CYBERSECURITY 101 · eBuilder Security

AI Anomaly Detection

Signature-based tools only block attacks they already recognise. AI anomaly detection learns what normal looks like across your systems and flags everything that does not fit, catching threats no one has described yet. Here is how it works and how to run it well.

WHAT IT IS

AI anomaly detection uses machine learning to spot activity that deviates from your normal behaviour. Instead of matching known attack signatures, it learns what normal looks like across your network, users and devices, then flags the unusual for a person to investigate. That lets it catch threats no signature has ever described, the category that tends to do the most damage.

WHY IT MATTERS

2.22 million dollars

The average breach-cost saving for organisations using security AI and automation extensively, the single biggest factor in the study.

IBM, 2024

About 258 days

The average time to identify and contain a breach, the window that fast anomaly detection is meant to shrink.

IBM, 2024

10 days

The global median time attackers stayed hidden before discovery in 2023, a gap early detection is built to close.

Mandiant M-Trends, 2024

WHAT HAPPENED

In 2013 attackers reached US retailer Target through a heating and ventilation contractor and installed malware on its checkout systems. The monitoring tools worked and raised alerts, but no one acted on them, and about 40 million payment card records and 70 million customer records were exposed. The lesson is blunt. Detection without a response is just noise. (US Senate Commerce Committee, 2014)

RED FLAGS TO WATCH FOR

Anomaly detection surfaces the unusual. These are the patterns that most often mean a real problem, and that no one should wave through.

- A valid login used at an odd hour or from a location the user never works from.
- A server reaching out to a country or system it has never contacted before.
- A user or account suddenly touching data far outside its normal job.
- A sharp spike in data leaving the network, especially outside business hours.

THE ONE RULE THAT STOPS IT

Treat anomaly detection as people, process and technology together, never a product you switch on and forget. It is only as good as the baseline it learns and the response it triggers.

THEN BUILD THE HABIT AROUND IT

- Let the model learn from weeks of your real traffic before you trust its alerts.
- Feed it network, identity, endpoint and log data so it sees the whole pattern.
- Tune relentlessly to cut false positives, because noise is how a real alert gets missed.
- Route every alert to someone who can act on it, at any hour.
- Retrain after any big change so a shifting baseline does not quietly erode your coverage.

MYTHS AND FACTS

MYTH

A firewall and antivirus already cover us.

FACT

Those match known signatures. Anomaly detection catches the new and unknown attacks they were never built to see.

MYTH

You just switch anomaly detection on.

FACT

It has to learn your environment first and be tuned continuously. An untuned model floods you with false positives or misses the real thing.

MYTH

Anomaly detection stops attacks on its own.

FACT

It produces signals, not verdicts. A person still has to investigate and respond, which is why the detection needs a team behind it.

THE LEGAL DUTY

Under NIS2, in force in Sweden as Cybersäkerhetslagen since 15 January 2026, continuous monitoring is a legal expectation and boards can be held personally accountable under Article 20.

GET HELP

eBuilder Security runs Sweden-based managed detection and response that pairs anomaly detection with a team to answer the alert. Read the full guide at the link below.

Read the Full Article →