From Detection to Disruption:

Why speed is table stakes in modern digital threat response

netcraft

Executive Summary

The attacks have gotten smarter. The infrastructure is industrialized. And the response models most organizations are running? They were built for a different problem.

Digital abuse has moved well beyond the stereotype “hoodie hacker,” rogue domains, and obvious phishing emails. Today's attackers run coordinated, multi-channel campaigns, combining cloned websites, fake social accounts, malicious apps, URL shorteners, deep links, and paid search ads into a single fraud pipeline. By the time one piece gets flagged, the rest of the campaign is already running and AI makes it happen faster than ever, eliminating the skill floor for threat actors.

This is no longer a domain management problem, or purely a security problem, or just a brand issue. It's a business risk — measured in customer harm, fraud losses, regulatory exposure, and trust that's difficult, time-consuming, and costly to rebuild.

The organizations pulling ahead aren't the ones detecting more threats; they're the ones reducing customer exposure the fastest. They've built connected operating models that bring legal, brand, security, fraud, IT, and digital teams around a shared view of risk, a common playbook and metrics that actually answer the right question: how quickly did we stop this from reaching customers?

That's the standard to which this paper is written.

Com Laude and Netcraft recently addressed this challenge from complementary angles at INTA, a conference focused on protecting against trademark infringement and brand impersonation. Com Laude handles domain portfolio management, brand protection, registrar expertise, and asset recovery. Netcraft provides internet-scale threat detection, evidence-led disruption, automated classification, and operational takedown across threats such as phishing, scams, impersonation, fake apps, malicious URLs, and 100+ more attack types.

Together, they move organizations from fragmented, reactive response to coordinated digital threat disruption.

The Problem Has Changed

Here's what a typical social engineering attack looked like five years ago: a typosquatted domain, a copied login page, and a poorly worded phishing email. Detectable. Containable. Annoying, but relatively simple in structure.

Here's what it looks like now: a customer sees a fake brand account on social media. They follow a shortened link. They land on a cloned site. They enter credentials through a mobile deep link.

They are redirected through a sophisticated chain of disposable infrastructure. All hidden behind geotargeting and criminal obfuscation techniques, then the whole campaign disappears before anyone's internal alert fires.

The attack surface has expanded — domains, subdomains, cloned websites, social media, messaging apps, fake customer support pages, malicious mobile apps, paid search ads, certificates, DNS records, email infrastructure, and other signals that can reveal suspicious or malicious infrastructure. Attackers don't limit themselves to the channels your monitoring covers. They use whichever path gets them to customers fastest.

There's a structural imbalance that's worth naming directly. Attackers specialize in one thing: getting a user to click, disclose, pay, download, or authenticate. Defenders have to protect the brand, manage digital assets, preserve customer trust, support legal enforcement, meet security requirements, and coordinate across multiple teams. That's not a fair fight by default.

AI and crime-as-a-service have made it worse. Convincing content is now generated in minutes, not days. Campaigns are localized, versioned, and tested at scale. Brand experiences are cloned with minimal technical skill. The result is a threat environment defined by speed, volume, and iteration.

In that environment, the leadership question shifts. It's no longer, "Can we find the abuse?"

It's, "Can we detect, classify, disrupt, and remove abuse before it causes meaningful harm to our customers?" This is a daunting task if you try to tackle it alone or with legacy technology.

Why Legacy Response Breaks Down

Most organizations are running response models designed for a slower era — manual review and submission, case-by-case escalation, cease-and-desist workflows, domain recovery proceedings. These tools still matter. For permanent recovery of a strategically important domain, for long-running infringement, for enforcement that needs a documented legal path — formal process is often exactly right.

The reality is formal process has an operational tempo that doesn't match how fast modern abuse moves. And for most active threats, the clock is already running.

But what about DMCA?

Where DMCA fits and where it falls short

DMCA notices — copyright-based takedown requests under the Digital Millennium Copyright Act — still have a role. Copyright-based takedown requests can help suppress cloned content, remove infringing search results, or create a documented legal path when you need one. But they were built for a different kind of problem than fast-moving phishing and fraud.

When a cloned site is actively stealing credentials or capturing payment details, copied content is secondary. The harm isn't the infringement — it's the infrastructure running behind it, harming consumers and impacting your brand. A URL scrubbed from search results feels like progress, but the site, the hosting, the redirects, and the downstream attack chain can all stay live. Victims can still reach it via various phishing lures. The campaign can still succeed.

This is where a lot of organizations get stuck. They reach for the tool they know — DMCA, trademark complaints, registrar reports, abuse forms — when what the situation actually needs is faster operational disruption. Those tools aren't wrong; they're just not the best path to fast remediation and meaningful outcomes. Legal process takes time, and time is exactly what attackers are counting on. It's also expensive — attorney hours spent on a takedown that could have been operationally disrupted in hours rarely pencil out.

Modern response models distinguish between legal removal, search suppression, immediate disruption, and full infrastructure takedown. DMCA earns its place when infringement or search visibility is genuinely the issue. For phishing, impersonation, and active fraud, the priority is reducing victim access fast, dismantling the infrastructure behind the attack, and escalating legally only when the risk or asset value makes that the right move.

A phishing site may monetize and disappear in hours, typically less than 24 hours, according to **Netcraft research**. A fake login page may capture credentials before the first internal alert is triggered. A social media impersonation campaign may run for days across platforms while teams figure out who owns the response. The eventual takedown may arrive in time to remove the remnants of something that has already achieved the threat actor's desired outcome.

Legacy models tend to fail in five specific ways:

1. ### Monitoring is fragmented.

   Domain teams see suspicious registrations. Security teams see phishing indicators. Fraud teams see customer complaints. Brand teams see social media impersonation. Rarely do they see the whole attack journey.

2. ### Triage is too manual.

   As volume scales, human-only review creates backlogs. Analysts spend hours separating signal from noise while active campaigns keep running.

3. ### Legal workflows are misapplied to operational threats.

   A cease-and-desist letter is not a fast-response tool. For disposable criminal infrastructure, it's often not the right tool at all.

4. ### Escalation paths are undefined.

   When responsibility sits across legal, security, fraud, IT, digital, and communications — with no single owner — response slows while teams determine accountability.

5. ### Metrics measure activity, not outcomes.

   Counts of domains flagged, notices sent, and cases opened don't tell leadership whether customer exposure actually decreased, or how fast.

Attackers know these gaps exist, and they exploit them. A modern response model is built to close them.

What Next-gen Impact Looks Like

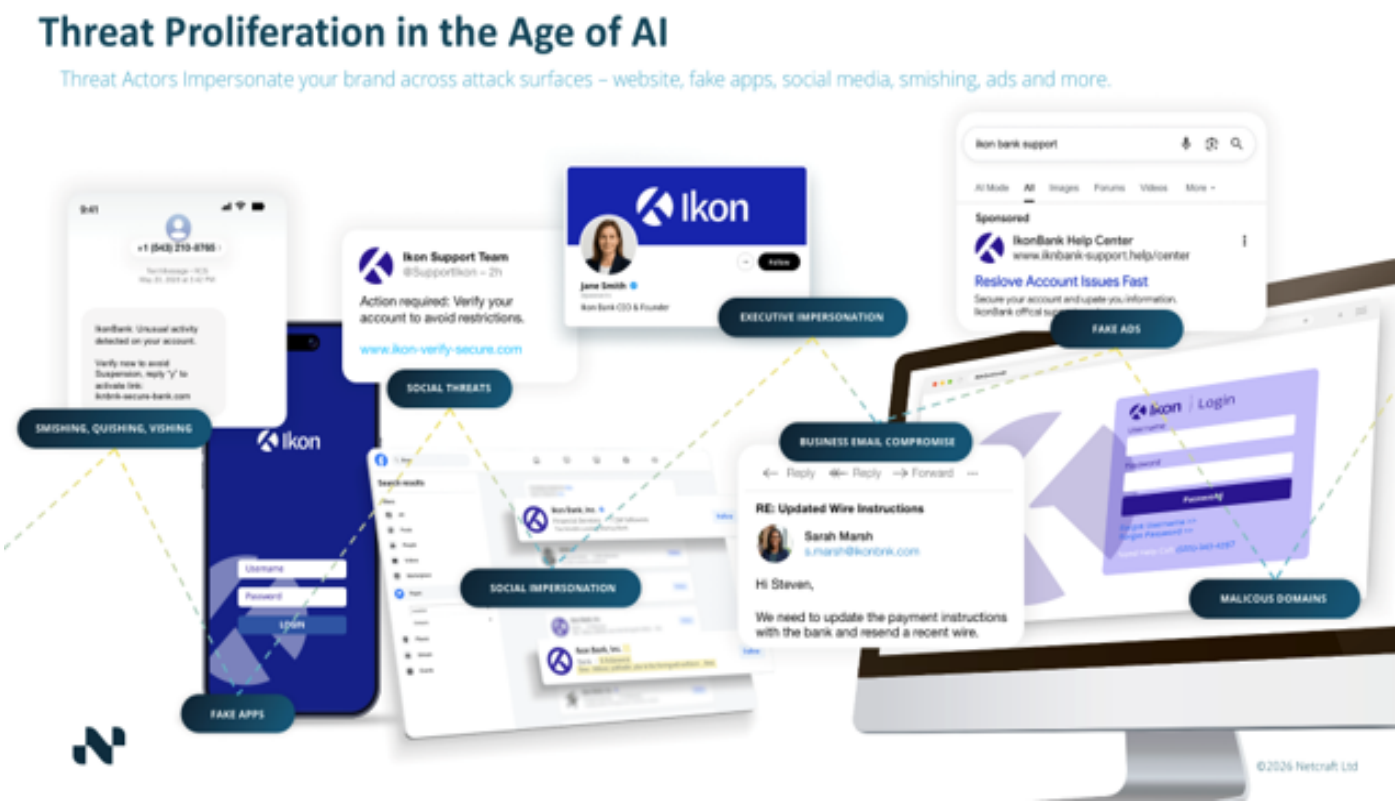

A mature digital threat response program has one objective: reduce the window between threat launch and effective disruption.

That requires five connected capabilities.

1 Detect broadly

Monitoring that only covers registered domains will miss most of the attack surface. Effective detection now requires visibility across domains, subdomains, websites, email infrastructure, certificates, social media, messaging channels, app stores, search results, URL shorteners, and other entry points attackers actually use.

The goal isn't to collect every signal — it's to see enough of the attack surface to identify campaigns early and understand how customers are being targeted before the volume spikes.

2 Classify accurately

Detection without classification is noise. Modern programs need the ability to distinguish between benign lookalikes, low-risk infringement, suspicious infrastructure, active phishing, malware, credential theft, payment fraud, fake apps, executive impersonation, and coordinated scam campaigns.

Accurate classification lets teams prioritize what matters most — and reserve human review for edge cases.

3 Disrupt immediately

Full takedown matters, but it isn't the only lever. Early disruption — browser-based blocking, registrar action, hosting provider notification, platform reporting, app store removal, malicious URL blocking — can reduce victim access while permanent removal is underway.

4 Take down quickly and persistently

Threats don't always stay down. Attackers move infrastructure, relaunch content, rotate domains, and reuse kits. Effective takedown requires evidence, provider relationships, escalation paths, and persistence — not a single notice and a checkbox.

Like with disruption, this is a critical area where speed has meaningful impact. The earlier a threat is disrupted, the fewer businesses and individuals are victimized. And with fewer victims, there is far less ROI for the threat actor.

A mature program tracks the full lifecycle: when a threat was first seen, how it was classified, what action was taken, whether it was removed, whether it resurfaced, and whether related infrastructure is part of a broader campaign.

5 Escalate legally when the risk demands it

Legal remedies belong in the mix — but they should be used selectively. For phishing and social engineering infrastructure, rapid disruption is the priority — full stop. Legal action belongs in the mix for high-value domains, persistent abuse, or assets worth recovering permanently. Not as the first move. As the right move, when the situation calls for it. The strongest programs don't choose between technical disruption and legal enforcement. They use the right tool for the risk in front of them.

How to Operationalize It

Technology alone won't solve this, but it is a critical element of a successful program built not only to stop threats today, but also to mitigate future attacks and support **long-term threat suppression**. Effective use of technology is amplified with internal governance, playbooks, clear ownership, and metrics that get to and report on ground truth.

Establish cross-functional ownership

Domain abuse and digital threats need a clearly accountable owner — one person or team responsible for outcomes — supported by a cross-functional working group. That group should include representation from legal, brand protection, information security, fraud, IT, digital, and communications.

The goal isn't another committee. It's making sure every relevant team shares the same view of the threat, agrees on response thresholds, and knows exactly who does what when an incident occurs.

Define response playbooks

Document playbooks for the most common scenarios: phishing sites, fake login pages, social media impersonation, fraudulent payment pages, fake apps, executive impersonation, high-value domain infringement.

Each playbook should define severity criteria, evidence requirements, escalation paths, communication responsibilities, and target response times. It should also distinguish between immediate disruption and full takedown, or legal escalation.

Measure outcomes, not activity

Executives need metrics that show whether risk is actually decreasing. The right measures include:

- Time from detection to final takedown
- Long-term threat suppression success from threat actors
- Customer exposure indicators (account takeovers, stolen credentials,
- Provider responsiveness and takedown success rates

These metrics tell leadership whether investment is improving resilience — not just increasing the number of cases processed.

Prepare for preemptive action

A growing opportunity is identifying suspicious infrastructure before malicious content even goes live. Criminal signals — for example MX records for BEC compromise — can surface likely abuse earlier.

Where registrar, hosting, and platform relationships support it, preemptive action can stop attacks before they ever reach customers.

Executive Checklist:
What to assess in the next 30—90 days

Leaders don't need to solve everything at once. Start by assessing whether the current operating model can keep pace with the speed and scale of modern abuse.

- [ ] **Reframe the risk.**

  Treat domain abuse and digital threats as material enterprise risks — not narrow legal, brand, or IT issues. Is it represented in the risk register? Is impact measured in terms of customer harm, fraud exposure, regulatory implications, and reputational damage?

- [ ] **Map the customer attack journey.**

  Get an executive-level view of how customers are being targeted across domains, subdomains, social media, messaging, search, apps, and cloned web experiences. Identify the most common paths from first contact to fraud.

- [ ] **Identify monitoring gaps.**

  Does current monitoring cover only domains, or the broader set of channels attackers use? Pay particular attention to subdomains, URL shorteners, app stores, messaging platforms, social media impersonation, and search-driven scams.

- [ ] **Test response speed.**

  Pick a recent incident and trace the timeline. When was it first visible? When was it detected? When was it classified? When was the first disruption action taken? When was it fully removed? Where did time get lost?

- [ ] **Clarify ownership.**

  Who is accountable for digital threat and abuse response? If ownership is split across legal, security, fraud, brand, and IT with no single owner, define one — and build a shared reporting model around them.

- [ ] **Review playbooks.**

  Do playbooks exist for the most common and highest-impact scenarios? Each should define roles, severity thresholds, escalation routes, evidence requirements, and response targets.

- [ ] **Align investment to outcomes.**

  Evaluate tools and partners based on measurable outcomes: faster detection, faster disruption, higher takedown success rates, lower recurrence, reduced customer exposure, better executive visibility.

Conclusion

Traditional brand protection hasn't failed. It's become one part of a larger, faster-moving problem.

As attackers use automation, AI-generated content, disposable infrastructure, and multi-channel campaigns, fragmented monitoring and slow case-by-case response will keep falling behind. The organizations that close that gap are the ones that detect earlier, classify faster, disrupt more quickly, remove threats persistently, and escalate legally when the risk demands it.

They connect legal, security, fraud, brand, IT, and digital teams around a common operating model. They measure outcomes rather than activity. And they work with partners who combine domain expertise, threat intelligence, automation, evidence, and operational disruption at scale.

For executives, the next step is practical: figure out where current monitoring stops, how quickly live threats can be disrupted, and which parts of the response still depend on slow or fragmented workflows.

From there, the path forward is clear: operational disruption first, legal escalation where it's earned, and a response model fast enough to protect customers before abuse becomes impact.

About Netcraft

Netcraft is the global leader in cybercrime detection and disruption, helping organizations detect, block, and take down digital threats before they cause harm. Combining decades of experience with extensive automation, AI and machine learning, and one of the world's most extensive threat detection data sets, Netcraft protects organizations from phishing, fraud, brand impersonation, scams, fake apps, malicious domains, and other forms of online abuse. Headquartered in Salt Lake City and London, Netcraft is trusted by leading global companies, financial institutions, and governments to disrupt cybercrime at scale and create a safer online experience for everyone. Learn more at **www.netcraft.com**.