



Contribution ID: 128

Type: **not specified**

Enhancing Homa Linux for Efficient RPC Transportation

Wednesday 15 November 2023 10:00 (30 minutes)

Homa, a unique transport protocol created specifically for hype-scale datacenters, provides optimized round-trip performance for request/reply messages. An in-depth evaluation of the Homa Linux module in contrast to TCP showed a considerable decrease in latency with RPC application benchmarks. Furthermore, our analysis of gRPC operating over Homa versus gRPC over TCP revealed significant benefits, specifically for smaller RPC messages (less than 20k), in both latency and throughput.

Despite these advantages, Homa's broader use as a standard RPC transport protocol is constrained by two main challenges:

1. Constraints of Message-based Interface: Homa's message-based interface poses hurdles for efficient pipelining, as it demands waiting for the complete message to ensure its full delivery to the application. This results in relatively low throughput for larger RPC messages (average size over 20k) compared with TCP.
2. Exclusive Unary RPC Support: Currently, Homa only accommodates unary RPC, where a client sends a single request message and waits for one response. The absence of support for Bidirectional Streaming RPC, which allows full-duplex message exchange between client and server, limits Homa's adaptability in certain situations.

This presentation aims to introduce our solutions to these identified obstacles, and offer practical guidance on improving the performance of Homa as a conventional RPC transport protocol. We will discuss optimal strategies to harness Homa's strengths to attain minimal RPC latency and maximum throughput. The session will incorporate real-world examples and demonstrations to clarify these concepts and highlight the benefits of using Homa as an efficient RPC transport protocol.

Primary authors: Dr LU, Xiaochun (Bytedance); ZHANG, Zijian (Bytedance)

Presenters: Dr LU, Xiaochun (Bytedance); ZHANG, Zijian (Bytedance)

Session Classification: eBPF & Networking

Track Classification: eBPF & Networking Track