Cloud 1 Solutions, Inc.
(844) 430-8820

AI SECURITY YOU CAN TRUST

The Complete Security Model Behind Iris

How a Governed AI Instance Protects Every Byte of Your Business Data

Presented by Cloud 1 Solutions, Inc.

What You Will Discover Inside

Application Security

Multifactor authentication on every account, with API keys, secrets, and passwords encrypted at rest.

Data Security

Modeled on NIST 800-171 — assessed and passed across all 110 control points, with ongoing penetration testing and vulnerability scanning.

User & Data Segregation

Policy-based access keeps HR, Finance, Operations, and every other department's data exactly where it belongs.

Business Continuity

Every system is backed up under a documented disaster recovery strategy — your data survives the unexpected.

INTRODUCTION

A Message to Business Owners About AI and Security

If you have heard the words “Artificial Intelligence” and felt a flicker of doubt, you are not alone. We hear it in nearly every conversation with a customer or prospect: concern that an AI tool could expose sensitive data to the wrong person, that an agent could delete something important with no way to undo it, or simply that “the AI” might do something nobody can see or explain.

Those concerns are valid, and they come from real incidents. But almost every horror story about AI going wrong traces back to the same root cause: an ungoverned agent, working independently, with no guardrails, no audit trail, and no boundary on what it is allowed to touch. That is not how Iris, our governed AI instance, is built.

This eBook is a plain-English walkthrough of the complete security model behind Iris — the standards it is measured against, the controls that protect your data, and the way access is restricted so that every department's information stays exactly where it belongs. Our goal is that by the end, you will understand not just that Iris is secure, but exactly how and why.

Our Promise:

No jargon, no hand-waving, and no “trust us.” Every claim in this eBook is backed by a specific control, standard, or practice — the same ones we would walk your own IT or security team through, line by line.

CHAPTER 1

Governed AI vs. an Independent Agent: The Real Risk

Before we walk through Iris's own controls, it helps to understand what most of the AI security fear is actually about. Most of the risk people associate with "AI" has nothing to do with AI as a technology — it comes from ungoverned agents operating with unrestricted access and no oversight. Give any agent, AI or otherwise, standing access to your systems with no guardrails, and the same list of risks appears every time.

Side by Side: What Changes with Governance

An AI Agent Working Independently	Our Governed Iris Instance
✕ Hackers or unauthorized actors can gain direct access to the agent	✓ Application security enforced by policies and roles — every action is permissioned
✕ Accidental or malicious deletion of production data — no guardrails to stop it	✓ Guardrails by design — no destructive actions are possible, ever
✕ API keys and secret credentials exposed in code, logs, or memory	✓ API keys, secrets, and passwords encrypted at the database level and at rest
✕ No role separation — HR staff or an HR agent can reach Financial data	✓ Policy-based access — HR never reaches Financial data unless explicitly allowed
✕ Rising, unpredictable token and compute costs as usage scales	✓ Predictable, centrally managed compute costs instead of per-agent sprawl
✕ No audit trail — no way to prove who (or what) touched what data, or when	✓ Full audit trail — every action logged, attributable, and reviewable
✕ No policy enforcement — the agent can do anything its access allows	✓ Data security modeled on NIST 800-171 controls, independently assessed
✕ Fragmented, siloed data with no historical record for reporting	✓ Centralized data lake captures history for trend and compliance reporting

This is the lens for everything that follows: Iris was built specifically to close every item in that left-hand column. The rest of this eBook walks through exactly how.

CHAPTER 2

Application Security: Every Action Is Permissioned

Application security governs how anyone — a human user or Iris itself — authenticates, and how the credentials behind every action are protected. This is the layer that stops an outside actor from simply logging in, and it is the layer that keeps your secrets secret even if a system is compromised.

What's In Place

- **Multifactor Authentication Required.** Every account, without exception, requires a second verification factor beyond a password before access is granted.
- **API Keys and Secret Codes Encrypted at Rest.** The credentials Iris uses to connect to your ERP, financial software, HR platform, and other systems are encrypted at the database level — never stored or logged in plain text.
- **Passwords Encrypted at Rest.** User passwords are encrypted using industry-standard methods, so even direct database access does not expose usable credentials.
- **All User Data Encrypted.** Beyond credentials, the data your team works with inside Iris — records, documents, and query results — is encrypted, both at rest and in transit.
- **Guardrails by Design.** Iris is architected so destructive actions — deleting or overwriting production data — are not something the AI is able to do. Iris reads, analyzes, and drafts; people retain the judgment calls, like approving a payment.

MFA	100%	0
REQUIRED ON EVERY ACCOUNT	OF SECRETS ENCRYPTED AT REST	DESTRUCTIVE ACTIONS PERMITTED

CHAPTER 3

Data Security: Verified, Tested, and Backed Up

Data security is about the standards Iris is measured against and the ongoing testing that proves those standards are holding, not just on day one, but continuously.

NIST 800-171

NIST 800-171 is a security control framework published by the National Institute of Standards and Technology for protecting sensitive information. It is a modern, evolving standard used well beyond government contracting — it is a recognized benchmark for what a serious data security program looks like. Iris has been assessed against all 110 individual control points defined in the standard, and passed every one of them.

Penetration Testing

Independent security professionals actively attempt to break into the environment — probing for exploitable weaknesses the way a real attacker would — on an ongoing basis, not as a one-time checkbox.

Vulnerability Scanning

Automated scanning continuously checks the infrastructure behind Iris for known vulnerabilities, so newly disclosed issues are identified and addressed quickly rather than discovered by accident.

Disaster Recovery and Backups

Every system behind Iris is backed up on a regular schedule, with a documented disaster recovery strategy in place. If the unexpected happens — hardware failure, a regional outage, or anything else — your data and your business are protected, and Iris is designed to be back online with minimal disruption.

110 / 110	24/7	DR
NIST 800-171 CONTROLS PASSED	VULNERABILITY SCANNING	BACKED UP WITH RECOVERY PLAN

CHAPTER 4

User Security: The Right Data, for the Right People, Every Time

This is the layer that answers the question we hear most often: “Can the wrong person — or the wrong department — see data they shouldn't?” With Iris, the answer is structural, not a promise.

How Access Actually Works

- **Users Are Added to Groups.** No individual user is granted standing access on their own.
- **Groups Are Managed by Policy.** A policy defines exactly what a group — Finance, HR, Operations, Executive, Sales, and so on — is permitted to see and do.
- **Policies Restrict Access Between Departments.** The classic example: HR staff, and Iris acting on their behalf, cannot see Financial data unless a policy explicitly grants it — and the same boundary applies between every other department and system you connect: Operations data stays out of Executive-only reporting unless permitted, Shipping data stays separate from Payroll, and so on.
- **This Extends System by System.** Because Iris connects across your ERP, financial software, HR platform, and document repositories, data separation isn't a single wall between two systems — it is enforced consistently across every connected data source, department, and role.
- **Every Access Decision Is Logged.** Whether access is granted or denied, the decision is recorded — producing a full audit trail that shows who (or what) touched which data, and when.

The Key Idea:

Individual users never have standing access on their own — it is always the policy that governs what any person, or Iris itself, is allowed to see. That single design choice is what keeps HR out of Finance, Operations out of Payroll, and every department's data inside its own boundary, automatically and without relying on anyone remembering to enforce it.

CHAPTER 5

The Complete Security Model, at a Glance

Security is never one control — it is what happens when every layer reinforces the others. Here is the full model in one place:

Application Security

- Multifactor authentication required on every account.
- API keys, secret codes, and passwords encrypted at rest.
- All user data encrypted, at rest and in transit.
- Guardrails by design — no destructive actions are possible.

Data Security

- Modeled on NIST 800-171 — assessed and passed across all 110 control points.
- Ongoing penetration testing by independent security professionals.
- Continuous automated vulnerability scanning.
- All systems backed up under a documented disaster recovery strategy.

User Security

- Users are added to groups; groups are managed by policy.
- Policies enforce separation between every department and system — HR from Finance, Operations from Executive, and beyond.
- Every access decision is logged, producing a complete, reviewable audit trail.

MFA REQUIRED EVERYWHERE	**110/110** NIST 800-171 CONTROLS	**100%** ACTIONS LOGGED & AUDITABLE
Encrypted DATA, KEYS & SECRETS AT REST	**DR + Backup** ON EVERY SYSTEM	**Isolated** YOUR DATA, YOURS ALONE

IN CLOSING

Security Built to Be Verified, Not Just Trusted

The most important thing to understand about Iris is that it is an isolated AI instance, built for your company alone. Your data is never pooled with another customer's data, never used to train a shared model, and never exposed to the open internet. Layered on top of that isolation is a security model spanning application security, data security, and user security — each one independently verifiable, not just promised.

If you have heard a concerning story about AI and data — deletion, exposure, or a tool with too much unchecked access — we would welcome the chance to show you, control by control, why that story does not describe Iris.

Ready to See Iris's Security Model First Hand?

Visit cloud1solutions.com to schedule your free AI security assessment.

Cloud 1 Solutions | Meet Iris | Your Company's Own AI Employee

(844) 430-8820