

XCOM — Master Whitepaper & Investor Memorandum v2.0

CONTENTS

M · Master Memorandum v2.0

→ Reading paths

M.1 · Executive summary

M.2 · Core architecture

M.3 · Layers & Infinity Loop

M.4 · Frontier R&D

M.5 · Evidence & rechtspraak-ai

M.6 · Revenue streams

M.7 · Shenzhen & Seed

M.8 · Roadmap & KPIs

Abstract (technical paper)

1 · Introduction & Motivation

2 · Background & Related Work

3 · Design Principles

4 · Architecture

5 · Deception Layers

6 · Autonomous Response & Frontier

7 · Threat Intelligence

8 · Evidence Discipline

9 · Governance & Ethics

10 · Status & Limitations

11 · Roadmap

12 · Business Model

13 · Conclusion

References

A · MCP Inventory (42)

B · Metric Definitions

C · Reproducibility

D · Equipment Procurement (附件 4.1)

E · Infinity Loop & Capital

XCOM: An Evidence-Bound, Agent-Native Cyber-Physical Deception Estate

Master Whitepaper & Investor Memorandum v2.0 — architecture, live measurements, commercial engine and frontier R&D roadmap

VERSION	Master Memorandum 2.0 (Part M) + technical paper 1.0 (§1–13, appendices A–E)
DATE	28 October 2026
AUTHOR	Peter W. “Jedi” Oldenburger — CyberSecurityAD BV (xcom.dev · agent-aix.com · cybersecurityad.com), Amsterdam
SNAPSHOT BASIS	Estate snapshot <code>2026-10-28_121035Z</code> (7th generation, 17 artifacts); ClickHouse live corpus query <code>2026-10-28T12:40Z</code>
STATUS TAGS	LIVE production · BETA design-partner use · R&D bench-validated · CONCEPT designed, not yet built
GRANT ANNEX	附件 4.1 — Equipment Procurement List, XCOM–Shenzhen Cyber-Physical Security Joint Laboratory: ¥4,800,000 total (¥1,500,000 municipal + ¥3,300,000 self-raised), ≈78% Shenzhen/CN bill-of-materials — reproduced in full in Appendix D , inside the ¥18,000,000 programme envelope reconciled in Appendix E.5
CAPITAL	€1,000,000 Seed Preferred at €4,000,000 pre-money — use of proceeds, matching-funds covenant and grant reconciliation in Appendix E.4–E.5
CONTACT	jedi@xcom.dev

1,219,833

CORPUS HITS

45,379

EXTERNAL REAL

99

ATTACKER SESSIONS

42

MCP ENDPOINTS

17,097

CREDENTIALS CAPTURED

49

DOMAINS

M The Cyber-Physical Deception & Defense Estate

This part is the strategic layer: one narrative for the seed investor, the Shenzhen municipal committee and the enterprise CISO. The technical paper (§1-13) is the evidential layer it summarises, and appendices A-E hold the machine-readable detail. Where the two ever disagree, the technical paper wins — a memorandum may compress, it may not inflate.

DOCUMENT	Master Whitepaper & Investor Memorandum — v2.0
LEGAL ENTITY	XCOM.DEV · KVK 99492334 · BTW NL005391698B11 · Sint Olofssteeg 4C, 1012 AK Amsterdam
PRINCIPAL	Peter W. “Jedi” Oldenburger — CEO & Founder
ECOSYSTEM BRANDS	xcom.dev · webcompliance.app · webcompliance.nl · agent-aix.com · rechtspraak-ai.nl
SNAPSHOT BASIS	Estate 2026-10-28_121035Z · ClickHouse 2026-10-28T12:40Z
AUDIENCES	Seed investors (EU/CN) · Shenzhen Municipal AI Grant Committee · enterprise CISOs

Reading paths

Four audiences, four shortest routes through this document. Each path is self-contained.

FOR VCS Architecture → Loop → Revenue → Capital §M.1, §M.2, §M.6, §M.7 then Appendix E — term sheet and the nesting proof	FOR SHENZHEN Frontier R&D → Lab → BOM §M.4, §M.7 then Appendix D — ¥4.8 M equipment, 采购方式, local content	FOR CISOs Deception layers → Loop → Pricing §M.3, §M.6 then §5 and §8 — what is proven, and how it is proven
---	--	--

Evidence discipline → Packs

§M.5 then §8 — provenance chain, redaction manifests, chain of custody

M.1 Executive summary

Modern cybersecurity is trapped in a dichotomy: deception platforms are static and log-heavy, while AI security vendors make unauditable, overblown claims. **XCOM.DEV** resolves this by operating a dual-engine cyber-physical defense platform.

The first engine is the **XCOM Core Estate**: a live, agent-native deception mesh processing **1.21 million cumulative attack-corpus hits** across 49 domains and 106 physical trap points. A **42-server Model Context Protocol (MCP) capability mesh** and hybrid AI routing classify threats against MITRE ATT&CK at ingestion, and execute autonomous response down to the packet layer (eBPF/XDP) and the silicon layer (voltage/thermal anti-tamper).

The second engine is **webcompliance.app**, the B2B SaaS GTM platform that deploys micro-honeypots globally and creates the **Infinity Loop** data moat that continuously trains the core AI. The evidence-bound telemetry discipline is productionised through **rechtspraak-ai.nl**, providing court-admissible, tamper-evident digital evidence.

To scale the cyber-physical frontier (RF, powerline, silicon), XCOM.DEV is raising a **€1.0 M Seed Round** and applying for the **Shenzhen Municipal AI R&D Grant (¥6.0 M)**, backed by **¥12.0 M** in self-raised matching funds, to establish the XCOM-Shenzhen Cyber-Physical Joint Lab.

M.2 The core architecture: an agent-native mesh LIVE

XCOM is not a software application; it is a queryable, autonomous nervous system.

M.2.1 · The 42-server MCP capability mesh. All sensing, analysis and action capabilities are decoupled into 42 Model Context Protocol servers, so autonomous agents (Agent-AIX) swap models and tools without rewriting core logic.

- **Sensing:** xcom-honeypot-archive, xcom-bluetooth, xcom-devtools, xcom-sdr R&D
- **Analysis:** xcom-mitre, xcom-analytics, xcom-intel, xcom-evidence
- **Action:** xcom-kali, xcom-playwright, xcom-firewall-intelligence, xcom-cloudflare

M.2.2 · Hybrid AI routing and compute fabric. Fleet-scale autonomy is economically infeasible on premium gated APIs alone, so tasks route by cost/privacy profile:

LANE	HOST / MODEL	TASK CLASS
Local / air-gapped	MI ARM64 · Ollama (Qwen / DeepSeek)	PII and hardware telemetry — never leaves local silicon
Volume / edge	OMEN 24-core · DeepSeek-V4	High-throughput triage and exploit drafting
Long-context / cloud	Kimi K3 · Qwen3.8-Max	Strategic threat-intel fusion
Gated frontier	Daybreak-class (authorized)	Edge-case red teaming under attestation

M.2.3 · Live corpus and MITRE classification. Queried from the native ClickHouse corpus — no estimation, no cache:

- **Total hits:** 1,219,833 (`atk_events_v1` + native live)
- **External real traffic:** 45,379 (synthetic and self-loop traffic is *labelled, not filtered*)
- **Attacker sessions reconstructed:** 99 across 7 attribution dimensions
- **Top techniques:** T1595 Active Scanning 65k · T1105 Ingress Tool Transfer 8.9k · T1190 Exploit Public-Facing App 5.7k

M.3 Deception layers and the Infinity Loop

M.3.1 · Web, physical and peripheral traps. 49 domains and 106 physical trap points (NFC/RFID, QR-mediated routes) absorb credential stuffing (17,097 redacted credentials) and PoC payloads (92 retained). Browser-mediated Web Bluetooth and camera *permission* traps are **proven at telemetry / permission-outcome level only** — no biometric or hardware-control claim is made anywhere in this document.

M.3.2 · The Infinity Loop (`webcompliance.app` international, `webcompliance.nl` local). Traditional honeypots need heavy enterprise sales; XCOM scales its sensor network through SaaS:

1. **Deploy** — companies monitor domain/email security and deploy XCOM micro-honeypots into their cloud perimeter.
2. **Ingest** — attack data flows into the central XCOM ClickHouse corpus.
3. **Immunise** — Agent-AIX analyses the global corpus and pushes real-time, personalised threat intelligence back to the tenant dashboard.

The result is a self-funding, globally distributed data moat in which every new B2B customer makes the core AI measurably smarter. The loop, its ten stages and its return edges are drawn out in [Appendix E.1](#).

M.4 Frontier R&D: the cyber-physical battlefield R&D

To defend critical infrastructure, XCOM extends its agents from the application layer down to the physical and silicon layers. This hardware programme is the core of the Shenzhen Joint Lab initiative.

M.4.1 • Network: eBPF/XDP AI takeover. Agents hook the eXpress Data Path for nanosecond-latency packet decisions: a payload classified as C2 is silently dropped at the NIC (`XDP_DROP`), or answered with synthetic TCP ACKs that trap the attacker in a fabricated session without provisioning infrastructure.

M.4.2 • RF spectrum: SDR / HackRF. Via `xcom-sdr`, agents drive software-defined radios to broadcast phantom BLE beacons and rogue Wi-Fi APs, exhausting attacker surveillance tooling, and detect/simulate GPS/PTP timing drift to test critical-infrastructure resilience against spoofing.

M.4.3 • Powerline: G3-PLC air-gap infiltration. Narrowband coupling over AC mains lets edge agents form a covert mesh inside RF-jammed or air-gapped premises: one sensor in a lobby socket can move bounded telemetry out of a restricted server room over the building's own wiring.

M.4.4 • Silicon anti-tamper.

- **Voltage fault-injection** — agents orchestrate precision VCC droops (ChipWhisperer-class) to audit our own edge nodes' secure-boot chains, proving supply-chain integrity rather than asserting it.
- **Kinetic anti-tamper ("Scorched Earth")** — on chassis intrusion, an edge node overrides thermal limits, disables PWM fans via IPMI and runs sustained tensor-matrix loops to destroy silicon-resident keys before extraction. Sacrificial nodes only, under policy approval ([§6.4](#)).

M.5 Evidence discipline and legal alignment (`rechtspraak-ai.nl`)

XCOM's core differentiator is **evidence-bound telemetry**. We publish our own limitations — 47.8% knowledge coverage, 20/29 edge-local parity — because a defense platform that cannot audit itself cannot audit anything.

That discipline is productionised through **rechtspraak-ai.nl**:

- **Immutable snapshots** — every event is hash-addressed and propagated through the pipeline with an immutable ID (§8.1).
- **Court-admissible** — evidence packs for Dutch proceedings and cyber-insurers use the same redaction manifests and signing discipline as the research estate.
- **Privacy by design** — credentials and PII redacted at rest; no biometric templates or camera frames by default, aligning GDPR with PIPL.

Status honesty. rechtspraak-ai.nl is R&D, not a shipped product: the evidentiary mechanism (§8) is live and measured, the judicial packaging around it is being built. No court has yet accepted an XCOM pack — that is the milestone, not the current state.

M.6 Commercialisation and revenue streams

PRODUCT LINE	DESCRIPTION	TARGET MARKET	ANCHOR
Webcompliance.app	B2B SaaS for compliance monitoring and distributed micro-honeypot deployment — drives the Infinity Loop	Global SMB / enterprise IT	€29–€299/mo
Agent-AIX Red Team	Scope-bounded, autonomous AI pentesting under written authorization	Enterprise / government	€18k/engagement
Intel Feeds	Anonymised, MITRE-mapped attack taxonomy from the 1.2 M+ hit corpus	MSSPs / SOCs	€2k/mo
Evidence Packs	Replayable, court-ready export bundles via rechtspraak-ai.nl	Legal counsel / insurers	€6k/pack

Two anchors, one line — reconciled. The technical paper (§12) prices the *managed* variant as **XCOM Grid · €490/sensor/mo**; this memorandum prices the *self-serve* variant at **€29–€299/mo**. Both are the same revenue line at different service depths (self-serve SaaS vs. fully managed estate), not two competing numbers. Neither figure is withdrawn.

M.7 Strategic expansion: the Shenzhen Joint Lab

Western cybersecurity is trapped in the cloud. XCOM is building a **cyber-physical mesh**: the 42-server software brain combined with Shenzhen-manufactured SDR, PLC and voltage-glitching hardware produces sovereign edge agents capable of physically defending restricted premises.

M.7.1 · Municipal grant application.

- **Project:** XCOM-Shenzhen Cyber-Physical Security Joint Lab (深圳信息物理安全联合实验室)
- **Requested:** ¥6,000,000 municipal AI R&D grant
- **Matching:** ¥12,000,000 self-raised (seed round + corporate reserves)
- **Compliance:** 100% data residency in the Shenzhen AZ (PIPL); offensive-adjacent tools restricted to owned/authorized systems; ≥70% local hardware BOM — actually ≈78% ([Appendix D](#))

M.7.2 · The €1.0 M Seed Round (term-sheet summary).

- **Issuer:** XCOM.DEV, Amsterdam, NL
- **Valuation:** €4 M pre-money / €5 M post-money
- **Use of proceeds:** 40% GTM (webcompliance scaling) · 30% core R&D · 30% Shenzhen WFOE capitalisation
- **Covenant:** seed proceeds + corporate reserves fulfil the ¥12 M matching requirement, with an FX-shortfall make-whole guaranteed by the founder

The full term sheet, the matching-funds covenant and the arithmetic that nests ¥4.8 M equipment inside the ¥18 M programme envelope are in [Appendix E.4–E.5](#).

M.8 24-month roadmap and KPI commitments

PHASE A · MONTHS 0–6 · JOINT LAB

Stand up the entity and the testbeds

- WFOE registration complete
- Edge/local parity **29/29**
- Knowledge coverage ≥ **75%**
- SDR / PLC / glitch testbeds operational

PHASE B · MONTHS 6–12 · COMPLIANCE

Prove residency and latency

- Shenzhen AZ live · PIPL audit passed
- CVE-to-lure median < **2 h**
- XDP authorized-surface rollout
- Coverage ≥ **85%**

Institutionalise the capability

- ≥ 2 patents filed
- ≥ 2 standards proposals
- ≥ 15 Shenzhen jobs created
- Coverage $\geq 90\%$

M.9 Conclusion

XCOM.DEV demonstrates that a deception estate can be simultaneously **autonomous** (reasoning over a 42-server MCP mesh), **evidence-bound** (replayable from immutable snapshots) and **cyber-physical** (extending from browser lures to packets, spectrum, powerlines and silicon). With the legal foundation in Amsterdam, the GTM engine scaling through `webcompliance.app`, the judicial pipeline maturing in `rechtspraak-ai.nl` and the hardware frontier launching in Shenzhen, XCOM is positioned to define the next decade of autonomous cyber defense.

Contact & data-room access: Peter W. Oldenburger, CEO — jedi@xcom.dev · +31 6 48 31 91 57 · xcom.dev · webcompliance.app · agent-aix.com · rechtspraak-ai.nl

END OF PART M · BELOW: THE TECHNICAL PAPER (V1.0) WITH ITS LIVE MEASUREMENTS AND APPENDICES

Abstract

Modern deception technology records packets but does not reason; modern security AI makes claims that cannot be audited. This paper presents **XCOM**, a live, agent-native deception estate that unifies (i) browser-mediated, physical, and — in research phases — radio-frequency and powerline deception surfaces; (ii) a **42-server Model Context Protocol (MCP) capability mesh** that lets autonomous agents query, correlate, and act over the entire estate; and (iii) an **evidence-bound telemetry discipline** in which every published number is traceable to an immutable snapshot artifact and every capability claim is stated at the narrowest level the evidence supports.

As of 28 October 2026, XCOM has processed **1,219,833 corpus hits** (1,129,971 `atk_events_v1` + 89,862 native live), of which **45,379 are externally-sourced real interactions**, reconstructed **99 attacker sessions** across seven attribution dimensions, captured **17,097 credentials** (redacted at rest), **92 proof-of-concept payloads**, and **118 AI prompt-injection attempts**, across **49 domains**, **106 physical trap points**, and **3 sensors**, ingesting **141 GB across 247 datasets** into **69 live datastores**. Attacks are classified live against MITRE ATT&CK directly from the native ClickHouse corpus — no estimates, no cache.

We describe the architecture in full, publish our own limitations (**47.8%** knowledge coverage; edge/local parity at **20/29** with a committed remediation date; peripheral traps proven at telemetry level only), and define a four-phase frontier roadmap extending autonomous defense from the application layer to the packet layer (eBPF/XDP), the electromagnetic spectrum (SDR/HackRF), the electrical grid (G3-PLC powerline agent deployment), and the silicon itself (voltage fault-injection auditing and kinetic thermal zeroization).

DOCTRINE • ORIGINAL DUTCH

“Onze honeypot-sensoren staan bewust open. Aanvallen worden niet geblokkeerd maar vastgelegd, geclassificeerd en gekoppeld aan MITRE ATT&CK. De cijfers hieronder komen rechtstreeks uit de native ClickHouse-corpus (geen schatting, geen cache).”

TRANSLATION

Our honeypot sensors are deliberately open. Attacks are not blocked but recorded, classified, and linked to MITRE ATT&CK. The figures come directly from the native ClickHouse corpus — no estimation, no cache.

1 Introduction and Motivation

1.1 The failure modes of current deception technology

Three structural failures define the current market:

1. **Static decoys.** Conventional honeypots emulate services but cannot form hypotheses, adapt lures, or produce court-admissible evidence chains. Telemetry frequently exists as alerts without durable, replayable records — a failure mode we measured inside our own pipeline and engineered against (§8).

2. Unauditable AI claims. Vendors routinely assert “hardware control,” “autonomous response,” or “AGI-class reasoning” with no verifiable artifacts. Buyers cannot distinguish capability from marketing. The industry’s credibility deficit is now the primary obstacle to enterprise adoption of AI security.

3. Broken agent economics. Fleet-scale autonomous defense is economically infeasible on premium gated API pricing alone. The 2026 open-weights frontier (Qwen, DeepSeek, Kimi families) changes this by one to two orders of magnitude per token, making continuous agent operation affordable at estate scale.

Simultaneously, the offense has industrialized. XCOM’s own surfaces have absorbed **118 AI prompt-injection attempts** and **17,097 credential submissions**, and the emergence of gated frontier cyber models (e.g., the GPT-5.6-Cyber / Daybreak class, August 2026) compresses exploit timelines from weeks to hours. Defenders need systems that *reason over evidence* and *act within authorized scope* – and that can prove both.

1.2 Contributions

- **C1 – An agent-native deception estate architecture** built on a 42-endpoint MCP capability mesh over a hybrid edge/local compute fabric (§4).
- **C2 – A live, MITRE-mapped attack corpus** with published counts, provenance, and exclusion rules (§7).
- **C3 – An evidence-bound claim discipline:** a formal practice of stating capabilities at the narrowest evidentially supported level, with published gaps, parity harnesses, and replayable snapshots (§8).
- **C4 – A cyber-physical frontier roadmap** spanning packet-layer AI takeover (eBPF/XDP), RF spectrum defense (SDR), powerline agent infiltration (PLC), and silicon-level anti-tamper (voltage glitch auditing, kinetic thermal zeroization) (§6).
- **C5 – A privacy-first trap design** for browser-mediated peripheral deception, including a signed-challenge protocol that eliminates client-forgable permission assertions (§5.2, §8.5).

1.3 Non-goals and scope

XCOM is a **defensive research and authorized-testing platform**. All offensive-adjacent capabilities described herein operate only on systems owned by CyberSecurityAD or under explicit written authorization. Nothing in this paper describes attacks on third-party systems. Where a capability is bench-validated only, it is tagged **R&D**; where designed but unbuilt,

CONCEPT.

2 Background and Related Work

Honeypots and deception. From early honeypot research to modern deception platforms, the field has converged on emulated services plus alerting. XCOM departs by treating the estate as a *queryable graph* rather than a log sink, and by adding browser-mediated peripheral traps (§5.2) and physical trap points (§5.3).

MITRE ATT&CK. Technique-level classification is industry standard; XCOM's distinction is that classification is performed *at ingestion* against the native ClickHouse corpus and exposed live, with self-loop and synthetic traffic **labeled rather than filtered**, preserving denominator honesty (§7.1).

Model Context Protocol (MCP). XCOM adopts MCP as its universal capability contract: every sensor, datastore, tool, and orchestrator is an MCP server with typed tools. This yields swappable models, auditable tool calls, and a topology graph that is itself a security asset (§4.2).

Gated frontier cyber models. The 2026 introduction of trusted-access cyber tiers (e.g., Daybreak Blue/Red) validates the thesis that reduced-refusal cyber capability requires identity verification, monitoring, and legal attestation. XCOM's internal governance mirrors and extends this model to autonomous agents (§9).

Fault injection and anti-tamper. Voltage glitching (ChipWhisperer-class) and thermal destruction are established in hardware security research and military zeroization practice respectively. XCOM's contribution is *agent-orchestrated* execution of both as continuous self-audit and last-resort defense (§6.3–6.4).

Powerline communication. ITU-T G.9903 (G3-PLC) and broadband PLC (HomePlug AV2-class) are mature for metering. Their use as an *agent deployment and exfiltration channel into RF-denied premises* is, to our knowledge, unexplored in commercial defense products (§6.5).

3 Design Principles

1. **Sensors stay open.** Deception surfaces do not block; they record. Blocking occurs only on authorized defensive surfaces (§6.2), never on research lures.
2. **Evidence over assertion.** Every published metric maps to a snapshot artifact. Every capability claim is stated at the narrowest level the evidence supports.

- 3. **Label, don't filter.** Synthetic, self-loop, and validation traffic are tagged (`traffic_class`) and excluded from threat KPIs — but never deleted, preserving reproducibility.
- 4. **Agents are untrusted callers.** All agent tool calls pass typed validation; all ingested attacker-controlled fields are treated as hostile input everywhere downstream (§8.5).
- 5. **Privacy by construction.** No biometric templates, no camera frames, no Bluetooth device enumeration by default; permission *outcomes* only, under signed challenge (§5.2).
- 6. **Publish the gaps.** Coverage ledgers, parity divergence, and guard failures are first-class dashboard citizens and appear in this paper unedited (§8.3, §10).

4 Architecture

4.1 Compute fabric LIVE

HOST	PLATFORM	ROLE
laptop	Linux x86_64 · 12 CPU · 30.7 GB · Docker 29.8.0	Primary ingestion, SSH jump, agent venv host
m1	Darwin arm64 · 8 CPU · 16 GB	Local AI inference (Ollama <code>:11434</code>), privacy-sensitive analysis
omen	Linux x86_64 · 24 CPU · 31.1 GB · Docker 29.8.0	MCP server farm, heavy simulation, ClickHouse

Edge layer: Cloudflare Workers + Caddy reverse proxies + tunnels across **49 hostnames / 4 origins** (caddy 9, tunnel 5, worker 2, none 33), absorbing **12,243 edge requests** in the snapshot window. **78 containers** and **1,652 discovered host services** are continuously inventoried.

4.2 The MCP capability mesh LIVE

42 registered MCP endpoints (35 first-party `xcom-*`, 7 third-party: `serena`, `prometheus`, `semgrep`, `wolfram`, `chrome-devtools`, `cloudflare-api`, `chromium-control-canvas`). Representative inventory (full list in **Appendix A**):

- **Sensing:** `xcom-honeypot-archive`, `xcom-bluetooth`, `xcom-devtools`, `xcom-sdr` R&D
- **Analysis:** `xcom-mitre`, `xcom-analytics`, `xcom-intel`, `xcom-evidence`, `xcom-bench`

- **Action:** xcom-kali , xcom-playwright , xcom-command , xcom-firewall-intelligence , xcom-cloudflare
- **Memory/state:** xcom-memory , xcom-memory-journal , xcom-data , xcom-sqlite , xcom-context7
- **Engineering:** xcom-code , xcom-git , xcom-compute , xcom-jupyter , xcom-lang , xcom-design , xcom-ops , xcom-ingest , xcom-enigma , xcom-firmware , xcom-dod , xcom-forum-training , xcom-web-knowledge

The mesh forms a **typed topology graph of 296 nodes and 179 edges across 7 edge kinds:**

mcp→host:port (62), dbcode→store (38), caddy→upstream (36), pipeline→dest (13), pipeline→source (13), worker-route→script (12), tunnel→origin (5), plus 18 datastore relationships. This graph is queryable by agents and diffed over time — topology drift is itself an alert.

4.3 Ingestion pipeline and data lake LIVE

247 datasets · 141 GB ingested · 88.1 GB live, by class:

CLASS	DATASETS	VOLUME
Telemetry	109	64.4 GB
Knowledge	22	39.2 GB
Other	55	35.9 GB
Training corpora	30	451 MB
Threat intel	22	35.4 MB
Forensic evidence	3	679 MB
Runtime	6	1.7 KB

Storage tiers: **ClickHouse** (live attack corpus, MITRE classification), **SQLite/DuckDB** (estate memory, journals, design snapshots), **Parquet** (weekly event exports, e.g. events-week42.parquet), **D1/object storage** (edge-side), **24 saved exports (11.5 GB)** as replayable evidence bundles (backup 14, dataset 7, db 2, sql 1).

4.4 Parity harness and snapshotting LIVE

A parity harness continuously diffs **Cloudflare edge runtime behavior against local runtime** across 29 probe routes. Current state: **20/29 converged; divergence published** with a committed remediation date of December 2026 (§10, §11). Snapshots are immutable, generationed (7 generations, 17 artifacts), and hash-addressed; every figure in this paper cites snapshot `2026-10-28_121035Z` or the ClickHouse query timestamp.

4.5 Hybrid AI routing LIVE

Agents route tasks by cost/privacy profile:

```
ROUTING = {  
  "pii_or_hardware_telemetry": "ollama/qwen-local (m1, air-gapped path)",  
  "volume_triage":             "deepseek-v4 (open-weights, self-hosted)",  
  "long_context_threat_intel": "kimi-k3 (2M ctx)",  
  "estate_orchestration":      "qwen3.8-max (Alibaba CN region for CN data)",  
  "edge_case_frontier":        "gated frontier tier (authorized, attested)",  
}
```

This yields 10–100× token-cost reduction versus premium-only routing while keeping sensitive classes (peripheral telemetry, credentials context) on local silicon.

5 Deception Layers

5.1 Web and application honeypots LIVE

49 domains expose credential portals, upload endpoints, admin consoles, and AI-facing surfaces. Snapshot counts: **1,114,044 honeypot requests, 283,981 events** (+1,603 v2-schema), **529 uploads, 17,097 credentials (redacted at rest), 118 AI prompt-injection leaks, 99 reconstructed sessions** with 7-dimension attribution (source IP, ASN, country, TLS/JA fingerprint, session graph, payload family, temporal cadence).

5.2 Browser-mediated peripheral traps LIVE TELEMETRY-LEVEL

Peripheral-themed routes exercise Web Bluetooth / media-capture *permission workflows* as lures. Recorded outcomes include `ble_denied (allowed=0)` and `face_commit (allowed=1)` events from the `cybersecurityad-admin` sensor.

Evidence-bound statement (from internal audit, September 2026): these traps are proven as *browser-mediated permission-outcome telemetry*, **not** as Bluetooth device enumeration, camera frame capture, or biometric processing. No images, templates, or GATT data are collected by default. Client-submitted `allowed` flags are **forgeable** and are therefore being replaced by a signed-challenge workflow.

```
# Signed peripheral challenge: bind nonce to sensor, route, session, expiry
payload = f"{sensor}:{route}:{nonce}:{expiry}"
sig = hmac.new(KEY, payload.encode(), hashlib.sha256).hexdigest()
# Ingestion rejects terminal-stage events lacking a valid, unreplayed challenge
```

Workflow stages are modeled separately (page viewed → API available → permission requested → denied/granted → device selected → interaction completed → artifact committed → validation event) and are **never collapsed into a single boolean**.

5.3 Physical traps `LIVE`

106 physical trap points (156 wall-touches recorded): NFC/RFID lures, QR-mediated routes, and wired sensor taps in controlled premises, feeding the same event graph as web telemetry.

5.4 RF spectrum surfaces `R&D`

HackRF-class SDR front-ends (`xcom-sdr`) broadcast phantom BLE beacons, rogue AP signatures, and phantom RFID populations to map proximity-attack tooling; receive-side pipelines detect GPS/PTP spoofing drift and non-IP C2 (LoRa/Zigbee-class). Bench-validated; field deployment gated on authorization review.

5.5 Powerline surfaces `R&D`

G3-PLC narrowband coupling over AC mains provides a covert, RF-independent channel for (a) deploying micro-agents into RF-denied premises and (b) exfiltrating bounded telemetry bursts from inside shielded zones. Bench prototype on isolated circuits; see §6.5.

6 Autonomous Response and Frontier Capabilities

Equipment substrate. Every capability in this section is backed by instrumentation that can be ordered today, not by a slide: broadband SDR platforms and an RF anechoic chamber (§5.4, §6.5), ChipWhisperer fault-injection benches and 8 GHz scopes (§6.3), G3-PLC coupling modules (§6.5), and 100G hardware TAPs for lossless packet mirroring (§6.2). The full bill of materials, with unit prices and local-content percentages, is published as [Appendix D](#).

6.1 Agent-AIX orchestration BETA

Events above threshold become typed agent tasks: hypothesis → tool calls over the MCP mesh → evidence bundle → proposal → (human- or policy-approved) action. All tool calls are journaled via `xcom-memory-journal`; all proposals carry the event IDs they depend on.

6.2 Packet-layer AI takeover R&D

On authorized defensive surfaces, agents hook the kernel data path (eBPF/XDP) for nanosecond-scale decisions: silent drop of classified C2, or synthetic response injection (“honeypot-on-the-wire”) that traps attackers in a fabricated session without provisioning infrastructure.

```
SEC("xdp_ai_takeover")
int xdp_decide(struct xdp_md *ctx) {
    __u32 score = agent_threat_score(ctx);    // shared map, local model
    if (score > 90) return XDP_DROP;          // silent drop + ClickHouse log
    if (score > 50) return XDP_TX;            // inject synthetic response
    return XDP_PASS;
}
```

Scope rule: **never** applied to open research lures (Doctrine §3.1).

6.3 Voltage fault-injection self-audit R&D

Programmable supply glitching (ChipWhisperer-class) is orchestrated by agents to continuously audit secure-boot and trust-anchor paths on XCOM's own edge silicon, proving supply-chain integrity rather than asserting it. Glitch parameters (droop depth, width, clock phase) are searched autonomously and results versioned in `xcom-firmware`.

6.4 Kinetic anti-tamper: thermal zeroization CONCEPT → R&D

On verified chassis intrusion or unauthorized BMC access on sacrificial/edge nodes, the agent executes a **physical zeroization** sequence: volatile key wipe → fan-curve disable via IPMI → power-limit and clock maximization → sustained tensor-core load until thermal trip, destroying silicon-resident key material. Peacetime use of the same engine: autonomous stress-testing of datacenter cooling failsafes.

```
memwipe --secure /dev/mem
nvidia-smi -pl 600 && nvidia-smi -gpc 100
ipmitool raw 0x30 0x45 0x01 0x00 # fans off (authorized sacrificial node only)
./thermal_runaway & curl -X POST $MESH -d "NODE_ZEROIZED_THERMAL"
```

6.5 Powerline agent infiltration R&D

A lobby-plugged sensor establishes a building-wide agent mesh over mains wiring, reaching RF-jammed or air-gapped rooms without new cabling. Exfiltration is bounded, encrypted, and journal-audited; deployment is limited to CyberSecurityAD premises and authorized pilot sites.

7 Threat Intelligence and Classification

7.1 Live MITRE ATT&CK corpus LIVE

Queried directly from native ClickHouse at 2026-10-28T12:40Z — no estimation, no cache:

METRIC	VALUE
Total hits	1,219,833 (atk_events_v1 1,129,971 + native live 89,862)
Last 24 h	3,515 (106 in the final hour)
External real	45,379 (self-loop & synthetic <i>labeled, not filtered</i>)
Frozen archive	2,213,896 rows (export 2026-10-18 06:21)

Observed techniques (top six):

TECHNIQUE	NAME	COUNT
T1595	Active Scanning	65,374
T1105	Ingress Tool Transfer	8,960
T1552	Unsecured Credentials	5,743
T1190	Exploit Public-Facing Application	5,681
T1110	Brute Force	908
T1592	Gather Victim Host Information	687

7.2 CVE and forum ingestion pipeline LIVE

Six monitored forums + NVD/vendor feeds → xcom-forum-training / xcom-intel scrape → xcom-mitre ATT&CK enrichment → xcom-kali exploit-availability validation → trap-config regeneration across all 49 domains. Target and measured median: **< 4 h from CVE publication to deployed lure**. 22 intel datasets (35.4 MB) and 30 training corpora (451 MB) currently retained.

7.3 Session reconstruction and the BLE chain LIVE

Example correlated chain: external scan (T1595) → peripheral lure request → ble_denied → credential submission (T1110) → agent task → Telegram alert + evidence export + ATT&CK enrichment. End-to-end autonomous latency: **< 15 s** from first probe to playbook proposal. 92 PoC payloads retained (curl 19, unknown-origin 73) as a full corpus in snapshot.

8 Evidence Discipline and Auditability

8.1 Provenance chain

One immutable event ID is assigned at ingestion and propagated Worker → datastore → stream → alert → export → agent task. Alerts lacking a durable record raise a pipeline-integrity alarm (this check was added after we observed notification/storage fragmentation in our own fleet).

8.2 Validation-traffic separation

QA and deployment-validation traffic is signed at source and stored with

`traffic_class=validation`; it is **excluded from all threat KPIs** and retained in a separate defensive-validation view. The September 2026 peripheral records (`ble_denied` , `face_commit`) are classified as validation traffic and are reported as such everywhere, including this paper.

8.3 Coverage and parity ledgers

Knowledge coverage: **119/249 declared sources (47.8%)**, with 4 missing, 35 stale, 63 partial, and 2,555 undeclared-in-store objects tracked publicly in-product. Parity: **20/29** edge/local convergence, divergence itemized. Both ledgers are investor- and auditor-visible by default.

8.4 Case study: the peripheral audit

An internal audit concluded that BLE/camera integration was proven at *telemetry* level only, that no hardware control or biometric capture existed, and that the correct public claim was therefore narrower than the exciting one. We published the narrower claim. This case is the canonical example of Doctrine §3.2 and is included in the data room in full.

8.5 Platform threat model (ourselves)

- **Attacker-controlled fields** (`kind` , `tag` , `email` , UA, device names) are untrusted everywhere: allowlisted types, parameterized queries, HTML/Markdown escaping per alert target, no shell interpolation, and explicit untrusted-data marking before any model or agent sees them.
 - **Privacy boundary**: permission outcomes and capability metadata only by default; any future frame/identifier capture requires documented purpose, lawful basis, retention limit, and access trail.
 - **Consent and context**: peripheral traps require user gesture + browser-native prompt, HTTPS-only, restrictive `Permissions-Policy` , default-deny on non-trap routes.
 - **Replay/forgery**: signed challenges with nonce replay caches; terminal-stage client assertions are never labeled as verified hardware access.
-

9 Governance, Privacy, and Ethics

- **Jurisdictions:** GDPR (EU operations); PIPL-aligned data-residency design for the planned CN region (Alibaba Cloud, CN-region routing for CN-sourced data).
 - **Authorization:** offensive-adjacent tooling executes only on owned or explicitly authorized systems; Daybreak-class gated models are used under their attestation regimes where applicable.
 - **Disclosure:** vulnerabilities in third-party software observed via traps are reported through coordinated disclosure; our own platform findings (including the peripheral audit) are published internally first and to partners under NDA.
 - **Human oversight:** autonomous actions are bounded by scoped permission profiles; kinetic and packet-layer capabilities require policy approval tiers above standard agent autonomy.
 - **Redaction:** credentials and PII redacted at rest; evidence exports carry redaction manifests.
-

10 Current Status and Honest Limitations

DIMENSION	STATUS
Corpus & classification	LIVE 1.22M hits, MITRE-mapped at ingestion
Web/physical deception	LIVE 49 domains, 106 trap points
Peripheral traps	LIVE at telemetry level; signed-challenge remediation in flight
Agent orchestration	BETA journaled, proposal-gated
XDP packet takeover	R&D bench; authorized-surface rollout Q2 2027
SDR / PLC	R&D bench prototypes
Voltage audit / thermal zeroization	R&D → CONCEPT ; sacrificial-node prototypes 2027 H2
Knowledge coverage	47.8% (target 75% Q4 2026, ≥90% 2027 H2)
Edge/local parity	20/29 (target 29/29 by Dec 2026)
Secrets management	Vault rollout in progress (currently flagged unconfigured in estate health)

We publish these limitations because a defense platform that cannot audit itself cannot audit anything.

11 Roadmap

- **2026 Q4:** parity 29/29; vault rollout; coverage → 75%; signed peripheral challenges live; 3 design partners; Shenzhen grant application filed.
- **2026 Q4 (grant track):** joint-lab equipment bill of materials locked at ¥4.8M with ≈78% Shenzhen/CN content — the physical substrate for every §6 capability, itemised in [Appendix D](#).
- **2027 Q1:** Agent-AIX beta GA; Evidence Packs productized; SDR field pilot (authorized sites); PLC bench-to-premises prototype.

- **2027 Q2:** CN-region deployment (PIPL-aligned); XDP takeover on authorized defensive surfaces; CVE→lure median < 2 h.
- **2027 H2:** voltage self-audit rig continuous operation; thermal zeroization prototype on sacrificial nodes; coverage ≥ 90%; Series A trigger (10 paying grids or €1M ARR-equivalent).

12 **Adoption and Business Model** (summary)

LINE	DESCRIPTION	ANCHOR
XCOM Grid	Managed deception estate (sensors, domains, traps, dashboards)	€490/sensor/mo
Agent-AIX Red Team	Scope-bounded autonomous pentest under written authorization	€18k/engagement
Intel Feeds	Anonymized ATT&CK-mapped taxonomy from live corpus	€2k/mo
Evidence Packs	Replayable, redacted, audit-ready export bundles	€6k/pack

Anchors validate with design partners Q4 2026. Non-dilutive funding pipeline: Shenzhen municipal AI grant; EU EIC Pathfinder.

Dual-engine structure. The four lines above split into a volume GTM engine (`webcompliance.app` , tenant micro-traps feeding the corpus) and a depth deep-tech engine (Agent-AIX, Intel Feeds, Evidence Packs via `rechtspraak-ai.nl`). The flywheel that connects them, the €1,000,000 Seed Preferred term sheet (€4 M pre / €5 M post) and the reconciliation between the seed round, the ¥6 M municipal request and the ¥4.8 M equipment subject are set out in [Appendix E](#).

13 **Conclusion**

XCOM demonstrates that a deception estate can be simultaneously **autonomous** (agents reason and act over a 42-server capability mesh), **evidence-bound** (every number replayable from snapshot; every claim stated at its proven level), and **cyber-physical** (extending from

browser lures to packets, spectrum, powerlines, and silicon). The estate's published imperfections — 47.8% coverage, 20/29 parity, telemetry-level peripheral proof — are not weaknesses of the platform but demonstrations of its core discipline: in an industry saturated with unverifiable claims, the scarcest and most valuable product is *trustworthy measurement*.

References

- [1] MITRE ATT&CK® knowledge base, technique taxonomy (T1595, T1105, T1552, T1190, T1110, T1592), accessed Oct 2026.
- [2] Model Context Protocol specification, 2025 ed.
- [3] ClickHouse documentation: real-time analytical corpus design.
- [4] Høiland-Jørgensen et al., *The eXpress Data Path: Fast and Flexible Packet Processing in the Linux Kernel* (XDP/eBPF).
- [5] ChipWhisperer fault-injection framework documentation (voltage glitching).
- [6] ITU-T G.9903 (G3-PLC narrowband powerline communication); HomePlug AV2 specification.
- [7] Great Scott Gadgets HackRF One documentation; GNU Radio project.
- [8] OpenAI, *Expanding Daybreak as the Cyber Defense Window Narrows*, 10 Aug 2026 (gated cyber-model access regimes).
- [9] Regulation (EU) 2016/679 (GDPR); PRC Personal Information Protection Law (PIPL).
- [10] Internal: XCOM peripheral-trap audit verdict, Sep 2026; estate snapshot 2026-10-28_121035Z ; ClickHouse corpus query 2026-10-28T12:40Z .
- [11] Internal: XCOM-Shenzhen Cyber-Physical Security Joint Laboratory, 附件 4.1 — Equipment Procurement List (设备购置清单), Oct 2026 — reproduced with unit prices in Appendix D.

A MCP Endpoint Inventory (42)

oraios/serena

prometheus

semgrep

wolfram

chrome-devtools-omen

cloudflare-api

chromium-control-canvas

xcom-analytics

xcom-bench

xcom-bluetooth

xcom-cloudflare

xcom-code

xcom-command

xcom-compute

xcom-context7

xcom-data

xcom-design

xcom-devtools

xcom-dod

xcom-enigma

xcom-evidence

xcom-firewall-intelligence

xcom-firmware

[xcom-forum-training](#)[xcom-git](#)[xcom-honeypot-archive](#)[xcom-ingest](#)[xcom-intel](#)[xcom-jupyter](#)[xcom-kali](#)[xcom-lang](#)[xcom-memory](#)[xcom-memory-journal](#)[xcom-mitre](#)[xcom-ops](#)[xcom-playwright](#)[xcom-sdr · R&D](#)[xcom-sqlite](#)[xcom-web-knowledge](#)[+3 provisioning endpoints](#)

B Metric Definitions

- **Corpus hit:** one row in `atk_events_v1` or the native live table; includes labeled synthetic/self-loop rows.
- **External real:** hits whose source is neither estate-owned nor labeled synthetic/self-loop.
- **Honeypot request:** HTTP request logged in `honeypot_logs` (snapshot count 1,114,044; corpus total differs by schema window — both are published).
- **Session:** attacker interaction cluster over the 7 attribution dimensions.
- **Wall-touch:** one physical interaction with a trap point.
- **Coverage:** declared sources with `have` status ÷ declared sources (119/249).

C Reproducibility

All figures cite snapshot `2026-10-28_121035Z` (artifacts: 17 files incl. `events-week42.parquet`, `xcom-memory-journal-snapshot.sqlite`, `xcom-design-snapshot.duckdb`) or ClickHouse query `2026-10-28T12:40Z`. Export bundle hashes and the parity-harness repository are available in the data room under NDA.

D Equipment Procurement List — 设备购置清单 (附件 4.1)

Grant auditors scrutinise the equipment annex harder than any other document: if the hardware does not match the technical claims in the paper, the application is flagged for *budget inflation* (预算虚高). This appendix therefore publishes the complete bill of materials that underwrites §6, at unit-price granularity, with the local-content percentage of every line.

PROJECT	XCOM — Shenzhen Cyber-Physical Security Joint Lab 深圳信息物理安全联合实验室
BUDGET SUBJECT	设备费 · Equipment
FUNDING SPLIT	Municipal fund requested ¥1,500,000 (150万元) + self-raised ¥3,300,000 (330万元)
TOTAL	¥4,800,000 (480万元) · 12 line items · 116 units
LOCAL CONTENT	Composite Shenzhen / CN bill-of-materials ≈78% (annex commitment: ≥70%)

#	EQUIPMENT · 设备名称	CORE SPECS / REFERENCE MODEL · 核心技术规格	QTY	UNIT ¥	TOTAL ¥	JUSTIFICATION & USE CASE · 购置理由	LOCAL BOM
一	Software-Defined Radio & RF Testbeds 软件无线电与射频测试床					540,000	
1.1	High-end wideband SDR platform 高端宽带 SDR 平台	Ettus USRP N320 / X310 · 1 MHz-6 GHz · 200 MHz instantaneous BW	2	85,000	170,000	Core RF deception node: GPS/PTP time-spoofing defence validation and wideband C2 signal capture. RF 欺骗核心节点	<10% · import
1.2	Edge SDR decoy nodes 边缘 SDR 诱饵节点	HackRF One / BladeRF 2.0 · bulk deployment at physical trap points	30	3,000	90,000	Generate large populations of phantom BLE / Wi-Fi fingerprints to burn attacker reconnaissance budget. 耗尽攻击者侦察资源	90% · SZ assembly
1.3	RF anechoic chamber + antenna array 射频微波暗室及天线阵列	Custom 3 m × 3 m chamber · wideband log- periodic (LPDA) array	1 set	280,000	280,000	Isolate external interference; compliant RF immunity testing and controlled signal injection. 隔离外部干扰的合规测试环境	100% · SZ custom
二	Hardware Fault Injection & Anti-Tamper Validation 硬件故障注入与防篡改验证					1,000,000	
2.1	Hardware fault-injection platform 硬件故障注入平台	NewAE ChipWhisperer- Husky + ChipWhisperer- Pro kits	3 sets	60,000	180,000	Validate edge-node secure-boot chains and trust anchors; emulate VCC voltage-glitch attacks (§6.3). 安全启动链验证	<20% · import
2.2	High-bandwidth oscilloscope 高带宽示波器	Keysight Infiniium UXR / DSOX · 8 GHz · 25 GS/s	2	250,000	500,000	Side-channel analysis; capture nanosecond voltage droop and electromagnetic leakage. 侧信道与纳秒级毛刺捕获	<10% · import
2.3	Programmable PSU + thermal-shock chamber 可编程电源与热冲击试验箱	High-precision linear supply · thermal/humidity cycling chamber, -40 °C to 150 °C	2 sets	160,000	320,000	Execute “Scorched Earth” thermal-runaway tests and silicon thermal- stress validation (§6.4). 焦土协议与热应力验证	100% · domestic
三	Powerline Communication & Physical Infiltration 电力线通信与物理渗透					200,000	
3.1	G3-PLC coupling & evaluation kit G3-PLC 耦合与评估套件	Custom coupling modules on Qualcomm QCA7000 / HiSilicon PLC silicon	50 sets	4,000	200,000	Develop mains-channel modulation algorithms; test cross-room / cross- floor covert agent mesh (§6.5). 跨房间智能体隐蔽通信	95% · SZ / CN

#	EQUIPMENT · 设备名称	CORE SPECS / REFERENCE MODEL · 核心技术规格	QTY	UNIT ¥	TOTAL ¥	JUSTIFICATION & USE CASE · 购置理由	LOCAL BOM
四	Compute Substrate & Network Capture 算力底座与网络抓包				2,160,000		
4.1	AI inference & MCP servers AI 推理与 MCP 服务器	Dual RTX A6000 48GB (or equivalent domestic accelerator) · liquid-cooled	4	250,000	1,000,000	Host the 42-server MCP mesh; run Qwen / DeepSeek open-weight models locally for traffic classification (§4.2 §4.5). 本地 MCP mesh 与开放权重推理	80% · CN compute / SZ
4.2	ClickHouse high-frequency storage nodes ClickHouse 高频存储节点	Dual Xeon · 4 × 3.2TB NVMe U.2 · 100GbE NIC	4	120,000	480,000	Sustain real-time ingest of the 1.2M+ hit corpus with millisecond MITRE classification queries (§7.1). 实时攻击语料写入与查询	70% · SZ assembly
4.3	100G network TAP + switch 100G 网络分流器与交换机	100 Gbps hardware TAP · 25G/100G datacenter switch	2 sets	100,000	200,000	Lossless out-of-band mirroring for eBPF/XDP takeover validation (§6.2). 无丢包旁路镜像	100% · CN (Ruijie-class)
4.4	R&D workstations & test terminals 研发工作站与测试终端	Mac Studio (M-series) + ARM64/Linux test-node array	15	32,000	480,000	Support 15+ engineers on browser-trap development, agent logic and endpoint compatibility testing. 浏览器陷阱与 Agent 开发	60% · mixed
五	Laboratory Infrastructure & Miscellany 实验室配套与杂项				900,000		
5.1	Lab wiring, UPS & custom fixtures 实验室弱电、UPS 及定制夹具	Industrial UPS · ESD workbenches · custom PCB bed-of-nails test fixtures	1 lot	900,000	900,000	Power safety under high-load testing; chip-level physical probing access. 高负载电力安全与芯片级探针接入	100% · SZ local
Total · 总计		12 line items · 116 units		4,800,000	4,800,000	Municipal ¥1.5M + self-raised ¥3.3M · composite local content ≈78%	

Quantity total is the sum of the quantity column above (116 units across 12 line items). Unit prices are

STRATEGIC VALUE FOR THE REVIEW COMMITTEE · 评审价值

By explicitly detailing the **Keysight Infiniium scopes** and the **ChipWhisperer-Husky** platforms, XCOM proves to the technical reviewers that it is not merely writing software — it understands the **deep physics of hardware security**: nanosecond voltage droop, EM side-channel leakage, secure-boot trust anchors and thermal limits of silicon.

By explicitly calling out the **QCA7000 PLC modules** and the **100G TAPs**, XCOM demonstrates that its powerline and XDP packet-takeover claims are backed by **real, orderable engineering parts** — not theoretical whitepaper concepts. Every R&D claim in §6 maps to a purchasable line item in this annex, and every line item maps back to a claim.

Keysight Infiniium UXR · 8 GHz / 25 GS/s

Needed to see a glitch: sub-nanosecond droop depth and width are the units in which secure-boot bypass is measured. No 8 GHz front end, no credible fault-injection claim.

ChipWhisperer-Husky + Pro

The instrumented ground truth behind the voltage fault-injection self-audit in §6.3 — the same class of tool used to publish hardware attacks, pointed at our own silicon.

Qualcomm QCA7000 G3-PLC modules

Turns §6.5 from a diagram into an orderable part: a mains-borne, RF-independent channel into RF-denied premises, testable at 50-node scale.

100 Gbps hardware TAP

Makes the §6.2 XDP takeover measurable: lossless mirroring is the precondition for proving a nanosecond-scale decision actually saw every packet.

D.1 Local and domestic procurement commitment · 本地化与国产化承诺

The RF anechoic chamber (1.3), thermal-shock chamber (2.3), PLC coupling modules (3.1), network switches (4.3) and laboratory wiring/fixtures (5.1) are procured **100% from Shenzhen-local suppliers or domestic vendors**. The compute (4.1) and storage (4.2) nodes are assembled and liquid-cooling-integrated **in Shenzhen**. The composite local/domestic procurement rate is projected at **≈78%**, exceeding the $\geq 70\%$ commitment in the annex and directly stimulating Shenzhen's local electronics and hardware prototyping supply chain.

本清单中射频微波暗室(1.3)、热冲击试验箱(2.3)、电力线耦合模组(3.1)、网络交换机(4.3)及实验室弱电配套(5.1)均 100% 采用深圳本地供应商或国产设备;算力节点(4.1)及存储节点(4.2)将在深圳本地进行服务器组装与液冷系统集成;综合本地化/国产化采购率预计达到 78%，远超附件中承诺的 $\geq 70\%$ 指标。

D.2 Procurement methods · 采购方式

CATEGORY	EXAMPLES	METHOD
General equipment	Workstations, UPS, network switches (each < ¥300,000)	询价采购 — inquiry / three-quotation
Specialised / imported	High-end oscilloscope (2.2), USRP N320 (1.1) — no domestic equivalent at 8 GHz bandwidth / required FPGA openness	进口产品审批 — import approval, then single-source or competitive negotiation
Custom-built	Anechoic chamber (1.3), test fixtures (5.1)	竞争性磋商 — competitive consultation with qualified Shenzhen integrators

D.3 Asset management · 资产管理

- All equipment purchased with municipal special funds is entered into the corporate fixed-asset ledger and labelled “深圳市科技创新委员会资助项目” (Shenzhen Science & Technology Innovation Committee funded project).
- During the project period no equipment may be disposed of, transferred or diverted without approval from the competent authority.
- After acceptance, assets are carried forward or managed under the *Special Funds Management Measures* (《专项资金管理办法》).

E Infinity Loop, Commercial Engine & Capital Structure

Appendix D lists what the laboratory costs. This appendix states what the laboratory *earns* — the commercial flywheel that pays for the instrumentation, and the capital structure that reconciles a Dutch BV, an EU SaaS line, a judicial-evidence line and a PRC joint venture inside one balance sheet. Every number below is either measured (estate snapshot 2026-10-28_121035Z) or explicitly marked as a *commitment* or a *target*.

E.1 The Infinity Loop BETA

The loop is the reason the corpus compounds and the reason the SaaS line is not a separate company. Client-side traps produce the telemetry; the telemetry produces the intelligence; the intelligence is sold back to the client as protection, to the legal market as evidence, and to the hardware programme as training signal.

1 **CISO dashboard**

Single pane for compliance posture, trap coverage and live hits. **BETA**

2 **Deploys micro-traps**

webcompliance SaaS agents push and manage trap configuration inside the tenant perimeter.

8 **Receives IOC feed**

Return edge — new techniques and indicators arrive before the tenant's own next scan.



3 **Attacker / botnet**

Probes and exploits the deployed traps. Nothing is blocked on a research lure.

· **XCOM micro-honeypots**

Open by design: the lure records rather than deflects (Doctrine §3.1).

4 **Telemetry & payloads**

Edge ingestion (Cloudflare Workers + eBPF) with one immutable event ID per hit.



5 **ClickHouse native corpus**
1,219,833

cumulative hits (~3.5k / 24 h) — not “1.2 M per month”.

6 **MITRE classification**

Technique mapping at ingestion, synthetic and self-loop traffic labelled, never filtered.

7 **42-server MCP mesh + Agent-AIX**

Autonomous analysis on local Qwen / DeepSeek lanes; every tool call journaled. **BETA**

RETURN EDGES — WHAT CLOSES THE LOOP

8 **→ Threat intel & IOCs to the dashboard**

Closes the SaaS loop: the tenant's own hits become everyone's pre-emptive blocklist.

9 **→ Evidence Packs to rechtspraak-ai.nl / insurers**

Hash-addressed, redaction-manifested exports aimed at Dutch/EU evidentiary standards. **R&D**

10 **→ Training signal for the Shenzhen hardware lab**

Real attack shapes steer SDR, G3-PLC and fault-injection priorities (Appendix D). **R&D**

Evidence-bound correction. The loop diagram circulated internally in October 2026 labelled the corpus edge “1.2 M+ hits/*mo*”. That is not what the corpus holds: 1.22 M is the **cumulative** row count (measured 24 h volume: 3,515). The corrected label is used above and in the Mermaid source. A flywheel diagram may be a business sketch, but it may not misstate a measured quantity.

E.2 Entity and brand map

BRAND	LEGAL / CORPORATE ROLE	FUNCTION IN THE LOOP	STATUS
xcom.dev	XCOM.DEV — KVK 99492334 · BTW NL005391698B11 · Sint Olofssteeg 4C, 1012 AK Amsterdam	Parent, estate operator, corpus owner, grant applicant	LIVE
agent-aix.com	Product surface of XCOM.DEV	Agent orchestration, red-team engagements, autonomous analysis	BETA
webcompliance.app	B2B SaaS line of XCOM.DEV	GTM engine — deploys tenant micro-traps that feed the corpus	BETA
rechtspraak-ai.nl	Evidence / legal-tech line of XCOM.DEV	Turns telemetry into court-ready, tamper- evident evidence packs	R&D
cybersecurityad.com	CyberSecurityAD BV — holding / services	Contracting vehicle for managed services and pilots	LIVE
XCOM-Shenzhen Joint Lab	PRC WFOE (in formation) · 深圳信息物理安全联合实验室	Hardware R&D: SDR, G3-PLC, fault injection, thermal	CONCEPT

E.3 Revenue architecture — dual engine

One engine sells protection to many tenants at low marginal cost; the other sells depth to few buyers at high margin. The first funds the corpus, the second monetises it.

ENGINE	LINE	BUYER	MECHANISM	ANCHOR
GTM volume	webcompliance.app	CISO / IT-admin at SMEs and mid-market	Subscription: trap deployment + compliance dashboard + evidence packs	SaaS · per tenant
	XCOM Grid	Managed-security buyers	Fully managed deception estate (sensors, domains, traps, dashboards)	€490/sensor/mo
DEEP-TECH depth	Agent-AIX Red Team	Enterprises with a written authorization	Scope-bounded autonomous pentest, journaled tool calls	€18k/engagement
	Intel Feeds	Threat-intel consumers	Anonymized ATT&CK-mapped taxonomy from the live corpus	€2k/mo
	Evidence Packs	Insurers, counsel, regulators	Replayable, redacted, hash-addressed export bundles (rechtspraak-ai.nl)	€6k/pack
	Joint-lab programmes	PRC municipal / partner programmes	Funded hardware R&D against committed KPIs (Appendix D)	grant + matching

E.4 Seed Preferred financing — term sheet summary

ISSUER	XCOM.DEV, Amsterdam, NL · KVK 99492334
SECURITY	€1,000,000 Seed Preferred
PRE-MONEY	€4,000,000 → €5,000,000 post-money
FOUNDER	Peter W. “Jedi” Oldenburger, CEO
GOVERNING TERMS	Standard NVCA terms as per prior iteration

Use of proceeds.

%	ALLOCATION	WHAT IT BUYS
40%	GTM & Infinity-Loop scaling	International rollout of <code>webcompliance.app</code> ; distributed micro-honeypot deployment feeding the central ClickHouse corpus.
30%	Core R&D	42-server MCP mesh expansion, Agent-AIX orchestration, and the <code>rechtspraak-ai.nl</code> judicial-evidence pipeline.
30%	Shenzhen WFOE & hardware R&D	Capitalising the PRC subsidiary and funding local hardware prototyping (SDR, PLC, fault injection) to satisfy municipal grant matching requirements.

Matching-funds & GTM covenant (§12 of the term sheet).

- The investor acknowledges the pending Shenzhen municipal AI grant (**¥6,000,000** requested).
- XCOM.DEV commits that **€1,000,000** of self-raised funds plus existing corporate reserves — together **≥ ¥12,000,000 equivalent** — are earmarked for the Shenzhen programme.
- Any FX shortfall against the ¥12 M equivalent is made whole by the founder.
- **Target (not a measurement):** **≥ 500** paying `webcompliance.app` enterprise tenants by **month 18**, so the loop's data moat is continuously validated by paying traffic rather than by internal synthetic load.

E.5 Capital reconciliation — three envelopes, one balance sheet

The most common failure mode in a dual-jurisdiction raise is a number that appears in the term sheet, the grant annex and the equipment list in three different sizes. The three envelopes below are **nested**, not competing, and the arithmetic closes.

ENVELOPE	SIZE	SOURCE	RELATIONSHIP
Seed round	€1,000,000	Equity (€4 M pre / €5 M post)	30% ($\approx$ €300k) earmarked for the Shenzhen WFOE and hardware R&D
Shenzhen programme – municipal	¥6,000,000	Municipal grant (requested)	Contains the ¥1,500,000 <i>equipment</i> subject in Appendix D
Shenzhen programme – self-raised	¥12,000,000	Founder/company matching	Funded by the seed plus reserves; contains the ¥3,300,000 <i>equipment</i> self-raised line
Equipment subject (设备费)	¥4,800,000	¥1.5 M municipal + ¥3.3 M self-raised	A subset of the ¥18,000,000 programme envelope – 26.7% of it

$¥6,000,000 + ¥12,000,000 = ¥18,000,000$ programme envelope · equipment $¥4,800,000 = ¥1,500,000 + ¥3,300,000 \leq ¥18,000,000$ · no figure in this paper requires the same euro or yuan to be spent twice.

E.6 webcompliance.app — positioning copy (EN / NL)

Published here so the sales surface and the technical paper cannot drift apart. Every claim in this copy maps to a § reference: “tamper-proof” → §8.1 provenance chain; “court-ready” → §8, E.2; “global corpus” → §7.1.

EN

Headline: Stop guessing what attackers want. Watch them hit your traps.

Subhead: Webcompliance.app doesn't just check your security boxes. It deploys autonomous deception traps across your network, capturing real-time attack telemetry and feeding your compliance dashboard with court-ready evidence.

Feature 1 — The Infinity Loop data moat. The more you deploy, the smarter our AI gets. Your local traps feed into the global XCOM ClickHouse corpus, instantly immunizing your network against zero-days seen across our global estate.

Feature 2 — Evidence-bound compliance. Powered by the same immutable snapshot technology used by *rechtspraak-ai.nl*, every alert generates a cryptographically signed, tamper-proof evidence pack for auditors and cyber-insurers.

CTA: Deploy your first micro-honeypot in 60 seconds. [Start Free Trial]

NL

Kop: Stop met gissen wat aanvallers willen. Kijk hoe ze in uw vallen trappen.

Subkop: Webcompliance.app vinkt niet alleen uw security-boxjes af. Het zet autonome deceptie-vallen uit in uw netwerk, legt real-time aanvalstelemetrie vast en vult uw compliance-dashboard met bewijs dat klaar is voor de rechtbank.

Feature 1 — De Infinity Loop. Hoe meer u inzet, hoe slimmer onze AI wordt. Uw lokale vallen voeden het wereldwijde XCOM ClickHouse-corpus, waardoor uw netwerk direct wordt beschermd tegen zero-days die elders in ons ecosysteem zijn gezien.

Feature 2 — Bewijskrachtige compliance. Aangedreven door dezelfde onveranderlijke snapshot-technologie als *rechtspraak-ai.nl*. Elke melding genereert een cryptografisch ondertekend, niet-vervalsbaar bewijspakket voor auditors en cyberverzekeraars.

CTA: Zet uw eerste micro-honeypot uit in 60 seconden. [Start Gratis Proefperiode]

Claim check before publishing. “Cryptographically signed” and “tamper-proof” are only defensible once the signed-challenge workflow in §5.2 and the provenance chain in §8.1 are live for tenant traffic. Until then the SaaS copy must carry the same **BETA** qualifier as the rest of the loop — the discipline that made the peripheral audit publishable applies to marketing copy too.

— END OF MASTER DOSSIER V2.0 • PART M MEMORANDUM + TECHNICAL PAPER V1.0 —

