

Data-Quality Impact Report — app_version 1.13.5

(Comparison)

Automated Data Analyst AI Agent · Hexa (read-only)

 **Illustrative sample — SYNTHETIC data.** Every figure below is fabricated for portfolio demonstration and is **not** real production data.
Method: mirrors dq_checks_1.13.5.sql (read-only) · **Self-check (Statement 2): 6/6 PASS ✓**

0 — Headline: before vs. after

Metric	Before	After	Change
Total connection legs	30,120	30,120	0
Successful legs	27,830	20,510	−7,320
Success rate %	92.4	68.1	−24.3
Failed legs	2,290	9,610	+7,320

Raw success rate 92.4% → true success rate 68.1% once the data-quality checks reclassify inflated successes. 'Change' = after – before.

1 — What the checks found, per check (Statement 1)

Check	Flagged	One-and-done A	Type B	Retained	Status
B2 Inactivity churn	5,180	3,400	1,300	480	ALERT
B7 One-and-done	3,020	3,020	0	0	ALERT
T1 Flap (<30s)	1,470	900	400	170	WATCH
#8 Self-replace bug	640	210	300	130	WATCH
B9 Country-hopping	260	120	90	50	OK
#5 Server IP valid	58	20	18	20	OK

One row per check: how many it flagged, split into One-and-done Type A / Type B / Retained.

2 — Which check caused how much of the drop (Statement 3)

Colour key: good / small watch bad / large

Check	Rule	Contribution (pts)
B2 Inactivity churn	no activity > 7 days	14.6
B7 One-and-done	single short session	6.8
T1 Flap (<30s)	dropped < 30s	2.1
#8 Self-replace bug	replace + reconnect < 60s	0.8
Other checks (sum)	—	0.0

Each flipped success is attributed to the FIRST check that catches it, so contributions sum to Statement 0's overall change. Redder = bigger contributor.

3 — Failure taxonomy (recovered from outbounds)

Colour key: good / small watch bad / large

Failure type	Events (legs)	% of failed legs
Handshake / probe timeout	3,910	40.7
Read timeout mid-session	2,540	26.4
TLS / cert reject	1,480	15.4
DNS resolve failure	980	10.2
Other probe failure	700	7.3

Read-out: most failed legs are early handshake/probe timeouts, followed by mid-session read timeouts — pointing at connection setup and keep-alive rather than auth.

4 — Bottom line

Raw reporting overstated success by ~24 points. The true drop is driven by inactivity churn and one-and-done users, not data corruption — so the priority is day-1 retention, while the failure taxonomy says the top *technical* fix is handshake/probe timeouts (40.7% of failed legs). No bot activity was detected.

