

# Watson & Walker

Software Pricing:  
The New Frontier in  
System Tuning

Frank Kyne  
frank@watsonwalker.com

# Welcome



- Who are we?
  - Watson & Walker was founded in 1985 by Cheryl Watson and Tom Walker.
    - Most famous for Cheryl Watson's Tuning Letter, published since 1991. Cheryl has spent her life working with performance, and loves nothing more than a good mystery.
  - There are now a total of 10 of us in Watson & Walker, a mixture of full and part time.
  - I joined Watson & Walker in 2014 as editor for the Tuning Letter and Cheryl's roadie.
    - Before that, I was a project leader in IBM's Redbooks group in beautiful Poughkeepsie, NY.
    - Before that, I was an MVS sysprog in IBM Ireland.



# Welcome



- Much of our recent growth has been driven by customers' focus on their z/OS software costs, and the bewildering array of products, actions, and offerings that can impact your bill.
- Given her vast performance experience, her reputation for integrity, her dedication to her customers, and her love of researching and problem solving, extending performance tuning to include software cost tuning was a natural fit.
  - In a brief moment of weakness, Cheryl confessed that if she had not gotten into IT, that she could have been a 'numbers gal'.
- Most of this session will be about looking at software pricing through a performance analysts eyes, but we will also throw in some good old-fashioned performance tips as well.



# Level Set

- Every other week we all see articles and webcasts about cutting MLC costs.
  - Many of these are focused on selling products or tools designed to drive down MSUs.
- We believe that reducing MLC costs is only one aspect of what should be a holistic exercise in delivering better value for money. It encompasses traditional performance tuning, an understanding of software pricing basics, exploitation of available pricing options, understanding that your actions can increase/decrease costs for other parts of your company, and driving better communications within your company.
- The objective of this session to get you as excited about software cost management as we are, and to help you see it as a technical, rather than a purely financial, challenge.

# SW Pricing Mgmt



- There are a few beliefs that you need to eject:
  - “SW pricing is boring”.
    - It is NOT! Managing your costs encompasses *all* of the knowledge that you have today, **plus** a new set of multi-dimensional considerations.
  - “Controlling software costs is simply a matter of limiting/reducing MSUs” It is NOT.
    - You need to stop thinking in terms of MSUs. Your execs probably don’t know or care what an MSU is. Your company budgets are created in terms of dollars, not MIPS or MSUs. And quite a few of the IBM pricing options break the direct connections between MSUs and the size of your bill, so you need to focus on \$s, not MSUs.
- There have been a lot of changes related to SW pricing over recent years, so keep an open mind and you might be pleasantly surprised with the results.






# SW Pricing Mgmt

- What can be achieved through prudent use of HW & SW functions and the right pricing options:

### Effect of Exploiting Available Options



# Traditional Performance Tuning

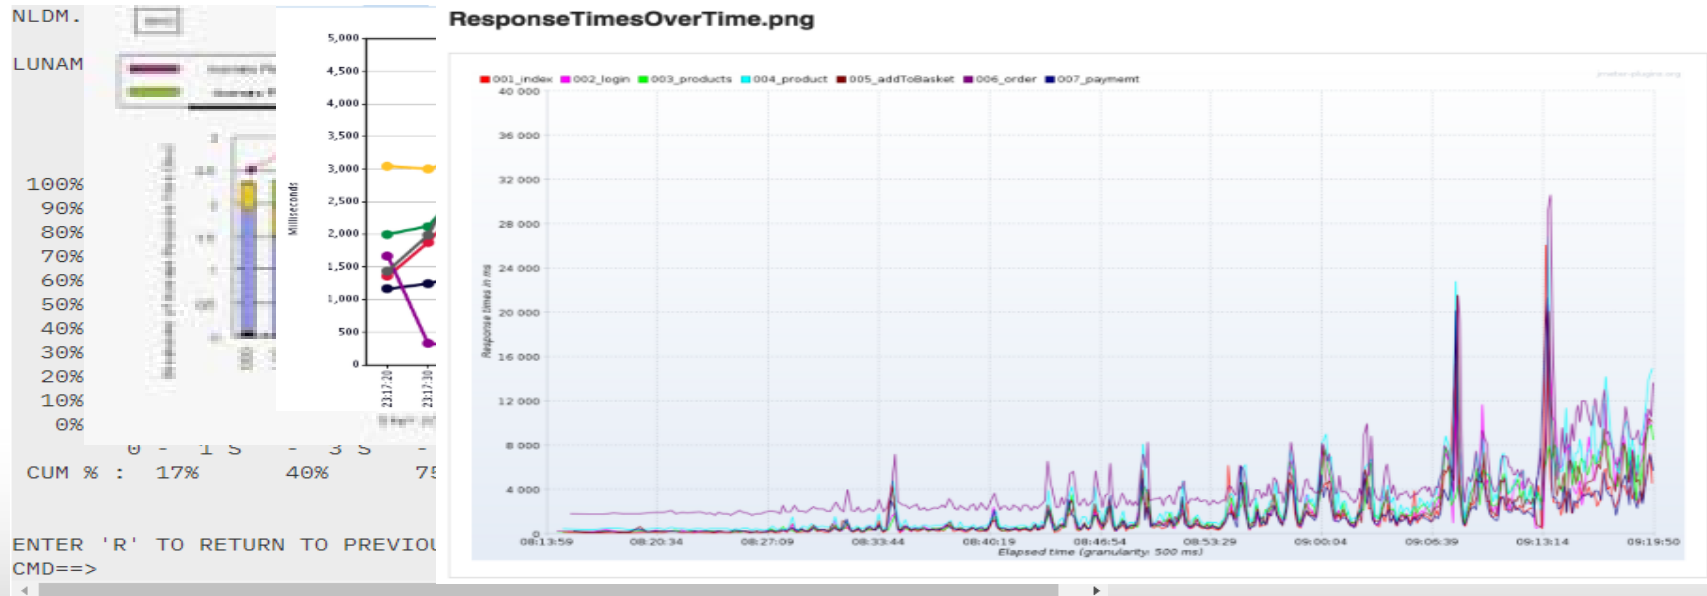


- The performance analyst role dates back nearly as far as z/OS itself.
- Early mainframes had limited supplies of capacity, high demands for that capacity, and very high costs.
  - The high costs drove the need for sharing (“virtualization”) – very few departments could afford their own S/360, so from Day 1 we have had:
    - Shared everything (CPU, memory, channels, disk, tape, network, etc).
    - Industry-leading metrics, so you could fairly apportion out the costs of providing the mainframe environment.
    - Performance as a core systems management discipline.
      - When something cost THAT MUCH, waste was not an option. In 1960, 1 (ONE!) MIPS cost \$10,000,000!



# Traditional Performance Tuning

- Traditional performance SLAs are based on transaction response times and batch turnaround.
  - How do we investigate/address performance problems?









# Traditional Performance Tuning

- Having identified the offending interval, we then need to understand the transaction flow, and the contribution of each response time component.





# Traditional Performance Tuning

- Why do we focus on the interval with the highest response time?
  - Because it is *probably* the one that is causing the most pain.
- We then drill down to see if the transactions are spending the bulk of their time *doing* something (what?) or *waiting* to do something (why is it waiting?)
- Eventually we determine and address (if possible) the root cause, or beg management for money to buy more resource....
- The specifics will vary depending on many things, **but the methodology (realize that there is a problem, drill down to find the cause, then address the issue) is consistent.**



# Traditional Performance Tuning

- The fix might be HiperDispatch tuning, WLM changes, faster DASD, larger buffers, better blocksizes, out of-spaceabend avoidance, smarter scheduling, GRS tuning, better/more efficient programs, new technology (e.g. zEDC), JES2 tuning, exploiting 1MB/2GB pages, recompiling application programs, turning on/off HSM compression, enabling zHyperWrite, tuning DFSORT, adding zIIPs, XCF tuning, more/less parallelism.... **You have a vast array of tools, but only a subset will address a given problem.**

# SCRT – The RMF of SW Pricing



- How about if we extend our key metrics to be:
  - Transaction response times.
  - Batch job turnaround.
  - Software dollars?
- Just as you have metrics to help with response time and batch job tuning, SW cost 'tuning' also has a set of metrics - to help you understand what is behind your bill.
- The methodologies that you use for traditional tuning are 100% relevant to tuning your software costs, and many of the actions that address response time or batch turnaround issues can also be applied to reducing your bill.
- *But*, your toolkit just got *much* bigger! 💰 - you now also have a host of pricing options.  
**Follow the money...**



# Quick Intro to Basics of Software Pricing



# SW Pricing Basics



There are three rules that hold true for just about anything to do with software pricing:

- **RULE 1:**

- There are multiple exceptions for every rule.
- Most exceptions have exceptions. Except in the exceptional case, where an exception has not been created yet. In this rare case, it is reasonable to have an expectation that an exception is under development.

- **RULE 2:**

- The one accurate answer to every question related to any aspect of software pricing is “It Depends”.

- **The Unwritten Rule 3**

# SW Pricing Basics - MLC vs IPLA



- MLC is Monthly License Charge.
  - Like renting – you pay for the product each month. The cost typically has some relationship to your capacity.
  - MLC tends to be the pricing metric for the most expensive products (z/OS, CICS, Db2, IMS, etc). As a result it usually gets the most focus.
- MLC is IBM's "bread & butter" – little or no room for horse trading, discounts are strictly controlled and documented.

# SW Pricing Basics - MLC vs IPLA



- IPLA is International Program License Agreement (aka One Time Charge).
  - Like purchasing – there is a large up-front cost – typically equivalent to 3 to 5 years of MLC.
  - However, if you want defect support and free upgrades to future versions, there is also a monthly “Subscription and Support” (S&S) charge.
- IPLA is used-car-salesman territory. Discounts can be as high as 90%. NO ONE should pay list price.
  - S&S can also be discounted, but discount is not necessarily the same percent.
  - The discount percentages depend on your negotiators skills and how desperate the sales team is to make a sale.
- IBM appears to be increasingly relying on IPLA to drive mainframe SW revenue growth..

# SW Pricing basics – Full vs Sub-Cap



- Nearly all ‘recent’ pricing metrics for MLC products (AWLC, CMLC, zNALC, etc) are based on ‘sub-capacity pricing’.
- IPLA products can be licensed at either full capacity or sub-capacity.
- Question: What exactly does *sub-capacity pricing* mean? Is the product charged based on:
  - The peak CPU usage of just that product?
  - The peak CPU usage of the LPARs that the product runs in?
  - The peak CPU usage of LPARs that some other product runs in?
  - The peak CPU usage of the CPCs that the product runs on?
  - The accumulated CPU usage of some or all of the above?
  - Something else?

# SW Pricing basics – Full vs Sub-Cap



It Depends!


- In general, MLC products are priced based on the combined peak R4HA of the LPARs where the product runs during that month.
  - However, the peak R4HA that is used to determine your bill might be reduced if you use price options such as Mobile Workload Pricing.
- A big advantage of Sub-Cap pricing is that it disconnects software costs from HW purchases, allowing you to purchase HW upgrades when you get the best price, but not paying associated software costs until you start using the additional capacity.
- We *strongly* recommend negotiating Sub-Cap licenses with your ISVs NOW, not when you are about to install the new CPC.
  - And get them to use the SCRT ISV support.



# SW Pricing basics – Full vs Sub-Cap



- For IPLA products, “Sub-Cap” means that the number of Value Units that you need is based on the highest R4HA to date.




- You only have to pay for your value units once (OTC); however, you have to pay for S&S on *all* the value units each month, regardless of the actual R4HA.
  - You also can't reduce the number of value units – you keep them *all*, or you throw them *all* away.

# SW Pricing Basics – R4HA



What is 'R4HA'? Rolling 4-Hour Average...

## R4HA vs Actual Consumption




The R4HA for an LPAR is reported by RMF, it is in the SMF type 70 records for that system, and it is used and reported on by IBM's Sub-Capacity Reporting Tool (SCRT).



# SW Pricing Basics – R4HA White Space

- *Generally* speaking (remember the exception rule), you pay for MLC SW based on its peak R4HA.
- Activity in the intervals *outside the peak* has no impact on your bill.
  - *Tuning* that work to reduce the MSUs will give you back capacity, but *not* reduce your software bill.
  - *Eliminating* work in those intervals will not save you a penny.
  - *ADDING* work in those intervals will also not *cost* you a penny.


R4HA vs Actual Consumption





# SW Pricing Basics - Pricing Curve

- IBM offers a 'bulk discount' on most software. The more you use, the less you pay for each additional MSU.





# SW Pricing Basics – Pricing Curve

- The *incremental* cost is the cost of the next MSU.
- The *average* cost is the cost across all your MSUs.
  - The average is ALWAYS higher than your incremental MSU cost.
- This means that:
  - *Increasing* peak R4HA by 10% will increase your bill by LESS than 10%.
  - *Decreasing* your peak R4HA by 10% will decrease your bill by LESS than 10%.
- This is a concept that is rarely factored into business cases when someone proposes to re-platform an application.
- Be wary of claims that infer that reducing your MSUs by x% by reducing your bill by the same percent.





# SW Pricing Basics – Your Bill

- Based on all this, you can see that your bill is influenced by:
  - Which products you have.
  - Which LPARs the products ran in.
  - The peak R4HA of each product.
  - If the products are MLC or IPLA (some are available as either).
- Given that your MLC bill typically goes up and down from month to month, how does IBM get all the information above in order to calculate your bill?
  - Sub-Capacity Reporting Tool (SCRT).

# SW Pricing Basics – Reporting to IBM



## How does SCRT know which products ran where?

- *Some* products (typically the newer/expensive ones) create SMF type 89 records that indicate that they are running in that LPAR in that interval.
  - This means that an interval is only included in the R4HA calculation for that product if it is actually running in that LPAR during that interval. **Think COBOL...**
- For products that *don't* create type 89 records, there are 'NO89' control statements in the SCRT job that you use to tell SCRT which CPCs/LPARs a product ran in during that month.
  - These NO89 statements are REALLY important. They are updated manually, so if you add or remove a product and don't tell the person that runs the SCRT job, your IBM bill will be wrong AND you could have an audit exposure.
  - Also, because the NO89 stmts only specify the LPAR name, SCRT assumes that the product was running in that LPAR in every interval in the month.
  - **NOTE:** If you discontinue using a product, you **MUST** notify IBM. Simply removing it from your NO89 statements will not stop you from being billed for it.

# SW Pricing Basics – Reporting to IBM



How does IBM know the peak R4HA for each product?

- The IBM type 70 records contain the R4HA for that system for that interval.

By combining the information from the type 70 and 89 records (and NO89 stmts) for each of your systems, SCRT is able to calculate the R4HA for each product for each interval.

Let's have a look at how that calculation works...

# SW Pricing Basics – Reporting to IBM



- In a **non-CMP** environment, SCRT determines peak for each CPC separately, and bill is based on sum of those peaks.

CPC <sub>1</sub>						CPC <sub>2</sub>					CPC <sub>3</sub>				AWLC SUM
LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	LP <sub>4</sub>	AWLC SUM		LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	AWLC SUM		LP <sub>1</sub>	LP <sub>2</sub>	AWLC SUM		
0:00	55	232	13	563	863	0:00	217	101	392	710	0:00	148	183	331	
1:00	64	481	49	246	840	1:00	276	392	384	1052	1:00	71	62	133	
2:00	60	454	15	255	784	2:00	235	382	65	682	2:00	179	288	467	
3:00	73	279	38	342	732	3:00	166	269	202	637	3:00	348	321	669	
4:00	75	257	37	671	1040	4:00	108	218	347	673	4:00	260	115	375	
5:00	52	442	32	329	855	5:00	369	86	122	577	5:00	450	123	573	
6:00	61	415	17	172	665	6:00	315	342	123	780	6:00	241	74	315	
7:00	75	406	12	168	661	7:00	366	293	155	814	7:00	148	340	488	
8:00	66	465	12	159	702	8:00	117	64	100	281	8:00	103	363	466	
9:00	68	374	18	390	850	9:00	154	264	347	765	9:00	446	155	601	
10:00	63	350	50	571	1034	10:00	266	83	220	569	10:00	229	399	628	
11:00	66	395	22	382	865	11:00	339	120	336	795	11:00	244	373	617	
12:00	52	459	24	263	798	12:00	342	247	318	907	12:00	304	211	515	
Peak					1040					1052				669	2761



# SW Pricing Basics – Reporting to IBM

- In a **CMP** environment, the R4HA is calculated across ALL CPCs in the multiplex, and your bill is based on *that* amount (plus an adjustment).

	CPC <sub>1</sub>					CPC <sub>2</sub>					CPC <sub>3</sub>			AWLC SUM	CMLC SUM	
	LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	LP <sub>4</sub>	AWLCS UM	LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	AWLCSU M	LP <sub>1</sub>	LP <sub>2</sub>	AWLCS UM				
0:00	55	232	13	563	863	0:00	217	101	392	710	0:00	148	183	331	1904	
1:00	64	481	49	246	840	1:00	276	392	384	1052	1:00	71	62	133	2025	
2:00	60	454	15	255	784	2:00	235	382	65	682	2:00	179	288	467	1933	
3:00	73	279	38	342	732	3:00	166	269	202	637	3:00	348	321	669	2038	
4:00	75	257	37	671	1040	4:00	108	218	347	673	4:00	260	115	375	2088	
5:00	52	442	32	329	855	5:00	369	86	122	577	5:00	450	123	573	2005	
6:00	61	415	17	172	665	6:00	315	342	123	780	6:00	241	74	315	1760	
7:00	75	406	12	168	661	7:00	366	293	155	814	7:00	148	340	488	1963	
8:00	66	465	12	159	702	8:00	117	64	100	281	8:00	103	363	466	1449	
9:00	68	374	18	390	850	9:00	154	264	347	765	9:00	446	155	601	2216	
10:00	63	350	50	571	1034	10:00	266	83	220	569	10:00	229	399	628	2231	
11:00	66	395	22	382	865	11:00	339	120	336	795	11:00	244	373	617	2277	
12:00	52	459	24	263	798	12:00	342	247	318	907	12:00	304	211	515	2220	
Peak					1040					1052				669	2761	2277



# SW Pricing Basics – Reporting to IBM



The SCRT job creates a CSV file that you must send to IBM by the 9th of each month, and that is input to the IBM process to create your bill for that month.

- If you discover an error **AFTER** you send the CSV file to IBM, getting them to accept a corrected CSV file can be a challenge.

# SW Pricing Basics




- The last point before we move on is that your IBM bill is probably based on a *number* of products.
- The cost of *each* product is calculated independently.
- The peak R4HA *for each product* is worked out based on the R4HA of the intervals that that product ran in each LPAR.
  - This means that each product could have a different peak R4HA interval.
- This is a **very important point**. Everyone talks about “tuning your peak R4HA” as if it was one number. In practice it rarely is, unless you have a very small number of LPARs.
  - When deciding where to focus your tuning efforts, you need to know which products are costing you the most, and then address the peak intervals for those products.

# SW Pricing Basics



- Not only might the peak intervals be different, the *relative* prices of the different products might be different as well, depending on where you are on the pricing curve.
- For example, in the chart below, 3 MSUs of CICS costs 50% more than 3 MSUs of z/OS, so you might want to focus on CICS if you are very small. But in the 46-175 MSU range, z/OS costs >2x as much per MSU as CICS.
- Always validate your assumptions before taking action.

## Relative Price Points



# This is where the fun starts!



- Now that we have all the easy stuff out of the way, let's look at some examples of how you might start tuning your software costs....



# Tuning Software Costs - Metrics

- There are two key sources of information to help you effectively improve the value for money for your company's z/OS investment:
  - SCRT Reports.
  - IBM Workload Pricer reports – for most customers, these are effectively your MLC invoices.
- By combining information from these sources, you can identify which products to focus your efforts on, and in which intervals.



# SCRT Reports

- In addition to the standard reports that are required by IBM, SCRT also produces a number of optional reports.
- The [SCRT User's Guide](#) describes all the reports. Perhaps the most interesting ones are:
  - **B5/E5** – this shows the number of MSUs for each MLC product that IBM will use when calculating your bill.
  - **P5** – shows the number of MSUs for each LPAR a product ran in, for the peak R4HA interval for that product.
  - **Q5** – for products that create type 89 records, this shows the % of intervals when the product was active in each LPAR.
  - **V5/W3** – for each product, shows the MSUs for every interval for every LPAR where that product ran.



# SCRT Reports

- The B5/E5 sections should be used to help track month to month changes in peak R4HA for each product.
  - Might highlight exception cases that can be excluded.

MLC Product Name	Number	MSU	Time	CPC1	CPC2	CPC3
z/OS	(All)	520	01 Nov 2017 - 02:00	102	241	177
z/OS V2	5650-ZOS	520	01 Nov 2017 - 02:00	102	241	177
DB2 for z/OS	(All)	520	01 Nov 2017 - 02:00	102	241	177
DB2 11 for z/OS	5615-DB2	520	01 Nov 2017 - 02:00	102	241	177
CICS	(All)	418	01 Nov 2017 - 02:00		241	177
CICS TS for z/OS V5	5655-Y04	418	01 Nov 2017 - 02:00		241	177
MQ for z/OS	(All)	520	01 Nov 2017 - 02:00	102	241	177
IBM MQ for z/OS V9	5655-MQ9	520	01 Nov 2017 - 02:00	102	241	177
IMS	(All)	279	01 Nov 2017 - 02:00	102		177
IMS V14	5635-A05	279	01 Nov 2017 - 02:00	102		177
COBOL	(All)	520	01 Nov 2017 - 02:00	102	241	177
Enterprise COBOL for z/OS V6	5655-EC6	5	31 Oct 2017 - 11:00		5	
IBM Enterprise Cobol for z/OS V4	5655-S71	520	01 Nov 2017 - 02:00	102	241	177
IPLA Product Name	Number	MSU	Time			
IPLA z/OS-Based	(All)	520	01 Nov 2018 - 02:00	102	241	177

This is the  
CMP  
version (B5)

Recommend  
specifying  
MLC\_ALL






# SCRT Reports

- The Q5 section might help you spot cases where a product was accidentally started in the wrong LPAR.

```
==Q5=====
```

PRODUCT GRID SNAPSHOT						
Product N:	Product ID			FPK1	FPK2	FPK3
z/OS	(All)			100.00%	82.90%	100.00%
z/OS V2	5650-ZOS			100.00%	82.90%	100.00%
DB2 for z/	(All)			100.00%	82.90%	100.00%
DB2 11 for	5615-DB2			100.00%	82.90%	100.00%
MQ for z/c	(All)			100.00%	0.13%	100.00%
IBM MQ for	5655-MQ9			100.00%	0.13%	100.00%
IMS	(All)			99.80%	82.90%	100.00%
IMS V14	5635-A05			99.80%	82.90%	100.00%
COBOL	(All)			100.00%	82.90%	100.00%


Does MQ *normally* run in this LPAR?



# Using the SCRT Info



Impact of MWP on R4HA



The SCRT report helps you find the interval that you want to focus on

# IBM SW Pricing Options



- Armed with information about the intervals you want to focus on, and your knowledge of your workloads, you can investigate whether some of the available IBM SW pricing options would help you:
  - MWP reduces the R4HA for all products in an interval by 60% of the MSUs used by qualifying mobile transactions.
    - If you can identify those transactions using the WLM classification rules, you can use WLM to gather, track, and report the MSU consumption of those transactions.
  - zWPC offers the same discount for transactions that originated in the cloud.
  - zCAP reduces the R4HA of products that a new application does NOT use by 100% of the MSU consumption of that new application. z/OS R4HA is reduced by 50%. The R4HA of products used by that application is not reduced at all.
    - zCAP is intended for collocated applications. Conceptually it is similar to zNALC/MzNALC, except that it doesn't require a dedicated LPAR.
- Only you can tell whether any of these would apply to you. You might have a lot of mobile work, but if it is all during the day, and your peak R4HA is in the middle of the batch window, it might not save you much.



# IBM SW Pricing Options - CMP

- A few thoughts on CMP...
- Unlike MWP, zCAP, etc, CMP was intended to be revenue- and cost-neutral.
- To achieve this cost-neutrality, an adjustment is applied to your CMP software bill, the objective being to make sure that, all other things being equal, you would pay the same after you switch to CMP as you did before the switch.
  - The adjustment immediately became known as the 'CMP uplift', giving the impression that it increases your software bill – that is not the intention. But I think that this is misleading, so we try to stick with 'adjustment'.
- The adjustment is effectively a (nearly) fixed cost that will be applied to your bill forever more. The amount of the adjustment is based on your costs over the 3 months before you move (but *you* get to pick when you move).
- If the uplift is a fixed cost, why do we care?



# IBM SW Pricing - CMP

- These are based on real customer numbers, over 5 months:

	Aug Base	Aug Diff	Aug Factor	Sep Base	Sep Diff	Sep Factor	Oct Base	Oct Diff	Oct Factor	Nov Base	Nov Diff	Nov Factor	Dec Base	Dec Diff	Dec Factor
z/OS	4697	121,101	41.27%	4814	120,927	40.66%	4653	117,077	40.10%	4685	119,658	39.26%	4721	126,185	41.23%
HSM	4697	9,647	35.14%	4814	9,591	34.42%	4653	9,272	33.97%	4685	9,423	33.07%	4721	10,065	35.16%
SDSF	4697	3,267	41.01%	4814	3,256	40.26%	4653	3,148	39.74%	4685	3,203	38.73%	4721	3,388	40.77%
C/C++	4252	6,581	44.35%	4814	6,518	43.19%	4164	6,466	44.04%	4143	6,706	44.02%	4147	7,121	46.72%
RMF	4697	5,483	35.79%	4814	5,465	35.19%	4653	5,297	34.75%	4685	5,395	33.90%	4721	5,752	35.99%
CICS V5	4262	66,488	22.93%	4260	71,361	24.62%	4148	72,031	25.37%	4252	75,216	24.99%	4367	77,335	25.16%
DB2	4262	58,092	22.70%	4260	62,399	24.40%	4145	63,030	25.18%	4249	65,830	24.80%	4364	67,724	24.98%
Cobol V4	639	2,777	25.93%	673	2,051	18.58%	758	1,277	10.77%	786	604	4.79%	788	423	3.34%
<b>Total</b>	4697	273,436	29.94%	4814	281,569	30.65%	4653	277,598	30.79%	4685	286,036	30.11%	4721	297,993	30.93%
Diff Max/Min of Diff		24,556													
Diff Max/Min of Base		161													

Depending on which months you select, the monthly adjustment could **vary** by as much as this, FOREVER

Note that the uplift % is different for every product



# IBM SW Pricing - CMP

- CMP offers fantastic technical benefits for anyone with more than one CPC, and especially for multi-site sysplexes.
- Common wisdom is that CMP is only attractive to growing customers, but I believe that the benefits might outweigh the additional cost even if your z/OS workload is declining.
- In any case, if you haven't moved to CMP yet, it is definitely something that you should investigate and then be smart about selecting your base period.



# IBM SW Pricing Options



- IBM's latest pricing initiative is "Container Pricing".
- Container Pricing is a set of pricing options that use new infrastructure in WLM, SMF, RMF, SDSF, z/OSMF, and SCRT. The infrastructure delivers the flexibility to have qualifying applications collocated with other applications, or running in dedicated LPARs, or both.
  - Container Pricing involves a Solution ID that is assigned in WLM (or in the SCRT job, for dedicated LPARs) and is used by IBM to determine how much to charge for that workload.
  - SCRT understands that the R4HA of all containers should be removed from the R4HA of all other work. The SCRT file contains the Total R4HA, the container R4HA, and the Solution ID.





# IBM SW Pricing Options

- The DevTest Container offering is the most popular of the Container Pricing options announced so far:
  - Fixes the MLC cost of your pre-production workloads.
  - Lets you grow the peak R4HA of that container up to 3X the current DevTest peak without increasing MLC costs of those LPARs.
  - Sysprogs will love you because you can now run those products in any/all of the DevTest LPARs with no change in the cost of the DevTest container.
  - Your Application Developers will no longer slash your tires.
  - The gotcha *might* be IPLA products – you must have sufficient Value Units *today* for the maximum container size that you choose.

# IBM SW Pricing Options



- The other interesting new pricing option is the New Application Solution.
- There are two flavors of this:
  - Original – Pricing for the container is based on z/OS on MLC - all other products used by the new application *must* switch to IPLA.
    - AND you must specify the container size and pay for that number of Value Units up front.
  - Spicy – Pricing for the container is based on the accumulated CPU time for the month.
    - Two rates – base, and 20% discount if you commit to paying for at least 25K MSUs/month.
    - Prices are based on new SCLC metric.

# IBM SW Pricing Options



- The New Application Solution option that lets you pay based on accumulated CPU time is a very interesting model:
  - If the application uses *any* CPU time, at any time in the month, you must pay for it.
    - Compare this to the R4HA white space, where work that only runs during those times is effectively free.
  - However, it is closer to what people expect from cloud pricing.
    - There is no 'it depends'. You pay \$x per MSU, regardless of when it runs or what else is running at the same time.
    - This predictability is attractive to some people.
  - However, you need to think about whether you need to have a way to limit the potential cost. You CAN do that today, with soft-capping. How would you achieve the same result if you switched to an accumulated CPU consumption model?



# IBM SW Pricing Options

- As with nearly all SW pricing options, we believe that there are some that will make sense for you, and others that will not.
- The only way to know for sure is to model it.
  - We recommend going back one year (to allow for seasonal workload variability) and comparing how much you *actually* paid to how much you *would have paid* if you used the new metric.
- See Container Pricing article in latest Cheryl Watson's Tuning Letter ([2018 No. 3](#)), for a detailed discussion about the pros and cons of each option.

# IBM SW Pricing Options AND YOU



- Why do you care about pricing options? Because they should influence how you configure your systems. Using MWP as an example:
  - CICS sysprogs should work with contract mgmt because:
    - If you use a shared CICS region, you only get to 'claim' the CPU time as reported in the CICS 110 records – this is typically about 80% of the total CICS CPU time.
    - If you use dedicated CICS regions for MWP-eligible txns, you get to claim the CPU time for the entire region.
    - Similarly for MQ, Db2, etc.
  - CICS, WLM, and Network techies, contract mgmt, and application architects should work together to agree a standard to allow the source of transactions to be identified by the WLM classification rules.

# IBM SW Pricing Options AND YOU



- How about zIIPs?
  - Db2 V11 and V12 make more of its processing zIIP-eligible IF (and only if) IIPHONORPRIORITY is set to YES.
  - But setting IIPHONORPRIORITY to YES can result in more zIIP-eligible work overflowing to general purpose CPs.
  - However, WLM APAR [OA50845](#) lets you control at the service class level whether zIIP-eligible work can overflow to a general purpose CP.
    - If you have zIIP-eligible CPU-intensive work that is not response time-critical, you can force *it* to only run on zIIPs while still letting Db2 to overflow to CPs if necessary.
  - You can use SMF 30 or 72.3 records to identify how much zIIP work ran on a GCP. With your new-found expertise with SCRT and WL Pricer, you can now put a real cost on that, and do a cost analysis of the benefit of adding zIIP capacity.

# IBM SW Pricing AND YOU



- More about zIIPs...
- IBM recently delivered PTFs for a load of DB2 and IMS tools that provide the ability to make some of their processing zIIP-eligible – YEAH!
- However, the function is turned OFF by default. Boo – well, maybe...
  - There is a CPU cost for the product to create the enclave that will be used for the zIIP work – thanks to Kathy Walsh for this tip. Unfortunately, the PTFs don't say anything about this.
- In one case, if you search hard enough, you will discover that you turn on the zIIP support by including a “//AUZIIP DD DUMMY” statement in the JCL of the associated started task.
- This is just one example of the valuable information that can be obtained by reviewing IBM APARs.
  - The only problem is that there are a LOT of them, and they are not always flagged in a manner that guarantees that you will find them.





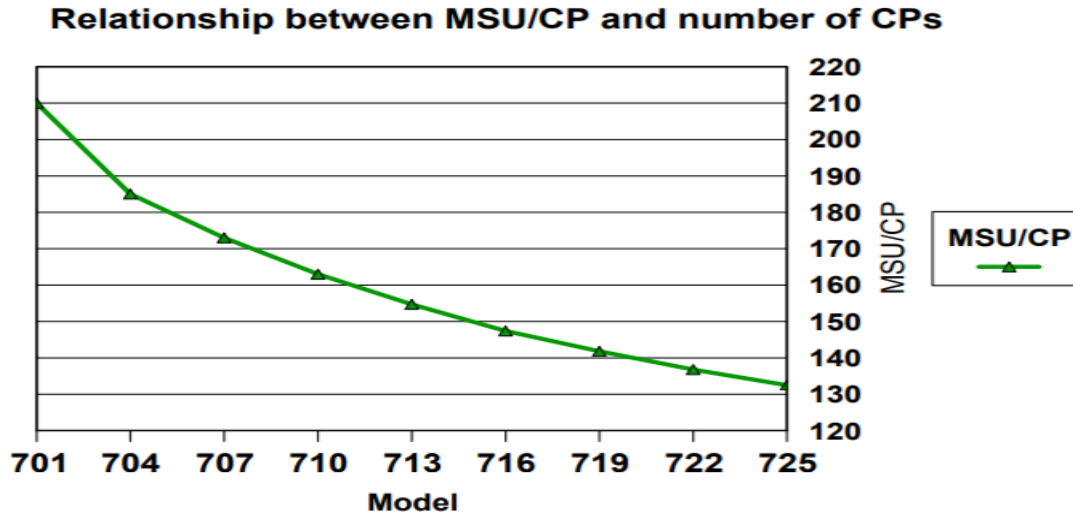
# IBM SW Pricing AND YOU

- Speaking of adding capacity... There is a long-held belief that a busy z/OS is a happy z/OS.
  - On more recent CPCs, that isn't the case.
  - According to an [IBM White Paper](#) by **Gary King**, every 10% increase in CPC utilization increases CPU time/txn by 3-5%.
    - Note that this effect is NOT factored into zPCR or zCP3K studies.
  - Various [Tuning Letter articles](#) and award-winning SHARE presentations by **Todd Havekost** show the significant advantages of running at lower utilizations.



# IBM SW Pricing AND YOU

- I just love this chart (courtesy of Todd)...



These values are based on the official MSU ratings, extracted from IBM LSPR reports.

The LSPR numbers are based on CPCs running at 90+% utilization.

- The reducing MSUs/CP reflects the increased overhead due to the volume of work on larger models. But what would happen if you put a 710's-worth of work on a 725?

# IBM SW Pricing AND YOU



- One of our customers, NASCO, recently upgraded from a z13 709 (12,371 Average RNI MIPS) to a sub-capacity z14 523 (12,584 Average RNI MIPS).
- Following the upgrade and configuration optimization, their peak R4HA MSUs dropped by 22%!
  - **NOTE: NOT EVERYONE WILL ACHIEVE THE SAME RESULTS.**
- Their RNI went from AVG-HI to AVG-LOW as a result of having so much additional cache (23 CPs vs 9) and more VH CPs.
  - Unfortunately, there are no tools to model how a change in the number of CPs will affect your RNI.
- When looking at your next CPC upgrade, PLEASE don't discount sub-capacity models without at least including them in your list of candidates.
  - If you would like help estimating the affect of changed CP numbers on your RNI, contact IBM or W&W.
- Remember that CPs are a 1-time expense, but your SW bill arrives every month. Compare the SW cost savings (including the MSU-based TTO discounts) over the life of the CPC to the cost of additional CPs, then find the sweet spot.

# My Hobbyhorse



Is cloud pricing REALLY as simple as we are lead to believe?

**Menu** **aws** **Contact Sales** **Products** **Solutions** **Pricing** **Getting Started** **Documentation** **More** **English** **My Account** **Create an AWS Account**

**Cloud**  
Amazon Web Services  
analytics, new  
These services  
and the hottest  
development

**Compute**  
Amazon EC2  
Amazon EC2 Auto Scaling  
Amazon Elastic Container Service  
Amazon Elastic Container Service for Kubernetes  
Amazon Elastic Container Registry  
Amazon Lightsail  
AWS Batch  
AWS Elastic Beanstalk  
AWS Fargate  
AWS Lambda  
AWS Serverless Application Repository  
Elastic Load Balancing  
VMware Cloud on AWS

**Storage**  
Amazon Simple Storage Service (S3)  
Amazon Elastic Block Store (EBS)  
Amazon Elastic File System (EFS)  
Amazon Glacier  
AWS Storage Gateway  
AWS Snowball  
AWS Snowball Edge  
AWS Snowmobile

**Networking & Content Delivery**  
Amazon VPC  
Amazon VPC PrivateLink  
Amazon CloudFront  
Amazon Route 53  
Amazon API Gateway  
AWS Direct Connect  
Elastic Load Balancing

**Developer Tools**  
AWS CodeStar  
AWS CodeCommit  
AWS CodeBuild  
AWS CodeDeploy  
AWS CodePipeline  
AWS Cloud9  
AWS X-Ray  
AWS Tools & SDKs

**Management Tools**  
Amazon CloudWatch  
AWS Auto Scaling  
AWS CloudFormation  
AWS CloudTrail  
AWS Config  
AWS OpsWorks

**Machine Learning**  
Amazon SageMaker  
Amazon Comprehend  
Amazon Lex  
Amazon Polly  
Amazon Rekognition  
Amazon Machine Learning  
Amazon Translate  
Amazon Transcribe  
AWS DeepLens  
AWS Deep Learning AMIs  
Apache MXNet on AWS  
TensorFlow on AWS

**Analytics**  
Amazon Athena  
Amazon EMR  
Amazon CloudSearch  
Amazon Elasticsearch Service  
Amazon Kinesis  
Amazon Redshift  
Amazon QuickSight  
AWS Data Pipeline  
AWS Glue

**Security, Identity & Compliance**

**AR & VR**  
Amazon Sumerian

**Application Integration**  
Amazon MQ  
Amazon Simple Queue Service (SQS)  
Amazon Simple Notification Service (SNS)  
AWS AppSync  
AWS Step Functions

**Customer Engagement**  
Amazon Connect  
Amazon Pinpoint  
Amazon Simple Email Service (SES)

**Business Productivity**  
Alexa for Business  
Amazon Chime  
Amazon WorkDocs  
Amazon WorkMail



**Desktop & App Streaming**  
Amazon WorkSpaces  
Amazon AppStream 2.0

**Internet of Things**

[https://aws.amazon.com/products/?nc2=h\\_q\\_l\\_p&awsml=ql-1](https://aws.amazon.com/products/?nc2=h_q_l_p&awsml=ql-1)

# Cloud Pricing



Google   

[All](#) [News](#) [Shopping](#) [Videos](#) [Images](#) [More](#) [Settings](#) [Tools](#)


About 24,600,000 results (0.60 seconds)

**What is Cloud Management? | Get Started w/ BMC Multi-Cloud**  
[www.bmc.com/cloud/management](#)  
Cloud Migration, Cost Control, Security, Discovery, Automation & Performance. Cloud Performance, Cloud Security, Cloud Cost Control, Cloud Automation, Cloud Visibility. Features: Optimize Cloud Resource Usage And Cost, Service Management, Keep Track Of Cloud Assets. Customer Stories · YouTube Videos · Contact Sales · Free Trials

**Manage Multiple Clouds | Visit HPE To Learn More | HPE.com**  
[www.hpe.com/Multiple/Clouds](#)  
Enterprise-class security, scalability and performance. Learn More Today! Discover Your Right Mix.

**Bill Management Done Right | in 2018 | stamplicom**  
[iwant.stamplicom/Bill/Approval](#)  
Watch our 5 minutes webinar to learn what is the Stamplicom hype is all about! Paperless AP Management. No Training Required. Unlimited Number of Users. Super Easy Implementation. Easy Intacct Integration. Simple ERP Integration. No Setup Fees. User-friendly Interface. Fast NetSuite Integration.

**Cloud Management Solutions | Optimize Cloud Deployment**  
[www.netscout.com/](#)  
Manages The Delivery Of Cloud Services. Cost-Effective And Scalable. Full Cloud Visibility. NETSCOUT Cloud Solutions · Service Provider Cloud · Cloud Smarter INFOGRAPHIC

**How to manage your cloud bill - Managed Data Center News**   
[resource.onlinetech.com](#) > Resources > Blog  
May 15, 2017 - Cloud bill management Part of the cloud's blessing and curse is your monthly cloud bill.

24.5 Million hits on managing your cloud bills? Maybe it is not as simple as some people like to think.

If you are interested in a short article about reality of cloud pricing, email us at [technical@watsonwalker.com](mailto:technical@watsonwalker.com).

# A Closing Thought



- Just about all customers have some sort of Enterprise License Agreement with IBM and other vendors.
- Depending on the type of agreement you have, any over-payments you make might be available to put towards other (specified) software, hardware(?), or services purchases. But, if you don't use it, you might lose it.
- On the other hand, if your actual bills for the year exceed the projected amount, IBM and others will kindly send you a lovely present – an invoice for the shortfall.
- So, one of your year-end objectives should be to see how your actual usage compares to your scheduled payments (called CBA), and if there is anything that you need that you can purchase with any credit you might have.



# Thank you!



- That's all folks!
- If you would like to discuss any of this with us, drop us an email ([technical@watsonwalker.com](mailto:technical@watsonwalker.com)) or have a look at our [SW Pricing workshop](#) or [SCRTPro offering](#).
- Thank you for coming. We hope you all have a great holiday season, and that we'll see you again some time soon.