MANAGEMENT REVIEW · RISK COMMITTEE

Quarterly threat review

A quarter-business-review of brand-abuse exposure: where the threat moved over three months, how you stand against the sector, the performance of the monitoring and takedown programme, and the regulatory evidence produced under DORA, NIS2 and eIDAS.

Prepared for **Sample Bank · BE-RETAIL-01 — risk committee & security leadership** · indicators defanged · EU-resident corpus

01 Executive summary

Across Q2, PhishNet tracked **3,470 malicious assets** impersonating the brand — up 18% on Q1 and the highest quarter on record. The mix shifted decisively toward **dedicated attacker infrastructure** and **mobile app clones**, both of which carry higher credential-capture risk than hijacked sites. The brand stayed the most-impersonated Belgian bank in the sample set by a wide margin. The takedown programme resolved 142 assets at a median 3.1 days, and the same evidence base produced **5 DORA evidence packs** and context for 2 NIS2 incident assessments. Net posture: exposure is rising structurally; the programme is keeping pace on removal but registration of new look-alikes continues to outrun takedowns.

Methodology. Passive aggregation of public threat-intelligence feeds (openphish, phishdestroy, threatfox, phishdb, krogza and others). No intrusive testing is performed and no personal data is processed; collection is GDPR-aligned and EU-resident end to end. Each asset is normalised, scored 0–100 for risk, SHA-256 hashed and timestamped on capture. Indicators in this document are defanged. Risk scores are PhishNet assessments derived from public feeds and are independently verifiable.

02 Quarter in numbers

Q2 2026

3,470

malicious assets tracked in Q2

+18% vs Q1 · ~68 Belgian brands in the corpus

74	142
HIGH-SEVERITY ASSETS	TAKEDOWNS RESOLVED
3.1d	**5**
MEDIAN RESOLUTION	DORA EVIDENCE PACKS

03 Trend across the quarter

NEW ASSETS / WEEK

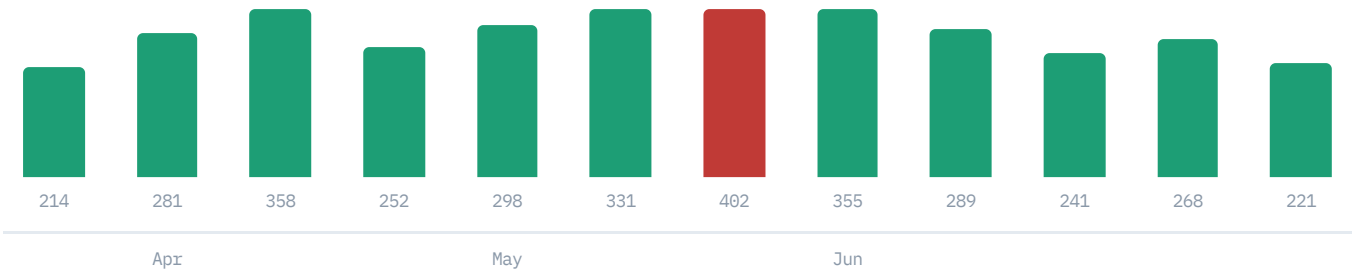

Weekly new-asset volume, Apr–Jun. The Week 7 peak (402) marks a coordinated app-clone campaign; the quarter never returned to the April baseline.

04 Sector benchmark — Q2 totals

SAME METHOD

Bank A (you)	3,470
Bank B	612
Bank C	288
Bank D	121
Bank E	24

Assets attributed per brand across the quarter. The brand absorbs the large majority of sector impersonation — being market leader means being primary target.

05 How the threat shifted this quarter

Dedicated infrastructure

+61%

Attacker-owned hosts over hijacked sites — purpose-built and harder to dismiss as noise.

App clones

+47%

Mobile-banking clone kits now dominate the high-severity set.

Smishing domains

+39%

SMS lures on low-cost TLDs with fast rotation, evading static blocklists.

06 Campaigns & infrastructure

- **April — credential harvest.** 4 hosts on 198.51.100.0/24 captured online-banking credentials; risk up to 89.
- **May — app-clone wave.** A single kit fingerprint produced 30+ mobile clones across shared hosting.
- **June — smishing surge.** Rapid-rotation SMS lure domains targeting retail customers.

Top hosting block

`198.51.100.0/24` — 19% of dedicated-infra assets this quarter.

Operator linkage

One kit fingerprint links 41% of high-severity clones to a single operator.

07 Monitoring & takedown programme

Assets monitored	Takedowns submitted	Resolved	Median resolution
3,470	198	142	3.1d

Resolution rate 72%. The gap is concentrated in two registrars with slow abuse handling; a standing escalation path is recommended (§09).

08 Regulatory evidence produced

DORA · NIS2 · EIDAS

DORA

ICT third-party & incident evidence captured with timestamped provenance for the resilience file.

NIS2

Threat context supporting reportable-incident assessment and supervisory review.

eIDAS

Qualified RFC 3161 timestamps on every cited finding — audit-grade to court-grade.

5 DORA ICT/incident evidence packs and context for 2 NIS2 assessments were produced this quarter, each timestamped with a qualified RFC 3161 token — the same intelligence serving both security and compliance.

09 Risk posture & outlook

01 **Exposure is structurally higher, not cyclically.** Fund continuous monitoring at sector-leader scale; snapshot-only coverage leaves months uncovered.

02 **Formalise registrar escalation.** Two registrars constrain the resolution rate — a standing relationship lifts throughput on the highest-volume sources.

03 **Operationalise kit fingerprints.** One operator drives most high-severity volume; standing detections catch new clones on day one.

04 **Carry the evidence into DORA/NIS2.** Brand-abuse findings already double as resilience evidence — one programme, two regulatory outcomes.

A Appendix · feed coverage, method & glossary

FEED	CONTRIBUTES	COVERAGE	CADENCE
phishdb	phishing URLs & kits	high	continuous
krogza	look-alike / app-clone domains	high	continuous
openphish	verified phishing feed	medium	continuous
phishdestroy	takedown & abuse corpus	medium	continuous
threatfox	IOC exchange (abuse.ch)	medium	continuous
closed channels	leaked-credential mentions	contextual	monitored

Methodology. Passive aggregation of public threat-intelligence feeds (openphish, phishdestroy, threatfox, phishdb, krogza and others). No intrusive testing is performed and no personal data is processed; collection is GDPR-aligned and EU-resident end to end. Each asset is normalised, scored 0–100 for risk, SHA-256 hashed and timestamped on capture. Indicators in this document are defanged. Risk scores are PhishNet assessments derived from public feeds and are independently verifiable.

Glossary

Live asset	A malicious URL, domain or IP confirmed reachable at observation time.
High-severity	Risk score ≥ 85 — typically dedicated infrastructure or working credential capture.
Dedicated infra	Attacker-controlled host stood up to target the brand, not a hijacked legitimate site.
Kit fingerprint	Shared code/asset signature linking multiple clones to one operator.
Defanged	Indicator rendered non-clickable (e.g. [.]) for safe handling.
MTTD	Median time to detect a new asset from first public appearance.
TLPT	Threat-led penetration testing — DORA-defined exercise the evidence base can support.
RFC 3161	Standard for a trusted timestamp token; eIDAS-qualified when issued by a qualified TSA.

A quarter your risk committee can act on.

Three months of brand-abuse intelligence that doubles as DORA and NIS2 evidence — EU-resident, every figure traceable to a timestamped source.

DIRECT CONTACT

stijn@galacticautomation.com

galacticautomation.com · phishnet.be

Stijn Van Hijfte

Galactic Automation BV · Zottegem, Belgium — Founder, PhishNet (EU-sovereign CTI)

Method. Passive aggregation of public threat-intelligence feeds. No intrusive testing, no personal data processed, GDPR-aligned. Indicators defanged; live as of report date.

CISSP · CRISC · IAAP Data Protection Fellow

SAMPLE report. All institutions, indicators, IPs (RFC 5737 TEST-NET) and domains (.example / .invalid) are illustrative and do not represent a real client or live threat. Risk scores are PhishNet assessments derived from public feeds.