ARCHIVE LOGS TO OPTIMIZE STORAGE & GAIN FULL VISIBILITY

To gain full visibility into modern cloud environments, businesses must collect an ever-growing avalanche of log data from a range of overly complex data sources. Retaining logs is key for real-time monitoring and troubleshooting, but it can quickly become expensive at high volumes, meaning that organizations must often choose which logs to index and which to archive.

With new business requirements to log everything all the time, it can be a challenge to store and analyze all this data effectively and cost-efficiently. The proliferating number of applications doesn't help, either. Another consideration is that the value of log data can transition from high to historical in a matter of weeks or days, which presents its own challenges when the storage cost of the data outweighs its potential value as a source of business insights.

With BMC Helix Log Analytics, you can archive and retain logs for a longer period of time at a cheaper cost. Before we dive into the details of the archival feature, let us look at some of the use cases that show us why archiving is important.

Why archiving log data is required

- **To meet regulatory standards**—For compliance reasons, archiving your logs ensures that you are fully protected. Data retention policies can vary from several months to several years, depending on the type of service you provide and the standard or regulation with which you need to comply. For instance, section 802 of the Sarbanes-Oxley Act (SOX) requires organizations to archive their data for at least seven years.

- **To identify patterns and trends**—Logs are essential for identifying and troubleshooting short-term problems but are less effective at identifying long-term trends. Older entries may get overwritten, deleted, or lost. Archives make it easier to identify patterns over a longer period than rolling log files.
- **To optimize log data storage**—Archiving log data by employing compression techniques and storing archived logs in a location that does not need to be optimized for quick access are effective ways to save storage space and reduce costs. Furthermore, since the data can be decompressed and loaded into active databases any time without any data loss, it can still easily be used for on-demand troubleshooting or any other operation.

Perform historical analysis and investigations

Having the ability to store and analyze enormous amounts of historical log data is vital for situations that do not necessarily need immediate query responses. These include things like running security investigations across large environments, conducting audits to adhere to strict compliance frameworks, and performing long-term analytics on high cardinality datasets.

For example, when you experience a security breach or receive a report of an insider threat, your security team will need to comb through weeks, if not months, of log events to identify malicious activity. An investigation of all the activity from a suspicious IP address may require scanning petabytes of data, assessing the timeline of activity from that IP, and generating reports for other teams (e.g., legal and executive).

Similarly, businesses operating in regulated industries—such as financial services, insurance, healthcare, and aviation—have stringent requirements around servicing audit requests from among vast amounts of historical log data.

E-commerce providers, digital content makers, sports and entertainment companies, and businesses using Internet of Things (IoT) devices frequently need to perform long-term analytics on high cardinality datasets, such as users, IP addresses, device IDs, or items purchased, among others.

The log archival solution provided by BMC Helix Log Analytics addresses these use cases by retaining logs for longer duration, and restoring them back for on-demand analysis so teams do not spend valuable time spinning up new solutions, finding data loss, or worrying about query capacity and associated costs.

Log archival solution from BMC

BMC Helix Log Analytics provides an easy-to-use solution to archive logs in multi-cloud, software-as-a-service (SaaS), and on-premises platforms to help you perform historical analysis and investigations.

The following diagram illustrates conceptual flow of log archival and restore. Logs are first saved in a hot storage for a predetermined retention period. After that, they are moved to cold storage for longer-duration retention. Logs stored in cold storage archives are not available for search. To analyze these logs, you need to restore and bring them back into hot storage. Post-analysis, the logs are then auto-archived. After the archive duration, logs are purged and unavailable for search.

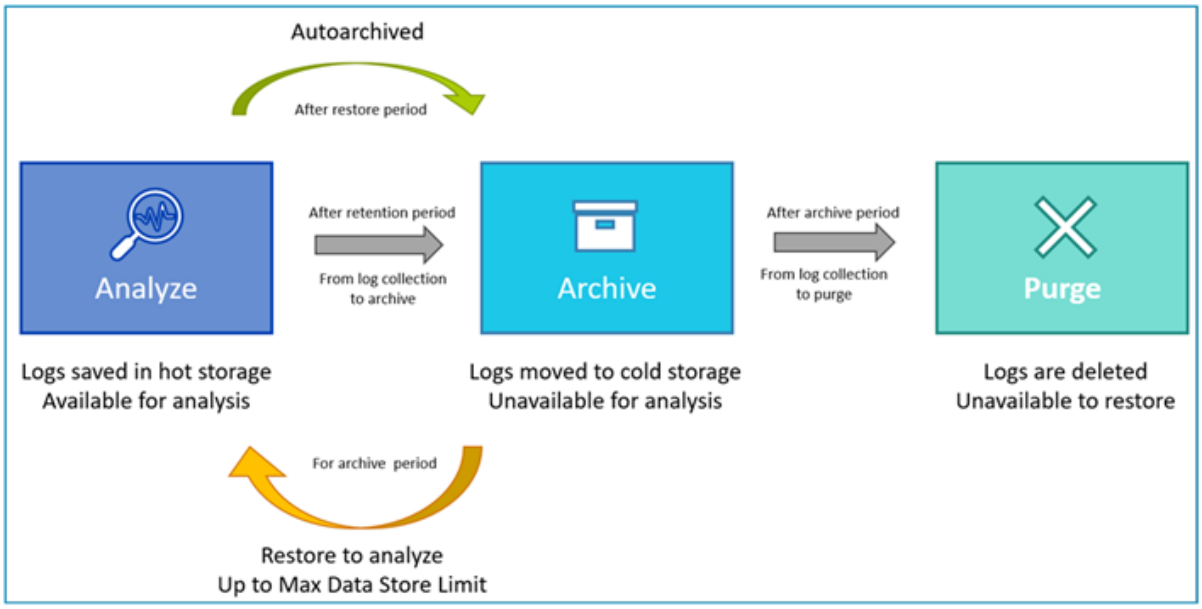

Figure 1. Conceptual overview of log archival and restore

Effortless configuration and data exploration

It is easy to configure this feature in your BMC Helix Log Analytics deployment by providing the application logs to be archived and the duration for which they need to be archived. Once enabled, all the logs' data residing in different indexes will move to cold storage following the specified hot-retention period. You can select the logs index to be used for troubleshooting and restore that data.

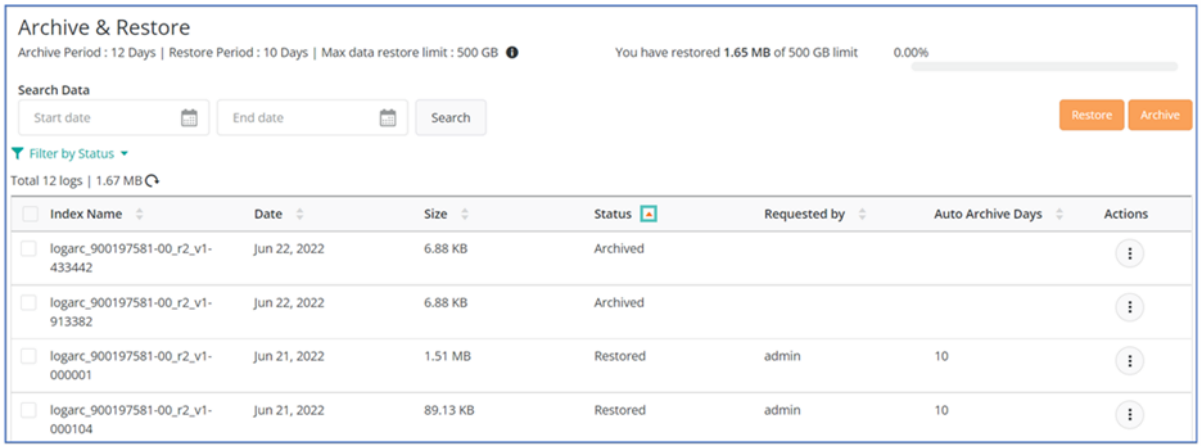

Figure 2. Configuration to archive and restore logs

The restored data is then available for analysis in the log explorer view, where you can query, search, and perform further action. The following diagram shows archived logs are restored and further analyzed in the log explorer.

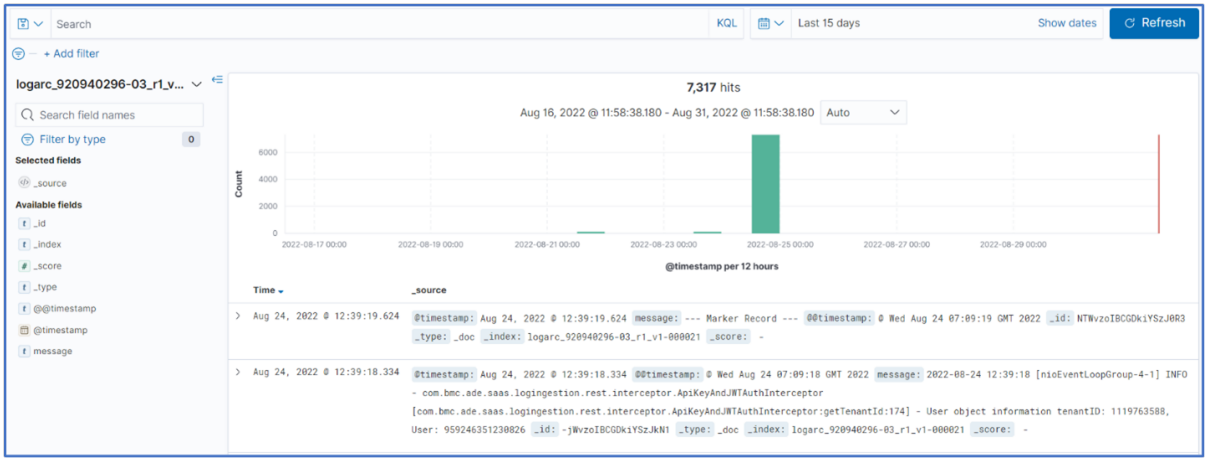

Figure 3. Analyzing restored logs in log explorer

Once your analysis is done, you can archive the data back, or have the restored data auto-archived once the restore period is over. This capability is accessible by administrator users who manage all the archival and restore operations.

The log-archival capability of BMC Helix Log Analytics is a cost-effective, cold-storage-based solution that helps organizations retain data for historical investigation and analysis and better meet compliance and regulatory standards, while continuing to use hot storage for real-time log streaming and alerting.

To find out more about log archival and restore, check out our BMC Helix Log Analytics product documentation and watch our overview video.

Related content

- Observability with Logs to Accelerate MTTR
- Make Your Data Smarter with Log Enrichment
- Kubernetes Observability with Logs