




PREDICT & PREVENT

The Next Era of Cybersecurity

*How generative, agentic AND predictive AI together
turn security operations from reactive to proactive*

Paul Jespersen
Co-Founder — PRE Security


Ingest | Correlate | Detect | Respond | Predict | Prevent



Exactly eight years ago, at this conference, I introduced a new solution no one here had heard of yet: **XDR**



2018 — the first XDR demo booth in Switzerland



2018 — introducing XDR from the ITLS stage

THEN

Tedious, manual threat hunting across silos



2018: XDR

ML correlates many alerts into one detection



NOW

Every vendor has an XDR product

WHAT IF YOU COULD PREDICT BAD ACTORS AND INCIDENTS BEFORE THEY CAUSE HARM?

\$4.99M

Average cost of a breach — a record, up 12%

IBM / Ponemon, Cost of a Data Breach 2026

247 days

To identify and contain — rising for the first time in 5 years


IBM / Ponemon, Cost of a Data Breach 2026

29 min

Average attacker breakout time. Fastest: 27 seconds.

CrowdStrike Global Threat Report 2026

The AI arms race stopped being theoretical **this year**



Human speed vs. machine speed: the math no longer works

ATTACKER

29 min

to break out laterally

4 min

to exfiltration (observed)

< hours

from disclosure to exploit

DEFENDER

55 min

before an analyst opens the alert

75 min

to investigate one alert

43 days

to fully patch (median)



28% of alerts never investigated

60% say one of those later proved material



40% have switched off a detection

...for lack of capacity to investigate it



76% report burnout


81% saw workloads rise last year



70% of junior analysts gone in < 3 years

Average tenure ≈ 2 years

We built an industry that starts the clock **after** the breach



The uncomfortable truth: an AI adversary runs the whole kill chain faster than a human SOC completes step one. More rules, parsers, or analysts don't change the shape of the problem.

3 kinds of AI. You need all three.



GENERATIVE

Understands & creates

Turns raw telemetry into meaning.
Writes what humans used to hand-craft.

"What does this mean?"



AGENTIC

Reasons & acts

Plans the investigation, decides, and
responds within guardrails.

"What should we do — and do it."



PREDICTIVE

Anticipates

Reads attacks as stories and surfaces
the incident that hasn't fired yet.

"What happens next?"

Generative AI: the beginning of the end of hand-built security

WHAT IT CHANGES

- Parsers now generated on the fly — any source, any vendor, any schema change
- Detections written as reasoned intent, not brittle and limited field-matching rules
- Manual Threat Hunting replaced with generative correlation and agentic triage
- Investigation in plain language: ask the question, get the evidence and explanation
- Logs become language — PRE Security's patented log2nlp foundation is a key enabler

WHERE IT STOPS

- It drafts and summarizes — it doesn't act
- It describes up to today — it doesn't know now or forecast tomorrow

REALITY CHECK

Gartner: security copilots are sliding into the **Trough of Disillusionment**. A chatbot on a SIEM was the on-ramp, not the destination.

Agentic AI: from cute assistant to assassin operator

THE AGENTIC SOC

- Triage every alert — including the 28% nobody reaches
- Investigates: plans, gathers evidence, writes the case
- Responds inside explicit autonomy boundaries
- Proposes new detections from what it finds
- Human on the loop — analysts keep containment authority

45–61%

faster investigations with an AI SOC agent (148-analyst controlled study)

Cloud Security Alliance, 2026

68% → 97%

investigation accuracy, manual vs. AI-assisted

Cloud Security Alliance, 2026

\$1.93M

saved per breach with extensive AI & automation;
65 days faster

IBM Cost of a Data Breach 2026

BUYER BEWARE

- 70% of large SOC's will pilot agents by 2028 — less than 15% will see measurable gains
- Gartner warns of "agent washing" four times in 1 report
- Most "agentic" products are assistants w a human bottleneck
- Anchor evaluation on mean time to contain
- Even better: move to predictive, proactive SecOps

Predictive AI: changing the game to Proactive SecOps



Think weather, not smoke alarm

A smoke alarm tells you the house is already on fire.

A forecast tells you where the storm lands tomorrow, with a confidence score and a recommendation to board the windows tonight.

Cyber has spent the past twenty years perfecting smoke alarms.



Behavior as narrative

Each step is weak evidence. The sequence is strong evidence — before exfiltration.



Posture forecasting

Yesterday, today, tomorrow — with the drivers and the explanations. No black box.



Reasoned Probability, not binary scores


Live risk score and confidence band on every entity.




Prevention that acts early

Cross the threshold → harden, isolate, verify. **Before** harm.

You already trust predictive AI — every day



The more you type, the better the answer gets.



The better your data, the better the predictions.

Traditional vs. AI Native SecOps

TRADITIONAL

Deterministic • Rigid • Pre-built • Reactive • Expensive

- ✗ Hand-coded parsers & detections
- ✗ Manual integrations and queries
- ✗ Rigid SOAR playbooks
- ✗ Alerts on known signatures
- ✗ Posture measured after the audit
- ✗ Clock starts at detection
- ✗ Ingestion and Token Based Pricing



AI NATIVE

Generative • Agentic • Predictive • Proactive • Cost Efficient

- ✓ Generated parsers & intent-level detections
- ✓ Plain-language investigation
- ✓ Autonomous triage & response, with guardrails
- ✓ Behavioral stories scored as they unfold
- ✓ Posture forecast for tomorrow
- ✓ Clock starts at prediction — and prevents
- ✓ Fixed, predictable, Asset based pricing



The Proactive Era

From reactive to proactive — the analysts have already called the next era

“

*Preemptive cybersecurity will soon be the new gold standard... DR-based cybersecurity will no longer be enough to keep assets safe from AI-enabled attackers. Organizations will need to deploy additional countermeasures that act **preemptively** and **independently** of humans to neutralize potential attackers before they strike.*

— Carl Manion, Managing Vice President, Gartner

GARTNER PREDICTS

<5%

of IT security spend was preemptive in 2024




50%

by 2030 — replacing standalone detection & response as the preferred approach

Gartner press release, Sept 2025

Our vision: **AI Native Predictive SecOps**



- 1 Generative builds understanding**
Every log becomes language. Parsers and detections are generated, not written.
- 2 Predictive reads the story**
Sequences scored as they form. Posture forecast for tomorrow.
- 3 Agentic acts on the forecast**
Harden, verify, isolate — early, within policy — and feed results back.

Our foundation: **log2nlp** — machine data as language

PRE Security patented IP — U.S. Patent 12,608,539 B2 (issued April 2026), with additional related patents pending

United States Patent

(10) Patent No.: US 12,608,539 B2

Peterson et al.

(45) Date of Patent: Apr. 21, 2026

(54) **ARTIFICIALLY INTELLIGENT SYSTEMS, METHODS AND MEDIA FOR CANONICALIZING COMPUTER SYSTEM LOGS INTO NATURAL LANGUAGE PROCESSED REPRESENTATIONS FOR THE PURPOSE OF DATA ANALYSIS**

(71) Applicant: **PRE Security Inc.**, Salt Lake City, UT (US)

(72) Inventors: **John Luis Peterson**, Morgan Hill, CA (US); **Paul Soren Jespersen**, Lopez Island, WA (US)

(73) Assignee: **PRE Security Inc.**, Salt Lake City, UT (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **19/216,504**

(22) Filed: **May 22, 2025**

(65) **Prior Publication Data**

US 2025/0292012 A1 Sep. 18, 2025

Related U.S. Application Data

(63) Continuation of application No. 19/075,168, filed on Mar. 10, 2025.

(60) Provisional application No. 63/565,961, filed on Mar. 15, 2024.

(51) **Int. Cl.**
G06F 21/57 (2013.01)
G06F 40/186 (2020.01)
G06F 40/58 (2020.01)

(52) **U.S. CL.**
CPC **G06F 40/186** (2020.01); **G06F 40/58** (2020.01)

(58) **Field of Classification Search**
CPC **G06F 40/186**; **G06F 40/58**
See application file for complete search history.

References Cited
U.S. PATENT DOCUMENTS

11,777,970 B1 * 10/2023 Wainer G06N 20/00 726/23
2021/0216928 A1 * 7/2021 O'Toole G06F 16/29
2022/0228959 A1 * 7/2022 O'Neill H04L 67/12
2023/0038671 A1 * 2/2023 Weber H04L 51/12
2023/0086863 A1 * 3/2023 Suleman G06F 21/57 726/22

FOREIGN PATENT DOCUMENTS

CN 115329382 A * 11/2022 G06F 40/30
CN 119773347 A * 5/2024 G06F 18/15

* cited by examiner

Primary Examiner—Hassan Mrbi
(74) **Attorney, Agent, or Firm**—Carr & Ferrell LLP;
Sean M. Cody; K. Brian Bathurst

(57) ABSTRACT

Provided herein is an exemplary system for canonicalizing computer system logs into natural language processed representations for data analysis, the system including a real-time data collector, a cyber security purpose-based large language model communicatively coupled to the real-time data collector, a multi-dimensional vector generator communicatively coupled to the cyber security purpose-based large language model and a vectorization index and a prediction engine communicatively coupled to the multi-dimensional vector generator.

20 Claims, 5 Drawing Sheets


```

graph TD
    V[VECTORIZE: Natural Language Alert] --> VI[VECTORIZATION INDEX]
    VI --> P[PREDICTION ENGINE]
    P --> R[Real Time Prediction]
    VI --> C[Current to alerts to natural language]
    C --> D[REAL TIME DATA COLLECTOR]
    D --> C
    D --> G[Generate Alerts]
    G --> V
    G --> S[Spoken Alerts]
    S --> V
    G --> C[Control Alerts]
    C --> V
    G --> D[Desktop Alerts]
    D --> V
  
```

The diagram illustrates a system architecture for canonicalizing computer system logs into natural language processed representations for data analysis. The process begins with a **VECTORIZE: Natural Language Alert** block, which feeds into a **VECTORIZATION INDEX**. The **VECTORIZATION INDEX** is connected to a **PREDICTION ENGINE**, which outputs a **Real Time Prediction**. The **VECTORIZATION INDEX** also feeds into a **Current to alerts to natural language** block, which is connected to a **REAL TIME DATA COLLECTOR**. The **REAL TIME DATA COLLECTOR** feeds into a **Generate Alerts** block, which outputs **Desktop Alerts**, **Spoken Alerts**, **Control Alerts**, and **Desktop Alerts**. These alerts feed back into the **VECTORIZE: Natural Language Alert** block.

Generate Alerts block details:


- Desktop Alerts**: Event: Windows Login Time (Source: Sysmon | Log: Security | ID: 1056 | Severity: High | Event Time: 12/28/2023/09:02:02)
- Spoken Alerts**: Event: Windows Login Time (Source: Sysmon | Log: Security | ID: 1056 | Severity: High | Event Time: 12/28/2023/09:02:02)
- Control Alerts**: Event: Windows Login Time (Source: Sysmon | Log: Security | ID: 1056 | Severity: High | Event Time: 12/28/2023/09:02:02)
- Desktop Alerts**: Event: Windows Login Time (Source: Sysmon | Log: Security | ID: 1056 | Severity: High | Event Time: 12/28/2023/09:02:02)



WHY IT MATTERS

- Attacks unfold like a story. Language is how you read a story. Words aren't what's interesting.
- Detections become questions of intent — resilient to new tools, vendors and tradecraft.
- It is the layer agents reason over and predictive models learn from.


What it looks like in practice: key AI Native capabilities



SignalGate™ AI Data Engineering

GENERATIVE


Dynamic AI driven Parserless™ data ingestion, prioritized, classified, and routed — with SIEM cost avoided in real time. (AI Native “CRIBL for Cyber”)



Predictive Posture Reports

PREDICTIVE


Yesterday, today and tomorrow's posture — with confidence and drivers, recommendations and actions.



ASO — Autonomous Security Operator including SOCGPT™

AGENTIC

Investigates, correlates across the estate, responds within policy – all in natural language with transparent processing and reasoning, explanations, and history.



Generative XDR+ & Predictions

ALL THREE

Traditional as well as Detections as intent and narrative; entities scored and forecast; anomaly, risk, and vulnerability factored in by design.

Five moves for the next 12 months

Whether or not you ever talk to PRE — this is how to stop defending at human speed

1



Make mean time to contain your anchor metric

It's the number the attacker's 29 minutes competes against.

2



Fix the data layer before the dashboard

Ask any vendor how they onboard a log source they've never seen.

3



Adopt agents with explicit autonomy boundaries

What acts alone, what asks first, what's forbidden — and the audit trail.

4



Become a predictive, proactive posture practice

Take one board metric from "last quarter" to "next month's forecast".

5



Bring the AI layer into scope


Inventory AI apps and agents.

The next era of cybersecurity
isn't just more Detect & Respond.

It's **Predict & Prevent.**



Paul Jespersen
Co-Founder — PRE Security
paul@presecurity.ai



Sources

- IBM & Ponemon Institute — Cost of a Data Breach Report 2026 (July 2026): \$4.99M global average; 247-day identify-and-contain; 1 in 4 malicious breaches AI-enabled at ~\$6M; \$1.93M savings and 65 days faster with extensive AI/automation; 20%+ reported breaches targeting AI models/apps.
- CrowdStrike — 2026 Global Threat Report (Feb 2026): 29-minute average eCrime breakout, 27-second fastest; +89% attacks by AI-enabled adversaries; 82% malware-free detections; GenAI tools prompt-injected at 90+ organizations.
- Verizon — 2026 Data Breach Investigations Report (May 2026): vulnerability exploitation 31% of breaches, #1 vector for first time in 19 years; 62% human element; third-party involvement 48% (+60%); shadow AI regular use 45% (from 15%); median full patch 43 days.
- Anthropic Threat Intelligence — Disrupting the first reported AI-orchestrated cyber espionage campaign (Nov 2025); MITRE ATT&CK Campaign C0062: 80–90% of tactical operations executed by AI across ~30 targets.
- Anthropic — Project Glasswing and Claude Mythos Preview cybersecurity assessment (Apr–May 2026): thousands of high/critical zero-days across every major OS and browser; 10,000+ found by partners in first month.
- Gartner — Hype Cycle for Security Operations, 2026 (June 2026); Validate the Promises of AI SOC Agents (2026); Top Cybersecurity Predictions (Mar 2026): AI SOC agents at Peak (1–5% penetration), AI assistants entering Trough; 70% of large SOCs to pilot agents by 2028, 15% measurable gains; 50% of IR effort to involve AI apps by 2028.
- Gartner — "In the Age of GenAI, Preemptive Capabilities, Not Detection and Response, Are the Future of Cybersecurity" (press release, 18 Sept 2025): preemptive cybersecurity to reach 50% of IT security spending by 2030, up from <5% in 2024; Carl Manion quote (slide 13).
- Prophet Security — State of AI in the SOC 2026 (n=250): 55-min average alert dwell, 75-min average investigation, 28% of alerts uninvestigated, 60% saw a skipped alert prove material, 40% disabled a detection for capacity.
- Tines — Voice of Security 2026 (n=1,800+): 81% workloads up, 44% of time manual/repetitive, 76% burnout. • Cloud Security Alliance / Dropzone AI controlled study (n=148): 45–61% faster investigations, accuracy 68%→97%. • ACM Computing Surveys: 70% of analysts ≤5 yrs experience leave within 3 years.