3AM.

In brief



Others guess. 3AM proves.

Every AI on-call tool reaches a guess. Banks need more than a guess before anything touches production.

Document	For	Published
3AM in brief	Everyone	November 2026

What it does.

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.

Proven, not probable

A live check that can only pass if the cause is real. No confidence score stands in for it.

Your data never moves

Runs inside your network, air-gapped if you need. Nothing to copy, nothing to send us.

Two keys, never the books

Two people approve every write. No setting gives it a key to your ledger.

Where the others stop.

Step	Typical AI on-call	3AM
Find the likely cause	Yes, with a confidence score	Yes, ranked against your estate's own history
Prove it	Hands you the hypothesis	A live check on your systems confirms it
Rehearse the fix	Suggests a command or a pull request	Dry-runs the exact fix before anyone is asked
Approve	One person, or set to run on its own	Two different people, on one plan that cannot change
Apply and verify	You run it, then check	Applies it, then proves the system recovered
The books	Not addressed	Never touched. A person runs a signed package

A real night, and what it proved.

Nothing has crashed. Nobody has been paged. But bad rows are being written right now, and only a handful of your people are even allowed to look. So: which 38? Why? And how do you fix them without touching the other 219?

A night, start to finish

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

1. **Asks all 257.** 38 said yes, 218 said no, 1 didn't answer.
2. **Proves why.** Bad rows traced to the release code that wrote them.
3. **Stops the damage.** Two people switch the flag off; new bad rows fall to zero.
4. **Rehearses the fix** on a clone of each server, row by row.
5. **Repairs in waves,** one server first, chosen at random.
6. **By morning:** 38/38 verified, 219 untouched, 0 unknown.

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

Only a proven cause can fix anything.

Hypothesis

A plausible cause, ranked against your estate's history.

Worth checking. Nothing changes.

Supported

The evidence agrees, from your logs, metrics and changes.

Your team is told. Still nothing changes.

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 it runs, and what a pilot looks like.

0	0	2	100%
inbound ports to open	bytes of incident data sent to us	approvers to raise autonomy	of actions on the ledger

Where does it run?	Wholly inside your network, air-gapped if you need. The only call to us is a signed licence check, and air-gapped it verifies offline.
Does our data leave?	No. Each question runs where the data lives and a small answer comes back. Nothing is copied, indexed or sent to us.
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.
Who can make it change something?	Two different people, approving one plan that cannot change after they sign. No admin setting removes the second.

Useful on day one. Sharper every incident.

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.



Your incidents never leave. Nothing changes without a human.

Letting software near production should worry you. So here is everything your security team will ask, in the order they'll ask it.

Document
Security and architecture overview

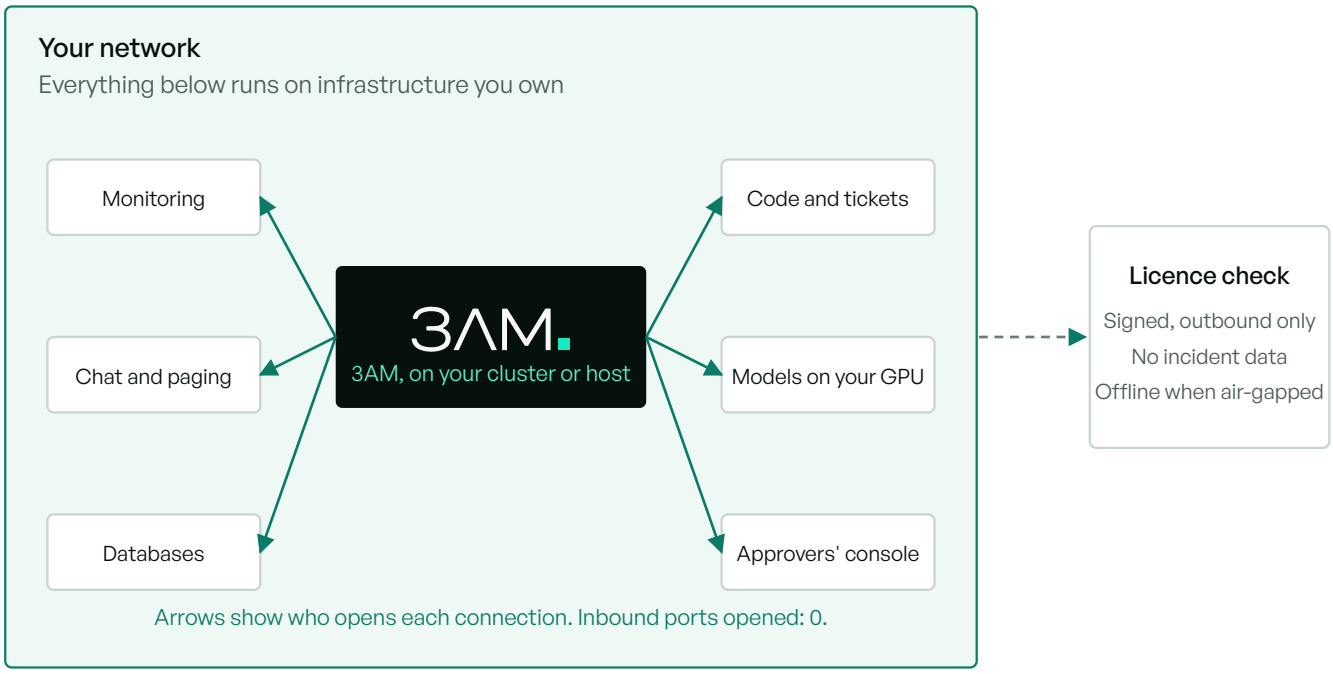
For
Security and risk

Published
November 2026

At a glance.

Built for environments where every action needs an owner and every decision needs a record.

0	0	2	100%
inbound ports to open	bytes of incident data sent to us	approvers to raise autonomy	of actions on the ledger



A Helm chart or a single host, inside your network. Models run on your GPU, or behind an endpoint you choose. Closed networks get an air-gap bundle with the weights included.

First, can anyone get in? No.

3AM runs inside your network and opens no inbound ports. Probe every one of them from outside, and none answer. It only ever calls out, and you decide how far.

	Connected	Air-gapped
What it calls out to	Calls out only to tools you connect, and a signed licence check.	Calls out to nothing. The licence verifies offline.
The licence	A signed check, outbound	Verifies offline
The models	On your GPU, or an endpoint you choose	Shipped in the bundle
Approvals	In the console, or in your chat tool	In the console
Inbound ports	None	None

- ✓ Helm chart or single-host bundle
- ✓ Air-gap bundle, weights included
- ✓ Outbound only, to tools you connect

Do we need to open inbound firewall ports?	No. Slack uses Socket Mode, Symphony its datafeed, and PagerDuty and other channels are polled. Every connection is opened from inside your network.
Can it run in an air-gapped environment?	Yes. The air-gap bundle ships the images and model weights together, so nothing needs to be fetched at install time.
How does it handle disaster recovery and several installs?	A licence can name your DR install, or carry grace hours for failover. On Enterprise and Strategic, Keychain sets up every install once, and that setup moves only between your own installs.

Next, what does it see? Never your secrets.

Logs are full of things that shouldn't travel. So passwords, tokens, keys and card numbers are redacted on the way in, before anything is written. Prompts are kept redacted too, with a hash for proof.

```
db.connect user=svc_3am [REDACTED] password=hunter2-prod! [REDACTED]
slack.post token=xoxb-2291-8812-Kq3vZ [REDACTED] channel=#approvals
payment.retry card=4111 1111 1111 1111 [REDACTED-PAN] amount=420.00
git.fetch Authorization: Bearer ghp_9fK2mQx7Lw0 [REDACTED]
ssh.exec key=-----BEGIN OPENSSH PRIVATE KEY----- [REDACTED]
```

Illustrative log lines. Each secret is removed before anything is written.

- ✓ No incident data sent to us
- ✓ Models on your GPU, or an endpoint you choose
- ✓ Card numbers caught by checksum

Does any of our data leave the network?	No. It never copies your data anywhere: each question runs where the data already lives, and only a small answer comes back. It reaches out only to the tools you connect, to a model endpoint if you use one instead of a GPU, and for a signed licence check, which you can turn off when air-gapped.
Which AI models does it use?	Open-weight models. With a GPU they run inside your install, and nothing calls out. Without one, you point 3AM at a model endpoint: your own, inside your network, or a hosted one. It won't run them on a CPU, where every investigation would wait.
What can your company actually see?	The licence service sees fingerprints, public keys and licences. It never sees a connection, a secret or a person. Air-gapped, it sees nothing at all.

Then, who can change anything? Two people, always.

Seeing is one thing. Changing is another. Every write takes two people, each with their own key. Code review can't grant that, and neither can an admin. Only their signatures can.

- ✓ Shadow by default, per-service autonomy
- ✓ Quorums, allowed hours, action budgets
- ✓ Mandatory dry run, one global halt

So locking takes a second. Unlocking takes days.

An attacker's first move is to loosen the rules, so that's the move that waits. Tightening them never does: any one person can do it, the moment something looks wrong.

Making it safer

Now

Any one person, immediately.

- ✓ Revoke a key
- ✓ Halt everything
- ✓ Lower what it may do

Making it riskier

Days

Several signatures, then a wait. Never the person who benefits.

- ✓ Let it do more
- ✓ Add an approver
- ✓ Upgrade the safety core

While it waits, everyone is alerted, and any one of them can cancel it. Your clock times it, not ours.

Suggested defaults. You set your own, above fixed floors.

Who holds the root of trust?	You do. At install, named people from your bank each enrol their own key and sign the first policy: who may approve, which tables are the books, and how many signatures each change takes. 3AM staff are never among them.

Not the books doesn't mean harmless

A failover on a payments database can stop a bank. So every action is rated by what it could do if hijacked, and the rating decides who must approve it. You set the ratings for your own estate.

The action	For example	Who must approve	Can it ever run alone?
Contained	Release a lock, clear a stuck job	Two people	Only once it has earned it
Data, not the books	Set-based fixes to tables you opt in	Two people	No
Disruptive	Failover, rollback, restarting a payment service	Two people and the business owner	Never
The books	Reversals and corrections	3AM can't run it. Your DBA does, with their own key	Never

Disruptive actions also stay out of your declared business-critical windows, and one person can stop everything at any moment.

Can it ever touch our books?

No, at any setting. 3AM holds no key to your ledgers, journals or balances. For those it hands your DBA a signed package, already rehearsed on a copy, and a person runs it with their own key. 3AM checks the result afterwards.

Documents can suggest. Only evidence can act.

A log line, a ticket or a row can hide an instruction, and it might change what the AI suggests. That's survivable, because no suggestion can pass a gate. Only a live check can authorise a fix.

Source	Tier	Can suggest	Can authorise
Code and schema	T1	✓	✓
Alert rules and live metrics	T2	✓	✓
Runbooks and wiki pages	T3	✓	—
Past postmortems	T4	✓	—

And whatever happened, you can prove it.

Every signal, check, approval and action is hash-chained, and your own key co-signs each seal, so nobody can quietly rewrite it. Not even us. Change one byte and the verifier shows where.

- ✓ Append-only, hash-chained, Ed25519-signed
- ✓ Independent verifier for edits, gaps, reorders
- ✓ Syslog to Splunk, QRadar or any SIEM

How do our auditors verify what happened?	Every event is written to an append-only, hash-chained ledger with signed segments. An independent verifier detects any edit, deletion, insertion or reordering, and the ledger exports to your SIEM.
Can we audit exactly what the AI was asked?	Yes. Every model call goes on the ledger: the model, the redacted prompt with its SHA-256, the output, tokens and latency. Whoever holds the original prompt can prove which one it was.

Now assume all of that fails. Including us.

Keys get stolen and vendors get hacked, and we're a vendor. So take each place an attacker could land, and ask what they'd get. None of them is enough to change your data or loosen the rules.

If an attacker lands in	They try to	What happens
Your Git or CI pipeline	Merge a malicious change	It's only a proposal until your approvers sign it.
The cluster 3AM runs in	Forge an approval	The signature fails, and opens an incident against 3AM itself.
Your identity provider	Add a new key	It means nothing until your existing approvers enrol it.
One of your sites	Reach beyond that site	It can't go beyond a plan your people have already approved.
3AM, the company	Push an upgrade	Upgrades need your signatures, and your people hold the root keys. Never our staff.

Your data and your rules stay as they were.

Built to be checked.

- ✓ Builds you can verify yourself
- ✓ SBOM with every release
- ✓ Source-free compiled build
- ✓ No shell access for the agent
- ✓ Signed, distroless images
- ✓ Signed licence, verified offline
- ✓ Read-only diagnostic credentials

Does it work with our single sign-on?	Yes, on every licence tier. It signs in through OIDC, including Entra ID, ADFS, Okta, Ping and Keycloak, or directly against LDAP or Active Directory over TLS. Your provider's MFA and access rules apply, and directory groups decide each person's role.
Who in our organisation can see what?	There are four roles: owner, admin, responder and viewer. A viewer can read incidents and the audit log, and can see connections but never their secrets. Someone in no mapped directory group gets no access.
How is the install itself locked down?	It runs as non-root with a read-only root filesystem, all Linux capabilities dropped, seccomp RuntimeDefault and no service-account token. An in-cluster model is reachable only from the core, by NetworkPolicy.
Could a log line or a ticket trick it into acting?	Text inside anything it reads is treated as data, never as instructions. A fix still needs a cause confirmed by a live check, and your approval policy still applies.

Have a vendor security questionnaire? Bring it to the demo. We'll walk your team through it line by line.

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

Write to us
hello@3am.si

Company
3AM Super Intelligence Inc.



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.

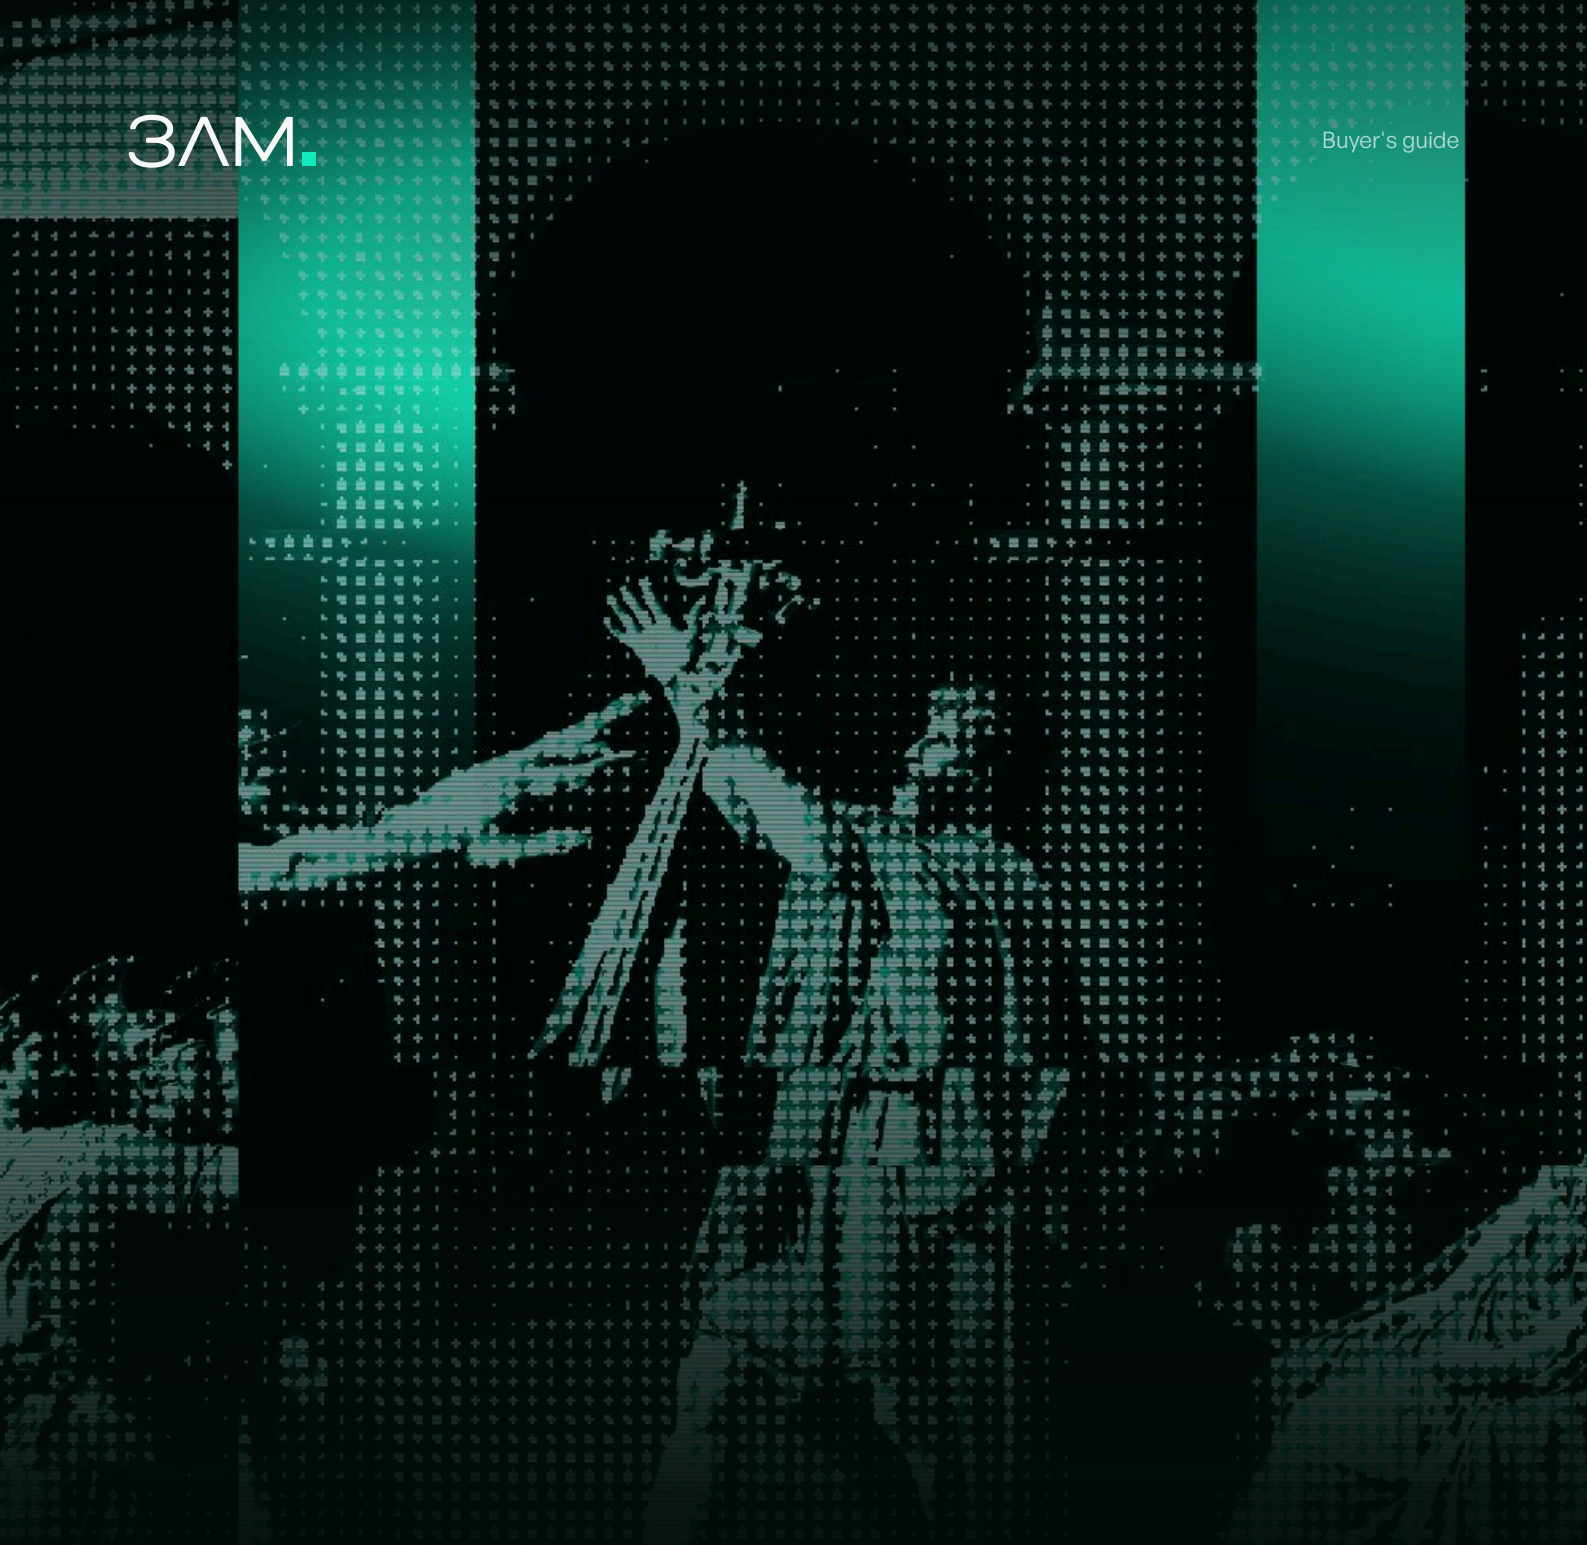
See it on your estate
www.3am.si/demo

Write to us
hello@3am.si

Company
3AM Super Intelligence Inc.

3AM.

Buyer's guide



Eight questions. Ask every vendor.

What a bank should ask any AI on-call before it goes near production, and how the field answers from its own public pages.

Document	For	Published
Choosing an AI on-call for a bank	Buying committees	November 2026

A bank is not a start-up.

An AI on-call for a bank will be judged by your security team, your risk function and your regulator before it ever helps an engineer. Most products were built for a different buyer.

$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

How to use this guide

1 **Ask all eight, of every vendor.** Each one is a question a bank's own reviewers will ask first.

2 **Read the public documentation, not the pitch.** If an answer is not written down, treat it as not there yet.

3 **Mark each answer** as documented, in part, or not found. Pages 5 and 6 do this for the field.

4 **Then run a pilot on your own incidents,** in shadow, before anything can change. Page 7 sets one out.

The eight questions, 1 to 4.

1 Runs entirely inside your network

Every part, management included, on infrastructure you own

Why it matters

Anything managed from outside is a path in from outside. Your examiners will ask who can reach the control plane.

A good answer

Every part runs in your network. The only call to us is a signed licence check; air-gapped, it verifies offline.

Ask to see

A network diagram of every connection: which side opens it, and what crosses it.

2 Runs air-gapped

No connection out at all

Why it matters

Some of the systems that matter most sit on networks with no route out.

A good answer

An air-gap bundle ships the images and models together. The licence verifies offline.

Ask to see

The air-gap install guide, and a demo with no connection out.

3 Never copies your telemetry

No ingest, no index, no cached results outside your systems

Why it matters

Every copy of telemetry is one more dataset to secure, classify, retain and explain to a regulator.

A good answer

Each question runs where the data lives and a small answer comes back. Nothing is copied out.

Ask to see

A list of everything the product stores, where it is kept and for how long.

4 Proves the cause before acting

A live check that confirms it, not a confidence score

Why it matters

A confidence score is a guess with a number on it. A change to production needs a cause that was checked.

A good answer

A live check confirms the cause. Only a confirmed cause can authorise a fix.

Ask to see

One past incident walked through: what was checked, on which systems, and what would have failed it.

The eight questions, 5 to 8.

5 Applies the fix and verifies it

Rehearsed first, then checked after it runs

Why it matters

A diagnosis alone leaves the riskiest part, the change itself, to tired people at 3 a.m.

A good answer

Every fix is dry-run first, applied with approval, then checked against the live system.

Ask to see

The rehearsal output for one fix, and the check that ran after it.

6 Two people approve every write

Not one person, and never switched off by an admin

Why it matters

Four eyes is the rule for every other change in a bank.

A good answer

Two different people approve one plan, and the plan cannot change after they sign.

Ask to see

In the demo, try to approve your own change, and to edit a plan after it was approved.

7 Holds no key to the books

A hard line around ledgers, journals and balances

Why it matters

A fix that turns an alert green by editing a balance is the worst outcome there is.

A good answer

No key to the books at any setting. A person runs a signed package with their own key.

Ask to see

Where the line is enforced, and what happens when someone tries to move it.

8 Tamper-evident audit trail

Hash-chained and signed, verifiable by your auditors

Why it matters

Months later, someone will ask exactly what happened and who approved it.

A good answer

Every step is hash-chained and signed, and exports to your SIEM.

Ask to see

An exported trail, and a verifier your own team runs, without the vendor.

How the field answers.

Nine products against the eight questions, every answer from the vendor's own public pages, read 11 November 2026.

Product	Inside	Air-gap	No copy	Proof	Fix	Two people	Books	Audit	Documented
3AM	✓	✓	✓	✓	✓	✓	✓	✓	**8 of 8**
Traversal Traversal AI SRE	◐	—	◐	◐	◐	—	—	—	0 of 8 4 in part
Resolve AI Resolve AI	—	—	◐	◐	◐	—	—	—	0 of 8 3 in part
incident.io AI SRE	—	—	◐	—	◐	—	—	◐	0 of 8 3 in part
Rootly Rootly AI SRE	—	—	◐	—	◐	—	—	—	0 of 8 2 in part
PagerDuty Paige SRE agent	—	—	—	—	◐	—	—	—	0 of 8 1 in part
Datadog Bits AI SRE	—	—	—	◐	◐	—	—	◐	0 of 8 3 in part
ServiceNow Now Assist for ITOM	◐	—	—	—	◐	—	—	◐	0 of 8 3 in part
NeuBird Hawkeye and Falcon	◐	◐	◐	—	◐	—	—	—	0 of 8 4 in part
Cleric Cleric	—	—	◐	◐	◐	—	—	◐	0 of 8 4 in part

✓ Documented ◐ In part — Not found in public docs

Where each one is strong.

Every product here is good at something. The question is whether it is good at what a bank must have.

Traversal AI SRE	Scale. Traversal indexes very large estates and counts large financial institutions among its customers. Traversal hands you a diagnosis. We prove the cause, then fix it.
Resolve AI AI SRE	Breadth. Resolve works across many tools and has well-known fintech and exchange customers. Resolve sends your incident to its cloud. Ours never leaves the bank.
incident.io Incident platform	Workflow. incident.io runs the human side of an incident very well: paging, channels, status pages. incident.io gives you a hypothesis. We give you a proven cause and a verified fix.
Rootly Incident platform	Speed to start. Rootly is quick to adopt and popular with fast-moving engineering teams. Rootly calls its answer a hypothesis. Ours is proven before it acts.
PagerDuty Incident platform	Reach. PagerDuty is the paging standard across large enterprises, banks included. PagerDuty's agent is 91% sure. Ours shows you the proof.
Datadog Observability	Coverage. If everything already lives in Datadog, Bits sees all of it, with no setup. Bits needs your data in Datadog. We go to the data where it already is.
ServiceNow ITSM	The record. ServiceNow is where most banks keep changes, incidents and the CMDB. Now Assist suggests a probable cause. We prove the actual one.
NeuBird AI SRE	Flexibility. NeuBird lets you bring your own model and deploy into your own cloud. NeuBird scores its answers. We prove ours.
Cleric AI SRE	Change safety. Cleric watches every deploy and catches regressions early. Cleric checks after the fix. We prove before it.

Running a fair pilot.

Put every shortlisted product on the same incidents, under the same rules, and score what your reviewers will score.

1 **Bring your own incidents.** Three past ones nobody could fully explain, and live traffic as it comes.

2 **Start in shadow.** Nothing changes. Compare each proposal with what your team actually did.

3 **Run the security review alongside,** with the vendor's documents and your own questionnaire.

4 **Decide on the record,** not on the demo: what was proven, what was guessed, and what it would have changed.

What to score

Measure	What to look for
Proven, or guessed	For each cause it named: was it checked on your systems, or given a score?
False alarms	A stated rate, measured on your estate, not a figure borrowed from someone else's.
Time to a proven cause	From the first alert to a cause your engineers accept.
The fix it would have run	Rehearsed and exact, or a suggestion for someone to work out.
Who would have approved	How many people, and whether any setting could remove one.
What your auditors can replay	Every step, verifiable without the vendor.

Put us side by side on your own incidents.

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

Write to us
hello@3am.si

Company
3AM Super Intelligence Inc.

Sources.

Read 11 November 2026. Every vendor answer in this guide comes from these public pages.

Traversal	Traversal, homepage ↗ Traversal, Security ↗ Traversal, What data leaves your environment (10 Aug 2026) ↗ Traversal, docs ↗
Resolve AI	Resolve AI, homepage ↗ Resolve AI docs, Security and administration ↗ Resolve AI docs, Satellite ↗ Resolve AI docs, Mitigation actions ↗
incident.io	incident.io, AI SRE ↗ incident.io docs, Security FAQs ↗ incident.io docs, AI usage ↗
Rootly	Rootly docs, AI SRE overview ↗ Rootly docs, Data privacy for Rootly AI ↗ Rootly docs, Investigation rules ↗ Rootly docs, Private Agent ↗
PagerDuty	PagerDuty docs, SRE agent ↗ PagerDuty docs, Service regions ↗ PagerDuty, Fall 2026 launch ↗
Datadog	Datadog docs, Bits AI SRE ↗ Datadog docs, Investigate issues ↗ Datadog docs, Bits remediation ↗
ServiceNow	ServiceNow Community, Now Assist for ITOM FAQ ↗ ServiceNow Community, Self-hosted AI platform (30 Apr 2026) ↗ ServiceNow Community, AI agents governance guide ↗
NeuBird	NeuBird docs, Security ↗ NeuBird docs, How it works ↗ NeuBird, Earned write access ↗
Cleric	Cleric docs, Security and data privacy ↗ Cleric, Technology ↗

Compared on 11 November 2026 from each vendor's public website and documentation, linked on every page. Products change; if something here is out of date, write to hello@3am.si and we will correct it. All product names and trademarks belong to their owners and are used only to identify them. No affiliation or endorsement is implied.

Figures on page 2: Splunk, The Hidden Costs of Downtime 2026; FCA, 20 December 2022; EIOPA, Digital Operational Resilience Act.