

# Update on monitoring at the MOC



Vinaykumar Kapalavai  
Tariq Anwar  
Jeremy Freudberg  
Raja Sambasivan  
Ata Turk  
Alina Oprea  
Ayse Kivilcim Coskun

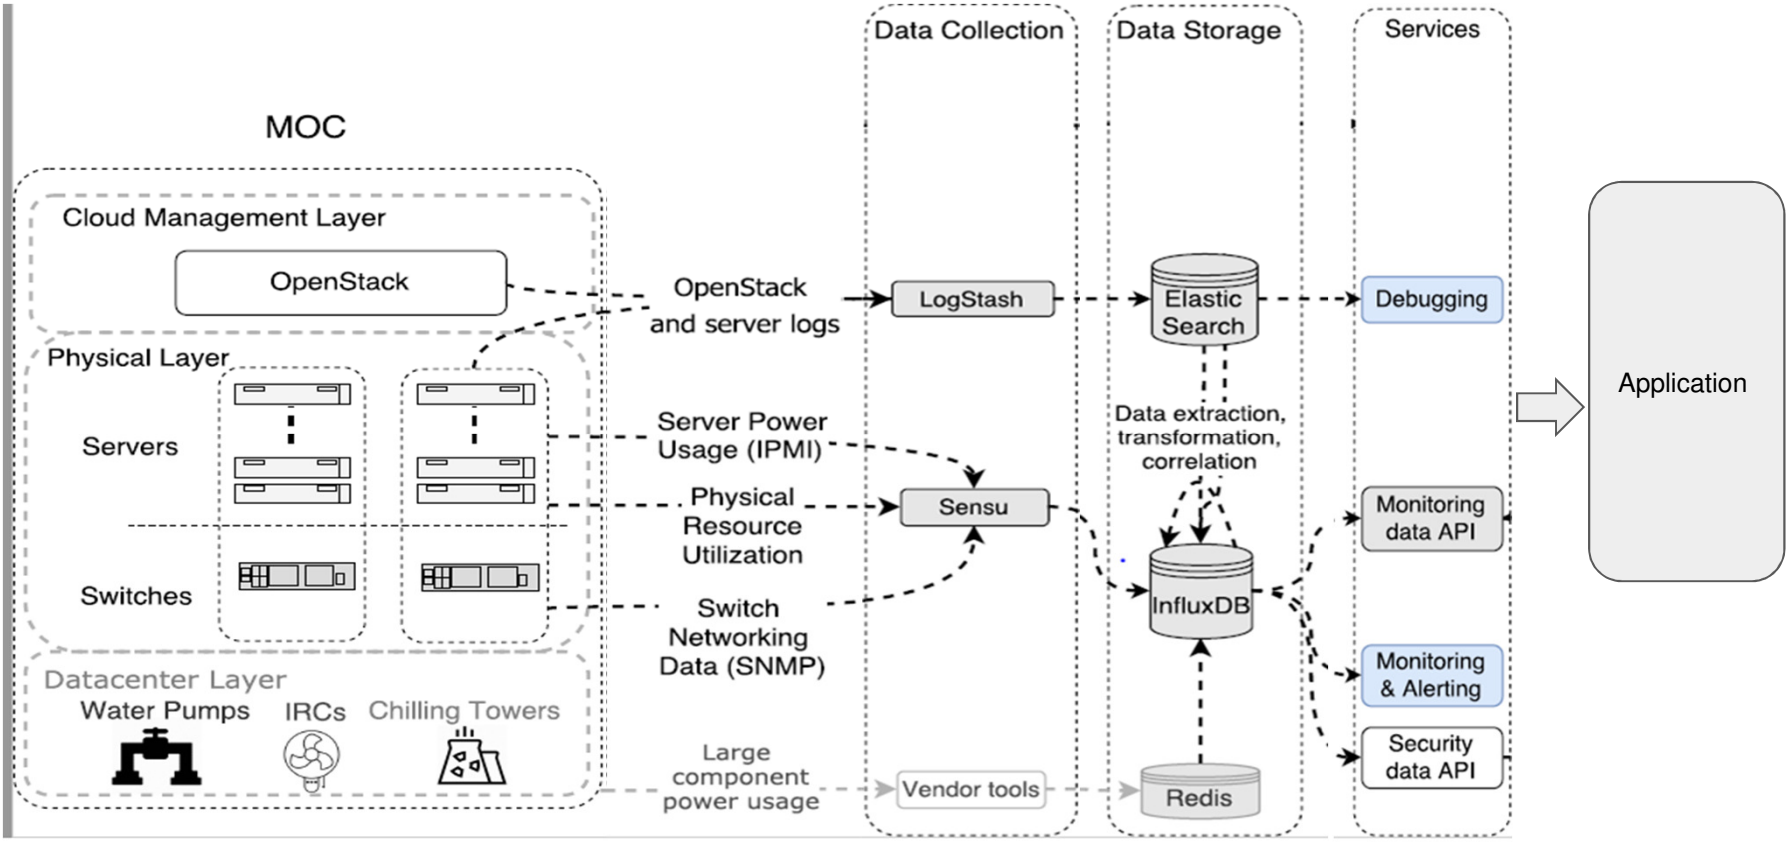
# Overview of Monitoring Efforts

- ❑ Infrastructure to provide visibility for Engineers
  - Proactively find problems
  - Alert and notify
  - To optimize our cloud(Capacity & Consumption)
- ❑ Give researchers access to data sets

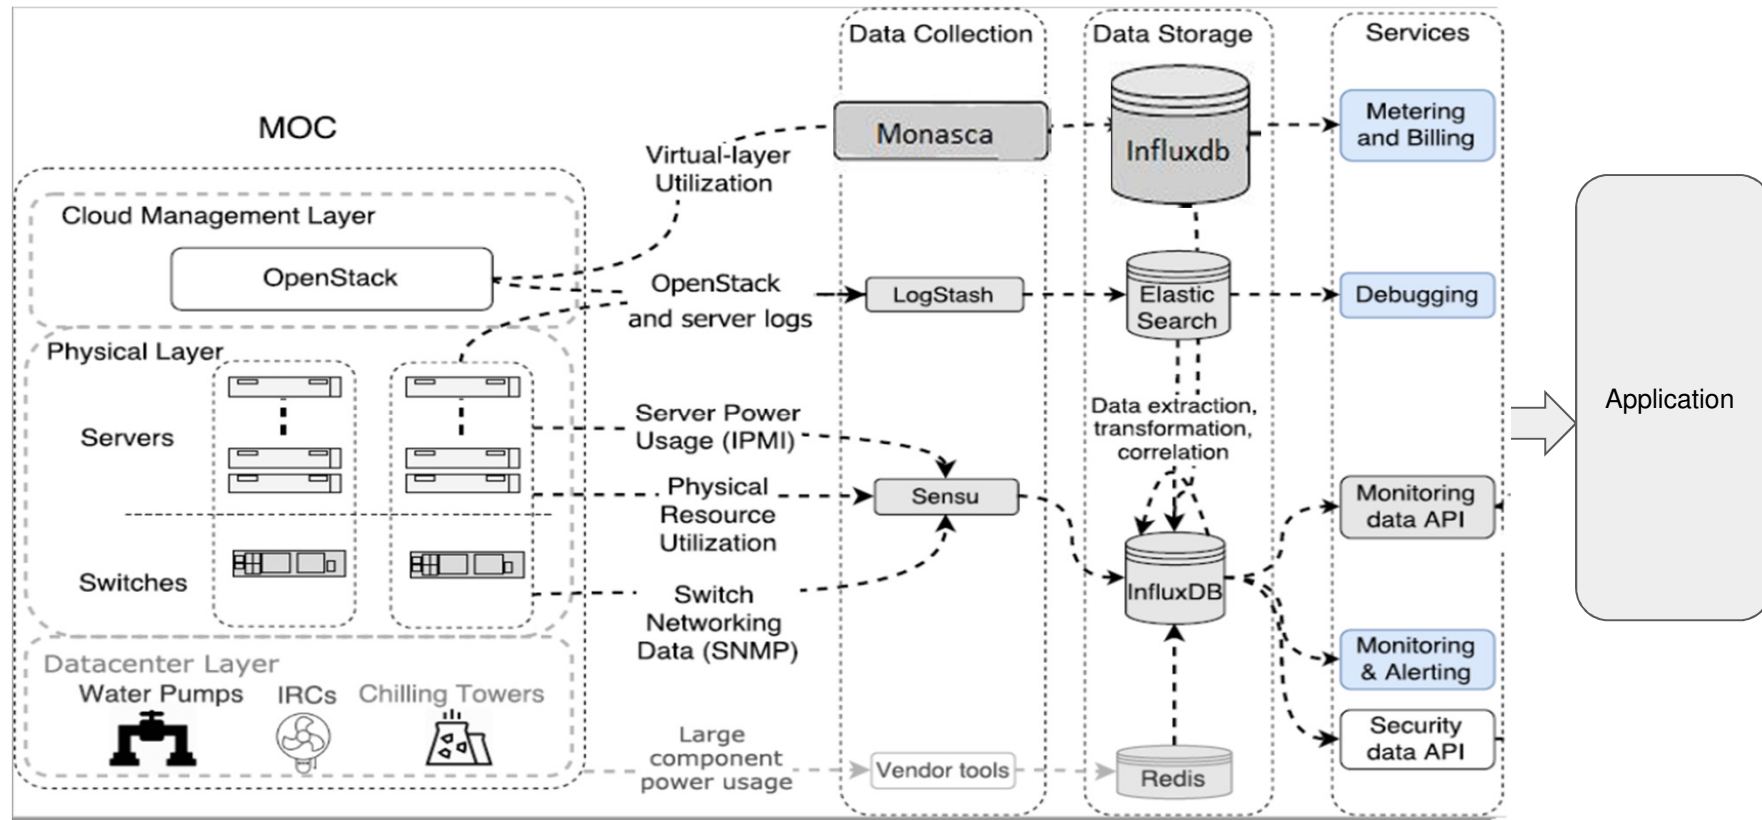
# MOCMON: Scalable cloud monitoring system

- ❑ Collects Data From Multiple Layers of Cloud
- ❑ Built out of open source tools
- ❑ Current statistics:
  - 162 metrics
  - ~78,317,280 data points per day
- ❑ Data archived in Openstack Swift

# Architecture



# 1 Effort #1: Virtual data collection



## Slide 5

---

1

Transition that shows virtual layer showing up in architecture diagram is super sudden. No one will notice it appearing.

Raja Sambasivan, 9/29/2017

# Overview of virtual layer information

- ❑ Virtual layer data – Ex: vcpu, memory utilization of vm

Last Year: Ceilometer

- ❑ Frame work to meter and collect physical and virtual layer resource utilization
- ❑ High performance overhead

This Year: Monasca

- ❑ Monitoring frame work to monitor all openstack components.
- ❑ Low performance overhead

4

## Performance testing Monasca

5

6

### □ Experimental setup:

- Deployed Monasca on kumo cluster (6 node cluster)
- Rally to stress test openstack cloud
- Benchmarked using rally workloads



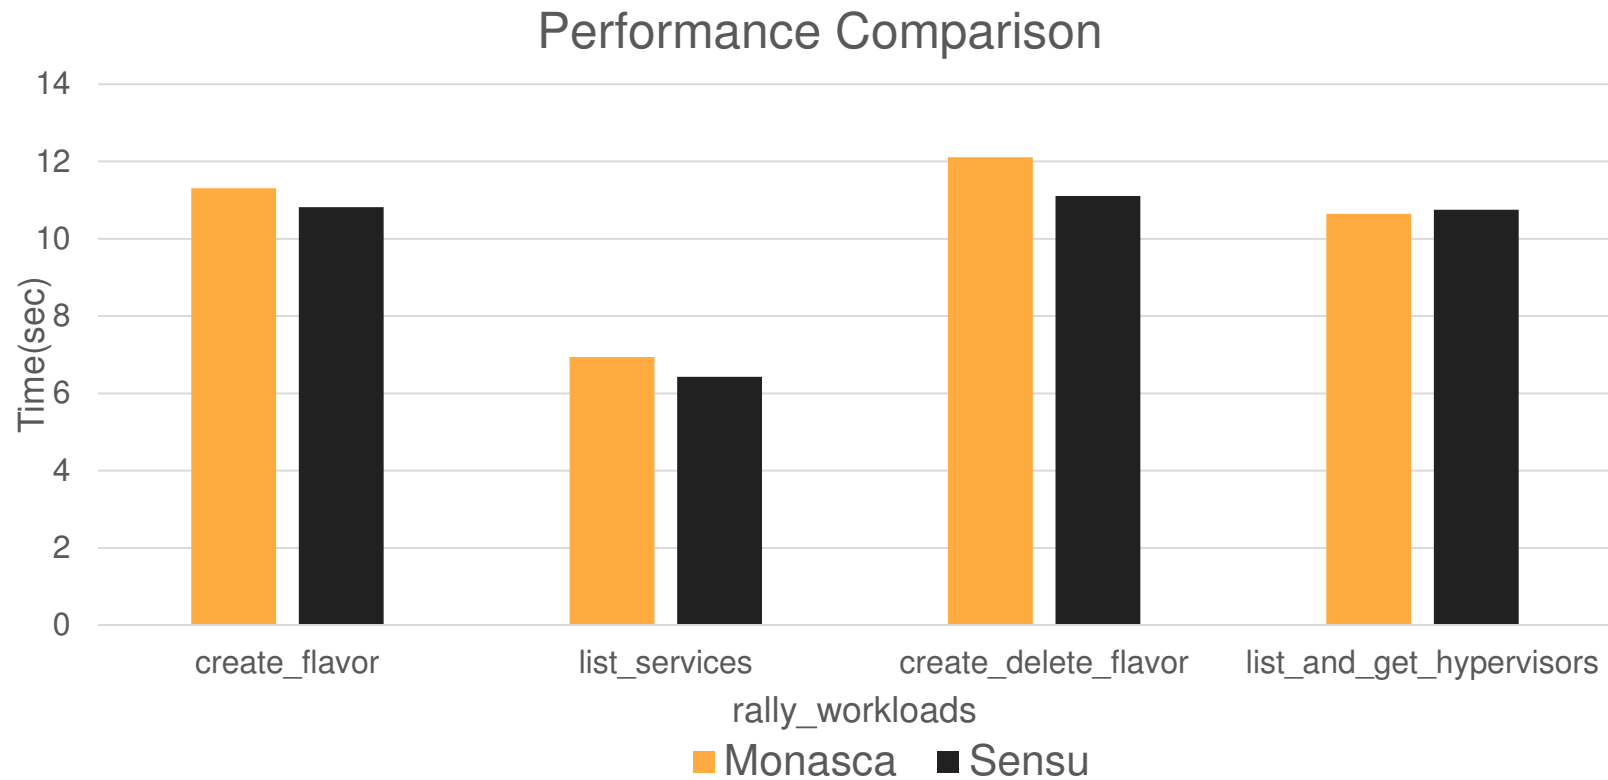
## Slide 7

---

- 6**      Make sure to explain what the sumo cluster is.  
Raja Sambasivan, 9/29/2017
- 4**      Make this slide "Performance testing Monasca." Then, on this slide you can state the experimental setup.  
Raja Sambasivan, 10/2/2017
- 5**      What is rally? What does it do? How many nodes were you benchmark run against? Need a lot more detail about the experimental setup here.  
Raja Sambasivan, 10/2/2017

7

# Results : Monasca overhead similar to Sensu



## Slide 8

---

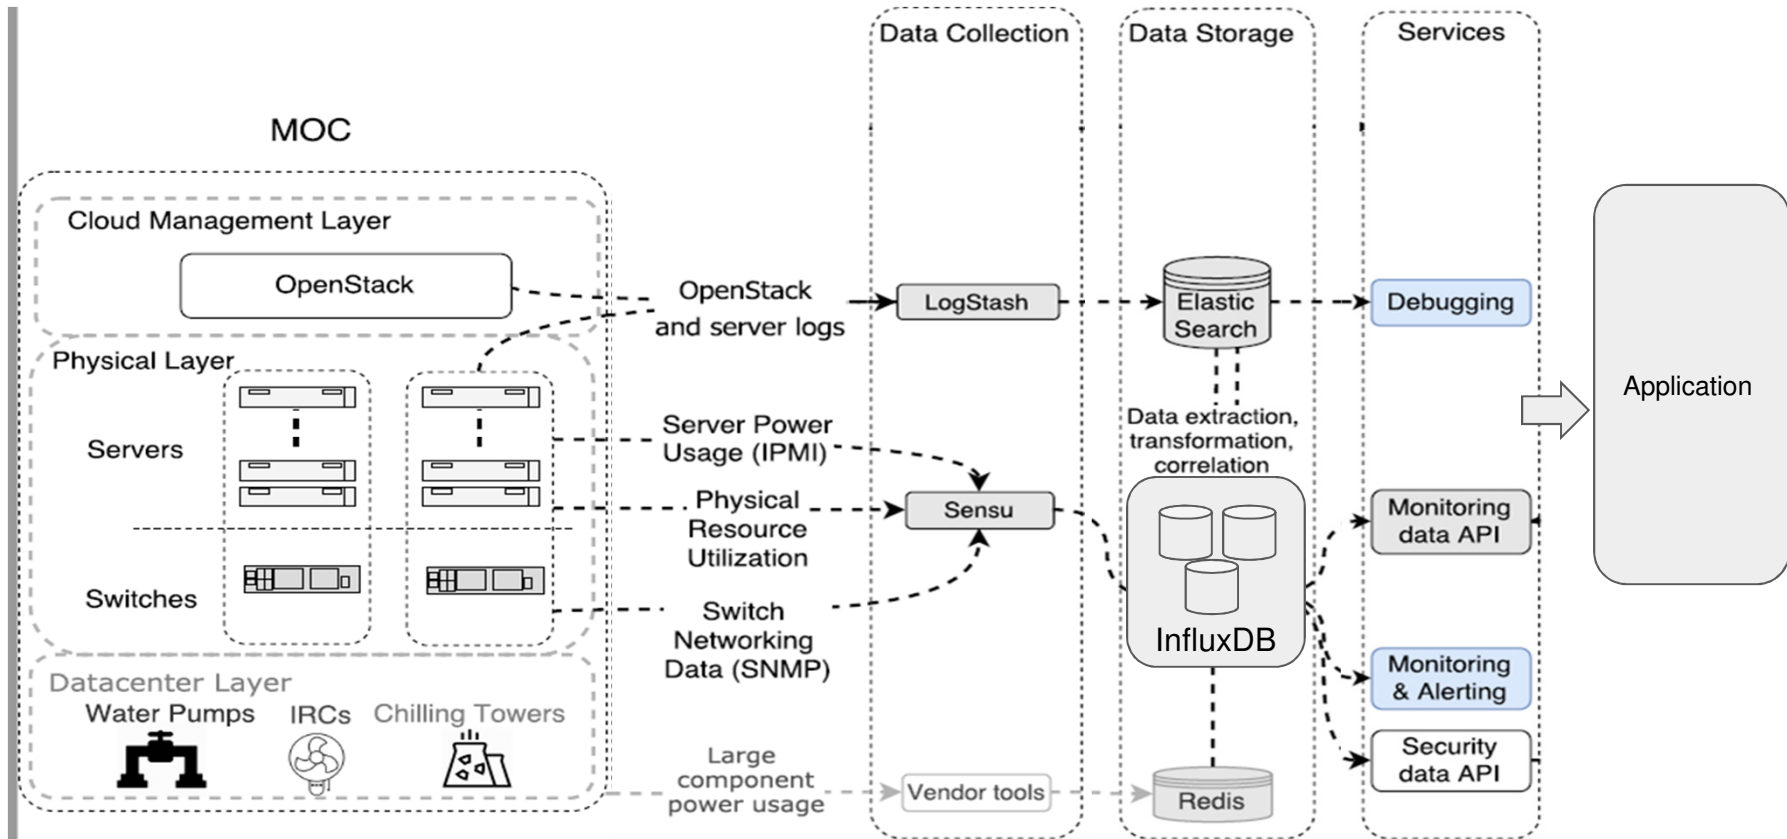
7 I like this slide, but font sizes are way too small. No one will be able to see. When presenting, I suggest animating in the axes to what you can state what they represent. Then, you can animate in the bars.

So, something like: "On this slide, I'm going to show you our experimental results. [ANIMATE] On the Y axis you'll see the time taken for each of the benchmarks. [ANIMATE] And on the X-axis, you see the individual benchmarks. [ANIMATE] Blue will represent Monasca and Yellow will represent Sensus. Here are the results."

Raja Sambasivan, 10/2/2017

8

# Effort #2: A more scalable storage database



**Slide 9**

---

**8**

**Make sure people will realize something is changing in this diagram**

Raja Sambasivan, 10/2/2017

## Evaluating clustered Influx DB

- ❑ Collected Monitoring Information Is Persisted In Time-Series Database
- ❑ Clustered InfluxDB Database
- ❑ Fault-Tolerant
- ❑ Highly Available

# How is the monitoring infrastructure being used?

## ❑ Infrastructure team:

- Obtain alerts on the consumption of hardware resources
- Query from the Kibana dashboards for debugging purposes

## ❑ Research team:

- Understanding what cloud providers can learn about tenants
- New techniques for identifying anomalies and problems

# Summary

- ❑ The importance of monitoring and current infrastructure of MOC
- ❑ The ongoing efforts in MOCMON
- ❑ The use cases for monitoring data