

# ENTERPRISE-WIDE LOG ANALYSIS: FROM OPERATIONS TO REGULATIONS

All employees that utilise technology, whether with malicious intent or not, leave traces of their activity in various logs, generated by IT components, such as firewalls, routers, operating systems, databases and even business applications. Such log records accumulate, creating mountains of data. At the same time, more organisations are starting to become aware of the value of collecting and analysing such data; it helps them keep an eye on the goings-on within the IT environment – the who's, what's when's, and where's of everything that happens. This also makes sense given the growing emphasis on data security and evolving regulatory compliance mandates such as PCI DSS and Sarbanes-Oxley, as well as "best practices" frameworks such as ITIL and ISO27002.

Of course, simply generating and collecting the logs is only half the battle. Being able to intelligently search and report on log data in order to detect, manage, or even predict, security threats, operational issues and to stay on top of regulatory requirements is the other half. Logs have traditionally been handled by reviewing them on their individual points of origin (such as servers or security devices) and only after a major incident. Such an approach is simply not working in this age of massive data breaches and stringent compliance requirements. It is not only inefficient and complex, but it can also cost a large company millions and take weeks, thus destroying, or severely reducing the positive effects of, such "log review".

Today, the call to action is shifting from merely having log data to centralised data collection, real-time analysis and in-depth reporting and searching to address IT security, operational and regulatory compliance issues. Thus, the main log-related goals of a company should be both to turn logging on and collect the logs and then to find a way to search and review log data from disparate points of origin across the system boundaries, across the IT infrastructure.

## What are these logs?

First, because logs contain information about IT activity, all logs generated within an organisation could have

More organisations are starting to become aware of the value of collecting and analysing log data to help them keep an eye on the goings-on within the IT environment

Dr Anton Chuvakin reports

**All logs (security, server, network, and others) make up one piece of the puzzle of IT infrastructure activity, so it makes sense that all log data is crucial to enterprise security, to regulatory compliance, and to IT operations.**

relevance to computer security and regulatory compliance. Some logs are directly related to computer security: for example, intrusion detection alerts are aimed at notifying users that known malicious or suspicious activity is taking place. Other logs, such as server and network device logs, are certainly useful to information security, but in less direct ways. Server logs, such as those from Unix, Linux, or Windows servers, are automatically created and maintained by a server; they represent activity on a single machine. Server logs are especially useful in cases of insider incidents; given that an insider attack or abuse might not involve any network access as well as not trigger intrusion detection systems and happen purely on the same system, server logs shed the most light on the situation. Relevant logged activities on a server include login success/failure, account creation and deletion, account settings and password changes, and file access, altering, or deletions, and usually contain the information on the user who performed the actions.

Network logs, on the other hand, describe data being sent and received over the network, so it makes sense that these are best-suited to assist in detecting and monitoring abnormal network activity. Unlike server logs, which are limited to one machine and indicate activity only, network logs indicate a connection on the network, a source, and a destination. Relevant information found in network logs include the time a message was sent or received, the direction of the message, which network protocol was used to transmit the message, the total message length, and the first few bytes of the message. On the other hand, such logs typically do not provide the information on the actual user who attempted the connection.

All logs (security, server, network, and others) make up one piece of the puzzle of IT infrastructure activity, so it makes sense that all log data is crucial to enterprise security, to regulatory compliance, and to IT operations. However, given the number of sources of logs and the varying information the logs contain, there are many different pieces of the IT infrastructure puzzle. While one can try to look at logs in siloed fashion, the logs will then fail to form the "big picture" of enterprise activity.

## **Compelling reasons in favour of centralised log collection and analysis**

First, logs from disparate sources reviewed in the context of other logs offer situational awareness which is key not only to managing security incidents but also to a company's day-to-day IT operations. Routine log reviews and more in-depth analysis of stored logs from all sources simultaneously, are beneficial for identifying security incidents, policy violations, fraudulent activity, and operational problems and for providing information useful for resolving such problems.

Moreover, when responding to an incident, one needs to review all possible evidence, which means all the logs from all the affected and suspect systems. One query across all logs saves time, and incident response, whether to internal and external security threats, requires quick access to all logs to figure out the details of the breach, especially if it involves more than one part of the IT infrastructure. For example, searching for a user or an IP address across 4,000 servers might take days if one has to login to each server, find the logs and perform the searches. If logs are centralised and optimised for searching, it will literally take seconds or less.

And we can't ignore the high degree of operational efficiency that accompanies having all logs in one place. Further, troubleshooting issues across all systems become a one-click process, as does running high-level trend reports across all systems in a business unit. To seek out log data from individual sources, administrators have to consider too many things and spend too much time piecing together information to be efficient. But a single point of control means that with the click of a button, all relevant information can be at their fingertips and with a tweak of one knob, all logging configurations can be updated as security policies and compliance mandates evolve.

What is no less relevant is a current onslaught of compliance mandates. Further pushing IT professionals towards a cross-device approach to log analysis, many compliance mandates put forth a broad call to action to examine and review logs, not specifically database logs, not only server logs, not just network logs...logs generally. For example, in the US, the Federal Information and Security Management Act (FISMA, via NIST SP 800-

Another critical benefit of centralised log collection and analysis is that it enables privileged user monitoring by removing the logs from the control of the privileged IT users, such as system and network administrators.

53 and NIST SP 800-92 documents) describes the broad need for log management in federal agencies and how to establish log management controls including the generation, review, protection, and retention of audit records, and steps to take in the event of audit failure. The above NIST documents are actually useful even outside the FISMA context as they contain recommended practices, some related to logging.

Worldwide, the Payment Card Industry Data Security Standard (PCI DSS) mandates logging specific details and log review procedures to prevent credit card fraud, hacking, and other related security issues in companies that store, process, or transmit credit card data. Requirement 10 requires that logs for all system components be reviewed at least daily and those from in-scope systems be stored for at least one year as well as protected. Thus, the above regulations call for central control over all log retention, stringent access control and other log protection for evidentiary purposes and even logging access to all logs. All of these are only possible when logs are centralised.

Another critical benefit of centralised log collection and analysis is that it enables privileged user monitoring by removing the logs from the control of the privileged IT users, such as system and network administrators. Abuse of system and data access or even data theft by trusted users is unfortunately all too common and logging is one way to curtail that. Log data integrity of log data in a centralised repository can also be guaranteed by the access control rules based on the “need to know” basis, logging all access and the use of cryptographic technologies, such as hashing and encryption.

In essence, the need to paint a total picture of IT infrastructure activity and the broad requirement of key regulatory compliance mandates to review logs generally means that IT professionals need to find a way to execute cross-boundary log analysis. Of course, this approach requiring centralised retention of logs from disparate sources, has benefits.

The key point to remember is that any information that can be gleaned from log data is always present in enterprise logs. However, the limiting factor to how well that information can be put to good use is how quickly and efficiently the log data that contain it can be retrieved, searched, and reported on. If a company's IT

staff cannot access the appropriate logs in time and as a result must spend all of its time fire-fighting security breaches rather than proactively preventing such breaches before they become major problems, this does not maximise efficiency and it leaves the company open to even more security threats as the team struggles to catch up.

Log data storage in a centralised, analysis-enabling repository allows those seeking information to bypass the time- and resource-consuming process of combing through each individual source of log data and piecing the information together afterwards. Instead of spending their time either searching for information, administrators have the data they need at their fingertips and also can proactively review any log data that might indicate abnormal activity and address security, compliance and operational issues before they become major company blunders. In order to maintain efficiency and effectiveness, enterprises must be able to break down log silos and allow the intelligent analysis of log data from disparate sources. ■

**Dr Anton Chuvakin** ([www.chuvakin.org](http://www.chuvakin.org)) is chief logging evangelist at LogLogic, a log management and intelligence company, and is a recognised security expert and book author.

A frequent conference speaker, he also represents the company at various security meetings and standards organisations.

He is an author of a book *Security Warrior* and a contributor to *Know Your Enemy II*, *Information Security Management Handbook*, *Hacker's Challenge 3* and *PCI Compliance*. He has also published numerous papers on a broad range of security subjects and has several blogs.

