ThreatModeler

WHITEPAPER

Threat Modeling for the CISO and Team

What each stakeholder gets from agentic, governed threat modeling, and the answer the CISO can take to the board.

WHAT THE BOARD IS ASKING

The CISO Owns the Answer

The CISO owns risk across a portfolio that is being built faster than ever, with AI, across more clouds and services. When the board asks what the organization's AI is doing for security, the answer has to be more than a prompt someone ran. It has to be a model you can defend.

Trust is the gap. In recent research, only 32% of enterprises said they trust AI-assisted threat modeling a lot or completely, and just 17% for safety-critical or regulated applications. The opportunity is real, but only if the outputs are consistent enough to standardize, auditable enough to defend, and integrated enough to operationalize across engineering and risk.

Hanover Research, 2026 (n=250).

ThreatModeler® Nexus™ turns secure design into outcomes each stakeholder can use, grounded in one shared model and governed enough to stand behind. This brief walks through what the CISO, the architect, and the developer each get from the platform, and how the same work serves the wider room of directors, project managers, and risk leaders.

That is what built-in confidence means in practice: every application secure by design, across the portfolio, at the speed teams build, with answers you can defend.

THREE VIEWS, ONE MODEL

One Platform, Three Surfaces

The same secure-design work looks different depending on where you sit. ThreatModeler Nexus presents three surfaces over one connected model, so each role gets what it needs without a handoff or a separate tool.

CISO

Enterprise risk view

A single, current view of risk across the portfolio, with governance over AI and decisions you can defend to the board.

ARCHITECT

Secure design control plane

Reference architectures, control selection, and design intent kept current as the system changes.

DEVELOPER

Invisible security in the IDE

Mitigation guidance where the work already happens, handed to developers and their coding agents through the MCP Server.

Underneath all three sits the Secure Design Graph, a connected, queryable record of components, data flows, trust boundaries, threats, and controls. One source of truth means the same question gets the same answer, for every role, with a record behind every decision.

The rest of this brief takes each surface in turn, then shows how the same model serves the wider room and how it holds together underneath.

THE ENTERPRISE RISK VIEW

For the CISO

The CISO needs to see and govern risk across everything, not one application at a time. ThreatModeler Nexus gives that view, and the governance to trust it.

A SINGLE VIEW OF RISK

Posture across the portfolio in one place, so coverage and exposure are visible rather than inferred from scattered reports.

GOVERNANCE OVER AI

A deterministic framework, bring-your-own-AI, RBAC, SSO, and a complete audit trail. AI operates within governed bounds, on approved security content.

DEFENSIBLE DECISIONS

Every accepted risk and control carries a record, so the answer to an auditor or the board is one query away, not one search through chat history.

COMPLIANCE MAPPED IN

Designs map to more than 180 frameworks, so regulatory alignment is part of the model rather than a separate, manual exercise.

The result is a posture the CISO can act on and an answer the CISO can stand behind. Budget aligns to real exposure, AI adoption proceeds under governance instead of around it, and a board conversation about security risk has evidence underneath it.

Proof points and the platform underneath them follow on the closing pages.

For the Architect

The architect is responsible for secure designs that hold up as the system changes. The hard part is scale and drift: keeping intent current, controls consistent, and outputs repeatable across many teams.

DESIGN INTENT, KEPT CURRENT

Intended design lives in the Graph and updates as architecture changes, so the model does not quietly go stale between reviews.

CONTROL SELECTION WITH CONTEXT

Threats and mitigating controls resolve against the real architecture, including what upstream and downstream components already cover.

REPEATABLE, NOT BESPOKE

Reusable patterns and templates apply consistently, so the same design gets the same analysis regardless of who runs it.

FIND WHAT IS MISSING

Because the model holds intended design, it surfaces the gaps, not just the findings already visible in the code.

For a system already in production with no current model, the System Mapping Agent reads the code to infer the architecture, and the Secure Design Graph grounds that inference in intended design. That is what turns a guess into ground truth, and what lets the architect reason about what should be there, not only what is.

The architect surface is where intended design is defined and kept honest. It is the step that can move, but cannot be skipped.

INVISIBLE SECURITY IN THE IDE

For the Developer

Security should meet developers where they already work, not pull them into another tool. ThreatModeler Nexus puts guidance in the IDE and the pipeline, in the language of the work.

IN THE TOOLS YOU USE

Mitigation guidance arrives in the IDE and in pull-request checks, handed to you or your coding assistant through the MCP Server.

LOW FRICTION BY DESIGN

No new surface to maintain. The model runs alongside the build instead of standing in front of it.

SPECIFIC, NOT GENERIC

Guidance is grounded in this system's design, so it reflects the real architecture rather than a generic list of what could go wrong.

BUILT FOR AI-NATIVE WORK

As code is written, including by AI, the model keeps up, grounding inferred architecture in intended design rather than falling behind it.

For the developer, secure design stops being a review that happens elsewhere and becomes part of the work in front of them. The threats that matter for what they are building show up at the moment they can act on them, with the context to fix them.

This is the surface that turns secure design from an event into a habit, without slowing the sprint.

THE WIDER ROOM

And the Stakeholders Around Them

The same model serves the people who plan the work, ship it, and answer for it. None of this is a separate report; each role reads one source of truth through its own lens.

Directors and managers	Dashboards, trends, and checklists that show coverage and validate that the right controls are in place, across the application portfolio.
Project managers	Design risks identified up front, with consistent security requirements linked to the components a project actually uses.
Compliance and risk leaders	Framework mapping, evidence, and version history, so an accepted control can be traced and defended without reconstructing the reasoning after the fact.

Because every view draws on the Secure Design Graph, a question asked by a project manager and a question asked by an auditor reach the same model, and the same answer. Reporting becomes a lens on the source of truth, not another document to keep in sync.

HOW IT HOLDS TOGETHER

One Model, Behind Every Answer

The three surfaces and the wider room all draw on one foundation. A multi-agent system builds and reasons over the Secure Design Graph, so the work stays consistent, current, and traceable.

The System Mapping Agent builds the architecture, from design artifacts or by reading the code of a system already in production. The Graph Agent reasons over the model to find what is missing. The Reporting Agent turns the result into something you can defend. Throughout, AI operates within a governed, deterministic framework, on your architectural context and approved security content, to produce explainable, auditable outcomes.

50%
less threat-modeling effort

5x
faster threat models

13
granted patents

Charles Schwab case study; 5x at a regulated healthcare provider. The Secure Design Graph carries 2,500+ security requirements, 1,500+ threats, and 180+ compliance frameworks.

When the board asks what your AI is doing for security, the answer is a governed model you can defend.

ThreatModeler

BUILT-IN CONFIDENCE

Give every stakeholder the same answer.

See how ThreatModeler Nexus turns secure design into outcomes the CISO, the architect, and the developer can each use, grounded in one model and governed enough to defend.

TALK WITH OUR TEAM

threatmodeler.ai

For more information, support, or inquiries:

Email
support@threatmodeler.com

Phone
+1 201 266-0510

Web
threatmodeler.ai