DEVOPS · DIAGNOSTIC DELIVERABLE

DevOps Maturity Assessment · DORA Report — SAMPLE

4-week assessment · DORA 2023 benchmark calibration

Client	Empresa Anônima S.A.
Engagement ID	#ANV-DORA-2604-1107
Period	Q2 2026 · April 13 – May 11, 2026
Issued by	Anuvia · Mila Vernazza, Principal Consultant
Classification	SAMPLE · Confidential · Not for distribution

This document is a sample preview of an Anuvia diagnostic deliverable. All figures, customer names, and findings are illustrative and anonymized. No real-client data is contained within.

anuvia.com.br · contato@anuvia.com.br

DORA METRICS DASHBOARD

Current state vs. DORA 2023 benchmarks

Measured over the 90-day window preceding kickoff. Source: GitLab pipelines, PagerDuty incident records, deployment ledger. Classification per the DORA 2023 *State of DevOps* Elite/High/Medium/Low banding.

Metric	Current	Tier	High threshold	Elite threshold
Deploy Frequency	Weekly (avg 1.4 / wk)	Medium	Daily	On-demand (multi/day)
Lead Time for Changes	8 days	Medium	<1 day	<1 hour
MTTR	4 hours (P50)	High	<1 day	<1 hour
Change Failure Rate	22%	Medium	0-15%	0-5%

Tier summary

- Current overall tier: **Medium** (3 of 4 metrics in Medium band).
- Closest to High: MTTR — already at the High threshold ceiling.
- Largest gap: Lead Time — 8× the High threshold.
- Highest CFR contributor: insufficient automated regression coverage on the checkout path (CFR-Δ +14 points).

*Target posture in 6 months: **High** across all four metrics. Roadmap on page 8.*

METRIC DEEP-DIVE · 01 / 04

Deploy Frequency — Weekly

Current	1.4 deploys / week (production) · 86 in last 90 days
Tier	Medium (DORA 2023)
Target (6 mo)	Daily — High tier (≥5/wk per service)
Gap explained	Batch-style release cadence; manual approval gate on every prod deploy; 6-hour smoke windo

Remediation actions

- Replace human approval on low-risk paths (toggled features, internal services) with policy-as-code gate (OPA).
- Add canary deploy strategy → eliminate 6h smoke window; rely on telemetry-based promotion.
- Decompose monolithic release pipeline into per-service deploys (currently 3 services share 1 pipeline).
- Adopt trunk-based development on the 2 highest-velocity services in the first 8 weeks.

METRIC DEEP-DIVE · 02 / 04

Lead Time for Changes — 8 days

Current	8 days (median commit → production)
Tier	Medium
Target (6 mo)	<24 hours — High tier
Gap explained	PR review queues average 2.4 days; integration test suite takes 38 min; manual UAT cycles ad

Remediation actions

- PR aging SLO: auto-nudge after 24h; round-robin reviewer assignment.
- Parallelize integration suite (currently sequential) → target 8 min wall time.
- Shift UAT to ephemeral preview environments per PR; reduce manual UAT to release-train cadence only.

METRIC DEEP-DIVE · 03 / 04

MTTR — 4 hours (P50)

Current	P50 4h · P90 11h · 28 production incidents in last 90 days
Tier	High (just within band)
Target (6 mo)	<1 hour P50 — Elite tier
Gap explained	Detection-to-diagnosis time dominates (avg 2.1h); runbook coverage on top-5 services only; ob

Remediation actions

- Service-level objective definitions + alarms for all production services (currently 60% coverage).
- Runbook coverage to 100% of services with traffic above defined threshold.
- Adopt structured incident command (IC role) — currently ad-hoc resolution.
- Investigate exemplar-driven debugging (OpenTelemetry exemplars) for tail-latency cases.

Incident timeline · last 90 days

Severity	Count	Median MTTR	Recurrence (root same?)
SEV1	2	1h 50m	0
SEV2	9	3h 40m	2
SEV3	17	5h 20m	5

METRIC DEEP-DIVE · 04 / 04

Change Failure Rate — 22%

Current	22% (19 of 86 production deploys triggered rollback or hotfix)
Tier	Medium
Target (6 mo)	<15% — High tier (stretch: <10%)
Gap explained	Checkout-path service shows 41% CFR vs. fleet 22% — outsized contributor. Insufficient regre

Remediation actions

- Expand contract-test coverage on payment-gateway integration; add 11 specifically-identified edge cases.
- Adopt feature-flag-based rollouts on checkout-path; 1% → 10% → 50% → 100% gates with auto-rollback on error budget burn.
- Pre-prod load profile in staging using sanitized production-traffic shadow.

CFR by service

Service	Deploys	Failures	CFR
checkout-api	22	9	41%
catalog-api	14	2	14%
identity-service	11	1	9%
search-service	18	4	22%
notification-worker	13	2	15%
billing-service	8	1	13%

PIPELINE & TOOLING ASSESSMENT

CI/CD · IaC · GitOps readiness

CI / CD maturity

Dimension	Current	Target
CI tool	GitLab CI (self-hosted runners)	Stay
Build time (P50)	12 min	≤6 min
Test layers	Unit · Integration · Smoke	Add contract + canary
Parallel jobs	Sequential	Fan-out by service
Test flakiness	8% job-flake rate	<1%

IaC adoption

Domain	Coverage	Gap
Networking (VPC, TGW, NAT)	62%	Legacy peering still click-ops
Compute (ECS, EKS)	94%	—
Data (RDS, ElastiCache)	78%	Parameter groups untracked
IAM / SSO	55%	Manual SSO group sync
Observability (alarms, dashboards)	31%	Mostly click-ops

GitOps readiness

- Argo CD partially adopted (1 cluster, 4 apps). Recommend extending to all clusters by W12.
- Application manifests live in 3 repos with inconsistent structure — consolidate before scaling Argo CD.
- App-of-Apps pattern recommended for the second wave.

INCIDENT RESPONSE REVIEW

Runbook · on-call · post-mortem culture

Runbook coverage

Service tier	Services	Runbook coverage	Last review
Tier-0 (revenue path)	5	5/5 · 100%	Avg 41 days ago
Tier-1 (customer-facing)	11	7/11 · 64%	Avg 6 months
Tier-2 (internal)	17	3/17 · 18%	>12 months

On-call rotation

- Single rotation (8 engineers, 1 week shifts). No secondary rotation → escalation gaps observed in 4 incidents.
- Avg page volume: 11 / week (high — investigate alert quality before adding people).
- Compensation policy informal — recommend codifying for retention.

Post-mortem culture

- 9 of 14 incidents (last 12 mo) have a post-mortem document. 5 missing — predominantly SEV3.
- Quality: 3/9 reach blameless standard; rest read as RCA-only without contributing-factors analysis.
- Action item closure rate: 47% within target date.

ROADMAP · 6 MONTHS TO HIGH TIER

Sequenced quarterly plan

The path from Medium to High tier converges on three structural moves: (1) shorten the path from commit to production via parallelization & policy-as-code gates, (2) reduce variance by feature-flag rollout discipline, (3) raise observability coverage to make MTTR collapse achievable.

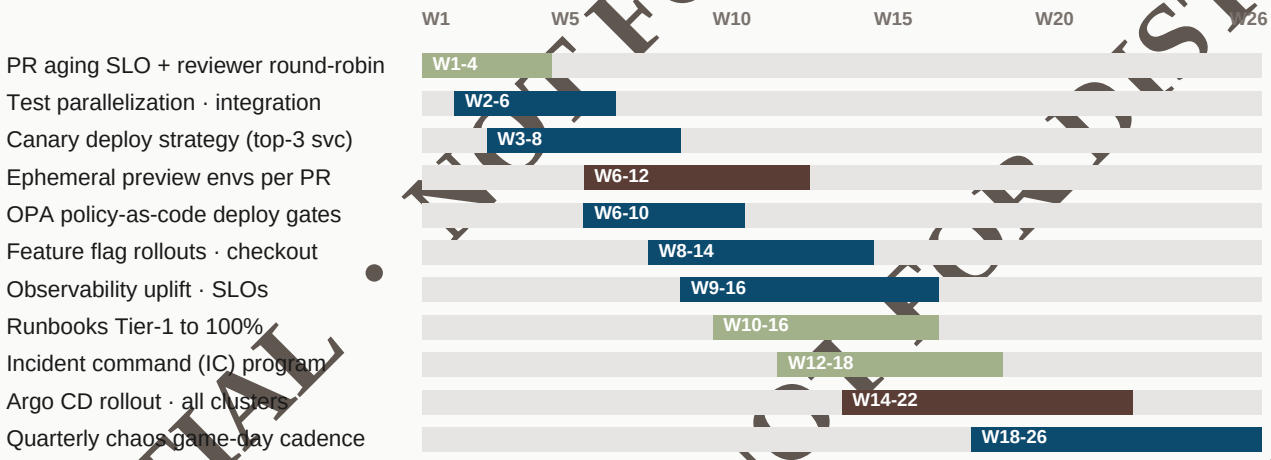

Target metrics (W26)

Metric	W0 (today)	W13	W26 target
Deploy Freq	Weekly	2-3 / day	Daily+
Lead Time	8 days	2 days	<24h
MTTR	4h	2h	<1h
CFR	22%	16%	<15%

APPENDIX A · DATA SOURCES

Methodology & evidence

Data sources

- GitLab CI/CD events · 90-day window · all pipelines
- PagerDuty incident archive · 90-day window
- Git history · all repos · 90-day commit log
- Production deploy ledger (custom audit channel)
- Interviews · 9 engineers + 3 EM + 1 VPE (12h total)

Calibration

DORA tier classifications use the 2023 DORA *State of DevOps Report* banding. We do not invent thresholds — the bands cited (Elite / High / Medium / Low) are verbatim from that survey. CFR is calculated as (rolled-back-or-hotfixed deploys / total prod deploys) per the same survey definition.

APPENDIX B · GLOSSARY

Glossary

Term	Definition
DORA	DevOps Research & Assessment · annual State of DevOps Report.
CFR	Change Failure Rate · % deploys requiring rollback or hotfix within a defined window.
MTTR	Mean Time to Restore · from incident detection to restored service.
Lead Time	From source commit to production deployment.
SLO	Service Level Objective · target performance threshold.
IC	Incident Commander · single coordinator during an active incident.
GitOps	Operational model where the desired state lives in Git; controllers reconcile reality.
OPA	Open Policy Agent · policy-as-code engine for admission control.
Argo CD	Kubernetes GitOps continuous-delivery controller.

— End of report — Anuvia · contato@anuvia.com.br · anuvia.com.br