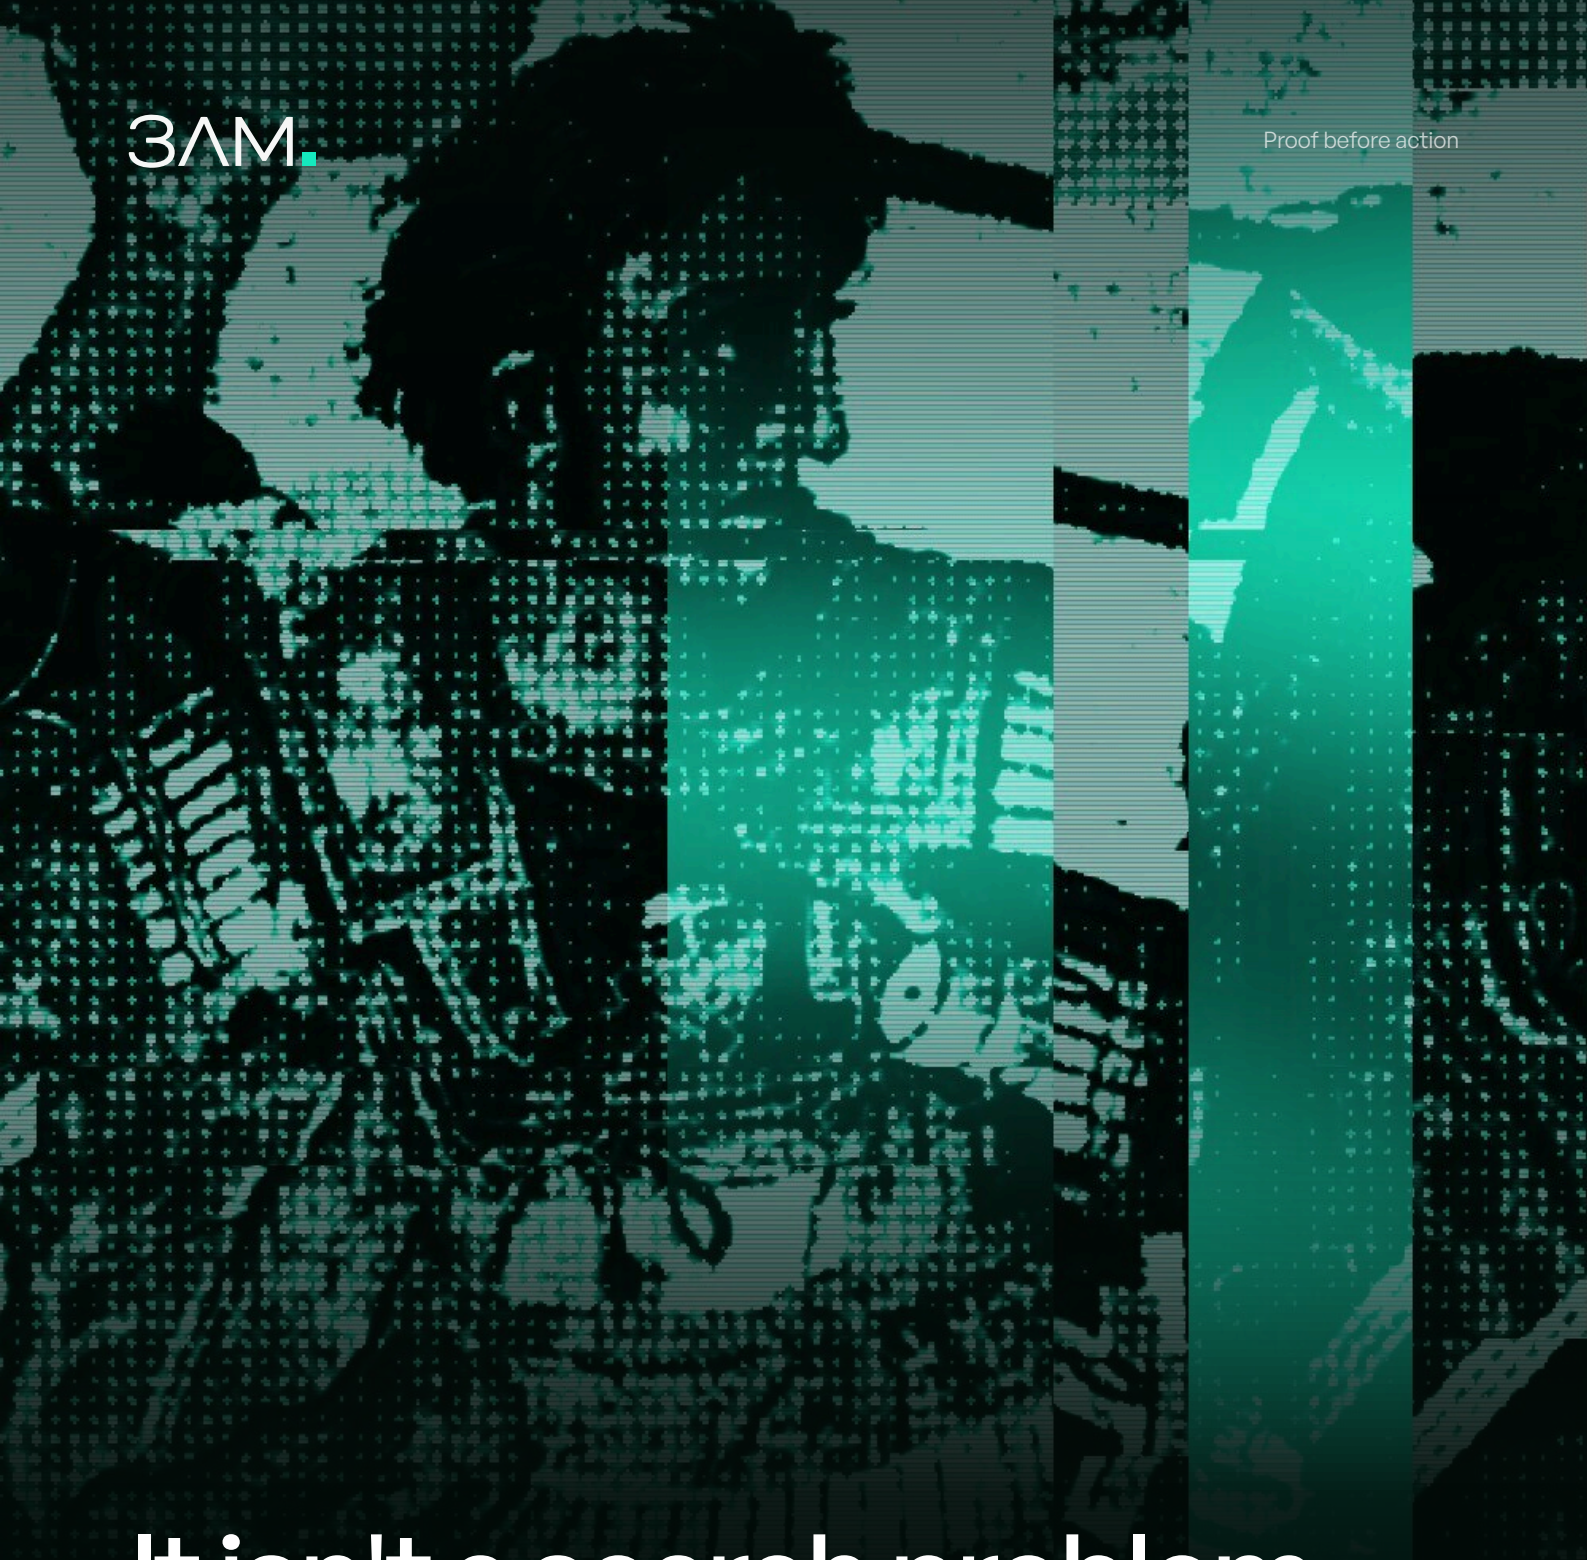


3AM.

It isn't a search problem. It's a proof problem.

Finding a likely cause is the easy half. Every tool on the market reaches a guess. A bank cannot change production on a guess, so 3AM does the other half: it proves the cause, then proves the fix.

Document
Proof before action

For
Engineering leads

Published
November 2026

Production outgrew the war room.

A bank's estate now spans mainframe, cloud and everything between. When it breaks, people still search for the cause by hand, while the clock and the regulator run.

$300M

lost each year to unplanned outages, on average, by a Global 2000 company

Splunk, The Hidden Costs of Downtime 2026

50%

rise in downtime costs across the Global 2000 in two years

Splunk, The Hidden Costs of Downtime 2026

£48.65M

fined to one UK bank by the FCA and PRA over a failed IT migration

FCA, 20 December 2022

17 Jan 2025

when DORA applied: EU banks now report major ICT incidents to their regulator

EIOPA, Digital Operational Resilience Act

It isn't a search problem. It's a proof problem.

Finding a likely cause is the easy half. Every tool on the market reaches a guess. A bank cannot change production on a guess, so 3AM does the other half: it proves the cause, then proves the fix.

Only a proven cause can fix anything.

Every cause 3AM considers carries a grade. The grade decides what it may do, and nothing short of the top grade can touch production.

1	**Hypothesis** A plausible cause, ranked against your estate's history.	**Worth checking. Nothing changes.**
2	**Supported** The evidence agrees, from your logs, metrics and changes.	**Your team is told. Still nothing changes.**
3	**Proven** A live check passed that could only pass if the cause is real.	**Now a fix may be drafted, rehearsed and put to two people.**

---------- Where a confidence score stops: below the top grade, every tool that stops at a score stops here.

How do you know the cause is right?	A live check on your systems that can only pass if the cause is real. A confidence score never stands in for it.
Can the AI raise a grade on its own?	No. The AI can suggest anything it likes, but a suggestion is only the bottom grade. Only evidence from your own systems moves a cause up, and only a proven cause can lead to a fix.

What counts as proof.

Two kinds of evidence can prove a cause. Both come from your live systems, and both could fail if the cause were wrong.

Proven

Caught in the act

bad rows traced straight to the code that wrote them.

Proven

Reproduced

switched on and off at random on a copy, and the fault follows the switch.

What does not count

Evidence	Why it cannot prove a cause	The most it can do
A confidence score	A score of 91% is still a guess.	Rank what to check first
Timing that looks suspicious	Two changes can land on the same servers at once.	Hypothesis
A pattern the failures share	It holds up, but cannot tell two causes apart.	Supported
A runbook or a past postmortem	Documents can suggest. Only evidence can act.	Point at a cause
The AI's own summary	Shown to approvers last, marked untrusted.	Explain, never decide

A fix is also a test

This is also a test of the diagnosis. If the copy doesn't behave as predicted, then 3AM was wrong about the cause, and the proof is thrown out.

One night, worked through.

A routine release went out tonight. 38 of your 257 databases are now wrong.

38	218	1
said yes. Now you know exactly where the problem is, and it's far from everywhere.	said no. Properly clean, and 3AM can show you the moment each one was checked.	didn't answer. It doesn't assume the best. That server stays unknown until it replies.

Then it works out what the 38 have in common, and writes it in a line you can read for yourself:

```
release = 2026.10.3 ∧ flag.accrual_v2 ∧ schema ≥ 57
```

The catch

Tonight there's a catch: an unrelated flag went out to the very same servers, so no statistic can tell them apart. 3AM waits at Supported until the bad rows settle it. They do. It's the release's new accrual code, which only runs behind one flag.

Rehearsed on a clone first

1,842 = 1,842	0 rows	9f2c = 9f2c	0.8 s
It changed exactly the rows it predicted	Run a second time, it changes nothing	Undone, the data is back exactly as it was	It never held its locks for long

By morning, you have your answer

38/38	219	0
verified, bad rows left: 0	untouched, hash 9f2c41e7 = 9f2c41e7	unknown, silent server answered

Example figures, from the walk-through at www.3am.si/how-it-works.

When nothing proves it.

Some nights the evidence will not settle. 3AM is built for those nights too.

1 **It holds the grade it has earned.** A cause that is only supported stays supported. It may draft a fix, never run one.

2 **It says so, in plain words.** Your team sees what was checked, what passed, and what is still open.

3 **It hands over to a person,** with every check, every server's answer and the evidence so far.

4 **It keeps watching.** When the evidence arrives, the grade moves, and the decision comes back to your approvers.

Numbers that belong to your estate

Your systems learn their own model of cause and effect. Nothing is shared between banks. Alerts come with a stated false-alarm rate, measured on your systems, and the shadow record shows what it has earned before it is allowed more.

Day 0

Connected

It maps your estate and can already prove a cause on a live system. Every service starts in shadow.

Weeks 2 to 6

Knows your normal

It learns your estate over full business cycles. Alerts come with a stated false-alarm rate.

After 10 to 30 incidents

Knows your estate

Recurring faults are proven faster, and the shadow record shows when it has earned more.

Bring us the incident nobody could explain.

See it on your estate
www.3am.si/demo

Write to us
hello@3am.si

Company
3AM Super Intelligence Inc.