SITUATIONAL REPORTS

Crafty Falcon and India's Financial Sector: AI-Driven Cyberattacks Expose a New Security Threat

Author: CNAWS Research Team • **Published:** November 06, 2026 • **Source:** Centre for New Age Warfare Studies



EXECUTIVE BRIEF & KEY FINDINGS

CERT-In has identified an active cyberattack campaign targeting Indian financial institutions. The use of autonomous AI tools to exploit payment infrastructure raises serious questions about the ability of conventional cybersecurity systems to respond to machine-speed intrusions.

Situation Overview

India's financial sector is confronting an escalating campaign of artificial intelligence-assisted cyberattacks that has exposed vulnerabilities in the country's rapidly expanding digital financial infrastructure. On 4 November 2026, the Indian Computer Emergency Response Team (CERT-In) issued a critical-severity advisory identifying a threat actor designated CRAFTY FALCON, which has been targeting non-banking financial companies (NBFCs), digital lending platforms, fintech companies, payment aggregators and cryptocurrency exchanges. The advisory identifies the use of publicly available AI-driven offensive security frameworks to automate intrusions, exploit payment systems and execute unauthorised financial transactions. More concerningly, CERT-In has confirmed that some institutions were successfully attacked again despite having undertaken remedial measures following earlier warnings. The campaign remains active,

with attackers reportedly exploiting weekends, holidays and other periods of reduced security monitoring to maximise their operational advantage.[1]

The development represents a significant shift in the cyber threat environment facing India's financial institutions. Although automated cyberattacks are not new, the increasing availability of autonomous AI frameworks allows attackers to combine reconnaissance, vulnerability identification, exploitation and post-compromise activity into compressed operational sequences. The principal concern is therefore not simply that artificial intelligence is being used to conduct cybercrime, but that the speed at which intrusions can be executed is beginning to challenge the assumptions underlying conventional defensive architectures.

How Crafty Falcon Is Targeting Financial Infrastructure

CERT-In's findings indicate that Crafty Falcon has been exploiting weaknesses across interconnected cloud services, application environments and payment interfaces. The attackers have reportedly compromised internet-facing systems, including testing and staging environments, before moving into production infrastructure through shared credentials, authentication tokens and network connections. In certain incidents, legitimate cloud management services were abused to execute commands and introduce malicious files. Once access was obtained, the attackers targeted payment and disbursement application programming interfaces (APIs), which allow financial institutions and their technology partners to initiate transactions. By exploiting these interfaces, malicious actors could initiate transfers through systems that were otherwise designed to facilitate legitimate financial activity.[2]

The October 2026 CERT-In advisory had already warned that attackers were using compromised payment credentials and trusted infrastructure to bypass controls associated with loan approvals, beneficiary verification and transaction limits. The November advisory provides further evidence of an evolving campaign. Attackers have reportedly performed micro-value transactions, sometimes involving as little as ?1, to test whether unauthorised fund transfers would succeed before attempting larger transactions. CERT-In has also identified the abuse of merchant settlement APIs and incidents involving the theft of cryptocurrency wallet credentials. These methods illustrate a particularly consequential vulnerability in modern financial infrastructure. When the systems responsible for authorising or transmitting transactions are compromised, fraudulent instructions may appear to originate from trusted applications rather than from an external attacker.[3]

The AI Dimension and the Compression of Attack Timelines

The latest advisory specifically identifies four publicly available offensive security frameworks associated with the campaign. These include ARTEX, an LLM-based multi-agent offensive security platform, PentAGI, an autonomous penetration-testing system, Cairn, an attack search engine, and CyberStrikeAI. Their significance lies in their ability to coordinate different stages of an intrusion, generate customised scripts and modify their approach in response to technical obstacles. Instead of requiring an operator to manually investigate each vulnerability and develop an appropriate exploit, AI-assisted frameworks can automate substantial portions of the process and pursue multiple avenues of access simultaneously.?

The implications extend beyond the capabilities of these individual tools. In a separate advisory issued on 15 October, CERT-In observed that agentic AI tooling had compressed the period between initial access and final impact from days to a few hours in financially motivated intrusions. Similar developments have been documented internationally. Research published by Google's Threat Intelligence Group and Mandiant on 30 October identified the operational use of ARTEX and other agentic frameworks in real-world cyber intrusions. Their investigations found that autonomous systems could conduct reconnaissance, test permissions and adapt attack procedures in rapid sequences, although these operations frequently generated substantial forensic evidence. This distinction is important. AI-enabled attacks

are not necessarily more sophisticated or difficult to detect than conventional intrusions. Their principal advantage is that they can complete damaging actions before defenders have sufficient time to identify, investigate and contain them. ?

For Indian financial institutions, this changes the relationship between detection time and financial exposure. Security mechanisms built around human-led investigation and sequential incident response may prove inadequate when an adversary can identify a vulnerable interface, obtain access and execute multiple transactions within minutes. The problem becomes especially serious when attackers operate through legitimate credentials and existing payment channels, where individual activities may not immediately appear malicious.

The China Connection and the Problem of Attribution

The international origins of some of the identified tools introduce an additional strategic dimension. Reuters reported on 5 November that security researchers had linked at least two of the frameworks identified by CERT-In to China. ARTEX, in particular, has reportedly been associated with attacks on South Korean financial institutions and with operations involving China-linked cyber threat actors. However, the availability of these frameworks across public or previously public repositories means that their use cannot independently establish the nationality, organisational affiliation or strategic motivation of the operators. **CERT-In has neither attributed Crafty Falcon to the Chinese state nor identified the geographical location of the group.**?

This distinction is essential to understanding the emerging threat. The proliferation of AI-assisted offensive frameworks is reducing the technical barriers that previously separated well-resourced state-linked cyber operators from financially motivated criminal groups. Capabilities that once required specialised personnel, considerable development time and sophisticated infrastructure can increasingly be assembled from accessible software components. Consequently, the presence of a tool associated with a particular country provides limited evidence about the identity of the actor using it. For India, the immediate security concern is the wider diffusion of advanced offensive capabilities rather than an unverified conclusion about Chinese sponsorship of the present campaign.

Strategic Implications for India's Economic Security

Crafty Falcon exposes a structural weakness created by the growing interdependence of India's financial institutions and technology service providers. Digital lending, automated disbursements, payment aggregation and cloud-based financial applications depend on extensive networks of APIs, credentials and third-party infrastructure. These connections improve efficiency and enable financial inclusion, but they also create pathways through which a compromise in one component can affect other systems. An attacker does not necessarily need to penetrate a major bank's core infrastructure to cause substantial financial losses. Compromising an intermediary with trusted access to payment services may provide an alternative route to the movement of funds.

This creates an important distinction between securing individual institutions and securing the broader financial ecosystem. A bank may maintain robust internal cybersecurity controls while remaining exposed to vulnerabilities originating from fintech partners, outsourced service providers or compromised payment interfaces. The consequences of such intrusions are therefore not confined to the immediate victim. Repeated attacks can increase operational costs, undermine confidence in digital financial services and place additional pressure on regulators responsible for maintaining financial stability. However, the publicly available evidence does not establish that Crafty Falcon has compromised India's core payment infrastructure or caused systemic financial disruption. The strategic concern is the potential for localised breaches to propagate through interconnected financial relationships.

The campaign also reveals a growing asymmetry between offensive automation and institutional response. Attackers can operate continuously and exploit temporary reductions in monitoring capacity, while defenders must balance security requirements against the need to process legitimate financial transactions without interruption. CERT-In's warning about repeated compromises following incomplete remediation suggests that restoring services after an incident is insufficient unless the underlying access pathways, compromised credentials and weaknesses in transaction authorisation are comprehensively addressed.

Assessment and Outlook

The immediate response requires financial institutions to move beyond conventional perimeter protection towards continuous monitoring of identities, cloud infrastructure and transaction behaviour. CERT-In has recommended round-the-clock monitoring, isolation of testing environments from production systems, stronger server-side validation of financial transactions, tighter controls over privileged credentials and automated detection of suspicious payment patterns. Particularly important are controls that identify small test transfers followed by unusually large disbursements, changes in beneficiary information and abnormal transaction volumes. Coordination between CERT-In, the Indian Cyber Crime Coordination Centre, financial regulators and affected institutions will be critical to containing attacks before stolen funds are dispersed across multiple accounts.?

CNAWS assesses that Crafty Falcon is an early warning of a broader transformation in financially motivated cyber operations. The decisive change is not the emergence of an entirely new method of cyber intrusion, but the increasing ability of relatively small groups to execute established attack techniques at unprecedented speed and scale. As autonomous frameworks become more capable and accessible, the interval between the discovery of a vulnerability and its exploitation is likely to narrow further. India's cybersecurity preparedness must consequently be assessed not only by its ability to prevent intrusions, but also by whether its financial institutions can detect, interrupt and contain malicious activity quickly enough to prevent financial losses.

Over the coming weeks, the principal indicators will be the frequency of repeat attacks, the emergence of additional affected institutions, evidence of increasingly autonomous exploitation and the effectiveness of enhanced transaction-monitoring measures. The Crafty Falcon campaign demonstrates that the security of India's digital financial ecosystem will increasingly depend on the capacity of defensive systems to operate at the same speed as the threats confronting them.

Endnotes

1. Indian Computer Emergency Response Team (CERT-In), “Escalation of AI-Driven Attacks on BFSI Entities by Threat Actor ‘CRAFTY FALCON,’” Advisory CIAD-2026-0058, November 4, 2026. Official advisory.

2. CERT-In, “Threat Actors Exploiting Payment and Disbursement APIs for Unauthorized Fund Transfers in BFSI Entities,” Advisory CIAD-2026-0046, October 9, 2026. Official advisory.

3. CERT-In, “Escalation of AI-Driven Attacks on BFSI Entities,” November 4, 2026. Official advisory.

4. CERT-In, “Escalation of AI-Driven Attacks on BFSI Entities,” November 4, 2026. Official advisory.

5. CERT-In, “AI Agents as Attack Vectors: Unauthorized Access by Agents under Evaluation and Misuse of Open-Source Agentic AI Tools Targeting Indian ICT Assets,” Advisory CIAD-2026-0049, October 15, 2026. Official advisory. See also Google Threat Intelligence Group and Mandiant, “ARTEX and Friends: Enter Agentic Offensive Frameworks,” Google Cloud Blog, October 30, 2026. Full research.

6. Munsif Vengattil and Ashwin Manikandan, “India Flags Chinese-Linked AI Tools in Hacking Campaign against Financial Firms,” *Reuters*, November 5, 2026. Full report.

7. CERT-In, “Escalation of AI-Driven Attacks on BFSI Entities,” November 4, 2026. Official advisory. See also Akashvani News, “I4C Cautions Banks and Fintech Firms against AI-Driven Cyberattacks,” November 5, 2026. Report.

Header Image Credits: Photo by Philipp Katzenberger on Unsplash