

Saving the MSUs used by zIIP-eligible workloads

Fabio Massimo Ottaviani – EPV Technologies

September 2026

1 Introduction

In 2006 IBM released the zIIP (z Integrated Information Processor), a specialty processor initially dedicated to run enclave SRB work.

In the years since, many other workloads have been made completely or partially zIIP-eligible; the most relevant are:

- any Java-based application;
- DDF requests from external network clients via TCP/IP querying the local Db2 on z/OS;
- most of DB2 DBM1 activities;
- Db2 MSTR log activities;
- Db2 parallel child tasks generated by complex SQL queries (e.g., Data Warehousing or Business Intelligence);
- Db2 utilities and statistics gathering;
- Python-based applications;
- z/OS XML System Services for parsing large XML files;
- Docker containers and Linux workloads running in z/OS Container Extensions (zCX).

However, eligibility is a necessary but not sufficient condition. A zIIP engine must be available with an acceptable delay, according to the “zIIP needs help” mechanism when the workload request arrives.

Fully leveraging zIIP technology is a key element of the strategy for any customer aiming to reduce hardware costs and, above all, software costs especially when the TFP Software Consumption Solution pricing option has been adopted and every MSU counts.

The goal must be to minimize the portion of zIIP-eligible workload that spills over to standard CPUs because a zIIP engine was not available at the time.

In this document, following a brief description of the “zIIP needs help” mechanism, we will discuss zIIP capacity planning, the Simultaneous Multi-Threading (SMT) technology, and the specific parameters you can configure to achieve this goal.