National Laboratories

SAND2019-1949C



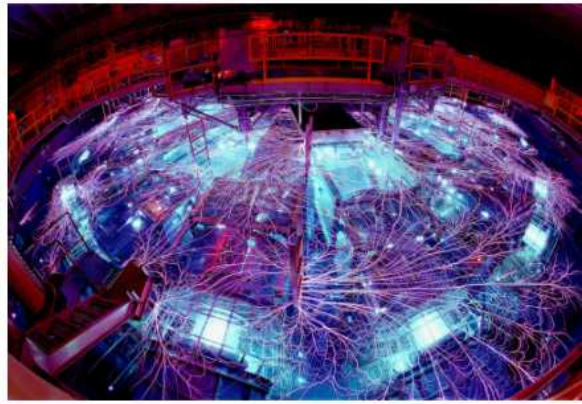

# Virtually the Same: Comparing Physical and Virtual Testbeds

**Jonathan Crussell**, Thomas M Kroeger, Aaron Brown, Cynthia Phillips

February 21st, 2019

U.S. DEPARTMENT OF ENERGY

oratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia,
neywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract

# Motivation

Testing systems is important

- Avoid deploying things that do not work
- Validate new designs
- Understand resiliency

# Motivation

Testing systems is important

- Avoid deploying things that do not work
- Validate new designs
- Understand resiliency

Testbeds are critical to this process

# Motivation

Testing systems is important

- Avoid deploying things that do not work
- Validate new designs
- Understand resiliency

Testbeds are critical to this process

Physical vs virtual testbeds

- Physical are more expensive to build and maintain
- Physical are not easily reconfigurable or portable

# Motivation

Sandia National Laboratories

Testing systems is important

- Avoid deploying things that do not work
- Validate new designs
- Understand resiliency

Testbeds are critical to this process

Physical vs virtual testbeds

- Physical are more expensive to build and maintain
- Physical are not easily reconfigurable or portable

Do artifacts from virtualization affect virtual testbeds?

# Goals

Discover where and how virtual and physical testbeds differ

- Virtualization artifacts
  - Higher network latency and lower throughput [19, 23, 25, 27]

# Goals

Discover where and how virtual and physical testbeds differ

- Virtualization artifacts
  - Higher network latency and lower throughput [19, 23, 25, 27]

Methodology:

- Run representative workloads on physical and virtual testbeds
- Collect, compare, and contrast metrics
  - Application-, OS-, and network-level

# Simple workload

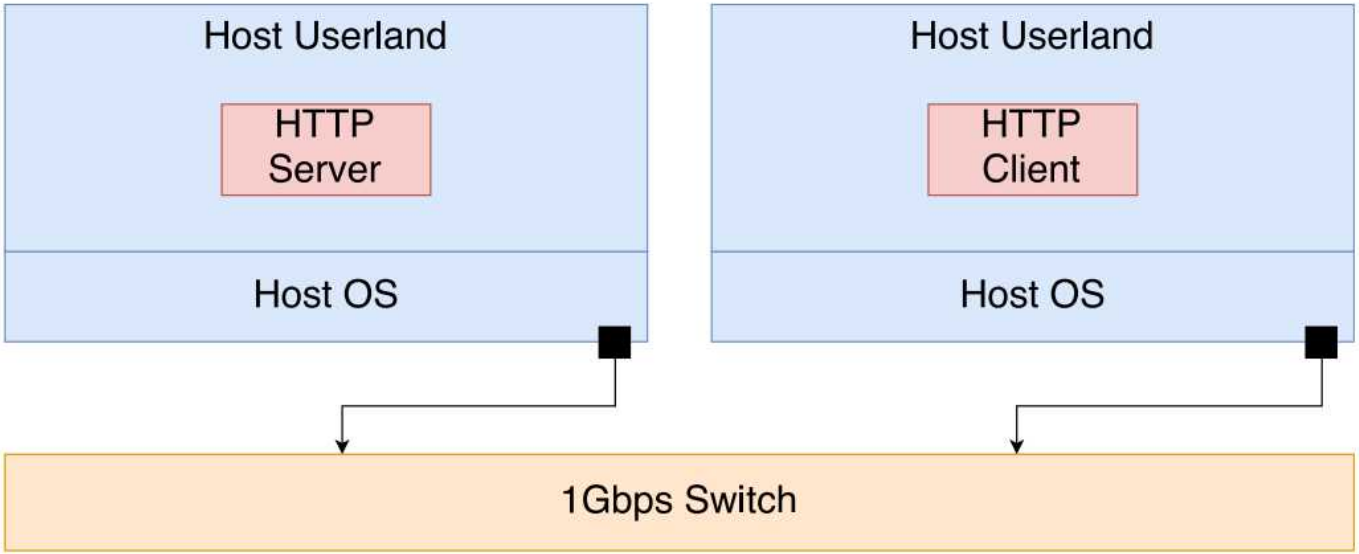

# Simple workload

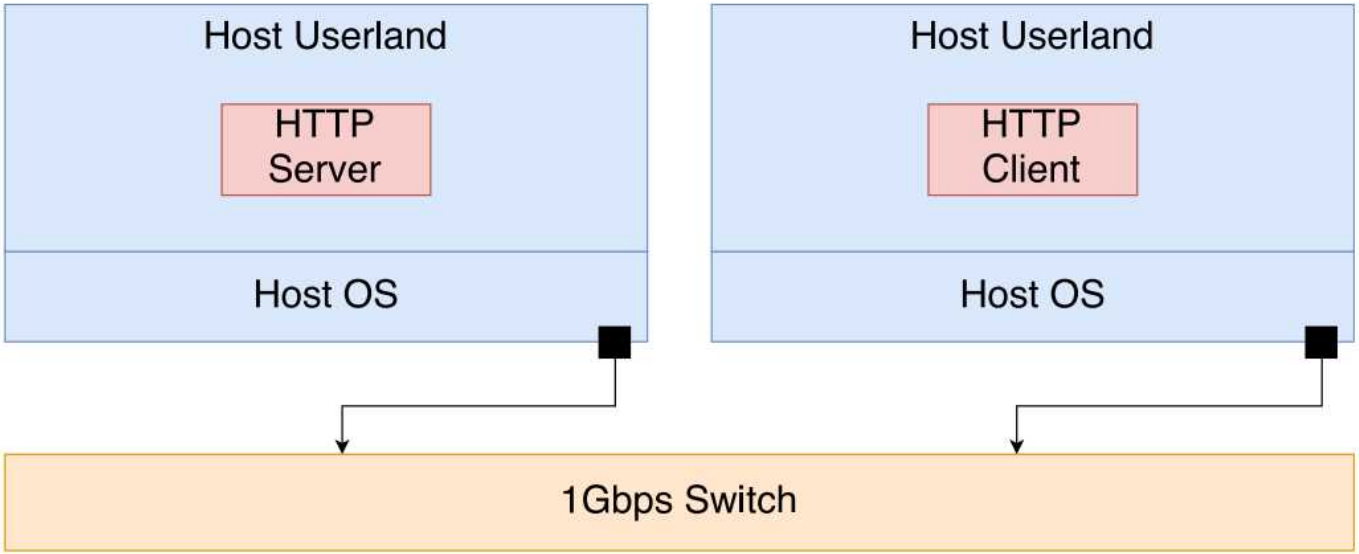


Client makes HTTP requests for 90 seconds

- Server response sizes tested: 500B, 1MB, 16MB

# How to construct virtual testbed?

Many parameters to explore:

# How to construct virtual testbed?

Sandia National Laboratories

Many parameters to explore:

- Virtual machine type?
  - *qemu/kvm*, for now

# How to construct virtual testbed?

Sandia National Laboratories

Many parameters to explore:

- Virtual machine type?
  - *qemu/kvm*, for now
- VM network drivers?
  - *e1000* and *virtio*, for now

# How to construct virtual testbed?

Sandia National Laboratories

Many parameters to explore:

- Virtual machine type?
  - *qemu/kvm*, for now
- VM network drivers?
  - *e1000* and *virtio*, for now
- VM Resources (VCPUs, Memory)?
  - 8 and 2GB, for now

# How to construct virtual testbed?

Sandia National Laboratories

Many parameters to explore:

- Virtual machine type?
  - *qemu/kvm*, for now
- VM network drivers?
  - *e1000* and *virtio*, for now
- VM Resources (VCPUs, Memory)?
  - 8 and 2GB, for now
- Virtual Switch?
  - *Open vSwitch*, for now

# How to construct virtual testbed?

Sandia National Laboratories

Many parameters to explore:

- Virtual machine type?
  - *qemu/kvm*, for now
- VM network drivers?
  - *e1000* and *virtio*, for now
- VM Resources (VCPUs, Memory)?
  - 8 and 2GB, for now
- Virtual Switch?
  - *Open vSwitch*, for now
- … *many more parameters*

# KVM Environment

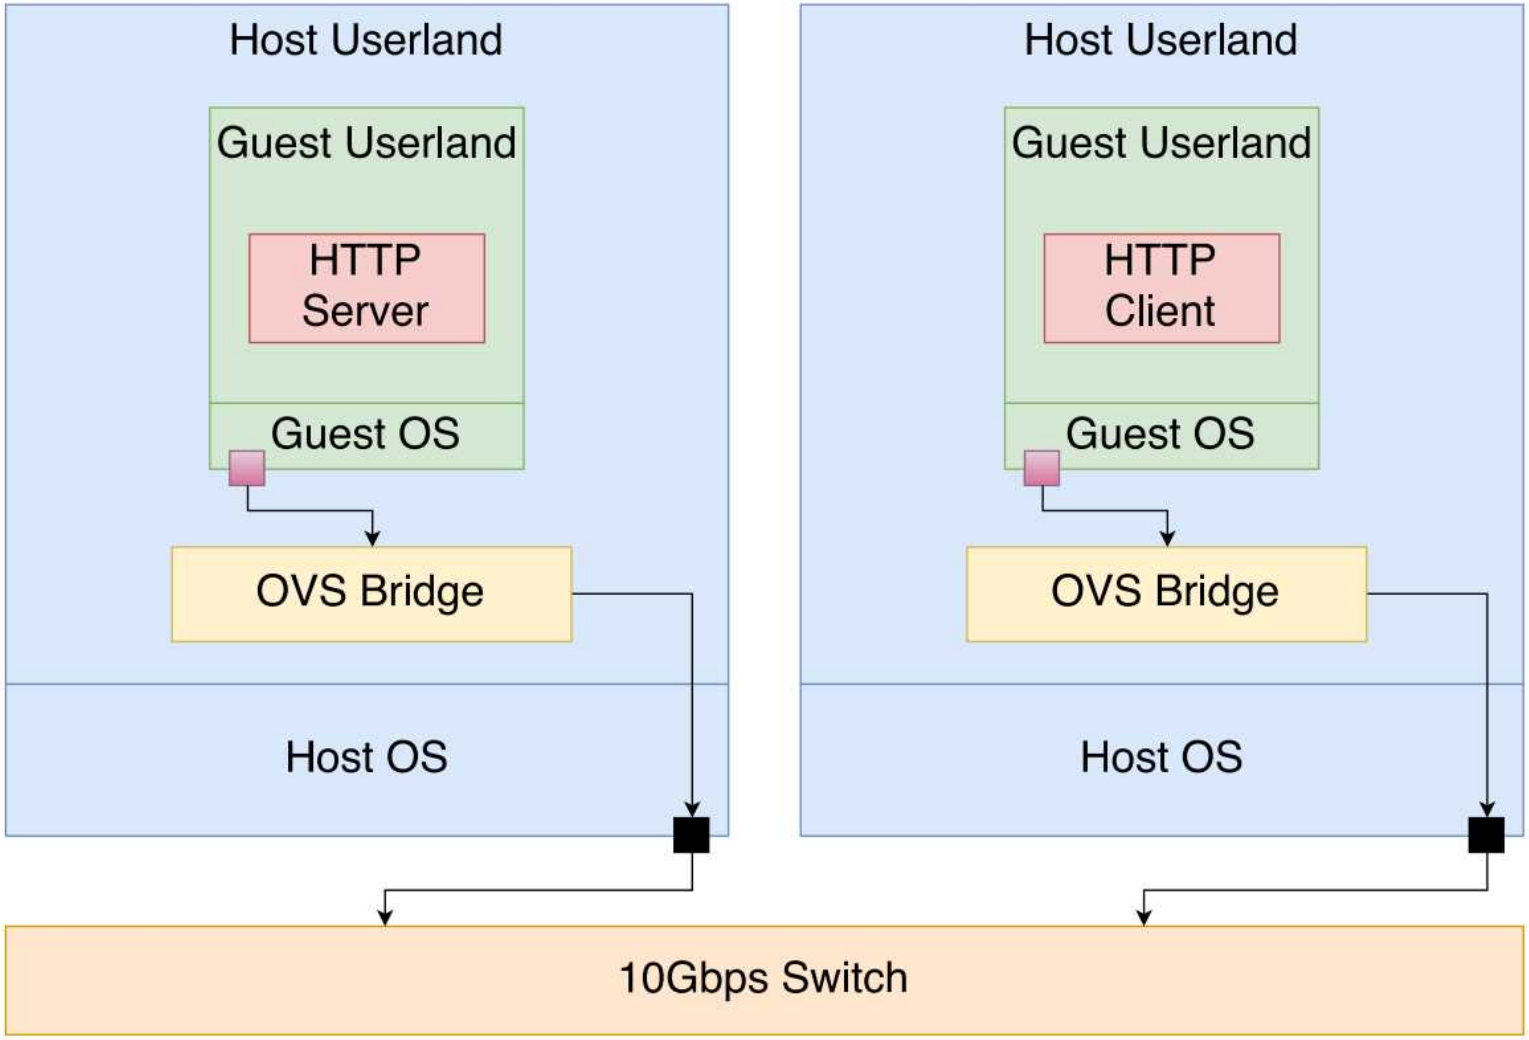

# Instrumentation

Application-level:

- Output from *ApacheBench*

OS-level:

- System-wide system call traces from *sysdig*

Network-level:

- Packet captures, processed with *tcptrace*
- Jitter from one-way latency measurements from *owping*

# Instrumentation

Application-level:

- Output from *ApacheBench*

OS-level:

- System-wide system call traces from *sysdig*

Network-level:

- Packet captures, processed with *tcptrace*
- Jitter from one-way latency measurements from *owping*

Performed experiments with/without instrumentation

- 5-13% decrease in workload performance, varies by workload/testbed

# Experiment Details

Sandia National Laboratories

We use *minimega* to orchestrate and instrument experiments

- Scriptable CLI to run experiments
- Command-and-control to orchestrate VMs
- … *more details at http://minimega.org*

Performed experiments with 1Gbps and 10Gbps interfaces

- See paper for 10Gbps results

# Result: Application-level Metrics

Sandia National Laboratories

| Size | Physical | e1000 | Virtio |
|---|---|---|---|
| 500B | $14420 \pm 74.3$ | $6476 \pm 707$ | $13590 \pm 139$ |
| 1MB | $112 \pm 0.012$ | $113 \pm 0.12$ | $113 \pm 0.006$ |
| 16MB | $6.97 \pm 0.006$ | $7.05 \pm 0.006$ | $7.09 \pm 0.032$ |

Table: Mean requests per second and confidence intervals for ApacheBench runs for 1Gbps tests.

Takeaways:

- e1000 has poor performance (found known bug leading to transmit-queue timeout)
- VMs outperform physical for larger payloads (*tc* parameters?)

# Result: OS-level Metrics

Sandia National Laboratories

| Size | Physical | e1000 | Virtio |
|---|---|---|---|
| 500B | 1.34 | 2.95 | 1.75 |
| 1MB | 547.44 | 395.02 | 325.27 |
| 16MB | 8345.82 | 5954.81 | 5072.02 |

Table: Mean number of read system calls per request for 1Gbps test.

Takeaways:

- Physical requires more reads for same amount of data
- Anecdotally, due to differences in offloading characteristics

# Result: Network-level Metrics

| Size | Physical | e1000 | Virtio |
|---|---|---|---|
| 500B | $5.00 \pm 0.08$ | $5.00 \pm 0.10$ | $5.00 \pm 0.12$ |
| 1MB | $67.7 \pm 2.19$ | $105.7 \pm 9.04$ | $77.3 \pm 10.2$ |
| 16MB | $834 \pm 46.3$ | $1527 \pm 817.40$ | $1087 \pm 706$ |

Table: Mean number of packets per request for 1Gbps test and standard deviation.

Takeaways:

- Higher variability for VMs
- Physical and virtio are similar, e1000 stands out

# Background: Markov Chains

Sandia National Laboratories

Markov chain:

- Graph: nodes represent state, edges represent transitions
- Edges have weights based on transition probability

# Background: Markov Chains

Example:

- User searching for an item in a file

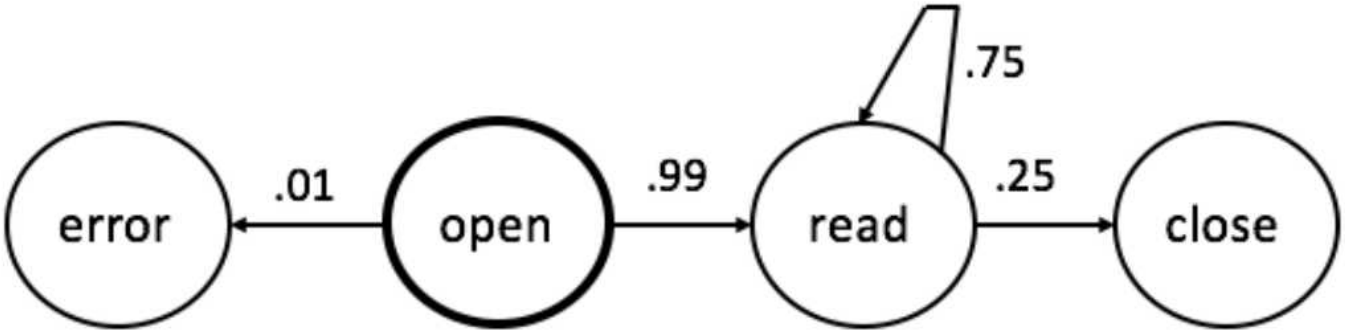

# Markov Chains

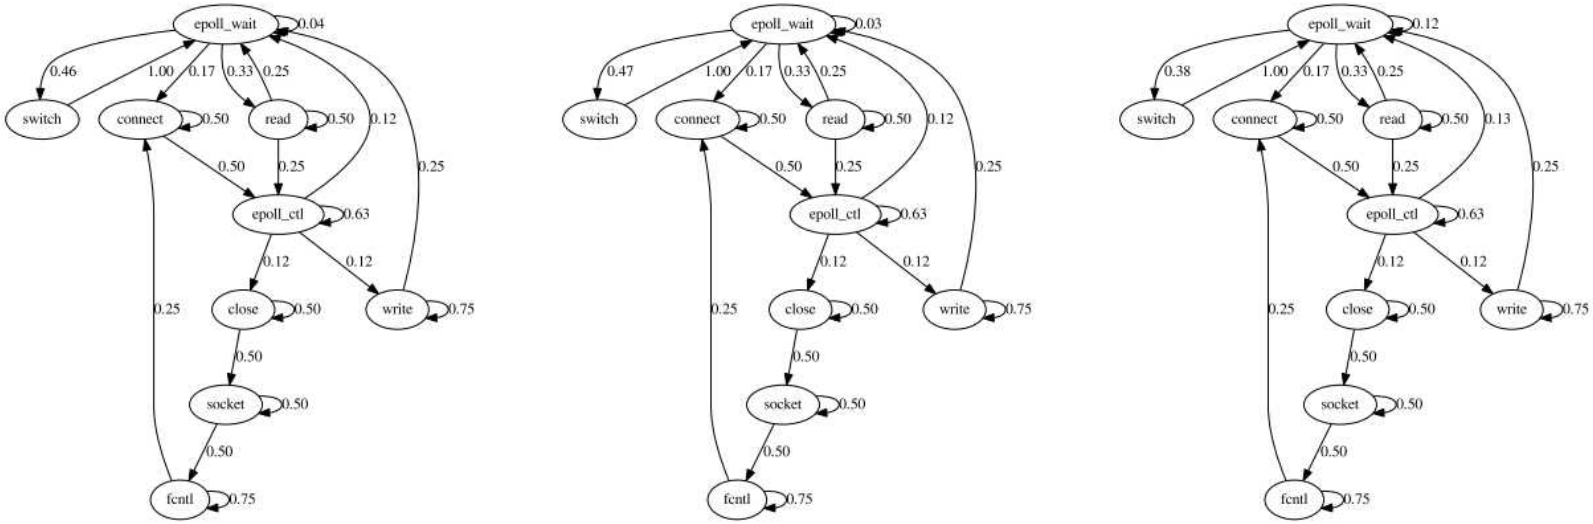


Figure: Client Markov chain for physical, e1000, and virtio. We dropped edges of weight less than $.001$ and renormalized edge weights.

# Conclusion

Sandia National Laboratories

Presented methodology to quantitatively compare physical and virtual testbeds

- Applied to simple HTTP workload
- Showed that virtual testbed behaves reasonably close to its physical counterpart (within 10%)

Questions/Comments?

Presenter: Jonathan Crussell
jcrusse@sandia.gov

# References


- 19: MENON, A., SANTOS, J. R., TURNER, Y., J ANAKIRAMAN, G. J., AND ZWAENEPOEL, W. Diagnosing performance overheads in the xen virtual machine environment. In Proceedings of the 1st ACM/USENIX inter- national conference on Virtual execution environments (2005), ACM, pp. 13–23.
- 23: RIZZO, L., LETTIERI, G., AND MAFFIONE, V. Speeding up packet i/o in virtual machines. In Proceedings of the ninth ACM/IEEE symposium on Architectures for networking and communications systems (2013), IEEE Press, pp. 47–58.
- 25: WANG, G., AND NG, T. E. The impact of virtualization on network performance of amazon ec2 data center. In INFOCOM, 2010 Proceedings IEEE (2010), IEEE, pp. 1–9.
- 27: WHITEAKER, J., SCHNEIDER, F., AND TEIXEIRA, R. Explaining packet delays under virtualization. ACM SIGCOMM Computer Com- munication Review 41, 1 (2011), 38–44.