How to Not buy the wrong CPC/CEC and what if you did?

Session Number: TECH_104s1
Tue Feb 24, 10:30am-11:30am
Room: **Salon 14**

Donald Zeunert
Independent Consultant
dzeunert@protechtraining.com

Agenda

- More MSUs isn't always the answer
- Determining Optimal CPC Model for your workload
- Don't upgrade too late.
- What if you bought the wrong model?
- Case Study – CPC Upgrade Plans / Results

More MSUs Solves ?

Customer 1 – More & Faster MSUs

- **Critical path -** ran much worse
- **Issues -** Entitlement and Latent demand

Customer 2 – Fewer Faster, more total MSU

- **Last Upgrade -** Faster caused erratic response
- **Planned Upgrade –** More MSUs would not solve issues
- **Solved –** Entitlement tuning and more Slower engines

More Slower vs Few Faster

Small / Medium CPCs with < 10 GCPs / CPC and with high LPAR to CPC GCP Engines ratios

- **More slower –** typically better than few faster.

Questions to Answer:

- **Con** - Will workloads suffer with slower GCPs ?
- **Pro** - Will LPARs have significantly lower Cycles / Instruction ?

No longer huge speed increase / generation

z13 -> z16 Slower GHz Chips than EC12

Family	GHz	% GHz Chg	MSU	Gen % Chg	EC12 %Chg
IBM Z® z17 A01 (9175) EC	**5.5**	6%	306	10.1%	**62.8%**
IBM Z® z16 A01 (3931) EC	5.2	0%	278	9.9%	47.9%
IBM Z® z15 T01 (8561)	5.2	0%	253	11.5%	34.6%
z Systems z14 (3906)	5.2	4%	227	8.1%	20.7%
z Systems z13 (2964)	5.0	-9%	210	11.7%	11.7%
zEnterprise EC12 (2827)	**5.5**	6%	188	**25.3%**	N/A
zEnterprise z196 EC (2817)	5.2	18%	150	**30.4%**	N/A
zEnterprise z10 EC (2097)	4.4	163%	115	**42.0%**	N/A
z9 Enterprise Class (2094)	1.7	N/A	81	N/A	N/A

z9-> EC12

- GHz increase related to MSU increase

z13 -> z17

- No GHz increase.
- **Where did MSUs come from?**

Why are lower GHz higher MSU?

Automatic

- Extra Cache (layers / memory)
  - z16+ 4x more L2 cache
    - Spare Core – L2 Cache as L3/L4
- Cross Drawer Efficiency
- Smarter prefetching logic
- Smarter branch prediction logic
  - Newer - Execute both branches
- Enhanced out-of-order execution

Changes Required

- New more efficient instructions
  - Needs recompile w/ new compilers
- Store Into Instruction Avoidance
  - Needs new compilers or manual BAL changes
- Smaller DAT - for 1mb / 2GB page, more fit more on chip
  - Define and use Large pages

Chip Cache Evolution

zEC12
L1 private 64k i + 96k d
L2 private 1 MB i + "L1+" 1 MB d
L3 shared 48 MB per CP chip
L4 shared 384 MB per **book**
6 cores + 1 L3 / CP chip
6 CP chips + 1 L4 / **book**
4 **books** (STAR) / CEC

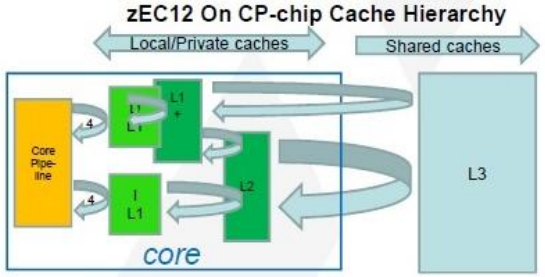

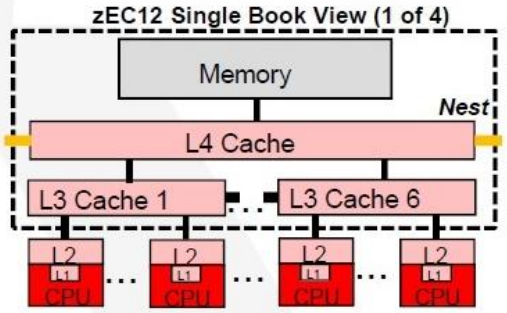

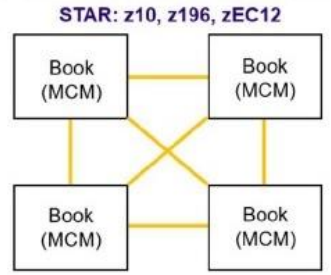

z13
L1 private 96k i, 128k d
L2 private 2 MB i + 2 MB d
L3 shared 64 MB / chip
L4 shared 480 MB / **node**
8 cores + 1 L3 / CP chip
3 CP chips + 1 L4 / **node**
2 **nodes** / **drawer**
4 **drawers** (NUMA) / CEC

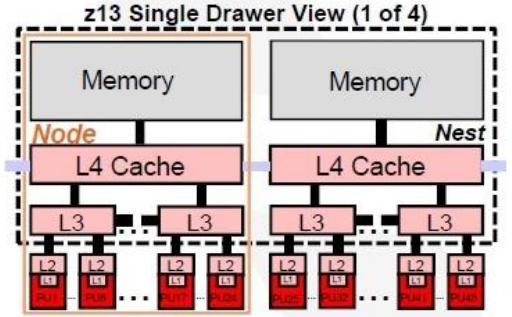

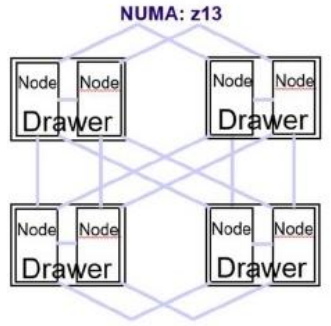

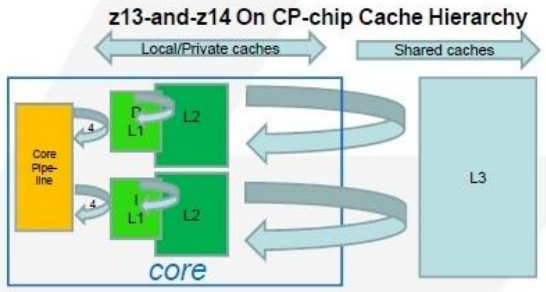

z14
L1 private 128k i, 128k d
L2 private 2 MB i, 4 MB d
L3 shared 128 MB / chip
L4 shared 672 MB / **drawer**
10 cores + 1 L3 / CP chip
3 CP chips / **cluster**
2 **clusters** + 1 L4 / **drawer**
4 **drawers** (numa) / CEC

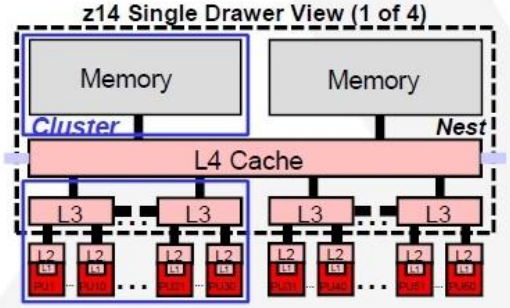

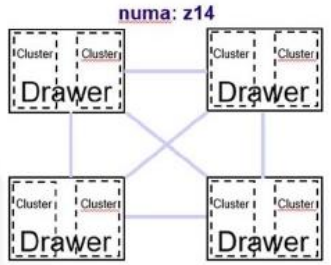

LPARs that cross drawers always a bad idea

- Really bad on z13
- Worse on z17 than z14+
- z17 great for drawers with spare engines as local cache used for others remote

MIPS Rating – Cache Hit Impact

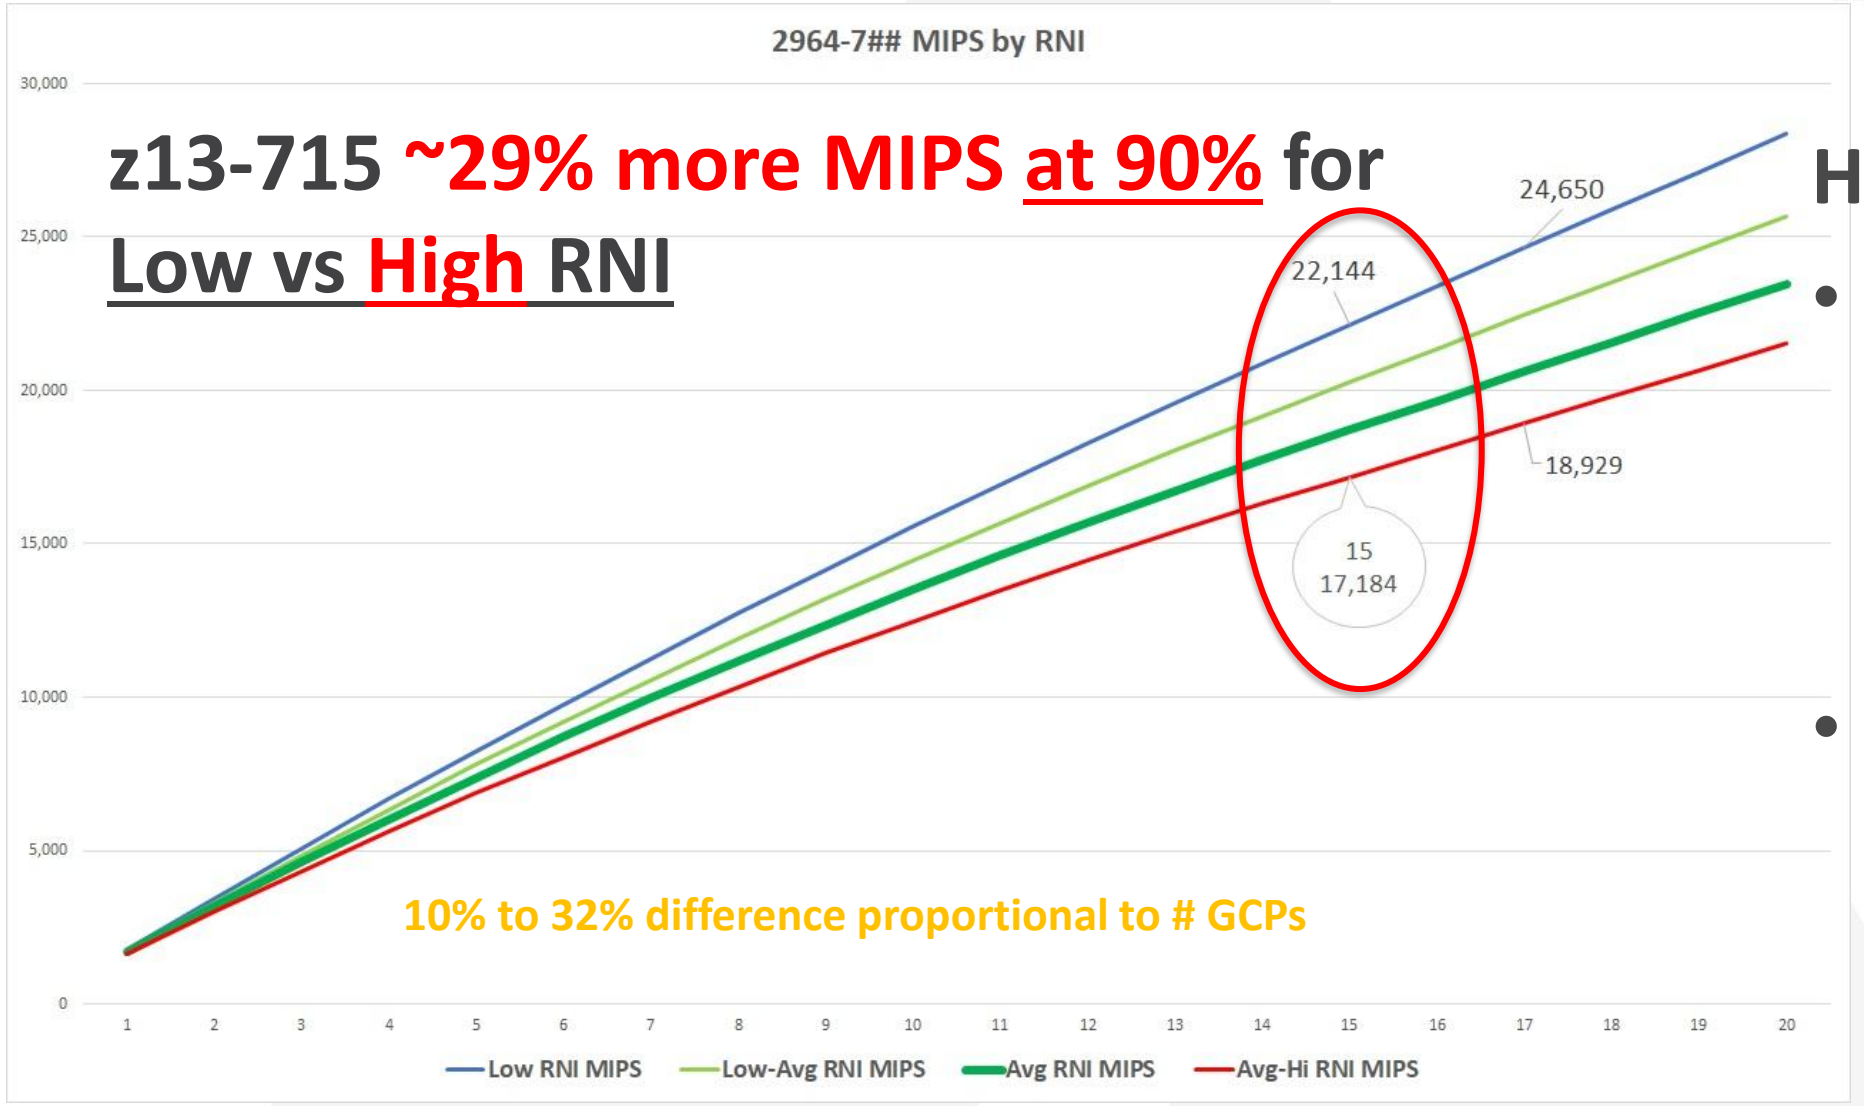

Higher RNI =

- Higher CPU secs / MSUs for same workload
- Fewer MSUs delivered per GCP

HiperDispatch – Improve Cache%

Vertical vs Horizonal Polarity

- **Vertical High – Physical CP virtually dedicated to LP**
- **Vertical Medium – split of entitlement < 2 across 2 CPs**
- **Vertical Low – LPs defined not needed to satisfy entitlement**
  - Parked Engines – VLs are parked when not needed
    - Unless 2 GCP LPAR
  - More likely to be in another Drawer

IBM z16™ 6## 28% Faster

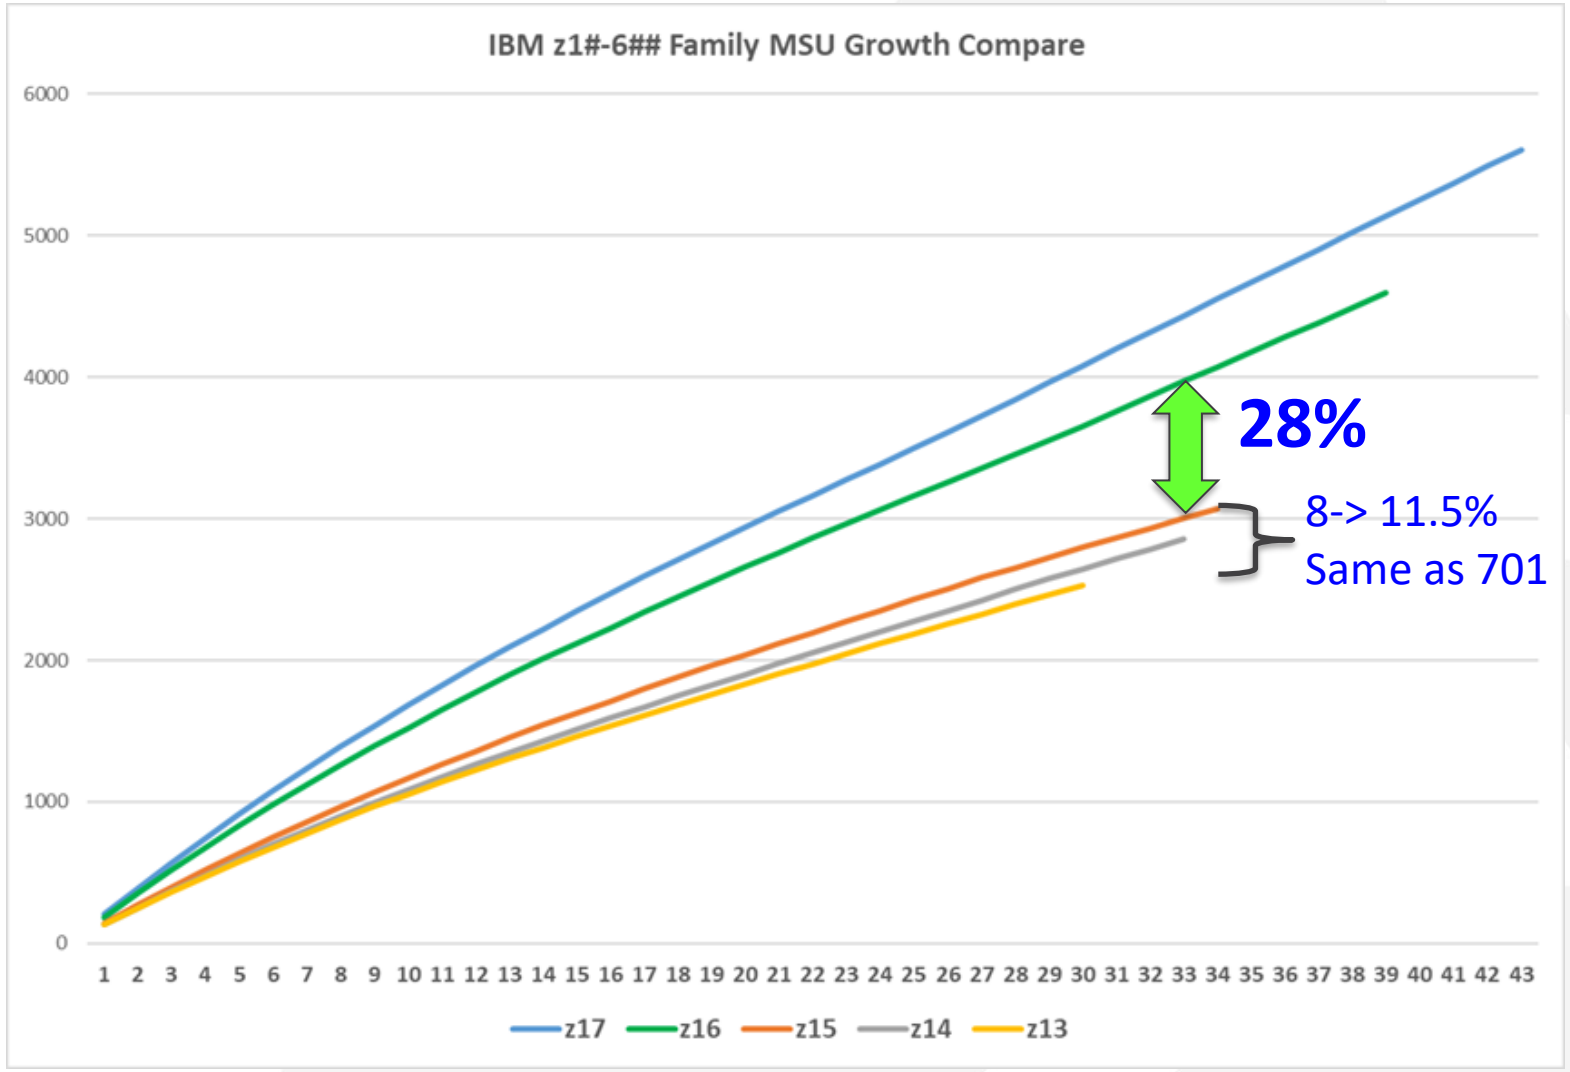

- More Cache is better.
- More GCPs EQ More Cache
  - IBM Wants 7## customers to consider 6##
- z16™ Faster z15™
  - Family is 9.9% faster
  - 6## is 28% faster

Current IBM z16™ 601 Very Fast

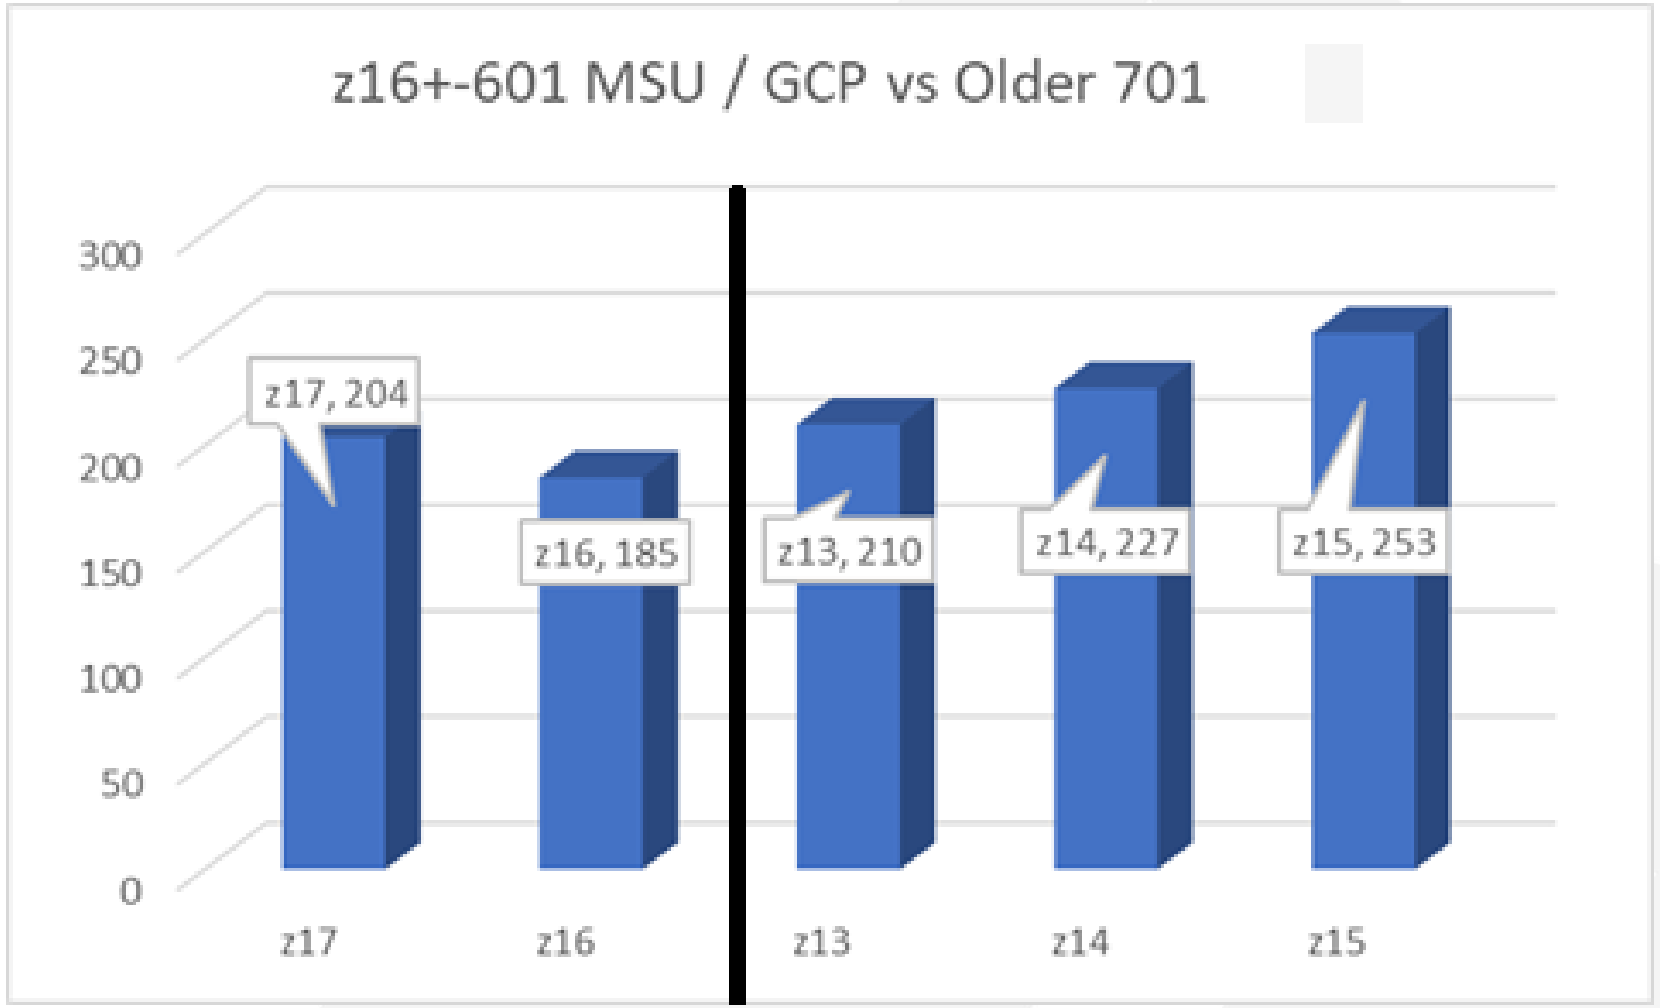

701s	% Faster IBM z16™ 601
z13	11.9%
z14	18.5%
z15	26.9%

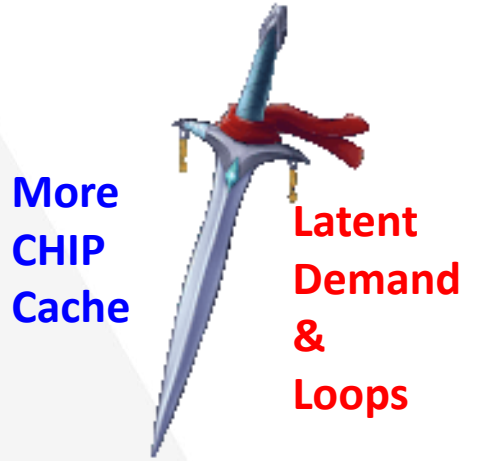

What is important to reducing IT Costs ?
CIO objectives and IT Budget Cost sources %

Shift gaze -
Hardware no longer $ issue

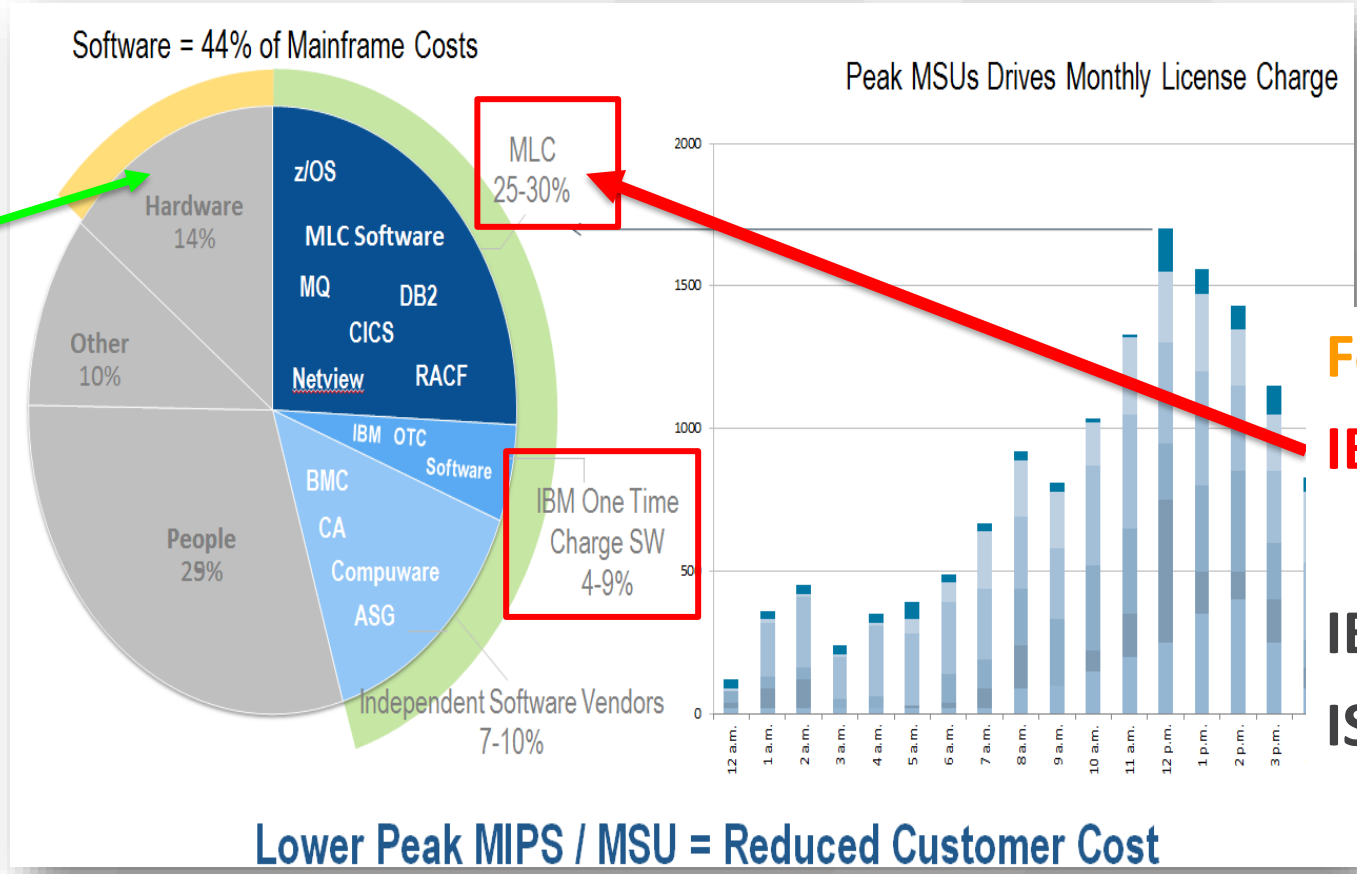

Meeting CIO objectives - With 80/20 Rule

Focus 20% of effort to get 80% of benefit of MF IT Cost reduction

Focus on 4HRA Costs

IBM MLC 25-30%

- Save 10% = 2.5%

IBM VUE and OTC 4-9%

ISVs 7-10%

- Save 10% = 0.7%

Determining Optimal CPC Model for your workload

Model Considerations

Number of LPARs

- Ratio of Logical to Physical Engines (GCP, zIIP)

GCP Speed vs Workload Requirements

- Online
- Batch

Cost

- Number of Drawers for optimal configuration

Fewer Faster

CONS

- **Loops consume more** w/o resource limits
- **Erratic Response –** LPARs can consume significantly more than its entitlement
- **Higher MSUs** for same workload, due to less Cache

PROS / Requirements

Critical Path Concerns

- **Online Single TCB Constraint**
  - CICS FOR
  - IMS Control Regions
- **Batch**
  - Single thread critical path

More Slower

CONS

- **Critical Path concerns**
  - Batch Pipes or equivalent
  - Avoid Thrashing
  - Reorg DBs faster or while
- **Extra Drawer(s) cost?**
  - Need 62% more 6x CPs than 7x for same capacity
  - Very Large LPARs - Number of GCPs and zIIPs may exceed a single drawer. Very bad on the z16 and newer models.
  - zIIPs - Too many for drawer, SMT-2 may be better option

PROS / Requirements

- **Loops limited by engine speed**
- **LPARs LP total MSUs available lower than faster GCPs**
- **More Cache – Lower RNI / MSUs for same workload**
- **7## zIIP, CF/ICF speed**
  - Higher Capacity @ 2x Limit
  - GCP Spins Slower – save MSU
    - Also zHiperLink relative speed

CF Spin (6## vs 5##) < MSU

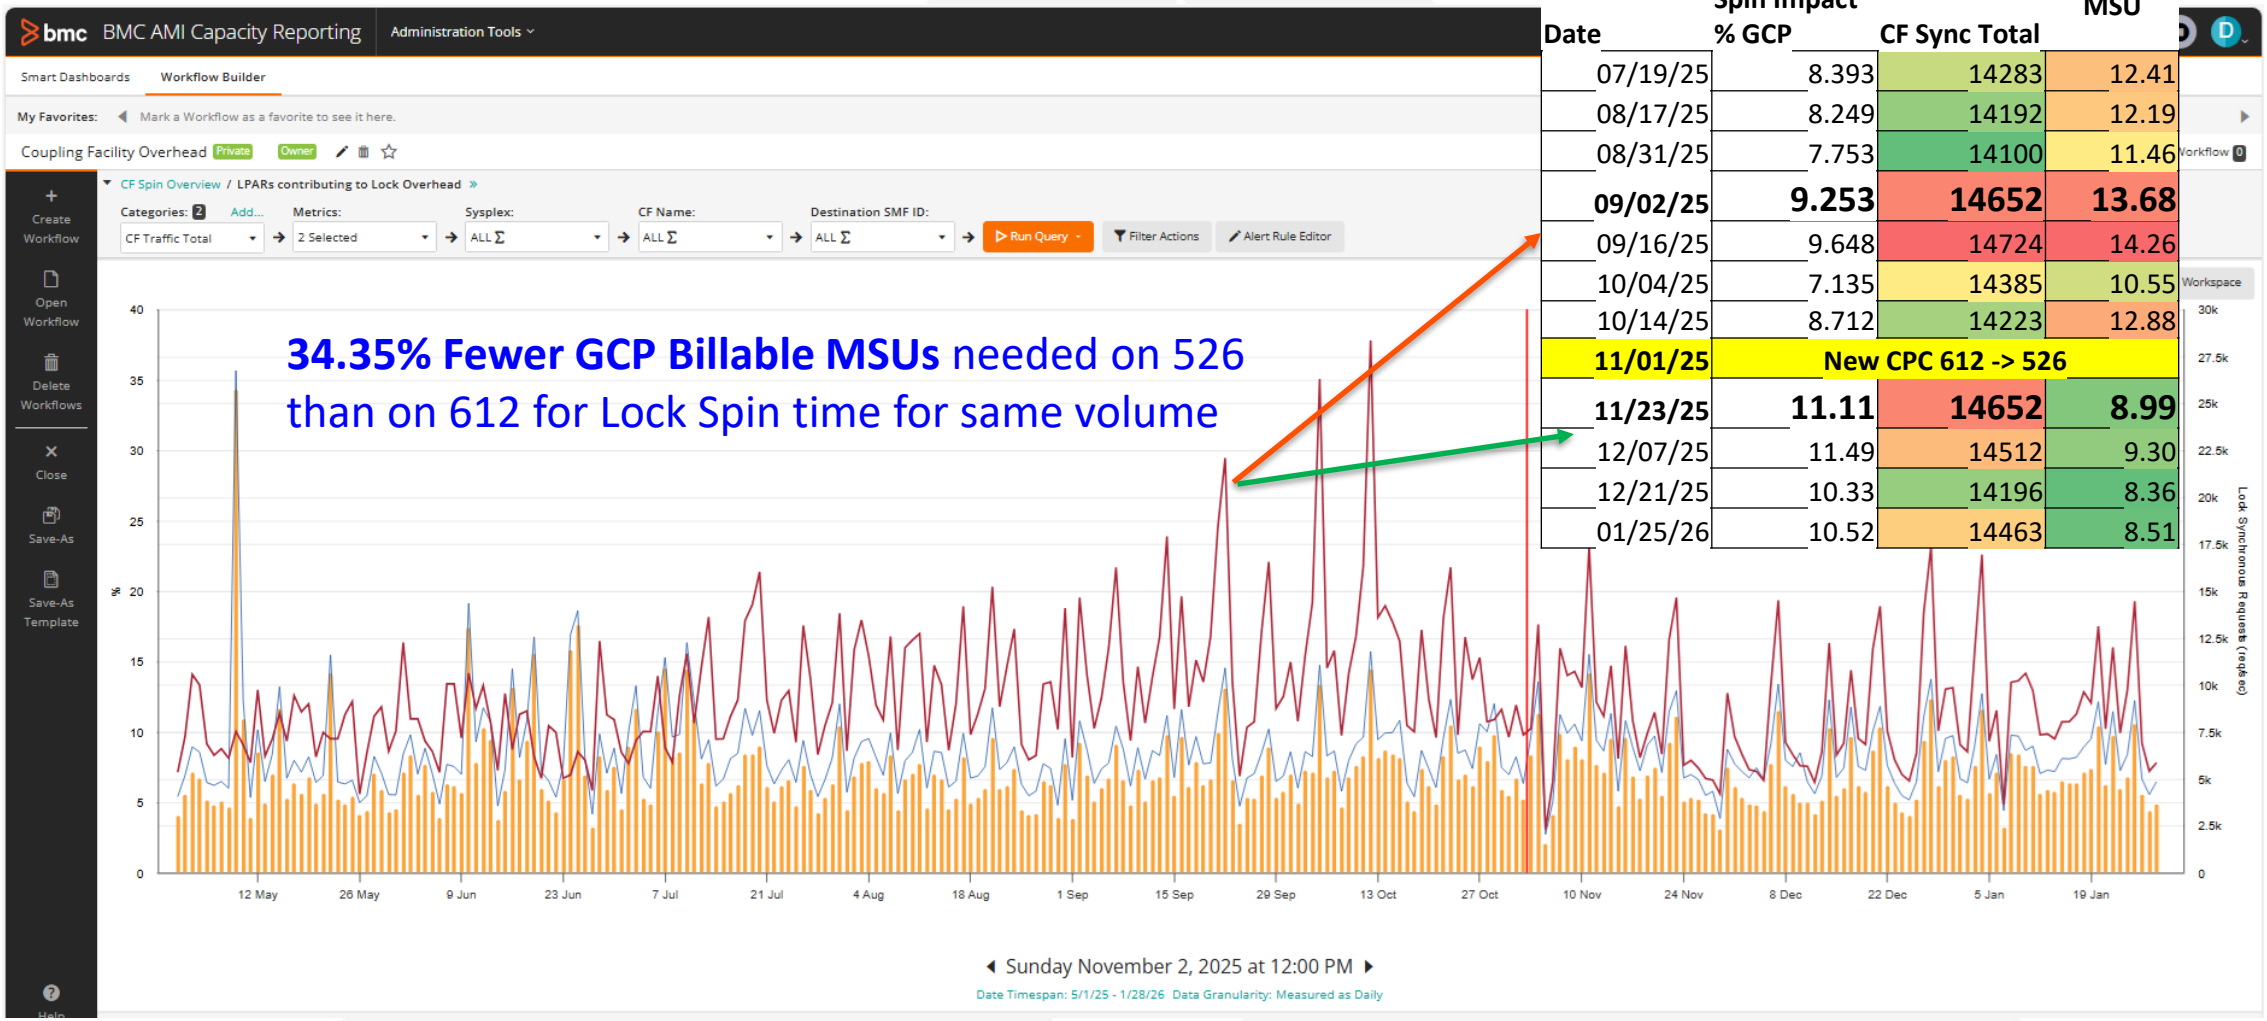

Date	Spin Impact % GCP	CF Sync Total	MSU
07/19/25	8.393	14283	12.41
08/17/25	8.249	14192	12.19
08/31/25	7.753	14100	11.46
09/02/25	**9.253**	**14652**	**13.68**
09/16/25	9.648	14724	14.26
10/04/25	7.135	14385	10.55
10/14/25	8.712	14223	12.88
11/01/25	**New CPC 612 -> 526**		
11/23/25	**11.11**	**14652**	**8.99**
12/07/25	11.49	14512	9.30
12/21/25	10.33	14196	8.36
01/25/26	10.52	14463	8.51

Risks Older Gen 6## to z16-6##

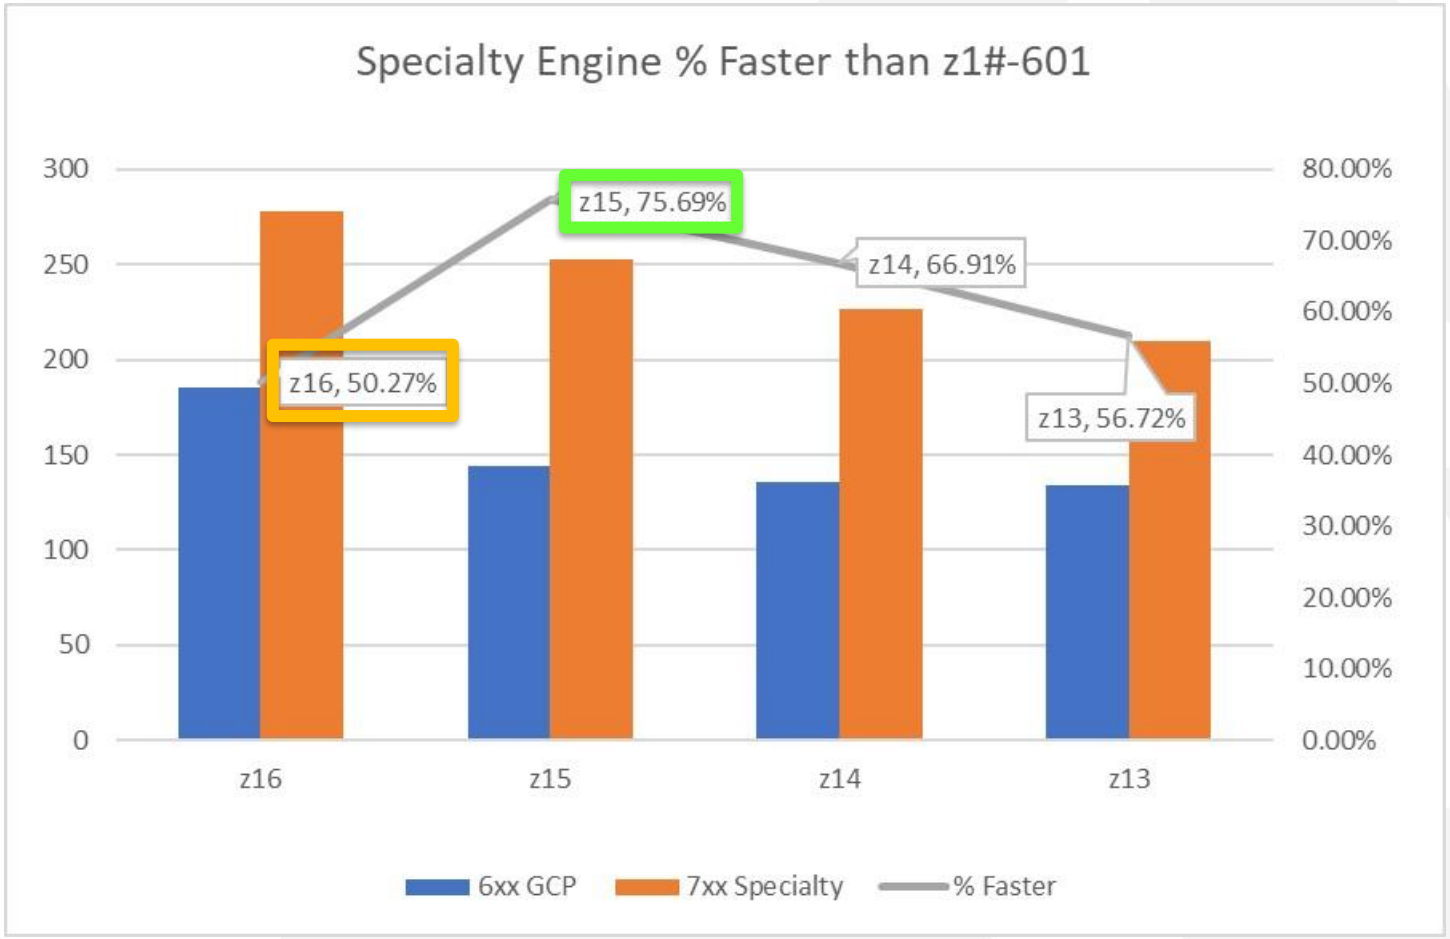

- **Lose some zIIP and CF Speed advantages**
- **Faster Engine concerns**
  - Loops Faster
  - Low entitlement 2 LP LPAR higher max w/o abs Cap

z15 6x to z16 6x - zIIP Risks

Model	#CP	MSU	%Incr	MSU /CP	zIIP	MSU	MSU/ zIIP	%Chg	Ratio
z16-610	10	1524	**30%**	152.40	20	3857	192.9	**11%**	2.53
z16-608	8	1259	**8%**	157.38	**20**	3857	192.9	**11%**	3.06
z16-608	8	1259	**8%**	157.38	16	3258	203.63	**-6%**	2.59
z16-607	7	1121	**-4%**	160.14	14	2953	210.93	**-15%**	2.63
z15-610	10	1170	N/A	117.00	20	3472	173.6	N/A	2.97

GCP to zIIP Ratio

- **2x = Drop**
- GT 2x = Ok

Z15 -> z16 Business Class

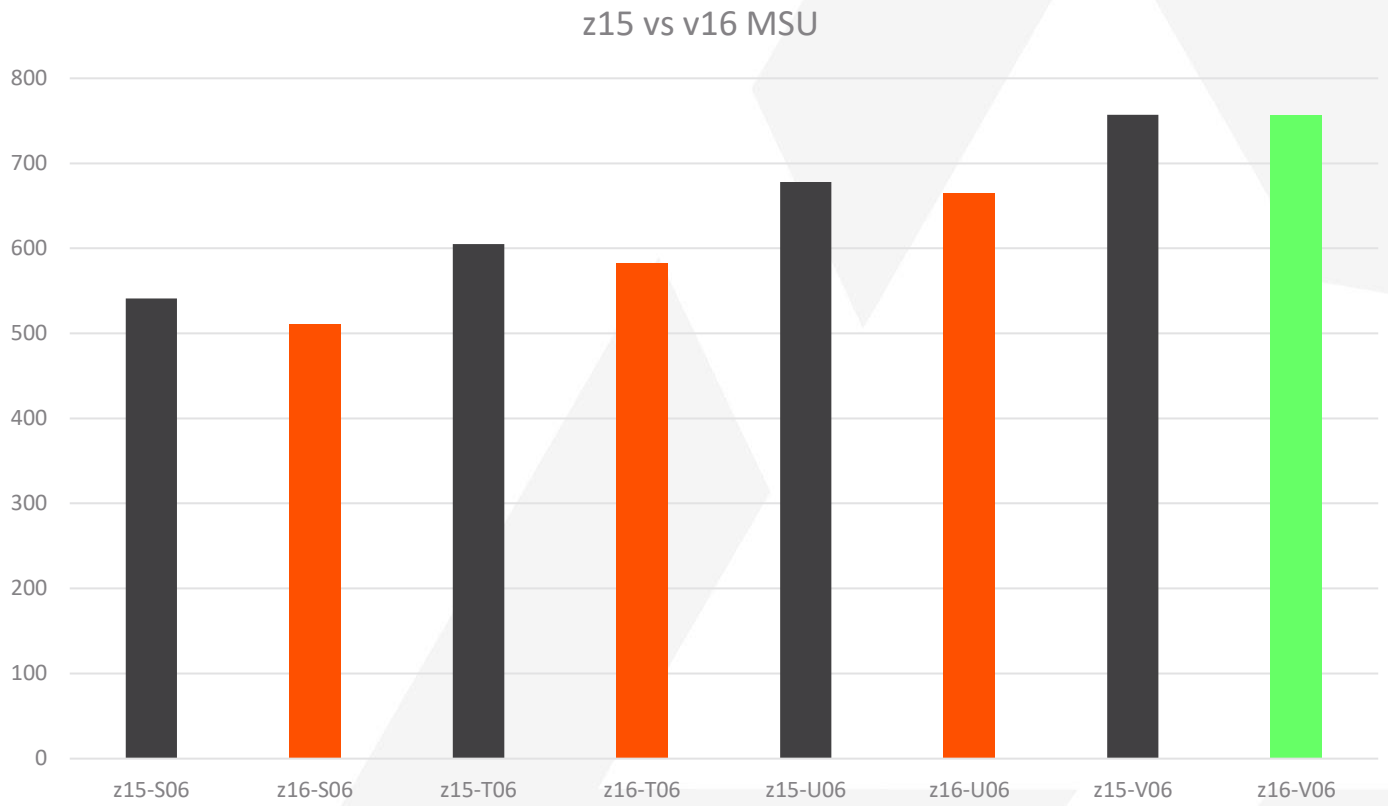

- Unlike Enterprise where objective was to give 7x option to migrate to 6x.
- Business Class is slightly slower to allow 6x ->6x to add more GCPs w/o significant more MSU

Critical Path Concerns

Online (CICS, IMS, DB2)

Address Space > 80% of 1 CP?

- **DB2 –** Non- issue, uses threads or Enclaves. Not single threaded
- **CICS –** Very little has to be single threaded. But it might be.
  - Not using Thread-safe
  - Not using MRO
- **IMS CTLRGN –** used to be issue, typically not currently

Batch

Critical Path – single threaded job stream?

- **Current batch window issue** need fast engines or change workload.
- **Batch Pipes** or equivalent could allow multiple slow engines to process in parallel

How to determine CP Constrained

1. **Identify jobs consuming high % of a GCP**
    1. GT 0.8 CPU Sec / Second or MSUs Used GT 80% of GCP MSU Rating
    2. Calculate ASID MSUs Used
2. **Determine Hot TCB % of Job**
    1. From SMF30_Highest_Task_CPU_Percent
        1. Note: Highest task could be a zIIP task, see SMF30_Highest_Task_CPU_Program for name
    2. Multiply Highest GCP TCB% by ASID MSU for single thread MSU
    3. Compare single thread MSU to proposed slower engine speed

Calculation Example

A CICS Region – Consuming 3 GCPs (from DB2 Threads and other TCBs).
Question how much is QR single thread TCB using?

1) **SMF 30(2) – 720 CPU secs in 900 (15 min) or 30(3 or 4) CPU Sec / Elapsed**
2) **SMF30_Highest_Task_CPU_Percent = 20%**
    1) SMF30_Highest_Task_CPU_Program = 'DFHSIP' (QR TCB)
3) **Single thread (QR) % of GCP (2700 * .20) / 900 = 60%**
4) **Multiply %GCP times MSU rating – z16-701 (278)**
    1) MSU Needed = (278 *.60) = 166.8
    2) Z16-601 – 185 MSU = 90% of GCP, concerning, split CICS AOR?
5) **Realize** your 7## is slower than 701 and Poor HiperDispatch not really getting rated capacity.

P.S. you could also use CICS statistics record QR TCB CPU

What if you bought the wrong CPC Model?

Typical Issues

PR/SM Settings not touched in years

Poor / High Logical to Physical Ratio

LPARs MSU consumption peak utilization different times of day

- Change LPAR Weights to match Online vs Batch needs

GDPS LPARs Weight inconsistent w/ normal consumption

HiperDispatch at SHARE

Zeunert's – Tune w/o new hardware

[Hiperdispatch - SLA improvements and MSU reductions]

[zIIP-A-Dee-Doo-That - So many new zIIP Concerns]

This presentation and others –

HiperDispatch - Tune by adding # of Engines (Sub-cap or current speed)

Upgrade Late Issues

Expect 2-5% more CPU consumed for same workload for every 10% increase in Utilization

Buy More MSU – Abs Cap

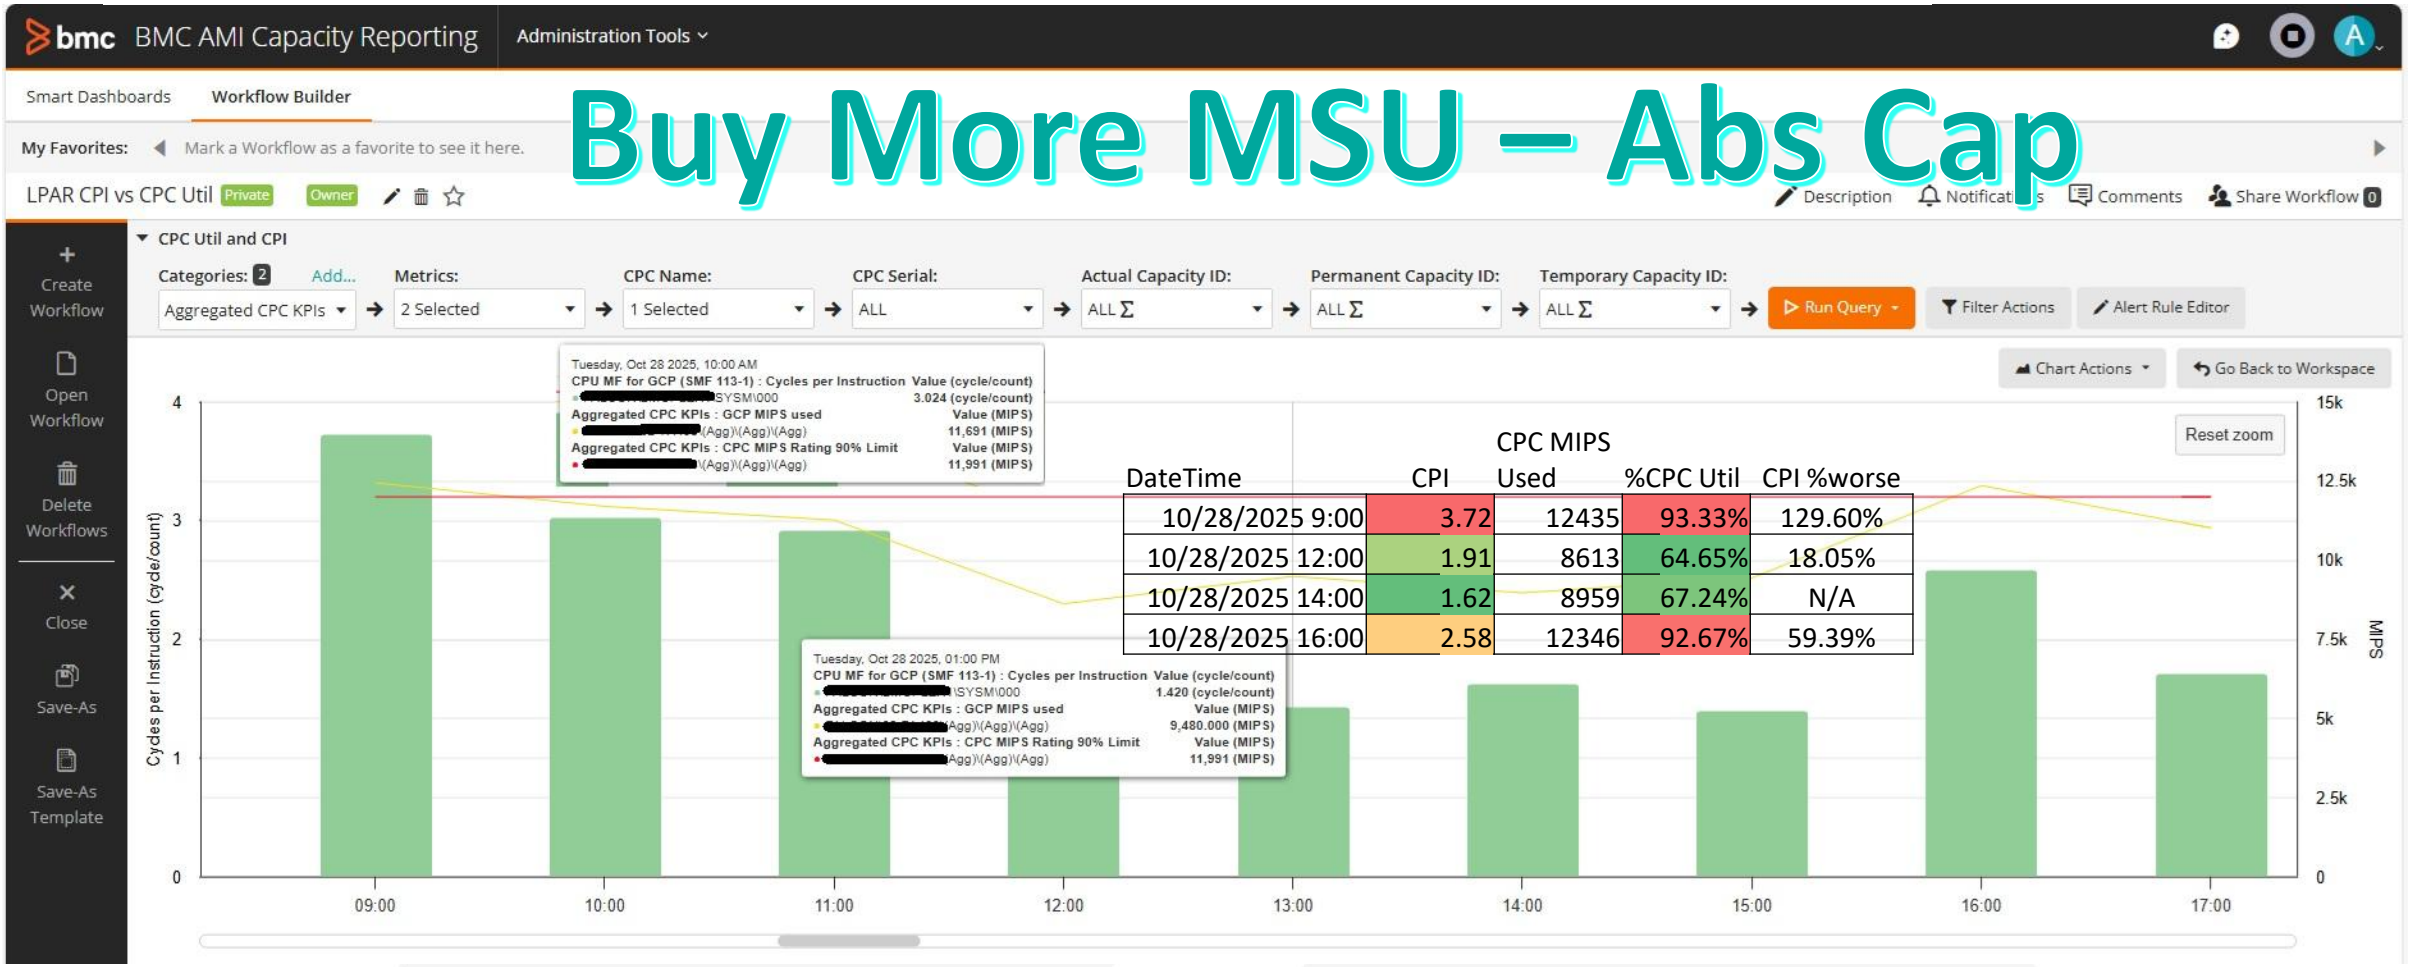

DateTime	CPI	CPC MIPS Used	%CPC Util	CPI %worse
10/28/2025 9:00	3.72	12435	93.33%	129.60%
10/28/2025 12:00	1.91	8613	64.65%	18.05%
10/28/2025 14:00	1.62	8959	67.24%	N/A
10/28/2025 16:00	2.58	12346	92.67%	59.39%

DZ- IF CPI > 2.5, investigate VL Usage, SIIS, Thrashing / MPL

Absolute Cap to <90% of Excess

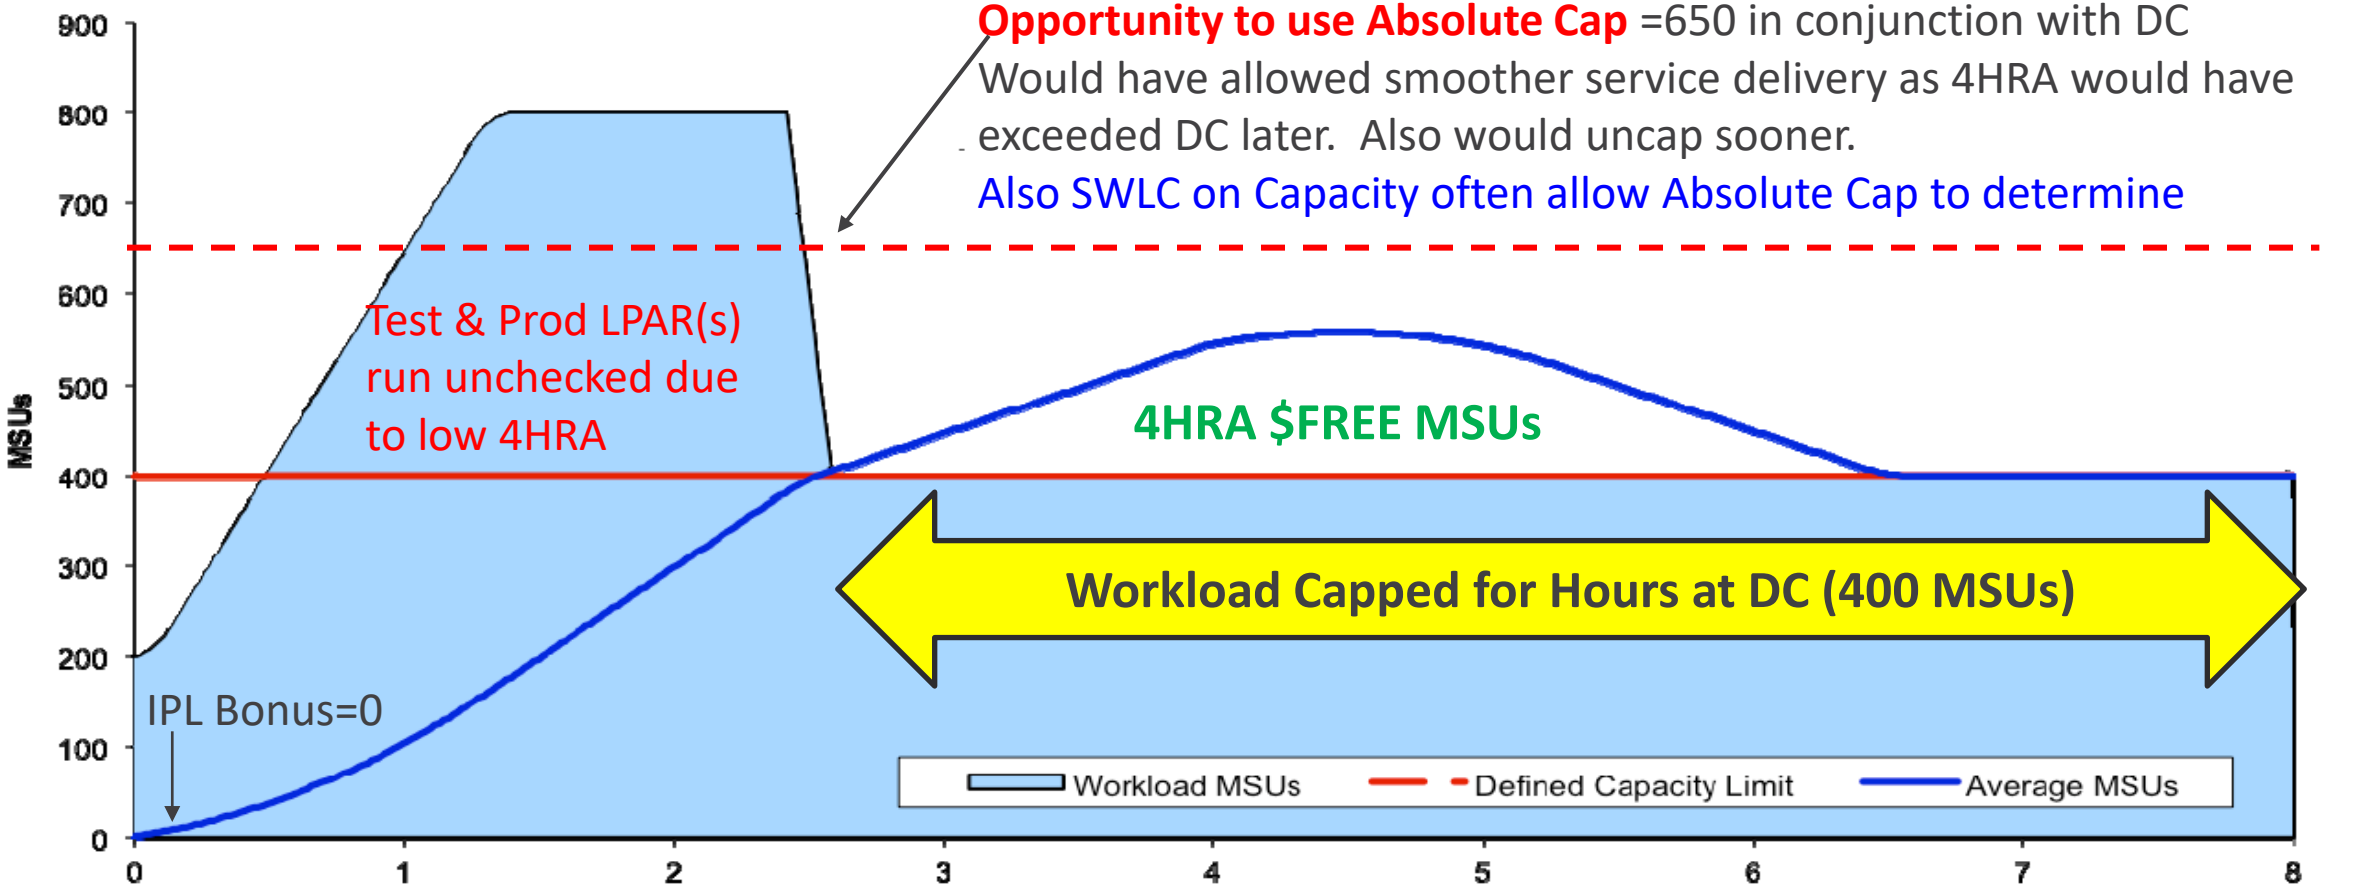

Technology Update Pricing EC12->z15

z15 TU - Quantity of z15 Full Capacity MSUs for a stand-alone server, or the sum of Full Capacity MSUs in an actively coupled Parallel Sysplex made up
entirely of z15 servers

TU7

MSU	Owned	z15
4	45	8.00%
46	315	16.00%
316	1315	18.00%
1316	2676	20.00%
2677	5476	24.00%
5477	999999	28.00%

z16- While specific "TU9" (Technology Update 9) pricing for the IBM z16 is not detailed as a distinct, public package, the z16 pricing is governed by the latest Tailored Fit Pricing (TFP) models, which prioritize cloud-like consumption over traditional, fixed-capacity models.

Technology Update Pricing

- **Encouraged customers** to migrate all CPCs to latest Technology level
- **AWLC / AEWLC Discount** for total MSUs purchased
  - Didn't have to be used
  - Absolute Cap for SW Contracts

Case Study – Upgrade Recommend

Disclaimers

ProTech Enterprise IT Training & Consulting –

- **is business partner w/ IBM, BMC Software Inc. and Rocket Software**

Anonymous Customers below, data shown w/ the tools that they own

- **In this Case – BMC AMI™ Capacity Management** was used for the analysis and data presentation

Customer 1 Case Study

Bought More MSUs on faster GCPs

Serious critical path batch issues on 2 of 10 LPARs

Source of Issue – Latent Demand

- **HiperDispatch not set up correctly**
- **Old CPC LPARs were borrowing 5 GCPs of VL from other LPARs not using their share**
- **New CPC Other LPARs used their share**
    - Critical path patch could only get to borrow 2 GCPs of VL

Borrow = Risk / Erratic

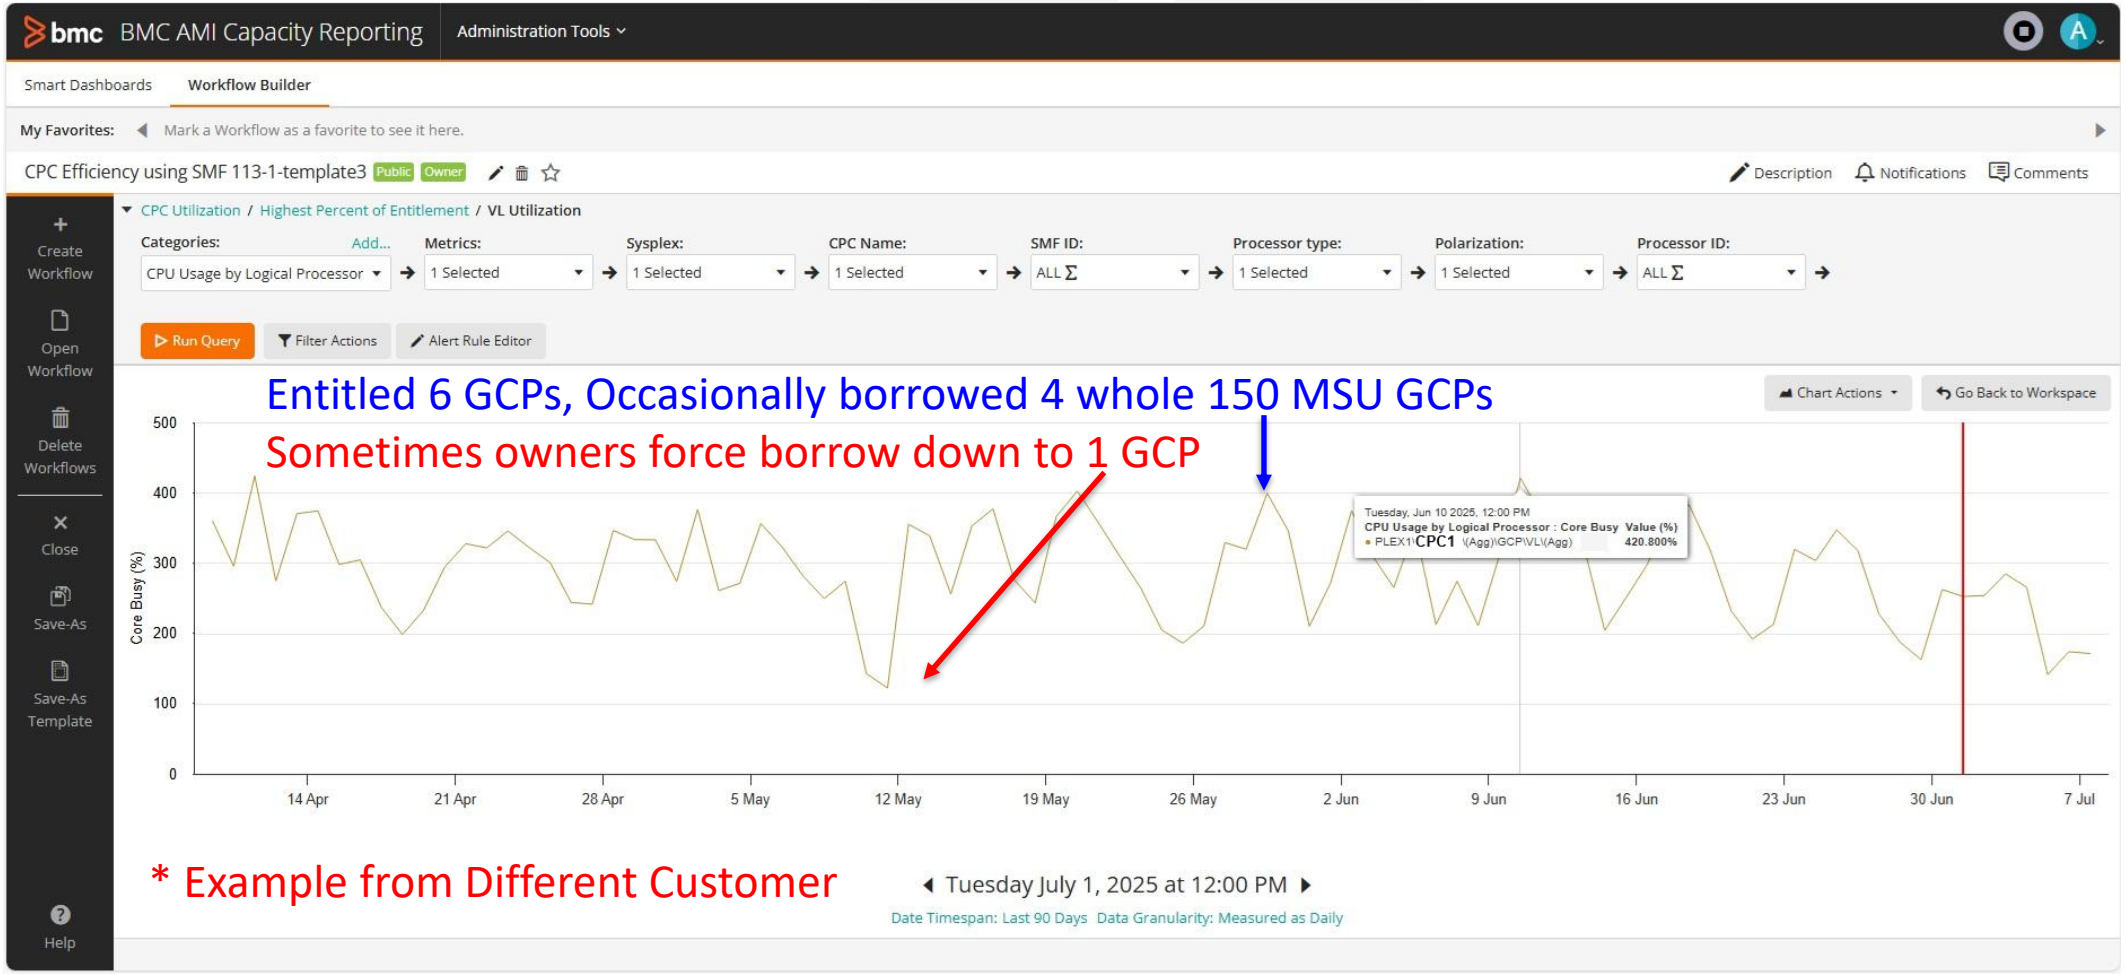

Customer 2 Case Study

Environment –

Erratic response time impacting productivity / revenue

Last Upgrade – Fewer Faster

z15-514 -> z16-612

- **90 % Faster** GCPs (77 -> 148)
- **66% More MSU** – (1070 -> 1774)

Customer Proposed Upgrade

z16-612 -> 617 (+564 MSU)

- **"Need" -** More MSUs for ProdPlex
- **Plan** more fast GCPs

Findings

- **ProdPlex needs 445 MSU of 564**
  - z16-617 Current Entitled Share only 131 MSU, very inadequate
- **HiperDispatch Tune –** yielded 78% of needed (348 of 445) MSUs

Logical / Physical Ratio = 6.9 x 1

- **37 LPARs –** Log 83 / Physical (12)
- **Upgrade Recommend –**
  - "More Slower" MSUs w/
  - z16-526 at +330 MSU
  - 40% Less MSUs than customer plan

Tuned to Reduce Upgrade Requirements

PRODPLEX Needed at least 445 MSUs of originally planned 617's 564 MSU, by default (current weights) they would have only received 131 MSU

Recommended – Entitled Share Increased to Average Used

Plex	Old % CPC	New % CPC	Old Ent MSU	New Ent MSU	Delta MSU
ProdPLEX	31.63%	51.27%	561	909	348
QACPLEX	21.40%	8.63%	380	153	(226)
DEMOPLEX	3.72%	4.06%	66	72	6
VM LPARs	29.77%	24.37%	528	432	(96)
Other	13.49%	11.68%	239	207	(32)
TOTAL	**100.00%**	**100.00%**	**1774**	**1774**	**(0)**

Gave MSUs 348 to ProdPlex most of 445 needed w/o CPC upgrade

QA Plex donated - was only supposed to have high share on Weekends – Automation to increase

z/OS LPARs Use of CPC

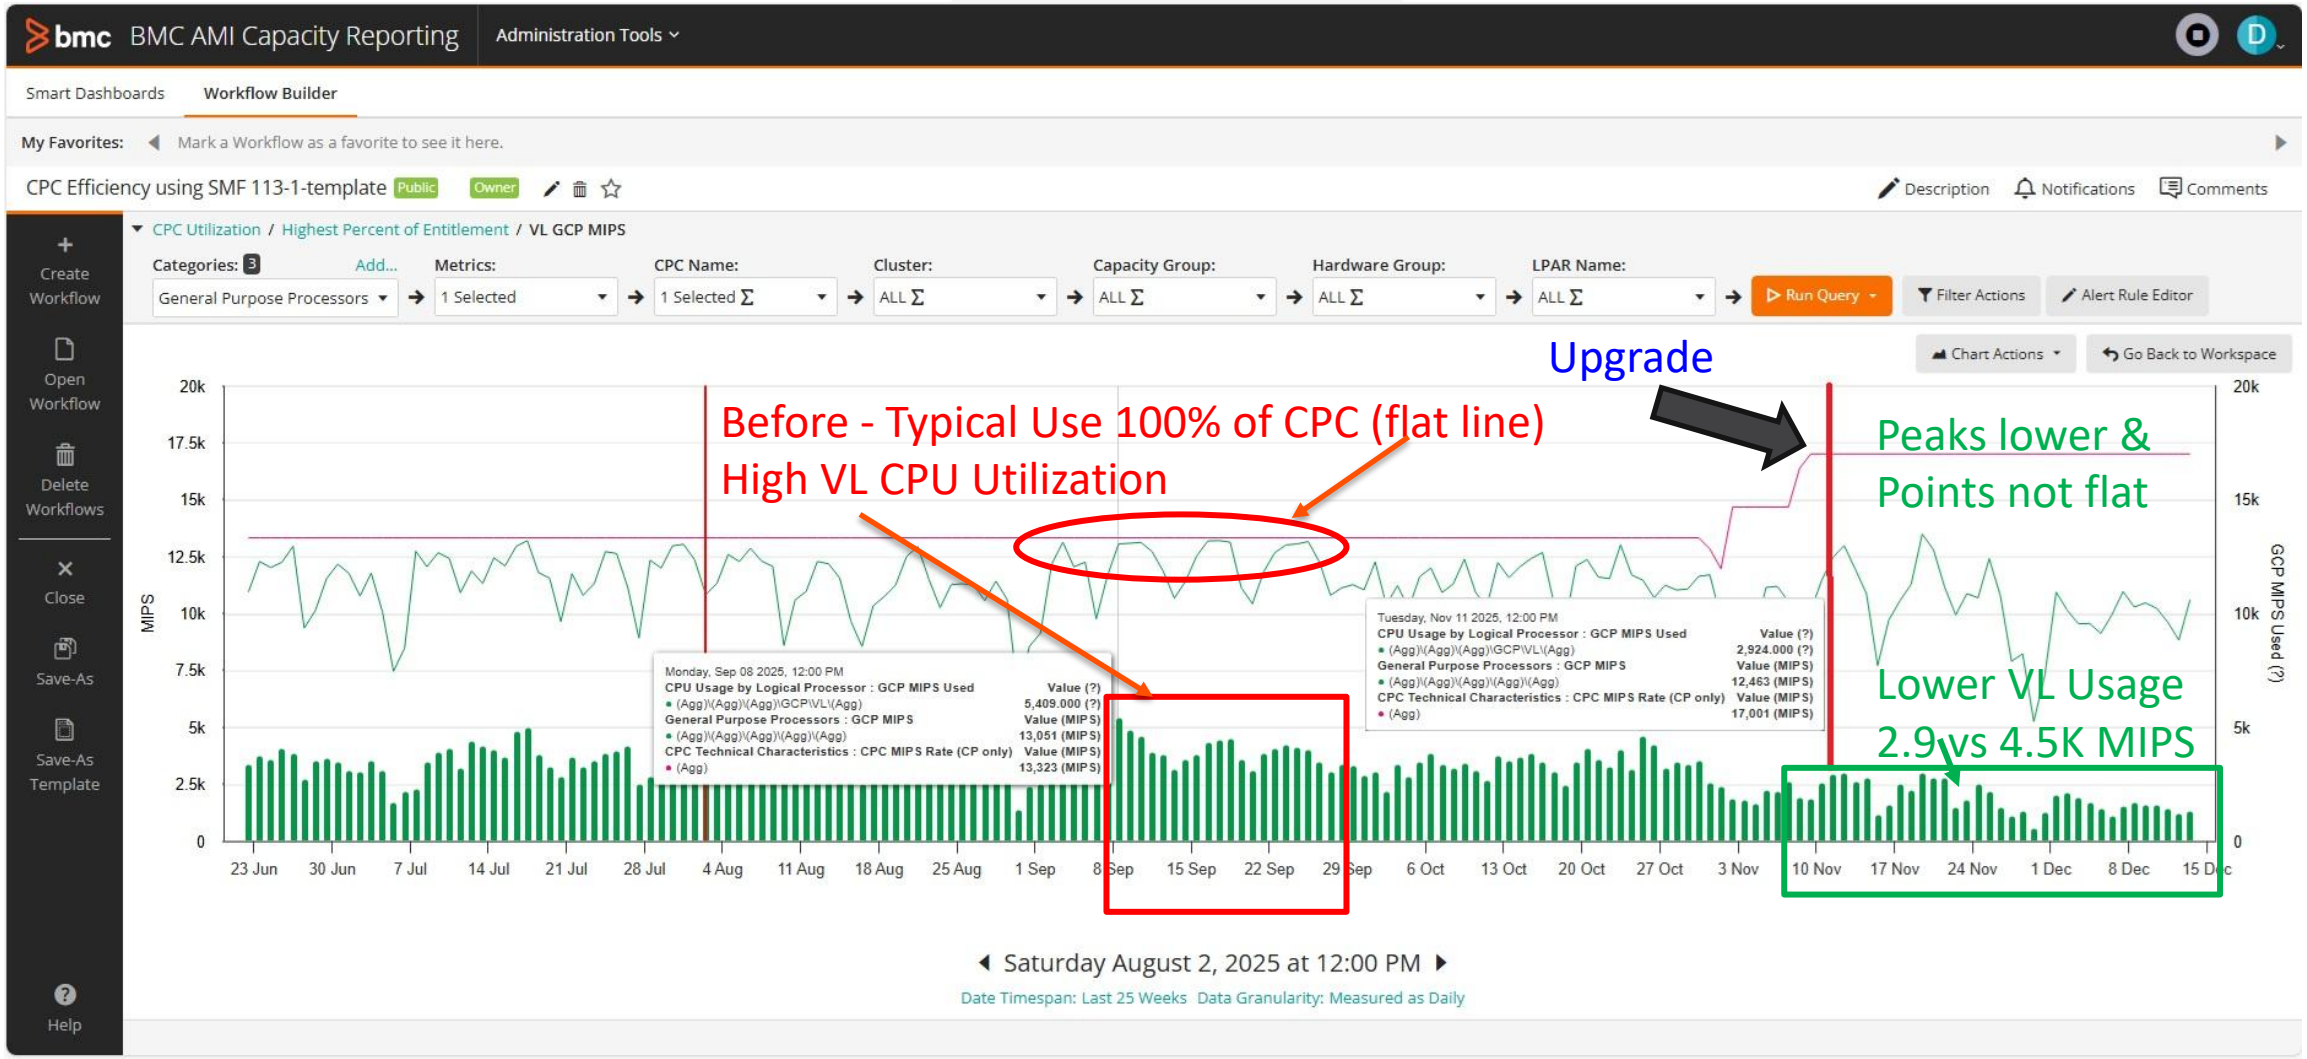

Efficiency Gains VH

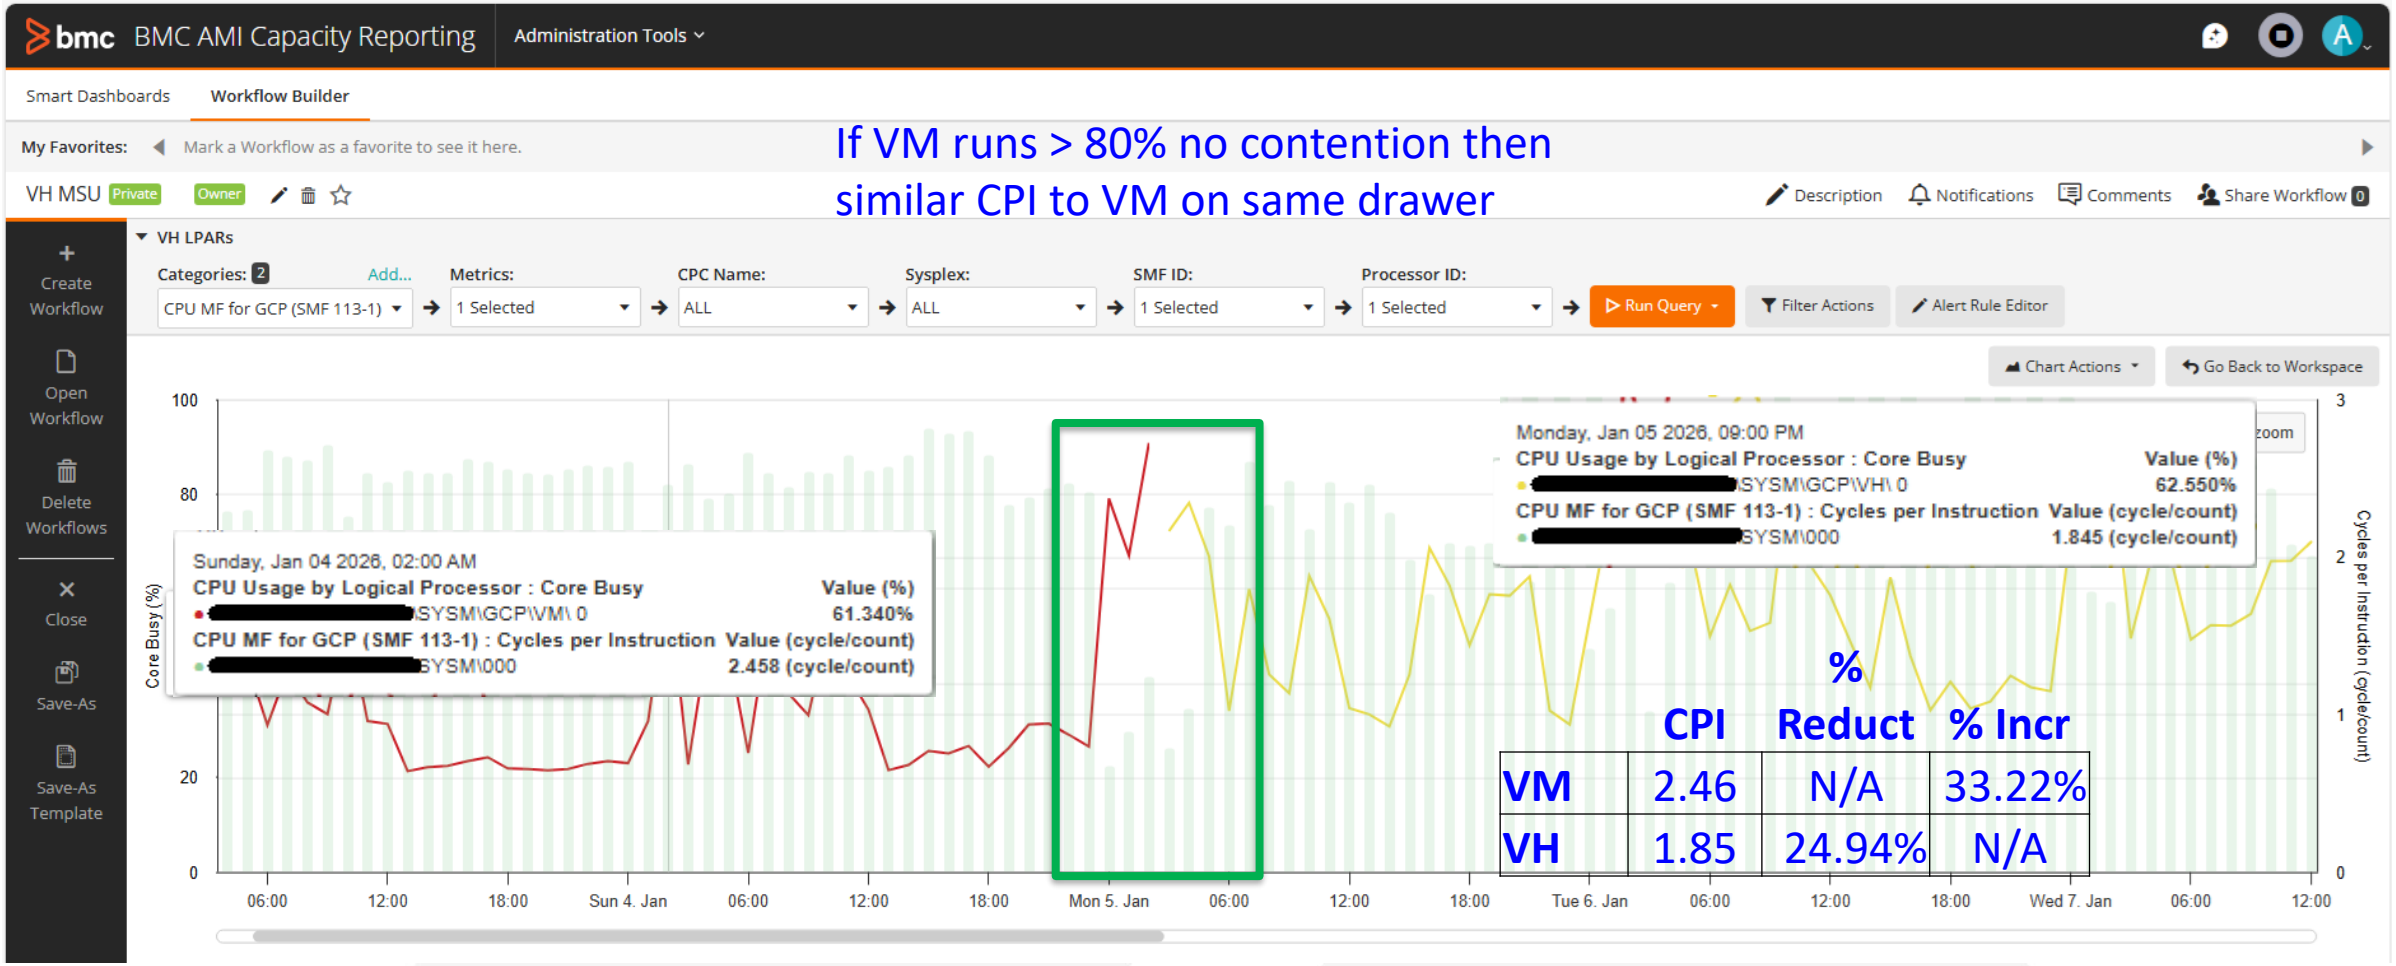

	CPI	% Reduct	% Incr
VM	2.46	N/A	33.22%
VH	1.85	24.94%	N/A

Cost Constraint

z16-617 Would have still fit in one Drawer

z16-529 similar MSU would not, require $$ 2nd Drawer

z16-526 Single drawer selected

- **Specialty engines reduced**
- **zIIP 8 reduced to 6 (25% less capacity)**
  - LPAR shares were adequate, need re-adjust as smaller total

Latent Demand - Upgrade

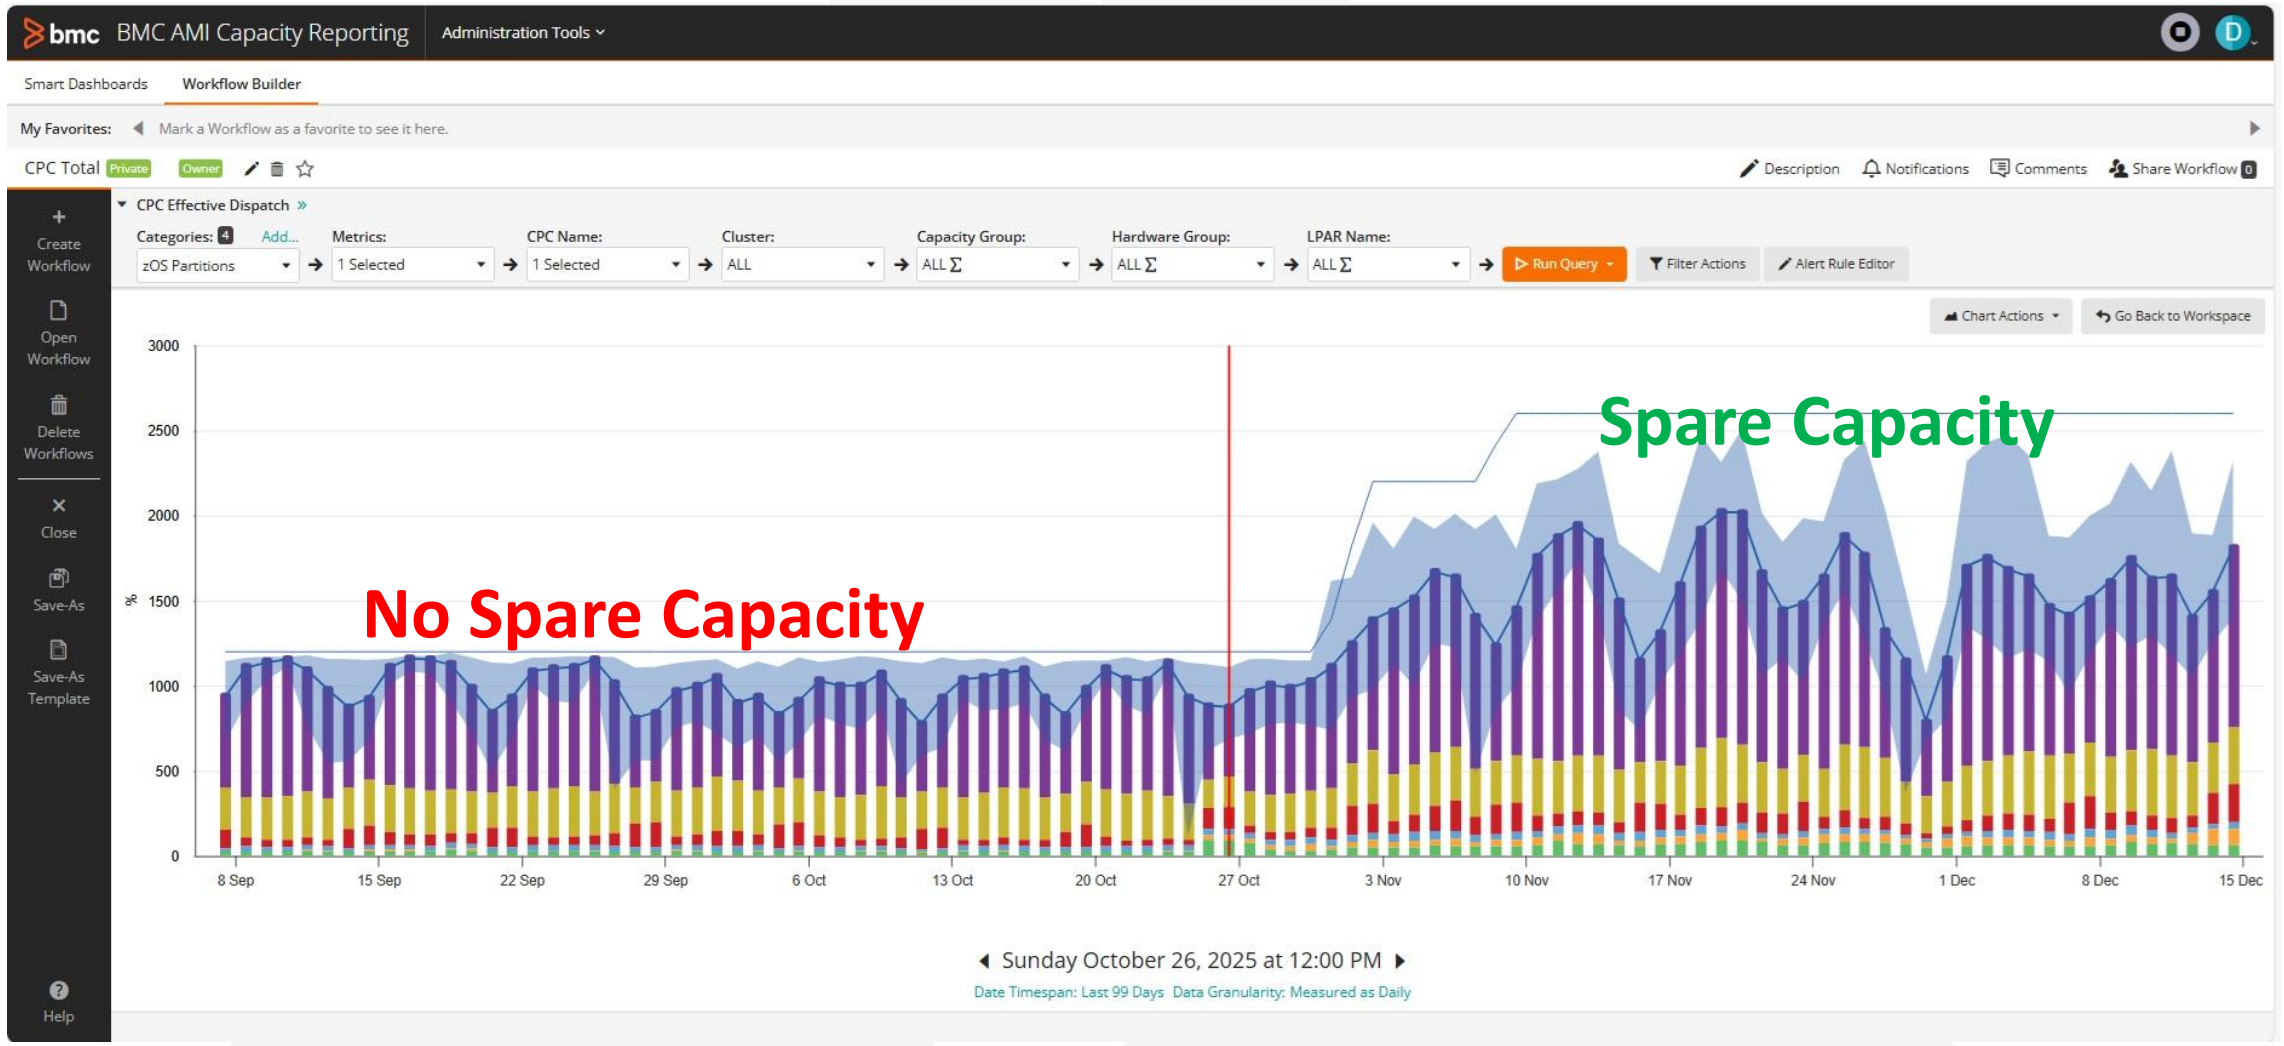

zIIP Capacity vs VL Usage

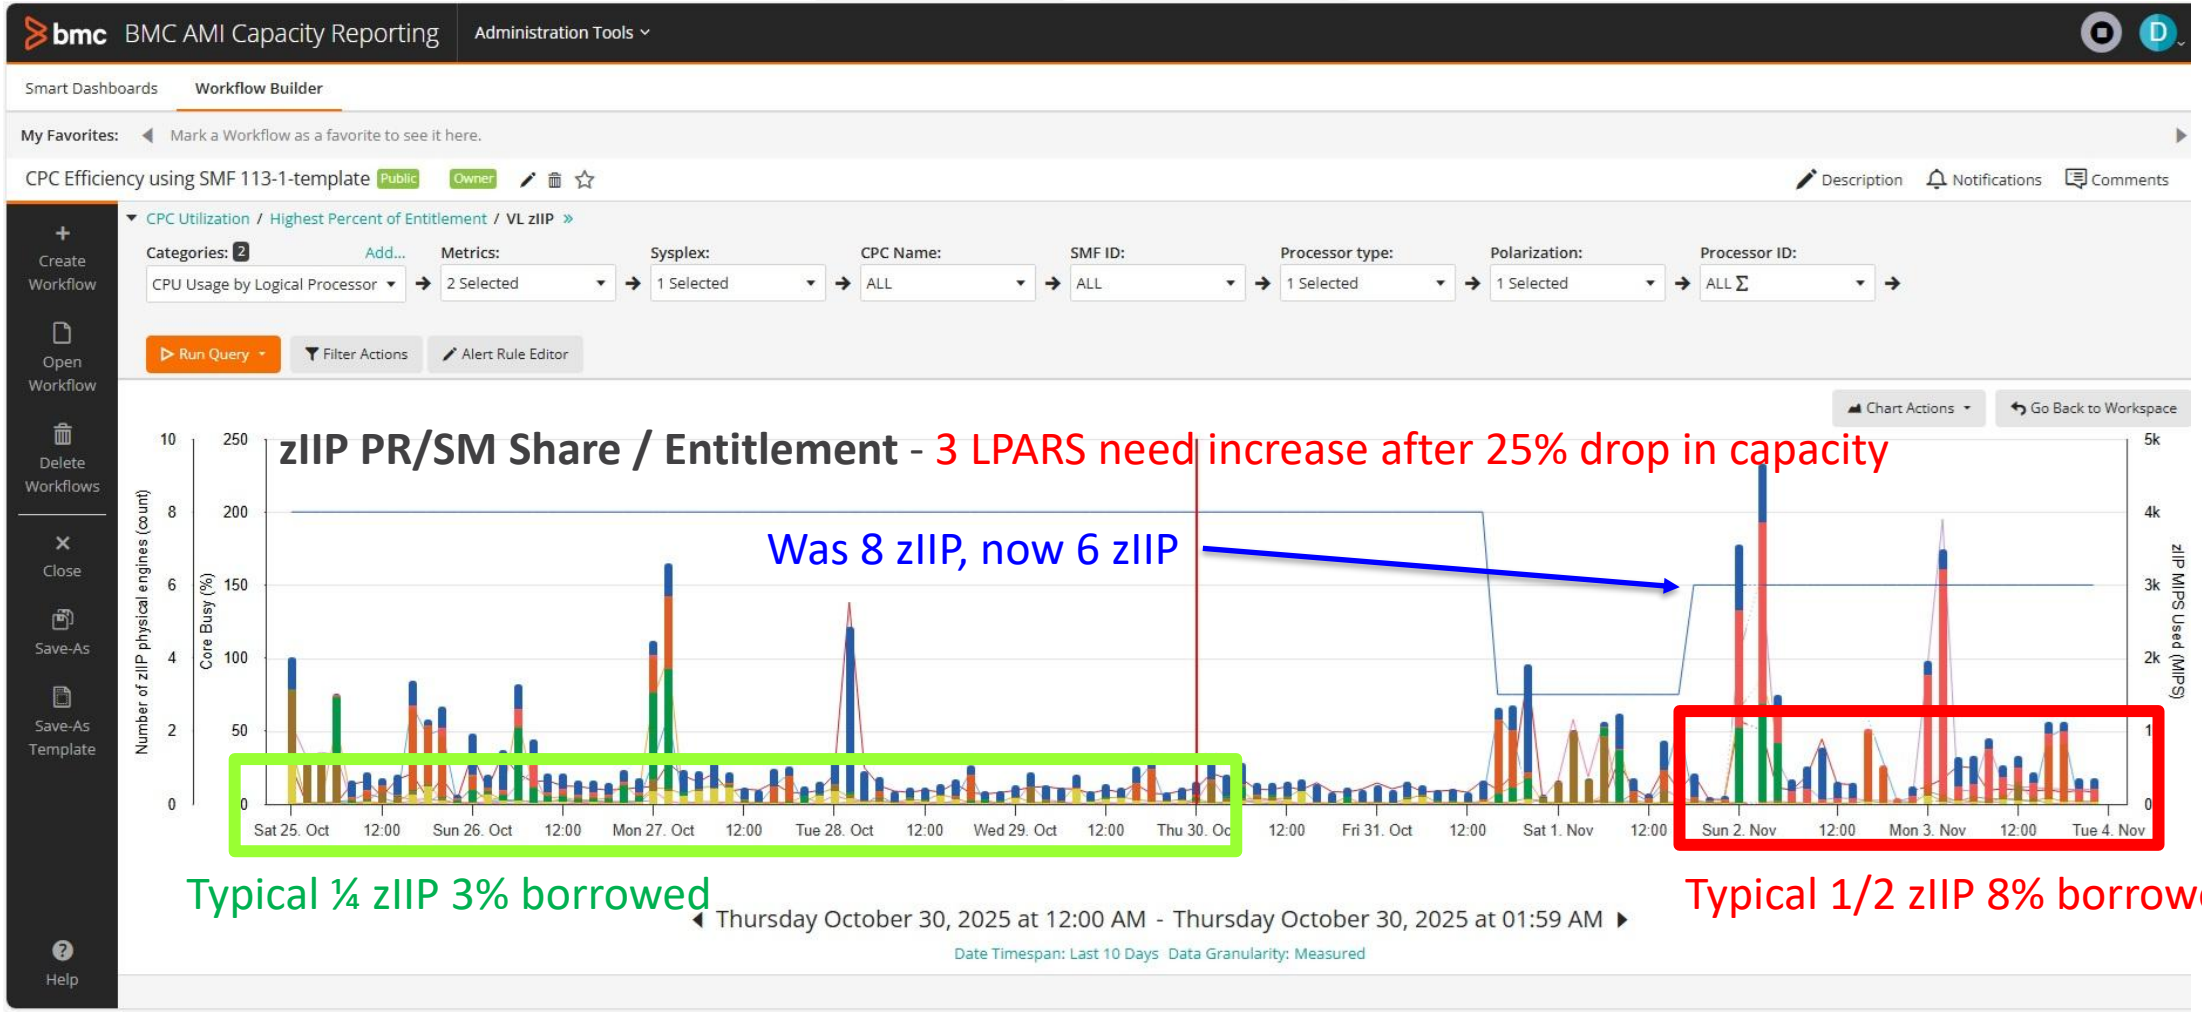

Z13+ Simultaneous MultiThreading (*SMT*)

zIIP 7## even slowed is faster than 5## GCP

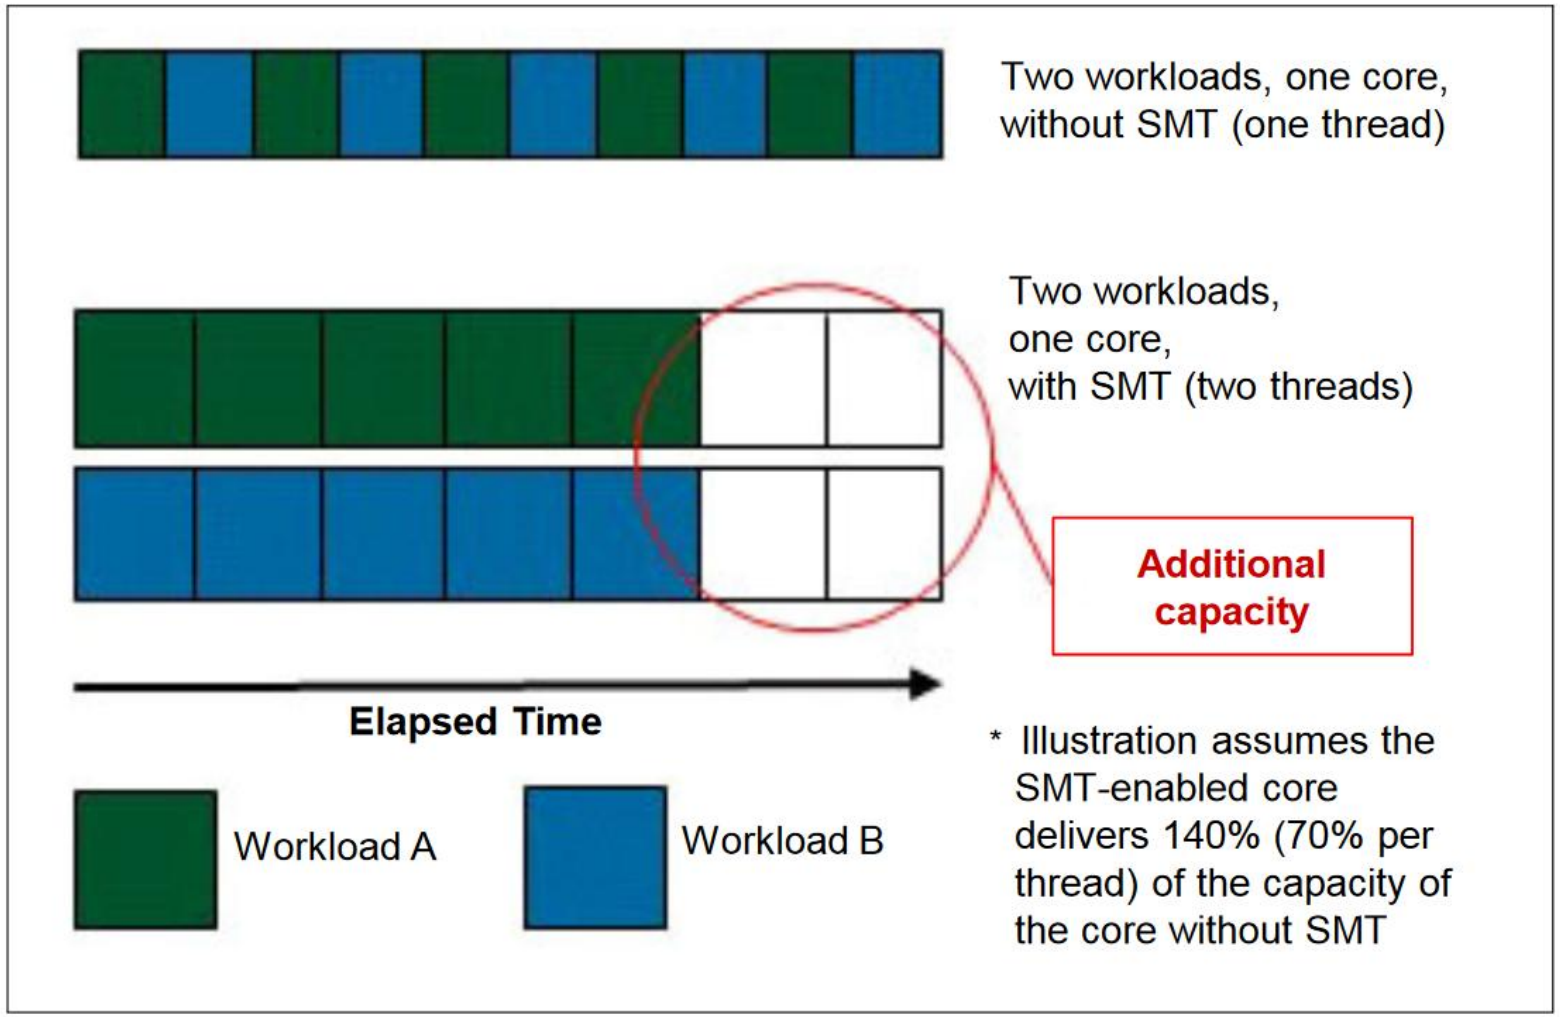

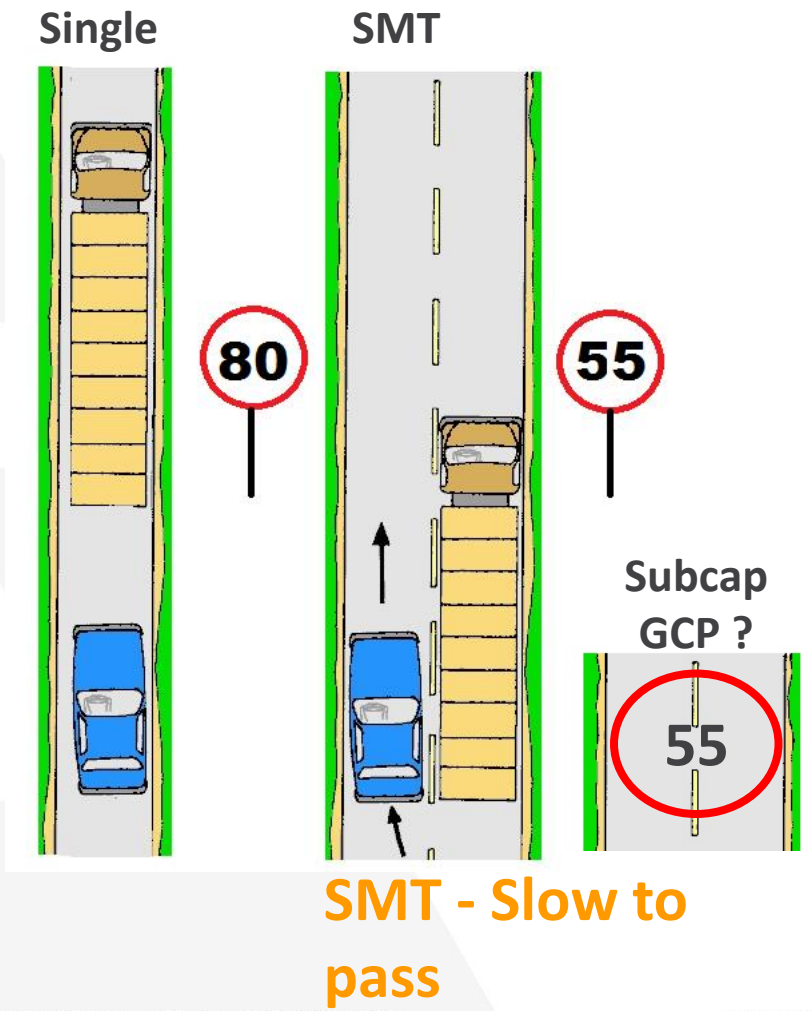

1 Address Loop Potential

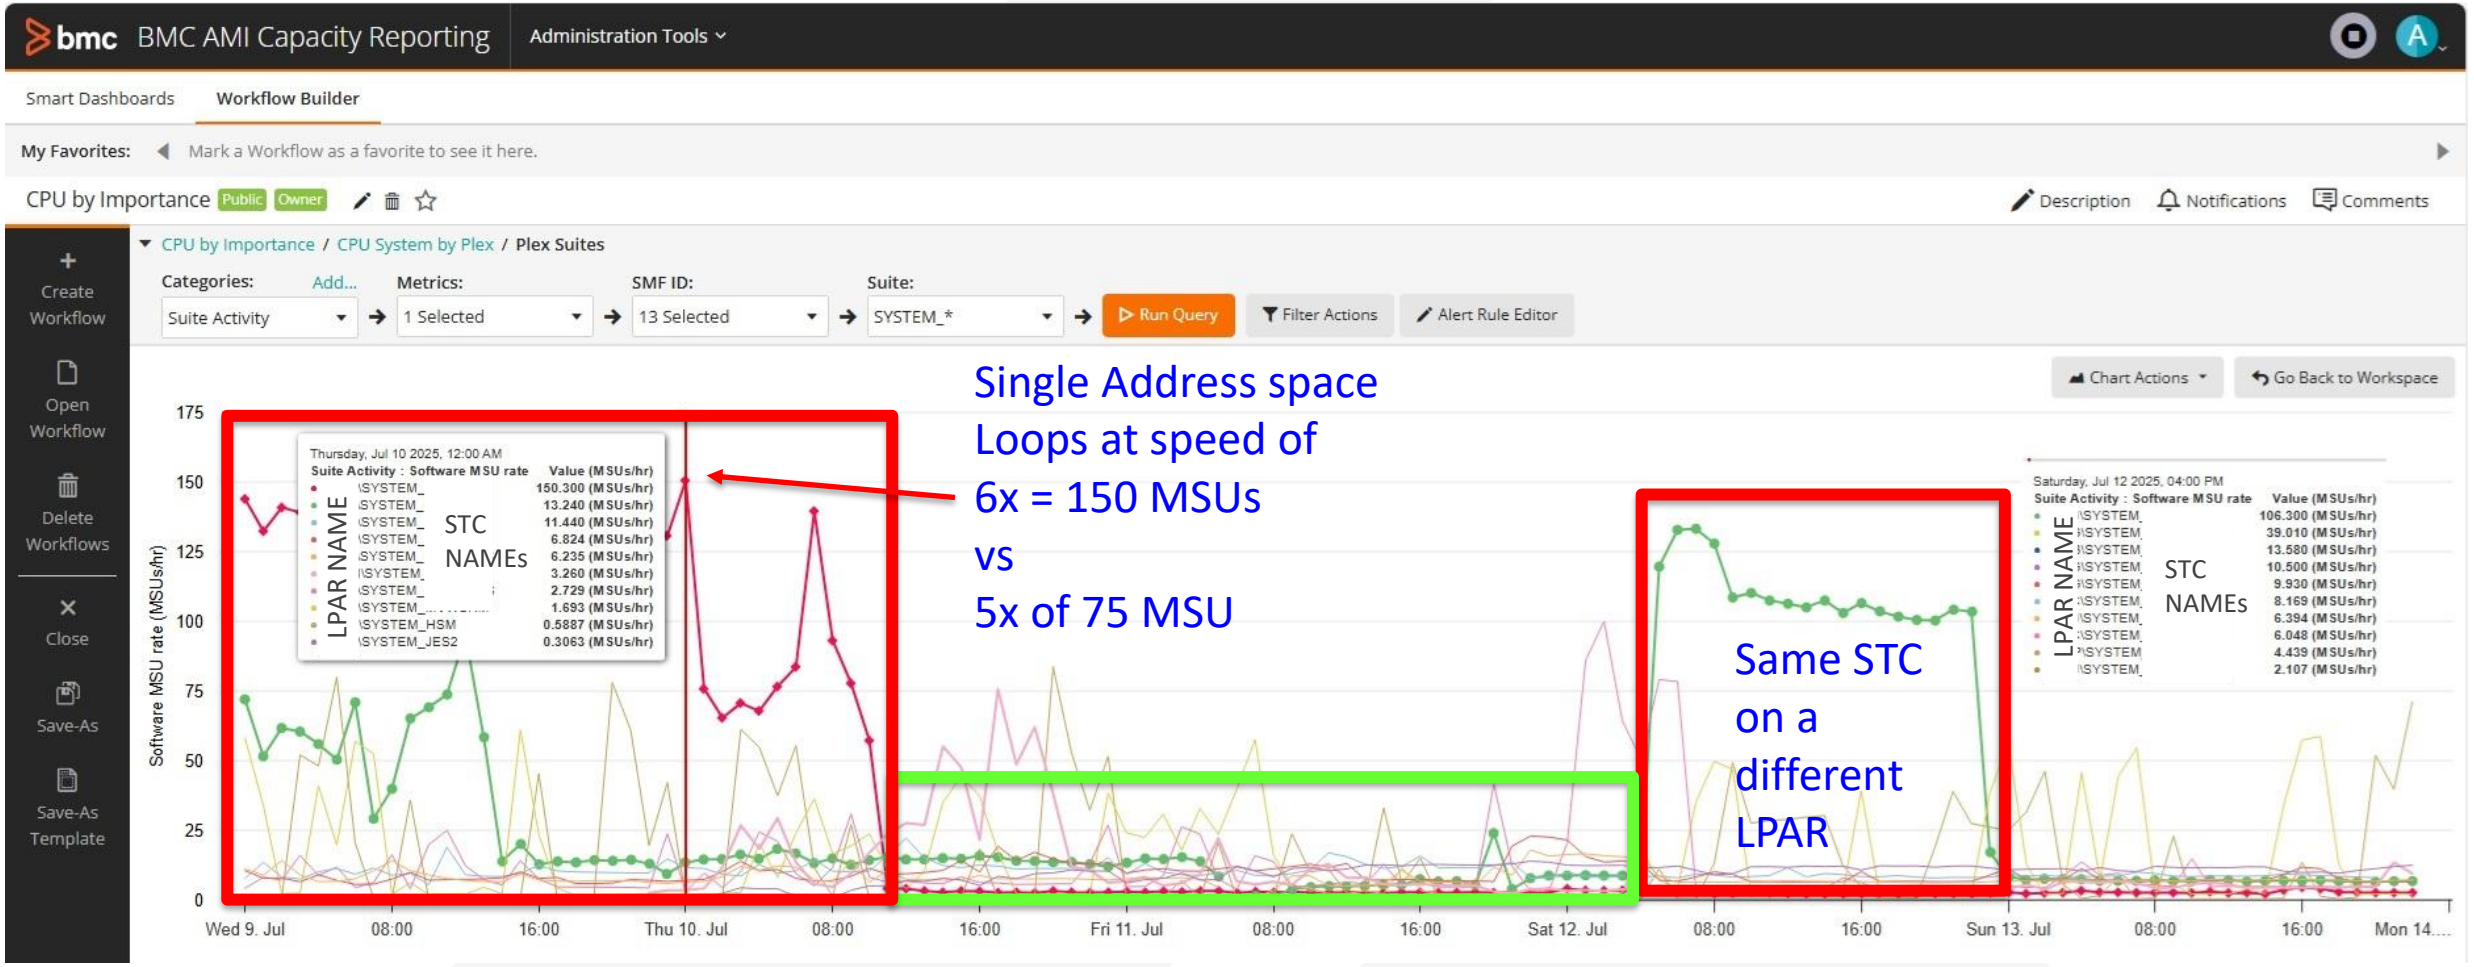

CICS Over Entitlement

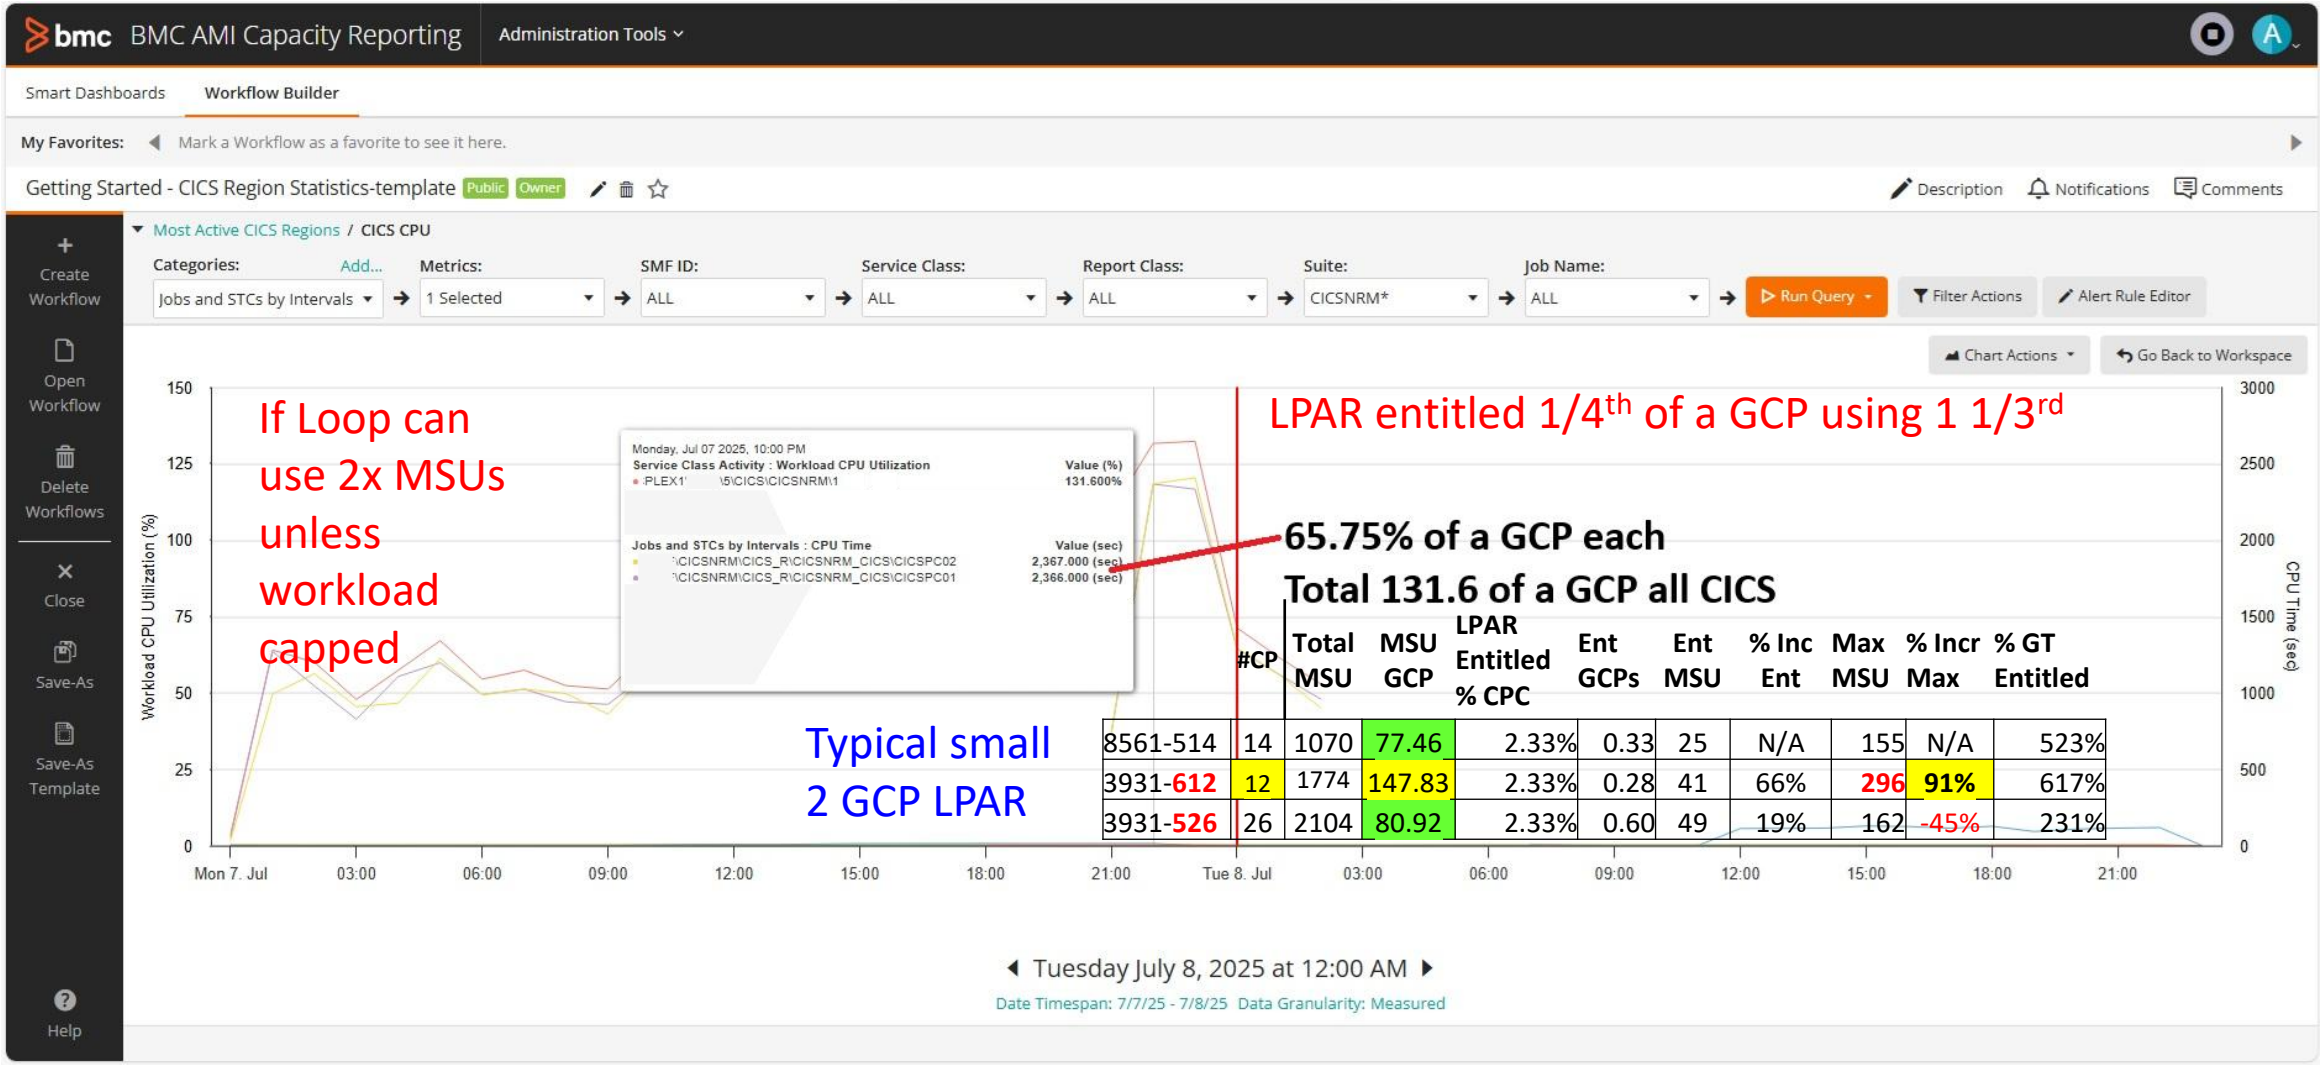

	#CP	Total MSU	MSU GCP	LPAR Entitled % CPC	Ent GCPs	Ent MSU	% Inc Ent	Max MSU	% Incr Max	% GT Entitled
8561-514	14	1070	77.46	2.33%	0.33	25	N/A	155	N/A	523%
3931-612	12	1774	147.83	2.33%	0.28	41	66%	296	91%	617%
3931-526	26	2104	80.92	2.33%	0.60	49	19%	162	-45%	231%

CICS suffer slower GCP?

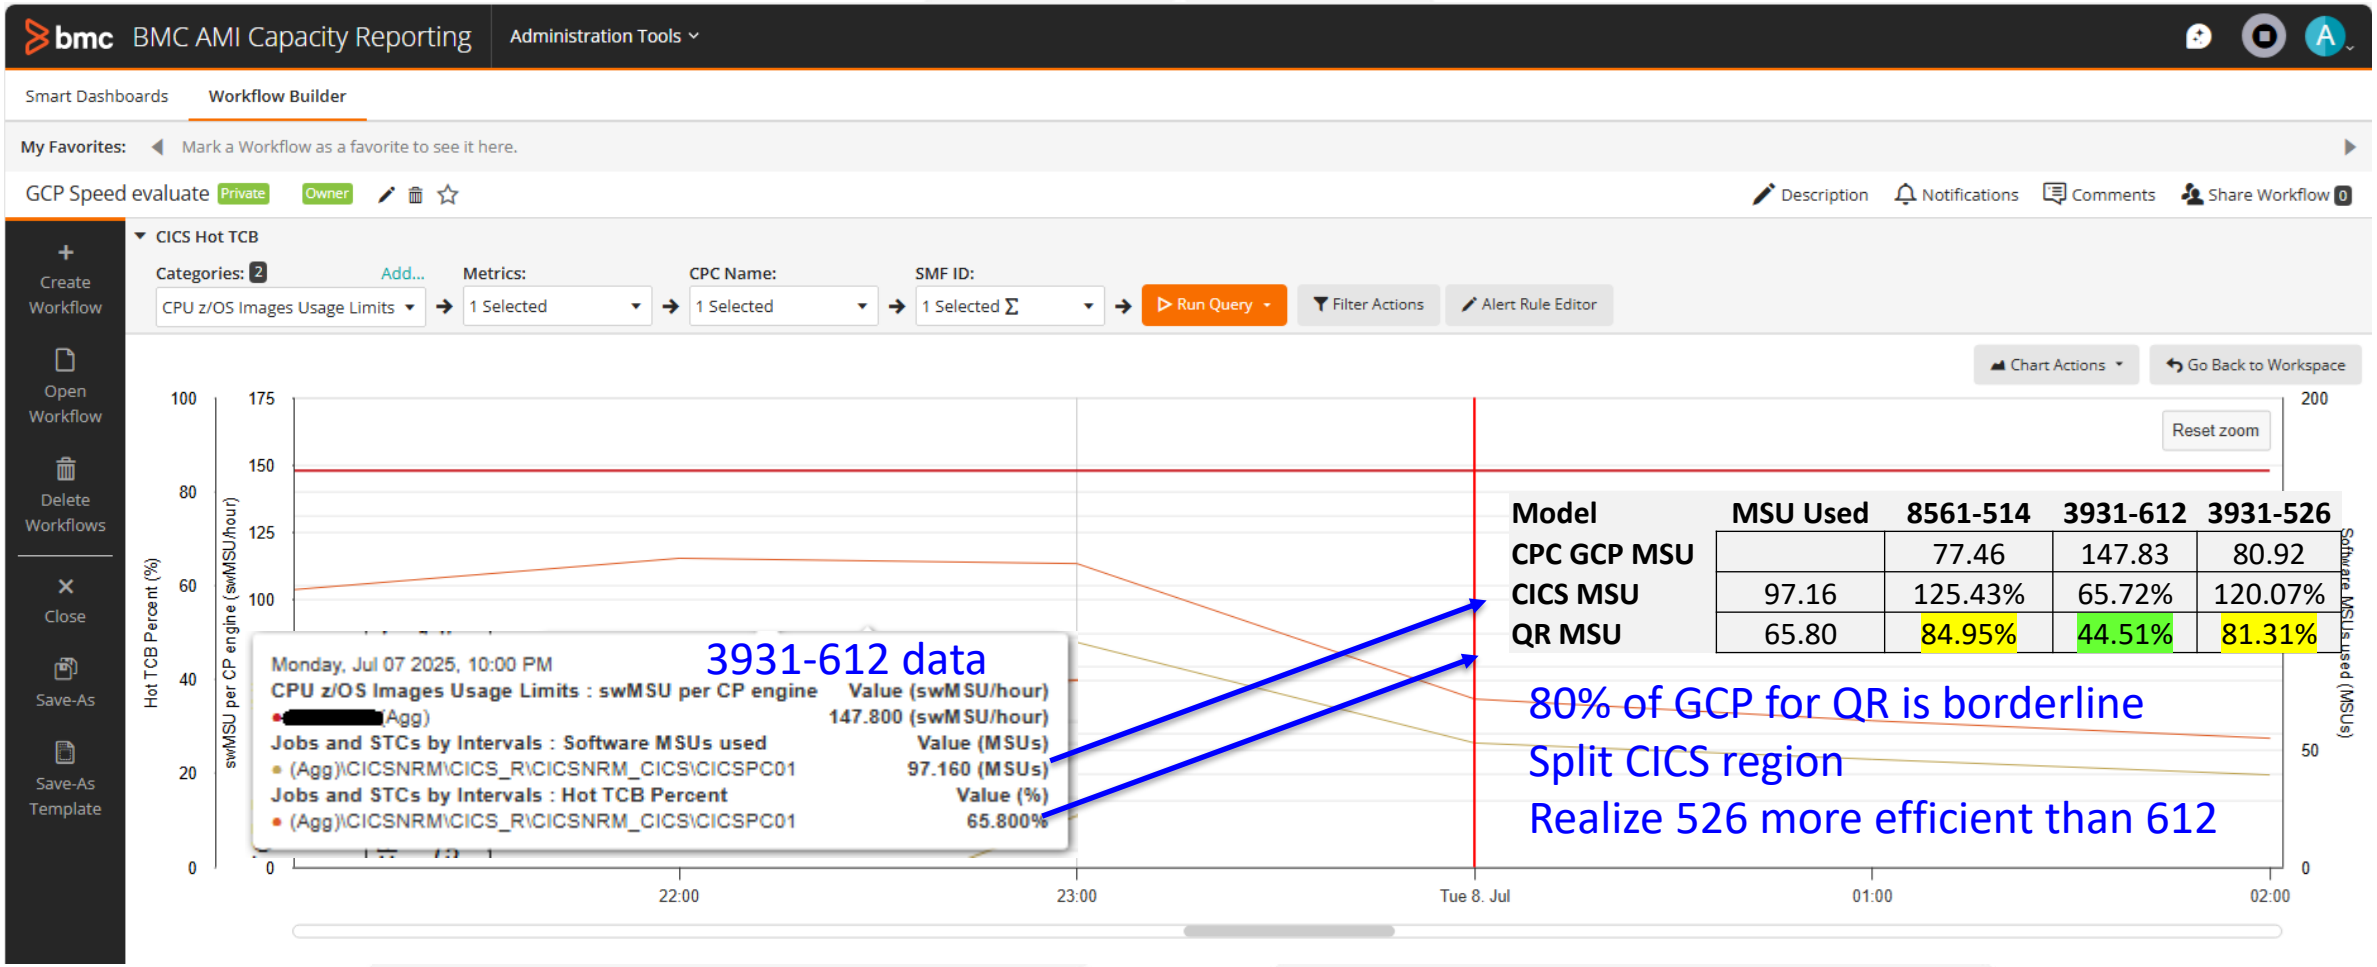

Model	MSU Used	8561-514	3931-612	3931-526
CPC GCP MSU		77.46	147.83	80.92
CICS MSU	97.16	125.43%	65.72%	120.07%
QR MSU	65.80	84.95%	44.51%	81.31%

80% of GCP for QR is borderline
Split CICS region
Realize 526 more efficient than 612

More Info

Course Code:	Area / Course	Days
PT30411	Mainframe Capacity Planning	2
PT3997	Understanding Workload Manager (WLM)	3
PT31682	z/OS Performance Tuning and Control with WLM	5

SHARE26 ORLANDO, FL 2026

SHARE Technology Expo
ProTech in Booth 120 (near IBM)
Mon & Tue 5-7pm
Wed 11am – 1pm

Cheryl Watson's Tuning Letter - (Aug 5th, 2025)
•**Paid Subscription** Required to access
https://tuningletter.planetmainframe.com/table-of-contents/2025-tuning-letters/

Planet Mainframe Republish (Oct 17, 2025)
https://planetmainframe.com/2025/10/dont-buy-more-mips-do-this-instead/

More info from IBM

Best Practice document - defining LPAR logical CPs and zIIPs

- URL: www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD106388

WLM LPAR Design Tool – for HiperDispatch

- **URL: https://ibm.biz/BdZTVw**

Danke

Grazie

Gracias

Merci

Tak

Dank u

Obrigado

謝謝
xièxie

감사합니다
gamsahabnida

有り難う
Arigatō

תודה רבה
Toda rava

Thank You

Ačiū

Dziękuję

Спасибо
Spasibo