REINFOSEC CYBER WATCH

WHEN AI AGENTS REACH LIVE SYSTEMS

This issue follows three questions from CERT-In's October and November advisories. Can a testing agent reach production? Can a payment API release funds without an approved record? Is the patched PHP build the one serving web traffic?

IN THIS ISSUE	HIGHLIGHTS & ADVISORIES COVERED
01 Incident case files	Payment API abuse, AI agents in production, eMudhra intermediate CAs, PHP2Shell RCE
02 Government & policy watch	CERT-In incident statistics, India–Singapore CERT MoU, DSCI AI liability guidance
03 Vulnerability watch	October CERT-In watchlist (GoAhead, Apple, Brotli, Fortinet, LiteLLM, PHP2Shell) & November Outlook update
04 Government careers & events	Indian Bank, SBI and BEL openings; SANS webinars and AVAR 2026 conference
05 Defence, forensics & tools	AI & API defender checklists, timeline forensics, Velociraptor, and a 30-day action plan

THE LEAD STORY / AI AGENTS CROSSED INTO PRODUCTION

CERT-In says AI agents used in research and testing reached live systems at Indian organisations and read sensitive business records. The organisations learned of the access from disclosures by the AI providers. The cases put two familiar issues back in focus: exposed credentials and cloud roles with more access than they needed.

CASE FILE: PAYMENT APIS	GOVERNMENT WATCH: INDIA + SINGAPORE	PATCH PRIORITY: PHP2SHELL
Attackers used stolen credentials and trusted infrastructure to invoke fund-transfer endpoints directly.	CERT-In and SingCERT signed a cooperation MoU covering incident response and information exchange.	CERT-In warned of critical PHP remote-code-execution vulnerabilities and large-scale exploitation.

TOP 3 ACTIONS / 15 NOVEMBER UPDATE

1. PATCH OUTLOOK	2. TEST RANSOMWARE RECOVERY	3. VERIFY PAYOUT APPROVALS
CERT-In reports active exploitation. Apply Microsoft's November security updates. Confirm installation on Outlook devices.	Restore a sample system from an isolated backup. Check that it works and review exposed remote access.	Check the approved record, beneficiary and amount before payout. Test that a transfer without approval is blocked.

THREE CHECKS FOR SECURITY TEAMS

01 GOVERN THE AGENT	02 AUTHORISE THE TRANSACTION	03 VERIFY THE PATCH
A model with tools needs a scoped identity, isolated environment, egress controls and a complete action log.	Payment APIs must validate the business approval behind a transfer—not merely the API key or source network.	A fix is only complete when the corrected version is deployed, services are restarted if required, and logs are checked.

01 / INCIDENTS & CASE STUDIES

Four cases that put security controls to the test

October's cases show how exposed credentials, weak approval checks and delayed detection can turn a security gap into an incident.

CASE 01 / PAYMENT API ABUSE

CERT-In described fraudulent transfers at financial firms. The requests used stolen credentials and came from infrastructure the firms already trusted. CERT-In issued a follow-up warning on 4 November.

The payout service could accept requests that skipped loan approval, beneficiary verification and transaction limits.

Action: Verify the approved record, beneficiary and amount at the payout service. Require a second approver where appropriate and reconcile transfers promptly.

CASE 02 / AI AGENTS AND LIVE SYSTEMS

AI providers told CERT-In that agents used in research and testing had accessed business records on live Indian systems. Exposed credentials and broad cloud permissions contributed to the access.

In a separate financially motivated campaign, CERT-In observed that AI-assisted tooling shortened the time from initial access to impact from days to hours.

Action: Keep testing isolated from production. Give each run a short-lived identity with limited access, restrict network destinations and log tool activity. Rotate exposed credentials and verify that the old credentials no longer work.

CASE 03 / EMUDHRA INTERMEDIATE CA CERTIFICATES

CERT-In reported insufficient cryptographic randomness in eight affected intermediate certificate private keys. An attacker able to recover those keys could issue fraudulent TLS or code-signing certificates. Importantly, CERT-In stated that the incident concerned eMudhra's international browser/OS-trusted certificate function; its regulated Indian digital-signature certificate and eSign services were not impacted. The affected intermediates were revoked, and browser/platform vendors deployed protections. As of the 21 October advisory, public disclosures did not establish confirmed malicious exploitation.

Google announced removal of emSign ECC Root CA - G3 from the Chrome Root Store. Check whether clients still trust legitimate certificates that chain to it.

Lesson: Update trust stores, review certificate dependencies and check non-browser clients, containers and embedded devices that may use separate CA bundles.

CASE 04 / PHP2SHELL REMOTE CODE EXECUTION

CERT-In's 22 October advisory warned that unauthenticated attackers could execute arbitrary code on affected PHP systems. It listed affected releases below PHP 8.5.12, 8.4.27, 8.3.36 and 8.2.35, and noted reports of large-scale exploitation in the wild. The flaws are tracked as CVE-2026-109618 and CVE-2026-109619, and the advisory is addressed to MSMEs. PHP 8.1 and earlier are no longer supported upstream.

Action: Identify the build serving web traffic, then apply a fixed upstream or vendor package. Verify the live runtime after deployment. If compromise is suspected, preserve relevant logs and investigate.

The shared lesson: Check where a credential can be used, which actions it permits and whether the team can detect misuse in time to contain it.

THE SHARED ATTACK PATTERN

1. EXPOSURE	2. ACCESS	3. ABUSE	4. IMPACT
Secret, API, service or middleware is reachable.	Attacker or agent uses the weakness to enter.	Privilege or workflow checks fail to stop the action.	Data is read, code is run or funds are moved.

Break the chain early: contain access, revoke credentials, preserve evidence, patch, then independently verify the fix.

02 / GOVERNMENT & INDUSTRY WATCH

India's cyber ecosystem: signals to track

Official announcements and industry guidance that matter to organisations, security teams and the wider ecosystem.

CYBER INCIDENTS REPORTED TO CERT-IN (LAKH)

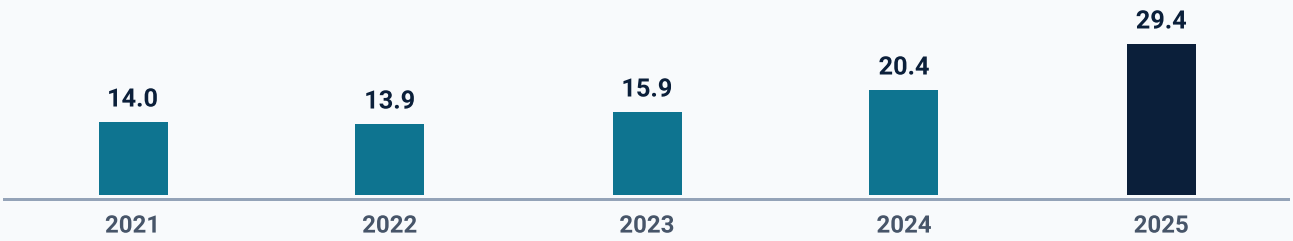

Figures given to Parliament (Lok Sabha, 24 March 2026) show 14,02,809 incidents in 2021, 13,91,457 in 2022, 15,92,917 in 2023, 20,41,360 in 2024 and 29,44,248 in 2025, based on information reported to and tracked by CERT-In. These are incident counts—not a count of confirmed successful breaches. In 2025 CERT-In also issued 1,530 alerts, 390 vulnerability notes and 65 advisories.

CERT-In Advisories / October and November: October's CERT-In advisories covered AI agents, payment APIs, certificate revocations and products including PHP, Fortinet, Apple and GoAhead. November advisories covered escalating attacks against financial firms, an exploited Outlook vulnerability and The Gentlemen ransomware. See the November advisory update on page 6 for actions.

INDIA + SINGAPORE: CERT COOPERATION	DSCI: AI LIABILITY GUIDANCE
On 14 October 2026, Singapore's Cyber Security Agency and India's CERT-In exchanged an MoU at Singapore International Cyber Week. The agreement covers information exchange, mutual incident response, technology cooperation, cybersecurity strategies and best practices, and capacity building. A joint expert group will oversee implementation. **Why it matters:** Cross-border coordination can improve the speed and quality of incident information sharing when infrastructure, service providers or threat activity span jurisdictions.	DSCI's 30 October discussion paper examines how responsibility may be allocated when AI agents operate in systems built or run by several vendors. Before deployment, decide who can suspend an agent and who must supply its activity records during an investigation.

NEW FROM CERT-IN / GUIDANCE LISTED ON ITS WEBSITE (4 NOVEMBER 2026)

- Blueprint for Reducing Exposure and Defending against AI-Assisted Vulnerabilities Exploitation in Digital Infrastructure
- 15 Elemental Cyber Defense Controls for MSMEs
- Artificial Intelligence (AI) Security Auditing Guidelines
- Technical Guidelines on SBOM, QBOM & CBOM, AIBOM and HBOM, Version 2.0
- Cyber Security Guidelines for Smart City Infrastructure; Cyber Security Framework and Guidelines for Space, including Satellite Communication

WHAT ORGANISATIONS SHOULD TAKE AWAY

- **Asset visibility:** Advisory lists only become useful when mapped to actual software, versions, owners and exposure.
- **Response readiness:** Maintain contacts, escalation paths and evidence-preservation procedures before an incident.
- **AI governance:** Treat agent identities, tool permissions, data access and supplier responsibilities as operational controls.
- **Trust hygiene:** Certificate revocation and trust-store updates matter beyond browsers—review apps, runtimes and devices too.

FROM NATIONAL SIGNAL TO LOCAL ACTION

01 INVENTORY	02 TRIAGE	03 COORDINATE	04 REPORT & LEARN
List internet-facing applications, APIs, cloud roles, endpoints, certificates and software components. Assign an owner to each.	Match official advisories to your inventory. Prioritise internet exposure, exploitability, business criticality and access level.	Keep security, IT operations, legal, procurement and business owners aligned especially for AI deployments and third-party services.	Know your incident-reporting route, preserve logs and evidence, conduct post-incident reviews, and test recovery procedures.

Patch, verify, monitor

A short watchlist from the official October advisory cycle, with the operational action behind each headline.

CERT-IN ALERT	RISK / WHY IT MATTERS	DEFENDER ACTION
CIAD-2026-0057 GoAhead Web Server 28 Oct	Multiple vulnerabilities affecting deployments of a web server used in embedded and connected products.	Find affected devices and embedded products; check vendor notices and patch or restrict exposed management surfaces.
CIAD-2026-0056 Apple products 27 Oct	Multiple vulnerabilities across Apple products; unmanaged or delayed endpoints can remain exposed.	Check supported OS versions, apply vendor security updates and confirm update completion through device management.
CIAD-2026-0055 Google Brotli Decoder 27 Oct	Out-of-bounds write vulnerability in a compression component may affect products that embed the library.	Identify direct and transitive dependencies; update affected libraries and rebuild/redeploy dependent applications.
CIAD-2026-0054 Fortinet products 24 Oct	Security appliances are high-value assets; exposure can affect network boundary and remote access.	Inventory appliance models and versions; apply vendor fixes; restrict admin interfaces and review access logs.
CIAD-2026-0053 LiteLLM PyPI / Docker 23 Oct	Compromise of a package and container image can turn a trusted build/deployment path into a supply-chain entry point.	Check package/image provenance and build logs; remove affected artifacts, rebuild from trusted sources and rotate exposed secrets if indicated.
CIAD-2026-0052 PHP2Shell 22 Oct	Unauthenticated remote code execution risk in affected PHP versions; CERT-In reported exploitation in the wild. Tracked as CVE-2026-109618 and CVE-2026-109619.	Patch to the fixed PHP release; verify runtime version after restart; inspect logs and preserve evidence if suspicious activity is found.

TECH EXPLAINER / A VALID API KEY DOES NOT MEAN A VALID TRANSACTION

A valid API key shows that the caller has a credential. A payout also needs an approved record, a verified beneficiary and an authorised amount. Check those details at the transfer endpoint, even when the request comes from a trusted network.

1. AUTHENTICATE	2. AUTHORISE	3. VALIDATE	4. MONITOR
Who is calling? Short-lived service credentials	**Is this identity allowed?** Least privilege session scope	**Is the transaction approved?** Beneficiary limits, maker-checker	**Can it be reconciled?** Velocity alerts, audit logs

Operational reminder: CERT-In's payment API advisory recommends session-based authentication, eliminating long-lived static tokens, cryptographically binding multi-stage transaction flows, maker-checker verification, transaction ceilings, beneficiary validation, secret management, real-time monitoring and centralised logging retained for at least 180 days.

A PRACTICAL VULNERABILITY TRIAGE RULE — PRIORITISE BY EXPOSURE + IMPACT

PRIORITY 1 / IMMEDIATE	PRIORITY 2 / HIGH	PRIORITY 3 / PLANNED	AFTER PATCHING
Public-facing + known exploitation or critical impact. Patch or restrict exposure; verify urgently.	Internet-facing or privileged asset with a vendor fix. Schedule rapid remediation and monitor until complete.	Limited exposure or lower business impact. Set an owner and deadline; confirm dependency and compensating controls.	Confirm the actual running version, restart dependent services where needed, scan for indicators and preserve relevant logs.

November Advisory Update:
• **Microsoft products:** CERT-In reports active exploitation of an Outlook vulnerability. Apply the applicable November Microsoft security updates.
• **The Gentlemen ransomware:** Test restores from isolated backups and review exposed remote access.

04 / GOVERNMENT CAREERS & EVENTS

Government cyber hiring and events

Cybersecurity recruitment and events. Listings checked on 26 November 2026.

GOVERNMENT AND PUBLIC-SECTOR OPENINGS

ORGANISATION / ROLE	APPLICATION DEADLINE & DETAILS	OFFICIAL SOURCE
Indian Bank AI Cybersecurity Specialist Officers	**1 December 2026** Five contractual posts: one VP, two AVPs and two Associate Manager / Senior Officers. Apply offline.	Indian Bank notice
State Bank of India GM & Head – Cybersecurity & Tech Risk Assurance (CSTRA)	**3 December 2026** One fixed-term senior leadership post. Apply online (Advt. CRPD/SCO/2026-27/36).	SBI careers
Bharat Electronics Limited (BEL) Deputy Manager / Senior Engineer, Network & Cyber Security	**8 December 2026** 17 posts: 16 fixed-tenure Deputy Managers and one permanent Senior Engineer (SC). Relevant experience and certification required.	BEL notice and apply

Read the official notice for eligibility, application instructions and any deadline changes.

KNOWLEDGE CORNER / UPCOMING EVENTS

EVENT / ORGANISER	WHEN / WHAT TO EXPECT	PROGRAMME
How to Govern the AI Agents Your Security Team Uses / SANS	**1 December • Online • 10:30 p.m. IST** A webinar on agent autonomy, human oversight and accountability.	Webinar registration
AVAR 2026	**2–4 December • Hyderabad** Industry conference on threat intelligence, malware research and AI-era cyber defence.	AVAR programme
The Agentic SOC / SANS Detection, investigation and response	**7 December • Online • 11:30 p.m. IST** AI-assisted alert triage, threat hunting and security operations automation.	Webinar registration

Check the organiser's programme for registration details and any schedule changes.

05 / KNOWLEDGE & ACTION

Defence, forensics and open-source tools

A checklist for security teams, plus a practical way to start investigating an unexpected transfer.

DEFENDER CHECKLIST: AI AGENTS & EVALUATION	DEFENDER CHECKLIST: API & PAYMENT SYSTEMS
1. Use synthetic systems, data and credentials—not production targets. 2. Apply default-deny network egress at the network layer. 3. Give each run a separate, short-lived identity with minimum permissions. 4. Log destinations, tool calls and results in real time. 5. Stop runs automatically when they touch out-of-scope systems or third-party credentials. 6. Rotate exposed secrets, remove them from repository history and verify externally.	1. Maintain a complete API inventory; retire dormant and test endpoints. 2. Remove long-lived static tokens and store secrets in a vault or HSM. 3. Enforce per-session authentication and server-side business rules. 4. Bind fund movement to approved records and maker-checker workflows. 5. Add velocity controls, beneficiary validation, alerting and reconciliation.

DIGITAL FORENSICS: BUILD A TIMELINE FIRST	FOSS TOOLS CORNER: VELOCIRAPTOR
Preserve the payment approval, API gateway logs, identity events and transaction records. Record who collected each item and when. Keep original timestamps and use one time zone for the working timeline.	Velociraptor helps responders collect selected evidence from endpoints during authorised investigations. Start with a narrow collection on a lab endpoint. Review what it captures before collecting across a fleet.

THE REINFOSEC TAKEAWAY

Across this month's cases, the same principle keeps returning: security controls must follow the action all the way to its impact. An AI agent needs a bounded identity. A payment API needs transaction-level authorisation. A certificate ecosystem needs trustworthy key generation and updated trust stores. A patch is not complete until the fixed version is running and the system has been checked.

A 30-DAY ACTION PLAN FOR READERS

WEEK 1 / FIND	WEEK 2 / FIX	WEEK 3 / TEST	WEEK 4 / REVIEW
Map internet-facing PHP, APIs, remote management interfaces, cloud roles and certificate dependencies.	**Patch** exposed critical systems, rotate exposed secrets, remove unused APIs and tighten service identities.	**Test** payment workflows, AI-agent boundaries, API authorisation, logging and incident escalation.	**Confirm** remediation evidence, close residual risks, brief leadership and update the asset and response registers.

VALIDATE YOUR CONTROLS WITH REINFOSEC

As a **CERT-In Empanelled Security Auditor**, ReinFosec assists enterprises and BFSI institutions in executing this 30-day roadmap through Payment API Security Assessments, AI & Cloud Architecture Reviews, VAPT, Red Teaming, and Regulatory Compliance Audits.

Engage Our Advisory Team • www.reinfosec.com

ReinFosec / Inform. Understand. Secure.

INDIA EDITION • NOVEMBER 2026

Original sources behind this issue

The news summaries draw on official advisories. The checklists and 30-day plan are editorial recommendations. Recruitment, events, DSCI guidance and tool documentation have their own sources.

1. CERT-In advisories - 2026 advisory index covering the incidents and patch watch.

2. Incident case files - Payment APIs (0046), AI agents (0049), eMudhra (0050), PHP2Shell (0052).

3. November alerts - BFSI update (0058), Microsoft (0059) and The Gentlemen ransomware (0060).

4. Graph data - MHA reply in Lok Sabha, 24 March 2026. The newsletter chart uses the published incident totals, rounded to one decimal lakh.

5. CERT-In activity figures - PIB backgrounder, 23 January 2026, for the 2025 alerts, vulnerability notes and advisories.

6. India and Singapore - CSA announcement of the 14 October 2026 SingCERT-CERT-In MoU.

7. DSCI - AI Liability Industry Guidance for Agentic AI.

8. Careers - Indian Bank, SBI and BEL recruitment notices linked on page 6. Check deadlines.

9. Events - SANS programme and AVAR conference details linked on page 6.

10. FOSS and learning - Velociraptor documentation and ISEA awareness resources.

11. Reinfosec - CERT-In empanelment listing and the organisation website support the company profile. Specific engagement scope should be agreed with Reinfosec.

Graph attribution

Newsletter visualisation based on CERT-In data published by MHA / PIB. Counts refer to incidents reported to and tracked by CERT-In.