2026 Threat Intelligence: Q1 Validation by Claude

Validation of predictions from Santosh Pandit's 2026 CTI Forecast (Parts 1 and 2, December 2025) against confirmed cyber incidents and APT activity from January to March 2026.

Executive Summary

14 Confirmed	7 Partial	3 Emerging	1 Not Yet Observed	25 Total Predictions

Legend: Confirmed = multiple independent sources documenting incidents matching the prediction. Partial = some elements validated, others await evidence or the supporting evidence predates Q1 2026. Emerging = early indicators consistent with prediction but insufficient for full confirmation. Not Yet = no supporting evidence found in Q1 2026.

Part 1: Threat Actor Predictions

Pandit identified 11 threat actors ranked across three tiers. Below is a validation of each actor's predicted 2026 activity against Q1 2026 evidence.

Threat Actor	Tier	Prediction (Dec 2025)	Status	Q1 2026 Evidence
Qilin	Tier 1	Continued aggressive operations; innovation in psychological pressure tactics; potentially incorporating AI for victim negotiation.	CONFIRMED	Qilin was the most active ransomware group in Q1 2026, topping Bitdefender, BitSight, and Purple Ops rankings. It claimed the most US victims in Jan-Feb 2026 and attacked Romania's Conpet (oil pipeline operator), alleging theft of roughly 1 TB of documents. It maintained consistent top-10 placement for over four consecutive months with multi-sector targeting across healthcare, government, energy, and professional services.
LockBit	Tier 1	Fifth-generation resurgence; CNI targeting; supply chain compromises; AI-enhanced lateral movement expected.	PARTIAL	LockBit's presence persisted but at reduced volume compared to peak 2024 activity. Affiliates migrated to other groups (DragonForce, Qilin). LockBit's leaked codebase was adopted by DragonForce and other emerging RaaS operators, validating the resilience prediction. Its direct operational volume lagged behind Qilin and Akira in Q1. CNI-specific targeting by LockBit-branded operations was not prominently confirmed in the quarter.
Lazarus Group	Tier 1	Hybrid criminal-state operations blending cybercrime with espionage; critical for TLPT exercises.	CONFIRMED	Lazarus attacked crypto platform Bitrefill on 1 March 2026, draining hot wallets and accessing 18,500 purchase records through a compromised employee laptop. This was a textbook hybrid revenue-generation and espionage operation. Continued fallout from the Feb 2025 Bybit heist ($1.5B stolen) persisted with ongoing laundering operations tracked into Q1 2026. The attack pattern matched prior Lazarus campaigns against Ronin Network, Harmony Bridge, and WazirX.

RansomHub	**Tier 2**	Despite disruption, continued operations likely under rebranded identities; resilient threat actor network that may splinter.	**PARTIAL**	RansomHub's infrastructure went down in April 2025. DragonForce publicly invited RansomHub affiliates to migrate. Elite affiliates scattered to Qilin, DragonForce, and other RaaS platforms, confirming the predicted splintering. However, no RansomHub-branded infrastructure or victim claims appeared in Q1 2026 tracking by Bitdefender, Cyble, or BitSight. The prediction of continued operations "under rebranded identities" is validated only through successor group activity rather than direct RansomHub operations.
Cl0p	**Tier 2**	Continued zero-day exploitation of file transfer applications; abundant attack surface from internet-facing file transfer services.	**CONFIRMED**	Cl0p remained in the top ransomware groups in Q1 2026. It continued exploitation of Oracle E-Business Suite flaws from late 2025, with 29 alleged victims listed on its leak site including Harvard University, Washington Post, and an American Airlines subsidiary. Blackpoint reported Cl0p published 43 global victims within a 24-hour period in January 2026. Cleo file-transfer exploitation continued alongside the Oracle campaign. Bitdefender ranked Cl0p consistently in the top 10 for US victims throughout the period.
Akira	**Tier 2**	Continued exploitation of firewalls and VPNs; industrialised RaaS ecosystem with no signs of slowing.	**CONFIRMED**	Akira maintained top-5 placement continuously from 2025 into Q1 2026 per Bitdefender, BitSight, and Cyble. It targeted iSMA CONTROLLI (building management systems, March 2026). CISA warnings about Akira's critical infrastructure targeting remained active. It ranked among the most active groups with an estimated 50+ US victims in January 2026 alone (Bitdefender).
DragonForce	**Tier 3**	Continued rise with potential AI adoption; aggressive affiliate recruitment; ambitions to join the top tier of ransomware operators.	**CONFIRMED**	DragonForce's rise was one of the strongest validated predictions. It ranked consistently in Bitdefender's top 10 for four or more consecutive months. Q1 2026 attacks included Salford City College (UK), Pride Solvents (US chemicals), STS Travel (Mexico), Fundacao Getulio Vargas (Brazil), and dozens more. It evolved to a 'cartel' RaaS model offering white-label ransomware. It recruited actively from RansomHub's collapsed affiliate network. Its affiliates (including Scattered Spider) conducted the high-profile UK retail attacks on M&S, Co-op, and Harrods.
Anubis	**Tier 3**	Destructive file-wiping with ideological motivation; irreversible damage without offering decryption.	**CONFIRMED**	Anubis maintained active operations in Q1 2026, consistent with its persistent niche profile. BitSight tracked an attack on A J Taylor Electrical (UK, 10 March). DeXpose confirmed an attack on Schlam Stone & Dolan LLP (US law firm, 27 March). The group's destructive, public-shaming approach continued. While

				the scale remained smaller than profit-driven groups, the actor profile and TTP pattern were validated as predicted.
Play	Tier 3	Likely to remain active; potential escalation in hybrid warfare contexts in regions of geopolitical tension.	**CONFIRMED**	Play ransomware maintained top-10 ranking per Bitdefender for over four consecutive months. It targeted Esquire Brands (footwear, January 2026) with data exfiltration threats. It remained active across government and CNI sectors in LatAm and Europe as predicted, with a steady operational tempo throughout Q1 2026.
SLH Alliance	Tier 3	Continued social engineering of IT service desks by native English speakers; "Trojan Horse" approach via trusted internal processes.	**CONFIRMED**	SLH's legacy was validated through ongoing affiliate activity. Bitdefender confirmed SLH/Scattered Spider/ShinyHunters led massive supply chain attacks in 2025, and their tactics continued into 2026. DragonForce's UK retail attacks used Scattered Spider social-engineering TTPs targeting help desks. The 2025-2026 ShinyHunters/Scattered Spider campaign compromised 760+ organisations through voice phishing (Vectra AI). Vishing attacks surged 442% (CrowdStrike H2 2024 data), with help desks confirmed as the primary social-engineering target.
Emotet	Tier 3	GOAT classification; potential return to frontline operations from current IAB role; resilient botnet infrastructure.	**NOT YET**	No prominent Emotet resurgence was documented in Q1 2026. Its infrastructure continued operating in the background, but no headline-making frontline campaigns materialised. The prediction was hedged ("could return") rather than definitive. The IAB role continued quietly. This remains a watch item for the rest of 2026.

Part 2: TTP System and Emerging Trend Predictions

Pandit forecast seven TTP Systems and seven cross-cutting trends. Below is how they mapped against Q1 2026 reality.

TTP System / Trend	Tier	Prediction (Dec 2025)	Status	Q1 2026 Evidence
TTP Systems (from Part 2 Priority TTP Innovations)				
Agentic AI Attack Chains	Tier 1	Transition from scripted to agentic AI; autonomous reconnaissance and lateral movement; time-to-exploit compression to 8-16 hours; high adoption for Tier 1 actors by mid-2026.	**PARTIAL**	The foundational precedent was set by the Anthropic-disclosed GTG-1002 incident (September 2025, publicly reported November 2025), where a Chinese state actor weaponised Claude Code to automate 80-90% of a large-scale espionage campaign targeting roughly 30 organisations. This confirmed that agentic cyberattacks are technically viable and operationally effective. However, the original prediction specifically targeted "high adoption by Tier 1 RaaS actors (LockBit,

				Lazarus) by mid-2026" and autonomous ransomware chains. No Q1 2026 ransomware incidents attributed to LockBit, Lazarus, or other RaaS groups have been documented as using agentic AI. Bitdefender confirmed attackers are using automation to shorten time-to-exploit for disclosed PoCs, and 48% of Dark Reading survey respondents identified agentic AI as the top 2026 attack vector. The OWASP Top 10 for Agentic Applications 2026 was published. The trajectory supports the prediction for mid-to-late 2026 adoption.
Professionalisation of Reputational Extortion	**Tier 1**	Multi-vector pressure (DDoS + leaks + media + regulatory complaints); regulatory weaponisation; media amplification; dedicated negotiation teams; supply chain disclosure weaponisation.	**CONFIRMED**	Approximately 74% of all ransomware incidents in Q1 2026 included data exfiltration alongside encryption, confirming multi-vector extortion as standard operating procedure. DragonForce directly contacted the BBC to amplify pressure on the Co-op breach, publicly contradicting the retailer's own disclosure narrative. Qilin and DragonForce both operated structured media-engagement campaigns. Unit 42's Global Incident Response Report confirmed that operational disruption is now intentionally weaponised to force higher payments. Ransomware actors increasingly used regulatory reporting deadlines as pressure tools against victims.
Lazarus AI-Enhanced Hybrid Operations	**Tier 1**	Dual-revenue fusion (espionage + ransomware simultaneously); AI arms race leadership 1-2 quarters ahead of RaaS groups; supply chain dominance in financial and critical sectors.	**PARTIAL**	Lazarus's hybrid model was strongly validated through the March 2026 Bitrefill attack (crypto theft combined with data access) and continued Bybit laundering operations. Supply chain targeting of crypto platforms was confirmed. However, specific evidence of Lazarus deploying agentic AI (as distinct from the Chinese GTG-1002 case) has not been publicly attributed in Q1. The AI arms race prediction remains plausible but unconfirmed for Lazarus specifically.
Living Off the Cloud (LotC) and API Hijacking	**Tier 2**	Cloud-native encryption via stolen API credentials; IAM permission revocation; cloud backup poisoning; no malware signatures, bypasses EDR.	**PARTIAL**	Identity-first compromise was confirmed as a dominant trend. Bitdefender observed ransomware groups prioritising credential theft (browser session tokens) over active exploitation. In one notable case, Marquis alleged that attackers leveraged SonicWall cloud backup API changes to conduct a breach. Tenable identified machine and non-human identities as the number one cloud breach vector for 2026. However, documented cases of pure LotC

				attacks (AWS KMS hijacking, IAM revocation) at scale were not yet prominently reported in Q1. Adoption appears in early stages as predicted.
Vishing-as-a-Service (VaaS) and Deepfake Identity	**Tier 2**	AI voice-cloning professionalised within RaaS ecosystems; CEO/CTO voice fraud; help desk becomes primary target; mainstream RaaS adoption by H2 2026.	**CONFIRMED**	Vishing attacks surged 442% in H2 2024, and the trend accelerated into 2026. Deepfake-enabled vishing grew over 1,600% compared to late 2024 (Breacher.ai). Voice cloning now requires just 3 seconds of audio (Microsoft VALL-E research). Fortune's 2026 deepfake outlook confirmed that AI-generated voice crossed the "indistinguishable threshold." Group-IB's High-Tech Crime Trends Report 2026 dedicated major sections to industrialised deepfake vishing. The FBI issued PSA250515 warning of AI voice impersonation of senior US government officials. The 2025-26 ShinyHunters/Scattered Spider campaign used vishing to compromise 760+ organisations. Help desks were confirmed as the primary social-engineering target across multiple sources.
Adversarial ML / Model Poisoning	**Tier 3**	"Snow Features" injection; silent model degradation; first major model integrity breach by late 2026. Pandit described this as "highly speculative."	**EMERGING**	Theoretical validation continued but operational evidence remains limited. OWASP's 2026 Agentic AI Top 10 included "Memory Poisoning" as a key threat category. Stellar Cyber reported 520 tool misuse and privilege escalation incidents and emerging memory poisoning attacks against agentic AI systems in 2026. Research showed a single poisoned agent could compromise 87% of downstream decisions in 4 hours. However, no confirmed large-scale "Snow Features" or silent EDR degradation incident was publicly disclosed in Q1 2026. This is consistent with Pandit's own medium-confidence assessment and late-2026 timeline.
Supply Chain Compromise (punishing bad behaviours)	**Tier 3**	Deliberate targeting of vendors with poor disclosure practices; cascading pressure; residual vulnerability amplification; vendor pre-vetting by attackers.	**CONFIRMED**	Supply chain attacks accelerated markedly in Q1 2026. Cl0p's continued Oracle E-Business Suite exploitation created cascading impact across downstream organisations (Harvard, Dartmouth, American Airlines subsidiary). Bitdefender confirmed "scaled attacks to take over vendor chains" as a defining behaviour of Q1 2026. Sinobi targeted an Indian IT services company, gaining access to Hyper-V servers, virtual machines, and customer backups (cascade effect). A Trivy supply-chain attack on Docker and GitHub

				was discovered. The Marquis/SonicWall breach exemplified targeting of weak third-party security partners. Blue Yonder (supply chain technology) was hit by Cl0p while still recovering from a prior breach.
Cross-Cutting Emerging Trends (from Part 2 Emerging Trends for 2026)				
AI as Universal Accelerant	--	AI will accelerate reconnaissance, lateral movement, exfiltration, and negotiation. 2026 attacks will be faster, more adaptive, and harder to interrupt.	**CONFIRMED**	MIT research found that roughly 80% of ransomware attacks in recent studies used AI for phishing, deepfakes, password cracking, and evasion. PromptLock, described as the first AI-powered ransomware, was discovered in August 2025 and its implications carried into 2026 threat assessments. Bitdefender confirmed that automation is shortening the time-to-exploit window for disclosed PoCs. Ransomware attack velocity increased approximately 30% over the prior nine-month average (Cyble), sustained over four consecutive months, indicating operational acceleration enabled in part by AI tooling.
Identity as the New Perimeter	--	Deepfakes, VaaS, API hijacking, and credential theft all target identity infrastructure. Zero Trust alone will be insufficient without extreme rigor.	**CONFIRMED**	Identity-first compromise became a dominant Q1 2026 pattern. Bitdefender stated that "more ransomware groups are focusing on identity-first compromise, prioritising credential theft over more active means of attack." Lazarus used stolen employee laptop credentials at Bitrefill. BYOVD (Bring Your Own Vulnerable Driver) for defence evasion resurged. Tenable identified non-human identities (NHIs) as the number one cloud breach vector. SecurityWeek predicted the "collapse of perimeter thinking" for 2026.
Reputational Damage as Primary Objective	--	Reputational damage and regulatory pressure become co-equal objectives with encryption, sometimes superseding ransom demands.	**PARTIAL**	The trend is clearly emerging but not fully dominant. The 74% data exfiltration rate confirms the dual-pressure model. DragonForce's outreach to the BBC demonstrated reputation weaponisation in practice. Ransom payments are declining (Bitdefender), suggesting actors rely more on reputational and data-leak threats. However, encryption still accompanies most attacks. Pure "reputation-only" campaigns without encryption are not yet standard in Q1.
Critical Infrastructure Focus	--	Geopolitical tensions will drive increased attacks on utilities, transportation, and communication infrastructure.	**CONFIRMED**	Healthcare led all ransomware-targeted industries in January 2026 with 27 incidents (BlackFog). University of Mississippi Medical Center shut down clinics after a February ransomware attack.

				Foster City (California) declared a state of emergency following a ransomware intrusion. LA Metro systems were breached. Peru's National Water Authority was attacked with 2 TB of data stolen. Stryker Corporation (medical devices, over $25B annual revenue) was hit by an Iran-linked wiper attack on 11 March 2026, described by US officials as the most significant wartime cyberattack against American targets. The FBI launched Operation Winter SHIELD to protect hospitals and CNI. VoltRuptor ICS/SCADA malware was deployed against critical infrastructure.
Ransomware Volume Surge	--	The 50% ransomware increase trend will continue; industrialisation and convergence of cybercrime into enterprise-scale operations.	**PARTIAL**	Ransomware attack velocity increased approximately 30% over the prior nine-month average (Cyble), sustained over four consecutive months. Bitdefender reported 750 to 800 US organisations were hit in January and February 2026 alone, with 53 ransomware groups active in the US. However, the growth rate is closer to 30% rather than the predicted 50%. Ransomware payments are simultaneously declining even as volume rises. On current trajectory, publicly named victims may reach 7,000 by year-end 2026, which would represent a fivefold increase since 2020 (Zero Networks).
N-Day Exploitation Dominance	--	Known vulnerabilities with existing patches will force revision of patching and vulnerability management practices across healthcare, defence, and aerospace.	**EMERGING**	Unpatched known vulnerabilities remain a primary entry vector. Cl0p's ongoing exploitation of Oracle EBS (disclosed in 2025, still exploited in 2026) is a clear example. Verizon DBIR 2025 found that 32% of ransomware attacks were orchestrated through vulnerability exploitation. New Citrix NetScaler CVEs demanded urgent patching in Q1. The SonicWall cloud backup API vulnerability was exploited. However, a definitive industry-wide revision of patching practices has not yet materialised. Regulatory proposals (CIRCIA, HIPAA updates) are in development but not yet enacted.
Disclosure Weaponisation	--	Poor disclosure practices will become an organisational liability; attackers will deliberately target vendors known for weak communication.	**EMERGING**	Patterns consistent with this prediction are visible but still nascent. DragonForce publicly contradicted Co-op's official breach narrative by contacting the BBC directly. Ransomware actors increasingly leverage regulatory reporting deadlines as pressure tools. Cl0p's mass-publishing of 43 victims in 24 hours was designed

				to overwhelm disclosure processes. However, systematic "vendor pre-vetting by attackers based on disclosure reputation" has not been explicitly documented in Q1.

Sources and Methodology

This assessment cross-references Santosh Pandit's 2026 CTI Forecast Parts 1 and 2 (December 2025, CC BY 4.0) against Q1 2026 incident data and analysis from the following sources: Bitdefender Business Insights, Cyble Knowledge Hub, BlackFog State of Ransomware 2026, Check Point Research, Breached.company, Purple Ops, SharkStriker, Bitsight Underground, PKWARE, Anthropic (AI-orchestrated attack disclosure), CoinDesk, SecurityWeek, SC Media, Dark Reading, Barracuda, Group-IB, Hoxhunt, Vectra AI, Stellar Cyber, LevelBlue, Picus Security, the Congressional Research Service, ProArch, and Zero Networks.

All claims are sourced from publicly available reports published between January and March 2026. Where evidence originates from 2025 incidents (notably the Anthropic-disclosed GTG-1002 agentic AI attack from September 2025), it has been clearly qualified as a foundational precedent rather than direct Q1 2026 confirmation.

Assessment criteria strictly follow: Confirmed requires multiple independent sources documenting incidents matching the prediction within Q1 2026. Partial indicates that some elements are validated but others await evidence or the primary supporting evidence predates Q1 2026. Emerging indicates early indicators consistent with the prediction but insufficient for full confirmation. Not Yet indicates no supporting evidence found in Q1 2026.