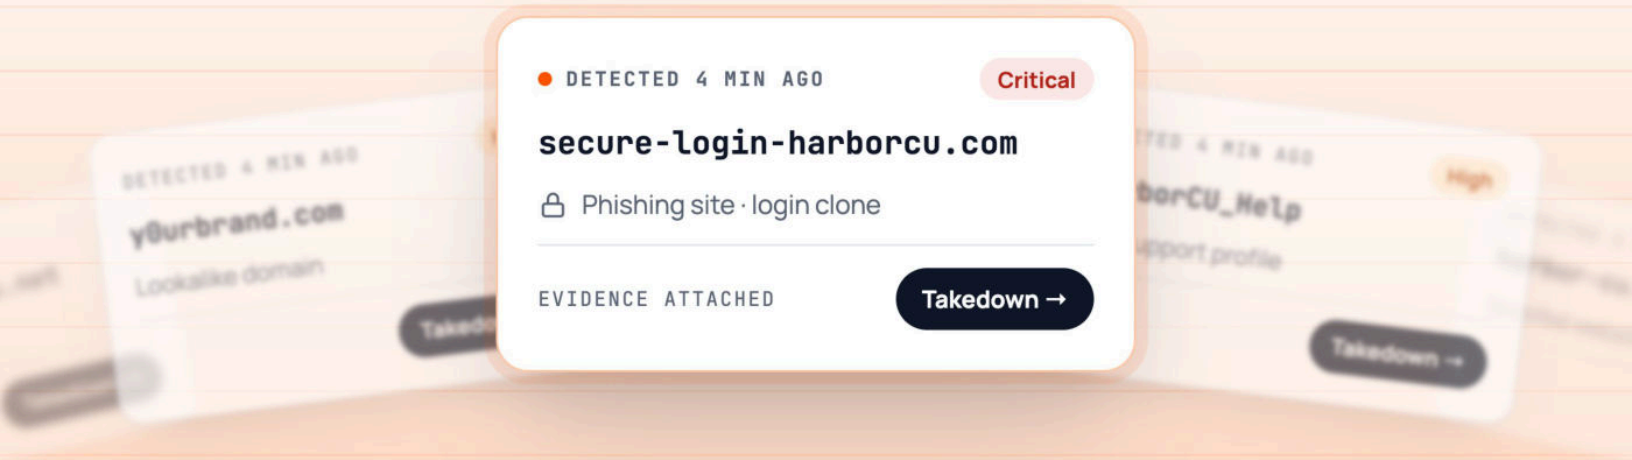
BRAND THREAT IDENTIFICATION & TAKEDOWN

Lookalike, Phishing, or Impersonation?

The 2026 Brand Threat Identification & Takedown Guide.

How to tell the three external brand threats apart, why detection without fast removal fails, and a program blueprint with the KPIs that prove it works.



INTRODUCTION

The AI-driven external threat landscape

Why brand impersonation became a board-level risk in 2026.

Attackers don't need to breach you to hurt you. They copy your brand outside your perimeter, on domains, sites, and profiles you don't control, and point your customers at them.

Generative AI has collapsed the cost of cloning a brand. Most security stacks watch the inside (email, endpoints, network), not the open web where impersonation lives.

#1

CEO CYBER CONCERN

The CEOs surveyed for the WEF Global Cybersecurity Outlook 2026 ranked cyber-enabled fraud and phishing as their top cyber concern for 2026, up from second in 2025.

Source: World Economic Forum, Global Cybersecurity Outlook 2026

THE FRAUD IS PERSONAL, AND GROWING

73%

personally affected, or know someone affected, by cyber-enabled fraud in the past year

Source: WEF Global Cybersecurity Outlook 2026

62%

of those cited phishing, vishing, or smishing

Source: WEF Global Cybersecurity Outlook 2026

971,181

phishing attacks observed in Q1 2026, up 13.8% quarter over quarter

Source: APWG Phishing Activity Trends Report, Q1 2026

CONTENTS

What's inside

Three threat vectors, the gap that lets them win, and a program to close it.

01 p. 04

The Lookalike Domain Vector

The infrastructure impersonation campaigns are built on.

02 p. 06

The Phishing Site Vector

Live pages that harvest credentials, cards, and MFA codes.

03 p. 08

The Impersonation Vector

Fake profiles that lure customers off your channels.

04 p. 09

The Takedown Gap

Why detection without fast removal fails.

05 p. 12

Building a Brand Protection Program

Five pillars and the KPIs that prove they work.

06 p. 13

2026 Brand Threat Comparison Matrix

All three vectors, side by side.

ALSO INSIDE Anatomy of an attack p. 05 Why speed decides p. 07 The resource reality p. 11 Where Sentrike fits p. 14



Use this guide to brief your board, align security and fraud teams, and benchmark your response times.

SECTION 01

The Lookalike Domain Vector

The infrastructure under most impersonation campaigns.

DEFINITION

Domains registered to resemble yours. They're the infrastructure most impersonation campaigns are built on, and often the earliest detectable signal.

5 SUBTYPES

Official domain **harborcu.com** Official

Typosquatting A dropped, added, or swapped letter.	harbrcu.com
Homoglyph / character swap Look-alike characters, such as a zero for an "o."	harb0rcu.com
Combosquatting Your name plus a trusted word like login, support, or pay.	harborcu-login.com
TLD swap Your exact name on a different extension.	harborcu.co
Subdomain abuse Your name in front of a domain the attacker owns.	harborcu.secure-verify.net

~20 days
BRAND-IMPERSONATING

~23 days
SQUATTED

Brand-impersonating domains averaged about 20 days, and squatted domains about 23 days, between detection and deregistration.

Source: arXiv 2025, "Registration, Detection, and Deregistration: Analyzing DNS Abuse for Phishing Attacks"

SECTION 01 · ANATOMY OF AN ATTACK

The attacker's playbook

From lookalike to losses in under two hours. A fictional campaign against Harbor Credit Union.

09:02 Registration

The attacker registers `harborcu-verify.net`.

DOMAIN **harborcu-verify.net** CREATED **09:02**
REGISTRANT **privacy-protected** DNS **live in minutes**

09:20 Cloning

AI rebuilds the login page and copy in minutes.



10:05 Delivery

An SMS blast goes out to scraped numbers.

> Your Harbor CU account is locked. Verify now: harborcu-verify.net/login

10:40 Harvest

Credentials and MFA codes are captured. Fraudulent transfers begin before the first customer calls support.

USER **m.ortiz** PASS ••••••••• MFA **482 913** Critical

INTERVENTION POINTS

Where it could have been stopped

09:02

Detect the registration. A new lookalike is the earliest signal you'll get.

09:20

Detect the clone and flag it in browsers, before the first SMS goes out.

After 10:05, every minute adds victims. The cheapest moment to act is before the lure is sent.

ATTACK CLOCK 09:02 REGISTERED ---------- 10:40 HARVEST 98 MIN

SECTION 02

The Phishing Site Vector

Where your customers' credentials and payments are actually taken.

DEFINITION

Live pages that harvest credentials, card data, or MFA tokens under your brand.

5 SUBTYPES

- **Credential-harvest login clones**
- **Real-time MFA-proxy pages**
- **Fake payment and "payment failed" pages**
- **Counterfeit storefronts**
- **Fake support and refund portals**

+10.1%

Phishing attacks rose 10.1% in Q2 2026, and June 2026 (425,808 attacks) was the highest monthly total since April 2023.

Source: APWG Phishing Activity Trends Report, Q2 2026

+40%

Smishing grew 40% from Q1 to Q2 2026.

Source: APWG Phishing Activity Trends Report, Q2 2026

EXECUTIVE ACTION PLAN — 3 CHECKS

- [x] Know your official domains and publish them to customers.
- [x] Monitor for copies of your login and payment pages, not just your domain name.
- [x] Treat any off-domain page collecting credentials in your name as an incident.

SECTION 02 · WHY SPEED DECIDES THE OUTCOME

Most of the damage happens in the first hours.

Phishing-site lifespans: short in the middle, dangerous in the tail.

MEDIAN LIFESPAN
~5.5 hrs

AVERAGE LIFESPAN
~54 hrs

DISTRIBUTION OF PHISHING-SITE LIFESPANS



Time online after detection (compressed scale; curve shape is illustrative, markers from the study below)

Many sites die fast, but the ones that persist do the most damage, and you can't know in advance which kind you're facing.

Source: Lim et al., "7 Days Later: Analyzing Phishing-Site Lifespan After Detected," WWW 2025

ANALYST PRO-TIP

Report the median and the worst case.

Measure your median and your worst-case (90th percentile) time to takedown. Averages hide the long-lived sites that cause the most loss.

SECTION 03

The Impersonation Vector

The lure that moves customers off your official channels.

DEFINITION

Fake profiles, accounts, and pages that pose as your brand or executives on social platforms, marketplaces, and messaging apps.

4 SUBTYPES

- **Fake customer-support accounts**
- **Executive impersonation and deepfakes**
- **Fraudulent ads and promoted posts**
- **Fake apps and messaging-app profiles**

43.8% of social-media threats in Q1 2026 were impersonation, and threat volume rose on every platform.

Source: APWG Phishing Activity Trends Report, Q1 2026

PUBLIC REPLY THREAD

JM **Jordan M.** @jordan_m · 2m
@HarborCU my card got declined twice today. Anyone else having issues?

H **Harbor CU Support**
@HarborCU_Help · 1m
So sorry, Jordan! We can fix this fast. Please DM us your account number and the code we just texted you.

Message @HarborCU_Help

High FAKE SUPPORT PROFILE

Not an official handle. The goal is to move the customer into private messages, away from your moderators.

REAL-WORLD NOTE The WEF Global Cybersecurity Outlook 2026 cites deepfake videos of Indonesia's president promising financial aid, used to defraud people across 20 provinces.

EXECUTIVE ACTION PLAN 3 CHECKS

- ✓ Publish your official handles.
- ✓ Monitor brand-name variants across platforms.
- ✓ Pre-agree an escalation path with legal and comms.

SECTION 04

Detection isn't defense.

The takedown gap: finding a fake is not the same as removing it.

Most organizations can find some fake sites. Few can remove them quickly. Manual takedown means:

Drafting abuse reports

Evidence, screenshots, and a written case for every provider.

Chasing registrars and hosts across jurisdictions

Different forms, policies, languages, and time zones.

Re-checking every few days

Until the site is actually gone, or until it comes back.

~18%

OF PHISHING SITES FLAGGED

In one large study, Google Safe Browsing flagged only about 18% of phishing sites, taking about 4.5 days on average.

Source: Lim et al., "7 Days Later: Analyzing Phishing-Site Lifespan After Detected," WWW 2025

The gap, drawn to scale

PHISHING DOMAINS · AVERAGE



Across phishing domains, the average time from detection to deregistration was **11.5 days**. Every one of those days, the domain can keep serving pages and sending mail.

Source: arXiv 2025, "Registration, Detection, and Deregistration: Analyzing DNS Abuse for Phishing Attacks"

SECTION 04 · CONTINUED

Manual vs automated takedown

The same five steps, on a very different clock.

	MANUAL Weeks → Days	AUTOMATED Hours → Minutes
Discovery	Periodic searches and customer reports	Continuous AI monitoring of domains, web, and social
Verification	Manual review and screenshots	Auto-collected evidence (screenshot, WHOIS, similarity, risk score)
Enforcement	None until removal	Browser safety warnings protect users in minutes
Takedown	Emails to abuse desks, follow-ups	Reports filed to registrars, hosts, and browser programs automatically
Tracking	Check back manually	Monitored until confirmed removed, with a full audit trail

THE AUTOMATED PIPELINE

DETECTION → VERIFIED REMOVAL

00 DETECTED

A new lookalike is found.

⚠ Lookalike alert
harborcu-verify.net
vs harborcu.com
Critical

01 VERIFIED

Evidence attached.

✓ Screenshot
✓ WHOIS
✓ Similarity
✓ Risk score
Takedown →

02 ENFORCED

Flagged by browsers.

⚠
Deceptive site ahead
This site may trick you into sharing passwords or card numbers.
Back to safety

03 REMOVED

Confirmed gone.

harborcu-verify.net
⊘
Site removed
Resolved

SECTION 04 · THE RESOURCE REALITY

Most teams can't staff this 24/7.

The work never stops, and the people to do it are scarce.

54%

cite insufficient knowledge or skills as their top hurdle to adopting AI for cybersecurity

Source: WEF Global Cybersecurity Outlook 2026

85%

of organizations with insufficient cyber resilience say they lack critical people and skills

Source: WEF Global Cybersecurity Outlook 2026

65%

of large companies name third-party and supply chain vulnerabilities as their greatest resilience challenge, up from 54%

Source: WEF Global Cybersecurity Outlook 2026

Brand protection is a continuous, high-volume, cross-jurisdiction task. It's a strong candidate for automation, or for delivery as a managed service by a trusted security partner.

PATH A

Automate in-house

Give your SOC or fraud team tooling that detects, verifies, enforces, and files takedowns without manual abuse-desk chasing.

PATH B

Buy it as a managed service

Have a trusted MSSP or MSP run brand protection around the clock, with reporting you can take to the board.

ANALYST PRO-TIP

Do the math on analyst hours.

If a single analyst spends more than a few hours a week chasing abuse desks, the cost of manual takedown already exceeds the cost of automating it.

SECTION 05

Building a Brand Protection Program

Five pillars, and the five KPIs that prove they work.

01
Monitor
Domains, web, and social

02
Verify
A consistent evidence standard

03
Enforce
Browser flagging

04
Remove
Registrar and host takedown

05
Report
Board and customer reporting

KPIs THAT MATTER

METRIC	WHY IT MATTERS	TARGET TO AIM FOR Sentrike platform data
Time to detect Registration or publication → alert	Every hour undetected is an hour the campaign can be staged and launched.	**4 min** average detection
Time to enforce Alert → browser warning	Browser warnings protect customers while removal is still in progress.	**4 min** average time to enforce
Time to takedown Alert → confirmed removal	The number that ends the harm. Track the median and the 90th percentile, not just the average.	**2.1 hrs** average time to takedown
Completion rate % of cases fully removed	A case that stalls at an unresponsive host is still harming customers.	**97.2%** completion rate
Automation rate % of steps handled without analyst time	Shows how much analyst capacity the program frees for real investigations.	**78.9%** automation rate

Set your own baseline first, then improve against it. Targets are averages, offered as benchmarks of what automation makes possible.

TARGETS: SENTRIKE PLATFORM DATA

SECTION 06

2026 Brand Threat Comparison Matrix

Three vectors, one campaign.

In live attacks these three vectors work together. The lookalike domain hosts the phishing site, and the impersonation profile drives traffic to it.

THE LURE
Impersonation profile → THE INFRASTRUCTURE
Lookalike domain → THE HARVEST
Phishing site

VECTOR	LOOKALIKE DOMAINS **The Infrastructure**	PHISHING SITES **The Harvest**	IMPERSONATION **The Lure**
Core intent	Borrow credibility	Capture credentials, cards, and MFA codes	Build trust and route victims
Methodology	Typo, homoglyph, combo, and TLD swaps	Cloned login, payment, and store pages	Fake support, executive, ad, and app profiles
AI impact	Bulk generation of variants	Pixel-perfect clones in minutes	Deepfakes and localized scripts at scale
Earliest signal	→ New registration	→ Page goes live	→ Profile created or ad launched
Key defense	**Continuous registration monitoring**	**Browser flagging + host/registrar takedown**	**Platform reporting + published official channels**

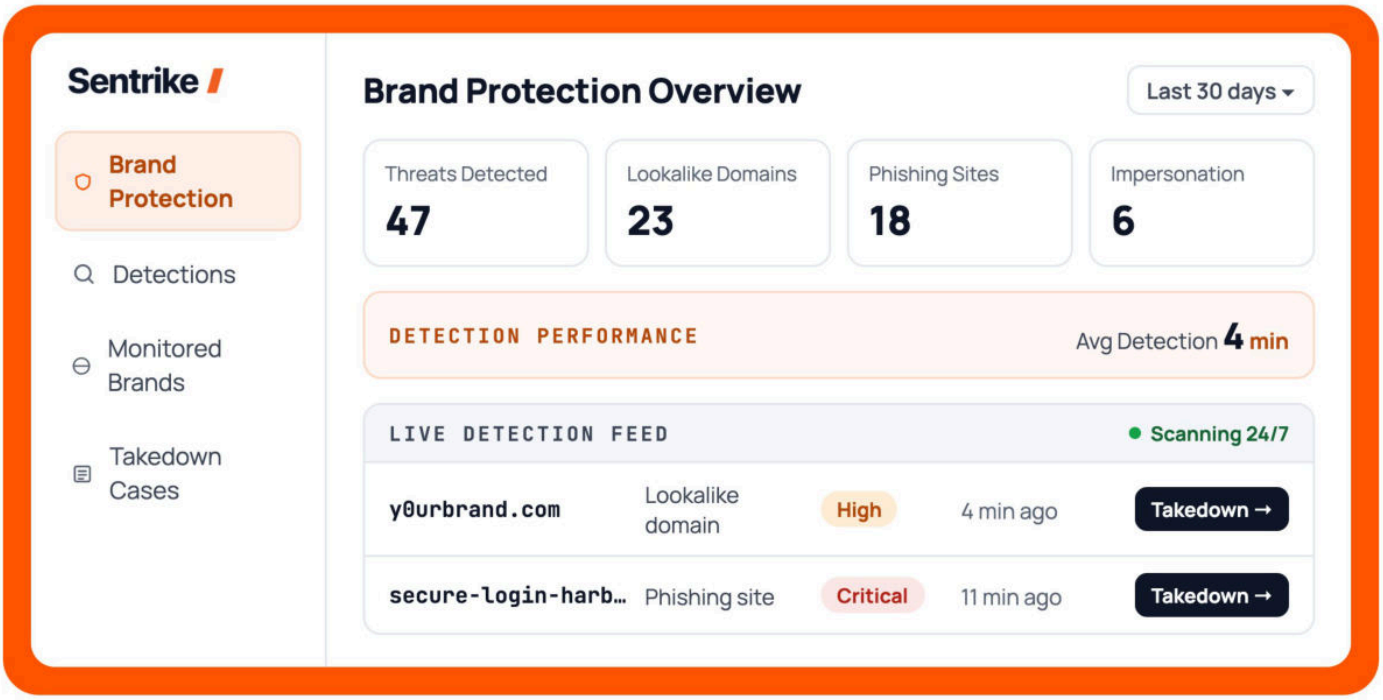
Tip: tear out this page for incident triage. When a report comes in, name the vector first, then work the row for its earliest signal and key defense.

WHERE SENTRIKE FITS

Proactive in minutes, not reactive in weeks.

Closing the takedown gap with autonomous AI agents.

Sentrike's AI agents monitor domains, websites, and social profiles 24/7. They detect impersonators in minutes, auto-collect evidence, trigger browser warnings, and file takedowns with registrars and hosts, tracking every case until it's removed.



Illustrative dashboard with fictional brands. Average detection time: Sentrike platform data.

FOR SECURITY TEAMS

No abuse-desk chasing

Evidence, filings, and follow-ups run automatically, with a full audit trail.

FOR FRAUD TEAMS

Fewer victims, faster

Browser warnings protect customers while the takedown completes.

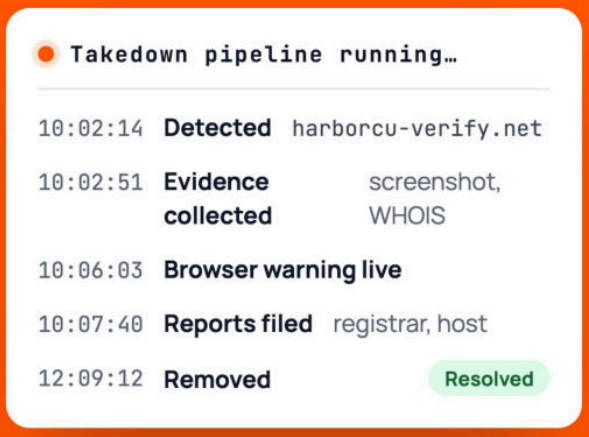
FOR PARTNERS

White-label, multi-tenant, resell-ready

Protect every client brand from one console, under your name.

The path to resilient brand protection.

Impersonation is now cheap, fast, and AI-assisted. The organizations that limit losses are the ones that shrink the time between registration and removal.



NEXT STEP

See your brand's exposure.

We'll show you the lookalikes and impersonators targeting you today.

Request access →

Schedule a call →

sentrike.ai/early-access

Sentrike

AI-powered brand protection and takedowns

SOURCES

- World Economic Forum, Global Cybersecurity Outlook 2026
- APWG Phishing Activity Trends Reports, Q1 and Q2 2026
- Lim et al., "7 Days Later: Analyzing Phishing-Site Lifespan After Detected," WWW 2025
- arXiv 2025, "Registration, Detection, and Deregistration: Analyzing DNS Abuse for Phishing Attacks"
- Sentrike platform data