HyperNext Data Centers

RESEARCH PAPER

HN-RP-009

Frontier AI and the Indian Banking System

Cybersecurity risks from Mythos-class AI and a framework of safeguards for the financial sector

An independent assessment prepared for the consideration of the Hon'ble Union Finance Minister.

This is what is Next.

Series	HyperNext Research
Paper	HN-RP-009
Issued	04 June 2026
Version	1.0
Classification	Public release
Citation	HyperNext Research, HN-RP-009

Frontier AI and the Indian Banking System

This paper is part of the HyperNext Research series. Methodology, assumptions, and source data are stated openly so other operators can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-009" is welcome.

Contents

1. What changed in April 2026

ABSTRACT

On 8 April 2026, Anthropic PBC released Claude Mythos Preview. In published evaluation across 198 manually-reviewed vulnerability findings, the model agreed with expert human severity assessment in 89 per cent of cases. Across ten fully-patched targets in internal Anthropic testing, the model achieved complete control-flow hijack. The implications for Indian banking are direct: UPI processed 21 billion transactions worth INR 27 lakh crore in December 2025 alone, the system depends on shared technology service providers serving up to 300 cooperative banks each, and the prevailing time-to-patch is measured in weeks. The window between vulnerability discovery and weaponisation has collapsed from months to hours. This paper sets out a seven-pillar framework of safeguards: indicative cost INR 4,200 to 5,800 crore over three years; equivalent to 3.5 to 4.8 per cent of the INR 1.2 lakh crore that the Indian Cyber Crime Coordination Centre estimates Indians lost to cyber fraud in 2025.

India's financial sector is among the most digitally integrated in the world and one of the most concentrated by vendor count. Both characteristics now register as risk.

Status and authorship

This paper is prepared by Harsh Macwan in personal capacity. It is an independent analysis. It does not represent the views of, and has not been prepared in consultation with, the Reserve Bank of India, the Indian Computer Emergency Response Team, the Computer Security Incident Response Team for the Financial sector, the National Critical Information Infrastructure Protection Centre, the Ministry of Finance, or any other authority of the Government of India. It draws on publicly available information: Anthropic's System Card and Alignment Risk Update; assessments by the UK AI Security Institute and Cloud Security Alliance; published policy responses from Japan, the United States, the European Union, and the United Kingdom; and publicly-issued instruments of the Reserve Bank of India and Government of India referenced in the bibliography. The author has not had access to Mythos itself.

What the paper recommends

The framework comprises seven pillars, sequenced over twenty-four months:

- **Pillar 1: Compress the patching cycle.** Mean time-to-patch (MTTP) of 48 to 96 hours by regulatory class, automated vulnerability management, supervisory penalties for non-compliance.
- **Pillar 2: Build sovereign defensive AI capability.** NAICD-BFSI as the national defensive AI capability for the financial sector, jointly with IndiaAI Mission compute and Project Glasswing access.
- **Pillar 3: Critical Technology Service Provider regime.** Direct supervisory authority over major shared technology providers, modelled on the European DORA framework.

- **Pillar 4: Authentication beyond static credentials.** Hardware-bound credentials for privileged access, phishing-resistant out-of-band for high-value transactions, device-binding for UPI.
- **Pillar 5: National Threat Intelligence and SBOM Platform.** Real-time intelligence sharing across the regulated entity base, central SBOM repository, FS-ISAC equivalent under CSIRT-Fin.
- **Pillar 6: Cyber-defence workforce at scale.** National curriculum through AICTE, three-tier certification ladder, bonded scholarship pipeline, 50,000 trained cybersecurity professionals for BFSI by 2030.
- **Pillar 7: Resilience, recovery, consumer protection.** Mandatory recovery time objectives, quarterly AI-augmented resilience exercises, Banking Sector Cyber-Insurance Facility, accelerated customer reimbursement.

Three actions are recommended for immediate Ministerial consideration: constitute an Inter-Ministerial Working Group within 90 days; issue an RBI Mythos-class threat advisory in parallel; initiate government-to-government engagement with Anthropic and peer laboratories for defensive partnership access. These three actions are addressed in Section 9. The substantive case for the framework is set out in Sections 2 through 8.

2. The threat to Indian banking, by the numbers

This section establishes the magnitude of the threat with reference to public data.

Mythos capability in measured terms

Anthropic's Mythos System Card reports the following measurements. SWE-bench coding benchmark: 93.9 per cent (the highest publicly recorded). USAMO mathematics: 97.6 per cent. CyberGym cybersecurity benchmark: materially superior to Claude Opus 4.6 (Anthropic does not disclose the specific delta). Internal Anthropic red-team evaluation of 198 manually-reviewed vulnerability findings: 89 per cent exact agreement with expert human severity assessment, 98 per cent agreement within one severity level. Across ten fully-patched targets, complete control-flow hijack.

The qualitative properties recorded by the UK AI Security Institute in its mid-April 2026 evaluation reinforce these numbers. AISI found that Mythos completed certain multi-step infiltration challenges that no other publicly evaluated AI model had previously completed, while not (yet) outperforming human red teams on individual cybersecurity tasks. The asymmetry is informative: Mythos is at the threshold of multi-step autonomous capability, with successor models projected to cross it.

Historical context: Mythos surfaced a 27-year-old TCP SACK vulnerability in OpenBSD, a 17-year-old NFS zero-day in FreeBSD, and a 16-year-old H.264 flaw in FFmpeg. These are codebases that have been examined by tens of thousands of expert security researchers over those years. The implication for Indian core banking software, which has not received equivalent independent review, is unfavourable.

Indian banking by exposure metrics

UPI: 21 billion transactions in December 2025 alone, INR 27 lakh crore value, 80 per cent of retail digital payment volume nationally. RTGS: clears settlement for inter-bank, foreign exchange, securities, and government payment flows; average daily value exceeds INR 12 lakh crore. Approximately 600 banks participate in UPI; settlement, addressing, and reconciliation are concentrated at NPCI.

Cyber incident base rate: CERT-In handled 29.44 lakh cyber incidents in 2025 across all sectors, with the BFSI sector among the most heavily targeted. The Indian Cyber Crime Coordination Centre reports financial loss from cyber fraud in 2025 of approximately INR 1.2 lakh crore. Post-Pahalgam 2025, CERT-In documented more than 1.5 million directed cyber-attacks on Indian websites, with the BFSI sector among the explicit target categories.

The C-Edge precedent of July 2024 is the most consequential supply-chain incident on record for the Indian banking sector. A misconfigured Jenkins server at Brontoo Technology Solutions, exploiting CVE-2024-23897 (a vulnerability publicly disclosed in January 2024 and left unpatched for seven months),

enabled a RansomEXX v2.0 deployment that disrupted payment services at approximately 300 cooperative and regional banks. Recovery took seven to ten days for the worst-affected institutions. NPCI intervened by temporarily disconnecting the affected banks to limit propagation.

The collapse of the patching window

In the C-Edge precedent, the adversary used a publicly-disclosed vulnerability that the defender had seven months to patch. Mythos-class capability eliminates that window. A Mythos-augmented adversary discovers the vulnerability before the defender does, develops a working exploit within hours, and weaponises it within days. The defender's first signal is the incident itself.

The arithmetic is unfavourable on present industry practice. Indian SCBs report mean time-to-patch for internet-facing systems of 9 to 28 days for Critical-rated vulnerabilities (industry benchmark surveys, 2024-2025). Cooperative banks and their service providers run weeks to months. Mythos-class adversary capability is measured in hours. The defensive cycle is approximately two orders of magnitude slower than the offensive cycle.

The seven-layer Indian banking attack surface

Layer	Description	Principal exposure
L7 Customer channel	Mobile banking, internet banking, third-party UPI, SMS/USSD, branch terminals.	Authentication bypass; 2FA circumvention; deserialisation in mobile SDKs; transaction-confirmation logic flaws.
L6 API and integration	Open Banking APIs, AA APIs, NPCI APIs (UPI, IMPS, NETC), card APIs, BBPS APIs.	Authorisation bypass; rate-limit evasion; request smuggling; insecure direct object references.
L5 Application	Core banking (Finacle, FLEXCUBE, BaNCS, Finzly), switching software, LOS, CRM.	Transaction processing logic; SQL/NoSQL injection variants; batch reconciliation flaws; privilege escalation.
L4 Middleware and database	Application servers, message queues, ESBs, RDBMSes, NoSQL stores, in-memory caches.	Long-tail middleware vulnerabilities (the class Mythos surfaced in FFmpeg and FreeBSD); deserialisation; cache poisoning.
L3 Operating system	Linux, IBM AIX, Microsoft Windows Server, mainframe OSes.	Kernel privilege escalation (class Mythos has demonstrated in Linux, OpenBSD, FreeBSD); container/hypervisor escape.
L2 Network and infrastructure	Routers, switches, firewalls, load balancers, VPN concentrators, ATMs, POS, HSMs.	Network device firmware; DevOps infrastructure (per C-Edge / Jenkins); HSM communication-protocol flaws.

L1 Vendor and supply chain	Technology service providers, software vendors, cloud providers, MSSPs, open-source dependencies.	Indirect compromise (the C-Edge transmission mechanism); SolarWinds-style update poisoning; open-source supply chain.

Mythos-class capability is broad-spectrum across these layers and most damaging where flaws compound between layers. The C-Edge incident illustrates: Layer 1 vulnerability (vendor) → Layer 2 exploitation (Jenkins server) → Layer 7 consequence (300 banks' customers).

3. Four scenarios with consequence estimates

The following scenarios are constructed against the documented Mythos capability set and against observed Indian banking adversary behaviour. Each scenario specifies adversary profile, vector, outcome, and consequence estimate. The scenarios are illustrative rather than predictive.

Scenario A: Supply-chain compromise of a core banking service provider

Adversary. Organised cybercriminal group, motivated by ransom, with Mythos-class capability via unauthorised channels.

Vector. Mythos-class capability directed at the publicly-discoverable software components used by a major core banking service provider serving cooperative banks. Within 72 hours, multiple exploitable vulnerabilities are surfaced. One (typically a deserialisation flaw in a job-scheduling component) is selected for weaponisation. Exploit delivered, remote code execution achieved on the service provider's build infrastructure, lateral movement into the production environment supplying core banking services.

Consequence estimate. Payment processing encrypted at 200 to 400 cooperative and regional banks. ATM withdrawal, UPI participation, and inter-branch settlement disrupted for the affected institutions. Recovery time of 7 to 10 days for the worst-affected, per C-Edge precedent. Direct loss to the financial system from disruption alone (excluding ransom): INR 800 to 1,500 crore. Estimated consumer harm in service interruption (lost income, missed payments, distress): INR 2,000 to 4,000 crore. Reputational consequence not quantified.

Distinction from C-Edge. In C-Edge, the defender had a 7-month window between CVE disclosure (January 2024) and exploitation (July 2024). The vulnerability was publicly known. In this scenario, the vulnerability is autonomously discovered by the adversary, never publicly disclosed; the defender has no advance signal through CVE channels, threat intelligence, or vendor advisories.

Scenario B: Authentication bypass in a major mobile banking application

Adversary. Sophisticated criminal group, possibly with state tolerance, prepared for a multi-month operation.

Vector. Mythos-class static and dynamic analysis of the publicly-distributed binary of one of India's four largest UPI applications. The model identifies a logic flaw analogous to the 2FA-bypass vulnerabilities documented in the Mythos red-team disclosures: the flaw allows the adversary to initiate transactions from compromised devices without triggering out-of-band confirmation. The adversary combines this with AI-generated voice-cloning calls (the technology was demonstrated in the wild in India by late 2025) to harvest device-level credentials.

Consequence estimate. Unauthorised transactions across 100,000 to 500,000 customer accounts before detection. Direct financial loss INR 1,500 to 3,000 crore. Liability allocation under the existing Limited Liability of Customers framework is unclear in this configuration: the customer has done nothing wrong, the regulated entity may also have followed prevailing secure-development practice, and yet the loss occurs at scale. Reputational damage to the affected application is substantial. Spillover effect on consumer trust in UPI as a category may produce a transient (one to three quarter) reduction in UPI transaction growth.

What is qualitatively new. Existing UPI fraud is dominantly social engineering. The Reserve Bank's limited-liability regime is calibrated to that pattern. A Mythos-discovered application-layer flaw is a different category: customer-blameless, entity-defensible, and at scale.

Scenario C: State-aligned APT targeting of UPI / RTGS infrastructure

Adversary. APT group attributed to a state actor adversarial to India, motivated by strategic disruption. CERT-In documented the targeting of Indian BFSI by such groups post-Pahalgam 2025; this scenario assumes the addition of Mythos-class capability to that adversary's toolkit.

Vector. Multi-stage, multi-month operation. Stage 1: Mythos-assisted reconnaissance of UPI and RTGS components, accumulating 50 to 100 zero-day vulnerabilities over 30 days. Stage 2: initial access at participant banks via commodity vectors (phishing, SIM-swap of privileged users); 30 to 60 days; intentionally non-disruptive. Stage 3: Mythos-class identification of privilege-escalation chains, lateral movement to RTGS communication infrastructure; 60 to 90 days. Stage 4: pre-positioning across multiple participant banks; capability dormant. Stage 5: activation on adversary signal correlated with strategic provocation.

Consequence estimate. Suspension or impairment of RTGS for hours rather than minutes. Cascading effects on inter-bank liquidity, foreign exchange settlement, securities settlement through CCIL, broader payment system. Macroeconomic consequence not quantifiable on existing data; a 24-hour RTGS outage would defer settlement on an estimated INR 12 to 18 lakh crore of inter-bank flows.

Comment. This is the scenario for which the Indian banking system is least prepared. The existing operational resilience framework, though materially strengthened by RBI direction over the past decade, has not been stress-tested at this adversary capability level. Mandatory quarterly cyber-resilience exercises against AI-augmented threat scenarios (Pillar 7) are principally directed at this scenario.

Scenario D: Insider-equivalent access via privilege-escalation chain

Adversary. Mid-tier criminal group with no privileged access, operating through commodity initial-access methods.

Vector. Phishing compromise of a customer-service representative at a Scheduled Commercial Bank. The compromised account has minimal privileges. Mythos-class capability identifies a privilege-escalation

chain across the bank's endpoint software stack (flaw in legacy desktop application combined with a kernel vulnerability in the deployed OS version). Escalation from CSR session to administrative access on the internal network; lateral movement to systems containing KYC data.

Consequence estimate. KYC data (Aadhaar, PAN, address, account details) exfiltrated for 5 to 20 million customers. Subsequent monetisation through identity-theft-driven account takeover at unrelated financial institutions. Long-tail consumer harm over 3 to 7 years; direct financial loss INR 4,000 to 12,000 crore depending on monetisation efficiency. Aadhaar-bound consequences include long-term reputational damage to the digital identity programme.

Aggregate consequence band

If any single scenario materialises at central case, the consequence is in the order of INR 1,500 to 4,000 crore in direct financial loss plus comparable indirect costs. If Scenario C materialises, the consequence is unquantifiable but plainly in the macroeconomic category. The framework's three-year incremental cost (INR 4,200 to 5,800 crore) is dominated by even a single mid-case scenario, and is materially smaller than the annual cyber-fraud loss baseline of INR 1.2 lakh crore.

4. The existing regulatory framework, and where it stops

India's BFSI cybersecurity regulatory architecture, examined against the Mythos-class threat, is mature in structural terms and is increasingly miscalibrated to the operational tempo of the threat. This section identifies the elements that work and the gaps that compromise the framework.

The elements that work

Board-level accountability. The 2023 IT Governance Master Direction locates accountability for technology and cybersecurity at the Board, with a Board-level IT Committee. This is consistent with the strategic-risk character of the Mythos-class threat.

Risk-based post-breach orientation. The progression from preventive controls (2011 IT Security Guidelines) to detection-response-recovery (2016 Cyber Security Framework and subsequent) anticipates the environment in which prevention cannot be relied upon as principal control. This is the correct framing for Mythos-class threat.

Source-code and escrow obligations. The 2024 PSO Master Direction's requirement that critical applications be supported by source-code availability or escrow enables the regulated entity, or a successor service provider, to perform independent vulnerability assessment of the kind that AI-assisted capability now makes feasible.

Mandatory incident reporting. The CERT-In six-hour reporting requirement establishes a culture of disclosure. Compliance is imperfect in practice but the framework is correct.

Sectoral CSIRT. CSIRT-Fin within CERT-In gives the BFSI sector a dedicated coordination point.

The gaps material to Mythos-class threat

Seven gaps are visible on the present analysis.

Gap 1: Patching cadence. Existing instruments require timely patching but do not specify maximum MTTP. Industry norm is weeks; Mythos-class adversary capability is hours. The C-Edge precedent (7 months unpatched) is the lower bound, not the median.

Gap 2: Vulnerability management beyond CVE channels. Existing supervisory expectations are built around the assumption that relevant vulnerabilities are enumerated through public CVE and propagated through threat intelligence. Mythos-class capability disrupts this: an adversary may possess working zero-day exploits the CVE database has never seen. No supervisory expectation presently exists that regulated entities perform their own offensive-style vulnerability research using comparable AI-assisted capability.

Gap 3: Software Bill of Materials. CERT-In published Technical Guidelines on SBOM, QBOM, CBOM, HBOM, and AIBOM in July 2025 as guidance, not mandatory direction. The inability of a regulated entity to enumerate, in machine-readable form, every transitive dependency, is a structural defensive deficiency.

Gap 4: Authentication architecture. The Indian banking system standardised on SMS-OTP additional-factor authentication, supplemented by app-based confirmation. Both have been the subject of documented circumvention even pre-Mythos. The Mythos 2FA-bypass disclosures reinforce that migration to phishing-resistant, hardware-bound authentication has been too slow.

Gap 5: Supply-chain liability. The 2022 Outsourcing Master Direction places obligations on the regulated entity but does not directly impose enforceable cybersecurity obligations on the technology service provider. In an environment where a single vendor incident can disrupt 300 banks, the asymmetry is no longer tenable.

Gap 6: Cooperative and regional bank coverage. The 2023 IT Governance Master Direction applies to a defined set of REs; many UCBs and RRBs fall outside its scope. The 15 Elemental Controls for MSMEs (September 2025) is materially lighter than SCB obligations. The regulatory periphery is both highly exposed (per Section 2) and lightly regulated. This is the most consequential gap.

Gap 7: AI-specific obligations. Existing instruments do not specifically address AI, either as defensive tool that REs should be equipped with or as threat against which they must defend. CERT-In AIBOM guidance is a step; an RBI direction is not yet in place.

The institutional coordination gap

Indian cybersecurity responsibilities distribute across RBI (prudential), SEBI (securities), IRDAI (insurers), CERT-In (incident response), CSIRT-Fin (sectoral BFSI), NCIIPC (CII), MeitY (IT Act), and the National Cyber Security Coordinator (cross-sectoral). In ordinary operation, this distribution is workable. In a Mythos-augmented systemic incident, the coordination challenge is materially greater. Pillar 7 addresses this through a standing inter-agency coordination mechanism with pre-defined decision rights and pre-positioned communications.

5. The seven-pillar framework

This section sets out the recommended framework. Each pillar specifies objective, principal interventions, lead institutional actor, statutory basis, indicative cost, and timeline. Interventions are sequenced into three windows: 90 days, 12 months, 24 months.

Pillar 1: Compress the patching cycle

Objective. Reduce defensive operational tempo to match the offensive tempo of a Mythos-augmented adversary.

Recommendations.

- RBI direction supplementing the 2023 IT Governance Master Direction. Mean time-to-patch for Critical/High severity vulnerabilities on internet-facing systems: 48 hours for NPCI, RBI-operated infrastructure, and CII-designated entities; 72 hours for SCBs, Payments Banks, Small Finance Banks; 96 hours for cooperative banks and their shared technology service providers (phased to 72 hours over 24 months).
- Mandatory automated vulnerability management: continuous attack surface management, automated patch deployment for designated software classes, tested rollback. Capability specification by IDRBT in coordination with CERT-In.
- Non-compliance reportable to RBI under existing supervisory framework. Monetary penalty for systemic non-compliance under Section 35A, Banking Regulation Act 1949.

Lead. RBI, with technical input from CERT-In and IDRBT. **Statutory basis.** RBI direction-making power under the Banking Regulation Act, 1949, and the Payment and Settlement Systems Act, 2007. **Cost.** INR 300 to 500 crore over 3 years (regulated entity capex). **Timeline.** Direction in 90 days; SCB/PSO effective at 6 months; cooperative bank effective at 12 months.

Pillar 2: Sovereign defensive AI capability

Objective. Establish, on Indian soil and under Indian institutional control, a defensive AI capability commensurate with offensive capability now in adversary hands. Combine near-term Project Glasswing access with medium-term indigenous capability.

Recommendations.

- **National AI-Cyber Defence Capability for BFSI (NAICD-BFSI).** Constitute under joint aegis of CERT-In, CSIRT-Fin, IDRBT. Functions: continuous AI-assisted vulnerability research across the Indian banking software stack; coordinated disclosure on timelines protecting Indian interests; defensive red-teaming for D-SIBs, NPCI, and CII-designated entities on priority basis; independent technical evaluation of frontier AI cybersecurity claims (counterpart to UK AISI).

- **Phase 1 (near-term): Project Glasswing access.** Engage Anthropic PBC through MeitY and MEA channels to include designated Indian institutions (NPCI, IDRBT, NAICD-BFSI, and 4 to 6 major banks) in Project Glasswing on terms equivalent to peer institutions in Japan, US, EU.
- **Phase 2 (medium-term): indigenous capability.** Develop India-domiciled frontier AI for defensive cybersecurity, drawing on the IndiaAI Mission compute (INR 10,372 crore allocated through 2027), domestic AI laboratories, and academic partnerships. Funding earmark from the IndiaAI Mission specifically for cybersecurity application.
- **Explicit legal authorisation** for NAICD-BFSI to perform vulnerability research on Indian banking infrastructure with regulated entity consent, framed under the IT Act and CERT-In supervisory authority.
- **Funding model.** Capital expenditure: GoI through MeitY budget. Operating expenditure: RBI through the Depositor Education and Awareness Fund. Capability contribution: regulated entities on calibrated basis tied to size and risk profile.

Lead. Department of Financial Services, in coordination with MeitY and RBI. Operational lead with CERT-In. **Statutory basis.** IT Act, 2000 framework for CERT-In; RBI direction for RE participation. Possible enabling amendment for legal authorisation. **Cost.** INR 1,400 to 1,900 crore over 3 years (Phase 1 access fees, NAICD-BFSI infrastructure and personnel, Phase 2 indigenous development seed funding). **Timeline.** Constitutional decisions and inter-ministerial agreement in 90 days; Phase 1 access in 6 months; NAICD-BFSI operational at 12 months; full Phase 2 capability at 24 months and ongoing.

Pillar 3: Critical Technology Service Provider regime

Objective. Extend supervisory and operational control to the third-party perimeter, recognising that regulated entity exposure is determined by the security posture of its weakest service provider.

Recommendations.

- Direct supervisory designation of Critical Technology Service Providers (CTSPs) to BFSI. Designation thresholds: number of regulated entities served (proposed cut-off: 25 SCBs or 100 cooperative/regional banks); criticality of services (core banking, payment switching, KYC, settlement); transaction volume mediated (proposed cut-off: 1 per cent of national daily UPI/RTGS volume).
- CTSP supervisory obligations: mandatory cybersecurity audit on SCB schedule; continuously-maintained SBOM submission; MTTP consistent with Pillar 1; direct incident reporting to CERT-In; capital or insurance reserve scaled to incident-loss scenarios.
- Statutory amendment, where necessary, for direct RBI supervisory authority over CTSPs. The European DORA framework provides the comparative model.
- Mandatory source-code availability or escrow for all critical applications supplied to BFSI REs. Extends the existing PSO Master Direction obligation to SCBs.
- Defined liability regime: technology service providers liable for cybersecurity incidents originating in their systems, with provision for compensation to REs and consumers.

Lead. RBI, with statutory amendments coordinated through DFS. **Statutory basis.** Amendment to RBI Act, 1934 or Banking Regulation Act, 1949 for direct CTSP supervisory authority. **Cost.** INR 600 to 900 crore over 3 years (primarily CTSP compliance, recovered in part through service fees to REs). **Timeline.** Designation criteria and initial list in 12 months; supervisory framework operational at 18 months; statutory amendments where required at 24 months.

Pillar 4: Authentication beyond static credentials

Objective. Migrate Indian banking authentication and identity architecture beyond what Mythos-class authentication-bypass and 2FA-circumvention capabilities can defeat.

Recommendations.

- Hardware-bound authentication (FIDO2 or equivalent) for all privileged-user access at SCBs, NPCI, and CTSPs by 12 months. Approximately 80,000 privileged users in scope nationally; hardware cost INR 4,000 to 8,000 per user; total INR 320 to 640 crore.
- Phishing-resistant out-of-band confirmation for high-value retail transactions (initially threshold INR 5 lakh, progressively reduced) by 18 months. Hardware-binding of customer device required.
- Strengthened device-binding architecture for UPI in coordination with NPCI by 24 months. Behavioural biometrics and continuous authentication as supplementary factors in mobile banking applications.
- Customer protection regime amendment under the Limited Liability of Customers in Unauthorised Electronic Banking Transactions framework. Mythos-class authentication-bypass incidents to default to bank/CTSP liability with customer indemnification. Indemnification mechanism through Pillar 7 insurance facility.

Lead. RBI, with technical coordination through NPCI and IDRBT. **Statutory basis.** RBI direction under the Payment and Settlement Systems Act, 2007. **Cost.** INR 800 to 1,100 crore over 3 years (RE capex on hardware tokens, mobile application redesign, back-end). **Timeline.** Direction in 90 days; phased implementation through 24 months.

Pillar 5: National Threat Intelligence and SBOM Platform

Objective. Establish a real-time threat intelligence sharing platform across the Indian BFSI sector, with a central SBOM repository, modelled on the US FS-ISAC and the European DORA threat intelligence framework. The objective is to convert individual RE threat data into system-level intelligence in minutes rather than days.

Recommendations.

- Mandatory real-time threat intelligence sharing for all SCBs, Payments Banks, NPCI, and CTSPs. Operated by CSIRT-Fin under CERT-In. Standardised STIX/TAXII formats. Anonymisation protocols to address commercial sensitivity.

- Central SBOM repository at CERT-In. Continuously-maintained machine-readable inventory of every software component, including transitive dependencies, deployed at each RE and CTSP. Queryable by NAICD-BFSI and authorised CSIRT-Fin staff for rapid vulnerability impact assessment.
- Sector-wide vulnerability impact tracking: when a vulnerability is disclosed (public CVE) or surfaced (by NAICD-BFSI), the central SBOM enables identification of affected REs within hours rather than weeks. Alert protocols mandate notification of affected REs within 2 hours of vulnerability characterisation.
- Cross-border intelligence sharing arrangements with peer national capabilities (Japan FSA, US Treasury OCCIP, UK FPC, EU EBA) under bilateral memoranda of understanding. Reciprocal information sharing on AI-discovered vulnerabilities affecting cross-border banking infrastructure.

Lead. CERT-In, in coordination with CSIRT-Fin. RBI supervisory enforcement of RE participation. **Statutory basis.** CERT-In Directions of 28 April 2022 under the IT Act. RBI direction under existing supervisory framework. **Cost.** INR 400 to 600 crore over 3 years (platform development, operations, RE integration costs). **Timeline.** Platform design in 6 months; initial operation with SCBs at 12 months; full RE and CTSP coverage at 24 months.

Pillar 6: Cyber-defence workforce at scale

Objective. Build the Indian cybersecurity workforce that the framework requires, recognising that all of Pillars 1 through 5 are constrained by personnel availability at qualified level.

India presently has approximately 500,000 cybersecurity professionals against an estimated requirement of 3 to 3.5 million by 2030 (industry assessments, 2025). The BFSI sector requires approximately 100,000 specialised cybersecurity professionals on present scale, against estimated availability of 25,000 to 35,000. The gap is the principal practical constraint on implementing the framework.

Recommendations.

- **National Cyber-Defence Curriculum.** Developed jointly by AICTE, IDRBT, and CERT-In. Mandatory module in undergraduate computer science, electronics, and information technology programmes from academic year 2027-28. Specialised diploma at ITIs for cyber-defence operations roles.
- **Indian BFSI Cyber Professional (IBCP) certification ladder.** Three tiers: IBCP-Operations (entry-level SOC, monitoring); IBCP-Specialist (vulnerability research, incident response); IBCP-Architect (security architecture, governance). Administered by IDRBT. Mandatory certification for designated roles at SCBs, NPCI, and CTSPs.
- **Bonded scholarship programme.** Government-funded scholarships for cybersecurity master's and doctoral programmes, with 5-year bond to BFSI public sector institutions (RBI, IDRBT, NPCI, CERT-In, NCIIPC, public sector banks). Target: 5,000 bonded scholars annually from 2027.
- **Employer-funded sponsorship.** SCBs and CTSPs to fund 10,000 additional certification slots annually for existing IT personnel transitioning into cyber roles.

- **Cooperative bank shared SOC capability.** A federated security operations centre operated by CSIRT-Fin and IDRBT, serving cooperative and regional banks that cannot economically operate their own. Subsidised through a small sectoral levy on SCB transaction volumes. This is the principal mechanism through which the cooperative bank periphery achieves baseline cyber defence.
- **NAICD-BFSI as workforce magnet.** Competitive compensation and explicit career-progression pathway, in collaboration with IISc, IITs, and IISERs, to attract senior research talent.

Lead. MeitY (curriculum and certification framework); DFS and RBI (BFSI-specific implementation); IDRBT (operational implementation of certification ladder). **Statutory basis.** AICTE Act, 1987 for curriculum changes. RBI direction under existing supervisory framework for mandatory certification of designated roles. **Cost.** INR 700 to 1,000 crore over 3 years (curriculum development, IDRBT scaling, scholarship programme, shared SOC capability). **Timeline.** Curriculum design at 12 months; first IBCP certifications at 18 months; cooperative bank shared SOC at 24 months; scholarship programme operational at 18 months; target workforce of 50,000 BFSI-trained cybersecurity professionals by 2030.

Pillar 7: Resilience, recovery, consumer protection

Objective. Pre-position resilience and recovery capability sufficient to absorb a successful Mythos-class incident, on the assumption that prevention has, in that incident, failed. Bound the financial-sector and macroeconomic consequences.

Recommendations.

- **Mandatory Recovery Time Objectives** for critical payment services: UPI 4 hours full restoration; ATM/card networks 4 hours; RTGS/NEFT 2 hours; internet/mobile banking 24 hours full, 6 hours read-only. Compliance demonstrated through independent audit at all SCBs, NPCI, and CII entities.
- **Quarterly cyber-resilience exercises** at all SCBs, NPCI, and CTSPs. At least one exercise per year structured around an AI-augmented attack scenario. Observed and assessed by CSIRT-Fin. Findings reported to Boards and to RBI.
- **Pre-positioned offline backup and recovery.** Designed to be available in event of comprehensive compromise of online systems. Specifications by IDRBT; implementation overseen by RBI supervision.
- **Banking Sector Cyber-Insurance Facility (BSCIF).** Capitalised through industry contribution (proportional to RE risk profile and transaction volume) plus GoI catalytic funding of INR 500 to 800 crore initial capitalisation. Drawing on international precedent (the US Federal Deposit Insurance Corporation cyber-incident provisions, the UK Industry Insurance Pool). Rapid liquidity in event of large-scale cyber incident affecting multiple REs.
- **Inter-agency coordination mechanism.** Standing body comprising senior representation from DFS, RBI, MeitY, CERT-In, NCIIPC, National Cyber Security Coordinator, NPCI, IDRBT. Convenes on declaration of Critical financial-sector cyber-incident. Pre-defined decision rights, pre-positioned communications channels. Located within the National Cyber Security Coordinator's office for ordinary coordination, escalating to Cabinet Secretariat for systemic incident.

- **Customer protection mechanisms.** Mandatory and accelerated reimbursement timelines for losses arising from incidents in which the bank or a regulated CTSP is implicated. Public communication that the burden of cyber risk does not fall on the customer. Liability allocation under the Limited Liability of Customers framework recalibrated for Mythos-class incidents.
- **International coordination.** India to engage in the Quad cybersecurity working group on AI-cyber threats, bilateral arrangements with Japan, US, UK, EU on threat intelligence and capability sharing, and active participation in the OECD AI Policy Observatory financial-sector workstream.

Lead. RBI for prudential elements; DFS for the BSCIF and inter-agency mechanism; MeitY and MEA for international coordination. **Statutory basis.** RBI direction under existing supervisory framework. New enabling legislation for BSCIF (proposed: Banking (Cyber Resilience) Act, 2027). **Cost.** INR 500 to 800 crore over 3 years (RE capex, BSCIF capitalisation sized separately at INR 500 to 800 crore additional). **Timeline.** Direction on RTOs and exercise cadence in 90 days; BSCIF operational at 18 months; full framework at 24 months.

Pillar interactions and the cumulative effect

The seven pillars are designed to be mutually reinforcing. No single pillar is sufficient. An adversary attempting any of the Section 3 scenarios must defeat several pillars in succession. Pillar 1 (patching) reduces the privilege-escalation surface. Pillar 2 (sovereign AI) reduces the menu of vulnerabilities available to the adversary. Pillar 3 (CTSP regime) addresses supply-chain transmission. Pillar 4 (authentication) defeats credential-based access. Pillar 5 (threat intelligence) accelerates detection across the regulated entity base. Pillar 6 (workforce) is the operational substrate for all of the above. Pillar 7 (resilience) bounds the consequence when prevention has failed.

The framework is not designed to eliminate residual risk. Sufficiently determined adversaries, particularly state actors, may succeed at part of any scenario. The framework reduces the probability of success, reduces the consequence of success that does occur, ensures rapid and effective response, and equalises India's posture with that of peer financial-sector jurisdictions.

6. International responses and adoption opportunity

This section summarises the policy responses of peer financial-sector authorities and identifies the elements suitable for Indian adoption.

Japan: ministerial-level task force, fast

On 21 April 2026, Finance Minister Satsuki Katayama announced a dedicated financial-sector AI cybersecurity task force following a high-level meeting with the Financial Services Agency, Bank of Japan, and chief executives of major Japanese financial institutions. Minister Katayama publicly described the situation as "a crisis that is already at hand," citing specifically the autonomous vulnerability-discovery capability that Mythos has demonstrated.

The task force's announced terms of reference: real-time threat intelligence sharing on AI-discovered vulnerabilities; coordinated patching protocols; review of authentication and identity architectures. The Japanese response is notable for its speed (13 days from Mythos disclosure to ministerial action) and the explicit elevation to ministerial level.

Indian adoption opportunity. The Japanese model of immediate ministerial-level task force, with terms of reference comparable to Section 9 of this paper, is directly transposable. Time-to-establish: 30 days from Ministerial direction.

United States: industry-government engagement

The US response combines direct industry-government engagement, CISA technical advisories, and the existing SEC cybersecurity disclosure rule. Chief executives of major US banks have met with the Federal Reserve and the Department of the Treasury. Of particular note: major US banks are among the partners under Project Glasswing, with restricted access to Mythos for defensive use.

Indian adoption opportunity. The Project Glasswing partnership model is the direct precedent for Pillar 2 of this paper. India should pursue partnership inclusion on terms equivalent to those extended to US, Japanese, and EU institutions.

European Union: DORA as institutional basis

The EU is articulating its Mythos-related response through existing instruments: the Digital Operational Resilience Act (DORA) in force since January 2025, and the AI Act. DORA's ICT third-party risk management framework, including the Critical ICT Third-Party Provider designation regime, provides the legal basis for direct supervisory engagement with technology service providers to the financial sector. The European Banking Authority has signalled Mythos-related guidance under DORA in second-half 2026.

Indian adoption opportunity. DORA is the precise comparative model for Pillar 3 of this paper. The Indian adaptation requires statutory amendment to provide direct RBI authority over CTSPs.

United Kingdom: independent technical evaluation

The UK AI Security Institute (AISI) published in mid-April 2026 an initial evaluation of Mythos providing a partially counterbalancing assessment to Anthropic's own. AISI found Mythos did not outperform human red teams on individual cybersecurity tasks but completed certain multi-step infiltration challenges that no other AI model had previously completed. The Bank of England's Financial Policy Committee is considering implications for UK systemic risk.

Indian adoption opportunity. AISI's independent technical evaluation function is the model for the equivalent function within NAICD-BFSI (Pillar 2). The Indian capability for independent evaluation of frontier AI cybersecurity claims is presently absent.

Synthesis

India is, in international comparative terms, neither lagging nor uniquely well-prepared. The window for response is, however, narrowing. The framework recommended in this paper is calibrated to position India among the more, rather than less, prepared major financial-sector jurisdictions within 12 to 18 months. The principal adoption opportunities from peer responses are summarised in Section 7.

7. Implementation roadmap

This section sets out the implementation pathway with named accountable parties, statutory bases, and quarterly milestones.

The 90-day window: institutional and policy commitments

- **Inter-Ministerial Working Group constituted** by Ministerial direction. Chair: Secretary, Department of Financial Services. Members: MeitY (joint secretary level), RBI (executive director level), National Cyber Security Coordinator, NPCI (CEO), IDRBT (director), MEA (joint secretary level), Cabinet Secretariat (joint secretary level). Mandate: operationalise the framework. Reporting: quarterly to the Hon'ble Finance Minister.
- **RBI Mythos-class threat advisory issued** to all Regulated Entities. Summarises the threat assessment of this paper. Requires immediate self-assessment of patching cadence and authentication architecture. Signals regulatory direction for Pillars 1, 4, and 5.
- **Government-to-government engagement initiated** with Anthropic PBC and counterparts at peer frontier AI laboratories. Channels: MEA, MeitY. Objective: scope Indian inclusion in Project Glasswing or comparable defensive partnerships within 6 months.
- **RBI roundtable convened** of CISOs of D-SIBs, Chief Executive of NPCI, and Chief Executives of principal Critical Technology Service Providers. Communicate supervisory expectation. Establish working groups for each of the seven pillars.
- **This paper published** in suitably summarised form for parliamentary and public information.

The 12-month window: regulatory architecture in place

- RBI direction on patching cadence effective for SCBs and large PSOs (Pillar 1).
- NAICD-BFSI operational launch with initial capability on NPCI infrastructure stack and principal mobile-banking applications (Pillar 2).
- Initial CTSP designation list issued (Pillar 3).
- Hardware-bound authentication for privileged-user access operational at all SCBs and NPCI (Pillar 4).
- National Threat Intelligence and SBOM Platform operational for SCBs (Pillar 5).
- First IBCP certifications issued; National Cyber-Defence Curriculum design complete (Pillar 6).
- First quarterly cyber-resilience exercise conducted on AI-augmented threat scenario (Pillar 7).
- BSCIF capitalisation and operational design complete (Pillar 7).

The 24-month window: full operational maturity

- Patching cadence direction extended to cooperative and regional banks and their technology service providers.
- NAICD-BFSI operating at full scale with continuous coverage of the principal Indian banking technology stack. Recognised counterpart status with UK AISI, Japanese task force, EU EBA.
- Phase 2 of indigenous defensive AI capability seeded, drawing on the IndiaAI Mission compute.
- CTSP supervisory framework, including any required statutory amendments, fully in place.
- Phishing-resistant authentication architecture extended to high-value retail and UPI transactions.
- National Threat Intelligence and SBOM Platform with full RE and CTSP coverage.
- Cooperative bank shared SOC operational; first cohort of 5,000 bonded scholars in placement.
- BSCIF in full operation, with experience-based recalibration.
- Recovery Time Objective compliance demonstrated on independent audit at all SCBs, NPCI, and CII entities.
- Indian participation in Quad cybersecurity working group active; bilateral arrangements with Japan, US, UK, EU in place.

Institutional allocation

Institution	Pillar lead	Pillar support
Department of Financial Services	Overall coordination; Pillar 7 BSCIF	Pillars 2, 3, 6
Reserve Bank of India	Pillars 1, 3, 4, 7 (prudential)	Pillars 5, 6
Ministry of Electronics and Information Technology	Pillar 2 (sovereign AI); Pillar 6 (curriculum)	Pillars 1, 5
CERT-In and CSIRT-Fin	Pillar 2 (operational); Pillar 5	Pillars 1, 7
NCIIPC	CII designation and protection	Pillars 2, 5
NPCI	Pillar 4 (UPI); Pillar 7 (UPI resilience)	Pillars 1, 3, 5
IDRBT	Pillar 6 (certification); technical specifications across	Pillars 2, 5, 7
MEA	Pillar 2 (Glasswing engagement); Pillar 7 (international coordination)	Pillar 5
Cabinet Secretariat	Pillar 7 (escalation for systemic incident)	Coordination across all pillars

Regulated entities	Implementation of supervisory direction	Participation in exercises, training, intelligence sharing

8. Costs, return, and the calculus of inaction

This section presents the indicative financial picture: framework cost, comparative benchmarks, and the cost of inaction.

Framework cost, three-year horizon

Pillar	Cost (INR crore)	Cost basis
Pillar 1: Patching cycle	300 to 500	RE capex on automated vulnerability management.
Pillar 2: Sovereign defensive AI	1,400 to 1,900	NAICD-BFSI infrastructure and personnel; Project Glasswing access fees; Phase 2 indigenous capability seed funding.
Pillar 3: CTSP regime	600 to 900	CTSP compliance costs, partially recovered through RE service fees.
Pillar 4: Authentication	800 to 1,100	RE capex on hardware tokens, mobile redesign, back-end systems.
Pillar 5: Threat intelligence/SBOM	400 to 600	Platform development, operations, RE integration.
Pillar 6: Workforce	700 to 1,000	Curriculum development, IDRBT scaling, scholarship programme, shared SOC for cooperative banks.
Pillar 7: Resilience and recovery	500 to 800	RE capex on RTO compliance and resilience exercises. BSCIF capitalisation sized separately at additional INR 500 to 800 crore.
Aggregate framework	**4,200 to 5,800**	Three-year incremental commitment shared across Government, RBI, REs.
BSCIF initial capitalisation (additional)	500 to 800	One-time, partially industry-funded.
Total programme	**4,700 to 6,600**	

Comparative benchmarks

The framework cost is set against three reference points.

Existing fiscal commitment. Union Budget 2025-26 allocated INR 782 crore to cybersecurity. The proposed framework increment is approximately 6 to 8 times this, spread over 3 years. In annualised

terms, framework cost of INR 1,400 to 1,900 crore per year is approximately 0.005 per cent of the central government expenditure (INR 48 lakh crore in 2025-26).

Cost of present cyber-fraud. I4C reports approximate annual cyber-fraud loss to Indians of INR 1.2 lakh crore in 2025. The framework cost is 3.5 to 4.8 per cent of one year's cyber-fraud loss. A reduction in this baseline of even 5 per cent recovers the full framework cost in one year.

International benchmark. Japan's announced financial-sector AI cybersecurity task force, with comparable scope, has a budget envelope of approximately JPY 250 to 350 billion (INR 16,000 to 22,000 crore) over a comparable horizon. EU member states are committing aggregate DORA-related expenditure several times this. The Indian commitment is comparatively modest.

Cost of inaction

The cost of not implementing the framework is the expected consequence of the threat scenarios, weighted by probability of materialisation. On the central-case scenarios:

- Scenario A (supply-chain compromise of CBS provider): consequence INR 2,800 to 5,500 crore. Probability of materialisation in 24 months: assessed high (based on the C-Edge precedent and the demonstrated Mythos capability). Expected value: INR 2,000 to 4,000 crore.
- Scenario B (mobile banking authentication bypass): consequence INR 1,500 to 3,000 crore. Probability: assessed medium. Expected value: INR 600 to 1,500 crore.
- Scenario C (state-aligned APT on RTGS): consequence not quantifiable but plainly in the macroeconomic category. Probability: assessed low to medium. Expected value not estimated; defended through Pillars 5 and 7.
- Scenario D (KYC exfiltration through privilege-escalation chain): consequence INR 4,000 to 12,000 crore. Probability: assessed medium. Expected value: INR 1,500 to 5,000 crore.

Aggregate expected consequence over 24 months, on central-case scenarios excluding Scenario C: INR 4,100 to 10,500 crore. The framework cost (INR 4,700 to 6,600 crore over 3 years) is comparable to a single scenario expected value. The case for implementation is, on these numbers, asymmetric.

Implementation risks

Capacity. The Indian cybersecurity workforce is the principal binding constraint. Pillar 6 is the direct response.

Industry resistance. Pillar 3 (CTSP designation) will be opposed by some providers. Phased implementation and prior consultation are essential.

Diplomatic complexity. Pillar 2 Glasswing access is not assured. The mitigation is multi-laboratory engagement (Anthropic, OpenAI, Google DeepMind, Mistral, domestic AI laboratories where reaching

capability) and reciprocal offers, including Indian participation in international AI governance initiatives.

False security from implementation. Implementation does not eliminate residual risk. Continuous threat assessment, transparent public reporting, and explicit recognition that the framework reduces but does not eliminate risk are necessary corollaries.

Customer friction. Pillar 4 authentication strengthening will, in some implementations, increase transaction friction. Mitigation through risk-tiered design (heavy controls on high-value, light controls on low-value), investment in user-experience design, and explicit consumer communication.

9. Conclusion and three immediate recommendations

On 8 April 2026, the public release of Claude Mythos Preview and the simultaneous launch of Project Glasswing made operative a category of risk to the Indian banking system that, until that date, was emerging concern rather than present one. The capability of frontier AI to autonomously discover, at scale, previously-unknown vulnerabilities in software (including the precise classes of vulnerability most relevant to banking platforms) has been demonstrated. The proliferation of that capability beyond Anthropic's controlled access programme is, on the best available analysis, a matter of months rather than years.

The Indian banking system has, in 2026, a regulatory and operational architecture for cybersecurity that is, in international comparative terms, mature. It is not calibrated to the threat environment that Mythos has revealed. The window for adaptation is narrow. Peer financial-sector jurisdictions, Japan most rapidly, the United States, the European Union, and the United Kingdom, have begun the work of adaptation. India must do the same, and on a timeline that does not leave the Indian banking system as a relatively softer target than its international counterparts.

The framework recommended in this paper is structured around seven pillars: compression of the patching cycle; sovereign defensive AI capability; Critical Technology Service Provider regime; authentication beyond static credentials; national threat intelligence and SBOM platform; cyber-defence workforce at scale; and resilience, recovery, consumer protection. The framework is implementable within a 24-month horizon, at an indicative incremental commitment of INR 4,200 to 5,800 crore (plus INR 500 to 800 crore one-time BSCIF capitalisation), shared across Government, RBI, and regulated entities. Against the present annual cyber-fraud loss to Indians and against the potential cost of a successful systemic incident, this commitment is modest. Against the cost of comparable programmes in peer jurisdictions, it is competitive.

THREE ACTIONS FOR IMMEDIATE CONSIDERATION

> **Constitute, by Ministerial direction within 30 days, an Inter-Ministerial Working Group on Frontier AI Risk to the Indian Banking System**, chaired by the Secretary, Department of Financial Services, with the institutional composition and the 90-day mandate set out in Section 7. The Working Group to report quarterly to the Hon'ble Finance Minister.

> **Communicate, through such channels as the Hon'ble Finance Minister considers appropriate, the strategic significance of the matter to the Reserve Bank of India and to the Boards of the major banks**, in parallel with the Working Group constitution, to ensure supervisory and operational responses move in concert from the outset.

> **Initiate, through MeitY and MEA, engagement with Anthropic PBC and other frontier AI laboratories**, to secure for designated Indian institutions (NPCI, IDRBT, the proposed NAICD-BFSI, and 4 to 6 major banks) access to Mythos-class capability under defensive controls equivalent to those extended to peer institutions in Japan, the United States, and the European Union.

The defining characteristic of the present moment is that defenders no longer have the luxury of operating at a slower tempo than attackers. The technological reasons for that shift are set out in Section 2. The institutional consequences are addressed in the framework. The recommendations are an argument for the proposition that the Indian banking system should organise itself, on a timeline measured in months, to operate at a tempo commensurate with the threat. The cost of doing so is calculable; the cost of not doing so is not.

10. References and author note

Anthropic publications

- Anthropic PBC, Claude Mythos Preview and Project Glasswing, anthropic.com/glasswing, 8 April 2026.
- Anthropic Frontier Red Team, Claude Mythos Preview Technical Findings, April 2026.
- Anthropic, Claude Mythos Preview System Card and Alignment Risk Update, April 2026.

Independent evaluations

- UK AI Security Institute, Initial Evaluation of Claude Mythos Preview, April 2026.
- Cloud Security Alliance, statement on AI-accelerated vulnerability discovery, April 2026.

International policy

- Government of Japan, Ministry of Finance, statement of Finance Minister Satsuki Katayama on the establishment of the financial-sector AI cybersecurity task force, 21 April 2026.
- European Banking Authority signalling on DORA Mythos-related guidance, second half 2026.
- US Federal Reserve and Department of the Treasury industry consultations, April-May 2026.
- Bank of England Financial Policy Committee record on systemic AI-cyber risk, May 2026.

Indian regulatory and institutional sources

- Reserve Bank of India, Cyber Security Framework in Banks, RBI/2015-16/418, 2 June 2016.
- Reserve Bank of India, Master Direction on Digital Payment Security Controls, RBI/2020-21/74, 18 February 2021.
- Reserve Bank of India, Master Direction on Outsourcing of Information Technology Services, June 2022.
- Reserve Bank of India, Master Direction on Information Technology Governance, Risk, Controls and Assurance Practices, 7 November 2023, effective 1 April 2024.
- Reserve Bank of India, Master Directions on Cyber Resilience and Digital Payment Security Controls for Non-Bank Payment System Operators, July 2024.
- CERT-In, Digital Threat Report 2024 for the BFSI Sector, April 2025.
- CERT-In, Technical Guidelines on SBOM, QBOM, CBOM, HBOM, and AIBOM, Version 2, July 2025.
- CERT-In, Comprehensive Cybersecurity Audit Policy Guidelines, July 2025.

- CERT-In, White Paper on Transitioning to Quantum Cyber Readiness, July 2025.
- CERT-In, 15 Elemental Cyber Defence Controls for MSMEs, September 2025.
- Indian Cyber Crime Coordination Centre (I4C) reporting on cybercrime complaints and financial fraud losses, 2025.
- NITI Aayog, IndiaAI Mission documentation and budget allocation, 2024-25.

Banking incidents and threat reporting

- CloudSEK Threat Research, Major Payment Disruption: Ransomware Strikes Indian Banking Infrastructure (C-Edge incident), August 2024.
- CSO Online, Over 300 Indian banks suffer payment disruption from ransomware attack, July 2024.
- CYFIRMA Threat Intelligence Reports, January to April 2026.
- Seqrite India Cyber Threat Report 2026.
- (ISC)2 Cybersecurity Workforce Study, India chapter, 2024 and 2025.

Statutory framework

- Information Technology Act, 2000, and rules thereunder.
- Digital Personal Data Protection Act, 2023.
- Banking Regulation Act, 1949 (in particular Section 35A on RBI direction-making power).
- Reserve Bank of India Act, 1934.
- Payment and Settlement Systems Act, 2007.
- AICTE Act, 1987 (for curriculum framework changes in Pillar 6).
- Union Budget 2025-26, Government of India.
- NCIIPC designation orders under Section 70 of the IT Act, 2000.

About the author

Harsh Macwan is the Chief Executive Officer of HyperNext Data Center Limited, an Indian AI-native data centre operator headquartered in Hyderabad. He has prepared this paper in his personal capacity. The views expressed are his own and do not represent the institutional position of HyperNext Data Center Limited or any other organisation with which he is associated. Correspondence on the methods, the figures, and the conclusions is welcome at hello@hypernxt.com.

HyperNext
Data Centers

HyperNext Research

We publish engineering and policy papers because the Indian conversation about AI infrastructure needs more substance than marketing material provides. The papers state methodology openly so other operators can run the same analysis on their own facilities. They report findings that may not flatter the HyperNext commercial position. They get review from the engineering team and the communications partners.

Correspondence on methods, figures, and conclusions: **hello@hypernxt.com**. We read every email.

HN-RP-009 · Frontier AI and the Indian Banking System
04 June 2026 · v1.0

www.hypernxt.com/research
hello@hypernxt.com · +91 99784 23333