THREAT REPORT: INDIRECT PROMPT INJECTION TARGETS AI AGENTS IN TECHNOLOGY AND FINANCE SECTORS

Visual Summary

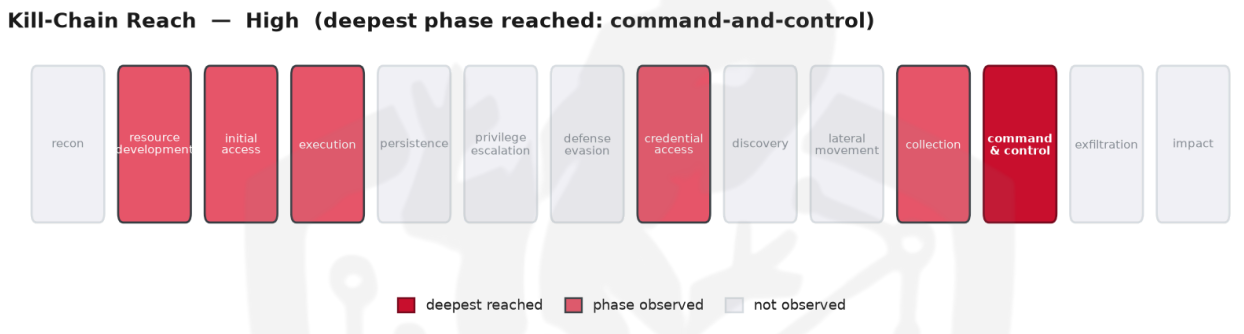

How far down the attack chain this campaign reached; deepest phase in crimson.

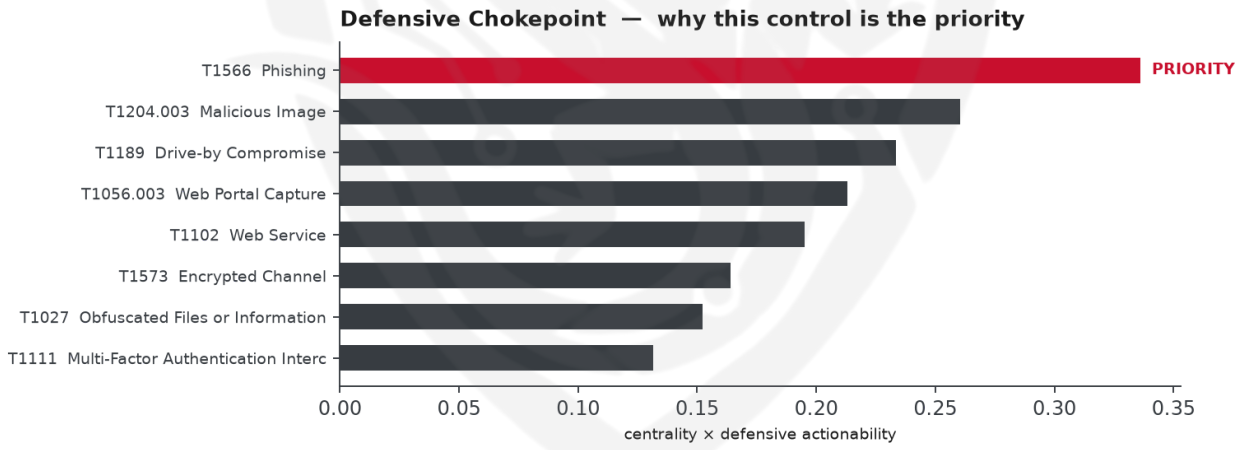

The control that most disrupts the attack (priority in crimson).

Summary

The observed cluster of activity targets the Technology and Finance sectors, utilizing various techniques to compromise systems. The priority defensive action should be taken to prevent further attacks, focusing on enabling attachment sandboxing and link rewriting. The threat actor's identity remains unknown, and the malware families and central CVE used in the attack are insufficiently understood. Priority should be given to defending against phishing attacks, which have been identified as a key technique used by the threat actor.

Threat Overview

The threat actor, whose identity is unknown, is using a range of techniques, including phishing, spearphishing, and drive-by compromise, to target organizations in the Technology and Finance sectors. The attack chain involves various techniques, but the central CVE and malware families used are not well understood. The victimology suggests that the threat actor is targeting specific sectors, but the targeted country and other details are insufficiently understood.

Attack Chain and Priority Control

The observed techniques used by the threat actor form a complex attack chain, involving techniques such as match legitimate resource name or location, spearphishing link, and deobfuscate/decode files or information. The priority technique is T1566 Phishing, which is the graph chokepoint. To defend against this threat, the priority control is to: Enable attachment sandboxing and link rewriting; block macro-enabled Office docs from the internet zone; deploy a user report-phish button.

Infrastructure and Corroboration

There is no available information on the hosting providers or ASNs observed, and no malware families have been corroborated by abuse.ch. Additionally, there are no indicators with an AbuseIPDB confidence of 80% or higher, with the highest confidence seen being 0%.

Technical Appendix

Ground-truth telemetry from the source pulse; analytical scores are secondary and clearly labelled.

Observed Techniques (ATT&CK, expert-tagged by source)

- T1036.005 Match Legitimate Resource Name or Location
- T1566.002 Spearphishing Link
- T1140 Deobfuscate/Decode Files or Information
- T1608.001 Upload Malware
- T1583.001 Domains
- T1056.003 Web Portal Capture
- T1102 Web Service
- T1608.005 Link Target
- T1566 Phishing
- T1204.003 Malicious Image
- T1027 Obfuscated Files or Information
- T1573 Encrypted Channel
- T1189 Drive-by Compromise
- T1111 Multi-Factor Authentication Interception
- T1071.001 Web Protocols
- T1204.001 Malicious Link

Analytical Scores

- Priority technique (graph chokepoint): T1566 Phishing (centrality 0.3618).
- Operational risk: High (score 0.636, reaches command-and-control).

Behavioral Signature Cluster

- Method: technique-overlap cosine vs 170 MITRE ATT&CK Groups (top overlap 0.5025; reference threshold tau = 0.6).
- These are behavioural resemblances for defensive playbook cross-referencing, not an identity assessment. Where the source pulse names an actor, that attribution is authoritative; the overlaps below neither confirm nor contest it.

Nearest historical profiles by behavioural overlap (resemblance only):

Historical Profile	Behavioural Overlap
Mustard Tempest	0.5025
Transparent Tribe	0.489
LazyScripter	0.4671
WIRTE	0.4606
BITTER	0.4599

Enrichment Signal

- Highest AbuseIPDB confidence among IOCs: 0%.

Indicators of Compromise (defanged — non-clickable)

The network and file indicators observed in this activity are listed below for blocklisting:

Type	Indicator
Domain	runners-daily-blog[.]com
Domain	bistro-reserve-now[.]net
Domain	consensus-protocol-v4[.]org
Domain	debank[.]auction
Domain	digital-asset-mart[.]org
Domain	edge-compliance-node[.]org
Domain	identity-breach-response[.]org

Type	Indicator
Domain	market-insight-global[.]com
Domain	py-lib-repository[.]dev
Domain	visual-media-rights-group[.]org
Domain	permits[.]global-transit-authority[.]org

Flame Cell // Adversary Pivot [BETA]

BETA. A hypothetical tabletop projection, anchored to documented ATT&CK behaviour and technique co-occurrence across 170 tracked groups. It is not a prediction; treat it as a war-game prompt and review before acting. Spot a pivot that looks wrong? norbert@salacti.com

- **Control applied:** T1566 Phishing (initial-access)
- **Observed here:** T1566 Phishing was the initial-access control point; blocking a means-of-entry rarely stops a determined actor, so the pivot projects forward to the execution stage they must still reach.
- **Projected pivot:** T1203 Exploitation for Client Execution (execution)
- **Basis:** of the 38+ groups that use Phishing, this pairing runs 1.6x above base rate, a disproportionately-coupled move rather than the obvious one.

The actor could pivot to T1203 Exploitation for Client Execution in an attempt to execute malicious code on the client-side, potentially by exploiting vulnerabilities in commonly used software, which may allow them to bypass the blocked phishing attempt and still achieve their objective. This technique may be particularly appealing as it could enable the actor to leverage existing vulnerabilities in client applications, such as Office or browsers, to gain execution and proceed with their campaign. To counter this potential pivot, the mandated defensive control of patching client applications, enabling exploit mitigations like ASLR, DEP, and CFG, and implementing application isolation would likely be effective.

Also plausible (same tactic): T1053 Scheduled Task/Job (n=43, lift 1.3), T1106 Native API (n=18, lift 1.6)

Detection and hardening for the pivot (T1203 Exploitation for Client Execution)

- **Telemetry:** EDR process + memory telemetry; Office/browser/reader child-process telemetry; Windows Error Reporting (application crashes)
- **Detect:**
  - Office, browser or PDF-reader processes (winword/excel/outlook/acrobat/chrome) spawning shells or LOLBINs
  - Sysmon 1 anomalous children of document/browser apps; unbacked-memory execution right after a client crash
  - WER application-crash events for client apps immediately followed by a new child process
- **Harden:**
  - Patch client software fast (browsers, Office, PDF, Java); retire end-of-life plugins

    - ASR rules blocking Office child processes; exploit protection (DEP/ASLR/CFG) and browser sandboxing
- **MITRE:** M1048 Application Isolation and Sandboxing, M1050 Exploit Protection, M1051 Update Software, M1040 Behavior Prevention on Endpoint · DS0009 Process, DS0011 Module, DS0015 Application Log