XM Cyber-Speedb Case Study

Resolving RocksDB Memory Management Challenges

Customer Profile

XM Cyber, a Schwarz Group Company, is aiming to change the way organizations approach cyber risk. With XM Cyber's Attack Path Management customers can continuously see their on-premise and cloud networks through the eyes of an attacker to identify and prepare for possible attacks. XM Cyber's products reduce the internal attack surface by uncovering hidden attack paths, cutting them off at key junctures, and eradicating risk with minimal effort.

No Stopping

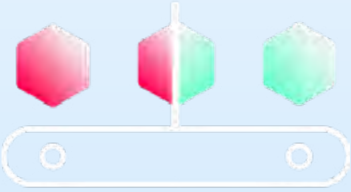

No Guesswork

No Blind Spots

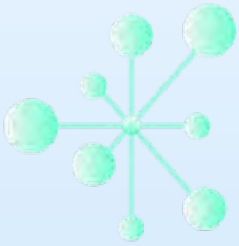

XM Cyber

The Challenge

To obtain a comprehensive view of all critical attack paths across hybrid networks, XM Cyber uses a combination of agentless and sensor-based approaches. In the latter case, OS-specific sensors are installed on endpoints. The sensors collect metadata including OS information, network, registry configuration, file versions, etc. to help formulate simulated cyber attacks. Session metadata is collected from all computers in the network in a manner that allows for investigating and identifying which sessions could be potentially compromised.

XM Cyber uses Apache Flink as its stream processing engine to handle the data and metadata generated by the XM Cyber sensors. Designed to process large scale data streams and deliver quick insights into processed data within the streaming application, Flink enables XM Cyber to find potential attacks in real-time by looking at metadata as it flows into the systems and checking it against existing data. However, as more session analysis was done and more queries were processed, XM Cyber has been faced with performance issues due to Flink's underlying data architecture.

Stateful Data operations in Flink are handled by RocksDB, a Key Value Storage (KVS) engine that uses a log-structured merge-tree (LSM) tree as its data structure. When a key-value state is registered, RocksDB maps it to a column family. The problem that XM Cyber encountered was that each newly added column family consumed significant system resources. As the number of column families increased, the RocksDB storage engine began to stutter, resulting in significant performance overhead. At a certain point, the challenge became a critical risk to the service level that the company prides itself on.

When the XM Cyber team researched the problem, the conclusion was that the problem was likely related to memory management in RocksDB when dealing with multiple column families. RocksDB column families share a common pre-write memory space (AKA write buffer manager). Hence, even when dealing with relatively small amounts of data, with the unique characteristics of the XM Cyber dataset, memory consumption increases significantly. Notably, XM Cyber noticed that the problem appeared when going over 10-20 GB of uncompressed memory.

Furthermore, RocksDB memory management issues have also led to performance degradation. As most of the RAM was used by the cache, large amounts of memory intended for indexes and read cache had to be swapped to disk, resulting in significant reduction of IOPS.

To address this issue, XM Cyber tried to use workarounds like splitting the major jobs into many smaller, vertical jobs, or building a hashmap on top of RocksDB. However, none of these solutions eliminated the need for XM Cyber to shard their data extensively, sometimes up to 8 partitions, resulting in mounting costs and inefficient resource allocation and memory usage. In addition, the need to continuously tune, customize and maintain the RocksDB engine resulted in increased allocation of precious engineering resources and time.

XM Cyber sought a different, innovative approach to overcome this challenge. Their initial research did not bring many alternatives, mostly very costly in-memory options and solutions that required highly accurate capacity planning in a field where the risk of error can be catastrophic. They were struggling to find an easy-to-implement and effective solution that could resolve their issues without adding more complexity.

Then they noticed an interesting post on Linkedin.

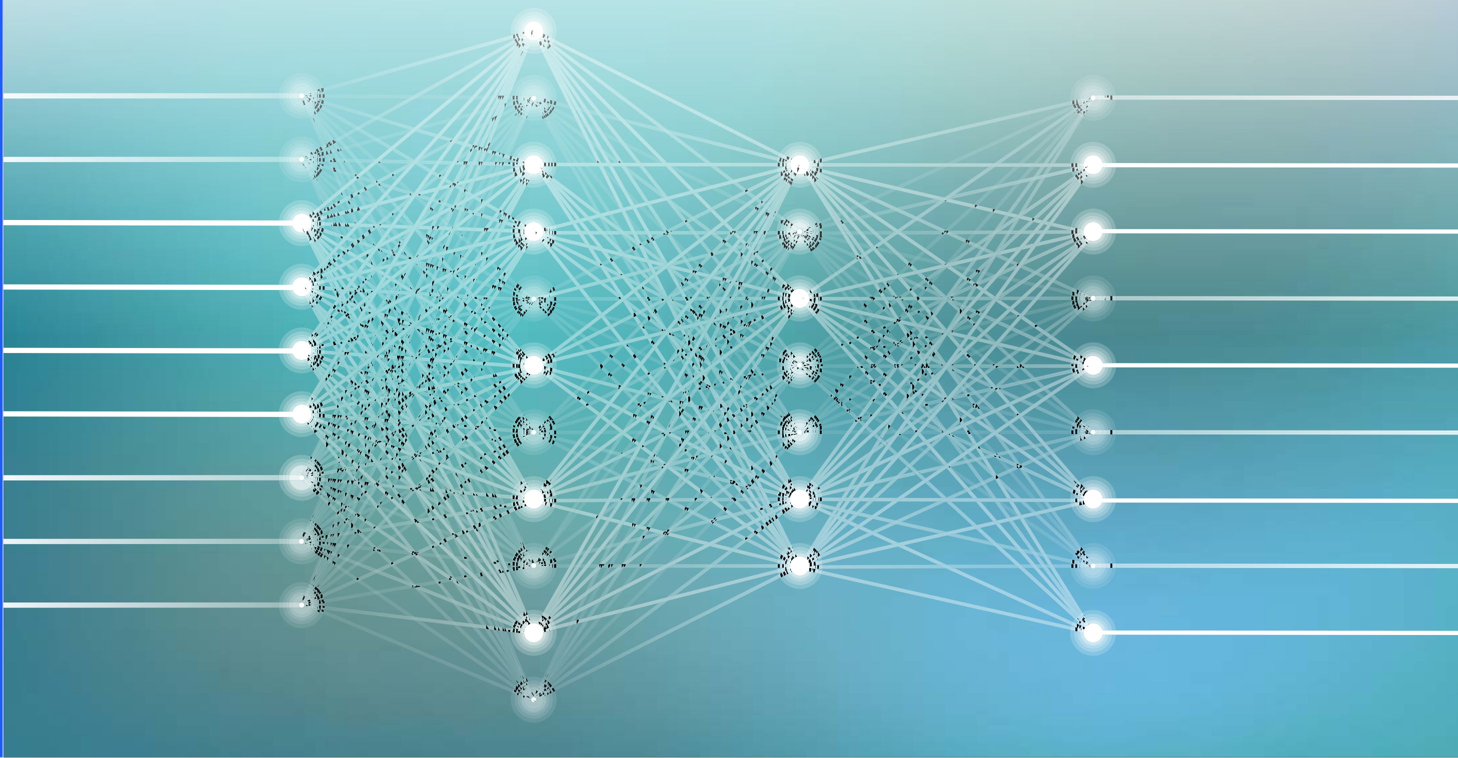

The Solution

Designed as a drop-in replacement for RocksDB that allows up to 10 times more OP/s while using less resources, Speedb allowed XM Cyber to leap over the aforementioned challenges and offer their customers an ideal solution that scales seamlessly. The main goal was to solve memory bottlenecks for small and medium sized customers and reduce parallelism with larger enterprise customers.

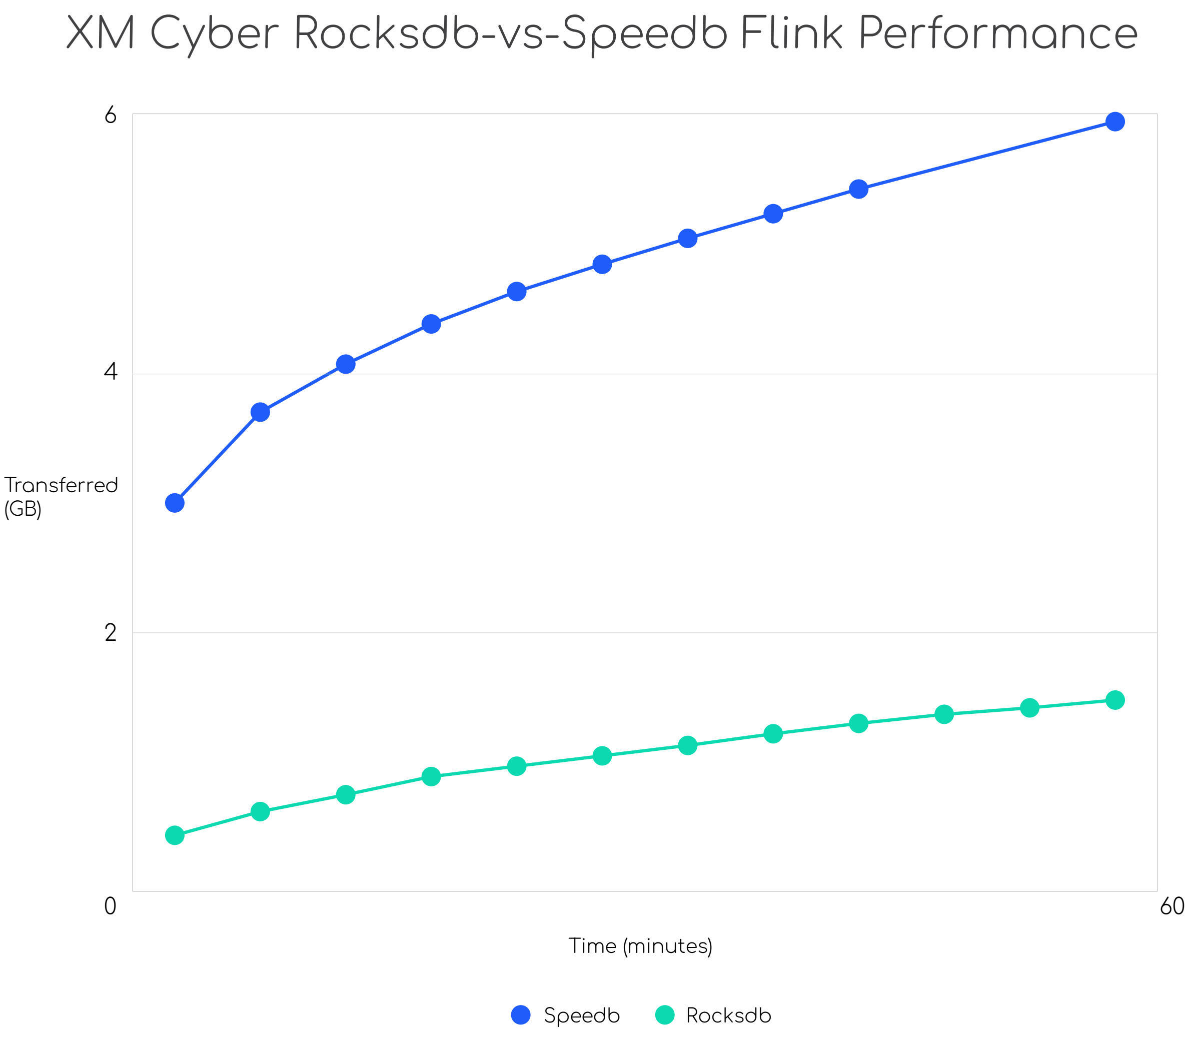

According to Yaron Shani, Senior Researcher and Tech Lead at XM Cyber, "We simply replaced a few lines in the Docker files, and in minutes it was ready." He adds that, "right from the start, Speedb took ownership of the problem. They provided us with excellent support and very quick response time and availability." Shani adds that, "Speedb's data engine alongside their expertise resulted in a very fast solution to the problems we encountered, and the performance improvements were instantaneous. Speedb is now deployed in our main build that goes to all customers, large and small. During the process of working together we discovered unique problems we were unaware of, like the fact that RocksDB does not trigger compactions well with lots of deletes."

Adding Speedb to the equation resulted in various benefits to XM Cyber and its customers. Prior to that, RocksDB's memory bottlenecks have led to delays in delivering simulated attacks while certain simulations could not have been carried out at all. According to Shani, "Some techniques that were not possible with RocksDB are now possible with Speedb. Our customers can get more features and better products with increased efficiency. In various cases we saw dramatic improvements of 10 times the performance we had before switching to Speedb. It's a very dramatic increase that allows us to give our customers better products and service than ever before."

Furthermore, with Speedb, XM Cyber was able to free developers from having to constantly deal with RocksDB sharding, tuning, and other time-consuming operational tasks, so they can focus on engineering. Shani mentions that "Less human resources are now required to invest in RocksDB problems. The amount of questions and problems I need to address with the team has dropped dramatically, and instead of spending highly specialized human resources we are now focused on providing our customers with more features and developing new attack techniques."

Shani emphasizes that some of the most exciting benefits of overcoming the continuous performance challenges they had will only be measurable in the future. "Relying on Speedb as our data engine gives more confidence to the managers and engineers to move forward with new technologies instead of maintaining old technologies." Shani explains that "for cyber security companies, the bigger the network, the bigger the problems. Working with Speedb eliminates the concerns that things might crash and gives us the ability to offer customers better scale, more features and more techniques."