TRACKING AND REPORTING SOLUTION CONSUMPTION LICENSE CHARGES

Simplified software pricing for qualified solutions



This document may be found on the web, www.ibm.com/support/techdocs.
Search for document number under the category of “White Papers”.

Version 1.1
May 14, 2019

Ed King IBM Corporation
eaking@us.ibm.com
Andrew Sica IBM Corporation
andrewsi@us.ibm.com

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area.

Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785, USA

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERC HANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

Performance data contained in this document was determined in a controlled environment. Therefore the results which may be obtained in other operating environments may vary significantly. No commitment as to your ability to obtain comparable results is any way intended or made by this release of information. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

© Copyright IBM Corporation, 2019

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

COPYRIGHT LICENSE:

This information might contain sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

TRADEMARKS:

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" (www.ibm.com/legal/copytrade.shtml).

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both, and referenced in this paper:
IBM®
IBM Z®
z/OS®
DB2
CICS
COBOL

A current list of IBM trademarks is available on the web at IBM copyright and trademark information - United States (www.ibm.com/legal/us/en/copytrade.shtml).

Java and all Java-based trademarks and logos are trademarks or registered trademarks of oracle, its affiliates, or both.

© Copyright IBM Corporation, 2019

Acknowledgements

The author of this document would like to acknowledge the following individuals for their contributions to this project:

Andrew Bates	Tailored Fit Pricing for IBM Z Offering Lead
Elianne Bravo	Solution Offering Manager, Tailored Fit Pricing for IBM Z
Horst Sinram	z/OS Workload & Capacity Management
Juergen Kuhn	z/OS Resource Measurement Facility
Kevin Banghart	Z Systems SW Licensing
Dan Peach	Z Software Pricing Strategy
Zachary Miller	Z Software Pricing Strategy
Scott Engleman	z/OS Offering Manager
Andrew Mead	WW Sales, Z System Software
Amy Tobenkin	WW Z Offerings
Cheryl Loughlin	User Experience Research & Design
Karina Campos	User Experience Research & Design

© Copyright IBM Corporation, 2019

Table of Contents

1. Introduction 5
2. Unique Value for Modern Applications 5
3. Built on Tailored Fit Pricing for IBM Z 6
4. How is consumption measured? 7
5. Interpreting the SCRT Report 8
6. Data Collection Requirements for Solution Consumption License Charges 10
7. Summary 11

© Copyright IBM Corporation, 2019

1. Introduction

z/OS® is known for its ability to run many concurrent workloads across one or more z/OS images, with very strong availability, reliability and security characteristics. The platform is known for its very large-scale transaction processing and its ability to support thousands of users and application programs, all simultaneously accessing system resources, terabytes of information in databases, and with very large-bandwidth communications. IBM Z® servers support the world's leading businesses, including worldwide banks, the world's largest insurers, many of the top US retailers and most of the world's largest airlines.

In this environment, there are different methodologies in place for licensing IBM software. Perhaps the most common approach today are those based on sub-capacity pricing. Traditional sub-capacity pricing for IBM Z software is based on the rolling 4-hour average (R4HA) of processor utilization as measured in "Millions of Service Units" (MSU). A rolling 4-hour average measures sustained use of CPC capacity and helps compensate for high and low utilization of short duration. Software costs for Monthly License Charge (MLC) products are based on the peak R4HA experienced during the month.

Under some scenarios and workloads, however, the use of the rolling 4-hour average to determine software license charges can lead to unpredictable software costs that are only indirectly related to overall business volume. For example, a high volume period of several hours can set the peak R4HA for an entire month, even when the majority of the month reflects comparatively low CPC utilization.

With Solution Consumption License Charges (SCLC) IBM introduces a new metric for determining MLC software costs for qualified workloads. SCLC is based on the MSU consumption of the New Application solution instead of the rolling 4-hour average. It delivers a true metered usage model, where the million service units (MSUs) that are consumed are charged at the same per-MSU rate, regardless of hourly peaks and spikes. This can deliver exceptional levels of pricing predictability.

Since SCLC measures the actual MSUs consumed by a software solution over a period of time, it can more directly reflect the total business volume of the workload. Volatile or unpredictable workloads can greatly benefit from not having a single point of high activity disproportionately impact monthly MLC costs.

2. Unique Value for Modern Applications

Many clients struggle to predict the capacity requirements of a new application. SCLC provides a variable model for MLC that can grow or shrink as your business demands in a pay as you go model. Additionally, SCLC provides for predictable costs so you can project future expenses and easily determine the value delivered by your new z/OS solution.

Consumption pricing can provide benefits over traditional sub-capacity pricing in the following scenarios:

- Adding new applications to a system. With traditional sub-capacity pricing, adding a new application to a system may impact the peak rolling 4-hour average, which can drive increased software license costs for all MLC products on the system. Solutions deployed under SCLC are isolated from existing workloads - so there is no direct impact

© Copyright IBM Corporation, 2019

to the rolling 4-hour average.

- Unpredictable workloads. When external factors dictate or influence workloads, it can be difficult to anticipate and plan for increases in processor activity. A large, but temporary, spike in processor activity can drive a higher than usual rolling 4-hour average, resulting in increased software license costs. This challenge can lead clients to consider alternatives for modern applications. SCLC provides a more cost effective model for these workloads, and utilizes the isolation provided by the Tailored Fit Pricing technical framework - keeping modern applications near the critical business data they rely on.

In short, the SCLC metric greatly delivers a predictable, transparent cost structure that can be directly correlated to actual MSU usage.

3. Built on Tailored Fit Pricing for IBM Z

IBM's goal for the Tailored Fit Pricing New Application Solution is that a system programmer can deploy a new solution workload on z/OS where it makes sense for the client's business.

The New Application Solution and MSU consumption model take advantage of the capabilities provided by Tailored Fit Pricing for IBM Z to isolate software solutions from traditional or unrelated workloads and provide standalone solution pricing within the highly virtualized IBM Z environment. Solutions are priced based on their business value, with no direct impact to the cost of unrelated workloads.

Tailored Fit Pricing for IBM Z is centered around two related concepts: a "container" and a "solution":

- Container - A container is a technical construct that:
  - Provides the mechanism for isolating workloads.
  - Defines the scope of an application or workload.
  - Implements the technical framework and behavior within z/OS and SCRT.
  - Note that despite the name, in this context, "container" is unrelated to virtualization technologies such as Docker.
- Solution - A solution is a contractual construct that:
  - Defines the individual workloads that are to be part of the solution.
  - Defines the pricing that is specific to the solution.

Solutions are defined through agreement with IBM and must meet certain qualifying conditions. Approved solutions are issued "solution ids" that are used to classify work via WLM and SCRT. IBM announced three different pricing solutions in November 2017 that utilize the Tailored Fit Pricing framework. For more information, see the IBM Z Software Pricing web page: https://www.ibm.com/it-infrastructure/z/software/pricing

SCLC does not require you to predetermine the container size to understand MLC costs. Instead, the MSU consumption pricing model enables a true pay-as-you-go model. MLC

© Copyright IBM Corporation, 2019

software costs will vary month to month relative to and correlated with the exact compute consumed by your business workload.

Like traditional sub-capacity pricing, MLC charges under SCLC are derived from information reported via the Sub-Capacity Reporting Tool (SCRT). SCRT, along with z/OS enhancements, determines the "Total MSU Consumed" value, on a per-solution basis. This value is the basis for the SCLC metric.

Solutions eligible for SCLC are isolated from impacting the costs of unrelated IBM software; as such, any Tenant Resource Groups (TRGs) and LPARs associated with the solution do not contribute to the traditional R4HA pricing metric. This removal of the Tailored Fit New Application Solution workloads from the R4HA is performed by SCRT.

For example, consider a single hour for an SCLC solution running in a single TRG (TRG1) on LPAR1:

Example hour during reporting period:

R4HA for LPAR1: 1000 MSU (pre-reduction)
R4HA for TRG1 on LPAR1: 100 MSU

MSU consumed in TRG1 on LPAR1: 120 MSU

Total MSU consumed for the month in TRG1: 35,000 MSU

SCRT will reflect a 900 MSU sub-capacity value for LPAR1 during this example hour. This reflects the adjustment to the rolling four-hour value of the LPAR, which is part of the traditional sub-capacity environment.

The new application solution is priced based on SCLC and MSU consumption. As such, SCRT reports the total MSU consumption for the month so IBM can determine the solution charges. In this example, the hour shown contributes 120 MSU to the monthly total MSU consumption of 35,000 MSU.

For additional information about Tailored Fit Pricing for IBM Z, including how to define and setup tenant resource groups for use with Tailored Fit Pricing, see the following whitepaper: https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102719

4. How is consumption measured?

An MSU is a measure of work performed on a system. A solution "consumes" an MSU when it performs equivalent work. To understand how consumption is measured, it is important to understand how Tailored Fit Pricing for IBM Z isolates workloads on the system. Workloads can be isolated in two different ways:

- Entire LPARs can be designated to be part of a container by specifying a special control statement to SCRT. In this case, SCRT removes the impact of all work performed on the LPAR from all other work on the system; the LPAR does not contribute to the rolling 4-hour average of the traditional workload.

© Copyright IBM Corporation, 2019

- Solutions co-located on an LPAR with unrelated work can be isolated through the use of Tenant Resource Groups (TRGs). Tenant Resource Groups are defined and associated with a solution via WLM. SCRT will remove the impact of any work associated with these TRGs from the rolling 4-hour average of the traditional workload.

Whether solutions are isolated as entire LPARs or through the use of TRGs (or a combination of both), SCRT calculates and reports the "Total MSU Consumed" on a per solution basis using information contained in the SMF 70 Subtype 1 records. SCRT aggregates SMF data on an hourly basis, computing the MSUs consumed for each hour. The hours are then summed to determine the total MSU consumed over the reporting period. Importantly, Solution Consumption is not measured on a per product basis. The solution is priced and reported as a whole.

Consumption is a measure of total CPU (on general purpose processors) scaled to a machine agnostic MSU value. Under the SCLC metric, if a particular solution consumes 50 MSUs in the first hour, 100 MSUs in the second hour and 50 MSUs in the third hour, the total number of chargeable MSUs for that three-hour period would be 200 MSUs. Hourly periods continue to be calculated this way over the entire month, providing a true usage-based cost model. SCRT computes the MSU consumption as follows:

For a full LPAR, MSU consumption is calculated as:

(SMF70EDT / 1000000) * 16 * SMF70CPA_scaling_factor / SMF70CPA_actual

- SMF70EDT is the effective dispatch time for the LPAR.
- SMF70CPA_scaling_factor and SMF70CPA_actual are the physical CPU adjustment factors based on the model capacity rating of the CPC. The SMF70EDT value may be adjusted by wait time in the case of dedicated processors or processors that are wait complete.

For Tenant Resource Group workloads, MSU consumption is calculated as:

***CP seconds* * 16 * SMF70CPA_scaling_factor / SMF70CPA_actual**

Where "*CP seconds"* is calculated as

SMF70_TRG_SUCP / *service units per second*

SMF70_TRG_SUCP is the service units consumed by general purpose processors by the Tenant Resource Group. *Service units per second* is calculated as

16000000 / SMF70ADJ

SMF70ADJ is the logical adjustment factor CPU rate.

5. Interpreting the SCRT Report

As in traditional MLC pricing, usage of the Sub-Capacity Reporting Tool (SCRT) is required for SCLC billing.

© Copyright IBM Corporation, 2019

SCRT reports on every solution defined to a machine (or to a group of machines if Country Multiplex Pricing is in effect) in a number of container related sections. The following sections are all reported per-container; if multiple containers are present, the sections will repeat for each container.

The container header section is first, followed by additional sections with more detailed information.

```
==CPS1============================================================

Solution ID                    Z111111-N31BB29-8FC80FDC07-NSDTZZZZ-91CC-465B-98CA-0565FF-E42B4F
Solution Name                  DevTest Solution
Peak Four Hour Rolling Average                                                  13893
Total MSU Consumption                                                         3581681
```

In this example, the container header section is identified by CPS1. The CPS1 identifier is generated by SCRT and maps one-to-one with the longer solution id. The container header section also reports the full solution id, the solution name, the peak four hour rolling average, and the total MSUs consumed by the solution on the machine. The "Total MSU Consumption" value provides the basis for SCLC billing.

Section T6 of the SCRT report shows individual contributions (by LPAR or TRG) to MSU consumption:

```
==T6==============================================================
CONTAINER MSU CONSUMPTION CONTRIBUTORS

CPS1          Z111111-N31BB29-8FC80FDC07-NSDTZZZZ-91CC-465B-98CA-0565FF-E42B4F

              TRG             Contribution to Total Consumed

LPAR1         *                                      2856297
LPAR2         TRG1                                    725384

CPS1                                                 3581681

* indicates the entire LPAR was configured as part of the Container
```

A detailed breakdown of MSU consumption can be found by examining the optional CPSTRACE output from SCRT. The CPSTRACE output contains two sections, "CPSTRACE 1" and "CPSTRACE 2". Both sections show Tailored Fit Pricing information on a per solution, per hour, per TRG or LPAR basis. The "CPSTRACE 1" section shows the Rolling Four Hour Average values, while the "CPSTRACE 2" section shows the MSU consumption values.

© Copyright IBM Corporation, 2019

==CPSTRACE 2==========

MSU Consumption

CPS1	DevTest Solution	Z111111-N31BB29-8FC80FDC07-NSDTZZZZ-91CC-465B-98CA-0565FF-E42B4F					
CPS1	Date Time	Processor	Partition -	LPAR1	LPAR2	CPC	Container
CPS1		(tttt-sssss)	TRG -	(lpar)	TRG1	Total	Total
CPS1	02 Aug 2018 - 00:00	3906-12345		5803	197	6000	6000
CPS1	02 Aug 2018 - 01:00	3906-12345		3330	810	4140	4140
CPS1	02 Aug 2018 - 02:00	3906-12345		1423	93	1516	1516
CPS1	02 Aug 2018 - 03:00	3906-12345		627	1024	1651	1651
CPS1	02 Aug 2018 - 04:00	3906-12345		2529	1437	3966	3966
CPS1	02 Aug 2018 - 05:00	3906-12345		3137	1402	4539	4539

In this example, we see that the container identified by CPS1 runs as a full LPAR solution on LPAR1 and in Tenant Resource Group TRG1 on LPAR2. The CPSTRACE 2 report shows the hourly MSU consumption per container, broken down by CPC, LPAR, and TRG. The sum of the hourly totals is the total MSU consumed for the reporting period.

6. Data Collection Requirements for Solution Consumption License Charges

IBM sub-capacity terms require the collection and reporting of relevant operating system data (generally, SMF 70 Subtype 1, SMF 89 Subtype 1, and SMF 89 Subtype 2 records) for the entirety of the report period. These requirements remain unchanged for SCLC, but the MSU consumption pricing model is more sensitive to unreported or missing data. As such, SCRT will report additional information regarding gaps in data, specifically at the sub-hour level and justification may be required for significant data gaps. SCRT will report this information in section H7 on a per container basis.

Section H7 indicates, on a per-container, per-LPAR basis, the number of minutes the solution could have been running, the number of minutes actually collected for the container, and the difference between the two, which is the number of missing (unreported) minutes. Section H7 is only produced for containers that are based on an MSU consumption metric. The solution id assigned to a container indicates whether or not the solution is consumption based.

Example 1:

New Application Solution SOLUTIONA is associated with Tenant Resource Groups TRG1 and TRG2 on LPAR1. Both TRGs run simultaneously for 3 hours during the report period. All SMF70 records are collected for the entire report period.

Total possible minutes: 180
Total collected minutes: 180
Total missing minutes: 0

Example 2:
New Application Solution SOLUTIONA is associated with Tenant Resource Groups TRG1 and TRG2 on LPAR1. TRG1 runs for 3 hours during the report period and TRG2 runs in a different 3 hours during the report period. All SMF70 records are collected for the entire report period.

Total possible minutes: 360
Total collected minutes: 360
Total missing minutes: 0

© Copyright IBM Corporation, 2019

Example 3:
New Application Solution SOLUTIONA is associated with Tenant Resource Group TRG1. TRG1 runs for 3 hours during the report period. A total of two fifteen-minute intervals are missing from the SMF70 records provided to SCRT. Any TRG, and thus any Tailored Fit Solution including SOLUTIONA, could have been running during the 30 minutes of missing SMF70 data. Those 30 minutes are included in the "Total possible minutes" and are indicated as missing.

Total possible minutes: 390
Total collected minutes: 360
Total missing minutes: 30

Section H7 of the SCRT reports this information, and may require justification for missing data over IBM designated thresholds:

==H7==========
MSU CONSUMPTION DATA COLLECTION

LPAR	Total Possible Minutes	Minutes Collected	Missing Minutes	Justification
LPAR1	44640	44580	60	(required)
LPAR2	44640	44580	60	(required)

Note, Section H7 does not report:

- LPARs that are not missing any minutes
- Hours that were already reported in Section H4 or Section H6
- Hours that were excluded through the use of an EXCLUDE control statement

This example shows that LPAR1 and LPAR2 were both missing a total of 60 minutes over the course of the reporting period, during which the associated solution could have running and accumulating MSU consumption.

The SCRT report is also enhanced with a new “detailed data” section that is enabled via the GENERATE_DETAILED_DATA control statement (or GUI or CLI equivalent). Section V6 shows an hour-by-hour breakdown of missing minutes that can be cross-referenced with your operational data to assist in understanding why the data was unreported.

==V6==========
SUB-HOUR DATA COLLECTION DETAILS - MINUTES MISSING

Date/Time	LPAR1	LPAR2
02 Aug 2018 - 20:00	0	30
03 Aug 2018 - 16:00	30	30
05 Aug 2018 - 09:00	30	0

7. Summary

Tailored Fit Pricing for IBM Z provides simplified software pricing for qualified solutions.

© Copyright IBM Corporation, 2019

With the addition of Solution Consumption License Charges (SCLC), IBM delivers a predictable and transparent method of measuring the cost of an application.

SCLC allows a client to easily understand the ROI of a New Application Solution using SCLC. As the workload grows and consumes MSUs, the client can predict the impact to their IBM software costs - something simply not possible with the rolling four-hour average.

© Copyright IBM Corporation, 2019