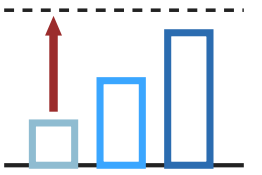
Principal Brief // two-minute read

DOMAIN & DNS | MONITORING & ALERTING | MISIDENTIFICATION



A Warning Is Not a Verdict

> *Security vendors can now warn you about a domain before it does anything. That warning is worth having, and it is the thinnest evidence your program will ever act on. Block it if you can unblock it. Do not let it name anyone.*

WHAT HAPPENED

Attackers register domains and obtain certificates before they attack, and some of that is visible from outside. Vendors sell the visibility as preemptive security. The best published system of its kind scores 99.87 percent precision, but it was tested against 4,358 of the most popular sites on the internet. Applied to the 140,000 new .com and .net domains registered daily, its error rate works out to roughly 570 wrong answers a day.

WHY IT REACHES YOU

Because the wrong answers land near you. The score attaches to a relative's new domain, counsel's registrar account, a counterparty's new subsidiary, a household employee's side business. Nobody writes an accusation. The domain is quietly blocked and the person is quietly treated as hostile, on evidence that would not survive being written down.

THREE THINGS TO TAKE FROM THIS

1

Earlier means thinner

The earlier a warning arrives, the less the domain has actually done. That is the price of the warning, not a defect in it.

2

Scale the act to the proof

A reversible, expiring block can run on a probability. Naming a person, exiting a vendor, or making a referral cannot.

3

Reversibility is a build, not a promise

A block is only reversible if it has an expiry date, a named owner, and a route to lifting it that is not a meeting.

DO THIS

- [ ] **YOU, WHEN SHOWN A WARNING**
  Ask what was observed, what was inferred, what would prove it wrong, and what happens if it is.
- [ ] **YOU, WITH YOUR FAMILY OFFICE**
  Approve expiring, device-scoped blocks at low confidence. Withhold any action that touches a person, and never block a whole hosting provider.
- [ ] **YOU, WITH YOUR SECURITY LEAD**
  Ask for lead time, precision by confidence band, and reversals. "Campaigns that never happened" is not a result.

NOT THE STORY YOU MAY HAVE HEARD

This brief does not say pre-attack detection is fake, and it does not say wait for the compromise. Early warning buys real decision time. It says the threshold for acting early and the threshold for accusing someone are different thresholds, and most of the damage comes from treating them as one.

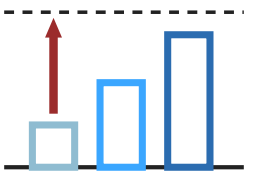
HOW SURE WE ARE

High that preparation is observable and that earlier means less certain. High on the published figures and the arithmetic drawn from them. Moderate on how commercial products compare, because none publish reproducible results.

THE FULL COMMAND BRIEF FOLLOWS: WHAT IS OBSERVABLE, THE BASE-RATE ARITHMETIC, THE ACTION LADDER, AND POSTURE BY FUNCTION.

CB-054

The Warning Arrives Before the Proof



"Earlier warning is bought with less certainty. The action has to be scaled to the proof."

> *Adversaries buy domains, stand up hosting, and obtain certificates before they attack, and some of that preparation is visible from outside. A vendor can therefore hand a principal, the protected person at the center of a security program, a warning about infrastructure that has not yet done anything. That warning is worth having. It is also, by construction, the least-evidenced thing the security program will ever act on, because the earlier the observation, the less completed hostile conduct exists to verify. The failure is not the warning. The failure is treating it as a finding.*

BLUF. Pre-delivery infrastructure signals are real and observable. MITRE ATT&CK classifies acquiring domains (T1583.001) and staging capabilities (T1608) as pre-compromise activity, and the traces are in registration records, DNS, certificate transparency logs, and hosting relationships. What the market sells on top of those traces is a probability, not a fact, and the public literature does not yet support the confidence with which it is sold. The best published system of its kind, DomainDynamics (Chiba, Nakano and Koide, 2025), reports 99.87 percent precision at a 0.41 percent false-positive rate, but its negative set was 4,358 domains drawn from the Tranco Top 10,000, about eighteen false positives in the whole test. Verisign's registry data puts new .com and .net registrations at 12.7 million in the second quarter of 2026, roughly 140,000 a day in two top-level domains. The same 0.41 percent applied to that stream is about 570 false positives a day. Meanwhile the window the capability sells is short: in a study of 15,126 newly registered phishing domains, the median lifetime was one day and 89.45 percent were active for under two days.

The governing frame: the evidentiary threshold must rise with the consequence of the action. A reversible, scoped control can act on a probability. Naming a person, ending a relationship, offboarding a vendor, or making a referral cannot act on the same evidence. A program that keeps those thresholds separate gets the benefit of early warning. A program that collapses them converts a score into an accusation, and does so against people close to the principal more often than against adversaries.

The point of an early warning is that the evidence is not finished. The response has to be built for that.

THE SIGNAL // WHAT IS ACTUALLY OBSERVABLE BEFORE DELIVERY

THE PRE-COMPROMISE RECORD · STANDING

MITRE ATT&CK places Acquire Infrastructure: Domains (T1583.001) and Stage Capabilities (T1608) in the Resource Development tactic and marks them PRE, meaning before compromise. A domain can be registered, delegated to nameservers, given hosting, issued a certificate, and connected to earlier infrastructure before it delivers anything. That sequence is the entire basis of pre-attack visibility, and it is well documented.

THE MEASURED SHAPE OF A HOSTILE DOMAIN · 2026

Agarwal and Vasek studied 15,126 newly registered phishing domains observed between 5 July 2023 and 23 May 2024. Average lifetime was 8.6 days and the median was one day, with 89.45 percent active for under two days. About 27 percent were re-registrations of domains with prior malicious history, which is a genuine early signal. Their population was domains already identified as phishing, so the study measures the life of a hostile domain rather than the accuracy of scoring one before first use. It is the best public sizing of the window, and the window it describes is often hours.

WHAT A CERTIFICATE PROVES · 2026

In the same population, 79.3 percent of newly registered phishing domains were issued TLS certificates, and 26 percent sat behind Cloudflare's reverse proxy. Certificate transparency logs are therefore a real observation channel and a certificate is not a signal of legitimacy. The corollary matters more for defense: because a quarter of hostile domains share the infrastructure of an enormous number of legitimate ones, blocking on a hosting or proxy relationship sweeps in bystanders at scale.

WHAT THE MARKET IS SELLING · SEP 2025

Gartner forecast on 18 September 2025 that preemptive solutions would grow from under 5 percent of IT security spending in 2024 to 50 percent by 2030. Read the category boundary carefully: Gartner's preemptive bucket covers predictive threat intelligence, advanced deception, and automated moving target defense. This brief addresses only the first, and only its domain and infrastructure form. A spending forecast for a broad category is evidence of where money is going, not evidence that this particular class of signal performs.

THE LAYER UNDER THE LABEL · JUN 2026

The NSA and CISA information sheet Selecting a Protective DNS Service reached version 1.5 in June 2026, and its version history records the reason: "Updated HYAS Protect to its new name, threatER EnforceDNS, in Table 1." The federal guidance that lists a principal's protective-DNS options had to be edited because a listed provider changed hands. HYAS was split in December 2025, its enforcement product going one way and its intelligence product the other. Intelligence and enforcement are separate layers with separate owners, however unified the marketing.

SO WHAT The preparation is genuinely visible, the interval is often measured in hours, and the only published accuracy figure was measured against a few thousand of the most popular sites on the internet. All three are true at once.

THE MECHANISM // EARLIER MEANS MORE AMBIGUOUS

THE TRADE THAT DOES NOT GO AWAY

A delivered phishing page can be inspected. A payload can be detonated. A command-and-control connection can be watched. Before those things happen, a classifier is working from preparation and resemblance, and legitimate operators also register domains, obtain certificates, change DNS, and buy cheap hosting behind large providers. The signal arrives earlier because it does not wait for hostile use, and that same choice removes the evidence that would make the conclusion certain. Better data and better models move the curve. They do not abolish the trade between lead time and evidentiary completion.

THE BASE RATE IS THE WHOLE PROBLEM

Precision measured on a curated negative set does not survive contact with the live registration stream. DomainDynamics reported a 0.41 percent false-positive rate against 4,358 benign domains from the Tranco Top 10,000, the easiest negatives in existence, which is about eighteen false positives in the test. Applied to roughly 140,000 new .com and .net registrations a day, the same rate is about 570 a day, in two top-level domains, before any other TLD is counted. Nothing about that arithmetic says the research is weak. It says a rate that looks excellent in a paper becomes an unmanageable queue, or an unnoticed pile of collateral blocks, when it meets default-deny at internet scale.

TWO DIFFERENT PRODUCTS, SOLD AS ONE

Protective DNS resolving against known-bad infrastructure is mature and endorsed: the NSA and CISA sheet describes blocking known malware and command-and-control domains, phishing sites including close lookalikes, and domains with known algorithmic-generation attributes. Predictive scoring of infrastructure that has not been used yet is a different product with a different error profile, and the same sheet says nothing about false positives at all. When a vendor presents both under one preemptive label, the second borrows the legitimacy of the first. Ask which of the two a given block came from.

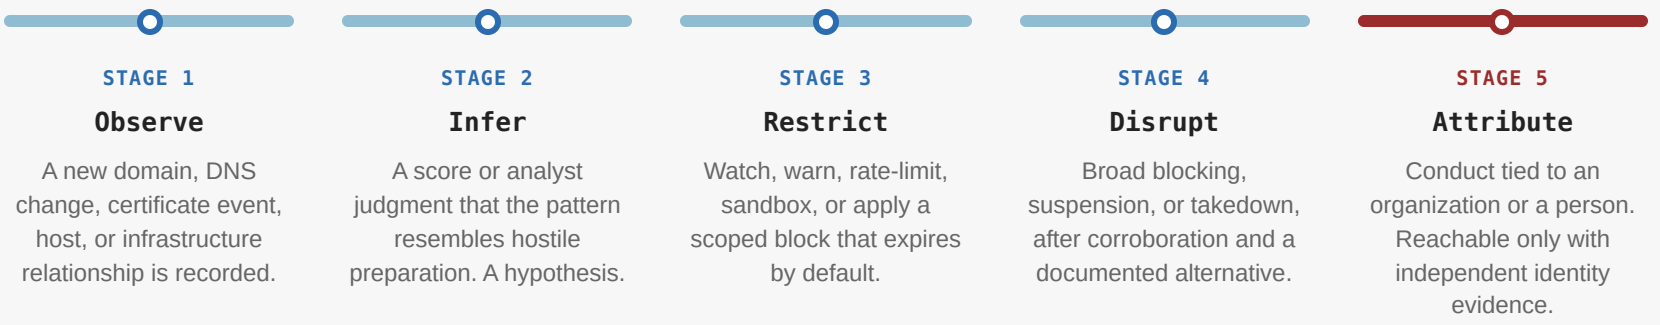
THRESHOLD COLLAPSE

A score built to prioritize an analyst's queue becomes a blocklist entry. The blocklist entry becomes a cluster. The cluster becomes an actor. The actor becomes a named person. Confidence appears to rise as the claim travels even though no new evidence entered the chain, and each handoff strips the qualifier that made the original number honest. This is the failure mode to design against, and it is a process failure rather than a modeling failure.

HOW IT REACHES A PRINCIPAL: PROXIMITY BECOMES SUSPICION

The family-office version of threshold collapse never names anyone. A score attaches to a relative's newly registered domain, to counsel's registrar account, to a counterparty's freshly stood-up subsidiary, or to a household employee's side business. Nobody writes an accusation. The domain is quietly blocked, the counterparty is quietly deprioritized, the relative is quietly discussed. Infrastructure proximity has become social suspicion with no document to challenge and no threshold anyone had to meet. As CB-053 shows, shared infrastructure now proves less than it used to, which makes the quiet version of this error more likely, not less.

FIGURE 1 // THE ACTION LADDER: FIVE STAGES, AND THE ONE THAT MUST NOT BE REACHED BY INFERENCE



THE WINDOW

Stage 3 is the only stage a principal or family office actually buys. Stages 4 and 5 belong to a platform, a registrar, a court, or law enforcement, and a vendor score is not a ticket to either.

▶ THE PRACTICAL TEST

Ask of any pre-attack warning: *what was observed, what was inferred, what would falsify it, and what happens if it is wrong?* A capability that cannot answer the fourth question has not been engineered for the third.

FIGURE 2 // EVIDENCE, OWNER, AND REVERSAL BY ACTION

ACTION	EVIDENCE THAT JUSTIFIES IT	WHO DECIDES, WHO CAN REVERSE	REVERSAL TARGET
MONITOR OR ENRICH	A single weak or novel signal is enough	Analyst or automation; no approval needed	Not applicable
WARN A USER OR ANALYST	Several signals, or one distinctive one	Analyst; human review stays in the loop	Immediate
SCOPED, EXPIRING BLOCK	A calibrated score, with fast reversal available	Security lead decides; service desk can reverse without escalation	Same business day, on request
BROAD BLOCK OR TAKEDOWN	Strong technical corroboration and a documented alternative route	Named owner with authority to accept business impact	Reviewed on a fixed date, not by inertia
NAME A PERSON OR ORGANIZATION	Independent identity evidence and the CB-246 validation gate	Counsel plus the principal's decision-maker	Not reversible in practice, which is why the gate exists

FIGURE 3 // A WORKED EXAMPLE: A LOOKALIKE DOMAIN, ONE DAY AT A TIME. CONSTRUCTED TO ILLUSTRATE THE RULE, NOT A CASE FILE.

MOMENT	WHAT IS ACTUALLY KNOWN	PROPORTIONATE ACTION	OUT OF BOUNDS
DAY 0, 09:00	A domain one character from the family office's own is registered at a registrar the office does not use, with privacy protection on the record. No hosting, no certificate, no mail.	Add to watch. Ask the office's registrar about a defensive registration of the remaining variants.	Telling anyone a threat has been identified.
DAY 0, 16:00	Nameservers are delegated and a certificate is issued. The certificate proves nothing: 79.3 percent of hostile new registrations get one.	Scoped block at the resolver for staff devices, expiring in fourteen days, with a ticket that names who can lift it.	Blocking the whole hosting provider or proxy network behind it.
DAY 1	A login page appears that copies the office's branding, and a staff member forwards an email pointing at it.	Now there is conduct. Extend the block, report to the registrar and the hosting provider, preserve evidence, notify staff and counterparties.	Naming whoever is behind it on the strength of registrar or hosting overlap.
DAY 1, THE ALTERNATIVE	The domain resolves to a small consultancy with the same initials that has been trading for six years. The overlap was the shared proxy.	Lift the block inside the expiry window, record the false positive, and feed it back to the vendor as a scored miss.	Leaving the block in place because removing it takes a meeting.

THE MEASUREMENT PROBLEM // A NON-EVENT PROVES NOTHING BY ITSELF

Campaigns that never materialize cannot be the measure of success, because that formulation cannot be audited. A flagged domain may never attack because it was blocked, because the operator abandoned it, or because the prediction was wrong, and the absence of an incident does not separate those three. Any metric that counts silence as prevention will report success under all of them.

What can be measured is the prediction against later evidence: how often each confidence band is confirmed, rejected, or left unresolved; how far the warning preceded first known weaponization or conventional listing; and what the control actually did, in blocked connection attempts, reversals, and analyst hours. Unresolved predictions have to stay in the denominator, because a vendor that drops them can count its confirmed hits, ignore its silent failures, and describe every non-event as a save.

The number a buyer most needs does not exist in public. Nobody publishes the precision of a new-domain risk score, by confidence band, measured against the live registration stream at 24 hours and again at 7 days. The academic figures are measured against curated benign sets, and the commercial figures are not published at all. A buyer should ask for that number, expect not to get it, and treat the shape of the refusal as information. A buyer should also be honest that scoring a vendor properly requires telemetry the buyer may not have, which is an argument for a pilot with a held-back sample rather than for a benchmark nobody can run.

FIGURE 4 // THE ARITHMETIC OF A GOOD FALSE-POSITIVE RATE

140K	0.41%	~570	1
New .com and .net registrations a day, from 12.7 million in the second quarter of 2026. VERISIGN DNIB, 23 JUL 2026	Best published false-positive rate, measured against 4,358 Tranco Top 10,000 domains. CHIBA, NAKANO & KOIDE, 2025	False positives a day if that rate meets the two-TLD registration stream. OBSCUREIQ CALCULATION FROM BOTH	Median lifetime in days of a newly registered phishing domain; 89.45 percent last under two. AGARWAL & VASEK, 2026

THE STEELMAN // THE CASE FOR ACTING EARLY ANYWAY

THE BEST CASE FOR THE OTHER SIDE

+ Waiting for weaponization spends the only cheap interval there is. A domain blocked before it is armed costs a resolver entry; a credential harvested from a family office costs considerably more.
+ Most defensive actions at Stage 3 are genuinely reversible, and the asymmetry favors acting: a lifted block is an inconvenience, a successful intrusion is not.
+ Public benchmarks are hard for real reasons. Vendors rely on proprietary telemetry, ground truth is contested, and adversaries adapt to any detection method that is published in detail.

All three land, and the answer is not to wait for compromise. It is that the reversibility the argument depends on has to be engineered rather than assumed. Reversibility means a stated expiry, a named owner, a route to reversal that does not require a meeting, and a record of what was blocked and why. Where those exist, acting on a probability is sound practice. Where they do not, "reversible in principle" describes a block that will still be in place in two years, attached to a business that never learned why its email stopped arriving. The steelman is an argument for building the reversal machinery, not for skipping it.

ASSESSMENT

FIGURE 5 // ANALYTIC CONFIDENCE BY CLAIM

Adversaries acquire and stage infrastructure before delivery, and some of it is externally observable.	**STRONG / MITRE ATT&CK, PUBLISHED RESEARCH**
Earlier classification carries less completed evidence of hostile intent, and that trade cannot be modeled away.	**STRONG / STRUCTURAL**
A false-positive rate measured on curated benign sets understates the load at registration-stream scale.	**STRONG / ARITHMETIC FROM PUBLISHED FIGURES**
The exploitable window is frequently hours, not days, for the commodity phishing class.	**STRONG / MEDIAN ONE-DAY LIFETIME, 15,126 DOMAINS**
Pre-attack intelligence is becoming a larger procurement category.	**MODERATE / DIRECTIONAL, FROM ONE FORECAST FOR A BROADER CATEGORY**
Precision of new-domain scoring by confidence band against the live registration stream.	**UNKNOWN / NOT PUBLISHED BY ANYONE, ACADEMIC OR COMMERCIAL**
Pre-attack detection is already a procurement baseline.	REJECTED / UNDER 5 PERCENT OF SECURITY SPENDING IN 2024
Vendor consolidation has demonstrated lower false-positive rates.	REJECTED / NO PUBLISHED EVIDENCE, ON EITHER SIDE
A campaign that does not materialize proves prevention.	REJECTED / NO COUNTERFACTUAL

POSTURE

FOR PRINCIPALS

→ When you are shown an early warning, ask what is proposed and whether it can be undone. If the proposal is to act against a person, an employee, a relative, or a counterparty, the answer is that infrastructure evidence does not reach that stage, and the CB-246 gate applies.

→ Ask which of the two products a block came from: known-bad enforcement, which is mature, or prediction about infrastructure that has not been used, which is not. They are priced together and they fail differently.

FOR FAMILY OFFICES

→ Grant standing permission for Stage 3 and withhold it for Stages 4 and 5. In practice that means resolver-level blocks on specific hostnames, scoped to managed devices, expiring in fourteen days unless renewed, with a named person who can lift one the same day and a ticket recording every reversal.

→ Register the defensive variants of the office's own domains and the principal's name, and put lookalike monitoring on them. That converts most of this problem from prediction into a list you already own.

→ Never block by hosting provider, proxy network, or autonomous system. A quarter of hostile new registrations sit behind one large proxy, and so do millions of legitimate businesses, including vendors, payment portals, and travel providers you depend on.

FOR SECURITY, LEGAL & PROCUREMENT

- → Make vendors define, in writing, what pre-attack means in their product, the unit of prediction, the confidence bands, the evaluation window, the ground truth, how false positives are treated, and what happens to unresolved cases. Then run a pilot against a held-back sample of your own traffic rather than accepting a benchmark that cannot be reproduced.
- → Separate intelligence from enforcement in the record: which system produced the judgment, which system acted, which team can reverse it. The December 2025 disposition of HYAS is the reason this is not academic, and the June 2026 revision of the NSA and CISA protective-DNS sheet is where it showed up in federal guidance.
- → Write the rule that an automated infrastructure score is never sufficient basis for an employment decision, a vendor offboarding, a counterparty exit, an insurance disclosure, or a referral to law enforcement. Those are the paths by which a probabilistic security signal becomes an adverse action against a named party, and they carry legal exposure that the security team does not own. This is intelligence analysis, not legal advice.
- → Set expiry and review on every predictive block. A stale prediction that nobody revisits becomes permanent adverse infrastructure history, and it will be read years later as a finding rather than as a guess that was never resolved.

CONFIDENCE High that pre-delivery infrastructure activity exists and is partly observable, which rests on MITRE ATT&CK's classification and on the published studies cited. High on the figures quoted from those studies and from Verisign's registry data. High on the base-rate arithmetic, which is division applied to two published numbers and is offered as an illustration of scale rather than as a measurement of any product. Moderate on the comparative performance of commercial products, because no vendor publishes reproducible evaluation. The one figure a buyer most needs, precision by confidence band against the live registration stream at 24 hours and at 7 days, is not published by anyone and is marked Unknown rather than Moderate. Gartner's spending figure is a forecast for a broader category and is treated as direction only. The Agarwal and Vasek population is domains already identified as phishing, so it sizes the window and the certificate and proxy prevalence, not the accuracy of scoring before first use; that adjacency is stated where the figures are used.

METHOD Revised edition. First published 2 May 2026; facts re-verified and analysis revised 10 September 2026. This edition argues from the evidentiary threshold rather than from the market: it sets out what is observable before delivery, what the published research does and does not measure, the arithmetic that turns a good false-positive rate into a large daily load, the five-stage action ladder with the stage a principal can actually buy, a worked lookalike example, and the measurement standard a buyer should hold a vendor to. Built from primary sources: MITRE ATT&CK, two peer-reviewed studies, Verisign registry data, NSA and CISA guidance, Gartner's forecast, and the parties' own acquisition announcements. Cross-read against CB-052, CB-053, and CB-246.

SOURCES MITRE ATT&CK, Acquire Infrastructure: Domains (T1583.001) and Stage Capabilities (T1608). Sharad Agarwal and Marie Vasek, "Examining newly registered phishing domains at scale", Journal of Cybersecurity 12(1), 2026, doi:10.1093/cybsec/tyag020. Daiki Chiba, Hiroki Nakano and Takashi Koide, "DomainDynamics: Lifecycle-Aware Risk Timeline Construction for Domain Names" (2025), with the extended version in Computers & Security. Verisign DNIB, second quarter 2026 (23 July 2026). NSA and CISA, Selecting a Protective DNS Service, version 1.5 (June 2026). Gartner on preemptive capabilities (18 September 2025). threatER on the HYAS Protect assets (11 December 2025) and Silent Push on HYAS Insight (15 December 2025).

Scope note: this brief covers warning drawn from domain, DNS, certificate, and hosting observables. It does not assess deception, moving target defense, or the rest of what is marketed as preemptive security. The 570-per-day figure is arithmetic applied to two published numbers to show the scale of the base-rate problem; it is not a measurement of any product's output. Figure 3 is constructed to illustrate the evidence rule and is not a case file. Acquisition announcements substantiate ownership, not efficacy. This is intelligence analysis, not legal advice.

OBSCUREIQ INTERNAL ANALYSIS // ARCHIVE COMM-054-PREATTACK-EVIDENCE-2026Q2 // XREF **CB-052 CB-053 CB-246** // ACCESS PUBLIC · PUBLIC-INTEREST RELEASE · REVISED 10 SEP 2026