KAIA-OR · WHITEPAPER · EVENT INTELLIGENCE

From volume to relevance

Relevance-first security intelligence: how KAIA-OR decides which hazards, incidents, disruptions and cyber events deserve your attention, how sure it is, and why each one matched.

DOCUMENT ID	WP-1
VERSION	1.0
DATE	3 November 2026
AUDIENCE	Security, risk, travel and operations leaders in enterprises and government; teams without a security function; technical evaluators; informed individuals and families

Contents

1. Executive summary 3

2. From volume to relevance 4
   1. The overload problem 4
   2. What relevance means 4
   3. How an event reaches you 5

3. Corroboration, confidence and the reason why 6
   1. What each event tells you 6
   2. Corroboration is not confidence 6
   3. Showing why, and what changed 7

4. Relevance and privacy together 8
   1. Matching on the device 8
   2. From personal alerts to a shared picture 8

5. AI under evidence 10
   1. Understanding events 10
   2. Answers checked against the evidence 10
   3. Authority stays with people 11

6. Honest scope and a buyer’s checklist 12

7. What this means for you 13
   1. Glossary 13

© Kaia Technologies, Inc.

Executive summary

Most people and organizations do not lack information about the events that threaten them. They lack a way to tell which events concern them, how sure anyone is, and what to do next. Alerts arrive from many channels, in many formats and with little context. The outcome is familiar: generic alarms that people learn to ignore, security teams that spend their day sorting reports instead of making decisions, and the one event that mattered noticed too late.

KAIA-OR is built on a different premise, which we call **relevance-first security intelligence**: one explainable picture of physical, cyber and travel risk, matched to what each person and organization is responsible for. What you tell KAIA-OR matters to you (your places, trips, sites, software and suppliers) decides what you see. Natural hazards, security incidents, infrastructure disruptions and cyber events come in from official and authoritative sources and are corroborated across independent sources. Each event is categorized and described by its severity, place, time window and confidence, together with its corroboration status. It is then matched to you, on your device, and shown with the reason it matched.

Four commitments run through the design:

- **Relevance.** You see the events that touch your own places, trips, sites, software and suppliers, not an unfiltered stream.
- **Explanation.** Each event shows why it matched and how well it is corroborated, and each travel advisory you follow shows what changed in each new revision.
- **Evidence before answers.** The assistant, Hey KAIA, answers by voice or text. Its answers are checked against the evidence and cite it, and model output grants no authority.
- **People decide.** KAIA-OR suggests a next step. It does not act on anyone's behalf without their authorization.

This paper explains each commitment at a conceptual level. It describes principles, the conceptual architecture and outcomes. It is not an implementation guide: it does not name sources, and it does not publish internal mechanics or the underlying data model. It ends with a checklist that any buyer can use to evaluate relevance, whichever vendor they are considering.

> ABOUT THIS PAPER
>
> This paper describes the design and the committed product scope of KAIA-OR. Capability availability depends on the edition and deployment. Statistics are quoted as market context, with their publisher and date under Sources, and are never a measure of any product's effect.

© Kaia Technologies, Inc.

1. From volume to relevance

The overload problem

Security information has never been more plentiful. Public bodies, operators, researchers and the press publish a great many reports, and any organization of size receives bulletins, travel notices and internal reports on top. The difficulty is no longer access. It is attention. Three patterns recur, whether the reader is one person, a family, a small team or a large organization.

- **Generic alarms.** An alert about a region, a sector or a technology says nothing about whether it touches a particular office, a particular trip or a particular product. When most alerts do not apply, people stop reading them, including the ones that do.
- **Fragmentation.** Physical, travel and cyber information live in different tools, owned by different people. Nobody connects a regional flood, a scam wave that exploits it and a colleague's trip into the affected city.
- **Uncertainty.** Early reports are often partial, and some are wrong. Without a clear statement of how well a report is corroborated, decision-makers either act on rumour or wait for a certainty that arrives too late.

Security teams end up sorting and removing duplicates instead of exercising judgement. Teams without a security function, individuals and families simply find out late, from the people they serve, from relatives or from the news.

The breadth of the risks involved is well documented. In a 2026 survey of risk experts, cyber incidents were the top global risk for the fifth year in a row, at 42% of responses; business interruption ranked third, and natural catastrophes fifth, at 21%.[1] Physical and cyber risk sit side by side on the same list, yet they are rarely presented in one place.

What relevance means

In KAIA-OR, relevance is defined by the person or the organization, not by the product. You tell KAIA-OR what matters, and that decides what you see.

WHAT YOU TELL KAIA-OR, AND WHAT MAKES AN EVENT RELEVANT TO IT

You add	For example	An event is relevant when
Places	Home, work, a town where relatives live	Its place and time window touch a place you saved
Trips	A destination and travel dates	It affects your destination or route during your dates, or the advice for that destination changes
Sites	An office, a shop, a warehouse, a plant, a post	It affects a site your organization is responsible for
Software	The products and services your organization relies on	It concerns a major breach or a newly exploited vulnerability in software you use
Suppliers	The organizations you depend on	A supplier is exposed to a critical event or a newly exploited vulnerability

© Kaia Technologies, Inc.

Events that touch none of these do not compete for your attention. Relevance-first does not mean less information. It means that what you see is chosen by criteria you set, with the reason shown. The same approach serves a person and an organization; chapter 3 explains how the two coexist.

How an event reaches you

Conceptually, each event follows the same six steps, from what you care about to the decision you make.

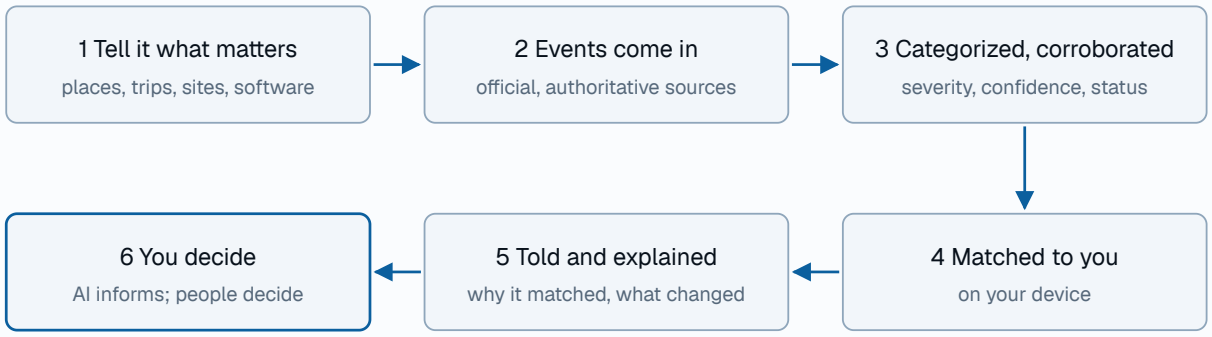


Figure 1. The six conceptual steps. Personal items stay on the person's device; an organization's sites go into its shared picture. The figure is conceptual and does not describe an implementation.

© Kaia Technologies, Inc.

2. Corroboration, confidence and the reason why

What each event tells you

An event in KAIA-OR is more than a headline. Each one carries a small set of attributes, presented in plain language, so that a reader can judge it without opening a separate report.

THE ATTRIBUTES SHOWN WITH EACH EVENT

Attribute	The question it answers
Category	What kind of event is this: a natural hazard, a security incident, an infrastructure disruption or a cyber event?
Severity	How serious is it, described in the same plain terms across categories?
Place	Where is it, or which area does it affect?
Time window	When did it begin, and over what period does it matter?
Confidence	How sure is KAIA-OR about the event as described?
Source	Where does the report come from?
Corroboration status	Do independent sources report the same event? Reports without independent corroboration are labelled as unconfirmed.
Why it matched	Which of your places, trips, sites, software or suppliers does it touch?

Corroboration is not confidence

- **Corroboration** asks whether independent sources report the same event. It is a statement about agreement between sources.
- **Confidence** asks how sure KAIA-OR is about what is described: the place, the timing and the severity, given what is known at that moment.

The two can diverge. A clearly described report from one authoritative source may be precise but awaiting a second source. Several brief reports may agree that something happened while leaving its extent unclear. Showing both lets a reader tell “probably happened, details unclear” apart from “clearly described, corroboration pending”, and brief leadership accordingly.

Early reports matter, so KAIA-OR does not hold them back until its details are settled. It shows them with their label, and the label changes as corroboration arrives or fails to arrive. Corroboration is evidence, not proof, and KAIA-OR presents it that way: an event is described as corroborated or unconfirmed, never as established fact.

> WHAT WE PUBLISH, AND WHAT WE DO NOT
>
> We publish the meaning of each attribute, not the internal mechanics that produce it. Publishing those would help anyone who wants to craft a report that looks corroborated, and would not help a reader judge a particular event.

© Kaia Technologies, Inc.

Showing why, and what changed

Each event in your view carries the reason it matched: the saved place it is near, the trip whose destination and dates it overlaps, the site it affects, the software it concerns or the supplier it involves. The explanation serves three purposes.

- **Trust.** You can check the reasoning instead of taking a verdict on faith.
- **Correction.** If a match is wrong because a place, trip or site is out of date, you can correct the input, rather than learning to ignore the product.
- **Briefing.** A security lead can pass an event to leadership with its reason, severity, confidence and corroboration status attached.

Much security information is not a new event but a change. The advice for a destination is revised. An incident escalates or subsides. A vulnerability that was theoretical is reported as exploited. For a travel advisory you follow, each new revision shows its previous and new publication dates and the fields that changed, so that nobody has to reread a long advisory to find the line that moved. Explanations appear in English, French (Canada), Turkish or Hebrew, with full right-to-left layout.

© Kaia Technologies, Inc.

3. Relevance and privacy together

Matching on the device

For personal use, the comparison between incoming events and a person's saved places and trips is made on that person's own device. The places and trips stay on the device. Exact location stays on the device too. When a person asks what is happening nearby, KAIA-OR uses one fresh position for that check and clears it when they stop. Notifications never name the place or the event, so a locked screen does not reveal where someone is or what they care about.

This is a design choice, and it is where relevance and privacy reinforce each other. A service that matches events centrally has to know where its users live, work and travel, and keep that knowledge up to date. A service that matches on the device does not. The more precisely KAIA-OR matches for a person, the less it needs to know about that person centrally. Relevance and privacy pull in the same direction instead of being traded against each other.

The same principle applies to the assistant. Hey KAIA works on the device first. Nothing is sent to a cloud AI service unless the person agrees first, and government organizations start with cloud AI off.

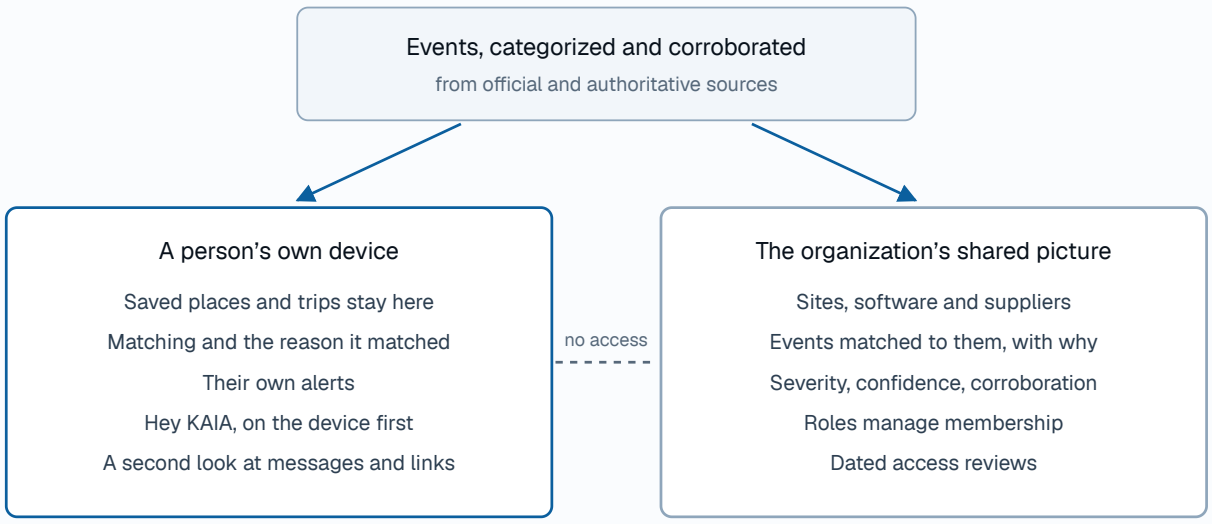


Figure 2. Two plain boxes: the device and the organization. Membership is not access. The organization sees its sites, not its people's whereabouts.

From personal alerts to a shared picture

KAIA-OR is one product from one person to a government agency. The Individual, Family, Team, Enterprise and Government packages share the same event intelligence and the same explanations. In an organization, two pictures coexist.

- **The shared picture** holds the organization's sites, software and suppliers, and the events that touch them, for the people responsible. An event that touches a shop, an office or a warehouse appears on the site it affects, with its severity, its corroboration status and why it matched. Events across a border or

© Kaia Technologies, Inc.

offshore that affect a site appear there too.

- **Each member's own KAIA-OR** holds that person's own places and trips, their own alerts and their own second look at messages and links. Each traveller sees what changed in the official travel advisory for their own trip.

Roles manage membership, not people. The organization sees its sites, not its people's whereabouts. Membership never grants access to another member's messages, devices, location or saved places. Access is accounted for: invitations are one-time and bound to an account, and access reviews are dated.

Across your supply chain. Suppliers exposed to critical events or to newly exploited vulnerabilities appear in the same picture as your sites, so that a physical event at a supplier and a cyber event in shared software are seen together.

Into your own tools. Events can flow into the organization's existing tools through a documented API and webhooks. Relevance-first intelligence then reaches the workflows a team already uses, instead of adding another screen.

© Kaia Technologies, Inc.

4. AI under evidence

KAIA-OR uses artificial intelligence in three places: to understand events, to explain them, and to answer questions through the assistant. In each, one rule applies: AI informs, and people decide. This chapter uses a few technical category terms, each paired with what it means for the reader.

Understanding events

Events are described in many vocabularies and languages. KAIA-OR maps them into one shared vocabulary: an **ontology** of hazards, incidents, disruptions, cyber events, places, organizations, software and suppliers, held as a **knowledge graph**. The outcome for the reader is consistency. The same kind of event is recognized whatever words a source uses, so a flood is a flood and a supplier is a supplier wherever the report came from.

Semantic reasoning works out what an event is and what it touches. Reasoning about how situations relate, unfold and change over time connects an outage to the storm that caused it, or a scam wave to the event it exploits, and follows an incident as it escalates or subsides. The outcome is fewer duplicate alerts and a clearer account of what changed.

Answers checked against the evidence

Hey KAIA answers by voice or text, in English, French, Turkish or Hebrew, and works on the device first. A **language model** is good at reading and writing natural language, and poor at knowing when it is wrong. KAIA-OR therefore pairs it with symbolic checks, an approach known as **neurosymbolic reasoning**. The language model proposes an answer. The answer is checked against the evidence held in KAIA-OR before you see it, and it cites that evidence. Where the evidence does not support a statement, the assistant says what it cannot tell instead of filling the gap. The outcome is an assistant whose answers you can check.

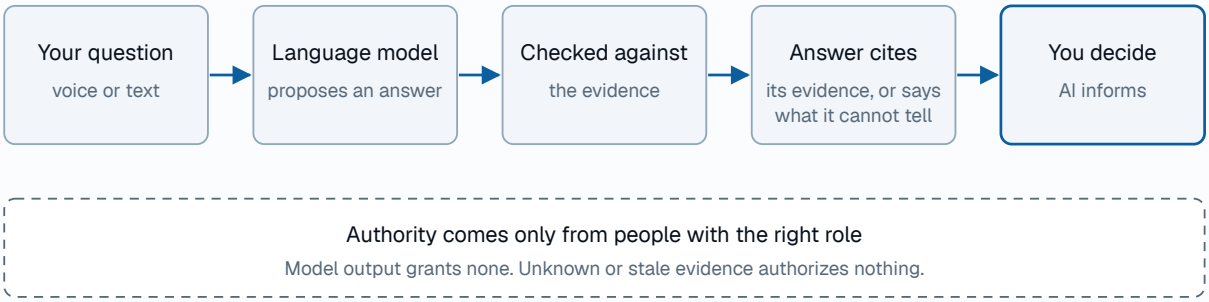


Figure 3. AI under evidence. The figure is conceptual; it names no model, engine or service.

On hazards, the assistant relays the publisher's own guidance. It never writes its own safety instructions. On messages and links, it suggests a next step, such as confirming through a channel you already trust, and the decision stays with you.

© Kaia Technologies, Inc.

Authority stays with people

Model output grants no authority. Nothing a model produces can approve an action, change a permission or widen what an agent may do. Authority comes only from people with the right role, through the product's own controls. This holds even when the content a model reads has been crafted to manipulate it. The assistant is designed to resist instructions smuggled into the content it reads and attempts at social engineering aimed at the AI itself.

Bounded scanning agents. Scanning agents carry out defined tasks within limits that the person or organization sets, for example looking for newly reported events that touch an organization's software or suppliers. They act only with your authorization. Unknown or stale evidence never authorizes an action. No agent takes a destructive step, sends a message or makes a purchase on anyone's behalf without explicit authorization.

These rules are what allow AI to be useful in security work without becoming another source of risk. The model helps people read faster and see connections. It never becomes the decision-maker.

© Kaia Technologies, Inc.

5. Honest scope and a buyer's checklist

Relevance-first intelligence is only as useful as it is honest about its limits. KAIA-OR is designed around four of them.

- **Coverage is broad, and it has limits.** No set of sources describes the whole world, or reports at the same speed in each place. KAIA-OR draws on official and authoritative sources worldwide and corroborates across independent sources. It does not claim to see all that happens, and an absence of alerts is not evidence that nothing has happened. When an event appears depends on when sources publish it and on corroboration, so KAIA-OR makes no promise about timing.
- **When a check cannot tell, it says so.** Some questions cannot be answered from the evidence available. KAIA-OR then says plainly that it cannot tell, and suggests how to find out, rather than offering a verdict it cannot support.
- **Publishers' guidance on hazards.** On natural hazards and public-safety events, KAIA-OR relays the publisher's guidance. It does not write its own safety instructions, and it does not give all-clear verdicts about places or trips.
- **Decisions stay with people.** KAIA-OR informs and suggests next steps. Responsibility for decisions stays with the person or the organization.

The questions below can be put to any vendor, including us. They separate relevance from volume.

A BUYER'S CHECKLIST FOR EVALUATING RELEVANCE

Ask	A strong answer
Who decides what is relevant: the vendor or you?	You define your places, trips, sites, software and suppliers, and you see what touches them.
Does each event show why it matched?	Yes, in plain language, next to the event, in the reader's language.
Are corroboration and confidence shown separately?	Yes. Agreement between independent sources is distinguished from certainty about the details.
Where does matching happen, and what does the vendor learn about where people are?	Personal matching happens on the person's device, and exact location stays there.
Are AI answers checked against evidence, and do they cite it?	Yes, before the reader sees them, and the assistant says what it cannot tell.
Can model output change permissions or trigger actions?	No. Authority comes only from people with the right role, and agents act only with authorization.
Does the organization's view stop at its sites?	Yes. It does not reach into its people's whereabouts, messages or devices.
How does the vendor describe its coverage?	Honestly: broad, with its limits stated, and with no promises about speed.

© Kaia Technologies, Inc.

What this means for you

Individuals and families. You see the hazards, incidents, disruptions and cyber events that touch the places you save and the trips you plan, each with why it matched and how well it is corroborated. Your places and exact location stay on your device. Ask Hey KAIA what an event means for you, by voice or text, in your language.

Teams without a security function. Without a security operations team, you see the events that touch your sites, with severity, corroboration and the reason they matched, and a suggested next step. Each member keeps their own alerts and their own second look.

Enterprise security, travel and cyber teams. Physical, cyber and supply-chain exposure appear in one shared picture, matched to your sites, software and suppliers. Each event carries severity, confidence, source and corroboration status, and events can flow into your own tools through the API and webhooks.

Government organizations. Your sites and missions appear in one picture your team can brief from, each event with its source, corroboration status and confidence. Government organizations start with cloud AI off. Model output grants no authority, and scanning agents act only with your authorization. Your officials decide.

To evaluate KAIA-OR, request early access on the KAIA-OR website. Enterprise and government teams can also choose Contact us on the KAIA-OR website.

Glossary

Term	Meaning in this paper
Corroboration status	Whether independent sources report the same event. Reports without independent corroboration are labelled as unconfirmed.
Ontology	A shared vocabulary of the kinds of events, places and assets KAIA-OR understands, so that the same event is recognized whatever words describe it.
Knowledge graph	The connected form in which that vocabulary and the relationships between events, places and assets are held, so that related events can be linked.
Language model	AI that reads and writes natural language. In KAIA-OR it proposes; it never decides.
Neurosymbolic reasoning	Pairing a language model with symbolic checks, so that answers are checked against the evidence before anyone sees them.

Sources

1. Allianz Commercial, Allianz Risk Barometer 2026, January 2026. Market context only.

© Kaia Technologies, Inc.

WP-1 · Version 1.0 · 3 November 2026

Capability availability depends on the edition and deployment. Capabilities are described as the committed KAIA-OR product scope.