INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-11 07:39 UTC

China-Linked Threat Groups Target AI Infrastructure in 58% of State-Sponsored Tech Sector Attacks

THREAT CAMPAIGN | HIGH

SCC Item ID	SCC-CAM-2026-0438
Type	Threat Campaign
Severity	HIGH
Affected Products	AI and technology firms globally, AI training data, ML infrastructure, intellectual property
Published	2026-06-10
Discovery Source	Gemini

Executive Summary

China-linked APT groups account for over 58% of state-sponsored cyberattacks against the technology sector, with five named groups actively targeting AI training datasets, machine-learning infrastructure, and proprietary intellectual property at technology firms worldwide, according to CrowdStrike's 2026 Technology Threat Landscape Report. The campaign reflects a sustained, strategic effort to acquire foreign AI research through cyber espionage rather than organic development. Organizations building or operating AI systems face elevated risk of long-term competitive damage, IP theft, and compromise of the foundational assets that underpin AI product differentiation. Note: The 58% figure and actor designations originate from CrowdStrike's 2026 Technology Threat Landscape Report (vendor T2 source). Independent verification against primary sources (CISA, MITRE, law enforcement) is not available in current intelligence.

Technical Analysis

This campaign is a persistent espionage operation, not a single CVE-exploitable vulnerability. Five China-nexus APT groups are conducting targeted intrusions against AI and technology firms globally (CrowdStrike 2026 Technology Threat Landscape Report). Mapped MITRE ATT&CK techniques include: T1591 (Gather Victim Org Information, reconnaissance against AI teams and infrastructure), T1078 (Valid Accounts, use of stolen or compromised credentials to maintain access), T1213 (Data from Information Repositories, exfiltration of training data and model artifacts from internal wikis, code repos, and data stores), T1537 (Transfer Data to Cloud Account, staging exfiltrated AI assets in adversary-controlled cloud storage), T1199 (Trusted Relationship, targeting third-party vendors and partners with access to AI development environments), and T1102 (Web

Service, using legitimate web services for command-and-control to blend with normal traffic). Primary targets include ML model weights, training datasets, experiment tracking systems (MLflow, Weights and Biases), model registries, and AI-related source code repositories. No CVE, CWE, or CVSS scoring applies, this is a campaign-level intelligence item with no single exploitable vulnerability as the root cause. Source: CrowdStrike 2026 Technology Threat Landscape Report (secondary tier, vendor-sourced; not independently verified against a primary URL in this session).

Action Checklist

1. Step 1: Containment. Audit access to AI development environments immediately. Identify all accounts with access to ML infrastructure (model registries, training data stores, experiment tracking systems) and enforce least-privilege access per NIST AC-6 (Least Privilege), grant only permissions required for the account's current operational role. Disable or suspend any accounts with access that cannot be justified by current operational need. Apply CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) to AI pipeline tooling.
2. Step 2: Detection. Review authentication logs for anomalous use of valid accounts (T1078) across AI development infrastructure. Query for unusual data repository access patterns consistent with T1213, specifically bulk reads or exports from model registries, training data buckets, or internal wikis outside business hours or from unfamiliar source IPs. Enable and review audit logs per NIST AU-2 (Event Logging) and AU-6 (Audit Record Review, Analysis, and Reporting) for ML platforms, cloud storage buckets containing AI assets, and code repositories. Apply D3-LAM (Local Account Monitoring) to detect lateral movement using compromised local accounts.
3. Step 3: Eradication. Rotate all credentials with access to AI infrastructure per D3-CRO (Credential Rotation), prioritizing service accounts, API keys, and developer accounts with repository or cloud storage access. Enforce MFA on all externally exposed AI development tools, cloud consoles, and remote access paths per CIS 6.3 (Require MFA for Externally-Exposed Applications) and CIS 6.4 (Require MFA for Remote Network Access). Review and harden third-party and vendor access to AI environments per NIST AC-20 (Use of External Systems), addressing T1199 (Trusted Relationship).
4. Step 4: Recovery. Validate that no unauthorized data transfers to external cloud accounts occurred by reviewing cloud egress logs and storage access logs aligned with T1537 detection. Confirm audit logging is active and complete across all AI development infrastructure per CIS 8.2 (Collect Audit Logs) and NIST AU-12 (Audit Record Generation). Revalidate account inventories against CIS 5.1 (Establish and Maintain an Inventory of Accounts) and CIS 5.3 (Disable Dormant Accounts). Monitor for re-entry using C2 over legitimate web services (T1102) by reviewing proxy and DNS logs for anomalous outbound patterns.
5. Step 5: Post-Incident. Assess whether separation of duties controls (NIST AC-5) are enforced across AI development, data access, and deployment pipelines to limit blast radius of a future compromise. Conduct a formal data inventory review per CIS 3.2 (Establish and Maintain a Data Inventory) scoping all AI training data, model artifacts, and experiment records as sensitive assets. Establish or update an information flow policy per NIST AC-4 (Information Flow Enforcement) to restrict cross-system movement of AI IP. Brief AI product and research leadership on targeted social engineering and trusted-relationship vectors (T1199, T1591).

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate immediately to executive leadership and legal counsel if cloud egress analysis confirms exfiltration of model weights, training datasets, or experiment records containing proprietary AI IP, as this may trigger trade secret notification obligations and, if export-controlled data or ITAR-relevant AI research is involved, mandatory USG reporting to CISA and the FBI Cyber Division per CISA's known China-linked APT notification guidance.
Recovery Notes	After containment and credential rotation, maintain elevated monitoring of ML pipeline egress and authentication logs for a minimum of 90 days, as China-linked APT groups attributed in the CrowdStrike report (e.g., Volt Typhoon, Salt Typhoon) are known to pre-position access and re-enter through previously undetected footholds weeks after initial remediation. Verify the integrity of all model artifacts in the registry by comparing current SHA-256 hashes against pre-incident baselines before resuming model training or deployment, as training data poisoning or model backdooring may have been a secondary objective beyond pure IP theft. Establish a weekly access review cadence for all AI infrastructure accounts for at least two quarters post-incident, specifically checking for new service account creation or IAM policy changes that could indicate re-establishment of persistent access.
Forensic Artifacts	AWS CloudTrail data event logs for S3 GetObject/PutObject on AI training data buckets — China-linked actors exfiltrating AI IP will produce high-volume GetObject sequences from a single IAM principal, often followed by PutObject events to external bucket ARNs (T1537), within a compressed timeframe inconsistent with normal developer activity patterns | MLflow or equivalent experiment tracking server access logs (default path: mlflow server logs, or application container stdout) — look for REST API calls to GET /api/2.0/mlflow/artifacts/get and GET /api/2.0/mlflow/runs/search with unusually broad search parameters, indicating automated bulk enumeration of experiment runs and associated model artifacts consistent with T1213 (Data from Information Repositories) | Git provider webhook and API audit logs (GitHub Enterprise audit log API endpoint or GitLab audit events API) — China-linked APT groups targeting AI IP frequently abuse developer access tokens to clone entire repository groups or download release archives; audit log entries for 'git.clone' events on repositories containing model training code, dataset preprocessing pipelines, or proprietary architecture definitions are primary indicators | VPC Flow Logs or equivalent NetFlow records for ML compute subnets showing sustained outbound TCP sessions on port 443 to Alibaba Cloud (AS37963), Tencent Cloud (AS132203), or anonymizing infrastructure, with session byte counts exceeding 100 MB from individual ML training nodes — these are atypical for normal inference or training workloads and consistent with staged exfiltration of model checkpoints | Linux kernel audit logs (auditd, /var/log/audit/audit.log) on ML compute nodes configured to capture SYSCALL events for 'open', 'read', and 'write' on model artifact file paths — an actor with valid account access conducting hands-on-keyboard reconnaissance of model weight files (typically .pkl, .pt, .h5, .onnx, .safetensors extensions) will produce a distinct read-then-compress-then-transfer sequence detectable via auditd rules targeting those file extensions in model registry mount paths

Per-Action IR Details

Step 1: Containment — Audit access to AI development environments immediately. Identify all accounts with access to ML infrastructure (model registries, training data stores, experiment tracking systems) and enforce least-privilege access per NIST AC-6 (Least Privilege). Disable or suspend any accounts with access that cannot be justified by current operational need. Apply CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts) to AI pipeline tooling.

NIST Phase: Containment

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

Controls: NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: Run 'net user /domain' and 'net localgroup' (Windows) or 'getent passwd' + 'getent group' (Linux) on MLflow, Jupyter Hub, and DVC hosts to enumerate all accounts with access to model registry directories and training data mount points. Cross-reference against your last known-good access list using a diff script. Use osquery ('SELECT * FROM users; SELECT * FROM logged_in_users;') on ML pipeline nodes to identify active sessions. Immediately disable unrecognized accounts via 'usermod -L ' (Linux) or 'Disable-ADAccount' (PowerShell) without deletion to preserve forensic state.

Evidence: Before disabling any account, capture: full output of 'last' and 'lastlog' on Linux ML hosts to establish account activity timeline; Windows Security Event Log Event ID 4624 (Logon) and 4648 (Explicit Credential Use) filtered to accounts with access to model registry paths (e.g., /opt/mlflow, s3://ai-training-*); cloud IAM access advisor reports showing last-used date for every IAM role/user with 's3:GetObject' or 'sagemaker:*' permissions; Jupyter Notebook server logs (typically ~/.jupyter/jupyter.log or /var/log/jupyter/) for session origins; and MLflow tracking server access logs showing experiment read/download events by account and source IP.

Step 2: Detection — Review authentication logs for anomalous use of valid accounts (T1078) across AI development infrastructure. Query for unusual data repository access patterns consistent with T1213 — specifically, bulk reads or exports from model registries, training data buckets, or internal wikis outside business hours or from unfamiliar source IPs. Enable and review audit logs per NIST AU-2 (Event Logging) and AU-6 (Audit Record Review, Analysis, and Reporting) for ML platforms, cloud storage buckets containing AI assets, and code repositories. Apply D3-LAM (Local Account Monitoring) to detect lateral movement using compromised local accounts.

NIST Phase: Detection Analysis

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

Controls: NIST AU-2 (Event Logging), NIST AU-6 (Audit Record Review, Analysis, And Reporting), NIST AU-3 (Content Of Audit Records), CIS 8.2 (Collect Audit Logs)

Compensating: For cloud storage (AWS S3): enable S3 server access logging and CloudTrail data events if not already active, then query CloudTrail logs using 'aws cloudtrail lookup-events --lookup-attributes AttributeKey=EventName,AttributeValue=GetObject' filtered to AI training buckets — flag any session downloading >1 GB or >500 objects in a single session. For on-prem Git repositories (GitLab/GitHub self-hosted): parse nginx or Apache access logs with 'grep -E "GET.*/archive|GET.*/raw" /var/log/nginx/access.log' to detect bulk repository archive downloads. Deploy the Sigma rule 'proc_creation_win_susp_data_exfil_via_cli.yml' via Chainsaw on Windows ML hosts to detect CLI-based bulk transfers. Use Sysmon Event ID 3 (Network Connection) filtered on rclone.exe, azcopy.exe, or aws.exe making outbound connections to non-corporate cloud endpoints.

Evidence: Collect before analysis: AWS CloudTrail logs for S3 GetObject/ListBucket events on AI training data buckets, filtered to the 90 days preceding detection (China-linked APTs typically conduct slow, low-volume staging before bulk exfiltration); Git repository access logs showing 'git clone --mirror' or archive download events (these leave distinct HTTP 200 responses on large .tar.gz payloads); MLflow artifact store access logs showing model download events (GET /api/2.0/mlflow/artifacts/get) from IPs outside corporate CIDR ranges; SSH auth logs (/var/log/auth.log or /var/log/secure) for successful logins to ML compute nodes from internal pivot IPs not associated with the legitimate developer; and Kubeflow/Kubernetes API server audit logs (kube-apiserver audit log) for 'get' or 'list' verbs against model serving namespaces by unexpected service accounts.

Step 3: Eradication — Rotate all credentials with access to AI infrastructure per D3-CRO (Credential Rotation), prioritizing service accounts, API keys, and developer accounts with repository or cloud storage access. Enforce MFA on all externally exposed AI development tools, cloud consoles, and remote access paths per CIS 6.3 (Require MFA for Externally-Exposed Applications) and CIS 6.4 (Require MFA for Remote Network Access). Review and harden third-party and vendor access to AI environments per NIST AC-20 (Use of External Systems), addressing T1199 (Trusted Relationship).

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication

Controls: NIST AC-20 (Use Of External Systems), NIST AC-17 (Remote Access), 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)

Compensating: Enumerate all AWS IAM access keys and service account tokens using 'aws iam generate-credential-report' then rotate every key associated with S3 buckets, SageMaker, or EC2 ML instances using 'aws iam create-access-key' / 'aws iam delete-access-key' — do not rotate without first confirming the new key is functional in the target service to avoid pipeline outages. For on-prem systems, use 'ssh-keygen' to regenerate SSH host keys on ML compute nodes if lateral movement via SSH is confirmed. Audit third-party vendor SSO federation trusts in your IdP (Okta, Azure AD): run 'Get-MsolServicePrincipal | Where-Object {$_.DisplayName -like "**"}' to identify external app registrations with delegated access to AI development tenants, and revoke any that cannot be validated against a current vendor contract. Enable TOTP-based MFA on Jupyter Hub using the nativeauthenticator plugin for teams without enterprise SSO.

Evidence: Before rotating credentials, capture: a full export of AWS IAM credential report ('aws iam generate-credential-report && aws iam get-credential-report') timestamped at incident time to document which keys were active and last used; the complete list of active API tokens from MLflow, Weights & Biases, or your experiment tracking platform's admin panel (these are often long-lived and overlooked); third-party vendor access logs from your VPN concentrator or zero-trust proxy showing connection frequency, source geolocation, and data volumes for each vendor session over the prior 60 days; and any OAuth token grant records from your Git provider (GitHub Enterprise/GitLab) showing active personal access tokens with 'repo' or 'read:packages' scope held by contractor or vendor accounts.

Step 4: Recovery — Validate that no unauthorized data transfers to external cloud accounts occurred by reviewing cloud egress logs and storage access logs aligned with T1537 detection. Confirm audit logging is active and complete across all AI development infrastructure per CIS 8.2 (Collect Audit Logs) and NIST AU-12 (Audit Record Generation). Revalidate account inventories against CIS 5.1 (Establish and Maintain an Inventory of Accounts) and CIS 5.3 (Disable Dormant Accounts). Monitor for re-entry using C2 over legitimate web services (T1102) by reviewing proxy and DNS logs for anomalous outbound patterns.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery

Controls: NIST AU-12 (Audit Record Generation), NIST AU-11 (Audit Record Retention), NIST AU-4 (Audit Storage Capacity), CIS 8.2 (Collect Audit Logs), CIS 5.1 (Establish and Maintain an Inventory of Accounts), CIS 5.3 (Disable Dormant Accounts)

Compensating: To detect T1537 (Transfer Data to Cloud Account), query AWS CloudTrail for 'PutObject' events to S3 bucket ARNs outside your organization's account IDs using: 'aws cloudtrail lookup-events --lookup-attributes AttributeKey=EventName,AttributeValue=PutObject' then filter results where 'requestParameters.bucketName' does not match your known bucket inventory. For T1102 C2 detection without a SIEM, capture a 24-hour DNS query log from your resolver (unbound query log or Windows DNS debug log) and run frequency analysis using 'sort | uniq -c | sort -rn' to surface low-frequency domains queried by ML hosts — flag any queries to Pastebin, GitHub Gist, Notion, or Slack API endpoints from compute nodes that have no business reason to reach those services. Use Wireshark or tcpdump on ML network egress points with filter 'host and port 443 and not (host)' to capture and baseline outbound TLS sessions.

Evidence: Before closing recovery, preserve: VPC Flow Logs or on-prem NetFlow records for the full suspected intrusion window showing all outbound connections from ML compute subnets, with byte counts — large outbound transfers to cloud provider IPs (AWS, Alibaba Cloud, Tencent Cloud) outside your normal CI/CD pipeline patterns are a primary T1537 indicator for China-linked actors; DNS query logs from ML hosts for the 90-day window preceding detection, as China-linked APTs frequently use DNS-over-HTTPS or domain-fronting for C2 that blends into legitimate traffic; proxy logs showing HTTP CONNECT or HTTPS CONNECT requests from ML nodes to unexpected destinations; and a hash-verified snapshot of current model registry contents (MLflow model artifacts, HuggingFace model cards, or SageMaker model packages) compared against your last known-good backup to detect model

tampering or poisoning as a secondary objective.

Step 5: Post-Incident — Assess whether separation of duties controls (NIST AC-5) are enforced across AI development, data access, and deployment pipelines to limit blast radius of a future compromise. Conduct a formal data inventory review per CIS 3.2 (Establish and Maintain a Data Inventory) scoping all AI training data, model artifacts, and experiment records as sensitive assets. Establish or update an information flow policy per NIST AC-4 (Information Flow Enforcement) to restrict cross-system movement of AI IP. Brief AI product and research leadership on targeted social engineering and trusted-relationship vectors (T1199, T1591).

NIST Phase: Post Incident

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

Controls: NIST AC-5 (Separation Of Duties), NIST AC-4 (Information Flow Enforcement), NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 3.2 (Establish and Maintain a Data Inventory), CIS 3.3 (Configure Data Access Control Lists), CIS 3.4 (Enforce Data Retention)

Compensating: To enforce separation of duties on AI pipelines without enterprise tooling, implement Git branch protection rules requiring two-reviewer approval for any merge to branches that trigger model training or deployment workflows — this prevents a single compromised developer account from pushing poisoned training data or exfiltrating model weights under the cover of a legitimate commit. Conduct the data inventory using a Python script that recursively hashes and catalogs all files in model registry directories and S3-equivalent buckets (e.g., 'find /opt/mlflow/artifacts -type f -exec sha256sum {} \;' output to a timestamped CSV), establishing a tamper-detection baseline. For the leadership briefing on T1591 (Gather Victim Org Information), pull your organization's public GitHub repositories, LinkedIn employee profiles for AI/ML roles, and conference presentation records — China-linked APTs including those named in the CrowdStrike report have used publicly available researcher profiles and open-source model contributions to identify and target specific individuals for spearphishing.

Evidence: For post-incident review, compile: a complete timeline of all model artifact upload and download events from the experiment tracking platform (MLflow, W&B, or Comet) for the full intrusion window to determine whether model weights or training checkpoints were exfiltrated; a diff of your data access control lists (S3 bucket policies, IAM policies, filesystem ACLs) between the pre-incident baseline and current state to identify any persistence mechanisms left via policy modification (T1098 — Account Manipulation); email and calendar logs for AI research staff who were targeted or whose credentials were involved, to identify spearphishing precursors consistent with T1591 reconnaissance; and a record of all third-party integrations (CI/CD webhooks, cloud marketplace apps, vendor API integrations) that had write access to AI development environments, to assess T1199 trusted-relationship exposure scope.

Detection Guidance

Focus detection on four behavioral patterns consistent with this campaign's mapped techniques. (1) Valid account abuse (T1078): Alert on successful authentications from new geolocations, IP ranges, or user agents for accounts with access to AI repositories, cloud ML platforms, or training data stores. Cross-reference NIST AU-6 audit review cadence. Apply D3-LAM (Local Account Monitoring) to flag privilege escalation or lateral movement using compromised accounts. (2) Bulk data repository access (T1213): Query data access logs for high-volume reads, exports, or clones from model registries, S3/GCS/Azure Blob buckets containing training data, MLflow or similar experiment tracking systems, and internal code repositories. Threshold-based alerts on export volume outside normal working patterns are a practical starting point. (3) Cloud data staging (T1537): Review cloud storage access logs for creation of new buckets or transfer of large volumes of AI-related data to accounts or regions inconsistent with your normal operations. (4) C2 via web services (T1102): Inspect proxy and DNS logs for repeated, low-and-slow outbound calls to legitimate cloud services (GitHub, Google Drive, OneDrive, Pastebin, Slack APIs) from systems that do not normally communicate with those services, particularly from ML training hosts or data pipeline servers. For IOC-based detection, monitor threat intelligence

platforms (CISA AIS, MITRE ATT&CK) for updates on these groups. Immediate focus should be on behavioral rules as described above.

Indicators of Compromise

Type	Value	Context	Confidence
DOMAIN	`not available`	No confirmed IOCs for this campaign are available from independently verified primary sources in this session. The CrowdStrike 2026 Technology Threat Landscape Report is the named source; specific IOCs were not independently extracted or verified here.	LOW

Framework Mappings

MITRE-ATTACK

- **T1102** — Web Service
- **T1591** — Gather Victim Org Information
- **T1078** — Valid Accounts
- **T1213** — Data from Information Repositories
- **T1537** — Transfer Data to Cloud Account
- **T1199** — Trusted Relationship

NIST-800-53R5

- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management

ISO-27001-2022

- **A.5.21** — Managing information security in the ICT supply chain

SOC2-TSC

- **CC9.2** — Manages risks associated with vendors and business partners

MITRE ATT&CK Mapping

Technique ID	Technique Name	Tactic
T1102	Web Service	Command-And-Control
T1591	Gather Victim Org Information	Reconnaissance

Technique ID	Technique Name	Tactic
T1078	Valid Accounts	Defense-Evasion
T1213	Data from Information Repositories	Collection
T1537	Transfer Data to Cloud Account	Exfiltration
T1199	Trusted Relationship	Initial-Access

Sources

Source	URL	Tier
Adversaries Leverage AI for Vulnerability Exploitation, Augmented ...	https://cloud.google.com/blog/topics/threat-intelligence/ai-vulnera...	T3
8 AI Cybersecurity Companies For 2026 - SentinelOne	https://www.sentinelone.com/cybersecurity-101/data-and-ai/ai-cybers...	T3
Top AI Security Vulnerabilities to Watch out for in 2026 - Cycode	https://cycode.com/blog/ai-security-vulnerabilities/	T3
Top 7 AI Security Risks - Sysdig	https://www.sysdig.com/learn-cloud-native/top-7-ai-security-risks	T3
Experts, tech gurus raise new concerns over AI and global security	https://www.youtube.com/watch?v=eAr8zzpgM7Y	T3

Generated 2026-06-11 07:39 UTC by TJS Security Command Center