

COGENT RESEARCH

Beyond the Human Horizon

Cogent Threat Report

Q4 2026



About this research

Every attack path to a crown jewel asset is a latent risk, and defenders have always prioritized those risks by asking what a human hacker would plausibly do. The compromise of Hugging Face in July 2026 challenged that assumption.

The swarm sustained lines of attack that no human team would have continued

The Hugging Face attack was carried out by a swarm of hundreds of coordinated AI agents. Investigators concluded that none of the individual techniques used were novel. What set the attack apart was continuous operation, parallel coordination across agent instances, and indifference to effort: the swarm sustained lines of attack that no human team would have continued.

That indifference opens a class of attack paths that human adversaries have never touched. Every environment contains paths that run too deep, return too little per step, or demand too many specialties to justify a human operator's time. Those paths have been protected by attacker economics alone. Agent swarms now make them exploitable.

The shift arrives against a backdrop of exponential growth in published vulnerabilities. Each new vulnerability is a potential link in a path, so path formation accelerates as disclosure volume climbs.

Enterprise complexity compounds the problem from the defender's side. Attack paths cross endpoint, identity, network, and cloud infrastructure, and no single security tool observes enough of the environment to reconstruct a complete path. The findings that form a path sit in separate consoles. This report quantifies the attack paths humans and machines can exploit: how fast they form, how deep they run, how long they persist, and why no single tool sees them.



The number of new vulnerabilities is exploding

HUMAN-VIABLE PATHS


Paths whose depth, per-step yield, and technique requirements fall within the envelope of documented human intrusion activity. A professional intrusion team could be expected to attempt them.

MACHINE-VIABLE PATHS

Paths that terminate at a crown jewel asset but exceed at least one human-economics threshold: too deep, too low-yield per step, or requiring more technique breadth than human operations sustain. Traversal is practical only for automated attackers.

THE HUMAN HORIZON

The set of thresholds that separates the two populations. Each threshold is calibrated against published red team engagement data and DFIR incident reporting, including median steps to objective and time on target in documented intrusions.





Key findings

Three new machine-viable attack paths appear for every human-viable one

01

In August 2026, 45 new attack paths opened up to crown jewel assets at the average enterprise: 11 viable for human attackers and 34 viable only for AI agents. Human-viable path formation grew 38% year over year, from 8 to 11 per month. Machine-viable formation surged 386%, from 7 to 34, as the offensive capability of AI models advanced.

Machine-viable paths run twice as deep and span five technique domains

02

Machine-viable paths run at least twice as deep as human-viable paths on four depth dimensions: a median of 8 asset hops versus 4, and 5 trust-boundary crossings versus 2. The median machine-viable path spans five distinct technique domains, and 64% span four or more, a breadth that has historically required a team of specialists.

83% of machine-viable attack paths persist for more than 30 days

03

Machine-viable paths outlast human-viable ones: 83% versus 55% were present in the environment for more than 30 days. These paths are built from medium and low severity vulnerabilities, configuration weaknesses, and identity conditions that severity-based remediation policies rank below their thresholds. Chaining low and medium severity vulnerabilities alone produced 7% of the critical paths in our dataset.

64% of attack paths cross domains that no single security tool observes

04

The majority of attack paths move between endpoint, identity, network, and cloud infrastructure, and reconstructing the median critical path required correlated data from five security tools. Every finding along these paths existed in at least one console. No single product held enough of the environment to join them into a path.



Difference between a human-viable attack path vs. a machine-viable attack path


A cyber attack using a swarm of AI agents can use the same techniques as those run by humans. What makes them different is continuous operation, parallel coordination across hundreds of agent instances, and indifference to effort.

Several hundred AI agents working in parallel can methodically try every technique on every asset, while a human adversary will develop hypotheses about what techniques may be successful, and prioritize a few of them. If they don't succeed after a few tries, they give up and try another path.

Human-viable attack paths have a higher yield per hop

Median yield per hop for four yield types across critical attack paths.


HUMAN-VIABLE MACHINE-VIABLE





Human-viable attack paths return more in the first steps

Cumulative yield against hop index for the median critical attack path. Documented human operations typically abandon unproductive lines of attack around hop 3, which is the depth at which 82% of crown-jewel paths accessible to machines have yet to return secrets or privileges.



For each hop, we measured four forms of yield: credentials and keys made readable, privilege gained, newly reachable assets, and access to sensitive data. Human-viable paths deliver high yield from the first hops. Machine-viable paths return a fraction of that value per step. In our assessment, low per-step yield is the reason these paths go unexploited, because an intrusion team operating under time and detection constraints has little incentive to continue a line of attack that returns almost nothing at each step. That constraint does not apply to automated attackers, for which the marginal cost of an additional hop approaches zero.



Three new machine-viable attack paths now appear for every human-viable one


At every organization, new attack paths are emerging all the time as credentials are issued, permissions change, software is deployed, infrastructure is reconfigured and new vulnerabilities are discovered.

In August 2026, the average enterprise saw 45 new attack paths emerge to crown jewel assets. These are latent risks. Unless an attacker discovers them, nothing happens. But for every one attack path a human adversarial team could exploit, three emerge that are only viable for AI exploitation.

There has been a significant increase in the number of machine-viable attack paths this year. Human-viable attack paths increased 38% from 8 to 11 per month. But machine-viable attack paths surged 386% to 34 per month. While the number of new vulnerabilities is driving an overall increase in the number of viable attack paths, the greater increase in machine-viable attack paths is due to the rapid advancement of AI models. AI agents are much more capable of executing cyber attacks today than they were a year ago.

The average enterprise has 45 new attack paths each month

New attack paths to crown jewel assets, August 2025 vs August 2026





Machine-viable attack paths run twice as deep and require a greater breadth of attack techniques

We measured attack path depth across four dimensions. Machine-viable paths were longer under every rule, and the difference was largest for action hops and trust-boundary crossings.


Those hops carry the highest detection risk for a human operator: each action and each boundary crossing increases the likelihood of triggering an alert that ends the operation, which is why documented human intrusions concentrate in short paths.

The four dimensions we assessed include:

1. Asset hops: a move onto a new machine or resource, such as a workstation reaching a file server.
2. Action hops: one discrete technique execution, and a single asset can host several in sequence: dumping a cached credential and then replaying it are two actions on one system.
3. Privilege hops: any transition that raises the attacker's effective identity tier, as when a standard user account is exchanged for a domain administrator credential.
4. Boundary hops: a crossing of a trust boundary, from one network segment into another or from the on-premises directory into the cloud tenant.

How one path is counted

A simplified crown-jewel path with an example of each hop type called out




One simplified path, scored under all four counting rules. Every counted hop is marked with a numbered circle. The color gives the counting rule, and the number is that hop's position within its count. Action-hop numbers match the step list above.



Machine-viable paths have significantly more depth

Four dimensions of attack path depth (median number of hops)



Technique domains per path

Median distinct ATT&CK-mapped domains



Paths spanning 4+ domains

Share of each population



Across these dimensions, machine-viable attack paths are at least twice as deep. They also differ in the breadth of tradecraft they require. The median machine-viable path spans five distinct technique domains, such as web exploitation, Active Directory abuse, and cloud IAM manipulation, and 64 percent span four or more. In documented human intrusions, operations of that breadth have typically required multiple specialists. An automated attack can execute all of them within a single operation, at no additional staffing or coordination cost.



Machine-viable attack paths persist longer in enterprise environments due to policy gaps


In our analysis, 55 percent of human-viable and 83 percent of machine-viable critical paths had been present in the environment for more than 30 days.

Machine-viable paths are disproportionately composed of findings that severity-based remediation policies rank below their thresholds: medium and low severity vulnerabilities, configuration weaknesses, and identity conditions that no SLA covers. Each component finding is individually and defensibly deprioritized. The paths those findings form inherit that deprioritization and remain standing.

Number of attack paths persisting more than 30 days


Share of critical paths present in the environment for more than 30 days

● HUMAN-VIABLE ● MACHINE-VIABLE



1 in 4 machine-viable paths is built entirely from findings below remediation thresholds

Machine-viable paths composed entirely of findings the organization's own remediation policy never requires fixing: sub-SLA severities, informationals, accepted configurations



of machine-viable paths contain no finding that current remediation policy requires fixing. Completing the remediation program as written does not affect them.

7% of critical paths were built by chaining low and medium severity vulnerabilities

Ranked individually, low and medium severity vulnerabilities fall below most remediation thresholds. Chained together, they produced 7 percent of the critical paths in our dataset. Severity-sorted remediation queues do not surface these combinations, because each component appears low risk in isolation.



The majority of attack paths span more domains than any single tool observes

Across our dataset, 64 percent of attack paths cross at least one domain boundary, moving between endpoint, identity, network, and cloud infrastructure.

Reconstructing the median critical path required correlated data from five security tools. This is a data aggregation problem faced by many enterprises with multiple point solutions covering different risk and asset types. The individual findings existed, in separate consoles, without the joins that make them a path.

Attack paths that are hybrid or cross-domain, crossing on prem, code, cloud, etc.

64%

Security tools whose data is required to draw the median critical path

5



Methodology

Cogent Research analyzed 43 million assets and 1.2 billion vulnerability, misconfiguration, and identity findings across the hybrid environments of over 50 Fortune 1000 organizations.

Attack path analysis drew on an average of 12 data sources per organization, spanning vulnerability scanners, endpoint detection and response, identity providers, cloud service providers, firewalls, and configuration management databases. Findings were normalized, deduplicated, and joined into a graph of assets, identities, and trust relationships for each environment.




An attack path is a traversable sequence of hops from an initial foothold to a crown jewel asset, with each hop validated against the environment's actual configuration, permissions, and controls. A path counts as new in the month its full kill chain becomes simultaneously present. Depth was scored under four counting rules: asset, action, privilege, and boundary hops.

A path is classified machine-viable when it exceeds at least one threshold of documented human intrusion economics covering depth, per-step yield, or technique breadth. Thresholds were calibrated against published red team engagement data and DFIR incident reporting. Paths inside every threshold are classified human-viable.

ABOUT

Cogent is vulnerability management at the speed of AI. Our AI agents find new vulnerabilities within minutes of disclosure, determine the best fix, and remediate at the level of autonomy you control.



cogent



To learn more, visit cogent.com