

REALM + 7AI

Your agents are only as good as the data they can reach

Realm is the SOC-aware security data pipeline: full control of your telemetry between source and destination. 7AI's agents take the non-human work of the SOC: they detect, investigate, respond, and hunt, with your team on the loop. Realm sits upstream: your SOC decides what to keep, what context to add, and what shape data arrives in. The result is agents that reach verdicts faster, hand back fewer inconclusive investigations, and work at their actual ceiling.

The center of gravity in the SOC has moved

Investigation and hunting are moving to where the data sits. That move does not settle the data question: what exists to be reached, and in what condition.

Most pipelines cannot answer it. They are transport layers: rules written once, a little filtering, hand-tuning when something breaks. That held up when there was one destination and requirements changed once a year.

It does not hold up now. Agents draw on the SIEM, data lakes, and cloud storage at once, and what each swarm needs from the data changes faster than anyone can hand-tune.

Agents need more than raw alerts: they reason against users, assets, policies, and history. Federated search reaches what exists. Which sources were connected, whether they are complete, and what was collected in the first place is decided upstream.

So the team has to run a pipeline that understands the data and makes decisions about it, and know the systems their agents query are complete before an investigation finds the gap.

What changes with Realm upstream

THE STATUS QUO

Static filters written once decide what every destination sees. Volume drops, and nobody can say what was cut until an investigation goes looking. Every new destination means re-plumbing collection, source by source.

WITH REALM

Realm reads what each source contains, field by field, and what each destination needs, then recommends what to keep, cut, normalize, enrich, and route, with the reasoning shown. Your team approves, and everything not sent downstream is retained in Data Haven, searchable and ready to be resupplied. **Collect once, cut with proof, keep everything.**

AT A GLANCE

Independent collection

Add or change destinations without re-plumbing a single source.

In-stream enrichment

IP intelligence and threat context attached before data lands anywhere.

Detection Integrity

Every reduction checked against the detections you run, mapped to MITRE ATT&CK.

Complete retention

Everything not sent downstream is retained in Data Haven, searchable directly.

Data composition

When volume moves, see the exact field and value driving the change.

0

Detections lost at Vensure while SIEM volume fell 83%.

What Realm does for your agents

The full picture, protected

- Detection Integrity:** Before any reduction ships, Realm checks it against the detections you run and protects the fields they depend on, mapped to MITRE ATT&CK. Nothing your agents work from goes silently missing.

Context for faster verdicts

- Enrichment:** Realm attaches IP intelligence and commercial threat intelligence while data is still in stream. Events arrive carrying the context an agent would otherwise assemble mid-investigation, so verdicts come faster and fewer end inconclusive.

AI-ready data

- Realm Lake format:** Every event lands in your lake as one consistent record that holds the raw log, the parsed fields, the observables, and the enrichments. There is no normalization project first. Your team can put an agent on it the week it lands, not six months later. The format is open JSON and destination-agnostic: proven in production today, and the same record works wherever your AI reads from next.

Key outcomes

- Lower cost, more context:** Cut SIEM ingestion while 7AI works from enriched, complete, consistent sources.
- Detection coverage held:** Reductions are validated against the detections you run before they ship.
- Better verdicts, fewer dead ends:** with more context in every event and full history in reach, more investigations resolve instead of escalating back to your team.
- Vendor independence:** Change SIEMs, storage, or analytics tools, even the SIEM itself, without re-plumbing collection.

If AI is going to run the SOC, AI should be preparing the data it runs on

7AI reads your data wherever you keep it: the SIEM, cheap storage, your lake. What it does not do is sit at the source: it does not govern what gets collected, what context those stores carry, or what reaches them in the first place. It reasons over what exists, and that is the right design.

Realm is what makes that design pay off. Federated search is only as strong as the stores it reads, and Realm fills the stores you control with data worth reading: raw preserved, context attached, consistent from source to store. Only Realm checks every reduction against the detections you actually run before anything ships. Your team makes the call, with a higher ceiling on what your agents can find. That is the whole idea behind Realm.

GLOBAL MANUFACTURER

One of the largest Microsoft Sentinel deployments in the world.

3×

return per \$1 spent

1,000+

KQL rules replaced

72 hrs

to replace them

Prove it on your data.

Pick one source and start a 7-Day Data Assessment. We show the reduction and the detection integrity, on your data.

[Start an assessment → realm.security](https://realm.security)

ABOUT REALM SECURITY

Realm Security gives security teams independent control over their telemetry before it reaches the SIEM, data lake, or AI agent. The platform understands what each source produces and what each downstream system needs, then guides teams on what to retain, cut, route. Headquartered in Boston, Realm deploys in about a week and typically cuts SIEM ingestion costs by 50% or more.



www.realm.security

LinkedIn: /company/realm-security

X: @Realm_Security