

Log Data Management: A Smarter Approach to Managing Risk

By Dominique Levin

Risk in today's corporate networks is all around. Not only must we worry about malicious attacks and viruses, corporate misuse and abuse can threaten a network's performance and jeopardize operations. Moreover, a lack of effective tools for monitoring and tracking network activity can lead to compliance and legal problems that cost time and money to remediate.

With business racing along at record pace, the demands put upon IT departments often prohibit proactive action against such threats. As a result, companies tend to be reactive to security breaches and network performance issues, fixing them after the fact—after downtime has already accumulated. That's why security issues, such as worms and viruses, internal or external fraud, and policy violations result in an average of 22 hours of downtime per year. Human error, system failures and natural disasters account for an additional 87 hours, the cost of which can add up to \$6.5 million per hour. Financial losses from IP theft alone amount to an average of \$1.3 million per company each year. Clearly, a more proactive stance for preventing such losses is needed.

To manage risk intelligently and cost-effectively, companies are looking to their network log data for answers. Log data provides a complete audit trail of user and system activity, while delivering critical decision support to mitigate security and performance incidents. Fortunately, log data is readily available in any data center, and companies need only find effective ways of collecting, aggregating and archiving it to put it to use.

What Your Log Data Does for You

Log data can provide a complete, independent record of network activity and user access to applications, servers and devices on a network. It can help corporations at risk to validate policies and perform change control audits in real time, and it can be used to trigger alerts to unusual or suspicious network behavior. Log data can even be used for root-cause analysis to aid in system recovery and damage cleanup. Perhaps most importantly, log data provides an audit trail of user logons and access for regulatory compliance and legal purposes.

Log data is extremely useful for companies looking to mitigate network risks and resolve issues quickly. By analyzing log data, IT gains insight about exactly who is using which applications and when, who has access to certain servers or applications and who has permission to make changes. Log data can be mined for information about accepted and denied connections, as well as what's happening during remote user sessions over VPNs. But the trick to using log data effectively is all in the collection and aggregation. Without access to complete data, searches can be ineffective and IT can waste time sifting through data sets that don't contain the necessary information. Any solution for log data management must therefore ensure the collection of complete data, in its raw form, so that searches are thorough and answers are guaranteed.

As it stands today, roughly 80 percent of Global 2000 companies still use homegrown scripts for mining log data, which are usually ineffective and slow and have difficulty adapting to disparate locations and formats. Since they are not standardized or automated, they require constant maintenance. Most companies lack the headcount to perform these tasks.

Another problem is that there is so much data to manage. 25 percent of all enterprise data is log data, and that percentage continues to grow. Global 2000 organizations can generate about ten thousand log data messages per second—the equivalent of two terabytes each month. Up to 94 percent of these log messages are seemingly normal, informational messages. Mining this amount of data is extremely time-consuming when using a script, and by the time IT locates the initial problem, damage could have spread throughout the network.

Besides homegrown scripts, some companies are using Security Information Event Management (SIEM) solutions. These may also prove ineffective and inefficient, as well as expensive to install and maintain. Most focus only on security events and ignore other issues, like misuse or impaired performance—a major oversight. Because they are incomplete, they cannot be used for compliance purposes or to improve availability. Lacking the back-end infrastructure, these solutions offer poor throughput and storage efficiency.

Compliance Groups Prove a Good Reference

To discover what's needed in a log management solution, companies need only look to the mandates and recommendations set forth by industry standards groups and new regulatory laws. Best practices frameworks developed in response to these mandates outline how to best take advantage of log data to protect a business and secure the network.

Although different industries answer to different compliance laws, every industry faces the same challenges in terms of collecting and managing data. In the finance industry, VISA CISP, FFIEC, GLBA, Sarbanes-Oxley (SOX) and Basel II require strict vigilance of IT departments in banks, brokerages and other financial institutions. Healthcare organizations must answer to HIPAA, as well as VISA CISP and SOX. Energy providers look to NERC and NISPOM for instruction on how to manage logs. In failing to meet the requirements outlined by these groups, companies can face fines or even criminal charges.

The basic tenets of log management are consistent across all best practices frameworks, such as CERT, COBIT and ISO, as well as statutes like HIPAA, SOX and GLBA.

They provide recommendations in four areas: authentication and authorization, configuration and change management, segregation of duties, and documentation. Here are some of the stipulations:

1. Companies should audit and monitor system and user activity logs for suspicious behavior, security breaches, unauthorized access and misuse
2. Companies should create a historical repository of events and retain complete and accurate log data for up to seven years.
3. No individual should have more rights than he or she needs to execute his or her assigned tasks.
4. No changes should be made without authorization.
5. A record of what changes are made should be maintained so that the state of a system or application at a previous time can be determined.
6. A single person should not have the right to configure IT systems as well as audit, initiate or approve incompatible activities in those systems. Similarly, development, testing and production environments should be segregated.
7. All entities must be held accountable.
8. Compliance should be documented and tested on an ongoing basis.
9. The audit trail should allow for testing of the internal IT control framework as well as substantiating regulatory compliance.

These recommendations serve as a blueprint for a log management solution that can actually help companies save money and reduce expenses by being more proactive about log data management.

Act, Don't React

Since most companies lack a complete log management solution, IT departments are in the habit of reacting only when a threat or performance issue occurs. For example, a server fails because of a virus that sneaks into the network from an unauthorized remote logon, and IT must scramble to pinpoint the point of entry and trace the path of the virus to all infected machines. This process is incredibly time consuming, and in the meantime, productivity and business continuity is at risk.

Why don't companies act sooner? Perhaps they feel the cost of implementing a complete solution is prohibitive. However, according to a log management survey recently performed by the SANS Institute, over half of Fortune 200 companies already spend more than \$250,000 a year managing their logs. Much of this cost can be reduced over time by deploying a log management infrastructure that automates the process of collecting, aggregating, analyzing and archiving log data from multiple network devices, and eliminates the need for much of the manual labor currently required to perform log management tasks. Some may also be concerned about deployment and maintenance. Finding a solution that installs easily and requires little support is critical to keeping costs down.

There are five main points to consider when choosing a log management solution to help your organization proactively manage the many risks that corporate networks experience today. They include:

- ▼ Data completeness
- ▼ Aggregation across dispersed network elements
- ▼ Automation
- ▼ Alerting and reporting features
- ▼ Secure, reliable archival

Data completeness

A thorough log management solution must address security, compliance and availability threats by providing IT with complete, accurate log data files that are readily searchable and accessible. Focusing only on warning

messages and alarms denies auditors the complete audit trail they need to validate security policies and test for compliance. It's important that companies keep a complete record of log data—unfiltered—in a secure archive while metadata is available for mining and analysis.

Aggregation across dispersed network elements

IT should have the ability to collect logs from disparate locations and transport the data reliably and securely over the network. This requires the ability to collect data from local and remote devices and parse and summarize a copy in real time. Data must be securely forwarded in its original, raw form to an archiving device.

Automation

Any log management solution should be completely automated, so that minimal maintenance and administration is required. Automation saves IT departments money and resources. Companies should look for an appliance-based solution that is ready to work out of the box; auto discovers network devices and interoperates with numerous platforms to avoid modification to the existing infrastructure. Automatic report generation can also save companies time and effort and keep key players in an organization abreast of any potential network issues.

Alerting and reporting capabilities

Alerting strategies that trigger automatic notification when unusual or suspicious activity occurs in the network are critical to enabling IT to be proactive about addressing network threats. Rules-based alerting allows administrators to define the state or policy deviation triggering the alert, and text-based alerts allow IT to set alerts on any text-string parameter found in log messages. Some solutions even offer automated statistical anomaly detection and intelligent threshold alerting, so that subtleties such as sudden changes in the ratio of successful to unsuccessful messages or a spike in denied packets on a particular port are detected early and can be addressed rapidly.

The best log analysis also provides real-time and historical log data mining, allowing IT to search log data at a granular level for information to aid in system recovery after security or performance incidents by investigating the best method of repair. In the case of network downtime, every second counts, and being able to find granular information through fast text-based searches can accelerate problem resolution by a factor of ten or more.


Secure, reliable archival

For legal and compliance reasons, any log management solution must provide central, secure log data archives, which are physically separate from the log data used for real-time analysis. That way, the raw log data is kept in tact, while a copy can be used for analysis. A high rate of compression helps to maximize transfer capacity, and the ability to connect to storage networks helps with storage capacity. Additional log data protection in the form of fail-over, fail-back and automatic backup capabilities must be employed to ensure data integrity.

A Jump Start on Risk Management

Say a travel firm loses 8 hours of bookings to downtime resulting from a virus attack that goes undetected long enough to crash the corporate firewall. The company would spend nearly \$200,000 to perform damage cleanup and could lose up to \$800,000 worth of online bookings. However, if log data anomaly detection mechanisms are in place before the incident, IT would receive an automatic alert, allowing administrators

to spring into action rapidly, pinpoint the problem, and reduce costs incurred by downtime. In addition to helping administrators reduce downtime, proactive log management protects corporate networks by reducing intellectual property theft, simplifying compliance activities, and reducing network misuse and abuse.

The reality is security breaches, downtime and performance issues will happen. However, mitigating the effects of these incidents is easier if the systems are in place before the problem arises. By arming the corporate network against risk with an effective log management solution, companies can reduce costs, improve productivity, and stop network incidents in their tracks. 

Dominique Levin is the Vice President of Product Management and Business Development at LogLogic, the log management visionary and leader.

