Virtualization and Containers

Comparing Linux Performance on VMs and OS Containers

Published on 01/2017

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Cloud Spectator is an analyst agency that specializes in cloud infrastructure performance analytics by measuring observable performance from a user perspective. In this study, Cloud Spectator worked with Virtuozzo to observe differences in virtual machine (VM) and container performance when deploying various hypervisors on identical environments using equivalently emulated production loads (more information on the physical host can be found in the Methodology). Cloud Spectator and Virtuozzo set up a total of three production environments, each with a different platform for virtualization: a commercial hypervisor, CentOS KVM, and Virtuozzo Containers™. Within each platform, Cloud Spectator deployed equivalent load generators (more information can be found in the Methodology) to saturate the physical host and emulate resource contention on a production environment. Cloud Spectator then used one VM/container to run benchmark tests to observe performance and performance variation of the loaded environment.

The testing was conducted over a 24-hour period, excluding setup time for the physical server, hypervisor installation, and load generation. The conclusions from the study highlight container performance advantages over virtual machines. Results showed better performance on the Virtuozzo container for both system performance (processor speed and memory bandwidth) and disk IOPS, as compared to its VM counterparts in the commercial hypervisor and CentOS KVM.

SYSTEM PERFORMANCE (PROCESSOR AND MEMORY)

Cloud Spectator evaluated processor and memory performance using the Geekbench 3 benchmark suite. Tests from the suite include tasks such as encoding and compression. The testing occurred over a 24-hour period.

SYSTEM PERFORMANCE KEY FINDINGS

- The Virtuozzo Containers™ demonstrated the highest system performance observed in the study.
- The Virtuozzo container scores higher in processor performance than the Commercial hypervisor VM by 1.5x.
- The CentOS KVM VM demonstrated the best performance stability over the test period, with a CV score of 1.9% variance. The Virtuozzo container showed similar stability with 2.4% variance.

DISK IOPS

Cloud Spectator evaluated disk performance using fio. The team conducted tests using sequential and random operations, looking separately at reads and writes. For details on file size, block size, queue depth, etc. please see the methodology. The testing occurred over a 24-hour period and disk testing ran after the system performance testing.

DISK PERFORMANCE KEY FINDINGS

- The disk IOPS on the commercial hypervisor performed consistently at around 20,000 IOPS.
- Virtuozzo containers showed high performance (average 50K IOPS) for all sequential and random operations except random reads, which produced an average 30K IOPS.
- KVM showed high performance (average 23K IOPS) for sequential operations, but very low performance (3K IOPS) for random operations.

METHODOLOGY

Cloud Spectator ran performance tests for 24 hours to measure performance and performance stability of the virtual processor, memory bandwidth, and disk IOPs on all three platforms. All three platforms were installed on equivalent physical environments. Only the Virtuozzo Containers™ environment was tuned to examine the maximum range of performance among the hypervisor-based VMs and the container; tuning the hypervisor environments or adjusting the Virtuozzo Containers™ environment may yield different results than shown in this report. The three tested platforms were:

1. A commercial hypervisor
2. CentOS KVM
3. Virtuozzo Containers™

All tests across platforms were conducted on the same physical hardware:

CPU	2x Intel Xeon E5-2620
RAM	64GB (4 x 16GB DDR4-2133 ECC REG)
RAID	LSI 9271-8i (8-port SAS2, 1GB)RAID0 over 8x HGST 450GB 15K RPM

TEST DIAGRAM

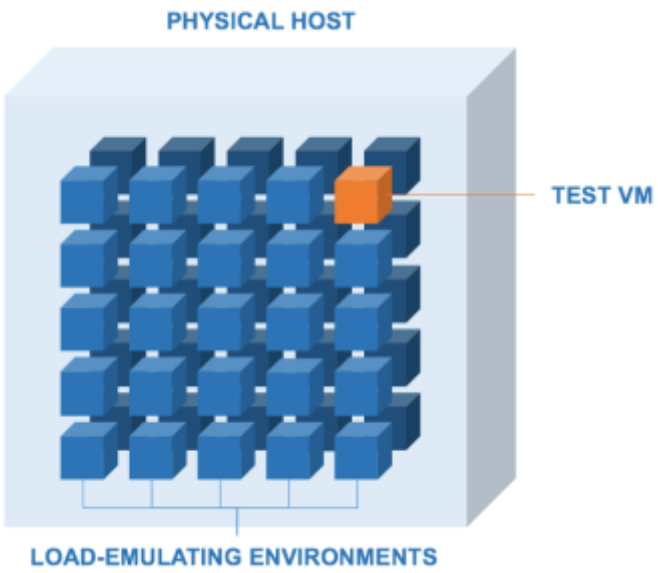

The Test Diagram illustrates a blueprint of the environment setup on each platform.

Within each physical host, 50 total load-emulating environments (VMs or containers) running httpress workloads were provisioned to saturate resources on the host. Cloud Spectator used a singular Test VM within each physical host environment to measure performance of the VM in a saturated production environment. This testing was performed continuously for 24 hours to collect information on performance and performance variation in response to resource contention.

LOAD GENERATION

The team used a modified version of httpress (original version can be found at https://bitbucket.org/yarosla/httpress/wiki/Home) to generate load on each server. The modified httpress can be found at https://github.com/virtuozzo/httpress.

Load generation command:

nohup ./httpress –p 32 –t 86400 –c 800 –R 120 $ADDRESS_MASK –r 4-53 &

- nohup: ignore a hangup signal to keep the process running upon logout

- ./httpress: the program running the load generation
- -p: the # of threads httpress will run
- -t: the amount of time in seconds the load will run (86400 seconds is 24 hours)
- -c: the number of concurrent connections in total (800 connections over 32 threads)
- -R: the number of requests per second (the rate is not fixed)
- -r: the IP address range of the VMs/containers on the server

TEST VM/CONTAINER

The two tested VMs and single container ran on 2 cores with 2GB RAM on CentOS 7.2.1511. Each test iterated between system tests and disk tests. The system tests measured processor speed through various tasks (e.g., encryption/decryption, compression/decompression) and memory bandwidth. The disk tests measured IOPS, latency and throughput of the file system.

SYSTEM TESTING: GEEKBENCH 3

Geekbench 3, a licensable product created by Primate Labs, is a cross-platform processor benchmark that can measure single-core and multi-core performance by simulating real-world workloads. The Geekbench 3 test suite is comprised of 27 individual tasks/workloads: 13 integer workloads, 10 floating point workloads, and 4 memory-bandwidth tasks. While processor and memory bandwidth are both performance factors that contribute to the final score provided by Geekbench 3, the test suite weighs processing performance much more heavily than memory bandwidth. Also, memory bandwidth is not necessarily affected by the amount of memory available for the VM, so VMs with larger amounts of memory may not exhibit larger bandwidth. For more information on Geekbench 3 and to see its individual workloads, please see http://www.primatelabs.com/geekbench/.

DISK TESTING: FIO

Fio is an open source I/O generator that spawns a number of threads and processes to conduct I/O requests. For this study, fio measured disk IOPS of the CentOS file system.

Flags:

- Block size: 4kb
- File size: 16mb
- Runtime: 60 seconds
- Direct I/O enabled
- Number of threads: 2 (one per core)
- I/O depth: 32
- Output format: JSON
- Group reporting enabled and time-based

UNDERSTANDING PERFORMANCE RESULTS

Data from the 24-hour period is aggregated into box-and-whisker plots, which illustrate the minimum, 5^{TH} percentile, median, 95^{TH} percentile, and maximum scores achieved in the test period. An example of the plot is displayed below:

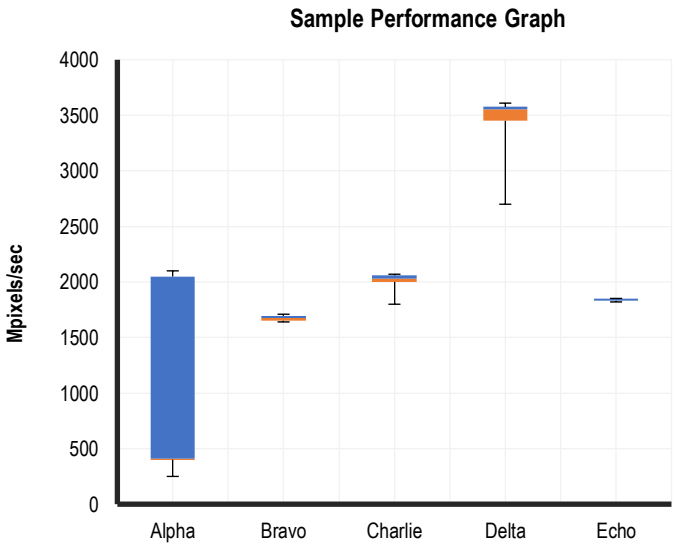

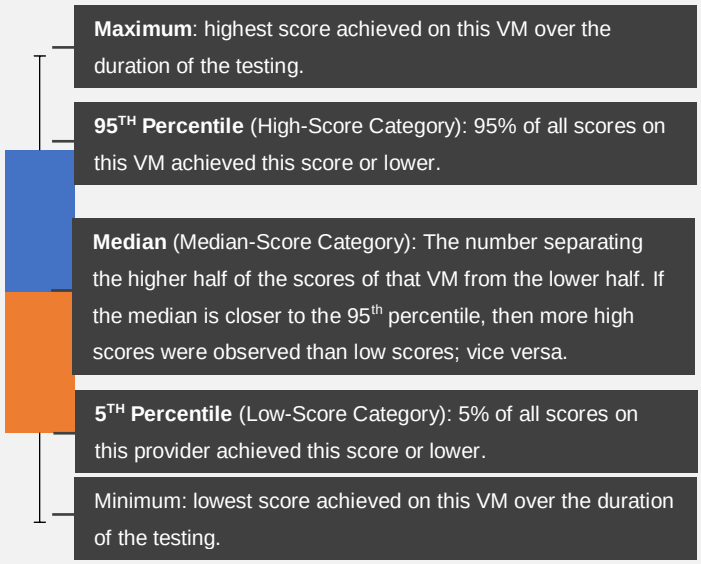

CONSIDERATIONS

- Stock configurations were used on each offering, and tuning will produce different results. No results in this report are reflective of tuning.
- While results in this report show performance advantages with Virtuozzo containers, containers and VMs should be considered with use case as well, not only performance. Use cases are not covered in this report.

PERFORMANCE RESULTS

SYSTEM PERFORMANCE

System performance is derived from 1) processor speed at completing integer and floating point tasks and 2) memory bandwidth. The tables on the right of the corresponding charts display the scores achieved by each VM or container. Processor speed is measured by running a variety of different individual tests to measure tasks such as encryption, compression, decompression, and image processing, and is split between integer and floating-point tasks to deliver a representative indicator of processor performance. Memory bandwidth is measured using STREAM, which can be installed as both an independent tool or as part of the Geekbench 3 suite. Memory bandwidth can directly affect processor performance as a potential bottleneck, and is especially important for in-memory databases.

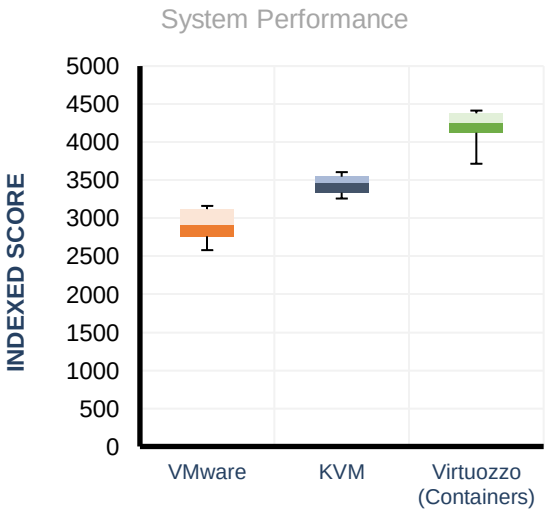

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	2581	2758	2911	3121	3162
KVM	3259	3336	3457	3546	3606
Virtuozzo	3716	4123	4252	4379	4415

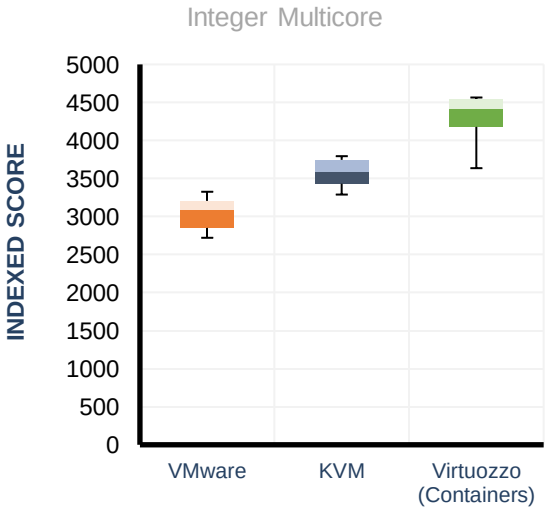

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	2721	2850	3089	3206	3326
KVM	3289	3431	3586	3746	3794
Virtuozzo	3636	4178	4417	4547	4565

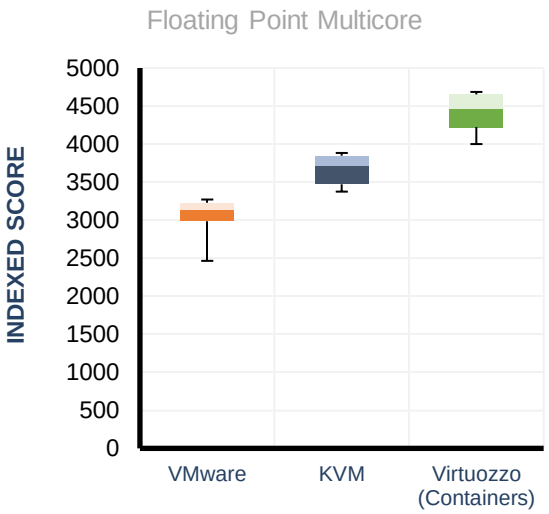

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	2465	2985	3140	3229	3271
KVM	3375	3471	3708	3847	3883
Virtuozzo	4000	4224	4461	4653	4685

Figure 1:
Mem Multi

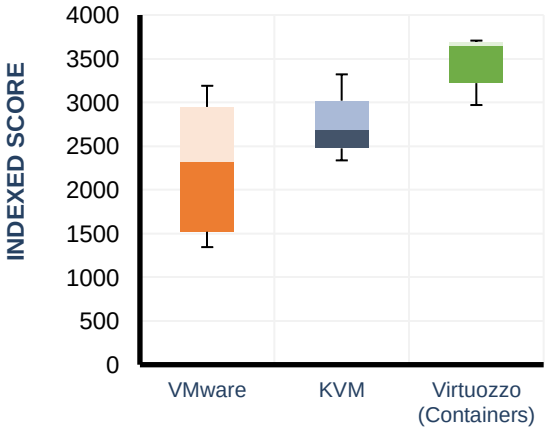

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	1345	1520	2324	2953	3190
KVM	2337	2474	2686	3020	3321
Virtuozzo	2970	3221	3653	3694	3707

DISK PERFORMANCE

Disk performance is not aggregated, and each type of operation is categorized independently. The study examined random and sequential operations for both 100% reads and 100% writes. The block size was set to 4KB and the direct I/O flag was active to bypass buffer/cache (more information on the disk testing can be found in the Methodology), to represent disk performance in relation to databases and file systems.

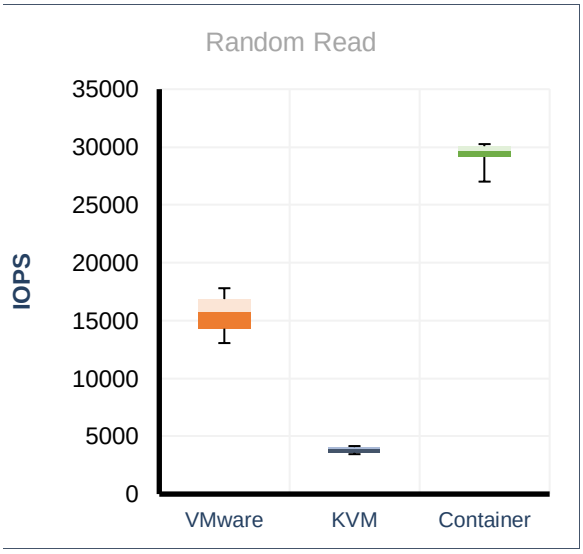

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	13053	14296	15766	16859	17795
KVM	3448	3565	3892	4064	4152
Virtuozzo	27009	29134	29711	30074	30252

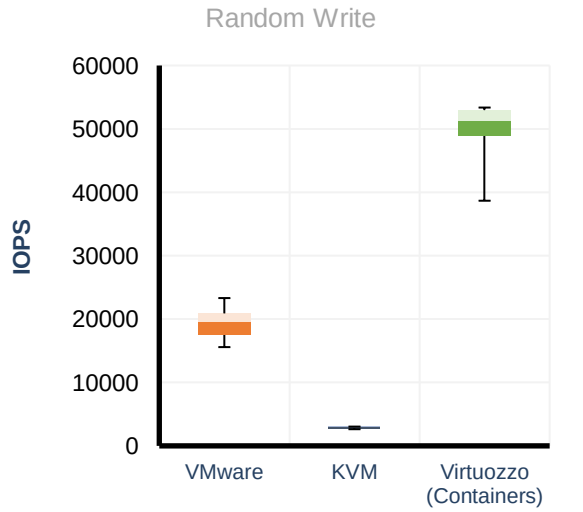

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	15581	17490	19505	20903	23317
KVM	2630	2702	2926	3010	3041
Virtuozzo	38669	48877	51341	52970	53366

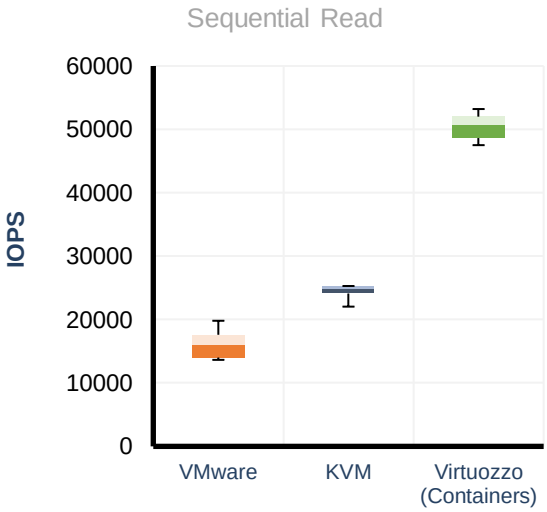

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	13645	14003	16094	17603	19814
KVM	22042	24111	24801	25185	25297
Virtuozzo	47514	48633	50672	52017	53212

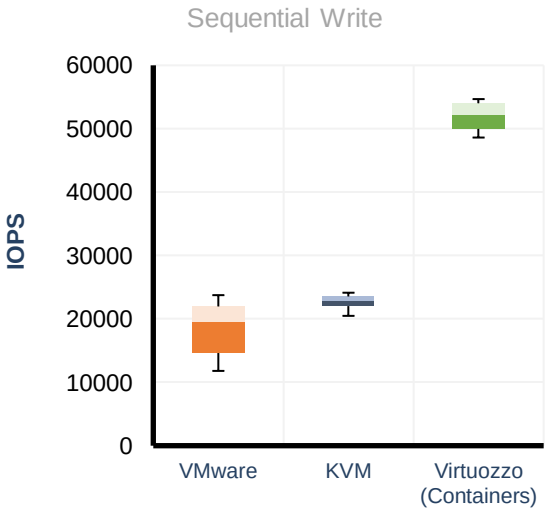

Platform	MIN	5TH	MED	95TH	MAX
Commercial Hypervisor	11798	14618	19570	21970	23750
KVM	20480	22023	22780	23553	24133
Virtuozzo	48608	49956	52236	54021	54672

CONCLUSION

All test machines were set up with the same amount and type of hardware, and configured with identical load generator environments. Despite the equivalence, though, each environment demonstrated performance and variance differences depending on the virtualization environment the test machine was provisioned on. The observed data shows that in the environments tested, the performance difference between containers and hypervisor-based solutions can be significant – in some cases as high at 50%. Therefore, the choice of virtualization technology and performance optimization can have an impact on performance for the end user.

ABOUT

Cloud Spectator is a cloud analyst agency focused on cloud Infrastructure-as-a-Service (IaaS) performance. The company actively monitors several of the largest IaaS providers in the world, comparing performance and pricing to achieve transparency in the cloud market. The company helps cloud providers understand their market position and helps business make intelligent decisions in selecting cloud providers and lowering total cost of ownership. The firm was founded in 2011 and is located in Boston, MA.

For questions about this report, to request a custom report, or if you have general inquiries about our products and services, please contact Cloud Spectator at +1 (617) 300-0711 or contact@cloudspectator.com.

For press/media related inquiries, please contact:
Ken Balazs
VP Sales & Marketing
kbalazs@cloudspectator.com