SAND2017-8992C

Virtually the Same? The Empirical Differences Between Physical and Virtual Networks

Jonathan Crussell, Tom Kroeger, Aaron Brown, Cindy Phillips

Sandia National Laboratories, California

August 11th, 2017

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

Quantifying Uncertainty in Emulations (que)

Sandia makes critical use of Emulytics^{TM}

- "Field test" deployments
- Advise procurements

Quantifying Uncertainty in Emulations (que)

Sandia makes critical use of Emulytics^{TM}

- "Field test" deployments
- Advise procurements

LDRD Goals:

- Discover where and how Emulytics differs from real world
  - Both in quantity and nature
  - Scoped by mission objectives

Quantifying Uncertainty in Emulations (que)

Sandia makes critical use of Emulytics^{TM}

- "Field test" deployments
- Advise procurements

LDRD Goals:

- Discover where and how Emulytics differs from real world
  - Both in quantity and nature
  - Scoped by mission objectives
- Use knowledge base to improve state-of-the-art:
  - Emulytics experimenter's handbook
  - Underpinnings of calibration tools

Quantifying Uncertainty in Emulations (que)

Sandia makes critical use of Emulytics^{TM}

- "Field test" deployments
- Advise procurements

LDRD Goals:

- Discover where and how Emulytics differs from real world
  - Both in quantity and nature
  - Scoped by mission objectives
- Use knowledge base to improve state-of-the-art:
  - Emulytics experimenter's handbook
  - Underpinnings of calibration tools

Methodology:

- Run representative workloads in both real world and emulations
- Collect, compare, and contrast metrics

Starting simple

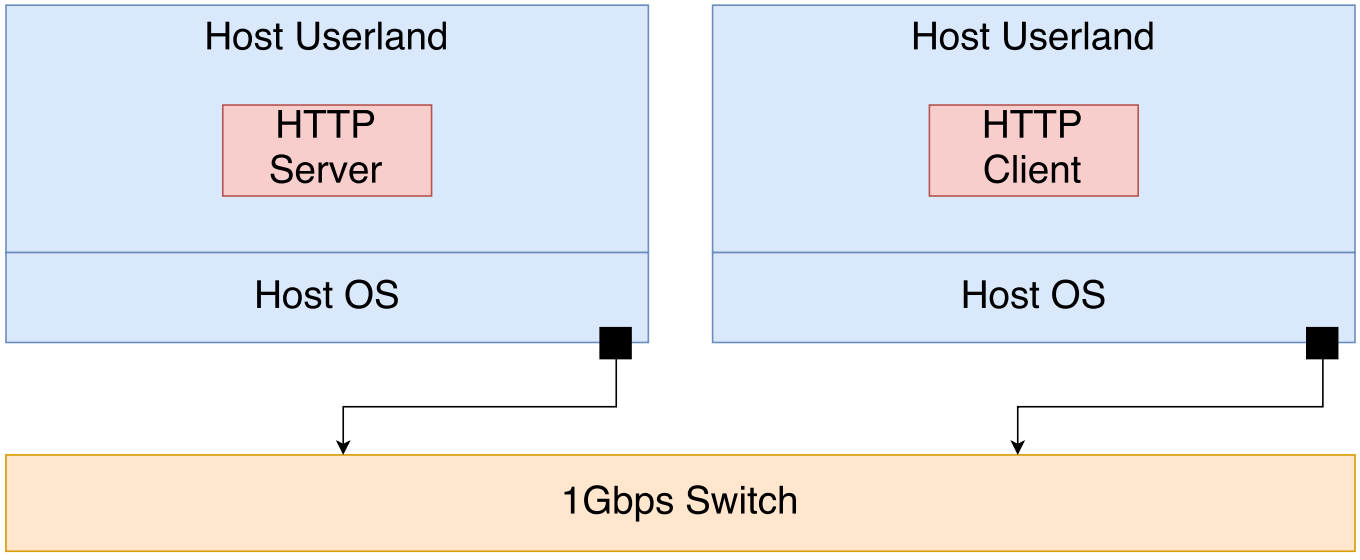

Starting simple

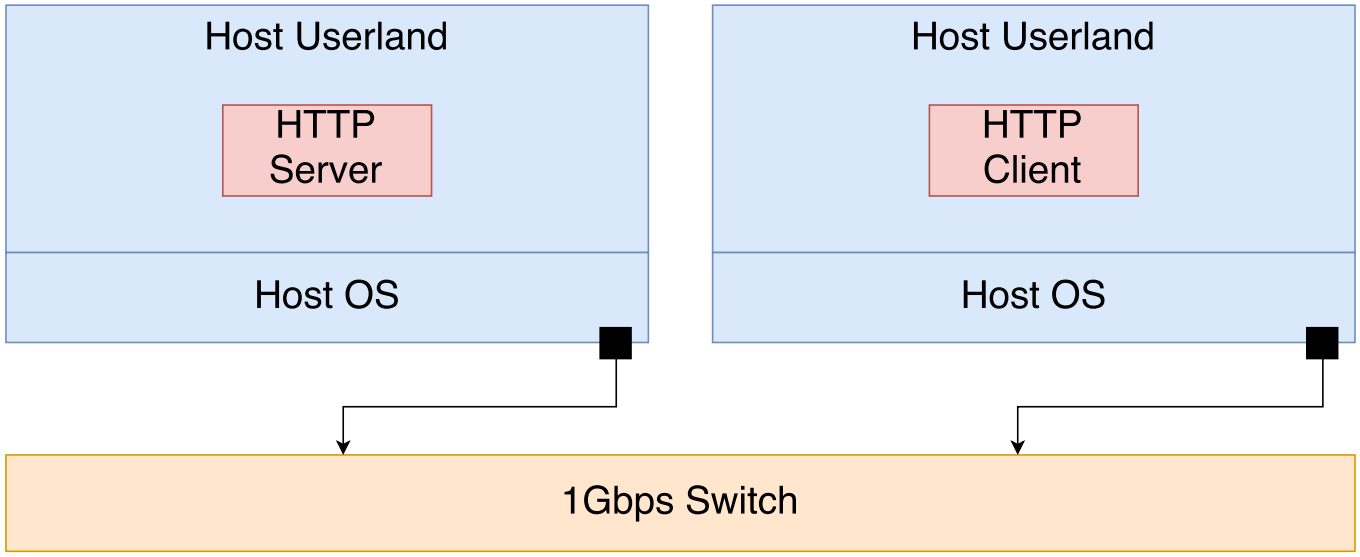

Client makes HTTP requests for 60 seconds

How to Emulate?

- VMs or Containers?

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?
- How many VCPUs to emulate?

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?
- How many VCPUs to emulate?
  - 8, for now

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?
- How many VCPUs to emulate?
  - 8, for now
- How many worker threads?

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?
- How many VCPUs to emulate?
  - 8, for now
- How many worker threads?
  - 1, 4, and 16, for now

How to Emulate?

- VMs or Containers?
- For KVM-based VMs, which network drivers?
  - e1000 and virtio, for now
  - *... many more*
- Should we disable offloading?
- How many VCPUs to emulate?
  - 8, for now
- How many worker threads?
  - 1, 4, and 16, for now
- *... many more parameters*

KVM Environment

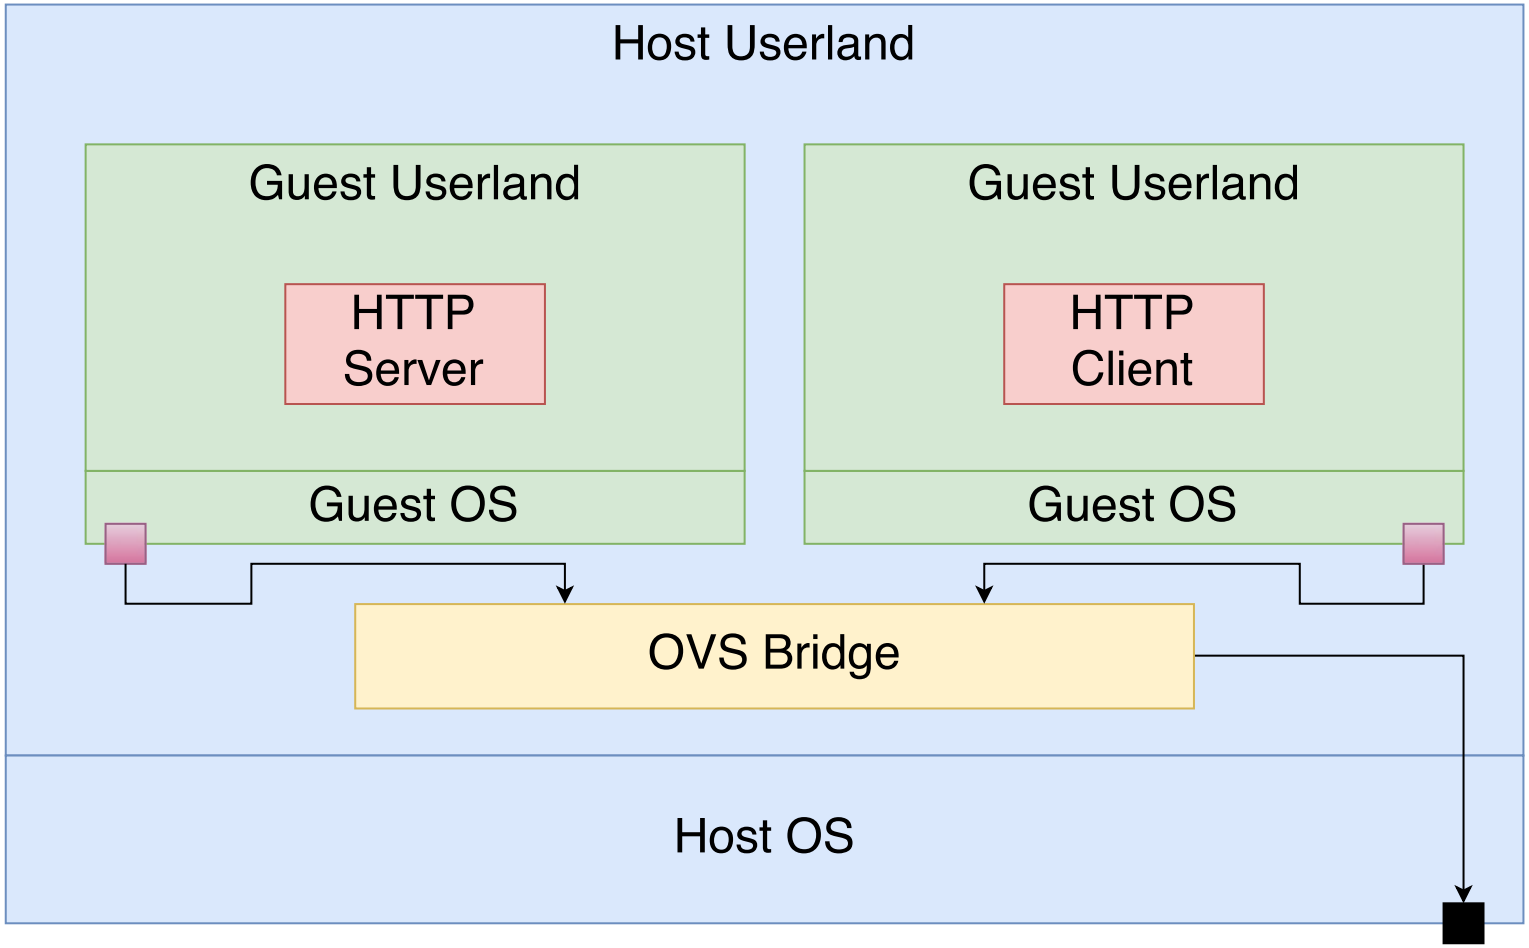

Container Environment

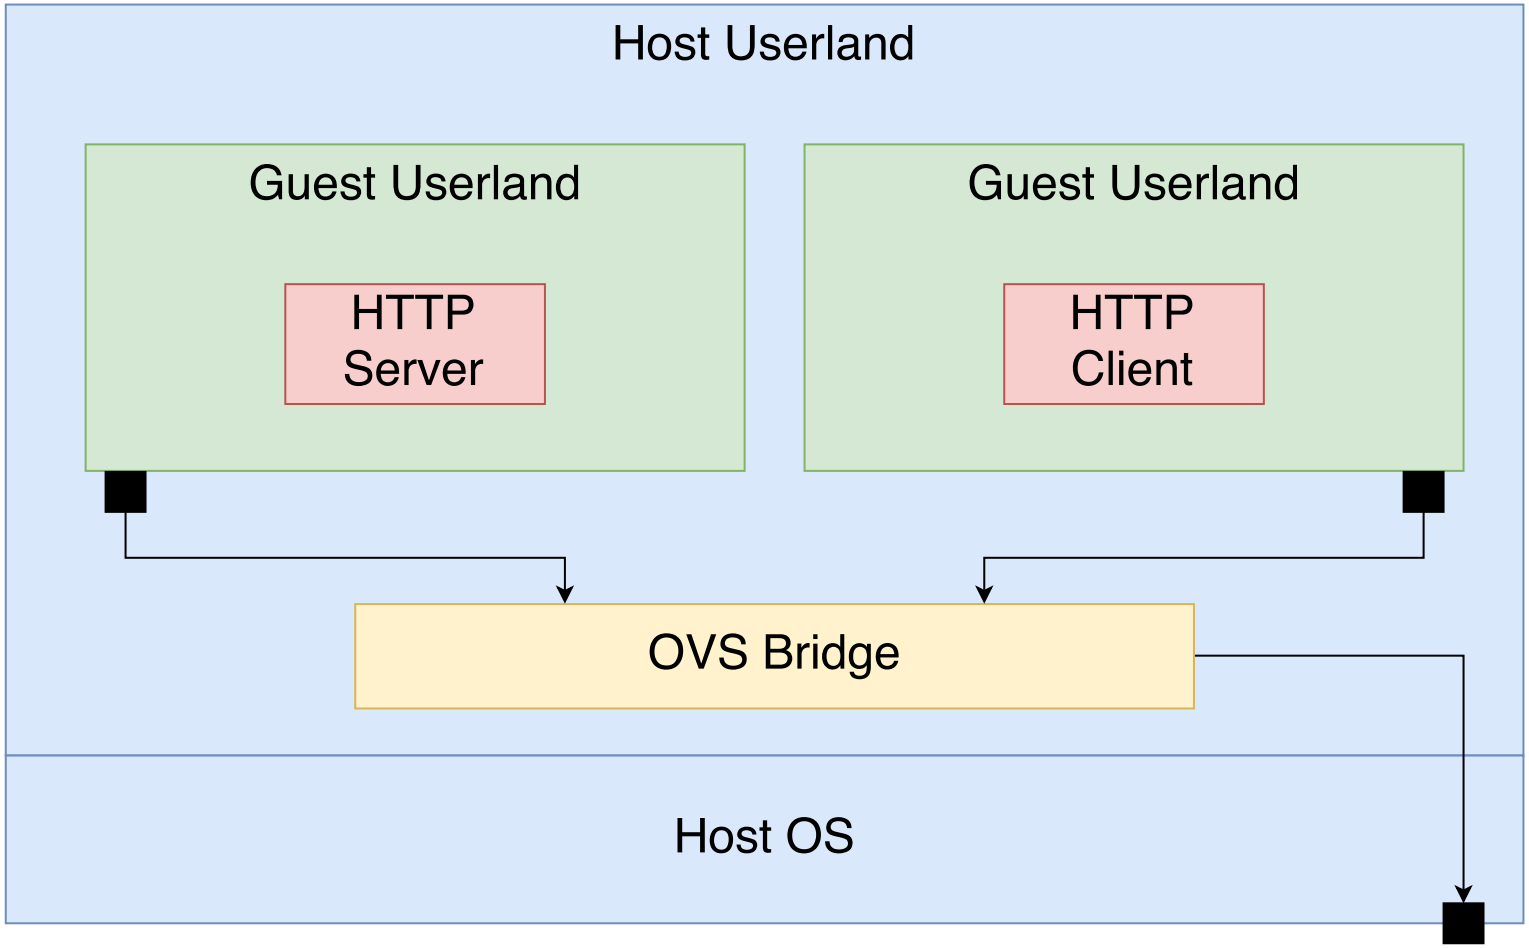

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	282 ± 6	549 ± 11	390 ± 20
virtio	N	916 ± 0	920 ± 1	934 ± 2
containers	N	888 ± 11	918 ± 4	931 ± 6
e1000	Y	914 ± 1	918 ± 1	929 ± 4
virtio	Y	916 ± 0	921 ± 1	933 ± 2
containers	Y	916 ± 0	920 ± 1	934 ± 2

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	282 ± 6	549 ± 11	390 ± 20
virtio	N	916 ± 0	920 ± 1	934 ± 2
containers	N	888 ± 11	918 ± 4	931 ± 6
e1000	Y	914 ± 1	918 ± 1	929 ± 4
virtio	Y	916 ± 0	921 ± 1	933 ± 2
containers	Y	916 ± 0	920 ± 1	934 ± 2

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	282 ± 6	549 ± 11	390 ± 20
virtio	N	916 ± 0	920 ± 1	934 ± 2
containers	N	888 ± 11	918 ± 4	931 ± 6
e1000	Y	914 ± 1	918 ± 1	929 ± 4
virtio	Y	916 ± 0	921 ± 1	933 ± 2
containers	Y	916 ± 0	920 ± 1	934 ± 2

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	282 ± 6	549 ± 11	390 ± 20
virtio	N	916 ± 0	920 ± 1	934 ± 2
containers	N	888 ± 11	918 ± 4	931 ± 6
e1000	Y	914 ± 1	918 ± 1	929 ± 4
virtio	Y	916 ± 0	921 ± 1	933 ± 2
containers	Y	916 ± 0	920 ± 1	934 ± 2

Multi-host environments

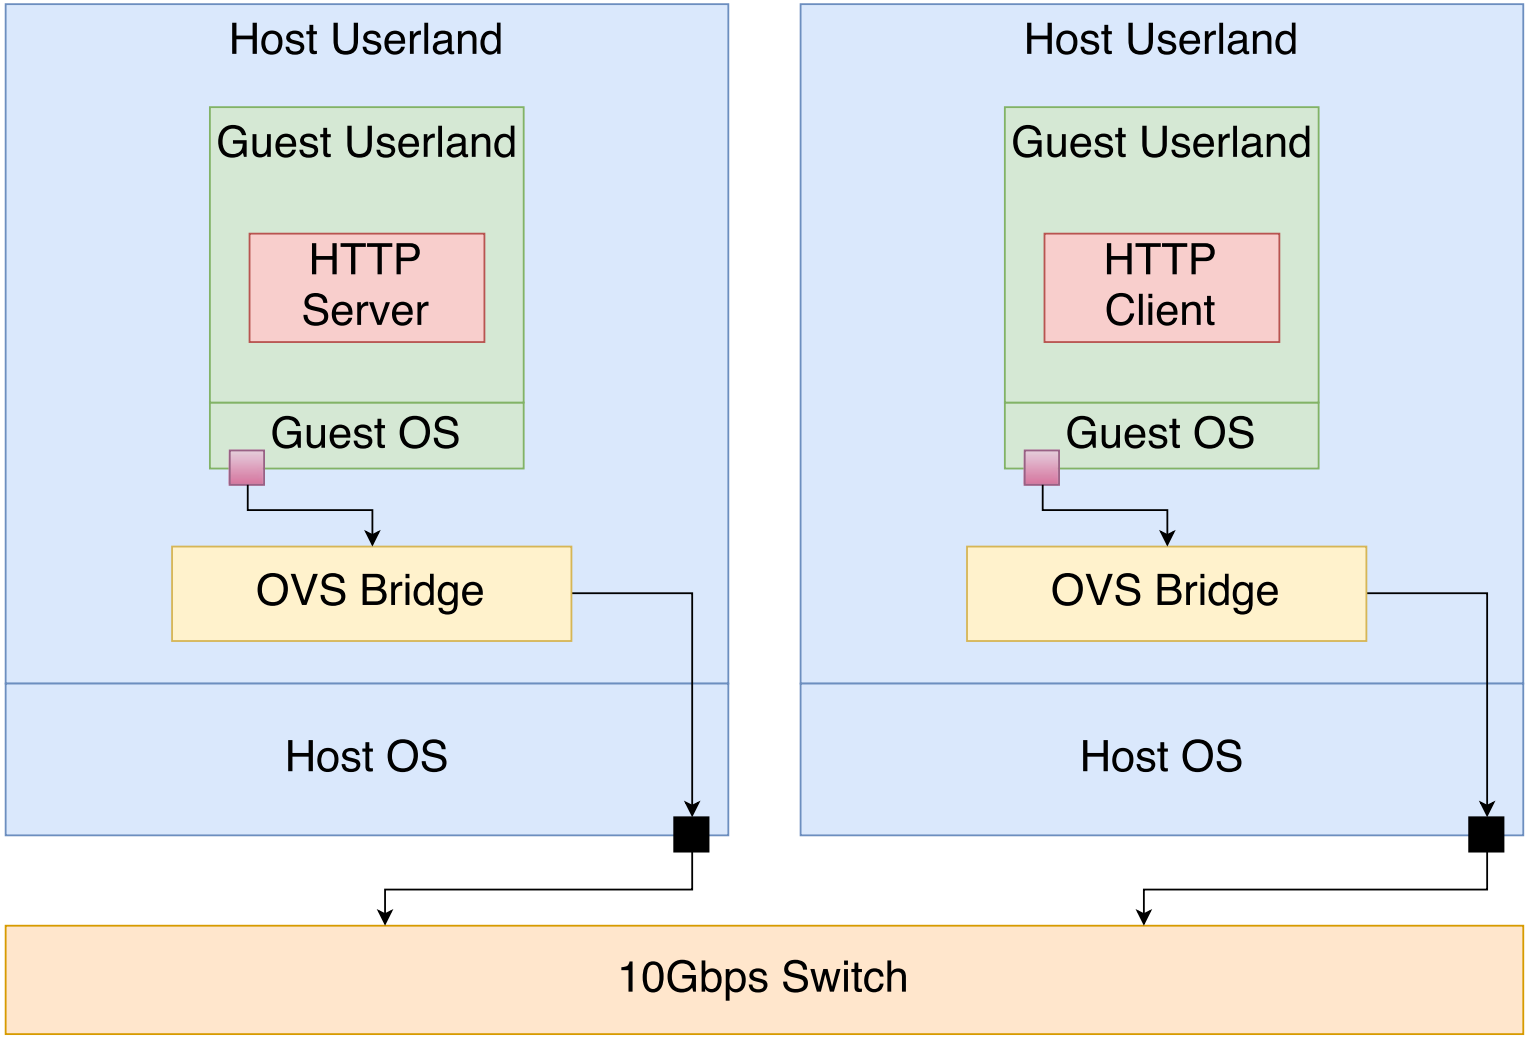

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	259 ± 13	429 ± 14	358 ± 15
virtio	N	916 ± 0	919 ± 1	931 ± 2
containers	N	916 ± 0	920 ± 1	931 ± 2
e1000	Y	903 ± 3	917 ± 1	890 ± 46
virtio	Y	915 ± 1	920 ± 1	932 ± 2
containers	Y	916 ± 0	921 ± 1	932 ± 4

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	259 ± 13	429 ± 14	358 ± 15
virtio	N	916 ± 0	919 ± 1	931 ± 2
containers	N	916 ± 0	920 ± 1	931 ± 2
e1000	Y	903 ± 3	917 ± 1	890 ± 46
virtio	Y	915 ± 1	920 ± 1	932 ± 2
containers	Y	916 ± 0	921 ± 1	932 ± 4

Number of Transactions

		Transactions		
Environment	Offloading	1 Worker	4 Workers	16 Workers
physical	N	883 ± 0	906 ± 1	919 ± 2
physical	Y	879 ± 0	906 ± 1	918 ± 2
e1000	N	259 ± 13	429 ± 14	358 ± 15
virtio	N	916 ± 0	919 ± 1	931 ± 2
containers	N	916 ± 0	920 ± 1	931 ± 2
e1000	Y	903 ± 3	917 ± 1	890 ± 46
virtio	Y	915 ± 1	920 ± 1	932 ± 2
containers	Y	916 ± 0	921 ± 1	932 ± 4

Anatomy of an HTTP request

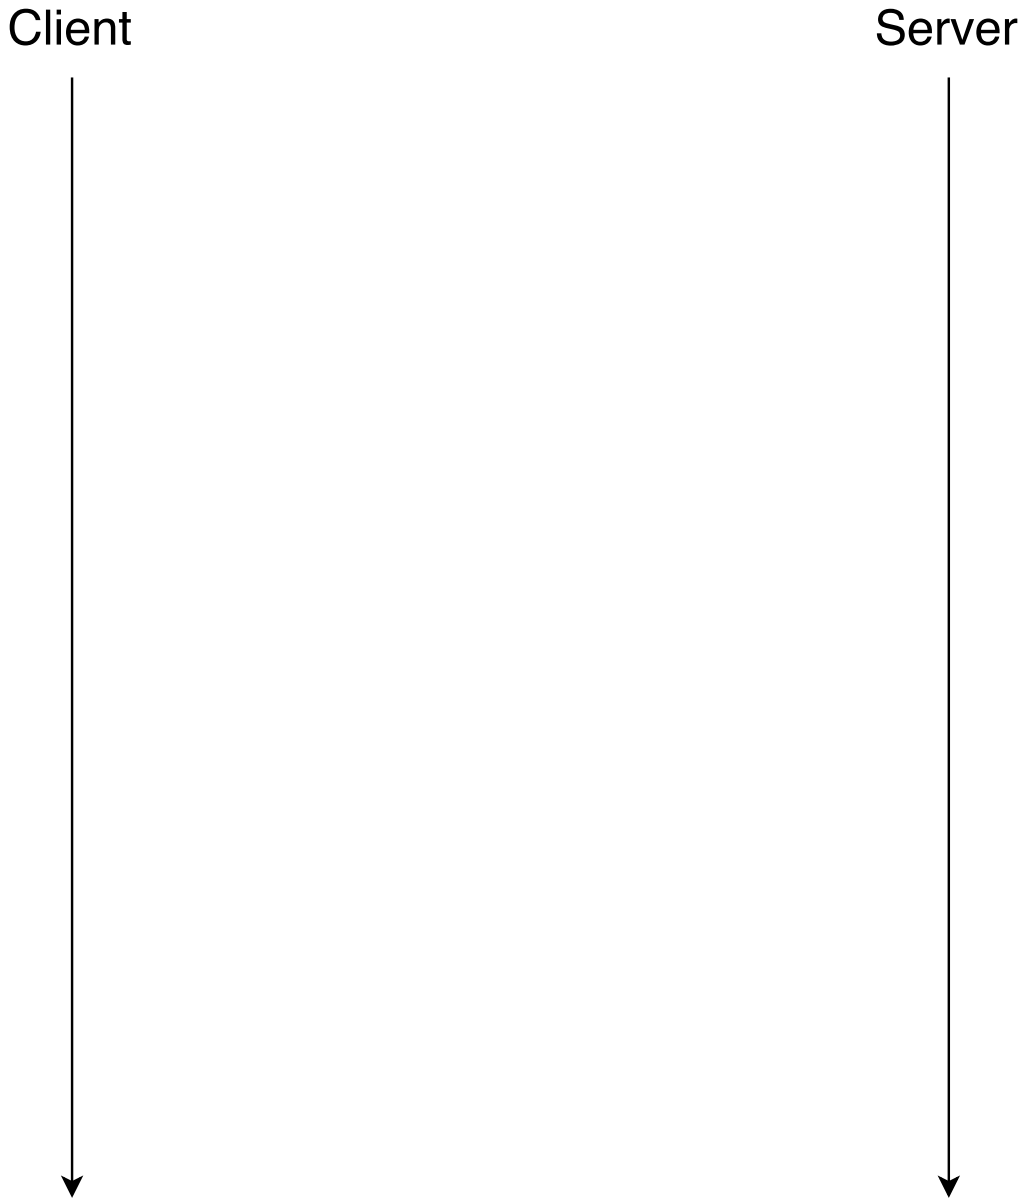

Anatomy of an HTTP request

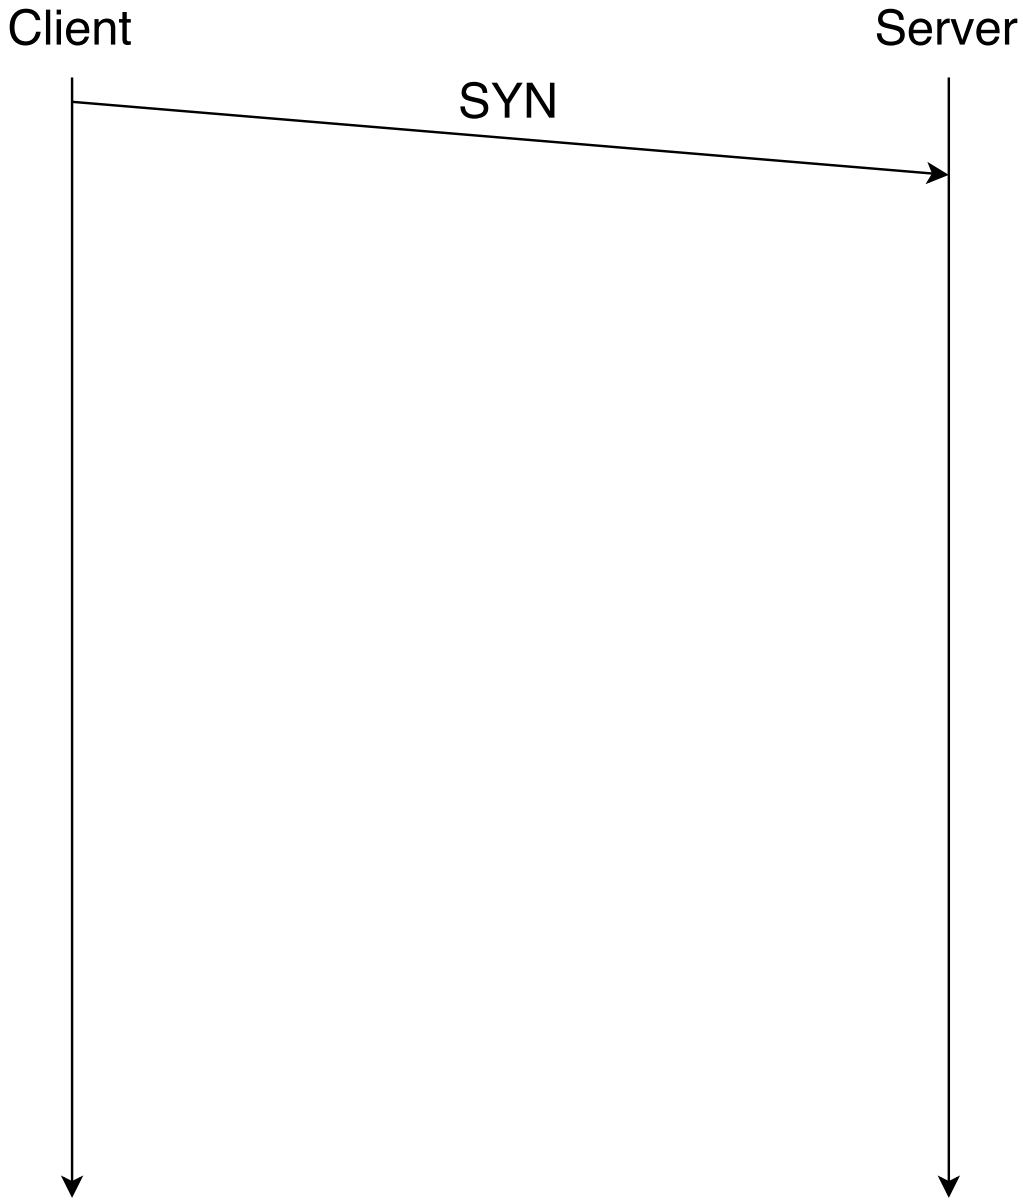

Anatomy of an HTTP request

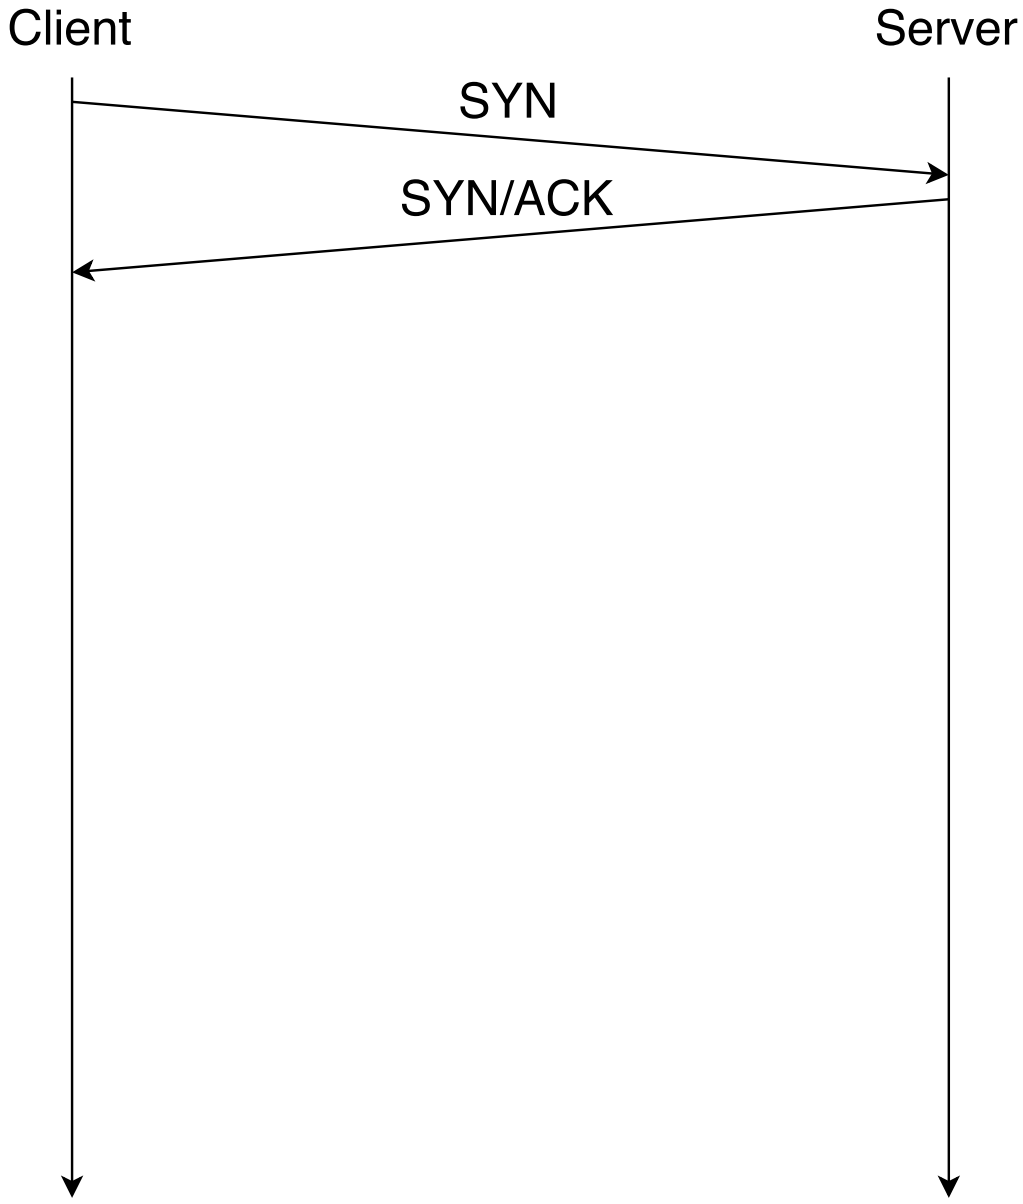

Anatomy of an HTTP request

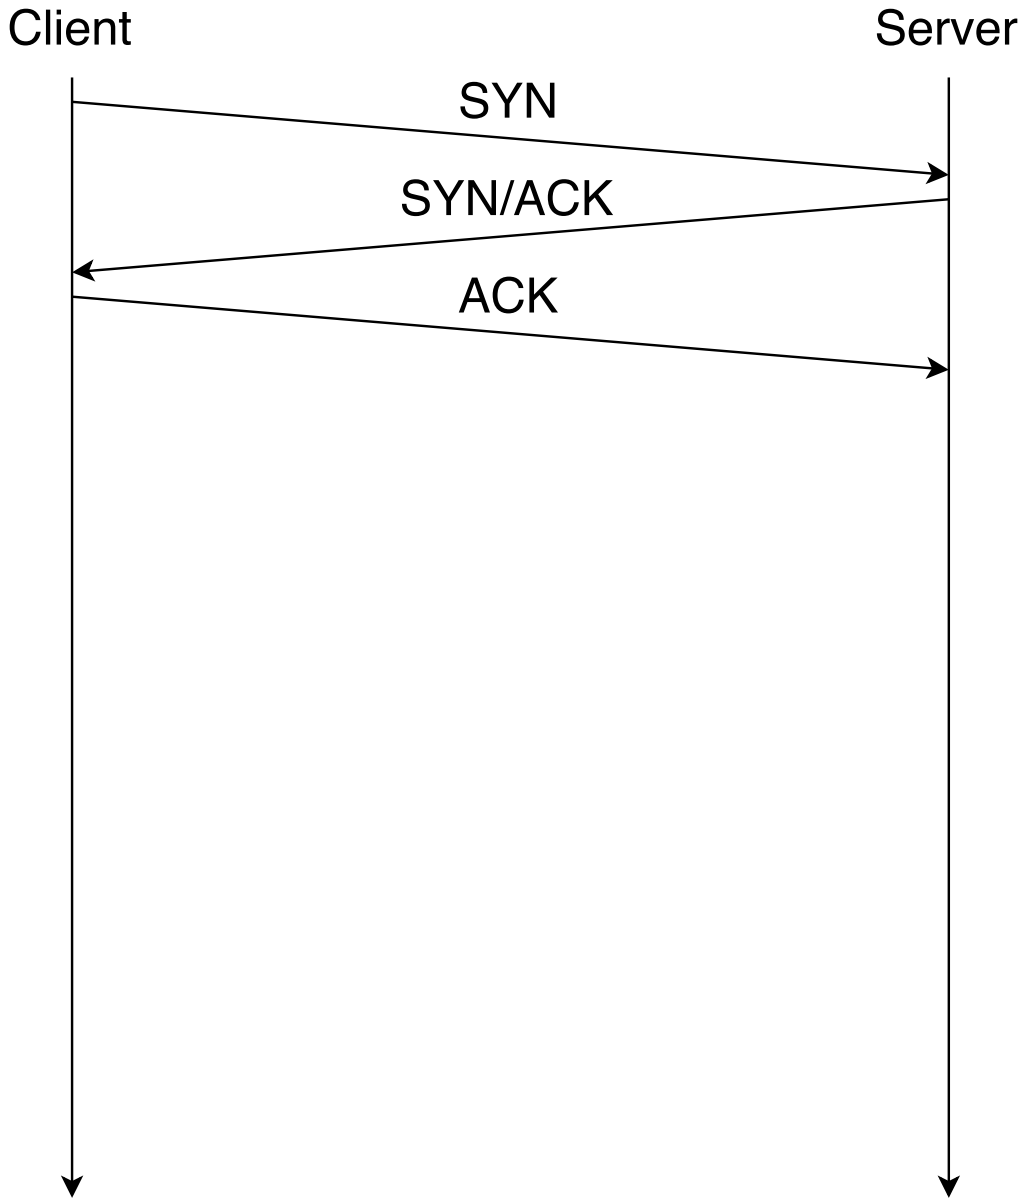

Anatomy of an HTTP request

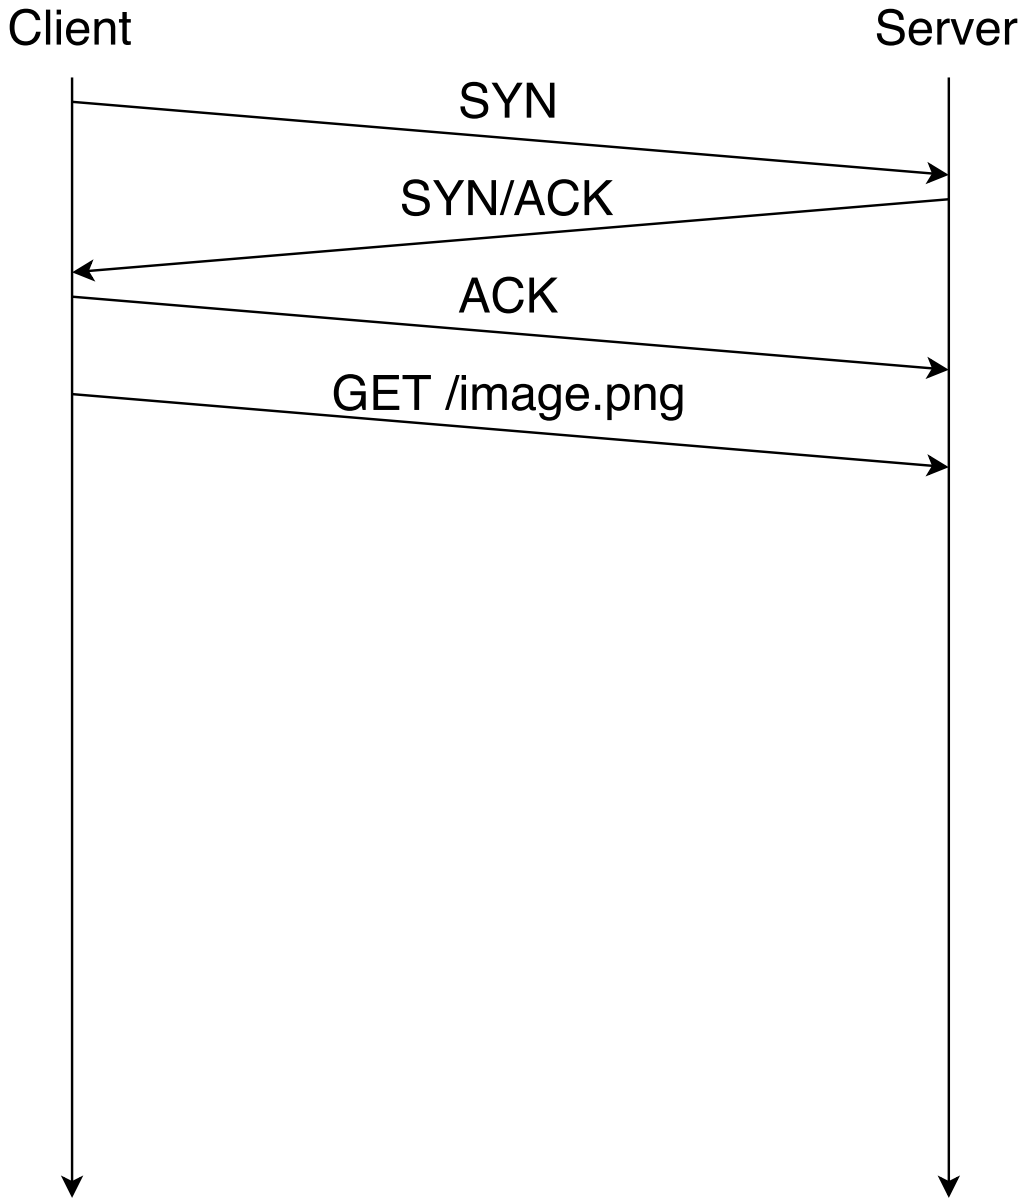

Anatomy of an HTTP request

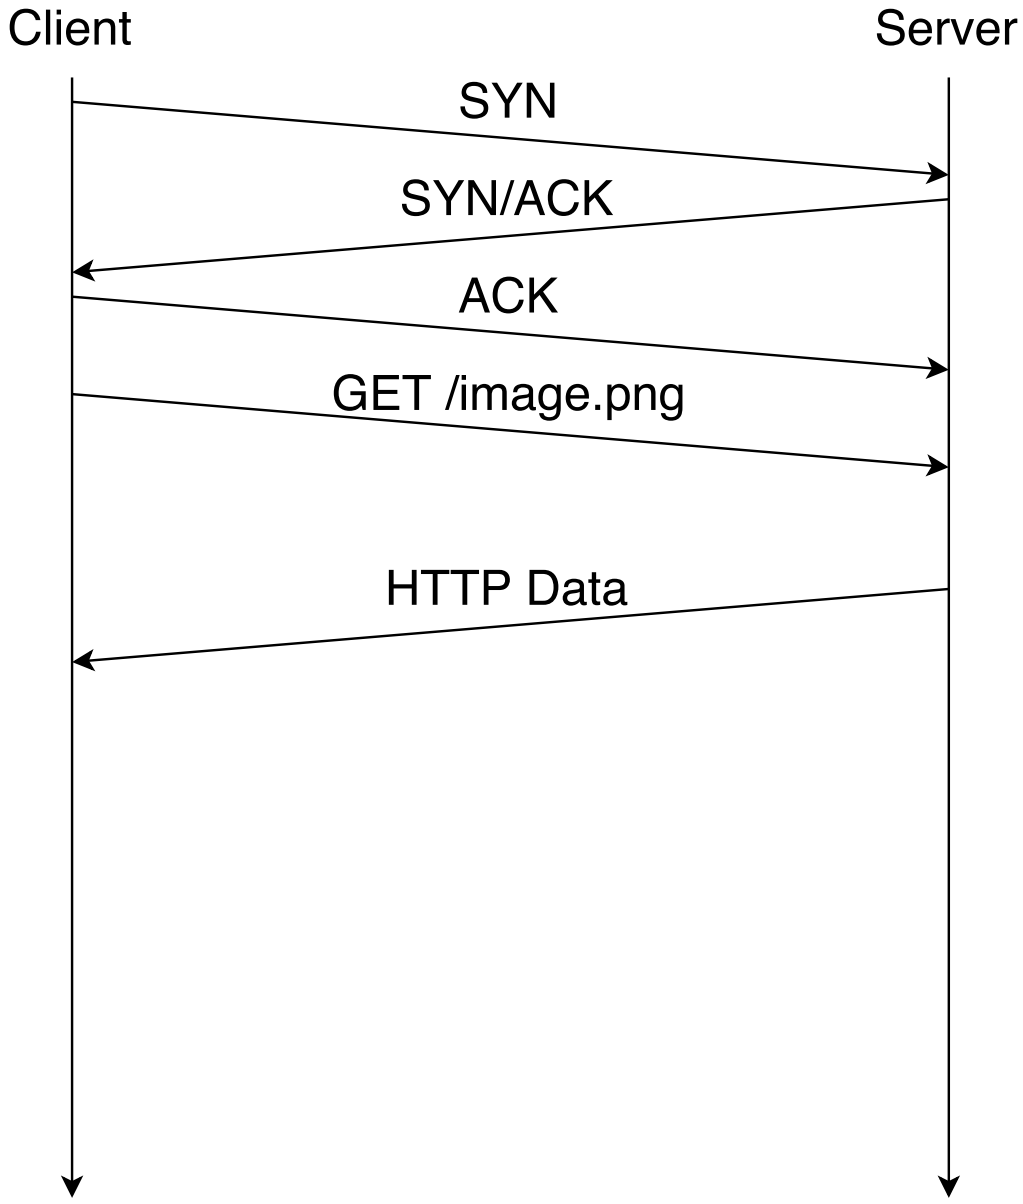

Anatomy of an HTTP request

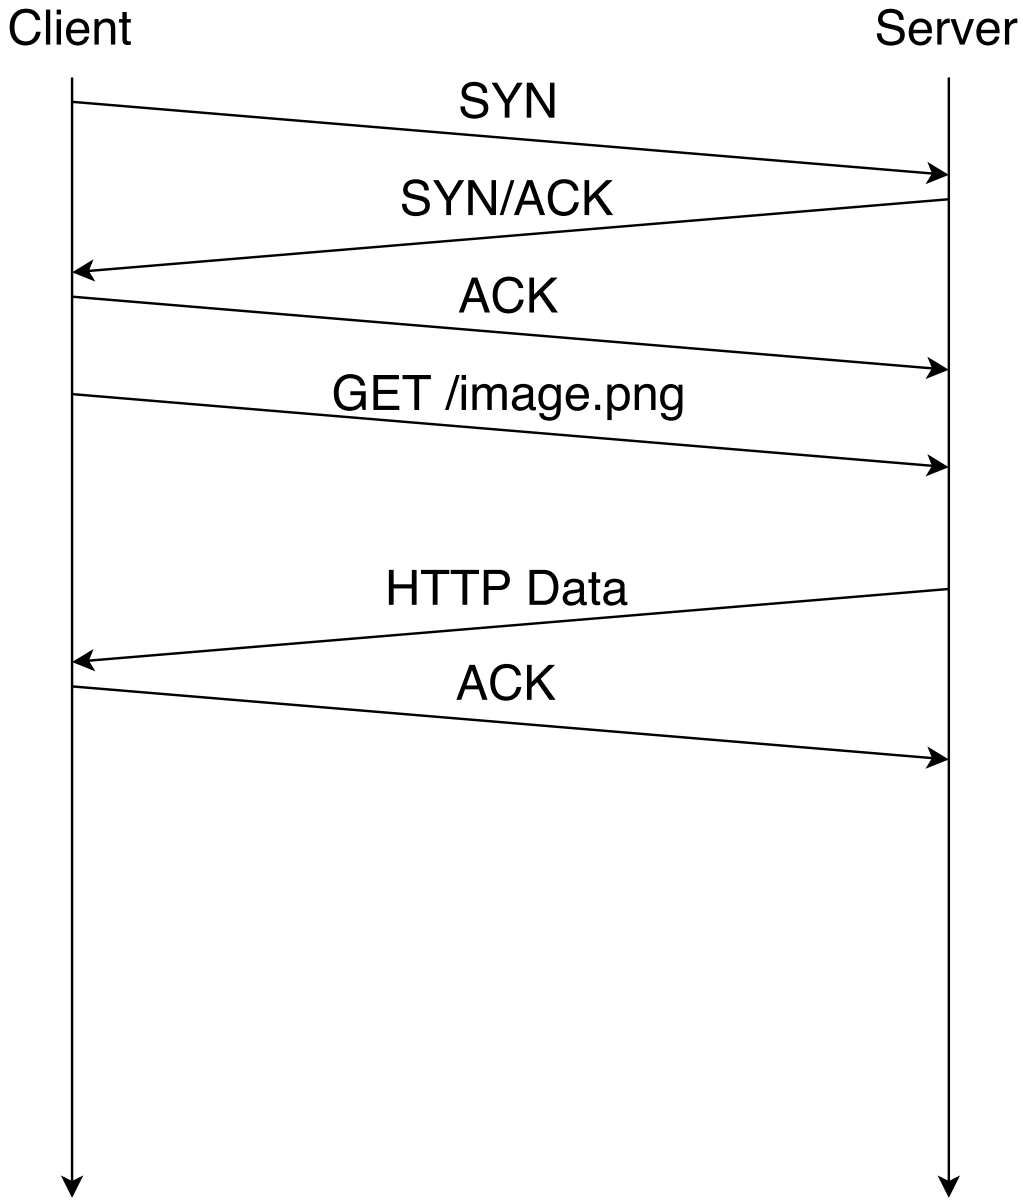

Anatomy of an HTTP request

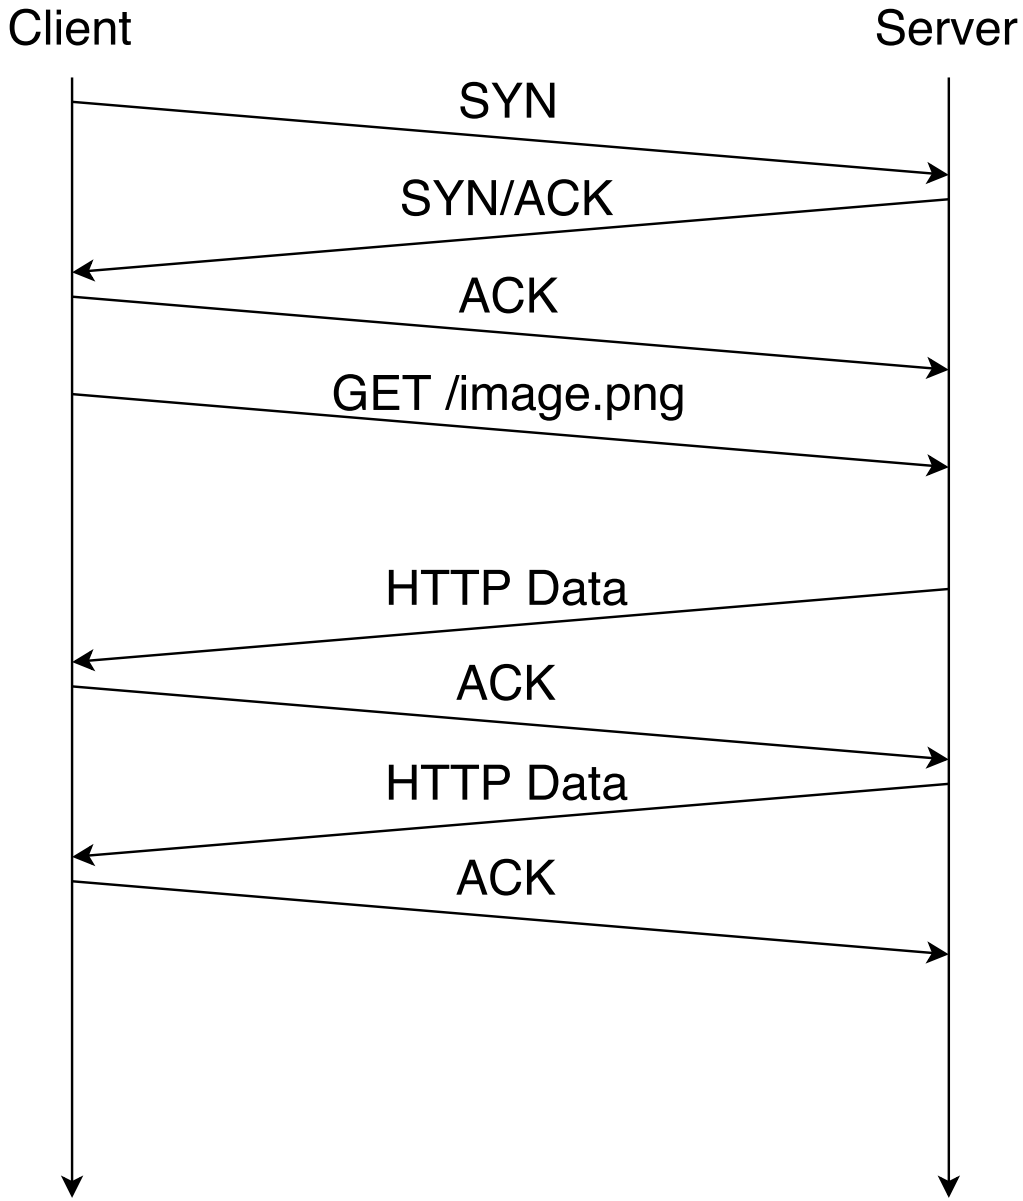

Anatomy of an HTTP request

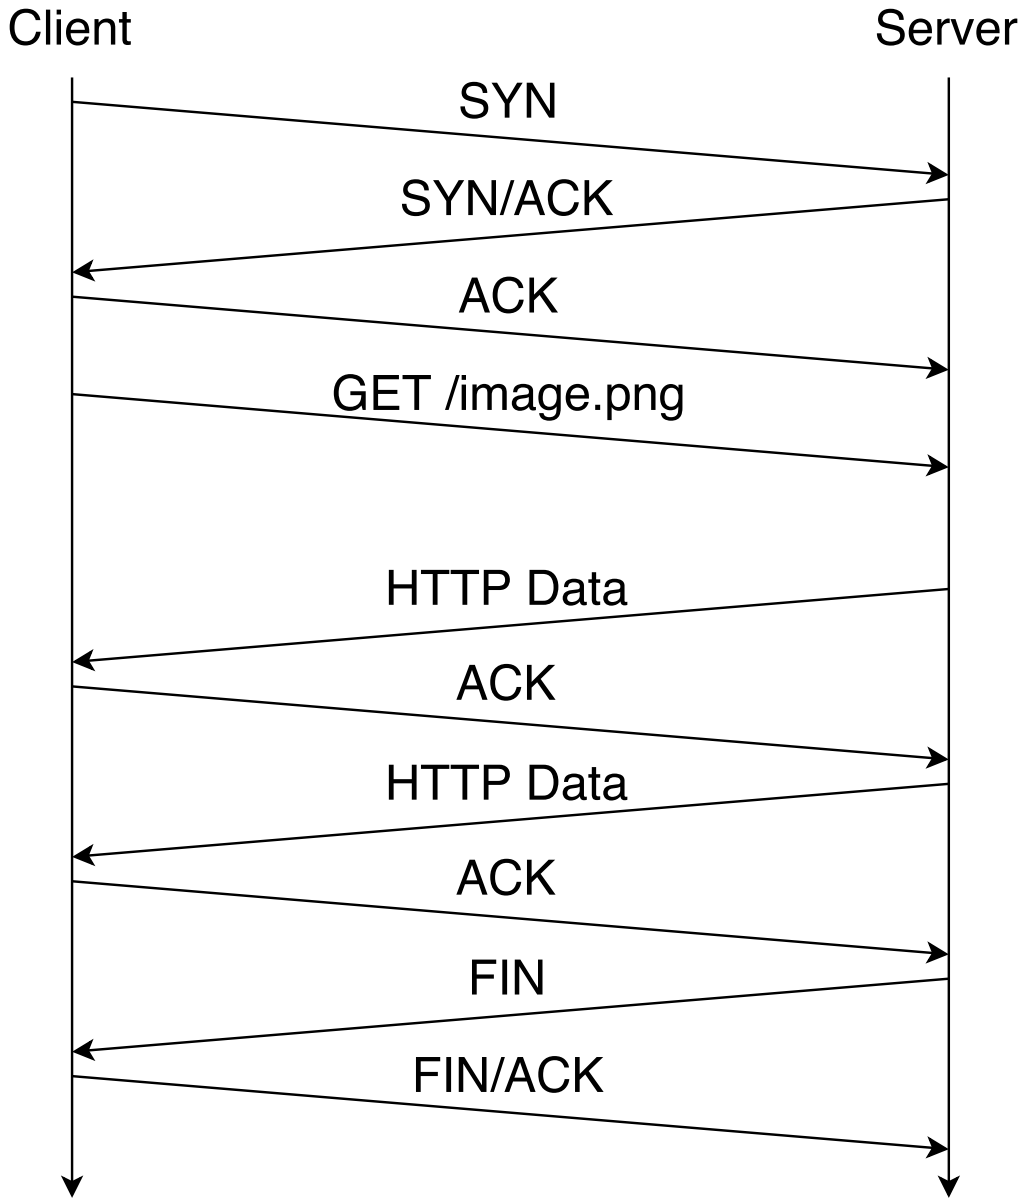

Anatomy of an HTTP request

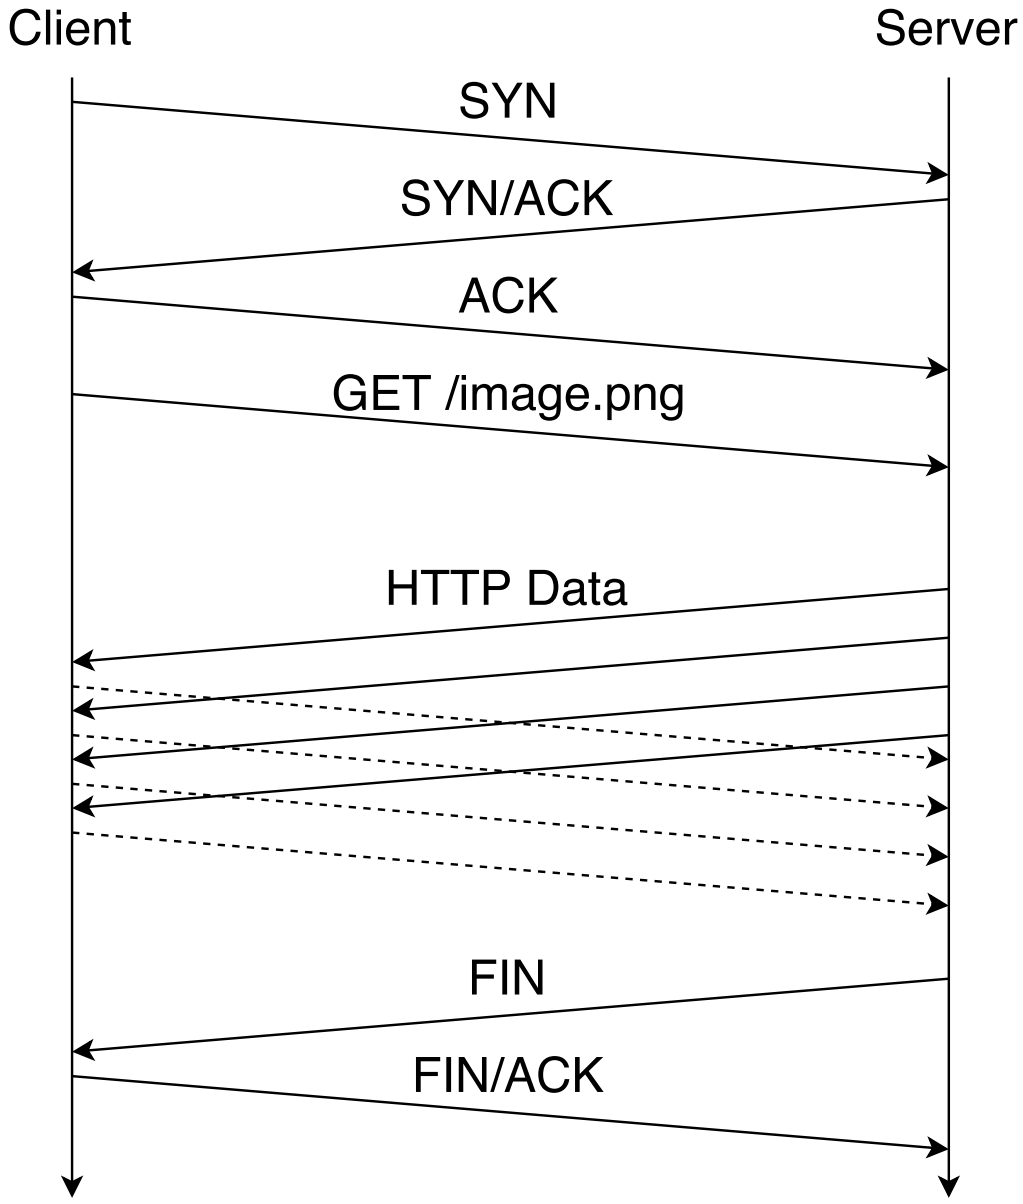

Anatomy of an HTTP request – Reordering

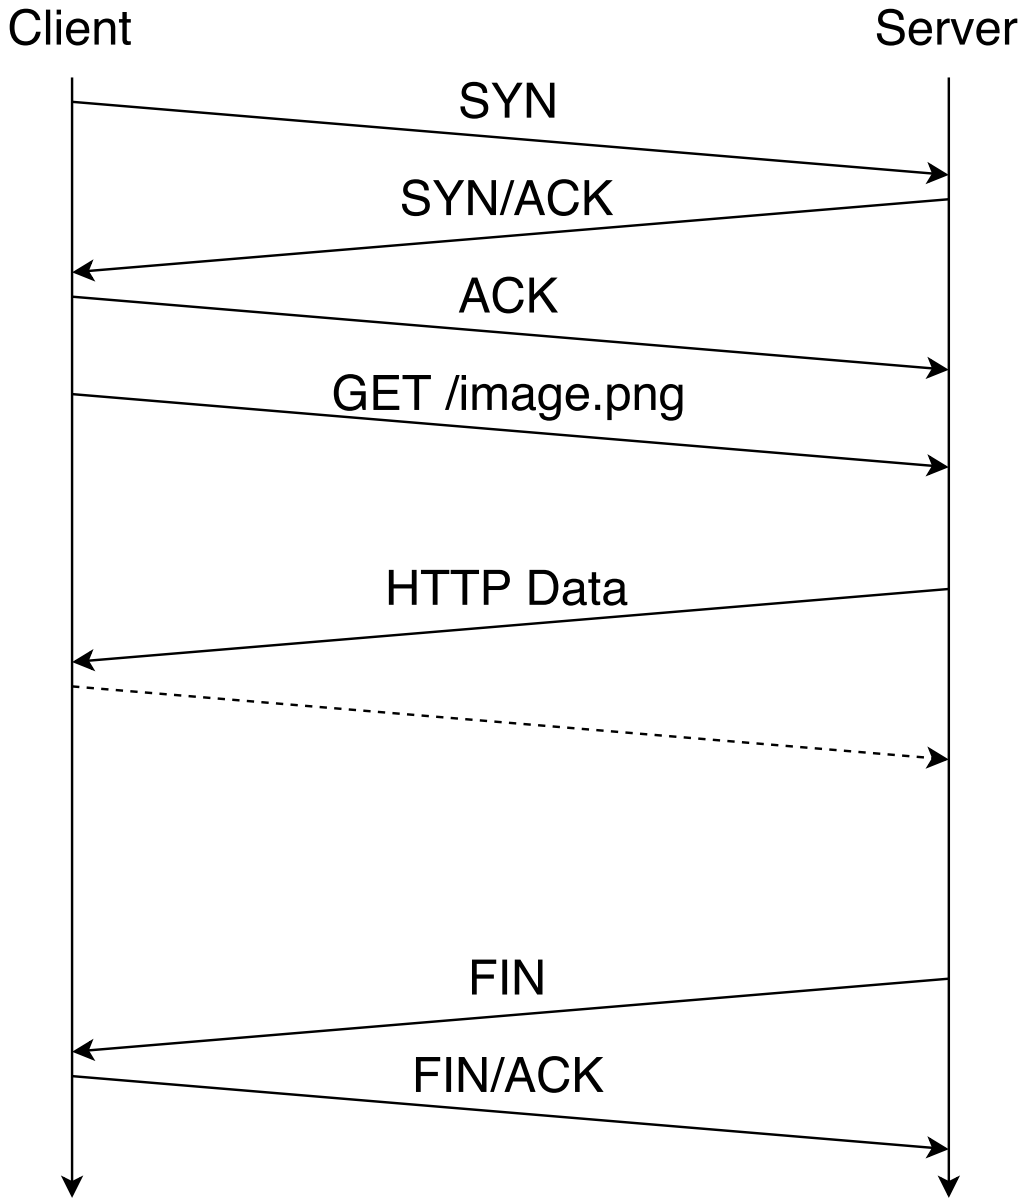

Anatomy of an HTTP request – Reordering

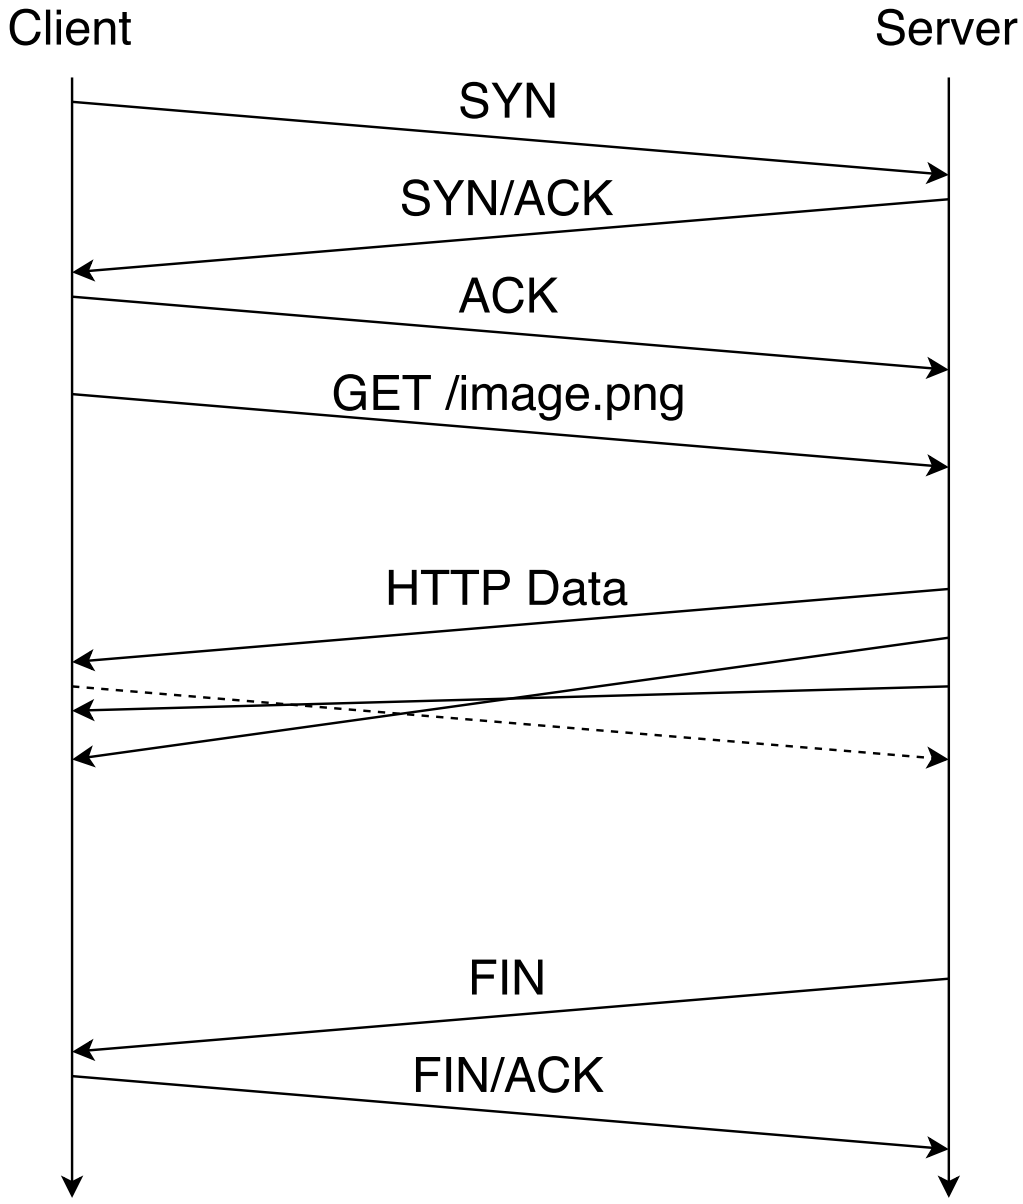

Anatomy of an HTTP request – Reordering

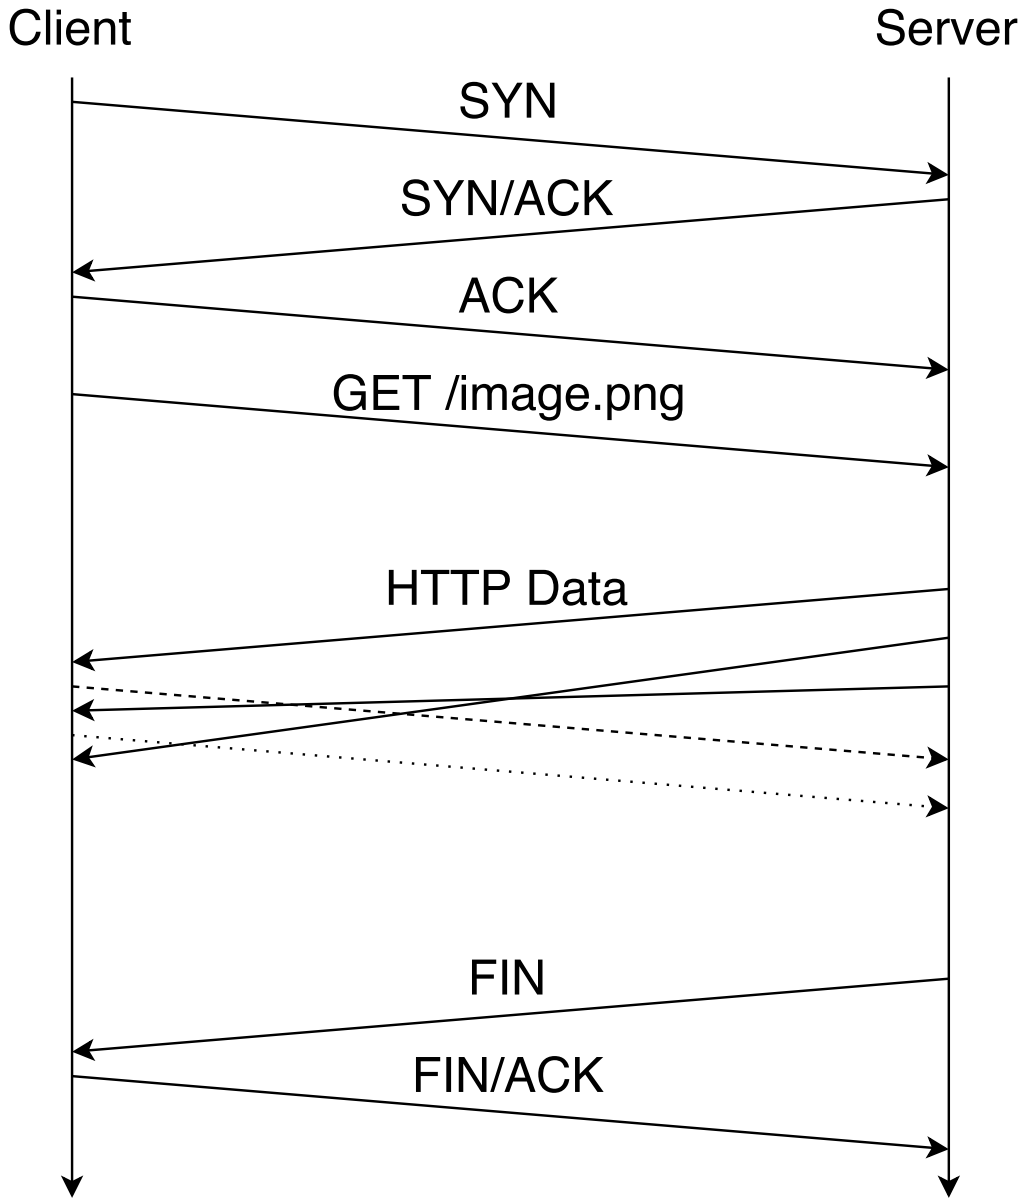

Anatomy of an HTTP request – Reordering

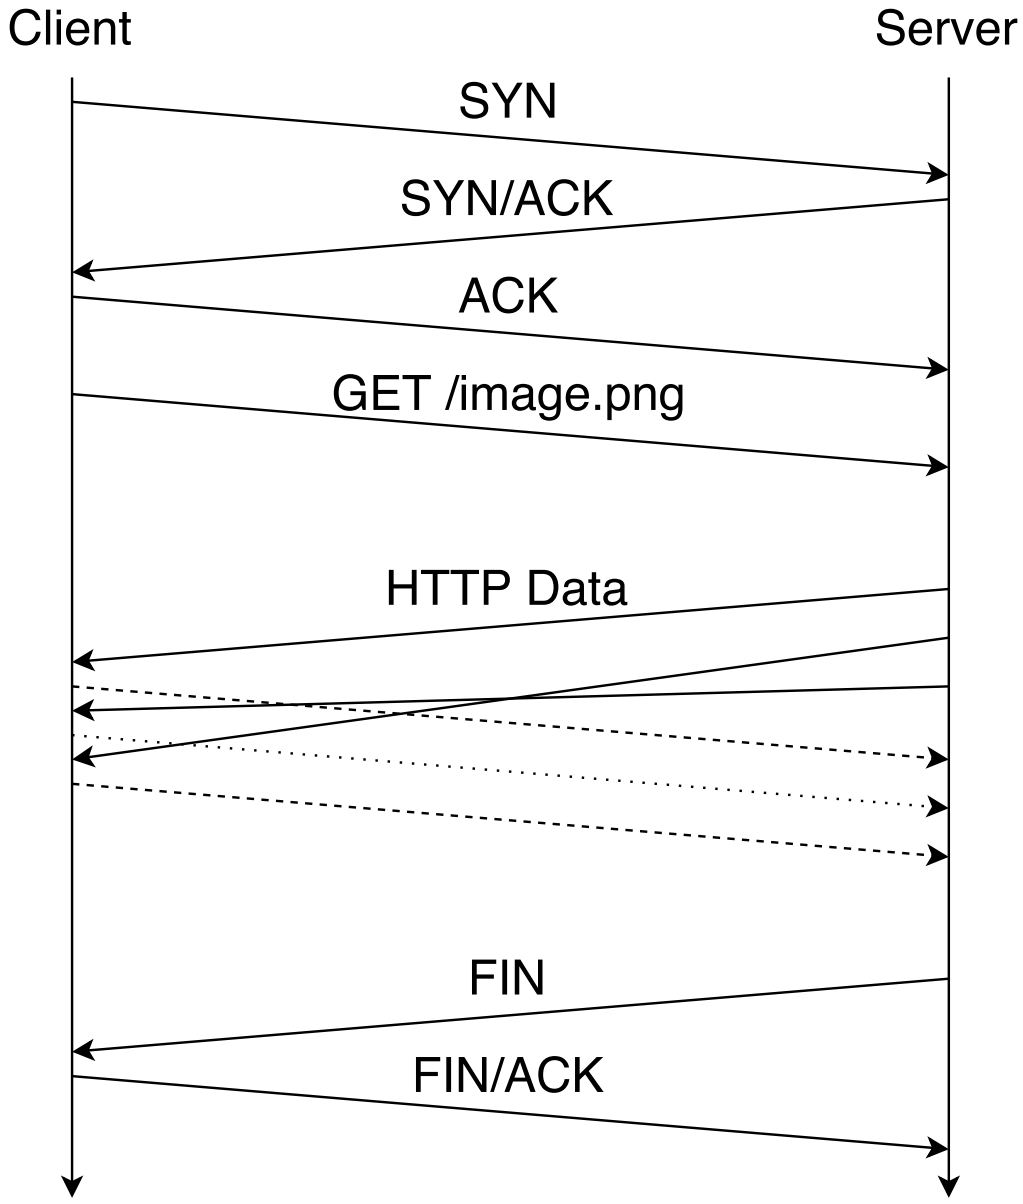

Reordering

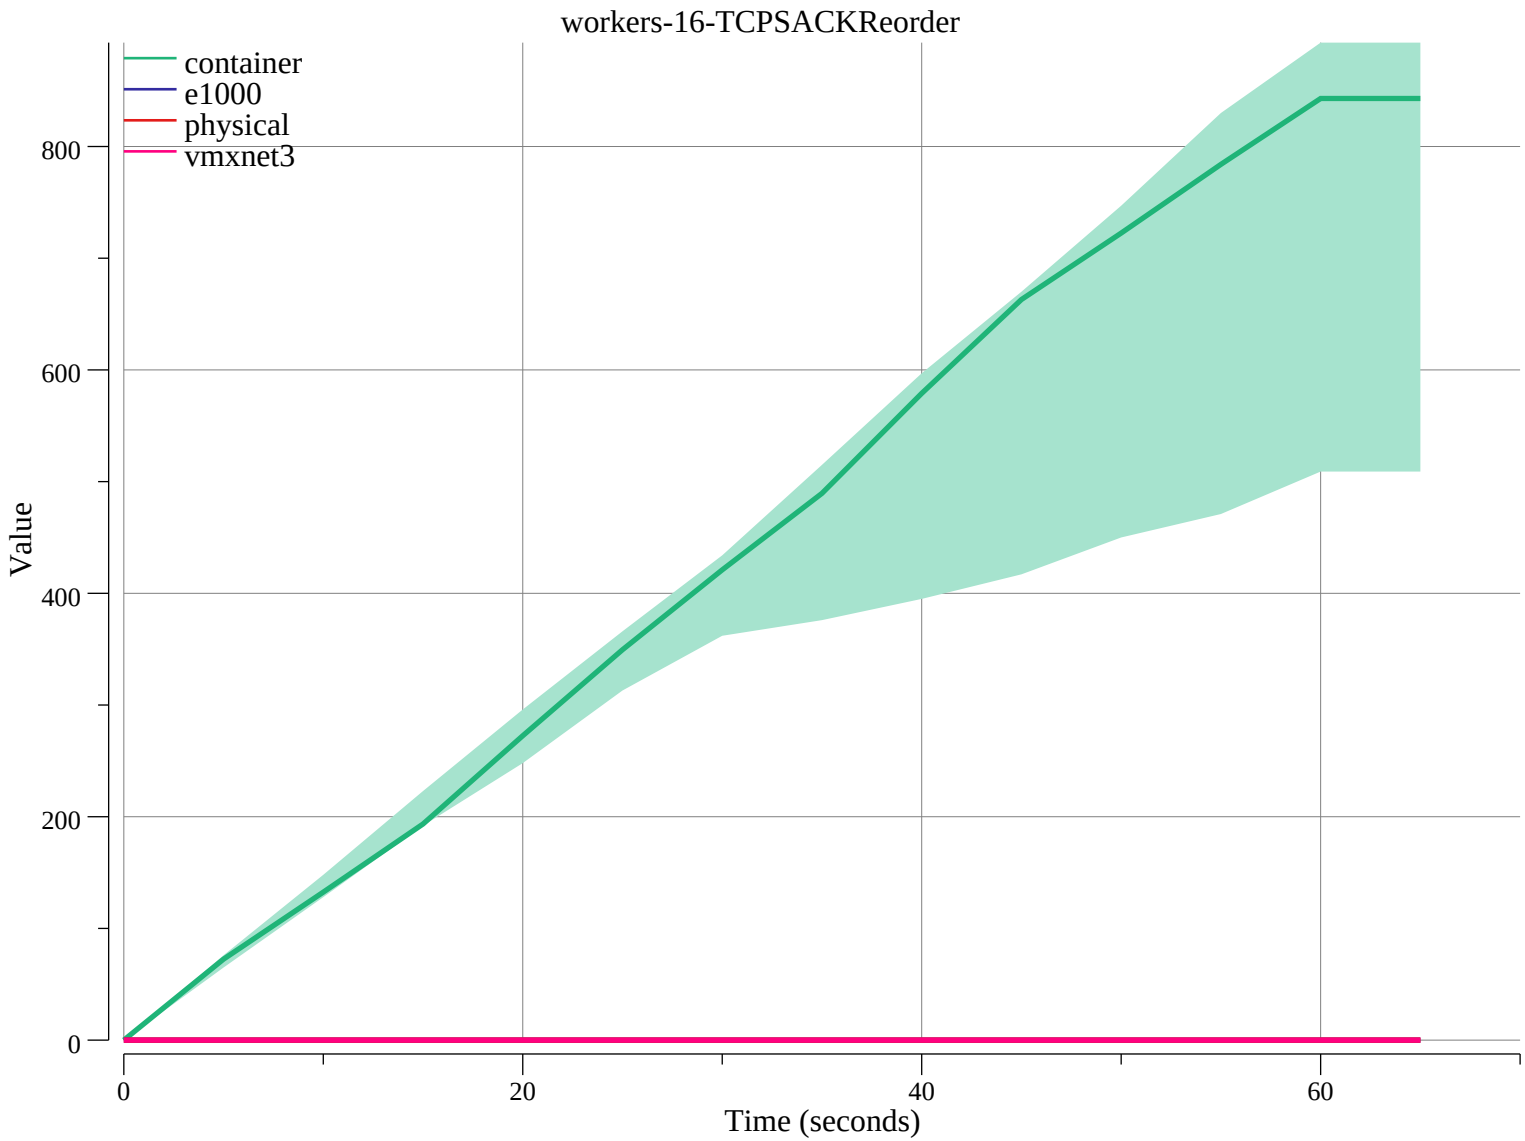
workers-16-TCPSACKReorder
container
e1000
physical
vmxnet3
Value
800
600
400
200
0
0
20
40
60
Time (seconds)

Anatomy of an HTTP request – Autocorking

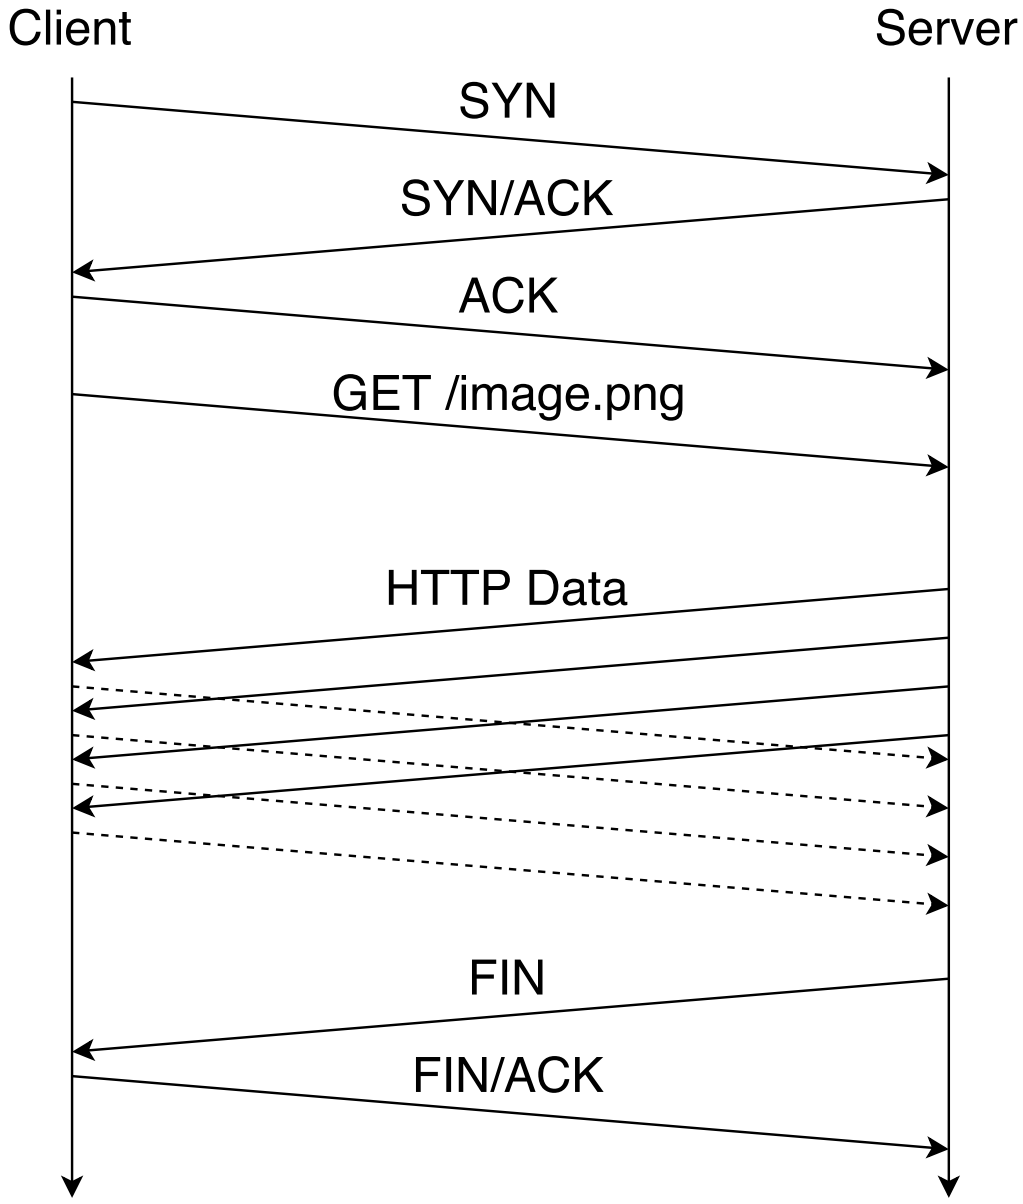

Anatomy of an HTTP request – Autocorking

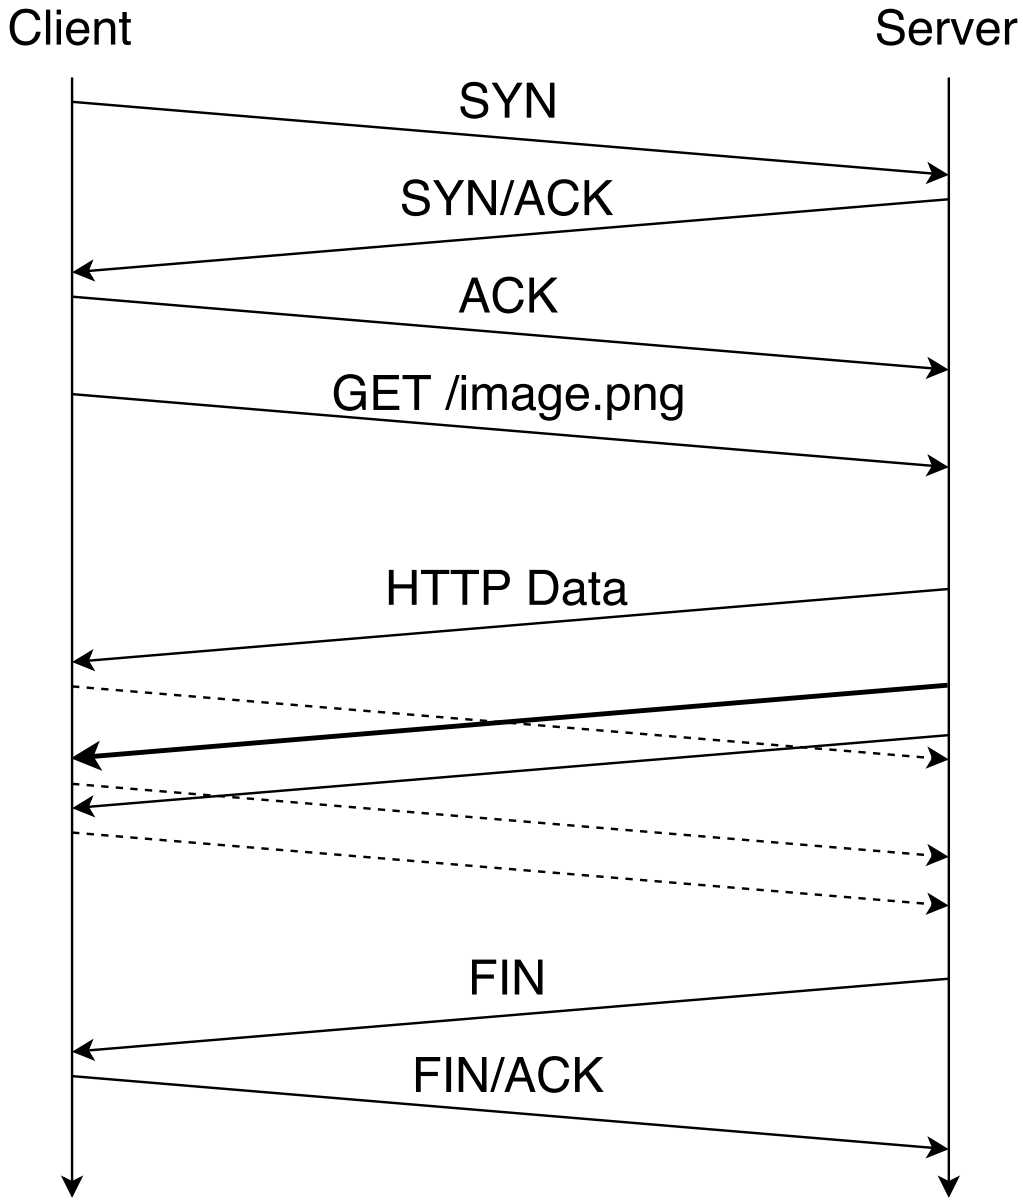

Autocorking

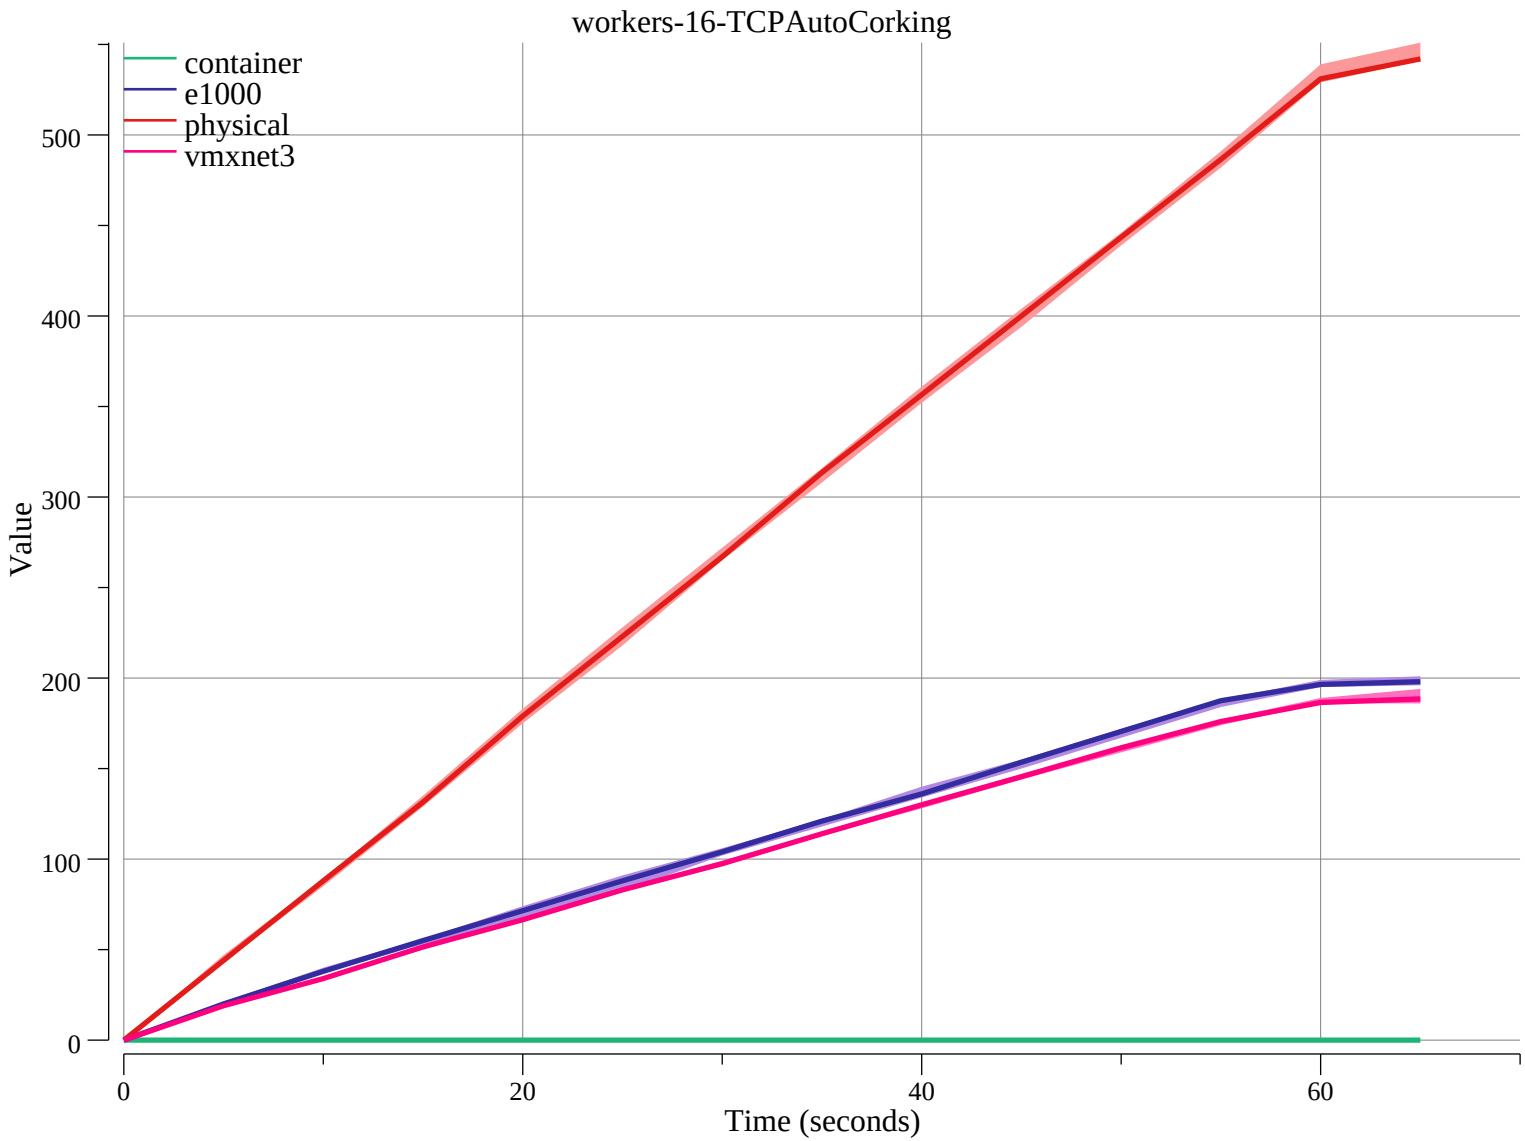

Next Steps

- Survey of Emulytics portfolio to guide future experiments
- Many more experiments to run
  - Varying resource contention
  - Varying network trunking
  - Varying workloads
  - ...

Conclusion

Questions/Comments?

Presenter: Jonathan Crussell
jcrusse@sandia.gov