us signal

us signal opencloud

Evaluating OpenCloud Against Other Hypervisors

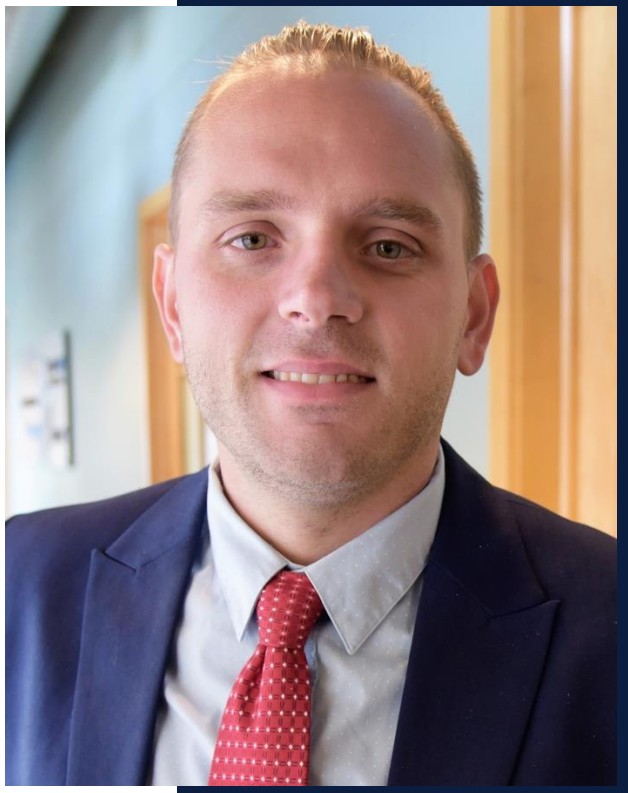

Daven Winans

VP of Engineering
& Innovation

- **15 years in the industry**
- **Managed Services**
- **Professional Services**
- **Circuit Design and Provisioning**
- **Technical Operations Center**
- **Cloud Engineering**
- **Monitoring Infrastructure**
- **Data Engineering**
- **Transport Engineering**
- **Innovation**

about

- Data Center Services Provider
- Founded in 2001, based in Grand Rapids MI
- 16 National Data Centers
- Dedicated pre and post sales architectural support
- Technical Operations Center staffed by engineers 24/7/365
- Network-powered cloud and data protection solutions
- Largest fully deployed and privately owned network in the Midwest

company timeline

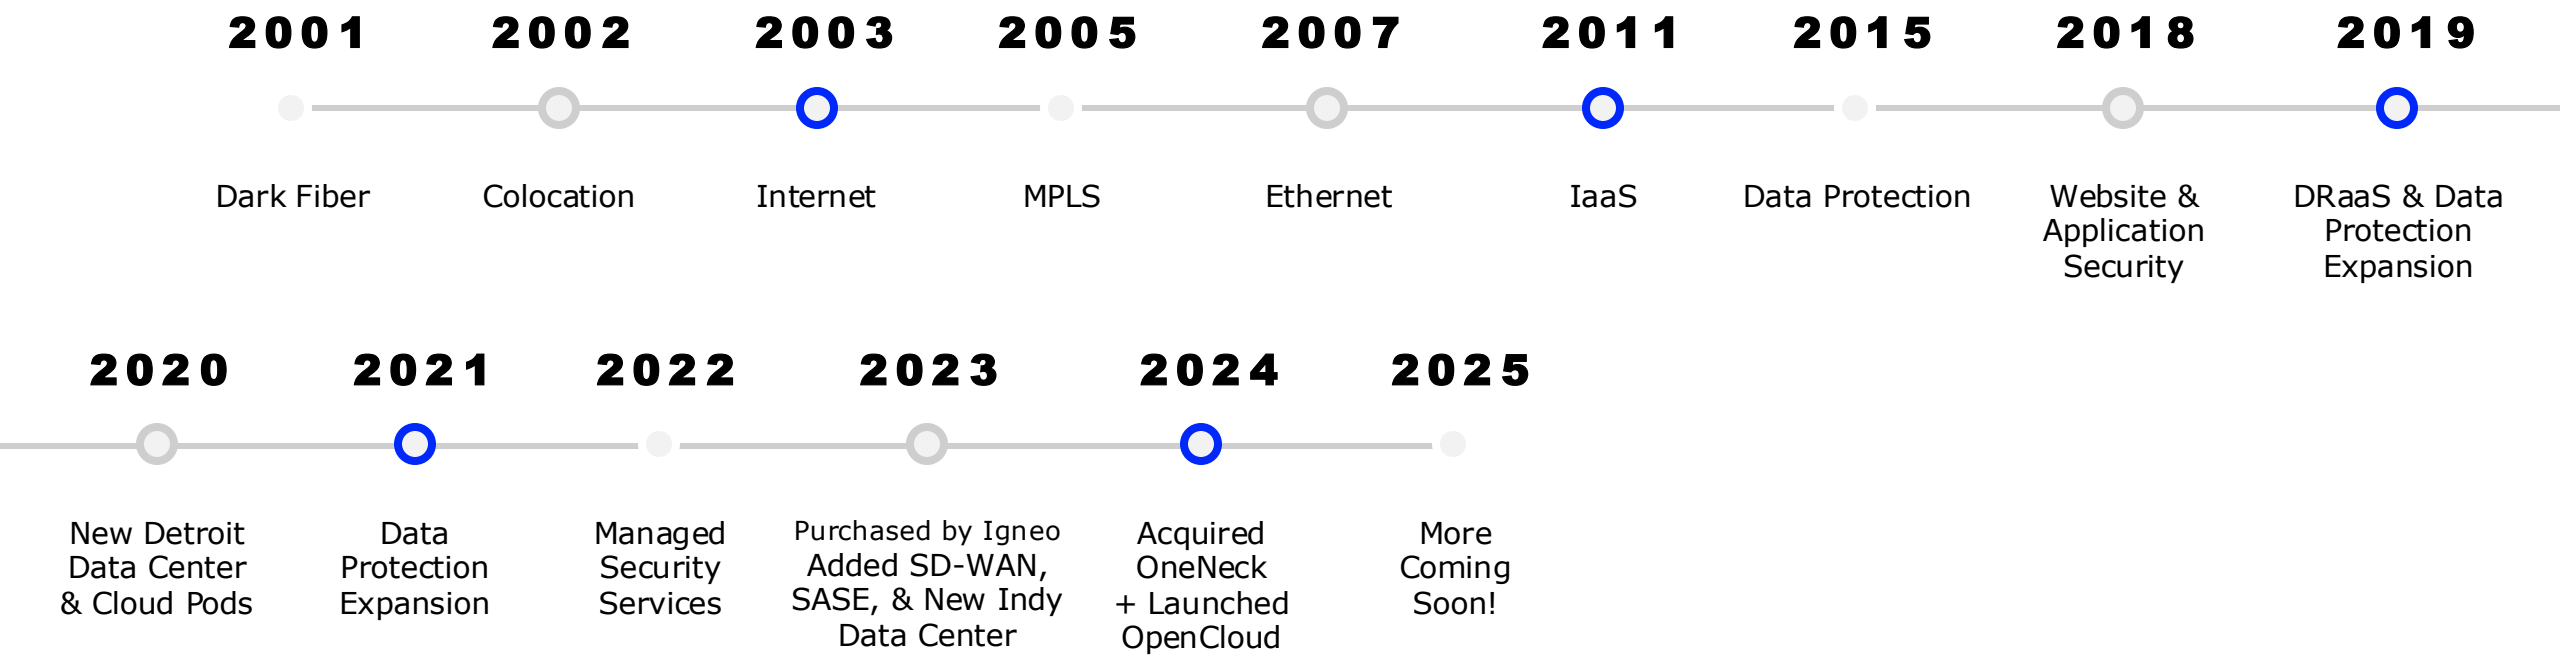

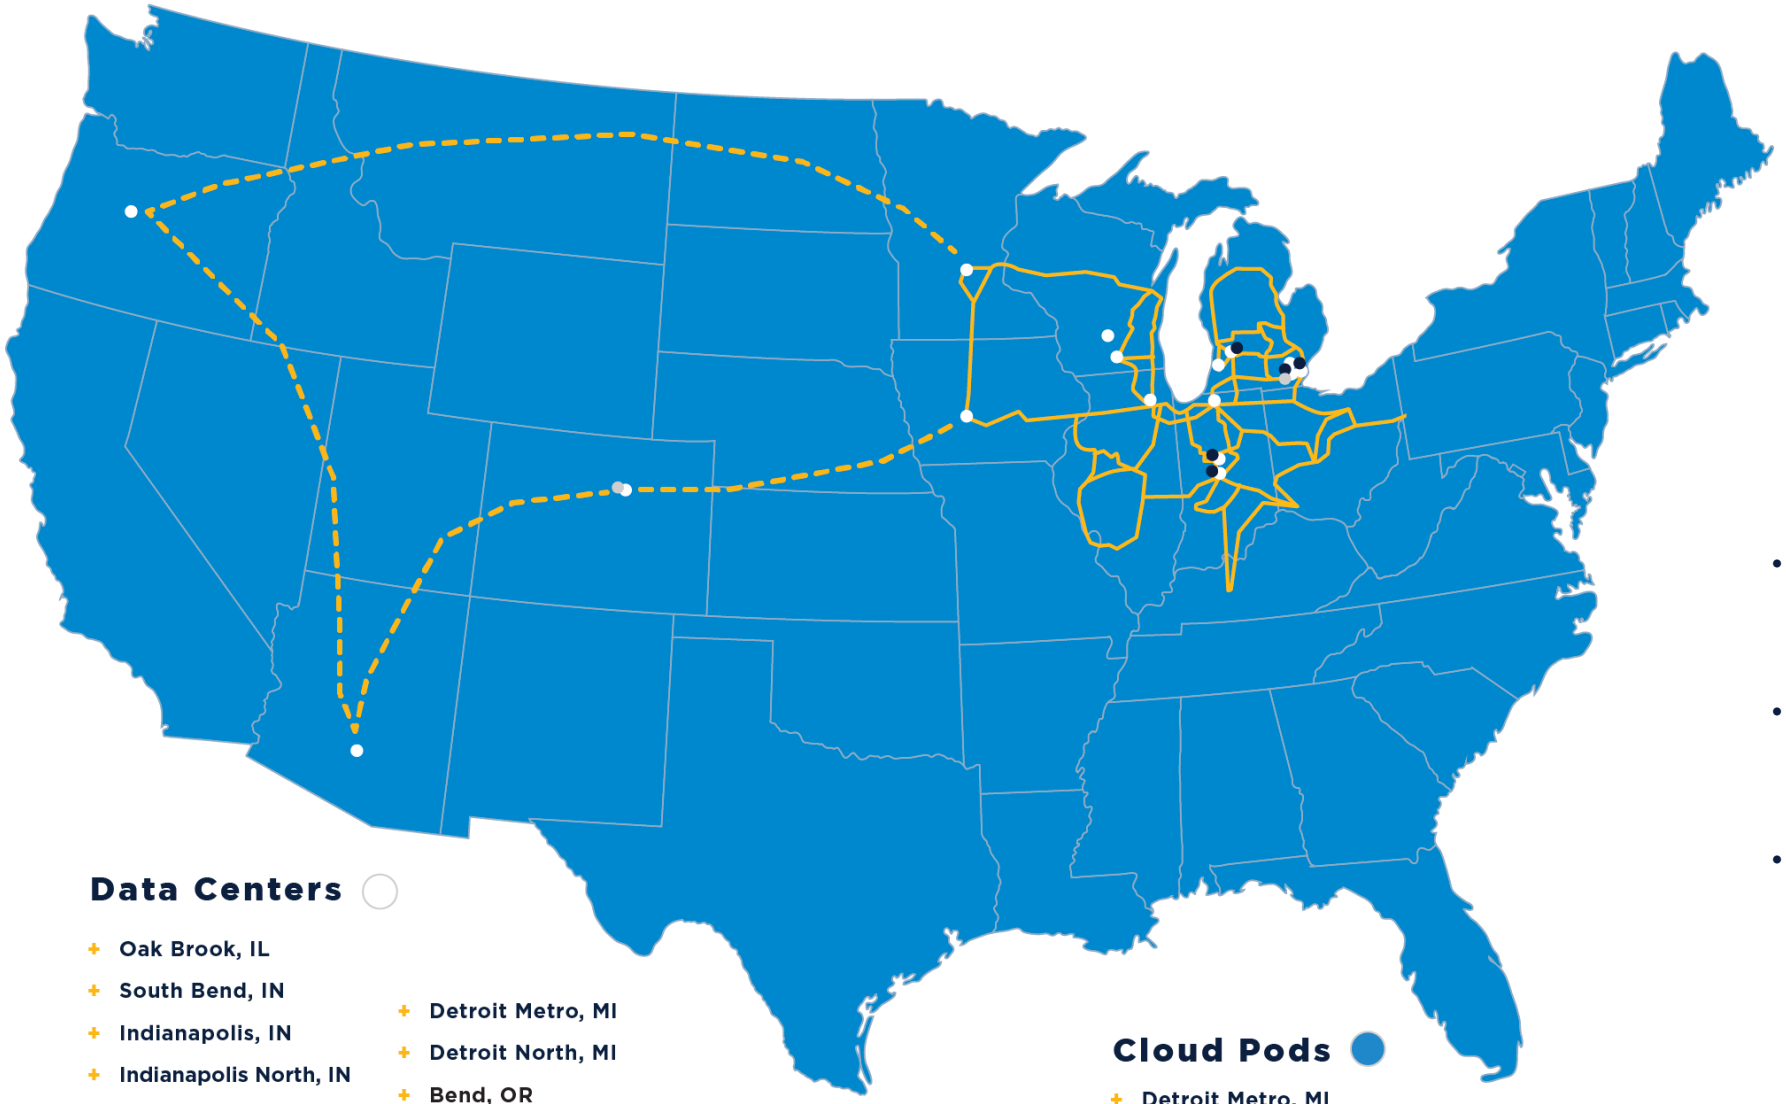

Data Centers

- Oak Brook, IL
- South Bend, IN
- Indianapolis, IN
- Indianapolis North, IN
- Madison, WI
- Grand Rapids East, MI
- Grand Rapids South, MI
- Southfield, MI
- Detroit Metro, MI
- Detroit North, MI
- Bend, OR
- Denver, CO
- Des Moines, IA
- Madison, WI
- Minneapolis, MN
- Phoenix, AZ

OpenCloud Pods

- Detroit North, MI
- Denver, CO

Cloud Pods

- Detroit Metro, MI
- Southfield, MI
- Grand Rapids East, MI
- Indianapolis, IN
- Indianapolis North, IN

Current US Signal Fiber

Future US Signal Fiber

- 9,500 miles of Fiber Optics with future additions in the works
- Fully protected core, distribution and edge networks
- Private access to AWS, Azure, Google, IBM, Oracle, SAP, Salesforce

technology trends

VMWare

- VMWare has been the go-to virtualization platform for nearly 20 years in both the on-premise and cloud IaaS spaces
- Broadcom acquired VMWare in late 2023, and in 2024 instituted massive price increases
- At least 20% of current VMWare customers will look to transition to a new platform going into 2025*

Hyperscalers

- Hyperscalers have mastered consumption modes & excel at reach, scalability, and service integration
- Large scale repatriation from public cloud for better economics and Service level agreement

*Source: https://www.forrester.com/blogs/predictions-2024-technology-infrastructure/

why create OpenCloud?

- We desired a solution that would deliver more predictability in pricing to our customers and to US Signal
- We wanted to deliver a true cloud experience instead of another virtualization option
- We wanted a consumption-based model in the market to give customers consider public cloud or repatriating from public cloud an option with the same “feel”

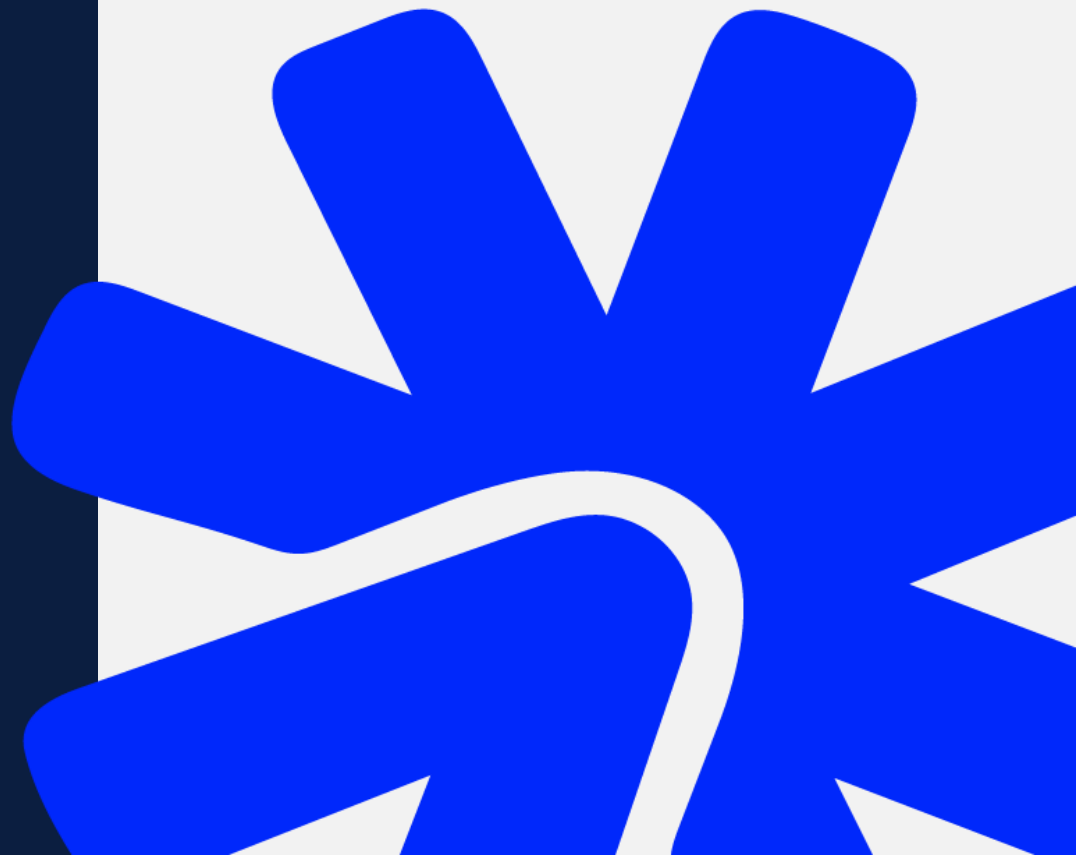

customer advisory board

Concerns:

- **Price** – *OpenCloud is competitively priced and insulated against the next VMWare price increase*
- **Performance** – *OpenCloud is built on Dell R650 servers, Arista switching, & all flash Infinidat storage*
- **Security** – *OpenCloud meets or exceeds USS standards*
- **Compliance** – *OpenCloud is projected to be SOC2 type2, PCI-DSS, & HIPAA compliant*
- **SLA** – *OpenCloud includes a 100% SLA*

Unconcerned with:

- Hardware
- Virtualization Platform

development process

Critera	VMWare (VCF)		Nutanix		KVM+CloudStack	
	Rating	Notes	Rating	Notes	Rating	Notes
Functionality - Platform						
Base Hypervisor	N/A	ESXI	N/A	AHV / KVM	N/A	KVM
Ability to migrate VMs from disparate vCloud Environments into Hypervisor environment.	Meets Expectations	Utilizing Zerto Integration or Acronis BMR. OVA/OVF upload can be troublesome.	Meets Expectations	On their website they say we can move VMs from ESXi to AHV. Need more info.	Meets Expectations	vcenter, backup tested successfully
Host High Availability	Meets Expectations	High Avaiability	Meets Expectations	N+1,2, RF2/3	Meets Expectations	
Automatic Cluster Balancing	Meets Expectations	DRS	Meets Expectations	Live Migration 5ms to 40ms	Meets Expectations	DRS
Non-disruptive VM migration	Meets Expectations	vMotion			Meets Expectations	
VDI Support	Meets Expectations	Horizon VDI	Meets Expectations	Nutanix VDI (can do HCI VDI or Horizon on ESXi as well)	Doesn't Meet Expectations	Claims that VDI support through XCp-ng but unclear if it really works
Storage Balancing (Non-HCI)						
Resource Allocation models for Mult-Tenancy	Meets Expectations	Flex/PayG/Reservation/Allocation			Meets Expectations	can set limits on resources provisioned and allocated
Streched Cluster Model Support					Meets Expectations	cannot break availability zone barrier
vTPM Guest Support		DI	Meets Expectations			
CPU Compatibility mode (Different CPU's in Clusters)	Meets Expectations	Vmware EVC	Doesn't Meet Expectations		Meets Expectations	
Storage Format		VMDK / RDM (RAW)	Meets Expectations	vmdk converted on upload	Meets Expectations	QCOW2
Hypervisor resource footprint			Meets Expectations	NCM/Prism/medium mgmt overhead per host. See sizer scenario	Exceeds Expecations	light mgmt overhead. seperate mgmt node(s)
Min nodes in a cluster	Meets Expectations	Non-VSAN 1 | VSAN =3	Meets Expectations	3	Meets Expectations	1
Max nodes in a Cluster	Meets Expectations	64	Meets Expectations	32	Doesn't Meet Expectations	16
Min VM size						
Max VM size						
Max VMs per host						
Max Operations		Configurable				
Nondisruptive Upgrades	Meets Expectations		Meets Expectations	Supported	Meets Expectations	management of upgrades not the easiest
Cluster Expansion Ability (More HCI related)						
vGPU Support		Supported		Supported		

evaluation process

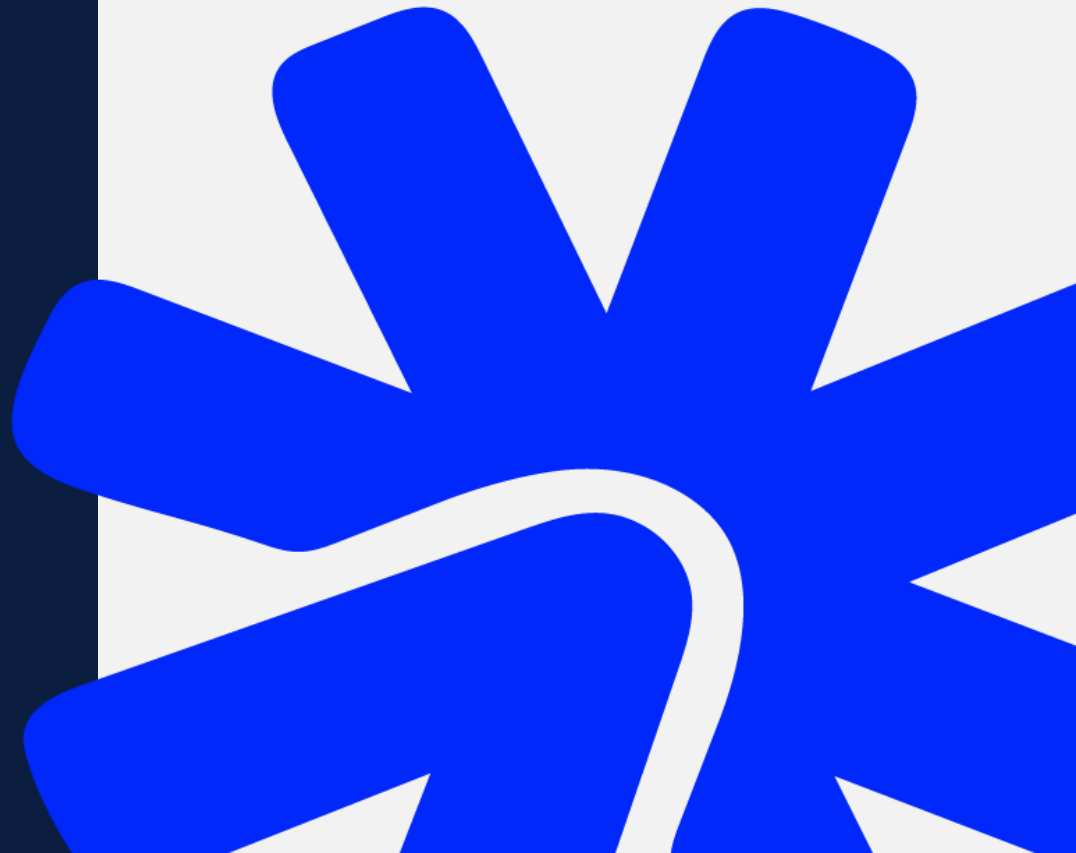

Who are we evaluating against?

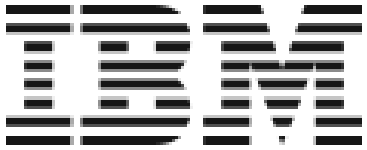

the toolset (hypervisor) – open source

Unixbench –

UnixBench is a benchmark suite that provides basic indicators of the performance of a Unix-like system. It runs multiple tests on a system and compares the scores from a baseline system to produce an index value, which is combined to make an overall index for the system

OpenSSL –

Used to test the performance of cryptographic algorithms.

Iperf –

IPerf is a tool for active measurements of maximum achievable bandwidth on IP networks.

Postgres/pgbench -

pgbench is a simple program for running benchmark tests on PostgreSQL. It runs the same sequence of SQL commands over and over, possibly in multiple concurrent database sessions, and then calculates the average transaction rate (transactions per second), as well as the average per-statement latency (execution time from the perspective of the client).

FIO -

FIO is a simple industry standard tool allowing for customizable IO testing. The power of FIO is simplicity and ability to tailor testing to specific workloads.

the toolset (hypervisor) – continued

UnixBench

Authors: Ben Smith, Rick Grehan, Tom Yager, Jon Tombs

Maintainers: David Niemi, Ian Smith, Anthony Voellm

FIO

Author: Jens Axboe

Maintainers: Jens Axboe & Vincent Fu

OpenSSL

Author: Mark Cox, Ralf Engelschall, Stephen Henson, Ben Laurie, Paul Sutton

Maintainers: The OpenSSL library project is managed co-equally by the Foundation and Corporation in conjuction with our various committees and communities. OpenSSL Software Foundation Inc Web: https://openssl-foundation.org

iperf

Author: Mark Gates, Alex Warshavsky

Maintainers: The Energy Sciences Network (ESnet)

pgbench

Author: Tatsuo Ishii

Maintainers: Greg Smith, Josh Kupershmidt, Emeric tabakhoff

the toolset (public) – Cloud Mercato

The Specification

In our quest for comprehensive cloud analysis, we've developed sophisticated methods to capture every crucial detail of virtual hardware. Our approach ensures that no aspect of the cloud infrastructure goes unexamined. Here's what we scrutinize:

•CPU topology, caches, frequency and feature
•Drives capacity
•Accelerated hardware
•Pricing

The Performance

Performance is at the heart of cloud computing, and our analysis digs deep into various aspects of cloud performance. We employ a range of benchmarks and tests to provide a complete picture of how cloud resources perform under different conditions:

•Raw performance from CPU, RAM, network
•Advanced workload performance

The Testing

- Specs score (score from the catalog specifications)
- CPU\RAM testing
- Network bandwidth
- Network latency
- Price/Performance value from the results above

https://www.cloud-mercato.com/

Performance Comparison - hypervisors

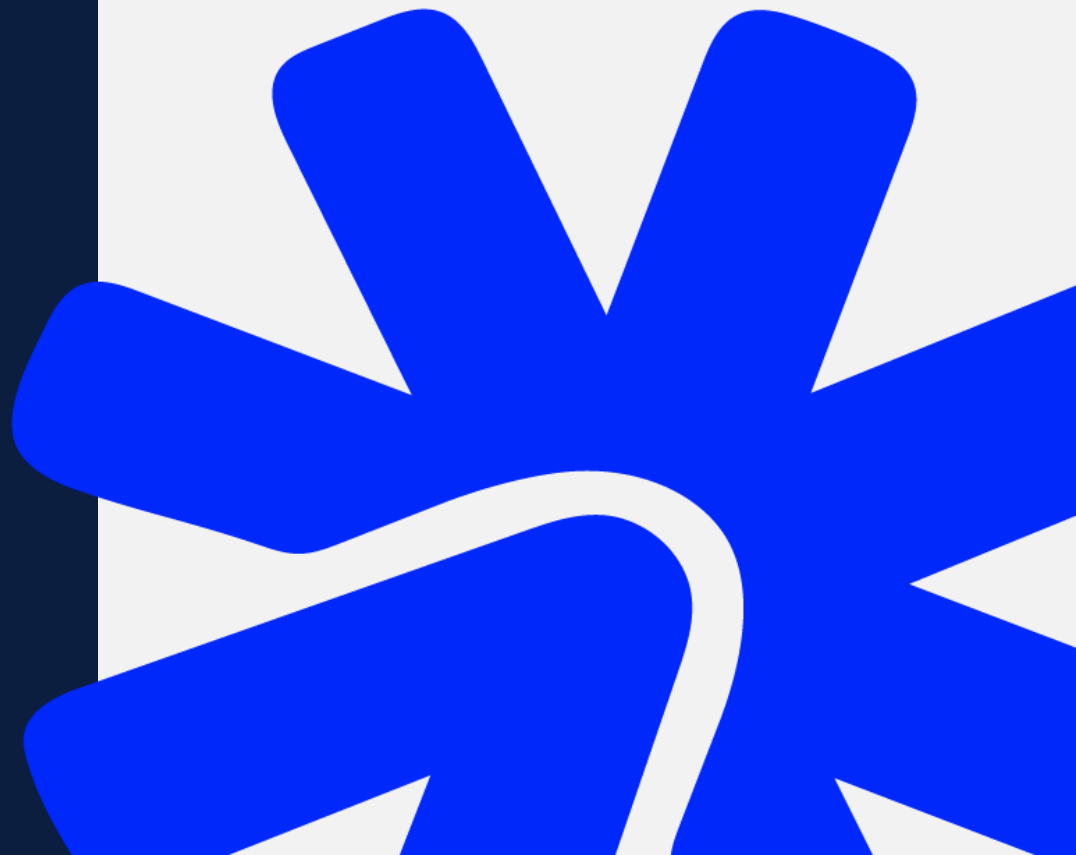

UnixBench - Compute

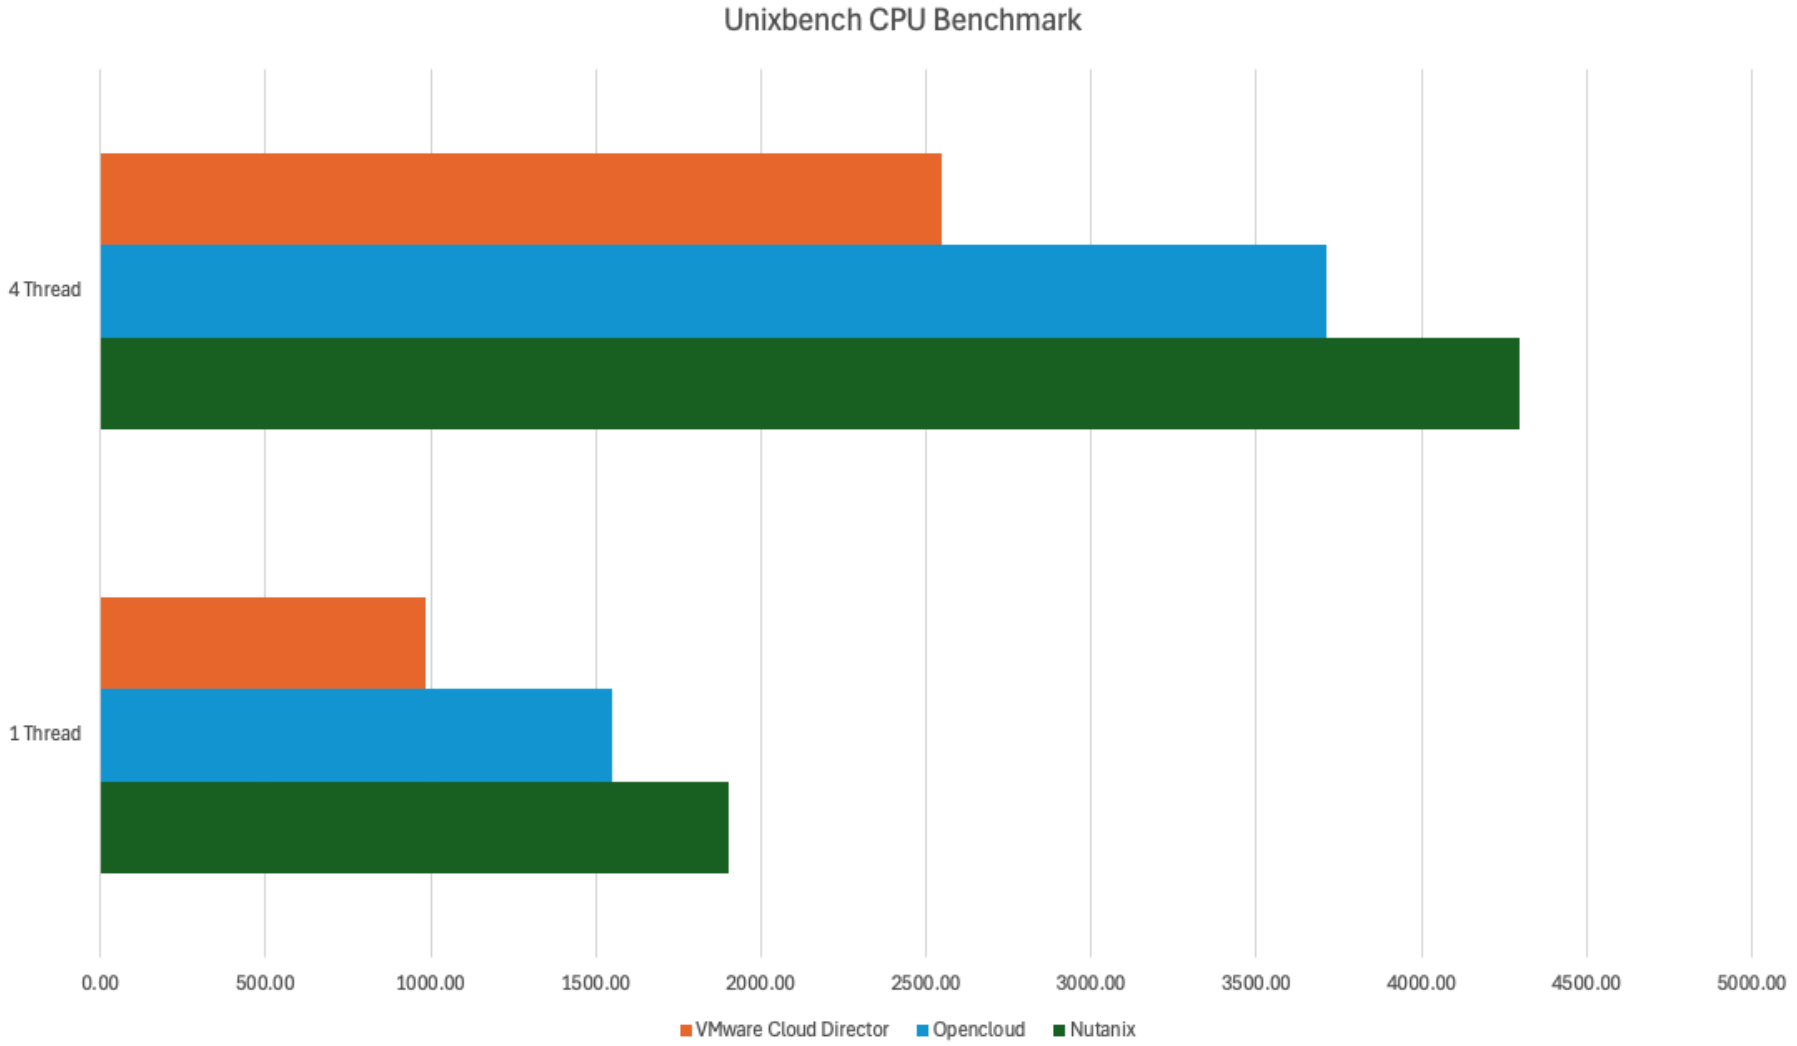

- UnixBench tests consist of two parts: single-process tests and multi-process tests. The number of processes in multi-process tests depends by default on the number of CPUs. The following tests are run:
  Dhrystone 2
- Double-Precision Whetstone
- Execl Throughput
- File Copy 1024 bufsize 2000 maxblocks
- File Copy 256 bufsize 500 maxblocks
- File Copy 4096 bufsize 8000 maxblocks
- Pipe Throughput
- Pipe-based Context Switching
- Process Creation
- System Call Overhead
- Shell Scripts (1 concurrent)
- Shell Scripts (8 concurrent)

-15% Nutanix +51% Vmware

OpenSSL Speed

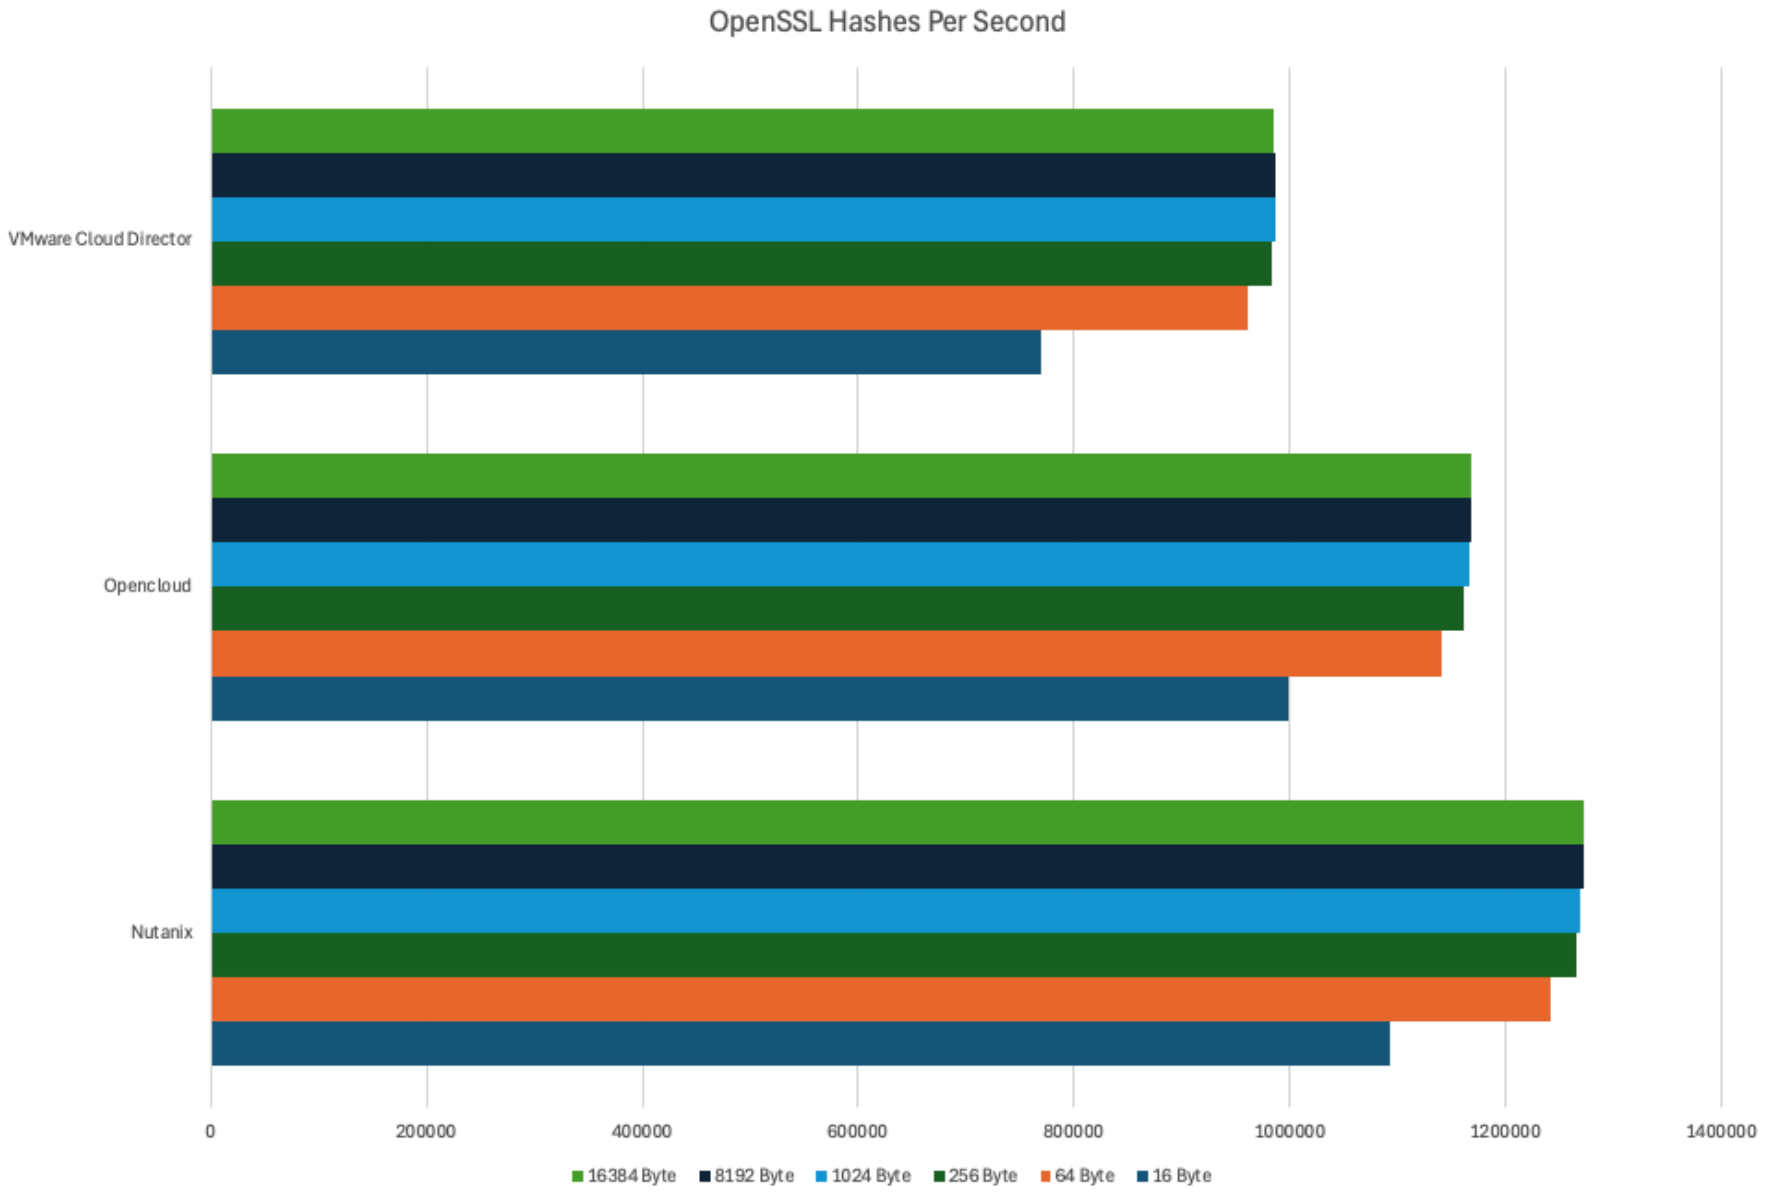

OpenSSL speed is a lighweight opensource tool provided with the openssl command-line tool. It times the encryption/decryption through different algorithm and block size. Its output delivers a good picture of how much a system is able to perform with the different encryption systems available in the industry

-8.2% Nutanix +20.2% Vmware

FIO - storage

The power of FIO is simplicity and ability to tailor testing to specific workloads.

Fio spawns a number of threads or processes doing a particular type of I/O action as specified by the user. fio takes a number of global parameters, each inherited by the thread unless otherwise parameters given to them overriding that setting is given. The typical use of fio is to write a job file matching the I/O load one wants to simulate.For our benchmarking we defined five workloads that are representative and include IOPS, Throughput & Latency.

256K block size Random R/W measuring Throughput

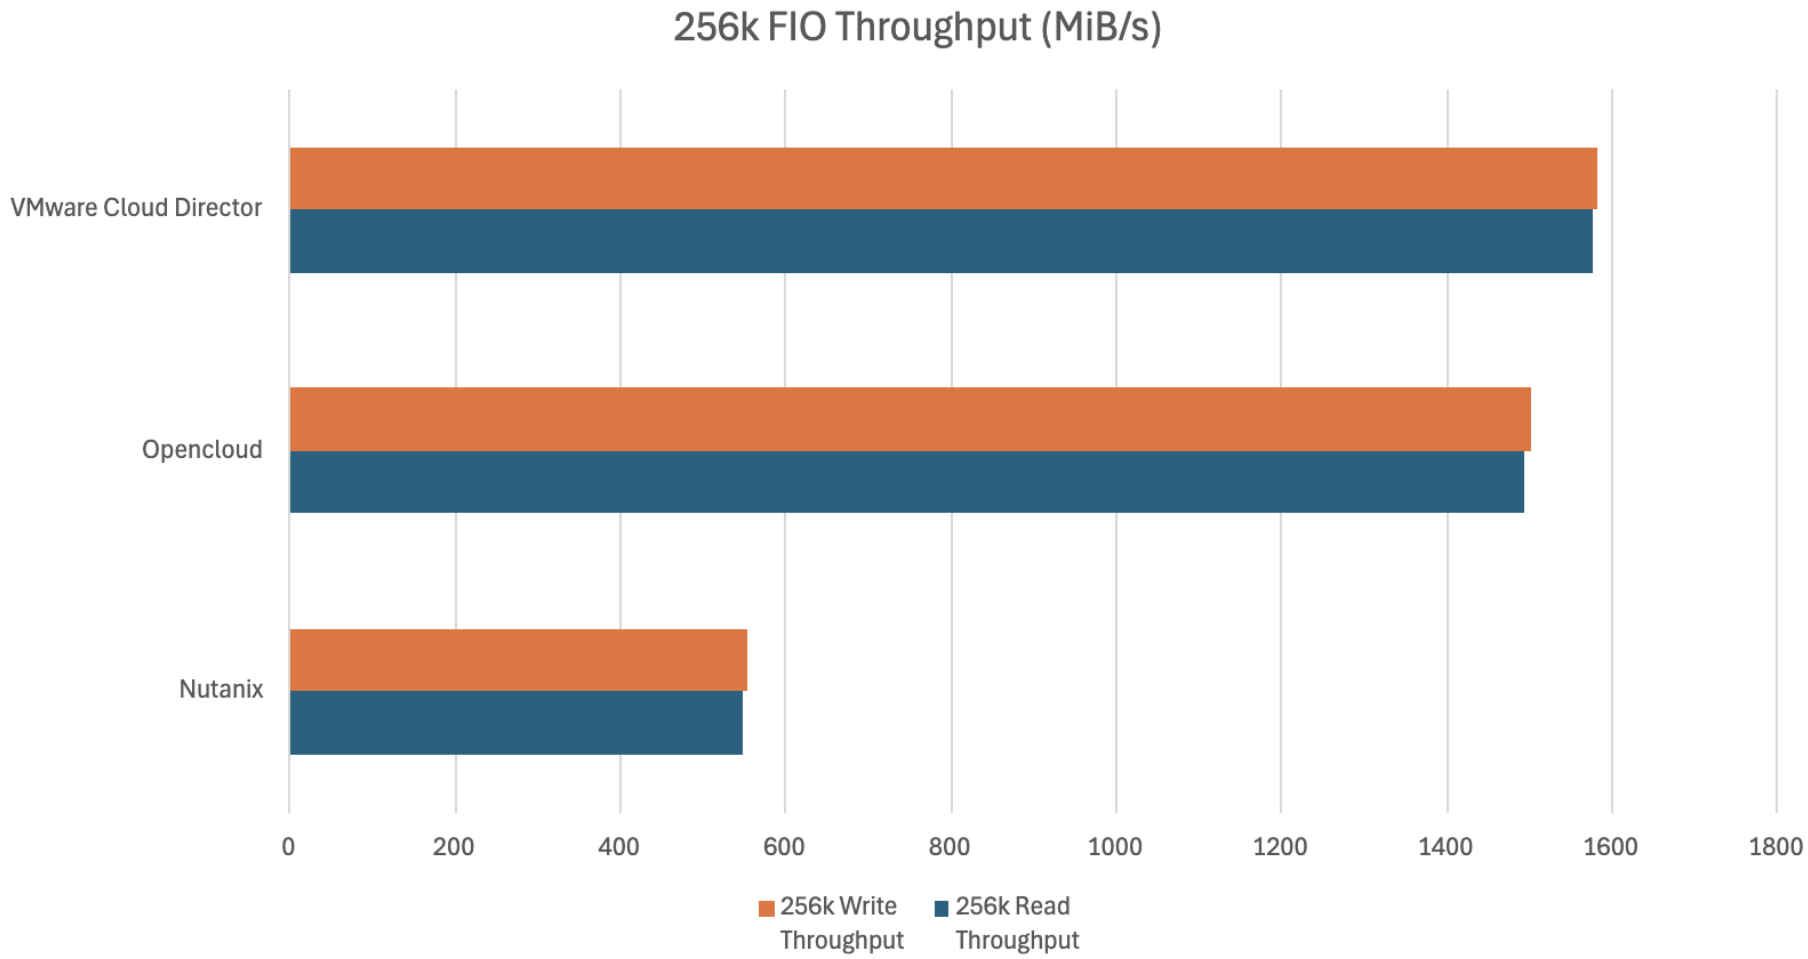

- 1 job per cpu
- queue depth of 64
  - Variables
    - bs=256K
    - iodepth=64
    - direct=1
    - rwmixread=50
    - size= 16G
    - ioengine=libaio
    - group_reporting
    - time_based
    - runtime=20
    - numjobs=4
    - name=256kraw-randrwxput
    - rw=randrw

+171% Nutanix -5.8% Vmware

1MB block size Random R/W measuring Throughput

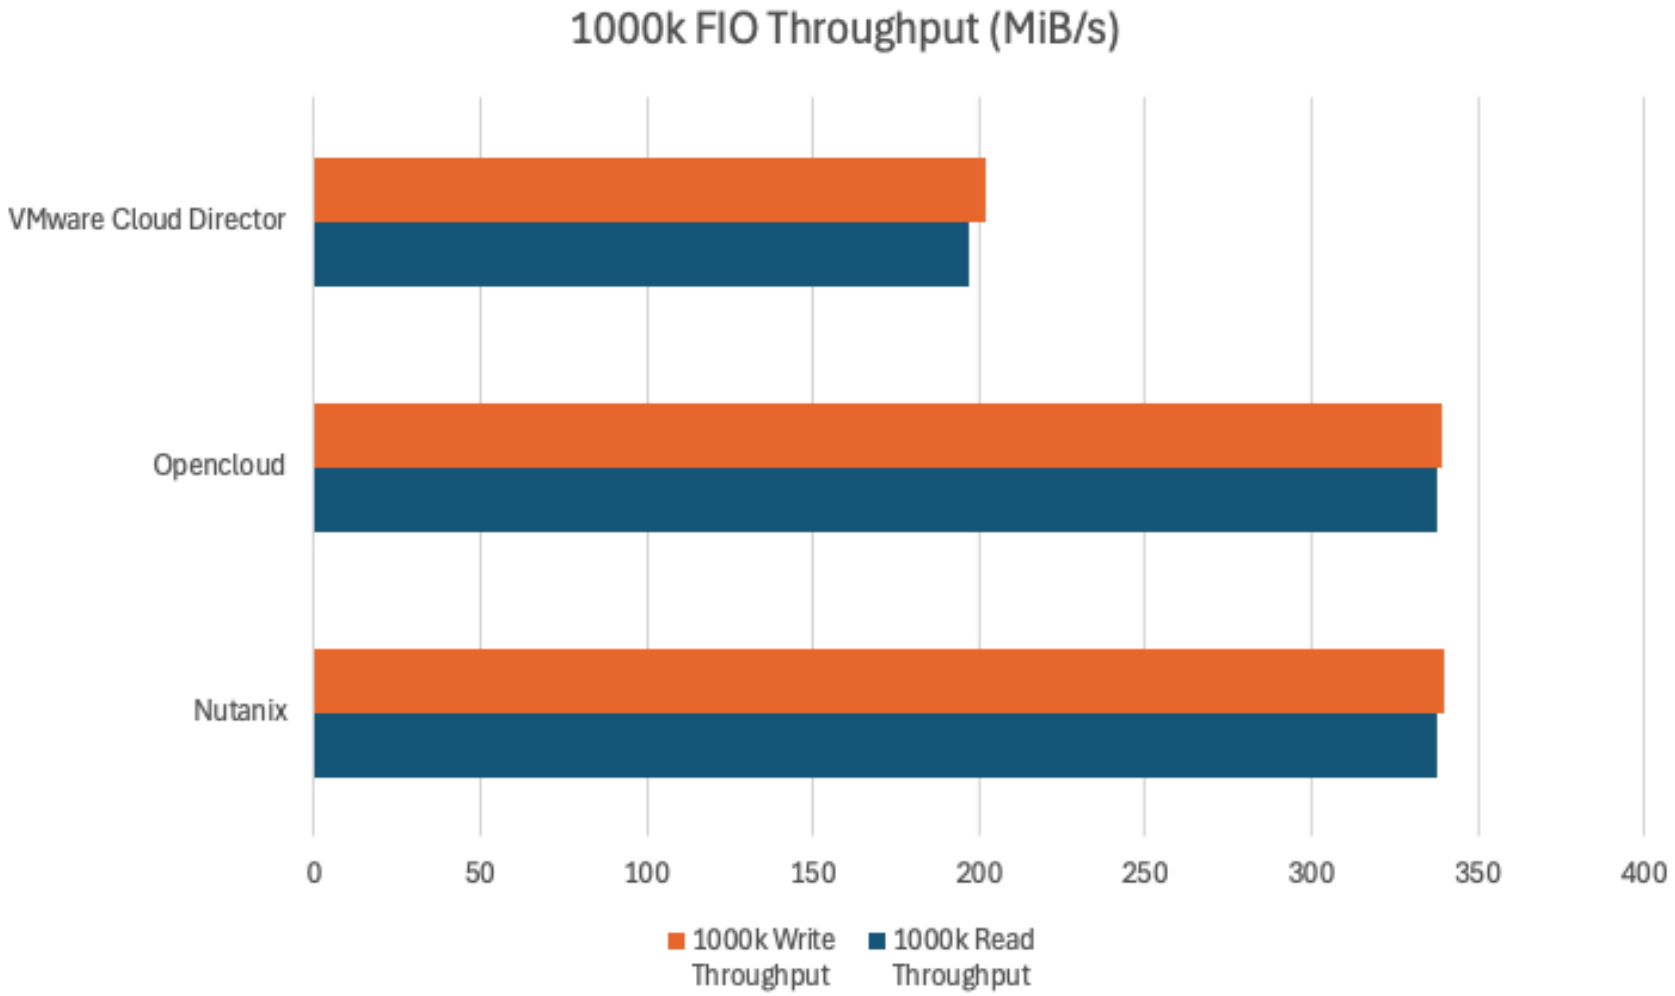

- single job
- queue depth of 1
- approximation of a contiguous operation
  - Variables
    - bs=1000K
    - iodepth=1
    - direct=1
    - rwmixread=50
    - size= 256G
    - ioengine=libaio
    - group_reporting
    - time_based
    - runtime=20
    - numjobs=1
    - name=1mraw-randrwxput
    - rw=randrw

+.15% Nutanix +69% Vmware

4K block size random R/W measuring IO

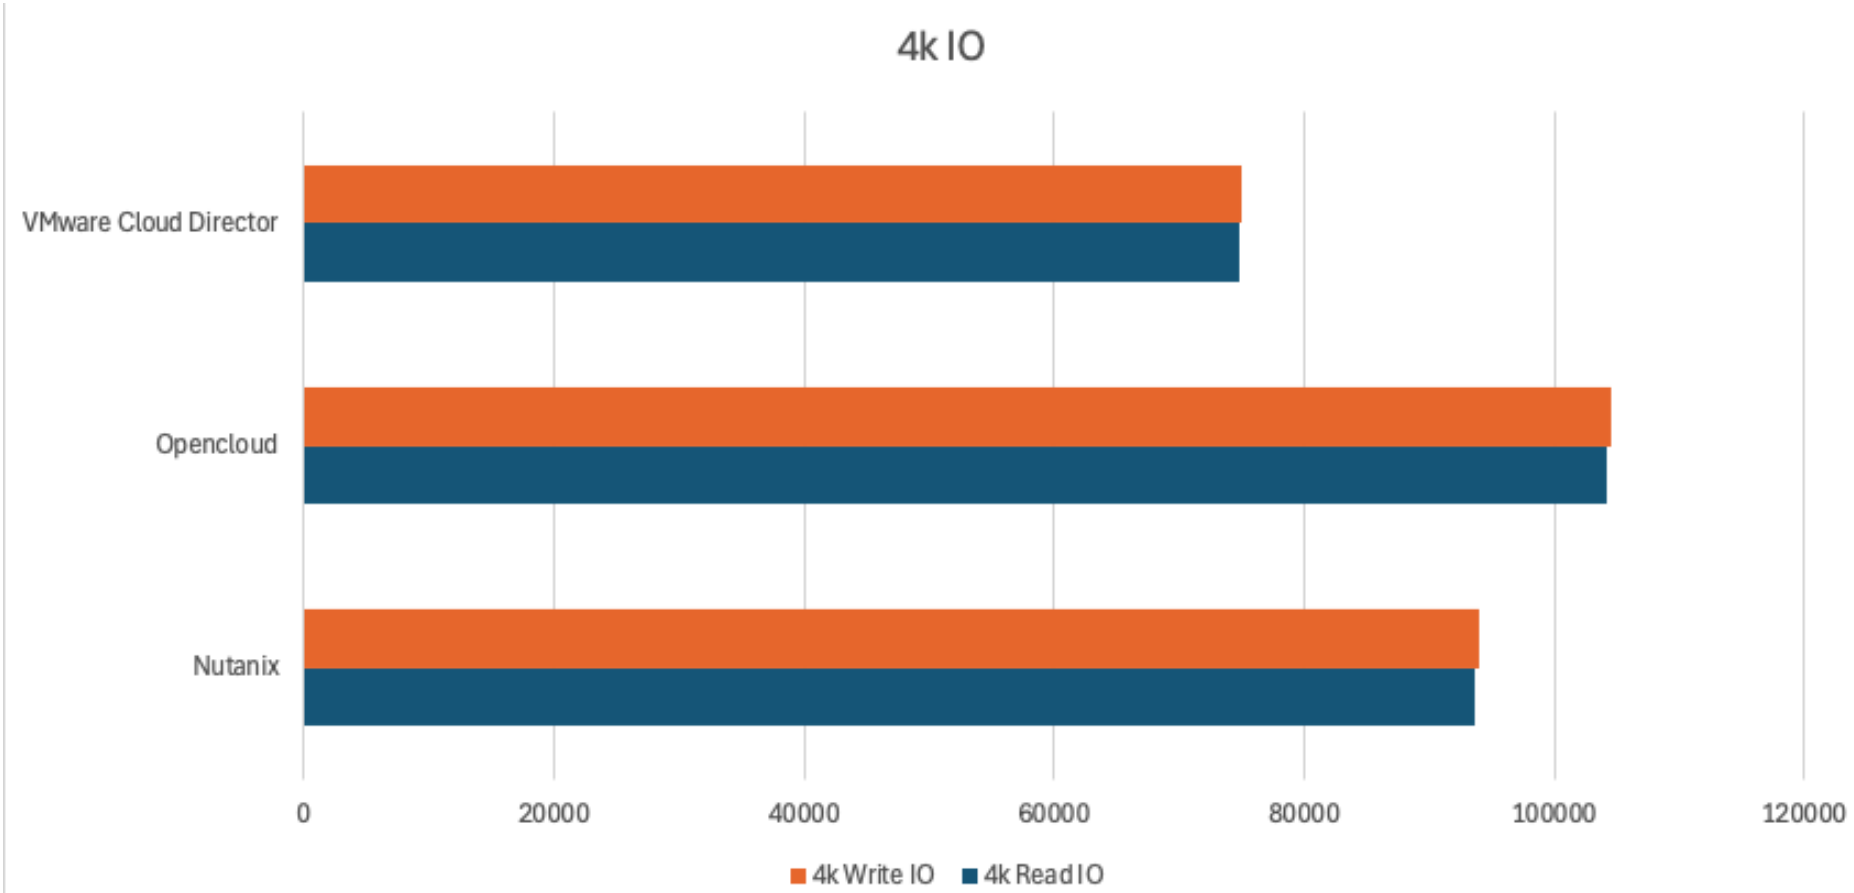

- single operation at a time
- one job per cpu
- disk intensive
- “real-life” scenario
  - Variables:
    - bs=4K
    - iodepth=1
    - direct=1
    - rwmixread=50
    - size= 16G
    - ioengine=libaio
    - group_reporting
    - time_based
    - runtime=20
    - numjobs=4
    - name=4kraw-randrwiops
    - rw=randrw

+11.4 Nutanix +39% Vmware

64K block size random R/W measuring IO

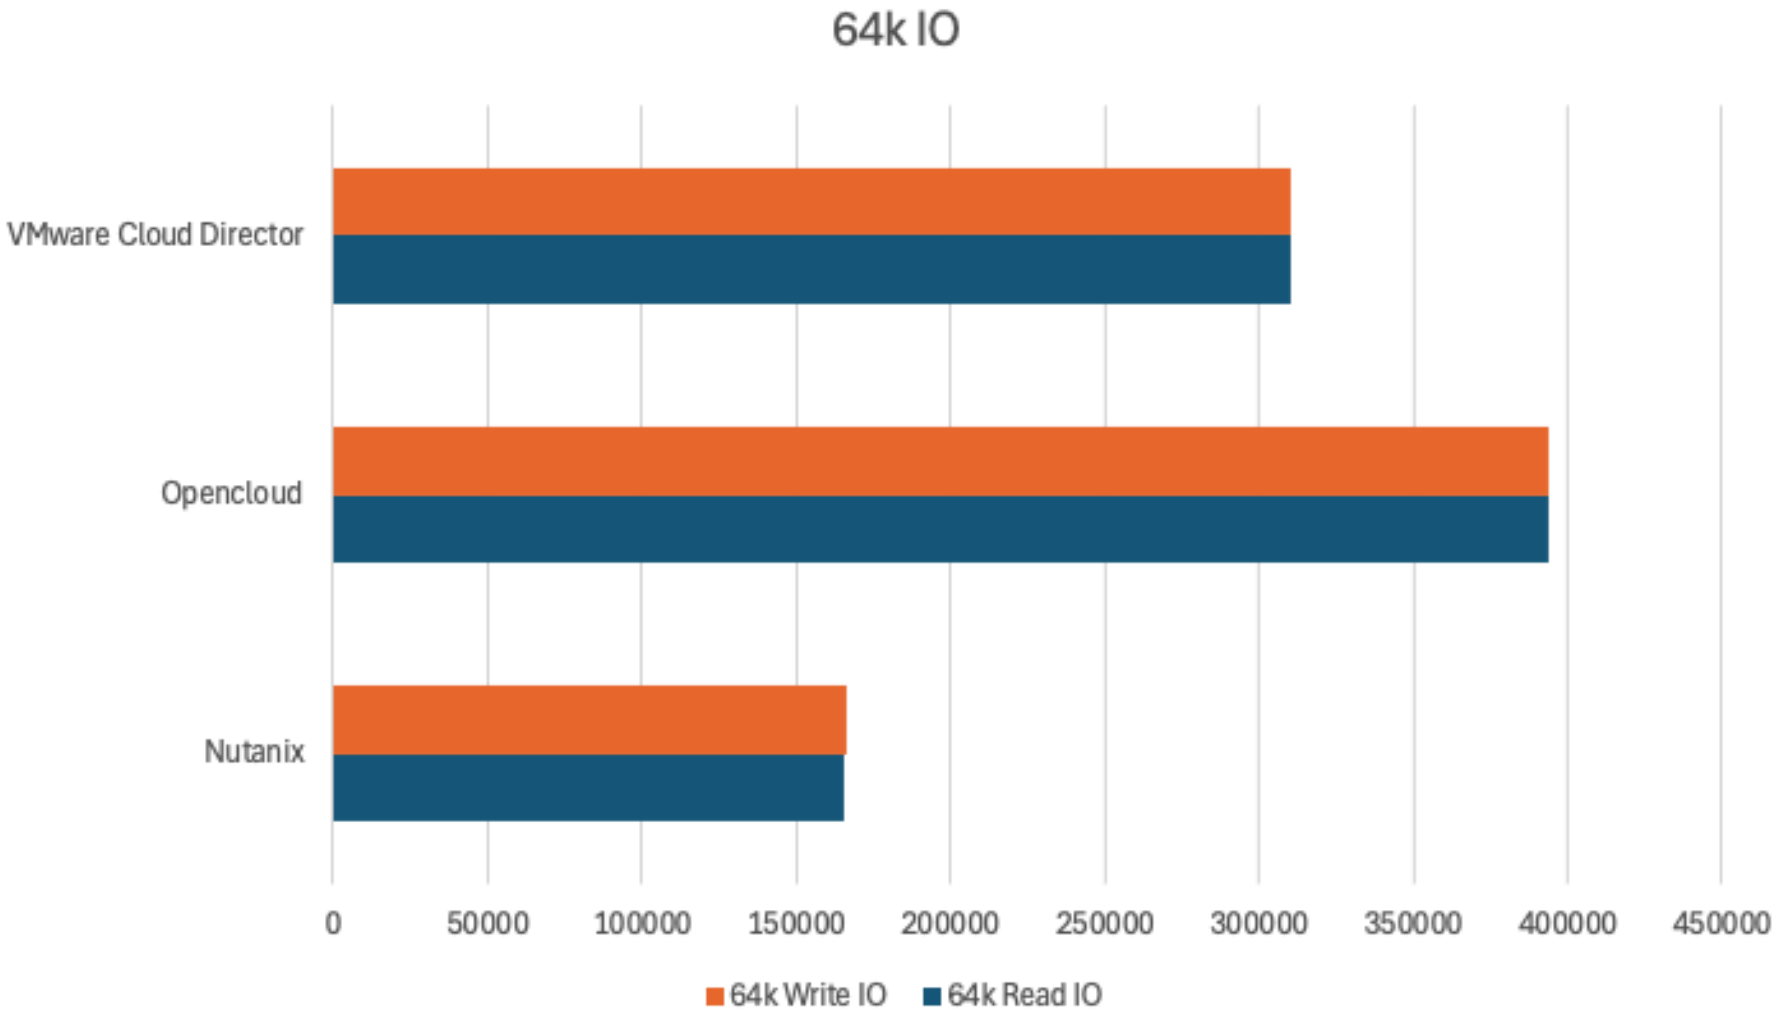

- 4 jobs per cpu
- queue depth of 16
- approximation of a “busy system”
  - Variables:
    - bs=64K
    - iodepth=16
    - direct=1
    - rwmixread=50
    - size= 16G
    - ioengine=libaio
    - group_reporting
    - time_based
    - runtime=20
    - numjobs=16
    - name=64kraw-randrwiops
    - rw=randrw

+138% Nutanix +27% Vmware

4K block size R/W measuring latency

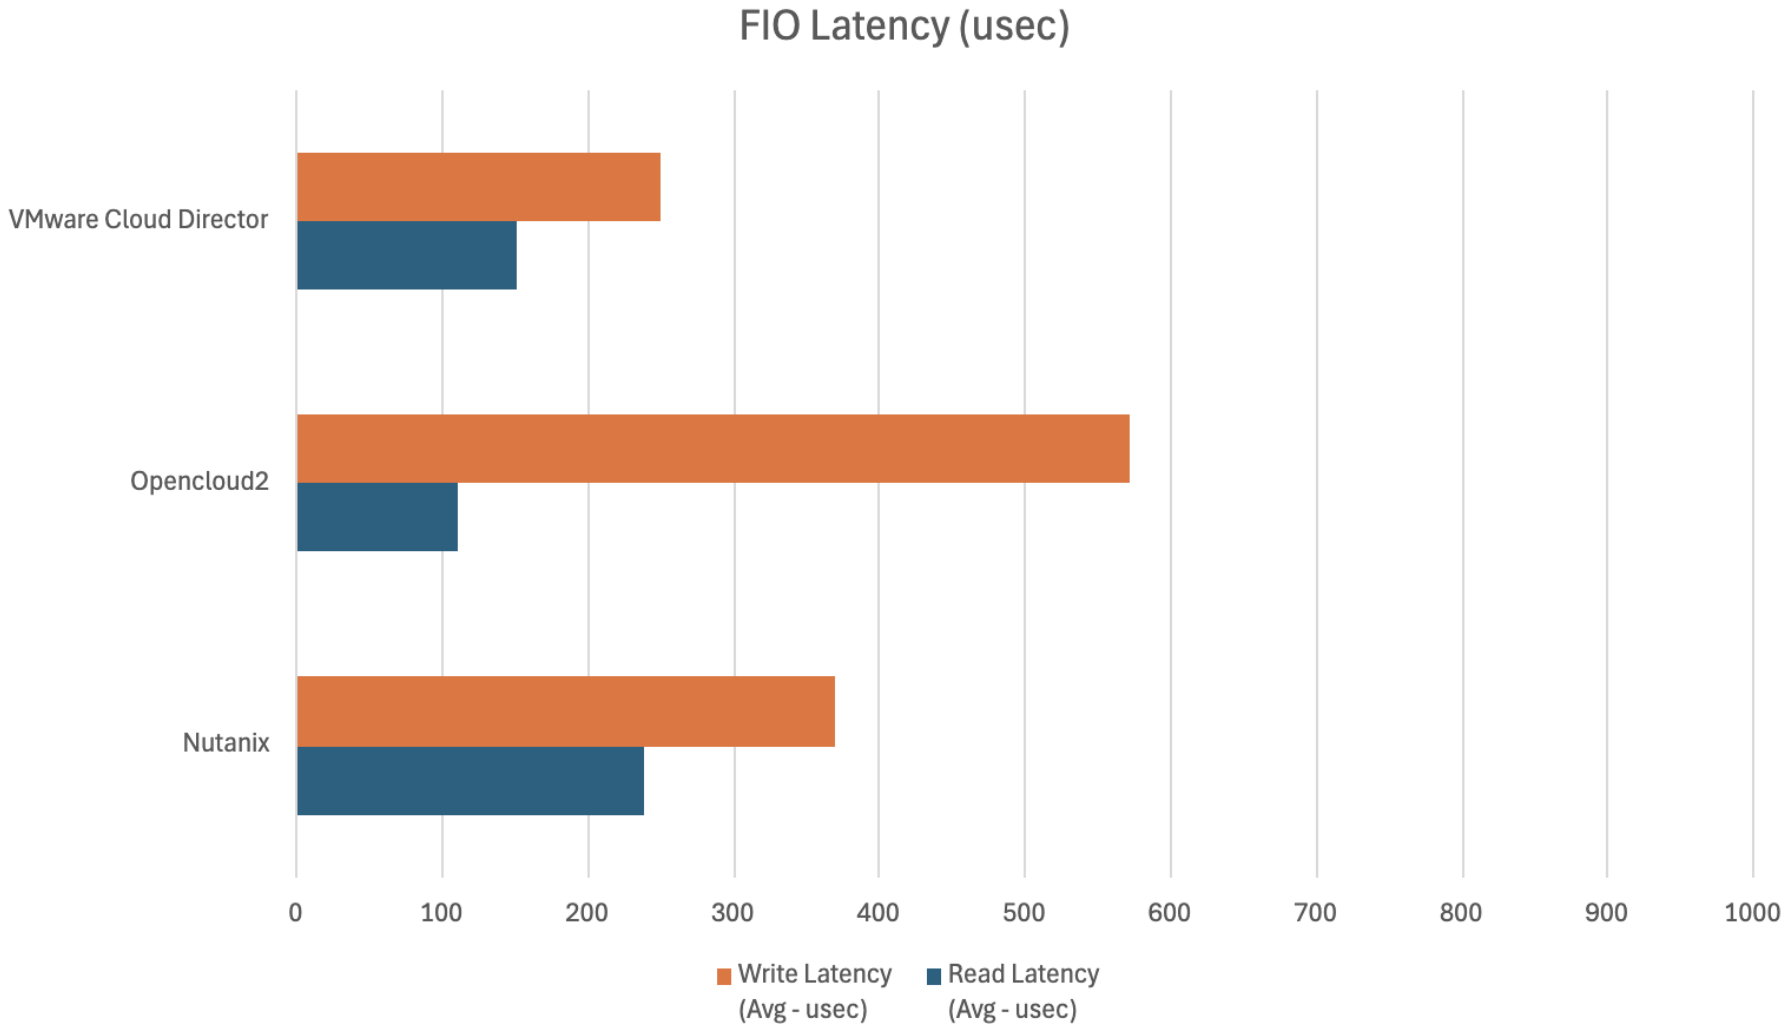

- single job
- queue depth of 1
- generic latency test
  - Variables
    - bs=4K
    - iodepth=1
    - direct=1
    - rwmixread=50
    - size= 16G
    - ioengine=libaio
    - group_reporting
    - time_based
    - runtime=20
    - numjobs=1
    - name=4kraw-randrwlatency
    - rw=randrw

iPerf - network

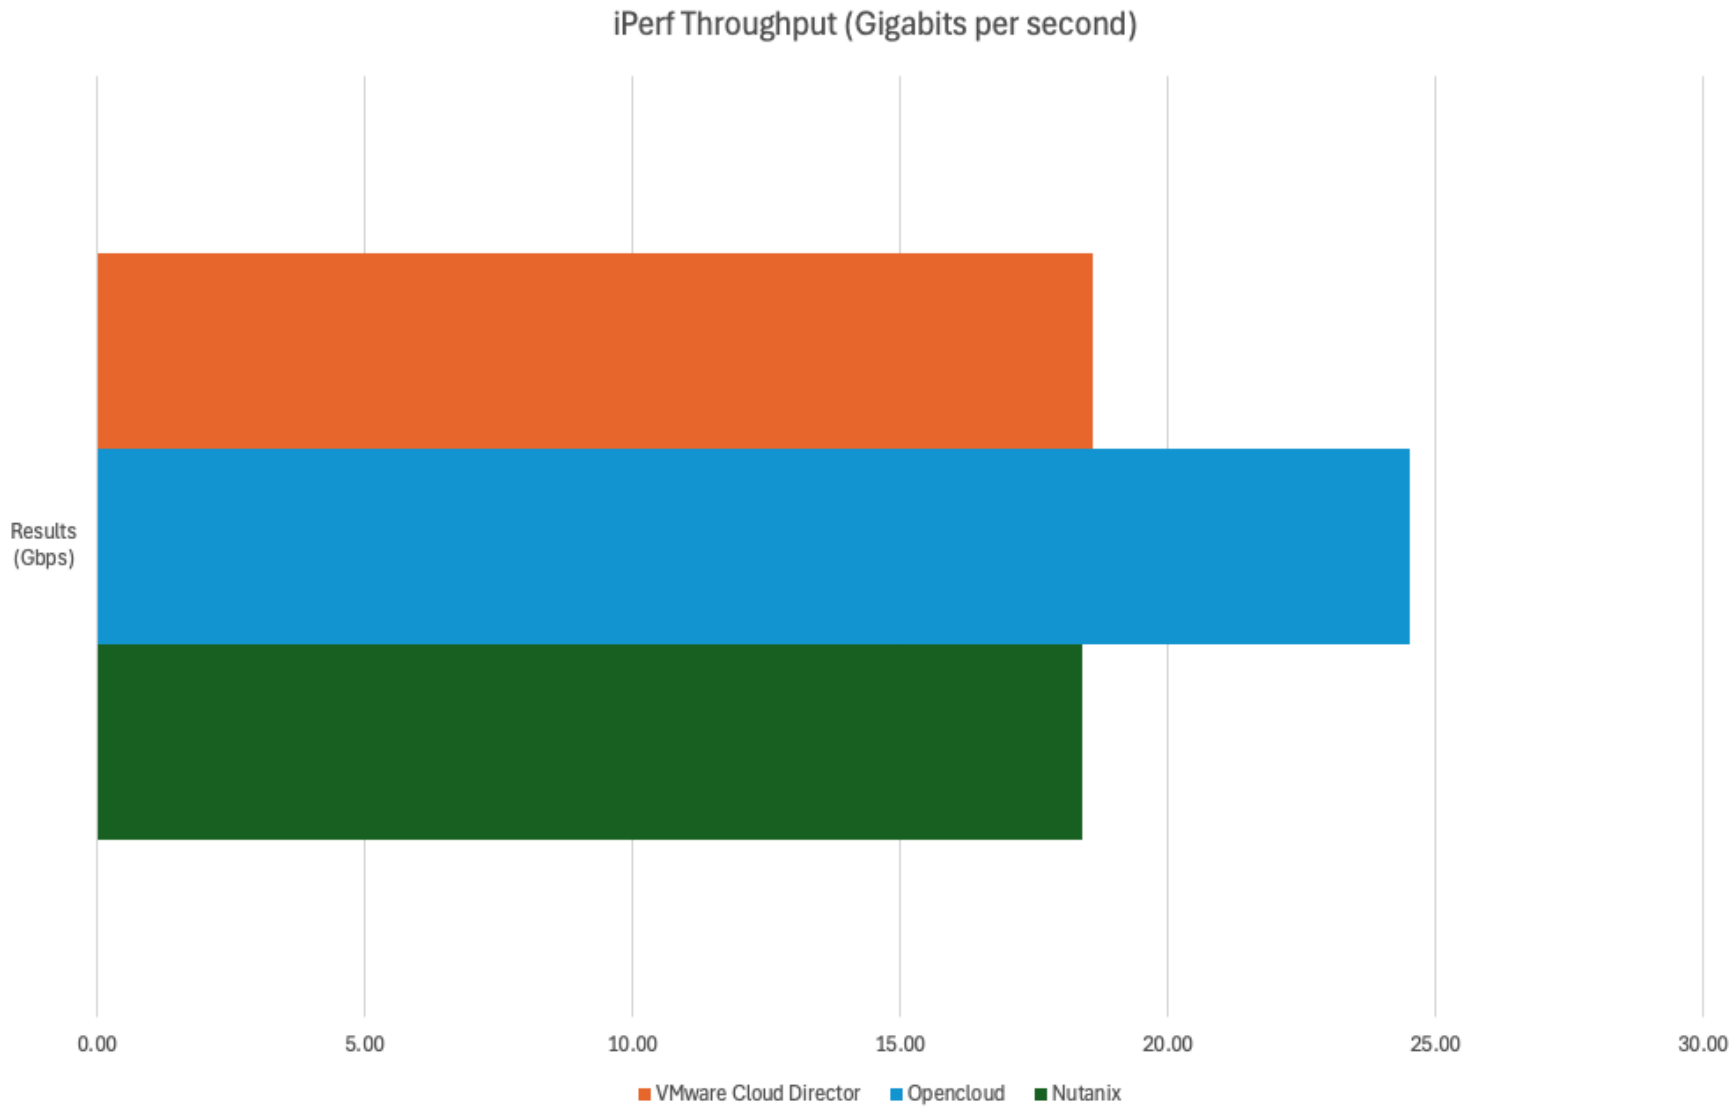

No complexity in this test, using anti-affinity rules we deploy our vm template on two different nodes & run the test with parameters to measure east-west

+33% Nutanix +32% Vmware

postgres/pgbench

By default, pgbench tests a scenario that is loosely based on TPC-B, involving five SELECT, UPDATE, and INSERT commands per transaction. However, it is easy to test other cases by writing your own transaction script files.

When executing pgbench, some of the key variables are:

- -c = number of clients
- -j = number of threads
- -t = number of transactions each client runs

For our testing we executed the following scenarios:

- **sudo -u postgres pgbench -c 10 -j 10 -t 10000 pgbench | tee /root/_pg-logs/10c-10t-10k.txt**
- **sudo -u postgres pgbench -c 25 -j 10 -t 10000 pgbench | tee /root/_pg-logs/25c-10t-10k.txt**
- **sudo -u postgres pgbench -c 50 -j 10 -t 10000 pgbench | tee /root/_pg-logs/50c-10t-10k.txt**

These variables and parameters will specifically measure transactions per second and latency across a varying number of how many transactions each client is running.

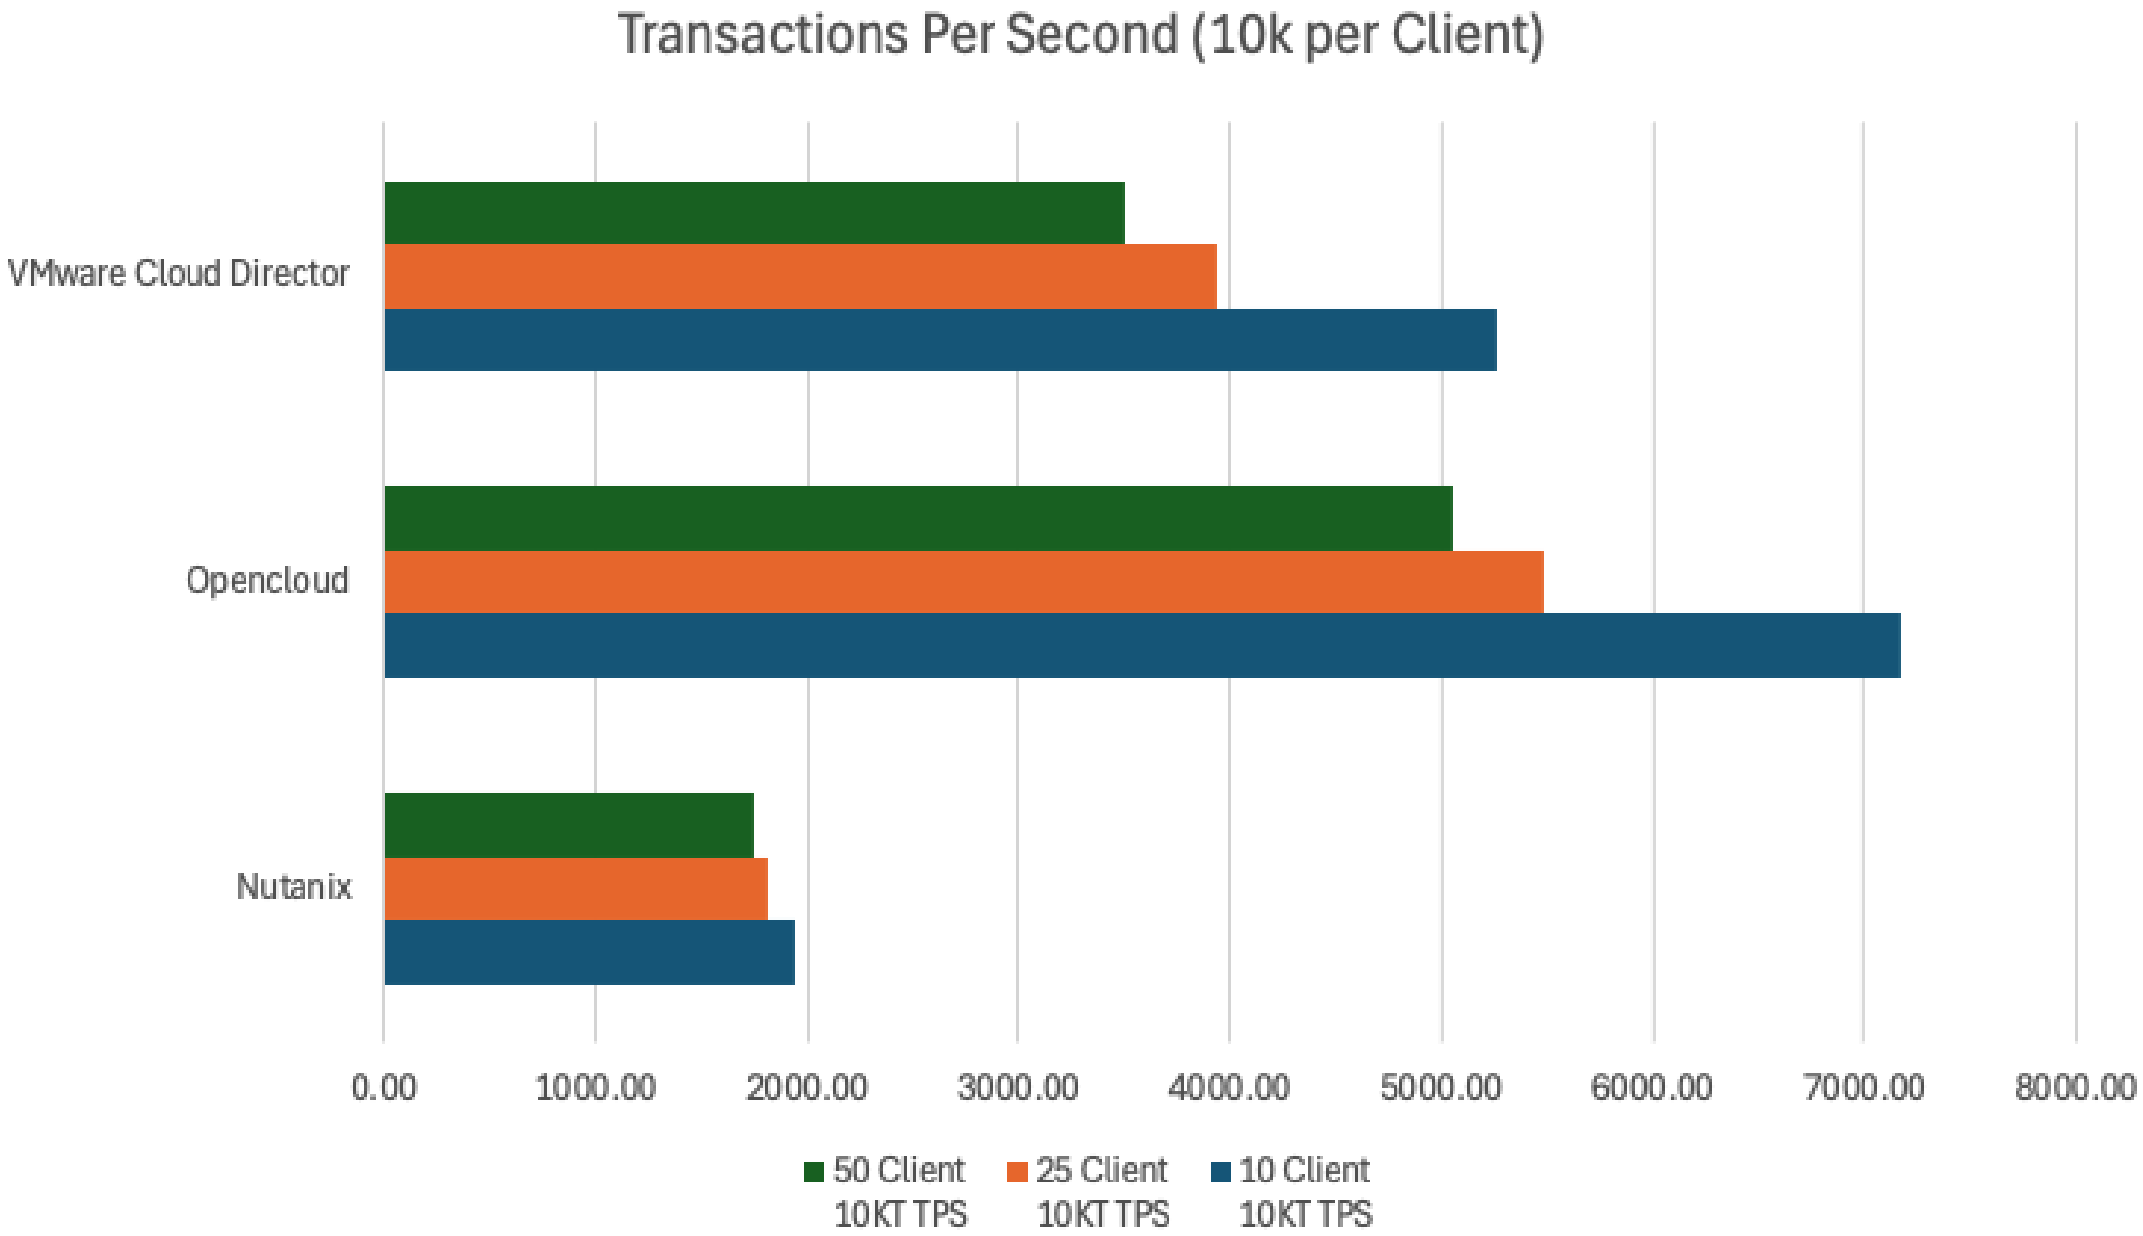

+220% Nutanix +40% Vmware

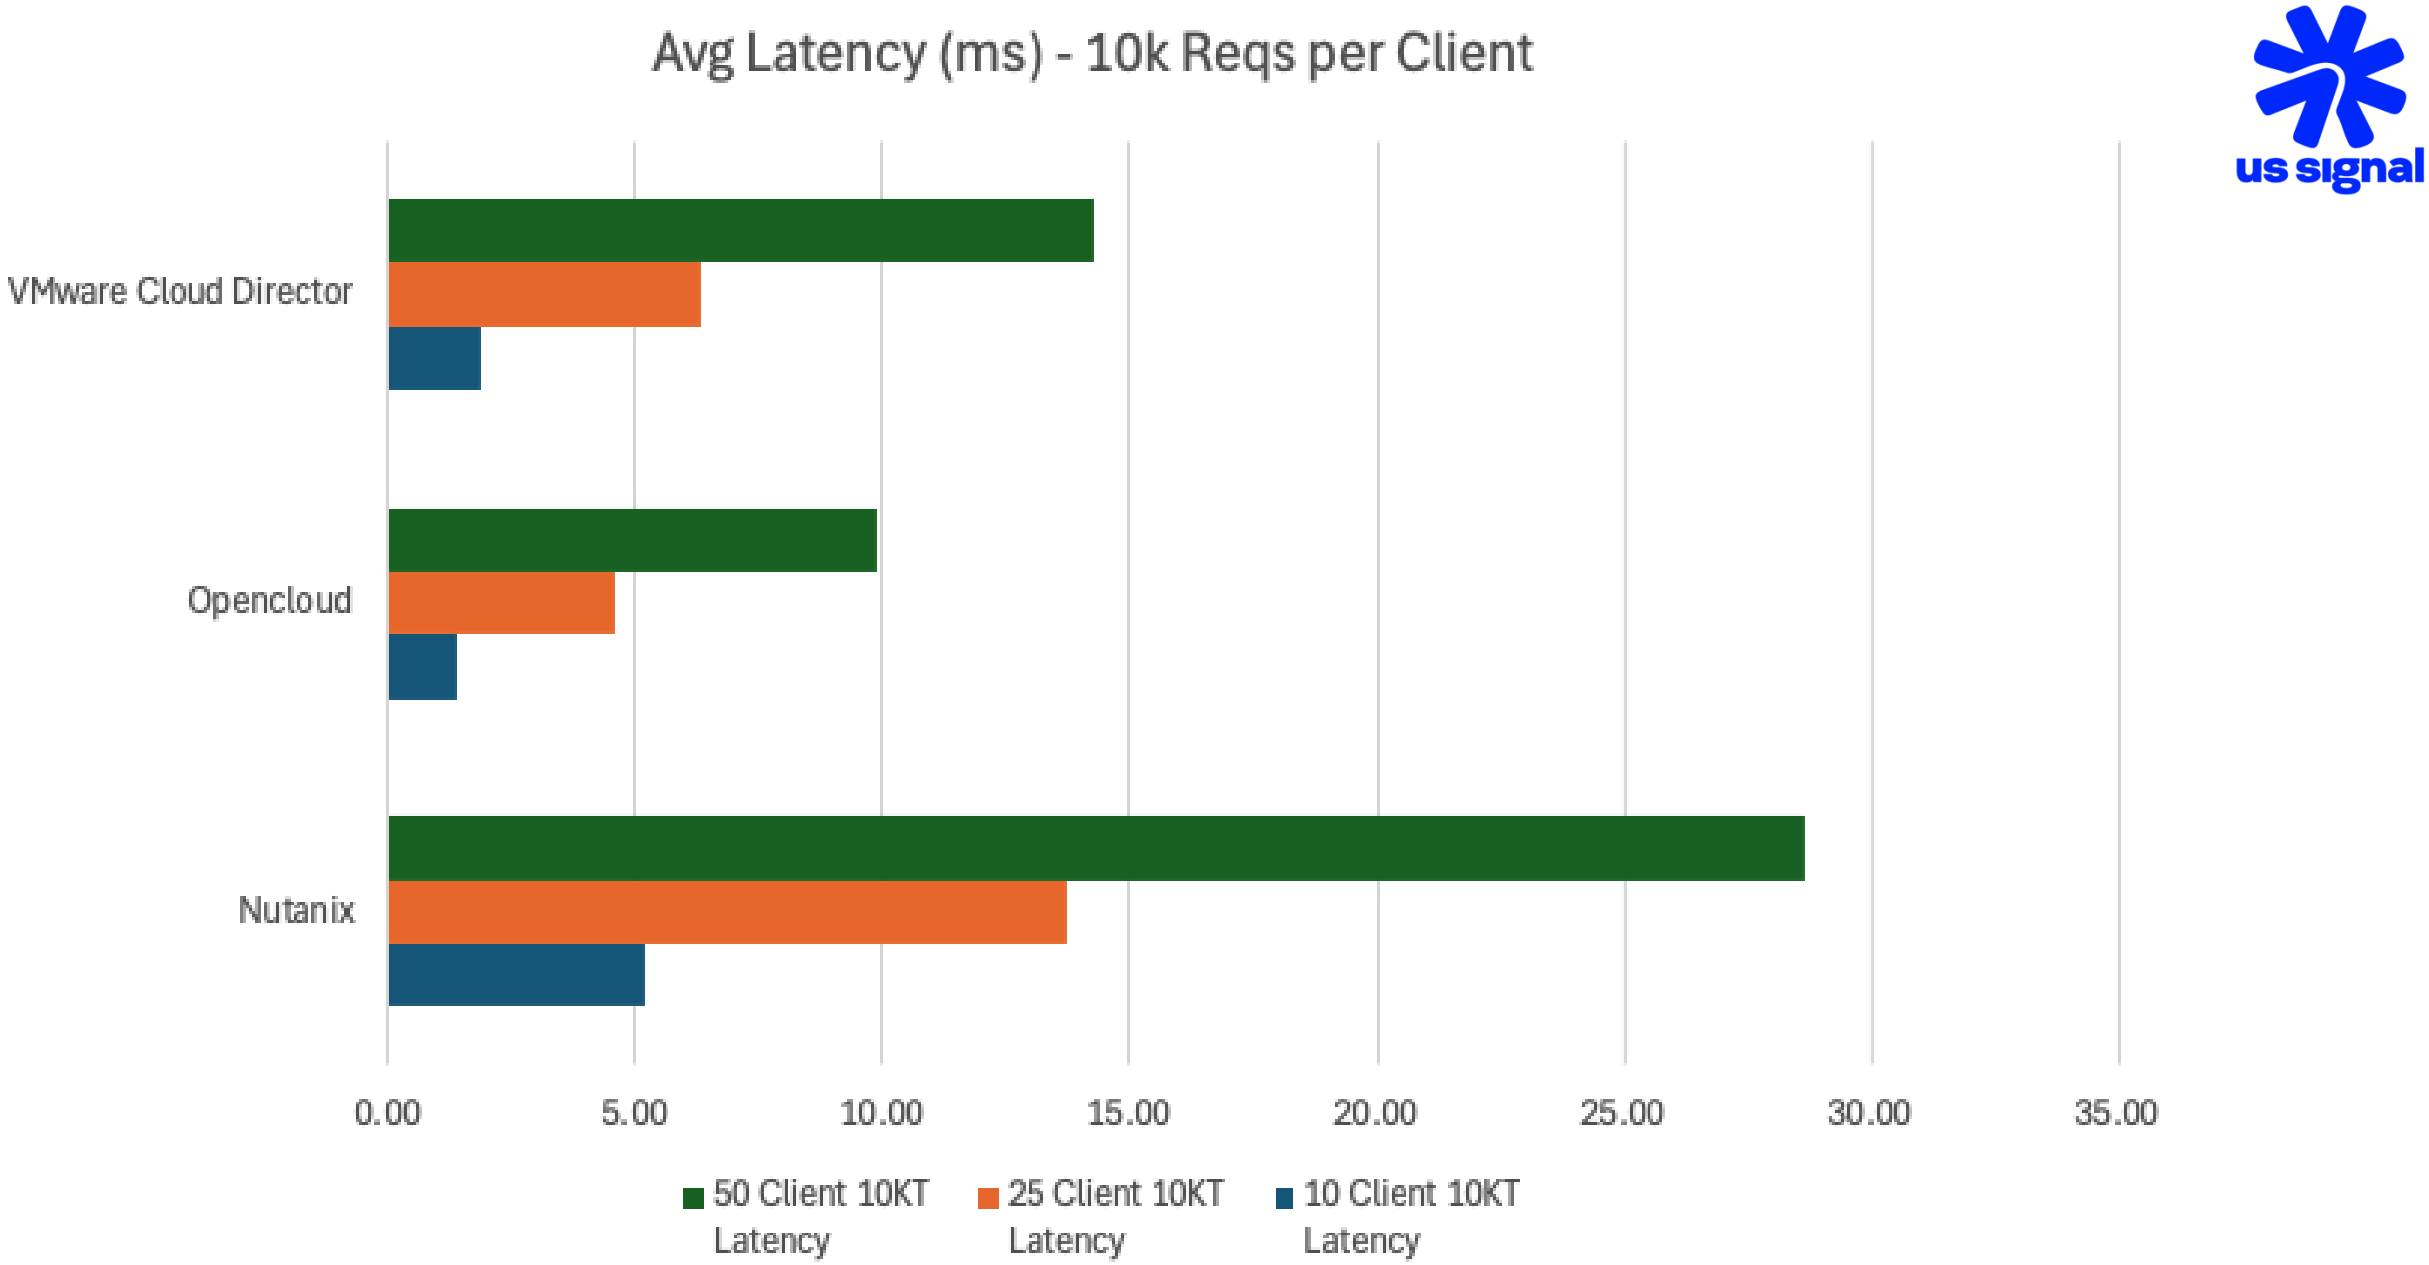

+68% Nutanix +29% Vmware

performance Comparison - public cloud

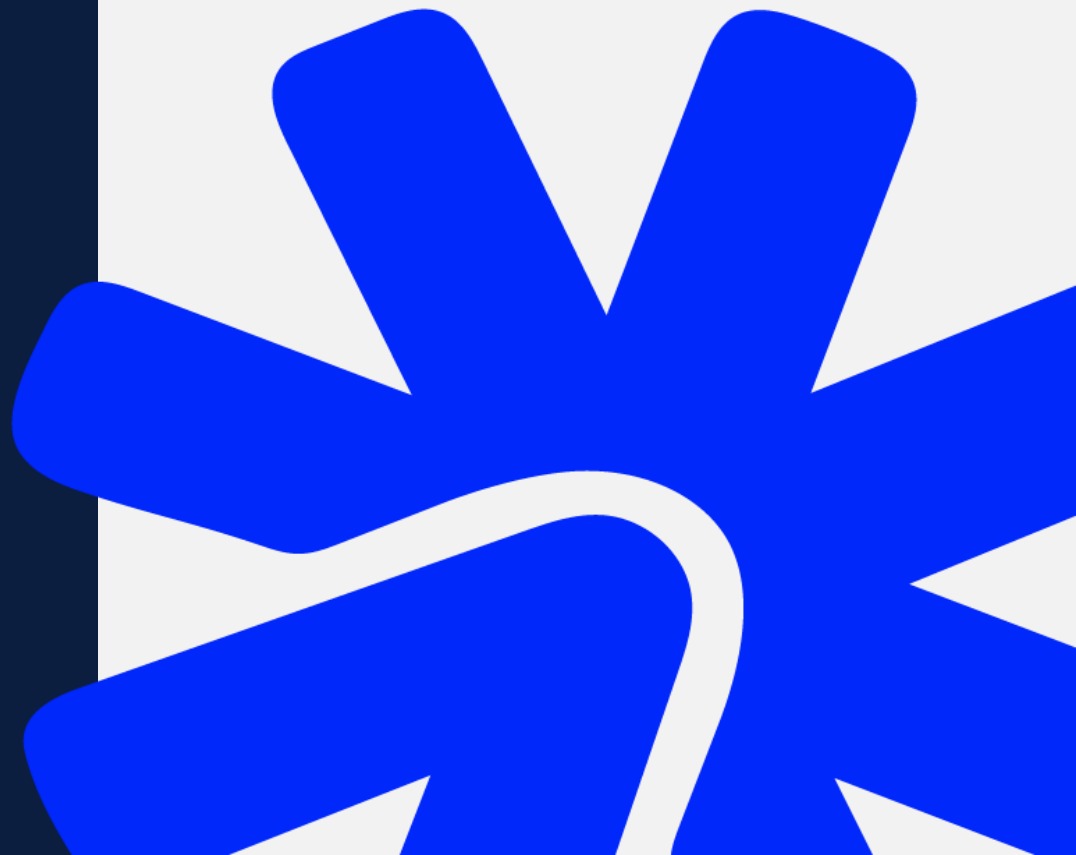

Spec Score

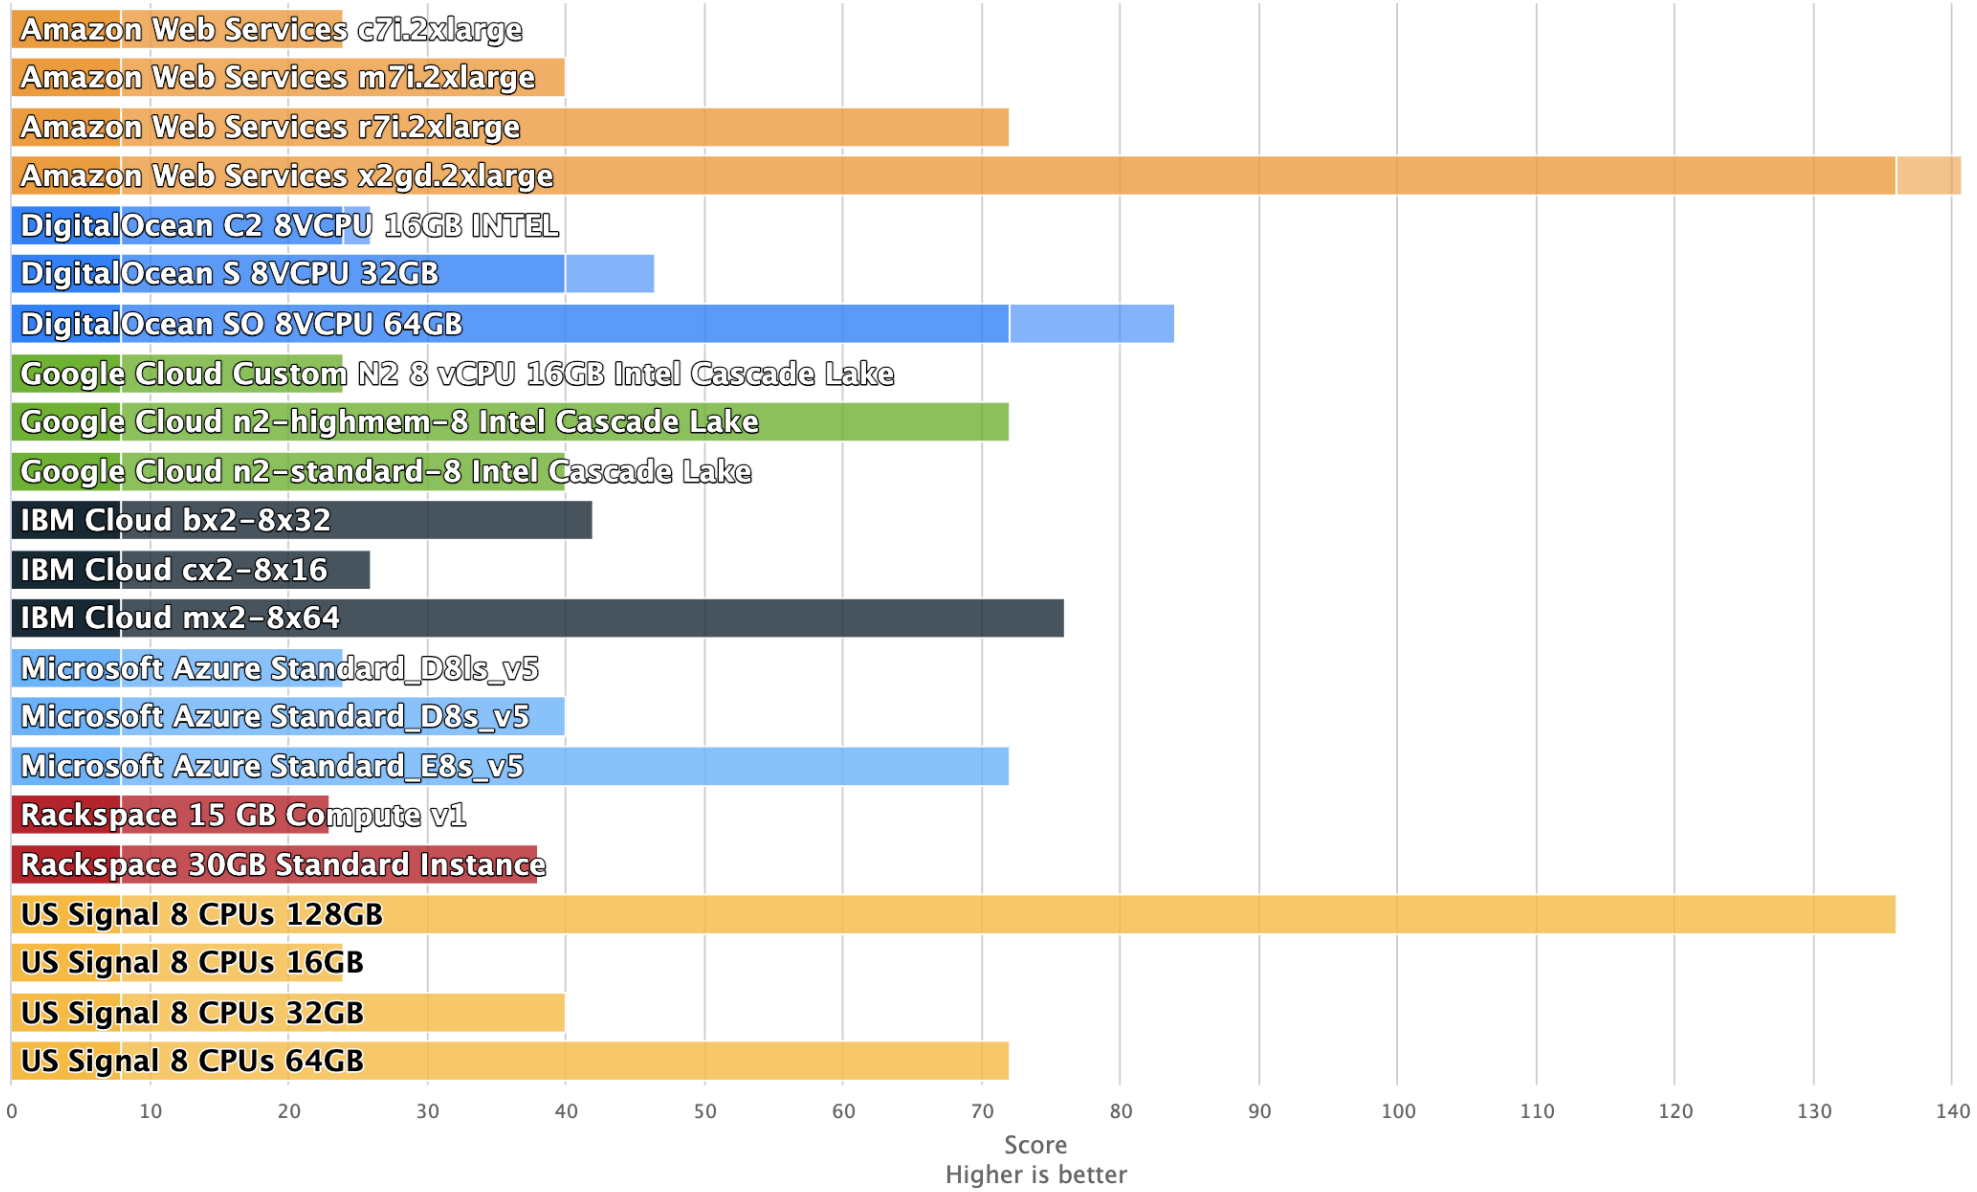
Amazon Web Services c7i.2xlarge
Amazon Web Services m7i.2xlarge
Amazon Web Services r7i.2xlarge
Amazon Web Services x2gd.2xlarge
DigitalOcean C2 8VCPU 16GB INTEL
DigitalOcean S 8VCPU 32GB
DigitalOcean SO 8VCPU 64GB
Google Cloud Custom N2 8 vCPU 16GB Intel Cascade Lake
Google Cloud n2-highmem-8 Intel Cascade Lake
Google Cloud n2-standard-8 Intel Cascade Lake
IBM Cloud bx2-8x32
IBM Cloud cx2-8x16
IBM Cloud mx2-8x64
Microsoft Azure Standard_D8ls_v5
Microsoft Azure Standard_D8s_v5
Microsoft Azure Standard_E8s_v5
Rackspace 15 GB Compute v1
Rackspace 30GB Standard Instance
US Signal 8 CPUs 128GB
US Signal 8 CPUs 16GB
US Signal 8 CPUs 32GB
US Signal 8 CPUs 64GB
0
10
20
30
40
50
60
70
80
90
100
110
120
130
140
Score
Higher is better

Geekbench 6 – general purpose

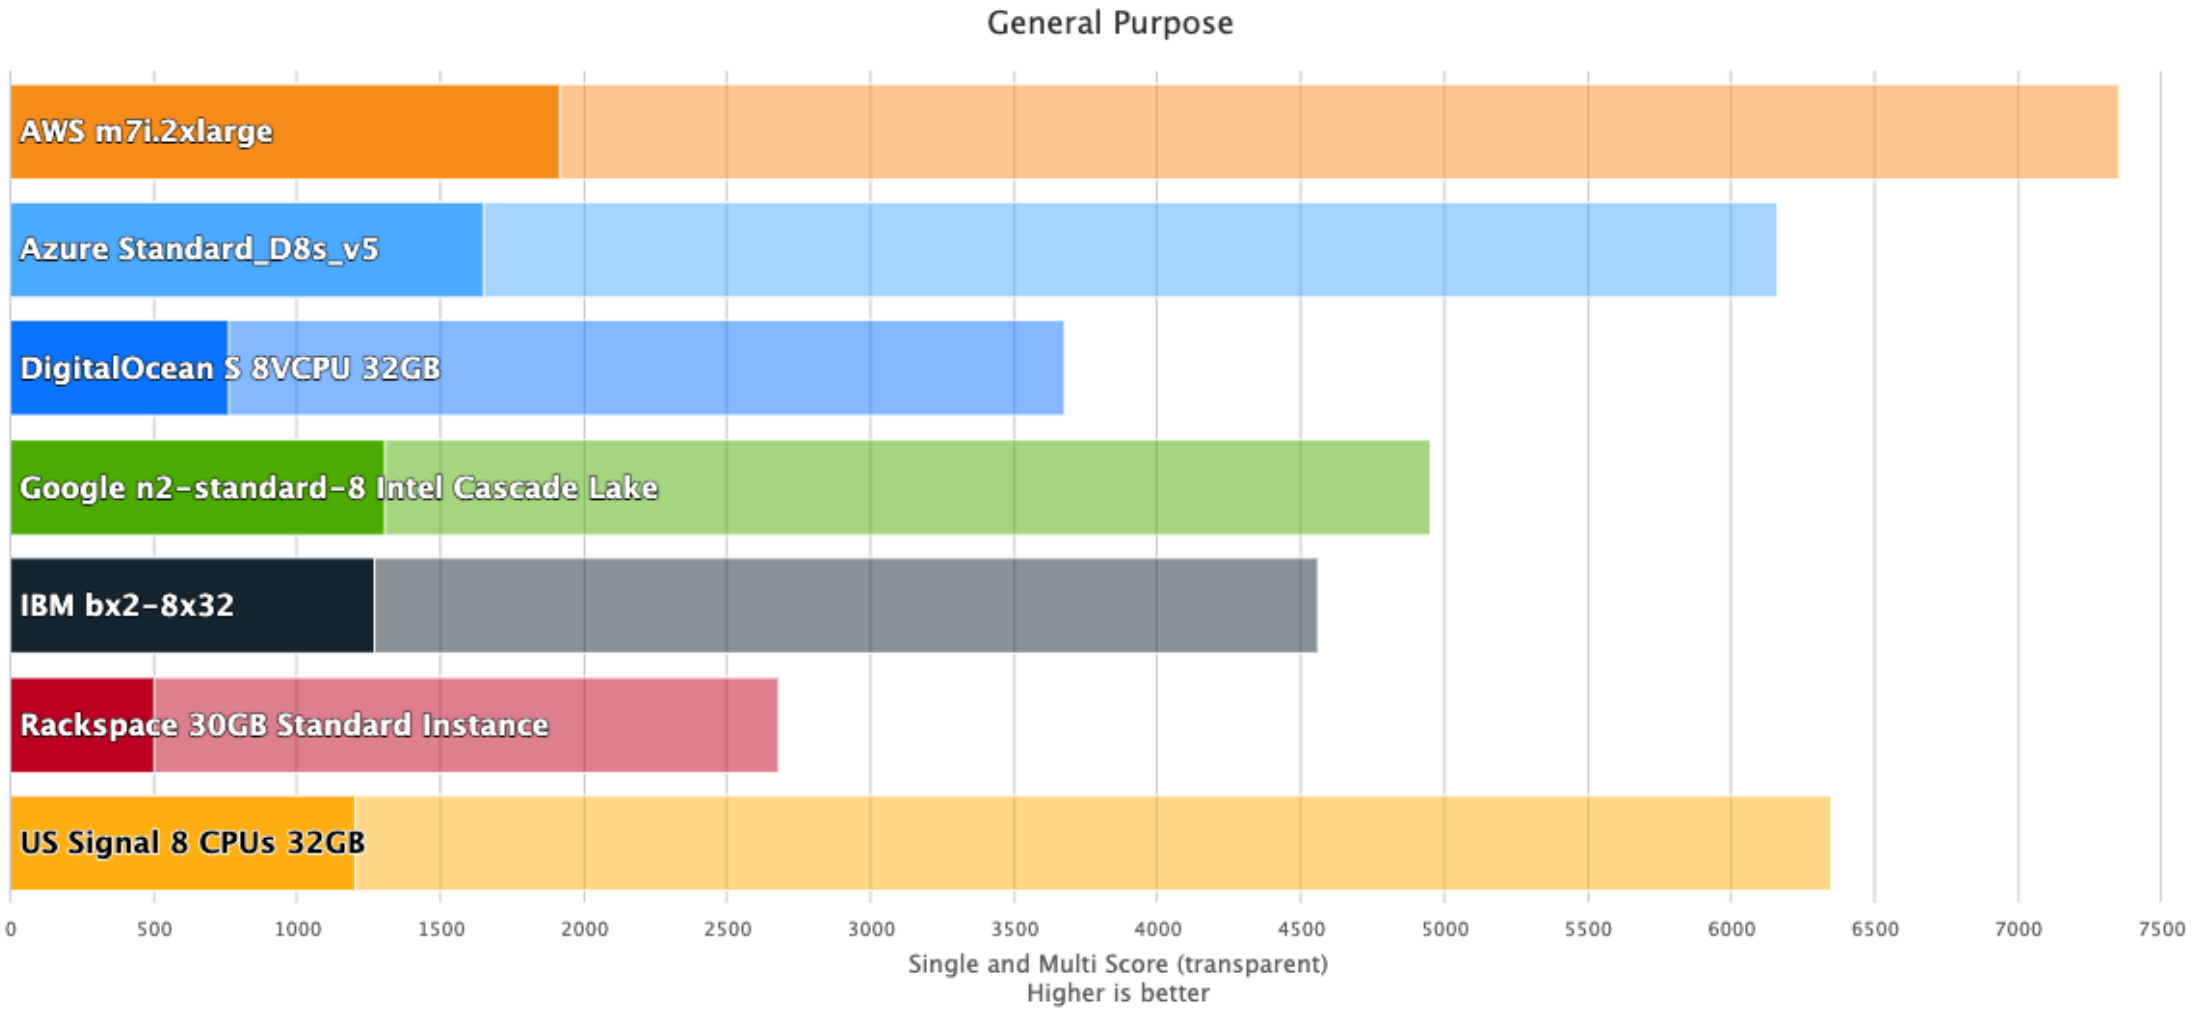

Geekbench 6 – compute optimized

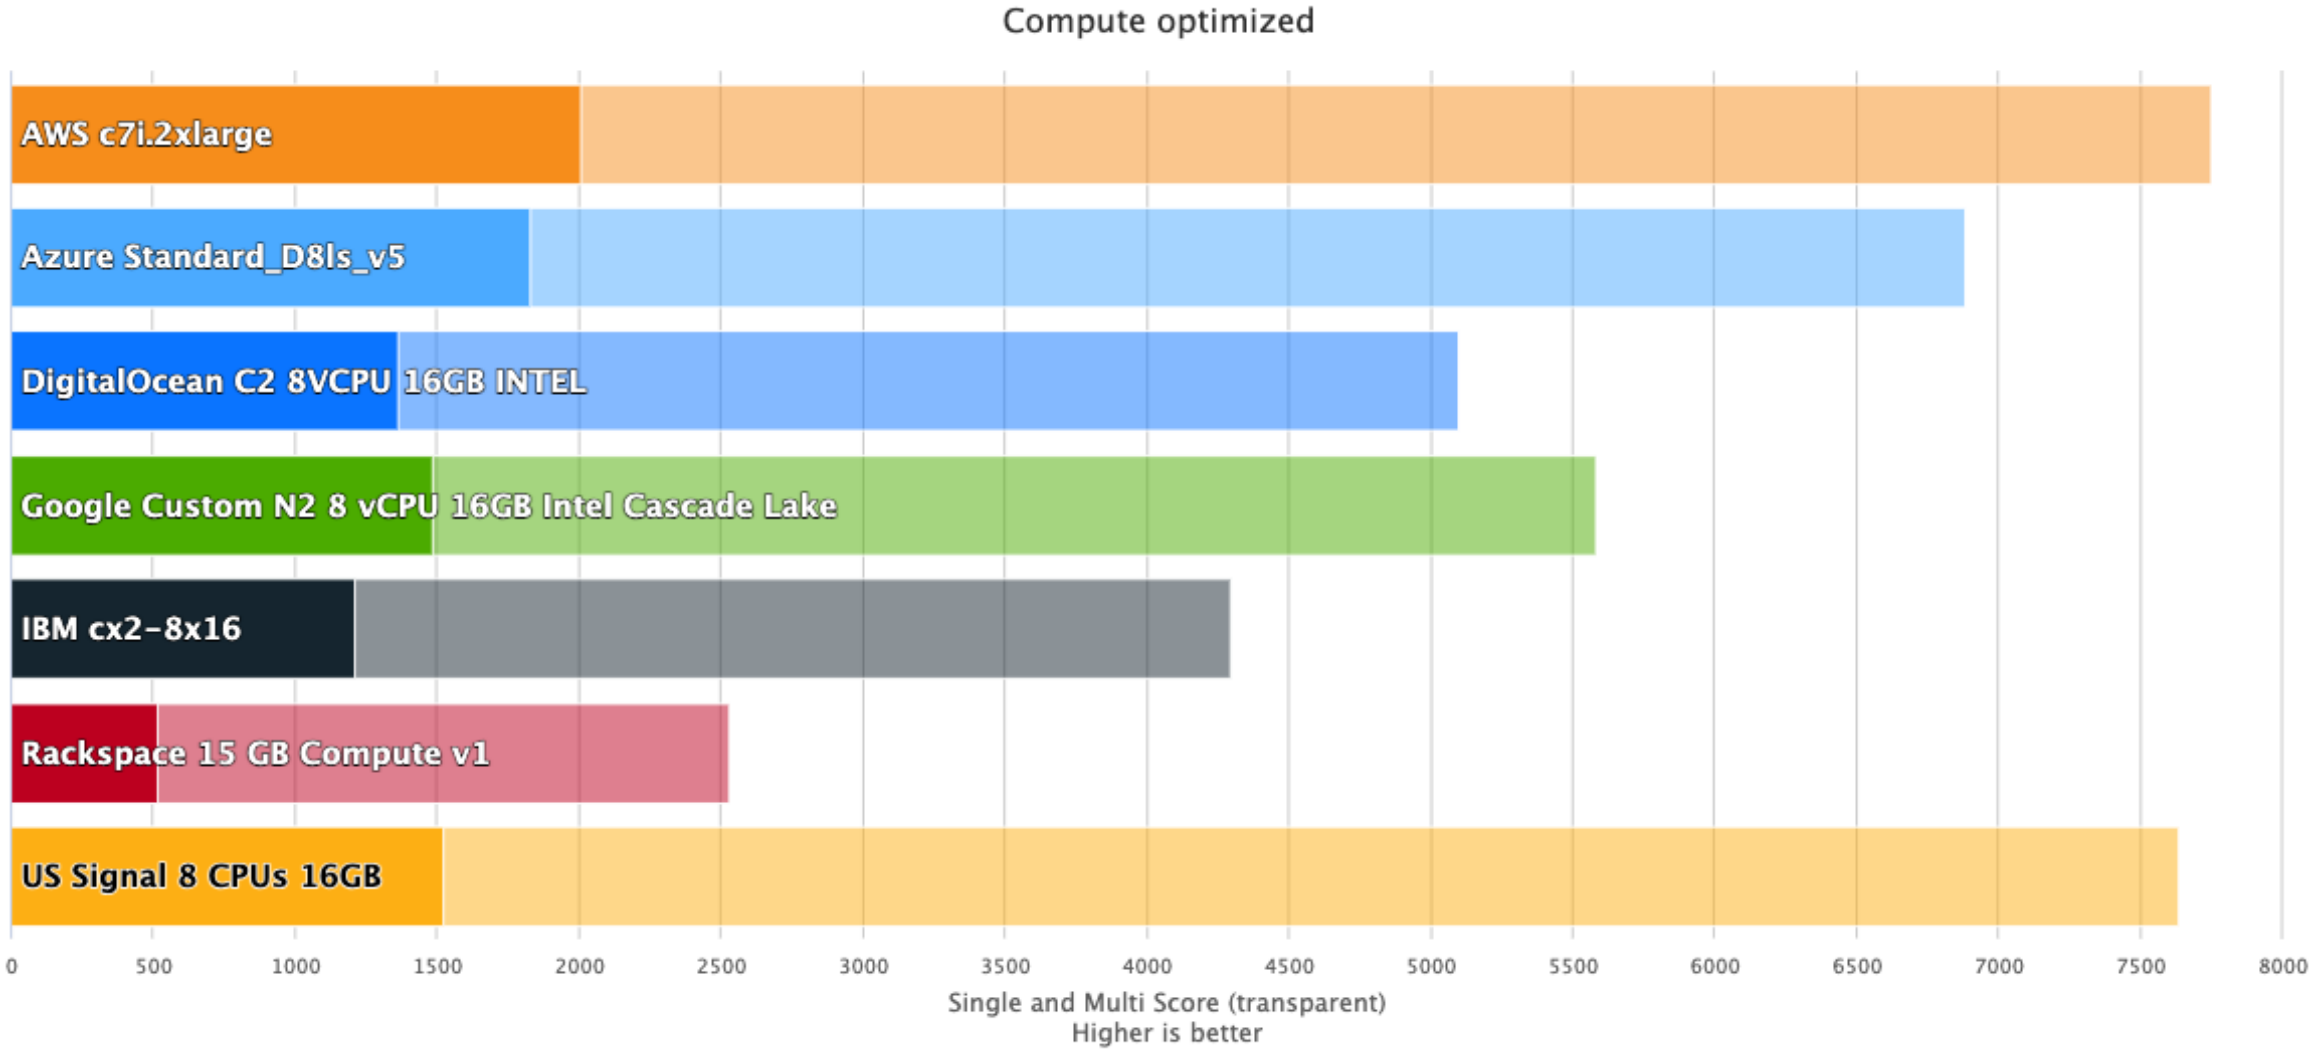

Geekbench 6 – memory optimized

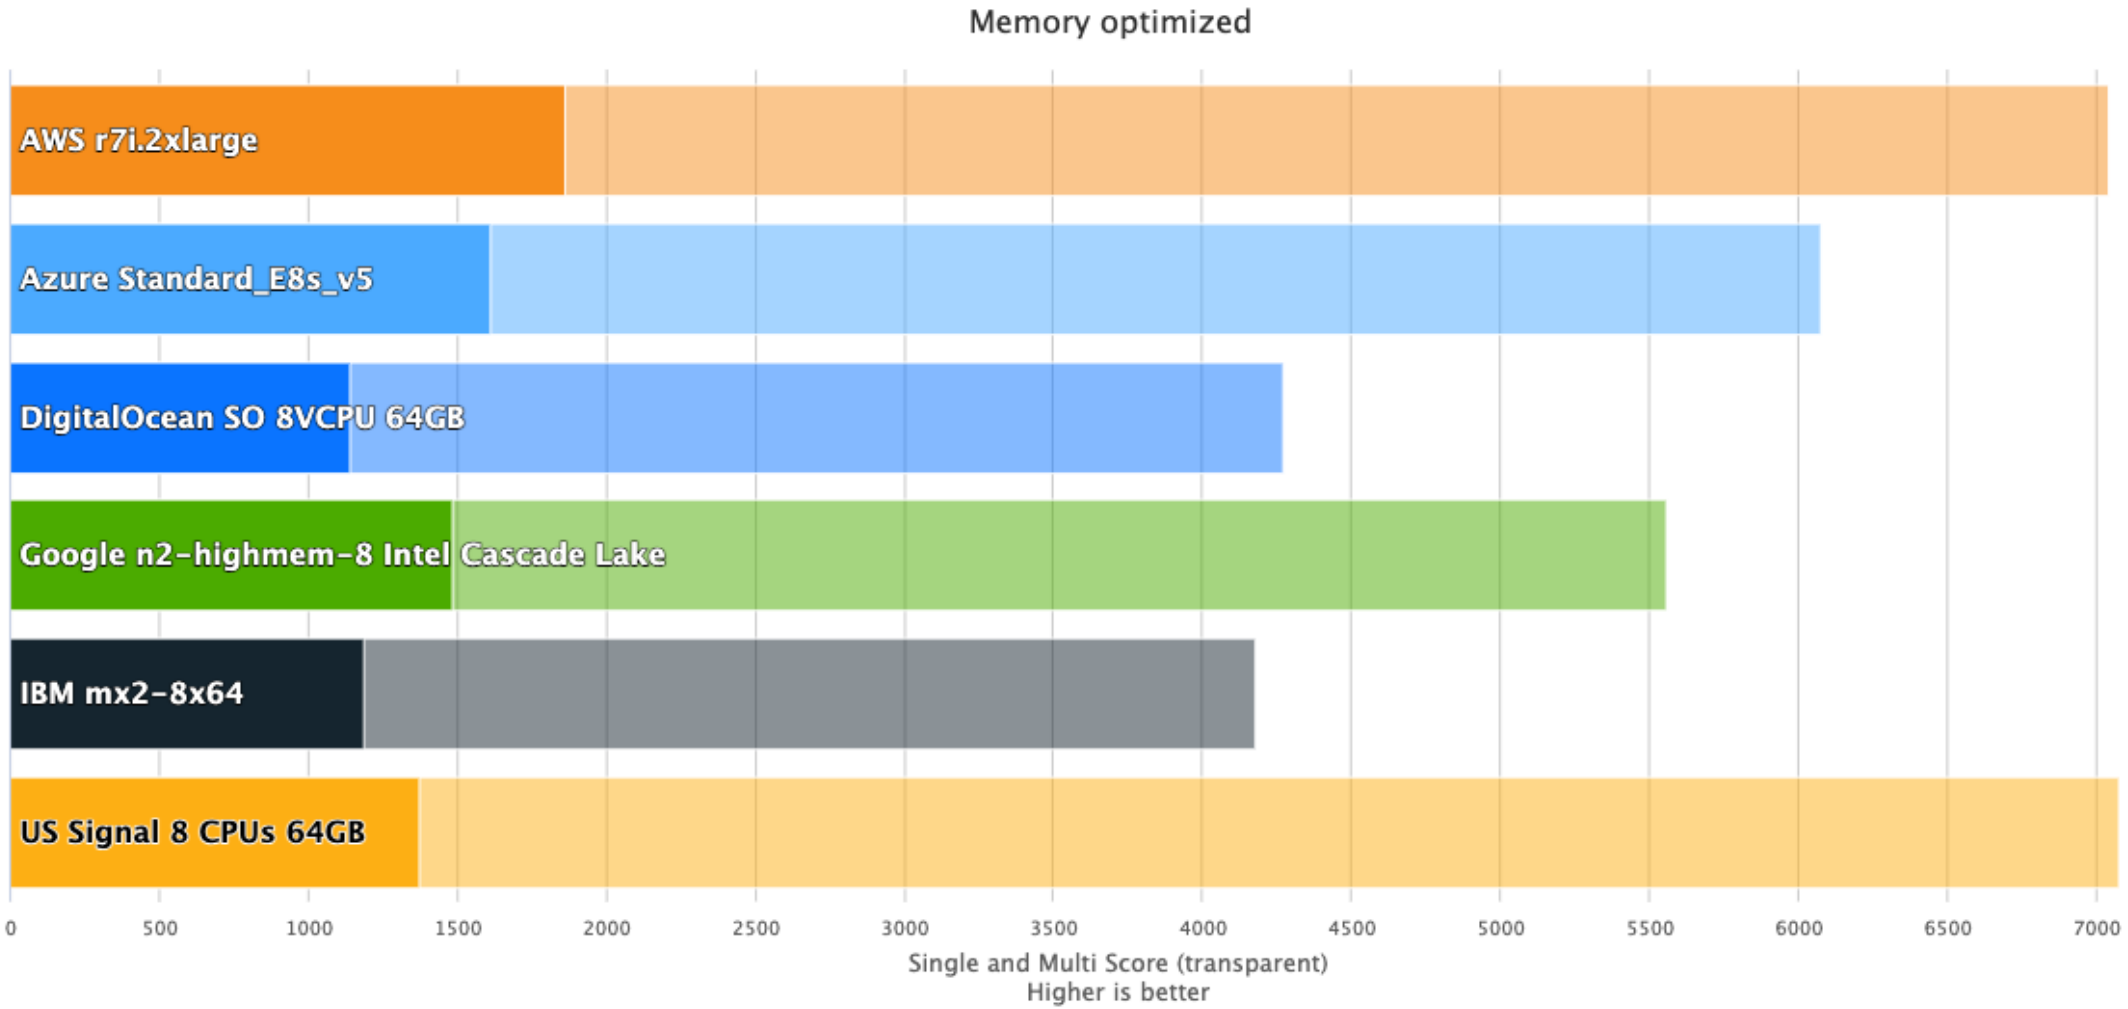

Geekbench 6 – price/performance

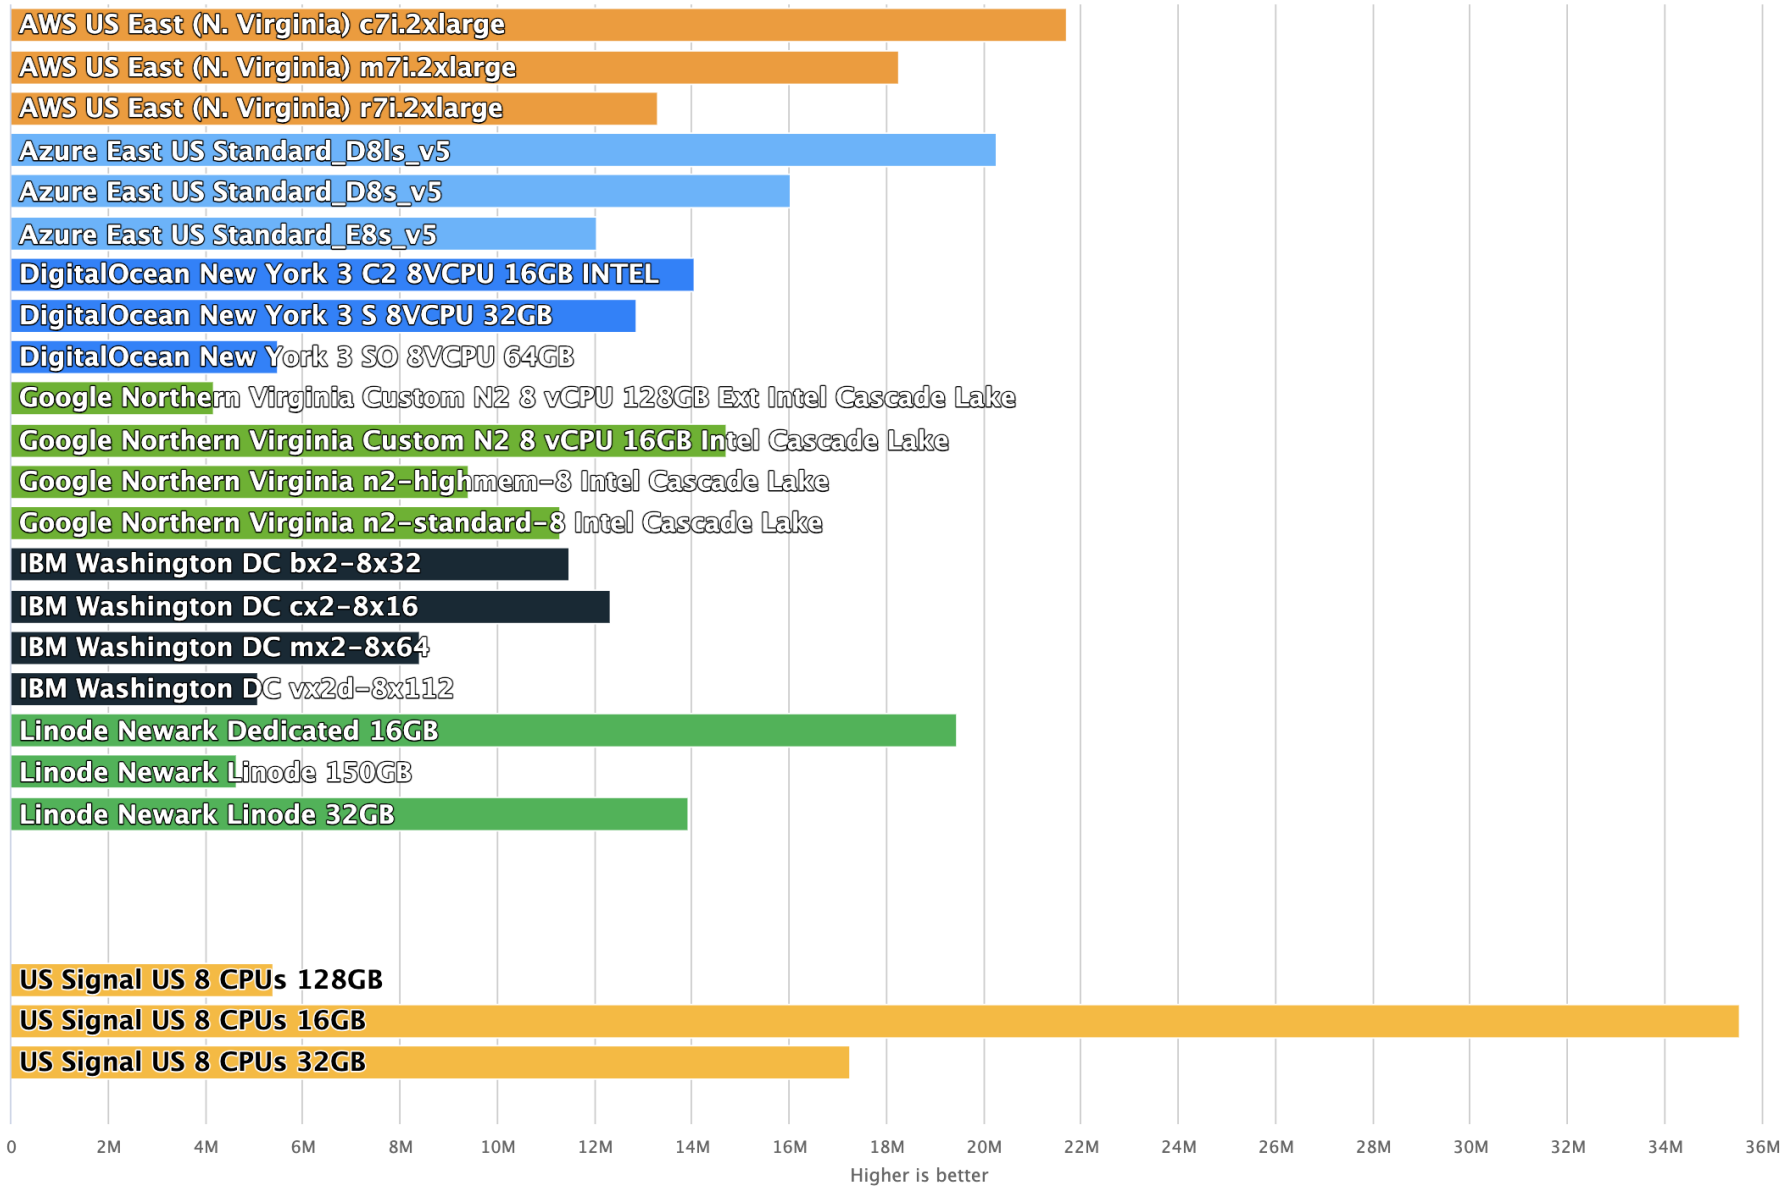

OpenSSL Speed

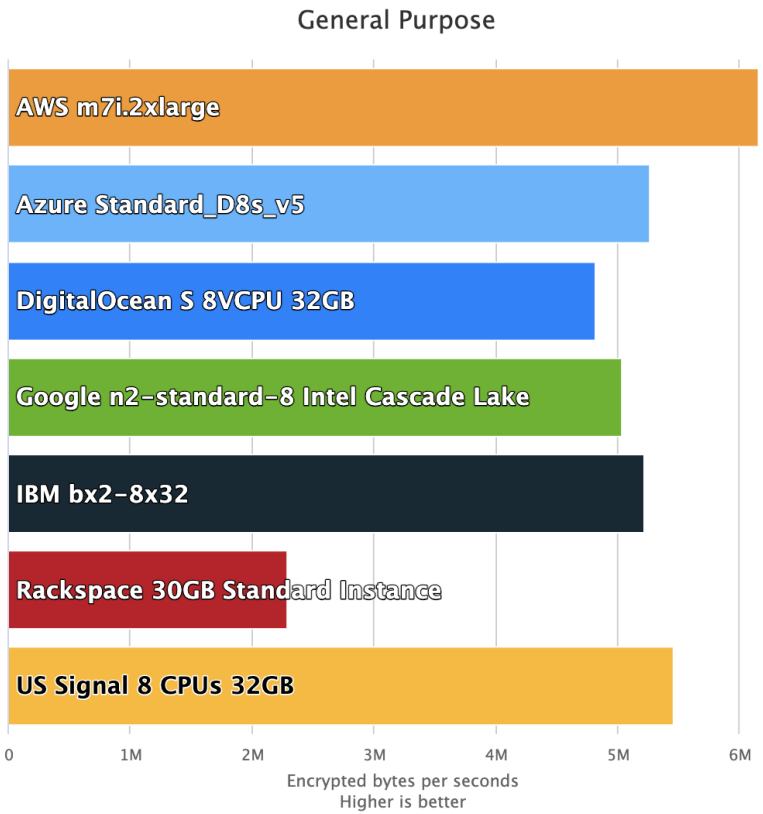

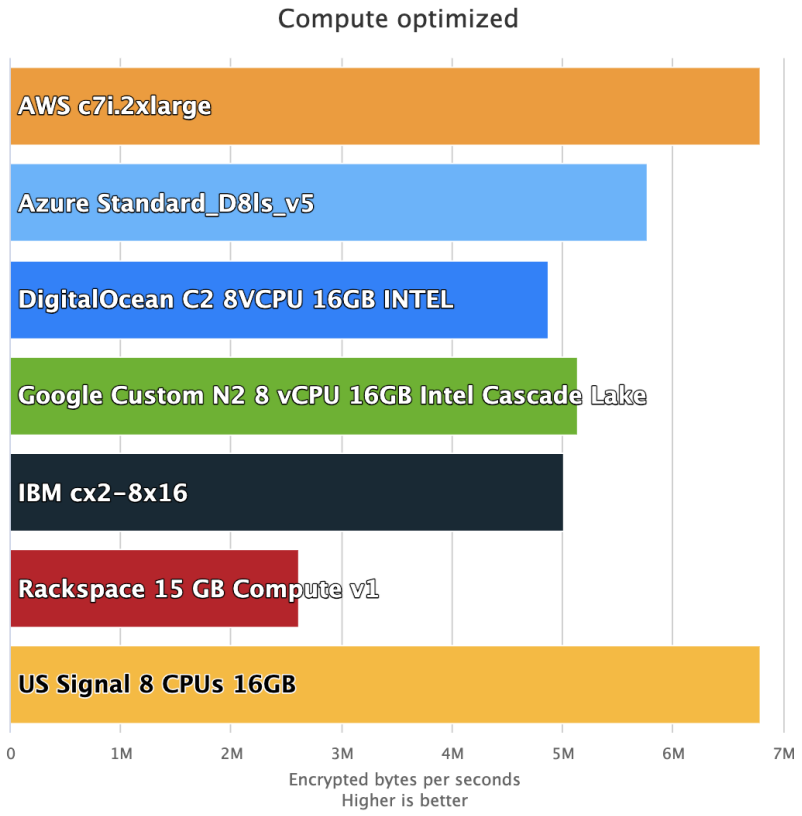

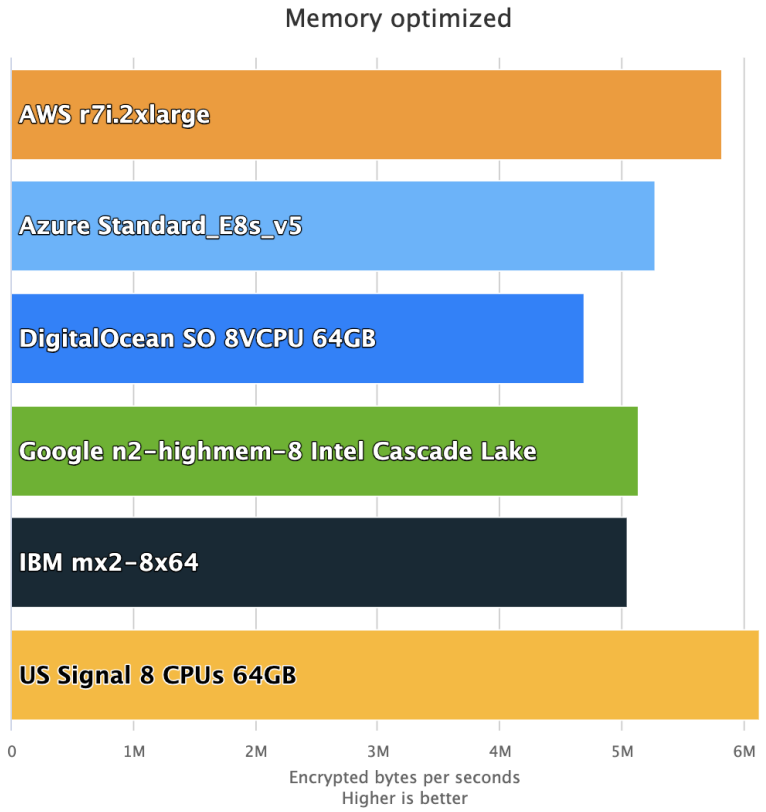

Network Latency – general purpose

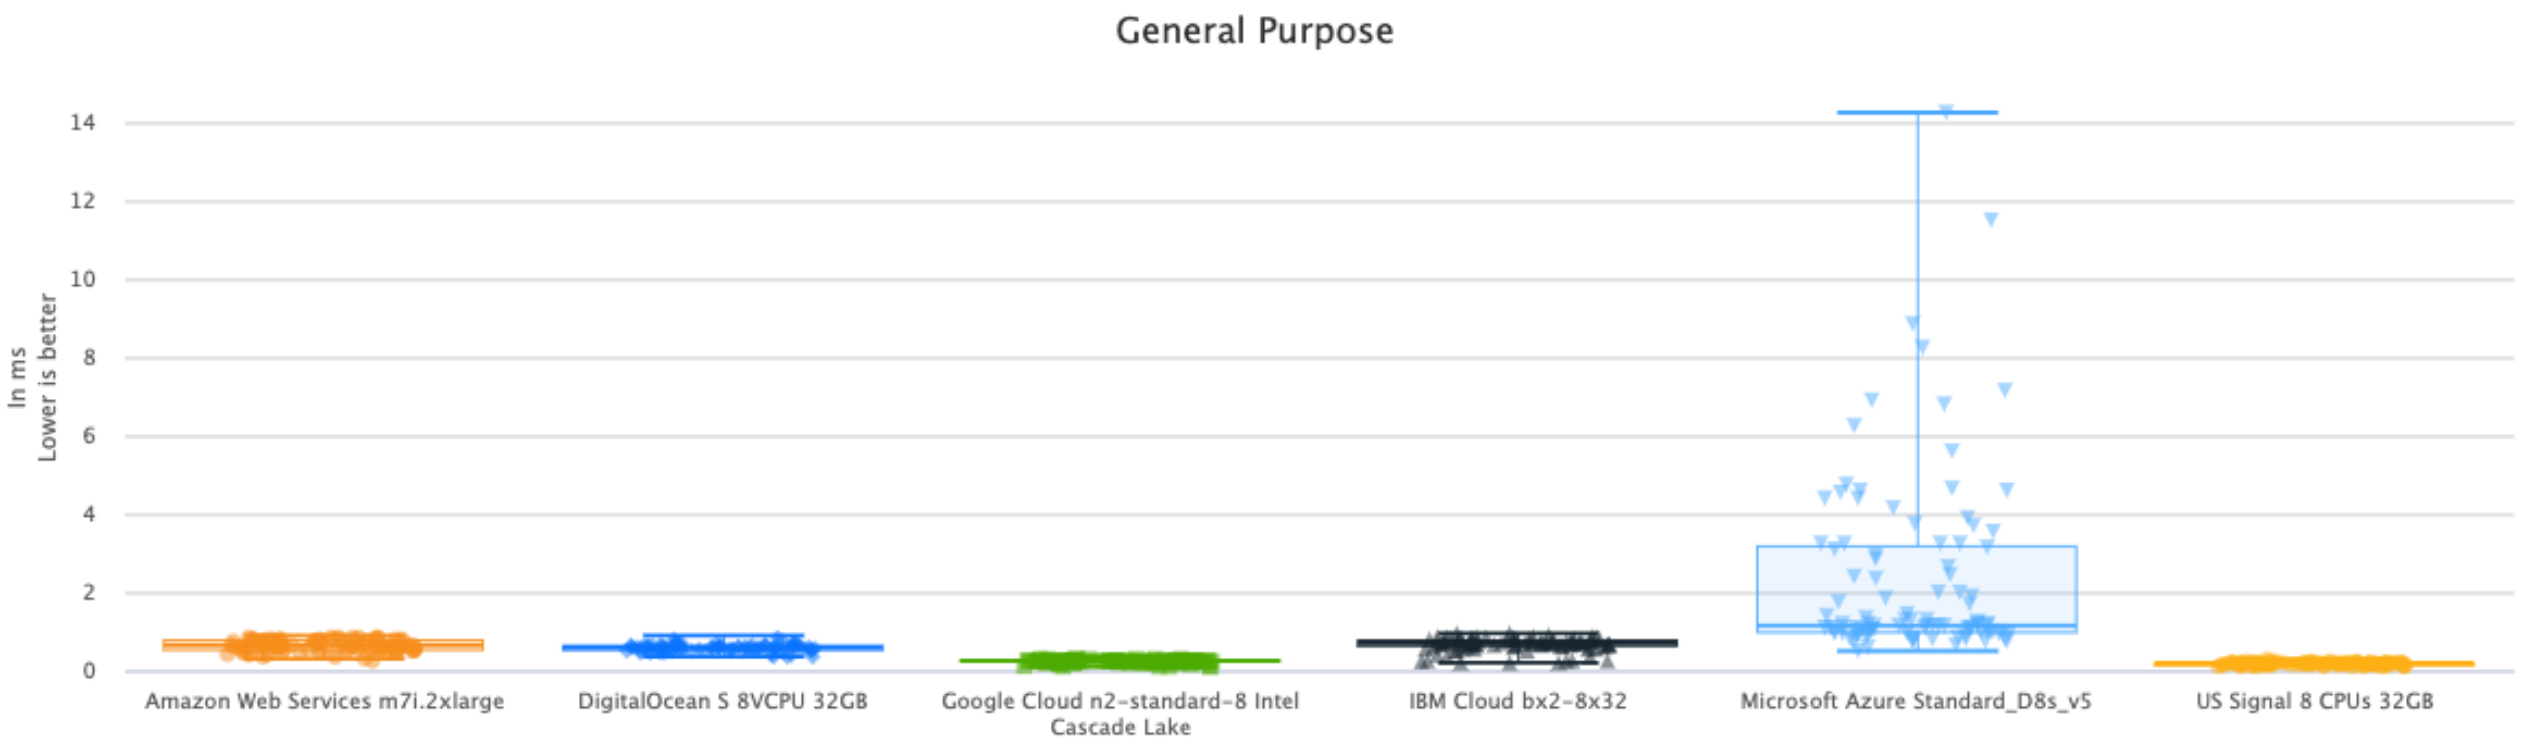

Network Latency – compute optimized

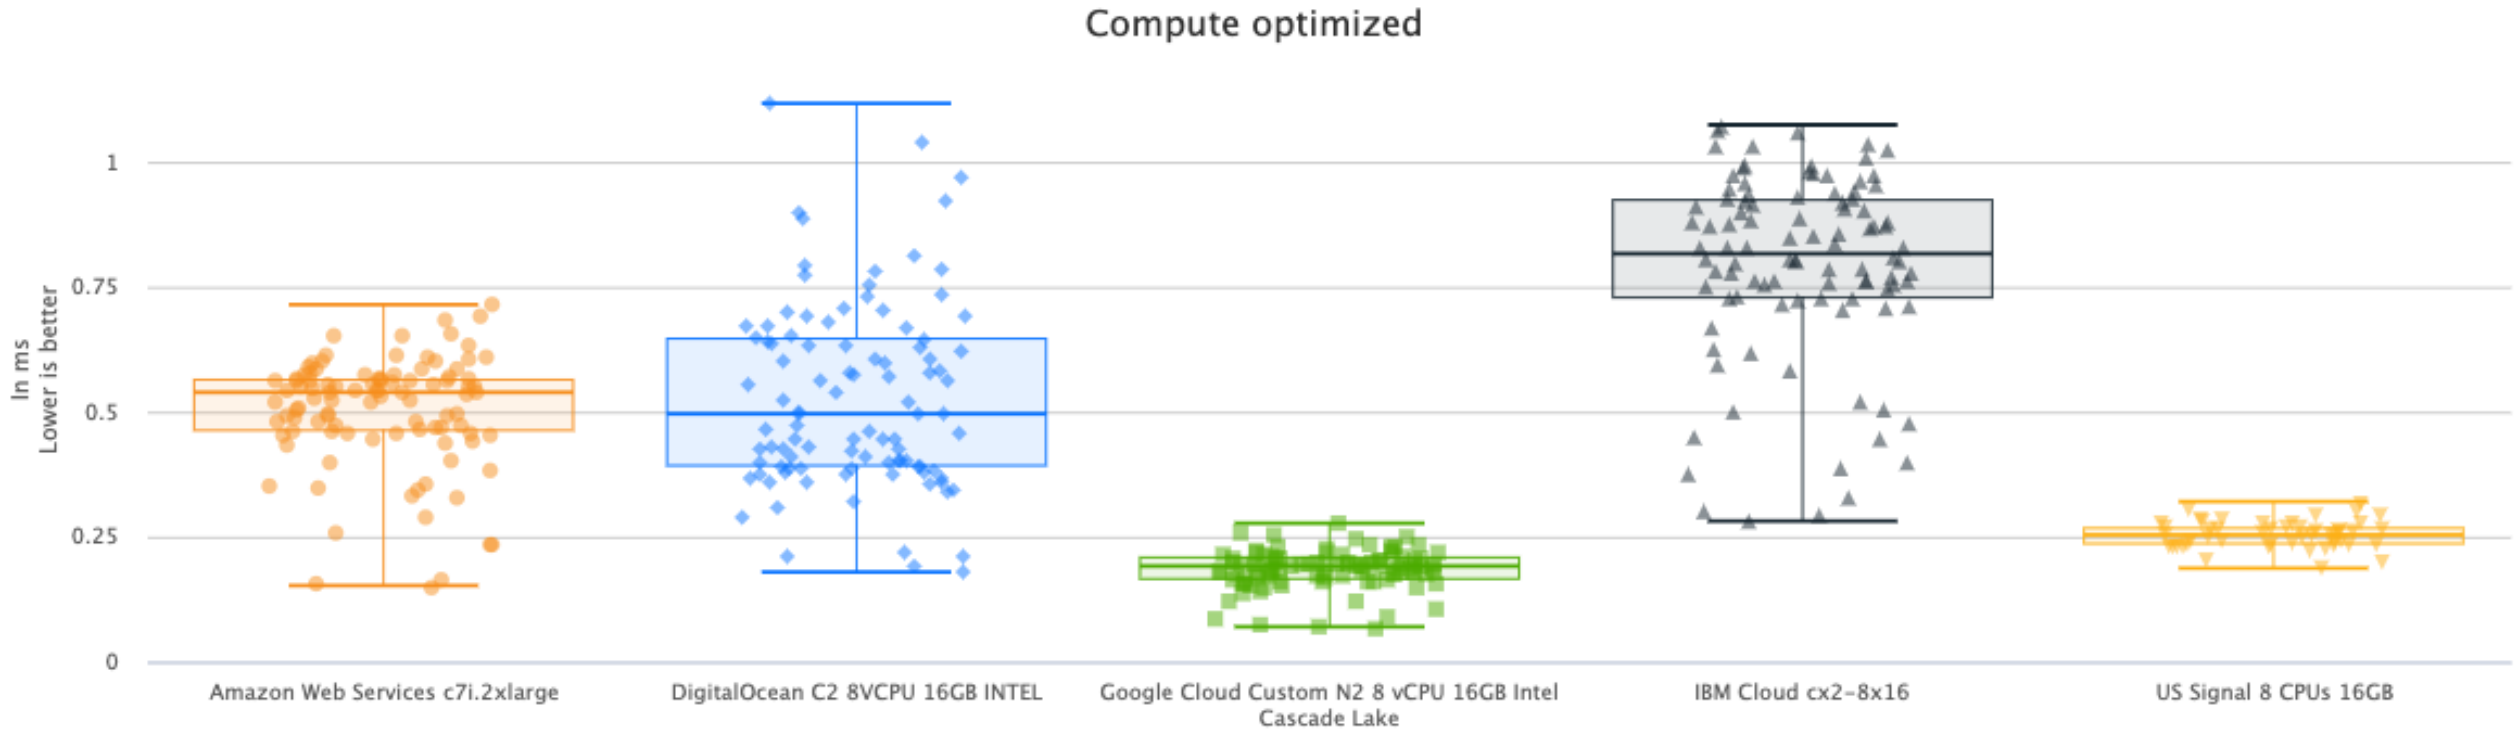

Network Latency – memory optimized

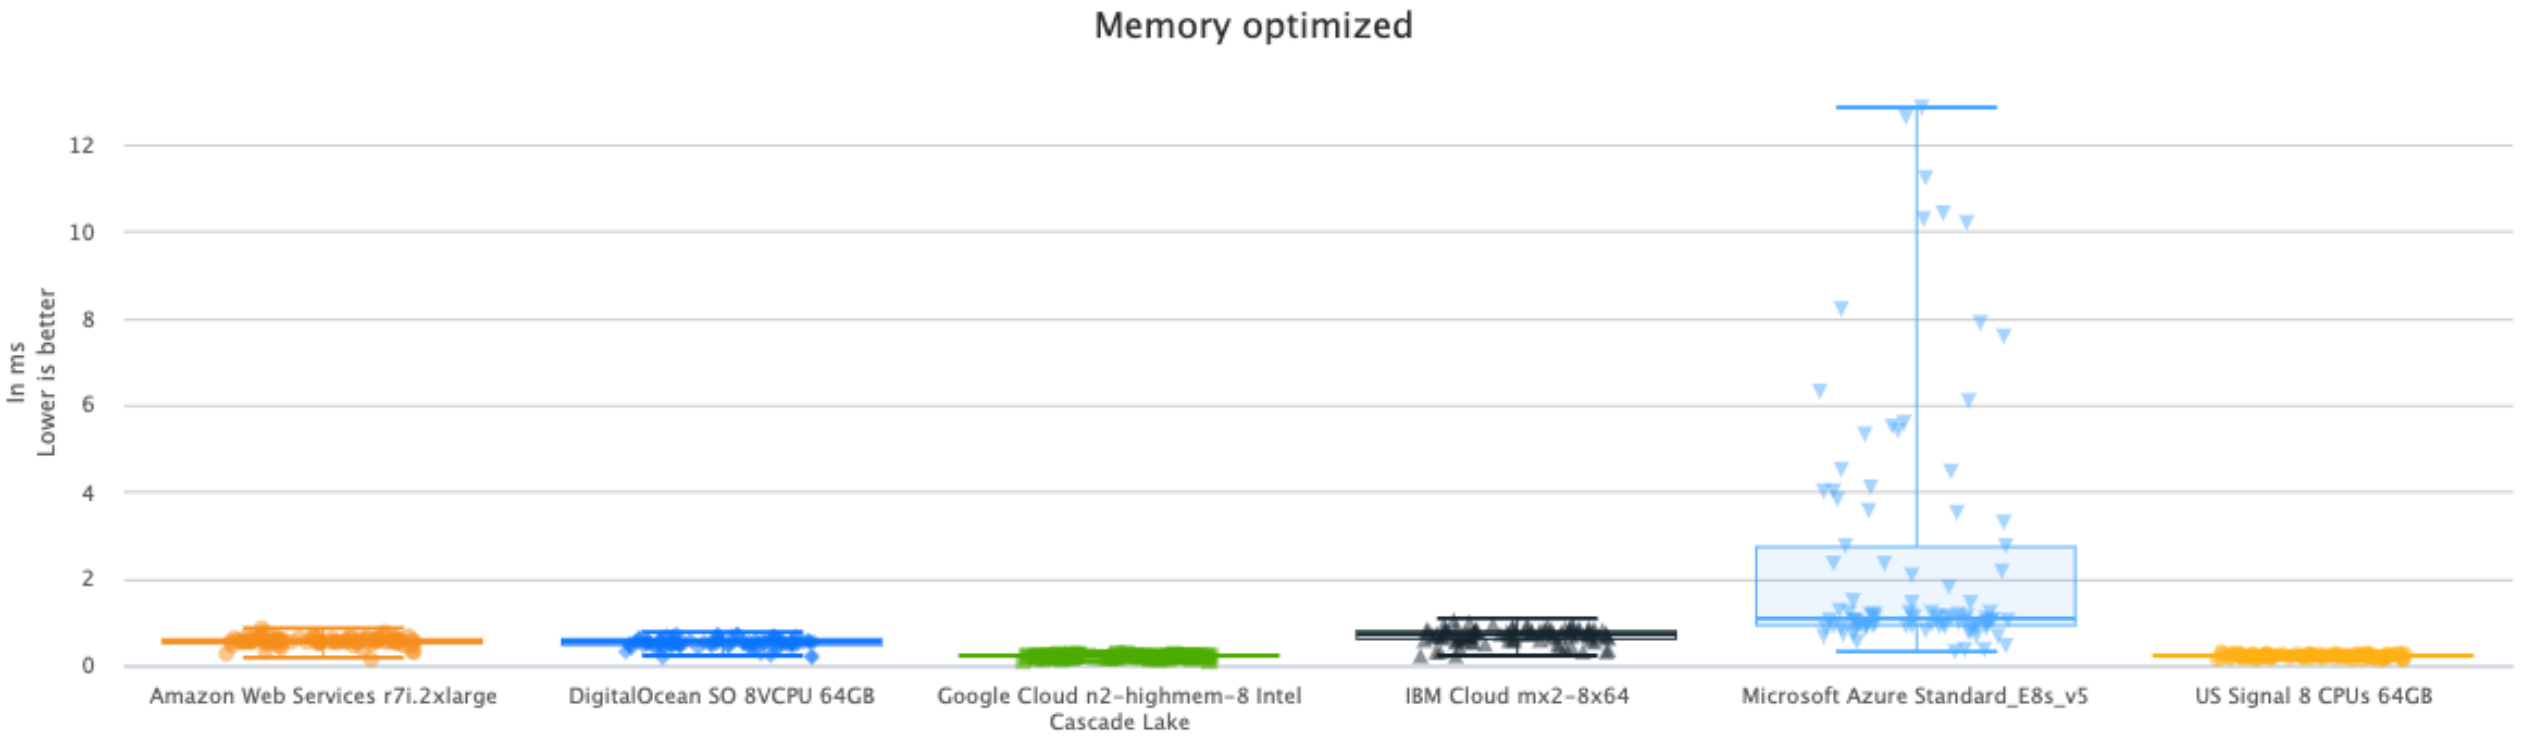

Network Bandwidth - general

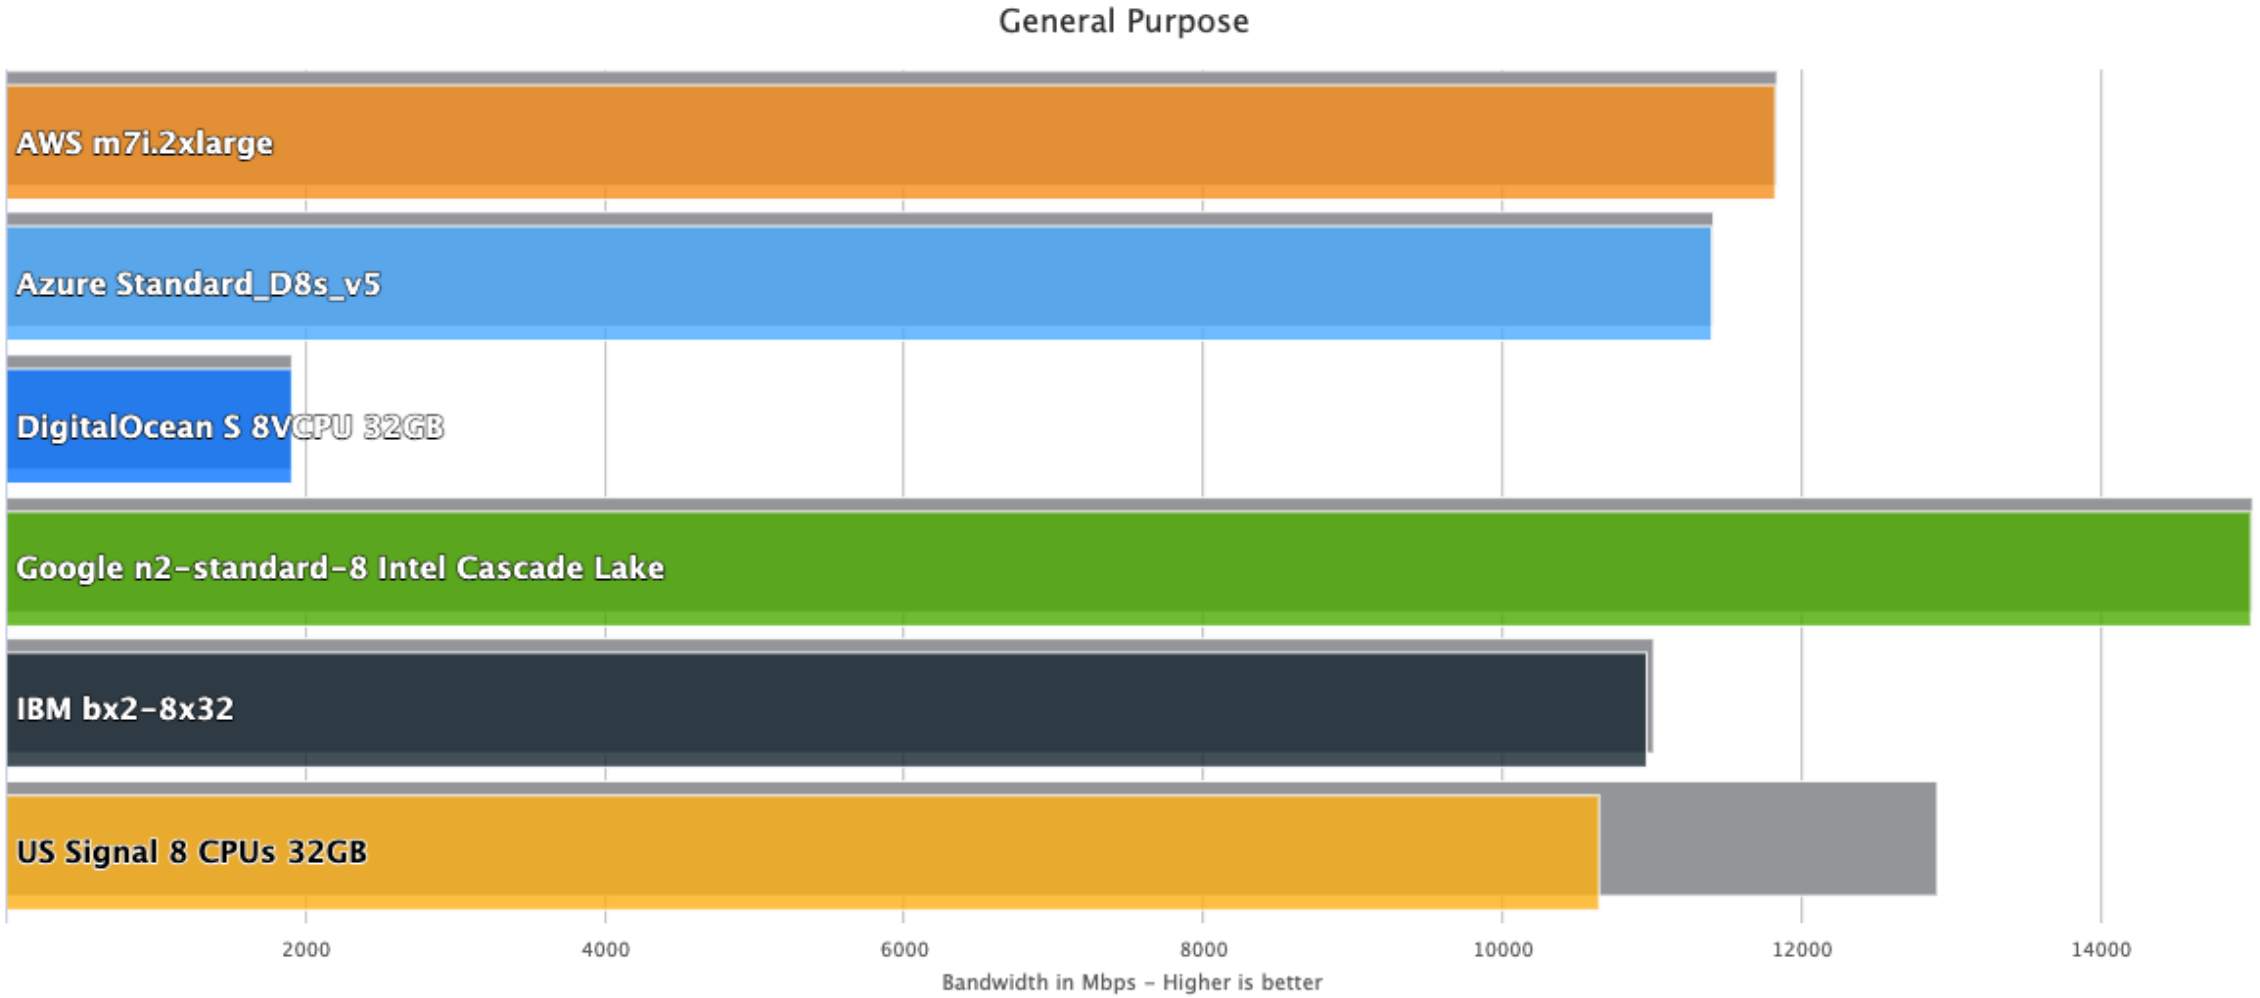

Network Bandwidth – compute optimized

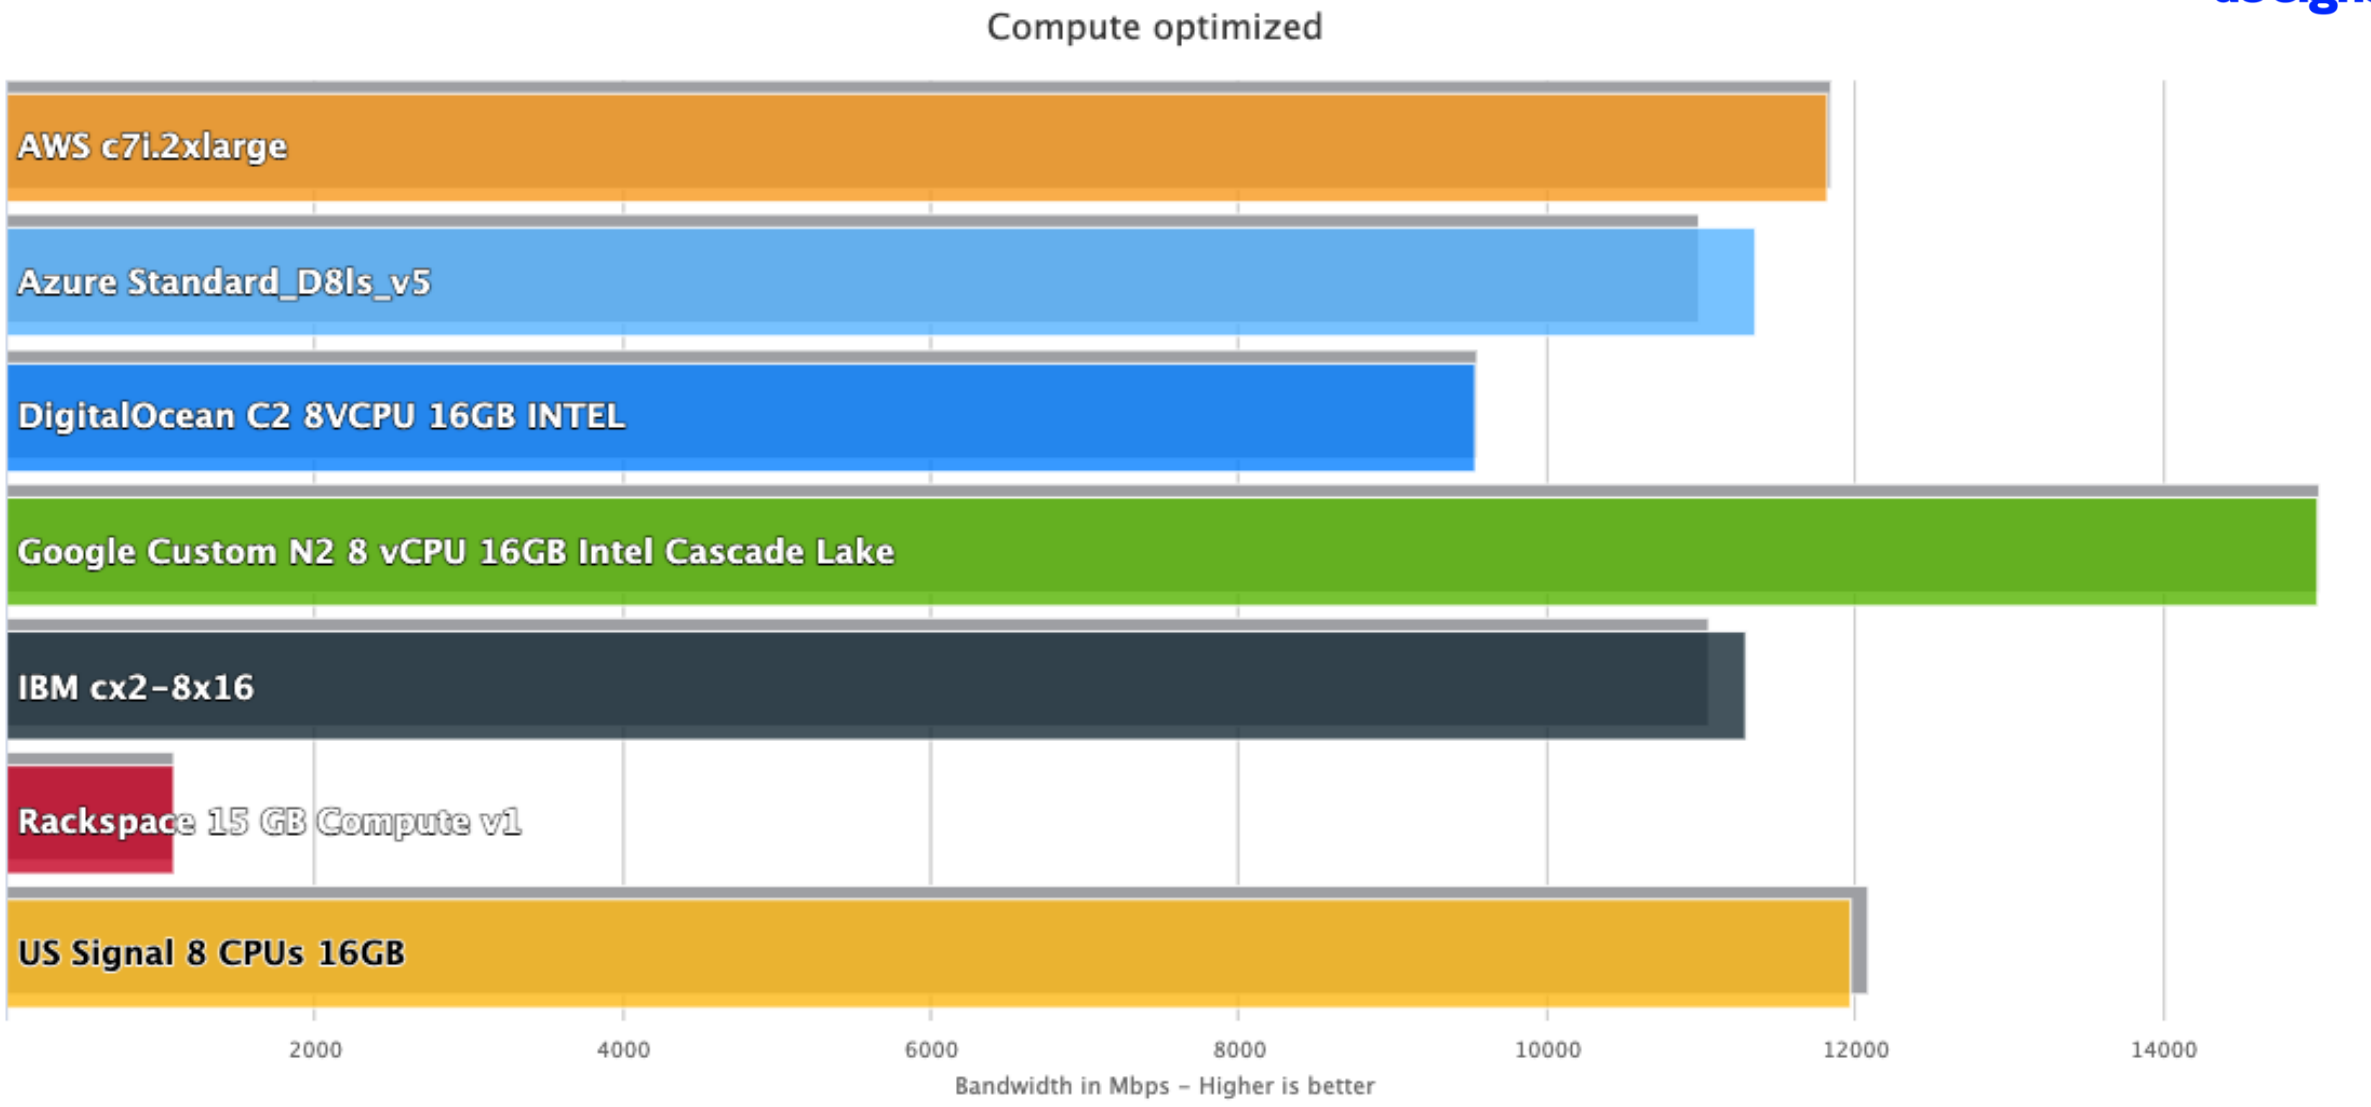

Network Bandwidth – memory optimized

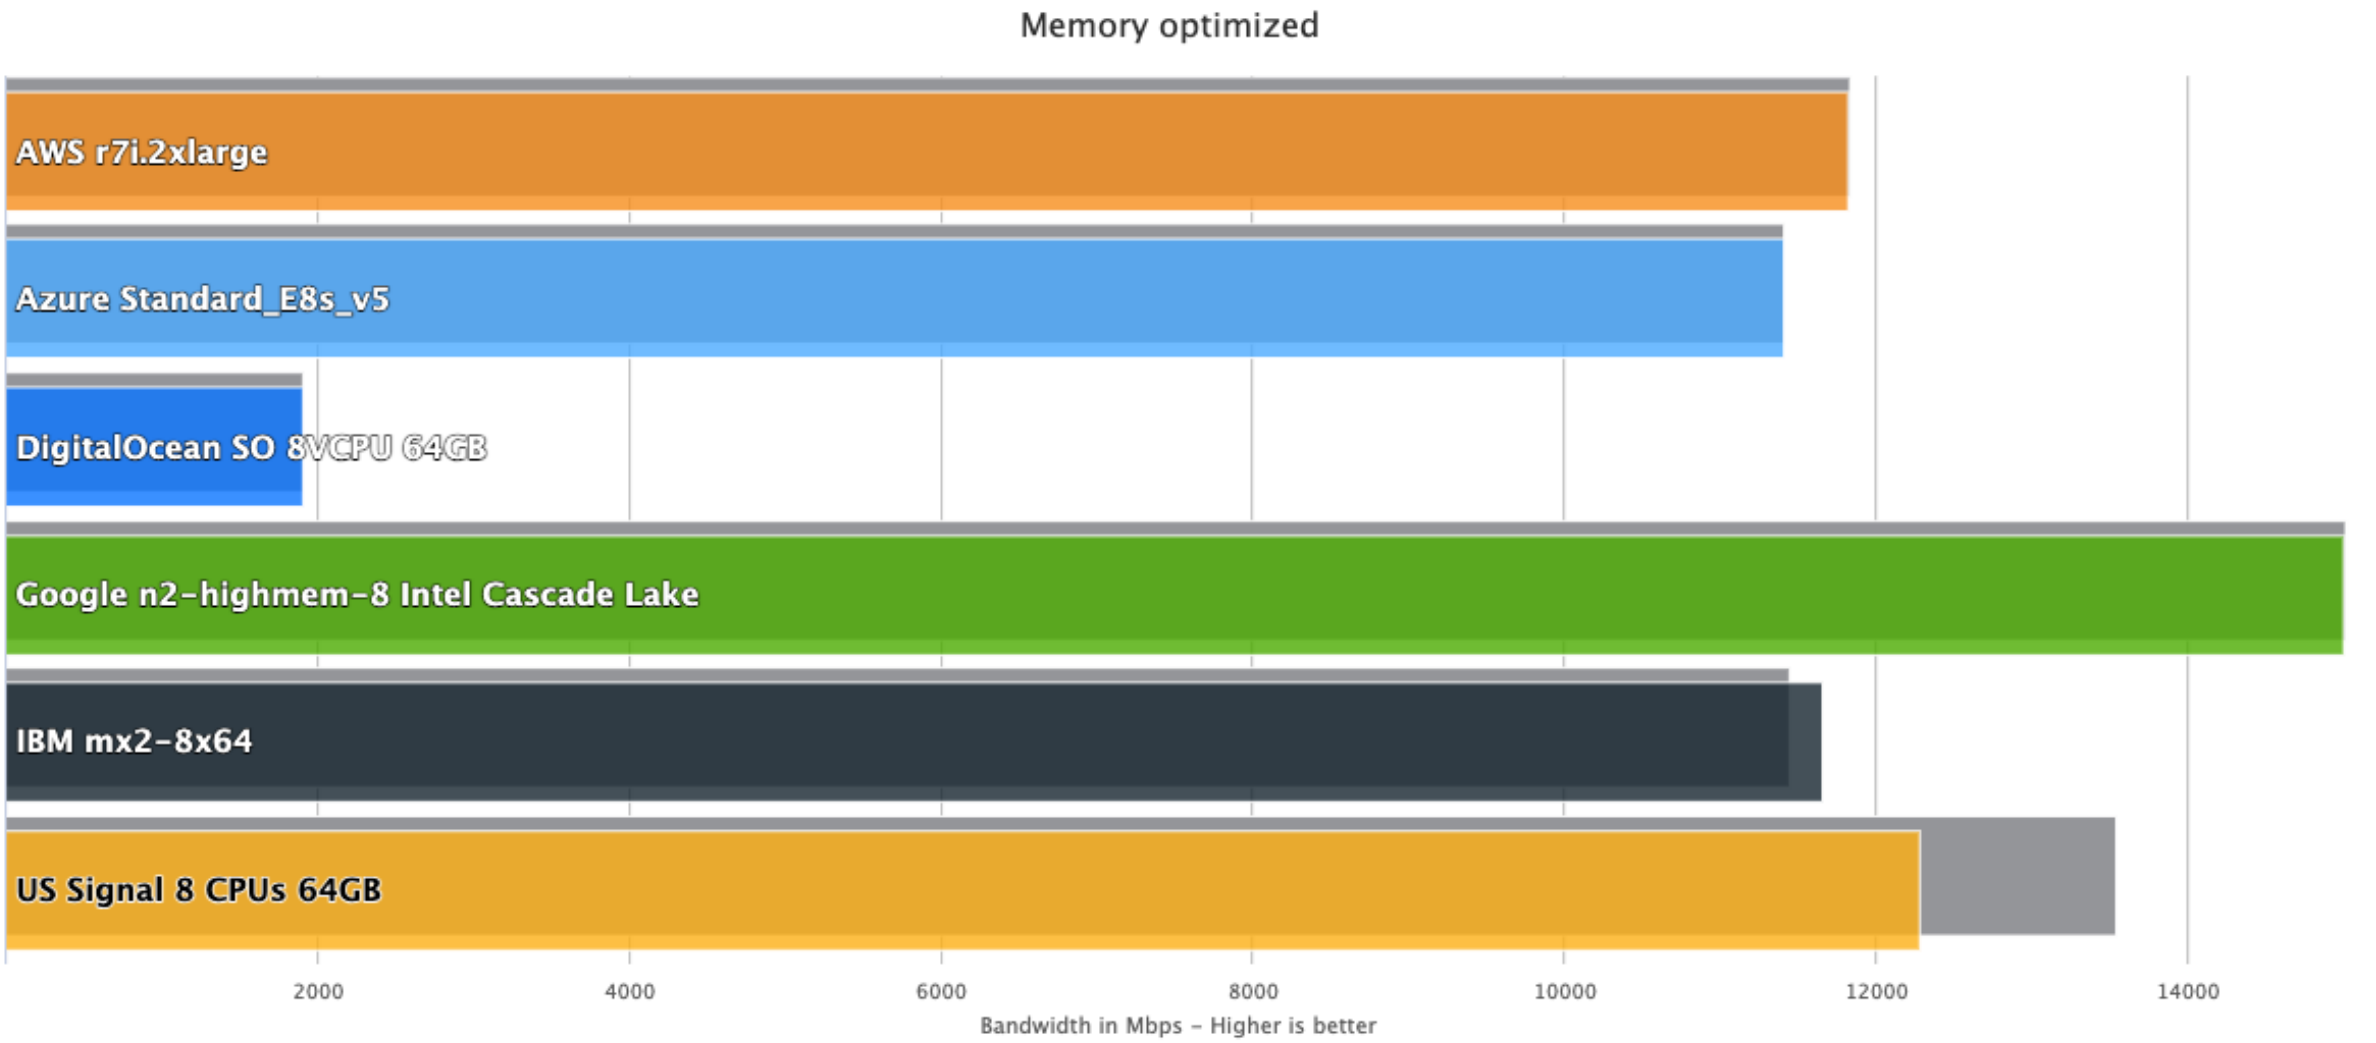

Storage - IOPS

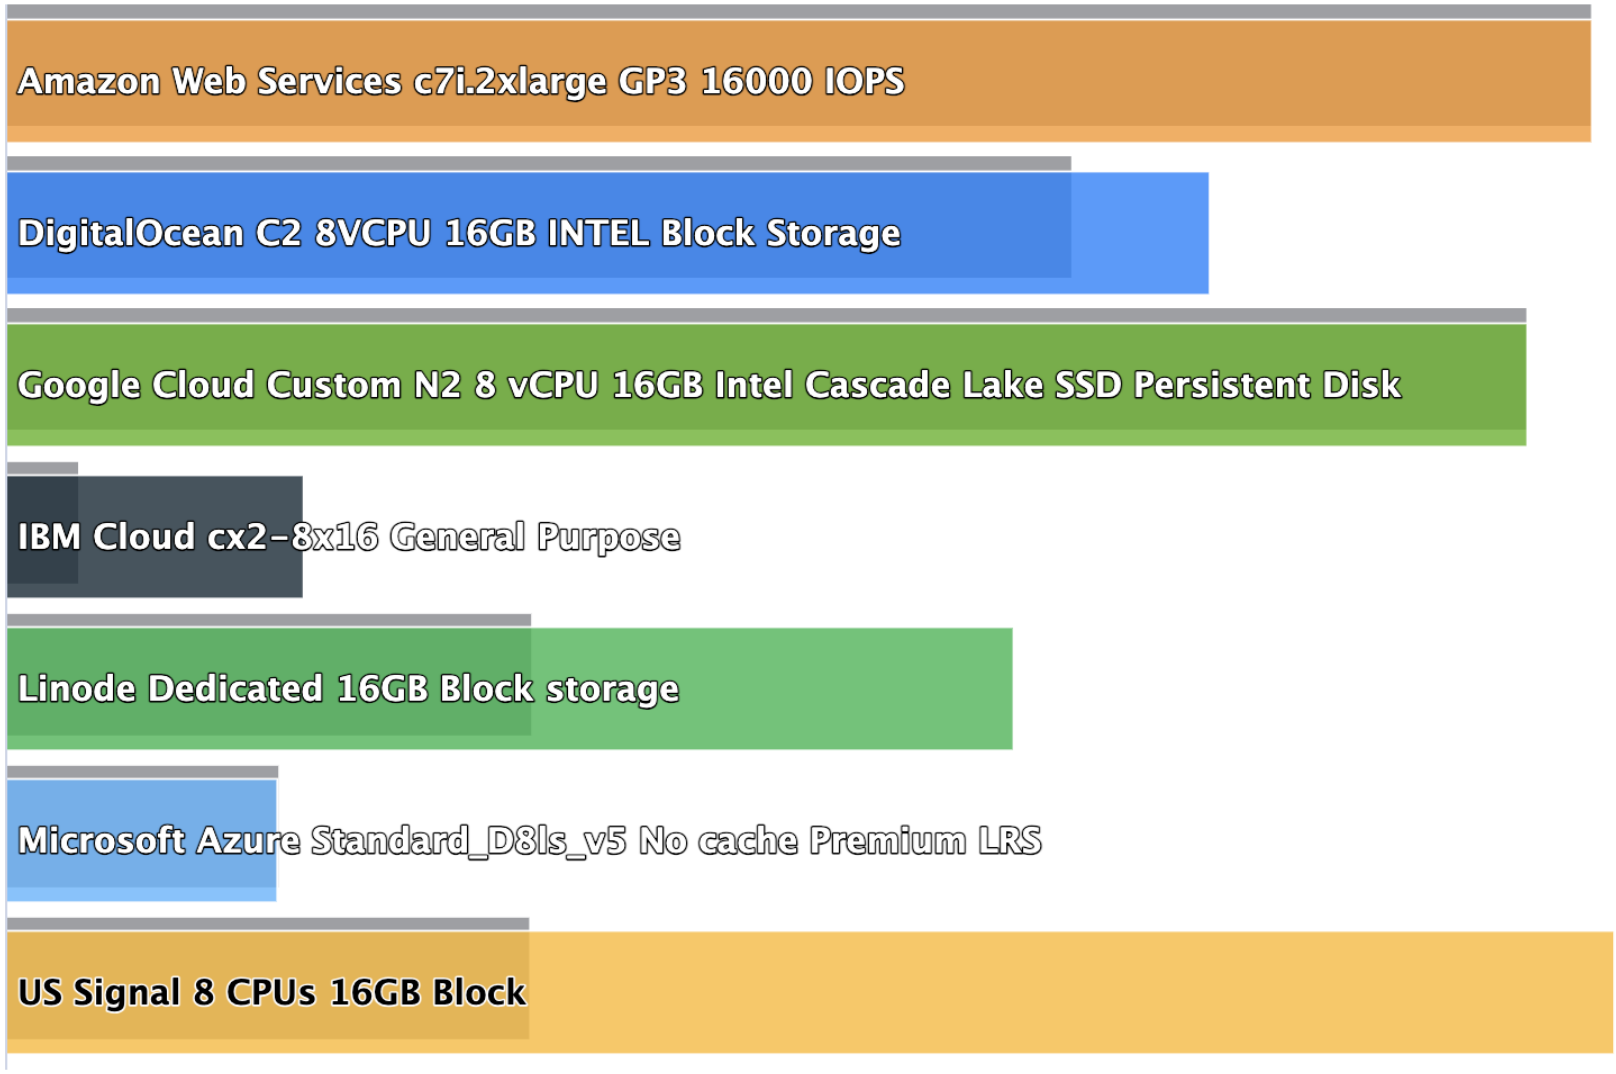

Storage – price/performance

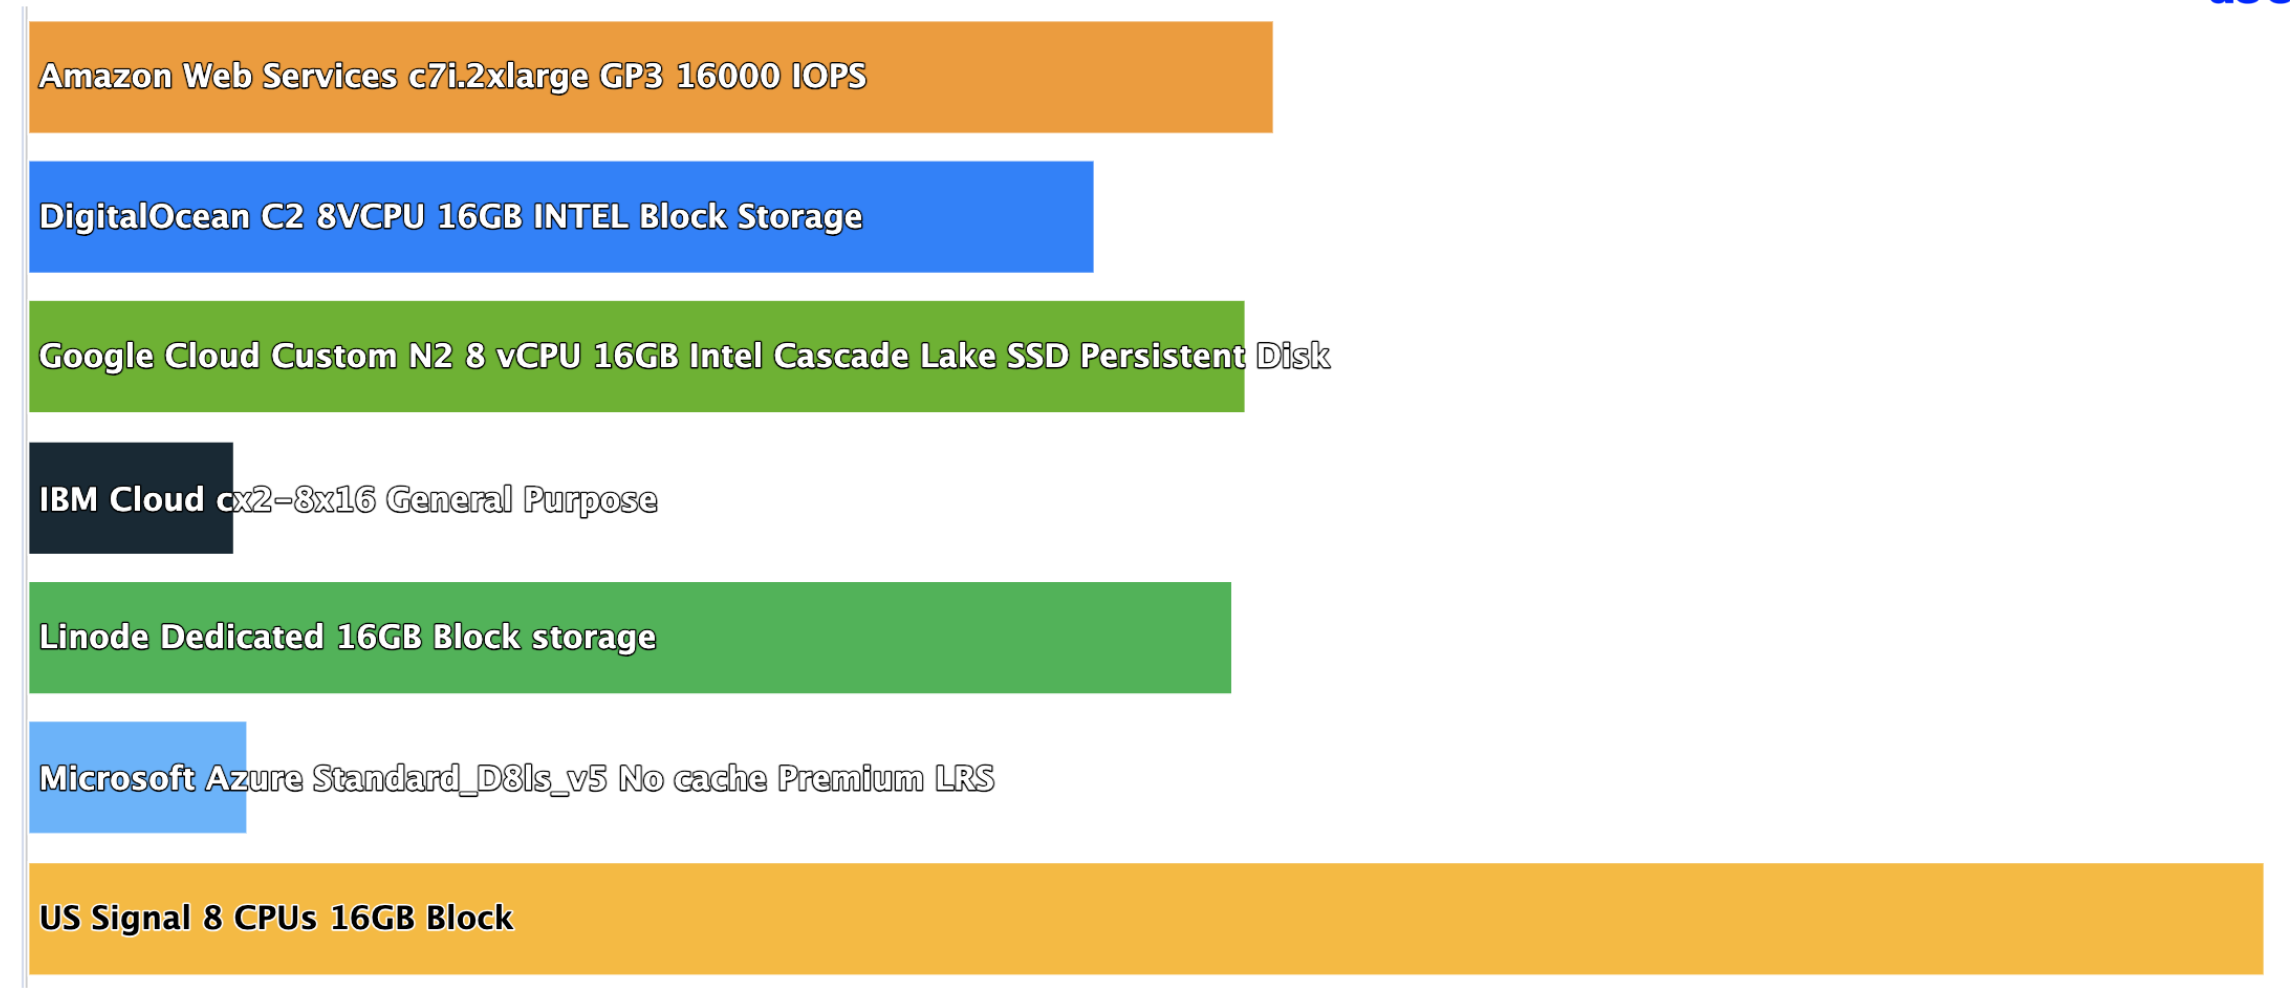

conclusion

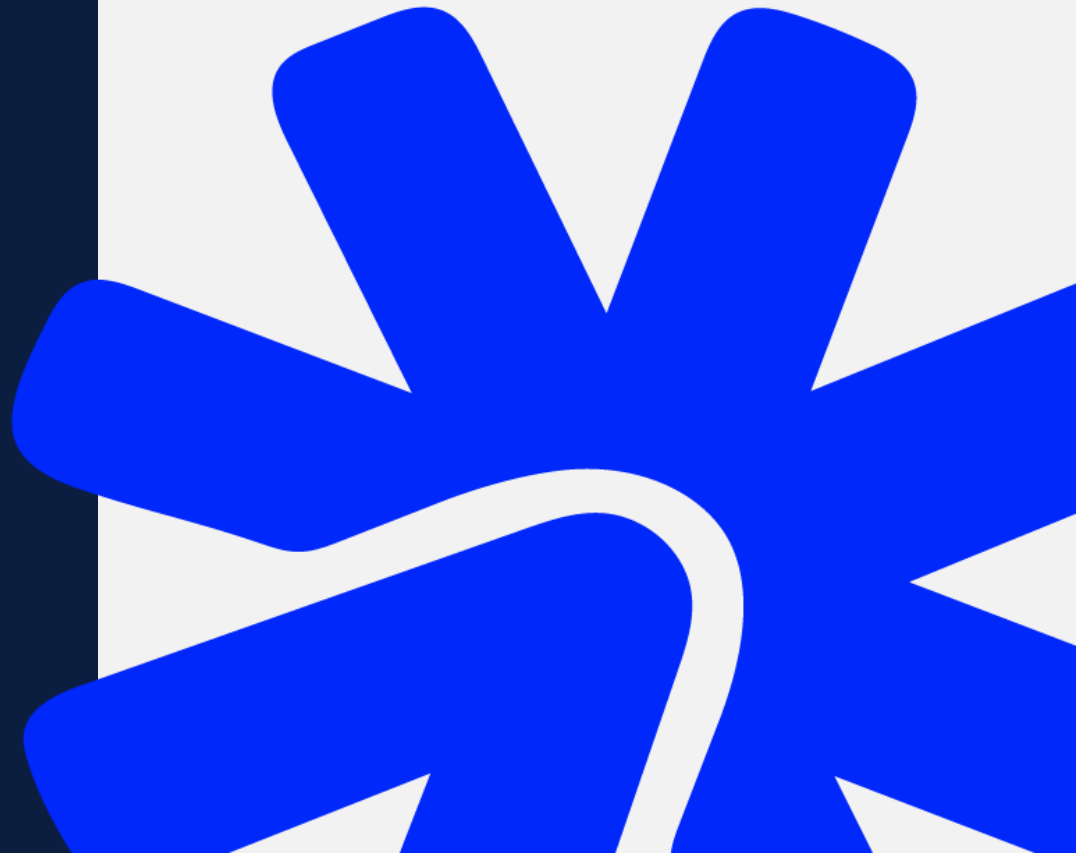

thank you!

Questions or Comments?

dwinans@ussignal.com
Linkedin.com/in/davenwinans