INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-05-09 18:46 UTC

Post-Quantum Cryptography Transition Risk for AI Infrastructure

SECURITY ANALYSIS | HIGH

SCC Item ID	SCC-STY-2026-0118
Type	Security Analysis
Severity	HIGH
Affected Products	AI infrastructure systems relying on classical cryptographic algorithms (RSA, ECC) for agent-to-tool and agent-to-agent communication
Published	2026-05-09
Discovery Source	Gemini

Executive Summary

AI infrastructure built on classical public-key cryptography (RSA, ECC) faces a credible, long-horizon threat from cryptographically relevant quantum computers, but the near-term risk is already operational. Adversaries conducting harvest-now-decrypt-later (HNDL) attacks may already be collecting encrypted AI agent traffic today, positioning themselves to decrypt it once quantum capability matures. NIST finalized post-quantum replacement standards in 2024 (FIPS 203/204/205), establishing a clear transition mandate; organizations that delay cryptographic modernization are accumulating silent, compounding exposure with every passing quarter.

Technical Analysis

The threat model here is structural, not episodic. Classical public-key algorithms, RSA and ECC, derive their security from computational problems (integer factorization, discrete logarithm) that a sufficiently capable quantum computer can solve efficiently via Shor's algorithm. NIST announced the final post-quantum standards in August 2024, with official FIPS publication following, producing three approved standards: ML-KEM (FIPS 203) for key encapsulation, ML-DSA (FIPS 204) for digital signatures, and SLH-DSA (FIPS 205) as a stateless hash-based signature alternative. NIST SP 800-131A Rev. 2 recommends transition timelines for RSA and ECC key exchange, removing any ambiguity about the strategic direction.

For AI infrastructure specifically, the attack surface extends beyond traditional encrypted channels. Multi-agent AI architectures introduce agent-to-tool and agent-to-agent communication layers, authentication handshakes, API calls, orchestration traffic, that typically rely on TLS with classical key exchange. Each of these channels represents an HNDL collection target. The mapped MITRE techniques illustrate the progression: T1040 (network sniffing) supports passive collection of encrypted traffic; T1557 (adversary-in-the-middle) enables

active interception of authentication handshakes; T1563 (remote service session hijacking) and T1021 (remote services abuse) describe what becomes possible once authentication material is compromised or decrypted.

The relevant CWEs ground this structurally: CWE-327 (use of broken or risky cryptographic algorithm) applies directly to continued RSA/ECC deployment; CWE-311 (missing encryption of sensitive data) captures scenarios where AI agent traffic is unencrypted or weakly protected at internal network segments; CWE-326 (inadequate encryption strength) applies to undersized key parameters that fall short of current NIST guidance even before quantum is factored in.

Cryptographic agility, the architectural capacity to swap cryptographic primitives without rewriting dependent systems, is the operational requirement that separates organizations positioned to transition from those that are not. AI systems built with hardcoded cipher suites, static TLS configurations, or undocumented cryptographic dependencies will face the highest remediation costs. The 2026 framing is accurate: this is no longer a distant planning exercise. HNDL is an active threat tactic, and the NIST standards are final.

Action Checklist

1. Step 1: Assess cryptographic inventory, conduct a discovery sweep of all AI agent communication channels (agent-to-tool, agent-to-agent, orchestration APIs) to identify which are using RSA or ECC key exchange; include TLS configuration audits across internal AI infrastructure, not just perimeter-facing services.
2. Step 2: Review controls against HNDL exposure, verify network segmentation isolates AI agent traffic from paths accessible to external adversaries; assess whether any AI system traffic traverses untrusted or semi-trusted network segments where passive collection is plausible; check logging coverage for T1040/T1557 detection on those paths.
3. Step 3: Evaluate cryptographic agility posture, determine whether your AI infrastructure components (orchestration frameworks, model inference APIs, inter-agent messaging layers) support cipher suite configuration without code changes; flag hardcoded cryptographic dependencies as technical debt requiring prioritized remediation.
4. Step 4: Map to NIST PQC transition requirements, review NIST SP 800-131A Rev. 2 and assess your organization's gap against ML-KEM, ML-DSA, and SLH-DSA adoption; engage AI platform vendors on their post-quantum cryptography roadmaps and request written commitment dates aligned to NIST guidance.
5. Step 5: Update threat model and brief leadership, add HNDL as an active threat scenario in your threat register, mapped to AI infrastructure data flows; brief the CISO and relevant board members that encrypted AI traffic collected today may be decryptable in the future, and frame the cryptographic transition as a risk mitigation action with a known cost and a foreseeable timeline.

IR / Forensic Enrichment

Triage Priority	STANDARD

Escalation Criteria	Escalate to urgent if network forensics (Zeek ssl.log, tcpdump PCAP analysis) reveals that AI infrastructure agent traffic is routed through or has traversed network segments accessible to external adversaries, or if any AI platform vendor confirms they have no PQC roadmap and no planned support for ML-KEM or ML-DSA, as these conditions materially increase HNDL exposure and may trigger supply chain risk notification obligations under contractual SLAs or sector-specific regulatory frameworks (e.g., FedRAMP, HIPAA if AI systems process PHI, financial sector guidance).
Recovery Notes	Post-transition recovery validation requires confirming that all AI agent communication channels (agent-to-tool, agent-to-agent, orchestration APIs) have successfully negotiated ML-KEM-based key encapsulation in TLS 1.3 sessions — verify via Zeek ssl.log `cipher` field showing FIPS 203-compliant algorithm identifiers or via `testssl.sh` post-migration scan confirming no RSA or classical ECDH cipher suites remain active on AI infrastructure endpoints. Monitor for cryptographic downgrade attempts for a minimum of 90 days post-migration by alerting on any TLS ClientHello or ServerHello negotiating RSA or pre-quantum ECDH cipher suites on AI infrastructure segments, as these may indicate misconfigured legacy agents or active downgrade attacks. Retain pre-migration PCAP baselines and post-migration ssl.log samples as evidence that the transition was completed within the organization's documented risk acceptance window and as a reference artifact for future audit of NIST SP 800-131A Rev. 2 compliance.
Forensic Artifacts	TLS handshake PCAPs from AI infrastructure network segments — specifically ClientHello and ServerHello records documenting cipher suite negotiation between AI agents, tools, and orchestration APIs; these confirm which RSA or ECC-based key exchange algorithms were active and for what duration, establishing the HNDL collection window | Zeek ssl.log files capturing `cipher`, `curve`, `cert_chain_fuids`, and `validation_status` fields for all TLS sessions on AI infrastructure VLANs — these provide a queryable, timestamped record of every classical cryptography negotiation across agent communication paths without requiring a SIEM | Cryptographic configuration files from AI orchestration frameworks and inference API servers — including nginx/HAProxy ssl_ciphers directives, Python ssl.SSLContext initialization parameters, Java SSLSocketFactory configurations, and environment variable dumps from containerized AI workloads — documenting the hardcoded or configurable cryptographic posture at the time of assessment | Network flow logs (NetFlow/IPFIX or firewall connection logs) for AI infrastructure segments covering a 90-day lookback period — used to identify anomalous external-facing connections or unexpected routing of AI agent traffic through untrusted segments that would indicate HNDL-relevant passive collection exposure has already occurred | Vendor-provided cipher suite support matrices and written PQC roadmap commitments from all AI platform providers — these serve as supply chain risk evidence under NIST SA-9 and establish which vendors represent residual RSA/ECC dependency that the organization cannot remediate unilaterally within the NIST SP 800-131A Rev. 2 deprecation timeline

Per-Action IR Details

Step 1: Assess cryptographic inventory — conduct a discovery sweep of all AI agent communication channels (agent-to-tool, agent-to-agent, orchestration APIs) to identify which are using RSA or ECC key exchange; include TLS configuration audits across internal AI infrastructure, not just perimeter-facing services.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: establishing visibility and asset inventory as a prerequisite to detecting and responding to cryptographic exposure in AI infrastructure

Controls: NIST SI-7 (Software, Firmware, and Information Integrity) — verify integrity of cryptographic configurations across AI orchestration layers, NIST CA-9 (Internal System Connections) — document and assess cryptographic posture of internal AI agent-to-agent and agent-to-tool connections, NIST SC-8 (Transmission Confidentiality and Integrity) — evaluate whether transmission protection mechanisms on AI infrastructure channels meet current standards, NIST RA-3 (Risk Assessment) — frame classical crypto dependency on RSA/ECC as a documented risk item with likelihood tied to HNDL threat horizon, CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory) — extend asset inventory scope to include AI agent communication endpoints and their cipher suite configurations, CIS 2.1 (Establish and Maintain a Software Inventory) — capture cryptographic library versions (e.g., OpenSSL, BoringSSL, PyCryptodome) in use across AI inference APIs and orchestration frameworks

Compensating: Use `nmap --script ssl-enum-ciphers -p 443,8443,8080 ` to enumerate TLS cipher suites on all AI API endpoints without a commercial scanner. For Python-based AI stacks (LangChain, AutoGen, CrewAI), grep source or installed packages: `grep -r 'RSA\|ECDH\|ecdsa\|rsa' /opt/ai_stack/ --include='*.py' --include='*.yaml' --include='*.json'`. Use `testssl.sh` (free, shell-based) against each internal AI service endpoint to produce cipher suite reports. Document findings in a shared spreadsheet mapping endpoint → algorithm → key size → library version.

Evidence: Before beginning the sweep, snapshot current TLS handshake data by running a short `tcpdump -i -w ai_tls_baseline_$(date +%Y%m%d).pcap 'port 443 or port 8443 or port 8080'` on the AI infrastructure network segment for 15–30 minutes during normal agent activity. This captures ClientHello/ServerHello records that document which cipher suites are actively negotiated between agents, tools, and orchestration APIs — the exact artifacts that confirm RSA/ECC key exchange is in use and would be present in HNDL-collected traffic.

Step 2: Review controls against HNDL exposure — verify network segmentation isolates AI agent traffic from paths accessible to external adversaries; assess whether any AI system traffic traverses untrusted or semi-trusted network segments where passive collection is plausible; check logging coverage for T1040/T1557 detection on those paths.

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: correlating network monitoring data and logging coverage to assess whether HNDL-relevant passive collection (T1040 Network Sniffing) is plausible on AI agent communication paths

Controls: NIST SI-4 (System Monitoring) — verify monitoring is deployed on internal network segments carrying AI agent traffic, specifically for passive collection indicators aligned with T1040, NIST AU-2 (Event Logging) — confirm logging is enabled and captures network flow data and TLS metadata on AI infrastructure segments, not solely perimeter egress points, NIST AU-12 (Audit Record Generation) — validate that AI orchestration frameworks (e.g., LangChain server, AutoGen runtime) generate audit records for inter-agent API calls, NIST SC-7 (Boundary Protection) — assess whether network boundaries segment AI agent traffic from zones accessible to adversaries capable of passive collection, CIS 8.2 (Collect Audit Logs) — verify that audit log collection is active on network devices and hosts in AI infrastructure segments, covering the paths where HNDL interception is plausible, CIS 4.4 (Implement and Manage a Firewall on Servers) — confirm host-based firewalls on AI inference and orchestration nodes restrict lateral traffic to known agent communication paths only

Compensating: Deploy Zeek (formerly Bro) on a network tap or SPAN port covering the AI infrastructure VLAN to generate `ssl.log` and `conn.log` files — these log cipher suite negotiated, certificate details, and connection volume per agent endpoint without a SIEM. Write a Sigma rule targeting T1040 indicators: processes like `tcpdump`, `wireshark`, or `tshark` spawned on AI infrastructure hosts (query via osquery: `SELECT name, cmdline, pid FROM processes WHERE name IN ('tcpdump','tshark','dumpcap','netsniff-ng');`). For T1557 (Adversary-in-the-Middle), use `arpwatch` on the AI segment to detect ARP cache poisoning attempts that could enable passive TLS interception.

Evidence: Capture Zeek `ssl.log` entries for all AI agent communication sessions — specifically the `cipher` and `curve` fields — to document which sessions are using RSA key exchange (cipher names containing `RSA`) or ECDH with classical curves (`secp256r1`, `secp384r1`). Collect `conn.log` to map source/destination pairs across agent-to-tool and orchestration API paths. Retrieve firewall and router ACL logs for the AI infrastructure segment to identify any traffic routing through untrusted intermediary segments. These artifacts establish the HNDL attack surface: which encrypted sessions were exposed on which paths and for how long.

Step 3: Evaluate cryptographic agility posture — determine whether your AI infrastructure components (orchestration frameworks, model inference APIs, inter-agent messaging layers) support cipher suite configuration without code changes; flag hardcoded cryptographic dependencies as technical debt requiring prioritized remediation.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: assessing the organization's technical readiness to execute cryptographic transition as a precondition for any future eradication of RSA/ECC dependency from AI infrastructure

Controls: NIST SI-2 (Flaw Remediation) — treat hardcoded RSA/ECC dependencies in AI orchestration frameworks as documented flaws requiring a tracked remediation plan with assigned ownership, NIST SA-9 (External System Services) — evaluate whether third-party AI platform vendors (e.g., model API providers, agent hosting platforms) expose cryptographic agility controls or lock cipher suite selection, NIST CM-6 (Configuration Settings) — document cipher suite configuration baselines for AI infrastructure components and identify which allow runtime reconfiguration vs. require code changes, NIST CM-3 (Configuration Change Control) — establish a change control process for cryptographic configuration updates to AI infrastructure that tracks before/after cipher suite state, CIS 4.6 (Securely Manage Enterprise Assets and Software) — manage cryptographic configuration of AI infrastructure through version-controlled configuration files, not hardcoded values, CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — incorporate cryptographic agility assessment into the vulnerability management process as a tracked finding class for AI infrastructure

Compensating: For Python-based AI frameworks, run `grep -rn 'ssl.PROTOCOL_TLS\|TLSv1\|RSA\|DHE\|ECDHE\|ssl_version\|ciphers=' /path/to/ai_framework/ --include='*.py'` to identify hardcoded TLS parameters. Use `python3 -c "import ssl; ctx = ssl.create_default_context(); print(ctx.get_ciphers())"` to enumerate the default cipher suite list for the Python runtime used by AI agents. For containerized AI workloads, inspect Dockerfile and `docker inspect` output for `SSL_CERT_FILE`, `OPENSSL_CONF`, or cipher suite environment variables. Record all findings in a prioritized remediation backlog with columns: component, hardcoded algorithm, configurable (Y/N), owner, remediation path.

Evidence: Before flagging components, extract and preserve the current cryptographic configuration state of each AI infrastructure component: export TLS configuration files (e.g., `nginx.conf` snippets with `ssl_ciphers` directives, Python `ssl.SSLContext` initialization code, Java `SSLSocketFactory` configurations). For message queue layers (e.g., RabbitMQ, Kafka used for agent messaging), pull `rabbitmqctl tls_info` or Kafka broker `ssl.cipher.suites` property values. These configuration snapshots serve as the baseline against which PQC migration changes will be validated and as evidence that the organization identified and tracked the RSA/ECC dependency prior to any future regulatory inquiry about the transition timeline.

Step 4: Map to NIST PQC transition requirements — review SP 800-131A Rev. 2 deprecation timelines and assess your organization's gap against ML-KEM, ML-DSA, and SLH-DSA adoption; engage AI platform vendors on their PQC roadmaps and request written commitment dates.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: translating identified cryptographic risk into structured improvement actions, updating policies and vendor relationships to reflect the NIST PQC transition mandate as a forward-looking corrective measure

Controls: NIST SI-2 (Flaw Remediation) — map RSA/ECC deprecation per SP 800-131A Rev. 2 timelines to a tracked remediation plan; treat the 2030 NIST deprecation deadline as a hard remediation target for AI infrastructure, NIST SI-5 (Security Alerts, Advisories, and Directives) — treat NIST FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) finalization (August 2024) as authoritative directives requiring organizational response, NIST SA-9 (External System Services) — require AI platform vendors to provide documented PQC roadmaps with written commitment dates as a supply chain risk management action, NIST RA-3 (Risk Assessment) — document the gap between current RSA/ECC posture and ML-KEM/ML-DSA adoption as a formal risk item with estimated exposure window tied to HNDL threat timeline, CIS 7.2 (Establish and Maintain a Remediation Process) — incorporate PQC algorithm adoption milestones (ML-KEM for key encapsulation, ML-DSA for digital signatures) into the organization's risk-based remediation strategy with assigned deadlines, CIS 7.4 (Perform Automated Application Patch Management) — plan for cryptographic library updates (OpenSSL 3.x with liboqs integration, AWS-LC-FIPS) as application-layer

patch events requiring the same rigor as security patches

Compensating: Download and review NIST SP 800-131A Rev. 2 and the FIPS 203/204/205 final standards directly from NIST (csrc.nist.gov — search-retrieved, recommend human validation). Create a gap matrix spreadsheet with columns: AI infrastructure component, current algorithm, NIST deprecation date, target PQC algorithm, vendor confirmation status, target migration date. For open-source AI frameworks (LangChain, LlamaIndex, AutoGen), check GitHub issue trackers and release notes for PQC or post-quantum keyword mentions to assess community roadmap maturity. Send vendor inquiry templates via email and retain written responses as supply chain risk evidence.

Evidence: Collect and archive current vendor documentation — TLS library versions, cipher suite support matrices, and any existing vendor security advisories — for all AI platform providers before initiating vendor outreach. This establishes the pre-outreach state of vendor PQC readiness and timestamps when the organization became aware of specific gaps. Retain all vendor responses to PQC roadmap inquiries as formal records under NIST IR-8 (Incident Response Plan) documentation requirements and for potential regulatory or audit reference. Archive SP 800-131A Rev. 2 deprecation schedule alongside your gap matrix to document the compliance timeline basis.

Step 5: Update threat model and brief leadership — add HNDL as an active threat scenario in your threat register, mapped to AI infrastructure data flows; brief the CISO and relevant board members that encrypted AI traffic collected today may be decryptable in the future, and frame the cryptographic transition as an insurance action with a known cost and an unknown but real deadline.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: integrating HNDL as a recognized threat scenario into organizational threat modeling and risk governance, and communicating cryptographic transition risk to senior leadership consistent with lessons-learned and continuous improvement obligations

Controls: NIST IR-8 (Incident Response Plan) — update the IR plan to include HNDL against AI infrastructure as a named threat scenario with defined detection indicators (T1040, anomalous traffic volume on AI segments) and response procedures, NIST IR-6 (Incident Reporting) — establish reporting thresholds and escalation paths for HNDL indicators, including who briefs the CISO and board when passive collection evidence is detected on AI infrastructure segments, NIST RA-3 (Risk Assessment) — formally document HNDL as an active risk in the organizational risk register, with likelihood tied to adversary quantum capability timeline and impact tied to confidentiality of AI agent traffic content, NIST PM-9 (Risk Management Strategy) — brief the CISO and board using NIST PQC transition framing: known cost (migration effort), unknown deadline (quantum maturity), confirmed adversary behavior (HNDL documented by NSA and CISA advisories), CIS 7.1 (Establish and Maintain a Vulnerability Management Process) — add HNDL exposure of AI infrastructure to the vulnerability management process as a standing risk item with quarterly review cadence until ML-KEM migration is complete, CIS 7.2 (Establish and Maintain a Remediation Process) — document the PQC migration as a formal remediation item in the risk-based remediation strategy, with CISO sign-off and board awareness as governance milestones

Compensating: Use the MITRE ATT&CK Navigator (free, browser-based) to create a layer file mapping T1040 (Network Sniffing) and T1557 (Adversary-in-the-Middle) to the AI infrastructure data flows identified in Steps 1–2; export as a PDF for the leadership brief. Draft the CISO/board brief as a one-page risk memo structured as: threat (HNDL targeting AI agent traffic), evidence basis (NSA/CISA public advisories on HNDL, NIST FIPS 203/204/205 finalization), current exposure (cipher suite inventory findings from Step 1), cost to mitigate (migration effort estimate), and consequence of inaction (traffic collected today decryptable post-quantum). Retain the signed brief acknowledgment as an IR-6 and IR-8 governance artifact.

Evidence: Before the leadership brief, compile a threat intelligence summary referencing public HNDL documentation: NSA CNSS Advisory U/OO/194427-22 on quantum computing risks to national security systems, CISA post-quantum guidance, and the NIST IR 8547 migration considerations document. Export the AI data flow diagrams showing which agent communication paths carry potentially sensitive payloads (model inputs, tool outputs, orchestration instructions) — these data flow artifacts define the HNDL blast radius and form the factual basis for the board brief. Document the date of the leadership briefing and any decisions made as a formal record under NIST IR-8 plan maintenance requirements.

Detection Guidance

Direct detection of HNDL collection is inherently difficult, passive traffic capture leaves no immediate signature. Focus detection engineering on the conditions that make collection possible and on anomalies that suggest active interception attempts.

For T1040 (network sniffing): monitor for promiscuous mode activation on network interfaces within AI infrastructure segments; alert on unexpected network tap or span port configurations; review NetFlow and packet capture metadata for sustained, high-volume outbound transfers from network appliances or monitoring hosts that lack a documented justification.

For T1557 (adversary-in-the-middle): monitor TLS handshake anomalies on AI agent communication paths, unexpected certificate changes, certificate authority substitutions, or cipher suite downgrades; deploy certificate pinning where feasible on high-value agent communication channels and alert on pin violations; inspect DNS response integrity for AI infrastructure service resolution.

For T1021 and T1563 (lateral movement and session hijacking within AI systems): log and alert on authentication events for AI orchestration APIs, particularly token reuse from unexpected source IPs or unusual access times; monitor for agent identity impersonation, requests claiming to originate from a known agent but deviating from established behavioral baselines (request volume, endpoint patterns, timing).

Policy audit priorities: identify any AI system components still negotiating TLS with RSA or ECC key exchange (check server hello cipher suite logs); flag services configured to accept TLS 1.1 or below; audit certificate lifetimes and renewal processes for AI infrastructure certificates to ensure they can be rotated rapidly if a key compromise is suspected.

Framework Mappings

MITRE-ATTACK

- **T1040** — Network Sniffing
- **T1563** — Remote Service Session Hijacking
- **T1021** — Remote Services
- **T1557** — Adversary-in-the-Middle

NIST-800-53R5

- **AC-17** — Remote Access
- **AC-3** — Access Enforcement
- **CM-7** — Least Functionality
- **IA-2** — Identification and Authentication (Organizational Users)
- **SC-13** — Cryptographic Protection

OWASP-TOP10-2021

- **A02:2021** — Cryptographic Failures

ISO-27001-2022

- **A.8.24** — Use of cryptography

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

HIPAA-SECURITY

- **164.312(e)(1)** — Transmission Security

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1040	Network Sniffing	Credential-Access
T1563	Remote Service Session Hijacking	Lateral-Movement
T1021	Remote Services	Lateral-Movement
T1557	Adversary-in-the-Middle	Credential-Access

Sources

Source	URL	Tier
gemini	https://securityboulevard.com/2026/05/the-2026-roadmap-to-post-quan...	T3
Assessing the impact of quantum computing on infrastructure	https://www.sciencedirect.com/science/article/pii/S0167404826000933	T3
The risk of quantum to classical cryptography - Project Eleven	https://blog.projecteleven.com/posts/the-risk-of-quantum-to-classic...	T3
Securing Cryptography in the Age of Quantum Computing and AI	https://arxiv.org/html/2603.06969v1	T2
What are the emerging threats in cybersecurity due to the adoption ...	https://www.researchgate.net/post/What_are_the_emerging_threats_in_...	T3

Generated 2026-05-09 18:46 UTC by TJS Security Command Center