


And Now for Something Completely Different: An insider's guide to Tailored Fit Pricing

John Baker
IntelliMagic


November 2019
Session **LK**



Agenda



- Introductions
- Review how we got here
 - Fundamentals of IBM pricing
- Tailored Fit Pricing
 - Two options
- Comparison
- Recommendations and Q&A



IntelliMagic Sessions at GSE-UK

Date	Time	Room	Session	Title	Presenters
Wednesday November 6	10:15 AM	Monza	DG	Flash Storage Technology: Get on the flash wagon or keep on spinning?	John Baker
Wednesday November 6	3:15 PM	Woodcote	LK	And now for something completely different: An insider's guide to Tailored Fit Pricing	John Baker
Thursday November 7	9:00 AM	Suzuka	AM	Road Rage against the machine: Understanding – and dealing with – z14 z15 traffic patterns	John Baker



How MLC Pricing is done today*

Sub-Capacity Pricing

- October 2000: IBM announces Workload License Charges
 - Sub-Capacity Pricing model at less than full machine capacity
- Two dimensions to determine costs:
 - Consumption (CPU/MSU) from SMF 70
 - Active Subsystems (IBM software products) from SMF 89
- IBM uses this report to generate monthly invoice

CPU time



- Accumulated dispatch time on one or more CPUs
- Typically reported in intervals (RMF type 7x)
 - SMF70PDT/EDT
- SMF70EDT is the basis for Tailored Fit Pricing
- All subsequent metrics are based on CPU time


Millions of Service Units

MSU

- Software measurement of machine capacity and consumption (similar to MIPS)
- Published values (see LSPR) – along with PCI (MIPS)
- Reported as MSUs per hour over four hour average (R4HA) for current software pricing (SMF70LAC)
- Converted from CPU time for Tailored Fit Pricing (SMF70EDT)

<https://www-01.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument>

Peak vs Peak




Sample: Pricing Tiers

Level	MSU Range	z/OS	DB2	CICS	IMS	MQ	Tier Rate	Accumulated MLC cost
Base	3	4000	5000	6000	10000	2000	27000	£27,000
0	4 to 45	350	300	300	600	150	1700	£98,400
1	46 to 175	300	150	150	300	100	1000	£228,400
2	176 to 315	200	100	115	230	75	720	£329,200
3	316 to 575	100	80	85	170	50	485	£455,300
4	576 to 875	90	60	65	130	40	385	£570,800
5	876 to 1315	60	55	60	120	35	330	£716,000
6	1316 to 1975	50	50	55	110	25	290	£907,400
7	1976+	40	50	55	110	25	280	

- Sample CPC @ 1000 MSUs: over **£600,000 per month**
 - Accumulated costs of Base to Level 4 plus 125 MSUs @ level 5 rate
- Average Rate: £600
- Incremental Rate: £330

Which R4HA should I care about?

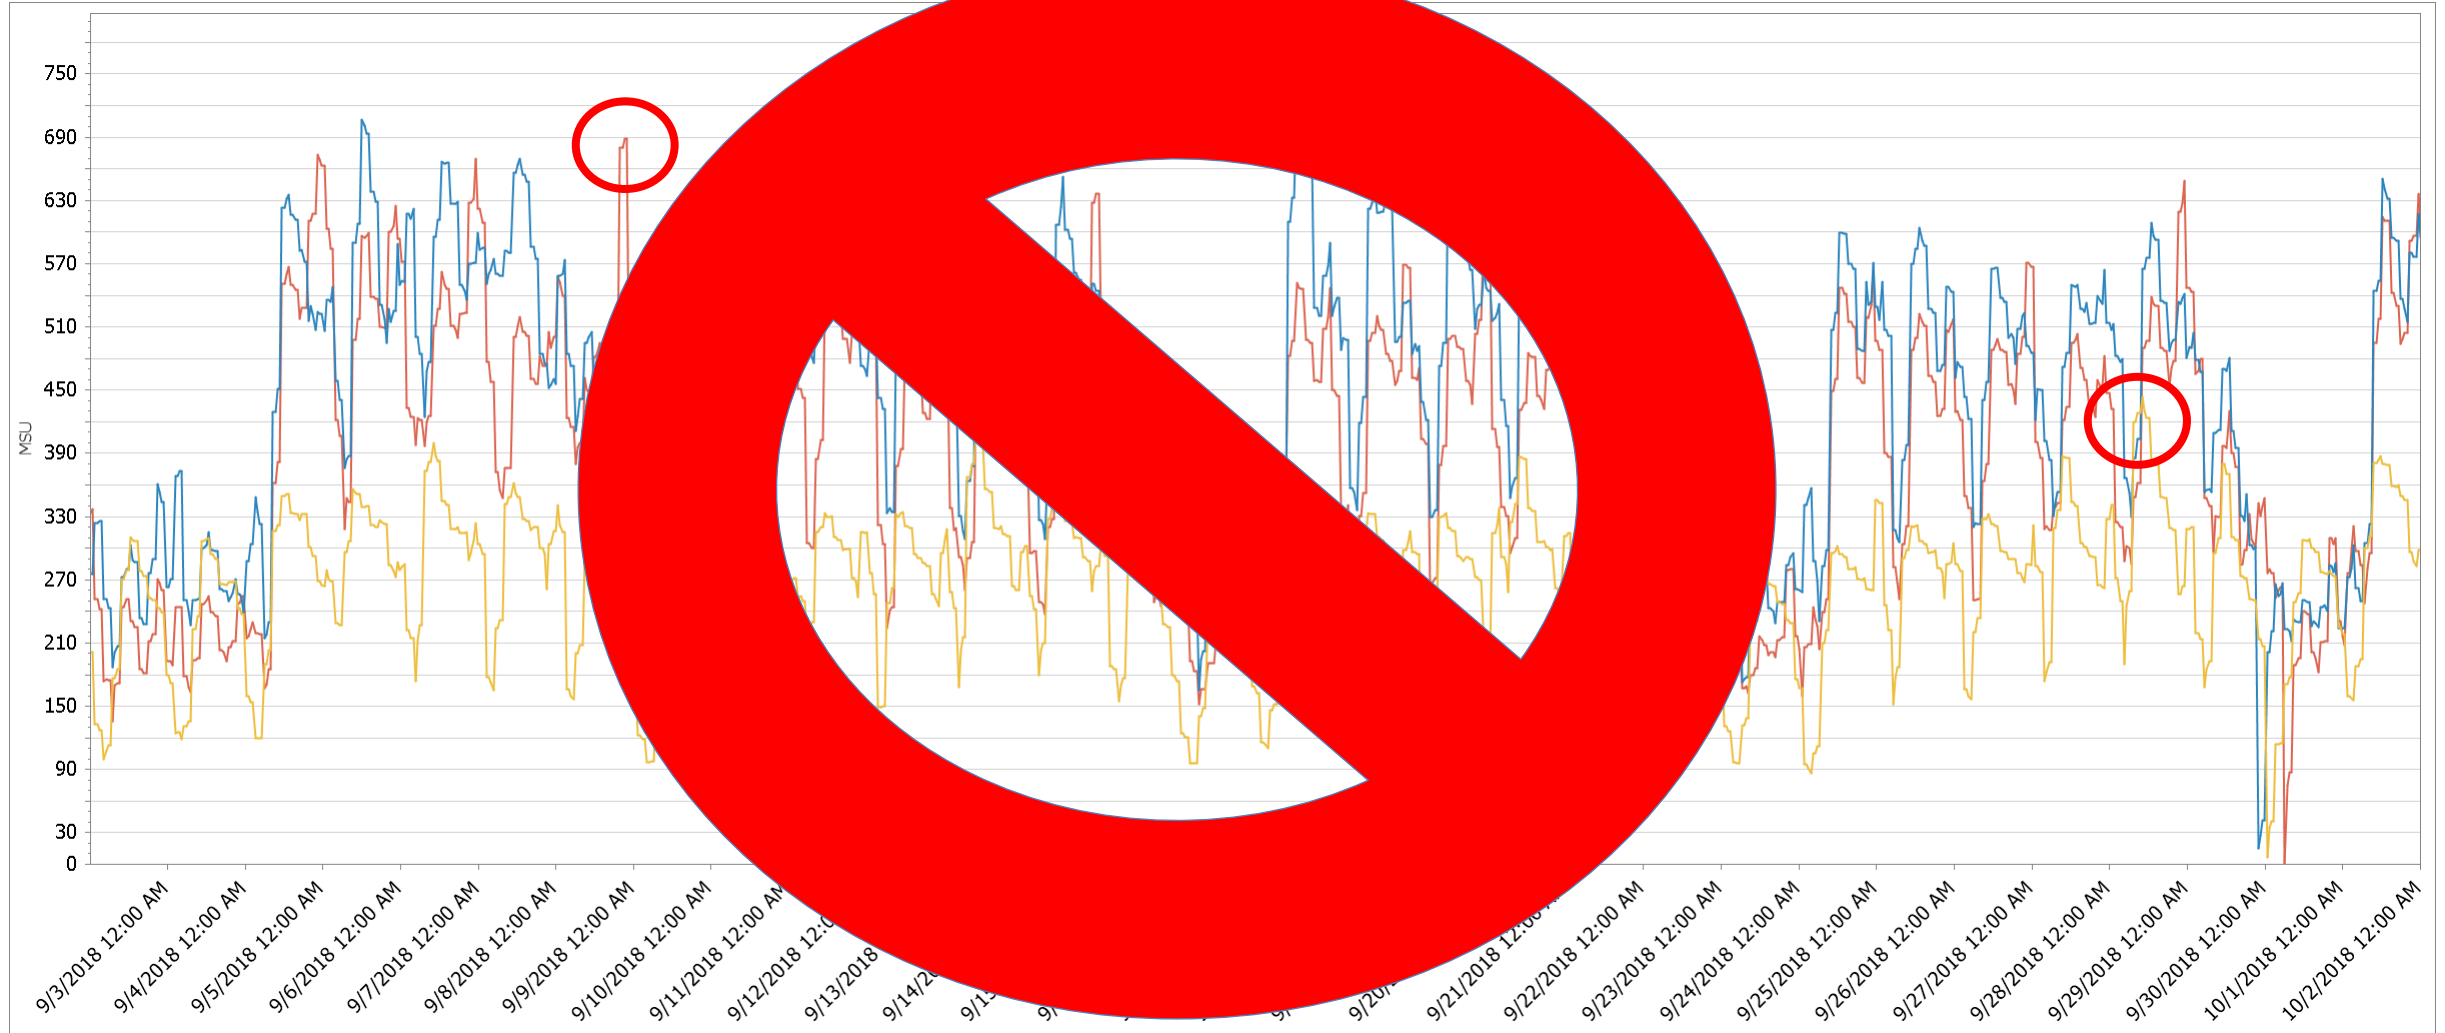
- Each hourly R4HA measurement from each LPAR where a subsystem is running is combined to create a peak R4HA by subsystem
- **It does not matter if the workloads on that LPAR actually use the subsystem(s) or not!**




Which R4HA Peak(s)?


- When combining LPARs to calculate your bill, individual peaks may be misleading
- The key to remember is to combine LPARs – by subsystem, by hour – for the entire billing period, before identifying for peaks
- You want: **“the peak of the sums”**
- **NOT: “the sum of the peaks”**

The Sum of the Peaks: **NO**



The Peak of the Sums: **YES**





Tailored Fit Pricing



1. Enterprise Capacity Solution
 - Blended full capacity pricing for everything
2. Enterprise Consumption Solution
 - Cloud Style “pay as you go”
 - Baselined on existing MSU consumed
 - Exchange R4HA for full consumption

<https://www.ibm.com/it-infrastructure/z/software/pricing-tailored-fit>

High Level (both options)

- No R4HA, sub-capacity or sysplex rules
- Requires z14+ and z/OS 2.2+
- 12 month 'baseline' of existing consumption
- Multi-year contract with committed MSUs
- "Aggressive growth pricing" beyond committed amount
- No plans to discontinue existing (R4HA) options

- Available: June 21, 2019

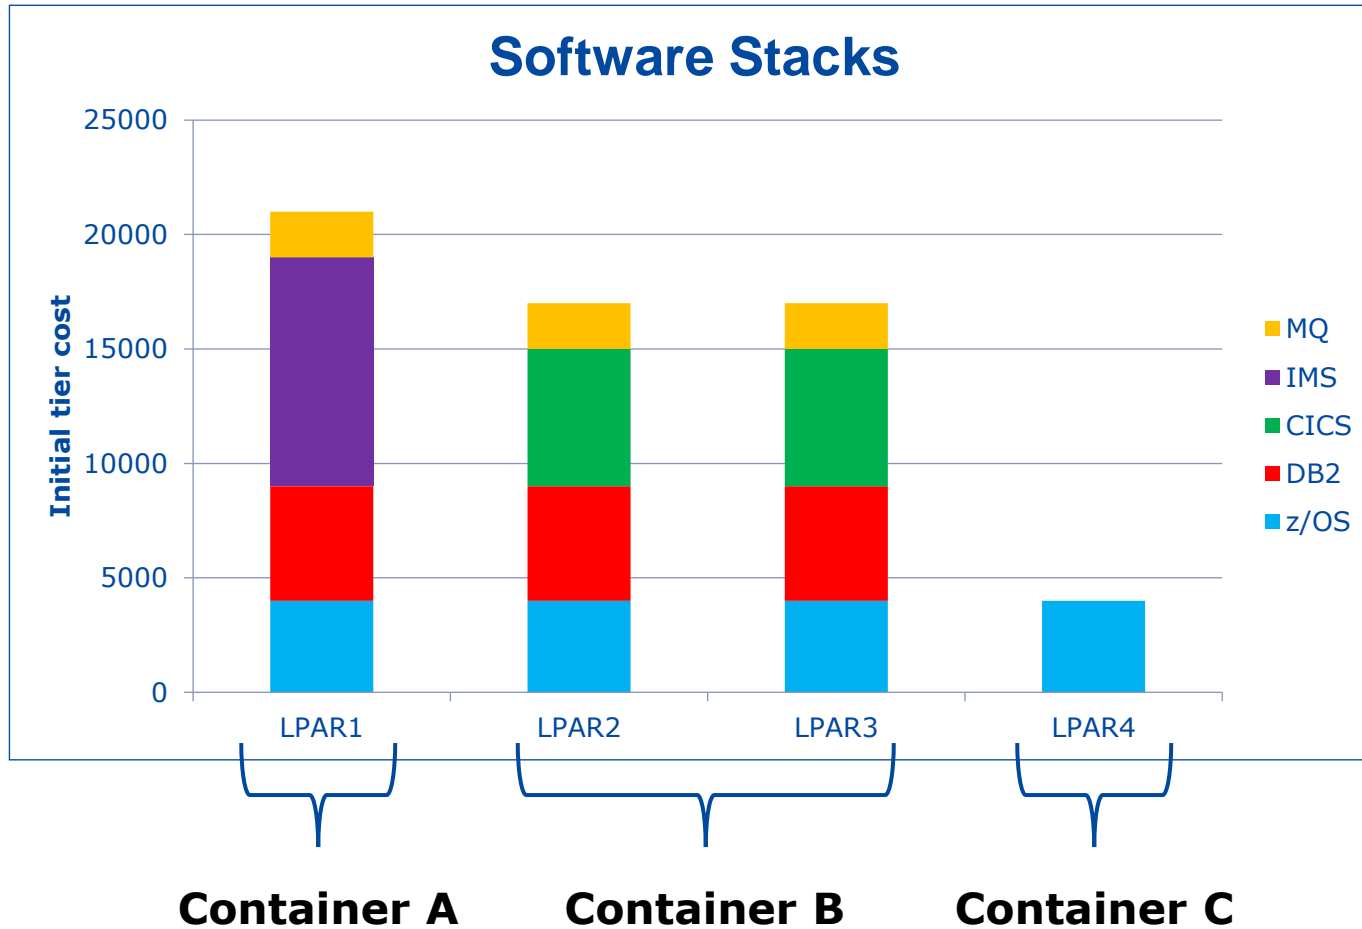
Capacity Solution

- Full Cap based on machine capacity, not usage
- MLC determined from baseline plus “an amount agreed to between IBM and the client to reflect estimated growth of the client’s individual workloads”
- IPLA license for total physical capacity
- All MLC and IPLA software can run anywhere*

Capacity Solution: Containers

- Container defined by number of engines
- With multiple containers, engines must be dedicated
- MSUs from machine model (LSPR)
- Blended 'zStack' prices covering hardware, OTC, MLC with single, unified pricing metric (\$/MSU)

Container Concepts



- Under Capacity option, engines must be dedicated to each container – *when more than one container defined*
- Consider single container (e.g. all Prod LPARs) to allow traditional CPU sharing
- Drawback is you will pay for full software stack on all LPARs in container

Consumption Solution

- Based on total annual CPU consumption converted to MSUs
- CPU consumption derived from Effective Dispatch Time (SMF70EDT) aggregated hourly
- Same container concepts but no need for dedicated engines as CPU time is accurately measured
- Annual entitlement, rather than monthly
- Monthly invoice 1/12 of annual, regardless of actual consumption (SCRT)
- Total consumption reconciled at year end


Consumption: Baseline

- Annual MSU entitlement derived from 12 month baseline
 - N7 section of SCRT (annual MSU consumption)
 - Plus any growth commitment
- MLC baseline derived from same 12 month charges
 - Likely R4HA-based using current pricing
- Monthly invoice 1/12 of annual
- Total consumption reconciled at year end
 - If you go under, no refund but can carry forward
 - If you exceed, annual 'true up' at discounted rate
 - Multi-year contract (e.g. 3 years)
 - If still under at end of contract, lost (use it or lose it)

A photograph of a lighthouse situated on a sand dune. The lighthouse has a white, textured body with a red-tiled lantern room and a weather vane on top. The background shows a vast, flat landscape under a blue sky with scattered clouds. The word "Contrast" is written in a large, purple, outlined font across the middle of the image.

Contrast


Comparison



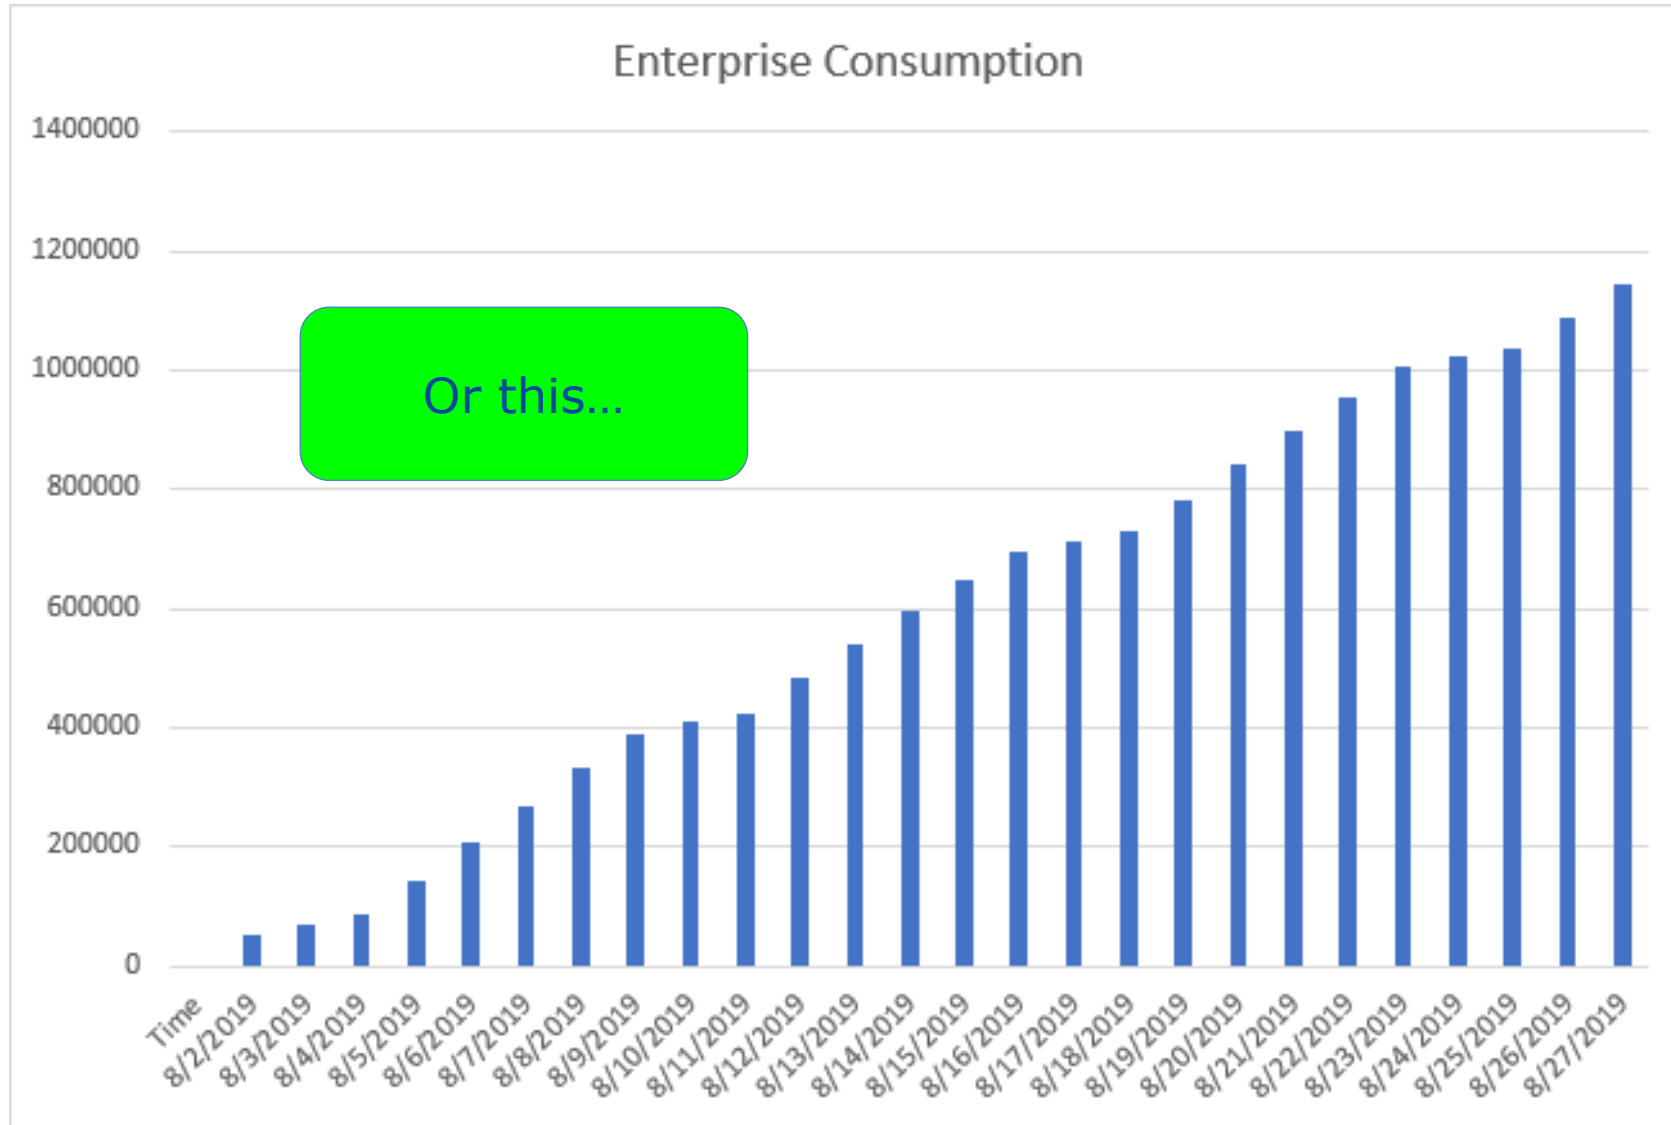
Comparison #2



MSU: Accumulated by Hour



MSU: Accumulated by Day



What about Capping?

Capping affects this




Consumption measures this

Caps reduce consumption *rate* but not absolute value


zIIP?

- More critical with consumption pricing
- zIIP 'crossover' (zIIP-eligible workloads on GP) only affected MLC during R4HA peaks
- With TF Enterprise Consumption, every GP CPU second counts

zIIP crossover



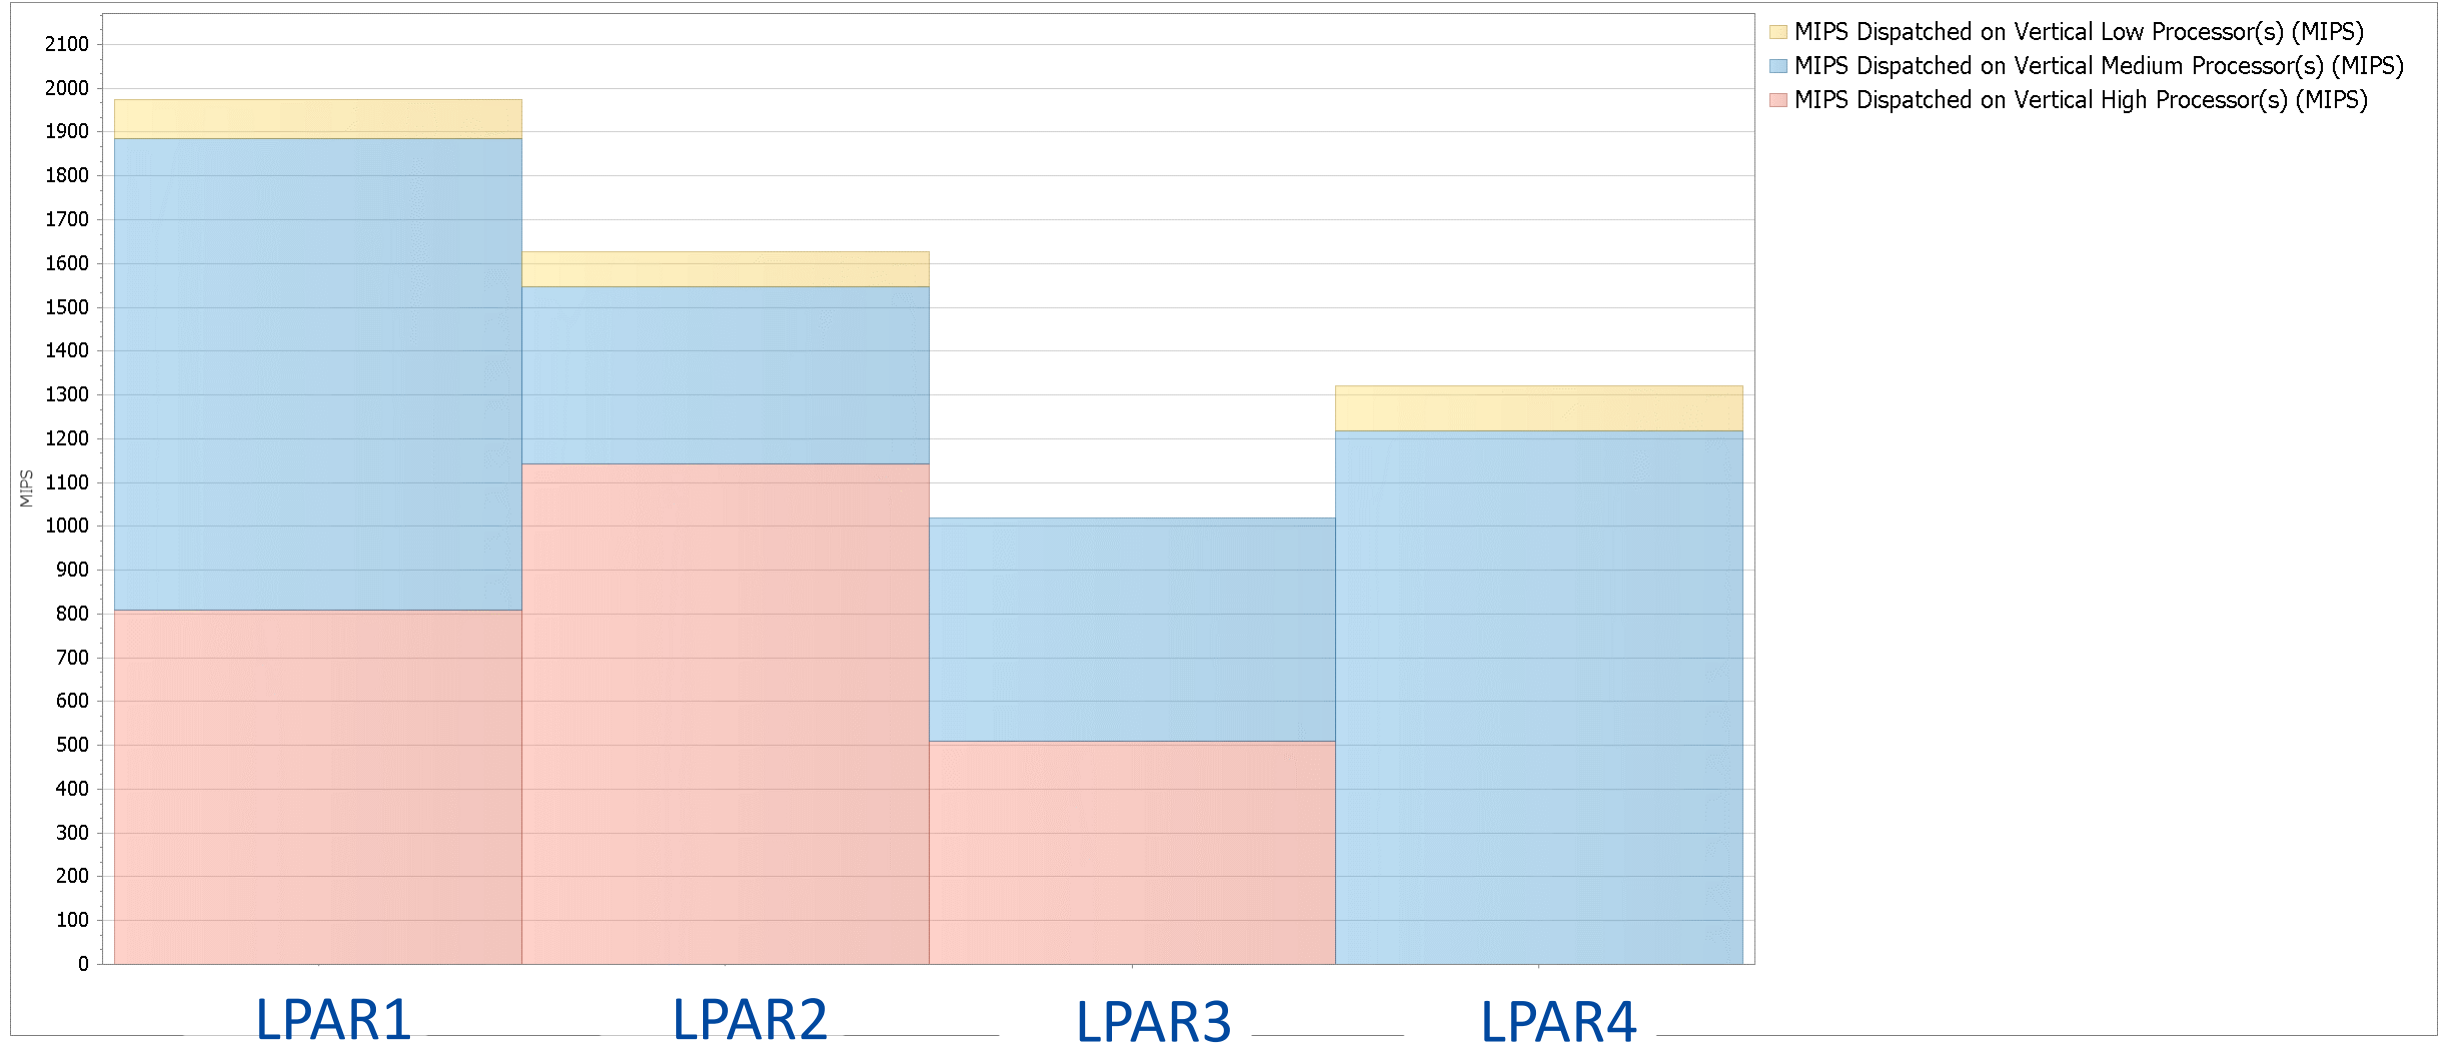
Don't start stuff you don't need




Hiperdispatch: Configuration matters!

- PR/SM dispatches LPARs using varying time slices
- Longer time slices improve probability of cache hits
- Cache misses elongate CPU time to do the same amount of work
- A Vertical High (VH) receives a 100 ms slice
- VM/VL receive $25 \text{ ms} * \text{physical} / \text{logical ratio}$
 - Minimum of 12.5 ms time slice
 - Try not to exceed 2:1 logical/physical ratio
- **Maximize VH use and avoid excessive logicals**


How costly is your configuration?




RNI: Polarity Matters



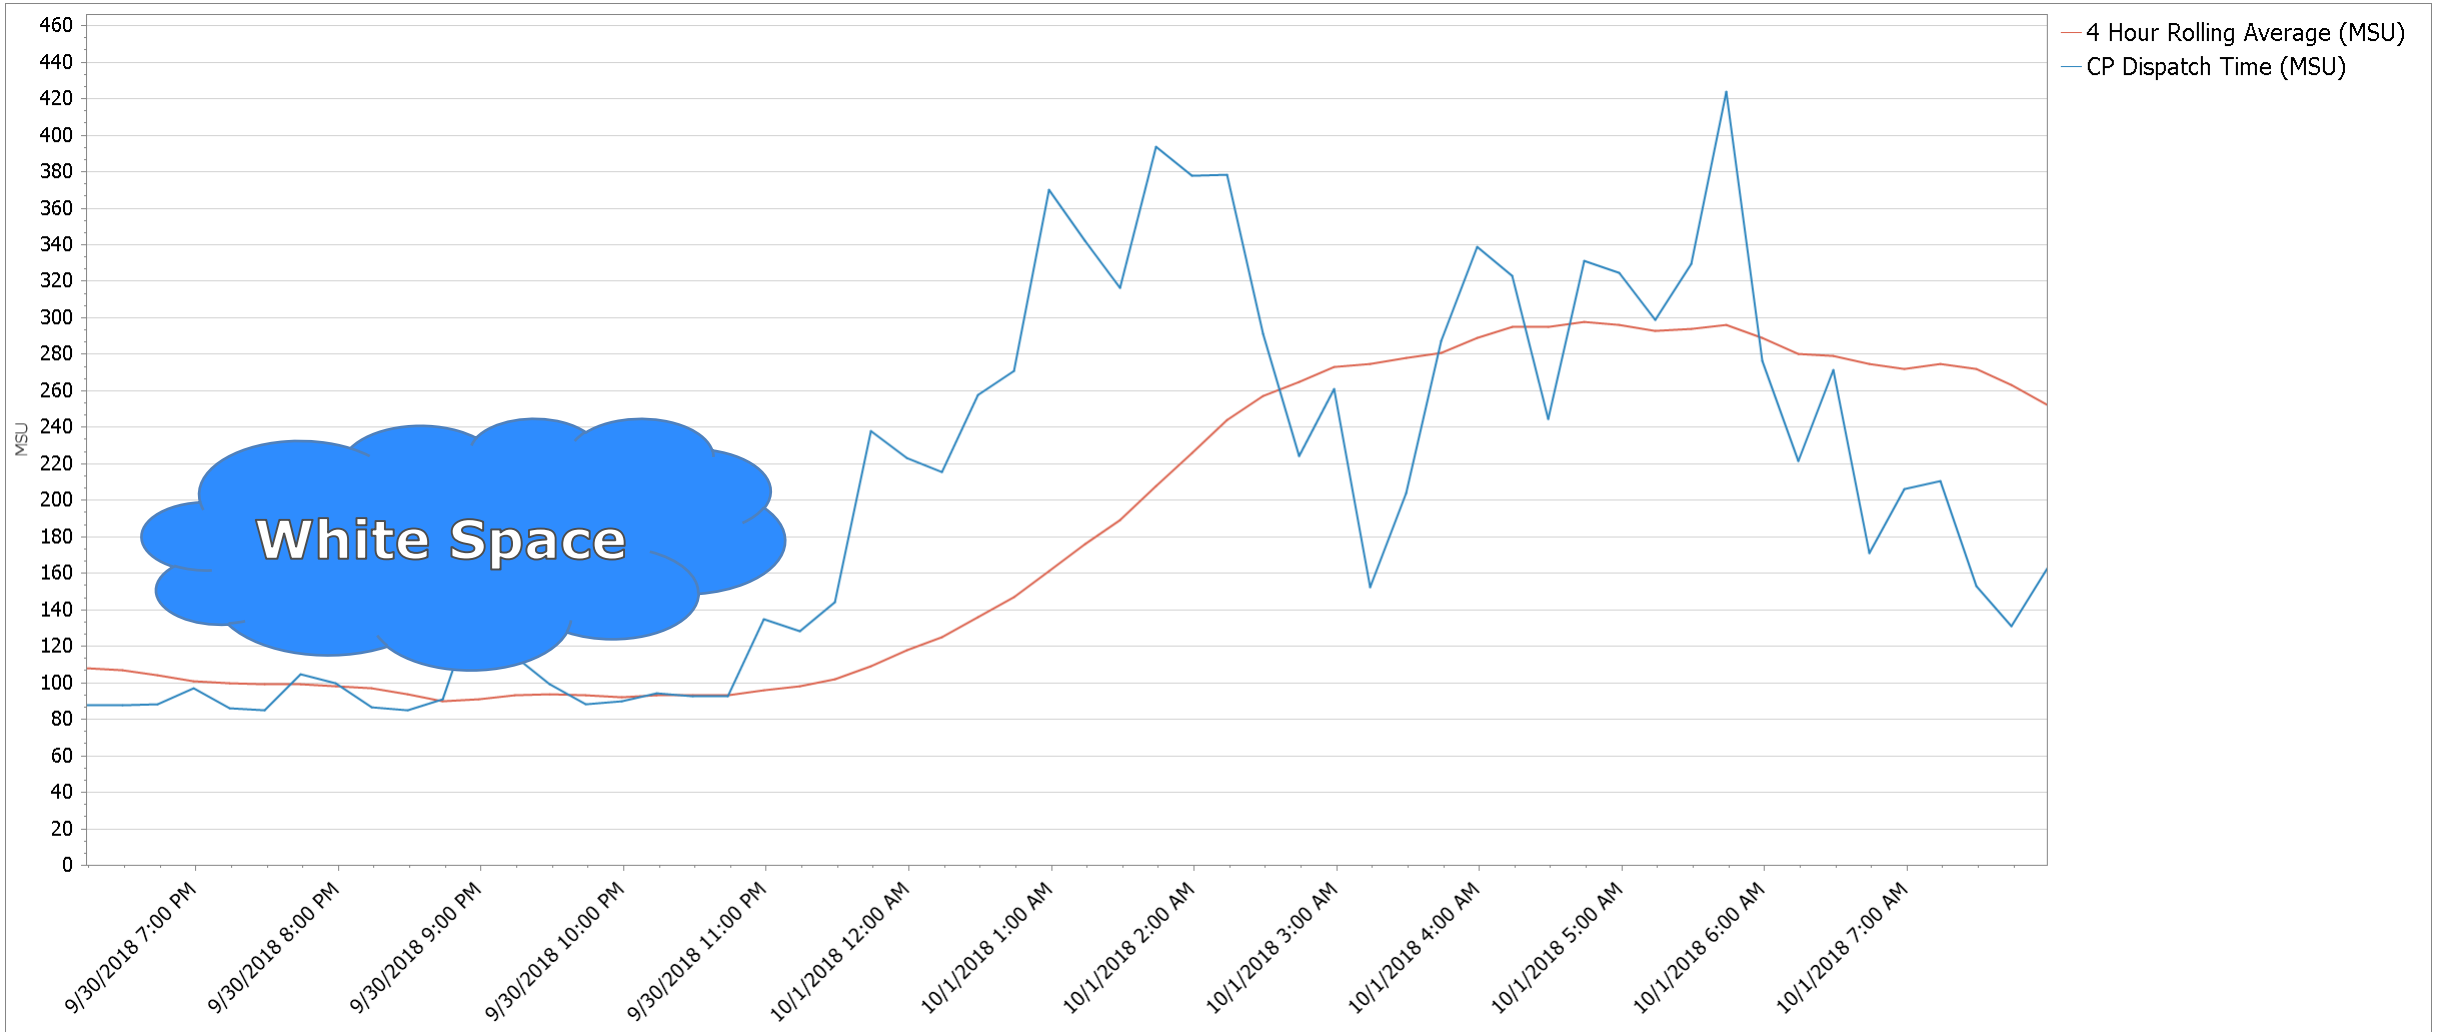
Flatten R4HA peaks



Identify Workloads



Fill in Valleys?



Observations

Tailored Fit Pricing for IBM Z delivers simple, flexible, and predictable cloud-like pricing, with economies of scale for all workloads on IBM z/OS

IBM United States Software Announcement 219-014
May 14, 2019

 ENUS219-014.PDF

Table of contents

- | | |
|---------------------|-----------------------------|
| ↓ Overview | ↓ Planned availability date |
| ↓ Key prerequisites | ↓ Description |

At a glance

Tailored Fit Pricing for IBM Z® is a transformational pricing option for IBM Z software. It offers simple, transparent, and predictable pricing for IBM Z software running on the z/OS® platform within a given country.

Tailored Fit Pricing introduces two comprehensive alternatives to the Rolling 4 Hour Average (R4HA)-based pricing model, for both new and existing workloads:

- The Enterprise Consumption Solution is a tailored consumption-based licensing model.
- The Enterprise Capacity Solution is a tailored full-capacity licensing model.

Both models dramatically simplify the existing pricing landscape, delivering flexible deployment options that are tailored to reflect the client's individual environments.

Both models include additional capacity for development and test environments as well as reduced pricing for all types of workload growth.

Recommendations



- Regardless of your chosen pricing model, you should aggressively pursue reducing your current R4HA software expense
 - Configuration, system and application measurement and tuning
- Are you growing?
 - If not, stay where you are (baseline is a commitment)
- Contract details matter!
- Consider independent services (IntelliMagic, Watson Walker)



IntelliMagic and Watson Walker have collaborated to create specific views that facilitate effective and efficient pricing plan decisions and MLC optimization work.

The IntelliMagic + Watson Walker partnership is a powerful duo representing the best available visibility and expert advice in this area.

For more information on the partnership see intellimagic.com/watsonwalker/



The alliance between IntelliMagic and Watson & Walker allows mainframe customers to benefit from IntelliMagic Vision's built-in intelligence and Watson & Walker's unrivaled services

A photograph of a lighthouse situated on a large sand dune. The lighthouse is a multi-tiered structure with a white base, a grey middle section with arched windows, and a red top section with a glass lantern room and a weather vane. The dune is golden-brown and stretches towards the horizon. In the distance, a few people and a dog are visible on the dune, and a coastal town is seen under a blue sky with scattered clouds. The word "Questions?" is written in a large, blue, sans-serif font across the middle of the image.

Questions?

Please submit your session feedback!

- Do it online at <http://conferences.gse.org.uk/2019/feedback/LK>
- This session is **LK**



1. What is your conference registration number?

This is the three digit number on the bottom of your delegate badge

2. Was the length of this presentation correct?

1 to 4 = "Too Short" 5 = "OK" 6-9 = "Too Long"

1 2 3 4 5 6 7 8 9

3. Did this presentation meet your requirements?

1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9

4. Was the session content what you expected?

1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9