Realtime
publishers

Log Management:
Best Practices for Security and Compliance
The Essentials Series

The Importance of Log Management to your Security and Compliance Practices

Eric Schmidt

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Table of Contents

Copyright Statement

The Importance of Log Management to your Security and Compliance Practices

Introduction

Every information technology system, application, or appliance that an Enterprise deploys shares a common thread no matter what type of operating system or application they are. What they share is that to one degree or another, details regarding the operations they perform are captured in log files. The log files that systems and applications create can contain a vast wealth of information about the health and daily activity the infrastructure; however, they are generally local to the system or applications that generate them. This highly distributed creation and storage of log files creates significant challenges to leverage them in a way that will benefit the enterprise from a security and compliance perspective. This Essentials series, which consists of three articles, will focus on the benefits of centralizing logs and best practices for leveraging them for troubleshooting, incident response, and maintaining compliance with existing and new regulations. In order to leverage log files we'll discuss how to develop an effective strategy for centralizing the collection of logs and the type of systems for which logs should be collected from.

Understanding Log Files

The first step to leveraging log files is to develop a general understanding of the types of data that is recorded by the various operating systems and applications that reside within the typical Enterprise Information Technology infrastructure. Operating systems of all types will log system and application activity as well as authentication and configuration changes. Each event that is logged, no matter what the actual log storage mechanism is, will contain data regarding the date and time the event was created and, when possible, the account that performed the action. Log files are created in several formats, from flat text files, those that adhere to standards like W3C for web servers, to those that are completely proprietary. Log file locations also vary greatly depending on the operating system or application. Unix/Linux operating systems have a standard /var/logs/ directory that most applications adhere to. Microsoft Windows has moved toward an xml based log system, however, third party vendors may not leverage the event log system, choosing instead to store text logs in different locations. In fact, there are some Microsoft services that don't leverage the built in log structure. Internet Information Server and Windows Firewall with Advanced Security utilize text files. In most cases, the degree of detail that logs can capture is also configurable and range in scope from nothing to everything. By default most operating systems and applications will log detail somewhere in the middle and vendors will only recommend logging everything for very short periods of time for troubleshooting purposes.

Log Forwarding

Modern operating systems, appliances, and network equipment all contain built in functionality to forward event logs another location, either for the purposes of centralization or simply to archive logs to other systems. Microsoft operating systems, for example now contain a built in log forwarding mechanism which is subscription based. In their implementation, one system subscribes to the events of another. For non-windows systems, network equipment, printers, and appliances, SYSLOG is the established standard for centralizing events.

Log File Uses

Log files can be used for multiple purposes with the most common use being by support staff to troubleshoot system, application, or configuration issues. Enterprises that limit the use of log files to troubleshooting are failing to take advantage of the additional benefits they provide, but may not be able to because the investment was not made in log centralization. Without an effective centralization strategy, properly leveraging the data in log files is a very labor intensive process. For example, log files can be a critical component of collecting activity when investigating a security incident. The files will likely contain the necessary information needed to answer the required questions of who, what, where and when. Taken one step further, centralization of logs enables events from different systems to be analyzed in one place. In the case of a security incident, multiple systems might be compromised and there may be common indicators in the logs that one could look for. The valuable insight can be used to detect what systems were impacted or where the actor/malware went once inside the infrastructure. Without centralized logs, the investigators would have to check logs on each system, which could be a very labor-intensive process. The next article will discuss in greater detail, the benefits of centralized logging and how tools can be leveraged to make detection of security incidents easier.

Compliance

Additionally, log files can be leveraged for maintaining regulatory compliance. Many organizations must comply with regulations like Sarbanes-Oxley which requires auditing of activities like the provisioning of user accounts or access to financial systems. There are a number of industries that are beginning to fall under additional regulatory requirements both within the United States and Internationally. Some of examples of these in the United States include Sarbanes-Oxley, HIPAA, Dodd-Frank, etc. Internationally, versions of Sarbanes-Oxley have been implemented in Japan and the European Union. This ever-growing list of regulations and their increased scope further necessitates that enterprises leverage the benefits of centralized log collection as a tool to help maintain their compliance. Logs play a key role in maintaining compliance because they provide supporting evidence during an audit. The auditor may take a sample of a particular activity (user provisioning, for example) and will require evidence of how those accounts were provisioned to verify that established processes and procedures were followed. The lack of centralized logging or tools to extract the right data for an auditor could result in wasted time collecting the right information (provided the individual systems still have it) or worst case, the required data has been overwritten, tampered with or lost.

Health and Troubleshooting

Log files also contain a vast wealth of information about the activity and health the IT infrastructure. Data points like system uptime, resource utilization, and user activity are things that are collected by log files. There are many examples where logs contained the necessary information to avoid a service disruption, but since the logs were not monitored the issues weren't caught before the service went down. A simple example of this would be a server with storage problems, either running out of available disk space or a hard drive that's failing. Systems will log these events, and in many cases, well before an actual failure. Given that a common practice by system administrators is to only review logs when there is a "real" problem the events may go unnoticed until a complete failure has occurred. If logs had been centralized and tools deployed to alert on these types of issues, it's likely that corrective action could have been taken in a controlled manner to avoid unexpected service disruptions. Detecting overall health problems also includes things that are less catastrophic than hardware failure, but potentially just as important. It may be a situation where a service is degraded and users notice it, but don't notify IT that they are seeing a problem. In this case the issue may go unnoticed by IT for an extended period of time. At some point, IT is made aware, and a review of the logs indicates that the problem has been "going on for weeks". If only those events had been forwarded to a central location where they could be detected and corrective action taken.

Centralization

Centralized log collection provides a tremendous amount of benefits, but there are also some challenges that need to be overcome. The first and most important is ensuring that the amount of available storage is sufficient to hold what will be collected. Failure to plan a sufficient amount of storage will nullify all the benefits that centralization affords so it is critical that a significant amount of planning and analysis be performed before embarking on centralizing logs. There are several factors to consider when planning centralized log collection is to first determine what systems to collect from. Services should be prioritized based on the sensitivity of the data that they hold. Those that contain authentication, financial, personnel, or company proprietary data should be considered as the most important. Centralizing log collection from these systems benefit both security and compliance monitoring. From there a review of the other services can be prioritized based on their relative importance to the most critical systems. From a security perspective one area that is often overlooked are users' workstations. If resources permit, it should be considered to collect workstation logs because from a security perspective these are generally the first systems to be compromised. The centralized collection of workstation logs, coupled with the ability to analyze what's collected, enables more rapid detection of a compromised system. This can help prevent the compromise of a single system from becoming a more widespread security incident. Workstation logging can also be used by support organizations to more proactively detect application and configuration issues within the environment. This helps to mitigate widespread user disruption do to issues caused by common activities like patch or application deployment.

Integrity

Another aspect of centralized logging pertains to the integrity of what is collected. One should assume that any log file can be tampered with or modified. This is one of the most common activities performed by malware and hackers to hide or mask their presence on a system. There are a couple ways that this can be avoided. The first is to make sure log events are streamed in real-time as they are created to the central repository. Solutions that provide this type of functionality eliminate the risk of an attacker compromising a system and then deleting the local logs to hide their activity. Real-time steaming of logs can negatively impact network performance in very large environments due to the sheer volume of traffic that it could generate. This can be mitigated by a properly architected log infrastructure, which will be discussed in more detail in third article in this series. Another way tampering can be detected is to leverage products that verify the integrity of collected events. This is generally done by using hashing algorithms to compare the source event to the one received by the collector. If the hashes of the source and destination logs match then one can be assured it wasn't tampered with.

Reporting and Alerting

Even though centralization provides a single repository for all log data there is still a major challenge with turning the collected data into useful, actionable information. Depending on the flexibility of the logging that a system or application provides to control how much is or is not logged, there may be a significant portion of the data that is of little or no use in support of compliance or security monitoring. In many cases this extra data creates noise that makes finding the important bits equivalent to searching for a needle in a haystack. The noise that logs contain also translates directly to wasted labor because it extends the amount of time required to identify and transform the collected data into actionable information. This is also true in the case of an audit where resources are wasted trying to identify the events or that the auditor requires. Noise also creates opportunities for critical information to be overlooked, or analyzed incorrectly. Therefore, when selecting a tool for log centralization, one must evaluate the tools ability to transform the collected data into useful information by eliminating the noise.

Conclusion

With the discussion of the importance of log forwarding and centralization concluded, the next article, "How to Leverage Your Logs to Secure Your Environment", dives deeper into the various ways that centralized logging can be fully utilized. Included will be real-world scenarios where log centralization coupled with appropriate tools to analyze, report, and alert on the collected events enables an Enterprise to better manage their information technology infrastructure.