ThreatModeler® + IriusRisk

Whitepaper

Build vs. Buy vs. AI

Navigating Threat Modeling in the AI Era

Executive Summary

AI and cloud-native development have accelerated how systems are built, while making them harder to secure over time. As a result, the traditional build-versus-buy decision for threat modeling now includes a third option: using AI directly.

Organizations typically consider three approaches:

- **Build an internal solution**
  Offers flexibility and control, and can work for narrow use cases or small teams. Over time, internal tools require sustained effort to maintain accuracy, governance, and alignment as architectures and frameworks evolve.
- **Use AI or LLMs directly**
  Provides fast, low-friction support for early exploration, learning, and ideation. These approaches excel at speed, but lack architectural grounding, repeatability, and governance required for long-term or audit-driven use.
- **Buy a purpose-built platform**
  Delivers consistent, governed threat modeling grounded in real architecture. Platforms reduce operational overhead and support scale, collaboration, and secure-by-design practices across complex environments.

This paper examines the tradeoffs across all three paths to help organizations choose an approach that remains effective as systems, teams, and security requirements continue to evolve.

Introduction: How the Build vs. Buy Debate Has Shifted in the AI Era

The question of whether to build internal security tools or buy established platforms has always involved careful tradeoffs. In recent years, that decision has grown more complex. AI and cloud-native architectures have expanded what teams can create on their own while also increasing the effort required to maintain accuracy, alignment, and long-term stability.

Earlier decisions were often framed around budget, staffing, and time to value. Today, organizations operate in environments where cloud services change quickly, infrastructure is increasingly software-defined, and distributed systems evolve continuously. These patterns introduce volatility that internal tools must absorb to remain reliable.

AI has shifted expectations further. Teams can now prototype quickly and explore ideas with far less friction. Yet this acceleration can create a widening gap between rapidly changing architectures and security processes that rely on manual updates or static models.

These shifts have expanded the decision beyond a simple build-versus-buy choice, introducing distinct approaches with very different tradeoffs in sustainability, governance, and scale.

This paper explores the tradeoffs across all three paths. The goal is to help teams think clearly about what each approach can offer, where each one introduces practical limits, and how to choose a path that will remain viable as their architecture and security needs evolve.

How AI Has Changed Expectations for Internal Tooling

AI has made it easier to begin building internal tools. Prototypes that once required significant engineering effort can now be created in a matter of hours. This has lowered the barrier to experimentation and encouraged teams to explore ideas that previously felt out of reach.

The ongoing work, however, remains substantial. Threat content, framework mappings, architectural coverage, and compliance alignment require continual updates. AI can assist with early drafts, but it does not maintain the accuracy or structure that long-term use demands. AI-generated outputs also vary, which is difficult for teams that depend on repeatable and traceable results.

Meanwhile, AI is accelerating the pace of development itself. Code and infrastructure change quickly, and internal tools that depend on manual updates often struggle to stay aligned with the systems they are meant to represent. AI helps get started, but sustaining accuracy and governance over time introduces responsibilities that must be considered before choosing a path.

These dynamics influence the build versus buy conversation in ways that did not exist even a few years ago.

Why This Matters for Threat Modeling

Threat modeling is most effective when it represents real architecture, produces consistent results, and can be reviewed and trusted over time. Lightweight tooling or AI-only workflows make these expectations harder to meet.

Models must stay aligned with actual cloud configurations rather than static diagrams that drift as environments change. Outputs also need to be repeatable, since teams rely on them for reviews, audits, and decision-making.

Structured governance is essential. Teams need clarity regarding version history, approvals, and rationale, especially when multiple groups contribute to the same system. Modern architectures evolve quickly, and emerging AI components bring new and unique risks that require more than informal or speculative suggestions.

These requirements help explain where DIY or AI-only approaches may fall short and why some organizations look for platforms that provide a more dependable foundation.

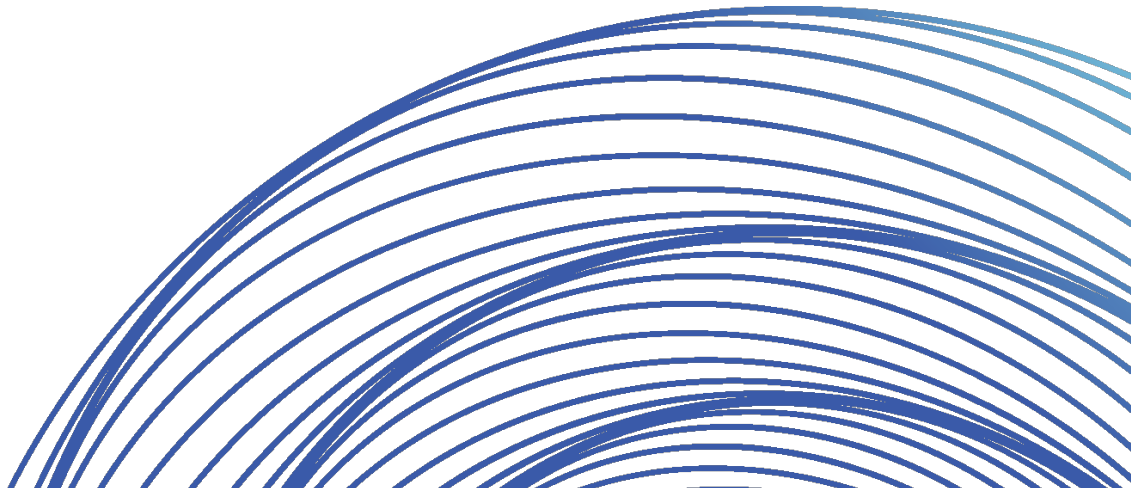

Approach: Build a Tool

Building internal tooling gives teams flexibility and direct control. Many organizations explore this path when they believe their needs are too specific or when they want to adapt workflows without relying on external vendors.

Why some teams choose to build

- Alignment with internal processes
- Ability to customize workflows and integrations
- Opportunity to experiment with AI-generated prototypes
- Desire to maintain full ownership of tools and data

Where AI can help

- Rapid prototyping and exploration
- Automating narrow tasks for individual teams
- Creating early proof of concept utilities

What building requires over time

- Regular updates to threat content and framework mappings
- Keeping models aligned with evolving cloud environments and security frameworks
- Added work to support collaboration and governance
- Development of versioning, permissions, and audit capabilities

Limitations that emerge as internal tools scale

- Architecture drift reduces accuracy unless it is actively managed over time
- AI-generated logic may not be consistent enough for audits
- Maintenance pulls security experts away from higher-value work
- As cloud architectures evolve, keeping models usable and understandable becomes increasingly complex

When building makes sense

Building is most suitable for small teams, narrow use cases, or short-lived prototypes that do not require long-term scale, governance, or sustained accuracy.

Approach: DIY AI Modeling

AI-only approaches appeal to teams that want immediate results without tool setup. They are helpful for early exploration and for individuals learning or documenting system behavior.

Why teams explore AI-only approaches

- Low barrier to entry
- Immediate feedback for exploration or learning
- Helpful for brainstorming or reviewing design concepts

Where AI-only methods provide value

- Early ideation
- Drafting initial threat lists
- Summaries for training or onboarding

Limitations of AI-only approaches

- Lack of architectural grounding and context
- Non-deterministic outputs are difficult to reproduce reliably
- No drift detection or alignment with cloud deployments
- No built-in governance or collaboration support

Where AI-only approaches break down

- Audit or compliance-dependent environments
- Multi-team engineering organizations
- Architectures that change frequently
- Workflows that require traceability or version history

The BYO LLM option

Running a model internally can help organizations address privacy considerations and maintain control over sensitive data. However, it does not resolve the core challenges related to architectural grounding, drift detection, or governance. These limitations remain even when the model is hosted within the organization's own environment.

Approach: Adopt an Enterprise Platform

Buying a platform appeals to organizations that need constant innovation, consistent accuracy, shared standards, and repeatable processes across teams and environments.

Why teams choose to buy

- Need for reliable modeling across distributed teams, frameworks, and systems
- Desire to reduce long-term internal maintenance and accelerate AI initiatives
- Requirements for auditability and structured governance
- Dependence on alignment across AI, cloud, and hybrid environments
- Need for a flexible and deterministic threat framework

Where buying provides value

- Modeling grounded in real cloud architecture
- Predictable and reviewable workflows
- Regular updates to threats, frameworks, and cloud architecture coverage

When buying is the strongest fit

A platform is most effective when organizations work at scale, operate across complex or fast-changing architectures, or maintain secure-by-design requirements that depend on accuracy and repeatability.

At a Glance: Build, AI, or Buy

Capability	Build (w/ or w/o AI)	AI-Only/ Prompt-Based	Buy a Platform
Accuracy	Depends on internal upkeep	No architectural grounding	Based on real cloud architecture
Repeatability	Varies	Non-deterministic	Consistent, governed workflows
Governance	Must be created internally	None	Built-in workflows
Maintenance	Continuous internal effort	Low effort but limited	Vendor maintained
Best Fit	Small or narrow use cases	Individual exploration	Multi-team and continuous modeling

The ThreatModeler Advantage

ThreatModeler offers a foundation that addresses the limitations often encountered in built or AI-only approaches. The platform connects directly to cloud environments, supports deterministic, repeatable workflows, and maintains threat and framework content on an ongoing basis. It also enables collaboration and uses AI within a governed structure that preserves traceability and accountability.

These capabilities support teams that must keep pace with rapid architectural change while maintaining stability and assurance.

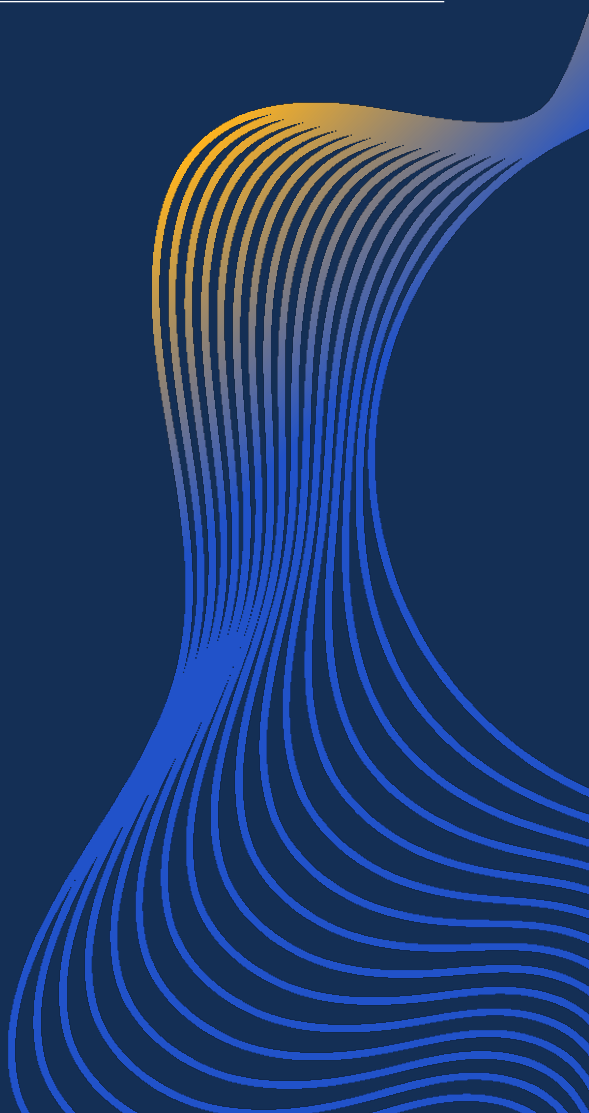

Conclusion: Choosing the Right Path in the AI Era

Each approach offers benefits and limitations. AI has broadened the possibilities for internal development, but it has not eliminated the need for sustained accuracy and governance. AI-only methods provide speed and flexibility, but are best suited for early exploration rather than long-term modeling. Platforms provide a more dependable structure for teams that require consistency, collaboration, and alignment with real systems.

The right choice depends on your architecture, your scale, and the level of assurance your organization needs. The goal is not to commit to a single category, but to choose the approach most likely to remain effective as your environment evolves.

Further Reading

This paper explores the tradeoffs between building, buying, and using AI for threat modeling. The following white papers expand on key considerations introduced here, including the responsible use of AI and how threat modeling practices evolve as organizations scale.

Operationalizing AI in Threat Modeling

Explores how AI can be applied within governed, repeatable threat modeling practices, and why standalone or prompt-driven approaches introduce challenges related to determinism, accountability, and long-term assurance. **Read more**.

The Threat Modeling Maturity Model

Outlines how threat modeling practices evolve over time, helping organizations assess their current state and understand when different approaches such as DIY, AI-only, or platform-based become necessary as scale, complexity, and assurance requirements increase. Learn more here.

Next Steps

If you are reviewing how Build, AI, or platform approaches can support your security and architecture objectives, **our team can help** explore practical examples and compare the paths based on your specific workflows and scale.