awidelabs

PostgreSQL Operations Intelligence Suite

# Make your PostgreSQL predictable

Monitoring, query optimization and automated operations in one platform — with AI recommendations and self-healing

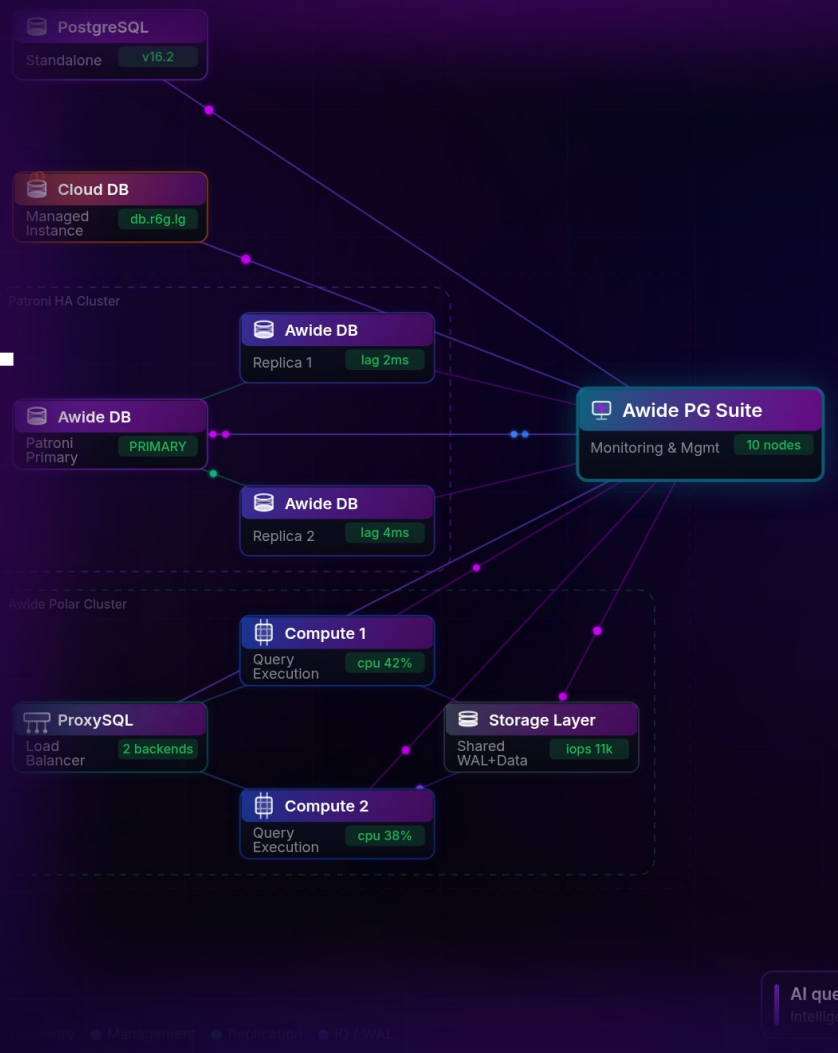

THE CHALLENGE

# Running PostgreSQL shouldn't feel this risky

PostgreSQL powers mission-critical systems. But in regulated, hostile or high-availability environments, a vanilla distribution turns database operations into a constant balancing act between security, stability and speed.

**Awide** moves your team from firefighting to foresight: answers before questions are asked, solutions before problems become incidents.

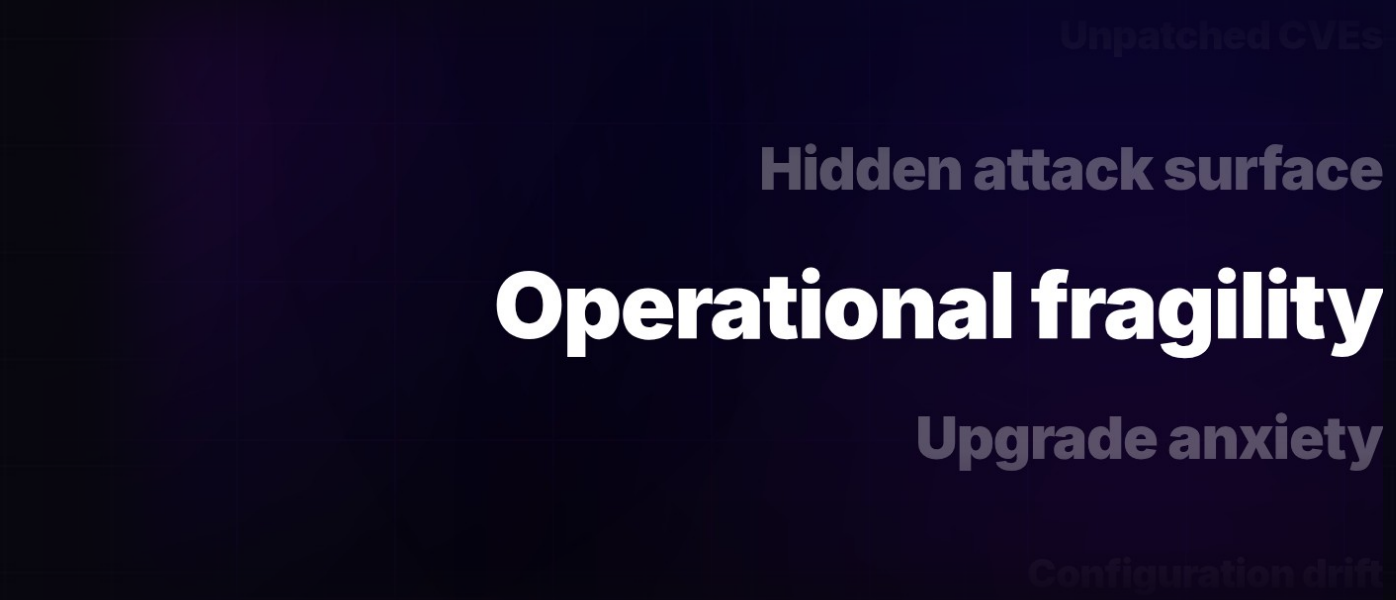


- Hidden attack surface
- Operational fragility
- Upgrade anxiety
- Configuration drift
- Unplanned downtime
- Unpatched CVEs



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

THE AWIDE ECOSYSTEM

# Three products, one PostgreSQL ecosystem

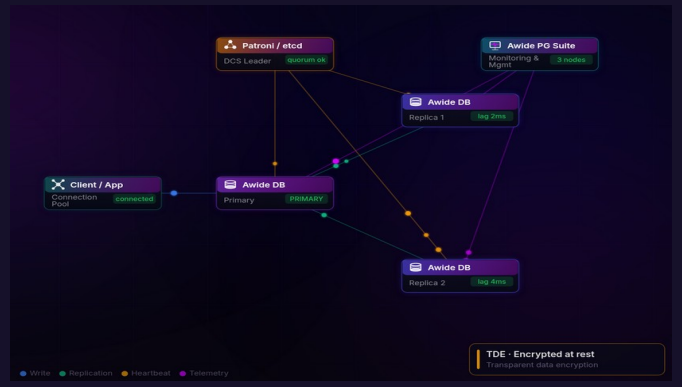


HARDENED DBMS

## Awide DB

Hardened PostgreSQL for regulated clouds, running inside the customer's own cloud account.

TDE · RBAC · 64-bit XID · pgvector · CIS

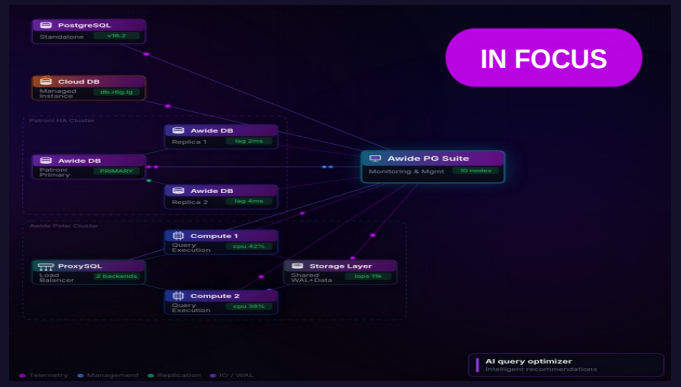


OPERATIONS PLATFORM

## Awide Suite

Full-stack observability, AI query optimization and zero-downtime tooling for every cluster.

8 modules · AI recommendations · Self-healing

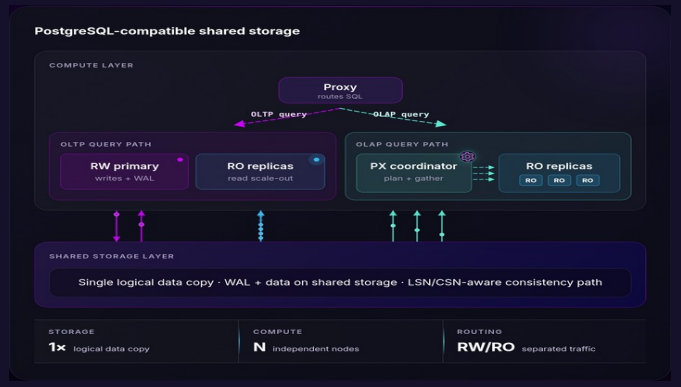


CLOUD-NATIVE DBMS

## Awide Polar

One write node, elastic read-only compute and a shared storage layer.

NVMe-oF over RDMA · HTAP without ETL · RPO = 0



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

OPERATIONS INTELLIGENCE SUITE

# From reactive to predictable operations

### Unified dashboard

Every cluster, instance and metric in one place instead of dozens of consoles

### Auto-configuration

Recommended settings tailored to each machine's CPU, RAM, disk and workload

### Smart monitoring

Health score, triggers and alerts delivered to messengers and email

### High availability

Patroni cluster management: switchover, role history, unified configuration

### Query optimization

Profiler, plans, locks and AI recommendations for slow queries

### Cost reduction

Up to 80% less reliance on expensive professional services



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SCALE

# A platform that grows with your database fleet

**5,000**
PostgreSQL instances per installation

**14–18**
supported PostgreSQL versions

**8**
modules in a single interface

**80%**
maximum savings on external experts

## WHAT YOU CAN CONNECT

PostgreSQL · Patroni HA clusters · AWS RDS · Docker containers · Awide DB · Awide Polar · ProxySQL

Agents are built for x86_64, ARM64 and IBM Z · 1,000+ instances require the high-availability deployment



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

ARCHITECTURE

# Lightweight agents, an event bus and one platform

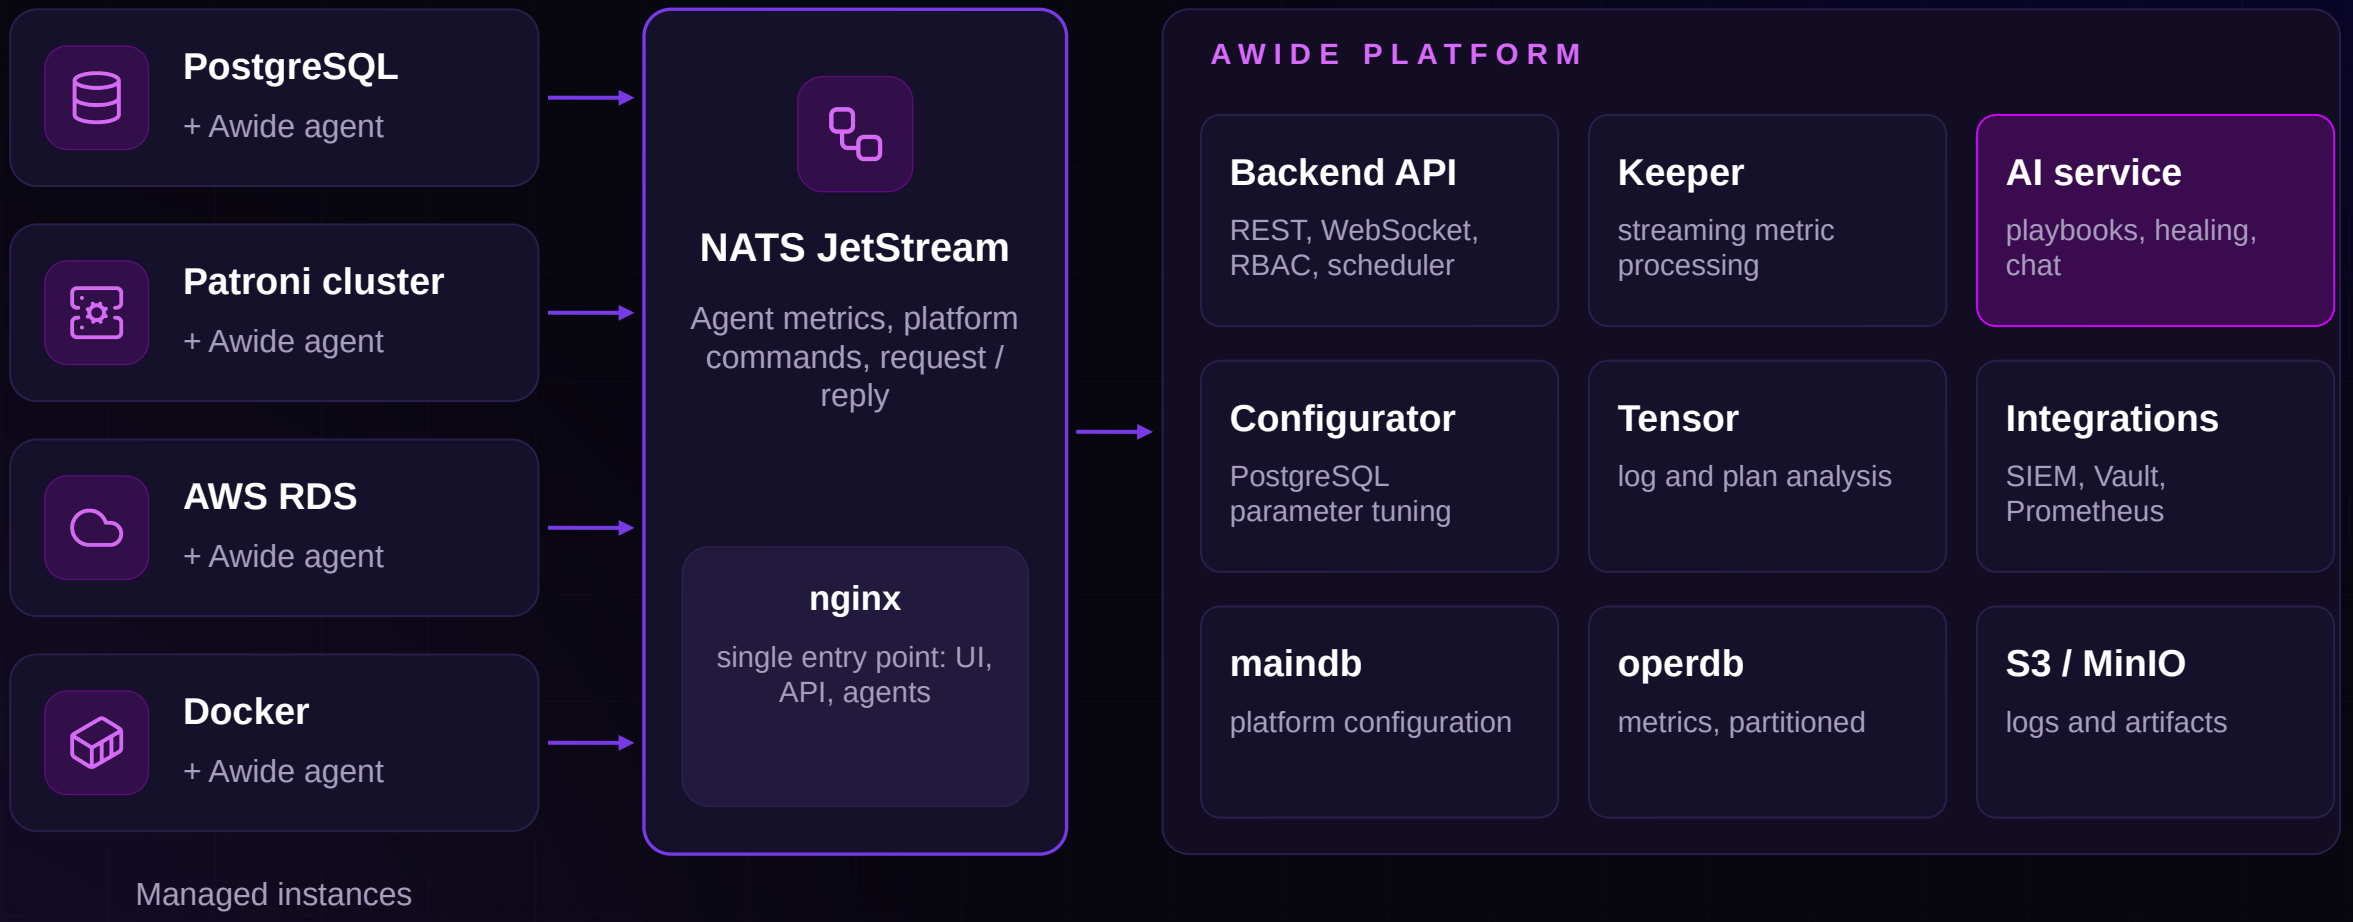




© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

POSTGRESQL DB MANAGEMENT

# A live view of every database

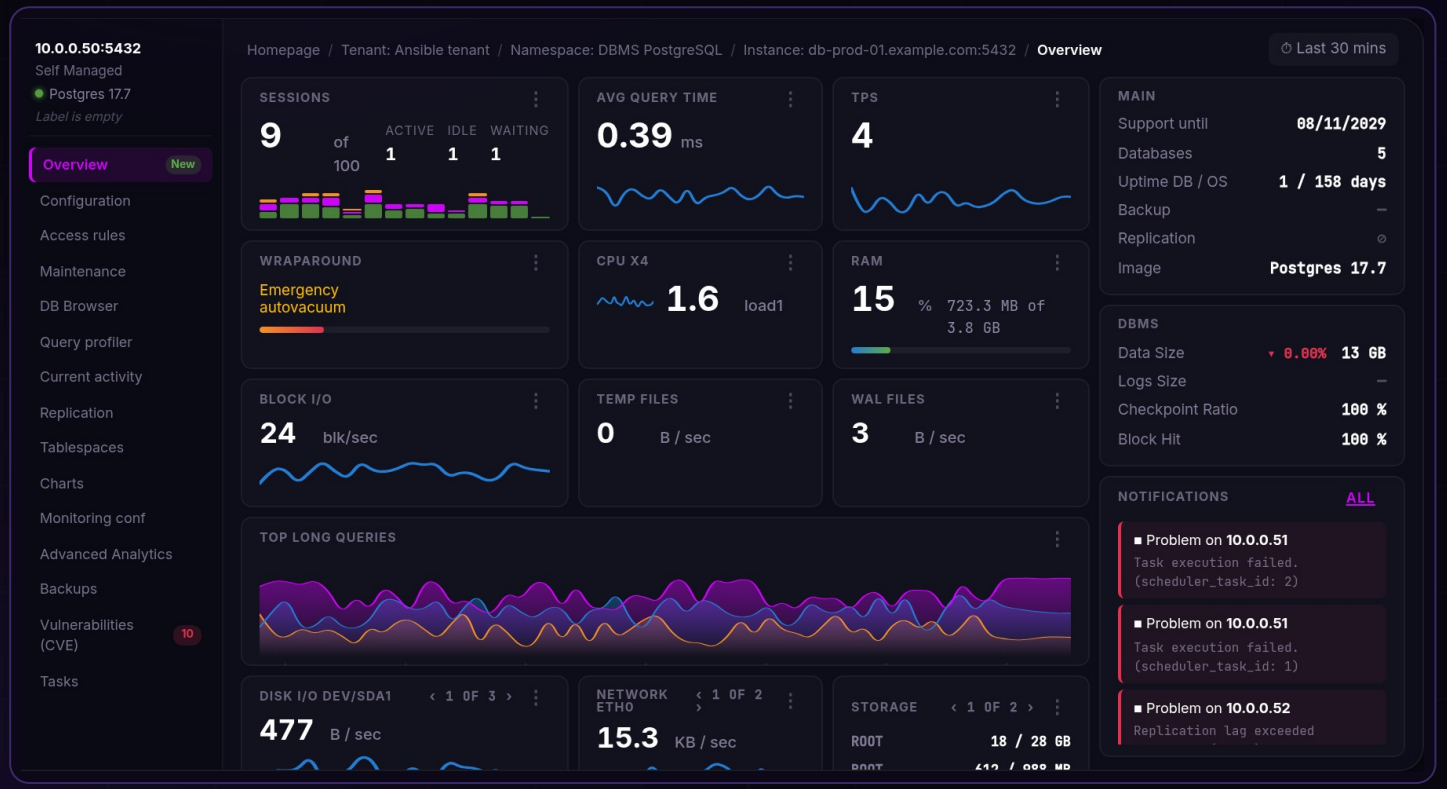


## Live dashboard

Sessions, TPS, replication lag, top queries, wraparound and autovacuum

## Health score 1–100

Daily score: cache, slow queries, bloat, disk, XID age, replicas

## Current activity

Sessions, locks and waits; cancel a query in one click

## Triggers and alerts

Email, Telegram, Mattermost, Slack, webhooks



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SQL QUERIES ANALYSIS

# Find the queries that hurt you

### Query profiler

pg_stat_statements and pg_store_plans: history, plans, EXPLAIN

### Log analysis

Problem queries, errors, checkpoints and VACUUM from PostgreSQL logs

### Locks and waits

Who blocks whom and where time is lost

### Schema audit

Bloat, unused and redundant indexes, FKs without indexes — with a report

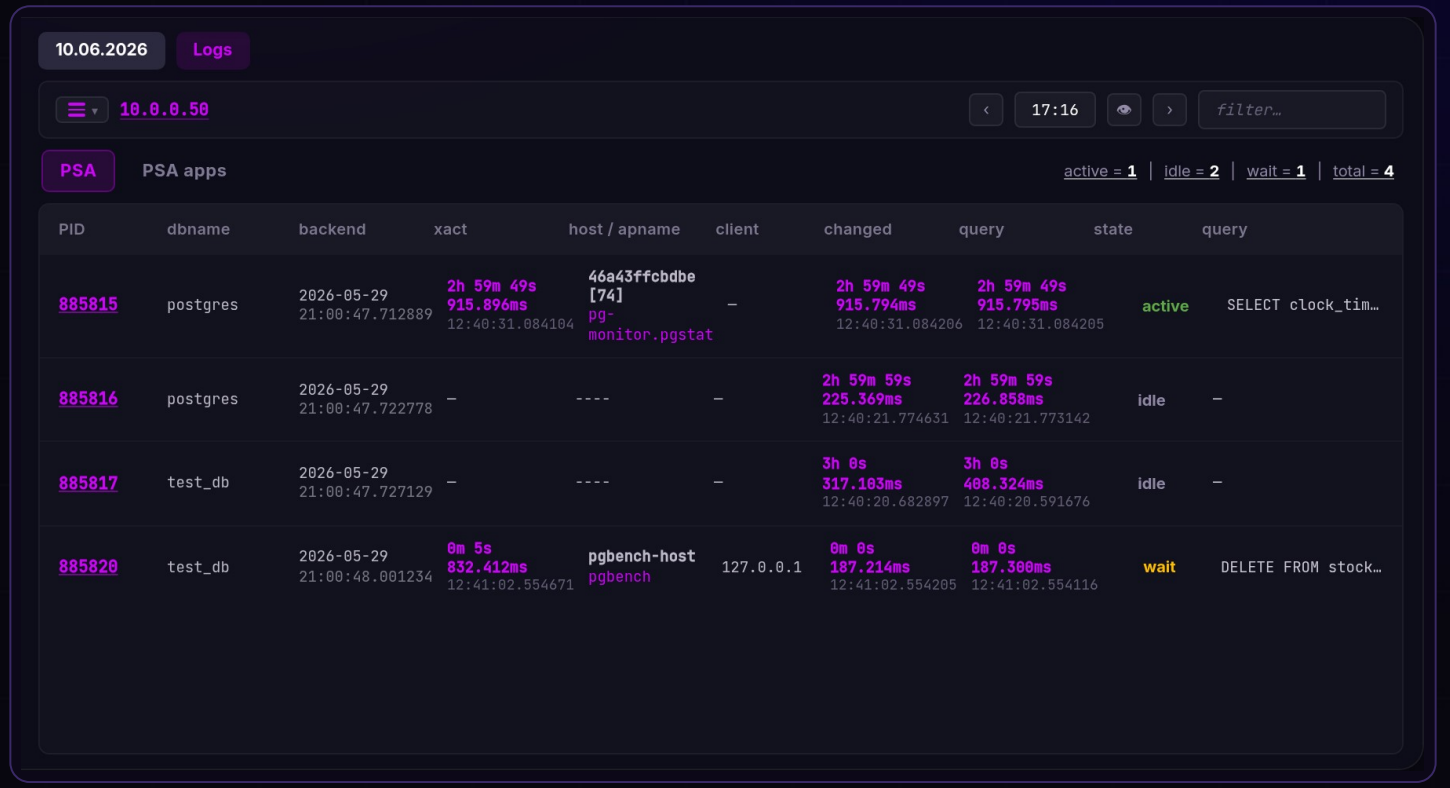




© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

AUTOMATION

# Configuration and maintenance by consistent rules

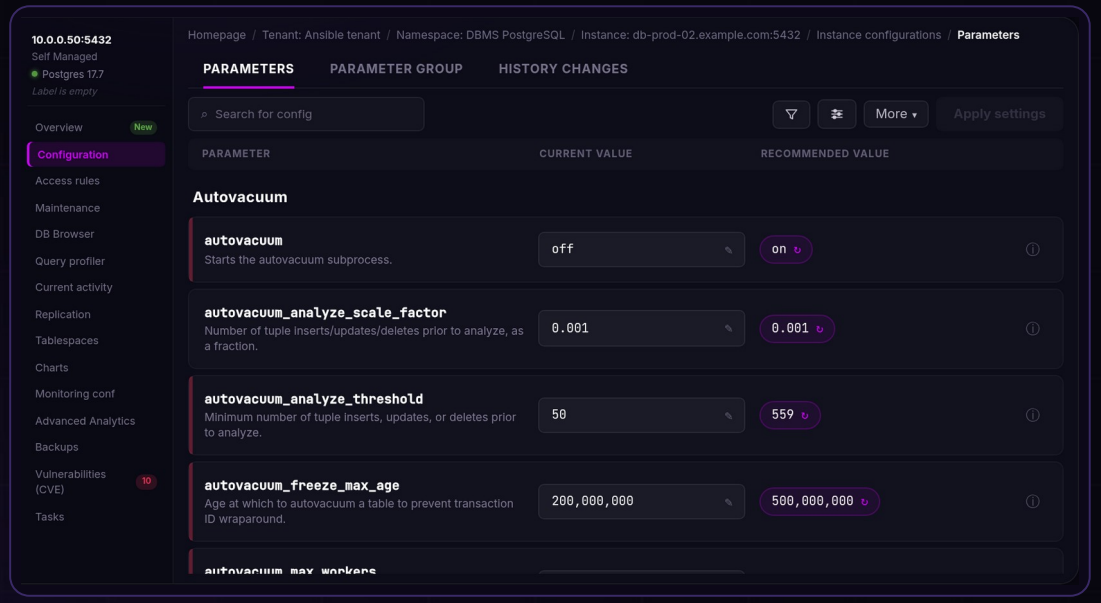


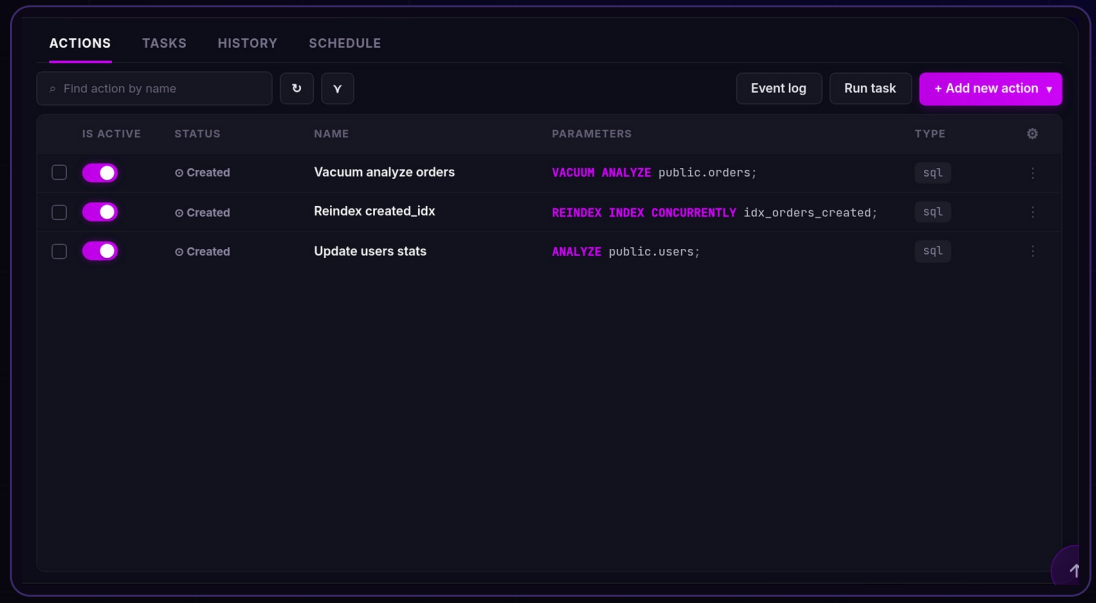


## Group parameters

- Reusable parameter groups for instances and clusters
- Preview, versioning, rollback and drift protection
- Auto-tuning for CPU, RAM, disk and workload

## Scheduled tasks and maintenance

- Run SQL or shell jobs: one-off, scheduled, recurring
- Auto-detect objects needing VACUUM, ANALYZE, REINDEX
- Job history, retries and role-based permissions



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

ENTERPRISE READINESS

# HA, security and integrations out of the box

## HA clusters

- Patroni: switchover, restart, adding nodes
- Cluster status and role history
- Replication and lag monitoring
- Awide Polar clusters with proxy nodes

## Security

- RBAC: 4 role levels and groups
- LDAP and Active Directory
- Event audit with SIEM export
- Secrets in HashiCorp Vault
- CVE checks for PG versions
- Data anonymization (pg_anon)

## Integrations

- Email, Telegram, Mattermost, Slack, webhooks
- Metric export to Prometheus / Grafana
- Backup monitoring and restore
- REST API, tokens and Swagger UI

**Multi-tenancy:** tenant → workspace → instance or cluster, with a visual topology and access control at every level



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

AI IN THE AWIDE PLATFORM

# From recommendations to self-healing

AI finds the problem, explains it, proposes a fix — and can safely apply it itself

AI monitoring

AI recommendations

Self-healing

AI assistant

AI MONITORING AND RECOMMENDATIONS

# AI checks your databases and explains what to fix

01

## Check

quick_check playbook runs in the lowest-load window; frequency adapts

02

## Severity

The model classifies the finding: info, warning or alert

03

## Drill-down

The agent picks a specialized playbook and gathers facts

04

## Recommend

A ready SQL script or configuration change, with an explanation

05

## Apply

Apply the fix in one click or mark the problem as resolved

SPECIALIZED DRILL-DOWN PLAYBOOKS

- Slow queries
- High CPU
- Low memory
- Connections
- Locks
- VACUUM
- Disk and IOPS
- 1C workloads
- Alert triage
- Config audit



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

AI RECOMMENDATIONS

# What a good AI recommendation looks like

*Illustrative example*

prod-db-01 · sales — Warning

## Bloat on orders

Dead tuples make up 34% of public.orders and autovacuum can't keep up with the update rate. A manual VACUUM reduces the dead tuple ratio without locking the table.

```
VACUUM (ANALYZE) public.orders;
```

Risk 2 of 5 · Regular | No table lock

Fix

prod-db-01 · sales — Alert

## Heavy query on orders

A sequential scan filtering public.orders by created_at takes 41% of total query time (mean 820 ms). An index removes the full scan.

```
CREATE INDEX CONCURRENTLY
ix_orders_created ON public.orders
(created_at);
```

Risk 2 of 5 · Regular | Reversible

Fix

reports-db · analytics — Warning

## Sorts spill to disk

Reporting queries write 12 GB/h of temp files: their sorts exceed work_mem = 4MB. Raising work_mem keeps sorts in memory.

```
work_mem: 4MB → 64MB
```

Risk 2 of 5 · Regular | No restart

Fix

EVERY CARD HAS

- Evidence with numbers
- One precise command
- Risk score and rollback
- Expected metric effect



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

RECOMMENDATIONS IN CONTEXT

# Recommendations right where you work

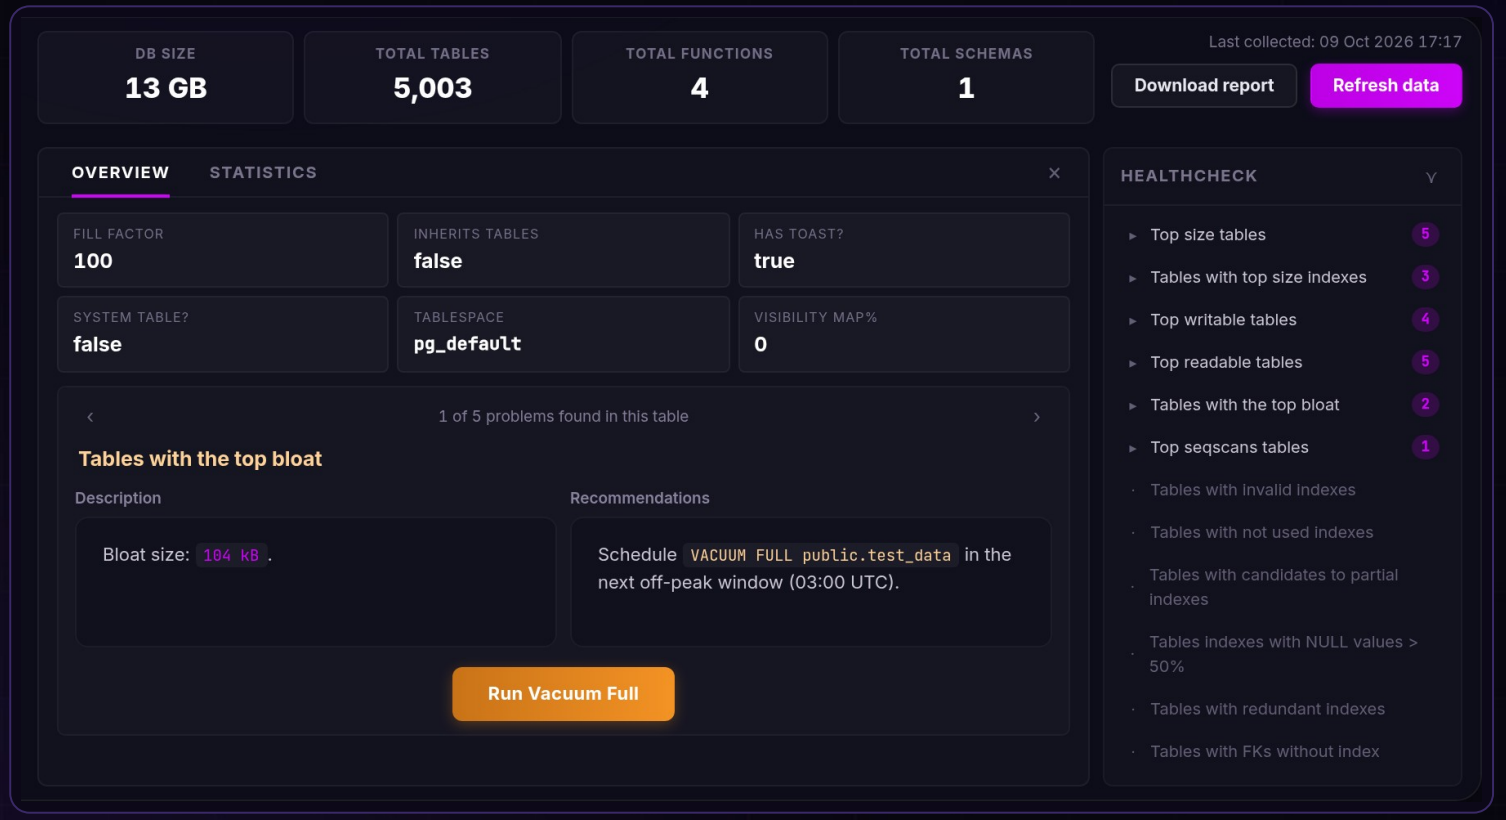


1

## The problem, located

Each finding is tied to a concrete object — a table, an index, a query

2

## The fix, explained

A plain-language description and a concrete recommendation

3

## One-click action

Apply the fix right away or schedule it for an off-peak window



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

HOW THE AI AGENT WORKS

# AI conclusions rest on real data, not guesswork

## Platform data

- Instance metrics
- PostgreSQL config
- Configurator baseline
- Applied profile
- Alerts
- PostgreSQL logs
- Scheduled tasks

## SQL sub-agents

- Database structure
- Queries and plans
- Sessions and locks
- Object state
- Ad-hoc read-only SQL

## Read-only diagnostics

- The validator allows only SELECT and EXPLAIN without ANALYZE
- FOR UPDATE and dangerous functions are blocked
- Queries are executed by the agent on the instance
- Playbooks are editable and can be reset to default
- Token usage is tracked for every run



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SELF-HEALING

NEW · RELEASE 7.1

# A closed loop: detect, fix, verify

1

## Detect

AI monitoring finds a problem and classifies its type

2

## Diagnose

Playbooks gather facts and build a recommendation

3

## Request

Up to 3 SQL commands with rollback, risk and expected effect

4

## Pre-check

Privileges, disk space, running maintenance, risk score

5

## Execute

The agent applies the command; locking ones wait for a window

6

## Measure

Metrics before and after decide the next step

OUTCOME OF EVERY STEP

### Improved

Proceed to the next planned step or close the session

### Degraded

Automatic rollback from a snapshot; escalation in manual mode

### No effect

Escalated to an engineer with a full log of actions and measurements



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SELF-HEALING

NEW · RELEASE 7.1

# What AI can fix automatically

## Problem classifier

- Vacuum lag
- Slow queries
- Locks
- Connections
- Memory
- Config drift
- Wraparound
- Observability

### ESCALATED TO A HUMAN ONLY

- Host saturation
- Disk space

## Table maintenance

VACUUM, VACUUM ANALYZE, ANALYZE, VACUUM FULL, extended statistics

## Indexes

Create, drop and rebuild — including CONCURRENTLY

## Configuration

A parameter or a batch of parameters, table-level autovacuum

## Sessions

Cancel a stuck query or terminate a session

*If there is no unambiguous command, AI does not invent a fix — it hands the decision to a human*



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SELF-HEALING SAFETY

NEW · RELEASE 7.1

# Layered protection at every step

### Two modes

Manual — an Apply button. Automatic — by rules per problem type for an instance, workspace or tenant

### Risk score 1–5

Max of platform and model scores: locking, scope, irreversibility, duration. Auto-run threshold: 4

### Second opinion

An independent review of the command's properties and the case against it. Not confirmed — no auto-run

### Hard bans

DML, TRUNCATE, dropping data, GRANT and roles, dangerous parameters and extensions, system schemas

### Limits and windows

Up to 3 commands, lock_timeout 15 s, locking operations only in a maintenance window

### Rollback

Snapshots of configuration, indexes and tables; automatic rollback when metrics degrade



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

MEASURABLE EFFECT

NEW · RELEASE 7.1

# Every fix is confirmed by metrics

1. **Baseline**
   Before the command, the platform records key metric values
2. **Observe**
   A pause so the workload reflects the change
3. **Re-measure**
   A 30-second window, up to 3 attempts; verdict: improved, no change or degraded

No false conclusions: a measurement is discarded under unrepresentative load (fewer than 5 TPS or 20 calls)

## EFFECT METRICS

- Dead tuple ratio
- Mean query time
- Lock wait time
- Connections used
- Cache hit ratio
- XID age



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

SELF-HEALING IN ACTION

NEW · RELEASE 7.1

# Solved by AI — or fixed by your operator

*Illustrative example*

## AI recommendations

AI | Triggers | Chat

prod-db-01 · sales — 2 h ago
**Bloat on orders**
Solved by AI — Dead tuple ratio 34% → 3%

prod-db-01 · sales — 5 h ago
**Heavy query on orders**
Fixed by user — Avg query time 820 → 35 ms

billing-db · core — now
**Stale statistics on payment_journal**
AI is fixing — Step 2 of 3 · since 14:05

etl-host-03 — 1 d ago
**Host CPU overload outside PostgreSQL**
Self-healing unavailable — Escalated to an engineer

## Bloat on orders

Solved by AI

**Problem resolved automatically**

### WHAT THE AGENT DID

| | |
|---|---|
| **Identified the problem** | Vacuum lag on public.orders |
| **Prepared a fix plan** | 1 command · risk 2 of 5 · no table lock |
| **Attempted a fix** | `VACUUM (ANALYZE) public.orders` |
| **Execution result** | Got better |

### REQUEST METRICS: BEFORE → AFTER

| | |
|---|---|
| Dead tuple ratio, % | 34.2 → 2.9 |
| Average query time, ms | 41 → 18 |



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

AI ASSISTANT

# Ask about your database, get a short, precise answer

### Plain-English questions

Alerts, incidents, CPU and memory load — with live platform data

### Knows the documentation

RAG over product docs: hybrid search (vector + BM25), multilingual embeddings

### Your model, your perimeter

Any OpenAI-compatible endpoint: your own server, vLLM, Ollama or a cloud API

### Always at hand

A side panel with AI, Triggers and Chat tabs

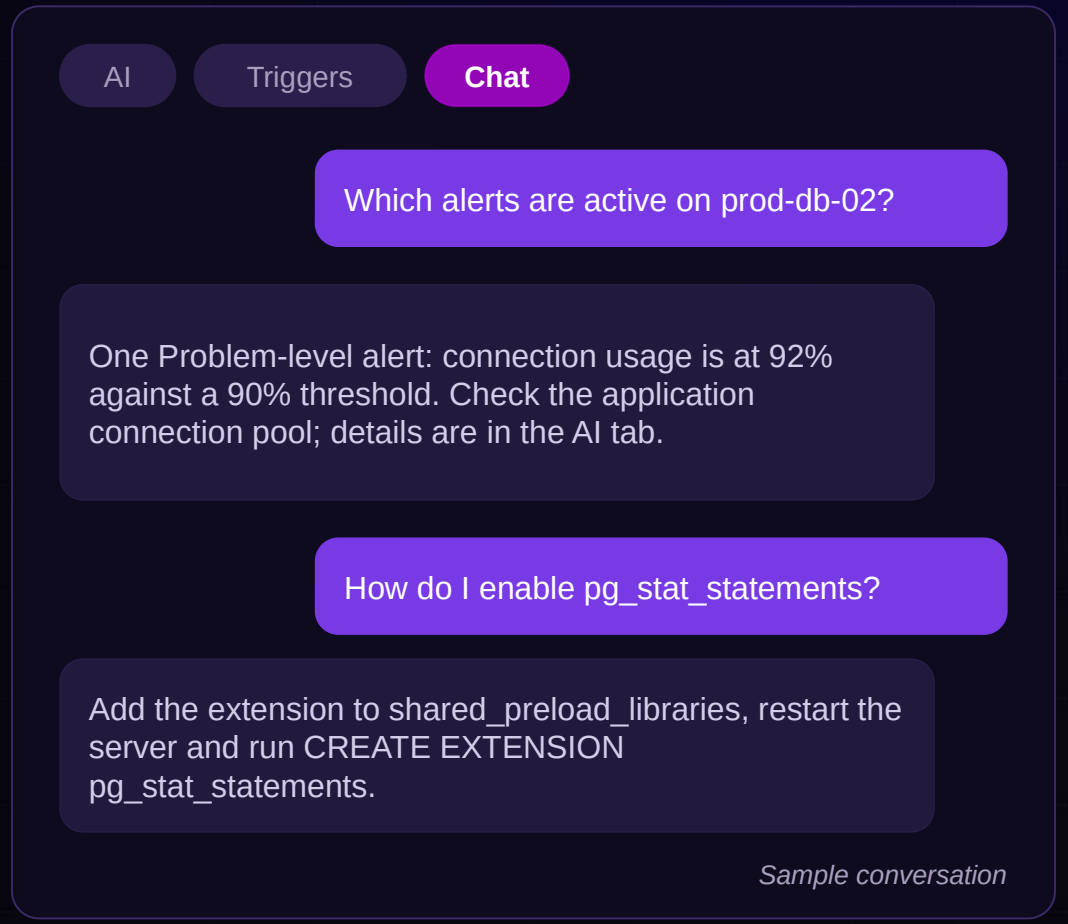


*Sample conversation*



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

DEPLOYMENT

# Deploy where your data lives

### Cloud marketplaces

AWS, Azure, GCP (coming soon) — runs inside your own cloud account

### Self-hosted

Docker Compose and a CLI installer: checks, secrets, TLS; upgrades with backup

### High-availability setup

HA on VMs via Ansible: NATS cluster, VIP for nginx, dedicated databases

## PLATFORM SERVER SIZING

| Instances | vCPU | RAM | SSD |
|---|---|---|---|
| **10** | 4 | 8 GB | 100 GB |
| **100** | 12 | 24 GB | 500 GB |
| **1,000** | 18 | 32 GB | 1–2 TB |
| **5,000*** | 32 | 64 GB | 5–10 TB |

* above 1,000 instances — high-availability deployment only

Licensed per CPU core

AI as a separate module



© 2026 Awide Labs Ltd. All rights reserved.

awidelabs

awidelabs

# From firefighting to foresight

Answers before questions are asked. Solutions before problems become incidents.

**Request a demo**



**awide.tech**
info@awide.io
docs.pgsuite.awide.tech