CSAI Foundation | Cloud Security Alliance

# EU Mandates System-Level Android Access for Rival AI Assistants

## Security and Identity Risks in the DMA's Interoperability Ruling

2026-07-17

AI-assisted Rapid Research

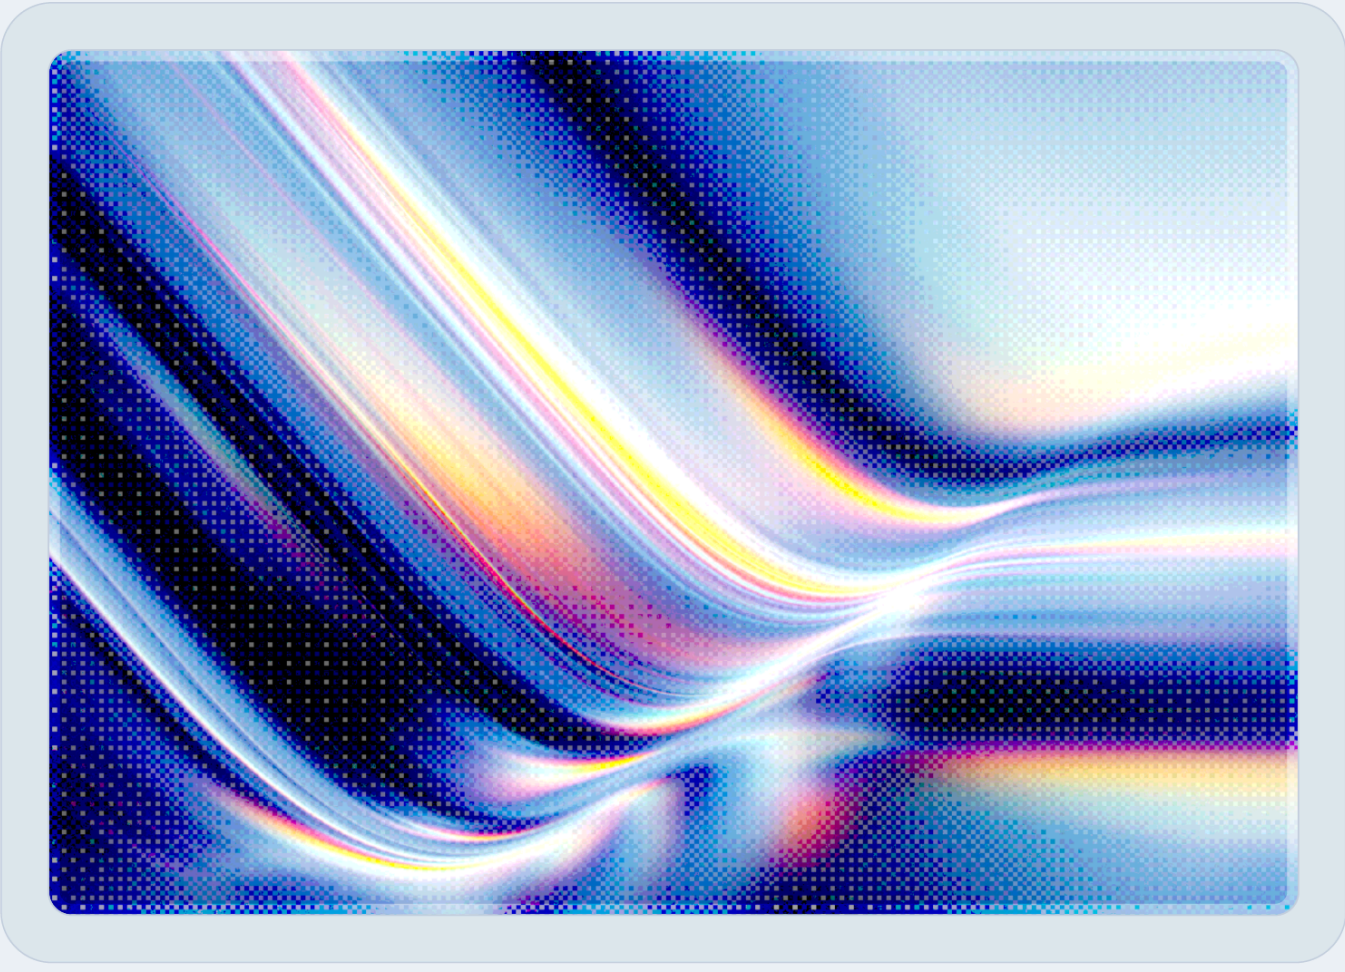

CSAI

CSA cloud security alliance®

## Key Takeaways

- On July 16, 2026, the European Commission issued binding specification decisions under Article 6(7) of the Digital Markets Act requiring Google to give rival AI assistants the same Android system access that Gemini currently holds exclusively, spanning four capability layers: invocation (wake words and system-wide entry points), context (screen contents, sensor data, and on-device app data), actions (screen automation and cross-app control), and resources (on-device models and background execution) [1][2].
- The decision effectively converts a set of highly sensitive device capabilities — camera, microphone, location, screen content, notifications, contacts, SMS, app-launch data, and cross-app databases — into a mandated, standing interoperability surface that any qualifying third-party AI assistant can request equal access to, rather than access Google alone controls and vets [2][3].
- Google must make most capabilities available with the Android 18 release by August 1, 2027, with concurrent wake-word support following in Android 19 by August 1, 2028; five of the most sensitive capabilities additionally require compliance with "objective and non-discriminatory" security, privacy, and integrity criteria, verified by independent parties alongside Google's own certification process, due by May 1, 2027 [2][3].
- Google has publicly objected that the ruling removes the OEM-level vetting that today filters which apps can request these permissions, arguing this "risks undermining vital privacy and security guardrails for millions of Europeans," while the Commission maintains that security and integrity were built into the decision from the outset [4][5].
- The structure of this mandate — multiple third-party AI agents granted broad, standing, device-level access with consent-based rather than vetting-based gating — closely parallels the over-privileging and attribution gaps CSA's own research has documented in enterprise agentic AI deployments, now applied at the scale of the entire EU Android user base.

# Background

Google has held gatekeeper status for Android under the Digital Markets Act since the regulation's core obligations took effect, and the Commission has spent the past year building out the specific technical measures that operationalize Android's general interoperability obligation for AI services [1]. Following a call for feedback in April 2026 on how to ensure interoperability with Android for AI assistants, the Commission published two related decisions on July 16, 2026: one requiring Google to share anonymized Google Search ranking, query, and click data with rival search engines beginning January 2027, and a second requiring Google to open Android's AI-relevant system capabilities to competing assistants [1][6]. This research note focuses on the second decision.

Today, Gemini benefits from integration points that are structurally unavailable to any app a user installs from the Play Store. It can be summoned by a dedicated wake word processed at the operating system's audio layer, activated from the long-press home button and navigation handle, and it can read on-screen content and take actions across other apps and system surfaces through system-level entry points that are unavailable to apps installed from the Play Store. The Commission's decision requires Google to extend comparable capabilities to any AI assistant that qualifies under the DMA's criteria: a third-party assistant will be able to register its own wake word and claim the same system-wide invocation surfaces, read screen content and sensor data on the same basis Gemini does, and execute structured actions across other apps — including OEM and Google's own apps — in the background [2][3]. The Commission has directed that this access be provided free of charge, documented, supported with testing and technical assistance, and made available to third parties on the same timeline Google makes new capabilities available to itself; access cannot be made cumbersome or conditioned on a rival assistant becoming the device's default [3].

Google's public response has centered on the trade-off between competition and security. Kent Walker, Google's head of global affairs, said the changes risk "undermining vital privacy and security guardrails for millions of Europeans," and the company has separately argued that the search-data-sharing portion of the same decision could expose user query data to companies outside EU data-protection jurisdiction without adequate anonymization or consent [4][5]. A senior EU official countered that the Commission "took integrity, security and privacy into utmost account" in designing the decision [5]. The decision does build in a security-certification requirement for the most sensitive capabilities, though the specifics of that regime are not yet public. Whether the loosening of Google's vetting authority represents an acceptable trade-off or an unaddressed gap will depend on how that certification regime is ultimately defined.

---

# Security Analysis

## A Regulator-Mandated Expansion of the Attack Surface

The practical effect of this decision is to take a set of device capabilities that security teams typically treat as maximally sensitive — always-on microphone access, camera access, precise location, full notification content, contact lists, SMS content, and the ability to read whatever is currently on screen — and convert them from a single-vendor integration (Gemini, built and reviewed by Google) into a standing interoperability surface that must be opened to any qualifying third party. Where an enterprise IAM team would normally treat a request for this combination of permissions as an extreme case requiring individualized risk review, the DMA decision requires it be granted categorically, to any assistant that meets the qualification criteria, across the entire population of EU Android devices. The number of entities with device-level sensor and screen access on a typical phone is set to grow from one (Google) to however many AI assistants a user chooses to install, each independently capable of camera, microphone, location, and cross-app action execution.

## Weakening the Existing Vetting Boundary

Google's specific objection — that OEMs today vet which apps can request Android's most sensitive permissions, and that this ruling removes that vetting layer — appears to describe a control being displaced, a characterization the Commission has not directly contested. The decision's proposed replacement is a certification regime: five of the most sensitive capabilities (screen automation, structured app integration, system integration, centralized app-data access, and context-aware intelligence) require an assistant to meet "objective and non-discriminatory" security, privacy, and integrity criteria, checked by independent parties in addition to Google's own certification, with Google's certification process due by May 1, 2027 [2][3]. This proposes to replace vetting with certification rather than eliminating oversight altogether, but as of this writing the substance of those criteria has not been published, meaning the industry does not yet know what technical bar a rival assistant must clear before receiving camera, microphone, and cross-app access equal to Gemini's. CSA's own research on shadow access has warned that governance and certification requirements which are not backed by identity-centric, continuously enforced technical controls tend to function as compensating stopgaps rather than genuine barriers to unauthorized access [7], and the DMA decision's current structure — a certification deadline nearly a year out, with criteria still to be defined — fits that pattern closely.

## Consumer-Scale Version of a Problem CSA Has Already Measured

CSA's 2026 survey on AI agent identity and access, drawn from 228 security and IT professionals, found that 74 percent of respondents agree AI agents typically receive more access than operationally necessary, that 68 percent cannot clearly distinguish AI agent activity from human activity in their own environments, and that 81 percent believe prompt manipulation could cause an AI agent to reveal sensitive credentials or tokens it holds [8]. Those findings describe enterprise environments where organizations at least retain the ability to inventory, monitor, and revoke agent access centrally. The DMA decision reproduces the same structural conditions — multiple AI agents holding broad, standing, sensor-and-action-level device access, without a single owner enforcing least privilege — but distributes it across hundreds of millions of individual consumer devices where no equivalent centralized IAM function exists. A malicious actor who can manipulate a third-party assistant's inputs (through a crafted webpage, a poisoned notification, or a manipulated app) inherits whatever standing access that assistant has already been granted; at consumer scale, with weaker vetting and no enterprise-style monitoring layer, the same attribution and containment gaps CSA's survey measured in corporate environments are likely to recur, and with far less recourse for the affected user.

## Cross-App Action Execution Raises the Stakes Beyond Data Exposure

Unlike a simple data-sharing obligation, the "actions" capability layer permits a qualifying assistant to imitate user interactions, read a controlled app's screen content, discover installed apps, and execute control in the background across third-party, OEM, system, and Google apps [2]. This moves the risk beyond passive data exposure into the territory of autonomous action: a compromised or maliciously designed third-party assistant with this level of access could, in principle, initiate transactions, send messages, or alter settings on a user's behalf without the same friction a conventional app installation would face. This is precisely the class of risk CSA's Context-Based Access Control guidance was written to address — the recognition that static, binary permission grants are inadequate once a system, whether an enterprise identity or a device-level AI agent, can act with the delegated authority of its user across multiple downstream systems, and that continuous, contextual risk evaluation is needed in place of a one-time consent screen [9].

---

# Recommendations

## Immediate Actions

Enterprises that permit employee use of Android devices, whether corporate-issued or BYOD, should begin tracking this decision now rather than waiting for the 2027 implementation date, since mobile device management policies that currently assume Google's own vetting is the only gate on system-level AI access will need to be revisited once that assumption no longer holds. Security teams should specifically inventory which third-party AI assistants employees have installed or are likely to install, and flag any assistant requesting camera, microphone, location, notification, contact, or cross-app permissions for the same scrutiny currently reserved for apps requesting device administrator privileges. Organizations should also begin monitoring the European Commission's DMA developer portal for the Alphabet specification proceedings, since the specific technical and security certification criteria referenced in the decision will be published there as they are finalized, and those criteria will be the only defined technical safeguard against the risks described above.

## Short-Term Mitigations

Enterprise mobile security programs should extend their existing zero-trust and conditional-access frameworks to treat AI-assistant permission grants on managed Android devices as a distinct, continuously evaluated risk category rather than a one-time installation decision, applying the same static- and dynamic-signal evaluation CSA's Context-Based Access Control guidance describes for other identity types. Where an organization's mobile device management platform supports it, security teams should require that any third-party assistant requesting Gemini-equivalent system access first demonstrate the independent security certification the DMA decision calls for, once that certification framework exists, before allowlisting the app for corporate devices. Organizations should also apply the identity-separation and attribution practices CSA recommends for enterprise AI agents — maintaining clear logs of which actions on a device were initiated by a human user versus an AI assistant acting with delegated permissions — to reduce the same accountability gap that CSA's identity and access survey found already affects two-thirds of enterprise AI agent deployments.

## Strategic Considerations

The certification criteria the European Commission and Google are expected to finalize over the next year, covering the five most sensitive Android AI-interoperability capabilities, will likely become a reference point beyond this single ruling: other regulators evaluating similar interoperability mandates,

and other mobile platform vendors facing comparable obligations, are likely to look to whatever technical bar is set here. CSA and its members have a window to engage with that process while the criteria are still being drafted, ahead of the May 2027 certification deadline, rather than responding after a framework is already locked in. More broadly, this decision is a useful forcing function for security teams to internalize a lesson CSA's enterprise research has already established: governance and consent requirements alone do not substitute for identity-centric, least-privilege, continuously monitored access controls once multiple autonomous AI agents hold standing access to sensitive systems. The DMA ruling applies that exact pattern to the Android ecosystem at a scale — hundreds of millions of consumer devices — that makes the underlying technical and policy questions considerably harder to walk back if the initial safeguards prove insufficient.

---

## CSA Resource Alignment

CSA's "Confronting Shadow Access Risks: Considerations for Zero Trust and Artificial Intelligence Deployments" is the most directly applicable prior CSA publication. Its core argument — that unintended or unmonitored access pathways expand fastest exactly where governance requirements substitute for enforced, identity-centric controls — describes the structural risk in the DMA decision's current form: a certification regime for the five most sensitive Android capabilities that is not yet technically defined, sitting alongside a broad standing grant of camera, microphone, location, and cross-app access to qualifying third parties.

CSA's "Identity and Access Gaps in the Age of Autonomous AI" supplies the empirical baseline for the risks this note describes at consumer scale: its finding that 74 percent of security professionals believe AI agents typically receive more access than necessary, and that 68 percent cannot clearly attribute agent actions versus human actions, describes exactly the attribution and over-privileging problem that multiple co-resident, system-level Android AI assistants will reproduce for the general public once the DMA's Android 18 and Android 19 timelines take effect.

CSA's "Context-Based Access Control for Zero Trust" offers the technical vocabulary and architecture — continuous, contextual risk evaluation in place of a static, one-time consent decision — that both enterprise mobile security teams and, potentially, the certifying bodies the Commission references should draw on when defining what "objective and non-discriminatory security, privacy, and integrity criteria" should actually require of a third-party assistant requesting Gemini-equivalent Android access.

# References


[1] European Commission. "Commission provides guidance to Google for AI interoperability on Android and sharing of Google Search data under the Digital Markets Act." Digital Markets Act (DMA), July 16, 2026.

[2] Techi. "EU Opens 11 Android Features to Rival AI Assistants." Techi, July 16, 2026.

[3] 9to5Google. "EU demands opening Android access and Search data to rivals, Google warns of privacy risks." 9to5Google, July 16, 2026.

[4] Bloomberg. "Google Must Give Gemini Rivals Equal Access to Android System, EU Says." Bloomberg, July 16, 2026.

[5] CNBC. "Google required to open up to AI, search engine rivals under EU-mandated changes." CNBC, July 16, 2026.

[6] Tech Times. "EU Gives Rival AI Assistants System-Level Android Access Google Reserved for Gemini." Tech Times, July 16, 2026.

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

[8] Cloud Security Alliance. "Identity and Access Gaps in the Age of Autonomous AI." Cloud Security Alliance, 2026.

[9] Cloud Security Alliance. "Context-Based Access Control for Zero Trust." Cloud Security Alliance, 2025.


 