LayerX

Buyer's Guide for AI Usage Control

A Detailed Guide on How to Evaluate and Choose the Right Security Solution to Discover, Govern, and Control AI Usage Across the Enterprise

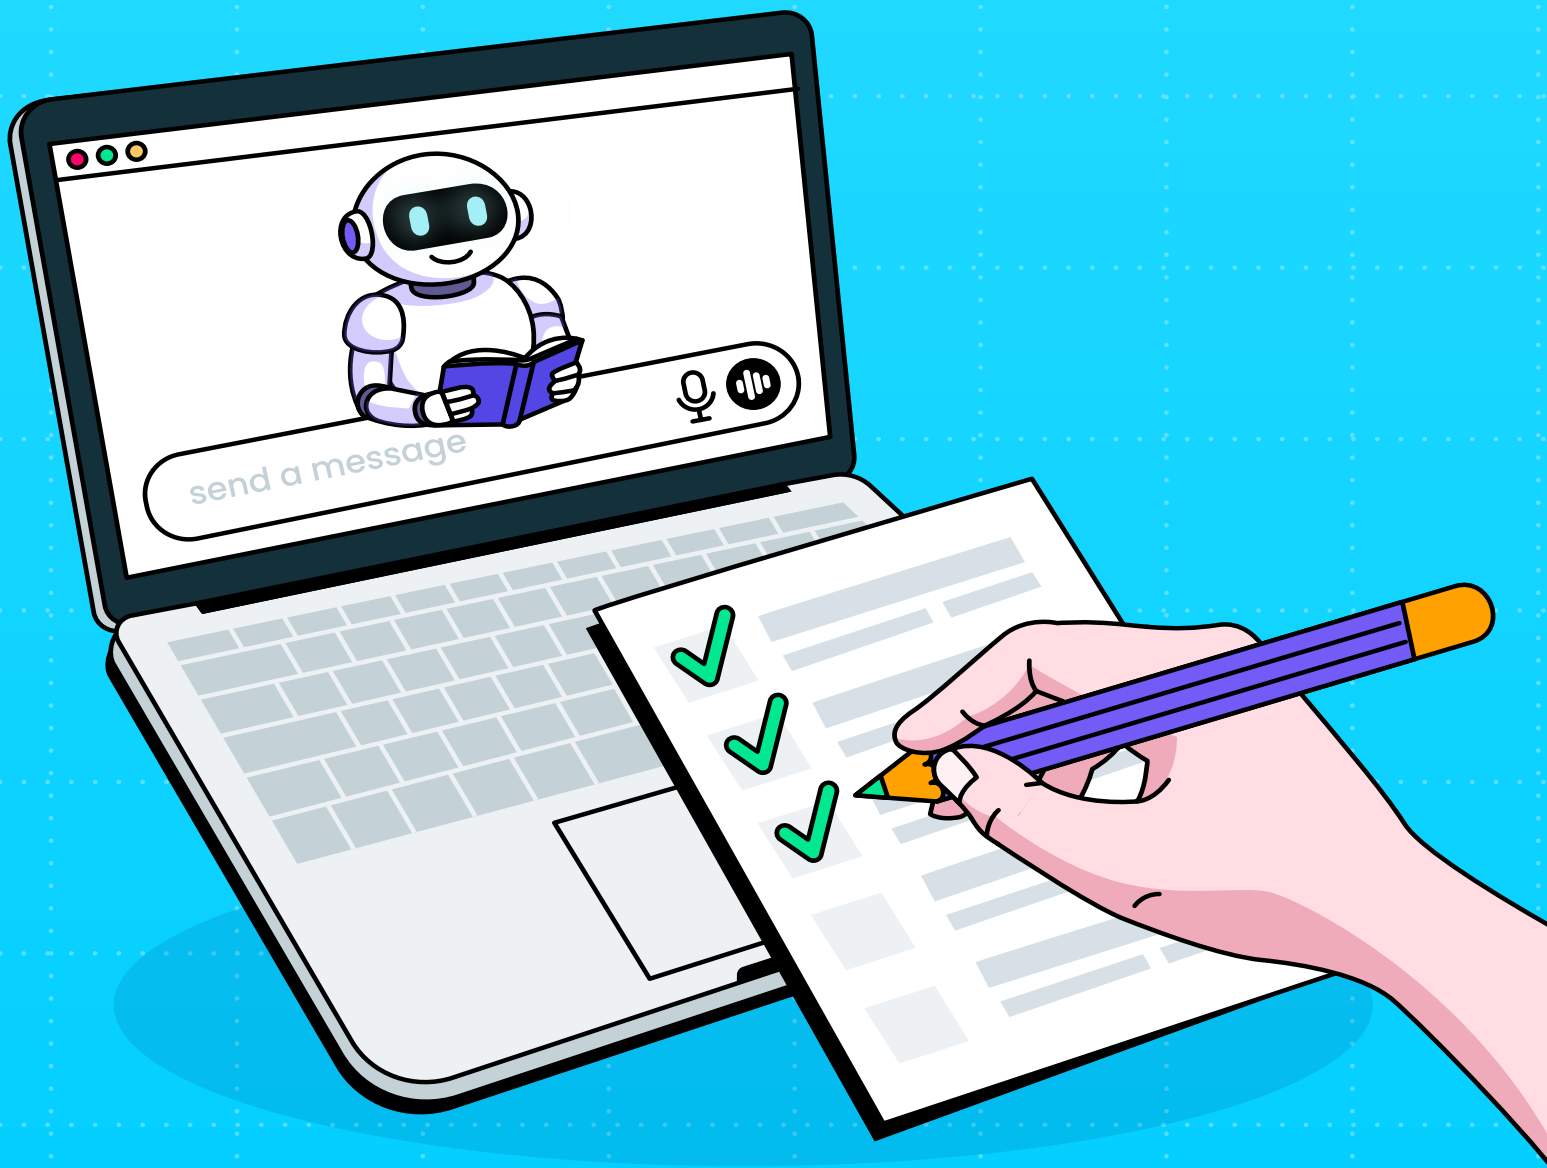

AI Usage Control (AUC) Tools Secure The Modern AI Workspace

Everybody's talking about AI but no one actually knows where it is. The rapid rate of AI adoption has led to a proliferation of AI tools, meaning that AI is everywhere but security managers can't see it anywhere.

This new 'shadow' AI economy is mostly hidden from sight and invisible to governance, exposing enterprises to data leakage, compliance violations, and AI security threats. And existing security stacks aren't really plugging the gap.

Almost overnight, AI has moved from a novelty to mission-critical infrastructure. What started as employees experimenting with standalone tools has evolved into AI capabilities woven directly into business processes across the enterprise. It has become a **way of doing business** across users, identities, devices, and services, often without clear visibility or control.

Most organizations have employee responsible use policies in place, but struggle to enforce them where AI is actually used. The result is fragmented oversight and limited accountability, even in supposedly tightly managed environments.

This gap has given rise to **AI Usage Control (AUC)** as a distinct security category.

As defined by Gartner, AI Usage Control is a foundational component of AI Trust, Risk, and Security Management (AI TRiSM). Its purpose is not simply to block AI or prevent isolated incidents, but to **discover how AI is being used, assess risk dynamically, enforce acceptable-use policies, and detect anomalous or unsafe behavior across the full spectrum of AI consumption.**

AI Usage Control represents a new control plane for the AI-powered enterprise, one that operates at the point of interaction rather than the perimeter and enables organizations to govern **how** AI is used, by **whom**, and under **what conditions**, without defaulting to blanket bans or reactive data-centric enforcement.

By combining comprehensive discovery, policy-driven enforcement, and anomaly detection across multiple AI usage vectors, AUC provides the structure security teams need to move from **reactive containment to proactive, scalable AI governance.** As AI continues to evolve and embed itself deeper into business workflows, this **governance-first approach** is becoming foundational to secure scalable AI adoption.

Understanding Your Options: What Is AI Usage Control (And What It's Not)

What Is AI Usage Control (AIUC)?

AI Usage Control (AUC) is a security and governance capability that helps organizations **discover, understand, and control how AI is used across the enterprise.** It governs AI interactions in real-time across users, apps, identities, and environments, enabling clear and enforceable standards for acceptable AI use as AI becomes embedded in daily workflows.

What AI Usage Control Includes

AI Usage Control enables security and governance teams to:

Discover AI usage across the enterprise
Continuously identify sanctioned and shadow AI across browsers, SaaS applications, extensions, endpoints, and emerging agentic workflows.

Govern acceptable AI use
Define and enforce AI usage policies based on risk tolerance, including which tools are allowed and under what users, identities, and conditions.

Enforce policies on AI interactions
Apply real-time controls on AI interactions such as prompts, uploads, copy/paste, logins, and extension behavior, without relying on static app lists or network-only controls.

Detect anomalous or risky behavior
Identify misuse and high-risk AI activity by analyzing interaction patterns, intent, and context, including data exposure, policy violations and prompt manipulation.

What AI Usage Control Is Not

AI Usage Control should not be confused with adjacent or legacy approaches:

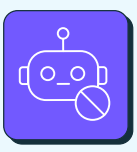

Not just blocking ChatGPT or individual AI tools
AIUC is not a deny list of known AI applications. It is designed to govern AI usage dynamically as new tools, features, and workflows emerge.

Not traditional data loss prevention (DLP)
While AIUC may inspect content, its primary purpose is governance and usage control at the point of interaction, not reactive data inspection after exposure has already occurred.

Not network-only visibility
AI usage frequently bypasses traditional network paths. AIUC is built to operate at the real-tim interaction points such as browsers, SaaS sessions, and endpoints.

Not model security or "AI you build"
AIUC does not secure internally developed models. It governs the risks of **AI you consume,** including third-party and embedded AI services used by employees.

AI Usage Control governs how AI is used, not just what data leaks.It represents a shift from tool- and data-centric defenses to interaction-centric governance that provides a scalable foundation for secure, compliant, and productive AI adoption.

Uncontrolled AI Usage in Enterprises Leads to Data Exposure and Compliance Violations

As AI becomes embedded across enterprise workflows, the primary risk is the **loss of governance over how AI systems are accessed, used, and acted upon.** Unlike traditional applications, AI tools operate through dynamic interactions, prompts, autonomous actions, and chained workflows that often bypass established controls. Without AI Usage Control, organizations lack the visibility and enforcement needed to manage this new operational layer.

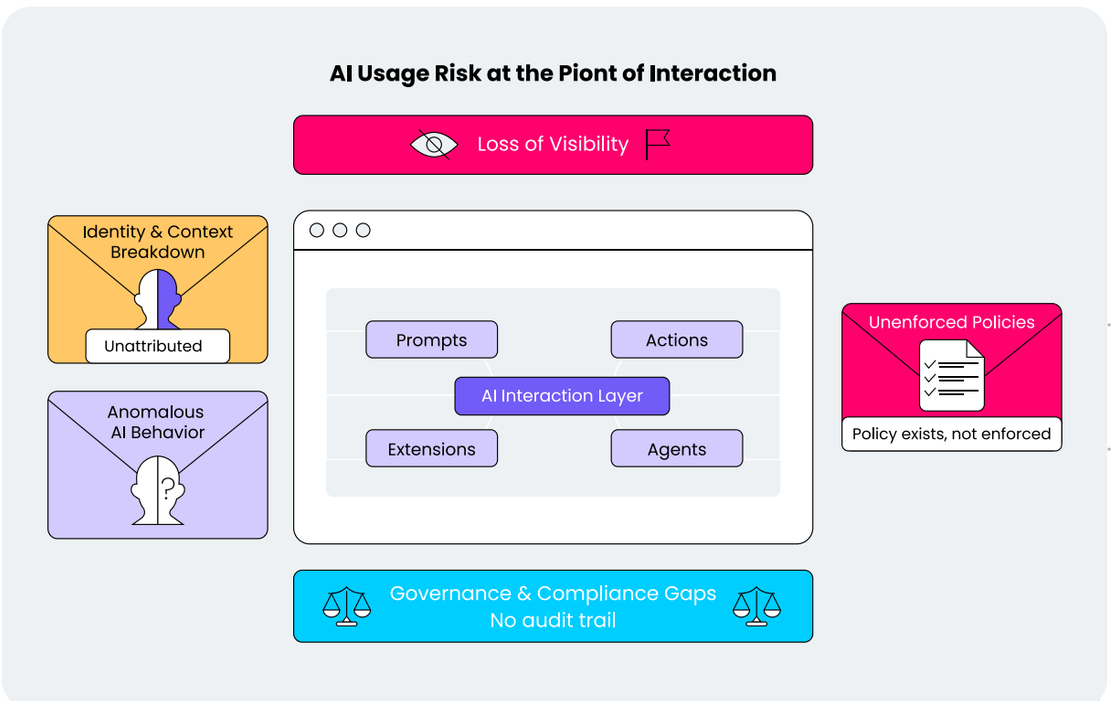

Failure to address this new layer of interaction leads to immediate and real-world risks to organizations:

#1 Loss of Visibility Into AI Usage and Behavior

AI tools are frequently consumed through browsers, embedded SaaS features, extensions, and automated agents, many of which fall outside traditional discovery mechanisms. When organizations cannot reliably identify which AI tools are in use, who is using them, and how they are being accessed, governance becomes impossible.

Key Risks

- Incomplete or inaccurate inventory of AI tools and features
- Inability to attribute AI activity to users, roles, or identities
- Blind spots created by browser-based, extension-driven, or embedded AI usage

#2 Unenforced or Inconsistent AI Usage Policies

Many organizations define acceptable AI use in policy documents but lack the technical means to enforce those rules in practice. As a result, AI usage policies remain aspirational rather than operational, applied inconsistently or not at all across tools and environments.

Key Risks

- Policies that cannot be enforced at the moment of interaction
- Inconsistent controls across sanctioned and unsanctioned AI tools
- Overreliance on manual processes or post-incident reviews

#3 Identity and Context Breakdown

AI interactions frequently occur outside enterprise identities, using personal accounts, unmanaged sessions, or automated agents. This breaks identity-based controls and removes the contextual signals needed to govern AI usage effectively.

Key Risks

- AI activity performed under personal or anonymous identities
- Loss of audit trails linking actions to accountable users
- Inability to apply role, device, or session-based policies

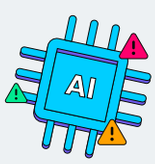

#4 Misuse, Abuse, and Anomalous AI Behavior

Without real-time inspection and behavioral analysis, organizations struggle to detect when AI is being misused intentionally or unintentionally. This includes excessive data exposure, policy circumvention, prompt manipulation, or automated actions that exceed the intended scope.

Key Risks

- Undetected misuse of AI tools for unintended or risky purposes
- Anomalous usage patterns that signal emerging threats or abuse
- Automated or agentic actions operating beyond defined guardrails

#5 Governance and Compliance Gaps

AI usage that cannot be discovered, monitored, or controlled introduces compliance and audit challenges, even in the absence of a confirmed data breach. Organizations may be unable to demonstrate adherence to internal policies or external regulatory requirements governing acceptable technology use.

Key Risks

- Inability to prove policy compliance or control effectiveness
- Gaps in auditability and reporting for AI-related activity
- Increased exposure during regulatory reviews, investigations, or incidents

Real-World Implications: What's at Stake

When AI usage goes ungoverned, organizations risk systemic loss of control. Employees and automated agents may take actions on behalf of the business that are misaligned with policy, compliance obligations, or risk tolerance, without clear attribution or visibility. Over time, this erodes trust in AI initiatives, forces reactive shutdowns, and slows adoption of legitimate AI-driven innovation.

Beyond operational risk, unmanaged AI usage can expose organizations to regulatory scrutiny, legal challenges, and reputational damage, often without a clear trigger event or breach. The result is fragmented governance, increased uncertainty, and stalled momentum in the organization's AI strategy.

Security vs. Productivity: A Bogus Dilemma

For years, security leaders have been forced into a false binary: lock everything down or let productivity run wild. AI tools have only intensified that tension. With employees generating content faster, automating tasks, and coding with AI-powered copilots, the productivity gains are undeniable. But so are the risks.

Many CISOs are rightly concerned: how do you secure AI usage without becoming the department of "no"?

The answer isn't in blanket bans or restrictive legacy policies. Blocking ChatGPT might check a compliance box, but it also sends users straight to their personal laptops, VPN-free, using unmonitored AI tools. That's not control. That's creating a shadow AI problem by design.

What's needed is nuance. The ability to say:

"Yes, you can use AI to automate that task but only in a sanctioned session, under a corporate identity, and without exposing sensitive IP."

This balance is only possible with context-aware AI usage and control that operates in real-time, at the exact moment of user interaction. It allows organizations to empower their teams with AI-driven efficiency, while ensuring sensitive data never leaves the guardrails.

Common Pitfalls When Evaluating AI Usage Control Solutions

As AI usage accelerates, many organizations rush to “add AI controls” by extending existing tools or adopting surface-level capabilities. In practice, this often leads to a false sense of security. The following pitfalls are common when evaluating AI Usage Control platforms.

Treating AI Usage Control as a Checkbox Feature

Some platforms position AIUC as a minor add-on to SSE, CASB, or DLP. These capabilities may detect access to a handful of known AI tools but lack the depth required to govern real AI interactions. AI Usage Control is not a feature, it is a dedicated control plane that must span discovery, context, and enforcement.

Relying Solely on Network Visibility

Network-based controls miss large portions of modern AI usage, including browser-native interactions, embedded SaaS AI features, personal accounts, and direct-to-cloud traffic. When visibility stops at the network layer, the most critical AI interactions remain invisible and uncontrolled.

Blocking AI Instead of Governing It

Blanket bans and binary controls are tempting but ineffective. Blocking AI drives users to personal devices, unmanaged browsers, and shadow tools, increasing risk rather than reducing it. Effective AIUC enables safe, role-aware usage through granular controls, not all-or-nothing enforcement.

Ignoring Browser Extensions and AI Brokers

AI-powered browser extensions and AI-enabled browsers often act as silent intermediaries, routing data to external models without clear visibility. Ignoring this layer leaves a major governance gap, as extensions can access prompts, page content, and sensitive data outside traditional controls.

Over-Indexing on “AI Detection” Without Enforcement

Visibility alone does not equal control. Many solutions stop at detecting AI usage or classifying data, with no ability to intervene at interaction time. Without real-time, context-aware enforcement, organizations can observe AI risk but cannot prevent it.

Avoiding these pitfalls requires rethinking AI security from the interaction layer outward. Effective AI Usage Control is purpose-built to govern how AI is used across tools, identities, and contexts, without slowing the business or forcing workarounds.

Consideration Criteria: Enumerating Requirements for AI Usage Control Solutions

How to Assess AI Usage Control (AIUC) Platforms

Evaluating AI Usage Control solutions requires a different lens than traditional security tools. Because AI risk emerges from how tools are used across users, identities, contexts, and interaction points, buyers must assess platforms based on their ability to deliver continuous visibility, governance, and control at the moment of AI interaction.

The following buyer-centric pillars outline the foundational requirements of an effective AI Usage Control strategy.

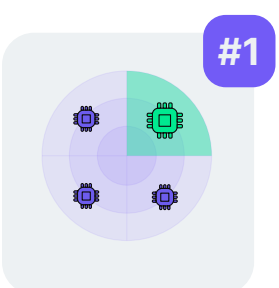

AI Discovery & Coverage

AI Usage Control begins with discovery. Without knowing which AI tools are being used, by whom, and under which identities, organizations cannot define or enforce meaningful control.

Objective:
Establish complete visibility into where and how AI is used across the organization.

Why It Matters:

AI adoption is highly decentralized. Employees use standalone AI tools, embedded SaaS copilots, browser extensions, and automated agents often without IT involvement. Most organizations underestimate the scope of AI usage, leading to blind spots where governance and accountability break down.

Key Requirements:

- **AI Tool Discovery**
  - Automatically detect browser-based AI tools (e.g., ChatGPT, Claude, Gemini, Perplexity)
  - Identify embedded AI features inside SaaS applications (e.g., email, CRM, collaboration tools)
  - Discover AI-powered browser extensions and AI browsers
  - Identify local desktop AI applications (e.g., ChatGPT, Copilot, etc.)
  - Detect API- and agent-driven AI workflows where applicable
  - Automatically detect new AI tools and usage patterns as they emerge
- **User & Identity Mapping**
  - Attribute AI usage to specific users
  - Distinguish between corporate and personal identities, not just the email address, but also whether the data is used for model training
  - Track usage of SSO across the organization
  - Track authenticated and unauthenticated AI interactions
- **Sanctioned vs. Shadow AI Identification**
  - Classify approved, tolerated, and unsanctioned AI tools
  - Identify Bring Your Own AI (BYOAI) usage introduced without IT awareness
- **Conversation Tracking:**
  - Discover all interactions with top GenAI platforms, including the prompt and the AI models response.
  - Identify the type of data being entered (e.g., source code, financial info, PII) and whether it includes sensitive or regulated information.

Outcome:
A complete, continuously updated inventory of AI usage across users, identities, tools, and environments forming the foundation for governance, policy enforcement, and risk management.

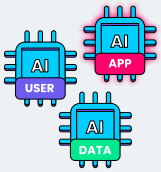

AI Risk Assessment & Contextual Awareness

AIUC is not about treating all AI equally. Different tools, access methods, and usage patterns carry different levels of risk and require differentiated governance.

Objective:
Prioritize risk based on the characteristics and behavior of AI tools, not assumptions.

Why It Matters:

Static allowlists and manual reviews cannot keep up with the pace of AI innovation. Without continuous risk assessment, organizations either over-restrict AI usage or leave high-risk behaviors unchecked. This helps organizations focus governance efforts where risk is highest, without stifling low-risk innovation.

Key Requirements:

- **AI Tool Risk Scoring**
  - Evaluate tools based on training policies (e.g., data reuse, model improvement)
  - Assess data retention and privacy posture
  - Account for access methods (browser, extension, API, embedded)
- **Behavioral Risk Signals**
  - Identify risky usage patterns based on actions, frequency, and context
  - Detect deviations from expected or approved usage
- **Continuous & Automated Assessment**
  - Automatically update risk scores as tools evolve
  - Adapt to changes in vendor policies, features, or usage behavior
- **Context-Aware Risk Evaluation**
  - Factor in user role, identity type, device posture, and session context

Outcome:
A dynamic understanding of AI risk that enables prioritization, informed policy design, and adaptive enforcement without relying on static assumptions.

Policy-Based AI Usage Governance

AI Usage Control translates governance intent into enforceable policies that guide how AI can be used safely and productively.

Objective:
Define and enforce acceptable AI usage across the organization.

Why It Matters:

Binary allow-or-block approaches fail in real-world environments. Organizations need flexible policies that reflect business roles, risk tolerance, and evolving AI use cases.

Key Requirements:

- **Granular Policy Definition:** Scope policies by -
  - User role and group
  - Corporate vs. personal identity
  - AI tool category or specific tool
  - Action type (prompting, uploads, copy/paste, responses)
  - Device type (govern AI usage across managed and unmanaged environments)
  - Session context (e.g., incognito browsing or unmanaged SaaS logins)
  - Geolocation/IP (e.g., restrict use from untrusted countries or networks)
  - Cross-domain activity (e.g., Salesforce → WeTransfer.com)
  - Cross-identity activity (e.g., corp → non-corp)
- **Flexible Enforcement Controls:** Match the enforcement level based on context.
  - **Allow:** Permit interaction if low risk.
  - **Monitor:** Record activity for audit without interruption.
  - **Warn:** Alert users in real time if their action may lead to a violation.
  - **Bypass with Justification:** Allow exceptions for high-trust users with policy-aware approvals and justification capture.
  - **Block:** Fully prevent risky actions or tool access.
  - **Redact:** Automatically mask or remove sensitive data (e.g., tokenize PII or obfuscate source code).
  - Apply rules uniformly across AI consumption paths

This tiered approach helps avoid productivity roadblocks while still safeguarding AI usage across the organization.

Outcome:
Clear, enforceable AI usage guardrails that align security, compliance, and productivity without forcing blanket bans.

#4

Real-Time Enforcement at Interaction Time

AIUC must operate at interaction time, when users or agents engage with AI tools, not after that.

Objective:
Control AI risk at the moment it occurs.

Why It Matters:

AI risk emerges during prompts, uploads, and responses. Post-event detection is too late to prevent misuse, policy violations, or unintended consequences.

Key Requirements:

- **Interaction-Level Inspection**
  - Inspect prompts, inputs, uploads, and responses in real-time
- **Contextual Detection**
  - Identify misuse, policy violations, or anomalous behavior
  - Account for intent, identity, and session context
- **Non-Disruptive Enforcement**
  - Enforce controls without breaking workflows
  - Avoid forcing users into workarounds
- **User Guidance**
  - Provide in-the-moment feedback when actions violate policy
  - Explain why an action was blocked or warned and offer links to approved AI tools or usage guidelines
  - Encourage compliant behavior rather than silent blocking

Outcome:
Immediate, context-aware enforcement that prevents risk while preserving user experience and productivity.

#5

Monitoring, Alerting & Auditability

AI governance requires visibility, traceability, and evidence especially as usage expands.

Objective:
Maintain accountability and support AI governance at scale.

Why It Matters:

Without monitoring and auditability, AI usage becomes impossible to govern across teams, tools, and time.

Key Requirements:

- **Comprehensive Activity Logging:** Record AI interactions, policy decisions, and enforcement actions
- **Anomaly Detection:** Identify unusual or risky usage patterns
- **Alerting & Reporting:** Generate actionable alerts for security teams and provide compliance-ready reports
- **Audit Trails:** Support investigations, audits, and regulatory reviews

Together, these capabilities enable organizations to move from reactive AI risk management to proactive, policy-driven governance, ensuring AI adoption remains secure, compliant, and scalable.

Outcome:
Continuous oversight and accountability that turns AI usage into a governed, auditable enterprise activity.

#6

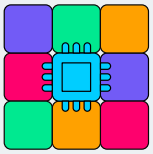

Architecture Fit & Operational Readiness

Objective:

Ensure the solution integrates seamlessly with your existing environment and works where AI usage actually occurs without adding operational burden.

Why It Matters:

Network-only or agent-heavy architectures struggle to cover AI usage that happens directly in browsers, SaaS apps, extensions and unmanaged environments. Your AIUC solution must work where the risk lives without disrupting users or requiring infrastructure overhauls.

Key Requirements:

- **Interaction-Level Coverage:** Visibility into AI usage where users interact with tools
- **Agentless or Low-Friction Deployment:** Minimal endpoint or infrastructure changes
  - Broad Environment Support
  - Works across managed and unmanaged devices
  - Supports multiple browsers and AI consumption paths
  - Covers all leading AI native apps via PWA
- **Security Stack Integration:** Seamless integration into your current stack with broad surface-level visibility and protection without any friction or compromise

Outcome:
Fast time-to-value with durable coverage without disrupting users or IT operations.

Deployment & Management

Objective:
Minimize operational overhead and ensure easy deployment across your environment.

Why It Matters:

Security solutions should protect, not burden. If a solution is hard to deploy or manage, it won't scale. You need instant time-to-value, tamper-resistance, and centralized control, especially in today's decentralized, browser-first environments.

Key Requirements:

- **Agentless Rollout:** No device-level installation needed; deploys instantly via existing browser infrastructure or MDM.
- **Centralized Policy Management:** Create, apply, and update policies from a single console across users, browsers, and devices.

- **Tamper-Proof Controls:** Resistant to user interference, uninstallation, or evasion, even in unmanaged or BYOD environments.
- **No Admin Overhead:** Minimal configuration effort. Can be integrated with SSO and directory services for easy onboarding.
- **Fast Time-to-Value:** Full coverage within hours, not weeks.

Outcome:
Low-lift deployment with high-impact protection, giving security teams more control without additional complexity.

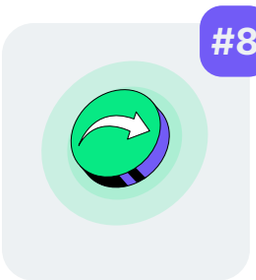

#8 Futureproofing & Adaptability

AI usage patterns, tools, and risks are changing rapidly. AIUC solutions must evolve just as fast.

Objective:
Ensure long-term governance as AI evolves.

Why It Matters:

Point-in-time controls quickly become obsolete in an AI landscape that changes weekly.

Key Requirements:

- **Tool-Agnostic Design:** Works across known and unknown AI tools
- **Adaptive Risk Modeling:** Evolves with new behaviors, agents, and usage patterns
- **Continuous Platform Expansion:** Regular support for new AI tools, extensions, workflows, and regulations
- **AI-Driven Detection:** Leverages machine learning to evolve with the threat landscape without relying solely on static rules.
- **Forward-Looking Governance:** Prepared for agentic AI, autonomous workflows, and emerging AI interaction models

Outcome:
A governance layer that scales with AI adoption, protecting today's usage and tomorrow's innovation.

Evaluation Checklist

AI Discovery & Coverage: Can we see all AI usage across the organization?

- ☐ Discovers browser-based AI tools and AI-native web apps
- ☐ Identifies embedded AI features within SaaS platforms
- ☐ Detects AI-powered browser extensions and AI browsers
- ☐ Maps AI usage to users and distinguishes corporate vs. personal identities
- ☐ Identifies if data used in an account, personal or business, is being used for model training
- ☐ Classifies sanctioned, tolerated, and shadow AI (BYOAI)

AI Risk Assessment & Context Awareness: Can we prioritize AI risk dynamically?

- ☐ Risk-scores AI tools based on training, retention, and privacy posture
- ☐ Accounts for access method (browser, extension, embedded, API)
- ☐ Detects risky or anomalous usage patterns
- ☐ Continuously updates risk as tools and behaviors evolve
- ☐ Incorporates user role, identity, device, and session context

Policy-Based AI Usage Governance: Can we define AI usage with precision?

- ☐ Policies scoped by user role and identity type
- ☐ Controls by AI tool category or specific service
- ☐ Enforcement by action type (prompt, upload, copy/paste, response)
- ☐ Governance across managed and unmanaged environments
- ☐ Consistent policy application across all AI consumption paths

Real-Time Enforcement at Interaction: Can risk be controlled the moment it occurs?

- ☐ Real-time inspection of prompts, uploads, and responses
- ☐ Context-aware detection of misuse and policy violations
- ☐ Graduated enforcement (allow, warn, guide, block)
- ☐ Non-disruptive enforcement that preserves workflows
- ☐ In-the-moment user guidance and education

Monitoring, Alerting & Auditability: Can AI usage be governed and audited at scale?

- ☐ Comprehensive logging of AI interactions and actions
- ☐ Anomaly detection for unusual AI behavior
- ☐ Actionable alerts for security teams
- ☐ Compliance- and audit-ready reporting
- ☐ Clear audit trails for investigations and reviews

Architecture Fit & Operational Readiness: Does the solution work where AI risk lives?

- ☐ Interaction-level visibility where users engage with AI
- ☐ Agentless or low-friction deployment
- ☐ Coverage across managed and unmanaged devices
- ☐ Support for all major browsers and AI-native apps
- ☐ Clean integration with existing security tools

Deployment & Management: Can this scale without operational burden?

- ☐ Fast rollout via browser infrastructure or MDM
- ☐ Centralized policy and configuration management
- ☐ Tamper- and bypass-resistant controls
- ☐ Minimal admin effort and ongoing maintenance
- ☐ Rapid time-to-value (hours, not weeks)

Futureproofing & Adaptability: Will this solution evolve with AI?

- ☐ Tool-agnostic coverage for known and unknown AI tools
- ☐ Adaptive risk modeling for new behaviors and agents
- ☐ Continuous support for new AI platforms and extensions
- ☐ AI-driven detection beyond static rules
- ☐ Readiness for agentic and autonomous AI workflows

AI Usage Control and the Future of AI Governance

As AI adoption accelerates, governance models must evolve beyond static controls and reactive risk management. AI Usage Control is emerging as a foundational pillar of modern AI governance, one that aligns closely with Gartner's AI TRiSM framework by addressing visibility, policy enforcement, and ongoing oversight across how AI is actually used in the enterprise.

The future of AI will be defined by distributed, agentic, and increasingly autonomous interactions. Agentic AI systems, MCP-style orchestration, and AI-powered browsers introduce workflows that act across tools, identities, and sessions with minimal human intervention. In this environment, one-time approvals, static allowlists, and tool-specific controls quickly become obsolete. Governance must shift from "which tools are allowed" to "how AI is allowed to operate, under what conditions, and with what intent."

AI Usage Control enables this shift through continuous discovery and real-time governance. Rather than relying on periodic reviews or manual risk assessments, AIUC provides ongoing insight into new tools, new usage patterns, and new interaction models as they emerge. This allows organizations to adapt policies dynamically, keeping pace with innovation without repeatedly resetting governance frameworks.

Critically, AIUC makes it possible to govern AI at scale without stifling innovation. By applying contextual, graduated controls at the moment of interaction, organizations can enable responsible AI usage across roles and teams while preventing misuse, policy violations, and unintended risk. Security becomes an enabler of AI adoption, not a barrier.

In this way, AI Usage Control should be viewed not as a point solution, but as a long-term control plane for AI governance, one that complements existing security investments while providing the visibility, context, and enforcement required for the next generation of AI-powered work.

The Bottom Line: AI Usage Control Solution Maintains the Balance Between AI Security and Productivity

AI adoption across the enterprise is inevitable and accelerating, but the risk no longer lies in whether employees use AI, it lies in how they use it. Traditional security controls were never designed to govern real-time AI interactions across users, identities, tools, and contexts, leaving organizations exposed as AI becomes embedded into everyday work.

AI Usage Control is no longer optional. It is a foundational governance layer for enabling AI safely at scale. The right AIUC platform goes beyond visibility, delivering continuous discovery, contextual risk awareness, real-time interaction control, and policy-driven enforcement without disrupting productivity or innovation.

Use this guide and checklist to evaluate solutions through an AIUC lens. Prioritize platforms that are built for interaction-level control, identity-aware governance, and continuous adaptation as AI tools and usage models evolve.

The organizations that succeed with AI won't be the ones that block it, but the ones that govern it effectively. Start building a secure, scalable AI future with LayerX.

How LayerX Can Help

Visibility Into All AI

Detect all AI tools and applications in use and gain full visibility into all user activity within them

Control Access to AI Tools

Restrict usage of shadow AI apps and secure access to sanctioned AI apps using corporate accounts

Prevent AI Data Leakage

Enforce last-mile AI security guardrails to stop users from sharing sensitive data with GenAI tools

Control AI Inputs and Output

Monitor and control data flowing into and out of AI tools to prevent sensitive or confidential data from being entered or exposed

LayerX is the only AI usage control platform that lets you control every prompt, agent action, and data exchange, across any channel, without changing your network architecture or disrupting user experience.

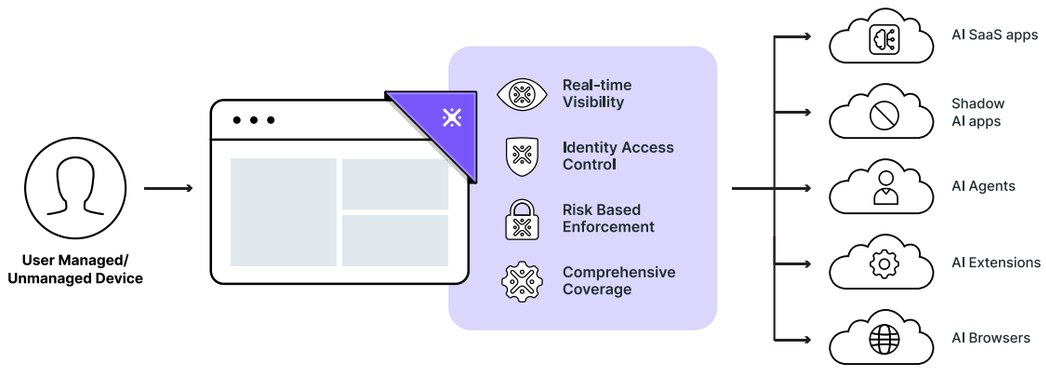

LayerX is a last-mile AI security platform that delivers unmatched visibility and control over every AI interaction, exactly where it happens. We support wide range of use cases:

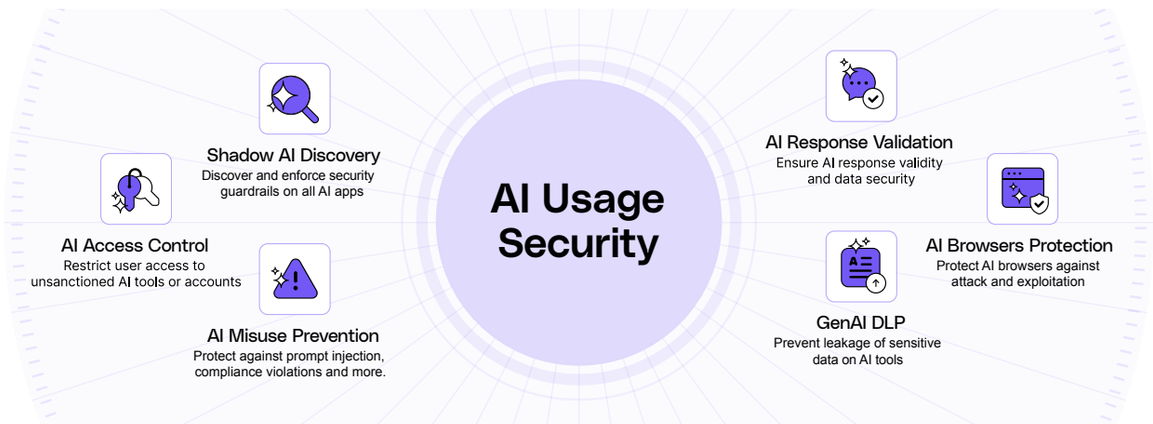

To learn more about how LayerX can help you prevent browser based data leakage, go to www.layerxsecurity.com and schedule a demo today!