

When to zIIP your workloads

Get the most from IBM Z by leveraging z Integrated Information Processors (zIIPs)



How are z Integrated Information Processors (zIIPs) beneficial?

A zIIP is a dedicated processor designed to operate asynchronously with the general processors in an IBM Z® system to help improve utilization of computing capacity and control costs. A zIIP is designed for select data (database), cloud and transaction (Java™) processing workloads.

zIIPs allow users to leverage additional processing power for new applications without affecting their total million service units (MSU) rating or machine model designation.

When workload is dispatched onto a zIIP processor, the processing cycles used do not contribute to the MSU count and therefore do not impact software usage charges. This makes adding new applications to the IBM Z extremely cost-effective especially compared to competing platforms such as distributed systems or public cloud.

While zIIPs exploitation for individual workloads may vary, hardware configuration updates to allocate zIIPs are minimal. Additionally, zIIP usage is transparent to the applications; no changes are needed within the application to run a zIIP eligible workload. Learn more about setting up zIIPs in the reference section of this paper.

Common workload scenarios

What workloads should run on zIIPs? Here are some suggested scenarios to exploit zIIP technology.

Db2® Distributed Relational Database Architecture (DRDA®)

Many IT organizations keep their critical system of record data on IBM Z and leverage other platforms for their system of engagement applications. An example of this is a web enabled customer service application that retrieves customer records from a master Db2 database on IBM Z.

If the Java Database Connectivity (JDBC) in TCP/IP workloads in which data requests from, and delivers to, another platform is performed through TCP/IP on general purpose CPs, it will contribute to MSU consumed and impact the software billing. Instead, zIIPs could be used to perform this task, thus reducing general purpose CP usage and related software billing charges. Most DRDA workloads can lower their general compute power by offloading at least 60% of their computing to zIIPs.¹

Gateway to Hybrid Cloud

IBM zIIPs can enable clients to rapidly and inexpensively transform, extend and deploy applications into a hybrid cloud environment. New applications can access cloud native services through z/OS Connect RESTful APIs using zIIPs. Modernizing existing z/OS® applications for a hybrid cloud environment can be started quickly using IBM z/OS Container Extensions and IBM zIIP processors, z/OS Connect Enterprise Edition for faster speed to value.

¹ https://www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/perf/src/tpc/db2z_authorizedziipuse.html

z/OS Container Extensions

z/OS Container Extensions (zCX) that are part of IBM z/OS V2R4 or higher, are designed to enable the ability to run almost any Docker container that can run on Linux® on IBM Z in a z/OS environment alongside existing z/OS applications and data without a separate provisioned Linux server. This extends the strategic software stack on z/OS, enabling developers to build new, containerized applications using Docker and Linux. zCX containers can be deployed directly onto z/OS, streamlining system administration.

Key benefits of zCX include:

- Integrating zCX workloads into existing z/OS workload management, infrastructure, and operations strategy thus taking advantage of z/OS strengths such as pervasive encryption, networking, high availability, and disaster recovery.
- Architecting and deploying a hybrid solution that consists of z/OS software and Linux on IBM Z Docker containers on the same z/OS system.
- Open access to data analytics on z/OS by providing developers with standard open API-compliant RESTful services.

z/OS Container Extensions address spaces are zIIP eligible and most zCX processing, including Docker containers deployed within zCX, can be dispatched on available zIIP processors on IBM z16™, IBM z15™ systems and z14®.² While some zCX processing will be dispatched onto standard processors, it is generally a much smaller percentage of processor cycles compared to what can be executed on zIIPs. An IBM internal test using the Acme Air benchmark on zCX, found up to 98% of the zCX CPU consumption was measured to be zIIP eligible.³

98%
of IBM z/OS
Container
Extensions CPU
consumption is zIIP
eligible³

IBM z/OS Connect EE

**Over
99%**
of z/OS Connect
Enterprise Edition is
zIIP eligible⁴

IBM z/OS Connect Enterprise Edition (z/OS Connect EE) provides RESTful API access to z/OS applications and data hosted in subsystems such as CICS®, IMS™, IBM MQ®, and Db2. It provides concurrent access, through a common interface, to multiple z/OS subsystems. z/OS Connect EE also provides the capability to allow z/OS-based programs to access any RESTful endpoint, inside or outside z/OS, through RESTful APIs with JSON formatted messages. The framework provides support for CICS, IMS, IBM MQ, and z/OS applications to call RESTful APIs.

² <https://www.ibm.com/support/z-content-solutions/container-extensions/>

³ Results were extrapolated from internal IBM benchmarks performed in a controlled environment using a single z14 z/OS 2.4 LPAR with TCP/IP inbound workload queuing (IWQ) for inbound traffic and two zCX containers: one running Node.js and one running a MongoDB database. zIIP eligibility is based on the CPU consumption of the work running on the zCX address spaces and the associated work on the TCP/IP and VTAM address spaces. Results may vary.

through z/OS Connect. z/OS Connect EE is a Java-based product and so over 99% of the product is eligible to be offloaded to zIIP.⁴

z/OS Connect EE and z/OS Container Extensions can accelerate the journey to hybrid cloud by extending and modernizing existing assets easily and affordably.

Data Virtualization Manager for z/OS

IBM Data Virtualization Manager for z/OS provides simplified virtual, integrated views of relational and traditional non-relational data residing on IBM Z. It enables users and applications read/write access to IBM Z data in place, without having to move, replicate or transform data.

With Data Virtualization Manager for z/OS you can...

- Readily access relational and non-relational IBM Z transactional data
- Access and update live IBM Z data via modern APIs such as SQL and RESTful APIs (when combined with z/OS Connect)
- Reduce the cost and delay of moving data to non-Z platforms
- Simplify development of applications accessing relational and non-relational data types including Db2, VSAM, IMS, ADABAS®, IDMS®, SMF and non-IBM Z data sources
- Modernize existing applications and reduce data movement off platform
- Reduce data copies that result in latency, cost and risk (security, governance, decision latency)

Up to
99%
of Data
Virtualization
Manager for z/OS
data transformation
activity is zIIP
eligible⁵

Data Virtualization Manager for z/OS can directly access z/OS data sets for bulk data access that bring significant MLC savings resulting from up to 99% of zIIPs eligible data transformation activity.⁵ The amount of zIIP offload activity will vary with different types of workloads depending on the data source, access method, and effective SQL operations in use.

Java for new or existing applications


Modernizing existing applications with Java code or writing new applications in Java allows customers to run their workloads on IBM zIIPs and enables them to leverage the legendary benefits of IBM Z scalability, security, performance and resiliency at a lower cost than running the application on a GP or other platforms. Co-locating these new applications with the systems of record on IBM Z also reduces latency and allows for the sharing of system resources for great efficiency and manageability. Many web application server workloads are good candidates for zIIPs, particularly if users already

⁴ IBM z/OS Connect Enterprise Edition Performance Reports <https://www.ibm.com/support/pages/node/585931>

⁵ <http://www.redbooks.ibm.com/redpapers/pdfs/redp5523.pdf>

have CICS and IMS. CICS and IMS applications can be converted to Java using a WebSphere® Liberty license that is already provided in CICS and IMS middleware.

1. WebSphere Liberty licenses in CICS or IMS are available at no extra cost
2. Existing WebSphere licenses on x86 can be freed up since the workload will move to z/OS



Again, if these workloads use data that originates on z/OS, applications executing on zIIPs will have the added benefit of improved performance with lower latency compared to running them on another platform. When the application is co-located with the data, data avoids moving over TCPIP and a physical network which can introduce delays.

Other use cases for zIIPs

Not only can GP cycles be offloaded to zIIPs for certain workloads. zIIP technology can be used to reduce the cost of IT operations costs and data processing. Here are some scenarios in which zIIPs can be used to lower GP usage resulting from encryption, data transmission and analytics.

Security network encryption offload to zIIP

With increasing security concerns IT organizations have begun encrypting data over the network in addition to encrypting data when at rest. Encrypting data over the network however can consume general processor usage, and MLC software costs. zIIPs can be used to process security (IPSec) network encryption and decryption. By running some of the IPSec network encryption instructions on zIIPs, organizations can reduce their general purpose CP usage and lower software costs.

System Data Mover of z/OS (SDM)

System Data Mover of z/OS and z/OS base elements properly invoked by the SDM is zIIP eligible. This includes DFSMS SDM of z/OS processing associated with zGM/ERC, including z/OS base elements properly invoked by z/OS DFSMS SDM.

⁵ [https://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/2cbf007bef04d4fe852580aa007f42e7/\\$FILE/Liberty%20zOS%20-%20Overview%20-%20CHARTS.pdf](https://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/2cbf007bef04d4fe852580aa007f42e7/$FILE/Liberty%20zOS%20-%20Overview%20-%20CHARTS.pdf)

XML processing

Extensible Markup Language (XML) is a simple, flexible text format to be used in data interchange on the Web and in other electronic communications. It is extendable with tags in a natural language hence the information structure can be easily discerned by both humans and computers. The data structure follows a noticeable and useful pattern, making it easy to manipulate and exchange data.

A number of XML-enabled products for IBM Z use z/OS XML System Services behind the scenes to gain advantage from its high-performance and consistent parsing. These include COBOL, PL/I, Java, CICS, IMS, Db2 pureXML®, etc. z/OS XML System Services provides the option to offload most of the XML parsing to zIIPs, enabling less GP usage.

Machine Learning (ML)

ONNX

Library calls and compiling AI models on Open Neural Network Exchange (ONNX) where the ONNX operators are defined to run directly on z/OS are zIIP eligible workloads.

Python

Up to 70% of Python AI and ML workloads are considered zIIP eligible.

Z AI Data Embedding library of z/OS when invoked by using the Java native application programming interface is considered zIIP eligible.

IBM Machine Learning for z/OS uses zIIP engines to bring artificial intelligence to transactional applications by offering an end-to-end machine learning platform that operationalizes predictive models. Additionally, IBM Open Data Analytics for z/OS (IzODA) can use zIIPs to run analytics and machine learning on very large, complex data sets.

IzODA allows you to federate your IBM Z environment with your existing analytics infrastructure across platforms to bring compute capabilities to your original data sources. This can avoid unnecessary data movement and duplication of data. It also enables users to enforce the privacy and security policies of their enterprise.

Capacity for System Recovery Boost

zIIPs can provide additional processing power to System Recovery Boost. This IBM z16 and z15 feature reduces the impact of downtime for both planned and unplanned outages. By providing additional boost capacity in the form of zIIPs, users can accelerate partition restart and catchup on workload backlog without charges for temporary incremental processing power or additional IBM software cost. An IBM test found IBM System Recovery Boost provided additional capacity (typically handled by

OOCoD during normal operations) to process a backlog during LPAR recovery twice as fast at no additional hardware or IBM software costs with an estimated savings of \$6.68 million over 3 years versus buying Capacity on Demand to speed up recovery.⁷

System Recovery Boost for SVC dumps, Middleware starts/restarts for client-selected middleware regions and HyperSwap Config Load is zIIP eligible on the IBM z16 family of systems and subsequent generations of systems.

z/OSMF offload to zIIP

zIIPs can also be leveraged in a z/OS Management Facility (z/OSMF) framework. Just as programmers and operators are able to update and access the z/OS operating system from the z/OSMF user interface that consumes general purpose CPs, users can offload much of this work to zIIPs. Depending on workload demands and a system's four hour rolling average, leveraging zIIPs can reduce MLC software costs. Up to 95% zIIP eligibility can be attributed to z/OSMF Java workloads.⁸


ISV software solutions

zIIPs specialty engines enable a vast number of independent party vendor (ISV) solutions to benefit from zIIPs savings. Many of these ISVs not only support running their solutions in zIIPs but exploit zIIPs technology to reduce CP consumption, resulting in lower IBM Z costs for their customers. Learn more about ISV IBM Z solutions in the [IBM Global Solutions Directory](#).

zIIPs can lower IT costs

zIIP technology allows a variety of workloads to exploit enterprise proven quality of services and the efficiencies of IBM Z at a significantly lower cost compared to workloads that rely on general compute power exclusively. Activating zIIPs on IBM Z systems can drastically minimize general purpose CP usage and MLC costs since existing Z workloads can be reconfigured to leverage zIIPs.

3 Year Total Cost of Ownership
for z15 T02 zIIPs versus x86 servers



⁷ This claim is based on IBM internal tests. Results may vary. The use case for this test is an enterprise running a transactional workload on z/OS with a service level agreement (SLA) of 24 million CICS transactions per hour. A 30-minute planned outage results in a backlog of 12 million transactions. The system is running z/OS on a z15 and configured to use 3 general processors (GPs) and 1 IBM z Systems Integrated Information Processor (zIIP). During IBM System Recovery Boost, the system uses 3 GPs and 8 zIIPs, including 7 additional zIIPs taken from the zIIP processor pool. Hardware costs are extrapolated based on the list price of Capacity on Demand (CoD), and software costs are based on a typical software acquisition bill of \$500/MIPS. Software costs are taken from IT Economics observations made during customer engagements. These are blended values and do not represent any particular circumstance

⁸ [http://www.03.ibm.com/support/techdocs/atsmastr.nsf/c6192fb3a432612485256d970082de57/04aacd547b6fa019862577bb001300f7/\\$FILE/WP101779_ZSW03197-USEN-00_V2.1.pdf](http://www.03.ibm.com/support/techdocs/atsmastr.nsf/c6192fb3a432612485256d970082de57/04aacd547b6fa019862577bb001300f7/$FILE/WP101779_ZSW03197-USEN-00_V2.1.pdf)

zIIPs can also help lower the cost of IT for distributed server environments. Java x86 workloads can be re-platformed to IBM Z to utilize zIIPs. By moving x86 workloads to IBM Z, distributed server footprint and administration overhead is decreased. More workloads can be run and administered from a single system with lower software license costs.

In a modeled cost comparison scenario for a Java application on zIIPs under z/OS CICS, on an IBM z16 or z15 zIIP environment can offer 77% savings in hardware, software and facilities costs over three years versus the workload on compared x86 servers.⁹ The chart and table illustrate the results on an IBM z15 T02 and would be equivalent on an IBM z16. By deploying the workload on an existing IBM Z system, distributed software and hardware costs are avoided. The Java workload runs IBM WebSphere Liberty, available with CICS, in the same LPAR as the database and offloads zIIPs eligible workload to purchased zIIPs. Due to the efficiencies of offloading eligible workload to zIIPs, the IBM Z scenario costs significantly less than the x86 scenario in which additional x86 servers and associated distributed software need to be purchased and deployed.

Cost Categories in x86 versus IBM Z 3 Year TCO Model

Costs	x86	z15 T02
Hardware	2 48-core x86 servers	3 zIIPs
Operating System	Red Hat® Enterprise Linux per socket software subscription and support	z/OS® per MSU
IBM WebSphere® Liberty	Per core licensing, software subscription and support	Included with CICS
System management and monitoring, virtualization software	Per core licensing, software subscription and support	SW per MSU
Floor space	Per square foot	Part of existing frame
Energy	Per kWh	Part of existing frame

Leverage zIIPs in your IT organization

⁹ An IBM IT Economics model was used to examine operating costs over three years for hardware, software, energy and floor space costs for a Java application being deployed on two x86 Cascade Lake 48-core servers (for a total of 96 cores) versus being deployed on an z15 T02 using three zIIPs. Data used by the Java workload resides on the z15 T02 in the same LPAR as the Java workload running on z/OS CICS. Eligible zIIPs workload is calculated at 80% based on IBM test results, [https://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/2cbf007bef04d4afe852580aa007f42e7/\\$FILE/Liberty%20zOS%20-%20Overview%20-%20CHARTS.pdf](https://www-03.ibm.com/support/techdocs/atsmastr.nsf/5cb5ed706d254a8186256c71006d2e0a/2cbf007bef04d4afe852580aa007f42e7/$FILE/Liberty%20zOS%20-%20Overview%20-%20CHARTS.pdf). The Java workload, on both x86 and z/OS, is running in WebSphere Liberty. IBM Z pricing, where applicable, is based on U.S. prices as of 07/01/2020 from our website and x86 hardware pricing is based on IBM analysis of U.S. prices as of 07/01/2020 from IDC. Software, energy and floor space costs are based on data from IBM IT Economics assessments for clients. For additional information on the use case model, contact the IBM IT Economics Team at IT.Economics@us.ibm.com.

If your IT organization is looking for ways to increase efficiencies and lower costs, ask for a no-charge IT Economics assessment at IT.Economics@us.ibm.com . IT Economics analysis can help identify workload candidates for zIIPs usage, lower software charges and reduced infrastructure costs.

Other resources about zIIPs

Move to hybrid cloud, accelerate digital transformation and implement new workloads on [IBM z Systems Integrated Information Processors](#)

Refer to the [IBM z Integrated Information Processor Exploiter \(zIIP\) List](#) for a compilation of known IBM software offerings that can exploit IBM zIIPs

How to set up and run zIIPs, [Planning considerations for running zIIPs](#)

About the authors

John Gustavson

CTO, Worldwide IBM IT Economics Consulting & Research

jgustav@us.ibm.com

Dr. J C Yao,


Chief Strategist, Worldwide IBM IT Economics Consulting & Research

jcy@us.ibm.com

Edward L. Linde II

Principal Product Manager, IBM zNext

ell11@us.ibm.com



*Copyright IBM Corporation 2023

IBM Corporation

New Orchard Road

Armonk, NY 10504

U.S.A.

Produced in the United States of America,

06/2023

IBM, ibm.com, IBM logo, IBM Z, IBM z16, CICS, Db2, DRDA, IMS, MQ, pureXML, WebSphere, z14, z15, and z/OS are trademarks or registered trademarks of the International Business Machines Corporation.

A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

InfiniBand and InfiniBand Trade Association are registered trademarks of the InfiniBand Trade Association.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the [OpenStack website](#).

Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

RStudio®, the RStudio logo and Shiny® are registered trademarks of RStudio, Inc.

TEALEAF is a registered trademark of Tealeaf, an IBM Company.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Worklight is a trademark or registered trademark of Worklight, an IBM Company.

Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.

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

The information contained in this documentation is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this documentation, it is provided "as is" without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this documentation or any other documentation. Nothing contained in this documentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM (or its suppliers or licensors), or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

References in these materials to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in these materials may change at any time at IBM's sole discretion based on market opportunities or other factors and are not intended to be a commitment to future product or feature availability in any way.

28035028USEN-04