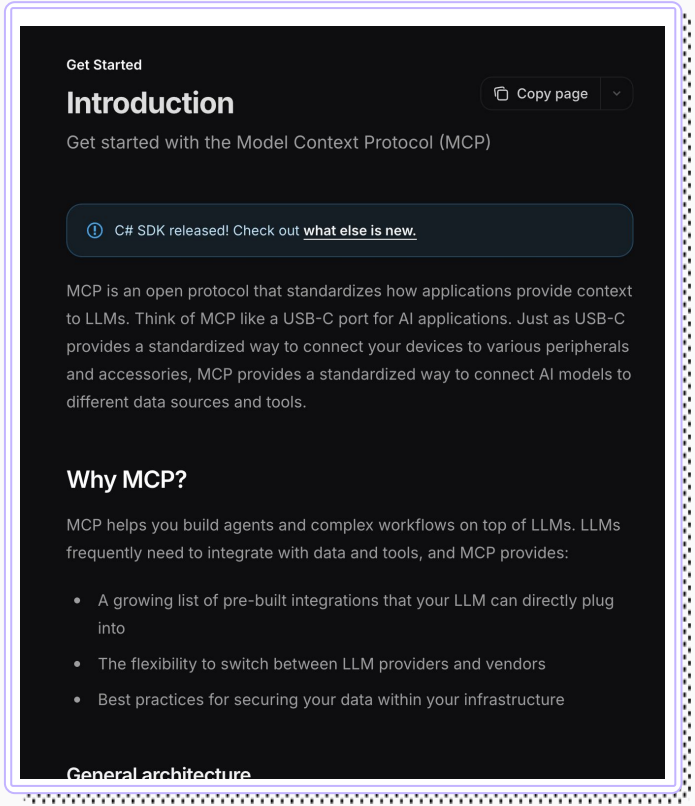

Custom tools: MCP servers

- MCP
  - Local MCP server (stdio)
  - Remote MCP server (jsonrpc)
- Local MCP makes it easy to add your own tools:
  - Write them in TS or Python, drop them in a folder
- As cloud providers add MCP servers, the agent can just use them.

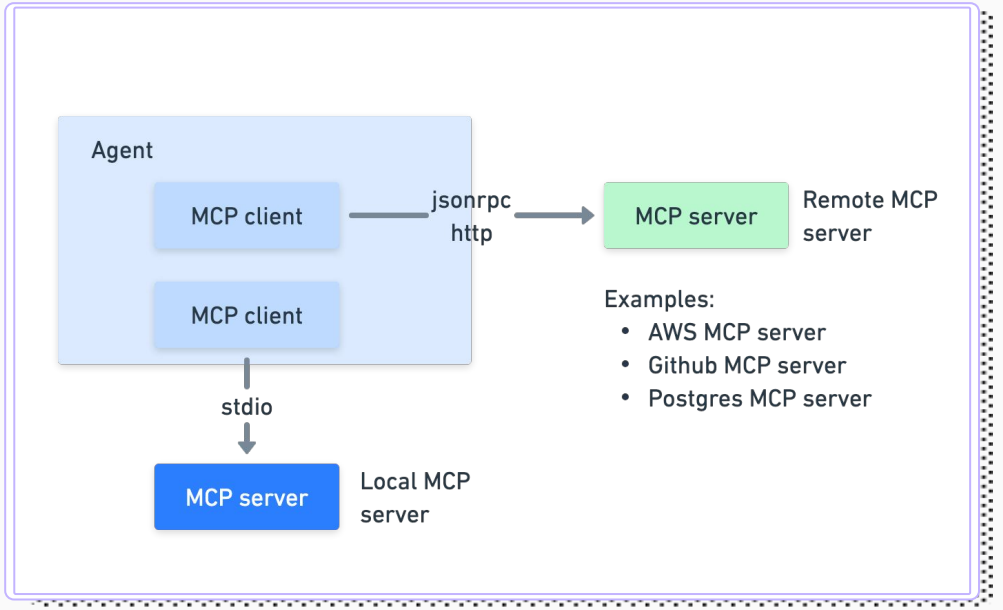

xata

Approval

Evaluate

xata

Recommend

Provide recommendations for end users to act on

Approval

Provide recommendations and the ability approve changes for the agent to acton

Autonomous

YOLO

xata

Fine Tune

Iterate

Models

- Uses Vercel AI SDK as an abstraction layer for multiple models
  - OpenAI
  - Anthropic
  - Deepseek
  - Gemini
  - Ollama
  - LiteLLM
- Reasoning models tend to do better on agentic use cases
  - Better answers/reasoning > speed
  - More expensive

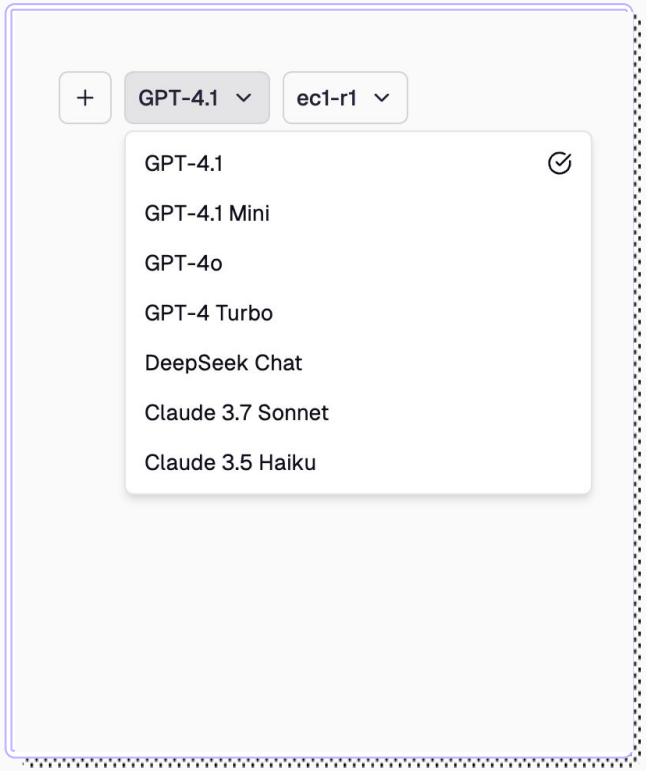

Evals

What

- LLM judges
- Repeatable test designed to measure some aspect of model behavior
- Crucial for tracking improvements, identifying regressions, and understanding limitations

When

- A playbook / prompt is changed
- A new model is release
- A new tool is added

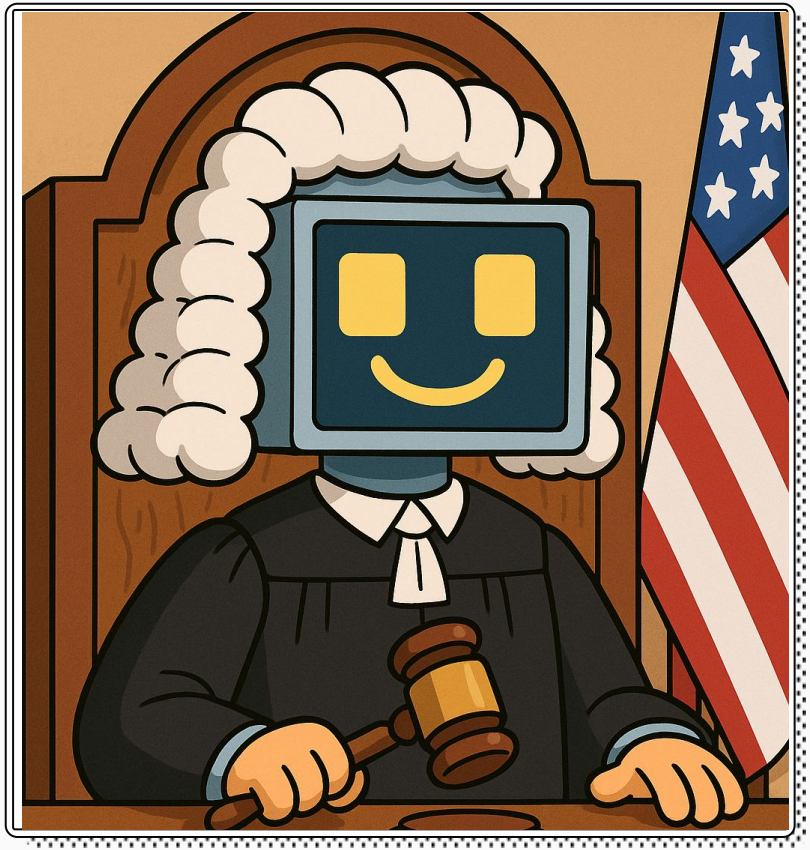

xata

Evals

LLM Judges score 👩‍⚖️ Did it identify the root cause? Is the feedback actionable? Is it concise or too much to read?

```
const finalAnswerJudge = (expectedAnswer = ''): LLMJudgeConfig => ({
  name: 'final_answer',
  prompt: ({ input, finalAnswer }) => `
  The following question was answered by an expert in PostgreSQL and database
administration:
  <question>${input}</question>
  <expertAnswer>${finalAnswer}</expertAnswer>
  ${expectedAnswer ? `<expectedAnswer>${expectedAnswer}</expectedAnswer>` : ''}
  Please determine whether expert passed or failed at answering the question
correctly and accurately. Please provied a critique of how the answer could be
improved or does not match the response of an expert in PostgreSQl and database
administration.
  `
});
```

xata

Challenges

Can I *really* trust an agent?

xata

Is my data safe and private?

xata

Is an agent cost effective for observability?

xata

Should I YOLO?

xata

What's next

Looking ahead

Agent Roadmap

- [x] Model tool / function calling support
- [x] Safely run arbitrary SQL
- [x] Better short & long term memory
- [x] Code & observability integrations
- [x] GitHub approval flows

xata

Xata Demo

Not a product pitch, just setting context - I swear

xata

Autonomous Postgres

- Store more context: metrics, logs, query stats, schema changes
- Recommendation engine: Slow queries first
- Apply to new CoW branch & test workflow
- Custom playbook & tooling support
- Non-functional requirements - evals, notifications, etc.

xata

Join in on the fun

Deploy and provide feedback

- Works with any Postgres
- Native support for:
  - AWS RDS, Amazon Aurora, GCP Cloud SQL

Contribute to the project

- We are friendly to outside contributions
- Gain experience with any of these technologies
  - Next.js / Typescript
  - Evals
  - LLM Memory
  - GitHub integration

xata

Postgres at scale

Thank you!

alex-francoeur

alexf@xata.io

xata.io