

Autonomous PostgreSQL

When AI Agents Become Your DBA

AI agents amplify DBA capabilities — they don't replace them



Vivek Singh

Principal Database Specialist, AWS



Shaily Porwal

Sr. Technical Account Manager, AWS

About the Speaker - Vivek Singh



20+ Years PostgreSQL Expertise



Principal Database Specialist at AWS



Trained 500+ AWS Professionals



Globee and Titan Platinum 2025 winner



Author & Thought Leader

Everything shown today is based on real implementations and proven results.

AWS & Technical Publications

- Improve PostgreSQL performance using pgstattuple extension
- Optimize the cost of Amazon RDS snapshots
- Best practices for migrating PostgreSQL to Amazon RDS & Aurora
- Deep dive on Aurora & RDS for PostgreSQL architecture
- Best practices for cross-Region read replicas
- AWS Whitepaper: Optimizing PostgreSQL on EC2 with EBS
- DZone: MERGE Command in PostgreSQL 15
- IoT-enabled smart robotic baggage monitoring system



About the Speaker - Shaily Porwal



15+ Years PostgreSQL Expertise



Senior Technical Account Manager at AWS



Trained 400+ AWS Professionals



About me

- 4x AWS Certified, including AWS Certified Solutions Architect - Professional
- Relational Database Service Core Subject Matter Expert and Relational Database Service PostgreSQL Subject Matter Expert
- Database Technical Field Community member with deep expertise in Relational Database Service and Aurora PostgreSQL, mentoring Technical Account Managers in database best practices
- Oracle Expert with extensive experience in Oracle-to-PostgreSQL migrations, having developed migration frameworks used across AWS internal teams
- Regular speaker and facilitator at enterprise customer workshops, webinars, and conference talks, including the DFW Relational Database Service/Aurora Operational Excellence Workshop with 82 participants and 56 enterprise customers
- Conducted multiple AWS internal enablement sessions on database services and AI/ML
- Passionate about bridging traditional database administration with modern AI-driven automation

01

THE CHALLENGE

Why traditional DBA workflows can't scale in the age of cloud-native databases



The DBA's Dilemma

Traditional DBA Time Allocation



Manual Monitoring

Reactive responses to alerts and issues



Human Error Risk

Error-prone remediation decisions



24/7 Availability

Always-on requirement drains teams



Scaling Expertise

Growing teams need growing knowledge



Routine Maintenance

Time-consuming repetitive tasks



Delayed Response

Critical issues wait for human attention



Can we make database administration autonomous and intelligent?

Research : Is This a Real Problem?

ARTICLE

How Much Time Are You Wasting on Manual, Repetitive Tasks?

by Katy Belsol

Content Center > Managing Work > Project management

Conversations about how automation can improve productivity and reduce costs for businesses are often coupled with stories of how automation is taking jobs away from workers. Yet according to a [report by McKinsey Global Institute](#), less than five percent of jobs consist of activities that can be fully automated.

Far more common will be the automation of constituent activities for a given role, according to McKinsey. In about 60 percent of occupations, at least one-third of the activities that make up a specific job could be automated. The good news is that most workers see this automation in a positive light and are keen to take advantage of automation that has the potential to make their work – and their lives – better.

redgate Products Solutions Customers Our company Support Learning & community

New 2025 State of the Database Landscape report | Get your copy →

Our company Overview Careers Contact us Blog **Newsroom** Leadership [Latest job opportunities](#)

February 11, 2025
Media Contacts
Mighana Shendrikar
redgen@athorworldwide.com

Redgate unveils 2025 State of the Database Landscape report: organizations double down on security practices amid multi-platform adoption

CAMBRIDGE, UK, February 11, 2025 – The growing urgency of safeguarding sensitive data, navigating regulatory compliance, and increasing adoption of multiple database platforms are just some of the concerns impacting IT leaders as they manage growing database complexities, according to the State of the Database Landscape, a major new survey from end-to-end Database DevOps provider Redgate.


38% of IT teams are concerned about data security and access controls when managing different technologies, an increase of 12% year-on-year. In fact, one in five (21%) say they won't adopt more than one database type, despite their advantages, simply because they're concerned about [security](#) and [compliance issues](#).


Embracing multiple database platforms – and dealing with complexity

Driven by the need for flexibility, cost-savings and serving different use cases, almost three quarters (74%) of IT teams now use two or more database platforms, with a quarter (25%) of respondents using more than four database platforms. However, the sheer complexity of the systems being managed highlights the crucial need for robust database management strategies to ensure data consistency, security and compliance. The research uncovered the top four most popular databases that organizations have introduced over the past two years: the open-source platforms PostgreSQL and MySQL, the non-relational database MongoDB, and cloud-based data platform Snowflake.

IDC August 2025 | Business Value White Paper, sponsored by Oracle

The Business Value of Oracle Autonomous AI Database

 Devin Pratt
Research Director,
Data Management, IDC

 Megan Szurley
Business Value Manager,
Business Value Strategy Practice, IDC

Canonical Products Solutions Partners Careers Company

BLOG / ARTICLE

 Canonical
on 26 March 2025

70% of IT teams spend more than 6 hours per week on security patching – new IDC research

[Research](#) [Security & compliance](#) [software supply chain](#)

Share on [f](#) [X](#) [in](#)

NEWSLETTER SIGNUP

Get the latest Canonical news and updates in your inbox.

* Work email:

*I agree to receive information about Canonical's products and services.

By submitting this form, I confirm that I have read and agree to [Canonical's](#)

Latest research from Canonical and The International Data Corporation (IDC) shows organizations struggle to apply fixes confidently under tough CVE patching mandates, and face other serious challenges in the open source software supply chain

Today, Canonical, the publisher of Ubuntu, published a research report in collaboration with IDC and co-sponsored by Google Cloud that uncovers new insights into the pressures and challenges organizations face with security patching and rising regulatory burdens. The report, titled *The state of software supply chains: Security challenges, opportunities and the path to resilience with open source software*, surveyed 500 organizations with more than 250 full-time employees to determine the most pressing issues these organizations face. Most notably, these are difficulties in vulnerability and patch management, insufficient visibility of software dependencies or software supply chain, and the trustworthiness of software sources.

Research Shows: The DBA Time Burden

70%

of IT teams spend 6+ hours/week on security patching alone

Gartner

75%+

of IT leaders report teams spend 5-25 hours/week updating legacy systems

IDC

40%

of technical workers spend 25%+ of their week on manual, repetitive tasks

Smartsheet

80%

of DBA time spent firefighting vs. proactive optimization

Redgate



The Good News: Automation can save 2-6 hours per DBA per day

02

THE SOLUTION

Agentic AI meets PostgreSQL — intelligent, adaptive, human-in-the-loop





This is Agentic AI for PostgreSQL



Observes

Continuously monitors your database health



Reasons

Understands context and patterns



Acts

Takes action with human approval

What Makes This Different?

Traditional Automation



Static Scripts

Hardcoded rules that never evolve



Predefined Conditions

Fixed thresholds miss context



No Learning

Same mistake, different day



Brittle

Breaks when conditions change

VS

Agentic AI



Context-Aware

Understands your full environment



Learns from Patterns

Improves with every incident



Natural Language

Talk to your database, naturally



Adaptive

Handles new situations gracefully

The MCP Foundation

The Problem



M×N Integration Problem

Every AI tool needs custom connectors



AI Isolation from Real-World

Models can't access live databases



Fragmented Tool Integration

Teams reinvent the wheel each time



Knowledge Silos & Context

Data trapped across multiple systems



Legacy System Gap

Old systems excluded from AI workflows



Model Context Protocol



Open Standard

Universal AI-to-system integration



Think: REST API for AI

Standardized protocol for agents



Vendor-Neutral

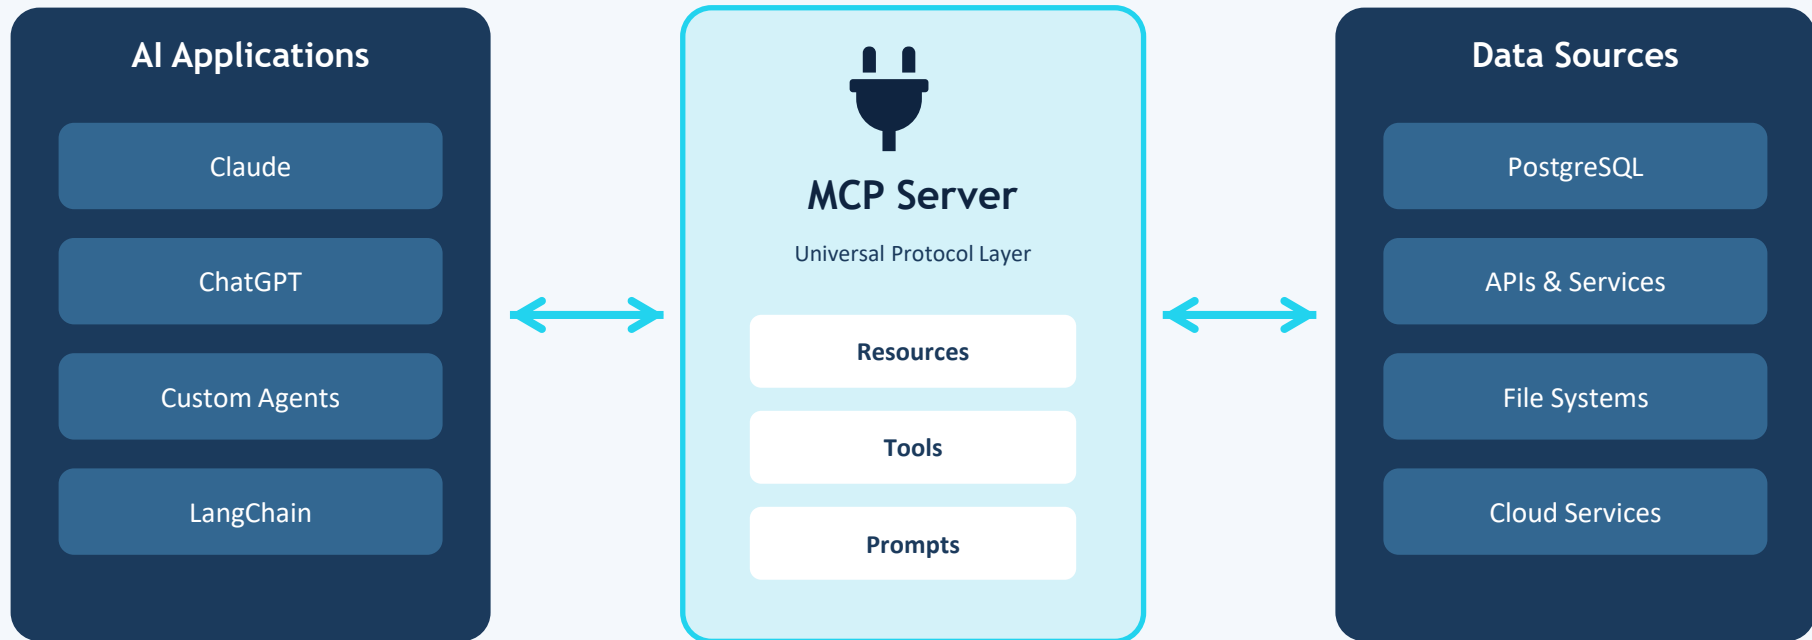
No lock-in to any AI provider



PostgreSQL-Friendly

Native database connectivity


How MCP Connects Everything



MCP eliminates the M×N integration problem with a single standardized protocol

Architecture: Open Source Path

OPEN SOURCE



Key Benefits

- ✓ Complete control over data and models
- ✓ Zero vendor lock-in — fully portable
- ✓ Run 100% on-premises if required
- ✓ Lower long-term cost at scale (2-4 GPU servers)

The Architecture – Open Source Stack

Ollama / vLLM (Local Models)



Llama 3.1 / Mistral / Qwen



Supervisor Agent — LangGraph



Action Agent — Execute



Health Check Agent — Monitor



MCP SERVER (Local)

Database Tools

analyze_query

check_health

create_index

vacuum_table



POSTGRESQL

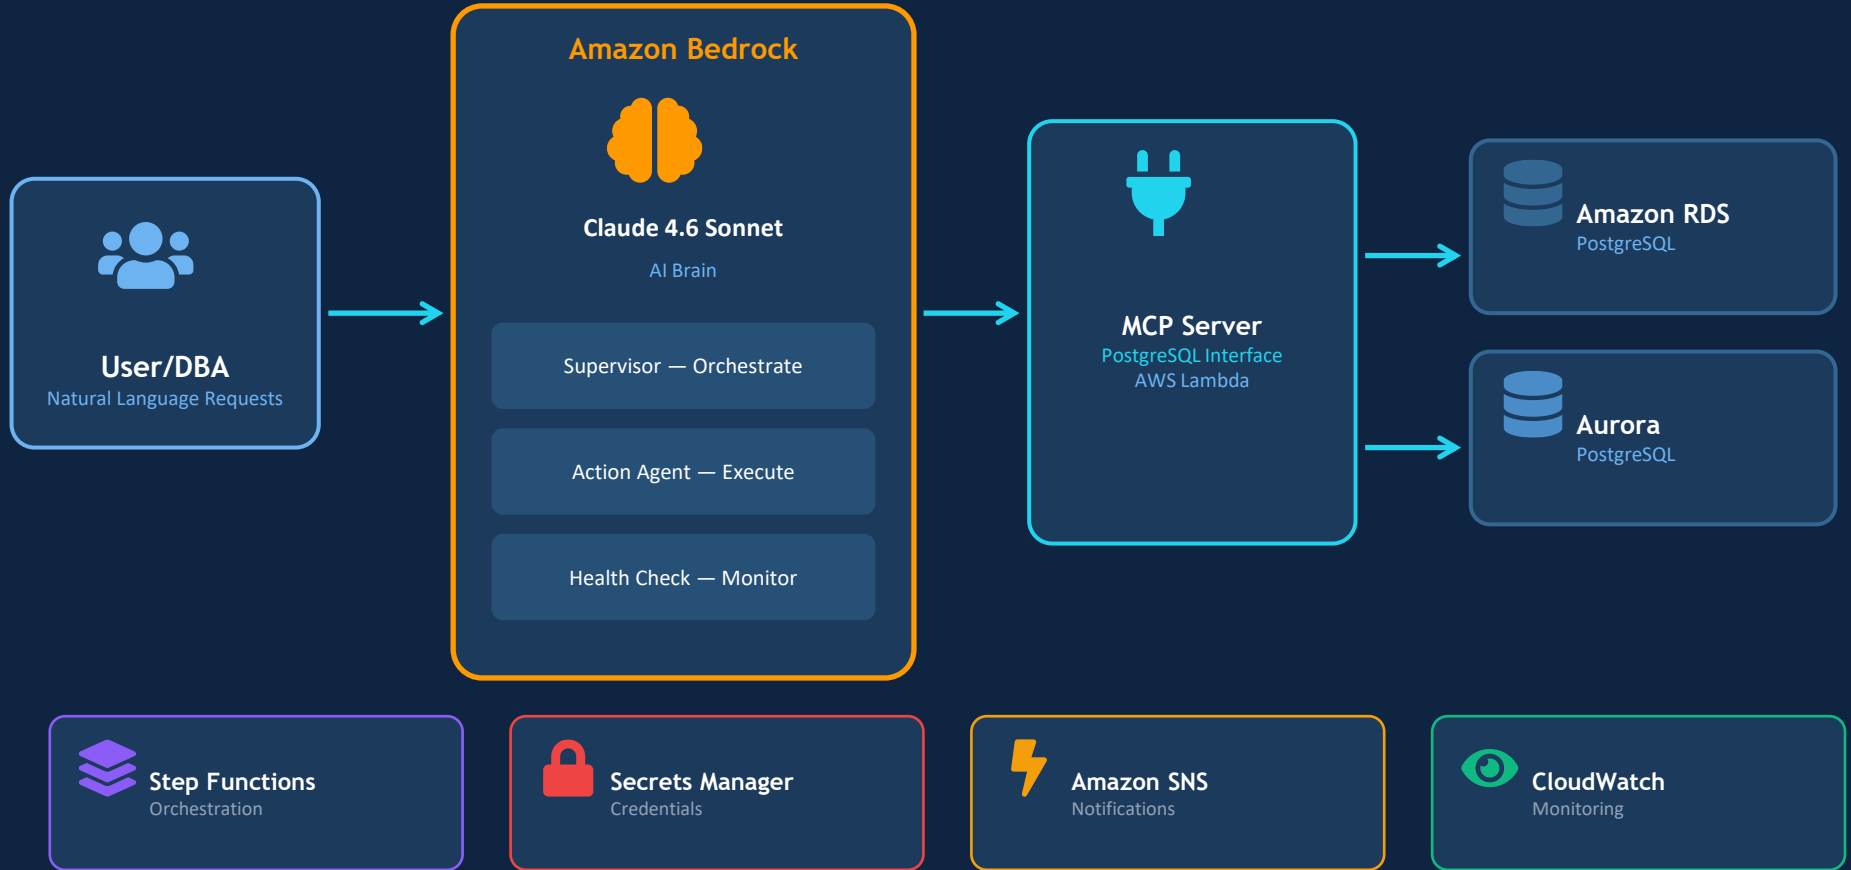
localhost:5432

customer_data

knowledge_base

pgvector

AWS Architecture – Agentic AI DBA for PostgreSQL



The Architecture – Agent Hierarchy

Region: us-east-1

AMAZON BEDROCK



Claude 4.6 Sonnet



Supervisor Agent — *Orchestrate*

Routes tasks to specialized agents



Action Agent — *Execute*

Performs database operations safely



Health Check Agent — *Monitor*

Continuous health assessment

MCP SERVER (Local)



Database Tools

Slow query

Analyze Query, Execution plan

Search knowledge base,
Missing index

Vacuum, Database
statistics



POSTGRESQL

localhost:5432

customer_data

knowledge_base


pgvector

Common Use Cases for Agentic AI DBA

Traditional Approach

- 1 Alert fires for slow query
- 2 DBA logs in, checks pg_stat_statements
- 3 Analyzes execution plan manually
- 4 Researches index options
- 5 Tests and implements fix

 30-60 minutes per incident



Agent-Led Approach

- 1 Agent detects slow query pattern
- 2 Automatically analyzes execution plans
- 3 Simulates index scenarios
- 4 Proposes fix with impact analysis
- 5 DBA approves, agent implements

 5-10 minutes per incident

MCP Tools Required:

`analyze_query_performance`

`get_execution_plan`

`simulate_index_impact`




`create_index_safely`

80-90%





Time Reduction

Intelligent Vacuum & Bloat Management

The Problem

-  Bloat accumulation in high-transaction tables
-  Manual vacuum scheduling is reactive
-  Autovacuum settings often suboptimal

Agent Solution

-  Monitors bloat metrics continuously
-  Predicts vacuum needs from transaction patterns
-  Schedules maintenance windows intelligently
-  Adjusts autovacuum parameters dynamically

MCP Tools Required:

`check_table_bloat`

Monitor bloat levels

`analyze_vacuum_history`

Pattern analysis

`schedule_vacuum_operation`

Smart scheduling

`tune_autovacuum_settings`

Dynamic tuning

40%

Reduction in Bloat-Related
Performance Issues



LIVE DEMO

Autonomous PostgreSQL Agent in Action



Natural language query
analysis



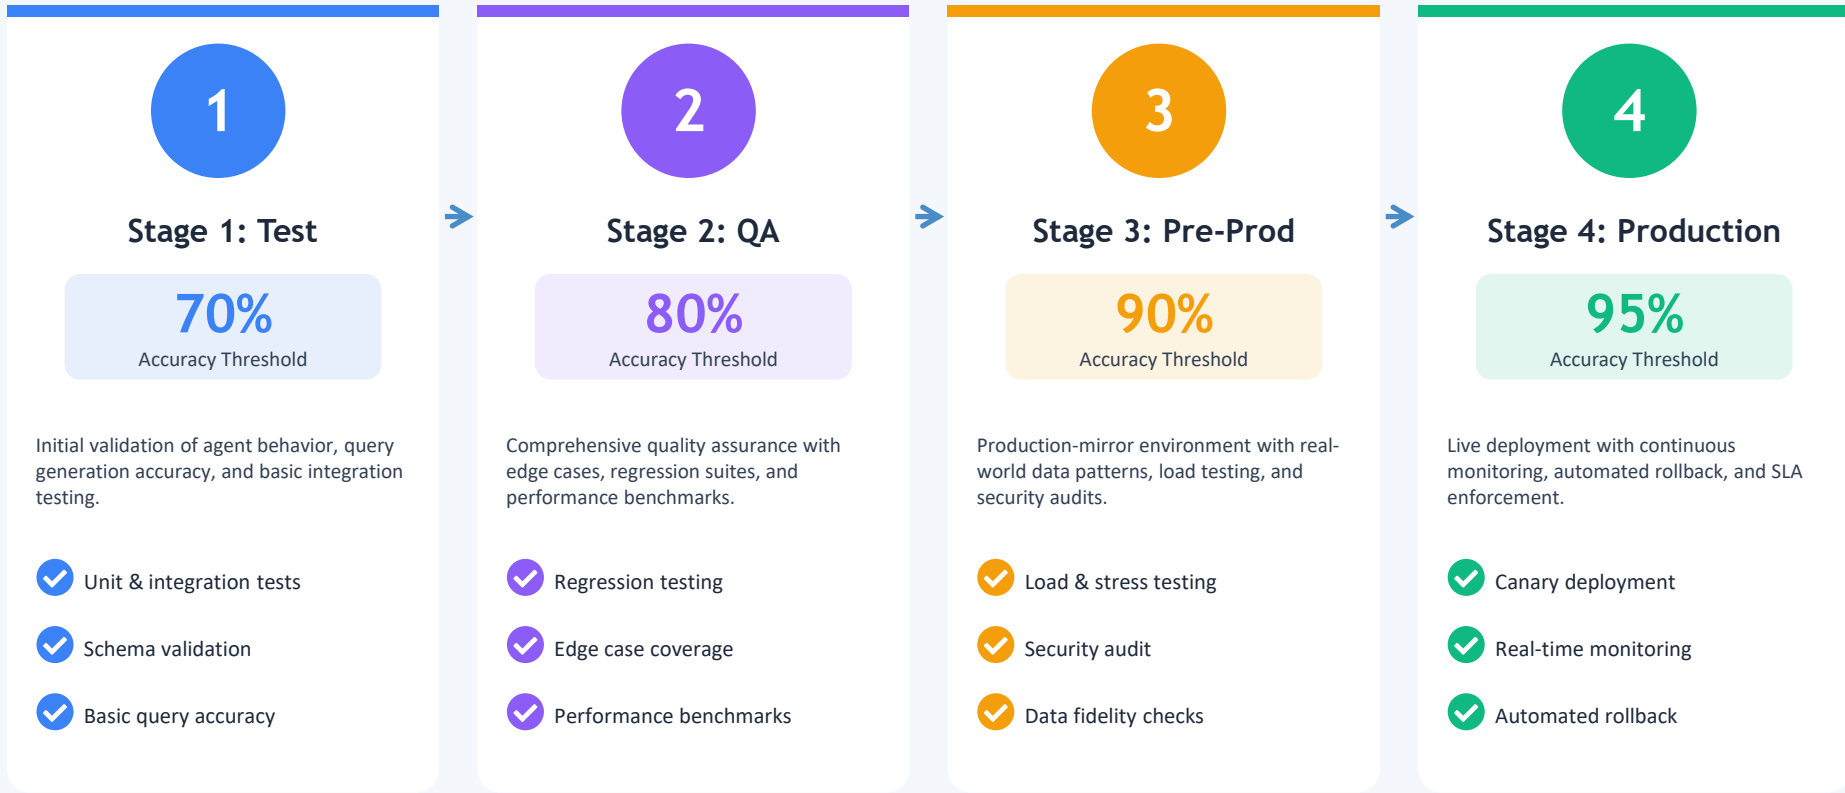
Real-time recommendations



One-click implementation

THE PROGRESSION

4-Stage Gate Deployment Process



Resources

Core Technologies & Protocols

Model Context Protocol (MCP): Learn about the open standard for connecting AI models to data sources in the [official Anthropic MCP Introduction](#) and explore the pre-built [PostgreSQL MCP Server on GitHub](#).

pgvector & HNSW: Deep dive into high-performance vector indexing with the [AWS Blog on HNSW Techniques](#) and Crunchy Data's guide to [HNSW Indexes in Postgres](#).

Practical Implementation & Tutorials

Building Database Agents: Follow this hands-on workshop on [Building AI Agents for Database Day-2 Operations](#) to automate routine DBA tasks.

Natural Language to SQL: Review Punit's guide on [Postgres MCP for querying in plain English](#) to understand schema discovery and tool-calling implementation.

Agentic Database Analysis: Explore the [Agentic Database Analysis & Migration \(ADAM\) solution](#) on AWS Marketplace for automated compatibility and performance analysis.

Operational Frameworks

Identify, Detect, and Resolve (IDR): Adapt the [NIST Incident Response Lifecycle](#) to database operations to reduce Mean Time to Recovery (MTTR) through automated detection and remediation.

AI-Driven Observability: Watch the AWS re:Invent session on [AI Agents for Databases: Discover, Recommend, Optimize](#) for real-world patterns used by AWS experts

GitHub: <https://github.com/shaily-2511/Autonomous-DBA/tree/main>

Thank You!



<https://www.linkedin.com/in/vivek-singh-4149aa19/>



Vivek Singh

Principal Database Specialist, AWS



www.linkedin.com/in/shailyporwal25



Shaily Porwal

Sr. Technical Account Manager, AWS