ETR vs ITR and the Basics of IBM's LSPR

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:

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

- **ETR vs ITR and the Basics of IBM's LSPR**

  - The concepts and formulas for External Throughput Rates (ETRs) and Internal Throughput Rates (ITRs) are the foundation 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. This session will also provide insights into IBM’s LSPRs which are used to represent IBM's assessment of relative processor

EPS: We do z/OS performance…

- **Pivotor** – z/OS performance reporting and analysis software and services
  - Not just SMF reporting, but analysis-based reporting based on expertise
  - www.pivotor.com
- **Education and instruction**
  - We teach our z/OS performance workshops all over the world
  - Want a workshop in your area? Just contact me.
- **z/OS Performance War Rooms**
  - Intense, concentrated, and highly productive on-site performance group discussions, analysis and education
  - Amazing feedback from dozens of past clients
- **MSU Reduction Exercises**
  - The goal is to reduce the MSU consumption of your applications and environment
- **Information**
  - We present around the world and participate in online forums
  - https://www.pivotor.com/content.html
    https://www.pivotor.com/webinar.html

z/OS Performance workshops available

During these workshops you will be analyzing your own data!

- WLM Performance and Re-evaluating Goals
  - May 12 - 16, 2025 (4 days)

- Parallel Sysplex and z/OS Performance Tuning
  - October 21 - 22, 2025 (2 days)

- Essential z/OS Performance Tuning
  - September 22 - 26, 2025 (4 days)

- Also… please make sure you are signed up for our free monthly z/OS educational webinars! (email contact@epstrategies.com)

Pivotor Software as a Solution (SaaS)

We do everything for you
so, you can focus on Analysis!

Because we also have access to your data, SaaS Includes:

- Formal yearly cursory review / discussion
- Ability to ask us performance questions, or for us to look at a particular problem or concern. (support@epstrategies.com)
- We can occasionally look in on your data and performance
- We can participate in performance debug with IBM, or other vendors

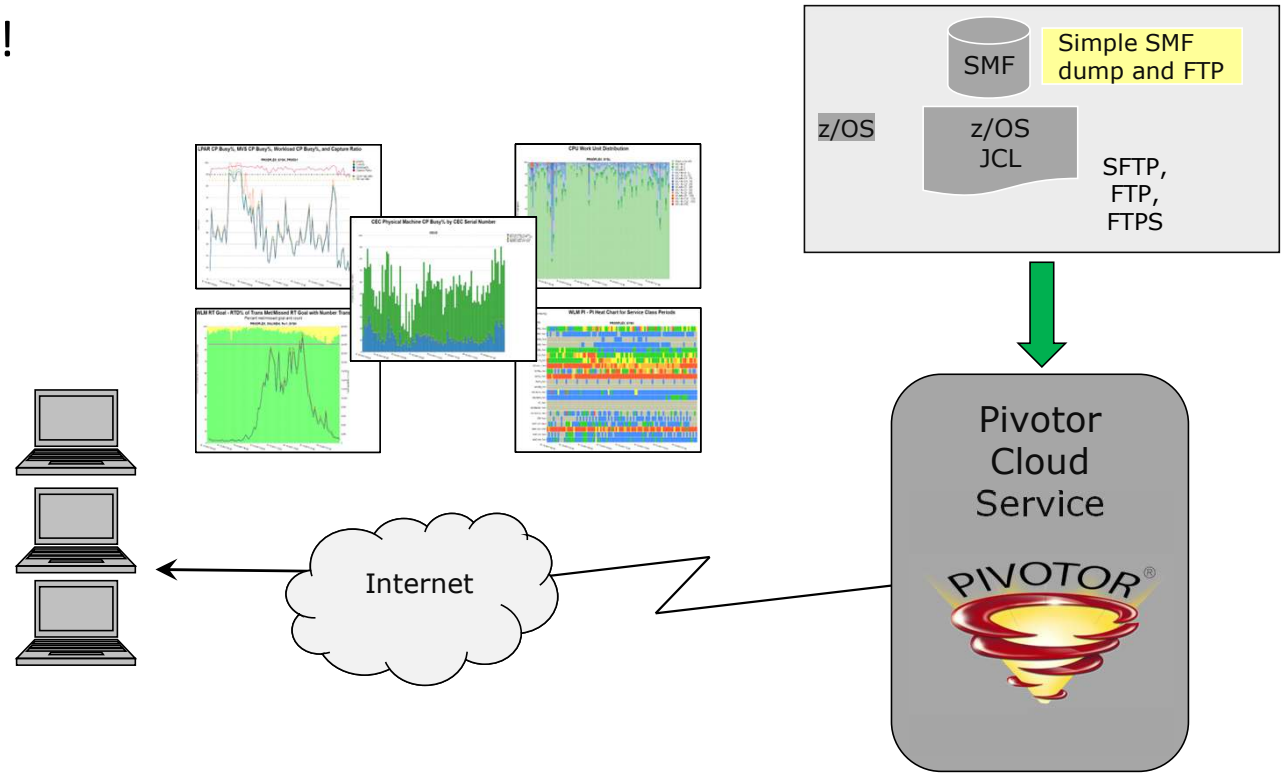

Like what you see?
Spring / Summer 2025 Webinars

- Free z/OS Performance Educational webinars!
  - The titles for our Spring / Summer 2025 webinars are as follows:
    - ✓ *Overseeing z/OS Performance Management With Your Outsourcer*
    - ✓ *Back to basics - Processor Consumption Analysis*
    - ✓ *Pivotor Pointers*
    - ✓ *WLM and CPU Critical Control*
    - ✓ *Back to Basics - Evaluating Latent Demand*
    - ✓ *Understanding SMF 98 Locking Measurements*
    - ✓ *Standard Measurements when Monitoring Transactions*
    - ✓ *Overseeing z/OS Performance Management with Your Outsourcer*
    - *ETR vs ITR and the Basics of IBM's LSPR*
    - *Processor Comparison Discussion*
    - *SMF 99 WLM Decision Making Traces*
    - *Optimization in the Application Trenches*
    - *Understanding SMF 98 Address Space Consumption Measurements*
    - *z/OS Performance Management in an AI World*
    - *Understanding z/Architecture Processor Topologies*

- If you want a free cursory review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html

Like what you see?

- If you want a free cursory review of your environment, let us know!
  - We're always happy to process a day's worth of data and show you the results
  - See also: http://pivotor.com/cursoryReview.html
- We also have a ***free*** Pivotor offering available as well
  - 1 System, SMF 70-72 only, 7 Day rolling retention
  - That still encompasses over 100 reports!
- We offer free bi-monthly z/OS performance educational Webinars
  - Sign our guestbook to get invited
  - Or send an email to
    - Peter.Enrico@EPStrategies.com

zPCR

- The purpose of this presentation is to provide an understanding of the concepts of LSPRs and how IBM determines relative processor capacities

- However, when doing a processor sizing exercise:
  - Use IBM's zPCR tool!
  - Uses IBM LSPRs internally

- zPCR is a free tool provided by IBM to assist in processor sizing
  - Provide capacity relationships for zSeries processors
  - Takes into consideration
    - Workload environment (as discussed in this presentation)
    - Particular LPAR configurations
  - IBM says the expected accuracy is ±5%
  - Inputs include
    - CPU activity measurements (SMF 70)
    - Workload activity measurements (SMF 72.3)
    - Processor counters (SMF 113)
    - LPAR configuration

This is an IBM LSPR Table excerpt

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
9175-701	1	2477	306	4.37	4.42	4.29
9175-702	2	4731	582	8.58	8.45	8.07
9175-703	3	6944	852	12.68	12.4	11.73
9175-704	4	9116	1112	16.67	16.29	15.25
9175-705	5	11213	1360	20.59	20.03	18.64
9175-706	6	13250	1600	24.46	23.67	21.94
9175-707	7	15229	1836	28.26	27.2	25.14
9175-708	8	17151	2060	32.02	30.64	28.25
9175-709	9	19020	2281	35.73	33.98	31.28
9175-710	10	20837	2497	39.38	37.22	34.22
9175-711	11	22602	2702	42.98	40.38	37.08
9175-712	12	24317	2894	46.54	43.44	39.86
9175-713	13	25984	3079	50.04	46.42	42.56
9175-714	14	27604	3256	53.49	49.31	45.18
9175-715	15	29178	3429	56.9	52.12	47.74

System z9 EC

(System z9 2094-701 = 1.00)

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
2094-701	1	560	81	1	1	1

The following is an excerpt of the IBM LSPR tables.

- LSPR = Large system Performance Reference
- These numbers represent the IBM assessment of the relative processor capacity of each processor model and type.
- These numbers will be explained more during this presentation
- Google “IBM LSPR” to find the relative processor ratings for all machines

Note the PCI relationship between the 9175-701 and PCI of 2094-701

If the Average rating for the 9175-701 is 4.42, then

(Average 9175-701 of 4.42) * (2094-701 PCI of 560)

= (9175-701 PCI of 2477)

Throughput Measures: ETRs, ITRs, ITRRs

Did a change make a positive improvement or a degradation?

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

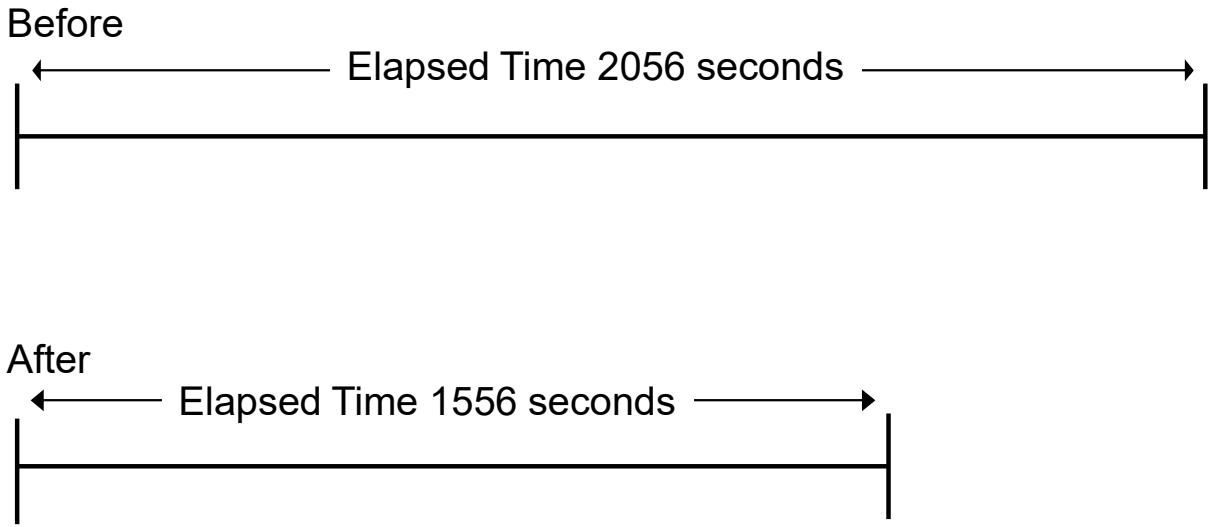

Understanding Workload Throughput

- 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

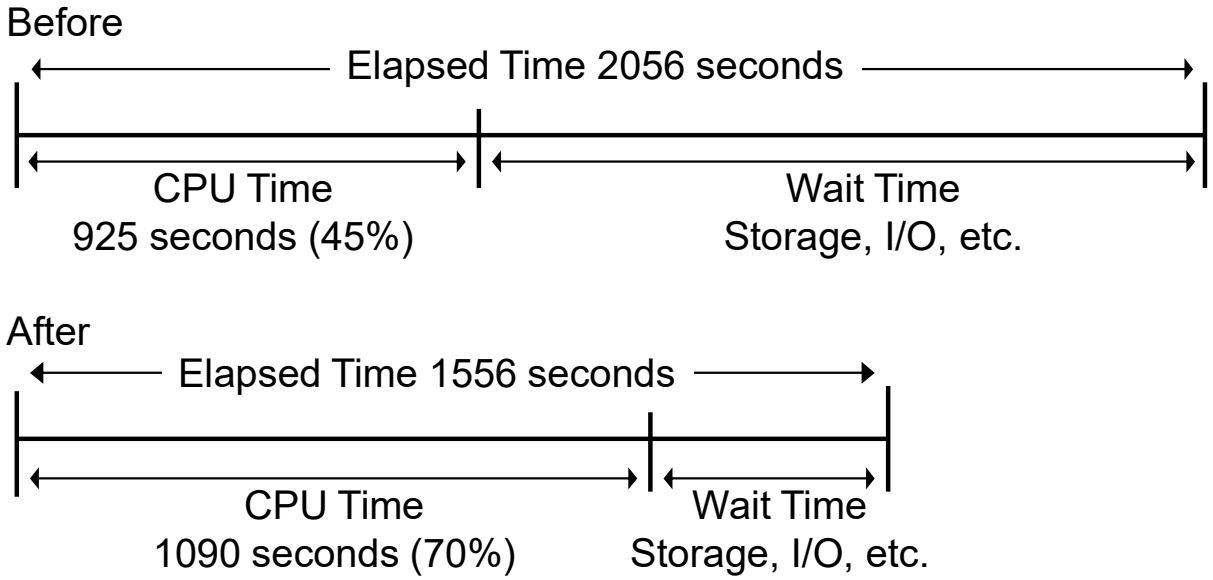

External Throughput Rate - ETR

- 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

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

- 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

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

  - 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:

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

  - Then

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

  - Thus

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

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

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

Example of Using ITR/ETR Relationship

- Since $$ITR = \frac{ETR}{CPUBusy\%}$$
  - 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

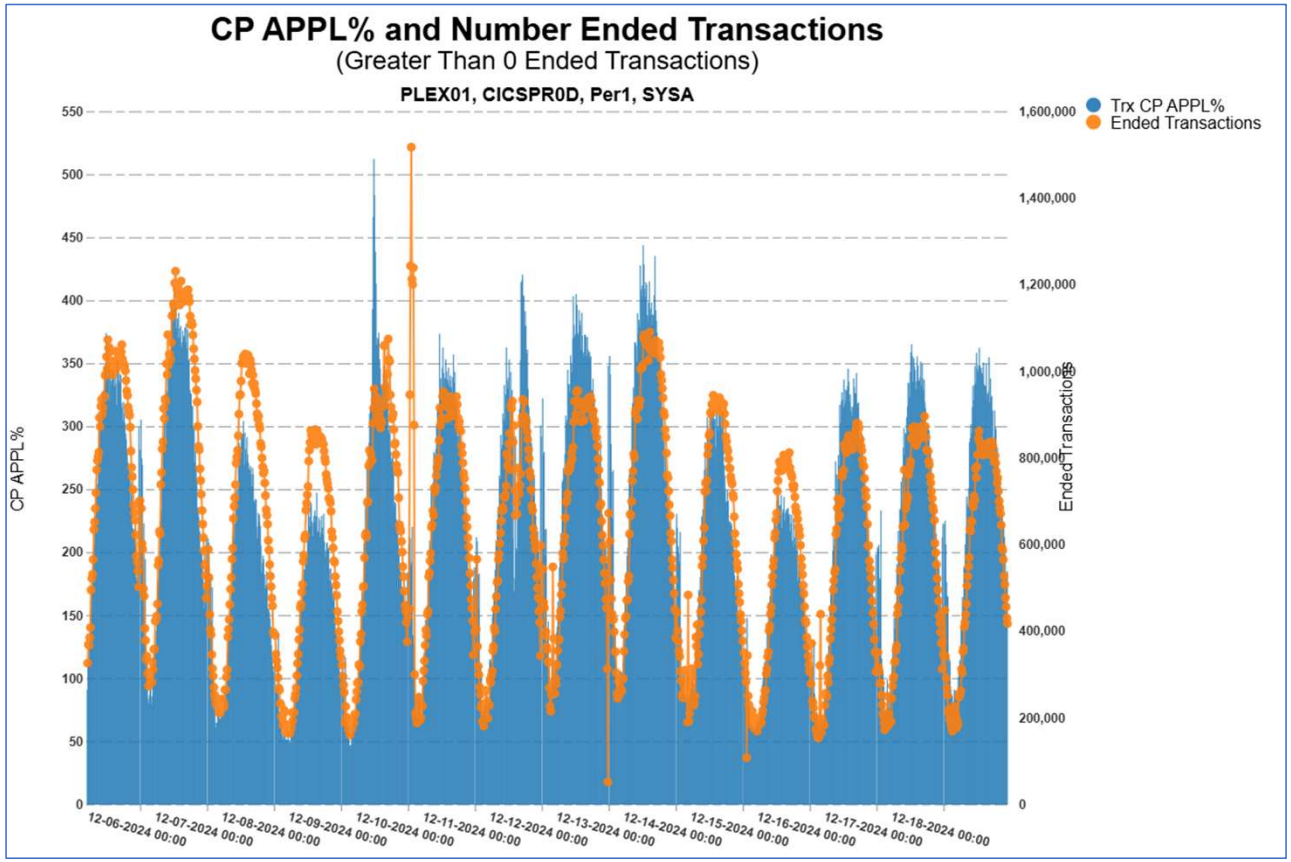

This double y-axis chart shows both the number of ended transactions for CICSPROD transactions, as well as the CP APPL% of these transactions.

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

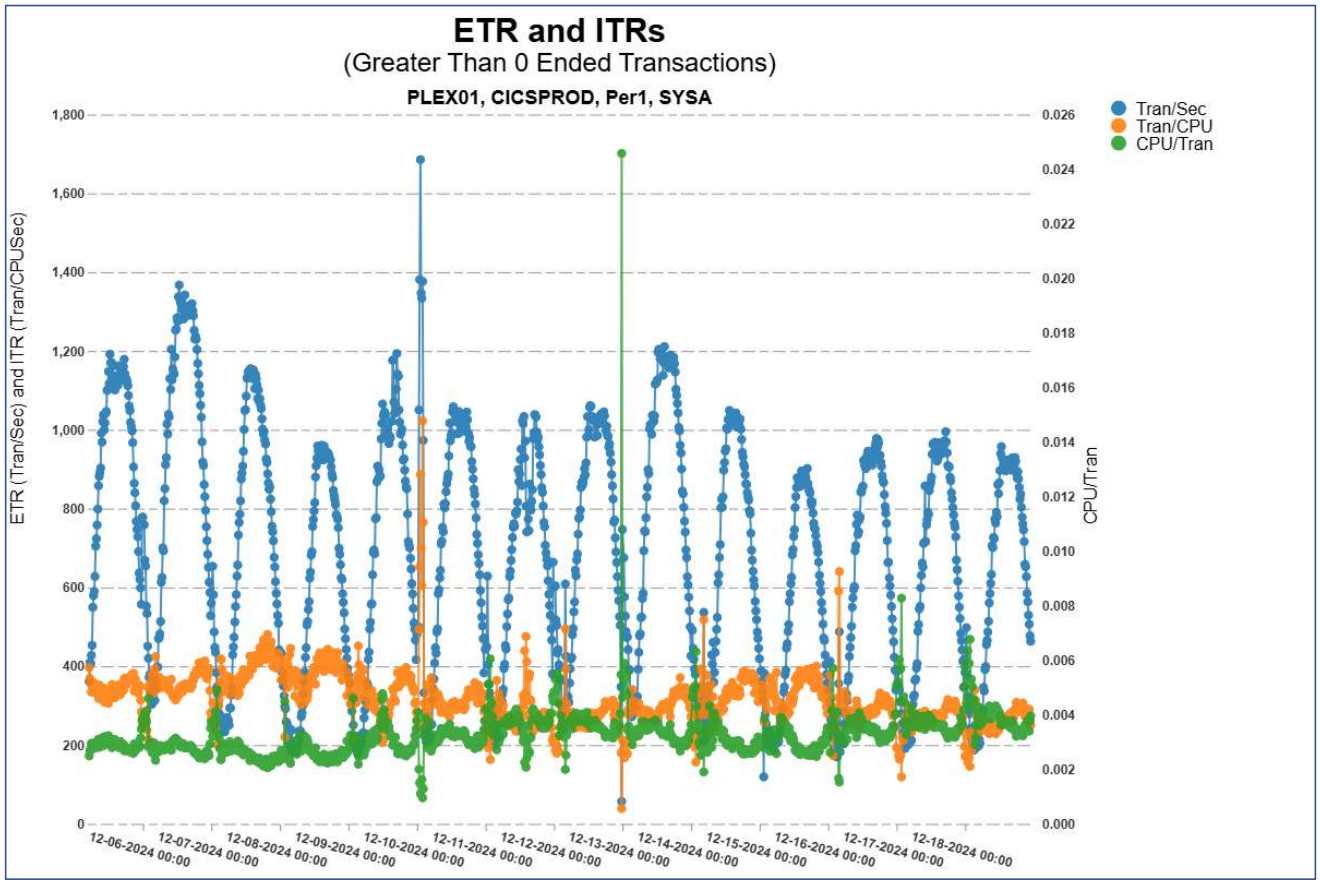

This chart shows the three ETR and ITR values:

- 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: CICS Transactions

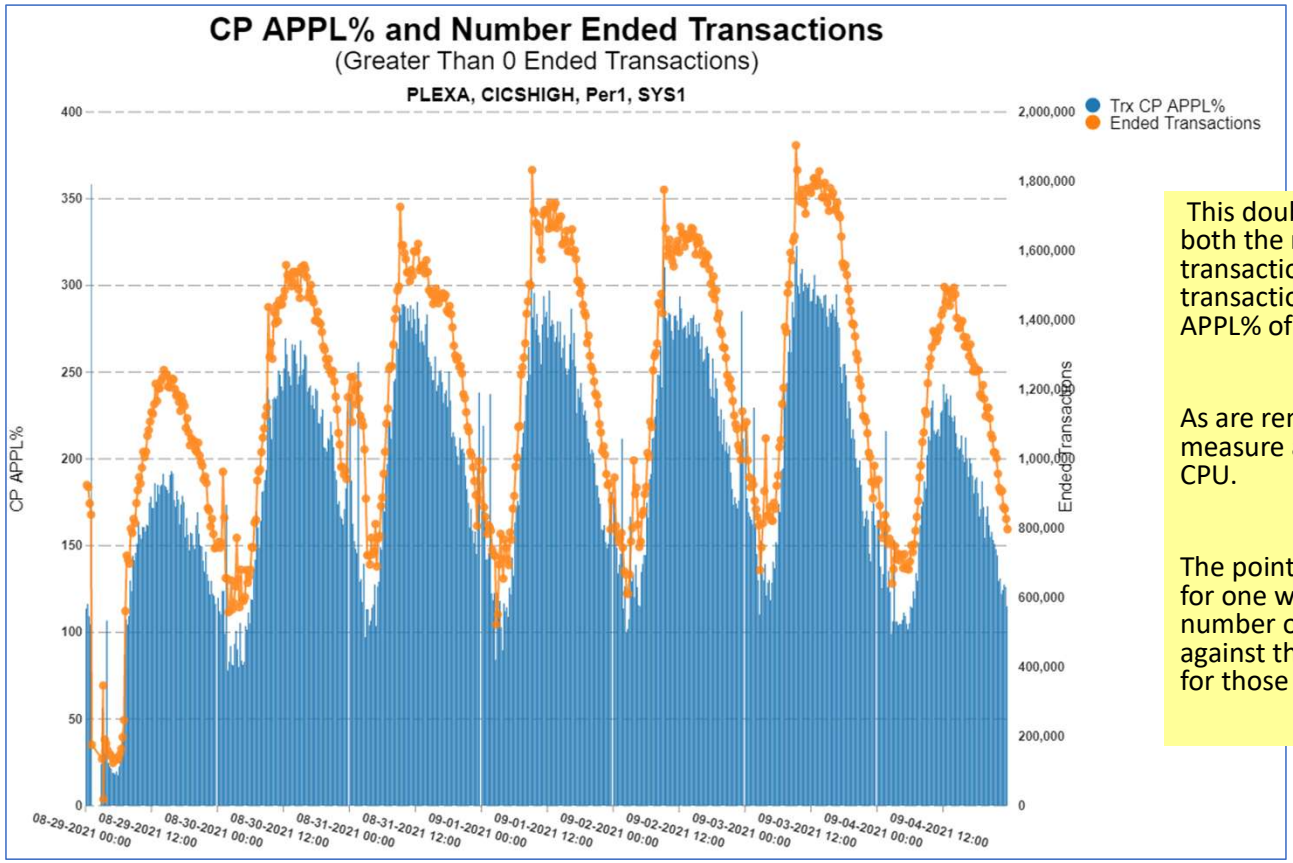

This double y-axis chart shows both the number of ended transactions for CICSHIGH transactions, as well as the CP APPL% of these transactions.

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

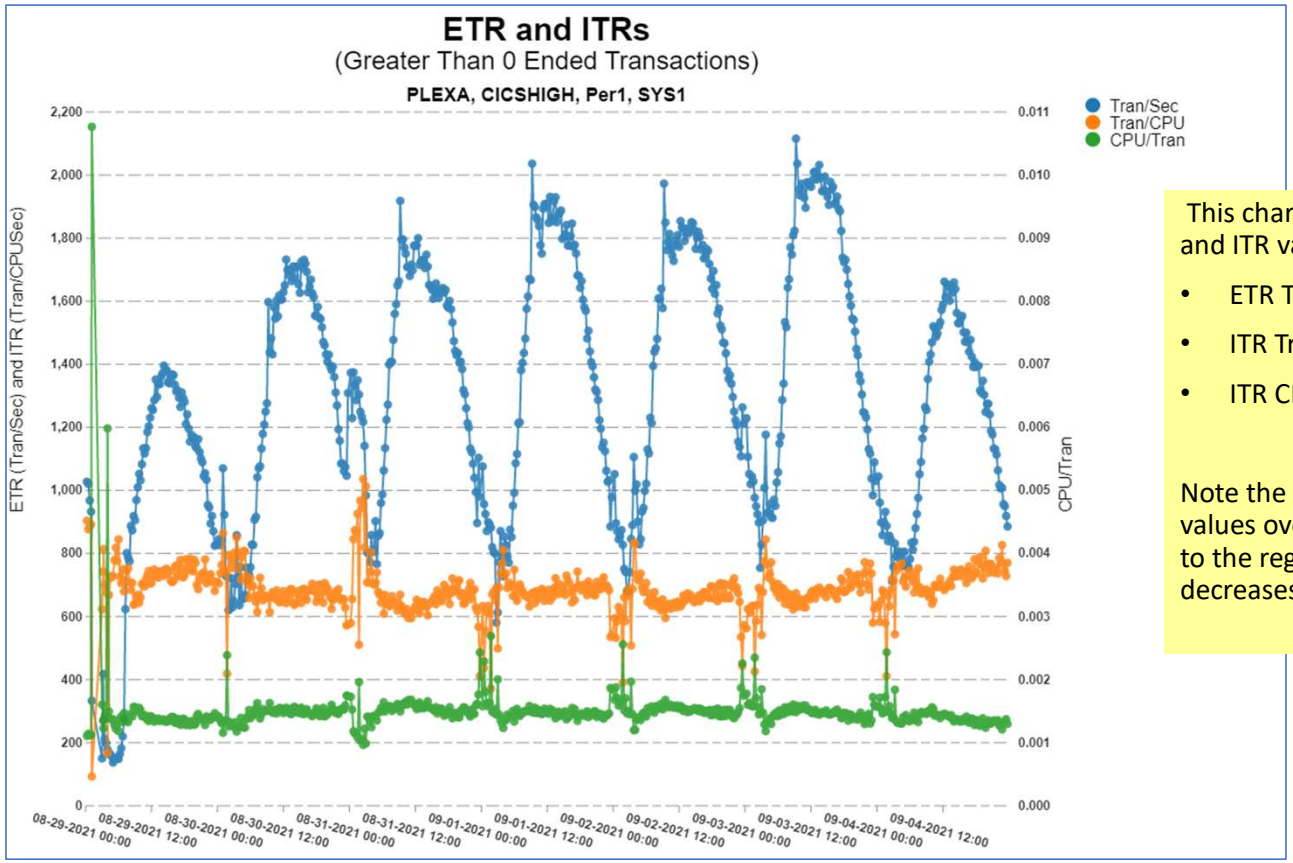

This chart shows the three ETR and ITR values:

- 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

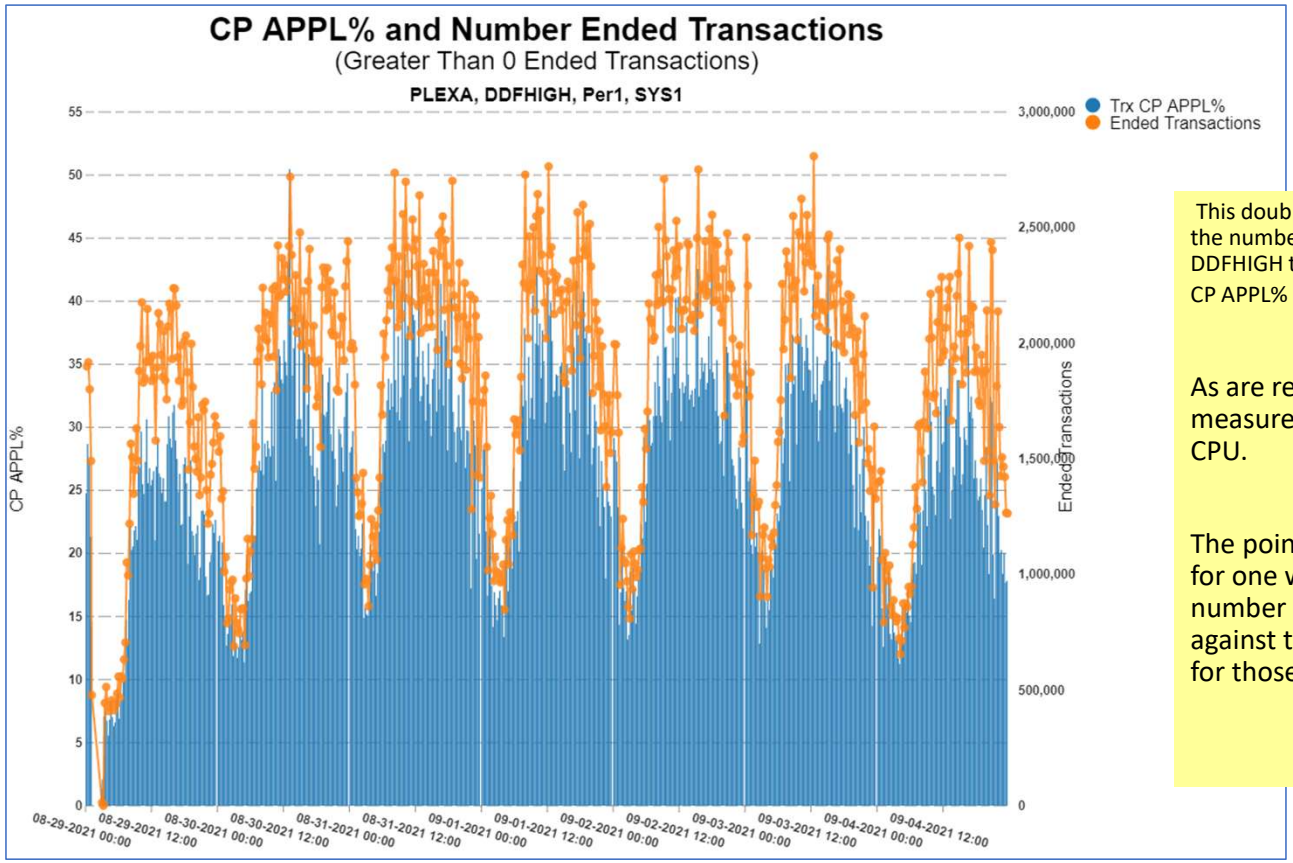

This double y-axis chart shows both the number of ended transactions for DDFHIGH transactions, as well as the CP APPL% of these transactions.

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: DDF Transactions

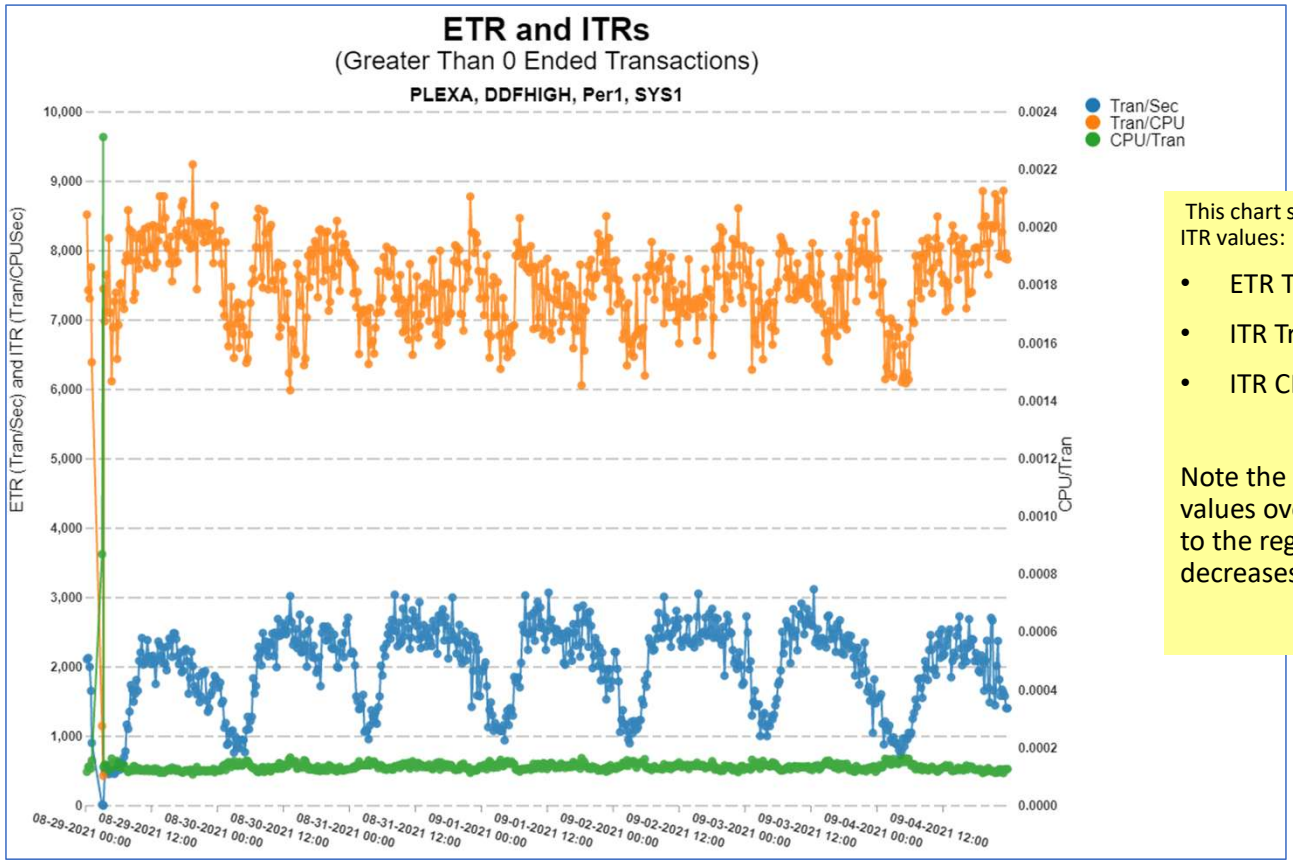

This chart shows the three ETR and ITR values:

- 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.

ITRs are the foundation of IBM's LSPRs for processor sizing

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
9175-401	1	296	37	0.52	0.53	0.51
9175-402	2	567	72	1.03	1.01	0.97
9175-403	3	835	106	1.52	1.49	1.41
9175-404	4	1100	138	2	1.96	1.85
9175-405	5	1355	168	2.48	2.42	2.26
9175-406	6	1603	199	2.94	2.86	2.67
9175-407	7	1845	228	3.4	3.3	3.07
9175-408	8	2082	257	3.86	3.72	3.45
9175-409	9	2316	286	4.31	4.14	3.84
9175-410	10	2545	314	4.76	4.55	4.22
9175-411	11	2769	342	5.21	4.95	4.59
9175-412	12	2988	369	5.65	5.34	4.96
9175-413	13	3201	395	6.08	5.72	5.31
9175-414	14	3408	420	6.51	6.09	5.66
9175-415	15	3611	445	6.93	6.45	6
9175-416	16	3809	470	7.35	6.8	6.33
9175-417	17	4004	493	7.76	7.15	6.66
9175-418	18	4199	517	8.16	7.5	6.98

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

Key Influences of Processor Performance and Capacity

- To understand LSPRs we need to ask this question:

  What are the key influences that result in variations of a particular processor's delivered capacity relative to a customer's environment and workload?

- Answer: As Gary King of IBM would say… there are three key influences:
  - Instruction complexity of one processor family to another
  - Path length of the code executed by customer applications and transactions
  - Usage of the Memory Hierarchy
- A machine's capacity will vary based on each of these three factors

Key Influence - Instruction Complexity

- Instruction complexity of one processor family to another
  - Types of instructions, the sequence in which they are executed, and the way they interact with the processor design
  - Many machine design alternatives affect instruction complexity
  - Each processor family has variations in the chip design
  - As processor evolve the way the chip executes instructions is enhance and geared towards the technology
- Examples include cycle time of CPU, how the instructions are wired to execute (using pipelining, branch prediction, out of order execution, etc.)
- Influences the workload
  - Once a customer workloads are on a processor, instruction complexity is relatively constant between customers and workloads
  - In other words, relative to the LSPRs and sizing, once a move is made to the new processor family instruction complexity does not vary much from one customer to the next.

Key Influence - Instruction Complexity

- Notice the PCI column showing the PCIs of the 701 series of each processor
  - PCI : Processor Capacity Index (term used by IBM instead of MIPS)

Processor	#CP	PCI**	MSU***	Low*	Average*	High*	
2094-701	1	560	81	1	1	1	z9 EC
2827-701	1	1,514	188	2.75	2.7	2.55	zEC12
2964-701	1	1695	210	3.18	3.03	2.75	z13
3906-701	1	1832	227	3.37	3.27	3.04	z14
8561-701	1	2055	253	3.72	3.67	3.46	z15 T01
3931-701	1	2253	278	4.03	4.02	3.85	z16 A01
9175-701	1	2477	306	4.37	4.42	4.29	z17

Key Influence – Path Length

- Path length of the code executed by customer applications and transactions
  - This relates to code executed by applications / jobs / transactions / etc.
  - Instruction count
- The actual path lengths executed by a workload will vary
  - From customer to customer, and from IBM synthetic workloads versus customer
  - From one customer's application environment versus another application environment of that same customer
    - Example: CICS / DB2 application versus a WAS / DB2 application
- Is sensitive to the configuration due to MP effects
  - Higher n-ways or difference in configuration may increase path lengths execute (which in turn influences the processor capacity relative to LSPRs)
    - Example: May have more locking in a higher MP environment, or queues may be longer, etc.
- But when move from one processor to another this generally does not change much for a specific customer
  - Whether the move is from one processor family to another
  - Or from one process in the same family to another

Key Influence – Memory Hierarchy

- Usage of the Memory Hierarchy
  - Heavily influenced by key factors result potentially wide variations in realized capacity
  - From one processor family to another there are many design alternatives
    - Levels of cache, scope of cache, latency, etc.
  - Configuration will influence usage of the memory hierarchy
    - LPAR configuration, competition between LPARs, options such as HiperDispatch, etc.
  - Exploitation by workloads will influence usage of the memory hierarchy
    - Transaction intensity, memory intensity, I/O intensity, application mixtures, competition of resource by applications, etc.
  - z/OS performance management and options
    - WLM management of resources, affinity nodes, IEAOPTxx opts, heap sizes, initiators, etc.
- Final result is that usage of memory hierarchy heavily influences a processor's delivered capacity and performance.
  - Workload performance sensitive to how deep into the memory hierarchy the processor must go to retrieve instructions and data
- So, for processor sizing, LSPRs have started focusing on this

LSPR Workloads and Usage of the Memory Hierarchy

- Usage of the measurement hierarchy influences processor capacity
  - RNI – A measure of the distribution of the activity to the shared processor caches
  - L1MP – Level 1 Misses per 100 instructions (basically an L1 cache miss %)

Low	**Relative Nest Intensity**	**High**
Batch	Application Type	Transactional
Low	I/O Rate	High
Single	Application Mix	Many
Intensive	CPU Usage	Light
High Locality	Data Reference Pattern	Diverse
Simple	LPAR Configuration	Complex
Extensive	Software Configuration Tuning	Limited

IBM LSPR Environment

- Primary workloads used to evaluate the zArchitecture running z/OS to exercise processor caches From these workloads ITRs are calculated, and ITRR values are use to derive LSPR values
  - **CB-L**
    - Represents a commercial batch job stream of large batch jobs with heavy CPU processing
    - CPU intensive commercial batch (Long Job Steps) and considered a Low I/O Workload (explained later)
  - **CB-J**
    - Represents a JavaBatch workload accessing a DB2 database and flat files
    - Simulates a production environment of a clearing bank that uses a collection of java classes
  - **WAS-DB**
    - Represents a WebSphere (WAS) application to a DB2 database with everything running on z
  - **ODE-B**
    - Represents an on-demand batch environment (which reflect the billing process used in the telecommunications industry
  - **OLTP-W**
    - Represents a web enabled production environment with a web-enabled front end to a traditional database
    - Web-enabled On-line Workload (WAS front end to a CICS/DB2 workload)
  - **OLTP-T**
    - Represents a traditional On-line Workload (IMS)
    - Moderate to heavy IMS transactions from DLI applications
  - **CICS/DB2**
    - Represents a traditional CICS/DB2 workload via an MRO model managed by CP/SM
  - **Note:** Other workloads for Linux and z/VM environments

There are 4 key values (Pivotor Example)

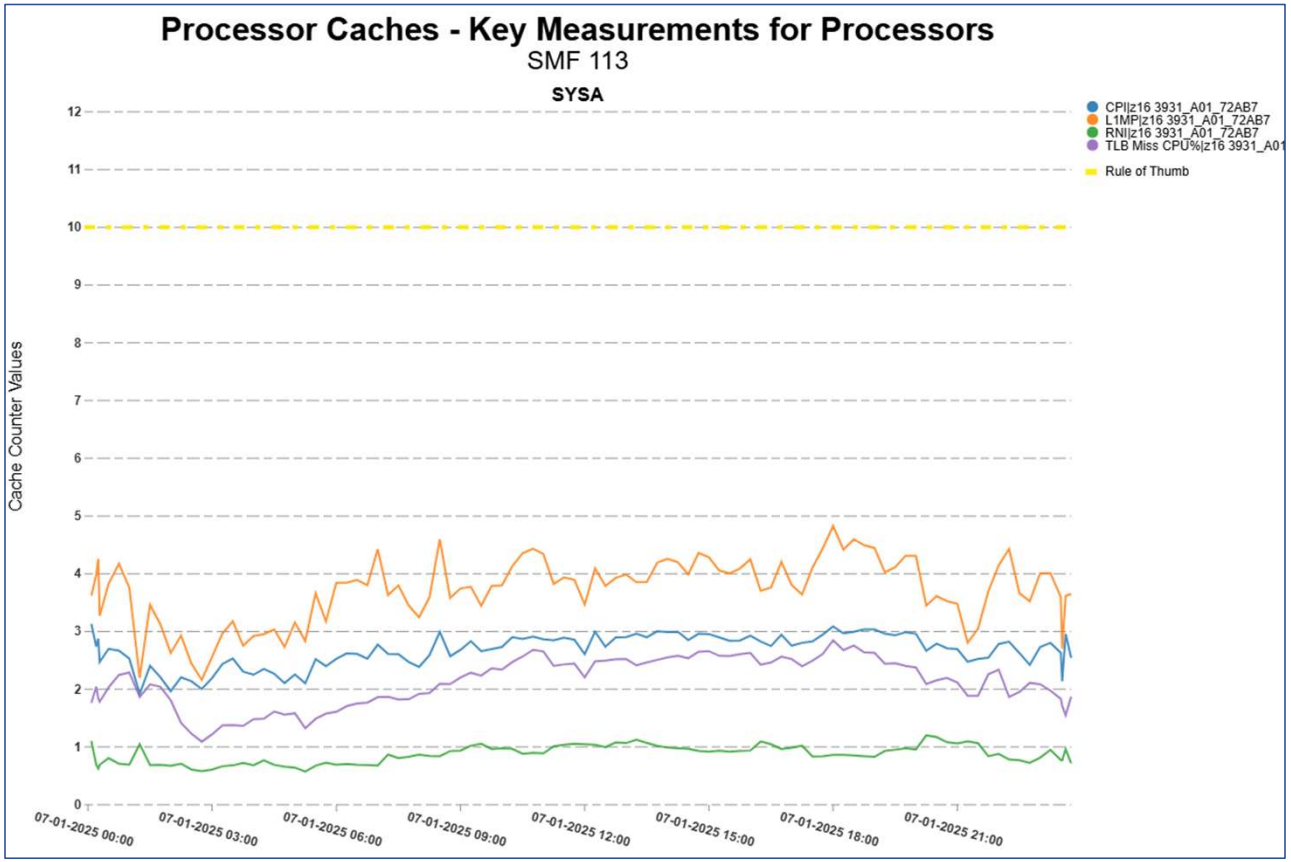

4 key calculated (CP CPU example):

- CPI: Cycles per Instruction
  - Used to gauge processor contention, as well as instruction mix consistency
- L1MP: L1 misses per 100 instructions
  - Think of this as a 'miss percentage'
- RNI: Relative Next Intensity
  - Workload 'signature' that gauges the pressures being put on the upper caches and memory
- TLB Miss CPU Percentage
  - Total percent of the CPU consumed by the LPAR that goes to dynamic address translation (DAT) due to a translation look-aside buffer miss

LSPR SMF 113 Based Workload Categories

- Traditional IBM workload categories are still being used to determine some base capacity values
  - (see previous slide)

- LSPRs use 3 LSPR categories based on the SMF 113 processor cache counter measurements
  IBM still uses the traditional workloads to evaluate LSPRs, but then puts them into context of SMF 113s
  - **LOW (RNI)**
    - Light usage of the memory hierarchy
  - **AVERAGE (RNI)**
    - Average usage by most customers of the memory hierarchy
    - Similar to the old LoIO mixed workload curve
  - **HIGH (RNI)**
    - Heavy usage of the memory hierarchy
    - Similar to old DI-mix workload curve
    - These replace the previous workload mixtures
    - Allows the LSPR capacity curves to be based on the underlying hardware sensitivities

LSPRs and SMF 113s and RNI Hint

- SMF 113 measurements are now used to provide guidelines / hints for LSPR and zPCR processor sizing
- This RNI Hint table was documented in the Large System Performance Reference (LSPR)
  - Document Number SC28-1187
- The next slide shows an example of an LSPR chart used for processor sizing
- Using the SMF 113 records you now need to calculate
  - L1MP - L1 Miss Per 100 Instructions
  - RNI – Relative Nest Intensity
- Note: This table and these guidelines are expected to change as more is learned from the SMF 113 records

L1MP	RNI	Workload Hint
<3%	>= 0.75	AVERAGE
	< 0.75	LOW
3% to 6%	>1.0	HIGH
	0.6 to 1.0	AVERAGE
	< 0.6	LOW
>6%	>=0.75	HIGH
	< 0.75	AVERAGE

LSPR Table Example - (z17 LSPRs)

Relating z15 LSPRs to z17 LSPRs

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

L1MP	RNI	Workload Hint
<3%	>= 0.75	AVERAGE
	< 0.75	LOW
3% to 6%	>1.0	HIGH
	0.6 to 1.0	AVERAGE
	< 0.6	LOW
>6%	>=0.75	HIGH
	< 0.75	AVERAGE

z15 LSPRs

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
8561-409	9	2067	255	3.98	3.69	3.32
8561-410	10	2270	280	4.39	4.06	3.64
8561-411	11	2469	305	4.79	4.41	3.96
8561-412	12	2663	329	5.19	4.76	4.28
8561-413	13	2852	352	5.58	5.1	4.58

z17 LSPRs

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
9175-701	1	2477	306	4.37	4.42	4.29
9175-702	2	4731	582	8.58	8.45	8.07
9175-703	3	6944	852	12.68	12.4	11.73
9175-704	4	9116	1112	16.67	16.29	15.25
9175-705	5	11213	1360	20.59	20.03	18.64

z17 Sub LSPRs

Processor	#CP	PCI**	MSU***	Low*	Average*	High*
9175-409	9	2316	286	4.31	4.14	3.84
9175-410	10	2545	314	4.76	4.55	4.22
9175-411	11	2769	342	5.21	4.95	4.59
9175-422	22	4976	612	9.77	8.89	8.27
9175-423	23	5170	635	10.17	9.24	8.6

Accuracy of IBM LSPRs, ITRRs, and zPCR

- LSPRs present themselves as an exact point

Remember that sizing numbers are just estimate of what you will receive

The points in the follow graph show the IBM LSPR values for z17 processors

9175 Processors 401 to 422

LSPR Estimates for:

- Low
- Average
- High

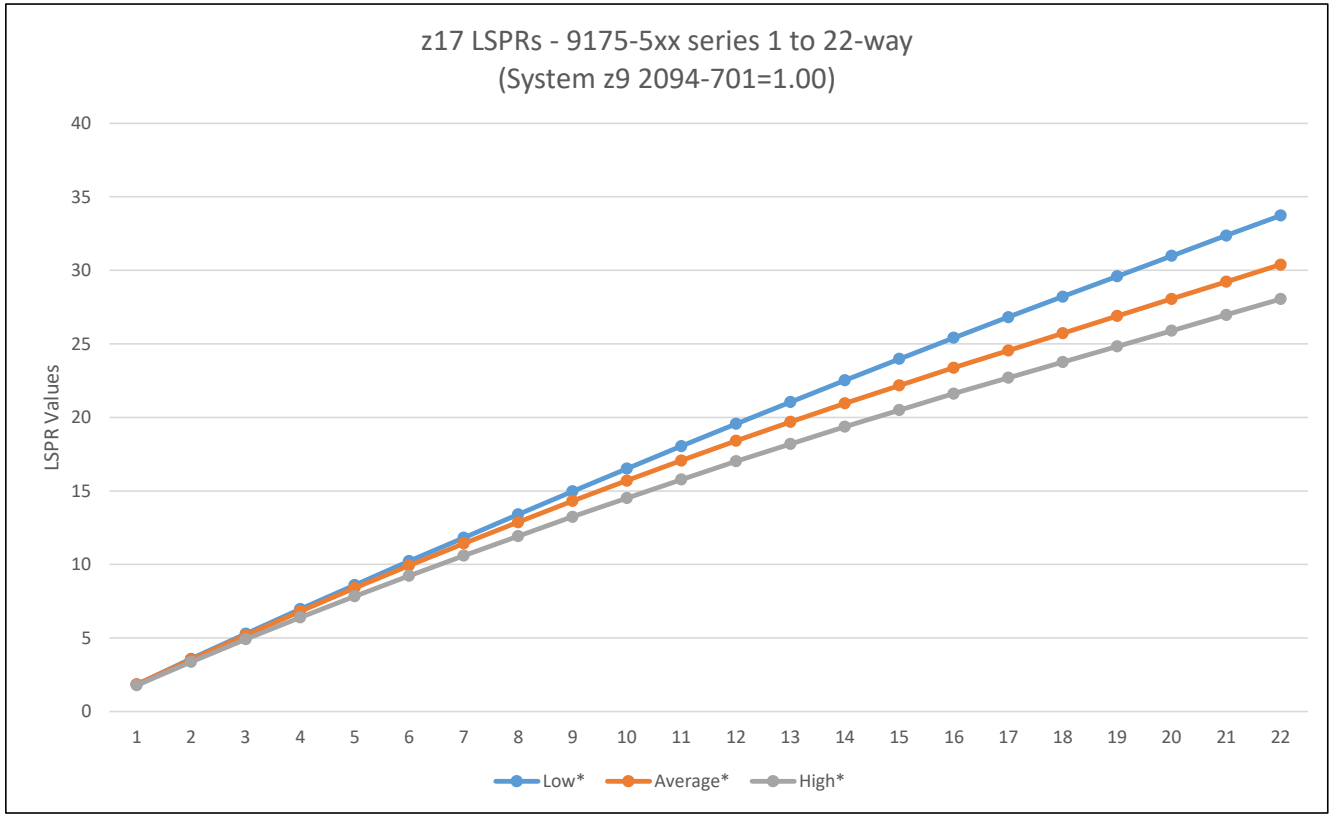

Accuracy of IBM LSPRs, ITRRs, and zPCR

Remember that sizing numbers are just estimate of what you will receive

The points in the follow graph show the IBM LSPR values for z17 processors

9175 Processors 401 to 422

LSPR Estimates for:

- Low
- Average
- High

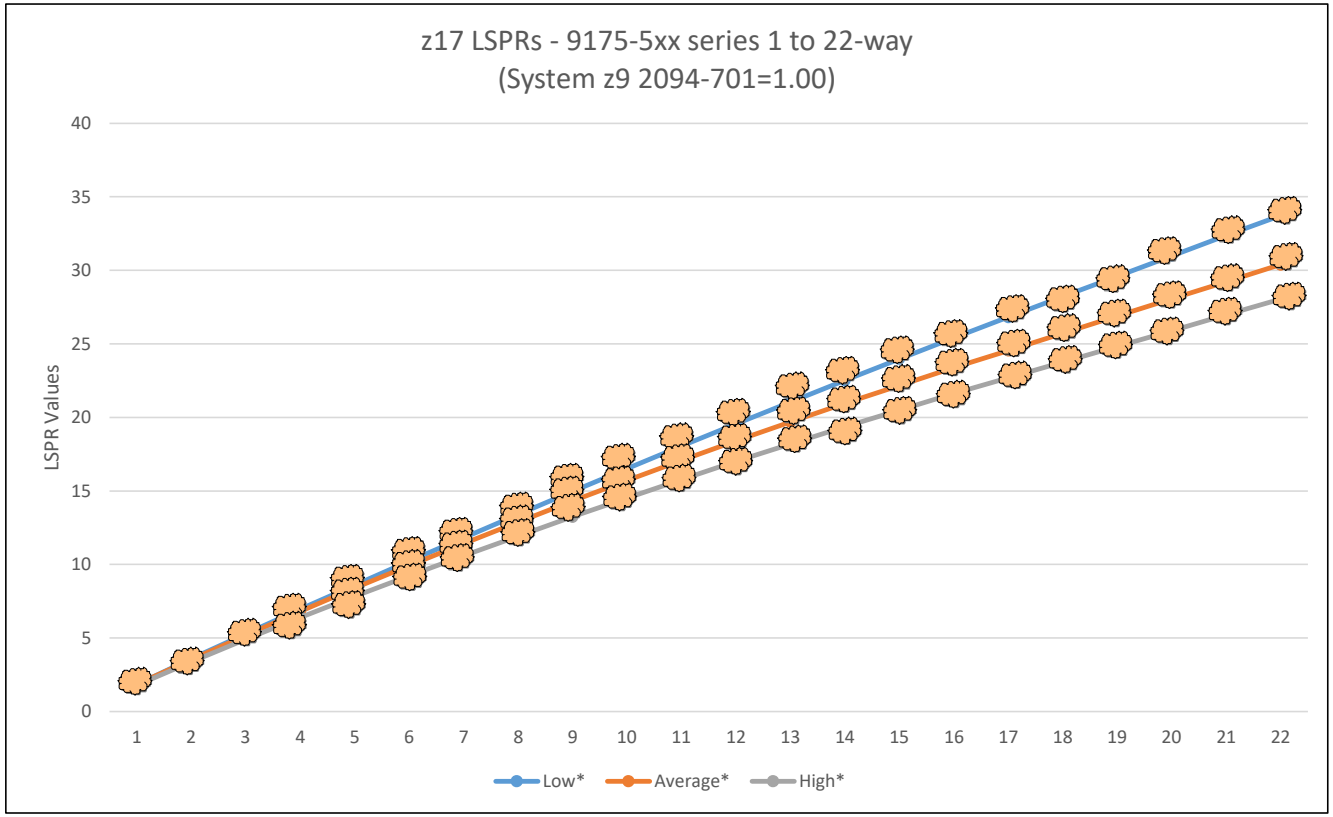

Reality Of An Upgrade When Values Are Estimates

- Since LSPR values, or your own sizing workload mix are only estimates
  - You may not always get what you expected (for many reasons)
  - Jumping from fuzzy to fuzzy

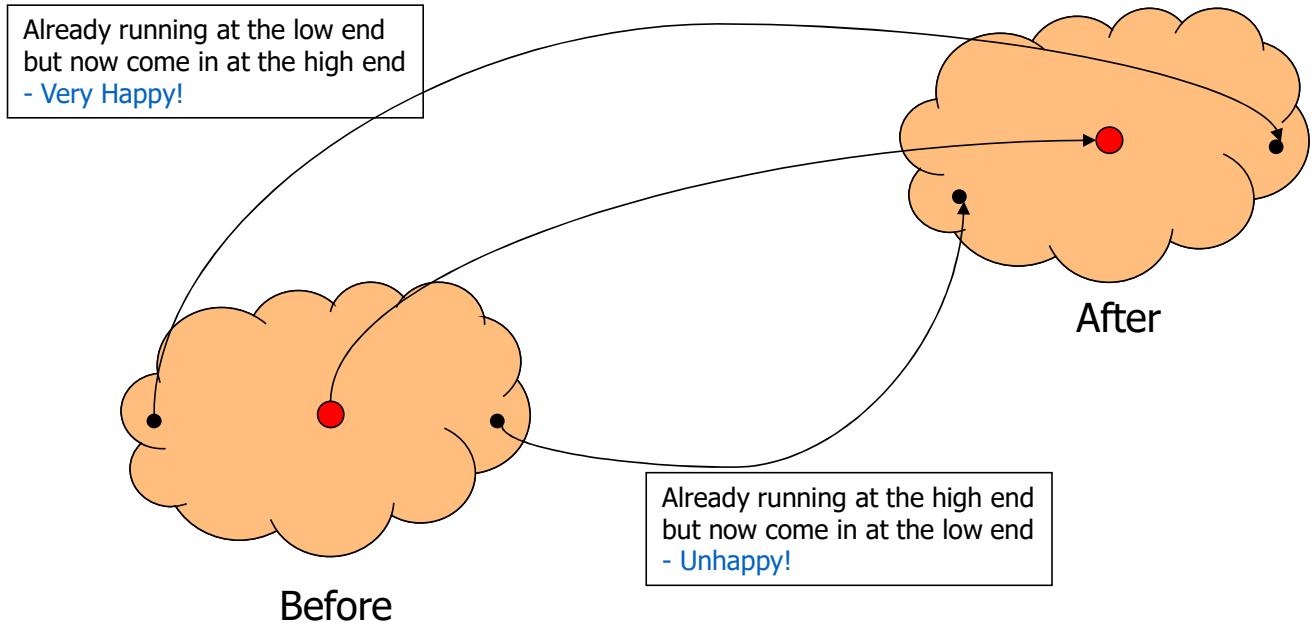

Also, it is Difficult to Anticipate all Changes

- When processor sizing, the LSPRs may help, but it is difficult to anticipate all changes
  - Workload change over time
  - Number of transactions could grow
  - Response time changes will impact end user behavior
  - Latent demand may exist
    - Especially true for single TCB type workloads that are limited to the capacity of a single engine
    - Affected as go to few but fast engines
  - Your workloads are not exactly like IBMs LSPR workloads
  - Each workload is affected differently based on WLM and other factors
    - High priority workload may do well, but low priority workload may suffer

Presentation Summary

- Remember – use the zPCR tool for processor sizing
- It is not necessary to know the details of LSPRs
  - **This presentation was meant to give you some background for understanding:**
- Different workloads perform differently on different configurations and machines
- Great formulas to use to evaluate changes to your environment
  - ETR – External Throughput Rate
    - Focuses on the environment capacity
  - ITR – Internal throughput Rate
    - Focuses on processor capacity
  - ITRRs – for relative improvements and degradations
  - LSPRs – IBM's assessment of a processor's capacity

Questions?