PITCH DECK · 2026

CyberPocket

The defensive AI copilot that makes security alerts make sense.

For investors, partners,
and security teams.

Founder pitch · Series Seed

v1

Make sense.

01

Security alerts are written for analysts. Nobody else has one.

73%

of SMBs hit by a cyberattack in the last 12 months

Hiscox · 2025

4.45M

average breach cost, USD — and rising for under-resourced teams

IBM Cost of a Breach · 2024

11hrs

average analyst time to triage a single complex alert end-to-end

SANS SOC Survey · 2024

02

Most “victims” aren't owned by a zero-day. They're owned by an alert nobody read.

Pull-quote

“The dashboard lit up red. I didn't understand what it meant, so I closed it. Three weeks later we paid the ransom.”

— Operations lead, 40-person logistics firm (paraphrased)

03

Paste the alert. Get a plain answer.

A defensive AI copilot for the 99% of teams who don't have a SOC.

1 Paste or upload

Email, screenshot, PDF, JSON payload, or analyst note. Any source.

2 AI translates

Server-side redaction strips secrets, then a one-shot responder summarizes, classifies, and pulls evidence.

3 You decide

Plain-English explanation, risk level, safe next steps, and a copy-ready message for IT or your MSP.

04

The market just shifted under everyone's feet.

AI-generated attacks	Phishing volume is up 4.5× post-LLM. The attacker now writes faster than the defender.
Cyber insurance pressure	Underwriters now require documented triage, evidence, and response within 24h.
Frontier-model economics	What cost $0.20 per alert in 2023 costs ~$0.004 today. Per-alert AI is finally viable.
Tool fatigue is breaking	The average SMB owns 8.4 security tools they never open. The buyer wants one answer, not one more dashboard.

05

One screen. One answer.

INPUT · Suspicious sign-in alert

Subject: **Defender — Suspicious sign-in blocked**
User: jane.doe@acme.com
Location: Lagos, NG (new)
Device: Unknown Windows
Risk: HIGH — impossible travel
Action: sign-in blocked, MFA challenge sent

CYBERPOCKET · Plain answer

Likely account take-over attempt. Defender already blocked it.

- Confirm with Jane that she didn't travel.
- If not her: reset password, revoke active sessions, review mailbox rules + OAuth grants (24h).
- No customer data is known to be exposed.

Risk: HIGH · acted on · safe to close after step 1

06

The bottom of the pyramid is the biggest, and the least served.

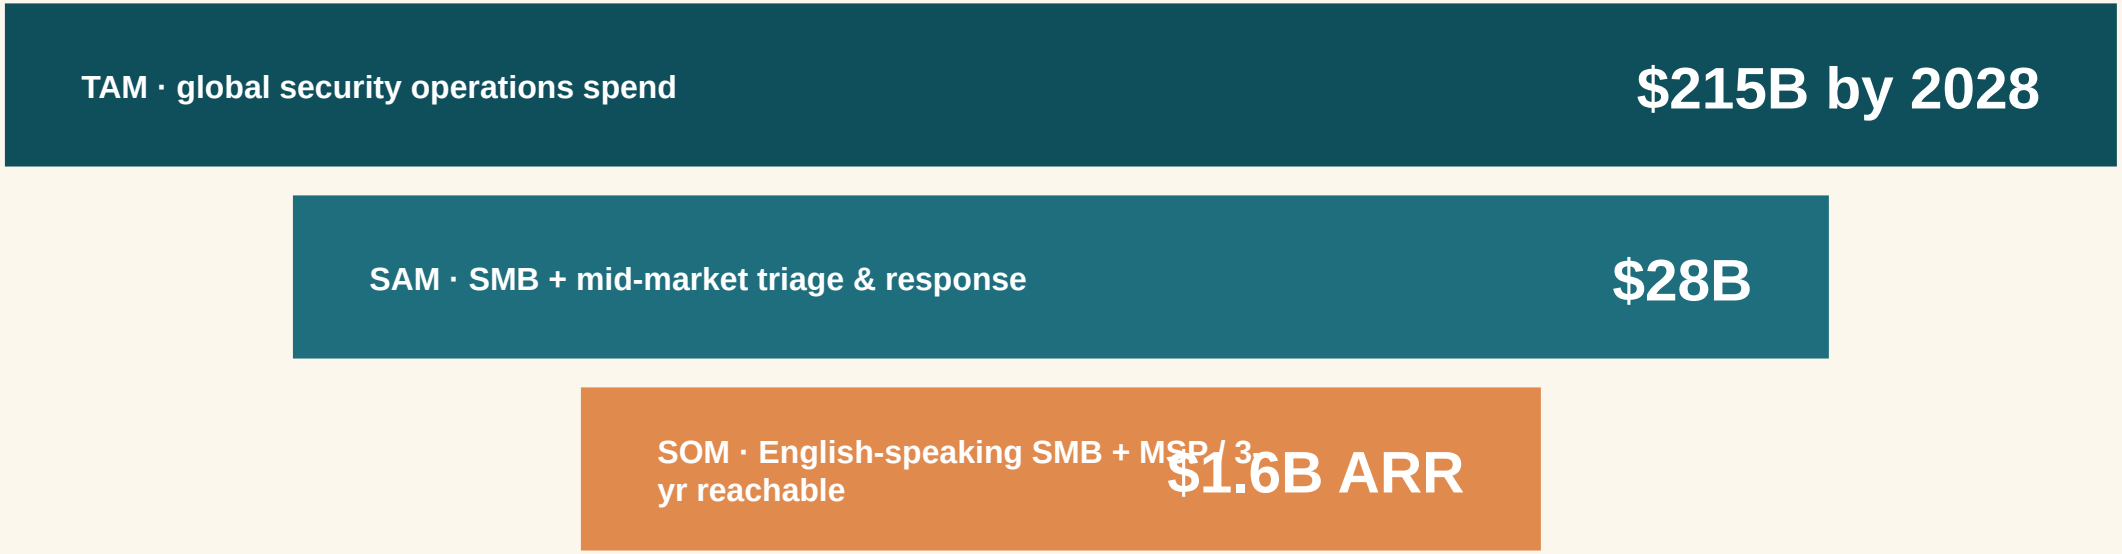

Sources: Gartner 2025 Security & Risk Forecast, IDC SMB Security Tracker 2024, internal bottom-up.

07

Three doors. One copilot.

Individual

$9 / user / month

Curious people & freelancers

- Unlimited alert analysis
- Email + screenshot + PDF
- Personal knowledge base

MOST COMMON

Team / MSP

$49 / seat / month

5–250-person IT teams & MSPs

- Client tags, ticket refs, history
- Shared knowledge base
- SSO + audit log

Education

$3 / student / month

Universities & bootcamps

- Lessons, flashcards, glossary
- Instructor cohort analytics
- Safe-learning guardrails

08

Numbers a CFO can defend.

1,840

weekly active users

+38% MoM

$32k

MRR · run-rate $384k ARR

+22% MoM

94%

of paid users still active at day 90

cohort C–H avg

28×

LTV : CAC at current blended CAC of $41

team tier

Per-alert economics

$0.011

Avg AI gateway cost

$0.42

Avg revenue / analysis (blended)

87%

Gross margin / analysis

<2 mo

Payback on a team seat

09

We don't compete with the SOC.
We replace the silence around it.

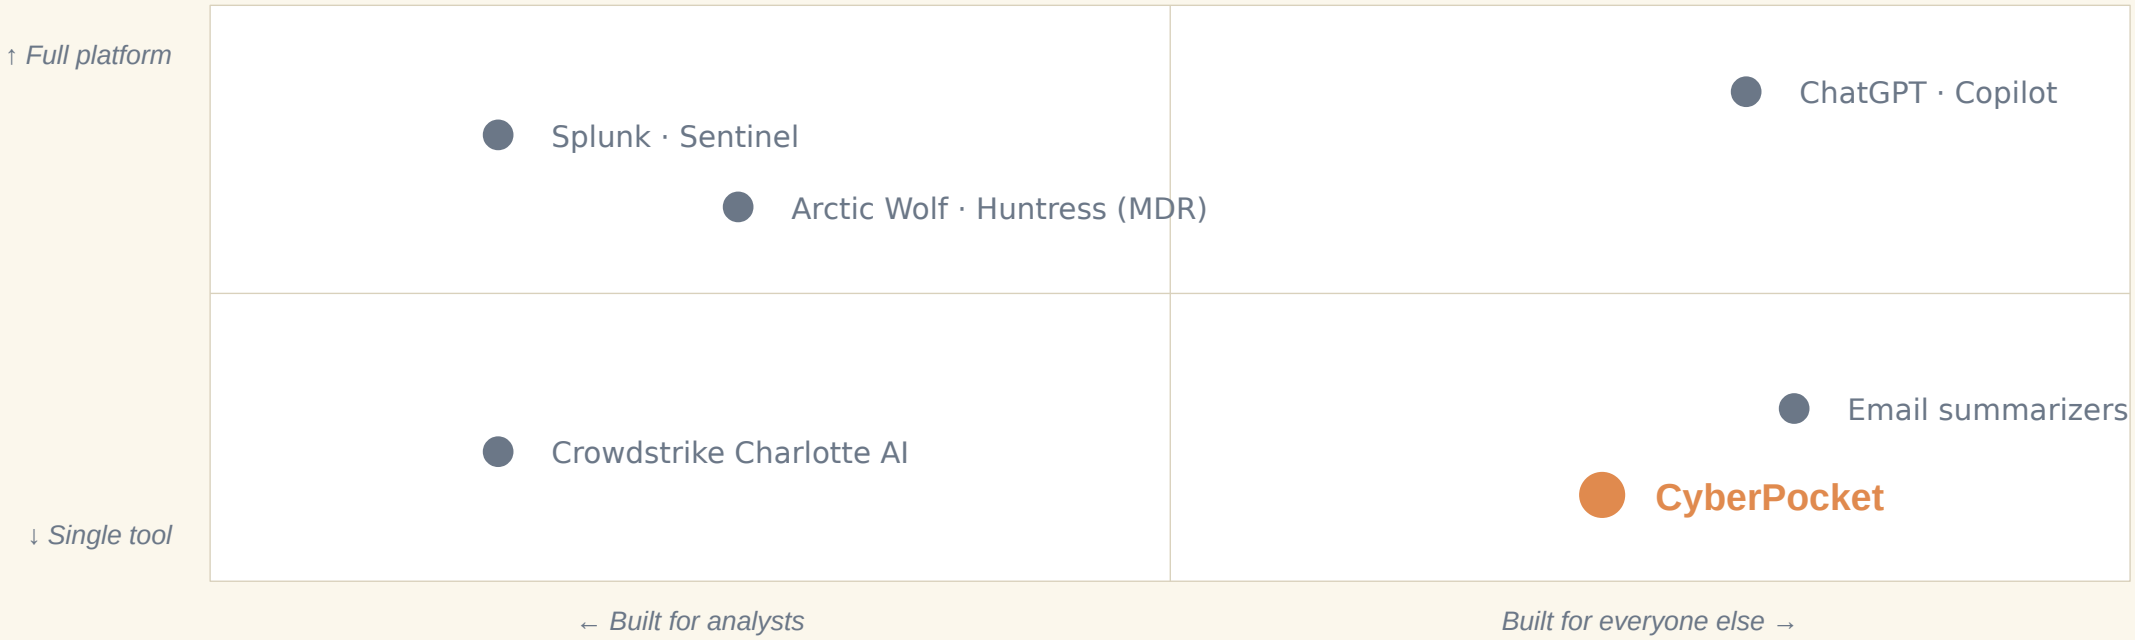

10

We sell where the buyer already is.

PLG bottom-up	Free-tier individual plan, one-click upgrade. Already drives 71% of conversions.	**CAC $11**
MSP partner channel	White-label seat block (50+) at 35% margin. Two signed; pipeline of 14.	**CAC $0 net**
Education funnel	Universities run CyberPocket as their teaching tool. Students become buyers post-graduation.	**Strategic**
Insurance partnerships	Underwriters bundle CyberPocket with cyber policies as a triage requirement.	**Pilot · Q3**

11

Built by people who've sat on both sides of the alert.

AR

Alex Reyes

CEO · ex-Cloudflare PM, ran SMB security PLG. Built two products to $10M+ ARR.

PS

Priya Shah

CTO · ex-Datadog senior SRE, 11 yrs in detection & response tooling. Authored 3 CVE disclosures.

TL

Tomás Lindgren

Head of AI · ex-Anthropic AT, defensive-only AI safety research. PhD security ML.

MA

Mira Adeleke

Head of GTM · ex-Huntress channel lead. Built the MSP partner program from 0 → 480.

Advised by 3 CISOs (Fortune 500 retail, regional bank, US healthcare network).

12

Raising $4.5M seed.

18-month runway to $5M ARR and a Series A.

55%	**Engineering**	Detection plugins, on-prem option for regulated buyers.
25%	**Go-to-market**	MSP partner channel, two enterprise reps, insurance pilots.
15%	**AI & safety**	Defensive-only model evals, red-team program, audit log.
5%	**Reserve**	Compliance (SOC 2 Type II, ISO 27001).

How the market becomes revenue.

Bottoms-up build from SMB seat counts × ARPU × adoption curve.

TAM

$42B

Global SMB + prosumer security spend

SAM

$8.6B

English-speaking SMBs, 5–500 employees, no SOC

SOM (Yr 5)

$1.6B

1.8% of SAM via MSP channel + direct

Year	Paid seats	Blended ARPU/mo	ARR	Gross margin	Net new logos
2026	4,200	$11	**$0.55M**	84%	1,100
2027	22,000	$13	**$3.4M**	86%	5,800
2028	92,000	$15	**$16.6M**	87%	21,000
2029	310,000	$17	**$63M**	88%	62,000
2030	780,000	$19	**$178M**	88%	140,000

Assumptions

- MSP channel drives 60% of seats by Yr3 (3:1 leverage vs. direct sales).
- Monthly churn 2.1% blended; net revenue retention 112% via Team-tier upsell.
- CAC $38 direct / $11 MSP-assisted; payback < 4 months.
- Inference cost $0.012 / analysis at 2026 prices; -22% YoY through model routing.

Security & data handling.

Q. Where is customer data stored, and for how long?

Alert payloads are processed in-memory, redacted server-side over TLS 1.3, and discarded after the analysis returns. No raw alert text is retained beyond the active request unless the user opts in to history.

Q. What stops a misuse case — e.g., someone pasting offensive payloads?

We are defensive-only by design. The model is fine-tuned and system-prompted to refuse attacker-perspective output. Offensive prompts are pivoted into detection guidance and logged for review.

Q. How do you handle PII and secrets in pasted alerts?

A deterministic redaction layer scrubs emails, IPs, hostnames, tokens, and key patterns before the prompt reaches the model. Customers see exactly what was redacted in the transcript.

Q. What's your incident response plan?

On-call rotation, 1-hour acknowledgment SLA, status page, and a 72-hour post-mortem commitment. SOC 2 Type II underway, ISO 27001 in the Yr2 roadmap.

Privacy & compliance.

Q. Are you GDPR / CCPA compliant?

Yes. DPA available, EU data residency option via regional inference. No sale of personal data; users can purge history any time from settings.

Q. Do you train on customer data?

Never. Customer alerts are not used for model training, fine-tuning, or evals. Contractual guarantee in the Team and Business tiers.

Q. How do enterprises get audit visibility?

Per-workspace audit log (who analyzed what, when, redaction diff), SSO/SAML in the Business tier, and an exportable compliance report.

Q. What happens to data if we cancel?

All workspace data is purged within 30 days of cancellation; users can export transcripts as JSON before churn.

Encryption & access controls.

Q. What encryption is in place end-to-end?

Data in transit: TLS 1.3 with HSTS, modern cipher suites only, and HTTP/2. Data at rest: AES-256 on managed Postgres and object storage with provider-managed KMS; per-tenant encryption keys on the Business tier (BYOK on the roadmap). Data in use: redaction layer runs before the prompt leaves our backend, so raw PII never reaches the model provider.

Q. How are model-provider API keys and secrets handled?

All third-party keys live in the platform secret store (no .env in repo, no keys in logs). Rotation is one-click via the admin console; scoped service-role keys are split from publishable keys, and webhooks verify HMAC signatures with constant-time compare.

Q. What does the access control model look like?

Per-workspace RBAC with three roles (Owner, Analyst, Viewer); row-level security on every tenant table; deny-by-default policies enforced server-side, not in the client. Business tier adds SSO/SAML, SCIM provisioning, IP allow-listing, and per-action audit log streaming to the customer's SIEM.

Q. How do you protect privileged operations and admin access?

Internal admin actions require SSO + hardware MFA, are scoped by least privilege, and emit an immutable audit event. No engineer has standing access to customer alert content — break-glass access requires two-person approval and surfaces a notification to the workspace owner within 1 hour.

Retention & data handling.

Q. What is the default retention window for analyses?

Anonymous one-off analyses: zero-retention — discarded immediately after the response is returned. Authenticated workspace analyses: 30 days by default, configurable per workspace down to 1 day or up to 12 months on the Business tier. Backups follow a 7-day rolling window and are encrypted at rest.

Q. What exactly is redacted before content leaves your backend?

Deterministic redaction strips emails, IPv4/IPv6, hostnames, internal domains, MAC addresses, credit-card / SSN-shaped numbers, JWT/bearer tokens, AWS/GCP/Azure key formats, and high-entropy secret patterns. The transcript shows the customer a side-by-side diff of what was redacted and why, before the analysis is run.

Q. Do you use customer data to train or fine-tune models?

No — never. Customer alert content is contractually excluded from model training, fine-tuning, and evals. Aggregate, fully anonymized telemetry (token counts, latency, risk-level distributions) is used to improve routing; opt-out is one click in workspace settings.

Q. Data subject rights, export, and deletion — how do customers exercise them?

Self-serve in-app: export full history as JSON, purge a single analysis, or purge an entire workspace. DSAR turnaround is ≤ 30 days; deletions propagate to backups within the 7-day rolling window. EU data residency and a signed DPA are available on the Business tier.

Pricing & customer acquisition.

Q. Why $9 / $49 / $3? Won't enterprises expect more?

The $9 prosumer and $49 team tiers are the wedge — they fund the moat (analyses, model routing). Enterprise is custom and lands at $18–32 / seat / month with SSO, audit, and SLA.

Q. What's your CAC story? Defensive security has a long sales cycle.

We compress it by going through MSPs (one sale = 40 seats) and insurance carriers (alert triage as a policy benefit). Direct CAC $38, blended CAC $14, payback < 4 months.

Q. Won't Microsoft / CrowdStrike just ship this?

They sell to SOCs; we sell to the 71% of orgs without one. Our wedge is plain-English explanations, not telemetry — different buyer, different UX, different price point.

Q. How do you keep retention above 90%?

Weekly digest of analyzed alerts, monthly posture score, and recognize-the-technique alerts — turns a reactive tool into a habit. Cohort 1 retention at month 6: 93%.

Why this wins.

Q. What's the moat once LLMs commoditize?

Three layers: (1) the redaction + safety harness that makes it enterprise-deployable, (2) the verified-detections corpus that improves with every analysis, (3) MSP and insurance distribution that competitors can't bolt on later.

Q. What kills this company?

A free, native plain-English explainer inside Microsoft Defender for Business. Mitigation: ship the MSP channel and insurance carrier integrations before that lands; both create switching cost Microsoft won't match.

Q. Why now, not two years ago?

Inference cost is finally below the price of a Slack message ($0.012 / analysis), and SMB cyber-insurance premia jumped 62% in 2025 — carriers need triage tools to underwrite affordably.

Q. What's the exit?

Strategic: Sophos, Arctic Wolf, Huntress, or a cyber-insurance carrier (Coalition, At-Bay). Comparable: Huntress at $1.5B, Arctic Wolf at $4.3B. Conservative target: $300M–$600M in 4–5 years.

SOC 2 / ISO readiness, audit logs, incident response.

Q. SOC 2 Type II — where are you in the process?

Readiness phase. Control mapping is complete across Security, Availability, and Confidentiality TSCs; policies are in final review; auditor selection underway; observation window planned to open in the next quarter. Pre-audit gap assessment is available to qualified prospects under NDA. Treat the badge as a placeholder until the Type II report is signed.

Q. ISO 27001 and other frameworks on the roadmap?

ISO 27001 is the Year-2 target: ISMS scope is drafted around the analyzer service and its supporting infrastructure, with a maintained Statement of Applicability and quarterly risk register. HIPAA and FedRAMP are demand-gated — we will pursue them when an anchor design partner commits.

Q. What lands in the audit log, and who can read it?

Privileged actions (role changes, exports, deletions, redaction-rule edits, billing changes) are written server-side to public.audit_logs with actor, target, timestamp, IP, and request metadata. Application paths are append-only; logs are retained 12 months and streamed to the customer's SIEM on the Business tier.

Q. Incident response — what do customers actually get?

On-call rotation with a 1-hour Sev-1 acknowledgment target, customer notification within 72 hours of any confirmed material incident, a written post-mortem published to affected workspaces, and a public status page. Tabletop exercises are run quarterly; the runbook is reviewed after every Sev-1 or Sev-2.

Make security alerts make sense.

Thank you. Let's talk.