



# **Non-Volatile Memory Express Performance with Linux on IBM LinuxONE Emperor II and LinuxONE Rockhopper II**

## **Results of Use Cases**

© Copyright IBM Corp. 2018

[www.ibm.com](http://www.ibm.com)

This document is provided “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranty of merchantability or fitness for a particular purpose. This document is intended for informational purposes only. It could include technical inaccuracies or typographical errors. The information herein and any conclusions drawn from it are subject to change without notice. Many factors have contributed to the results described herein and IBM does not guarantee comparable results. Performance numbers will vary greatly depending upon system configuration. All data in this document pertains only to the specific test configuration and specific releases of the software described.

## Table of Contents

About this publication.....	3
Remarks .....	3
Introduction .....	4
Objectives of this white paper .....	5
Summary.....	6
NVMe Protocol Overview.....	8
LinuxONE disk I/O general flow description .....	9
Description of the system under test (SUT).....	11
Hardware and software configuration for the SUT .....	12
IBM LinuxONE Server.....	12
PostgreSQL test.....	14
Db2 Warehouse test .....	15
PostgreSQL performance test .....	16
Db2 Warehouse performance test .....	18
Case study results.....	20
Transaction throughput.....	20
Observations .....	20
Conclusion .....	21
Price-Performance .....	22
Disclaimer.....	22
Observations .....	23
Conclusion.....	24
References .....	24
Notices .....	25
Trademarks.....	28
Terms and Conditions .....	29

## **About this publication**

This White Paper examines the setup, tuning, and performance aspects of direct attached Non-Volatile Memory Express (NVMe) storage on IBM LinuxONE Emperor II and IBM LinuxONE Rockhopper II using Linux on IBM Z.

### **Authors**

IBM Linux on IBM Z Performance Evaluation team:

**Martin Kammerer**

IBM Systems Performance Linux on IBM Z team:

**Dr. Muhammad Usman Karim Khan, Dr. Stefan Wirag**

### **Remarks**

The web-links referred in this paper are up-to-date as of October, 2018.

## Introduction

Introductory information is provided in this topic about the setup, tuning, and performance aspects of of direct attached Non-Volatile Memory Express (NVMe) storage on IBM LinuxONE™ Emperor™ II and IBM LinuxONE Rockhopper™ II using Linux® on IBM Z®.

More information about IBM Emperor II and Rockhopper II can be found at

<https://www.ibm.com/it-infrastructure/linuxone>

<https://www.ibm.com/us-en/marketplace/linuxone-emperor-ii>

<https://www.ibm.com/us-en/marketplace/linuxone-rockhopper-ii>

The latest NVMe specifications can be found at

<https://nvmexpress.org/resources/specifications>

## ***Objectives of this white paper***

The high performance characteristics of NVMe storage have the potential to improve response times and throughput of I/O intense use cases.

The expectation is that in servers where NVMe storage is large enough to contain all data NVMe will be a fast and inexpensive storage. This could include applications like MongoDB, ScyllaDB, Apache Cassandra, IBM Db2® Warehouse, Apache Spark, Apache Hadoop, PostgreSQL or MariaDB and Docker containers and KVM host/guests. In case the NVMe storage is too small to contain all data NVMe will accelerate functions like Db2 Warehouse external sort, Oracle Database cache devices, Apache Hadoop scratch disk and in general database specifics like log file writes or cold start.

This paper will demonstrate the advanced performance when using NVMe storage with a set of performance comparisons for certain use cases.

The selected use cases are

- PostgreSQL database workload
- Db2 Warehouse database workload

Performance criteria of I/O throughput metrics are compared when running the workloads and the obtained results were used to calculate Price-Performance numbers. Measurements are provided for each use case individually.

## **Summary**

The performance evaluation of database using direct-attached NVMe SSD storage versus external storage shows that it depends on the database disk I/O characteristic, the external storage type, and the number of connections to the external storage whether the workload performance improves using NVMe storage or not.

Measuring the throughput of PostgreSQL 10 with the YCSB benchmark on LinuxONE Emperor II using 4 direct-attached NVMe SSDs in RAID 10 setup versus an IBM FlashSystem™ 900 storage server connected by 2 FCP Express32S adapters, we obtain similar transactions processed per second (throughput). Since the YCSB benchmark is emitting a continuous stream of queries against the database, the PostgreSQL database also creates continuous disk I/O requests, which are processed by the I/O subsystem and the storage in parallel. Hence, PostgreSQL throughput depends upon the number of I/O requests that the storage can process at a time and therefore the number of direct-attached NVMe SSDs and the used RAID setup for them versus the number of used I/O adapters and connections to the storage server. Therefore, we cannot say that direct-attached NVMe SSDs always have a performance advantage compared to external storage and vice versa for a workload with a steady disk I/O.

Measuring the throughput of Db2 Warehouse 2.11 with the Big-Data-Insights (BDI) benchmark, which is TPC-DS like, we obtain similar throughput with sufficient memory (i.e., very low disk access) using 4 direct-attached NVMe SSDs in a RAID 10 setup compared to using an IBM Storwize® V7000 storage server with HDD storage. However, when the memory size is constrained, the underlying database storage must be accessed more often and therefore, we notice up to 2x better throughput. The performance advantage with NVMe storage when being memory-constrained results from the workload disk I/O characteristics. In the benchmark scenario, Db2 Warehouse had I/O requests in burst. Due to the lower latency of the NVMe SSDs, these I/O requests could be served quicker by the NVMe SSDs than the V7K storage server. Therefore, we can say that direct-attached NVMe SSDs have a performance advantage for workloads with I/O requests in burst.

The price-performance evaluation for the PostgreSQL and the Db2 Warehouse scenarios show that direct-attached NVMe SSDs have a compelling price-performance advantage when comparing the purchase costs of NVMe SSDs with the IBM Adapter for NVMe versus the purchase costs of the external storage servers with the same amount of raw storage capacity.

For the PostgreSQL scenario the price-performance advantage of direct-attached NVMe SSD storage is up to 92% versus an IBM FlashSystem 900 storage server attached by 2 FCP Express32S adapters.

For the Db2 Warehouse scenario the price-performance advantage of direct-attached NVMe SSD storage is up to 79% with sufficient memory and up to 90% with constrained memory versus a IBM Storwize V7000 storage server with HDD storage.

Bottom line, for database workloads you can expect a price-performance advantage using direct-attached NVMe SSDs storage compared to an externally attached storage server. For database workloads you can expect a performance advantage if the workload has I/O requests in burst. If the database workload has a steady disk I/O it depends on the number of direct-attached NVMe SSDs and the used RAID level versus the used storage server type and the number of used I/O adapters whether the NVMe storage or the external storage provides a better workload throughput.

## **NVMe Protocol Overview**

The NVMe protocol is an open collection of standards and interfaces that fully exposes the benefits of non-volatile memory in all types of computing environments, from mobile to data center. It is designed to deliver high bandwidth and low latency storage access.

IBM has a long history of commitment to industry standards, and the NVM Express organization is no exception. As a member of the NVMe workgroup, IBM aims to contribute to the standards and to lead the enablement of enterprise storage systems that use the new drives, network protocols, and I/O architectures.

The NVM Express consortium develops the NVM Express specification and describes how NVMe differentiates itself from other interfaces, such as SATA and SAS, as follows:

- *NVM Express (NVMe) is an optimized, high-performance scalable host controller interface designed to address the needs of Enterprise and Client systems that utilize PCI Express-based solid-state storage. Designed to move beyond the dark ages of hard disk drive technology, NVMe is built from the ground up for non-volatile memory (NVM) technologies. NVMe is designed to provide efficient access to storage devices built with non-volatile memory, from today's NAND flash technology to future, higher-performing, persistent memory technologies.*
- *There are several performance vectors that NVMe addresses, including bandwidth, IOPs, and latency. For example, the maximum IOPs possible for Serial ATA was 200,000, whereas NVMe devices have already been demonstrated to exceed 1,000,000 IOPs. By supporting PCI Express and Fabrics such as RDMA and Fibre Channel, NVM Express can support much higher bandwidths than SATA or SAS (e.g., a PCI Express Gen3 x4 delivers 4 GB/s). Finally, next generation memory technologies may have read access latency under a microsecond, requiring a streamlined protocol that enables an end-to-end latency of under 10 microseconds, including the software stack.*
- *NVMe is a completely new architecture for storage, from the software stack to the hardware devices and systems.*

## Linux disk I/O general flow description

For Linux NVMe storage adds a completely new way of executing disk I/O to the existing FICON®/ECKD™ and FCP/SCSI protocols. *Figure 1.1* shows an overview of Linux disk I/O flow.




Figure 1: Overview of Linux disk I/O flow

Disk I/O in Linux is processed in these steps

- The application program is issuing reads and writes (system calls) and the Linux block device layer is entered
- The Virtual File System (VFS) dispatches requests to different devices and translates the data buffers to sector addressing
- The Logical Volume Manager LVM is used to define the physical to logical device relation
- Multipath sets the multipath policies
- The device mapper (dm) holds the generic mapping of all block devices and performs 1:n mapping for logical volumes and/or multipath devices
- The page cache contains all file I/O data
- Direct I/O bypasses the page cache
- The I/O schedulers merge, order and queue requests, start device drivers via (un)plug device
- The appropriate device driver communicated with the hardware layer depending on the device type

- Direct Access Storage Devices (dasd) are handled by the Linux specific Linux device driver, using the FICON/ECKD protocol
- Small Computer System Interface (SCSI) devices are handled by the Linux common code SCSI device driver, followed by the Linux specific Linux drivers Z Fiber Channel Protocol (zFCP) for wrapping SCSI with FCP and Queued Direct I/O (qdio) for transferring the data to the FICON Express card, using the FCP/SCSI protocol
- NVMe SSD devices are handled by the Linux common code NVMe device driver, using the NVMe protocol

## Description of the system under test (SUT)

The SUT was based on a Linux system executing in LPAR with access to direct-attached NVMe SSD storage, FICON connections to storage servers and network connections to another LPAR hosting workload drivers. **Note:** The selection and purchase of the SSDs used in this feature is the responsibility of the client. IBM can only comment on the reliability and performance of SSDs that have been tested. The client assumes all risk in, and IBM is not responsible for, the use of SSDs as the functionality and/or performance may vary.

The following NVMe Solid State Drives (SSDs) have been tested by IBM:

- Intel® PN SSDPE2KX010T801 (1TB)
- Intel PN SSDPE2KX040T801 (4TB)

In this documentation we refer to them as “1TB NVMe SSD” and “4TB NVMe SSD”.

In our case studies we used 4x 4TB NVMe SSDs, each plugged into an IBM Adapter for NVMe (Feature Code=0435).

Although NVMe support is included in Linux since kernel level 3.19 the Linux distributions for LinuxONE did not automatically build the NVMe driver because the hardware was not available. In all recent SUSE, RedHat, and Ubuntu distributions this step will be included. NVMe SSD storage serves as block device for all kind of data including dump and swap. IPL from NVMe SSD storage is not possible with the current hardware support.

## ***Hardware and software configuration for the SUT***

This topic describes the hardware and software that was used for this case study.

### **IBM LinuxONE Server**

The case study involved the configuration of LPARs in an IBM LinuxONE Emperor II, Model 3906.

To determine if the NVMe SSD devices are recognized by Linux, run the following command:

```
$> lspci
0000:00:00.0 Non-Volatile memory controller: Intel Corporation Express Flash NVMe P4510
0001:00:00.0 Non-Volatile memory controller: Intel Corporation Express Flash NVMe P4510
0002:00:00.0 Non-Volatile memory controller: Intel Corporation Express Flash NVMe P4510
0005:00:00.0 Non-Volatile memory controller: Intel Corporation Express Flash NVMe P4510
```

*Figure 2: 4 NVMe drives are recognized on this system and can be used.*

Afterwards the parted command is used in Linux and create partitions. Once the partitions are created, one can use the NVMe SSDs as block storage devices, or use software-based RAID to create a RAID configuration. Below, we show an excerpt of creating RAID-10 based upon the 4 NVMe SSD drives we have on the SUT:

```
$> apt install mdadm
$> mdadm --create /dev/md0 --chunk=256 --level=10 \
--raid-devices=4 /dev/nvme0n1p1 /dev/nvme1n1p1 \
/dev/nvme2n1p1 /dev/nvme3n1p1 --verbose
$> mkfs.xfs /dev/md0
$> mount /dev/md0 /nvmeData
```

*Figure 3: first install mdadm and then create RAID-10 based upon the 4 NVMe drives.*

Afterwards the mounts can be verified with the usual commands, for example:

```
$> lsblk
NAME            MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
dasda           94:0    0   45G  0 disk
|-dasda1        94:1    0  201M  0 part  /boot/zip1
|-dasda2        94:2    0    2G  0 part  [SWAP]
`-dasda3        94:3    0 39.1G  0 part  /
nvme0n1         259:0    0   3.7T  0 disk
`-nvme0n1p1     259:5    0   3.7T  0 part
  `--md0         9:0     0   7.3T  0 raid10 /nvmeData
nvme2n1         259:1    0   3.7T  0 disk
`-nvme2n1p1     259:7    0   3.7T  0 part
  `--md0         9:0     0   7.3T  0 raid10 /nvmeData
nvme3n1         259:2    0   3.7T  0 disk
`-nvme3n1p1     259:8    0   3.7T  0 part
  `--md0         9:0     0   7.3T  0 raid10 /nvmeData
nvme1n1         259:3    0   3.7T  0 disk
`-nvme1n1p1     259:6    0   3.7T  0 part
  `--md0         9:0     0   7.3T  0 raid10 /nvmeData
```

Figure 4: 4x4TB devices verified by the output of the lsblk command.

## PostgreSQL test

The LPAR setup consisted of:

- 30 dedicated cores processors
- 16 GB storage
- 1x 25GbE RoCE Express2 card
- 2x FCP Express32S cards supporting FCP devices
- 4x IBM Adapter for NVMe
- 4x 4TB NVMe SSDs

Storage subsystem setup:

- IBM FlashSystem 900

The software used was:

- Ubuntu 18.04.1 LTS Server Edition,
- PostgreSQL 10

As workload driver:

The LPAR setup consisted of:

- 16 dedicated cores processors
- 64 GB storage
- 1x 25GbE RoCE Express2 card

The software used was:

- SUSE Enterprise Linux Server 12 Service Pack 3
- Yahoo! Cloud Serving Benchmark (YCSB) 0.15.0

## **Db2 Warehouse test**

The LPAR setup consisted of:

- 30 dedicated cores processors
- 256 GB storage
- 4x FICON Express16S cards supporting FCP devices
- 4x IBM Adapter for NVMe
- 4x 4TB NVMe SSDs

Storage subsystem setup:

- IBM Storwize V7000 with HDD

The software used was:

- Ubuntu 18.04.1 LTS Server Edition
- Db2 Warehouse 2.11
- IBM Java 1.8
- IBM internal Big Data Insights (BDI) benchmark

## PostgreSQL performance test

The purpose of this test is to determine the performance of the PostgreSQL database server under heavy transactional workload stress. The PostgreSQL server stores the data in either the NVMe devices, or in the IBM FlashSystem 900 storage. A transactional workload driver actually sends a request to the database server to search for a particular row/record within a table using an ID or key of the record. This request can either be a read request to send data back to the driver or an update/write request to write data to the table.

When a read-request is submitted to the server, the row/record is searched if it resides in the memory. If not, then the data must be fetched from the external storage, in our case NVMe or FlashSystem 900. Moreover, if the memory size increases, the I/O traffic decreases and vice versa.

The YCSB program suite was used to drive all workloads for this set of tests in another LPAR using the 25 Gbps network.

The Scenarios are described by these parameters

- Number of users (= threads) fixed to 256
- Record size: 64 KB
- Tests with FlashSystem 900 involves using 2 FCP Express32S adapters, using a 5 TB LUN
- NVMe devices configured as RAID-10
- The throughput is measured in transactions per second
- The PostgreSQL database size was 256 GB
- The measurements ran for 25 minutes
- Experiment 1: Read-Only (95% read)
- Experiment 2: Read-Mostly workload (95% read, 5% update)
- Experiment 3: Update-Heavy workload (50% read, 50% update)
- Every test first loads the database and then runs the above experiments




Figure 5: Postgre benchmark configuration showing FlashSystem and NVMe attachment

## Db2 Warehouse performance test

In comparison to a transactional workload, an analytical workload will search for a large set of data within the database and then return the result to the workload driver. The workload driver actually sends a “query” to the database server. A query is a composition of multiple conditions, and the database returns the set of row/records to the driver, which satisfy these conditions. The complexity of the query determines the amount of work the database server has to do. If the size of the database is larger than the memory size, the database server must search for the rows/records that satisfy the query conditions on the external disks and thus cause I/O traffic. The purpose of these tests is to determine the impact of the external storage on the performance of such queries (i.e., time in finding all the rows/records satisfying the conditions). For these tests, the database either resides on the RAID-10 NVMeS or the IBM Storwize V7000 system backed by HDDs.

For these tests, we used Big-Data-Insights (BDI) benchmark to test the performance of a server running Db2 warehouse. BDI is based upon TPC-DS Decision Support Benchmark. In these tests, there are multiple parallel drivers, and at one time each pushes a single query to the Db2 warehouse database. The drivers fetch these queries from a set, with queries categorized as simple, intermediate and complex queries. Thus, there is a mixture of simple, intermediate and complex queries running simultaneously on the database server. These queries will cause I/O traffic, and the latency of a read or write request to the external drives would determine the performance of the database server.

The workload drivers resided for this set of tests in the same LPAR.

The Scenarios are described by these parameters

- Number of parallel users fixed to 5, performing predefined SQL queries
- Record size: average 6 KB
- The throughput is measured in queries per hour
- Tests with FlashSystem 900 involves using 4 FICON Express16S adapters, using a 1 TB LUN
- NVMe SSD devices configured as RAID-10
- Scaling Linux memory from 128 GB to 192 GB and 256 GB
- The Db2 Warehouse database size was 500 GB




Figure 6: Db2 benchmark configuration showing Storwize V7000 and NVMe attachment


## Case study results

For the measurement series discussed in this chapter all of the described configurations were applied to the SUT.

### Transaction throughput

This topic shows the throughput rate for the executed workloads on the SUT.


### Observations



The I/O performance with the YCSB benchmark on PostgreSQL 10 shows

- equal throughput

using direct-attached NVMe SSD storage compared to using an IBM FlashSystem 900 storage server connected by 2 FCP Express32S adapters.



The throughput rate when running the BDI benchmark on Db2 Warehouse 2.11 showed

- equal throughput with sufficient memory
- 2x better throughput with constrained memory

using direct-attached NVMe SSD storage compared to using a Storwize V7000 storage server with HDD storage.

## Conclusion

For these particular workloads we can state that

- If I/O requests are completely saturating adapters and storage system as we hit in the YCSB benchmark on PostgreSQL 10 with a big number of small I/Os, the workload performance advantage depends upon the number of I/O requests that the storage can process at a time and therefore the number of direct-attached NVMe SSDs and the used RAID setup for them versus the number of used I/O adapters and connections to the storage server.
- In case of larger I/O requests or lower I/O frequency which we reached with the BDI workload on Db2 Warehouse 2.11, the throughput rates with NVMe are considerably higher than those of an attached storage server due to the lower latency of the NVMe SSDs.

## ***Price-Performance***

This topic shows the price-performance comparisons for the executed workloads on the SUT.


### **Disclaimer**

Performance and price-performance result is extrapolated from IBM internal tests running the

- YCSB benchmark workloads read-only, read-mostly, update-heavy on PostgreSQL on LinuxONE Emperor II native LPAR using direct-attached NVMe SSD storage versus IBM FlashSystem 900 storage and purchase costs for the measured storage options. Storage cost were extrapolated from purchase costs for 4x 4TB NVMe SSDs and 4x IBM Adapter for NVMe versus the purchase costs for an IBM FlashSystem 900 with 17.1 TB storage.
- BDI benchmark, which is based on TPC-DS, on Db2 Warehouse on LinuxONE Emperor II native LPAR using direct-attached NVMe SSD storage versus IBM Storwize V7000 HDD storage and purchase costs for the measured storage options. Storage cost were extrapolated from purchase costs for 4x 4TB NVMe SSDs and 4x IBM Adapter for NVMe versus the purchase costs for an IBM Storwize V7000 (Model 624) with 16.2 TB HDD storage.

Results may vary.


## Observations



The price-performance with the YCSB benchmark on PostgreSQL 10 is

- up to 92% better

using direct-attached NVMe SSD storage compared to using an IBM FlashSystem 900 storage server connected by 2 FCP Express32S adapters.



The price-performance when running the BDI benchmark on Db2 Warehouse 2.11 is

- up to 79% better with sufficient memory
- up to 90% better constrained memory

using direct-attached NVMe SSD storage compared to using a Storwize V7000 storage server with HDD storage.

## Conclusion

For these particular workloads we can state that NVMe SSD storage has a compelling price-performance advantage over I/O attachments with storage servers.

## References

The following documents are referenced in this white paper:

IBM Emperor II and Rockhopper II

<https://www.ibm.com/it-infrastructure/linuxone>

<https://www.ibm.com/us-en/marketplace/linuxone-emperor-ii>

<https://www.ibm.com/us-en/marketplace/linuxone-rockhopper-ii>

NVMe specifications

<https://nvmexpress.org/resources/specifications>

IBM Storage and the NVM Express Revolution (IBM Redbooks®)

<https://www.redbooks.ibm.com/redpapers/pdfs/redp5437.pdf>

Linux on Z Disk I/O Performance

[http://www.ibm.com/developerworks/linux/linux390/perf/tuning\\_diskio.html#dio](http://www.ibm.com/developerworks/linux/linux390/perf/tuning_diskio.html#dio)

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing 2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106-0032, Japan

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:**

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new

editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Software Interoperability Coordinator, Department 49XA  
3605 Highway 52 N  
Rochester, MN 55901  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any

other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. \_enter the year or years\_. All rights reserved.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

©Copyright IBM Corporation 2018  
IBM Systems  
Route 100  
Somers, New York 10589  
U.S.A.  
Produced in the United States of America,  
09/2018

## ***Trademarks***

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation and might also be trademarks or registered trademarks in other countries.

Db2®  
ECKD™  
FICON®  
IBM®  
IBM LinuxONE Emperor II™  
IBM FlashSystem®  
IBM LinuxONE Rockhopper II™  
IBM Z®  
Redbooks®  
Storwize®

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

The NVM Express® design mark is a registered trademark of NVM Express, Inc.

TPC, TPC Benchmark and TPC-C are trademarks of the Transaction Processing Performance Council.

Other company, product, or service names may be trademarks or service marks of others.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed.

Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

## ***Terms and Conditions***

Permissions for the use of these publications is granted subject to the following terms and conditions.

**Personal Use:** You may reproduce these publications for your personal, non-commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of the manufacturer.

**Commercial Use:** You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by

the manufacturer, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THESE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.