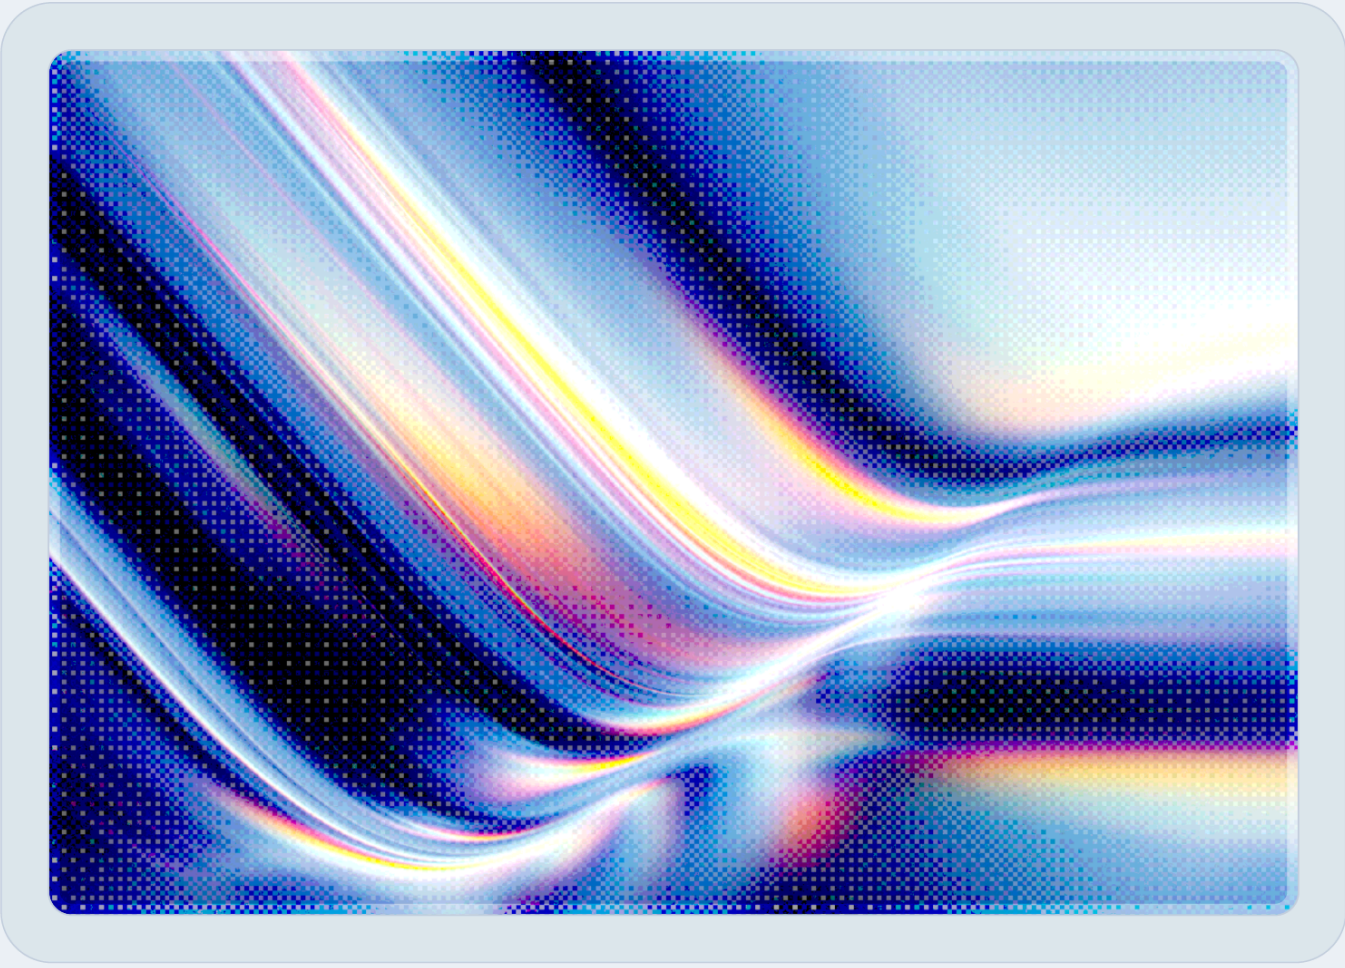
CSAI Foundation | Cloud Security Alliance

Sanctioned Channels Only: Sweden's Government Block on Direct ChatGPT and Claude Access

Security Implications of Routing Public-Sector AI Use Through Contracted Platforms

2026-11-02

AI-assisted Rapid Research



CSAI

CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

On 26 October 2026, the Government Offices of Sweden (Regeringskansliet), the body that supports the prime minister and the ministries, blocked browser access to OpenAI's ChatGPT and Anthropic's Claude from staff work computers [1][2][3]. The measure is not a ban on the underlying models. Staff were directed to Microsoft Copilot, which the Government Offices approved after a security review, and some employees retain access to OpenAI and Anthropic language models through Intric, a Stockholm-based AI assistant platform that won a Government Offices procurement in July 2026 [1][3][4]. The decision is therefore better understood as a control on *channels and contracts* than as a judgment about specific models.

The stated rationale centers on the absence of contractual agreements rather than on any technical flaw in the models. The Government Offices' director of digitalisation warned that AI services used without an agreement in place could have "consequences for our information security" [2]. The move is also not isolated: Sweden's Tax Agency (Skatteverket) and the National Board of Institutional Care (SiS) had already introduced similar restrictions, and the Swedish Enforcement Authority (Kronofogden) is reported to be considering the question [1][3].

The underlying policy pattern offers a reusable template for organizations handling regulated or confidential information. It consists of blocking consumer endpoints, approving a reviewed enterprise tool, and routing remaining frontier-model use through a contracted intermediary. Blocking alone, however, is necessary but not sufficient. Organizations adopting this pattern should verify where inference actually occurs on sanctioned platforms, close residual paths such as personal devices and API keys, and give staff a sanctioned option capable enough that they do not route around controls.

Background

What the Government Offices Did

Sveriges Radio's news service Ekot first reported the plan on 23 October 2026, and SVT confirmed it the same day through an emailed statement from the Government Offices [1][3]. According to that statement, the web browser interfaces for ChatGPT and Claude would no longer be reachable from Government Offices work computers from Monday, 26 October. The email also stated that all employees

have access to Microsoft's AI tool and that, "within the framework of Intric," employees still have access to the ChatGPT and Claude language models [1]. The Government Offices described Copilot as having been reviewed and having had its security raised before approval [1][3].

Euronews reported that the decision followed a review earlier in 2026 of AI tool use at the Swedish Transport Administration, which found "problematic" uses of AI platforms in government work [2]. Officials were concerned that sensitive information entered into services without a data protection agreement could be retained and used for model training [2]. Mikael Wettercrantz, director of digitalisation at the Government Offices, characterized such unsanctioned platforms as shadow IT [2][3]. CSA was unable to locate a public copy of the Transport Administration review, so its findings are cited here only as described in press coverage. Neither OpenAI nor Anthropic had publicly commented on the decision at the time of reporting [3].

Coverage has been inconsistent on scope. Most outlets, including SVT, Euronews, and The Next Web, confine the block to the Government Offices plus the separately reported restrictions at the Tax Agency and SiS [1][2][3]. A small number of secondary sources described a ban spanning all central, regional, and municipal authorities; CSA found no primary-source support for that broader claim, and the Government Offices' own statement to SVT describes a measure limited to its own work computers [1]; CSA therefore treats the broader characterization as unsupported. Separately, Euronews reported that the Swedish government had already placed DeepSeek on its list of unapproved AI tools [2]. The measure also does not affect members of the public.

The Sanctioned Channels

The approved routes are as significant as the blocked ones. Copilot is covered by an agreement between Microsoft and the Swedish government and passed the Government Offices' security screening [2]. Intric announced on 8 July 2026 that it had won the Government Offices procurement for an AI platform on which employees can build and use AI assistants on their own data and under their own governance [4]. Intric states that the platform runs on Swedish infrastructure without dependence on foreign cloud services, that the agreement covers both cloud and on-premises installation, and that the Government Offices handle information that may be classified under Swedish public access and secrecy law [4]. Neither the Government Offices nor Intric has publicly disclosed how many employees can reach OpenAI and Anthropic models through Intric, or how those model calls are hosted [3][4].

Political and Institutional Context

The decision lands in a country where government use of consumer chatbots has already been publicly contested. In August 2025, Prime Minister Ulf Kristersson told the business daily Dagens industri that he used tools including ChatGPT and Mistral's Le Chat "quite often" for a second opinion, prompting criticism from AI researchers and commentators; his office responded that security-sensitive information was not used [5]. Internationally, the Swedish measure follows an established pattern. In 2023, employers including Apple, Samsung, and several large banks restricted staff use of ChatGPT over confidentiality concerns [6], and the U.S. Space Force temporarily paused use of web-based generative AI tools pending a data security review [7]. What distinguishes the Swedish case is that it pairs the block with explicitly named, contracted alternatives that still provide access to the same frontier models.

Security Analysis

The Risk Is the Contract, Not the Model

The most important analytical point is that the Government Offices did not conclude that ChatGPT or Claude models are unsafe. Staff can still reach those same model families through Intric [1][3]. What differs between the blocked and permitted paths is the legal and operational wrapper around the data. A browser session on a consumer or individually purchased account is governed by the provider's consumer terms, which the organization neither negotiated nor can audit. Those terms can change: in 2025, Anthropic updated its consumer terms so that Claude Free, Pro, and Max users would have chats and coding sessions used for model training unless they opted out, with data retention extended to five years for users who allowed training [8]. Commercial and enterprise offerings from major providers are generally subject to different terms, under which customer business data is by default not used for training [8][9]. An individual civil servant is poorly placed to verify which regime applies to a given session, which account tier is in use, or what settings are enabled.

For a government body, the exposure extends beyond training. The Government Offices handle information that may be classified under the Public Access to Information and Secrecy Act, a point Intric explicitly ties to the scope of its procurement [4], and disclosure to an external service provider without appropriate arrangements can itself raise legal questions. A procured contract can specify data location, retention, sub-processors, incident notification, and audit rights; a personal browser session cannot. Seen this way, the block functions as an enforcement mechanism for a procurement decision: data may flow to AI only where a contract exists that the organization's lawyers and security staff have reviewed.

Comparing the Three Access Paths

The table below summarizes how the three paths described in public reporting differ on the dimensions most relevant to security teams. Cited entries for the sanctioned paths reflect public statements rather than independent verification. Uncited entries reflect typical enterprise deployment characteristics rather than disclosed details of the Government Offices' configuration.

Dimension	Direct browser access (blocked)	Microsoft Copilot (approved)	Intric platform (limited access)
Contractual basis	Individual consumer or self-service terms	Agreement between Microsoft and the Swedish government [2]	Government Offices procurement awarded July 2026 [4]
Organizational review	None reported	Security review and hardened configuration [1][3]	Procured platform; review details not public
Training on inputs	Depends on account tier and user settings [8]	Typically governed by enterprise agreement (terms not public)	Governed by platform and upstream model agreements
Hosting and residency	Provider-determined	Not publicly detailed	Swedish infrastructure claimed for platform; upstream model hosting not disclosed [4]
Organizational visibility	Typically minimal; outside organizational identity and logging controls	Typically tenant-integrated (configuration not public)	Platform-level governance over assistants and data [4]
Who can use it	All staff (prior to block)	All staff [1]	Limited, undisclosed subset [1][3]

The Inference Location Question

The Intric arrangement raises a question that organizations adopting a similar model should ask explicitly. Intric states that its platform runs on Swedish infrastructure without depending on foreign cloud services [4], yet the Government Offices also states that employees can reach ChatGPT and Claude models through it [1]. These statements are compatible, because a platform can be hosted domestically while forwarding prompts to externally hosted models, but they describe different layers. If prompts and retrieved documents are sent to a model endpoint operated by, or on infrastructure contracted by, a U.S.-headquartered provider, the data crosses an organizational and potentially jurisdictional boundary at the moment of inference. Public sources do not reveal whether Government Offices traffic to these models goes to provider-operated APIs, to European-region deployments through a cloud marketplace, or to some other arrangement. CSA does not suggest that the arrangement is deficient; the point is that "sovereign platform" claims typically describe the orchestration layer, and security teams need contractual and technical confirmation of where the model layer runs, what it retains, and under which jurisdiction.

Residual Paths and the Displacement Problem

A network-level block on two consumer domains addresses one path among several. Staff may still reach generative AI through personal phones and home computers, through less prominent chatbot services not yet blocklisted, through browser extensions and productivity add-ons that embed model calls, or through API keys used in scripts. Reporting has not addressed whether personal devices are in scope [3]. Several large employers restricted ChatGPT in 2023 [6], and broader shadow IT experience suggests that blocking without a usable alternative tends to displace rather than eliminate usage. If that pattern holds, the Government Offices' provision of Copilot to all staff is arguably the more consequential half of the decision. The durability of the control will depend on whether the sanctioned tools are good enough for the work staff are trying to do. If they are not, usage may migrate to channels that are less visible than the browser sessions that were just blocked.

There is also a concentration effect worth noting. Making Copilot the default tool for all staff reduces exposure to unmanaged services but increases dependence on a single vendor's AI stack and configuration. The Intric path partially offsets this by providing a second, domestically hosted orchestration layer, and the Swedish government's earlier placement of DeepSeek on its list of unapproved AI tools suggests that approved lists in the Swedish public sector are actively curated [2]. Organizations replicating this approach should treat the approved list as a living register with its own review cadence rather than a one-time decision.

Implications Beyond Sweden

The Swedish measure is notable less for its novelty than for its clarity. It articulates a principle many organizations hold implicitly: AI use is permitted, but only through channels with a reviewed contract, organizational identity, and logging. The fact that the Tax Agency and SiS had already moved in the same direction, and that the Enforcement Authority is weighing it, suggests the approach may spread across Swedish public administration [1][3]. Similar reasoning applies to regulated private-sector organizations in finance, healthcare, and critical infrastructure, where confidentiality obligations attach to the data rather than to the tool. In CSA's assessment, because the restriction does not target any vendor's model, it is also unlikely to be resolved by vendor capability improvements alone. It is more likely to be resolved by vendors offering contractual terms that public bodies can accept.

Recommendations

Immediate Actions

Organizations handling confidential or regulated information should first establish an accurate inventory of which generative AI services their staff currently reach. Secure web gateway, DNS, and CASB logs typically reveal consumer chatbot traffic, and identity provider logs show which sanctioned tools are actually in use. Where consumer endpoints are blocked, the block should be accompanied by a named, supported alternative and a short explanation, as the Government Offices did by directing staff to Copilot when announcing the measure [1]. Security teams should also confirm the training and retention settings on any AI service already under enterprise contract, because consumer and commercial terms can differ materially [8][9].

Short-Term Mitigations

Within the following quarter, organizations should formalize an approved AI services register that records, for each sanctioned channel, the contractual basis, the data classifications permitted, the location of the orchestration and inference layers, retention and training terms, and the review owner. For intermediary platforms that broker access to third-party models, procurement and security reviews should explicitly request disclosure of upstream model endpoints, regions, and sub-processor agreements, rather than relying on the platform's own hosting claims. Routing model access through a central AI gateway — whether a commercial platform or an internal service — allows consistent identity binding, prompt and response logging, data loss prevention, and rapid revocation of a provider if its

terms or legal status change. Organizations should also extend controls to residual paths by restricting AI-related browser extensions, monitoring for model API keys in code repositories, and setting clear policy for personal-device use with official information.

Strategic Considerations

Over the longer term, the Swedish case argues for treating AI access as a procurement and identity problem rather than a URL-filtering problem. Blocklists tend to trail the growing number of available services; an allowlist approach that permits only contracted, identity-bound channels avoids that maintenance burden, though it requires a functioning approval process. Leadership should weigh the concentration risk created when a single enterprise assistant becomes the default for all staff against the shadow AI risk the block is meant to reduce, and should consider maintaining at least two independently contracted model paths for resilience. Finally, organizations should invest in the usability of sanctioned channels and in staff training on data classification, since the long-run effectiveness of any block depends on employees having a legitimate route that meets their needs. Public bodies, in particular, may benefit from shared procurement frameworks so that smaller agencies can adopt sanctioned channels without each negotiating terms independently.

CSA Resource Alignment

The behavior the Government Offices is trying to curb maps closely to the problem CSA's Identity and Access Management Working Group describes in **Shadow Access and AI** [10]. That report examines how AI can both create and help address Shadow Access, including the unintended and ungoverned access paths that arise as AI is adopted. Direct browser use of consumer chatbots is a clear instance: organizational data reaches an external service under an identity the organization does not manage and through a path its access controls do not observe. Restricting AI use to contracted, identity-bound channels, as Sweden has done, is a practical application of the report's emphasis on bringing AI-related access back under governance.

The companion guide **Confronting Shadow Access Risks: Considerations for Zero Trust and Artificial Intelligence Deployments** [11] provides the architectural framing for the sanctioned-channel model. Its Zero Trust orientation — continuous verification, least privilege, and explicit policy at each access point — corresponds to routing AI traffic through a gateway where identity, data classification, and destination can be evaluated per request. The inference-location question raised above is a Zero Trust question at heart: trust in the orchestration platform should not be extended implicitly to the upstream model endpoint without separate verification.

For control selection, CSA's **AI Controls Matrix (AICM) v1.1** [12] offers the most direct mapping. Its domains covering identity and access management, data security and privacy, supply chain management, and governance and risk management address the elements of an approved AI services register: vetted providers, documented data handling terms, access controls on model interfaces, and monitoring of AI usage. Organizations can use AICM's shared-responsibility orientation to determine which controls they must implement themselves and which they should require contractually from intermediary platforms and model providers. **AI Organizational Responsibilities: Core Security Responsibilities** [13] complements this by clarifying which internal functions own AI data governance and vendor oversight decisions, a question that becomes concrete when an organization must decide which channels to sanction and who reviews them.

References

[1] SVT Nyheter (G. Motte and M. Ståhl). "AI-tjänster spärras på Regeringskansliet." SVT, 23 October 2026 (updated 24 October 2026).

[2] Benton, J. "Swedish government blocks ChatGPT and Claude. Here's why." Euronews, 26 October 2026.

[3] Stan, A. M. "Sweden's Government Offices block ChatGPT and Claude on work computers." The Next Web, 26 October 2026.

[4] Intric. "Intric wins procurement with the Swedish Government Offices for sovereign AI platform." Intric, 8 July 2026.

[5] The Guardian. "'We didn't vote for ChatGPT': Swedish PM under fire for using AI in role." The Guardian, 5 August 2025.

[6] McGlauflin, P. "Apple, Goldman Sachs, and Samsung among growing list of companies banning employees from using ChatGPT at work." Fortune, 19 May 2023.

[7] Freedman, L. F. "Web-Based Generative AI Tools Use Temporarily Prohibited by U.S. Space Force." The National Law Review, 12 October 2023.

[8] Anthropic. "Updates to Consumer Terms and Privacy Policy." Anthropic, August 2025.

[9] OpenAI. "Enterprise privacy at OpenAI." OpenAI, accessed 2 November 2026.

[10] Cloud Security Alliance. "Shadow Access and AI." CSA Identity and Access Management Working Group, 2025.

[11] Cloud Security Alliance. "Confronting Shadow Access Risks: Considerations for Zero Trust and Artificial Intelligence Deployments." CSA, 2024.

[12] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2025.

[13] Cloud Security Alliance. "AI Organizational Responsibilities: Core Security Responsibilities." CSA, 2024.