浪潮云 **Inspur Cloud Information Technology Co., Ltd**

TPC Express Benchmark™ V
Full Disclosure Report

InspurCloud ICP Edge ICP5220A4

running

Inspur Cloud Platform V3.7

TPCx-V Version	2.1.9
Report Edition	First
Report Submitted	August 21, 2024

First Edition - August 2024

Inspur Cloud Information Technology Co., Ltd (InspurCloud)*, the Sponsor of this benchmark test, believes that the information in this document is accurate as of the publication date. The information in this document is subject to change without notice. The Sponsor assumes no responsibility for any errors that may appear in this document.*

The pricing information in this document is believed to accurately reflect the current prices as of the publication date. However, the Sponsor provides no warranty of the pricing information in this document.

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary because of these and other factors. Therefore, the TPC Express Benchmark™ V should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. No warranty of system performance or price/performance is expressed or implied in this report.

InspurCloud and the InspurCloud Logo are trademarks of Inspur Cloud Information Technology Co., Ltd and/or its affiliates in the U.S. and other countries. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between InspurCloud and any other company

TPC Express Benchmark™ V, TPCx-V, and tpsV, are registered certification marks of the Transaction Processing Performance Council.

The InspurCloud products, services or features identified in this document may not yet be available or may not be available in all areas and may be subject to change without notice. Consult your local InspurCloud business contact for information on the products or services available in your area. You can find additional information via InspurCloud's web site at https://en.inspur.com. Actual performance and environmental costs of InspurCloud products will vary depending on individual customer configurations and conditions.

Abstract

InspurCloud conducted the TPC Express Benchmark™ V (TPCx-V) on the InspurCloud ICP Edge ICP5220A4. The software used included Inspur Cloud Platform V3.7. This report provides full disclosure of the methodology and results. All testing was conducted in conformance with the requirements of the TPCx-V Standard Specification, Revision 2.1.9.

The benchmark results are summarized in the follow table.

Hardware	Software	Total System Cost (USD)	tpsV	USD/tpsV	Availability Date
InspurCloud ICP Edge ICP5220A4	Inspur Cloud Platform V3.7	$56,041	4,640.00	$12.08	August 21, 2024

Executive Summary

The Executive Summary follows on the next several pages.

浪潮云	InspurCloud ICP Edge ICP5220A4		TPCx-V 2.1.9 TPC Pricing 2.9.0 Report Date Aug. 21, 2024
Availability Date	TPCx-V Throughput	Price/Performance	Total System Cost
August 21, 2024	**4,640.00 tpsV**	**$12.08 USD / tpsV**	**$56,041 USD**

System Under Test Configuration Overview

Virtualization Software	Guest VM OS	Processor Description	Memory Size
Inspur Cloud Platform V3.7	InLinux 23.12 LTS	AMD EPYC 9754 2.25GHz 128-core	1,536 GB

1x InspurCloud ICP Edge ICP5220A4 with:

- 2x AMD EPYC 9754 2.25GHz 128-core 360W Processor
- 24x 64GB DDR5 RECC 4800B 2R*4
- 2x SSD 960G M.2
- 3x MegaRAID LSI-9560-16i(8G)
- 12x SSD 12.8T U.2PCIe 2.5in D7-P5620
- 1x PCIE 200G 2Port QSFP112(MCX755106AS-HEAT)

浪潮云

InspurCloud ICP Edge ICP5220A4

TPCx-V	2.1.9
TPC Pricing	2.9.0
Report Date	Aug. 21, 2024

Description	Part Number	Key	Unit Price	Qty	Extended Price	3 yr. Maint. Price
Server Hardware						
InspurCloud ICP Edge ICP5220A4	B0.00.00.00012.00	1	$4,149.00	1	$4,149.00	
AMD EPYC 9754 2.25GHz 128-core 360W Processor	C0.20.00.00049.00		$3,320.00	2	$6,640.00	
2U Passive CPU Heat Sink for AMD Socket SP5 Processors	SNK-P0083P		$41.00	2	$82.00	
Middle Cooling Fan for 2U Hyper-S Systems 80x80x38mm 13.5K RPM	FAN-0209L4-1		$28.00	4	$112.00	
64GB DDR5 RECC 4800B 2R*4(M321R8GA0BB0-COKZJ)	C0.50.03.00004.00		$277.00	24	$6,648.00	
1600W redundant single output power supply with inp	PWS-1K63A-1R		$207.00	2	$414.00	
PCIE 200G 2Port QSFP112(MCX755106AS-HEAT)	C0.43.05.00005.00		$1,660.00	1	$1,660.00	
MegaRAID LSI-9560-16i(8G)	C0.44.00.00015.00		$830.00	3	$2,490.00	
Keyboard and Mouse	n/a		$40.00	1	$40.00	
Monitor	n/a		$290.00	1	$290.00	
3-yr 24x7 w/ 4hr Maintenance						(included)
				Subtotal	**$22,525.00**	**$0.00**
Server Storage						
SSD 960G M.2PCe .2(22110)PM983(MZ1LB960HAJQ-000V7)	C0.31.04.00046.00	1	$180.00	2	$360.00	
SSD 12.8T U.2PCIe 2.5in D7-P5620	C0.31.00.00070.00	1	$2,213.00	12	$26,556.00	
				Subtotal	**$26,916.00**	**$0.00**
Server Software						
InLinux 23.12 LTS		1	$0.00	49	$0.00	
InspurCloud Inspur Cloud Platform (ICP) V3.7 Subscription Edition - 3 Years		1	$6,600.00	1	$6,600.00	
				Subtotal	**$6,600.00**	**$0.00**
				Total	**$56,041.00**	**$0.00**

Pricing: 1 = InspurCloud

.

Audited by Doug Johnson, InfoSizing

Three-Year Cost of Ownership:	**$56,041**
TPCx-V Throughput:	**4,640.00**
$ USD/tpsV:	**$12.08**

Prices used in TPC benchmarks reflect the actual prices a customer would pay for a one-time purchase of the stated Line Items. Individually negotiated discounts are not permitted. Special prices based on assumptions about past or future purchases are not permitted. All discounts reflect standard pricing policies for the listed Line Items. For complete details, see the pricing section of the TPC Benchmark Standard. If you find that the stated prices are not available according to these terms, please inform the TPC at pricing@tpc.org. Thank you.

浪潮云

InspurCloud ICP Edge ICP5220A4

TPCx-V	2.1.9
TPC Pricing	2.9.0
Report Date	Aug. 21, 2024

Guest VM Details

Database Manager	VM Memory (Total)	vCPUs (Total)	DB Initial Size	Configured Customers	Active Customers
PostgreSQL 13.12	1,379.91 GiB	696	29,762.7 GB	2,320,000	2,320,000

Transaction Response Times (in seconds)

Transaction Type	Min	Avg	90th%	Max
Broker-Volume	0.000	0.002	0.003	0.045
Customer-Position	0.000	0.002	0.004	0.063
Market-Watch	0.000	0.003	0.005	0.053
Security-Detail	0.001	0.003	0.004	0.088
Trade-Lookup	0.000	0.015	0.019	0.089
Trade-Order	0.001	0.004	0.005	0.084
Trade-Result	0.002	0.005	0.007	0.101
Trade-Status	0.000	0.002	0.003	0.063
Trade-Update	0.003	0.022	0.027	0.084
Data-Maintenance	0.000	0.004	0.008	0.040
Market-Feed	0.000	0.002	0.002	0.030

Transaction Mix

Transaction Type	Transaction Count	Mix Percentage
Broker-Volume	13,287,881	3.900%
Customer-Position	51,107,546	15.000%
Market-Watch	57,922,004	17.000%
Security-Detail	54,514,941	16.000%
Trade-Lookup	30,664,450	9.000%
Trade-Order	34,412,526	10.100%
Trade-Result	34,075,008	10.001%
Trade-Status	61,329,294	18.000%
Trade-Update	3,407,135	1.000%
Data-Maintenance	3,840	N/A
Market-Feed	230,399	N/A
Transaction Total		340,720,785
Measurement Interval		02:00:00
Business Recovery Time		00:04:03
Redundancy Level Details		Redundancy Level 1 (via RAID 10)
Auditor		Doug Johnson, InfoSizing

Table of Contents

Clause 0 – Preamble

0.1 TPC Express Benchmark™ V Overview

The TPC Express Benchmark™ V (TPCx-V) measures the performance of a virtualized server platform under a demanding database workload. It stresses CPU and memory hardware, storage, networking, hypervisor, and the guest operating system. TPCx-V workload is database-centric and models many properties of cloud services, such as multiple VMs running at different load demand levels, and large fluctuations in the load level of each VM. Another unique characteristic of TPCx-V is an elastic workload that varies the load delivered to each of the VMs by as much as 16x, while maintaining a constant load at the host level.

The TPCx-V kit is available from the TPC (See www.tpc.org/tpcx-hs for more information). Users must sign-up and agree to the TPCx-V User Licensing Agreement (ULA) to download the kit. Re-distribution of the kit is prohibited. All related work (such as collaterals, papers, derivatives) must acknowledge the TPC and include TPCx-V copyright. The TPCx-V Kit includes: TPCx-V Specification document, TPCx-V Users Guide documentation, and all software necessary to set up the benchmark environment and execute the benchmark load.

The purpose of TPC benchmarks is to provide relevant, objective performance data to industry users. To achieve that purpose, TPC benchmark specifications require that benchmark tests be implemented with systems, products, technologies and pricing that:

Are generally available to users;

Are relevant to the market segment that the individual TPC benchmark models or represents (e.g., TPCx-V models and represents multiple concurrent operating and application environments running on a platform);

Would plausibly be implemented by a significant number of users in the market segment the benchmark models or represents.

The use of new systems, products, technologies (hardware or software) and pricing is encouraged so long as they meet the requirements above. Specifically prohibited are benchmark systems, products, technologies or pricing (hereafter referred to as "implementations") whose primary purpose is performance optimization of TPC benchmark results without any corresponding applicability to real-world applications and environments. In other words, all "benchmark special" implementations that improve benchmark results but not real-world performance or pricing, are prohibited.

The rules for pricing are included in the TPC Pricing Specification.

Further information is available at www.tpc.org.

Clause 1 – General Items

1.1 Test Sponsor

This benchmark was sponsored by Inspur Cloud Information Technology Co., Ltd.

1.2 Configuration Diagrams

The priced configuration diagram is shown above in the Executive Summary. The measured configuration diagram is shown below in Figure 1-1. In addition, any differences between the priced and the measured configurations are described.

1.2.1 Measured Configuration Diagram

Figure 1-1 Measured Configuration

1.2.2 Differences Between the Priced and the Measured Configurations

The measured configuration included the client system to drive the workload.

1.3 Hardware Setup Steps

Detailed instructions for installing and configuring the hardware used in the System Under Test (SUT) are included in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

1.4 Software Setup Steps

Detailed instructions for installing and configuring the software used in the SUT are included in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

Clause 2 – Database Design, Scaling, & Population

This section provides details of the process used to create the database environment.

2.1 Database Creation Steps

Detailed instructions for creating the database environment used in the SUT are included in the Supporting Files. Also included is the output captured from running setup.sh. Please see the Supporting Files Index for a summary of the files available.

Table 2-1 provides details on the distribution of tables, partitions, and logs across all media.

Disk #	Controller	Drives	Partition	RAID	Size	Use
1	N/A	960 GB M.2	/dev/nvme0n1	RAID1 (software)	894 GiB	Boot
2	N/A	960 GB M.2	/dev/nvme1n1			
3	MegaRAID 1	4x 12.8 TB SSDs	/dev/sda	RAID10	23 TiB	DB Data DB Log
4	MegaRAID 2	4x 12.8 TB SSDs	/dev/sdb	RAID10	23 TiB	
5	MegaRAID 3	4x 12.8 TB SSDs	/dev/sdc	RAID10	23 TiB	

Table 2-1 Distribution of Tables, Partitions, and Logs Across Media

2.2 Database Load Methodology

InspurCloud used the setup.sh script provided with the TPCx-V benchmark kit to load the databases. The necessary data is generated with the required properties and loaded it into the databases. The output from the script is available in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

Clause 3 – Transactions

All transaction implementation details are handled by the TPC's TPCx-V benchmark kit. Therefore, the TPCx-V Standard Specification, Revision 2.1.9 does not have any disclosure requirements for this clause.

Clause 4 – SUT, Driver, & Network

4.1 Network Configuration Description

The client machine used to drive the workload was connected to the SUT by a 200 GbE network as depicted in the Measured Configuration Diagram.

Clause 5 – Benchmark Kit

5.1 Version

InspurCloud used the required TPC-provided benchmark kit for this benchmark. Table 5-1 shows the version of the kit InspurCloud used.

TPCx-V Benchmark Kit Version
2.1.9

Table 5-1 Benchmark Kit Version

5.2 Modifications

The script load_tables.sh had “wait” statements added to control concurrent table population.

Clause 6 – Performance Metrics & Response Times

6.1 VGenDriver Configuration

6.1.1 Customer Emulator (CE)

A TPCx-V Customer Emulator (VCE) process is created by invoking vce.jar. The number of VCE processes is controlled by the configuration parameter NUM_DRIVER_HOSTS in the vcfg.properties file. The number of CE threads used to present the CE load to the SUT is controlled by the configuration parameter NUM_CE_DRIVERS.

Table 6-1 summarizes the configuration of VGenDriverCE used for this benchmark. Additional configuration details can be found in vcfg.properties.

VCE Processes	10
Total CE Threads	180

Table 6-1 VGenDriverCE Configuration

6.1.2 Market Exchange Emulator (MEE)

A TPCx-V Market Exchange Emulator (VMEE) process is created by invoking vmee.jar. The number of VMEE processes is controlled by the configuration parameter NUM_VMEE_PROCESSES in the vcfg.properties file.

Each MEE has one thread pool for handling Trade-Result transactions and another thread pool for handling Market-Feed Transactions. The size of these thread pools is controlled by the configuration parameters MEE_TR_POOL and MEE_MF_POOL, respectively.

Table 6-2 summarizes the configuration of VGenDriverMEE used for this benchmark. Additional configuration details can be found in vcfg.properties.

VMEE Processes	1
MEEs	160
Total Trade-Result Threads	800
Total Market-Feed Threads	160

Table 6-2 VGenDriverMEE Configuration

6.2 Overall Throughput

The TPCx-V Standard Specification:

- Defines Nominal Throughput as 2.00 tpsV per 1,000 Active Customers
- Requires Measured Throughput to be between 80% and 102% of Nominal Throughput
- Sets Reported Throughput to:
  - Measured Throughput when it is less than Nominal Throughput
  - Nominal Throughput when Measured Throughput is between Nominal Throughput and 102% of Nominal Throughput

Table 6-3 summarizes the overall throughput results for this benchmark.

Measured Throughput	4,732.64 tpsV
Reported Throughput	4,640.00 tpsV

Active Customers	2,320,000
80% Nominal	3,712.00 tpsV
Nominal Throughput	4,640.00 tpsV
102% Nominal	4,732.80 tpsV

Table 6-3 Overall Throughput Results & Nominal Throughput Summary

6.3 Measured Throughput by Group

Table 6-4 shows the measured throughput for each Group over the Measurement Interval. The TPCx-V Standard Specification requires each Group's measured throughput to be within 2% of its expected value.

Tile	Group	Expected	tpsV	Delta
1	1	118.31	117.95	-0.30%
1	2	236.63	234.45	-0.92%
1	3	354.94	352.76	-0.61%
1	4	473.26	478.04	1.01%
2	1	118.31	117.94	-0.31%
2	2	236.63	234.47	-0.91%
2	3	354.94	352.76	-0.61%
2	4	473.26	477.95	0.99%
3	1	118.31	117.92	-0.33%
3	2	236.63	234.46	-0.92%
3	3	354.94	352.81	-0.60%
3	4	473.26	477.99	1.00%
4	1	118.31	117.96	-0.30%
4	2	236.63	234.43	-0.93%
4	3	354.94	352.71	-0.63%
4	4	473.26	477.95	0.99%

Table 6-4 Measured Throughput by Group

6.4 Test Run Graph

Figure 6-1 shows the throughput versus elapsed wall clock time for the Trade-Result transaction.

Figure 6-1 Test Run Graph

6.5 Transaction Input Parameter Mix Percentages

Table 6-5 shows the mix percentages over the Measurement Interval for key transaction input parameters.

	Setting	Mix	Required Range		
Customer-Position			Min	Target	Max
By Tax ID	True	50.00%	48.00%	50.00%	52.00%
Get History	True	50.00%	48.00%	50.00%	52.00%
Market-Watch			Min	Target	Max
Security Chosen By	Watch List	60.00%	57.00%	60.00%	63.00%
	Account ID	35.00%	33.00%	35.00%	37.00%
	Industry	5.00%	4.50%	5.00%	5.50%
Security Detail			Min	Target	Max
Access LOB	True	1.00%	0.90%	1.00%	1.10%
Trade-Lookup			Min	Target	Max
Frame to Execute	1	40.00%	38.00%	40.00%	42.00%
	2	30.00%	28.50%	30.00%	31.50%
	3	20.00%	19.00%	20.00%	21.00%
	4	10.00%	9.50%	10.00%	10.50%
Trade-Order			Min	Target	Max
By Third Party	True	10.01%	9.50%	10.00%	10.50%
By Company Name	True	40.00%	38.00%	40.00%	42.00%
Buy On Margin	True	8.00%	7.50%	8.00%	8.50%
Rollback	True	0.99%	0.94%	0.99%	1.04%
LIFO	True	35.00%	33.00%	35.00%	37.00%
Trade Quantity	100	25.00%	24%	25%	26%
	200	25.00%	24%	25%	26%
	400	24.98%	24%	25%	26%
	800	25.01%	24%	25%	26%
Trade Type	Limit Buy	20.00%	19.8%	20%	20.2%
	Limit Sell	10.00%	9.9%	10%	
	Market Buy	30.00%	29.7%	30%	30.3%
	Market Sell	30.00%	29.7%	30%	30.3%
	Stop Loss	9.99%	9.9%	10%	10.1%
Trade-Update			Min	Target	Max
Frame to Execute	1	45.01%	43%	45%	47%
	2	32.98%	31%	33%	35%
	3	22.01%	20%	22%	24%

Table 6-5 Transaction Input Parameter Mix Percentages

Clause 7 – Transaction & System Properties

7.1 Atomicity

The following atomicity tests were conducted on all Tier-B VMs using the xVAudit.Atomicity application provided with the TPCx-V benchmark kit.

- Commit Test
- Rollback Test

The results of these tests are available in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

7.2 Consistency

The following consistency conditions were tested on the initial population of all Tier-B VM databases using the xVAudit.Consistency application provided with the TPCx-V benchmark kit. NOTE: these conditions are all also re-evaluated at the conclusion of the Business Recovery test.

- Consistency Condition 1
- Consistency Condition 2
- Consistency Condition 3

The results of these tests are available in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

7.3 Isolation

The following isolation tests were conducted on all Tier-B VMs using the xVAudit.Isolation applications provided with the TPCx-V benchmark kit.

- P1 Test in Read-Only
- P1 Test in Read-Write
- P2 Test in Read-Write

The results of these tests are available in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

7.4 Data Accessibility

Data Accessibility tests the SUT's ability to maintain database operations with full data access after the permanent irrecoverable failure of any single Durable Medium containing database tables, recovery log data, or database metadata.

7.4.1 Redundancy Level

Table 7-1 shows the redundancy level, as defined in the TPCx-V Standard Specification, provided by the SUT.

Redundancy Level
Level 1 – via RAID 10

Table 7-1 Redundancy Level

7.4.2 Durable Media Technologies

Table 7-2 shows the combinations of Durable Media technologies that were tested. All unique combinations (as defined by the specification) that contained database data or logs were tested.

Contents	Durable Media Type	Bus Type	Array Redundancy	Controller
DB Data and Log	SSD	SAS	RAID10	LSI MegaRAID

Table 7-2 Tested Durable Media Combinations

7.4.3 Test Description

Validation of Redundancy Level 1 was accomplished by performing the following steps.

1) The current number of completed trades, *count1*, was determined.
2) A test run was started using the same configuration as was used in the measured run.
3) The Data Accessibility Throughput Requirements were met for at least 20 minutes.
4) The failure was induced by physically removing a drive that contained both database data and database log. Because the array was RAID protected, the test run continued.
5) After a few minutes, a new drive was inserted into the disk enclosure to replace the failed drive.
6) The array began the necessary recovery process.
7) The test run continued for at least 20 minutes.
8) The test run terminated gracefully.
9) The new number of completed trades, *count2*, was determined.
10) The number of Trade-Results successfully completed (*count2* – *count1*) was verified to be equal to the number of successful Trade-Result transaction reported by the driver.
11) Successful completion of the drive recovery process was confirmed.

7.4.4 Data Accessibility Graph

Figure 7-1 shows the measured throughput versus elapsed time for the Data Accessibility test.

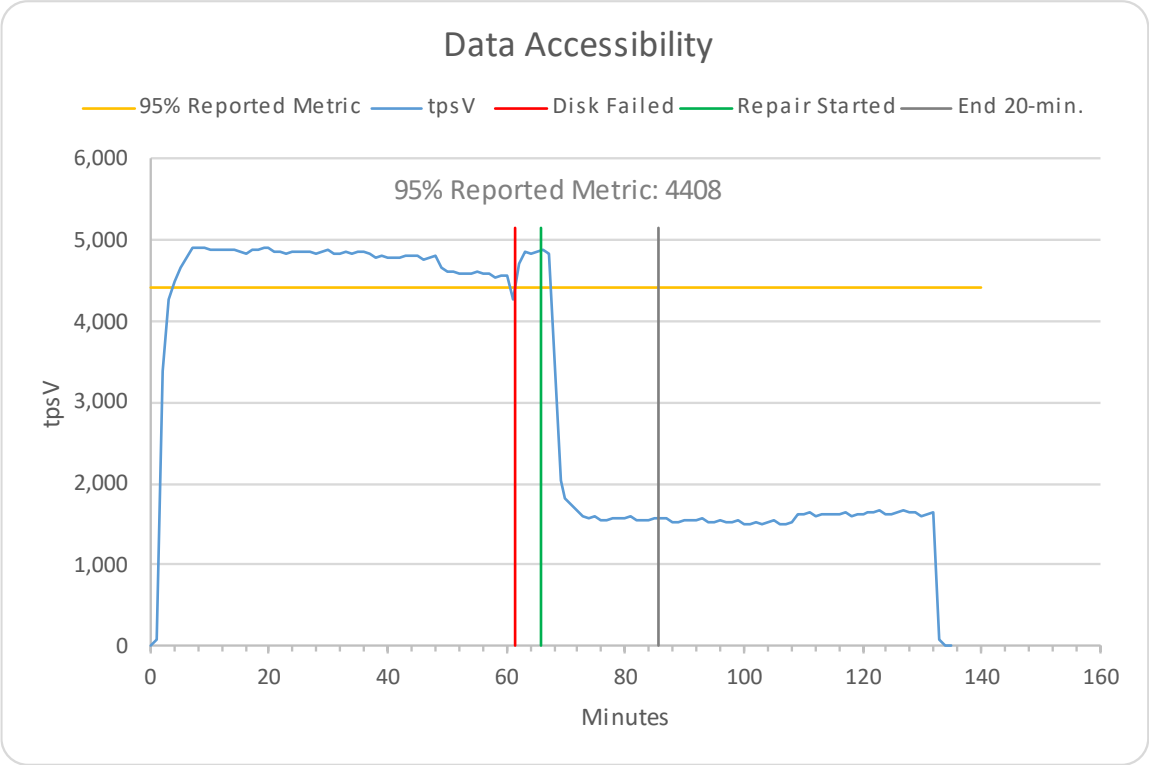

Figure 7-1 Data Accessibility Test Run Graph

7.5 Business Recovery

Business Recovery tests the SUTs ability to recover from a Loss of Processing failure as defined in the TPCx-V Standard Specification and restore certain operational criteria.

7.5.1 Test Description

Business Recovery was evaluated by performing the following steps.

1) The current number of completed trades, *count1*, was determined.
2) A test run was started using the same configuration as was used in the measured run.
3) The Durability Throughput Requirements were met for at least 20 minutes.
4) The failure was induced by instantaneously powering off Tile 1 Group 1 VM 3.
5) The test run was terminated.
6) Tile 1 Group 1 VM 3 was powered back on; Postgres was started and began automatic database recovery. The timestamp in the Postgres log for when the service started is considered the start of Database Recovery. The timestamp in the Postgres log for when the database was ready to accept connections is considered the end of Database Recovery.

7) A test run was started using the same configuration as was used in the measured run. The time when the driver started submitting transactions is considered the start of Application Recovery.
8) The run proceeded until a 20-minute window existed such that the first minute of the window and the entire window both had a tpsV that was at least 95% of the Reported Throughput. The time of the beginning of the window is considered the end of Application Recovery.
9) The test run terminated gracefully, and it was verified that the driver did not report any errors.
10) The new number of completed trades, *count2*, was determined.
11) The number of Trade-Results successfully completed (*count2* – *count1*) was verified to be equal to or greater than the number of successful Trade-Result transaction reported by the driver. In the case of an inequality, it was verified that the difference was less than or equal to the maximum number of Trade-Result transactions that could be simultaneously in-flight from the SUT to the driver.
12) Consistency of all databases was verified.

7.5.2 Business Recovery Times

Table 7-3 summarizes the key times associated with the Business Recovery test.

Event	Elapsed Time
Database Recovery	00:00:03
Application Recovery	00:04:00
Business Recovery	00:04:03

Table 7-3 Business Recovery Test Times

7.5.3 Business Recovery Time Graph

Figure 7-2 shows the measured throughput versus elapsed time for the Business Recovery test.

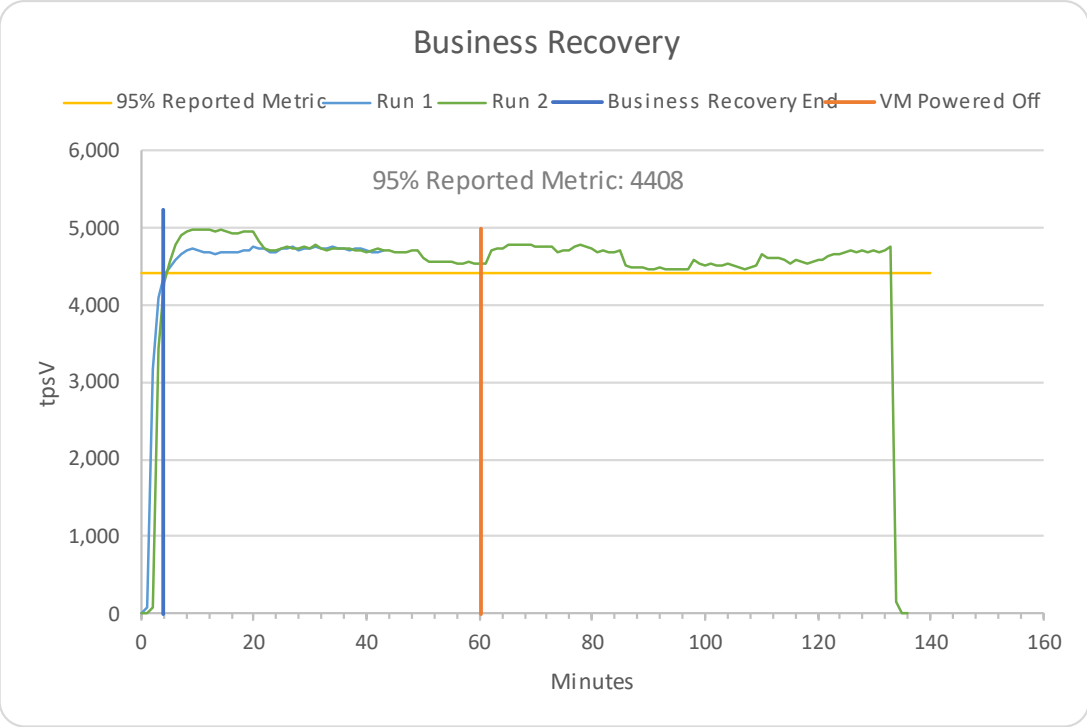

Figure 7-2 Business Recovery Time Graph

Clause 8 – Pricing

8.1 Business Day Space Calculations

To satisfy the requirements in Clauses 5.6.6.4 and 5.6.6.5 of the Standard Specification, it was verified that the file systems containing the database data and database log had at least 10% free space before and after the performance test. Details are available in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

8.2 Pricing Related Metrics

Table 8-1 contains all pricing related metrics. The total solution, as priced, will be generally available on the Availability Date.

Pricing Related Metrics	
Total Price	$56,041
Performance Metric	4,640.00 tpsV
Price/Performance Metric	$12.08 USD/tpsV
Availability Date	August 21, 2024

Table 8-1 Pricing Related Metrics

8.3 Additional Pricing Details

All additional pricing disclosure items, such as line item details and pricing calculations, are included in the Executive Summary.

Letter of Attestation

InfoSizing
The Right Metric For Sizing IT

TPC
Certified Auditor

Jiatong Shen
Inspur Cloud Information Technology Co., Ltd
No.1036 Inspur Road
Jinan City
China

August 18, 2024

I verified the TPC Express Benchmark™ V 2.1.9 performance of the following configuration:

Platform:	InspurCloud ICP Edge ICP5220A4
Virtualization Software	Inspur Cloud Platform V3.7
Guest VM OS:	InLinux 23.12 LTS

The results were:

Performance Metric	**4,640.00 tpsV**
Configured Customers	2,320,000
Active Customers	2,320,000
Tile Count	4

Server	**1x InspurCloud ICP Edge ICP5220A4**		
CPUs	2 x AMD EPYC 9754 2.25 GHz 128-Core Processor		
Memory	1,536 GB		
Storage	***Qty***	***Size***	***Type***
	2	960 GB	M.2 SSD (RAID-1)
	12	12.8 TB	SSD (RAID-10)

In my opinion, these performance results were produced in compliance with the TPC requirements for the benchmark.

The following verification items were given special attention:

- All TPC-provided components were verified to be version 2.1.9
- 1 modification was made to the TPC-provided kit (see Audit Note below)
- All databases were properly scaled and populated
- Each Group contributed the appropriate overall load to the SUT
- The mandatory network between the driver and the SUT was configured
- The ACID properties were met
- Input data was generated according to the specified percentages

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

- All 90% response times were under the specified maximums
- The measurement interval was 120 minutes
- The implementation used Redundancy Level 1
- The Business Recovery Time of 00:04:03 was correctly measured
- The system pricing was verified for major components and maintenance
- The major pages from the FDR were verified for accuracy

Additional Audit Notes:

The script load_tables.sh had "wait" statements added to control concurrent table population.

Respectfully Yours,

Doug Johnson, Certified TPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

Supporting Files Index

Clause	Description
Introduction	Database Tunable Parameters OS Tunable Parameters config.out file, detailing the full VM Configuration Hardware and Software Configuration Driver Software Configuration SUT Software Configuration Driver Hardware Configuration SUT Hardware Configuration
Clause 2	Output of setup.sh
Clause 4	Modified source file
Clause 5	File system space for Database growth
Clause 6	Outputs of ACID applications
Clause 10	VGenDriver Configuration VGenLoader parameters CE VGenLogger Output DM VGenLogger Output MEE VGenLogger Output

Third-Party Price Quotes

All components are available directly through the Test Sponsor **(**Inspur Cloud Information Technology Co., Ltd).

vcfg.properties

This file (included here for easy reference) is also included in the Supporting Files. Please see the Supporting Files Index for a summary of the files available.

```
: '
/*
 * Legal Notice
 *
 * This document and associated source code (the "Work") is a part of a
 * benchmark specification maintained by the TPC.
 *
 * The TPC reserves all right, title, and interest to the Work as provided
 * under U.S. and international laws, including without limitation all patent
 * and trademark rights therein.
 *
 * No Warranty
 *
 * 1.1 TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, THE INFORMATION
 *     CONTAINED HEREIN IS PROVIDED "AS IS" AND WITH ALL FAULTS, AND THE
 *     AUTHORS AND DEVELOPERS OF THE WORK HEREBY DISCLAIM ALL OTHER
 *     WARRANTIES AND CONDITIONS, EITHER EXPRESS, IMPLIED OR STATUTORY,
 *     INCLUDING, BUT NOT LIMITED TO, ANY (IF ANY) IMPLIED WARRANTIES,
 *     DUTIES OR CONDITIONS OF MERCHANTABILITY, OF FITNESS FOR A PARTICULAR
 *     PURPOSE, OF ACCURACY OR COMPLETENESS OF RESPONSES, OF RESULTS, OF
 *     WORKMANLIKE EFFORT, OF LACK OF VIRUSES, AND OF LACK OF NEGLIGENCE.
 *     ALSO, THERE IS NO WARRANTY OR CONDITION OF TITLE, QUIET ENJOYMENT,
 *     QUIET POSSESSION, CORRESPONDENCE TO DESCRIPTION OR NON-INFRINGEMENT
 *     WITH REGARD TO THE WORK.
 * 1.2 IN NO EVENT WILL ANY AUTHOR OR DEVELOPER OF THE WORK BE LIABLE TO
 *     ANY OTHER PARTY FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO THE
 *     COST OF PROCURING SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS
 *     OF USE, LOSS OF DATA, OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT,
 *     INDIRECT, OR SPECIAL DAMAGES WHETHER UNDER CONTRACT, TORT, WARRANTY,
 *     OR OTHERWISE, ARISING IN ANY WAY OUT OF THIS OR ANY OTHER AGREEMENT
 *     RELATING TO THE WORK, WHETHER OR NOT SUCH AUTHOR OR DEVELOPER HAD
 *     ADVANCE NOTICE OF THE POSSIBILITY OF SUCH DAMAGES.
 */
'
################################################################################
#
# VM Configuration
#
# The specification defines 1 to 6 Tiles. Each Tile contans 4 Groups.
# Each Group contains 3 VMs
#
VM_GROUPS = "4"
VM_TILES = "4"

#
################################################################################

################################################################################
#
# Runtime Configuration
#
```

```
# RUN_ITERATION_SEC: the combined runtime for all load phases. This value is
# divided  by the number of phases to determine the run duration for each phase.
#
# For a valid run, RAMPUP_SEC has to be >= 720 seconds. Included in Ramp-up is
# DRIVER_SCALEUP_SEC, which is the time to graudally log in CE threads and
# start submitting transactions. We are at full load after DRIVER_SCALEUP_SEC.
# A 30-60 second DRIVER_SCALEUP_SEC is usually adequate. After transactions
# start executing at full load, it takes 6 minutes for limit-order Trade-Results
# transactions to reach their steady-state throughput. So you want the
# difference between RAMPUP_SEC and DRIVER_SCALEUP_SEC to be at least 6 minutes
#
# DRIVER_RAMPDN_SEC: the number of seconds to ramp down the load at the end
# of the final measurement phase before terminating the run.
RUN_ITERATION_SEC = "7200"
DRIVER_SCALEUP_SEC = "60"
RAMPUP_SEC = "720"
DRIVER_RAMPDN_SEC = "60"

# USE_60_SEC_POLLING_INTERVAL and VCE_POLLING_INTERVALS_PER_PHASE are mutually
# exclusive means of controlling the frequency of stats polling during each
# phase of the run. If a 60-sec polling interval is desired, simply set
# "USE_60_SEC_POLLING_INTERVAL = 1" and make sure the phase duration is a
# multiple of 60 seconds. If another polling interval is desired, set
# "USE_60_SEC_POLLING_INTERVAL = 0" and set VCE_POLLING_INTERVALS_PER_PHASE to
# the number of polling intervals desired in each phase. (For example, for the
# default 12-minute phase duration, setting VCE_POLLING_INTERVALS_PER_PHASE = 12
# will result in 60-second polling, the same as setting
# "USE_60_SEC_POLLING_INTERVAL = 1". Setting
# "VCE_POLLING_INTERVALS_PER_PHASE = 24" would result in 30-sec polling.)
# The value of VCE_POLLING_INTERVALS_PER_PHASE is ignored when
# USE_60_SEC_POLLING_INTERVAL is set to 1. Comment out both for no polling.
USE_60_SEC_POLLING_INTERVAL = 1
VCE_POLLING_INTERVALS_PER_PHASE = 12

# NUM_RUN_ITERATIONS: the number of times to run a full set of all load phases
# NUM_RUN_PHASES: the number of load phases in a single run iteration
NUM_RUN_ITERATIONS = "1"
NUM_RUN_PHASES = "10"

#
################################################################################

################################################################################
#
# VDriver Configuration
#
# VDriver (prime) hostname and RMI listening port
VDRIVER_RMI_HOST = "drivervm"
VDRIVER_RMI_PORT = "30000"
#
# Script for executing load balancing during RAMP_UP (comment out or leave as empty
# string if not using one)
# You can append any options to the script to the end of this string
# LOAD_BAL_SCRIPT = /opt/VDriver/scripts/rhel6/activate_load_balancing.sh"
#
################################################################################

################################################################################
```

```
#
# VCe Configuration
#
# NUM_DRIVER_HOSTS: the number of CE *processes* (i.e. how many invocations of
# vce.jar) that you want to drive load against the SUT.  A value of 1 usually
# suffices, unless you need to drive the load from multiple driver systems
#NUM_DRIVER_HOSTS = "5"
# by sjt
NUM_DRIVER_HOSTS = "10"

# Default and index-specific VCe driver hostnames and ports for RMI
# communication between processes (These let the VDriver process know where to
# contact the VCE processes to send benchmark control commands). There must be
# one host/port pair combination for each NUM_DRIVER_HOSTS (additional entries
# are ignored).
VCE_RMI_HOST[] = "drivervm"
VCE_RMI_PORT[] = "30100"

# Indexes for VCE start from 1
VCE_RMI_PORT[1] = "30100"
VCE_RMI_PORT[2] = "30101"
VCE_RMI_PORT[3] = "30102"
VCE_RMI_PORT[4] = "30103"
VCE_RMI_PORT[5] = "30105"
VCE_RMI_PORT[6] = "30106"
VCE_RMI_PORT[7] = "30107"
VCE_RMI_PORT[8] = "30108"
VCE_RMI_PORT[9] = "30109"
VCE_RMI_PORT[10] = "30110"
VCE_RMI_PORT[11] = "30111"
VCE_RMI_PORT[12] = "30112"
VCE_RMI_PORT[13] = "30113"
VCE_RMI_PORT[14] = "30114"
VCE_RMI_PORT[15] = "30115"
VCE_RMI_PORT[16] = "30116"
VCE_RMI_PORT[17] = "30117"
VCE_RMI_PORT[18] = "30118"
VCE_RMI_PORT[19] = "30119"
VCE_RMI_PORT[20] = "30120"

# NUM_CE_DRIVERS: the total number of CE threads that you want to drive load
# against the SUT VMs. If you are using multiple DRIVER_HOSTS, you can specify
# the  number of CEs to start on each host by using the indexed version of this
# key. Otherwise, the CEs per host are distributed evenly between hosts.
#NUM_CE_DRIVERS[] = "400"
# by sjt
NUM_CE_DRIVERS[] = "180"

# Indexed version. Index values start from 1
#NUM_CE_DRIVERS[1] = "2"

# Debugging property; when client thread waits longer than
# CONN_WAIT_DELAY_MSEC_THRESHOLD msec to get a connection to a SUT VM, it will
# print to the console how long the wait was
# CONN_WAIT_DELAY_MSEC_THRESHOLD = "1000"

################################################################################
```

```
##########################################################################################
#
# VMEE Configuration
#
# The number of VMEE processes the VDriver should talk to. Each VMEE spawns
# a number of "mee" threads, each of which is dedicated to a single
# Tile/Group/vconnector process
# Typically, a single VMEE process on a single system is enough, but you can
# run multiple processes, but we have to run them on the same host
NUM_VMEE_PROCESSES = "1"

# These settings specify the host name and port number a given VMEE is
# listening on. vDriver will use these to connect to the VMEE processes. If
# starting the VMEE processes manually (i.e. not using the provided script),
# the values specified here must match those used on the VMEE command line
# (-rh and -rp) when starting a given VMEE process.
#
# Unindexed value - used as a default if a given indexed value is not specified.
VMEE_RMI_HOST[] = "drivervm"
VMEE_RMI_PORT[] = "30200"
#
# Indexed values (1 to (NUM_VMEE_PROCESSES) will be used if they exist).
#VMEE_RMI_HOST[1] = "drivervm"
VMEE_RMI_PORT[1] = "30200"
VMEE_RMI_PORT[2] = "30201"
VMEE_RMI_PORT[3] = "30202"
VMEE_RMI_PORT[4] = "30203"
VMEE_RMI_PORT[5] = "30204"


# These settings specify individual MEE configuration options. The MEE
# threads are divided between the VMEE processes.  There is a 1-1
# mapping between vconnector processes on Tier A VMs and MEEs. The
# VMEE process will have one MEE for each vconnector process
#
# MEE_TXN_HOST - must match VMEE_RMI_HOST; host name the MEE will listen on
#              (for connections from SUT SendToMarket in a vconnector process)
# MEE_TXN_PORT - port number the MEE will listen on (for connections from SUT
#              SendToMarket in a vconnector process)
# MEE_MF_POOL - Size of the Market-Feed thread pool (should be 1 for TPCx-V)
# MEE_TR_POOL - Size of the Trade-Result thread pool (adjust this based on load)
#
# The indexes used for these parameters are [tile][group][vconn], indicating
# the vconnector (index) in a given group on a given tile that the MEE is
# connected to.
#
# Unindexed value - used as a default if a given indexed value is not specified.
MEE_TXN_HOST[] = "drivervm"
MEE_TXN_PORT[] = "30300"
MEE_MF_POOL[] = "1"
MEE_TR_POOL[] = "5"
#
# (Indexed values will be used if they exist. Add more entries for additional
# tiles.)
#
# Tile 1 Group 1
# MEE_TXN_HOST[1][1] = "drivervm"
# MEE_TXN_PORT[1][1][1] = "31101"
```

```
# Tile 1 Group 2
# MEE_TXN_HOST[1][2] = "drivervm"
# MEE_TXN_PORT[1][2][1] = "31201"
# Tile 1 Group 3
# MEE_TXN_HOST[1][3] = "drivervm"
# MEE_TXN_PORT[1][3][1] = "31301"
# Tile 1 Group 4
# MEE_TXN_HOST[1][4] = "drivervm"
# MEE_TXN_PORT[1][4][1] = "31401"


##########################################################################################

##########################################################################################
#
# VConnector Configuration
#
# VConnector is the process on the Tier A VM1 that receives transactions from
# the CE and MEE drivers, and submits them to the VM2 and VM3 databases

#
# Number of times to retry a failed DB transaction before reporting failure
NUM_TXN_RETRIES = "25"

# The "vconnector" is the process on the Tier A VM (VM1) that receives
# transactions from the driver and submits them to the database. There can be
# be one or more vconnector processes on each Tier A. NUM_VCONN_PER_GROUP
# is the number of VConnector processes running on each Tier A VM (The
# requests will be distributed across all of these processes). Each process
# is multi-threaded, and one process may be enough. But if you see odbc
# contention issues on the Tier A VM1, increase this value
NUM_VCONN_PER_GROUP = "10"

# Default VConnector hostnames and ports
VCONN_RMI_HOST[] = "vm1"
VCONN_RMI_PORT[] = "33000"
VCONN_TXN_HOST[] = "vm1"
VCONN_TXN_PORT[] = "34000"
# The common case is to set an unindexed CONN_DSN_LABELS[]  = "PSQL2,PSQL3
# and VCONN_NUM_DBS[] = "2" to cover the whole SUT
VCONN_DSN_LABELS[]  = "PSQL2,PSQL3"
VCONN_NUM_DBS[] = "2"

# Index-specific hostnames and ports. Add more entries for additional tiles.
# All host/port entries are of the form VCONN_RMI_HOST[tile][group][index]
# The harness will automatically increment "index" if there are multiple
# VConnector processes per group (i.e. NUM_VCONN_PER_GROUP > 1) unless values
# for every tile/group/index are specified here. So the options for specifying
# these values are:
#
#   To automatically increment port numbers for multiple VConnector processes:
#
#   VCONN_RMI_HOST[1][1] = "vm1"
#   VCONN_RMI_PORT[1][1][] = "42000"    (VCONN_RMI_PORT[1][1][1] = "42000",
#                                VCONN_RMI_PORT[1][1][2] = "42001", ...)
#   VCONN_TXN_HOST[1][1] = "vm1"
#   VCONN_TXN_PORT[1][1][] = "44000"    (VCONN_TXN_PORT[1][1][1] = "44000",
#                                VCONN_TXN_PORT[1][1][2] = "44001", ...)
```

```
#   Or, in the case of 3 VConnector processes per group, to specifically assign
#   values for each port (in this example, for Tile 1 Group 1):
#   VCONN_RMI_HOST[1][1] = "vm1"
#   VCONN_RMI_PORT[1][1][1] = "51100"
#   VCONN_RMI_PORT[1][1][2] = "32109"
#   VCONN_RMI_PORT[1][1][3] = "25432"
#   VCONN_TXN_HOST[1][1] = "vm1"
#   VCONN_TXN_PORT[1][1][1] = "41100"
#   VCONN_TXN_PORT[1][1][2] = "11243"
#   VCONN_TXN_PORT[1][1][3] = "27211"
#

VCONN_RMI_HOST[1][1] = "vm1"
VCONN_TXN_HOST[1][1] = "vm1"
VCONN_RMI_HOST[1][2] = "vm4"
VCONN_TXN_HOST[1][2] = "vm4"
VCONN_RMI_HOST[1][3] = "vm7"
VCONN_TXN_HOST[1][3] = "vm7"
VCONN_RMI_HOST[1][4] = "vm10"
VCONN_TXN_HOST[1][4] = "vm10"
VCONN_RMI_HOST[2][1] = "vm13"
VCONN_TXN_HOST[2][1] = "vm13"
VCONN_RMI_HOST[2][2] = "vm16"
VCONN_TXN_HOST[2][2] = "vm16"
VCONN_RMI_HOST[2][3] = "vm19"
VCONN_TXN_HOST[2][3] = "vm19"
VCONN_RMI_HOST[2][4] = "vm22"
VCONN_TXN_HOST[2][4] = "vm22"
VCONN_RMI_HOST[3][1] = "vm25"
VCONN_TXN_HOST[3][1] = "vm25"
VCONN_RMI_HOST[3][2] = "vm28"
VCONN_TXN_HOST[3][2] = "vm28"
VCONN_RMI_HOST[3][3] = "vm31"
VCONN_TXN_HOST[3][3] = "vm31"
VCONN_RMI_HOST[3][4] = "vm34"
VCONN_TXN_HOST[3][4] = "vm34"
VCONN_RMI_HOST[4][1] = "vm37"
VCONN_TXN_HOST[4][1] = "vm37"
VCONN_RMI_HOST[4][2] = "vm40"
VCONN_TXN_HOST[4][2] = "vm40"
VCONN_RMI_HOST[4][3] = "vm43"
VCONN_TXN_HOST[4][3] = "vm43"
VCONN_RMI_HOST[4][4] = "vm46"
VCONN_TXN_HOST[4][4] = "vm46"


#
#
################################################################################

################################################################################
#
# VDm Configuration
#
# VDm hostname and RMI listening port
VDM_RMI_HOST = "drivervm"
VDM_RMI_PORT = "30001"
#
```

```
# The Data-Maintenance transaction is supposed to run once every 60 seconds
VDM_REQ_INTERVAL_SEC = "60"
#
##############################################################################

##############################################################################
#
# Group-specific Load Configuration
#
# Set CUST_CONFIGURED and CUST_ACTIVE for each Tile/Group with the index
# parameters below. SCALE_FACTOR and INIT_TRADE_DAYS are not typically
# changed from their defaults; the unindexed parameters should suffice
CUST_CONFIGURED[] = "5000"
CUST_ACTIVE[] = "5000"
SCALE_FACTOR[] = "500"
INIT_TRADE_DAYS[] = "125"

# Group-specific values
CUST_CONFIGURED[1] = "58000"
CUST_ACTIVE[1] = "58000"
#
CUST_CONFIGURED[2] = "116000"
CUST_ACTIVE[2] = "116000"
#
CUST_CONFIGURED[3] = "174000"
CUST_ACTIVE[3] = "174000"
#
CUST_CONFIGURED[4] = "232000"
CUST_ACTIVE[4] = "232000"
#
#GROUP_PCT_DIST_PHASE[1] = "1.0"
GROUP_PCT_DIST_PHASE[1] = "0.10,0.20,0.30,0.40"
GROUP_PCT_DIST_PHASE[2] = "0.05,0.10,0.25,0.60"
GROUP_PCT_DIST_PHASE[3] = "0.10,0.05,0.20,0.65"
GROUP_PCT_DIST_PHASE[4] = "0.05,0.10,0.05,0.80"
GROUP_PCT_DIST_PHASE[5] = "0.10,0.05,0.30,0.55"
GROUP_PCT_DIST_PHASE[6] = "0.05,0.35,0.20,0.40"
GROUP_PCT_DIST_PHASE[7] = "0.35,0.25,0.15,0.25"
GROUP_PCT_DIST_PHASE[8] = "0.05,0.65,0.20,0.10"
GROUP_PCT_DIST_PHASE[9] = "0.10,0.15,0.70,0.05"
GROUP_PCT_DIST_PHASE[10] = "0.05,0.10,0.65,0.20"


# Use DB_CONN_BUFFER_PCT_GROUP to modify the initial number of connections
# opened by the CEs to each Tier A VM for each group (the index value indicates
# the group number). Use values greater than 1.0 to increase the number of
# connections (up to the theoretical maximum) and values less than 1.0 to
# decrease the number of initial connections.
DB_CONN_BUFFER_PCT_GROUP[1] = "1.5"
DB_CONN_BUFFER_PCT_GROUP[2] = "1.5"
DB_CONN_BUFFER_PCT_GROUP[3] = "1.5"
DB_CONN_BUFFER_PCT_GROUP[4] = "1.5"

#
##############################################################################

##############################################################################
#
```

```
# Misc Configuration Parameters
# These values are unlikely to need to be modified
#

# Log names:
# CE log file names
CE_MIX_LOG = "CE_Mix.log"
CE_ERR_LOG = "CE_Error.log"

# MEE base file names for logging purposes.
MEE_LOG = "MEE_Msg"
MEE_MIX_LOG = "MEE_Mix"
MEE_ERR_LOG = "MEE_Err"

# VDm log file names
VDM_TRANSACTION_LOG = "DM_Txn"
VDM_MESSAGE_LOG = "DM_Msg"

RESULT_DIR = "results"
LOG_DIR = "."
SORT_MIX_LOGS = "0"
SORTED_LOG_NAME_APPEND = "sorted"
LOG_SAMPLE_SEC = "60"
# VGEN_INPUT_FILE_DIR = ""
DEBUG_LEVEL = "0"
SUPPRESS_WARNINGS = "1"
CHECK_TIME_SYNC = "0"
COLLECT_CLIENT_LOGS = "1"

TIME_SYNC_TOLERANCE_MSEC = "1000"

# CE_EXIT_DELAY_SEC is the number of seconds the user wants to wait to allow
# "cleanup" before final exit. This is mostly in case there are "retries" going
# on that need to have time to time out before a final exit.
CE_EXIT_DELAY_SEC = "10"

# NUM_TXN_METRICS is the number of metrics created for report purposes
NUM_TXN_METRICS = "5"
NUM_TXN_TYPES = "12"

CE_MIX_PARAM_INDEX = "1,2"
# BrokerVolumeMixLevel,CustomerPositionMixLevel,
# MarketWatchMixLevel,SecurityDetailMixLevel,
# TradeLookupMixLevel,TradeOrderMixLevel,
# TradeStatusMixLevel,TradeUpdateMixLevel
#CE_MIX_PARAM_1 = "0,0,0,0,0,1000,0,0"
CE_MIX_PARAM_1 = "39,150,170,160,90,101,180,10"
# CE_MIX_PARAM_2 = "59,130,180,140,80,101,190,20"
# TXN_TYPE
# "-1" = EGEN-GENERATED MIX
# "0" = SECURITY_DETAIL
# "1" = BROKER_VOLUME
# "2" = CUSTOMER_POSITION
# "3" = MARKET_WATCH
# "4" = TRADE_STATUS
# "5" = TRADE_LOOKUP
# "6" = TRADE_ORDER
# "7" = TRADE_UPDATE
```