WEEKLY THREAT BRIEF

Signal & Noise

RADICL
radicl.com
vSOC PUBLICATION
ISSUE 01 · WEEK OF JUL 06, 2026

INCIDENT SPOTLIGHT

A Fake Vendor, a Forged Approval, Zero Dollars Out

An invoice for $64,294.12 hit a client's accounts-payable inbox, already approved by one of their own executives. Except the executive never approved it, the vendor didn't exist, and the same forged paperwork had been rehearsed on another company days earlier. We shut it down before a cent moved.

A client's accounts-payable team forwarded us an invoice that felt slightly off. Good instinct. But a bad feeling doesn't stop a wire transfer, and what stops the wire is what happened next: our analysts didn't take the email at its word, they pulled the thread.

The ask was simple, pay a $64,294.12 invoice from a consulting vendor. The pressure was quieter than that: the email came wrapped in a thread that showed one of the client's own executives already signing off on the work. To a busy AP team on a deadline, this didn't read as a request. It read as a formality. That's the whole point of the play.

It fell apart on the second look. The email failed DMARC outright, and it hadn't come from any corporate vendor domain at all. It came from a student account at a vocational school, riding legitimate Microsoft infrastructure to clear the spam filters. The W-9 was worse: hand-edited in a Linux annotation tool, with two different company addresses stacked on top of each other where the attacker had swapped in their own payment details. And the approval thread everyone was trusting? Fabricated. The executive never sent a word of it.

We didn't stop at the email, because this kind of fraud usually rides a compromised insider. We pulled sign-in and mailbox activity on both the executive being impersonated and the AP inbox itself. Both were clean with no strange logins, no hidden forwarding, no inbox rules quietly burying the evidence. The attacker had never been inside. They'd just scraped a name off the internet and built a story around it.

Then came the part no single company could see on its own. That forged W-9 had been quietly run through a public malware scanner four days before it reached our client. This wasn't a one-off email, it was a template, and our client wasn't the first name on the list. What looked like one company's bad day was one stop on a route.

The ending is the boring kind, which is exactly how we like it: fraud confirmed, client warned off, not a dollar out the door. The invoice, the vendor, and the attacker's payout account are all documented now, so when this campaign resurfaces somewhere else across the businesses we protect, we won't be meeting it for the first time.

IN THIS ISSUE

THIS WEEK IN CYBER HISTORY

JUL 9 · 2015 · OPM BREACH

The U.S. Office of Personnel Management revealed a China-linked intrusion had exposed the background-investigation files of 21.5 million people, including SF-86 security-clearance records. Still the defining breach of the cleared workforce.

CASE AT A GLANCE

What it was
Fake-vendor invoice fraud (BEC), $64,294.12

How it arrived
Forged approval thread, sent from a compromised student account on Microsoft infra

The tell
Failed DMARC, hand-edited W-9, fabricated exec approval

Scope
One flagged invoice, traced to a shared fraud template

Confirmed impact
Zero dollars transferred; exec & AP accounts confirmed clean

Response
Fraud documented; client warned; IOCs filed fleet-wide

WHY THIS MATTERS

On its own, this is just one strange invoice, the kind that gets dismissed, or worse, paid. What no finance team can see from inside its own inbox is that the same forged document was tried on someone else first. That pattern only exists across companies, which means **only a SOC watching all of them at once can see it**. One flagged email became intelligence that now protects every client we have. That's the difference between getting lucky once and knowing the playbook.

MEET THE ANALYST

Dylan Haase

SOC ANALYST

Dylan spent his early career in military intelligence before earning a computer science degree from the Colorado School of Mines. He brings blue team expertise to RADICL as a SOC Analyst defending Defense Industrial Base contractors.

Drawn to the full picture of how attackers operate and how defenders catch them, he pursues red team skills on his own time, ranking in the **top 1% on TryHackMe** and pairing genuine offensive chops with his defensive work.

> "To defend the Defense Industrial Base, you have to think like the people trying to break in. So I practice both.
>
> DYLAN HAASE, SOC ANALYST · vSOC

CASE SPOTLIGHT

TRUE POSITIVE · MALICIOUS · REPORTED

Not every threat starts with a client alert. Proactive research turned up a legitimate company's public cloud-storage bucket that an outside actor had repurposed as free hosting for cross-platform malware and a scam page impersonating China's Ministry of Public Security. No client touched it. We fingerprinted the tooling and filed an abuse report to pull it down.

TRUE POSITIVE · NOT MALICIOUS · REMEDIATED

A network scanning tool had been quietly deployed across a client's environment, on many machines, with no sign anyone authorized it. It's a legitimate utility, but also a favorite of intruders mapping a network, and unsanctioned software at that scale is its own risk. We ran down every instance and cleaned them up.

TRUE POSITIVE · MALICIOUS · CONTAINED

A finance executive reported her laptop acting strange. We traced it to a browser hijacker that had gone unflagged by hiding in a less common startup mechanism, on a machine with access to company financials. Rather than just remove it, we called for a full rebuild.

EYE ON THE THREATSCAPE

01 FBI Seizes 13 China-Linked Sites Targeting Clearance Holders

The DOJ and FBI took down 13 fake consulting-firm domains that suspected Chinese intelligence used to lure current and former U.S. security-clearance holders with well-paid "consulting" gigs. Filings say the campaign leaned on AI-generated content and hiring platforms.

02 First Fully AI-Run Ransomware Attack Documented

Sysdig documented JadePuffer, the first ransomware run end-to-end by an AI agent. It even fixed its own botched login attempt in 31 seconds. More alarming is a flaw in its code: the AES key is pushed to STDOUT and never persisted or transmitted, so victims are out of luck even if they pay.

03 Five Eyes: AI Could Breach Defenses in "Months, Not Years"

The Five Eyes alliance issued a rare joint call to "act now," warning that AI models capable of overwhelming government and corporate defenses are months, not years, away. One official cautioned that under-invested SMBs will be "sitting ducks."

VSOC TIP OF THE WEEK

Before paying an invoice or changing vendor banking details, pick up the phone and confirm with the supposed approver directly using a number you already have. Attackers often add contact information that routes to them as a social-engineering technique. **They can even use AI to clone your contact's voice.** If you see anything like this, **send it to the vSOC.**