OpsLab Infrastructure Audit Report

SAMPLE REPORT

Generated: 2026-07-24 08:22 UTC | Banking / Payment Processing | Confidential — for authorised recipients only

6.5% of incidents caused 80% of downtime — and a 7.1% chance of any outage running past 6 hours

This is a real run on anonymized bank incident data — not a mockup. Get a free audit: t.me/OpsLaboratory_bot

CRITICAL ISSUES DETECTED

802	199 min	2664.8 h
Incidents	Avg. Resolution	Downtime

Your Network at a Glance

Infrastructure type	Banking / Payment Processing
Analysis period	344 days (8245 h)
Not matched in data	atm network, internet banking, payment gate, processing center

Next Expected Window

Expected incidents	7d: ~23 (14–33), 30d: ~99 (80–119)

Summary

Over 343.5 days, 802 incidents caused 2664.8 hours of downtime, with extreme concentration: 6.5% of events accounted for 80% of total impact. The incident rate doubled after a regime shift in February 2022 and remains elevated, with ~23 incidents forecast for the next 7 days. Heavy-tailed recovery times make multi-hour outages likely, requiring urgent focus on the top few critical failures.

Key Issues

R1. Severe downtime concentration: few incidents dominate impact — CRITICAL

R2. Incident rate doubled after February 2022 changepoint — HIGH

R3. Heavy-tailed downtime: high probability of multi-hour outages — HIGH

Recommended Actions

1) Analyze top downtime incidents and harden recovery [Immediate]

2) Investigate the February 2022 rate increase trigger [Short-term]

3) Strengthen major-incident escalation for processing center [Long-term]

Full technical analysis on the next page.

Risk Analysis

#	Issue	Severity	Downtime
R1	**R1 — Severe downtime concentration: few incidents dominate impact** A handful of exceptional incidents dominate total downtime, so targeted improvements on those top causes can yield disproportionately large reliability gains. (2022-03-30 00:50:42) (2022-07-01 14:43:18) (2022-08-13 20:25:52) services: Процессинговый центр, События мониторинга, ККС 015 Переводы Pareto analysis shows strong concentration (Gini=0.896, concentration=strong). Only 6.5% of incidents cover 80% of total downtime, and the top 3 incidents alone contribute 26.7%. Probable cause: inadequate recovery processes for specific critical failures, such as the multi-day processing center outage.	CRITICAL	—
R2	**R2 — Incident rate doubled after February 2022 changepoint** The sustained higher incident rate since early 2022 suggests a permanent systemic change that continues to strain operations. Changepoint (CUSUM+permutation) detected a regime shift around 2022-02-19: daily incident rate rose from 1.63 to 3.27 (+100.2%, p=0.004). non-parametric comparison test (Mann-Whitney U) confirmed the shift with medium effect (p=3.5e-05, effect size (Cliff's δ)=+0.369, 95% CI [+0.219, +0.524]). Probable cause: an undetected system change or external factor that has not been reversed.	HIGH	—
R3	**R3 — Heavy-tailed downtime: high probability of multi-hour outages** Even seemingly short incidents have a non-trivial chance of escalating into multi-hour outages, indicating weak containment. Tail-risk analysis (GPD fit, Peaks-Over-Threshold) shows heavy distribution (shape=+1.098, 95% CI [+0.707, +1.528]). P(downtime > 1.0 h) = 25.9%, P(downtime > 6.0 h) = 7.1%. Probable cause: lack of fast failover or escalation for critical services, allowing minor issues to cascade into prolonged outages.	HIGH	—
R4	**R4 — Prolonged MTTR reveals slow recovery from major incidents** The average recovery time is inflated by a few crippled recoveries, masking generally fast resolution and hiding improvement opportunities. (2022-03-30 00:50:42) services: Процессинговый центр Mean MTTR is 199 min, but median is only 14 min with P95=552 min, driven by a few extremely long incidents (worst 17522 min). This skew indicates that while most issues resolve quickly, a subset of critical failures suffer from inadequate recovery processes. Probable cause: missing escalation paths or runbooks for large-scale processing center failures.	MEDIUM	—

#	Issue	Severity	Downtime
R5	**R5 — High incident-rate forecast signals ongoing operational load** If the current incident rate persists, the team will face roughly 100 incidents per month, making it harder to focus on long-term fixes. Incident-rate forecast (Poisson, λ=3.29/day, post-changepoint segment) projects 23.0 incidents in the next 7 days (95% PI [14, 33]) and 98.7 in the next 30 days (PI [80, 119]). Probable cause: the elevated baseline rate established after the February 2022 shift remains stable, driving consistent inflow of events.	MEDIUM	—

Early Warning — Which Event Comes First

Cross-correlation of event streams (time bin: — min)

Early warning not available: at least two event streams with enough overlap are required.

Regime Shifts — When the Behaviour Changed

CUSUM with a permutation test on the daily series: incident_rate_daily

Date	Before (per day)	After (per day)	Change	p-value
2022-02-19	1.63	3.27	+100.2%	0.004
2022-04-01	3.27	2.19	-33.1%	0.006
2022-05-30	2.19	3.28	+50.0%	0.006

Tail Risk — Probability of a Very Long Outage

Peaks-Over-Threshold, Generalized Pareto Distribution (GPD) fit on long incidents

GPD shape ξ = +1.098 (95% CI [+0.707, +1.528]); fitted on 81 events longer than 3.8 h. Tail behaviour: heavy — very long outages are materially more likely; plan SLA and redundancy for this.

Outage longer than	Probability	95% CI
1 h	25.94%	[22.81%, 28.93%]
2 h	16.33%	[13.96%, 18.95%]
6 h	7.06%	[5.19%, 8.77%]

Recommendations

1) Analyze top downtime incidents and harden recovery [Immediate · Medium]

Conduct a detailed root cause analysis and post-mortem for the top 3 downtime incidents (e.g., processing center outage on 2022-03-30). Implement specific infrastructure or process changes to prevent recurrence and reduce recovery time.

2) Investigate the February 2022 rate increase trigger [Short-term · Medium]

Review change logs, deployments, and external service updates around 2022-02-19 to identify what caused the permanent doubling of incident frequency. Candidate areas: payment gateway, processing center, or vendor integrations (Sberbank, NSPK).

3) Strengthen major-incident escalation for processing center [Long-term · High]

Establish an accelerated escalation procedure and dedicated runbook for critical infrastructure (especially processing center) to ensure that major incidents are resolved within a target MTTR of 60 minutes, reducing the current long recovery tail.

Top Problematic Time Windows

Time Window	Incidents	% of Total	Avg. Duration
Wed 10:00	17	2.1%	108 min
Mon 10:00	17	2.1%	71 min
Tue 00:00	17	2.1%	43 min
Thu 14:00	14	1.7%	50 min
Tue 22:00	12	1.5%	754 min

Data Quality Grade: B	Analysis is reliable with minor caveats — see observations above.

Data Quality Note

WARNING: Column 'Actual Resolution Time' is 100% empty; Column 'Checkpoint Time' is 100% empty; Column 'Parent Unique Id' is 99% empty; Column 'Contr Type Alias' is 100% empty; Column 'Fault Chronology' is 70% empty; Column 'Fault Damage' is 67% empty; Column 'Fault Size' is 100% empty; Column 'Fault Loss' is 100% empty; Column 'Fault Reasons' is 100% empty; Column 'Fault Processes' is 100% empty; Column 'Affected Sys' is 100% empty; Column 'Op Telecom' is 97% empty; Column 'Op PS' is 95% empty; Column 'Maj Part' is 96% empty; Column 'Inf Aggr' is 96% empty; Column 'Vendor' is 89% empty; WARNING: 311 time gaps >5 min detected — interpolation may affect quality; WARNING: Missing timestamp/metric rate: 55.9% (>30% threshold; periodicity result may be biased); WARNING: No peaks with FAP <= 0.1 (best FAP: 1); WARNING: Missing timestamp/metric rate: 55.9% (>30% threshold; periodicity result may be biased) [ACF]; WARNING: No lag cleared the 0.047 significance bound [ACF]; WARNING: change_date 2022-02-19 auto-detected by changepoint module (not supplied by client); WARNING: Best silhouette 0.169 below threshold 0.25 — incidents look heterogeneous, no common pattern detected; WARNING: Actual Resolution Time: only 0 points (< 20) — skipped; WARNING: Checkpoint Time: only 0 points (< 20) — skipped; WARNING: Parent Unique Id: only 12 points (< 20) — skipped; WARNING: Contr Type Alias: only 0 points (< 20) — skipped; WARNING: Fault Loss: only 0 points (< 20) — skipped; WARNING: Fault Reasons: only 0 points (< 20) — skipped; WARNING: Affected Sys: only 0 points (< 20) — skipped; WARNING: 83 segments found; showing the 5 busiest plus an aggregate bucket; WARNING: Named critical services (payment gate, ATM network, internet banking, processing center) not found in data; check service names and date range.

METHODOLOGY

Statistical Methods Applied

Method	Purpose	α / Threshold
Lomb-Scargle	Detect recurring failure cycles in time-series data	FAP ≤ 0.10
Mann-Whitney U (two-sided)	Compare incident rates / metrics before and after a change date	α = 0.05
Cliff's δ + Bootstrap CI	Quantify the magnitude and uncertainty of the before/after change	95% CI, 1 000 resamples
K-Means + Silhouette	Group incidents by behaviour (duration, time of day, weekday)	k □ [2..6], S ≥ 0.25
Gini coefficient + Lorenz	Measure concentration of downtime among incidents	—

Significance level α = 0.05 for all hypothesis tests. Tests run only when minimum data requirements are met (see Limitations section).

Before vs After — Change Impact

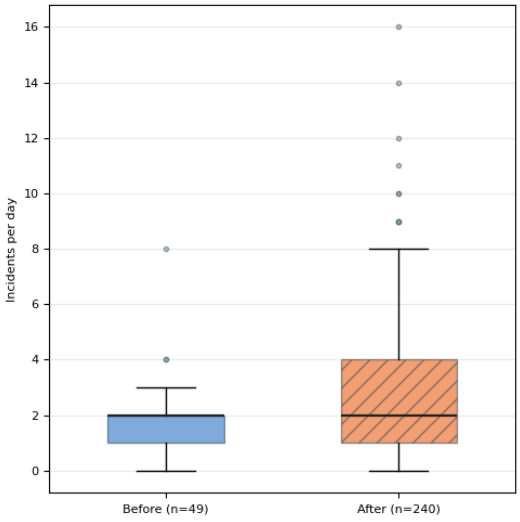

Lorenz Curve — Downtime Concentration

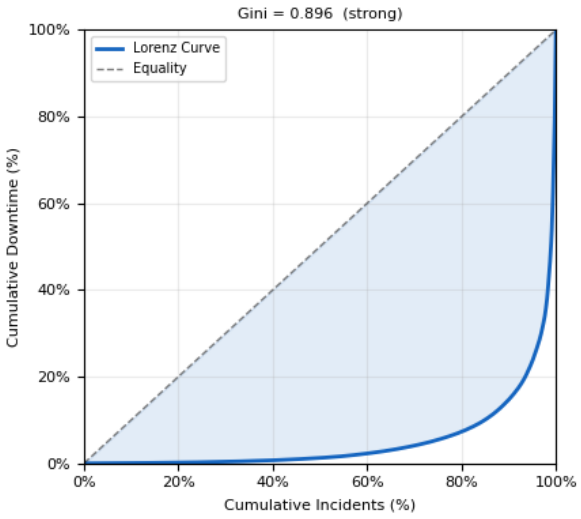

Pareto Chart — Downtime by Incident

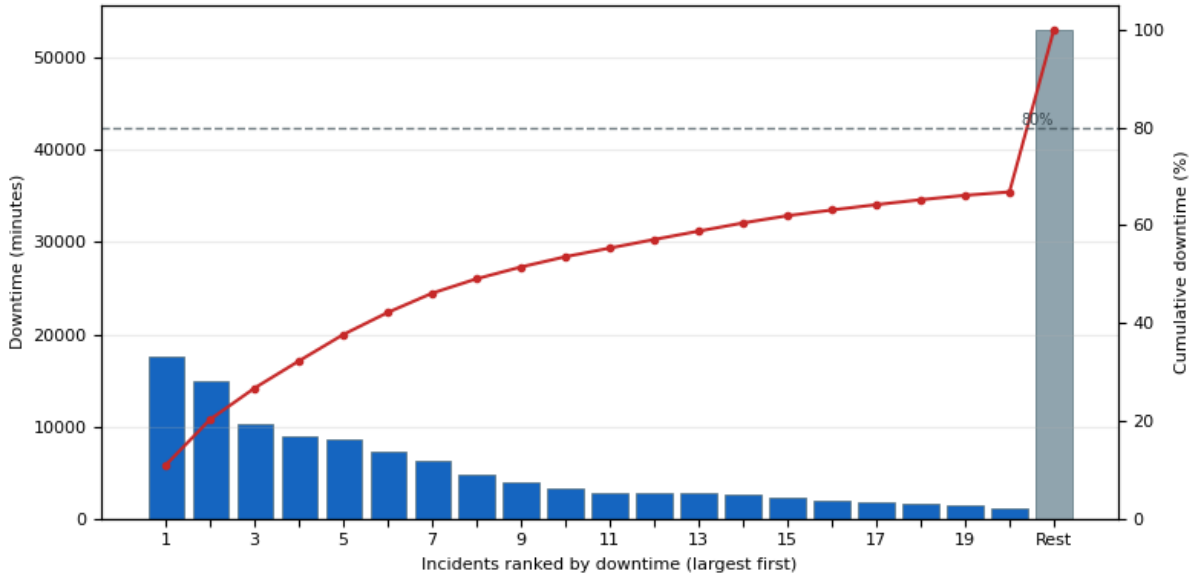

Brief Glossary of Methods Used

p-value	Probability of observing results at least as extreme as measured, assuming no real effect. p < 0.05 is statistically significant (< 5 % chance it is random).
95% CI	95 % Confidence Interval — the range that contains the true value in 95 out of 100 repeated experiments. Wider CI = more uncertainty.
Cluster	A group of incidents with similar behaviour (duration, time of day, day of week). Incidents in the same cluster likely share a root cause.
Gini coefficient	Measures how unevenly downtime is distributed. 0 = all incidents equal; 1 = one incident caused all downtime. Gini > 0.6 = strong concentration.
FAP	False Alarm Probability — how likely a detected cycle is random noise. FAP < 0.01 = strong evidence of a real recurring pattern.