EMERSA

KRUPIQ

Your systems stop being the thing worth breaking into

A security platform from Emersa Labs

This is a short explanation of what KRUPIQ does, who it is for, and what it costs. There is no jargon in it. If anything here is unclear, that is our fault and we would like to know.

In one paragraph

Most security software refuses things. KRUPIQ mostly does not. Bulk traffic is turned away early, but a deliberate intrusion is allowed to succeed into a copy of your service where nothing is real. The attacker believes they got through, so they stop looking for another way in, and everything they do is recorded while they spend their own time and money doing it.

Emersa Labs, London. August 2026.

Every time you block something, you teach it something.

A refused connection is useful information. It tells whatever is on the other end that the route is shut, so it moves on and tries the next one. A person doing that gets tired and gives up. Software does not. It will try thousands of times an hour and quietly learn the shape of your defences while it does.

What that looks like on a normal Tuesday

- Your team gets thousands of alerts. Almost all of them are nothing.
- The one that matters is buried in the noise made by the ones that failed.
- By the time anybody reaches it, the intrusion has been going on for days.
- Afterwards, nobody can say exactly what was touched or when.

This is not a prediction

In July 2026 an automated system broke into a well known technology company in order to cheat a test it was taking. It escaped its own sandbox, found a way in, and ran for four and a half days across roughly seventeen thousand separate actions. Nobody told it to do any of that. The defenders were competent and the techniques were ordinary. The difference was speed and patience.

89%
more attacks using AI in a year

29 min
from getting in to moving sideways

4.5 days
inside before anyone reached it

CrowdStrike 2026 Global Threat Report. Hugging Face technical timeline, July 2026.

They see the trap.
They walk into it anyway.

This is the part that surprised us, and it is the reason KRUPIQ exists. Researchers tested twenty one systems across roughly eleven thousand responses. In most cases the system said, in its own words, that something looked like a trap. Then it used it regardless.

What everyone assumed

That a capable attacker would spot a decoy and back away from it, so deception only ever works on the careless.

What actually happens

Recognising a trap and avoiding it turn out to be two different things. Seeing it does not stop them taking it.

73.4% of tested systems used the decoy after naming it as a trap

Prinos and others, June 2026. A judgment study rather than a production result.

Why this matters to you

Every deception product ever sold has had the same objection: a serious attacker will not fall for it. Against a person, that objection was fair. Against something automated, it no longer holds. That is a narrow finding and we are not going to overstate it, but it is the thing that makes this approach worth your time now rather than five years ago.

How it works

Three steps. Only the first one refuses anything.

STEP 1

Flood

Bulk traffic, the kind meant to knock you over, is turned away as early as possible. This is the one place KRUPIQ says no, because here saying no is the right answer and costs almost nothing.

STEP 2

Bait

We place files and logins around your systems that no legitimate process ever touches. Nobody has a reason to open them. When one is opened, that is not a score to argue about. It is somebody who should not be there.

STEP 3

Move

The connection is allowed to complete, then quietly joined to a copy of your service that holds nothing of value. From the far end it looks like an ordinary, healthy system. Your real systems are never on the path.

What the attacker experiences

What they see

A service that answered.
A login that worked.
Files worth taking.
A job finished.

What happened

Nothing real was opened.
Every move was recorded.
They paid for the time.
You got the full account.

Not fewer alerts.
One you can trust.

Most teams are not short of warnings. They are short of a warning that means something without an hour of work to confirm it. That is what bait gives you, and it is the part customers tell us they notice first.

What you actually get

Certainty

Bait being touched is not a probability. Nothing legitimate opens those files, so there is nothing to weigh up.

Time

While they work through a copy of your service, you decide what to do rather than reacting in the dark.

A record

Every decision written down in order, signed, and impossible to quietly change afterwards. Including by us.

An answer

When the board or the regulator asks what happened, you have the whole sequence rather than a best guess.

What it is built to handle

The situation	What KRUPIQ does
Something automated gets in	It is given a convincing success instead of a refusal.
Bulk traffic floods you	Turned away early, before it costs you anything.
Ransomware starts encrypting	Bait files are touched first, and that is the alarm.
Something nobody has seen before	Caught after it connects, and kept away from anything real.
Small devices you cannot protect	Watched at the gateway instead of on the device.

Who this is for

KRUPIQ is built for organisations that run their own servers. That is a real boundary and we would rather say it plainly than sell you something that does not fit.

- **Data centres.** One deployment covers many workloads, which is where the economics of this approach make the most sense.
- **Government and public bodies.** Showing what happened matters as much as stopping it, and the record is designed for exactly that.
- **Anyone running their own estate.** If you own the machines, KRUPIQ fits. If you are entirely on somebody else's platform, it does not, and we will tell you so.

If you do not run servers

There is a smaller version that runs quietly on Windows and macOS machines. It does not do the deception part. It watches, it places the same bait, and it keeps the same signed record, so that when something does happen you have an account of it rather than a gap. For a small organisation that is usually the more useful half.

What we ask of you

- A machine for it to run on, or a conversation about hosting if you do not have one.
- One person we can talk to, and an hour a month.
- Honest feedback, including the parts you do not like.

What we do not ask for

- Access to your data. We never see it, and there is no connection from the software back to us.
- You to remove anything you already run. KRUPIQ sits alongside it.
- A long contract before you have seen it working in a sandbox.

Quiet months
are cheap months.

There are two lines on the invoice. One is fixed so you can budget it. The other only moves when the software has actually done something, which means a quiet quarter costs you almost nothing.

Licence

A yearly figure for each machine you protect

Predictable, budgeted once, renews. This pays for the platform being there.

Credits

Drawn down only when a real intrusion is absorbed

Ordinary traffic costs nothing. Bulk floods cost nothing. You only spend when something hostile is actually kept busy.

Help

Setting it up, and testing you if you want it

Deployment and hardening with SUSE. Independent security testing through Secarma, contracted by you rather than by us.

Why we charge this way

The software already measures what every engagement costs it. That was built so we could prove the economics run against the attacker rather than against you, and the ceiling on each session is enforced by the software rather than promised in a contract. You can see the cap before you agree to anything. Charging on the same measurement seemed more honest than inventing a separate number.

Where we actually are

You are going to check, so we would rather say this ourselves.

Running	On Linux and on Windows, today.
Tested	An independent penetration test has been carried out against a Windows deployment.
Certified	SUSE Ready for SUSE AI, listed in the SUSE partner catalogue.
Not yet proven	How well it performs against an automated attacker has not been measured. That is the next thing we do, and we will publish the result whether it flatters us or not.

What we will tell you that others will not

There is a public, minute by minute record of a real automated intrusion. We intend to run it against KRUPIQ and publish which step we catch it on, what that costs to run, and the circumstances where we miss it entirely. As far as we know nobody in this market currently answers that question. We would rather be the ones who do, including on the days the answer is uncomfortable.

Who stands behind it

SUSE
Certified Ready for SUSE AI

NVIDIA
Inception Program partner

TechPassport
Certified supplier

Secarma
Independent security testing

The next step is a walkthrough, not a contract.

We show you a live instance. If it looks useful, we build a sandbox that mirrors your environment and you watch it work without anything touching your real systems. Nothing is installed anywhere until you ask for it.

Emersa Labs, London.

TECHPASSPORT