

# Mainframe Capacity and Performance Challenges in an Outsourced Environment

Steven Thomas, Principal ITBI Evangelist, SMT Data  
Sth@SMTData.com

### Challenges:

For the customer:

- You can outsource the work, but not the responsibility.
- Lack of transparency makes it hard to stay in control.
- Skills have been outsourced.

For the outsourcer:

- Balancing contracts, customer satisfaction and P/L.
- Relating external costs to revenue.

For both:

- Ensuring the partnership is win-win.
- Being agile and fighting inertia.



### Topics:

- What do I pay for?
- How is it measured?
- What do I need to follow up on?
- What tradeoffs and pitfalls do I need to consider?
- What data do I need?



# What do I pay for?

# What do I pay for?

## Many cost drivers

Hardware: CPU, storage, memory, net, etc.

Software: MLC, IPLA, vendor, etc.

Facilities, people, power, cooling etc.

Resilience: Backup, security, etc.


SLA: Availability, performance, etc.

Services: Run and change the business.

Terms and Conditions: Payment terms, regulatory compliance, ownership, design authority, term and termination, location, etc.

'Intangibles': Agility, complexity

And more...



## Many cost drivers

Hardware: CPU, storage, memory, net, etc.

Software: MLC, IPLA, v

Resilience: Backup, security, e

Services: Run and change the bu

SLA: Availability, performance, e

Terms and Conditions: Payment, design authority, term and termi

'Intangibles': Agility, complexity

Facilities, people, power, cooling etc

And more...

Outsourcers often base charges on CPU.

- Other cost drivers are often baked into the CPU price.
- Simple for the customer
- Complex for the outsourcer
- Contracts are often fuzzy on:
  - What units are used for billing (e.g. cpu seconds, MIPS, MSU)
  - How those units are measured or calculated
  - Reporting and data sharing
  - How tradeoffs between cost and performance are made

## **Kinds of CPU: CP, zIIP, IFL**

General Processors ('GP', 'GCP', or simply 'CP')


- Can run all workloads

Integrated Information Processors ('zIIP')

- A CP with microcode to limit functionality
- Only for certain types of workloads (Java, some DB2, ML, Python, etc.)
- A 'commercial' not a 'technical' invention, allows IBM to provide a discount for certain workloads


Integrated Facility for Linux ('IFL')

- Like a CP, but with certain instructions disabled so it can only run Linux
- Also a way of providing a discount for selected workload




### Measuring CPU – CPU Seconds

- The number of seconds that a CPU has been busy on a certain activity
- Measured and reported directly by the operating system in SMF
- Not normalized – faster hardware will take fewer CPU seconds to perform the same work
- If used for billing purposes, then the contract normally includes a price adjustment mechanism for hardware or configuration changes
- Typically measured using SMF70PDT or SMF70EDT




### Measuring CPU – MIPS

- Originally 'Million Instructions per Second', but no longer directly related to number of instructions
- A way of normalizing CPU usage across different hardware models and configurations
- In theory, the same workload uses roughly the same number of MIPS no matter what configuration it is run on
- Cannot be measured directly
- Calculated from CPU seconds and a configuration dependent MIPS factor (more on this later)
- Can be computed for all processor types (CP, zIIP, IFL)



### Measuring CPU – MSU


- Originally ‘Million Service Units’
- Calculated based on factors determined by IBM
- Like MIPS, a way of normalizing CPU usage across different CPU models and hardware configurations
- Used by IBM (and others) to charge for software
- In theory, the same workload uses roughly the same number of MSU’s no matter what configuration it is run on
- But IBM may adjust the factor for different HW for commercial reasons
- Only defined for CP processors



### Measuring CPU – MSU

The term 'MSU' can mean different things depending on the context:

- The actual hardware MSUs consumed during a period – roughly corresponds to 5-5.5 MIPS
- The actual software MSUs consumed during a period – roughly corresponds to 8-8.5 MIPS (iMSU, ACTMSU, TFP MSU)
- The rolling four-hour average of consumed software MSUs as logged by the operating system every few minutes. This is reported in SMF70LAC.
- The hourly average of these individual rolling four-hour averages. This is reported in the SCRT report.
- The monthly peak of these hourly averages. This is the number on the SCRT report that IBM uses to bill software (WLC / MLC).



## How is it measured?

### iMSU versus MIPS

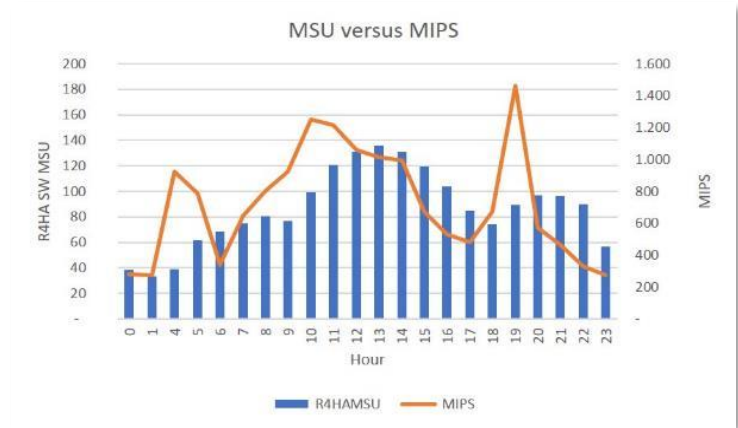
- iMSU and MIPS measure the same thing, but there are between 8.0 and 8.5 MIPS per iMSU depending on:
  - MIPS calculation method
  - Machine model
- iMSU and MIPS have the same peaks at the same time
- MIPS is also defined for zIIP and IFL, but MSU is only defined for CP
- iMSU as used in TFP is fully defined by fields in the SMF (SMF70EDT, SMF70CPA\_actual, SMF70CPA\_scaling\_factor)
- MIPS are based on fields in the SMF (e.g. SMF70EDT or SMF70PDT) and a MIPS rate from a table.



## How is it measured?

### R4HA MSU versus MIPS

- R4HA flattens spikes
- R4HA peaks are to the later than the corresponding MIPS peaks
- R4HA peaks and MIPS peaks can be at completely different times



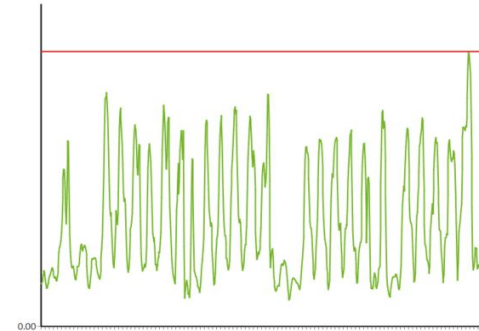
### R4HA MSU versus iMSU and MIPS

R4HA MSU	iMSU	MIPS
Measured directly	Computed from CPU seconds, method defined by IBM	Computed from CPU seconds, various methods
Applies only to CP	Applies only to CP, but can be 'simulated' on zIIP and IFL	Applies to all processor types
Only measured at the LPAR level (SMF70LAC)	Can be computed at all levels (job, transaction, etc.)	Can be computed at all levels
Flattens out spikes	Can be spikey	Can be spikey

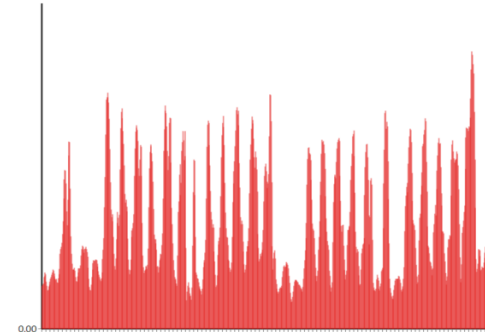
### Where are we headed? Tailored Fit Pricing (TFP)

- Not usually available for outsourcers, but IBM has hinted that it will be in the future
- Cost based on the total number of MSU hours used during the month for a group of LPARS
- Up-front commitment and discount on growth
- Peaks become less interesting
- Capping becomes less relevant
- Cost optimization involves understanding what is going on 24x7 and optimizing the big consumers
- Less need to worry about which software products are running on which LPARS
- Much simpler cost allocation and billing for outsourcers – you pay for what you use

R4HA



TFP



### Calculating MIPS


MIPS = CPU seconds \* configuration dependent MIPS factor

For contractual purposes, the MIPS factor is generally determined based on:

- 'Effective' machine model – determined by model and number of *shared* CPU's
- MIPS per CPU delivered by that model – determined by a standard table (e.g. Watson Walker, IBM LSPR)

Many other factors can affect the actual MIPS factor such as

- Workload characteristics
- Vertical affinity
- Overall system load
- And many others



### Finding the 'effective' machine model

In theory, the model can be read from SMF70, e.g. 9175-624

But

- The nominal model only includes CP's, and the MP overhead is also affected by number of zIIPs.
- Processors can be dedicated to specific LPARS, reducing the number of shared processors.
- zIIPs always run full speed, even on models where the CP's run half speed (kneecapped).
- Many contracts are silent on how the MIPS factor is determined.
- zPCR can be helpful in calibrating your choices.



## How is it measured?

### Finding the MIPS rate

- Use a table (here thanks to Watson Walker) to determine the MIPS/CP for the effective model.
- Use 'Avg RNI MIPS/CP'.
- Note that the Watson Walker numbers are very similar to the IBM PCI numbers (from LSPR)
- Watson Walker tables also include S/W MSUs and MIPS/MSU ratios, which can be useful.
- For example, a model 9175-721 has 21 processors and a MIPS/CP of 1816.1 and a MIPS/MSU ratio of 8.5.

Type-Model	# of CPs	IBM PCI	Avg RNI MIPS	Avg RNI MIPS/CP	S/W MSUs	MIPS/S/W MSU
9175-721	21	38141	38138.2	1816.1	4471	8.5
9175-722	22	39621	39618.2	1800.8	4645	8.5
9175-723	23	41099	41096.0	1786.8	4818	8.5

## How is it measured?


### Calculating MIPS

For billing purposes, MIPS are normally calculated as an average over one hour:

$$\text{MIPS} = (\text{cpu seconds per hour} / 3600) * \text{MIPS per processor}$$

e.g. if we use 100,000 cpu seconds in an hour on a z17 model 721 we get:

$$(100,000 \text{ cpu seconds} / 3600) * 1816.1 = 50,447 \text{ MIPS}$$



### Best Practice

Have a clear definition of how CPU is measured in your contract

- If it is based on MIPS then the contract must specify:
  - Which MIPS table is used (e.g. LSPR)
  - How the model is determined (e.g. based on the number of shared CP's)
  - Use 'average' RNI
  - MIPS are computed as an hourly rate
  - The source of the CPU seconds used in the calculation (typically SMF70PDT but could be SMF70EDT)
  - How zIIP usage is handled
- If it is based on MSU then the contract must specify:
  - Software MSU (not Hardware MSU)
  - iMSU or R4HA MSU
  - How zIIP usage is handled



## How is it measured?

### Best Practice

Ensure access to data that will allow you to follow up on all the CPU usage:

- SMF data or corresponding reporting
- Specification of how the invoice is calculated each month




### Examples of outsourcer billing models

Outsourcer costs are typically driven by monthly 4HRA peak MSU (SW) and monthly MIPS peak (HW).

But outsourcers have many creative ways of billing their customers:

- Peak MIPS per hour per month
- Peak R4HA MSU (or MIPS) per month
- Peak hour during working hours only
- Percentile Peak (e.g. 95<sup>th</sup> Percentile Peak)
- Peaks calculated independently or combined peak across multiple LPARs
- Peaks for CP and zIIP calculated together or separately
- Fixed capacity (e.g. based on weight, lcp's, capping)
- Total or average usage per month (e.g. TFP)
- Combinations of the above e.g. CP usage based on MIPS peak and zIIP usage based on fixed capacity



### **Billing model best practice**

- Work towards alignment of interest
- Look for win-win, a shared environment and shared resources should be a benefit for both outsourcer and customer
- Ensure transparency through access to data and reporting
- Understand historic usage patterns and future plans when considering alternate billing models
- Avoid complexity that doesn't add value
- Ensure clear communication about changes from both sides
- Have a clear process for validating the invoices
- Understand the relationship between cost and quality of service



**What (else) do I  
need to follow up  
on?**

# What do I need to follow up on?

## Things to follow up on:

- What has changed that can affect cost or performance?
  - Machine configuration
  - LPAR configuration
  - Usage patterns
  - Capping
  - Workload Manager
  - zIIP eligible
  - RNI
  - Vertical affinity
  - Response time
  - Cost drivers




# What do I need to follow up on?

## Monitoring the machine configuration

- Hardware model – to determine the MIPS rate
- Total number of processors (CP's and zIIP) – to determine the MIPS rate and LCP to PCP ratio
- Based on SMF70

Daily Total # Physical Processors by Type



## Monitoring the LPAR configurations

LPAR capacity is determined by three things:

- Number of logical processors per LPAR (CP's and zIIP) – determine the maximum capacity
- LPAR weight – to determine the guaranteed capacity
- Capping - keeps the CPU usage below a certain level (more on this later)

**Max LPAR Capacity =**

*number of logical processors \* MIPS rate per processor*


**Guaranteed LPAR Capacity =**

*LPAR weight / sum of all weights \* total CPC MIPS rate*


# What do I need to follow up on?

## LPAR avg and max usage:

### versus weight



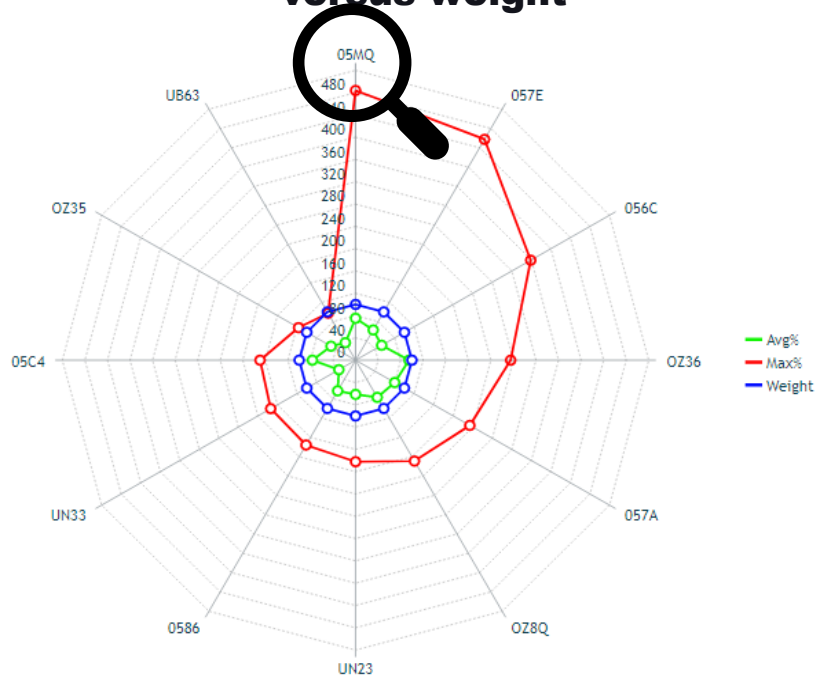
### versus LPC max




# What do I need to follow up on?

## LPAR avg and max usage:

### versus weight




### versus LPC max



# What do I need to follow up on?

## Usage versus configuration


- The MIPS usage is often going way over the guaranteed capacity (weight).
- This (currently) doesn't matter because the other LPARS (or customers) on the machine aren't using their capacity, so this LPAR may use the spare capacity up to its max.
- The spikes are primarily driven by low importance workload. (SMF72)
- WLM, if set up correctly, will help some, but a long period where this LPAR is held to its weight will be a challenge!
- Further analysis can be done by workload, service class, report class, job name, transaction, etc. to understand what exactly is driving the spike. (SMF30, 101, 110, and others)
- Historical analysis will tell us if this is an anomaly or normal.



# What do I need to follow up on?

## Commercial considerations

- What is the customer paying for:
  - Max capacity?
  - Guaranteed capacity?
  - Total MIPS usage?
  - Peak MIPS usage?
- What are the tradeoffs, who makes them and how are they regulated in the contract?
  - More available capacity leads to higher peaks being driven by low priority work
  - Is the customer better off with less capacity?
  - How will Work Load Manager (WLM) handle this?
  - Who is responsible if the LPAR configuration leads to performance problems?




### Capping

Capping determines the LPAR's maximum capacity

- Soft capping is based on your R4HA, so you can go over your capping for short periods
- Hard capping holds you below your limit

Capping can also be applied to a specific workloads (resource groups).



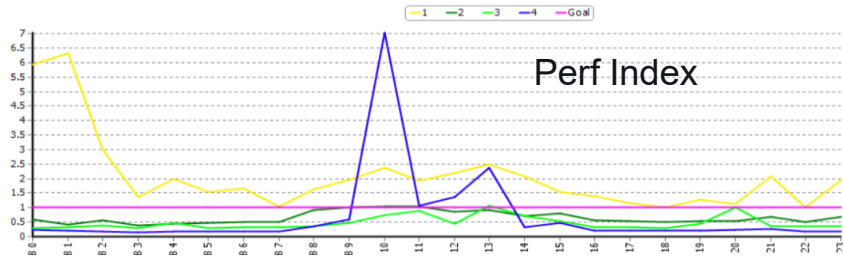
Workload Manager (WLM) determines what gets prioritized when you don't have enough capacity. SMF72



# What do I need to follow up on?

## Work Load Manager (WLM)

- WLM prioritizes resources based on 'goals' and 'importance'.
- Performance Index indicates whether the goals were met (1), missed (>1), or overachieved (<1).
- Velocity is an indication of how often work is delayed due to resource constraints. 100% = no delays.
- WLM can also tell us what resources are constrained.
- The outsourcer is dependent on the customer providing input regarding what is important to the business
- Running 'unconstrained' lets you worry less about WLM can be more expensive.



# What do I need to follow up on?

## zIIP eligible

- zIIP usage is typically much cheaper than CP usage
- But if there is not enough zIIP capacity, then work that is 'eligible' to run on a zIIP will be run on a CP (IIPCP crossover).
- Be aware of your billing model. If you are paying for peak usage, then zIIP eligible workload running on a CP in off peak periods may not impact your cost.



# What do I need to follow up on?

## Relative Nest Intensity (RNI)

There are multiple layers of cache on a mainframe. Layers closest to the CPU are the fastest.

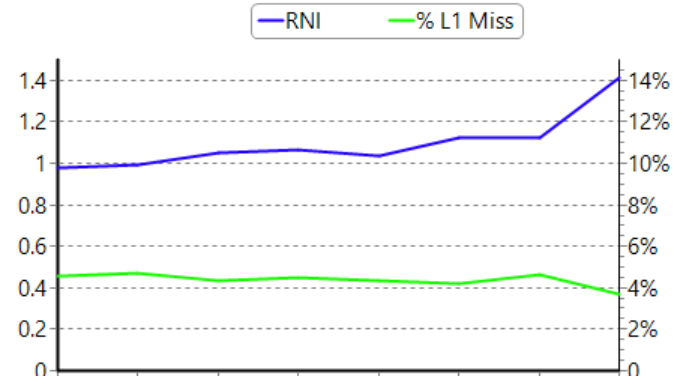
The deeper into the cache or 'nest' you go to find data, the more CPU time is spent waiting.

This depth is referred to as the Relative Nest Intensity (RNI).


SMF 113 gives insight:

- An increased RNI can be due to changes in your usage patterns or applications
- But can also be due to configuration changes
- Significant changes in RNI patterns should be a basis for dialog between the customer and the outsourcer

Hourly RNI (Relative Nest Intensity) and % Level 1 Cache Miss



Hourly L1 Cache Sourcing in % of Total



## Vertical Affinity

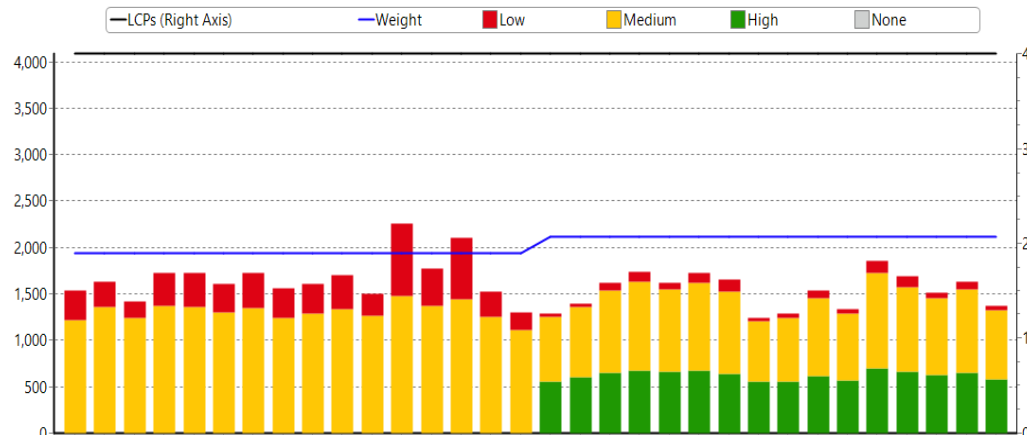
Vertical affinity determines how likely a logical cpu is to be scheduled on the same physical location.

High vertical affinity is good, because it increases the likelihood of finding your data already in cache = low RNI

Vertical affinity is determined by the system based on LPAR weight.

But the tradeoffs can be complex for the outsourcer


Use the vertical affinity reporting as a basis for dialog between the customer and the outsourcer



# What do I need to follow up on?

## Response Time

- Response times should be monitored for subsystems like CICS (SMF 110), IMS (56FA), DB2 (SMF 101)
- Often included in the SLA and WLM goals.
- Slow response times can be related to delays (that can be seen in WLM)
- The ability to understand root causes (application change or config change) is key in the outsourcer/customer relationship.




# What do I need to follow up on?

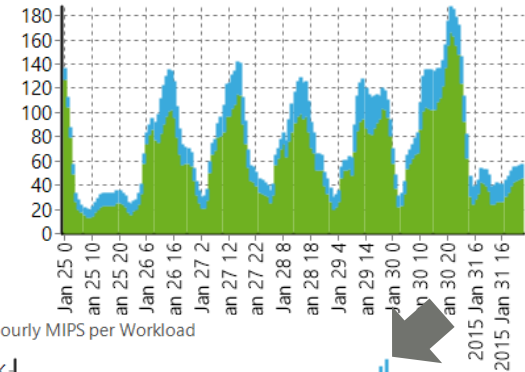
## Cost Drivers – for the customer

- What are you paying for?
- Which workloads, jobs, transactions, etc. are your top consumers relative to that metric? (SMF 72, 30, 110, 101 etc.)
- What applications or business areas do they represent?
- What can you do about it?
  - Reschedule
  - Change configuration
  - Optimize
  - Change user behavior

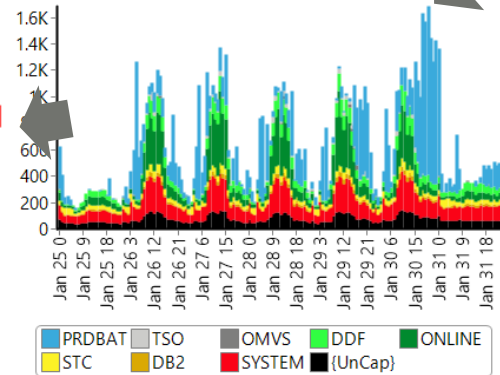
Top-10 MIPS-Users for Selected Period



Hourly R4HA MSU Usage per Measured System



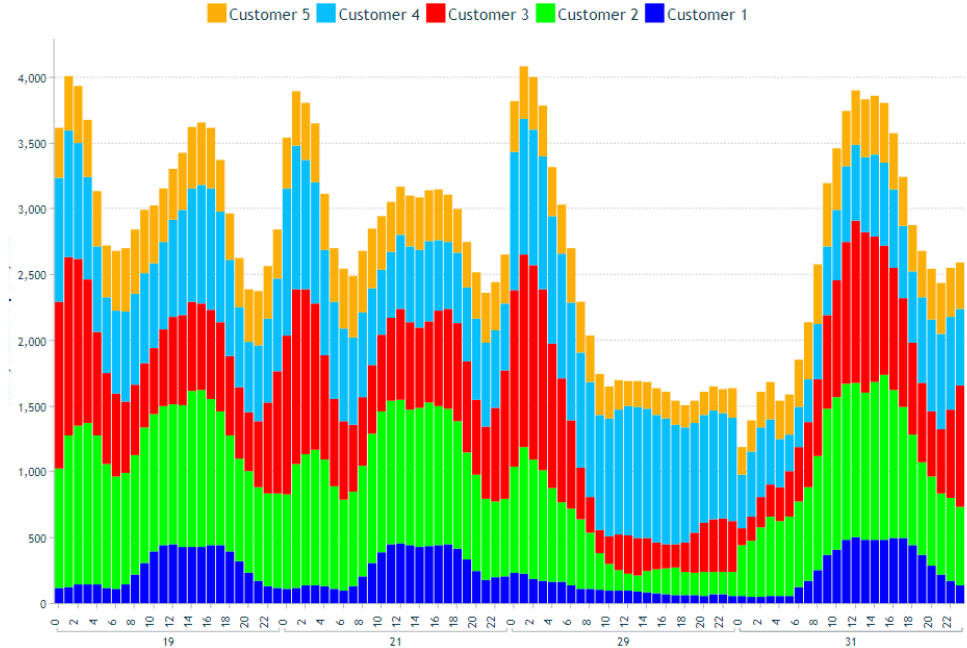
Hourly MIPS per Workload



# What do I need to follow up on?

## Cost Drivers – for the outsourcers

- What are you paying for?
- How are you charging your customers?
- How does each customer contribute to the cost and revenue?
- How to evaluate a potential new customer's impact on cost and revenue?
- How to find a win-win with the customer, e.g. optimization that saves you and the customer money.



# Summary

## Summary

An outsourced environment is often a shared environment, so there will be compromises


Ensure a common understanding of:

- What is the customer are paying for?
- How is it measured or calculated?
- How can it be monitored?
- And what you can do about it?
- What are the tradeoffs?

Have access to good data and reporting tools – ideally ‘self service’.

Look for alignment of interest.

Have access to an independent ‘trusted advisor’ if you don’t have skills in-house.



# Questions?

# Booth 104 at SHARE 2026 in Orlando!



# Your feedback is important!

## Submit a session evaluation for each session you

[www.share.org/evaluation](http://www.share.org/evaluation)

