unosecur

WHITEPAPER

Cyber Resilience in Japan:

A CISO Guide to Identity Security



Table of contents

Executive summary 02

Part 1. Japan's 2026 incident data: what it shows, and what it does not 03

Part 2. Who holds the keys: four access patterns 04

Part 3. What Japan's authorities now expect 07

Part 4. Assess: six identity questions for security leaders 10

Part 5. Respond: when identity is involved in an incident 12

Part 6. A 90-day plan: from access to evidence 13

How Unosecur helps 14

Next step 16

About Unosecur 16

About the data and sources 16

Executive summary

In Japan, more of 2026's reported incidents come down to access: who or what could log in, with how much power, and whether anyone was watching. Between 1 April and 16 October 2026, the Personal Information Protection Commission (PPC) received 2,094 leak reports attributed to unauthorized access. Of those, 830 arrived in the first 16 days of October, compared with 303 in all of September.

Incidents have different causes, and many are still being investigated. Across the public cases we reviewed, four access patterns keep appearing:

1. **Stolen and poorly managed credentials.** Of 407 unauthorized-access acts that led to arrests in 2025, 399 used someone else's ID and password.

2. **Administrator and contractor accounts.** Askul reported that its 2025 ransomware attack began with leaked credentials for a contractor's administrator account that had no multi-factor authentication (MFA).

3. **Shared platforms and suppliers.** In one public incident dataset, 22% of 2026 incidents came through a contractor or cloud provider.

4. **Tokens, APIs and AI agents.** Non-human credentials are a growing route. The clearest AI agent cases so far are outside Japan.

Japan's authorities now point to the same controls. The PPC, the National Cybersecurity Office (NCO) and METI's supply chain security scheme all name MFA, limited administrator accounts and regular access clean-up. The Financial Services Agency (FSA) has made phishing-resistant MFA mandatory for securities firms. Buyers, auditors and regulators will increasingly ask for evidence that these controls are in place.

Every one of those controls assumes an organization can list who and what has access. For most enterprises, that list is the hard part. This guide helps security leaders build it:

- **Part 1** sets out what Japan's 2026 data shows, and what it does not.
- **Part 2** describes the four access patterns.
- **Part 3** maps what Japanese authorities now expect.
- **Part 4** is a six-question self-assessment covering human, non-human and AI agent identities.

- **Part 5** covers response when identity is involved in an incident.
- **Part 6** is a 90-day plan that turns access into evidence, with or without a tool.

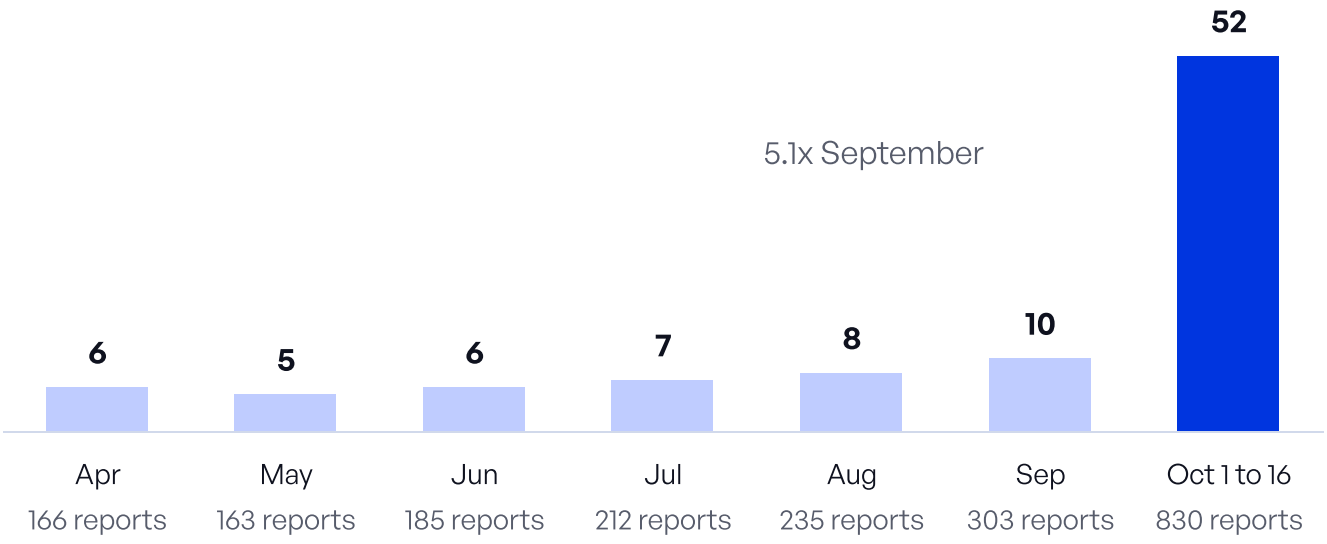
The question this guide asks: can you show today who and what holds the keys to your systems?

Part 1. Japan's 2026 incident data: what it shows, and what it does not

Leak reports caused by unauthorized access rose steadily through 2026, then jumped in October. The PPC received 166 such reports in April and 303 in September, then 830 in the first 16 days of October.

Early October ran at about five times September's daily rate

Leak reports to the PPC attributed to unauthorized access, per day received, 2026



Preliminary figures; counts include duplicate reports from companies and their contractors.

PPC reference material of 28 October 2026, as reported by Business+IT · April to 16 October 2026

Police data points the same way. Ransomware reports reached a record 123 in the first half of 2026, after 226 across all of 2025 (The Japan Times, reporting National Police Agency data).

Police recognized 7,190 unauthorized-access acts in 2025, up 34% on 2024 (joint report by the National Police Agency (NPA), the Ministry of Internal Affairs and Communications (MIC) and METI).

October 2026 brought several large disclosures. These included skyticket (about 14.6 million items), Daiichi Kosho (about 8.7 million) and Bookoff Group (up to about 6.4 million), as reported by Impress Watch. A ransomware attack on cloud provider IDC Frontier disrupted systems used by 495 companies and local authorities (South China Morning Post).

Read the numbers with care

- **Not every leak is an attack.** By count, most PPC reports concern misdelivered or misdirected mail. This guide concerns the access-driven share. Reports flagged as a likely improper purpose, such as unauthorized access, were 26.6% of operators' reports in April to June 2026, the highest of the last five quarters.
- **More is being counted.** Leak reporting became mandatory in April 2022, so part of any rise reflects reporting rather than attacks. The October spike appears in more than one dataset, but it should not be read as a clean trend.
- **Most public notices do not say how attackers got in.** Where a route is not public, this guide says so and does not guess.

What the data does show: the weak points Japanese authorities now name (credentials, administrator accounts and supplier access) keep recurring in the cases where a route is known.

Part 2. Who holds the keys: four access patterns

In the public cases where a route is known, attackers rarely needed to break anything. They used access that already existed. Four patterns account for most of what we reviewed.

1. Stolen and poorly managed credentials

Stolen credentials dominate the police record. Of 407 unauthorized-access acts that led to arrests in 2025, 399 used another person's ID and password. Police traced those credentials to four main sources: weak password setup or management (84 cases), phishing sites (77), other people (75), and

former employees or acquaintances who could know them (61) (NPA, MIC and METI joint report).

The financial impact follows the login. In 2025, a wave of phishing led to the takeover of brokerage accounts: 17,668 unauthorized accesses and about ¥740.8 billion in unauthorized trades, according to FSA data cited by the NPA.

The identity lesson: a valid login is no longer proof of who is behind it. Strong authentication and prompt removal of unused credentials matter as much as detection.

2. Administrator and contractor accounts

Askul reported that its October 2025 ransomware attack began with leaked credentials for an administrator account used by an outsourcing contractor, with no MFA applied. About 740,000 records were affected, and the company later booked a ¥5.2 billion special loss.

The NPA's account of ransomware shows what typically follows entry. Attackers explore the network, try to seize administrator privileges, disable security tools and search for backups. The NPA also advises administrators to delete or change IDs promptly when someone stops using a system, citing cases involving former employees and contractors.

The identity lesson: administrator rights held by people outside the organization, or by people who no longer need them, are among the most valuable keys an attacker can find.

3. Shared platforms and suppliers

About one in five 2026 incidents in the JapanHacked dataset (113 of 507, or 22%) came through a supplier: a contractor or cloud provider was breached rather than the organization's own systems. When a platform is shared, one failure spreads. The IDC Frontier attack disrupted 495 organizations at once. An outage is not the same as data theft, so each customer's disclosure should be read on its own terms.

The Information-technology Promotion Agency (IPA) ranks supply chain and contractor attacks second among organizational threats for 2026. Verizon's 2026 Data Breach Investigations Report found third parties involved in 48% of breaches.

The identity lesson: a supplier's access is your access. It needs an owner, a scope and an end date.

4. Tokens, APIs and AI agents

Non-human credentials are a fast-growing route. On 7 October 2026, the PPC added an API abuse case to its warning guidance. After logging in, an attacker changed a URL or member ID value to read other users' data, helped by sequential IDs and no rate limits.

Outside Japan, attackers used stolen OAuth tokens linked to Salesloft Drift, an AI chat agent integration, to export data from many Salesforce instances in August 2025 (Google Threat Intelligence Group).

OWASP's 2026 list of risks for agentic applications includes identity and privilege abuse. In the sources reviewed, we found no Japanese incident tied to a hijacked AI agent.

The identity lesson: service accounts, API keys, OAuth grants and AI agents act with real permissions, often broader than intended and rarely reviewed.

Why these patterns are growing

- **Credentials are easier to get.** Phishing reports to JPCERT/CC rose every quarter of 2025, reaching 12,397 in the fourth quarter.
- **Attackers use legitimate access.** CrowdStrike found that 82% of its 2025 detections were malware-free. A login does not look like an attack to a malware scanner.
- **More third parties hold access.** Outsourcing, cloud services and SaaS integrations all add identities that sit outside normal reviews.
- **The basics lag.** In the NPA's ransomware data, VPN devices were over 60% of intrusion routes, and only 18% of victims had a continuity plan that anticipated a cyberattack.

Identity is not the only door. Verizon's 2026 report found vulnerability exploitation overtook credential abuse as the top initial access route. Patching and access control have to advance together, and some risks, such as flaws in an application's authorization logic, are fixed in the application itself.

Part 3. What Japan's authorities now expect

Three authorities now name the same access controls, with different force. The PPC's items are examples of safeguards under the existing legal duty to take necessary and appropriate measures, not new obligations. The NCO's text is an advisory. METI's supply chain scheme is buyer-driven: buyers specify a tier in contracts, and the IPA plans to open ★3 and ★4 applications in March 2027.

Control	PPC (7 Oct 2026 preview of guideline examples)	NCO advisory (9 Oct 2026)	METI supply chain scheme (draft criteria)
Multi-factor authentication	MFA, including phishing-resistant MFA, for outside access, admin access and important personal data	Introduce MFA	Always for cloud services holding important information (★3); admins and highly sensitive systems over the internet (★4)
Passwords	Ban ID and password reuse in principle; lock IDs after repeated failures	Ban easily guessed passwords	No reuse across devices and services; 10+ characters; a procedure when a leak is found
Admin accounts	Limit the staff, devices and situations that use privileged accounts	Minimal admin users and rights	Limit who holds admin IDs and list them, including external staff
Joiners, movers, leavers	Update access on transfers and retirements; periodic access-rights inventory	Remove access for leavers, movers and unused accounts	Disable departed or unused IDs promptly; keep an access-rights ledger and inventory

Control	PPC (7 Oct 2026 preview of guideline examples)	NCO advisory (9 Oct 2026)	METI supply chain scheme (draft criteria)
Supplier and contractor access	Not in the preview	Check contractors continuously; require breach reporting in contracts	List externally managed systems connected to you (★3); check suppliers that can reach your systems yearly and revoke access at contract end (★4)
Logs and detection	Retain and analyze authentication and access logs; disable accounts to contain an intrusion	Monitor logs	Keep authentication server logs for 6 months and review them monthly (★4)
Context-aware access	Evaluate location, time and device state per request; allow only approved devices	Move toward zero trust	Not in the draft

The wider rule stack

• **Active Cyber Defense framework** (enacted May 2025). Since 1 October 2026, designated critical infrastructure operators (258 across 16 sectors as of 1 July) must report specified intrusions. These operators are expected to ask their suppliers for logs and forensic findings.

• **Unified Standards for Cybersecurity Measures for Critical Infrastructure** (in effect 1 October 2026). A common baseline for governance, risk management, incident response and supply chain security, to be reflected in sector standards.

• **FSA supervisory guidelines for securities firms.** Phishing-resistant MFA is mandatory, with firms asked to finish by the end of June 2026 in principle.

• **Amended Act on the Protection of Personal Information (APPI)** (promulgated 17 July 2026). Introduces administrative surcharges for the first time, taking effect no later than July 2028.

Why this is now a management issue

- **Failure is expensive.** In the NPA's survey, more than half of ransomware victims spent ¥10 million or more on investigation and recovery.
- **Buyers will ask for proof.** Critical infrastructure operators pass requirements down to suppliers, and METI's tiers will appear in contracts.
- **Management is accountable.** The FSA tells securities firms that security is management's responsibility, and the NCO has said attacks are no longer only an IT department problem.

The visibility gap

Every control in the table assumes you can list who and what has access. That list is the hard part, and authorities ask for it directly. METI's draft criteria require lists of user and administrator IDs, a list of cloud services that lack MFA support, and a yearly review of those lists.

A buyer, auditor or regulator asking about these controls will ask questions like these:

- Who holds administrator rights today, including contractors, service accounts and AI agents?
- Which leavers and movers still have access, and which accounts have gone unused?
- Which partner, guest and contractor identities can reach our systems, and when does that access end?
- Which service accounts, tokens and app grants exist, who owns each, and are they covered by MFA?
- How much of the access we have granted is actually used?

The incident data cannot show how often a missing inventory caused a breach, because most public notices do not say how attackers got in. The data does show that the weak points authorities name keep recurring. An accurate, current inventory of access is what lets an organization address them and prove it has.

Part 4. Assess: six identity questions for security leaders

Use these six questions to judge whether you could answer a buyer, auditor or regulator today. Score each from 1 to 3. The point is not the total. It is to find the questions where you would be guessing.

Question	1. Unknown	2. Partial	3. Evidence-ready	Related expectations
Q1. Can we see every identity? Employees, contractors, service accounts, API keys, OAuth grants and AI agents, across cloud, SaaS and on-premises	No single list; each team keeps its own	Human identities listed; non-human and AI agent identities partly known	One current inventory covering all identity types, each with an owner	METI ID lists; PPC access-rights inventory
Q2. Where is privilege excessive or inherited?	Admin holders not fully known	Admin roles listed, but group inheritance and unused rights not reviewed	Admin rights listed, including external staff; granted access compared with used access	PPC and NCO limits on privileged accounts; METI admin ID list
Q3. What can suppliers and contractors reach?	Supplier access granted case by case, not tracked	Supplier accounts listed; scope or end dates missing	Every external identity has a scope, an owner and an end date, reviewed yearly	NCO contractor checks; METI ★3 and ★4 supplier criteria

Question	1. Unknown	2. Partial	3. Evidence-ready	Related expectations
Q4. Which accounts and credentials are dormant?	Leavers removed when someone remembers	HR leavers checked; unused service accounts and keys not checked	Leavers, movers and unused accounts found and removed on a set cycle	PPC, NCO and METI joiner, mover and leaver items
Q5. How do identities connect across environments?	Unknown	Known for key systems through federation or directory sync	Access paths across cloud and on-premises mapped, including what one identity could reach if misused	METI externally managed systems list
Q6. Can we investigate identity activity with context?	Authentication logs not kept or not reviewed	Logs kept; reviewed only after an incident	Logs kept for at least 6 months and reviewed monthly; alerts on unusual sign-ins and privilege changes	PPC log analysis; METI ★4 log review

Reading your score: a 1 on any question is where to start, whatever the total. A 2 usually means the data exists but no one can produce it quickly. A 3 means you could hand over the evidence this week.

For AI agents, add one more question to Q1 and Q2: which agents exist, which tools and data each can call, and on whose authority?

Part 5. Respond: when identity is involved in an incident

When an incident involves a valid account, the first questions are about access. Which identity was used, what could it reach, and what else might it have changed? The steps below are general guidance. They complement your incident response plan and legal advice; they do not replace them.

Before an incident

- Name an owner for every administrator account, service account and supplier account, so someone can act quickly.
- Keep authentication logs long enough to investigate. METI's ★4 draft asks for 6 months.
- Agree in supplier contracts how and when suppliers report a breach that could affect you.

During an incident

1. **Contain the identity.** Disable or restrict the accounts involved, as the PPC's examples describe. End active sessions, and revoke tokens and app grants tied to them. Rotate exposed keys and secrets.
2. **Look for persistence.** Check for newly created accounts, privilege changes, new MFA devices and new app consents made by the affected identity.
3. **Size the blast radius.** List what else that identity, and anything it controls, could reach. Include supplier and service accounts that share the same access paths.
4. **Report on time.** Under the APPI, a qualifying leak requires a preliminary report to the PPC promptly and a confirmed report within 30 days, or 60 days where an improper purpose such as unauthorized access is involved. Affected individuals must also be notified. Designated critical infrastructure operators have separate reporting duties under the Active Cyber Defense framework.
5. **Restore on least privilege.** When access is restored, grant only what is needed, rather than putting everything back as it was.

After an incident

- Add the route used to your access reviews, so the same kind of key is checked everywhere.
- Share what is confirmed and what is not. Disclosures that separate the two keep trust with customers and regulators.

Part 6. A 90-day plan: from access to evidence

A security team can build the evidence authorities ask for in three phases, with or without a tool. Each phase follows the controls in Part 3 and raises your scores in Part 4.

Days 1 to 30: find out who holds the keys

- List every administrator identity: employees, contractors, service accounts and shared accounts.
- List AI agents and app integrations with access to company data, and the permissions each holds.
- List cloud services that do not support MFA, and administrator accounts without it.
- Name an owner for each administrator account, service account and agent.

Outcome: a first inventory that answers Q1 and Q2.

Days 31 to 60: close the obvious gaps

- List contractor, partner and guest accounts, app grants and tokens, with an end date for each.
- Compare HR leaver and mover records with live accounts, and disable what is no longer needed.
- Remove accounts and keys that have gone unused for a defined period.
- Apply MFA to administrator and remote access first, with phishing-resistant methods where possible.

Outcome: fewer standing keys, and supplier access with scope and end dates (Q3 and Q4).

Days 61 to 90: turn it into evidence

- Record each access review: who reviewed what, when, and what changed.
- Review authentication logs at least monthly for suspicious sign-in attempts, as METI's ★4 draft asks.
- Map access paths for your most sensitive systems, including what a single compromised identity could reach.
- Prepare answers for buyers who ask about ★3 or ★4, and track the PPC's final guideline date.

Outcome: a review record you can hand to a buyer, auditor or regulator (Q5 and Q6).

After day 90

Make the cycle routine. Agree review frequencies, keep owners current as people change roles, and repeat the self-assessment each quarter.

How Unosecur helps

Unosecur's Unified Identity Fabric (UIF) builds one live view of every human, non-human and AI agent identity across multi-cloud, SaaS, identity providers and on-premises systems. This gives security teams the inventory that each part of this guide depends on.

- **Connect.** Native API connectors link identity providers, cloud accounts and SaaS applications across 100+ integrations, without agents. On-premises systems require a lightweight agent.
- **See.** Human, non-human and AI agent identities appear in one identity graph. It shows the permissions each identity holds and uses, its access paths, and what it could reach if misused.
- **Prioritize.** Findings are risk-scored, so teams see first the actions that remove the most risk.

- **Remediate.** Least-privilege and time-boxed access policies, including templates for temporary contractor access. Uno AI, part of UIF, helps teams understand identity risk and work through remediation, starting with least-privilege recommendations.

- **Prove.** Timestamped records of access reviews and remediation actions, with identity controls mapped to frameworks such as ISO 27001, SOC 2, NIST and PCI-DSS. METI's draft criteria cite ISO/IEC 27001 controls for most requirements.

AI agents. Unosecur treats AI agents as identities to govern like any other. It discovers agents across cloud and SaaS, traces the tools, roles and data each can reach, and compares granted access with used access. For agents that use the Model Context

Protocol (MCP), Unosecur's Ark AI checks each tool call against policy before it reaches the server.

What a buyer or regulator asks	How Unosecur helps answer it
Who holds admin rights, including service accounts and AI agents?	One inventory of human, non-human and AI agent identities, with permissions held and used
Which accounts are dormant, orphaned or over-privileged?	Risk-scored findings on dormant administrators and over-privileged accounts
Which accounts are not protected by MFA?	Visibility into blind spots such as shadow admins, service accounts and AI agents without MFA
Which contractors have access, and until when?	Policy templates for time-boxed contractor access
Show us the access review record	Timestamped evidence of reviews, escalations and remediation
Has a login or an agent been misused?	Behavioral baselines, alerts and audit trails, with detections mapped to MITRE ATT&CK

Not every risk in this guide is an identity problem. Flaws in an application's authorization logic are fixed in the application, and unpatched devices are fixed by patching. Identity visibility makes those fixes easier to prioritize, but it does not replace them.

Next step

Start with the question this guide asks: who and what holds access to your systems today? Bring your Part 4 scores, or your first 30-day inventory, to a conversation with Unosecur's team, and we will focus it on your administrator, contractor and non-human identities.

About Unosecur

Unosecur is redefining identity security for the agentic AI era, protecting human and non-human identities, including AI agents, at runtime. Your AI agents already have access. Unosecur shows you how far it reaches. Discover every agent, trace its access paths, and uncover risky access across your enterprise. Within 15 minutes of agentless onboarding, gain real-time visibility across multi-cloud, SaaS, and agentic workflows, uncovering toxic combinations, and privilege anomalies. Turn observed usage into just-enough privilege, so you can scale AI without scaling identity risk.

About the data and sources

Incident counts and shares from JapanHacked (Aawa Technologies, licensed CC BY 4.0) were calculated by Unosecur from 1,650 incident rows read on 6 November 2026. JapanHacked holds selected incidents from public sources, not every incident, and its classifications are its own. Read these figures as patterns within that dataset. All other figures come from the official and reported sources below. Links are carried over from the security team's brief and should be checked before publication.

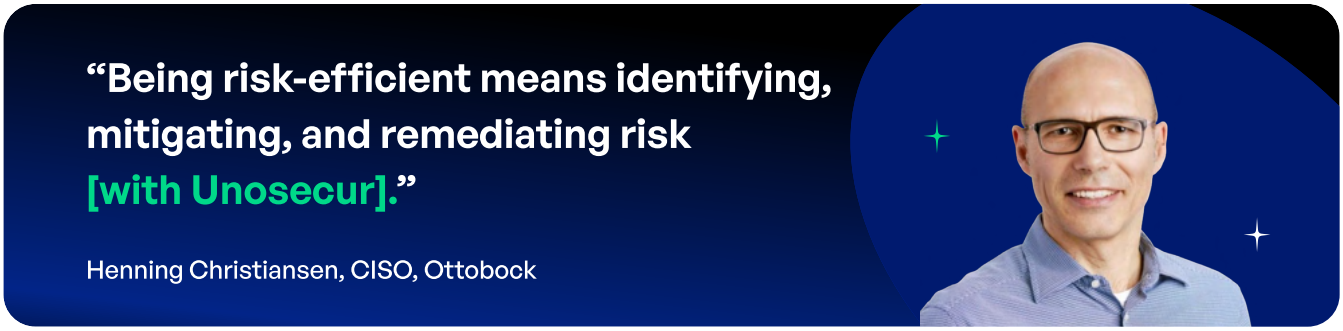
- **Official sources:** NPA, MIC and METI joint report on unauthorized access (March 2026); NPA, Threats in Cyberspace in 2025; PPC alert of 7 October 2026; PPC revision overview; METI draft supply chain requirement list; FSA talking points (18 November 2025); Cabinet Secretariat FY2027 budget request.

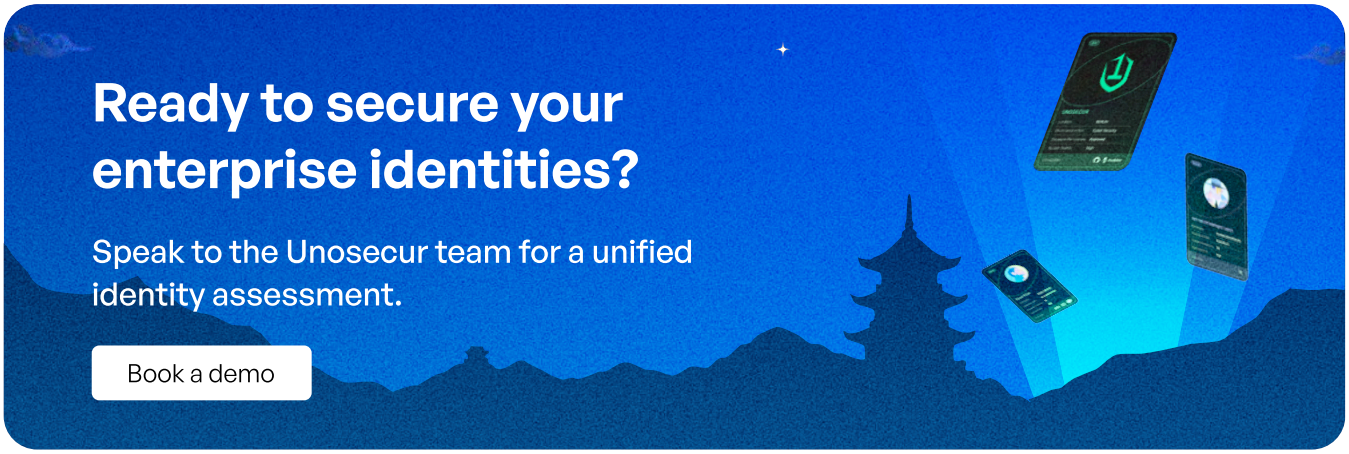
- **Reported from official material:** Business+IT on PPC data; The Japan Times on the NPA half-year figure; ITmedia and ScanNetSecurity on the NCO advisory; Internet Watch and Mynavi on METI's scheme; ITmedia on IPA's 10 Major Security Threats; HLC on the Active Cyber Defense framework; A&O Shearman on the APPI amendment.

• **Incidents and trends:** IT Leaders and Netshop Forum on Askul; South China Morning Post on IDC Frontier; Impress Watch on the October disclosures; Google Threat Intelligence Group on Salesloft Drift; OWASP Top 10 for Agentic Applications 2026; Verizon 2026 Data Breach Investigations Report; CrowdStrike 2026 Global Threat Report.

Trusted by security-first teams **worldwide**

Rakuten Symphony | KFINTECH | Zeotap | ottobock. | aikido





GDPR | AICPA SOC | ISO 42001 Certified | ISO 27001 Certified | STAR LEVEL ONE | HIPAA COMPLIANT