

# Cloud WorkBench

## Infrastructure-as-Code Based Cloud Benchmarking

Joel Scheuner, Philipp Leitner, Jürgen Cito, Harald Gall

University of Zurich, Switzerland  
 {lastname}@ifi.uzh.ch



**University of  
Zurich<sup>UZH</sup>**



12/16/2014



# Motivation (1)

Infrastructure-as-a-Service (IaaS)




# Motivation (1)

## Infrastructure-as-a-Service (IaaS)



Google Compute Engine




# Motivation (1)

## Infrastructure-as-a-Service (IaaS)

>22 Instance Types




Google Compute Engine



# Motivation (2)

## Performance Variations




**Figure 6: Disk Write Bandwidth (KB/Sec)**

[1] K. Salah, M. Al-Saba, M. Akhdhor, O. Shaaban, and M. Buhari. Performance Evaluation of Popular Cloud IaaS Providers. In 6th Int. Conf. on Internet Technology and Secured Transactions (ICITST), pp. 345–349, Dec. 2011.

# Motivation (2)

## Performance Variations



**Figure 6: Disk Write Bandwidth (KB/Sec)**

[1] K. Salah, M. Al-Saba, M. Akhdhor, O. Shaaban, and M. Buhari. Performance Evaluation of Popular Cloud IaaS Providers. In 6th Int. Conf. on Internet Technology and Secured Transactions (ICITST), pp. 345–349, Dec. 2011.

Even for services with the same specification!

# Motivation (3)

- Benchmark  
(i.e. performance test)
- Problems
  - Time-consuming
  - Error-prone



# Related Work

## Cloud Experiment Automation

- CloudBench** (1) M. Silva, M. Hines, D. Gallo, Q. Liu, K. D. Ryu, and D. Da Silva. **CloudBench: Experiment Automation for Cloud Environments**. In IEEE Int. Conf. on Cloud Engineering (IC2E), pp. 302–311, March 2013.
- CloudCrawler** (2) M. Cunha, N. Mendonça, and A. Sampaio. **A Declarative Environment for Automatic Performance Evaluation in IaaS Clouds**. In 6th IEEE Int. Conf. on Cloud Computing (CLOUD), pp. 285–292, June 2013.
- Expertus** (3) D. Jayasinghe, J. Kimball, S. Choudhary, T. Zhu, and C. Pu. **An Automated Approach to Create, Store, and Analyze Large-scale Experimental Data in Clouds**. In 14th IEEE Int. Conf. on Information Reuse and Integration (IRI), pp. 357–364, August 2013.



# Related Work

## Cloud Experiment Automation

- CloudBench** (1) M. Silva, M. Hines, D. Gallo, Q. Liu, K. D. Ryu, and D. Da Silva. **CloudBench: Experiment Automation for Cloud Environments**. In IEEE Int. Conf. on Cloud Engineering (IC2E), pp. 302–311, March 2013.
- CloudCrawler** (2) M. Cunha, N. Mendonça, and A. Sampaio. **A Declarative Environment for Automatic Performance Evaluation in IaaS Clouds**. In 6th IEEE Int. Conf. on Cloud Computing (CLOUD), pp. 285–292, June 2013.
- Expertus** (3) D. Jayasinghe, J. Kimball, S. Choudhary, T. Zhu, and C. Pu. **An Automated Approach to Create, Store, and Analyze Large-scale Experimental Data in Clouds**. In 14th IEEE Int. Conf. on Information Reuse and Integration (IRI), pp. 357–364, August 2013.

Our approach: **Provisioning via Infrastructure-as-Code**

# Research Questions

I *How can existing IaaS cloud benchmarks be described in a modular and portable manner?*

# Research Questions

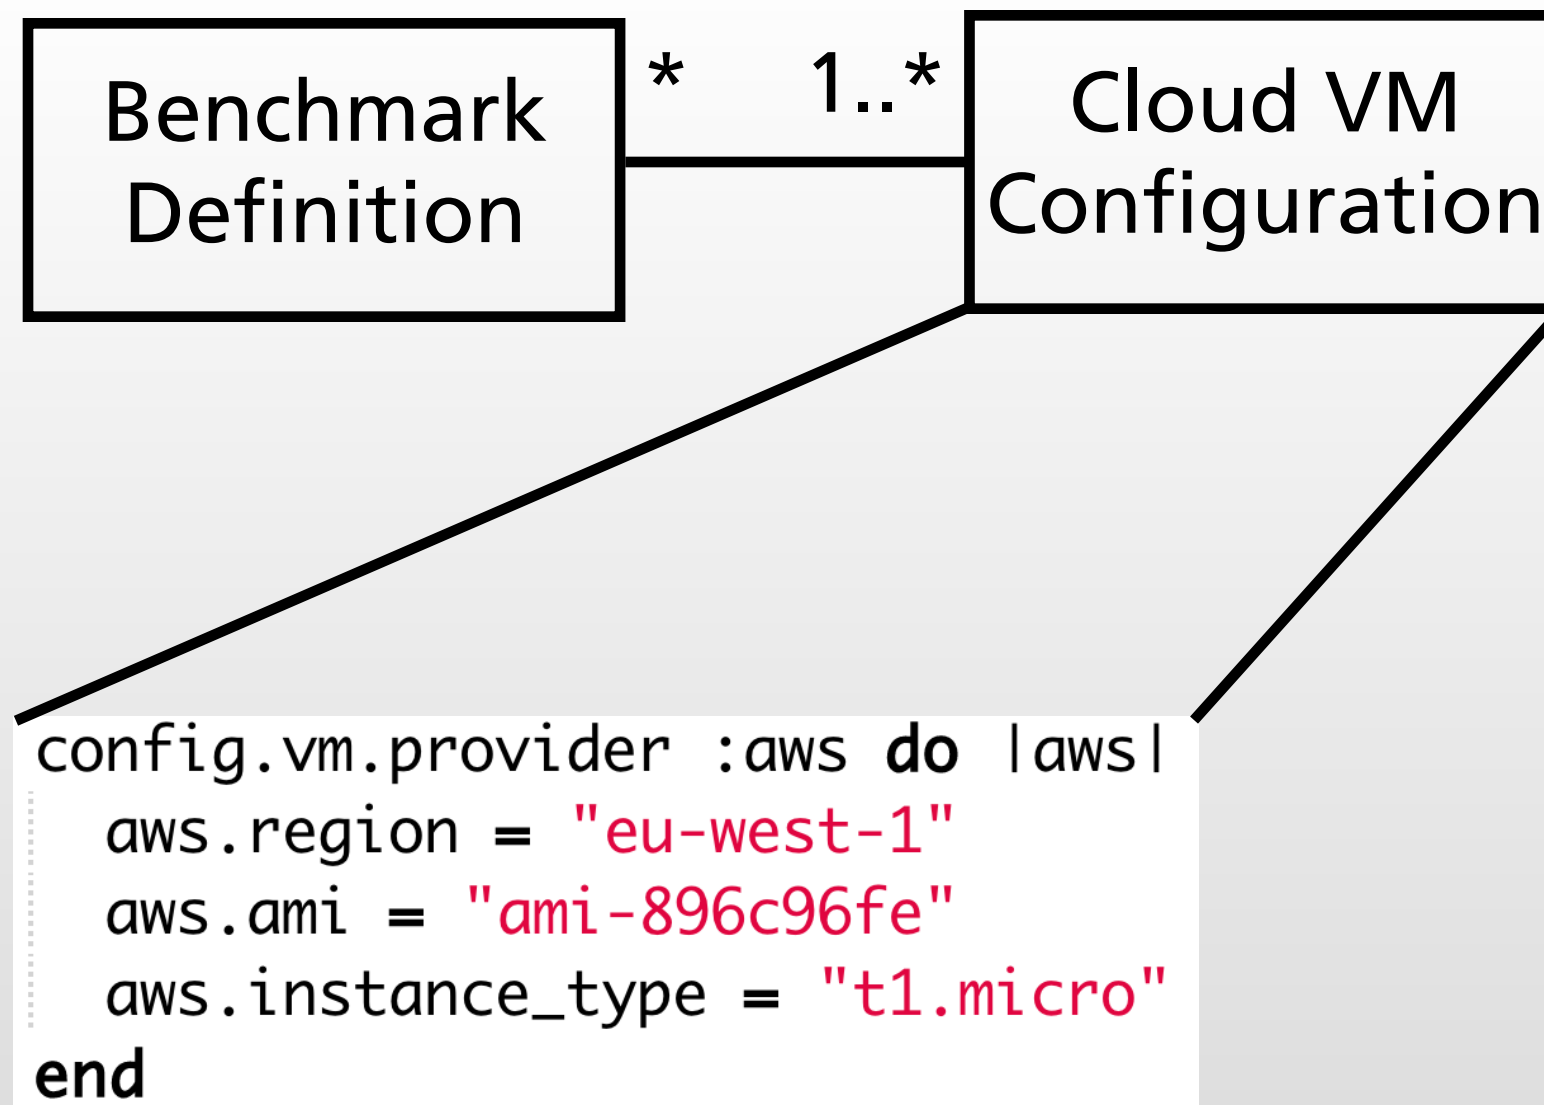
I *How can existing IaaS cloud benchmarks be described in a modular and portable manner?*

II *How can such benchmarks be periodically scheduled and executed in cloud environments in a fully reproducible way, and without manual interaction?*

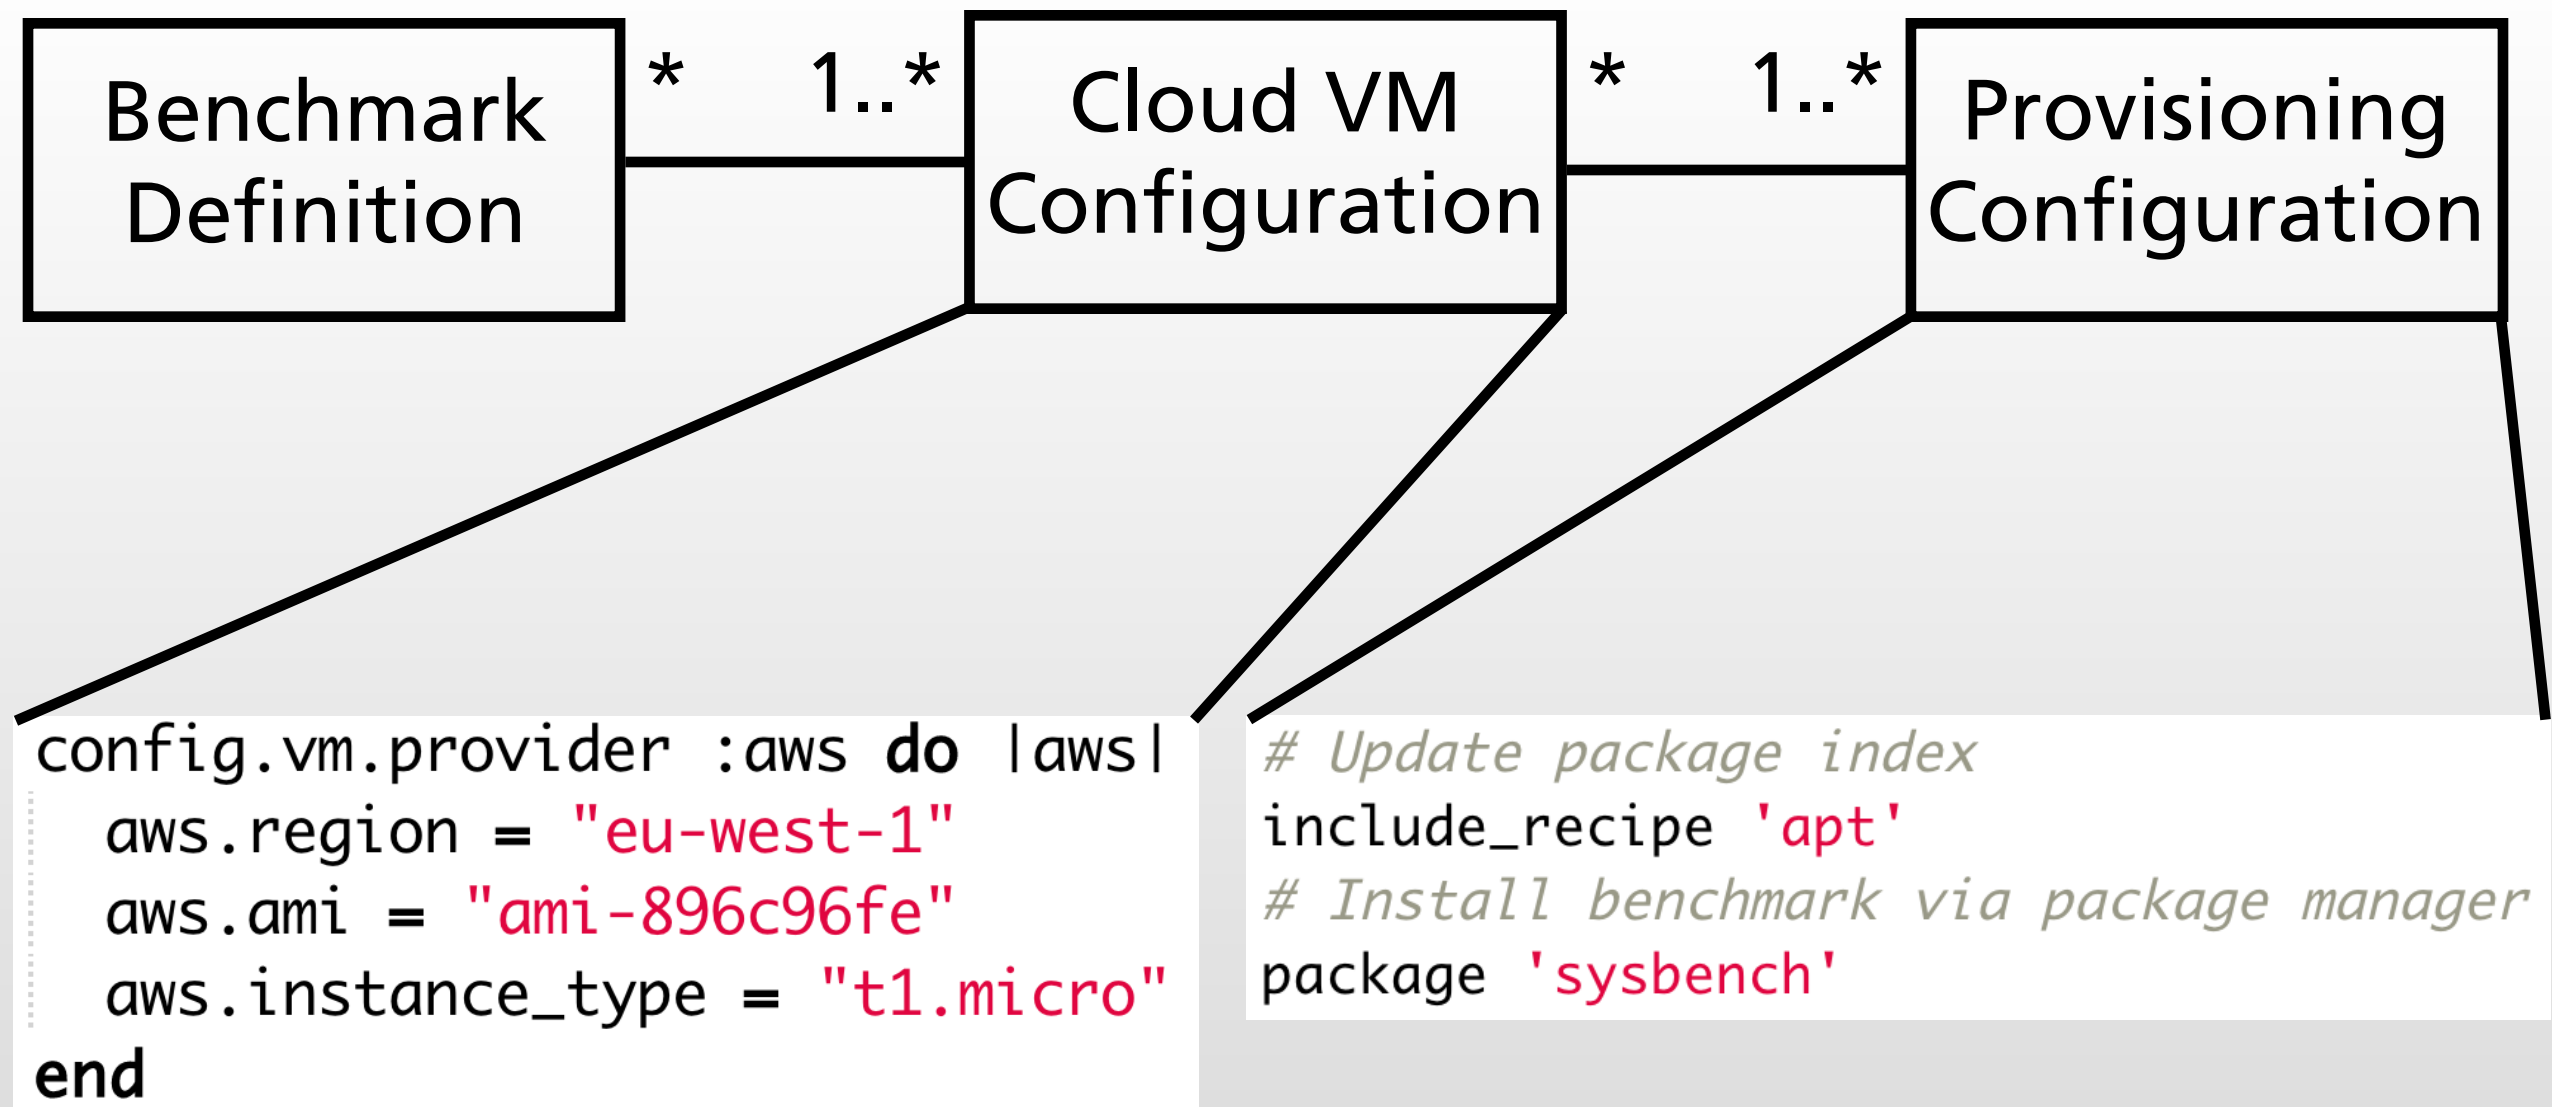
# Benchmark Definition

Benchmark  
Definition

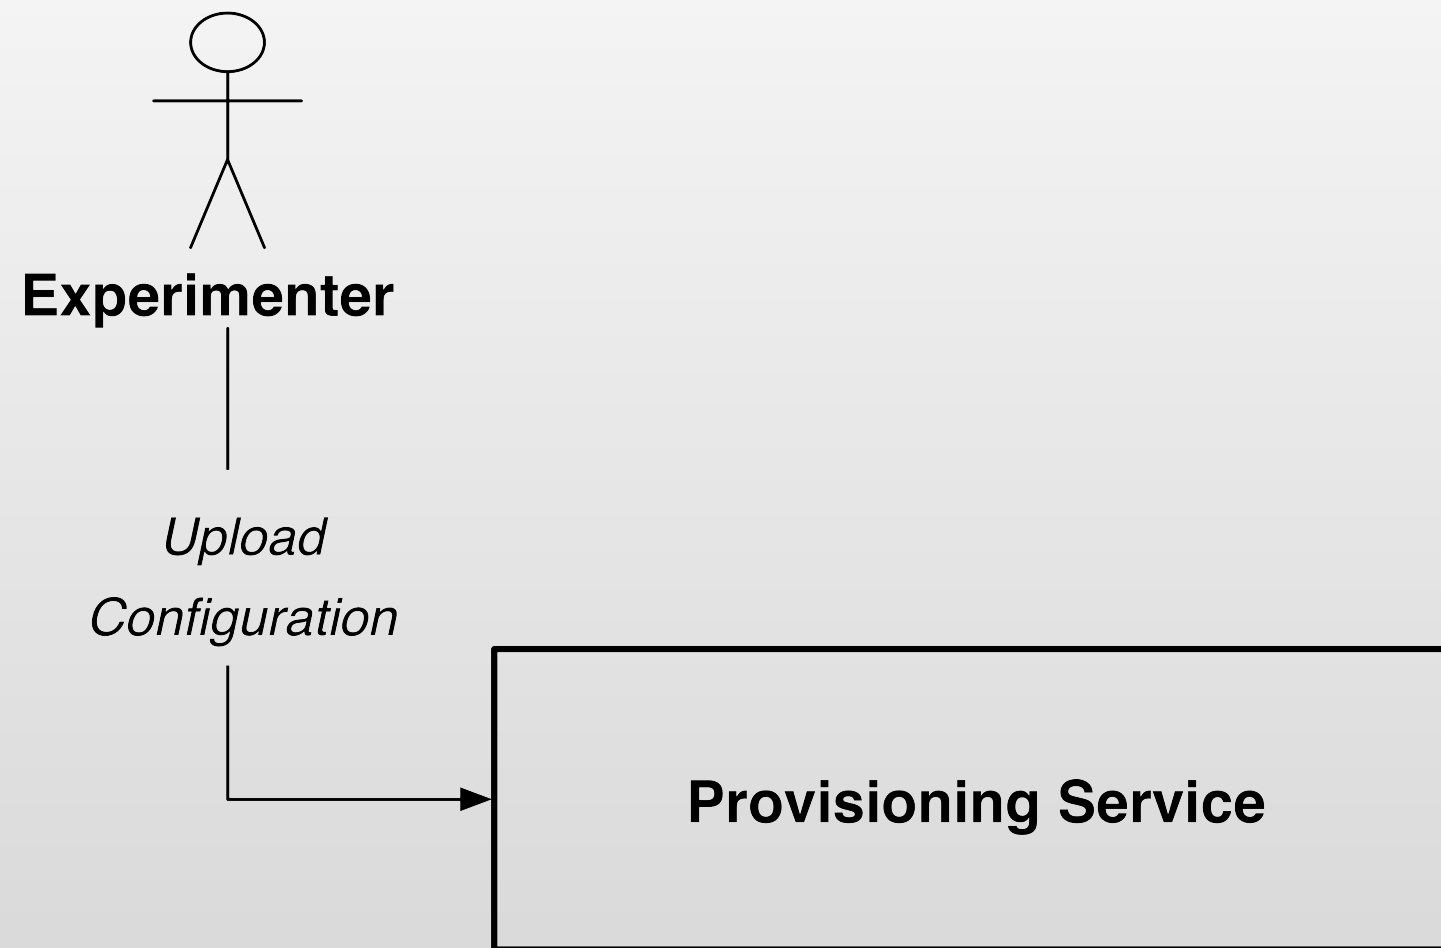
# Benchmark Definition



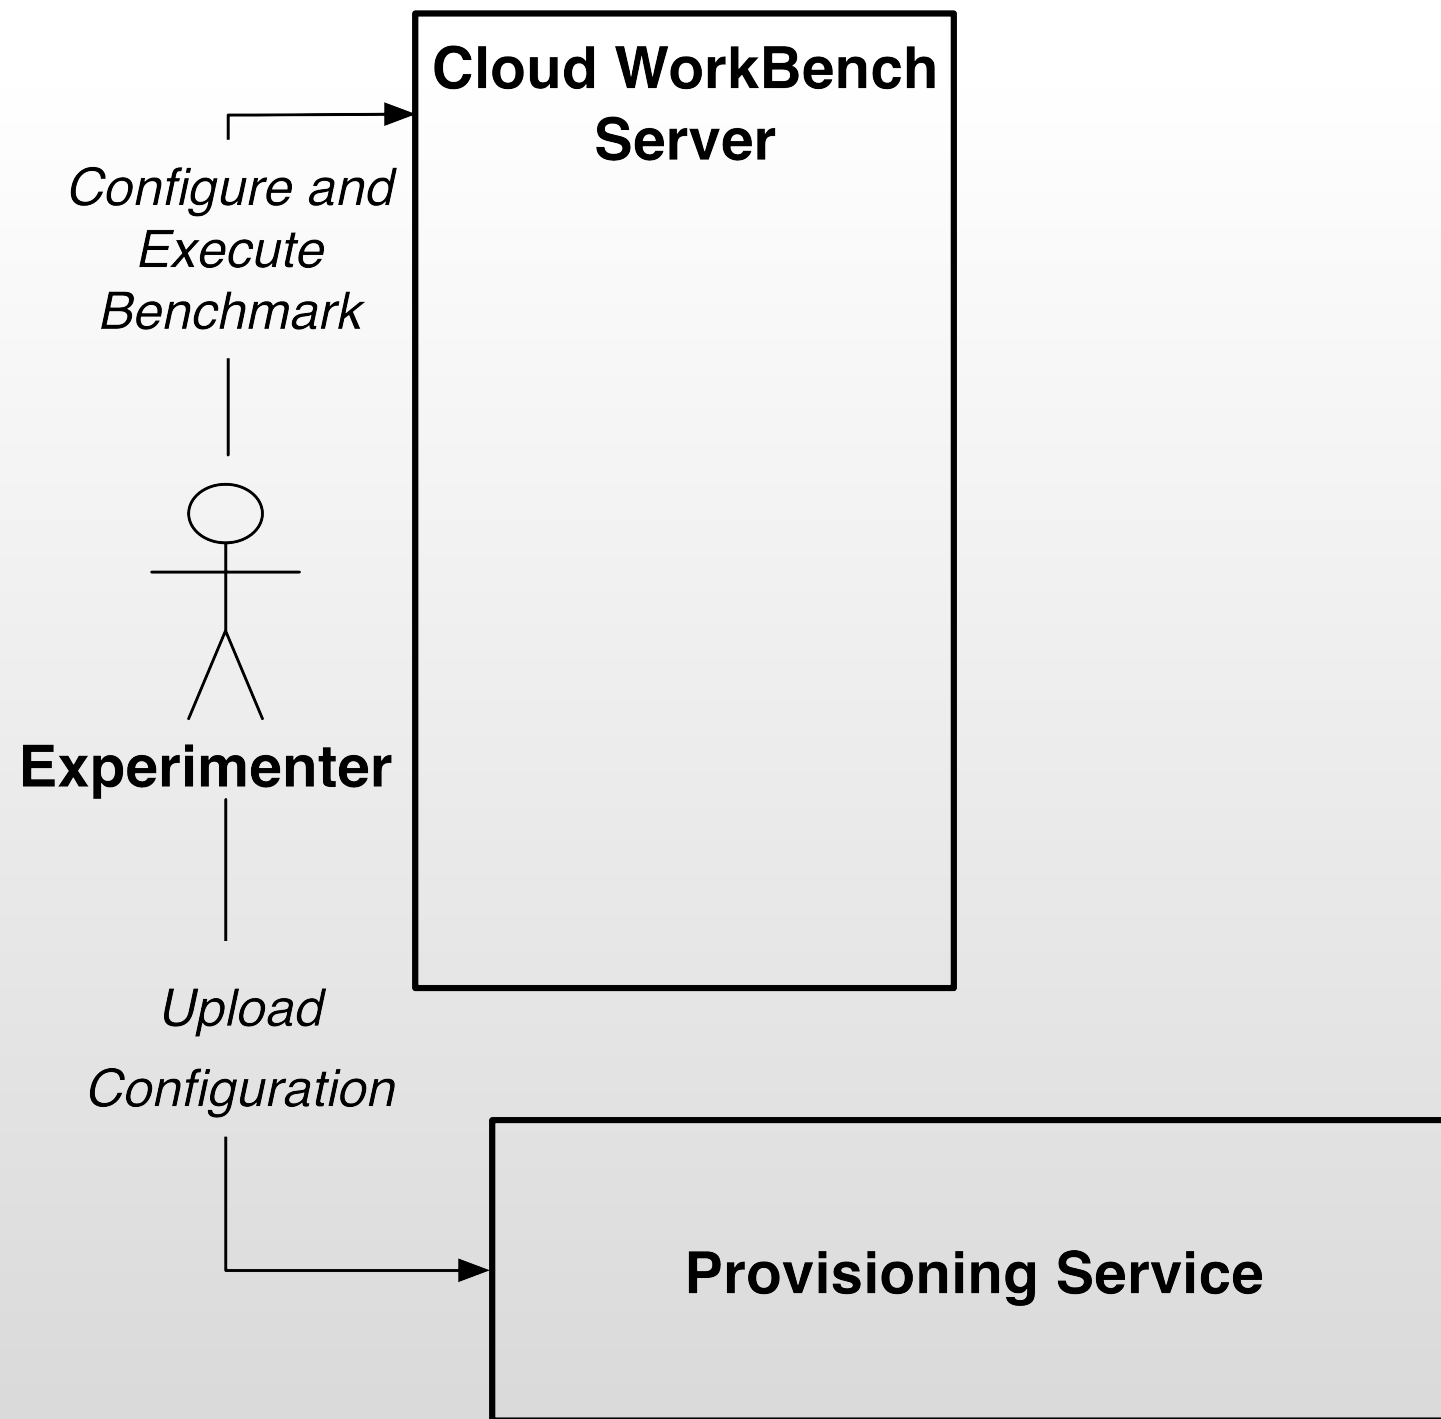
# Benchmark Definition



# Architecture Overview

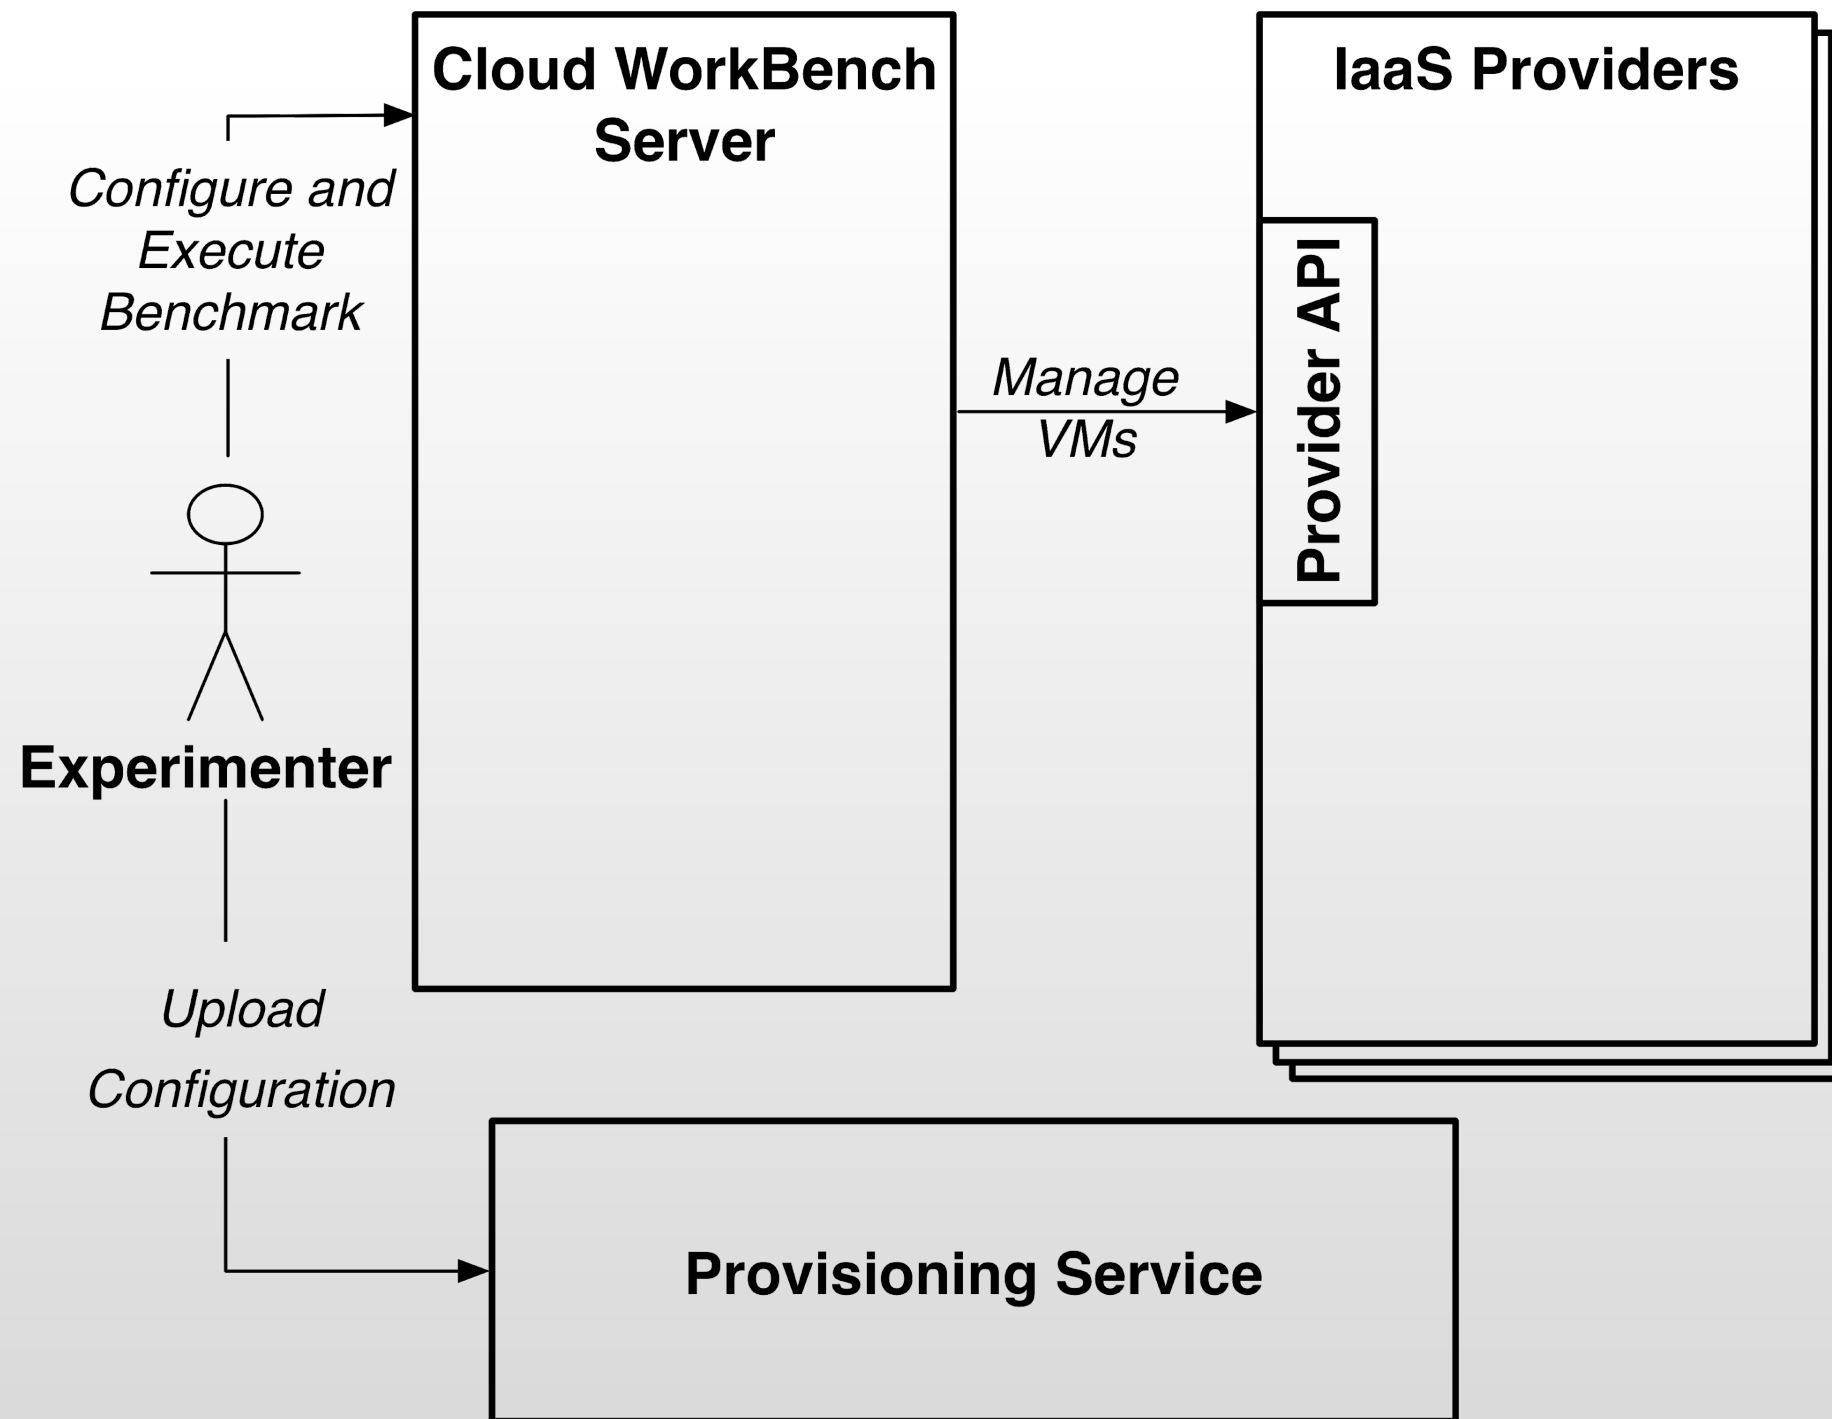


# Architecture Overview

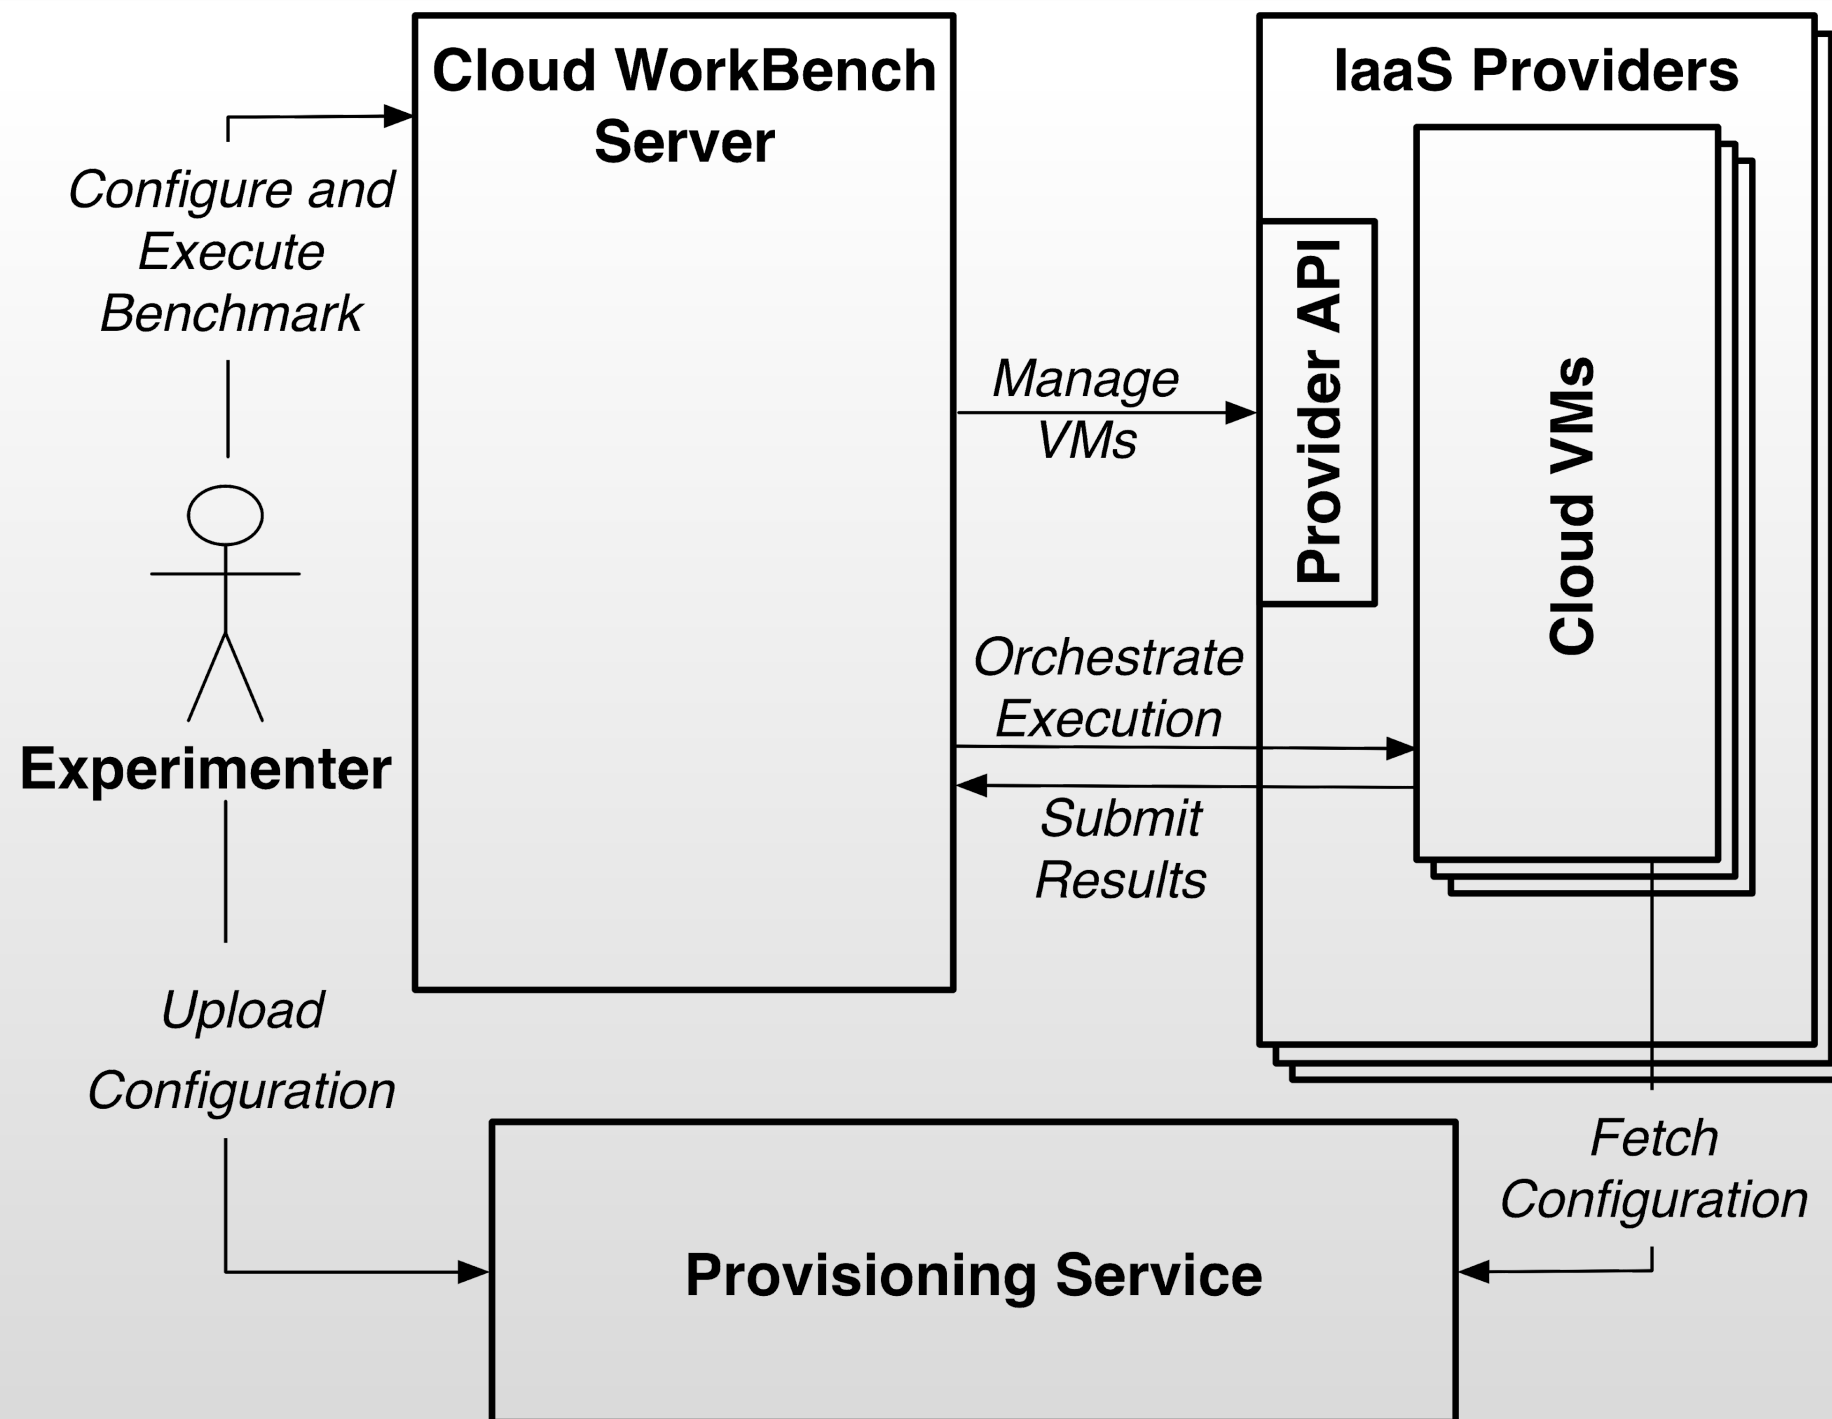




# Architecture Overview



# Architecture Overview



# Implementation (1)

The screenshot shows the 'Edit Benchmark Definition' page for a benchmark named 'google-test'. The interface includes a sidebar with navigation links for Dashboard, Definitions, Executions, Schedules, and Virtual Machines. The main content area contains form fields for Name, Timeout, Provider name, and Vagrantfile. The Vagrantfile is a Ruby script for configuring a VM on Google Cloud. On the right, there are buttons for Delete, Start Execution, and a Schedule section with a cron expression and a Metrics section with a button to create new metric definitions.

**Edit Benchmark Definition** `google-test`

**Name:** google-test

**Timeout for running benchmark:** 1h

**Provider name:** Google

**Buttons:** Delete, Start Execution

**\* Vagrantfile**

```
1 Vagrant.configure(VAGRANTFILE_API_VERSION) do |config|
2   config.vm.provider :google do |google, override|
3     google.image = "debian-7-wheezy-v20140619"
4     google.machine_type = "n1-standard-1"
5     google.zone = "europe-west1-a"
6   end
7
8   config.vm.provision "chef_client", id: "chef_client" do |chef|
9     chef.add_recipe "fio-benchmark@0.3.0"
10    chef.json =
11      {
12        "fio" => {
13          "metric_definition_id" => "seq. write",
14          "config" => {
15            "size" => "10m",
16            "refill_buffers" => "1"
17          }
18        }
19      }
20    end
21  end
22 end
```

**Schedule**

15 5,17 \* \* \*

— 15 minutes past 5am and 5pm of every day

**Buttons:** Edit Schedule, Deactivate Schedule

**Metrics**

**+ Create New Metric Definition**

seq. write  
KB/s (ratio)

cpu  
model name (nominal)

**Update Benchmark**

# Implementation (1.1)

```
1 Vagrant.configure(VAGRANTFILE_API_VERSION) do |config|
2   config.vm.provider :google do |google, override|
3     google.image = "debian-7-wheezy-v20140619"
4     google.machine_type = "n1-standard-1"
5     google.zone = "europe-west1-a"
6   end
7
8   config.vm.provision "chef_client", id: "chef_client" do |chef|
9     chef.add_recipe "fio-benchmark@0.3.0"
10    chef.json =
11    {
12      "fio" => {
13        "metric_definition_id" => "seq. write",
14        "config" => {
15          "size" => "10m",
16          "refill_buffers" => "1"
17        }
18      }
19    }
20  end
21 end
22
```

# Implementation (1.2)



## Edit Benchmark Definition [google-test](#)

\* Name

google-test

Timeout for running benchmark

1h

\* Provider name

Google

\* Vagrantfile

```
1 Vagrant.configure(VAGRANTFILE_API_VERSION) do |config|
2   config.vm.provider :google do |google, override|
3     google.image = "debian-7-wheezy-v20140619"
4     google.machine_type = "n1-standard-1"
5     google.zone = "europe-west1-a"
6   end
7
8   config.vm.provision "chef_client", id: "chef_client" do |chef|
9     chef.add_recipe "fio-benchmark@0.3.0"
10    chef.json =
11      {
12        "fio" => {
13          "metric_definition_id" => "seq. write",
14          "config" => {
15            "size" => "10m",
16            "refill_buffers" => "1"
17          }
18        }
19      }
20  end
21 end
```

# Implementation (1.3)

do |chef|

Start Execution

Schedule

15 5,17 \* \* \*

— 15 minutes past 5am and 5pm of every day

Edit ScheduleDeactivate Schedule

Metrics

Create New Metric Definition

seq. write  
KB/s (ratio)

cpu  
model name (nominal)

# Implementation (2)

The screenshot displays the WorkBench web interface for a benchmark execution. The browser address bar shows the URL `cloud-workbench.io/benchmark_executions/396`. The interface includes a sidebar with navigation links: Dashboard, BENCHMARK (Definitions, Executions, Schedules), and OTHERS (Virtual Machines). The main content area is titled "Execution from 29. June 2014 13:32 of google-test". It features three summary cards: "FINISHED" with a green checkmark, "3 minutes" for Benchmark Duration, and "7 minutes" for Execution Duration. Below these is a timeline of events for June 29, 2014, at 13:32, including "Created", "Started preparing", "Finished preparing", "Started running", and "Finished running". The "Started preparing" event is expanded, showing a log of system changes and file operations. On the right, there are buttons for "Start Execution" and "Delete", a section for "1 Virtual Machine" (google, default, cwb-38-396), a "Schedule" section (15 5,17 \* \* \*), and a "Metrics" section with a "Create New Metric Definition" button.

Execution 2014-06-29 of x

cloud-workbench.io/benchmark\_executions/396

WorkBench

Sign out

Dashboard

Execution from 29. June 2014 13:32 of google-test

**FINISHED**

Status

**3 minutes**

Finished

Benchmark Duration

**7 minutes**

Finished

Execution Duration

29. June 2014

13:32 Created

13:32 Started preparing

Download Start live refresh +

```
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] owner changed to 0
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] group changed to 0
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] mode changed to 755
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] created file /usr/local/cloud-benchmark/start.sh
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] updated file contents /usr/local/cloud-benchmark/start.sh
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] owner changed to 0
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] group changed to 0
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] mode changed to 755
[2014-06-29T11:35:21+00:00] INFO: Chef Run complete in 86.557784584 seconds
[2014-06-29T11:35:21+00:00] INFO: Running report handlers
[2014-06-29T11:35:21+00:00] INFO: Report handlers complete
```

13:35 Finished preparing

13:35 Started running

13:38 Finished running

Start Execution

Delete

1 Virtual Machine

google default cwb-38-396

Schedule

15 5,17 \* \* \*

— 15 minutes past 5am and 5pm of every day

Edit Schedule Activate Schedule

Metrics

Create New Metric Definition

# Implementation (2.1)

29. June 2014



13:32 Created



13:32 Started preparing

Download

Start live refresh



```
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] owner changed to 0
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] group changed to 0
[2014-06-29T11:35:21+00:00] INFO: cookbook_file[/usr/local/cloud-benchmark/fio_log_parser.rb] mode changed to 755
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] created file /usr/local/cloud-benchmark/start.sh
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] updated file contents /usr/local/cloud-benchmark/start.sh
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] owner changed to 0
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] group changed to 0
[2014-06-29T11:35:21+00:00] INFO: template[/usr/local/cloud-benchmark/start.sh] mode changed to 755
[2014-06-29T11:35:21+00:00] INFO: Chef Run complete in 86.557784584 seconds
[2014-06-29T11:35:21+00:00] INFO: Running report handlers
[2014-06-29T11:35:21+00:00] INFO: Report handlers complete
```



13:35 Finished preparing



13:35 Started running



13:38 Finished running

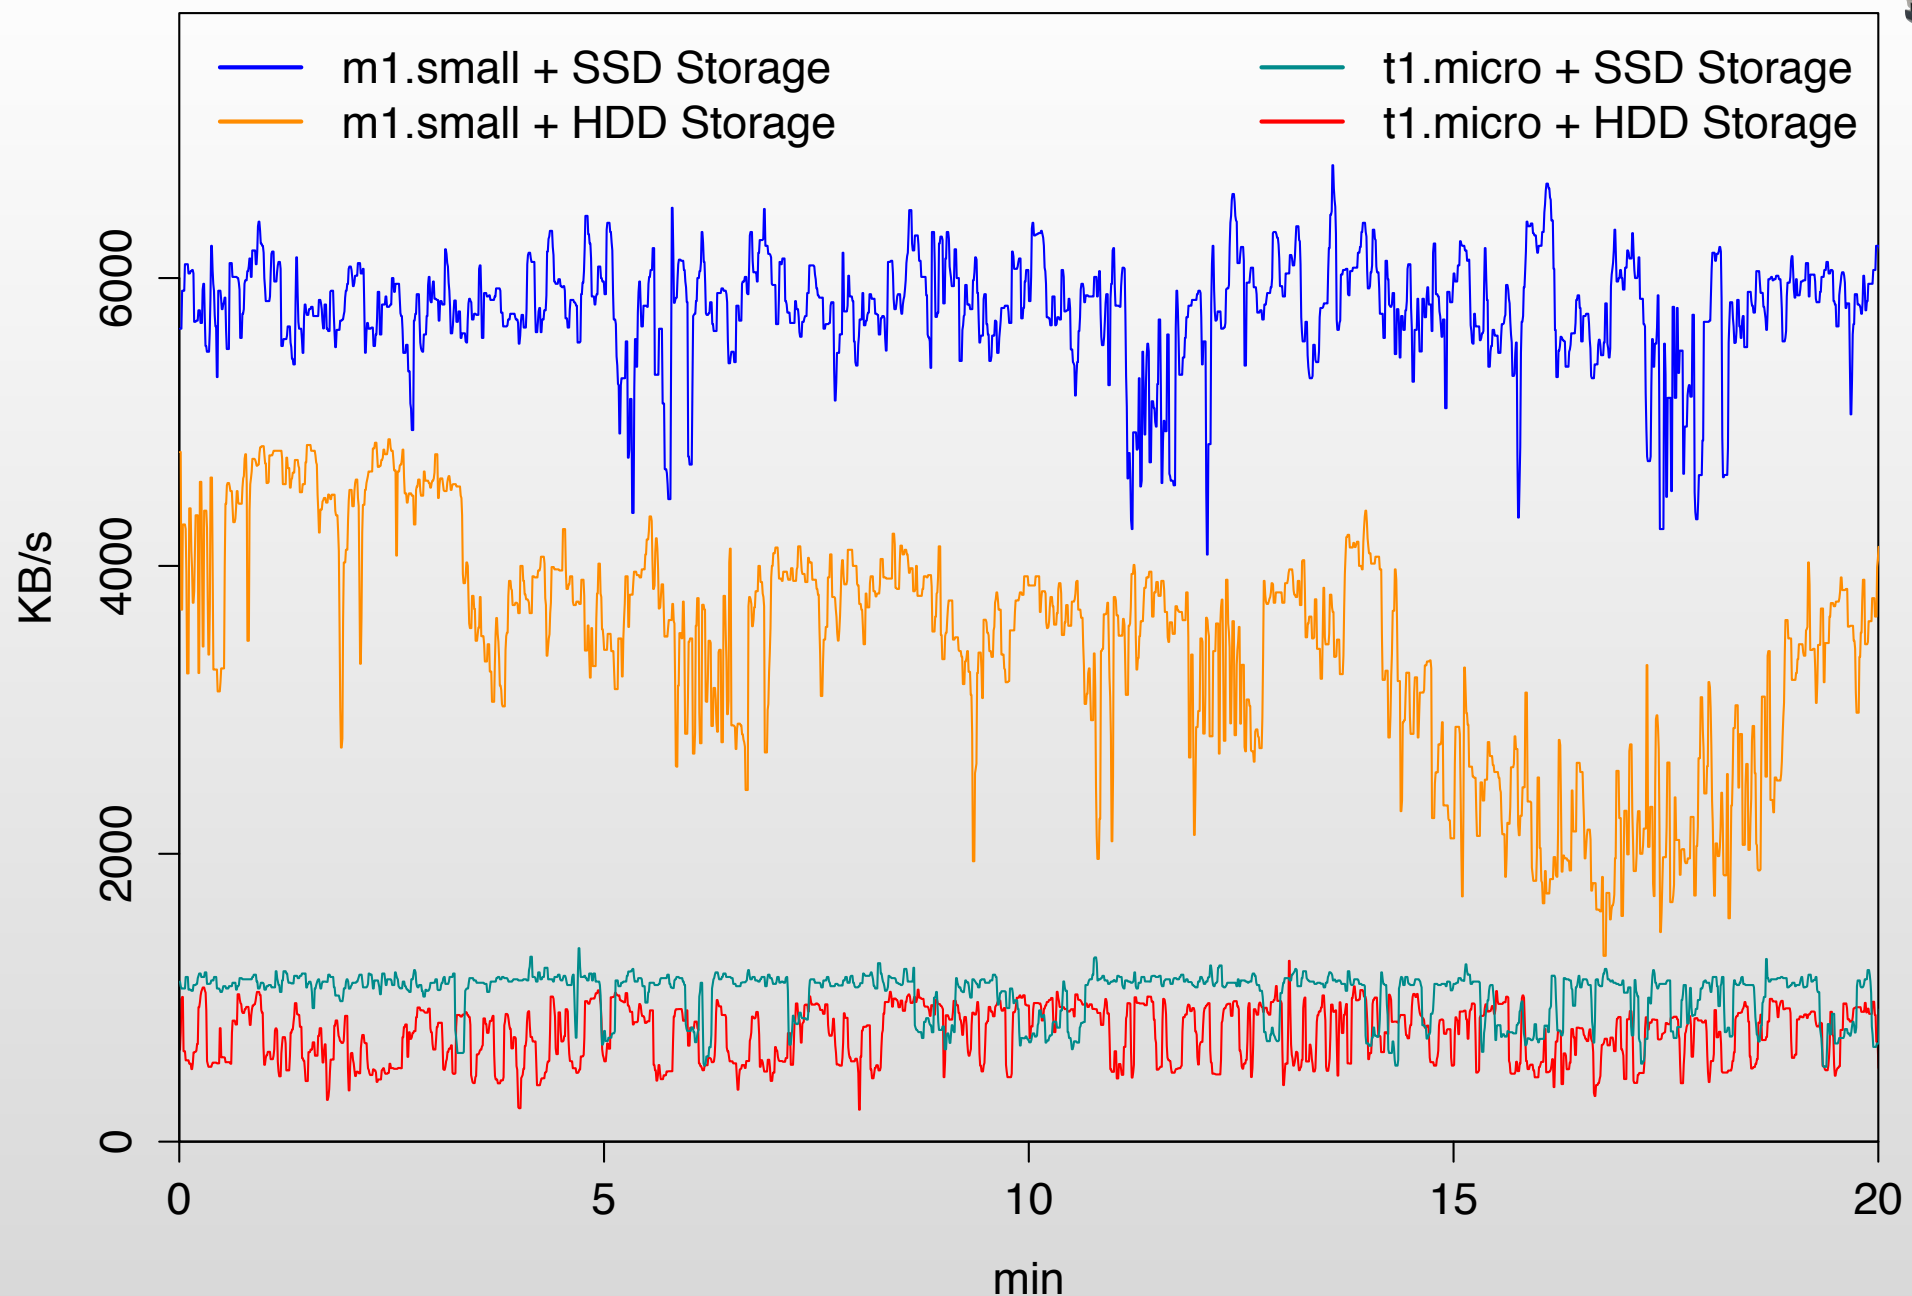


13:38 Started postprocessing



# Case Study

## Sequential Disk Write Speed



# Conclusion

- Cloud WorkBench:  
Open source cloud experiment automation tool
- Infrastructure-as-Code for cloud benchmarking
  - Easily configurable
  - Reproducibly executable

# Cloud WorkBench

The screenshot displays the Cloud WorkBench web interface. At the top, the 'WorkBench' logo and a 'Sign out' link are visible. The left sidebar contains navigation links for 'Dashboard', 'BENCHMARK' (Definitions, Executions, Schedules), and 'OTHERS' (Virtual Machines). The main content area shows an execution from May 30, 2014, for the 'www-app-compile' benchmark on 'aws-eu-large' instances. Three summary cards indicate the status: 'FINISHED' (green), '6 minutes' benchmark duration (blue), and '15 minutes' execution duration (orange). Below these, a list of execution steps is shown with green checkmarks: 'Created', 'Started preparing', 'Finished preparing', 'Started running', and 'Finished postprocessing'. A terminal window displays log output for the 'compile-benchmark.rb' script, showing file permissions and Chef Run completion. On the right, there are buttons for 'Start Execution' and 'Delete', and a section for '1 Virtual Machine' with details like 'aws', 'default', and 'i-da30f39a'. Below that, a 'Schedule' section shows a cron job '15 3,7,11,15,19,23 \*\*\*' and buttons for 'Edit Schedule' and 'Activate Schedule'. A 'Metrics' section at the bottom right has a 'Create New Metric Definition' button.

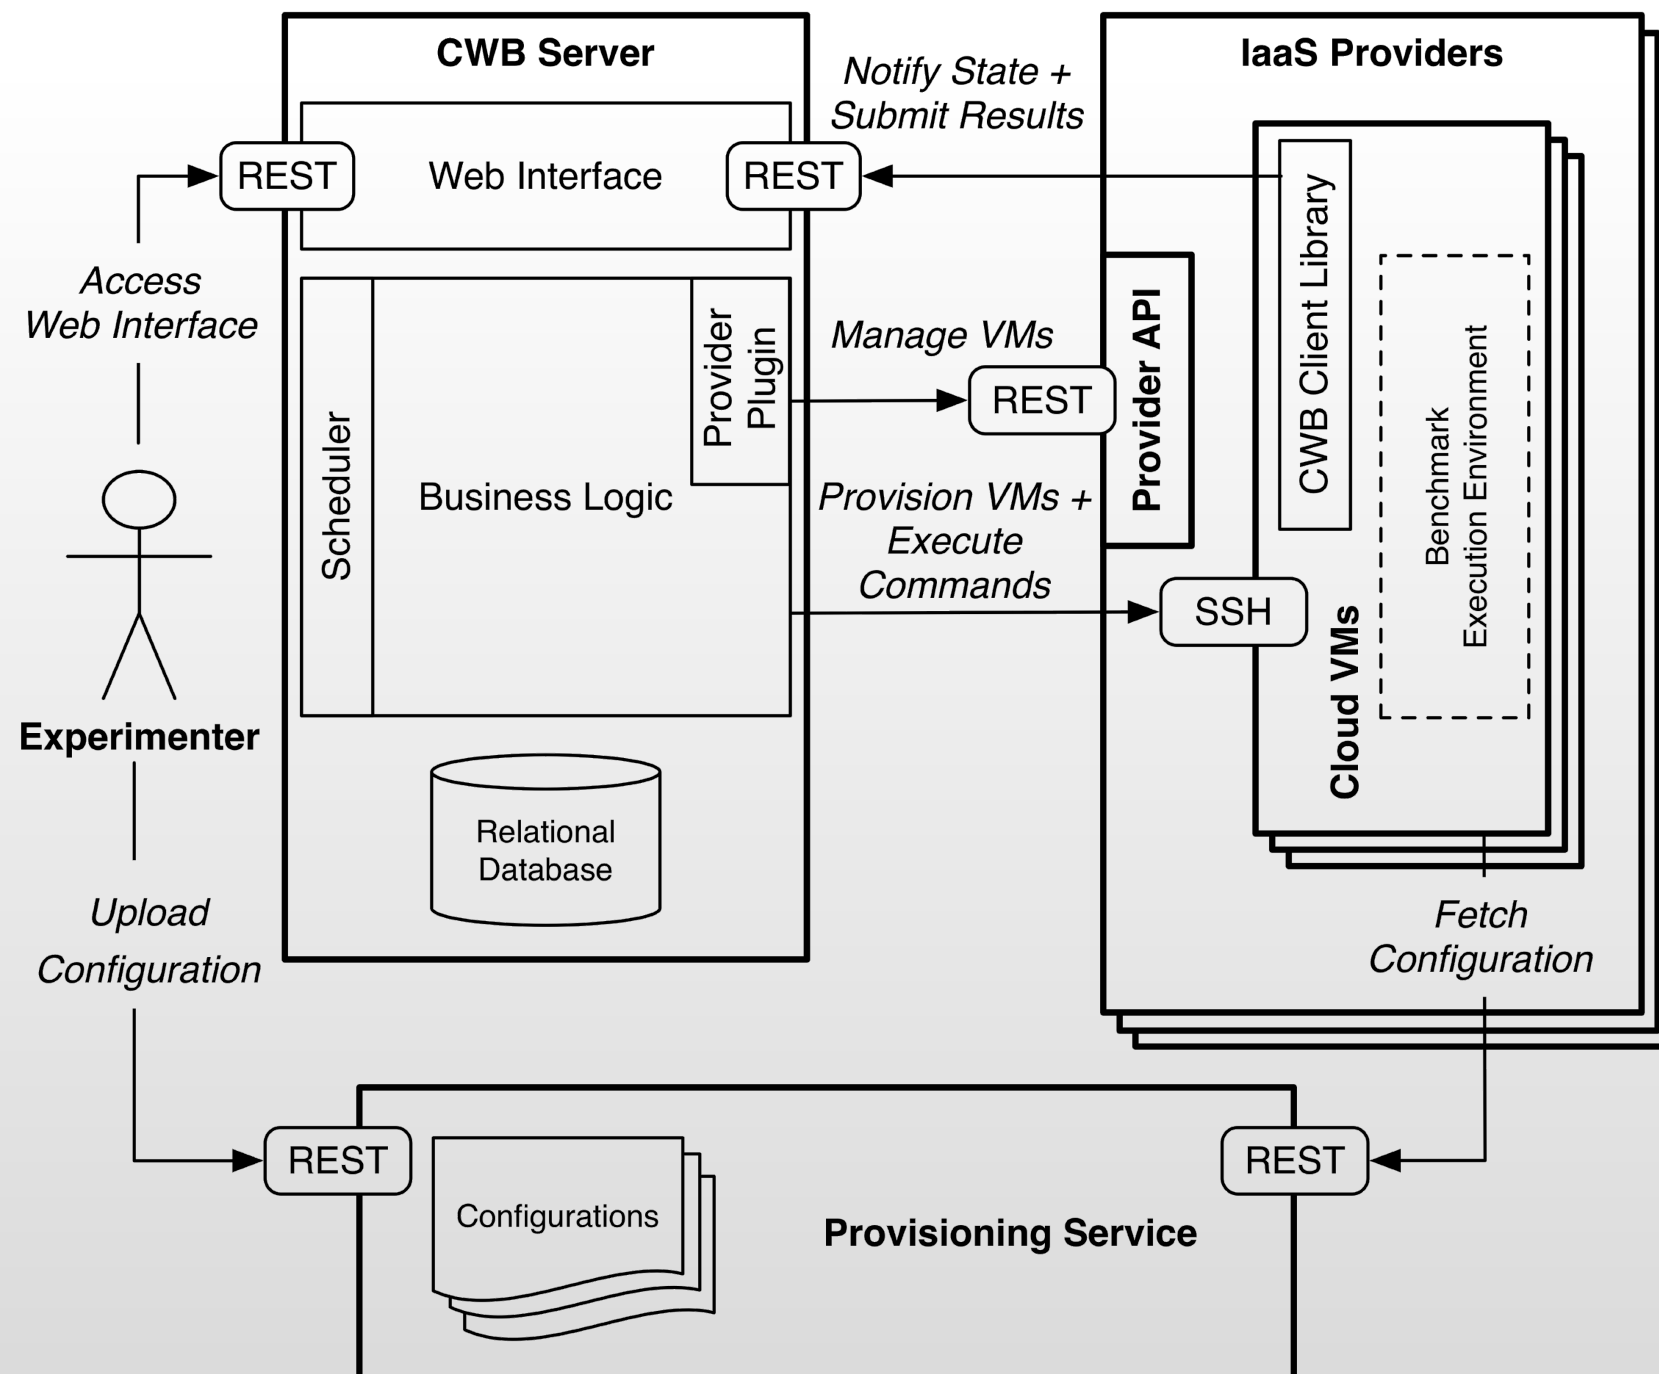
**Open Source**

<https://github.com/sealuzh/cloud-workbench>

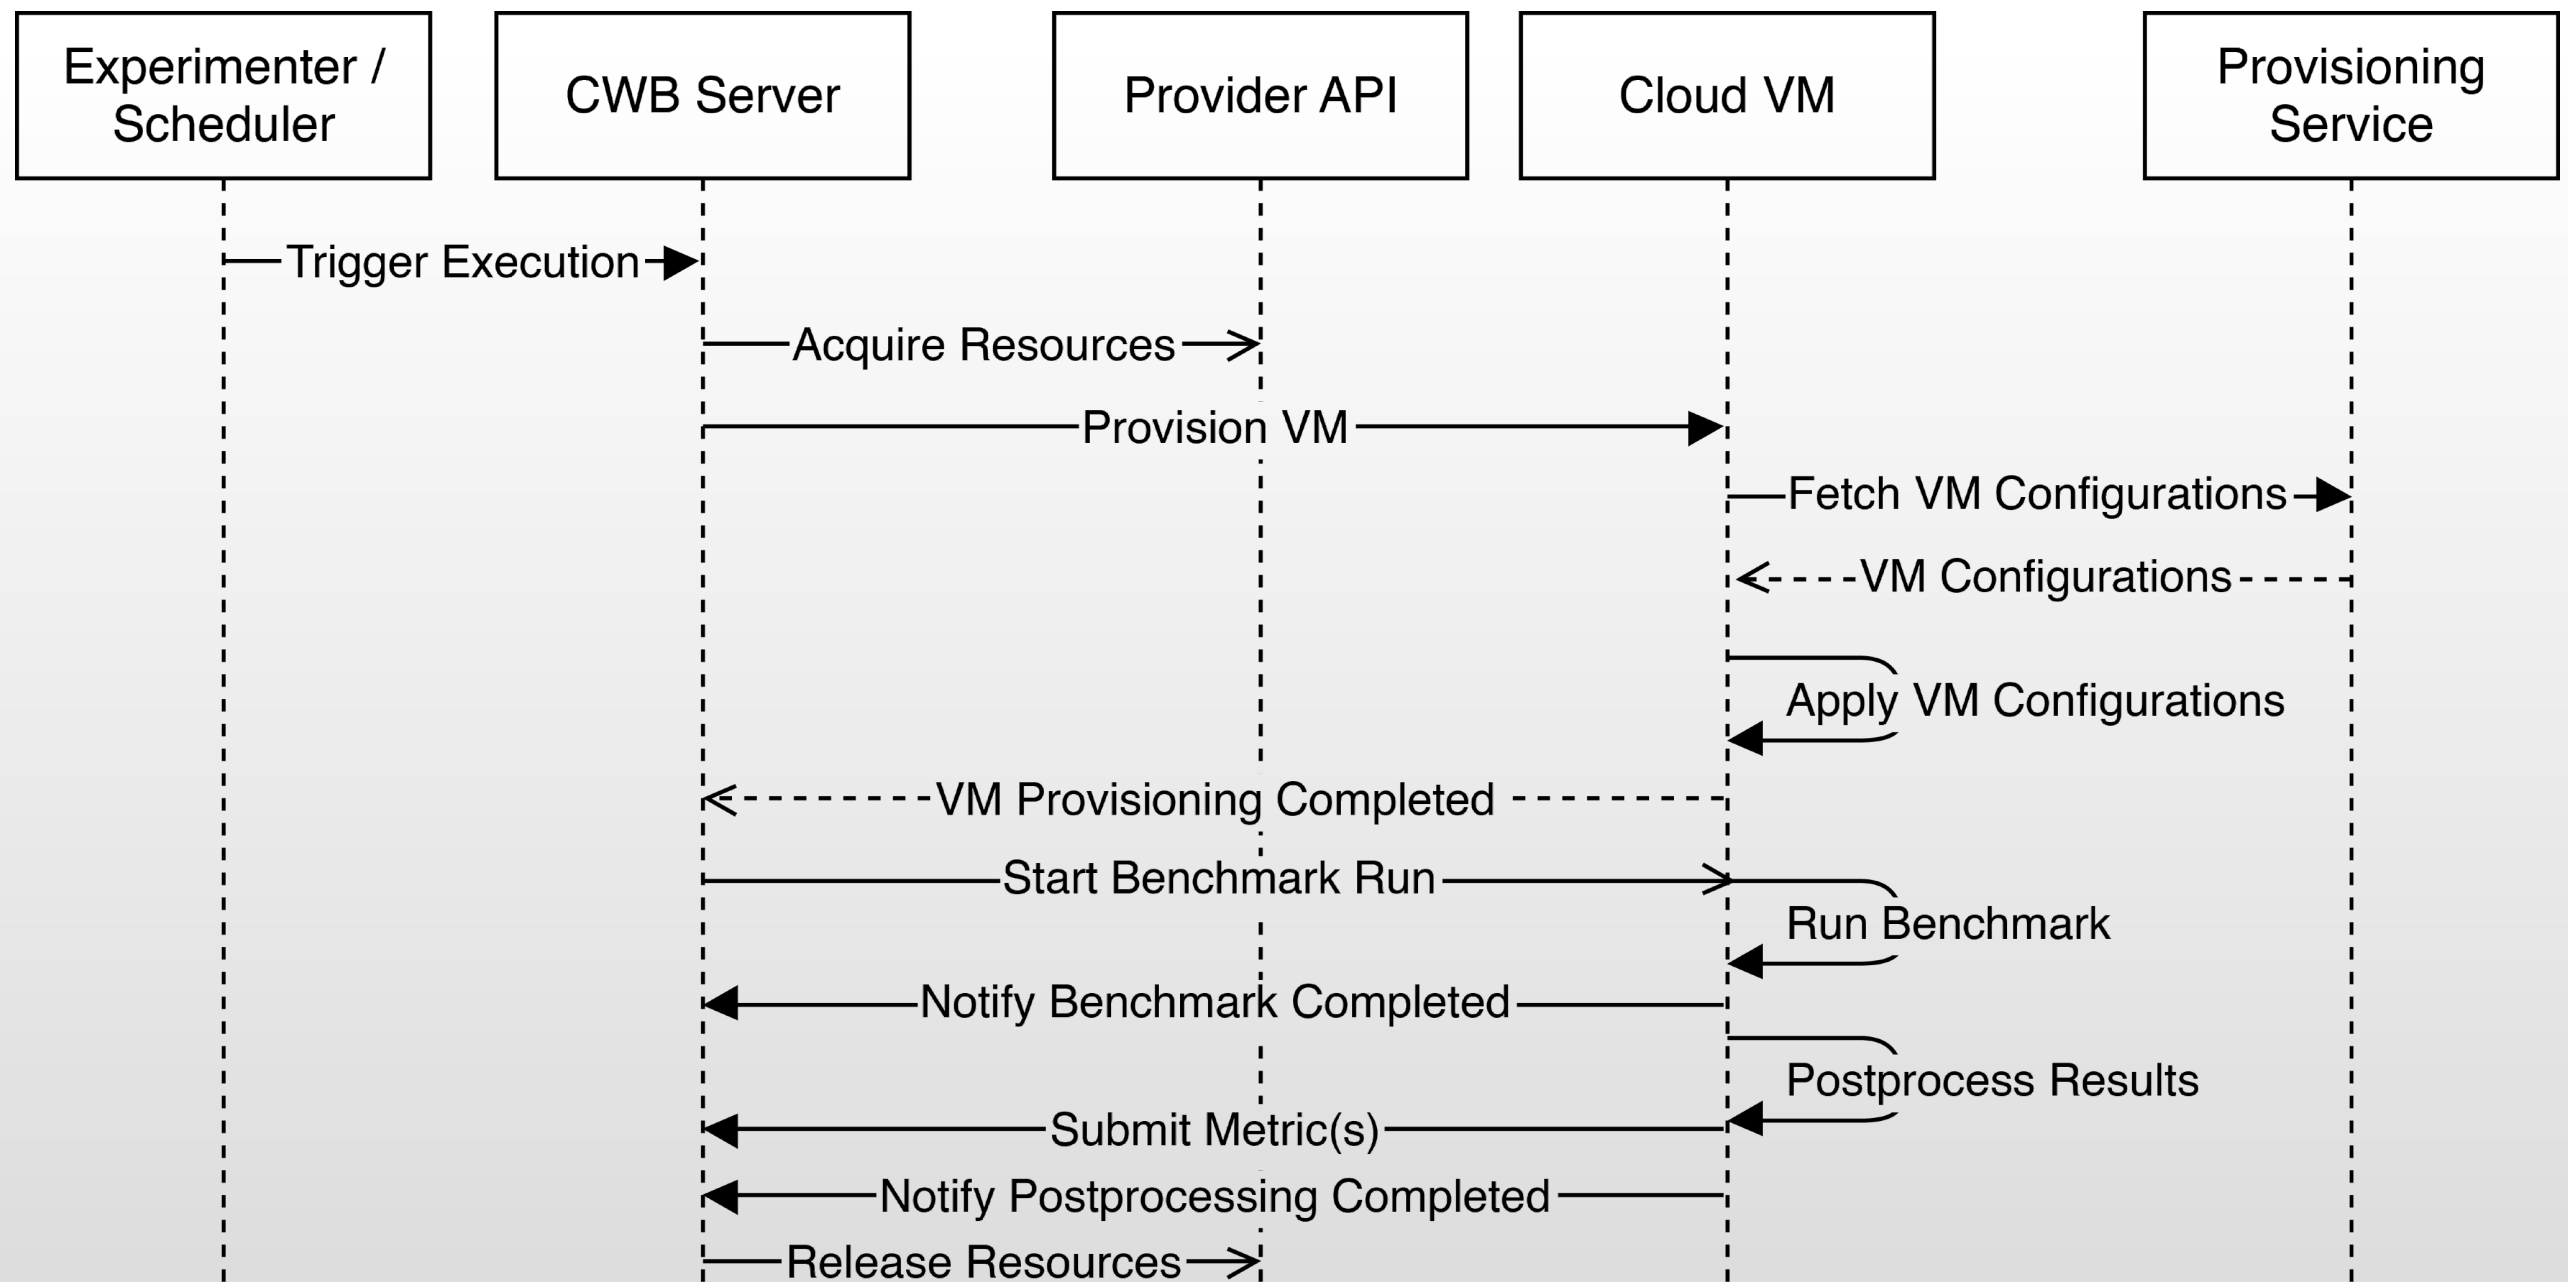
Further results: „Patterns in the Chaos“ <http://arxiv.org/abs/1411.2429>



# Architecture Overview



# Benchmark Execution



# Future Work

- Support additional cloud providers  
(Currently: Amazon EC2 + Google Compute Engine)
- Integrate statistical analysis capabilities
- Support the entire benchmarking lifecycle via a single web tool  
(define, execute, analyze)

