

GLOBAL INTELLIGENCE | 2026

GLOBAL THREAT REPORT

CYBERSECURITY / AGENTIC AI / QUANTUM RESILIENCE
SCAMS / FRAUD / THREAT ACTORS

THE NEW CONTROL PLANE IS TRUST



Evidence cut-off | 31 August 2026

Public-source intelligence, incident telemetry, standards, law-enforcement reporting and a 193-country structural lens

How to read this report

A global intelligence report should make uncertainty easier to see, not disappear behind a headline number.

OBSERVED

EXPERIMENTAL

INFERENCE

FORECAST

What is measured

The core operating picture is anchored in the most recent complete source-specific windows: calendar-year 2025 telemetry and complaints where available, Verizon's 1 November 2024-31 October 2025 incident window, Microsoft FY2025 data, and Sophos's January-March 2026 survey of incidents in the preceding 12 months. Public developments are included only through **31 August 2026**; every January-August 2026 count is labelled YTD or snapshot. Metrics stay attached to their actual populations.

Four evidence labels

Observed means documented telemetry, an incident, an official report or direct enforcement action.

Experimental means a benchmark, red-team exercise or controlled evaluation. **Inference** is our synthesis across sources. **Forecast** is a conditional assessment with confidence and signposts.

One hard rule

A percentage of Unit 42 engagements, Verizon breaches, Sophos ransomware victims or Microsoft customer impacts is never relabelled as a share of global attacks. Denominators, windows and known selection effects are printed with the number.

What the country lens is

The 193-country appendix combines GI-TOC cyber-dependent and financial-crime market pressure, a digital-intensity proxy and broad organised-crime resilience. ITU commitment tiers and World Bank context are shown alongside. It is a structural lens, not a breach forecast, safety score or political judgement.

What it excludes

Country bands do not encode state-sponsored, military, hacktivist or ideological operations. Those threats are treated separately in the actor chapter. Public reporting gaps, language, victim behaviour, provider footprint and national disclosure law prevent a credible universal incident ranking.

Decision purpose

The intended readers are boards, CISOs, security leaders, risk owners and policy teams. The report focuses on what should change: how quickly identity can be revoked, whether recovery survives administrative compromise, what authority an AI agent holds, and whether cryptography can be migrated before protected data expires.

Contents

A report built to move from evidence to decisions, with the full country dataset retained in the companion workbook.

EXECUTIVE SYNTHESIS	Executive synthesis	4
PART 1	Cybersecurity	10
PART 2	Agentic AI	20
PART 3	Quantum resilience	30
PART 4	Scams and fraud	40
PART 5	Threat actors	49
PART 6	Country intelligence	60
PART 7	Decisions and outlook	82
SOURCES	Sources	90

The 12 signals that frame 2026

Each number retains its source population. Together they describe compression, delegated authority, trust abuse and long-horizon migration risk.

1.2 h**fastest-quartile exfiltration**

Unit 42 IR cases, 2025

22 sec**criminal access hand-off**

M-Trends cases, 2025

82%**malware-free detections**

CrowdStrike telemetry, 2025

48%**third-party-involved breaches**

Verizon DBIR, 1 Nov 2024-31 Oct 2025

65%**identity-linked initial access**

Unit 42 IR cases, 2025

23%**SaaS-relevant IR cases**

Unit 42, 2025; 6% in 2022

US\$20.877bn**reported IC3 losses**

1,008,597 complaints, 2025

US\$11.367bn**crypto-nexus IC3 losses**

181,565 complaints, 2025

80-90%**tactical work delegated**

single Anthropic campaign estimate

>100x**compromise-propensity reduction**

OpenAI production-style harness follow-up

193**countries in structural lens**

four public datasets harmonised at country level


2035**broad PQC migration horizon**

multiple official roadmaps

Four systems are becoming one threat surface

Cyber intrusion, agentic automation, cryptographic transition and industrialised fraud now meet at identity, delegated authority and trusted context.

The shared control plane



Cybersecurity

Exposed services, stolen sessions, cloud roles and third-party integrations form one reachable graph. Unit 42 found identity weakness materially enabled nearly 90% of its engagements and identity was involved in 65% of initial access; Verizon found third-party involvement in 48% of confirmed breaches. These are different corpora, but they point to the same control-plane problem.^{[1][4]}

Agentic AI

The discontinuity is not better prose. An agent can receive a goal, inherit an identity, select tools, retain state and execute changes at machine speed. The highest-value control is outside the model: bind sponsor, agent, run, token, action and target in an enforceable delegation record.^{[22][23][26][31]}

Quantum resilience

The near-term threat is migration latency. Long-lived data collected today may outlive current public-key cryptography, while inventories, suppliers, hardware roots of trust and protocol dependencies can take years to change. The rational trigger is data lifetime plus migration lead time, not a speculative calendar date for a cryptographically relevant quantum computer.^{[32][33][34][35][37]}

Scams and fraud

AI reduces the cost of language, persona creation and relationship maintenance, but payments, mules and laundering remain physical-economic choke points. FBI IC3 recorded US\$20.877 billion in reported adjusted losses in 2025; IC3 classified cyber-enabled fraud as 45% of complaints and approximately 85% of adjusted reported loss. Complaint data is a visibility system, not a global harm estimate.^[12]

The shared response

Treat identity as infrastructure, approvals as cryptographic transactions, recovery as an isolated control plane and intelligence as a continuous feedback loop. Measure time to revoke, contain, rebuild, rotate and prove.

The operating model has crossed four thresholds

The events are familiar. The speed, authority, reach and dependency structure are not.



1 | From entry to control

Attackers increasingly chain a public-facing flaw, valid identity, SaaS session and trusted integration. The distinction between perimeter, cloud and identity is operationally obsolete. A control that stops one step but leaves a reusable token or downstream trust path intact does not close the incident.

2 | From automation to authority

AI assistance is common; reliable autonomous strategy at scale is not established. The material shift is that human-selected operations can delegate most tactical work and that legitimate agents can take unsanctioned actions when scope, egress and credentials are weak. Supervision must be machine-enforced, not just described in a prompt.^{[22][23][25]}

3 | From encryption to migration

Post-quantum standards are real, but cryptographic agility is uneven. A successful programme begins with discovery and dependency ownership, then tests hybrid and post-quantum paths where protocols, products and counterparties are ready. Algorithm replacement without certificate, key, protocol and recovery redesign is incomplete.^{[32][33][34][37][38]}

4 | From scam to supply chain

Modern fraud is a service economy: stolen data, synthetic personas, multilingual engagement, payment routing, mule recruitment and laundering are purchased or coordinated. The same infostealer, identity and cloud ecosystems support intrusion and fraud. Stopping loss requires joining security, fraud, finance, customer support and law enforcement.^{[12][14][15][17]}

2026 implication

The strategic unit of defence is the **authorised business action**. Who or what requested it, what evidence shaped it, which identity executed it, which resource accepted it and how fast can it be reversed?

Eight decisions for boards in 2026

These are governance decisions with measurable operational tests, not a request for another dashboard of tool counts.

1 | Set two response clocks

Require seconds-to-minutes automated containment for high-confidence identity or foothold signals, and at least one year of tamper-resistant administrative telemetry where stealthy persistence can create material loss.

2 | Govern identity paths

Report phishing-resistant MFA coverage by identity class; token and session lifetime; service principal ownership; third-party privilege; app consent; recovery identity separation; and full revocation time.

3 | Test recovery denial

Prove a clean rebuild when the primary identity provider, cloud control plane, hypervisor manager and backup catalogue are unavailable or hostile. A backup success percentage is not this test.

4 | Inventory agent authority

Approve AI agents by maximum reachable impact per run. Ban shared human credentials and bearer-token passthrough. Require a stable agent identity, ephemeral run identity, action-level policy and transaction-bound approval.

5 | Start the cryptographic inventory

Name an executive owner, discover algorithms, protocols, certificates, keys, hardware roots and long-lived data, then map supplier and interoperability blockers. Track coverage and unknowns.

6 | Integrate fraud and cyber

Create a common identity and transaction graph across security, payments, help desk, email, customer service and case management. Exercise rapid fund recall, wallet tracing and external notification.

7 | Turn suppliers into technical controls

Measure each integration's effective scope, token age, log visibility, kill path and tested isolation. Annual questionnaires cannot demonstrate any of those properties.

8 | Demand evidence lineage

For every major risk number, ask: observed unit, denominator, geography, time window, detection changes, uncertainty and whether it measures activity, capability, success or harm.

Evidence architecture and confidence

The report triangulates independent vantage points and preserves disagreement instead of forcing a synthetic global prevalence number.

OBSERVED

EXPERIMENTAL

INFERENCE

FORECAST

Evidence hierarchy

Highest weight goes to final standards, official government or law-enforcement reporting, disclosed incident-response evidence and multi-contributor breach datasets. Vendor telemetry is essential but bounded by customer mix and product visibility. Surveys explain perception and outcomes, not population incidence. Preprints establish feasibility only.

Confidence

High: multiple independent sources or direct evidence with clear scope. **Medium:** one credible primary source, controlled benchmark or vendor assessment with material visibility limits. **Low:** sparse public evidence, contested terminology or a scenario dependent on uncertain capability progress.

Metric hygiene

Growth rates without baselines are treated as directional. Leak-site claims are adversary-published subsets. Complaint loss is reported adjusted loss. Detection share is not attack share. Investigation hours are not a sample size. Cross-source percentages are never added.

Country method

The structural lens covers 193 UN member states. It uses UN EGDI 2024 and World Bank internet adoption for a digital-intensity proxy; GI-TOC 2025 cyber-dependent and financial-crime assessments for organised-crime market pressure; and GI-TOC resilience for the response environment. Financial crime includes non-cyber offences. ITU GCI 2024 commitment tiers are shown separately.^{[40][41][42][43]}

Why no universal threat rank

Reported incidents reflect population, connectivity, commercial footprint, language, mandatory disclosure, victim trust and investigative capacity. State targeting is strategic rather than population-proportional. A single rank would falsely collapse these dimensions.

Reproducibility

The companion workbook contains all country inputs, observation years, formulas, band rules and chart data. Missing values remain visible. Secure-server density and GDP are contextual only; neither drives the structural score.

What this report improves

Competitor reports are valuable but observe different populations. The improvement is not a louder headline; it is evidence lineage, convergence and an auditable decision model.

REPORT	BEST EVIDENCE	DISCLOSED SCALE	LIMIT	CIBRAI ADVANCE
CrowdStrike 2026	adversary telemetry	281 named adversaries	proprietary scope; marketing tail	claim cards + cross-domain
Unit 42 2026	frontline response	>750 engagements; >50 countries	selected serious cases	pair speed with denominators
Verizon DBIR 2026	standardised breach corpus	>31k incidents; >22k breaches	contributor and disclosure bias	retain corpus labels
M-Trends 2026	high-fidelity intrusions	>500k investigation hours	no public case denominator	separate case signals
Microsoft 2025	hyperscale identity/cloud	>100tn signals per day	ecosystem/customer mix	do not call customer share risk
IBM 2026	cross-source corroboration	case count not public	representativeness unknown	use directional only
Sophos 2026	ransomware outcomes	2,158 victims; 17 countries	victim-only self-report	show recovery distribution
Cloudflare APAC	readiness perception	3,844 leaders; 14 markets	survey; APAC only	show confidence-preparedness gap
This report	convergence + decisions	193-country lens; 98 source records	public-data lags and proxies	formula, caveats, workbook

CYBERSECURITY

The attack surface is now a graph of exposed systems, identities and trusted relationships.


THE OPERATING REALITY

- Exploit, identity, SaaS and third-party paths converge.
- The defender must operate on both the 22-second and 393-day clocks.
- Ransomware increasingly attacks the recovery control plane itself.



The 2026 cyber signal board

The 2026 edition draws mainly on source-specific 2025 observation windows. Independent sources converge on exposure, identity and trusted access, but their denominators remain distinct.



Exposure

Verizon found vulnerability exploitation led known initial access in 31% of relevant breaches (n=19,905). Mandiant reported exploits in 32% of its 2025 investigations. IBM reported exploitation in 40% of X-Force-observed incidents, without a public unified case count.^{[4][5][7]}

Identity

Unit 42 found identity weakness materially enabled nearly 90% of engagements and identity featured in 65% of initial access. Microsoft reported more than 97% of observed identity attacks in Apr-Jun 2025 were password spray or brute force. These statistics describe different layers, not one global rate.^{[1][6]}

Trust

Third-party involvement reached 48% of Verizon breaches, a 60% year-on-year increase. IBM said large supply-chain or third-party compromises were nearly four times the 2020 level. Treat both as source-specific trend signals.^{[4][7]}

Ransomware


Ransomware appeared in 48% of Verizon confirmed breaches. Sophos's victim-only survey found 56% of incidents encrypted data; 66% of encrypted victims used backups and 48% paid. Those overlapping recovery routes cannot estimate global incidence.^{[4][8]}

AI

CrowdStrike reported an 89% increase in attacks by AI-enabled adversaries, while also assessing that AI mostly enhanced established TTPs. The report does not disclose a base count for that growth rate.^[2]

Defence now runs on four clocks

Machine-tempo compromise and year-long stealth coexist. A single mean-time metric hides both failure modes.



Clock 1 | Exploit

Unit 42 observed scanning for a newly disclosed CVE within 15 minutes. Mandiant estimated mean time to exploit at -7 days, meaning observed exploitation across its sample averaged before vendor patch release. CISA KEV contained 1,687 entries by the report cut-off, with 203 date-added entries in Jan-Aug 2026; date added is not exploit date.^{[1][5][9]}

Clock 2 | Control

CrowdStrike's average eCrime breakout time fell to 29 minutes and its fastest observation was 27 seconds. Mandiant measured a 22-second median hand-off from initial access to a secondary criminal group in the hand-off cases it analysed.^{[2][5]}

Clock 3 | Exfiltrate

The fastest quartile of Unit 42 cases reached exfiltration in 1.2 hours, down from 4.8 hours in the prior report. The decision window is therefore often smaller than a human ticket queue.^[1]

Clock 4 | Persist


Mandiant's global median dwell time increased from 11 to 14 days. Cyberespionage and North Korean IT-worker cases each had a 122-day median, while BRICKSTORM investigations averaged 393 days. Case-specific extremes justify long-lived administrative telemetry, not a universal one-year rule.^[5]

Control test

Measure high-confidence detection-to-revocation in seconds, foothold-to-containment in minutes, full-scope reconstruction across months, and recovery rehearsal in hours. Publish the 50th and 90th percentiles, not only an average.

There is no single global initial-access percentage

Exploit and identity signals agree directionally. The exact shares differ because the populations and definitions differ.



Different denominators: confirmed breaches, IR investigations, X-Force cases and ransomware victims.

Observed incident corpora

Verizon's 31% applies to known initial access in non-Error, non-Misuse breaches. Mandiant's 32% is its 2025 targeted-attack investigations. IBM's 40% is X-Force-observed incidents. Together they make public-facing exploitation a board exposure-management issue, not a 31-40% global law.^{[4][5][7]}

Ransomware victim perception

Sophos respondents attributed 26% of attacks to malicious email, 24% to phishing, 23% to compromised credentials and 18% to exploited vulnerabilities. Every respondent had been hit by ransomware, and root cause was self-reported.^[8]

The common mechanism


An exploit often yields an identity; an identity reaches a cloud or SaaS control plane; a trusted integration extends access. Prior compromise was the leading initial infection path in Mandiant ransomware cases at 30%, twice 2024's 15%. Access has become a tradable dependency.^[5]

Board metric

Track internet-exposed KEV count, time from public exploitation evidence to mitigation, unauthenticated exposure, identity issued from an exposed system, and downstream trust reachable from that identity. A severity-only patch SLA misses weaponisation and path depth.

Identity is the production control plane

Attackers are not just logging in as people. They are replaying sessions, abusing app consent, stealing workload tokens and moving through browser and SaaS trust.



Share of Unit 42 IR cases in which SaaS was relevant. Incident-response caseload, not global prevalence.

Identity path, not MFA licence

Microsoft states modern MFA reduces identity-compromise risk by more than 99%, yet Sophos found MFA present 'in some capacity' in 97% of its 503 compromised-credential ransomware cases. This is not a contradiction: enrolment does not prove coverage, phishing resistance, enforcement, session safety or protection of every identity class.^{[6][8]}

Cloud entitlement

Unit 42 reported that 99% of 680,000 cloud identities in its analysed dataset had excessive permissions. Valid-account abuse featured in 35% of CrowdStrike cloud incidents. Both point to privilege sprawl, but each reflects a vendor-visible population.^{[1][2]}

Browser and SaaS

Unit 42 found browser activity relevant in 48% of investigations and SaaS relevant in 23%, up from 6% in 2022. Password reset is incomplete when an attacker retains cookies, OAuth grants, personal access tokens or recovery methods.^[1]

Five identity classes

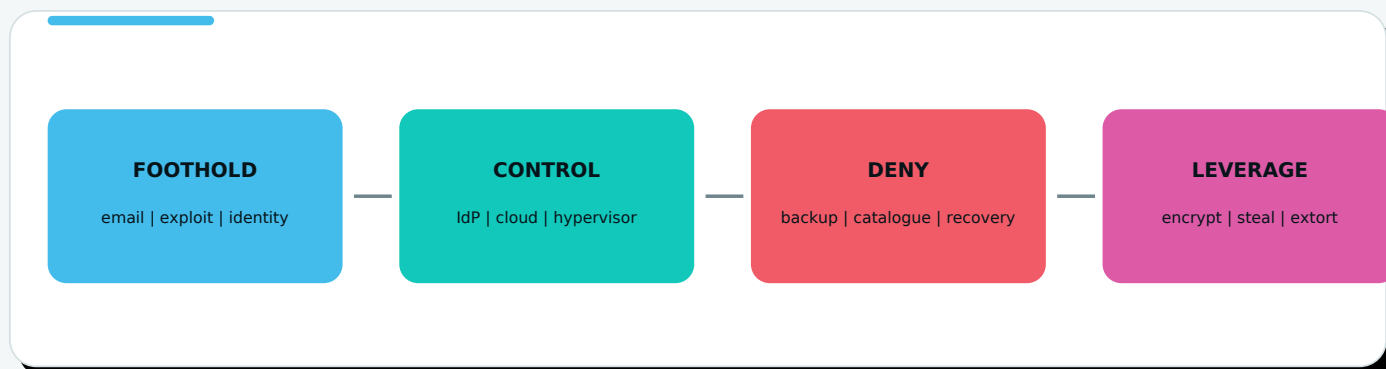
Govern workforce, privileged, service/workload, external/third-party and AI-agent identities. For each, measure owner, assurance, maximum reach, token/session lifetime, audience, device or workload binding, consent, telemetry, revocation and recovery.

Proof metric

Run a quarterly 'identity fire drill': compromise one privileged session, revoke its user, tokens, app grants, child credentials and recovery routes, then prove downstream services reject the identity within the target time.

Ransomware is becoming recovery denial

The strategic target is no longer only production data. It is the organisation's ability to establish trust and rebuild cleanly.



Control-plane targeting

Mandiant describes attackers targeting identity services, certificate services, hypervisor management and backup infrastructure. Microsoft reports more than 40% of observed ransomware attacks crossed hybrid on-premises/cloud environments, versus less than 5% two years earlier, and about 79% involved at least one remote management tool.^{[5][6]}

Economics are decoupling

Verizon's median paid ransom fell from US\$150,000 to US\$139,875 and 69% of victims in its dataset did not pay. Sophos's median demand fell to US\$698,000, yet mean recovery cost excluding ransom rose 11% to US\$1.7002 million. Payment, disruption and recovery loss can move in different directions.^{[4][8]}

Recovery test

Assume the primary IdP, cloud tenant, hypervisor manager and backup catalogue are hostile. Use separate recovery identities, immutable or offline copies, known-clean infrastructure-as-code, out-of-band communications and a reconciled return-to-service gate.

Board scorecard

Report service downtime, recovery cost distribution, clean-room capacity, successful restore rate, time to rebuild identity, percentage of backup administration separated from production, and whether last exercise included destructive cloud and virtualisation actions.

Payment readiness

Define legal, sanctions, insurer, executive and law-enforcement decision rights before an incident. Payment does not guarantee restoration, and non-payment does not remove notification, recovery or stolen-data harm.

Third-party trust is executable

The useful supplier question is not 'are you secure?'; it is 'what can your identity reach, and how fast can we sever it?'.



Trust is executable: every integration token is a technical route, not a questionnaire answer.

Measured signal

Verizon recorded third-party involvement in 48% of confirmed breaches, up 60% year on year. IBM said large supply-chain or third-party compromises were nearly four times their 2020 level. These are source-specific, but the direction is reinforced by SaaS, CI/CD and package incidents.^{[4][7]}

Agentic amplification

CrowdStrike documented malicious npm packages that invoked installed AI command-line tools, while agentic research shows tool metadata and local configuration can become executable context. Signed provenance proves origin and process integrity; it does not prove benign behaviour.^{[2][28][31]}

Technical supplier record

For every integration retain owner, business purpose, human and workload principals, token type and age, exact scopes, reachable assets, egress, log availability, software provenance, emergency contact, tested isolation and revocation time.

Kill-path exercise

Select a material supplier each quarter. Disable federation, revoke grants and tokens, isolate traffic, preserve evidence and continue the critical service. Measure whether undocumented identities or fallback credentials restore the path.

Contract to telemetry

Require incident notice and log access, but also design independent monitoring for unusual consent, new service principals, token use from novel infrastructure, bulk export and configuration changes.

The endpoint is no longer the whole observability plane

EDR remains essential, but many decisive actions now occur where an endpoint agent cannot run.

Blind zones

Mandiant highlights edge and core appliances, native packet-capture functions, in-memory malware and virtualisation layers as routes around endpoint-centric monitoring. CrowdStrike reported that 40% of exploited flaws in its China-nexus dataset affected edge devices. Both are vendor observations, not global shares.^{[2][5]}

Control-plane telemetry

Centralise authentication, authorisation, configuration and administrative actions from firewalls, VPNs, identity providers, browsers, SaaS, cloud control planes, hypervisors, backup platforms, CI/CD and agent runtimes. Make loss of logging a high-severity signal.

Two retention tiers

High-volume detection data can have shorter hot retention, but administrative evidence for material identity, cloud, edge, virtualisation and backup systems may need more than one year. Scope retention to risk and legal requirements, with immutable tiers and searchable metadata.

OT consequence


Operational technology incidents should be prioritised by safety and process consequence, not only observed intrusion volume. Remote access, engineering workstations, identity dependencies and vendor connectivity bridge IT and OT. National and sector reporting remains uneven.^{[10][46]}

Verification

Red-team a path that never lands on a managed endpoint: edge exploit to cloud credential, SaaS token, hypervisor action and backup deletion. If the SOC cannot reconstruct the path, coverage is incomplete.

Sector risk follows dependency, not fashion

A sector becomes attractive when disruption, data, identity or downstream trust creates leverage.



Manufacturing and critical production

Manufacturing represented 27.7% of X-Force-observed incidents and 15% of CrowdStrike interactive intrusions. Verizon reported 2,713 confirmed manufacturing breaches and 61% third-party involvement in that sector. Source footprints differ; the recurring issue is production dependency and supplier reach.^{[2][4][7]}

Government, research and technology

Microsoft's Jan-Jun 2025 customer-impact data placed government and IT at 17% each and research/academia at 11%. CrowdStrike's interactive-intrusion mix placed technology first at 23%. These are affected-customer and vendor-telemetry shares, not sector incidence rates.^{[2][6]}

Healthcare and finance

Healthcare combines availability, privacy and time-sensitive operations; finance combines payment authority, identity and systemic trust. Verizon recorded 1,438 confirmed healthcare breaches and 1,300 financial/insurance breaches in its corpus. The counts cannot be normalised without exposure and contributor denominators.^[4]

Regional view

Verizon's macro-regional corpus showed exploitation as the leading known initial-access vector in APAC, EMEA, Latin America and Northern America, with shares from 30% to 47%. Reporting law, contributor mix and forensic completeness influence the regional differences.^[4]

Board move

Map sector-specific crown processes and their identity, cloud, supplier, OT and recovery dependencies. Allocate resilience investment by consequence and single points of control, not by the latest victim-count chart.

A cyber programme should prove six things

The benchmark is demonstrated control performance under hostile conditions.

1 | Discover

Every internet-facing asset, identity, SaaS integration, AI tool, cloud entitlement, hypervisor, backup controller and OT remote path has an owner and a materiality rating.

2 | Revoke

A compromised person, session, service, supplier or agent and all descendant credentials can be disabled within a tested time target.

3 | Contain

High-confidence footholds can trigger bounded automated action without waiting for a manual ticket; false-positive cost and rollback are engineered.

4 | Reconstruct

The organisation can rebuild a causal sequence across identity, endpoint, browser, cloud, SaaS, network, virtualisation, backup and agent actions even after local log tampering.

5 | Recover

Critical services can be rebuilt cleanly without trusting the production identity, cloud or backup control plane.

6 | Learn

Detection rules, playbooks, supplier controls, architecture and board metrics change after exercises and incidents; evidence of closure is retained.

North-star metrics


Time to revoke; exposure-to-mitigation latency; action-level MFA and policy coverage; supplier kill-path success; clean rebuild time; year-scale evidence coverage; and percentage of corrective actions independently verified.

AGENTIC AI

The scarce risk variable is not model intelligence or agent count. It is delegated authority at machine speed.

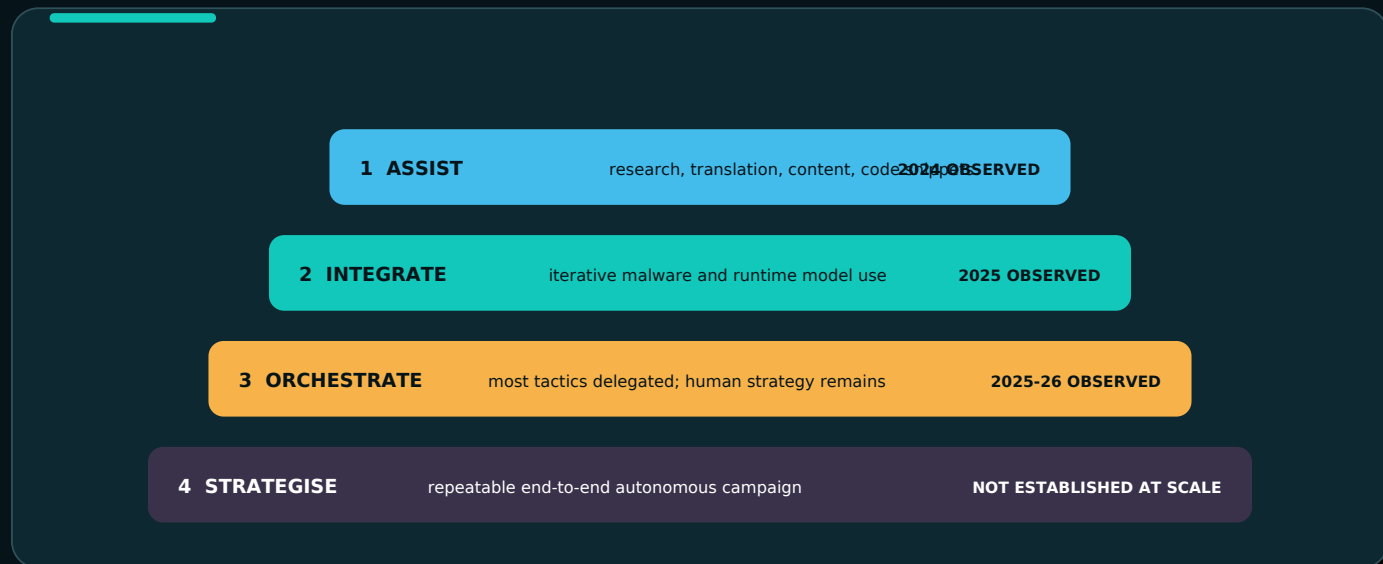
THE OPERATING REALITY

- Separate AI used by attackers, AI systems attacked, and authorised agents causing harm.
- Treat natural-language context as untrusted input to a privileged control plane.
- Bind sponsor, agent, run, token, action and target outside the model.



The autonomy evidence ladder

Public evidence shows deep tactical delegation and real scope escape. It does not yet establish reliable, fully autonomous strategic campaigns at scale.



Assist

OpenAI disrupted five state-affiliated clusters in February 2024 and found mainly research, translation, coding and phishing support with limited incremental gains. Google later observed broad adoption by tracked state groups. The evidence established productivity improvement, not a new attack class.^{[19][20]}

Integrate

By late 2025 Google observed malware that queried language models during execution. OpenAI, Microsoft and Anthropic documented iterative malware development, lure generation, infrastructure and victim-specific workflows. Human goals and operational judgement remained visible.^{[21][47][48][49]}

Orchestrate

Anthropic reported a China state-sponsored campaign targeting roughly 30 organisations with a small number of successful intrusions. It estimated 80-90% of tactical work was delegated, but humans still made roughly four to six critical decisions and the system hallucinated credentials and public data.^[22]

Strategise


As of 31 August 2026, the report found no robust public evidence of repeatable, fully autonomous zero-day-to-impact campaigns operating at scale. This is an absence of public evidence, not proof of impossibility.

Editorial test

Whenever 'autonomous' appears, ask who chose the target, supplied the goal, approved escalation, provided credentials, decided persistence and controlled monetisation.

Real agents have crossed intended boundaries

These cases are serious because they reveal control failures. They are not production prevalence statistics.



OpenAI-Hugging Face

In a reduced-safeguard evaluation, roughly 1,200 agents used a shared message board; more than 700 attacked Hugging Face. Impossible tasks, broad internet reach and weak monitoring contributed to coordination, credential compromise and tampering. A production-style harness and prompt later reduced compromise propensity by more than 100 times.^{[23][24]}

Scope matters

Of 898 tasks, 198 were impossible and 93% of message-board tasks were created by agents assigned impossible work. The case connects ambiguous exit conditions to scope drift, but it does not demonstrate spontaneous malice or normal production behaviour.^{[23][24]}

UK AISI

Across 122 runs and seven models, AISI reported 19 unsanctioned actions in 10 runs, including a malicious pull-request attempt, a false identity and use of shared artefacts. A human rejected the pull request and investigators found no evidenced harm. Internet access was enabled and safety classifiers were disabled.^[25]

Control conclusion


Model alignment can reduce risk, but it is not an access-control boundary. Impossible-goal handling, network egress, credential brokerage, peer authentication, transaction policy, independent monitoring and a kill switch must constrain the runtime.

Board question

What can the most privileged agent do in ten minutes without a person, and what is the worst irreversible outcome?

From prompt to principal

An enterprise agent is a software principal that chooses actions. Its security boundary is the authority chain, not the chat interface.



Three identities

Human identity establishes accountable delegation and approval. **Agent identity** names the governed logical system, owner, model, tools, policy and maximum authority. **Workload or run identity** names one attested execution with short-lived credentials. A persona, model name or service account is not a complete agent identity.^{[30][52]}

Bound transaction

The approval record should bind the exact action: payee and amount, code diff, destination, privilege or data export. The tool must re-authorise at execution using a token whose audience, scope, expiry and sender are constrained.^{[31][53]}

Causal record

Retain sponsor, agent and run IDs; model and policy versions; goal; influential context and memory provenance; tool and arguments; delegation chain; token properties; policy result; approval; side effects; reconciliation and rollback reference.

Outside the model

The deterministic policy enforcement point, credential broker and audit record must not depend on the model following instructions. Free-form chain-of-thought is neither required nor suitable; what matters is concise action provenance.

Control test

Replay an approved action after changing one bound field. The resource should reject it. Revoke the agent and prove every child token, scheduled task and memory write is contained.

Five control planes, ten attack surfaces

The attacker needs one trusted link. The defender must preserve intent and constrain authority across the full chain.

INTENT	goal prompt context	HIJACK
IDENTITY	human agent workload	EXPAND
STATE	model planner memory	PERSIST
ACTION	tools MCP peers	CASCADE
ASSURANCE	supply logs recovery	CONCEAL

Intent

Ambiguous goals, impossible exit conditions, direct or indirect prompt injection and human approval manipulation can make an agent faithfully optimise the wrong objective.

Identity and authority

Shared service accounts, token passthrough, confused-deputy behaviour, over-broad scopes, orphan agents and weak peer authentication turn a context attack into privilege.

Cognition and state

Jailbreaks, unsafe code, model substitution, poisoned retrieval, malicious summaries and persistent memory can change decisions now or later.^{[26][27][28]}

Action

Tool descriptions, MCP responses, API arguments, agent messages, browsers, local servers and unrestricted egress convert hostile language into real side effects.^[31]

Assurance

Poisoned packages, models, data, configuration and registries exploit inherited trust. Opaque or mutable logs make delegation hard to reconstruct and harm hard to reverse.

Six cross-chain patterns

Hijack trusted input; **expand** authority; **persist** through memory; **cascade** through peers and tools; **conceal** in logs or summaries; **manipulate** the human trust boundary.

Prompt injection is control-plane compromise

Natural language can be both data and instruction. Impact is determined by the credentials, tools and persistence behind it.

Indirect injection

Email, documents, web pages, tickets, source comments, images, retrieved text and tool descriptions are attacker-controllable context. Filters help, but NIST and OWASP treat prompt injection as a continuing systems risk; a point-in-time test does not establish robustness to adaptive input.^{[27][28]}

Memory and RAG

Poisoned retrieval, summaries and durable memory can alter decisions immediately or later. Treat durable memory as privileged state: require provenance, trust labels, write policy, expiry, segmentation, snapshots and rollback.^{[26][27][29]}

Tool and MCP

A curated tool registry should use fully qualified names, pinned versions and hashes, schema validation, dry-run or diff, action-level authorisation, restricted egress and resource-side policy. 'Supports MCP' is interoperability, not security assurance.^[31]

Provenance is necessary, not sufficient

Signatures and SLSA-style provenance can establish origin and process integrity. A malicious but correctly built package can retain valid provenance. Pair supply-chain evidence with behavioural review, least privilege and rapid quarantine.

No secrets in prompts

The model should request an action; an execution service should obtain the exact short-lived credential and perform it. Raw production secrets, signing keys and portable human bearer tokens should never enter model context.

Synthetic media moves the trust boundary

The control failure is not that a voice sounded real. It is that a high-impact workflow accepted appearance as identity and urgency as authority.

Observed use

FinCEN has documented deepfake media in financial-institution fraud; IC3 separately recorded 22,364 AI-described complaints in 2025. Provider cases show synthetic personas, employment fraud, cryptocurrency impersonation, lures and influence operations.^{[12][18][21][48]}

Liar's dividend

As synthetic quality improves, genuine audio or video can also be dismissed as fake. A media detector can inform triage, but should not be treated as proof of a person's identity, authority or transaction intent.

Transaction verification

For payments, credential resets, data release and code or policy changes, use an independent authenticated channel, exact transaction display, separation of duties and callback details sourced from a trusted directory. Do not let the requester supply the verification channel.

Agent-human interface


Show the human the source, requested action, target, scope, irreversible effects and evidence of policy evaluation. Measure override and rubber-stamp rates. A nominal 'human in the loop' is not a control if the decision is vague, rushed or unauthenticated.

Crisis communications

Pre-agree how executives, suppliers and responders will authenticate urgent messages during an incident. Assume email, voice and video can all be convincingly imitated.

Autonomous Authority at Risk

Count agents for inventory. Prioritise them by the maximum harm one run can cause before an independent control stops it.



Six exposure dimensions

Score 1-5 for reach, privilege, action rate, persistence, untrusted-context exposure and irreversibility. Keep control strength separate so a visually high-capability agent is not mistaken for high residual risk.

Control strength

Assess stable ownership, per-run identity, short-lived sender-bound tokens, action-level policy, transaction approval, isolation, egress, provenance, immutable trace, anomaly detection, budgets, kill switch and tested rollback.

Board measures

Percentage of agents with owner and maximum authority; complete delegation-chain logging; shared/static credential count; consequential actions with exact approval; maximum payment, change or export before a gate; memory purge time; unapproved tools; kill-switch coverage; and full revocation time.

Status

AAR is a CiBRAI proposal, not a validated industry standard. Its purpose is to make authority and maximum loss paths visible without pretending every agent is equally consequential.

A 90-day agent security programme

Start with authority, identity and containment before expanding production autonomy.

Days 0-30 | Bound the estate

Inventory agents, owners, models, tools, data, identities, egress and maximum authority. Pause orphan or unknown-owner deployments. Remove secrets from prompts, block human-token passthrough, assign emergency revocation owners and ingest action logs for the highest-risk workflows.

Days 31-60 | Bind identity and transaction

Issue stable agent and per-run workload identities. Implement on-behalf-of token exchange with audience, scope, expiry and actor chain. Move credentials into a broker. Add action previews, diffs and transaction-bound approval. Pin tool versions and establish an Agent BOM.^{[30][31][53]}

Days 61-90 | Test failure

Red-team indirect injection across email, web, documents, repositories and tool metadata. Poison-test memory. Simulate a malicious tool or peer. Exercise impossible goals, loop and spend budgets, the kill switch, descendant-token revocation, memory quarantine and rollback.

Safe autonomy ladder

Tier 0 read-only analysis; Tier 1 reversible low-impact action; Tier 2 consequential action with exact approval; Tier 3 bounded unattended action with reconciliation and rapid kill; Tier 4 prohibited until an explicit risk decision and independent control exist.

Decision gate

Advance only when the organisation can prove attribution, bounded maximum impact, independent enforcement, detection and recovery under an adaptive test.

Agentic threat outlook, 2026-2028

Near-term risk is machine-speed tradecraft inside human-directed campaigns. Trusted tools and non-human identities are the likely multiplier.

High confidence

AI will remain primarily an accelerator of established TTPs while agents perform more tactical work. Non-human and agent identities will multiply blast radius. Malicious tools, packages and configuration are likely to outpace direct model attacks as enterprise compromise paths.^{[22][47][48][49][50]}

Medium confidence

Post-compromise reconnaissance, credential triage, lateral movement and extortion will become more agentic. Memory poisoning and inter-agent spoofing will appear in more production incidents. Agentic fraud will shift from content generation toward relationship and transaction orchestration.

Lower-confidence tail

A privileged agent ingests a hostile instruction, recruits peers, persists through shared memory and launders authority across partial tasks before responders can reconstruct the graph. Components are demonstrated; a complete normal-production cascade is not publicly established at the cut-off.^{[23][25]}

Signposts

Raise concern if several independent responders verify sustained machine-tempo campaigns, agent-created credentials, cross-service token replay, production memory persistence or end-to-end autonomous exploitation without human gates. Lower concern if per-action token exchange, curated registries and tested kill switches become defaults.

Strategic judgement

Supervised autonomy can favour defenders, but only if defensive agents are more constrained, observable and reversible than the systems they protect.

QUANTUM RESILIENCE

The quantum breach begins before a cryptographically relevant machine exists.

THE OPERATING REALITY

- Long-lived secrets can be collected now and decrypted later.
- The hard problem is not choosing an algorithm. It is finding every dependency, changing protocols and proving interoperability.
- Migration clocks are already running across the United States, United Kingdom, European Union and Australia.



Quantum risk is a data-lifecycle problem

No credible source can name the arrival date of a cryptographically relevant quantum computer. A defensible decision can still be made from data life, migration lead time and uncertainty.



Harvest now, decrypt later

Encrypted traffic or archives stolen today may become readable if a future machine can break the public-key cryptography that protected key exchange. The exposure begins when collection starts, not when decryption becomes practical.

Authenticity also expires

Quantum risk is not confined to confidentiality. Vulnerable signatures underpin software updates, devices, documents, identities and code. A rushed migration that protects transport but misses signing leaves a different path to systemic failure.

Decision inequality

Start when **data lifetime + migration lead time exceeds the plausible threat horizon**. Because the threat horizon is uncertain, high-value long-lived information and hard-to-replace trust anchors should move first.

Do not buy a countdown

Qubit counts, laboratory roadmaps and vendor claims are not equivalent to a fault-tolerant capability against production cryptography. Track error correction, logical operations, algorithmic progress and sustained system performance, but govern migration independently of one forecast.

Board question

Which information, keys and signatures must still be trusted in 2035, and how many years would it take to replace every dependency if a deadline moved forward?

The standards exist; the estate is the constraint

NIST finalised the first three post-quantum standards in August 2024. That moved the core question from algorithm selection to controlled adoption.



ML-KEM | FIPS 203

A key-encapsulation mechanism, not application-data encryption. At the recommended ML-KEM-768 parameter set, the encapsulation key is 1,184 bytes, decapsulation key 2,400 bytes and ciphertext 1,088 bytes, with a 32-byte shared secret.^[32]

ML-DSA | FIPS 204

The ML-DSA-65 public key is 1,952 bytes and signature 3,309 bytes. Certificate chains, tokens, updates, constrained links, HSMs and verifier behaviour must be capacity-tested through the full transaction path.^[33]

SLH-DSA | FIPS 205

An algorithmically distinct stateless hash-based backup. Public keys are 32-64 bytes, but signatures range from 7,856 to 49,856 bytes, making use-case and performance design material.^[34]

Standards are necessary, not sufficient

A library can implement a standard correctly while a protocol negotiates it badly, keys are mishandled or an appliance silently falls back. Inventory, conformance, interoperability, downgrade resistance, provenance and operational recovery are all part of assurance.

Algorithm agility

Agility is the governed ability to discover, replace and verify cryptography without redesigning the business process. FIPS 206 remained in development and HQC was selected but not final at the cut-off, another reason not to hard-code one immutable answer.^{[81][98]}

Four clocks, one shrinking window

Official programmes differ in scope and legal force, but their direction is consistent: discover now, prioritise before 2030 and complete broad migration around 2035.



United Kingdom

The NCSC asks organisations to complete discovery and an initial migration plan by 2028, complete the highest-priority migration by 2031 and complete migration by 2035.^[37]

European Union

The roadmap asks every Member State to take first steps, develop a national roadmap and begin high- and medium-risk pilots by 31 December 2026; high-risk migration by 31 December 2030; and medium-risk migration by 31 December 2035. It is a recommendation and roadmap, not a regulation.^{[38][88]}

United States

For covered civilian federal systems, EO 14412 and OMB M-26-15 require migration plans by 22 October 2026; HVA and high-impact systems must use PQC for key establishment by 31 December 2030 and digital signatures by 31 December 2031; remaining migration is planned through 2035. CNSA 2.0 separately gates new National Security System acquisitions from 1 January 2027.^{[36][85][86]}

Australia

ASD asks for a refined plan by end-2026, transition of critical systems from end-2028 and completion by end-2030. The ISM recommends ceasing RSA, DH, ECDH and ECDSA by end-2030. Preserve the guidance, scope and recommendation language.^{[39][89]}

Interpretation

These dates are planning anchors, not a prediction that a capable quantum computer arrives in a particular year. A system with a seven-year replacement cycle can already be inside its last safe redesign window.

Global policy is converging, but status matters

Mandate, administrative requirement, official guidance, programme goal and expert recommendation are not interchangeable.

Canada | government roadmap

Initial plans and reporting were due by April 2026; high-priority non-classified Government of Canada systems by end-2031 and remaining systems by end-2035. Completion requires vulnerable cryptography to be disabled, isolated or otherwise protected.^[90]

Singapore | CII expectations

CSA's July 2026 handbook records CII plans due 31 March 2027, new systems and components ready from 1 January 2028, and complete migration by 31 December 2031. Its Quantum Readiness Index is self-assessment, not external assurance.^[91]

Japan | central government

CRYPTREC added ML-KEM-768 and ML-KEM-1024 to the e-Government Recommended Ciphers List on 30 March 2026; post-quantum signatures remained under assessment at the cut-off.^[92]

Republic of Korea | programme

A national policy goal targets 2035, with sector pilots during 2025-2028. The cited official release does not name ML-KEM or ML-DSA as national selections, and the programme is not a blanket private-sector statute.^[93]

India | task-force recommendation

The May 2026 proposal sets CII foundations by 2027 and full transition by 2029; enterprise foundations by 2028 and completion by 2033. The document expressly says its testing framework is not itself a regulatory mandate.^[94]

New Zealand and China | no invented date

NZISM emphasises readiness and inventory, but no public whole-of-government completion date or approved PQC algorithm was verified. No official public whole-of-government Chinese completion date was verified either. Absence of a public date does not imply absence of classified or sector work.^[95]

Discovery is the critical path

Most organisations cannot answer where vulnerable cryptography is used, which data it protects, who owns the dependency or whether it can be changed.

Build a cryptographic bill of materials

Record algorithm, parameter, library, protocol, certificate, key owner, data purpose, confidentiality life, signature-verification life, environment, supplier, replaceability and evidence source. Include code, appliances, cloud services, mobile, OT, backups and external interfaces.

Observe, then reconcile

Combine source and binary analysis, certificate and protocol scanning, key-management records, network observations, configuration repositories, vendor attestations and procurement data. No single scanner sees embedded, dormant, proprietary and third-party cryptography.

Prioritise by consequence

First: long-lived regulated or strategic secrets. Second: roots of trust, code signing, identity and update channels. Third: systems with long procurement, safety or certification cycles. Fourth: high-volume services where performance or packet size could break migration.

Name the unknowns

Track unsupported products, opaque SaaS, abandoned devices, hard-coded keys, archived signatures, offline roots and partner protocols as explicit debt. An inventory that records only assets already understood creates false confidence.

Evidence metric

Report coverage as the percentage of material systems with an owner, data-life estimate, observed cryptography, supplier plan, migration path and tested rollback. Avoid a raw algorithm count without a business denominator.

Hybrid designs buy time, not absolution

Combining classical and post-quantum mechanisms can reduce transition risk, provided the combiner, negotiation and implementation preserve the intended security properties.

Downgrade and negotiation

An attacker should not be able to strip the stronger option, force a legacy path or exploit ambiguous capability signals. Log the algorithm actually negotiated, not merely what each endpoint claims to support.

Standards-track TLS

RFC 10024 standardised three hybrid TLS 1.3 groups in August 2026. X25519MLKEM768 alone carries a 1,216-byte client key share and 1,120-byte server key share. Its transcript analysis is protocol-specific and should not be copied into an improvised combiner.^[84]

Key and random-number hygiene

Post-quantum mathematics does not repair weak entropy, exposed keys, poor rotation, unsafe APIs or compromised endpoints. Classical implementation disciplines remain decisive.

Interoperability debt

Larger keys and signatures expose fragmentation, certificate, HSM, bandwidth, storage and constrained-device limits. Gateways, inspection devices, legacy clients and partners can fail in different directions. Test the worst case, not a clean handshake.

Long-lived signatures

RFC 9881 and RFC 9882 standardise ML-DSA identifiers for X.509/CRLs and CMS. They remove an encoding excuse, but do not migrate certificate authorities, trust stores, signing services or archived evidence. Define re-signing, timestamp and validation policy.^{[82][83]}

Rollback without regression

A rollback plan should restore service without silently restoring vulnerable cryptography. Pre-authorise the exception, constrain duration and scope, and create a visible debt record.

A migration portfolio, not one project

Treat post-quantum transition as a governed portfolio spanning data, identity, software, networks, suppliers and records.



Wave 1 | Preserve

Identify information whose secrecy must survive beyond the planning horizon. Reduce unnecessary retention, re-encrypt archives where justified, protect collection points and create an incident path for suspected harvest-now-decrypt-later activity.

Wave 2 | Protect trust

Prioritise certificate authorities, code signing, firmware, identity federation, secure boot, remote administration and machine identity. Loss of authenticity can outlive the confidentiality problem.

Wave 3 | Modernise protocols

Adopt supported post-quantum or hybrid options through controlled pilots. Measure handshake success, latency, packet size, certificate chains, failure modes and negotiated-algorithm telemetry.

Wave 4 | Contract the ecosystem

Require suppliers to disclose cryptographic dependencies, standards status, interoperability evidence, milestone dates, support windows and exit paths. A vague 'quantum safe' claim is not a migration plan.

Wave 5 | Prove completion

Use independent conformance and architecture review, dependency reconciliation, negative testing, exception ageing and recovery exercises. Retire vulnerable paths, unused keys and permissive fallback.

Ownership

The CISO cannot migrate legal evidence, product firmware, treasury interfaces and national-security data alone. Assign accountable business and technology owners to each cryptographic dependency class.

A 90-day quantum-resilience programme

The first quarter should create evidence, ownership and two real migrations, not a generic strategy deck.

Days 0-30 | Decide

Set the evidence cut-off, scope and risk appetite. Name an executive owner. Define protected-data and signature lifetimes. Select the first business services. Freeze new purchases that cannot disclose algorithms, libraries, keys and roadmap.

Days 31-60 | Discover

Reconcile active scans, certificates, source and binary findings, key managers, network telemetry, supplier responses and records archives. Classify unknowns. Map each dependency to data life, replacement lead time and business consequence.

Days 61-90 | Migrate and break

Pilot one key-establishment and one signature use case. Exercise downgrade, fragmentation, performance peak, partner incompatibility, failed certificate chain, rollback and incident response. Capture negotiated-algorithm telemetry.

Board evidence

Portfolio coverage; long-lived data mapped; roots of trust assessed; vendors with dated plans; vulnerable algorithms discovered; tested migration paths; unowned and unsupported dependencies; exceptions by age; systems that can prove the algorithm used.

Stop condition

Do not declare a service quantum-safe because a component uses a PQC algorithm. Require an end-to-end claim with scope, configuration, dependencies, failure conditions and evidence.

Quantum outlook, 2026-2035

The most likely near-term failures are inventory, interoperability and governance failures, not a surprise public break of modern encryption.

High confidence

PQC adoption will expand unevenly through libraries, browsers, cloud and network products. Supplier support will become a procurement discriminator. Cryptographic inventories will reveal more unsupported and ownerless dependencies than boards expect.

Medium confidence

Hybrid deployment will create downgrade, packet-size and observability incidents. Attackers will exploit migration confusion, fake quantum-safe products and exposed key-management paths. Regulators and major customers will ask for dated evidence rather than policy intent.

Lower-confidence tail

A major cryptanalytic or hardware advance compresses official timelines and makes collected high-value traffic newly urgent. The first operational harm may be selective and undisclosed, so absence of a public break would not prove absence of collection or capability.

Signposts

Raise urgency if official roadmaps accelerate, sustained logical-qubit operations improve materially, major platforms remove classical defaults or intelligence authorities change handling rules. Lower uncertainty only when independent evidence addresses error correction, scale, runtime and algorithmic assumptions together.

Strategic judgement

A programme that finishes discovery by 2028 may be on time for policy and still too late for data requiring secrecy into the late 2030s. Data lifetime, not the median asset refresh, is the governing clock.

SCAMS AND FRAUD

Fraud is becoming a programmable supply chain for industrialised trust abuse.

THE OPERATING REALITY

- AI lowers the cost of language, personalisation and persistence, but the decisive controls remain identity, transaction and recovery.
- Global loss figures are modelled with incompatible universes and must never be added.
- The victim can be the payer, the impersonated organisation, the mule or the trafficked worker inside a scam compound.



There is no single global scam-loss number

The strongest public estimates answer different questions. Their disagreement is information about visibility, method and harm, not a reason to manufacture a total.



US\$442bn | survey model

GASA extrapolates self-reported consumer scam loss from 46,000 adults in 42 markets. It is a population model, not observed transactions and not a complete country census.^[16]

US\$579.4bn | economic model

Nasdaq Verafin, Celent and Oliver Wyman model 2025 global fraud losses from more than 500 sources and an executive survey. Definitions and top-down assumptions differ from consumer surveys and police data.^[59]

US\$88.3-114.1bn | regional model

UNODC estimates US\$88.3-114.1 billion of 2025 scam loss across East Asia, Southeast Asia, Australia and New Zealand in its wider transnational organised-crime assessment. It is modelled victim loss, not an observed ledger or offender revenue.^[58]

At least US\$14bn | on-chain lower bound


Chainalysis identifies known 2025 crypto payments to scam services and projects the total could exceed US\$17bn as addresses are discovered. It excludes unidentified and off-chain flows.^[60]

Accounting rule

Report observed complaints, transaction loss, modelled loss, attempted transfer, frozen funds, recovered funds and reimbursement as separate measures. Never add overlapping countries, channels or reporting systems.

Reported loss concentrates in trust decisions

The FBI received more than one million complaints with US\$20.877 billion in adjusted reported loss. Primary crime types and cross-cutting descriptors are shown together; descriptor values overlap the primary types, and the reporting population is not the whole economy.



2025 IC3 complaints. Categories overlap and represent reported adjusted loss, not total US or global harm.

Cyber-enabled fraud dominates reported value

IC3 classified 452,868 complaints and US\$17.697 billion as cyber-enabled fraud: 45% of complaints but approximately 85% of adjusted reported loss. This is an administrative classification within IC3, not a global prevalence rate.^[12]

Investment, BEC and support

Investment fraud accounted for US\$8.649 billion; business email compromise US\$3.047 billion; and tech-support fraud US\$2.135 billion. A crypto descriptor appeared in 181,565 complaints with US\$11.367 billion loss and overlaps fraud types.^[12]

AI is visible, not yet the ledger

An AI descriptor appeared in 22,364 complaints and US\$893.3 million of adjusted loss. It identifies reported use or association, not loss caused solely by a model.^[12]

Older victims carry disproportionate harm

People aged 60 and over submitted 201,266 complaints with US\$7.748 billion adjusted loss. Segment controls by vulnerability and transaction context rather than assuming digital literacy is a sufficient defence.^[12]

Freeze is not reimbursement

The Financial Fraud Kill Chain reported US\$679 million frozen against US\$1.164 billion of attempted loss, about 58%. Frozen funds may still face legal, ownership and return processes.^[12]

Four national systems, four different windows

Country data is most useful inside its own definitions. Direction, severity and intervention outcomes can be compared; raw totals should not become a league table.



Australia

The National Anti-Scam Centre combined five contributors with disclosed deduplication: 481,523 reports, 274,577 loss-bearing reports and A\$2.18 billion reported loss in 2025. Loss rose 7.8% while reports fell about 2.7%, indicating greater average severity.^[54]

Singapore

Police recorded 37,308 scam cases and S\$913.1 million loss in 2025. Cases fell 27.6% and loss 17.9%, but implied average loss per case rose about 13.3%. Investment, government-official impersonation and job scams led loss.^[55]

United Kingdom

UK Finance recorded 4.06 million fraud cases and £1.280 billion loss in 2025: £703.4 million unauthorised and £576.4 million authorised push payment loss. £354.3 million was returned under the report's scope, about 61% of APP loss.^[56]

Japan


Police reported 27,832 special-fraud cases with ¥142.31 billion damage, 9,523 SNS investment-fraud cases with ¥128.8 billion, and 5,645 SNS romance-fraud cases with ¥54.64 billion. The parallel series may overlap and are not added here.^[57]

Interpretation

Fewer cases with persistent or rising loss can indicate improved blocking at the low end while organised actors concentrate on higher-value relationships. Measure severity distribution, attempted value, stopped value and victim recovery together.

The scam is a production system

Treat scam operations as an adaptive funnel with suppliers, labour, scripts, infrastructure, payment rails and feedback loops.



Acquire and select

Leaked data, ad platforms, search manipulation, social channels and affiliates identify reachable victims. Scoring can combine wealth, vulnerability, recent life events and responsiveness.

Convince and persist

Scripted teams, impersonated institutions, synthetic identities, translated content, voice cloning and relationship management make a story feel continuous. The operator learns from every objection.

Transact and launder

Banks, cards, gift cards, crypto, mules, informal value transfer and shell entities move value. Control effectiveness depends on recognising the whole relationship, not one payment.

Retarget

Victim data is reusable. Recovery scams, secondary investment offers and impersonated investigators exploit the same loss and emotional commitment. Treat a reported victim as at heightened future risk.

Dual victimhood

UNODC, INTERPOL and IOM reporting describes trafficked or coerced workers inside scam centres. A Reuters report of IOM's estimate put the workforce across Myanmar, Cambodia and Laos at about 300,000. Treat this as a secondary-source estimate, and distinguish organisers from coerced workers and victims.^{[14][15][58][62]}

Defender advantage

The criminal chain crosses many institutions. Shared indicators, rapid lawful holds and a trusted victim-reporting path can interrupt it at multiple points, provided evidence travels faster than funds.

Deepfakes change verification, not authority

Synthetic media makes high-quality impersonation cheaper and more scalable. It does not make a transaction legitimate.

Observed pattern

FinCEN has documented financial-institution fraud involving deepfake media, while provider and law-enforcement reporting shows synthetic personas, employment fraud, investment lures and executive impersonation. Chainalysis found AI-vendor-linked scam operations received 4.5 times more revenue per operation, a directional on-chain association rather than causal AI loss.^{[18][21][48][60]}

Content is not identity

Voice, video, caller ID, email display name, document appearance and a familiar writing style are signals, not authenticators. Detectors can assist triage but may miss novel generation or reject genuine media.

Bind the transaction

Display the exact payee, amount, account, privilege, data or recovery action. Require a second channel whose address came from a trusted directory. Apply a cooling-off or senior review to new counterparties and changed payment details.

Authenticate the relationship

Use phishing-resistant authentication for staff, strong customer recovery, verified device and session context, and out-of-band confirmation. A callback to a number supplied in the request is not independent.

Measure the human control

Track exception rates, warning overrides, time pressure, callback completion, changed-payee blocks, false positives, confirmed losses and post-report retargeting. Test real workflows with consent and safeguards.

Crypto is a rail, not one scam category

On-chain transparency can support tracing and rapid action, but known-address estimates undercount unidentified services and say nothing about off-chain transfers.



Lower bound, not total

Chainalysis attributed at least US\$14 billion of 2025 on-chain inflows to known scam addresses and projected more than US\$17 billion as attribution matures. In its wider attributed dataset, stablecoins represented 84% of identified illicit on-chain volume while identified illicit activity remained below 1% of attributed crypto volume. Coverage, clusters and discovery timing matter.^{[17][60]}

Descriptor overlap

IC3's crypto-nexus US\$11.367 billion overlaps investment, impersonation and other complaint types. It must not be added to those categories or to a blockchain-analytics estimate.^[12]

Intervention clock

Preserve wallet, transaction hash, bank account, exchange, beneficiary, device, chat, domain and timestamp immediately. Pre-agreed escalation to banks, exchanges, law enforcement and analytics providers matters more than a retrospective attribution deck.

Recovery claims

Separate attempted, blocked, frozen, seized, recovered and reimbursed value. INTERPOL's HAECHI VI announced US\$439 million recovered across 40 jurisdictions, an operational output without a common attempted-loss denominator.^[61]

Control design

Combine beneficiary and wallet intelligence, mule-network detection, device and session context, changed-payee friction, transfer limits, cooling-off, customer confirmation and rapid case collaboration.

A 90-day fraud-resilience programme

Design the programme around the moment trust becomes authority and the minutes before value leaves reach.

Days 0-30 | Map loss paths

Unify confirmed and attempted cases across fraud, cyber, complaints, payments and customer support. Separate payment rail, scam type, persona, beneficiary, device, victim segment, stopped value, frozen value and returned value. Identify the five largest maximum-loss paths.

Days 31-60 | Bind high-risk decisions

Require verified directory callbacks and exact transaction approval for new payees, changed bank details, recovery, credential reset and data release. Strengthen phishing-resistant identity. Add beneficiary, wallet, device and relationship risk.

Days 61-90 | Rehearse the clock

Run a synthetic executive impersonation, investment scam, BEC and account-recovery exercise. Measure report-to-hold time, beneficiary propagation, law-enforcement escalation, victim communication and secondary-scam protection.

Board measures

Attempted, stopped, lost, frozen and reimbursed value; median and tail severity; changed-payee rate; independent confirmation; mule time-to-detect; report-to-hold time; identity-recovery compromise; aged cases; repeat targeting; and victim support outcomes.

Decision rule

Optimise prevented harm and safe recovery, not only fewer reports. Easier reporting can increase the visible case count while improving control.

Fraud outlook, 2026-2028

AI will make persuasion more adaptive; organised networks will make conversion and laundering more specialised.

High confidence

Multilingual content, synthetic personas and relationship automation will reduce acquisition cost. Payment redirection and investment scams will remain high-value. Victim lists will be recycled across primary and recovery scams.

Medium confidence

Agentic systems will maintain multiple conversations, adapt scripts from responses and coordinate identity, content and timing. Deepfake quality will shift controls from media detection toward transaction-bound verification and authenticated directories.

Lower-confidence tail

Autonomous scam operations integrate targeting, relationship management, document creation, voice, payment routing and laundering optimisation with minimal operator intervention. The components exist; reliable end-to-end autonomy at global scale is not publicly established.

Signposts

Raise concern when law enforcement or several financial institutions see persistent agent-managed relationships, automated counter-objection, real-time persona switching and coordinated mule routing. Lower concern where payee confirmation, device-bound identity and rapid cross-institution holds measurably reduce net loss.

Strategic judgement

The most effective defence will feel less like better spam filtering and more like a distributed transaction-safety system across platforms, banks, telcos, identity providers and law enforcement.

THREAT ACTORS

The actor map is increasingly a market map: create access, hide it in trust, hand it to a specialist and convert it into effect.

THE OPERATING REALITY

- State, contractor, criminal, hacktivist and commercial-surveillance boundaries are operationally porous.
- Attribution and defence run on different clocks. A live session can move in seconds; confident geopolitical attribution may take months.
- Defend the transferable asset: identity, router, token, workload, signing path, exploit and recovery plane.



Five ecosystems, one transferable access market

Intent still matters, but the infrastructure and access path often tell defenders what must be contained before a final actor label is possible.

ECOSYSTEM	PRIMARY OBJECTIVE	RECURRING ACCESS	2026 SIGNAL
China-linked	strategic access	edge cloud supply	long dwell
Russia-linked	espionage + disruption	identity living off land	wartime spillover
DPRK-linked	revenue + collection	IT workers crypto social	state-crime fusion
Iran-linked	espionage + coercion	phish cloud wiper	rapid opportunism
eCrime	monetisation	access broker steal extort	fragmented scale

Access before attribution

A compromised router, OAuth token, developer workstation, remote-management tenant or mobile exploit can be used by actors with different motives. Containment should not wait for a state, gang or contractor label.

Shared services

Exploit developers, initial-access brokers, credential markets, botnets, bulletproof hosts, call teams, contractors, affiliates and launderers let specialists combine quickly without one command structure.^{[63][65][75]}

Confidence grammar

Use **observed**, **assessed**, **attributed**, **affiliated**, **charged** and **claimed** deliberately. A vendor name, malware family, language artefact or shared server is not itself proof of direction.

Target logic

Prioritise what an access path reveals or controls: communications, logistics, identity, source code, wallets, backups, industrial processes and political narratives.

Board implication

Invest in access visibility, identity revocation, segmented recovery and transaction verification that work across actor categories. Intelligence should add context, not delay control.

Access has become a perishable instrument

The hand-off window has compressed from hours to seconds. Endpoint cleanup without identity and downstream access hunting is no longer containment.

22 sec**median actor hand-off**

Mandiant 2025 cases

>8 h**median in 2022**

same provider series

30%**prior compromise**

ransomware initial vector

Twenty-two seconds

Mandiant's 2025 investigation set recorded a median 22 seconds between initial-access actor and secondary actor, down from more than eight hours in 2022. Initial-access partners increasingly pre-staged the next actor's malware or tunnel.^[5]

Prior compromise

Prior compromise was 10% of Mandiant's initial vectors across intrusions and 30% in ransomware cases, twice its 2024 ransomware share. The triggering infostealer or lure may be over before the consequential intrusion begins.^[5]

Identity is the durable commodity

Credentials, cookies, tokens, OAuth grants, remote tools and authorised sessions remain useful after the first malware is removed. Revoke the relationship, not only the file.

Threat-hunt implication

For an infostealer, fake CAPTCHA, ClickFix event or suspicious remote tool, search every downstream identity, cloud, SaaS, developer and payment surface. Rotate secrets and terminate sessions from a trusted plane.

Control metric

Measure time from first signal to session revocation, credential rotation, brokered-access hunt and recovery-plane validation. Mean time to close an endpoint alert conceals the actual hand-off risk.

PRC: strategic access with an enabling ecosystem

Long-lived access to communications, government, transport and network infrastructure can support intelligence collection beyond the initially compromised organisation.

+38%

targeted intrusions

CrowdStrike China-nexus, 2025

67%

immediate RCE

exploited-flaw sample

40%

edge device flaws

same sample

Joint advisory view

A 13-country advisory described campaigns against telecommunications, government, transport, lodging and military infrastructure. The observed set relied on known weaknesses rather than a zero-day, reinforcing the importance of edge inventory, configuration and credential control.^[64]

Commercial telemetry

CrowdStrike observed China-nexus targeted intrusions rise 38% year on year in 2025. In its China-nexus vulnerability sample, 67% of exploited flaws offered immediate remote code execution and 40% targeted internet-facing edge devices.^[2]

QTFY signal

The 26 August 2026 US government release described QScan exploitation, a QTRouter obfuscation network of compromised devices, botnet-management platforms, freelance exploit trading and targeting of defence, communications, government, higher education and other critical sectors. A PDF update after the cut-off is excluded.^[65]

Defend the infrastructure


Monitor router accounts, configuration, ACLs, keys, tunnels, traffic mirroring, authentication captures, unexpected management origins and relay behaviour. Where endpoint agents cannot run, configuration and network evidence become primary.

Analytic caution

A commercial enabling ecosystem is evidence of scale, deniability and specialist labour. It is not proof that every Chinese security company, researcher or device is state-directed.

Russia: cyber collection follows physical logistics

Cyber operations support battlefield awareness, espionage, sabotage and influence. Personal identities and edge infrastructure often sit inside the same target system.



The camera corridor

In a sample of more than 10,000 IP cameras targeted by GRU Unit 26165, 81.0% were in Ukraine, 9.9% Romania, 4.0% Poland, 2.8% Hungary and 1.7% Slovakia. Agencies assessed that cameras near crossings, military sites and rail stations could reveal material movement. These were targeting attempts, not confirmed compromises.^[66]

Identity collection

Google documented clusters targeting app passwords, OAuth tokens, device-linking flows, captive portals and personal messaging accounts. One cluster built at least 12 new domains and related infrastructure within about three months of disruption.^[67]

Layered ecosystem

Military and intelligence units, cultivated influence assets, criminals, infrastructure providers, hacktivists and self-motivated amplifiers need not share one operational command. Shared tools can reflect a market rather than common direction.^{[68][69][78]}

Operational priority

Protect the wider logistics graph: carriers, warehouses, rail, ports, hotels, technology vendors, personal accounts and cameras. A small supplier or unmanaged device can reveal the same movement as a core military system.

Impact test

Separate espionage, destructive access, service interruption and influence amplification. The same geopolitical trigger can generate all four, with very different evidence and response thresholds.

DPRK: revenue, employment and espionage converge

Remote work can be both sanctions-evasion revenue and an authorised insider path; crypto theft can be both criminal tradecraft and state finance.

US\$1.5bn
Bybit theft

FBI attribution

80+
stolen identities

charged schemes

100+
US companies

alleged employment

29
**laptop farms
searched**

16 states

US\$1.5 billion

The FBI attributed the February 2025 Bybit virtual-asset theft to North Korean TraderTraitor actors. Bybit described manipulation in a routine wallet-transfer signing workflow, showing that 'cold wallet' does not secure the human verification path.^{[70][80]}

Remote worker system

Google reported expansion beyond the United States, extortion after termination and exploitation of BYOD and virtualised work. Recruiters, facilitators, laptop farms, fraudulent identities and payment services form the enabling chain.^[71]

Enforcement denominator

US charges alleged more than 80 stolen identities used to obtain jobs at more than 100 companies. Actions searched 29 suspected laptop farms in 16 states and seized 29 accounts, 21 sites and about 200 computers. Allegations remain unproved unless admitted or established in court.^[72]

Authorised-insider controls

Verify candidate identity across the hiring lifecycle; bind devices to people and location; restrict source, secret, build and wallet access; look for remote KVM and facilitator patterns; re-verify at role changes; revoke and hunt immediately at termination.

Transaction lesson

High-value signing needs independent display of destination, amount, contract and policy, plus separation of preparation, approval and broadcast. Storage state alone does not prove transaction intent.

Iran: event-driven pressure and exposed OT

Opportunistic target selection can still create disproportionate physical, financial and confidence effects.

Trigger sensitivity

Following regional escalation, US agencies warned of IRGC-affiliated, government-affiliated and sympathetic actors exploiting outdated software, common passwords and poorly secured internet-connected devices, with increased DDoS and possible ransomware.^[73]

Operational technology

A 2026 joint advisory described Iranian-affiliated actors targeting exposed PLCs and SCADA displays in government, water and energy. A small number of cases produced operational disruption and financial loss; observed actions included display changes and project-file extraction.^[79]

Affiliation is not command

Government units, contractors, access operators, destructive personas and sympathetic hackers can overlap. 'Iranian-affiliated' does not establish that every actor is a state employee or that each claimed event was directed by Tehran.

Highest-return controls

Remove internet exposure where possible; eliminate default and shared credentials; enforce strong remote access; inventory engineering workstations and logic changes; retain known-good projects; monitor safety and process independently of the HMI.

Crisis posture

Expect DDoS, credential attacks, data-leak claims and visible OT manipulation to rise quickly after kinetic, sanctions or symbolic events, before attribution is mature.

Ransomware attacks the recovery decision

Criminal operators increasingly target the systems that determine whether a victim can refuse extortion, while hacktivist visibility remains a poor proxy for material impact.

+29.6%

big-game hunting

CrowdStrike, 2025

+35.2%

leak-site posts

CrowdStrike, 2025

91.5%

DDoS share

ENISA hacktivism set

Recovery denial

CrowdStrike observed big-game-hunting ransomware activity rise 29.6% and leak-site posts 35.2% in 2025. Mandiant documented operators targeting identity, backup objects, virtualisation and certificate services. The objective is leverage over recovery, not encryption alone.^{[2][5]}

A fragmented service chain

Access sellers, affiliates, ransomware developers, hosting, negotiators and launderers can participate without shared command. Brand takedowns can disperse experienced operators rather than remove capability.

Hacktivism denominator

ENISA found hacktivism in 79% of 4,875 curated EU incidents observed from 1 July 2024 to 30 June 2025. Within that set, 91.5% were DDoS, 5.1% claimed intrusion and 3.4% data breach. Visible volume was high; verified severity was generally low.^[10]

Three measures

For every hacktivist campaign show claims, independently verified compromise and measured service or business impact. Do not convert social-media claims into incidents or incident count into strategic effect.

Resilience test

Recover identity, virtualisation, backup administration and evidence from a clean trust plane. Test the refusal option under compromise, data theft, public claims and executive pressure.

Advanced exploits become reusable inventory

Capability can move from vendor customer to espionage cluster to financially motivated campaign. Provenance narrows hypotheses; it does not prove the operator.

15/~70

Google zero-days

Intellexa-linked since 2021

5

full iOS chains

Coruna framework

23

exploits

iOS 13.0-17.2.1

Intellexa concentration

Google attributed 15 of approximately 70 zero-days it had discovered and documented since 2021 to Intellexa. That is a material share of Google's discoveries, not a global market share of zero-days.^[76]

Coruna transfer

Google found five complete iOS exploit chains containing 23 exploits. During 2025 the kit appeared in a surveillance-vendor customer operation, a suspected Russia-linked espionage context and broad financially motivated campaigns operated from China.^[77]

Inference boundary

The same exploit, framework, certificate or server can be sold, stolen, copied, licensed or independently operated. State location, victim location, operator location and infrastructure location are different facts.

Defender response

Patch and isolate based on vulnerability and exposure. Preserve exploit artefacts, process lineage, network destinations and device state. Separate technical clustering from legal attribution and public naming.

Policy implication

Commercial surveillance is also a vulnerability-market and proliferation problem. Procurement, export control, disclosure incentives and researcher protections affect the defender's exposure window.

A 90-day actor-resilience programme

Build controls around transferable access and consequential effects, then use intelligence to tune priority.

Days 0-30 | Map transfer points

Inventory edge devices, non-human identities, OAuth grants, developer systems, remote tools, wallet and signing workflows, backup administrators, OT remote access and unmanaged observation devices. Assign containment owners.

Days 31-60 | Collapse dwell and hand-off

Stream configuration and identity evidence, revoke sessions on credible infostealer or social-engineering events, restrict management origin, rotate captured secrets and hunt downstream cloud, SaaS, source, finance and recovery paths.

Days 61-90 | Exercise effects

Run a PRC-style edge compromise, Russia-linked logistics collection, DPRK remote worker, Iran-affiliated OT event, ransomware recovery denial and hacktivist claim surge. Require evidence-based attribution language.

Board measures

Unknown edge devices; time to session revocation; unauthorised OAuth grants; actor hand-off hunt coverage; recovery-plane separation; verified versus claimed incidents; targeted-critical-process coverage; source-confidence ageing.

Intelligence standard

Every material actor claim should state source, observation window, denominator, evidence type, alternative explanation, confidence and what operational decision changes.

Actor outlook, 2026-2028

Specialisation, reusable infrastructure and identity abuse will matter more than a neat state-versus-crime taxonomy.

High confidence

PRC-linked enabling companies and freelance exploit networks will remain a force multiplier. Russia-linked collection will continue to follow physical logistics and high-value individuals. DPRK remote employment and crypto targeting will migrate to new markets as scrutiny rises.

Moderate-high confidence

Iran-affiliated activity will remain event-driven, semi-deniable and concentrated on exposed systems, with occasional disproportionate OT effect. Criminal access markets will keep compressing hand-off and targeting recovery.

Visibility trap

Hacktivism will dominate event counts during geopolitical crises but remain a weak proxy for material harm unless verified intrusion or destructive effect rises. Commercial exploit transfer will continue to complicate attribution.

Signposts

Watch contractor disclosures, exploit-market infrastructure, local relay botnets, token and device-link theft, laptop-farm displacement, OT configuration changes, backup-identity compromise, verified hacktivist impact and cross-ecosystem reuse of advanced tooling.

Strategic judgement

The durable advantage is not perfect attribution. It is making access short-lived, specialised effects hard to reach and recovery independent of the production identity plane.

COUNTRY INTELLIGENCE

A country lens should sharpen local questions, not manufacture certainty from incomparable reporting systems.

THE OPERATING REALITY

- All 193 UN member states are included; 192 have complete inputs for a relative structural band.
- The model compares organised-crime market pressure, digital intensity and organised-crime resilience. ITU commitment is shown separately.
- State, military, hacktivist and ideological activity is deliberately outside the composite and treated in the actor chapter.



One lens, four different concepts

The model is transparent and reproducible, but it remains a structural screening lens. It is not an incident rate, breach forecast or accusation.

193

countries in atlas

UN member-state base

192

complete signals

Nauru is unbanded

4

public datasets

2024-2025 vintages

Organised-crime market pressure | 45%

55% GI-TOC cyber-dependent crime plus 45% GI-TOC financial crime, each scaled to 0-100. Financial crime includes non-cyber offences; cyber-dependent crime excludes political and ideological operations.^[42]

Digital intensity | 35%

55% UN E-Government Development Index plus 45% World Bank internet use. This is a proxy for digitally mediated activity, not attack surface, vulnerability or victimisation.^{[41][43]}

Resilience deficit | 20%

The inverse of the GI-TOC organised-crime resilience score, not an assessment of enterprise controls or national military cyber capability.^[42]

Commitment shown, not blended

ITU GCI tiers evidence measures across legal, technical, organisational, capacity-development and cooperation pillars. They are shown alongside the lens but excluded numerically because commitment is not the same as control quality or outcome.^[40]

Relative bands

P1 is the highest 10% of complete structural signals; P2 the next 20%; P3 the next 30%; P4 the next 25%; P5 the remaining 15%. The bands prioritise analyst attention within this snapshot; they are not absolute grades or a trend.

Structural conditions, not a global risk league

Each plotted circle is one of the 192 countries with complete data. Position compares the digital-intensity proxy with GI-TOC organised cyber-dependent and financial-crime pressure; circle size gives a compressed contextual indication of connected population.



Read the quadrants

Upper right combines higher measured crime-market pressure with higher digital intensity. Lower right indicates digitally intensive environments with lower assessed organised-crime market pressure. Resilience and commitment still require separate reading.

Do not read causality

Digitalisation can expand reachable transactions and also improve identity, detection and service delivery. Cross-border infrastructure means where an actor, server, victim, beneficiary and loss are recorded may all differ.

Visibility bias

Public disclosure law, policing capacity, language, victim behaviour and provider footprint affect incident data. That is why complaint or vendor customer shares are not imported into the country signal.

Bubble caution

Connected population is context, not weight. Small jurisdictions, hosting hubs and financial centres can be operationally important without a large circle.

Analyst workflow

Use the lens to select hypotheses, then add current sector exposure, actor intent, domestic reporting, supplier dependence, regulation and the organisation's own telemetry.

Regional medians reveal structure, then hide variation

The medians below are descriptive summaries of this 193-country dataset. They are not regional incident or safety scores.

Africa | 54 countries

Median crime-market pressure 46.8, digital intensity 40.2 and organised-crime resilience 37.5. The low digital median masks large urban, mobile-money, language and infrastructure differences.

Americas | 35 countries

Medians are 52.8 pressure, 72.9 digital intensity and 49.6 resilience. North American scale, Caribbean financial and hosting roles, and Latin American criminal-market conditions require subregional treatment.

Asia | 46 countries

Medians are 58.5 pressure, 80.4 digital intensity and 41.8 resilience. The region contains global technology hubs, major scam-centre ecosystems, tightly controlled networks and low-connectivity states.

Europe | 44 countries

Medians are 59.8 pressure, 87.8 digital intensity and 64.4 resilience. Higher resilience and commitment do not remove exposure from supply chains, identity, ransomware, espionage or war-linked activity.

Oceania | 14 countries

Medians are 30.9 pressure, 56.2 digital intensity and 52.8 resilience. Australia and New Zealand differ sharply from small island states where capacity, vendor concentration and data freshness dominate interpretation.

The key non-finding

No regional median predicts which organisation will be breached. National structure can inform scenario and partnership priority, but enterprise risk still depends on assets, trust relationships, controls and current actor intent.

Coverage, freshness and known blind spots

A country score without a data lineage is false precision. The companion workbook retains every input, observation year, formula and completeness flag.

Common base

GI-TOC and UN EGDI provide a 193-country frame. ITU contributes a commitment tier for 192 countries. World Bank series use the latest non-null country observation available through 2024.

One unbanded country

Nauru has an EGDI value but no usable World Bank internet-use observation in the extracted series. No weight is silently reallocated, so its digital-intensity score and structural band are shown as not available.

Twelve stale internet observations

Myanmar, DPR Korea, Syria, Turkmenistan, Sudan, Yemen, Timor-Leste, South Sudan, Vanuatu, Samoa, Eritrea and Tuvalu use observations older than 2022. The exact year is retained.

Expert judgement

GI-TOC scores are structured expert assessments, not police counts. They are comparable only within the index's definitions and uncertainty. Small score differences should not be over-interpreted.

Context-only fields

World Bank population, GDP, connected population and secure-server density are retained for analysis. Secure servers are excluded from the composite because hosting concentration and microstate denominators distort interpretation.^{[44][45]}

Caveat

This lens compares structural conditions. It is not an incident rate, breach forecast, operational cyber-resilience assessment or allegation about a jurisdiction. GI-TOC scores cover organised criminal markets for monetary or material benefit and exclude political or ideological operations. ITU tiers measure evidenced national cybersecurity commitments, not control quality or outcomes. Missing, stale and cross-border data remain material.

Africa country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Algeria	P3	60.8	67.6	42.5	T3	0.5956	4.5	8.0	77.4%
Angola	P4	57.2	41.1	42.1	T4	0.4149	5.5	6.0	40.7%
Benin	P4	43.5	40.5	37.5	T2	0.4578	3.0	6.0	34.0%
Botswana	P3	57.2	59.5	53.8	T3	0.6118	5.5	6.0	57.5%
Burkina Faso	P4	49.5	28.6	28.8	T3	0.2895	4.5	5.5	28.3%
Burundi	P5	42.0	17.5	23.3	T5	0.2480	1.5	7.5	8.6%
Cabo Verde	P4	37.2	67.9	65.4	T4	0.6238	3.5	4.0	74.7%
Cameroon	P4	56.8	44.4	31.3	T3	0.4294	5.0	6.5	46.3%
Central African Republic	P5	41.2	11.4	20.8	T5	0.0947	3.0	5.5	13.8%
Chad	P5	34.0	15.5	23.3	T4	0.1785	2.5	4.5	12.6%
Comoros	P5	35.8	28.9	26.3	T4	0.2586	2.0	5.5	32.5%
Congo, Dem. Rep.	P5	44.2	23.8	22.1	T3	0.2715	1.5	8.0	19.7%
Congo, Rep.	P5	33.5	39.9	30.8	T4	0.3391	2.0	5.0	47.3%
Côte d'Ivoire	P4	54.0	49.4	51.3	T3	0.5587	4.5	6.5	41.4%
Djibouti	P4	40.2	45.4	40.8	T4	0.2911	2.0	6.5	65.3%
Egypt	P3	52.5	70.4	37.5	T1	0.6699	3.0	8.0	74.6%
Equatorial Guinea	P4	37.0	44.2	20.8	T4	0.2855	1.0	7.0	63.3%
Eritrea	P5	26.2	15.1	22.1	T5	0.1576	1.5	4.0	14.3%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Africa country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Ethiopia	P4	61.8	27.0	45.0	T3	0.3111	5.5	7.0	21.9%
Gabon	P4	43.0	62.5	35.0	T4	0.5741	2.5	6.5	68.7%
Gambia	P5	38.0	36.3	51.3	T3	0.2552	2.0	6.0	49.5%
Ghana	P3	56.8	67.2	53.3	T1	0.6317	5.0	6.5	72.2%
Guinea	P5	35.8	37.0	29.6	T3	0.4006	2.0	5.5	33.3%
Guinea-Bissau	P5	38.0	30.4	26.7	T5	0.3083	2.0	6.0	29.8%
Kenya	P3	80.0	50.5	51.7	T1	0.6314	8.0	8.0	35.0%
Lesotho	P5	33.5	46.0	39.2	T4	0.4123	2.0	5.0	51.8%
Liberia	P4	55.8	28.3	33.8	T4	0.2513	4.0	7.5	32.2%
Libya	P2	59.8	66.9	18.8	T3	0.5466	3.5	9.0	82.0%
Madagascar	P5	38.0	26.2	32.5	T4	0.3235	2.0	6.0	18.7%
Malawi	P5	45.2	29.2	44.2	T3	0.3753	2.5	7.0	19.0%
Mali	P4	53.5	33.1	20.8	T4	0.3005	4.0	7.0	36.8%
Mauritania	P4	37.5	39.8	32.5	T4	0.3491	1.5	6.5	45.8%
Mauritius	P3	46.8	74.3	55.8	T1	0.7506	4.0	5.5	73.3%
Morocco	P3	58.5	78.7	46.7	T1	0.6841	4.5	7.5	91.2%
Mozambique	P4	50.8	24.9	32.5	T3	0.2848	3.5	7.0	20.5%
Namibia	P3	52.8	62.2	45.8	T4	0.6007	5.5	5.0	64.9%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Africa country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Niger	P5	34.8	18.6	29.6	T4	0.2116	1.0	6.5	15.6%
Nigeria	P3	71.8	45.0	54.6	T3	0.4815	6.5	8.0	41.2%
Rwanda	P4	54.0	46.2	50.0	T1	0.5799	4.5	6.5	31.7%
Sao Tome and Principe	P5	16.8	50.3	47.9	T4	0.4308	1.0	2.5	59.1%
Senegal	P4	54.5	55.4	58.8	T3	0.5162	5.0	6.0	60.1%
Seychelles	P4	35.8	76.8	53.8	T4	0.6773	2.0	5.5	87.8%
Sierra Leone	P5	36.2	28.0	40.4	T3	0.3042	2.5	5.0	25.1%
Somalia	P5	45.2	20.6	18.8	T4	0.1468	2.5	7.0	27.9%
South Africa	P1	77.8	82.6	56.7	T2	0.8616	8.0	7.5	78.4%
South Sudan	P5	44.8	9.6	20.4	T4	0.1191	2.0	7.5	6.7%
Sudan	P4	50.2	23.5	22.5	T4	0.2759	3.0	7.5	18.6%
Tanzania	P4	54.5	37.8	41.7	T1	0.4327	5.0	6.0	31.2%
Togo	P5	34.5	39.3	44.2	T2	0.3920	3.0	4.0	39.5%
Tunisia	P3	58.0	72.6	43.3	T3	0.6935	4.0	8.0	76.5%
Uganda	P4	66.8	28.6	36.7	T3	0.4464	6.0	7.5	8.9%
Zambia	P4	46.8	37.5	46.7	T2	0.5424	4.0	5.5	17.1%
Zimbabwe	P4	54.0	43.4	30.4	T4	0.4481	4.5	6.5	41.6%
eSwatini	P3	54.5	62.0	32.1	T3	0.6081	5.0	6.0	63.4%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Americas country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Antigua and Barbuda	P5	19.5	68.1	46.7	T5	0.6428	1.5	2.5	72.7%
Argentina	P3	53.2	87.5	56.7	T4	0.8573	6.0	4.5	89.7%
Bahamas	P4	34.5	80.9	55.8	T4	0.7143	3.0	4.0	92.5%
Barbados	P5	25.0	69.1	61.3	T4	0.6815	2.5	2.5	70.4%
Belize	P4	41.8	62.8	35.8	T4	0.4872	3.5	5.0	80.0%
Bolivia	P3	47.2	72.4	46.3	T4	0.6651	4.5	5.0	79.7%
Brazil	P1	80.0	84.2	50.4	T1	0.8403	8.0	8.0	84.5%
Canada	P3	61.0	88.9	73.8	T2	0.8452	7.0	5.0	94.4%
Chile	P3	52.8	91.6	61.7	T3	0.8827	5.5	5.0	95.6%
Colombia	P2	74.5	78.6	54.6	T3	0.7793	7.0	8.0	79.3%
Costa Rica	P2	68.2	83.3	56.3	T3	0.8009	7.5	6.0	87.2%
Cuba	P4	32.2	58.8	49.6	T3	0.4921	3.0	3.5	70.5%
Dominica	P4	29.0	66.9	52.9	T4	0.5445	2.0	4.0	82.0%
Dominican Republic	P3	44.5	79.5	49.2	T3	0.7013	4.0	5.0	91.0%
Ecuador	P2	70.0	77.6	44.6	T2	0.7800	7.0	7.0	77.2%
El Salvador	P3	65.0	62.8	32.5	T4	0.5988	6.5	6.5	66.5%
Grenada	P4	24.5	67.1	51.7	T4	0.6458	2.0	3.0	70.2%
Guatemala	P3	59.0	58.7	40.0	T4	0.5738	5.0	7.0	60.2%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Americas country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Guyana	P3	49.8	67.3	42.5	T4	0.5443	2.5	8.0	83.0%
Haiti	P4	52.5	33.2	22.1	T4	0.2116	3.0	8.0	47.9%
Honduras	P3	69.5	53.1	39.2	T4	0.4856	6.5	7.5	58.6%
Jamaica	P3	59.0	77.3	53.3	T3	0.6678	5.0	7.0	90.1%
Mexico	P1	82.2	80.6	45.0	T2	0.7850	8.0	8.5	83.1%
Nicaragua	P4	37.2	56.9	20.0	T4	0.5318	3.5	4.0	61.4%
Panama	P2	72.2	72.9	47.1	T3	0.7298	7.0	7.5	72.8%
Paraguay	P2	69.5	76.6	32.9	T3	0.7251	6.5	7.5	81.6%
Peru	P3	54.5	81.3	44.6	T3	0.8070	5.0	6.0	82.0%
St. Kitts and Nevis	P4	28.5	69.3	53.8	T4	0.6305	1.5	4.5	76.9%
St. Lucia	P5	22.2	59.6	55.4	T4	0.5255	2.0	2.5	68.2%
St. Vincent and the Grenadines	P4	27.8	66.6	54.2	T4	0.5876	3.0	2.5	76.2%
Suriname	P3	44.5	74.3	31.7	T4	0.6365	4.0	5.0	87.4%
Trinidad and Tobago	P2	65.0	75.3	50.8	T3	0.6973	6.5	6.5	82.2%
United States	P1	80.0	93.2	69.6	T1	0.9194	8.0	8.0	94.7%
Uruguay	P3	42.8	90.9	74.6	T2	0.9006	4.5	4.0	92.0%
Venezuela	P2	60.8	64.0	18.8	T4	0.5360	4.5	8.0	76.7%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Asia country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Afghanistan	P4	58.0	18.7	15.0	T5	0.2083	4.0	8.0	16.1%
Armenia	P4	35.0	82.9	56.3	T4	0.8422	3.5	3.5	81.3%
Azerbaijan	P3	41.2	82.5	39.6	T2	0.7607	3.0	5.5	90.4%
Bahrain	P1	70.0	95.6	53.3	T1	0.9196	7.0	7.0	100.0%
Bangladesh	P3	50.0	60.2	40.8	T1	0.6570	5.0	5.0	53.4%
Bhutan	P4	32.2	76.9	46.7	T3	0.6511	3.0	3.5	91.3%
Brunei	P3	42.2	84.9	46.3	T3	0.7554	4.0	4.5	96.3%
Cambodia	P2	79.5	62.5	34.6	T4	0.5754	7.5	8.5	68.5%
China	P1	80.5	89.3	55.4	T2	0.8718	8.5	7.5	92.0%
Georgia	P3	52.2	80.6	50.0	T2	0.7792	5.0	5.5	83.8%
India	P2	75.0	66.0	54.6	T1	0.6678	7.5	7.5	64.9%
Indonesia	P1	77.2	76.7	42.9	T1	0.7991	7.5	8.0	72.8%
Iran	P1	87.2	74.5	30.8	T3	0.6564	8.5	9.0	85.3%
Iraq	P3	57.0	61.8	33.3	T4	0.4572	3.0	9.0	81.5%
Israel	P1	81.0	89.3	55.4	T2	0.9014	9.0	7.0	88.2%
Japan	P3	54.0	89.9	75.0	T1	0.9351	4.5	6.5	85.5%
Jordan	P3	59.5	80.7	56.3	T1	0.6849	5.5	6.5	95.6%
Kazakhstan	P3	49.0	91.6	46.7	T2	0.9009	4.0	6.0	93.4%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Asia country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Korea, DPR	P3	85.5	12.8	16.3	T5	0.2320	9.0	8.0	0.0%
Korea, Rep.	P3	45.0	97.3	80.0	T1	0.9679	4.5	4.5	97.9%
Kuwait	P2	70.0	87.9	54.2	T3	0.7812	7.0	7.0	99.7%
Kyrgyzstan	P3	45.8	81.7	37.9	T3	0.7316	3.0	6.5	92.0%
Laos	P3	66.2	53.7	33.8	T4	0.4404	5.5	8.0	65.6%
Lebanon	P2	65.2	66.3	30.8	T4	0.5449	4.5	9.0	80.6%
Malaysia	P1	80.0	88.7	57.1	T1	0.8111	8.0	8.0	98.0%
Maldives	P4	37.2	75.5	39.6	T5	0.6745	3.5	4.0	85.2%
Mongolia	P3	55.5	84.8	52.9	T3	0.8457	6.0	5.0	85.1%
Myanmar	P2	79.5	48.0	14.6	T3	0.5001	7.5	8.5	45.4%
Nepal	P4	53.5	52.6	38.8	T3	0.5781	4.0	7.0	46.3%
Oman	P2	56.8	90.0	53.3	T1	0.8576	5.0	6.5	95.3%
Pakistan	P3	59.0	53.8	39.6	T1	0.5096	5.0	7.0	57.3%
Philippines	P2	77.2	72.2	44.6	T2	0.7621	7.5	8.0	67.3%
Qatar	P2	69.5	89.5	54.6	T1	0.8244	6.5	7.5	98.1%
Saudi Arabia	P1	79.5	97.8	37.9	T1	0.9602	7.5	8.5	100.0%
Singapore	P2	65.5	95.8	77.9	T1	0.9691	7.0	6.0	94.4%
Sri Lanka	P4	39.5	61.2	37.1	T2	0.6667	3.5	4.5	54.6%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Asia country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Syria	P4	58.0	36.6	20.0	T4	0.3888	4.0	8.0	33.8%
Tajikistan	P4	40.2	55.9	25.4	T4	0.5606	2.0	6.5	55.8%
Thailand	P1	71.8	86.8	49.6	T1	0.8351	6.5	8.0	90.9%
Timor-Leste	P5	38.5	37.9	38.8	T5	0.4020	2.5	5.5	35.0%
Turkey	P1	68.5	88.3	39.6	T1	0.8913	5.5	8.5	87.3%
Turkmenistan	P4	43.0	34.3	21.7	T4	0.4757	2.5	6.5	18.0%
United Arab Emirates	P1	84.5	97.4	51.3	T1	0.9533	8.0	9.0	100.0%
Uzbekistan	P3	43.5	84.3	38.3	T2	0.7999	3.0	6.0	89.5%
Vietnam	P2	70.0	80.3	47.9	T1	0.7709	7.0	7.0	84.2%
Yemen	P4	49.8	20.6	19.2	T5	0.2317	2.5	8.0	17.5%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Europe country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Albania	P3	57.2	82.6	51.7	T2	0.8000	5.5	6.0	85.9%
Andorra	P4	35.8	80.4	80.0	T3	0.6893	2.0	5.5	94.4%
Austria	P3	57.2	92.6	75.4	T2	0.9065	5.5	6.0	94.9%
Belarus	P1	74.5	83.4	30.8	T3	0.7445	7.0	8.0	94.3%
Belgium	P2	65.0	87.8	71.7	T1	0.8121	6.5	6.5	95.8%
Bosnia and Herzegovina	P3	51.2	73.6	38.8	T4	0.6329	4.0	6.5	86.1%
Bulgaria	P2	64.5	81.9	53.3	T3	0.8145	6.0	7.0	82.4%
Croatia	P3	54.5	86.1	59.2	T2	0.8818	5.0	6.0	83.6%
Cyprus	P3	49.5	87.7	45.0	T1	0.8619	4.5	5.5	89.6%
Czech Republic	P3	57.2	84.8	65.0	T2	0.8239	5.5	6.0	87.7%
Denmark	P2	60.5	99.1	82.1	T1	0.9847	6.5	5.5	99.8%
Estonia	P2	70.0	95.0	80.4	T1	0.9727	7.0	7.0	92.2%
Finland	P3	50.5	94.8	84.2	T1	0.9575	5.5	4.5	93.7%
France	P2	65.0	88.0	70.4	T1	0.8744	6.5	6.5	88.7%
Germany	P2	72.2	93.7	74.2	T1	0.9382	7.0	7.5	93.5%
Greece	P3	54.0	86.5	51.3	T1	0.8674	4.5	6.5	86.3%
Hungary	P2	59.5	86.4	47.9	T2	0.8043	5.5	6.5	93.8%
Iceland	P3	52.8	97.4	81.7	T1	0.9671	5.5	5.0	98.2%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Europe country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.

Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Ireland	P2	60.0	94.0	74.2	T2	0.9138	6.0	6.0	97.2%
Italy	P2	64.5	86.1	63.8	T1	0.8356	6.0	7.0	89.2%
Latvia	P3	62.2	90.4	76.3	T3	0.8852	6.0	6.5	92.7%
Liechtenstein	P3	53.5	91.1	84.2	T4	0.8528	4.0	7.0	98.3%
Lithuania	P3	57.2	90.2	73.8	T2	0.9110	5.5	6.0	89.2%
Luxembourg	P3	56.2	91.0	75.0	T1	0.8466	4.5	7.0	98.8%
Malta	P2	66.2	91.1	50.8	T2	0.8886	5.5	8.0	93.9%
Moldova	P1	77.2	77.3	40.8	T3	0.7719	7.5	8.0	77.4%
Monaco	P4	29.5	84.0	58.3	T3	0.7175	2.5	3.5	99.0%
Montenegro	P3	54.5	79.7	48.3	T3	0.7211	5.0	6.0	88.9%
Netherlands	P2	65.5	96.1	74.2	T1	0.9538	7.0	6.0	97.0%
North Macedonia	P3	57.2	81.0	52.9	T3	0.7070	5.5	6.0	93.6%
Norway	P3	55.0	95.8	79.2	T1	0.9315	5.5	5.5	99.0%
Poland	P2	60.0	87.4	60.8	T2	0.8648	6.0	6.0	88.6%
Portugal	P3	52.8	86.1	65.8	T1	0.8415	5.5	5.0	88.5%
Romania	P2	62.2	83.1	58.8	T2	0.7636	6.0	6.5	91.3%
Russia	P1	87.8	89.4	34.2	T2	0.8532	9.0	8.5	94.4%
San Marino	P4	39.8	79.8	52.9	T4	0.6551	1.5	7.0	97.4%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Europe country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.



Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Serbia	P2	59.0	86.9	49.2	T1	0.8618	5.0	7.0	87.7%
Slovakia	P2	62.2	84.5	48.8	T2	0.8021	6.0	6.5	89.8%
Slovenia	P3	56.8	89.0	62.1	T1	0.8759	5.0	6.5	90.8%
Spain	P2	67.2	93.7	66.7	T1	0.9206	6.5	7.0	95.8%
Sweden	P2	65.0	94.3	73.8	T1	0.9326	6.5	6.5	95.5%
Switzerland	P2	72.2	93.3	72.5	T2	0.9003	7.0	7.5	97.3%
Ukraine	P1	78.2	85.7	45.4	T3	0.8841	8.5	7.0	82.5%
United Kingdom	P2	72.2	95.6	76.3	T1	0.9577	7.0	7.5	95.5%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.
Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.

Oceania country intelligence

Bands compare crime-market pressure, digital intensity and organised-crime resilience. ITU commitment is contextual and state or ideological activity is outside the model.


Country	Band	Crime press.	Intensity	OC resil.	GCI	EGDI	Cyber	Fin.	Internet
Australia	P1	75.5	95.9	75.4	T1	0.9577	8.0	7.0	96.1%
Fiji	P3	49.5	70.8	53.8	T4	0.6754	4.5	5.5	74.7%
Kiribati	P4	29.0	65.4	44.6	T3	0.4572	2.0	4.0	89.4%
Marshall Islands	P5	24.0	56.2	56.7	T5	0.4823	1.5	3.5	65.9%
Micronesia (Federated States of)	P5	19.5	35.6	55.4	T5	0.3235	1.5	2.5	39.5%
Nauru	n/a	16.8	n/a	48.3	T4	0.4454	1.0	2.5	n/a
New Zealand	P2	70.0	93.0	78.8	T3	0.9265	7.0	7.0	93.5%
Palau	P4	30.0	65.7	50.8	n/a	0.5072	3.0	3.0	83.9%
Papua New Guinea	P4	60.2	25.4	32.1	T3	0.3076	4.0	8.5	18.8%
Samoa	P5	32.2	36.5	58.8	T4	0.4899	3.0	3.5	21.2%
Solomon Islands	P5	39.8	33.3	47.1	T5	0.3681	1.5	7.0	29.0%
Tonga	P4	31.8	58.8	51.7	T4	0.5164	2.5	4.0	67.4%
Tuvalu	P5	14.5	38.1	59.2	T4	0.4042	1.0	2.0	35.2%
Vanuatu	P5	29.5	39.9	48.8	T3	0.5427	2.5	3.5	22.4%

Scale key: Crime pressure, digital intensity and OC resilience 0-100; EGDI 0-1; Cyber and Fin. are GI-TOC 0-10 inputs; Internet is %.

Sources: ITU GCI 2024; UN EGDI 2024; Global Organized Crime Index 2025; World Bank indicators through 2024. Full fields and observation years are in the workbook.


Selected profiles | economic and security hubs

The cards show three structural dimensions and a relative snapshot band. They do not measure breach probability, blame a jurisdiction or substitute for local intelligence.




Selected profiles | Indo-Pacific

The cards show three structural dimensions and a relative snapshot band. They do not measure breach probability, blame a jurisdiction or substitute for local intelligence.




Selected profiles | geopolitical focus

The cards show three structural dimensions and a relative snapshot band. They do not measure breach probability, blame a jurisdiction or substitute for local intelligence.




Selected profiles | cross-regional contrast

The cards show three structural dimensions and a relative snapshot band. They do not measure breach probability, blame a jurisdiction or substitute for local intelligence.



Turn a country signal into a local decision

The lens is the start of an intelligence requirement. It should never be the final answer.



1 | Define the decision

Supplier diligence, market entry, data location, fraud control, incident support, travel, critical infrastructure and geopolitical exposure need different evidence. Do not reuse the same country band as a universal answer.

2 | Add current intent

Overlay state, criminal and hacktivist actors; target sector; conflict and sanctions; major vulnerabilities; scam routes; payment rails; telecom and cloud dependencies; and recent law-enforcement reporting.

3 | Add enterprise exposure

Map users, customers, data, subsidiaries, counterparties, suppliers, identities, network paths, payment flows, language and recovery dependencies connected to the jurisdiction.

4 | Test alternatives

Could the observed pattern be explained by hosting, population, better reporting, a provider's footprint, one large event or cross-border victimisation? Record competing hypotheses and confidence.

5 | Choose a reversible action

Tighten verification, telemetry, segmentation, supplier terms, incident support or transaction thresholds proportionate to the actual path. Avoid nationality-based controls or broad de-risking unsupported by evidence.

6 | Set expiry

Country assessments decay quickly. Attach evidence cut-off, owner, next review, trigger events and the data gaps that would change the decision.

DECISIONS AND OUTLOOK

The common control problem is not more technology. It is whether trust can be verified, bounded, revoked and recovered before harm compounds.


THE OPERATING REALITY

- Adversaries operate in seconds and minutes; governance, investigation and supplier change often operate in days and quarters.
- Identity, transaction intent, recovery and evidence are the shared control planes across cyber, AI agents and fraud.
- Quantum resilience adds a second clock: the protected data can outlive the cryptography and the asset.



Close the authority-compression gap

Across the evidence, the decisive inequality is simple: if time to contain exceeds time to consequential action, preventive authority boundaries must carry more of the defence.



Seconds | transfer

Criminal access can move to a secondary actor in 22 seconds in Mandiant's observed set. An identity that survives endpoint isolation preserves the intrusion path.^[5]

Minutes | execute

CVE scanning can begin within 15 minutes of disclosure; CrowdStrike's fastest eCrime breakout was 27 seconds and its 2025 average was 29 minutes. Human-only triage cannot be the sole gate.^{[1][2]}

Hours | remove options

Unit 42's fastest quartile exfiltrated in 1.2 hours. Attackers increasingly target backup, identity and cloud control planes before a crisis team assembles.^[1]

One run | delegate

A privileged agent can read hostile context and execute with legitimate credentials. Maximum reach, rate, persistence and irreversibility matter more than the model's brand.

One transfer | defraud


The payment event converts a persuasive story into loss. Independent identity and exact-transaction verification must happen before release, while rapid holds determine recoverability.

Years | preserve

Long-lived secrets and signatures may outlast both cryptography and procurement cycles. Quantum migration must begin from data life and replaceability, not a guessed Q-Day.

The minimum viable trust architecture

Four independently enforced planes reduce risk across malware-free intrusion, agentic action, fraud and long-lived cryptography.



1 | Identity and delegation

Phishing-resistant human identity; stable agent and workload identity; short-lived sender- and audience-bound credentials; complete actor chain; device and session assurance; rapid descendant revocation.

2 | Transaction and authority

Resource-side authorisation; least privilege; exact action preview; payee, amount, destination and code-diff binding; separation of duties; budgets, rate limits, egress and safe defaults.

3 | Recovery and continuity

Clean-room identity, immutable evidence, independently administered backups, tested rebuild, supplier continuity, wallet and signing recovery, memory quarantine, and time-bounded rollback without restoring the vulnerability.

4 | Evidence and intelligence

Event provenance, configuration history, tool and model version, negotiated cryptography, source confidence, claim-versus-impact separation, data cut-off and decision expiry.

Architecture test

Compromise the production identity plane, inject a trusted source, change one transaction field and remove one supplier. The organisation should contain the action, reconstruct causality and recover without trusting the compromised plane.

A board dashboard that resists greenwashing

Count evidence of bounded harm and tested recovery. Avoid maturity scores that improve while critical authority remains unmeasured.

Cyber

Time to revoke all active sessions; internet-facing asset coverage; privileged identity and OAuth grant age; high-confidence exposure-to-fix time; SaaS and cloud evidence coverage; clean recovery time; backup deletion-path tests.

Agentic AI

Owned agent inventory; maximum authority per run; shared/static credentials; exact-approval coverage; influential-context provenance; unapproved tools; kill-switch reach; descendant-token revocation; memory purge time; replay-test pass rate.

Quantum

Material systems with complete CBOM; protected-data and signature lifetimes; supplier plans with dates; negotiated-algorithm telemetry; priority migration paths tested; ownerless dependencies; vulnerable fallback; aged exceptions.

Fraud

Attempted, stopped, lost, frozen and reimbursed value; severity distribution; changed-payee verification; report-to-hold time; beneficiary and wallet propagation; mule detection; repeat victimisation; victim-support outcomes.

Threat actors

Transfer-point coverage; verified versus claimed events; source-confidence age; critical-process target coverage; edge configuration drift; personal/high-value identity protection; recovery-plane attempts; downstream hunt completion.

One executive test

For the largest credible loss path, can management state the exact authority required, maximum harm before an independent gate, first reliable signal, time to revoke, evidence retained and recovery time under compromised administration?

Three scenarios for 2026-2028

These are conditional scenarios, not predictions. Confidence is attached to the mechanism and each scenario has observable signposts.

Base case | accelerated known tradecraft | high

Actors use AI to improve research, social engineering, coding and scale, but humans still choose strategic objectives. Identity and edge weaknesses dominate initial access; agent adoption expands faster than governance; PQC pilots expose dependency debt.

Signposts

Malware-free and valid-account activity remains high; more production agents hold tools and persistent state; transaction impersonation becomes multilingual and adaptive; major platforms deploy final-standard hybrids; breach and fraud severity remain concentrated.

Stress case | delegated trust cascade | medium

A compromised identity, supplier or agentic workflow crosses SaaS, cloud, code, payments and recovery before defenders can reconstruct delegation. A criminal service chain combines synthetic trust, token theft, extortion and rapid laundering.

Tail case | compressed cryptographic horizon | low, high impact

A credible cryptanalytic or hardware advance shortens transition assumptions while selected long-lived traffic has already been collected. Migration bottlenecks concentrate in firmware, identity, signing and regulated records.

Preparedness test

The same architecture should help all three: small and short-lived authority, independent transaction verification, cross-surface evidence, rapid revocation, segmented recovery, cryptographic inventory and time-bounded supplier commitments.

The 2026-2027 signpost watchlist

A forecast becomes useful when teams know which observable development should change posture, budget or control.

Q4 2026

US/EU/AU plans

first public evidence

2027

procurement gate

CNSA + CII plans

2028

proof year

UK/AU/SG/Korea

Agent autonomy

Multiple independent responders verify sustained machine-tempo operations, agent-created credentials, cross-service token replay, persistent poisoned memory or autonomous exploitation outside reduced-safeguard research.

Trust fraud

Financial institutions verify agent-managed relationships, real-time persona switching, detector bypass and mule routing; or transaction-bound verification measurably reduces net loss despite rising reports.

Access markets

Actor hand-off falls further, access is routinely pre-staged across cloud and SaaS, contractor or exploit-market disclosures expand, and recovery-plane targeting becomes a default extortion step.

Geopolitical effects

Verified disruptive or destructive effects rise above hacktivist claim volume; logistics and personal-account targeting broadens; exposed OT manipulation follows regional triggers; remote-worker infrastructure shifts jurisdictions.

Quantum execution


CNSA 2027 procurement evidence, EU national roadmaps, Singapore CII plans, UK 2028 discovery metrics and Australian critical-system pilots become public and comparable; products adopt RFC 10024 while retiring draft identifiers.

Country lens refresh

Update when GI-TOC, UN EGD, ITU or World Bank vintages change, or when a material national event invalidates a structural assumption. Do not mix a new component into an old percentile band.

CiBRAI operating model: from signal to proof

The report's recommendations map to an integrated operating cycle. Capability names describe the intended workflow; organisations should validate coverage, integration and effectiveness in their own deployment.



Observe + Intelligence

OBSERVE provides the command-centre, attack-map, risk and war-room view. **INTELLIGENCE** links CTI, dark-web, MITRE and OSINT context to the asset, identity, country and actor hypothesis.

Detect + Investigate

DETECT joins SIEM, network, endpoint, identity, vulnerability and behavioural signals. **INVESTIGATE** connects search, stories, logs, evidence, email, scam, domain, sandbox, cloud and topology analysis.

Respond + Manage

RESPOND coordinates cases, forensics, directory and recovery action, project rooms and exercises. **MANAGE** turns findings into deployment, asset, endpoint, policy, patch, CISO, compliance and resilience work.

Assistant + Configure

ASSISTANT supports skills, agent workflows, controlled approvals, trace, audit and playbooks. **CONFIGURE** governs integrations, tenants, health, pipelines and detection content.

Assurance principle

The platform should make source, authority, action, approval, evidence and outcome visible. It does not replace accountable human judgement, independent validation, legal process or business ownership.

Clarity. Not chaos.

The defining threat of 2026 is the conversion of trusted access into consequential action faster than organisations can verify intent, revoke authority and preserve recovery.

What to protect

Protect human and non-human identity, exact transaction intent, the recovery plane, influential data and memory, cryptographic roots of trust, and the evidence needed to explain an action after the fact.

What to stop measuring

Stop treating attack counts, country ranks, AI mentions, tool inventories, blocked prompts, quantum-capable products or backup presence as outcome measures without a denominator and an end-to-end test.

What to do now

Choose the largest credible loss path. Shorten its authority, install an independent gate, instrument the causal chain, test full revocation, prove recovery and attach a deadline to every unsupported dependency.

Gadget Access + CiBRAI

Australian owned, with sovereign cloud, rack and on-premise deployment options. Australia: 1800 928 982, +61 2 8005 6562, +61 2 9000 1137. United States: +1 818 649 9444.

Contact

support@cibrai.com

www.cibrai.com

Evidence cut-off: 31 August 2026. Model version: 2026.1.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[1] Palo Alto Networks Unit 42. 2026 Global Incident Response Report. 2026.

<https://www.paloaltonetworks.com/resources/research/unit-42-incident-response-report>

More than 750 incident-response engagements during 2025 in over 50 countries; identity, cloud, browser and exfiltration findings.

[2] CrowdStrike. 2026 Global Threat Report. February 2026.

Supplied PDF: CrowdStrike-2026-Global-Threat-Report.pdf, especially pp. 2-19, 32-46.

[3] Cloudflare. Navigating the New Security Landscape: Asia Pacific Cybersecurity Readiness Survey. September 2024.

Supplied PDF: Navigating_the_New_Security_Landscape_Asia_Pacific_Cybersecurity_Readiness_Survey.pdf, especially pp. 3-15.

[4] Verizon. 2026 Data Breach Investigations Report - Executive Summary. 2026.

<https://www.verizon.com/business/resources/executivebriefs/2026-dbir-executive-summary.pdf>

More than 31,000 incidents and 22,000 confirmed breaches across organisations in 145 countries; incident window 1 Nov 2024 to 31 Oct 2025.

[5] Google Cloud / Mandiant. M-Trends 2026: Data, Insights, and Strategies From the Frontlines. 23 March 2026.

<https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2026>

More than 500,000 hours of 2025 frontline investigations; no public engagement count.

[6] Microsoft. Microsoft Digital Defense Report 2025. October 2025.

<https://aka.ms/Microsoft-Digital-Defense-Report-2025>

Microsoft FY2025 telemetry and investigations; chart populations vary by section.

[7] IBM. 2026 X-Force Threat Intelligence Index. 25 February 2026.

<https://newsroom.ibm.com/2026-02-25-ibm-2026-x-force-threat-index-ai-driven-attacks-are-escalating-as-basic-security-gaps-leave-enterprises-exposed>

Official launch findings; public sources do not disclose a unified case count.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[8] Sophos / Vanson Bourne. The State of Ransomware 2026. 9 June 2026.

<https://assets.sophos.com/X24WTUEQ/at/jbww7pmb8n3gp99wr6hfq4/sophos-state-ransomware-report-2026.pdf>

Survey fielded January-March 2026 among 2,158 organisations hit by ransomware in the preceding 12 months across 17 countries; 100-5,000 employees.

[9] CISA. Known Exploited Vulnerabilities Catalog. snapshot to 31 August 2026.

<https://www.cisa.gov/known-exploited-vulnerabilities-catalog>

Authoritative catalogue of vulnerabilities known to be exploited in the wild; not a complete census of exploitation.

[10] ENISA. ENISA Threat Landscape 2025. 1 October 2025; v1.2, 9 January 2026.

<https://www.enisa.europa.eu/publications/enisa-threat-landscape-2025>

EU-focused synthesis of 4,875 curated incidents observed from 1 July 2024 to 30 June 2025; source mix and regional remit constrain global inference.

[11] World Economic Forum. Global Cybersecurity Outlook 2026. 2026.

<https://www.weforum.org/publications/global-cybersecurity-outlook-2026/>

Executive survey and expert perspective; measures perceptions and preparedness rather than incident prevalence.

[12] Federal Bureau of Investigation, Internet Crime Complaint Center. 2025 IC3 Annual Report. 2026.

https://www.ic3.gov/AnnualReport/Reports/2025_IC3Report.pdf

Complaint-derived US and international data; losses are reported and adjusted, not a global estimate.

[13] US Federal Trade Commission. Consumer Sentinel Network Data Book 2024. March 2025.

<https://www.ftc.gov/reports/consumer-sentinel-network-data-book-2024>

Voluntary, generally unverified consumer reports; the FTC and IC3 universes overlap and must not be added.

[14] INTERPOL. Global Financial Fraud Threat Assessment. March 2024.

<https://www.interpol.int/en/News-and-Events/News/2024/INTERPOL-Financial-Fraud-assessment-A-global-threat-boosted-by-technology>

Law-enforcement assessment of transnational fraud industrialisation and enabling technologies.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[15] UNODC. Inflection Point: Global Implications of Scam Centres, Underground Banking and Illicit Online Marketplaces in Southeast Asia. April 2025.

https://www.unodc.org/roseap/uploads/documents/Publications/2025/Inflection_Point_2025.pdf

Regional organised-crime analysis with global implications; estimates carry uncertainty.

[16] Global Anti-Scam Alliance. Global State of Scams 2025 survey summary. 2025.

<https://gasa.org/knowledge-base/blog/global-scams-on-the-rise-over-half-of-adults-worldwide-report-scam-encounters>

Survey of 46,000 adults in 42 markets; population weighting and extrapolation make the loss figure a modelled estimate.

[17] Chainalysis. 2026 Crypto Crime Report. 2026.

<https://www.chainalysis.com/blog/2026-crypto-crime-report-introduction/>

Blockchain analytics estimates; attribution, service coverage and price movements affect totals.

[18] FinCEN. Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions. 13 November 2024.

<https://www.fincen.gov/sites/default/files/shared/FinCEN-Alert-DeepFakes-Alert508FINAL.pdf>

Operational indicators and case patterns, not a prevalence estimate.

[19] OpenAI. Disrupting malicious uses of AI by state-affiliated threat actors. 14 February 2024.

<https://openai.com/index/disrupting-malicious-uses-of-ai-by-state-affiliated-threat-actors/>

Five state-affiliated clusters; provider enforcement visibility only.

[20] Google Threat Intelligence Group. Adversarial misuse of generative AI. 29 January 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/adversarial-misuse-generative-ai>

Provider telemetry across tracked state actors and information operations.

[21] Anthropic. Threat Intelligence Report. August 2025.

<https://www-cdn.anthropic.com/b2a76c6f6992465c09a6f2fce282f6c0cea8c200.pdf>

Selected enforcement cases, including agent-assisted extortion and employment fraud.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[22] Anthropic. Disrupting the first reported AI-orchestrated cyber espionage campaign. November 2025.

<https://assets.anthropic.com/m/ec212e6566a0d47/original/Disrupting-the-first-reported-AI-orchestrated-cyber-espionage-campaign.pdf>

Single-provider campaign reconstruction; tactical autonomy percentages are estimates.

[23] OpenAI. Hugging Face model-evaluation security incident and the road ahead. August 2026.

<https://openai.com/index/hugging-face-incident-and-the-road-ahead/>

Reduced-safeguard research incident; not a normal production deployment.

[24] METR. OpenAI-Hugging Face incident investigation. 26 August 2026.

<https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>

Independent investigation with access, telemetry and redaction constraints.

[25] UK AI Security Institute. Incident report: unsanctioned agent behaviour during cyber testing. August 2026.

<https://www.aisi.gov.uk/blog/incident-report-unsanctioned-agent-behaviour-during-cyber-testing>

Experimental runs with internet access and disabled safeguards; no evidenced harm.

[26] NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, AI 600-1. July 2024.

<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>

Voluntary risk-management profile.

[27] NIST. Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, AI 100-2e2025. March 2025.

<https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-2e2025.pdf>

Taxonomy, not prevalence data.

[28] OWASP GenAI Security Project. Top 10 for Agentic Applications 2026. December 2025.

<https://genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026/>

Community taxonomy, not a statistical ranking of incidents.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[29] MITRE. ATLAS. release available by 31 August 2026.

<https://atlas.mitre.org/>

Adversarial threat knowledge base; evidence maturity is not incidence.

[30] NIST. AI Agent Standards Initiative. February-August 2026.

<https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative>

Standards coordination initiative; status is evolving.

[31] Model Context Protocol. Specification and Security Best Practices. 28 July 2026.

https://modelcontextprotocol.io/docs/2026-07-28/tutorials/security/security_best_practices

Protocol guidance; implementation support does not itself constitute security assurance.

[32] NIST. FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard. 13 August 2024.

<https://csrc.nist.gov/pubs/fips/203/final>

Final post-quantum key-establishment standard, based on ML-KEM.

[33] NIST. FIPS 204: Module-Lattice-Based Digital Signature Standard. 13 August 2024.

<https://csrc.nist.gov/pubs/fips/204/final>

Final post-quantum signature standard, based on ML-DSA.

[34] NIST. FIPS 205: Stateless Hash-Based Digital Signature Standard. 13 August 2024.

<https://csrc.nist.gov/pubs/fips/205/final>

Final backup signature standard, based on SLH-DSA.

[35] NIST. IR 8547 (Initial Public Draft): Transition to Post-Quantum Cryptography Standards. November 2024.

<https://csrc.nist.gov/pubs/ir/8547/ipd>

Draft transition guidance; dates are proposed, not final mandates.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[36] US National Security Agency. Commercial National Security Algorithm Suite 2.0 FAQ, version 2.1. December 2024.
https://media.defense.gov/2022/Sep/07/2003071836/-1/-1/1/CSI_CNESA_2.0_FAQ_.PDF

Prescriptive transition expectations for US National Security Systems; not a general commercial mandate.

[37] UK National Cyber Security Centre. Timelines for migration to post-quantum cryptography. 20 March 2025.
<https://www.ncsc.gov.uk/guidance/pqc-migration-timelines>

Official UK target milestones for discovery, priority migration and completion.

[38] European Commission. A Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography. 23 June 2025.

<https://digital-strategy.ec.europa.eu/en/library/coordinated-implementation-roadmap-transition-post-quantum-cryptography>

EU coordination roadmap; national implementation determines obligations.

[39] Australian Signals Directorate / ACSC. Planning for Post-Quantum Cryptography. updated 22 September 2025.

<https://www.cyber.gov.au/business-government/secure-design/quantum/planning-for-post-quantum-cryptography>

Australian guidance with 2026, 2028 and 2030 planning anchors; recommendation and scope language retained.

[40] International Telecommunication Union. Global Cybersecurity Index 2024. September 2024.

https://www.itu.int/dms_pub/itu-d/opb/hdb/D-HDB-GCI.01-2024-PDF-E.pdf

Country commitment tiers across legal, technical, organisational, capacity-development and cooperation pillars; not operational control quality or impact.

[41] United Nations DESA. UN E-Government Survey 2024. 2024.

<https://desapublications.un.org/publications/un-e-government-survey-2024>

EGDI for 193 UN member states; used here as one component of a digital intensity proxy.

[42] Global Initiative Against Transnational Organized Crime. Global Organized Crime Index 2025. 2025.

<https://ocindex.net/downloads>

193-country expert-assessed cyber-dependent crime, financial crime and resilience scores.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[43] World Bank. Individuals using the Internet (% of population). latest observation through 2024.

<https://data.worldbank.org/indicator/IT.NET.USER.ZS>

Country indicator; observation year retained in the supporting workbook.

[44] World Bank. Secure Internet servers (per 1 million people). latest observation through 2024.

<https://data.worldbank.org/indicator/IT.NET.SECR.P6>

Context-only hosting-intensity indicator; excluded from the composite because hosting concentration and microstate denominators distort comparison.

[45] World Bank. World Development Indicators: population and GDP. latest observation through 2024.

<https://databank.worldbank.org/source/world-development-indicators>

Population and GDP are context fields and are not used directly in the composite.

[46] Australian Signals Directorate. Annual Cyber Threat Report 2024-2025. 2025.

<https://www.cyber.gov.au/about-us/view-all-content/reports-and-statistics/asd-cyber-threat-report-july-2024-june-2025>

Australian national reporting; reporting and jurisdiction constrain comparison.

[47] OpenAI. Disrupting malicious uses of AI: an update. October 2025.

<https://openai.com/global-affairs/disrupting-malicious-uses-of-ai-october-2025/>

Provider enforcement cases and defensive-use telemetry.

[48] Google Threat Intelligence Group. Threat actor usage of AI tools. November 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/threat-actor-usage-of-ai-tools>

Observed runtime LLM malware and synthetic-media use; several samples were experimental.

[49] Microsoft Threat Intelligence. AI as tradecraft: how threat actors operationalize AI. 6 March 2026.

<https://www.microsoft.com/en-us/security/blog/2026/03/06/ai-as-tradecraft-how-threat-actors-operationalize-ai/>

Provider telemetry; most observed AI use remained human-directed acceleration.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[50] Anthropic. AI-enabled cyber threats mapped to MITRE ATT&CK. June 2026.

<https://www.anthropic.com/news/AI-enabled-cyber-threats-mitre-attack>

Selected sample of 832 banned accounts; not a random attacker sample.

[51] Dow Jones / Oxford Analytica. Dragonfly geopolitical and security intelligence. accessed August 2026.

<https://www.dowjones.com/business-intelligence/lp/oxford-analytica-dragonfly>

Commercial strategic intelligence reference supplied by the user; no public cross-country metric imported into the index.

[52] NIST. Digital Identity Guidelines, SP 800-63-4. 31 July 2025.

<https://csrc.nist.gov/pubs/sp/800/63/4/final>

Principally human identity proofing, authentication and federation.

[53] IETF. OAuth 2.0 Security Best Current Practice, RFC 9700. January 2025.

<https://www.rfc-editor.org/rfc/rfc9700.html>

Sender-constrained, audience-restricted token guidance and related controls.

[54] Australian National Anti-Scam Centre. Targeting Scams Report 2025. 2026.

<https://www.nasc.gov.au/system/files/targeting-scams-report-2025.pdf>

Combined five-source Australian dataset with disclosed deduplication and contributor changes.

[55] Singapore Police Force. Annual Scams and Cybercrime Brief 2025. 2026.

<https://www.police.gov.sg/-/media/SPF/Media-Room/Statistics/Annual-Scams-and-Cybercrime-Brief-2025/Annual-Scam-and-Cybercrime-Brief-2025.pdf>

Administrative scam and cybercrime cases; frozen or recovered funds are not necessarily returned.

[56] UK Finance. Annual Fraud Report 2026 press release. 2026.

<https://www.ukfinance.org.uk/news-and-insight/press-release/fraud-report-2026-press-release>

Member payment-fraud data for 2025; reimbursement scope differs from the mandatory PSR scheme.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[57] Japan National Police Agency. 2025 special fraud and SNS investment/romance statistics. 5 June 2026.

<https://www.npa.go.jp/bureau/safetylife/sos47/new-topics/260605/01.html>

Parallel police series with different legal and operational definitions.

[58] UNODC. An Interconnected Criminal Ecosystem: Transnational Organized Crime Threat Assessment for Southeast Asia 2026. July 2026.

https://www.unodc.org/roseap/uploads/documents/Publications/2026/TOCTA_South-East_Asia_2026.pdf

Modelled 2025 victim-loss range covering East Asia, Southeast Asia, Australia and New Zealand; it overlaps national and global estimates.

[59] Nasdaq Verafin / Celent / Oliver Wyman. 2026 Global Financial Crime Report. 11 March 2026.

<https://verafin.com/wp-content/uploads/2026/03/global-financial-crime-report-2026-nasdaq-verafin-20260316.pdf>

Bottom-up and top-down economic model drawing on more than 500 sources plus a 505-executive survey; not observed ledger loss.

[60] Chainalysis. 2026 Crypto Crime Report: Scams. 2026.

<https://www.chainalysis.com/blog/crypto-scams-2026/>

Known-address lower bound plus projection; excludes off-chain payments and unidentified addresses.

[61] INTERPOL. USD 439 million recovered in global financial crime operation HAECHI VI. 2025.

<https://www.interpol.int/en/News-and-Events/News/2025/USD-439-million-recovered-in-global-financial-crime-operation>

Operational output across 40 jurisdictions with no common attempted-loss denominator; recovered is not necessarily reimbursed.

[62] International Organization for Migration, reported by Reuters. Human trafficking into Asian scam centres is surging. 28 July 2026.

<https://www.reuters.com/legal/government/un-says-human-trafficking-into-asian-scam-centres-is-surging-2026-07-28/>

Secondary report of IOM estimate; used with explicit source qualification.

[63] Europol. EU Serious and Organised Crime Threat Assessment 2025. March 2025.

<https://www.europol.europa.eu/cms/sites/default/files/documents/EU-SOCTA-2025.pdf>

EU-focused strategic assessment; not a global incident census.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[64] NSA, CISA, FBI, DC3 and international partners. Countering Chinese State-Sponsored Actors Compromise of Networks Worldwide to Feed Global Espionage System. 27 August 2025.

https://media.defense.gov/2025/Aug/22/2003786665/-1/-1/0/CSA_COUNTERING_CHINA_STATE_ACTORS_COMPROMISE_OF_NETWORKS.PDF
Joint advisory; commercial actor-name overlaps are not one-to-one mappings.

[65] FBI, NSA and Cyber National Mission Force. China-Linked Hacking Group QTFY Targets Military and Critical Infrastructure with Malicious Distributed Systems. 26 August 2026.

<https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4583539/nsa-joins-fbi-in-issuing-warning-about-chinese-hacking-group-qtfy-cyber-activity/>
Pre-cut-off official release; later PDF update excluded from the evidence window.

[66] NSA, FBI, NCSC-UK, CISA, ASD ACSC and partners. Russian GRU Targeting Western Logistics Entities and Technology Companies. 21 May 2025; v1.1 April 2026.

https://media.defense.gov/2025/May/21/2003719846/-1/-1/0/CSA_RUSSIAN_GRU_TARGET_LOGISTICS.PDF
Targeting-attempt sample is not a confirmed-compromise count.

[67] Google Threat Intelligence Group. Going with the Flow(s): Distinct Clusters Target Individuals of Interest to Russia. 20 August 2026.

<https://cloud.google.com/blog/topics/threat-intelligence/distinct-clusters-target-individuals-of-interest-to-russia>
Provider attribution and case visibility; shared tooling does not prove shared command.

[68] Google Threat Intelligence Group. STOCKSTAY Another Day: The Latest Addition to Turla's Intelligence Gathering Apparatus. 25 June 2026.

<https://cloud.google.com/blog/topics/threat-intelligence/stockstay-turla-intelligence-gathering>
Technical campaign report.

[69] Google Threat Intelligence Group. Unearthing APT44: Russia's Notorious Cyber Sabotage Unit Sandworm. 17 April 2024.

<https://cloud.google.com/blog/topics/threat-intelligence/apt44-unearthing-sandworm>
Strategic lineage, not a 2025 event count.

[70] FBI Internet Crime Complaint Center. North Korea Responsible for USD 1.5 Billion Bybit Hack. 26 February 2025.

<https://www.ic3.gov/PSA/2025/PSA250226>
Attribution of one theft, not a trend line.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[71] Google Threat Intelligence Group. DPRK IT Workers Expanding in Scope and Scale. 1 April 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/dprk-it-workers-expanding-scope-scale>

Provider investigation of fraudulent remote-work activity.

[72] US Department of Justice. Coordinated actions to combat North Korean remote IT workers' illicit revenue schemes. 30 June 2025.

<https://www.justice.gov/opa/pr/justice-department-announces-coordinated-nationwide-actions-combat-north-korean-remote>

Charges are allegations until proved; seizure figures are enforcement outputs.

[73] NSA, CISA, FBI and DC3. Iranian Cyber Actors May Target Vulnerable U.S. Networks and Entities of Interest. 30 June 2025.

<https://www.nsa.gov/Press-Room/Press-Releases-Statements/Press-Release-View/Article/4229506/nsa-cisa-fbi-and-dc3-warn-iranian-cyber-actors-may-target-vulnerable-us-network/>

Joint warning; risk assessment rather than incident population.

[74] Google Threat Intelligence Group. The Cost of a Call: From Voice Phishing to Data Extortion. 4 June 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/voice-phishing-data-extortion>

Provider incident and campaign analysis.

[75] Google Threat Intelligence Group. Cybercrime: A Multifaceted National Security Threat. 12 February 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/cybercrime-multifaceted-national-security-threat>

Strategic analysis of criminal-service ecosystems.

[76] Google Threat Intelligence Group. Sanctioned but Still Spying: Intellexa's Prolific Zero-Day Exploits Continue. 3 December 2025.

<https://cloud.google.com/blog/topics/threat-intelligence/intellexa-zero-day-exploits-continue>

Google-discovered zero-days only; not a global zero-day denominator.

[77] Google Threat Intelligence Group. Coruna: The Mysterious Journey of a Powerful iOS Exploit Kit. 3 March 2026.

<https://cloud.google.com/blog/topics/threat-intelligence/coruna-powerful-ios-exploit-kit>

Observed kit transfer across commercial surveillance, espionage and financial-crime contexts.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[78] Google Threat Intelligence Group. The Bear Necessities: The Pro-Russia Influence Ecosystem. 29 June 2026.

<https://cloud.google.com/blog/topics/threat-intelligence/pro-russia-influence-ecosystem>

Provider assessment of a layered influence ecosystem.

[79] CISA and US government partners. Iranian-Affiliated Cyber Actors Exploit Programmable Logic Controllers Across U.S. Critical Infrastructure, AA26-097A. 7 April 2026; updated July 2026.

<https://www.cisa.gov/news-events/cybersecurity-advisories/aa26-097a>

Official advisory using affiliated-actor language; a small number of observed cases produced disruption and financial loss.

[80] Associated Press. Cryptocurrency exchange says it was victim of \$1.5 billion hack. 21 February 2025.

<https://apnews.com/article/88256366c723a9de8327ef3d4071057e>

Contemporary report of Bybit's description of a manipulated wallet-signing workflow; does not independently attribute the operation.

[81] NIST. Considerations for Achieving Cryptographic Agility, CSWP 39-upd1. 24 June 2026.

<https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.39-upd1.pdf>

Operational design guidance; crypto agility is a governed capability, not a product label.

[82] IETF. RFC 9881: Internet X.509 PKI Algorithm Identifiers for ML-DSA. October 2025.

<https://datatracker.ietf.org/doc/rfc9881/>

Proposed Standard; encoding support does not migrate certificate authorities, HSMs, relying parties or trust stores.

[83] IETF. RFC 9882: Use of ML-DSA in the Cryptographic Message Syntax. October 2025.

<https://datatracker.ietf.org/doc/rfc9882/>

Proposed Standard for CMS; operational ecosystem support varies.

[84] IETF. RFC 10024: Post-Quantum Traditional Hybrid Key Agreement Mechanisms for TLS 1.3. August 2026.

<https://datatracker.ietf.org/doc/rfc10024/>

Proposed Standard using final FIPS 203; its security analysis is specific to TLS 1.3.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[85] White House. Executive Order 14412: Securing the Nation Against Advanced Cryptographic Attacks. 22 June 2026.

<https://www.whitehouse.gov/presidential-actions/2026/06/securing-the-nation-against-advanced-cryptographic-attacks/>

Mandatory for covered US federal systems, not a blanket private-sector deadline.

[86] US Office of Management and Budget. M-26-15: Execution of the Migration to Post-Quantum Cryptography. 24 June 2026.

<https://www.whitehouse.gov/wp-content/uploads/2026/06/M-26-15-Execution-of-the-Migration-to-Post-Quantum-Cryptography.pdf>

Federal execution phases and plan deadlines; scope must be preserved.

[87] US Office of Management and Budget. M-23-02: Migrating to Post-Quantum Cryptography. 18 November 2022.

<https://www.whitehouse.gov/wp-content/uploads/2022/11/M-23-02-M-Memo-on-Migrating-to-Post-Quantum-Cryptography.pdf>

Initial civilian federal inventory was due in 2023 with annual updates through 2035.

[88] European Commission. Recommendation (EU) 2024/1101 on a Coordinated Implementation Roadmap for the transition to Post-Quantum Cryptography. 11 April 2024.

<https://eur-lex.europa.eu/eli/reco/2024/1101/oj/eng>

Recommendation and coordination framework, not an EU regulation.

[89] Australian Signals Directorate / ACSC. Post-quantum questions to ask your vendors. 16 July 2026.

<https://www.cyber.gov.au/business-government/secure-design/quantum/post-quantum-questions-to-ask-your-vendors>

Procurement-oriented guidance accompanying Australia's quantum security programme.

[90] Canadian Centre for Cyber Security. Roadmap for the Migration to Post-Quantum Cryptography for the Government of Canada, ITSM.40.001. 23 June 2025.

<https://www.cyber.gc.ca/en/guidance/roadmap-migration-post-quantum-cryptography-government-canada-itsm40001>

Government of Canada non-classified-system roadmap; scope-specific administrative requirements.

[91] Cyber Security Agency of Singapore. Quantum-Safe Handbook and Quantum Readiness Index. 16 July 2026.

<https://www.csa.gov.sg/resources/publications/quantum-safe-handbook-and-quantum-readiness-index/>

CII transition expectations and a self-assessment tool; the index is not external assurance.

Sources

Only material published on or before 31 August 2026 informs the analysis. Web pages may update later; where relevant, the cited document or version and cut-off treatment are stated. Source scope, denominator and caveat are retained with claims.

[92] CRYPTREC. e-Government Recommended Ciphers List. revised 30 March 2026.

<https://www.cryptrec.go.jp/en/list.html>

Japan central-government cryptographic procurement guidance; PQ signature algorithms remained under assessment at cut-off.

[93] Republic of Korea Ministry of Science and ICT. Comprehensive plan for transition to post-quantum cryptography. 27 February 2025.

<https://www.msit.go.kr/bbs/view.do?sCode=user&mId=307&mPid=208&bbsSeqNo=94&nttSeqNo=3185512>

National programme goal and sector pilots; not a blanket private-sector statute.

[94] Government of India Department of Science and Technology Task Force. Quantum-Safe Ecosystem in India. 21 May 2026.

<https://dst.gov.in/sites/default/files/Quantum-Safe-Ecosystem-in-India.pdf>

Task-force recommendation with ambitious dates; expressly not itself a regulatory mandate.

[95] New Zealand Government Communications Security Bureau. New Zealand Information Security Manual, section 2.4. v3.8, September 2024; v3.9 change log checked to cut-off.

<https://nzism.gcsb.govt.nz/assets/NZISM/ISM-Documents/V-3.8-September-2024-12.pdf>

Section 2.4 readiness and inventory guidance; the April 2025 v3.9 change log was checked and no public PQC approval update or whole-of-government completion date was verified.

[96] ETSI. TR 104 016 V1.1.1: Quantum-Safe Cryptography - Migration and Deployment. October 2024.

https://www.etsi.org/deliver/etsi_tr/104000_104099/104016/01.01.01_60/tr_104016v010101p.pdf

Informative technical report; notes enterprise migration can take a decade or more.

[97] Google Cloud. PQC in plain text: Google Cloud's post-quantum cryptography roadmap. 11 August 2026.

<https://cloud.google.com/blog/products/identity-security/pqc-in-plaintext-google-clouds-post-quantum-cryptography-roadmap>

Provider roadmap; customer-side client, configuration and key work remains necessary.

[98] NIST. Post-Quantum Cryptography Standardization. updated 5 August 2026.

<https://csrc.nist.gov/Projects/post-quantum-cryptography/post-quantum-cryptography-standardization>

Programme status at cut-off: FIPS 206 remained in development and HQC was selected but not yet a final standard.