How Should I Think About Tailored Fit Pricing (TFP)?

Frank Kyne – Watson & Walker

John Baker – IntelliMagic

Session 14848 – August 2022

Agenda

- A Few thoughts on TFP/Contract tips
- Extreme(?) real world customer experiences
- What is 'TFP for Hardware'?

- Visibility into Top CPU Consumers and Drivers of CPU Growth

- Awareness of Common CPU Optimization Opportunities
  - Infrastructure
  - Address Space
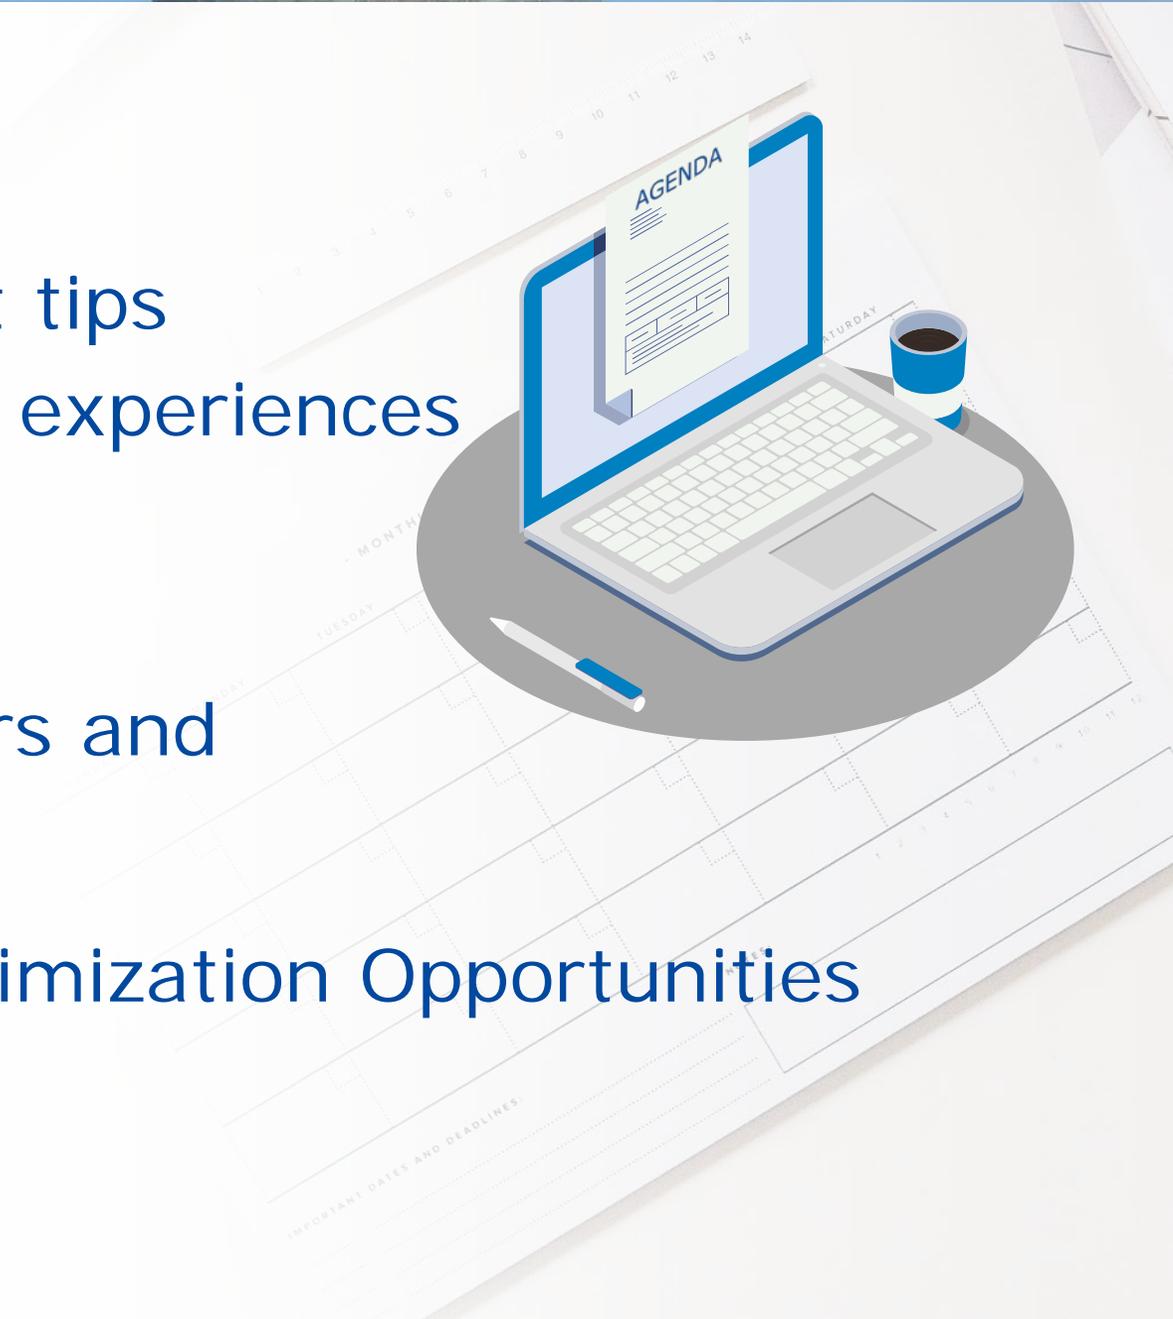
TFP – Is it a good fit for *your* site?

What is TFP – a Recap

- 'Traditional' MLC software bills have been based on the peak Rolling 4-Hour Average MSU consumption each month.
  - Monthly MLC software bill is calculated using Peak R4HA for each product and <u>IBM price list</u>.
- Tailored Fit Pricing is based on your total MLC for 12 months before you started TFP (plus price increases), plus a (discounted) per MSU cost for every MSU over the number of MSUs you consumed in that 12 month baseline period.
  - Agreement includes a baseline MLC that you must pay (plus price increases), <u>regardless of whether you use the corresponding baseline MSUs or not</u>.
  - In return for that commitment, the price per MSU for 'growth MSUs' (that is, those above the baseline) is discounted.
-

Understanding the TFP 'Growth Discount'

CMLC Incremental vs Average Costs

$ per MSU

- CMLC price per MSU
- Avg cost per MSU at Max

MSUs: 4-45, 46-175, 176-315, 316-575, 576-875, 876-1315, 1316-1975, 1976-2499, 2500-3499, 3500-4999, 5000-6999, 7000-9999, 10000-13999, 14000-18999, 19000-24999, 25000+
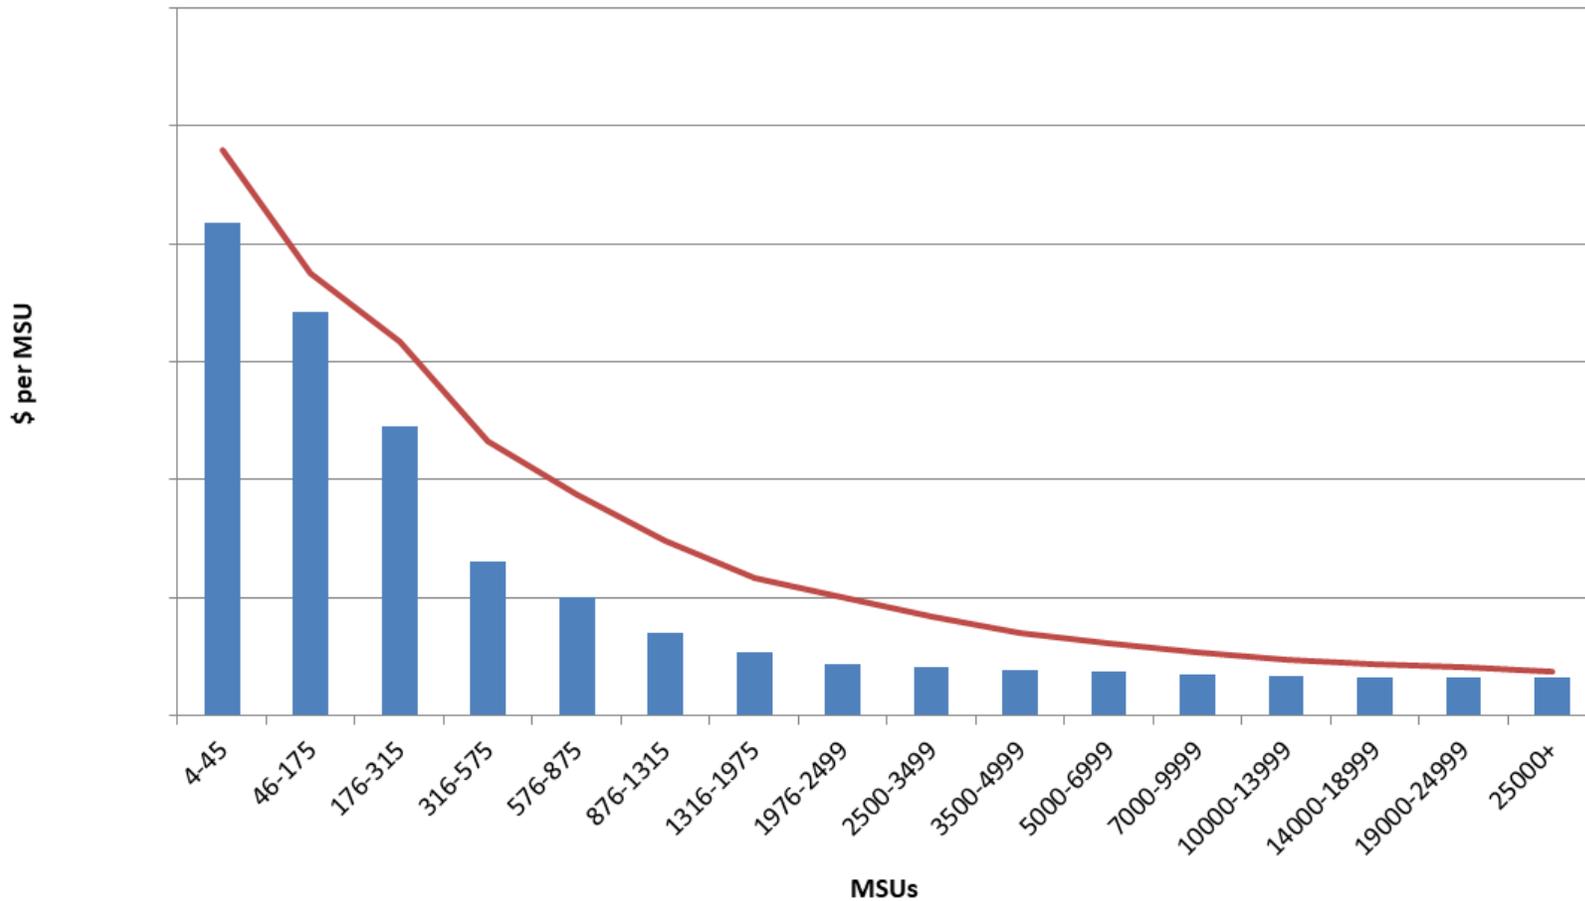
On R4HA metric, the price for additional (incremental) MSUs reduces as you grow, and is always less than the average.

Objective of TFP Growth discount is that growing should cost less on TFP than on R4HA.

Cost of growth

- IBM mentions 50% discount for growth MSUs.

  – This doesn't necessarily mean that the price/MSU for growth MSUs is 50% of the baseline price/MSU.

- You should aim for a growth price/MSU that results in growth costing 50% of what it *would have cost if you stayed on current R4HA-based metric.*

  1. **Model the annual cost of increasing your peak R4HA by x%.**

  2. **Model the annual cost of increasing your accumulated MSUs by x%.**

  – If 2) is not less than 50% of 1), go back to IBM.

Evaluating TFP – is it a good fit for your site ?

- Generally speaking, TFP is more financially attractive for sites whose MSU consumption is *in*creasing.

- Because you **must** pay the baseline MLC, regardless of whether you use the baseline number of MSUs or not, TFP is less likely to be a good fit from a financial point of view if you have a <u>realistic</u> expectation of declining MSU consumption.

- <u>However</u>, there are technical advantages to TFP that *might* be appealing, even if you end up paying more than you would have if you stayed on your current metric.

Evaluating TFP – is it a good fit for your site ?

- Technical advantages of TFP:
  - You can run every IBM MLC and (optionally) IPLA product in every LPAR (<u>if that makes sense</u>) without increasing your MLC costs (but allow for a little MSU growth for additional instances).
    - You should work with IBM when the TFP modeling is done to make sure you understand all the implications.
  - TFP is required if you want to use the 'TFP for Hardware' offering.
  - TFP is a pre-req for 'Flexible Capacity for Cyber Resiliency' on z16.
  - Don't get huge bill if you have one really high R4HA spike.
    - Especially important for IPLA S&S costs.
  - *Might* allow you to cater for same amount of growth with fewer additional value units for IPLA products.

TFP Considerations

- IBM prefer if you use R4HA-based DevTest container for non-production LPARs.
  - Consumption pricing is not suited to DevTest. You want to *encourage* developers to perform comprehensive testing. Penalizing them for doing so is counter-productive.
    - It is also counter-productive to screw down DevTest LPARs *so tight* that the developers spend their day looking out the window! Do you LIKE paying people to do nothing?
  - Having LPARs that contain a mix of production and DevTest work is not really ideal for TFP.

- The 'standard' configuration is one container for production and one (R4HA) for DevTest, but it is possible to negotiate more.
  - For example, if you have traditional production LPARs and zNALC production LPARs, it might make sense to have these in separate containers.

TFP Considerations

- Management of IPLA product costs (value units) when using a TFP agreement is somewhat complex and needs detailed analysis! You do not *have* to include the IPLA products in your TFP agreement, but it *may* make sense for you to do it. Be aware that you must have enough VUs for your IPLA products to cover running in **ALL** LPARs in the TFP container. This may be a problem if you only run some products in a subset of LPARs today.

- ISV products will probably still be licensed based on R4HA metrics, so if you have a lot of ISV products then you may not be able to take advantage of some of the benefits of an MSU consumption model such as being able to remove capping that is used to limit your peak R4HA for software licensing purposes.

- While you can't cap cumulative MSU consumption, the reality is that you don't *not* run production work today – capping simply changes *when* the work will run. In our experience, removing caps on *production* LPARs has little or no impact on cumulative MSUs.

TFP Considerations

- There is **no reliable RoT** to say definitively if TFP would cost you more or less than staying on your current pricing metric.

- The *only* way to know for sure is to model various scenarios (x% growth, flat, x% decline) n years out into the future using a tool that produces prices that have been validated against your IBM Pricing Reports.

TFP – Contract considerations

- What if TFP turns out to be more expensive ?

  – You can return to your previous billing metric (e.g. CMLC) at any stage by giving IBM notice. <u>This should be clearly detailed in your contract.</u>

  – Request a period (6 – 12 months) after the start of your TFP contract where you measure both R4HA and MSU consumption so you can verify that TFP is, at least, no more expensive then if you had stayed on an R4HA metric.

    - Typically, you will need to create two SCRT reports every month, a Multiplex one and a TFP one. IBM will need to calculates bills based on both reports for the agreed period.

    - Agree how IBM will compensate you in the event of TFP being the more expensive model. This could be in the form of a credit etc.

TFP – Contract Considerations

- IPLA considerations:
  - If you include your IPLA products in your TFP agreement you need to get IBM to:
    - Detail the calculations used to convert your existing VU entitlements into an equivalent annual MSU consumption figure so you clearly understand what MSU consumption you are entitled to use for each IPLA product before you would need to acquire more VUs. You should have a table covering your IPLA products for each TFP container.
    - Detail the calculations used to work out how many VUs you would need to acquire in the event your MSU consumption exceeds your entitlements for each IPLA product.

TFP – Contract Considerations

- The contract should clearly describe the formula used to adjust your Price per MSU if you were to *discontinue* an MLC product.
    - **TIP** – make sure to save a copy of your current/old CMP and DevTest contracts so you can calculate the cost if you were to consider dropping back to CMP some time in the future (the CMP uplifts are only documented in the CMP contract).
- The contract should clearly describe the formula used to adjust your Price per MSU if you were to *add* an MLC product or a new chargeable feature on an existing MLC product.

Customer experiences with TFP

Customer experiences with TFP

- We have some clients that think TFP is the best thing ever, and others that are less convinced.

- Let's have a look at two real world, extreme, examples.

- And remember, the 'extreme' happens every day, but once-in-a-1000 years events only happen weekly.

Cust_1 Background

- The first customer is in the finance industry.
- They have a history of making the most of IBM's technology and pricing options to let them support a growing workload without growing costs:
  - Major applications are heavy zIIP users.
  - Used z New Application License Charging to minimize z/OS MLC costs for qualifying applications.
  - Fully optimized environment designed to get the best value from Mobile Workload Pricing.
- They had a slow but steady workload growth and had done all they could to minimize the cost of that growth, resulting in MLC costs starting to trend upwards.

Cust_1 Profile (pre-TFP)

Cust_1 CMP MLC and MSU Consumption
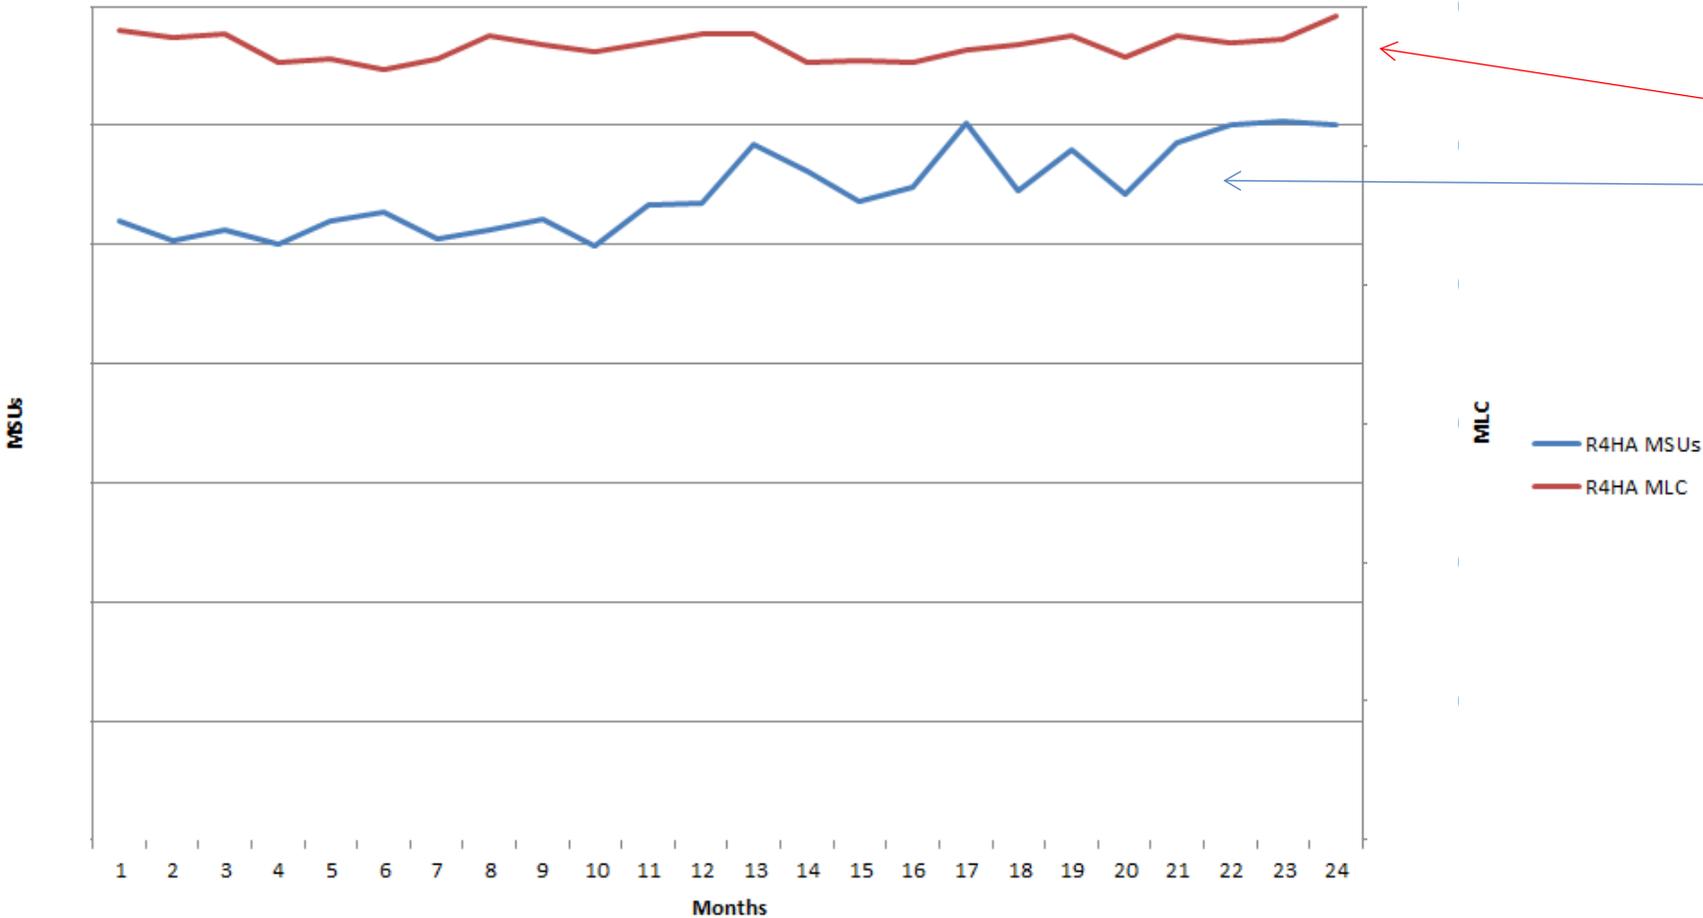
Prior to moving to TFP, MLC was increasing slowly, MSU trend was increasing steadily

Cust_1 Move to TFP

- At this point, they switched to TFP, based on the last 12 months of MLC and Cumulative MSUs.

- Then they had a stroke of good fortune – a competitor went out of business.

- Many of that other company's customers moved their business to Cust_1, driving CPU consumption up far faster than planned.

Cust_1 Post-TFP Profile

Savings!

Cust_1 TFP MLC and MSU Consumption

R4HA MSUs

R4HA MSUs

Cumulative R4HA MLC

Cumulative TFP Cost

MLC
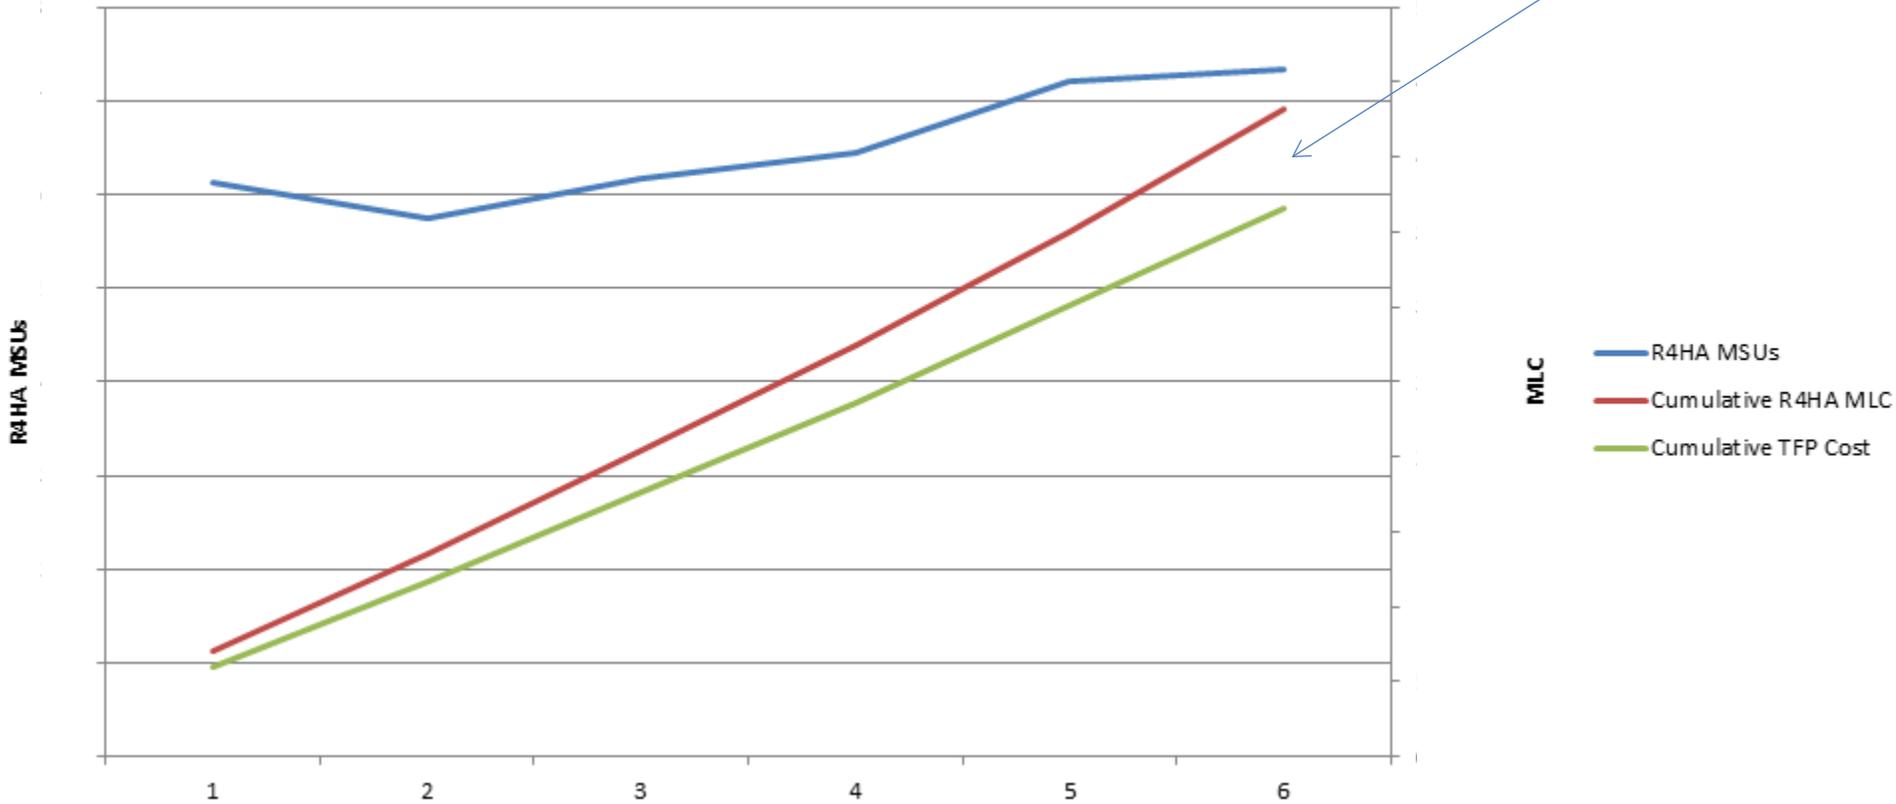
Because of the much faster than expected organic MSU growth, the discount on growth MSUs resulted in a significantly lower MLC than if they had stayed on CMP.

Cust_2 Profile

- Second customer was also experiencing slow but steady growth.
- Little exploitation of 'special' software pricing options (no zNALC, MWP, zCAP, little zIIP exploitation, etc).   As a result, software costs were increasing in line with workload volume growth.

- When presented with discounted costs for future growth offered by TFP, why would you not sign up?

Cust_2 Baseline Period

Cust_2 TFP Baseline Period

Nice, predictable, workload volumes with steady long term upward trend.
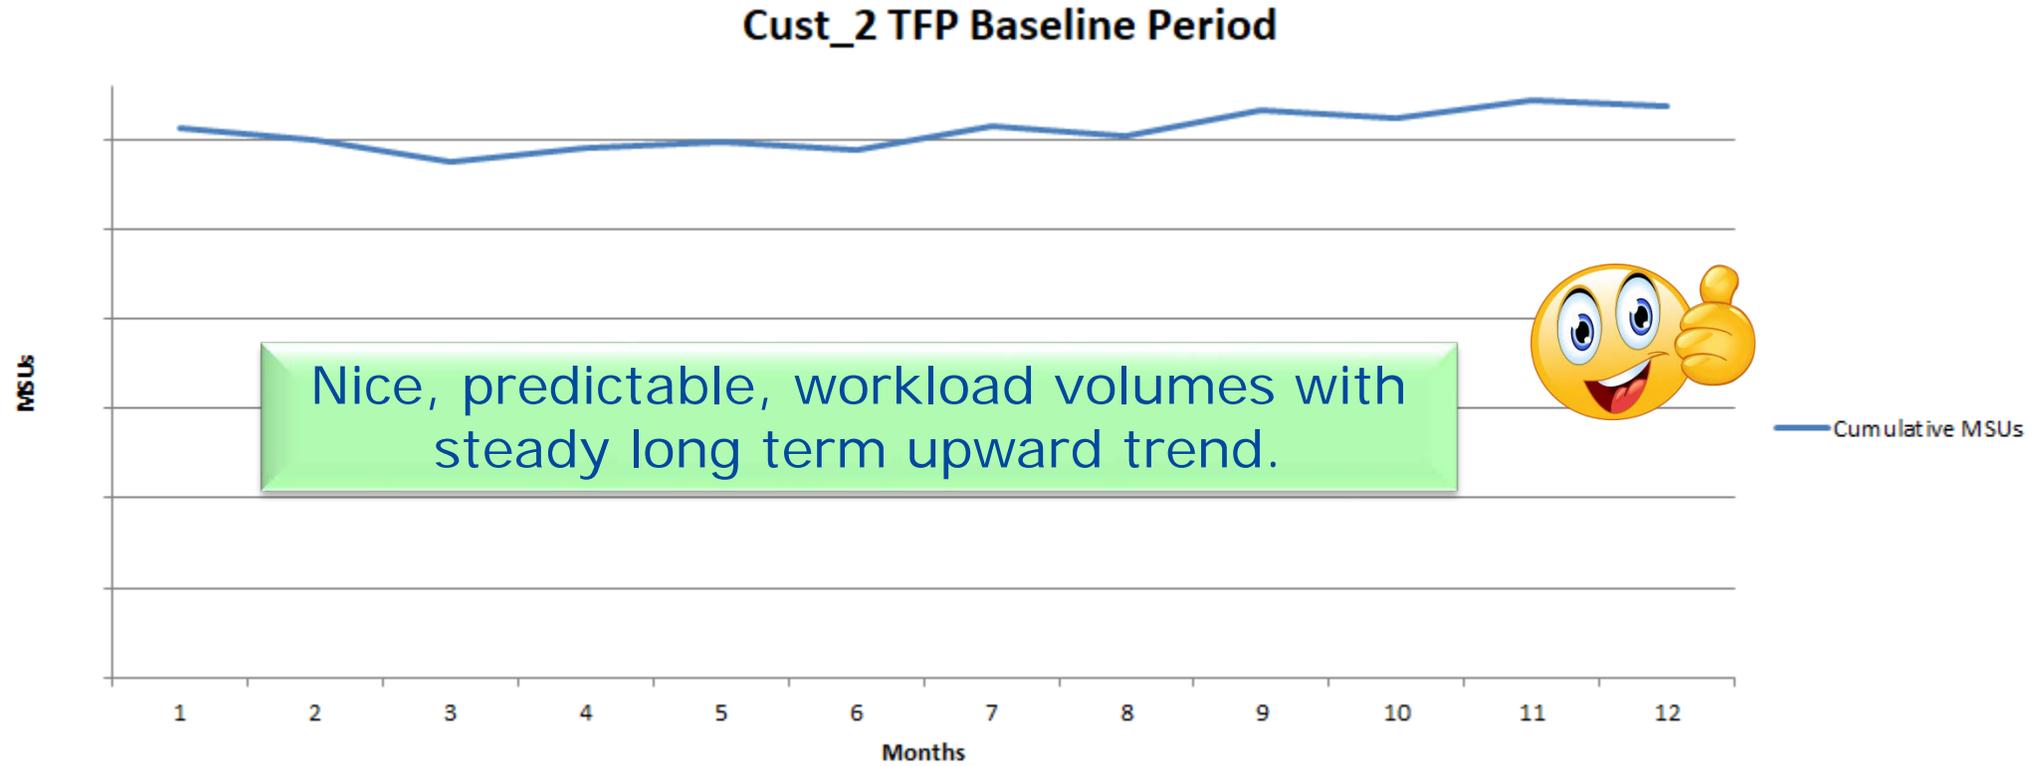
Cust_2 Profile

- And then COVID arrived.

- And customers left. Client was in a industry that was particularly impacted by the global downturn in everything.

- And the return to normal was always going to start 'next month'. If that turned out to be correct, reverting to their previous R4HA-based pricing metric would not make sense.

    – Plus, having reverted to the previous metric, there is a required waiting period before you can move forward to TFP again.

Cust_2 Pre- and Post-TFP Switch
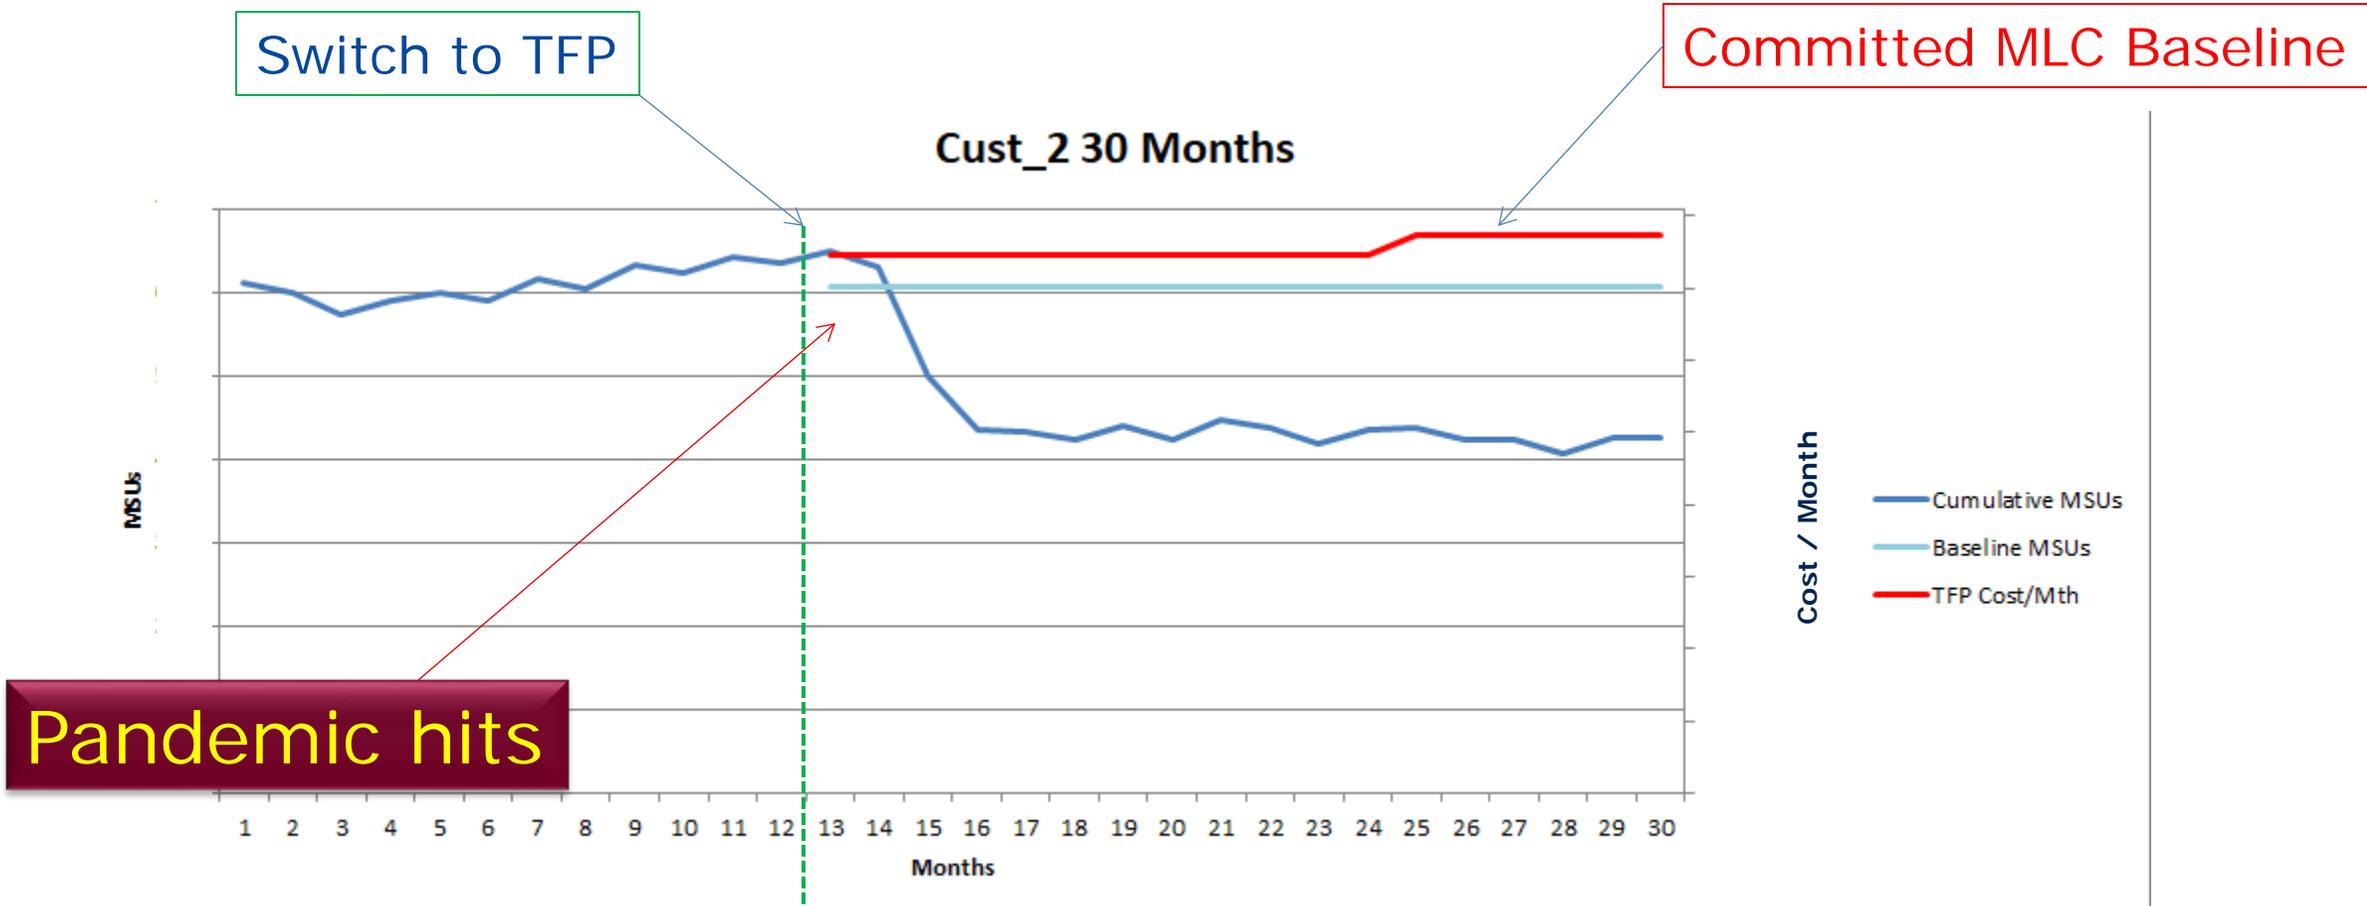
Cust_2 Profile

- The 'that could never happen' happened – a huge, prolonged, and unpredictable reduction in production MSU consumption.

- But they were now committed to the MLC baseline.

- The MSU credit from 2020 could be carried forward to 2021 – but 2021 wasn't any better.

- And the MLC price increase still applied.

- And the end of the ELA was looming.

- And the contract states that any unconsumed MSU 'credit' is lost at the end of the contract period.

Summary

- Both clients were impacted by events that were completely outside their control and unpredicted by nearly anyone.
    - One client was fortunate with their timing.
    - The other was unfortunate.

- Hopefully these were once-in-a-lifetime events.
    - But such events are becoming more frequent.
    - When deciding on TFP, an informed decision should be made about the value of the potential benefit versus the potential cost if MSU consumption drops significantly.

What is TFP for Hardware?

TFP for Hardware

- Very briefly, TFP for Hardware (TFP-HW) is an offering in the Capacity on Demand family of capabilities.

- Rather than you explicitly turning additional capacity on and off, TFP-HW delivers some number of GCPs on top of your normal capacity.

  - Must be purchased in units of whole GCPs, can't buy xxx MSUs.

  - These GCPs are online *all the time.*

  - You don't own them:

    - Pay a subscription fee regardless of whether you ever actually use that capacity or not. The fee includes Maintenance costs.

    - If your total used MSUs in a 15-minute interval exceeds the capacity of your 'normal' GCPs, you pay a per-MSU usage fee.

TFP for Hardware - Considerations

- From a SW pricing perspective, any capacity used on the additional GCPs is treated just like any other MSUs – no dearer, no cheaper.

- Just like all the other CoD offerings, there will be a point where it would have been cheaper to just purchase the additional CPs. This is not intended for use to constantly add capacity to a chronically-undersized CPC.

- If you have IPLA products whose entitlement requirement is not covered by a TFP agreement, that complicates matters.

TFP for Hardware - Considerations

- Speak to your ISVs for their position on this.
- Only available on z15 or later.
- Any supported z/OS release.
- Minimum contract term is 12 months.
- Technical benefit is that additional capacity is available instantly if required for unexpected workload spikes – no delays while you activate OOCoD.
- Depending on how many additional GCPs you get, there may be an MSU reduction due to the efficiency benefits of running at lower utilization.