



Log360 Cloud's
log storage

M E C H A N I S M



Table of **contents**

Overview

ManageEngine Log360 Cloud is a cloud-based SIEM solution designed for analyzing, correlating, auditing, and storing log data from both on-premises environments and cloud sources such as AWS, GCP, Azure, and Microsoft 365.

Storage Tier

Log360 Cloud offers [customizable tiers](#) for search and archival storage. Retention [policies](#) can be configured to meet the specific requirements of each organization, and storage retention can also be customized based on log categories. For example, compliance-related logs can be retained for a longer period than application performance logs, depending on organizational needs.

Search Storage

Search Retention refers to the number of days logs can be searched in Log360 Cloud. This setting can be customized based on the purchased plan by navigating to the **License** page, selecting the **Edit** icon under **Search Retention**, updating the desired values, and saving the changes.

Customers can also use the **Reload Archive Logs** feature to search through archived log data. By default, 50 GB of search storage is included, and additional search storage is priced from **\$10 per GB**.

Archival retention

Archival retention specifies the number of days logs are stored in Log360 Cloud. Similar to search retention, this can be customized based on the purchased plan. To configure this setting, go to the **License** page, hover over **Archival Retention**, click the **Edit** icon, update the values, and save the changes. By default, 150 GB of archival storage is provided, and additional archival storage is priced from **\$0.25 per GB**.

Data Categories for Storage

In addition to log data, Log360 Cloud also stores the following types of data:

- Stats data
- Audit data
- Alerts and correlation data

Stats data: The statistics data logs encompass information related to the [usage statistics](#) of uploaded logs in Log360 Cloud. This includes details such as the date of uploads, data size, and the number of logs originating from configured on-premises environments or cloud sources. By default, this data has a retention period of one year.

Audit data: Audit data logs contain crucial [audit information](#) regarding operations performed within the Log360 Cloud account, specifically additions, modifications, and deletions. This data includes details such as the user responsible for the action, the time stamp of the action, the module in which the action occurred, status messages, and additional information like changed values. The default retention period for audit data is one year.

Alerts and correlation data: Log360 Cloud's [alerting](#), incident management, and correlation modules play a vital role in identifying and mitigating security threats at an early stage. The system matches logs with configured alert profiles and correlation rules, maintaining relevant data for better accessibility.

Fault tolerance mechanism

The system employs a fault over script designed to automatically initiate a failover process in the event of a primary server failure. In such instances, the cluster IP associated with the primary server is released. The secondary server, which consistently pings the virtual IP, detects a lack of response, and promptly takes over the virtual IP. Subsequently, it assumes the role of the primary server, responding to queries from the application.

If the primary server fails to respond to the health check performed by the cluster monitoring schedule, the system checks if the secondary server has successfully taken over as the primary server. If affirmative, the previous primary server is marked as **server down** and relegated to the end of the chain as the secondary server. At this point, the system either restarts the data server or removes it from the configuration.

Backup strategies

Here are two data backup strategies to consider implementing in your organization:

- Active/live backup:** The data cluster, running MySQL, is organized in a chain formation with multiple nodes. This architecture ensures the maintenance of multiple copies of user data in real-time, enhancing data redundancy and availability.
- Passive backup:** Weekly full backups and daily incremental backups are configured to capture the entirety of data clusters systematically. These backups are stored on dedicated backup servers. Additionally, the system provides an option for a monthly comprehensive backup, ensuring a robust backup strategy for data recovery and system resilience.

Disaster recovery and business continuity plan

A disaster is defined as a significant physical or functional failure on a large scale. The business continuity plan is a strategic framework that ensures the survival and sustained growth of the business, even in the face of disasters. This plan meticulously outlines countermeasures for all potential threats that could disrupt regular business processes.

While each component within a data center is meticulously designed with redundancy, the disaster recovery site extends this high availability standard on a geographical scale. It facilitates active replication from the main site and is brought online in the event of irrecoverable failure at the main site. In situations where there are issues with the WAN in the main domain controller or failures in the incoming path into the DC, instead of switching the applications, a process known as backhauling is implemented.

Log collection and upload flow in agent

The agent responsible for log collection follows a structured process. It collects logs from associated devices based on the configured monitoring interval for Windows devices or in real time for Syslog devices. These logs are then written to raw log files located in the agent installation folder. Subsequently, these raw log files are compressed into ZIP files and uploaded to the cloud using the HTTPS protocol every five minutes. In instances where network issues arise or the cloud storage reaches full capacity, these ZIP files are stored on the agent machine and retried for upload during subsequent upload cycles.

Log processing workflow in Log360 Cloud

Upon reception of logs by Log360 Cloud, they are transferred to Zoho logs storage for accessibility in searches. Subsequently, the data undergoes processing by various customer-configured processors such as [alerts](#), [correlation](#), [log forwarding](#), and [threat analytics](#). The data remains stored until the aforementioned retention period elapses.



ManageEngine Log360 Cloud, a unified cloud SIEM solution with integrated CASB capabilities, helps enterprises secure their network from cyberattacks. With its security analytics, threat intelligence and incident management capabilities, Log360 Cloud helps security analysts spot, prioritize and resolve threats in both on-premises and cloud environments. The solution is highly scalable and helps drive down infrastructure and storage costs. For more information about Log360 Cloud, visit <https://www.manageengine.com/cloud-siem/>.

[▶ Demo](#)[◀ Sign up](#)