

# Issues That Arise in MLC Software Pricing

Cheryl Watson

Watson & Walker, Inc.

[technical@watsonwalker.com](mailto:technical@watsonwalker.com)

# Major Issues With Software Contracts

- Background
  - We review dozens of software contracts each year
  - We teach software pricing classes on-site and public
  - We consult on software contracts (IBM and ISVs) and audits
  - We have offering to run SCRT and provide recommendations to reduce the R4HA and other costs (SCRTPro)
  - ISV contracts often keep customers from expanding their capacity
- Agenda
  - MIPS vs MSUs
  - Increasing Capacity to Reduce MIPS
  - Country Multiplex Pricing (CMP)
  - Container Pricing

# MIPS vs MSUs

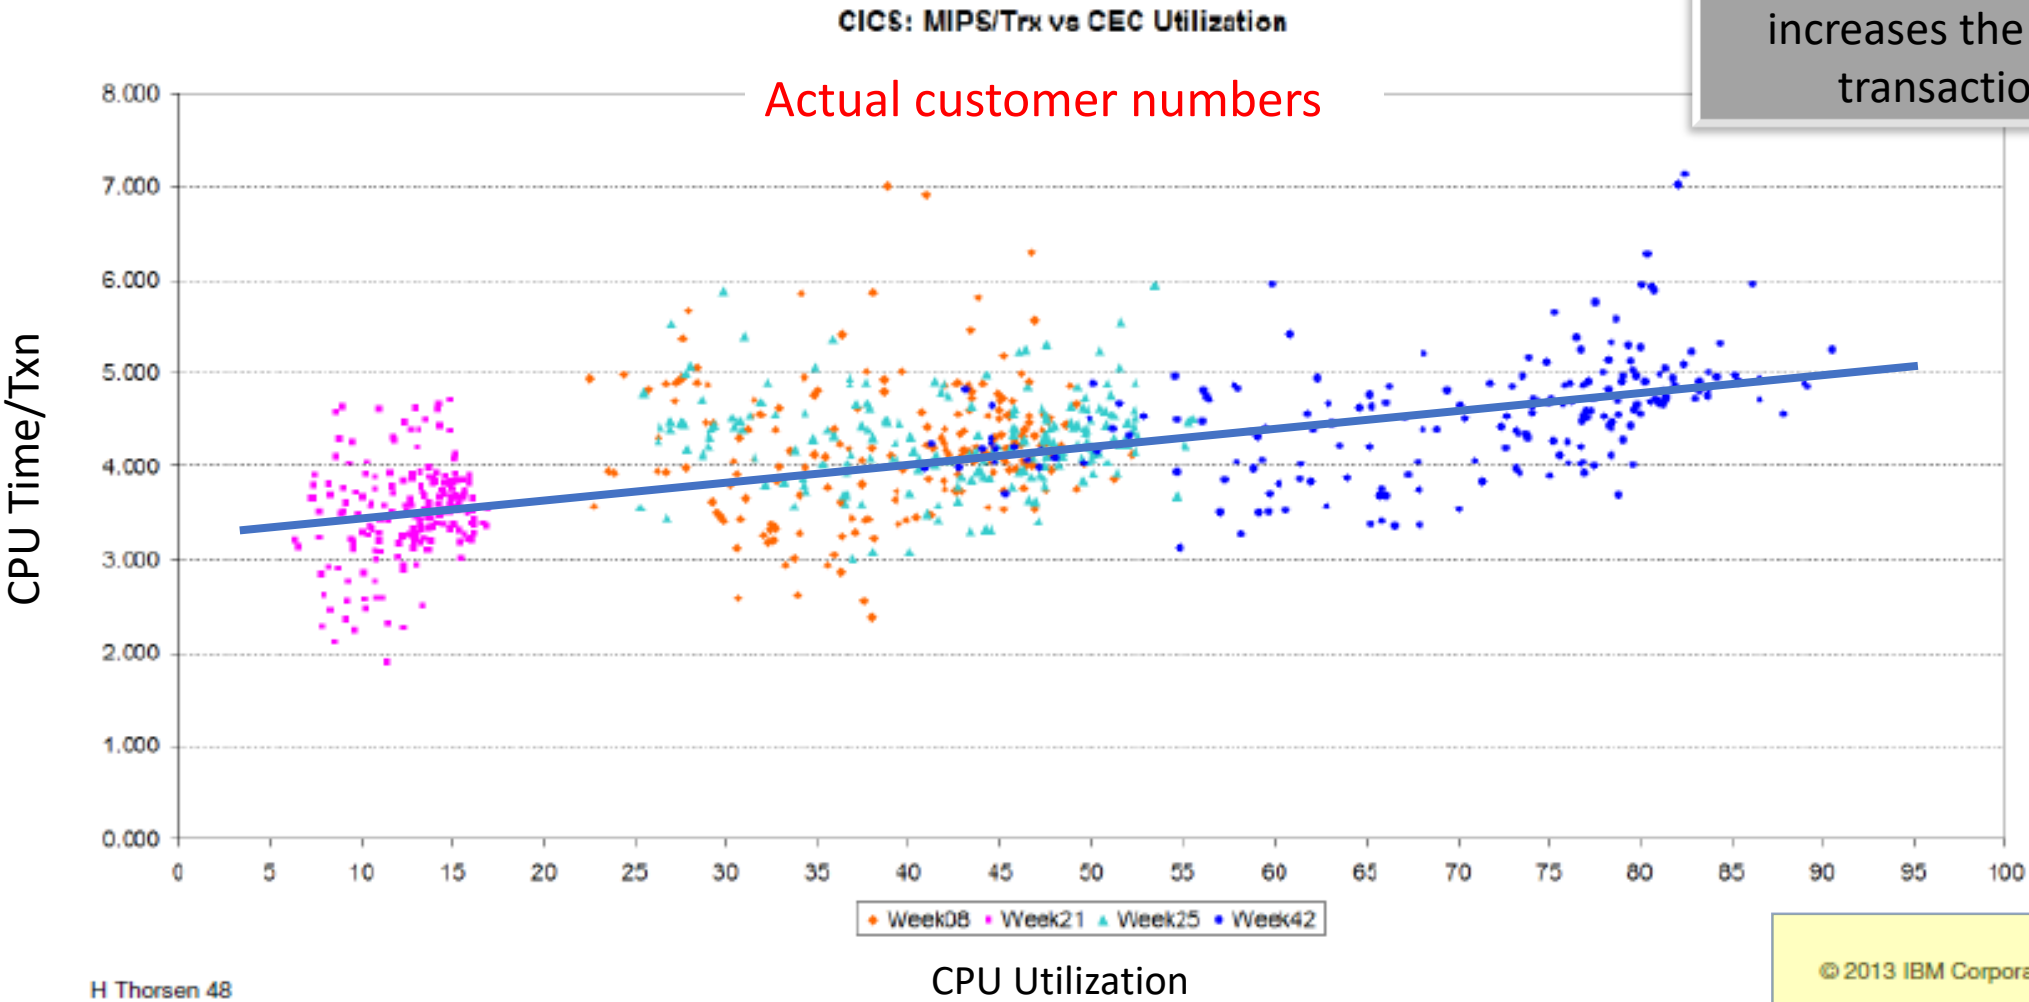
- Too many contracts specify MIPS without defining them
- Which MIPS? Gartner/IBM/Cheryl Watson/calculated from MSUs?
- Which MSUs? R4HA/SCRT N5 or P5/ISV SCRT/CMP/Container?
- Solution: change contracts to use MSUs; specify how MSUs are determined; **use ISV SCRT facility!**
- Enterprise Systems Media: *Enterprise Executive 2016: Issue 4* – [Mainframe Software Audits](#)

# Increasing Capacity to Reduce MIPS

- CPU time used increases by 3-5% as the CPU utilization increases by 10%
- Solution: add excess capacity (e.g. go from 90% busy to 50% busy can reduce the peak MSUs by up to 20%)
- Problem: all IBM full-capacity and all ISV full-capacity products are show stoppers for this solution
- One example: USAA went from four zEC12-711 to z13-711 and MSUs went up; upgrading to z13-716s reduced MIPS by 9000 MIPS; upgrading to z13-726s reduced MIPS by another 4000 MIPS; See SHARE in Providence session [21045](#) by Todd Havekost

# Less is more

IBM Rule of thumb is that every 10% increase in CPU utilization increases the CPU time per transaction by 3-5%



© 2013 IBM Corporation

# Use of Excess Capacity

- When you add engines, the amount of cache increases, the CPU utilization goes down, the RNI decreases, and the cost of sub-capacity software goes down because the peak R4HA goes down
- This is one reason that sub-capacity processors are becoming more popular
- More and more companies are attempting this option
- Big complaint is ISVs who won't support sub-capacity pricing
- Solution: support sub-capacity pricing; **use ISV SCRT facility!**

# Country Multiplex Pricing

- MANY customers want to go to CMP
- Result is that the R4HA is reduced
- IBM adjusts pricing during invoicing with a “base factor” so that their revenue doesn’t go down
- But ISVs that use IBM’s SCRT simply see the MSUs go down
- Most ISV contracts make no mention of CMP
- Solution: **use ISV SCRT facility and request AWLC report instead of CMLC!**

# CMP – Different R4HA Calculation Method

- When using CMP, your peak R4HA is calculated by summing the MSUs for LPARs across ALL CPCs, *not* on a CPC-by-CPC basis.
- The *worst* case is that the CMP R4HA will be the same as the pre-CMP R4HA. In practice, it should nearly always be less.

AWLC SUM = 1040 + 1052 + 669 = 2761  
 CMP SUM = 2277

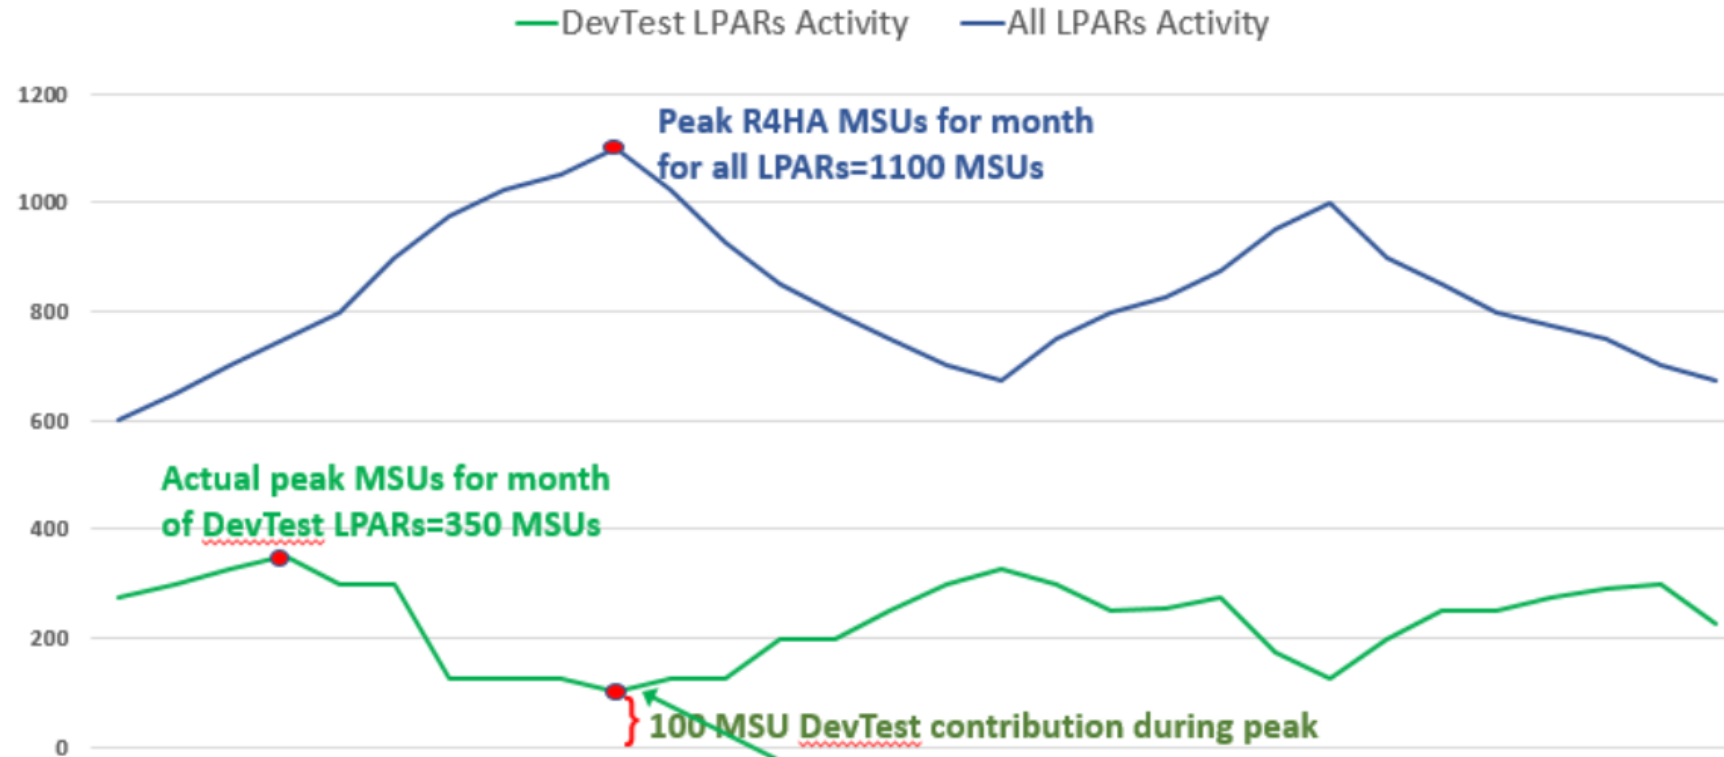
	CPC <sub>1</sub>						CPC <sub>2</sub>						CPC <sub>3</sub>			AWLC SUM	CMLC SUM	
	LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	LP <sub>4</sub>	AWLCS UM		LP <sub>1</sub>	LP <sub>2</sub>	LP <sub>3</sub>	AWLCS UM		LP <sub>1</sub>	LP <sub>2</sub>	AWLCS UM				
0:00	55	232	13	563	863		0:00	217	101	392	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



# Container Pricing

- LOTS of customers are interested in new container pricing, primary Dev/Test
- IBM says “Container Pricing **does not directly impact** the cost of unrelated workloads.”
  - But it indirectly impacts them; primarily due to the CPU utilization
  - In the next slide, look at what happens when you let the developers go wild during peak periods
- The cost per MSU for traditional workload may increase (see last slide)
- If ISV uses IBM SCRT for charging; they may need to change
- Solution: **use ISV SCRT facility!**

# DevTest Solution Example



In this example, actual peak for DevTest is seen as 350 MSUs so up to 1050 MSUs permitted at no extra MLC cost (subject to HW MIPs and zIPLA SW coverage being available)

Contribution of DevTest to peak R4HA used for billing. \$ value calculated by excluding from SCRT run and calculating delta versus BAU total \$ cost based on announcement letter terms – this becomes the MLC \$ cost for DevTest workload going forward.

# Container Pricing


- Traditional pricing may be best suited to applications that run outside the peak.

Peak MSUs for new application is 100 MSUs.

Contribution of new application to Peak R4HA is nearly nothing.

Running it in a container (with a 100 MSU peak) would presumably cost more than treating it the same as all your traditional work.


New Application Using R4HA Whitespace



# Container Pricing


- Container pricing may be best suited to applications that use a lot of capacity *in the peak*

New Application Using Container




OR

Traditional



Container



- If you treated it as traditional work, the peak R4HA for ALL work would be 310 MSUs.
- If you use Container Pricing for the new application, your Peak R4HA for everything else would be 210 MSUs, PLUS something (???) for the container?

# It's a New World

