The New Era of Phishing

Evolution, Trends, and Key Lessons from Over 2 Million Successful Takedowns

AXUR

Executive Summary

Phishing has evolved into one of the most persistent and sophisticated cyber threats today. It's no longer limited to fake emails — phishing now spans multiple channels, including social media, SMS, paid ads, and even deepfakes. Attackers are also getting smarter, avoiding direct mentions of brands in domains or HTML code to bypass traditional detection systems.

Companies that still treat phishing as just an IT issue may be underestimating its true impact. Phishing directly threatens brand reputation and the security of customers and partners, driving up costs related to customer support, reimbursements, and restoring trust.

Key Shifts in Phishing Tactics

→ Omnichannel Attacks:

Phishing may start with an email, but it just as easily begins with an SMS, a social media message, a paid ad, or even a phone call.

→ AI-Powered Phishing:

Threat actors are using advanced AI tools to build highly convincing fake pages that evade traditional detection methods.

→ Mobile-Optimized Scams:

Fake sites are increasingly designed for mobile devices, where URLs are harder to spot, and users act faster.

→ Weaponizing Paid Ads:

Fraudsters are using online ads to target users actively searching for specific products or services. By delivering malicious links at the exact moment someone is looking for a deal, they dramatically increase click-through rates.

Fighting Back: Monitoring & Takedown

Combating phishing requires more than just passive monitoring — rapid detection must be paired with fast action to minimize damage.

Advanced AI, including Vision Language Models (VLMs), can analyze phishing campaigns in real time, detecting threats before they spread.

But detection is just the first step. A critical part of any response strategy is the takedown process, which removes fraudulent websites, fake social profiles, and malicious ads from the web.

Phishing is constantly evolving — and getting harder to detect and stop. Traditional methods alone aren't enough.

The following pages explore how advanced monitoring, AI-driven detection, and automated takedown workflows are changing the game — delivering real protection for brands and consumers alike.

Read the full content →

Phishing: A Direct Threat to Brands

The most common type of phishing attack is the one that impersonates a trusted institution — like a bank, retailer, or government agency — to trick the victim. From this perspective, phishing is inherently tied to the brands and legitimate identities that attackers exploit to make their scams believable.

Phishing is undeniably a major concern for corporate networks and end users alike, who are bombarded with massive volumes of malicious emails. In 2023, Google reported blocking 15 billion spam and phishing messages every single day.

However, phishing today is no longer limited to just one communication channel. Attacks can start with an email, an SMS, a phone call, or even a paid ad on social media. These scams often guide victims from one channel to another, creating a seamless experience that maximizes the chances of stealing sensitive data.

At this point, it's safe to say phishing has evolved into a multichannel, multivector threat — far beyond just email-based attacks.

What hasn't changed, however, is how heavily phishing relies on copying a brand's visual identity — leveraging logos, designs, and messaging to deceive customers.

For companies that haven't yet recognized phishing as a direct brand threat, there's a major opportunity to shift that mindset.

By treating phishing the same way they would brand abuse, counterfeit products, or unauthorized offers, companies can gain a clearer understanding of how their customers and partners are being targeted online.

At the same time, they can adopt proactive measures — like fast, automated takedowns — to minimize phishing's impact and create a safer, cleaner digital experience. This proactive stance protects customers from falling for scams and protects the brand from reputational damage, not to mention the costs tied to customer support, remediation, and trust recovery after an incident.

The first step toward fighting phishing effectively is understanding that it's not just the recipient's problem — it's the brand's problem too.

Phishing Fundamentals

While traditional phishing often brings to mind fake emails impersonating a financial institution, this narrow view no longer reflects today's reality.

Phishing is better understood as a broader category of fraud where criminals misuse a brand's identity and craft a convincing story to lure victims into visiting a fake website or taking some other action that leads to data theft.

From this perspective, phishing shouldn't be seen as something confined to a single channel (like email), nor should it always be assumed that the goal is financial information.

Phishing also isn't limited to consumer-facing attacks. Threat actors frequently impersonate suppliers, partners, and vendors to gather sensitive information or steal corporate credentials — ultimately targeting corporate networks for ransomware attacks, data theft, and other types of large-scale breaches.

With the widespread adoption of cloud office suites and Software-as-a-Service (SaaS) platforms, phishing campaigns now target credentials that provide direct access to critical business assets — often using the same tactics that work on consumers.

Most Targeted Sectors in 2024

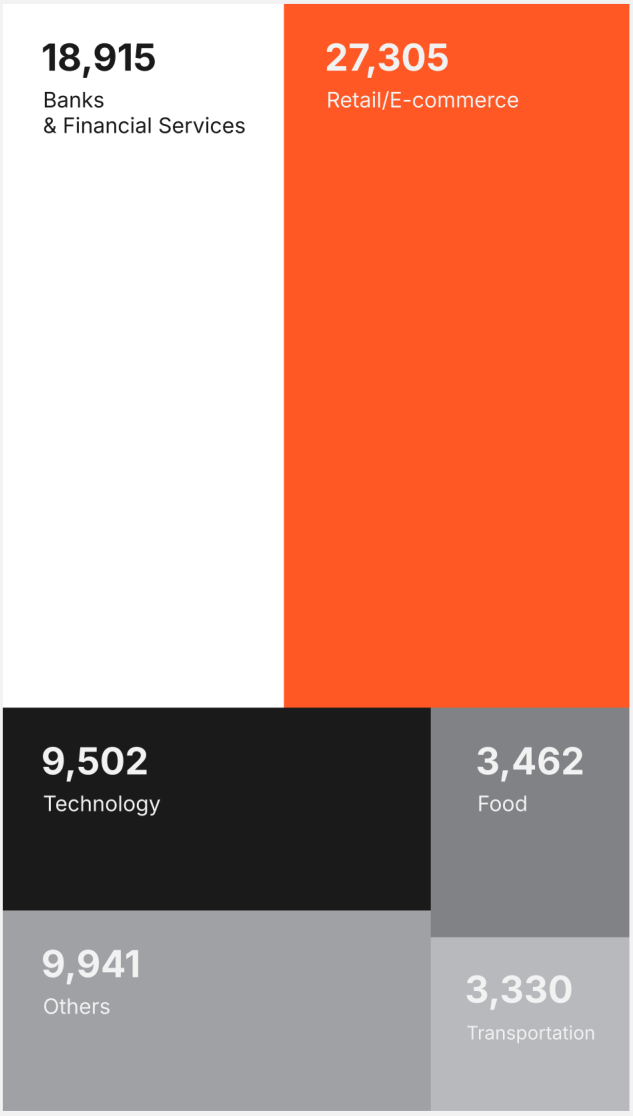

Based on detections from the Axur platform, the industries facing the highest volume of phishing attacks in 2024.

Types of Phishing

Phishing is no longer limited to a single channel or method.

By Channel

→ **Smishing:** SMS

→ **Vishing:** Voice calls

→ **Quishing:** QR codes

→ **Fake Ads/malvertising:** Fraudulent ads

By Technique

→ **Spear phishing:** Crafted messages targeting specific individuals or organizations

→ **Fake News:** Articles designed to drive traffic to phishing sites

→ **Pharming:** Manipulating DNS settings to redirect users to malicious sites

Phishing Losses for Brands

For direct victims of phishing — whether companies or individual consumers — the fraud can cause significant disruption, from recovering stolen funds to canceling credit cards and resetting passwords.

→ Companies whose brands are impersonated in phishing scams also suffer both tangible and intangible losses.

→ Customer dissatisfaction after falling victim to a phishing scam, often resulting in long-term damage to trust.

→ Increased customer support costs from assisting phishing victims who believed they were interacting with the legitimate brand.

→ Lost sales and missed opportunities stemming from customer or partner insecurity about the brand's online safety.

Essential Technologies

→ Mobile Phishing Detection
It's critical to simulate mobile access during detection processes. Many phishing pages are designed to only load on smartphones, blocking access from desktops and laptops. On top of that, phishing attacks targeting social media and mobile users often incorporate multiple filtering layers, specifically designed to block less sophisticated crawlers — making it even harder for bots to detect these threats.

→ AI-Powered Large-Scale Analysis
Millions of pages and artifacts must be analyzed daily to detect phishing campaigns. AI can perform advanced visual inspections, detecting brand logos, layout similarities, and color schemes that mimic legitimate pages. This analysis happens in a fraction of the time it would take a human, without sacrificing detection quality.

Phishing and External Cybersecurity

Companies don't have to rely solely on end-user security solutions to block phishing attacks. With an External Cybersecurity approach to fraud prevention, businesses can actively monitor brand misuse, identify phishing campaigns, and take proactive action to mitigate the impact of malicious campaigns.

Phishing Has Become More Sophisticated — and Harder to Detect

If phishing attacks were still relying on the same tricks from decades ago, phishing wouldn't remain one of the most prevalent cybercrimes in the world today. What we see instead is a constant evolution of tactics. The techniques that prove successful for cybercriminals quickly become trends.

Unfortunately, defensive strategies can't only focus on the latest techniques. If an old trick becomes effective again due to gaps in modern defenses, attackers are more than happy to bring those older tactics back into rotation.

One example of this is phishing emails with malicious content hidden in attachments. Moving suspicious elements into a different file type — like PDF, DOC, or XLS — to evade spam filters was widely used in the early 2010s. After fading for a while, this method resurfaced and became effective again.

In other words, old tricks are often recycled when changes in detection algorithms — usually designed to fight newer threats — inadvertently reopen old vulnerabilities.

That said, there are several newer phishing trends that have emerged and are shaping today's threat landscape: fewer direct brand mentions, mobile-first attack strategies, paid ads as a distribution channel, and the use of intermediary pages to conceal the scam's final objective.

These trends don't exist in isolation. On the contrary, it's common for a single phishing attack to combine one or more of these strategies at the same time.

Fewer Direct Mentions of Brands

Since phishing relies on impersonating well-known companies or entities, most phishing attacks could previously be identified by analyzing newly registered domains and page URLs.

In many cases, this made it possible to anticipate phishing campaigns before they were even launched, as there was often a gap between domain registration and the start of the campaign to promote the fake site.

Domain Percentage

70% of malicious domains now avoid using brand-related keywords.

Mention in HTML Code

18% do not mention the brand anywhere in the HTML code.

To avoid detection through these techniques, criminals are deliberately avoiding the use of brand-related keywords in domain names.

This doesn't mean phishing attacks have stopped exploiting brands to deceive consumers. They've simply changed how they create the connection to the brand, making it less obvious to traditional detection algorithms.

In the simplest cases, the brand name is simply moved to elements like subdomains.

Due to the focus on mobile-targeted attacks, the absence of the brand name in the domain doesn't significantly reduce the effectiveness of the scam. The smaller size of the address bar on smartphones makes it harder for users to see the full URL, giving attackers more flexibility to manipulate the visible address.

To completely remove any textual mention of the brand, criminals often rely on visual references, either by embedding the brand's logo directly or using elements that evoke the brand's visual identity — such as product images or photos of store fronts.

Artificial intelligence is one of the most effective tools for addressing this challenge. Through machine learning, AI can recognize these visual similarities and link phishing websites to specific brands, even when the brand name itself is never mentioned.

Multichannel Phishing with a Mobile Focus

Mobile devices have always been convenient communication tools, but their functional limitations once made it difficult to imagine a fully mobile-based fraud scheme. Smartphones changed that.

On top of that, people now spend significantly more time on their phones. According to Statcounter, 61% of all web traffic today comes from mobile devices.

As a result, many phishing campaigns are now specifically designed with mobile users as the primary target.

Mobile as the Primary Target

→ Promoting scams via SMS or other channels that exist exclusively in the mobile environment

→ Spreading malicious apps

→ Using targeted ads on social media and search engines to ensure phishing campaigns are only shown to mobile users

→ Implementing filters on phishing sites that redirect traffic from non-mobile devices (like laptops) to legitimate pages, making it harder for security teams to analyze the malicious site

Criminals often kick off these multichannel phishing attacks by sending an SMS that promotes a link, a phone number, or a malicious app — creating a seamless cross-channel fraud scheme.

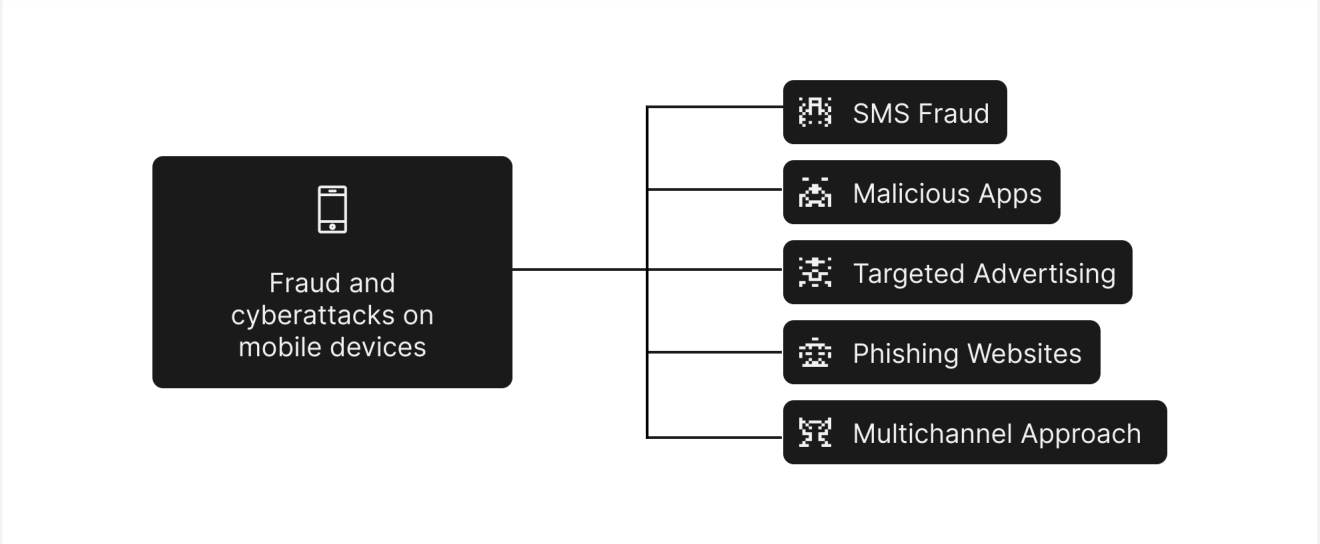

The widespread connectivity of mobile devices creates countless opportunities for fraudsters. The most effective response is an equally comprehensive anti-phishing strategy — one capable of detecting and mitigating all forms of phishing, whether they originate via SMS, email, online ads, or fake apps uploaded to any app store.

How to Ensure Visibility into Phishing Campaigns

Since criminals use filters to limit phishing page exposure to security systems, companies need solutions that can bypass these criminal counterintelligence tactics.

1. Access using proxy systems across multiple regions.

2. Simulate the technical characteristics of mobile devices, including graphics capabilities and touchscreen interactions.

3. Attempt to collect phishing pages from other device categories.

4. Capture a screenshot of the phishing page, exactly as it appeared during a successful access attempt.

Advertising as a Phishing Channel

The practice of pushing malicious content via paid ads isn't new — it's been referred to as malvertising for years, a term originally tied to the injection of malicious code (typically Javascript) into online ads.

These malicious scripts would force ads to load even if the visitor never interacted with them, or they could perform other unwanted actions, such as secretly mining cryptocurrency. The term "malvertising" also reflects the fact that malware could be directly delivered via these ads.

However, what we're seeing today doesn't always fit that definition. Some of the new phishing ads contain no malicious code at all. Instead, they're simply published on social media platforms or as paid search results, both of which severely restrict the technical capabilities and formats allowed for ads.

Because these ads need to convince the user to click voluntarily, the strategy once again closely resembles phishing — relying heavily on unauthorized brand use to appear trustworthy.

The damage caused by unauthorized brand use in ads can be even greater than traditional phishing. Beyond harming consumers, these fake ads can directly compete with legitimate brand campaigns — especially when criminals bid on the same search keywords. Over time, frequent exposure to fraudulent ads can make consumers more hesitant to trust the brand at all.

Not every fraudulent ad is purely a phishing attempt aimed at stealing credentials. Criminals also use ads to promote fake apps or set up payment fraud schemes — offering fake discounts or limited-time deals designed to steal payment information or credit card data directly.

In one way or another, the core logic behind unauthorized brand use in malicious ads is exactly the same as traditional phishing. The goal remains to capture the victim's attention using a trusted brand. Credential theft is also a frequent outcome, whether directly or indirectly through fake apps designed to support the scam.

For these reasons, trying to draw hard distinctions between different types of fraud isn't particularly helpful. In practice, all these threats share common characteristics and can be mitigated using the same core strategies: monitoring, threat intelligence, and takedown.

How Can You Detect Phishing in Ads?

Online ads can be highly targeted, reaching users based on their recent interests, such as products they've searched for. This level of targeting means that brand monitoring must extend to these specific channels, providing broader visibility into the ads being shown to potential victims.

What Changes in Ad Takedowns?

On social media platforms, ads must be linked to a business account, which is essentially a profile within the platform itself. This means there are at least three elements that need to be included in a takedown request:

Active Ad	Linked Profile	Removed Site
The ad that is being promoted on social media	The social media profile associated with the ad	The destination website promoted by the ad

To ensure brand protection, Axur's Takedown notifies all parties involved in the scam's advertising flow.

Fake News

Fake news and disinformation campaigns are a challenge for society as a whole. However, just as the creators of these campaigns exploit the credibility of traditional media outlets, phishing scammers have long exploited the credibility of trusted brands.

From this perspective, it's perhaps not surprising that phishing operators have also realized they can incorporate the concept of fake news into their scams.

The typical modus operandi involves creating or manipulating news stories that mention the targeted brand, then publishing the "article" on a fake news website — usually designed to mimic the colors and layout of well-known news platforms.

This tactic offers three key advantages for attackers:

→ Conceals the scam
When phishing campaigns are promoted through ads, the "news-like" appearance of the intermediary page helps mask the scam and reduces the chances that the ad platform itself will block the content.

→ Expands the narrative
The appearance of legitimate news content allows scammers to build more elaborate narratives — fake promotions, clearance sales from companies in bankruptcy, and auction events are all examples designed to trigger an urgent response from victims. Consumers are less prepared to recognize these narratives compared to older phishing tactics, like fake password reset requests.

→ Boosts credibility
By encountering a forged news story supposedly published on a trusted portal, the phishing victim becomes more likely to trust the next step in the scam's storyline — even if that next step doesn't directly mention the brand itself.

Fake news is also a powerful vehicle for introducing AI-generated deepfakes into phishing campaigns.

To make the scam even more convincing, attackers may include fabricated audio clips of executives or even manipulated videos.

Given the wide range of possibilities this tactic enables, it's crucial for threat intelligence teams to closely monitor these evolving fraud techniques.

How Can Phishing Combine All These Techniques?

1. The attacker creates a fake news page designed to look like a legitimate news website, publishing a story about Apple.

2. The fake news page includes a link to another page or app, which will then attempt to steal data from Apple customers. Since the victim already came from the news site, the phishing page or app no longer needs to mention the brand directly.

3. The attacker designs ads and SMS messages to promote the fake news site. Since the entire scheme is designed to target mobile users, the page can filter out traffic from non-mobile devices, blocking less sophisticated collection systems.

How to Identify the Most Advanced Phishing Scams

Phishing is one of the most common cyber threats, impacting businesses across all industries. For companies that still don't have a monitoring solution in place, the scale of the problem can be surprising. Searching for keywords and brand-specific terms can certainly help detect and mitigate a significant number of fraud attempts.

However, no two companies are the same, and organizations need the flexibility to gather intelligence and leverage monitoring tools to find the best way to customize their processes and policies for dealing with these evolving threats.

How LLMs Are Redefining Threat Detection

Digital threat detection has evolved significantly in recent years, and Large Language Models (LLMs) are playing a key role in this transformation.

Traditionally, the identification of phishing and fraud relied on predefined rules and keyword-based analysis — an approach that worked well in the past but struggles to keep up with the ever-evolving tactics used by cybercriminals.

By applying Vision Language Models (VLMs), companies can analyze not just the text content of malicious sites, but also their visual and semantic structure. These models are capable of detecting advanced impersonation patterns, identifying fraud indicators on pages that deliberately avoid direct brand mentions, and extracting insights that previously required manual review.

By processing massive volumes of data every day, VLMs enhance detection accuracy, significantly reducing false positives while spotting threats that would evade traditional systems. In addition, the use of generative AI enables the creation of detailed descriptions of suspicious pages, providing valuable context and making investigations faster and more efficient.

Application in Phishing and Fraud Detection

Axur has integrated these capabilities into Clair VLM (Cyber Lens for Anomaly and Impersonation Recognition), a proprietary model based on VLMs, trained on more than 15 years of digital threat data. Clair inspects millions of websites daily, analyzing both visual elements and structural data from each page.

Clair detects brand impersonation attempts, flags requests for credentials, payments, or passwords, and generates detailed descriptions of the analyzed pages.

This process is fully automated, ensuring efficient detection while minimizing the need for manual intervention.

Unlike traditional approaches that rely solely on keywords or domain blocklists, the integration of VLMs and threat analysis offers a broader, more contextualized view of attacks — delivering greater accuracy and more comprehensive protection against digital fraud.

Axur's advanced Phishing & Domain Intelligence combines collected intelligence and AI-extracted signals to detect the most complex and innovative fraud schemes.

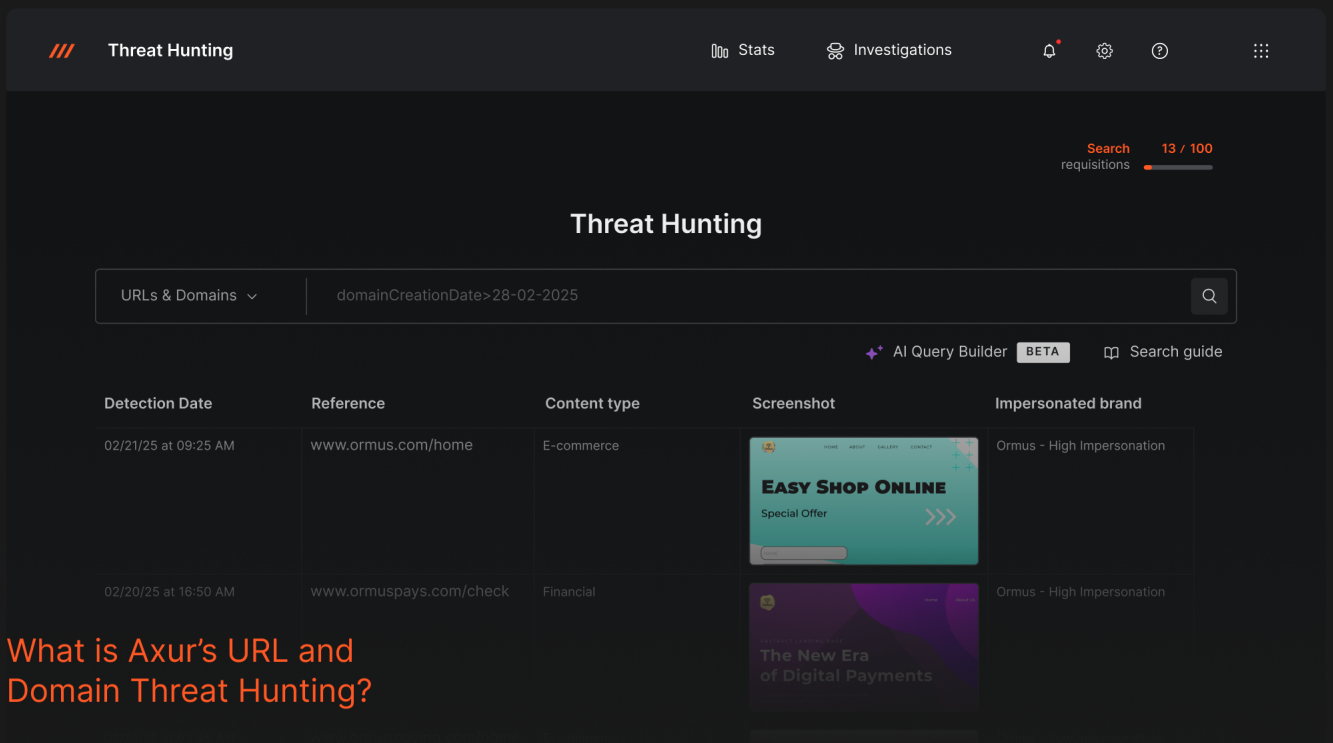

What is Axur's URL and Domain Threat Hunting?

Threat Hunting is a discovery and intelligence tool designed to expand visibility into digital fraud without disrupting existing processes or the automated handling of incidents.

For users, Threat Hunting works like a search engine, but instead of returning general web results, it delivers intelligence on digital fraud and other cyber threat signals.

One of its key categories is URLs and Domains, which is specifically designed for investigating phishing campaigns. It allows analysts to find suspicious pages based on criteria such as URL, detection date, and identified brands.

The results provide the analyst with the website address, a screenshot, the IP address, the page's HTML code, and more than 20 additional collected signals to support deeper investigations.

AI-Powered Threat Intelligence at Scale

One of the key advantages of Axur's Phishing & Domain Detection is its advanced analysis capabilities, powered by the Clair VLM model. Instead of relying solely on objective factors — like domain registration dates or the first time a page was detected — Clair processes and interprets critical attributes to uncover potential fraud

Language AI

Identifies the languages present on the page, helping prioritize investigations based on the scam's target audience.

Content type AI

Recognizes whether the page resembles login screens, payment forms, or e-commerce pages, even when there are no explicit brand mentions.

Credential requested AI

Checks for fields designed to collect emails, usernames, and other sensitive information.

Password requested AI

Identifies the presence of specific fields for login credentials and authentication.

Payment requested AI

Analyzes both text and structural elements to identify pages trying to capture financial data.

By combining these analyses with machine learning, the model enables more precise and context-aware detection, ensuring analysts can uncover fraud that might otherwise slip past conventional methods.

For every signal processed by the Clair model, analysts can see the exact result for each detection criterion, providing full transparency.

The model also identifies which brands appear on the page and measures how closely the page's visual identity matches each identified brand.

By combining these criteria, analysts can pinpoint pages that were recently created, show a high visual similarity to a specified brand, and include one of the key fraud indicators (such as payment, password, or credential requests).

Technical criteria — like HTTP codes, IP addresses, and other infrastructure markers — can also be used to map phishing campaigns that share the same hosting or backend infrastructure.

With this comprehensive view, analysts can deep dive into each phishing campaign, uncovering patterns and connections that go beyond what automated systems detect. Because this is a real-time search, analysts can experiment with different criteria on the spot, adjusting their approach as needed.

Whatever is discovered can inform decisions and help refine the automated detection and fraud-fighting criteria for even greater precision over time.

But identifying phishing is only the first step. Once a threat is detected, it's critical to act fast to prevent victims from falling for the scam and to minimize the impact on the brand. This is where takedown comes in — removing fraudulent pages and malicious profiles before they can do more damage.

How Takedown Fights Phishing and Other Digital Fraud

Takedown refers to the process of removing harmful or fraudulent content from the web. If the content is a social media profile, a takedown results in the suspension or deletion of that profile. If it's a website or domain, the takedown occurs when the hosting provider or domain registrar cancels the service, effectively bringing the page down.

Unauthorized brand use can be classified as an illegal activity. In these cases, the most common type of complaint is an intellectual property violation report, which requests the removal of profiles or pages that misuse the brand's name, logo, or other distinctive elements for personal gain.

The DMCA (Digital Millennium Copyright Act) is a U.S. copyright law that grants safe harbor protections to digital service providers if they promptly remove content upon receiving a valid copyright infringement notice.

In addition, Internet Service Providers (ISPs) have their own terms of service and acceptable use policies, which establish rules and limits on how their platforms can be used.

Since phishing and many other types of digital fraud violate these policies, providers are typically obligated to take action and remove fraudulent content once they've been properly notified.

When a company actively monitors its brand and identifies fraudulent content online, sending takedown notifications becomes a natural and essential next step in its anti-fraud strategy.

Once malicious content is taken down, the threat is neutralized, preventing customers from associating fraudulent narratives with the legitimate brand.

A consistent takedown strategy can also discourage criminals from targeting a specific brand. Fraudsters invest time and resources into their scams — and each time a phishing campaign is swiftly taken down, they lose hours of work. As a result, criminals tend to favor brands that don't actively pursue takedowns, as those fraudulent pages stay online longer.

That said, there are important technical aspects to consider to ensure takedowns are handled correctly, and companies need to understand how to measure the effectiveness of this approach.

Technical Aspects of Takedown

Takedown requests can technically be submitted by anyone and are usually free — often requiring nothing more than an email to the provider. However, this apparent simplicity is misleading, as several factors impact how effective the request will be. If the takedown request is submitted incorrectly, it can lead to delays or even be ignored altogether, even when the reported content clearly violates platform policies.

The first step is to understand how the provider interprets violations. What's obvious to the brand — a fraudulent page or post — isn't always as clear to the provider. The provider will need to classify the incident as a violation of their terms of service in order to justify the removal.

It's essential to gather solid evidence to support the takedown notification, ensuring the provider or platform understands how the content violates their policies.

That evidence also needs to be submitted through the correct channel, using the appropriate language (both in terms of the local language and the technical terminology) and maintaining the proper tone. Automated takedown processes tend to have higher success rates and are typically handled faster by platforms and providers.

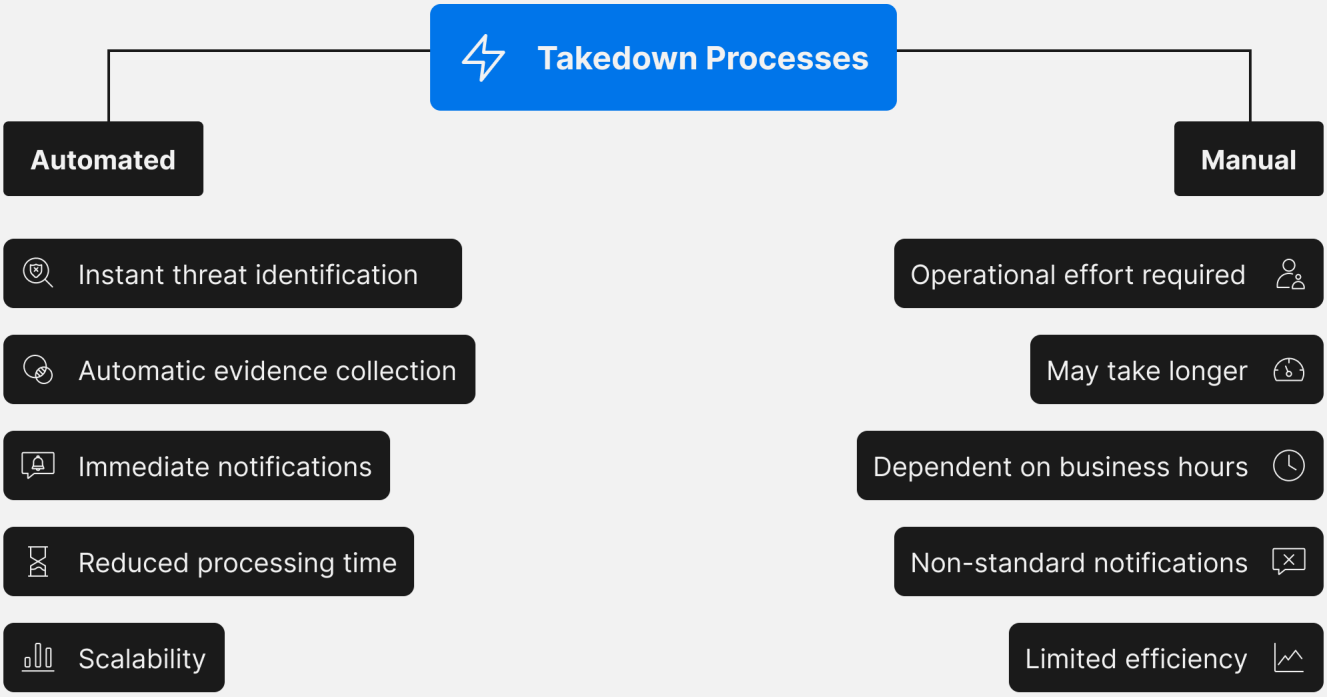

In phishing cases, Axur also reports the fraudulent page to institutions that manage web safety filters (Web Safe Reporting), which causes browsers to warn users about the page — even if the hosting provider hasn't taken it down yet. This additional step helps reduce victim exposure to the malicious content.

How Web Safe Reporting Works

Phishing Identification

Axur detects a phishing attempt and submits a takedown request.

Web Safe Notification

Axur reports the phishing page to web safety institutions.

Alert Generation

These institutions generate alerts for web browsers.

Users Alerted

Users are warned about the phishing threat even before the page is taken down.

In most cases, cybercriminals have to involve at least one well-known and reputable service provider when running a digital fraud scheme. For example, it's extremely difficult to reach a large audience without using a reputable email provider or a popular social media platform. If scammers send phishing emails from an unknown or suspicious provider, the messages are far more likely to be flagged as spam.

That's why fraudsters tend to favor larger, more reputable providers — which, in turn, are the same providers that are typically responsive to properly filed takedown requests.

Axur is recognized as a trusted entity and has priority access to dedicated channels, speeding up the takedown process across multiple providers.

See it in Action: How Automation Reduces Threat Exposure Time

Axur's technology detected a phishing page mimicking a client's website and automatically generated a ticket at 10:47 PM. Thanks to automation:

→ By 10:55 PM, a takedown request was automatically submitted, and the threat was transferred to the incident management workflow.

→ At the same time, 15 partner entities were notified via Web Safe Reporting, limiting exposure and reducing access to the fraudulent page.

→ Also at 10:55 PM, the first notification was sent to the ISP, followed by a reinforcement notification shortly after.

→ By 11:00 PM, Axur received the provider's response, confirming the takedown process had started.

→ At 11:33 PM, the domain was confirmed as offline, completing the takedown process.

With this automated approach, the phishing site was detected, analyzed, and removed in just 46 minutes, significantly reducing the risk of victims accessing the malicious page.

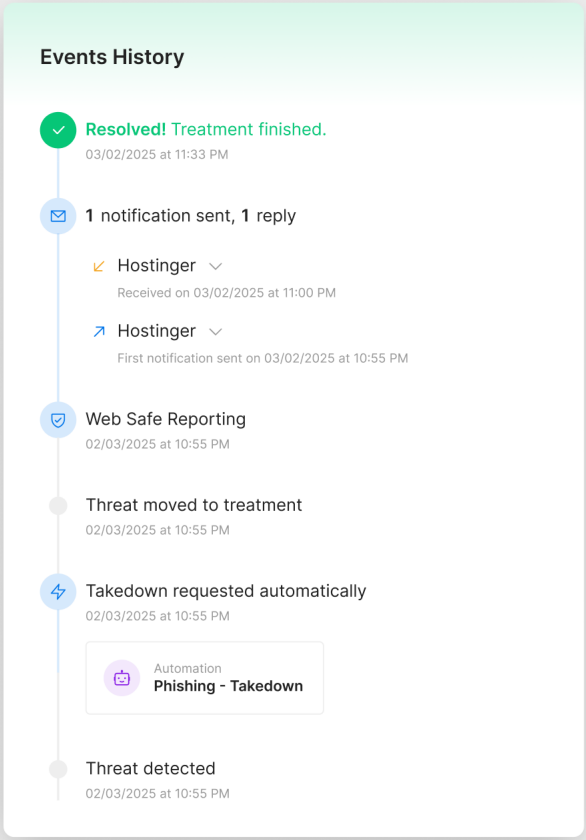

What Can Be Taken Down Through Takedown Requests

- Web hosting services hosting pages that infringe on trademarks (for phishing or SEO manipulation)
- Malware distribution, including trojans disguised as legitimate software
- Creation of fake profiles or accounts impersonating others (including brands) without clearly indicating they are parody accounts
- Mass email campaigns sent without recipients' consent (spam)
- Ads using deceptive images or text, misleading users into believing they come from a trusted brand or individual
- Bot activity or campaigns artificially boosting content visibility
- Storage or distribution of stolen corporate data or private information
- Piracy and Copyright-Protected Data

What Can't Be Taken Down Through Takedown Requests

→ Negative reviews or opinions about the brand

If there are indications of bots or artificial content promotion, those inorganic posts may be reported.

→ False information or accusations that do not violate platform guidelines

Many platforms have strict guidelines regarding fake news, especially for advertising content or promoted posts, which are more regulated than standard user posts.

→ Parked domains or domains with no content

Domains that resemble the brand name but have no hosted content cannot be taken down immediately. However, monitoring these domains so they can be taken down if they become malicious is the recommended approach.

Takedown Success Metrics

One of the key advantages of Takedown as an anti-fraud strategy is that it comes with reliable metrics to demonstrate its effectiveness and track every step of the process.

At Axur, the entire takedown workflow is fully transparent and can be monitored step by step through our platform.

% Success Rate

The success rate reflects the percentage of takedown requests that result in content being removed by the provider.

A successful takedown means the fraudulent content is taken offline, effectively neutralizing the threat.

Axur's current success rate is 98.9%, meaning nearly every request results in successful removal.

First Notification Time

To speed up the takedown process, the first notification to the provider must be sent as quickly as possible.

The time to first notification reflects how fast evidence is collected and the correct channel is triggered to start the removal process.

At Axur, the median time to send the first notification **is just 4 minutes** across nearly all types of incidents.

Uptime

Uptime measures how long reported fraudulent content typically remains online before providers remove it. This timeframe can vary significantly depending on the provider and may also include reinforcement actions, such as resending the notification.

The median uptime for fraudulent content reported **by Axur is 9 hours**, meaning most fraudulent pages stay online for only a few hours before being taken down.

Persistence

In some cases, fraudulent content may be reinstated by attackers who are determined to salvage the time and effort they invested in the campaign. That's why it's critical to keep monitoring the content and resubmit takedown requests if the fraudulent page or profile comes back online — ensuring the takedown has a lasting effect.

Axur guarantees that removed content stays offline for at least 15 days following a successful takedown. Any additional notifications required within that period are handled at no additional cost.

Takedown Viability and Cost-Effectiveness

One of the biggest challenges with takedown is scale. Given the high volume of online fraud, companies often struggle to keep up with notifications, leaving a significant portion of threats unaddressed.

The key to solving the scale challenge is automation. At Axur, 86% of takedowns are fully automated. You can define custom criteria to trigger takedown requests automatically — or manually select incidents and start the process with a single click. After that, the system handles the rest until the fraudulent content is taken down.

Axur's automation framework integrates directly with over 400 ISPs worldwide, including web hosting providers, domain registrars, and social media platforms. This also ensures that multiple notifications are sent when a single fraudulent campaign spans multiple points of contact — such as a fake profile, a malicious web page, and a fraudulent ad.

With AI-powered parameters, Axur can automatically trigger takedown requests for high-risk cases, such as those involving direct brand mentions or specific circumstances (for example, when a fake profile reaches a certain number of followers). This ensures 24/7 protection for your brand.

A key advantage is that takedown metrics are fully transparent, making it easy to calculate the return on investment for this strategy based on fraudulent content successfully removed. Axur charges only for successful takedowns.

This makes takedown one of the most complete and cost-effective anti-fraud strategies, especially when combined with clear performance metrics and automation technology that supports systematic and large-scale enforcement.

5 Myths About Takedown You've Probably Heard

1. We Guarantee Removal Within a Fixed SLA

Reality: No company can guarantee the removal of malicious content within a fixed timeframe, because the actual removal depends on third parties (providers, social networks, hosting companies, registrars). What can be guaranteed is a response SLA — meaning how quickly the company will initiate the takedown process.

2. We Remove 100% of Reported Content

Reality: Not all content can be removed. Some providers ignore requests, claim free speech protections, or require legal action. The real differentiator is the ability to escalate cases, leverage alternative channels (direct contacts, partnerships, legal arguments), and persist until action is taken.

3. We Can Take Down Any Malicious Domain

Reality: Fraudulent domains may be registered with registrars that resist requests or are protected by local laws and policies. In some cases, removal is impossible without formal legal action or lengthy administrative procedures.

4. We Don't Need Additional Evidence to Request Removal

Reality: Most providers require concrete evidence, including screenshots, logs, and direct links. Simply claiming something is phishing without proof often results in delays or outright rejection.

5. We Detect Every Threat Before It Causes Harm

Reality: No company can guarantee 100% preemptive detection. Some threats only become visible after they go live. The real differentiator is the speed of response and the accuracy of analysis once a threat is detected.

How Takedown Protects Your Brand

Takedown is the process of reporting malicious or unauthorized content to cloud providers, hosting companies, and social media platforms. These reports can result in the removal of that content, taking fraudulent pages, fake profiles, or phishing sites offline.

With takedown, businesses, government agencies, and financial institutions can regain control over how their brands are used online.

Unauthorized brand use and illegitimate reproductions — including phishing campaigns — can be effectively removed through takedown, regardless of which channel the fraud is using.

Axur has the world's best Takedown.

Axur's Takedown

98.9% Success Rate

Pages taken down after reporting

86% Automated Takedowns

Compatible with 400 Providers and Platforms

9 hours Median Uptime

Median uptime until takedown

500,000 Takedowns Per Year

Scalable for Any Brand

Axur Guarantee

15 days

Monitoring continues after takedown to ensure the fraud remains neutralized.

4 minutes

Median time to first notification, including phishing cases, reducing the fraud's exposure.

Phishing, Digital Fraud, and Data Leaks

Takedown addresses a wide range of incidents where your brand is linked to malicious activities.

- Fake Websites
- Fraudulent Apps
- Fake Social Media Profiles
- And More...

Track Every Action

Access all artifacts and evidence collected for the takedown process:

- Domain/WHOIS
- Emails
- Screenshots
- Source code

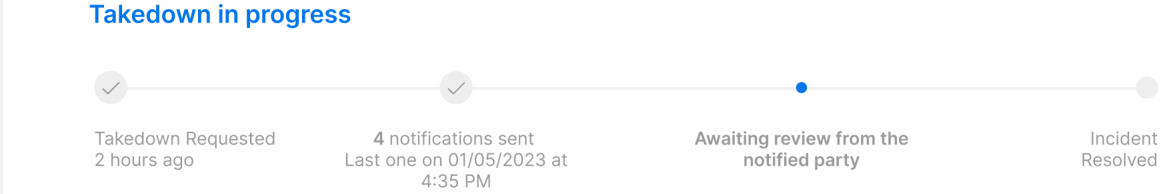

Protect Your Customers

30 Security Entities Notified

In addition to notifying the hosting provider, our Web Safe Reporting mechanism alerts phishing filters to help protect end users.

Experience the World's Best Takedown

REQUEST A DEMO

ISO/IEC 27001 Information Security Management CERTIFIED

CleanDNS Trusted Reporter

Discover all our solutions at **axur.com**

AXUR