RELIAQUEST

Top 5 Log Sources You Should Be Ingesting but Probably Aren't

Joe Partlow, CTO

Best Practices for Ingesting & Monitoring a High Volume of Custom Log Sources

Logging Tips That Save Money and Enhance Security

Logging and monitoring all relevant events from across the IT environment has ups and downs — some common log sources are fairly easily ingested and parsed, while others are difficult to manage at scale, creating visibility challenges.

Logs from servers, firewalls, Active Directory, intrusion detection systems (IDS) and endpoint tools are usually the easiest to get and first ingested. Many other sources are invaluable for incident response (IR), but rarely ingested because of the level of effort involved.

To maximize benefits of logging efforts, enterprise organizations must evaluate and adapt existing processes to fit current needs and threats, and consider logging additional — often overlooked — sources that could provide a huge benefit for IR and HUNT exercises. This document details these key activities and identifies their threats to assist IT leaders in maximizing logging and monitoring return on investment (ROI), but is by no means an exhaustive list of all recommended logging activities.

Assessing and Improving Current Log Activities

When planning which log sources to bring into the security information and event management (SIEM) system or other security tool for analysis, always consider the level of effort involved. Many common log sources, such as syslog or Windows events, are relatively straightforward and have many supported ingestion options bundled into the security tool. However, custom or uncommon log sources often present challenges for accessing actionable data.

> **The first challenge is the problem of log source volume.**

The first challenge is the problem of log source volume. Since many of the tools have architectural considerations and pricing associated with the volume of data ingested, start by figuring out what data you actually need to bring in and how to get it.

In the past, IT professionals and tools handled much less data, so they could point everything to a logging and monitoring platform that alerts on actionable data. Today, security and analytics teams deal with many more data sources and massive event volumes, which makes the old method both ineffective and impractical. Best practice now suggests reviewing the log source events, performing extensive filtering to ingest only the actionable events, and archiving the non-actionable events to long-term storage solutions.

When ingesting events from custom log sources, make sure to parse out potentially non-standard fields in the record layout. In-house or custom applications developed recently often create logs in a standard format such as JavaScript Object Notation (JSON) or Syslog. However, many older legacy applications have highly specialized, multi-line or obscure formats that need difficult regex expressions to parse. If working with an application log whose program was not developed in house, it may be difficult to have third-party developers change the format or add additional fields to the log message because it may not be included in the statement of work or could incur additional development costs.

Top 5 Log Sources You Should be Ingesting but Probably Aren't

The prioritization of which often-overlooked log sources to ingest depends on many factors, so we've listed the below items alphabetically, rather than ranking in order of importance. As a leading IT security provider, ReliaQuest works with enterprise organizations to develop and prioritize these and other processes to ensure consistent security, recognizing there is no effective one-size-fits-all approach.

1. CLOUD PLATFORM LOGS

More and more enterprises are turning to cloud providers such as Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP), or cloud services provided by Salesforce, Dropbox or Office 365 to house their data and applications. Unfortunately, many of these cloud platforms do not have consistent logging formats and require different parsers or methods of logging the events from the various applications contained on the platform.

Building these parsers to scale to the number of events is a challenge for most security and analytics teams — but much like web applications, effectively pre-filtering the data before ingesting will prevent overwhelming your SIEM or logging tool by only handling the actionable events.

> **CASB solutions are essential for incident response and forensic investigations since monitoring and alerting on unauthorized access to cloud services is often critical to detecting potential insider threats.**

While typically not an all-encompassing enterprise platform, Cloud Application Security Broker (CASB) solutions provide granular auditing capabilities at the application or service level and need to have the same logging and monitoring considerations as the full cloud platforms. CASB solutions are essential for incident response and forensic investigations since monitoring and alerting on unauthorized access to cloud services is often critical to detecting potential insider threats.

Important items to watch to ensure visibility for cloud services and platforms are:

Administrative log ins to console portals
Remote user access and actions to sensitive data folders or storage buckets
Underlying cloud infrastructure changes such as instance creation/deletion or routing/access changes
Platform user or role permission changes
Unauthorized cloud application usage, i.e. Dropbox or file-sharing sites
Office 365 users emailing attachments to external email providers
Cloud-based DLP application alerts for sensitive content or files containing personally identifiable information (PII)

2. DATABASE LOGS

Database auditing and logging poses a challenge for most enterprises, as database administrators often avoid enabling any feature that could potentially affect server performance—rightfully so—and enabled auditing typically results in adverse performance.

Auditing individual databases and tables is also very difficult given the large number of database servers present in a normal enterprise environment. Security teams may also struggle to gain access and visibility into operations occurring in databases created by third parties that may have placed restrictions on viewing the data or table structures.

For sufficient visibility into these databases without enabling auditing functions on the database tables, consider one of the following options:

1. If Database Activity Monitoring (DAM) is present, ingest and correlate built-in rules and alerts into the SIEM, since it performs many of the same restrictive functions as a firewall or web application firewall (WAF).

   OR

2. Create stored procedures that watch for specific actions. For example, build a script that watches for unencrypted PII stored in a field every 30 minutes, and write an event log with the record ID, date and time of the violating record entry to trigger an alert. These scripts are usually very low impact and can look for very specific scenarios, unlike the all-or-nothing approach for table auditing.

Important scenarios to watch for with any collection method include:

- System Administrator (SA) other administrative log ins
- Connections from unauthorized hosts and users (servers typically are only accessed by specific service accounts and not directly by users)
- Authorization failures
- Built-in scripts or functions that perform or enable direct Operating System actions
- Unauthorized server startup/shutdown commands
- DDL, DML, or DCL schema operations
- Queries returning or modifying abnormally large numbers of records
- Unauthorized commands such as **INSERT, UPDATE or DELETE** into lookup or administrative tables

3. DOMAIN NAME SYSTEM (DNS) LOGS

Logs from DNS servers provide a wealth of information about what sites users visit and show if any malicious applications reach out to command and control sites.

DNS has also been successfully used as a tunneling protocol for exfiltrating data since firewalls typically allow it out. These logs are challenging because of the volume of data, the traditionally difficult way to export them, and their multi-line format.

Consider using Microsoft's new Analytical Event Logging method, which uses a more standard logging format, instead of the old method of turning on debugging and importing the flat file. The new Analytical logs have major performance gains over the debug method, and the events are stored in the common Windows Event Log format.

DNS has also been successfully used as a tunneling protocol for exfiltrating data since firewalls typically allow it out.

Useful events to watch from a security standpoint include:

- DNS queries made using transmission control protocol (TCP) rather than the common user datagram protocol (UDP)
- Any non-standard record types such as TXT, or responses that contain large amounts of nonsense characters
- DNS queries either to non-standard ports other than port 53, or requests at strange times when a typical user would not be browsing
- Requests from Internal RFC 1918 IP addresses not on the company domain
- Zone transfers to unauthorized DNS servers
- Requests to uncommon country extensions such as .ru (Russia) or .cn (China), especially if you do not conduct business in those countries
- Substantial amounts of lookup failures
- Long or random hostname queries sent to the same or small set of domain addresses
- Requests for unique hostnames that aren't commonly requested (i.e. other than the Top Alexa list) or that are very young from a registration standpoint

4. PHYSICAL SECURITY LOGS

Obtaining event logs from physical security applications such as camera systems, biometric/card access readers, or alarm systems is extremely valuable for cases involving insider threat. Combining these events with other evidence correlated from servers, workstations, firewalls, VPN or remote access devices is essential to demonstrate a person's presence at the time of violation or whether their credentials were stolen.

One of the biggest challenges in obtaining these types of logs is the separation between the physical security team and IT security team. These two teams rarely work together because enterprises often outsource physical security personnel monitoring these systems, and most access control systems are on a closed legacy system.

However, it's not impossible to ingest logs from these systems and efforts should focus on the following types of events:

> **One of the biggest challenges in obtaining these types of logs is the separation between the physical security team and IT security team.**

- Remote log in with corresponding badge entries
- Unauthorized physical access to remote, unmanned facilities
- Audit of authorized employees accessing potentially unauthorized areas of the company
- Excessive biometric or card failures
- Visitor/contractor access to unauthorized areas
- After-hours alarm triggers or excessive open-door time alerts

5. WEB SERVER LOGS

According to vulnerability and exploit-tracking efforts such as the Verizon Data Breach Investigations Report (DBIR), the highest percentage of breaches trace back to holes in public-facing web applications — typically the log sources teams have the least visibility into.

Web applications often have access to highly sensitive customer account information and should top everyone's to-monitor list, but high visitor volume of an average e-commerce site can make dealing with these logs unwieldy.

Parsing web server logs is challenging as well because they are often in a multi-line or custom format, and possibly logged in a non-standard way to a text file or database as opposed to the native web server log, such as Internet Information Services (IIS) or Apache. If using standard web server logs, be sure to enable all the relevant fields since the default World Wide Web Consortium (W3C) layout in IIS doesn't capture some critical elements, such as page size and cookie values.

By using extensive filtering before ingesting into the SIEM or logging tool, IT teams can exclude requests for static images or style content and reduce non-actionable events by up to 90 percent. Logging events from a WAF could also possibly achieve the same result without ingesting the raw server/application logs since it already watches for potentially malicious actions.

By using extensive filtering before ingesting into the SIEM or logging tool, IT teams can exclude requests for static images or style content and reduce non-actionable events by up to 90 percent.

Other scenarios that benefit from alerts or anomaly reporting include:

- Excessive 500 or 404 request status errors
- Successful 200 request status on unauthorized file extensions (.zip, .inc, .txt, etc.) or extremely low-usage files
- Very long requests that could indicate buffer overflow attempts
- Known injection (SQL, CSRF, etc.) patterns contained in the query string parameter
- Unusually high page load times or page size
- Access to administrative portals outside of regular business hours
- Abnormal or blank User Agent values
- Page requests by IP addresses instead of the domain name
- Requests for known web shells
- Base64 or other encoded data in the requested URL that could indicate an exfiltration attempt
- Invalid or uncommon referrers or page request methods (POST vs. Get, TRACE, etc.)

What's Next?

The log sources above represent the next step in improving visibility into the enterprise security environment. It's often helpful to create a roadmap with all possible log sources and work with the other affected business units to set priorities, taking into account the level of effort for ingesting them and the potential risk mitigated by gaining visibility into these events.

Having the security team work with the data or application owners ahead of time ensures they can review the actionable event types together and see what elements the source owners might need visibility into as well.

Starting these processes early can also improve results, as gaining access to the data or getting the data sent to your aggregation or collection layer is often the most time-consuming step toward improving logging and monitoring.

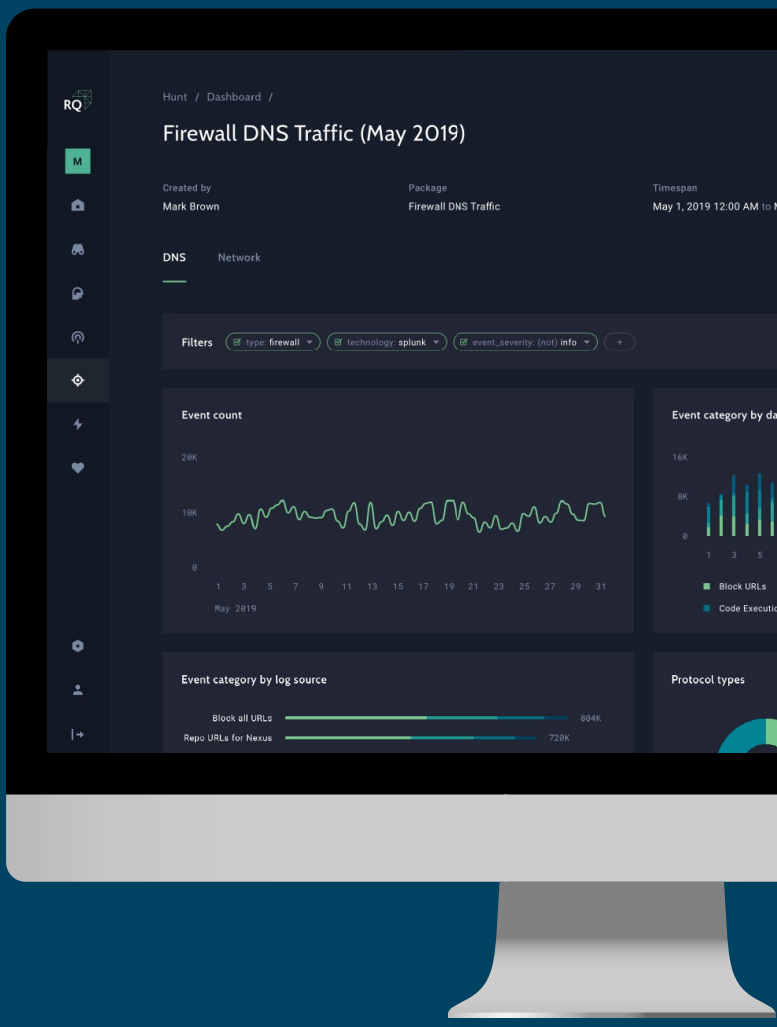

How ReliaQuest GreyMatter Improves Visibility

With the current pace of business change, shift to cloud, and growing tool sprawl, it is inevitable that critical data will not be centrally collected by the SIEM. ReliaQuest fortifies the world's most trusted brands against cyber threats with GreyMatter, its platform for proactive security model management. It does this by unifying and integrating existing SIEM, EDR, multi-cloud, and third-party apps, to deliver a centralized, transparent view across the environment. With GreyMatter, customers are able to increase enterprise visibility and automate their threat detection and response. The platform's analytics provide actionable reporting and metrics that measure ongoing improvement of the security model.

LEARN MORE ABOUT RELIAQUEST GREYMATTER

RELIAQUEST

Make Security Possible™

(800) 925-2159 **www.reliaquest.com** **info@reliaquest.com**