





---

# Low Latency in the Cloud, with OSS

Mark Price  
@epickrram

---


# Properties of Financial Systems





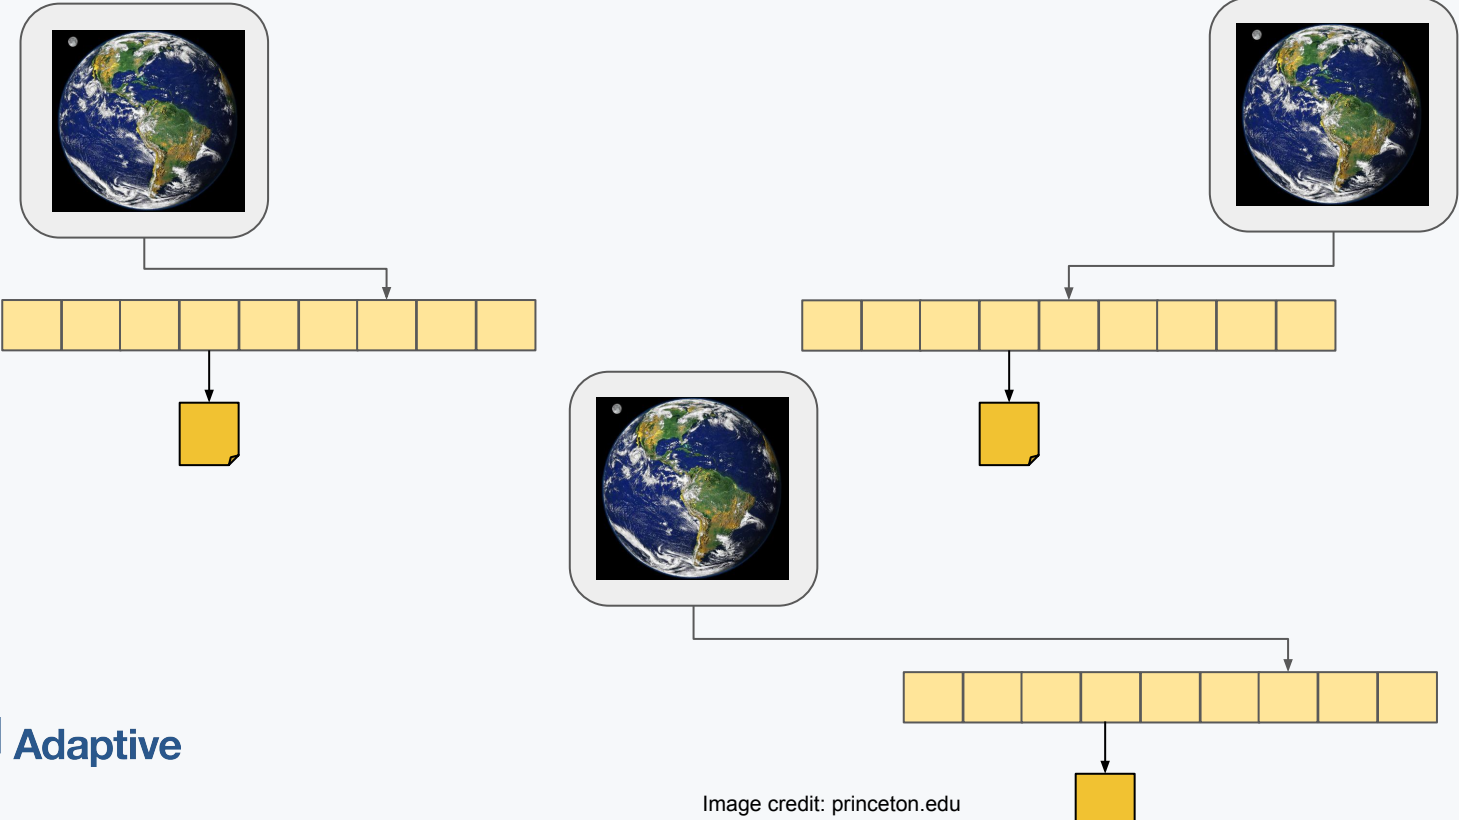
# Hydra

## Exchange Architecture

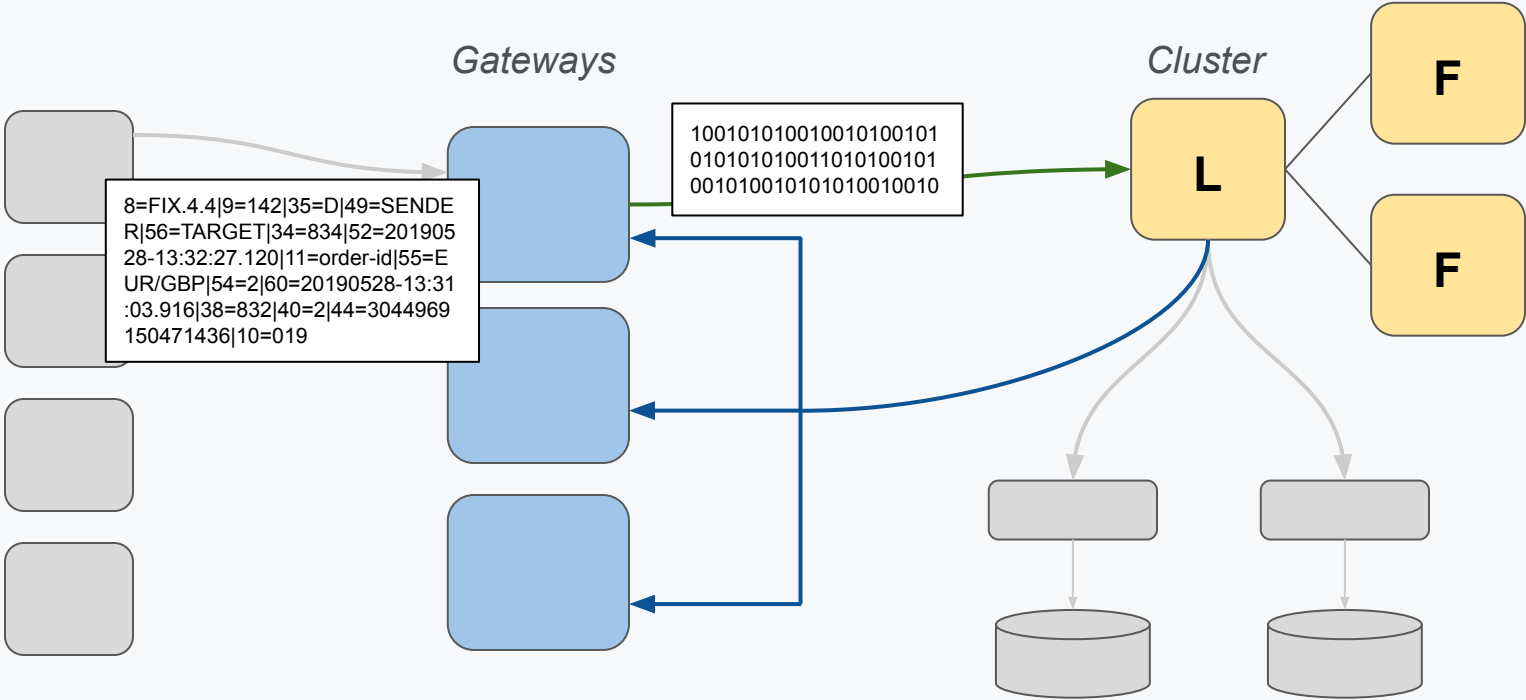


---

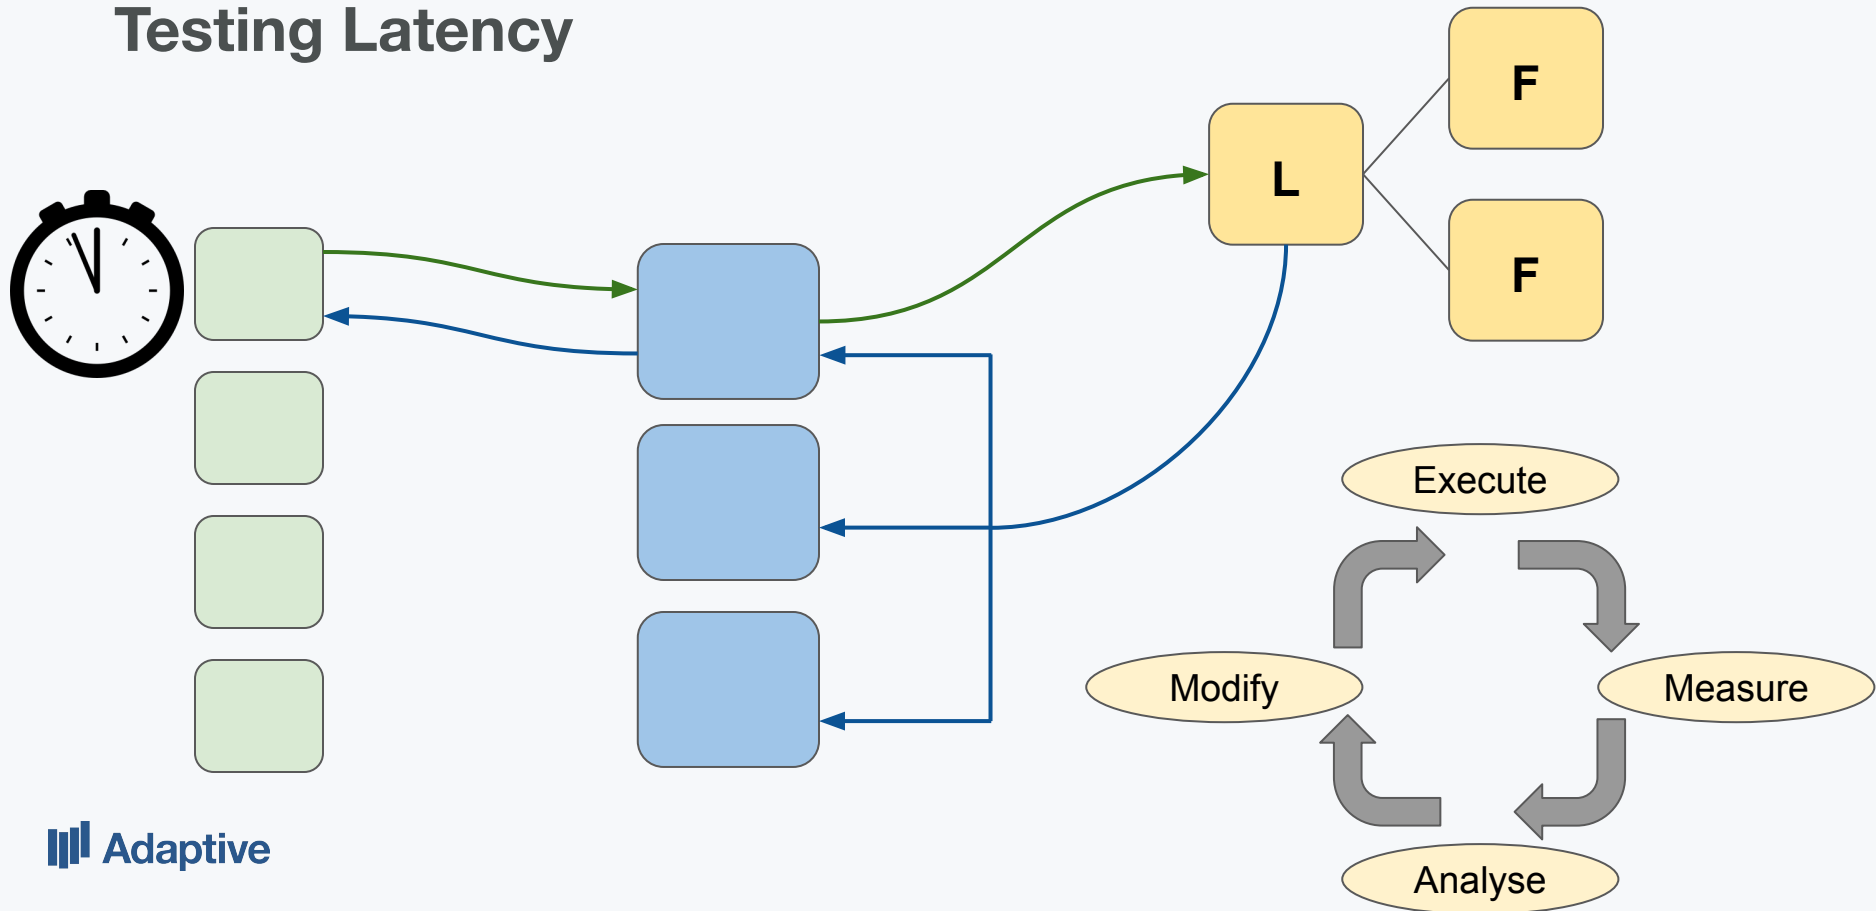
# Durability & Redundancy




# Speed



# Testing Latency




# Measurement in an Ideal World




---

# Measurement in the Cloud

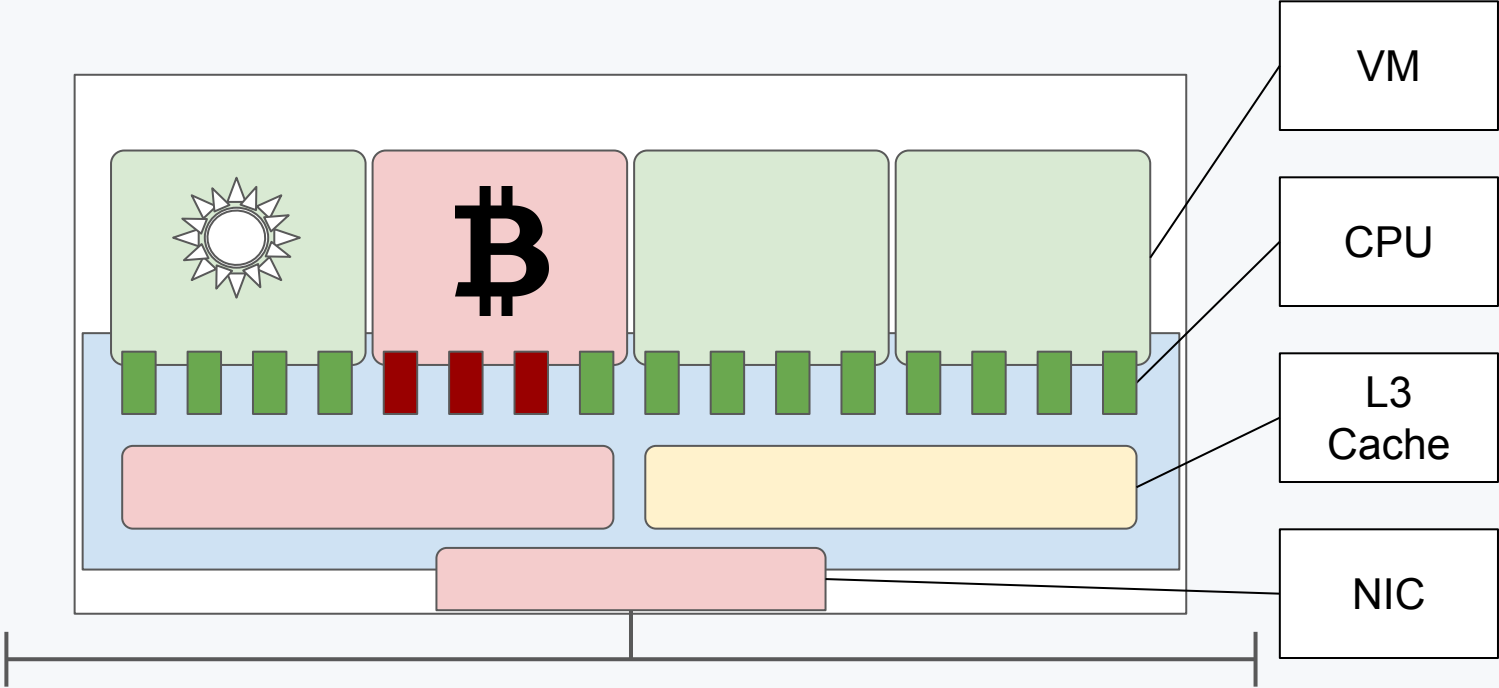




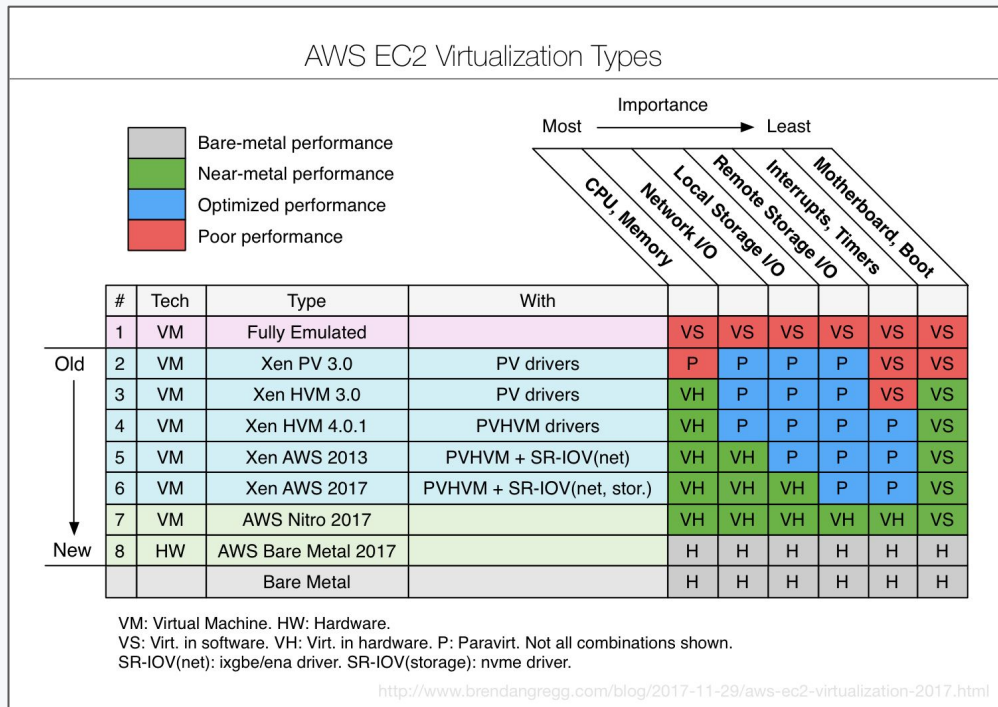
# Latency Measurements



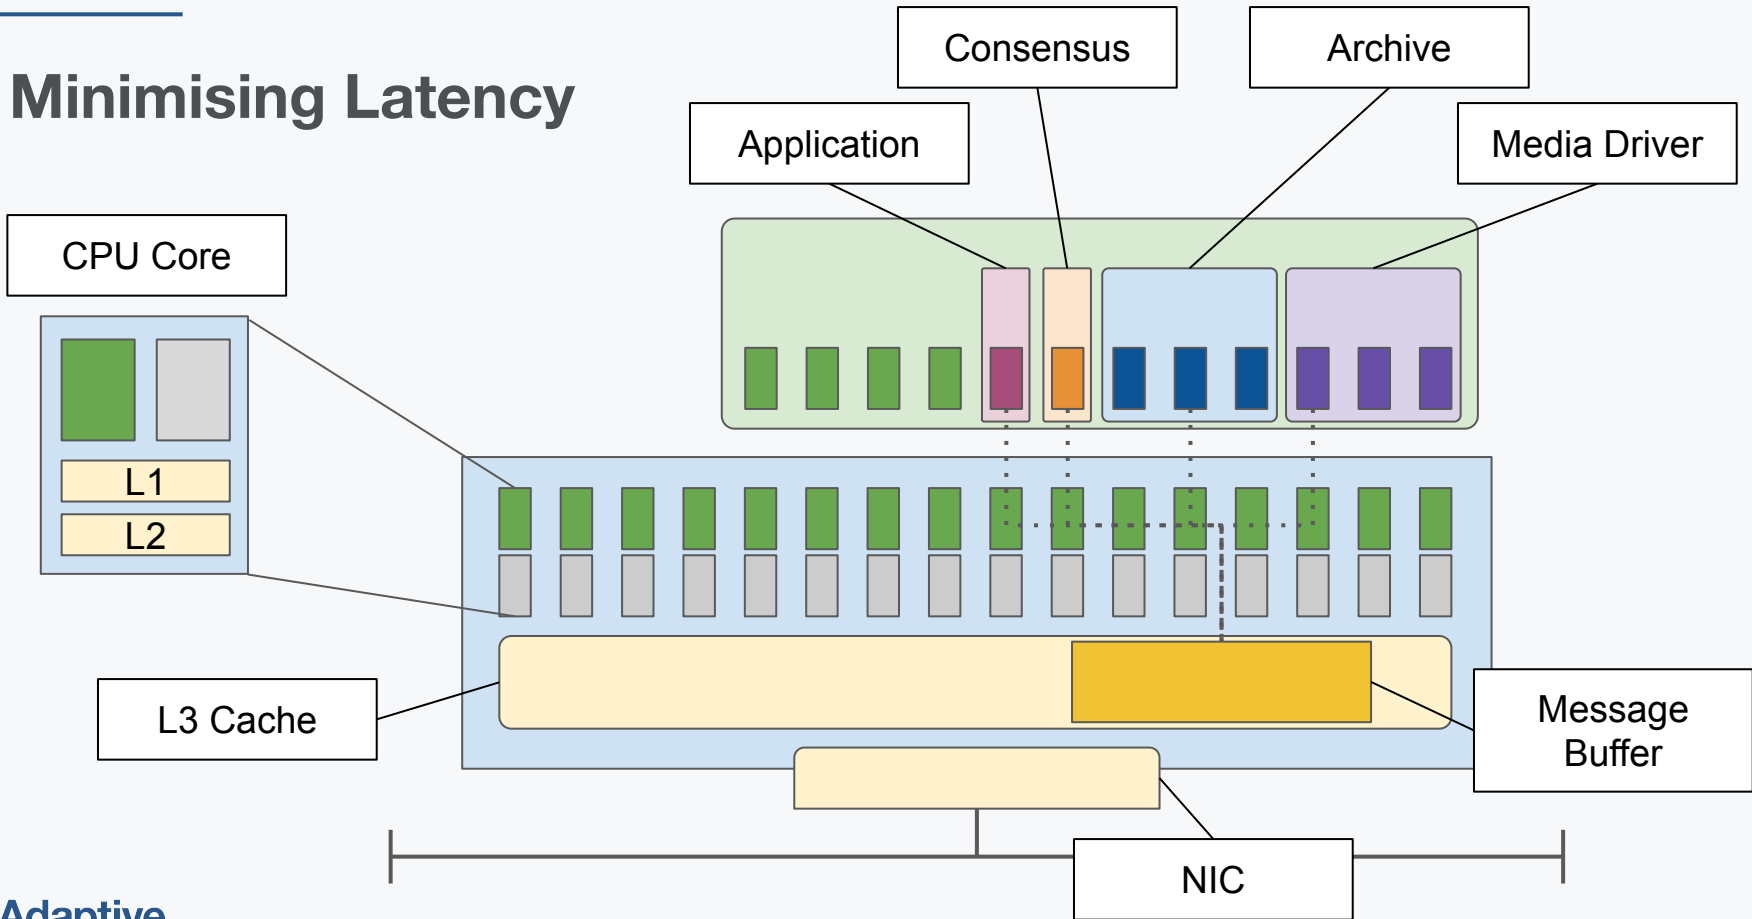
# Cloud Tenancy



# Evolution



# Minimising Latency



---

## Low Latency?

- 5,000 requests/s
- 99.9% lat
- 99
- 
- 

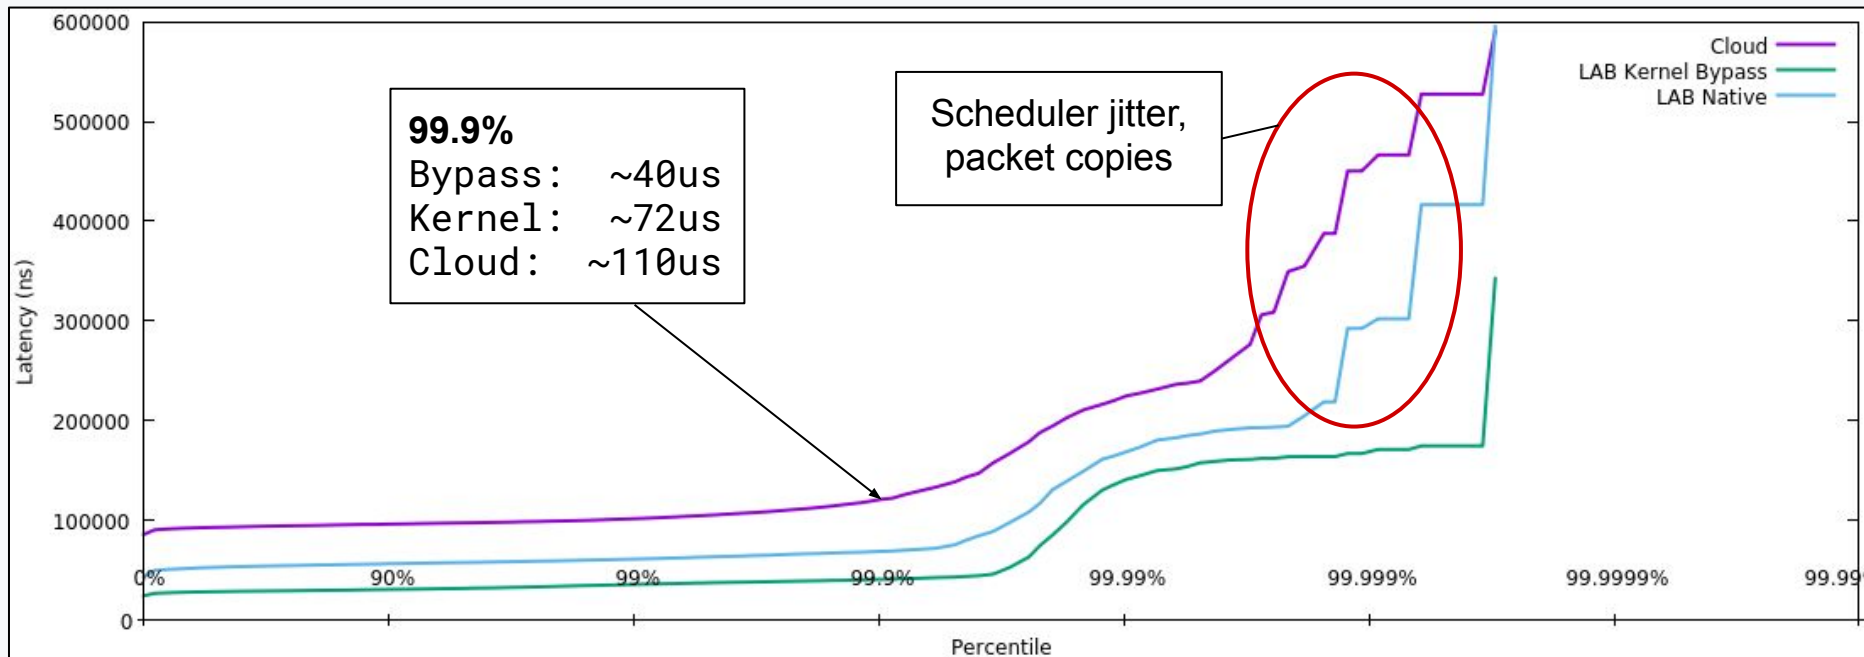
**UNREALISTIC  
WORKLOAD  
WARNING!**

---

## Low Latency?

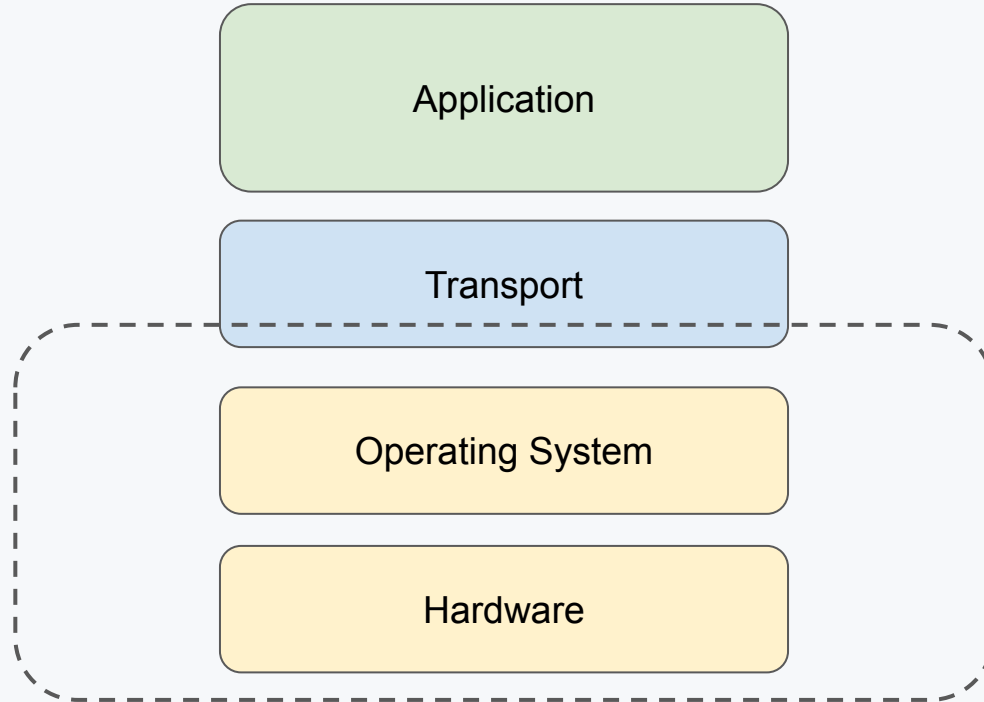
- 5,000 requests/sec
- 99.9% latency of **120** microseconds
- 99.99% latency of **223** microseconds
- Max latency of **590** microseconds
- vs max **342** microseconds on lab hardware

# RTT Latency



---

# The Future





---

## Hydra ❤️ Open Source

- Aeron (Messaging/Clustering)
- Artio (FIX)
- HDRHistogram (Metrics)
- + many more!

---

# Performance Test Hardware

## Adaptive Lab

Dell PowerEdge R540  
Intel(R) Xeon(R) Gold 5118 CPU @ 2.30GHz  
Caches: [32KB, 1MB, 16MB]  
RAM: 32GB DDR4 2.40GHz  
NVMe-attached SSD  
SolarFlare SFC9250 NIC  
Arista Switch  
CentOS 7

## AWS

m5d.metal  
Intel(R) Xeon(R) Platinum 8175M CPU @ 2.50GHz  
Caches: [64KB, 1MB, 32MB]  
RAM: 384GB  
NVMe-attached SSD  
ENA NIC  
Amazon Linux 1

---

Q & A

Thanks!