VendRisk360
ASSESS · MONITOR · STAY AHEAD

WHITEPAPER

SaaS-to-SaaS Breaches: The Integration Tokens Nobody Is Watching

The fastest-growing path into an organization is not the front door. It is the OAuth token connecting one vendor to another. A due-diligence and monitoring guide for app-to-app risk.

Published 27 September 2026

Written for Heads of TPRM, CISOs, identity and SaaS security leads, procurement, internal audit

Online www.vendrisk360.com/resources/whitepapers/saas-to-saas-integration-risk/

Contents

How SaaS-to-SaaS integrations and non-human identities became a leading breach path, what recent incidents teach, and how to bring OAuth connections, integration tokens and their scopes into vendor due diligence, contracts and continuous monitoring.

When a vendor is breached, the instinct is to ask what happened inside their environment. Increasingly, the more important question is what that vendor was connected to. Modern SaaS products are woven together by integrations: your marketing platform reads your CRM, your CRM syncs to your data warehouse, your support desk posts to your chat tool, an analytics add-on holds a token into all of them. Each connection is authorized once, usually through OAuth, and then persists silently, often with broad access, frequently forgotten. These connections are now a leading way that a compromise of one vendor becomes a compromise of your data held in another.

The uncomfortable feature of this risk is that the breached system is often not yours and not even your primary vendor's. An attacker who obtains the tokens held by a small, connected add-on can read the data those tokens unlock in a major platform, without ever touching the platform's own defenses. The primary vendor's security posture, the one you assessed, was never the weak point. This paper treats SaaS-to-SaaS integration as a first-class category of third-party risk: why it has grown, what recent public incidents teach, and how to bring integration tokens and their scopes into due diligence, contracts and monitoring. It describes the risk and the defenses; it is not an attack guide.

KEY TAKEAWAYS

- The connection is the exposure. An integration authorized once can hold standing access to your data for years. When the connected app or its provider is compromised, that access is what the attacker inherits.
- Non-human identities now vastly outnumber human ones. Tokens, service accounts, API keys and connected apps are the majority of identities in most estates, and they are the least governed. Vendor research consistently puts machine identities many times above human accounts.
- Recent breaches followed the integration, not the front door. Several of the most consequential 2024 and 2025 incidents spread through stolen OAuth tokens and connected applications rather than through the primary platform's own controls.
- Scope is everything. An integration with read-write access to all records is a different risk from one scoped to a single object. Most are over-scoped because the broad option was the easy one at setup.
- Standing tokens rarely expire on their own. Unlike a password a leaver loses, an integration token keeps working until someone revokes it. Offboarding a vendor must include revoking its connections.
- This risk is assessable. Ask which of your systems a vendor connects to, with what scopes, how its tokens are stored and rotated, and how it would tell you if they were exposed. Then monitor for the connections you did not authorize.

Why app-to-app risk grew so fast

Three forces combined. SaaS adoption multiplied the number of platforms holding your data. Every platform built an app marketplace and an API, so connecting them became a few clicks. And the connection model, OAuth, was designed for convenience: authorize once, and the connected app receives a token it can use repeatedly without a human present. That is exactly what makes integrations useful and exactly what makes them risky. The token is a standing credential, held by a third party, that does not need your user to be logged in, does not trigger your multi-factor prompts, and does not expire when your employee changes their password.

The governance did not keep pace. Organizations invested heavily in governing human identities, with single sign-on, multi-factor authentication and joiner-mover-leaver processes. The non-human identities, the tokens and service accounts and connected apps, grew faster and were governed less. Industry research repeatedly finds machine identities outnumbering human ones by a large multiple, and most organizations cannot produce a complete inventory of the connected applications holding tokens into their core systems. An asset you cannot list is one you cannot assess.

What recent incidents teach

The lesson of the last two years is consistent: attackers followed the integration. The public record includes several incidents that turned on connected-app tokens rather than on breaking the primary platform. Descriptions here are at the level of public reporting and are included for the control lessons; organizations should verify details against primary sources.

INTEGRATION-LINKED INCIDENTS AND THE CONTROL EACH POINTS TO

Pattern (as publicly reported)	Control lesson
OAuth tokens held by a connected sales-engagement/chat integration were stolen and used to reach many customers' CRM data	An add-on's tokens can unlock a major platform; assess and monitor the connected app, not only the platform
A vendor's support-system compromise exposed session artifacts that were then used against that vendor's customers	A vendor incident can flow downstream through tokens; require prompt notice and rotate on any exposure
A legacy or test OAuth application with broad permissions was abused to reach corporate email	Dormant and over-scoped connected apps are a standing liability; inventory and remove them
A developer-platform breach forced mass rotation of customer secrets and tokens stored there	Secrets a vendor holds on your behalf are your exposure; know what they store and how it rotates
Analytics and BI connectors with warehouse tokens were implicated in customer data access	Data-warehouse connectors are high-value; scope them tightly and monitor their use

None of these required defeating the primary platform's own security. In each, a token or connected app was the path. That is the shift a third-party risk program has to absorb: the vendor you assessed may be sound, and the exposure may still run through what it is connected to.

THE BREACHED SYSTEM IS OFTEN NOT THE ONE THAT LEAKS YOUR DATA

In an integration breach, the compromised component is frequently a small connected app or a vendor's peripheral system, while the data that leaves belongs to a major platform you rated highly. Assessing only the primary platform misses the path entirely. The unit of assessment has to include the connection.

Bringing integrations into due diligence

Integration risk is assessable with questions that elicit evidence, not assurances. Add an integration section to assessments for vendors that connect to systems holding sensitive data, scaled to criticality.

INTEGRATION DUE-DILIGENCE QUESTIONS

Area	What to ask for
Connection inventory	Which of our systems the vendor connects to, and the OAuth scopes or API permissions each connection holds
Least privilege	Whether connections can be scoped down, and the justification for any broad read-write or admin scope
Token storage	How the vendor stores integration tokens and secrets, and whether they are encrypted and access-controlled
Token rotation	How often tokens are rotated, and whether they can be revoked and reissued without service disruption
Their fourth parties	Which subprocessors or connected apps the vendor in turn grants access to your data
Exposure notification	How, and how quickly, the vendor would notify you if integration tokens were exposed
Offboarding	How connections and tokens are revoked when the relationship ends

A vendor that can answer these has thought about the connection as a risk. A vendor that treats "we use OAuth" as a sufficient answer has not.

Scope and standing access: the two levers that matter most

Two properties determine how much an integration breach costs you. The first is scope. A connection granted read-write access to everything, because that was the default at setup, hands an attacker who obtains its token the same breadth. The same integration scoped to the one object it needs limits the blast radius to that object. Most integrations are over-scoped, and tightening them is among the highest-value, lowest-effort controls available.

The second is standing access. A human credential is bounded by the human: it fails when they leave, when they change their password, when multi-factor challenges them. An integration token has none of those bounds. It keeps working until someone deliberately revokes it, which means a

forgotten integration from a vendor you stopped using three years ago may still hold live access. Offboarding a vendor is not complete until its connections are revoked, and that step is routinely missed because the connection is invisible in the contract-centric view of the relationship.

Monitoring for the connections you did not authorize

Due diligence assesses the integrations you know about. Monitoring is how you find the ones you do not. Connected apps appear without a formal decision: a team enables a marketplace add-on, a trial integration is never removed, a vendor switches on a new connection by default. An outside-in and inside-out monitoring posture should treat a newly appearing connection into a sensitive system the same way it treats a new exposed service, as a signal to confirm, attribute and act on. The question a monitoring program should be able to answer is not only "how are our vendors' external postures changing" but "what now holds a token into our data that did not last month."

HYPOTHETICAL EXAMPLE: A TOKEN OUTLIVES THE RELATIONSHIP

An organization trials a niche analytics add-on, granting it a read token into the data warehouse, and decides after a month not to proceed. The subscription lapses; the OAuth connection is never revoked. Eighteen months later the add-on's provider is breached and its stored customer tokens are taken. The organization's warehouse token is among them, still valid, still scoped to read every table, because nobody revoked it when the trial ended. The primary warehouse platform was never breached. The exposure was a standing token from a vendor the organization no longer used. A revoke-on-offboarding step, and monitoring that flagged the dormant connection, would each have closed it.

Contract levers

The connection should be governed by the contract, not only by the setup screen.

- **Scope minimization:** the vendor holds the least access needed, and broad scopes require written justification
- **Token protection:** integration tokens and secrets are encrypted, access-controlled and rotated on a defined cadence
- **Exposure notification:** prompt notice if integration tokens or secrets are exposed, aligned to your incident regime
- **Subprocessor flow-down:** connected apps and subprocessors the vendor grants access to are disclosed with equivalent obligations
- **Revocation on exit:** all connections and tokens are revoked on termination, with confirmation

- [ ] **Right to review:** you may review the scopes and connections the vendor holds into your systems

What examiners and boards increasingly ask

Supervisory attention to identity and outsourcing already reaches this territory. Operational resilience and outsourcing regimes expect firms to understand and control how third parties access their systems and data, and identity-security guidance increasingly names non-human identities specifically. The board-level question is simple to ask and hard to answer without an inventory: which third parties hold standing access into our critical systems, at what scope, and how would we know if that access were abused.

WHAT AUDITORS AND EXAMINERS ASK

- Show your inventory of connected applications and integration tokens into systems holding sensitive data, with their scopes.
- For a material vendor, which of our systems does it connect to, and how are its tokens stored and rotated?
- How are integrations revoked when a vendor is offboarded, and how do you confirm it happened?
- How would you learn that a vendor's integration tokens had been exposed, and how fast?
- How do you detect a new connection into a sensitive system that was not formally authorized?

Practical checklist

- [ ] Vendors that connect to sensitive systems have an integration section in their assessment
- [ ] Each integration's OAuth scopes or API permissions are recorded and justified against least privilege
- [ ] Integration tokens the vendor holds are covered by storage, rotation and exposure-notification terms
- [ ] Connected apps and their subprocessors are recorded in the nth-party inventory
- [ ] Offboarding includes revoking every connection and confirming it
- [ ] Monitoring flags new or dormant connections into sensitive systems for confirmation and action
- [ ] The board can be shown which third parties hold standing access into critical systems

How VendRisk360 supports this

VendRisk360 lets you treat SaaS-to-SaaS integration as the risk category it has become, and assess the connection, not only the vendor.

Artifact-based assessments across 30+ control domains and configurable questionnaires let you add an integration section, connections, OAuth scopes, token storage and rotation, exposure notification and revocation, to any vendor that connects to systems holding sensitive data, scoped by criticality tier. Because the questions ask for evidence, a vendor's answers on scope and token handling become reviewable facts rather than assurances. Findings, remediation and formal risk acceptance run through **multi-stage sign-off with segregation of duties**, so a decision to accept a broad integration scope has an owner and a review date. See Comprehensive Vendor Risk Assessment Services.

The connected apps and providers a vendor in turn grants access to are fourth parties, and **[nth-party intelligence](/solutions/nth-party-intelligence/)** records them from SOC report subservice organizations and your own relationship data, so a connector that sits behind several of your vendors surfaces as a concentration and data-flow question. Where the connection changes, **[Continuous Monitoring](/services/continuous-monitoring/)** applies the confirm-attribute-act discipline: a confirmed, attributed, material signal, including a vendor's changing external posture, is routed to the relationship owner and can trigger an early reassessment or a targeted evidence request, rather than adding to an unread feed. Offboarding runs through the platform's exit workflow with a checklist and audit trail, so revoking a vendor's access is a tracked step rather than a forgotten one.

When an examiner asks which third parties hold access into your systems and how you govern it, the **examiner package export** bundles the assessment, evidence, sign-offs and audit trail, and board and executive reporting summarizes program status. Customer data is protected by **row-level-security tenant isolation**, encryption in transit and at rest, SSO (SAML/OIDC), mandatory MFA, custom roles and a full audit trail. See Security and Board and executive reporting.

How to work with VendRisk360. On the Vendor Lifecycle Management Platform, your team manages vendors, sends integration-focused due diligence through the vendor portal, performs the review and signs off, with every step tracked, and runs offboarding as a governed workflow. With Comprehensive Vendor Risk Assessment Services, you onboard the vendor and certified assessors collect the evidence, perform the review scaled to the vendor's tier with a second-expert quality review, and follow findings through remediation, while you keep final approval. Continuous Monitoring Services add the outside-in view between assessments, and Report-Specific Reviews are available standalone or alongside either. Assessors hold certifications such as CISSP, CISA, CISM, CRISC and ISO/IEC 27001 Lead Auditor and Lead Implementer.

To see integration risk assessed and monitored end to end, book a demo.

About VendRisk360

VendRisk360 is an independent company providing a third-party risk management platform and expert services for regulated organizations: banks, credit unions, fintech and payments, healthcare and SaaS. The platform runs vendor lifecycle management, risk-tiered assessments, continuous monitoring, nth-party intelligence and board-ready reporting from one governed record per vendor, combining point-in-time, evidence-based review with continuous outside-in monitoring. Customers run it themselves on the Vendor Lifecycle Management Platform, or add Comprehensive Vendor Risk Assessment Services, Continuous Monitoring Services and Report-Specific Reviews delivered by certified VendRisk360 assessors.

Learn more at **vendrisk360.com** or write to **info@vendrisk360.com**.

This paper is general guidance for practitioners, not legal advice. Regulatory requirements change and vary by jurisdiction and institution; confirm your obligations with counsel and your supervisors.