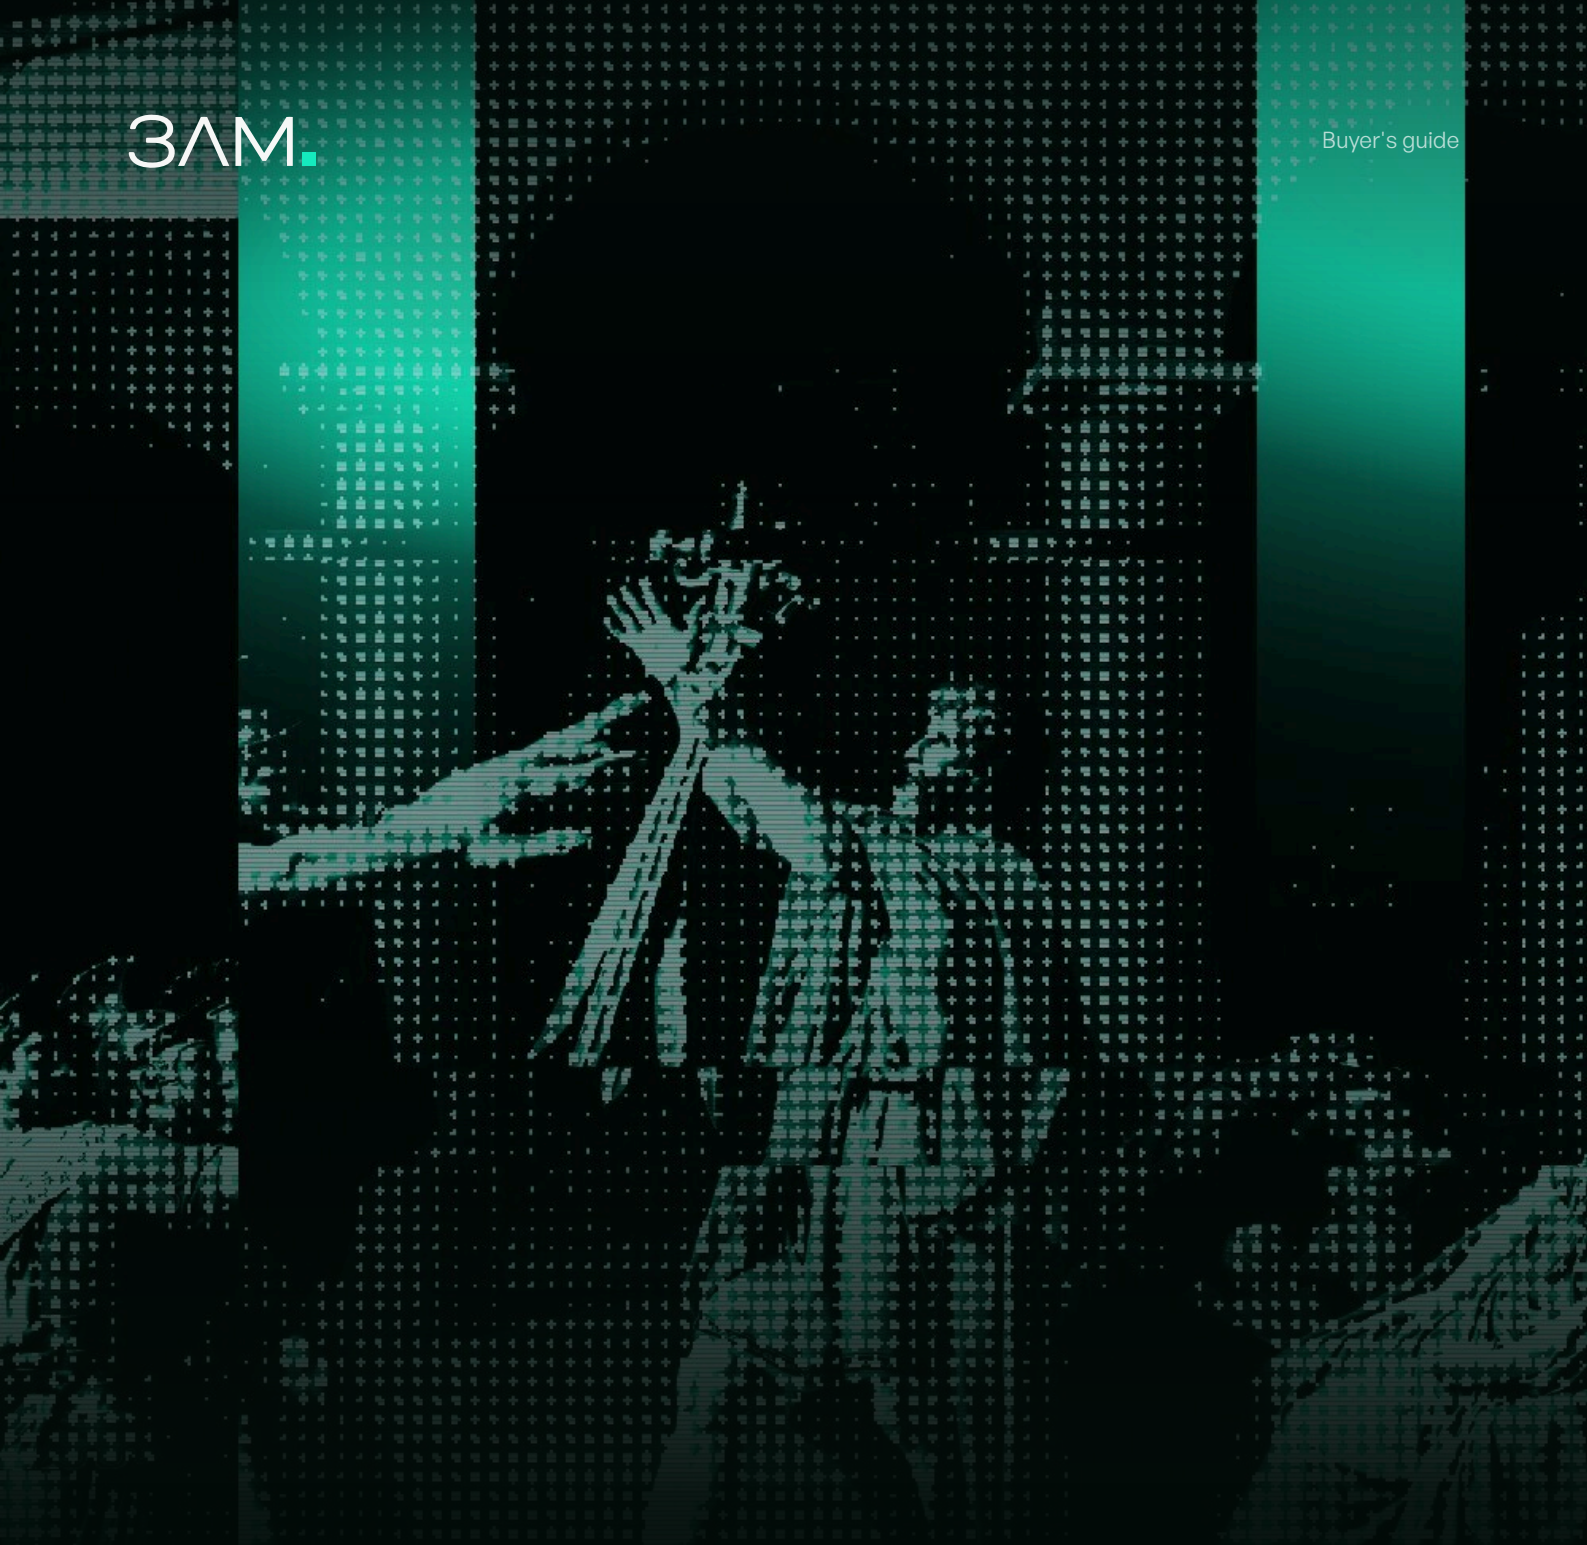
3AM.

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.