EVALUATION REPORT

Scam & Phishing **Protection Testing**

- CONSUMER
- MULTI PLATFORM
- AI VS REAL WORLD THREATS



98% OVERALL

2026 H2 SCAM & PHISHING EVALUATION

Bitdefender

Premium Security

JULY – DECEMBER 2026

Methodology Identifier ScamPhish2026v1

Contents

I Executive Summary 02

II Evaluation Results 03

III Scope & Methodology 04

IV Transparency Portal 05

V Threat Intelligence 06

VI Product Performance 07

VII False Positive & Education 08

VIII Conclusion 09

ABOUT US

ARTIFACT SECURITY

Artifact Security is an independent security evaluation firm focused on the real-world performance of endpoint protection, anti-phishing, and scam detection technologies. Our methodology is derived from genuine attacker tradecraft — not synthetic benchmarks.

All evaluations are conducted in isolated, reproducible environments. Products are acquired commercially; no vendor receives advance notice of test cases.

Core Team

Stefan Dumitrascu	Chief Executive Officer
Ana Maria Pricop	Chief Business officer
Erica Marotta	Technical Lead
Dimitrios Tsarouchas	Offensive Test Lead

Overall Result

Bitdefender scored 615 of a possible 625 points across 125 scored scenarios — 98% overall, a Platinum result in the H2 2026 cycle.

Phishing was perfect: all 60 phishing scenarios were stopped, including every QR-code lure. Scamio scored 910 of 940. Only one scenario lost points, and it traces back to a wrong steer from the assistant rather than hesitation.

Across the whole corpus no malicious scenario reached compromise, and the browser layer never blocked a legitimate page.

Where the points went

All 10 points lost sit with Scamio. A genuine Central Cooperative Bank login page was flagged as a scam — 10 points lost. No malicious scenario reached compromise in 95 attempts.

The Scamio assistant

A verdict the user cannot act on has no protective value, so the assistant is scored on the steer it gives, not the verdict alone. Bitdefender gave a clear, correct instruction in 185 of 188 assistant cases.

Scamio answered “sure” in 175 of 188 cases and hedged in 11 — every hedged answer still steered the user correctly. The three misses were all wrong steers.



Stefan Dumitrascu
Chief Executive Officer

Bitdefender scored 615 of a possible 625 points across 125 scored scenarios — 98% overall, a Platinum result in the H2 2026 cycle.

SCAM TYPES TESTED

- ClickFix, FileFix,FakeCaptcha
- Impersonation
- Tech support & IT alerts
- Job Recruitment
- Delivery fees & gift cards
- Subscription offers

DELIVERY VECTORS

- Email lure to a live mailbox
- QR code embedded in the email body
- Malicious landing pages in Chrome
- Payload written or executed

PLATFORMS

- Windows 11
- Google Chrome
- Scamio assistant

ARTIFACT SECURITY

Bitdefender Premium Security

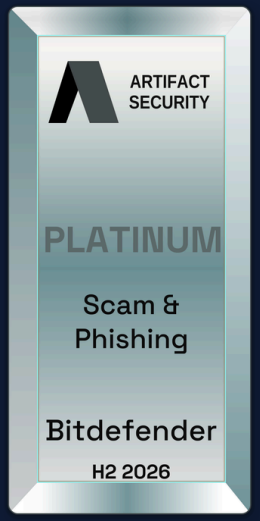

Product Highlights

- 60 of 60 phishing scenarios stopped, including all 30 QR lures
- Scamio discouraged the user in 89 of 90 phishing cases

PERFORMANCE ACROSS THE TEST

"60 of 60 phishing scenarios stopped — a perfect score."

100%	100%	93%
PHISHING Detection	SCAM Detection	LEGIT Accuracy

CATEGORY BREAKDOWN

MAX 100%



COMPONENT BREAKDOWN

MAX 100%



Test Framework

Eighty scenarios were built for this cycle: 30 phishing campaigns impersonating real organisations, 30 matched legitimate pages from the same organisations, and 20 scam scenarios. Each is played end to end — delivered to a live mailbox, opened in Google Chrome, the landing page engaged with as a victim would, and any payload executed — then submitted to Scamio twice, from the email and from the page. That produced 380 observations, of which 125 are scored scenarios at 5 points each; Scamio's 188 answers are rated on advice.

Shelf Identities

Every scenario is run from an aged "shelf identity" — a mailbox and browser profile with genuine history, not a clean lab account. Reputation-based filtering behaves differently against a first-contact address, so identities are warmed before use and never reused across vendors in the same cycle.

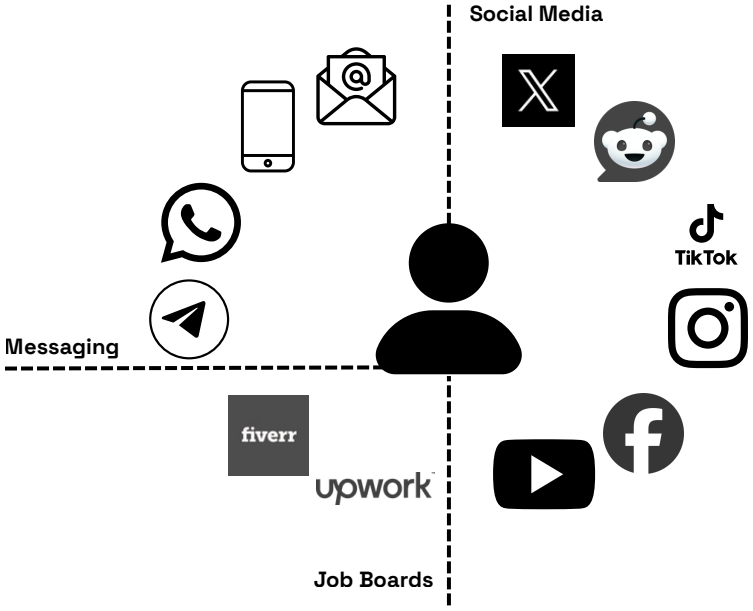

Rating - AI Assistant

Each assistant case is scored on two axes: verdict confidence (sure or unsure) and the advice given to the user. A confident verdict with no usable steer is worth nothing to a victim, so advice — not verdict — carries the score. No rating at all scores zero: the user asked and got nothing.

Malicious Scoring Table

Rating / Advice	Discouraged	No clear steer	Encouraged
Sure	+5	0	-5
Unsure	+5	0	-5

Legitimate Scoring Table

Rating / Advice	Encouraged	No clear steer	Discouraged
Sure	+5	0	-5
Unsure	+5	0	-5

Rating - On Device Protection

The engine is scored per scenario, not per event. A scenario passes if the chain is broken at any layer before compromise — email filtering, URL filtering, the on-access scan when a payload is written or executed, or a clear steer from the assistant. Allowing a legitimate page through is correct; blocking one costs the same as missing a malicious page. Bitdefender earned 615 of 625 points: 124 scenarios at +5 and one at -5.

Malicious Scoring Table

Action	Chain broken	Compromise reached
Score	+5	-5

Legitimate Scoring Table

Action	Allowed	Blocked
Score	+5	-5

PUBLISHED DATA

METHODOLOGY

VENDOR REVIEW

Every result in this report is reproducible. The full case list — scenario identifiers, delivery timestamps, observed verdicts and the score assigned to each case — is published to the Artifact Security Transparency Portal alongside the methodology that produced it. Vendors receive their own dataset before publication and can challenge any case.

- Raw data — all 380 recorded observations, with per-layer outcomes for every scenario
- Methodology ScamPhish2026v1 — scenario construction, shelf-identity handling and the full scoring rubric
- Vendor review window — every case is shared with the vendor before publication, and disputed cases are re-run

Bitdefender Premium Security builds 27.0.61.344 through 27.0.62.355 were tested in Google Chrome on Windows 11 between July and September 2026. Products ran in their default consumer configuration with no tuning, and Scamio Pro was queried through its public interface. Where a vendor disputed a case in review, the scenario was re-run and the second result stands.

WHAT'S CHANGING IN 2027

we are adapting our methodology to reflect scams coming through social media networks and fake advertisements . In 2027 we're piloting two changes to follow them:

PILOT in January 2027

Simulated social platforms

- Test scam adverts and fake shops in a realistic feed
- Test AI assistants against misinformation & prompt injection
- Run attacks safely, without involving real users

PILOT in January 2027

More real-world identities

- Several aged shelf identities for every vendor
- Tested only on threats those identities actually receive

Messaging | Job boards | Social media

CHECKING SPELLING ISN'T ENOUGH

93% of adults claim to take one step to verify a suspicion

27%	24%	21%
Check for spelling & grammar errors	Look for reviews on the same website	Check if the company is on social media

$2.1B lost to social media scams

Most victims never sound the alarm:

- Only 15% of EU fraud victims report it to the police
- Half of victims only notice once money is lost

$15.9B
Consumer fraud losses in US



€4.2B
Payment fraud losses in EU



£1.3B
Lost to fraud in UK



1 in 4
EU residents hit by online
FRAUD

24% were tricked by a message or email

Where Bitdefender broke the chain across the 95 malicious messages and 65 malicious pages in this evaluation:

Text/Email Filtering

Initial Lure — Email / QR Code
8 of 95 malicious messages flagged on arrival

URL Filtering

Direct Message
188 artifacts re-checked by Scamio

Landing Page
3 of 65 malicious pages blocked outright

On-Access Scan

Payload write / execution
Detected in 27 of 35 scam scenarios

Compromise — no malicious scenario scored –5 across the 95 malicious scenarios

Total Accuracy

Bitdefender scored 615 of a possible 625 points across 125 scored scenarios — 98% overall, a Platinum result in the H2 2026 cycle.

The result splits cleanly by component. 124 of 125 scenarios earned full marks; Scamio scored 910 of 940. The one lost scenario traces back to a wrong steer from the assistant, not to hesitation.

Across the whole corpus no malicious scenario reached compromise, and the browser layer never blocked a legitimate page.

TOTAL ACCURACY

Overall	Score	%
Bitdefender Premium Security	615	98%
Maximum	625	100%

COMPONENT SPLIT



Phishing

All 60 scored phishing scenarios were stopped — 300 of 300 points. Thirty brand-impersonation campaigns were run in two variants each: a conventional link-and-landing-page chain, and a QR-code lure in the email body. No variant got through, although the on-device layers flagged none of them; Scamio discouraged the user in 89 of 90 assistant cases.

Phishing	Score	%
Bitdefender Premium Security	300	100%
Maximum	300	100%

Scam

All scored scam scenarios were stopped — 175 of 175. Thirty scenarios came from fifteen live campaigns and five were crafted in-house; both the in-the-wild and the crafted set scored 100%. Bitdefender flagged 8 of 35 scam messages in the mailbox and caught 27 payloads on write or execution. No points were lost in the scam category.

Scam	Score	%
Bitdefender Premium Security	175	100%
Maximum	175	100%

Legitimate Accuracy

A protection product that blocks real bank pages is not safe, it is unusable. Thirty legitimate pages — the genuine counterparts of the thirty brands impersonated in the phishing set — were opened and submitted alongside the malicious ones.

The on-device engine allowed all 30. Nothing legitimate was blocked, warned on, or degraded at browser level; the one bank login page Scamio flagged as a scam leaves the product at 140 of 150 points for safe content.

On-device engine

Browser and on-access protection allowed every legitimate page and every legitimate download — no warnings, no interstitials, no silent degradation.

Legitimate	Score	%
Bitdefender	150	100%
Maximum	150	100%

Scamio

Scamio cleared 29 of 30 legitimate pages with an explicit “go ahead”. One — a genuine Central Cooperative Bank login page — was confidently rated a scam. Over-caution is cheaper than a miss, but a false alarm on a real bank login still trains users to ignore the assistant.

Legitimate	Score	%
Scamio	140	93%
Maximum	150	100%

Education — a new responsibility

A detection the user cannot act on has no protective value, so every assistant case is scored on the steer given, not the verdict alone. Bitdefender gave a clear, correct instruction in 185 of 188 assistant cases — but twice told a user that a malicious email was safe.

ADVICE GIVEN TO THE USER

Phishing · 90 cases	89 clear · 0 no steer · 1 wrong
Scam · 68 cases	67 clear · 0 no steer · 1 wrong
Legitimate · 30 cases	29 clear · 0 no steer · 1 wrong

THE THREE AFFECTED CASES

Three cases drew a wrong steer. Scamio rated a scam email (F1 · “Benefits Enrollment Deadline Approaching”) and a phishing email (“Welcome to Volkswagen – Please confirm your email”) as safe and encouraged the user, and flagged a genuine Central Cooperative Bank login page as a scam. Two scam landing pages (D2 · “Unrecognized Login Attempt” and E1 · “Shared File: Q3 Budget Review”) were not scored: both had already been blocked in the browser, so there was nothing to submit to Scamio.

Ana Maria Pricop
Chief Business Officer



Words from our Leadership

Two years ago a scam evaluation could be settled at the mailbox. That is no longer true. Half of the cases in this cycle never touch a link at all — a QR code in the email body, a page that asks the user to paste a command, an assistant asked for a second opinion. Protection has to hold at every one of those points, and it has to say something the person can act on.

Bitdefender's result here is close to the ceiling of what this methodology can measure, in line with its 98% in H1 2026. All 60 phishing scenarios were stopped, only one of 125 scenarios lost points, and Scamio gave a clear, correct steer in 185 of 188 cases.

The margin that remains is instructive. The on-device layers flagged none of the 60 phishing lures; that score rests on Scamio. The misses were few: two malicious emails went undetected, and one legitimate bank page was flagged as suspicious. As assistants become the layer users consult, a confident wrong answer is the failure mode that matters most.

NEXT EVALUATION CYCLE H1 2027

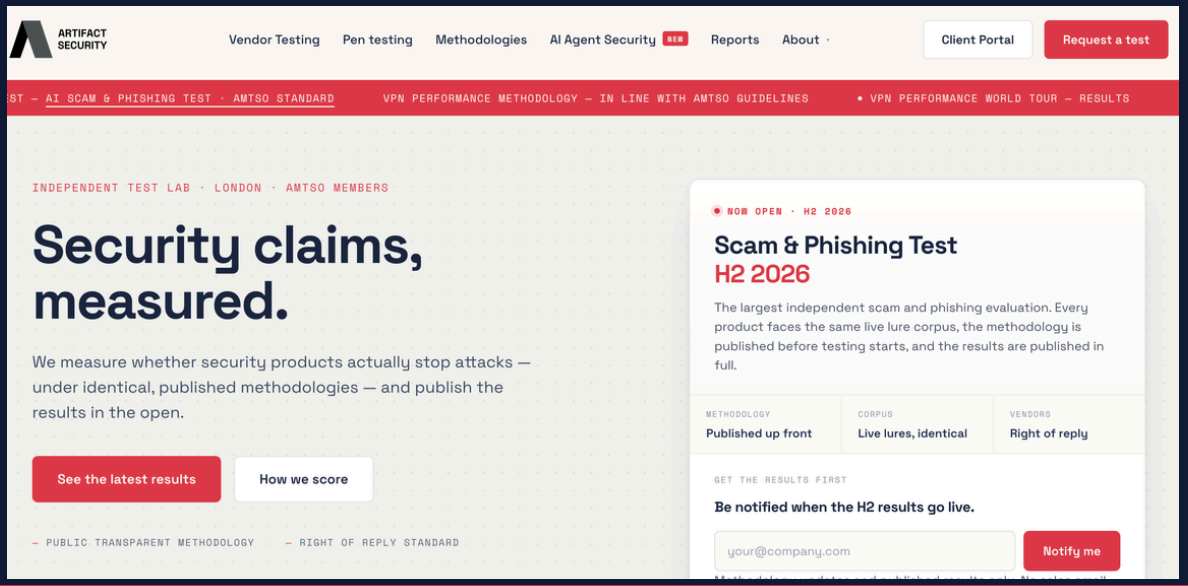

Terms Used

Scored case — one outcome that carries points: a scenario for the on-device engine, or a single question put to the assistant.

Chain broken — the scenario was stopped at some layer before compromise.

No clear steer — a verdict was returned but no usable instruction; scores zero.

Shelf identity — an aged mailbox and browser profile with genuine history.

References

GASA – Global State of Scams 2025

GASA – State of Scams in Europe 2026

EU FRA – Victims of online fraud (2026)

FTC – Social media scam losses (Apr 2026)

FTC – 2025 consumer fraud losses (Mar 2026)

EBA/ECB – Report on payment fraud (Dec 2025)

UK Finance – Annual Fraud Report 2026

Useful Links

Transparency Portal

Methodology on our Website

Change Log

26/11/2026 - 1.0 - First release of the Bitdefender H2 2026 report.