ZETALY

Understanding Z Pricing to Control your Software Costs

What are monthly licencing charges?

Evolution of Mainframe cost control?

How to reduce software cost with ZETALY Automated Capacity?

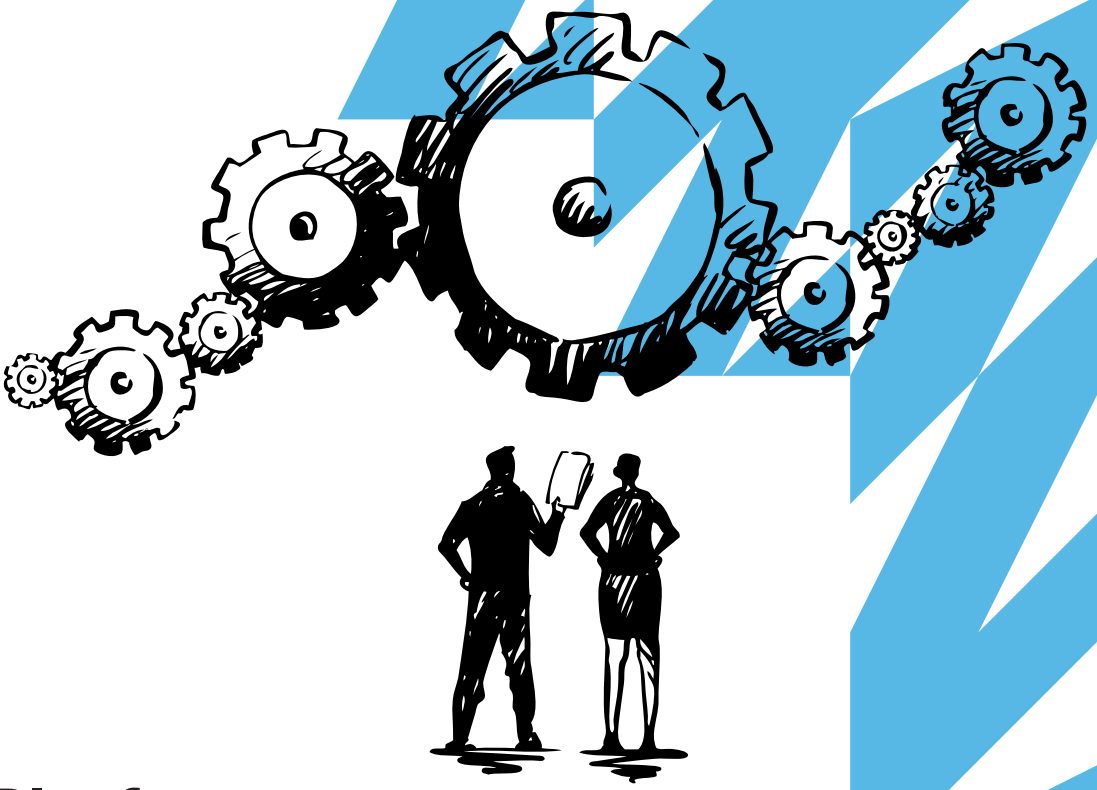

Z Platform Cost Control

Monthly Licensing Charge (MLC) is a prominent z/OS pricing methodology which bills users recurring fees based off of a variety of licensing terms. This broad pricing categorization is comprised of several pricing metrics, each with their own licensing guidelines. But, for the purpose of this guide, we will focus on one pricing metric : Advanced Workload Licensing Charges (aWLC).

The aWLC metric promotes cost-savings thanks to greater invoicing flexibility, and most MLC products are eligible for this pricing metric. Under the aWLC model, monthly invoicing is based off of mainframe usage and not a fixed, potentially outdated, price. Under this pricing model, users pay for the amount of system consumption they experienced in the month, and are not locked into a static cost. aWLC is the “pay for what you use” pricing metric, thus eliminating unnecessary fees.

This guide will help you learn the ins-and-outs of Monthly Licensing Charges and how to control aWLC through resource managwement software such as soft capping.

ZETALY

What are Monthly Licensing Charges?

Put simply, Monthly Licensing Charges (MLC) are recurring costs for renting z/OS software. This broad pricing methodology is comprised of several, varied pricing metrics that users leverage to meet the specific needs of their z/OS environment. Common MLC pricing metrics include :

- **Advanced Workload Licensing Charges**
- **Parallel Sysplex Licensing Charges**
- **Country Multiplex Licensing Charges**
- **And many more**

Most MLC licenses offer sub-capacity options, which means software invoicing is based on logical partition (LPAR) consumption rather than consumption for the whole processor. When it comes to sub-capacity z/OS environments, MLC pricing is always determined by Million Service Unit (MSU) consumption.
For the aWLC pricing metric, billing level for each machine (CPC) is determined by the peak of average usage consumption. This figure is known as the Peak Rolling 4-hour Average (R4HA). The highest R4HA MSU value, for any LPAR, determines aWLC billing level :

- **Base WLC: 3 MSUs**
- **Level 0: 4 - 45 MSUs**
- **Level 1: 46 - 175 MSUs**
- **Level 2: 176 - 315 MSUs**
- **Level 3: 316 - 575 MSUs**
- **Level 4: 576 - 875 MSUs**
- **Level 5: 876 - 1315 MSUs**
- **Level 6: 1316 - 1975 MSUs**
- **Level 7: 1976+ MSUs**

In order to qualify for z/OS sub-capacity savings, you must submit a sub-capacity report to IBM every month which outlines MSU usage. This report is used for software invoicing, and is generated with a Sub-Capacity Reporting Tool (SCRT). SCRT provides licensed capacity for each LPAR and produces a report outlining key information which IBM uses for invoicing.

Evolution of mainframe cost control

There are various cost control products on the market to help optimize performance and reduce mainframe cost. Some popular cost control approaches include :

Software performance optimization

Make adjustments to software performance in areas such as design, source code, build, compile, etc. in order to improve efficiency and reduce resource consumption.

Software urbanism

An overall audit of software utility followed by adjustments to eliminate inefficiencies and redundancies.

Hardware limitation (Hard capping)

Hard capping virtually limits machine capacity to reduce consumption costs. Hard capping sacrifices machine performance in order to reduce monthly costs.

Batch processing and job scheduling Batch processing and job scheduling

Regulate non-priority workloads to later in the day to reduce high billing peaks.

ZETALY

Soft Capping

Software to regulate MSU peaks and reduce sub-capacity charges has taken on several forms, and over the years soft capping products have incrementally improved on the flaws of previous iterations.

In light of new z/OS tools such as Workload Manager (WLM), users were able to manage resources based on workload types. This allowed users to define a performance goal for each workload and prioritize resources based on budget. This led to two of the soft capping methods used today :

Defined Capacity

Fixed MSU limit. Defined Capacity restricts MSU consumption to a static limit. However, this soft capping method has proven ineffective because there is no way to account for variation in machine activity without in-depth SCRT review and manual implementation.

Group Capacity Limit

LPARS are assigned an MSU limit as a group. However, Group Capacity Limit also inhibits mainframe performance. In instances where one LPAR reaches MSU limit, the entire group is throttled which can lead to performance bottlenecking.

What these cost control options have in common, and the cause of their shared inefficiencies, is this: both Defined Capacity and Group Capacity Limit reduce R4HA values by restricting MSU consumption to a fixed limit. Without an intuitive way to account for variable workload management, soft capping products have historically restricted mainframes — effectively sacrificing machine performance for cost-reduction.

That is, until the release of the newest evolution in soft capping known as ZETALY Automated Capacity. This product accounts for workload variability and reduces MLC cost at no detriment to machine performance.

How to reduce software cost with ZETALY Automated Capacity

The newest iteration of soft capping software, and the one best equipped for controlling cost without negatively impacting machine performance, is ZETALY Automated Capacity (ZETALY AC). ZETALY AC addresses the issues of previous soft capping software by automatically and dynamically optimizing capacity limits.

Unlike other soft capping options, ZETALY AC does not focus on reducing R4HA. Instead, ZETALY AC reduces the SCRT value during the billing peak window by dynamically managing machine performance according to a predetermined minimum and maximum sum of Defined Capacities for each CPC (known as CPCMIN and CPCMAX). Essentially, ZETALY AC is designed to ensure the sum of all Defined Capacities, every hour, stays under a predetermined limit in order to avoid high billing peaks.

ZETALY AC employs two cost-control measures to keep your Z Platform within budget without affecting day-to-day performance :

Automatically set Defined Capacity

Prevent performance capping with variable MSU consumption limits. These limits adjust automatically based on partition needs.

Share MSU resources among LPARs

ZETALY AC automatically manages resources among LPARs by borrowing capacity from low-priority LPARs, and giving it to high-priority ones.

Free space on the CPC, which is normally wasted, is automatically aggregated by ZETALY AC into a non-active bank known as a “Phantom LPAR.” The MSU stored in the Phantom LPAR is then tapped according to the MSU needs of each LPAR.

ZETALY AC offers z/OS users a valuable win/win: ZETALY AC reduces MLC, at no detriment to machine performance, by dynamically optimizing Defined Capacity limits.