IBM Performance Report

zHyperLink for Db2 Active Logging with SAP Core Banking on IBM z17 & DS8K G10

Authors:

Fimy Hu
Veng Ly

Document Owner:
Fimy Hu

Table of Contents

Figures

Disclaimers

Neither this document nor any part of it may be copied or reproduced in any form or by any means or translated into another language, without the prior consent of the IBM Corporation. IBM makes no warranties or representations with respect to the content hereof and specifically disclaims any implied warranties of merchantability or fitness of any particular purpose. IBM assumes no responsibility for any errors that may appear in this document. The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

Performance data and customer experiences for IBM and non-IBM products and services contained in this document were derived under specific operating and environmental conditions. The actual results obtained by any party implementing any such products or services will depend on many factors specific to such party's operating environment and may vary significantly. IBM makes no representation that these results can be expected or obtained in any implementation of any such products or services.

The information in this document is provided "as-is" without any warranty, either expressed or implied. IBM expressly disclaims any warranties of merchantability, fitness for a particular purpose or infringement.

Trademarks

IBM, IBM logo, AIX, Db2, DS8800, FICON, FICON Express, High Performance FICON, OSA, OSA Express, zHyperLink, System Storage, IBM Z, IBM z16, IBM z17, z/Architecture, z/OS, and zSeries, are trademarks or registered trademarks of the International Business Machines Corporation in the United States and other countries.

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

Other products may be trademarks or registered trademarks of their respective companies. Information is provided "AS IS" without warranty of any kind. Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, 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.

Acknowledgements

The authors of this document would like to acknowledge the following individuals for their support and contributions:

For review SAP internal team
- Paul DeLessio
- Gina Skelton
- Jasmi Thekveli

Support from system programming teams
- Andrea Fuga
- Nick Crimmins
- Jeff Torres
- Joe Stein
- Ernie Balloni
- Chris Van Wagner

SME Consultation:
- Akiko Hoshikawa: Distinguished Engineer – Db2 for z/OS and AI Product Development
- Robert Stjohn: STSM – IBM Z Performance and Perf. Tools
- Nick Clayton: STSM – Solution Architect for DS8000, Enterprise Storage Development
- Beth Peterson: STSM – DS8000 Storage CKD I/O Architecture/Development

Feedback

Please send comments or suggestions for changes to fshu@us.ibm.com

1.0 Introduction

To manage the growing customer base of a commercial bank the capacity of the banking system must also grow to handle the increased volume of transactions. The system must also continue to ensure timely completion of overnight account settlement processing within a designated window of time. Db2 for z/OS customers can improve both transactional latency and overnight batch processing speed by leveraging zHyperLink 2.0 for IBM Z.

zHyperLink is a short-distance, direct I/O connection between the mainframe and physical storage systems that provides significant improvement in latency over high-performance FICON. It has three modes of operation: read, write, or both. zHyperLink Read assists in speeding up transaction processing by reducing the latency of retrieving data from the direct access storage device (DASD). zHyperLink Write improves the latency and throughput of Db2 active logging, which may help to reduce the total elapsed time to complete log-bound workloads.

This paper details the results of a performance study on the impact of exploiting zHyperLink Write to reduce the total batch elapsed time for the SAP Core Banking Account Settlement workload. This workload simulates the overnight account balancing of 100 million bank accounts.

The study evaluates the performance of high performance FICON (zHPF) and zHyperLink Write across multiple test configurations spanning the two latest generations of mainframes, the IBM z16 and IBM z17, and DASD, the DS8900 (Generation 9) and DS8A00 (Generation 10).

The results show significant benefits to utilizing zHyperLink Write for Db2 batch workloads. These results are conducted as part of a system stress test. Measurement results with zHyperLink may vary depending on differences in workload, software and hardware.

2.0 zHyperLink Write Overview

zHyperLink Write provides short direct I/O connections between an IBM Z CEC and an IBM System Storage DS8000. Its latency is significantly improved over traditional FICON or high-performance FICON (zHPF). The short direct links provides a lower latency I/O path for Db2 active logging. This allows for synchronous processing of I/O requests during which the processor will wait for completion of the log write request, instead of asynchronous I/O which would have inefficiency from task dispatching, cache disruption, queueing delays, and I/O interrupts.

The high-speed link, together with the overhead saved by processing I/O synchronously, allows zHyperLink Write to achieve significantly lower Db2 active logging latency which in turn can improve overall throughput and elapsed time for Db2 log-bound workloads. Additionally, zHyperLink Write synchronous I/O processing is eligible for 100% offload to IBM Z Integrated Information Processors (zIIP). In a system configured with sufficient zIIP capacity, the General Processors(GCP) will not incur any additional processing costs from utilizing zHyperLink Write.

Figure 1 shows the layout of a site configured for zHyperLink (zHL). The CECs are connected to the DS8000 with both FICON and zHL links. The zHyperLink optical links have a maximum range of 150 meters and must be directly connected between the CEC and storage system. The zHL adapters on IBM's DS8000 10th generation storage systems and IBM z17 are built on PCIe Gen4 technology. PCIe Gen4 offers twice the data transfer rate of PCIe Gen3. Gen4 operates at 16 GB/s. Migrating from PCIe Gen3 to PCIe Gen4 can provide significant performance enhancements.

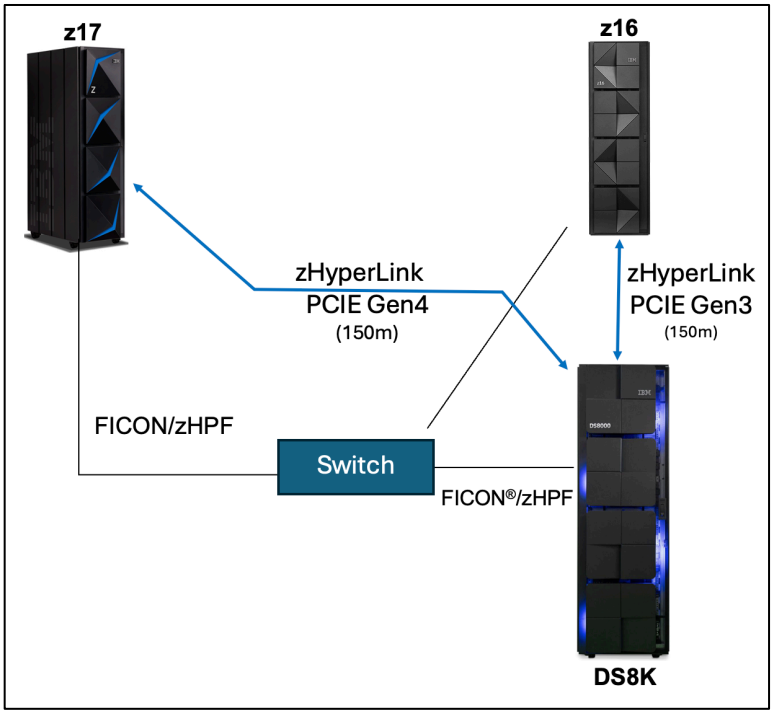

Figure 1: zHyperLink Single Site Layout

3.0 Test Overview

3.1 Measurement Scenarios

The performance of zHyperLink Write was evaluated across three test scenarios. The first evaluates the performance differences between utilizing High Performance FICON (zHPF) and zHyperLink (zHL) for Db2 active logging. The second scenario compares zHyperLink performance between the latest two generations of mainframes, the IBM z16 and IBM z17. The third scenario compares the performance of zHyperLink between the two latest generations of storage devices, the IBM DS9850 and DS8A50.

The table below shows a summary of each test scenario. All aspects other than the components highlighted in blue were kept constant for each test scenario, e.g. total processors, system memory, software levels, and workload level.

Test Scenario:	**zHPF vs. zHyperLink**	**zHyperLink with IBM z16 vs. IBM z17**	**zHyperLink with DS8950 vs. DS8A50**
CEC Model:	IBM z17	IBM z16 vs. IBM z17	IBM z16
I/O:	zHPF vs. zHL	zHL	zHL
DASD Unit:	DS8A50	DS8A50	DS8950 vs. DS8A50

3.2 Workload Characteristics

The test workload chosen for this study was the SAP Banking Services (SBS) 9.0 Account Settlement workload. This workload simulates the typical mass account balancing that takes place during the night processing of a retail bank and calculates the account balance, interest and charges for each bank account in the database. It typically runs overnight and often has a critical window of time in which processing must be completed before the bank opens for business the following day. This logging intensive workload matches the profile of a workload that would benefit most from the improved Db2 active logging latency that zHyperLink Write provides. For this study the workload was configured with 256 concurrent batch jobs to process the 100 million accounts in the test database.

3.3 Key Performance Indicators

To provide performance insights, and gauge the system resource consumption the following Key Performance Indicators (KPIs) were captured for the measurements of this study:

- Db2 for z/OS Active Logging Rate (MB/second)
- Workload Throughput Rate (Millions of accounts processed per hour)
- Average GCP & zIIP utilization for Db2 database server
- Average GCP and zIIP processor time per processed account (milliseconds)
- Total Batch Elapsed Time (hours)

4.0 Test Environment

4.1 SAP Core Banking with zHyperLink Landscape

The SAP on Z solution environment is a three-tier system landscape consisting of a database layer, application layer, and presentation layer. Figure 2 below shows the physical layout of the test environment. The Db2 for z/OS database server and SAP application servers are hosted on logical partitions (LPAR) within the same IBM z16 or IBM z17 CEC. Database and application servers communicate via TCP/IP over Open System Adapter (OSA) devices. The presentation server serves as the application front-end client and drives the workload transactions. The performance of the presentation server is not measured or considered for the purposes of this study.

Four IBM System Storage units were utilized for this test configuration. The Db2 database tables are stored across two DS8900 units and are connected to the CEC via FICON links. Depending on the test scenario the Db2 active logging volumes are hosted on either a DS8950 or DS8A50 and are connected to the CEC via both FICON links and two 80m zHyperLink Express links.

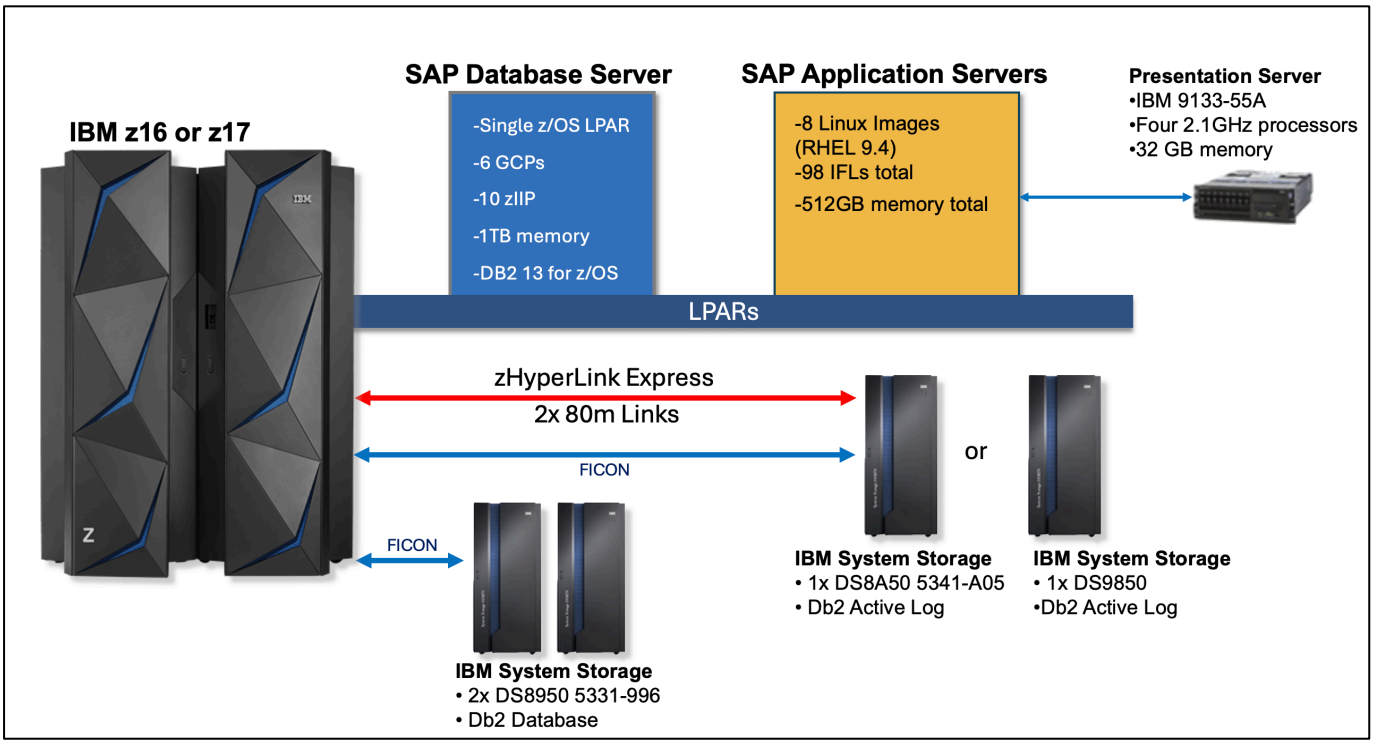

Figure 2: Test Environment

4.2 Hardware Environment

System Z Database Server
The Db2 database server operated within an LPAR on an IBM z16 or IBM z17
Configured with 6 dedicated GCPs, 10 zIIPs, and 1TB of memory

Database Storage
SAP SBS 9.0 database with 100 million banking accounts
Stored across two IBM System Storage D8900s

Database Active Log & Connectivity
The database log volumes resided on one IBM System Storage DS8A50/DS8950
4x FICON Express16SA Host Adapters
1x zHyperLink Express Host Adapters (2x 80m Links)

SAP Application Servers
Eight native Linux LPARs on an IBM z16 or IBM z17
Each LPAR was configured with 12 IFLs and 64GB of memory

SAP Presentation Server
The SAP presentation server was hosted on an IBM Power 750
31 3.6 GHz processors
125GB memory

4.3 Software Environment

Database server:
z/OS release 3.1
Db2 for z/OS 13

Application Server:
RHEL 9.2
IBM Data Server Driver for CLI -- Db2 Connect 11.01.0000
SAP Netweaver 7.5 SP6
SAP Banking Services 9.0 SP4
SAP Kernel 7.49 Level 823

5.0 Measurement Results and Analysis

5.1 IBM z17 zHyperLink Write vs zHPF

This test scenario evaluates the performance results between two configurations: Db2 active logging with high performance FICON (zHPF) and Db2 active logging with zHyperLink Write (zHL Write). Figure 3 below shows the logging rate in MB/second for the zHPF and zHL Write configurations.

Enabling zHL Write increased the logging rate from 214 MB/sec to 594 MB/sec, by 177%. This is achieved through a combination of zHL Write's synchronous I/O, which avoids queueing delays and I/O interrupts, as well as the lower I/O latency of the specialized 80m zHyperLink physical connections.

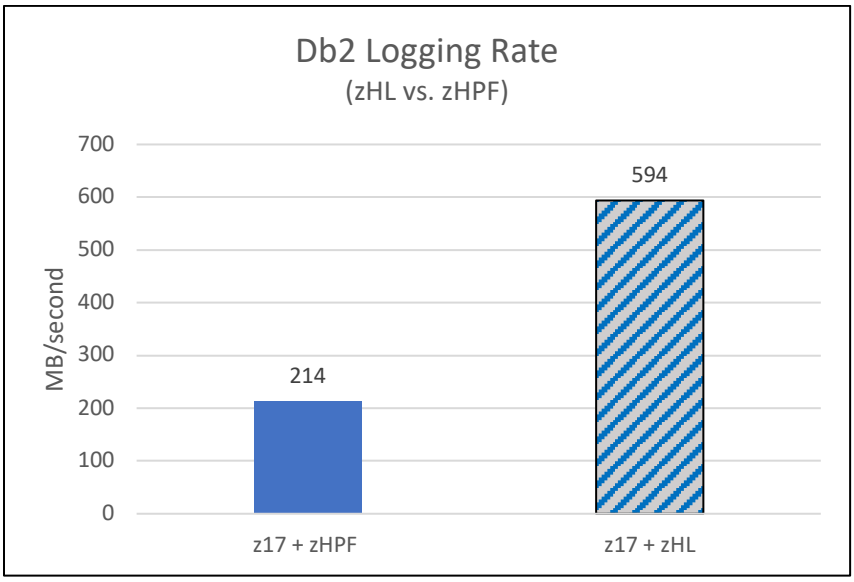

Figure 3: IBM z17 zHL vs. zHPF – Db2 Active Logging Rate

The higher logging rate indicates that enabling zHL Write alleviates the constraints to the Db2 logs which also improves the average workload throughput rate. Figure 4 shows the average workload throughput rate measured in millions of accounts processed per hour. The zHPF baseline configuration had a throughput rate of 17.29 million accounts per hour. Switching to zHyperLink Write increased the average throughput rate by 176% to 47.75 million accounts per hour. The improvement to throughput tracks closely with the improvement to the Db2 active logging rate.

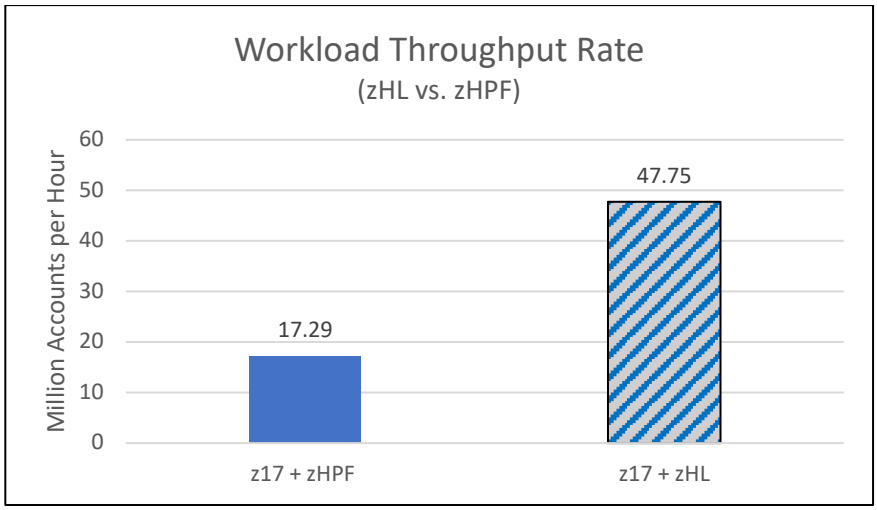

Figure 4: IBM z17 zHL vs. zHPF – Workload Throughput Rate

By alleviating the constraints on Db2 logging and increasing the workload throughput rate the system is also capable of processing the 100 million bank accounts more quickly. This requires additional processing capacity, which is reflected in the average processor utilization, shown in figure 5. GCP utilization in the chart is shown in blue while zIIP utilization is shown in green. The LPAR was configured with a total of 6 GCPs and 10 zIIPs for both zHPF and zHL configurations.

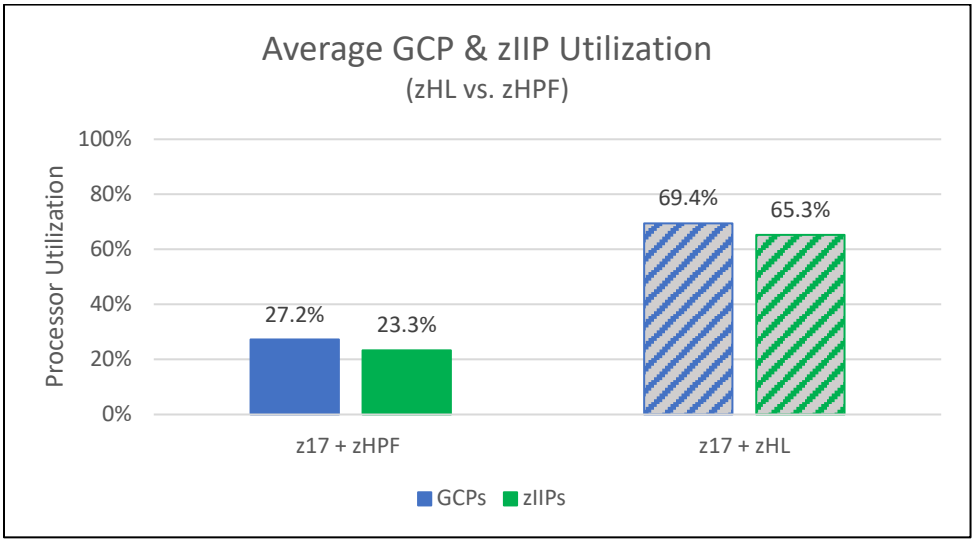

Figure 5: IBM z17 zHL vs. zHPF – Average GCP and zIIP Utilization

The baseline zHPF measurement had an average GCP utilization of 27.2% and zIIP utilization of 23.3%. The zIIP utilization is driven by the workload activity. SAP application servers access Db2 for z/OS through the Distributed Relational Database Architecture (DRDA). Up to 60% of the Db2 processing for DRDA requests can be offloaded from GCPs to zIIPs.

The zHL Write enabled graph shows the additional capacity that was required to support the configuration and increased throughput rate. The GCP utilization increased by 155% to 69.4%. The zIIP utilization increased by 180% to 65.3%. This increase in zIIP utilization accounts for both the higher throughput rate as well as the zHyperLink Write synchronous I/O processing cost, which is entirely offloaded to zIIPs.

Although the zHyperLink Write configuration requires more system capacity, the total cost of processing the workload is lower. Figure 6 shows the average GCP and zIIP time spent to complete processing for each bank account, in milliseconds. Each account processed in the zHPF baseline took 0.34 ms of GCP time and 0.49 ms of zIIP time. In comparison, the zHyperLink Write configuration took 0.31ms of GCP time, an 8.8%. The average zIIP processing time remained 0.49 ms per account. This is due to the additional zHyperLink Write I/O processing that was offloaded to zIIPs.

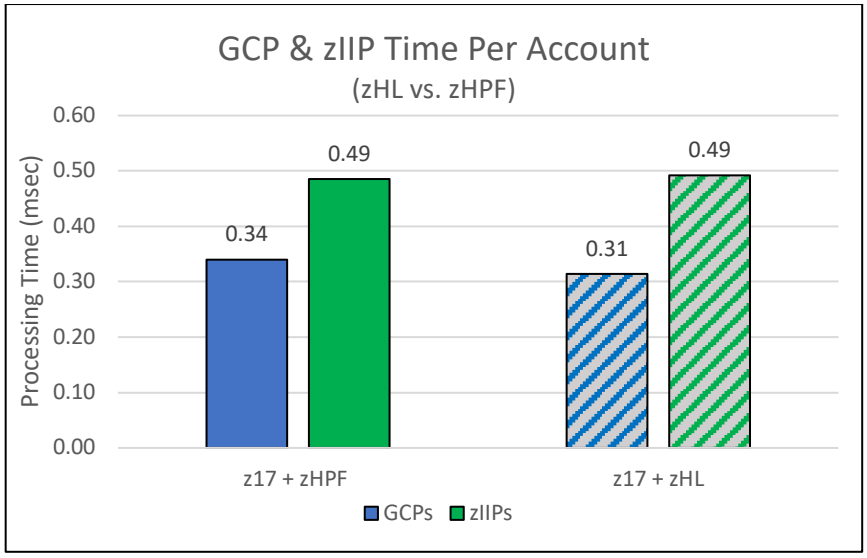

Figure 6: IBM z17 zHL vs. zHPF – GCP and zIIP Time per Account

In addition to the overall lower processing cost, the most significant benefit of zHyperLink Write is the reduction in the total batch elapsed time. Figure 7 shows the total batch processing time of the account settlement of 100 million bank accounts. The boost in the Db2 active logging rate, which drives a higher throughput rate, decreased the total batch elapsed time by 64% from 5.8 hours down to just over 2 hours. By utilizing zHyperLink Write the system can process the same amount of work in less time and at lower cost.

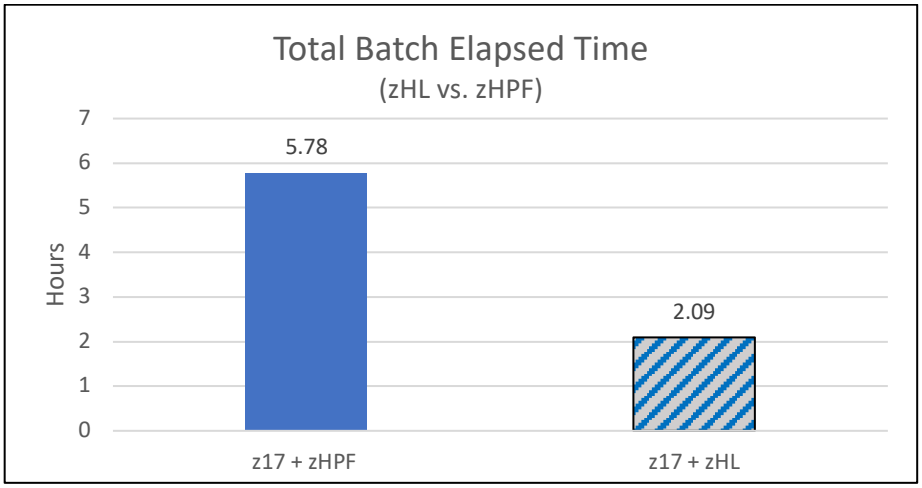

Figure 7: IBM z17 zHL vs. zHPF – Total Batch Elapsed Time

5.2 zHyperLink Write on IBM z17 vs IBM z16

The following performance data and analysis compare the results of two zHyperLink Write enabled configurations on the IBM z16 and IBM z17. Both utilize a DS8A50 for Db2 logging and identical hardware and software levels apart from the mainframe generation.
Figure 8 below shows the logging rate for the IBM z16 at 537 MB/s and IBM z17 at 594 MB/s. The z17 has 10% higher logging rate.

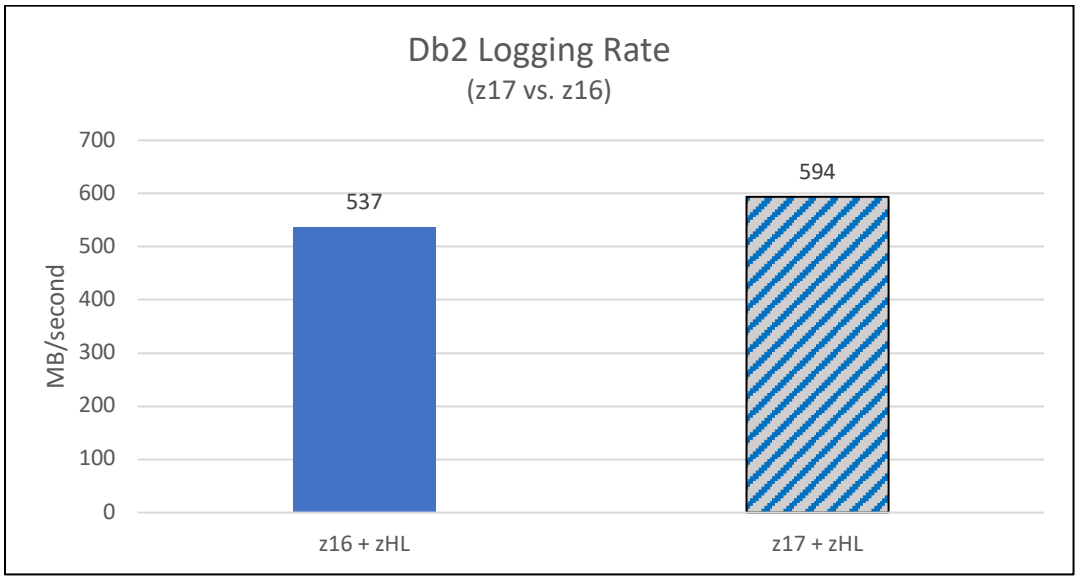

Figure 8: zHyperLink z17 vs. z16 – Db2 Active Logging Rate

The 10% improvement to the Db2 active logging rate also increases the workload throughput rate of the system. Figure 9 below shows the average throughput rate of processing the 100 million bank accounts. The baseline z16 has a throughput rate of 42.97 million accounts per hour while the z17 has a throughput rate of 47.75, an increase of 11%.

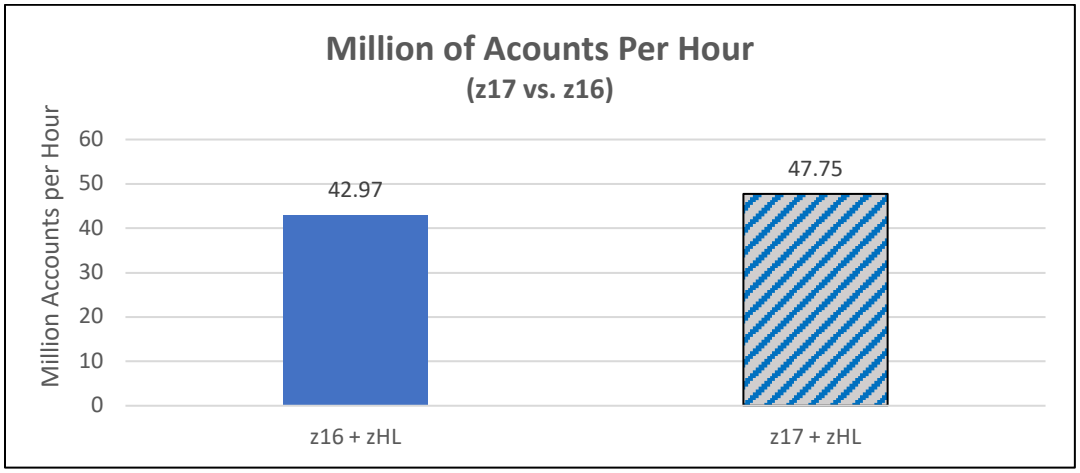

Figure 9: zHyperLink z17 vs. z16 – Workload Throughput Rate

Handling a higher throughput rate requires additional system capacity and usually results in a higher average processor utilization. However, due to improvements in the IBM z17, the average GCP and zIIP utilization of the z17 is lower than that of the z16 despite having an 11% higher throughput rate.

Figure 10 below shows the average GCP and zIIP utilization of the IBM z16 and IBM z17 configurations. The z16 baseline measurement had an average utilization of 75.7% on GCPs and 67.4% on zIIPs. Processing the same workload on the IBM z17 requires 8% lower GCP utilization and 3% lower zIIP utilization.

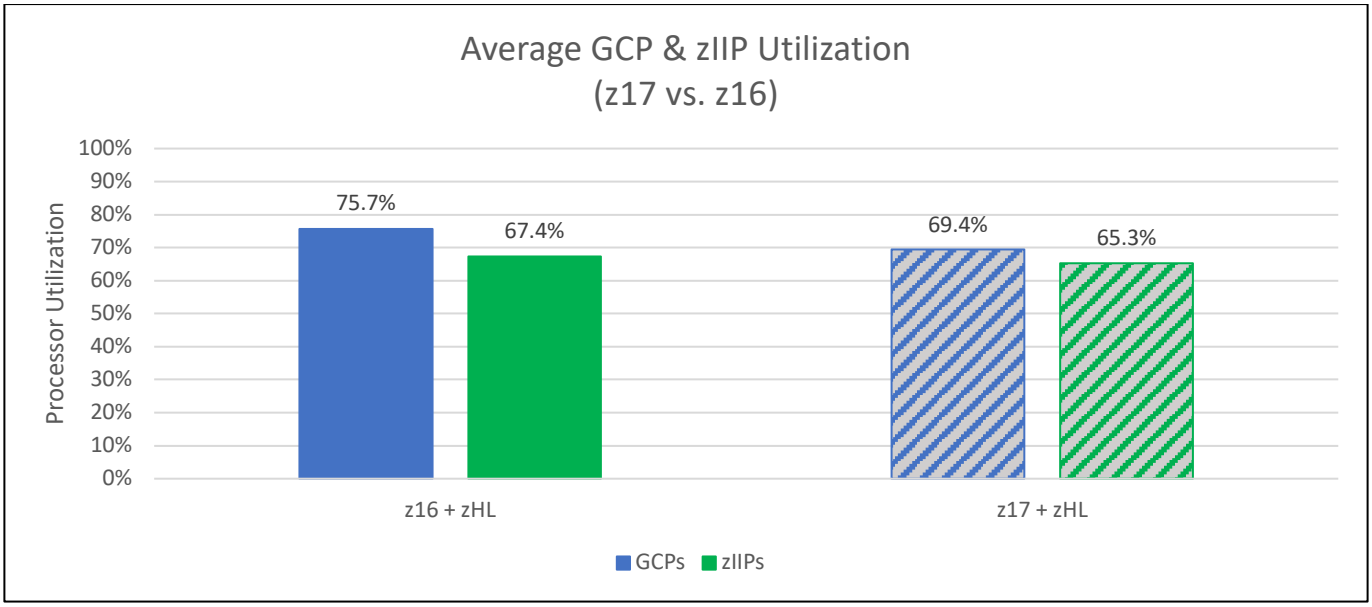

Figure 10: zHyperLink z17 vs. z16 – Average GCP and zIIP Utilization

The 8% reduction in average GCP utilization on the IBM z17 includes the additional processing required by the 11% increase in throughput rate. Accounting for the higher throughput rate the average time per account processed is even lower than 8%.

Figure 11 below shows the average processor time spent on GCPs and zIIPs per account. Each account processed on the IBM z16 required an average processing time of 0.38 ms on a GCP and 0.56 ms on a zIIP. The same workload executed on the IBM z17 requires 18% less processor time per account on GCPs and 12% less time on zIIPs.

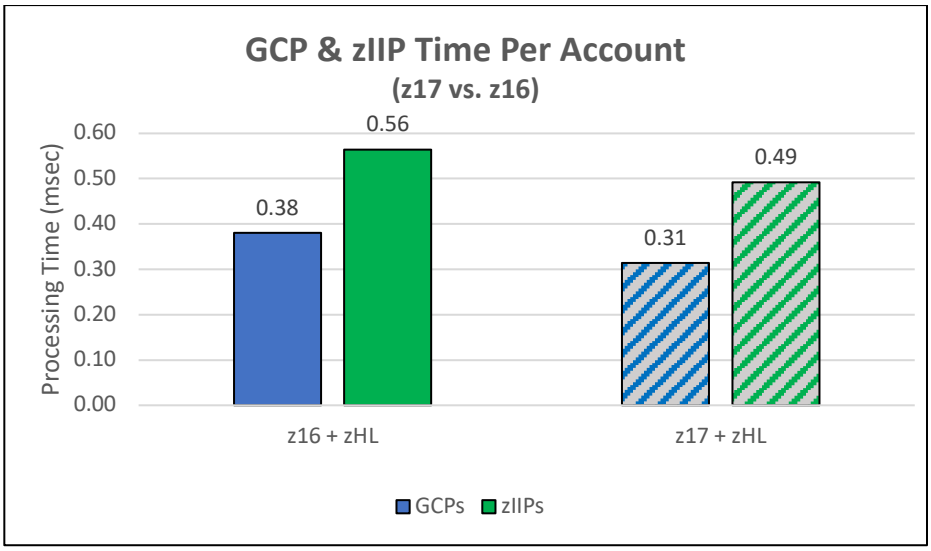

Figure 11: zHyperLink z17 vs. IBM z16 – GCP and zIIP Time per Account

The 11% improvement in workload throughput rate on the IBM z17 also resulted in a lower batch elapsed time. Figure 12 shows the total batch elapsed time of the IBM z16 and IBM z17 in hours. The IBM z17 has a 10% shorter batch elapsed time, reducing the time from 2.33 hours to 2.09 hours.

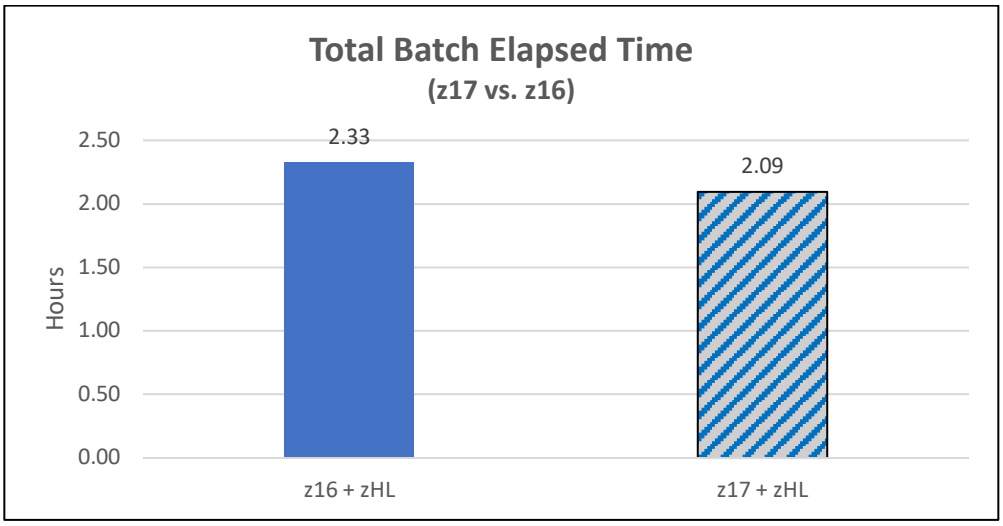

Figure 12: zHyperLink z17 vs. IBM z16 – Total Batch Elapsed Time

5.3 zHyperLink for Db2 Log Writes on DS8A50 vs DS8950

This test scenarios compares between two zHyperLink Write configurations on an IBM z16 utilizing two different generations of IBM System Storage. The baseline configuration utilizes a DS8950 (Gen 9) for the Db2 active logging volumes. The test configuration utilizes a DS8A50 (Gen 10) for the logging volumes. Otherwise, the two configurations utilize all the same hardware and software levels and load.

Figure 13 shows the logging rates of the two zHyperLink enabled configurations. The baseline DS8950 has a logging rate of 383 MB/second. Switching to the DS8A50 improves the logging rate by 40% up to 537 MB/second.

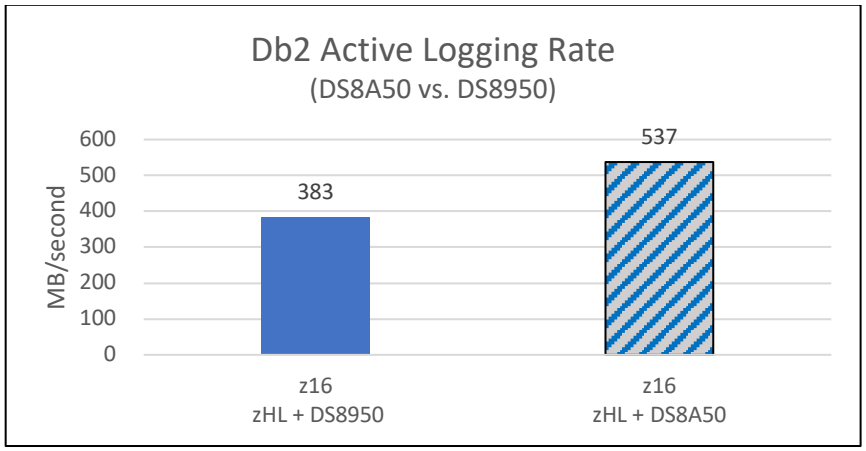

Figure 13: zHyperLink DS8A50 vs. DS8950 – Db2 Active Logging Rate

The 40% increase in Db2 active logging rate helps to improve the workload throughput rate, shown in Figure 14. The workload throughput rate shows similar gains of 39% increasing it up to 42.97 million accounts per hour.

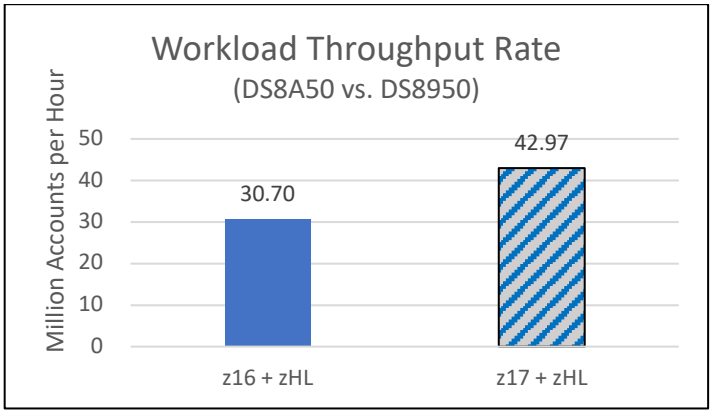

Figure 14: zHyperLink DS8A50 vs. DS8950 – Workload Throughput Rate

The higher throughput rate allows the system to complete processing of the same load in a shorter time. As a result, the average processor utilization is increased, as shown in Figure 15.

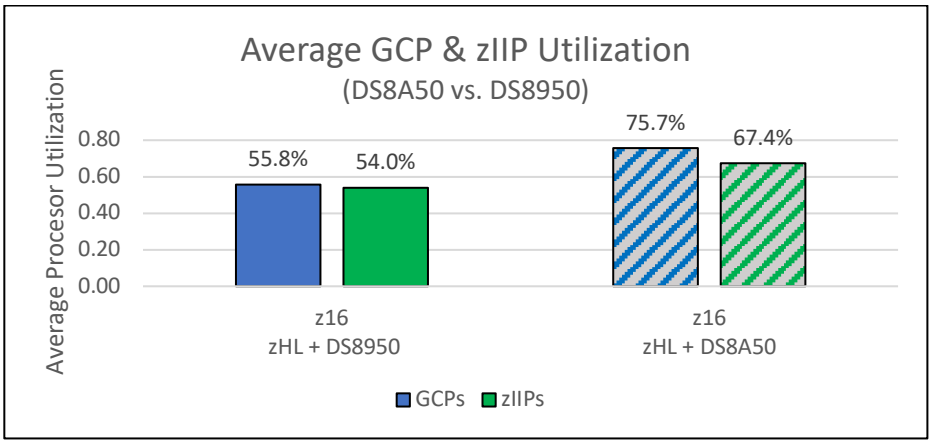

Figure 15: zHyperLink DS8A50 vs. DS8950 – Average GCP and zIIP Utilization

GCP utilization is shown in blue while zIIP utilization is shown in green. The zIIP utilization is driven by a combination of the offloaded DRDA eligible processing and zHyperLink Write synchronous I/O.

The DS8950 baseline had an average GCP utilization of 55.8% and zIIP utilization of 54%. The DS8A50 configuration required 35% more GCP capacity and 25% more zIIP capacity, resulting in an average utilization of 75.7% and 67.4% respectively.

Although the DS8A50 configuration required more GCP and zIIP capacity, the cost per account processed remains the same or lower. Figure 11 shows the average GCP and zIIP time per account processed. The baseline configuration required 0.39 ms of GCP time and 0.63 ms of zIIP time per account. The DS8A50 configuration maintains the required GCP time at 0.38 ms per account and reduces the required zIIP time by 11% to 0.56 ms.

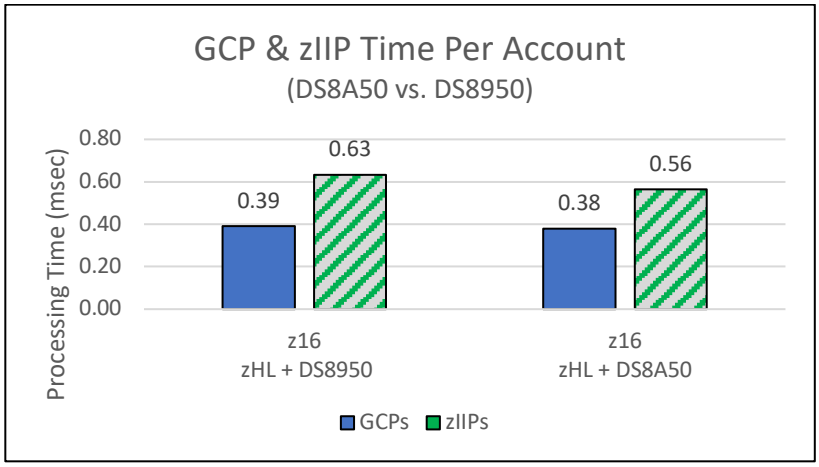

Figure 16: zHyperLink DS8A50 vs. DS8950 – GCP and zIIP Time per Account

In addition, the 40% improvement in throughput rate enabled by the DS8A50 performance results in a lower total batch elapsed time. Figure 17 shows the total batch elapsed time for the DS8950 and DSA850. Utilizing the DS8A50 over the DS8950 shortened the total batch time by 28.5% from 3.26 hours, down to 2.33 hours.

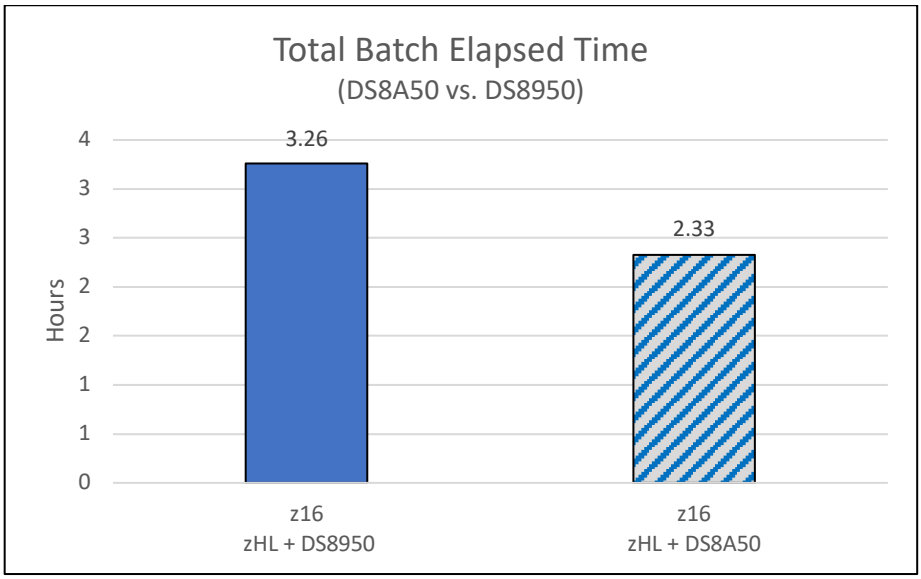

Figure 17: zHyperLink DS8A50 vs. DS8950 – Total Batch Elapsed Time

6.0 Conclusion

Utilizing zHyperLink Write for Db2 active logging can significantly reduce the total batch elapsed time without increasing the total cost of processing. Measurements using the SAP Core Banking Account Settlement workload show that switching Db2 logging I/O from High Performance FICON (zHPF) to zHyperLink Write (zHL) can reduce the total batch time to complete processing of 100 million bank accounts by up to 64% from 5.8 hours to 2.1 hours.

Figure 18 below shows an overview of the total batch elapsed time for a series of configurations with incremental upgrades. The baseline configuration, an IBM z16 configured with zHPF and a DS8950, took 5.85 hours to complete processing of the 100 million bank accounts. For comparison, the equivalent z17 with zHPF measurement (section 5.1) took 5.78 hours, a 2% improvement over the z16.

Enabling zHyperLink Write for the IBM z16 resulted in a substantial reduction in total batch elapsed time of 44%. Upgrading the logging storage system to a DSA850 further reduced the batch elapsed time by 28%. Finally, upgrading the CEC to an IBM z17 reduced the batch elapsed time by an additional 10%.

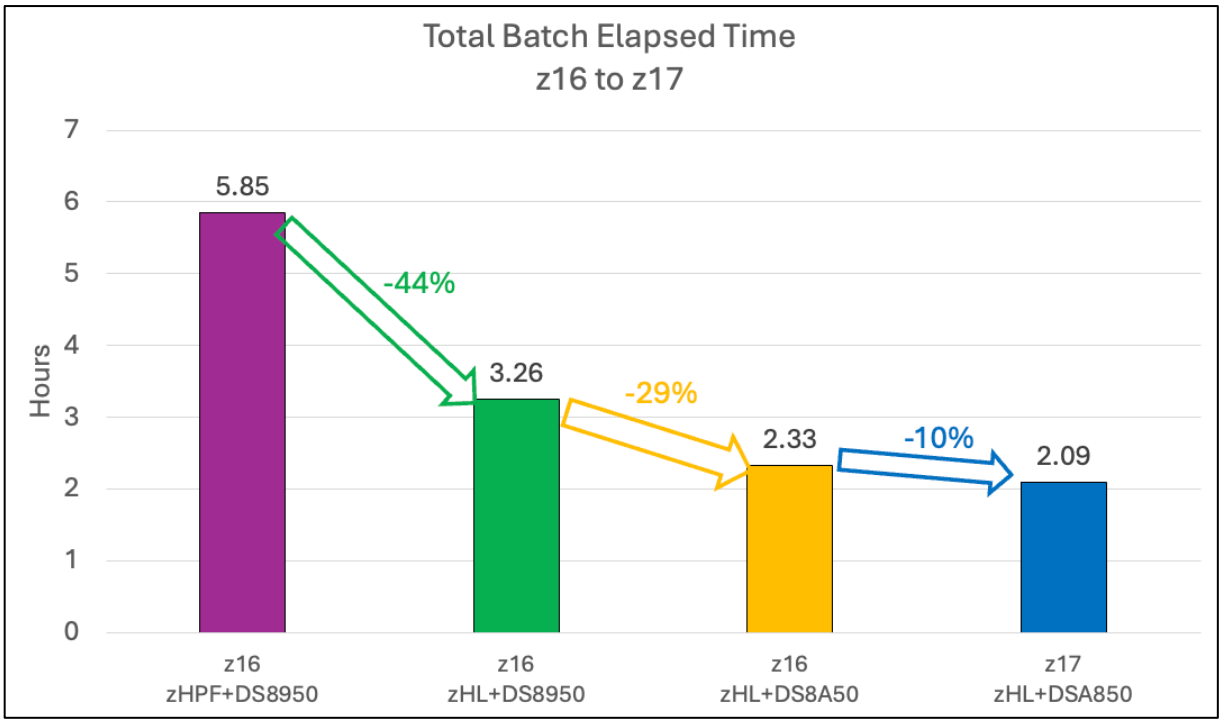

Figure 18. Total Batch Elapsed Time Overview

Additionally, zHyperLink Write can reduce the total elapsed time of Db2 batch workloads without impacting the total cost of processing. Figure 19 below shows the cost in processor time (milliseconds) per account for both GCPs and zIIPs across the same four configurations shown in figure 18.

Each account processed in the baseline z16 configuration required 0.42 ms of GCP time and 0.63 ms of zIIP time. Switching from zHPF to zHL reduced the GCP cost per account by 7%. Upgrading the logging storage system to a DS8A50 reduced the GCP time per account by an additional 2% and zIIP time by 11%. Finally upgrading the CEC from an IBM z16 to IBM z17 further reduced the GCP time per account by 18% and zIIP time by 12%.

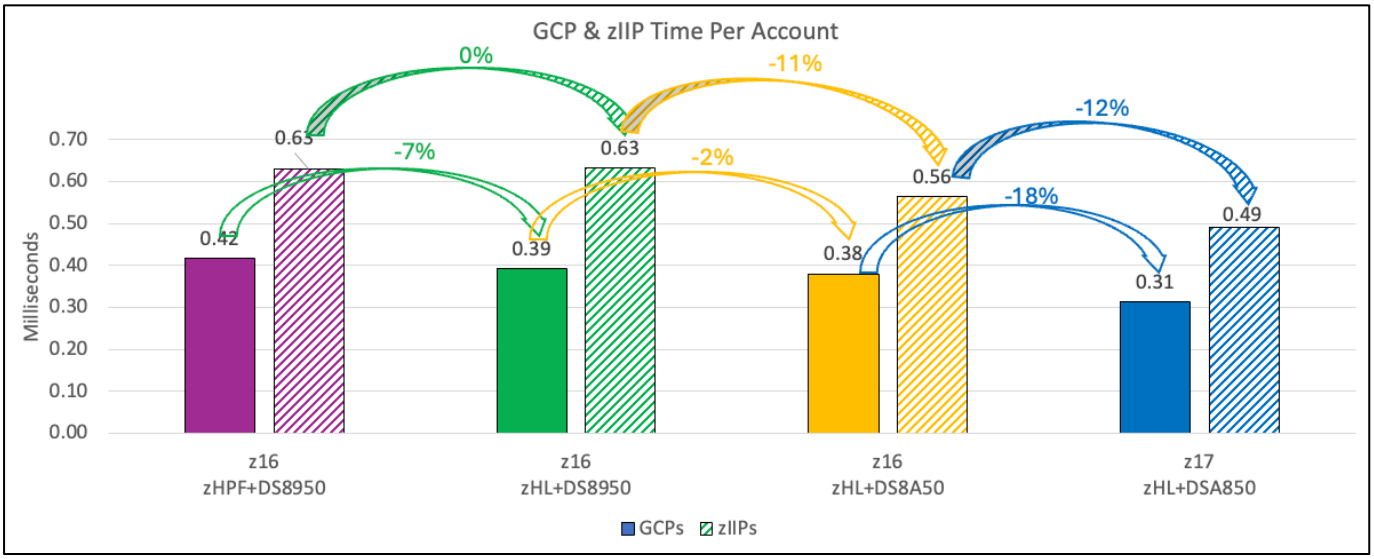

Figure 19. GCP and zIIP Time Cost per Account

zHyperLink Write can shorten the elapsed time of a logging-bound batch workload without incurring additional processor costs, and in some cases may lower, processor costs. The impact zHyperLink Write has on a given workload will depend on the deployed hardware and software configuration and whether the workload is constrained by the Db2 active logging rate.

7.0 References

1. Getting Started with IBM zHyperLink for z/OS
   https://www.redbooks.ibm.com/redpapers/pdfs/redp5493.pdf

2. SAP on Z IBM Z Reference Architecture - SAP for Banking
   https://www.sap.com/documents/2015/07/8aab7888-5b7c-0010-82c7-eda71af511fa.html

3. IBM Db2 13 for z/OS
   https://www.redbooks.ibm.com/abstracts/sg248527.html

4. IBM Db2 13 for z/OS Performance Topics
   https://www.redbooks.ibm.com/redbooks/pdfs/sg248536.pdf

5. IBM z17 Technical Guide
   https://www.redbooks.ibm.com/abstracts/sg248579.html

6. IBM Storage DS8A00 Architecture and Implementation Guide
   https://www.redbooks.ibm.com/abstracts/sg248559.html