

# Scalable Distributed Table Storage Experiments on OpenCloud

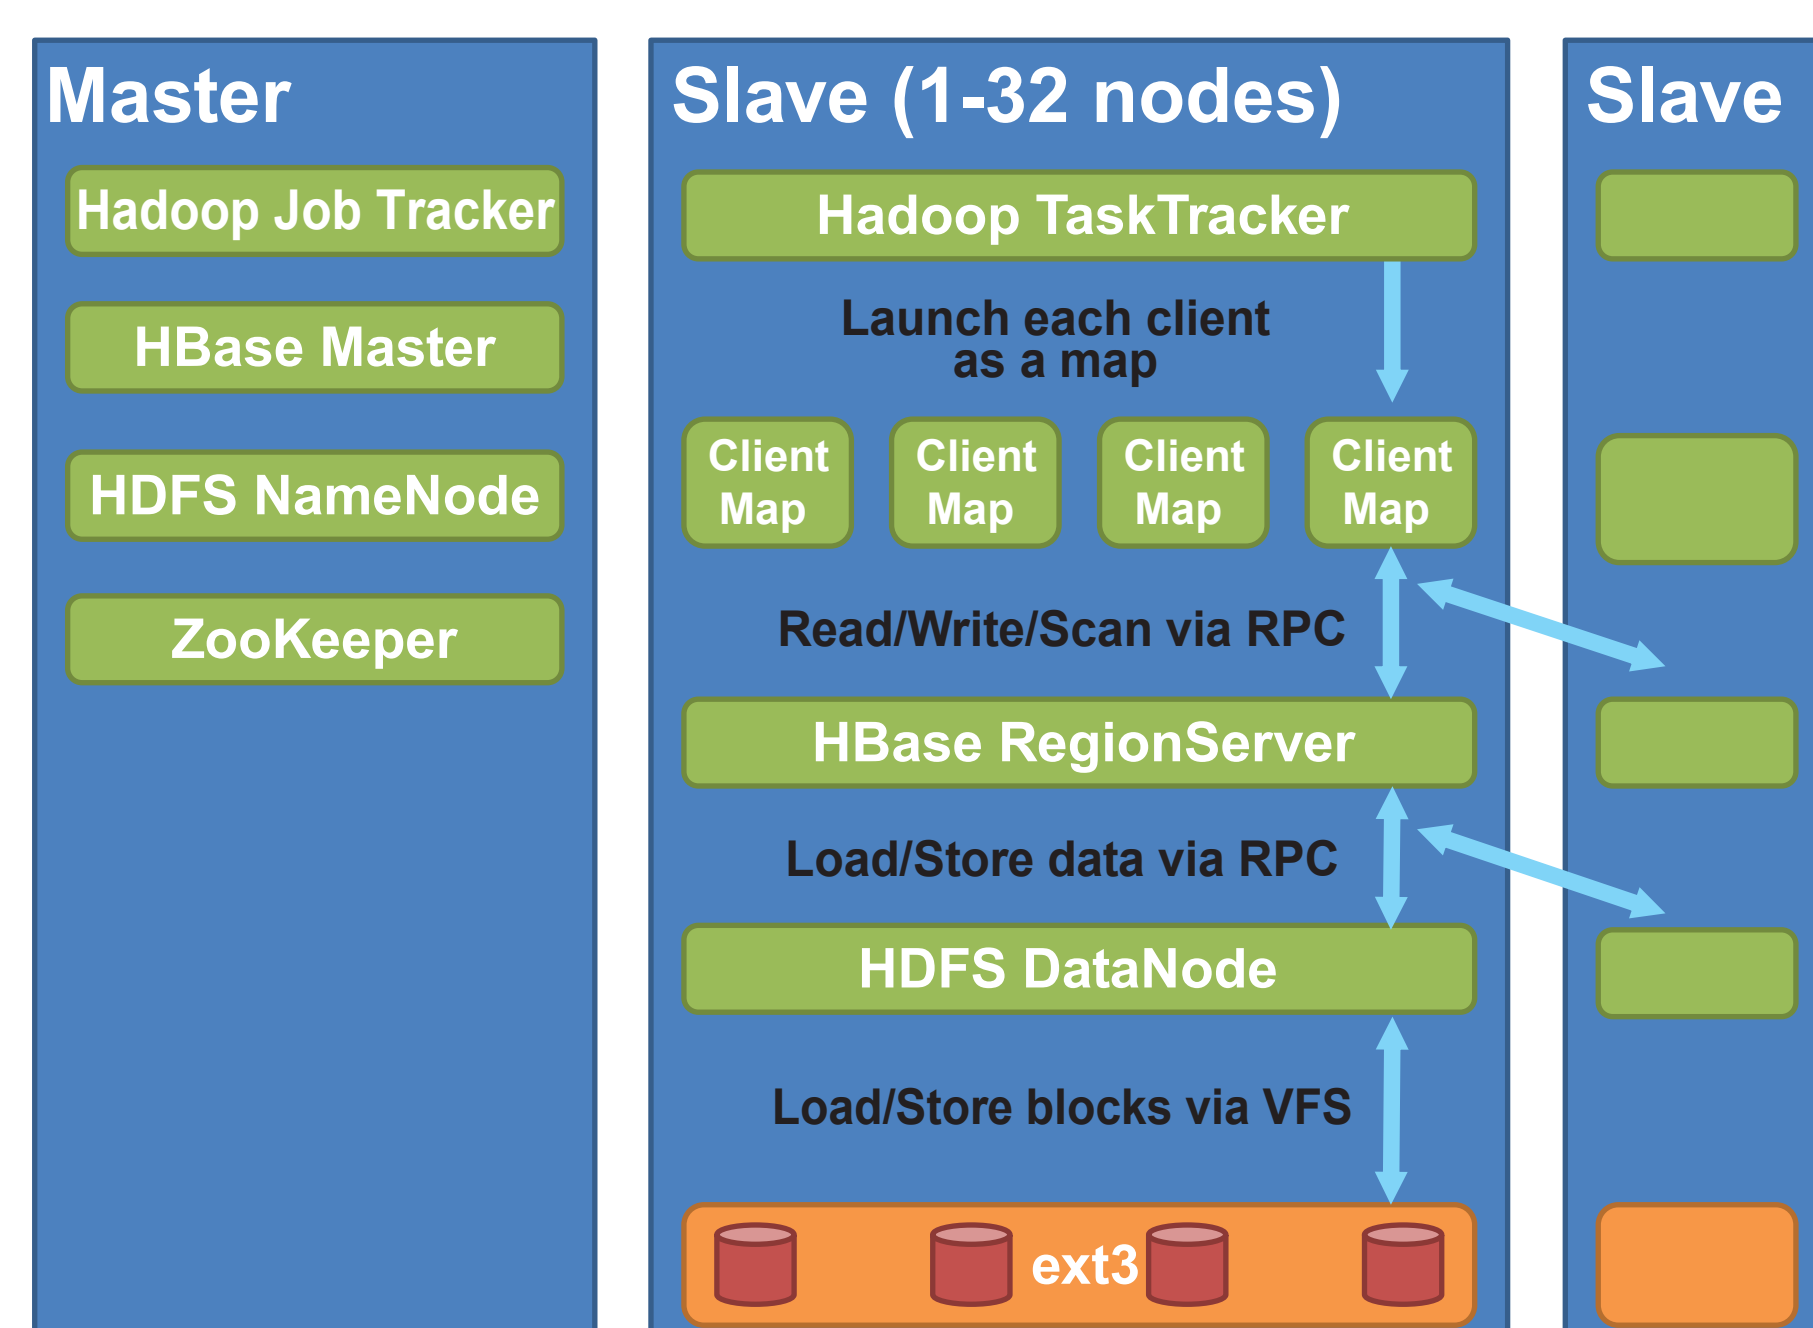
Sarun Savetsila, Garth Gibson, Wittawat Tantisiriroj, Lin Xiao, Milo Polte

## Overview

- **Mission**
  - Understand limitation of existing open source distributed table storage systems
  - Look for opportunities to improve scalability & speed
- **Why HBase?**
  - Part of Hadoop Apache project
  - Most popular open source distributed table storage software
  - Well integrated with Hadoop
  - Getting much faster with more BigTable-like implementation techniques

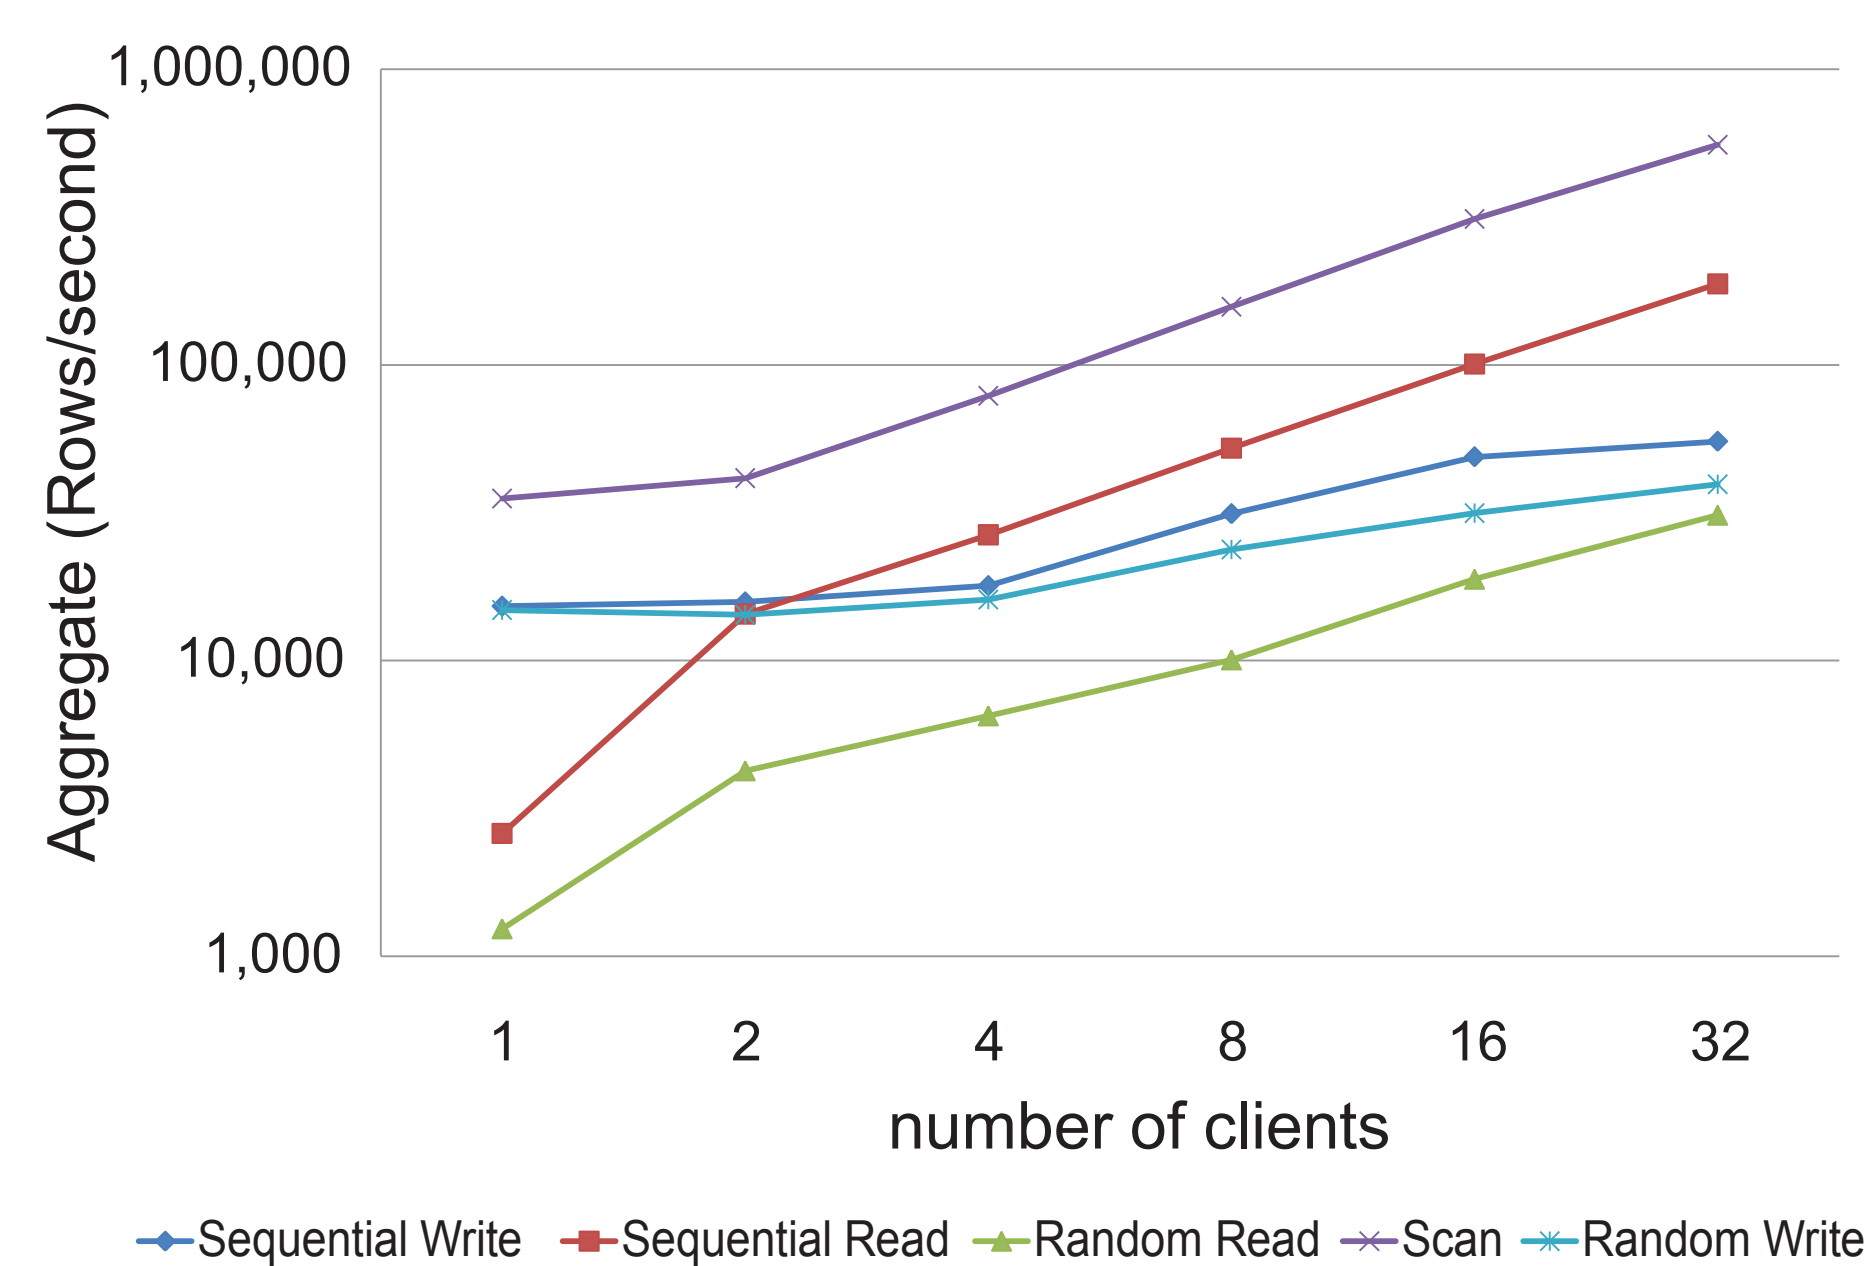
## Standard Benchmark for HBase

- Mimics BigTable paper benchmark
- A single master, N tablet servers, and N clients
- Each client performs 1 million 1 KB row operation
- Benchmark Phases:
  - seq write → seq read → rand read → scan → rand write




- Simple homogeneous concurrent access

## Experimental Result



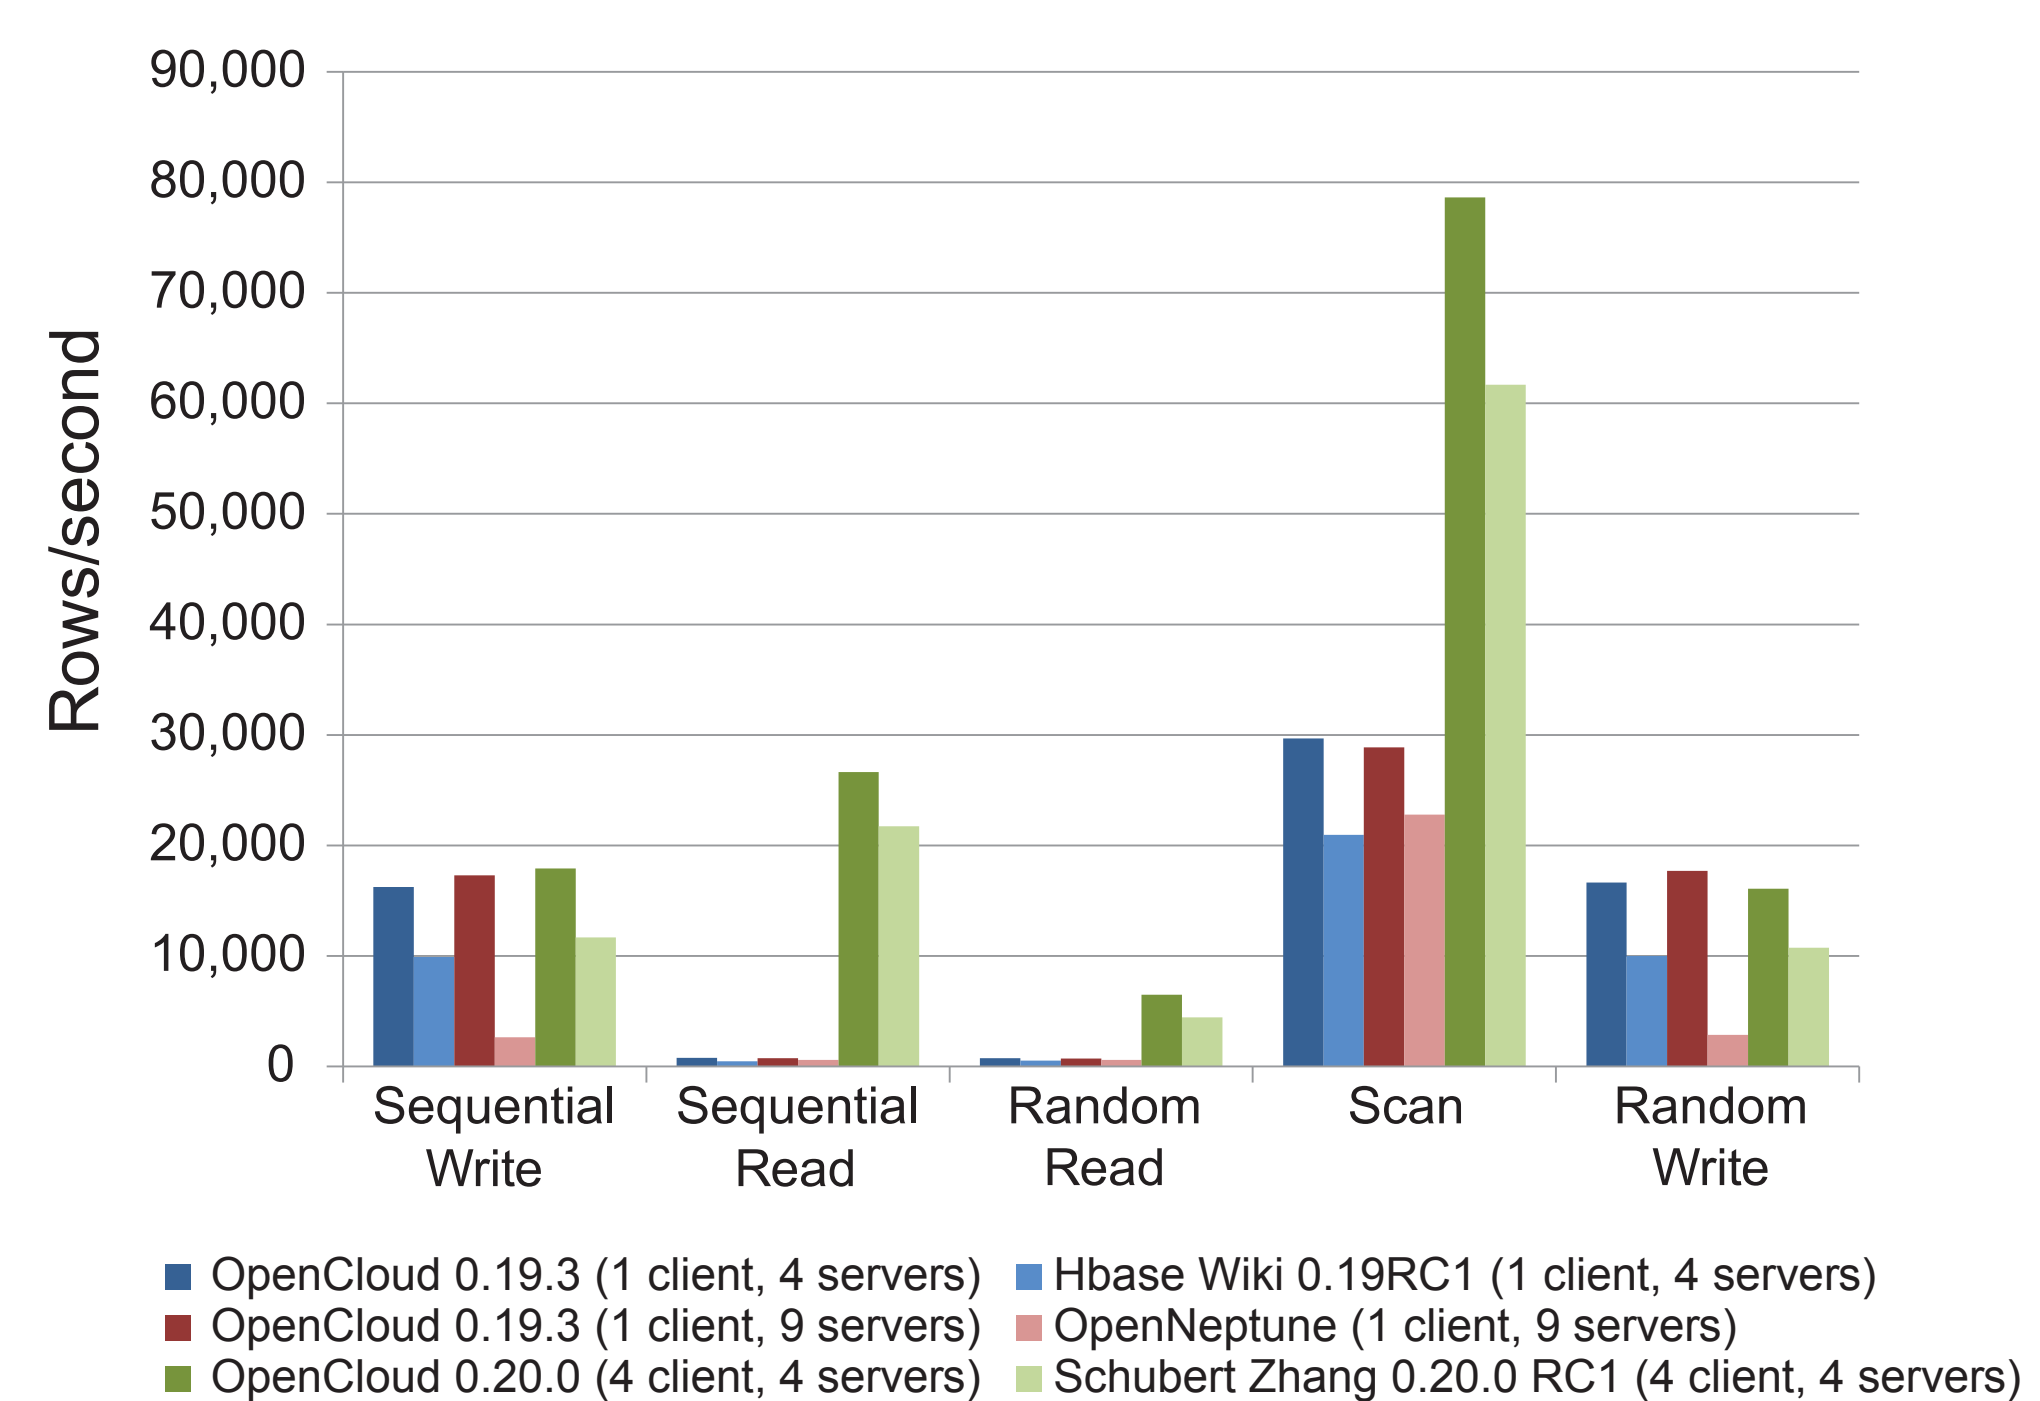
- Log-scale: mostly linear
- Strange behaviour for sequential writes

## Continuous Improvement @ HBase



- HBase is maturing
- Better storage representation, Hfile, compression, etc.

## OpenCloud Is Not A Bottleneck



- Our setup and platform are compatible with others

## Where Are We Going?

- Understand all performance issues
- Larger testing
  - Beyond OpenCloud: > 60+ nodes
  - Amazon EC2, NSF @ LANL
- Closely related project: Hypertable
- Better Benchmark
  - Yahoo's GridMix, Blog-queries benchmark
- Improve implementation
  - App-specific indices Bulk insertion, etc.