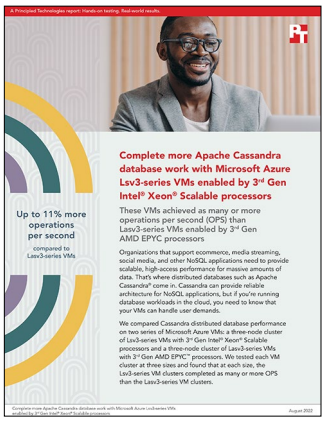

The science behind the report:

Complete more Apache Cassandra database work with Microsoft Azure Lsv3-series VMs enabled by 3rd Gen Intel® Xeon® Scalable processors

This document describes what we tested, how we tested, and what we found. To learn how these facts translate into real-world benefits, read the report Complete more Apache Cassandra database work with Microsoft Azure Lsv3-series VMs enabled by 3rd Gen Intel® Xeon® Scalable processors.

We concluded our hands-on testing on June 23, 2022. During testing, we determined the appropriate hardware and software configurations and applied updates as they became available. The results in this report reflect configurations that we finalized on June 15, 2022 or earlier. Unavoidably, these configurations may not represent the latest versions available when this report appears.

Our results

To learn more about how we have calculated the wins in this report, go to http://facts.pt/calculating-and-highlighting-wins. Unless we state otherwise, we have followed the rules and principles we outline in that document.

Table 1: Results of our testing.

Three-node cluster	Operations per second	Lsv3-series advantage
L8s_v3 VMs	194,820	N/A
L8as_v3 VMs	174,363	11.73%
L16s_v3 VMs	339,648	N/A
L16as_v3 VMs	313,023	8.50%
L32s_v3 VMs	470,115	N/A
L32as_v3 VMs	467,172	0.62%

System configuration information

Table 2: Detailed configuration information on the Lsv3 VMs we tested.

System configuration information	3x L8s_v3	3x L16s_v3	3x L32s_v3
Tested by	Principled Technologies	Principled Technologies	Principled Technologies
Test date	06/21/2022	06/21/2022	06/21/2022
CSP / Region	Microsoft Azure / South Central US	Microsoft Azure / South Central US	Microsoft Azure / South Central US
Workload and version	Cassandra and cassandra-stress 4.0.4	Cassandra and cassandra-stress 4.0.4	Cassandra and cassandra-stress 4.0.4
Workload-specific parameters	100% writes, Split load with a cassandra-stress instance on each node	100% writes, Split load with a cassandra-stress instance on each node	100% writes, Split load with a cassandra-stress instance on each node
Iterations and result choice	3 runs, median	3 runs, median	3 runs, median
Server platform	L8s_v3	L16s_v3	L32s_v3
BIOS name and version	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021
Operating system name and version/build number	Ubuntu 20.04.4 LTS	Ubuntu 20.04.4 LTS	Ubuntu 20.04.4 LTS
Date of last OS updates/patches applied	06/15/2022	06/15/2022	06/15/2022
Processor			
Number of processors	1	1	1
Vendor and model	Intel® Xeon® Platinum 8370C	Intel® Xeon® Platinum 8370C	Intel® Xeon® Platinum 8370C
Core count (per processor)	32	32	32
Core frequency (GHz)	2.80	2.80	2.80
Stepping	6	6	6
Hyper-threading	Yes	Yes	Yes
Turbo?	Yes	Yes	Yes
Number of vCPUs per VM	8	16	32
Memory module(s)			
Total memory in system (GB)	64	128	256
NVMe® memory present?	No	No	No
Total memory (DDR+NVMe RAM) (GB)	64	128	256
General hardware			
Storage: NW or Direct Att / Instance	Direct Att	Direct Att	Direct Att

System configuration information	3x L8s_v3	3x L16s_v3	3x L32s_v3
Local storage			
OS			
Number of drives	1	1	1
Drive size (GB)	30	30	30
Drive information	Premium SSD LRS	Premium SSD LRS	Premium SSD LRS
Data drive			
Number of drives	1	2	4
Drive size (TB)	1.92	1.92	1.92
Drive information	NVMe, 400,000 IOPS (instance max: 12,800 IOPS)	NVMe, 400,000 IOPS (instance max: 25,600 IOPS)	NVMe, 400,000 IOPS (instance max: 51,200 IOPS)
Network adapter			
Vendor and model	Mellanox Technologies MT27800 Family	Mellanox Technologies MT27800 Family	Mellanox Technologies MT27800 Family
Number and type of ports	1x 100Gb	1x 100Gb	1x 100Gb

Table 3: Detailed configuration information on the Lasv3 VMs we tested.

System configuration information	3x L8as_v3	3x L16as_v3	3x L32as_v3
Tested by	Principled Technologies	Principled Technologies	Principled Technologies
Test date	06/23/2022	06/23/2022	06/23/2022
CSP / Region	Microsoft Azure / South Central US	Microsoft Azure / South Central US	Microsoft Azure / South Central US
Workload and version	Cassandra and cassandra-stress 4.0.4	Cassandra and cassandra-stress 4.0.4	Cassandra and cassandra-stress 4.0.4
Workload-specific parameters	100% writes, Split load with a cassandra-stress instance on each node	100% writes, Split load with a cassandra-stress instance on each node	100% writes, Split load with a cassandra-stress instance on each node
Iterations and result choice	3 runs, median	3 runs, median	3 runs, median
Server platform	L8as_v3	L16as_v3	L32as_v3
BIOS name and version	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021	Microsoft Corporation Hyper-V UEFI Release v4.1, 7/22/2021
Operating system name and version/build number	Ubuntu 20.04.4 LTS	Ubuntu 20.04.4 LTS	Ubuntu 20.04.4 LTS
Date of last OS updates/patches applied	06/15/2022	06/15/2022	06/15/2022

System configuration information	3x L8as_v3	3x L16as_v3	3x L32as_v3
Processor			
Number of processors	1	1	1
Vendor and model	AMD EPYC™ 7763	AMD EPYC 7763	AMD EPYC 7763
Core count (per processor)	64	64	64
Core frequency (GHz)	2.50	2.50	2.50
Hyper-threading	Yes	Yes	Yes
Turbo?	Yes	Yes	Yes
Number of vCPUs per VM	8	16	32
Memory module(s)			
Total memory in system (GB)	64	128	256
NVMe® memory present?	No	No	No
Total memory (DDR+NVMe RAM) (GB)	64	128	256
General hardware			
Storage: NW or Direct Att / Instance	Direct Att	Direct Att	Direct Att
Local storage			
OS			
Number of drives	1	1	1
Drive size (GB)	30	30	30
Drive information	Premium SSD LRS	Premium SSD LRS	Premium SSD LRS
Data drive			
Number of drives	1	2	4
Drive size (TB)	1.92	1.92	1.92
Drive information	NVMe, 400,000 IOPS (instance max: 12,800 IOPS)	NVMe, 400,000 IOPS (instance max: 25,600 IOPS)	NVMe, 400,000 IOPS (instance max: 51,200 IOPS)
Network adapter			
Vendor and model	Mellanox Technologies MT27710 Family	Mellanox Technologies MT27710 Family	Mellanox Technologies MT27710 Family
Number and type of ports	1x 100Gb	1x 100Gb	1x 100Gb

How we tested

Testing overview

We used the cassandra-stress workload to compare Apache Cassandra® database performance of two types of storage-optimized three-node VM clusters:

- Lsv3-series VMs featuring 3rd Gen Intel Xeon Platinum 8370C processors
- Lasv3-series VMs featuring 3rd Gen AMD EPYC 7763 processors

Creating the Cassandra VMs

This section contains the steps we took to create our Cassandra VMs, as well as our client VM for remotely running the cassandra-stress tests.

1. Log into the Azure Portal, and navigate to the Virtual Machines service.
2. To open the Add VM wizard, click Add.
3. On the Basics tab, perform the following steps:
    a. From the appropriate drop-down menu, choose your Subscription.
    b. From the appropriate drop-down menu, choose your Resource group.
    c. Name the Virtual Machine.
    d. From the appropriate drop-down menu, choose your Region. We used South Central US.
    e. Leave the Availability options set to No infrastructure redundancy required.
    f. From the Image drop-down menu, choose Ubuntu Server 20.04 LTS - Gen 2.
    g. Leave Azure Spot instance set to No.
    h. Select the instance size you wish to use. We used Standard_L{8,16,32}s_v3 and Standard_La{8,16,32}s_v3.
    i. Choose a new Username and Password for the Administrator account.
    j. Leave Public inbound ports set to Allow selected ports.
    k. For Select inbound ports, choose SSH (22).
4. On the Disks tab, perform the following steps:
    a. For the OS disk type drop-down menu, choose Premium SSD.
    b. Leave the default Encryption type.
5. On the Networking tab, perform the following steps:
    a. From the appropriate drop-down menu, choose your Virtual network.
    b. To create a new Public IP, choose Create new.
    c. Leave the rest of the settings at their defaults.
6. On the Management tab, perform the following steps:
    a. From the appropriate drop-down menu, choose your Diagnostics storage account.
    b. Leave the rest of the settings at their defaults.
7. On the Advanced tab, leave all defaults.
8. On the Tags tab, add any tags you wish to use.
9. On the Review + create tab, review your settings, and click Create.
10. Repeat steps 1 through 9 to create two more Cassandra VMs and one client VM.

Configuring Ubuntu 20.04 and installing Cassandra on the three Cassandra VMs

1. Log into the each of the Cassandra VMs via SSH.
2. Update the system, and install Java:

```
sudo apt update
sudo apt install openjdk-11-jdk net-tools nmon
```

3. Install Cassandra:

```
echo "deb http://www.apache.org/dist/cassandra/debian 40x main" | sudo tee -a /etc/apt/sources.
list.d/cassandra.sources.list
curl https://downloads.apache.org/cassandra/KEYS | sudo apt-key add -
sudo apt-get update
sudo apt-get install cassandra cassandra-tools
```

4. Create and mount storage for the Cassandra database:

```
sudo mkfs.xfs /dev/nvme0n1
sudo mkdir /data
sudo sh -c "echo '/dev/nvme0n1  /data xfs  defaults 0  0' >> /etc/fstab"
sudo mount -a
```

5. Configure Cassandra:

```
sudo mkdir -p /data/db
sudo mkdir -p /data/log
sudo chown cassandra /data/db
sudo chgrp cassandra /data/db
sudo chown cassandra /data/log
sudo chgrp cassandra /data/log
sudo sed -i 's#/var/lib/cassandra/data#/data/db#g' /etc/cassandra/cassandra.yaml
sudo sed -i 's#/var/lib/cassandra/commitlog#/data/log#g' /etc/cassandra/cassandra.yaml
sudo sed -i 's!# rpc_interface: eth1!rpc_interface: eth0!g' /etc/cassandra/cassandra.yaml
sudo sed -i 's/rpc_address: localhost/#rpc_address: localhost/g' /etc/cassandra/cassandra.yaml
sudo sed -i 's/listen_address: localhost/#listen_address: localhost/g' /etc/cassandra/cassandra.yaml
sudo sed -i 's/# listen_interface: eth0/listen_interface: eth0/g' /etc/cassandra/cassandra.yaml
```

6. Update the seeds IPs in /etc/cassandra/cassandra.yaml with the IP addresses of the three Cassandra VMs.
7. Restart Cassandra service:

```
sudo service cassandra restart
```

Running the tests

We ran a workload by running one instance of cassandra-stress on each Cassandra VM. Each instance writes a portion of a database table.

1. Log into the client VM via SSH.
2. Create a shell script clean_up.sh with the following content:

```
#!/bin/bash
    IPs=(<IPs of the three Cassandra VMs>)
    for host in ${IPs[@]}; do
ssh $host 'nodetool flush ; nodetool cleanup ; sync'
done
cqlsh --request-timeout=120 -e "drop keyspace keyspace1" <IP of the first Cassandra VM>
    for host in ${IPs[@]}; do
ssh $host 'nodetool flush ; sleep 10 ; nodetool cleanup ; nodetool clearsnapshot --all ; nodetool
flush ; nodetool cleanup ; sync'
done
for host in ${IPs[@]}; do
ssh $host 'sudo service cassandra restart'
done
```

3. Create a shell script run_test.sh with the following content:

```
#!/bin/bash
ssh azureuser@<IP of the 1st Cassandra VM> "cassandra-stress write n=100000000 cl=one -mode native
cql3 -schema keyspace="keyspace1" -pop seq=1..100000000 -log file=~/test.log -node <IPs of all the
three Cassandra VMs>" > /dev/null &
ssh azureuser@<IP of the 2nd Cassandra VM> "cassandra-stress write n=100000000 cl=one -mode native
cql3 -schema keyspace="keyspace1" -pop seq=100000001..200000000 -log file=~/test.log -node <IPs of
all the three Cassandra VMs>" > /dev/null &
ssh azureuser@<IP of the 3rd Cassandra VM>  "cassandra-stress write n=100000000 cl=one -mode native
cql3 -schema keyspace="keyspace1" -pop seq=200000001..300000000 -log file=~/test.log -node <IPs of
all the three Cassandra VMs>" > /dev/null &
```

4. Clean the Cassandra database before each test run:

```
./clean_up.sh
```

5. Run the test:

```
./run_test.sh
```

6. Run each test three times, and capture the median performance.

Read the report at https://facts.pt/OJpQ8Gc▶

This project was commissioned by Intel.

Principled Technologies®

Facts matter.®