The Prompt That Could Break Your Business:
Are You Ready for the New AI Risks?

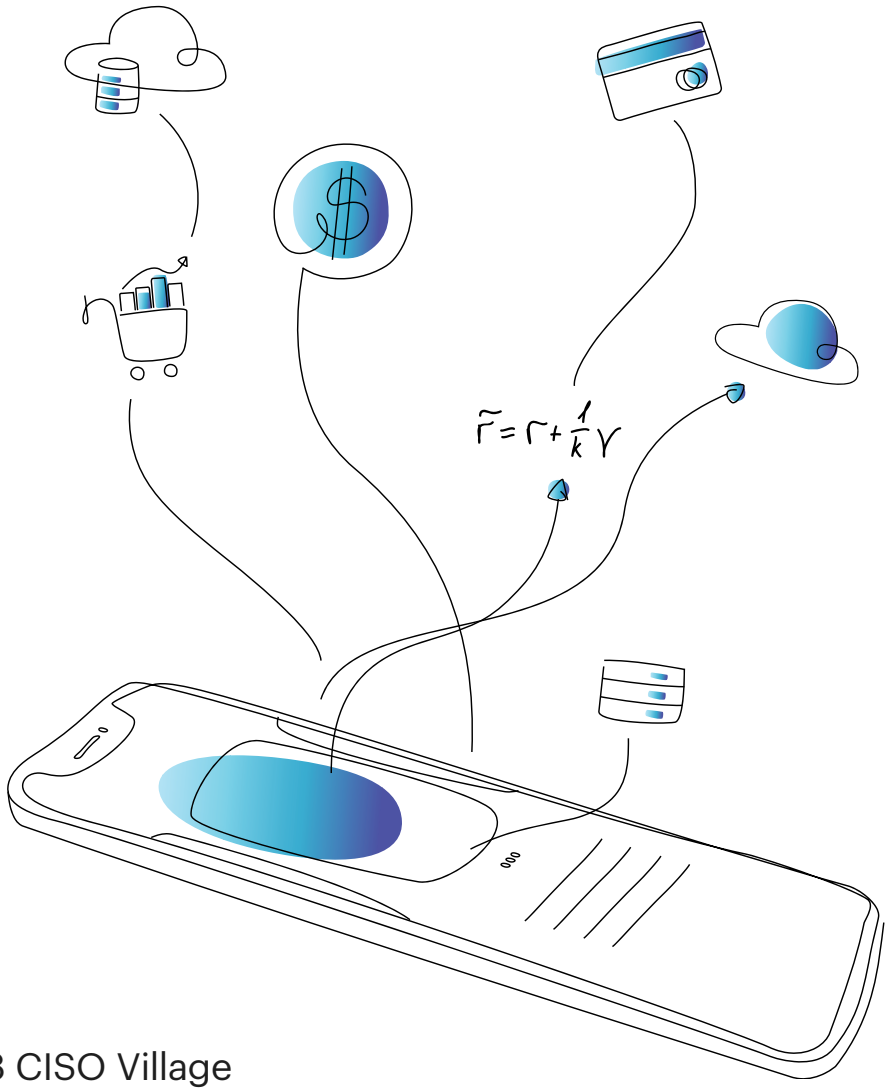

By the Team8 CISO Village
May 2025

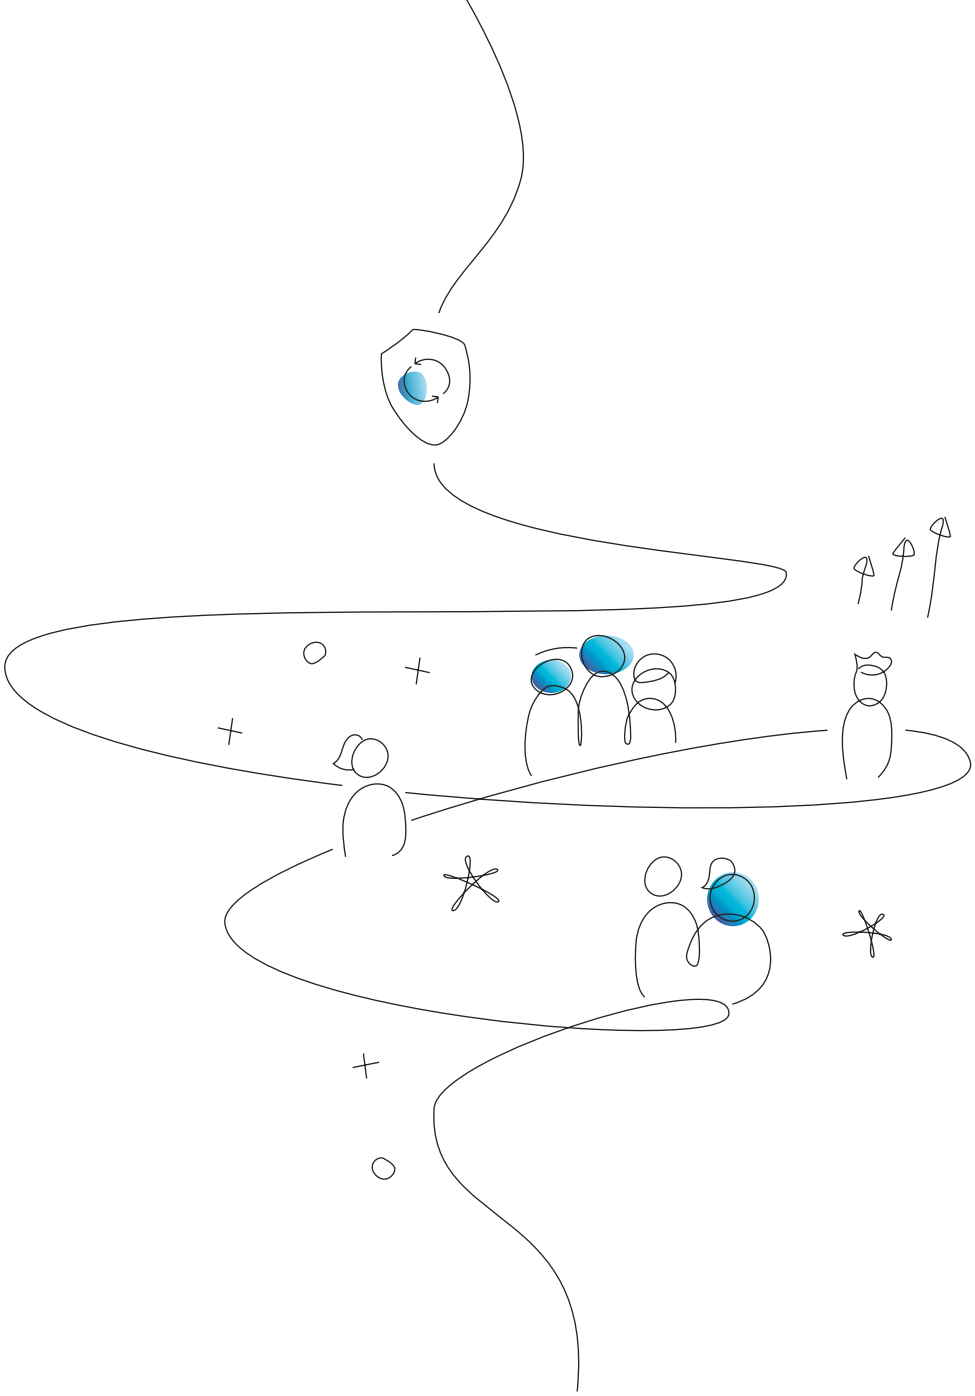

The Team8 CISO Village is a community of CISOs from the world's leading enterprises. The primary focus of the Village is to facilitate collaboration among the world's most prominent companies with the goal of sharing information and ideas, conducting intimate discussions on industry and technology trends and needs, and generating value and business opportunities for all parties.
By helping Team8 to identify real pain points and understand the requirements of large organizations, members of the Village are first in line to leverage solutions that are purpose-built by Team8's portfolio companies to support their needs.

To contact the Team8 CISO Village, please email **cisovillage@team8.vc**

WRITTEN BY

Ross Young
CISO-in-Residence
Team8

CONTRIBUTORS

Andrew Wilder
Chief Security Officer
at Vetcor

Michael Calderin
CISO at YAGEO Group

Devin Rudnicki
Chief Information Security Officer
at Fitch Group

Pieter Vanlperen
CISO at Own

Duane Gran
Director of Information
Security at Converge
Technology Solutions

Rock Lambros
Founder and CEO
of RockCyber, LLC

Kristen Beneduce
Deputy CISO at Nextdoor

Yabing Wang
VP, CISO & CIO, Justworks

Kurby Brown
Senior Director of Security &
Compliance at Turnitin

CONTENTS

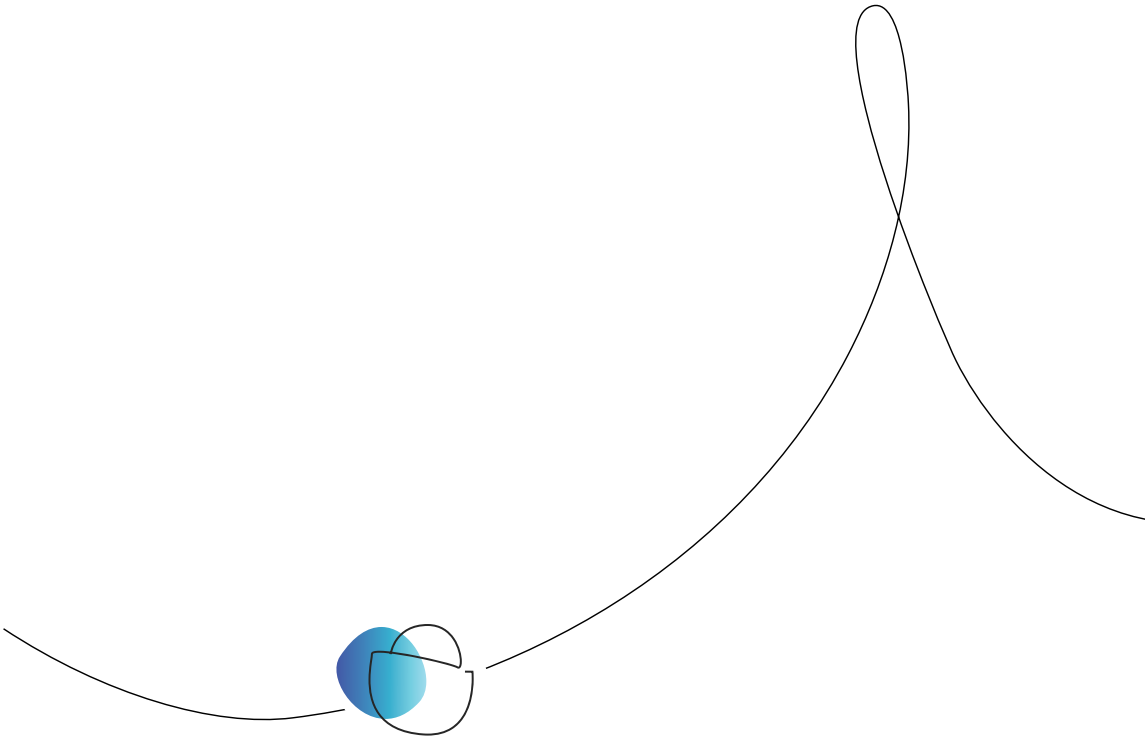

INTRODUCTION

"We didn't know we needed to check for that."

These nine words have launched a thousand post-incident reviews.

Imagine this: After months of successfully using AI to revolutionize customer service, you discover a devastating data leak. Sensitive customer information has been quietly siphoned away through seemingly harmless prompts. Your immediate reaction? "But we trusted our AI provider's security!"

Here's the uncomfortable truth every CISO must face: AI security isn't just your provider's responsibility—it's a shared, high-stakes challenge that demands executive oversight. And in 2025, the risks have never been higher.

Consider the consequences of poor AI governance:

- Knight Capital's $440M catastrophe – a faulty trading algorithm that led to financial ruin in under an hour.
- Amazon's AI hiring bias – an automated recruiting tool that systematically downgraded resumes with words like "women's" or "girl's."
- Intranet-exposing chatbots – generative AI tools deployed without guardrails, giving employees access to sensitive salary and HR data.

Regulators are taking notice. Liability for AI-driven failures doesn't just fall on engineers—it scales with executive negligence. Fines, reputational damage, and regulatory scrutiny are now business risks tied directly to leadership decisions.

AI isn't just another IT challenge. It's a boardroom issue. Investing in governance today isn't just about avoiding disaster—it's about securing the future of your enterprise. Are you ready to learn the most important tips to running safe AI?

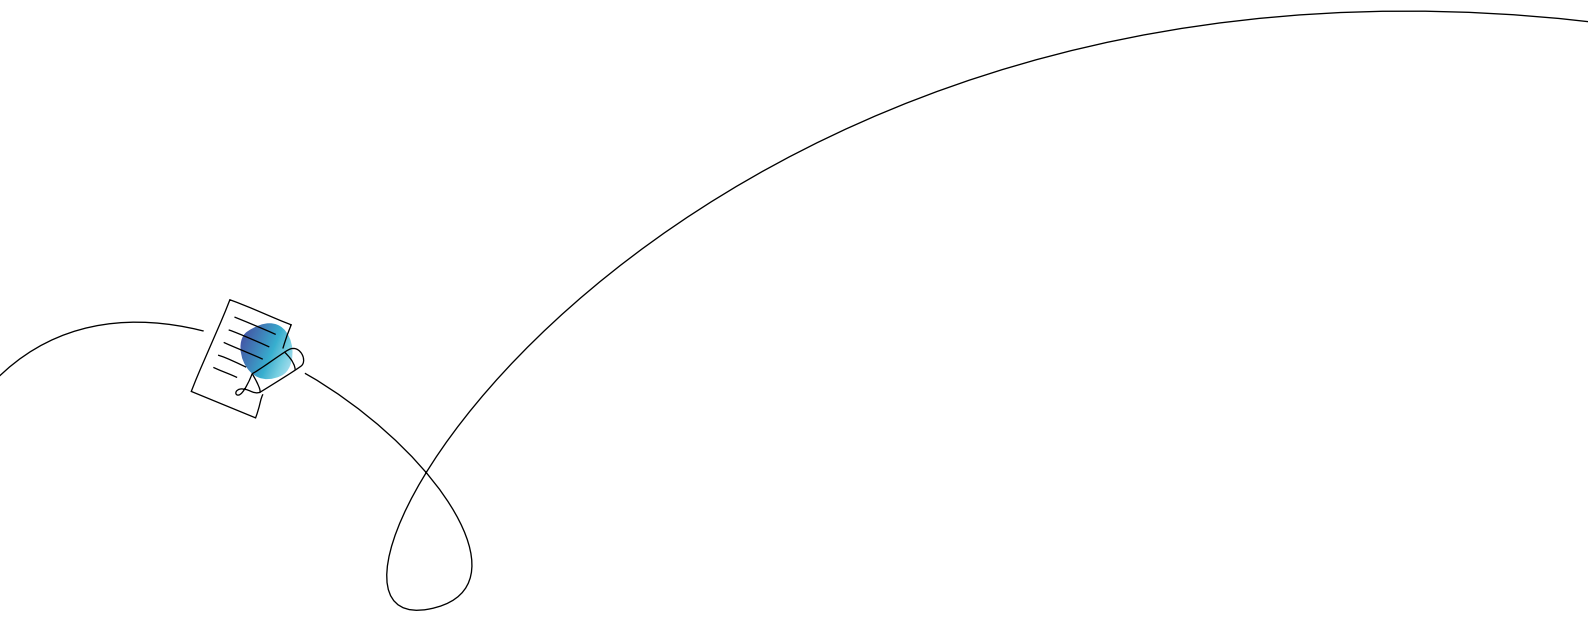

CISO Tip #1
Create a Cross-Functional AI Steering Committee

AI is transforming industries at breakneck speed, but without the right guardrails, it can expose your organization to unintended risks—security gaps, compliance failures, and even reputational damage. The solution? A cross-functional AI Steering Committee.

Think of AI governance like city planning. You wouldn't build roads without traffic laws, zoning regulations, and emergency services in place. An AI Steering Committee ensures your organization's AI initiatives are innovative yet controlled, compliant yet agile.

Why AI Governance Matters

Effective AI governance isn't about slowing innovation—it's about accelerating high-value AI solutions while blocking hype, unnecessary risks, and ethically misaligned applications. It:

- Aligns AI initiatives with corporate strategy, compliance, and ethical standards
- Streamlines AI adoption by providing clear frameworks for security and risk diligence
- Enables organizations to move fast while maintaining control over AI-related risk

AI STEERING COMMITTEES ENSURE YOUR ORGANIZATION'S AI INITIATIVES ARE INNOVATIVE YET CONTROLLED, COMPLIANT YET AGILE

Who Should Be at the Table?

To govern AI effectively, the right leaders must be engaged at the right level. At a minimum, an AI Steering Committee should include:

- Core Members (Voting): Legal, Privacy, Security, Compliance, Risk, Finance, IT, and Product/Engineering leadership (if AI is embedded in products)
- Informed Stakeholders (Non-Voting): Business units using AI (e.g., HR for hiring tools, Marketing for personalization, Customer Service for AI chatbots)

CISOs play a critical role in balancing security, risk, and innovation. While they don't necessarily need to lead the committee, they must be key influencers ensuring AI risk management remains a priority.

How to Structure AI Governance

Not all organizations need a brand-new AI committee—don't reinvent the wheel if existing governance structures (e.g., Data, Risk, or Ethics Committees) can incorporate AI oversight. However, if AI adoption is accelerating, an AI-focused governance body can help manage risk effectively.

Organizations typically structure AI governance in one or more of the following ways:

- AI Steering Committee – Defines AI policies, sets strategic goals, and ensures ethical alignment
- Review Board – Evaluates AI use cases, balancing business goals, speed, and risk
- Ethics Committee – Focuses on responsible AI principles like transparency and fairness
- Education & Advocacy Group – Ensures AI literacy, security training, and compliance awareness

Ensuring Effective AI Governance

With AI use cases emerging across the enterprise—from third-party AI tools (e.g., Salesforce AI, MS Copilot) to proprietary AI models embedded in products—CISOs must ensure governance frameworks evolve in real-time.

Effective committees:

- Map out AI usage across the enterprise (through SaaS discovery tools and user engagement)
- Create decision frameworks to separate high-value AI applications from risky distractions
- Provide continuous feedback loops to refine AI policies and security standards

Bottom Line: Don't Be the Department of Slowing AI Adoption

The right AI governance framework doesn't slow AI adoption—it makes it safer, smarter, and more strategic. Whether through an existing governance committee or a dedicated AI Steering Committee, CISOs must ensure AI adoption aligns with corporate risk management.

AI security isn't just an IT issue—it's an enterprise-wide responsibility. The question is: Does your organization have the right governance in place?

CISO Tip #2
Leverage an AI Shared Responsibility Model

Imagine moving into a luxury high-rise. The property owner ensures the foundation is solid, the walls are strong, and the utilities run smoothly—just like your AI provider secures their base model.

But here's the catch:
Would you leave your front door unlocked just because the building has security guards?

EVERY ENTERPRISE NEEDS AN AI SHARED RESPONSIBILITY MODEL THAT DEFINES WHAT YOUR PROVIDER SECURES AND WHERE YOUR DUTY BEGINS

Too many organizations assume their AI provider has security fully covered. The reality? AI security is a shared responsibility. If you don't know where your accountability begins, you're leaving the door wide open to risk.

That's why every enterprise needs an **AI Shared Responsibility Model**—a clear framework that defines:

- What your provider secures
- Which risks require joint oversight
- Where your responsibility begins

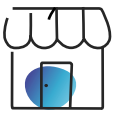

Provider
Model Development
Training Safety
Infrastructure Safety
API Security
Model Monitoring
Base Model Safety

Shared Responsibilities
Risk Management
Compliance
Incident Response
Privacy Protection
Documentation
Quality Assurance

Customer
Prompt Security
Output Validation
Access Control
Integration Security
Data Sanitization
Usage Monitoring

Just like building management and tenants must cooperate on security protocols, multiple stakeholders in AI—providers, customers, and internal teams—must work together to ensure robust protection. These responsibilities should be clearly defined in contracts and internal agreements, leaving no room for ambiguity.

Ignore this, and you're not just gambling with data—you're putting your enterprise, reputation, and regulatory standing at risk

CISO Tip #3

AI Security Isn't a Reboot—It's an Evolution

We're in the middle of an AI gold rush, and security teams are making a critical mistake: treating AI like traditional software.

AI is a completely different beast—one that requires a new security playbook. A single cleverly crafted prompt could:

- Extract confidential data from your systems
- Generate harmful content that damages your brand
- Bypass carefully constructed security controls
- Manipulate AI behavior in ways you never anticipated

But here's the good news: 30 years of modern security principles aren't obsolete. Many traditional security methods can still mitigate AI risks—if they're adapted correctly.

✓ **DLP & IDS:** Data Loss Prevention (DLP) and Intrusion Detection Systems (IDS) still help mitigate AI-driven data leaks—just extend them to detect prompt injections.

✓ **Identity & Access Management (IAM):** Least-privilege access is more critical than ever. AI systems act like users, so non-human identities (NHI) must be governed just as strictly as human accounts.

✓ **Data Tagging & Governance:** AI security starts with **knowing your data**—tagging and controlling it to ensure models only access what they should.

The **AI threat landscape is evolving fast,** but that doesn't mean security teams need to start from scratch. The best AI risk strategy? Lean into existing security approaches while adapting to this new frontier.

CISO Tip #4
Apply the 4 Tactical Phases of AI Governance

Phase 1: AI Discovery & Policy Foundation (Weeks 1-4)

- Identify Every AI System in Use—whether approved or shadow AI.
- Draft AI Governance Policies that align with regulations and business needs.
- Assemble a Cross-Functional AI Governance Committee—Legal, IT, Security, and Business leaders.
- Launch AI Awareness Training to educate employees on risks and responsibilities.

Key Outcome: A clear AI governance framework, an inventory of AI tools, and a baseline for risk management.

Phase 2: Building Governance Into the Business (Months 2-3)

- Extend Existing Security Controls (DLP, IAM, logging) to detect AI-specific threats like prompt injection.
- Monitor AI Usage & Compliance—track evolving regulations and security vulnerabilities.
- Set Up Continuous Feedback Loops to refine governance based on real-world AI use.

Key Outcome: A flexible, adaptive AI governance framework that evolves with technology and regulation.

Phase 3: Operationalizing AI Security & Compliance (Months 4-6)

- Embed AI Governance into Procurement & Development Workflows to enforce compliance from the start.
- Implement Dashboards for AI Risk & Performance Monitoring to detect drift, bias, and adversarial attacks.
- Strengthen AI Security Controls—MFA, role-based access, anomaly detection, content filtering.

Key Outcome: AI governance moves from policy to practice, preventing security from becoming an afterthought.

Phase 4: Future-Proofing AI Governance (Ongoing)

- Track Emerging AI Threats—deep fakes, synthetic fraud, AI-powered insider threats.
- Audit Explainability & Accountability—make AI decisions traceable and justifiable.
- Engage with Industry & Regulators to influence best practices and stay ahead of compliance shifts.

Key Outcome: A resilient AI security strategy that scales with innovation—without blocking business progress.

CONCLUSION

The Bottom Line: Your Call to Action

The AI Shared Responsibility Model isn't just another framework – it's your survival guide in the AI age. As these systems become more powerful and ubiquitous, implementing this model isn't optional – it's essential for survival.

Remember: Your AI provider builds the fortress, but you control the gates. You decide who enters, what they can do, and how they do it. Security isn't a spectator sport – it's a shared mission requiring constant vigilance from everyone involved.

Don't wait for an AI security breach to take this seriously. The question isn't whether you need this model – **it's whether you'll implement it before or after a crisis forces your hand.**

Next Steps for CISOs

1. Share this paper with your security team
2. Create a cross-functional AI Steering committee
3. Adopt an AI Sharing Model
4. Apply the 4 tactical phases to create quick adoption

Enhance your understanding of securing AI systems by exploring the insights shared on the Team8 Blog. For additional guidance, feel free to connect with Ross Young, Team8's CISO in Residence, at ross.y@team8.vc

Share this whitepaper with your security team and start the conversation about securing AI today!

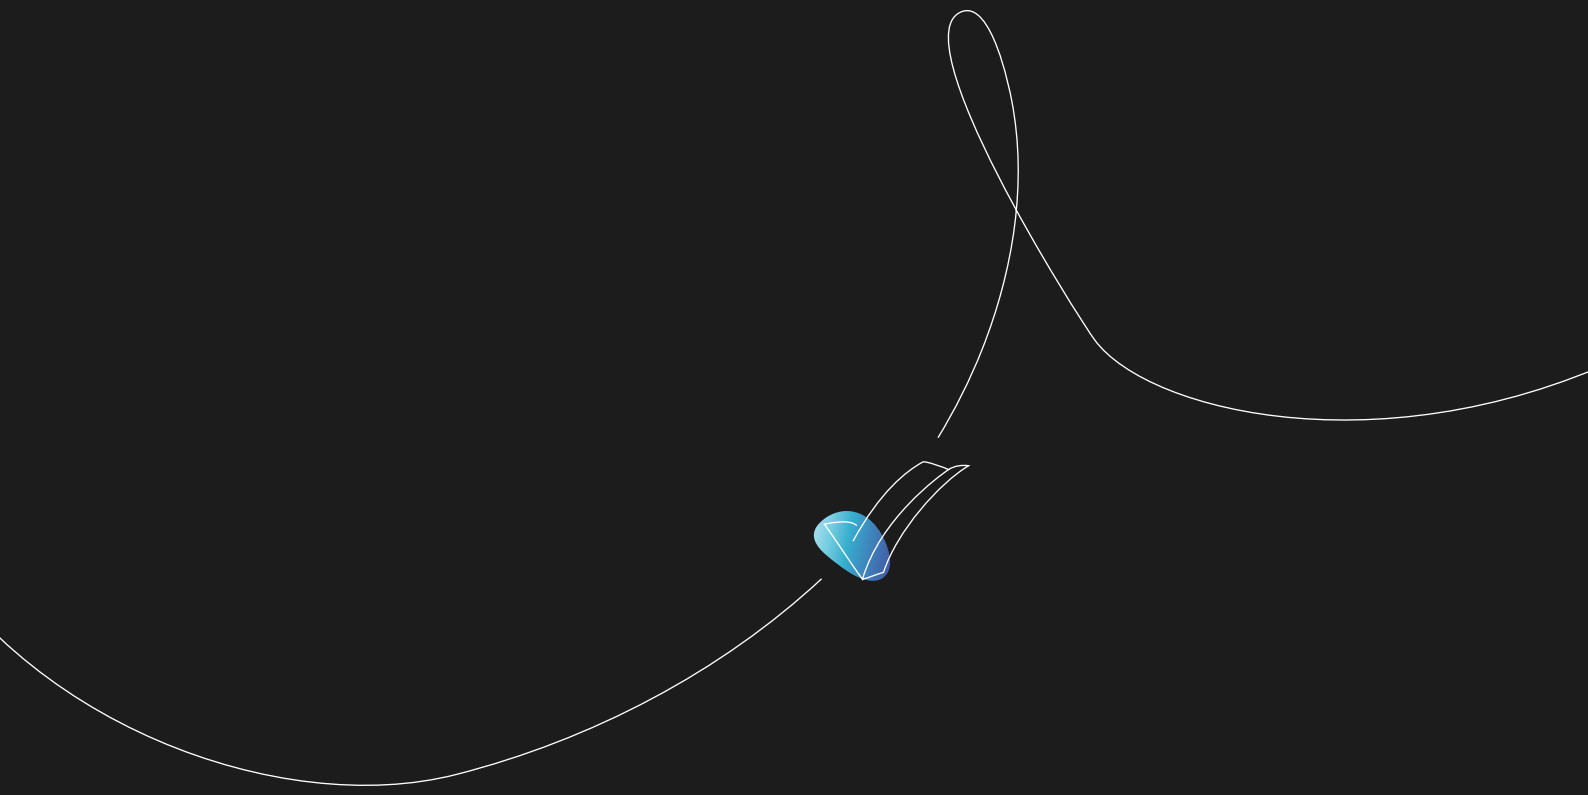

For more information

Contact us at: cisovillage@team8.vc | www.team8.vc

TEAM8™