Improve Performance of Go Program

Hongchao Deng, hongchao.deng@coreos.com

About Me

Work Experience

- Cloudera, Big Data
- CoreOS, Container

Open Source Experience

- etcd, ZooKeeper
- Kubernetes

CoreOS

Modern, minimal OS
Self updates with atomic and seamless workflow

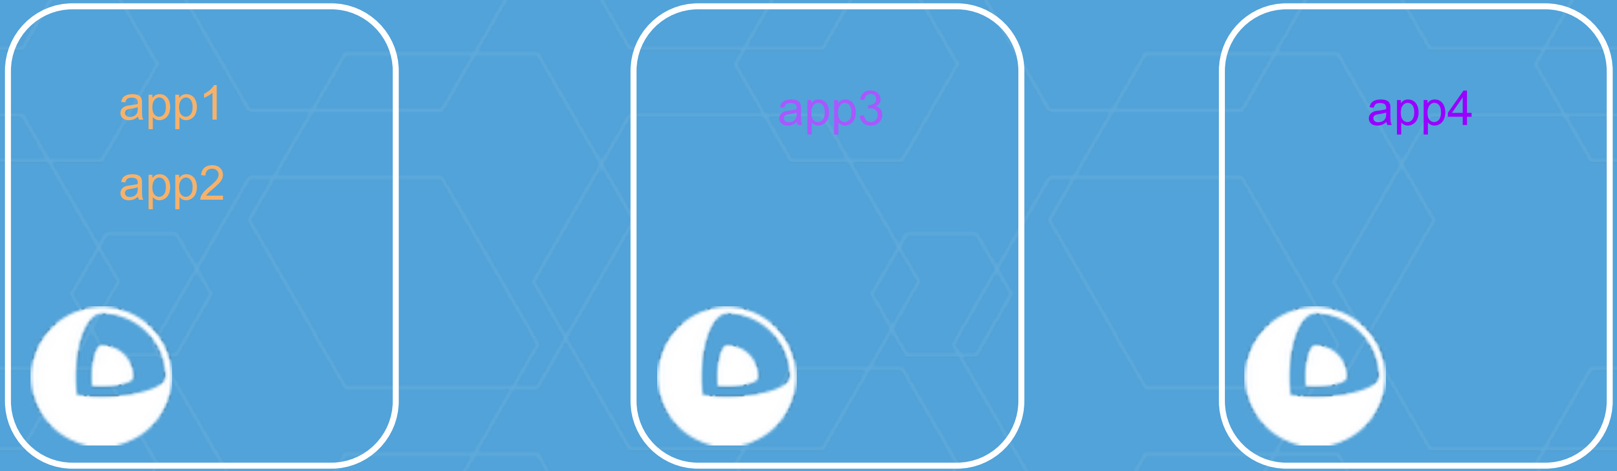

app1
app2

app3

app4

CoreOS

Modern, minimal OS

Self updates with atomic and seamless workflow

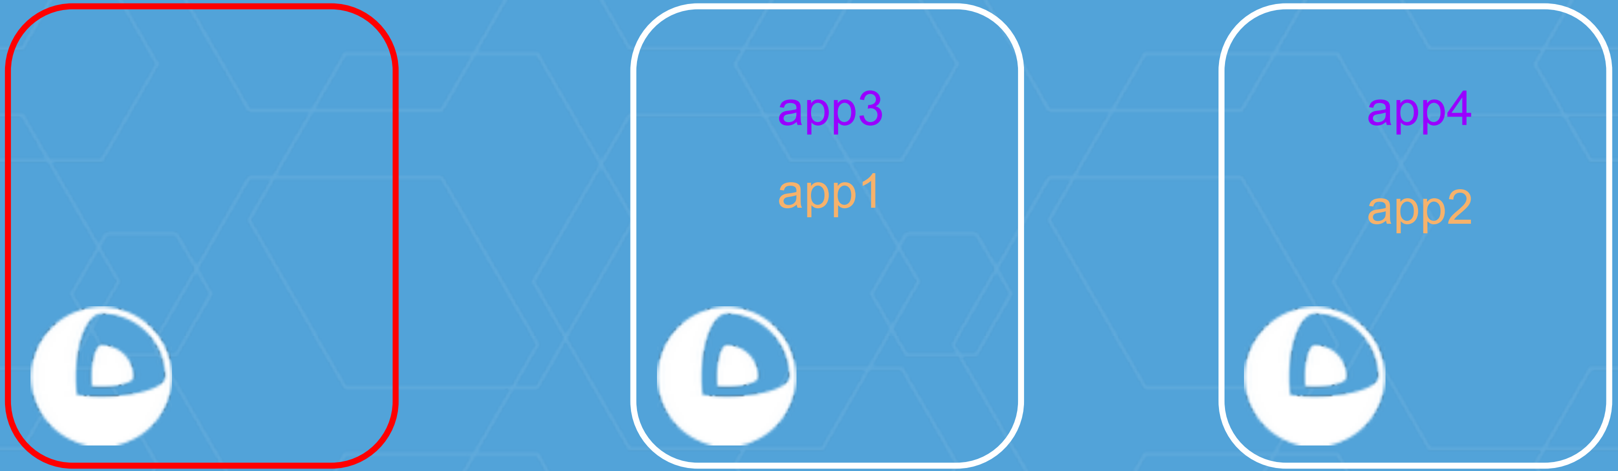

Updating...

CoreOS

Modern, minimal OS

Self updates with atomic and seamless workflow

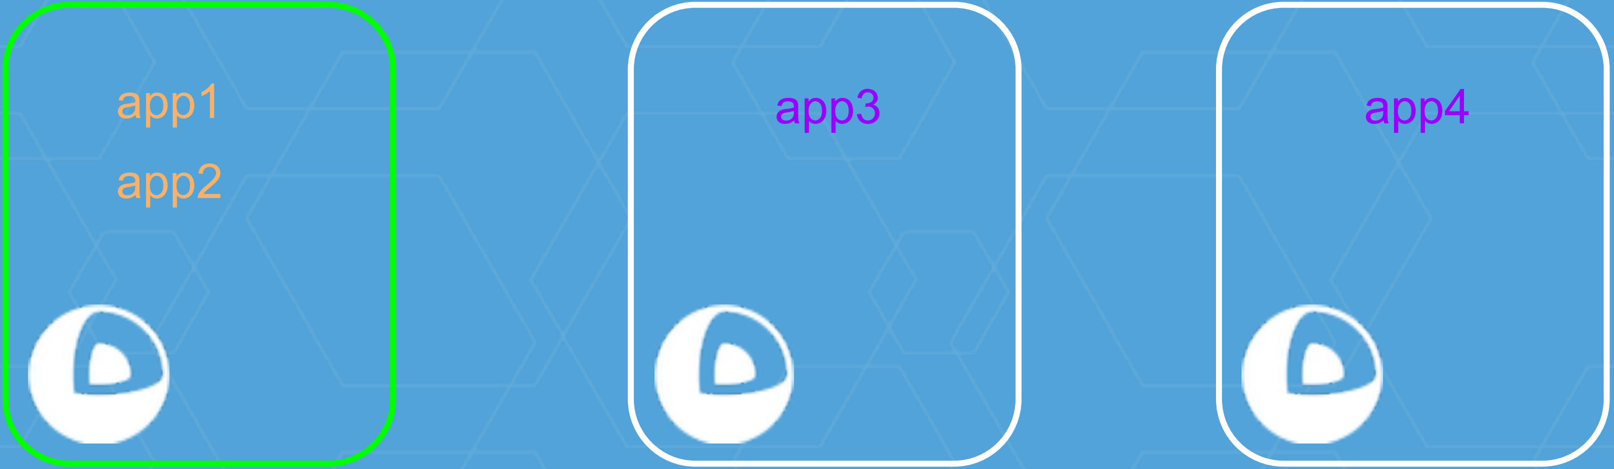

app1

app2

app3

app4

Updated!

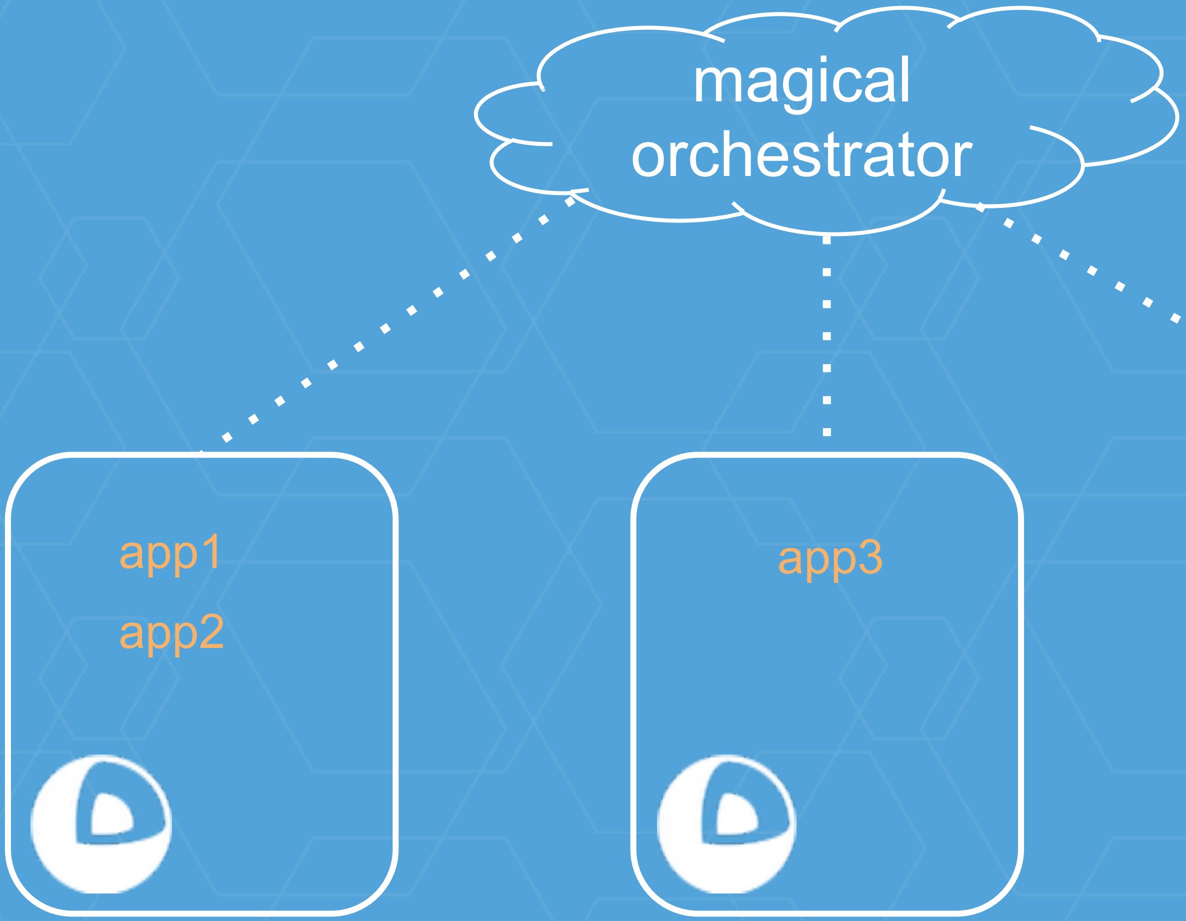
magical orchestrator
app1
app2
app3
app4

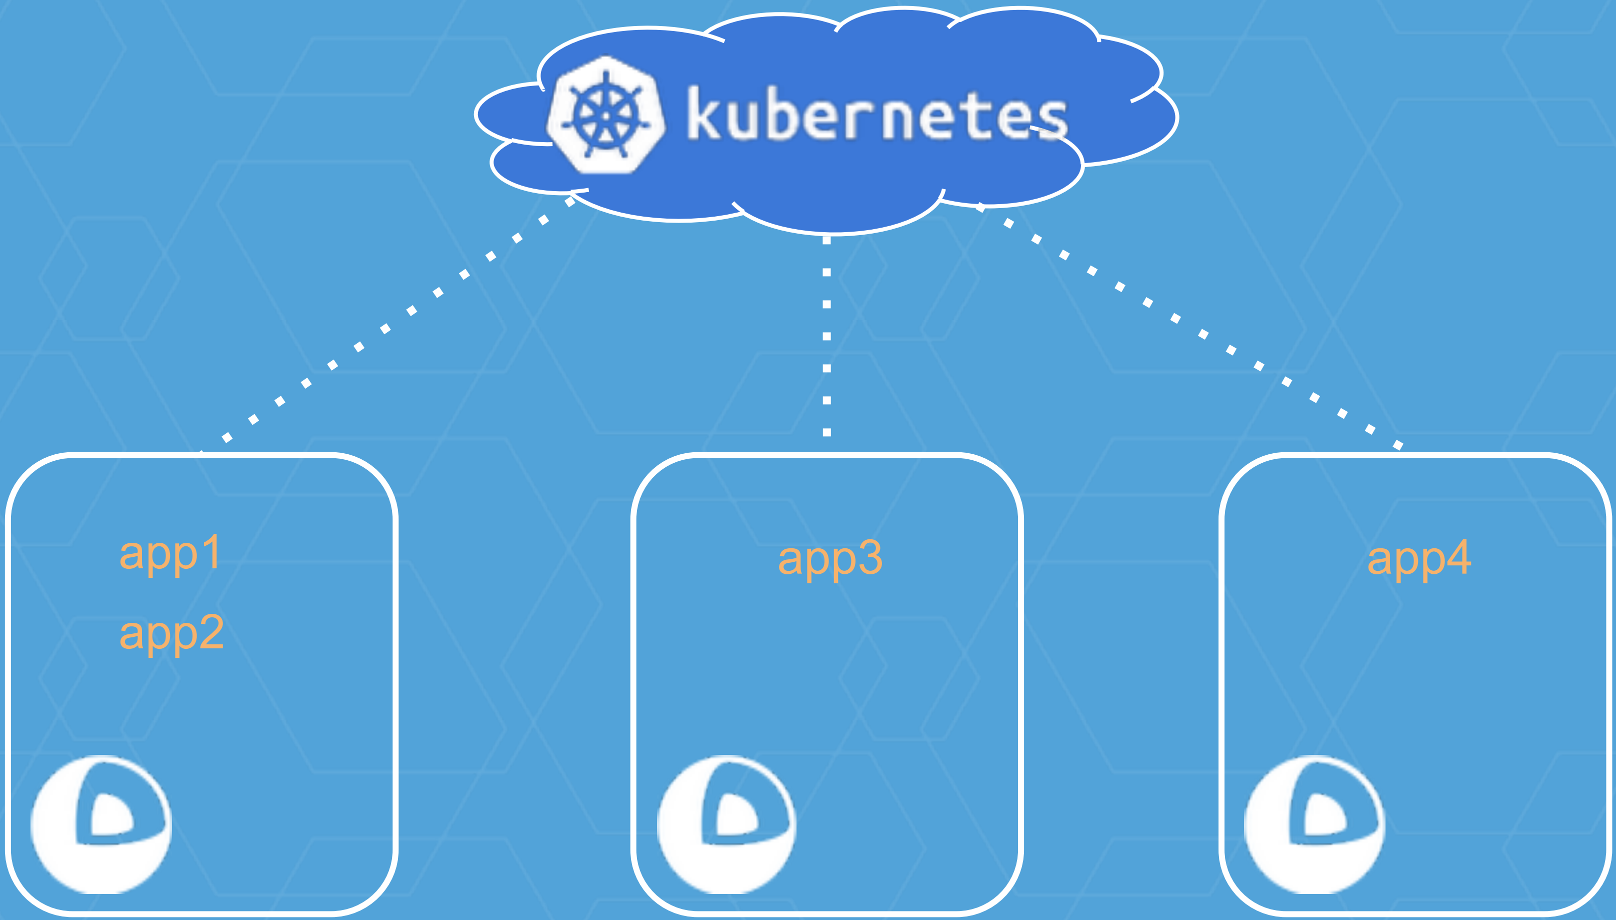

app1

app2

app3

app4

Chubby

At the Heart of Google

The Chubby lock service for loosely-coupled distributed systems

Mike Burrows, *Google Inc.*

Abstract

the Chubby lock ser-
grained lock-
for

example, the Google File System [7] uses a Chubby lock
to appoint a GFS master server, and Bigtable [3] uses
Chubby in several ways: to elect a master, to allow the
master to discover the servers it controls, and to permit
clients to find the master. In addition, both GFS and
Bigtable use Chubby as a well-known and available loca-
a small amount of meta-data; in effect they
root of their distributed data struc-
locks to partition work (at

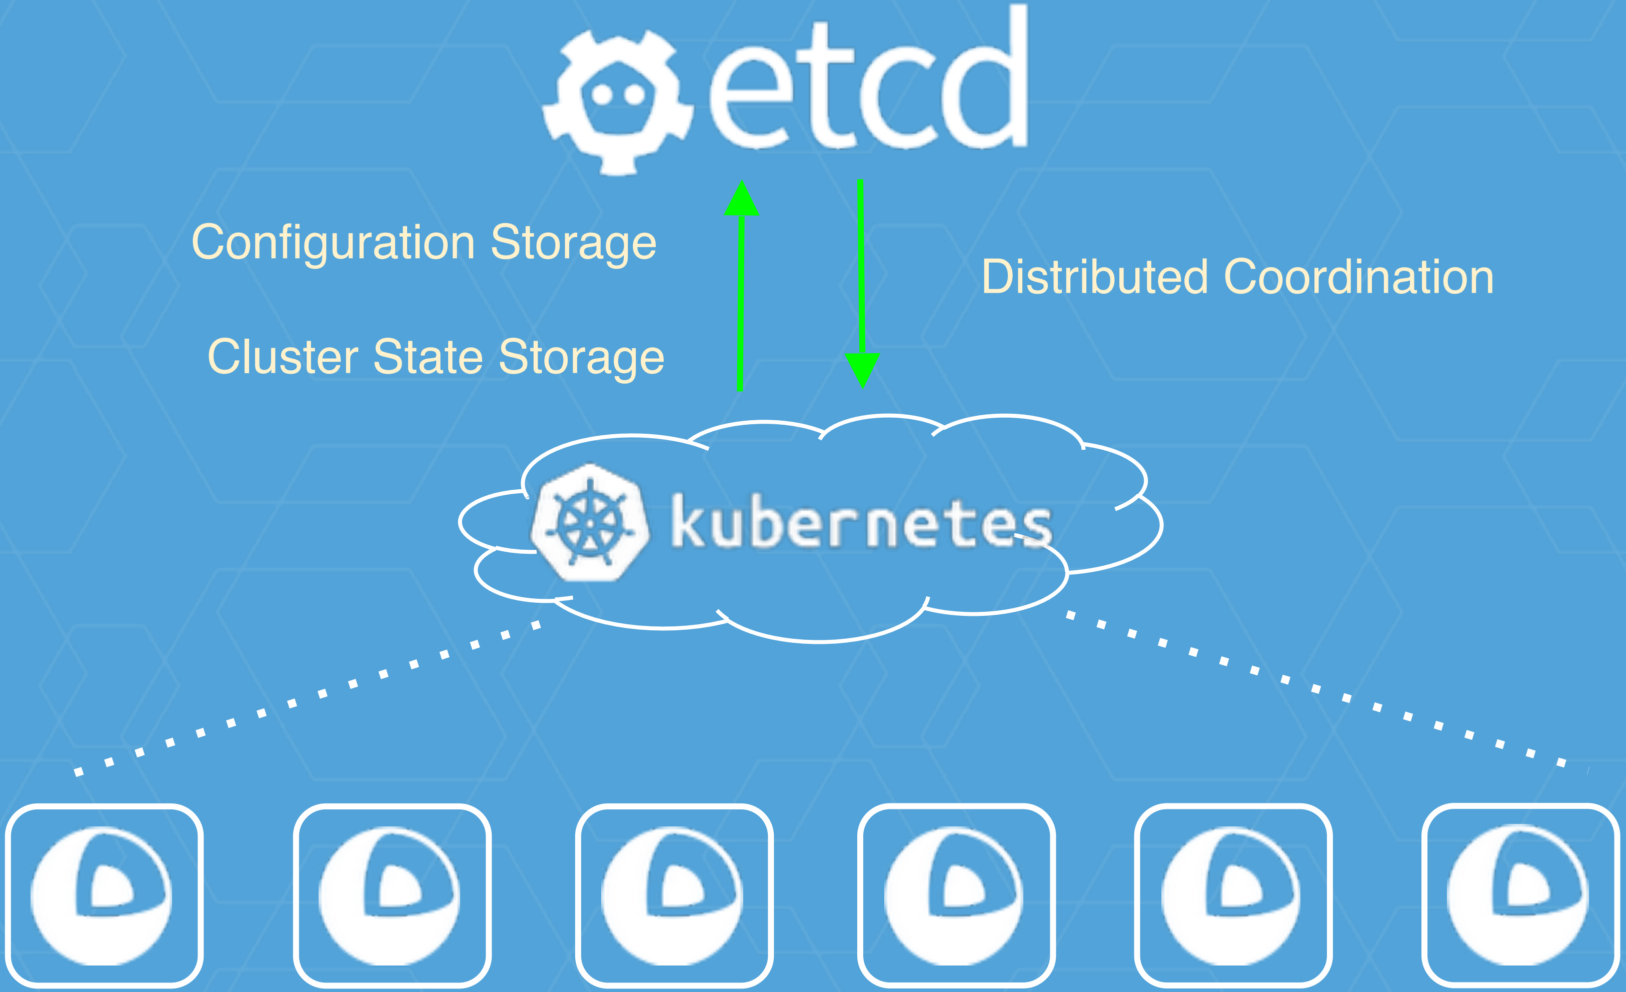

Configuration Storage

Cluster State Storage

Distributed Coordination

Core Mission

Reliable and Highly Performant Infrastructure For Everyone

OPEN SOURCE

CoreOS

etcd

rkt

flannel

kubernetes

ENTERPRISE

TECTONIC

QUAY

Performance Is Paramount

Without high performance, we couldn’t achieve our goal.

Improving performance is part of daily job:

- Overall system runs faster
- Memory footprint becomes smaller

A Core Day:
Improve Performance of Go Program

Tasks completed!

What Did We Learn?

Finding Bottlenecks Is Easy in Go

Benchmarking, Profiling

Focus on Real Problems

Tips to Optimize Performance

Algorithm

- Can I reduce the Big O time complexity?

System Internals

- Is my program cache-friendly?
- Are there too many temp objects and pressing GC?

Parallelism

- Can I parallelize my program, or a function?

Avg Throughput != SLO

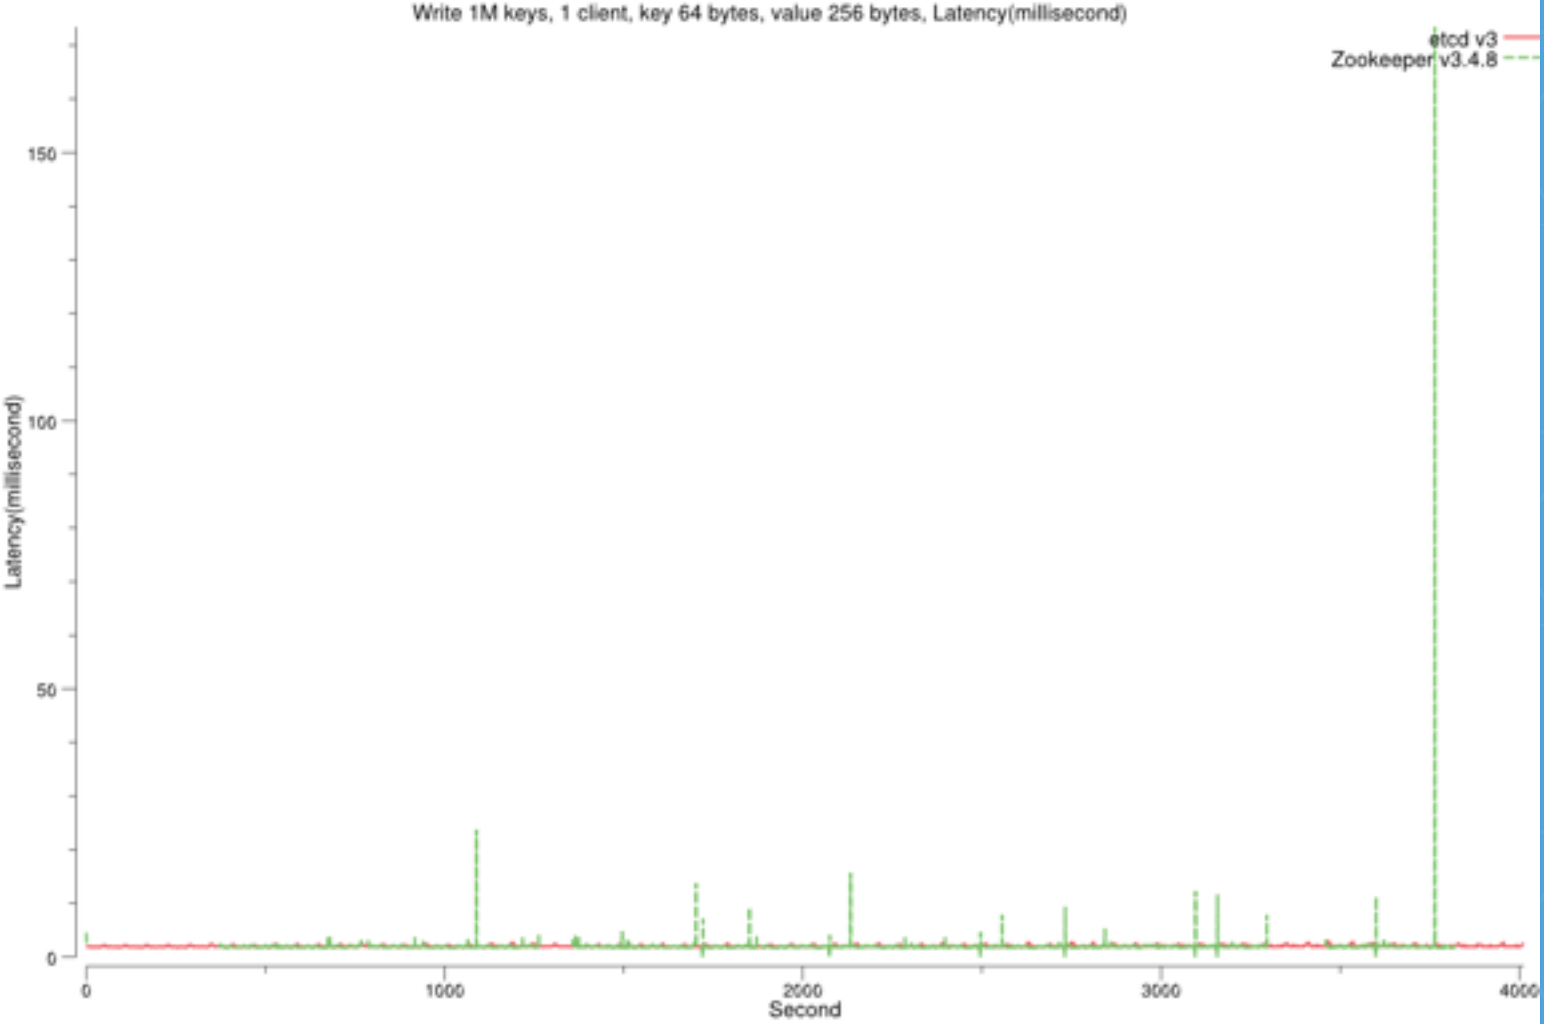

https://github.com/coreos/dbtester

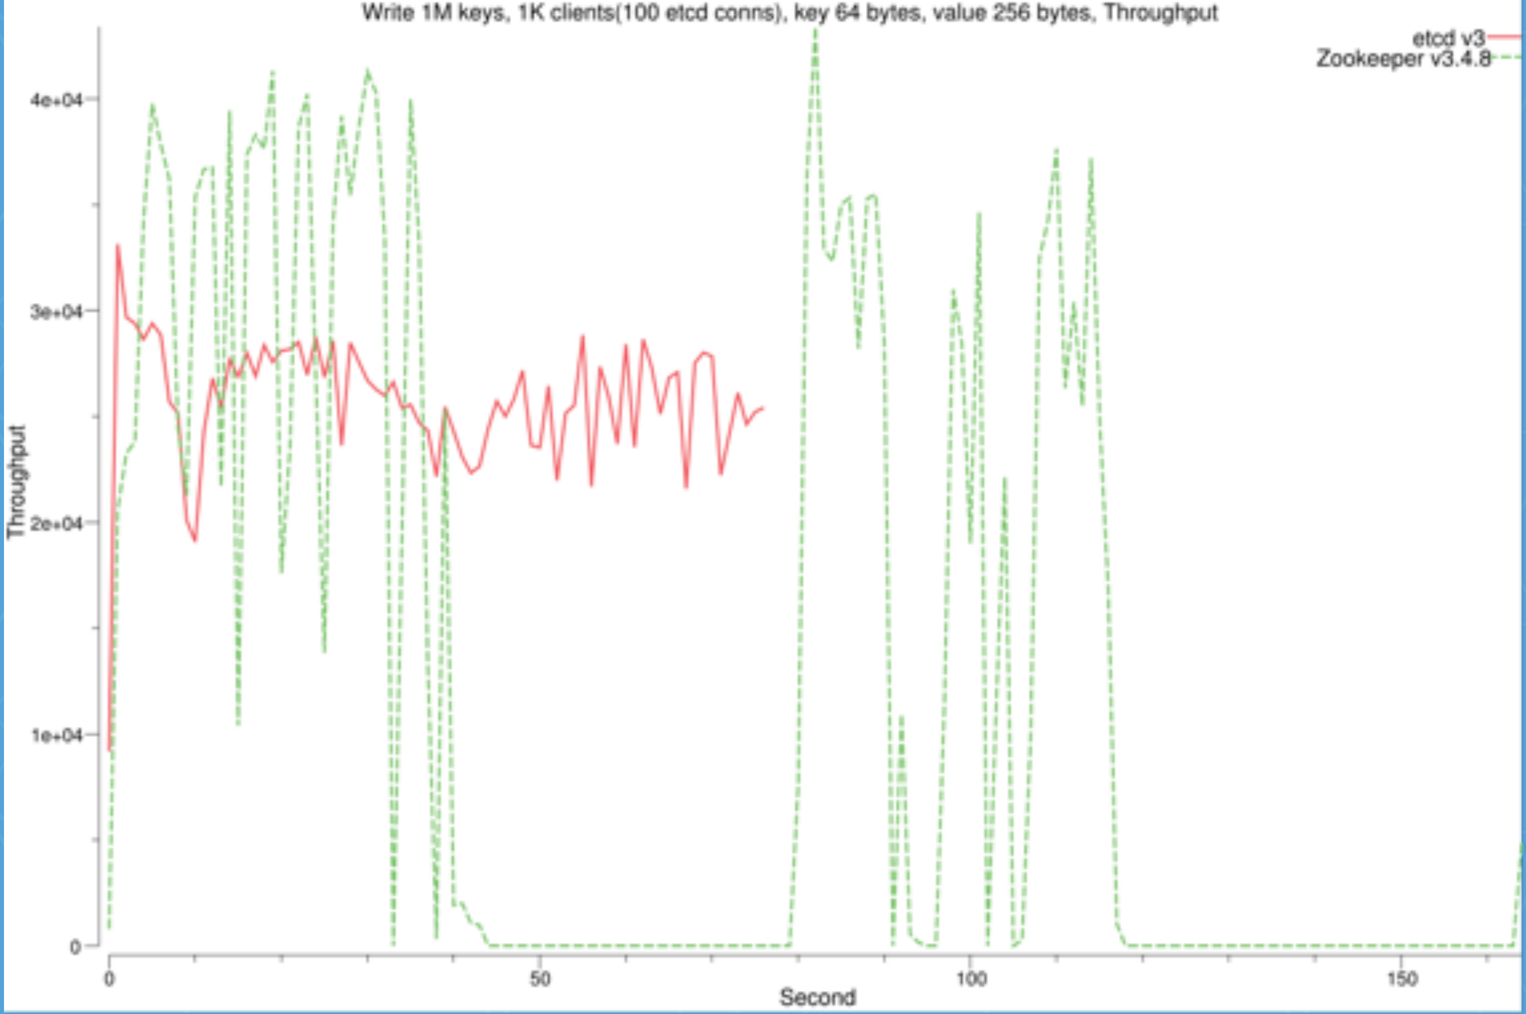

https://github.com/coreos/dbtester

> “I think it would be fairly easy to make the generated programs 20% smaller and 10% faster.”
> — khr

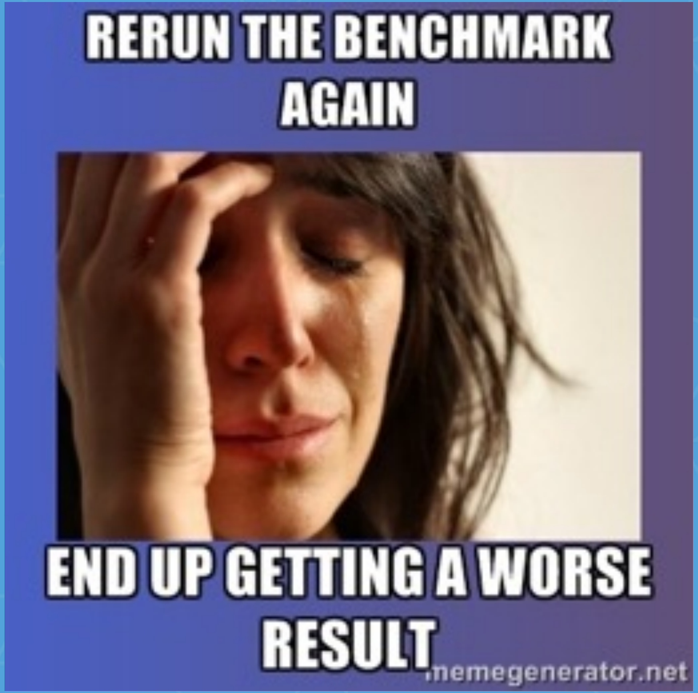

普

二

文

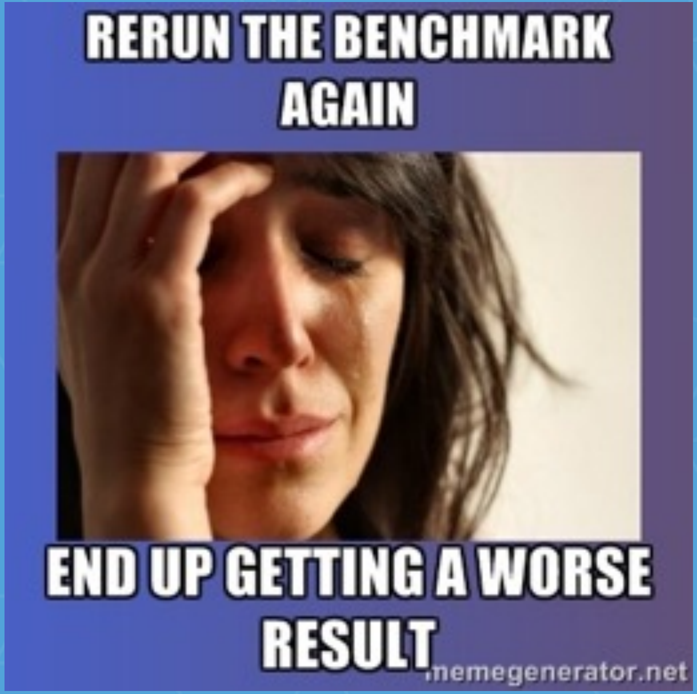

> “I think it would be fairly easy to make the generated programs 20% smaller and 10% faster.”
> — khr

Distributed Systems

End-to-end Testing

The only guarantee of cloud: Your instance will be gone

End-to-end Testing

Density Test

- How fast can your cluster run?

Failure Test

- How resilient is your cluster?

Conformance Test

- Are all your clusters configured the same and correctly?

Simulated Testing

Simulated Testing

Agility and Efficiency

Reproducibility

Get Rid of Noise

Focus on Key Component

Understand Your Use-case

Low latency? Consistency?

Design For growth

How many users? How much QPS/SLO?

CoreOS

私人, Hongchao Deng, hongchao.deng@coreos.com

招聘, coreos.com/careers

商务合作, sales@coreos.com

CoreOS

etcd

kubernetes

rkt

TECTONIC