ultraviolet

THREAT ADVISORY

TanStack Supply Chain Attack

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

Published Date:
May 20, 2026
TLP:GREEN

Executive Snapshot

On 11 May 2026, the threat actor group TeamPCP compromised 42 TanStack npm packages by chaining three GitHub Actions vulnerabilities to hijack the project's legitimate CI/CD pipeline. The attackers then published 84 malicious package versions carrying valid SLSA Build Level 3 provenance attestations, making them indistinguishable from legitimate releases by standard verification methods. The payload, a variant of TeamPCP's Mini Shai-Hulud credential-stealing worm, propagated autonomously by using harvested npm tokens to infect additional packages, ultimately reaching over 170 packages and 400 malicious versions across npm and PyPI within five hours. This attack is the latest escalation in a campaign that has compromised Trivy, Checkmarx KICS, LiteLLM, Bitwarden CLI, and numerous other trusted developer tools since March 2026. Additionally, TeamPCP's partnership with the Vect ransomware group created a direct pathway from supply chain credential theft to destructive ransomware operations. The open-sourcing of the Shai-Hulud framework has further expanded the threat by enabling copycat actors to fork and modify the toolchain independently. Enterprise organizations with any exposure to affected packages or to TeamPCP's earlier campaign waves should treat this as an active incident requiring immediate response.

- **Audit and rotate credentials immediately.** Inventory all development and production environments for compromised package versions published on or after 11 May 2026, and extend the review to packages affected by earlier TeamPCP waves including Trivy, KICS, LiteLLM, and Telnyx SDK. Treat any system that installed an affected version as compromised and rotate all accessible credentials, including cloud provider keys, SSH material, Kubernetes service accounts, CI/CD tokens, and package registry tokens.

- **Harden CI/CD pipeline trust boundaries.** Restrict or eliminate the use of pull_request_target workflows in GitHub Actions, enforce cache isolation between fork and base repository workflows, and separate publishing credentials from general-purpose build runners. Implement controls that prevent OIDC tokens from being extracted by untrusted code executing within the same runner environment.

- **Implement dependency consumption controls.** Pin dependencies by cryptographic hash rather than version, enforce minimum release age policies that quarantine newly published package versions before they enter build pipelines, and deploy real-time behavioral analysis tooling capable of detecting anomalous preinstall or postinstall scripts, unexpected binary downloads, and outbound connections to unknown infrastructure at install time.

- **Treat supply chain compromise as a ransomware precursor.** Validate that offline or air-gapped backups are current and provably isolated from production networks, given that TeamPCP-sourced credentials have already been used in confirmed Vect ransomware deployments where a fundamental encryption flaw renders affected files permanently unrecoverable. Ensure incident response playbooks explicitly address the scenario in which a supply chain compromise serves as the initial access vector for a follow-on extortion or data destruction campaign.

TIDE Team Analysis

On 11 May 2026, the threat actor group TeamPCP executed a coordinated supply chain attack against the TanStack open-source project, compromising 42 packages across the @tanstack npm namespace and publishing 84 malicious versions in a window of approximately six minutes. The attack leveraged a chained exploitation of three GitHub Actions vulnerabilities: a pull_request_target misconfiguration that granted fork-originated pull requests write access to the base repository's cache, GitHub Actions cache poisoning across the fork-to-base trust boundary, and runtime extraction of an OpenID Connect token directly from the GitHub Actions runner process memory. No npm credentials were stolen, and the legitimate publish workflow was not directly modified. Instead, the attacker hijacked TanStack's own CI/CD pipeline to publish credential-stealing payloads that carried valid SLSA Build Level 3 provenance attestations, making this the first documented npm supply chain attack to produce malicious packages indistinguishable from legitimate releases by provenance verification alone. The malicious versions were detected publicly within 20 to 26 minutes by an external researcher at StepSecurity, but the damage was already propagating.

The payload deployed through these compromised packages is a variant of TeamPCP's Mini Shai-Hulud toolchain, a self-propagating credential stealer that first appeared under the TeamPCP banner in late April 2026 during attacks against SAP npm packages. It is important to distinguish Mini Shai-Hulud from the original Shai-Hulud worm first identified by ReversingLabs on 15 September 2025. The original Shai-Hulud compromised hundreds of npm packages across two waves in September and November 2025, but its authors were never publicly identified. While TeamPCP adopted the Shai-Hulud branding and the campaigns share a conceptual model of worm-like propagation through stolen registry tokens, multiple research firms including OX Security and Wiz have noted that definitive attribution linking the original Shai-Hulud to TeamPCP has not been established. There are also meaningful technical differences: the original used TruffleHog for secret scanning and exfiltrated exclusively to GitHub repositories, while Mini Shai-Hulud uses dedicated command-and-control infrastructure alongside GitHub-based exfiltration and does not rely on TruffleHog. The thematic continuity appears to be deliberate branding rather than confirmed tooling lineage.

Upon installation, the Mini Shai-Hulud payload downloads the Bun JavaScript runtime as a living-off-the-land binary and executes a 2.3 MB obfuscated payload that systematically harvests CI/CD tokens, cloud provider credentials, Kubernetes service account tokens, HashiCorp Vault secrets, SSH keys, password vault data from applications such as 1Password, and package registry tokens. Stolen npm tokens are then used to publish additional malicious versions of any package the victim has write access to, enabling autonomous worm-like propagation through the npm ecosystem. The malware only executes on Linux systems, exits if it detects Russian language settings, and requires a minimum of four CPUs, indicating it is specifically designed to target CI/CD runners and cloud workloads rather than individual developer workstations. If the malware cannot achieve its objectives, it includes a destructive fallback mechanism capable of wiping files in the user's home directory, a feature inherited from the Shai-Hulud 2.0 wave in November 2025.

The blast radius extended well beyond TanStack. Within five hours of the initial compromise, TeamPCP published over 400 malicious versions across 172 distinct packages spanning both npm and PyPI. Confirmed victims include Mistral AI, UiPath, OpenSearch, and Guardrails AI. The @tanstack/react-router package alone accounts for approximately 12.7 million weekly downloads. OpenAI publicly disclosed that two corporate employee devices were impacted, resulting in unauthorized access to a limited subset of internal source code repositories and the exfiltration of credential material. OpenAI subsequently accelerated its macOS certificate rotation and deployed additional supply chain security controls, including minimum release age policies for package managers. However, organizations should note that Wiz's further analysis identified a bug in the payload delivered to the @uipath and @mistralai npm packages that renders the malware non-functional in those specific cases, which materially narrows the effective blast radius for consumers of those particular packages compared to the headline count of affected

namespaces.

This incident represents the latest escalation in a sustained campaign rather than an isolated event. TeamPCP emerged in late 2024 as a cloud-native exploitation crew targeting misconfigured Docker APIs, Kubernetes clusters, and Redis servers for cryptomining and credential theft. The group escalated dramatically in March 2026 with a cascading supply chain campaign that compromised Aqua Security's Trivy vulnerability scanner, Checkmarx KICS, the LiteLLM AI gateway, and the Telnyx Python SDK, crossing five software ecosystems including GitHub Actions, Docker Hub, npm, PyPI, and OpenVSX. Each wave has built on the technical sophistication of the prior one. The May 2026 TanStack wave introduced a capability that fundamentally undermines existing software supply chain verification: the ability to publish malicious packages with valid, machine-verifiable provenance attestations. Organizations that rely on SLSA provenance as a trust signal must now recognize that provenance verification is a necessary but no longer sufficient control. Following the TanStack attack, TeamPCP open-sourced what appears to be the complete Shai-Hulud offensive framework on GitHub, and independent threat actors have already begun forking and modifying the code, broadening the threat surface beyond TeamPCP itself.

The strategic risk to enterprise organizations is compounded by TeamPCP's partnership with the Vect ransomware-as-a-service group, though this dimension requires careful assessment. In late March 2026, Vect announced a formal partnership with TeamPCP on BreachForums, explicitly stating the intent to deploy ransomware across organizations affected by TeamPCP's supply chain compromises. Vect also operationalized a mass affiliate mobilization program, offering affiliate keys to BreachForums' approximately 300,000 registered users. At least one confirmed Vect ransomware deployment using TeamPCP-sourced credentials has been reported, and on 15 April 2026 Vect listed its first victim with data attributed to the TeamPCP Trivy campaign. However, Vect's leak site listed only two victims as of late April 2026, and one of those listings remains unconfirmed by third parties. More critically, Check Point Research disclosed that Vect 2.0 contains a fundamental encryption implementation flaw that discards decryption nonces for all file segments except the last, rendering any file larger than 128 KB permanently unrecoverable by both victim and attacker. JUMPSEC's independent analysis concluded that TeamPCP can get ransomware through the door, but the monetization pathway is broken, effectively making Vect a destructive wiper masquerading as extortion. The data destruction risk from a Vect deployment is real and arguably worse than functioning ransomware, but the "industrialized ransomware pipeline" narrative should be tempered by the operational immaturity of the extortion side of this partnership.

For enterprise security leaders, this incident demands a reassessment of software supply chain risk posture across several dimensions. Organizations should audit all development and production environments for the presence of compromised package versions published on or after 11 May 2026, treating any system that installed an affected version as potentially compromised regardless of whether active exploitation was observed. Credential rotation must be treated as an immediate operational priority, covering not only npm and PyPI tokens but also cloud provider credentials, SSH keys, Kubernetes service accounts, and any secrets accessible from CI/CD runners. Beyond immediate triage, organizations should implement controls such as dependency pinning by hash rather than version, minimum release age policies that prevent automatic adoption of newly published versions, and network segmentation of CI/CD infrastructure to limit lateral movement from compromised build environments. Organizations that incorporated Trivy, Checkmarx KICS, LiteLLM, or Telnyx SDK into their pipelines during the March 2026 window should assume compromise and rotate all credentials in those environments regardless of whether they were directly affected by the May TanStack wave.

The TanStack compromise exposes systemic weaknesses in the trust model that underpins modern software delivery. The attack did not require phishing a maintainer, stealing a password, or exploiting a code vulnerability in the target application. It exploited the implicit trust relationships between GitHub Actions workflows, shared caching infrastructure, and OIDC-based publishing mechanisms, which are architectural trust boundaries that most organizations neither monitor nor control. Security teams should evaluate their exposure to GitHub Actions cache

poisoning, restrict the use of pull_request_target workflows, enforce artifact integrity verification beyond provenance attestation alone, and consider isolating CI/CD publishing credentials from general-purpose build runners. The pace of TeamPCP's operations shows no indication of slowing, copycat actors are already deploying modified versions of the open-sourced Shai-Hulud framework, and the group's estimated 300 GB trove of harvested credentials could enable future compromises at any time. The threat model has shifted from isolated package compromise to identity-driven propagation through trusted infrastructure, and the window between initial supply chain compromise and downstream exploitation continues to narrow.

Why It Matters

The TanStack compromise is not an anomaly. It is the latest and most technically sophisticated entry in a pattern of software supply chain attacks that has been escalating in frequency, scale, and consequence for half a decade. The 2020 SolarWinds intrusion demonstrated the model at its most devastating: attackers inserted a backdoor into the Orion platform's legitimate update mechanism, and approximately 18,000 organizations, including U.S. government agencies, Cisco, Deloitte, Intel, and Microsoft, installed the compromised build over a period of months before the intrusion was discovered. The average cost to each affected organization was estimated at $12 million. The 2021 Kaseya VSA compromise showed how a single vulnerability in a managed service provider's remote administration tool could cascade into ransomware deployment across 800 to 1,500 downstream businesses, forcing a major Swedish grocery chain to close hundreds of storefronts and disrupting schools, pharmacies, and government services across multiple countries. In both cases, the attackers exploited the implicit trust that organizations place in their vendors' software delivery pipelines, turning that trust into the primary attack vector. The TanStack incident follows this same structural logic but applies it to the open-source dependency layer that now underpins virtually all modern application development, and it does so with the added capability of producing malicious packages that pass provenance verification, a control that did not exist during SolarWinds or Kaseya but has since been promoted as a core defense.

The scale of the open-source attack surface compounds the risk considerably. Sonatype's 2026 State of the Software Supply Chain report identified more than 454,000 new malicious packages published across npm, PyPI, Maven, NuGet, and Hugging Face in 2025 alone, bringing the cumulative total of known malware to over 1.2 million packages. ReversingLabs recorded a 73 percent year-over-year increase in malicious open-source packages from 2024 to 2025, the steepest annual rise on record. Over 99 percent of detected open-source malware in 2025 targeted the npm ecosystem specifically, and the average npm project pulls in 79 transitive dependencies, meaning a single compromised package can propagate through entire application stacks within hours. The September 2025 Shai-Hulud worm proved that self-replicating malware could spread autonomously through the npm registry by harvesting maintainer tokens and republishing infected versions of every package those maintainers controlled. The November 2025 second wave compromised 795 packages with a combined download count exceeding 100 million. State-linked actors have joined financially motivated groups in exploiting this surface: Sonatype identified more than 800 Lazarus Group-associated malicious packages in 2025, concentrated overwhelmingly in npm, and in March 2026 the axios npm package compromise was attributed by Google to UNC1069, a financially motivated North Korean threat actor. The open-source dependency layer is no longer an incidental risk. It is the primary battleground for both criminal and state-sponsored supply chain operations.

What makes the current threat landscape qualitatively different from earlier supply chain incidents is the convergence of worm-like propagation, CI/CD trust exploitation, and direct monetization through ransomware partnerships. SolarWinds and Kaseya were each singular events carried out by well-resourced actors against specific vendor platforms. The TeamPCP campaign model, by contrast, is repeatable, scalable, and increasingly commoditized. The group has compromised targets across five software ecosystems in a matter of weeks, the Shai-

Hulud offensive framework has been open-sourced for anyone to fork and deploy, and the partnership with the Vect ransomware group creates a direct pipeline from credential theft to extortion or data destruction. Mandiant's M-Trends 2026 report places third-party compromise as the second most common cloud-compromise initial access vector at 17 percent, and Palo Alto Networks' Unit 42 has characterized the post-September 2025 period as a fundamental inflection point where npm supply chain attacks evolved from isolated nuisance incidents into systematic campaigns that weaponize the trust powering modern software development. For Security Leadership, the implication is that software supply chain security can no longer be treated as a developer hygiene problem or a compliance checkbox. It is a first-order enterprise risk that sits at the intersection of application security, CI/CD infrastructure defense, identity management, and ransomware preparedness, and the organizations that survive the next incident will be the ones that recognized that convergence before it arrived at their door.

How to Respond

- Strictly adhere to cybersecurity Fundamentals and ensure all personnel undergo annual phishing and social engineering training. Speak with your UltraViolet Cyber TAM Representative to schedule a live phishing engagement.
- Inventory and assess NPM, Bun, and TanStack usage throughout your enterprise infrastructure.
- Perform annual tech refresh reviews to gain a holistic understanding of your infrastructure. Speak with your UltraViolet Cyber TAM Representative to schedule a RedTeam or PurpleTeam engagement to gain insight into the vulnerabilities in your environment.

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

---