TRACE3

FORENSIC LOG REQUIREMENTS

Prepared By

Patti Hallock

Regional Solutions Architect

Infrastructure Solutions

Midwest

Issued

Date Issued:

03/06/2026

Document Version:

Final

Table of Contents

Introduction

For CISOs, the real cost of an incident often comes after detection, when the question is no longer whether an incident occurred but if the scope, impact, and exposure can be proven quickly. Too often, investigations stall because the right signals weren't collected, retained long enough, or correlated across systems. A single alert might tell you something suspicious occurred, but forensics is about reconstructing the full story: which account was used, from where, on what device, what executed, what changed, and what data was touched or taken.

A thorough forensic investigation depends on visibility across the entire attack path. Endpoint telemetry shows process execution and persistence. Identity and directory logs reveal sign-ins, privilege changes, and token abuse. Network signals expose command-and-control, lateral movement, and exfiltration paths. Email and collaboration logs often capture the initial foothold and the data escape route. Cloud control-plane audits prove who changed what in your environment, while data-plane logs show what was accessed.

In this paper, we'll break down the log sources and signals you need for real-world incident forensics, how they map to common attacker behaviors, and the practical requirements (e.g., time sync, centralized collection, and retention) that make the difference between a confident conclusion and uncertainty.

Required Logs/Signals

The required signals span several domains that, taken together, allow you to understand attacker intent and business impact without guesswork. Endpoint data shows what executed and what changed on a system. Identity telemetry tells you who authenticated and which privileges or tokens were used. Network activity reveals where systems communicated and what data likely moved. Email, SaaS, and cloud audit logs then close the loop by showing how access was gained and which data was accessed or shared. The goal isn't to collect "the kitchen sink", it's to collect the right, high-fidelity signals that can be correlated into a single timeline using durable identifiers (user, device, session/token, and IP) and retained long enough to cover realistic dwell time.

Endpoint

Start with endpoint telemetry because it provides the most direct evidence of execution and persistence. Forensics-grade endpoint signals include full process lineage (parent/child), command-line and script content where available, module loads or injection indicators, file and registry/config changes, local logon events, and outbound connections from the host. In large environments, this must be paired with strong asset identity (device IDs that persist through reimaging where possible), coverage across servers and workstations, and off-host log shipping so evidence survives an attacker trying to wipe traces. Without endpoint visibility, investigations collapse without being able to prove what actually ran or whether persistence remains.

Identity

Next are identity and directory signals, which are essential because so many enterprise incidents are initiated by breached but valid accounts. These signals should capture authentication outcomes

(success/failure), MFA results, conditional access or risk decisions, and the contextual attributes that matter in investigations (source IP, device posture, location, user agent, and session/token indicators). Pair that with audit trails for directory changes (e.g., new accounts, password resets, MFA method changes, OAuth consent grants, service principals, role/group membership updates, and privileged access activations) so you can distinguish normal admin activity from adversary-driven control. In practice, identity logs are what let you answer the executive questions quickly, which accounts were used, what access they had, and whether the attacker could still be operating with stolen tokens or newly granted privileges.

Network

Network signals are the connective tissue that lets you confirm where activity traveled, especially when endpoints are missing coverage or attackers live off the land. At minimum, you want DNS logs (queries/responses), web proxy/secure web gateway logs (URLs, SNI, user/device mapping, uploads), and firewall logs (allow/deny, NAT, egress destinations) because these three routinely expose command-and-control, payload retrieval, and exfiltration paths.

You can add NetFlow/IPFIX to quantify unusual data movement and highlight beaconing patterns, and layer in IDS/IPS/NDR detections to catch protocol anomalies and lateral movement that may never produce a high-confidence endpoint alert. In large enterprises, the key is consistency, meaning if only some segments route through the proxy, or DNS is split without central logging, you create blind spots that attackers will naturally exploit.

Email/SaaS/Cloud

Email, SaaS, and cloud signals are where enterprise incidents often begin and where the real impact (data access and sharing) becomes visible.

For email, you want message trace and gateway telemetry (delivery, attachments, URL rewrites/clicks, authentication results) plus mailbox audit signals that capture rule creation, forwarding, delegation, OAuth app access, and mailbox export/search behaviors.

For collaboration platforms and SaaS, prioritize audit trails for file access, mass download, external sharing, permission changes, link creation, and admin actions because “quiet exfiltration” is frequently just a legitimate download at scale.

In cloud (IaaS/PaaS), split your requirements into control-plane (who changed IAM, policies, logging, security groups, keys, deployments) and data-plane (who read/downloaded objects, queried data stores, accessed secrets). Control-plane logs tell you how the attacker expanded capability; data-plane logs tell you whether data exposure likely occurred.

Centralization of Logs

In an enterprise, forensic readiness depends on centralizing these signals into a SIEM or security data lake so they can be correlated quickly and consistently across identity, endpoint, network, email/SaaS, and cloud. The value isn’t just storage and retention, it’s correlation and fidelity by normalizing key identifiers (user, device, hostname, IP, tenant/subscription, session/token where available), enforcing consistent timestamps, and retaining data long enough to connect early-stage access to later-stage impact. Done well, this also hardens you against evidence loss.

Threat actors can wipe local logs or disable point solutions, but they have a much harder time erasing off-host, access-controlled, and ideally immutable centralized telemetry. The practical outcome is faster scoping, higher-confidence containment decisions, and the ability to move from isolated alerts to high-signal detections. This is so you can move from "low-and-slow" uncertainty to confident, evidence-backed detection that ties scattered signals into a single, actionable narrative.

Data Retention

Data retention is where some enterprise investigations quietly fail because the evidence window was too short to reconstruct scope with confidence. Attackers don't operate on your ticketing timeline as they can blend into normal admin and user activity, and you often discover the true entry point weeks (or months) after the first evidence of malicious behavior. NIST frames log management as a planning discipline (what to collect, protect, and retain) precisely because logs are only useful if they're still there when you need them and if they're protected from tampering. (NIST Publications)

A practical enterprise retention strategy usually needs tiers. Keep "hot" data long enough to support rapid scoping and threat hunting, then move older data into lower-cost archival that's still retrievable for legal, regulatory, or deep forensics.

Security leaders are caught off-guard when they realize different platforms keep different windows, and "retained" doesn't always mean easily queried. For example, Microsoft Defender XDR retains data for 180 days in the portal, but Advanced Hunting query access is 30 days unless you stream data to Sentinel (or another store) to extend investigative reach.

To improve forensic readiness before an incident, enterprises should treat retention as a risk decision backed by measurable requirements, such as expected dwell time, regulatory obligations, and business impact.

The playbook is straightforward:

1. Define your minimum investigation window (what you want searchable vs. merely retrievable) and ensure alignment with regulatory and contractual expectations.
2. Align retention across identity, endpoint, network, email/SaaS, and cloud audit logs.
3. Centralize storage off-host with strong access controls.
4. Add immutability where feasible to reduce evidence-tampering risk.

In practice, retention should be treated as a compliance and forensic control, not just a storage decision. That means defining policy minimums by data type, keeping key logs easily searchable, and using immutable storage where scrutiny is most likely. The final step is often overlooked, which is confirming you can retrieve and correlate that data quickly when an incident is unfolding.

Regulatory Requirements

Regulatory requirements can drive your retention strategy because they often specify minimum retention windows and, in some cases, how records must be preserved. For example, NYDFS 23 NYCRR 500.6 requires covered entities to maintain audit-trail records for no fewer than five years to support reconstruction and detection of cybersecurity events. In the payment space, PCI DSS v4.0.1 (section 10.5, page 252) requires audit-log history be retained for at least 12 months, with the most recent three months immediately available for analysis. And for securities firms, SEC Exchange Act

Rule 17a-4 includes multi-year record preservation requirements (commonly three or six years depending on record type) and has historically required preservation in a non-rewriteable, non-erasable manner for electronic records.

Payment Card Industry - PCI

PCI adds a handful of signals that are easy to overlook in a general forensics write-up because they're more prescriptive and sometimes extend beyond pure cybersecurity telemetry.

First are the PCI-defined audit events inside the CDE. These go beyond the basic admin activity, but explicit coverage for access to the logs themselves, failed logins, changes to authentication credentials and privileges, creation/deletion of system-level objects, and the start/stop/initialization of logging (which is critical for proving an attacker didn't simply turn evidence collection off).

PCI also leans hard on tamper-evidence, so it's worth calling out file integrity/change-detection alerts focused specifically on log stores and audit pipelines, not only on operating system binaries and configurations.

In addition, PCI introduces two enterprise-relevant signal categories that broaden forensics into customer-impact and assurance. For e-commerce, PCI emphasizes detection of payment page tampering and skimming. This refers to monitoring changes to scripts and security-impacting HTTP headers as received by the consumer browser, which often requires telemetry like synthetic checks, script/header integrity alerts, and (where implemented) CSP-style reporting.

PCI also expects organizations to ensure prohibited sensitive authentication data isn't leaking into debug/trace/error logs or even crash dumps, making logging hygiene a forensic control, meaning you want evidence, but not evidence that creates a compliance breach.

Finally, PCI uniquely pulls in physical access records tied to the CDE (badge/door access logs, visitor logs, and camera retention requirements), which can be decisive when you need to rule in/out insider actions or tampering with point-of-sale/CDE systems.

Conclusion

Forensic readiness isn't just a technical luxury. It is a foundational control that separates guesswork from decisive incident response. To reconstruct attacker actions with clarity, enterprises must prioritize collecting high-fidelity signals across endpoint, identity, network, email, SaaS, and cloud platforms, ensuring these signals are time-synchronized, centrally stored, and retained long enough to cover realistic dwell times. Centralization, normalization, and immutability turn ephemeral alerts into durable evidence chains that can withstand both regulatory scrutiny and adversarial tampering. Retention strategy should be treated as a measurable risk decision, not a cost center, and aligned with both investigative needs and compliance frameworks like NIST, PCI, SEC, and NYDFS. Ultimately, it's not just about having logs but about having the right logs, structured for correlation, preserved for resilience, and immediately accessible when seconds count. This shift in mindset (from reactive detection to proactive forensic capability) will define the maturity of modern security programs in the face of evasive adversaries.

Forensic Log Collection Checklist (by Domain)

The following checklist is intended to help organizations begin their log collection journey by identifying the core signals required for effective incident forensics and investigation readiness.

Domain	Purpose	Required Logs	Recommended Logs
Endpoint (Workstations & Servers)	Prove execution, change, and persistence	Process creation with parent and child relationships Command-line arguments and script content File creation, modification, and deletion events Registry or configuration changes tied to persistence Local authentication and logon events Outbound network connections from the host Durable device identifiers	Module and DLL load events Process injection or memory abuse indicators Sensor disablement or tampering attempts Off-host log forwarding
Identity & Directory (IdP, IAM, AD, Cloud Directory)	Prove authentication, access, and privilege changes	Authentication success and failure events MFA challenges and outcomes Conditional access or risk decisions Source IP address Device context, user agent, and location Session or token usage	Account creation, deletion, and disablement Password resets and MFA method changes Group and role membership updates Privileged role activations and approvals OAuth consent grants and service principals
Network	Confirm communications and	DNS query and response logs	NetFlow or IPFIX records IDS, IPS, or NDR alerts

Domain	Purpose	Required Logs	Recommended Logs
	potential data movement	Web proxy or secure web gateway logs Firewall allow and deny events Egress traffic destinations	East-west traffic visibility
Email	Capture initial access and mailbox abuse	Message trace and delivery logs Sender and recipient metadata Email authentication results Attachment analysis results URL rewrite and click activity	Mailbox rule creation and modification Auto-forwarding and delegation changes OAuth application access Mailbox search and export activity
SaaS & Collaboration Platforms	Identify data access and sharing	File access and download events Mass download indicators External sharing activity Link creation events Permission and ownership changes	Tenant configuration changes API access and token usage Administrative role actions
Cloud Control Plane (IaaS/ PaaS)	Track configuration and privilege changes	IAM role and policy changes Security group and network updates Logging configuration changes Key and secret lifecycle events Resource creation and deletion	Alerts when logging is disabled Privilege escalation indicators

Domain	Purpose	Required Logs	Recommended Logs
Cloud Data Plane (IaaS/ PaaS)	Prove data access and exposure	Object storage reads and downloads Database query and export activity Secret and key access events Sensitive API calls	Behavioral baselines for abnormal access
Data Access & Movement (Cross-Domain)	Establish scope and exposure	File reads from sensitive repositories Database and analytics exports External data sharing events	DLP alerts CASB or SSPM audit logs
Log Management & Integrity	Preserve evidence and trust	Centralized log collection Consistent time synchronization Off-host log storage	Immutable or write-once storage Access logging for log repositories Alerts on logging changes

Minimum Log Retention Guidance

Retention Tier	Minimum Window
Hot and searchable	90–180 days
Archived and retrievable	12–24 months
Regulated audit logs (e.g., NYDFS)	Five years, where required