

BEYOND THE CHECKLIST:


Modernizing compliance in an AI-driven world





- 3 Executive overview
- 5 Changes in an AI-powered world
- 9 Maintaining compliance with AI
- 15 Stay a step ahead with Cloudflare


GenAI changes everything



65% of survey respondents say their organizations regularly use GenAI ([McKinsey](#))

800+
new regulations

> 800 measures for regulating AI are under consideration in **> 60 countries** worldwide ([Boston Consulting Group](#))



71% of countries have passed legislation to secure the protection of data and privacy ([UNCTAD](#))

Executive overview



Generative artificial intelligence ([GenAI](#)) transforms virtually every aspect of modern business – in a good way for the most part – but raises equally impactful new risks and challenges. Progress always does, but rarely under the same watchful eye of regulators and consumers worldwide.


For CISOs, the growing development and use of large language models ([LLMs](#)) blurs the lines — and the geographic borders — between industry standards, individual countries' localization laws, and companies' own internal policies for guarding sensitive data. To keep pace with new regulations governing use of tools like ChatGPT, organizations need flexible, modern [compliance](#) strategies that promote innovation, [protect data](#), and play by the rules – even when the rules overlap or conflict.




Justyna Kucharczak
Sr. Solutions Marketing
Manager, Data




David Liu
Sr. Product Marketing
Manager, AI



GenAI stands to enable the automation of up to 70% of business activities across virtually all occupations between now and 2030 ([McKinsey](#))



Only 28% of executives surveyed say their organization is fully prepared for new regulation ([Boston Consulting Group](#))



Eclipsing the old 'checkbox' mentality



Instead of thinking of compliance as a 'necessary evil,' or growing list of boxes to check during audits, modern CISOs view compliance programs as a way to streamline workflows and make their business run faster — like GenAI tools themselves.

At a minimum, flexible, forward-looking cybersecurity and governance strategies:



Secure companies' development and consumption of AI applications and the output they produce



Balance centralized use and local control of data worldwide



Promote innovation while side-stepping risk



Fulfill the intent behind all the new rules — protecting privileged data to protect the reputations of a business and its customers

This ebook explores the challenges CISOs face in modernizing their governance strategies and outlines new goals and best practices for making compliance a foundation for growth — an opportunity versus more boxes to check — as they weave AI deeper into the fabric of business.

Changes in an AI-powered world




Everything has changed. New risk, new regulations, new requirements for demonstrating you can avoid the new threats and follow the rules. Regulators and individual countries continue to move fast but CISOs can't afford to wait for laws to take effect before expanding their compliance programs.

Evolving usage brings new risk

GenAI brings the same game-changing economies of scale to cyberattackers as it does to enterprises, only attackers don't abide by the rules. CISOs must find ways to bolster existing defenses against AI-led attacks and prevent LLMs from becoming new vectors for inbound and outbound threats.

GenAI adds foreseeable risk from:

- Tools introducing compromised code
- Algorithms conveying a noticeable bias
- Sensitive data being inputted into LLMs by users that could train LLMs
- LLMs producing inaccurate output (unsubstantiated news or marketing claims, misleading medical misdiagnosis, etc.) that subject businesses to fines, financial loss, and damage to their brand



71% of IT leaders are concerned that generative AI will introduce new security threats to their data
([Salesforce](#))

Modern compliance strategies must prevent new tools from generating disinformation, privacy violations, discriminatory content, and third-party risk without hindering innovation. Two pivotal usage trends make it harder to balance data sovereignty and privacy as AI gains popularity:



Companies exposing more sensitive data to AI:

In the early stages of implementation, most organizations limited GenAI application to less sensitive data used in marketing campaigns, sales due diligence, and processing monitoring tool alerts. Once they see what AI can do, more businesses reveal different types of data to new applications, including sensitive data such as invoices, recruiting data used in HR, employee policy handbooks, and other data used in core business analytics.



AI acting on its own:

For example, accounting departments might prompt a GenAI tool to search and display a company's projects or billable hours for the current month. They might then instruct the tool to send out invoices or reminders as customer due dates approach. Now, instead of just training on data, AI initiates an action that needs to be tracked, documented, and above all kept private.

Evolving risk brings new rules: CISOs top 3 challenges



The convergence of technology, global data sharing, and evolving regulatory frameworks creates complex challenges for organizations maintaining data compliance.

CISOs face three overarching challenges in modernizing their strategies to safeguard GenAI:

1

Keeping up with new regulations as their threat landscape evolves.

GenAI's value and risk continue to mature against a backdrop of economic and geopolitical tensions and fast-moving regulatory environments.

2

Balancing risk, productivity, and adherence.

This challenge includes unleashing innovation throughout a company while adhering to privacy and localization laws.

3

Managing the AI compliance journey.

Companies need to manage the maturity of their compliance strategies programmatically to leverage, and cope with AI's game-changing potential.

The regulatory stakes keep rising

The rules and regulations surrounding AI change with uncharacteristic speed and scale befitting AI. Since OpenAI introduced ChatGPT in Q4 2022, research published by the [Boston Consulting Group](#) in 2023 found more than 800 regulatory measures governing AI already underway in more than 60 countries worldwide.

Policies for maintaining data privacy and sovereignty must stay flexible as GenAI continues to transform operations, and proposals for new legislation morph and change.

AI is a powerful tool that can provide substantial benefits to individuals and society, but it can also be a "black box." The inner workings of AI — and the results it produces — are often hard to fully comprehend. Governments have a clear interest in protecting against AI failures, such as bias in algorithms or outcomes that are incorrect. ([Boston Consulting Group](#))

Lack of transparency makes it harder to stay current

The lack of transparency surrounding AI starts with the technology itself as companies build new models that grant internal LLMs wide-ranging access to sensitive data and intellectual property. Maintaining visibility to limit interactions proves challenging because LLMs differ from other applications in two fundamental ways:

- **Models are non-deterministic by design** and may produce a variety of outputs even when given the same input.
- **Training data becomes part of the model itself** making it hard to control how information gets shared in response to a user prompt. Limited visibility into how models get developed and train makes it much harder to track, manage, and demonstrate compliance.

The overall lack of transparency into these processes paves the way for new attack vectors targeting LLMs to fly under the radar of legacy rules-and-signatures-based security.




Visibility gaps promote 'shadow AI'

Many businesses also lack transparency into how employees interact with GenAI tools and data. Security and compliance teams don't automatically get notified when someone prompts GenAI or builds new applications without engaging IT.

Or, they know who's using AI but have no way to track what types of data people enter into prompts. IT might know, for example, that the Marketing team uses AI to build content but not notice when someone in Sales inputs non-public information about clients' unique requirements or monthly recurring revenues (MRR).

For more information on risks associated with LLMs check out the [OWASP](#) top 10 classes of LLM vulnerabilities.



Privacy and productivity vie for priority

To realize the full potential of GenAI, forward-looking strategies must strike the right balance between risk, productivity, and adherence to compliance laws and policies. After weighing the various risks, CISOs must develop programs that:

- Dedicate resources to oversee ongoing development, deployment, and documentation of GenAI applications
- Track and control the output generated by AI tools
- Ensure tools and applications do not:
 - Expose proprietary code
 - Fall victim to model tampering
 - Generate unstable code
 - Violate copyrights
 - Introduce vulnerabilities
 - "[Hallucinate](#)" and fabricate wrong, incomplete, or biased answers

Case in point



Samsung [banned the use of ChatGPT](#) after an engineer accidentally uploaded internal source code to the tool.



ChatGPT [suffered a data breach](#) that compromised 100,000+ accounts exposing names, email and payment addresses, and credit card information.


Beyond checking boxes: Responsible AI programs



Progressive companies have begun developing Responsible AI programs that outline principles that guide the design, development, deployment and use of AI. These initiatives aim to establish a framework for making decisions that ensure fairness, reliability, privacy, transparency, inclusiveness, and accountability.

Moving to a modern compliance

	Cost	Efficiency	User experience
Traditional Compliance Practice	Multiple siloed tools and homegrown solutions for data protection, data security, data sovereignty, and data privacy	Manual process to collate logs and meet audit requirements across disparate toolsets	Slow app performance due to data localization constraints
Modern GenAI Compliance Strategy	Consolidated platform-centric approach — one tool for data protection, data security, data sovereignty, and data privacy to eliminate tool sprawl and reduce costs	Automatically monitor, alert, and log GenAI application usage across the organization in a centralized platform	Automated localization decided by software for better user experience and for automatic compliance



Maintaining compliance with AI: 4 modern best practices



Today's regulatory landscape unfolds before a backdrop of economic and geopolitical change. With the rules governing the use of GenAI still evolving, CIOs and CISOs must update policies surrounding when and how their companies use LLMs on a regular basis.

Evolving compliance mandates may well outpace the growth of security and compliance budgets. Following these four best practices helps CISOs do more with their team's existing resources — fulfilling the ultimate promise of AI — while keeping privacy and productivity in lock step.



1

Always know your data



Make sure key players on your team understand:

- **What type of data gets used to train tools**

A clear view of where data exists at all times helps prevent intellectual property, PII, and privileged customer data from inadvertently being used to train LLMs. A comprehensive view should also include augmentative data used to enhance models over time.

- **Where and how data gets collected and stored**

Tracking should also provide a view of where and how data gets protected (for example, the data could be on-prem or in multiple cloud locations).

- **Where the most sensitive data resides**

In this scenario, 'crown jewel' assets would include AI training data and vector embeddings (data used for Retrieval Augmented Generation).

- **How information flows throughout your organization**

Compliance and security teams need a way to understand which users have access to different data sets and by what methods. This might include adopting and applying least privilege authorization models in keeping with Zero Trust.

- **Whether privacy programs meet worldwide certification standards**

Programs must balance industry-wide standards like ISO 27701 that defines parameters for protecting PII with a company's own cybersecurity best practice and governance policies.



Case in point

Cloudflare AI experts witnessed several attacks on LLMs in which threat actors tried to steal data, credentials, and cloud access privileges, or use models to spread crypto-mining malware.

To avoid negative outcomes like these, [Applied Systems](#) runs its own instance of ChatGPT in a Cloudflare isolated browser to block oversharing of data within the tool.



Keep data privacy and security in balance



Placing too high a premium on performance and user experience (UX) sometimes comes at the expense of maintaining sovereignty over your company's data. New localization laws governing the handling and transfer of information within or beyond national borders adds to CISOs' dilemma.

Localization constrains AI training and data storage

Each government sets its own rules for preventing the misuse and exposure of citizens' data to other governments and threat actors. Global GenAI implementations may benefit from storing datasets on servers outside an organization's control — in foreign countries or in the cloud for example — but add to the risk of violating data privacy and localization policies.

Some local laws stipulate that data collected within a nation's borders be processed and maintained within that country. Localization rules vary from country to country which limits multi-national organizations' ability to apply consistent strategies on a worldwide basis, and to process and store data centrally or cost-effectively.

For example, GDPR laws in EMEA and the California Consumer Privacy Act (CCPA) in the US both outline rules for localization but GDPR requires an opt-in for the collection of personal data while CCPA does not. These subtle differences may lead to complex workflows and dangerous workarounds that offset the benefits of using LLMs.

The right compliance strategy allows companies to access data centrally for analysis and decision-making and also store data within the borders of multiple countries where they do business.

Beyond checking the box

Forward-looking GenAI compliance strategies ensure your use of LLMs and handling of data generated by AI:

- **Complies with rules restricting cross-border transfer.** Challenges include consuming SaaS and cloud services, automating data backups, and sharing data between company branches in multiple regions.
- **Enables secure storage and processing of data within each country.** This requires constant vigilance to stay current with evolving local laws and support third-party hosting and processing services in multiple jurisdictions.
- **Meets requirements for encryption and other security measures.** Where requirements vary, maintaining multiple workflows can drive up cost and impact monitoring and security tool performance as data volumes increase.
- **Maintains full visibility and flexibility.** The right strategy lets CISOs choose where to decrypt, analyze, store and maintain logs and encryption keys
- **Leverages real-time intelligence.** Real-time knowledge and tracking of emerging threats and vulnerabilities must take place across all jurisdictions.

Cloudflare's comprehensive approach to compliance, data privacy and localization offers a variety of tools and features to help businesses navigate local compliance regulations without impacting web and application performance. Data gets logged at edge data centers — in memory only (vs. on disk) — and securely transported over a private backbone. Cloudflare offers automatic jurisdiction restrictions when storing vector embeddings for AI Retrieval Augmented Generation applications so your business remains compliant with localization policies like FedRAMP and GDPR.

[Learn more](#)

3

Augment legacy defenses to stay compliant



GenAI tools lack the inherent protections we expect from conventional applications, and legacy controls only get you so far. Adding sophisticated or enhanced DLP capabilities increases security teams' ability to define and enforce granular rules to deal with AI. Development and security leaders can also add new layers of domain-policy-specific content monitoring to ensure GenAI prompts and outputs meet policy guidelines.

People – and point tools – can't keep up

AI consumes and churns out too much data for security and compliance teams to process on their own. Modern risk management strategies call for investing in third-party compliance solutions to streamline policies and automate enforcement and audit reporting across security tools.

New technologies help to identify problems but can't always act to resolve them. Tools need a foundational cybersecurity platform like the [Cloudflare connectivity cloud](#) to write, apply and enforce new policies as well as audit data.

At a minimum, processes and schedules for assessing the risk inherent in training and using LLMs should include:

- **Investing in data governance software solutions** that automate and monitor governance & data lineage
- **Augmenting solutions with AI-powered data masking** that automatically veils information fed to LLMs as well as the resulting output
- **Implementing an AI gateway** to centralize the monitoring, management and security of all GenAI apps (make sure the solution includes logs of input prompts to detect the use of proprietary code, personal or customer data or IP)

Through 2025, attacks leveraging generative AI will force security-conscious organizations to lower thresholds for detecting suspicious activity, generating more false alerts, and thus requiring more — not less — human response. ([Gartner](#))



4 Update strategies and processes regularly



Cybersecurity and development best practices always include regular reviews, assessments, and adjustments. AI intensifies the challenge by accelerating the creation of new applications and content as well as new threat vectors and privacy regulations.

The most progressive companies have begun forming cross-functional AI councils to educate and establish company-wide policies for using AI responsibly. These policies aim to ensure complete data visibility that keeps up with new code, inclusion of more sensitive data, and local and industry regulations.

Cross-functional governance committees or AI councils should oversee all activity and build plans to conduct regular risk and vendor assessments as well as more frequent internal audits. As new proprietary code, personal data/customer data, and IP get added to LLMs, sound assessment practices help organizations understand and control:

- **What data gets exposed.** Publicly available GenAI models and tools like ChatGPT cannot be used in applications if it means training on PII and other non-public data.
- **Monitoring of LLM usage.** Current privacy and AI regulations — and common-sense best practices — include conducting regular risk and impact assessments. To guide the assessment process, University of Oxford researchers created [capAI](#), an independent tool for helping organizations determine their organizations' conformity with [The EU Artificial Intelligence \(AI\) Act](#) that took effect in 2024.
- **Strategies for excluding bias and discrimination** from data sets used to train AI models.
- **Ways to protect algorithms** from unauthorized use by malicious actors. This may mean investing in new technologies and enhancing existing controls such as Data Loss Prevention (DLP).

Go beyond checking the boxes to formulate a cross-functional AI team

☐ AI leaders or head of innovation:

- Lead the implementation of AI and/or how it's brought in to the company
- Evaluate AI products and how they get used by different departments

☐ Sales and marketing leaders:

- Most common users of GenAI applications. Chatbots, content generation, sales email generation, customer sentiment analysis, etc...

☐ IT:

- Make sure there no shadow AI gets implemented by different departments
- Review integration with their existing tech stack

☐ Finance:

- Users of AI since a lot of data entry can be automated
- Review that proper cost limits are implemented

☐ Internal audit teams:

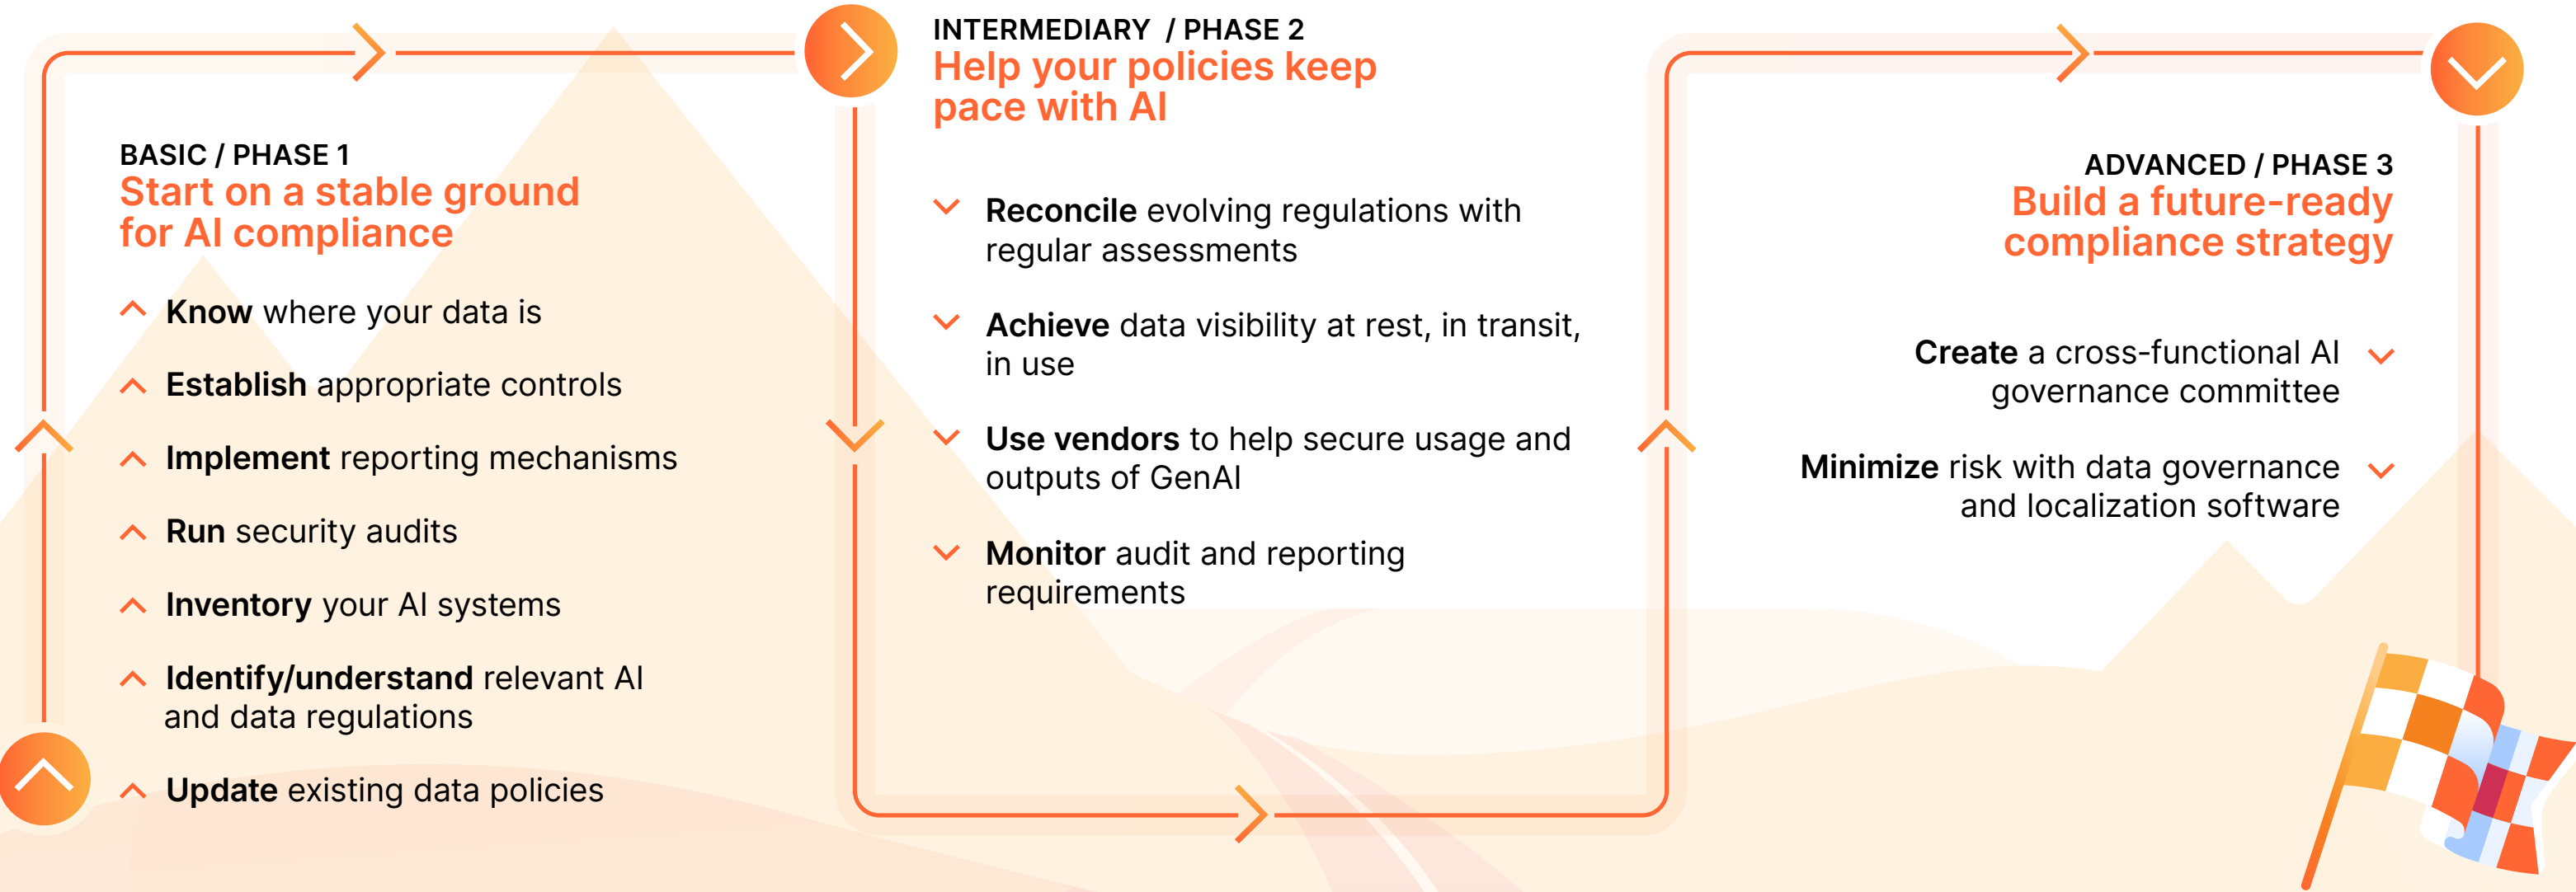
- Provide insights into compliance risks
- Ensure alignment with organizational policies and standards



Stay in control as you manage the journey



AI's potential seems limitless and businesses naturally like to move fast. Those responsible for avoiding risk from cyberattacks, privacy violations, and bad press should come together to build a roadmap for their AI compliance journey.



Things continue to change fast so stay flexible. In the next section we'll see how the Cloudflare connectivity cloud equips CISOs to build and enforce consistent policies that eliminate weak links across global data stores and avoid undue [complexity](#) from operating too many tools to run too many environments

Stay a step ahead with Cloudflare

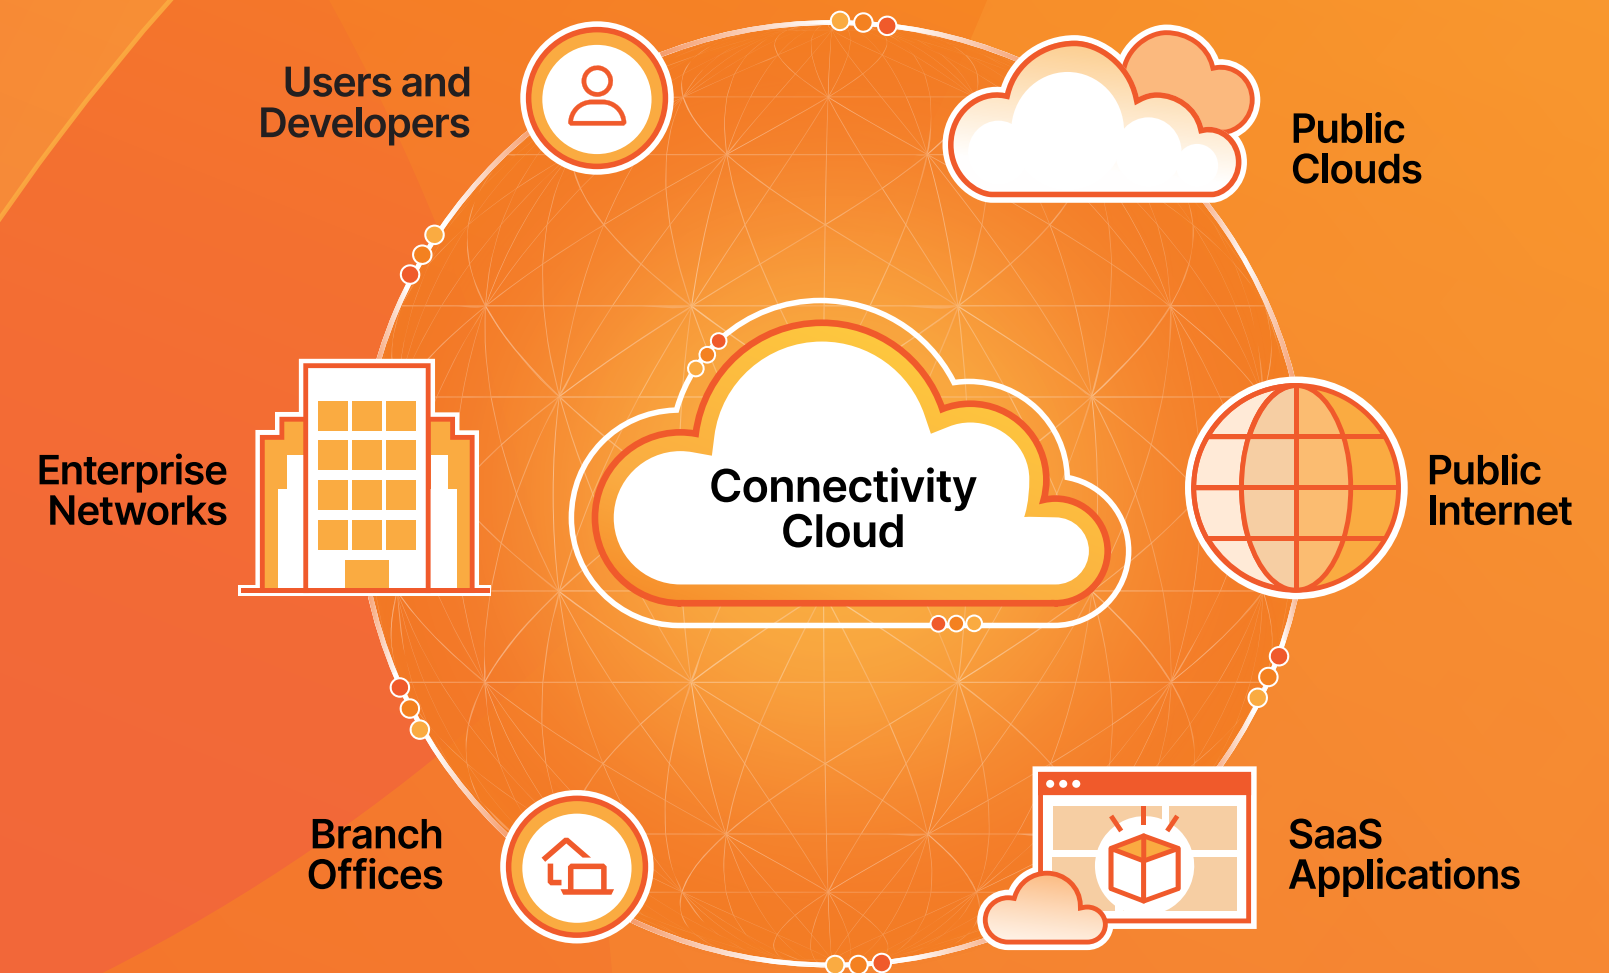


As attackers use LLMs to craft modern social engineering, phishing, and ransomware attacks, cybersecurity vendors and CISOs must integrate GenAI into security stacks just as quickly.

The Cloudflare connectivity cloud consolidates the power of multiple data protection point solutions onto a single platform making it easy to keep pace with fast-changing compliance demands.

Benefits of the Cloudflare connectivity cloud to GenAI compliance include:

- **Unrivalled privacy for exponentially better scale** across 120+ countries with built-in controls for safe handling of PII
- **Complete visibility, control, analytics and reporting** across users, apps, devices, and networks
- **Scalable localization** that enables centralized processing and restriction of traffic inspection to particular regions
- **Programmable network architecture** integrates with compliance automation platforms
- **Reporting and logs provide detailed audit trails** that help meet all audit requirements with geographic-specific regulations related to PII, IP, healthcare, and financial data
- **Maintaining speed and consistency** while lowering cost



Watch webinar

Cloudflare streamlines compliance management to help CISOs scale and protect their data from the game-changing potential of GenAI at the same time.



Advantages include:

The Cloudflare Data Localization Suite allows centralized control of data

Cloudflare’s approach supports data sovereignty and localization at scale with tools and features that help businesses adhere to local compliance regulations. The platform’s [Data Localization Suite](#) features inspection that takes place in seven major regions — the EU, US, FedRAMP Moderate, Canada, Australia, Japan, and India.

Comprehensive visibility and flexible control equip compliance professionals to choose whether to store logs in our core data centers or at third-party locations, and to direct logs to any storage provider in any region.

Firewall for AI keeps time on the side of defenders


GenAI compliance strategies aim to identify abuses before they reach LLM models. To this end, Cloudflare’s [Firewall for AI](#) capability adds a powerful protection layer that can be deployed in front of LLM applications to detect vulnerabilities and provide model owners with full visibility.

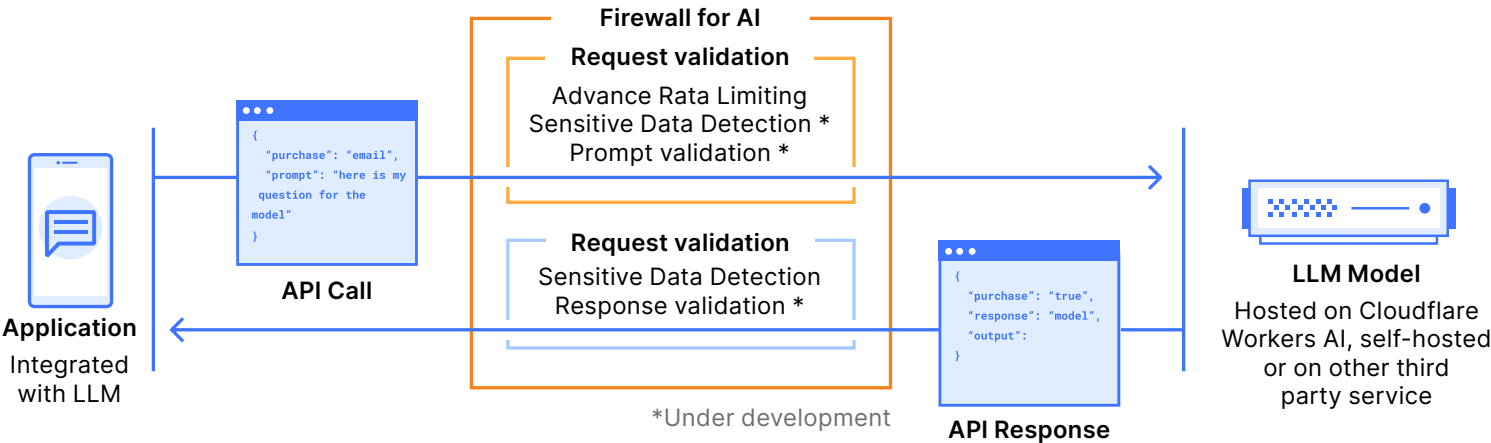
An advanced [Web Application Firewall \(WAF\)](#) tailored to apps that use LLMs, Firewall for AI combats the next wave of attacks targeting the functionality, critical data, and trade secrets held within GenAI tools. Cloudflare’s toolkit includes traditional WAF capabilities like Rate Limiting and Sensitive Data Detection along with new protections that help spot attempts to exploit tools and extract data.

Firewall for AI runs as close to the user as possible to identify attacks early. Cloudflare analyzes prompts submitted by users and API requests containing LLM prompts for IoCs and known attack signatures. The added protection works to prevent volumetric attacks such as DDoS campaigns and prevent model abuses like “prompt injection” and requests designed to generate false or negative responses.

Workers AI secures distributed workforces

Cloudflare’s Workers AI equips organizations to run machine learning models on the Cloudflare network from their own code that runs near users anywhere in the world.

“
We use Cloudflare for everything – storage, cache, queues, and most importantly for training data and deploying the app on the edge, so I can ensure the product is reliable and fast. It’s also been the most affordable option, with competitors costing more for a single day’s worth of requests than Cloudflare costs in a month. The power of the developer platform and performance tools made it an easy choice to build with Cloudflare.”
Bhanu Teja Pachipulusu
[Founder, SiteGPT.ai](#)




Try it now



See what it means to consolidate your strategies for securing innovation and implementation of generative AI tools while staying in compliance and control of your data.


Learn more

Talk with an expert


Recommended resources



[Cloudflare announces Firewall for AI](#)



[How to secure Generative AI applications](#)



[Dispelling the Generative AI fear: how Cloudflare secures inboxes against AI-enhanced phishing](#)



This document is for informational purposes only and is the property of Cloudflare. This document does not create any commitments or assurances from Cloudflare or its affiliates to you. You are responsible for making your own independent assessment of the information in this document. The information in this document is subject to change and does not purport to be all inclusive or to contain all the information that you may need. The responsibilities and liabilities of Cloudflare to its customers are controlled by separate agreements, and this document is not part of, nor does it modify, any agreement between Cloudflare and its customers. Cloudflare services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied.

© 2024 Cloudflare, Inc. All rights reserved. CLOUDFLARE® and the Cloudflare logo are trademarks of Cloudflare. All other company and product names and logos may be trademarks of the respective companies with which they are associated.