ultraviolet

THREAT ADVISORY

The Dangers of Long Dwell Malware

Services Performed By:
UltraViolet Cyber TIDE Team
tide@uvcyber.com

Published Date:
September 23, 2026
TLP:GREEN

Executive Snapshot

Long dwell malware refers to tools built not for speed, but for stealth and deeper penetration of targets. Imagine code designed to sit undetected inside a network for months or, in extreme cases, years while quietly collecting intelligence, harvesting credentials, or preparing for a later, more damaging action. This category represents one of the more monumental shifts in the current threat landscape: as detection tooling has matured, sophisticated actors have responded by prioritizing stealth and legitimate-looking behavior over speed and destructiveness. The result is intrusions that frequently outlast an organization's log retention windows, meaning by the time an incident is discovered, the evidence needed to fully understand it may already have aged out.

Malware like BRICKSTORM, GOLDVEIN.JAVA, and Daxin (and the associated “Stupig”) are examples of long dwell attack vectors. Furthermore, tooling like Cobalt Strike, which has legitimate use cases for penetration testing, is commonly abused for quiet and malicious operations. As such, in this advisory, we will briefly illustrate the range of tradecraft seen in long dwell operations today, from repurposed commercial frameworks to purpose-built nation-state implants.

Recommended priorities for monitoring this class of threat:

- Consider extending log and network flow retention well beyond standard windows, since dwell times for this malware class routinely exceed a year, and the earliest evidence of access is often the first thing to age out.
- Treat network appliances, hypervisors, and virtualization management platforms as high-value assets requiring dedicated monitoring, since they often sit outside standard endpoint coverage entirely.
- Maintain independent, out-of-band visibility into appliance configuration and traffic, since devices used for long dwell persistence can't always be trusted to self-report accurately.
- Build historical log review into routine practice, not just incident response, since a single confirmed indicator may point to a compromise far older than initial evidence suggests.
- Recognize that legitimate administrative tools and frameworks can be repurposed for long-term persistence, and account for that risk in how those tools are governed.

TIDE Team Analysis

When we look toward stealthily conducted malware campaigns, long dwell malware comes into focus. This type of malware succeeds by minimizing the signals that traditional detection depends on, such as unusual files, obvious new connections, and abrupt behavioral changes. The four examples below span a spectrum from repurposed commercial tooling to bespoke nation-state implants, but all share this core design philosophy that aims to undermine current security strategies.

To illustrate how a legitimate, professionally engineered tool becomes a long dwell threat once it escapes its intended use case, we look at Cobalt Strike. Originally built for authorized red-team engagements, cracked and leaked versions of its Cobalt Strike's "Beacon" payload circulate widely, which gives actors of varying skill levels access to a mature toolkit for command and control, lateral movement, and credential access. Its "Malleable C2" feature lets operators disguise Beacon traffic as ordinary web activity, and, because Beacon can run reflectively in memory without writing to disk, it resists traditional file-based detection. Cobalt Strike held the top spot in Mandiant's most-frequently-observed malware rankings for five consecutive years, which is a testament to how effective disguised, legitimate-tool abuse remains a long-term persistence strategy.

On the other hand, malware like BRICKSTORM represents a more purpose-built approach, targeting infrastructure that most detection tooling simply cannot reach. According to a recent joint malware analysis report from CISA, the NSA, and Canada's Cyber Centre, this backdoor has been used by PRC-linked actors against VMware vSphere and Windows environments to achieve extended persistence, steal credentials, and maintain covert command and control. The malware layers multiple forms of encryption around its communications, uses DNS-over-HTTPS to help disguise C2 lookups, and can act as a SOCKS proxy to support further lateral movement once established. Because it runs on appliances that don't support conventional endpoint tooling, and because affected devices often have minimal onboard storage for forensic analysis, dwell times for BRICKSTORM-linked intrusions have reportedly stretched beyond a year in some cases (a length that often overwhelms typical log retention practices entirely).

A separate, long dwell mechanism. shows up in campaigns using fileless, in-memory execution paired with disguised network communication. An example of this is GOLDVEIN.JAVA, which is a Java-based downloader that beacons to a remote server using traffic crafted to resemble a normal encrypted handshake. It then reports execution results back to its operator while remaining hidden within otherwise unremarkable web traffic. Rather than persisting through a heavy footprint on disk, it runs entirely in memory, which lets it slip past detection tools built around file scanning. It topped Mandiant's list of most frequently observed malware families in a recent year of major incident-response investigations, which is the same year Cobalt Strike Beacon dropped out of the top spot it had held for five consecutive years. It has typically surfaced as part of broader extortion campaigns following exploitation of enterprise application vulnerabilities, used less as a standalone tool and more as an early-stage downloader that establishes a foothold before additional payloads or follow-on access are introduced

This leads to a discussion on Daxin and its recently discovered companion backdoor Stupig, which together represent one of the clearest illustrations of just how long "long dwell" can mean. Daxin avoids creating new, detectable, network connections altogether by opting to hijack legitimate existing traffic already flowing through a compromised server. This is a technique that defeats most detection logic built around spotting new or unusual connections. Stupig, which can be found running alongside Daxin, achieves execution before a user even logs in, leaving no standard logon event behind. Compile timestamps on both tools suggest the underlying intrusion in at least one recent case may have gone unnoticed for well over a decade, challenging the notion that long dwell operations also have fixed time limits.

Why It Matters

Long dwell malware inverts a core assumption many security programs are built on: that a compromise, once it happens, will be discovered within a reasonably bounded window. Standard practices, such as 90-day log retention, quarterly access reviews, and annual configuration audits were designed around that assumption. When dwell times regularly exceed a year, as seen with BRICKSTORM, or run into the range of over a decade, as suspected with Daxin, that assumption doesn't just weaken but instead fails outright. Stakeholders need to understand that a "clean" scan by itself or an absence of active alerts in a console is not the same as an absence of compromise.

The financial and operational stakes compound with time. The longer an actor sits inside an environment, the more they typically learn about it; conducting reconnaissance to determine which systems matter, where sensitive data lives, how the organization responds to anomalies, and more. With this knowledge, they can also prepare for a later, higher-impact action like mass data exfiltration, a destructive attack, or simply continued quiet intelligence collection. By the time discovery happens, the scope of what needs to be reviewed, and potentially disclosed, can be far larger than a typical incident.

This threat category further exposes a structural gap in how many organizations allocate security investments. Endpoint detection has matured significantly, but the infrastructure layer (e.g., network appliances, hypervisors, virtualization management platforms) often lacks equivalent insight. BRICKSTORM's success stems directly from targeting this blind spot. Organizations that have invested heavily in endpoint coverage while leaving infrastructure-layer monitoring thin are, in effect, leaving their most persistence-friendly systems the least watched.

Furthermore, there is a vendor and supply chain dimension worth flagging. Cobalt Strike's abuse shows that even legitimate, well-intentioned tools can become long-term threats once leaked or cracked versions spread. Any tool with broad legitimate adoption and strong evasion capability by design is a candidate for this kind of repurposing. What this means is vendor risk conversations should include how a tool's own capabilities could be weaponized against the organization using it, not just against its intended targets.

Finally, this category disproportionately affects the credibility of incident response itself. When forensic evidence has aged out of retention, or when the compromised system is an appliance that can't support standard IR tooling, "we don't have evidence of that" can mean either "it didn't happen" or "we can no longer tell." Stakeholders, boards, and regulators increasingly expect organizations to know the difference, but long dwell malware, by design, makes that distinction hard to reach.

How to Respond

- Extend log retention and network flow data collection well beyond 90 days, since confirmed dwell times for threats like BRICKSTORM regularly exceed a year, and standard retention windows are often too short to preserve evidence of the initial access point.
- Invest in dedicated monitoring and configuration visibility for infrastructure that sits outside typical endpoint coverage, such as network appliances and hypervisor management platforms, since this layer is increasingly the preferred target for long dwell operations.
- Maintain independent, out-of-band records of appliance configuration and traffic baselines, since compromised infrastructure can't always be relied upon to accurately report its own state.
- Treat any confirmed indicator of long dwell activity as reason to review historical logs and configuration records well beyond the immediate incident window, given how far these compromises can extend backward in time.
- Govern and monitor the use of legitimate administrative tools and frameworks, such as Cobalt Strike, since long dwell operations frequently rely on repurposed legitimate tooling to blend into normal activity.

What UltraViolet Cyber is Doing

- Tracking new CVEs and high impact vulnerabilities, analyzing and deploying public Proof-Of-Concept code against custom built targets.
- Proactively enabling custom detections based on the collected artifacts, tactics, techniques, and procedures identified in this activity.
- 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