CSURFACE

WHITEPAPER · ATTACK SURFACE MANAGEMENT

The real attack surface and what the official inventory does not record

Discovery: how to identify what subdomain enumeration platforms do not reveal — Shadow IT, inherited infrastructure, undocumented APIs and the digital supplier chain.

CAPABILITY Attack Surface Management	AUDIENCE CISO · Security Engineering	EDITION 2026
READING 11 pages	CLASSIFICATION Public	SOURCE CSURFACE Platform

EXECUTIVE SUMMARY

Asset discovery is an intelligence problem

The thesis of this whitepaper, on a single page — the essentials before the technical detail.

INCIDENTS FROM UNMANAGED ASSETS

74%

of organizations reported incidents caused by unknown or unmanaged assets[1]

ATTACKS VIA EXPOSED ASSET

69%

suffered an attack originating from an internet-exposed asset that was unknown or unmanaged[2]

DISCOVERY WITH MACHINE LEARNING

2–3×

more assets than subdomain enumeration — and, on dispersed surfaces, far beyond

Most Attack Surface Management (ASM) tools start from a fragile premise: that enumerating subdomains from public lists is sufficient to cover the attack surface. Practice demonstrates the opposite. Most of the real surface lies outside those lists — **Shadow IT contracted without the knowledge of the IT department, infrastructure inherited from mergers and acquisitions, undocumented APIs and exposed staging environments**.

Independent market research converges on this diagnosis. A global survey by Trend Micro with more than two thousand security leaders indicates that **74% of organizations suffered incidents caused by unmanaged assets**[1]; a study by Enterprise Strategy Group records that **69% were the target of an attack originating from an exposed and unknown asset**[2]. What remains outside the inventory constitutes the portion of the surface over which there is no observation — the operational blind spot on which no control can act.

This whitepaper presents the CSURFACE approach: **Discovery**, which combines passive collection of public signals, correlation through a stack of Machine Learning engines, Shadow IT detection and mapping of the digital supplier chain — with no agent, no installation and no manual work per asset.

The thesis. Asset discovery is an intelligence problem. Correlating signals through Machine Learning reveals what isolated sources — such as public DNS — do not reach. And what remains outside the inventory remains without protection.

What this document covers

- Why the attack surface grew — and why the tools did not keep up.
- The four structural gaps of wordlist enumeration.
- The four techniques that make up Discovery, with emphasis on the correlation engine.
- A representative use case and the expected outcome indicators.

THE LANDSCAPE · PART 01

The surface expanded — the tools did not keep up

What the attack surface is composed of in 2026 and why the traditional inventory is already born incomplete.

ATTACK SURFACE · ~2014

Small and stable

Public DNS | Email servers | VPN

→

ATTACK SURFACE · 2026

Dispersed across the entire operation

Public DNS | Email | VPN | SaaS & Shadow IT | Undocumented APIs | Multiple clouds | M&A legacies | Supplier chain | Ephemeral & low-code environments

74%
of organizations suffered incidents caused by unknown or unmanaged assets[1]

69%
were the target of an attack originating from an exposed and unknown asset[2]

~40%
of corporate infrastructure remains invisible to the IT team[4]

Why the surface grows ungoverned

The external attack surface is not an organized inventory — it accumulates. Each product team publishes an application, each business unit contracts a SaaS, each acquisition incorporates infrastructure inherited from another company. The absence of formal decommissioning processes means that this accumulation is rarely reversed.

According to Unit 42 (Palo Alto Networks), an organization's surface incorporates, on average, **more than 300 new services per month**[5]. A manually maintained inventory does not keep up with this pace: it becomes incomplete as it is completed.

The result is a structural gap between the surface the organization *believes* it has and the one it actually exposes. Assets absent from the inventory remain without observation, without prioritization and without response — and constitute the attacker's preferred entry vector.

Why this matters. The inventory is the foundation of every security program. Without knowledge of what exists, prioritization, testing and incident response operate on an incomplete basis — and what was left out remains exposed.

The cost of not seeing. The average cost of a data breach in the United States reached **$10.22M** in 2025, according to the IBM Cost of a Data Breach 2025[3]. In the United States, breach exposure arises from state breach-notification laws, FTC enforcement, and state statutes such as the CCPA (statutory damages of $100–$750 per consumer per incident). Continuous discovery reduces the probability of an incident by eliminating blind spots — and each reduction in probability corresponds to a direct and proportional reduction in the expected financial exposure.

CURRENT LIMITATIONS · PART 02

Why subdomain enumeration is not enough

The ASM platforms on the market are based on subdomain enumeration — a method that resolves the obvious case and fails where the risk concentrates.

What subdomain enumeration platforms do

The ASM platforms available on the market are based, in essence, on subdomain enumeration: lists of common prefixes (**mail.**, **dev.**, **staging.**, **api.**), permutations of the company name and queries to public sources.

This method locates **mail.empresa.com** and **staging.empresa.com**. It does not locate:

- **api-gtw-prd-v3-fc8a.empresa.com** — automatically generated identifier.
- **tenant-b3b9e0.empresa.com** — multi-tenant subdomain.
- A domain acquired in a merger, with a distinct naming convention.
- SaaS hosted by third parties — **empresa.atlassian.net**, **empresa.salesforce.com**.

The blind spot is structural. Subdomain enumeration locates only what someone previously imagined would exist. The relevant risk lies precisely in what was not anticipated — and which, for that reason, was not sought.

The four gaps of subdomain enumeration

1. **Shadow IT in third-party SaaS.** The marketing, HR and sales departments contract tools that create public URLs linked to the company. No generic list identifies them.
2. **Proprietary naming conventions in the cloud.** Each team provisions resources with its own naming pattern, to which a generic list has no access.
3. **Legacies from mergers and acquisitions.** The domains of acquired companies remain active for years, without governance — and subdomain enumeration does not know the incorporated organization.
4. **Ephemeral environments.** Pull requests generate preview environments with unique URLs, and CI/CD pipelines create hostnames at each execution.

Each of these gaps represents, in practice, an entry point absent from any inventory — uncataloged, unmonitored and therefore uncorrected.

THE CSURFACE APPROACH · PART 03

The correlation engine — when WHOIS is not enough

Passive collection of public signals and correlation through Machine Learning — the first two layers of the Discovery.

Whitepaper 2026
CSURFACE

DISCOVERY · THE DISCOVERY PROCESS

1. **Passive collection** — Public signals, without touching the infrastructure.
2. **Active collection** — Active and non-intrusive verification of the candidates.
3. **ML correlation** — Estimates the probability of ownership of each asset.
4. **Agentic decision** — The correlation layer resolves the ambiguous cases and delivers ready-to-use inventory.
5. **Technology stack** — Active identification of the technologies of each asset.
6. **Vulnerabilities and exposure** — Mapping of the flaws and exposures of each asset.

WHOIS has ceased to be proof of ownership. Currently, most domains operate under privacy protection or redaction of registration data, which keeps the real holder hidden. A WHOIS query, in isolation, no longer establishes to whom an asset belongs — determining ownership has come to require inference, not a simple query.

01. Passive collection of signals

The discovery begins in an entirely passive way. CSURFACE gathers and cross-references a broad set of public information available on the internet — a universe of signals considerably broader than the public DNS on which traditional platforms rely.

No packet is sent to the organization's infrastructure at this stage: the collection is non-intrusive and generates no noise on the analyzed assets. This broad evidence base is the raw material on which the correlation engine operates.

The broader and more diverse the signal base, the greater the correlation engine's ability to distinguish what belongs to the organization from what does not — with the consequent reduction of incorrect attributions and expansion of the real coverage of the surface.

02. The correlation engine — a Machine Learning stack

If WHOIS no longer answers the question of ownership, CSURFACE answers it by correlation. The public information collected feeds a **stack of Machine Learning engines** — each with its own characteristics, trained for a distinct reading of the data — that operate in combination.

The objective is common to all: to identify patterns and estimate the **probability of an asset belonging to the organization** analyzed. The final decision results from the consensus of the stack — it is this collective arrangement, and not an isolated model, that sustains the precision.

Each attribution comes accompanied by the **probability** and the evidence that supports it — it is an auditable result, which the security team can verify and contest.

Measured precision. In combined operation, the correlation stack operates **with high confidence**. With this level of confidence, the platform identifies the ownership of domains, organizations, web applications, FQDNs, IP addresses and other asset classes.

WHAT CSURFACE OBSERVES · THE PLATFORM'S OWN DATA

122k
external assets discovered across 68 organizations analyzed by the platform

1,068
related organizations — subsidiaries and affiliates — revealed by the discovery

up to 173
related organizations mapped from a single root domain

THE CSURFACE APPROACH · PART 03

Shadow IT, supplier chain and the measurable result

The two remaining techniques — and what the approach changes, side by side with traditional enumeration.

03. Shadow IT detection

The platform identifies the third-party SaaS connected to the organization — tools contracted by business units that create a public presence linked to the company. This Shadow IT, invisible to the official inventory, is recognized by the correlation of multiple public indicators that associate it with the client.

04. Mapping of the supplier chain

For each asset discovered, the platform maps the external APIs called, the third-party CDNs and JavaScript, the shared cloud providers and the embedded libraries — the digital supplier chain that expands the surface without appearing in any inventory.

A dedicated specialist. For the large security platforms, the central product is the scanner and vulnerability management; the ASM module was added later, as an accessory component that merely feeds the pre-existing engine — and in which the speed and precision of discovery are not a priority. CSURFACE adopts the opposite position: discovery is the product. The platform was conceived as an instrument dedicated to keeping the client aware of their exposure in the shortest time and with the greatest possible precision.

Measurable differentiators

CRITERION	SUBDOMAIN ENUMERATION	CSURFACE · ML DISCOVERY
Typical surface coverage	Partial	**Broad, context-aware**
Shadow IT detection	Limited or nonexistent	**Active**
Manual effort per asset	3–10 hours	**Zero**
Frequency	One-off / scheduled	**Continuous, 24/7**
Supplier chain depth	One level, when present	**Up to three levels**
Required input	Known asset list	**Only the root domain**

Ranges observed in discoveries conducted by the CSURFACE platform; see the methodological note on page 10. The result varies according to the size and dispersion of each organization's surface.

APPLICATION AND RESULTS · CASE 01

An ERP startup: from 84 to 570 assets

What a vulnerability management inventory left out — and which continuous discovery made visible and prioritizable.

INVENTORY VIA VULNERABILITY MANAGEMENT

84

assets — what the company believed to be its entire external exposure.

→

MAPPED BY CSURFACE

570

assets on the real surface — **nearly 7×** the known inventory.

The case

An ERP software startup tracked its external exposure through a traditional vulnerability management platform. The inventory recorded **84 assets**, and the team believed that this number represented the entirety of its attack surface.

Within a few minutes of discovery with CSURFACE, the real surface revealed itself to be composed of **570 assets**. Among what was exposed and off the radar:

WARNING **Expired digital certificates**
On services still accessible over the internet.

WARNING **End-of-life (EoL) software**
Legacy assets without vendor support or security patches.

CRITICAL **Invisible vulnerabilities**
On assets the inventory did not even know — hence never assessed.

CRITICAL **Exposed sensitive data**
Configuration errors that exposed data from internal systems.

CRITICAL **Exposed administrative panels**
High risk; they should be restricted to the internal network.

Together, these findings configured a risk profile significantly higher than the level acceptable to the organization — without any of them having been cataloged or assessed.

Why it went unnoticed

A traditional vulnerability management platform examines only what it is told: it enumerates a supplied list, while contextual discovery maps the real surface. Everything that was not on the list of 84 assets remained, by definition, invisible.

A recurring pattern. According to Unit 42 (Palo Alto Networks), an organization's attack surface incorporates, on average, **more than 300 new services per month**[5]; and more than **25% of exposures** involve critical IT infrastructure — including administrative login pages accessible over the internet[5]. A static inventory becomes progressively incomplete as the surface evolves.

The result

The immediate gain was **visibility**: from 84 presumed assets to 570 real ones, the risk ceased to be invisible and became prioritizable. Effective exposure management presupposes, as a prior condition, complete knowledge of the surface to be managed.

APPLICATION AND RESULTS · CASE 02

A large bank: entire companies outside the inventory

Growth through acquisitions, a constantly changing surface and excessive reliance on traditional tools.

The starting point

The bank came to CSURFACE with three open questions:

No visibility of the external surface
Lack of control over what was actually exposed.

Excessive reliance on the traditional
The existing tools only saw the already known inventory.

Risk in the chain and in acquisitions
Each acquired company expanded the surface without governance.

What the discovery revealed

A large bank — recognized for its work in asset management and for a history of major acquisitions in recent years, from logistics operations to hotel chains — coexisted with a surface in permanent change. The corporate environment, the multiple clouds and each new acquisition became the responsibility of the cybersecurity team.

After a few weeks with the CSURFACE platform, the bank identified **structures of entire companies — inside and outside the country — that were not on its radar**. Some of them presented vulnerabilities and exposures that required urgent intervention.

The volume and dispersion exceeded internal estimates: assets distributed across multiple clouds and entire corporate structures that no team had mapped, with no designated owner, no monitoring and no established response process.

Remediation within the critical window. On one of the assets, the platform detected a vulnerability with an impact of **9.8 (CVSS)** on the same day as its public disclosure. The security team was alerted and remediated it still within the critical period — when the flaw was already being exploited on the internet. A potentially severe incident was avoided.

Adoption by governance. The information security governance (ISG) teams began to use the platform to document and understand organizations and assets over which, until then, there was no visibility.

The reading

In an organization that grows through acquisition, the attack surface changes faster than any manual inventory process. Continuous discovery ceases to be a convenience and becomes a condition for governing risk.

APPLICATION AND RESULTS · COMPLEMENTARY CASE

Anatomy of a perfect phishing — avoided

How an unmanaged subdomain and an outdated SPF record created the conditions for a phishing using the bank's legitimate infrastructure.

The hijacking of a domain authorized in the SPF

The risk arises from a subtle distinction between **two different domains**: the **bank's own subdomain**, which contains the SPF record, and the **domain of a third party**, which that record authorizes to send emails. When the second is abandoned, the first continues to trust it.

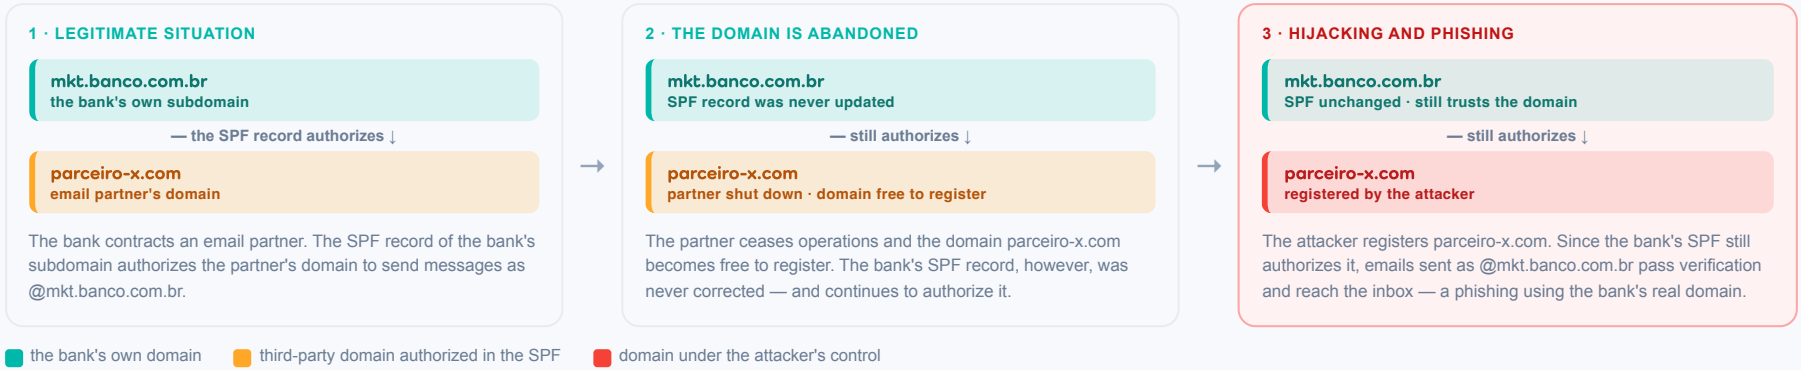

The composite vector — hijacked SPF combined with subdomain takeover

The platform identified, at the same bank, **subdomains susceptible to takeover** — pointing to deactivated cloud services and, for that reason, claimable by third parties. Combined, the two flaws would allow an almost undetectable phishing: an email sent from a legitimate bank domain — via hijacked SPF — containing a link to another legitimate bank domain — via hijacked subdomain. To the account holder, every signal would point to the real bank.

The outcome. Neither of the two flaws was visible to the bank's traditional tools — both lived on assets outside the inventory. CSURFACE flagged them, and the bank corrected the records and the subdomains before they were found. The real SubdoMailing campaign — which reached the point of sending millions of emails per day from hijacked legitimate domains[6] — confirms that the threat is real.

METHODOLOGY AND SOURCES

The origin of each number

How this document distinguishes public market research from metrics observed by the platform.

Methodological note

This whitepaper combines two natures of data, kept distinct throughout the text:

Market statistics. Attributed, at the point of use, to public third-party research and identified by a numeric marker. The complete references appear alongside.

Platform metrics. Indicators such as the broad coverage of the surface and the material gain in discovered assets result from discoveries conducted by the CSURFACE platform itself. They vary according to the size and dispersion of each organization's surface.

No statistic is presented without a source. The cases described are based on real engagements and are anonymized — no organization is identified.

Educational material. This document is intended for informational purposes. The methodologies described reflect the state of the CSURFACE platform on the date of publication.

References

1. **Trend Micro.** Global survey on security incidents caused by unmanaged assets, conducted with more than 2,000 security leaders. 2025.
2. **Enterprise Strategy Group.** Research on the exploitation of internet-exposed assets that are unknown or unmanaged. Reported by SC Media.
3. **IBM.** Cost of a Data Breach Report 2025 — average cost of a data breach in the United States: $10.22M.
4. **Gartner.** Analyses on corporate infrastructure visibility and External Attack Surface Management (EASM).
5. **Unit 42 (Palo Alto Networks).** 2024 Attack Surface Threat Research — exposures of internet-connected infrastructure and the growth pace of the attack surface.
6. **Guardio Labs.** SubdoMailing — investigation into the hijacking of subdomains and SPF records for large-scale phishing campaigns. 2024.

The market statistics cited are public in nature. The percentages may vary according to the methodology and the sample scope of each study.

NEXT STEPS

Discovery begins with a single domain

Discovery is one of the capabilities of the CSURFACE platform — a Continuous Exposure Management platform. In an integrated platform, it brings together continuous discovery of the external surface, analysis of the digital supplier chain, exploitability validation, threat intelligence with dynamic prioritization and monitoring of leaked credentials. It operates in an entirely external manner — with no agent and no installation — with optional cloud, WAF and CIEM integrations for organizations that require greater analytical depth.

Receive the preliminary analysis

Provide your company's domain and receive, at no cost, a snapshot of the exposed assets discovered by the platform.

Use the risk calculator

Estimate the annual loss expectancy (ALE) and the VaR for your sector, with a FAIR methodology calibrated by market benchmark.

Read the next whitepaper

Digital Supplier Chain — the Polyfill.io case and how to map dependencies across three levels.

See your real surface — discovered with the root domain as the only input

Receive free preliminary analysis

© 2026 CSURFACE · Educational material. Market statistics with public sources indicated on page 10.