Tailored Fit Pricing: How To Manage Workloads in a World Without Capping

Nathan Brice
IBM

November 2019
Session OO

Important Disclaimer

IBM’s statements regarding its plans, directions and intent are subject to change or withdrawal without notice at IBM’s sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

The information in this document is provided “as is” without any warranty, express or implied, including without any warranties of merchantability, fitness for a particular purpose and any warranty or condition of non-infringement. IBM products are warranted according to the terms and conditions of the agreements under which they are provided. It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

- Performance. Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user’s job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.
- Customer Examples. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.
- Availability. Not all offerings are available in every country in which IBM operates. This document is current as of the initial date of publication and may be changed by IBM at any time.
- Trademarks. IBM and the IBM logo are trademarks of International Business Machines Corporation, registered in many jurisdictions. Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates. Other company, product and service names may be trademarks, registered marks or service marks of their respective owners.

Highlights

- Tailored Fit Pricing is a revolutionary new pricing model that eliminates the need for workload capping and provides a complete alternative to the Rolling 4 Hour Average
- Transitioning to a model where measurement of overall MSU consumption is critical represents more than just a change in how you pay for software on IBM Z. Several stakeholders are affected and various considerations around how workload is managed must be made
- Effective use of tooling can help you: IBM Z Decision Support for Capacity Planning delivers pre-defined dashboards to give visibility in to current consumption levels, forecast future consumption and provide insight into where resource optimization can take place

What is **Tailored Fit Pricing**?

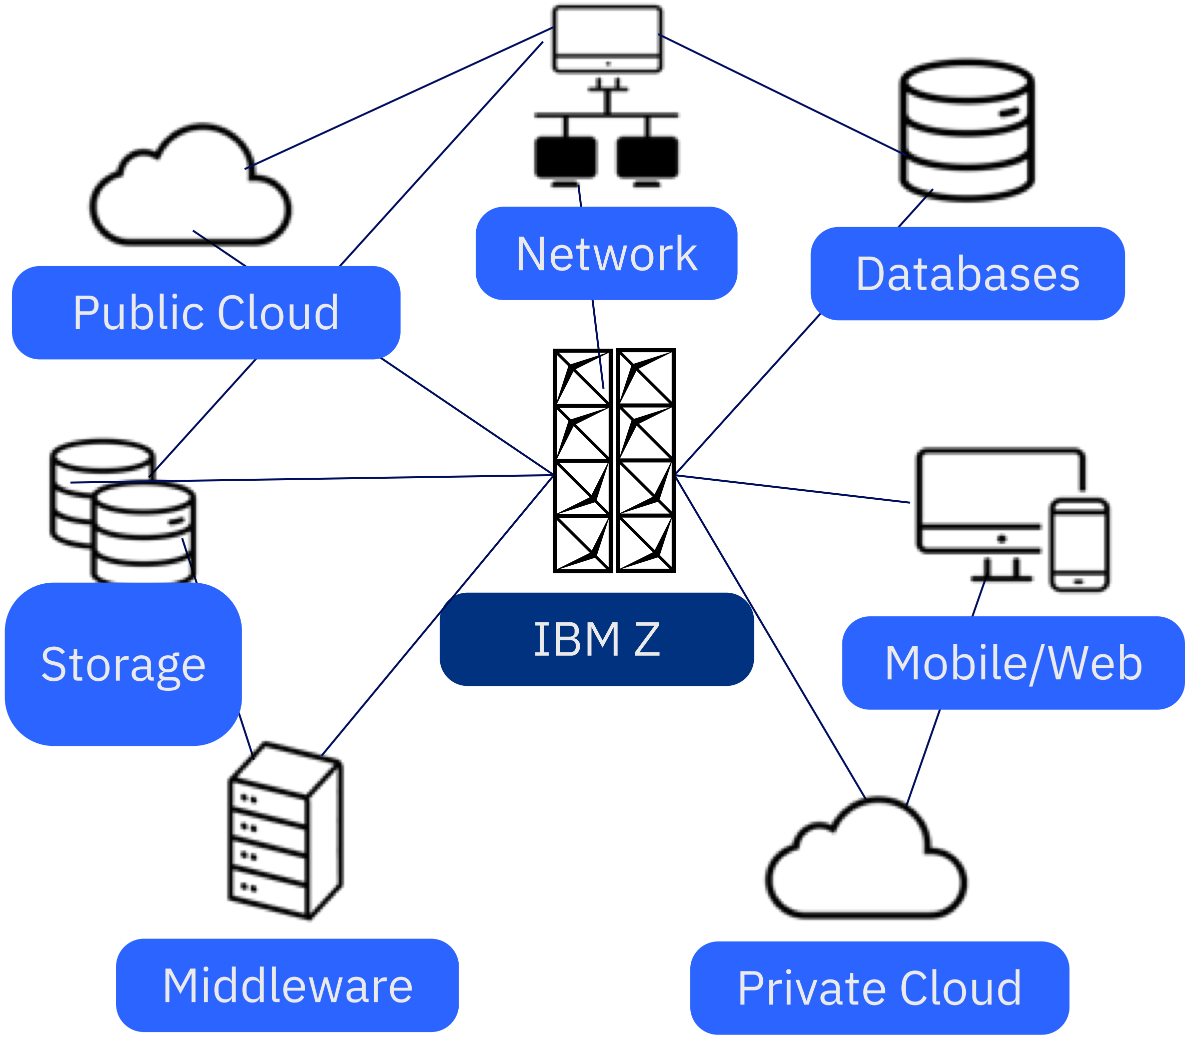

Digital Transformation is leading to Hybrid Cloud pattern of deployment across enterprises

IBM Z is a critical part of this infrastructure

Unpredictable demand in era of Hybrid Cloud

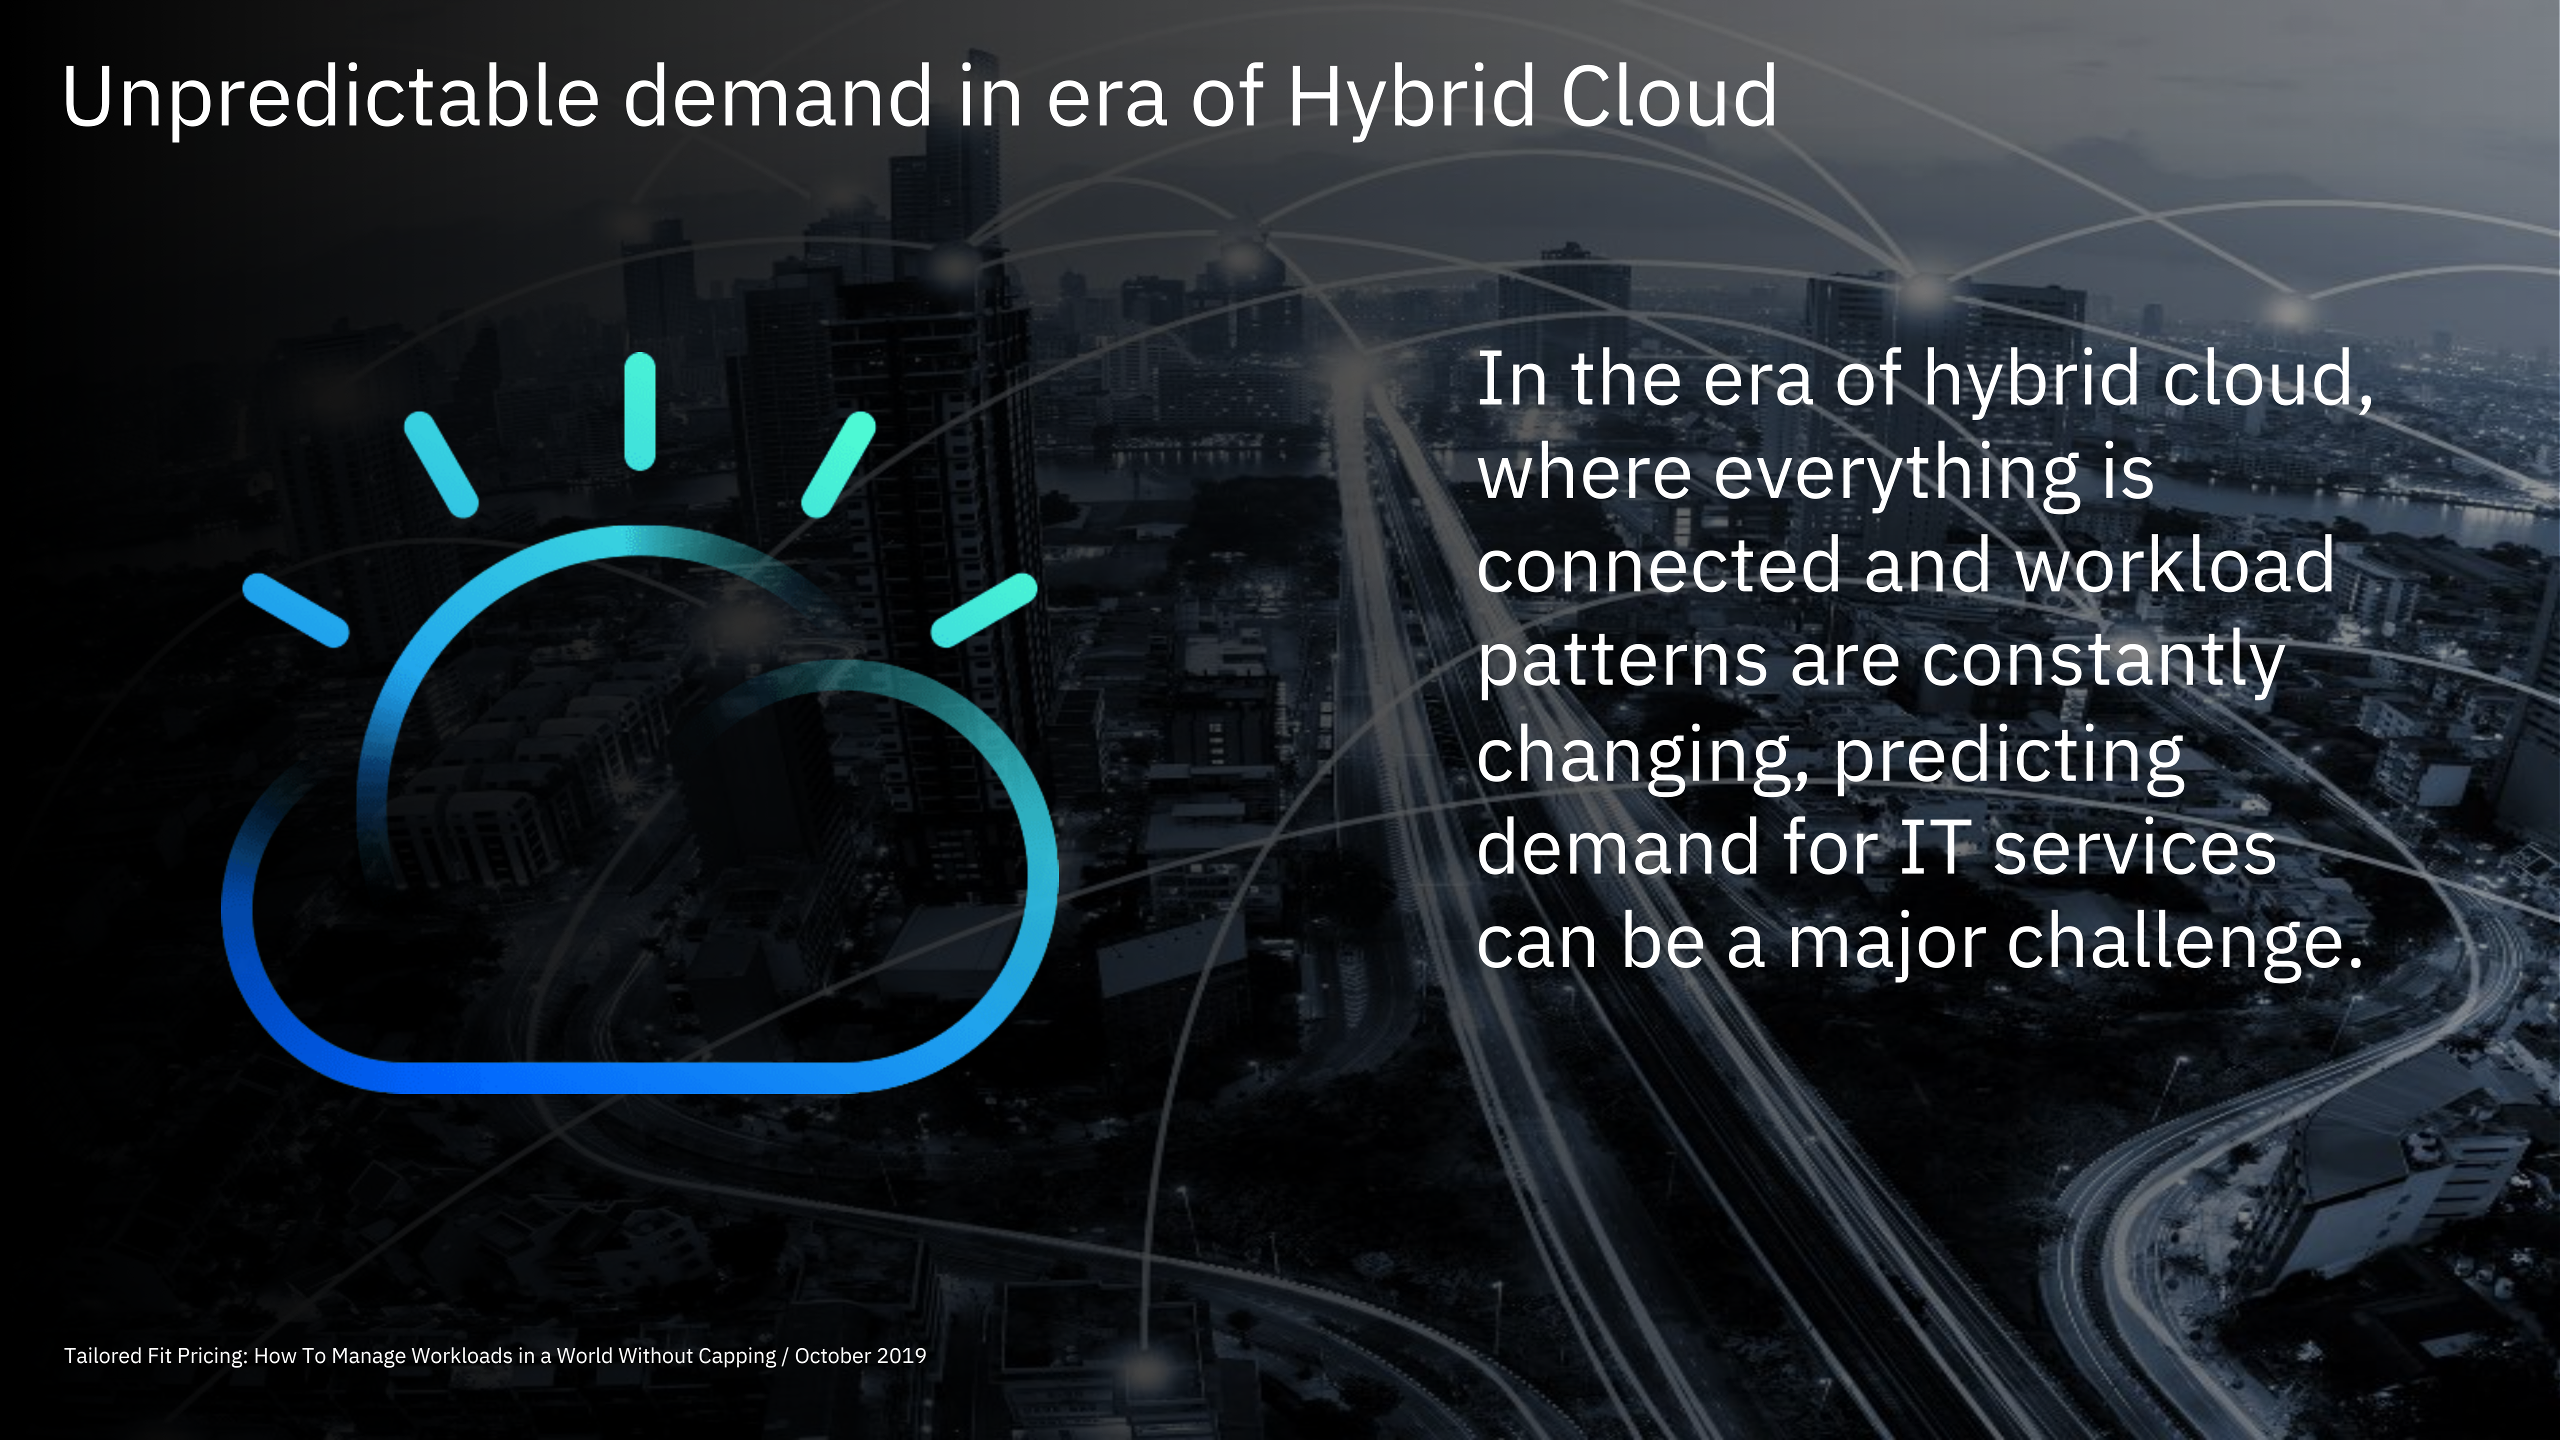

In the era of hybrid cloud, where everything is connected and workload patterns are constantly changing, predicting demand for IT services can be a major challenge.

Evolution of Z Software Pricing

1970 - 1999

PAST

Full Capacity

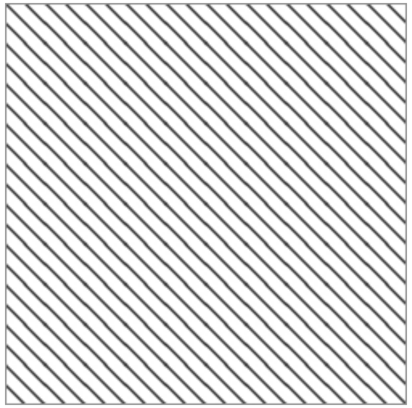

- Simple way to charge for z/OS-based software

1999 - 2019

PRESENT

Sub-Capacity

(R4HA)

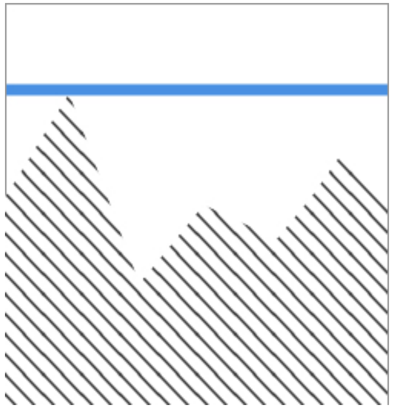

- Modelled on 90% utilization
- As system size increases, align product value to less than full capacity

2019 onwards

FUTURE

Tailored Fit Pricing

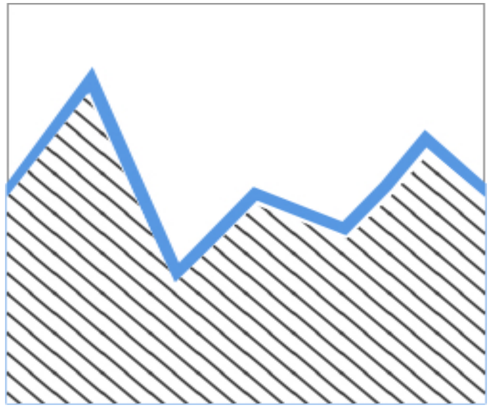

- Remove R4HA
- Align value to the workload for the amount of system resources it consumes

Models of Tailored Fit Pricing

Enterprise Capacity

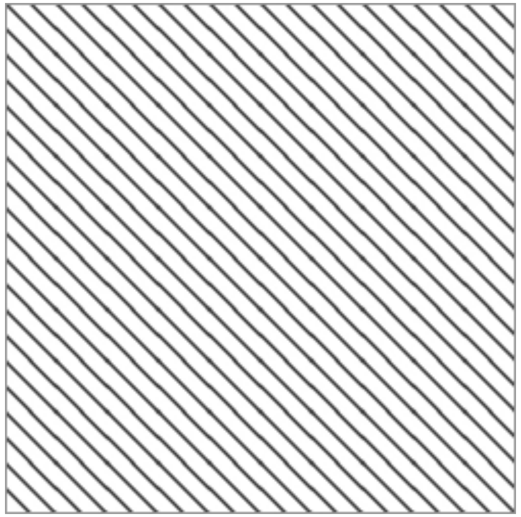

Enterprise Consumption

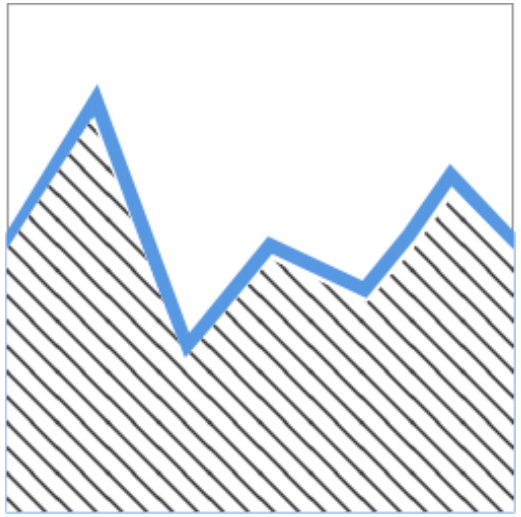

An MSU consumption pricing model that allows clients to:

- Take full advantage of the hardware they own
- Peak and spike without ‘penalty’
- Smooth seasonal variations over the entire year
- Grow at a highly competitive per MSU price
- Pay for workloads with price consistency

We’re going to focus more on the Enterprise Consumption model today

Overall Usage Determines Charges, Not Peaks

Peak and spike as the business demands…without blowing the budget

- In both workloads below, the total MSUs consumed is 10,000 over the same period of time
- Under R4HA, the single spike of 2,400 MSUs sets the price for the entire month
- **Under consumption, you pay for exactly what you use at the same rate**

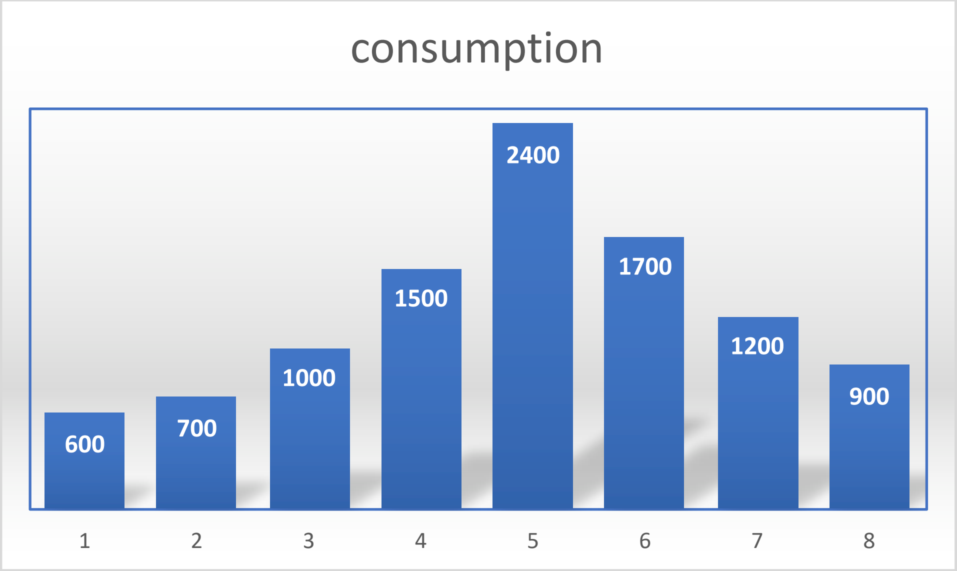

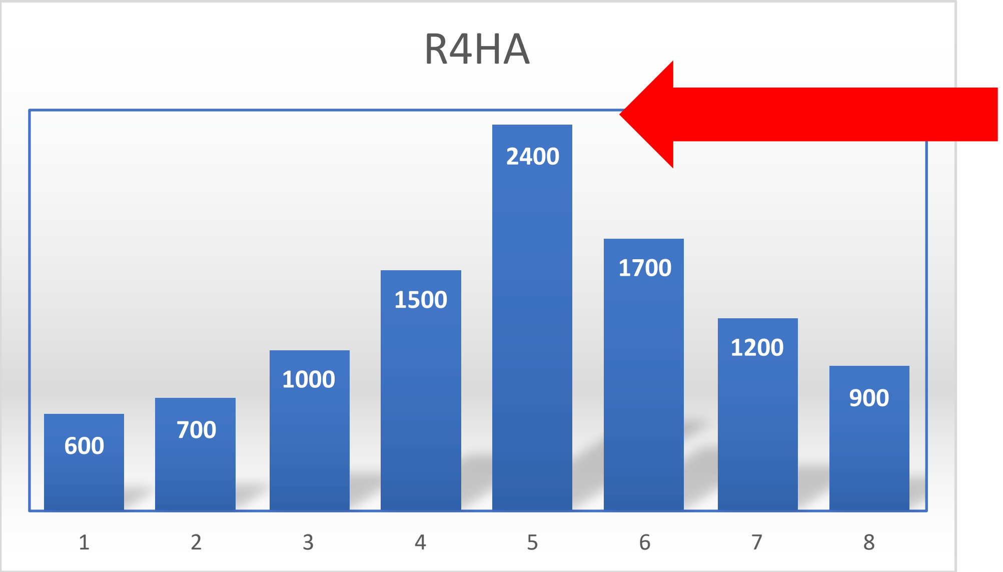

This determined the price under R4HA

What exactly is MSU consumption?

Capacity reference (capacity markers to measure entitlement)

- **Full-cap** = size of the machine, based on HW MSU ratings; Sub Capacity Reporting Tool (SCRT) not required
- **Sub-cap** = combination of R4HA peak, LPARs where run, metric type, etc., as reported by E5/B5

Consumption

- **Total MSUs** as reported by K5 or N7 sections

In SCRT:

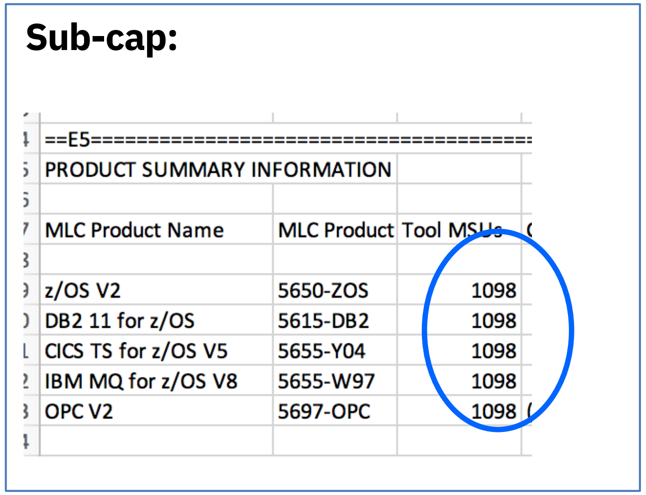

Sub-cap:

==E5=======================================		
PRODUCT SUMMARY INFORMATION		
MLC Product Name	MLC Product	Tool MSUs
z/OS V2	5650-ZOS	1098
DB2 11 for z/OS	5615-DB2	1098
CICS TS for z/OS V5	5655-Y04	1098
IBM MQ for z/OS V8	5655-W97	1098
OPC V2	5697-OPC	1098

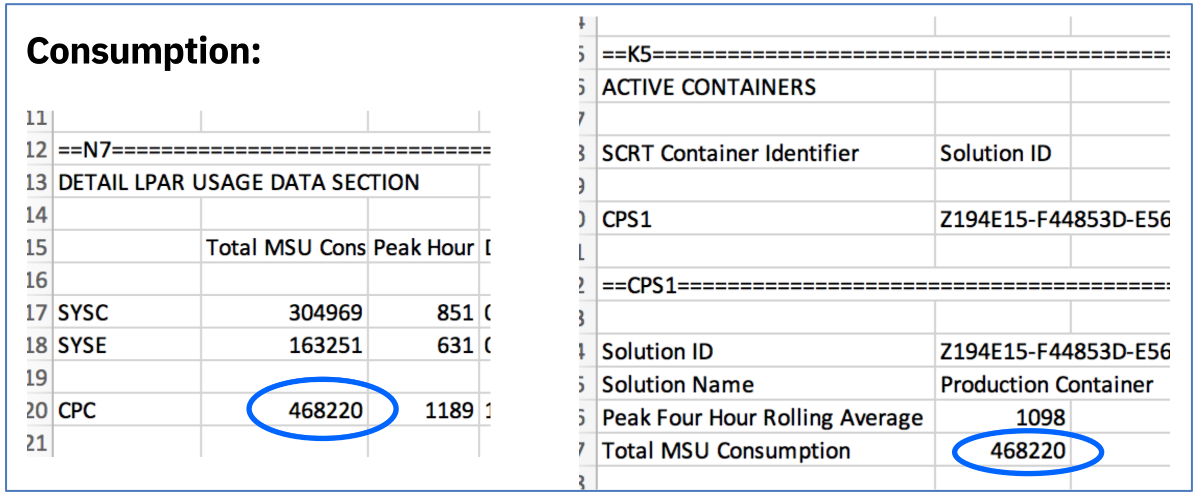

Consumption:

==N7================================		
DETAIL LPAR USAGE DATA SECTION		
	Total MSU Cons	Peak Hour
SYSC	304969	851
SYSE	163251	631
CPC	468220	1189

==K5===	
ACTIVE CONTAINERS	
SCRT Container Identifier	Solution ID
CPS1	Z194E15-F44853D-E56
==CPS1===	
Solution ID	Z194E15-F44853D-E56
Solution Name	Production Container
Peak Four Hour Rolling Average	1098
Total MSU Consumption	468220

Defining the MSU Baseline

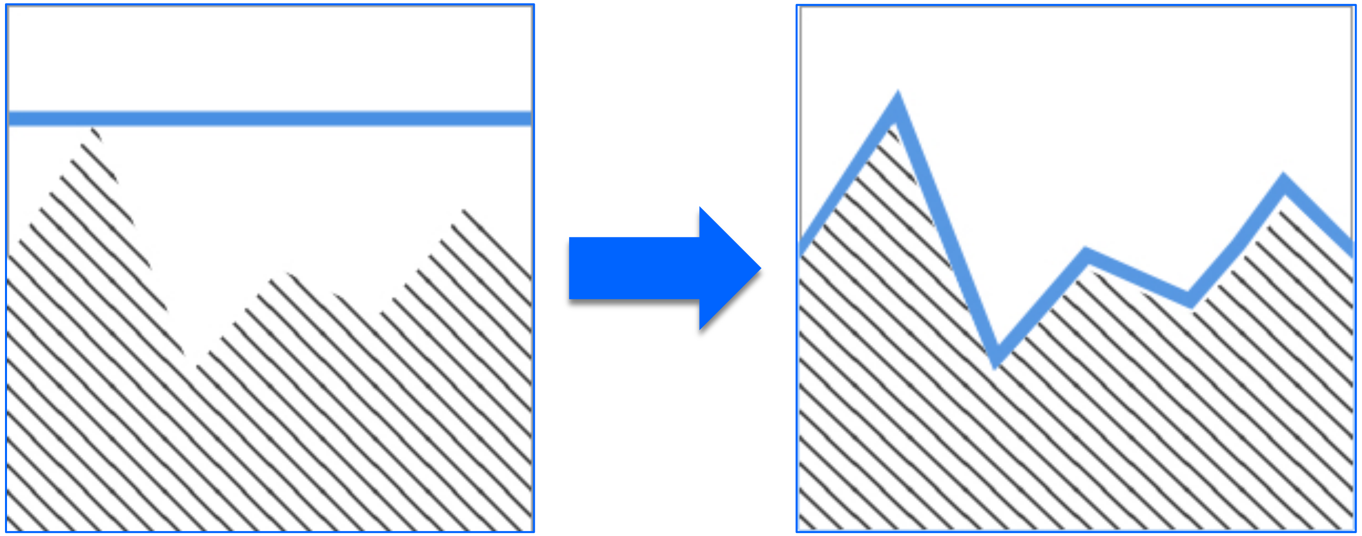

The baseline should account for your previous MLC and IPLA usage. Growth beyond that is calculated at a reduced rate

Baseline established upon 12 months of SCRT reporting

This should allow seasonal variations to be accounted for plus considerations for growth

Therefore important to understand what workloads you have running within your environment to start and allocate to the container

Typical Enterprise Consumption Solution

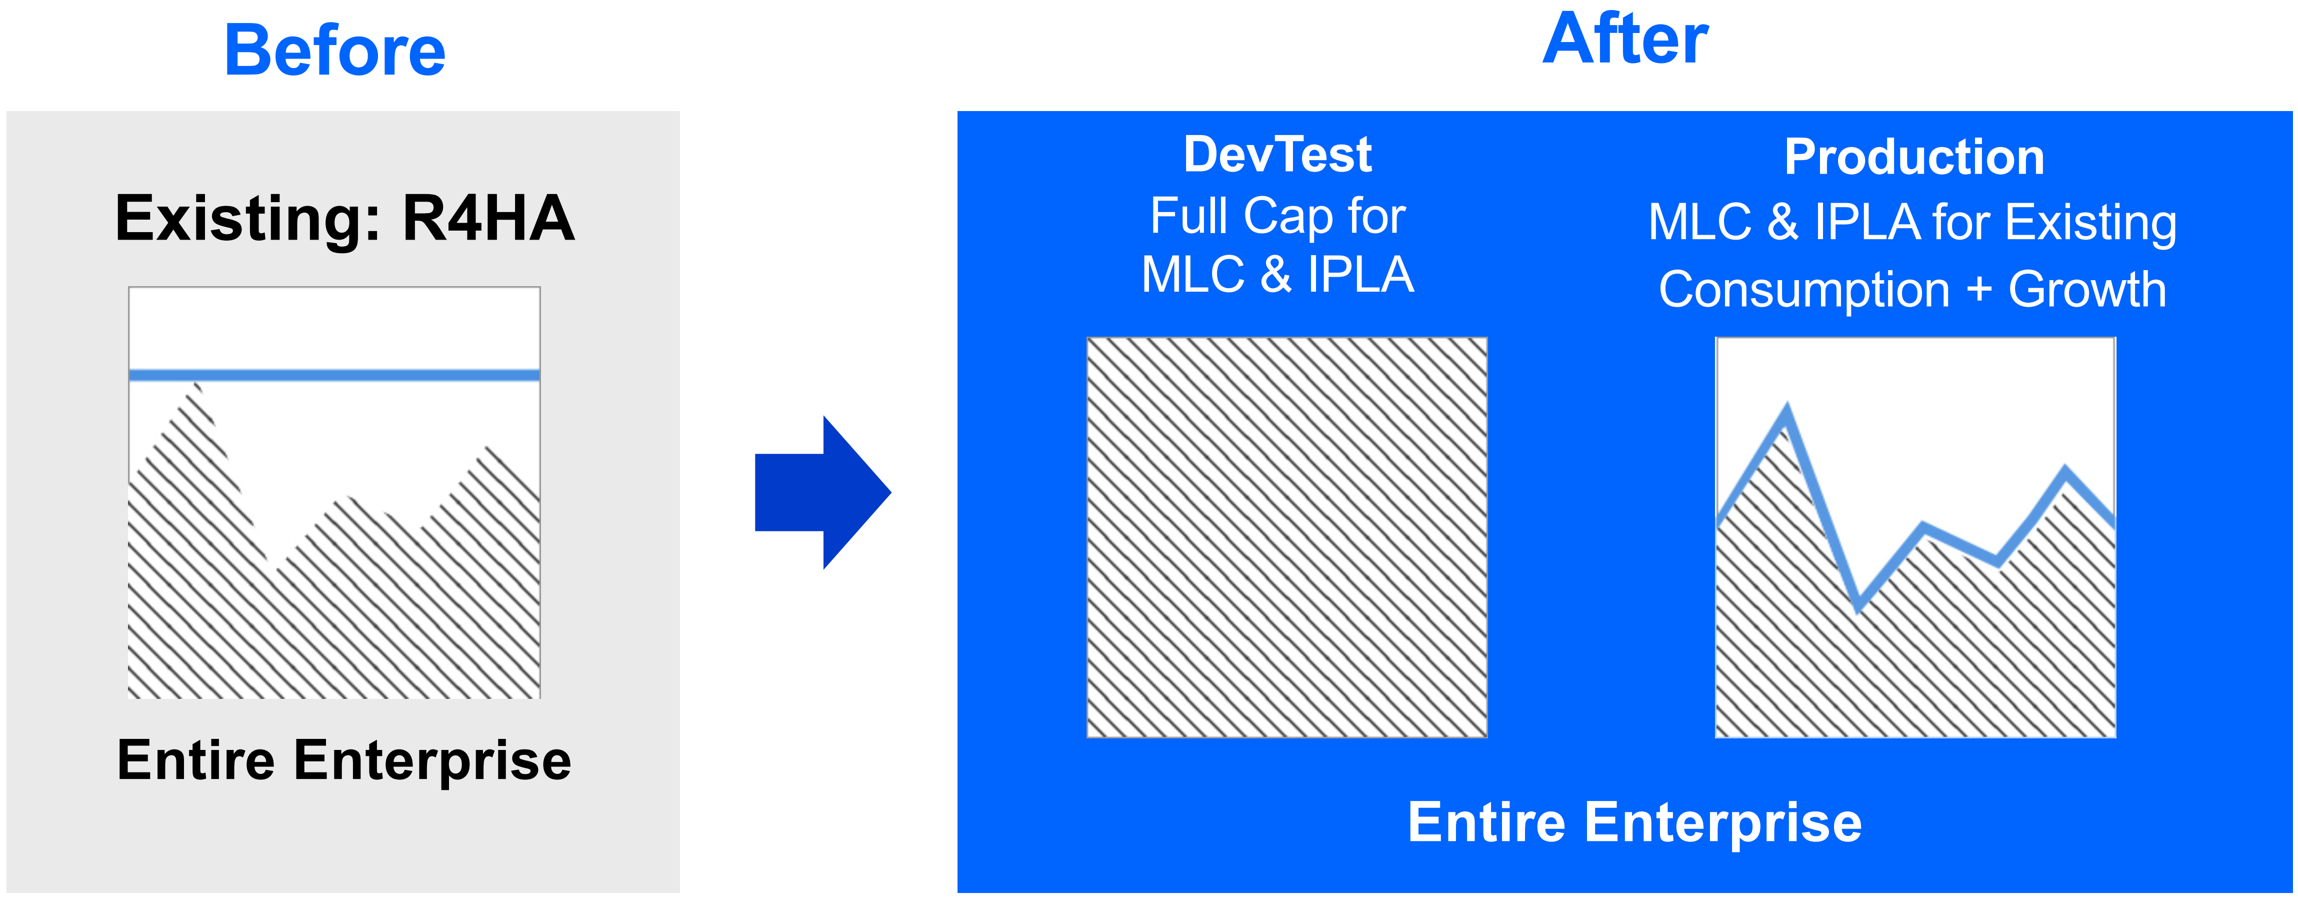

Before

Existing: R4HA

Entire Enterprise

After

DevTest
Full Cap for
MLC & IPLA

Production
MLC & IPLA for Existing
Consumption + Growth

Entire Enterprise

What are the challenges today for managing **cost** and **performance**?

Key User Personas

Gemma
CIO

- Finds it difficult to figure cost and believes IBM Z is too expensive compared to other platforms
- Thinks there is a shortage of IBM Z skills and talent but knows her business is dependent on quality of service levels

Dan
IT Architect

- Configures system to handle peaks
- Projects technology trends to make choices and depends on performance testing understand impact of changes
- Doesn't know how much an application costs

Carl
Capacity Planner

- Has trouble forecasting demand for next 3-4 years
- Has to reduce service towards the end of the month
- Spends too much time on admin, tuning, instead of bringing on new workload

Under the R4HA…

Controlling cost is vitally important. We don’t want any surprises on operational cost so we can accurately manage our budgets

Gemma
CIO

I work to ensure the applications that must run within the peak are optimized but I don’t have time to look at anything outside

Dan
IT Architect

My focus is on ensuring the peak period is kept in control, including capping and planning long term to manage operational costs

Carl
Capacity Planner

…there are consequences

I use several tools to track costs within the R4HA peak and implement various capping strategies.

It’s easier for me to understand the costs on our public cloud contract. We should deploy new work there

My applications are affected by our capping processes. Sometimes our SLAs are impacted and workloads are not taking full advantage of the capabilities of our mainframe

R4HA Reports

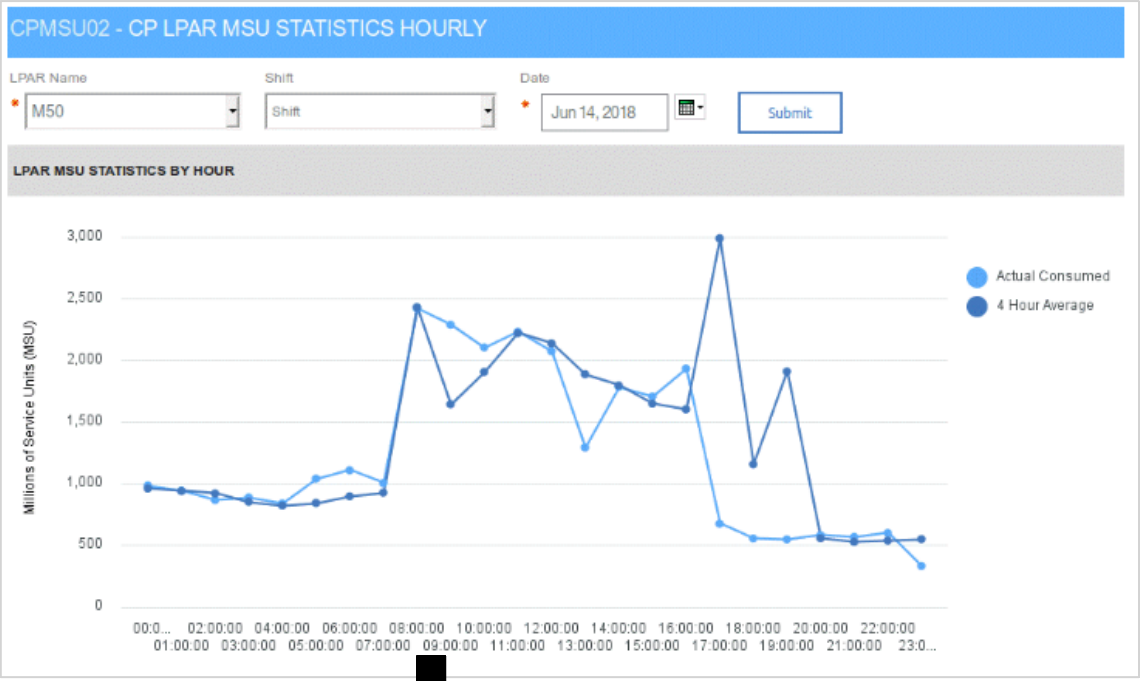

CPMSU02 - CP LPAR MSU STATISTICS HOURLY

LPAR Name * M50

Shift: Shift

Date * Jun 14, 2018

Submit

LPAR MSU STATISTICS BY HOUR

Identifying the peaks earlier can avoid surprises prior to SCRT reports are generated…

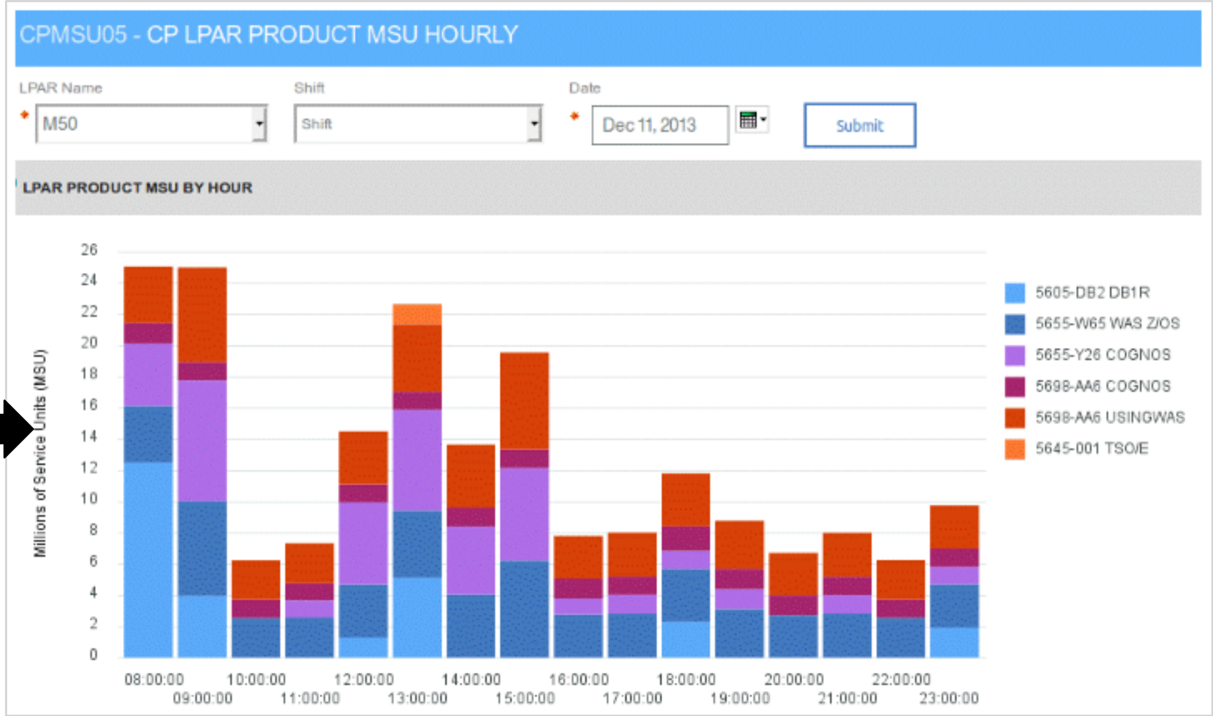

CPMSU05 - CP LPAR PRODUCT MSU HOURLY

LPAR Name * M50

Shift: Shift

Date * Dec 11, 2013

Submit

LPAR PRODUCT MSU BY HOUR

…including drill down to a product's contribution to the R4HA

Capping can negatively impact workload

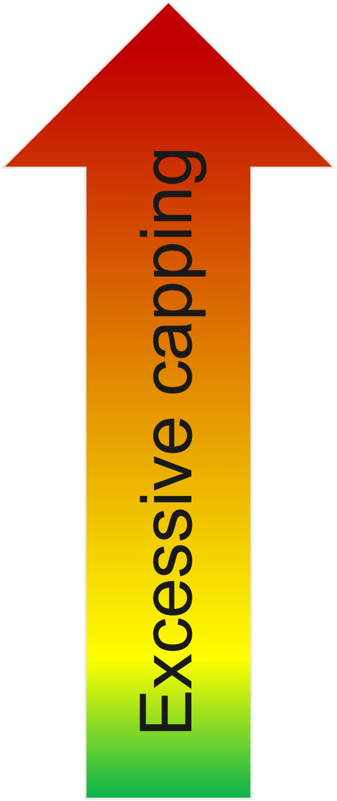

Excessive capping

- System outages
  - Resources not being freed in timely fashion
  - Storage shortages
  - Work (e.g. Service Request Blocks (SRB)) backed up, common storage shortage
- Important work displaced or Service levels missed
- Less important work displaced
- Increased response times or CPU delays

Capping in a "Consumption based" installation ?

We used to cap workloads to limit our costs. Is there any benefit to this now we’ve switched to an MSU consumption model?

- Everything running in one specific 4 hour window each month impacted the cost. Outside the window you could perceive the workload as “free”
- The various capping algorithms helped clients to limit the MSU *peak* usage – with all negative consequences described earlier.
- In a "Consumption based" installation capping is irrelevant – as the pricing is not derived from the peak, but rather from every MSU consumed. Thus, controlling peaks alone (as capping does) is not helpful.
- As there can‘t be a "control via capping" in a "Consumption based" installation – other methodologies need to be applied.

What if we removed capping?

Maximize the hardware, minimize the batch window:

- Let’s assume the nightly batch requires a total of 10,000 MSUs to complete
- Let’s assume the machine is rated at 2,500 MSUs, but capped at 1,800 for the R4HA
- **By removing unnecessary soft caps, batch windows can be dramatically reduced**

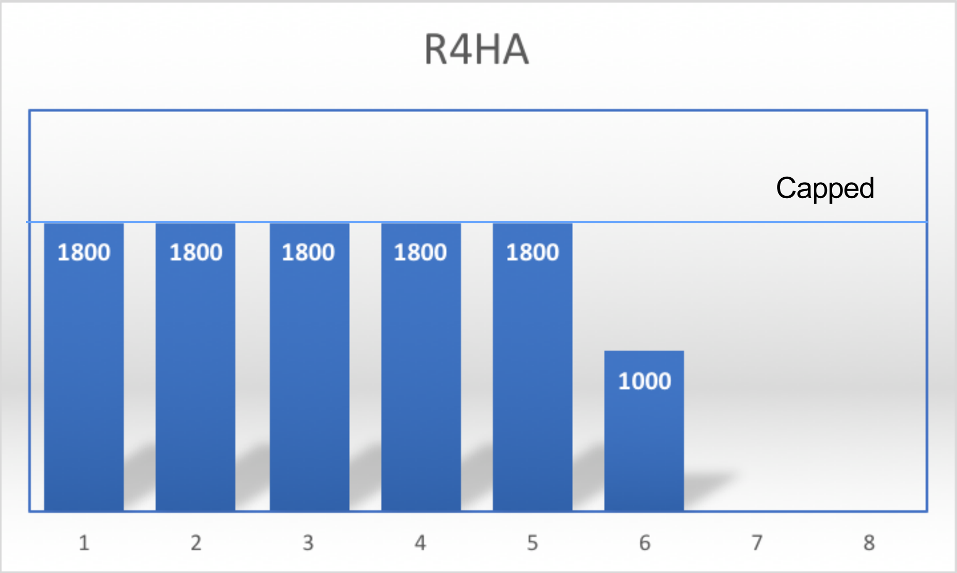

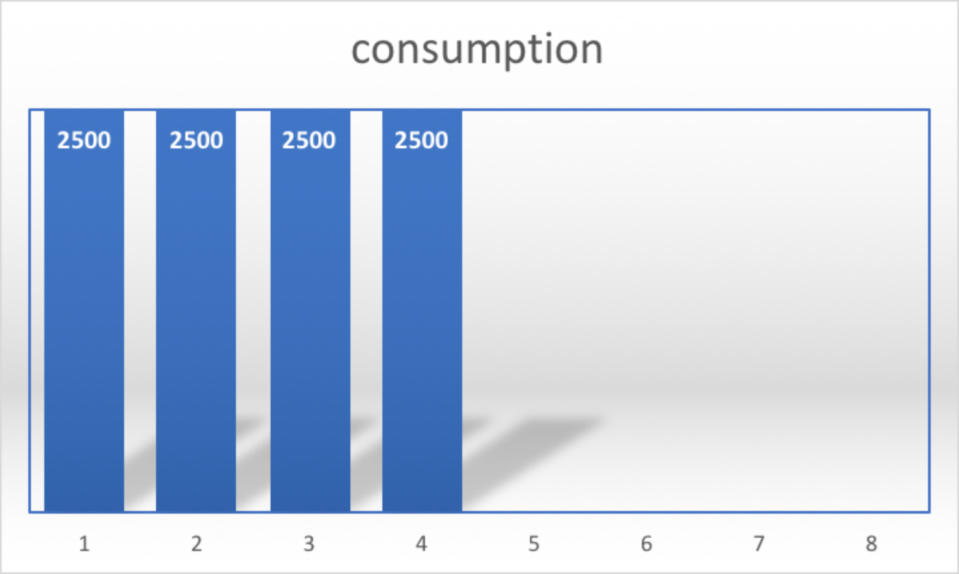

How does transitioning to Tailored Fit Pricing affect **you**?

Under Tailored Fit Pricing new challenges await

We need to understand the level of our current MSU consumption & future use so we can understand impacts of workload changes

My applications are being driven more by external sources. Without capping how do we ensure we are not impacted by spiky workloads

I need to assist in projecting the needs for growing workloads on the mainframe and fit in with our MSU allocation

Challenge #1: Measuring (and forecasting) MSU consumption

We need to understand the level of our current MSU consumption & future use so we can understand impacts of workload changes

- The transparency of consumption based pricing is attractive to many customers
- Forecasting future consumption can help with quantifying costs, tie back to business decisions and avoid surprises
- When forecasting knowing that you may exceed your allocated baseline can be a good thing (or maybe a bad thing!)

The need for clear views of the current state

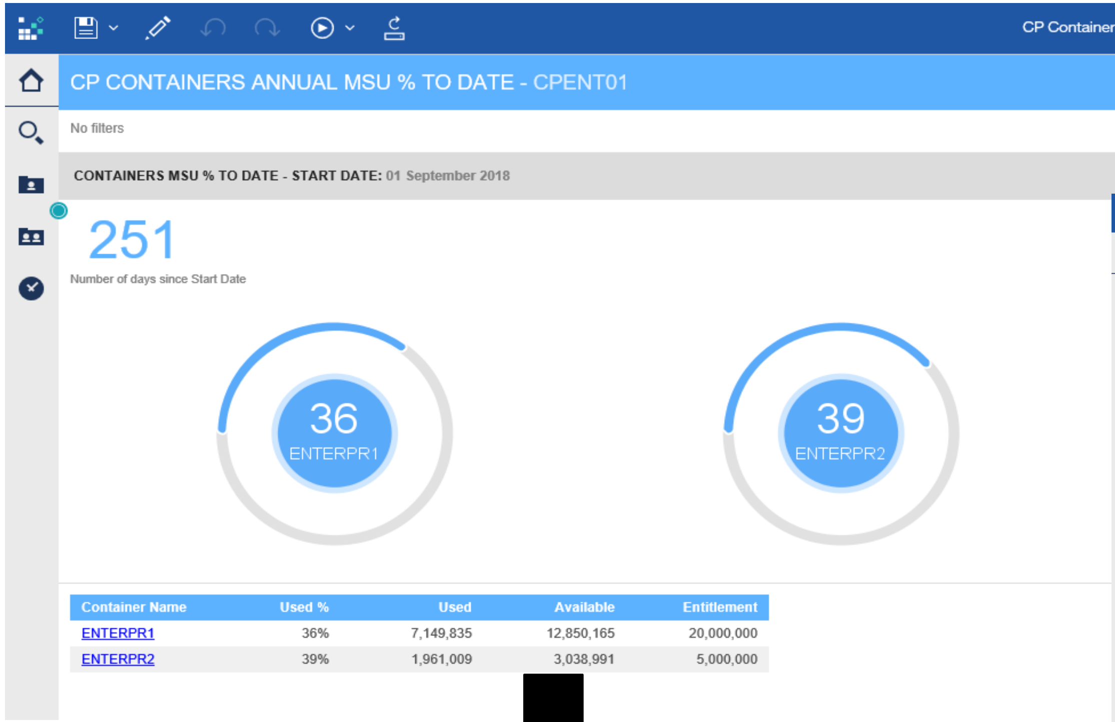

CP CONTAINERS ANNUAL MSU % TO DATE - CPENT01

No filters

CONTAINERS MSU % TO DATE - START DATE: 01 September 2018

251

Number of days since Start Date

Container Name	Used %	Used	Available	Entitlement
ENTERPR1	36%	7,149,835	12,850,165	20,000,000
ENTERPR2	39%	1,961,009	3,038,991	5,000,000

Overall consumption so far versus annual projections

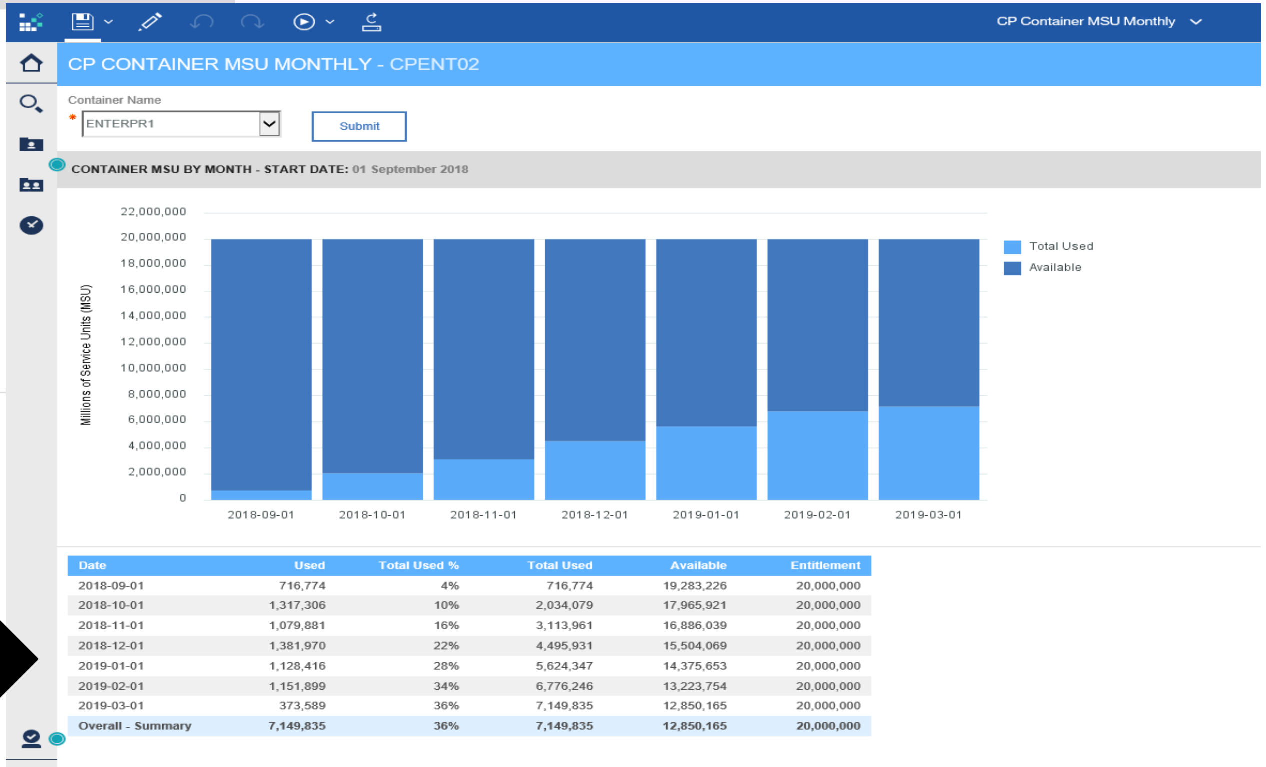

CP CONTAINER MSU MONTHLY - CPENT02

Container Name

* ENTERPR1

CONTAINER MSU BY MONTH - START DATE: 01 September 2018

Date	Used	Total Used %	Total Used	Available	Entitlement
2018-09-01	716,774	4%	716,774	19,283,226	20,000,000
2018-10-01	1,317,306	10%	2,034,079	17,965,921	20,000,000
2018-11-01	1,079,881	16%	3,113,961	16,886,039	20,000,000
2018-12-01	1,381,970	22%	4,495,931	15,504,069	20,000,000
2019-01-01	1,128,416	28%	5,624,347	14,375,653	20,000,000
2019-02-01	1,151,899	34%	6,776,246	13,223,754	20,000,000
2019-03-01	373,589	36%	7,149,835	12,850,165	20,000,000
Overall - Summary	**7,149,835**	**36%**	**7,149,835**	**12,850,165**	**20,000,000**

Where the workload is running

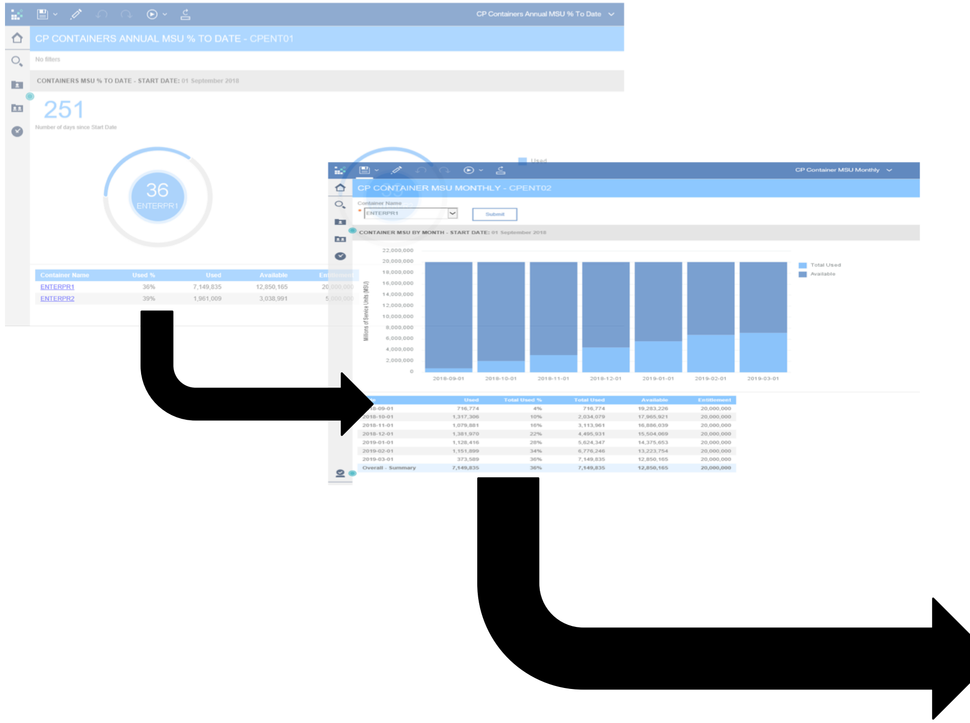
CP Containers Annual MSU % To Date
CP CONTAINERS ANNUAL MSU % TO DATE - CPENT01
No filters
CONTAINERS MSU % TO DATE - START DATE: 01 September 2018
251
Number of days since Start Date
36
ENTERPR1

Container Name	Used %	Used	Available	Entitlement
ENTERPR1	36%	7,149,835	12,850,165	20,000,000
ENTERPR2	39%	1,961,009	3,038,991	5,000,000

CP Container MSU Monthly
CP CONTAINER MSU MONTHLY - CPENT02
Container Name
ENTERPR1
Submit
CONTAINER MSU BY MONTH - START DATE: 01 September 2018
Millions of Service Units (MSU)
22,000,000
20,000,000
18,000,000
16,000,000
14,000,000
12,000,000
10,000,000
8,000,000
6,000,000
4,000,000
2,000,000
0
2018-09-01
2018-10-01
2018-11-01
2018-12-01
2019-01-01
2019-02-01
2019-03-01
Total Used
Available

Date	Used	Total Used %	Total Used	Available	Entitlement
2018-09-01	716,774	4%	716,774	19,283,226	20,000,000
2018-10-01	1,317,306	10%	2,034,079	17,965,921	20,000,000
2018-11-01	1,079,881	16%	3,113,961	16,886,039	20,000,000
2018-12-01	1,381,970	22%	4,495,931	15,504,069	20,000,000
2019-01-01	1,128,416	28%	5,624,347	14,375,653	20,000,000
2019-02-01	1,151,899	34%	6,776,246	13,223,754	20,000,000
2019-03-01	373,589	36%	7,149,835	12,850,165	20,000,000
Overall - Summary	7,149,835	36%	7,149,835	12,850,165	20,000,000

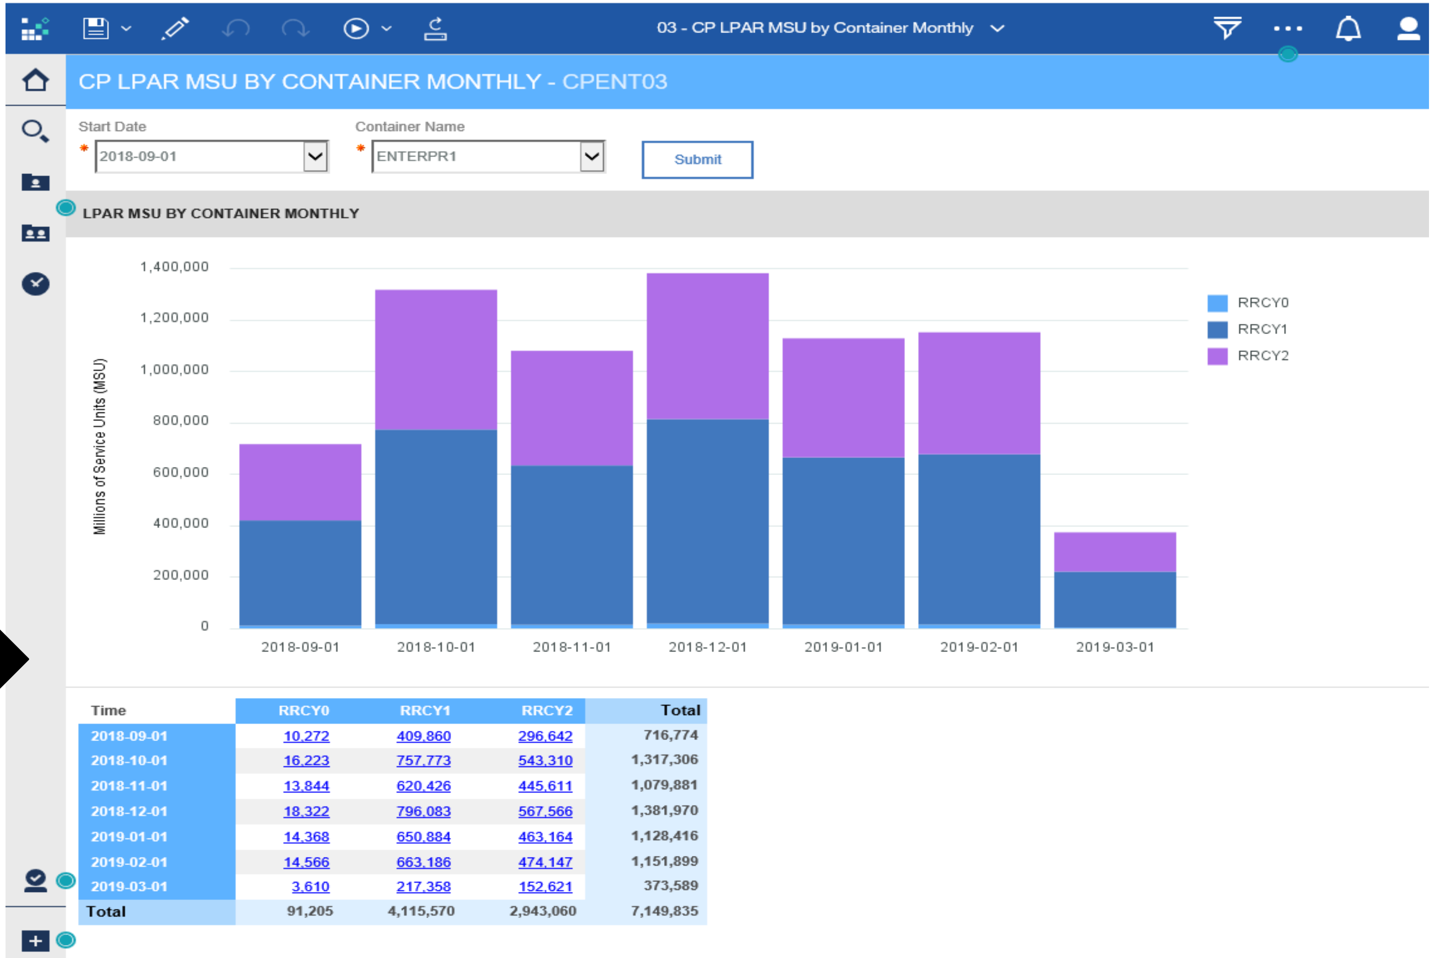

Time	RRCY0	RRCY1	RRCY2	Total
2018-09-01	10,272	409,860	296,642	716,774
2018-10-01	16,223	757,773	543,310	1,317,306
2018-11-01	13,844	620,426	445,611	1,079,881
2018-12-01	18,322	796,083	567,566	1,381,970
2019-01-01	14,368	650,884	463,164	1,128,416
2019-02-01	14,566	663,186	474,147	1,151,899
2019-03-01	3,610	217,358	152,621	373,589
Total	91,205	4,115,570	2,943,060	7,149,835

Now can see on an LPAR or Service Class basis what workload was running and how much did it consume

And which workloads are contributing

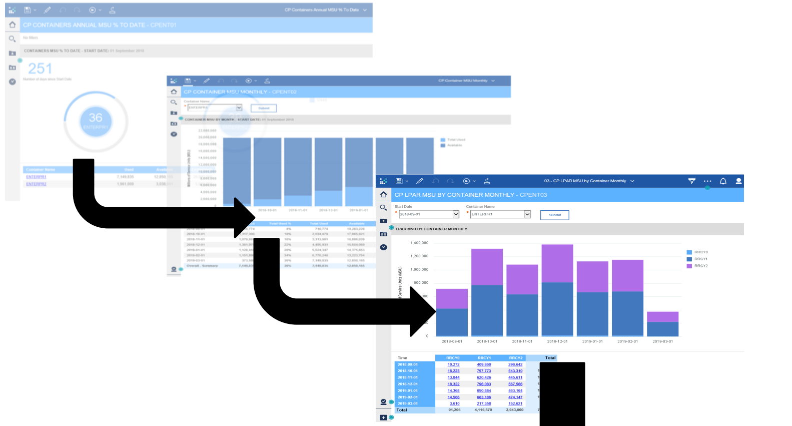

Drilling down further into the jobs or products that are consuming MSUs

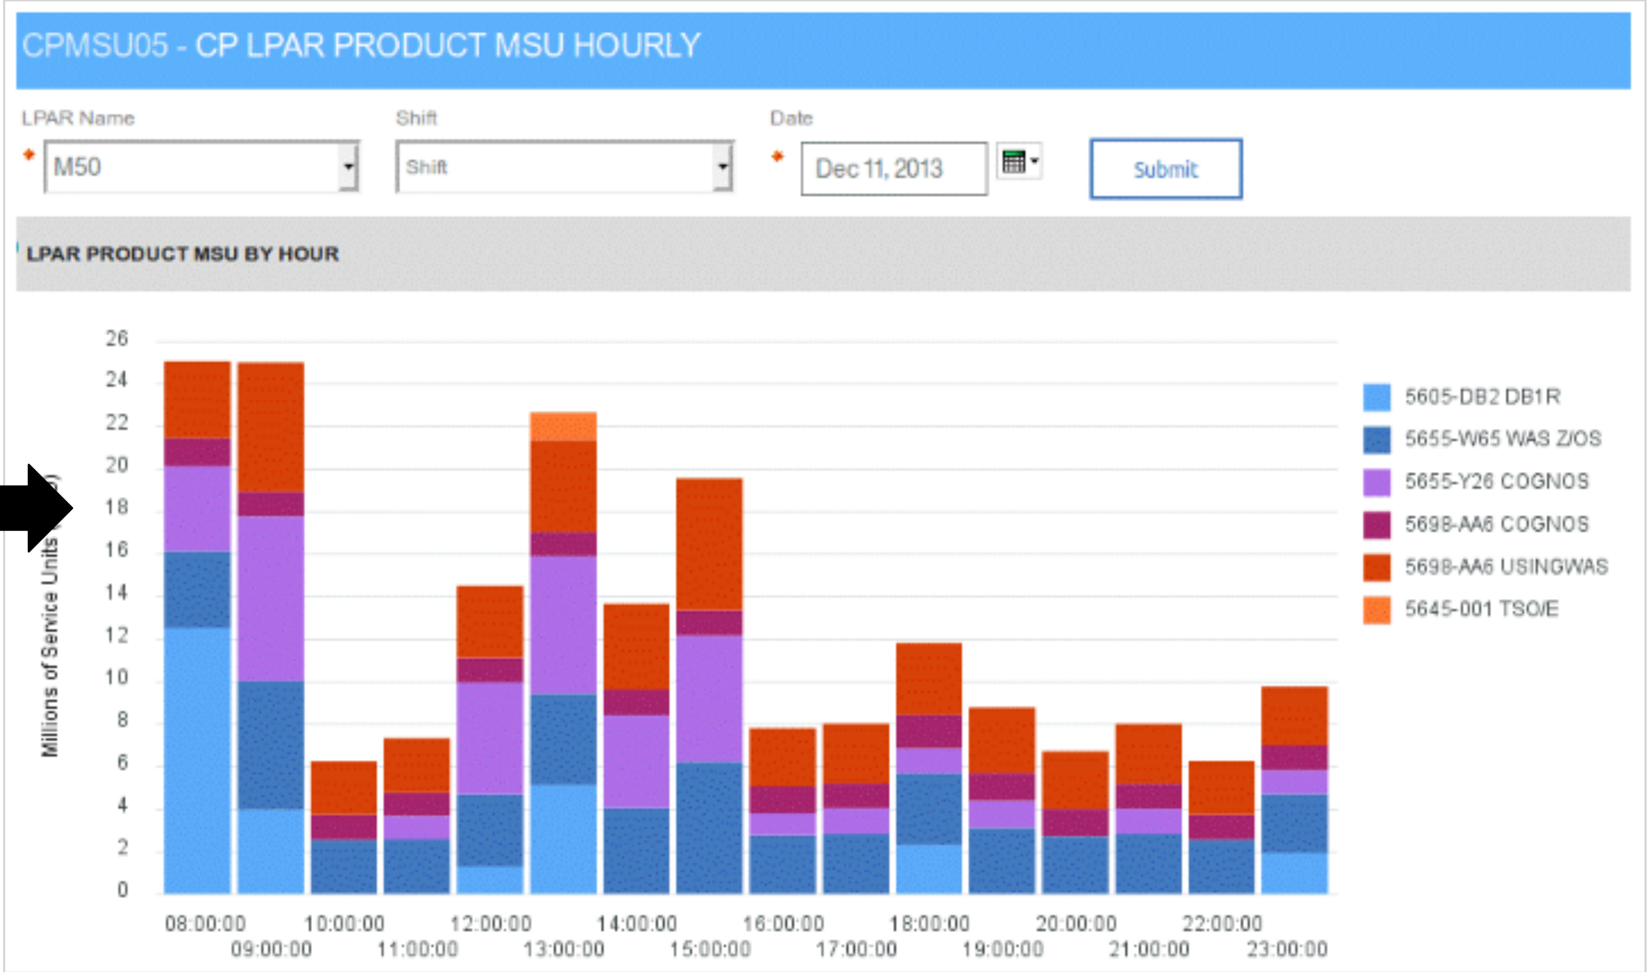

CPMSU05 - CP LPAR PRODUCT MSU HOURLY

LPAR Name * M50

Shift Shift

Date * Dec 11, 2013

Submit

LPAR PRODUCT MSU BY HOUR

Forecasting is critical to avoid surprises…

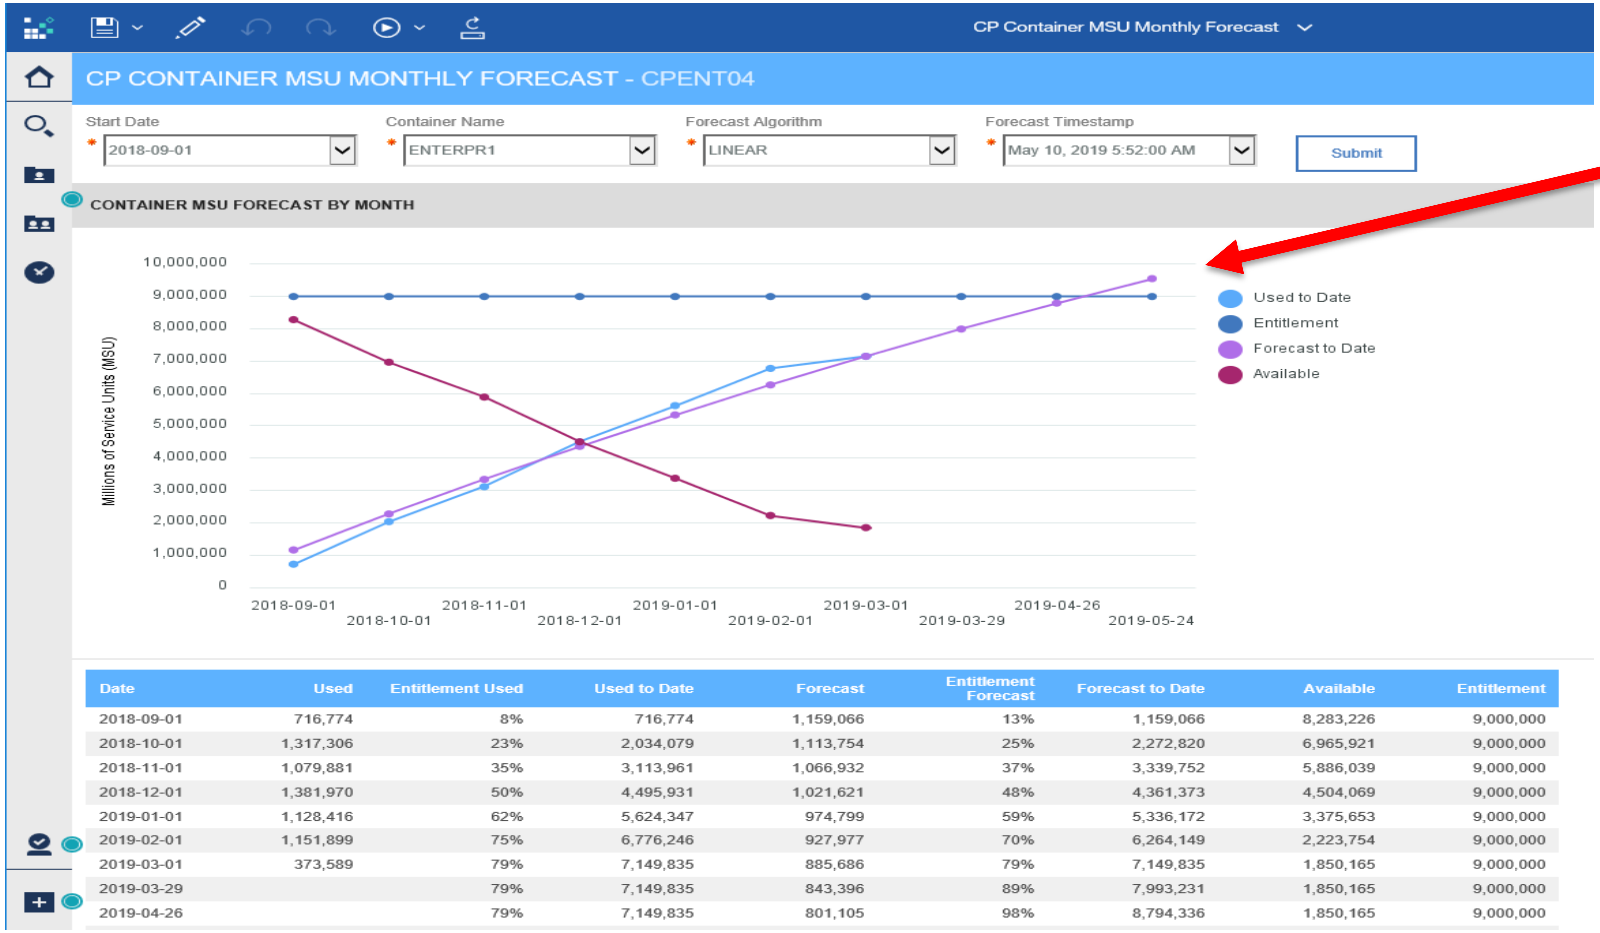

CP CONTAINER MSU MONTHLY FORECAST - CPENT04

Start Date: 2018-09-01

Container Name: ENTERPR1

Forecast Algorithm: LINEAR

Forecast Timestamp: May 10, 2019 5:52:00 AM

CONTAINER MSU FORECAST BY MONTH

Date	Used	Entitlement Used	Used to Date	Forecast	Entitlement Forecast	Forecast to Date	Available	Entitlement
2018-09-01	716,774	8%	716,774	1,159,066	13%	1,159,066	8,283,226	9,000,000
2018-10-01	1,317,306	23%	2,034,079	1,113,754	25%	2,272,820	6,965,921	9,000,000
2018-11-01	1,079,881	35%	3,113,961	1,066,932	37%	3,339,752	5,886,039	9,000,000
2018-12-01	1,381,970	50%	4,495,931	1,021,621	48%	4,361,373	4,504,069	9,000,000
2019-01-01	1,128,416	62%	5,624,347	974,799	59%	5,336,172	3,375,653	9,000,000
2019-02-01	1,151,899	75%	6,776,246	927,977	70%	6,264,149	2,223,754	9,000,000
2019-03-01	373,589	79%	7,149,835	885,686	79%	7,149,835	1,850,165	9,000,000
2019-03-29		79%	7,149,835	843,396	89%	7,993,231	1,850,165	9,000,000
2019-04-26		79%	7,149,835	801,105	98%	8,794,336	1,850,165	9,000,000

The forecast report indicates when we believe MSU consumption will exceed the baseline for the current year

It is for the business to decide if this is a good or bad thing! The key is that you have the visibility to make an informed decision

….and tie costs back to the business

We’re being asked by the business to support new sales events, can we estimate the impact on workloads?

We know what applications will be affected and can see the workload growth so costs can be identified and incorporated into our chargeback process

I can tie this back into my capacity planning dashboards to know that we can support these changes

Challenge #2: Ensuring workloads can be managed effectively

My applications are being driven more by external sources. Without capping how do we ensure we are not impacted by spiky workloads

- Capping gave customers a feeling of reassurance and removing it to exploit the full value of the hardware BUT concern can be that consumption can run away and end up costing more
- Existing tooling focused to R4HA cost control or based of SCRT / RMF report may not have provide the right level of granularity here

Actions to track changes in consumption

- Comparing past days with actual days can help to understand if the consumption is within expectation or not.

- For example: Comparing the same final days of each month, or maybe days of the week across the month to know when typical peaks occur

- Using this benchmarking can help in setting understanding of application performance. Comparing “wrong” days can lead to incorrect assumptions

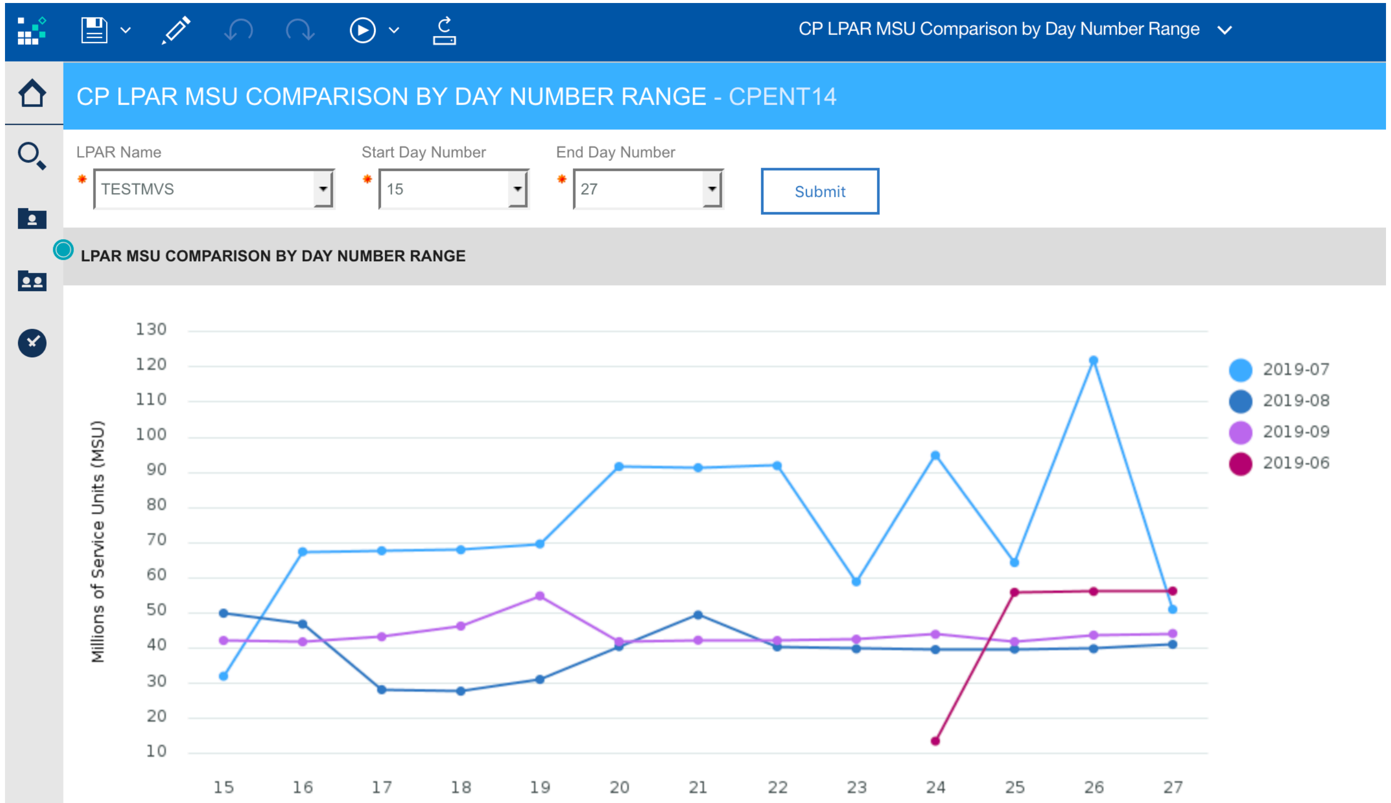

CP LPAR MSU Comparison by Day Number Range

CP LPAR MSU COMPARISON BY DAY NUMBER RANGE - CPENT14

LPAR Name * TESTMVS

Start Day Number * 15

End Day Number * 27

Submit

LPAR MSU COMPARISON BY DAY NUMBER RANGE

Day to day comparison of workload MSU and drill down to hourly levels

Actions to track changes in consumption

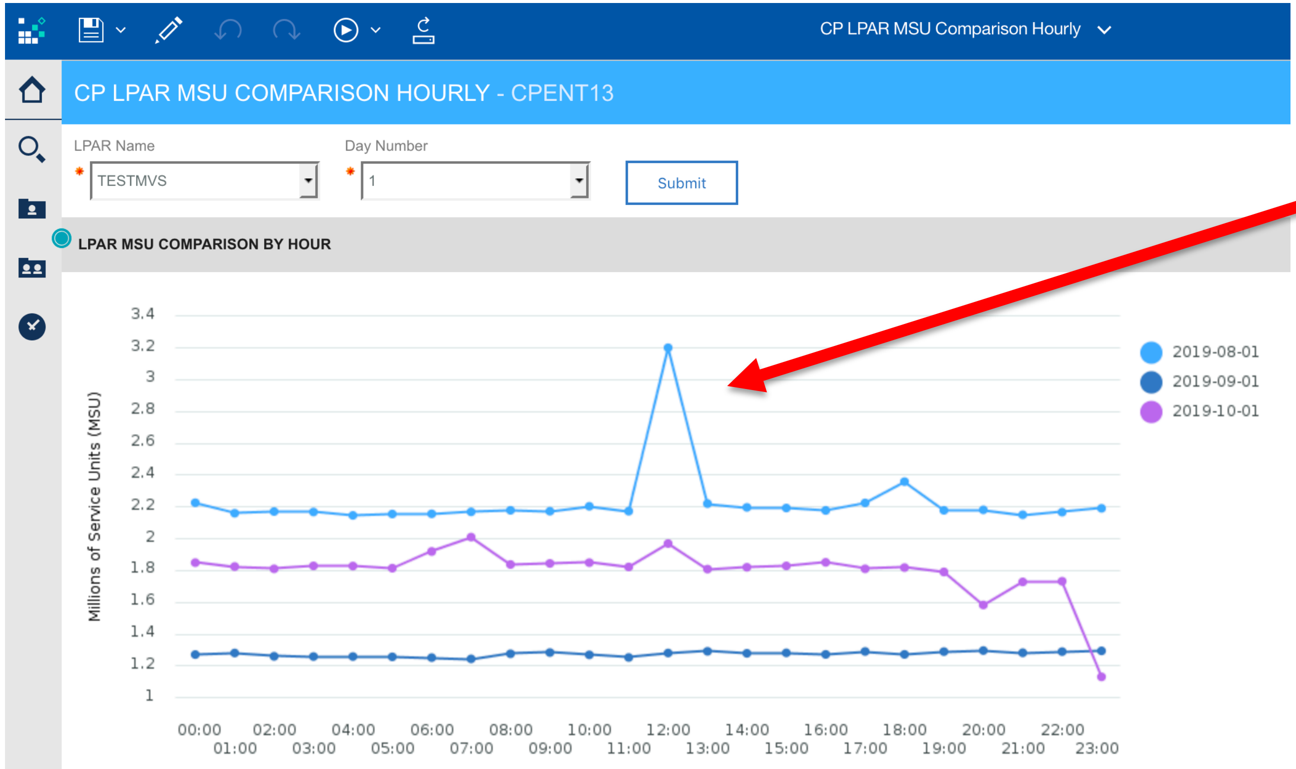

An obvious spike occurs on one LPAR during a given hour. This can be investigated through drill-down

Ensure the right alerts are in place

Our monitoring and other tooling has been tuned to alert and throttle based on experience with the R4HA.

When our application changes are applied, the behavior may change and drive up MSU. I need to consider the impacts

Check what alerting and exceptions you are checking for today

Challenge #3: Optimize use of MSU allocation

I need to assist in projecting the needs for growing workloads on the mainframe and fit in with our MSU allocation

- Note we are focusing now on workload 24/7/365, not just a 4-hour peak period!
- There are many options we take advantage of here:
  - Are we making effective use of zIIP capacity?
  - Are older applications running efficiently exploiting hardware and compiler updates?
  - Would we benefit from a database reorg or looking at our backup policy?
  - Health check of subsystems and looking at resources
- A holistic approach will identify the areas ripe for optimization

Need to focus on ALL workloads

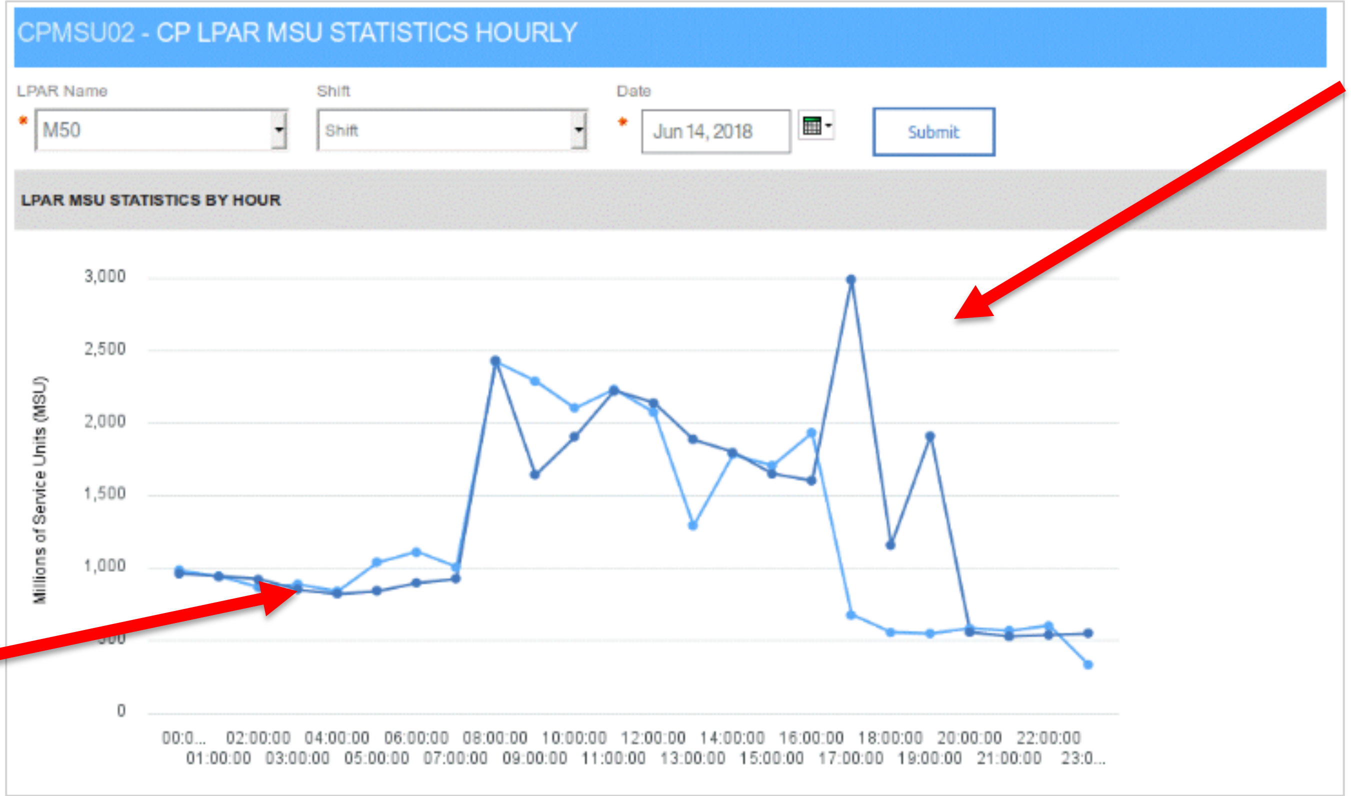

CPMSU02 - CP LPAR MSU STATISTICS HOURLY

LPAR Name: * M50

Shift: Shift

Date: * Jun 14, 2018

LPAR MSU STATISTICS BY HOUR

Previously, focus would be on optimizing workload within the 4 hour peak

Now workload outside of the peak period is just as important to manage

Drill-down to identify the highest consumers

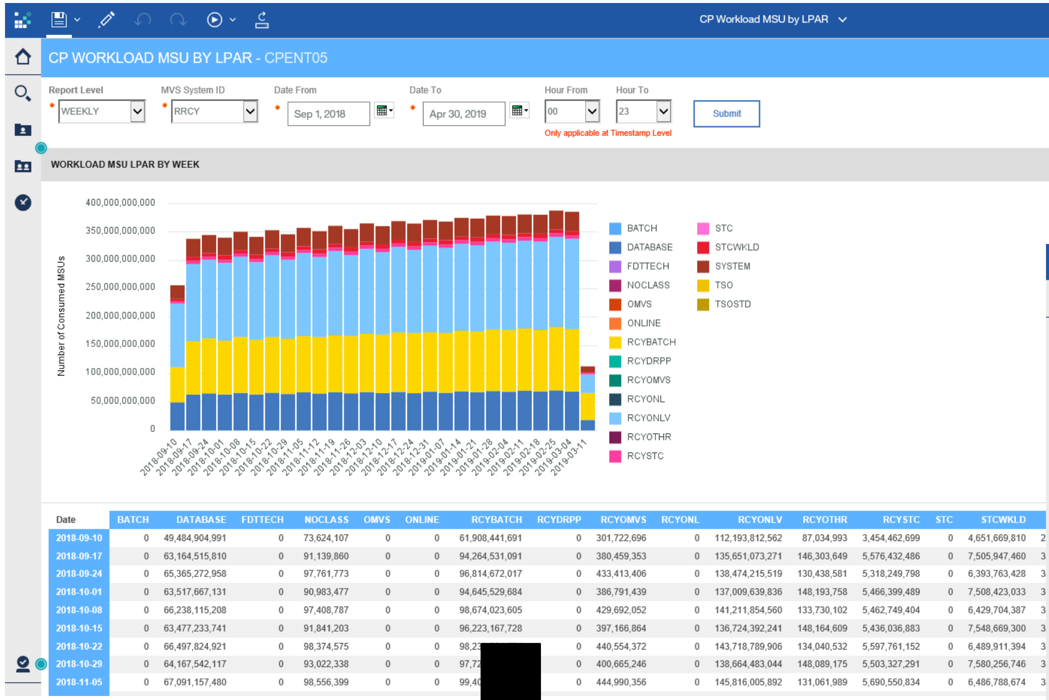

CP WORKLOAD MSU BY LPAR - CPENT05

Report Level: WEEKLY
MVS System ID: RRCY
Date From: Sep 1, 2018
Date To: Apr 30, 2019
Hour From: 00
Hour To: 23
Only applicable at Timestamp Level
Submit

WORKLOAD MSU LPAR BY WEEK

Date	BATCH	DATABASE	FDTTECH	NOCLASS	OMVS	ONLINE	RCYBATCH	RCYDRPP	RCYOMVS	RCYONL	RCYONLV	RCYOTHR	RCYSTC	STC	STCWKLD	
2018-09-10	0	49,484,904,991	0	73,624,107	0	0	61,908,441,691	0	301,722,696	0	112,193,812,562	87,034,993	3,454,462,699	0	4,651,669,810	2
2018-09-17	0	63,164,515,810	0	91,139,860	0	0	94,264,531,091	0	380,459,353	0	135,651,073,271	146,303,649	5,576,432,486	0	7,505,947,460	3
2018-09-24	0	65,365,272,958	0	97,761,773	0	0	96,814,672,017	0	433,413,406	0	138,474,215,519	130,438,581	5,318,249,798	0	6,393,763,428	3
2018-10-01	0	63,517,667,131	0	90,983,477	0	0	94,645,529,684	0	386,791,439	0	137,009,639,836	148,193,758	5,466,399,489	0	7,508,423,033	3
2018-10-08	0	66,238,115,208	0	97,408,787	0	0	98,674,023,605	0	429,692,052	0	141,211,854,560	133,730,102	5,462,749,404	0	6,429,704,387	3
2018-10-15	0	63,477,233,741	0	91,841,203	0	0	96,223,167,728	0	397,166,864	0	136,724,392,241	148,164,609	5,436,036,883	0	7,548,669,300	3
2018-10-22	0	66,497,824,921	0	98,374,575	0	0	98,23	0	440,554,372	0	143,718,789,906	134,040,532	5,597,761,152	0	6,489,911,394	3
2018-10-29	0	64,167,542,117	0	93,022,338	0	0	97,72	0	400,665,246	0	138,664,483,044	148,089,175	5,503,327,291	0	7,580,256,746	3
2018-11-05	0	67,091,157,480	0	98,556,399	0	0	99,40	0	444,990,356	0	145,816,005,892	131,061,989	5,690,550,834	0	6,486,788,674	3

MSU Consumption by Workload and Service Class

Details on the consumption based on daily / hourly / timestamp granularity

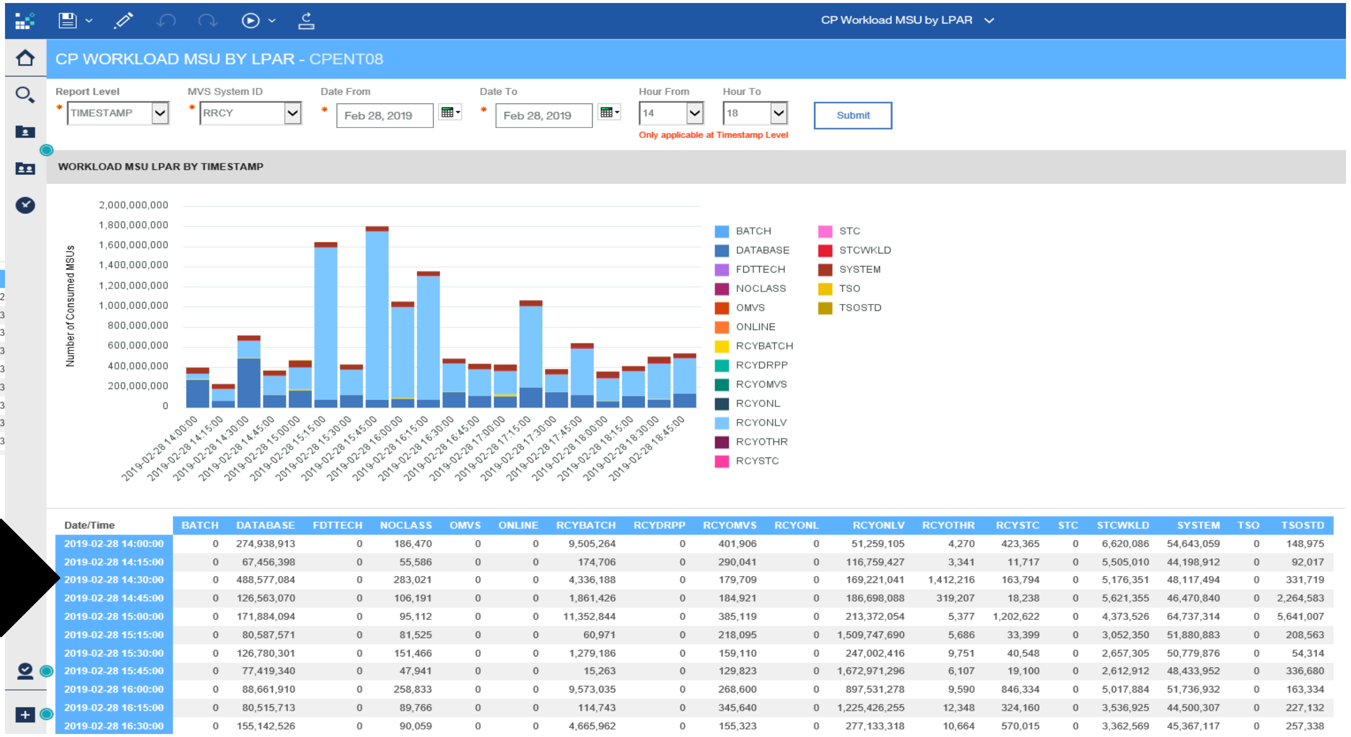

CP WORKLOAD MSU BY LPAR - CPENT08

Report Level: TIMESTAMP
MVS System ID: RRCY
Date From: Feb 28, 2019
Date To: Feb 28, 2019
Hour From: 14
Hour To: 18
Only applicable at Timestamp Level
Submit

WORKLOAD MSU LPAR BY TIMESTAMP

Date/Time	BATCH	DATABASE	FDTTECH	NOCLASS	OMVS	ONLINE	RCYBATCH	RCYDRPP	RCYOMVS	RCYONL	RCYONLV	RCYOTHR	RCYSTC	STC	STCWKLD	SYSTEM	TSO	TSOSTD
2019-02-28 14:00:00	0	274,938,913	0	186,470	0	0	9,505,264	0	401,906	0	51,259,105	4,270	423,365	0	6,620,086	54,643,059	0	148,975
2019-02-28 14:15:00	0	67,456,398	0	55,586	0	0	174,706	0	290,041	0	116,759,427	3,341	11,717	0	5,505,010	44,198,912	0	92,017
2019-02-28 14:30:00	0	488,577,084	0	283,021	0	0	4,336,188	0	179,709	0	169,221,041	1,412,216	163,794	0	5,176,351	48,117,494	0	331,719
2019-02-28 14:45:00	0	126,563,070	0	106,191	0	0	1,861,426	0	184,921	0	186,698,088	319,207	18,238	0	5,621,355	46,470,840	0	2,264,583
2019-02-28 15:00:00	0	171,884,094	0	95,112	0	0	11,352,844	0	385,119	0	213,372,054	5,377	1,202,622	0	4,373,526	64,737,314	0	5,641,007
2019-02-28 15:15:00	0	80,587,571	0	81,525	0	0	60,971	0	218,095	0	1,509,747,690	5,686	33,399	0	3,052,350	51,880,883	0	208,563
2019-02-28 15:30:00	0	126,780,301	0	151,466	0	0	1,279,186	0	159,110	0	247,002,416	9,751	40,548	0	2,657,305	50,779,876	0	54,314
2019-02-28 15:45:00	0	77,419,340	0	47,941	0	0	15,263	0	129,823	0	1,672,971,296	6,107	19,100	0	2,612,912	48,433,952	0	336,680
2019-02-28 16:00:00	0	88,661,910	0	258,833	0	0	9,573,035	0	268,600	0	897,531,278	9,590	846,334	0	5,017,884	51,736,932	0	163,334
2019-02-28 16:15:00	0	80,515,713	0	89,766	0	0	114,743	0	345,640	0	1,225,426,255	12,348	324,160	0	3,536,925	44,500,307	0	227,132
2019-02-28 16:30:00	0	155,142,526	0	90,059	0	0	4,665,962	0	155,323	0	277,133,318	10,664	570,015	0	3,362,569	45,367,117	0	257,338

Look for opportunity to reduce MSU usage

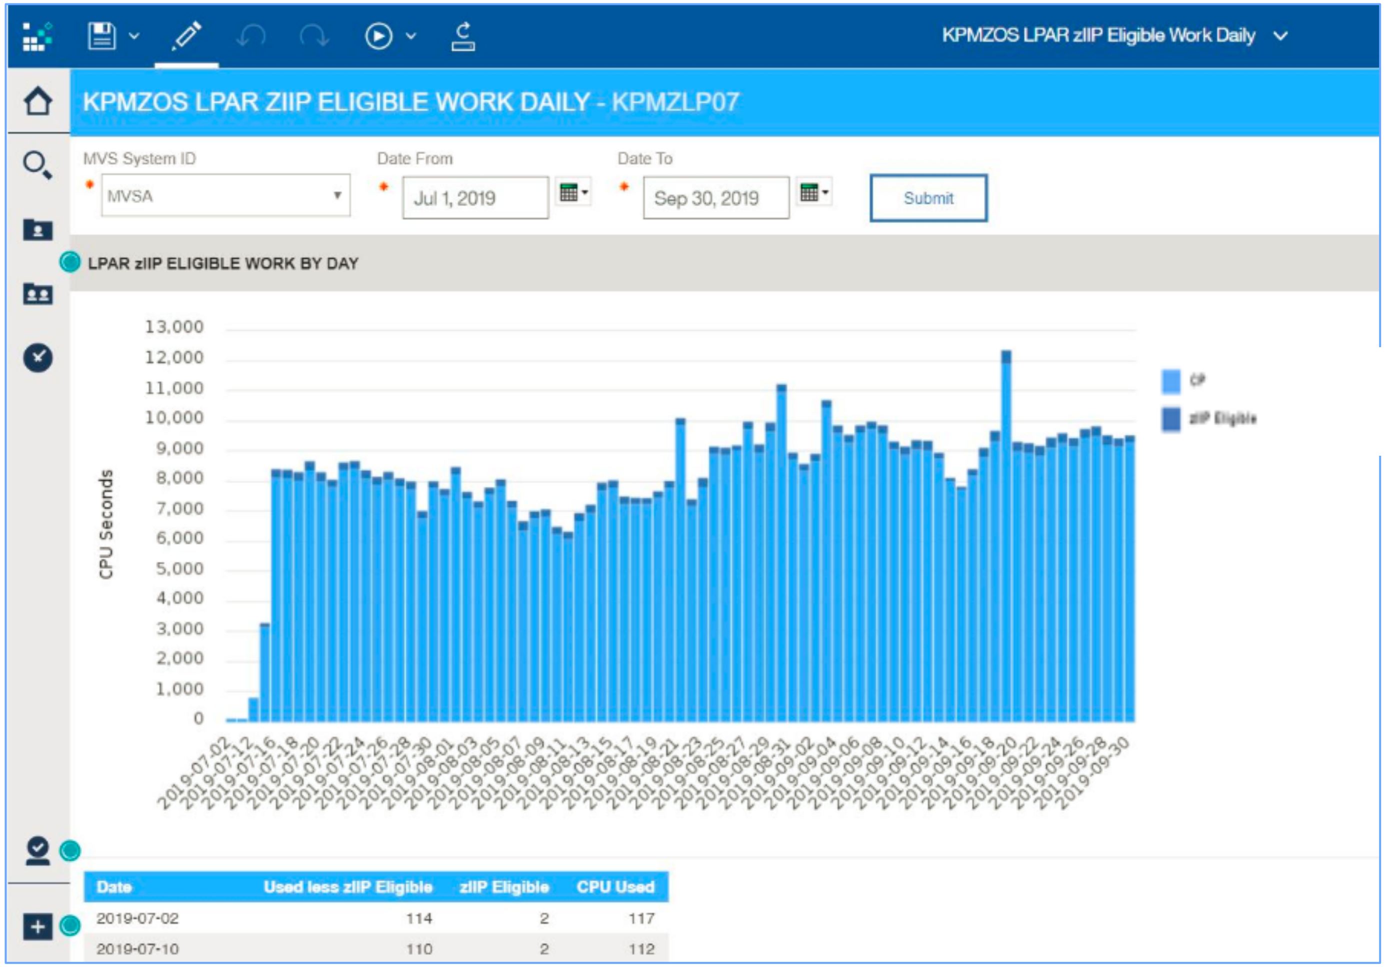

KPMZOS LPAR ZIIP ELIGIBLE WORK DAILY - KPMZLP07

MVS System ID: MVSA

Date From: Jul 1, 2019

Date To: Sep 30, 2019

Date	Used less zIIP Eligible	zIIP Eligible	CPU Used
2019-07-02	114	2	117
2019-07-10	110	2	112

Understand levels of zIIP offload

Tailored Fit Pricing: How To Manage Workloads in a World Without Capping / October 2019

What solutions are **available today** to support you?

Supporting Your Journey to Tailored Fit Pricing

Planning

Managing

Optimizing

Planning

Understanding your current applications and workloads including which products are running on each LPAR

Key tools:

- Tivoli Asset Discovery
- IBM Z Decision Support for Capacity Planning

Managing

Ongoing analysis to visualize current MSU consumption levels per container and forecast future consumption

Key tools:

- **IBM Z Decision Support for Capacity Planning**
- IBM OMEGAMON

Optimizing

Maximizing the value of your deployments to drive efficiency and performance improvement

Key tools:

- Compilers
- Specialist tools (for example: IMS Buffer Pool Analyzer, IBM Z Batch Resiliency)
- IBM Application Discovery and Delivery Intelligence

Managing with IBM Z Decision Support for Capacity Planning

Announced September 2019: ibm.biz/IZOIAnnounce

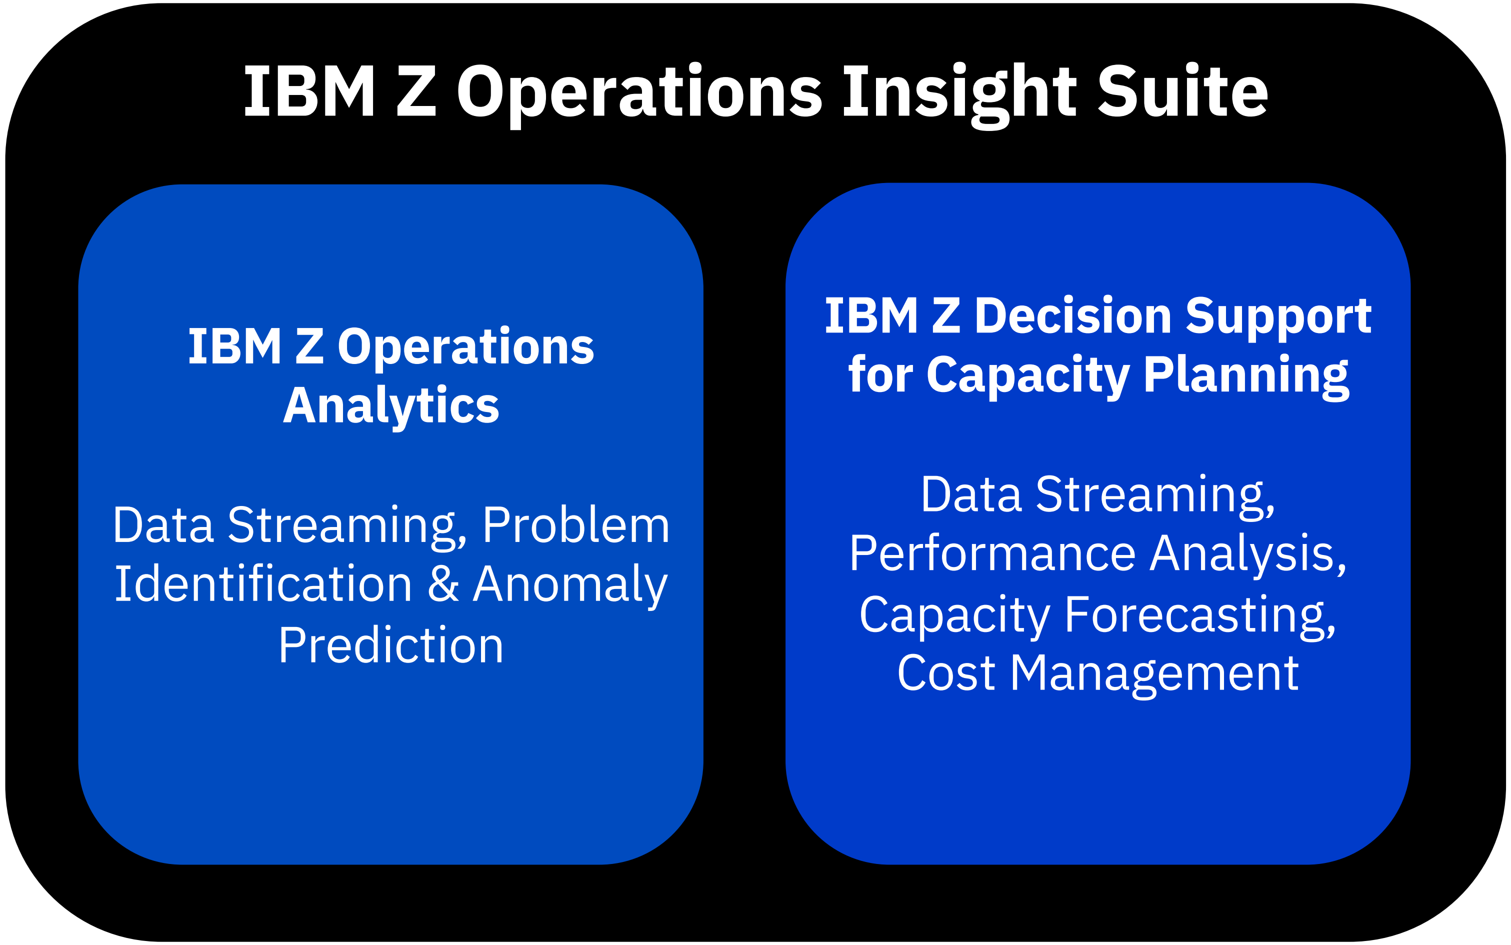

IBM Z Operations Insight Suite

IBM Z Operations Analytics

Data Streaming, Problem Identification & Anomaly Prediction

IBM Z Decision Support for Capacity Planning

Data Streaming, Performance Analysis, Capacity Forecasting, Cost Management

Visualize on the platform of your choice

IT Operational Analytics platforms

IBM

splunk>

elastic stack

Multiple Stakeholders – Single Source of Truth

Z IT Operator

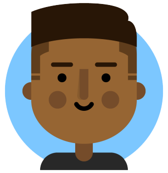

IT Operations Manager

Z Capacity Planner

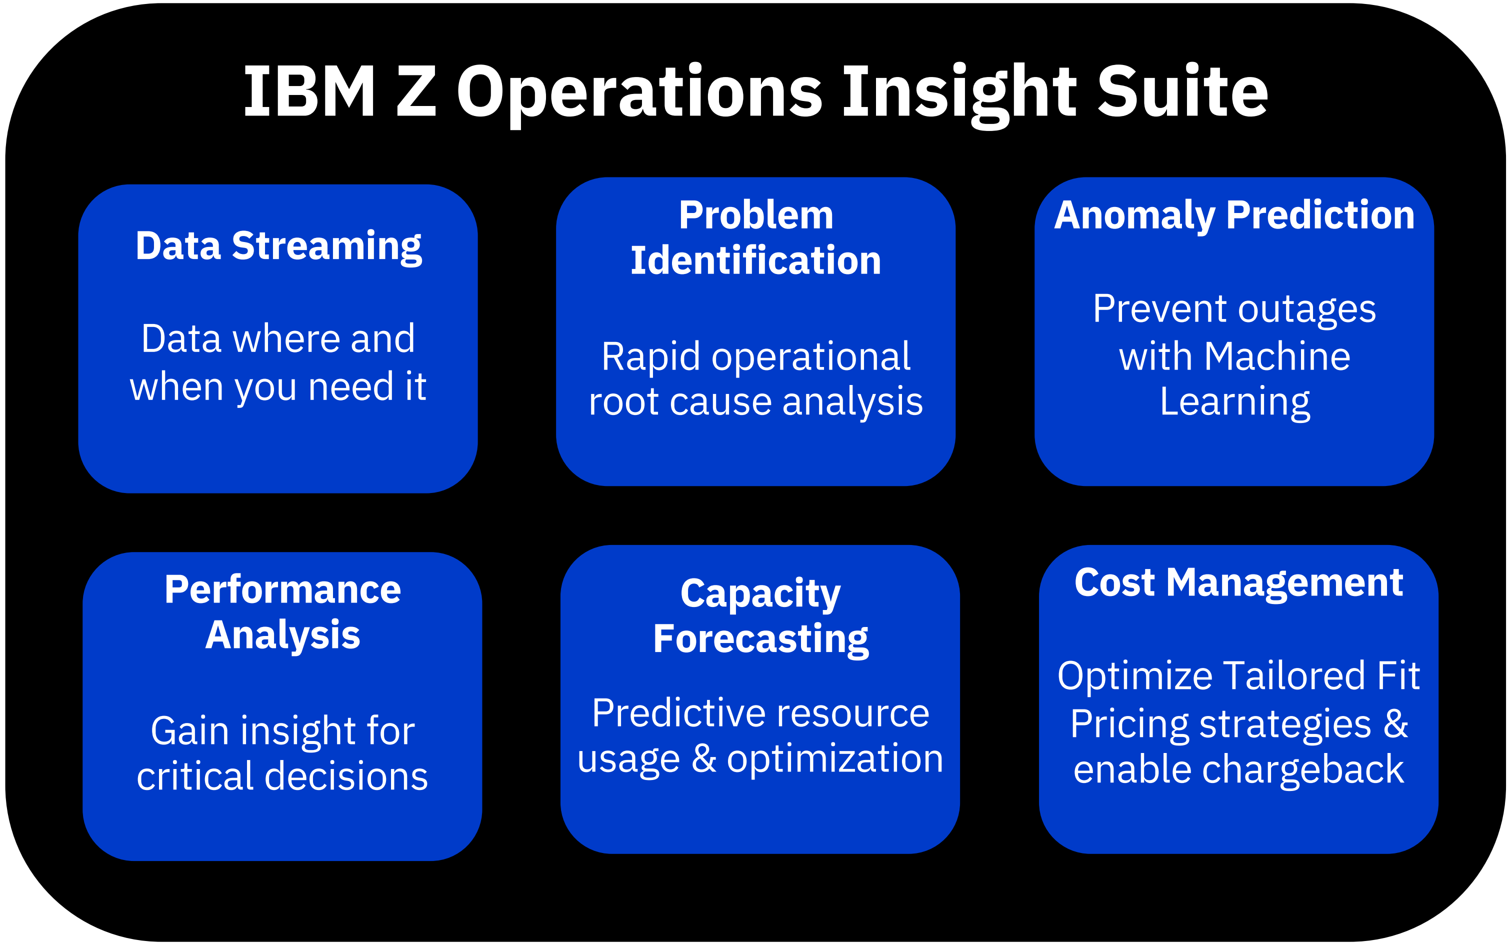

IBM Z Operations Insight Suite

Data Streaming

Data where and when you need it

Problem Identification

Rapid operational root cause analysis

Anomaly Prediction

Prevent outages with Machine Learning

Performance Analysis

Gain insight for critical decisions

Capacity Forecasting

Predictive resource usage & optimization

Cost Management

Optimize Tailored Fit Pricing strategies & enable chargeback

Z SME

Line of Business Owner

Head of Mainframe

Learn More

IBM Z Operations Insight Suite on IBM Marketplace

The latest updates and information about IBM's leading solution for managing IBM Z as part of a hybrid cloud environment

ibm.biz/IZOIInfo

Ibm.biz/IZDSCapPlanInfo

Tailored Fit Pricing: How to manage workload in a world without capping

Discover how IBM Z Decision Support for Capacity Planning can help manage, forecast and optimize your MSU consumption

ibm.biz/IZDSCPandTFP

Achieve Operational Excellence

Learn how to optimize complex hybrid cloud environments with efficient systems management and operational insights

ibm.biz/OpExcellence

IBM Z Software Newsletter: Operations & Management Edition

Subscribe to the quarterly newsletter for operation, systems programmers and administrator to get the latest news, tips, blogs and trials in one place

ibm.biz/ZOperations

Summary

Tailored Fit Pricing is revolutionary new pricing model for software on IBM Z that eliminates the need for capping and provides a complete alternative to the R4HA.

To be successful requires a change in how you look at managing workloads. Instead of capping to contain costs, optimizing existing applications is the key and this cannot be done without a clear view on current consumption levels

With a clear view on usage levels of workloads, accurate forecasts can be made that tie back to business drivers and your capacity planners are empowered to model workload growth on the mainframe

IBM Z Decision Support for Capacity Planning, as part of the IBM Z Operations Insight Suite, gives unique visibility into current and future MSU consumption and provide guidance into how to optimize workloads to control costs

Reach out to us for a deep-dive discussion on how you can leverage the current capabilities or join one of our Proof of Technology workshops

Please submit your session feedback!

- Do it online at http://conferences.gse.org.uk/2019/feedback/oo
- This session is OO

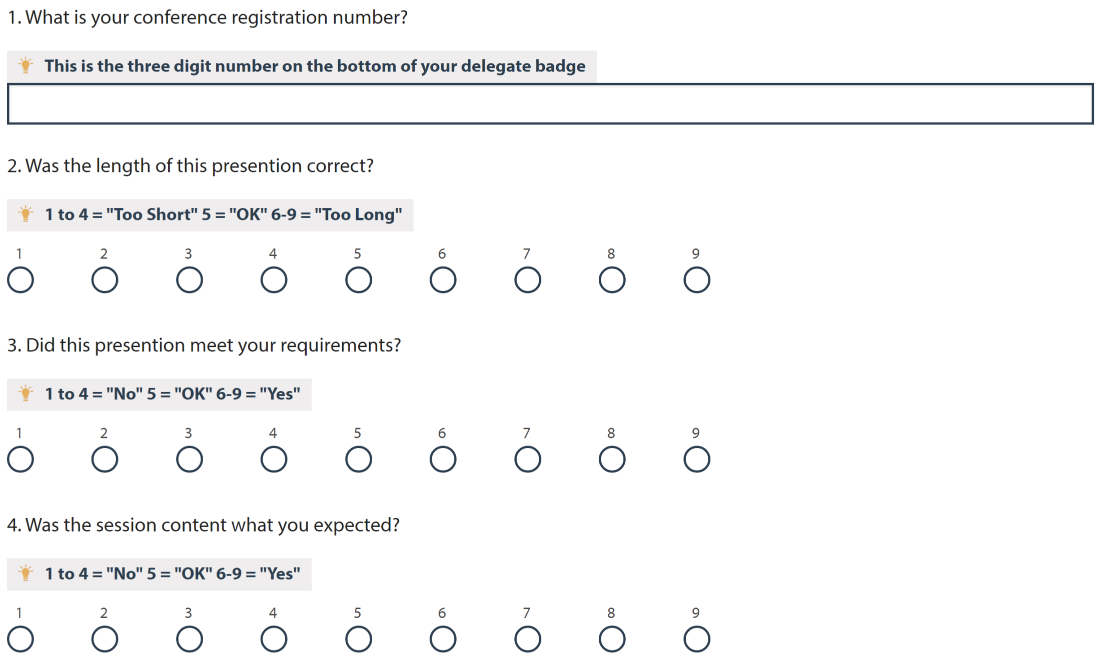

1. What is your conference registration number?

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

2. Was the length of this presention correct?

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

1	2	3	4	5	6	7	8	9
○	○	○	○	○	○	○	○	○

3. Did this presention 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
○	○	○	○	○	○	○	○	○