

FCI UNIFIED CONTROL — AI GOVERNANCE

Full Content Inspection. Total Control.

Where the **AI** policy meets the network.

CONTENTS

01	AI arrived before any policy did	2
02	Every AI policy comes down to four decisions	3
03	Turn policy into an enforcement rule	4
04	Example of an AI Governance rule	4
05	Scenarios: what the rules do	5
06	Coding assistants and FCI Frontier AI Defense	5

OVERVIEW

Companies already use AI at work. Employees draft emails with ChatGPT, summarize contracts with Gemini, and let coding assistants write software. Much of it happens on personal accounts, with company data pasted in, and none of it came through IT. The AI usage policy written in response is a document, and a document does not stop a prompt.

FCI Unified Control turns that policy into rules the platform enforces. Company traffic passes through Trinity Cyber’s Full Content Inspection™ (FCI) platform, which reads the content of each session as well as its destination. The security team writes the policy as an AI rule that names the allowed providers, accounts, and models, what a session may do, and the users and devices it applies to. The platform enforces that rule on every department and team while the traffic is still moving. Policy Logs record every match, so a rule shows what it would do before it blocks anything.

SECTION 01

AI arrived before any policy did

Someone in accounts payable pastes a vendor contract into a free chatbot for a quick summary. A developer installs a coding agent with full access to the code repository. Marketing turns on an AI plug-in inside the document editor. Most of it happens on personal accounts, in browser tabs, and inside applications with no website to block.

Verizon’s 2026 Data Breach Investigations Report counts 45 percent of employees as regular AI users on their corporate devices, up from 15 percent a year earlier; 67 percent of them use non-corporate accounts, and source code is what they submit most often.¹ IBM’s 2026 Cost of a Data Breach Report found shadow AI incidents (AI tools adopted without approval) at 43 percent of breached organizations, up from 20 percent, and 68 percent had no AI governance policy in place.²

A ban moves the work to phones and home networks, where nobody sees it. Blanket permission sends company data to whatever model, tool, or upload the service offers, under accounts nobody can audit. The right policy sits in between, and a web filter cannot enforce it, because a filter sees only the website. Whatever enforces the policy has to read the session itself: the account, the model, and what the session does.

BY THE NUMBERS

45%

of employees are regular AI users on corporate devices, up from 15 percent a year earlier.

67%

of them use non-corporate accounts. Source code is what they submit most often.

68%

of breached organizations had no AI governance policy in place.

BETTER WAY TO ANSWER AI USE

A BAN

Moves the work where nobody sees it

Phones and home networks. The policy holds on paper and fails on the network.

BLANKET PERMISSION

Sends company data anywhere the service offers

Any model, tool, or upload, under accounts nobody can audit.

[Inspect the session]

FCI Unified Control sits in between

The account, the model, and what the session does are conditions in one rule, enforced while the traffic is still moving.

SECTION 02

Every AI policy comes down to four decisions

Any AI usage policy, stripped down to what it decides, comes to four questions: which providers and accounts, which models and applications, what a session may do, and who may do it from which devices. Most policies answer them in a document that employees are asked to read and follow. FCI Unified Control turns each answer into a rule. Because the platform reads the content of every session, not just its destination, it can establish the account a request is signed in to, the model it calls, the application it comes from, and the capability it invokes, and whatever the platform can establish about a session, a rule can enforce. The four decisions below are the whole policy; the rest is enforcement.

01

THE POLICY DECIDES

Which providers and accounts

THE RULE ENFORCES

The rule allows or blocks each AI provider separately and allows it only on approved company accounts, so a personal ChatGPT subscription fails from the office network.

02

THE POLICY DECIDES

Which models and applications

THE RULE ENFORCES

The rule limits use to approved models and applications. A model released tomorrow, or a preview feature someone switches on, stays blocked until the security team approves it, and approving it takes one edit.

03

THE POLICY DECIDES

What the session may do

THE RULE ENFORCES

The rule blocks tool use, code execution, web search, or file upload, in any combination, and leaves ordinary chat alone.

04

THE POLICY DECIDES

Who, and from which devices

THE RULE ENFORCES

The same rule carries user identity and device posture (whether the device is managed, encrypted, and up to date), so allowing a provider also says who may use it and from which devices. A default-deny rule beneath it blocks every AI session the allow rule does not cover.

UNIFIED POLICY
All Unified Policy capabilities

An AI rule is a Unified Policy like any other. The same rule can also require drive encryption and a current operating system, a corporate network, a business-hours schedule, a file type on an upload, or any identity provider attribute, combined with ALL OF, ANY OF, and NOT.

SECTION 03

Turn policy into an enforcement rule

The first two steps record matches and block nothing, and the third turns the rule on.

1

LOG AND CONTINUE

Seeing what already happens

The platform sorts destinations into 79 domain categories, Generative AI among them, so it recognizes a session to an AI service as it happens, encrypted or not. Log and Continue is the action that records a match and blocks nothing. One Log and Continue rule turns a day of ordinary work into an inventory nobody had to compile, covering the services, teams, devices, accounts, and models actually in use.

2

POLICY BUILDER

Writing the rule

The four answers become conditions in the policy builder's AI Governance category, combined with user identity, device posture, and network conditions through ALL OF, ANY OF, and NOT. Instead of blocking a whole category, a rule can allow it with conditions attached. Because a rule names accounts, models, applications, and capabilities rather than specific tools, the rule covers a new tool the day it appears.

3

ALLOW OR BLOCK

Switching it on and observing the results

Changing the action from Log and Continue to Allow or Block turns the rule on, and the platform enforces it before any data reaches the provider. Policy Logs record every match down to the session, so the record of what a rule did is exact. No department or team beneath a rule set to Mandatory can loosen it.

SECTION 04

Example of an AI Governance rule

The security team writes one allow rule. It lets Data Science use approved models and applications on managed devices, under company accounts, with code execution off. A default-deny rule beneath it blocks every other AI session.

ALLOW RULE

Data Science: approved models and applications, managed devices, company accounts, code execution off.

Enforcement: Mandatory · Action: Allow

ALL OF

- User Identity → Groups → Exact → "Data Science"
- Device Posture → Managed Device → True
- AI Governance → Account → Company
- AI Governance → Application → Approved
- AI Governance → Model → Approved

NOT

- AI Governance → Capability → Code execution

Group, device, account, application, and model must all match.

ALLOW

Data Science

Approved coding assistant, managed laptop, company account. Every condition matches.

ALLOWED

Marketing

Same assistant on a personal subscription. Neither group nor account matches; default-deny takes over.

BLOCKED

Editor plug-in

Calls the same provider from a document editor. Application not approved; attempt logged.

BLOCKED

THE LEVEL YOU MANAGE

Security team

Writes two rules once. Every department and team beneath inherits them. Default-deny sits beneath the allow rule. What no rule covers is blocked.

DEFAULT-DENY

Every AI session the allow rule does not cover.

Sits beneath the allow rule. Policy Logs record each attempt.

BLOCK

SCENARIOS: WHAT THE RULES DO

● Allowed ● Blocked ● Rule change

SCENARIO	OUTCOME
A data scientist opens an approved coding assistant on a managed laptop under the company account.	● The platform allows the session, because group, device, account, application, and model all match.
Someone in Marketing tries the same assistant on a personal subscription.	● The platform blocks the session. Neither the group nor the account matches, so the default-deny rule takes over.
A plug-in nobody inventoried calls the same provider from inside a document editor.	● The platform blocks the session. It sees the plug-in as one more AI session, the application is not approved, and Policy Logs record the attempt.
Marketing is approved to use the assistant.	● Security adds those users to the group. Nothing else changes.

CODING ASSISTANTS

One rule covers the whole assistant

Coding assistants read and write source code directly. For assistants that run through their maker’s own service, such as GitHub Copilot, Cursor, and Windsurf, one rule allows or blocks the whole assistant. Assistants that call a model provider’s general API, such as Claude Code, OpenAI Codex, and Gemini CLI, announce which application they are in every request, and the same account, model, and capability conditions govern them.

THE OTHER SIDE

FCI Frontier AI Defense

Governing how a company uses AI is one side of the platform. The other is stopping attacks built with it. FCI Frontier AI Defense runs on the same Full Content Inspection™ engine, beneath every rule a company writes, and removes AI-generated malware, exploits, phishing, and command-and-control from live traffic before they reach the network. Countermeasures target adversary techniques rather than signatures, so a payload an AI produced an hour ago is caught the same way as one seen a thousand times, at line speed, with no agents to deploy.

SEE IT ON
LIVE TRAFFIC

Ready to see it on live traffic?

VISIT US AT trinitycyber.com

SOURCES

- 1. Verizon, *2026 Data Breach Investigations Report*, May 2026.
- 2. IBM and Ponemon Institute, *Cost of a Data Breach Report 2026*, July 2026.