

Query Parallelism – Enhanced with zIIP Integrated Processors



“... It’s not just about adding more processors to a resource pool. It is, more importantly, about the ability to get more complicated, interrelated business processing done, when it needs to be done, and at an incrementally lower cost. This is what zIIP does for DB2 workloads.”

Highlights

- **Help improve price / performance of parallel processing on IBM System z™ with the IBM System z9™ Integrated Information Processor (zIIP)**
- **zIIP can lower general purpose processor usage for eligible workloads, which can result in lower overall costs**
- **Leveraging the zIIP processor can enable you to exploit centralized concurrent sharing of data by On-Line Transaction Processing (OLTP), On-Line Analytical Processing (OLAP) and Batch applications, reducing the need to manage multiple databases**
- **Eligible zIIP workload engines do not have IBM software charges associated with their capacity**

Redirecting query processing to zIIP processors

The distinction between data warehousing and OLTP applications is blurring due to dynamic business requirements and the mission critical nature of business focused workloads. Data warehousing and analytic applications are increasingly accessing operational data or near real-time data.

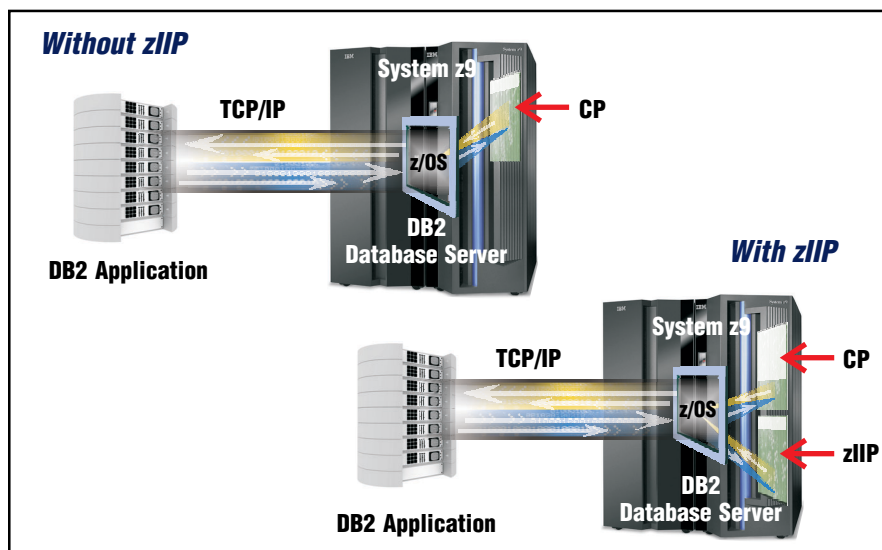
With data at the core of today’s most critical business issues, the IBM mainframe equipped with a System z9 Integrated Information Processor (zIIP) can help you orchestrate information resources in an economical way.

Historically, System z has been the highly available “best of breed” server for a variety of processing requirements, from transaction based OLTP and Batch, to query and reporting workloads. The addition of the zIIP engine further strengthens the System z9 position as the data serving hub of an enterprise.

The zIIP is an optional specialty engine available on the System z platform that enhances the integration of other workload types such as Data Warehousing (DW) or Business Intelligence (BI) on to the platform. The ability of z/OS® to redirect a significant portion of the parallel query processing to the zIIP processors enhances the overall cost effectiveness of the System z, as IBM does not impose IBM software charges on eligible zIIP capacity.

Query Parallelism

Warehousing applications include workloads that are typically composed of queries with widely varying elapsed times that result from accessing large amounts of data. DB2® has long had the capability of splitting up that work to take advantage of multiple processors that are available to the workload. DB2’s likelihood of breaking this into separate segments increases as queries access more data. This segmentation is referred to as query parallelism.



These segments can then be executed on any engine available to the DB2 workload. Spreading the processing over multiple engines can significantly reduce the query elapsed time. Adding zIIP engines to the pool of processors available to the DB2 workload enables the use of more cost-effective cycles to handle the processing.

Parallel queries that are typical of Business Intelligence workloads are well suited to exploit the zIIP engines. To demonstrate this, a study was completed within IBM in which parallel query workloads were executed on two configurations with varying processors available to the workload. The first configuration consisted of two general purpose engines; the second configuration consisted of one general purpose engine and one zIIP engine. The results from this study are summarized in the figures below.

zIIP Capabilities

The zIIP engine has the same functionality as a general purpose processor on a System z platform. This zIIP engine can be used to redirect processing for a variety of workloads from the general purpose processors, including TCP/IP DRDA, parallel queries, DB2 utilities, IPSec, and z/OS XML services.

zIIP:

- Supports On/Off Capacity on Demand (On/Off CoD), Capacity Upgrade on Demand for the nondisruptive addition of one or more zIIPs and Capacity BackUp (CBU) for emergency situations.
- Provides a full capacity processor independent of the capacity of the general purpose processors.
- Is designed to operate asynchronously with general purpose processors.
- Is managed by PR/SM™ in a logical partition with dedicated or shared processors.

- Is ordered as a hardware feature, and is unique to System z. The number of orderable zIIPs varies by the server model and configuration.

The strengths of System z facilities such as Workload Manager (WLM) allow for the logical separation and prioritization of the different workload types within the same system (LPAR) or multi-system (Parallel Sysplex®) environment containing general purpose engines and zIIPs. Service levels and priorities for multiple business units as well as their individual application processing requirements may be defined and automatically managed by WLM.

Summary

The system can successfully redirect eligible processing of query workloads from general purpose processors, while maintaining consistent overall elapsed times for a parallel query workload.

Leveraging the zIIP processor allows you to gain even more value from your System z9 investment. Now you can extend the security, availability, and workload management benefits found on the System z9 environment, to all your business data. The System z9 mainframe helps minimize the need to maintain duplicate databases and provides better security to protect both your data and your business.

In **Figure 2**, the dark bars represent the processing cycles consumed by the query workload on the general purpose processors, and the blue colored bar reflects the processing that was redirected to the available zIIP engine. Leveraging the zIIP processors lowered the processing demands on the general purpose processors for this query workload.

Figure 2:
Relative utilization of processors





Figure 3 reflects the elapsed time for the scenarios. The first bar represents the elapsed time for the workload on general processors, without using parallel processing. The second bar represents the same query workload when using parallelism on general purpose processors. The last bar represents the query elapsed time when leveraging a zIIP engine to complete the processing. Clearly, parallelism has the potential to significantly reduce the elapsed time for long running queries, and leveraging the zIIP engine maintained the improvement in the elapsed time for the workload with minimal overhead.

Figure 3:
Relative elapsed times of queries



Copyright IBM Corporation 2007

IBM Corporation
New Orchard Road
Armonk, NY 10504 U.S.A.
05/07

All Rights Reserved

References in this publication to IBM products or services do not imply that IBM intends to make them available in every country in which IBM operates. Consult your local IBM business contact for information on the products, features, and services available in your area.

IBM, the IBM logo, the e-business logo, CICS, DB2, Parallel Sysplex, PR/SM, System z, System z9, z/OS, z/VM, z/VSE and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.