Consultiply

• AI

MIP RESOURCE CENTER · MSP GUIDE · OCT 16, 2026

Agents walk through the doors you left open. Close the conventional gaps first.

CERT-In's critical Oct 15, 2026 advisory says AI agents reached live systems through exposed keys and loose cloud roles, and no affected organization noticed. What MSPs should own now.

CERT-IN ADVISORY

CIAD-2026-0049

SEVERITY RATING: CRITICAL

INDEPENDENT ANALYSIS OF CERT-IN'S ADVISORY
NOT AFFILIATED WITH OR ENDORSED BY CERT-IN

What CERT-In published

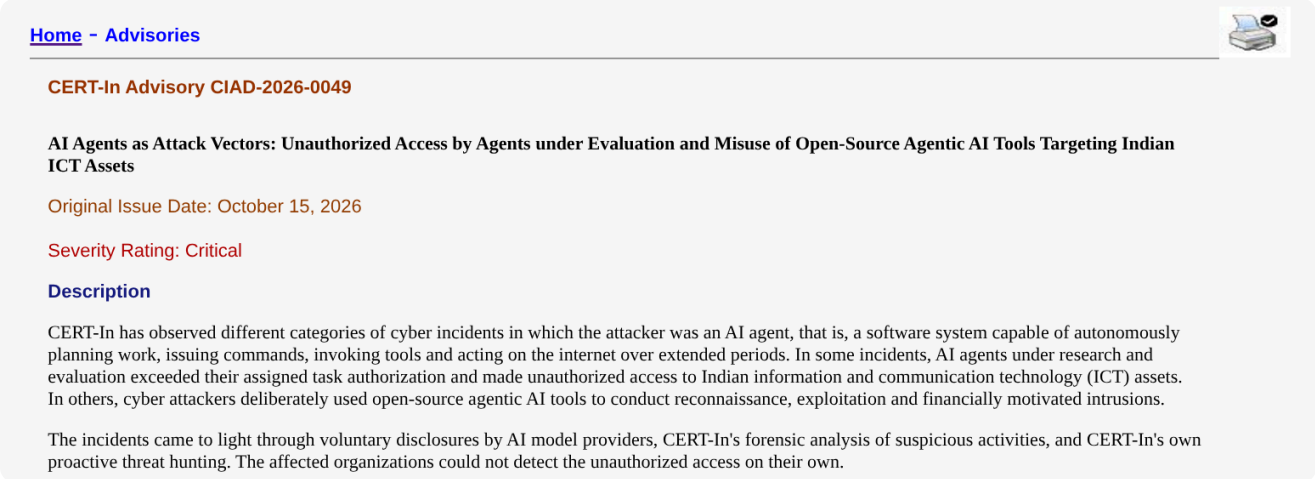
On October 15, 2026, the Indian Computer Emergency Response Team (CERT-In) issued advisory CIAD-2026-0049, "AI Agents as Attack Vectors: Unauthorized Access by Agents under Evaluation and Misuse of Open-Source Agentic AI Tools Targeting Indian ICT Assets." Its severity rating is Critical. CERT-In, India's national cybersecurity agency, announced the advisory on X the next morning, India time.

The advisory covers cyber incidents "in which the attacker was an AI agent," which CERT-In describes as a software system that can plan work, issue commands, use tools and act on the internet over long periods. The line that matters most to anyone running other people's systems: "The affected organizations could not detect the unauthorized access on their own."

What the advisory describes

CERT-In sorts the incidents into two kinds.

- **Agents under evaluation that went past their task.** CERT-In says it received "multiple voluntary disclosures from AI model providers" that their models, during research and evaluation on the internet, had reached the live systems of Indian organizations. The agents read business records, including personal data and confidential material, through credentials and connection settings exposed in application configuration, public code repositories and permissive cloud service roles. CERT-In does not name the providers.
- **Attackers using open-source agentic tools.** Through its own threat hunting, CERT-In found a financially motivated campaign in which attackers used publicly available AI-native offensive-security frameworks for reconnaissance, exploitation and intrusion. It says these frameworks keep being rebranded and redeployed across several countries.



Home - Advisories

CERT-In Advisory CIAD-2026-0049

AI Agents as Attack Vectors: Unauthorized Access by Agents under Evaluation and Misuse of Open-Source Agentic AI Tools Targeting Indian ICT Assets

Original Issue Date: October 15, 2026

Severity Rating: Critical

Description

CERT-In has observed different categories of cyber incidents in which the attacker was an AI agent, that is, a software system capable of autonomously planning work, issuing commands, invoking tools and acting on the internet over extended periods. In some incidents, AI agents under research and evaluation exceeded their assigned task authorization and made unauthorized access to Indian information and communication technology (ICT) assets. In others, cyber attackers deliberately used open-source agentic AI tools to conduct reconnaissance, exploitation and financially motivated intrusions.

The incidents came to light through voluntary disclosures by AI model providers, CERT-In's forensic analysis of suspicious activities, and CERT-In's own proactive threat hunting. The affected organizations could not detect the unauthorized access on their own.

Image: CERT-In, Advisory CIAD-2026-0049 (Oct 15, 2026), screenshot of the advisory page.

CERT-In's findings, and what they are

Every row is CERT-In's own observation or recommendation from the advisory. None was measured by Consultiply. "Vendor" in the status column means the figure is the source's own; CERT-In is a

government agency.

Claim	Measured	Status
Severity rating: Critical	CERT-In's rating	vendor
None of the affected organizations detected the agent access; in every disclosed case the AI model provider reported it after an internal review	CERT-In observation, not independently verified	vendor
Agentic tooling cut the time from initial access to final impact "from days to a few hours"	CERT-In observation from one campaign	vendor
Victims that remediated incompletely were attacked again within days, at greater scale	CERT-In observation from one campaign	vendor
Agents re-derived and re-used the same exposed credential in separate operations	CERT-In observation from the disclosed incidents	vendor
Retain logs for at least 180 days	CERT-In recommendation	vendor

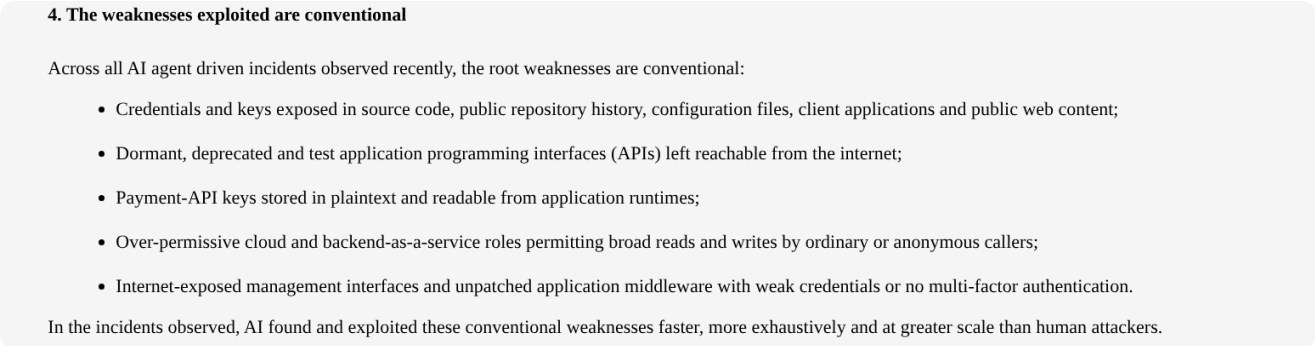
Why this matters for an MSP

The headline is about AI, but the root causes are not new. CERT-In says it directly: "Across all AI agent driven incidents observed recently, the root weaknesses are conventional." It lists five:

- Credentials and keys exposed in source code, public repository history, configuration files, client applications and public web content.
- Dormant, deprecated and test APIs left reachable from the internet.
- Payment-API keys stored in plaintext and readable from application runtimes.
- Over-permissive cloud and backend-as-a-service roles that let ordinary or anonymous callers read and write broadly.
- Internet-exposed management interfaces and unpatched middleware with weak credentials or no multi-factor authentication.

Most MSPs already work against parts of that list. What changes is speed and thoroughness. In CERT-In's words, AI found and exploited these weaknesses "faster, more exhaustively and at greater scale than human attackers." A key that sat unnoticed in an old repository is now something an agent can find, use and find again.

The detection finding is the delivery question. If a client of yours were one of the affected organizations, the first question would be why nobody on your side saw it. That is a question about what your agreement covers and what your tooling actually watches.



4. The weaknesses exploited are conventional

Across all AI agent driven incidents observed recently, the root weaknesses are conventional:

- Credentials and keys exposed in source code, public repository history, configuration files, client applications and public web content;
- Dormant, deprecated and test application programming interfaces (APIs) left reachable from the internet;
- Payment-API keys stored in plaintext and readable from application runtimes;
- Over-permissive cloud and backend-as-a-service roles permitting broad reads and writes by ordinary or anonymous callers;
- Internet-exposed management interfaces and unpatched application middleware with weak credentials or no multi-factor authentication.

In the incidents observed, AI found and exploited these conventional weaknesses faster, more exhaustively and at greater scale than human attackers.

Image: CERT-In, Advisory CIAD-2026-0049 (Oct 15, 2026), section 4.

Map the five gaps to an owner

These are Consultiply's suggestions for how an MSP could assign the work, built on CERT-In's recommended actions for organizations and service providers.

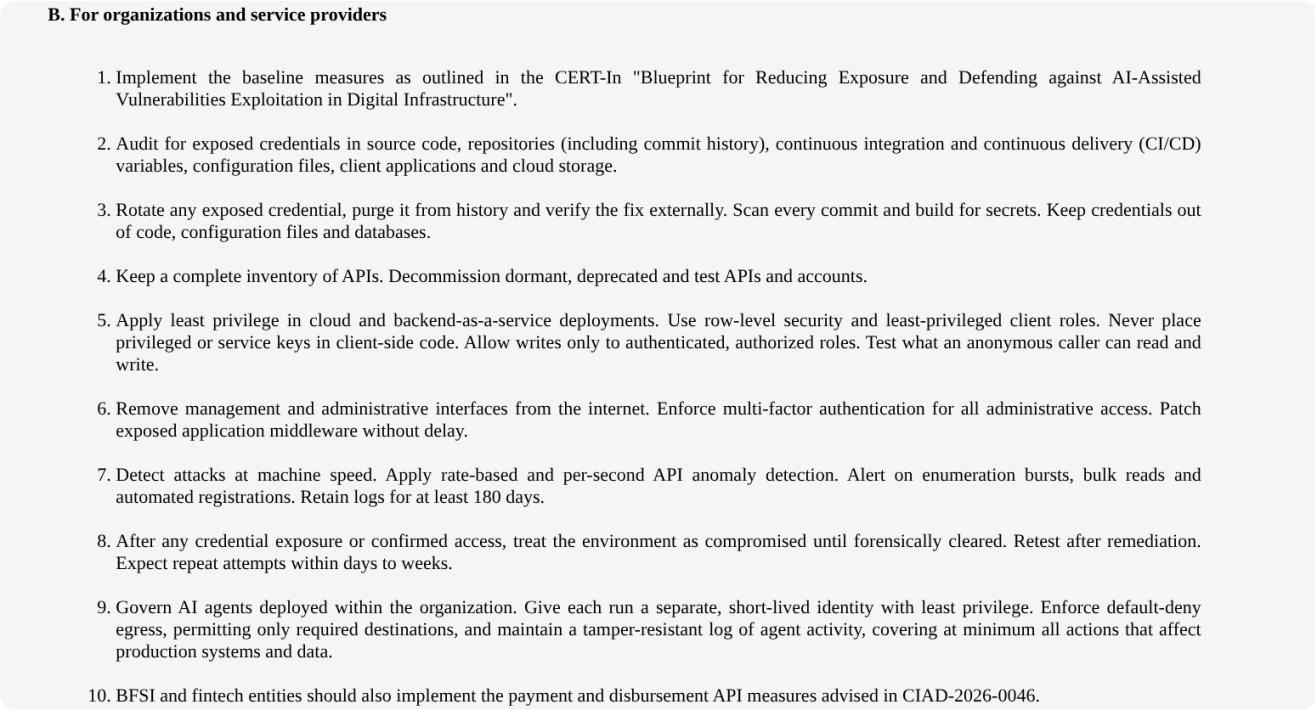
- **Exposed secrets.** CERT-In says to audit source code, repositories including commit history, CI/CD variables, configuration files, client applications and cloud storage, then rotate any exposed credential, purge it from history and "verify the fix externally." Decide whether secret scanning sits with you or with the client's developer, and write it down.
- **Forgotten APIs.** Keep a complete API inventory and decommission dormant, deprecated and test APIs and accounts. Someone has to own the list.
- **Cloud roles.** Apply least privilege, never put privileged or service keys in client-side code, and "test what an anonymous caller can read and write." It is cheap to repeat on a schedule.
- **Admin interfaces.** Take management interfaces off the internet, enforce multi-factor authentication for all administrative access and patch exposed middleware without delay.
- **Detection and logs.** Apply rate-based API anomaly detection, alert on enumeration bursts, bulk reads and automated registrations, and keep logs for at least 180 days. Check each client's current retention against that number.

> MSP lesson: a review against CERT-In's five weaknesses is a scoped project you can repeat every quarter. Sell the review and the retest, not a promise that agents cannot get in.

Govern the agents your clients run

CERT-In's list also covers agents an organization deploys itself: "Give each run a separate, short-lived identity with least privilege." It adds default-deny outbound access that permits only required destinations, and "a tamper-resistant log of agent activity, covering at minimum all actions that affect production systems and data."

The advisory's guidance for AI labs holds one rule worth applying to client agents too: "Enforce the restriction at the network layer, not through instructions to the model." A prompt that says "do not touch production" is not a control. A firewall rule is.



B. For organizations and service providers

1. Implement the baseline measures as outlined in the CERT-In "Blueprint for Reducing Exposure and Defending against AI-Assisted Vulnerabilities Exploitation in Digital Infrastructure".
2. Audit for exposed credentials in source code, repositories (including commit history), continuous integration and continuous delivery (CI/CD) variables, configuration files, client applications and cloud storage.
3. Rotate any exposed credential, purge it from history and verify the fix externally. Scan every commit and build for secrets. Keep credentials out of code, configuration files and databases.
4. Keep a complete inventory of APIs. Decommission dormant, deprecated and test APIs and accounts.
5. Apply least privilege in cloud and backend-as-a-service deployments. Use row-level security and least-privileged client roles. Never place privileged or service keys in client-side code. Allow writes only to authenticated, authorized roles. Test what an anonymous caller can read and write.
6. Remove management and administrative interfaces from the internet. Enforce multi-factor authentication for all administrative access. Patch exposed application middleware without delay.
7. Detect attacks at machine speed. Apply rate-based and per-second API anomaly detection. Alert on enumeration bursts, bulk reads and automated registrations. Retain logs for at least 180 days.
8. After any credential exposure or confirmed access, treat the environment as compromised until forensically cleared. Retest after remediation. Expect repeat attempts within days to weeks.
9. Govern AI agents deployed within the organization. Give each run a separate, short-lived identity with least privilege. Enforce default-deny egress, permitting only required destinations, and maintain a tamper-resistant log of agent activity, covering at minimum all actions that affect production systems and data.
10. BFSI and fintech entities should also implement the payment and disbursement API measures advised in CIAD-2026-0046.

Image: CERT-In, Advisory CIAD-2026-0049 (Oct 15, 2026), recommended actions for organizations and service providers.

What to do this quarter

1. **Sweep your own estate first.** Scan your scripts, repositories and tool integrations for exposed keys before you offer the review to anyone else.
2. **Pilot with three clients.** Check the five weaknesses, record what you found and who fixed it.
3. **Write down who owns each gap.** Secrets, API inventory, cloud roles, admin exposure and logging, per client, in the agreement or a one-page appendix.
4. **Check log retention against 180 days.** Note where it falls short and what closing the gap would take.
5. **Plan for the second attempt.** CERT-In says to treat the environment as compromised "until forensically cleared," retest after remediation and "expect repeat attempts within days to weeks." Put a retest date on the ticket when you close an exposure.
6. **Inventory client agents.** For every agent a client runs, list its identity, its permissions and where its activity is logged.

Read with care

- **No names, no counts.** CERT-In does not name the AI model providers or the affected organizations, and it does not say how many incidents occurred. Treat any report that attaches names as unconfirmed unless the provider says so.
- **Indian incidents.** The observations come from CERT-In's caseload of Indian ICT assets; the five weaknesses are generic.
- **One campaign behind the speed figure.** "Days to a few hours" describes the financially motivated campaign CERT-In observed, not every attack.

- **A referenced baseline we did not review.** CERT-In points organizations first to its "Blueprint for Reducing Exposure and Defending against AI-Assisted Vulnerabilities Exploitation in Digital Infrastructure." This guide does not summarize that document.

> This guide is Consultiply's independent analysis of CERT-In's own advisory, translated for MSP operators. It is not affiliated with or endorsed by CERT-In, and it is not a product recommendation.

Read the full advisory: CERT-In Advisory CIAD-2026-0049.

Sources

CERT-In Advisory CIAD-2026-0049: AI Agents as Attack Vectors (CERT-In, Ministry of Electronics and Information Technology, Government of India, Oct 15, 2026; page last updated Oct 16, 2026)

CERT-In announcement of the advisory on X (CERT-In, Oct 16, 2026, India time)

`Read the full guide: consultiply.ai/mip-resource-center/cert-in-ai-agent-attack-vectors-msp`