

The Need for Speed

Why Scalyr doesn't use keyword indexes

Introduction:

The world runs on code.

High growth, cloud-native, application-first organizations are fast-changing, dynamic, and equate system uptime and performance with revenue and customer satisfaction. Downtime is money out the door for the business. While nothing is more satisfying to an engineer than solving a hard problem and creating a new feature or application, nothing is more frustrating than spending their day finding and fixing problems in the code that they or someone else wrote, especially under time pressure. Scalyr was founded by engineers for engineers to provide real-time, granular visibility to application and infrastructure problems. We reduce the mean time to WTF.

Sound too good to be true?

Like anything in the modern world, the right architecture can create breakthrough results. While other tools in the observability space use keyword indexing, a technology optimized for searching long form, seldom-changing documents, at Scalyr we reimaged the problem to come up with a unique approach to solving the problem while building our product. From the beginning, we designed every aspect of our architecture to be performant, at scale, at processing and searching machine data.

Here's how we do it, and why it matters to you.

Unlike traditional log management solutions and observability tools that rely on keyword indexes, we took a different approach. Instead of wasting compute resources and, more importantly, time to build and maintain one or more indexes, we immediately parse incoming data and ingest it into a more efficient NoSQL columnar database. This brute-force approach gives our customers the performance, responsiveness and usability that they need to quickly troubleshoot application and modern infrastructure problems. This gives customers several key advantages when choosing Scalyr over legacy tools:

1. Our customers can use their data immediately.

We give our users a real-time view into their data and give them a true Live Tail so they can see the effect of code changes in a matter of seconds. Keyword indexes can't do this because they are too busy building (and rebuilding) the index.

2. Our customers can run blazing fast ad hoc searches

On any and all log entries, across any field and/or value that they see in their logs. Most of the time we only need to search in the relevant column. 96%+ of searches return results in less than one second. Keyword indexes can only provide this level of performance for the preselected terms, not for ad hoc searches which are much more common when troubleshooting applications and modern infrastructure.

3. We provide a simple yet powerful point-and-click user experience that is easy-to-use and easy-to-learn.

Facet-based search enables powerful data exploration that lets users go in, out and deep into their data without the need for queries. Keyword indexes can't provide the performance needed across all the incoming data to do this, nor can they handle a large volume of queries from hundreds of users without a dedicated large infrastructure requiring constant fine-tuning.

See for yourself how our backend technology powers search speeds that exceed 1.5TB/sec.

Give it a try - scalyr.com/signup