Top 10 2020 – Cloud Compute

North America

Cloud Mercato

I. Introduction

At the start of this year 2020, cloud computing is a large and mature industry, always growing and creating new models of computers' resources consumption. Despite a large adoption by enterprise and governments, from a consumer prospective, the cloud market still tends to appear to be hard to decrypt.

To answer to this problematic, Cloud Mercato designed this report gathering the 10 most interesting cloud service providers in North America. This document aims to brings a neutral and objective evaluation of CSPs' capabilities. It provides a state-of-art of cloud industry made by a third-part analysis team.

II. Executive summary

Across providers, there's a multitude of offering and ways to evaluate cloud vendors, in this document we focus on classical cloud computing services. Virtual machines, containers, VPS, whatever the name is, the principle is selling on-demand an amount of CPU and RAM with a network availability directly on the Internet. This kind of service is one of the most basic IaaS offers provided by CSP and generally called "Compute".
Due to its essential nature and being also one main component of VM offering, block storage is part of evaluation in our study.

The subject of this document is efficiency measurement of compute services by the following questions:

- How do they perform?
- How much do they cost?
- Which is the most valuable?

III. Methodology

According to the nature of analyzed components, definition of performance varies. CPU and storage express their efficacity very differently, so tests and analysis are way different from one component to the other. Cloud Mercato regularly launches tests on many cloud products and for this report we picked few of our existing usual methodologies.

1. Provider criteria

Among vendors tested by Cloud Mercato during year 2019, 10 providers have been selected for this ranking. More candidates were available and we picked only the best performers filling the following requirements:

- Presence with several datacenters in North America
- Offer VMs with at least 2 up to 16vCPU
- Offer CPU/RAM ratio at least 1:2 up to 1:4
- Offer block storage with volume at least 100 up to 500GB powered by SSD
- Offer hourly billing option without engagement

2. Setup

The ranking is based on 4 categories of VMs, from 2 to 16 vCPU. The table below describes the specifications that we attribute to each category. We tried to match with a CPU/RAM ratio of 1:2, but RAM may vary across providers.

SIZE	vCPU	RAM (GB)	STORAGE (GB)
Small	2	4	100
Medium	4	8	150
Large	8	16	200
Extra large	16	32	500

If a provider allows detachable volume, virtual machines are equipped with SSD block storage as an extra volume else root volume is used. All instances had been tested with Ubuntu 18.04. Appendix brings an accurate definition of virtual hardware used for each provider.

For each VM type, at least three instances were provisioned simultaneously. For confirmation or validation, more copies could have been launched.

3. Test software

TEST	SOFTWARE
Compute	Geekbench 5
Storage IOPS	FIO
Storage bandwidth	FIO

a. Compute

CPU performance was collected using Geekbench 5. This well-known suite runs a large number of tests covering a large spectrum of computing domains. Topics tested are:

- Integer
- Floating point
- Cryptography

Around twenty metrics are reported by Geekbench plus a score linked to each of them, for single and multi-thread workloads. In the context of this report we use:

- **Single score**: As CPU power evaluation
- **Multi score**: As multi-task power evaluation

b. Storage

Storage performance are measured by two metrics: IOPS and bandwidth. For both, Flexible I/O tester has been used with two different scenarios. In both we used the following parameters:

- A number of thread equal to the number of vCPU
- A direct connection to the device without file system
- Flags by-passing buffers and cache such as IO_DIRECT
- Read and write access (not mixed)

i. IOPS

This scenario aims to reveal what is the best rate in terms of block handling. So, we applied a small block size of 4K allowing the maximum number of blocks handled simultaneously. A random access to drive is applied, removing predictability of a sequential operation.

ii. Bandwidth

IOPS captures the transaction rate, but it doesn't reflect the maximum throughput available by volume. The nature of the scenario with small blocks and random access already represent a bottleneck for this metrics. To collect the maximum bandwidth, we access sequentially to drive with big blocks of 1M reducing block processing and allocation time.

4. Calculations

The performance reported in this document are averages, so to gauge the variability we also bring standard deviation but we don't use it in calculations. For information, iin a normal distribution, the range between AVG-STDDEV and AVG+STDDEV represents 68.2% of the population.

To aggregate different kind of values into understandable comprehensive scores, Cloud Mercato used few formulae:

Name	Type	Calculation
Compute score	Compute	Multi thread performance / Monthly price
Storage score	Storage	AVERAGE(IOPS, Bandwidth) / Monthly price
Overall score	Overall	AVERAGE(Compute score, Storage score)

IV. Overview

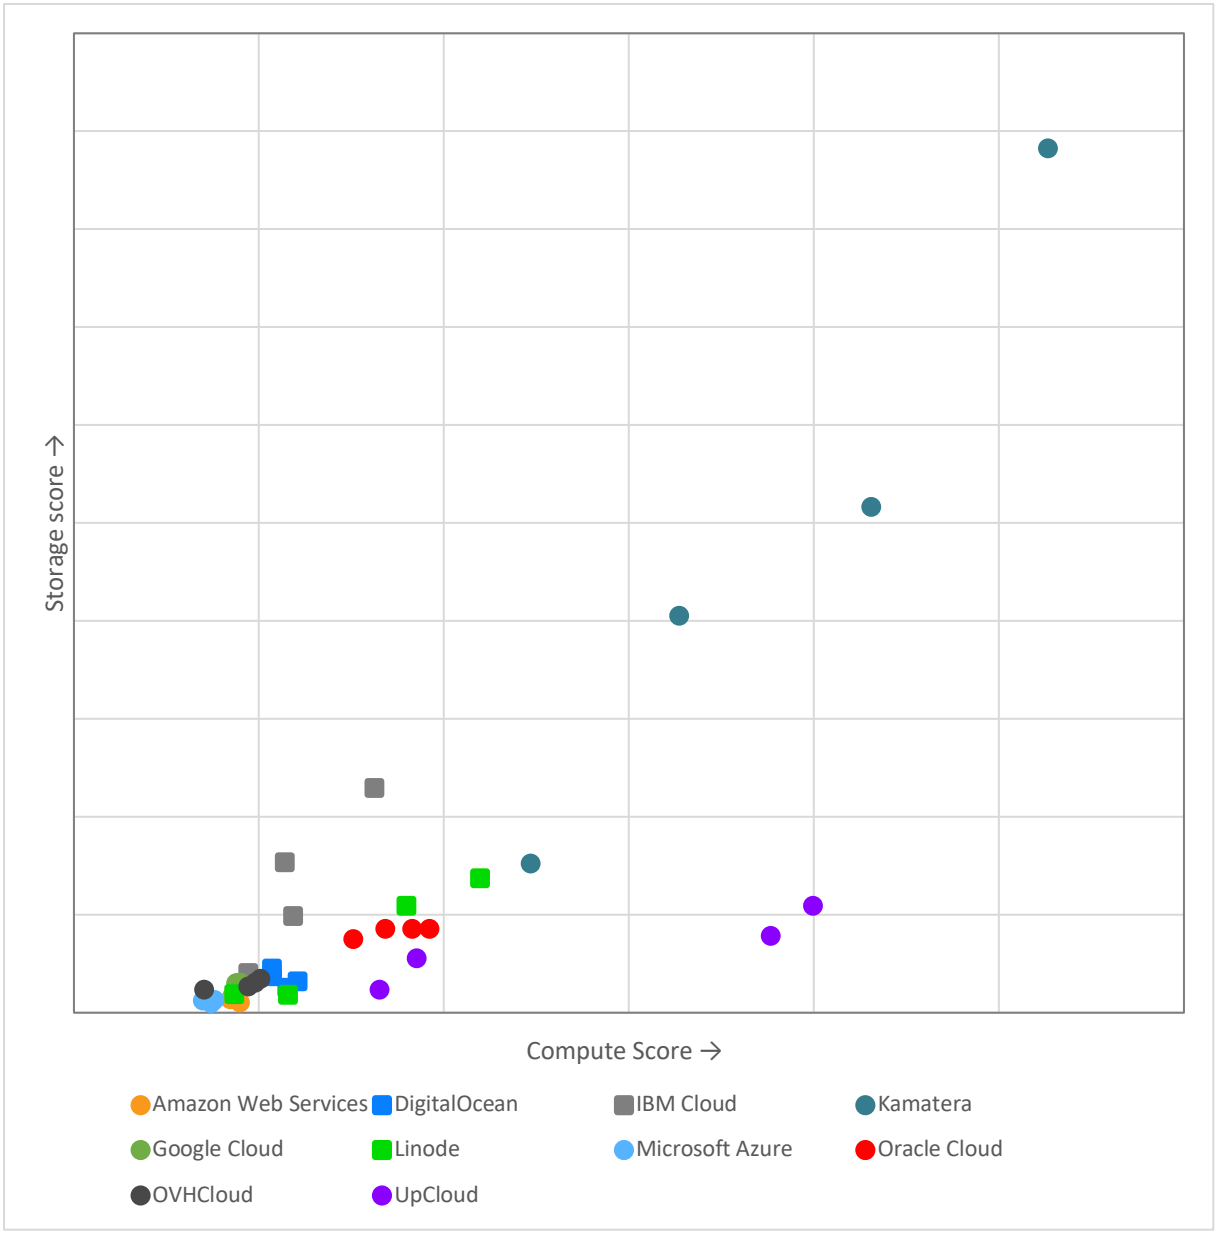
Storage score →
Compute Score →
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

V. Performance

For each category of products analyzed in this document, vendors provide almost the same specification. Catalogs all claim good vCPU, RAM and SSD storage but behind announced virtual hardware, they aren't equal in term of performance. Efficacity is not a value given by catalog or at least in a too subjective manner.

To clarify this question, Cloud Mercato launched a benchmark suite on virtual machines and block storage. This section outlines the average performance discovered during our testing.

1. Compute

a. Overall

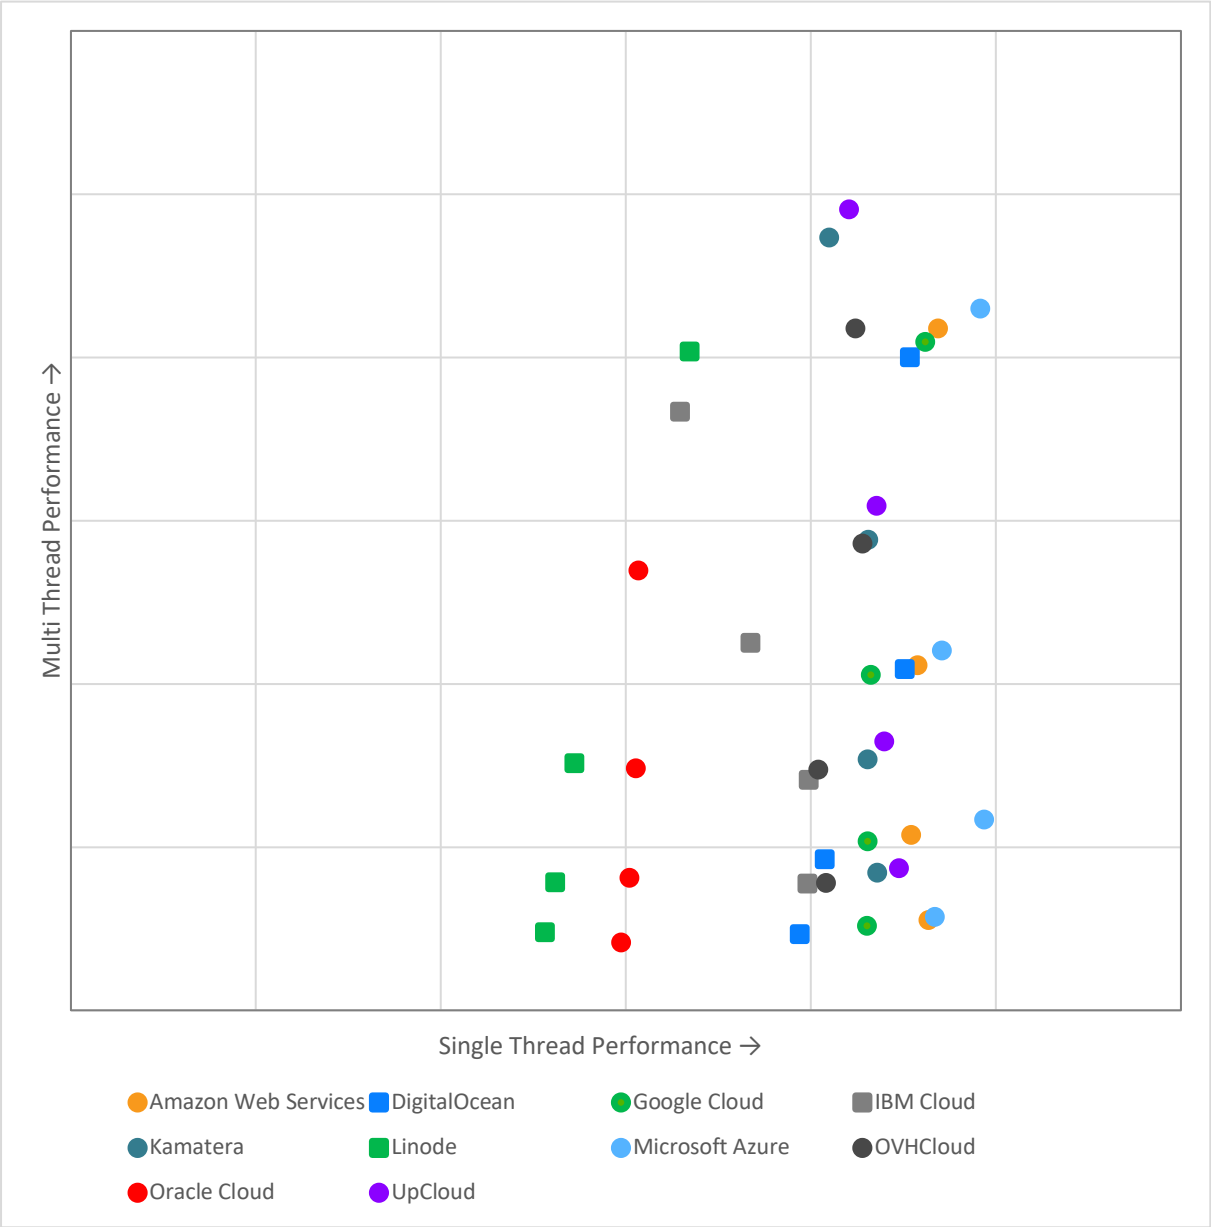

b. By category

i. Small

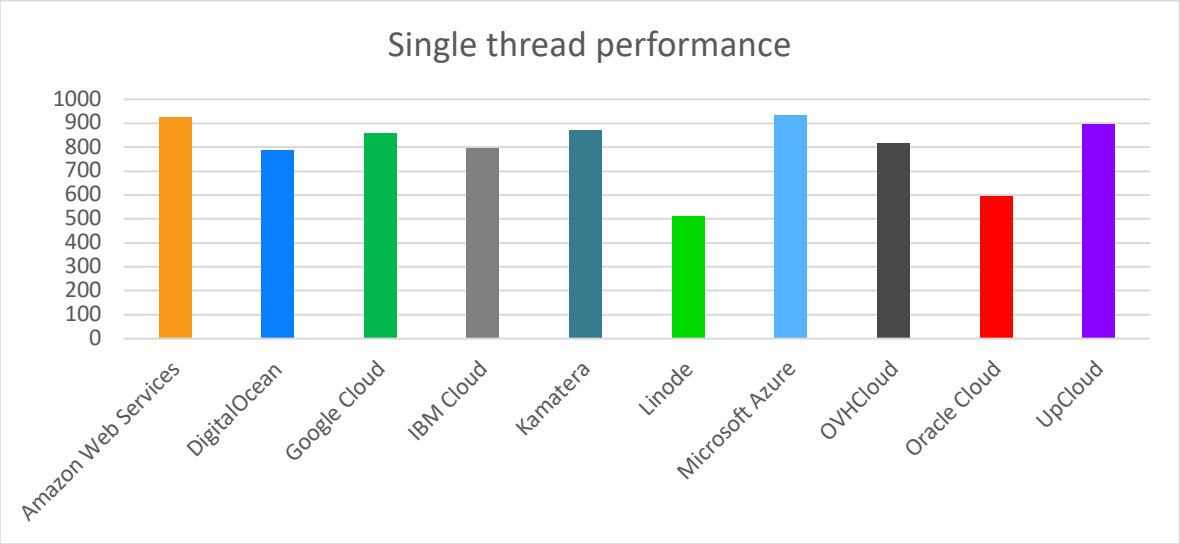

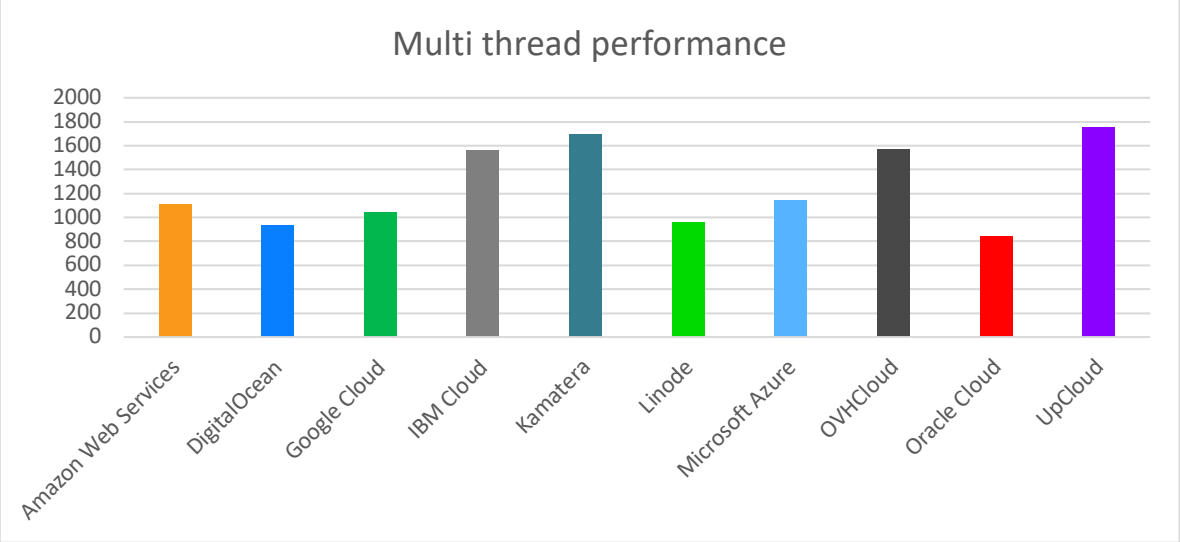

	Single thread		Multi thread	
	Mean	Deviation	Mean	Deviation
Amazon Web Services	926	14,63	1110	11,94
DigitalOcean	787	24,30	937	19,57
Google Cloud	860	12,15	1044	10,93
IBM Cloud	796	121,68	1561	235,10
Kamatera	871	32,77	1690	76,53
Linode	512	89,59	961	184,46
Microsoft Azure	933	37,39	1146	37,47
OVHCloud	816	37,92	1569	45,10
Oracle Cloud	594	7,21	840	11,09
UpCloud	895	30,85	1749	69,31

ii. Medium

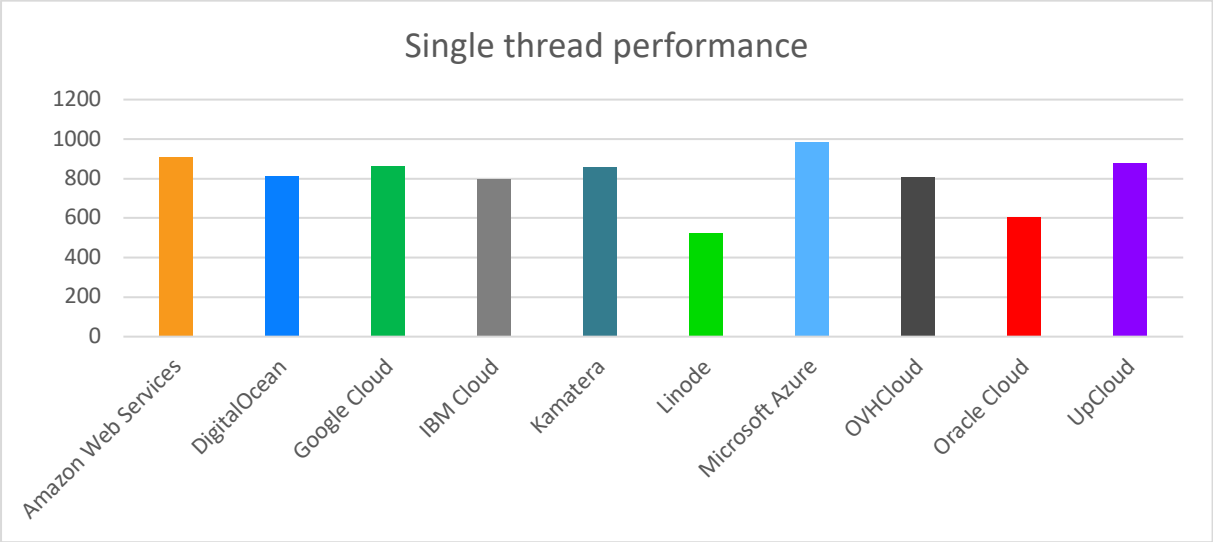

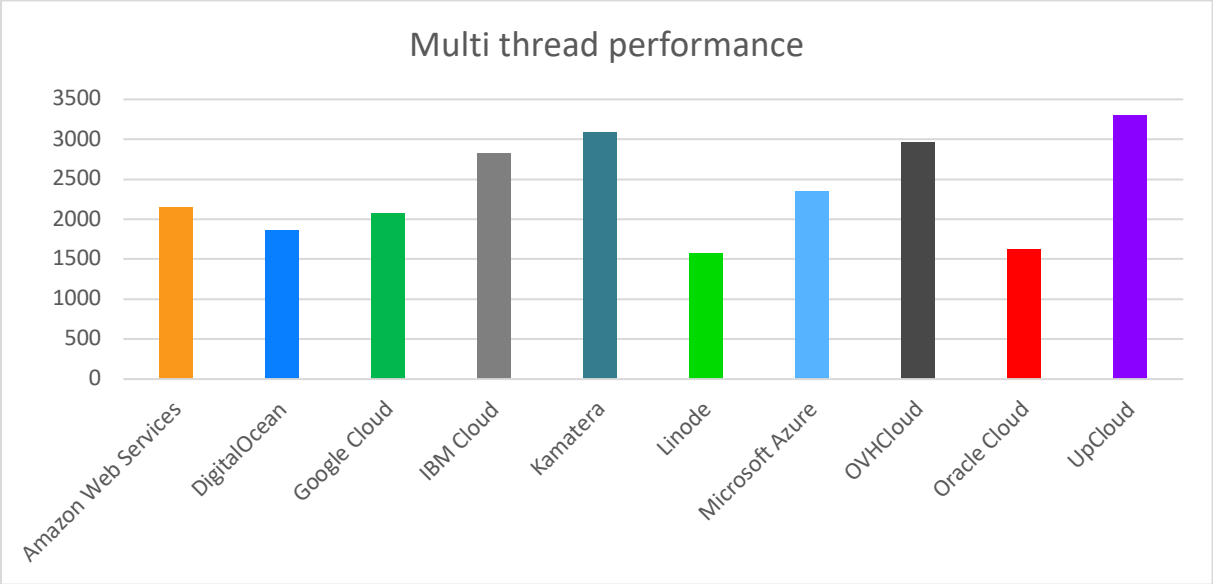

	Single thread		Multi thread	
	Mean	Deviation	Mean	Deviation
Amazon Web Services	908	36,77	2152	75,92
DigitalOcean	814	38,77	1855	45,81
Google Cloud	861	12,84	2078	22,44
IBM Cloud	797	41,43	2825	383,91
Kamatera	860	39,91	3083	278,76
Linode	523	50,33	1573	416,60
Microsoft Azure	987	39,00	2345	57,40
OVHCloud	807	15,38	2958	89,45
Oracle Cloud	603	7,51	1625	20,70
UpCloud	879	16,08	3299	68,39

iii. Large

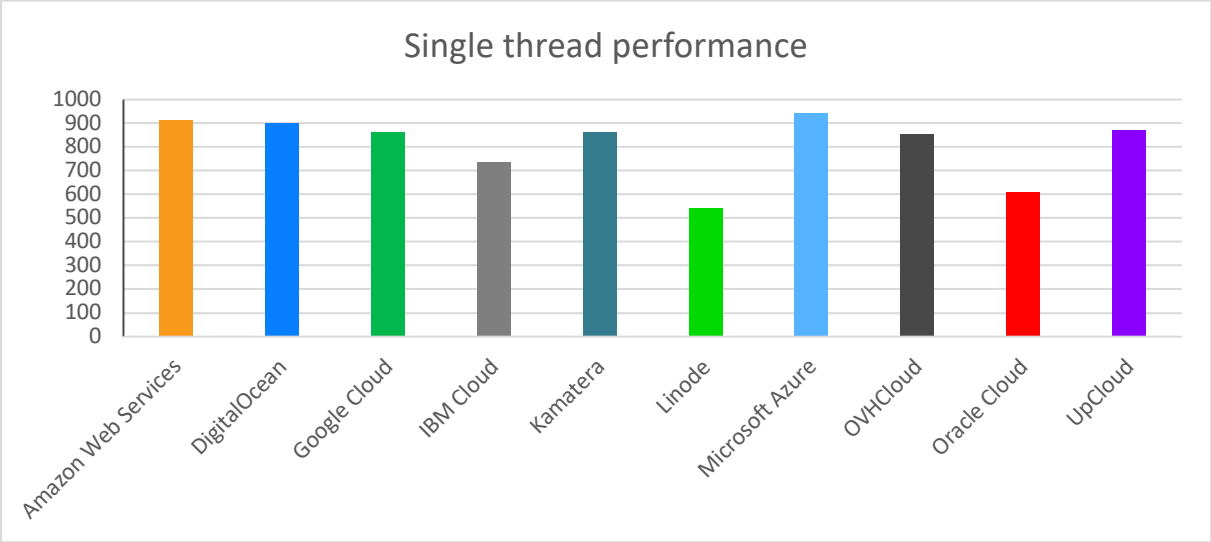

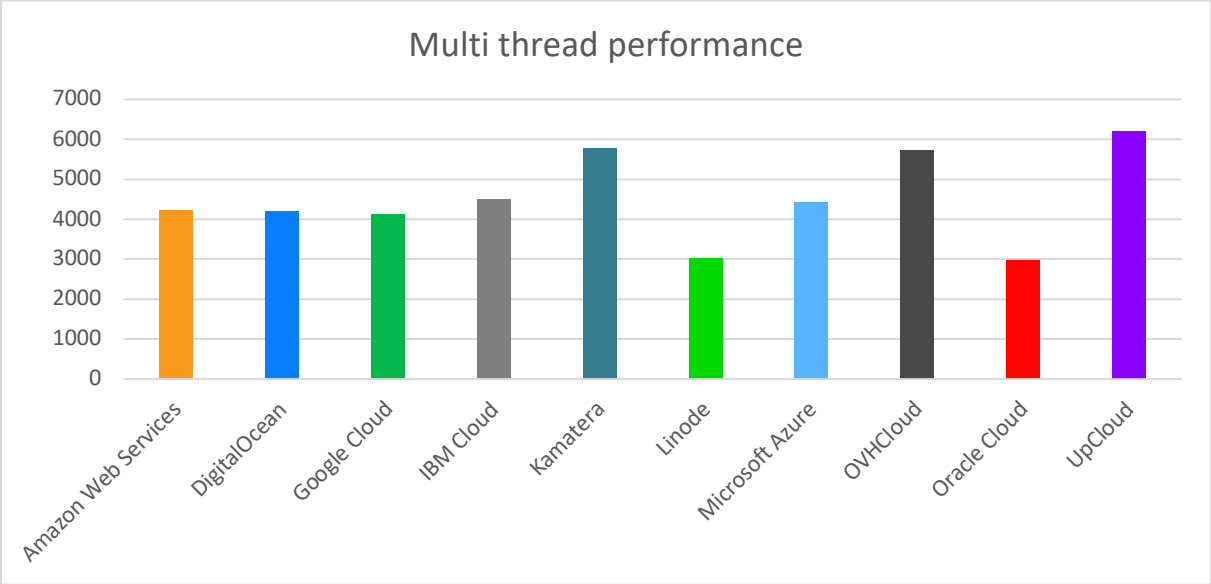

	Single Mean	Deviation	Multi Mean	Deviation
Amazon Web Services	915	18,63	4231	60,23
DigitalOcean	901	8,27	4188	29,42
Google Cloud	864	3,81	4114	12,90
IBM Cloud	734	15,79	4507	177,54
Kamatera	862	38,62	5774	651,62
Linode	543	22,58	3031	465,41
Microsoft Azure	941	27,48	4414	115,94
OVHCloud	856	31,43	5727	161,52
Oracle Cloud	610	4,32	2967	161,56
UpCloud	870	30,95	6187	385,91

iv. Extra Large

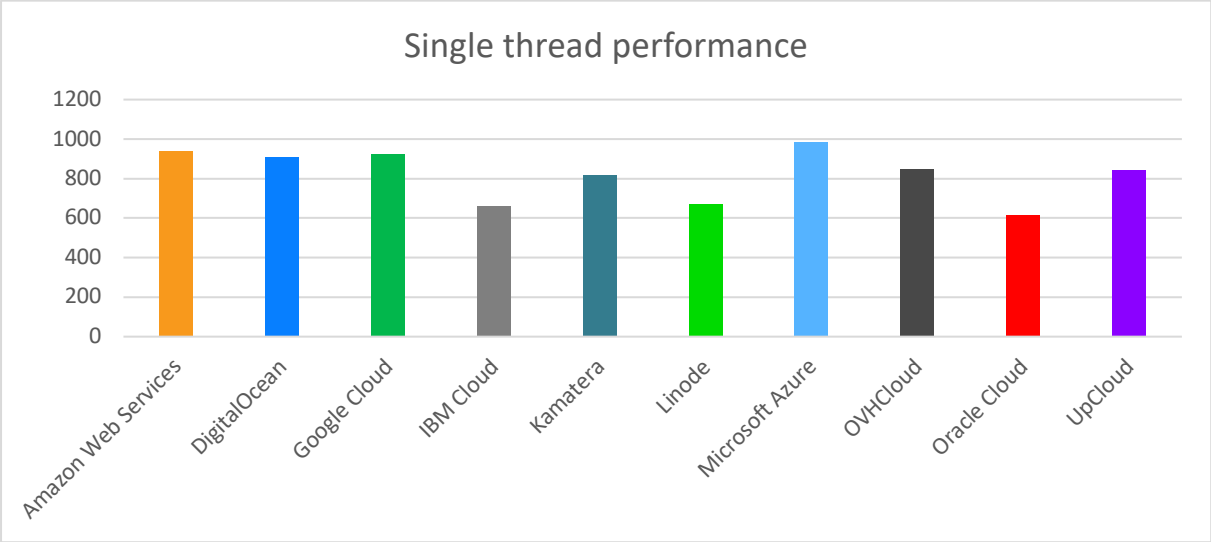

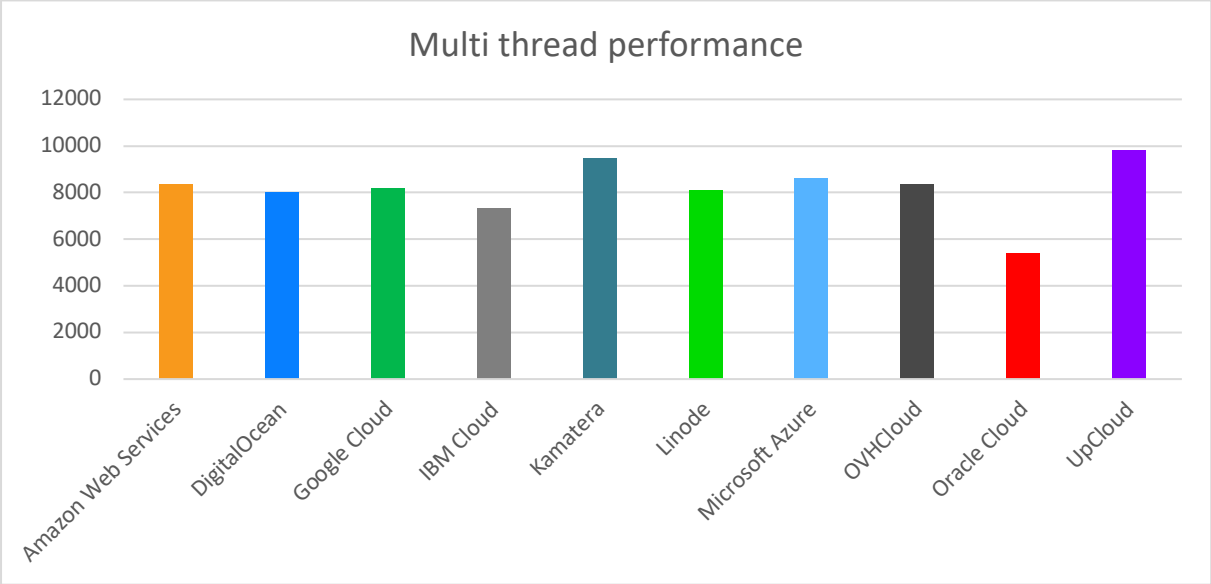

	Single thread		Multi thread	
	Mean	Deviation	Mean	Deviation
Amazon Web Services	937	35,35	8360	274,65
DigitalOcean	906	6,23	8006	46,11
Google Cloud	923	7,65	8193	28,16
IBM Cloud	658	4,29	7339	255,88
Kamatera	819	79,94	9475	2013,48
Linode	668	16,33	8081	335,80
Microsoft Azure	983	24,69	8601	149,50
OVHCloud	848	28,57	8363	1126,93
Oracle Cloud	613	2,28	5390	51,28
UpCloud	841	50,98	9816	1438,56

2. Storage IOPS

a. Overall

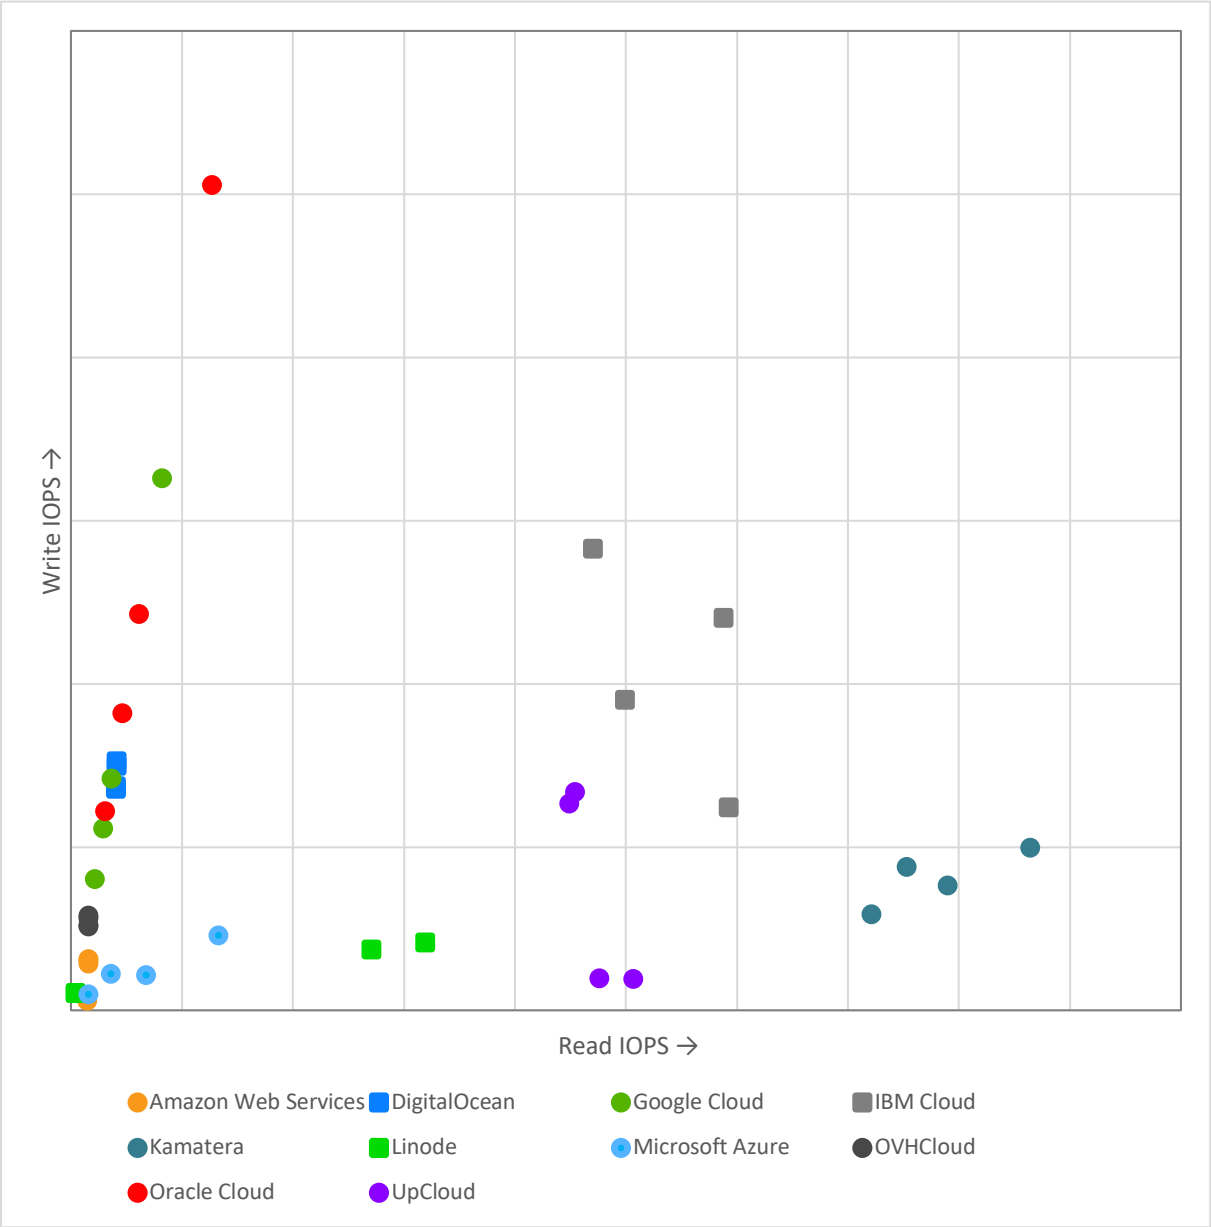

b. By category

i. Small

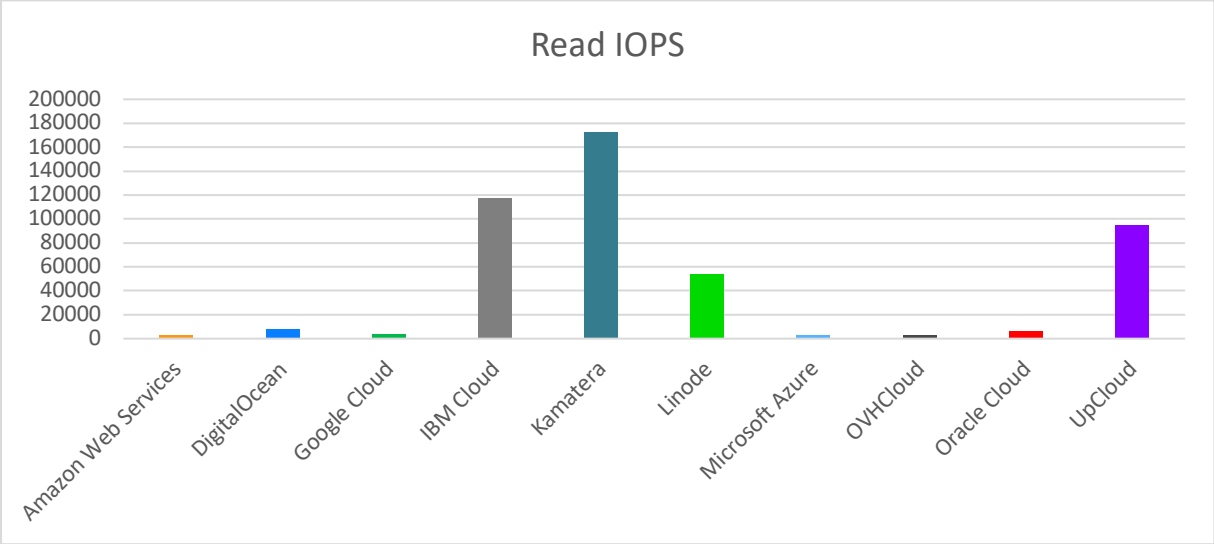

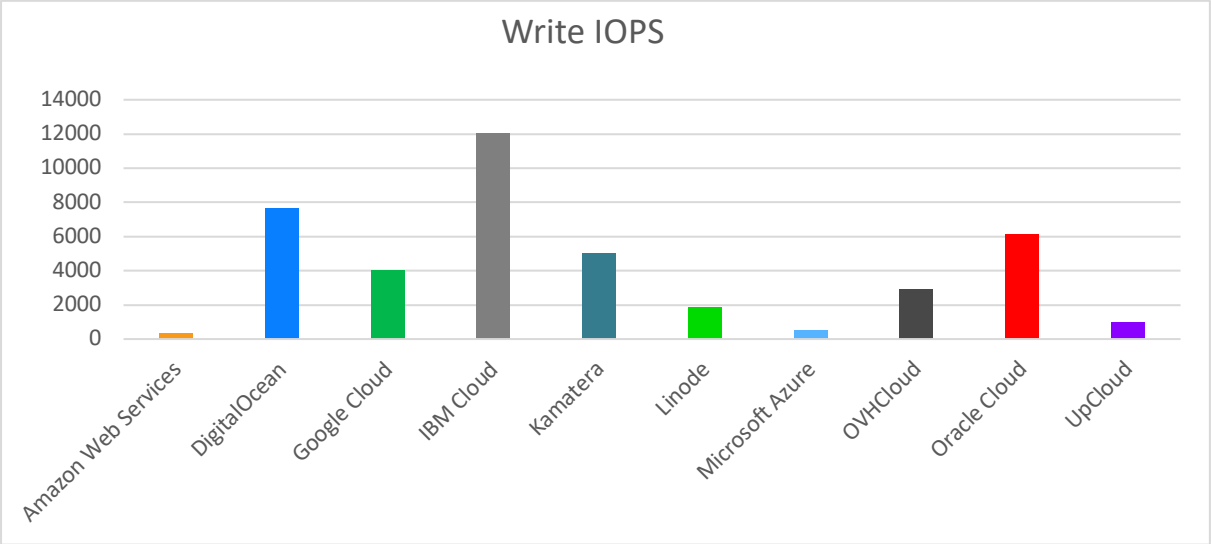

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	2891	230,94	310	5,42
DigitalOcean	8170	913,77	7642	1247,68
Google Cloud	4172	370,58	4030	51,67
IBM Cloud	117448	16600,73	12032	3764,59
Kamatera	172756	16255,45	4997	4890,64
Linode	54030	30541,84	1869	527,08
Microsoft Azure	3118	942,51	510	0,71
OVHCloud	3055	0,71	2919	346,79
Oracle Cloud	6125	51,81	6101	17,65
UpCloud	95159	50,98	983	1438,56

ii. Medium

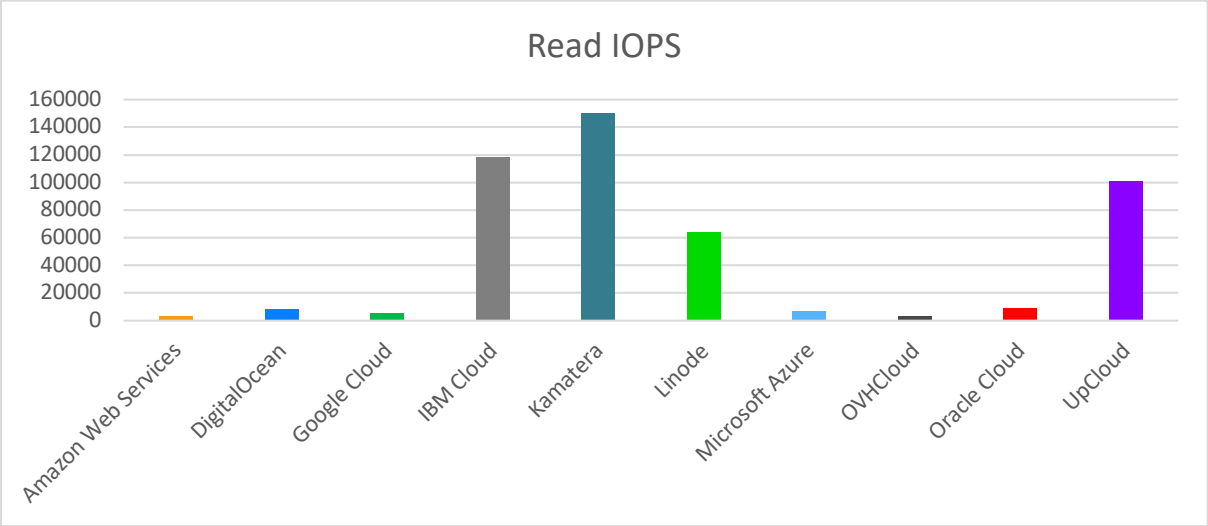

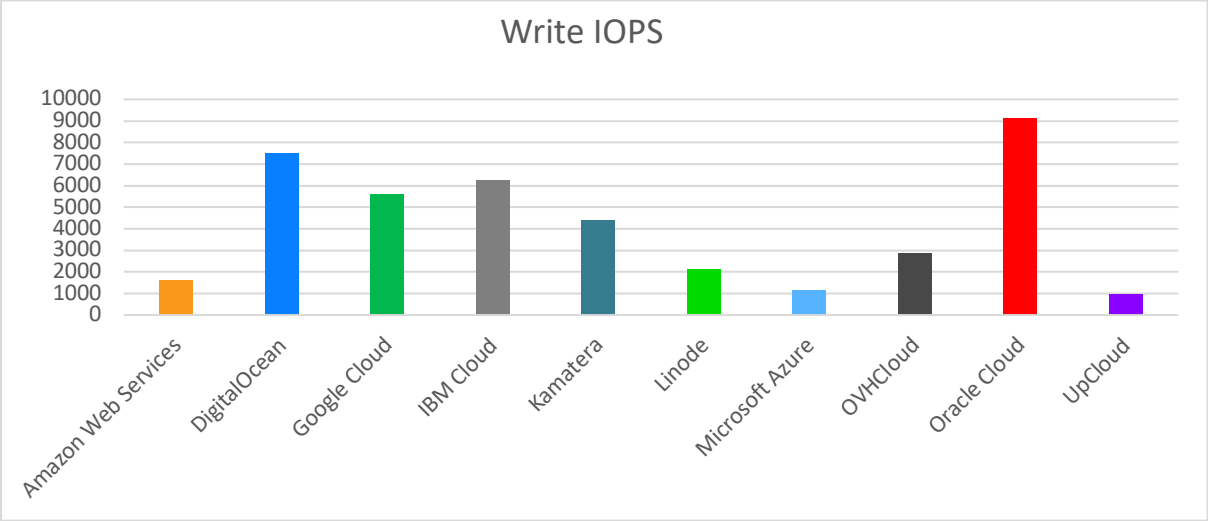

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	3045	2,08	1588	492,42
DigitalOcean	8169	898,61	7500	1 569,63
Google Cloud	5681	351,90	5573	74,20
IBM Cloud	118381	3 188,58	6227	5 617,44
Kamatera	150444	4 007,06	4404	4 100,21
Linode	63773	17 645,25	2094	910,50
Microsoft Azure	7149	1 467,44	1122	1,00
OVHCloud	3055	0,71	2843	493,22
Oracle Cloud	9130	74,75	9108	38,79
UpCloud	101211	2 126,93	980	15,43

iii. Large

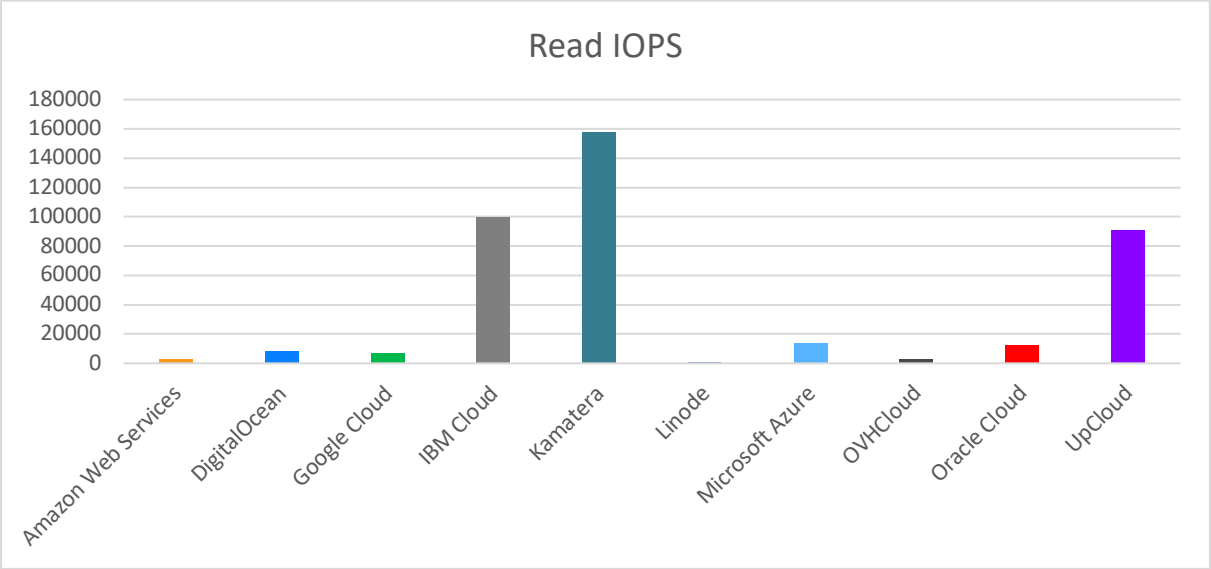

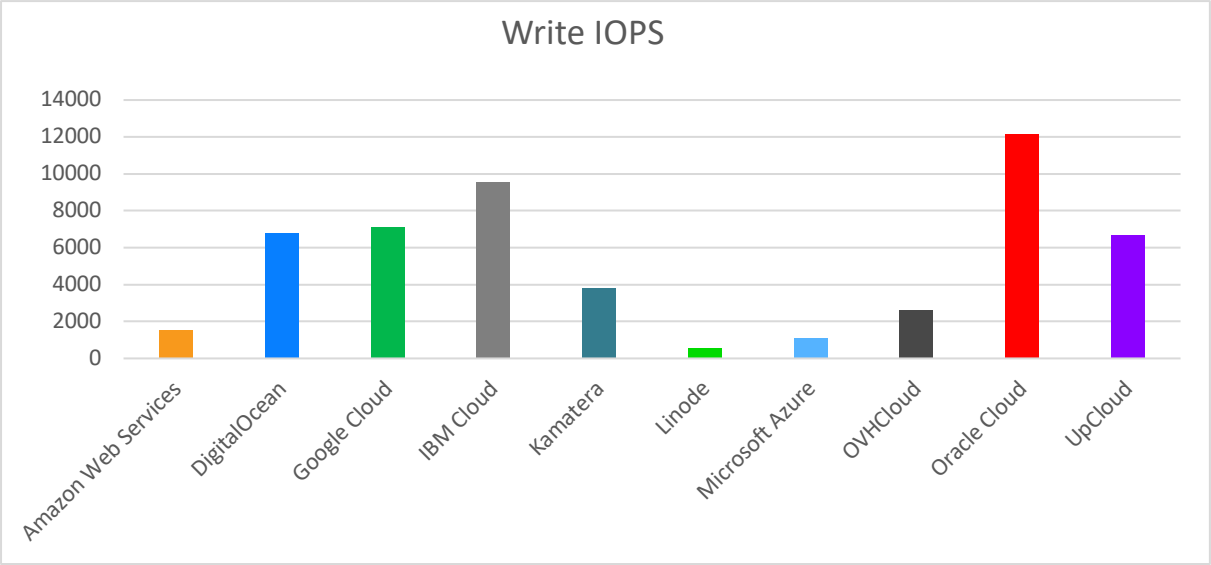

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	3041	6,07	1540	317,22
DigitalOcean	8015	785,71	6799	514,48
Google Cloud	7193	276,83	7121	94,91
IBM Cloud	99772	2366,20	9531	3239,37
Kamatera	157894	15110,07	3837	3712,67
Linode	816	387,78	542	177,20
Microsoft Azure	13474	3278,77	1088	59,06
OVHCloud	3055	0,71	2628	652,38
Oracle Cloud	12156	85,74	12150	48,69
UpCloud	90704	6700	13384,33	627,53

iv. Extra Large

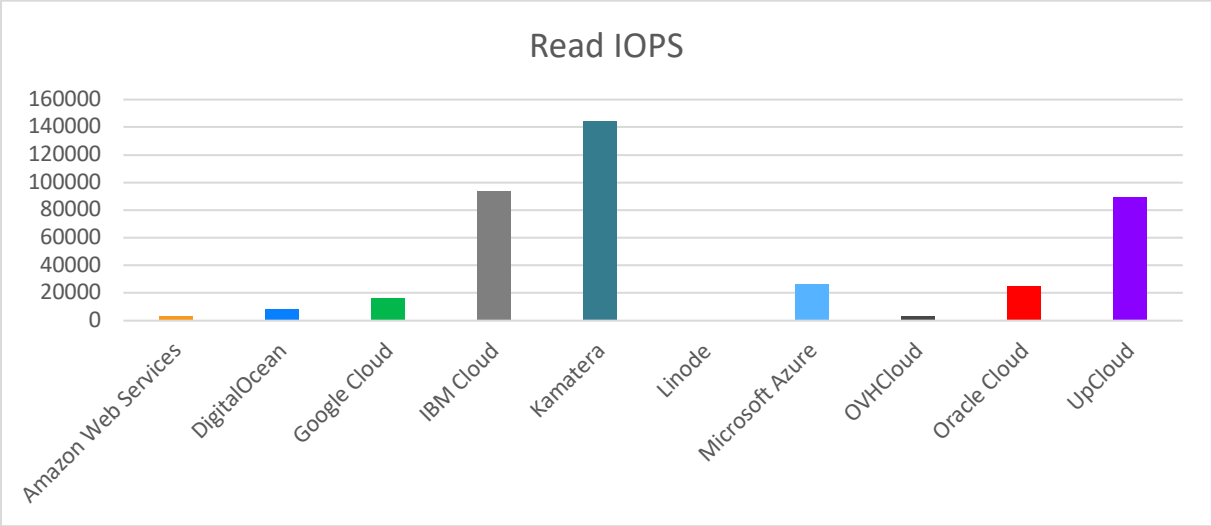

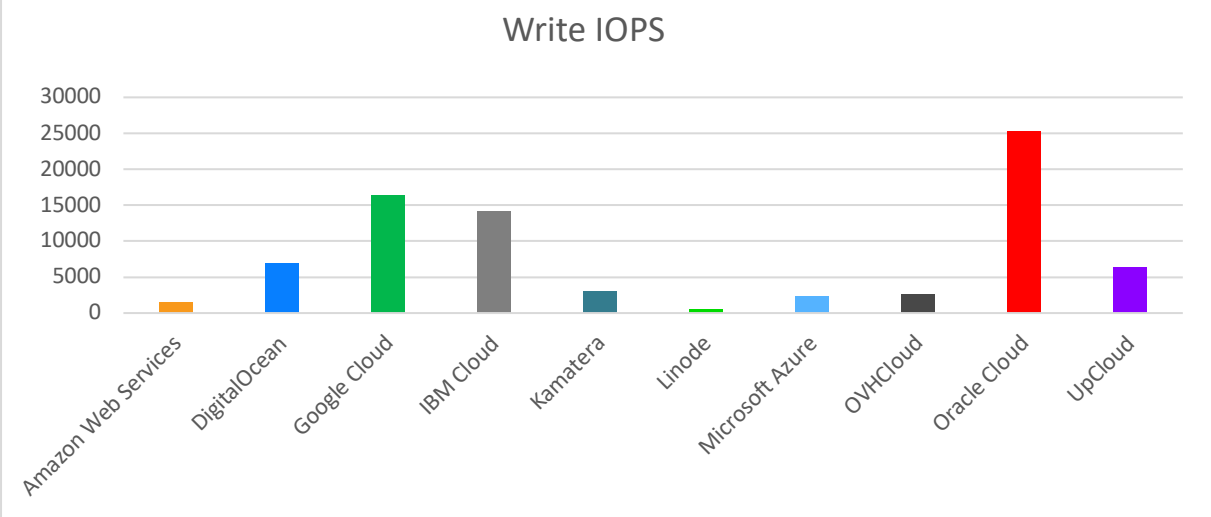

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	3045	3,11	1435	250,90
DigitalOcean	8039	849,49	6885	331,98
Google Cloud	16314	17,68	16312	21,23
IBM Cloud	93955	734,45	14159	4282,56
Kamatera	144191	20053,72	2959	3391,63
Linode	762	348,45	535	132,32
Microsoft Azure	26542	7352,95	2310	89,78
OVHCloud	3054	1,00	2582	717,87
Oracle Cloud	25277	89,96	25282	112,17
UpCloud	89687	11478,14	6353	1044,05

3. Storage Bandwidth

a. Overall

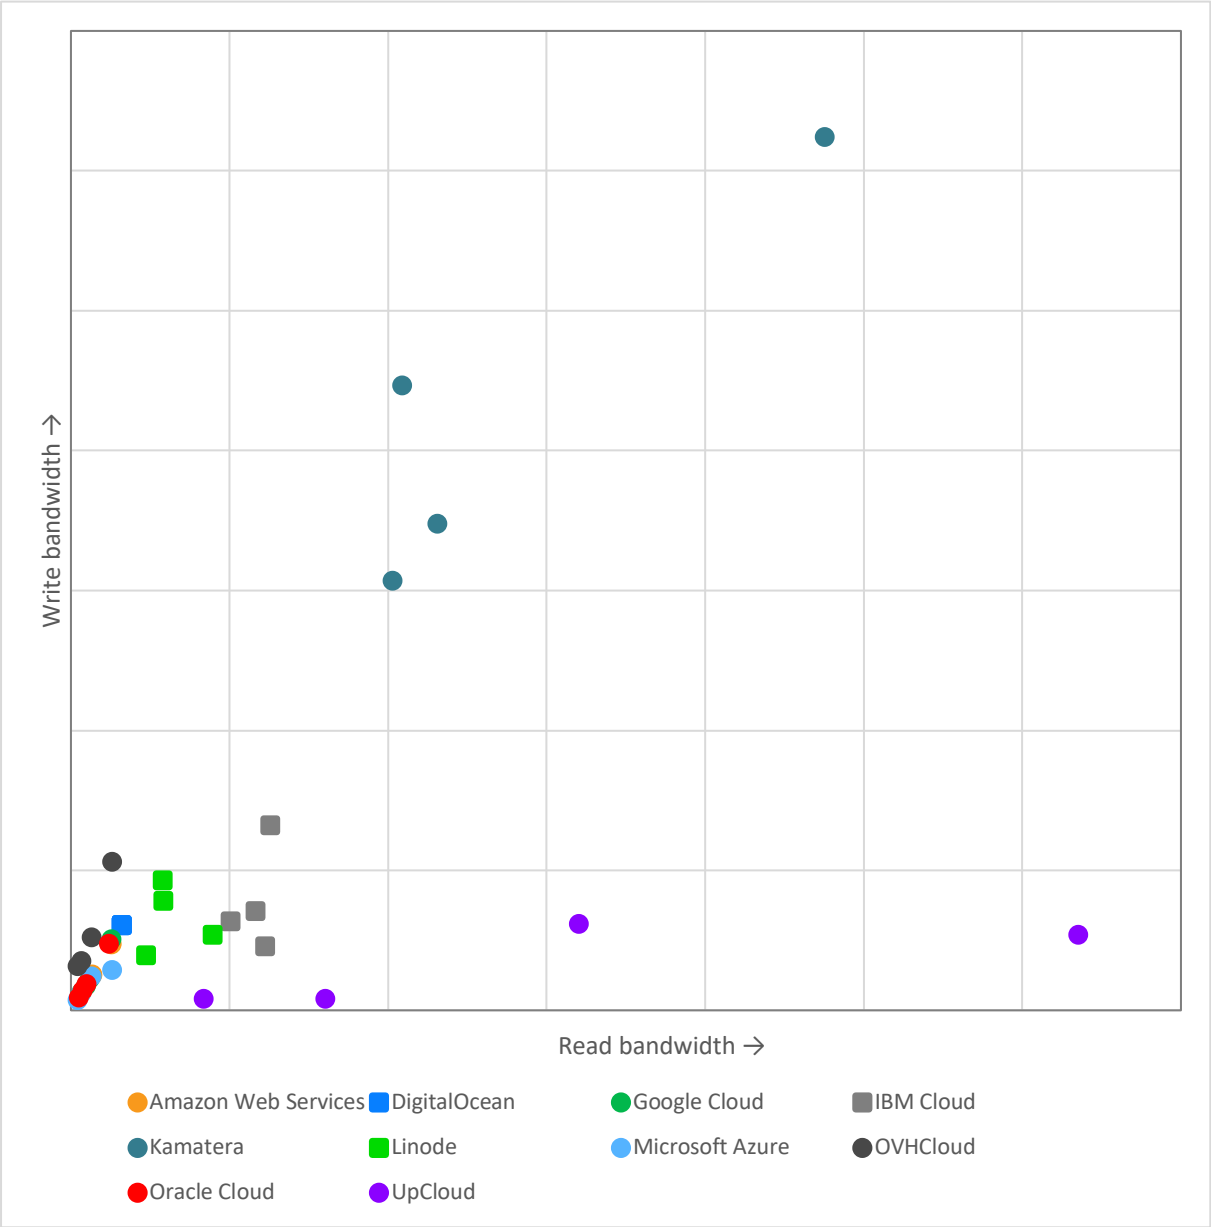

b. By category

i. Small

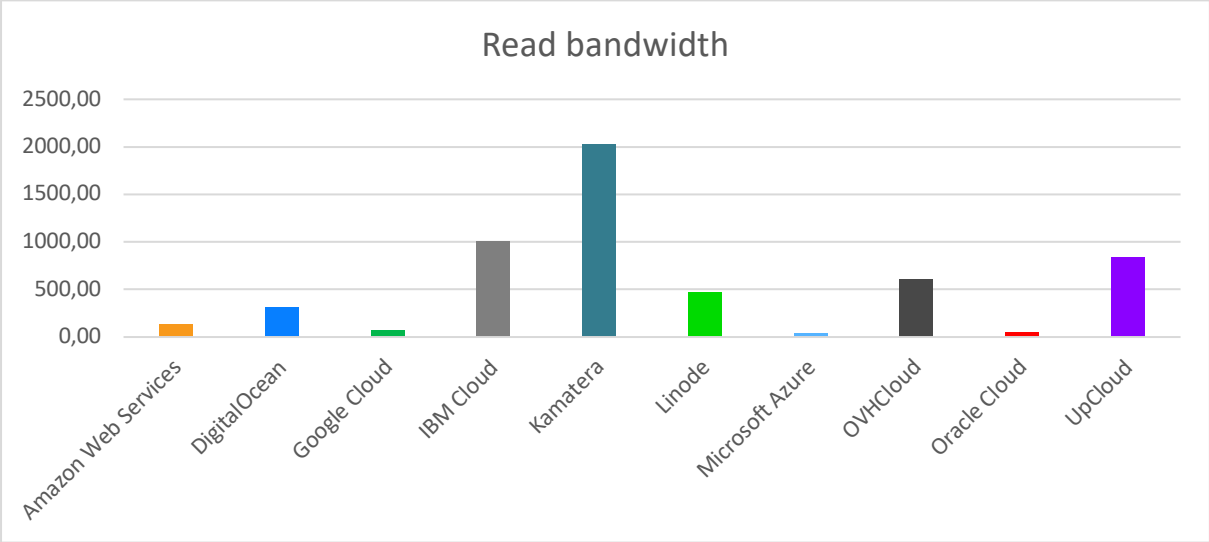

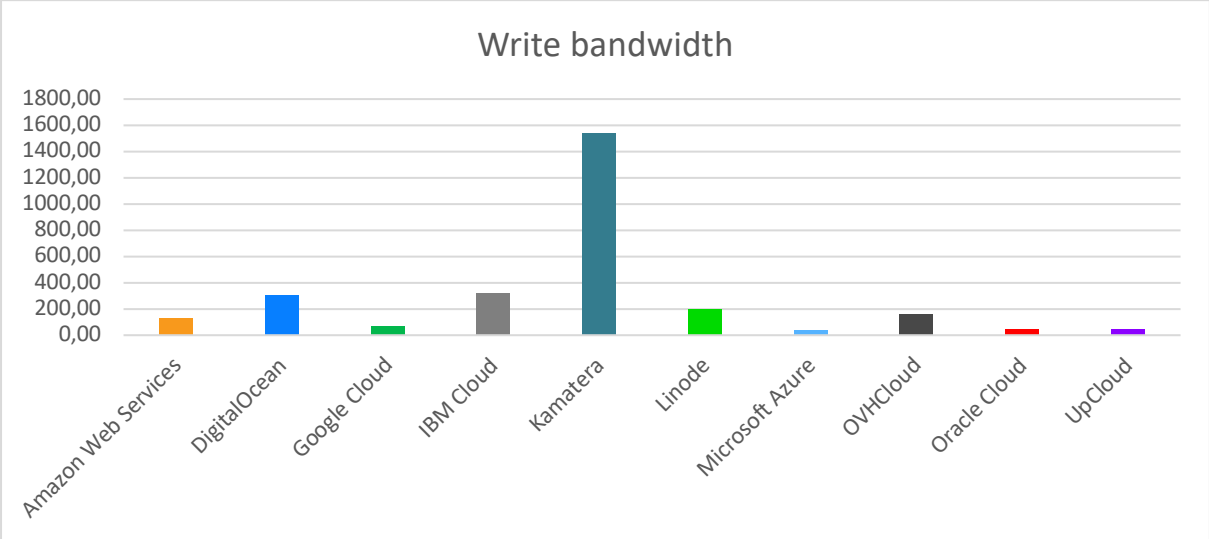

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	130,19	0,03	129,20	1,55
DigitalOcean	315,80	35,33	305,83	1,84
Google Cloud	64,74	2,99	64,90	3,35
IBM Cloud	1005,76	93,54	318,83	37,29
Kamatera	2025,32	1006,40	1536,26	920,27
Linode	471,46	39,44	197,34	22,59
Microsoft Azure	37,47	7,24	38,77	7,47
OVHCloud	600,72	269,74	158,61	44,85
Oracle Cloud	47,72	0,22	47,69	0,15
UpCloud	833,42	41,74	24,57	0,47

ii. Medium

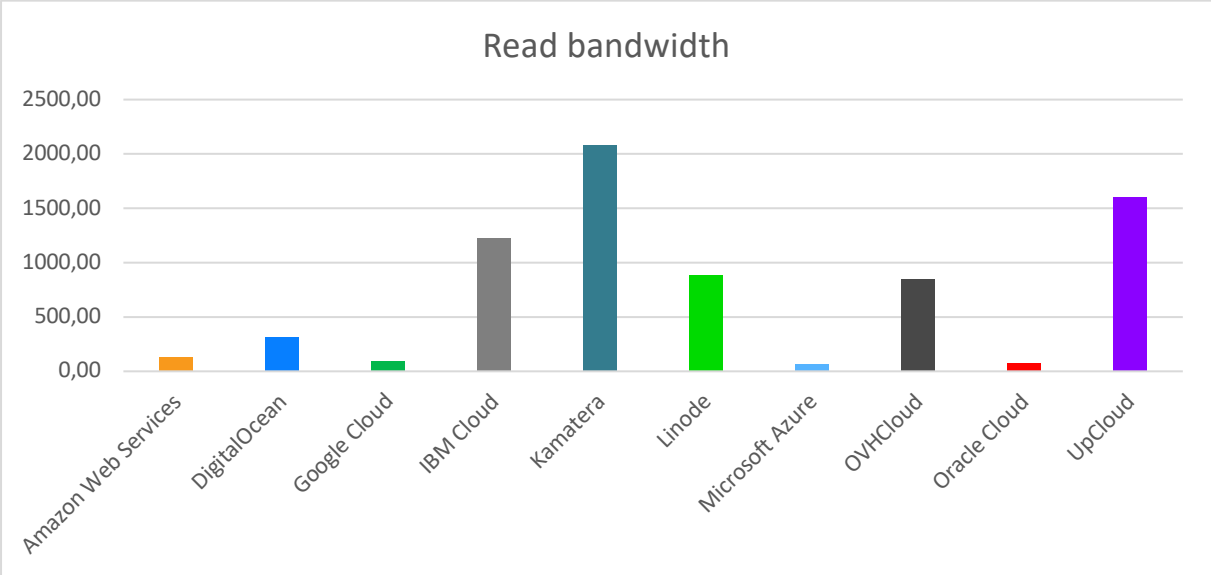

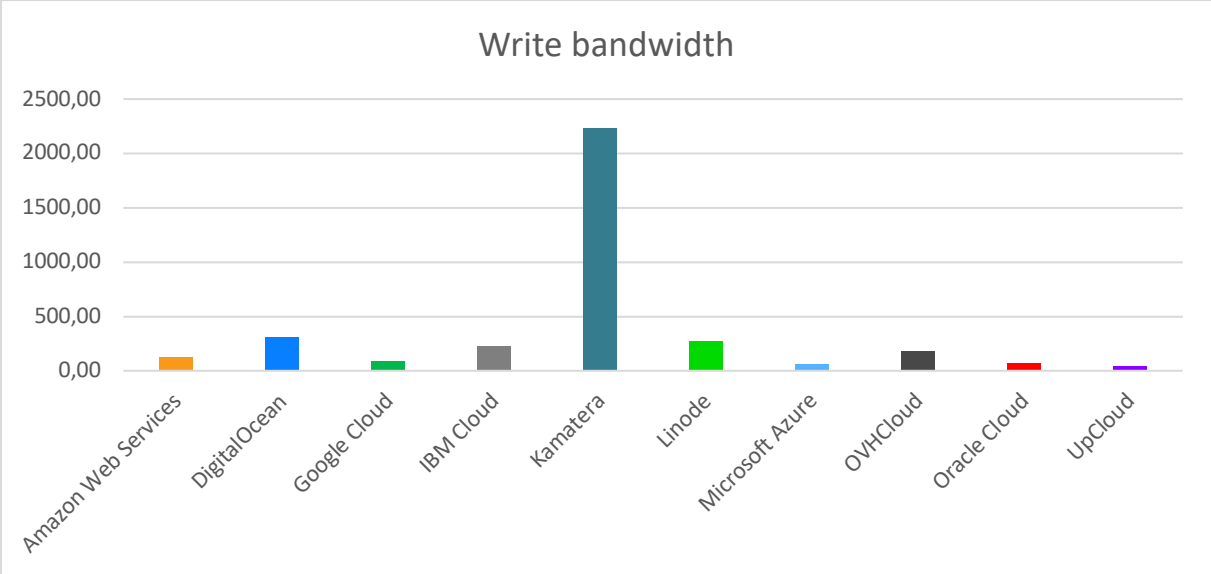

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	129,72	0,92	128,14	3,20
DigitalOcean	317,44	38,18	305,79	1,56
Google Cloud	89,97	3,74	89,21	3,30
IBM Cloud	1222,35	20,58	229,39	58,61
Kamatera	2084,21	1003,20	2235,32	1225,29
Linode	888,77	126,57	272,31	18,08
Microsoft Azure	63,83	0,02	63,66	0,78
OVHCloud	847,39	114,75	176,98	36,56
Oracle Cloud	71,74	0,28	71,71	0,29
UpCloud	1603,25	9,71	42,86	0,59

iii. Large

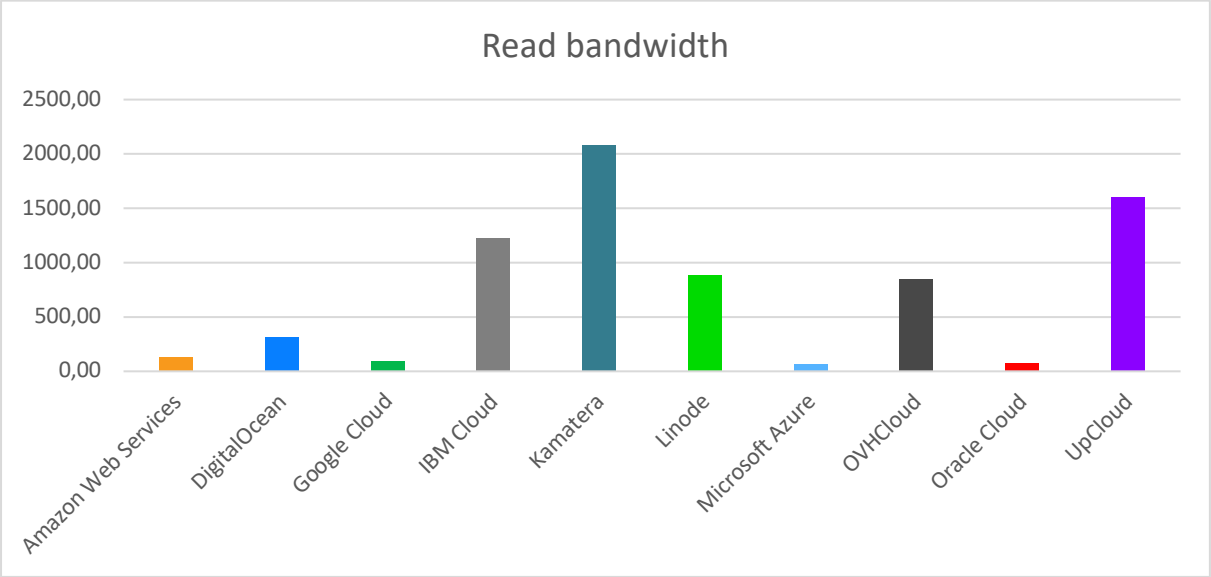

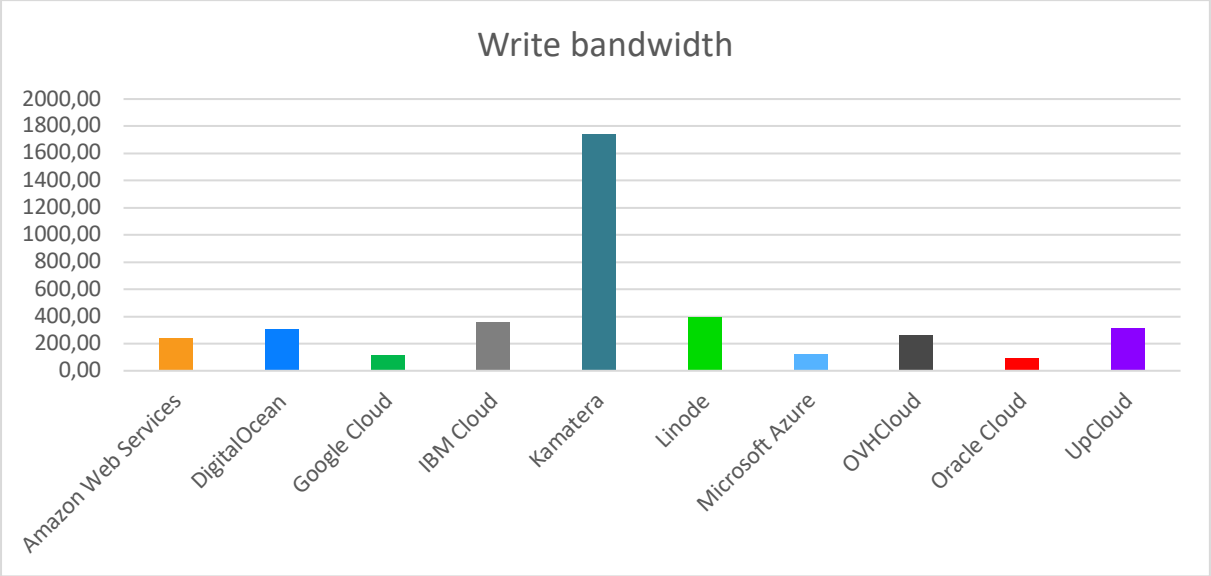

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	254,14	0,57	239,56	19,60
DigitalOcean	318,09	39,85	305,71	2,19
Google Cloud	113,95	2,34	114,09	2,77
IBM Cloud	1160,12	31,41	355,30	92,57
Kamatera	2309,25	988,70	1739,33	2081,12
Linode	581,18	147,68	392,86	104,35
Microsoft Azure	127,57	0,04	121,47	0,62
OVHCloud	722,32	178,42	262,21	124,05
Oracle Cloud	95,37	0,30	95,25	0,16
UpCloud	3199,33	31,70	309,73	8,31

iv. Extra Large

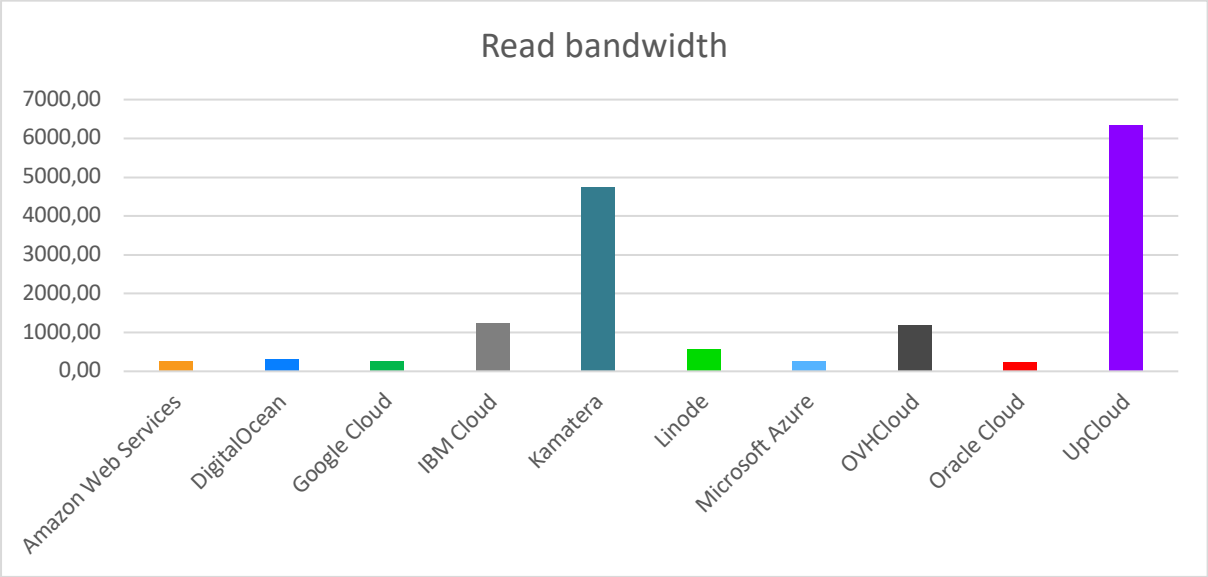

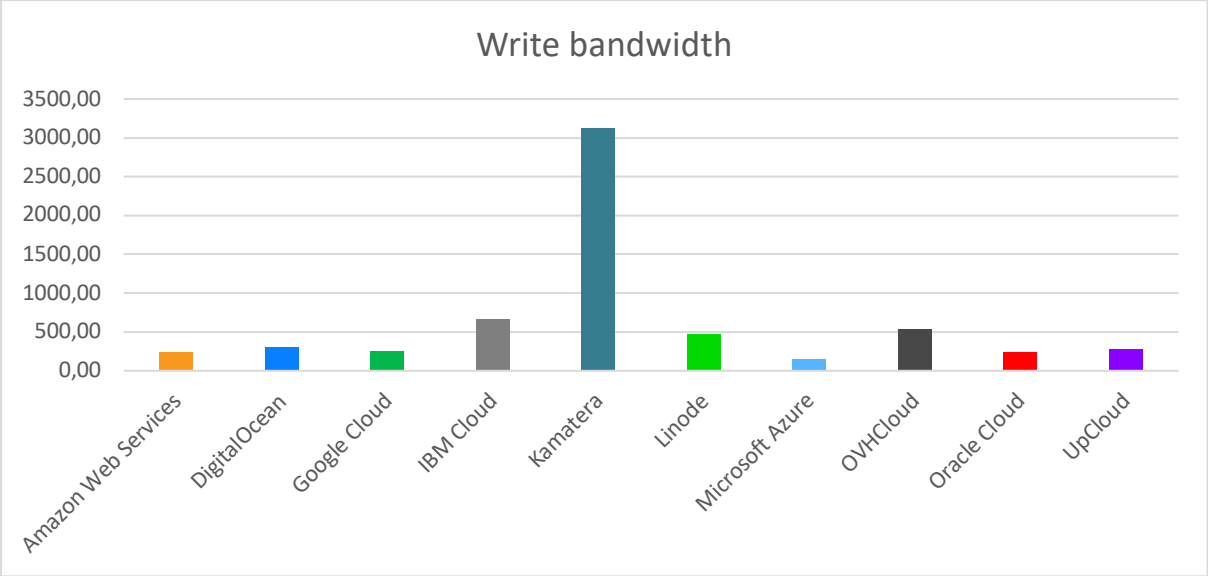

	Read Mean	Deviation	Write Mean	Deviation
Amazon Web Services	254,34	0,12	238,00	21,13
DigitalOcean	317,89	39,69	306,41	2,66
Google Cloud	254,62	0,01	254,59	0,02
IBM Cloud	1253,46	6,00	662,54	9,90
Kamatera	4752,55	354,16	3122,12	1729,90
Linode	576,74	144,91	464,70	121,94
Microsoft Azure	255,13	0,09	145,96	0,07
OVHCloud	1193,48	50,45	531,30	40,37
Oracle Cloud	238,51	0,76	238,38	0,43
UpCloud	6347,61	271,36	38,83	12,12

VI. Pricing

In comparison of similar products, price is one of the main metrics differentiating the value of an offer. Like performance, pricing is unequal across providers, we can observe a large gap between the smaller price and the most expensive virtual machine. Same thing for volume, where block storage varies from the most competitive to the less affordable pricing.

This section summarizes the cost of VMs and block storage for all products studied in the document.

To improve readability and give a better understanding of presented data, we convert hourly pricing into a monthly one by multiplying by 730. Vendors may apply custom discount after a long-term usage or for a monthly billing, this is not reflected in this document.

1. Virtual machines

a. Overall

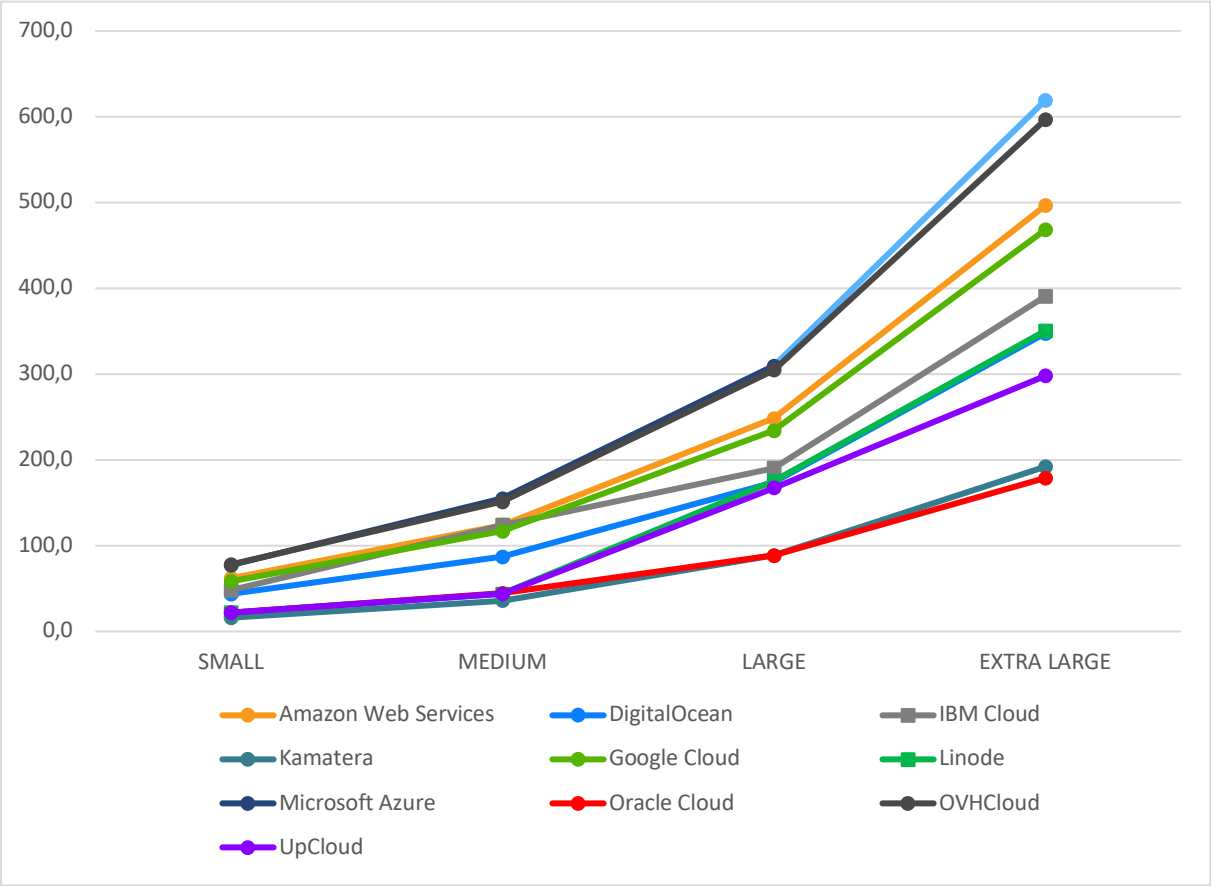

	SMALL	MEDIUM	LARGE	EXTRA LARGE
Amazon Web Services	62,1	124,1	248,2	496,4
DigitalOcean	43,8	86,9	173,7	347,5
IBM Cloud	48,2	124,1	190,5	390,6
Kamatera	16,1	35,8	88,3	192,0
Google Cloud	58,4	117,0	234,1	468,2
Linode	21,9	43,8	175,2	350,4
Microsoft Azure	77,4	154,8	309,5	619,0
Oracle Cloud	21,9	44,5	88,3	178,9
OVHCloud	78,1	151,1	304,4	596,4
UpCloud	21,9	43,8	167,2	297,8

b. By category

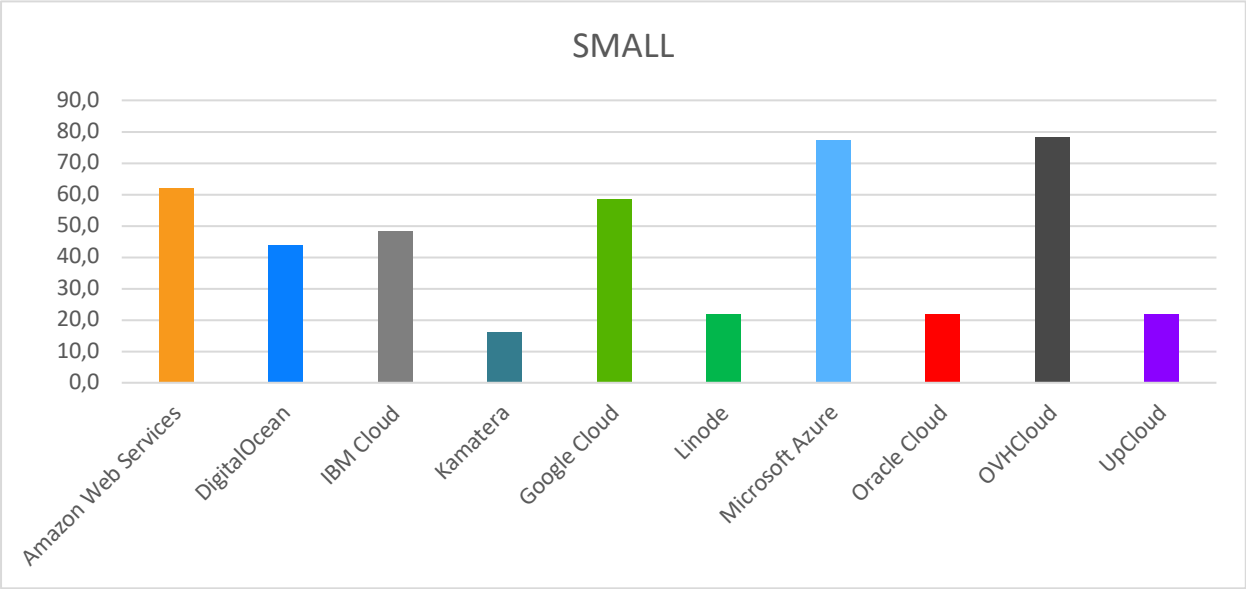

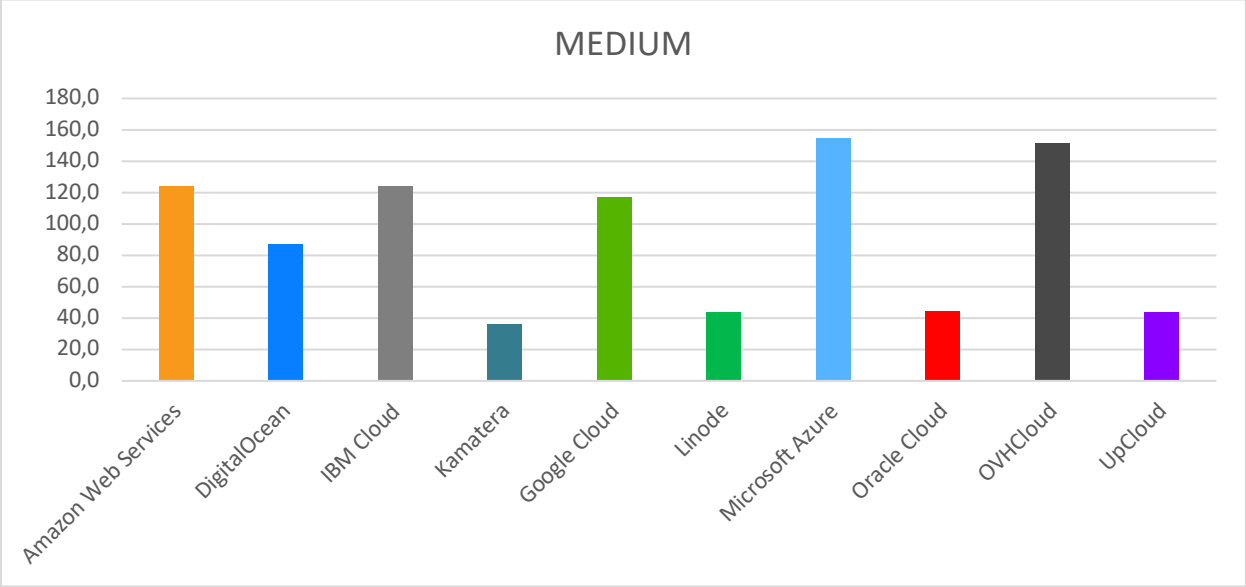

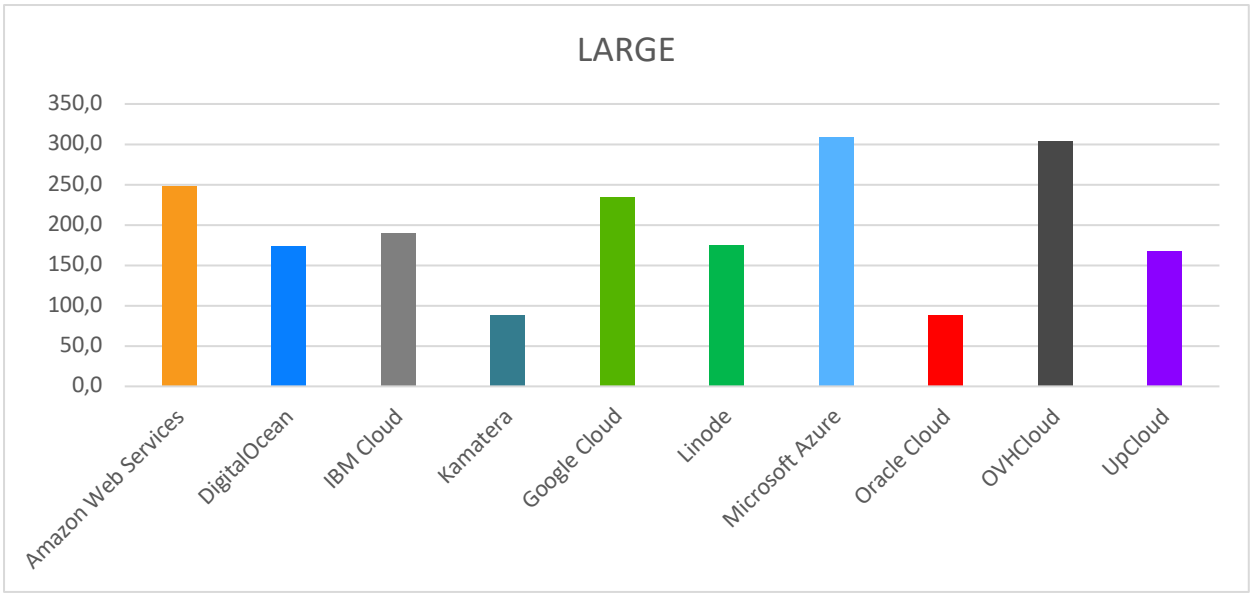
LARGE
350,0
300,0
250,0
200,0
150,0
100,0
50,0
0,0
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

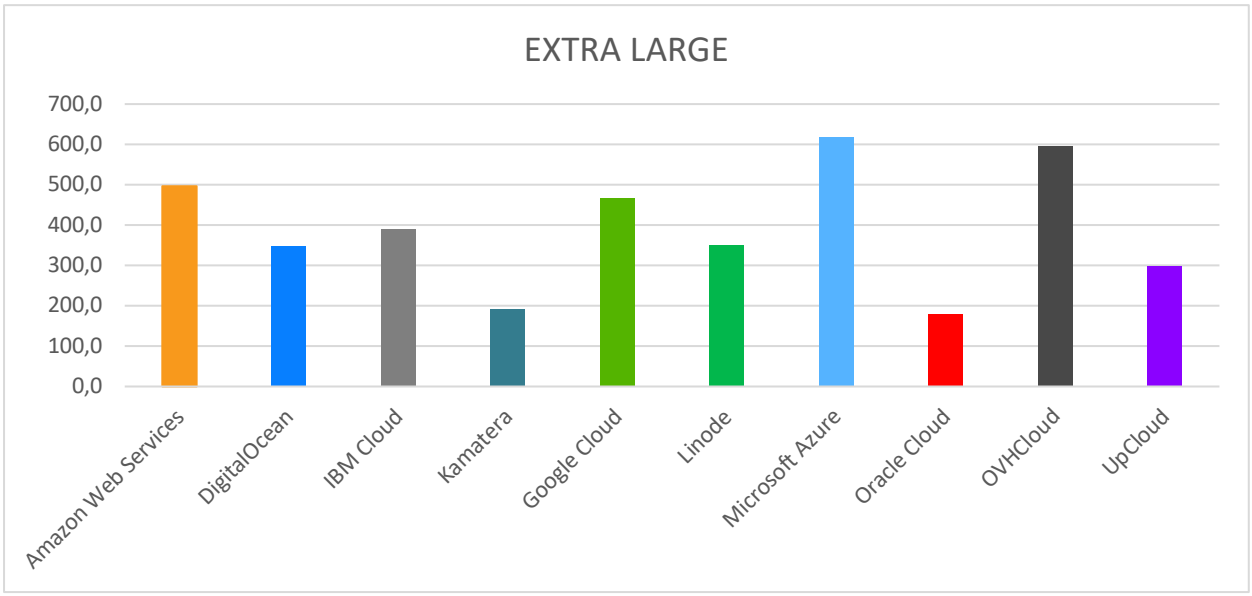
EXTRA LARGE
700,0
600,0
500,0
400,0
300,0
200,0
100,0
0,0
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

2. Storage

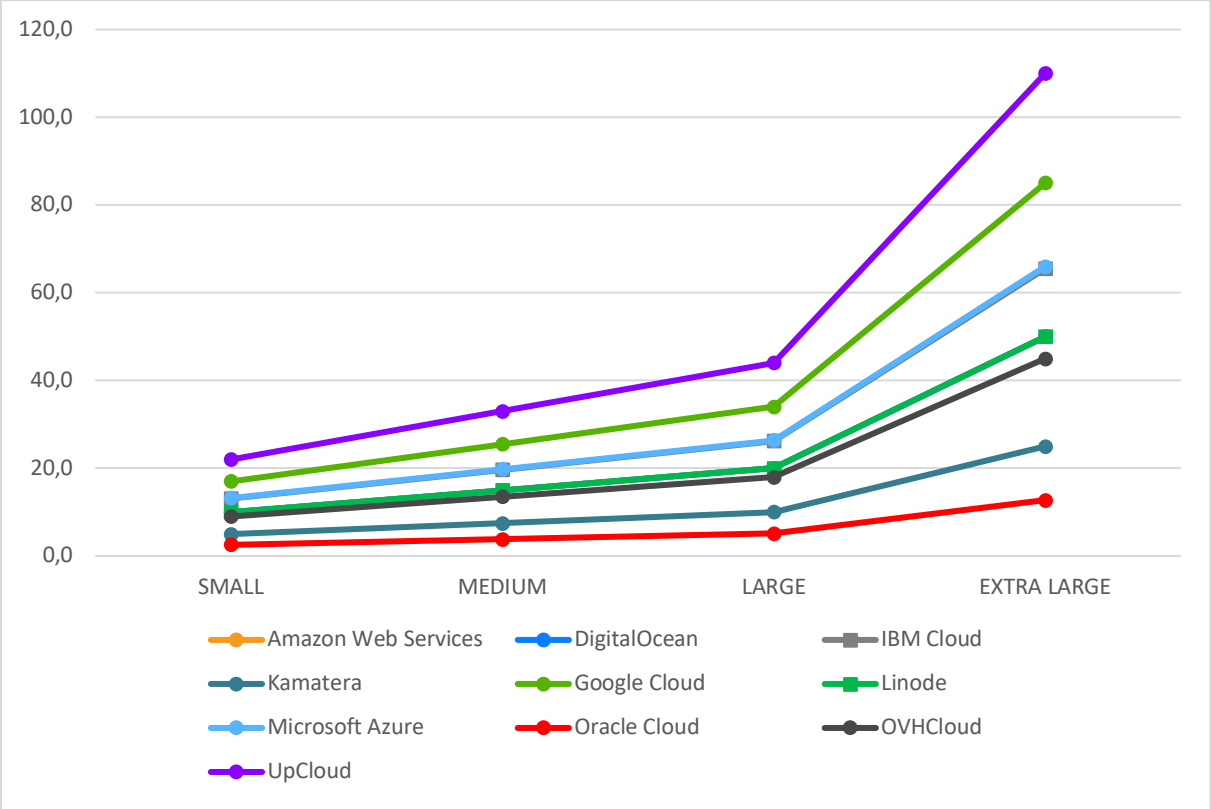

	SMALL	MEDIUM	LARGE	EXTRA LARGE
Amazon Web Services	10,0	15,0	20,0	50,0
DigitalOcean	10,0	15,0	20,0	50,0
IBM Cloud	13,1	19,7	26,2	65,5
Kamatera	5,0	7,5	10,0	25,0
Google Cloud	17,0	25,5	34,0	85,0
Linode	10,0	15,0	20,0	50,0
Microsoft Azure	13,2	19,8	26,4	66,0
Oracle Cloud	2,6	3,8	5,1	12,8
OVHCloud	9,0	13,5	18,0	45,0
UpCloud	22,0	33,0	44,0	110,0

VII. Price/Performance value

Consumers tend to search costs that match their budget but as seen in Performance section, a gap of 4x exists between the least and the best performer. Across providers, there's no correlation between performance and price, consumers cannot expect to behave the instinctive way where he/she pays more to obtain better. Moreover, compute and storage can draw a completely different profile.

This section outlines price/performance ratio for each category. It concludes this analysis by providing high-level scores aggregating metrics with values representing the best effort by dollar spent.

1. Overall

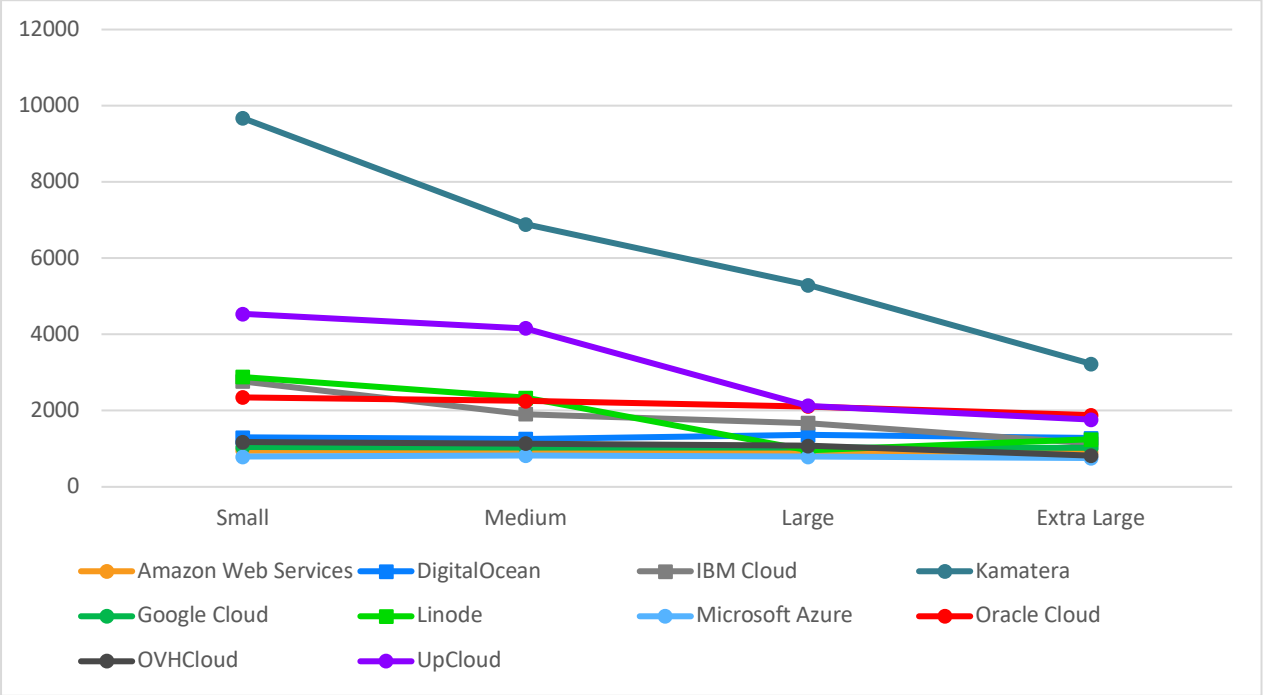

2. By category

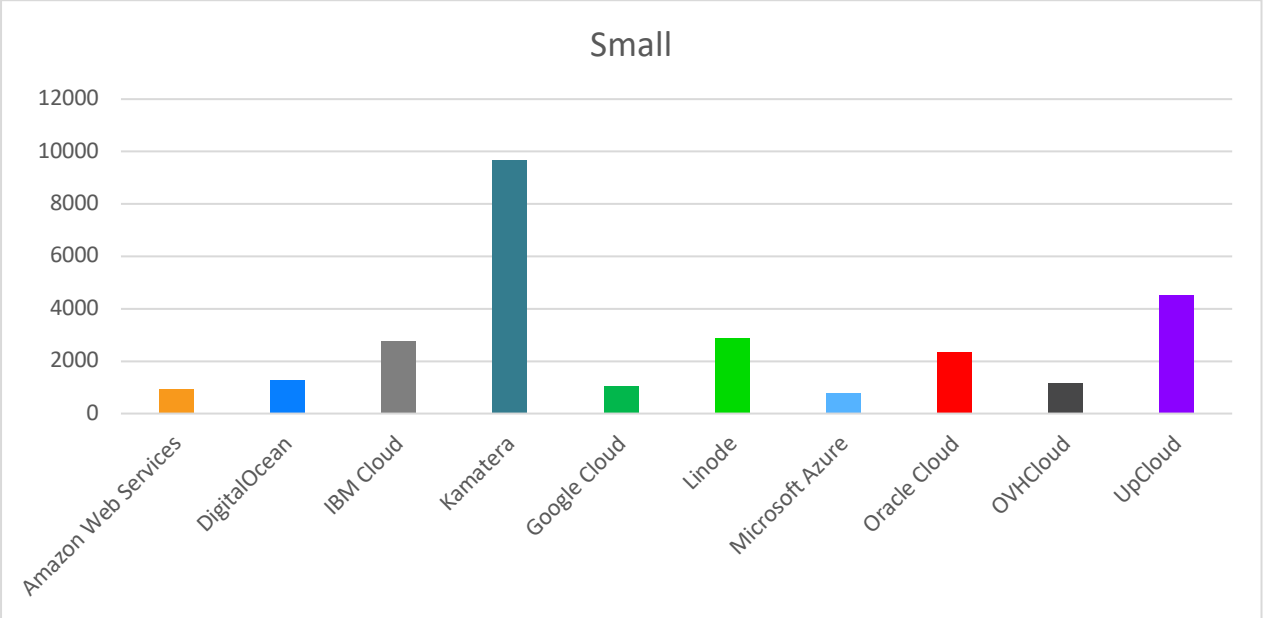

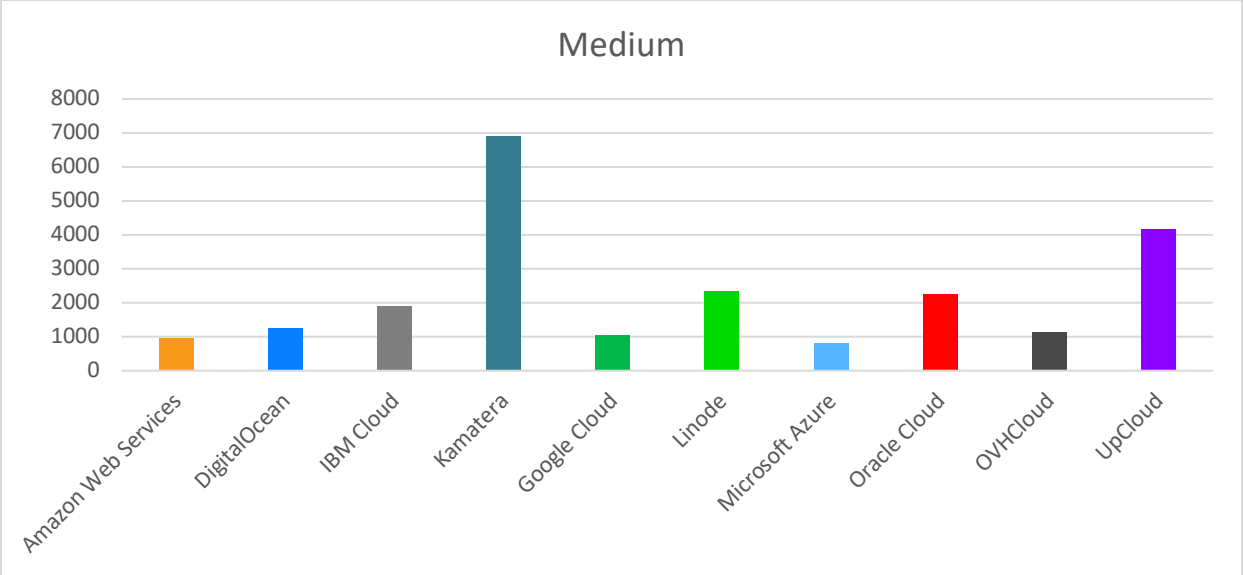
Medium
8000
7000
6000
5000
4000
3000
2000
1000
0
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

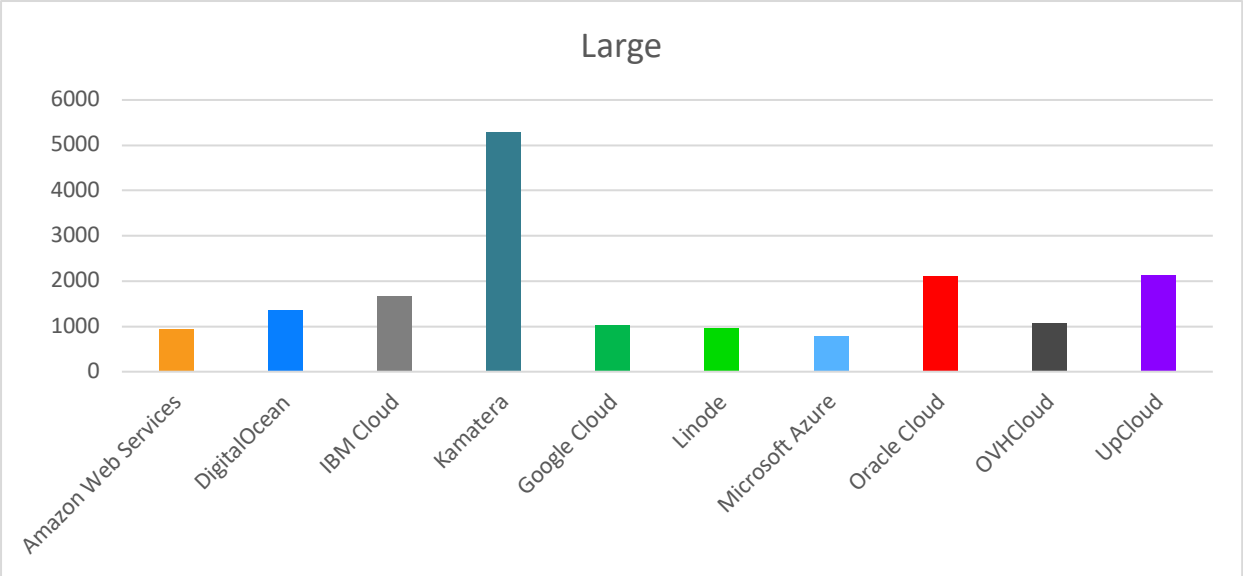
Large
6000
5000
4000
3000
2000
1000
0
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

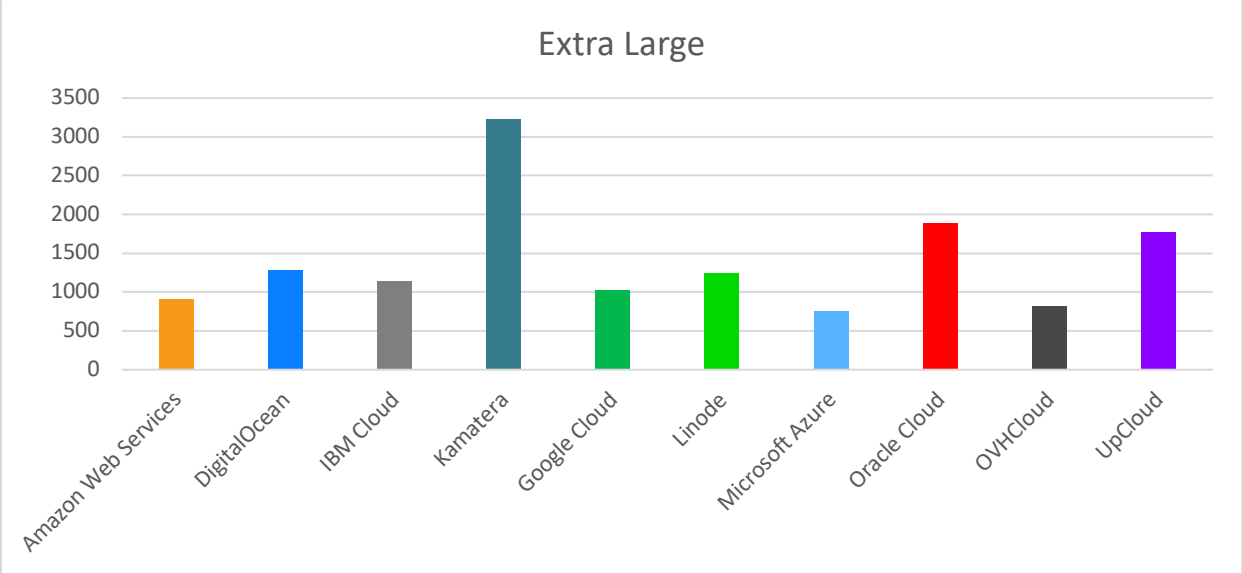
Extra Large
3500
3000
2500
2000
1500
1000
500
0
Amazon Web Services
DigitalOcean
IBM Cloud
Kamatera
Google Cloud
Linode
Microsoft Azure
Oracle Cloud
OVHCloud
UpCloud

VIII. About Cloud Mercato

Cloud Mercato is a research and consultancy firm dedicated to study of cloud market. We are a team expert working in cloud benchmark since 2013 helped by our automated software which actively collect and monitor key metrics for all kind of products in the cloud industry. Our neutral and objective approach helps customer to get better insights on their possibilities and vendors to know how to place and compete between all these innumerable services.

For any inquiries about our services, question about this report or any custom request, please contact Cloud Mercato at contact@cloud-mercato.com.

cloud-mercato.com

If you are a Cloud Service Provider and you are interested to be included in our research and documents, do not hesitate to contact us.

i. Appendix

1. Server specifications

	NAME	vCPU	RAM	Storage Type
Amazon Web Services	c5.large	2	4	General Purpose SSD
	c5.xlarge	4	8	
	c5.2xlarge	8	16	
	c5.4xlarge	16	32	
DigitalOcean	Optimized 2CPU 4GB	2	4	Block storage
	Optimized 4CPU 8GB	4	8	
	Optimized 8CPU 32GB	8	16	
	Optimized 16CPU 32GB	16	32	
IBM Cloud	2 Cores 4GB	2	4	Portable Storage SAN
	4 Cores 8GB	4	8	
	8 Cores 16GB	8	16	
	16 Cores 32GB	16	32	
Kamatera	2ACPU 4GB	2	4	Cloud Block Storage
	4ACPU 8GB	4	8	
	8ACPU 16GB	8	16	
	16ACPU 32GB	16	32	
Google Cloud	N2 2vCPU 4GB	2	4	Standard Persistent SSD
	N2 4vCPU 8GB	4	8	
	N2 8vCPU 16GB	8	16	
	N2 16vCPU 32GB	16	32	
Linode	Linode 4GB	2	4	Block Storage
	Linode 8GB	4	8	
	Linode 32GB	8	32	
	Linode 64GB	16	64	
Microsoft Azure	Standard F2s v2	2	4	No cache Premium LRS
	Standard F4s v2	4	8	
	Standard F8s v2	8	16	
	Standard F16s v2	16	32	
Oracle Cloud	VM.Standard.E2.1	2	4	Block volume
	VM.Standard.E2.2	4	8	
	VM.Standard.E2.4	8	16	
	VM.Standard.E2.8	16	32	
OVHCloud	C2-7	2	7.5	High-speed
	C2-15	4	15	
	C2-30	8	30	
	C2-60	16	60	
Upcloud	2xCPU-4GB	2	4	MaxIOPS
	4xCPU-8GB	4	8	
	Custom 8xCPU-16GB	8	16	

	Custom 16xCPU-32GB	16	32	

2. Server additionnal features

- DigitalOcean:
  - Each instance has a root SSD volume given for free. Their sizes have a ratio of 12.5GB per vCPU, then 25, 50, 100 and 200GB
- Linode:
  - Each instance has a root SSD volume given for free. Their sizes have a ratio of 20GB per GB of RAM, then 80, 160, 640 and 1280GB
- Microsoft Azure:
  - Each instance has an extra SSD volume given for free. Their sizes have a ratio of 8GB per vCPU, then 16, 32, 64 and 128GB
- OVHCloud:
  - Each instance has a root SSD volume given for free. Their sizes 50GB for all
- UpCloud:
  - UpCloud non-custom instances have a root SSD volume given for free. They have a size of 80GB for 2xCPU-4GB and 160GB for 16xCPU-8GB

3. Volume additionnal features

- Amazon Web Services:
  - Performance of General Purpose SSD depends of volume size. It increases linearly by 100 IOPS per GB with a maximum of 16,000 IOPS
  - This volume type work in a burstable mode meaning that maximum IOPS are limited in the time
- Microsoft Azure:
  - Performance of Premium LRS depends of volume size. More accurately, of the total size of volume attached.

4. CPU specifications

	MODEL NAME	FREQUENCY	REALEASE
Amazon Web Services	Intel Xeon Platinum 8124M Intel Xeon Platinum 8275CL	3.0	
DigitalOcean	Intel(R) Xeon(R) CPU E5-2697A v4	2.6	Q1 2016

	Intel(R) Xeon(R) Platinum 8168	2.7	Q3 2017
IBM Cloud	Intel(R) Xeon(R) Gold 6140	2.3	Q3 2017
	Intel(R) Xeon(R) CPU E5-2683 v4	2.1	Q1 2016
Kamatera	Intel(R) Xeon(R) Gold 6150	2.7	Q3 2017
Google Cloud	Intel(R) Xeon(R) CPU		
Linode	AMD EPYC 7601	2.2	Q3 2017
	AMD EPYC 7501	2.0	
	Intel(R) Xeon(R) Platinum 8168	2.7	
Microsoft Azure	Intel(R) Xeon(R) Platinum 8168	2.7	Q3 2017
Oracle Cloud	AMD EPYC 7551	2.0	Q3 2017
OVHCloud	Intel Core Processor (Haswell, no TSX)	2.4	
UpCloud	Intel(R) Xeon(R) Gold 6136	3.0	Q3 2017

5. Test scripts

The pieces of code below are part of our methodology and runnable in a Linux command line environment, you can copy them as they are and must set the following variable to make them operate correctly:

- cpu_number: The number of vCPU available on machine
- rw: Access mode 'read' or 'write'
- device_path: Absolute path to the raw device, i.e. /dev/vdb

Feel free to reproduce our results from these snippets.

a. Storage IOPS

```
fio --numjobs=$cpu_number \
        --bs=4k --rw=rand$rw \
        --ioengine=libaio --iodepth=32 \
        --direct=1 --invalidate=1 --end_fsync=1 \
        --time_based --runtime=60 --timeout=60 \
        --filename=$device_path \
        --group_reporting --output-format=json --name=fio
```

b. Storage bandwidth

```
fio --numjobs=$cpu_number \
        --bs=1m --rw=$rw \
        --ioengine=libaio --iodepth=32 \
        --direct=1 --invalidate=1 --end_fsync=1 \
        --time_based --runtime=60 --timeout=60 \
        --filename=$device_path
        --group_reporting --output-format=json --name=fio
```