ultraviolet

THREAT ADVISORY

SHORT DWELL ATTACKS: WHEN SECONDS DECIDE THE OUTCOME

November 04, 2026 **TLP:GREEN**

UltraViolet Cyber TIDE Team
tide@uvcyber.com

FROM RESEARCH TO REALITY: AI-SUPPORTED RAPID-FIRE ATTACKS ARE ALREADY HERE

November 04, 2026 TLP:GREEN UltraViolet Cyber TIDE Team, tide@uvcyber.com

Executive Snapshot

The time between initial network access and meaningful attacker impact has collapsed dramatically. What once took days or weeks now often takes hours, and in some documented cases, seconds. The broader research community has recently documented agentic AI offensive tooling accelerating this trend, with autonomous frameworks running execution bursts in as little as two seconds, and separate research tracing a human-supervised AI framework directly to live ransomware operations. Neither case represents a fundamentally new attack technique; both represent old, familiar techniques executed at a pace defenders haven't had to plan around before.

This matters because detection and response programs built around multi-day investigation windows are increasingly out of step with how these intrusions actually unfold. The encouraging counterpoint is that speed carries a cost: both documented cases left behind distinctive, machine-paced behavioral signatures that remain detectable with the right focus and the right visibility.

Priorities for detecting and deterring short dwell attacks:

- Ensure firewall, VPN, and network device logs are being generated and retained consistently, since these often capture the earliest signs of credential abuse and lateral movement before an endpoint is ever touched.
- Maintain active, centralized collection of event logs from domain controllers and identity infrastructure, since this is frequently where the first indicators of compressed attacks appear.
- Review and tighten log retention on internal systems that sit outside typical endpoint coverage, since short dwell attacks leave little time to go looking for evidence after the fact.
- Keep internal asset and account inventories current, so that unusual activity on a given system or account can be recognized quickly rather than requiring lengthy investigation.
- Maintain close coordination with your security partner on alerting thresholds and escalation paths, since fast attacks require fast, well-practiced handoffs between internal teams and outside detection and response support.

TIDE Team Analysis

Attackers are increasingly optimizing for speed over thoroughness, and the incentive behind that shift is straightforward: the longer an intrusion lingers, the more chances defenders get to catch it. This isn't a gradual efficiency gain happening quietly in the background; it reflects a broader, observable move across the criminal ecosystem toward compressing every stage of an intrusion, from gaining a foothold to acting on it, into as little time as possible.

The most significant recent development pushing this trend further is the emergence of agentic AI offensive frameworks. ARTEX is currently the clearest, most thoroughly documented example. Google Threat Intelligence Group and Mandiant published joint research on ARTEX on October 30, 2026, describing an open-source, multi-agent, LLM-driven penetration testing and command-and-control platform observed in real intrusions since at least late August 2026. An operator hands the framework a target and a goal, and its agents independently plan and carry out the steps needed to get there, with little ongoing direction required. The framework has already surfaced in campaigns run by both financially motivated extortion crews and suspected state-linked actors.

The techniques involved aren't new; what's changed is the timeline. Individual execution bursts have been clocked at two to 25 seconds. One documented password-spray campaign ran with a median gap of exactly one second between login attempts, with blocked source addresses replaced within minutes. In one case, the framework hit a script error mid-exfiltration, trimmed the script, and relaunched it roughly ten seconds later, without any human stepping in. None of this is a new category of attack; it's old attacks running at a pace no human operator could sustain.

A second, independently documented case shows a different flavor of the same shift. Sophos's Counter Threat Unit investigated an AI-orchestrated framework originally built to test endpoint detection evasion, and traced that infrastructure directly to live Aur0ra ransomware operations run by a group tracked as GOLD OGRE. Here, the AI wasn't left to run on its own. Operators supplied credentials, network access, and objectives, then let an AI coding assistant handle the technical execution with comparatively light oversight. One especially telling detail: an operator redirected the AI away from repeatedly brute-forcing domain administrator accounts, instructing it to pursue quieter methods instead, a clear sign of a human actively coaching the AI toward stealth in real time. Two independent research firms corroborated the findings after separately accessing the same exposed attacker infrastructure, and the operation has already been tied to several confirmed victims.

Together, these cases sketch a spectrum rather than a single pattern. ARTEX sits at the fast, largely autonomous end: effective, but mechanically noisy in ways a cautious human wouldn't be. Aur0ra sits closer to the other end: a human still steering, but using AI to compress decisions that used to take hours. The reassuring thread connecting both is that neither is actually invisible. ARTEX leaves distinctive timing patterns and self-identifying artifacts because it trades stealth for velocity, and even the more deliberately stealth-tuned Aur0ra operations were still caught through conventional security research once investigators had the right visibility. The underlying techniques remain ones defenders have built detection content against for years. Faster and more automated doesn't yet mean undetectable.

Why It Matters

The compression of attack timelines challenges a core assumption built into many security programs: that there will be meaningful time between initial compromise and significant impact to detect, investigate, and respond. When that window shrinks to hours, or in documented cases, seconds, the quality of an organization's visibility into its own environment becomes the deciding factor in whether an intrusion gets caught early or discovered only after damage is done. This is precisely why the partnership between an organization and its security provider matters more now than it did a few years ago. Strong detection and response capability on the endpoint side only goes as far as the data available to it, and a well-rounded defense depends on both sides doing their part: a provider actively hunting across the telemetry it has visibility into, and an organization maintaining the logging, retention, and hygiene on its own infrastructure that makes that hunting effective.

This trend also reshapes the economics of attack. When the gap between buying stolen access and deploying ransomware against it collapses to 22 seconds, the entire criminal ecosystem around initial access brokerage becomes dramatically more efficient and scalable. Fewer things slow an attacker down, which means organizations and their security partners alike have less room for gaps in visibility, wherever those gaps happen to sit.

The AI dimension adds urgency, but not despair. Current research shows two different ways AI is already being operationalized in real attacks today, one largely autonomous and fast, one human-guided and more deliberate. Both are confirmed and actively occurring, not speculative future risks. At the same time, both cases demonstrate something genuinely reassuring: speed and stealth remain in tension with each other. Every documented example of AI-accelerated intrusion left behind detectable artifacts, whether distinctive timing patterns, self-identifying tool signatures, or exposed infrastructure that led researchers directly back to the operation. The same credential theft, lateral movement, and privilege escalation methods that security teams have built detection content against for years were simply executed faster, not replaced with something unrecognizable.

This is where a layered partnership earns its value. Endpoint-based detection and response can catch a great deal, but short dwell attacks often leave their earliest fingerprints in places endpoint tooling doesn't see: firewall logs, VPN authentication records, network device telemetry, and internal event logs. An organization that keeps that data flowing and retained gives its security partner a far better chance of catching an intrusion in its opening seconds rather than its aftermath. Put simply, speed favors whichever side closes its visibility gaps first, and that's a responsibility best carried together.

How to Respond

- Confirm firewall, VPN, and network device logs are actively collected, centralized, and retained long enough to support rapid investigation, since these sources often show the earliest signs of an intrusion before it ever reaches a monitored endpoint.
- Maintain consistent event log collection from domain controllers, authentication servers, and other identity infrastructure, since credential-based techniques like Kerberoasting and privilege escalation frequently show up here first.
- Keep internal network and asset inventories current and accurate, so unusual access to a given system or account can be recognized quickly during a fast-moving investigation.
- Establish and regularly test a clear escalation path between internal teams and your security provider, since short dwell attacks leave little room for delayed handoffs or unclear points of contact.
- Maintain tested, isolated backup and recovery capability as a first-line control, since short dwell attacks may leave little to no warning window before meaningful impact occurs.

What UltraViolet Cyber is Doing

- Proactively enabling custom detections based on the collected artifacts, tactics, techniques, and procedures identified in this activity. UltraViolet is also maximizing detection coverage proactively through Equinox.
- Tracking the latest threat actor intelligence, as well as the corresponding impact on industries around the world, through TIDE's Adversary Intelligence platform.
- Performing hypothesis driven threat hunts based on threat actor behavior and artifacts. UVCyber customers will be informed of the results through secure channels.
- Parsing available victim dump data for any social, financial, business, or technical relations to UVCyber Clients and partner organizations.
- Aggregating threat intelligence from myriad sources and applying the most up-to-date knowledge to proactive threat hunting and response.

About UltraViolet Cyber

UltraViolet Cyber is a leading tech-enabled managed security services provider, delivering unparalleled cybersecurity expertise that fills technology and talent gaps across Global 2000 and Federal Government customers. Founded and operated by security practitioners from the national intelligence community, UltraViolet Cyber connects offensive security, application security, detection and response, and security engineering to deliver a differentiated approach to cybersecurity operations. Transforming customers' security programs, UltraViolet Cyber's flagship security-as-a-service solution, UV Lens, removes complex operational silos, replacing them with integrated security capabilities. UltraViolet is headquartered in McLean, Virginia with technology centers across the world.

443.351.7630 / info@uvcyber.com | UltraViolet Cyber | @uv_cyber