

Scalyr 200 Terabytes per Day, Benchmark Report

The purpose of this report is to summarize the series of benchmarking tests that Scalyr conducted to validate service quality and performance at ingestion rates of 200 TB/day. It is intended to help the reader understand how Scalyr evaluates its SaaS Log Analytics service and how the system performs at large scale.

Success Criteria

These tests were performed in conjunction with Scalyr's 'Sonic Boom' announcement of our new architecture. At the outset of the benchmarking project, success criteria were defined as follows:

- The load generator can emit logs at a sustained rate of 2.3 GB/sec (200 TB/day)
- Scalyr ingests 100% of that log data with zero loss
- 99th percentile ingestion delay (until logs are available for querying) is less than 2 seconds
- Simultaneously, the cluster processes queries at a rate equivalent to normal usage
- Queries return complete results with no missing data
- 90% of queries complete in less than 1 second
- Resource usage is consistent with our cost targets.

The detail below describes how we executed the benchmarking test and how we met the success criteria using Scalyr's new architecture.

Test Setup

The test is set up to generate, ingest, and parse a massive volume of logs, while simultaneously searching the ingested and accumulated data. We designed the tests based on production usage patterns to strive for realism and stress the system in the ways users stress it.

Fig 1 shows a conceptual diagram of the major functions and data flow within Scalyr. Scalyr is SaaS, and users are only concerned with the input and output, but to explain the performance considerations, we want to show what happens inside the service. We ran the load test on the same infrastructure we use for our customers.

