CSAI Foundation | Cloud Security Alliance

Laundry Bear's Zimbra Zero-Day: A Year of Russian Espionage

2026-07-24

AI-assisted Rapid Research

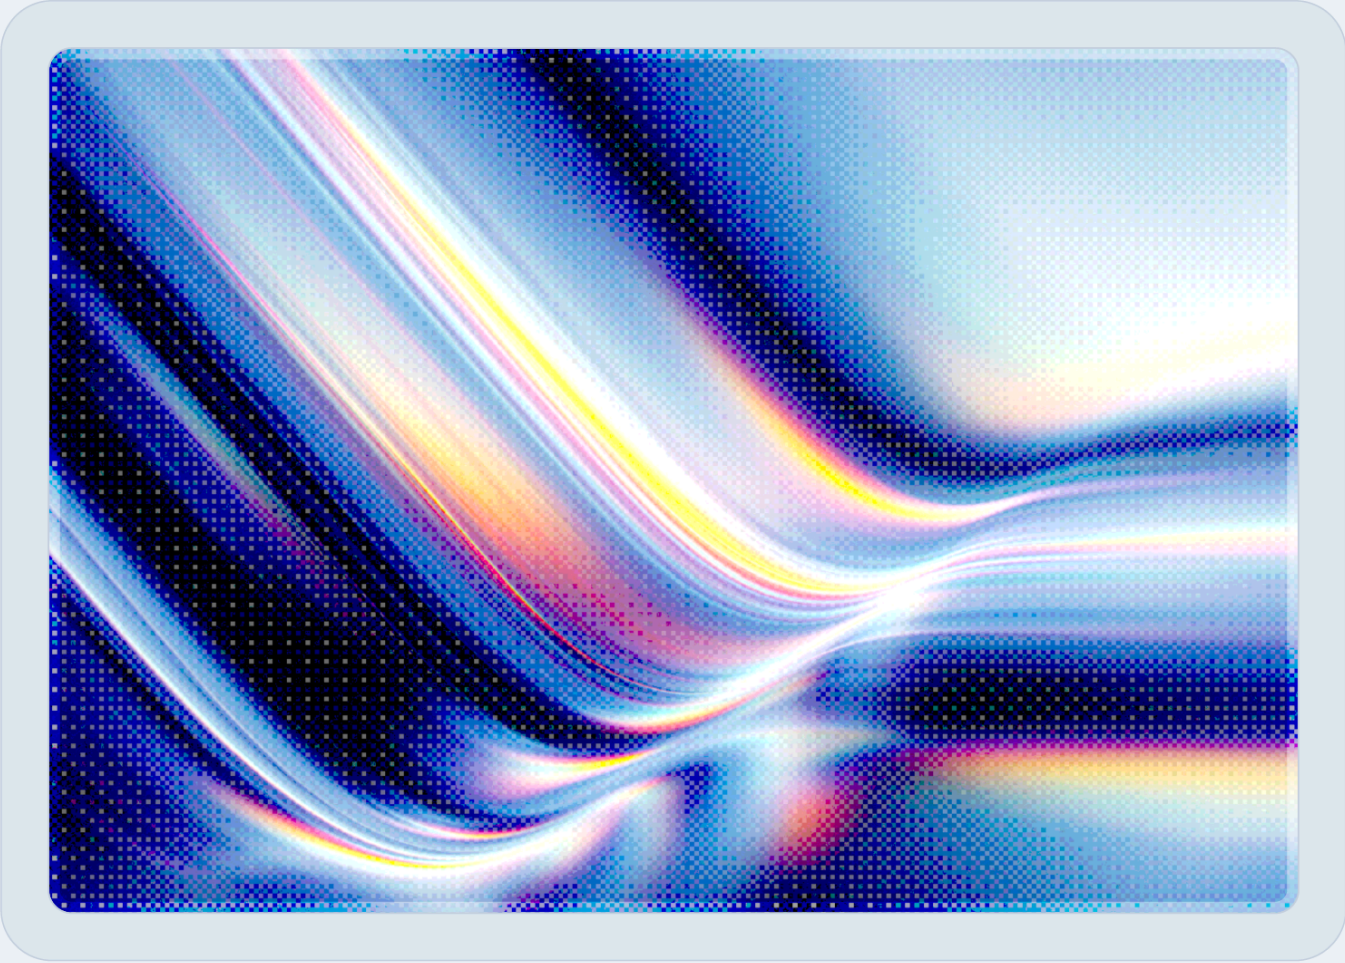

CSAI

CSA cloud security alliance®

Key Takeaways

A Russian state-supported group tracked as LAUNDRY BEAR by Dutch intelligence services, and also known as Void Blizzard, CL-STA-1114, and TA488, exploited a previously undisclosed cross-site scripting flaw in the Zimbra Collaboration Suite webmail client for roughly five months before a patch existed, and has continued targeting unpatched servers since [1][2][3]. A joint advisory issued July 23, 2026 by CISA, the NSA, Palo Alto Networks Unit 42, and Proofpoint, and co-signed by fifteen additional governments, details how the group used a single previewed email to steal ninety days of correspondence, an organization's complete address book, browser-saved passwords, and two-factor authentication recovery codes from government, defense, and critical infrastructure targets across NATO member states and Ukraine [1][3]. The exploitation chain, now assigned CVE-2025-66376, required no click and no attachment; the vulnerability lived in how Zimbra's HTML sanitizer handled CSS import directives embedded inside an SVG image, and the group's custom JavaScript payload, tracked by some researchers as Ulej and by Proofpoint as ZimReaper, used that gap to plant a persistent IMAP credential on the mailbox before the victim ever noticed anything unusual [1][4][5][12]. Because the flaw affected both the versions Zimbra shipped before its November 2025 fix and, in practice, any server administrators had not yet upgraded, the campaign illustrates how quickly a narrowly scoped webmail vulnerability can become a lasting espionage foothold when patch cycles for on-premises collaboration software lag behind active nation-state interest.

Background

Zimbra Collaboration Suite is a self-hosted email and calendaring platform commonly used by government agencies, universities, and mid-sized enterprises that prefer to keep mail infrastructure on their own networks rather than migrate to a cloud provider such as Microsoft 365 or Google Workspace. That preference is often driven by data residency requirements, cost, or long-standing internal customization, and it is precisely the profile of organization the joint advisory says LAUNDRY BEAR pursued: agencies and companies across the defense, transportation, financial services, scientific research, energy, technology, education, media, and law enforcement sectors, along with operators of nuclear installations and members of national defense industrial bases, concentrated in NATO member states, Ukraine, and the Commonwealth of Independent States, with additional activity in Africa and the United States [1][2].

LAUNDRY BEAR itself is not a new actor. Dutch intelligence services, the AIVD and MIVD, first named the group publicly in May 2025 after tracing a 2024 breach of Dutch police systems back to a previously untracked Russian-nexus cluster; Microsoft published a parallel disclosure the same week under the name Void Blizzard, assessing the group had been active globally since at least 2024 with a disproportionate focus on NATO members and Ukraine [6][7]. That earlier tradecraft centered on cloud email: password spraying, adversary-in-the-middle phishing proxies, stolen credentials, and session-token replay used to abuse legitimate Exchange Online and Microsoft Graph APIs, enumerating mailboxes and cloud-hosted files before automating bulk collection. The Zimbra campaign documented in the July 2026 advisory represents a notable evolution in capability. Rather than relying on credential theft and legitimate cloud APIs to reach mail already hosted by a major provider, the group invested in discovering and weaponizing an unpatched vulnerability in a self-hosted platform, then built a bespoke JavaScript payload to operate inside it. Analysts tracking the intrusion set under different names, Proofpoint's TA488 and Unit 42's CL-STA-1114 among them, caution in the advisory that the mapping between these clusters and LAUNDRY BEAR or Void Blizzard may not be fully one-to-one. A related but distinct data point adds to that naming complexity: the security vendor Seqrite has separately assessed, with only medium confidence, that a January 2026 phishing intrusion against a Ukrainian maritime and hydrology agency, exploiting the same CVE-2025-66376 vulnerability and tracked by Seqrite as Operation GhostMail, may be the work of the long-running GRU-linked group APT28; Seqrite's report does not identify that intrusion as part of the broader LAUNDRY BEAR/Void Blizzard campaign described in the joint advisory, and names no overlap with TA488 or CL-STA-1114 [13]. That attribution uncertainty does not diminish the operational significance of either campaign, but it is a reminder that threat actor naming across vendors remains an imperfect science, and defenders should treat the technical indicators in the advisory as the actionable core of the disclosure regardless of which label, or which related campaign, ultimately proves correct.

The vulnerability at the center of the campaign, CVE-2025-66376, is a stored cross-site scripting flaw in the Zimbra Classic UI webmail client caused by insufficient sanitization of CSS `@import` directives when they appear inside email content. Synacor, Zimbra's maintainer, patched the flaw in November 2025 in Zimbra Collaboration 10.0.18 and 10.1.13, and CISA added it to the Known Exploited Vulnerabilities catalog on March 18, 2026 after confirming active exploitation. Notably, LAUNDRY BEAR appears to have been exploiting the flaw as a genuine zero-day for at least five months before that patch existed, with the earliest confirmed activity dating to July 2025, meaning the group had a substantial undisclosed window of operation before Synacor, or anyone else, knew the flaw existed [1][2][3].

Security Analysis

The exploitation chain is notable for how little it asks of a victim. Traditional phishing depends on convincing someone to click a link, open an attachment, or enter credentials into a spoofed page. LAUNDRY BEAR's Zimbra campaign instead required only that a target open or preview a malicious email inside a vulnerable webmail session, a distinction the joint advisory itself emphasizes as a meaningful departure from the group's earlier, more conventional phishing tradecraft [1][4]. Messages arrived from adversary-controlled Proton Mail accounts or from addresses the group had already compromised, using deliberately generic lures that appear designed to avoid raising suspicion before the email was even opened. The payload itself was concealed inside an SVG image element, with the malicious script disguised using a tag-splitting technique that broke the executable markup across CSS `@import` directives and HTML comments, a fragmentation technique consistent with an attempt to slip past Zimbra's HTML sanitizer and signature-based detection tools that look for intact, recognizable script patterns; the script was further obfuscated with base64 encoding and XOR encryption before delivery [4][5].

Once triggered, the payload, which some researchers have identified as a custom tool named Ulej (Russian for beehive) while Proofpoint separately tracks the same tool as ZimReaper, executed with the full session privileges of the victim's webmail account and proceeded through a methodical collection sequence using Zimbra's own SOAP API rather than any external exploitation framework, making its activity harder to distinguish from ordinary webmail usage [4][5][12]. It harvested CSRF tokens and any passwords saved in the browser, extracted two-factor authentication scratch codes directly through the platform's own account-recovery APIs, and packaged the victim's most recent ninety days of email into a compressed archive for exfiltration. Critically, the malware also used its access to create a new application-specific password labeled "ZimbraWeb" and to enable IMAP access on the account by flipping the `zimbraPrefImapEnabled` setting to true, giving the group a durable, credential-based foothold that would survive well after the initial script had finished running and the victim had closed the offending email. From that position the malware brute-forced the organization's Global Address List to map the broader user directory, extending the group's reach beyond the initial victim's mailbox. Command-and-control and data exfiltration relied heavily on DNS, with the malware issuing lookups against long, randomly generated subdomains, a technique that blends into ordinary DNS traffic and is difficult to catch without tooling purpose-built to flag anomalous query entropy. Unit 42's tracking of the associated infrastructure identified nine IP addresses and nine domains supporting the operation, with an average server lifespan of just over a month, consistent with deliberate infrastructure cycling to stay ahead of takedown and blocklisting efforts [11].

Two aspects of this campaign deserve particular attention from defenders beyond the immediate patch requirement. First, the targeted theft of two-factor authentication scratch codes shows that the group treats backup MFA mechanisms, not just primary passwords, as a collection priority; scratch codes that sit dormant in an account's settings for emergency use are just as valuable to an attacker with script execution inside the session as the credentials they are meant to back up. Second, the creation of a persistently named application password is a deliberate persistence mechanism that outlives the vulnerability itself: patching CVE-2025-66376 stops new exploitation, but it does nothing to revoke an IMAP credential the attacker already planted on a compromised account before the patch was applied. Organizations that upgrade Zimbra without separately auditing for artifacts of prior compromise risk leaving that access in place indefinitely.

Recommendations

Immediate Actions

Organizations running Zimbra Collaboration Suite should confirm they are on version 10.1.13 or later, and ideally the current 10.1 build, since Zimbra 10.0 reached the end of its support lifecycle on December 31, 2025 and will receive no further security fixes regardless of exposure to this or future vulnerabilities [14]. Administrators should search Zimbra's `/opt/zimbra/log/audit.log` for `CreateAppSpecificPassword` calls, with particular attention to any credential named "ZimbraWeb," and for `GetScratchCodesRequest` SOAP calls or unexplained changes to the `zimbraPrefImapEnabled` setting, all of which are direct indicators tied to this campaign's known tradecraft. Any account showing these artifacts, or any account that previewed a suspicious message during the exploitation window beginning in July 2025, should have its passwords reset, its active sessions invalidated, and its two-factor authentication scratch codes regenerated. Security teams should also block the C2 domains and IP addresses published in the joint advisory and add detection rules for the long, randomly generated subdomain lookups characteristic of the malware's DNS-based exfiltration channel.

Short-Term Mitigations

Because the campaign specifically targeted static two-factor scratch codes, organizations should move toward phishing-resistant, hardware-bound authentication such as FIDO2 security keys for high-value accounts rather than relying on backup codes that a script running inside an active session can extract wholesale. IMAP and POP3 access should be disabled by default and enabled only where there is a

documented business need, since the malware's persistence mechanism depended entirely on being able to turn IMAP on for an account that likely did not need it. Given that the root cause was a sanitizer bypass hidden in nested CSS and HTML markup, defenders should not treat patching alone as sufficient; a compensating control such as a web application firewall rule set tuned to flag SVG-embedded script fragments and unusual `@import` patterns in inbound mail can catch similar bypass techniques against this or related webmail platforms before a vendor patch exists. Audit log retention should be extended to cover the full disclosed exploitation window, since the July 2025 start date means some organizations will need to look back nearly a year to fully rule out earlier compromise.

Strategic Considerations

This campaign is consistent with a broader pattern our team has observed in on-premises collaboration software, though a systematic platform-by-platform comparison is beyond this note's scope: on-premises webmail and document-collaboration servers are attractive, enduring espionage targets precisely because they sit outside the rapid patch cadence and centralized telemetry that cloud-hosted equivalents benefit from, and because a single sanitizer flaw can be weaponized quietly for months before public disclosure. Organizations that continue to operate on-premises Zimbra, or similar self-hosted collaboration infrastructure, for data residency or cost reasons should budget for an accelerated patch SLA on that infrastructure specifically, treating it as a persistently high-risk asset class rather than routine IT plumbing. More broadly, the group's evolution from cloud API abuse using stolen credentials to custom zero-day exploitation of on-premises software suggests LAUNDRY BEAR is investing in vulnerability research capability, not just social engineering, and defenders in the sectors named in the advisory, government, defense, energy, and the broader NATO-aligned defense industrial base, should treat this as a signal to build threat-informed monitoring around the group's evolving toolset rather than assuming its earlier, better-known cloud-focused tradecraft remains its primary method of entry.

CSA Resource Alignment

This campaign's reliance on phishing as an initial delivery mechanism, executed by a state-sponsored actor with clear nation-state resourcing, connects directly to CSA's research note AI-Weaponized Phishing: Nation-State Quality at Commodity Scale, which documents how government-backed actors from Russia and other states have increasingly industrialized phishing lure generation and target reconnaissance. While the Zimbra campaign's lures were described by investigators as deliberately generic rather than AI-tailored, that CSA analysis provides useful context for why phishing-delivered zero-click exploitation is becoming a repeatable playbook for state actors rather than a one-off

technique, and it names the same Russia-nexus tracking conventions, including APT28-linked infrastructure, that this advisory and the related Seqrite reporting also raise as a possible, if unconfirmed, overlap.

The campaign's systematic theft of two-factor authentication scratch codes and creation of durable application-specific credentials also maps closely to the findings in CSA's EvilTokens: Device-Code Phishing Renders MFA Irrelevant, which examined a separate Russian-aligned campaign against government, NGO, defense, and energy targets that similarly bypassed multi-factor authentication not by defeating it outright but by capturing the tokens and codes that MFA implementations depend on. Both campaigns illustrate the same underlying lesson for defenders: static backup codes and long-lived application credentials are frequently weaker links than the primary authentication factor they are meant to support, and organizations should extend MFA hardening efforts to cover these secondary mechanisms explicitly.

Finally, the technical and governance gaps this campaign exposed, unsanitized input handling in a security-critical parsing path, insufficiently monitored session and credential lifecycle management, and delayed patch application on internet-facing infrastructure, map to multiple domains within the CSA AI Controls Matrix (AICM) v1.1, particularly its Threat and Vulnerability Management and Application and Interface Security domains. Organizations using the AICM as an assessment baseline should confirm their control implementation addresses timely patching of self-hosted collaboration software, monitoring for anomalous application-credential creation, and periodic review of MFA recovery mechanisms, each of which this campaign demonstrated to be a viable and exploited attack path [8][9][10].

References

[1] CISA, NSA, Palo Alto Networks Unit 42, and Proofpoint. "Russian State-Supported Cyber Actors Conduct Phishing Campaign Targeting Users of Zimbra Collaboration Suite (AA26-204A)." CISA, July 23, 2026.

[2] The Hacker News. "Russian Espionage Group Exploited Zimbra Zero-Day to Steal Mail and 2FA Codes." The Hacker News, July 2026.

[3] CyberScoop. "Russian espionage group using novel Zimbra exploit to steal sensitive data from Western countries." CyberScoop, July 2026.

[4] Security Affairs. "US Agencies Warn of Laundry Bear Campaign Targeting Unpatched Zimbra Servers." Security Affairs, July 2026.

[5] GBHackers. "Russian LAUNDRY BEAR Hackers Exploit Zimbra Zero-Day to Steal 90 Days of Emails." GBHackers, July 2026.

[6] Microsoft Threat Intelligence. "New Russia-affiliated actor Void Blizzard targets critical sectors for espionage." Microsoft Security Blog, May 27, 2025.

[7] The Record. "Dutch intelligence unmasks previously unknown Russian hacking group 'Laundry Bear'." The Record from Recorded Future News, May 2025.

[8] Cloud Security Alliance. "AI-Weaponized Phishing: Nation-State Quality at Commodity Scale." CSA AI Safety Initiative, June 14, 2026.

[9] Cloud Security Alliance. "EvilTokens: Device-Code Phishing Renders MFA Irrelevant." CSA AI Safety Initiative, May 20, 2026.

[10] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.

[11] Unit 42, Palo Alto Networks. "Russian Global Webmail Espionage." Unit 42, July 23, 2026.

[12] Proofpoint. "TA488 Targets Zimbra Mailservers with Half-Click Exploits." Proofpoint Threat Insight, July 23, 2026.

[13] Seqrite. "Operation GhostMail: Russian APT Exploits Zimbra Webmail to Target Ukraine State Agency." Seqrite Labs, March 17, 2026.

[14] Zimbra. "Product Lifecycle." Zimbra by Synacor, 2026.