

FCI UNIFIED CONTROL — AI GOVERNANCE FOR MSPS

Full Content Inspection. Total Control.

Multi-tenant **AI governance**, flexible by design.

CONTENTS

01	AI arrived before any policy did	2
02	Every AI policy comes down to four decisions	3
03	Three steps set the policy once for every client	4
04	Mandatory and Standard rules	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 the company’s managed service provider (MSP). Each client now needs an AI usage policy that something enforces, and MSPs need a way to set that policy once instead of once per client.

FCI Unified Control does both. Client traffic passes through Trinity Cyber’s Full Content Inspection™ (FCI) platform, which reads the content of each session as well as its destination. The platform is built for hierarchical multitenancy, which means each client is its own tenant and every tenant sits beneath its MSP in one hierarchy, with its sites and teams beneath it.

An MSP writes an AI rule once, at the top of that hierarchy. Every client beneath inherits it, and the platform enforces it while the traffic is still moving. No client can loosen a rule set to Mandatory, and every client can tune a rule set to Standard. Policy Logs show what every rule did, client by client.

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. The marketing team 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, and 67 percent of them use non-corporate accounts.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 Companies expect their MSP to handle the problem. In Kaseya’s 2026 State of the MSP report, 48 percent of MSPs name AI and automation as the top client need for 2026, and 61 percent say most or all of their clients turn to them for cybersecurity advice.³

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.

43%

of breached organizations had shadow AI incidents, up from 20 percent.

48%

of MSPs name AI and automation as the top client need for 2026.

SECTION 02

Every AI policy comes down to four decisions

Any AI usage policy, stripped down to what it decides, comes to four questions. FCI Unified Control turns each answer into a rule, and the MSP decides which answers hold for every client and which each client may tune.

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. A Mandatory rule fixes the provider list; each client adds its accounts in a Standard baseline.

02

THE POLICY DECIDES

Which models and applications

THE RULE ENFORCES

The rule limits use to approved models and applications. With the approved list set to Mandatory, a model released tomorrow, or a preview feature someone switches on, stays blocked at every client until the MSP 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. A Mandatory file-upload block holds for everyone, while a client can switch code execution on for the one team that needs it.

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.

BEYOND AI CONDITIONS

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

Three steps set the policy once for every client

Each step happens once, and every client inherits the result.

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 covers the whole client base, and a day of ordinary work becomes a per-client inventory of the services, teams, devices, accounts, and models each client actually uses.

2

POLICY BUILDER

Writing the rule once

The four answers become conditions in the policy builder's AI Governance category, beside device posture, network, file, user identity, and schedule, combined with ALL OF, ANY OF, and NOT. 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, with the path the rule took through the hierarchy to reach that client. When a client asks what its AI policy did last night, the answer is a filtered view and an export.

KEY POINT
An MSP writes an AI rule once, at the top of the hierarchy. Every client beneath inherits it, and the platform [enforces] it while the traffic is still moving.
Every rule can also be managed at the client level, so a single client's policy can be set and tuned on its own.

SECTION 04

Mandatory and Standard rules

Every rule an MSP writes is either Mandatory or Standard. A Mandatory rule applies to every client beneath it exactly as written and cannot be loosened at any lower level of the hierarchy. A Standard rule is a baseline: it applies to every client the same way, and the MSP can adjust it for an individual client, site, or team without touching the rest.

MANDATORY

In the scenarios that follow, the MSP blocks file upload to AI services and blocks any AI provider, model, or application outside the approved list. No client can loosen either rule.

STANDARD BASELINE

The MSP allows AI sessions on managed, encrypted devices under company accounts, with code execution off. Each client may tune it. A default-deny rule beneath both blocks every other AI session.

SCENARIOS: WHAT THE RULES DO

Allowed

Blocked

Mandatory

SCENARIO	OUTCOME
An engineer at Client A opens an approved coding assistant on a managed laptop under the company account.	Allowed The platform allows the session with code execution on, because Client A's engineering team overrode the baseline. Marketing at Client A runs the same session with code execution off.
Someone at Client B tries the same assistant on a personal subscription.	Blocked The platform blocks the session and logs it under Client B. The account does not match the baseline, so the default-deny rule applies.
A plug-in nobody inventoried calls an AI provider from inside a document editor at Client B.	Blocked The platform blocks the session. It sees the plug-in as one more AI session, and the application is not on the approved list.
Client B asks to lift the file-upload block for its finance team.	Mandatory The block is Mandatory, so Client B cannot change it. The Details view lists it as inherited, and the request goes to the MSP.
A Client A laptop's drive encryption is switched off in the middle of an AI session.	Blocked The platform checks posture continuously, so the session stops matching the baseline and the default-deny rule takes over.

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 clients use 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 an MSP writes, and removes AI-generated malware, exploits, phishing, and command-and-control from live traffic before they reach a client's 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 CLIENT
TRAFFIC

Ready to see it on client 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.
- 3. Kaseya, *2026 State of the MSP Report*, April 2026; survey of 1,061 MSPs, November 2025.