

SOC360.AI

Adoption Book

A practical guide for security and business leaders: what the platform does, how to get value in the first 30 days, who does what, and an honest evaluation of what is production-ready today.

Contents

Read sections 1-3 for the business case, 4-6 to run an adoption programme, 7 for the honest evaluation.

- 1** What problem this solves
- 2** How the platform works, in plain terms
- 3** What you get out of it
- 4** Your first 30 days
- 5** Who does what — roles and daily rhythm
- 6** How to use each area of the product
- 7** Capability evaluation: ready, preview, or planned
- 8** What you must supply from outside
- 9** Measuring success
- 10** Frequently asked questions

1 · What problem this solves

Security teams are not short of alerts. They are short of clear, defensible answers.

Most organisations already own good security tools: a SIEM, endpoint protection, cloud security, a vulnerability scanner, a ticketing system. Each one produces its own list. Nobody owns the single view across them, so leadership questions go unanswered for days.

The everyday pain

Hundreds of daily alerts with no business meaning. Four tools telling four versions of the truth. Board reports assembled by hand in spreadsheets. No evidence trail when the auditor or regulator asks. AI tools spreading through the business with no register and no owner.

What SOC360.AI is

An aggregation, correlation and reporting layer on top of the tools you already run. It pulls your data in on a schedule, joins related signals into a small number of business actions, tracks deadlines, and produces the evidence and reports leadership needs.

What it deliberately is not

It is not a second place to investigate and close alerts. Assignment, triage and closure stay in your SIEM or ITSM. SOC360.AI mirrors their state so nobody works twice, and nobody argues about which system is right.

2 · How the platform works, in plain terms

Five steps that repeat automatically, around the clock.

Step 1 — Collect

Every few minutes the platform signs in to your connected tools with credentials you provide and pulls new records — incidents, vulnerabilities, identity risk, cloud findings.

Step 2 — Normalise

Different vendors describe the same thing differently. Records are rewritten into one shared shape so they can be compared.

Step 3 — Correlate

The platform extracts real observables — addresses, machine names, user accounts, domains, file hashes, CVE references — and joins records that share them, including matches against threat intelligence feeds.

Step 4 — Act and track

Joined signals become a small number of actions and cases with owners, severity, and response clocks. Notifications go out when a deadline is at risk.

Step 5 — Report and retain

Reports are generated on a schedule as PDFs, stored privately, and shared by secure link. Data is kept for the period you configure, with snapshots you can restore from.

Everything is separated per customer and per organisation at the database level, so one tenant can never see another tenant's data, regardless of what the interface requests.

3 · What you get out of it

The outcomes worth buying, stated as outcomes rather than features.

Fewer things to look at

Related signals collapse into one action with shared evidence, instead of arriving as separate alerts from separate tools.

A defensible board answer

A Board View and scheduled PDF reports that show movement versus the previous period, upcoming deadlines, and who is accountable.

No duplicated work

Two-way ticket synchronisation with ServiceNow and Jira keeps your existing process as the single place work is done.

Evidence on demand

Every change is written to an audit trail; reports and snapshots give you a dated record for auditors and regulators.

Control over AI adoption

A register of AI systems in use, EU AI Act risk classification with two-person approval, supplier questionnaires, and policy drafting.

Portability and recovery

Configurable retention per data type, scheduled snapshots, dry-run restore previews, and full export.

4 · Your first 30 days

A realistic sequence. Each week ends with something you can show to someone senior.

Week 1 — See it working

- Create your organisation and invite two or three colleagues with the right roles.
- Load the sample data set so everyone can see what a populated platform looks like.
- Walk the Start here group: Command Center, What Needs Action, Board View.
- Agree who owns adoption and book a 30-minute weekly review.

Week 2 — Connect your first real source

- Pick the single tool with the best data — usually the SIEM or endpoint platform.
- Create a read-only service account in that tool and record the credential in Connect Your Tools.
- Run the connection test, then let it run for 48 hours and compare counts with the source.
- Clear the sample data once you trust the real numbers.

Week 3 — Make it operational

- Connect ServiceNow or Jira so actions turn into tickets your team already works.
- Turn on threat intelligence feeds; CISA KEV needs no key and is a good first feed.
- Set your response deadlines and choose where notifications should arrive.
- Review What Needs Action daily with the operations team for a full week.

Week 4 — Prove the value

- Schedule the monthly board report and send it to the leadership audience.
- Set retention and snapshot policy for each data type, then run a dry-run restore.
- Start the AI inventory with the AI tools the business already uses.
- Review the first month: signals collected, actions raised, deadlines met, time saved.

5 · Who does what

Adoption fails when nobody owns the daily rhythm. Assign these four roles by name.

ROLE	OWNS	RHYTHM
Security leader (CISO)	Board View, monthly report, risk acceptance, AI governance sign-off	Weekly review, monthly board pack
Operations lead	What Needs Action, cases, deadlines, ticket quality	Daily, 15 minutes
Platform administrator	Connections, users and roles, retention, branding	Weekly, plus on change
Compliance / risk	Evidence, audit trail, AI Act classification, supplier assessments	Monthly, plus audit cycles

The daily 15 minutes

Open What Needs Action. Work the list top down — it is already ordered by urgency, number of joined signals, and confidence. Push anything real into your ticketing system. Check the deadline panel for anything about to breach. That is the whole routine.

The weekly 30 minutes

Open the Board View with the leadership sponsor. Compare against last week: what moved, what did not, what is overdue and why. Decide one improvement for the coming week — a new source connected, a rule tuned, or an owner assigned.

6 · How to use each area

The menu is grouped by what you are trying to do, not by which system the data came from.

Start here

Command Center for today's picture, What Needs Action for the work queue, Board View for leadership, My Exposure for where you are most at risk. → Use daily. If you only open one page, open What Needs Action.

Daily operations

Cases, incidents, alerts and deadlines, vulnerabilities, identity risk, Microsoft 365 and cloud posture, and ticket synchronisation. → Use when you need the detail behind an action, or to check a specific domain.

Threat picture

Live threat map, global coverage map, and threat intelligence feeds matched against your own records. → Use for context and for briefing non-security audiences; never as proof of attacker identity.

Reporting & evidence

Scheduled and on-demand reports, custom dashboards, export, retention and recovery. → Use monthly for the board pack and whenever an auditor asks for evidence.

AI governance (preview)

AI inventory, EU AI Act classification, AI security assessment, supplier questionnaires, control frameworks, policy drafting. → Use to get ahead of AI regulation; start with the inventory, everything else builds on it.

Customers & organisations

Customer portfolio, guided customer onboarding, organisation structure. → Use if you are a service provider or run several business units.

Setup & administration

Connect your tools, setup guide, users and roles, demo and live data, branding, administration. → Use at the start, then only when something changes.

7 · Capability evaluation

An honest assessment so you can plan around what exists today. Ready = working on real data. Preview = usable but still maturing. Needs input = blocked on something you supply.

CAPABILITY	WHAT IT DOES TODAY	STATUS
Tenant separation	Every record scoped to an organisation and enforced in the database	Ready
Access control	Sign-in, roles, audit trail on every change	Ready
Source collection	Scheduled pull from Sentinel, Splunk, Elastic, CrowdStrike, Defender for Cloud, Microsoft 365	Ready
Correlation engine	Joins records by shared address, host, account, domain, hash and CVE with evidence	Ready
Action grouping	Correlated signals collapsed into a ranked action queue	Ready
Cases	Durable cases rebuilt from correlations with owner and status	Ready
Deadlines and alerting	Rules, response clocks, webhook and in-app notification	Ready
Ticket synchronisation	Two-way with ServiceNow and Jira, including inbound webhooks	Needs input
Threat intelligence	CISA KEV, ThreatFox, AlienVault OTX working; OpenCTI and MISP available	Ready
Reporting	Scheduled PDF build, private storage, secure link delivery	Ready
Email delivery of reports	Blocked until a verified sending domain exists	Needs input
Retention and recovery	Per-type retention, archive before delete, snapshots, dry-run restore	Ready
Copilot assistant	Answers grounded strictly in your own tenant data	Ready
AI governance suite	Inventory, EU AI Act classification, assessments, supplier questionnaires, policies	Preview
Data residency choice	Recorded as a stated preference; storage location is not yet regional	Preview
Custom dashboards	Build and save views over live records	Preview

Anything not listed above does not exist in the product. Placeholder content such as industry benchmarks was removed rather than left in place, because a number nobody can defend is worse than a blank.

8 • What you must supply from outside

Short list, but adoption stalls without it. Assign an owner and a date to each line.

YOU SUPPLY	WHY	TYPICAL OWNER
Read-only service accounts for each security tool	So the platform can collect your data without write access	Security engineering
ServiceNow or Jira credentials and a project	So actions become tickets your team already works	IT service management
A verified sending domain	So reports can be emailed rather than shared by link	IT / messaging team
Threat intelligence keys (optional)	CISA KEV needs none; commercial feeds need a key	Threat intelligence lead
Agreed response deadlines	So the clocks and escalations reflect your policy	Security leadership
Named owners per area	So every action has somebody accountable	Security leadership

What you do not need

- No servers to buy, patch or scale — the platform runs on managed global infrastructure.
- No agent to deploy on endpoints for the core product; collection is through your existing tools' interfaces.
- No inbound access to your network; all collection is outbound over encrypted connections.

9 · Measuring success

Track these six numbers from day one. They are the ones an executive sponsor will ask about.

MEASURE	STARTING POINT	TARGET BY DAY 90
Sources connected	0	3 or more, covering detection, vulnerability and identity
Raw signals per week	Unknown	Counted and trending
Actions per week	Hundreds of alerts	A reviewable queue, worked to zero weekly
Deadline breaches	Not tracked	Tracked, and falling month on month
Time to produce the board pack	Days of manual work	Automatic, reviewed in under an hour
AI systems registered	0	Every AI tool in business use, with an owner

A simple value statement for your sponsor

"We reduced a weekly flood of disconnected alerts into a short, ranked list of actions with owners and deadlines; we produce the board report automatically instead of by hand; and we can show an auditor dated evidence for every claim we make."

10 · Frequently asked questions

Does this replace our SIEM?

No. It sits above the SIEM and the other tools, joins their output, and reports on it. Detection and investigation stay where they are.

Will our team have to work in two places?

No. Actions are pushed into ServiceNow or Jira and their state is mirrored back, so the team keeps working where they already do.

Where is our data held?

In a managed cloud database, separated per organisation, with private storage for generated reports. Regional storage selection is recorded as a preference today and is not yet enforced by location.

How long is data kept?

For the period you set per data type. Older records are archived before deletion, and snapshots let you restore a prior state after a dry-run preview.

Is our data used to train an AI model?

No. The assistant sends only the question and the tenant context needed to answer it, and the content is not used for model training.

What happens if a source goes down?

Collection is incremental and resumes from where it stopped. Connection health and the last successful collection time are shown in the interface.

Can we white-label it for our customers?

Yes. Logo, colours and product name are configurable per account, and service providers can manage a portfolio of customers.

How quickly can we see value?

A populated demonstration takes minutes. A first real source usually takes a day or two once the service account exists; the useful picture arrives within two weeks.