

THREAT RESEARCH REPORT

Infrastructure Fingerprinting of a Multi-Region Phishing Operation

How a modular phishing platform hides active campaigns behind seemingly inactive domains, regional paths, and token-gated access

INSIDE THIS REPORT

Table of Contents

01 Region-Wide Phishing Platform Identified 3

Executive Summary 3

02 Phishing Infrastructure Initial Discovery 4

Server-Side Fingerprinting 5

03 Infrastructure Quality Assessment 6

Infrastructure Snapshot 6

04 Regional Targeted Brands 8

05 Application-Level Fingerprinting and Operational Infrastructure 9

Case Study: Rewards-Themed Campaign Delivery Through Regional Path Activation 9

Case Study: Conclusions 14

06 Infrastructure Ownership Breakdown 15

07 Comparing the Modus Operandi to Known PhaaS Groups 16

08 Conclusions 17

09 Recommendations and Attribution Limitation 17

Detection and Hunting 17

Attribution Limitation 18

Contact Us 19

At a glance. Starting from a handful of phishing URLs, the Check Point Exposure Management Research Team built a server-side fingerprint that narrowed a broad Internet footprint to **322 IP addresses** and expanded into **thousands of rotating domains**. The domains look inactive. The phishing content hides behind regional paths, locale parameters, and token-gated access screens.

01 / OVERVIEW

Region-Wide Phishing Platform Identified

Executive Summary

The Check Point Exposure Management Research Team investigated a large-scale phishing infrastructure designed to conceal active phishing content behind seemingly inactive domains.

Instead of hosting the phishing page directly from the root domain, the infrastructure relied on regional activation mechanisms, localized content delivery, token-gated access controls, and dynamically generated backend configuration. These techniques significantly reduce visibility through traditional domain monitoring approaches and increase the difficulty of identifying active campaigns through conventional inspection techniques alone.

~19K

Initial Shodan results for the GoFrame response banner

322

Unique IP addresses after adding the Caddy fingerprint

10,187

Associated domains in the September 11, 2026 window

90%

Of those domains in the highest malicious category

Starting from a small number of phishing URLs, the investigation uncovered a repeatable infrastructure fingerprint based on a combination of GoFrame HTTP Server, Caddy, concealed regional paths, and recurring API structures. Applying this fingerprint reduced a broad Internet-facing footprint to 322 unique IP addresses, which were subsequently expanded into thousands of related domains spanning multiple countries, languages, industries, and impersonated organizations. The observed domain population demonstrated substantial churn between collection windows, indicating a continuously rotating infrastructure rather than a static campaign inventory.

Application-level analysis revealed that seemingly unrelated phishing campaigns shared common architectural components. Multiple deployments used region-specific activation logic, tokenized backend APIs, session-controlled workflows, dynamically generated configuration objects, and reusable user interface templates. Despite differences in branding, language, and victim targeting, the underlying application behavior remained consistent. These similarities strongly suggest a shared platform architecture capable of supporting localized phishing campaigns across multiple regions and industries.

The investigation further demonstrated that the platform supports multiple localization mechanisms. Some campaigns hide phishing content behind country-specific paths, while others dynamically select content through country and locale parameters. Additional deployments introduce token-gated validation layers and multilingual access-control screens before exposing the phishing page. These variations indicate a modular design in which routing, localization, and delivery components can be adapted independently while preserving a common technical foundation.

02 / DISCOVERY

Phishing Infrastructure Initial Discovery

The investigation began with a small set of URLs associated with an active phishing campaign. The initial hosts exposed several interesting application paths:

```
<generic-domain-name>/[Region Geolocation] e.g. /uk /gr /de /<randomized-path>/api  
/<randomized-path>/api/input /<template>_[Region]/header.html  
/<template>_[Region]/footer.html
```

This behavior was observed internally across various Check Point customers, with phishing sites resembling the brand while the phishing lure itself was hidden in the website directory. At first glance, the obvious pivot was the domains themselves. Several related domains followed a similar naming pattern and shared the same authoritative nameserver pairs. That seemed promising, until the full scope revealed that these authoritative nameservers belong to a CDN provider and serve approximately 5 million domains, making the infrastructure overlap too broad to be a reliable indicator on its own.

However, a closer examination of domains using the CDN provider's infrastructure revealed a recurring pattern: similar naming conventions and typosquatting domains impersonating various brands and government entities. These commonalities warranted further investigation.

We focused on a known phishing site targeting an Italian government entity that we had previously observed. Its domain used Italian gov-themed domains and transliterations combining iterations of 'it' and 'gov', incorporating subdomain terms associated with payments, reimbursements, and similar services. Other malicious domains using similar tech stacks followed the same pattern, using local-language terms written in the Latin alphabet to target entities in different countries. We confirmed that the page hosted on the Italian-themed domain was served from the **/it** directory. This correlation persisted across the network: for example, Germany-themed domains served their phishing pages from **/de**.

At this stage of the investigation, the recurring directory structure and regional naming conventions suggested that the sites might share the same underlying phishing kit, although they were not necessarily operated by the same threat actor. Subsequent findings provided further evidence supporting this assessment.





Figure 1 A blank 'Not Found' page is returned when the concealed directory is not accessed

● Server-Side Fingerprinting

Typically, when the phishing directory is unknown, a sophisticated phishing kit provides researchers and threat hunters with no obvious indication that active content is hidden elsewhere on the domain. The root path may return a generic "Not Found," "404," or other error responses, even while a fully functional phishing page remains accessible through a concealed directory.

The most common way to identify the hidden path is to capture the complete phishing URL at the distribution layer. However, visibility is limited when URLs are distributed through SMS, private email messages, WhatsApp, or similar closed communication channels. In these cases, discovery generally depends on a recipient reporting the full URL or on related traffic being observed. Without either source, the concealed phishing operation may remain undetected through conventional domain-based monitoring and requires infrastructure-related analysis.

As no other notable signals were visible at this stage, we pivoted to the underlying technology stack. Shodan data for port 80 revealed an uncommon combination of HTTP response headers: the server identified itself as a GoFrame HTTP Server operating behind a Caddy intermediary. The response also included GoFrame's Trace-Id header, typically containing a 32-character hexadecimal value. Because this behavior is native to GoFrame, the header is not independently distinctive; its hunting value emerges when combined with Caddy, stable response behavior, session cookies, concealed application paths, and repeated phishing API structure.



```
// 80 / TCP {} ↗ 2037417530 ⓘ  
  
HTTP/1.1 302 Found  
Content-Type: text/plain; charset=utf-8  
Date: Thu, 03 Sep 2026 18:01:29 GMT  
Location: https://43.165.1.140/  
Server: GoFrame HTTP Server  
Trace-Id: 781c135e85e1d11880e4205ed2351c08  
Via: 1.1 Caddy  
Transfer-Encoding: chunked
```

Figure 2 HTTP response banner exposing parts of the server stack

03 / ASSESSMENT

Infrastructure Quality Assessment

GoFrame is a legitimate general-purpose web framework, yet it appears to have a comparatively limited Internet-facing footprint. At the time of our investigation, an initial Shodan query for the relevant GoFrame response banner returned approximately 19,000 results. For an HTTP server fingerprint, this is a relatively small population and makes GoFrame useful as an initial hunting signal.

The signal became considerably more interesting when combined with the rest of the observed server stack. Requiring the GoFrame HTTP response characteristics together with the **Caddy intermediary and associated response-header pattern reduced the population to approximately 322 distinct IP addresses.**

Neither GoFrame nor Caddy is malicious or actor-specific. However, this reduction demonstrates why the combination is valuable for infrastructure discovery. Rather than searching millions of generic web servers, the server-side fingerprint produces a comparatively small candidate population that can then be examined and compared with Check Point's own data.

● Infrastructure Snapshot

To evaluate how much of the candidate infrastructure remained visible, we compared the service response against our snapshot taken on September 24, 2026 at 16:00 UTC. The query applied the same GoFrame HTTP Server response characteristic across the candidate IP set.

Of the 322 candidate IP addresses evaluated, 229 were associated with at least one domain in the current snapshot, representing 71.1% of the candidate population. The remaining 93 IP addresses, or 28.9%, produced no matching domain records under the same filter at that observation time. This demonstrates that the underlying infrastructure is dynamic.

Candidate IP population	Currently domain-associated	No current domain association
322	229 71.1%	93 28.9%



Figure 3 Overlap between service-responsive candidate infrastructure and domain-associated infrastructure observed on September 24, 2026 at 16:00 UTC

METHODOLOGY NOTE

The domain query was restricted to the predefined candidate IP population. This snapshot measures continued domain association within that candidate set and does not identify newly emerged infrastructure outside it.

To determine whether the identified infrastructure reflected ordinary shared-hosting activity or a broader malicious network, we expanded the analysis to associated domains. An earlier collection window on September 11, 2026 identified more than 10,000 associated domains. The most recent snapshot, collected on September 24, 2026 at 16:00 UTC, contained 5,785 unique domains associated with 229 of the candidate IP addresses. The difference between these populations shows substantial domain and IP association churn.

Of the 10,187 domains in the previous collection window, 9,280 (90%) fell within the highest malicious category, as they follow typical typosquatting and other malicious behavior. Such a distribution is uncommon for a general-purpose hosting environment, where benign and malicious content would normally coexist in significantly more balanced proportions. The same is seen in the most current snapshot.

04 / TARGETING

Regional Targeted Brands

The Check Point Exposure Management Research Team then analyzed the most recent sample of 500 domains from the cluster and identified the following targeted industries:

Southeast Asia

Strong concentration of banking, financial services, and customer loyalty or rewards program impersonation themes

Europe

Transportation, parking management, taxation and public administration, logistics and parcel delivery services, and regulatory enforcement-related themes

Latin America

Financial institutions, municipal and public payment services, government agencies, postal services, social benefit programs, and rewards redemption campaigns

Israel

Digital mobility and parking services, law enforcement notifications, and postal service impersonation themes

Australia

Travel, aviation loyalty programs, and retail rewards ecosystems

Africa and Asia

Postal and courier services, transportation and traffic enforcement authorities, tolling and road payment systems, banking and financial institutions, telecommunications providers, and digital service platforms

We also detected generic, repeated keywords related to parking, fines, payments, banks, rewards, card collection, postal and parcel fees, and government-associated tax, refund, and violation themes.




Figure 4 Infrastructure discovery flow: 5,785 total domains (high-confidence phishing)

05 / CASE STUDY

Application-Level Fingerprinting and Operational Infrastructure

Case Study: Rewards-Themed Campaign Delivery Through Regional Path Activation

We investigated a randomly selected phishing domain from our identified cluster. It was linked to a major financial group operating across Central America, with significant operations in Costa Rica, Panama, and other countries, including the United States.

During the investigation, we identified a recently observed domain, ***canjes*[.]click**, that exhibited the same concealed content delivery model previously observed across unrelated phishing campaigns.

The domain name itself provided an initial clue regarding the likely target. The Spanish word **"canjes"** is commonly associated with points redemption, rewards, and loyalty-program exchanges. Combined with the abbreviation, the naming convention strongly suggested a possible attempt to impersonate a large Central American bank. This resembled previously observed regional lures that combined local-language terminology with recognizable financial or government entities.

To evaluate this hypothesis, several regional directory identifiers consistent with the bank's operating footprint were tested. Initial attempts using regional paths such as **/cr** (Costa Rica) returned the same generic **"Not Found"** response observed at the root of the domain, similar to the initially observed Italian-targeted use case.




Figure 5 *canjes*[.]click/cr returns a 'Not Found' error

However, when the **/pa** path was requested, the site initialized a custom security-verification interface, indicating that phishing content was not exposed directly through the root domain and instead required access through a region-specific path.




Figure 6 *canjes*[.]click/pa custom validation layer to block security crawlers

The investigated infrastructure employed an intermediate validation layer before delivering the operational phishing workflow. The root page returned a generic **"Not Found"** response, while the regional path /pa exposed a localized access-control interface. The page presented a human-verification step and, upon interaction, requested a dynamically generated session token from the backend before redirecting the visitor to the campaign URL.

```

147     if (eatToken) {
148         callback(eatToken);
149     } else {
150         var callbackName = 'jsonpCallback_' + Date.now();
151         window[callbackName] = function(data) {
152             delete window[callbackName];
153             callback(data.token || null);
154         };
155         var script = document.createElement('script');
156         script.src = window.location.origin + '/eat?callback=' + callbackName;
157         script.onerror = function() {
158             delete window[callbackName];
159             callback(null);
160         };
161         document.head.appendChild(script);

```

Figure 7 Dynamically generated session token section

Interestingly, the website implemented multi-language support covering a broad range of locales, including English, Spanish, French, German, Portuguese, Arabic, Hebrew, Japanese, Korean, Hindi, Russian, and Chinese.

```

121
122   for (let lang in langMap) {
123     if (userLang.startsWith(lang)) {
124       return langs[langMap[lang]] || langs[1];
125     }
126   }
127   return langs[1];
128 }
129
130 const texts = {
131   title: { text: "安全验证|Security Check|Verificación de Seguridad|Vérification de Sécurité|Sicherheitsüberprüfung|セキュリティ",
132   message: { text: "请完成下方验证以继续访问|Please complete the verification below to continue|Por favor complete la verificación",
133   checkbox: { text: "我不是机器人|I'm not a robot|No soy un robot|Je ne suis pas un robot|Ich bin kein Roboter|私はロボットではない",
134   success: { text: "验证成功! 正在跳转...|Verification successful! Redirecting...|Verificación exitosa! Redirigiendo...|Vérification réussie! Redirigeant...",
135   footer: { text: "由安全服务提供保护|Performance & security by Cloudflare|Rendimiento y seguridad por Cloudflare|Performance et sécurité par Cloudflare",
136 };
137 document.title = getLangText(texts.title.text, texts.title.isML);
138 document.getElementById('pageTitle').textContent = getLangText(texts.title.text, texts.title.isML);
139 document.getElementById('pageMessage').textContent = getLangText(texts.message.text, texts.message.isML);
140 document.getElementById('checkboxLabel').textContent = getLangText(texts.checkbox.text, texts.checkbox.isML);
141 document.getElementById('success').textContent = getLangText(texts.success.text, texts.success.isML);
142 document.getElementById('footerText').textContent = getLangText(texts.footer.text, texts.footer.isML);
143
144 function getOrFetchToken(callback) {
145   var urlParams = new URLSearchParams(window.location.search);
146   var eatToken = urlParams.get('_eat');
147
148   if (eatToken) {
149     callback(eatToken);
150   } else {
151     var callbackName = 'jsonpCallback_' + Date.now();
152     window[callbackName] = function(data) {
153       delete window[callbackName];
154       callback(data.token || null);
155     };
156     fetch('/api/token?callback=' + callbackName);
157   }
158 }
159
160 // Redirect to the next stage of the workflow
161 window.location.href = window.location.protocol + '//' + window.location.hostname + '/next-stage?token=' + token;
162
163 // Clean up the window
164 delete window['_eat'];
165 delete window['_token'];
166 delete window[callbackName];
167
168 // End of the script
169
190 lines, 8550 characters selected

```

Figure 8 Various supported languages

Additionally, the verification process was not a traditional CAPTCHA challenge. Instead, the page generated or retrieved an `_eat` session token and appended it to the redirect URL before granting access to the next stage of the workflow.




Figure 9 Main phishing lure page




Figure 10 Exfiltration overlay

This delivery model introduces an additional layer that was not observed in all previously documented examples. Unlike the earlier observed campaigns, where access to the phishing application was obtained directly through a regional directory, the bank-themed campaign employed an intermediate access-control page that generated a session token before exposing the phishing workflow. While similar campaigns relied on concealed regional paths and conditional content delivery, the token-gated verification page may represent either a newer evolution of the platform or a modular component that can be selectively enabled across different deployments.

Both use a tokenized API (`/api?token=<UUID>`) that returns a JSON configuration object with an identical schema controlling rewards, regional settings, and multi-step card/OTP verification flows. This strongly indicates a shared operator or toolkit, with only branding, language, and reward catalog swapped per targeted bank and region.

Beyond the regional path and token-gated access-control layer, the bank-themed campaign reveals a deeper architectural overlap with a separate rewards-themed phishing operation targeting a Philippine bank. In both cases, the phishing application is delivered as a Vue-based single-page app that communicates with a tokenized backend API of the form:

```
https://*canjes*[..]click/JsIRZLhQrF/api?token=<UUID>
https://b*****[.]com/xtvjuQQNDV/api?token=<UUID>
```

Each API call returns a JSON configuration object containing a field named **custom**, which is itself a JSON-encoded string. When decoded, the **custom** object in both campaigns shares an **identical schema**:

```
{
  "address_msg": "",
  "black_msg": "Error en el pago. Intente con otra tarjeta de crédito o American Express.",
  "customInput": [],
  "customOp": [
    { ... PIN page template ... },
    { ... Email code page template ... }
  ],
  "deny_c_msg": "",
  "deny_d_msg": "",
  "error_card_msg": "Error en la tarjeta. Por favor, verifique e ingrese correctamente la información de su tarjeta.",
  "goods": {
    "addressTheme": "1",
    "cardTheme": "1",
    "feeType": "1",
    "format": "1",
    "homeTheme": "1",
    "list": [],
    "payTheme": "1",
    "points": "6158",
    "theme": "1",
    "unit": "B/."
  },
  "goodsVersion": 5,
  "pay_amount": "",
  "pay_msg": "",
  "white_msg": ""
}
```

Figure 11 JSON schema for *canjes*[.]click/JsIRZLhQrF/api?token=<UUID>

```
{
  "address_msg": "",
  "black_msg": "We're unable to verify this card at this time. Please try using a different card.",
  "customOp": [],
  "deny_c_msg": "Please redeem before 10:00 AM due to scheduled system maintenance.",
  "deny_d_msg": "We couldn't verify your card. Check that your mobile number is up to date and your card is active.",
  "error_card_msg": "...",
  "goods": {
    "addressTheme": "1",
    "cardTheme": "1",
    "feeType": "1",
    "format": "1",
    "homeTheme": "1",
    "list": [ /* 10 reward items, iPhones/PS5/etc. */ ],
    "payTheme": "1",
    "points": "8413",
    "theme": "1",
    "unit": "P"
  },
  "goodsVersion": 5,
  "hiddenDefaultButtons": ...,
  "pay_amount": "",
  "pay_msg": "...",
  "white_msg": ""
}
```

Figure 12 JSON schema for b*****[.]com/xtvjuQQNDV/api?token=<UUID>

During the investigation, we identified another domain impersonating a large consumer-facing transportation and service platform that exhibited a different implementation of the same regional content-delivery concept observed elsewhere in this infrastructure cluster. Rather than exposing localized phishing content through region-specific URL paths, the site accepted explicit country and locale parameters. Successful responses were observed for multiple Northern European countries, including Finland, Sweden, Denmark, Norway, and Germany, as well as language-country locale pairs corresponding to those regions. This indicates that the platform supports geographic and language-aware content delivery, allowing localized phishing experiences to be served dynamically based on user context.

While the implementation differs from the path-based activation model observed in previous case studies, the underlying operation remains consistent: exposing region-specific phishing workflows only after a country or language context has been established. The variation in routing logic is consistent with a modular phishing platform, where localization and traffic-routing components can be adapted to

different impersonated organizations and geographic regions while preserving the same underlying operational model.




Figure 13 Identical custom validation layer observed previously




Figure 14 Phishing lure for a different targeted brand

● Case Study: Conclusions


The structural equivalence of the **custom** configuration object, combined with the shared tokenized API pattern and API path framework, strongly indicates a common underlying toolkit or operator. The platform appears to support modular regionalization, allowing the threat actor to reuse the same core exfiltration logic while swapping branding, language, currency, and reward catalog per targeted institution and country.

06 / HOSTING

Infrastructure Ownership Breakdown

Hosting Concentration by Provider / ASN

~75% of observed IPs sit on two Chinese hyperscalers

**Tencent Cloud ~65%**

AS132203, AS45090

Alibaba Cloud ~10%

AS45102

Other providers ~25%

VPS/hosting: : DigitalOcean, IONOS, Contabo, M247, NTT America + regional networks

Smaller VPS/hosting: FranTech Solutions, Hostwinds, HostPapa/RackNerd, CTG Server

©2026 Check Point Software Technologies Ltd.

Figure 15 Hosting provider and ASN ownership

The IP dataset attributed to this phishing-as-a-service (PhaaS) operation shows that it is centralized. An overwhelming majority of the observed IP addresses reside within two major Chinese cloud ecosystems: Tencent and Alibaba. Infrastructure associated with Tencent ASNs (AS132203, AS45090) accounts for nearly two thirds of all IPs, while Alibaba's AS45102 contributes an additional 9–10%.

This heavy reliance on large commodity cloud providers offers the operators significant operational advantages: rapid provisioning, global network reach, and a high degree of plausible deniability behind reputable brands.

Beyond the core Chinese cloud backbone, we observe a long tail of smaller VPS and hosting providers supporting the operation, including FranTech, Hostwinds, HostPapa/RackNerd, CTG Server, and Contabo. These platforms are typically associated with low-cost, low-friction hosting and have historically been favored by threat actors for staging, redirector infrastructure, and short-lived campaigns. Together, this peripheral ecosystem suggests a layered infrastructure model: core phishing content and customer-facing PhaaS services deployed on large cloud providers, backed by auxiliary VPS nodes used for redirection, obfuscation, and operational resilience.

07 / ANALYSIS

Comparing the Modus Operandi to Known PhaaS Groups

We cannot conclusively attribute this infrastructure to a specific phishing kit family, service provider, operator, or threat group based on the currently available evidence. No administrative panel, operator identifiers, source-code provenance, or other attribution artifacts were recovered during this investigation. However, analysis of multiple independently observed campaigns revealed a consistent operational model.

Across unrelated lures, brands, languages, and regions, phishing content was commonly concealed behind regional activation mechanisms rather than exposed directly on the root domain. The delivery architecture supported multiple localization methods:

- Region-specific directories such as /it, /de, and /pa
- Explicit country parameters such as country=FI and country=DE
- Locale-based targeting through language-country pairs such as fi-FI, sv-SE, and de-DE
- Session-gated delivery using dynamically generated access tokens
- Multilingual user interfaces capable of adapting phishing content to different regions and languages


Although the victim-facing brands, themes, and languages varied significantly, the underlying application architecture remained highly consistent. Multiple campaigns exposed tokenized backend APIs, dynamically generated session identifiers, region-aware configuration objects, and structurally equivalent application workflows. The strongest overlap occurred within the application layer itself, where separate campaigns used equivalent configuration schemas, shared API design patterns, and similar content-delivery logic.

These findings suggest that the observed infrastructure was not constructed as a collection of independent phishing sites. Instead, it appears to operate as a reusable, modular platform capable of supporting numerous localized campaigns while preserving a common technical foundation.

08 / TAKEAWAYS

Conclusions

We broke the concealed phishing platform down into five layers, demonstrating repeatable methods for identifying related deployments even when the phishing content itself is hidden.



The most durable indicators are not the domains themselves, but the platform's repeated behaviors: concealed regional delivery, tokenized access mechanisms, randomized API paths, reusable configuration schemas, localization frameworks, and session-controlled phishing workflows. These characteristics allow us to identify related deployments even as the visible infrastructure changes over time.

09 / GUIDANCE

Recommendations and Attribution Limitation

Detection and Hunting

Domains returning generic 404 or Not Found responses may still host active phishing content under concealed paths or after conditional routing. These are usually not detected through traditional phishing discovery methods and depend on capture at the distribution layer, or on infrastructure clustering once hunting artifacts are found.

Hunting logic should combine multiple signals, such as:

- | | |
|--|--|
| ☐ Server response behavior combined with server technology | ☐ Stable Trace-Id formatting and repeated response structure |
| ☐ Concealed regional directories | ☐ Country or locale query parameters |
| ☐ Token-generation endpoints and unique redirect parameters | ☐ Randomized API paths followed by /api |
| ☐ UUID-like session tokens | ☐ Repeated configuration-object field names |
| ☐ Reused scripts, normalized JavaScript, analytics identifiers, and asset-directory structures | |

HOW TO APPLY

Use the server fingerprint to identify candidates. Use application-level artifacts to raise confidence and assign related deployments to a technical family.

● Attribution Limitation

Application-layer equivalence, together with other metrics, can support a shared phishing toolkit finding. It cannot independently establish operator identity. Named attribution should require additional evidence such as panel access, developer identifiers, marketplace advertising, source-code origin, administrative infrastructure, or independently verified law enforcement and intelligence reporting.

Phishing platforms like this one hide in plain sight, so we watch for the signs that give them away. Check Point Exposure Management monitors for lookalike domains and brand impersonations across the web, social media, and app stores, so your security team sees abuse as it appears. Our proprietary Phishing Beacon sends a signal within seconds of a clone of your website going live. Once a threat is confirmed, our in-house takedown team works to remove malicious domains, phishing pages, and fake profiles.

[See how Check Point Exposure Management protects your brand from impersonation →](#)

GET IN TOUCH

Contact Us

ISRAEL

Tel: +972-73-226-4555

5 Shlomo Kaplan Street
6789159

USA

Tel: 1-800-429-4391

100 Oracle Parkway, Suite 800
Redwood City, CA 94065

UK AND IRELAND

Tel: +44-20-7628-4211

85 London Wall, 4th Floor
London, EC2M 7AD

SINGAPORE

Tel: +65-6435-1318

78 Shenton Way, #09-01 Tower 1
Singapore 079120

PHILIPPINES

Tel: +63-2-8465-9200

Unit 2005, 20th Floor, Zuellig Building,
Makati Avenue, corner Paseo de
Roxas, Makati City 1223, Metro Manila

JAPAN

Tel: +81-3-6205-8340

Toranomon Kotohira Tower 25F, 1-2-8,
Toranomon, Minato-ku, Tokyo 105-
0001

About Check Point Exposure Management

Check Point's exposure management changes the game.

We combine billions of internal telemetry points with billions of external signals from the open, deep, and dark web to deliver a unified intelligence fabric. This provides clear visibility across the full attack surface, including brand risk.

The industry is moving from fragmented feeds to real context and real priorities. We support that shift through active threat validation, confirmation of compensating controls, and deduplication across tools, so teams can focus on what actually matters.

With safe-by-design remediation, fixes aren't just assigned, they're implemented. Every fix is validated before enforcement, enabling measurable risk reduction without downtime.

Gartner predicts organizations adopting continuous threat exposure management with mobilization will see 50% fewer successful attacks by 2028. We're leading that shift with action, not just tickets, and Fortune 500 organizations across major industries already rely on Check Point Exposure Management.

For more information visit: checkpoint.com/exposure-management



EXPOSURE MANAGEMENT

© Check Point, 2026. All Rights Reserved.