halcyon

The Executive Guide to Ransomware Tactics

OUTMATCHED: 5 Tools That Won't Deliver Ransomware Resilience

One Opportunity Is All It Takes

Ransomware doesn't need stability; it just needs opportunity. Even with major RaaS takedowns and arrests, ransomware still thrives on the smallest openings. **EDR evasion** is a tactic seen in almost every ransomware attack; however, it is not the only one. Other tools in the security arsenal are also targeted. Social engineering, careless mistakes, and insider threats keep handing ransomware plenty of opportunities served on a silver platter.

It is only a matter of time before a motivated or opportunistic threat actor will discover openings in your defenses that merely require the right timing or circumstance. As they move along the attack chain, these tools can become minor inconveniences at best for threat actors. Ransomware resilience starts by understanding how attackers bypass defenses and slamming those doors before they stroll through.

1 Backups: Recovery Plans That Never Recover

Backups and recovery are cybersecurity's version of eating vegetables. Everyone knows they should do it, but nobody wants to think about it until they're due for that annual physical checkup. But in ransomware, backups aren't optional; they're survival. And yes, the phrase "Back That Thing Up" is sticker-famous, which is fun until it becomes the entire recovery strategy. Threat actors know this full well, which is why, when it comes to data extortion and destruction, backup and recovery become public enemy number one.

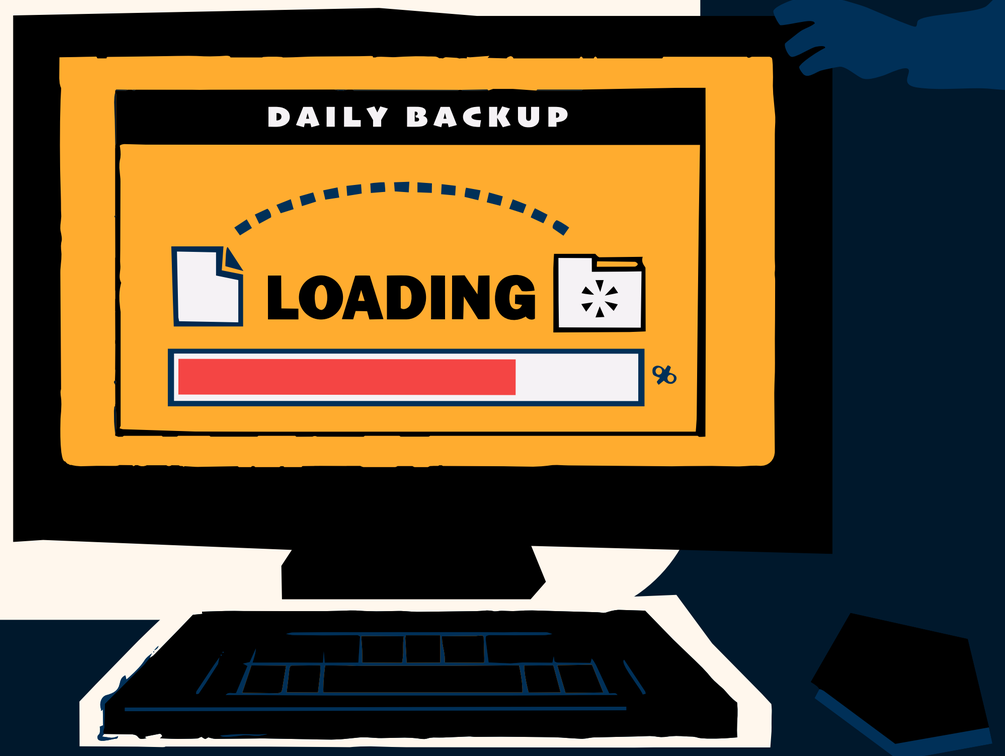

The tried and tested go-to for threat actors is deletion or disabling backups.

Deleting or Disabling Backup Data:

The tried and tested go-to for threat actors is deletion or disabling backups. This can come in many different forms, but the most common include deleting snapshots, erasing backup files, or adjusting retention policies and compromising credentials/API access to backup targets to orchestrate these deletions. Think this is a thing of the past? Consider this. Research by Veeam has found that **93% of ransomware attacks target backup storage**, and in about 75% of those cases, attackers succeed in crippling the ability to recover using backups. If your plan starts and ends with backups, ransomware groups already know your playbook.

The good news is that there are some quick ways to address this. Air-gapped backups preferably stored off site can be lifesaving if all else fails. While this can be a tedious process, it can also be the difference between months of possible downtime or just a few hours, at best. Enforcing least-privilege access is just as important to prevent general access. Finally, backing up to immutable storage ensures data cannot be tampered with or deleted.

Deleting Volume Shadow Copies & System Recovery Points:

Instead of targeting backing up systems directly, ransomware attackers use native OS tools to destroy local recovery data so endpoint or system restore won't help. Tools like *vssadmin*, *diskshadow*, and *wbadmin* are abused to remove point-in-time recovery artifacts. These techniques are simple in their application and need minimal commands to do so. Shadow copies disappear, backup metadata gets wiped, and recovery services get shut off, because recovery is bad for business... their business.

Prevention comes down to three things: limiting who can manage recovery features, detecting suspicious backup/restore tampering early, and keeping immutable recovery points that attackers can't delete. If ransomware can erase your restore options, it will, enthusiastically.

Compromising Cloud / Network Backup Configurations:
Cloud-first doesn't mean ransomware-last, and here's one you'll want to watch for. Cloud backups aren't magic; they're simply settings. And threat actors with admin access love changing settings. Ransomware groups have also gone cloud native as well, targeting common storage locations like S3, Azure Blog Storage and Azure Storage Explorer. So, what more often than not ends up in these cloud storages? According to **CSO Online**, more than half (54%) of organizations using AWS ECS task definitions and 52% using GCP CloudRun have secrets embedded in configurations. Around 3.5% of AWS EC2 instances contain secrets in user data. With stolen creds in hand, attackers don't waste time: **cloud backups get wiped, retention gets revoked, and exfiltration kicks off immediately.**

Protect cloud consoles with strict RBAC and phishing-resistant MFA since attackers love cloud admin access almost as much as you do. Split duties so backup controls aren't sitting behind the same keys used for everyday work. Finally, enforce immutable snapshot locks so backups can't be deleted early, even by an admin having a very bad day.

2 SIEM and Log Management: Audit Trails Edited by Ransomware

Security Information and Event Management (SIEM) has been the gold standard for roughly 20 years for log aggregation, centralization, and correlation. When a security tool has this much insight, it's basically wearing a neon sign that says, "*Disable me, and quick!*"

Clearing Logs:

One of the quickest ways around SIEM is to simply clear the Windows Event Logs. No logs means no ingestion, slower triage, and a timeline held together with vibes and assumptions. Native tools built into your OS, designed to archive, export, or clear logs, can be leveraged by threat actors to rapidly clear log data. RaaS groups such as Akira have been documented using this method to hide their tracks.

Addressing this is straightforward. Forward logs in real-time, alert on *wevtutil cl* usage, and log disruptions, sudden gaps in log volume and missing heartbeat signals. Ship logs to immutable storage attackers can't edit and lock down who's allowed to touch logging policies, because "everyone's an admin" is not a security strategy.

Turn Off The Pipes That Feed SIEM:

If clearing the logs isn't an option, why not just stop the logs from ever reaching the SIEM in the first place by disabling the pipelines.

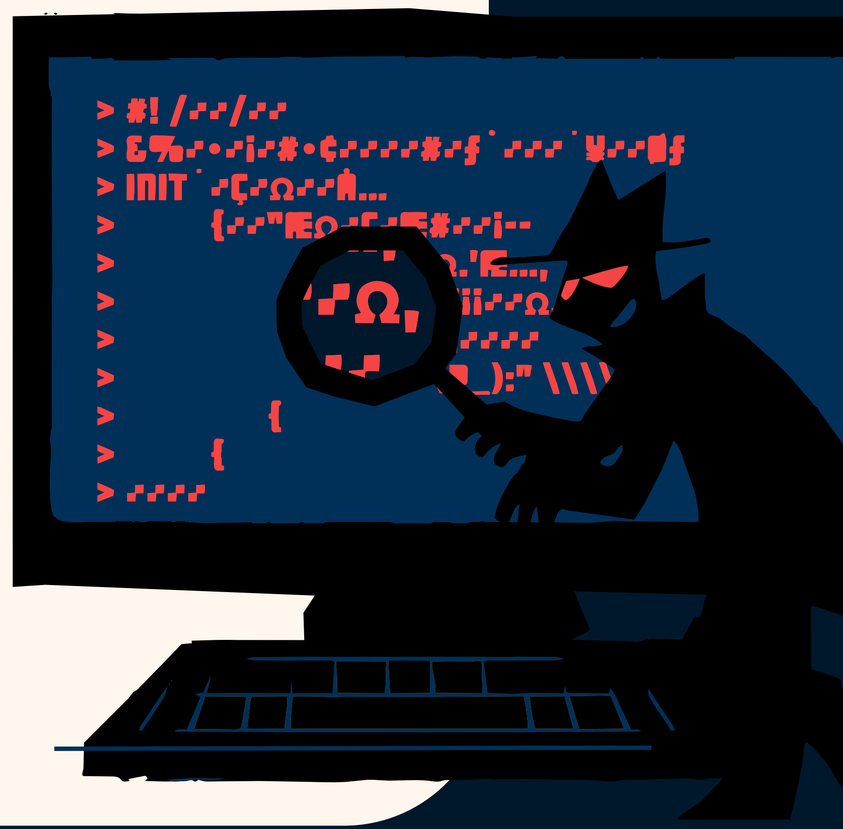

How is this done? Through a four-step process:

1. Identify the shipping log's source (log forwarding, agents, and scheduled tasks)
2. Stop the forwarder/shipper service (Splunk Universal Forwarder, Elastic Winlogbeat, agents)
3. Prevent restarts (disable startup mode, remove reschedule restart tasks, change service configs)
4. Operate in the darkness (leverage LOTL tools, stage ransomware payloads, exfiltrate data)

The best defense here is to alert on forwarder stoppages, protect log shipper services by restricting access, and monitor and require change control. Forward logs to immutable storage quickly and, lastly, use redundancy to ship logs via multiple paths to avoid single points of failure.

Quiet the Sensor Layer (ETW/EventLog): Typically, the most dangerous scenario is this bypass. It's a silent, nonchalant dashboard masking a very active attacker. Defenders often think log evasion means threat actors clearing event logs after the fact. Modern ransomware operators go further by disrupting the sensor and telemetry pipeline itself, so logs never get generated, forwarded, or remain complete and reliable enough for detection and response. **In short, instead of deleting evidence, they prevent the evidence from ever existing.**

Event Tracing for Windows (ETW) is a core Windows telemetry mechanism used by security products like EDR/XDR, monitoring pipelines, detection rules that depend on network visibility, and threat hunting. Tampering with these pipes doesn't just mean fewer logs; it means reduced detections, weaker correlations, missing log context and incomplete timelines. Quieting the sensor layer doesn't always look like dramatic log wipes; it often shows normal-looking logs with missing details, fewer process events than expected, and suspiciously low alert volume while bad things are clearly happening. ETW tampering is dangerous because it makes reality and telemetry disagree. So, defenders feel confident right up until the ransomware note proves otherwise.

The best way to avoid this is to watch for telemetry drop-offs. Like the quiet before a storm, silence can be deceptive. Least privilege and immediate log forwarding help spot sensor disruption early, because **zero alerts doesn't mean you are under zero attacks**. Finally, keep an eye out for mismatches between activity and log data and cross-check known bad patterns appearing, such as privilege escalation, EDR tampering, and log forwarding stoppage.

3 MFA: The Three Doors in Your Two-Factor Wall

Multi-Factor Authentication (MFA) is incredibly effective at stopping basic credential attacks. At its core, MFA adds a second proof-of-access (usually a one-time code), so a stolen password alone isn't enough to get in. Unfortunately, ransomware operators have graduated from basic tactics with the rise of Ransomware-as-a-Service groups. Unsurprisingly, MFA bypass techniques have evolved fast, and the list is longer than most defenders want to admit.

Help Desk / Service Desk Social Engineering: Why bypass MFA the hard way, when you can just request a new factor like its IT room service? In 2023, several phone calls were placed to the **MGM help desk**, impersonating employees, requesting access to passwords and two-factor codes.

Once these were granted, attackers had the access needed to cause one of the largest attacks targeting the hospitality sector to date. Social engineering isn't new, but it remains overwhelmingly effective. Threat actors impersonating employees or technicians can either be granted temporary codes or add new rogue devices to the approved MFA list by creating a sense of urgency.

Addressing this is simple in theory, so long as policy is followed. Implement strict help-desk verification, manager approvals, and verified device posture for any MFA changes/resets. Enrollment holds and alerts will add friction, but that friction is precisely what keeps companies from landing on a **'Top 5 Worst Breaches'** roundup.

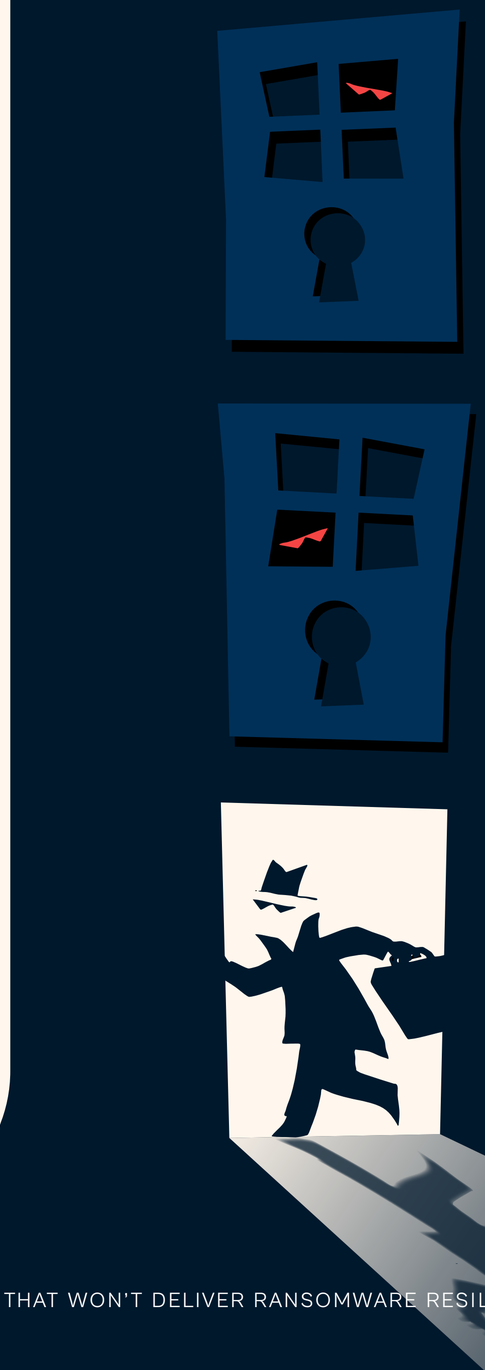

MFA Fatigue (Push Bombing):

What's worse than one random MFA text sent without solicitation? Several. And to add to the annoyance, a social engineering follow-up from an IT team in a foreign country claiming it's legitimate. The catch? That 'IT person' isn't on your payroll, and the only thing they're troubleshooting is how fast they can bypass your MFA. Think it can't happen to you? Try receiving this alert at 1 AM and throughout the remainder of the morning with a simple approve/deny link. One quick approval click is all that's needed, and the incessant prompts go away. This is exactly what happened to Uber in 2022. This resulted in threat actors gaining access to additional authentication data, and Uber's bug bounty showcasing vulnerabilities that had yet to be addressed.

So, what's the best way to prevent late-night MFA fatigue? Move high-risk users to phishing-resistant MFA, such as FIDO2/WebAuthn, and reduce reliance on push-based approvals. Alternatively, disable approve/deny approvals when possible and require number matching or stronger prompts with alerting on excessive push attempts. Finally, train users to deny unsolicited push requests immediately.

Adversary-in-the-Middle (AiTM) / Reverse Proxy Phishing:

The most modern type of MFA bypass isn't social engineering-based; it's an adversary-in-the-middle (AiTM) reverse-proxy attack that lets the user authenticate for the attacker. The victim lands on a convincing phishing site that quietly proxies the real login flow, captures the username/password, and then watches the user complete MFA who thinks everything is normal. Once the login succeeds, the identity provider issues a session cookie/token, the proxy steals it, and the attacker replays that session from their own machine. No new MFA prompt is required. From there, they pivot into email, files, and SSO apps, set persistence (mail rules, OAuth grants, new MFA enrollment), and move straight into ransomware's favorite activities: recon, exfiltration, and encryption.

Stopping this type of bypass isn't impossible, but it does require some preventive measures. As previously mentioned, phishing-resistant MFA is still the go-to here. FIDO2/WebAuthn, passkeys, and hardware-backed authenticators bind authentication to the real domain/device, making token replay far harder. Conditional access, such as blocking risky geographies, requiring a compliant device posture, and limiting sessions by IP, device, or risk score, will make life harder for threat actors. Lastly, reduce the potential blast radius by enforcing least privilege, monitoring new MFA enrollments, and alerting on mailbox forwarding and rule creation.

4 Vulnerability Management: Patch Tuesday, Ransomware Wednesday, Recovery Friday

Patch management reduces attack surface, but ransomware operators thrive on the exceptions list. They look for the one server that couldn't be rebooted, the one appliance no one owns, or the one critical update that got deprioritized. In other words: they win by being patient, punctual, and precise.

Exploiting Unpatched Remote Management Software:

This attack path is still incredibly common today. Security teams keep finding abandoned RMM tools, mostly because RMM sprawl is real and hearing "we forgot that was installed" is more common than it should be. Once lost and now found, these tools are notorious for having several unpatched vulnerabilities waiting to be discovered like ancient sunken treasure. CVE-2024-57727, used against a vulnerable SimpleHelp RMM was just one of several exploitable vulnerabilities used to gain access to and around patch management policies. These vulnerabilities can also be packaged into further bypasses, commonly used to get around EDR.

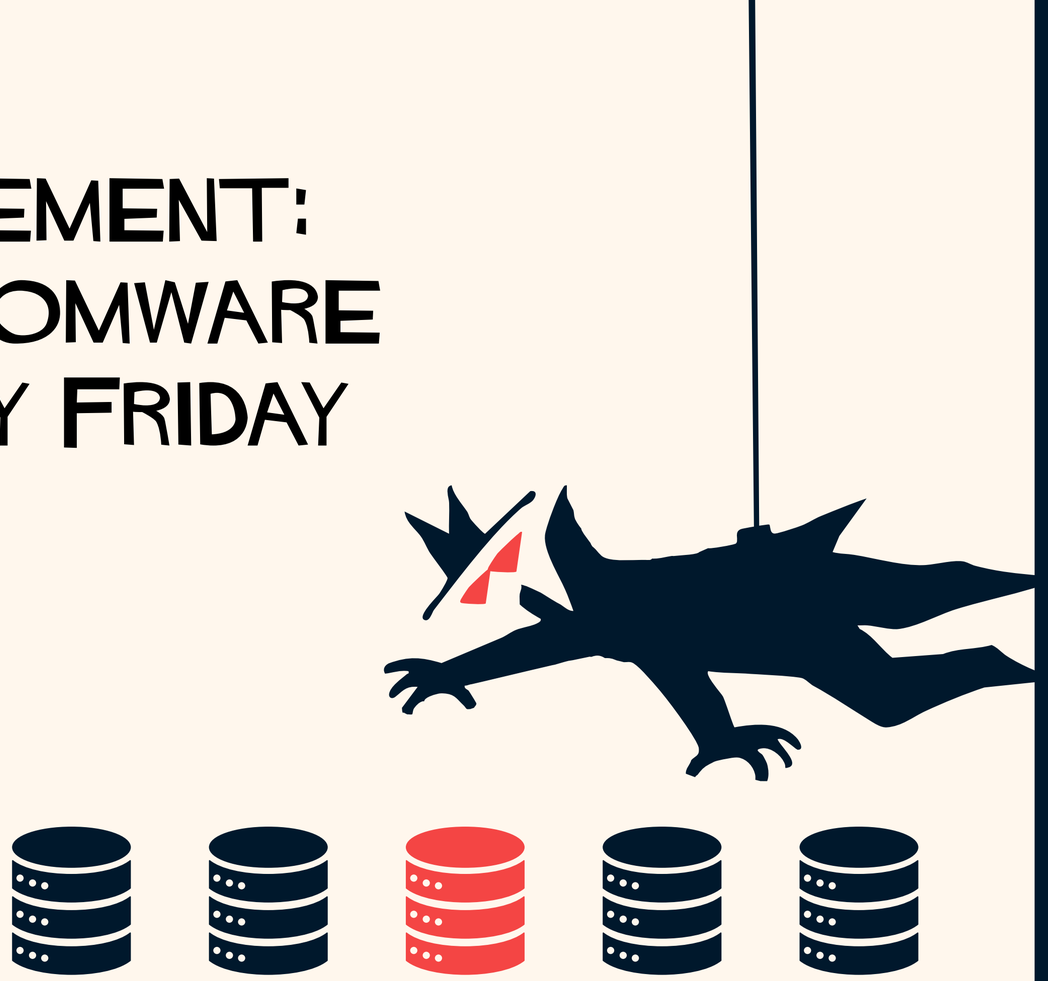

Attackers look for the one server that couldn't be rebooted, the one appliance no one owns, or the one critical update that got deprioritized. In other words: they win by being patient, punctual, and precise.

Defensive measures here include patching RMM tools immediately; “we’ll get to it later” is basically just a ransomware scheduling assistant. Keep remote management off the open internet (VPN/zero-trust only), and segment it away from critical systems. Then alert on suspicious logins, because privileged tools shouldn’t be used like public Wi-Fi.

Exploiting Virtualization and Infrastructure Vulnerabilities: Ransomware operators don’t always start with endpoints. They can go for virtual platforms running everything, using unpatched hypervisor and host vulnerabilities to get RCE or escalate privileges. Critical VMware ESXi-related bugs have been actively exploited and flagged on CISA’s Known Exploited Vulnerabilities list, meaning this isn’t theoretical. If your hypervisor is behind on updates, it’s not a minor to-do task for later; it’s an attack bypass path.

Keep virtualization platforms current on critical updates and prioritize KEV-listed vulns like your uptime depends on it, because it does. Pair patching with segmentation so attackers can’t pivot across hosts, and audit configs regularly to make sure management consoles aren’t one network bridge away from public access.

Exploitation of Other High-Risk Unpatched Services: This technique is basically scan, exploit, repeat. Attackers target any internet-facing service with a known vuln that gives them execution or a pivot point. **Fortinet** bypass vulnerabilities have been exploited to seize control of network gear, **GoAnywhere MFT** bugs have been used to deploy Medusa ransomware, and SharePoint issues like **ToolShell** have been leveraged for ransomware delivery. If it’s unpatched and reachable, ransomware treats it like an open door with a welcome sign.

Run fast patch cycles for high-severity vulnerabilities, especially anything exposed externally; those are ransomware’s favorite shortcuts. Shrink your attack surface by removing unsupported software, limit permissions on critical services, and automate exploit-aware alerts, so you’re not learning about exposures from the ransom note.

5 Data Loss Prevention: Taking the Back Roads Past DLP

Data Loss Prevention (DLP) helps prevent unauthorized sharing of sensitive data across endpoints, email, and cloud services. The challenge is ransomware operators are built for evasion, and they're experts at moving data in ways that don't look unauthorized until it's too late. So DLP can be part of your security stack and still miss the critical moment that matters.

Encrypting or Compressing Data Before Exfiltration:
Attackers often encrypt or compress data before exfiltration, because many DLP tools prefer their data in a nice, readable, clear text. Wrap it up in a zip or encryption layer and, suddenly, content-based rules can't tell if it's sensitive customer records or cat photos. Ransomware groups use this to quietly stage and steal sensitive data before, or alongside, encryption to maximize payment opportunities.

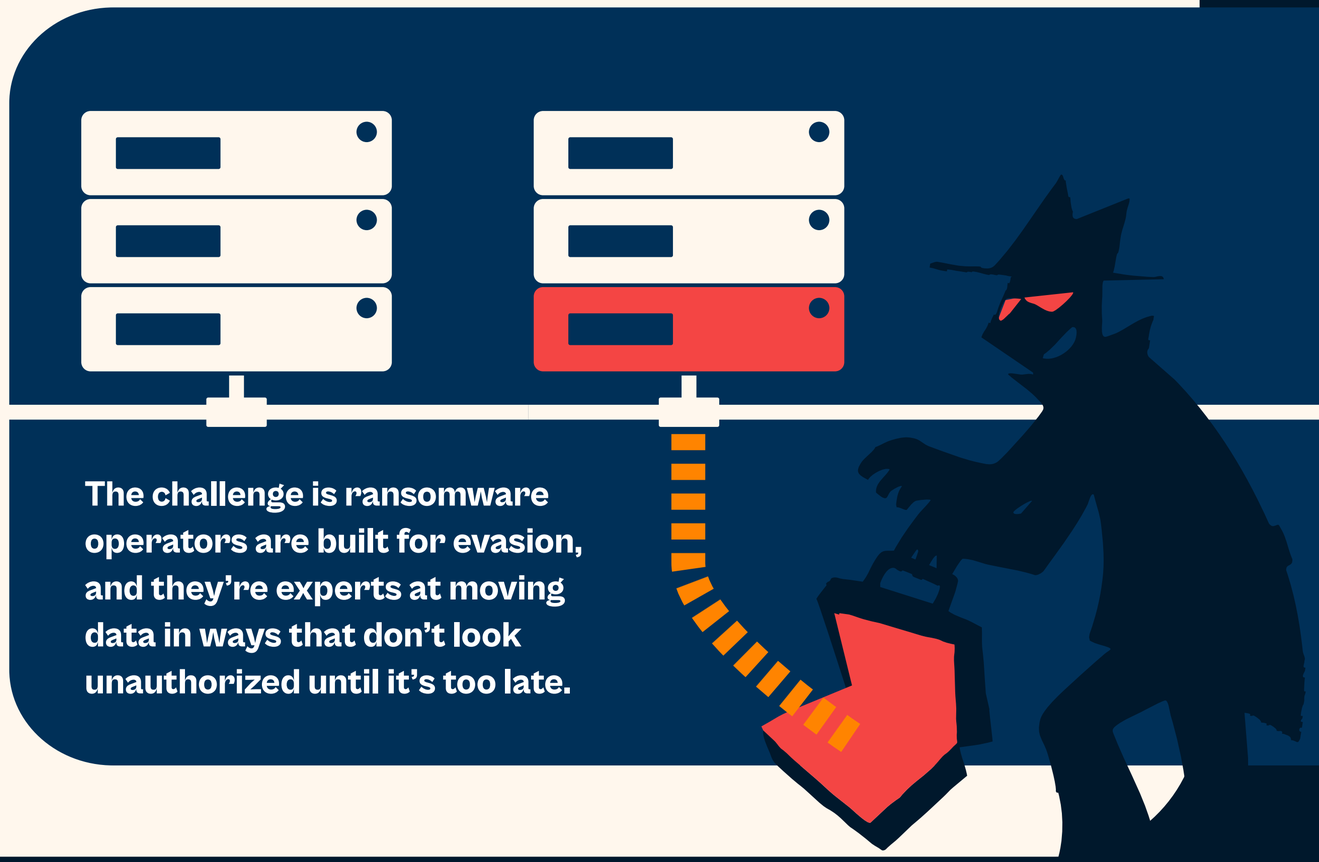

Configure DLP to inspect metadata and file attributes, not just clear text content, and alert on unusual compression/encryption activity. Monitor suspicious data movement and stop exfiltration before it leaves the environment using tools that target this directly. If you're not sure where to begin, here is a **Halcyon DXP blog** to help you get started. If attackers insist on zipping it, the goal is making sure they can't ship it.

Leveraging Legitimate File Transfer Tools or Cloud Services:
One of the easiest ways to bypass DLP is to exfiltrate data using tools that already look normal in enterprise environments like sync apps, file transfer clients, and remote admin utilities. Attackers commonly use rclone, MEGASync, and FTP/SFTP tooling, and may rename executables to blend in and avoid signature-based alerts. Ransomware operators use this to move sensitive files fast.

Lock down known file transfer utilities (and their "definitely-not-rclone.exe" cousins), restrict unapproved cloud/FTP destinations, and watch for suspicious endpoint transfers.

The goal is simple: stop attackers from using legitimate tools as illegitimate data movers. If it looks like a bulk upload and smells like exfiltration, it probably isn't your coworker uploading their favorite TikTok influencers' new content.

Avoiding Traditional DLP Inspection Points:
DLP is great at catching data exiting the front door. Attackers prefer back doors, encrypted tunnels, weird protocols, and covert channels that don't get scanned the same way. Ransomware crews use these stealth paths to move data out without lighting up content-based policies. After all, no alert is the best alert for any good ransomware actor.

Cover endpoint egress, not just perimeter flows, and correlate DLP with network anomaly detection to catch suspicious uploads early. **Attackers will always find a quieter channel, so the win is detecting the behavior, not just the payload.**

Where Opportunity Meets Its Match

Your backups can vanish, your SIEM can be silenced, your MFA can be phished, your patches can lag, and your DLP can be tunneled around like a speed bump. Welcome to modern ransomware warfare. These aren't amateurs stumbling through your network. They're professionals who've done their homework, tested their methods, and know exactly which doors to pry open. They don't need your defenses to collapse; they just need one opportunity. And, guess what? They're very good at finding it.

Ready to combat them with equal expertise?
Ransomware is a business problem.
Resilience is the answer.

See what that means for your organization.

halcyon