



ETR vs ITR

Peter Enrico

Email: Peter.Enrico@EPStrategies.com

Enterprise Performance Strategies, Inc.

3457-53rd Avenue North, #145

Bradenton, FL 34210

<http://www.epstrategies.com>

<http://www.pivotor.com>

Voice: 813-435-2297

Mobile: 941-685-6789



z/OS Performance
Education, Software, and
Managed Service Providers



Creators of Pivotor®

Contact, Copyright, and Trademark Notices



Questions?

Send email to Peter at Peter.Enrico@EPStrategies.com, or visit our website at <http://www.epstrategies.com> or <http://www.pivotor.com>.

Copyright Notice:

© Enterprise Performance Strategies, Inc. All rights reserved. No part of this material may be reproduced, distributed, stored in a retrieval system, transmitted, displayed, published or broadcast in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Enterprise Performance Strategies. To obtain written permission please contact Enterprise Performance Strategies, Inc. Contact information can be obtained by visiting <http://www.epstrategies.com>.

Trademarks:

Enterprise Performance Strategies, Inc. presentation materials contain trademarks and registered trademarks of several companies.

The following are trademarks of Enterprise Performance Strategies, Inc.: **Health Check**[®], **Reductions**[®], **Pivotor**[®]

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries: IBM[®], z/OS[®], zSeries[®], WebSphere[®], CICS[®], DB2[®], S390[®], WebSphere Application Server[®], and many others.

Other trademarks and registered trademarks may exist in this presentation

Abstract



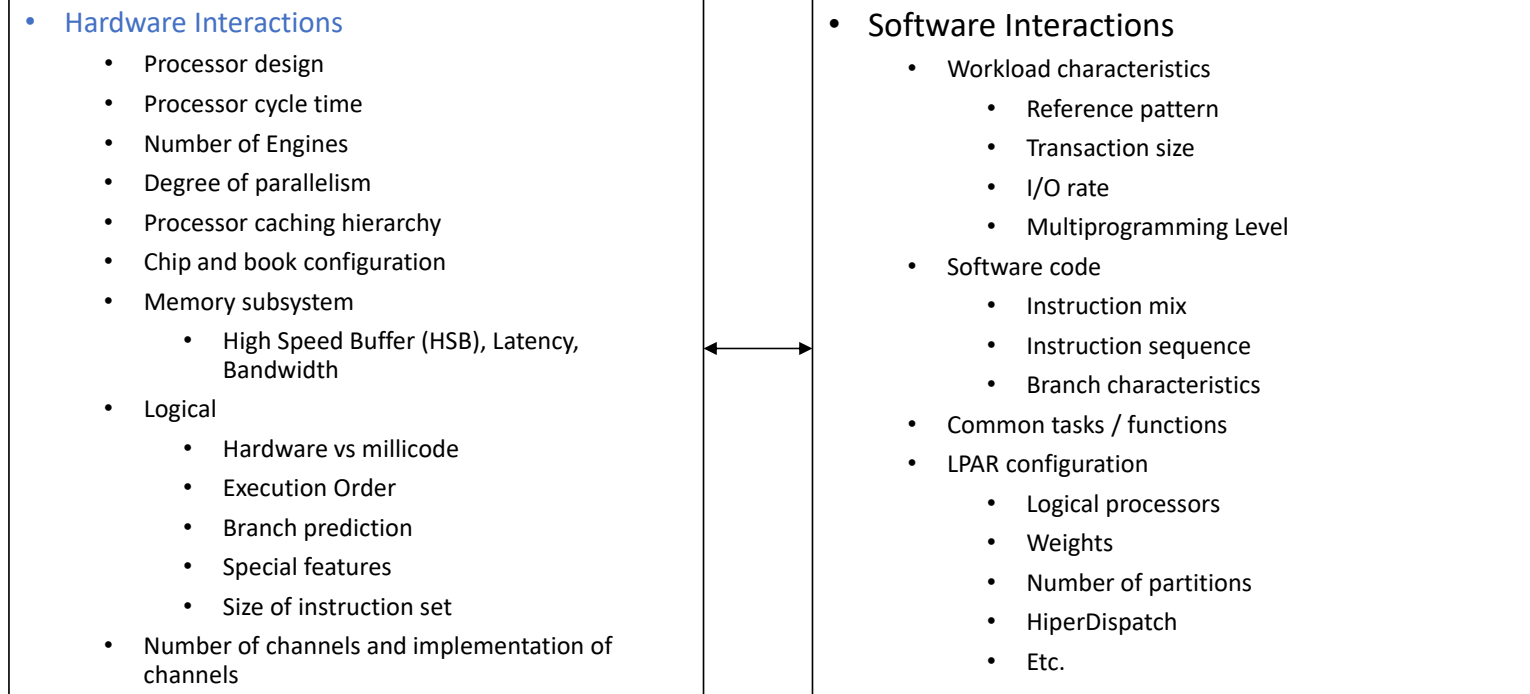
- **ETRs vs ITRs**

- The concepts and formulas for External Throughput Rates (ETRs) and Internal Throughput Rates are the foundations for IBM's evaluation of the capacity of the mainframe processors and of the Large System Performance Reference (LSPRs) tables.
- During this session, Peter Enrico will review the concepts of ETRs and ITRs, and he will show you how they are used to evaluate not just processor capacities, but also how they can be used to evaluate changes to your workloads.

Factors That Influence Processor Performance



- Workloads behave / perform differently on different processors
- Influenced by both hardware and software interactions



Throughput As A Measure of CP Power



- Given the intricacy of the processor complex and range of work to be performed
 - Traditionally 'throughput' has been a good measure of CP power

- Throughput
 - A logical measure of the amount of useful work that can be completed in a unit of time
 - Units of work – could be jobs or transactions
 - Units of time – is the time it took to complete the units of work

- Common measures or throughput (discussed in this presentation) include:
 - ETR – External Throughput Rate

 - ITR – Internal Throughput Rate

 - ITRR – Internal Throughput Rate Ratio

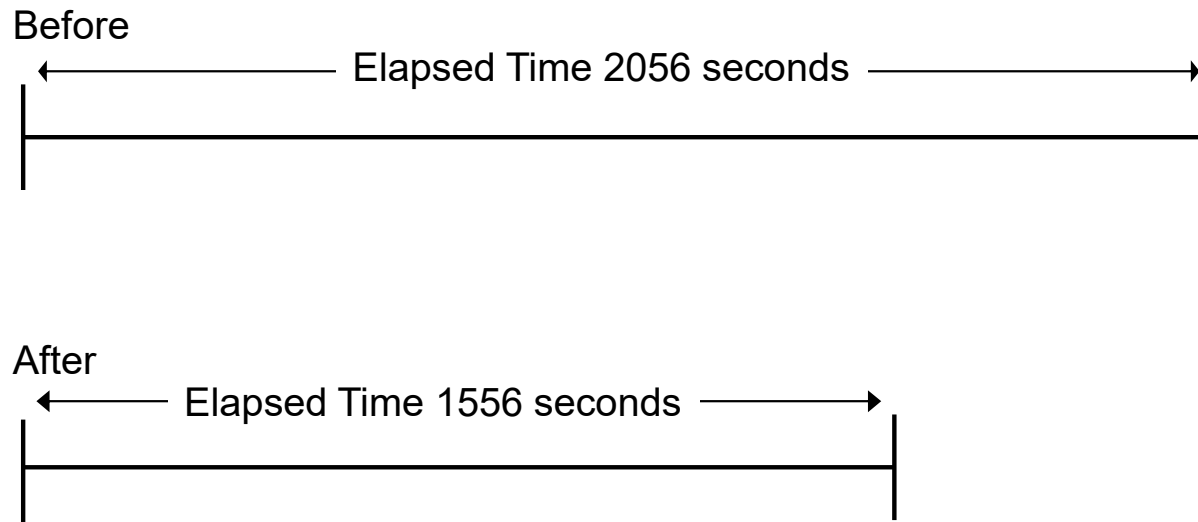
 - Control Program Constants
 - Example: z/OS SRM component has the SRM Constant
 - Not discussed in this presentation

 - MIPS – Millions of Instructions Per Second

Why Are We Interested In Throughput Rates?



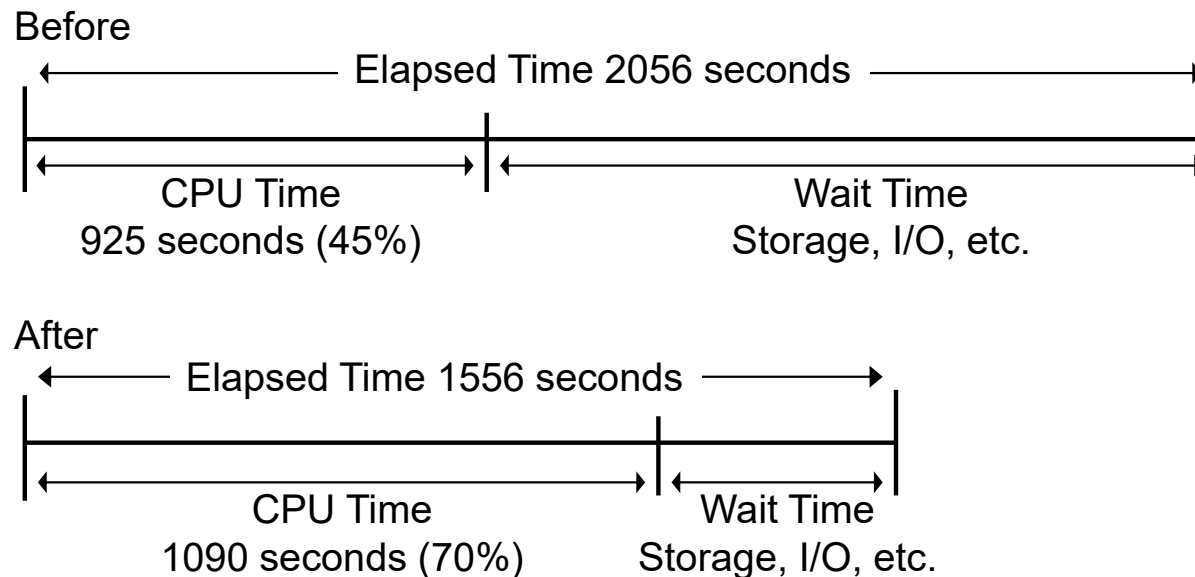
- Using the below figure, answer the following two questions:
 - Which is the better system to run the workload?
 - Which system has the better processor to run the workload?
 - Note: Assume that 1,000 transactions are run in the elapsed time



Understanding Workload Throughput



- Which is the better system to run the workload?
 - For improved transaction throughput – The After System
- Which system has the better processor to run the workload?
 - For less usage of the CPU = The Before System
- Need to consider ETRs and ITRs



External Throughput Rate - ETR



- A measure that focuses on system capacity
 - A measure of throughput as defined as by the number of transactions per wall clock second
 - Helps to answer the question”
 - ‘Which system better processes a workload’ for transaction throughput

$$\text{ETR} = \frac{\text{Units of Work}}{\text{Elapsed Time}} = \frac{\text{Units of Work}}{\text{Second}}$$

- Inverse of ETR formula is ‘Average transaction response time’

$$\text{Avg Trans Response Time} = \frac{\text{Elapsed Time}}{\text{Units of Work}}$$

- Previous Example:

- Before ETR = 1000 trans /2056 elapsed sec = 0.486 transactions / second
- 😊 After ETR = 1000 trans /1556 elapsed sec = 0.643 transactions / second

Internal Throughput Rate - ITR



- A measure that focuses on processor capacity

- A measure of throughput as defined as by the number of transactions per CPU second
- Since this is a busy time measurement, it helps to answer the question:
 - ‘Which processor better processes a workload?’
 - Useful when comparing processors

$$\text{ITR} = \frac{\text{Units of Work}}{\text{Processor Busy Time}} = \frac{\text{Units of Work}}{\text{CPU Second}} = \frac{\text{ETR}}{\text{Utilization}}$$

- Processor Time should include system overhead
 - On n-way machine, should include busy time of all processors
 - Attempts to factor in only processor as the performance factor

- Previous Example:

- 😊 Before ITR = 1000 trans / 925 CPU sec = 1.081 transactions / CPU second
- After ITR = 1000 trans / 1090 CPU sec = 0.917 transactions / CPU second

ETR vs ITR



- ETR - used to characterize system capacity
 - Since it is an elapsed time measure
 - It encompasses the performance of the processor, the operating system, and all the external resources
 - e.g. disk, cache, storage, network, operations, etc
 - All resources are potential inhibitors
 - The highest ETR achieved is the processing capability of the system
- ITR - used to characterize processor capacity
 - Since only based on CPU time
 - It encompasses the performance of just the processor
 - When measured, all external resources must be adequate
 - Thus, whenever two processors are compared, they must be measured at the same utilization
 - Could be used to evaluate the efficiency of a workloads use of CPU

ITR and ETR Relationship



- The main purpose of computing ITR is to help normalize out unequal utilizations that are represented in ETR
 - ITR is simply a measure of ETR normalized at full processor utilization

- Proof:

- Since:

$$\text{ITR} = \frac{\text{ETR}}{\text{CPU Busy \%}}$$

- Then

$$\text{CPU Busy \%} = \frac{\text{Processor Busy Time}}{\text{Elapsed Time}}$$

- Thus

$$\text{ITR} = \frac{\left(\frac{\text{Units of Work}}{\text{Elapsed Time}} \right)}{\left(\frac{\text{Processor Busy Time}}{\text{Elapsed Time}} \right)} = \frac{\text{Units of Work}}{\text{Processor Busy Time}}$$

- Thus, if you know the ETR for a workload (Avg Trans/Sec) and you know the CPU Busy %
 - Then you can easily calculate ITR



Example of Using ITR/ETR Relationship

• Since


$$\text{ITR} = \frac{\text{ETR}}{\text{CPUBusy \%}}$$

- It is simple to derive an ITR from a standard ETR without having to run a special test where the processor is the only limiting factor
 - Useful when evaluating an application, or system change
- Below example: Was a 50% increase in CPU% and a 22% decrease in efficiency of the CPU by the workload worth an 18% improvement in throughput?

	Before Change	After Change	% Change
Elapsed Seconds	900	900 ↔	
Processor Seconds	540	810 ↑	
Transaction Count	1100	1300 ↑	18%
CPU Utilization (%)	60%	90% ↑	50%
ETR	1.22	1.44 ↑	18%
ITR	2.04	1.60 ↓	-22%

Installation must decide the value of the change.

Example: CICS Transactions




This double y-axis chart shows both the number of ended transactions for CICSHIGH transactions, as well as the CP APPL% for these transactions.

As are reminder, APPL% is a measure as a percentage of 1 CPU.

The point is, this chart shows for one week compares the number of ended transactions against the CPU consumption for those transactions

Example CICS Transactions




This chart shows the three ETR and ITR values:

- ETR Transactions/Sec
- ITR Transactions/CPU Sec
- ITR CPU / Transaction

Note the stability of the ITR values over the week relative to the regular increases and decreases in the ETR.

Example: DDF Transactions




This double y-axis chart shows both the number of ended transactions for DDFHIGH transactions, as well as the CP APPL% for these transactions.

As a reminder, APPL% is a measure as a percentage of 1 CPU.

The point is, this chart shows for one week compares the number of ended transactions against the CPU consumption for those transactions

Example: DDF Transactions



This chart shows the three ETR and ITR values:

- ETR Transactions/Sec
- ITR Transactions/CPU Sec
- ITR CPU / Transaction

Note the stability of the ITR values over the week relative to the regular increases and decreases in the ETR.

Expressing Relative Capacity with ITRs



- ITRs can be used to characterize processor capacity
 - Since the primary factor is CPU consumed
- Thus, ITRs are useful for determining capacity relationships between processor
 - ITRs used must be for identical systems and workloads
 - Done by calculating the ITR Ratio (a.k.a. ITRR)

$$\text{ITRR} = \frac{\text{ITR for CPU-B (After)}}{\text{ITR for CPU-A (Before)}}$$

- This has been the traditional basis for IBM's LSPR methodology...
 1. Calculate the ITRs for all processors for all workloads
 2. Select a base processor to relate all other processors to
 - *IBM uses zSeries z9 2094-701 as its current base processor*
 3. Calculate ITRRs for all processors for all workloads relative to the base processor
 4. Publish results

Why ETRs, ITRs and ITRRS matter



- IBM evaluates each new processor by measuring their achieved ITRs
 - Different workloads are used to gain insights to how well a process may perform for a given typical workload
- LSPR - Large System Performance Reference
 - Represent IBM's assessment of relative processor capacity in an unconstrained environment for the specific benchmark workloads and system control programs
 - For LSPR to be reliable it must be sensitive to various workload environments
 - Thus the reason the SMF 113 records are so important
 - IBM's processor evaluation may differ from yours
 - Differences between the specified workload characteristics and your operating environment
 - Differences between the specified system control program and your actual system control program
 - I/O constraints in your environment

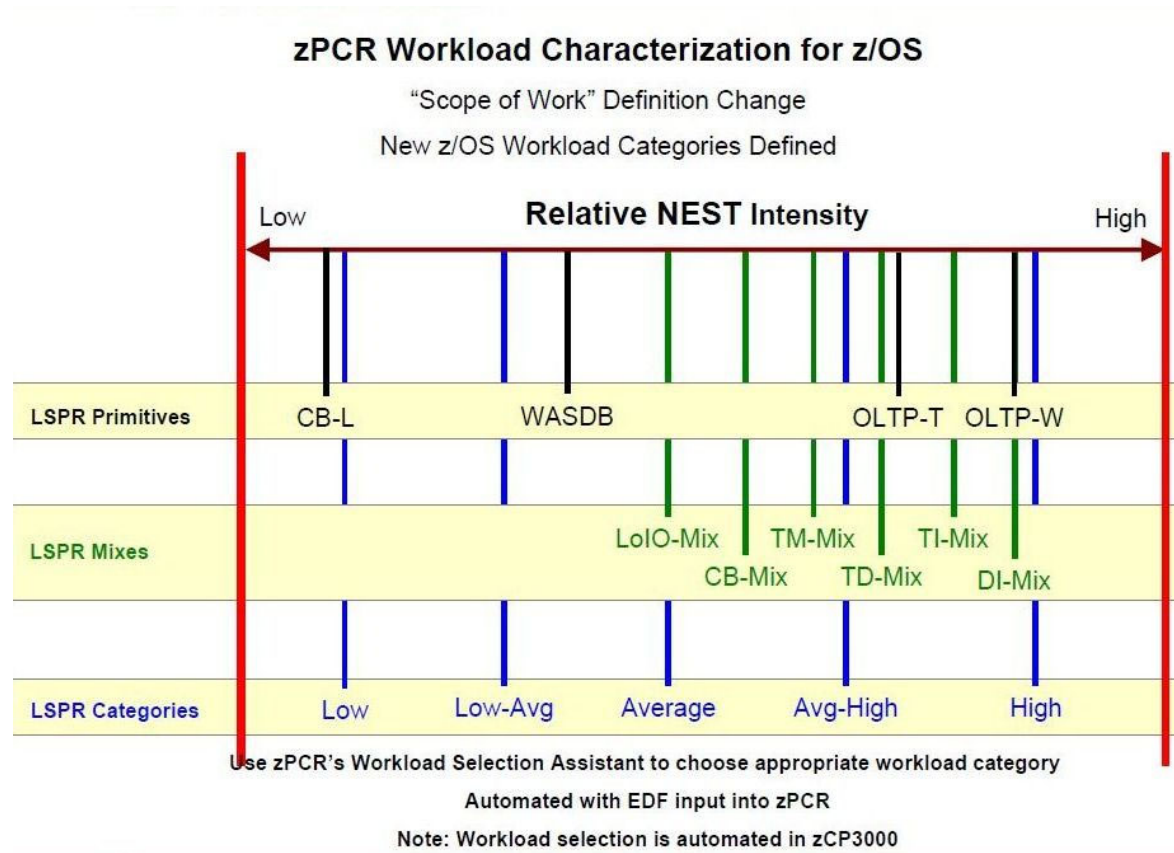
IBM LSPR Environment




- IBM runs a series of 'LSPR Workloads' on new processor models
 - Uses measured ITR values to evaluate processor performance
 - Note: Not every workload is run on every machine. Instead, many ITR values are projected
- Primary workloads used to evaluate the zArchitecture running z/OS
 - CB-L – Commercial Batch Long Job Steps
 - CPU intensive commercial batch (Long Job Steps)
 - Considered a Low I/O Workload (explained later)
 - CB-J – Java Batch
 - The JavaBatch workload reflects the batch production environment of a clearing bank that uses a collection of java classes working on a DB2 database and a set of flat files in z/OS
 - WAS-DB – WebSphere Application Server Database
 - WebSphere Application Server access to
 - ODE-B – One Demand Environment - Batch
 - The ODE-B workload reflects the billing process used in the telecommunications industry
 - OLTP-W
 - Web-enabled On-line Workload (CICS/DB2 with web front end)
 - OLTP-T
 - Traditional On-line Workload (IMS)
 - Formally known was the IMS workload

Relative Nest Intensity and IBM LSPRs

(chart source: John Burg SMF 113 sessions)



LSPR Table Example - (z14 LSPRs)



All values relative to (System z9 2094-701 = 1.00)

Processor	#CP	PCI**	MSU***	MSUps****	Low*	Average*	High*
3906-701	1	1832	227	182	3.37	3.27	3.04
3906-702	2	3464	427	342	6.55	6.19	5.61
3906-703	3	5050	620	496	9.66	9.02	8.08
3906-704	4	6590	808	646	12.68	11.77	10.45
3906-705	5	8082	990	792	15.64	14.44	12.74
3906-706	6	9528	1162	930	18.55	17.02	14.96
3906-707	7	10927	1326	1061	21.41	19.52	17.11
3906-708	8	12283	1487	1190	24.21	21.94	19.18
3906-709	9	13597	1642	1314	26.96	24.29	21.2
3906-710	10	14869	1793	1434	29.66	26.56	23.14
3906-711	11	16101	1939	1551	32.31	28.76	25.03
3906-712	12	17294	2077	1662	34.91	30.89	26.85
3906-713	13	18450	2213	1770	37.47	32.96	28.62
3906-714	14	19570	2346	1877	39.97	34.96	30.33
3906-715	15	20655	2476	1981	42.43	36.9	31.99

Presentation Summary



- ETR – External Throughput Rate
 - Focuses on the environment capacity
- ITR – Internal throughput Rate
 - Focuses on processor capacity
- Two great formulas to use to evaluate changes to your environment



Questions?