ZEROFOX INTELLIGENCE

Brief

TeamPCP Resurfaces to Sell Old Data: Assessing the Impact

B-2026-07-28a

Classification: TLP:CLEAR
Criticality: Medium
Intelligence Requirements: TeamPCP, Software Supply Chain Attacks, Self-Propagating Worm

July 28, 2026

Scope Note

ZeroFox Intelligence is derived from a variety of sources, including—but not limited to—curated open-source accesses, vetted social media, proprietary data sources, and direct access to threat actors and groups through covert communication channels. Information relied upon to complete any report cannot always be independently verified. As such, ZeroFox applies rigorous analytic standards and tradecraft in accordance with best practices and includes caveat language and source citations to clearly identify the veracity of our Intelligence reporting and substantiate our assessments and recommendations. All sources used in this particular Intelligence product were identified prior to 12:00 AM (EDT) on July 22, 2026; per cyber hygiene best practices, caution is advised when clicking on any third-party links.

| Brief | TeamPCP Resurfaces to Sell Old Data: Assessing the Impact

| Key Findings

- TeamPCP has resurfaced, likely in an effort to repurpose previously harvested datasets to identify and pursue further intrusions; the group is less likely to conduct a net new intrusion campaign and is likely currently focused on disclosing older breaches.
- On July 21, 2026, ZeroFox observed new dark web activity potentially linked to threat group TeamPCP, following its relative operational silence since at least May 2026. Xploitrs (also spelled xpl0itrs), the group's alleged associate, claimed to have leaked data tied to RapidFort, a U.S.-based cybersecurity company.
- TeamPCP's attacks can be typically categorized into two stages: credential acquisition and credential operationalization. The majority of the data points stolen during TeamPCP's credential acquisition stage lose value once the credentials are rotated, revoked, or otherwise neutralized, making the loss severe but neutralizable.
- The real and lasting impact of TeamPCP's campaign almost certainly comes from the group's credential operationalization activities.

Introduction

ZeroFox has observed a potential resurfacing of the TeamPCP threat group, following its relative operational silence since at least May 2026. However, the group is unlikely to be conducting new attacks at the moment; instead, it is likely currently focused on disclosing and repurposing older breaches. On July 21, 2026, Xploitrs (also spelled xpl0itrs), an alleged associate of TeamPCP, advertised data allegedly tied to RapidFort, a U.S.-based cybersecurity company on dark web forums Spear and PwnForums. Xploitrs claimed that the breach was part of the "CanisterWorm" campaign with TeamPCP.[1] Based on TeamPCP's reputation, the data breach is likely legitimate. Xploitrs said the data dates back to March 2026, indicating the sale of previously harvested access rather than a fresh compromise.

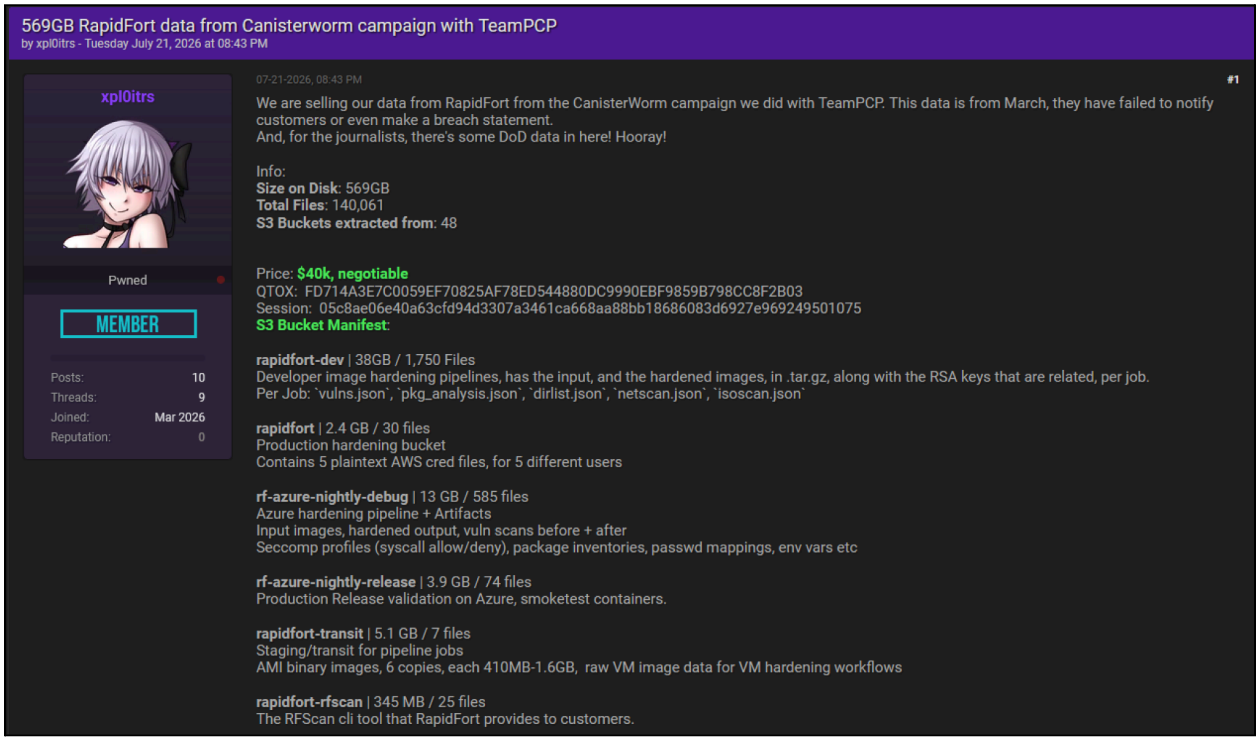

Xploitrs selling data allegedly tied to RapidFort on PwnForums
Source: ZeroFox Intelligence

TeamPCP's core relationships and tradecraft appear intact: Xploitrs has allegedly remained a consistent partner across several strained or unsuccessful collaborations,

[1] https://cloud.zerofox.com/intelligence/advanced_dark_web/107700

and TeamPCP is likely seeking additional collaborators or rebuilding infrastructure to obscure the now widely documented tactics, techniques, and procedures (TTPs) and indicators of compromise (IoCs) tied to its campaign. Whether the group is preparing a successor campaign remains unconfirmed. ZeroFox assesses that TeamPCP is likely still validating and operationalizing the credentials harvested during the initial wave of compromises to identify access that remains active. The group is also likely sharing or selling other datasets within the cybercriminal ecosystem.

TeamPCP dominated the software supply chain threat landscape through the first half of 2026, conducting credential theft at scale through its self-propagating CanisterWorm malware. That operational tempo came to a near standstill around May 2026. Since then, the group's observable footprint has been confined to a single post on the dark web forum Breached on June 13, 2026, and the disclosure of the RapidFort leak on July 21, rather than any new infections.

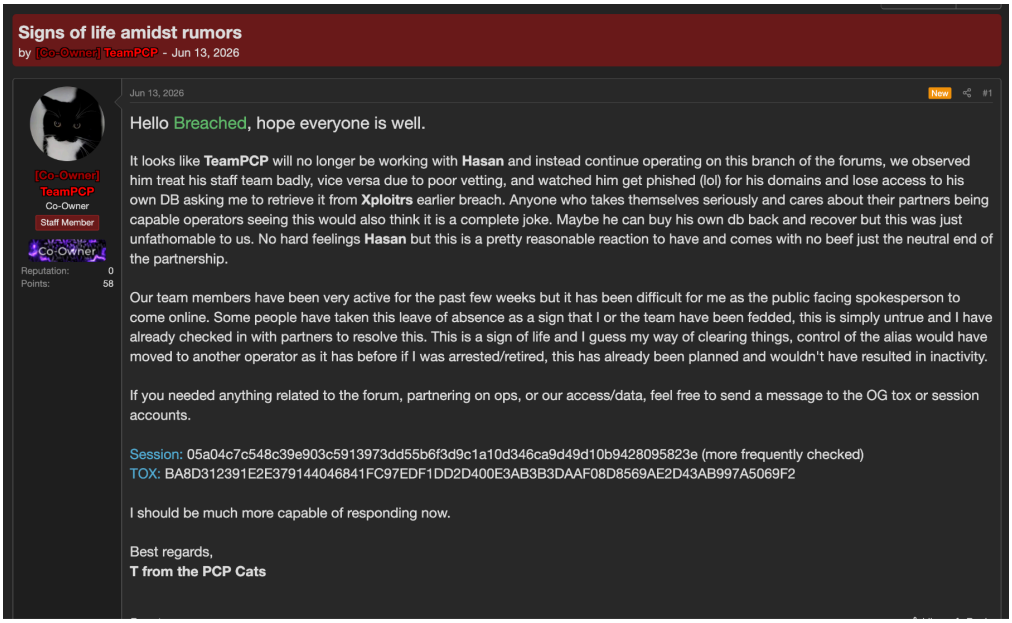

TeamPCP's post on Breached
Source: ZeroFox Intelligence

- In the June post, TeamPCP announced a split from prominent threat actor HasanBroker following alleged operational disputes and claimed that the group has been active for weeks, in response to rumors of arrests.

This brief examines the impact of the credentials TeamPCP has already harvested and offers ZeroFox's assessment of the group's likely next moves as it adapts its initial-access tradecraft.

Stages of TeamPCP Attacks

TeamPCP's attacks can typically be categorized into two stages:

- **Credential acquisition:** Supply chain compromise of trusted open-source continuous integration/continuous delivery (CI/CD) environments, as well as developer and security tooling projects to obtain credentials and initial access
- **Credential operationalization:** Use of those stolen credentials to access cloud infrastructure and organizational CI/CD and production environments, enabling additional credential harvesting, lateral movement, and data exfiltration

TeamPCP (tracked under multiple aliases, including PCPcat and ShellForce) reportedly began activity in 2024 as a cloud-native cryptomining worm. Researchers first observed TeamPCP as a distinct cluster in late November 2025, and the group allegedly compromised roughly 59,000 Next.js servers by December.[2][3] Beginning in late February and early March 2026, the group used its compromise-harvest-propagate model on developer and CI/CD credentials. Within weeks, it had poisoned trusted open-source tooling across GitHub Actions, npm, PyPI, Docker, and OpenVSX (including that of Trivy, Checkmarx KICS, LiteLLM, and Telnyx).[4]

- The group reportedly reused the stolen credentials for breaches at Cisco, Mistral AI, and a GitHub internal haul with roughly 4,000 repositories. By mid-May, TeamPCP reportedly open-sourced Mini Shai-Hulud, a variant of its original self-propagating, cross-ecosystem (npm + PyPI) credential-stealing worm, and triggered a copycat wave. Mini Shai-Hulud gave rise to copycat derivatives, such as Miasma, IronWorm, and Megalodon.

[2] hXXps://thehackernews[.]com/2026/02/teampcp-worm-exploits-cloud.html
[3] hXXps://cybersecuritynews[.]com/pcpcat-hacked-next-js-servers/
[4] hXXps://www.ic3[.]gov/CSA/2026/260702.pdf

Economy of Stolen Data

The data stolen during TeamPCP's credential acquisition stage loses value once the credentials are rotated, revoked, or otherwise neutralized, offering threat actors a limited window for abuse. The real and lasting impact almost certainly stems from what TeamPCP does with that access during operationalization: exfiltrating data that cannot be "unleaked."

TeamPCP rapidly weaponizes stolen credentials to enumerate repositories, modify packages, and propagate malicious releases before remediation occurs.

- The abuse window for TeamPCP-related intrusions in certain incidents spanned minutes. For example, malicious PyTorch Lightning packages remained available for approximately 42 minutes before removal, and trojanized Nx Console extensions were remediated within roughly 11 minutes.[5][6] On the other hand, the Trivy GitHub Action compromise reportedly remained active for approximately 12 hours before containment.[7]

TeamPCP and its associates most frequently stole organizations' identity/access credentials, followed by cloud access keys during its supply chain phase (roughly February–May 2026). Notably, the two most frequently stolen categories are also the most readily neutralized once rotated.

[5] hXXps://thehackernews[.]com/2026/04/pytorch-lightning-compromised-in-pypi.html
[6] hXXps://nx[.]dev/blog/nx-console-v18-95-0-postmortem
[7] hXXps://github[.]com/aquasecurity/trivy/discussions/10425

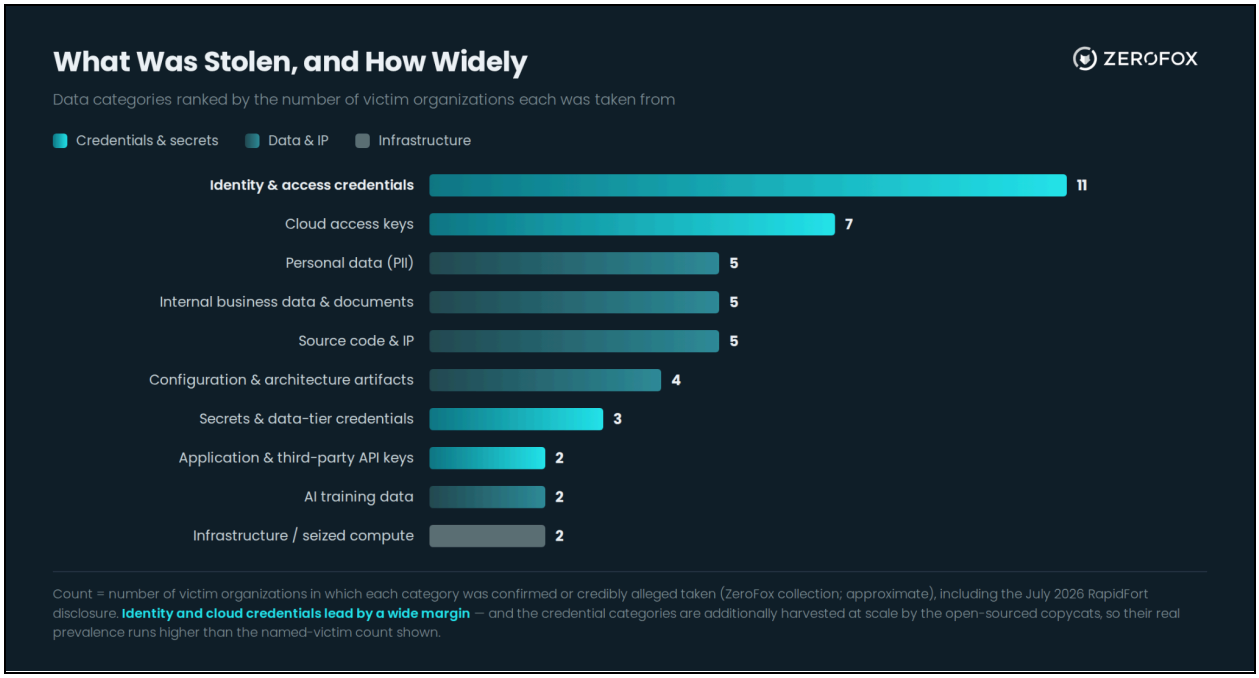

Categories of data most frequently stolen by TeamPCP

Source: ZeroFox Intelligence

ZeroFox assesses that the exfiltration of data—such as personally identifiable information (PII), source code, and other sensitive information—during credential operationalization activities is very likely to pose significant long-term risk, as it cannot be unleaked once exposed.

- Exposed source code and proprietary data, internal business processes, product roadmaps, or sensitive operational details are very likely to be sold on the dark web to other threat actors or competitors for corporate espionage.
- Stolen source code is likely to facilitate reconnaissance and downstream vulnerability discovery. Recent advances in artificial intelligence (AI)-assisted code analysis are likely lowering the technical barriers to identifying exploitable weaknesses, making source code exposure increasingly valuable to a broader range of threat actors.
- Similarly, the exposure of customer and employee PII is very likely to create long-term privacy, regulatory, and reputational consequences. AI training datasets are likely to contain proprietary information, unique organizational

knowledge, or sensitive records that lose their exclusivity once obtained by unauthorized parties.

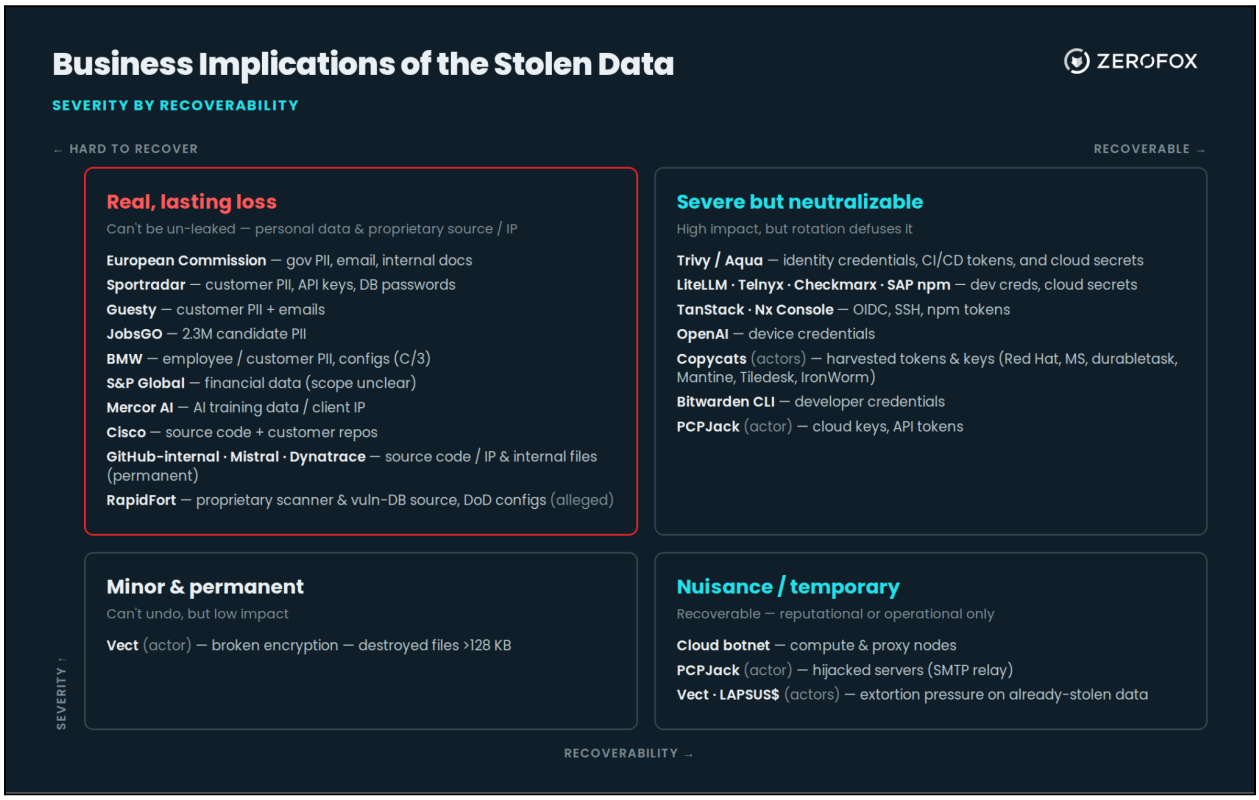

Implications of stolen data
Source: ZeroFox Intelligence

The theft of configuration and architecture artifacts very likely serves as an intelligence and reconnaissance enabler, though cross-account identity and access management (IAM)/AssumeRole templates can provide near-direct access rather than a direct source of compromise. Such information can reveal internal network structures, deployment patterns, technology stacks, security controls, and operational dependencies that are likely to be useful for future intrusion activity. This is illustrated by the RapidFort leak, which exposed Azure Kubernetes Service (AKS) kubeconfigs, Terraform states, and CI/CD platform configurations as well as plaintext cloud credentials.

- However, infrastructure changes, architectural updates, and routine software development practices are very likely to blunt the long-term impact.

- Meanwhile, compromised infrastructure and compute resources are likely to be leveraged to conceal cybercriminal infrastructure. Rotating affected hosts and re-establishing trusted infrastructure are very likely to mitigate malicious use of infrastructure.

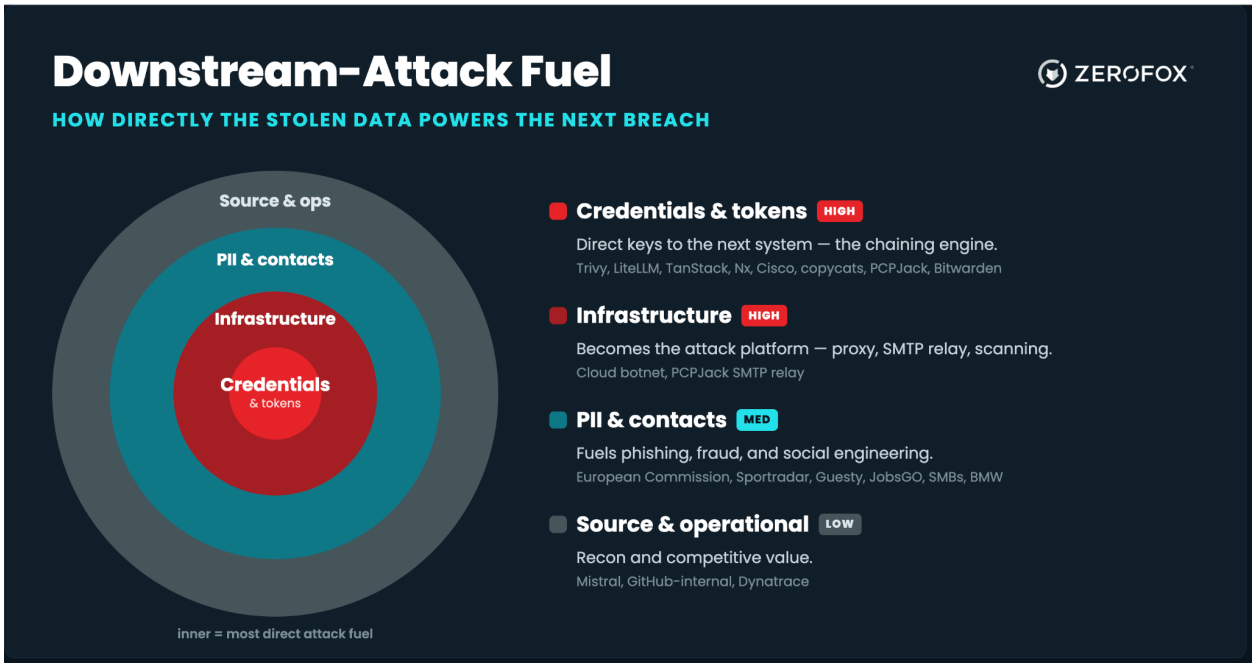

Downstream impact of stolen data
Source: ZeroFox Intelligence

What Is Next?

ZeroFox assesses that the July 2026 data breach post by TeamPCP's affiliates likely indicates ongoing operational intent. Xploitrs had earlier publicly claimed future victims via their X account, mirroring the advance claims made prior to the Dynatrace compromise. The victims named included Spotify, the U.S. Department of the Treasury, OpenAI's ChatGPT, and Trustpilot.[8]

- Additionally, TeamPCP is very likely interested in combining supply chain compromise and credential theft with ransomware deployment, as indicated by its allegedly unsuccessful association with Vect ransomware group. Added

[8] hXXps://x[.]com/xploitrsturtle2/status/2067447055832817944

ransomware capabilities are very likely to significantly increase the operational and financial impact on affected organizations through data-for-ransom extortion or disruptive encryption-based attacks.

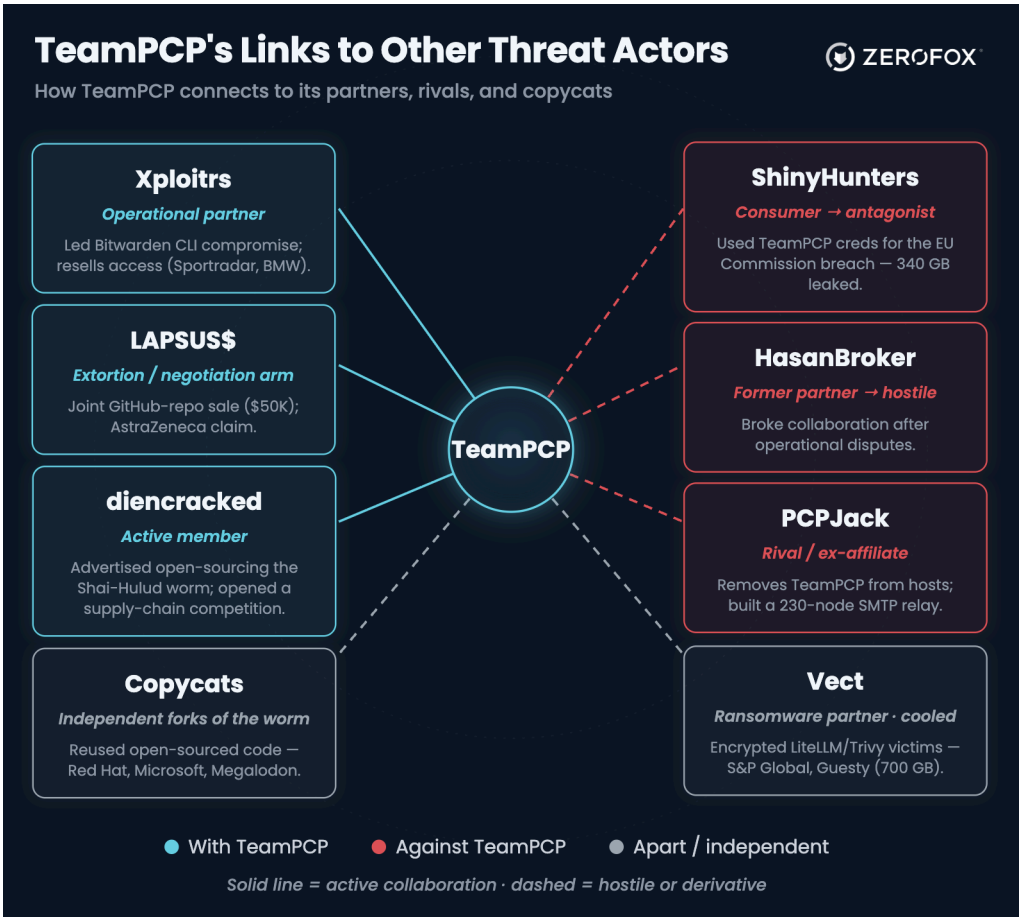

TeamPCP's alleged relationships with other threat actors
Source: ZeroFox Intelligence

TeamPCP is likely to move to next initial access vectors for its future campaigns, given npm's announcement of a v12 release in July 2026—which would lock down install-time script execution (unless explicitly allowed) to prevent supply chain attacks.[9]

- In the event of new campaigns, the group is very likely to increasingly target PyPI and similar ecosystems (such as RubyGems and CRAN)—where package-installation and import-time execution mechanisms remain

[9] hXXps://github[.]blog/changelog/2026-06-09-upcoming-breaking-changes-for-npm-v12/

comparatively less constrained—to distribute next-generation worms. ZeroFox assesses that the likely IoCs would be unusual .pth hooks or startup scripts in high-impact packages.

- Groups conducting TeamPCP-derived copycat activities are likely to continue pursuing credential acquisition as their primary operational objective. While earlier campaigns relied heavily on compromised packages and open-source tooling, more recent activities such as Miasma demonstrate a shift toward repositories, AI-assisted coding tools, integrated development environment (IDE) configurations, and other trusted developer workflows.

Conclusion

ZeroFox assesses that the primary risk from TeamPCP now likely stems from the continued operationalization of stolen access, data, and tooling, which can facilitate the exposure of source code, intellectual property, and other sensitive organizational information. As npm v12 introduces restrictions on install-time script execution, groups conducting TeamPCP-derived copycat activities are likely to increasingly target alternative ecosystems and trusted developer workflows, including PyPI, AI-assisted coding tools, repositories, and development environments. The potential financial impact on victim organizations is also likely to increase significantly if TeamPCP succeeds in incorporating ransomware capabilities into future campaigns.

- Read the complete TeamPCP threat actor profile **here**.
- The **Federal Bureau of Investigation (FBI) has also published** the IoCs and TTPs related to TeamPCP.

Appendix A: Traffic Light Protocol for Information Dissemination

	Red	Amber
WHEN SHOULD IT BE USED?	**Sources may use** TLP:RED when information cannot be effectively acted upon by additional parties and could lead to impacts on a party's privacy, reputation, or operations if misused.	**Sources may use** TLP:AMBER when information requires support to be effectively acted upon but carries risks to privacy, reputation, or operations if shared outside of the organizations involved.
HOW MAY IT BE SHARED?	**Recipients may NOT share** TLP:RED with any parties outside of the specific exchange, meeting, or conversation in which it is originally disclosed.	**Recipients may ONLY share** TLP:AMBER information with members of their own organization and its clients, but only on a need-to-know basis to protect their organization and its clients and prevent further harm. **Note that** TLP:AMBER+STRICT restricts sharing to the organization only.

	Green	Clear
WHEN SHOULD IT BE USED?	**Sources may use** TLP:GREEN when information is useful for the awareness of all participating organizations, as well as with peers within the broader community or sector.	**Sources may use** TLP:CLEAR when information carries minimal or no risk of misuse in accordance with applicable rules and procedures for public release.
HOW MAY IT BE SHARED?	**Recipients may share** TLP:GREEN information with peers and partner organizations within their sector or community but not via publicly accessible channels.	**Recipients may share** TLP:CLEAR information without restriction, subject to copyright controls.

Appendix B: ZeroFox Intelligence Probability Scale

All ZeroFox intelligence products leverage probabilistic assessment language in analytic judgments. Qualitative statements used in these judgments refer to associated probability ranges, which state the likelihood of occurrence of an event or development. Ranges are used to avoid a false impression of accuracy. This scale is a standard that aligns with how readers should interpret such terms.

Almost No Chance	Very Unlikely	Unlikely	Roughly Even Chance	Likely	Very Likely	Almost Certain
1-5%	5-20%	20-45%	45-55%	55-80%	80-95%	95-99%

Appendix C: ZeroFox Intelligence Threat Actor Reputation Scale

Untested	Moderately Credible	Well-regarded	Prominent
Has garnered no reputation; credibility cannot be determined.	Has made up to 10 transactions; has been active on forum for at least three months.	Has at least 10 transactions; has been active on forum for three months to one year.	One of the most well-known and credible threat actors on the site; long-term, established presence on the forum of more than one year.