IBM z16 FICON Express32S Performance

(Addendum: SAP Capacity)

March 2023

Brian Murphy
Dinesh Kumar
Su KR Han

IBM Systems
© 2022-23, International Business Machines Corporation

Table of Contents

© Copyright 2022-23 IBM Corp.

Notices, Disclaimer and Trademarks

Copyright © 2022-23 by International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM Corporation. Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information may include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs, or services available in all countries in which IBM operates or does business. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein. The performance data contained herein was obtained in a controlled, isolated environment. Actual results that may be obtained in other operating environments may vary significantly. While IBM has reviewed each item for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Statements regarding IBM's future direction and intent are subject to change or withdraw without notice and represent goals and objectives only. The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

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

IBM, FICON, zHPF, System Storage, zSystems, z/OS, zEnterprise, and DS8900 are trademarks of International Business Machines Corporation in the United States, other countries, or both. Other company, products or service names may be trademarks or service marks of others.

Acknowledgements

The authors would like to thank the following colleagues for their expertise, insight, and contributions:

- Francis Gassert – IBM Systems, Poughkeepsie, NY
- Elke Nass – IBM Systems, Boeblingen, Germany
- Joachim Von Buttlar – IBM Systems, Boeblingen, Germany
- Brian Lupini – IBM Systems, Poughkeepsie, NY
- Michael Fee – IBM Systems, Poughkeepsie, NY
- Shelly Black – IBM Systems, Poughkeepsie, NY
- Andrea Harris – IBM Systems, Poughkeepsie, NY
- Yamil Rivera – IBM Systems, Poughkeepsie, NY
- Nicholas VanMater – IBM Systems, Poughkeepsie, NY
- Joseph Stein & Entire Infrastructure Team –IBM Systems, Poughkeepsie, NY

The authors would also like to thank all the reviewers of this paper for their helpful comments.

Introduction

This whitepaper assumes a familiarity with the general concepts of Z Systems. Readers unfamiliar with these topics should consult the IBM Z Mainframe Capabilities website: https://www.ibm.com/it-infrastructure/z/capabilities/networking

In this paper, “FICON Express32S” and “FEx32S” are used interchangeably unless otherwise noted in the context of the text. When zHPF is used by itself, it refers to the zHPF Protocol. When FICON is used by itself, it refers to the FICON Protocol. When FCP is used by itself, it refers to the FCP Protocol.

The IBM z16 supports the following FICON features:

- FICON Express32S (FEx32S) – new build
- FICON Express16SA (FEx16SA) – carry forward only
- FICON Express16S+ (FEx16S+) – carry forward only

The FICON features conform to the following architecture:

- Fibre Connections (FICON)
- High Performance FICON on Z (zHPF)
- Fibre Channel Protocol (FCP)

The FICON features provide connectivity between any combination of servers, directors, switches, and devices (control units, flash drives, tapes, and printers) in a Storage Area

© Copyright 2022-23 IBM Corp.

Network (SAN). The Z systems run special *System Assist Processor(s) (SAPs)* that are dedicated for meeting IO operation demands of the zOS LPARs and application workloads running in them. The *IO Subsystem capacity* or *SAP capacity* of Z systems is highly dependent on the number of LPARs running in the system, number of FICON features installed and used, number of IO devices being used and the technology of switches and external storage system which are part of the Storage Area Network (SAN).

Summary

The IBM z16 supports the new PCIe (PCI Express) Gen4 links with x16 lanes that provide increased data transmission and IO Subsystem (SAP) capacity.

The IBM z16 runs the new Telum™ processor with 7nm technology and dual-chip module design. IBM z16, with its industry first integrated on-chip AI accelerator, delivers latency-optimized inferencing designed to enable customers to analyze real-time transactions, at scale.

As compared to IBM z15, the IBM z16 has a massively redesigned CPC drawer architecture with greater cache access per core and higher capacity buses to provide greater processing power to application workloads and their IO requirements:

- The maximum IO processing capability on an IBM z16 with 5 SAPs (Max39, 1-drawer model) is over *2.3 M IOs/sec*, a significant 40% increase compared to an IBM z15 with 4 SAPs (Max34, 1-drawer model).
- The IBM z16 has no 5 physical drawer model as offered in IBM z15. Hence, this SAP capacity study did not focus on comparison between IBM z16 Max200, 4-drawer model vs. the IBM z15 Max190, 5-drawer model due to different cabling.
- The maximum IO processing capability on an IBM z16 with 24 SAPs (Max200, a 4-drawer model) is over *5.6 M IOs/sec*, a 12% increase compared to an IBM z15 with 16 SAPs (Max145, 4-drawer model).
- Due to IBM z16's massively redesigned CPC-drawer architecture, the trend in SAP capacity growth of IBM z16 vs. IBM z15 is very different compared to that of IBM z15 vs. z14.

Additional zHPF, FCP, and FICON product information is available on the IBM Z websites: https://www.redbooks.ibm.com/redbooks/pdfs/sg245444.pdf

https://www.redbooks.ibm.com/redpieces/pdfs/sg248951.pdf

© Copyright 2022-23 IBM Corp.

The Massively Re-designed IBM z16 and its Improvements

The IBM z16 is a breakthrough machine that has been massively redesigned when compared to the previous machine generations of z13, z14 and IBM z15. At the brain of it is the new Telum™ processor chip fabricated with state-of-the art 7 nanometer (7nm) technology with 17 metal layers that runs at 5.2GHz system frequency (Figure 1). The new Telum is an 8-core processor chip with each core having a 32MB private L2 cache located right next to the core providing applications and workload ultrafast access to this cache. In addition, these L2 caches from each of the 8-cores form a Virtual L3 cache through the high-bandwidth on-chip interconnect ring buses (Figure 1, 2). Each chip is part of a Dual Chip Module (DCM) and 4 such DCMs connected with a star-network of high bandwidth buses together come to form a single system drawer (Figure 2, 3). The Virtual L3 caches of each chip in-turn together come to form a Virtual L4 cache for the whole drawer comprising 4 DCMs (Figure 2). The largest configuration of the machine is formed by 4 such physical system drawers (Figure 3) instead of the 5 physical drawers offered before in IBM z15.

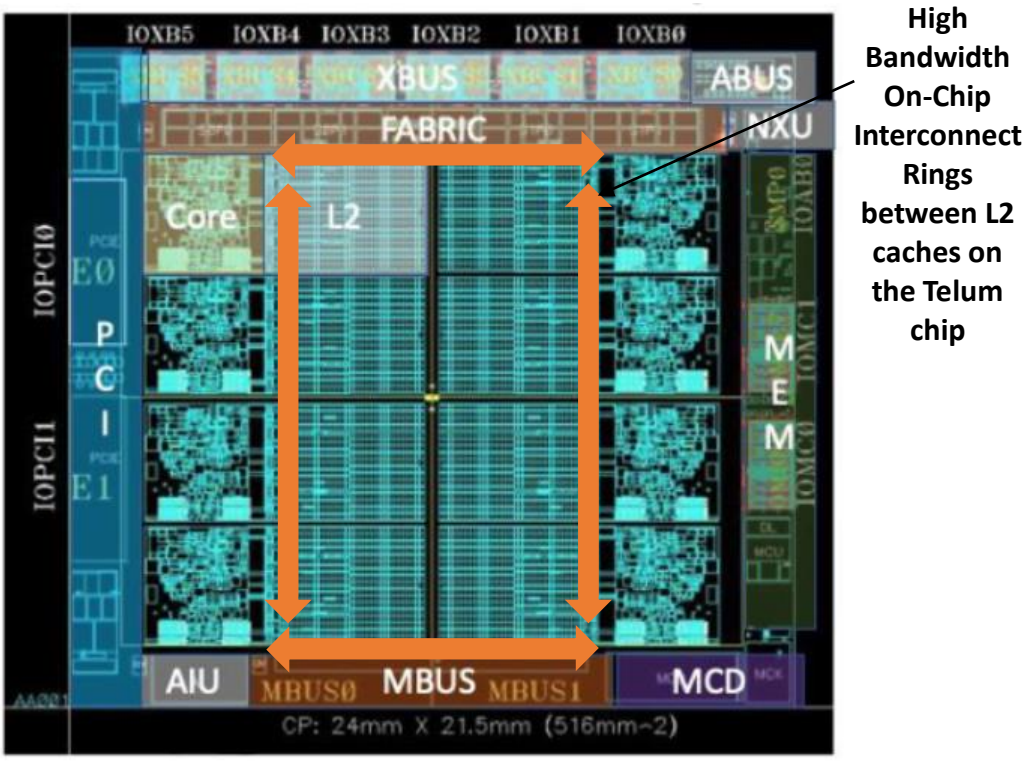

Figure 1 – IBM z16's new Telum™ processor with 7nm technology

This entire system of private L2, Virtual L3 and Virtual L4 caches on a drawer is managed by a new 'System Coherency Manager Fabric' that has been specifically

© Copyright 2022-23 IBM Corp.

introduced in the new IBM z16's design to replace the functions of the 'System Controller' (SC) chip in IBM z15. The SC chip of IBM z15 has not been carried over into the new IBM z16 design.

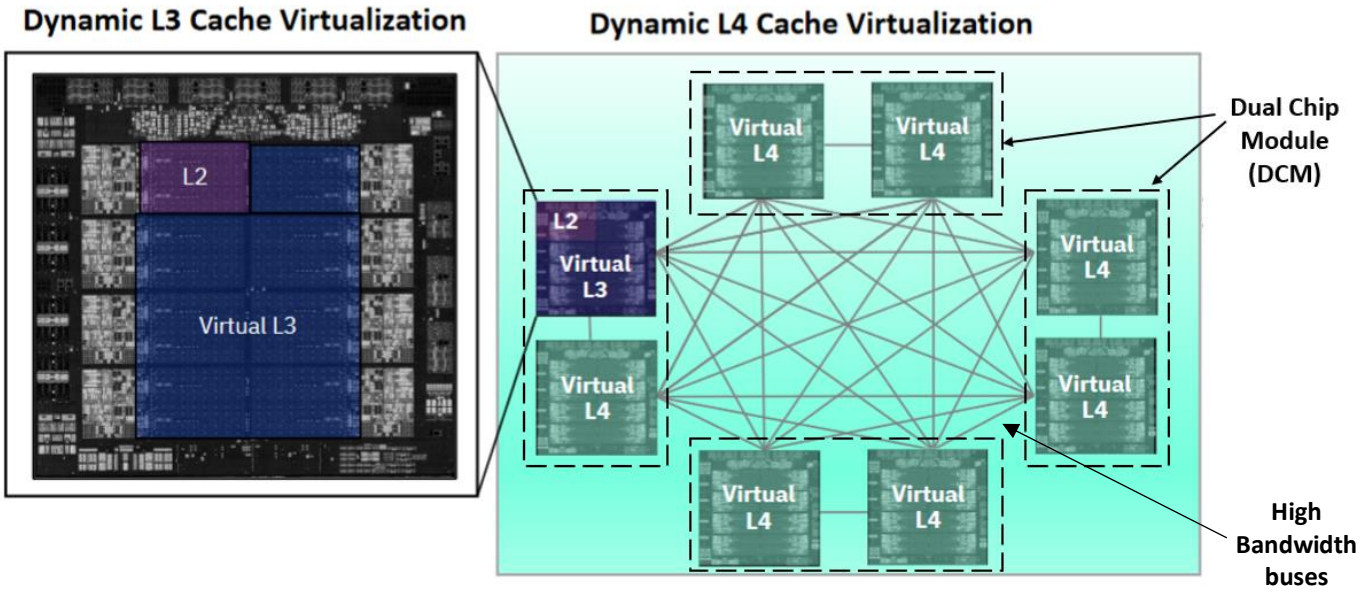

Figure 2 – IBM z16's On-Chip and On-Drawer Dynamic L3 and L4 Cache Virtualizations

The IBM z16 chip has new on-chip accelerator functions for AI (Artificial Intelligence) and Coupling (Figure 3). The system can hold 10 TB (Terabyte) memory DIMMs per drawer with a total 4-drawer machine capacity of 40 TB memory. The new Telum chip is a functionally lower latency chip due to improved event rate distribution which results in increased L2 cache hits, compared to IBM z15. Compared to IBM z15, a single drawer of IBM z16 is also a functionally lower latency drawer due to fully interconnected chips through the new star-network of high bandwidth buses (Figure 2, 3). Specifically, the high bandwidth off-chip buses/links between the 8 chips (4 DCMs) on a drawer offer an on-DCM bandwidth in excess of 100 GB/sec per direction and an off-DCM bandwidth of 6x33 GB/sec per direction. All caches in the drawer contribute to bottom line performance for the LPAR (Logical Partition), 2nd level guests and application workloads running within them. This new CP-cache-memory nest architecture is enhanced by in-drawer memory bandwidth in excess of 100 GB/sec which contributes to the massive boost in single drawer IO Read/Write capacity and performance which will be quantified in the later sections.

A special system operations traffic, which is part of the IBM z16 hardware architecture, utilizes a low-latency point-to-point network within the system. The different drawers within an IBM z16 system are connected via high bandwidth off-drawer buses/links that operate at 40 GB/sec per link pair.

© Copyright 2022-23 IBM Corp.

Due to the new design of IBM z16 described above, the more powerful IBM z16 has a CPC drawer capacity growth of at least 25% and an overall box capacity growth of 17% over the previous generation IBM z15.

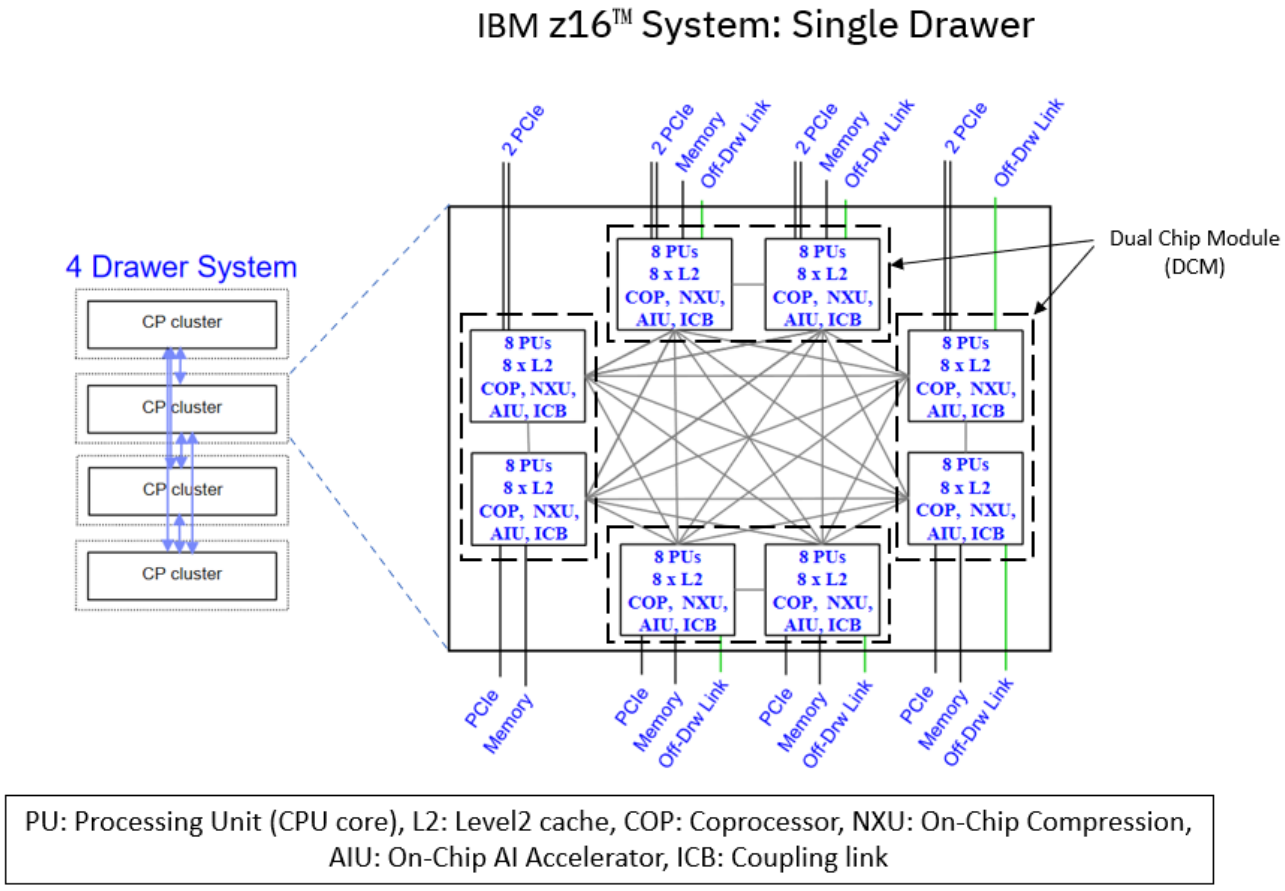

Figure 3 – Fully interconnected 4-drawer IBM z16 system with single drawer view

The single-threaded performance of IBM z16 has improved by 11% over IBM z15, and the SMT (Simultaneous Multi-Threading) and SMT 2nd level guest performance has also improved by ~11% over the previous generation IBM z15.

In addition to the hardware design changes and performance increments described above, significant firmware and software improvements have been made in the new IBM z16. The IBM proprietary z/OS operating system has been given a performance uplift to support the design changes made in the hardware. Specifically, the Real Storage Manager (RSM) component of z/OS has been optimized to support the new CP-

© Copyright 2022-23 IBM Corp.

cache-memory nest design of the system. The z/OS hypervisor PR/SM™ (Processor Resource/System Manager) that facilitates the running of LPARs has also been given improvements in its algorithms and efficiency. There is a new and dynamic SAP (System Assist Processors) placement algorithm for the placement of SAPs over the chip cores. Specifically, the placement of SAPs is now done through a dynamic placement algorithm and the position of SAPs on chip cores may vary from DCM to DCM and drawer to drawer for any given configuration.

IBM z16 Models Offered

The hybrid cloud ready IBM z16 can have a total of only 4 physical drawers, and no 5 physical drawer model is offered for IBM z16, unlike in IBM z15. This also eliminates an extra drawer-bus cable connecting the 4th physical drawer of IBM z16. Such a cable previously connected the 4th drawer on IBM z15 to the 5th physical drawer of IBM z15. Figure 4 below shows the various IBM z16 feature models offered to the customers: Max39, Max82, Max125, Max168 and Max200. These are all water-cooled CP (A01) full-scale models, as part of the IBM z16 offering. Note that the 4-drawer model, Max200, is a 4 physical drawer model and not a 5 physical drawer model.

# of Drawers	Machine Type	Model	Feature	CP Chips	DCMs	CPs	IFLs	Unassigned IFLs	zIIP	ICFs	IFPs	Std SAP	Add'l SAP	IBM Spares
1	3931	A01	Max39	8	4	0-39	0-39	0-38	0-25	0-39	2	5	0-8	2
2	3931	A01	Max82	16	8	0-82	0-82	0-81	0-54	0-82	2	10	0-8	2
3	3931	A01	Max125	24	12	0-125	0-125	0-124	0-82	0-125	2	15	0-8	2
4	3931	A01	Max168	32	16	0-168	0-168	0-167	0-110	0-168	2	20	0-8	2
4	3931	A01	Max200	32	16	0-200	0-200	0-199	0-132	0-200	2	24	0-8	2

Figure 4 – IBM z16 feature models offered: Max39, Max82, Max125, Max168, Max200

Due to this significant change in the total number of drawers offered in IBM z16, the SAP capacity of the Max200, 4-drawer model is not to be compared to the Max190, 5-drawer model that was offered for IBM z15. With the Max200 model, the maximum number of cores offered in IBM z16 is 200, compared to 190 cores offered in IBM z15.

Figure 5 below shows the total number of frames offered for each IBM z16 feature model. It also shows the total number of PCIe fanout card slots available for each of the models having different number of CPC drawers. Each fanout card slot can connect to a switch card of a single IO domain in an IO drawer. Each IO domain has a total of 8 PCIe slots for 8 FICON Express or other PCIe cards to be plugged in.

© Copyright 2022-23 IBM Corp.

Model	Feature	CPC Drawers	CP Chips	DCMs	Frames	Fanouts
A01	Max39	1	8	4	1-3	12
A01	Max82	2	16	8	1-3	24
A01	Max125	3	24	12	1-3	36
A01	Max168	4	32	16	2-4	48
A01	Max200	4	32	16	2-4	48

Figure 5 – IBM z16 feature models: # of frames and PCIe fanout connectivity for the IO domains in an IO drawer

RMF IO Queueing Activity Report

In IBM z16, the single XSAP core in any system configuration is now SMT enabled, unlike in IBM z15. The 'Std SAP' column in Figure 4 includes the 1 XSAP core for each of the feature models. Since the XSAP core is also now SMT enabled along with all other SAP cores, the RMF 'IO Queuing Activity' report now shows a total of *2N* IOP (IO Processor) or SAP threads when the total number of SAPs in the IBM z16 model is *N.* Figure 6 below shows an RMF snapshot for the 5 Std SAP, 1-drawer, Max39 model with 10 IOP threads.

1 I/O Q U E U I N G A C T I V I T Y

PAGE 1

z/OS V2R4 SYSTEM ID P34 START 03/28/2022-03.22.17 INTERVAL 000.00.59

RPT VERSION V2R4 RMF END 03/28/2022-03.23.16 CYCLE 1.000 SECONDS

-TOTAL SAMPLES = 59 IODF = CE CR-DATE: 01/31/2022 CR-TIME: 08.09.18 ACT: POR

INPUT/OUTPUT PROCESSORS

0 IOP	-INITIATIVE QUEUE- ACTIVITY RATE	AVG Q LNGTH	IOP UTILIZATION % IOP BUSY	% CMPR BUSY	% SCM BUSY	I/O START RATE	INTERRUPT RATE	% I/O REQUESTS RETRIED -- ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	RETRIES / SSCH ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY
0 00	1832.958	0.00	1.28	0.00	0.00	1832.958	3432.340	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
01	1833.161	0.00	0.57	0.00	0.00	1833.161	43.303	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
02	1831.504	0.00	1.03	0.00	0.00	1831.504	3588.779	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
03	1832.570	0.00	0.47	0.00	0.00	1832.570	49.678	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
04	1839.976	0.00	1.12	0.00	0.00	1839.976	4001.894	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
05	1833.026	0.00	0.43	0.00	0.00	1833.026	50.929	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
06	1831.944	0.00	1.24	0.00	0.00	1831.944	3683.790	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
07	1833.246	0.00	0.49	0.00	0.00	1833.246	32.904	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
08	1831.927	0.00	1.01	0.00	0.00	1831.927	3684.838	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
09	1832.823	0.00	0.49	0.00	0.00	1832.823	36.320	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
0 SYS	18333.13	0.00	0.81	0.00	0.00	18333.13	18604.77	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00

Figure 6 – RMF IO Queuing Activity Report's changes in IBM z16

© Copyright 2022-23 IBM Corp.

In Figure 7, it is to be noted that the ‘IOP Utilization’ section in ‘IO Queuing Activity’ report of IBM z16 not only shows the ‘% IOP Busy’ field for the IO activity busy period (utilization), but also the ‘% CMPR Busy’ and ‘% SCM Busy’ fields.

```
                                I/O   Q U E U I N G   A C T I V I T Y                                                    PAGE    1
           z/OS V2R5              SYSTEM ID RSC2            DATE 09/30/2021          INTERVAL 15.00.032
                                  RPT VERSION V2R5 RMF      TIME 23.44.33            CYCLE 1.000 SECONDS
TOTAL SAMPLES =    900   IODF = 07   CR-DATE: 09/14/2019   CR-TIME: 16.05.52     ACT: ACTIVATE

-------------------------------------------------------------------------------------------------------------------------------
                                                INPUT/OUTPUT PROCESSORS
-------------------------------------------------------------------------------------------------------------------------------
     -INITIATIVE QUEUE-  ----------- IOP UTILIZATION ------------   -- % I/O REQUESTS RETRIED --   -------- RETRIES / SSCH ---------

 IOP   ACTIVITY  AVG Q  % IOP  % CMPR  % SCM I/O START INTERRUPT          CP    DP    CU    DV            CP    DP    CU    DV
         RATE    LNGTH  BUSY   BUSY    BUSY    RATE      RATE      ALL  BUSY  BUSY  BUSY  BUSY     ALL  BUSY  BUSY  BUSY  BUSY

  00   1395.107   0.00   1.20   4.30    4.30  1384.433  3573.426   0.0   0.0   0.0   0.0   0.0    0.00  0.00  0.00  0.00  0.00
  01   1401.048   0.00   0.78   4.30    4.30  1390.375  1307.718   0.2   0.2   0.0   0.0   0.0    0.00  0.00  0.00  0.00  0.00
 SYS   2796.15    0.00   1.00   4.30    4.30  20824.67  21871.27   1.4   1.4   0.0   0.0   0.0    0.01  0.01  0.00  0.00  0.00
```

Figure 7 – RMF IO Queuing Activity Report’s changes in IBM z16

‘% CMPR Busy’ denotes the ratio of number of times IOP was found busy with EADM (Extended Asynchronous Data Mover) compression work (compress or decompress) to the total number of IO Processor samples.

‘% SCM Busy’ denotes the ratio of number of times IOP was found busy with SCM (Storage Class Memory) work to the total number of IO Processor samples.

The reader is also advised to consult the ‘RMF Report Analysis’ document at the URL below for information on the new ‘Extended Asynchronous Data Mover (EADM) Activity Report’ related to the EADM facility for EADM devices or subchannels:

https://www-40.ibm.com/servers/resourcelink/svc00100.nsf/pages/zOSV2R5SC342665/$file/erbb500_v2r5.pdf

EADM subchannels are like IO subchannels (devices) in the way that IO requests can be issued. However, they do not have channel paths or device numbers assigned, and they are not defined in the IO configuration. They are created automatically during IPL. For this reason, the presence of EADM devices in an IBM z16 configuration can impact the ‘% IOP Busy’ numbers.

© Copyright 2022-23 IBM Corp.

Significant SAP Capacity Increments in IBM z16

With the new PCIe Gen4 links for the PCIe Gen3 IO fanout cards and the massive redesign of the CPC drawer architecture, the trend in SAP capacity growth of IBM z16 vs. IBM z15 is very different compared to previous growth trends of IBM z15 vs. z14. Specifically, focus has been placed on significantly boosting the IO (SAP) capacity of a single drawer IBM z16 machine configuration.

It should also be noted that compared to the previous IBM z15 FEx16SA whitepaper released in January 2020, for the purpose of this IBM z16 FEx32S SAP Capacity addendum, an updated and more efficient benchmarking setup was introduced and fine-tuned individually for the IBM z15 and IBM z16 machine models. All performance data shown here was captured with this updated benchmark, due to above-mentioned changes.

The maximum IO processing capability on an IBM z16 with 5 SAPs (Max39, 1-drawer model) is over *2.3 M IOs/sec*, a significant 40% increase compared to an IBM z15 with 4 SAPs (Max34, 1-drawer model). This increase was measured when running 4 z/OS (version 2.4) LPARs.

The maximum number of drawers offered for IBM z16 has been brought down to 4 physical drawers from 5 physical drawers in IBM z15. Hence, SAP capacity of the IBM z16 Max200, 4-drawer model cannot be compared to the IBM z15 Max190, 5-drawer model.

The maximum IO processing capability on an IBM z16 with 24 SAPs (Max200, 4-drawer model) is over *5.6 M IOs/sec*, a 12% increase compared to an IBM z15 with 16 SAPs (Max145, 4-drawer model). This increase was measured when running 20 z/OS (version 2.4) LPARs.

All results in this whitepaper were obtained with a configuration having a mix of FICON Express channels of different generations/speeds, and each channel executing zHPF small block IO operations of 4Kbytes size.

The IOs/sec numbers stated above were obtained at approximately 99% SAP (IOP) utilization levels. However, for customer production environments and data-center operations, IBM recommends a maximum SAP utilization limit of 70%.

Figure 8 below shows the relative increase in IO (SAP) capacity of various IBM z16 feature models over the IBM z15 models shown. The absolute values of the bars in ‘IOs/sec’ unit are shown at 70% SAP (IOP) utilization levels.

© Copyright 2022-23 IBM Corp.

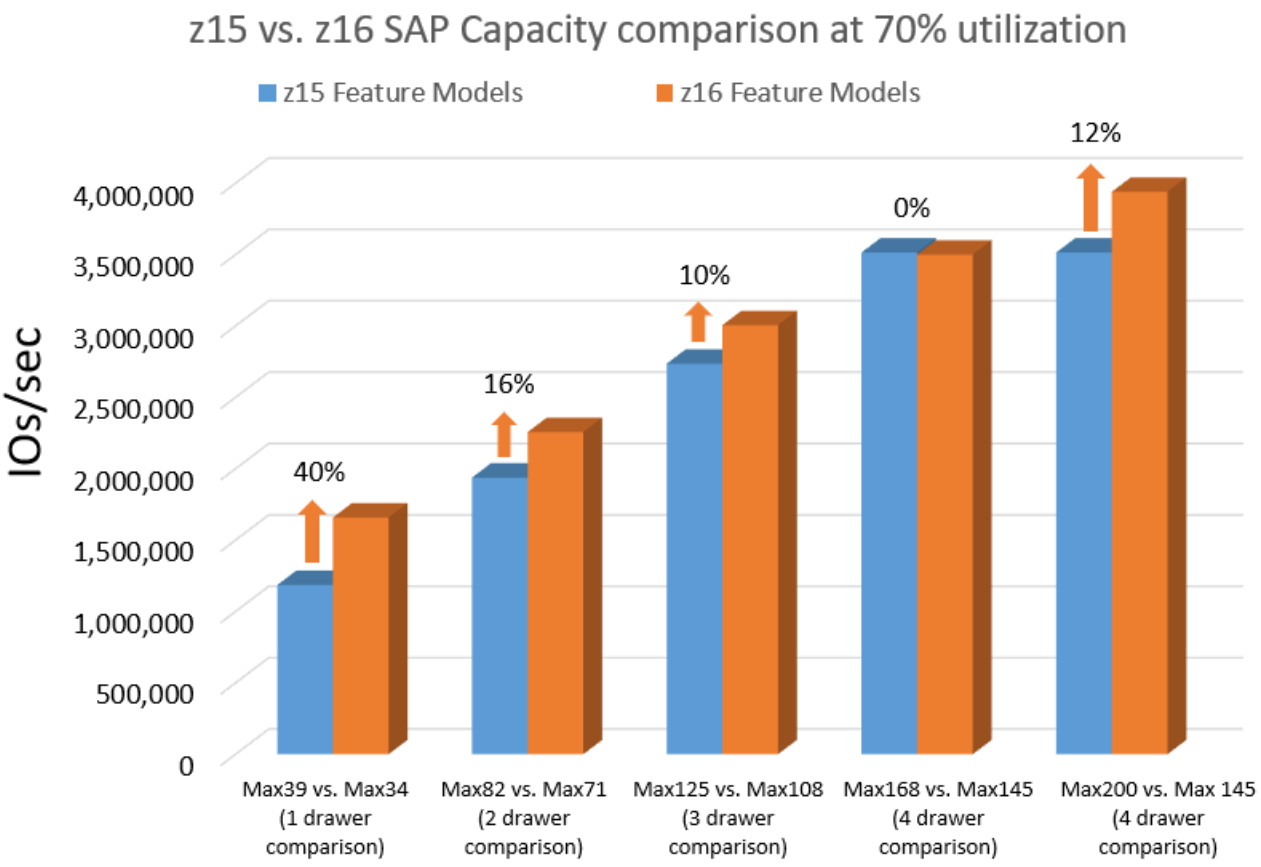

Figure 8 – SAP Capacity improvements at 70% SAP utilization levels for IBM z16 Max39, Max82, Max125, Max168 and Max200 models. Notably, 1-drawer IBM z16 vs. IBM z15 is up by a fantastic 40% due to complete redesign of IBM z16

As shown in Figure 8 above, at *70% SAP (IOP) utilization levels:*

- *IBM z16 with 5 SAP cores (Max39, 1-drawer model) attained SAP capacity of 1.66 Million IOs/sec with 4 LPARs. This is a significant improvement, 40% higher than SAP capacity of IBM z15 with 4 SAP cores (Max34, 1-drawer model).*
- *IBM z16 with 10 SAP cores (Max82, 2-drawer model) attained SAP capacity of 2.26 Million IOs/sec with 8 LPARs. This is a significant improvement, 17% higher than SAP capacity of IBM z15 with 8 SAP cores (Max71, 2-drawer model).*

© Copyright 2022-23 IBM Corp.

- *IBM z16 with 15 SAP cores (Max125, 3-drawer model) attained SAP capacity of 3.0 Million IOs/sec with 12 LPARs. This is 10% higher than SAP capacity of IBM z15 with 12 SAP cores (Max108, 3-drawer model).*
- *IBM z16 with 20 SAP cores (Max168, 4-drawer model) attained SAP capacity of 3.5 Million IOs/sec with 16 LPARs. This is comparable to the SAP capacity of IBM z15 with 16 SAP cores (Max145, 4-drawer model).*
- *IBM z16 with 24 SAP cores (Max200, 4-drawer model) attained SAP capacity of 3.94 Million IOs/sec with 20 LPARs. This is a significant improvement, 12% higher than SAP capacity of IBM z15 with 16 SAP cores (Max145, 4-drawer model).*

System Dynamics: Variations and Scaling

One of the major components of SAP capacity and performance analysis of a new System Z machine as a whole box, is to see the scaling trends in utilization levels of LPAR CPs, internal PCIe bus, FICON Express channels and IO devices when SAP (IOP) cores are driven to utilization levels of IBM's recommended 70%. These scaling trends indicate how the system dynamics of the whole machine will evolve and perform when IO workload intensity on the SAP cores is increased in a ramp-up fashion. Keeping this concept in mind, here we observe scaling trends for parts of the new IBM z16 other than SAP cores, when the SAP cores themselves are ramped-up for their IOs/sec operational intensity. Figure 9 below shows 'Average Busy Utilization (%)' scaling trends of:

- CP cores of the 4 LPARs A, B, C and D
- PCIe bus
- FICON Express channels
- IO devices

as the SAP cores (IOP threads) of a Max39 1-drawer model (with 4 CPs per LPAR) are performance tested under the IBM controlled environment, configuration, and setup of the above-listed entities. A sufficient number of these entities/resources (infrastructure to drive SAPs) were configured to prevent them from being a 'first' bottleneck. We see that throughout the operation range of 1 to 1.66 Million IOs/sec, until 70% SAP utilization levels are reached, the SAP cores themselves scale in a very linear fashion. The linear scaling is better than any non-linear (quadratic, cubic, any function that scales faster than linear) or exponential scaling. This observation validates the *robust* performance of SAP cores of the new IBM z16 in handling high

© Copyright 2022-23 IBM Corp.

amounts of IO workload assigned to them by the LPAR CPs and the entire IO subsystem of IBM z16.

Moreover, in Figure 9 we also observe the *fantastic* 'close-to linear scaling' of the other parts of the machine setup (listed above). This observation separately validates very *robust* and *durable* performance-scaling of the entire IBM z16 machine box when SAP utilization levels reach 70% in a ramp-up and linear fashion.

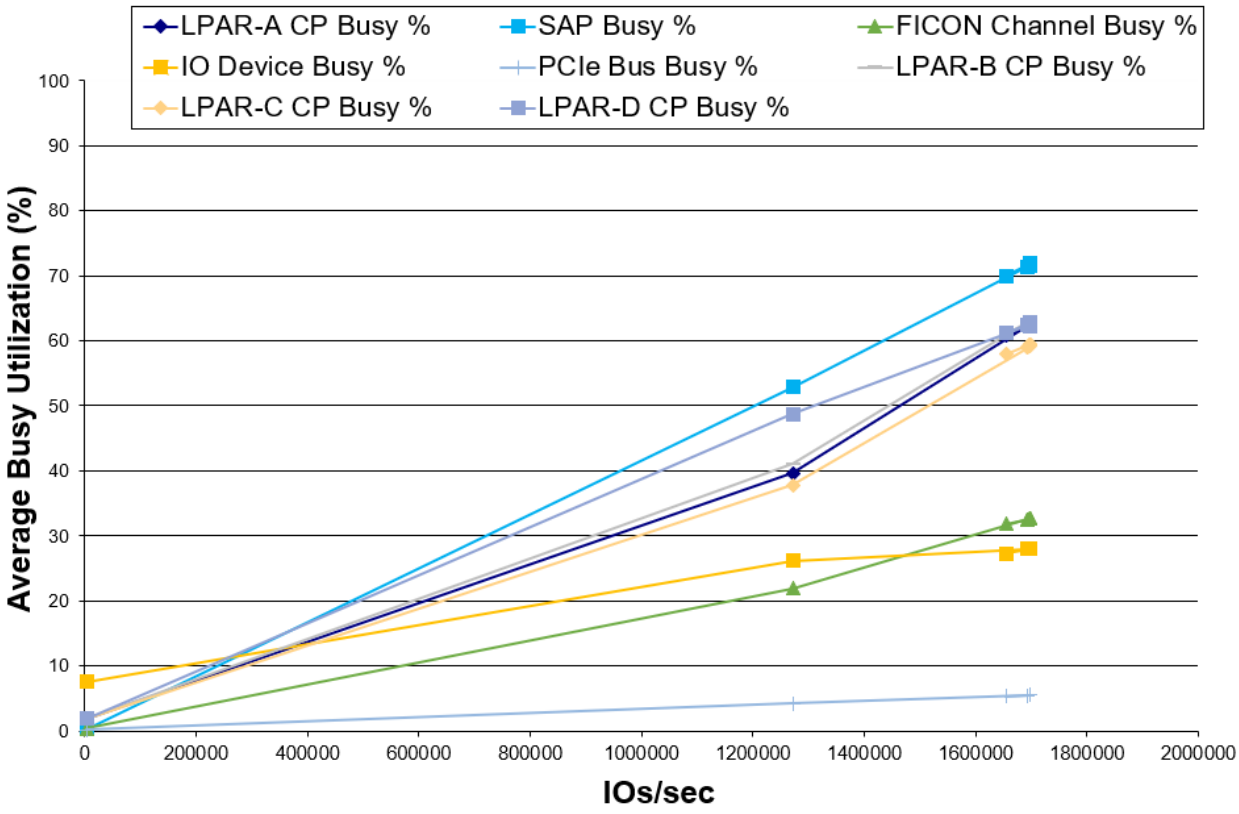

Figure 9 – Max39 (1-drawer) model's fantastic 'close-to linear scaling' of the CPs of 4 LPARs, System Assist Processors (SAPs), PCI Express (PCIe) bus, FICON channels and IO devices when IOs/sec ramp-up to 1.66 Million IOs/sec at 70% SAP utilization

Conclusions and Recommendations for the Customer

Here are some conclusions and trade-off insights about the IO and SAP capacity performance of IBM z16 and some general recommendations for the customer. The customer must think and verify 'if and why' these recommendations are applicable to their situation:

© Copyright 2022-23 IBM Corp.

- The new IBM z16 has more of the CP and cache nest performance packed in the same physical space of a given single drawer configuration, when compared to IBM z15.
- CP cores of the Telum chip of IBM z16 are more powerful than in IBM z15 for processing IO operations. Hence, for IO intensive workloads, customer should avoid using more CPs than required for the best IO performance of an LPAR.
- The capacity of a single zOS (version 2.4) LPAR executing zHPF small block (4K bytes) IO operations on the IBM z16 with 5 SAPs (Max39 1-drawer model) is now 1.2 Million IOs/sec. This is a significant increase of 26% when compared to 950K IOs/sec in IBM z15 (running zOS version 2.3 LPAR). Thus, based on IBM's internal benchmark used, more LPARs can be run in an IBM z16 single drawer configuration to attain greater IO throughput levels than IBM z15.
- The IO or SAP capacity of a given single drawer and multiple drawer configuration of IBM z16 is highly dependent on:
    - the Input/Output Definition File (IODF),
    - the number and generation of FICON Express channels used,
    - the IO domains in use,
    - the SAN switch technology,
    - the conditions and signal strengths of the SAN switches and links (errors produced),
    - the number of IO devices of the external storage system used, and
    - the transmission speed and technology of the host adapter (16G, 32G, etc.) of the external storage system.
- The customer is strongly recommended to keep the size of an LPAR, in terms of number of CPs, such that it fits within the scope of a single CPC drawer.

© Copyright 2022-23 IBM Corp.

Copyright IBM Corporation 2022-23
IBM Corporation
New Orchard Rd.
Armonk, NY 10504
U.S.A
Produced in the United States of America
2/2023
All Rights Reserved
....
PAI/O is a trademark of Performance Associates, Inc.