SACR

JULY 2026

THE FUTURE OF DIGITAL RISK PROTECTION (DRP) IS EXTERNAL TRUST OPERATIONS

SEAN SOSNOWSKI

outtake

SACR®

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Us

Software Analyst Cyber Research (SACR) is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000 readers and followers**, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

Acknowledgements

Practitioners and CISOs

We're excited to share our new report on the evolution of Digital Risk Protection (DRP). The report introduces External Trust Operations: an emerging operating model that moves DRP beyond isolated alerts and takedowns toward coordinated, evidence-backed campaign closure.

The Author

Sean Sosnowski serves as the Research Director for Security Operations and Cloud Security at SACR, where he leads research on SOC strategy and operations, detection engineering, and the evolving role of automation and agentic AI in security workflows. Drawing on a decade of intelligence experience in the U.S. Marine Corps he has authored several analytic reports on emerging technologies and threats regarding sensitive national security and military operations.

Table of Contents

Key Actionable Summary

SACR is introducing its first report on Digital Risk Protection. The wording around External Trust Operations as the term for this emerging operating model marks the point where Digital Risk Protection moves from artifact-level detection to campaign-level accountability.

Digital Risk Protection (DRP) programs were built to find visible abuse of known brands, domains, executives, products, and customer-facing assets. That mission remains necessary, but the external threat has outgrown an alert-and-takedown operating model. One campaign can now move across social accounts, cloned sites, fraudulent ads, messaging channels, synthetic personas, mobile apps, and payment or credential-harvesting flows. Each artifact may look like a separate case even when the same adversarial network connects them. AI increases the pace by making convincing content, localization, variation, and relaunches cheaper and easier to produce.

The category is resetting around a broader operating requirement: External Trust Operations, where the protected object expands beyond a domain, account, or listing to the business entity and the public trust relationships being exploited. The campaign becomes the unit of analysis, and the work from discovery and assessment through response, verification, and recurrence monitoring becomes one connected workflow measured by verified campaign closure.

Research Basis

This analysis combines current public threat reporting and vendor positioning with SACR's ongoing research across digital risk protection, brand protection, threat intelligence, external attack surface management, fraud, narrative intelligence, and executive protection. Public sources establish the threat environment, SACR's analysis defines the category and buyer implications, and Outtake examples show how one vendor is approaching the operating problem.

The accompanying market map is a convergence map built from vendors' primary market motions and functional relevance to External Trust Operations. Placement identifies a contribution to the workflow without implying validation, parity, endorsement, or leadership.

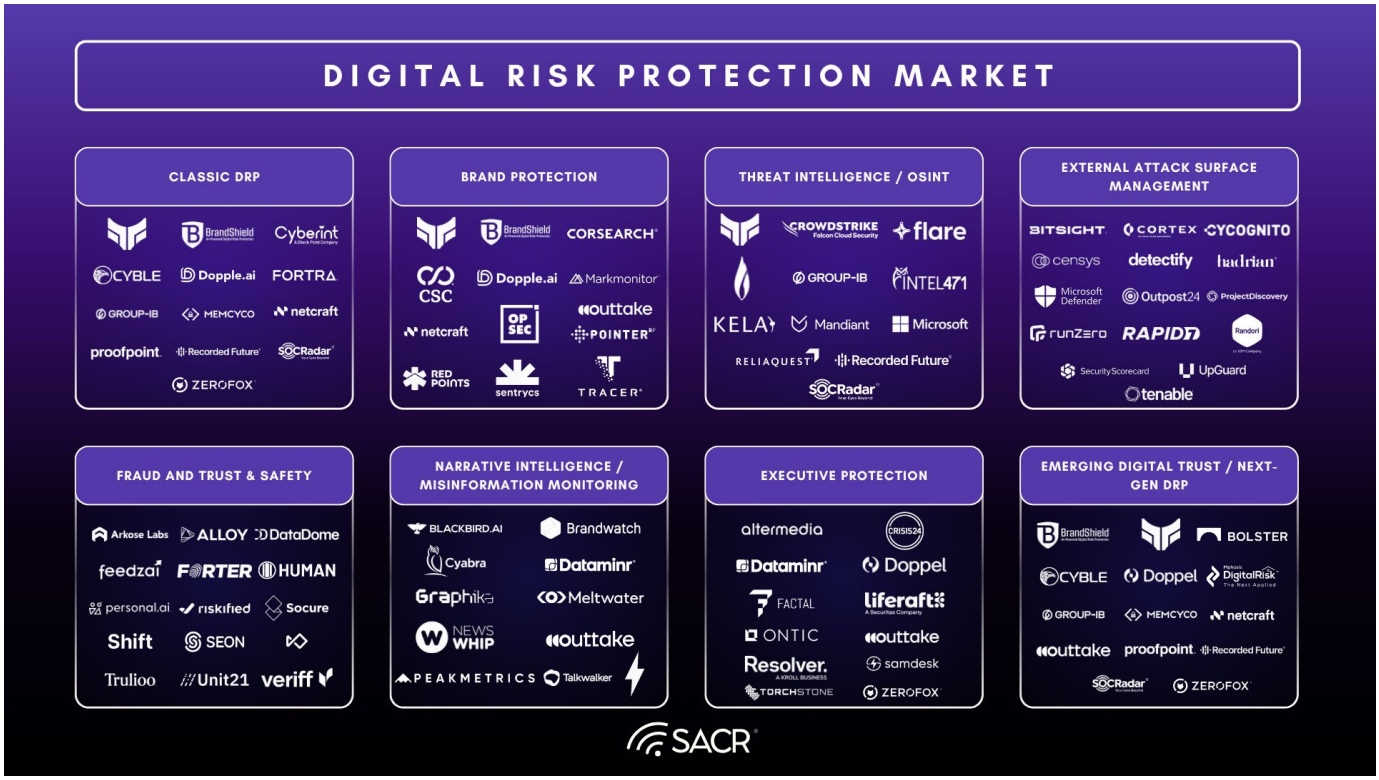

Classic DRP capabilities still matter, but the buyer's question is whether the program can progress from isolated artifact detection to coordinated, evidence-backed closure. The lasting requirement is an operating model that connects external abuse to business context, brings the right internal owners into a shared response, and proves that the active threat pathway was closed.

CISOs should evaluate whether a platform and program can:

1. Model the business entities and trust relationships that require protection, including brands, executives, employees, products, locations, subsidiaries, partners, domains, apps, official accounts, and customer-facing services.
2. Discover abuse beyond static watchlists, explain why a signal matters, and connect related accounts, content, infrastructure, channels, and monetization paths into a campaign.
3. Route evidence and response across security, fraud, legal, communications, brand, customer support, and executive protection without forcing each team to reconstruct the case.
4. Measure verified campaign closure, recurrence, unresolved infrastructure, analyst effort, and business harm. A completed takedown request is one response action, not proof that the threat has ended.

Outtake sponsored this research and appears later as an illustrative product profile. SACR's category definition and buyer framework apply across the market. Buyers should still validate production performance, integration depth, remediation success, recurrence reduction, and business impact rather than treating product direction as proof of outcomes.

Why DRP Is Resetting

DRP became valuable because organizations lost control over how their brands, domains, executives, products, and customer-facing services appeared across the public internet. Attackers could register confusing domains, publish phishing pages, impersonate executives, clone applications, and misuse marketplaces faster than most teams could monitor manually. By bringing those assets into a repeatable monitoring, investigation, escalation, and takedown workflow, DRP gave enterprises a practical response to a rapidly expanding problem.

That mandate remains necessary, but a visible artifact rarely represents the full case. A fake executive account may lead to a cloned support page, moving victims into a messaging channel where a synthetic persona continues the conversation. It may also share infrastructure with fraudulent advertisements, lookalike domains, credential-harvesting pages, or payment accounts. What appears to be a collection of unrelated alerts can often be one coordinated operation, allowing a watchlist to find the entry point without revealing the full path to harm.

The Public Enterprise Is an Attack Surface

Security architecture usually begins with the internal enterprise: identities, endpoints, cloud services, SaaS applications, networks, and data. The organization's public identity is different because it is distributed across platforms and intermediaries the enterprise does not control, including registrars, app stores, advertising networks, marketplaces, search engines, messaging services, review sites, and social networks.

The organization knows which identities, accounts, domains, applications, and support channels are legitimate, but customers, employees, investors, partners, and automated systems must often infer legitimacy from what they encounter. That asymmetry creates an external trust gap that attackers exploit with impersonation, synthetic media, manipulated narratives, and fraudulent infrastructure.

Enterprise AI agents now face the same problem as they retrieve web content, process email and documents, summarize public sources, and make

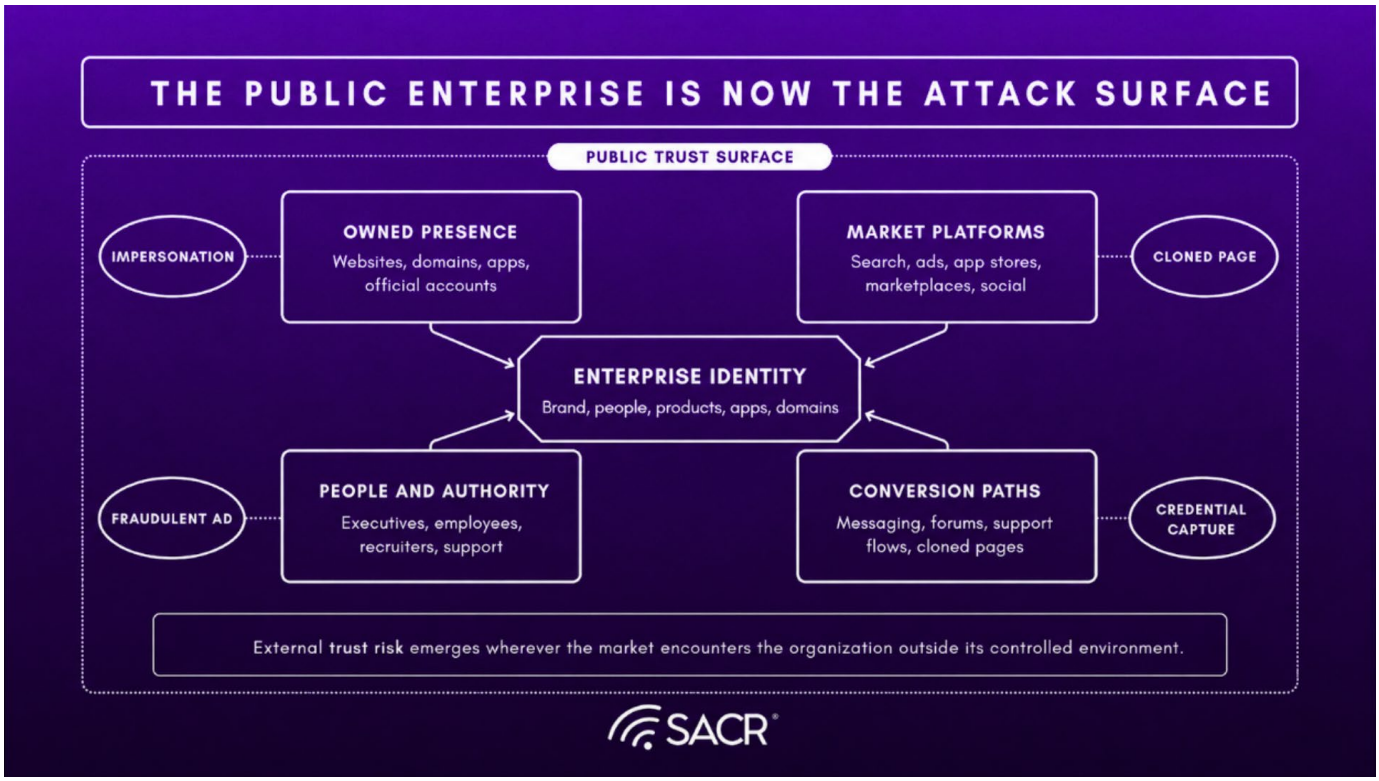

recommendations from external information. Manipulated pages, fake entities, poisoned documents, and adversarial instructions can enter internal workflows through systems that trust public content. This does not turn DRP into a complete AI security category, but it raises the importance of external legitimacy as an input to automated decisions.

AI Changes the Economics and Tempo

AI changes the economics of deception by letting adversaries produce more convincing material, vary it across channels, and localize it for specific targets at far lower cost. Phishing copy, profile images, cloned voices, translated scripts, altered imagery, and fraudulent advertisements can be tested and revised quickly, increasing the pressure on manual review.

Current threat reporting reinforces this shift. The UK's National Cyber Security Centre assessment through 2027 concludes that AI will continue to improve reconnaissance, vulnerability research, social engineering, malware development, and the processing of stolen data. Verizon's 2026 Data Breach Investigations Report found that generative AI bolstered 15 percent of observed attack techniques and reported higher click rates for mobile phishing.

Together, these findings point to an acceleration of familiar attack methods rather than a separate class of external threat. Abuse will arrive in more forms and move across channels more quickly, leaving programs built around static lists and isolated alerts to validate fragments while the campaign continues elsewhere.

Crossing Ownership Boundaries

External trust threats cross organizational boundaries because each team encounters a different expression of the same campaign. Security may discover a fake domain while fraud sees customer losses; brand and legal teams may focus on an impersonating account and its takedown; communications, customer support, and executive protection may see reputational damage, victim reports, or a synthetic persona. Each signal creates a legitimate local task, but the campaign becomes visible only when those perspectives are connected.

This fragmented ownership allows for campaigns to persist through multiple remediation actions. Removing a malicious domain does not automatically trigger a takedown of the social account directing traffic to it. A shared case model can connect the entity under attack, related artifacts and infrastructure, required evidence, actions already taken, and the conditions that must be met before closure to the fragmented teams that are responsible for each individual step of the process.

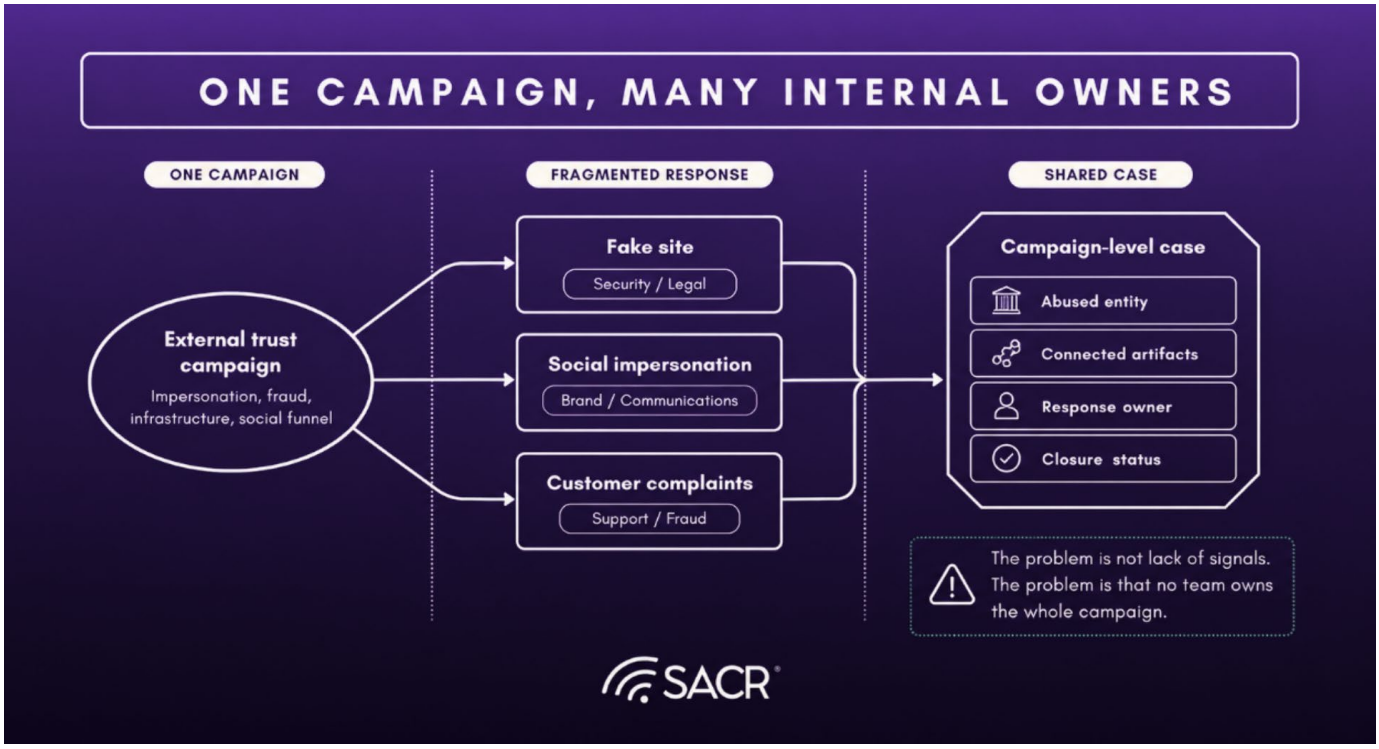

Defining External Trust Operations

External Trust Operations is the operating model emerging within the DRP category reset. It connects business entities, campaign analysis, coordinated response, and verified closure so that external abuse can be managed as an operating problem rather than a collection of unrelated alerts.

The model gives DRP a clearer purpose: protect the public trust surface surrounding the organization's brands, people, products, locations, subsidiaries, partners, customer-facing services, domains, apps, accounts, and narratives, while understanding how adversaries exploit the relationships between those entities and the people or systems that rely on them.

Protected Entities Give Context

Traditional DRP programs commonly begin with known assets such as brands, domains, executives, keywords, trademarks, and logos. A protected-entity model adds the operational context required to determine what is legitimate, why a signal matters, and how apparently different artifacts may relate to the same target.

Products combine official names and imagery with SKUs, authorized sellers, marketplaces, regional availability, and known counterfeit patterns.

Executives combine aliases, verified accounts, public appearances, imagery, organizational relationships, and common impersonation scenarios.

Locations combine signage, geocoordinates, local listings, regional languages, and customer communications.

This context helps a platform judge relevance and connect fragments that use different text, imagery, or infrastructure. A suspicious profile, cloned login page, advertisement, messaging channel, and payment destination may first become a coherent campaign through their relationship to the same executive, product, customer workflow, or brand.

Campaign as the Unit of Analysis

Artifact-level response remains necessary because teams still need to report accounts, remove pages, block domains, preserve screenshots, and escalate listings. The campaign view provides the context needed to determine whether those individual actions address the full threat or only one visible part of it.

By connecting domains, accounts, personas, and other artifacts, campaign analysis reveals behavior that individual nodes cannot show on their own. A low-reach account may justify monitoring in isolation, yet the same account can require an urgent cross-functional response when it is tied to a cloned support page, active advertising, victim reports, and a payment channel.

The Operating Sequence

The operating sequence connects six activities that are often split across tools and teams. Each stage should preserve enough context and evidence for the next stage to act without reconstructing the case.

INTERPOL's Spotlight Cybercrime analysis illustrates the connection stage in practice. Its Project Rapid example shows investigators starting with a single suspicious URL and identifying related phishing sites through lookalike visuals, shared hosting, and similar creation dates. The example shows why discovery becomes more valuable when it expands an artifact into an evidence-backed campaign view.

Verified Campaign Closure

Takedown is one of several remediation actions, but isn't a good measure of a campaign being closed effectively. A case reaches verified closure when the organization has enough evidence to conclude that the active threat pathway has been disrupted, related assets have been addressed or consciously accepted, the right owners have completed their actions, and recurrence monitoring is in place.

Closure must reflect the residual risk of each campaign. An account can disappear while its supporting domain network remains active, a phishing page can be removed after credentials have already been collected, and a counterfeit listing can return through another seller, so the closure standard must account for what remains capable of causing harm.

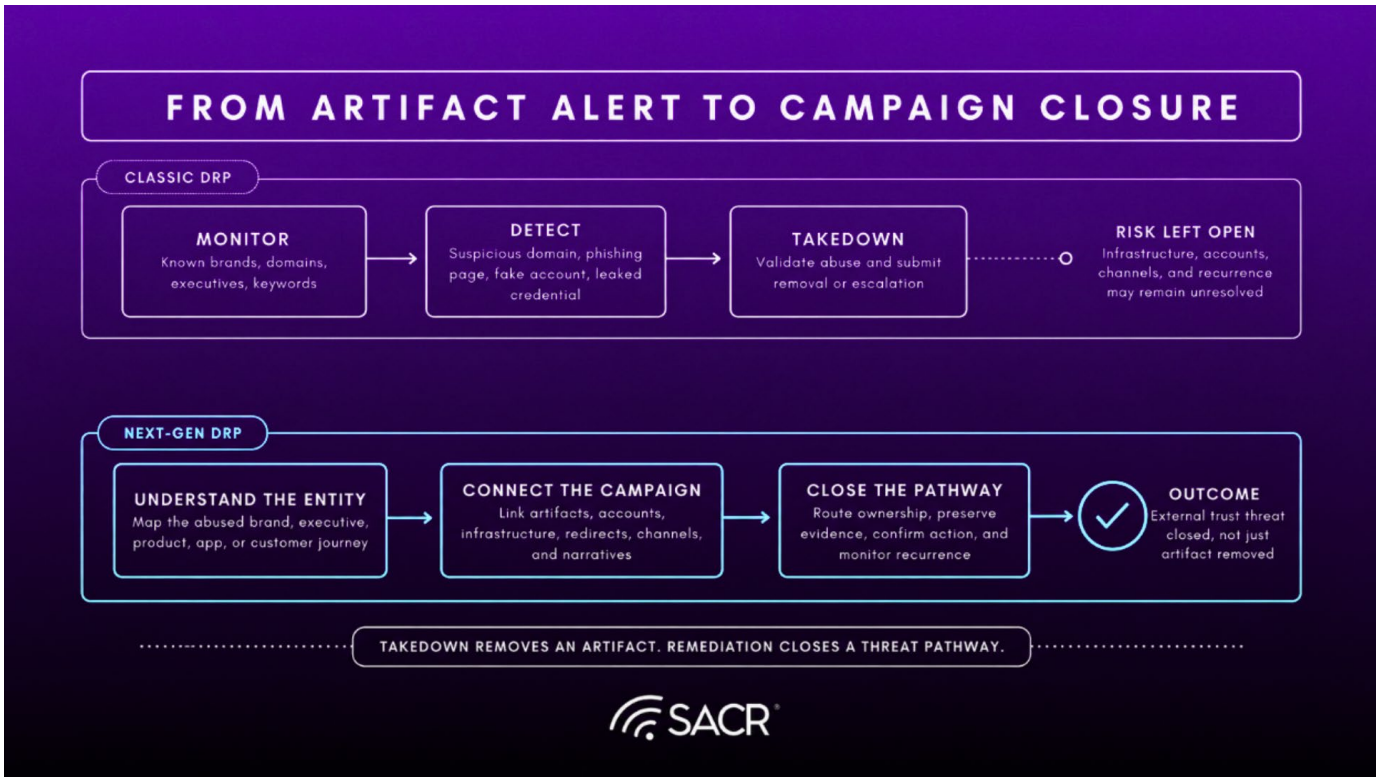

Performance should be measured across the full campaign lifecycle rather than through alert volume or takedown throughput alone:

- Discovery and analysis: time to discover unknown abuse, determine relevance and severity, and connect related artifacts into a campaign.
- Coordination and response: cross-functional handoff time, analyst effort per case, and time to remediation.
- Outcomes and residual risk: percentage of cases with verified closure, recurrence after action, unresolved campaign infrastructure, and customer, employee, executive, or brand harm.

Together, these measures show where the operating model breaks. Fast takedown can coexist with weak campaign understanding, high alert volume can conceal poor prioritization, and a low backlog can reflect cases that were closed before the threat was resolved.

Market Convergence and Category Boundaries

The next generation of DRP draws on markets that developed around different control points:

External monitoring and removal: classic DRP and brand protection contribute monitoring, investigation, takedown, trademark protection, counterfeit detection, and marketplace response.

Adversary and infrastructure context: threat intelligence, open-source intelligence, and external attack surface management add underground-community, campaign, infrastructure, and unmanaged-asset context.

Harm, trust, and people-centered response: fraud, trust and safety, narrative intelligence, and executive protection contribute monetization, synthetic-media, reputational, impersonation, and customer-abuse perspectives.

These categories overlap because an external campaign can move through several control points without fitting cleanly into any one of them. A platform may detect the infrastructure, another may understand the fraud path, and a third may own the relationship needed for removal or customer protection. Their value depends on whether the organization can connect those perspectives into one case.

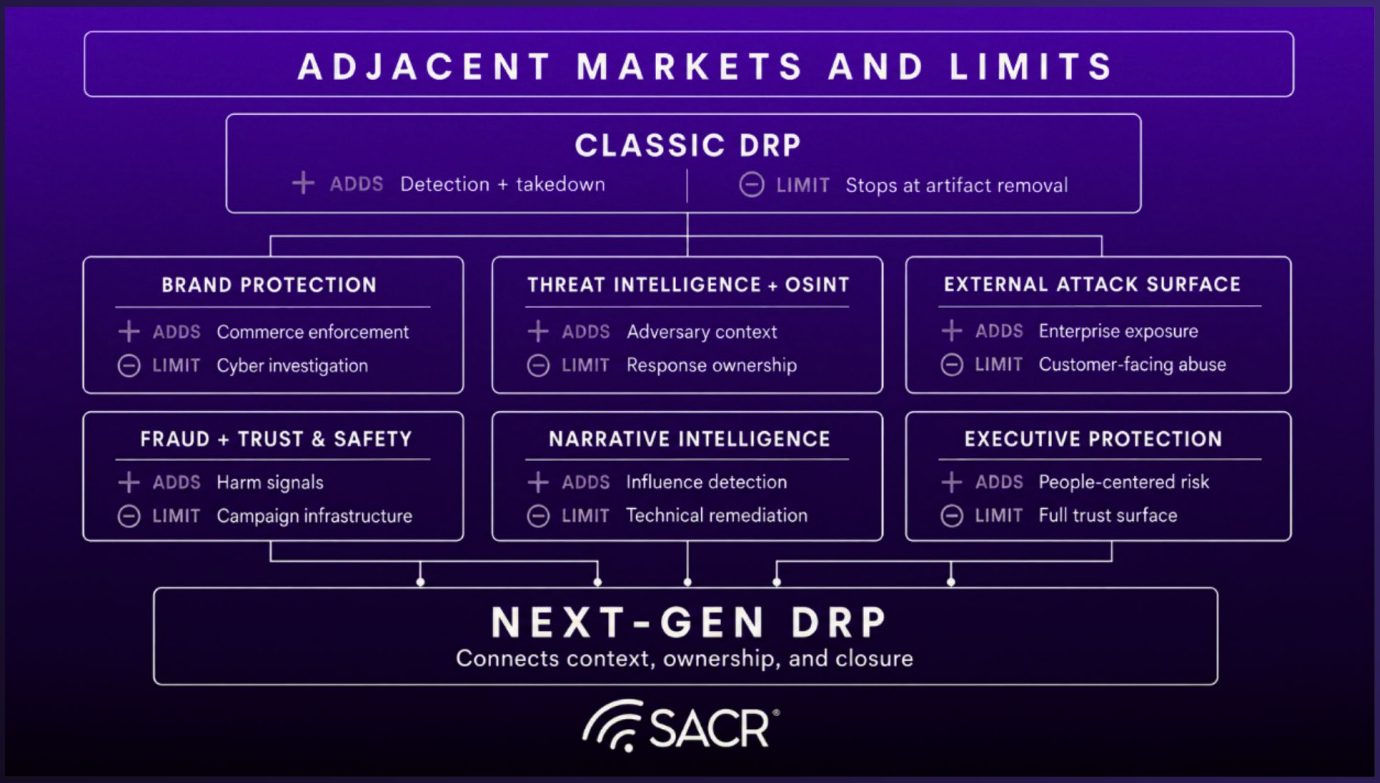

The market is therefore unlikely to collapse into a single product class. Organizations will continue to use security operations, threat intelligence, fraud, legal, communications, brand protection, and customer-support systems, while DRP vendors compete to become the connective layer around the external entity and campaign. That layer should assemble context, preserve evidence, coordinate action, and maintain an outcome record across the systems teams already use.

Vendors may appear in multiple containers because this is a convergence market. The important point is not to force every vendor into a single box; it is to show how different control points are converging around the same external trust problem.

How DRP Risks Show in Reality

Executive Impersonation and Fraud

Consider an executive impersonation campaign that develops into a customer fraud funnel. The organization first encounters a social account using a senior executive's name and image, but the account is only the entry point: it directs users to a cloned customer-support page, moves the conversation into Telegram or WhatsApp, and ultimately sends victims toward a credential-harvesting or payment flow.

A conventional workflow may find the first account through a watchlist, keyword match, or customer report, confirm the impersonation, and submit a takedown through security, legal, or brand protection. Closing the case when that account is reported or removed leaves the rest of the campaign free to continue through related profiles, fraudulent advertisements, messaging channels, cloned pages, redirects, and payment infrastructure. The adversary can also relaunch under a different name or shift victims to another platform.

An External Trust Operations workflow begins with the executive and the connected customer journey. Authorized accounts, official domains, known imagery, login flows, brand relationships, and public context establish what legitimate activity looks like; profile similarity, linked domains, forms, calls to action, infrastructure, and customer reports help determine whether a new signal belongs to the threat.

The investigation then expands from the first account into a campaign view that connects profiles, advertisements, domains, sites, messaging channels, infrastructure, personas, and monetization paths while showing which relationships are supported by evidence. The organization can prioritize the campaign by reach, velocity, customer and credential risk, fraud exposure, executive impact, and recurrence, then route a shared evidence base to the teams responsible for technical response, fraud intervention, platform action, audience warning, and executive protection.

Closure requires evidence that the organization addressed more than the initial account. The case remains active until related assets have been assessed, the fraud path has been contained or consciously accepted, continued victim reports have been checked, and recurrence monitoring is in place.

outtake

Outtake.ai

Outtake is a useful proof point for the shift toward entity-centric external trust defense. Its platform models protected entities, uses specialized agents to interpret multimodal external content, applies signal-based triage, maps related artifacts through threat graphs, coordinates platform-specific takedowns, and preserves context for response and recurrence analysis. The company is best understood as an illustrative solution profile rather than the sole definition of the category.

Why Outtake Fits the Thesis

Outtake organizes protection around brands, people, locations, and products, enriching each entity with relevant context such as authorized accounts, imagery, aliases, geocoordinates, and product information. This gives the platform richer context than static keyword lists alone.

This maps to the report's broader argument that the protected object is the organization's public trust surface, which Outtake describes as its Global Digital Trust Footprint. Entity context helps establish what is being impersonated, how the activity connects to a broader campaign, and who should act.

Search agents use entity descriptions, visual context, transcripts, text overlays, authorized social accounts, and web research to determine relevance. In one example, context-based filtering narrowed roughly 3,000 collected items to 389 relevant findings, illustrating how entity-aware reasoning reduces broad collection to an actionable set.

Agentic Search and Multimodal Analysis

Outtake's search layer looks beyond pre-approved watchlists and keyword strings across traditional web domains, social media, deep web, dark web, and customer-reported mobile threats such as smishing. Specialized agents interpret text, images, video transcription, logos, visual similarity, and indirect references before connecting findings to protected entities. Content Wall then aggregates agent-found material into a controlled view, allowing analysts to review external content without visiting each source platform directly.

Intelligent Workflows and Remediation

Outtake turns raw findings into explainable triage and response paths using signals such as external links, Telegram references, profile-image similarity, login forms, MX records, credential-harvesting JavaScript, and infrastructure clues. These signals inform risk scoring and help route related artifacts to the appropriate owners while preserving them within a common case.

Flagged queues and takedown workflows support graduated automation, moving from manual review toward automated enforcement as the buyer gains confidence in precision. Buyers should still validate production performance, approval requirements, and how closure is verified after action.

Recon Agent and Campaign Graphing

Recon Agent is the clearest product example of the shift from artifact monitoring to adversarial network understanding. In one example, an investigation began with a single suspicious social media account and expanded into a graph of 27 nodes and 31 edges across various social accounts, communication channels and servers, as well as a dedicated website and live app.

The significance is the operating pattern: a single suspicious artifact becomes the starting point for mapping related accounts, communities, infrastructure, and monetization paths. This moves investigation beyond isolated takedowns and gives analysts a stronger basis for prioritization, evidence packaging, disruption, and recurrence monitoring.

Adaptive Learning and Persistent Trust Context

Analysts can label findings, edit learnings, and feed those decisions into future triage. An observer-agent concept can also recommend backfill when new learning would alter the classification of existing alerts. Outtake's Digital Reservoir preserves the resulting entity context, suspicious variants, fraudulent assets, adversarial infrastructure, evidence, decisions, takedown outcomes, recurrence signals, and prior cases. Together, these capabilities support adaptive case quality and persistent external trust context.

Digital Trust Kill Chain, Provenance, and Intake

Outtake describes its Digital Trust Kill Chain as extending the MITRE ATT&CK approach beyond traditional cyberattacks to map eight stages across four zones:

1. **Zone 1: Discovery**
    a. Reconnaissance: The attacker profiles an organization's public attack surface before building the campaign.
    b. Infrastructure Setup: The attacker creates and stages the assets needed to support the operation, including fake domains, applications, and accounts.

2. **Zone 2: Exposure**
    a. Trust Exploitation: The staged infrastructure is made credible through brand impersonation, fake endorsements, and fraudulent use of trademarks, logos, and other recognizable signals.
    b. Target Engagement: The attacker initiates targeted contact with the intended victims through phishing, social media, messaging platforms, and malicious advertising.

3. **Zone 3: Compromise**
    a. Engagement Capture: The initial engagement becomes a controlled interaction between the victim and the staged malicious infrastructure.
    b. Value Transfer: The attacker leverages the controlled interaction to gain value through direct financial transaction, account takeover or compromise, or data exfiltration.

4. **Zone 4: Harm**
    a. Impact and Fraud: The attacker's operation produces a business or human impact, including consumer scams, reputation damage, data exposure, or risk to physical operations.
    b. Monetization: The attacker converts and launders the operational proceeds into financial return.

The analytical value that the Digital Trust Kill Chain brings is that it connects what can initially appear as separate artifacts into a progressing operation. Outtake also emphasizes that intervention becomes more expensive as the campaign moves from the Discovery and Exposure zones to Compromise and Harm. Buyers should treat Outtake's Digital Trust Kill Chain as an effective vendor-developed model for evaluating campaign maturity and response coverage.

Actor Tracing connects related activity, infrastructure reuse, and other actor context to support executive protection, adversary dossiers, remediation, and recurrence monitoring. Buyers should distinguish this form of provenance and adversarial-network analysis from definitive attribution, which remains difficult.

Outtake Intake structures customer- or user-reported scams, impersonation, smishing, and suspicious external activity into intelligence, alerts, and takedown workflows. This gives organizations another early-signal path for detecting and responding to external trust abuse.

Integrations and actionability

Outtake uses an open API model to move digital trust context into systems such as Microsoft Sentinel, ServiceNow, Jira, and OpenCTI. This supports SOC escalation, threat intelligence, workflow management, legal review, fraud investigation, and operational response without creating another disconnected queue. Buyers should distinguish customer-deployed integrations from directional or roadmap support.

Analytical assessment

Outtake is compelling because it aligns with the category's most important shifts: from keywords to entities, from single alerts to threat graphs, from monitoring to investigation, from takedown to remediation workflows, from static configuration to adaptive context, and from brand-only protection to broader external trust defense.

Outtake's category-shaping terms are useful, but the category should not become dependent on one vendor's language. Evidence is the second risk. Product examples show capability direction, while claims about reduced triage time, takedown success, recurrence reduction, or business harm reduction still require customer evidence, anonymized case studies, or measurable before-and-after data. Scope is the third risk. Outtake's model touches DRP, threat intelligence, executive protection, fraud, brand protection, narrative intelligence, and digital trust. That breadth is strategically interesting, but buyers should not assume that one platform replaces every adjacent category. The stronger framing is that Next-Gen DRP complements existing tools by connecting external trust signals around business entities and response workflows.

Buyer Scenario: Executive Impersonation Fraud Funnel

A practical buyer scenario helps make the category shift concrete. Imagine an executive-impersonation campaign that turns into a customer fraud funnel. The organization first sees a suspicious social account using a senior executive's name and image. The account links to a cloned customer-support page, pushes users toward a Telegram or WhatsApp channel, and redirects victims into a credential-harvesting or payment-redirection flow.

A traditional DRP workflow may find the visible artifact through a watchlist, customer report, or keyword match. An analyst validates that the account, page, or domain abuses the brand or executive identity. Legal, brand protection, or security submits a takedown request. The case may be treated as closed when the artifact is removed or reported. The problem is that related social accounts, messaging channels, fraudulent ads, redirect paths, cloned pages, and monetization infrastructure may remain active or reappear later.

A Next-Gen DRP workflow starts with the protected entity. The platform understands the executive, the associated brand, the customer login flow, authorized accounts, domains, products, and known public context. It can detect suspicious social profiles, domains, images, videos, ads, or messaging channels even when the attacker avoids exact keywords. It uses signals such as profile-image similarity, external links, login forms, Telegram or WhatsApp calls to action, infrastructure reuse, and content similarity to assess relevance and severity.

The platform then connects the campaign. It links the fake profile, cloned page, messaging channel, domain, ad, infrastructure, and monetization path into a single case or graph. It prioritizes the case against customer harm, credential-harvesting risk, fraud exposure, executive exposure, reach, velocity, and recurrence. It routes evidence to security, fraud, legal, communications, executive protection, customer support, and external platforms based on each team's role. It verifies whether visible artifacts were removed or contained, whether related infrastructure remains active, and whether recurrence appears after takedown. It also learns from the case by using analyst feedback, takedown outcomes, recurrence patterns, and entity context to improve future triage.

Buyer Action Plan: First 90 Days

A useful evaluation begins with the organization's external trust problem, tests vendors against realistic campaigns, and moves a limited set of high-value use cases into production. The first 90 days should build evidence progressively rather than treating a feature demonstration as proof of operating performance.

First 30 Days: Define the External Trust Baseline

The first month should establish what the organization is protecting, how work moves today, and which measures will show whether the program improves.

- Define the highest-priority entities and the context required to protect them, including aliases, imagery, official accounts, domains, geographies, business relationships, customer journeys, and common abuse patterns.
- Map discovery, investigation, escalation, takedown, verification, and recurrence monitoring across security, fraud, legal, brand, communications, customer support, and executive protection. Identify where cases stall, split into duplicate work, or close without verification.
- Identify the systems that need external trust context and establish baseline measures for alert volume, analyst effort, severity decisions, campaign connection, handoff time, remediation, recurrence, and premature closure.

Within 60 Days: Test Discovery, Reasoning, and Campaign Connection

The second month should test whether the platform can move from isolated signals to an explainable campaign view under realistic operating conditions.

- Run current or representative scenarios in which adversaries avoid exact keywords, rely on images or video, use multilingual content, change platforms, or separate the initial lure from the eventual fraud or credential path.
- Require the platform to explain relevance and campaign connections through inspectable entity context, content, infrastructure, relationships, and evidence. The campaign view should distinguish supported relationships from weak or inferred links.
- Confirm that evidence, analyst decisions, indicators, screenshots, infrastructure details, and action history can move into each owner's systems. Validate deployed integrations, remediation approvals, and the boundary between observable actor context and unsupported attribution.

Within 90 Days: Operationalize Remediation and Verified Closure

The final month should place a bounded workflow into production and test whether the organization can prove that its actions changed the active threat.

- Launch a limited production workflow around a small set of high-priority entities and abuse scenarios, with clear response ownership, evidence requirements, approval gates, and closure conditions.
- Measure discovery, severity decisions, campaign connection, handoff time, analyst effort, remediation, recurrence, unresolved infrastructure, and verified closure against the baseline established in the first month.
- Reopen apparently closed cases to check for surviving assets, platform migration, continued victim reports, or new infrastructure. Use the results to improve entity context, workflow rules, recurrence monitoring, and recurring governance.

Market Implications

Source coverage will remain part of DRP buying decisions, but it will carry less weight on its own as campaigns become more connected and operationally complex. Buyers will increasingly assess whether a platform can model protected entities, discover unknown abuse, interpret multimodal content, connect campaigns, explain prioritization, preserve evidence, coordinate workflows, support remediation, and monitor recurrence.

Competitive pressure will come from several vendor lineages. Established DRP and brand-protection providers bring collection breadth, takedown experience, and mature customer workflows; threat-intelligence and external-attack-surface vendors contribute infrastructure and adversary context; fraud and trust-and-safety platforms understand customer abuse and monetization; and newer AI-native vendors can make search, context assembly, and investigation more adaptive. Buyers should judge the combined operating outcome, using vendor origin as context rather than as a proxy for capability.

Category boundaries will continue to blur as vendors use terms such as digital risk, digital trust, external threat intelligence, brand protection, fraud intelligence, executive protection, and AI-native DRP for overlapping functions. Clear operational requirements will matter more than the label attached to the product.

The longer-term direction is External Trust Operations, an environment that maintains reliable context about the organization's public entities, understands threats against the trust relationships around them, and coordinates responses across security and business teams. DRP is an important route into that future because it already owns many of the external signals and remediation relationships, but the market will earn the expansion only by showing measurable reductions in analyst effort, time to campaign understanding, recurrence, and business harm.

Conclusion

The DRP category reset is an operating-model change. Organizations still need to find fake domains, phishing pages, impersonating accounts, leaked credentials, counterfeit listings, fraudulent apps, and other external abuse. They also need to understand how those artifacts connect, which business entity and trust relationship are under attack, who must respond, and whether the campaign remains active after action.

External Trust Operations provides that structure. It organizes the program around protected entities, campaign analysis, coordinated response, evidence continuity, and verified closure. The model also gives CISOs a more useful way to evaluate vendors because it connects technical capability to operational performance and business harm.

A DRP program is moving in this direction when it can discover unknown abuse, explain why it matters, map the wider campaign, route action with the required evidence, and prove that the active threat pathway has been closed. If those steps still depend on manual reconstruction across disconnected tools and teams, the organization has an External Trust Operations gap. The strongest vendors will help the organization close external trust threats faster, with better context, less coordination burden, and stronger evidence that the risk is actually resolved.

Software Analyst®
Cyber Research

business

personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

CISO	VENDOR
SECURITY TEAMS	INVESTORS