Secure GenAI Adoption with Zero Trust:

Secure Use of Public GenAI Applications

WHITEPAPER

Table of Contents

Introduction

Generative AI (GenAI) is transforming how governments operate, enabling them to improve productivity, streamline processes, and better serve constituents. However, to harness GenAI's transformative potential while mitigating its inherent risks, agencies must apply Zero Trust principles. This paradigm ensures that no entity—human or machine—is trusted by default, ensuring continuous visibility and stringent verification at every interaction .

This white paper is the first in the "Secure GenAI Adoption with Zero Trust" series, a comprehensive strategy designed to support government agencies in safely navigating the GenAI landscape. The series includes three phases:

- Phase 1, outlined in this document, focuses on securing public GenAI applications to address risks like data leakage and unauthorized/unsanctioned AI usage ("shadow AI").
- Phase 2 will explore adopting Agentic AI tools to boost employee productivity securely.
- Phase 3 will focus on safely deploying GenAI systems for citizen services, ensuring that government systems and data remain protected.

Each phase emphasizes a proactive, layered approach to balancing innovation with robust governance and security.

Secure Use of Public GenAI

Overview

Governments are increasingly aware of the transformative potential of Generative AI (GenAI) for their operations and the services they provide to citizens. This technology offers a pathway to significant productivity gains and the evolution of citizen services through diverse applications. These range from understanding public sentiment and providing AI-powered chatbots for citizen and IT support to facilitating language translation and automating internal processes like drafting job descriptions, summarizing meetings, and creating public announcements.

Early adopters within government are already witnessing improvements in employee experience and satisfaction. The emergence of publicly accessible Large Language Models (LLMs), like ChatGPT, has spurred experimentation across the public sector as organizations seek to understand and leverage AI capabilities. This widespread interest underscores the opportunities for enhancing efficiency and service delivery through the integration of these advanced AI tools.

However, the integration of GenAI, particularly through public LLMs and third-party models, introduces considerable security challenges. Unauthorized use of AI tools ("shadow AI") can expose sensitive citizen data, business records, or intellectual property. The risk is further amplified in workflows involving Retrieval Augmented Generation (RAG) or Model Content protocol (MCP) and AI Agents, potentially compromising sensitive data and posing national security risks through the potential for state-sponsored actors or malicious entities to exploit these vulnerabilities for espionage, sabotage, or the disruption of critical infrastructure. Moreover, GenAI presents a broad attack surface that traditional security measures, often relying on restrictive binary controls or lacking comprehensive visibility across different environments, are ill-equipped to effectively manage.

To leverage GenAI's potential, agencies should embrace a Zero Trust approach with robust security, visibility, and user simplicity.The following steps outline a process agencies can take to harness GenAI, while proactively mitigating the risks of data leakage and preventing undue burden on security teams:

1. Establish AI Governance Frameworks and Policies
2. Tightly Integrate User Experience and Training
3. Choose the Right Architecture and Prioritize Security
4. Implement Data Protection from the Start
5. Use a Layered Approach to Protection

Let's dive into these steps in more detail.

1. Establish AI Governance Frameworks and Policies

To fully leverage the benefits of GenAI, agencies must implement robust security measures that directly address risks without hindering user productivity. This section explores how agencies can adopt a Zero Trust approach to GenAI applications while ensuring security controls do not impede a seamless user experience.

Developing AI governance frameworks and policies is essential for securing GenAI adoption within state agencies. This often involves creating a dedicated task force or governance body to oversee policy development and implementation. For example, the Alabama GenAI Task Force serves as a model with its collaborative, cross-functional team approach. Agencies should also leverage established Zero Trust frameworks, such as CISA's Zero Trust maturity model and NIST 800-207, along with AI-specific security frameworks like the NIST AI Risk Management Framework (AI RMF), which emphasizes core functions such as governance, mapping, measurement, and management, or Gartner's TRISM, to guide their efforts. By adopting a focused task force and utilizing these proven frameworks, agencies can accelerate the secure integration of GenAI technologies across departments.

To support this process, Zscaler provides insights that help agencies track AI usage across their environments, assess potential risks tied to GenAI applications, and identify instances of data leakage. Leveraging Zscaler's reports allows agencies to access critical data on how GenAI tools are currently being used.

Datapoints Provided by Zscaler to Support Creation of AI Governance Framework and AI Usage Policy

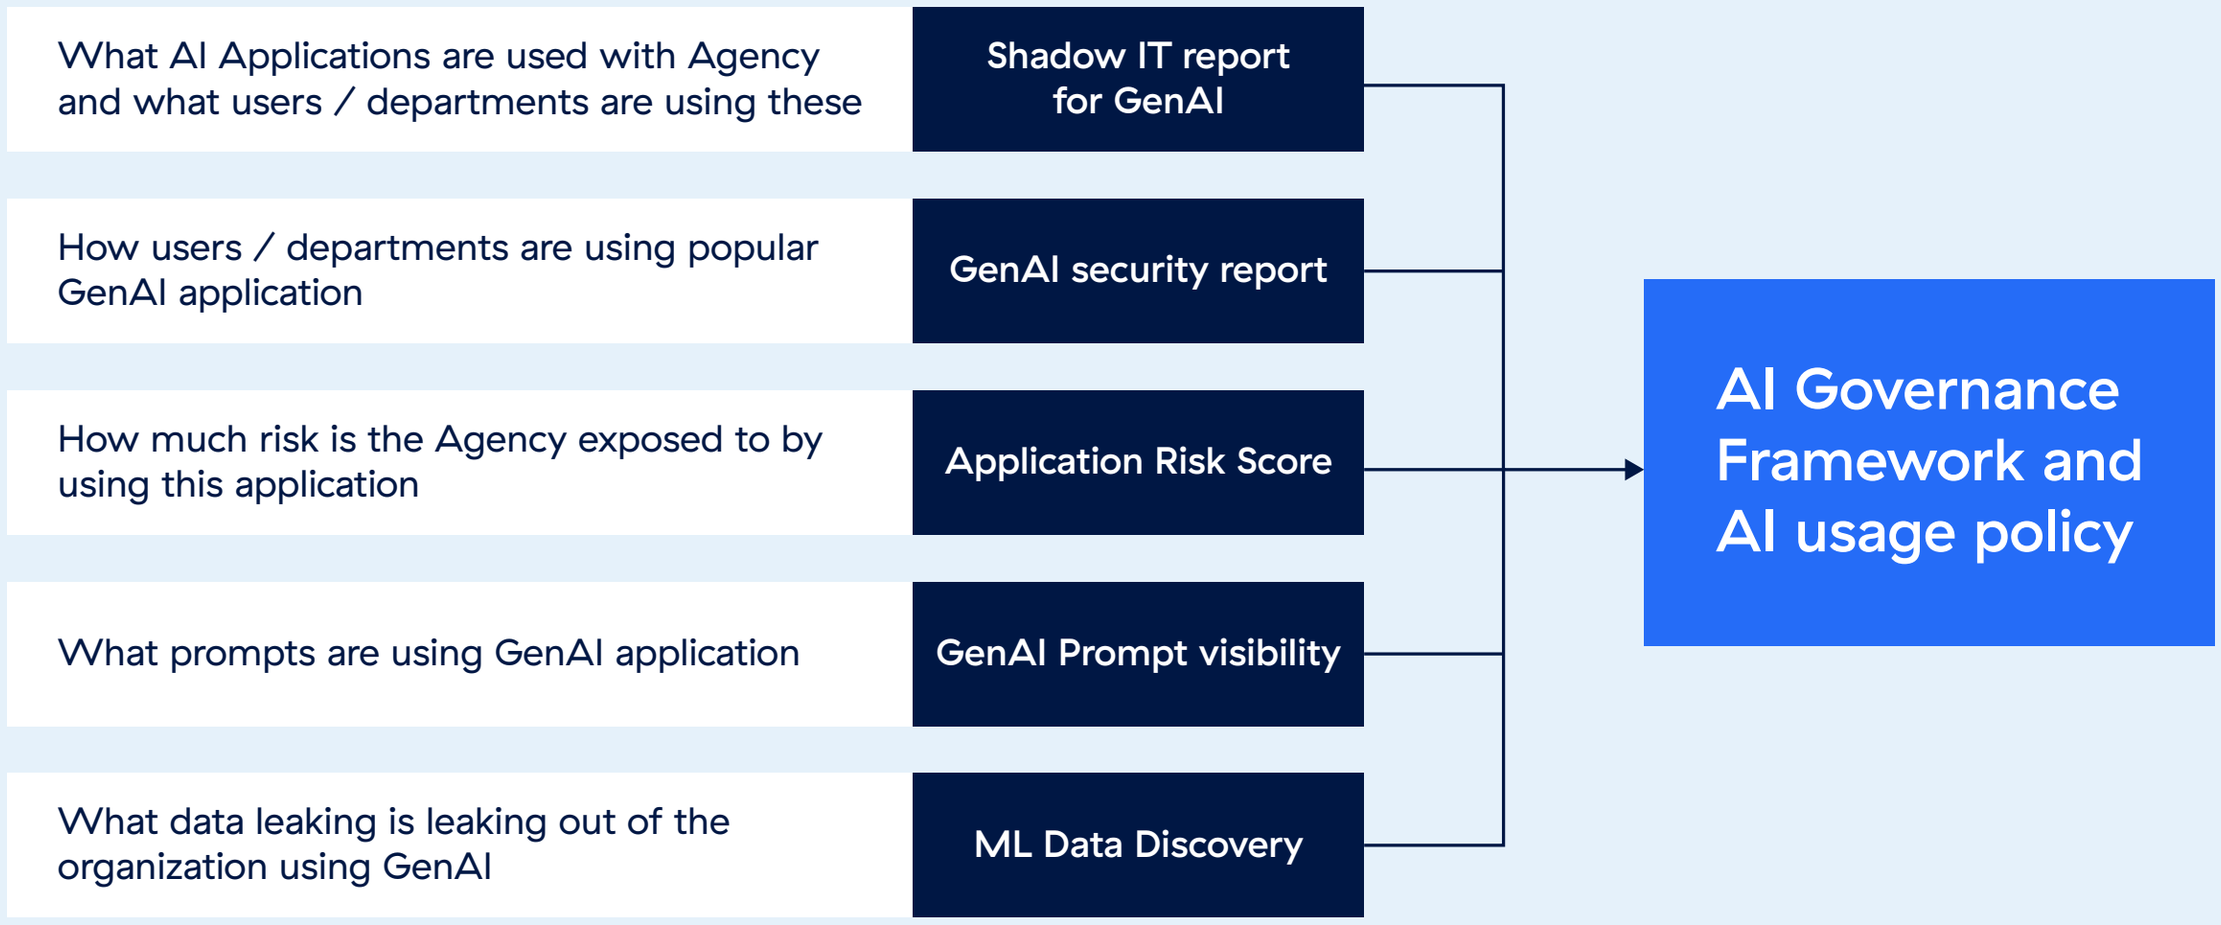

Understanding Current AI Usage

Understanding current AI usage is a key step in creating governance frameworks. By analyzing which GenAI applications are being utilized, how they are being applied, and the associated risk factors, agencies can identify where policies are most needed. This data-driven approach ensures the framework remains relevant, actionable, and tailored to effectively address the agency's unique challenges and opportunities.

Zscaler provides detailed AI usage reports that offer transparency into which GenAI applications are being used across agencies and the extent of their usage. These insights can be segmented further to show usage patterns within specific departments or sub-agencies, giving organizations a clearer view of their AI usage landscape.

Insights into Shadow AI Usage

Shadow IT app discovery for GenAI

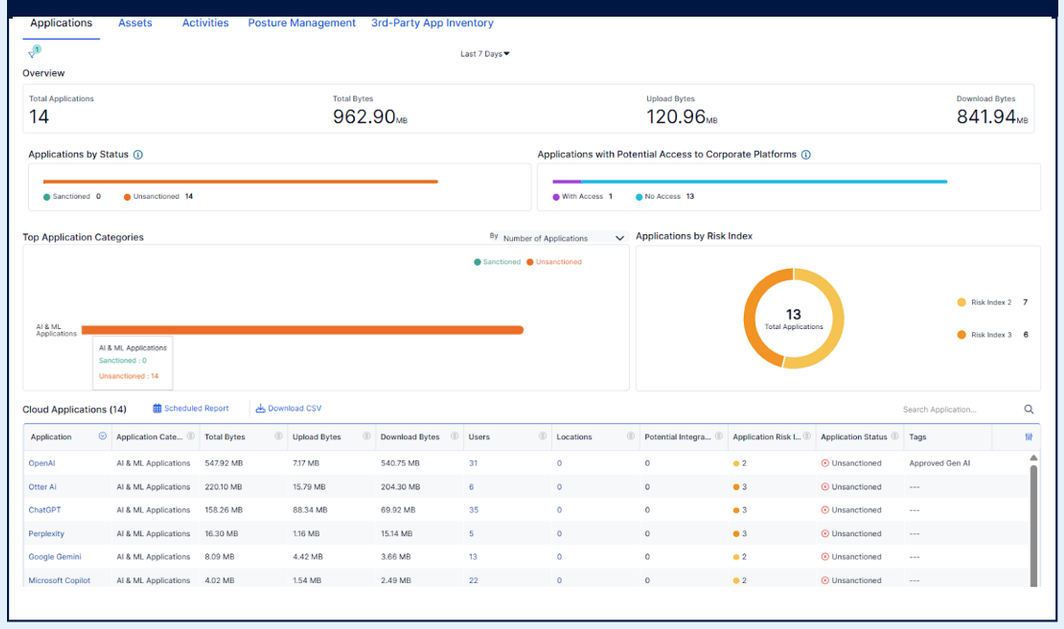

GenAI Security Report

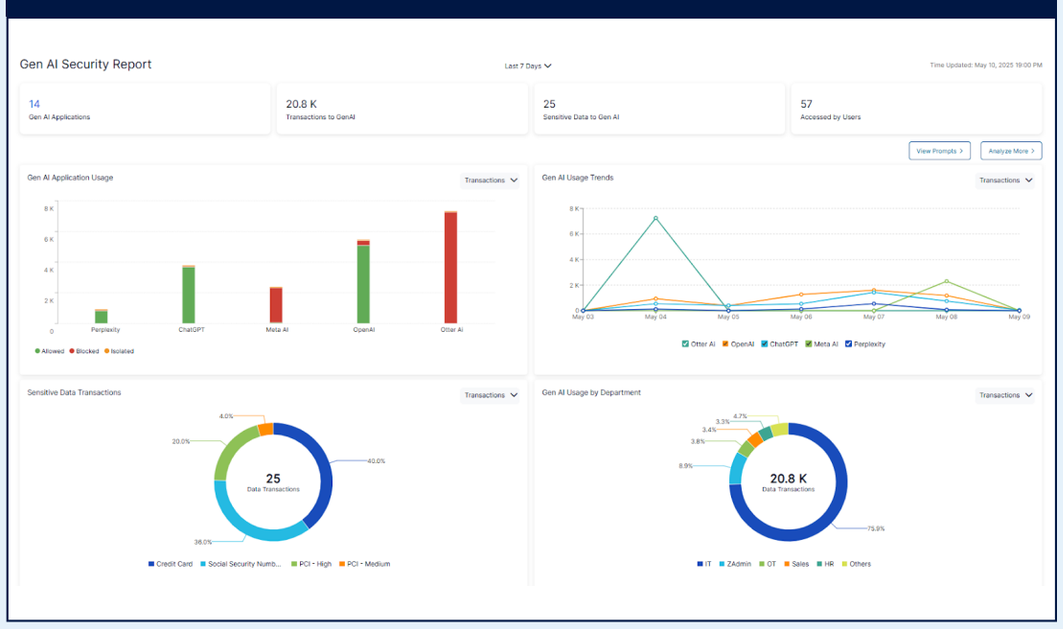

With this visibility, agencies gain the ability to dive deeper into the associated risk factors tied to these applications. Zscaler's ThreatLabz team, in coordination with third-party threat intelligence, evaluates these risks and assigns them aggregated scores ranging from 1 to 5, simplifying risk consumption for decision-makers. Agencies also have the flexibility to customize these scores based on their unique priorities and requirements. Risk assessments can include key factors such as security vulnerabilities or regulatory compliance issues, enabling policymakers to focus resources on the areas most relevant to their mission and security needs. An example of few of the risk factors are shown in the report below such as security vulnerabilities, or non-compliance with regulation which allows the agency policy makers to prioritize the areas that are important to the respective agency.

Risk Associated with Shadow AI Usage

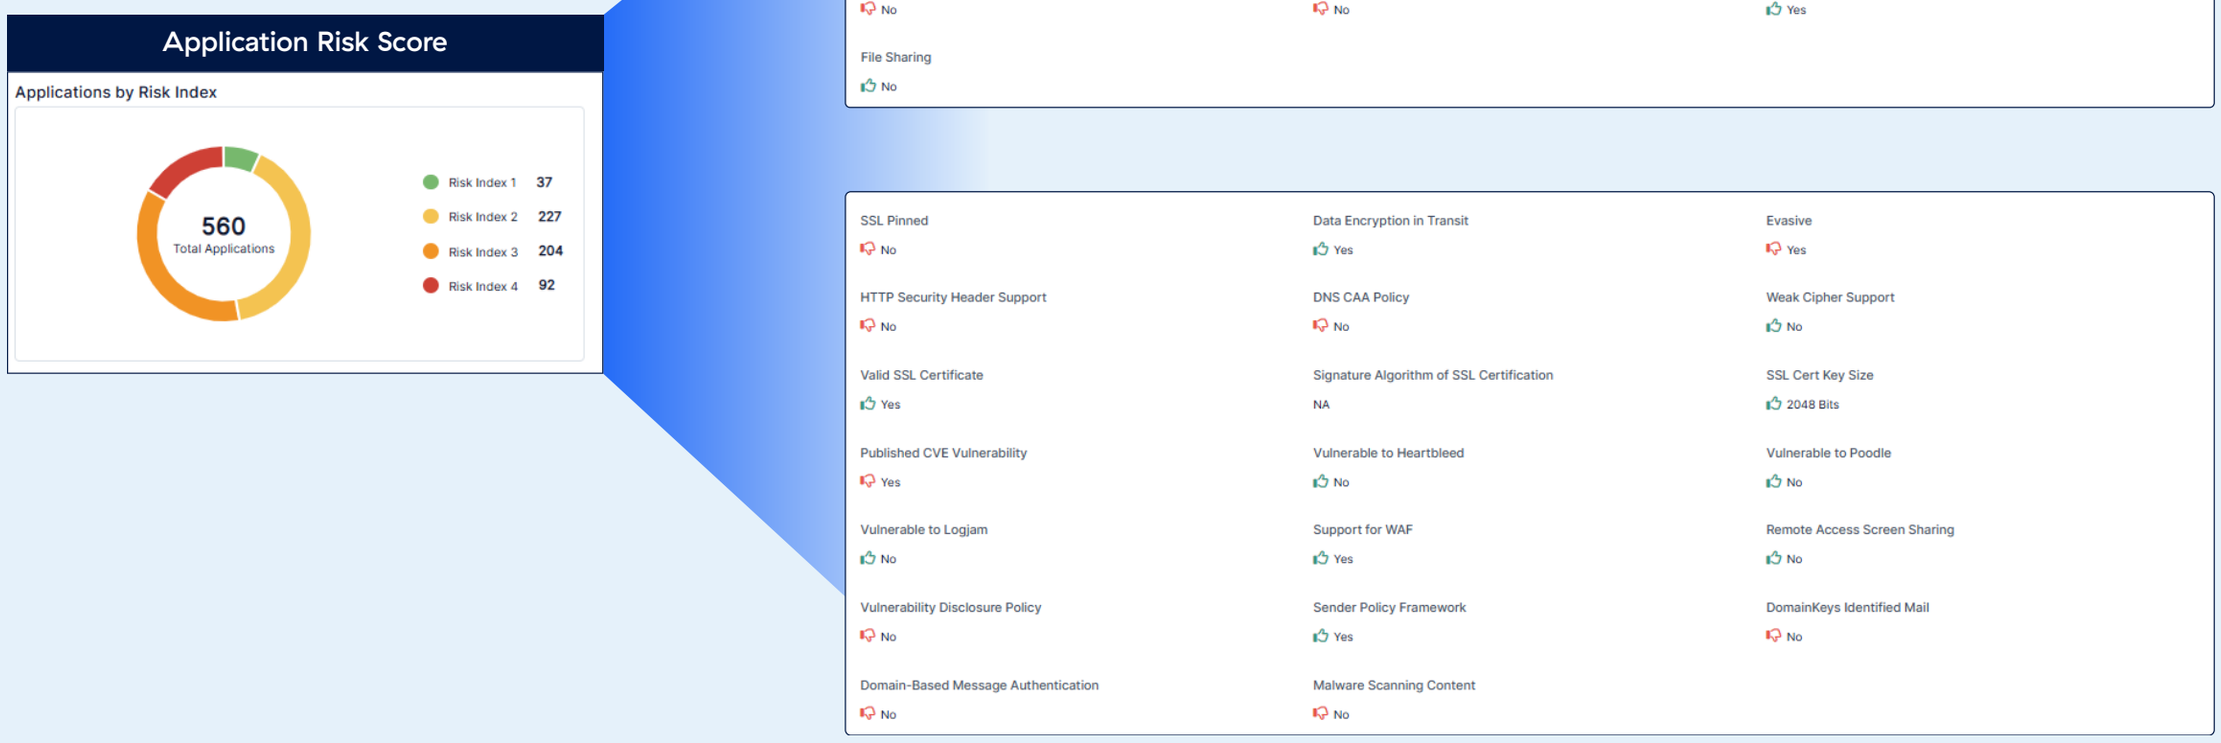

Detailed Insights into User Interactions with GenAI Applications

Zscaler goes beyond application–level visibility by providing granular insights into each transaction, prompt, and user interaction within GenAI applications. This includes detailed data on what users input not just through file transfers, but also via methods like keyboard entries, clipboard activities, and other supported inputs. These insights are invaluable for agencies, helping them better understand the type of data being shared, refine security policies, and ensure compliance with governance standards. Additionally, this level of visibility is essential for audit purposes and can be seamlessly exported to the agency's SIEM for comprehensive tracking and analysis.

Report of Current Data Leakage Outside the Agency

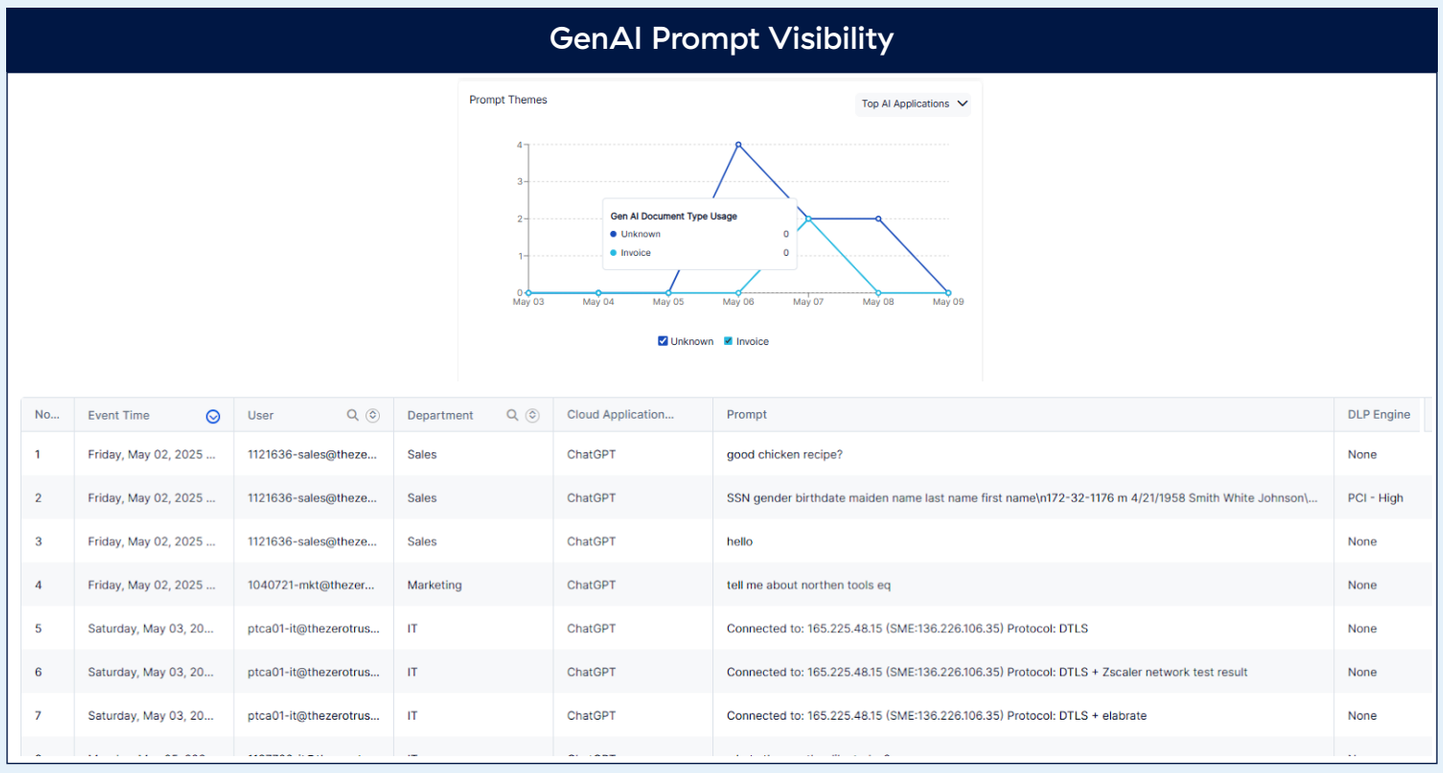

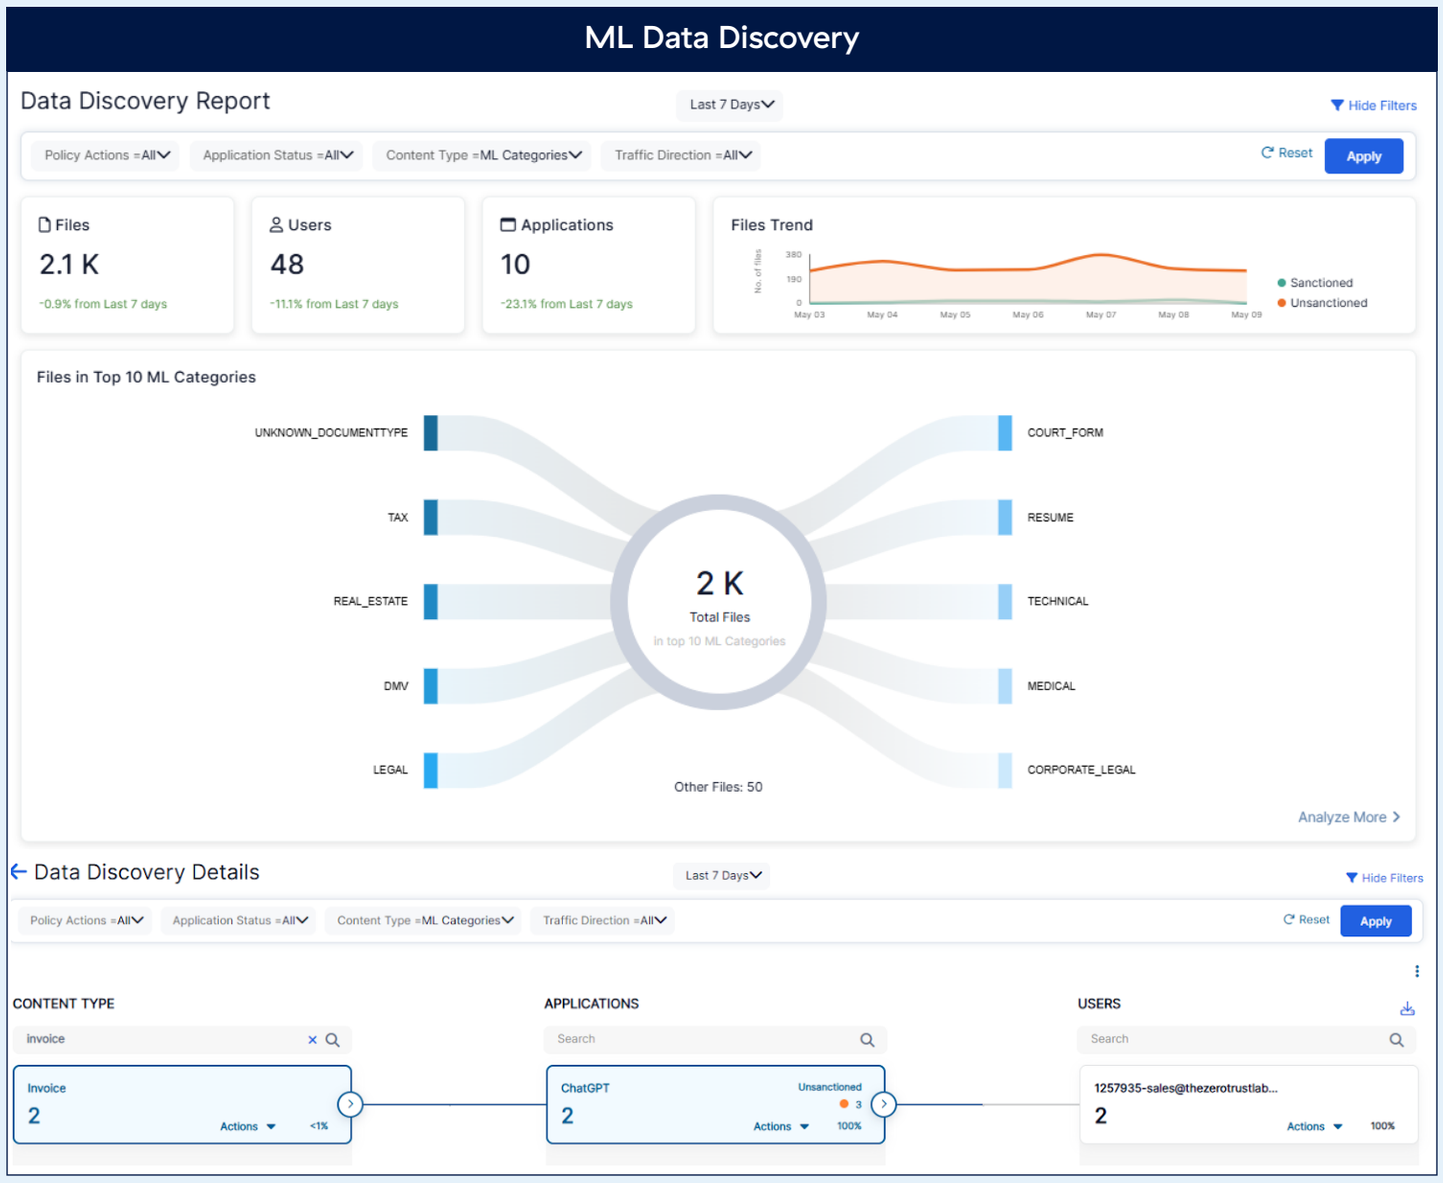

Unknown Data Visibility

Zscaler further enhances visibility by identifying data that agencies may not be aware is leaking through GenAI applications. Using AI/ML–powered capabilities, Zscaler's ML Discovery report goes beyond traditional DLP "monitor–only" rules to proactively detect and classify sensitive data being shared with public GenAI tools. This enables data owners and security administrators to pinpoint unknown or unrecognized data leaks and address them before they become critical issues.

This deep data visibility empowers agencies to proactively identify high-risk data that could be exposed to public LLMs. It also helps establish or refine ownership of sensitive information, develop usage policies, and implement tailored guidelines to protect key datasets.

By combining insights on users, applications, application risks, prompts, and data patterns, Zscaler supports the creation of specific policies and procedures that align with organizational goals. These insights drive resource allocation and help define roles and responsibilities within the Zero Trust governance framework, allowing agencies to take a forward-thinking approach that balances innovation with definition comprehensive risk mitigation strategy.

2. Tightly Integrate User Experience and Training

User experience and training play a central role in the secure and successful adoption of Generative AI (GenAI) within state agencies. To ensure smooth adoption, it's essential that security measures and user training are designed in ways that allow users to remain productive while still offering strong protections. Introducing yet another tool or application, especially those that might introduce new overhead tools for users to learn should be avoided whenever possible. Additionally, effective security controls must be paired with continuous user education to maximize their impact. Platforms should integrate seamlessly with existing workflows and channels while incorporating user interaction and feedback mechanisms. This will help agencies align with frameworks like the NIST AI Risk Management Framework (AI RMF) from the start.

Here are some key platform capabilities that support this approach:

Seamless GenAI Access

The primary goal of GenAI tools is to free users from repetitive tasks and enable them to focus on work that benefits from human judgment. Security measures for GenAI should not disrupt user workflows. Zscaler facilitates this by eliminating the need for additional software or managed browsers. For example,

- **Zscaler Single Agent**
  The same Zscaler agent that ensures secure access to public and private applications also manages GenAI controls, delivering seamless access without introduction of additional tools.

- **Secure Agentless Access**
  Cases,users can use their native browser and their existing workflow(example tile via IDP app portal) to access secured GenAI applications w/o needing an agent.

- **Flexible Security Controls**
  Instead of relying on just “allow or block” options for AI use, Zscaler offers cloud-based browser isolation. This capability redirects users accessing GenAI applications to an isolated browser environment hosted in Zscaler’s cloud. This allows users to maintain a native browser experience while applying advanced security measures such as preventing clipboard activity, printing, or file uploads. This design ensures that security policies are enforced without disrupting the user experience, all managed through a unified platform and single Zscaler agent to simplify administration.

These controls can be deployed with minimal impact on existing infrastructure or endpoints, enabling agencies to implement security policies while preserving a seamless user experience and keeping the administrative effort at a minimal.

Integrated User Training and Feedback

Continuous education on secure GenAI usage and violations is essential, especially given GenAI's rapid evolution. Training should be regular, ongoing and integrated directly into the user's native workflow and tools. Zscaler supports this through dynamic notifications, when a resource is blocked, isolated, or flagged for content violations, users receive custom alerts. For example, if an unsanctioned GenAI application is blocked, Zscaler suggests approved equivalents, helping redirect the users behavior while maintaining productivity. In data use violation scenarios, Zscaler integrates with familiar tools like email and slack, making it easier for users to provide justifications or receive tailored feedback in the tools they already use.

By building user training into security workflows, agencies can establish a strong governance foundation for GenAI applications. This approach not only ensures that users understand how to safely interact with the technology but also helps create a scalable framework for handling GenAI–related incidents and refining AI usage policies across the organization.

User Training and Feedback with Zscaler

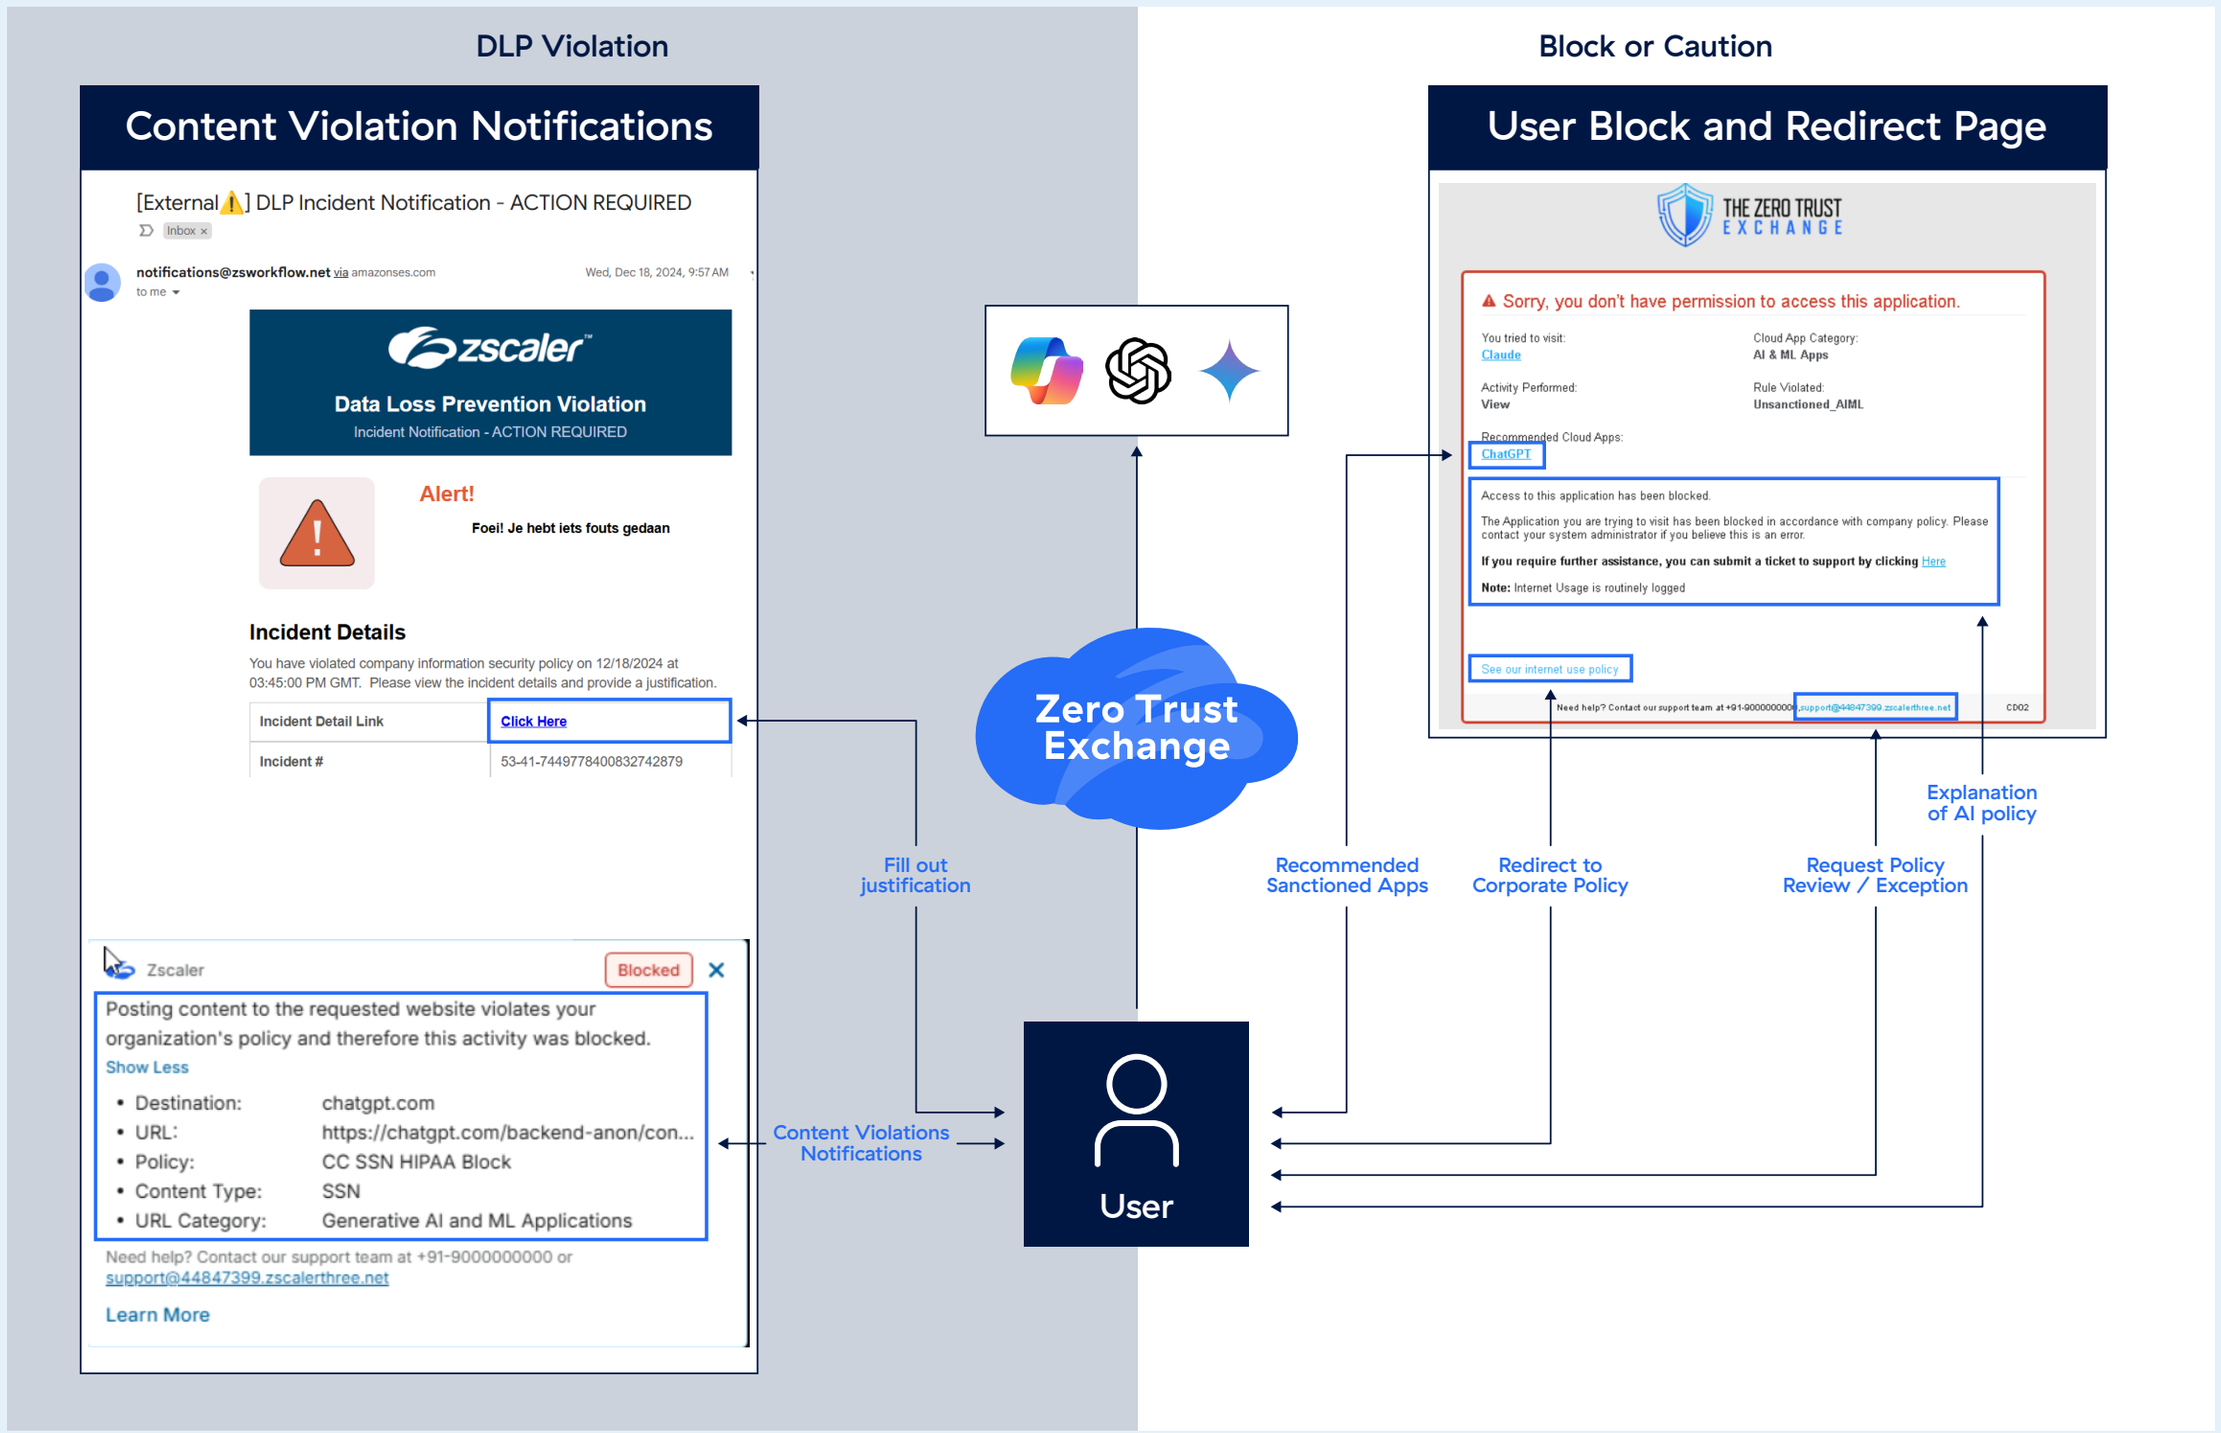

3. Prioritize Security and Choose the Right Architecture

As the attack surface expands with the growing adoption of AI tools, it's crucial for organizations to choose a platform that not only protects against AI–powered threats but also enables the secure use of GenAI applications, built on a Zero Trust architecture, and with FedRAMP authorization. Zscaler provides a comprehensive solution that combines security, usability, and compliance for agencies adopting GenAI.

Agencies can start leveraging Zscaler's secure GenAI capabilities with minimal changes to their existing infrastructure. Traffic can be forwarded to the platform either through the Zscaler App, which is compatible with major operating systems including Windows, Mac, iOS, Android, and Linux, or via branch device, or simply a browser for agentless deployment. Once connected, Zscaler ensures that users are authenticated and authorized before accessing any resources on the internet including GenAI applications. Acting as a proxy, Zscaler inspects all TLS sessions established with GenAI applications, providing agencies full visibility and granular control over how users interact with these tools. This architecture empowers agencies to provide deep insight into how users are using GenAI applications, and provide multiple industry leading controls to secure the Agencies data.

HTTPS and Secure Websocket Inspection with Zscaler

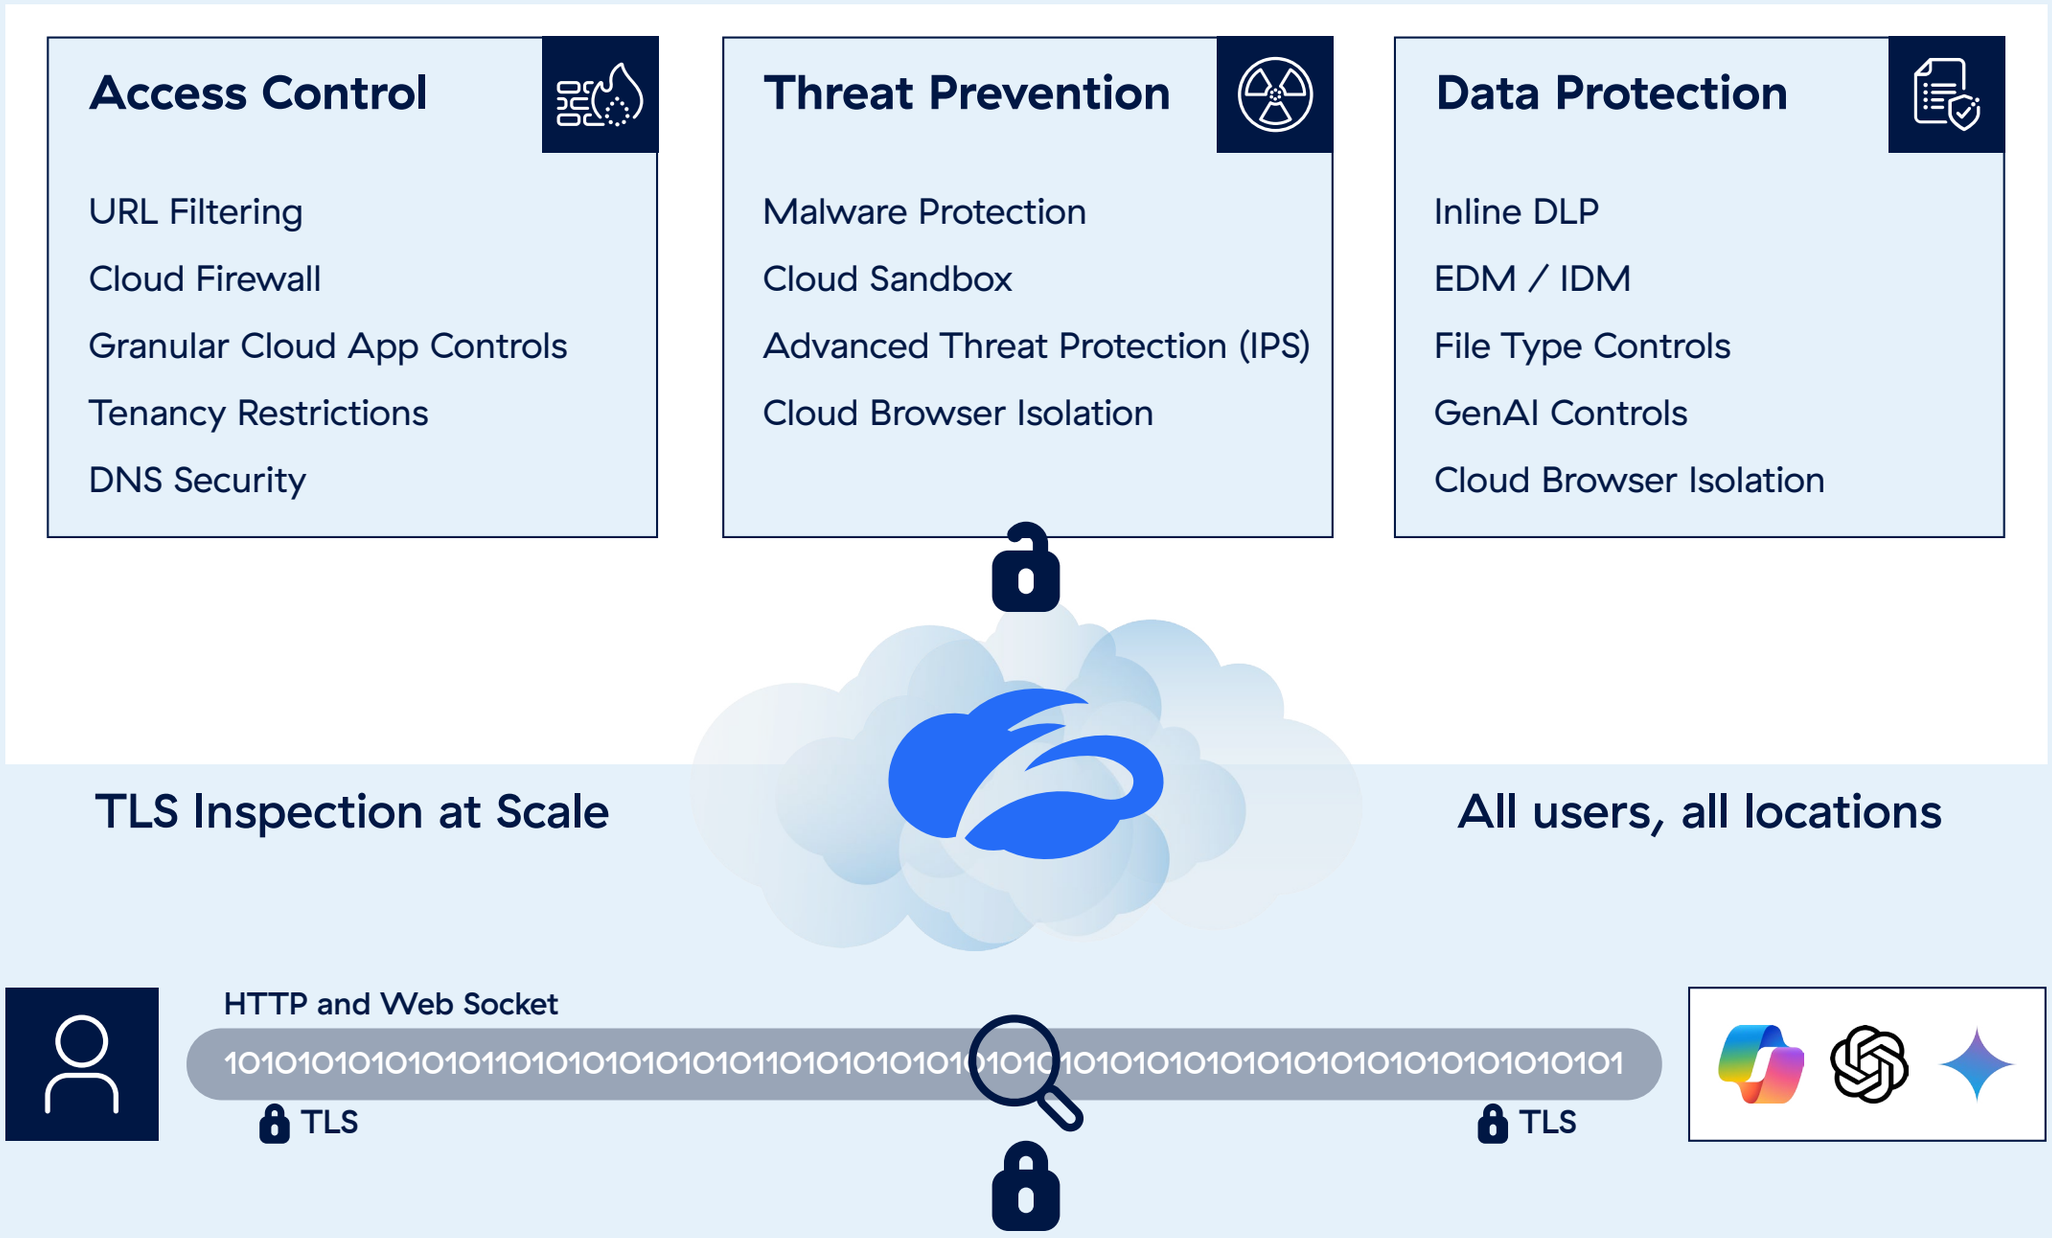

Automate GenAI Application Discovery and Management

With TLS inspection in place, agencies gain access to the full suite of Zscaler's capabilities, including granular control over GenAI and machine learning applications. A key advantage lies in Zscaler's AI & ML Applications category, curated by the ThreatLabz team. This category encompasses a wide range of artificial intelligence applications, including popular tools like ChatGPT, Gemini, MetiAI, Claude, and others.

Using this category, agencies can enforce policies to block unknown or unvetted GenAI applications by default, ensuring only approved tools are accessible. As new applications emerge, they are automatically added to these categories, saving agencies the effort of manual discover and push updates. Additionally, agencies have the flexibility to expand or tailor this list by adding custom domains to better align with their specific needs. Zscaler also provides dedicated categories such as "General AI & ML Applications" and "Generative AI and ML Applications," which, when paired with the broader "AI Cloud Applications" list, deliver significant coverage to reduce the security risks GenAI applications pose. This layered approach enables agencies to effectively manage access to hundreds of applications being developed and released every week.

Broad Category and AI Application Specific Selection

URL Categories for Wide Net

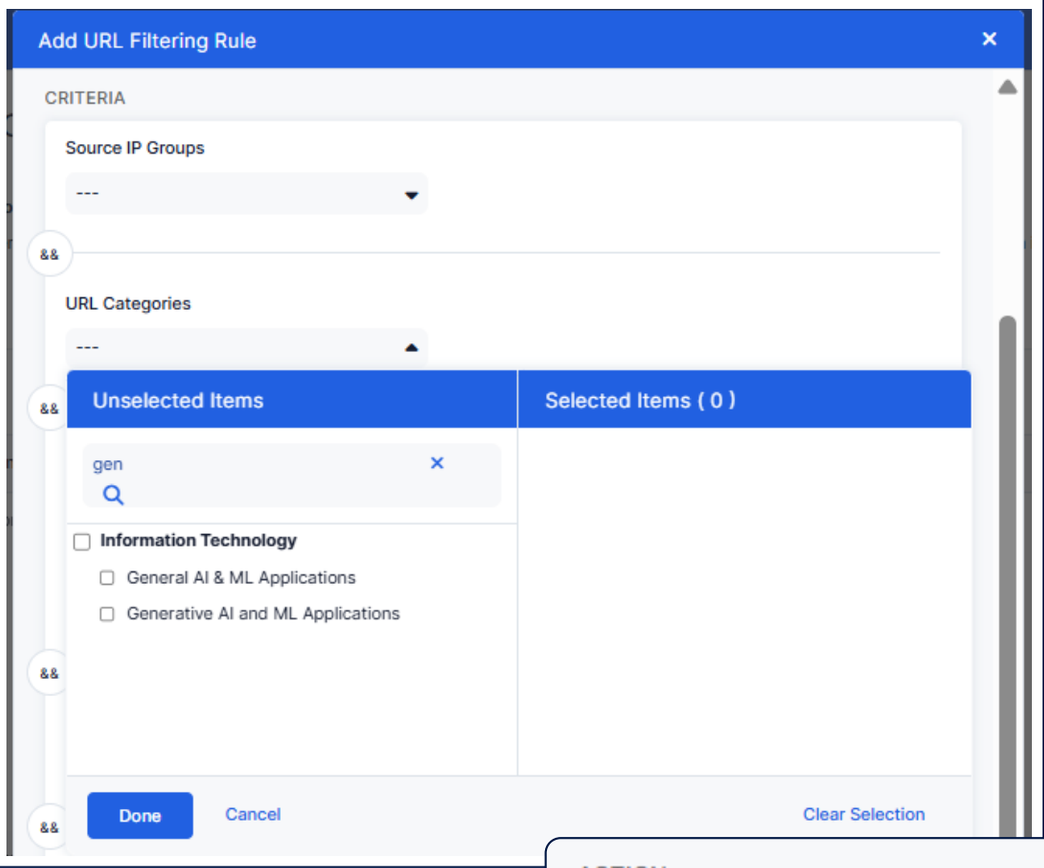

GenAI Application for Granular Controls

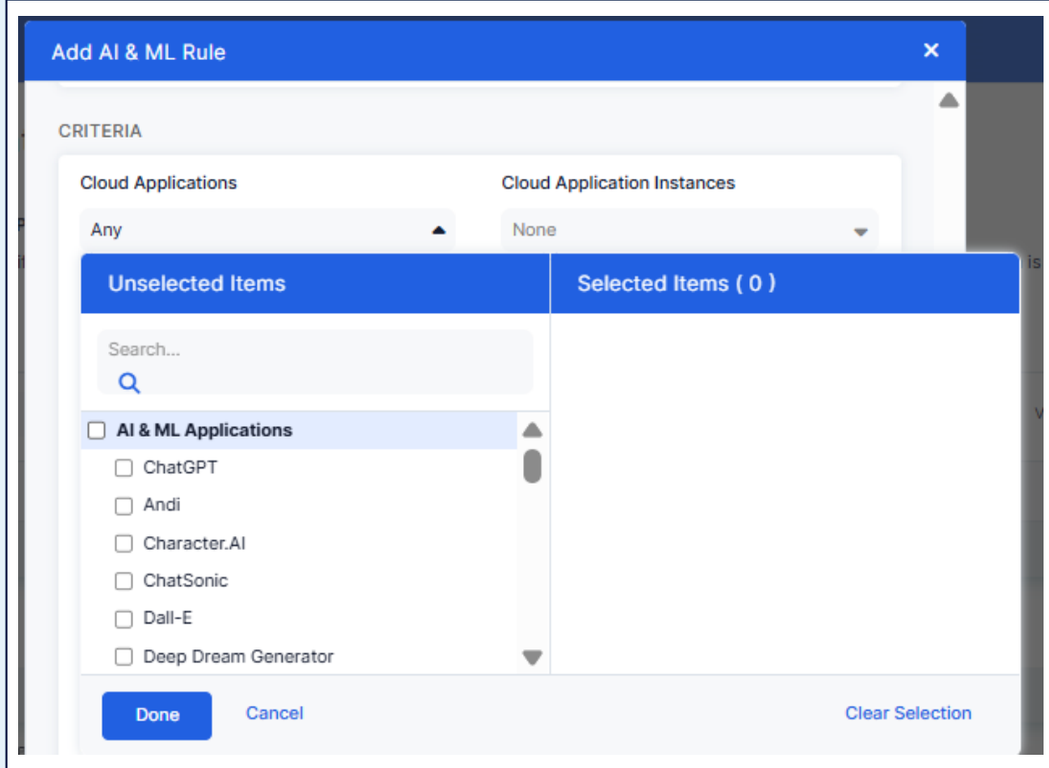

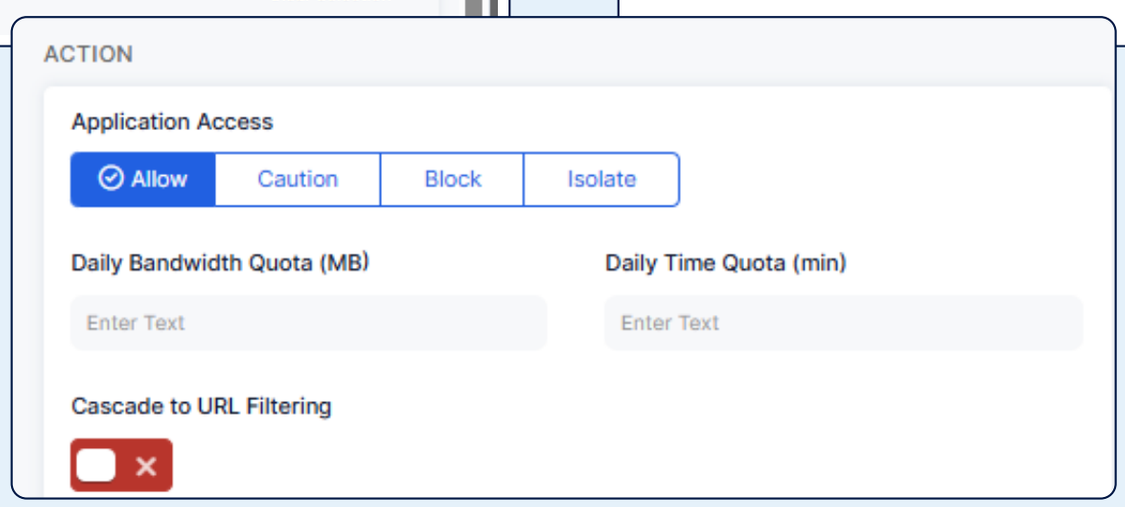

Granular Controls for SaaS, Web, and AI Applications

GenAI Application Granular Controls

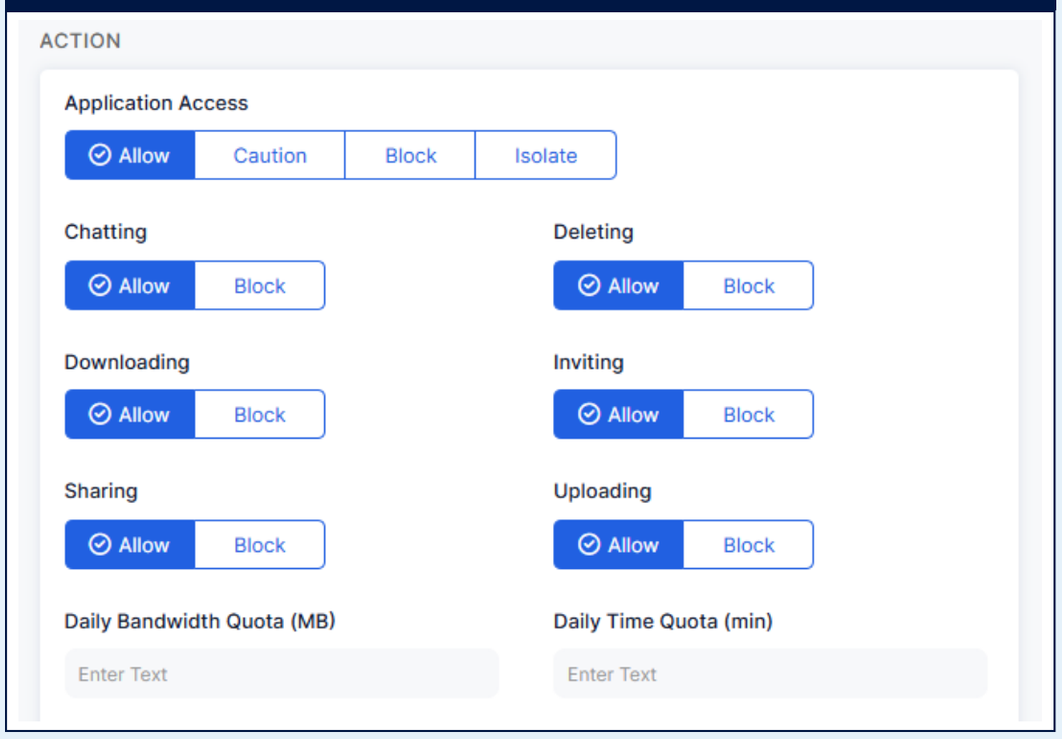

GenAI Application File Upload Controls

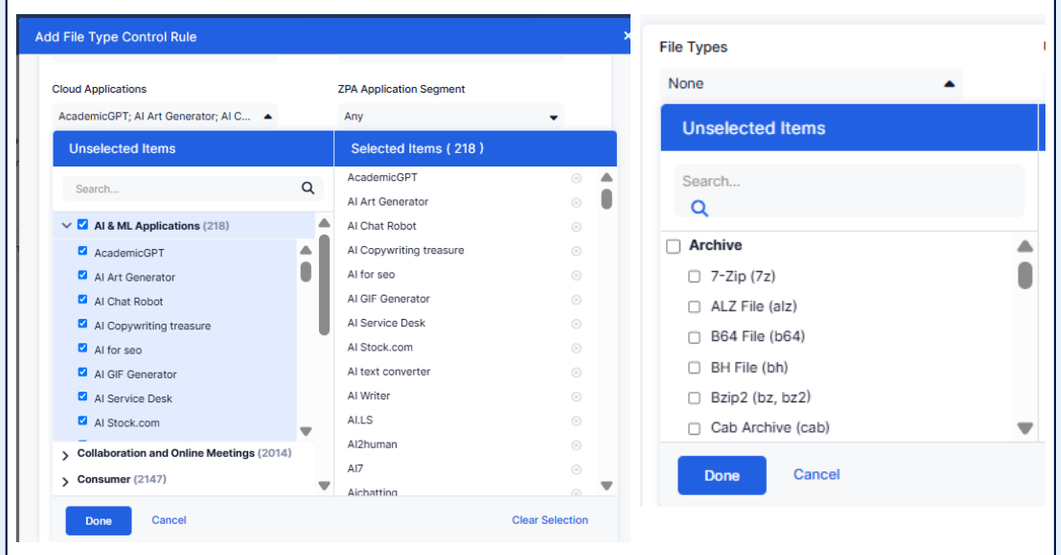

Allow Sanctioned Apps via SaaS Application Security Control

On top of maintaining a comprehensive list of AI applications, , Zscaler provides granular controls of how users will interact with genAI applications. These controls are incredibly simple to apply, very powerful, and consolidated within a single platform. The left side of the image shows a few of the examples of granular controls that can be applied, in case a chatGPT security policy can include granular controls such as allow chat, but blocking file uploads, or restricting sharing of chats. Agencies can apply these, department wide, or even at each user level. These granular controls can be further refined by restricting the file types that users are allowed to upload to GenAI applications as shown on the right. This control of files can also include restricting uploads of encrypted documents.

Restrict Access to Enterprise Instances of GenAI Applications

Agencies should strongly consider use of enterprise-grade versions of GenAI applications to ensure better security and control. Enterprise versions, such as ChatGPT Enterprise, give agencies full ownership and control of their business data and conversations, without the corporate data contributing to model training. These solutions are SOC2 compliant and provide encryption both in transit and at rest. Additionally, they simplify user management with features like team-based access, domain verification, Single Sign-On (SSO), and usage insights, enabling large-scale secure deployment.

Enterprise instances of GenAI applications should be paired with SSO to maximize security and provide agencies with greater visibility and control over application usage. With SSO in place, agencies can enforce policies that block access to non-enterprise versions of GenAI applications. For example, Zscaler's tenancy control for ChatGPT ensures that only approved tenants can be accessed while others are automatically restricted. In addition, agencies can put controls in place at the Identity and Access Management (IAM) layer using whitelisting to ensure enterprise versions are the sole instance of GenAI use, and to guarantee that access occurs over secure environments like Zscaler's cloud platform. To further extend secure access, enterprise GenAI instances can also be made available to unmanaged or BYOD devices using Zscaler's agentless BYOD access.

A simple "allow-all or block-all" approach is insufficient in today's GenAI landscape. Agencies must adopt a layered security strategy with granular controls tailored to different application interactions. Consolidating these capabilities into a unified platform not only streamlines deployment but simplifies adherence to the core principles of Zero Trust, ensuring least privilege access, continuous visibility, and comprehensive protection for every GenAI interaction.

Access Control to Sanctioned Instances of AI Applications

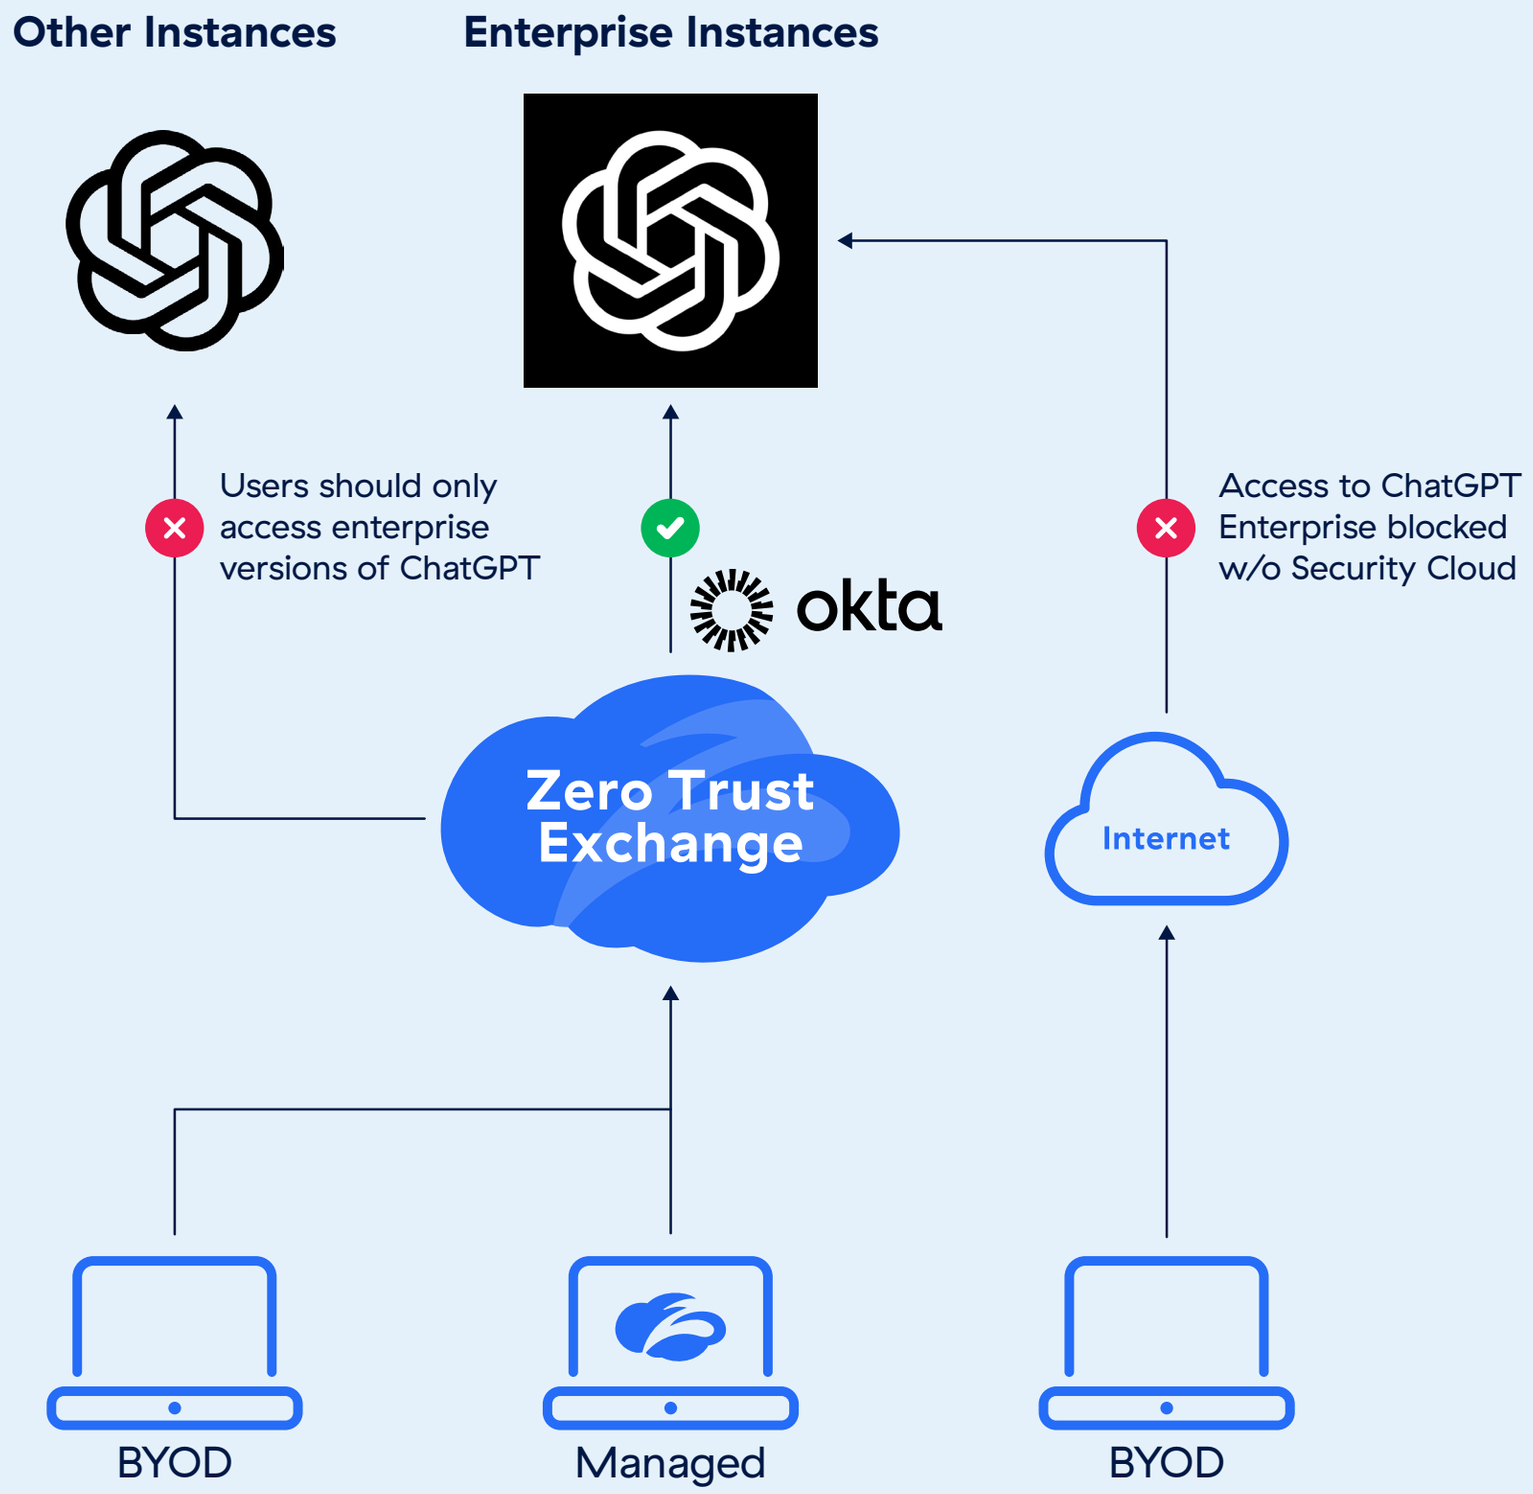

Reduce Risk from Unsanctioned GenAI Applications

When access is needed for GenAI applications that are not sanctioned(they lack enterprise licensing and Single Sign-On (SSO)) these GenAI applications should be treated as high-risk. Data uploaded to such applications may be used to train the GenAI models, potentially exposing sensitive information. To address this heightened risk, agencies must implement additional layers of security controls to ensure stricter oversight of data interactions.

Zscaler offers an effective solution to manage this risk through its Zero Trust Browser. This tool enables agencies to provide secure access to unsanctioned GenAI applications with advanced controls, such as limiting actions like file transfers, printing, and clipboard use. Additionally, the Zero Trust Browser prevents GenAI applications from running code directly on the user's browser, instead rendering interactions on isolated pages. This helps protect against fingerprinting, 3rd-party cookie tracking, and other vulnerabilities—all while allowing users to continue using the same browser deployed by the agency.

This approach can be implemented in two ways: with the Zscaler Unified Agent or using an agentless model. For agency-owned devices, an agent-based deployment is recommended to ensure all traffic is routed through Zscaler's enforcement platform. In situations where an agent cannot be installed, Zscaler's agentless option provides a secure alternative, ensuring controlled access to GenAI applications without compromising security.

Granular Controls to Secure Isolated AI Applications While Balancing User Experience

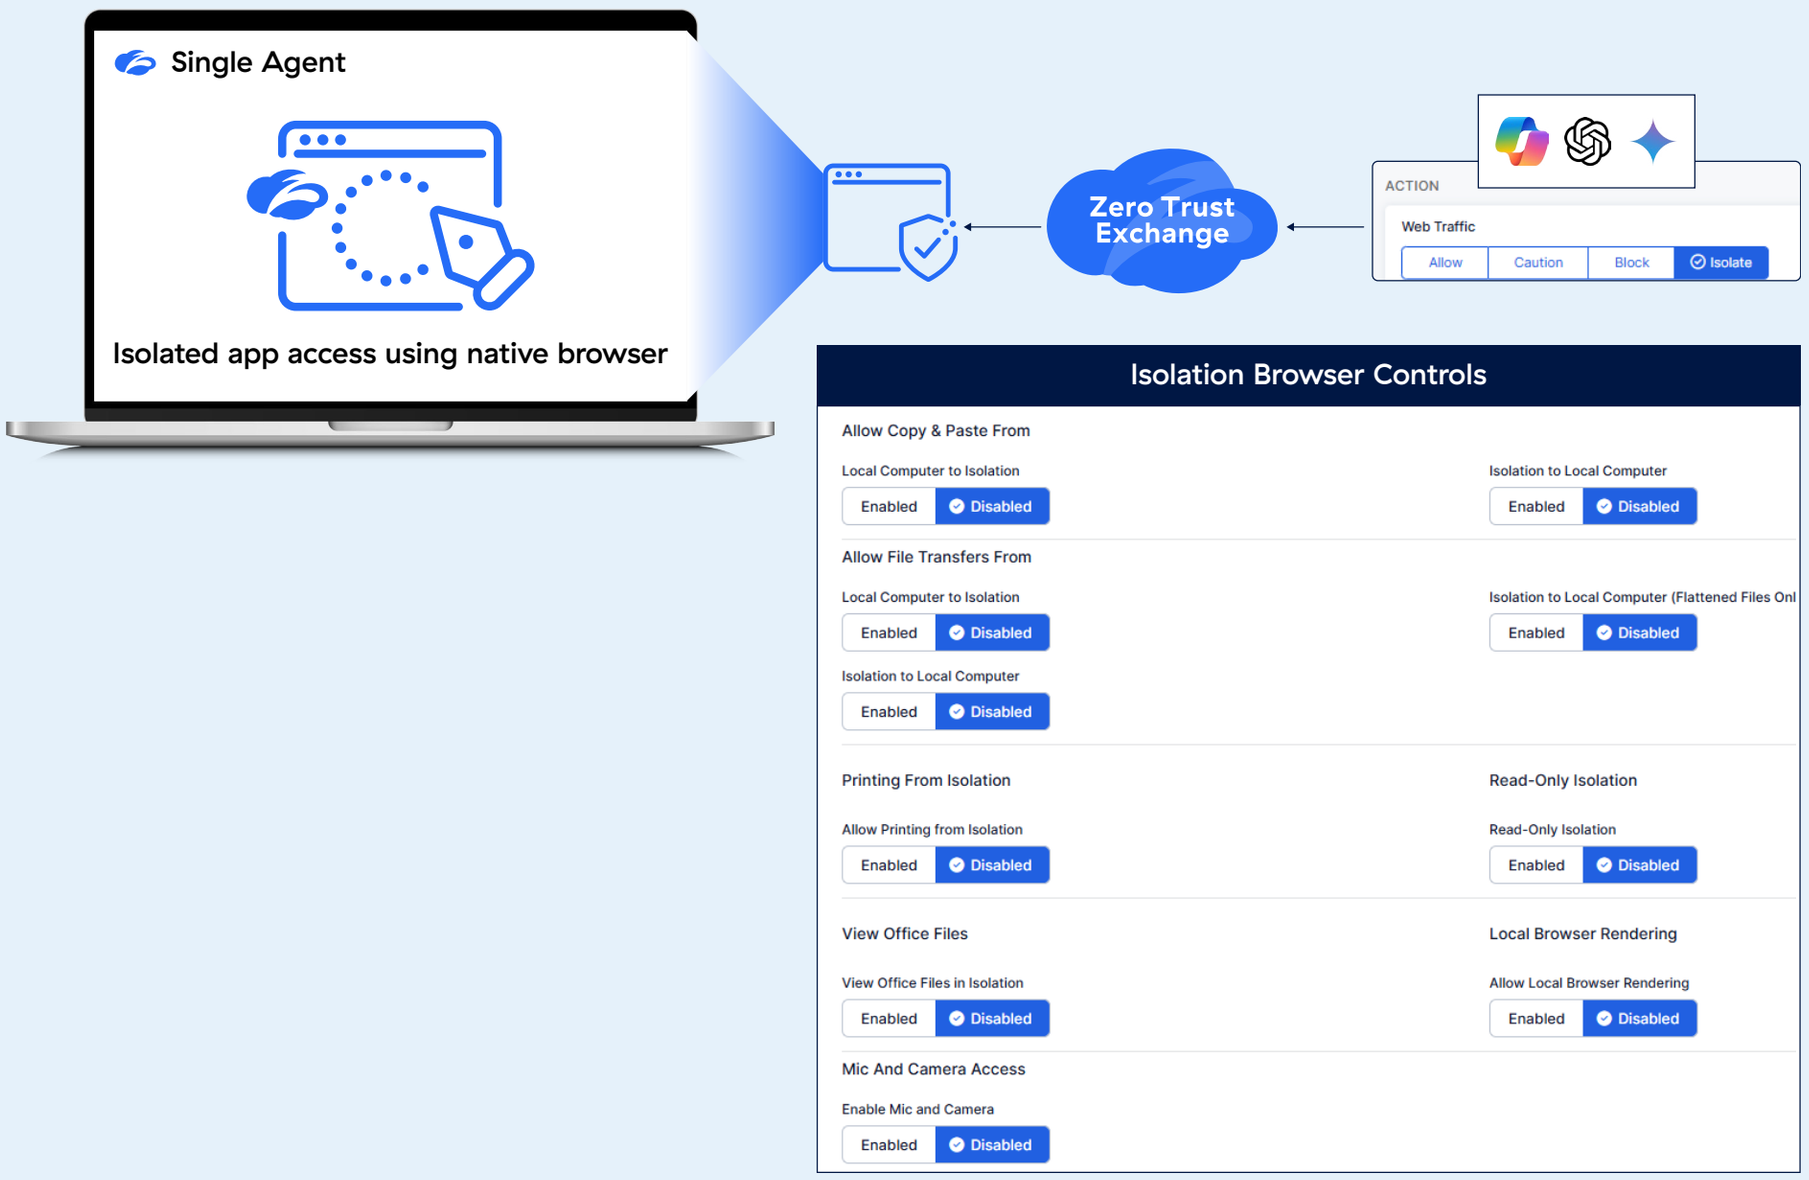

4. Implement Data Protection from the Start

Failing to implement strong data protection from the beginning of GenAI adoption can result in data breaches, violations of privacy regulations, and a loss of public trust—ultimately undermining the success of these tools. The conversational and user-friendly nature of public GenAI applications increases the risk of users unintentionally exposing sensitive government data. Simple actions like copying and pasting information or uploading files can, without careful oversight, leak confidential details due to context or integration with other systems. This highlights why embedding robust data protection measures should be a core part of any public GenAI adoption strategy for state and local governments.

Zscaler enables agencies to tackle these risks head-on with its advanced Data Loss Prevention (DLP) capabilities. Built to protect sensitive information from the start, Zscaler's DLP solution for GenAI identifies and blocks the sharing of confidential data—whether it's through a prompt, file upload, or misuse—before it can reach public GenAI models. This proactive approach ensures agencies can embrace GenAI while safeguarding sensitive information and maintaining compliance.

Accelerate DLP Adoption

Starting a data protection journey can feel like a challenging endeavor for many organizations, especially when balancing the need to grant access to GenAI tools with implementing strong guardrails. Zscaler addresses this challenge by offering a streamlined platform built to support lean teams, enabling rapid adoption of GenAI with effective data protection controls. This approach ensures agencies can scale their security framework efficiently across diverse departments and user bases.

For agencies that already have inline rules applied to other internet destinations, extending those policies to GenAI applications is straightforward. Zscaler also integrating existing DLP engines and dictionaries used for other channels directly into AI & ML applications, reducing redundancy and accelerating deployment. If an agency is starting from scratch, Zscaler provides predefined dictionaries that can be applied to GenAI applications in just a few clicks to prevent sensitive data from being leaked. Additionally, known documents or datasets can be protected using EDM/IDM capabilities, and tagging from Microsoft Information Protection (MIP) can further safeguard encrypted or classified data from exposure.

To refine policies further, Zscaler's machine learning (ML) discovery capabilities identify previously unknown sensitive information and data leaks within GenAI applications, enabling agencies to continuously evolve their protection strategy. Whether by fine-tuning existing dictionaries or creating custom detection rules using regex or keywords, agencies can tailor these to fit their needs. Zscaler also integrates with data backup solutions like Rubrik, simplifying data identification and protection.

Accelerating DLP Implementation with Zscaler

Implement Day O

Agency specific data with EDM and IDM

Pre-built dictionaires that should be used by Government Agencies

- ABA Bank Routing Numbers
- Corporate Finance Document
- Corporate Legal Document
- Court Document
- Credentials and Secrets
- Credit Cards
- Diseases Information
- Driver's License (United States)
- Drugs Information
- Financial Statements
- Immigration Document
- Insurance Document
- Invoice Document
- Legal Document
- Medical Document
- Medical Information
- Real Estate Document
- Social Security Numbers (US)
- Tax Document
- Tax Identification Number (US)
- Transportation and Motor Department Document
- Treatments Information

MIP / AP Labels

Continuous Monitoring & Visibility

Identify Unknown Data Leakage and Apps

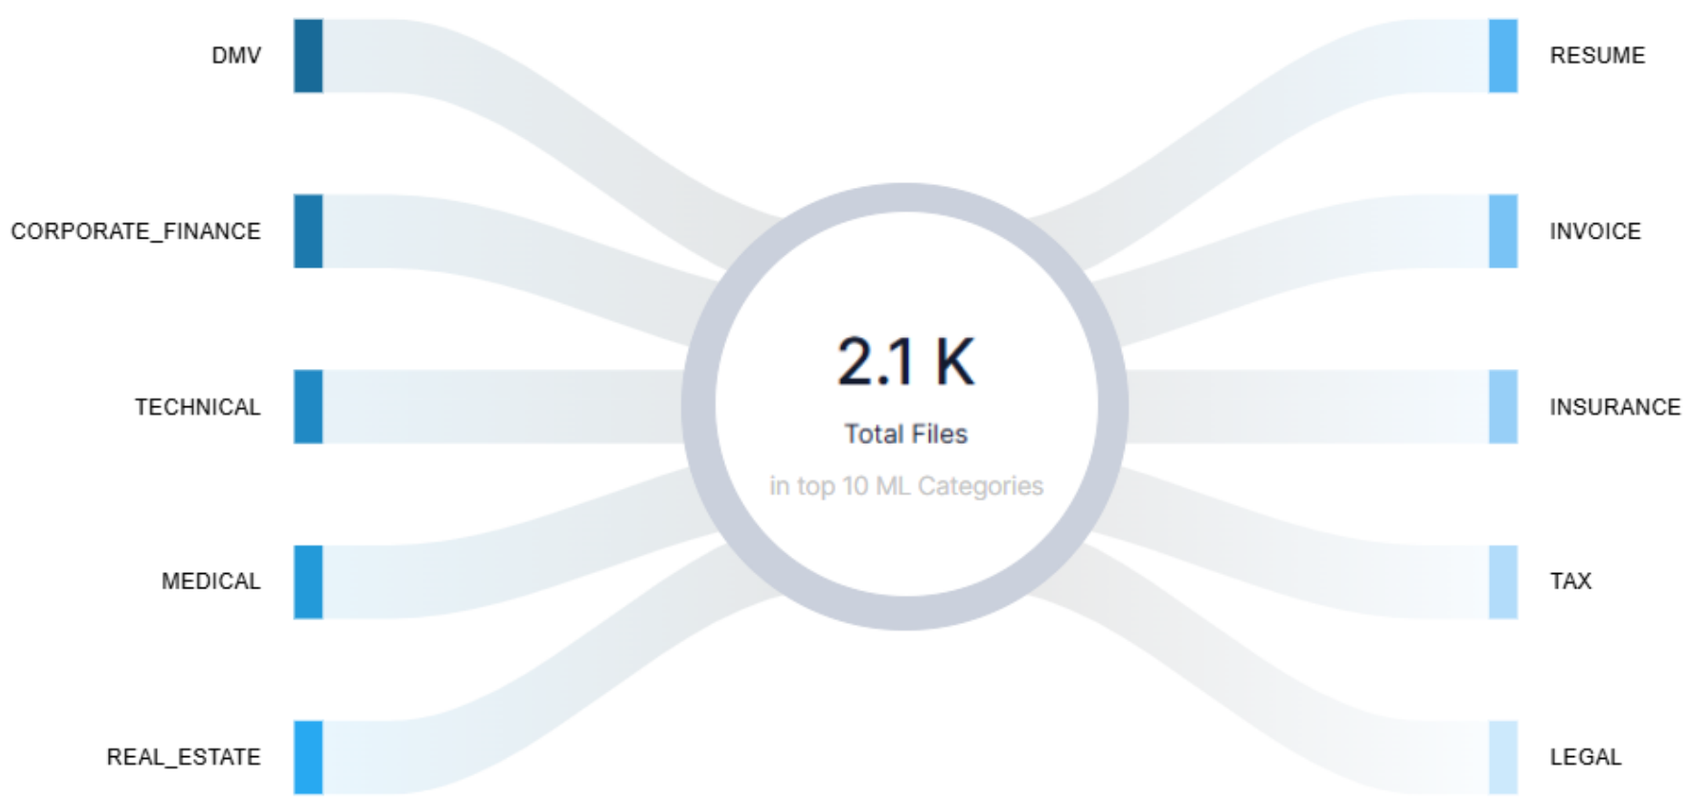

Data Captured from Incidents

User Input and Feedback

Refine and Tune | As Needed

Build Custom Dictionary Regex / Keyword

Single & multi word keywords with proximity

Expand EDM + IDM to data backup solutions

Real-time policy enforcement and granular visibility enable IT teams to secure sensitive data without added complexity or manual oversight. This streamlined approach facilitates confident, rapid adoption of GenAI tools, leveraging their productivity benefits while ensuring compliance and public trust, aligning with the "Never Trust, Always Verify" principle of Zero Trust.

Make DLP governance easy

One common challenge in implementing Data Loss Prevention (DLP), especially in large agencies or shared services organizations, is the volume of incidents that SOC teams and data owners need to manage. These incidents can range from requiring employee follow-ups for justification, reinforcing user training, handling exceptions, or maintaining an audit trail. Without an efficient system, this can quickly become overwhelming.

Workflow Automation simplifies this process by delivering a centralized solution for managing GenAI-related data protection incidents. It provides a complete view of all incidents in one place, including the metadata and details of the specific actions or data that triggered the violation. This centralization allows admins to quickly review, prioritize, and remediate incidents as needed.

A key feature of Workflow Automation is its ability to group incidents based on shared characteristics and assign priorities. These groups can then be allocated to specific admins for focused resolution. Automation plays a significant role here by enabling workflows that notify or train end users involved in incidents, request justifications, or escalate issues to managers or data owners for approval. Automated workflows can also trigger actions to remediate incidents without manual intervention.

By leveraging Workflow Automation in DLP, agencies can significantly cut down resolution times, reduce operational burdens on the SOC, and gain actionable insights into areas of risk. These insights can be used to further refine policies or enhance training programs, ensuring users are better equipped to operate securely while reducing the likelihood of future incidents.

Streamline Incident Management with Case Management & User Coaching

Easily manage, assign and escalate incidents

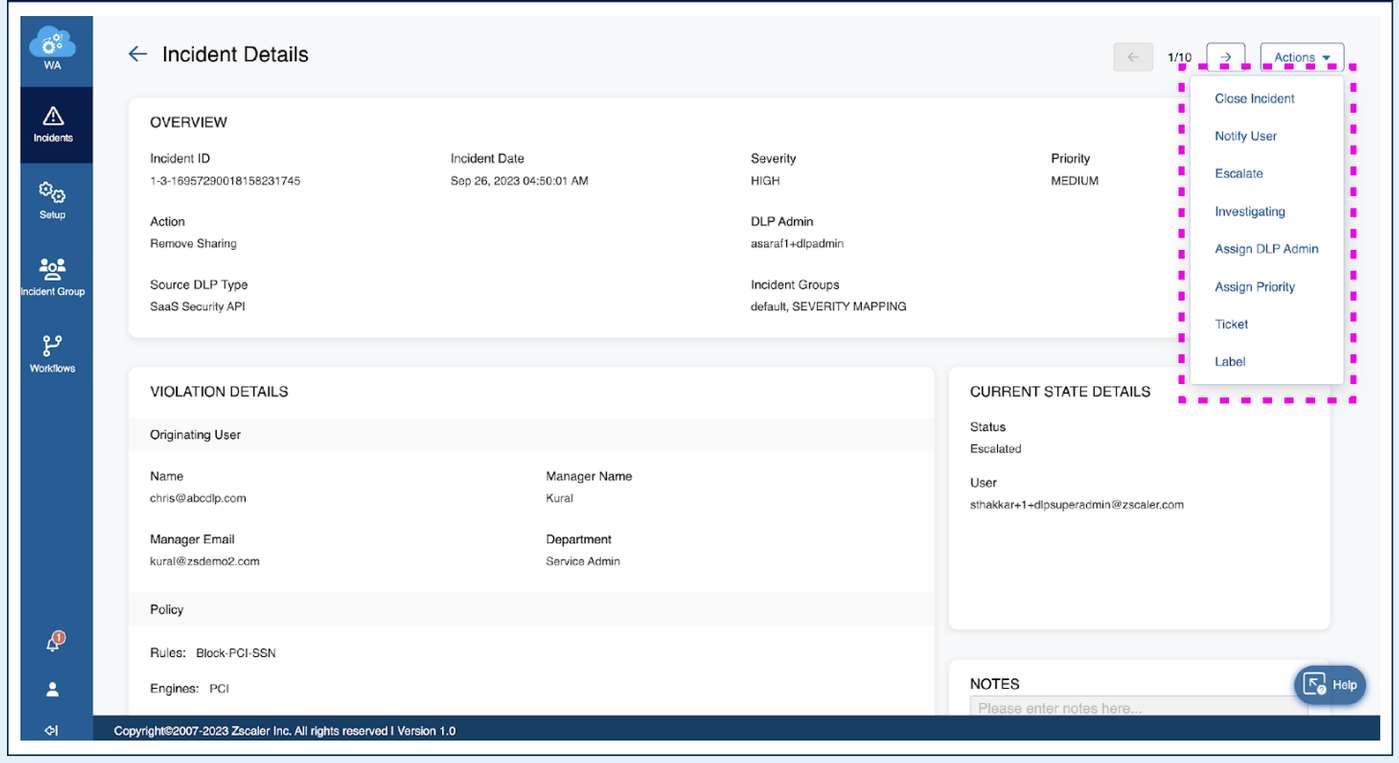

Define routines for workflows

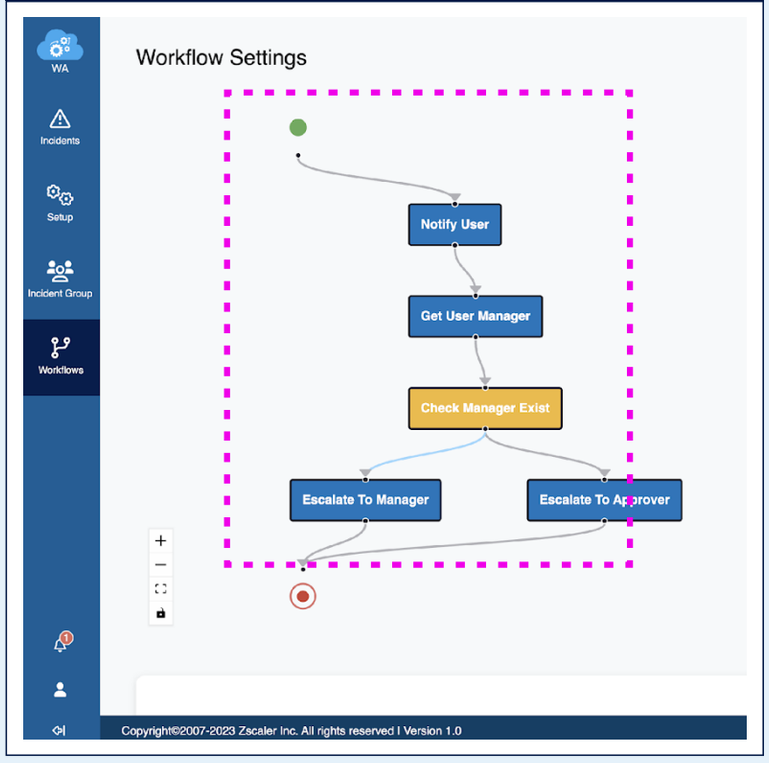

5. Bringing it All Together and Using a Layered Approach

State and local governments are adopting generative AI (GenAI) to unlock new efficiencies and improve services, but doing so securely is essential. With thousands of GenAI tools available, along with risks like data leakage and unsanctioned use, agencies need a clear strategy that prioritizes security, integrates Zero Trust principles, and still enables productivity. A layered approach simplifies this process by grouping applications based on risk, applying tailored security controls, and automating incident management to reduce pressure on IT teams. This strategy helps agencies protect sensitive data, streamline operations, and empower users to safely take advantage of GenAI applications all within a scalable and manageable framework.

Implement Layered Controls

In this section, we'll explore how agencies can bring together the various elements of secure GenAI adoption using a layered approach. With thousands of GenAI tools already available and new ones launching every week, managing policies and incidents can quickly become overwhelming without a well-thought-out strategy.

A layered approach simplifies this process by organizing access and implementing data controls tailored to risk levels. This method not only reduces the workload for security administrators but also significantly minimizes data leakage risks and cuts down the number of incidents that IT and security teams need to address. By taking this structured approach, organizations can securely and effectively harness the power of GenAI while maintaining operational efficiency.

As discussed earlier, tools like Shadow IT app discovery, GenAI discovery reports, and GenAI prompt visibility provide valuable insights into how AI policies should evolve and how security controls can be customized to meet changing needs. These insights form the foundation for a practical, layered approach to managing GenAI applications.

A useful way to implement this approach is by categorizing GenAI applications into three buckets: high-risk, medium-risk, and low-risk. High-risk applications should be blocked outright to prevent exposure to unnecessary vulnerabilities. Medium-risk applications can be accessed with heightened security controls, such as browser isolation and stricter data protection measures. Low-risk applications can be allowed native access but with restrictions focused on the specific content or actions users can take.

Layered Approach to Securing AI Applications

This structure allows agencies to adopt a Zero Trust approach to GenAI. Under this model, unknown, newly released, or unapproved applications are blocked by default. Approved but unsanctioned applications are isolated with additional security layers, while fully sanctioned applications benefit from a more seamless user experience with tailored safeguards. To make this easier to implement and manage, agencies can use tools like custom application labels and risk profiles. These allow security teams to define preset policies that automatically apply to applications based on their assigned risk. By simply labeling an application, the appropriate policies are enforced, minimizing administrative effort while maintaining robust control.

Automating Incidents workflows

Another critical layer to consider is incident management. It's essential for agencies to reduce the number of incidents that Security Operations Center (SOC) or data administrators must handle manually. Medium and low–severity violations, for example, should be logged for audit purposes and automatically closed without requiring significant manual intervention. However, since these still represent policy breaches, users should be notified and asked for justification—a step that is invaluable in reinforcing user training and fostering accountability.

With Zscaler, content inspection policies for GenAI allow agencies to define the severity level of violations, which are then passed to workflow automation tools. This feature enables admins to design workflows tailored to the severity of each incident. Additional attributes like severity and other shared characteristics can be used to categorize incidents into groups, and these groups can be tied to automated workflows. This approach simplifies how incidents are processed, ensuring that violations are addressed appropriately while significantly easing the burden on SOC teams.

Closing Thoughts

Government agencies need to be at the forefront of leveraging generative AI (GenAI) applications to transform operations, empower employees, and better serve citizens. However, its adoption must be underpinned by a Zero Trust architecture. By ensuring that every user, device, and interaction is verified, monitored, and controlled—no matter the location or application—agencies can confidently secure GenAI initiatives with strong data protection, clear governance, and streamlined user experiences at the core of their strategy.

Zscaler enables government agencies to embrace the productivity benefits of GenAI with a secure, layered approach that simplifies governance, streamlines deployment, and embeds robust security into every interaction. By establishing AI governance frameworks, automating GenAI application discovery and management, controlling use of GenAI application instances, and implementing advanced DLP capabilities from the start, agencies can dramatically reduce risks and scale their adoption strategies with minimal burdens on IT and security teams.

As the GenAI landscape continues to evolve, agency leaders are encouraged to take a strategic, phased approach to adoption. Start by securing access to public genAI applications, Securely unlock greater productivity with Agentic AI (future document). Lastly, we will explore how to securely extend GenAI capabilities to citizen-centered services, ensuring systems remain secure at every step. With Zscaler, agencies can implement these phases with confidence, accelerating innovation while maintaining the highest standards of data security and compliance.

Please reach out to your account team or contact us to schedule a workshop specific to your organization.

About Zscaler

Zscaler (NASDAQ: ZS) accelerates digital transformation so customers can be more agile, efficient, resilient, and secure. The Zscaler Zero Trust Exchange™ platform protects thousands of customers from cyberattacks and data loss by securely connecting users, devices, and applications in any location. Distributed across more than 150 data centers globally, the SSE-based Zero Trust Exchange™ is the world's largest in-line cloud security platform. Learn more at **zscaler.com** or follow us on Twitter **@zscaler**.

+1 408.533.0288 **Zscaler, Inc. (HQ) • 120 Holger Way • San Jose, CA 95134** **zscaler.com**

zscaler™

Zero Trust Everywhere