RED HAT SUMMIT

BOSTON, MA
JUNE 23-26, 2015

Containers versus Virtualization

Jeremy Eder, Principal Performance Engineer
Scott Herold, Principal Product Manager, RHEV
June 2015

redhat.

Agenda

- Venting
- Tech Overview
- Workload Classification
- Cold War ?
- Performance Data Roundup

TOP 5 MISCONCEPTIONS ABOUT CONTAINERS

1. Containers are new.
2. Containers equal virtualization.
3. Containers are universally portable.
4. Containers are secure by default.
5. Containers are not enterprise-ready.

VIRTUALIZATION AND CONTAINERS

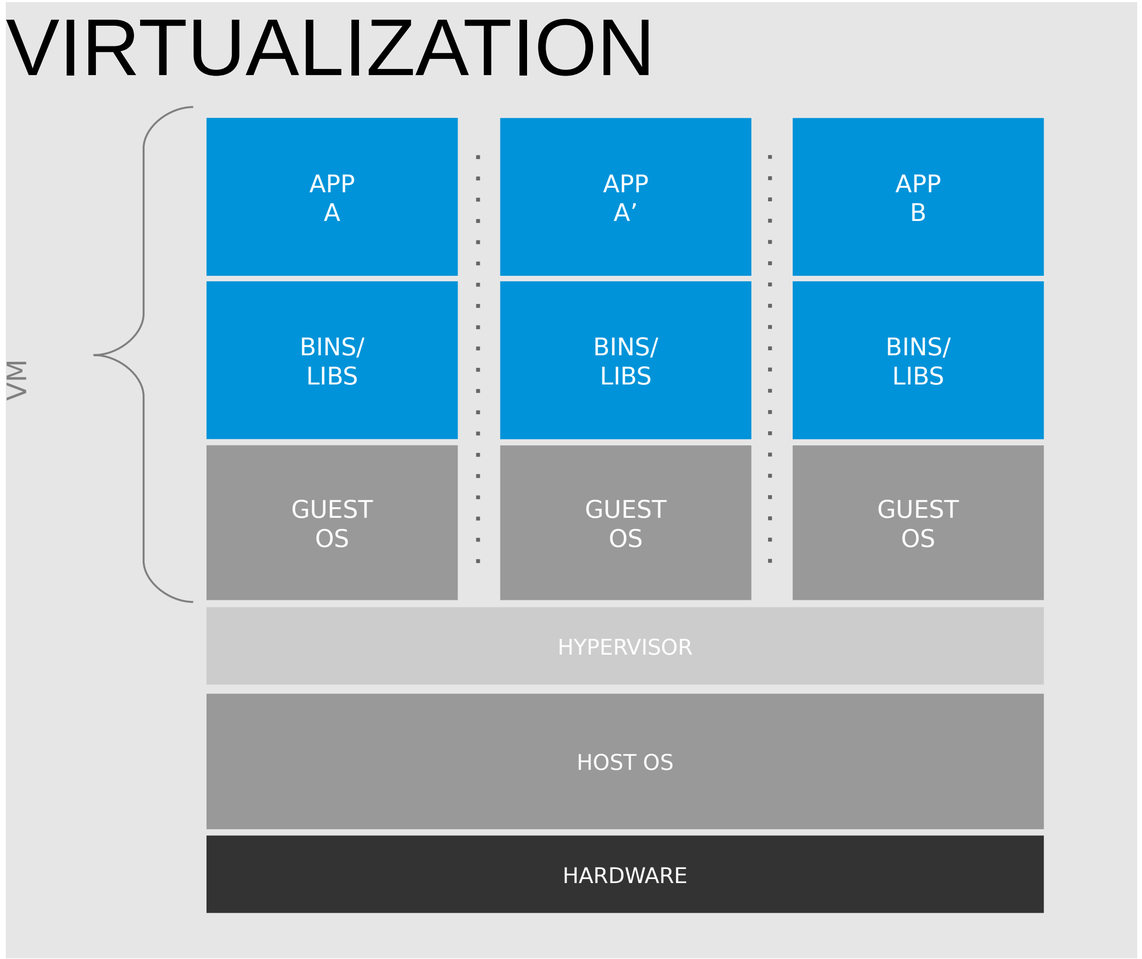

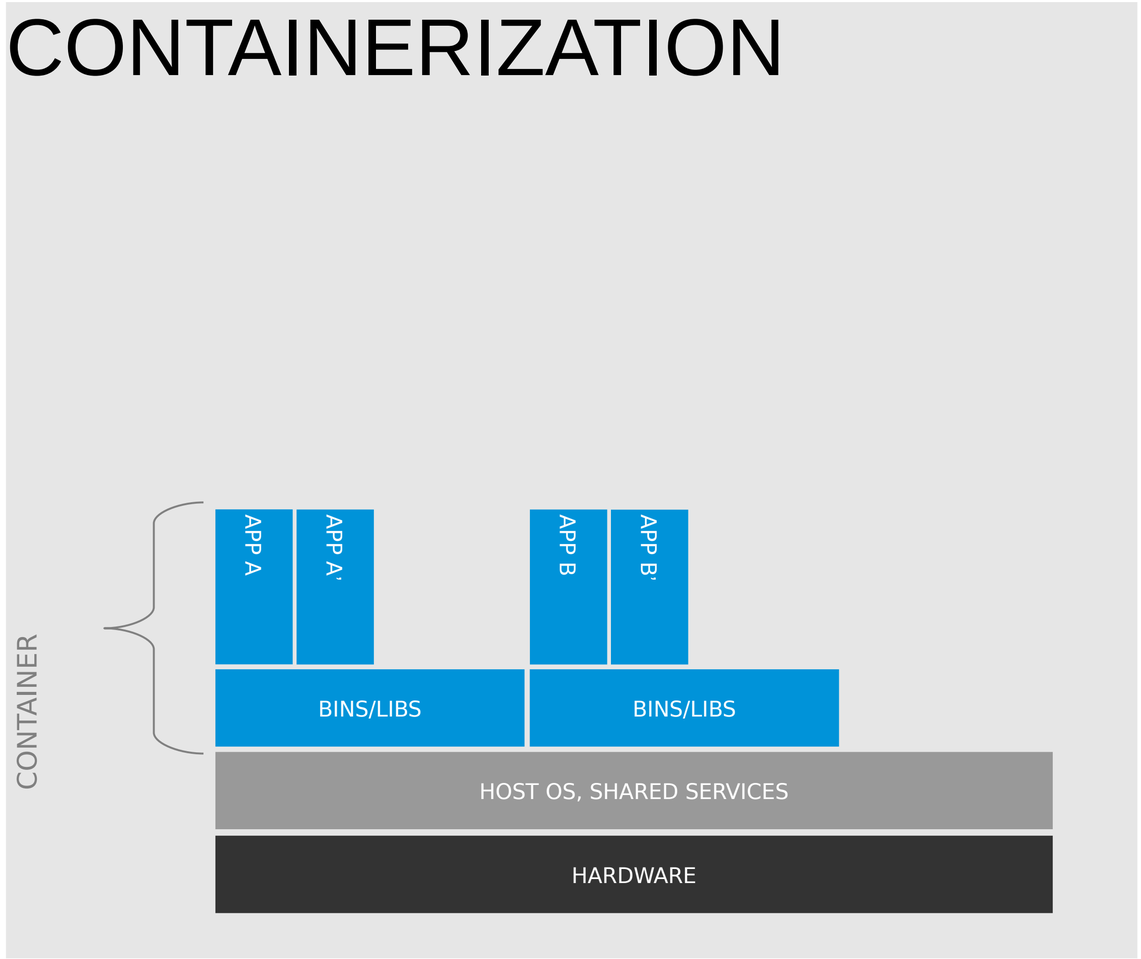

RED HAT ENTERPRISE VIRTUALIZATION

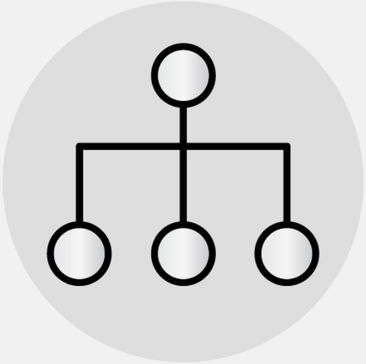

Centralized Management of KVM Hypervisor

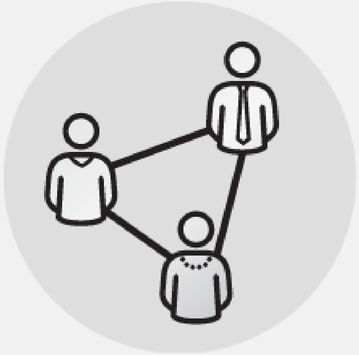

Self-Service User Portal

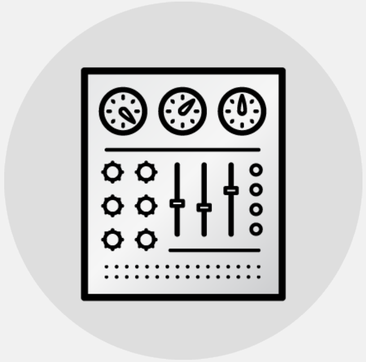

VM Workload Management

Differentiating Features

WHAT ABOUT DENSITY?

"For every VM, you can run 10 billion containers." -- Internet

How many containers will you run on one OS instance ?

- 1
- 10
- 50
- 100
- > 100

WHAT ABOUT DENSITY?

“This may be the most misleading stat ever.” -- Me

Get off my lawn!

Because it’s ALL about the workloads

- Some don’t care where they run
  - Batch workloads
- Some care greatly
  - Security/Isolation
  - Uptime
  - Performance

What is a workload? Subsystems

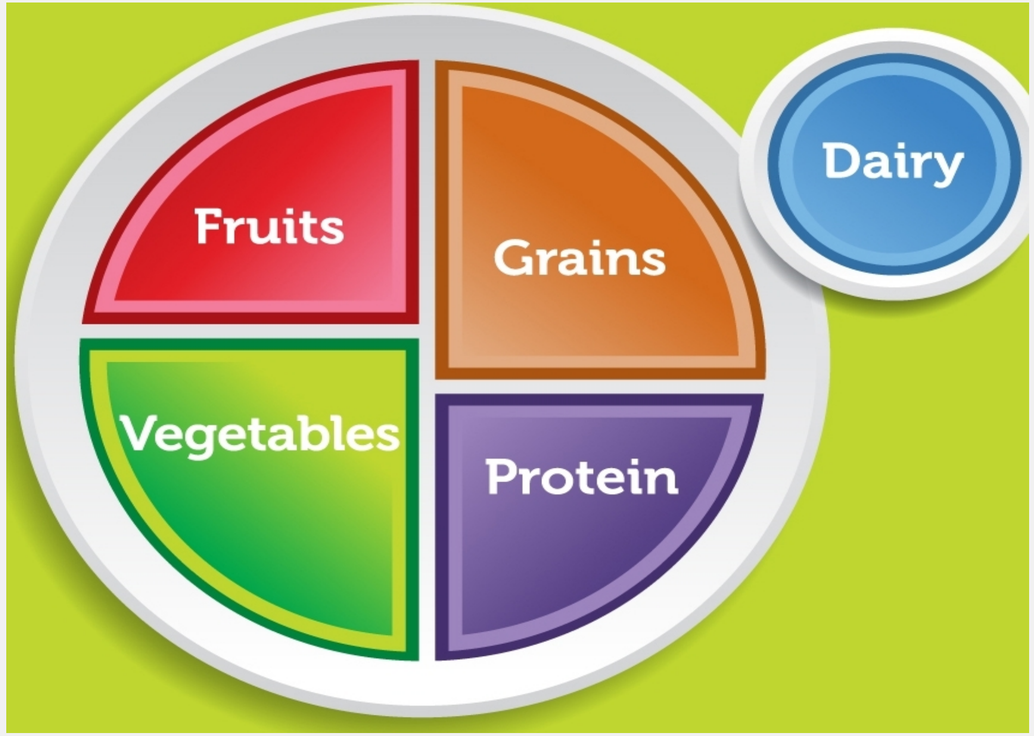

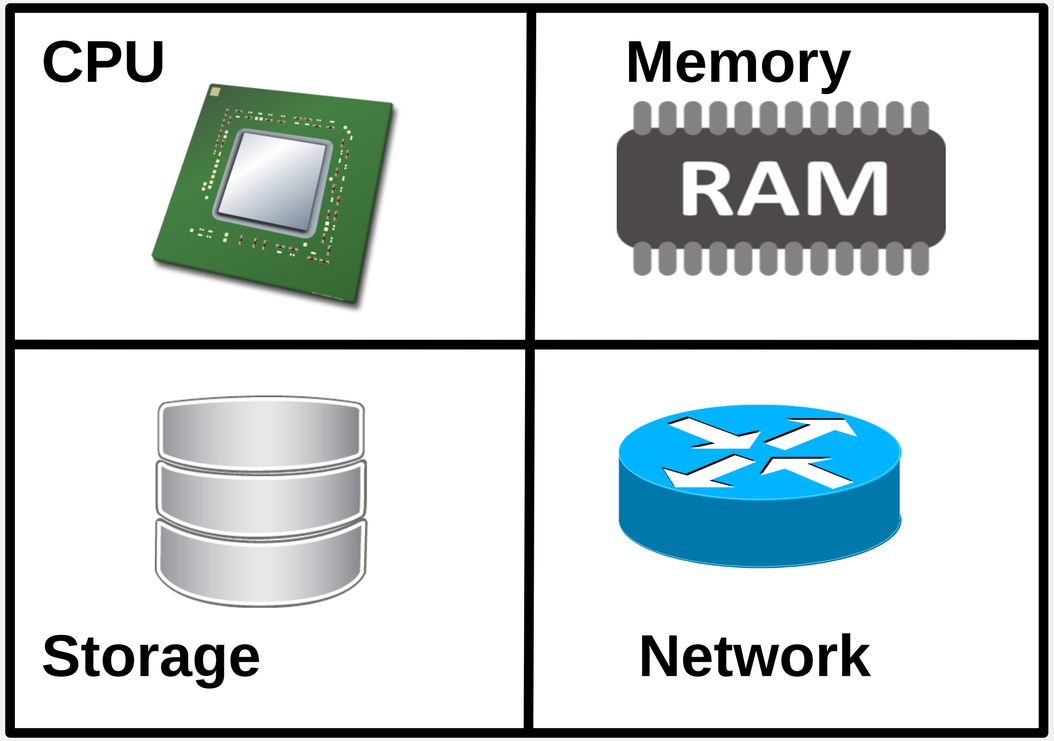

What is a workload? Requirements

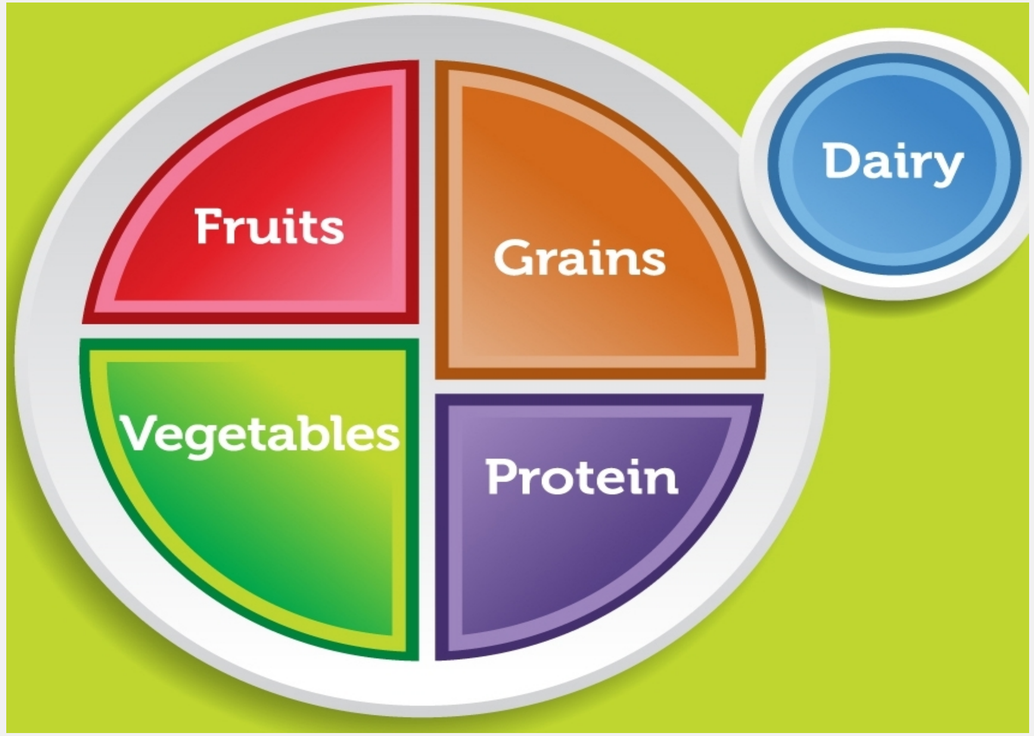

Security 	**Stability** 
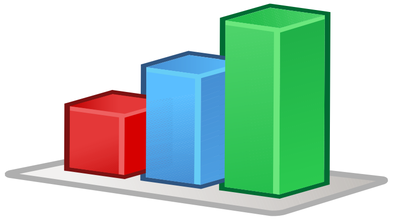 **Performance**	 **Support**

Culture, Control

I WANT CHANGE

Code Down (Dev)

versus

Infra Up (Ops)

I WANT STABILITY

WHEN WILL YOU MAKE SOMETHING THAT MATTERS?

WHEN WILL YOU MAKE SOMETHING COOL?

CONTAINERS VERSUS VIRTUALIZATION:

NEW COLD WAR ?

Minimum Overheads

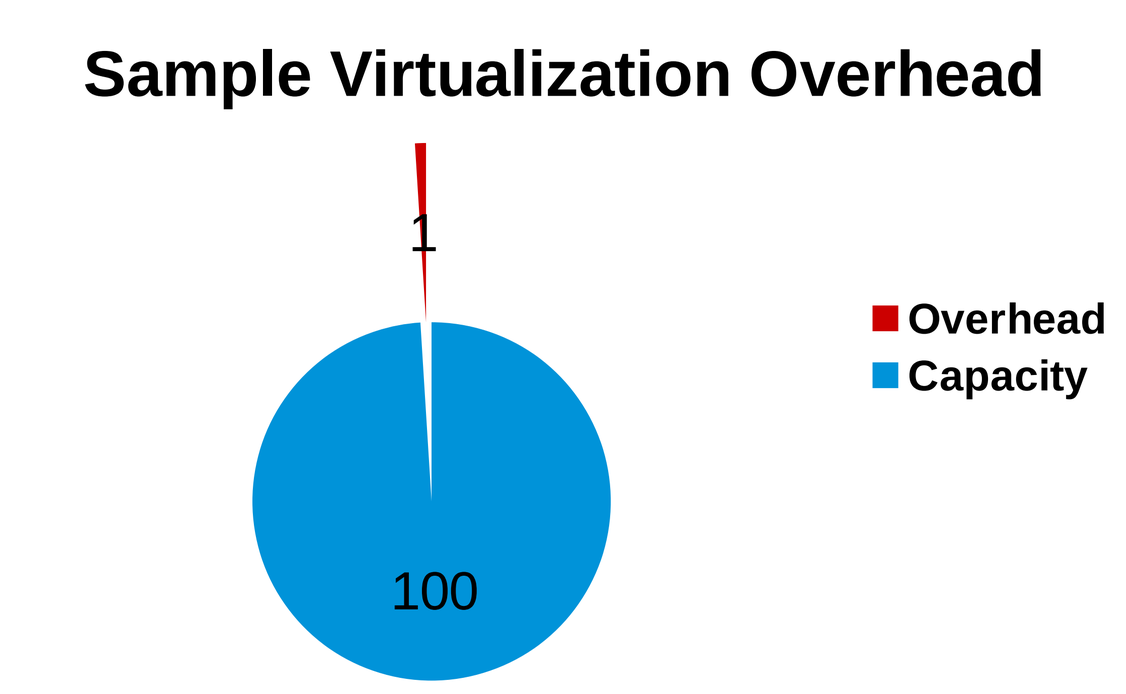

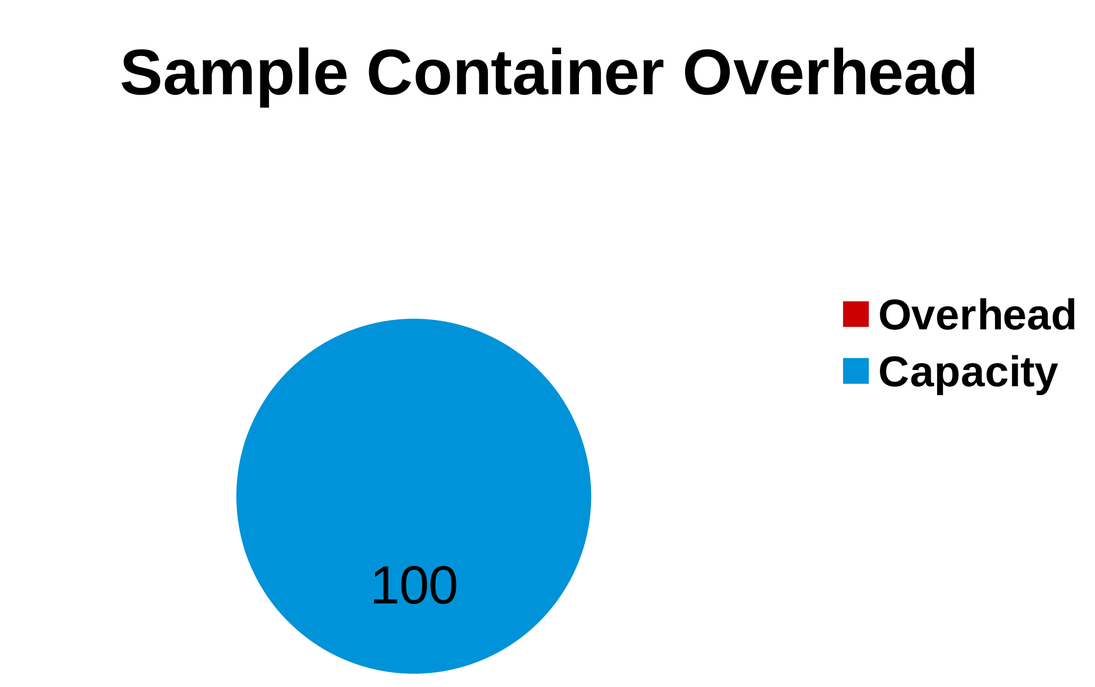

Maximum Overheads

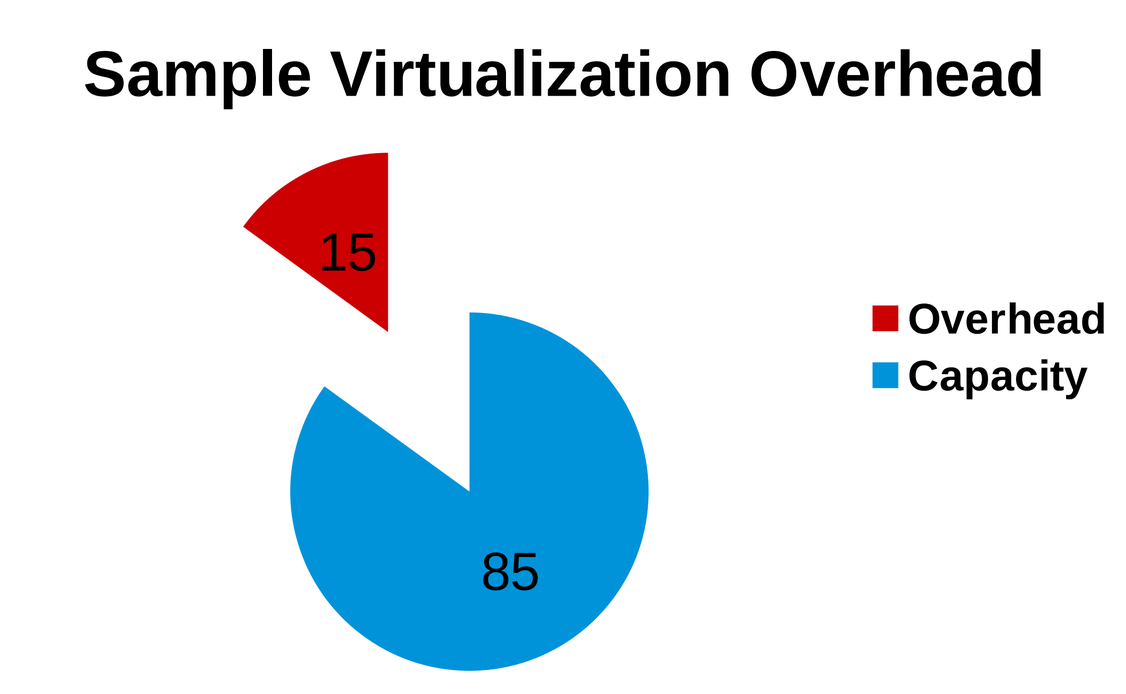

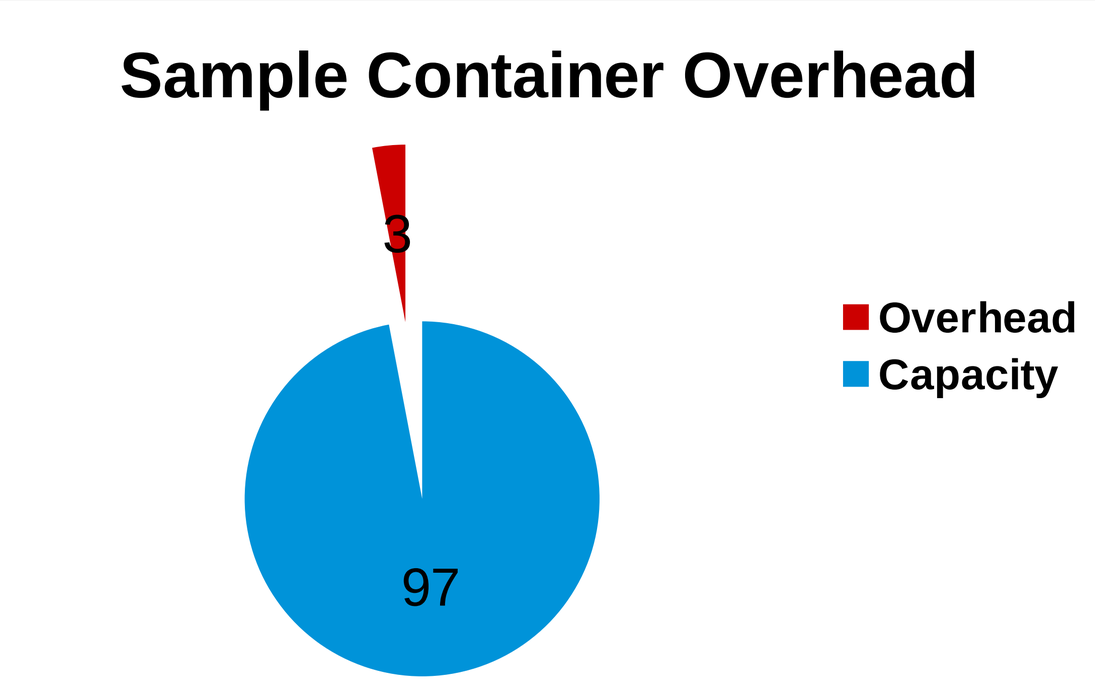

Reducing Overhead in VMs

Workload	Mitigation
CPU-intensive	•CPU Pinning •Avoid syscalls •Setup NUMA topology in-Guest
Memory-heavy	•Use hugepages •NUMA Pinning •Setup Hugepages in-Guest
Network (Latency)	•SR-IOV •PCI Passthrough •Busy Poll
Network (Throughput)	•Not normally an issue
Storage (Latency)	•Increase threads •virtio-blk-dataplane coming soon
Storage (Throughput)	•Not normally an issue

CONTAINERS VERSUS VIRTUALIZATION:

PERFORMANCE DATA ROUND-UP

CPU Tests: MPI LINPACK

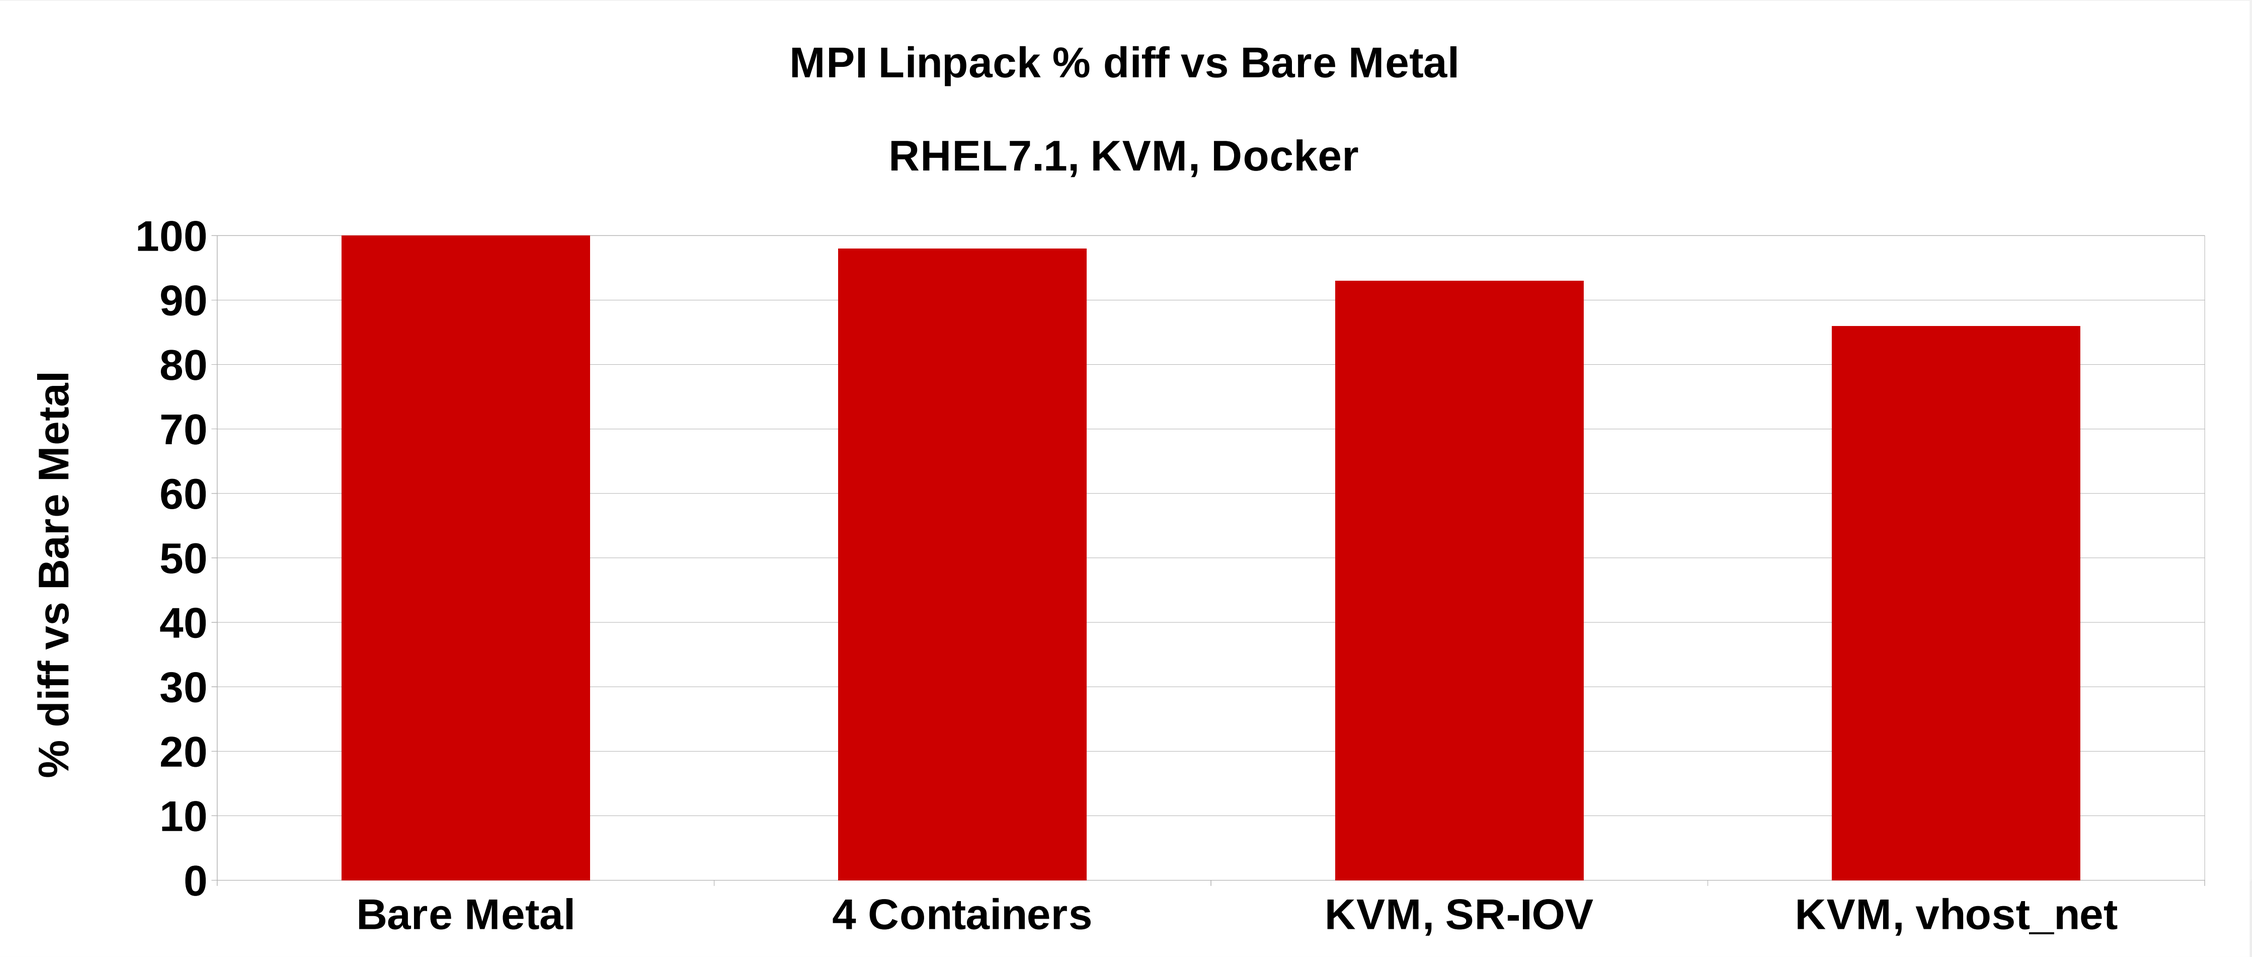

I/O Tests: fio … Bare Metal, KVM, Atomic, Docker

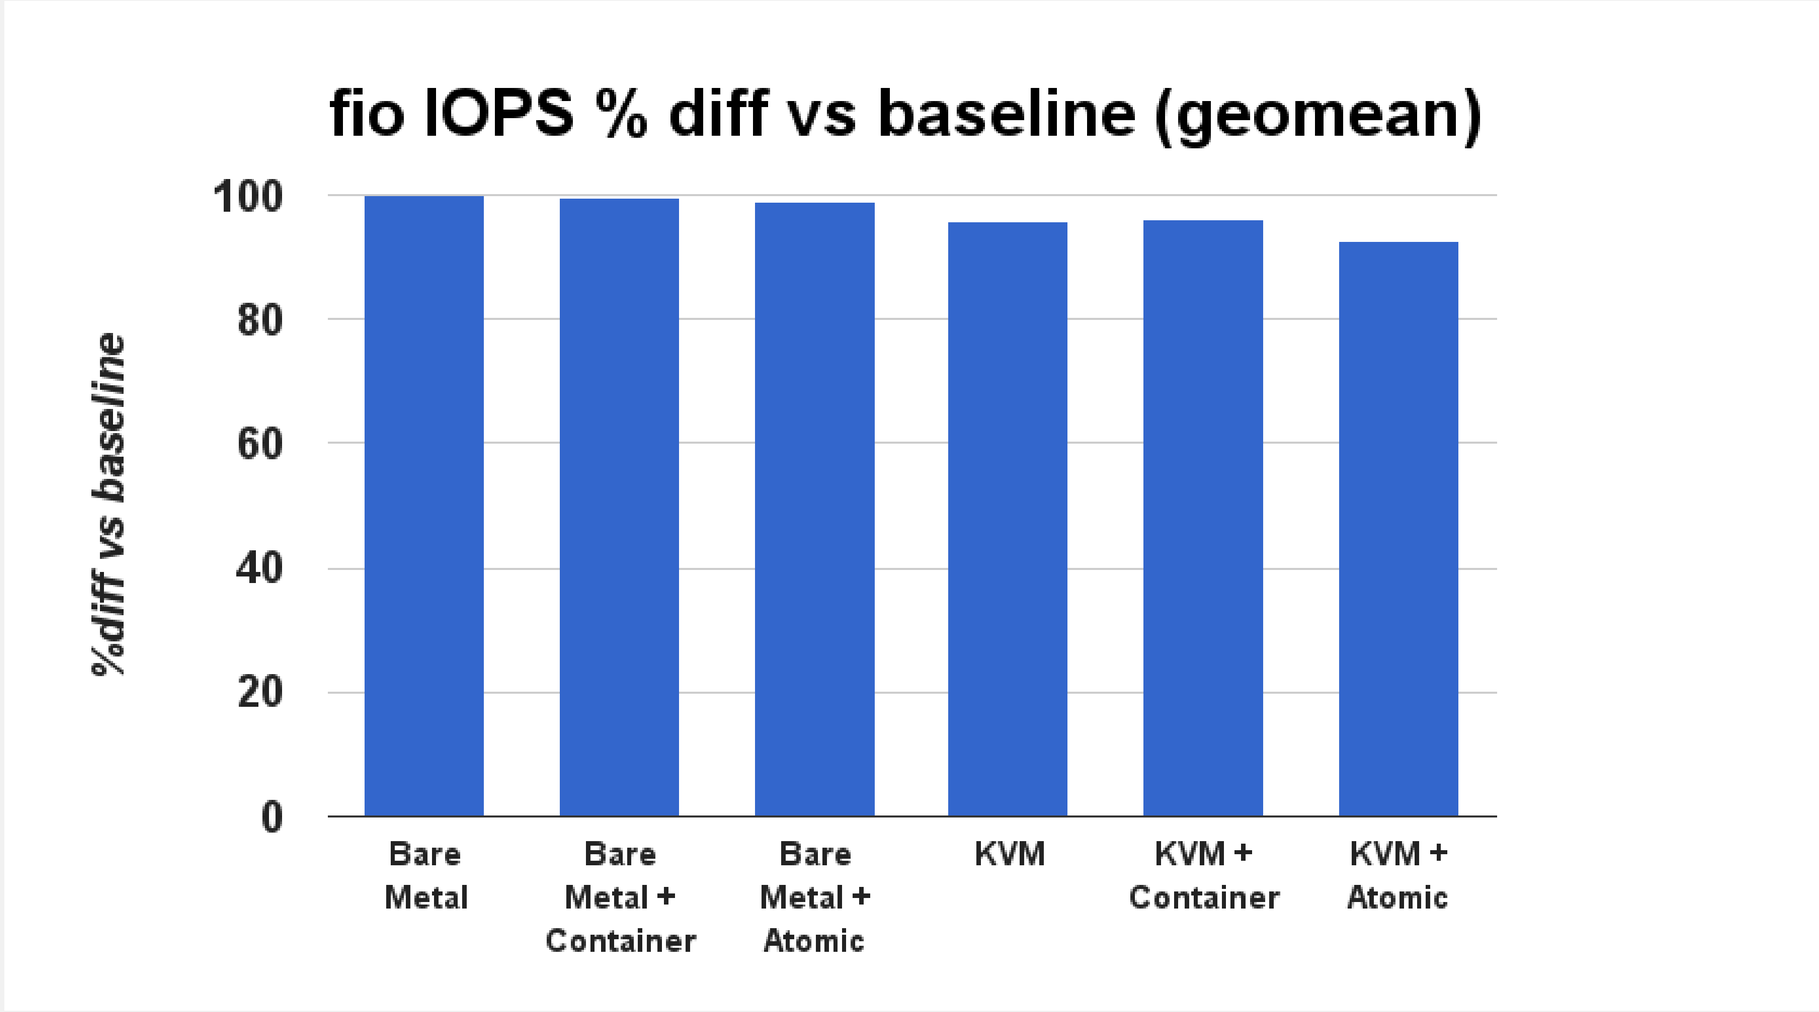

Application Tests: Business Analytics

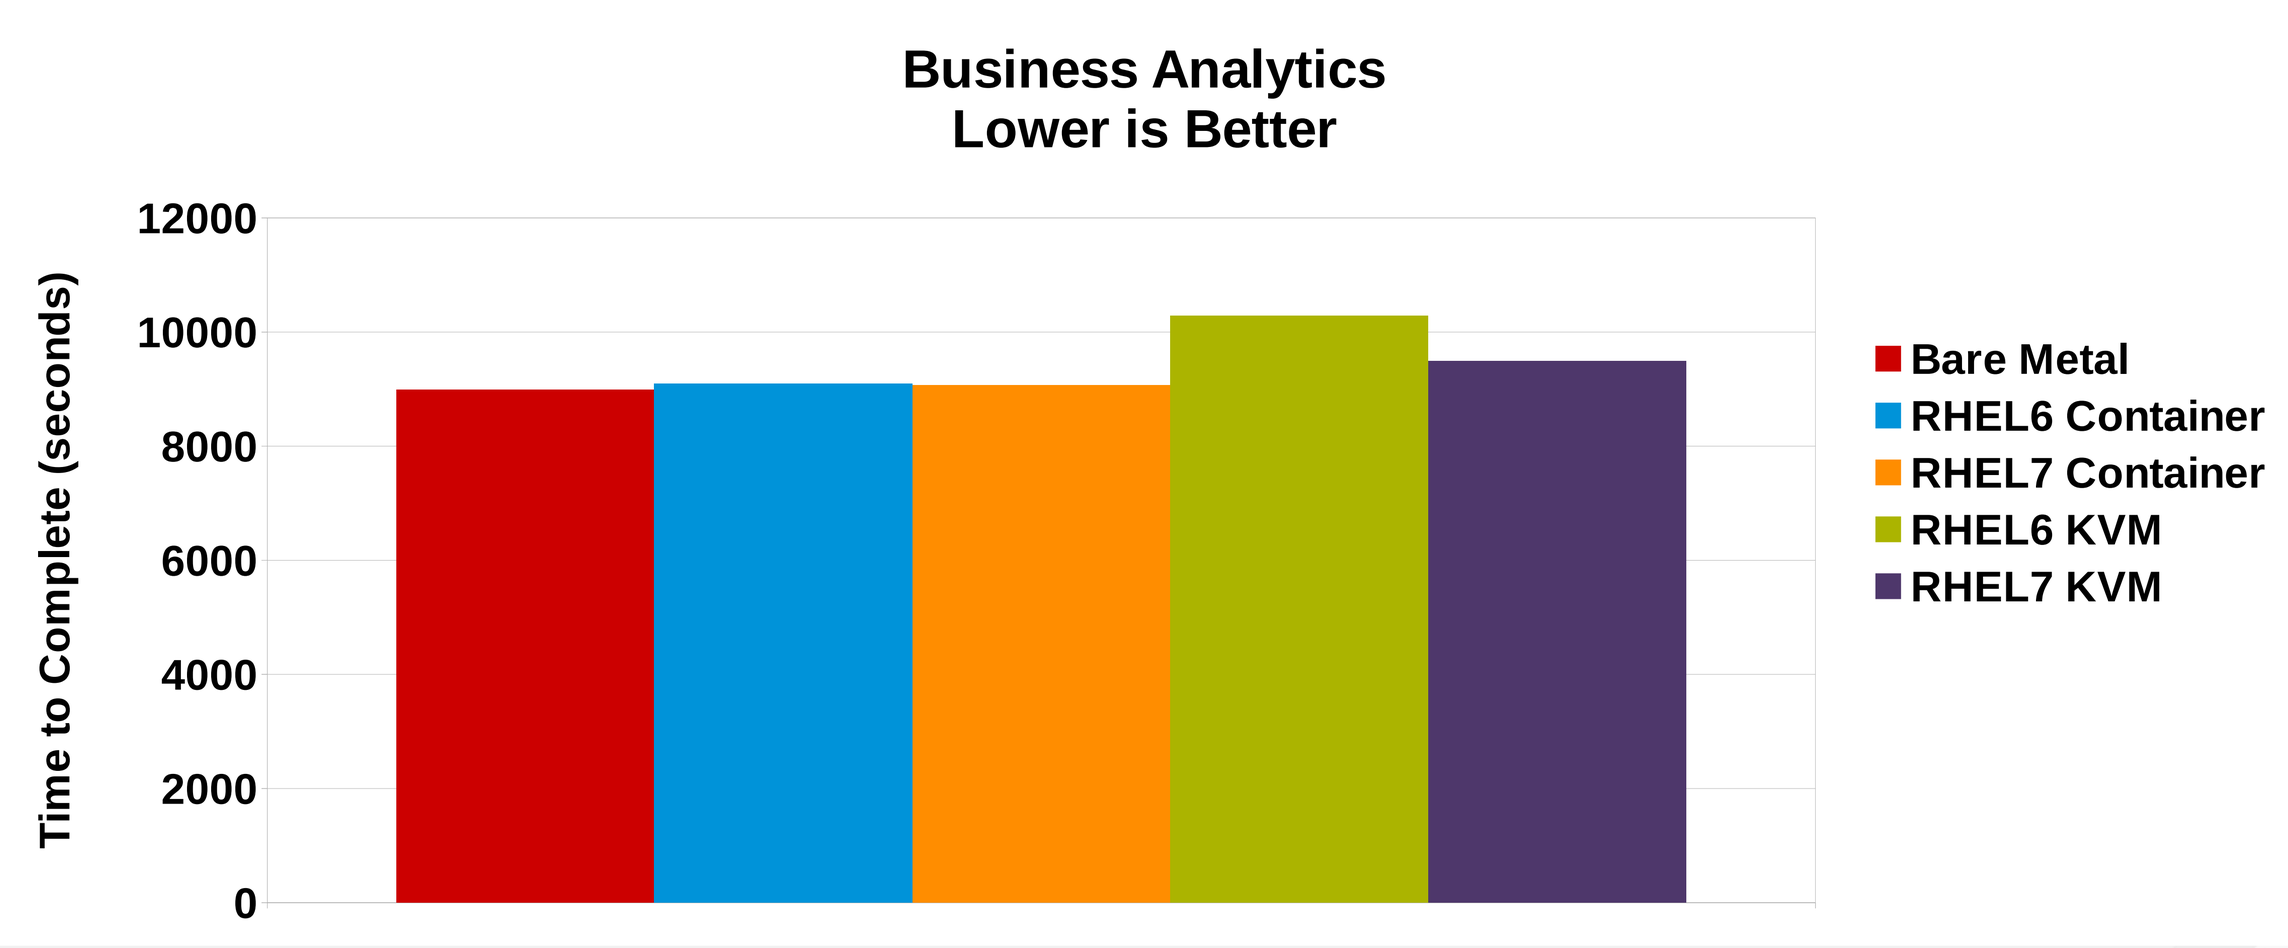

Network Latency and Throughput

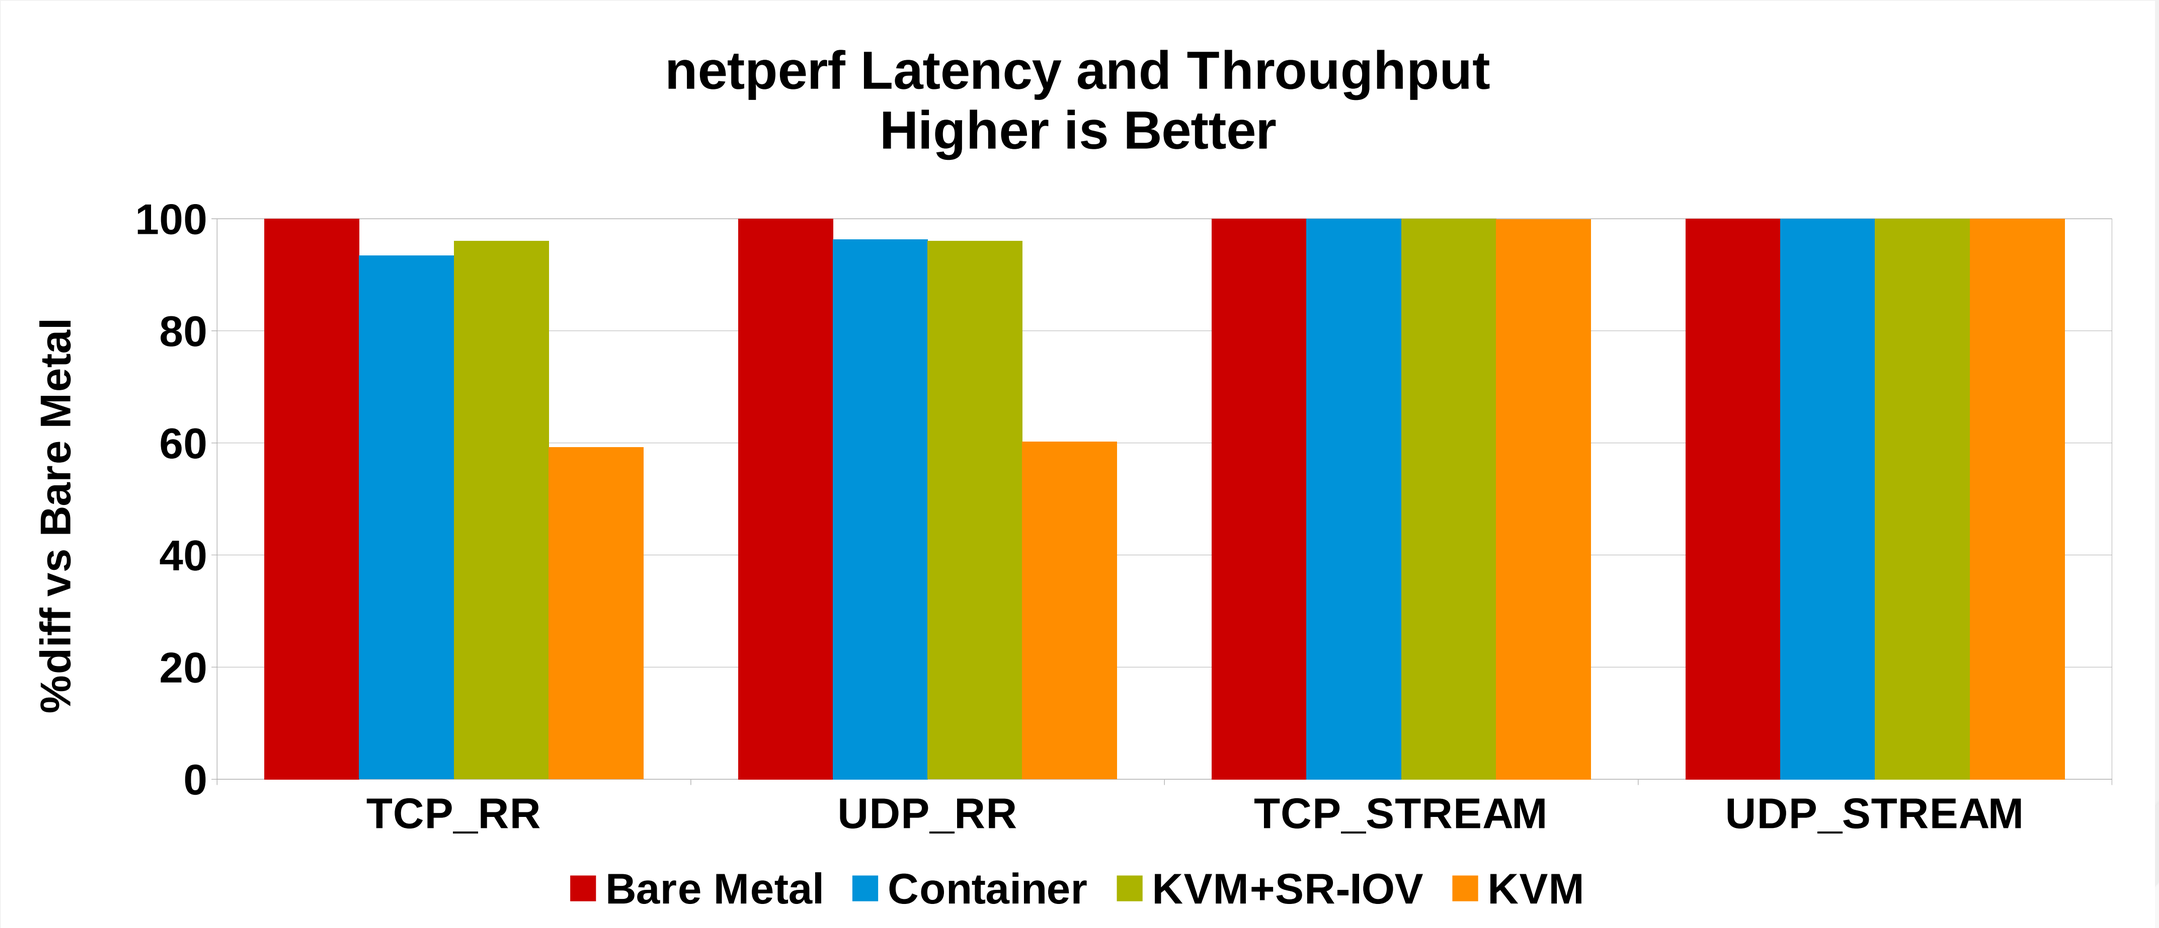

Large OLTP Database, BM vs Container vs KVM

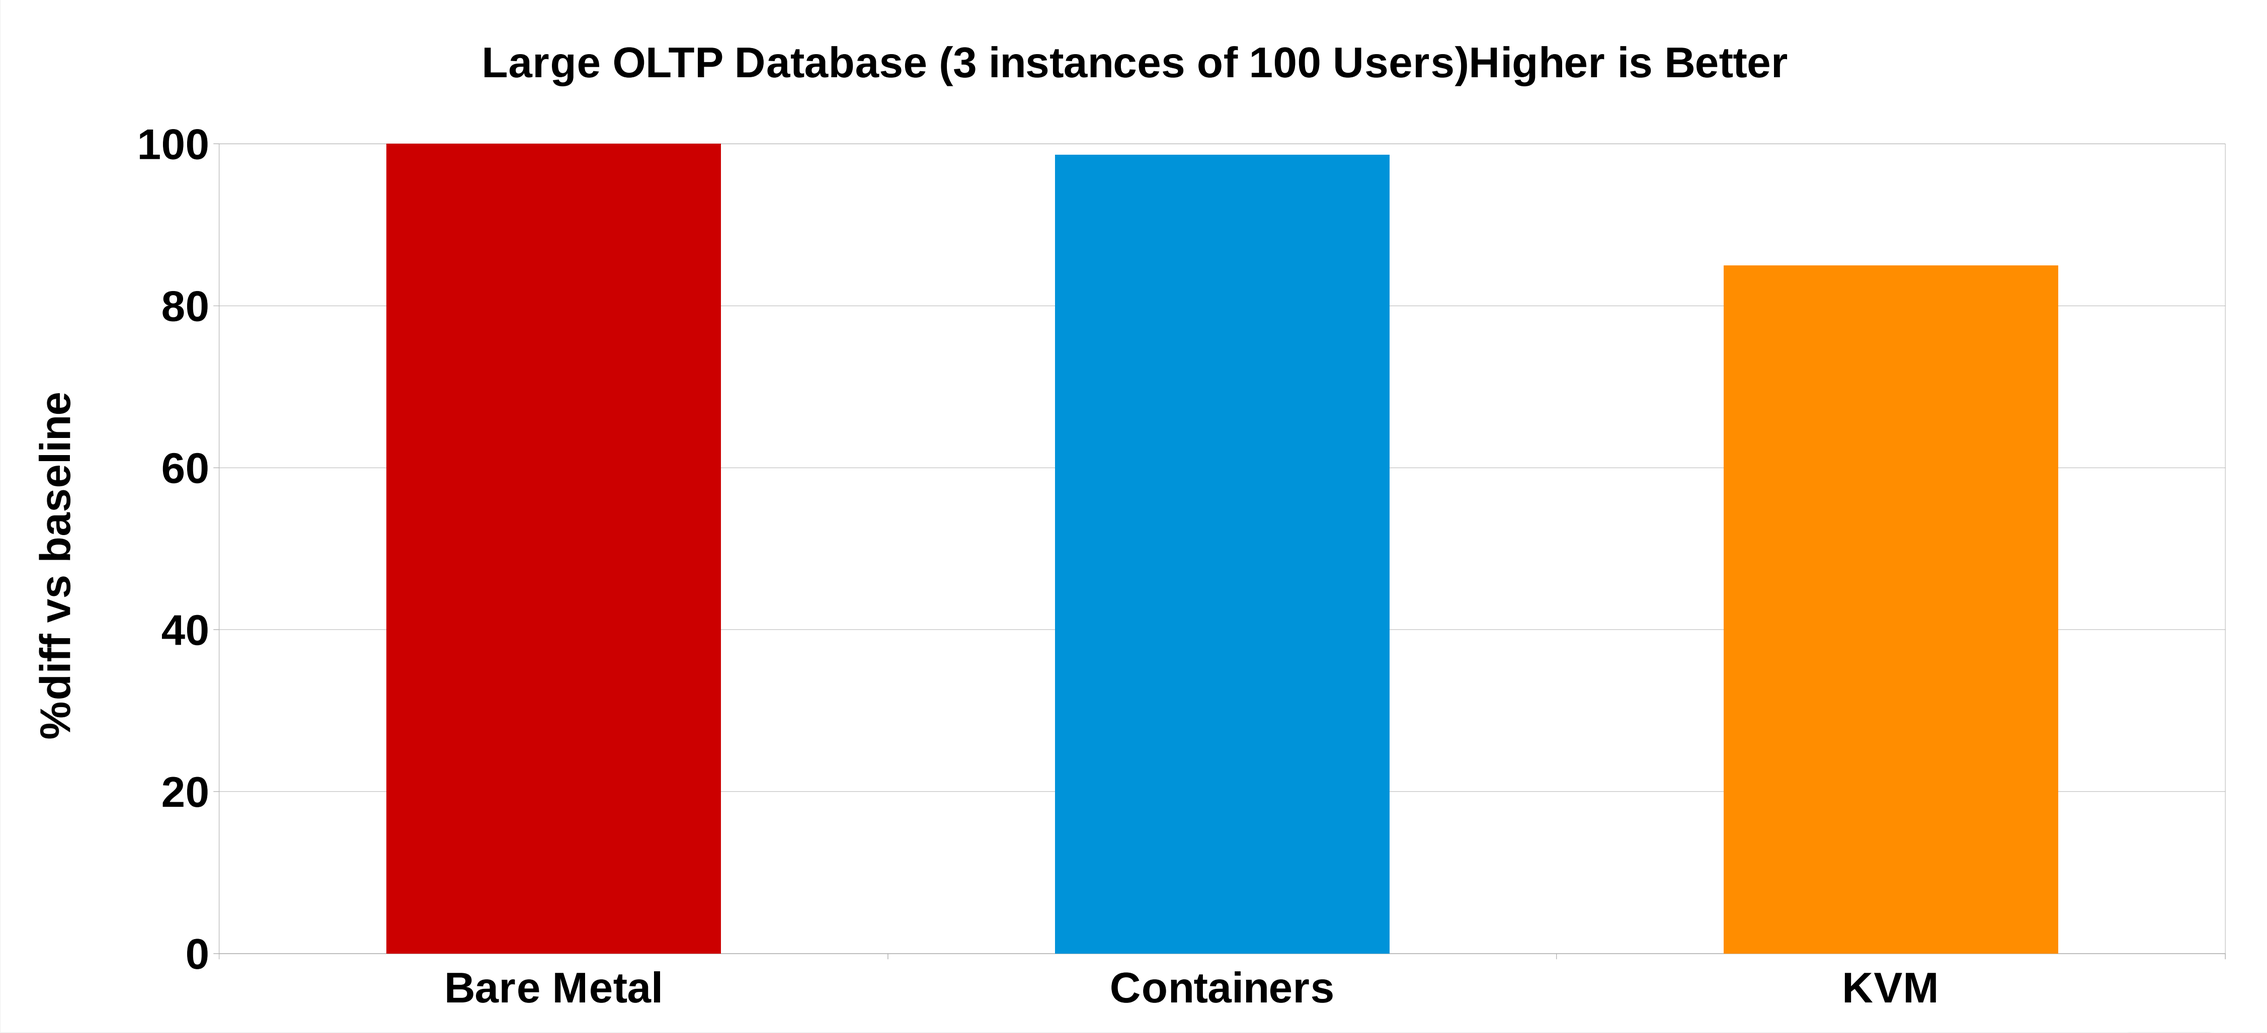

Network Function Virtualization (NFV) Throughput and Packets/sec (RHEL7.x+DPDK)

208Mpps+ INTO KVM DPDK

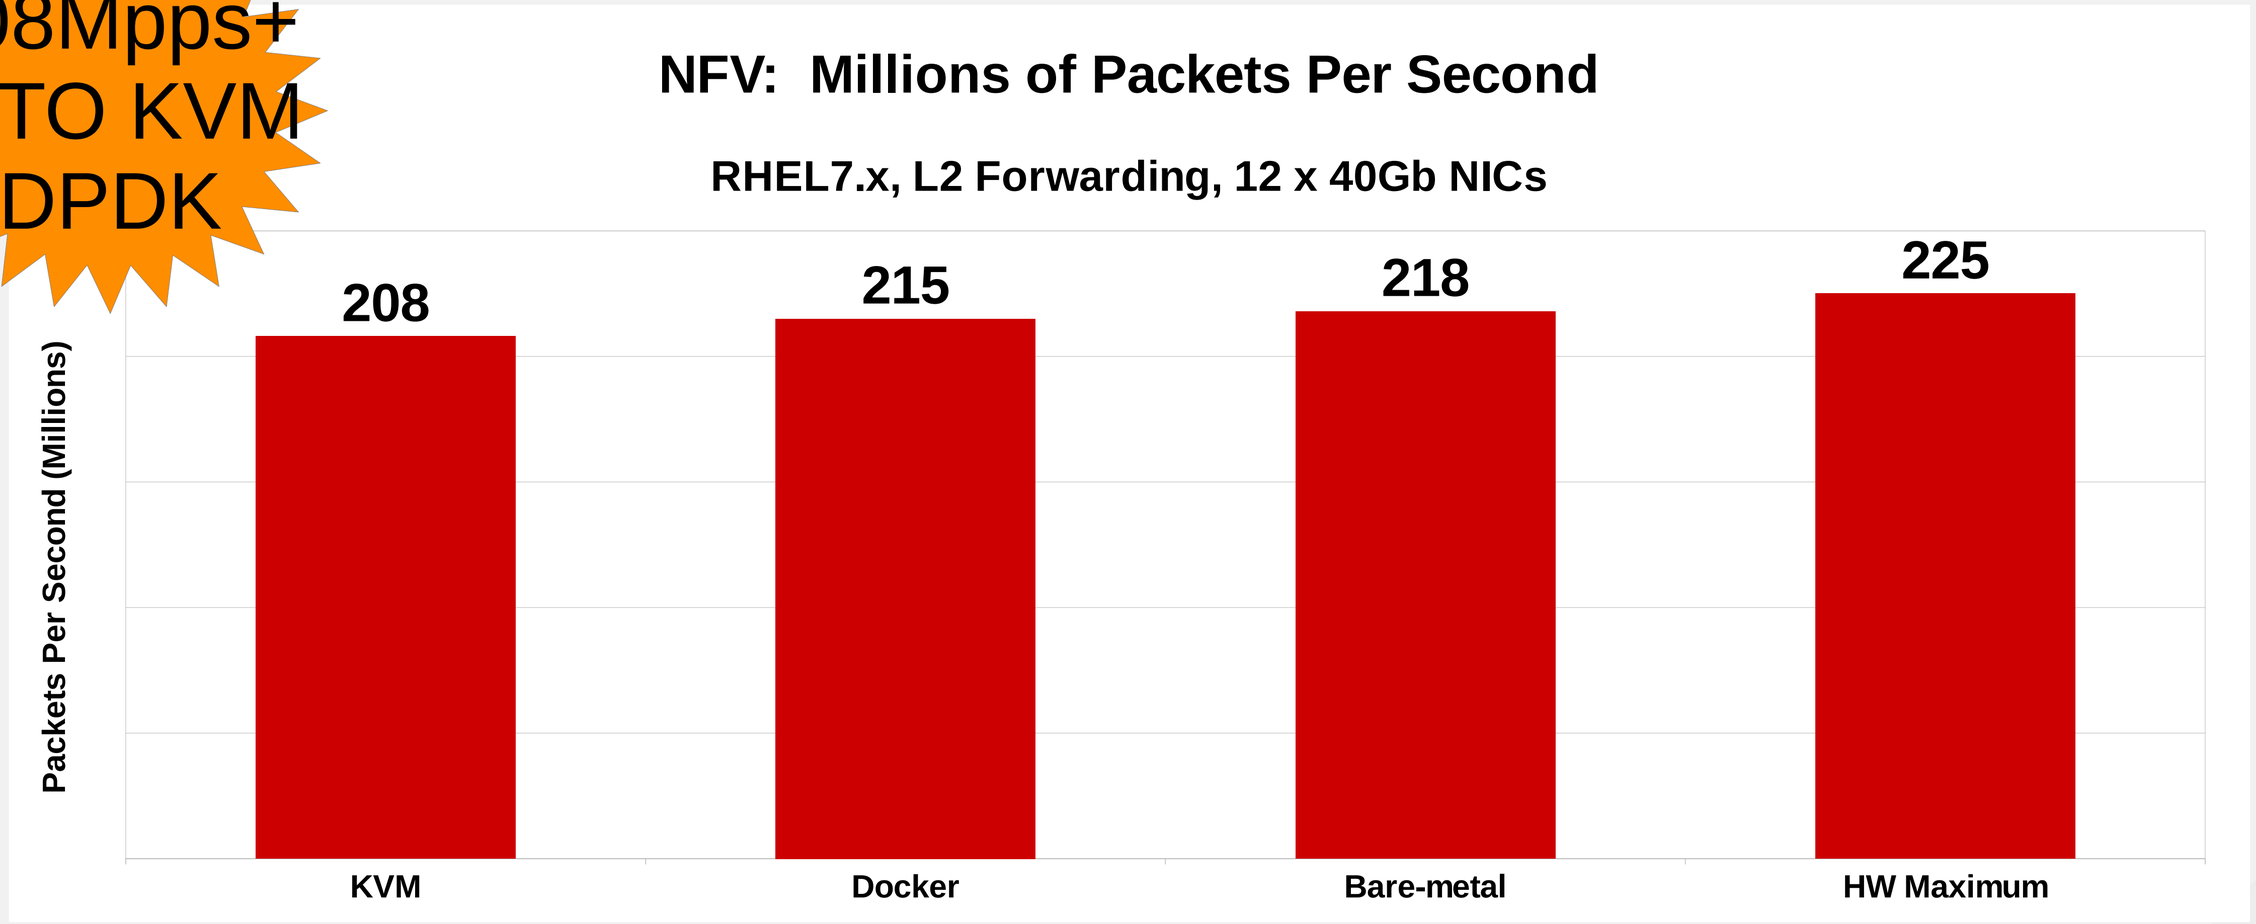

RHEL7.1 + Solarflare OpenOnload Bare Metal / KVM / Containers

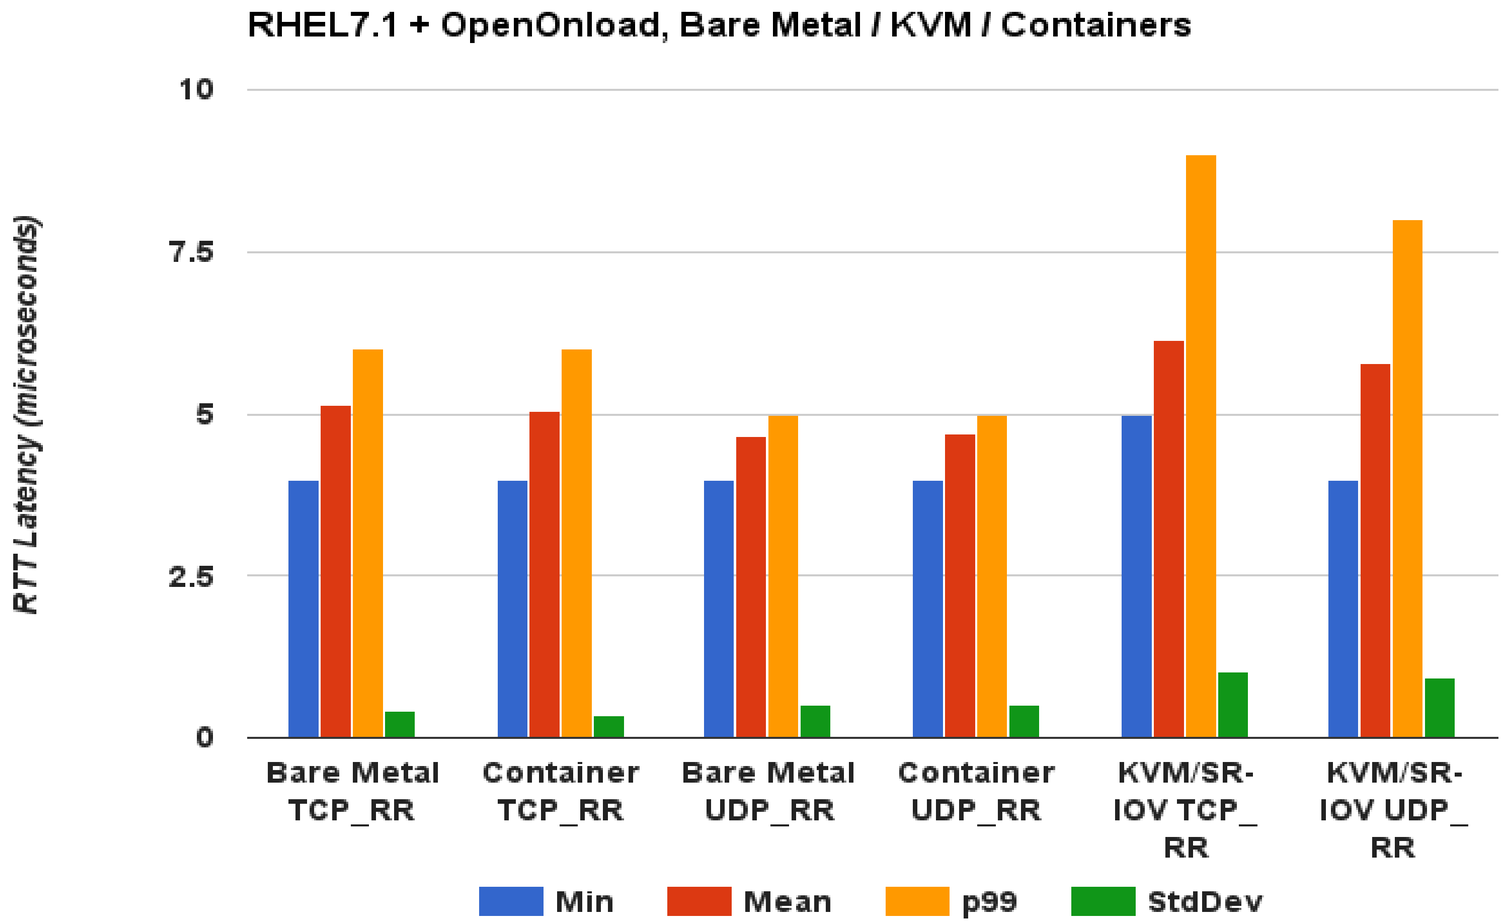

- Lower is better
- Alternative kernel-bypass mechanism to DPDK

Workload Classification

Workload	Bare Metal	Containers	KVM
CPU-bound			
Memory Intensive			
Disk Latency			
Disk Throughput			
Network Latency			
Network Throughput			
Security			
Uptime (Live Migration)			
Deployment Speed			
Alternative OS			

LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.