MAKE YOUR DATA SMARTER WITH LOG ENRICHMENT

Logs are a key pillar of the underlying data that feeds an observability solution and artificial intelligence for IT operations (AIOps), so, the data and insights derived from monitoring logs is as valuable as any other data type. Effective log analysis aids understanding of a system's performance and health to, help IT operations (ITOps) teams and site reliability engineers (SRE) identify issues as they emerge and quickly track down the cause of failures.

While log data is desirable and helps you understand what has occurred to cause a problem, it can often be cryptic, difficult to interpret and use, contain sensitive data, or lack the relevant context, all of which makes problem analysis difficult for a business analyst.

Consider a case where an SRE engineer or IT data analyst reports that a threat actor has been targeting their company's line of business for the last three months until two weeks ago. They need to investigate whether their company data was compromised. The analyst would gather logs from multiple applications or sources, or access them from the logs repository. These could include application, firewall, network, and system logs, and more, each containing a variety of useful information for investigation.

However, the analyst cannot triage correctly without contextual information. Searching the logs by a vulnerable host's name is not possible if the logs contain only IP addresses but no hostnames because the volatile, dynamic IP data may change every day or week, leading to incorrect and misleading summary and detail information—an issue that's exponential when investigating a three-month span. The only way to effectively search is by capturing the host name in real time. The further away we get, from the time the logs originated, the more inaccurate the information becomes. Likewise, there can be many other examples where problem analysis is difficult as the underlying data logs lack relevant information and context to debug any issue.

BMC Helix Log Analytics enriches log data by adding necessary context in real-time for enhanced observability and diagnosis. By enabling enrichment to log data (e.g., converting IP addresses to host names), it makes log data more useful for search, analysis, and other operational needs. You can enrich logs by connecting to multiple different enrichment sources like DNS, LDAP, GeoIP, and CSV and use them to define policies.

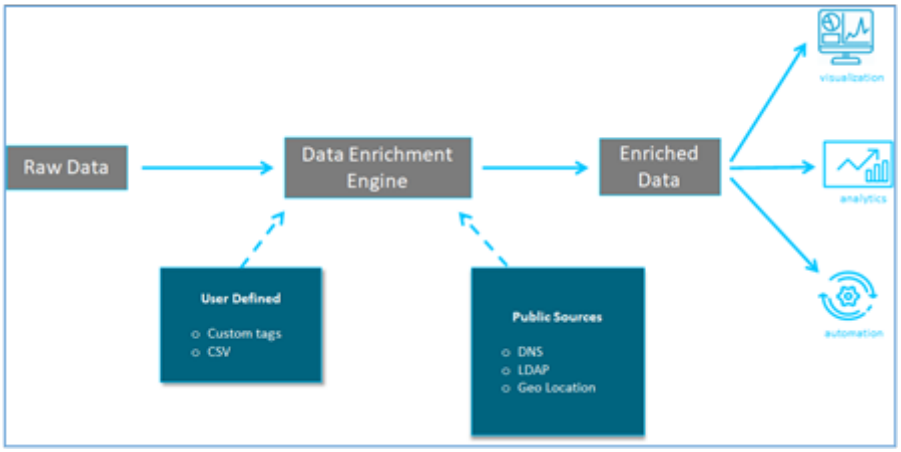

Figure 1. Log enrichment conceptual overview

The example below illustrates how log enrichment adds meaningful context (status text) to the audit logs generated, allowing SREs and DevOps engineers to audit user transactions and logins and troubleshoot whether a user login failure is due to invalid credentials or an internal server error.

Consider a case where an application is upgraded, and as a result, a new audit log is generated that lists the users' login status, but doesn't tell you whether the user login is successful or failed.

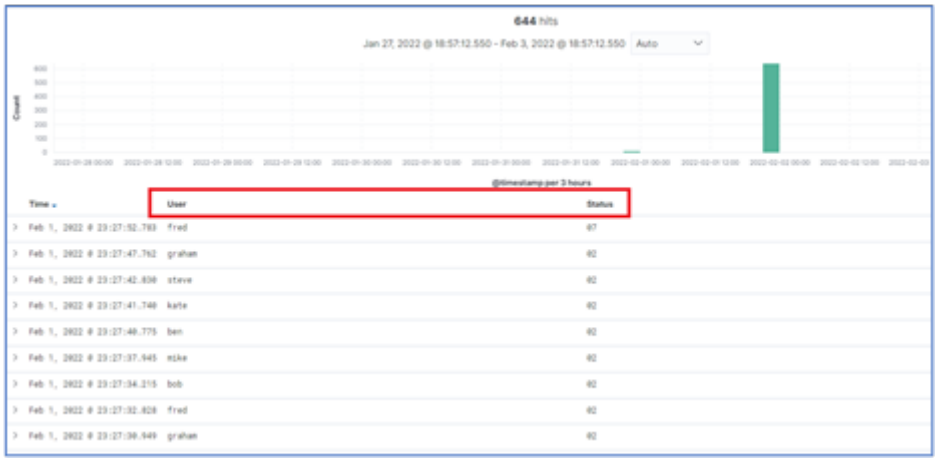

Figure 2. Log data before enrichment

So, this log data needs to be enriched, and for the given example, we would use a CSV file, which maps the status code with the status description and can therefore be used to enrich the logs.

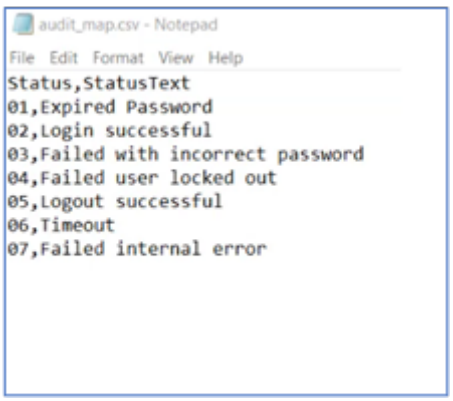

Figure 3. CSV file to enrich logs

Log enrichment is a two-step process. First, you must define and upload the type of enrichment source (CSV in this case) and then map the source field in the raw data and enrichment fields to be added to the log data.

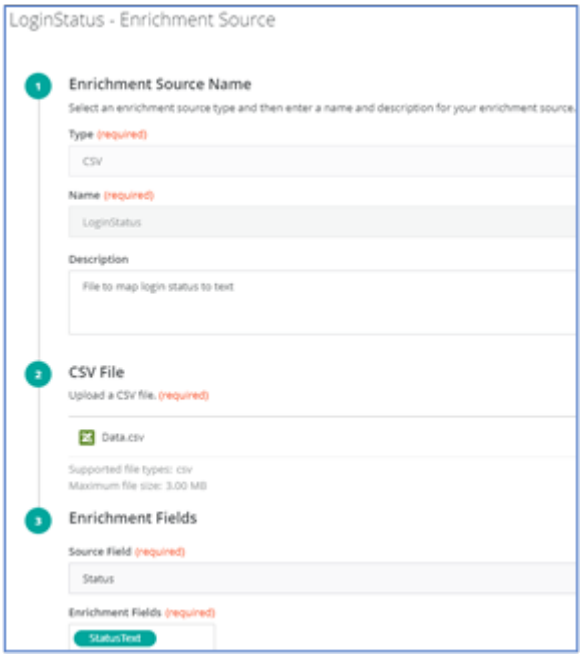

Figure 4. Configuration for enrichment source

Once the CSV enrichment source has been defined, you must define an enrichment policy by providing the condition used to trigger it using the fields present in the logs. One or more enrichment source is then associated to the policy, providing the mapping to the source field and target enrichment fields.

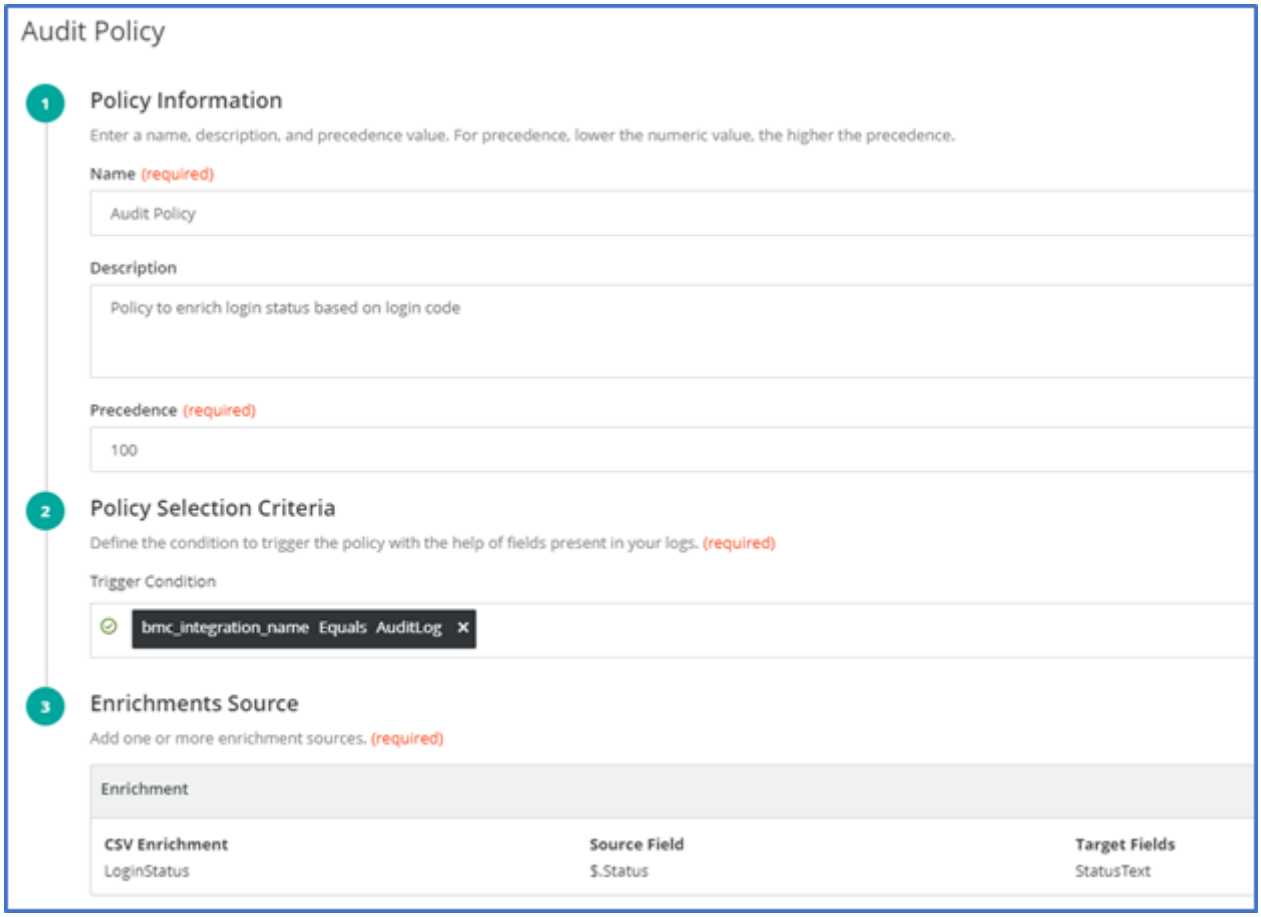

Figure 5. Configuration for enrichment policy

After the enrichment policy for the audit is enabled, the logs are enriched with the “Status Text” field, which provides the audit status and meaningful context to an analyst or SRE troubleshooting the application issues. Further, the analyst may choose to create an alert and be notified whenever a user's login status shows a failure.

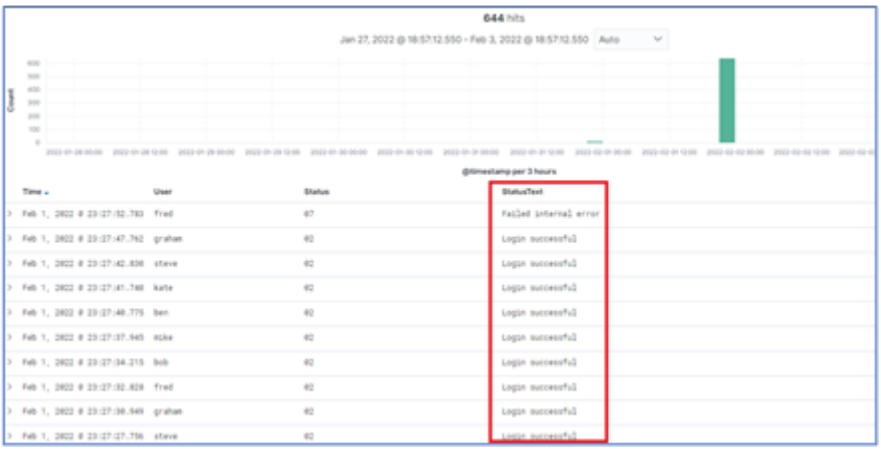

Figure 6. Logs after enrichment

Applying log enrichment on log data alongside the other advanced capabilities of BMC Helix Log Analytics can be invaluable for managing, maintaining, and troubleshooting IT systems; identifying performance or configuration issues; and meeting operational objectives and service level agreements (SLAs).

To learn more about BMC Helix and BMC Helix Log Analytics capabilities, watch our overview video here or visit www.bmc.com/helix or our documentation site.

Related Content

- Observability with Logs to Accelerate MTTR