

# Co-op Project at Magnum Research Trading System Profiling and Optimization

Cheung Sze Yuen

Supervised by Prof. Wang Wei, Ryan Liu

Sponsored by Magnum Research Limited (AQUMON)



## Introduction

Magnum Research Limited, a Hong Kong-based wealth management firm, aims to maximize investor returns through its automated platform. The system efficiently executes manual or automated orders for diverse users, managing assets primarily in China, Hong Kong, and the US markets. However, as user demands and data flow increase, performance bottlenecks arise, and limited visibility necessitates manual probing. To address these challenges and reduce costs, a novel trading system is required. It should offer easy extension and maintenance, faster execution, and standardized profiling to evaluate performance effectively.

## Objectives

1. **Develop** standardized profiling and benchmark process for analysis and optimization while maintaining performance.
2. **Optimize** new trading system based on benchmarks for future requirements.
3. **Enhance** system for specific use-cases and known issues.
4. **Evaluate** efficiency gain of optimizations.

## Methodology

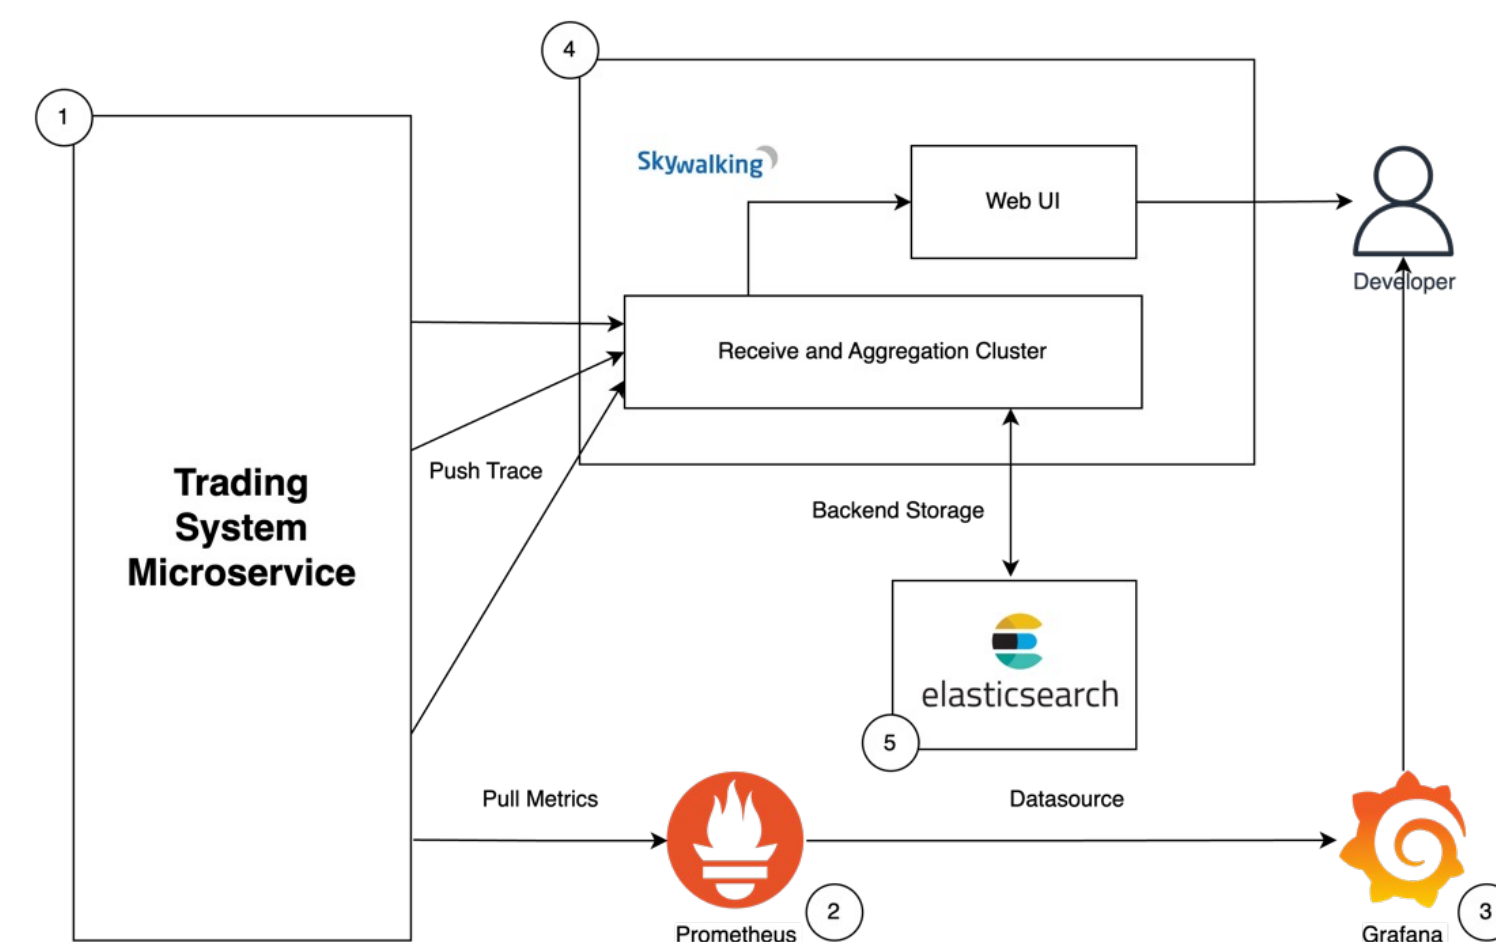
### Profiling system for key performance metrics

Our mechanism measures mainly 2 performance metrics below, for each metrics, we will record the average speed as well as the tail latency. These metrics are displayed real-time in a dashboard to holistically monitor the trading system:

1. **Order placing speed**
2. **Callback speed**

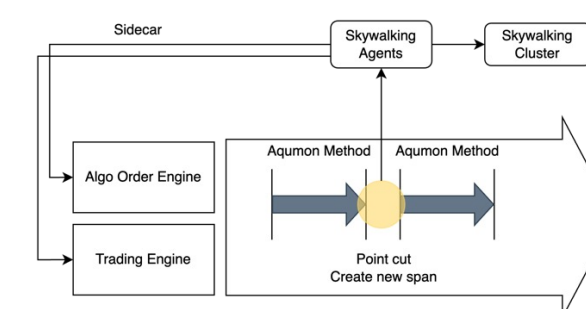
The monitoring system should primarily track the mentioned metrics and gather other valuable information like traces without heavily impacting the performance of the core system. Hence, we encourage a **lightweight** design. Users should be able to identify performance bottlenecks without compromising the overall flexibility and scalability of the monitoring system. Therefore, the system should support **custom metrics** and **custom trace spans** to cater to our specific requirements.

## Methodology



Monitoring system architecture  
1. The trading system micro-service includes multiple services e.g. Java Spring Applications, Redis, MySQL, Kafka  
2. Prometheus will actively pull the custom metrics which are annotated in the applications  
3. Grafana serves as a visualization tool to view the metrics  
4. Apache Skywalking serves as application monitoring tool specialize in distributed tracing and topology visualization, which could generate actionable insights of performance for us  
4.1. Receive and aggregation cluster are responsible for receiving and processing telemetry data from system components and handle data transformation, storage, and analysis tasks  
5. Elasticsearch act as the backend system for Skywalking since it does not have its own backend, using Elasticsearch as it is more scalable in a long-term and isolated from the core database.


Since analyzing method call time is a key focus for us in understanding the trace data, we utilized Spring Aspect-Oriented Programming to inject an Aspect code into the common package that is utilized by all microservices. This injected Aspect intercepts all method calls within the package `'com.magnumresearch.aqumon.*'` (our source code package) and assigns the method name as the span name.




## Result

Our non-invasive monitoring system, leveraging Prometheus, established crucial metrics and provided valuable insights to pinpoint and optimize bottlenecks. The automated test suite enabled easy load testing and comparison of latencies. Focusing on one task, we achieved significant results in optimizing low-level endpoints, **reducing tail latency by over 94%** and **median latency by nearly 20%**. The order placing response time **reduced 36% in median** and **77% on P95 and P99**.


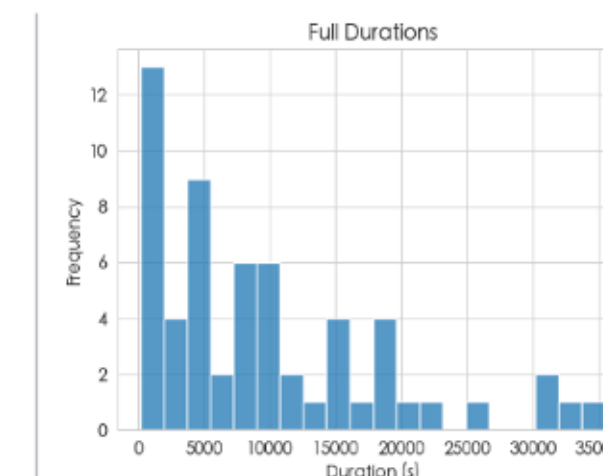
### Order Placing Speed



### Low Level Endpoint Speed



Before



## Conclusion

Our novel monitoring approach using a non-invasive system and Prometheus has optimized our trading system, reducing latency significantly. It provides valuable insights, identifies infrastructure issues, and ensures stability without impacting core trading. This system enhances **observability** and **reliability**, paving the way for further improvements to achieve overall success and stability.