Science of Performance

Dr. Mourad Bouache

Performance Engineering
Verizon Media

June 24th, 2020

Storage Performance Development Kit (SPDK)
Persistent Memory Development Kit (PMDK)
Intel® VTune™ Profiler

intel

Virtual Forum

Performance Engineering Group (PEG)

verizon

Drive next-generation technologies and optimize existing footprint

Hardware Review

CAR-Wash
Hardware request evaluation
Server HW Modifications
Technical consulting
Benchmarking methodologies
Elastic cloud initiative

Data Mining & Analysis

Server utilization (Y!CSI)
Server optimization (YUP)
Hardware scanning
System abnormality reporting
Infrastructure planning support

R&D

Vendor management
Next-generation technology
TCO model
Server configuration selection
Other initiatives

Hardware Lab

Hardware certification
SKU Tool
Access to hardware
Performance benchmarking
Servers & network devices

AGENDA

1. **Performance**
   Benchmarking + Profiling
2. **Hardware and Software**
   CPU, Memory, IO, GPU, FPGA
3. **Cloud Performance**
   Scale Out Performance
4. **AI + Performance Engineering**
   Artificial Intelligence + Performance

Performance career: Mourad Bouache

verizon media group

PhD Computer Architecture

Post Doc
Paris - France
Illinois - USA
Waterloo - Canada

Yahoo
(35 data centers)

Oath
Verizon
(6 data centers)

Performance Research

Performance @ Industry

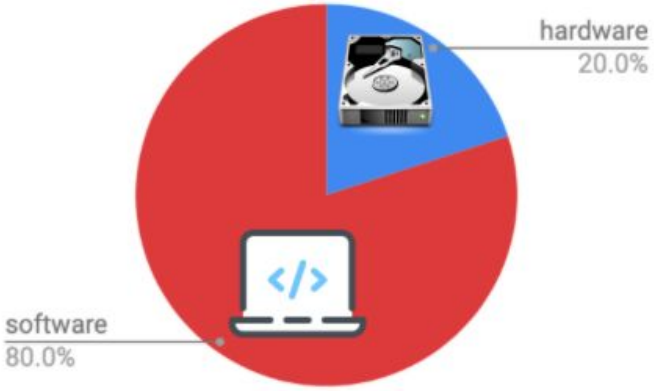

Interaction: SW/HW

Computer Architecture

2003

John Hennessy
David Patterson

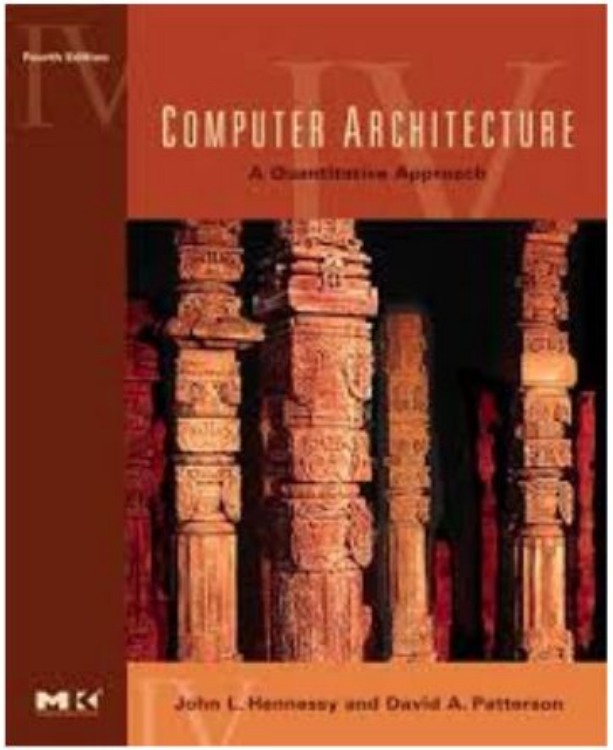

Berkeley
UNIVERSITY OF CALIFORNIA

David Patterson

Facebook @scale 2018

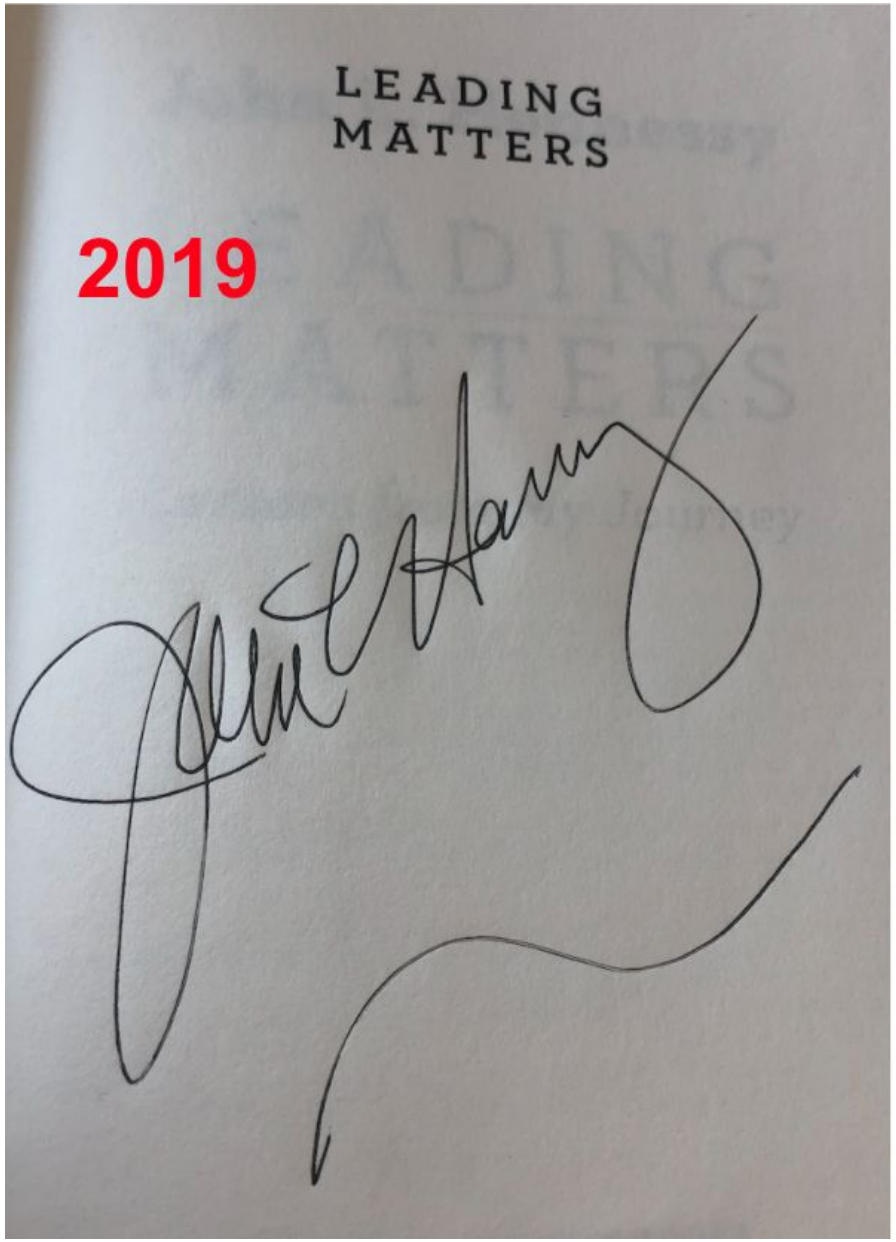
LEADING
MATTERS
2019


John Hennessy

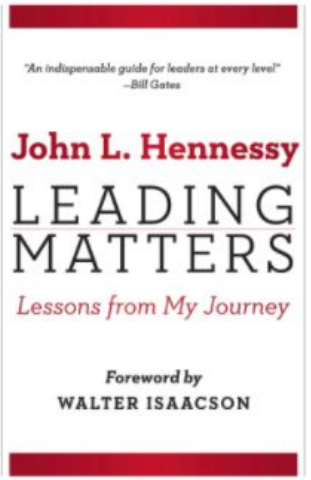
"An indispensable guide for leaders at every level"
—Bill Gates
John L. Hennessy
LEADING
MATTERS
Lessons from My Journey
Foreword by
WALTER ISAACSON

Where is Performance?

verizon media group

Software Performance

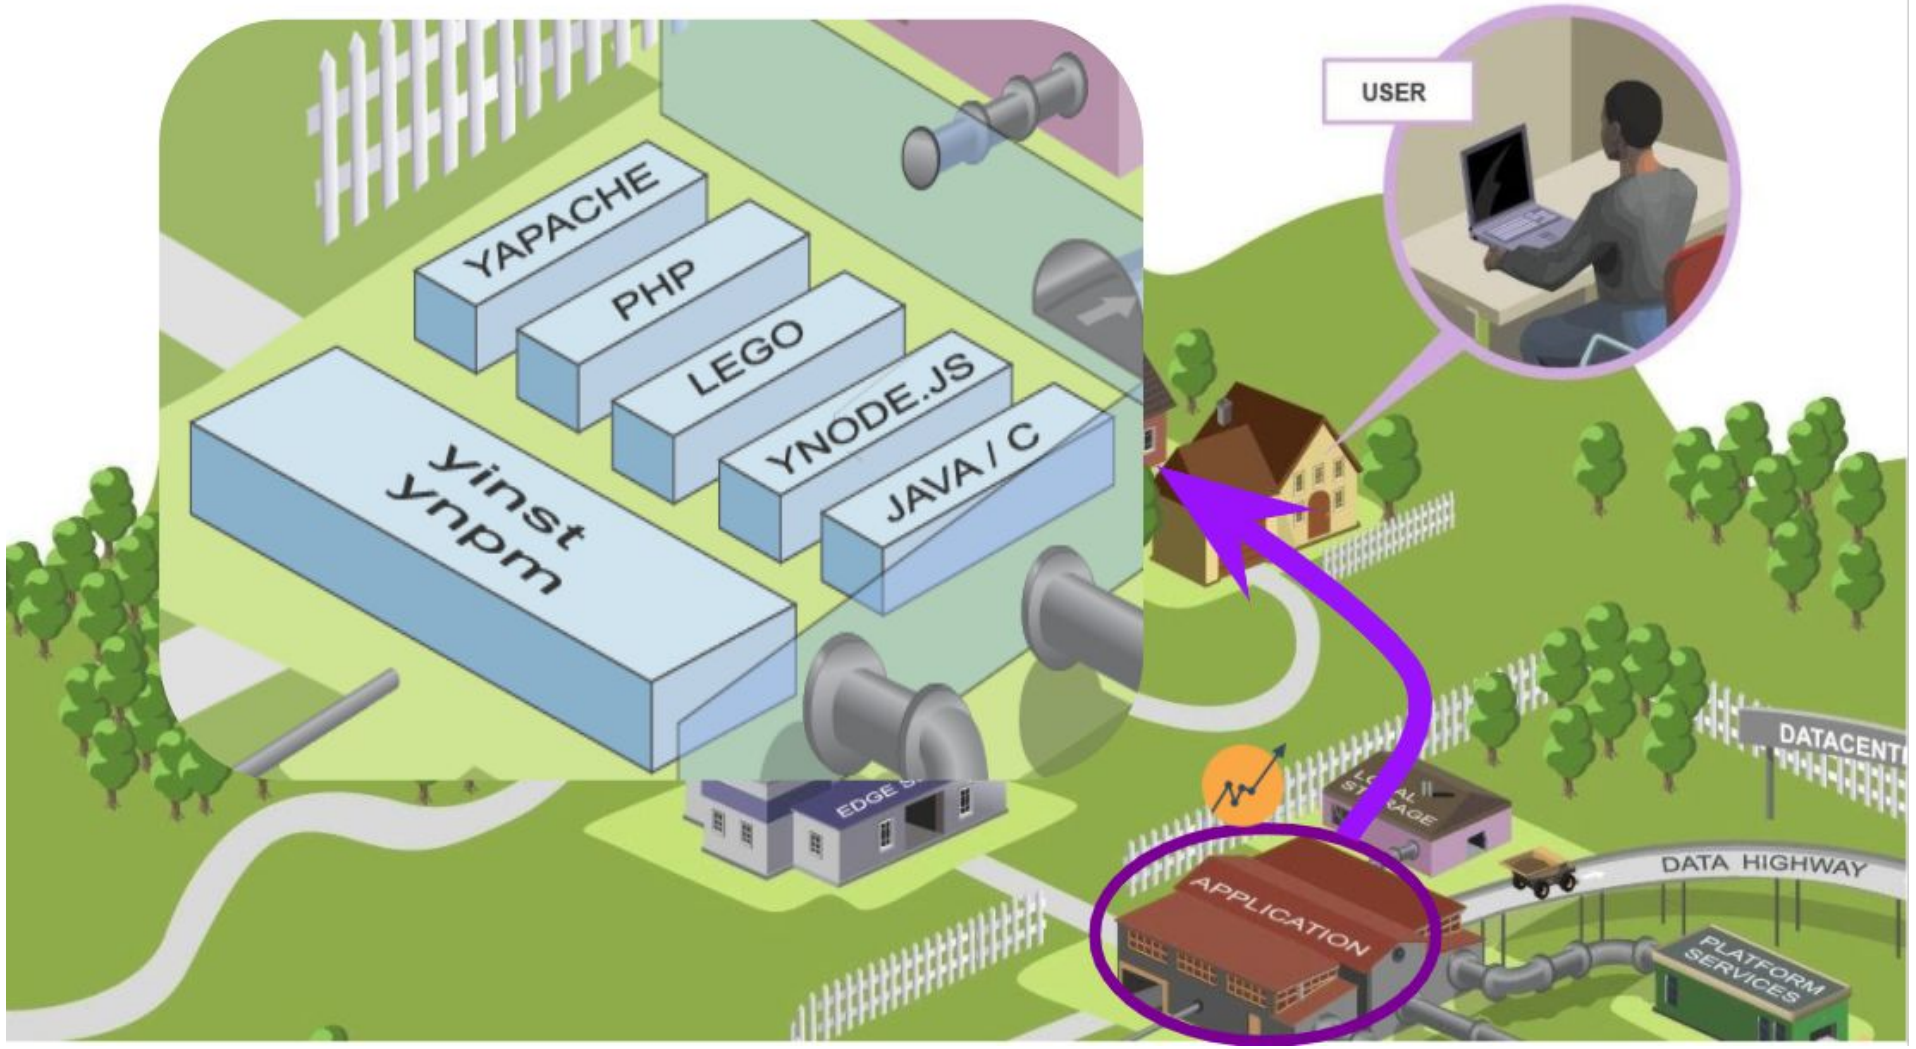

Performance study tells us:

- Why is some hardware better than others for different programs?
- What factors of system performance are hardware related? Do we need a new machine or new Hardware?
- How does the machine's instruction set affect the performance?

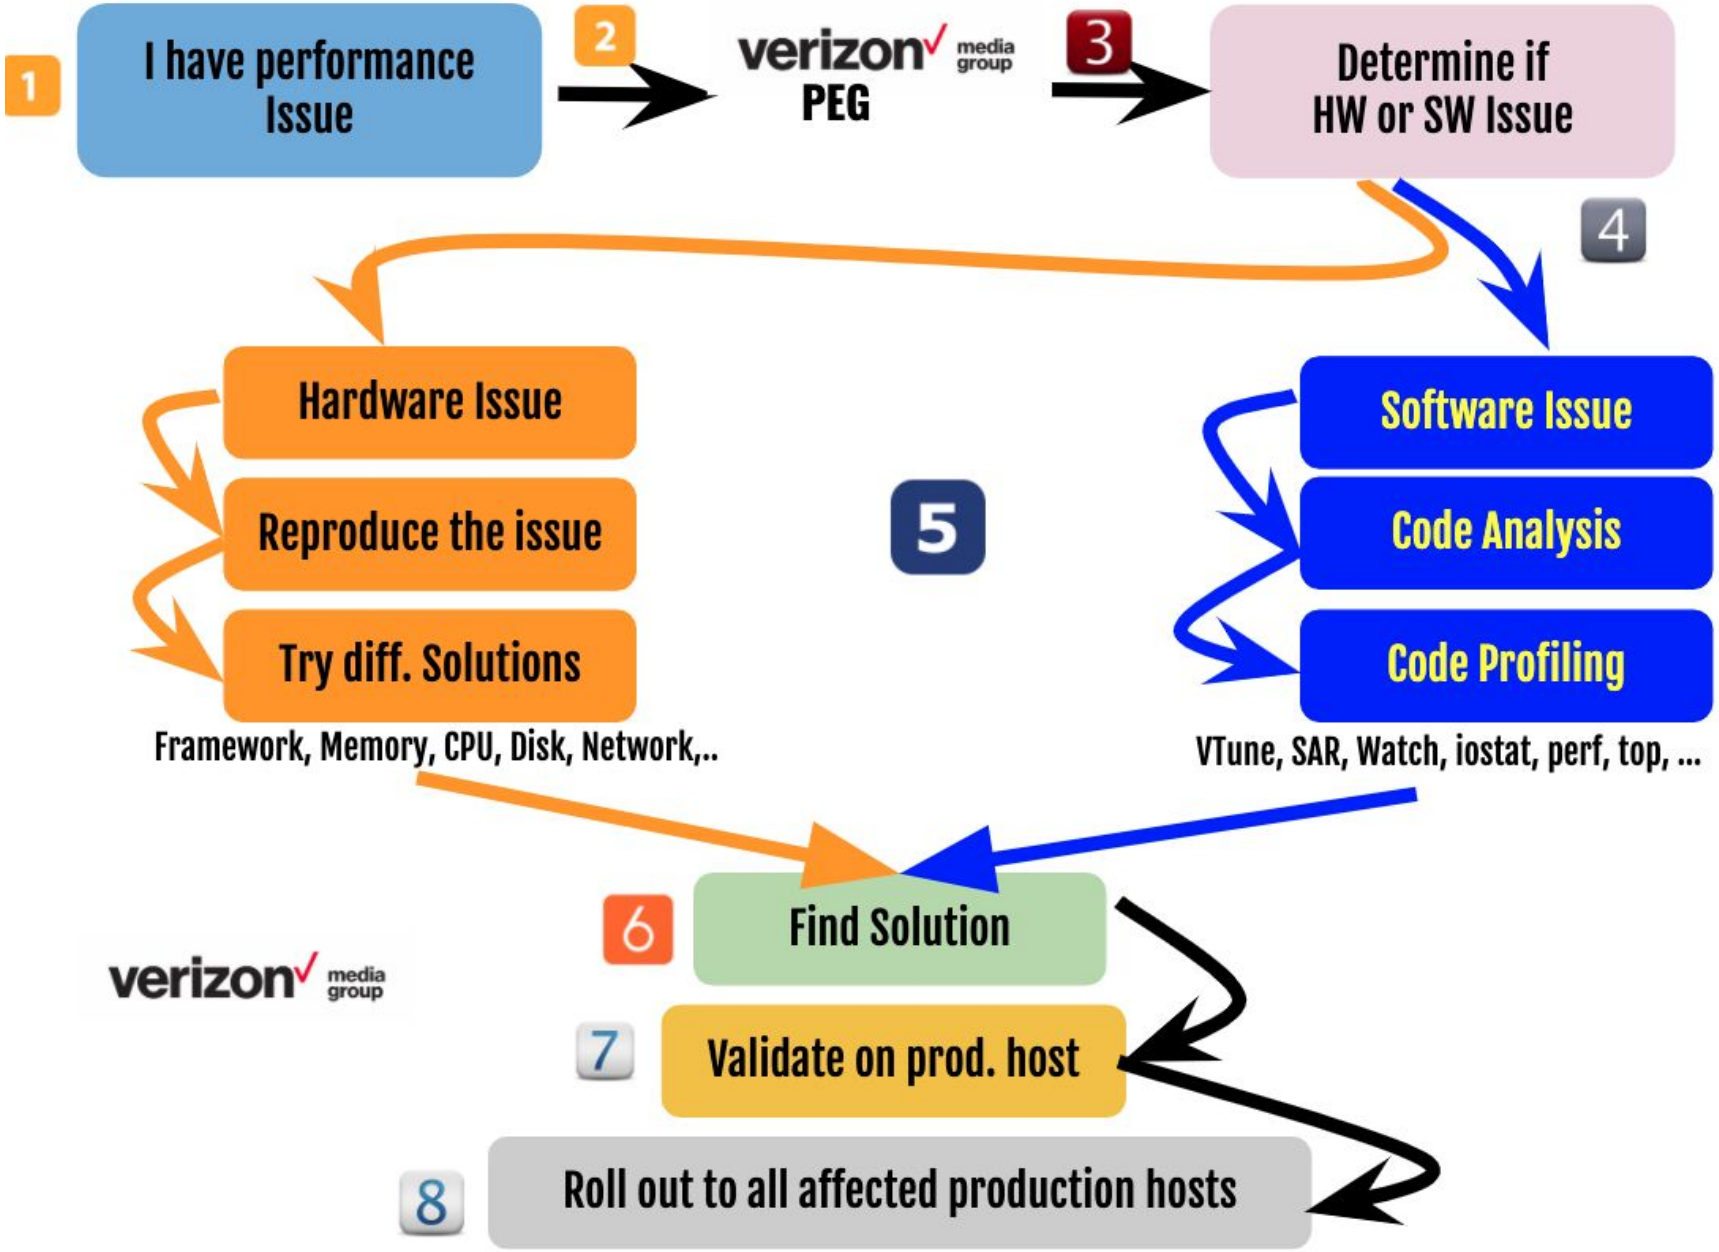
1
I have performance Issue
2
verizon media group
PEG
3
Determine if HW or SW Issue
4
Hardware Issue
Reproduce the issue
Try diff. Solutions
Framework, Memory, CPU, Disk, Network,..
5
Software Issue
Code Analysis
Code Profiling
VTune, SAR, Watch, iostat, perf, top, ...
6
Find Solution
7
Validate on prod. host
8
Roll out to all affected production hosts
verizon media group

Boeing 777

Boeing 747

AIRFRANCE

Lufthansa

DC 8-50

Concorde

Airline example

plane	speed (mph)	range (miles)	passengers	time (hr)	throughput (pxmph)
Boeing 777	610	4630	375	6.5	228750
Boeing 747	610	4150	470	6.5	286700
Concorde	1350	4000	132	3.0	178200
DC 8-50	544	8720	146	7.4	79424

What you want to achieve?

Speed

Long flight

Capacity/ Throughput

Psngr x speed

In one given criteria we can perform better than other criteria.

Two aspects of "Performance"

- Latency (response time/ execution time)

- how long does it take to execute my job?
- how long must I wait for the database query?

- Throughput (task per unit time)

- how many jobs can the machine run at once
- what is the average execution rate?
- how much work is getting done?

Performance definition

Performance is in **units of things-per-second**

- *Bigger is better*

if we are primarily concerned with response time

For a given programme:

- **performance(x)=1/execution_time(x)**
- **execution_time(x) = #inst×CPI×Cycle**

 **performance(x) =**  **execution_time(x)**

Performance of Computers

Performance is defined for ***a program and a machine***.

How to compare computers? Need benchmark programs:

- **Real applications**
- **Modified applications**
- **Kernels**

- **Benchmarks,** are measurements to evaluate the performance of a function, operation or business relative to others.

Hardware Performance

- CPU
- Memory
- Network
- Disk - Input Output

CPU Generations: Tick-Tock

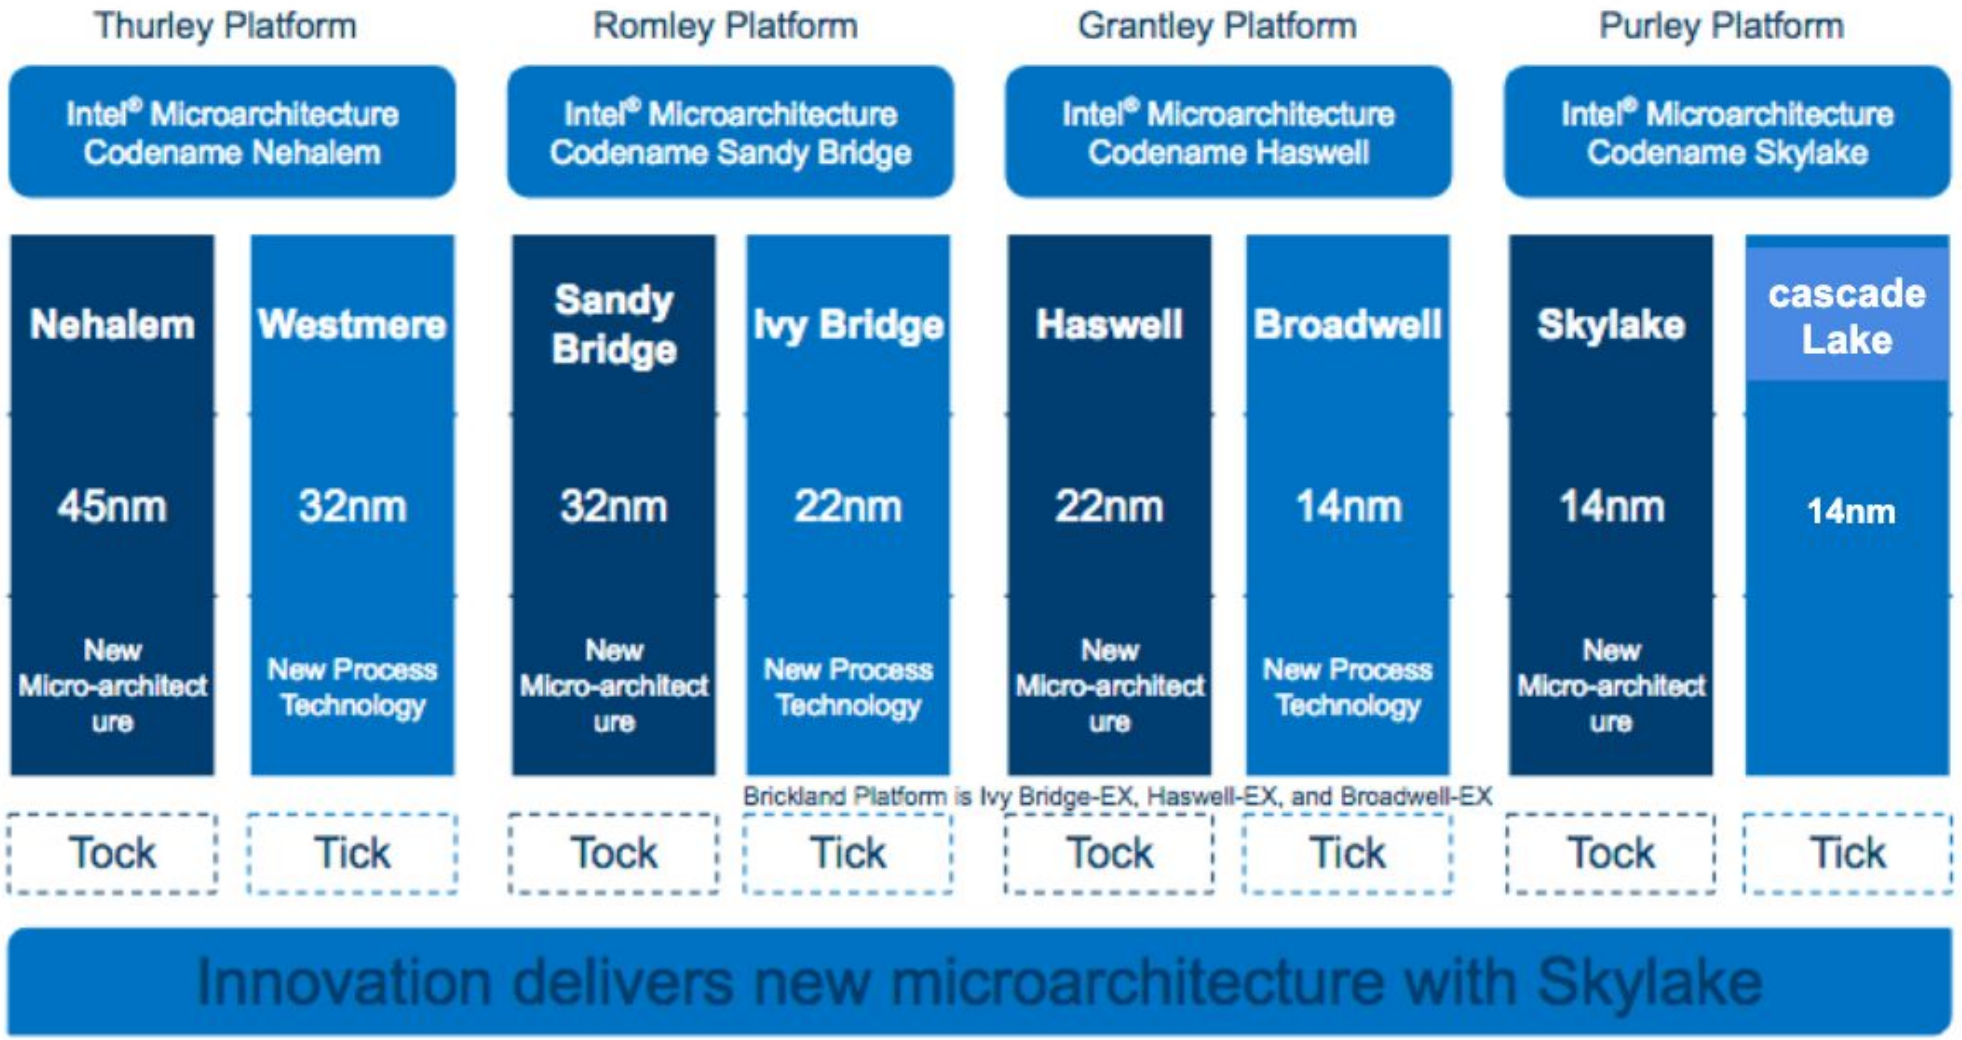

CPU Frequency

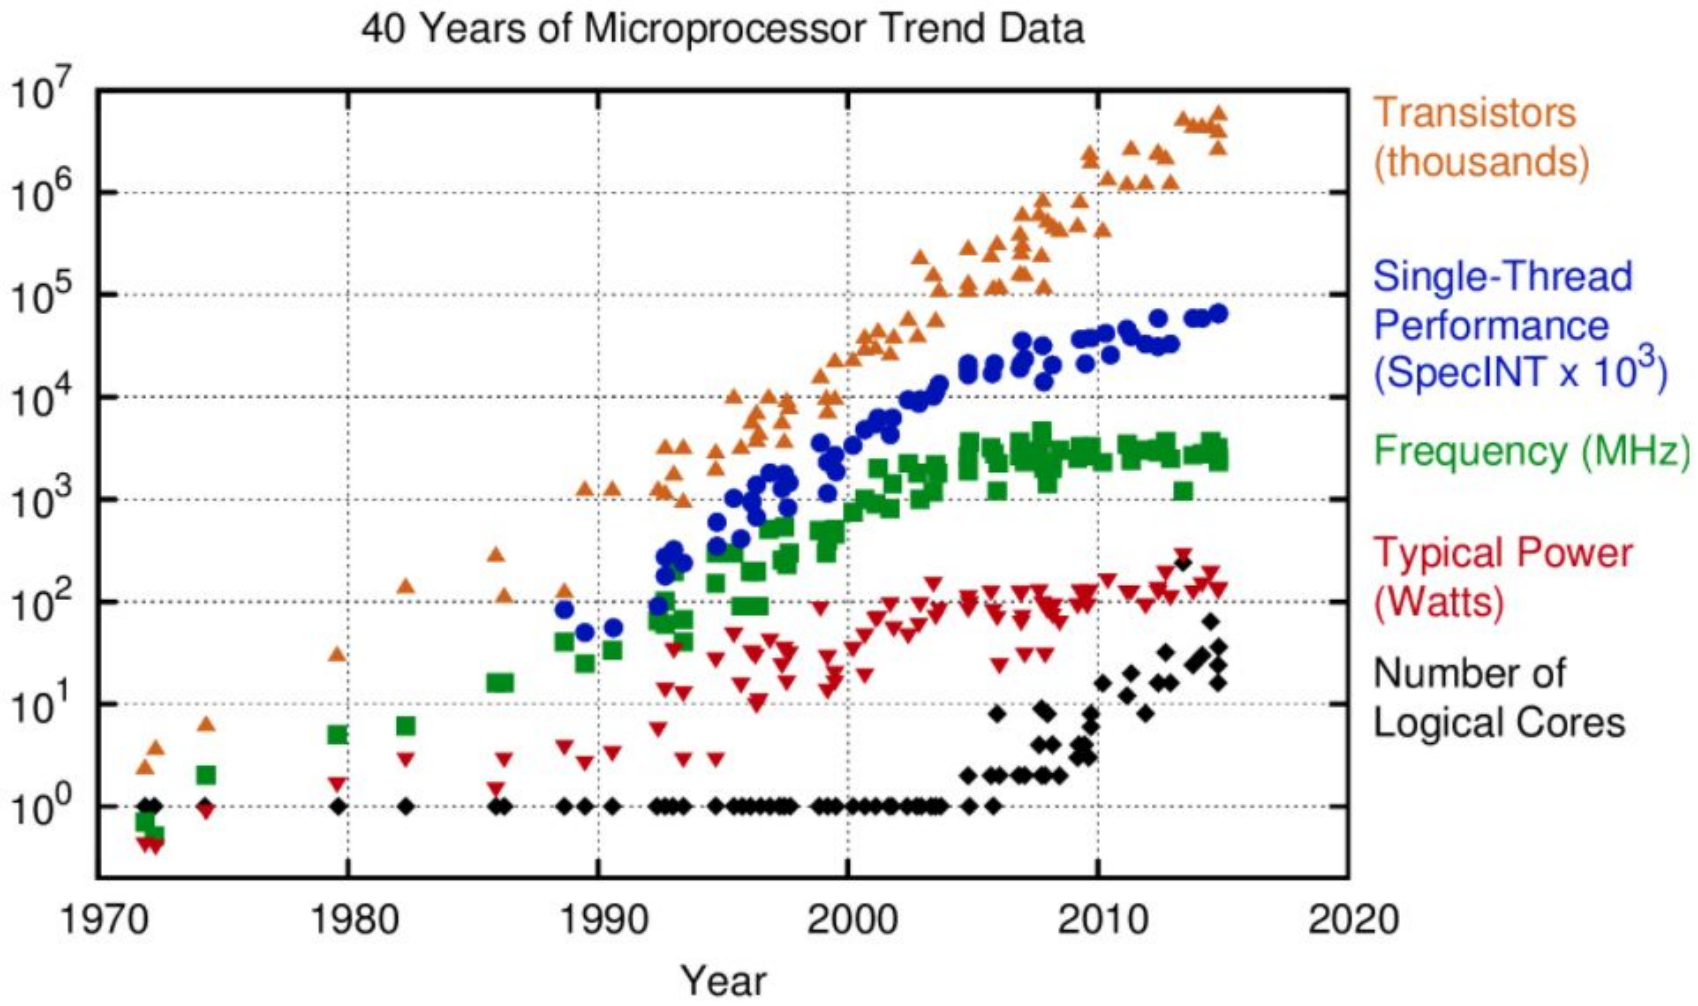

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

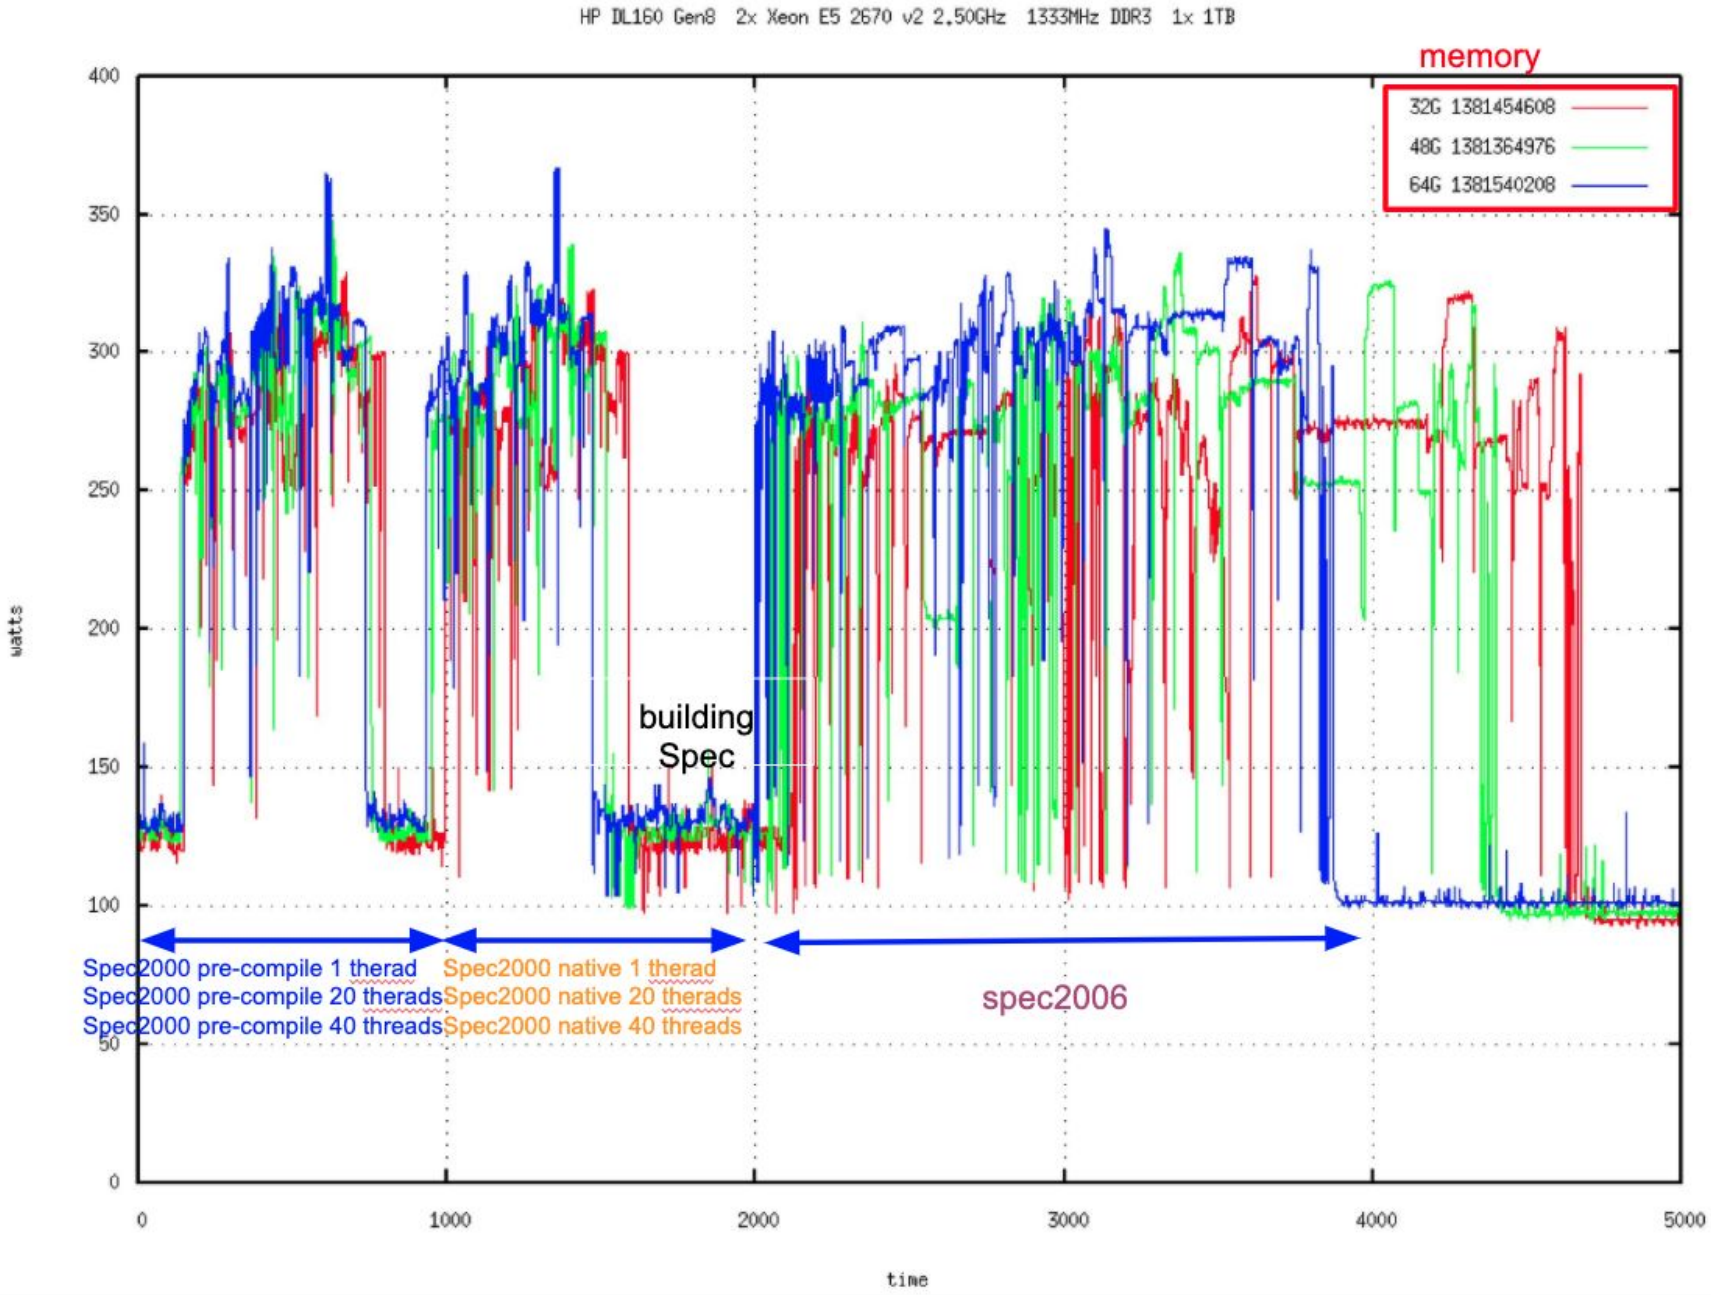
HP DL160 Gen8 2x Xeon E5 2670 v2 2.50GHz 1333MHz DDR3 1x 1TB
memory
32G 1381454608
48G 1381364976
64G 1381540208
watts
time
building Spec
Spec2000 pre-compile 1 therad
Spec2000 pre-compile 20 therads
Spec2000 pre-compile 40 threads
Spec2000 native 1 therad
Spec2000 native 20 therads
Spec2000 native 40 threads
spec2006

Hardware certification Process

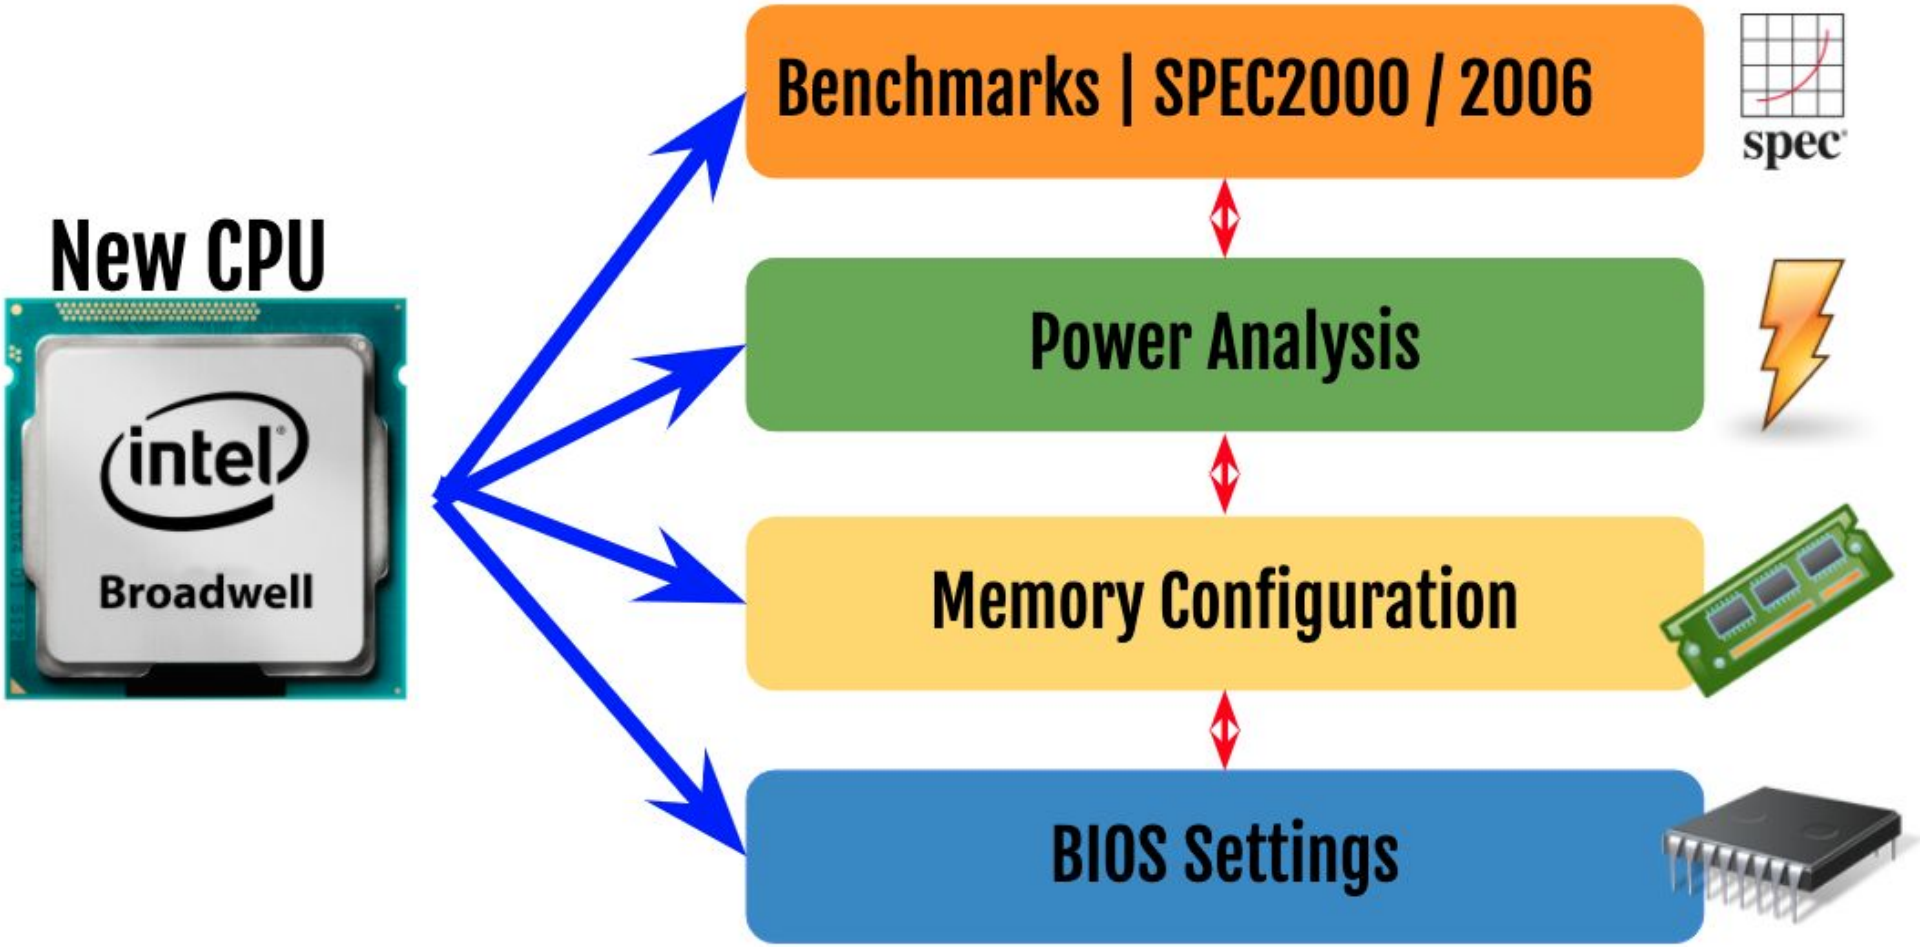

Server component

CPU Performance

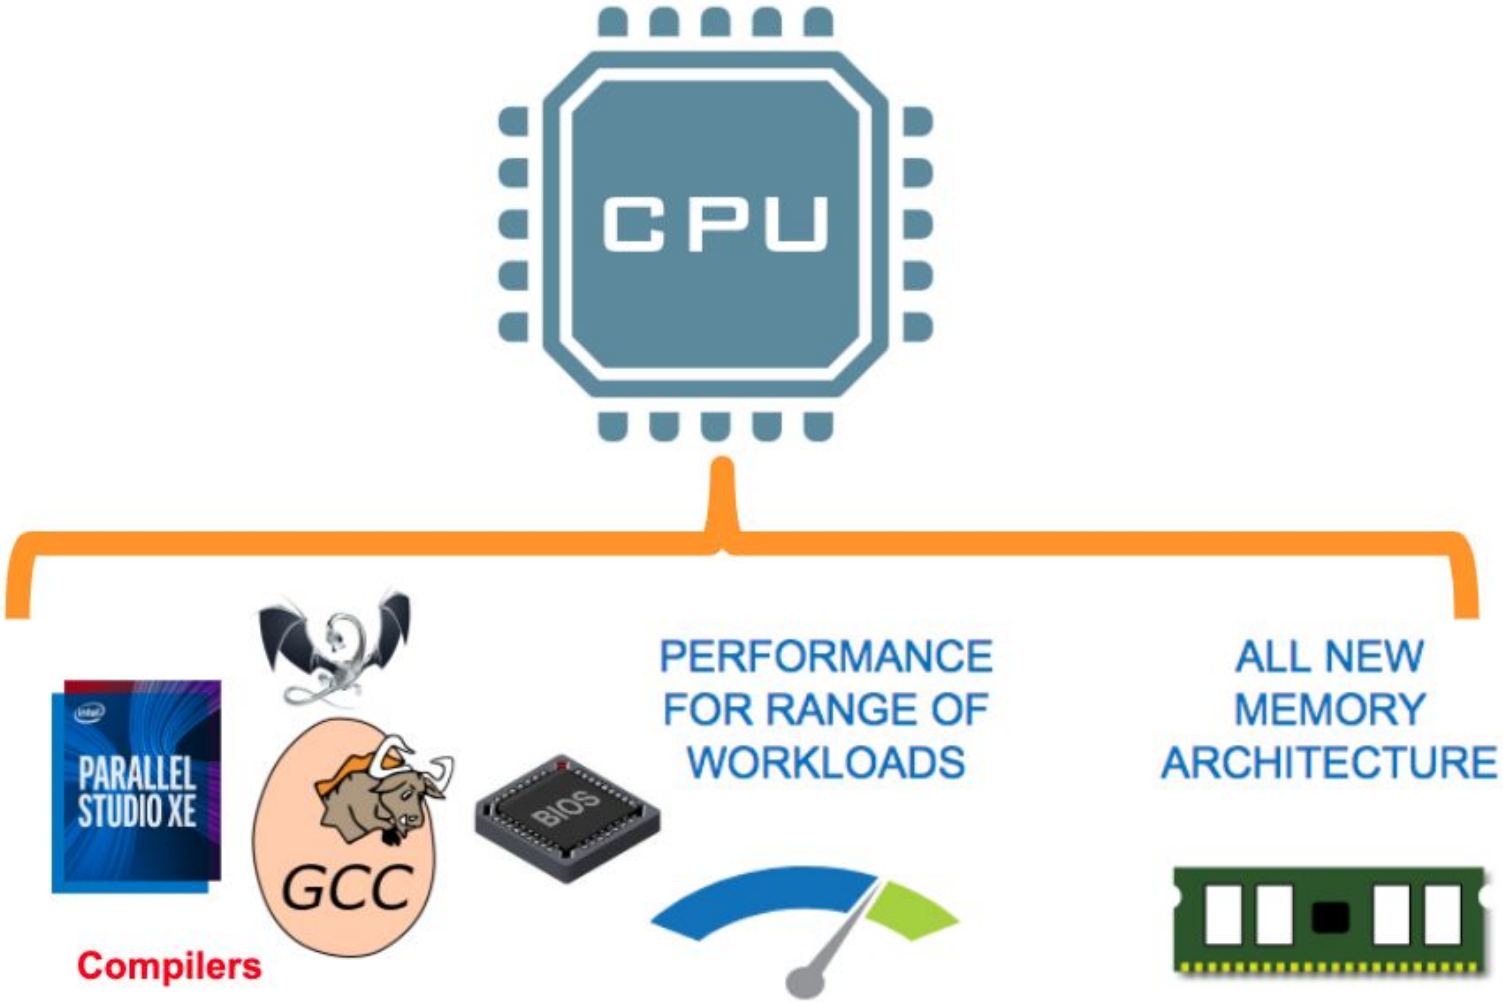

Skylake CPU

Spec	Purley with Skylake CPU
CPU TDP	70-205W
Socket	Socket P
Cores	Up to 28C with Intel® HT Technology
Memory	6 channels DDR4 per CPU DIMM
	2133, 2400, 2666 2DPC No 3 DPC support
Intel® UPI	UPI: 2-3 channels per CPU (9.6, 10.4 GT/s)
PCIe*	PCIe* 3.0 (2.5, 5.0, 8.0 GT/s)
	48 lanes per CPU support: x16, x8, x4

Memory Performance

Memory is a limited resource available to the application and poor use of that resource will affect performance.

DIMM Speed. Faster DIMM speeds deliver lower latency, particularly loaded latency.
1600 MT/s has about 20% < loaded latency than 1333 MT/s.

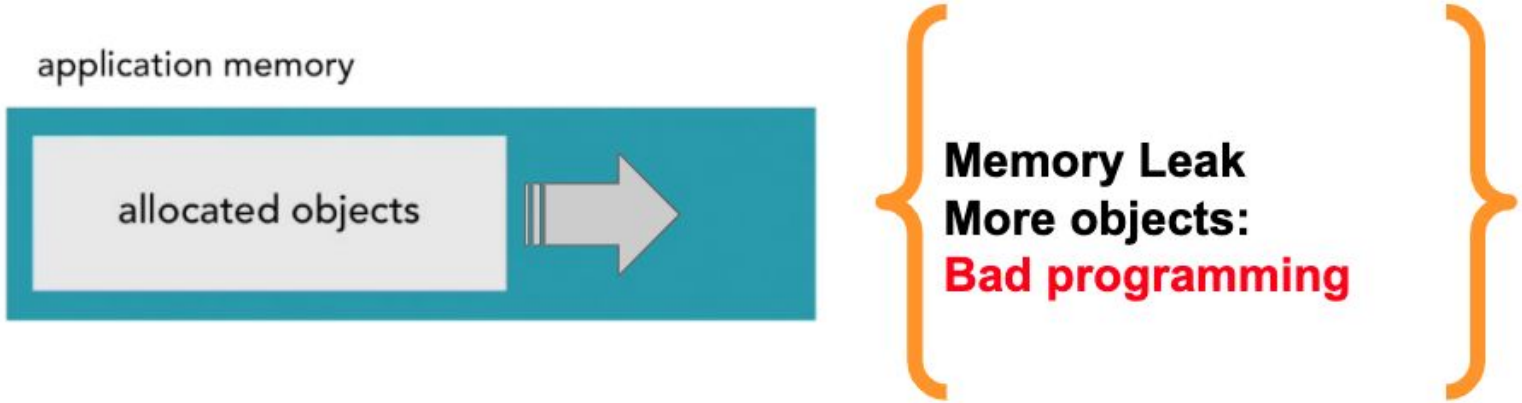

Memory Efficiency

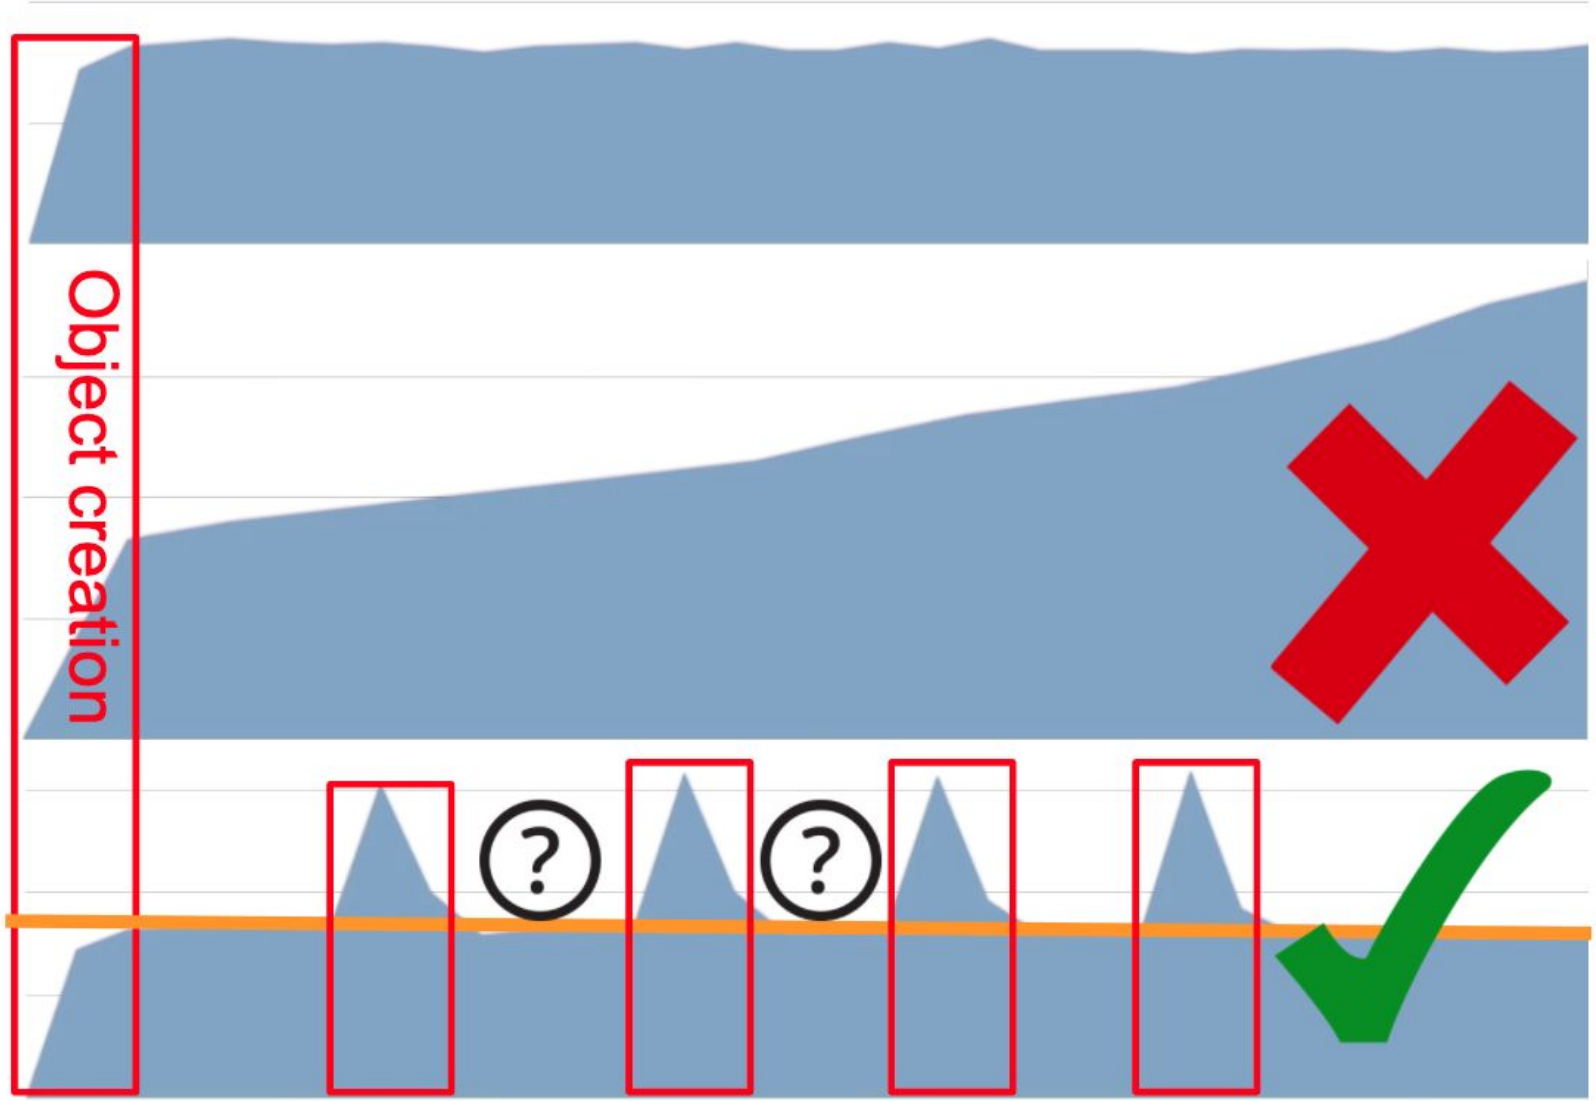

Memory Performance

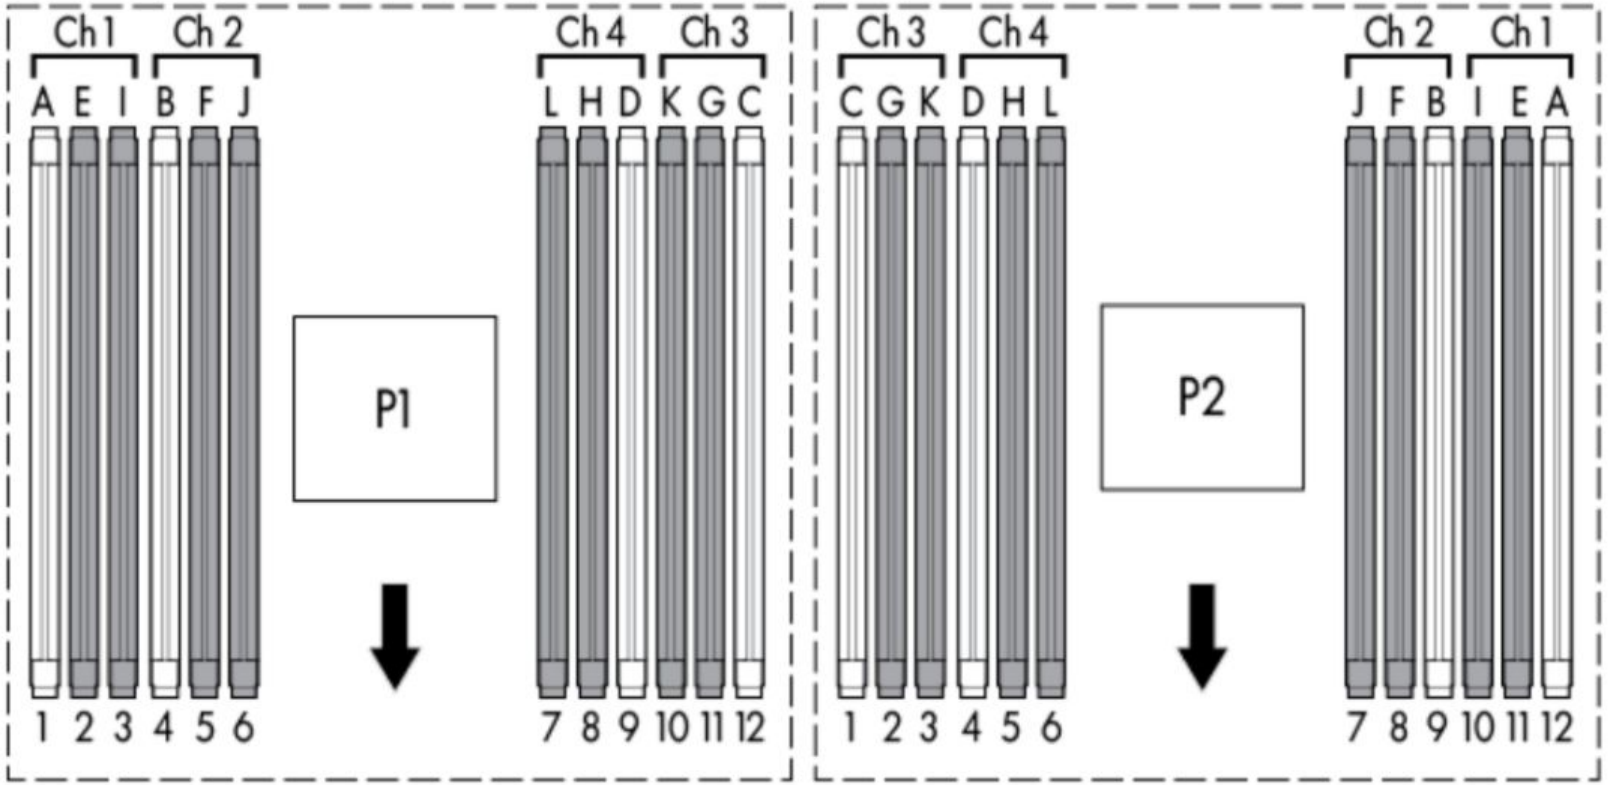

Important: Accurate memory population

Memory Performance

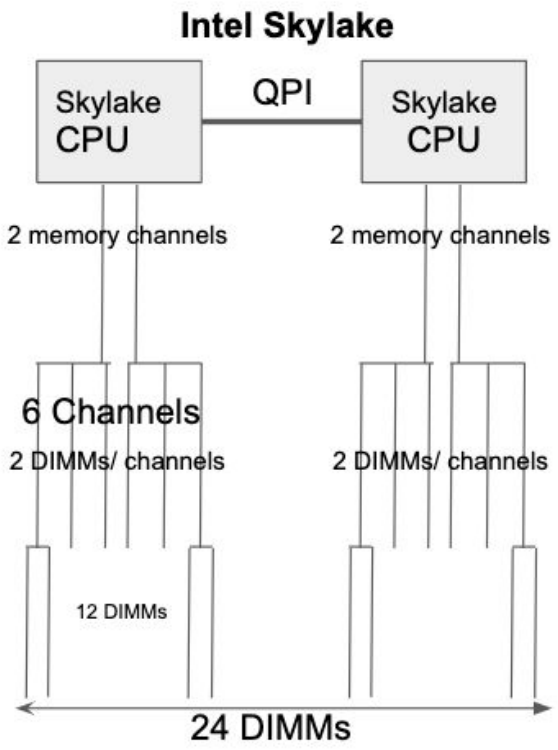

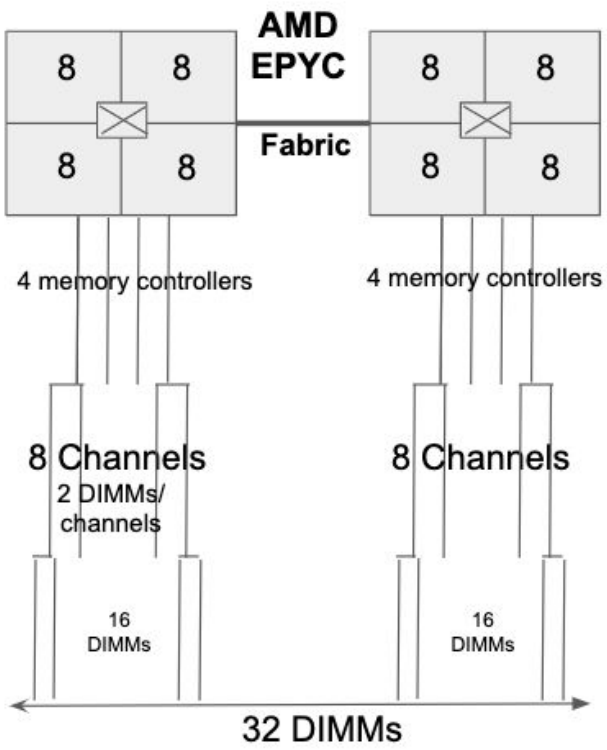

42GB/sec bi-dir BW per link

More ~33% better performance in AMD with extra 8 DIMMS/ 4 memory channels

Memory Performance

DDR Version	DDR1	DDR2	DDR3	DDR4	DDR5
Released Date	2000	2003	2007	2012	2020
Data Rate (MT/s)	200-400	400-800	800-2133	1600-3200	3200-6400

I/O Performance

Minimize disk Reads and Writes

main memory reference: 100ns
read 1MB: 20,000ns-100,000ns

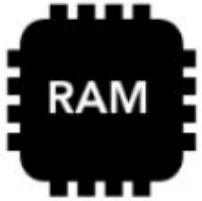

drive seek: 4,000,000-10,000,000ns
read 1MB: 2,000,000ns-20,000,000ns

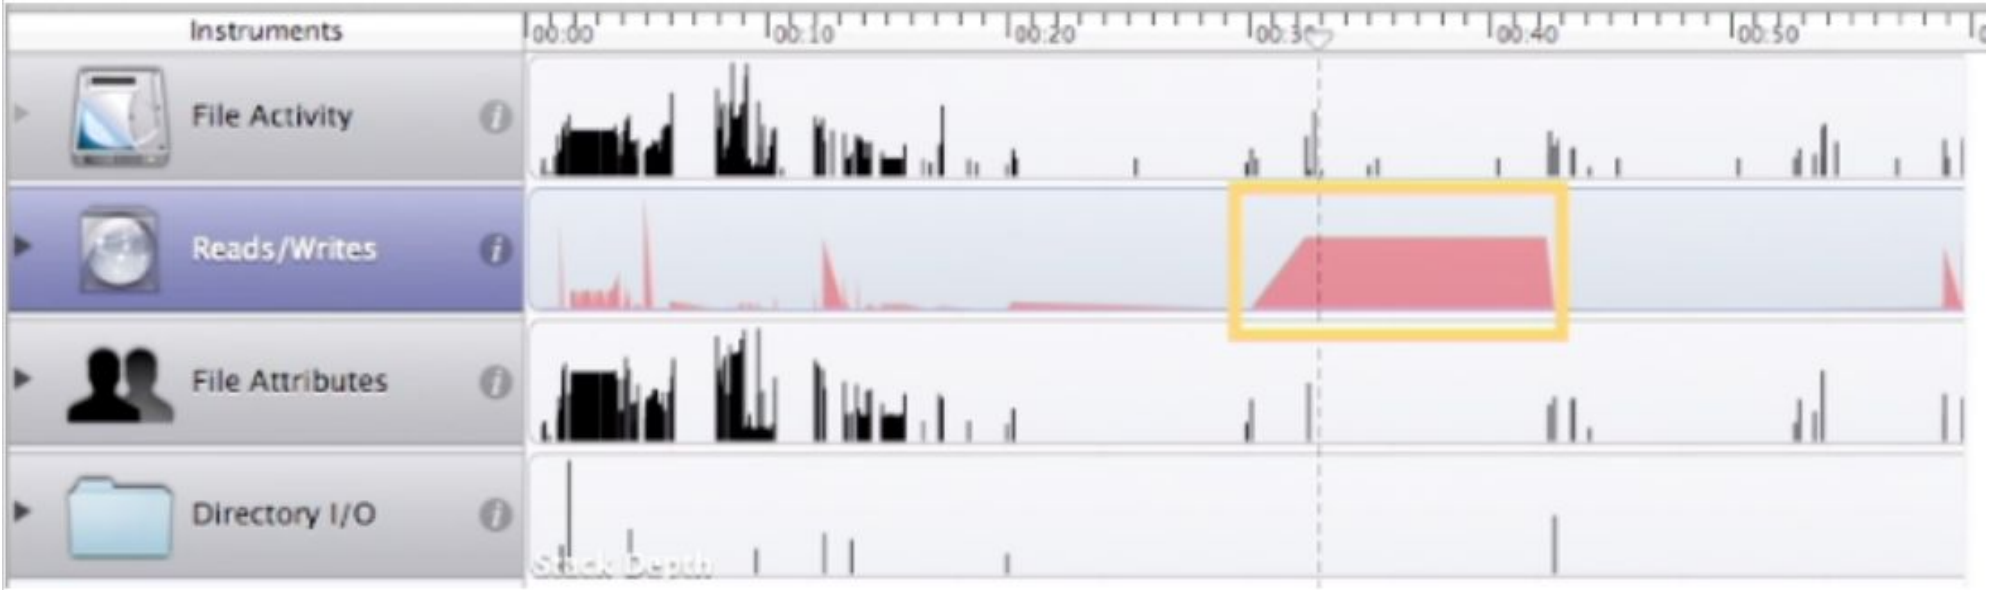

I/O Performance

HDD vs. SSD

- Price, capacity, speed, durability, form factor, noise, …

SATA and NVMe (Non-Volatile Memory Express)

NVMe works with PCI Express (PCIe) to transfer data to and from SSDs.

NVMe: Enables rapid storage.

HDD : Large-capacity storage.

Flash memory (SSD) : mobile devices such as smartphones, tablets, USB drives and SD cards.

Use SSD?

Network Efficiency (Performance)

Minimize Network Latency

(at least, repeated amounts of)

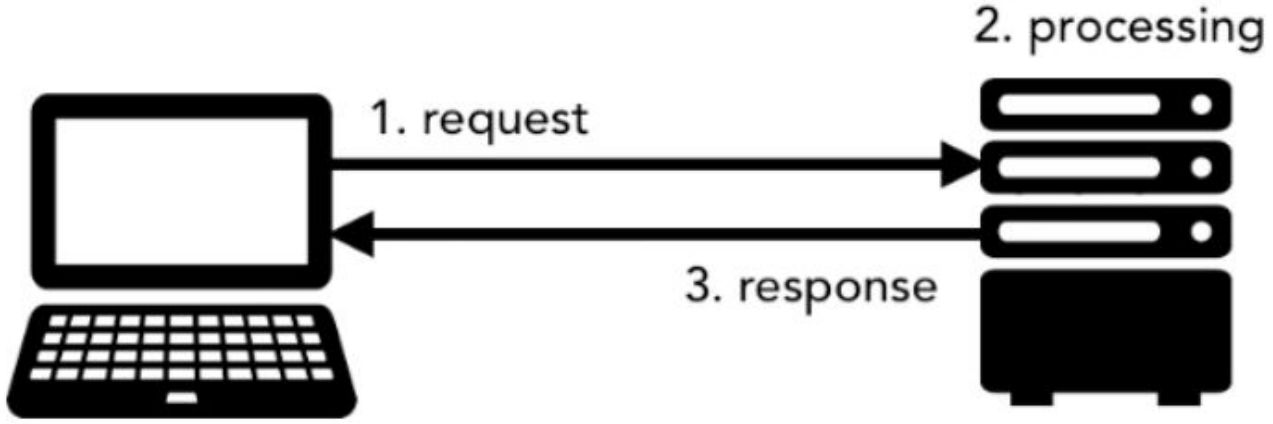

1 + 2 + 3 = latency (delay)

Network Performance

verizon media

- 10GbE/ 25/ 50 and 100GbE

  Latency - Throughput

- Intel FPGA Programmable Acceleration Card for Networking Application.
  - *Live video streaming, virtual and augmented reality (AR and VR).*
  - *5G applications*

Accelerators: GPU + FPGA

Deep Learning/ AI Projects:

- User behavior ranking features
- Ad Click Prediction
- Personalization

Access to Data

- Search: next generation mobile search with AI assistant
- Mail: Leverage Yahoo Database
- Publisher
- Ads and video

GPU + FPGA

- Parallel Processing
- architecture and programming language
- Customize the computer chip

Training + Inference

- GPU / FPGA

Inference:

Xeon CPU

FPGA

Software Performance

- Monitoring
- Benchmarking vs. Profiling
- Profiling Tools
- Real Issue: Solved
- Performance Tools

Algorithmic efficiency

"Could these statements have been written to accomplish exactly the same result, but do it faster?"

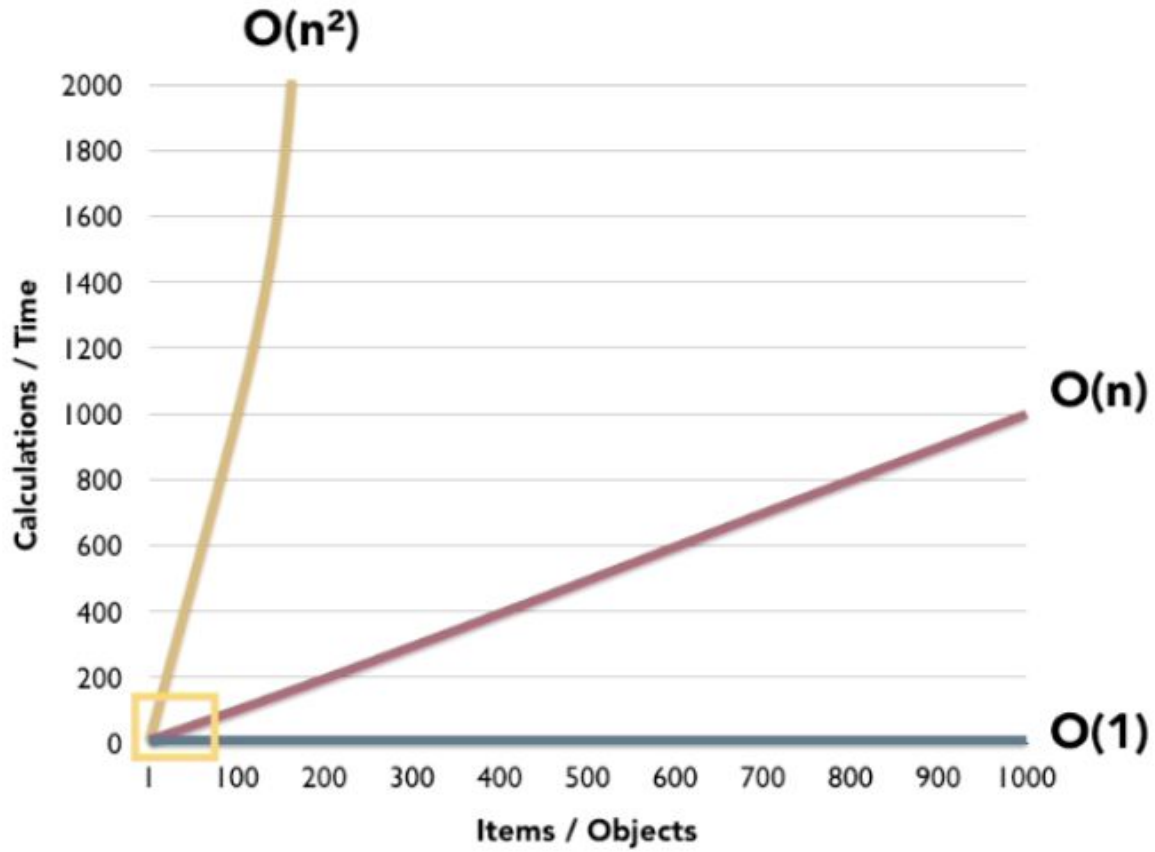

Performance

- Monitoring
- Profiling
- Benchmarking
- … simulating, emulating …

Benchmarking or profiling?

- **Benchmarking** is evaluating the performance of a function, operation relative to others.

Benchmarking = competing

- **Profiling** doesn't compare multiple versions of the same operation/function, it tells us where the slowest parts of our code base are.

Profiling = observation

Profiling

- Find and eliminate waste, wait time, ...
- Find areas to tune, …
- Understand system limits
- Solve performance issues
- Locate bottlenecks and latency
- Build scalable/modular architectures (system/hardware)

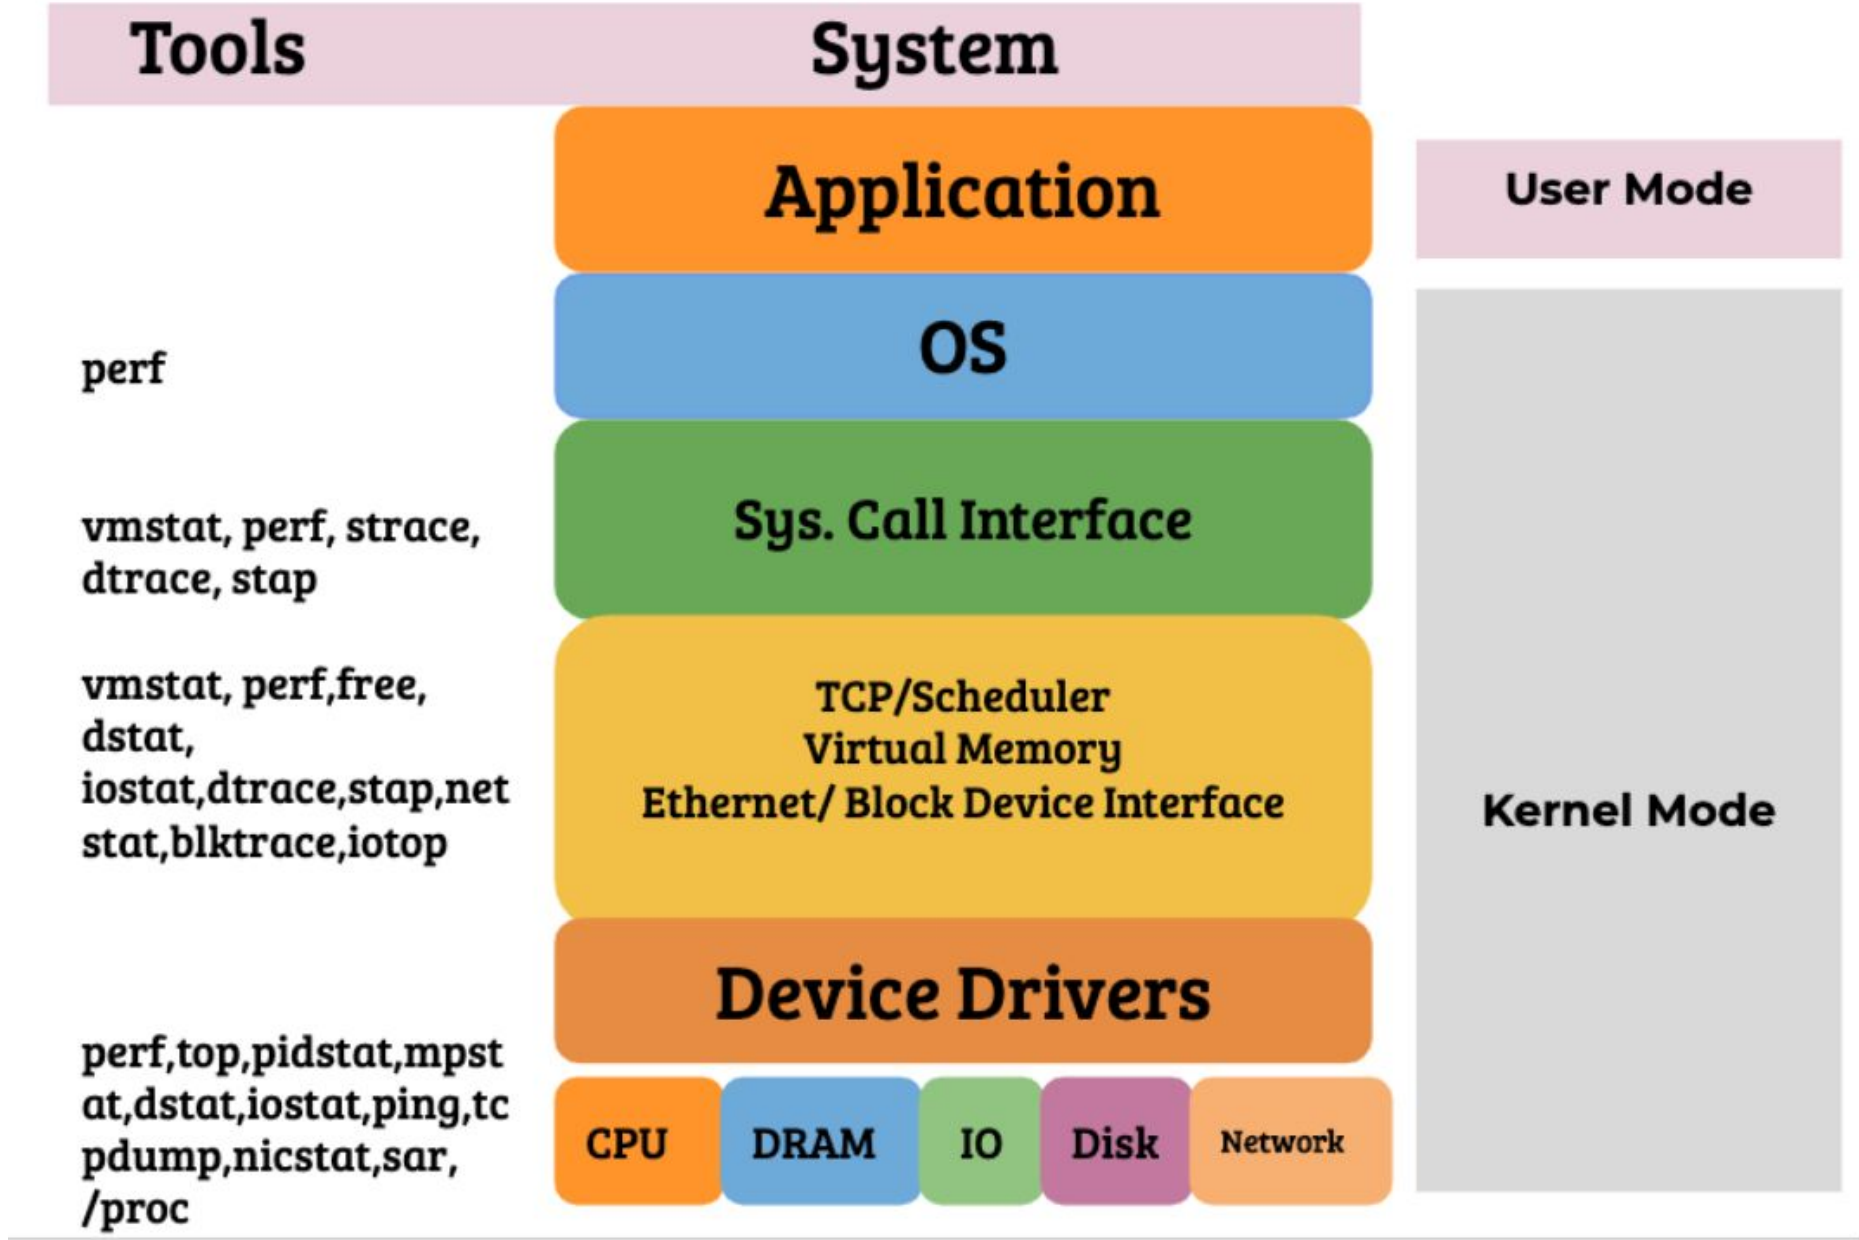
Tools
System
Application
User Mode
OS
perf
Sys. Call Interface
vmstat, perf, strace,
dtrace, stap
TCP/Scheduler
Virtual Memory
Ethernet/ Block Device Interface
vmstat, perf,free,
dstat,
iostat,dtrace,stap,net
stat,blktrace,iotop
Kernel Mode
Device Drivers
perf,top,pidstat,mpst
at,dstat,iostat,ping,tc
pdump,nicstat,sar,
/proc
CPU
DRAM
IO
Disk
Network

Intel VTune Amplifier

Mourad Bouache

https://www.youtube.com/watch?v=cjiUXZh0ZCo

HPC Seminar Series ~ Intel VTune™ Training

Stanford High Performance Computing Center

Scale Out Performance

Challenges: Scalable horizontally or vertically

Increase throughput

Scale vertically: More CPU #core, memory?

> Cost

Drive the company Capacity Planning:

- Plan the largest server and datacenter capacity
- Data Center expansion plan

Cloud Performance

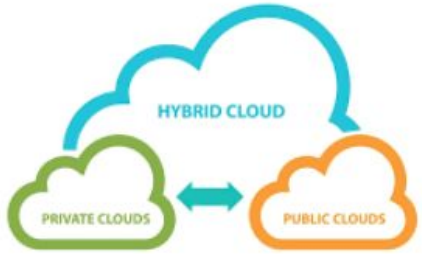

Virtualized compute, network, and storage resources

- Private
- Public
- Hybrid: Combines workloads.

Performance:

Latency, Memory and I/O bottlenecks, SLA, Elasticity and Security

seasonal capacity changes

Virtualization and Containerization

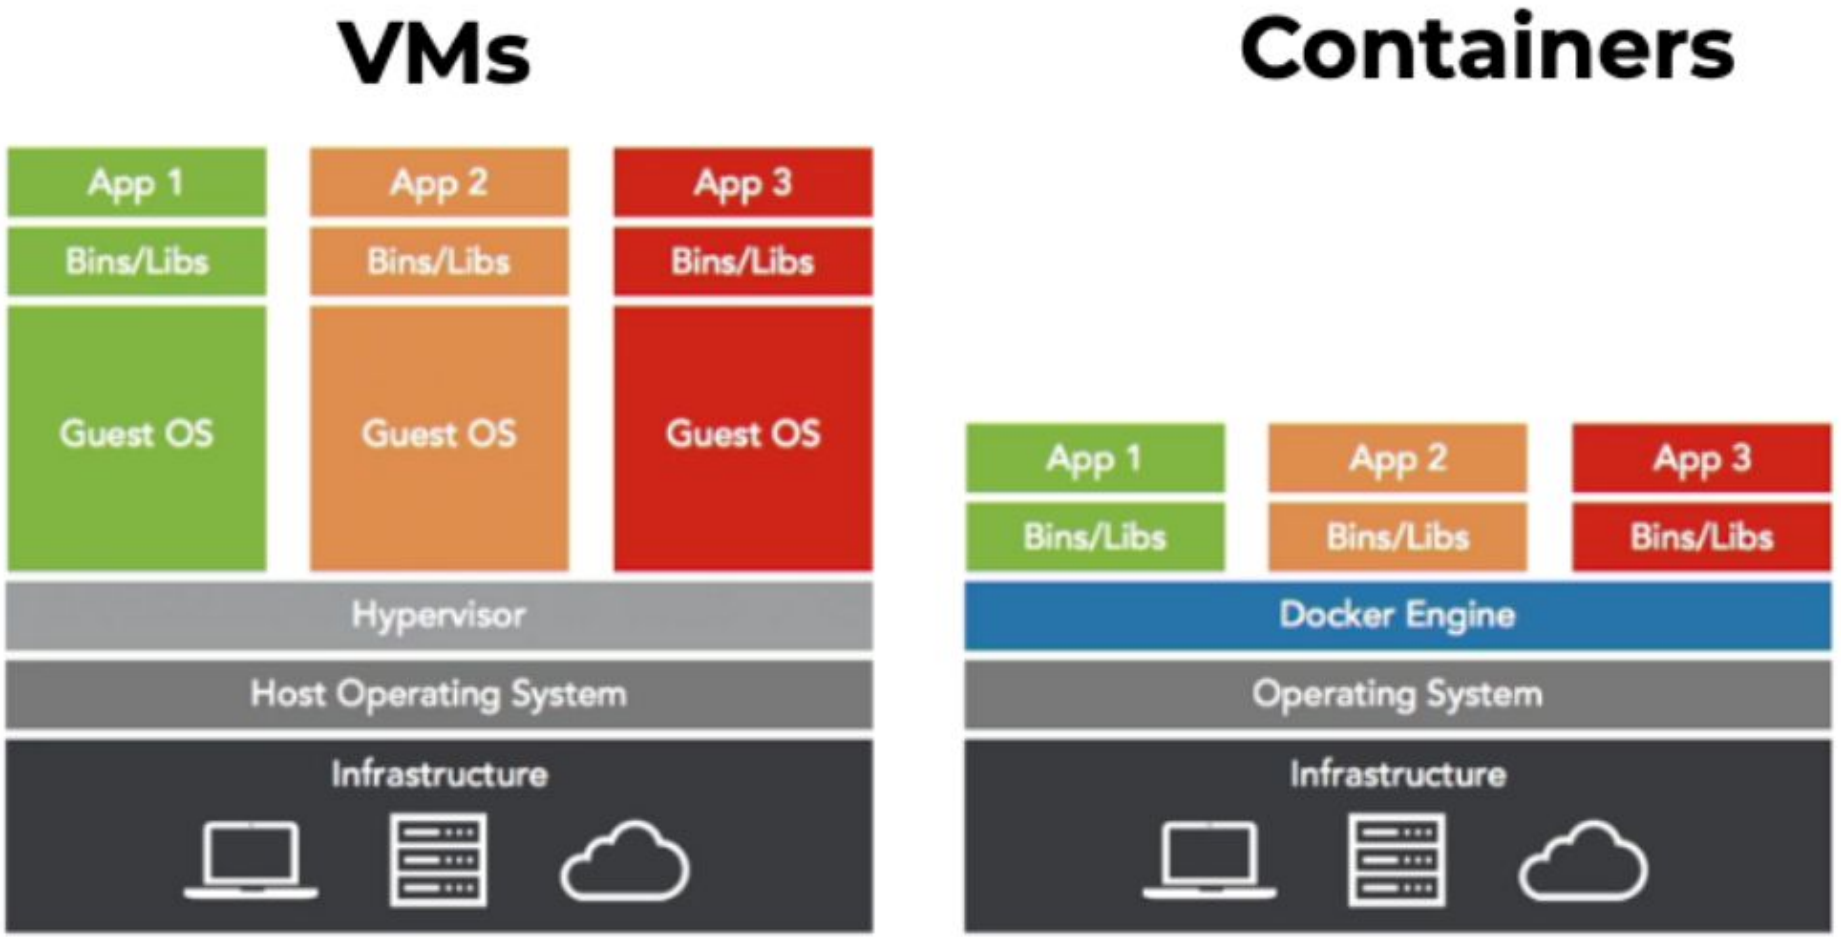

Performance

Value	Containers	Virtual Machines
Boot speed	✓	
Footprint size	✓	
Maturity		✓
Security		✓
Ease of patching	✓	
Developer agility	✓	

Containers within VMs give you the best of both worlds!

Performance Eng. + Artificial Intelligence

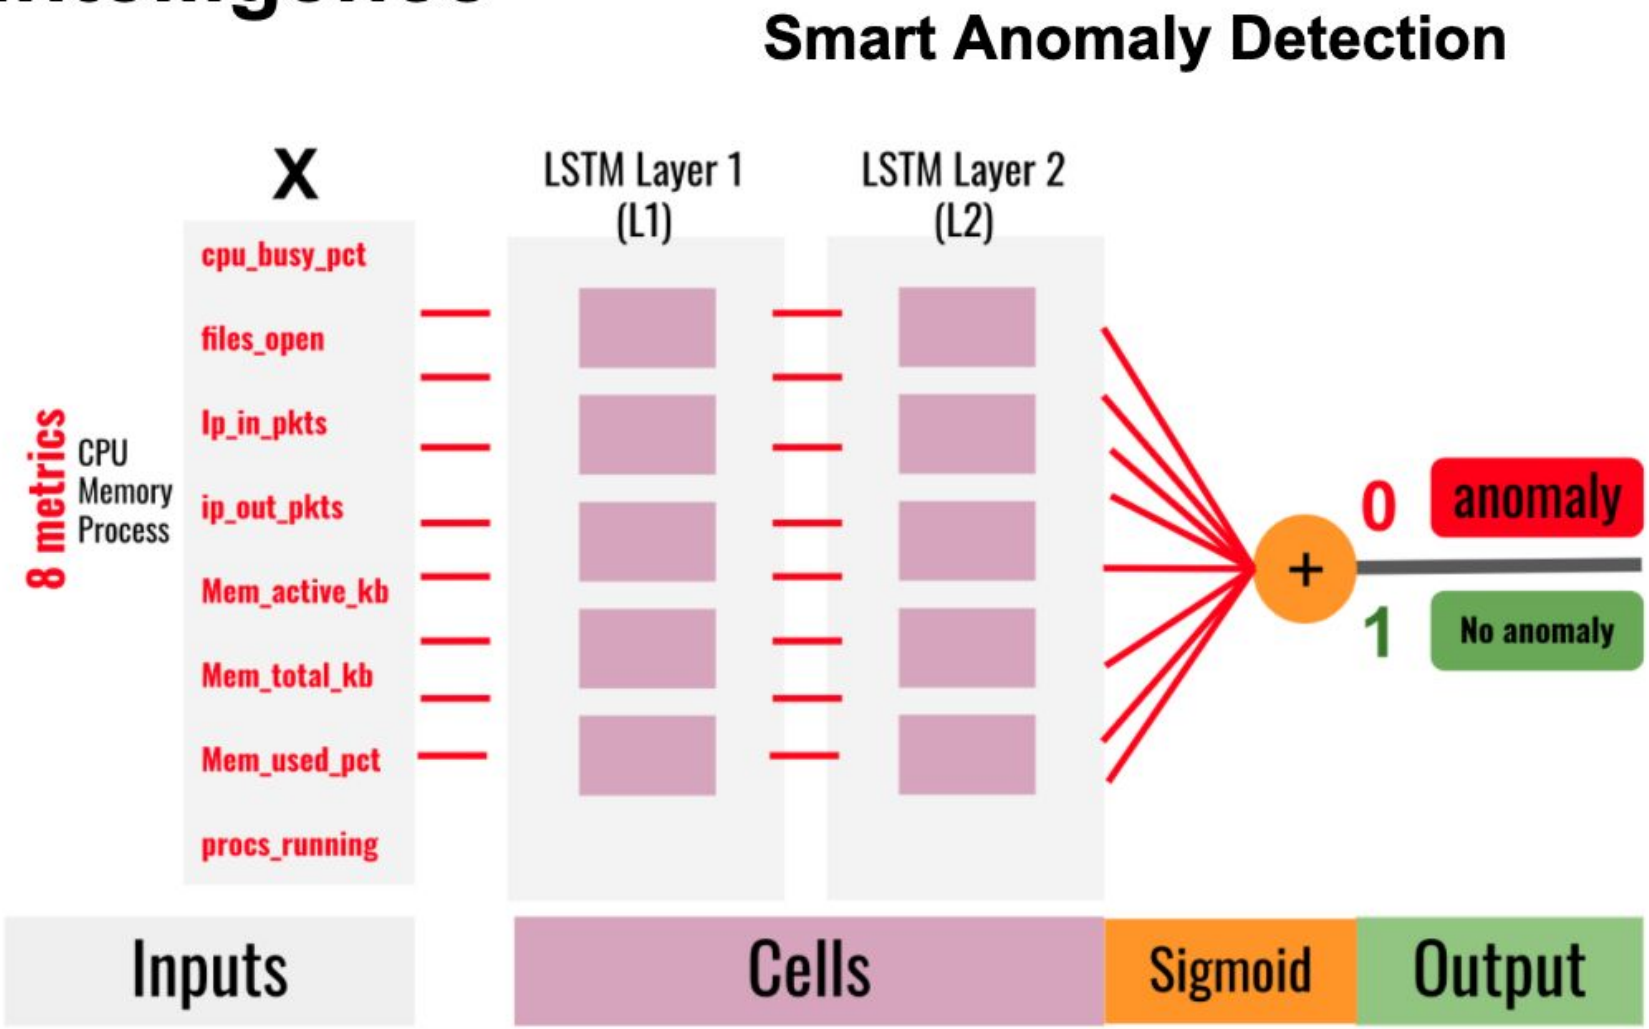

Need help!

If you are interested in **PERFORMANCE** contact:

- bouache@verizonmedia.com

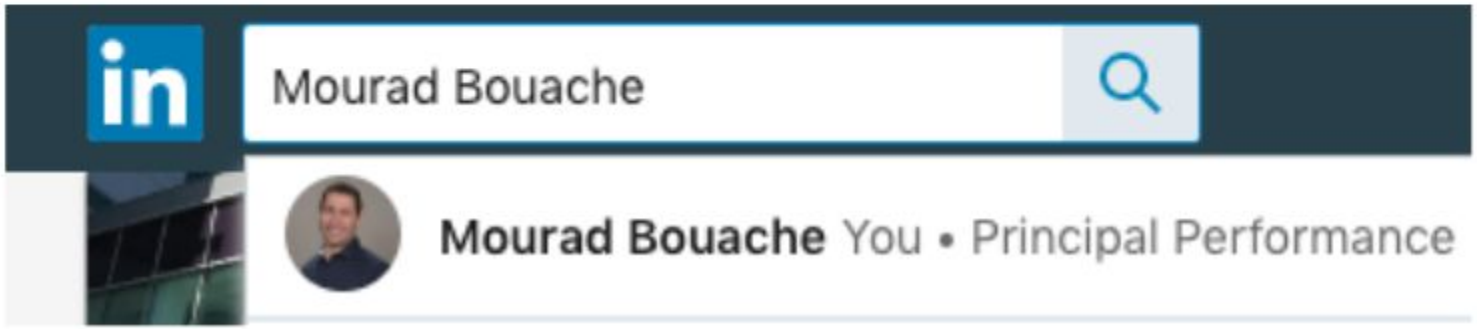

Thank you!