

Inside the Future of Time-Series Databases with Tobias Oetiker (H1)

“We must learn from the past to build a brighter future.” Many well-known quotes echo the idea that to build something better, we first need to understand what came before. With our goal of making an impact in the time-series space, we decided to do just that.

We had the opportunity to speak with the creator of the first open-source time-series database, RRDTool. This database was developed by Swiss software engineer Tobias Oetiker in 1999.

We wanted to learn more about his story, what motivated him to build RRDTool, and his perspective on the future of time-series databases. Along the way, we also discovered that RRDTool wasn’t his first successful project.

The Evolution of Time-Series Databases from MRTG to RRDTool (H2)

During our interview with Tobias, we learned that his earlier project was MRTG. He created it in 1995 while working as a system administrator at De Montfort University in Leicester. His role involved monitoring and managing the university’s network.

To better understand and troubleshoot recurring issues with a slow and unstable internet connection, Tobias developed MRTG.

MRTG is still available today and works by recording network traffic and displaying it on a webpage. While it was effective for network monitoring, Tobias wanted to build something more flexible—something that could support a wider range of time-series use cases.

This led him to the idea of a standalone time-series database.

That idea became reality in the summer of 1999, when Tobias participated in a work exchange in San Diego with CAIDA. There, he continued developing his concept and released the first version of RRDTool.

With that, he created the first open-source time-series database—one whose legacy continues today.

How Modern Time-Series Databases Compare Including RRDTool and InfluxDB (H2)

Although RRDTool is open source and didn't bring Tobias significant financial gain, it earned him strong recognition in the computer science community. His work played an important role in the evolution of big data storage and analysis.

RRDTool has been used in a wide range of applications, including metric monitoring, satellite tracking, and weather data collection.

Even today, it offers some advantages compared to newer database solutions. While it may require more setup and planning upfront, once data is properly stored, it's straightforward to retrieve and work with.

In contrast, some modern databases like InfluxDB make it easy to ingest data quickly. However, retrieving and analyzing that data later can become more complex if it wasn't structured correctly from the start.

Another key advantage of RRDTool is its efficiency for long-running systems. It automatically removes old data, preventing storage from growing indefinitely. Many modern databases follow a more traditional model where data continues to accumulate unless manually managed.

Over time, we've seen major improvements in data storage, monitoring, and troubleshooting—often with free or open-source tools. Still, one question remains:

How do we extract the most value from all this data?

Where Time-Series Data Is Headed with Automation and Advanced Analytics (H2)

Tobias shared valuable insights on how time-series databases have evolved, but we were especially interested in his thoughts on what comes next.

He explained that the real value of time-series data lies in understanding what “normal” looks like. This baseline allows us to detect issues, compare behavior, and respond when something goes wrong, especially as systems scale and data volumes grow.

While we've made strong progress in collecting and storing data, Tobias pointed out a key gap. We still struggle to identify what truly matters within massive datasets.

He has considered building a new version of RRDTool, but only if it includes stronger analytical capabilities with minimal user input.

In his view, the future of time-series data lies in smarter analytics and automation.

Unlocking More Value from Time-Series Data with Smarter Tools (H2)

It was a privilege to speak with someone as experienced and thoughtful as Tobias. His insights gave us a lot to reflect on.

One interesting takeaway is how open-source projects like RRDTool and Graphite often begin with a specific goal but evolve into widely used tools with broader applications.

At Shapelets, our focus has always been on building practical solutions for the technology community. That's why we released Khiva in 2018, an open-source library of time-series algorithms optimized for both GPU and CPU environments.

Our goal is to empower developers, engineers, and data scientists.

Looking ahead, we strongly agree with Tobias's perspective. As data volumes continue to grow, the real challenge is not collecting data, but finding the most meaningful insights within it.

That's the problem we're focused on solving.

We aim to provide tools that help users discover valuable patterns and insights within large time-series datasets. By combining our platform with advanced algorithms, we believe we're helping move the field forward and contributing to the future of time-series data.