INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-28 07:05 UTC

AI-Driven Scams on the Rise; Indian Government Rolls Out New Cyber Safety Guidelines

GOVERNANCE | MEDIUM

SCC Item ID	SCC-GOV-2026-0086
Type	Governance
Severity	MEDIUM
Affected Products	General public, Indian government ministries, digital platform users
Published	2026-06-27
Discovery Source	Gemini

Executive Summary

Industry reporting and threat intelligence indicate accelerating AI-driven fraud targeting Indian government entities and the general public, including deepfake impersonation, voice cloning, and AI-generated phishing campaigns. Organizations operating in or partnering with Indian government entities face elevated social engineering risk as threat actors use AI to scale and personalize attacks that bypass traditional awareness training. The primary business risk is financial fraud, credential compromise, and reputational damage from successful impersonation of executives or trusted institutions.

Technical Analysis

This advisory describes a governance-level response to an escalating social engineering threat landscape, not a discrete technical vulnerability. No CVE is associated. Relevant CWEs are CWE-1021 (Improper Restriction of Rendered UI Layers) and CWE-693 (Protection Mechanism Failure), reflecting the breakdown of user-facing trust mechanisms under AI-augmented deception. MITRE ATT&CK techniques in scope: T1566 (Phishing), T1566.001 (Spearphishing Attachment), T1566.002 (Spearphishing Link), T1534 (Internal Spearphishing), T1585 (Establish Accounts), and T1588 (Obtain Capabilities). The threat model centers on adversaries using generative AI to produce high-volume, linguistically convincing phishing lures; voice cloning to impersonate executives or government officials; and deepfake video to authenticate fraudulent requests. No patch exists. Mitigation is procedural, detection-oriented, and governance-driven. Source confidence is medium; the advisory originated from secondary discovery with no confirmed primary authoritative URL from an official Indian government cybersecurity authority.

Action Checklist

1. Step 1: Containment, Immediately brief communications and finance teams on AI-generated voice and video fraud. Establish a verbal code word or secondary verification channel for any out-of-band wire transfer, credential reset, or executive instruction received by phone or video call. This directly addresses T1566 and T1534 abuse vectors.
2. Step 2: Detection, Review email gateway logs for anomalous sender patterns, lookalike domains, and AI-generated lure characteristics (unusual linguistic fluency, generic salutations, urgency framing). Flag inbound calls requesting credential or financial action for secondary verification. Reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) for log review cadence. Detection is behavioral and procedural; no specific IOC signatures are available for this threat class.
3. Step 3: Eradication, No software patch applies. Enforce strict out-of-band verification procedures for any financial or access-control request. Review and tighten email authentication controls (SPF, DKIM, DMARC) to reduce spoofing surface supporting T1585 and T1566 techniques. Reference CIS 6.3 (Require MFA for Externally-Exposed Applications) to reduce credential takeover risk downstream of successful phishing.
4. Step 4: Recovery, Validate that MFA is enforced on all externally exposed applications and administrative accounts per CIS 6.3, CIS 6.4, and CIS 6.5. Confirm no accounts were compromised during any recent unsolicited credential or verification request. Monitor for anomalous login patterns using local account monitoring consistent with D3-LAM. Apply D3-MFA (Multi-factor Authentication) as a standing recovery safeguard.
5. Step 5: Post-Incident, Conduct a social engineering awareness exercise specific to AI-generated threats, including deepfake video and voice cloning scenarios. Audit whether existing security awareness training addresses AI-augmented lures. Map identified gaps to NIST AC-17 (Remote Access) and NIST AU-2 (Event Logging) to ensure procedural controls are codified and logged. Document findings and update the incident response playbook to include AI-specific social engineering scenarios.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to executive leadership, legal counsel, and (if operating under Indian DPDP Act or partnering with Indian government entities) the relevant ministry CISO if any employee confirms a completed wire transfer, credential submission, or access grant following an AI-generated voice or video request, as these actions constitute potential financial fraud and may trigger regulatory notification obligations.
Recovery Notes	Verify MFA enforcement and out-of-band verification procedures are operational across all finance, HR, and executive assistant roles before declaring recovery complete. Monitor authentication logs for anomalous login patterns — particularly off-hours access, new device enrollments, or MFA bypass attempts — for a minimum of 30 days following the advisory, as AI-driven phishing campaigns described in the Indian government advisory operate as sustained waves rather than single-event attacks. Update the security awareness training curriculum to include AI-generated deepfake and voice cloning scenarios within 30 days and re-test targeted staff with a simulated AI-lure exercise to validate behavioral change.

Forensic Artifacts	Email gateway delivery logs (.eml raw headers) for inbound messages exhibiting AI-generated lure characteristics: near-perfect grammar in non-native languages, urgency framing, generic salutations, and sender domains registered within 90 days of the campaign — these headers preserve DKIM/SPF/DMARC validation state that is destroyed if messages are deleted | VoIP/PBX call detail records (CDRs) for inbound calls to finance and executive assistant extensions during the campaign window, capturing caller ID, call duration, and timestamp — critical for correlating voice-cloned executive impersonation attempts against specific staff targets | DMARC aggregate reports (XML, from the `rua` reporting address) for the 30 days preceding the advisory, documenting which external IPs were sending mail using your organization's domain without valid authentication — direct evidence of spoofing infrastructure supporting AI-generated phishing | Authentication and sign-in logs (Microsoft 365 Entra sign-in logs or Google Workspace Admin security report) for all accounts that received suspicious credential requests, preserving source IP, user agent string, MFA challenge result, and session token issuance — must be captured before any forced password reset or session revocation destroys the token evidence | Any recorded video call or voicemail artifacts still resident on collaboration platforms (Microsoft Teams, Zoom, WhatsApp) that may constitute deepfake or voice-cloned synthetic media — these are the primary forensic evidence of the AI-generation capability used and are subject to platform retention policy deletion within days

Per-Action IR Details

Step 1: Containment — Immediately brief communications and finance teams on AI-generated voice and video fraud. Establish a verbal code word or secondary verification channel for any out-of-band wire transfer, credential reset, or executive instruction received by phone or video call. This directly addresses T1566 and T1534 abuse vectors.

NIST Phase: Containment

Reference: NIST 800-61r3 §3.3 — Containment Strategy

Controls: NIST AC-17 (Remote Access), CIS 6.3 (Require MFA for Externally-Exposed Applications)

Compensating: Distribute a one-page brief via secure internal email (not by phone or video call) to finance, HR, and executive assistants detailing AI voice/video fraud indicators: unexpected urgency, requests to bypass normal approval chains, calls from unknown or spoofed numbers claiming to be senior leadership. Implement a mandatory callback procedure using a pre-registered number from the internal directory — not a number provided by the caller. No tooling required; enforce via policy memo and immediate team huddle.

Evidence: Before briefing teams or enacting any new verification channel changes, capture: (1) records of any recent unsolicited wire transfer requests, credential reset demands, or executive instructions received by phone or video call (pull call logs from PBX/VoIP system if available); (2) email gateway logs for the prior 30 days showing inbound messages referencing finance, wire transfer, credential reset, or executive action from external senders; (3) any voicemail recordings or video call recordings still resident on collaboration platforms (Teams, Zoom, Google Meet) that may constitute synthetic media evidence. These artifacts are volatile if call recording retention is short.

Step 2: Detection — Review email gateway logs for anomalous sender patterns, lookalike domains, and AI-generated lure characteristics (unusual linguistic fluency, generic salutations, urgency framing). Flag inbound calls requesting credential or financial action for secondary verification. Reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) for log review cadence. No specific event ID or IOC pattern is confirmed from the available source material.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis

Controls: NIST AU-6 (Audit Record Review, Analysis, and Reporting), NIST AU-2 (Event Logging), CIS 8.2 (Collect Audit Logs)

Compensating: For teams without a SIEM, use native email gateway export (Microsoft 365: `Get-MessageTrace` via Exchange Online PowerShell; Google Workspace: Admin Console > Reports > Email Log Search) to pull the last 30 days of inbound mail. Filter for: (1) sender domains registered within the past 90 days (check via `whois` or MXToolbox); (2) display name spoofing where the display name matches an internal executive but the SMTP domain does not; (3) subject lines containing terms like 'urgent', 'wire', 'reset', 'verify', 'invoice', or 'CEO'. No SIEM required — a spreadsheet pivot on these fields is sufficient for a 2-person team.

Evidence: Capture the following before any gateway rule changes or sender blocks that would alter mail flow: (1) raw email headers (`.eml` files) from suspicious inbound messages, preserving DKIM signature validity status, SPF alignment result, and DMARC disposition; (2) full email gateway delivery logs showing sender IP, PTR record, and authentication results for the review window; (3) VoIP/PBX call detail records (CDRs) for inbound calls to finance and executive assistant extensions, capturing caller ID, call duration, and timestamp. DKIM and SPF validation data is especially volatile if messages are deleted or headers are stripped by downstream filtering.

Step 3: Eradication — No software patch applies. Enforce strict out-of-band verification procedures for any financial or access-control request. Review and tighten email authentication controls (SPF, DKIM, DMARC) to reduce spoofing surface supporting T1585 and T1566 techniques. Reference CIS 4.4 (Implement and Manage a Firewall on Servers) for perimeter hygiene and CIS 6.3 (Require MFA for Externally-Exposed Applications) to reduce credential takeover risk downstream of successful phishing.

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST SI-2 (Flaw Remediation), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 4.2 (Establish and Maintain a Secure Configuration Process for Network Infrastructure)

Compensating: Use MXToolbox (https://mxtoolbox.com/SuperTool.aspx) or `dig`/`nslookup` to audit your domain's SPF record for overly permissive `+all` or `?all` mechanisms; tighten to `-all`. Verify DKIM is publishing a valid selector via `dig TXT ._domainkey.`. Check DMARC policy via `dig TXT _dmarc.` — escalate policy from `p=none` to `p=quarantine` or `p=reject` if reporting data confirms legitimate mail flows are covered. These are zero-cost DNS changes executable by a 2-person team within one change window. Note: CIS 4.4 (firewall on servers) does not directly govern email authentication hardening — the SPF/DKIM/DMARC actions are the operationally relevant eradication steps for this threat; CIS 4.4 cited in the original step is a force-fit and has been replaced with CIS 4.2 above, which governs secure configuration of network infrastructure including mail routing controls.

Evidence: Before enforcing DMARC policy changes or blocking spoofed sender domains, preserve: (1) current SPF, DKIM, and DMARC DNS records (capture with `dig TXT` and timestamp the output — DNS changes are not logged by default and prior state is lost immediately on update); (2) DMARC aggregate reports (`rua` destination) for the past 30 days showing which sending sources are authenticated vs. failing — these confirm blast radius of any spoofing already in progress; (3) any identified lookalike or newly registered domains targeting your organization (document via WHOIS before any takedown request removes registration data).

Step 4: Recovery — Validate that MFA is enforced on all externally exposed applications and administrative accounts per CIS 6.3, CIS 6.4, and CIS 6.5. Confirm no accounts were compromised during any recent unsolicited credential or verification request. Monitor for anomalous login patterns using local account monitoring consistent with D3-LAM. Apply D3-MFA (Multi-factor Authentication) as a standing recovery safeguard.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST AC-2 (Account Management), NIST AC-7 (Unsuccessful Logon Attempts), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.4 (Require MFA for Remote Network Access), CIS 6.5 (Require MFA for Administrative Access), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: For Microsoft 365 environments, run `Get-MsolUser -All | Where-Object { $_.StrongAuthenticationMethods.Count -eq 0 }` (Azure AD PowerShell) to enumerate accounts lacking MFA enrollment. For on-premises Active Directory, use `Get-ADUser -Filter * -Properties LastLogonDate, PasswordLastSet` to identify accounts with recent logon activity against which AI-driven phishing may have been attempted.

Cross-reference any account with a password change in the past 14 days against the list of employees who reported suspicious calls or emails. Free tool: Entra ID (formerly Azure AD) sign-in logs are available at no additional cost and show MFA success/failure, location, and device compliance.

Evidence: Before revoking any suspected compromised sessions or forcing password resets (which destroy session token evidence), capture: (1) active session tokens and sign-in logs for accounts that received suspicious credential requests — in Microsoft 365 pull via `Get-AzureADAuditSignInLogs` or the Entra portal, preserving IP address, user agent, and MFA method fields; (2) for Google Workspace, export the Admin Console > Security > User Account Activity report for the same accounts; (3) any browser-saved credential entries or password manager exports the user may have submitted to a phishing page (interview the user before resetting — user recall of the exact URL submitted to is volatile and lost after reset).

Step 5: Post-Incident — Conduct a social engineering awareness exercise specific to AI-generated threats, including deepfake video and voice cloning scenarios. Audit whether existing security awareness training addresses AI-augmented lures. Map identified gaps to NIST AC-17 (Remote Access) and NIST AU-2 (Event Logging) to ensure procedural controls are codified and logged. Document findings and update the incident response playbook to include AI-specific social engineering scenarios.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity

Controls: NIST AU-2 (Event Logging), NIST AU-6 (Audit Record Review, Analysis, and Reporting), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Use the SANS Security Awareness Maturity Model (free framework) to score current training against AI-specific threat scenarios: deepfake video calls impersonating Indian government ministry contacts, voice-cloned executive fraud calls, and AI-generated phishing emails with native-language fluency targeting Hindi or regional-language speakers. Develop two tabletop scenarios: (1) a finance employee receives a WhatsApp video call from a deepfake CFO requesting an urgent wire transfer; (2) an IT help desk receives a voice-cloned call from a 'ministry official' requesting a VPN credential reset. Document gaps and map to updated IR playbook. No tooling budget required — use Google Meet or Zoom for tabletop delivery. Note: NIST AC-17 (Remote Access) cited in the original step does not directly govern awareness training or logging — AU-2 and AU-6 from the knowledge base are the operationally correct post-incident control mappings retained above.

Evidence: For post-incident documentation, preserve: (1) all reported suspicious call and email records from the incident window as training case studies — anonymize before use; (2) current security awareness training curriculum documentation showing last update date, to establish a baseline gap against AI-generated threat scenarios; (3) the updated IR playbook version with AI social engineering scenarios added, signed and dated, to satisfy audit evidence requirements under NIST AU-10 (Non-Repudiation) for policy change accountability.

Detection Guidance

No confirmed IOCs are available from the source material. Behavioral detection should focus on: inbound emails with high linguistic fluency but mismatched sender domains (review email gateway and DMARC reports); unsolicited calls or video requests from individuals claiming executive or government authority, particularly those requesting wire transfers or credential resets; anomalous account creation or access attempts following communication spikes (reference NIST AU-6 for review cadence and AU-3 for audit record content requirements). Apply D3-LAM (Local Account Monitoring) to detect unauthorized account activity following a suspected social engineering event. Apply D3-ALA (Authentication Log Analysis) to monitor for post-phishing access patterns. No specific log query or event ID can be confirmed from available source material at medium confidence.

Framework Mappings

MITRE-ATTACK

- **T1566.002** — Spearphishing Link
- **T1566** — Phishing
- **T1585** — Establish Accounts
- **T1534** — Internal Spearphishing
- **T1566.001** — Spearphishing Attachment
- **T1588** — Obtain Capabilities

NIST-800-53R5

- **AT-2** — Literacy Training and Awareness
- **SC-7** — Boundary Protection
- **SI-3** — Malicious Code Protection
- **SI-4** — System Monitoring
- **SI-8** — Spam Protection
- **CA-7** — Continuous Monitoring

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **A.5.34** — Privacy and protection of personal information

CIS-V8

- **14.2** — Train Workforce Members to Recognize Social Engineering Attacks

HIPAA-SECURITY

- **164.308(a)(5)(i)** — Security Awareness and Training

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1566.002	Spearphishing Link	Initial-Access
T1566	Phishing	Initial-Access
T1585	Establish Accounts	Resource-Development
T1534	Internal Spearphishing	Lateral-Movement
T1566.001	Spearphishing Attachment	Initial-Access
T1588	Obtain Capabilities	Resource-Development

Sources

Source	URL	Tier
Assessing the information security posture of online public services ...	https://www.sciencedirect.com/science/article/pii/S0740624X25000255	T3
Cyber security in government for secure public services	https://www.oneadvanced.com/resources/cyber-security-in-government-...	T3
How to tackle the most common cyber threats government faces	https://www.globalgovernmentforum.com/events/how-to-tackle-the-most...	T3
Digital Government: Building a 21st Century Platform to Better Serve ...	https://obamawhitehouse.archives.gov/sites/default/files/omb/egov/d...	T1
Significant Cyber Incidents | Strategic Technologies Program - CSIS	https://www.csis.org/programs/strategic-technologies-program/signif...	T3

Generated 2026-06-28 07:05 UTC by TJS Security Command Center