COHESITY
REDLab

Ransomware vs. the last line of defense

Observed tradecraft, detection coverage, and operational outcomes against 53 strains—January 2025 through May 2026

FIRST EDITION - 2026

Foreword

We built REDLab to answer a simple question: when real world ransomware reaches your backup systems, what actually happens?

Seventeen months later, the answer is more complicated—and more consequential—than we expected. Some of what we found confirmed widely held assumptions about how these attacks unfold. Some of it didn't. A few findings genuinely surprised us, and we've spent considerable time stress-testing them before putting them in writing.

This report is the result of that work. It covers 53 ransomware families detonated in a controlled environment against the backup infrastructure that organizations rely on when everything else has failed and recovery is the only option. We mapped every detonation to a common framework. We recorded what happened. We repeated it.

If your mental model of how ransomware interacts with backup infrastructure was formed more than a year ago, parts of this report will challenge it. If you work in security, infrastructure, or risk, that challenge is worth your time.

Amol Sarwate, Director of Security Research and REDLab, Cohesity

> **When real world ransomware reaches your backup systems, what actually happens?**

Report at a glance

Field	Value
Reporting period	January 2025 – May 2026
Dataset	53 ransomware families; 15 additional advisories covering ransomware, stealer, backdoor, and exploitation tradecraft
Testing environment	Cohesity REDLab, an air-gapped detonation facility with one-way logging
Platforms exercised	Cohesity DataProtect; Cohesity NetBackup
Workloads exercised	Windows Server, virtualized infrastructure, NAS, SQL, Oracle, Microsoft 365
Detection signals analyzed	Client Offline, Job Metadata, Image Entropy, Threat Scan, Rapid Threat Hunt
Framework	MITRE ATT&CK Enterprise v15
Author	Cohesity REDLab research team

Executive summary

Public reporting on ransomware has concentrated on endpoint encryption activity and the economic impact to victims. Reporting on operational outcomes at the data protection tier, the layer that determines whether a compromised organization can actually recover, remains limited. This report converts seventeen months of controlled detonation evidence into a quantitative view of how 53 contemporary ransomware families behave when they reach the backup systems designed to protect critical data.

Cohesity REDLab is the data protection industry's only dedicated source of actionable intelligence on how malware interacts with backup infrastructure. This team of experts runs an air-gapped facility that executes live malware against production-grade Cohesity DataProtect and Cohesity NetBackup deployments. Between January 2025 and May 2026, REDLab detonated 53 ransomware families and adjacent malware tools, mapped each one to the MITRE ATT&CK framework, and recorded backup-tier behavior under controlled conditions.

This report converts that proprietary detonation evidence into the industry's first quantitative resilience benchmark for enterprise backup infrastructure.

It's written for the people who answer for recovery when an incident hits: CISOs, CIOs, infrastructure and security leaders, architects, and the underwriters who price their risk. Read it as a field report. The methodology and the raw detonation records behind every claim are available on request for independent researchers, customers, and partners.

Key Judgments

Summary of REDLab's principal findings, January 2025 – May 2026

KJ-1. **REDLab observed** one of three outcomes in each malware detonation: backup job failed immediately; backup job completed but with ransomware-encrypted or corrupted data; backup job completed without impact. REDLab assesses that the most dangerous outcome is the successful backup job that preserves ransomware-encrypted or corrupted data. In this scenario, additional security capabilities and controls (like anomaly detection, threat protection, and others) are essential to proving recoverability.

KJ-2. **REDLab assesses** with high confidence that ransomware operators treat enterprise backup infrastructure as a primary target rather than a secondary objective. Across the reporting period, multiple commodity ransomware families tried to disrupt native operating system backup processes alongside primary workloads.

KJ-3. **REDLab observed** that Cohesity DataProtect backup operations continued without disruption across the ransomware families detonated in the testing environment. In every case, scheduled backups executed successfully and a recoverable snapshot was preserved.

KJ-4. **For Cohesity NetBackup**, whenever communication was disrupted in backup processes, REDLab observed at least one backup-tier anomaly signal.

KJ-5. **The MITRE ATT&CK Stealth tactic** (Defense Evasion pre v19.1) is the most heavily represented in the dataset. T1027 (Obfuscated Files or Information), T1070 (Indicator Removal), and T1036 (Masquerading) appear in a majority of the 20 strains detonated in 2026 that carry full technique-ID mappings; Process Injection (T1055) is also prominent. This partially aligns with REDLab's December 2025 in-the-wild MITRE technique frequency analysis, in which T1055 appeared in 58.55% of analyzed samples.

KJ-6. **REDLab anticipates** with moderate confidence that stealth tradecraft documented in 2025 advisories—including BRICKSTORM and GRIDTIDE—will be adopted by financially motivated ransomware operators over the next 12 to 18 months, extending dwell times across multiple backup cycles and increasing the operational importance of integrity validation against historical recovery points.

KJ-7. **The exploit-to-ransomware window** is collapsing, with campaigns like React2Shell, ToolShell, Shai-Hulud, CLOP/Oracle EBS, and Medusa weaponizing disclosed CVEs within hours to weeks.

Recommended mitigations

The following mitigations are derived from observed REDLab outcomes.

1. **Treat the backup tier as in-scope from initial compromise.** Threat models, IR runbooks, and tabletop exercises should assume the data protection plane is targeted within the first hours of an intrusion—not discovered as an afterthought during recovery.
2. **Enforce snapshot immutability.** REDLab observed no mutations of immutable snapshots by any strain in the dataset. No account reachable from the workload plane should be able to shorten retention, delete snapshots, or invoke an override of any kind.
3. **Isolate backup agent runtime from the protected workload OS.** Architectures that execute the backup agent inside the workload OS inherit the workload's blast radius.
4. **Surface backup-tier anomaly signals in the SOC console.** Job Metadata, Image Entropy, and Client Offline anomalies are operationally useful only when visible in the same console used to triage EDR and SIEM alerts. Siloed backup alerting isn't alerting; it's noise nobody sees.
5. **Implement scan-before-restore as a gating policy.** Recovery without remediation re-introduces malicious artifacts. Threat scanning or an equivalent "cleanliness" step should be required before restoring any workload following a confirmed compromise. During our tests, enforcing this policy was highly effective at preventing the restoration of snapshots containing malware.
6. **Establish routine threat hunting against historical recovery points.** Given the hundreds-of-days dwell times observed in BRICKSTORM and GRIDTIDE tradecraft, and the expectation that similar techniques will be adopted by financially motivated operators, sub-minute blast-radius scoping against older snapshots using current IOC feeds should be a baseline operational practice—not a reactive one.

How to read this report

- Section 1 explains REDLab's methodology and lab architecture.
- Section 2 maps the 53-strain dataset onto MITRE ATT&CK and an Anti-Backup Kill Chain that synthesizes the techniques attackers used to reach the backup tier.
- Section 3 details the platform's capabilities to signal the presence of malware—Client Offline, Job Metadata, Image Entropy, Threat Scan, and Rapid Threat Hunt—and which signals are fired against which strain category.
- Section 4 introduces the REDLab Backup Resilience Scorecard, a framework defenders can apply to their own architecture, and crosswalks it against common cyber insurance control language.
- Section 5 looks ahead, drawing on REDLab's seventeen months of advisories to outline the threats most likely to reshape backup architecture decisions in the next 12 to 18 months.

Section 1 — Methodology and lab architecture

This section documents REDLab's testing methodology and lab architecture so that the data presented in subsequent sections can be evaluated on its merits. Per REDLab practice, every detonation referenced in this report is logged with a timestamp, sample hash, platform version, workload configuration, and resulting anomaly records. Independent researchers, customers, and analysts may request the underlying detonation record for any strain referenced.

1.1 Lab architecture

REDLab pairs two Cohesity products with representative workload servers. Both products run current generally availability software, so published findings reflect deployed customer configurations.

Cohesity REDLab architecture

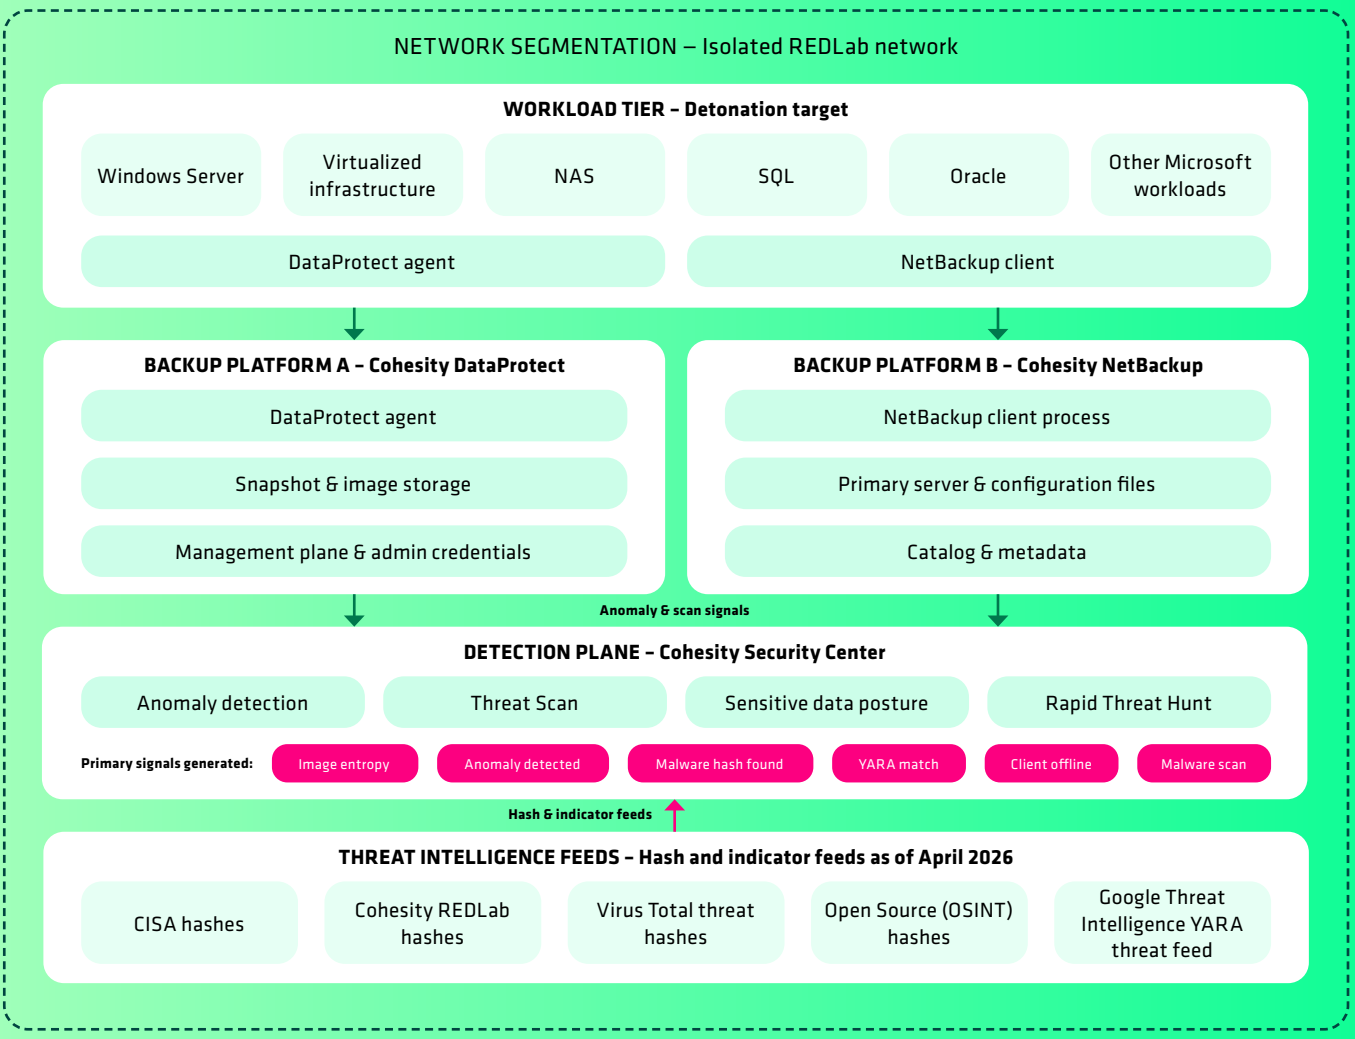

Layer	Component	Role in testing
Backup platform A	Cohesity DataProtect with DataProtect Agent on workload hosts	Measures agent operations, client integrity, snapshot integrity, anomaly signals, and recovery.
Backup platform B	Cohesity NetBackup with NetBackup Client on workload hosts	Measures agent operations, client integrity, snapshot integrity, anomaly signals and recovery.
Workload tier	Windows Server, virtualized infrastructure, NAS, SQL, Oracle, Microsoft 365	Detonation target. Encryption activity, file-extension changes, and ransom notes occur here.
Detection plane	Cohesity Security Center – Anomaly Detection, Threat Scan, Custom YARA scanner, Rapid Threat Hunt, and Google Threat Intelligence Sandbox	Surfaces backup-tier signals.
Threat intelligence feeds	CISA, Cohesity REDLab, Google Threat Intelligence, Open Source Intelligence (OSINT), Custom File Hashes, YARA Rules	Provides indicator and hash feeds from CISA, Google Threat Intelligence, OSINT, Cohesity REDLab feeds and 6500+ YARA rules.
Network segmentation	Air-gapped from corporate network and internet, with one-way logging out of the lab	Prevents detonation escape; ensures captured data is uncontaminated by external traffic. Specific details of each test are noted in **monthly newsletters**.

1.2 Attack surface mapping

Before each detonation, REDLab maps the attack path a ransomware operator would have to traverse to reach the backup tier. The mapping informs where to instrument for impact, and provides a shared vocabulary for the architectural realities of attacks that target the data protection plane. The layers of mapping include:

- **Backup clients on protected workloads.** The agent or client process running on the protected host. This is the most frequent point of first contact between ransomware and the backup tier, typically as a side effect of mass file encryption.
- **Primary backup servers and configuration files.** The control plane that the backup client communicates with. Multiple 2025–2026 strains in the dataset deliberately encrypted these files to disrupt recovery.
- **Snapshot and image storage.** The protected copy. Immutability and least-privilege management are intended to render this layer unreachable; REDLab tests whether attackers can circumvent those controls.
- **Management plane and administrative credentials.** Accounts and consoles that operate the backup platform. Credential harvesting and abuse of backup administrative accounts are recurring patterns across the dataset.
- **Catalog and metadata.** Records that describe protected assets and recovery paths. Tampering at this layer can poison recovery without modifying backup images.

1.3 Sample acquisition

REDLab acquires malware samples from CISA, Google Threat Intelligence, Ransomware.live, customer submissions, and its internal reverse-engineering pipeline. Malware detection signals are extracted, curated, and validated before they are integrated into daily feeds within the product. Every strain that customers can scan against has been tested or characterized in the REDLab environment.

We selected 53 of the commonly observed ransomware families and adjacent malware tools detonated between January 2025 and May 2026. The dataset is presented below by category.

Per-strain MITRE ATT&CK technique tables and outcome detail appear in Section 2 and in the Appendix.

Category	Families, tools, and vulnerabilities (REDLab 2025–2026 dataset)
Ransomware — high-volume / commodity	LockBit 5.0, Conti, DarkSide, Fog, HelloKitty, Locky (variant tracked as Lockis), DireWolf, Nuclear, MedusaLocker v3, Nitrogen, SafePay, BackLock
Ransomware — emerging / first observed in reporting period	Gagakick, Interlock, Gunra, Termite, CS173, NightSpire, Sinobi, Devman, Deep, WarLock, LeakDB, KaWaLocker, VanHelsing, Dharma/CrySiS (Find variant), Trigona (Trial_recovery variant), Payload, Kyber, VECT 2.0, Zollo (MedusaLocker variant), ShinySp1d3r
Ransomware — additional detonations (newsletter dataset expansion)	NetBackup only: MountLocker, CryptMIC, INC, Abyss, Cactus, LuckBit, Raf, Trinity, Scrypt, Snatch, Lynx, Prince, Nebula, SatanLocker. August 2025: MakOp, ViceSociety, Cephalus, Handala. March 2026: Sarcoma, Novalock, SilentAnonymous.
Stealers and access brokers tracked alongside the ransomware kill chain	Lumma Stealer, Agent Tesla
Espionage backdoors with ransomware-relevant tradecraft	BRICKSTORM, GRIDTIDE
Initial-access and exploitation campaigns covered in REDLab advisories	React2Shell (CVE-2025-55182); ToolShell (CVE-2025-49704 / 49706 / 53770) → WarLock / AK47C2; Shai-Hulud npm worm; CL0P / Oracle EBS (CVE-2025-61882, CVE-2025-61884); Medusa (CVE-2025-10035); CopyFail; AI Prompt Injection

Notes:

- REDLab distinguishes between the 53 ransomware families that constitute the core detonation dataset and the 15 additional advisories (cross-referenced in Section 2.5) that provide context for ransomware, stealer, espionage, and exploitation tradecraft. Quantitative claims in this report refer to the 53-family ransomware dataset unless otherwise specified.
- Cohesity DataProtect was added to REDLab testing in August 2025. The 14 strains detonated between January and July 2025 were therefore evaluated against Cohesity NetBackup only, while the 39 strains detonated from August 2025 through May 2026 were evaluated against both DataProtect and NetBackup. DataProtect outcomes and the Cohesity Security Center anomaly signals (Client Offline, Job Metadata, Image Entropy) refer to those 39 strains.
- The seventeen consecutive monthly newsletters from January 2025 through May 2026 broaden the dataset beyond the commodity and emerging families emphasized above. It now includes hacktivist-aligned tooling (Handala), an encryption-flaw wiper whose multi-nonce implementation prevents decryption even after payment (VECT 2.0), a post-quantum-branded operation targeting Windows and VMware ESXi (Kyber), VPN- and edge-exploitation-led double extortion (Cactus), and additional Phobos, Makop, Babuk, GlobeImposter, and HiddenTear lineages. Across this wider set, REDLab's central finding is unchanged: backup-tier anomaly signals and DataProtect agent isolation held regardless of endpoint-evasion sophistication.

1.4 Measurement points

Each detonation captures the same five data points for cross-strain comparability. The aggregated result of those five points underpins the Backup Resilience Scorecard introduced in Section 4. The measurement points are:

1. **DataProtect Agent and NetBackup Client operation status.** Did the next scheduled backup execute successfully, and did it produce a recoverable snapshot? Did the Client process continue to function and are the configuration files intact?
2. **Anomaly signals generated.** Which Cohesity anomaly signals fired - Client Offline, Job Metadata, Image Entropy – and at what severity, and on which jobs?
3. **Threat Scan and Rapid Threat Hunt response.** Did the Threat Scan identify malicious artifacts within backup images? Did Rapid Threat Hunt successfully scope affected clusters, objects, and file artifacts?
4. **Recovery validation.** Could a representative workload be restored from the most recent backup, to a known-clean state without reintroducing the malware?

1.5 Reproducibility

Every test conducted in REDLab is logged with timestamp, sample hash, platform version, workload configuration, and resulting anomaly records. Independent researchers, customers, and partners may request the detonation record for any strain referenced in this report. REDLab publishes a hash for each tested sample to enable independent verification of family attribution. Future editions of this report will disclose any change in lab configuration, platform version, or sample acquisition pipeline.

Section 2 — Attack pattern taxonomy

Section 1 covered how REDLab runs its tests. This is what the tests showed. Two patterns run through the 53-strain dataset. First, threat actors increasingly treat the backup tier as a target in its own right. Second, the techniques that beat endpoint defenses still leave fingerprints at the backup tier.

2.1 MITRE ATT&CK distribution

Every strain detonated in REDLab is mapped to the MITRE ATT&CK framework. Across the 53-family dataset, the most heavily populated tactics are Defense Evasion (labeled Stealth in the newsletters), Execution, Discovery, Privilege Escalation, and Persistence. Full technique-ID tables are published only for the 20 strains detonated in 2026; the 33 strains detonated across 2025 are characterized at the behavioral level, so the technique-frequency figures below are drawn from the 2026 MITRE-mapped subset. The most frequently documented techniques are T1027 (Obfuscated Files or Information, 14 of 20 strains), T1070 (Indicator Removal, 13 of 20), and T1036 (Masquerading, 11 of 20), with T1059 (Command and Scripting Interpreter), T1547 (Boot or Logon Autostart Execution), and T1129 (Shared Modules) each appearing in half. T1055 (Process Injection) and T1562 (Impair Defenses) appear in a substantial minority (8 of 20 and 6 of 20); software packing, Hidden Window, and Hide Artifacts each appear in only three to four of the 20. Across the 2025 detonations the dominant behavioral themes are consistent: data encryption for impact, recovery inhibition through shadow-copy deletion, defense evasion, and discovery.

This distribution is consistent with REDLab's December 2025 in-the-wild MITRE technique frequency analysis, in which T1055 (Process Injection) appeared in 58.55 percent of analyzed files, with T1071 (Application Layer Protocol) and T1036 (Masquerading) ranking second and third. REDLab assesses with moderate confidence that this overlap supports the dataset's external relevance.

Ransomware targets enterprise backup infrastructure

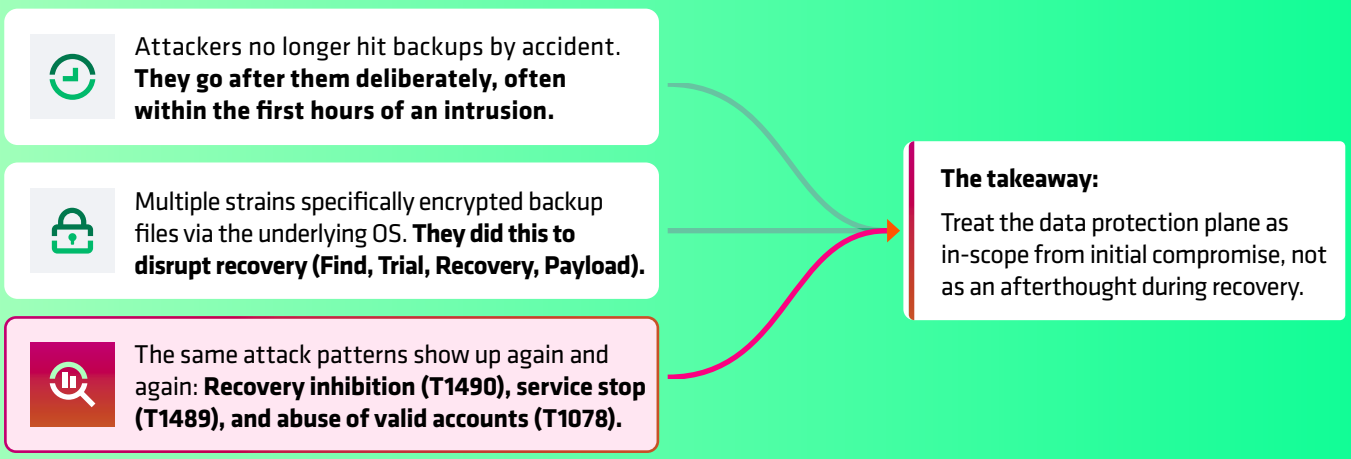

2.2 The Cohesity Anti-Backup Kill Chain

In addition to the ATT&CK techniques above, we observed a set of technique categories aimed at the backup tier—discovery, credential access, recovery-mechanism inhibition, and data/configuration impact. Not every strain exhibited every step, and the steps did not always occur in a fixed sequence, but we synthesize the pattern as a five-step Anti-Backup Kill Chain that captures the range of tradecraft observed.

Step	Action	MITRE ATT&CK	REDLab observation
1	**Backup process discovery**	T1057 (Process Discovery), T1518 (Software Discovery)	Strains enumerate running processes for backup agents, scheduling services, and management consoles prior to encryption.
2	**Credential harvesting targeting backup administrative accounts**	T1003 (OS Credential Dumping), T1539 (Steal Web Session Cookie)	Operators harvest stored credentials and session tokens with access to backup administrative consoles or service accounts.
3	**Recovery infrastructure sabotage**	T1490 (Inhibit System Recovery)	Abuse Windows-native binaries like vssadmin.exe, wbadmin.exe, and bcdedit.exe to target system restore infrastructure, purging Volume Shadow Copies, clearing backup catalogs, and disabling Windows Boot Configuration Data (BCD) recovery modes to trap the OS in an unrecoverable state. This technique was not observed in REDLab's 2024 dataset for commodity strains.
4	**File system and production data ransom**	T1485 (Data Destruction), T1486 (Data Encrypted for Impact)	Encrypt user-defined extensions on local NTFS volumes while simultaneously issuing commands to overwrite Master Boot Record (MBR) sectors or systematically zero-out critical data.
5	**Snapshot tamper attempts**	T1485 (Data Destruction), T1562 (Impair Defenses)	Where the platform exposed snapshot mutation paths, strains attempted to delete or rewrite recovery points before the next scheduled job. REDLab observed no successful mutation of immutable Cohesity snapshots across the dataset.

REDLab observed many MITRE ATT&CK Techniques, but there are several others relevant to backup and recovery operations that were not recorded in this research period. These techniques are listed in the Appendix.

2.3 Observed outcome categories

REDLab classifies each of the 39 strains detonated from August 2025 onward—when DataProtect and Cohesity Security Center anomaly detection were in the lab—into one of three outcome categories defined from the perspective of the backup platform. The 14 strains detonated January–July 2025 were assessed against NetBackup client integrity only.

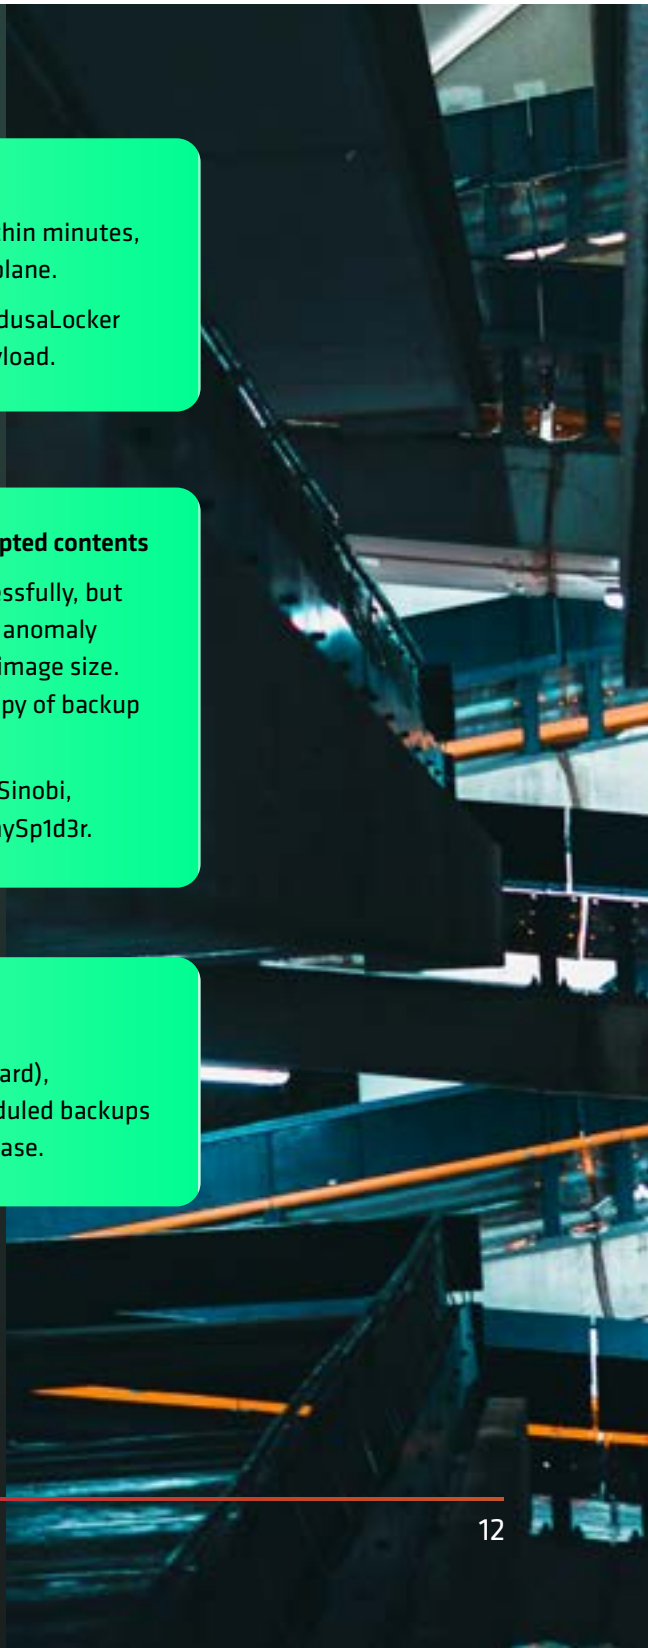

Category A

Communication disrupted

Ransomware disrupts communication; the Client Offline anomaly fires within minutes, and the responding analyst receives a notification in the Cohesity control plane.

Representative strains observed in Category A: Gagakick, LockBit 5.0, MedusaLocker v3, Nuclear, SafePay, Nitrogen, Find (Dharma), Trial_recovery (Trigona), Payload.

Category B

Client data encrypted; configuration intact; backup job succeeds with corrupted contents

Ransomware encrypts protected data. The next backup job executes successfully, but ingests encrypted ransomware artifacts. When enabled, the Job Metadata anomaly fires on deviations in file count, data transferred, deduplication ratio, and image size. The Image Entropy anomaly fires because encrypted blocks raise the entropy of backup images outside the established baseline.

Representative strains observed in Category B: Gunra, CS173, NightSpire, Sinobi, Devman, VanHelsing, DireWolf, HelloKitty, Kyber, VECT 2.0, Zollo, and ShinySp1d3r.

Category C

DataProtect backup successful; recovery point preserved

Across all the 39 strains detonated against DataProtect (August 2025 onward), DataProtect Agent backup operations continued without disruption. Scheduled backups executed successfully and a recoverable snapshot was preserved in every case.

2.4 Behavioral clustering — Endpoint evasion vs. backup-tier signal

Beating the endpoint does not mean beating the backup tier. The crews that spent the most effort evading endpoint detection still threw off clear signals once their encryption reached the backups.

- **BYOVD (Akira).** Akira affiliates load signed (but vulnerable) kernel drivers to disable endpoint detection and response prior to encryption. The technique succeeds at the endpoint, resulting in data being encrypted. The Image Entropy anomaly fires on the resulting backup.
- **Living-off-the-land double extortion (Qilin AGENDA / AGENDA.RUST).** Qilin uses Living off the Land (LOTL) techniques to evade detection. The group uses pre-installed administration tools like PowerShell, WMI, and PsExec to disable defenses, harvest credentials, and move laterally across compromises. Job Metadata anomalies were flagged based on unusual deviations in backup job attributes.
- **Restart Manager abuse (Devman).** Devman uses Windows Restart Manager to stop file-locking processes, framing the encryption step as a routine maintenance event for many security tools. The downstream backup ingests encrypted files, and Image Entropy fires.
- **LOLBin and LOLBAS execution chains (Nitrogen).** Nitrogen executes via signed Windows utilities and avoids dropping conventional binaries to disk. Backup integrity is unaffected by the execution mechanism; the encrypted state of protected files is what appears in the backup.
- **Rootkit-level concealment (DireWolf).** DireWolf uses T1542 (Pre-OS Boot), T1564 (Hide Artifacts), and T1562 (Impair Defenses) to suppress in-host telemetry. Endpoint visibility is lost; backup-tier signals continue to fire.

2.5 Cross-references to REDLab advisories

The 53 detonations are the empirical core of this report. Around them sit seventeen months of REDLab advisories that add context on the initial-access and stealth tradecraft feeding ransomware crews or shared with them. Each advisory below is cataloged in the Appendix with its publication date and indicators where available.

- **Update on Akira Ransomware BYOVD**. Exploitation of compromised firewall appliances and signed-but-vulnerable drivers to defeat endpoint detection and response prior to encryption.
- **Qilin LOTL Double Extortion**. AGENDA and AGENDA.RUST malware combined with credential theft, data exfiltration, and public-shaming pressure tactics.
- **Recovering from Kyber Ransomware on VMware ESXi Hypervisors**. Coordinated Windows-and-ESXi operation against the Kyber family; the ESXi variant's "post-quantum" Kyber1024 branding in fact relies on ChaCha8 with RSA-4096 key wrapping. Includes recovery guidance for encrypted /vmfs/volumes datastores.
- **SharePoint to Ransomware: ToolShell and AK47C2**. CVE-2025-49704 / 49706 / 53770 chained to deploy AK47C2 web shells and WarLock and AK47-linked ransomware.
- **CL0P Attack Against Oracle E-Business Suite**. Mass exploitation of CVE-2025-61882 and CVE-2025-61884 with extortion correspondence directed at executives.
- **React2Shell**. CVE-2025-55182, an unauthenticated remote code execution flaw in React Server Components, weaponized within 24 hours of disclosure.
- **Supply Chain Compromise – Shai-Hulud**. Self-replicating npm worm that harvests maintainer tokens and republishes malicious package versions.
- **BRICKSTORM Backdoor**. Modular stealth backdoor associated with long-dwell-time espionage; DNS-over-HTTPS command and control; tradecraft now observed in financially motivated intrusions.
- **GRIDTIDE**. C-based backdoor that uses Google Sheets API as covert command and control; observed across 53 organizations in 42 countries.
- **Tracing Lumma Stealer Through the Kill Chain**. Credential and access-broker behavior that consistently precedes ransomware deployment.
- **Agent Tesla Returns**. Credential-stealing remote access trojan distributed via phishing and exploit kits; included in CIS Top 10 malware ranking for Q3 2025.
- **Emerging Attack Patterns in the Wild**. REDLab in-the-wild MITRE technique frequency analysis (T1055 at 58.55%; T1071; T1036).
- **Medusa (CVE-2025-10035), CopyFail, AI Prompt Injection**. Newer advisories that inform the forward outlook in Section 5.

Section 3 — Detection signals from the backup pipeline

For this research, the backup platform's most useful role is as a sensor. It sees every protected file, every scheduled job, every shift in deduplication ratio. It records state changes endpoint tools miss, and it sits outside the endpoint-evasion arms race: it describes the resulting state of your data, not the runtime behavior of a process the malware is busy concealing. Endpoint tooling tells you about now. The backup tier can tell you about then.

Cohesity Data Cloud exposes five backup-tier detection signals. This section describes each signal, the conditions under which it fires, and the strain categories it detected in the REDLab dataset.

3.1 Client offline anomaly

The Client Offline anomaly fires when platform health checks identify that a backup client has stopped communicating in a pattern consistent with ransomware-induced failure rather than scheduled maintenance. In REDLab testing, Client Offline fired against every Category A strain—those that disrupted communication of the backup workload and/or encrypted primary data. Severity is reported as High in Security Center. The failed job ID and asset name are surfaced for triage.

Representative strains: Find, Trial_recovery, Payload (April 2026 detonations); Gagakick, LockBit 5.0, MedusaLocker v3, Nuclear, SafePay, Nitrogen (earlier detonations in the reporting period).

3.2 Job metadata anomaly

The Job Metadata anomaly is a machine-learning-driven detection that fires on deviations across four backup-job attributes: file count, data transferred, deduplication ratio, and image size. The signal does not depend on signature lookup or behavioral analysis of the malware. It depends on the statistical distribution of the customer's own backup jobs, learned over time. When ransomware encrypts protected files, those attributes shift outside the learned distribution and the anomaly fires, typically with three or four sub-anomalies surfaced together.

Representative strains: VanHelsing, DireWolf, HelloKitty (February–April 2026); Gunra, CS173, NightSpire, Sinobi, Devman (earlier in the reporting period).

3.3 Image entropy anomaly

The Image Entropy anomaly is REDLab's most evasion-resistant detection signal. It operates on the physical property of encrypted data: when a file is encrypted, its entropy rises, and a backup image of that file inherits the higher entropy. Image Entropy compares each image's entropy against the established baseline and triggers when deviation exceeds the configured policy threshold. The mechanism is agnostic to how malware unpacked, hid its loader, or evaded endpoint detection. Image Entropy fired across REDLab testing regardless of whether the strain employed BYOVD, LOTL, packing, or other in-host evasion techniques.

Representative strains: VanHelsing, DireWolf, HelloKitty, Kyber, VECT 2.0, Zollo, ShinySp1d3r, and every other Category B strain in the dataset.

3.4 Threat Scan

Threat Scan performs pre-restore inspection of backup snapshots against the REDLab-curated threat library and customer-supplied YARA rules. The control can help identify the root cause of the incident to improve trust in remediation and recovery. Threat Scan is configurable to run on a schedule, on demand, or as a gating step in the recovery workflow. In REDLab testing, Threat Scan correctly identified malicious artifacts within backup images for every Category A and Category B strain in the dataset.

3.5 Rapid Threat Hunt

Rapid Threat Hunt searches across clusters, objects, and file artifacts for known indicator-of-compromise hashes. It draws on five hash feeds in parallel: CISA, Cohesity REDLab, Google Threat Intelligence, Open Source (curated from community sources), and Custom File Hashes uploaded by the customer. A defender can use a single SHA-256 hash to perform a complete blast-radius assessment without recalling backups.

In a representative April 2026 detonation against the Payload ransomware family, REDLab triggered Rapid Threat Hunt with the Payload SHA-256 hash. The hunt scanned 11 production-scale clusters, identified 5 affected clusters, matched 1 hash, surfaced 16 affected objects, and completed in 18 seconds. REDLab observed comparable results in equivalent runs against VanHelsing, Find, and Trial_recovery, as well as the Kyber, VECT 2.0, Zollo, and ShinySp1d3r families.

What matters here is recovery economics, not detection for its own sake. Without automated blast-radius scoping, working out which recovery points are safe to restore is a manual, serial slog: recall a backup, scan it, repeat, for hours or days while the business stays down. A sub-20-second query against a single indicator collapses that work, so responders restore only verified-clean recovery points on the first attempt. That is what shortens the recovery time objective (RTO) and the downtime cost boards and cyber insurers actually measure.

Hash feed	Source
CISA	U.S. Cybersecurity and Infrastructure Security Agency
Cohesity REDLab	Cohesity REDLab curated, lab-validated feed
Google Threat Intelligence	Google Threat Intelligence (Mandiant)
Open Source	Aggregated open-source threat intelligence (MalwareBazaar, Ransomware.live)
Custom File Hashes	Customer-supplied YARA and hash uploads

Hash feed counts update daily.

3.6 Signal coverage by outcome category

The summary below maps outcome categories to the signals that fire and to the operational implication for security operations runbook design:

Outcome category	Primary signal	Secondary signal	Operational implication
A – Communication disrupted	Client Offline (High)	Threat Scan; Rapid Threat Hunt	Failure mode is unambiguous; alerting is immediate.
B – Data encrypted; configuration intact; job succeeds with corrupted contents	Job Metadata (High); Image Entropy (High)	Threat Scan; Rapid Threat Hunt	Highest risk category for organizations without backup-tier anomaly detection. Backup appears successful but contents are corrupted; backup-tier signals are essential.
C –DataProtect backup successful; recovery point preserved	Recovery validated	Threat Scan; Rapid Threat Hunt as defense in depth	Clean recovery point is the operative defense; signals confirm rather than alert.

Section 4 — Backup resilience scorecard and architectural guidance

This section turns the dataset into something you can act on: a scorecard and a short set of design principles. Both are built for the leaders evaluating their own data protection architecture, and for the conversations they must hold with boards, regulators, and cyber insurers.

4.1 The REDLab backup resilience scorecard

The scorecard rates a data protection architecture across four dimensions, each scored from 1 (initial) to 5 (optimized). The dimensions correspond directly to outcome categories observed in REDLab testing.

Dimension	What it measures	1 – Initial	5 – Optimized
Anti-backup TTP resistance	Architectural resistance to the techniques in REDLab's Anti-Backup Kill Chain—process discovery, credential harvesting, shadow-copy deletion, configuration tampering, and snapshot tampering.	Backup administrative accounts share credentials with workload administrative accounts; configuration files writable by workload accounts; snapshots mutable.	Compute and storage separated; snapshots immutable; least-privilege management plane; using VMware based backups if possible.
Detection coverage	Backup-tier signal availability for detecting ransomware-driven backup deviation prior to restore.	No anomaly engine; no entropy analysis; no indicator-of-compromise hash hunting.	Client Offline, Job Metadata, Image Entropy, Threat Scan, and Rapid Threat Hunt integrated with the security operations console.
Recovery Remediation Validation	Architectural assurance that a recovery point is clean prior to use in restore.	Existence of backups is the only assurance.	Mandatory pre-restore Threat Scan; YARA and hash indicator validation; quarantine workflow for suspicious snapshots.
Time-to-validated-recovery	Elapsed time from initial anomaly signal to validated clean restore.	Days to weeks; backups recalled without integrity validation.	Minutes to hours; Rapid Threat Hunt scopes blast radius in seconds; recovery point validated by Threat Scan prior to restore.

4.2 Architectural principles

Four design principles follow from what we measured. Each rests on a specific result in the dataset rather than on best-practice boilerplate.

1. **Snapshot immutability is required, not optional.** Snapshots subject to deletion or rewriting by the same plane that protects production data constitute a single point of failure. REDLab observed no successful mutation of immutable Cohesity snapshots by any strain in the dataset.

2. **The backup management plane must be isolated.** Backup administrative accounts must not share authentication or authorization with workload administrative accounts. Strains in the dataset that achieved backup-tier impact consistently did so through abuse of shared credentials rather than by compromising the backup platform directly.

3. **Backup anomalies must be integrated into the security operations console.** Job Metadata, Image Entropy, and Client Offline anomalies are operationally useful only when displayed in the same console used to triage endpoint detection and response and security information and event management alerts.

4. **Scan-before-restore must be enforced as policy.** Recovery without integrity validation reintroduces malicious artifacts. Threat Scan and Rapid Threat Hunt should be configured as gating steps in the recovery workflow following any confirmed compromise.

4.3 Crosswalk to cyber insurance control language

The scorecard dimensions map to control questions that appear consistently in cyber insurance application questionnaires across major carriers. The mapping below is presented for use by Chief Information Security Officers and Risk Officers preparing renewal materials.

Scorecard dimension	Common underwriter control language	Where REDLab data supports the control
Anti-Backup TTP Resistance	Immutable backups; segregated backup credentials; least-privilege backup administration.	Empirical evidence that immutability and credential isolation defeat the Anti-Backup Kill Chain in the 53-strain dataset.
Detection Coverage	Anomaly detection on backup data; alerting integrated with security operations.	Per-strain documentation of which signals fired and at what latency.
Recovery Remediation Validation	Validated, tested recovery; clean recovery procedures.	Threat Scan and Rapid Threat Hunt can accurately detect threats targeting and compromising backups. Demonstrates that threat hunting on backups can help ensure that recoveries can use clean snapshots.
Time-to-Validated-Recovery	Documented recovery time objective; tabletop exercises.	Sub-20-second Rapid Threat Hunt blast-radius scoping; sub-minute anomaly signals across the dataset. The faster threats can be detected, the faster that a clean recovery can occur.

4.4 Self-assessment checklist

The checklist below maps each scorecard dimension to a binary yes-or-no question intended for use by chief information security officers in preparation for board reporting or cyber insurance renewal.

- Are all backup snapshots immutable for the policy retention window, with no override path invocable from a workload-side account?
- Are backup administrative credentials and consoles isolated from workload administrative credentials and consoles, with separate authentication and authorization paths?
- Does the backup platform produce anomaly signals on file count, data transferred, deduplication ratio, image size, image entropy, and client connectivity, and are those signals visible in the security operations center?
- Is Threat Scan or equivalent recovery remediation configured as a gating step in the recovery workflow following confirmed compromise?
- Can the security team scope blast radius across the entire backup footprint within minutes of receiving a single SHA-256 hash?
- Are these controls documented in language that maps to cyber insurance application questionnaires and not solely to the technical runbook?

Section 5 — Forward outlook, 12 to 18 months

The dataset is a snapshot in time. Seventeen months of advisories point to where things are heading. What follows is our read on the next 12 to 18 months, drawn from the campaigns we analyzed across the period, with a confidence level attached to each call.

5.1 Continuation of defense evasion dominance

REDLab's December 2025 in-the-wild MITRE technique analysis identified three techniques that hold consistent share across the year—T1055 (Process Injection), T1071 (Application Layer Protocol), and T1036 (Masquerading). The December dataset showed sharper increases in T1119 (Automated Collection), T1056 (Input Capture), and continued elevation of T1055. REDLab assesses with moderate-to-high confidence that this distribution is consistent with a tooling shift among ransomware affiliates toward greater automation, broader credential capture, and increased reliance on in-process execution. Defenders should expect these techniques to be weaponized at scale through 2026.

5.2 Compression of the exploit-to-ransomware window

Multiple campaigns documented during the reporting period demonstrate compression of the interval between public vulnerability disclosure and ransomware deployment. React2Shell (CVE-2025-55182) was disclosed on December 3, 2025, and exploited within 24 hours. The ToolShell chain (CVE-2025-49704 / 49706 / 53770) was used to deploy WarLock ransomware in July 2025. The Shai-Hulud npm worm trojanized maintainer accounts at machine speed. The CL0P / Oracle E-Business Suite campaign (CVE-2025-61882 / 61884) escalated to extortion within weeks of disclosure. Medusa exploitation of CVE-2025-10035 followed the same pattern. The Copy Fail Linux kernel privilege-escalation flaw (CVE-2026-31431), disclosed on April 29, 2026, extends the same dynamic to the source systems that feed the backup tier: a deterministic, race-free local exploit that grants root on essentially every major Linux distribution shipped since 2017 and is escapable from an unprivileged container. Because the corruption occurs only in the page cache, on-disk file-integrity monitoring does not detect it—but any backdoors or modified binaries an attacker writes after gaining root are captured faithfully in the next snapshot, so an unpatched, compromised Linux host becomes a compromised recovery point.

REDLab assesses with high confidence that this compression makes backup retention and integrity validation an operational concern of the vulnerability management program. The longer a defender takes to identify compromise, the more recovery points are at risk of containing compromised credentials, malware, misconfigurations, and other persistence mechanisms. Frontier models are accelerating this reality for cyber defenders: vulnerabilities are going to be exploited faster than ever, you need to ensure your snapshots are clean and don't contain backdoors that lead to reinfection.

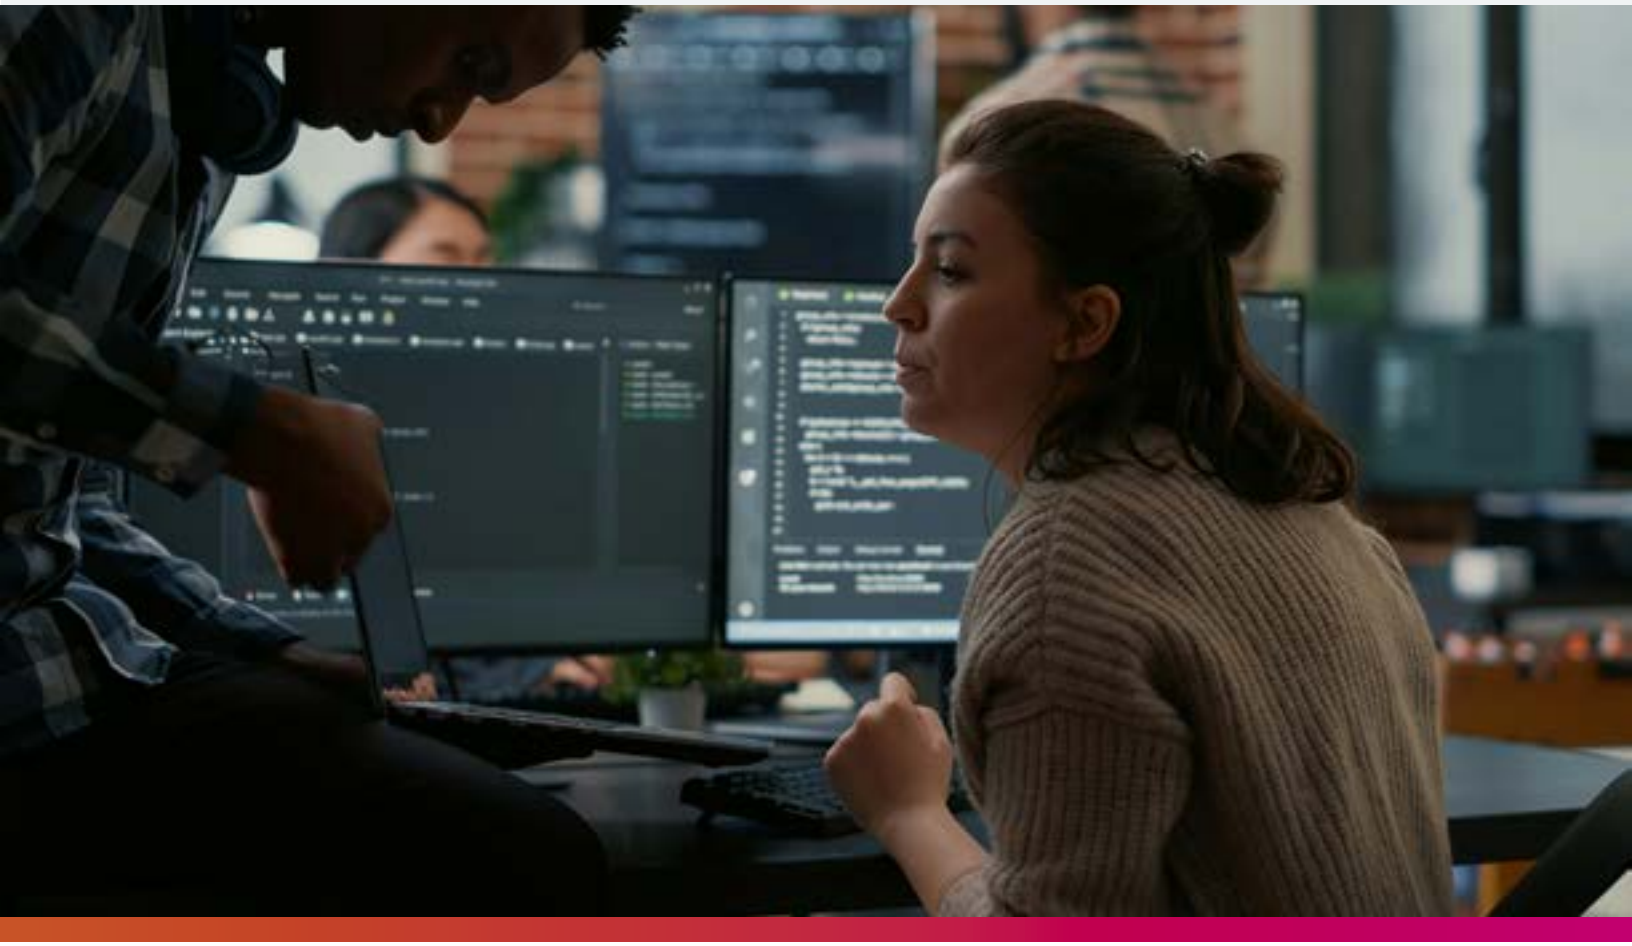

5.3 Supply chain as initial access vector

Shai-Hulud and the CL0P / Oracle E-Business Suite campaign were not one-offs. We assess with moderate-to-high confidence that supply chain compromise, whether at the package registry, the SaaS layer, or the build system, is now a leading way in and a route for lateral movement. Build-system backups deserve particular attention in Threat Scan and YARA-based validation.

5.4 AI-assisted malware authorship

REDLab advisories on CopyFail and AI Prompt Injection, together with **Palo Alto Unit 42's assessment that a large language model** likely contributed to the Shai-Hulud worm script, indicate emerging AI-assisted malware authorship and AI-targeting attacks against AI systems themselves. REDLab assesses with moderate confidence that detection signals grounded in the physical state of data—entropy, metadata distribution, and integrity validation—retain effectiveness against AI-authored malware to a greater degree than signature-based pattern matching. The same logic applies to the post-quantum branding now appearing in the ransomware dataset. Kyber markets a Kyber1024 key-encapsulation scheme, yet its ESXi variant in fact relies on ChaCha8 with RSA-4096 wrapping—a reminder that cryptographic claims in ransom notes are a marketing surface, not a reliable indicator of capability. Whether the underlying cipher is classical or post-quantum, the entropy and metadata signatures of encrypted backup images remain observable at the backup tier, and Kyber confirmed this in testing by triggering the NetBackup Image Entropy anomaly while DataProtect VMware backups continued without disruption.

Section 6 — REDLab testing roadmap

The REDLab testing roadmap is built around four threat shifts we expect to define the 2026–2027 attack landscape.

- **Hypervisor-centric ransomware continues to escalate.** Attackers are increasingly targeting virtualization infrastructure directly—taking down entire server fleets in a single strike—while cross-platform campaigns detonate simultaneously across Windows, Linux, and ESXi environments. Kyber, one of the families in the dataset, is a concrete instance of this trajectory: a coordinated ESXi-and-Windows operation that encrypts /vmfs/volumes datastores and defaces the management interface while a companion Windows payload destroys local Volume Shadow Copies and backup services.
- **AI-accelerated attack chains are compressing timelines from days to hours.** This includes polymorphic payloads, kernel-level evasion that disables endpoint defenses before encryption starts, and early-stage autonomous propagation that reduces the window for defenders to respond.
- **Identity and cloud layers are replacing the endpoint as the primary attack surface.** API-based ransomware now targets cloud backup storage directly. Credential-theft chains compromise backup administrators before any payload is deployed. Encryption-free extortion models exfiltrate data silently, bypassing file-based detection entirely.
- **Emerging cryptographic and destructive threats round out the picture.** Post-quantum algorithms are appearing in active ransomware operations, and wiper campaigns are designed to destroy backup infrastructure rather than end-user data. Both shifts are represented in the dataset: Kyber advertises a Kyber1024 post-quantum scheme, while VECT 2.0 behaves as a de facto wiper—a flaw in its multi-nonce ChaCha20 implementation renders files larger than roughly 131 KB permanently unrecoverable even when the correct key is supplied, so payment cannot guarantee restoration.

REDLab will monitor the threat landscape continuously and update this roadmap throughout the year as new attack techniques and threat actors emerge. Customers, partners, and researchers may submit ransomware samples and indicators for inclusion through the channels listed on the **REDLab landing page.**

Appendix

Strain reference summary

The table below summarizes selected strains discussed in this report. Per-strain technique tables, hash indicators of compromise, and full detonation telemetry are maintained in the REDLab newsletter archive and are available on request.

Strain	Family / Group	First REDLab detonation	Outcome category	Notable tradecraft
Kyber	Kyber Ransomware Family	May 2026	B	Claimed post-quantum crypto (Kyber1024 + X25519); ESXi variant in fact ChaCha8 with RSA-4096. Coordinated Windows + VMware ESXi targeting; .#~~~ extension; READ_ME_NOW note; Rust-based modular RaaS.
VECT 2.0	Vect Ransomware Family	May 2026	B	Multi-nonce ChaCha20 flaw renders files larger than ~131 KB permanently unrecoverable even with the key, behaving as a de facto wiper; !!!_READ_ME_!!! note.
Zollo	MedusaLocker family	May 2026	B	MedusaLocker variant; hybrid RSA + AES with double extortion; incrementing .zollo6 / .zollo10 extensions; READ_NOTE.html note.
ShinySp1d3r	ShinyHunters Ransomware Group	May 2026	B	Identity-led access (SaaS account takeover, MFA fatigue, helpdesk impersonation); randomly generated extensions; R3ADME_1Vks5fYe.txt note; Tox-based contact.
VanHelsing	VanHelsing Ransomware Group	April 2026	B	Targets local and network shares; .vanhelsing extension; README.txt ransom note.
Find	Dharma / CrySiS	April 2026	A	AES + RSA per-file keys; .FIND extension; info.txt ransom note.
Trial_recovery	Trigona	April 2026	A	Encrypts local, network, and mounted volumes; .-encrypted extension; how_to_decrypt.txt.

Strain	Family / Group	First REDLab detonation	Outcome category	Notable tradecraft
Payload	Payload Ransomware Group	April 2026	A	Targets ESXi and virtual disks; Babuk-derived cryptography; RECOVER_ payload.txt.
DireWolf	DireWolf Hacking Group	February 2026	B	Rootkit-level concealment; Pre-OS Boot; Hide Artifacts; .direwolf extension.
HelloKitty	HelloKitty	February 2026	B	Re-emerged in 2025; .crypted extension; targets Windows enterprise environments.
KaWaLocker	KaWaLocker	February 2026	B	Profile-aligned with DireWolf-style Windows enterprise targeting.
Locky (Lockis variant)	Locky family	February 2026	B	Long-running Locky lineage; commodity tooling with continued operational impact.
WarLock / AK47C2	WarLock and AK47-linked	Q3 2025 advisory and ongoing	A	Delivered via ToolShell SharePoint exploitation chain (CVE-2025-49704 / 49706 / 53770).
LockBit 5.0	LockBit	2025	A	Continued anti-backup focus including configuration file targeting.
MedusaLocker v3, Nuclear, SafePay, Nitrogen, Gagakick	Multiple families	2025	A	Active backup-process termination and configuration tampering.
Gunra, CS173, NightSpire, Sinobi, Devman	Multiple families	2025	B	Encrypt protected data without disabling NetBackup Client; entropy and metadata anomalies fire.
Conti, DarkSide, Fog, BackLock	Multiple families	2025	A	Established high-volume and RaaS families included for dataset breadth; recovery inhibition and configuration-file targeting consistent with the Anti-Backup Kill Chain.

Strain	Family / Group	First REDLab detonation	Outcome category	Notable tradecraft
Interlock, Termite, Deep, LeakDB	Multiple families	2025	B	Emerging families first observed in the reporting period; encrypt protected data for impact, surfacing Image Entropy and Job Metadata anomalies at the backup tier.
MountLocker, CryptMIC, INC, Abyss, Cactus, LuckBit, Raf, Trinity, Scrypt, Snatch, Lynx, Prince, Nebula, SatanLocker	Multiple families	Jan–Jul 2025	NetBackup Only	Detonated against NetBackup before DataProtect joined REDLab (August 2025). Commodity and emerging families; double extortion, shadow-copy deletion, and recovery inhibition.
MakOp, ViceSociety, Cephalus, Handala	Multiple families	August 2025	B	First month of dual DataProtect + NetBackup testing. Includes hacktivist-aligned Handala (re-tested March 2026); DataProtect backups preserved.
Sarcoma, Novalock, Silent-Anonymous	Sarcoma; GlobeImposter; HiddenTear	March 2026	B	DataProtect Agent backups unaffected; NetBackup Job Metadata and Image Entropy anomalies fired.

Advisory cross-references — observed in REDLab advisories; not counted in the 53-strain detonation dataset				
Akira (BYOVD update)	Akira	Q4 2025 advisory and ongoing	A	Compromised firewalls for initial access; signed-but-vulnerable drivers to defeat endpoint detection and response.
Qilin (AGENDA / AGENDA. RUST)	Qilin	Q4 2025 advisory and ongoing	A	Living-off-the-land double extortion; high-profile victims including Asahi and VW Group France.

Additional MITRE ATT&CK techniques

T1490 – Inhibit System Recovery Adversaries delete or disable recovery mechanisms - Volume Shadow Copies, backup catalogs, restore points - to prevent victims from recovering without paying.

T1078 – Valid Accounts Adversaries obtain and abuse legitimate credentials to access backup systems and management consoles, bypassing security controls that would flag malicious tooling.

T1098 – Account Manipulation Adversaries modify account permissions, add unauthorized credentials, or alter roles within backup platforms, cloud tenants, and identity control planes to maintain persistent access.

T1562.001 – Impair Defenses: Disable or Modify Tools Adversaries disable or tamper with security software—backup anomaly detection, malware scanning, alerting, and audit logging—to operate undetected before and during an attack.

T1489 – Service Stop Adversaries stop backup agents, VSS, database services, and hypervisor processes to ensure files are unlocked and available for encryption or deletion.

T1485 / T1486 / T1561 – Data Destruction / Data Encrypted for Impact / Disk Wipe Adversaries destroy, encrypt, or wipe backup data directly—sometimes using management interfaces such as BMC or MDM—to eliminate recovery options alongside production data.

T1021 – Remote Services *(sub-techniques: .001 RDP, .002 SMB, .004 SSH, .007 Cloud Services)* Adversaries use remote access protocols to move laterally into backup servers and storage infrastructure using valid or compromised credentials.

T1552 / T1003 – Unsecured Credentials / OS Credential Dumping Backup systems frequently store credentials in configuration files or local databases, making them a high-value target for credential harvesting that enables broader infrastructure compromise.

About Cohesity REDLab

Cohesity REDLab is the data protection industry's only dedicated source of actionable intelligence on how malware interacts with backup infrastructure. This team of experts runs an air-gapped facility that executes live malware against production-grade Cohesity DataProtect and Cohesity NetBackup deployments. REDLab maps every detonation to MITRE ATT&CK, and feeds the resulting findings back into Cohesity's threat intelligence capabilities. The threat library that powers Threat Scan and Rapid Threat Hunt in customer environments is curated by REDLab.

REDLab publishes a monthly newsletter and a continuously updated stream of advisories at **www.cohesity.com/redlab** and **www.cohesity.com/trust/redlab/advisories**. Press representatives, analysts, and customers seeking to validate data in this report, request custom strain testing against their own environments, or schedule an executive briefing may contact the REDLab team through their Cohesity account representative.

About this report

This is the first edition of the REDLab Threat Intelligence Report. The dataset covered in this edition spans detonations conducted between January 2025 and May 2026. Subsequent editions will refresh the dataset, broaden workload coverage, and incorporate new strains and tradecraft as they emerge. Changes to methodology, lab configuration, and dataset composition will be disclosed explicitly in each edition.

Comments, corrections, sample submissions, and briefing requests should be directed through Cohesity account teams or to the REDLab landing page at **cohesity.com/redlab**.

Disclosure and Trademarks