

Kubernetes for Stateful Workloads Benchmarks

Baremetal Like Performance for
For Big Data, Databases And AI/ML



Executive Summary

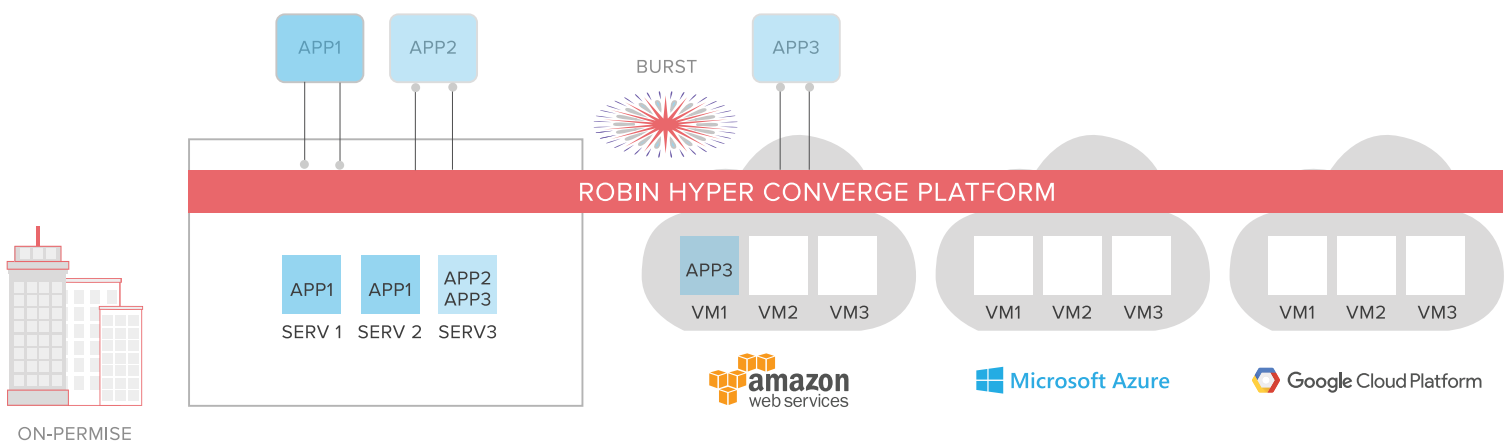
Customers are actively evaluating stateful workloads for containerization and management using Kubernetes but have not been able to find adequate solutions to get baremetal performance for data-heavy applications. The results in a recent CIO survey highlighted the confidence in Kubernetes as a viable platform for stateful applications.

Robin completed the below performance, scale and key workload-specific benchmark for containerized Stateful workloads that are managed using Kubernetes. Performance and Scale are the key parameters for CIOs in making a choice for modernizing their environment data infrastructure.

Storage Performance

The first parameter that the customers consider when containerizing data-intensive applications is the storage performance of an application. This benchmark showcases that Robin delivers near baremetal like performance when compared to bare-metal. Using Robin entails gaining multiple beneficial features such as consolidation of multiple apps on fewer resources or within the same server, cost savings with better hardware utilization, and an addition of a host of features such as snapshots, cloning, backup, scale and upgrade with 1-click.

Packaging more applications within the same hardware resources without performance degradation is critical to large enterprises that use big data, databases, and AI/ML apps across physical and cloud environments. Robin enables auto-monitoring resource allocation across on-premise and public cloud depending on the needs of the applications. Robin recalibrates the resource allocation as and when needed.



Environment

Common environment details used shown below.

Tool Used

Vdbench

Command Used

```
[john@server01]# vdbench -f config io=seq|random bs=16k|32k|64k|128k rw=0|100 rate=max etime=60000 thr=8 lun1=/dev/sdx
```

Workloads

Sequential and random workloads

SSD Benchmark

Below are the test environment details within which the tests were performed.

Server Details:

- » Software : Robin Hyper Converged Platform
- » SSD Device : ata-SAMSUNG_MZ7KM960HAHP-00005
- » Memory : 16GB
- » CPU : 4 cores

SSD Test Results

I/O Type	Sequential Write				Sequential Read				Random Write				Random Read			
	Block Size	16k	32k	64k	128k	16k	32k	64k	128k	16k	32k	64k	128k	16k	32k	64k
Robin	434.0	463.4	478.9	491.7	487.7	507.7	519.5	524.7	425.3	455.8	477.5	488.9	477.7	494.1	500.4	501.9

HDD Benchmark

Below are the test environment details within which the tests were performed.

Server Details:

- » Software : Robin Hyper Converged Platform
- » HDD Device : ATA- HGST_HUS724020AL
- » Memory : 16GB
- » CPU : 4 cores

HDD Test Results

I/O Type	Sequential Write				Sequential Read				Random Write				Random Read			
	Block Size	16k	32k	64k	128k	16k	32k	64k	128k	16k	32k	64k	128k	16k	32k	64k
Robin	109.7	115.8	121.0	121.2	153.7	153.7	152.7	153.8	2.4	4.6	9.2	15.7	4.4	8.5	16.0	28.4

Application Workload Performance

The earlier section shows that Robin offers near baremetal device performance. Now, at the Application level, we chose SQL and NoSQL applications to showcase that Robin is inline with bare-metal application performance.

Stateful application or workloads have unique infrastructure requirements and Robin brings to the table a host of features that help manage the enterprise environment in an efficient way with near bare-metal performance.

Oracle SLOB

Tests were performed with different replication factors with Robin Hyper Converged Kubernetes platform. The benchmark results show that even with application protection enabled using Robin replication factor, the performance is near bare-metal.

Test Environment

Server Details:

- » Number of nodes : 3
- » Software : Robin Hyper Converged Platform
- » Memory : 128GB/Node
- » CPU : 20 cores/Node

Oracle SLOB Test Results

- » All data is from AWR reports
- » DB Size: 50GB
- » Test duration: 2 hrs

SSD	Baremetal		Robin Replication-1		Robin Replication-2		Robin Replication-3	
	IOPS	MB/s	IOPS	MB/s	IOPS	MB/s	IOPS	MB/s
100% Write	6436.90	75.2	5872.10	69.6	5156.10	67.5	4982.9	65.2
100% Read	47158.40	368.6	46401.20	362.6	46419.0	363.1	46446.50	363
OLTP, 70% Reads, 30% Writes	9521.7	96.3	10640.8	107.6	9902.4	102.2	9450.1	99.2

HDD	Baremetal		Robin Replication-1		Robin Replication-2		Robin Replication-3	
	IOPS	MB/s	IOPS	MB/s	IOPS	MB/s	IOPS	MB/s
100% Write	26.1	4.1	510.1	7.1	497.6	6.9	499.6	7
100% Read	290.3	2.3	310.1	2.5	309.8	2.5	306.4	2.5
OLTP, 70% Reads, 30% Writes	269.1	3.3	353.8	4.1	353.8	4.0	343.8	3.9

Cassandra YCSB Results :

Tests were performed processing 1 Billion records with different workloads, on a 3 node setup using Robin Hyper Converged Kubernetes platform. The benchmark results show the performance is near baremetal even with such a huge amount of data.

Test Environment:

Server Details:

- » Number of nodes : 3
- » Software : Robin Hyper Converged Platform
- » Memory : 128GB/Node
- » CPU : 20 cores/Node

Cassandra YCSB run Test Results

YCSB Benchmark for Cassandra (1 Billion Records)				OPS	
Workload	Type	#Records	Cassandra Replication	Baremetal	Robin
Workload A	LOAD	100000000(1B)	3	78150	77762
Workload A	RUN	100000000(1B)	3	77727	75954
Workload B	RUN	100000000(1B)	3	62074	64743
Workload C	RUN	100000000(1B)	3	89114	86217
Workload D	RUN	100000000(1B)	3	88049	85264
Workload E	RUN	100000000(1B)	3	79612	77104

Application Scaling

Scaling to meet the demand is a fundamental requirement for cloud native applications and this need is a complex one to meet in the case of distributed or stateful applications. ROBIN is architected to scale horizontally with ease.

Bringing up data-heavy application clusters is a huge effort and takes a lot of time - hardware, software installation, cluster configuration, and performance tuning - A repetitive cycle for each cloudera installation - As opposed to this Robin was able to bring up a 500 node cluster in less than one hour - bringing the time-line down from many weeks to 54 minutes.

Application Scaling - Cloudera

To showcase the high level of scalability with Robin, we brought up a Cloudera cluster with 500 nodes.

- » # of application : 1
- » # of cloudera nodes : 500
- » 491 DataNodes
- » 1 Gateway Node
- » 2 Name Nodes
- » 1 CM Server
- » 3 zookeepers
- » 1 node for utilities
- » 1 node for hive
- » Total cores : 6400 configured
- » Memory : 47.65 TiB
- » Storage configured 1.45 PiB
- » 100 physical servers used
- » App creation time : 54 mins
- » Cloudera Version: CDH-5.12.0