



# High-throughput Small File Access for Large-scale Machine Learning Applications


Hiroki Ohtsuji<sup>1</sup>, Erika Hayashi<sup>1</sup>, Takuya Okamoto<sup>1</sup>, Eiji Yoshida<sup>1</sup>,  
Osamu Tatebe<sup>2</sup>

1)Fujitsu Limited, 2)University of Tsukuba

# High-throughput Small File Access for Large-scale Machine Learning Applications

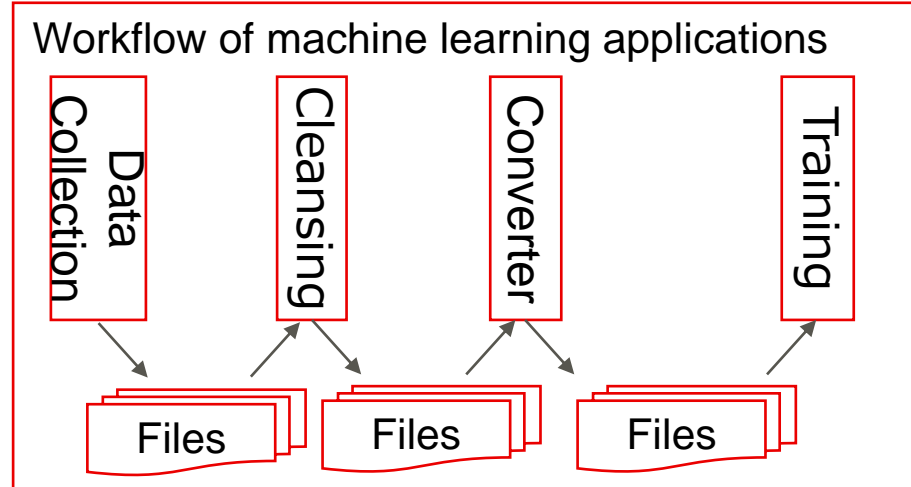
Hiroki Ohtsuji<sup>1</sup>, Erika Hayashi<sup>1</sup>, Takuya Okamoto<sup>1</sup>, Eiji Yoshida<sup>1</sup>, Osamu Tatebe<sup>2</sup>

1)Fujitsu Limited, 2)University of Tsukuba



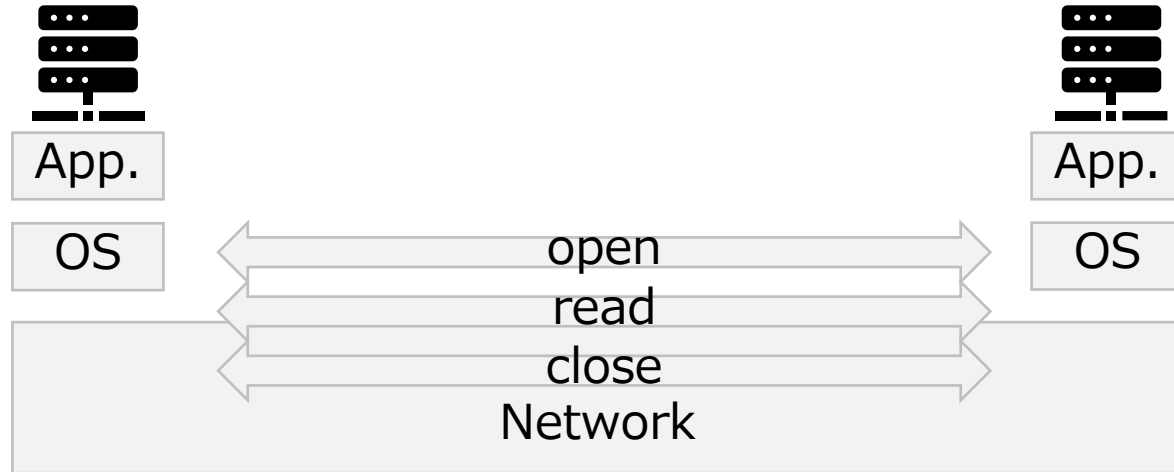
- Machine learning applications consist of multiple data processing programs
  - Proprietary tools for data collection and conversion scripts etc.
  - There is no common data format like TFRecord of TensorFlow
- Huge number of small files will be generated by a workflow

High-throughput remote small file access is required to handle the **entire workflow** of large-scale machine learning applications



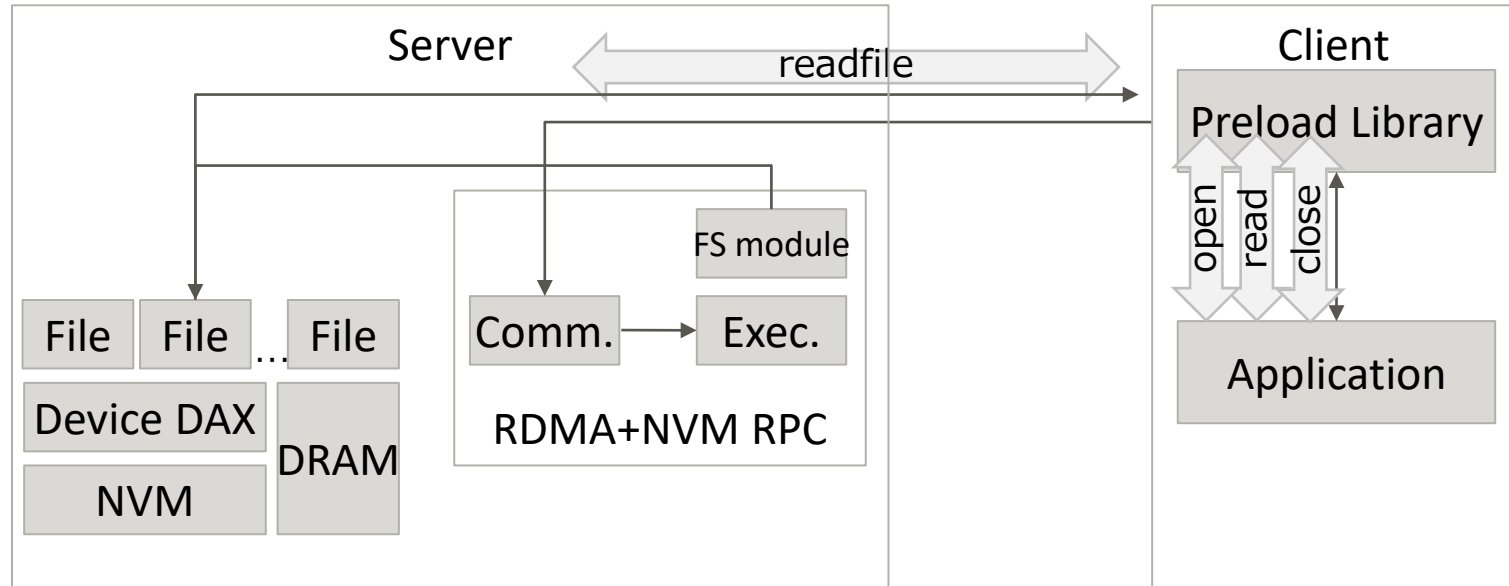
# Overhead of Remote Small File Access

- Conventional file access operation requires at least three operations for each file access
  - Open, read, close
- Remote file access operations involve multiple intermediate layers

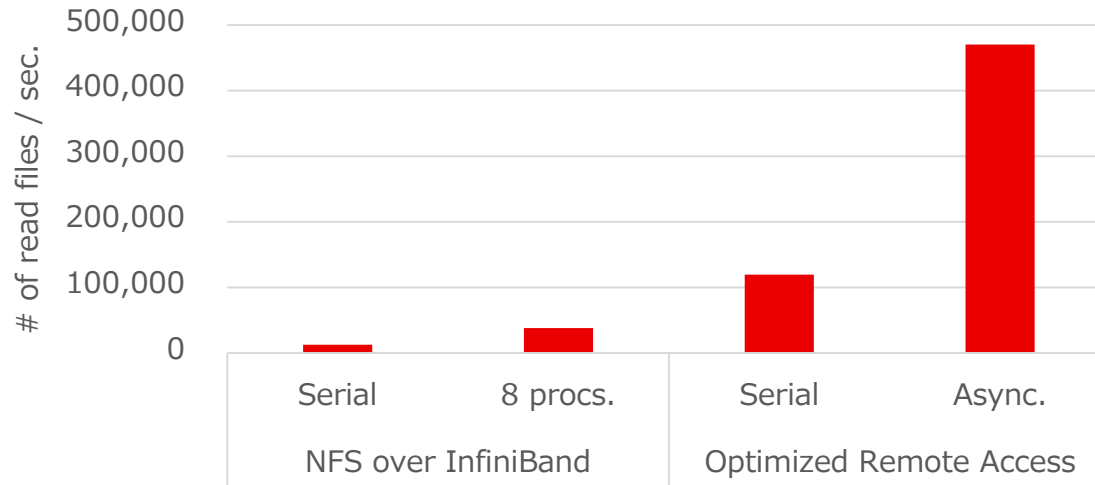


# The Lightweight Remote File Access

- The application calls multiple ordinary operations, but they are processed as a single remote operation
  - The lightweight remote file access layer handles them as readfile operations to minimize the number of network communications



- Reading 1KB files from the remote node
  - Throughput comparison between NFS over InfiniBand and the lightweight remote file access



Accelerated small file access operations by a factor of 10

## ○ Conclusion:

- The lightweight remote file access layer can provide 10 times faster throughput compared to the conventional implementation
  - Reduced the number of intermediate layers of remote operations
  - Readfile like operation reduced the number of network operations

## ○ Future work:

- Applying the lightweight remote file access to a full featured filesystem
  - Evaluating the performance with real workflows.

# Thank you

Contact:

ohtsuji.hiroki [at] fujitsu.com

