Country Multiplex Pricing
The Gift That Keeps Giving, But To Whom?

Frank Kyne – Watson & Walker

frank@watsonwalker.com

Session 21601

Welcome

- Thank you for attending this session – we hope to give you lots of useful tips to bring home.
- Who are we?
  - Watson & Walker Inc established 1986; Cheryl has been working on IBM mainframes since 1965; Frank joined Watson & Walker in 2014.
  - Publish Cheryl Watson's Tuning Letter (since 1991).
    - Now available to subscribers online at www.watsonwalkerpublications.com
  - Teach classes, consult, have 2 software products, SCRTPro Service Offering
  - z/OS evangelists, Subject Matter Experts in Software pricing, Parallel Sysplex, and Workload Manager.
- If you have questions, please ask as we go along.
- Disclaimer – in this presentation we focus on IBM sub-capacity MLC products. You should always look at ALL your costs (IBM, ISV, MLC, OTC, sub-cap, full-cap, HW, and Maint) when making contract or configuration decisions.

Agenda

- Intro to the problem
- Intro to the solution
- Digging a little deeper – making sure that you get the maximum value from CMP
- Questions?

Valuable tip

- There is an old saying that might help you maintain your sanity in the world of software pricing:

  “God, grant me the serenity to accept the things I cannot change, the courage to change the things I can, and the wisdom to know the difference.”

- The things I *cannot* change – Most companies today are in business to make as much profit as possible.
- The things I *can* change – Use sub-capacity instead of full-capacity pricing, use large pages, use zEDC, maximize use of zIIPs, optimize HiperDispatch topology, use Mobile Workload Pricing, use sub-cap pricing for ISV products (based on new SCRT ISV support), and hundreds of other potential actions.

The Problem

- What ‘challenges’ is Country Multiplex Pricing (CMP) meant to address?
  - Some IBM software pricing mechanisms effectively encourage customers to configure systems in ways that make no technical sense (‘Shamplexes’).
  - Pre-CMP pricing mechanisms penalized customers for moving workloads between CPCs and discouraged full exploitation of dynamic workload routing.
  - The qualification criteria for Sysplex Aggregation were complex to understand and difficult to enforce/manage.
  - The requirement that certain systems had to be in the same sysplex in order to ‘aggregate’ the associated CPCs limited the distance between data centers.
  - Single Version Charging and IPLA migration limits interfered with normal business cycles.

  - Let’s look at these in a little more detail…..

The Basics – How Your IBM Software is Calculated

- How IBM arrives at your (pre-CMP) monthly z/OS-based software bill:
  - DO THIS For each product:
    - For each CPC:
      - Identify the peak Rolling 4-Hour Interval (in MSUs) in the month for that product (A).
    - For each CPC or aggregated group:
      - Sum the 'A' values for that CPC or group.

	CPC1					CPC2					CPC3			AWLC SUM
	LP1	LP2	LP3	LP4	AWLC SUM	LP1	LP2	LP3	AWLC SUM		LP1	LP2	AWLC SUM	
0:00	55	232	13	563	863	217	101	392	710	0:00	148	183	331	
1:00	64	481	49	246	840	276	392	384	1052	1:00	71	62	133	
2:00	60	454	15	255	784	235	382	65	682	2:00	179	288	467	
3:00	73	279	38	342	732	166	269	202	637	3:00	348	321	669	
4:00	75	257	37	671	1040	108	218	347	673	4:00	260	115	375	
5:00	52	442	32	329	855	369	86	122	577	5:00	450	123	573	
6:00	61	415	17	172	665	315	342	123	780	6:00	241	74	315	
7:00	75	406	12	168	661	366	293	155	814	7:00	148	340	488	
8:00	66	465	12	159	702	117	64	100	281	8:00	103	363	466	
9:00	68	374	18	390	850	154	264	347	765	9:00	446	155	601	
10:00	63	350	50	571	1034	266	83	220	569	10:00	229	399	628	
11:00	66	395	22	382	865	339	120	336	795	11:00	244	373	617	
12:00	52	459	24	263	798	342	247	318	907	12:00	304	211	515	
Peak					1040	+			1052	+			669	2761

The Basics – How Your IBM Software is Calculated

- Then use the software pricing curve to calculate the price for each product.
  - For aggregated CPCs, one price is calculated for each *set* of CPCs.
  - For non-aggregated CPCs, one price is calculated for each CPC.

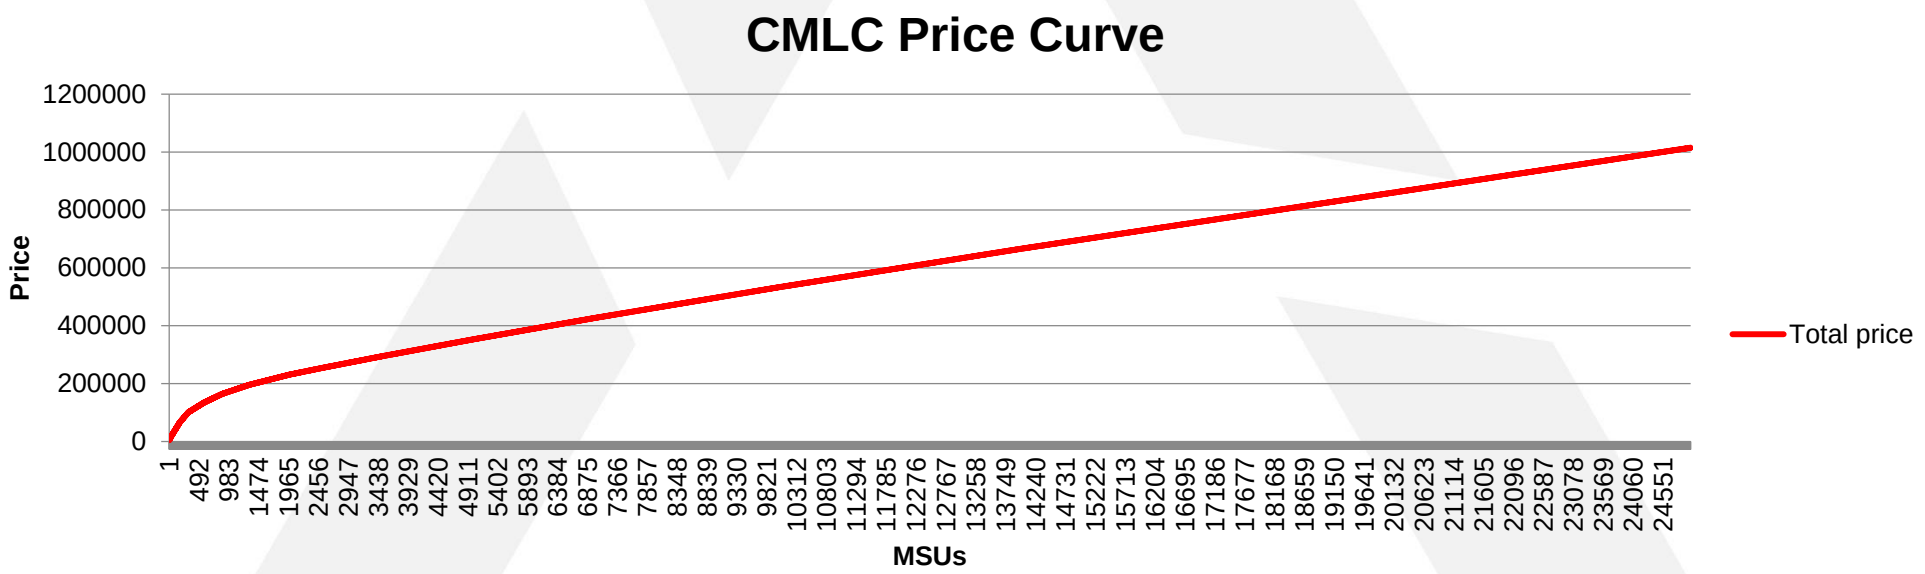

- Because the price for additional MSUs decreases with the total number of consumed MSUs in that CPC or aggregate, it is financially attractive to have all CPCs in *some* aggregate, and to have as few aggregates as possible.

Software Pricing Curve

NO ONE wants to be at this end of the pricing curve

Incremental Price per MSU

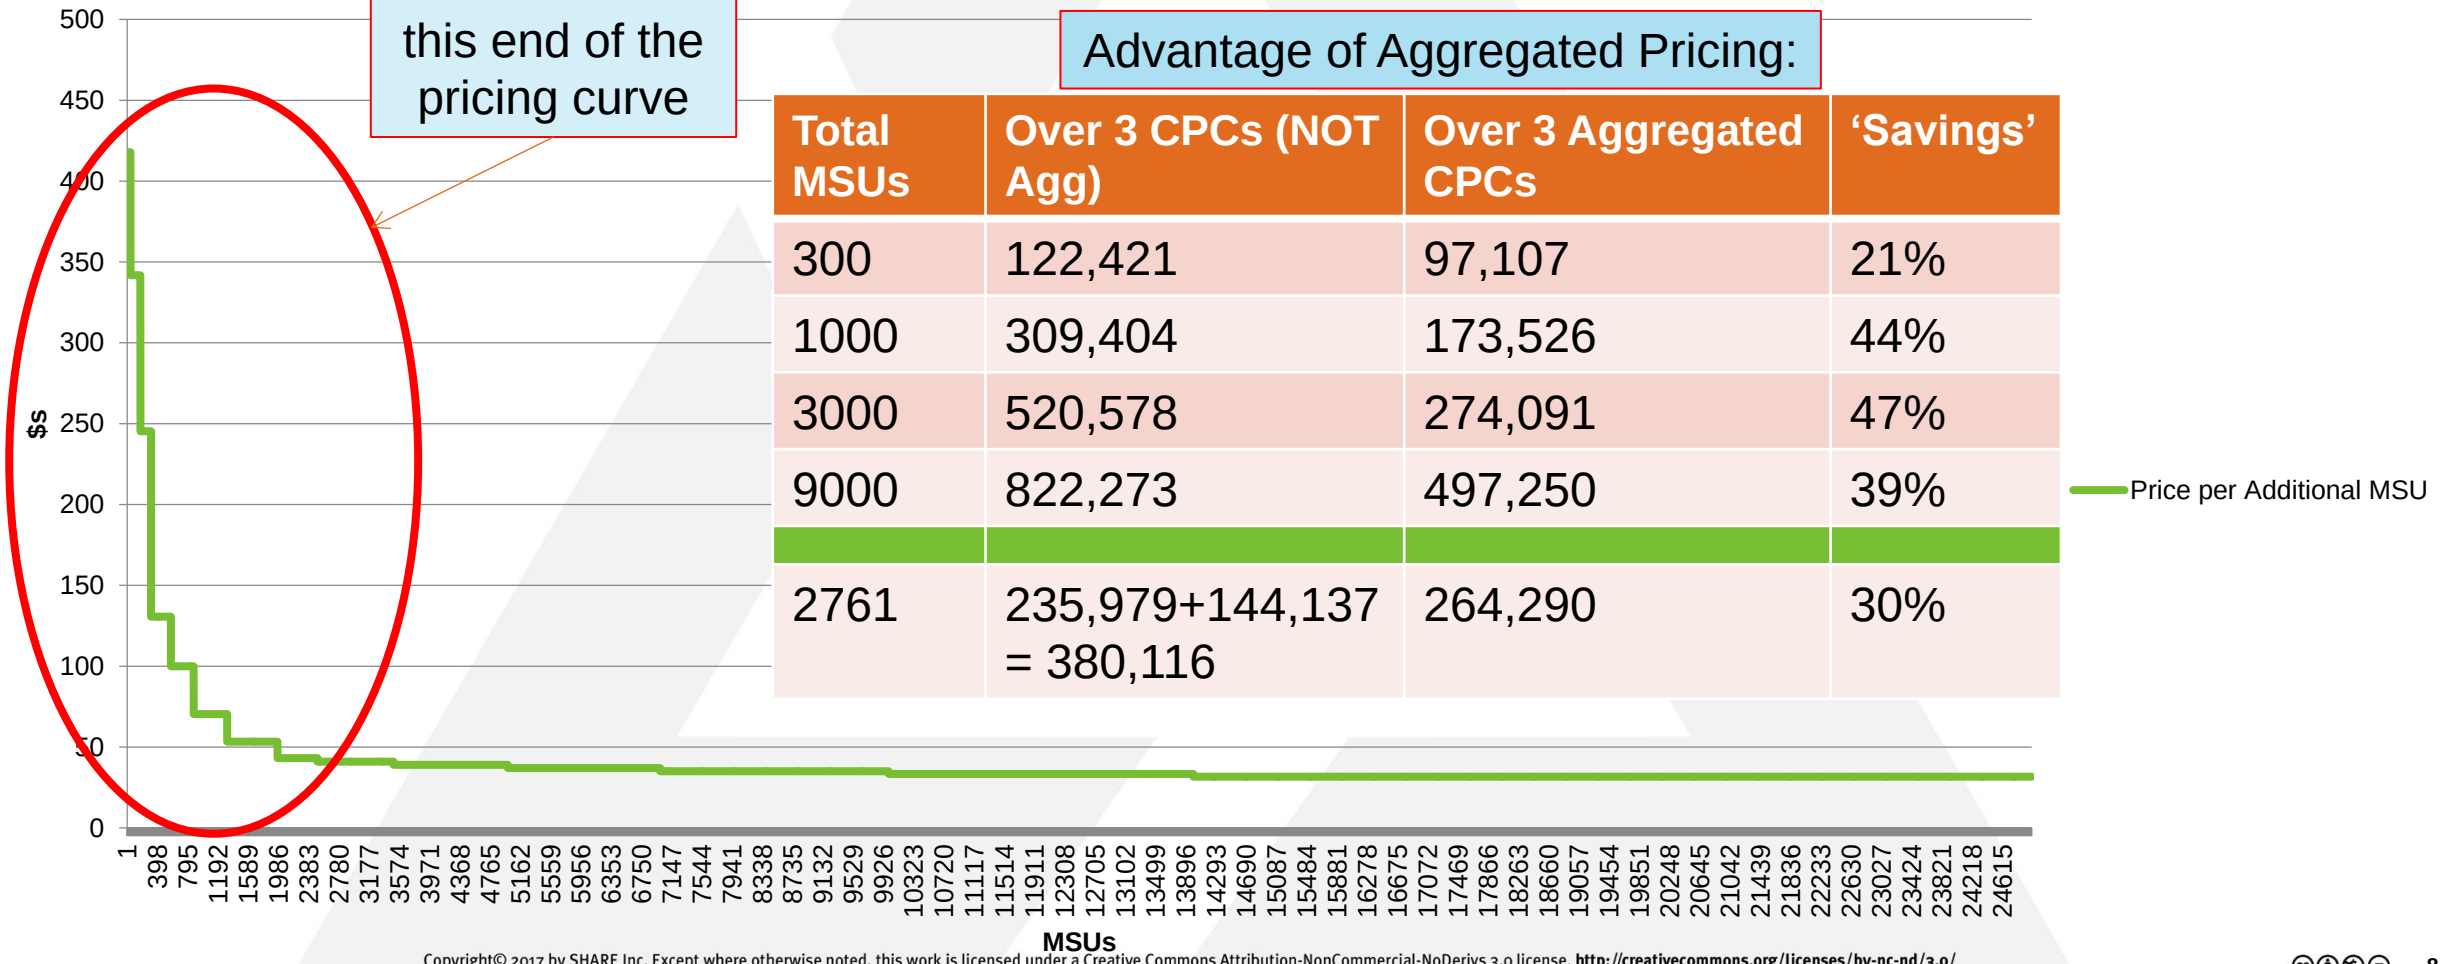

Advantage of Aggregated Pricing:

Total MSUs	Over 3 CPCs (NOT Agg)	Over 3 Aggregated CPCs	'Savings'
300	122,421	97,107	21%
1000	309,404	173,526	44%
3000	520,578	274,091	47%
9000	822,273	497,250	39%
2761	235,979+144,137 = 380,116	264,290	30%

Software Pricing Curve (An Aside)

An important aspect of the price curve is how it interacts with growth or shrinkage. In this example, a 10% increase in MSUs would result in a cost increase of just 6.57%. On the flip side, for the same reason, DEcreasing your MSU usage by 10% will NOT decrease your costs by 10%.

MSU Summary for PricePlex ===>	PRODPLEX		Recent Mnths	MLC Average	+ Growth %	MSU Incr	Total MSUs	New Disc %	MLC Incr	MLC Total	MLC Incr %
MLC Products	Product ID	Lic Type	Average MSUs	$ 813,617.06	10.0%	MSU Incr	Total MSUs	8.4%	$ 53,470.45	$ 867,087.51	6.57%
z/OS V2 (Traditional)	5650-ZOS	AWLC	2687	$ 239,680.36	10.0%	269	2956	8.4%	$ 10,634.80	$ 250,315.16	4.44%
z/OS V2 DFSMS dsshsm	5650-ZOS	AWLC	2687	$ 21,564.19	10.0%	269	2956	8.4%	$ 1,118.67	$ 22,682.86	5.19%
z/OS V2 DFSMS rmm	5650-ZOS	AWLC	2687	$ 10,554.79	10.0%	269	2956	8.4%	$ 559.34	$ 11,114.13	5.30%
z/OS V2 DFSORT	5650-ZOS	AWLC	2687	$ 3,543.34	10.0%	269	2956	8.4%	$ 278.44	$ 3,821.78	7.86%
z/OS V2 SDSF	5650-ZOS	AWLC	2687	$ 6,261.50	10.0%	269	2956	8.4%	$ 322.79	$ 6,584.29	5.16%
z/OS V2 C/C++ without Debug	5650-ZOS	AWLC	755	$ 6,814.90	10.0%	76	831	8.4%	$ 395.42	$ 7,210.32	5.80%
z/OS V2 Infoprint Server	5650-ZOS	AWLC	766	$ 5,123.22	10.0%	77	843	8.4%	$ 240.51	$ 5,363.74	4.69%
non-z/OS (AWLC)											
CICS TS for z/OS V5	5655-Y04	AWLC	2661	$ 216,050.38	10.0%	266	2927	8.4%	$ 16,719.67	$ 232,770.05	7.74%
DB2 11 for z/OS	5615-DB2	AWLC	2661	$ 190,351.43	10.0%	266	2927	8.4%	$ 14,802.10	$ 205,153.53	7.78%
IBM MQ for z/OS V8	5655-W97	AWLC	2429	$ 91,950.73	10.0%	243	2672	8.4%	$ 7,160.66	$ 99,111.38	7.79%
Tivoli NetView for z/OS V6	5697-NV6	AWLC	2687	$ 12,415.92	10.0%	269	2956	8.4%	$ 970.22	$ 13,386.13	7.81%
IBM Enterprise Cobol for z/OS V4	5655-S71	AWLC	203	$ 5,765.22	10.0%	20	223	8.4%	$ 267.84	$ 6,033.06	4.65%
FWLC (Workload License Charge)											
ACF/SSP Version 4 MVS	5655-041	FWLC	1	$ 1,891.72			1			$ 1,891.72	
IBM Library for REXX/370	5695-014	FWLC	1	$ 1,286.02			1			$ 1,286.02	
Transforms to AFP	5655-N60	FWLC	1	$ 363.33			1			$ 363.33	

The Problem – Aggregation Savings

- You can see that being able to avail of aggregation can result in a significantly reduced MLC software bill.
- Parallel Sysplex Aggregation lets you aggregate multiple CPCs IF they meet a list of requirements.
  - This is GOOD, because it reduces your costs.
  - This is BAD, because the potential savings can be so large that they incent companies to create less-than-ideal configurations, such as placing test, development, and production systems in the same sysplex, or co-mingling completely unrelated systems in the same sysplex, mixing production and test LPARs on the same CPC, and so on.
    - What would *your* bean-counters go for? A guaranteed savings now, or a reduced chance of a sysplex-outage later?
  - It is also common to see people burning valuable system programmer time just to ensure that the aggregation criteria continue to be met. This delivers *zero* technical or business advantage, and is a pointless waste of rare skills.

The Problem – Moving Workloads

Parallel Sysplex enables Single System Image capability. Single interface to users and customers, front ending multiple transaction managers, database managers, systems, CPCs, and even sites. In a "PlatinumPlex", any work can run on any system, delivering optimal availability, performance, and resource utilization.

Parallel Sysplex was only *needed* by a relatively small number of customers when it was first introduced. Today, it is a must for any company wishing to compete and offer a discernibly different level of customer service.

In addition to the availability benefits of Parallel Sysplex, the automatic workload routing and shared message queue capabilities can send work to whichever system is most likely to be able to meet the performance objectives of that work. This protects service levels and optimizes resource utilization.

The Problem – Moving Workloads

However, while sysplex provides the technical ability to float and shift workloads between systems, CPCs, and even sites, the use of the peak R4HA for each CPC to determine your software bill can act as a financial disincentive to exploit this capability.

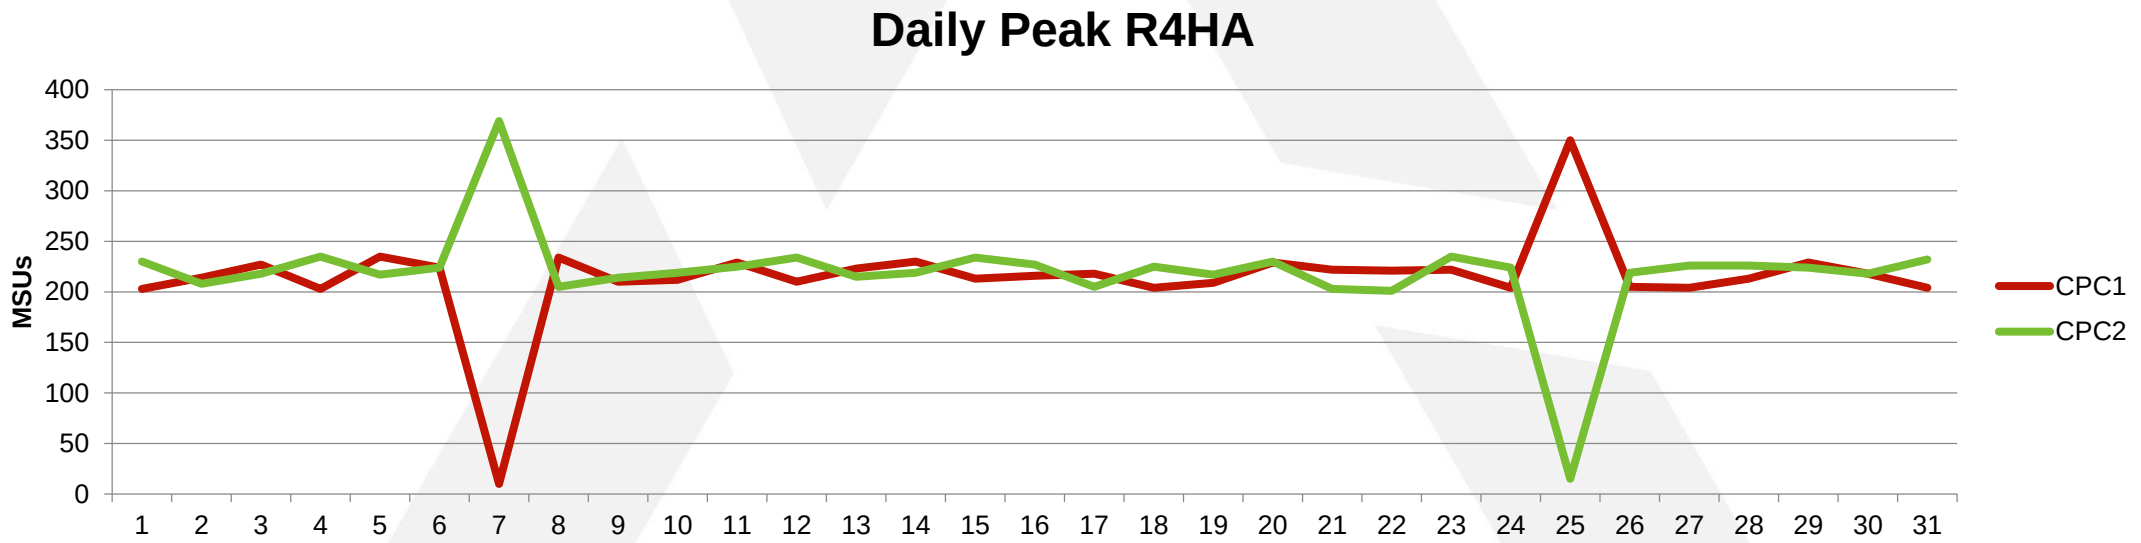

- In this example, the combined CPC consumption never exceeded about 460 MSUs. However, the bill for this month would be for 710 MSUs (peak of 350 MSUs for CPC1 plus 360 MSUs for CPC2).

When doing a permanent move of a sizeable workload between CPCs, you would need to do it at 23:59 on the 1st of the month to avoid having to pay for that workload twice in that month.

The Problem – Meeting Aggregation Criteria

- The Parallel Sysplex Aggregation criteria are:
  - The qualifying LPARs must all be in the same sysplex.
  - There must be at least one Coupling Facility that *all* LPARs in the PrimaryPlex are connected to.
  - There must be at least one (non-XCF) structure that is used by *all* LPARs in the PrimaryPlex.
  - LPARs belonging to the PrimaryPlex must exist across two or more machines.
  - LPARs belonging to the PrimaryPlex must:
    - Generate 50% or more of the MVS-based workload (both zIIPs and General Purpose CPs) on each CPC where the PrimaryPlex exists.
    - For the 8-hour prime shift.
    - For 5 consecutive days.

- How do you have an environment with multiple sysplexes (sandbox, test, development, acceptance, and production) and still meet all these requirements?

The Problem – Distance Limitations

- For resiliency (or for business or historical or skills purposes), you might want your production and backup data centers to be as far apart as possible.
- For optimum performance and minimum overhead, you want to limit a multi-site sysplex to span not more than 10km.
- To meet the aggregation requirements, all aggregated CPCs *must* be in the same sysplex (which imposes a maximum distance of 100 km).
  - Putting them close together is good for performance, but increases the risk of a single event impacting both locations.
  - Putting them further apart improves resiliency, but impacts performance and system overhead.
  - If <100km not possible, both sites cannot be in the same aggregate. This can result in unnatural configurations, purely because of cost considerations.

The Problem – SW Contract Terms

While not directly related to sysplex aggregation, another annoying restriction was the way software was billed during migrations:

- Single Version Charging (SVC) provided a fixed period (typically 12 months) to compete the migration from one version of a product to another. But if you did not complete the migration in the agreed period, you would be billed for both versions, for the MSUs used by that version. Some customers struggled to complete the migration in the available time, especially bearing in mind business schedules.
- Customers only had 6 months (‘grace period’) to complete the migration to new versions of IPLA products.
- You had 90 days from the Code 20 date to compete your migration from an old CPC to a replacement one without incurring additional SW costs.

The Solution – CMP

- Those were some of the challenges faced by z/OS customers.

- On 28 July, 2015, IBM announced Country Multiplex Pricing (CMP) – a new software pricing option intended to address these challenges and remove artificial constraints to customers growing their z/OS environment intelligently.
  - IBM US Software announcement 215-230.

- CMP effectively behaves as if all z/OS LPARs in a country are in a single CPC when calculating your peak R4HA.
- Let's have a look at the changes CMP introduced….

CMP – Shamplex Elimination

- When a customer moves to CMP, ALL of their z/OS CPCs in that country effectively become part of a single aggregation group.
  - There is no longer a requirement for a certain set of LPARs to account for more than xx% of the used capacity.
  - This eliminates the incentive to create OR MAINTAIN sysplexes that contain systems that do not logically belong in the same sysplex.
  - Prior to CMP, disaggregating the PrimaryPlex would likely result in increased software bills. ***After*** migrating to CMP, breaking up the PrimaryPlex should have no impact on the software bill, all other things being equal.

CMP – Freely Move Workloads Between CPCs

- When using CMP, your peak R4HA is calculated by summing the MSUs for LPARs across ALL CPCs, *not* on a CPC-by-CPC basis.
- This means that moving a workload from one CPC to another should have zero impact on your software bill – note that CMP line below is no higher on days 7 or 25, so the bill would be for 460 MSUs rather than 710 MSUs if not using CMP.
  - This means that you can now move workloads when it suits *your business*, not when it fits in with how your R4HA is calculated.
- You can also better exploit queue sharing and dynamic workload routing.

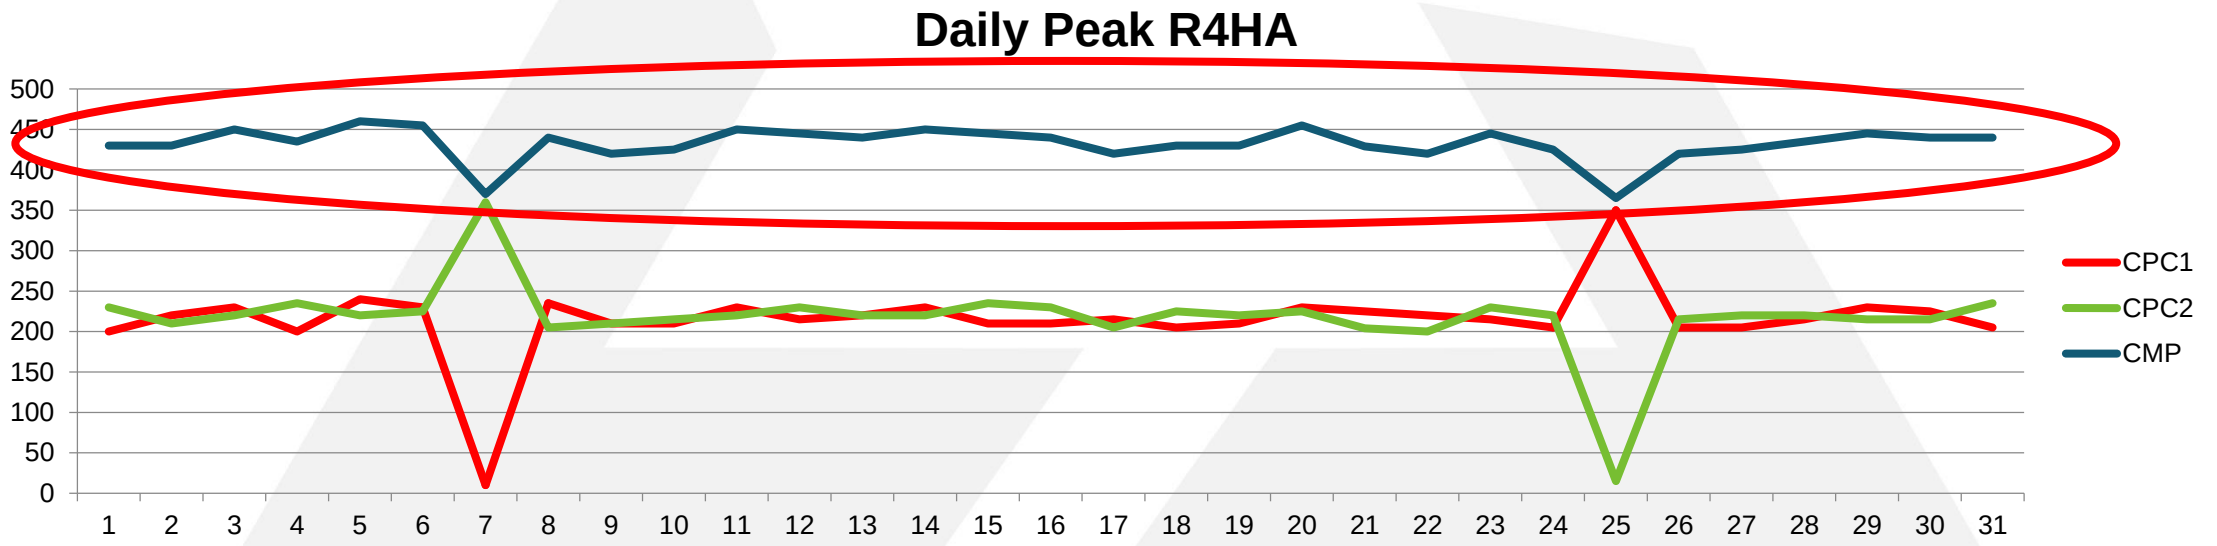

CMP – Eliminate Sysplex Agg Qualification Criteria

- Because all CPCs owned by the enterprise in that country are automatically included in the one aggregation group **AFTER** you move to CMP, post-CMP sysplex aggregation becomes irrelevant.
- All of the effort to ensure that you continue to meet the sysplex aggregation criteria can be eliminated as soon as you move to CMP.
- You can choose to maintain your current sysplex topology, or to adjust it, whichever makes the most sense from a technical or business perspective. Whichever path you take should have nearly no impact on your software bills.
  - **But your *pre*-CMP sysplex aggregation status is still *very* important, as we shall see. Don't even consider disaggregating until after you complete your move to CMP.**

CMP – Eliminate Distance Limitations

- Because all CPCs owned by the enterprise in that country are automatically included in the one aggregation group AFTER you move to CMP, post-CMP sysplex aggregation becomes irrelevant.
  - If sysplex aggregation becomes irrelevant, then the need to be in a sysplex in order to qualify for aggregation also becomes irrelevant.
- This should *not* affect the need for sysplex – if anything, sysplex should deliver even more value because sysplexes should be more homogenous, and therefore deliver better availability.
- It also doesn't eliminate the desirability of multi-site sysplexes – Parallel Sysplex and HyperSwap still enable transparent site swaps.
- But it SHOULD eliminate nearly every reason for creating very long distance sysplexes.

CMP – Migration Restrictions

- Part of the announcement of CMP was Multiplex Version Measurement (MVM). This was a replacement for Single Version Charging, Migration Pricing Option (MPO), and the IPLA 6-month grace period limitation.
  - It eliminated the 12-month SVC limit, the 24-month MPO limit, and the 6-month IPLA migration period limit.
  - Lets you keep multiple versions for as long as you like[1], but pay as if all the MSUs were used by the later version – this avoids having two products on the steeper part of the pricing curve.
- When CMP was announced, MVM was only available to customers that signed up for CMP. On 14 February 2017, IBM announced Multi-Version Measurement.
  - This is effectively just a renaming of Multiplex Version Management.
  - BUT, MVM is now available to all customers, regardless of whether they use CMP or not.

[1] MVM does not change existing product service practices – at some point, products will still run out of support.

CMP – Reduced Prices for Very Large Sites

- CMP adds more tiers. PSLC had 5 tiers (top was 316+ MSUs), AWLC has 9 tiers (top was 1976+ MSUs), CMLC has 17 tiers (top is 25,000+ MSUs).
  - CMLC and AWLC $/MSU identical up to 2499 MSUs, then CMLC continues to drop.

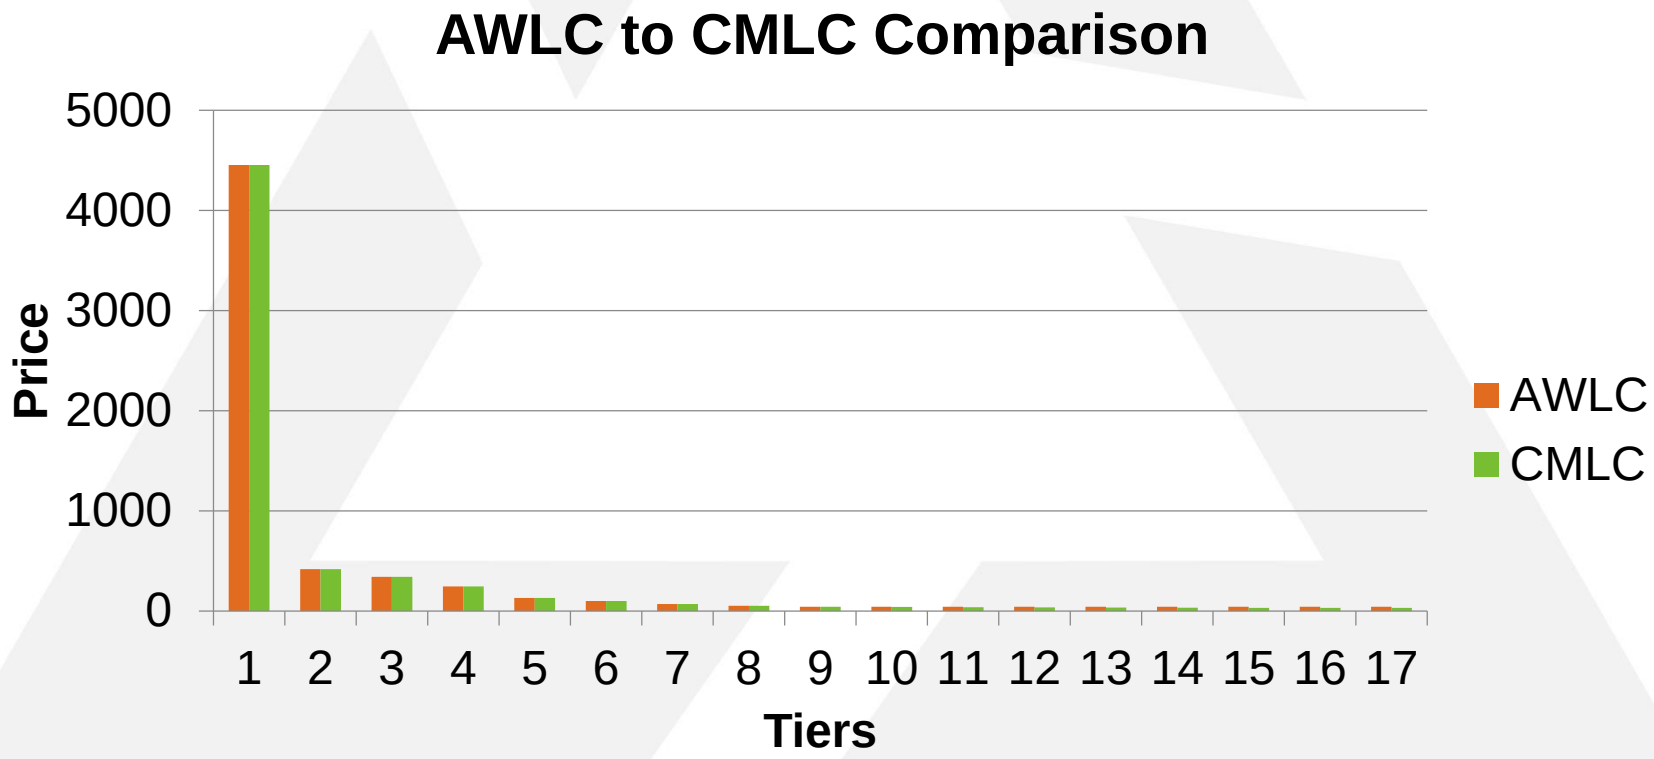

CMP – Different R4HA Calculation Method

- 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

	CPC1					CPC2					CPC3				AWLC SUM	CMLC SUM
	LP1	LP2	LP3	LP4	AWLC SUM		LP1	LP2	LP3	AWLC SUM		LP1	LP2	AWLC SUM		
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

CMP – Move Capacity Between CPCs

- Many sites use one or more forms of soft capping to control software costs.
- When you add a dynamic capping product, you can balance the need to control costs with the requirement to meet Service Level Agreements.
  - Automatically adjust caps so that important workloads in need of help get more capacity, but costs are protected by taking an equivalent amount of capacity away from 'less deserving' LPARs *on the same CPC*.

- With CMP, you could potentially achieve the same effect at the cross-CPC level. Because MSUs are summed across all CPCs, you could INcrease the Defined Capacity of an LPAR on CPC B, DEcrease the Defined Capacity of an LPAR on CPC A by an equivalent amount, and there should be no impact on your software bill because the total available MSUs across the multiplex hasn't changed.

CMP – What's The Catch?

- Does all this seem too good to be true?
- Well, it's all true, but what have we *not* spoken about so far? The cost…
- CMP is intended to remove all the constraints associated with the old sysplex aggregation and reduce the cost *of future growth*.
  - You would correctly assume that the new method of calculating the Peak R4HA MSU amount would normally result in a smaller number than the previous method.
  - And because CMLC $ per MSU are equal to or lower than AWLC, you would expect that this will result in lower bills for you.
  - However, IBM applies an 'uplift' that ensures that at the point you switch to CMP, the CMLC price would be identical to the AWLC price (all other things being equal). IBM is very clear that CMP is **NOT** intended to reduce your current bills. It is intended to reduce the cost of *growth*.
- Let's have a look at how this works (**you will be paying this uplift for the next x YEARS, so pay attention!**).

CMP – Financials

- Prior to moving to CMP, IBM calculates two baselines for each product:
  - Using your SCRT reports, IBM calculates what the Peak R4HA MSUs for the 3 months prior to switching to CMP would have been if the CMP rules were applied. The average of those 3 months is called the **MSU Base**.
    - This value probably will be different to the values that were used to calculate your bill for those 3 months, but it is consistent with how your bill will be calculated after you move to CMP.
  - The other baseline is the average of the *billed* amount ($s) for each of the prior 3 months – this is called the **MLC Base**.
- The % difference between the MLC base and what the price would have been based on the CMP rules and CMLC tiers is calculated – this is called the **MLC Base Factor** (but most people call it the 'uplift').

CMP – Financials

- How does all this get factored into working out your price?
  1. Calculate the list price of the total reported Multiplex peak MSUs using the CMLC metric.
  2. Calculate the current list price of the **Base MSUs** using the CMLC metric.
  3. Multiply the result of #2 by the **MLC Base Factor**.
  4. Add the values in #1 (actual usage) and #3 (your uplift) to determine the CMLC price for that period.

- This calculation takes place EVERY MONTH until IBM changes the rules, or you move to some other pricing metric.

- Let's look at an example (based on an example by IBM's David Chase). For the sake of 'simplicity', I have ignored any TTO or other discounts.

CMP – Financials

MLC Base Factor = 71,888/307,334
= 23.391%

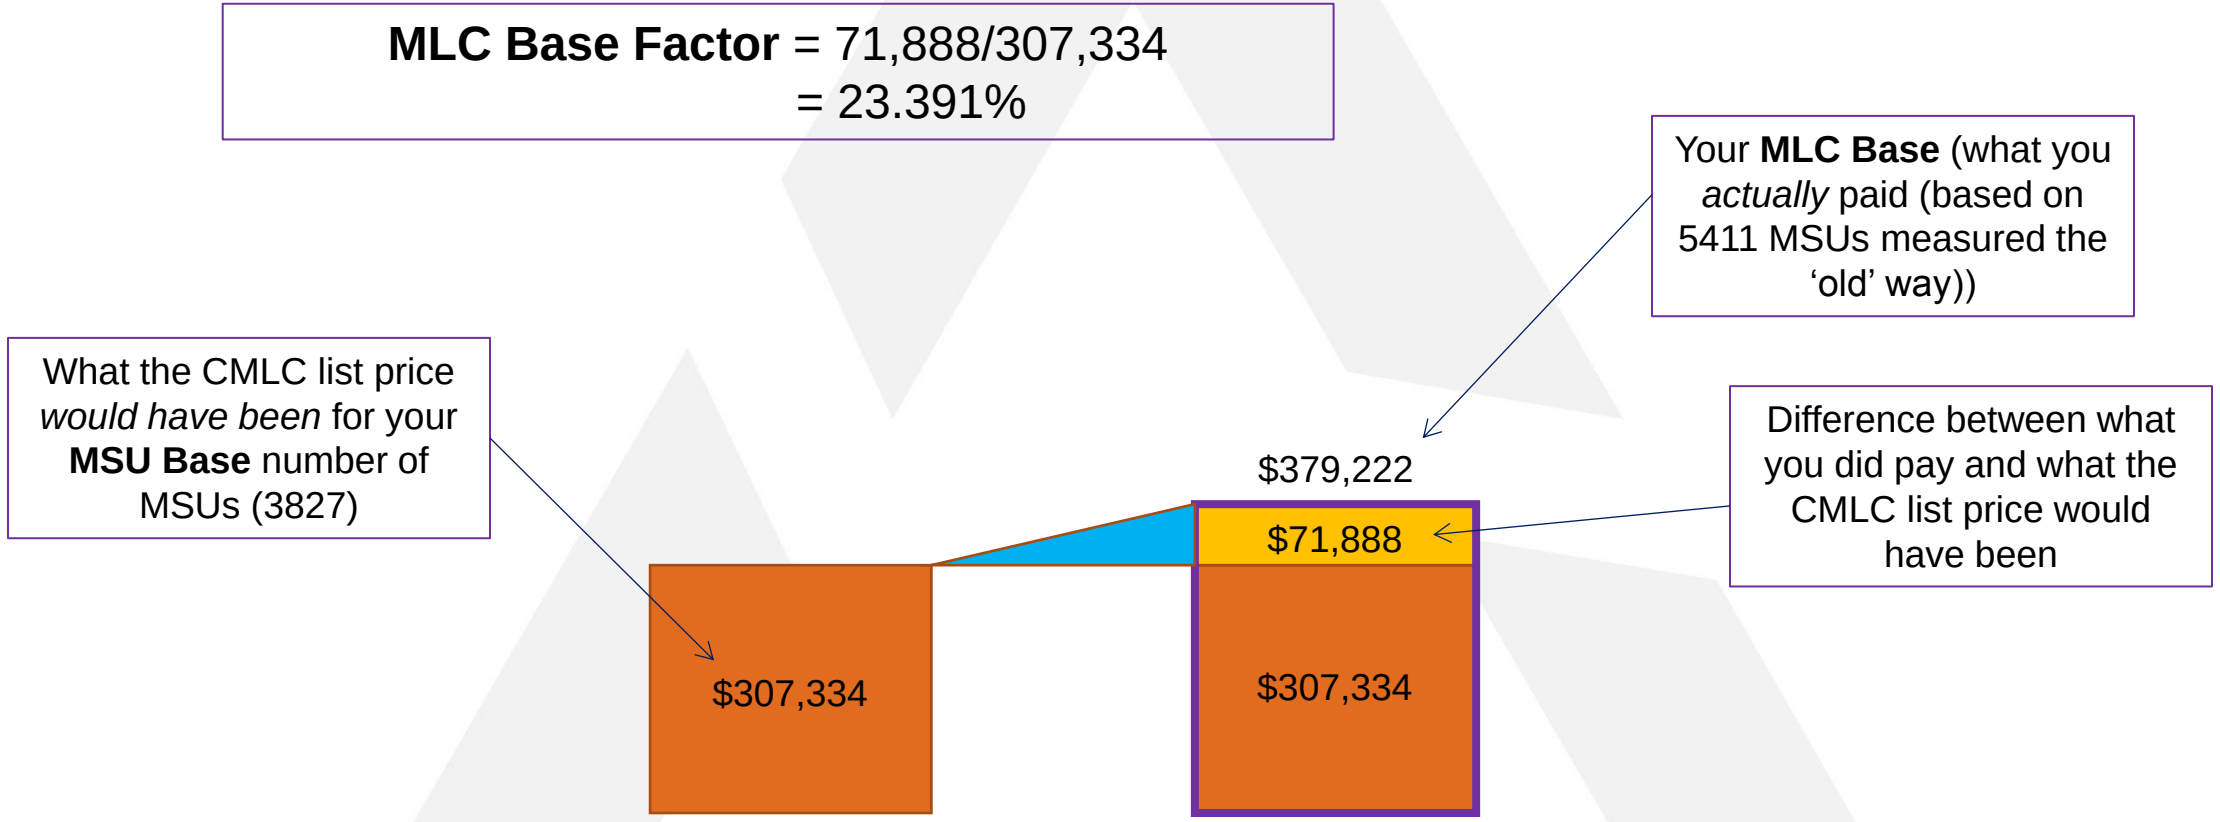

- How will all this affect my bill when I move to CMP?

CMP – Growth Scenario

4 Add uplift to current usage CMLC price

1 Price current usage

2 Get current CMLC price for MSU Base

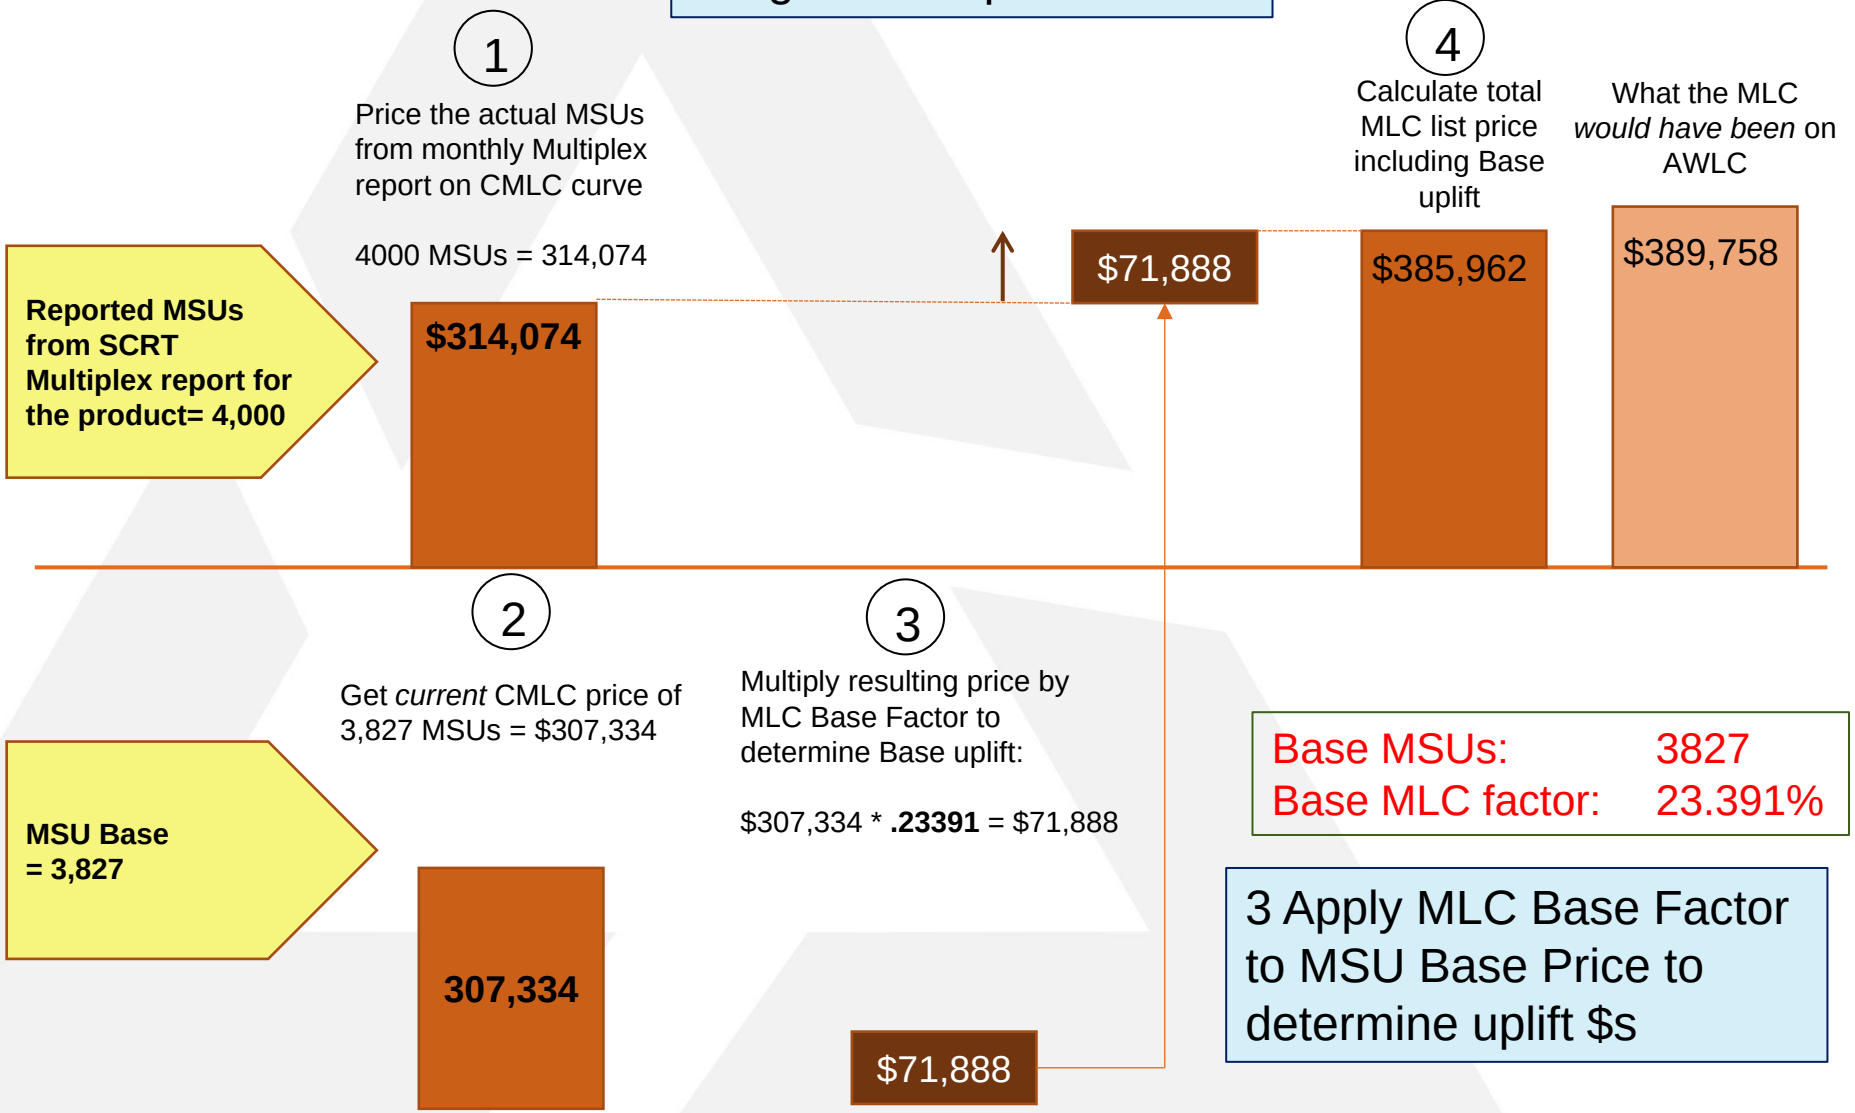

Base MSUs: 3827
Base MLC factor: 23.391%

3 Apply MLC Base Factor to MSU Base Price to determine uplift $s

CMP – Shrinking Scenario

4 Add uplift to current usage CMLC price

1 Price current usage

2 Get current CMLC price for MSU Base

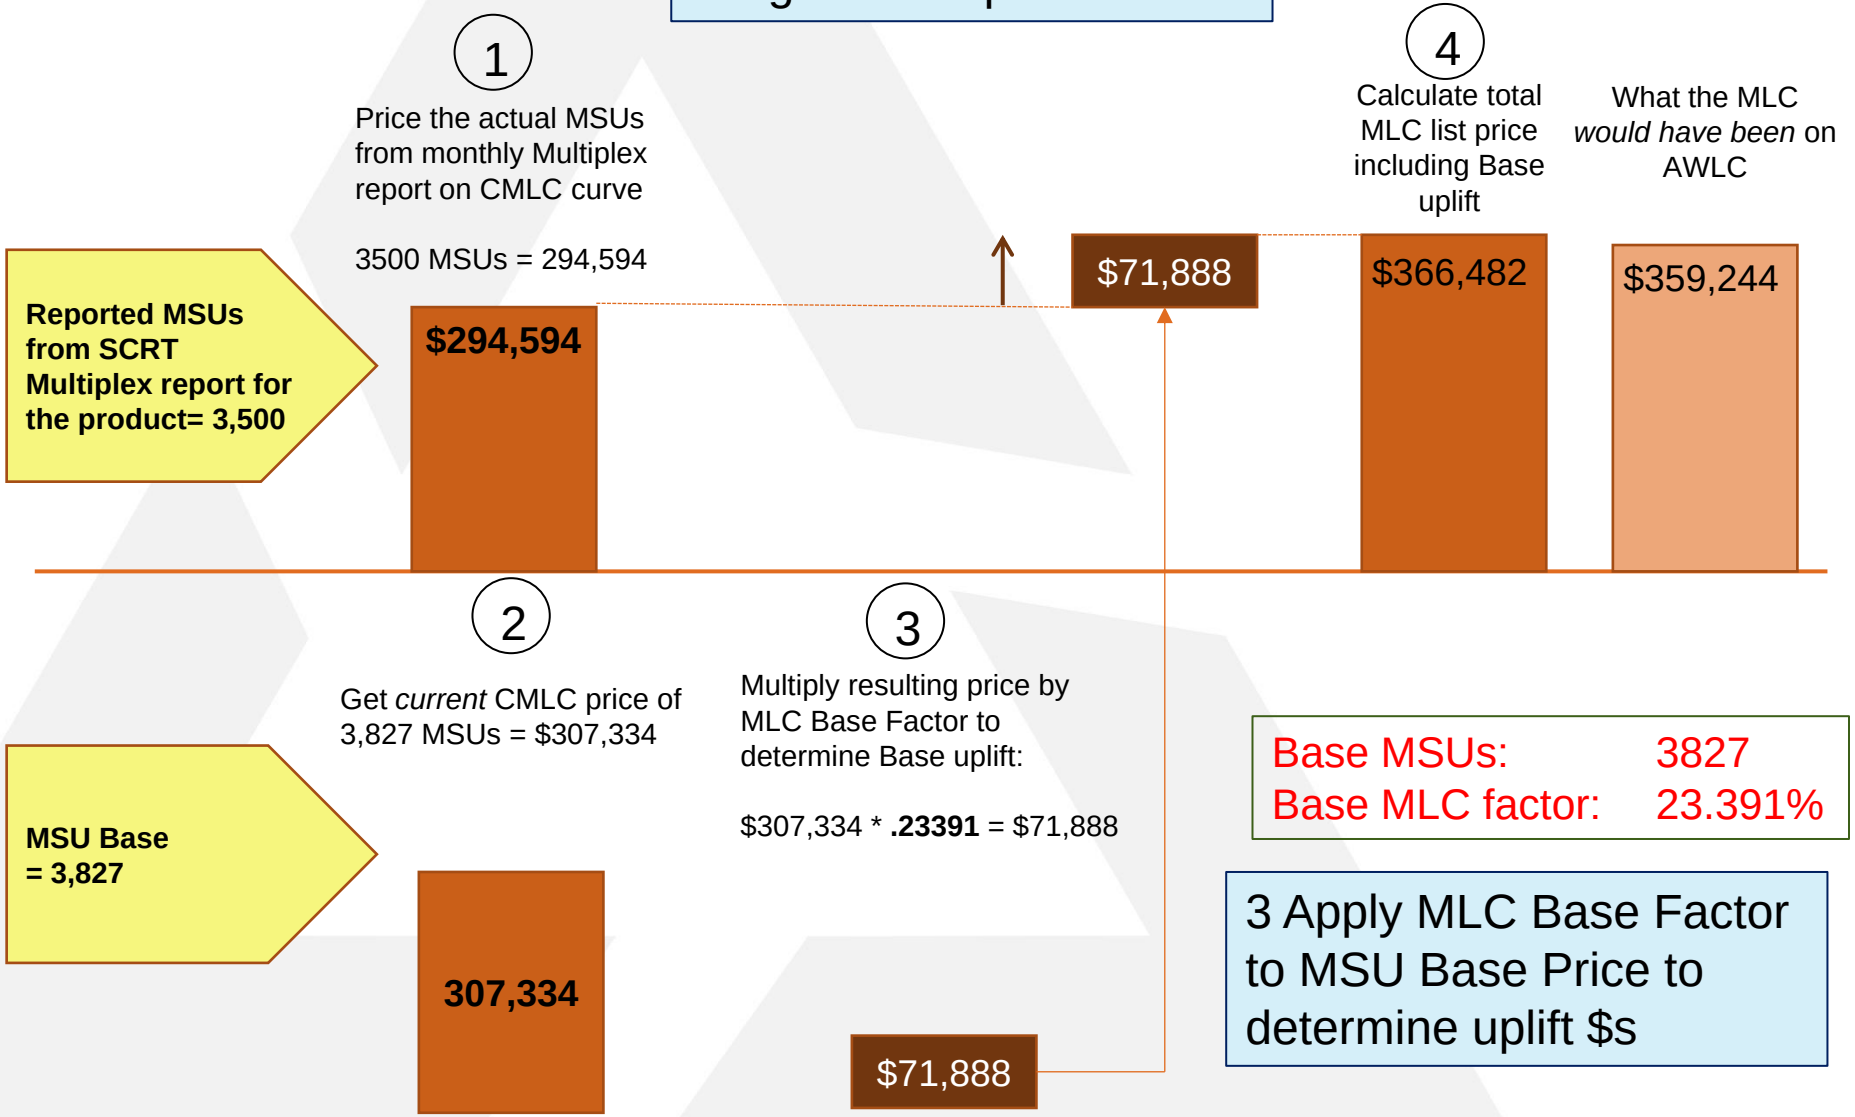

Base MSUs: 3827
Base MLC factor: 23.391%

3 Apply MLC Base Factor to MSU Base Price to determine uplift $s

CMP – Financials

- How does all this make growth cheaper?
  - The uplift is a fixed percent, based partially on the Base MSU value. As your actual MSU consumption grows, the uplift does *not* grow. You could view it as like paying the pre-CMP price for the pre-CMP capacity you were using, and the (*possibly* lower) CMLC price for the additional capacity you are using.
    - The CMLC price per MSU continues to decrease after 2500 MSUs – with AWLC, the decreases stop at 2500 MSUs.
    - Growing customers with a Base MSU of just under 2500 will get the best value from moving to CMP.
  - The Peak R4HA used as the basis for your bill is calculated across ALL CPCs, so very likely is lower than the Peak R4HA would have been using the pre-CMP methodology. This means that as you grow, the difference between peak R4HA based on the CMP calculation and the peak R4HA based on the pre-CMP calculation is likely to increase. In other words, your growth is resulting in a smaller R4HA increase than would have been the case, so you are paying for fewer additional MSUs for the same amount of growth.

CMP – Financials

- What hurts you on an ongoing basis is that uplift factor. So let's look at that a little more closely…. It is based on:
  - **MLC Base –** that is, how much you paid prior to moving to CMP. That amount is based on the sum of the peak R4HA for each CPC or aggregate group.
  - **MSU Base –** Your simultaneous peak R4HA across all CPCs (peak of the sums).

- If you want to minimize the uplift, you need to understand how these numbers affect the MLC Base Factor

- Let's look at some grossly over-simplified examples to illustrate different scenarios…

CMP – Financials – Example 1

- Our base case - 3 CPCs, 1 aggregation group, utilization of each CPC moves independently of the other CPCs, resulting in a large difference in the R4HA as measured by CMP and pre-CMP.

Interval	CPC1	CPC2	CPC3	Peak AWLC	Peak CMLC
1	500	250	100		850
2	250	100	500		850
3	100	500	250		850
Peak R4HA	500	500	500	1500	850

  - AWLC = $205,579
  - CMLC = $162,276
  - $ Difference = $43,353
  - MSU Difference = 650
  - MLC Base Factor = 43,343/162,276 = 26.715%

CMP – Financials – Example 2

- What if we try to minimize MSUs during 3 month baseline period? Same config as Example 1, but softcap all CPCs to 450 MSUs.

Interval	CPC1	CPC2	CPC3	Peak AWLC	Peak CMLC
1	450	250	100		800
2	250	100	450		800
3	100	450	250		800
Peak R4HA	450	450	450	1350	800

- AWLC = $197,573
- CMLC = $157,229
- $ Difference = $40,344
- MSU Difference = 550
- MLC Base Factor = 40,344/157,229 = 25.659%

A little bit better (Example 1 was $43,353)

CMP – Financials – Example 3

- Rather than trying to minimize the peak R4HA, how about minimizing the *difference* between the CMP and pre-CMP R4HA values, by distributing work more evenly across the CPCs, for example?

Interval	CPC1	CPC2	CPC3	Peak AWLC	Peak CMLC
1	400	250	200		850
2	300	400	170		870
3	200	250	400		850
Peak R4HA	400	400	400	1200	870

- AWLC = $187,608
- CMLC = $164,225
- $ Difference = $23,383
- MSU Difference = 330
- MLC Base Factor = 23,383/164,225 = 14.238%

That seemed to be more effective

CMP – Financials – Example 4

- Let's try bringing the utilizations still closer, to prove that wasn't an anomaly

Interval	CPC1	CPC2	CPC3	Peak AWLC	Peak CMLC
1	300	250	300		850
2	300	400	170		870
3	300	400	300		1000
Peak R4HA	300	400	300	1000	1000

- AWLC = $173,526
- CMLC = $173,526
- $ Difference = $0
- MSU Difference = 0
- MLC Base Factor = 0/173,526 = 0%

Nobody will be *this* perfect. BUT it proves the concept of minimizing the difference between CMP R4HA and non-CMP R4HA is worth pursuing.

CMP – Financials – Example 5

- Let's see how much difference it would make if all CPCs were *not* in the same aggregation group prior to moving to CMP.

Interval	CPC1	CPC2	CPC3	AWLC	CMLC
1	500	250	100		850
2	250	100	500		850
3	100	500	250		850
Peak R4HA	500	500	500	1500	850

- AWLC = $298,474 (173,526 + 124,948)
- CMLC = $162,276
- $ Difference = $136,198
- MSU Difference = 650
- MLC Base Factor = 136,198 / 162,276 = **83.929**%

Just to be clear – this is what you are paying today. CMP does not increase your costs. But it will not provide sysplex aggregation prices if you were not aggregated when you moved to CMP.

CMP – Financials – A Little More Detail

- Notice that uplift varies by product. Products purchased after migrating to CMP (e.g. DFSrmm) and FWLC products have NO uplift.

Month Jan 2017			CPC01	CPC02	CPC03	Totals	MLC Disc%	CMP Adjust	MLC Total	MLC Perc
MLC Products	Product ID	Lic Type	MSUs	MSUs	MSUs	Total MSUs	8.4%	$ 26,397.97	$ 863,315.05	100.0%
z/OS V2 (Traditional)	5650-ZOS	CMLC	894	665	1249	2808	8.4%	$ 6,239.99	$ 250,095.49	29.0%
z/OS V2 DFSMS dsshsm	5650-ZOS	CMLC	894	665	1249	2808	8.4%	$ 596.78	$ 22,601.89	2.6%
z/OS V2 DFSMS rmm	5650-ZOS	CMLC	894	665	1249	2808	8.4%		$ 10,772.43	1.2%
z/OS V2 DFSORT	5650-ZOS	CMLC	894	665	1249	2808	8.4%	$ 106.62	$ 3,761.05	0.4%
z/OS V2 SDSF	5650-ZOS	CMLC	894	665	1249	2808	8.4%	$ 204.68	$ 6,594.40	0.8%
z/OS V2 C/C++ without Debug	5650-ZOS	CMLC	894			894	8.4%	$ 95.44	$ 7,613.70	0.9%
z/OS V2 Infoprint Server	5650-ZOS	CMLC		665		665	8.4%	$ 219.12	$ 5,026.87	0.6%
non-z/OS (CMLC)										
CICS TS for z/OS V5	5655-Y04	CMLC	869	665	1249	2783	8.4%	$ 9,033.09	$ 231,859.60	26.9%
DB2 11 for z/OS	5615-DB2	CMLC	869	665	1249	2783	8.4%	$ 6,932.87	$ 203,282.39	23.5%
IBM MQ for z/OS V8	5655-W97	CMLC	869	665	1051	2585	8.4%	$ 2,163.96	$ 98,584.82	11.4%
Tivoli NetView for z/OS V6	5697-NV6	CMLC	894	665	1249	2808	8.4%	$ 589.96	$ 13,386.39	1.6%
IBM Enterprise Cobol for z/OS V4	5655-S71	CMLC	219			219	8.4%	$ 215.46	$ 6,194.95	0.7%
FWLC (Workload License Charge)										
ACF/SSP Version 4 MVS	5655-041	FWLC	1			1			$ 1,891.72	0.2%
IBM Library for REXX/370	5695-014	FWLC	1			1			$ 1,286.02	0.1%
Transforms to AFP	5655-N60	FWLC	1			1			$ 363.33	0.0%

CMP – Positively Influencing the Uplift Percent

- So, what have we learned?
  - Having non-aggregated systems *prior* to moving to CMP <u>really hurts</u>!
    - Although the CMP price will not be higher than the pre-CMP price, this illustrates the importance of optimizing sysplex aggregation prior to moving to CMP.
  - Closing the gap between the CMP-calculated Peak R4HA and the pre-CMP R4HA can potentially deliver larger ongoing savings than trying to minimize the pre-CMP MSUs.
    - *In*creasing utilization in an interval is typically easier/less painful than trying to reduce it.
  - The more complex your configuration, especially if different LPARs have different software stacks, the harder it is to achieve this.
    - But focusing on just your most expensive LPARs and products can still deliver worthwhile savings.
  - If your z/OS environment is shrinking, CMP will likely cost you more than other options. You need to decide if the technical advantages outweigh the financial cost.

CMP – Evaluation and Preparation

- How do you know if CMP is right for you?
  - Try to focus on the technical and business benefits from removal of the sysplex aggregation criteria.
    - MVM used to be another advantage of CMP, but now it is available to everyone.
  - Unless your z/OS workloads are shrinking, in the worst case, CMP should not result in *higher* bills than you are paying today (for the same number of MSUs).
  - If you are growing, CMP should save you some money, but don't expect the savings to be significant. Again – focus on the technical and business benefits .
  - The one area where you might be able to influence the ongoing cost is the **MLC Base Factor** – if you can reduce that, you will save money every month for the coming months and years.

CMP – Evaluation and Preparation

- Before you make any commitments, you should determine the MSU Base, the MLC Base, and the MLC Base Factor *for each product*.
- To do this, you can:
  - Ask IBM to calculate them for you.
  - Use our SCRTPro service to determine them for you.
  - Work it out yourself.

CMP – Evaluation and Preparation

- Regardless of which option you choose, you have to create a multiplex SCRT file for (at least) the last 3 months.
  - If your workloads fluctuate significantly by time of year, consider performing the calculations for every month in the last year.
- To get a multiplex SCRT report:
  - SCRT job must process the SMF records from ALL the systems that will be in the multiplex – in other words, 'every' z/OS in that country. Don't forget that you will need to create NO89 statements that also reflect *all* your systems in that country.
  - Finally, add the following keyword after your //SPECIAL DD * statement:

    `Country_Multiplex_Pricing`
- You will also need your IBM bill for each month that you create an SCRT report for.

CMP – Evaluation and Preparation

- If you want to calculate the key metrics yourself:
  - You can calculate your **MSU Base** for each product by taking the peak R4HA for that product from the multiplex SCRT reports for 3 consecutive months and getting an average of those numbers.
  - You can calculate your **MLC Base** for each product by getting the cost for each product from your IBM bills. Make sure that you have the bills that cover the same set of systems that were input to your SCRT job.
  - Now, the fun part – to calculate the all-important **MLC Base Factor**, you need to calculate what the bill would have been using CMLC List prices. To do that, you need the MSU Base value for each product, and you need access to the IBM price files so you can calculate the CMLC List price.

CMP – Evaluation and Preparation

- With the MLC Base Factor in hand for each of your products (and ideally for a number of 3-month periods), you should determine if it is large enough to put some effort into reducing it.
  - If it is small (say 2%), the effort to reduce it would probably cost more than any savings you might manage.
  - If it is big (50%!), then it probably *is* worth investing some time to determine the best/most achievable way to reduce it.
  - If it is somewhere between these two, you need to decide whether savings are possible, and significant enough to justify the effort.

CMP – Evaluation and Preparation

- Some things you can do to minimize the MLC Base Factor:
  - Try to minimize the number of CPCs that are *not* in a sysplex aggregation group.
    - Can you move systems off those CPCs and onto a CPC in an aggregation group?
  - Try to minimize the number of aggregation groups.
  - Identify products that are used in more than one aggregation group – is it possible to limit them to one group?
  - Try to minimize the difference between the CMP R4HA value and the pre-CMP R4HA value.
    - Exploit dynamic workload routing to get better balance across your CPCs.
    - Intelligent use of dynamic capping might be able to shift workload from larger LPARs to smaller ones, resulting in a more balanced configuration.
    - If you have a complex configuration (many LPARs, many CPCs, different SW stacks), concentrate on the big hitters (most expensive products in the largest LPARs).

CMP – Requirements

- Customers that currently use sysplex aggregation must have submitted a valid Sysplex Verification Package within the last 12 months.
- ALL CPCs owned by that enterprise in that country running z/OS or z/TPF *must* be included in the agreement.
  - CPCs running *only* z/VSE, z/VM, or zLinux are ignored.
- Must be running z/OS V1 or V2.
- CMP requires SCRT 23.10.0 or later (the current level is 24.11.3).
- SCRT 24.2.0 or 24.11.0 are required if you want to use MVM with a sub-capacity product.

CMP – Requirements

- At the time of signing the CMP contract, all CPCs must be no older than the current generation minus 2.
  - For example, if z13 is the latest generation, all CPCs must be z114/z196 or later.
  - There is an exemption for customers that sign an agreement before 12/31/17 – even though n-2 is now zEC12, customers with z114/z196 can switch to CMP.
    - But remember that z114/z196 cannot be in the same sysplex or the same common timing network as a z14.
  - Any machines which were in the N to N-2 range when added to a Multiplex will remain in the Multiplex even when they eventually become “N-3” or older
  - New machines added to the Multiplex must be within the current N to N-2 range
  - Machines older than N-2 may not be added to an existing Multiplex (even if others of the same generation are already there).

References

- CMP home page:
  - https://www-03.ibm.com/systems/z/resources/swprice/cmp.html
- CMP FAQ:
  - https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=ZSQ03088USEN
- MVM info:
  - https://www-03.ibm.com/systems/z/resources/swprice/mvm.html
- Enterprise Executive (2017 No. 3) article:
  - Country Multiplex Pricing: What You Need to Know by **Cheryl Walker** and **Alan Murphy**
- Watson & Walker SCRTPro Service to help clients manage and optimize their z/OS software portfolio
  - http://watsonwalker.com/software/scrtpro/

Related Sessions

- Utilizing the Sub-Capacity Reporting Tool for Independent Software Vendors, by **Andrew Sica**, Monday 15:15-16:15
- What's New With the Sub-Capacity Reporting Tool, by **Andrew Sica**, Wednesday 08:30-09:30
- Have You Tried All the Options to Reduce Your z/OS Software Costs?, by **Al Sherkow**, Monday 13:45-14:45
- z/OS Pricing, SoftCapping, and Capacity Product Shootout, multiple vendors, Tuesday 10:00-11:00
- VSP: Finding MSU Savings by Mining Your Performance Data?, **Peter Enrico**, Tuesday 13:45-14:45
- The (Virtual) Cheryl and Frank zRoadshow, by **Cheryl Watson** and Frank Kyne, Friday 10:00-11:00

Summary

- For nearly every growing z/OS customer, CMP should be a positive move.
  - Eliminates all the sysplex aggregation restrictions and financial enticements to create configurations that make no technical sense.
  - Effectively extends all the *good* aspects of sysplex aggregation to all CPCs in the country.
    - For example, adding xx MSUs of usage for a product should cost the same no matter which CPC it occurs on.
  - Eliminates disincentives to fully exploit sysplex dynamic workload routing functions.
  - Enables full use of installed capacity without the cost inhibitors.
- If your z/OS trajectory is on a *downward* slope, you need to carefully evaluate whether CMP would make sense for you – it is designed for growing sites, not shrinking ones.
- HOWEVER, the CMP uplift will impact your bills until IBM changes the rules or bring out a replacement offering – i.e. for A LONG TIME.
  - This is why it is critical that you plan and manage your migration to CMP many months in advance.

Z End

- Thank you for coming!
- If you have any questions, please email us at technical@watsonwalker.com
- Please remember to complete an evaluation – Session number 21601.