

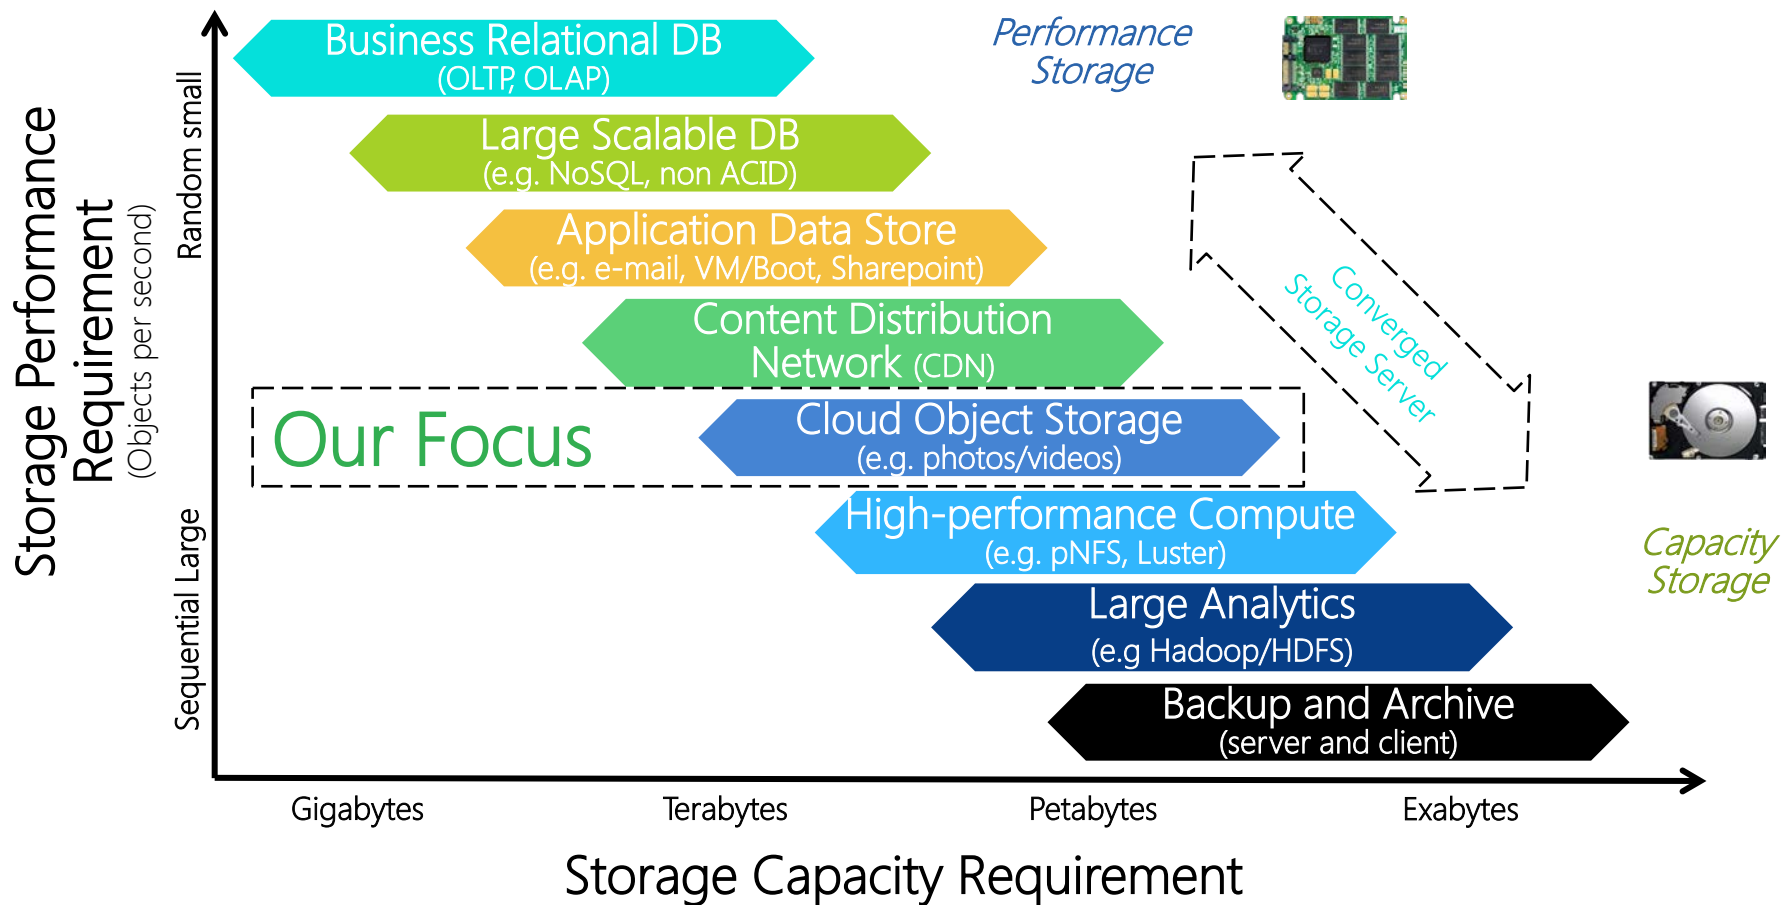
Know the **Performance** of your **Object Storage Services**

Qing Zheng

REINS Group
School of Software
Shanghai Jiao Tong University

Cloud Infra. Tech.
(DCST, STO, SSG)
Intel Asia-Pacific R&D

Usage Patterns



Different usage patterns drive different data solutions

Object Storage



No Namespaces

No Locks

Named Blobs

Data & Metadata

GET

PUT

POST

HEAD

DELETE

Implementations

Amazon

Rackspace

Google

Microsoft

Openstack

Eucalyptus

Ceph

Gluster FS

Hadoop FS

Problem

No benchmark tool available so far ...

Perl scripts, Specific apps

Object Storage Benchmark

Compare & evaluate different h/w & s/w stacks

Characterize systems and guide h/w & s/w configurations

Identify performance bottleneck & drive optimization

Find common issues, shared h/w expectation, future R paths

We present COSBench ...

COSBench

Cloud Object Storage Benchmark

```
<?xml version="1.0" encoding="UTF-8" ?>
<workload name="swift-sample"
description="sample benchmark for swift">

  <storage type="swift" />
  <auth type="swauth"
config="username=test:tester;password=testing;
url=http://192.168.10.1:8080/auth/v1.0" />
  <workflow>
    <workstage name="main">
      <work name="main" workers="8"
rampup="100" runtime="300">
        <operation type="read" ratio="80"
config="containers=u(1,32);objects=u(1,50)" />
        <operation type="write" ratio="20"
config="containers=u(1,32);objects=u(51,100);siz
es=c(64)KB" />
      </work>
    </workstage>
  </workflow>
</workload>
```

Firefox

COSBench Controller

127.0.0.1:18088/controller/index.html

GA Release version: 2.0.0.GA

COSBENCH - CONTROLLER WEB CONSOLE

Controller Overview

Name: not configured URL: not configured

Driver	Name	URL	Link
1	driver1	http://127.0.0.1:18088/driver	view details
2	driver2	http://127.0.0.1:18088/driver	view details

There are 2 drivers attached to the controller.

Active Workloads

Id	Name	Submitted-At	State	Link
w6	demo	Aug 3, 2012 2:56:48 PM	processing	view details
w7	demo	Aug 3, 2012 2:56:52 PM	queuing	view details

There are currently 2 active workloads.

[submit new workloads](#)

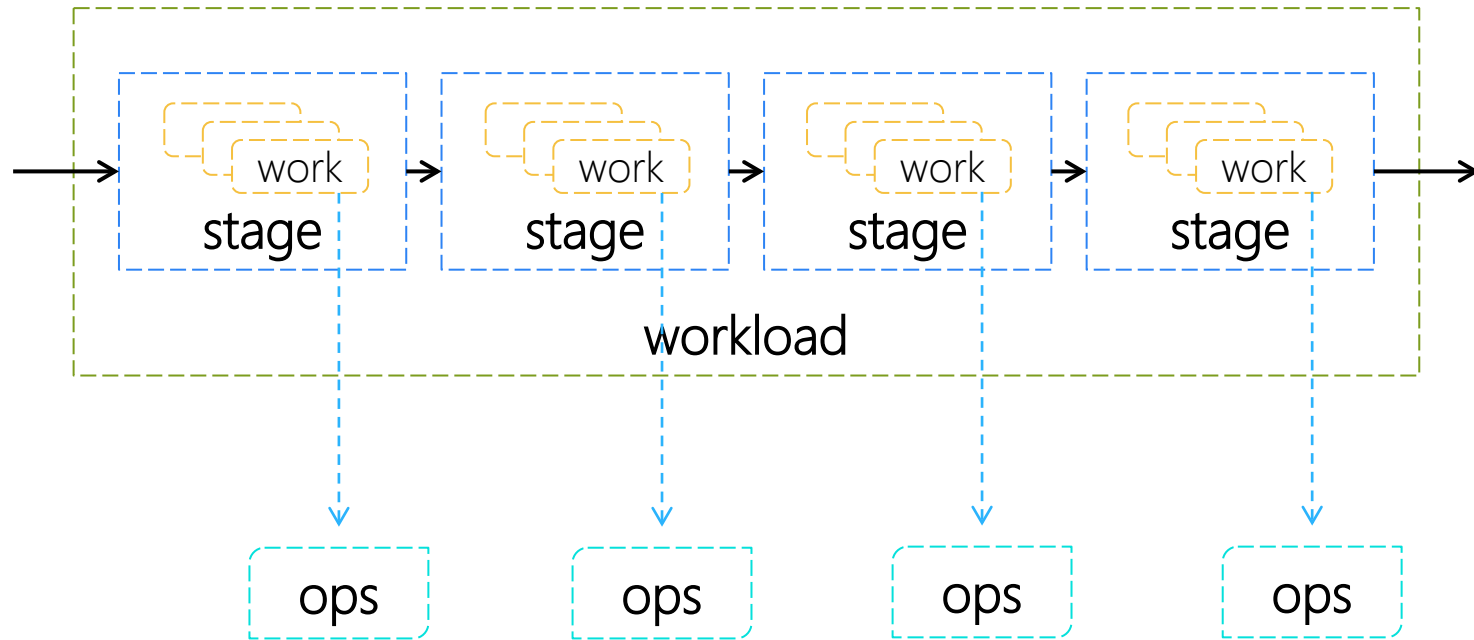
History Workloads

[view performance matrix](#)

Id	Name	Duration	Op-Info	State	Link
w4	demo	Aug 3, 2012 2:52:51 PM - 2:53:37 PM	prepare, read	finished	view details
w5	demo	Aug 3, 2012 2:53:37 PM - 2:54:23 PM	prepare, read	finished	view details

Software Solution Group - System Software Division - System Optimization Technology Center

Workload Model



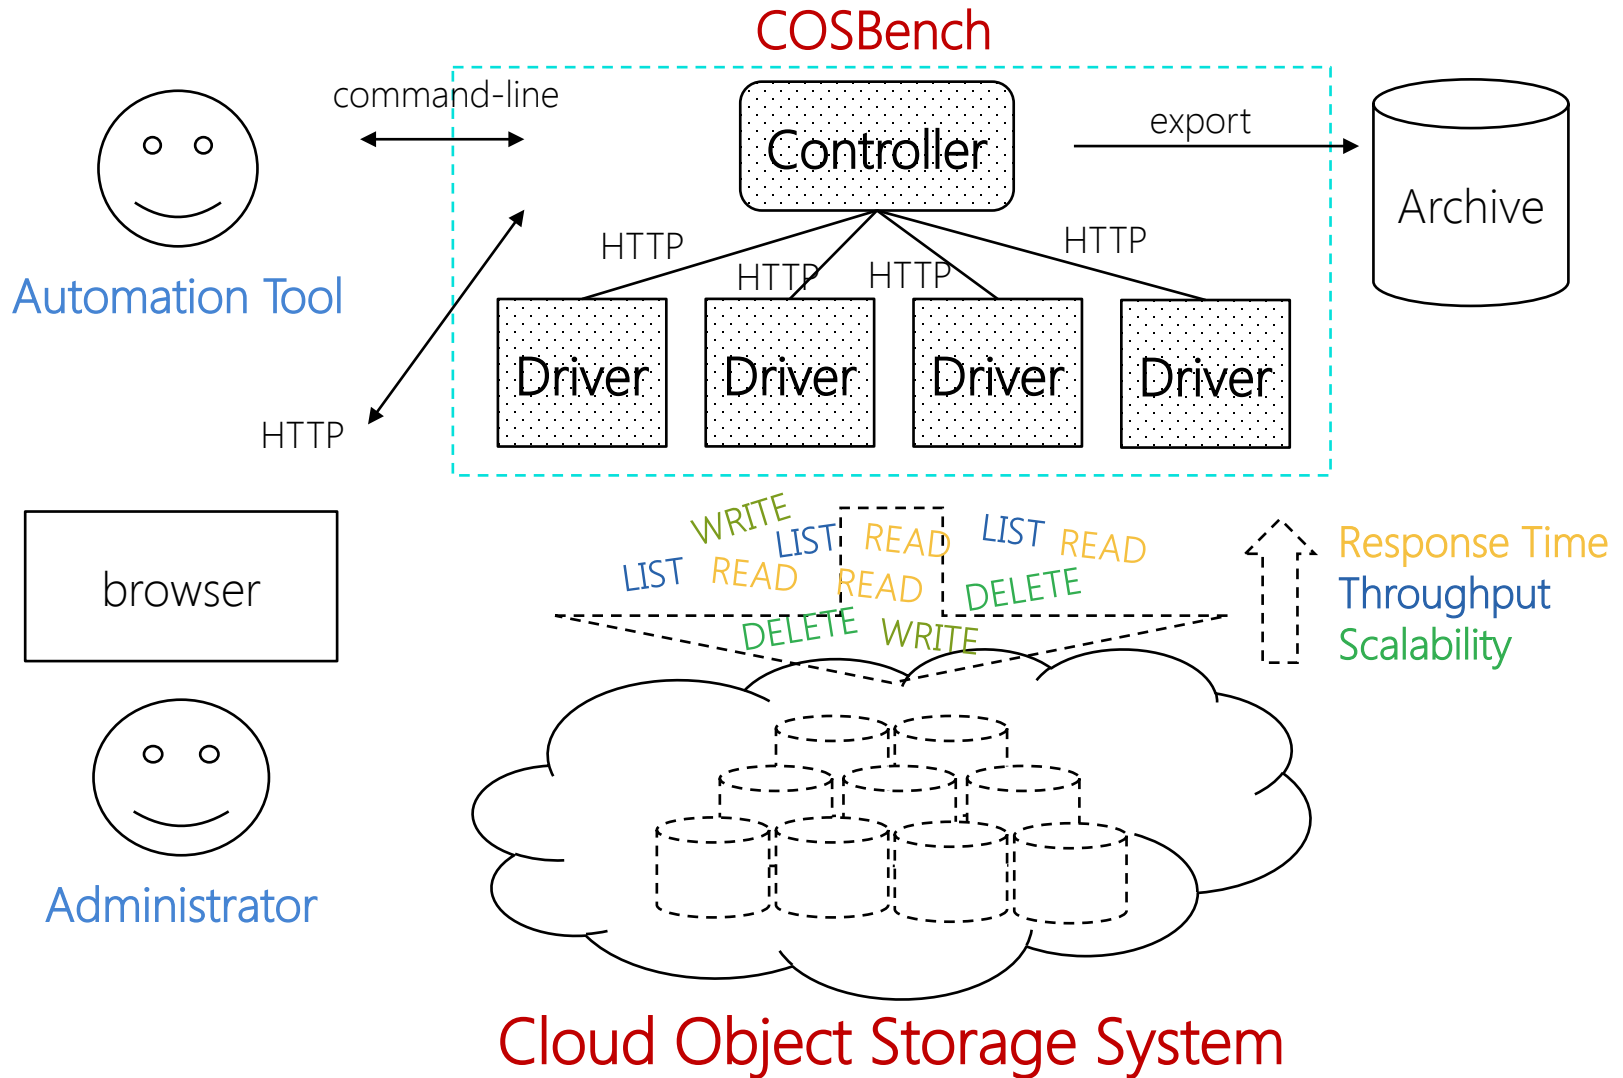
Operation

Ratio, Object/Container Path, Object Size, Other Params

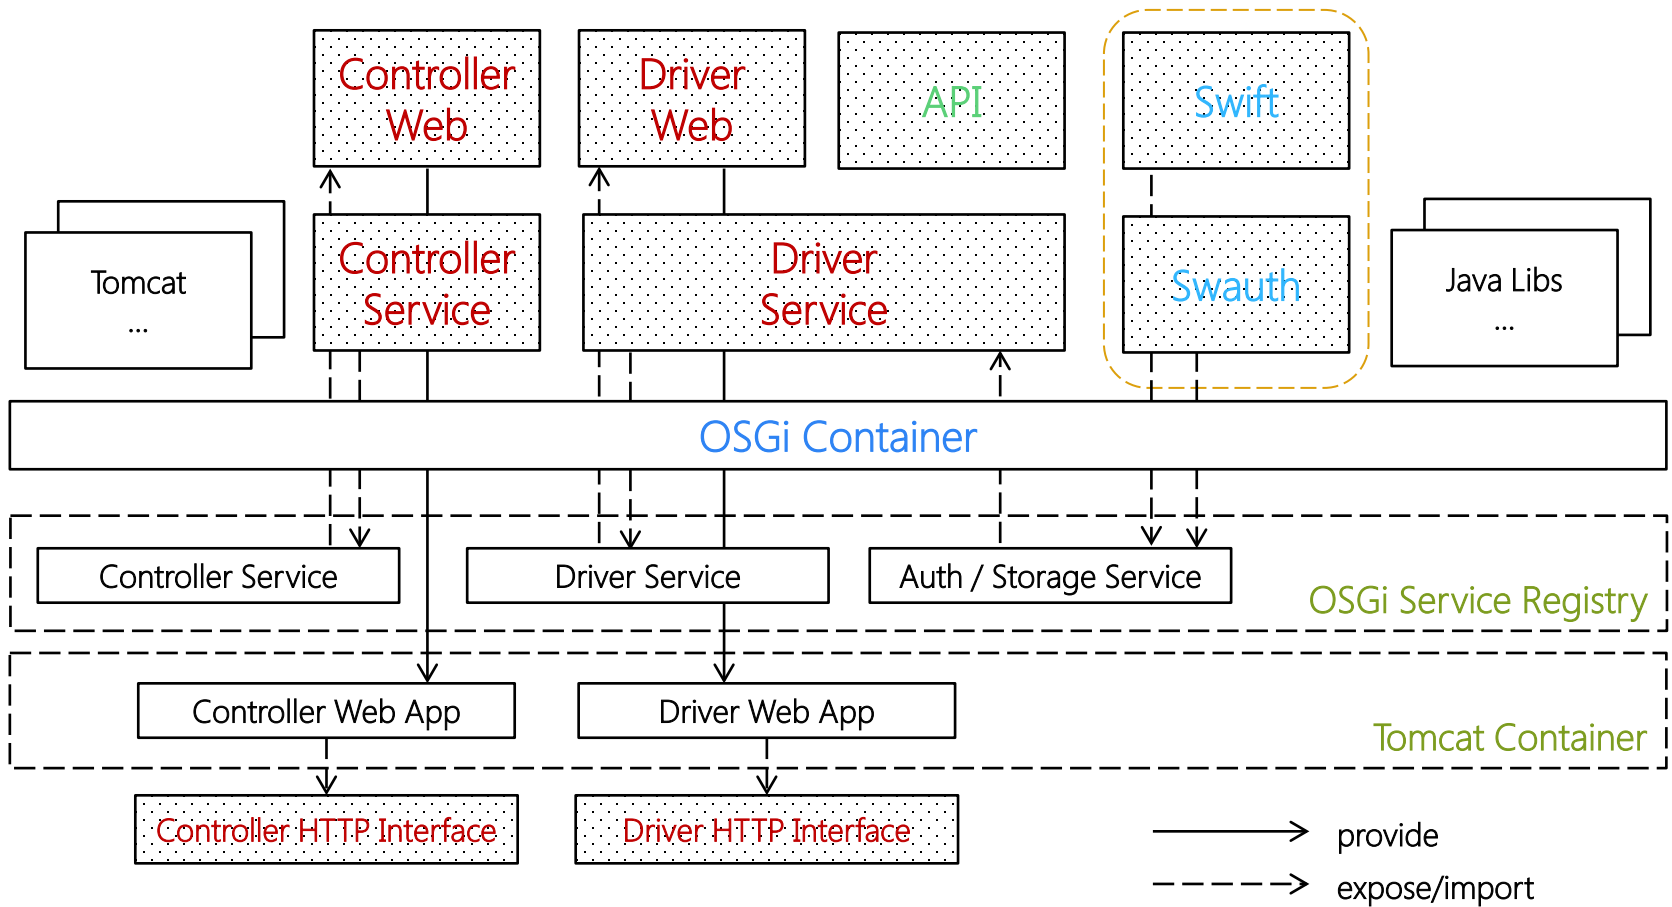
80% read, 800 containers, 10000 objs

20% write, 800 containers, 10000 objs, 32 ~ 64KB in size

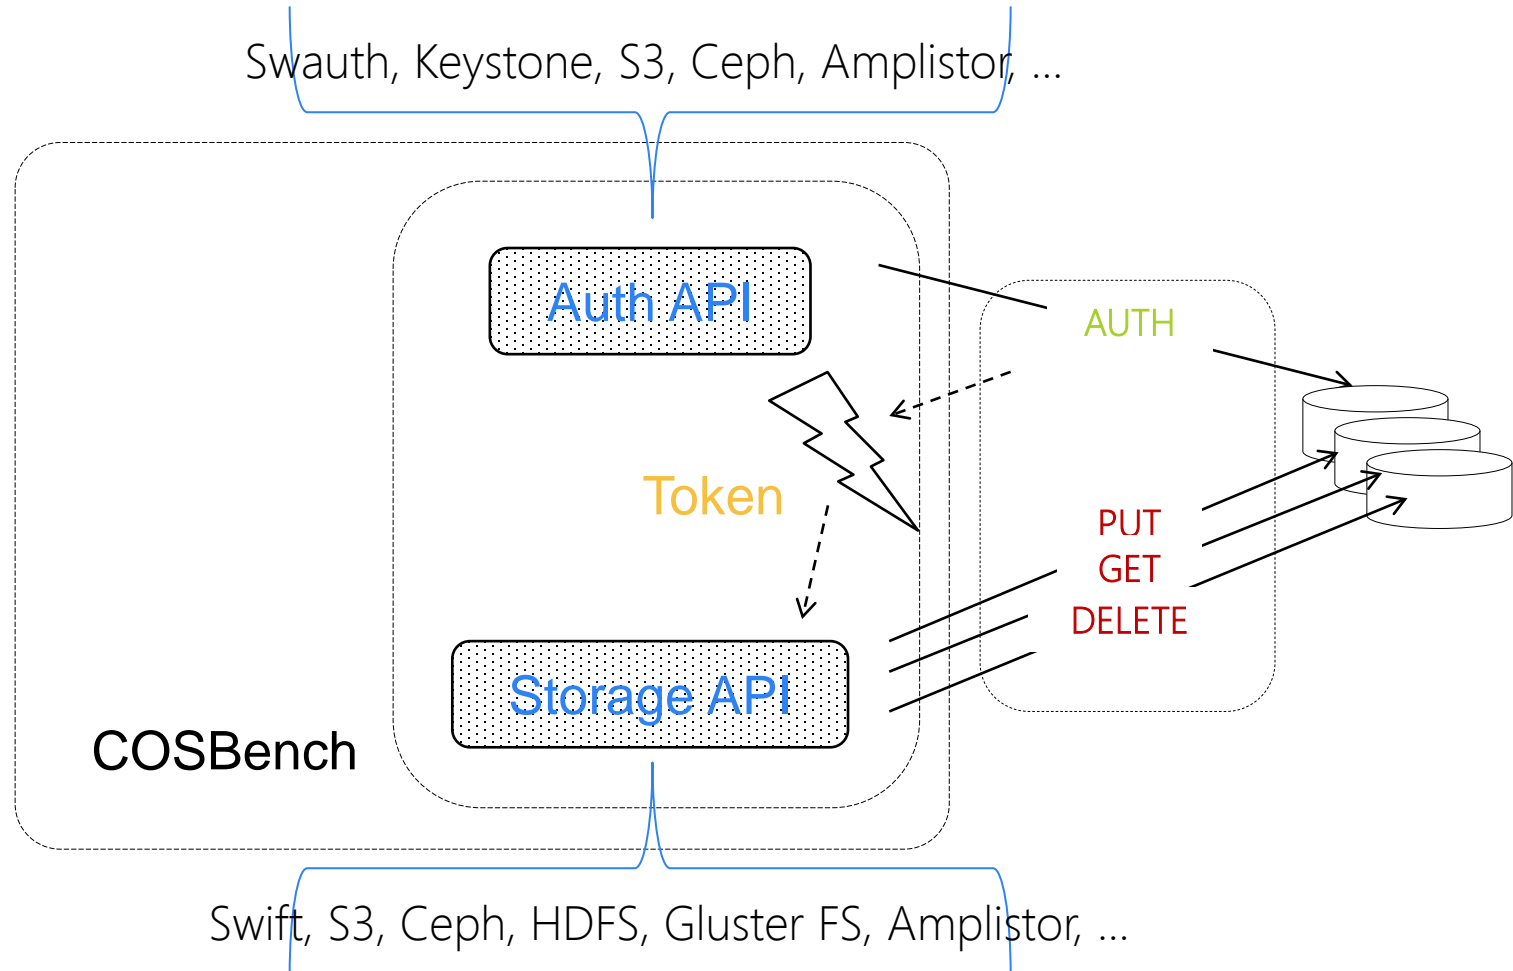
Overview



Modular Architecture

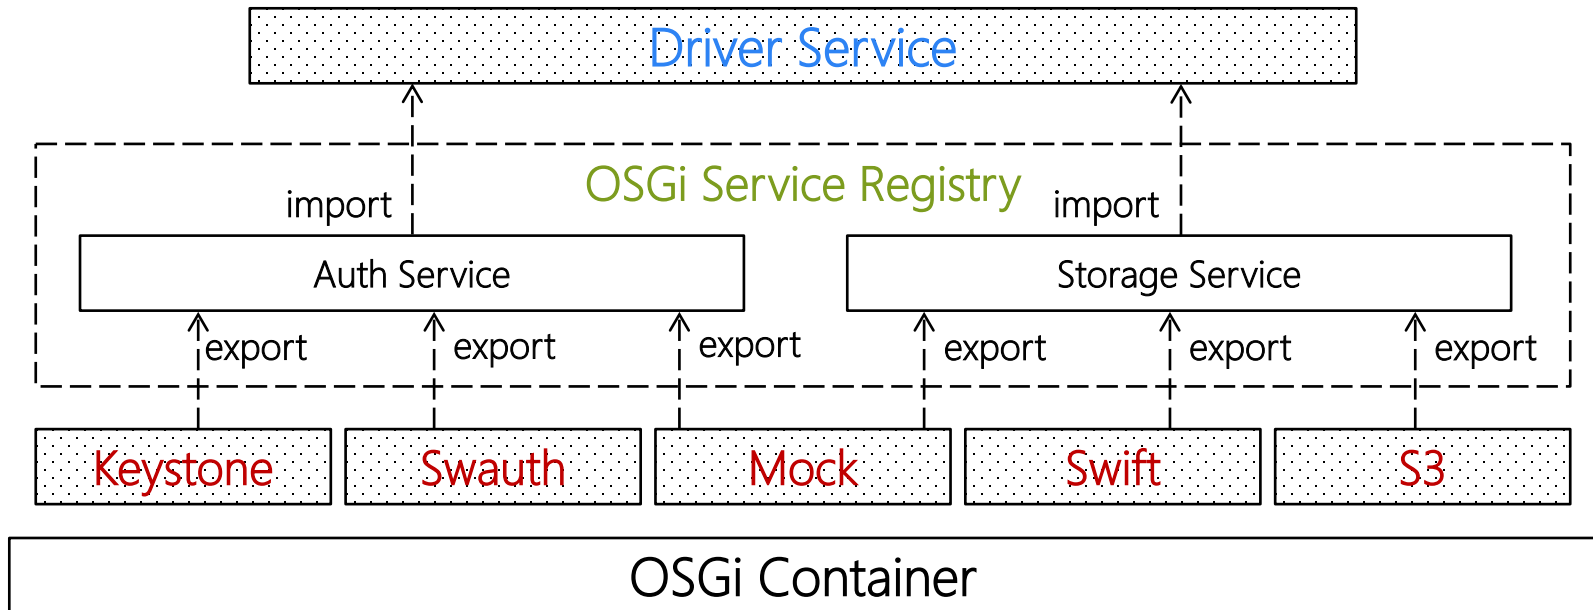


Extensible API



Workload Configuration

```
<operation type="writ", ratio="30",  
config="containers=u(1,100); objects=u(1,100);  
sizes=u(64,128)KB; object-expiration=10" />
```



Functionality

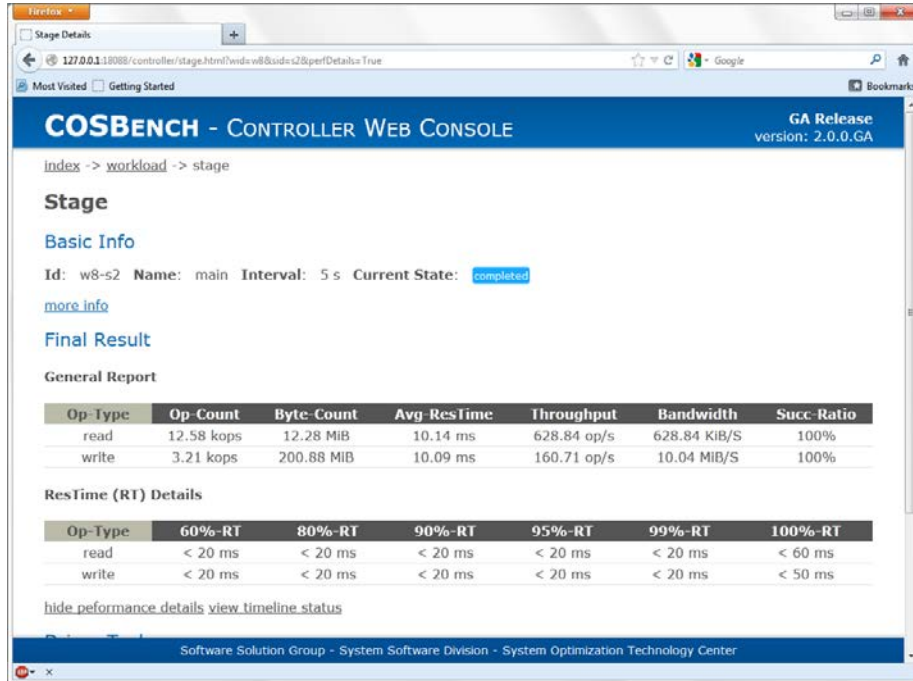
The screenshot shows the 'Workload Submission' page of the COSBENCH Controller Web Console. The page title is 'COSBENCH - CONTROLLER WEB CONSOLE' with a 'GA Release version: 2.0.0.GA' badge. The main content area is titled 'Submission Results' and displays a green 'Success' message: 'your workload has been accepted!'. Below this, there is a 'Workload Submission' section with a text input field for 'workload config:', a 'Browse...' button, and a 'submit' button. A table titled 'Active Workloads' lists two workloads: w4 and w5, both named 'demo', submitted on Aug 3, 2012 at 2:52:51 PM and 2:52:54 PM respectively. The states are 'processing' and 'queued'. A footer bar contains the text 'Software Solution Group - System Software Division - System Optimization Technology Center'.

Online Workload Submission

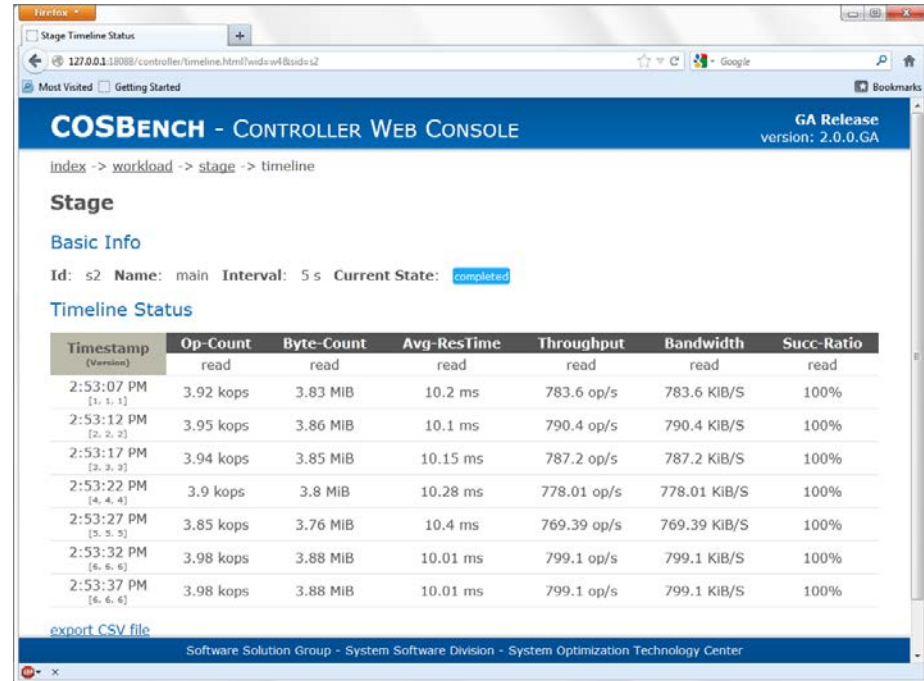
The screenshot shows the 'Workload Details' page of the COSBENCH Controller Web Console. The page title is 'COSBENCH - CONTROLLER WEB CONSOLE' with a 'GA Release version: 2.0.0.GA' badge. The main content area is titled 'Workload' and displays 'Basic Info' for workload 'w8' named 'demo'. The 'Current State' is 'processing' and the 'Current Stage' is 'main'. It also shows 'Submitted At: Aug 3, 2012 3:09:33 PM', 'Started At: Aug 3, 2012 3:09:33 PM', and 'Stopped At: N/A'. Below this is a 'Snapshot' section with a 'General Report' table. The table has columns for 'Op-Type', 'Op-Count', 'Byte-Count', 'Avg-ResTime', 'Throughput', 'Bandwidth', and 'Succ-Ratio'. The data rows are: 'read' (3.15 kops, 3.07 MIB, 10.07 ms, 628.4 op/s, 628.4 KIB/S, 100%) and 'write' (826 ops, 51.62 MIB, 10.09 ms, 164.99 op/s, 10.31 MIB/S, 100%). A note states 'The snapshot was taken at 3:09:51 PM with version 1.'. Below the report is a 'Stages' section with a box containing 'prepare --> main'. A footer bar contains the text 'Software Solution Group - System Software Division - System Optimization Technology Center'.

Real-time Status Monitoring

Functionality

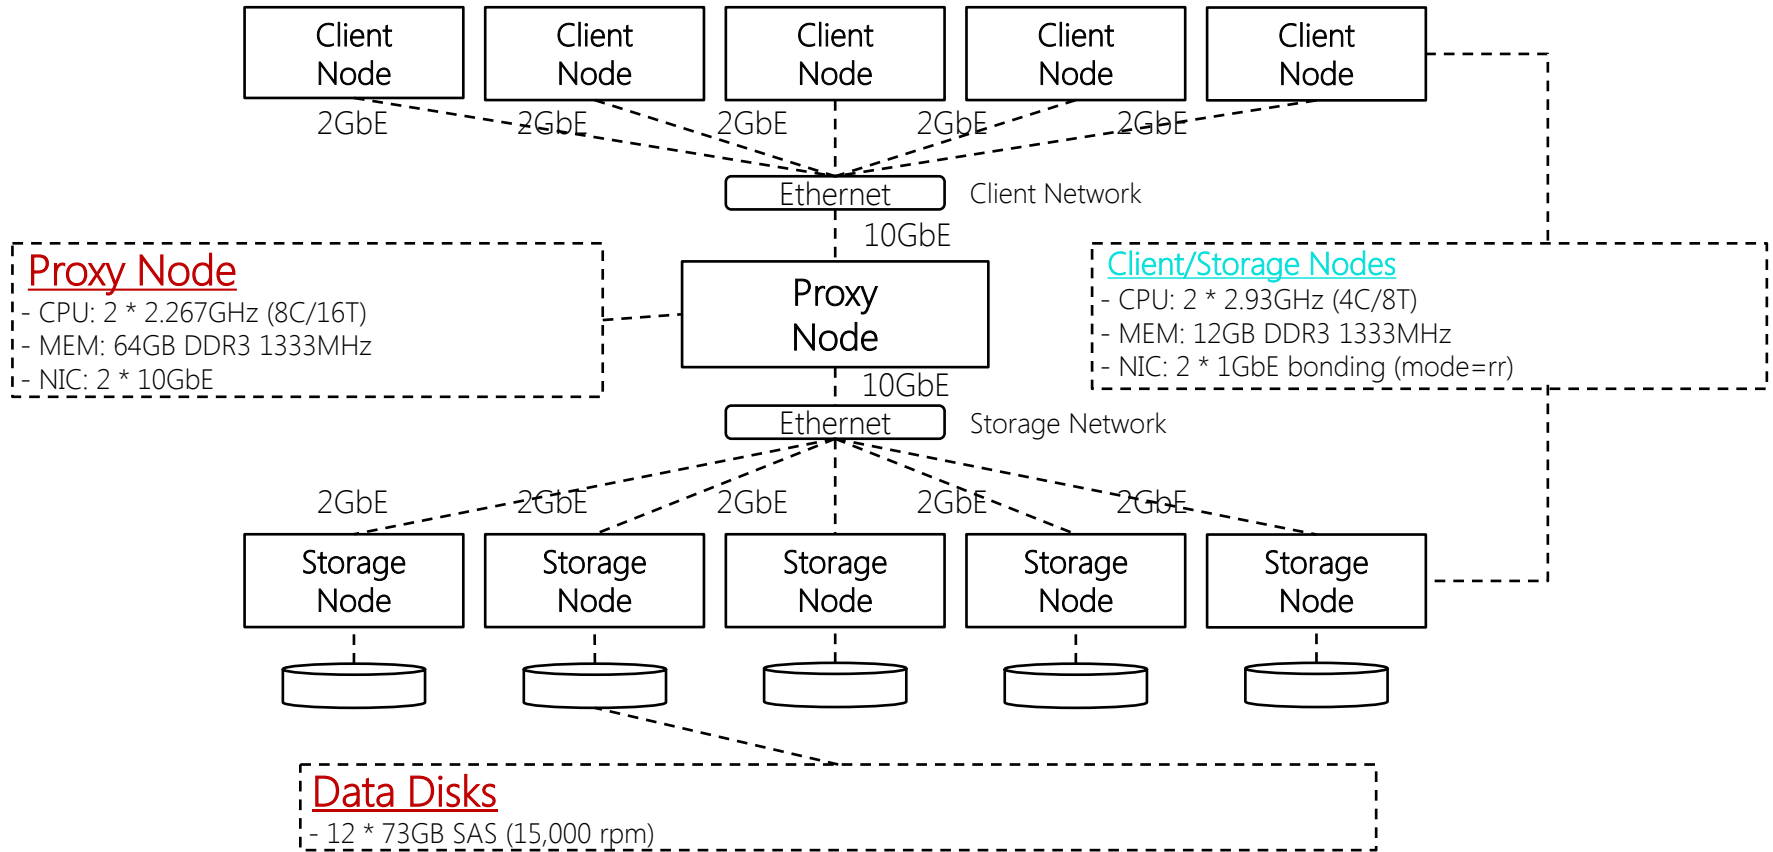


Workload Report

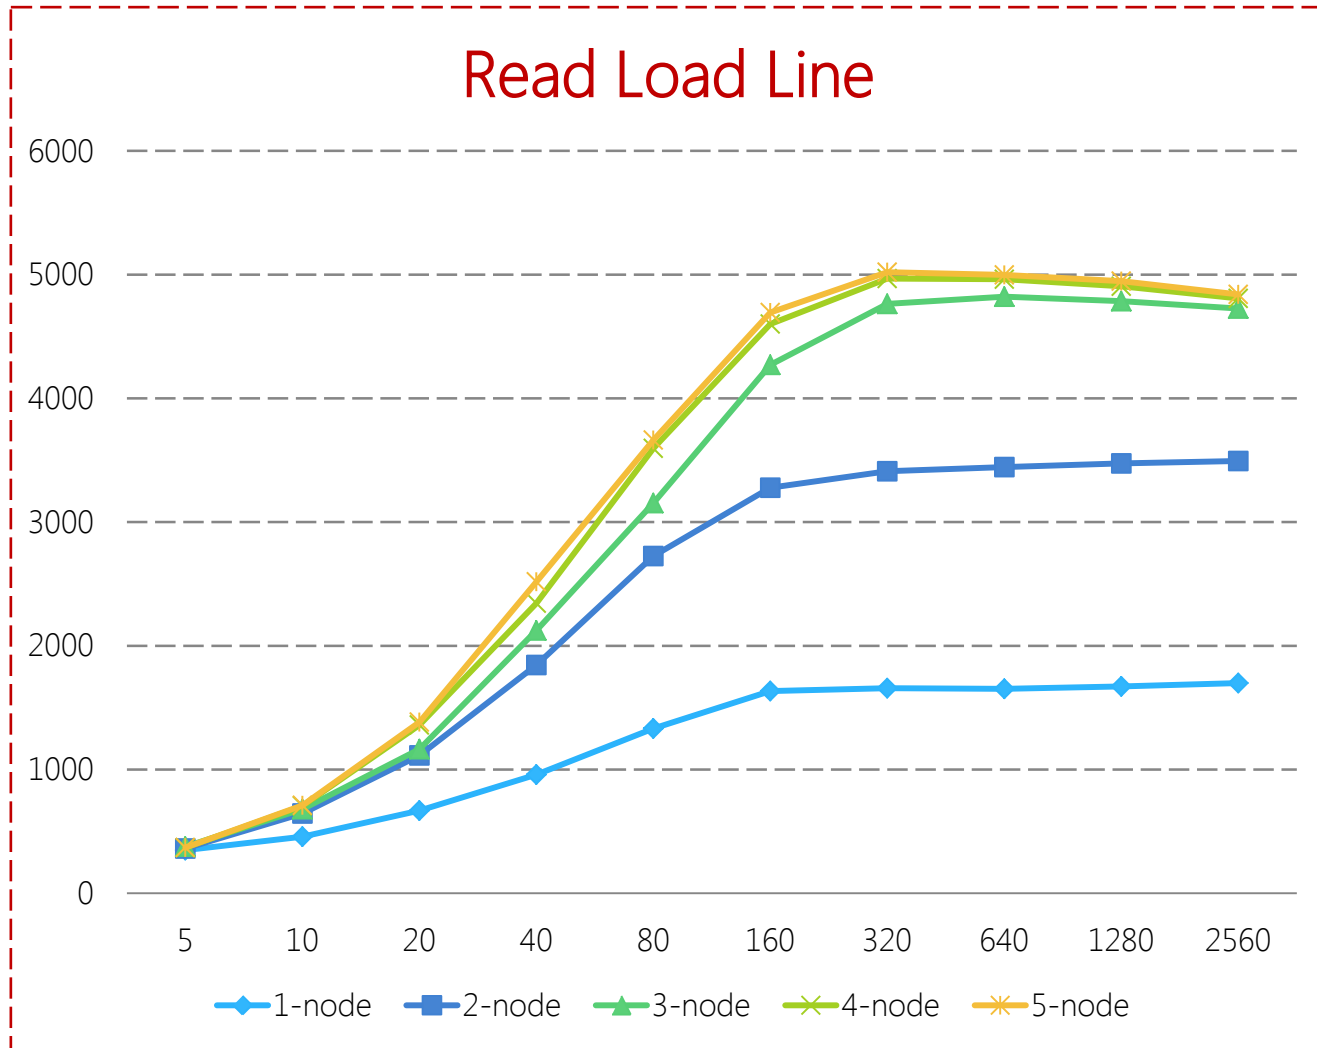


Timeline Plotting


Swift



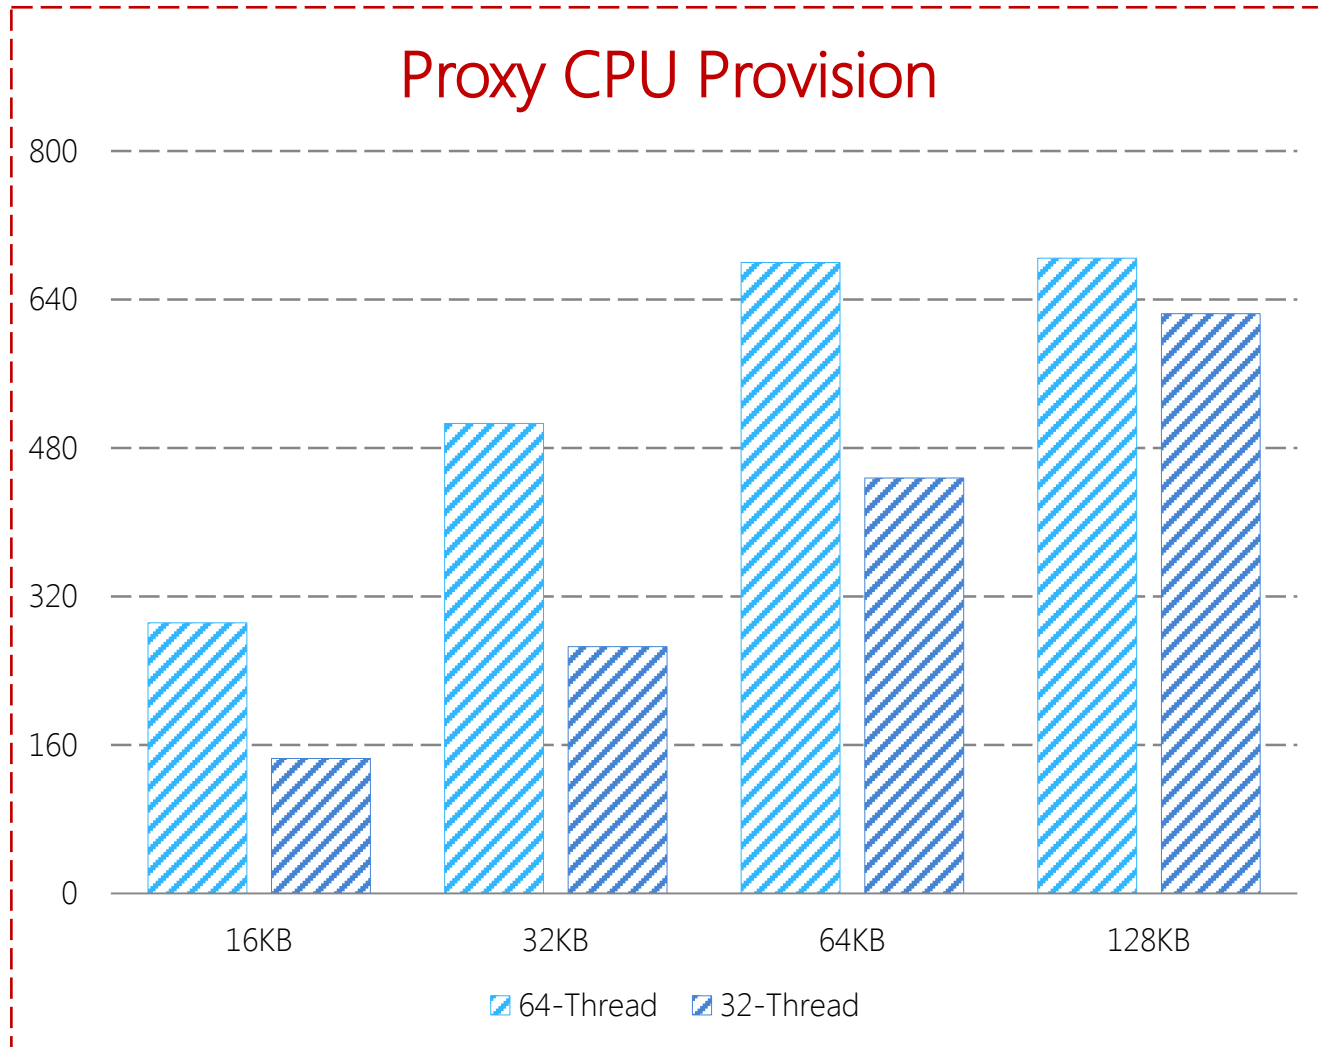
128KB Reads



Small Reads



Small Reads



Thank you!

Open Source (*Apache 2.0 License*)

<https://github.com/intel-cloud/cosbench>

Shared among 20+ groups of people inside/outside Intel
thanks to them, cosbench is getting better & better