From Intrusion to Permission Abuse: SaaS Identity Infrastructure as a Strategic Attack Surface

Colin Cammack

April 8, 2026

Introduction

Software-as-a-Service (SaaS) platforms have become foundational to modern commercial infrastructure. Historically, organizations operated applications through local systems consisting of locally managed internal infrastructure like servers and databases. Today, many organizations rely on cloud-hosted software from web browsers. In this model, infrastructure is operated by the service provider rather than the organization itself. This centralizes authentication and administrative control within provider-managed environments.

While SaaS platforms increase efficiency and accessibility, they also control sensitive activity in a small number of cloud ecosystems. Defense contractors, universities, corporations, and government agencies are now conducting critical decision-making and data sharing inside SaaS environments. These come in the form of email systems, collaboration platforms, and document-sharing services.

Threat actor behavior has evolved with this shift. Attackers used to rely on malware or software exploits to attack vulnerable systems. Now they are increasingly focused on abusing identity infrastructure and permission systems inside SaaS environments. These attacks usually leverage OAuth permissions and tokens, API keys, and trusted third-party integrations. These tools look like legitimate platform activity. This makes it a major issue because abuse can allow adversaries to access sensitive systems without triggering intrusion detection systems. SaaS identity infrastructure has become a strategic attack surface. It analyzes the shift in attacker behavior, reviews publicly documented incidents involving token and permission abuse, and identifies common patterns that make these attacks difficult to detect and attribute.

The Structure and Importance of SaaS Platforms

SaaS applications are hosted in the cloud and accessed through the internet instead of being installed locally on a device. They centralize data, authentication, collaboration, and system administration within infrastructure controlled by the service provider.

SaaS systems rely on identity based access that allows authenticated users and services to interact with cloud resources. Components of this architecture include OAuth permissions, API keys, and authentication tokens. These systems allow organizations to connect many different applications and automate workflows across platforms.

SaaS provides a significant advantage for companies. Companies can deploy software rapidly, scale infrastructure easily, and allow teams to collaborate from anywhere. However, centralization always creates strategic risk. If attackers gain access to tokens or integration permissions, they could interact with systems as trusted entities without exploiting any software vulnerability.

Many institutions conduct sensitive operations inside SaaS platforms. This creates an immense risk because unauthorized access to these environments can reveal internal communications, strategic planning, and operational data.

The Shift from Intrusion to Permission Abuse

Traditional cyber attacks often relied on direct intrusion. Attackers exploited vulnerabilities, deployed malware or used user credentials from phishing campaigns. These typically showed detectable signals from abnormal login activity or compromised endpoints. More recent incidents are showing that threat actors are shifting toward a different strategy of abusing cloud permissions and identity tokens instead of breaking into systems directly.

Tokens, OAuth permissions, and API keys allow applications to access cloud resources without requiring a user to repeatedly log in. These are designed to improve usability and enable automation between services. However when attackers get access to these tokens, they can operate within the system as if they were legitimate services.

These services rely on identity infrastructure. Because of this, the activity appears normal in platform logs. The attackers are not exploiting software vulnerabilities or bypassing authentication systems in obvious ways that can be detected. They are using trusted applications that already exist within the platform. This approach creates a major challenge for defenders. Traditional security tools are designed to detect malicious code or unauthorized logins instead of subtle abuse of legitimate permissions.

Case Studies of SaaS Permission Abuse

Salesforce OAuth Token Compromise

The Salesforce OAuth token compromise attributed to UNC6395 represents a shift from traditional intrusion toward identity-layer exploitation of SaaS integrations. Instead of targeting vulnerabilities in Salesforce, the attacker leveraged connection between Salesforce and the third-party application Drift. [5], [6]

In a standard OAuth integration, a user authorizes a third-party application to access Salesforce data. This generates access and refresh tokens that allow the application to act on behalf of the user without requiring repeated authentication. These tokens function as persistent identity credentials. They often have broad scopes such as reading customer data or interacting with internal workflows.

UNC6395 obtained these OAuth tokens through compromise of systems associated with the third-party integration Drift. Because OAuth tokens are bearer credentials, possession of them is sufficient for access. The attacker did not need to bypass authentication controls such as passwords or multi-factor authentication. They also did not trigger a login event. Instead, they interacted directly with Salesforce as a trusted integration. [5]

The scale of the incident was created by the multi-tenant nature of SaaS ecosystems. Drift maintained authorized connections across hundreds of independent Salesforce environments. By compromising tokens associated with Drift, the attacker could access data across more than 700 organizations. This demonstrates how SaaS integrations create dangerous trust relationships. A single compromise can cascade across environments.

The attack also highlights a critical detection challenge. Because all activity was performed using valid OAuth tokens, the logs reflected normal API usage. Security tools designed to detect login-based threats were ineffective, as no authentication boundary was crossed during the attack.

Downstream impacts were confirmed by major technology providers including Cloudflare, Zscaler, and Palo Alto Networks. The involvement of these organizations proves the risk posed by integration-layer compromises. This is especially dangerous when they intersect with companies responsible for securing internet infrastructure and enterprise environments. [5]

Microsoft Storm-0558 Token Abuse

Another major incident involved a threat actor known as Storm-0558. This actor conducted a campaign targeting Microsoft cloud email environments in 2023. Attackers obtained a Microsoft signing key used for authentication systems. With this key, they were able to generate valid access tokens for other accounts without triggering security protections such as multi-factor authentication. [1], [2]

With the tokens, the attackers accessed cloud email environments belonging to multiple organizations. These organizations included government agencies and managed service providers. The attackers were using valid authentication tokens instead of logging in through traditional mechanisms, so their activity appeared legitimate in platform logs.

The incident showed how control over identity infrastructure can enable large-scale access to cloud systems without exploiting software vulnerabilities.

Although Storm-0558 did not rely on a third-party SaaS integration in the same way as the Salesforce case, it demonstrates the same core principle: control over identity infrastructure can provide adversaries with trusted access at scale without requiring direct exploitation of the victim environment.

Google Gemini API Key Exposure

A third, more recent example involves the exposure of Google Cloud API keys that were publicly embedded in websites and applications. Researchers discovered thousands of these keys available in publicly accessible code. Since they have been around, developers have treated these keys as low-risk identifiers. However, when Google's Gemini generative AI services were enabled within a project, the same keys could be used to authenticate requests to AI systems. [3], [4]

Attackers who collected these keys could send requests to Gemini services which could generate responses or create large usage costs billed to the victim's account. This incident proves how changes in cloud service permissions can change previously low-risk credentials into powerful access tokens.

The significance of this incident was not because API keys were exposed in public code. Publicly embedded cloud API keys have existed for years and are treated by developers as low-risk identifiers instead of sensitive credentials. What changed was the capability attached to those keys. Once generative AI services such as Gemini were enabled within the associated Google Cloud projects, the same exposed keys could be used to authenticate requests to AI infrastructure. This transformed allowed exposure into dangerous access.

This illustrates a broader weakness in cloud ecosystems. The significance of a credential can change over time without the credential actually changing. This means risk is determined by exposure, permissions, and services. A key that used to have limited functionality can become valuable when new services are activated or when platform architecture changes.

Although this case differs from OAuth token theft or signing-key compromise, it reflects the same pattern of adversaries increasingly seeking access through legitimate identity and service mechanisms instead of the traditional method of software exploitation alone.

Common Patterns in SaaS Identity Attacks

Analysis of these incidents reveals several recurring characteristics. First, access is usually non-interactive. Attackers do not log in through traditional authentication systems. They use tokens, API keys, or permissions to access systems programmatically. Second, multi-factor authentication protections can be bypassed because tokens represent authenticated sessions or trusted services. Third, activity appears legitimate within platform logs. This is because the requests originate from valid authentication mechanisms instead of suspicious login attempts. Fourth, attackers can have persistent access over long periods of time. This happens if the compromised tokens remain valid or are automatically refreshed. Finally, these create challenges for attribution and response because the activity resembles normal platform usage rather than an obvious intrusion.

Conceptual Model of SaaS Identity Abuse

To illustrate the relationships between the major themes discussed in this paper, the following concept map provides a visual model of the SaaS identity attack surface. The diagram organizes the analysis into four connected components: the structure of Software-as-a-Service platforms, the shift in attacker behavior toward identity and permission abuse, documented case studies of SaaS compromise, and the common patterns observed across these incidents. Together, these elements demonstrate how modern cloud platforms have created a new attack surface centered on identity infrastructure such as OAuth permissions, API keys, and authentication tokens.

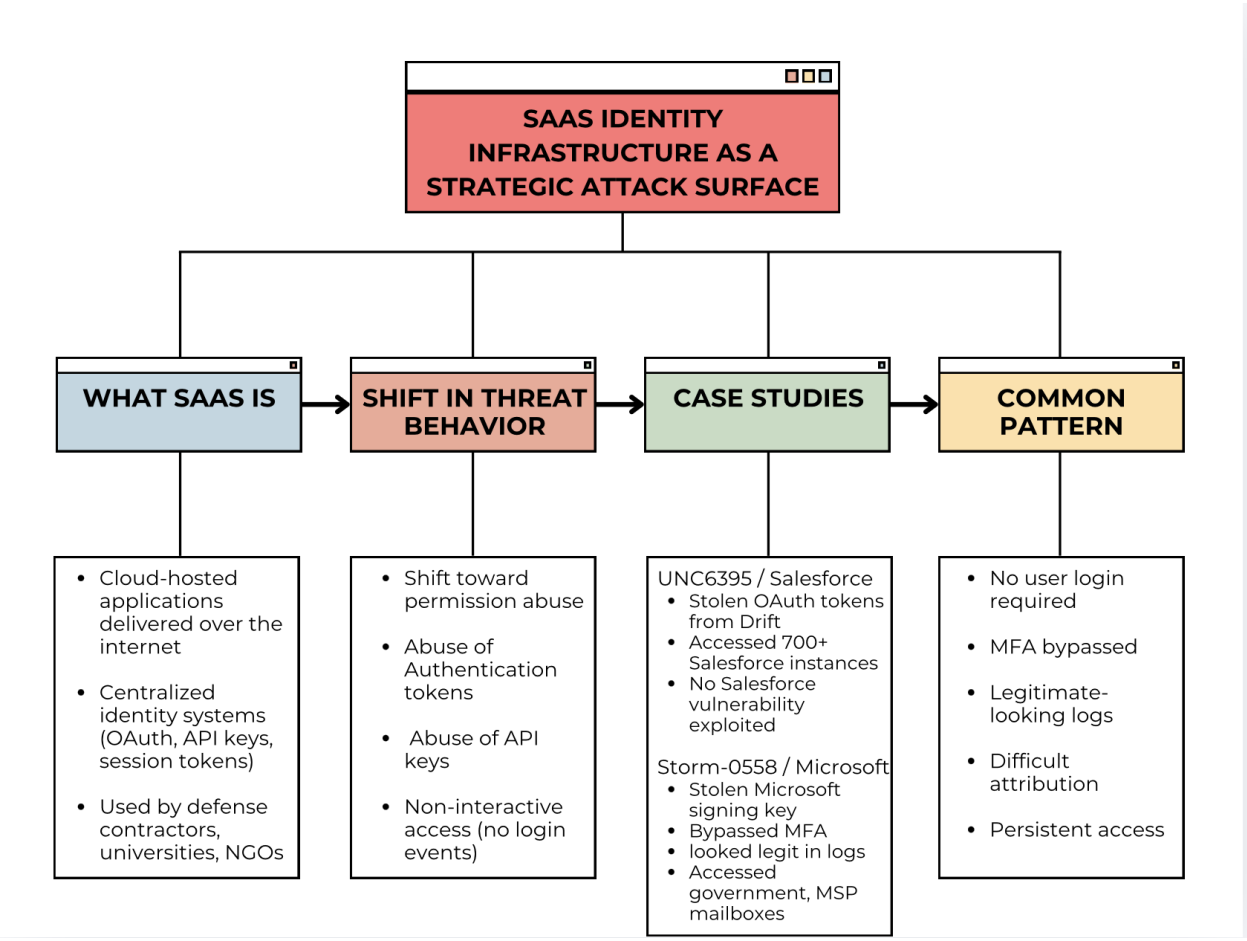

National Security Implications

The concentration of organizational activity in SaaS platforms creates larger strategic risks. When adversaries gain access to tokens or integrations, they can observe internal communications and operational planning of institutions. Persistent access to these sensitive platforms like email systems and document repositories could enable intelligence collection on political, corporate, or military processes.

The centralization of sensitive activity within a small number of cloud providers creates strategic chokepoints. A successful compromise of identity infrastructure within widely used SaaS ecosystems could provide adversaries with access to numerous organizations simultaneously. As governments and institutions continue migrating operations to cloud platforms with SaaS, securing identity infrastructure will only become more critical.

Conclusion

The fast adoption of SaaS platforms has transformed how organizations are structured. It changes how they collaborate and manage operations. This transformation has also created a new type of cybersecurity risk centered on identity infrastructure.

Recent incidents have shown a shift in attacker behavior from traditional intrusion to the abuse of tokens and API keys. These mechanisms are components of SaaS ecosystems, and their misuse can allow adversaries to operate in systems while appearing the same as authorized users or services.

As organizations continue to migrate sensitive workflows into cloud ecosystems, identity infrastructure determines whether modern institutions remain secure or strategically exposed. OAuth permissions, authentication tokens, and trusted integrations are not just implementation details. They are becoming part of the terrain of modern cyber conflict.

Sources

1. Microsoft Threat Intelligence. (2023, July 14). *Analysis of Storm-0558 techniques for unauthorized email access*. Microsoft Security Blog. https://www.microsoft.com/en-us/security/blog/2023/07/14/analysis-of-storm-0558-techniques-for-unauthorized-email-access/

2. Microsoft Security Response Center. (2023, September 6). *Results of major technical investigations for Storm-0558 key acquisition*. Microsoft Security Response Center Blog. https://www.microsoft.com/en-us/msrc/blog/2023/09/results-of-major-technical-investigations-for-storm-0558-key-acquisition/

3. Google Cloud. (2025). *Best practices for managing API keys*. Google Cloud Documentation. https://cloud.google.com/docs/authentication/api-keys-best-practices

4. Arntz, P. (2026, February 27). *Public Google API keys can be used to expose Gemini AI data*. Malwarebytes Labs. https://www.malwarebytes.com/blog/news/2026/02/public-google-api-keys-can-be-used-to-expose-gemini-ai-data

5. Google Threat Intelligence Group. (2025, August 26). *Widespread data theft targets Salesforce instances via Salesloft Drift OAuth compromise*. Google Cloud Blog. https://cloud.google.com/blog/topics/threat-intelligence/data-theft-salesforce-instances-via-salesloft-drift

6. ITPro. (2025, September 15). *How to check if you've been affected by Salesforce attacks — and stop hackers dead in their tracks*. ITPro. https://www.itpro.com/security/cyber-attacks/fbi-flash-warning-salesforce-attacks-salesloft-drift