

PostgreSQL Health & Performance Pack

Discovery questionnaire — score each item honestly; low scores are the point: they show where the program will hurt.

Pack	PostgreSQL Health & Performance Pack
Document	Discovery questionnaire (24 questions)
Response scales	Maturity 1–5 · Yes / Partially / No · Written
Date	2026-09

AnovaCloud — Engineering the AI-Native Enterprise. Frisco, Texas. This pack is a working instrument, not a brochure: complete it with your team, score honestly, and bring it to your discovery call. © 2026 AnovaCloud.

HOW TO USE

Run this like a workshop, not a survey

Complete it with the stakeholders named in each section's intro. One scribe, one room (or call), 90 minutes. Score from evidence — monitoring data, documents, demos — not opinions. Any item scoring 1–2 is a program risk: it needs an owner and a close date before you commit budget. Bring the scored questionnaire to your discovery call; it compresses weeks of fact-finding into one session.

1 · INSTANCE & CONFIGURATION**Instance & configuration**

Defaults are not a configuration strategy.

Q1. What PostgreSQL major version is running, and how far behind current stable?

Evidence: `SELECT version()` output per cluster.

Q2. Have `postgresql.conf` settings been sized for this hardware/workload (shared_buffers, work_mem, maintenance_work_mem, max_connections)?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: config file with sizing rationale.

Q3. Is connection pooling (PgBouncer/pgpool) in place, and what is the connection churn rate?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: pooler stats; connections/sec on the primary.

Q4. Are checkpoints tuned (checkpoint_timeout, max_wal_size) — or do you see checkpoint spikes in I/O graphs?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: I/O graph around checkpoints.

2 · WORKLOAD & QUERY PROFILE**Workload & query profile**

Slow queries are a workload problem before they are a database problem.

Q5. Is `pg_stat_statements` enabled, and is there a ranked top-query list by total time?

☐

Yes

☐

Partially

☐

No

Evidence: top-20 queries by total_time.

Q6. What share of top-query time is sequential scans on large tables that should be indexed?

Evidence: EXPLAIN ANALYZE samples of top offenders.

Q7. Are N+1 patterns, missing pagination, or unbounded result sets present in application code?

☐

1

☐

2

☐

3

☐

4

☐

5

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: code review notes or ORM query logs.

Q8. p95/p99 latency for the critical read and write paths — measured, not estimated?

Evidence: APM or proxy latency percentiles.

3 · STORAGE, VACUUM & BLOAT**Storage, vacuum & bloat**

PostgreSQL's storage engine needs active housekeeping.

Q9. What is the estimated bloat percentage on the largest 20 tables?

Evidence: bloat-estimate query output.

Q10. Is autovacuum keeping up (n_dead_tup staying bounded), or are manual VACUUMs routine?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: dead-tuple trend over 7 days.

Q11. Are long-running transactions or idle-in-transaction sessions blocking vacuum?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: pg_stat_activity snapshot during peak.

Q12. Is table partitioning used where it should be (time-series, large append-mostly tables)?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: partitioning design or its absence, justified.

4 · HA/DR & BACKUPS

HA/DR & backups

An untested backup is a rumor.

Q13. What is the HA topology (streaming replicas, quorum, failover mechanism) and measured RPO/RTO?

Evidence: architecture diagram + last drill timings.

Q14. Are base backups + WAL archiving (pgBackRest/WAL-G) configured and monitored?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: backup job history, last 30 days.

Q15. When was the last successful PITR restore drill to a point in time?

Evidence: drill log with timings. Never = score 1.

Q16. What is the replication lag distribution (p50/p99) under peak write load?

Evidence: lag metrics during peak.

5 · SECURITY & ACCESS**Security & access**

The database holds the crown jewels; act like it.

Q17. Is authentication certificate/SCRAM-based with least-privilege roles per application?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: role inventory and grant review.

Q18. Is data encrypted in transit (TLS) and at rest, with managed key rotation?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: encryption configuration + rotation log.

Q19. Are superuser and replication credentials vaulted, rotated, and audited?

1	2	3	4	5
---	---	---	---	---

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: vault/rotation policy.

Q20. Is audit logging enabled for DDL and sensitive-table access?

☐	Yes	☐	Partially	☐	No
--------------------------	-----	--------------------------	-----------	--------------------------	----

Evidence: sample audit log entries.

6 · OBSERVABILITY**Observability**

You cannot tune what you cannot see.

Q21. Are key PostgreSQL metrics (cache hit, checkpoints, locks, replication, bloat) on dashboards with alerts?

☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: dashboard screenshots + alert rules.

Q22. Is slow-query logging sampled and reviewed on a cadence (not just when something breaks)?

☐ Yes ☐ Partially ☐ No

Evidence: review cadence + last review notes.

Q23. Can you attribute database load to application, query, and user?

☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5

Maturity 1-5 (1=Ad hoc, 2=Repeatable, 3=Defined, 4=Managed, 5=Optimized)

Evidence: example attribution for last week's top load.

Q24. Is there a runbook for the top 5 production incidents (failover, bloat crisis, runaway query)?

☐ Yes ☐ Partially ☐ No

Evidence: runbook links.

SCORING

Read the pattern, not just the average

Average each section. Sections averaging below 3 are your critical path — sequence the project plan so those risks close first. A single 1 on a load-bearing question (downtime budget, rollback, access control) outweighs a 4 average elsewhere. Record owners and close dates below.

Lowest-scoring section:

Single biggest risk:

Owner:

Close date:
