

Benchmarks in the Cloud

Steve Newman
Scalyr



Scalyr Logs

Juggling too many monitoring tools? Not getting full use of your logs? Scalyr Logs is a game-changer. Analyze, graph, aggregate, drill down. Powerful, fast, and easy to use.

[More Details →](#)

[Get Started →](#)

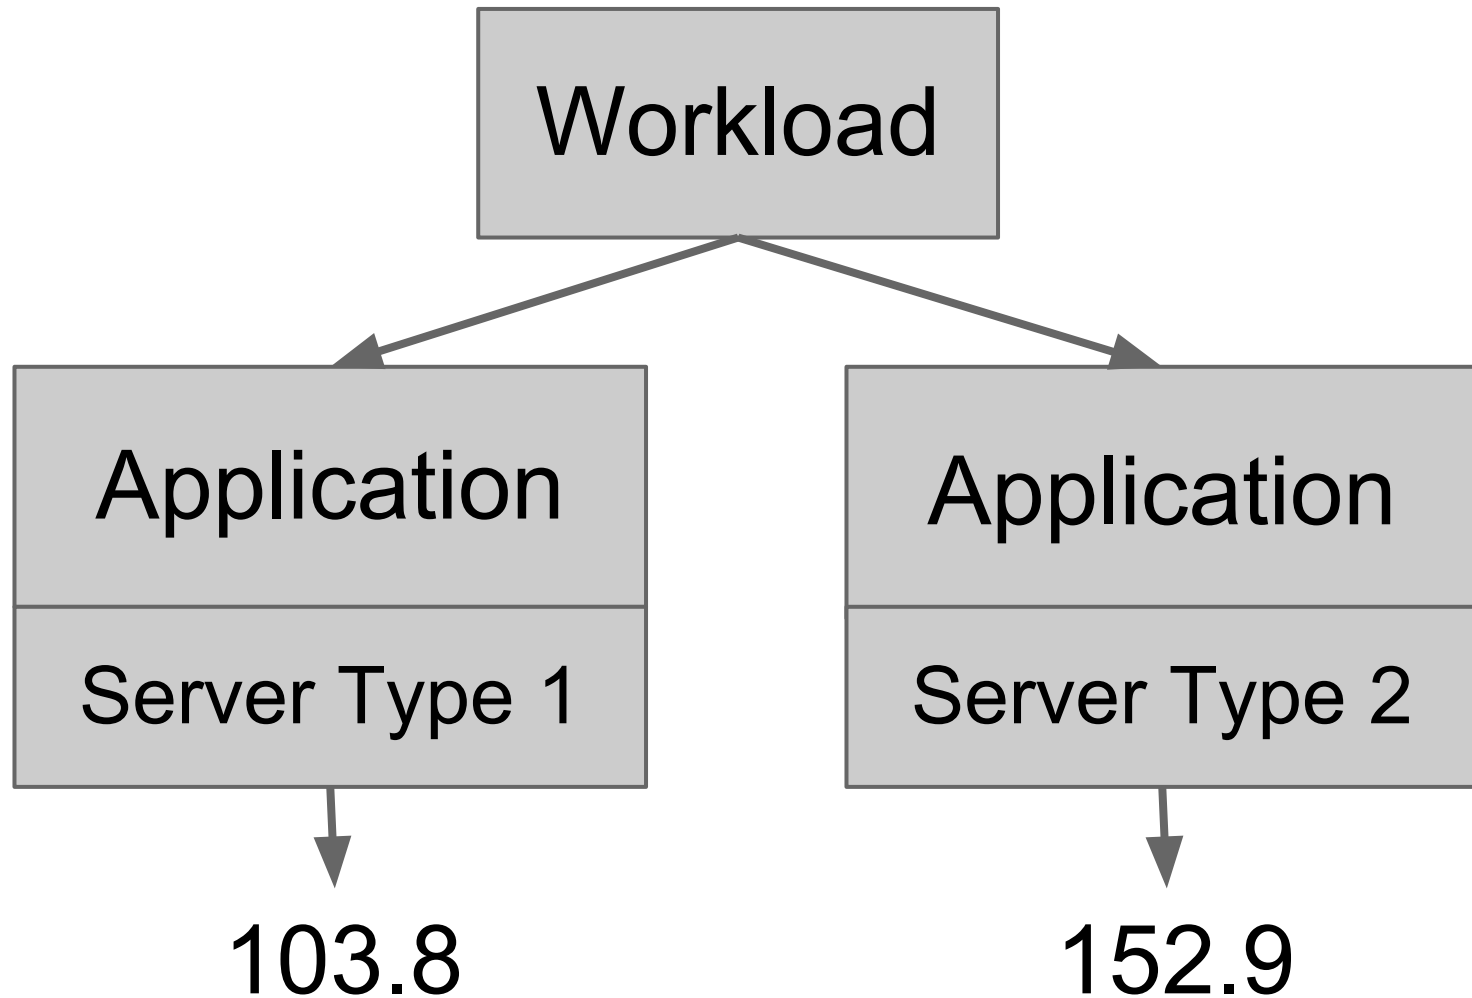
How I'll Spend Seven Minutes

- Engineering decisions require data
- Good decisions require good data
- In the cloud, that means multiple measurements

Importance of Perf Benchmarks

- Validate design assumptions
- Evaluate alternatives
- Capacity planning / load tests

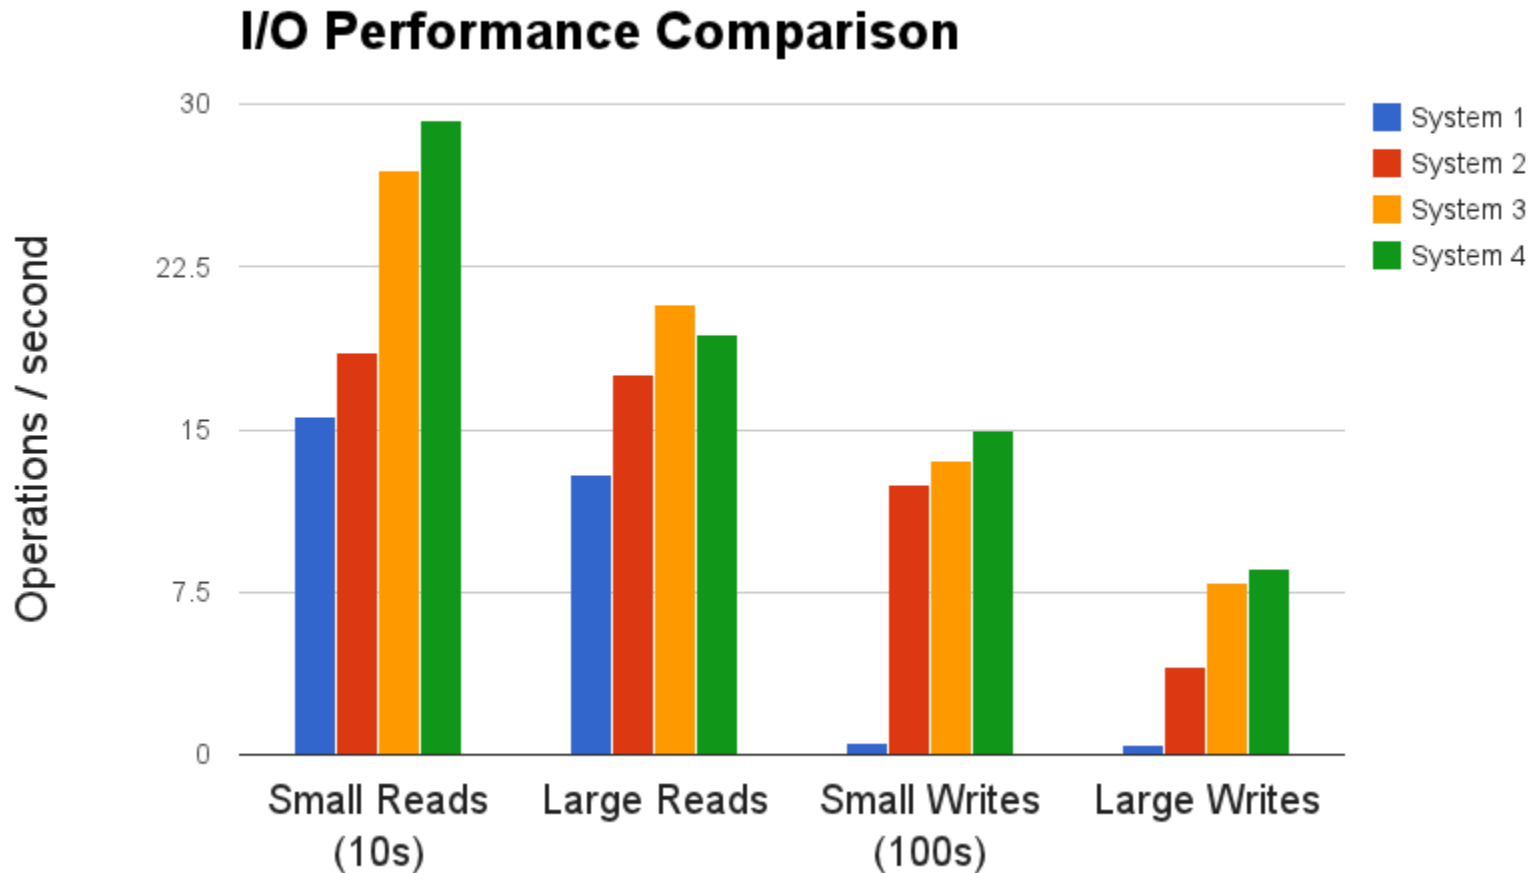
Controlled Experiments



What Are You Measuring?

- Behavior of an application
- With a particular tuning
- Under some workload
- On a specific server type
- And system configuration
- From a particular hosting provider
- In a particular data center
- On a specific machine
- At a point in time
- Under some state of the network and supporting services

Some Variables are Outside Your Control



Test Many, Trust None

- Run each test many times
 - On different instances
 - At different times
- Aggregating across instances eliminates instance variability
- Comparing instances measures instance variability

Wrapup

- Further reading -- EC2 I/O analysis:
<http://blog.scalyr.com/2012/10/16/a-systematic-look-at-ec2-io/>
- Questions / comments:
steve@scalyr.com



Scalyr Logs

Juggling too many monitoring tools? Not getting full use of your logs? Scalyr Logs is a game-changer. Analyze, graph, aggregate, drill down. Powerful, fast, and easy to use.

[More Details →](#)

[Get Started →](#)