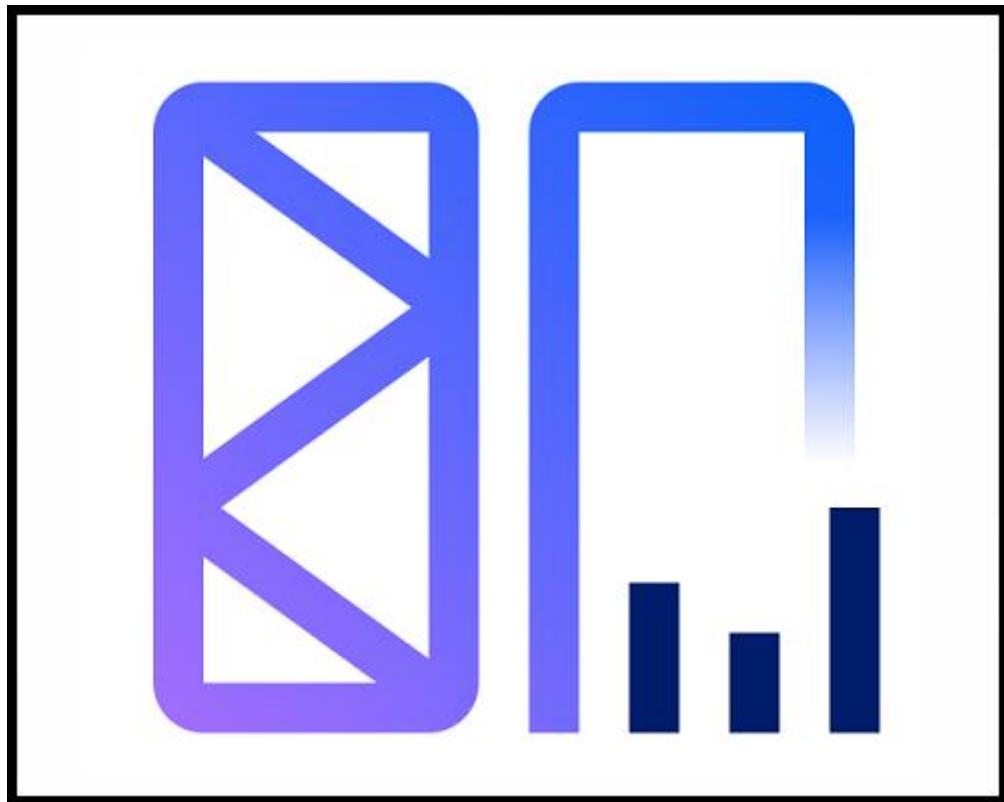


IBM z Processor Capacity Reference

for

IBM Z and **IBM LinuxONE**



© Copyright IBM Corp. 2003, 2025

Version 9.7.4

v974 zPCR UG 2025a03.docx

July 7, 2025

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

CICS®	Parallel Sysplex®
Db2®	System z®
DB2®	System z9®
eServer™	System z10®
IBM®	S/390®
IBM Z®	System/390®
IBM zEnterprise®	WebSphere®
IBM z Systems®	z9®
IBM z13®	z13®
IBM z13s®	z13s®
IBM z14®	z14®
IBM z15®	z15®
IBM z16™	z16™
IBM z17™	z17™
LinuxONE™	z/Architecture®
LinuxONE Emperor™	z/OS®
LinuxONE Rockhopper™	z/VM®
IMS™	z/VSE®
MVS™	zEnterprise®
OS/390®	z Systems®

The following are trademarks or registered trademarks of other companies.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

InstallShield 2023 Premier Edition is a registered trademark of Flexera Software, Incorporated.

All other products may be trademarks or registered trademarks of their respective companies.

Note to U.S. Government Users

Documentation related to restricted rights - Use, duplication or disclosure is subject to restrictions set forth in GSA SDP Schedule Contract with IBM Corp.

Table of Contents

Introduction.....	1
Getting Started.....	7
Registration.....	9
Usage Tips.....	10
User Controls.....	12
Logo Window.....	19
Control Panel	21
Preferences	28
Reference-CPU.....	30
How the Scaling-Factor Is Used	33
Provisional Reference-CPU.....	34
Workloads	37
LSPR Processor Capacity Ratios.....	42
IBM GP Multi-Image Table.....	42
IBM IFL Multi-Image Table.....	43
LinuxONE Multi-Image Table.....	44
LSPR Capacity Ratio Tables Overview.....	45
LSPR Table Control	52
Find Similar CPCs	53
Favorites List	55
IBM Z Single-Image Table	58
Charts and Graphs	59
LPAR Configuration Capacity Planning.....	63
LPAR Host and Partition Configuration.....	64
Creating LPAR Configuration from EDF.....	69
Creating LPAR Configuration from RMF	70
Creating LPAR Configuration from a zPCR Study	71
LPAR Host Processor	72
Power-Save Mode	75
Absolute Capping	75
Nonstandard zAAP or zIIP Configurations	76
zAAP on zIIP Capability	76
LPAR Host Hardware Designations	76
Partition Definition	95
EDF Overview	106
EDF Interval Selection.....	110
Get Partitions from EDF	113
RMF Overview	123
RMF Interval Selection	125
Get Partitions from RMF	128
Create Host and Partitions from Study	136

zAAP/zIIP Capacity Considerations.....	139
zAAP (Application Assist Processor).....	139
zIIP (Integrated Information Processor)	140
zAAP on zIIP Capability	140
General Comments.....	141
Host Summary Report.....	142
Partition Detail Report	147
Reserve Capacity	159
Book and Drawer Considerations	161
System Recovery Boost	164
Absolute Capping	166
SMT Benefit	168
LCP Alternatives.....	176
zAAP/zIIP Loading	179
HiperDispatch.....	183
Topology Report.....	186
<i>Partition View Controls</i> group box.....	187
<i>Partition Summary by Pool</i> group box	188
Modify SCP/Workload	189
Calibrate Capacity to LPAR Host.....	192
Utilized Capacity Report.....	194
Control Panel - Advanced Usage	197
Host Capacity Comparison Summary.....	209
Host Capacity Comparison.....	212
Partition Capacity Comparison	215
Optimizing LPAR Configuration Capacity.....	216
Host Margin-of-Error.....	219
Partition Margin-of-Error	221
Copy Partitions to an LPAR Configuration	223
Move Partitions between Configurations	228
Accuracy of LPAR Capacity Projections	235
CP Calculator	239
zAAP Capacity Estimator Input.....	240
zAAP Capacity Estimator Report.....	245
Support	249
External Study File	250
Definition of Terms	251

Introduction

IBM z Processor Capacity Reference (IBM zPCR) is a PC-based productivity tool that runs on Windows and macOS. It is designed to provide capacity planning insight for **IBM Z** and **IBM LinuxONE** processors running various SCP/workload environments using various LPAR partition configurations. Capacity results are based primarily on IBM's published LSPR data for **IBM Z** families. Generally, we have found the accuracy of using **IBM zPCR** with an appropriate z/OS workload category to be well within the range of $\pm 5\%$ for reasonable LPAR configurations.

IBM zPCR (zPCR) is based on a single version of LSPR data, normally published at the time of a new IBM mainframe processor announcement. The most recent LSPR data, identified as **z/OS-3.1**, adds the **IBM z17** to the entire set of recent IBM mainframe families. The specific families included are shown in the table below.

IBM Z Processor Families				UG Key
High End	M/T	Mid-Range	M/T	
z17-ME1	(9175)			z17
z16-A01	(3931)	z16-A02/AGZ	(3932)	→ z16
z15-T01	(8561)	z15-T02	(8562)	→ z15
z14	(3906)	z14-ZR1	(3907)	→ z14
z13	(2964)	z13s	(2965)	→ z13
z12	(2827)	zBC12	(2828)	→ z12
z196	(2817)	z114	(2818)	→ z11
z10 EC	(2097)	z10 BC	(2098)	→ z10
z9 EC	(2094)	z9 BC	(2096)	→ z9

IBM LinuxONE Processor Families				UG Key
High End	M/T	Mid-Range	M/T	
Emperor 5	(9175)			→ z17
Emperor 4	(3931)	Rockhopper 4	(3932)	→ z16
LinuxONE III	(8561)	LinuxONE III	(8562)	→ z15
Emperor II	(3906)	Rockhopper II	(3907)	→ z14
Emperor	(2964)	Rockhopper	(2965)	→ z13
		Rockhopper	(2828)	→ z12

Note: **UG Key** identifies the processor family group names used throughout this User's Guide documentation.

LSPR capacity values used by **zPCR** are based on the most recent set of LSPR data, which includes various workload environments run under z/OS-3.1. Additional LSPR measurements are made for z/VM (a multiple Linux guest workload) and for native Linux (a single workload). For **zPCR** capacity planning purposes, the z/OS measurement data is considered to be representative for z/VM, z/VSE, KVM, and Linux.

For z/OS, LSPR capacity ratios representing three workload categories are published for each processor model. To increase granularity for capacity planning purposes, **zPCR** defines two additional workload categories. The assigned names are:

- **Low** represents light use of the processor's memory hierarchy
- **Low-Avg** Defined midway between **Low** and **Average** (**zPCR** only)
- **Average** represents average use of the processor's memory hierarchy
- **Avg-High** Defined midway between **Average** and **High** (**zPCR** only)
- **High** represents heavy use of the processor's memory hierarchy

The capacity numbers representing these workload categories are based on LSPR workload primitives that are measured. An analysis of various production workload environments is then matched to specific mix combinations of these primitives to define the workload categories. These workload categories may also be assumed to be representative for z/VM and Linux production workloads.

z/OS and z/VM production workloads running on **IBM z17, z16, z15, z14, z13, z13, z12, z11, and z10** processors can be characterized using **CPU MF** (CPU Measurement Facility hardware counter data). The resulting data is then summarized into an EDF (**Enterprise Data File**) to be read by **zPCR**. EDFs that include **CPU MF** data will be used by **zPCR** to automatically map the production workload of a partition to one of the 3 measured workload categories mentioned above.

When CPU MF is not available, the default workload **Average** will be assigned.

The workload categories can be displayed in the **LSPR Processor Capacity Ratios** tables. These same workloads can be defined to individual partitions when using **zPCR**'s **LPAR Configuration Capacity Planning** function.

Two purposes are served by **zPCR**:

1. LSPR Processor Capacity Ratios Tables

These tables provide capacity ratios for various SCP/workload environments displayed side-by-side. Specific workload categories and the order presented is user controlled. Capacity ratios are provided for General Purpose CPs (support any SCP type) and IFL CPs (support z/VM with Linux guests and native Linux only).

Two forms of processor **LSPR Processor Capacity Ratios** tables are provided, **Multi-Image** and **Single-Image**. Capacity data is entirely based on a single version of z/OS.

- The **Multi-Image LSPR** table provides capacity relationships for processors running with multiple partitions that are assumed to exploit all the available RCPs. While the table is based on z/OS LSPR data, it can also be assumed to be representative for z/VM and Linux environments. The table can be presented with either General Purpose CPs or IFL CPs. Capacity relationships portrayed assume configurations of multiple partitions deemed typical for the N-way of the model. Only General Purpose, IFL, and ICF partitions are represented (zAAP and zIIP specialty engines are not considered). The same workload is assumed to be running in every General Purpose or IFL partition. A 1-way processor is considered to be running with five partitions. Detail concerning the specific partition configurations represented for the other N-way models is not provided. However, as the number of RCPs increases, the number of partitions increases, the average number of partition LCPs increases, and the LCP to RCP ratio diminishes.

Note: The **Multi-Image LSPR Processor Capacity Ratios** table serves as a generalization of capacity expectation for processors with typical, multiple-partition configurations, all running the same SCP and workload. The **LPAR Configuration Capacity Planning** function (discussed below) should be considered as a more accurate way to assess the capacity expectation for a specific LPAR Host including its RCPs and types, each partition including SCP/workload environment, mode, LCPs, weight, and capping assignment. Also considered is the association of zAAP or zIIP LCPs with a parent General Purpose partition.

- The **Single-Image LSPR** table relates capacity for a single copy of the SCP (1 partition) managing all the RCPs. The table can be presented with General Purpose CPs or IFLs. While (internally) the table contains values up to the maximum CPs that can be configured, capacity values are only displayed for reasonable single-image configurations (i.e., 30 CPs maximum).

The Single-Image values represent all the capacity data that is considered measured for LSPR purposes. These values are used as the capacity basis for both the **Multi-Image LSPR** table and the **LPAR Configuration Capacity Planning** function.

Note: The **Single-Image LSPR Processor Capacity Ratios** table serves only as a representation of capacity expectation for processors running a single partition with a particular SCP and workload. This table is intended primarily to support the **LPAR Configuration Capacity Planning** function, described below.

2. LPAR Configuration Capacity Planning

This function is designed to project the capacity expectation for a specific partition configuration on a specific LPAR host processor. Capacity results are provided for each individual partition and for the LPAR host as a whole.

The LPAR host processor is selected and configured with General Purpose CPs, zAAPs, zIIPs, IFLs, and/or ICFs, where supported. Then each partition is defined, specifying type (General Purpose, IFL, or ICF), SCP/workload, LP configuration (dedicated/shared with number of CPs), and weight/CAP assignments. zAAP and zIIP partitions are always associated with a parent General Purpose partition. IFL partitions may be associated with a parent General Purpose z/VM partition.

The LPAR host can be assigned as any **IBM Z** or **LinuxONE** processor model. z/OS, z/VM, z/VSE, KVM, Linux, zAware*, zACI*, SSC, or CFCC may be defined as the operating system to any appropriate partition.

* Note: zAware can be only be assigned to an **IBM z12** or **z13** partition. zACI can be only be assigned to an **IBM z13** partition. SSC is intended to replace these SCPs on **IBM z13** and later processor models.

The **LPAR Configuration Capacity Planning** function allows multiple LPAR configurations to be defined. Existing LPAR configurations can be cloned and subsequently modified. Windows are available with which to make direct comparisons of their capacity. In cases where the name would be a duplicate, a repetition index indicator is appended.

Partition configurations can be created in **zPCR** several ways:

- Automatically from a z/OS or z/VM generated EDF
- Automatically from a z/OS generated RMF CPU Activity report
- Manual entry
- Cloning a previously created LPAR host and its partition configuration

zPCR uses the concept of a **Reference-CPU** for the purpose of scaling capacity results to a common base. A General Purpose 1-way processor model and a scaling-factor are required (default and typical settings are provided). The scaling-factor is always interpreted to be the productive capacity of a shared single-partition configuration. A scaling-metric may also be supplied.

The assignment of a commonly accepted processor model and scaling-factor/metric can help in understanding capacity results and associating them with various hardware configurations and capacity tables. A **Reference-CPU** setting that is currently popular is the **2094-701** rated at **593 MIPS**.

The scaling-factor assigned to the **Reference-CPU** is adjusted internally, based on the function being accessed.

- For **LSPR Single-Image Capacity Ratios** table purposes, the scaling-factor is used directly, representing the capacity of a shared single-partition configuration on that particular 1-way **Reference-CPU** model.
- For **LSPR Multi-Image Capacity Ratios** table purposes, the scaling-factor is adjusted to represent the capacity of a typical shared 5-partition configuration on the 1-way **Reference-CPU** model. For information concerning this adjustment, see the **LSPR FAQ** paper, included with **zPCR**.
- For the **LPAR configuration Capacity Planning** purposes, the scaling-factor is adjusted to represent the capacity of the entire 1-way **Reference-CPU** model, without regard to any partitioning costs. This is the necessary starting point for the algorithms as implemented.

When comparing capacity between configurations, it does not matter what **Reference-CPU** settings are used. Capacity ratios are unaffected as long as the same **Reference-CPU** settings are used for each configuration being compared.

zPCR results shown on windows can be captured as tables and graphs for notes, presentations, or handouts. Complete studies can be captured for future reference.

zPCR currently supports all IBM **Z** processors (starting with the z9) and **LinuxONE** processor (starting with the z12). Features common to these models are separate RCP pools for each CP type (i.e., General Purpose, zAAP, zIIP, IFL, and ICF).

The number of partitions that may be defined depends on the processor model:

High end models

- IBM z17, z16, z15, z14, and z13: up to 85 partitions
- zEC12, z196, z10EC, and z9 EC: up to 60 partitions

Mid-range models

- IBM z16, z15, z14-ZR1, and z13s: up to 40 partitions
- zBC12, z114, z10 BC, z9 BC: up to 30 partitions

For capacity planning purposes, zAAP, zIIP, IFL, and ICF LCPs associated with a parent General Purpose partition are considered as separate partitions. However, associated partitions do not count against the maximum number of partitions supported on the defined LPAR host.

The multi-image and single-image LSPR data is generally updated only when a new processor announcement is made. However, the algorithms supporting the **LPAR Configuration Capacity Planning** function may change more frequently, as additional experience is gained. **To obtain legitimate capacity comparisons, all LPAR configurations being compared should be run using the same zPCR version**, since the supporting algorithms are subject to change.

IBM zPCR User's Guide

Users should assure that the most recently available version of **zPCR** is being used. You can verify the version being run is current by clicking **Check for updates** under **Help** on the ***Control Panel*** window menu-bar.

Getting Started

zPCR is a PC-based productivity tool that runs under Microsoft's Windows or macOS. It is one of a family of tools produced and maintained by your Capacity Planning Support (CPS) team, part of IBM's Washington Systems Center (WSC) in Herndon, Virginia.

PC System (minimum requirement)

Processor	
Intel	Core i3 or better (64-bit required).
Apple	M1 or better (Rosetta 2 required)
OS	
Intel	Microsoft Windows
Apple	Apple macOS
Memory	1 GB is required to install the tool; more is recommended
Graphics	1280x1024; higher recommended

The currently supported environment for **zPCR** is Windows (64-bit) and macOS.

- **Windows:** **zPCR** has been successfully installed and run under Windows 11 and 10. Problems reported with prior Windows versions will not be addressed.
- **macOS:** **zPCR** has been successfully installed and run on Intel and Apple Silicon based Mac hardware. To install and run **zPCR** on an Apple Silicon based Mac, the Rosetta 2 feature must be installed. macOS problems can be addressed only if they can be recreated on the latest version or the version prior to the latest.

Obtaining zPCR

IBM Employees can obtain **zPCR** and other CPS tools via the IBM Intranet at:

<https://w3.ibm.com/w3publisher/wsc-ibm-z/capacity-and-performance/capacity-sizing>

IBM Business-Partners can obtain **zPCR** and other CPS tools via the Internet at:

<https://www.ibm.com/partnerworld/techdocs/prs1762>

IBM Clients can obtain **zPCR** via the Internet at:

<https://www.ibm.com/support/pages/node/6354029>

Questions and Feedback

Contact Capacity Planning Support via ...

- E-mail zpcr@us.ibm.com

Installation on Windows

zPCR is packaged with **Installshield**. To install the tool, execute the package file that you downloaded to your PC from your web site named above. It can be installed to an Administrator ID or a Standard ID. It must be run from the user ID to which it was installed.

zPCR and its supporting Java are intended to be installed to a single folder (the default is **C:\CPSTools\zPCR**). Prompts are provided during installation to specify a different target folder. **The tool should not be installed to C:\Program Files or to C:\Program Files (x86)**. An application icon can optionally be placed on the desktop.

When performing a **zPCR** upgrade as a Standard user, you may be prompted for an Administrator password. This requirement can be eliminated by first uninstalling **zPCR**, and then doing a fresh **zPCR** install.

Installation on macOS

zPCR is packaged in the **pkg** format. To install the tool, execute the package file you downloaded to your Mac from the website named above. ***Note:** when trying to execute the file for the first time, you may receive a message that the **pkg** file cannot be opened because it is from an “unidentified developer”. If you encounter this message, go to **System Preference → Security & Privacy** section in the Mac System Preference panel and click “Open Anyway” to allow zPCR to be installed. Further details can be found in the [Installation Guide for CPS Tools on macOS](#) document.

Once installation is complete, **zPCR** will be found and can be executed from the Applications folder.

In the zPCR install folder, a Documentation sub-folder has been created which includes:

- **NEWS File**
- **zPCR User Guide**
- **Abstract**
- **QuickStart Guide**
- **LSPR FAQ**
- **LSPR Workloads**
- **LSPR Document**

* The **zPCR User's Guide PDF** (this document) is not included with the zPCR install package; rather, it is available as a separate download. If copied into zPCR's **Documentation** folder as “**zpcr_usrguide.pdf**”, it will be accessible from that menu.

* The **LSPR Document PDF** (form SC21-1187) is not included with the zPCR install package; rather, it is available as a separate download. If copied into zPCR's **Documentation** folder as “**LSPRDOC.pdf**”, it will be accessible from that menu.

The following folders are also available for your use, found in C:\Documents\....\zPCR Defaults. Note: For zPCR version 9.4a and prior, these folders were found in the zPCR install folder.

- ✓ **Study Files** Default folder for Study files. Includes a **Single Partition Sample Study** file and a **Multi-Partition Sample Study** (these studies are used for examples shown in this document).
- ✓ **EDF Files** Default folder for EDFs. Includes sample EDFs.
- ✓ **RMF Files** Default folder for RMF files. Includes sample RMF files.
- ✓ **Outputs** Default folder for HTMP and CSV output files.

Java Requirement

zPCR requires the **IBM 64-bit Java runtime environment**. Special versions of the install material for **zPCR** are available to IBM Clients and to IBM Business-Partners that include the necessary Java. You can find instructions for obtaining/installing an appropriate version of the IBM Java runtime environment at the web site from which you obtained **zPCR**. As time progresses, newer versions of Java will be made available.

The IBM Java runtime environment for **zPCR** is installed without replacing or removing any other Java that may already be installed, and does not change your "System" Java.

Note: When other Java runtime packages are installed and configured with Java security enabled, a fresh install of **zPCR** may fail to execute properly. This has been observed in a very small number of **zPCR** installations (less than 0.05%). Should this situation occur, one possible solution is to 1) de-install the other Java, 2) install **zPCR** and its associated Java, and 3) reinstall the other Java.

User's Guide

This **zPCR User's Guide** is distributed in Adobe PDF format. It has been created to be compatible with Acrobat Reader 5.0 and later. The current version of Acrobat Reader for Windows is available for free from:

<http://www.adobe.com>.

The related online help consists of the user's guide text without the figures, implemented as Java Help. This help is context sensitive to the window currently being viewed.

Registration

A user registration process has been implemented to assist in monitoring the distribution and use of **zPCR**. Registration is required for continued usage. **You should be connected to the internet in order to register.**

Until your registration process is completed, a registration form will appear each time **zPCR** is started. **zPCR** may be used up to 3 times without submission of the registration information. After that, the registration process must be completed before the function of the tool can be accessed.

Fill in the requested fields (e.g., name, company name, geographical location, and e-mail address) and click the **Register (Internet)** button. **The primary value of providing a valid e-mail address lies in our ability to notify you of any critical news relating to zPCR usage and/or updates.** Use of your e-mail address will be limited to this purpose only.

There may be cases where a company's firewall will prevent internet registration. In this case click **Register (e-mail)**. This will:

- Attempt to initiate a properly addressed e-mail for you
- Invoke a dialog box with instructions to copy the encoded registration request information into the e-mail note.

Send the e-mail as addressed and wait for an e-mail response (generally within a few hours). Note: **Do not modify the Subject line, as its text is used to trigger an automated registration response**. Once the registration response is received, copy it into the Registration Response Code area and click **Complete Registration**.

Note: If **zPCR** had been terminated, restart it and the registration process will resume the dialog box with instructions to paste the Registration Response.

Registration is only required once, the first time that **zPCR** is started. Once registered, access to the tool is unlimited. Occasionally, as major versions of **zPCR** become available, your registration will be renewed. The renewal is done automatically, if internet registration is allowed. Otherwise, the e-mail registration process is necessary.

Capability to modify or remove your registration is provided under **Registration** on **zPCR**'s **Control Panel** window menu-bar.

Execution

Start execution of **zPCR** by clicking on its program icon, located on your desktop or in the program folder that was designated when it was installed. Multiple copies of **zPCR** can be run at the same time. When loaded, icons are provided on the **Control Panel** window to:

- Set the **Reference-CPU**.
- Select one of the **LSPR Multi-Image Capacity Ratio** tables to review.
- For the **LPAR Configuration Capacity Planning** function, select a specific LPAR configuration to define or review.

Check for Updates

At startup, **zPCR** will check if the version is the most current. If not, the version that is current will be reported (does not necessarily mean an update is required). You must be internet connected for this check to function. Newer versions can be obtained from your **zPCR** download site.

This check can also be done manually by clicking **Check for Updates** under **Help** on the **Control Panel** window.

Usage Tips

Drive\directory path\filename.ext references are currently limited to 255 bytes. This is a Java limitation.

Distorted Windows: If the windows displayed by **zPCR** appear to be distorted or have text displays that tend to get cut off, it is probably due to your Windows display settings. **zPCR** will not tolerate settings for larger fonts. In Windows 10, go to **Display Settings** and set the font size to a smaller percentage.

Low Resolution Displays: **zPCR** is designed for a display resolution of 1280 x 1024 or higher. While a display resolution of 1024 x 768 may function, some windows may be

less than appealing. At lower resolutions, several of **zPCR**'s windows may require the entire vertical or horizontal dimension. If your Windows task bar is always on top, portions of some windows may not be visible without moving the **zPCR** window. This situation can be corrected by going to Windows task bar properties and unchecking "Keep the taskbar on top of other windows".

Erratic Behavior and Abend Situations: Most of the reported **zPCR** execution problems have been traced to outdated graphics drivers and Java activities that invoke "Hardware acceleration". To identify this as the problem, you should turn "Hardware acceleration" off (go to Desktop properties, click on the Settings tab, click the **Advanced** button, click the Troubleshoot tab, and move the slider to "None"). If this solves the problem, you should try to find a more current graphics adapter driver version, or continue to run with the accelerator slider at a reduced setting.

User Controls

A variety of user interfaces are used to control the function of **zPCR**.

Menu-bar: Provides various controls, including the ability to capture output, save a study, and view related documentation and help. These menu-bar items are discussed with the description for each window.

Toolbar Icons: Most of **zPCR**'s windows include smart icons on the toolbar, providing a fast path to various common functions. These icons should be familiar to windows users. If you allow the mouse to dwell over a toolbar icon, a description of the icon's function will appear.

Push button: Click the button.

Entry field: Click on the field to gain focus, key in desired data, and press **Enter**.

Dropdown list : Click the icon at the right end of the entry field, and then click on the desired selection. Selection is limited to the items contained in the list.

Spin buttons : Shown as a pair of buttons at the right end of the entry field. Click the upper or lower button to cycle forward or backward through the predefined entries. Selection is limited to the items contained in the list.

Checkbox : Allows the activation of a feature or function. Click on the box to activate or de-activate.

Radio button : Provides the ability to select one of a number of mutually exclusive choices. Left click on the desired selection.

J-table: Many of **zPCR**'s windows present data in a J-table format. J-tables contain rows and columns of data. Each row/column intersection (cell) may be a presentation field , or one of a number of input field types. In **zPCR**, all J-table user input fields have a white background, while all fixed fields have a shaded or colored background.

To select a row in a J-table, click on any fixed field. You must select a row in order to **Clone** or **Delete** it.

To gain access to a J-table input field, click on the input field. If the field is a

Dropdown list : The selection list will appear. Click on the desired selection.

Entry field : Double click on the cell, key in the desired data, and press **Enter**.

Checkbox : Click the box to set on or off.

Output Capability

- **HTML:** Information from most **zPCR** tables and windows may be output in **HTML** format using the **Output to HTML** toolbar icon. **HTML** can be used as input to various tools, preserving the column and row format, fonts, and coloring. With spreadsheets, **HTML** provides the ability to do further analysis. With most document processing applications, the formatting aspects of **HTML** are preserved, providing a more polished product over that of simple text.

Many PC applications are able to render **HTML** output. It may be more convenient, however, to 1) double click the **HTML** file to open it in its default application, 2) stripe the desired area and copy it to the Windows clipboard, and 3) paste it into the open application where it is desired (i.e., a note, spreadsheet, presentation, etc.).

Note that the precision of numbers in HTML output will be the same as shown on the originating window.

- **CSV:** Comma separated variable format is available with the **LPAR Configuration Capacity Planning** function, for the **Partition Detail Report** and the **Utilized Capacity Report** windows using the **Output to CSV** toolbar icon. The CSV file can be loaded into a spreadsheet for further analysis. While the numeric fields in HTML output carry the same precision that is displayed on the originating window, CSV data carries the full precision used by the internal algorithms.

Note: CSV output does not normally respect the various national languages. All output fields are truly separated by commas. **zPCR** has been modified to determine what character is being used as a decimal place indicator. If it is found to be a comma, then zPCR's CSV output will use a semicolon as the separator character.

Alternatively, the national language can be temporarily set to English (United States) before generating standard CSV output.

- **JPG** and **PNG:** Graphs from **zPCR** can be captured as files in **JPG** format or **PNG** format, or they may be copied to the Windows clipboard.
- **Bitmap:** A picture bitmap of any **zPCR** window can be easily captured (you may want to resize the window so as to include all the information desired before capture). With focus on the desired window, press **Alt-PrintScreen**. The entire window will be copied to the windows clipboard. The clipboard contents can then be pasted into a note or an application such as Microsoft's **Paint**.

Directory Path Settings

zPCR's current **Directory Path Settings** can be found on the **Preferences** window. There are 4 folders (listed below). Each is designated for certain input/output file types.

- **Saved Studies** is the target for ".zpcr" and ".zpcrzap" files.
- **EDF Input Files** is the target for ".edf" files.
- **RMF Input Files** is the target for ".rmf" and ".txt" files.
- **HTML and CSV Output Files** is the target for ".html", ".csv", ".jpg", and ".png" files.

For each I/O related to one of the file types listed above, when done to/from a folder different from that displayed in **Directory Path Settings**, the setting will be updated to reflect the new folder. Path settings can also be changed from the **Preferences** window. Default folder settings can be restored from the **Preferences** window.

Saved Studies

Most **zPCR** inputs can be captured as a saved study. Items saved include the **Reference-CPU** setting and all **LPAR Configuration Capacity Planning** function inputs, including the LPAR host and the individual partition definitions. (Note: A maximum of 130 unique General Purpose, IFL, and ICF partition definitions is allowed, regardless of the limitation on how many can be simultaneously active. zAAP, zIIP, IFL, and ICF partitions associated with a parent GP partition do not count against the maximum.)

zPCR studies can be saved from the **Control Panel** window, by using the **Save** toolbar icons or clicking on **File** on the menu-bar. Studies can also be saved from the various windows in the **LPAR Configuration Capacity Planning** function. The file extension **zPCR** will be assigned.

Saved study files can be loaded into **zPCR** in any of the following ways:

- From the **Control Panel** window, click **File, Load** on the menu-bar.
- Drag and drop a study file  icon onto the **Control Panel** window.
- Drag and drop a study file onto the **zPCR** desktop  icon.
- Double click on a **zPCR** study file  icon.

Saved study files can also be referenced for the purpose of copying one or more partitions into an existing LPAR configuration.

Study files are intended to be supported for "N minus 1" **zPCR** versions. That means, when using **zPCR** version 9.x, a study file from **zPCR** version 8.x should load successfully. When loading older studies, it is possible that adjustments will be made, due to changes in tables or processing methods.

Termination

zPCR is terminated from the **Control Panel** window by clicking on the **Exit** toolbar  icon, or by clicking **File** → **Exit** from the menu-bar (**Ctrl+E** from the keyboard). If an LPAR configuration has been started or modified, a prompt will appear to allow the study to be saved. Clicking on **File** → **Fast Exit** from the menu-bar (**Ctrl+Q** from the keyboard) will bypass the prompt to save a study.

Customization (Preferences)

zPCR can be customized to initialize with your own preset settings: From the **Control**

Panel window menu bar, click **Edit → Preferences** or click the **Preferences** toolbar  icon. Customization settings include:

1. **Reference-CPU Settings**, processor model, scaling-factor and scaling-metric
2. **LSPR Table Control** settings
 - **Workload Categories** displayed in LSPR tables
 - **Selected Families** default list
3. **Partition Table View** - Set the default order that partitions are displayed.
4. **Directory Path for**
 - **Saved Studies**, default folder for loading or saving study files
 - **RMF Input Files**, default folder for loading RMF reports
 - **EDF Input Files**, default folder for loading EDFs
 - **HTML and CSV Output Files**, default folder for generated outputs

Multiple zPCR Invocations

If running multiple **zPCR** invocations, each can be differentiated by referring to the title bar. All primary **zPCR** windows will display the filename of the current study file on the window's title bar. Where subdirectories are involved, only the drive letter, a separator, and the actual filename are displayed. This helps to assure that the filename is visible.

Multiple Monitor Support

If using a PC with multiple monitors set to “Extend these displays”, you can control where **zPCR** windows will display. **zPCR** will initially open on the primary monitor. Subsequent windows will open on the same monitor where the controlling window was displayed. Window positioning on the 2nd monitor remains the same as for a single monitor. This support can be useful to visually isolate two **zPCR** invocations from each other.

In cases where PDF documentation is displayed, the PDF will generally open on the same monitor where it was previously displayed.

Maintaining and Using Multiple zPCR Versions

It is recommended that you always use the latest version of **zPCR**. However, it is possible to have multiple versions installed and usable at the same time.

Preserve access to the former version

1. Make a copy of the **zPCR** installation folder: Before installing the new version, make a copy of the entire folder containing the version to be retained. The best way to do this is to right click on the folder and drag it to the same location or a new location. Then select **Copy** from the popup.
2. Rename the copied folder to represent the former version.
3. Create a desktop shortcut: Open the copied folder, right click on **zPCR.exe**, drag it to the desktop, and select **Create shortcuts here**.
4. Rename the desktop shortcut to better represent the former version.
5. Associate the **zPCR** icon with the desktop shortcut.
 - Right click the desktop icon, and select **Properties**.
 - On the **General** tab, click **Change Icon**.
 - In the **Change Icon** window, click **Browse**
 - Navigate to the **zPCR** install folder and click on the **zPCR.ico** file.
 - Click **OK**.
 - Click **Apply**.
 - Click **OK**.

Install the new version

The former install directory will be the target for the new version. Proceed with a normal install of the new version.

If you want to install the new version to a new folder, you must first uninstall the former version. Uninstall will not affect the steps taken above to retain a former version.

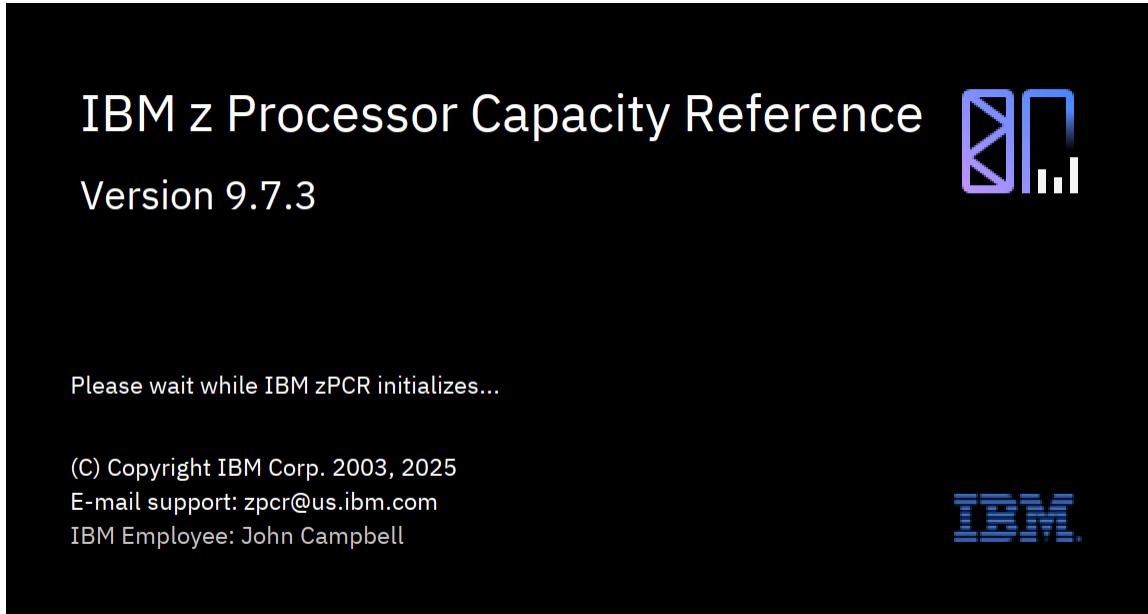
Documentation

Many documentation sources are included available with **zPCR**, including:

- **zPCR User's Guide** (this document) - an Adobe PDF which may be printed or viewed directly on your PC. Figures showing examples of the various **zPCR** windows are included. Bookmarks provide quick access to the various sections. If the **zPCR User's Guide** PDF is downloaded and copied to the **zPCR Documentation** folder (do not rename), it can be accessed by clicking **zPCR User Guide** under **Documentation** on the **Control Panel** menu-bar.
- **zPCR Online Help** - text from the **zPCR User's Guide**, implemented as **Java Help** to be context sensitive to the window currently being viewed. Help can be accessed via the  icon, by pressing PF1, or by clicking on **Help → Context Help** on the menu-bar.
- **zPCR QuickStart Guide** - an Adobe PDF document available via a button on the **Control Panel** window. Use this as a guide for defining a current and an alternate LPAR configuration, and make capacity comparisons between them.

- **LSPR Document** - current version of **IBM's LSPR Document** (SC28-1187) in PDF format. If downloaded and copied to the **zPCR Documentation** folder, this document can be accessed by clicking **LSPR Document** under **Documentation** on the **Control Panel** menu-bar. This document may be found on the following website.
[**IBM Z Large Systems Performance Reference**](#)
- **LSPR FAQ** - an overview (shown as Frequently Asked Questions) of recent changes to LSPR tables, particularly concerning the capacity basis for the **Multi-Image LSPR Processor Capacity Ratios** table. Click **LSPR FAQ** under **Documentation** on the **Control Panel** menu-bar.
- **LSPR Workloads and Capacity Planning Considerations** - information and considerations concerning the LSPR workloads and their use with **zPCR**. Click **LSPR Workloads** under **Documentation** on the **Control Panel** menu-bar. This document has several sections, including:
 1. **LSPR Data and zPCR Usage Considerations** - discusses the current LSPR data and the workload category names for which data is provided. This section includes a discussion concerning the assignment of the various workload categories to workload primitive and workload mix names used previously in **zPCR**.
 2. **LSPR Workload Categories** - provides rationale concerning the choice of capacity planning workloads. Click **LSPR Workloads** under **Documentation** on the **Control Panel** menu-bar.
 3. **LSPR Workload Primitive Descriptions** - a short description of each of the historic LSPR workload primitives. Some of these are used to define the workload categories that are currently published.
- **zAAP/zIIP Capacity Considerations** - a special section of the **zPCR User's Guide** offers insight on zAAP and zIIP capacity (see [**zAAP/zIIP Capacity Considerations**](#)). A white paper concerning zAAPs is also provided. Click **zAAP White Paper** under **Documentation** on the **Control Panel** menu-bar.

Logo Window

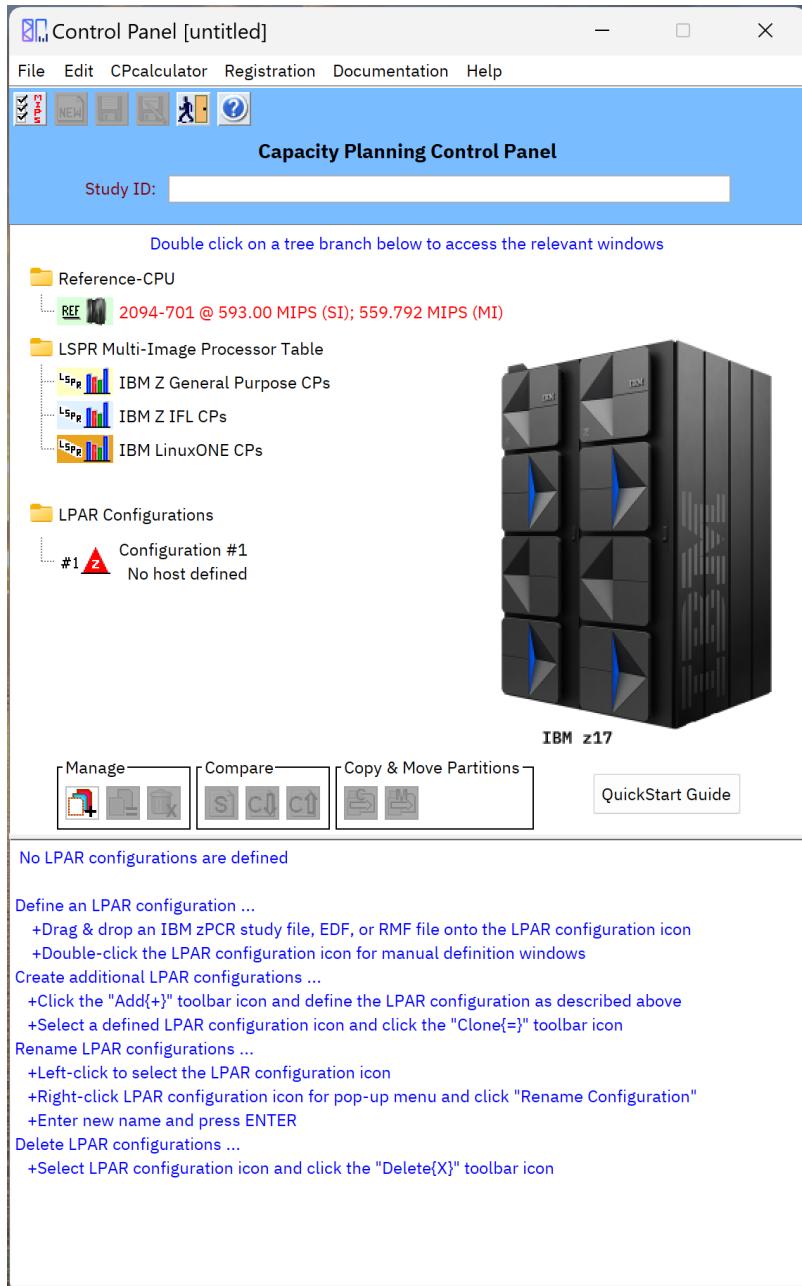


When **zPCR** execution is initiated, the **Logo** window is displayed. This window identifies the application, provides version information (also available under **Help → About** on the menu-bar of the **Control Panel** window).

Registration is required to use **zPCR**. Upon initial use after install, you will be presented with a registration window requesting a few simple entries. The primary purpose of registration is to assure that you can be notified should there be a problem of serious concern. Complete the form and click the **Register** button. Once registration has successfully been completed, the registered user's name appears at the bottom of the **Logo** window. For more detail concerning registration, see [Registration](#) under "Getting Started".

To access the function of **zPCR**, wait for the **Control Panel** window to appear. The "Please wait" message indicates a short delay while the Java code is dynamically compiled for execution.

Control Panel



The **Control Panel** window is displayed immediately following the **Logo** window (once the Java code is compiled). This is the primary window to access all the function of **zPCR**, including

1. Provide a description for the Study ID.
2. Set the **Reference-CPU** metrics 

ITRR Scaling-Factor and Metric are used throughout all **zPCR** function.

3. Review **LSPR Multi-Image capacity** for **IBM Z Processors**

IBM Z GP CPs  (z/OS, z/VM, z/VSE, KVM, or Linux assumed).

IBM Z IFL CPs  (z/VM, KVM or Linux assumed).

IBM LinuxONE CPs  (z/VM, KVM or Linux assumed).

1. Define or review an **LPAR Partition Configuration** 

Use the **LPAR Configuration Capacity Planning** function to:

- Define an **LPAR Host** processor with **General Purpose CPs**, **zAAP CPs**, **zIIP CPs**, **IFL CPs**, and **ICF CPs** (see [LPAR Host Processor](#)).
- Define partition types of **General Purpose**, **IFL**, or **ICF**. **zAAP** and **zIIP** LCPs are always associated with parent GP partition. **IFL** and **ICF** LCPs can be associated with a parent z/VM GP partition (see [Partition Definition](#)).
- View the **Host Summary Report**; for an overview of the LPAR configuration and the capacity results (see [Host Summary Report](#)).
- View the **Partition Detail Report** to review individual partition capacity. Partition definition changes can also be made from this window (see [Partition Detail Report](#)).
- Additional capability: (see [Control Panel - Advanced Usage](#))
 - Compare capacity between two LPAR configurations.
 - Move or copy partitions between LPAR configurations.

In **zPCR**, IBM mainframe processor capacity planning requirements can be satisfied by using:

1. The **LSPR Multi-Image Capacity Ratios** table, which generalizes on capacity for processors assumed to be running typical partition configurations. Use this table to identify models of a desired approximate capacity.

To help generalize capacity relationships for processors running only a single partition, an **LSPR Single-Image Capacity Ratios** is available. This table ignores any effect on capacity due to configurations that involve more than one partition. Note that the **Single-Image** table is limited to showing only reasonable single-image partition configurations (the limit is currently 30 CPs).

2. The **LPAR Configuration Capacity Planning** function provides a more precise capacity projection for any particular LPAR Host processor running with any specific partition configuration, each partition running a specific SCP and workload.

All contemporary IBM Z processors run multiple partitions. To fairly show capacity relationships for **IBM Z** platforms, the **Multi-Image Capacity** Ratios table should be used.

The **LPAR Configuration Capacity Planning** function provides the most accurate view of capacity since it understands the specific LPAR configuration in question.

Note: The **LPAR Configuration Capacity Planning** function should be viewed as providing the most reliable capacity information, since its results are based on the specific LPAR host processor and its precise partition configuration, including SCP/workload assignments that are representative of the actual production workload.

The **LSPR Multi-Image Capacity Ratios** table is not referenced by the **LPAR Configuration Capacity Planning** function of **zPCR**. Its use is limited to the generalization of capacity relationships between processors, assuming average LPAR configurations that are considered typical for each processor model. The table is actually built using **zPCR's LPAR Configuration Capacity Planning** function.

An extended **LSPR Single-Image Capacity Ratios** table is used exclusively by the **LPAR Configuration Capacity Planning** function as the initial source for computing capacity for specific LPAR configurations.

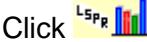
Reference-CPU

Capacity values in **zPCR** are always displayed relative to a single **Reference-CPU** which may be set to any 1-way General Purpose model. The capacity rating can be set to any desired scaling-factor and scaling-metric.

On the **Reference-CPU** tree, click  to set the processor model and its capacity metrics (see [Reference-CPU](#)).

Multi-Image LSPR Capacity Ratio Tables

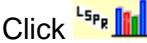
On the **LSPR Multi-Image Processor** tree:

- Click  to review capacity for **General Purpose CPs**.
z/OS, z/VM, z/VSE, KVM, or Linux workload assumed.
- Click  to review capacity for **IFL CPs**.
z/VM, KVM or Linux workload assumed.
- Click  to review capacity for **LinuxONE CPs**.
z/VM, KVM or Linux workload assumed.

For information concerning LSPR workload categories (see [Workloads](#)).

Single-Image LSPR Capacity Ratio Tables

Right click on the **LSPR Multi-Image Processor** tree, and click **Switch to LSPR Single-Image Table**. The icons will provide access to the Single-Image view.

- Click  to review capacity for **General Purpose CPs**.
z/OS, z/VM, z/VSE, KVM, or Linux workload assumed.
- Click  to review capacity for **IFL CPs**.
z/VM, KVM or Linux workload assumed.

Note that the switch to Single-Image is temporary. After the LSPR table has been viewed, Multi-Image will be restored to the tree.

LPAR Configuration Capacity Planning

Specific partition configurations may be defined for up to 10 LPAR host processors.

On the **LPAR Configurations** tree:

- Click  #1  to define or review the 1st LPAR configuration.
- Additional icons will appear as more LPAR configurations are added.

Note: The **LPAR Configuration Name** shown on the **Control Panel** has 2 levels,

The 1st level has default text which can be user modified from the **Control Panel** (see [Renaming LPAR Configurations](#)). This level is displayed on all report windows.

The 2nd level is generated automatically when the LPAR host processor becomes defined. This name is based on the processor hardware model with the real GP CP model information. Appended is the number of other real CP types defined ("A=" for zAAPs; "I=" for zIIPs; "F=" for IFLs; "C=" for ICFs). This level is displayed only on the **Control Panel**.

When a duplicate name would result, a repetition index is appended. Whenever the LPAR host is changed, a new name is generated.

When a defined LPAR configuration is selected, a summary table is displayed at the bottom showing CP counts and capacity values. If there are more real CPs than can be used, two additional rows will appear representing Reserve Capacity. For details [see Reserve Capacity](#).

Saved Studies

zPCR inputs that can be captured in a ***.zpcr** study file, includes the **Reference-CPU** settings, including the scaling-factor and metric, and all **LPAR Configuration Capacity Planning** function definitions. Study files can be reloaded at a later time for review or further analysis. Using the menu-bar, save a study by clicking on **File**, and then on **Save as** (or **Save**, if this is already a named study). **Save as** will prompt for a file name. The default file extension assigned is **zPCR**.

When a saved study is loaded, the study filename appears on the title bar of all the primary windows.

New Study

A new study can be initiated by clicking **New** under **File** on the menu-bar. This is equivalent to terminating and restarting **zPCR**. All inputs and controls (including the **Reference-CPU** settings) will be returned to their **Preferences** state, or, if preferences have not been set, to their default state.

Customizing zPCR Initialization

zPCR can be customized via the **Preferences** window to always open with your own preset settings for the following items:

1. **Reference-CPU Settings**, processor model, scaling-factor and scaling-metric
2. **LSPR Table Control** settings
 - **Workload Categories** displayed in LSPR tables
 - **Selected Families** default list
3. **Partition Table View Settings** (default order partition types are displayed)
4. **Directory Path for**
 - **Saved Studies**, default folder for loading or saving study files
 - **RMF Input Files**, default folder for loading RMF reports
 - **EDF Input Files**, default folder for loading EDFs
 - **HTML and CSV Output Files**, default folder for generated outputs

To activate customization, click on **Edit → Preferences** on the menu-bar; the **Preferences** window will open (see [Preferences](#)). Customization can only be invoked prior to defining partitions.

The **Control Panel** window can be positioned anywhere on the desktop (dual displays included). The new position will be used each time this window is returned to. When using multiple **zPCR** invocations at the same time, this feature makes it convenient to keep track of them.

Updating your Registration

You can update your registration (for example, you have a new email address), or delete your registration, by clicking **Registration** on the menu-bar. If you delete registration, **zPCR** (and any other CPS tools installed) will cease to run until you complete a new registration.

Documentation Sources

Some of the information sources found under **Documentation** on the menu-bar, include:

- ✓ [zPCR News file](#) carries information concerning changes in this **zPCR** version.
- ✓ [LSPR FAQ](#) provides access to a paper with information concerning the use and interpretation of the Multi-Image and Single-Image LSPR capacity ratio tables. Also discussed is the impact of HiperDispatch and capacity considerations for z/VM on z10 and later processors.

The [IBM Z Large Systems Performance](#) web site by IBM development covers z/OS and z/VM version measurement data each IBM Z family. It can be found at:

www.ibm.com/support/pages/ibm-z-large-systems-performance-reference

- ✓ [LSPR Document](#) (form SC28-11187) is a PDF available via the above website. It provides insight for LSPR and its use. It also contains IBM Z capacity ratio tables for all recent LSPR versions. If ported into **zPCR**, it can be accessed via the menu-bar.

The version of the LSPR Document included with **zPCR** is most useful from a discussion standpoint, but the tables in the appendix could possibly exclude some of the latest capacity data. In lieu of using the tables in this document, one should reference **zPCR**'s **Multi-Image Ratios** table, which always includes the latest data. In addition, **zPCR** gives you the flexibility of setting the basis of the capacity values by setting the **Reference-CPU** and its scaling-factor and scaling-metric.

- ✓ **LSPR Workloads** displays pertinent information concerning the LSPR workload categories included in the LSPR tables. Also offered are considerations related to their usage for capacity planning purposes.

The **QuickStart Guide** button provides access to a short paper describing the process using **zPCR** to define a current and an alternate LPAR configuration and make capacity comparisons between them.

Toolbar Icons

Click the **Preferences** toolbar icon to go to the **Preferences** window.

Click the **New** toolbar icon to start a new study.

Click the **Save** toolbar icon to save the current study.

Click the **Save as** toolbar icon to save the current study with a new name.

Click the **Exit** toolbar icon to exit **zPCR**.

Click the **Help** toolbar icon to access context sensitive help for this window.

Notes

The **Sample zPCR StudyStudy - Advanced Usage.zpcr** study file, included with the **zPCR** package, is the source used for the examples shown for the **Control Panel** and following chapters.

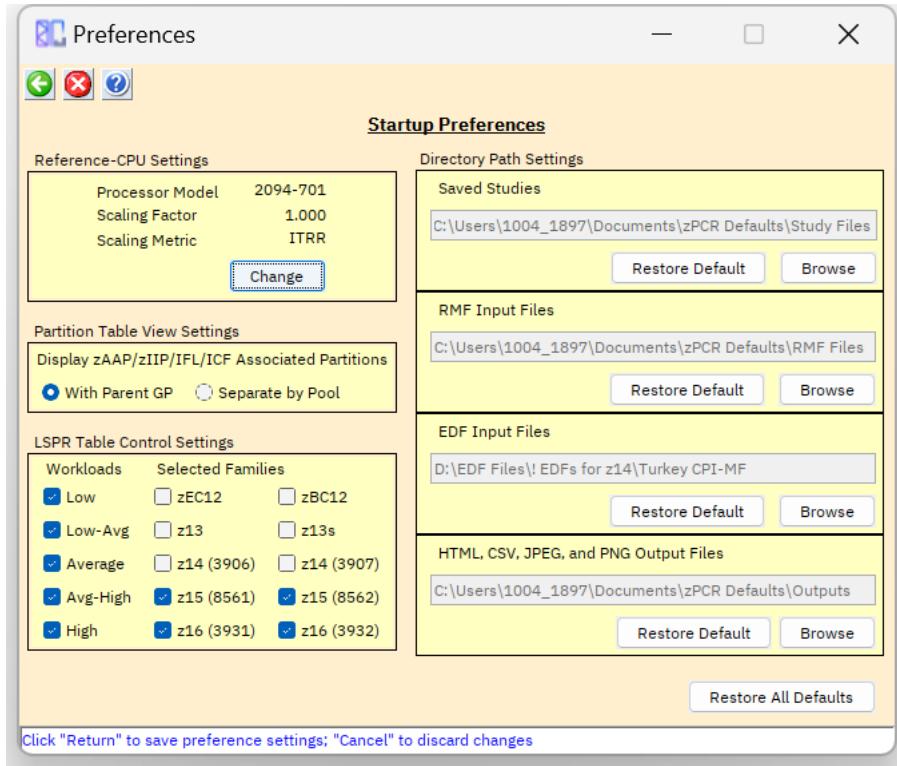
The **Sample zPCR StudyStudy - Basic Usage.zpcr** study file included with the **zPCR** package, is the source used for the examples shown in this and following chapters.

zPCR study files can be loaded by dragging them onto the **Control Panel** window.

Menu-bar

File	
New	Start a new (untitled) study.
Load	Open a previously saved study.
Save	Save this study (must already be a titled study).
Save as...	Save this study with a new name.
Up to 10 filenames	Recent study files are listed for possible loading.
Exit (Ctrl+E)	Terminate zPCR execution.
Fast Exit (Ctrl+Q)	Terminate zPCR execution immediately.
Edit	
Preferences	Personalize zPCR startup settings (see Preferences).
CPcalculator	Capacity planning calculator function (see CP Calculator)
zAAP Capacity	Capacity estimator for migration to a zAAP or zAAP on zIIP configuration (see zAAP Capacity Estimator Input)
Registration	
Remove	Delete your CPS tool registration (disables registration for zPCR and any other CPS tools installed).
Documentation	Various supporting documentation, including:
	What's New?
	User Guide
	LSPR FAQ
	LSPR Workloads
	LSPR Document
	HiperDispatch Consideration
	zAAP/zIIP Considerations
	zAAP White Paper
	Concerning Accuracy
	Obtain CP3KEXTR
	Obtain CP3KVMXT
Help	
Context Help (F1)	A tool bar icon is also available.
QuickStart Guide	Provides guidance on using zPCR .
Check for updates	Verifies that the latest version of zPCR is being used.
Toolbar Icons	Summarizes the LPAR Configuration Toolbar icons used.
Privacy Policy	Displays the zPCR Privacy Policy.
Third Party Notices	Display zPCR 3 rd party notices.
About	Display Logo window information including the zPCR version.

Preferences



The **Preferences** window provides the capability to establish your own initialization conditions at **zPCR** startup. These settings will be used whenever a new study is started.

The following preference items may be set:

1. **Reference-CPU Settings** group box:

To set the **Reference-CPU**, click the **Change** button. The **Reference-CPU** window will be presented, from which to make your selection, setting the processor model, a scaling-factor and a scaling-metric. This preferences setting can alternatively be made using the normal **Reference-CPU** window via the **Update zPCR Startup Preferences on Return** checkbox (see [Reference-CPU](#)).

2. **Partition Table View Settings** group box.

Used to set the default order in which partitions are displayed in various tables.

With Parent GP: Each GP partition is immediately followed by any associated zAAP, zIIP, IFL, and ICF partitions.

Separate by Pool: All GP partitions are listed together, followed by all zAAP partitions, all zIIP partitions, all IFL partitions, and finally all ICF partitions.

Regardless of which default setting is selected, the order can be temporarily switched while the window is displayed.

3. **LSPR Table Control Settings** group box:

Check the LSPR workload categories that are to be displayed in the **LSPR Table** windows.

Check the specific processor families that will be displayed when **Selected Families** is used to view the **LSPR Table** windows. Note that only the most recent processor families are available for selection.

4. **Directory Path Settings** group box:

The default location for **Saved Studies**, **RMF**, **EDF**, and general output such as **HTML**, **CSV**, **JPG**, and **PNG** can be defined.

Click **Browse** to open a standard windows dialog from which to make the **drive:\directory** selection.

Click **Restore Default** to restore the default setting.

Click the **Restore All Defaults** button to return all preference settings to their original default values. The **Preferences** window remains open in case there is a desire to change any specific setting before exiting the window.

Note: **Directory Path Settings** are not saved with in **zPCR** study file. When a study file is loaded, its folder may affect the **Saved Studies** folder. The remaining 3 folder settings will remain those set the last time **zPCR** was used.

The Preference settings for the **Reference-CPU** are used each time a new **zPCR** study is started. Once a study has been started, the **Preferences** window can no longer be accessed. However, the preferences setting for the **Reference CPU** can be changed at any time from the **Reference-CPU** window by checking the box **Update zPCR Startup Preferences on Return**.

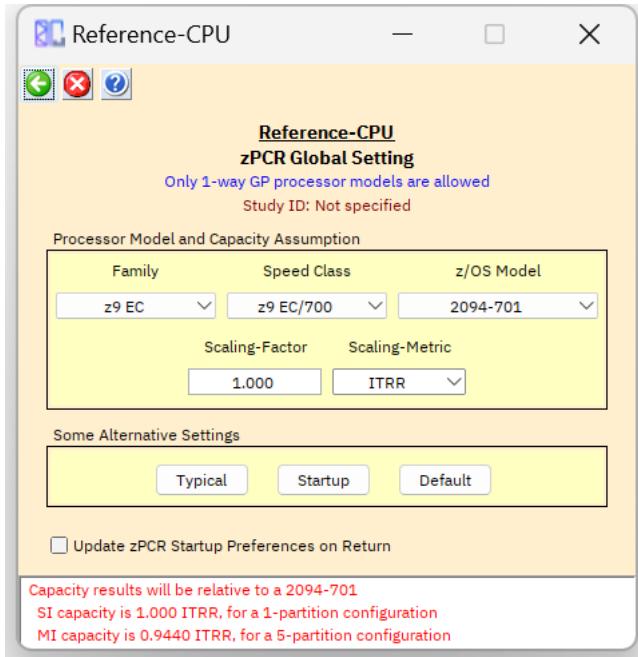
Whenever the **Reference-CPU** window is open, the current preferences setting can be restored by clicking **Startup**. Whenever the **LSPR Table Control** window is open, its current preferences setting can be restored by clicking **Restore Startup Settings** under **Settings** on the menu-bar.

When loading a previously saved study, all settings revert to those that were in use at the time the study was saved.

Click the **Return** toolbar icon to accept your preference changes and return to the **Control Panel** window. Your preference setting will take effect immediately.

Click the **Cancel** toolbar icon to return to the **Control Panel** window without accepting any preference setting changes that were made.

Reference-CPU



The **Reference-CPU** window is accessed from the **Control Panel** by clicking the  icon on the Reference-CPU tree. The example below sets the **Typical** setting of **2094-701 @ 593 MIPS**, which will be used through the remainder of this document.

The **Reference-CPU** can be set to any **IBM Z** processor 1-way model. Only models representing General Purpose CPs (as known to z/OS) can be selected. If intending to represent IFL capacity (**IBM Z** or **LinuxONE**), the **Reference-CPU** should be set to an equivalent full speed GP 1-way model. Specific controls on the window are discussed below.

The **Reference-CPU** model and its scaling-factor/metric can be changed at any time. When changed, capacity values in all of the **LSPR Capacity Ratios** tables and the **LPAR Configuration Capacity Planning** will be updated to reflect the change. [It is critical that all capacity results that are to be compared, be obtained using consistent Reference-CPU metrics.](#)

While default and typical **Reference-CPU** metrics are provided in **zPCR**, the processor model chosen and its scaling-metric remain a purely arbitrary decision to be made by the user. When changing the processor model, no attempt is made to adjust the scaling-factor; rather it is initially set to **1.00** and the scaling-metric is set to **ITRR** (Internal Throughput Rate Ratio). These may then be user modified.

Processor Model and Capacity Assumption group box

- **Family** Select the processor family from the dropdown list
- **Speed Class** Select the processor speed for GP CPs from the dropdown list.
- **Model** The **Reference-CPU** processor can only be set to 1-way **Models**. Since there is only a single 1-way model for each **Speed Class**, no selection is necessary.
- **Scaling-Factor** Enter the capacity value assumed for this processor. The value entered will be assumed as the capacity of a shared 1-partition configuration (default scaling-factor is 1.00). Regardless of the scaling-factor value entered, excessive decimal places will be dropped so as to preserve a maximum of six significant digits. Note that, while some processor families can be run in **Power-Save** mode, the scaling-factor always represents the **Reference-CPU** at full power.
- **Scaling-Metric** Set the metric assumed for the capacity value, using the dropdown list or by keying in up to 12 characters. The default scaling-metric is **ITRR** (Internal Throughput Rate Ratio).

The **LSPR Multi-Image Capacity Ratios** table and the **LPAR Configuration Capacity Planning** function support all IBM mainframe families and models. Therefore, any General Purpose processor 1-way model can be assigned as the **Reference-CPU**.

Some Alternative Settings group box

- **Default** assigns **2094-701** at **1.00 ITRR**
- **Typical** assigns **2094-701** at **593 MIPS**
- **Startup** assigns startup preferences **Reference-CPU** values

A message box at the bottom of the **Reference-CPU** window will display the status of any exception condition.

Click the **Return** toolbar icon to accept your changes and return to the calling window.

Click the **Cancel** toolbar icon to return to the calling window without accepting any changes.

The new **Reference-CPU** settings can be saved as the **Preferences** setting by checking the **Update zPCR Startup Preferences on Return** checkbox. The **Preferences** setting will be updated when **Return** is clicked.

The normal **Reference-CPU** is always limited to 1-way processor models. However, when viewing **LSPR Capacity Ratios** tables, any processor model in the table can be set as a **Provisional Reference-CPU** (temporary setting). For detail see [Provisional Reference-CPU](#).

For detail on how the **Reference-CPU** scaling-factor is interpreted for the various **zPCR** functions, see [How the Scaling-Factor Is Used](#).

LPAR Configuration Capacity Planning Considerations

A 1-way processor model is always required for the **Reference-CPU**. For **LPAR Configuration Capacity Planning** purposes, a 1-way processor model provides a more consistent and understandable capacity perspective, particularly when small partitions are defined on large N-way processors, running workloads with lesser N-way efficiency. If the **Reference-CPU** were allowed to be set as an N-way processor, the MP effects of various workloads would affect the capacity perspective for partitions, such that, though correct, would be difficult to understand.

Most data processing installations exploit processor models with multiple engines (N-way is greater than 1). And most installations configure multiple active partitions. An installation often associates some capacity value and metric (i.e., MIPS) with the processor running their production work. In such cases, one would like to assign the current LPAR host model as the **Reference-CPU**, setting the scaling-factor to the associated capacity value. Since the **Reference-CPU** is limited to 1-way processor models, this is not possible. To accommodate such situations, a **Calibrate** function has been provided (see [Calibrate Capacity to LPAR Host](#)). You will need to start with a 1-way model as the **Reference-CPU**. Once the LPAR host's partition configuration has been defined, you can then calibrate the current **Reference-CPU** model's scaling-factor such that the LPAR host capacity result is the desired value. Use any capacity scaling-factor/metric that will help to make capacity values that are presented relevant for the study. Once calibrated, the **Reference-CPU** settings must not be changed when comparing new LPAR configurations to the current one.

How the Scaling-Factor Is Used

The **Reference-CPU** window establishes the **processor model** and **scaling-factor/metric** used for all capacity values displayed in **zPCR**. It is always set in terms of the LSPR Single-Image table. The chart below is provided to help in understanding how the **Reference-CPU** settings are interpreted by the various **zPCR** functions.

Reference-CPU Examples of how the scaling-factor is interpreted by zPCR functions		
Reference-CPU window (button for auto-defining the setting)	Default	Typical
Reference-CPU model (1-way General Purpose models only)	2094-701	2094-701
Scaling-factor (capacity of a shared 1-partition configuration)	1.00	593
Scaling-metric	ITRR	MIPS
Value used for LSPR Multi-Image table	0.944	559.792
Value used for LSPR Single-Image table	1.000	593.000
Value used for LPAR Configuration Capacity Planning function assumed to represent capacity of the entire CPC	1.0197	604.682
Regardless of the scaling-factor value entered, excessive decimal places will be dropped so as to preserve a maximum of six significant digits.		
The above relationships apply to 1-way processors only. For 1-way processors, partition configurations are always predicable (1 LCP only) while not for N-way models. This is one reason why the Reference-CPU setting is limited to 1-way processor models.		

Buttons are provided on the **Reference-CPU** window to automatically populate the input fields with the settings shown below:

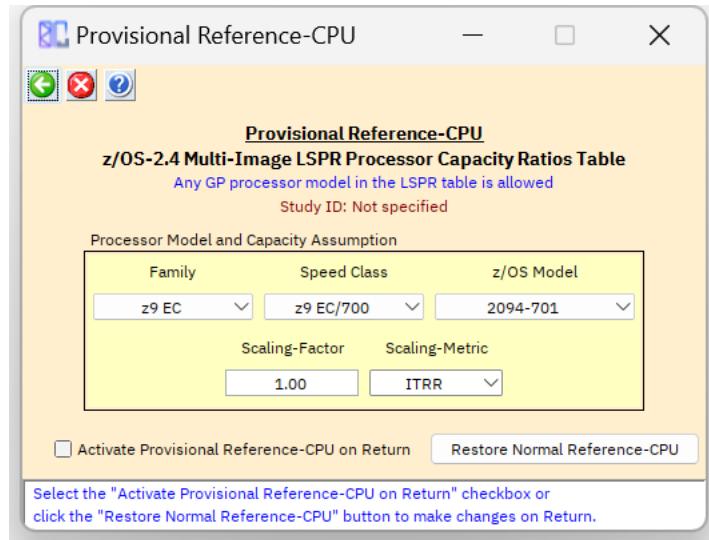
- Default** Assigns the **2094-701 at 1.00 ITRR**
- Typical** Assigns the **2094-701 at 593 MIPS**
- Startup** Assigns the startup setting, if defined in user preferences

Note: IBM's **LSPR Document** (SC28-1187) publishes only **Multi-Image ITR Ratios**, all based on the **2094-701 rated at 1.00** (rather than the 0.944 shown above). To reproduce these ITR ratios in **zPCR**, you must define the **Reference-CPU** in such a way that the 2094-701 comes up as 1.00 in the Multi-Image table. This can be done by setting the scaling value to **1.000 / 0.944**, or **1.059322**. To reproduce the **PCI values**, the **Reference-CPU** must be set as **2094-701 at 593 MIPS**.

Capacity results for any specific LPAR configuration can vary considerably from those provided in the **LSPR Multi-Image table**. The difference lies in the degree to which the specific LPAR configuration deviates from the typical LPAR configuration assumed for that same processor model in the table. The use of zAAP or zIIP specialty engines will also contribute to deviation, since these CP types are not included in the typical configurations.

Provisional Reference-CPU

When viewing the **LSPR Capacity Ratio** table for General Purpose CPs, a **Provisional Reference-CPU** may be set, serving as a temporary **Reference-CPU** setting only for that window while it is open. Upon closing, the normal **Reference-CPU** settings are restored. Click the **Provisional Reference-CPU** button at the bottom of the **Processor Capacity Ratios** window to open the **Provisional Reference-CPU** window.



The **Provisional Reference-CPU** window is aware of which **LSPR Processor Capacity Ratios** table window is calling it; a title line will so indicate. Entry fields, similar to those on the **Reference-CPU** window, are described below.

Each **LSPR Table** window invocation will default to using the **Normal Reference-CPU** settings. Once a **Provisional Reference-CPU** has been established, it will remain in effect until either the **LSPR Table** window is closed or it is manually changed via the **Provisional Reference-CPU** window.

Unlike the **Reference-CPU**, the **Provisional Reference-CPU** can be set to any GP processor in the LSPR table, including N-way models, thus allowing any specific model to model comparisons.

When viewing IFL or LinuxONE models, click the **Provisional Reference-CPU** button must be used.

The **Provisional Reference-CPU** can only be set for General Purpose models. To get an equivalent capacity perspective for IFL CPs, one should choose an equivalent full speed model from General Purpose CP table.

Fast Path for setting the Provisional Reference-CPU

When viewing the **GP LSPR Processor Capacity Ratios** window, double-click any processor model displayed in the window. The **Provisional Reference-CPU** window will open with appropriate values already populated in the **CPU family** and **Processor Model** fields. Only the scaling-factor and scaling-metric need be entered (if desired).

When viewing IFL or LinuxONE models, the **Provisional Reference-CPU** button must be used.

Processor Model and Capacity Assumption group box

- **Family** Select the processor family from the dropdown list.
- **Speed Class** Select the processor speed for GP CPs from the dropdown list.
- **Model** Select the specific processor model from the dropdown list. The list includes all N-way models.
- **Scaling-Factor** Enter the assumed capacity value (default is 1.00). The value entered will be assumed as the capacity of a shared 1-partition configuration. Regardless of the scaling-factor value entered, excessive decimal places will be dropped so as to preserve a maximum of six significant digits.
Note that, while some processor families can be run in **Power-Save** mode, the scaling-factor always represents the **Reference-CPU** at full power.
- **Scaling-Metric** Set the metric assumed for the capacity value, using the dropdown list or by keying it in. The default is **ITRR** (Internal Throughput Rate Ratio).

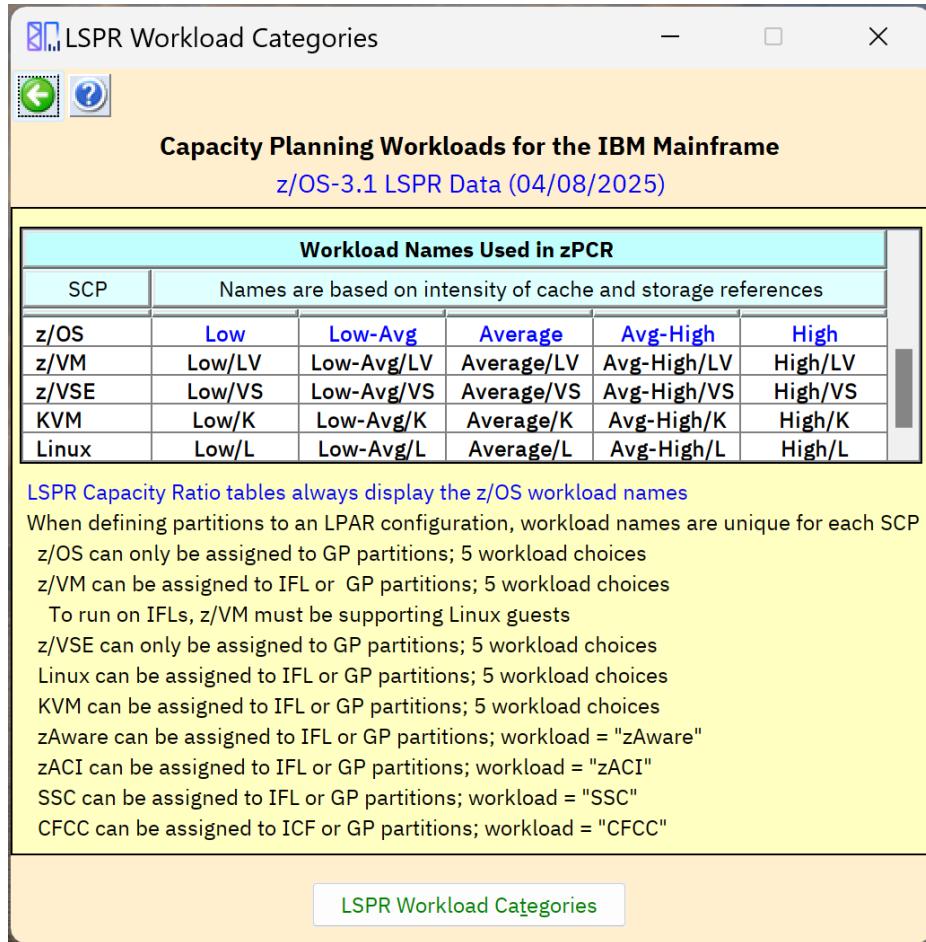
Check **Activate Provisional Reference-CPU** to confirm that the temporary setting should be applied when the **Return** toolbar icon is clicked. Otherwise, the current setting will remain in effect.

Click the **Restore Normal Reference-CPU** button to restore the normal **Reference-CPU** values. Note that normal **Reference-CPU** values are always restored when exiting the **LSPR Table** window.

Click the **Return** toolbar icon to accept the settings and return to the LSPR table window.

Click the **Cancel** toolbar icon to reject any changed settings and return to the LSPR table window.

Workloads

A screenshot of a software window titled "LSPR Workload Categories". The window contains a table titled "Workload Names Used in zPCR" showing workload categories for various SCPs (z/OS, z/VM, z/VSE, KVM, Linux). It also includes a note about the ratio tables and specific assignment rules for each SCP.

Workload Names Used in zPCR					
SCP	Names are based on intensity of cache and storage references				
	Low	Low-Avg	Average	Avg-High	High
z/OS	Low/LV	Low-Avg/LV	Average/LV	Avg-High/LV	High/LV
z/VM	Low/VS	Low-Avg/VS	Average/VS	Avg-High/VS	High/VS
z/VSE	Low/K	Low-Avg/K	Average/K	Avg-High/K	High/K
KVM	Low/L	Low-Avg/L	Average/L	Avg-High/L	High/L
Linux					

LSPR Capacity Ratio tables always display the z/OS workload names
 When defining partitions to an LPAR configuration, workload names are unique for each SCP
 z/OS can only be assigned to GP partitions; 5 workload choices
 z/VM can be assigned to IFL or GP partitions; 5 workload choices
 To run on IFLs, z/VM must be supporting Linux guests
 z/VSE can only be assigned to GP partitions; 5 workload choices
 Linux can be assigned to IFL or GP partitions; 5 workload choices
 KVM can be assigned to IFL or GP partitions; 5 workload choices
 zAware can be assigned to IFL or GP partitions; workload = "zAware"
 zACI can be assigned to IFL or GP partitions; workload = "zACI"
 SSC can be assigned to IFL or GP partitions; workload = "SSC"
 CFCC can be assigned to ICF or GP partitions; workload = "CFCC"

[LSPR Workload Categories](#)

The **Workloads** window is accessed from any of the **LSPR Processor Capacity Ratios** tables via the **Workload Categories** button. Its purpose is to show the workload category names used in **zPCR** for each of the major SCPs.

Each of the z/OS workload categories is represented on every **IBM Z** processor model. These z/OS workload categories are always displayed in the **LSPR ITR Ratio Tables**. These z/OS workload categories are also considered to be representative for z/VM, Linux, and z/VSE. When General Purposed CPs are being displayed, the capacity values are considered applicable for z/OS, z/VM, Linux, and z/VSE. When IFL CPs are being displayed, the capacity values are considered applicable for z/VM or Linux (z/OS and z/VSE can only be run in a GP partition).

General information about actual z/OS LSPR workload primitives used to generate the workload category data is included in the **LSPR Workloads** document. However, no specific information is provided about which workload primitives contribute to any specific workload category.

For **zPCR**'s **LPAR Configuration Capacity Planning** function, the table provides names for all of the workloads available for assignment to z/OS, z/VM, Linux, and z/VSE partitions.

- **z/OS** can only be defined to GP partitions; **Average** is the default assignment.
- **z/VM** can be defined to IFL or GP partitions; **Average/LV** is the default assignment.
- **z/VSE** can only be defined to GP partitions; **Average/VS** is the default assignment.
- **KVM** can be defined to IFL or GP partitions; **Average/K** is the default assignment.
- **Linux** can be defined to IFL or GP partitions; **Average/L** is the default assignment.
- **zAware** can be assigned to IFL or GP partitions on z12 and z13 processors. **zAware** is also the sole workload assignment. **SSC** is intended to replace **zAware**.
- **zACI** can be assigned to IFL or GP partitions on z13 processors. **zACI** is also the sole workload assignment. **SSC** is intended to replace **zACI**.
- **SSC** can be assigned to IFL or GP partitions on z13 and later processors. **SSC** is also the workload assignment.
- **CFCC** can be defined to ICF or GP partitions; **CFCC** is also the sole workload assignment.

z/OS-3.2 supports z17, z16, and z15

z/OS-3.1 supports z17, z16, z15, and z14

z/OS-2.5 supports z17, z16, z15, z14, and z13

z/OS-2.4 supports z17, z16, z15, z14, z13, and z12

z/OS-2.3 supports z16, z15, z14, z13, and z12.

z/OS-2.2 and **2.1** support z15, z14, z13, z12, z11, and z10.

z/OS-1.13 and prior are supported by **zPCR** for various older processor families

z/VM-7.4 supports z17, z16, and z15

z/VM-7.3 supports z17, z16, z15, and z14

z/VM-7.2 supports z16, z15, z14, and z13

z/VM-7.1 supports z15, z14, z13, and z12

z/VM-6.4 and **5.4** support z10 and z9

KVM and **Linux** can support any number of LCPs (IFL or GP)

zAware is limited to a maximum of 16 LCPs (IFL or GP)

zACI is limited to a maximum of 16 LCPs (IFL or GP)

SSC can support a maximum of 208 LCPs (IFL or GP)

CFCC supports all IBM z models with a maximum of 16 LCPs (IFC or GP).

LSPR Workloads

Various LSPR workload primitives are measured in order to fully understand the total capacity envelope (or bounds) that can be expected for each processor model. Combinations of primitives are then used to develop capacity values for the 5 LSPR workload categories carried in **zPCR**. The association a production workload to one of these workload categories provides a reliable way to evaluate capacity expectation when planning for a processor replacement.

Relative Nest Intensity (RNI) is used to understand the memory use of a workload. The term **Nest** refers to a processor's memory sub-system, beyond the 1st level High Speed Buffer all the way out to Central Storage. **Intensity** refers to the amount of impact the **Nest** has on capacity for each workload category. **Relative** is used because each different processor hardware design will affect the absolute **Nest Intensity** of any given workload.

The five z/OS LSPR workload categories available in **zPCR** are listed in the table below. The calculated **RNI** in combination with **L1MP** (level 1 cache miss percent) is used select either the **Low**, **Average**, or **High** workload category. Note that **CPU MF** data must be captured in order to compute a production workload's **RNI** and **L1MP**.

Name	Characterization
Low	Light use of the processor's memory hierarchy
Low-Avg	Mid-way between Low and Average
Average	Average use of the processor's memory hierarchy
Avg-High	Mid-way between Average and High
High	Heavy use of the processor's memory hierarchy

These 5 workload categories are expected to satisfy most IBM mainframe capacity planning requirements. All of these categories are available in **zPCR**, in the **LSPR Multi-Image** tables, the **LSPR Single-Image** tables, and for the **LPAR Configuration Capacity Planning** function.

While the 5 workload categories are derived primarily from z/OS measurement data, they are also representative for z/VM, Linux, z/VSE, KVM, Linux production workloads. Similarly named workload categories are also available for each of these SCPs when assigned to a partition.

Window Controls

Click the **Return** toolbar icon to return to the calling window (usually the **Control Panel** window).

Click the **LSPR Workload Categories** button for discussion concerning how to select between the workload categories available in **zPCR**.

Choosing a workload category to represent production work

The selection of an appropriate workload category to best represent a given production workload running in a partition is very important when projecting capacity requirements for a potential replacement processor.

- **z/OS** can be defined to General Purpose partitions (and any associated zAAP or zIIP partitions). When ***CPU MF*** data is not available, ***Average*** will default for the assigned workload category. The workload used to represent a z/OS partition can be manually designated as any of the 5 workload categories.

CPU Measurement Facility (***CPU MF***) data provides a reliable method for assigning an LSPR workload category to a partition. ***CPU MF*** data can be obtained from z/OS systems running on z10 and later processors. ***CPU MF***, which is captured as SMF-113 records, must be turned on for each partition where the information is desired. SMF is post-processed into an EDF format that ***zPCR*** can read. For each partition where ***CPU MF*** data was captured, the most representative LSPR workload category will be assigned and ***CPU MF*** will be shown as the ***Method Used***. zAAP and zIIP partitions will be assigned the same workload category as its parent GP partition.

- **z/VM** can be defined to General Purpose or IFL partitions. When ***CPU MF*** data is not available, ***Average/LV*** will default for the assigned workload category. The workload used to represent a z/VM partition can be manually designated as any of the 5 workload categories.

CPU Measurement Facility (***CPU MF***) data provides a reliable method for assigning a workload category to a z/VM partition. ***CPU MF*** data can be obtained from z/VM systems running on z10 and later processors. ***CPU MF***, captured in VM Monitor records, must be turned on for each partition where the information is desired. The VM Monitor data is post-processed into an EDF format that ***zPCR*** can read. For each partition where ***CPU MF*** data was captured, the most representative LSPR workload category will be assigned and ***CPU MF*** will be shown as the ***Method Used*** for the assignment.

When z/VM is running in a General Purpose partition, while also supporting IFLs in an associated partition, ***CPU MF*** data is captured separately for each. The workload category assignment could be different for these two partitions.

- **z/VSE** can be defined to General Purpose partitions. The default workload category assignment for z/VSE is ***Average/VS***. The workload used to represent a z/VSE partition can be manually designated as any of the 5 workload categories.
- **KVM** can be defined to General Purpose or IFL partitions. The default workload category assignment for Linux is ***Average/K***. The workload used to represent a KVM partition can be manually designated as any of the 5 workload categories.
- **Linux** can be defined to General Purpose or IFL partitions. The default workload category assignment for Linux is ***Average/L***. The workload used to represent a Linux partition can be manually designated as any of the 5 workload categories.
- **zAware** can be defined to General Purpose or IFL partitions on z12 and z13 processors. The workload category assignment is zAware. zAware is intended to be replaced by SSC.
- **zACI** can be defined to General Purpose or IFL partitions on z13 processors. The workload category assignment is zACI. zACI is intended to be replaced by SSC.

- **SSC** can be defined to General Purpose or IFL partitions on z13 and later processors. The only workload category assignment is SSC. SSC is intended to be the new support structure for zAware and zACI on these models.
- **CFCC** can be defined to General Purpose or ICF partitions. The only workload category assignment is CFCC. Note that CFCC capacity is not based on the z/OS LSPR data, but rather determined by another measurement and analysis process. Each hardware family is supported by its own unique CFCC level.

LSPR Processor Capacity Ratios

IBM GP Multi-Image Table

LSPR Capacity Ratio Table

Workload Graph Help

z/OS-2.4 LSPR Data (04/04/2023)

LSPR Multi-Image Capacity Ratios
IBM Z General Purpose CPs

Values are applicable for z/OS; representative of z/VM, KVM, and Linux

Capacity basis: 2094-701 @ 559.792 MIPS for a typical multi-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

IBM Z Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
3931-636	36W	=	4,281	42,148	39,133	36,521	34,629	32,923
3931-637	37W	=	4,385	43,172	40,084	37,408	35,464	33,713
3931-638	38W	=	4,489	44,192	41,032	38,294	36,298	34,500
3931-639	39W	=	4,593	45,207	41,978	39,180	37,131	35,286
z16/700								
3931-701	1W	=	278	2,255	2,254	2,253	2,204	2,157
3931-702	2W	=	529	4,422	4,357	4,294	4,165	4,044
3931-703	3W	=	771	6,527	6,405	6,288	6,063	5,853
3931-704	4W	=	1,008	8,570	8,400	8,237	7,900	7,589
3931-705	5W	=	1,232	10,583	10,351	10,129	9,683	9,274
3931-706	6W	=	1,449	12,565	12,258	11,965	11,413	10,910
3931-707	7W	=	1,659	14,518	14,122	13,748	13,094	12,499
3931-708	8W	=	1,866	16,442	15,945	15,478	14,725	14,042
3931-709	9W	=	2,061	18,336	17,727	17,157	16,309	15,540
3931-710	10W	=	2,253	20,203	19,470	18,788	17,846	16,995
3931-711	11W	=	2,442	22,041	21,173	20,371	19,339	18,407
3931-712	12W	=	2,625	23,851	22,838	21,907	20,788	19,778
3931-713	13W	=	2,793	25,635	24,465	23,398	22,195	21,110
3931-714	14W	=	2,953	27,392	26,057	24,846	23,561	22,403
3931-715	15W	=	3,108	29,122	27,612	26,251	24,887	23,659
3931-716	16W	=	3,258	30,826	29,132	27,615	26,175	24,878
3931-717	17W	=	3,404	32,505	30,619	28,939	27,425	26,062
3931-718	18W	=	3,550	34,175	32,099	30,261	28,673	27,242
3931-719	19W	=	3,703	35,836	33,575	31,583	29,918	28,420
3931-720	20W	=	3,857	37,488	35,046	32,903	31,162	29,595

Processor models in table = 2,797; In this view = 317; Currently selected = 5

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Normal Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
Select multiple processors with Ctrl+LeftClick or Shift+LeftClick; For flag explanation, position mouse on indicator

The **Multi-Image Capacity Ratios** table for **IBM Z General Purpose CPs** is accessed from the Control Panel by clicking the icon on the **LSPR Multi-Image Processor Table** tree.

LSPR Processor Capacity Ratios Table

IBM IFL Multi-Image Table

LSPR Capacity Ratio Table

Workload Graph Help

z/OS-2.4 LSPR Data (04/04/2023)

LSPR Multi-Image Capacity Ratios
IBM Z IFL CPs

Values are representative of z/VM, KVM, and Linux

Capacity basis: 2094-701 @ 559,792 MIPS for a typical multi-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

IBM Z Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
z16/700								
3931-7xx I1	1W IFL	=	2,255	2,254	2,253	2,204	2,157	
3931-7xx I2	2W IFL	=	4,422	4,357	4,294	4,165	4,044	
3931-7xx I3	3W IFL	=	6,527	6,405	6,288	6,063	5,853	
3931-7xx I4	4W IFL	=	8,570	8,400	8,237	7,900	7,589	
3931-7xx I5	5W IFL	=	10,583	10,351	10,129	9,683	9,274	
3931-7xx I6	6W IFL	=	12,565	12,258	11,965	11,413	10,910	
3931-7xx I7	7W IFL	=	14,518	14,122	13,748	13,094	12,499	
3931-7xx I8	8W IFL	=	16,442	15,945	15,478	14,725	14,042	
3931-7xx I9	9W IFL	=	18,336	17,727	17,157	16,309	15,540	
3931-7xx I10	10W IFL	=	20,203	19,470	18,788	17,846	16,995	
3931-7xx I11	11W IFL	=	22,041	21,173	20,371	19,339	18,407	
3931-7xx I12	12W IFL	=	23,851	22,838	21,907	20,788	19,778	
3931-7xx I13	13W IFL	=	25,635	24,465	23,398	22,195	21,110	
3931-7xx I14	14W IFL	=	27,392	26,057	24,846	23,561	22,403	
3931-7xx I15	15W IFL	=	29,122	27,612	26,251	24,887	23,659	
3931-7xx I16	16W IFL	=	30,826	29,132	27,615	26,175	24,878	
3931-7xx I17	17W IFL	=	32,505	30,619	28,939	27,425	26,062	
3931-7xx I18	18W IFL	=	34,175	32,099	30,261	28,673	27,242	
3931-7xx I19	19W IFL	=	35,836	33,575	31,583	29,918	28,420	
3931-7xx I20	20W IFL	=	37,488	35,046	32,903	31,162	29,595	
3931-7xx I21	21W IFL	=	39,131	36,512	34,222	32,403	30,767	
3931-7xx I22	22W IFL	=	40,765	37,973	35,539	33,642	31,936	
3931-7xx I23	23W IFL	=	42,390	39,429	36,855	34,878	33,102	
3931-7xx I24	24W IFL	=	44,006	40,881	38,170	36,113	34,266	

Processor models in table = 1,301; In this view = 200; Currently selected = 5

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Normal Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
Select multiple processors with Ctrl+LeftClick or Shift+LeftClick; For flag explanation, position mouse on indicator

The **Multi-Image Capacity Ratios** table for **IBM Z IFL CPs** is accessed from the **Control Panel** by clicking the  icon on the **LSPR Multi-Image Processor Table** tree.

LSPR Processor Capacity Ratios Table

LinuxONE Multi-Image Table

LSPR Capacity Ratio Table

Workload Graph Help

z/OS-2.4 LSPR Data (04/04/2023)

LSPR Multi-Image Capacity Ratios
IBM LinuxONE CPs

Values are representative of z/VM, KVM, and Linux

Capacity basis: 2094-701 @ 559,792 MIPS for a typical multi-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LinuxONE Processor **Features** **Flag** **MSU** **Low** **Low-Avg** **Average** **Avg-High** **High**

Emperor 4 3931/400

Processor	Features	Flag	MSU	Low	Low-Avg	Average	Avg-High	High
3931-401_L	1W GP CP	=	35	281	280	280	274	268
3931-7xx L1	1W IFL	=		2,255	2,254	2,253	2,204	2,157
3931-7xx L2	2W IFL	=		4,422	4,357	4,294	4,165	4,044
3931-7xx L3	3W IFL	=		6,527	6,405	6,288	6,063	5,853
3931-7xx L4	4W IFL	=		8,570	8,400	8,237	7,900	7,589
3931-7xx L5	5W IFL	=		10,583	10,351	10,129	9,683	9,274
3931-7xx L6	6W IFL	=		12,565	12,258	11,965	11,413	10,910
3931-7xx L7	7W IFL	=		14,518	14,122	13,748	13,094	12,499
3931-7xx L8	8W IFL	=		16,442	15,945	15,478	14,725	14,042
3931-7xx L9	9W IFL	=		18,336	17,727	17,157	16,309	15,540
3931-7xx L10	10W IFL	=		20,203	19,470	18,788	17,846	16,995
3931-7xx L11	11W IFL	=		22,041	21,173	20,371	19,339	18,407
3931-7xx L12	12W IFL	=		23,851	22,838	21,907	20,788	19,778
3931-7xx L13	13W IFL	=		25,635	24,465	23,398	22,195	21,110
3931-7xx L14	14W IFL	=		27,392	26,057	24,846	23,561	22,403
3931-7xx L15	15W IFL	=		29,122	27,612	26,251	24,887	23,659
3931-7xx L16	16W IFL	=		30,826	29,132	27,615	26,175	24,878
3931-7xx L17	17W IFL	=		32,505	30,619	28,939	27,425	26,062
3931-7xx L18	18W IFL	=		34,175	32,099	30,261	28,673	27,242
3931-7xx L19	19W IFL	=		35,836	33,575	31,583	29,918	28,420
3931-7xx L20	20W IFL	=		37,488	35,046	32,903	31,162	29,595
3931-7xx L21	21W IFL	=		39,131	36,512	34,222	32,403	30,767
3931-7xx L22	22W IFL	=		40,765	37,973	35,539	33,642	31,936
3931-7xx L23	23W IFL	=		42,390	39,429	36,855	34,878	33,102

Processor models in table = 906; In this view = 201; Currently selected = 5

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Normal Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
Select multiple processors with Ctrl+LeftClick or Shift+LeftClick; For flag explanation, position mouse on indicator

The **Multi-Image Capacity Ratios** table for **LinuxONE** CPs is accessed from the **Control Panel** by clicking the  icon on the **LSPR Multi-Image Processor Table** tree. The **LinuxONE** model names in this table are intentionally unique from the **IBM Z** model names.

LSPR Processor Capacity Ratios Table

LSPR Capacity Ratio Tables Overview

The Multi-Image table and the Single-Image table both provide **LSPR Processor Capacity Ratios** representing the capacity relationship of each processor, relative to the current **Reference-CPU** for each workload category displayed. The workload categories that are displayed are user controlled from the **LSPR Table Control** window under **Settings** on the menu-bar.

The window's title area specifies **General Purpose CPs** or **IFLs**, and includes the **Reference-CPU** basis for the capacity values.

The LSPR table of processors columns are displayed as follows:

- **Processor** Specific processor model identification
- **Features** N-way and other notable information
- **Flag** Exception or additional information indicator
- **MSU** Rating used for software pricing

Note: MSU ratings are applicable for General Purpose CPs only.

LSPR Workload Category

- **Low** Relative capacity for Low RNI workload
- **Low-Avg** Relative capacity for Low-Avg RNI workload
- **Average** Relative capacity for Average RNI workload
- **Avg-High** Relative capacity for Avg-High RNI workload
- **High** Relative capacity for High RNI workload

Several exception settings are provided for the purpose of changing the perspective of the capacity values displayed in the table. For detailed discussion, see [LSPR Table Control](#) below.

Note: For all z10 and later processor models, **z/OS capacity is represented with HiperDispatch turned ON**. A white paper concerning **HiperDispatch** is available at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101229>

LSPR Multi-Image Capacity Ratios Table

Three separate tables are available:

- **IBM Z** General Purpose CPs
- **IBM Z** IFL CPs
- **LinuxONE** CPs (includes IFL CPs and a single GP CP)

Capacity projections in the **LSPR Multi-Image Capacity Ratios** table assume a multiple-partition configuration deemed typical for each processor N-way model. As the N-way of the LPAR host increases, the number of partitions increases, and the number of LCPs defined to the primary partitions increases, while the LCP:RCP ratio diminishes.

For a 1-way processor model, a shared 5-partition configuration is assumed. Therefore, the **Reference-CPU** single-image scaling-factor is adjusted down to consider the cost of managing 5 partitions over that of managing 1 partition. The adjustment is simply the multiplication of 0.944 times the scaling-factor assumed for a single partition. The **Reference-CPU** processor model will appear in the **Multi-Image Table** with this adjusted capacity value.

Any of the 5 LSPR workload categories can be displayed in the **LSPR Multi-Image Capacity Ratios** table. While these capacity values can be considered to accurately represent z/OS partitioned environments, they may also be viewed as being reasonable for z/VM, Linux, and z/VSE partitioned environments. Capacity values in each workload category column assume that the same workload is running in every partition.

Capacity values in the LSPR Multi-Image Capacity Ratios table represent typical (or average) partition configurations, and are therefore a generalization of processor capacity.

The capacity values for any specific LPAR configuration can deviate significantly (higher or lower) from those shown in the LSPR Multi-Image Capacity Ratios table due to the degree it differs from the typical configuration represented in the table.

Capacity differences can be due to any of the items listed below.

1. The following items may vary from that assumed as a "Typical" LPAR configuration:
 - Number of partitions defined
 - Number of LCPs defined to each partition
 - Relative weights assigned to each partition
 - Capping of partitions
 - The workload category assigned to each partition may vary
 - Use of zAAP or zIIP CPs in support of z/OS
 - Additional partition types defined
 - SCPs other than z/OS (z/VM, z/VSE, KVM, Linux, CFCC) may be assigned to some partitions.

Important Note: When multiple partition types are defined (i.e., GP, IFL, and ICF), contention will exist between the real CP types. If the LPAR host configuration includes multiple books or drawers, that contention will be reduced, thus improving available capacity. The effect of contention is not reflected in the **LSPR Multi-Image table**. **zPCR's LPAR Configuration Capacity Planning Function** must be used to reveal the effect of such contention on capacity..

2. When zAAP CPs are configured to support a portion of the z/OS workload content, z/OS capacity will be impacted. Capacity results for other partitions will also be impacted slightly. Capacity projections for the zAAP CPs are provided.
3. When zIIP CPs are configured to support a portion of the z/OS workload content, z/OS capacity will be impacted. Capacity results for other partitions will also be impacted slightly. Capacity projections for the zIIP CPs are provided.
4. When IFLs are configured to run z/VM, KVM, Linux, zAware, zACI, or SSC, capacity results for other partition types will be impacted. Capacity projections for the IFL partitions are provided.

5. When ICFs are configured to run CFCC, the capacity results for other partition types will be impacted. Capacity projections for the ICF partitions are provided.
6. When drawers on IBM z17, z16, z15, z14, and z13 processors or books on z12, z11, z10, and z9 processors are configured in excess of that needed for the RCPs required (i.e., unused drawers/books), capacity results for all partition types will be impacted slightly.
7. For z10 and later processor models, the **HiperDispatch** effect of keeping workload content on the same RCPs is inherent in the LSPR capacity data. However, any potential parking of LCPs is ignored. If parked LCPs were removed from the configuration, the reduced LCP:RCP ratio will result in slightly improved capacity.

The impact of any of these items (or some combination thereof) on capacity can be significant, thus providing capacity results that are quite different from those observed in the **LSPR Multi-Image Capacity Ratios** table.

To get a more reliable capacity assessment for any specific LPAR configuration, **zPCR**'s **LPAR Configuration Capacity Planning** function should be used. Any legitimate LPAR configuration can be modeled, with capacity results provided for each partition, for each RCP type (i.e., CP pool), and for the LPAR host as a whole.

The **LSPR Multi-Image Capacity Ratios** table plays no role in **zPCR**'s **LPAR Configuration Capacity Planning** function. Rather, the **LPAR Configuration Capacity Planning** function is based on the **LSPR Single-Image Capacity Ratios** table. This function is then used to compute the capacity values shown in the **LSPR Multi-Image Capacity Ratios** table.

LSPR Single-Image Capacity Ratios Table

Two separate tables are available:

- **IBM Z** General Purpose CPs
- **IBM Z** IFL CPs
- **LinuxONE** CPs (includes IFL CPs and a single GP CP)

Capacity projections in the **LSPR Single-Image Capacity Ratios** table assume a single shared partition configuration, running the indicated SCP/workload. Capacity projections in this table should be considered as having minimal LPAR overhead and no contention between partitions for sharing RCP resource. This table is a reflection of the z/OS workload environments that are considered measured for LSPR purposes.

This table is intended to show reasonable Single-Image partition configurations, and is therefore limited to displaying a maximum of 30 CPs.

The Single-Image LSPR table is not intended to imply what is supported by any given SCP. Rather, it is an indication of the capacity available if it could be supported.

LSPR data is currently based on z/OS-2.4 measurement data. While the capacity ratios are based on z/OS, they can also be considered representative for other z/OS versions, for various z/VM versions, and for all Linux, KVM, and z/VSE versions.

SCP names used when defining partitions:

- **z/OS:** When a non-measured version is specified, an “*” is appended to the name indicating that it is other than the measured version. z/OS version numbers are used to enforce rules concerning features that are supported.
- **z/VM:** Version 7.2 is assumed to be the measured version (even though z/OS ITRRs are used). When any other version is specified, an “*” is appended to the name indicating that it is other than the version considered measured. z/VM version numbers are used to enforce rules concerning features that are supported.
- **Linux, KVM, and z/VSE:** Version numbers are not considered. Any features supported must be considered outside of **zPCR**.

For all **IBM Z** and **LinuxONE** processor models included in the analysis, you must assume that the capacity relationships represent LPAR-mode, since that is the only way that these processors can operate.

The values contained in the **LSPR Single-Image Capacity Ratios** tables are used as the capacity basis for the algorithms in **zPCR's LPAR Configuration Capacity Planning** function.

Concerning both the Multi-Image and the Single-Image tables

Beneath the table is listed a quantification of the processors included in the current view, in the entire list box, and those that have been selected.

LSPR capacity ratios are available for General Purpose CPs and IFLs. To determine capacity for zAAPs, zIIPs, and ICFs, you must use the **LPAR Configuration Capacity Planning** function. LPAR configurations with any valid mix of General Purpose CPs and specialty engines can be characterized with this planning function.

One should not try to compare capacity values between the **LSPR Multi-Image Table** and the **LSPR Single-Image Table**. For the **Multi-Image Table**, the LPAR configuration grows increasingly complex as the N-way of the host increases, while the LPAR configuration represented by the **Single-Image Table** remains constant (a single partition) regardless of the N-way of the host. Therefore, direct comparison of results from the multi-image table to those from the single-image table is a worthless exercise. Since all contemporary processors run with multiple images, the **LSPR Multi-Image Capacity Ratios** table should be the preferred reference to generalize on capacity. The **LPAR Configuration Capacity Planning** function should be used to refine the capacity expectation for any specific LPAR host and its partition configuration.

LSPR Table Control Window

Whenever an **LSPR Table** window is displayed, an associated **LSPR Table Control** window also appears beside it, providing quick specification of the processor models, workload categories, and capacity exceptions to be displayed in the table. See [LSPR Table Control](#) for detailed information.

LSPR Table (continued)

Selecting Processors

Individual processor models can be selected from the table for the purpose of reducing the number being viewed to only those of interest. In addition, processors which are selected can be used to generate a variety of graphs.

No processor are selected when the table is initially displayed. To select processors, you must hold the **Ctrl** key while clicking anywhere on that processor row. To unselect an already selected processor row, hold the **Ctrl** key and click again on that processor row. If you click without holding the **Ctrl** key that row becomes the sole selection. To select a contiguous range of processors, click on the first processor model, and then, holding the **Shift** key, click on the last processor model.

Controlling the Processor Table Sequence

When viewing **All Families** or **Selected Families**, the processor table may be sorted on the **Processor Name** column, the **MSU** column, or on any **Workload** column by clicking the underlined column title. The first click will show ascending sequence, the second click will show descending sequence, and the third click will restore the default table sequence. An indicator is displayed next to the column title to indicate when processors are sorted and the sort order.

Whenever a processor table is in a sorted sequence, the processor family headers no longer serve a purpose, and therefore, are not displayed.

When viewing **Favorites**, the LSPR table cannot be sorted. However, the order of the processor models in the **Favorites** list can be changed using controls at the bottom of the **LSPR Table Control** window; the LSPR table will immediately be presented in that order.

Selected processors are always displayed on graphs in the same order as they are displayed in the table.

Controls on the LSPR Table Window

Click the **Provisional Reference-CPU** button on either the **Multi-Image** or the **Single-Image Capacity Ratios** table window to temporarily change the **Reference-CPU** model, scaling, and/or metrics while viewing that **LSPR Table** window. Unlike the **Reference-CPU**, the **Provisional Reference-CPU** can be set to any N-way processor model in the table. When the **LSPR Table** window is changed or closed, the normal **Reference-CPU** settings are restored.

Alternatively, when viewing **General Purpose CP** capacity, you can double click on any specific processor in the table to open the **Provisional Reference-CPU** window with that processor already selected. This technique does not work when viewing the IFL Engine capacity table, since the models displayed there cannot be named as a **Reference-CPU** processor model.

Note: Setting the **Provisional Reference-CPU** will cause z14, z13, zEC12 and z196 processor models, if being displayed in **Power-Save** mode, to revert to **Full Power**. If **Power-Save** mode is desired, it can be restored from the **LSPR Table Control** window by clicking **Settings** → **Capacity Exceptions** on the menu-bar.

Click the **Workload Categories** button to view the **Workloads** window, which will display a table of the 5 LSPR workload categories and their naming convention when defined to partitions running SCPs other than z/OS.

Click the **Copy Selected to Favorites** button to copy selected processor models to the **Favorites** list. Processor models selected that are already in the **Favorites** list will be ignored. If the selected processors will cause the **Favorites** list to exceed its maximum of 20, a dialog will appear and only the ones that add up to the maximum will be added.

Click the **Return** toolbar icon to return to the **Control Panel** window. Both the **LSPR Capacity Ratio Table** window and the **Table Control** windows will be closed.

Click the **HTM** toolbar icon to create an HTML file with the report tables.

When displaying the Single-Image LSPR Table with General-Purpose CPs, the table can be output as a flat file (this capability is in support of the **zTPM** tool). From the menu-bar, click **File** → **Save as PRN** to create the **PRN** file output

Menu-bar

File	A special case only for the Single-Image table when displaying General-Purpose CPs. The LSPR table can be output in a flat file format (this capability is in support of the zTPM tool). Click Save as PRN to create the PRN file output.
Workload	
Describe	Display a short description for the 5 LSPR workload categories.
Graph	Create charts for processors that have been selected.
Workload	Generate capacity bars showing capacity for a specific LSPR workload category.
All Workloads	Generate capacity bars showing capacity for all of the published LSPR workload categories side-by-side.
Response Time	Generate line graph showing relative internal processor response time based on transaction rate or on utilization. A specific workload category must be selected. See discussion below.
Help	

Information Concerning the Response Time Graphs

Two graphs are available:

- Processor response time based on utilization
- Processor response time based on transaction rate

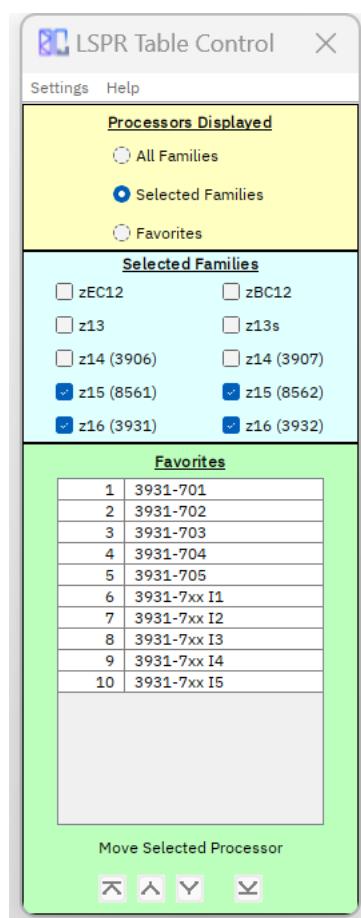
For these charts, the currently defined **Reference-CPU** is assumed to have a transaction service time (or response time) of 10 mille seconds at a low transaction rate or a low utilization. Then, for each of the processors selected, the relative (theoretical) response time is plotted, showing the change as the transaction rate or utilization increases. The response time plotted is based solely on CPU time (no other potential delays assumed).

If the **Reference-CPU** is changed, that processor is then assumed to have a transaction service time (or response time) of 10 mille seconds. The **Provisional Reference-CPU** may be used to temporarily change the **Reference-CPU** setting. The original setting will be restored when the **LSPR Capacity Ratio Table** window is closed.

LSPR Table Control

Whenever an **LSPR Capacity Ratio Table** window is displayed, the **LSPR Table Control** window appears on the right. Its purpose is to quickly select what is to be displayed in the LSPR table. The settings are common to the **Multi-Image** table and the **Single-Image** table, for both **GP CPs** and **IFL CPs**. The primary settings (with a few exceptions) remain persistent throughout the **zPCR** invocation. Settings that are active at the time will be saved with a **zPCR** study. Settings to be used by default for a new study can be defined as **zPCR** preferences.

Controlling the LSPR Table window



Processors Displayed group box

- **All Families** – All **IBM Z** families included in the current LSPR table will be displayed. This is the only setting that will allow viewing of the older z800, z900, z890, and z990 families.
- **Selected Families** – Only the **IBM Z** families that are checked under **Selected Families** will be displayed. Checkboxes are provided for only the more current **IBM Z** families. Checkboxes may only be altered when viewing **Selected Families**.
- **Similar CPCs** – Find processors with similar capacity or MSU rating (available for Multi-Image General Purpose LSPR table only). See [Similar CPCs](#) for a detailed discussion.
- **Favorites** - Only the **IBM Z** processor models that have been copied to the **Favorites** list will be displayed. To be selected, there must be at least one processor in the **Favorites** list. The **Favorites** list may contain both GP and IFL models.

Selected Families group box provides checkboxes to control which **IBM Z** families will be visible when **Processors Displayed** is set to **Selected Families**. Check the box for each of the processor family you wish to be displayed in the LSPR table. The LSPR Table window is updated dynamically as boxes are checked or unchecked. At least one processor family must be checked.

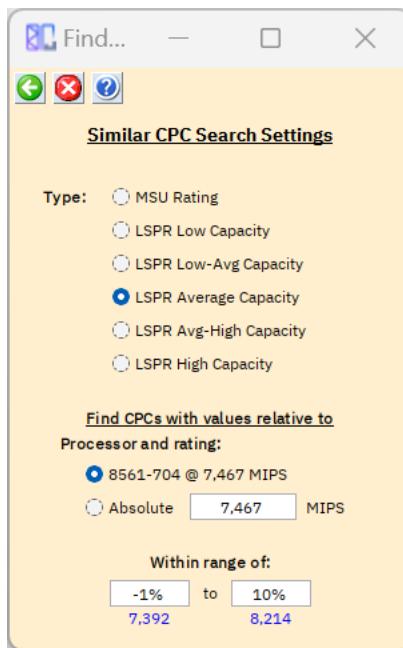
Favorites group box displays a user-specified set of processor models in the LSPR table when **Processors Displayed** is set to **Favorites**. When displaying **Favorites**, all the processors presented in the LSPR table are considered to be selected (i.e., graphs will include all of them). See [Favorites List](#) for detail.

Find Similar CPCs

The **Find Similar CPCs** function provides the capability to search for processors of similar capacity or similar MSU rating. The function is available only when viewing the **Multi-Image General Purpose LSPR Table** window (it is not available for IFLs, for the **LinuxONE LSPR Table**, or for any of the **Single-Image LSPR Tables**).

The search is limited to the set of processors currently being displayed. Select either **All Families** or **Selected Families** under **Processors Displayed**. For **Selected Families**, check the processor families that are to be included. The **Similar CPCs** radio button is always disabled when viewing the **Favorites List**. It is also disabled when displaying the **Single-CP** capacity exception case.

To access **Similar CPCs** capability, a single processor model must be selected to be used for the comparison. Right click on that processor row and select **Show Similar CPCs** from the pop up menu. Once selected, the **Similar CPCs** radio button on the **LSPR Table Control** window is enabled and selected and the **Find Similar CPCs** window appears.



A search may be made for a specific LSPR workload category (based on the current **Reference-CPU** settings) or on a MSU rating. This selection is made with the radio buttons under **Type** at the top of the window.

The selected processor becomes a reference point for the search, identified under **Processor and rating**. The rating shown and scaling-metric will be based on the **Type** selection above.

The row beneath the selected processor affords the capability to specify an **Absolute** value for the search. The field is automatically set to the value for the selected processor. Click on the field to manually enter an **Absolute** value. Once a value has been entered, that value will persist.

Under **Within range of**, percent values are used to control the scope of the search. Click on a field to replace the default values. Values may be between -50% and 100%; the minimum must be at least 1% less than the maximum.

Click the **Return** tool bar icon to initiate the search. The **Multi-Image LSPR Table** window will now display the search results, sorted based on the search **Type** requested. Click **Cancel** to skip the search and restore the previous **Processors Displayed** view.

LSPR Capacity Ratio Table

Workload Graph Help

z/OS-2.4 LSPR Data (04/04/2023)

LSPR Multi-Image Capacity Ratios
IBM Z General Purpose CPs (Similar CPCs)
 7 CPCs in the range of 7,392 to 8,214 MIPS for "Average" are displayed
 Capacity basis: 2094-701 @ 559,792 MIPS for a typical multi-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

IBM Z Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
3931-438	38W	=	912	8,458.3	7,923.4	7,452.2	7,125.7	6,826.6
8561-704	4W	=	914	7,883.1	7,669.2	7,466.6	7,083.6	6,737.9
3931-439	39W	=	933	8,654.8	8,107.9	7,626.1	7,290.4	6,983.1
8561-512	12W	=	941	8,420.9	8,039.1	7,690.4	7,258.9	6,873.3
8561-608	8W	=	966	8,487.0	8,177.8	7,890.3	7,450.0	7,056.1
3931-510	10W	=	968	8,470.1	8,178.6	7,906.5	7,545.1	7,215.4
3931-606	6W	=	980	8,351.3	8,175.0	8,005.9	7,626.7	7,281.8

Processor models in table = 2,797; In this view = 7; Currently selected = 1

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Normal Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
 Select multiple processors with Ctrl+LeftClick or Shift+LeftClick; For flag explanation, position mouse on indicator

The window title is updated to reflect the **Similar CPCs** view and to show the search range that was applied and to show the number of hits. Columns remain available for sort. Graphs are available. Processor models may be selected and copied to the **Favorites List** (you may want to change the sort order before copying). The **Settings** menu-bar item on the **LSPR Table Control** window is disabled while this window is displayed.

If the search finds no matching CPCs, a dialog is presented.

Note: The **Provisional Reference-CPU** setting is not directly available from this window. To apply a **Provisional Reference-CPU** setting, copy the search results to the **Favorites List**, display **Favorites**, and then apply the **Provisional Reference-CPU** setting.

To exit the **Similar CPCs** view without leaving the **LSPR Table** window, select any of the other **Processors Displayed** radio buttons. To exit the **LSPR Table** window entirely, click the **Return** tool bar icon on that window (all **Similar CPCs** settings will be restored to the default values).

To restart **Similar CPCs** while remaining in the LSPR Table window, first choose either **Selected Families** or **All Families**, and then choose **Similar CPCs**. In this case, previous user settings will persist.

When selecting the **Similar CPCs** radio button, if the previously selected processor is no longer found in the displayed LSPR table, a dialog will suggest the selection of a different processor that is included in the table, or to start from the **All Families** view, which includes every one of the IBM mainframe processor models.

Settings defined for **Find Similar CPCs** are not saved when the **LSPR Multi-Image Table** window is closed.

Favorites List

The **Favorites** list is user-controlled, with a maximum of 20 processor model entries. The following function is available:

- **Copy** a processor to the **Favorites** list – From the **LSPR Table** window, select one or more processor rows and click the **Copy Selected to Favorites** button. Multiple processor models can be selected using Ctrl-Click (or Shift-Click). Once processor models have been selected, you may also right click and select **Copy to Favorites** from the pop-up menu.
- **Remove** a processor from the **Favorites** list – From the **Favorites** list, right click on a processor model and select **Remove** from the pop-up menu. Multiple processor models can be selected using Ctrl-Click (or Shift-Click). If currently displaying **Favorites**, the LSPR table will be dynamically updated.
- **Reorder** the **Favorites** list by selecting a single processor model, and, using the control buttons below, move the entry to the desired location in the list. If currently displaying **Favorites**, the LSPR table will be dynamically updated.

Note that when displaying Favorites from the Single-Image table, any processor models with more than 30 CPs will be excluded. This is because the Single-Image table is intended to represent reasonable single partition configurations only.

Menu-bar

Settings	See discussion below
Help	

Settings (on the menu-bar) provides some less frequently needed controls for the **LSPR Table** window, including:

- **Workloads Displayed** – By default, all 5 LSPR workload categories are displayed in the **LSPR Table** window. The workload categories to be displayed when starting a new **zPCR** study can be set from the **Preferences** window. These can be refined to any specific workload categories by unchecking/checking the appropriate box. Note that the order in which the workloads are displayed cannot be changed.
- **Capacity Exceptions**: There are 6 capacity exception conditions (listed below). Capacity exceptions are retained only while an **LSPR Table** window is open. For each capacity exception requested, relevant information will appear in the title area of the **LSPR Table** window. All 6 settings will be restored to unchecked when the **LSPR Table** window is closed. If the LSPR table is in a sorted order when any of these items are checked, it is set to unsorted order before returning control to the table (this is because capacity values will change, making the sorted order no longer applicable). **Capacity Exception** settings cannot be defined as **zPCR** preferences.
 1. **Single-CP Capacity**: Checking this box will change the capacity values of all processors displayed in the **LSPR Table** window to represent the capacity of a single CP (by dividing processor total capacity by the number of CPs). When displaying single-CP capacity, capacity values are displayed in **brown** rather than **red**. Note that this setting applies to both the **Multi-Image** and the **Single-Image LSPR** table.
 2. **z196 Power-Save Mode**: Checking this box will change the MSU and capacity values for all z196 processor models to represent **Power-Save** mode (a **green** background will be displayed for all the values that were changed).
 3. **zEC12 Power-Save Mode**: Checking this box will change the MSU and capacity values for all zEC12 processor models to represent **Power-Save** mode (a **green** background will be displayed for all the values that were changed).
 4. **z13 Power-Save Mode**: Checking this box will change the MSU and capacity values for all z13 processor models to represent **Power-Save** mode (a **green** background will be displayed for all the values that were changed).
 5. **z14 Power-Save Mode**: Checking this box will change the MSU and capacity values for all z14 processor models to represent **Power-Save** mode (a **green** background will be displayed for all the values that were changed).
 6. **IBM z17, z16, z15, z14, z13 SMT Benefit** (when viewing IFL models): Checking this box will increase capacity values for those IFL processor models based on a percent improvement expected when exploiting Simultaneous Multi-Threading (SMT). The percent improvement is selected in 1% increments between 0% and 60% using the spin button on the right. The spin button is initially set to the **zPCR** default **SMT Benefit** value for IFLs. Whenever an **SMT Benefit** has been applied, the title line “**IFL CPs**” will indicate the **SMT Benefit** percent (a **yellow** background will be displayed for all the values that were changed).

While **Capacity Exception** settings are immediately applied, the **LSPR Table** window cannot be manipulated until the **Capacity Exceptions** dialog has been closed.

Notes

1. The **Reference-CPU** scaling-factor is always considered to be full power. Only the affected processor families in the LSPR table will be displayed in **Power-Save** mode.
2. The z14, z13, zEC12, and z196 processor families will default to displaying **Full Power** whenever:
 - An **LSPR Table** window is opened
 - The **Provisional Reference-CPU** is set or changed
3. The IBM z17, z16, z15, z14, and z13 IFL processor families will default to displaying capacity without **SMT Benefit** applied. Use the **Capacity Exceptions** setting to factor in expected improvement due to the exploitation of SMT.

- **Restore Default Settings:** Clicking this menu item will perform the following:
 - Restore the **Workloads Displayed** to the 5 LSPR workload categories
 - Check all the processor families boxes under **Selected Families**.
 Processor models copied to the **Favorites** list will be unaffected.
- **Restore Startup Settings:** Clicking this menu item will set all items as defined on the **zPCR Preferences** window. Processor models copied to the **Favorites** list will be unaffected
- **Clear Favorites List:** Clicking this menu item provides the ability to erase the current Favorites List. A dialog will ask to confirm this action.

The **LSPR Table Control** window can be closed by clicking **[X]** on the title bar. If closed, it can be redisplayed by clicking the **Table Controls** button at the bottom of the **LSPR Table** window.

Both the **LSPR Table** window and the **LSPR Table Control** window can be repositioned anywhere on the desktop (dual displays included). The new position will be used each time these windows are opened. When using multiple **zPCR** invocations at the same time, this makes it convenient to keep track of them.

Important Note: When multiple partition types are defined (i.e., GP, IFL, and ICF), contention will exist between the real CP types. If the LPAR host configuration includes multiple books or drawers, the degree of contention will be reduced, thus improving available capacity. The effect of contention is not reflected in the **LSPR Multi-Image table**. **zPCR's LPAR Configuration Capacity Planning Function** must be used to reveal the effect of such contention on capacity.

IBM Z Single-Image Table

LSPR Capacity Ratio Table

File Workload Graph Help

HTM

z/OS-2.4 LSPR Data (04/04/2023)

LSPR Single-Image Capacity Ratios

IBM Z General Purpose CPs

Values represent LSPR data for z/OS

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LSPR Workload Category

IBM Z Processor	Features	Flag	MSU	LSPR Workload Category				
				Low	Low-Avg	Average	Avg-High	High
3931-628	28W	=	3,452	36,553	34,163	32,065	30,641	29,337
3931-629	29W	=	3,553	37,626	35,127	32,939	31,475	30,135
3931-630	30W	=	3,655	38,685	36,077	33,799	32,296	30,921
z16/700								
3931-701	1W	=	278	2,415	2,413	2,412	2,359	2,309
3931-702	2W	=	529	4,740	4,680	4,621	4,489	4,364
3931-703	3W	=	771	7,032	6,908	6,788	6,566	6,357
3931-704	4W	=	1,008	9,293	9,100	8,914	8,591	8,291
3931-705	5W	=	1,232	11,522	11,255	11,000	10,567	10,167
3931-706	6W	=	1,449	13,712	13,357	13,020	12,480	11,983
3931-707	7W	=	1,659	15,863	15,407	14,976	14,333	13,742
3931-708	8W	=	1,866	17,977	17,407	16,871	16,127	15,446
3931-709	9W	=	2,061	20,054	19,357	18,706	17,865	17,097
3931-710	10W	=	2,253	22,095	21,259	20,483	19,548	18,695
3931-711	11W	=	2,442	24,100	23,113	22,205	21,178	20,243
3931-712	12W	=	2,625	26,070	24,922	23,872	22,757	21,742
3931-713	13W	=	2,793	28,013	26,702	25,509	24,306	23,212
3931-714	14W	=	2,953	29,930	28,454	27,116	25,827	24,655
3931-715	15W	=	3,108	31,821	30,177	28,695	27,320	26,070
3931-716	16W	=	3,258	33,686	31,873	30,245	28,785	27,459
3931-717	17W	=	3,404	35,527	33,542	31,768	30,222	28,821
3931-718	18W	=	3,550	37,342	35,184	33,262	31,634	30,157
3931-719	19W	=	3,703	39,133	36,801	34,730	33,019	31,468
3931-720	20W	=	3,857	40,900	38,391	36,172	34,378	32,753
3931-721	21W	=	4,012	42,643	39,956	37,588	35,712	34,015
3931-722	22W	=	4,166	44,363	41,496	38,978	37,022	35,252
3931-723	23W	=	4,320	46,059	43,012	40,343	38,307	36,466
3931-724	24W	=	4,474	47,732	44,503	41,683	39,568	37,657
3931-725	25W	=	4,629	49,383	45,971	43,000	40,806	38,825
3931-726	26W	=	4,783	51,011	47,415	44,293	42,021	39,971
3931-727	27W	=	4,936	52,618	48,836	45,562	43,214	41,096
3931-728	28W	=	5,090	54,203	50,235	46,809	44,384	42,198
3931-729	29W	=	5,243	55,766	51,611	48,033	45,533	43,280
3931-730	30W	=	5,396	57,308	52,966	49,235	46,660	44,342

Processor models in table = 1,878; In this view = 120; Currently selected = 5

Provisional Reference-CPU Workload Categories Copy Selected to Favorites Table Controls

Normal Reference-CPU is active; double click any processor row to set it as a Provisional Reference-CPU
Select multiple processors with Ctrl+LeftClick or Shift+LeftClick; For flag explanation, position mouse on indicator

The **LSPR Single-Image Capacity Ratios** window can be viewed as follows:

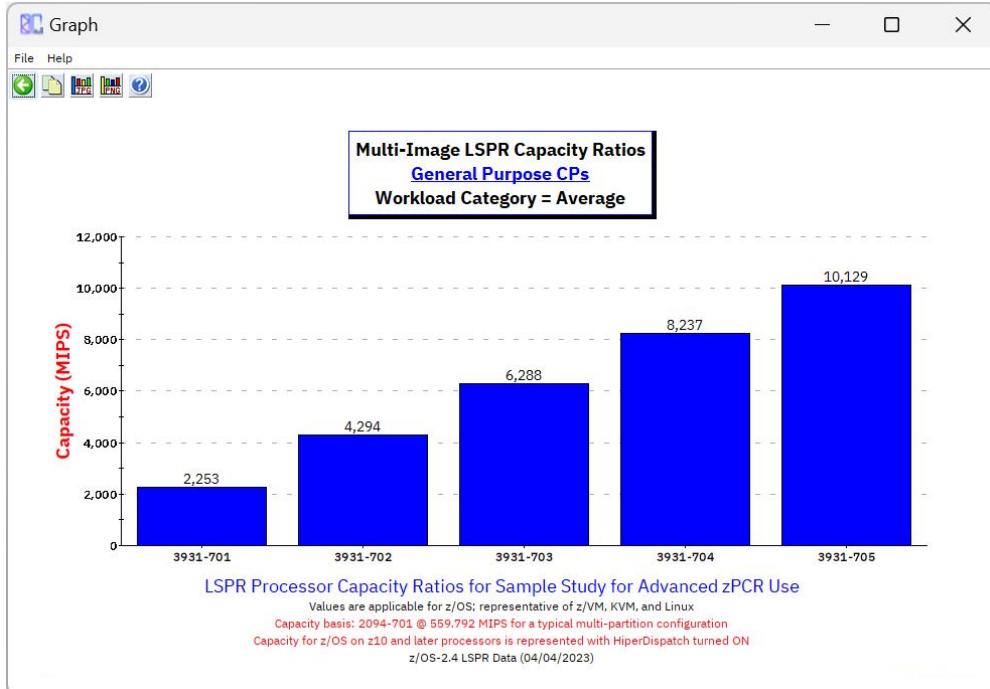
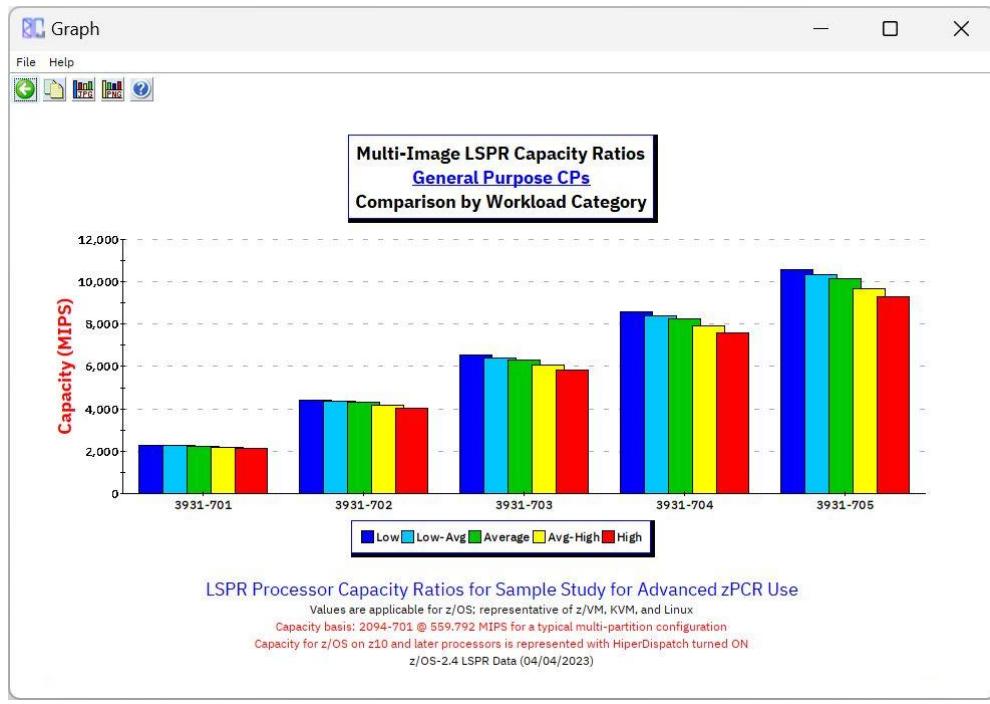
1. On the **Control Panel** right click the **LSPR Multi-Image Processor Table** tab.
2. Click the popup that reads **Switch to LSPR Single-Image Table**.
3. Then select the table to be viewed.

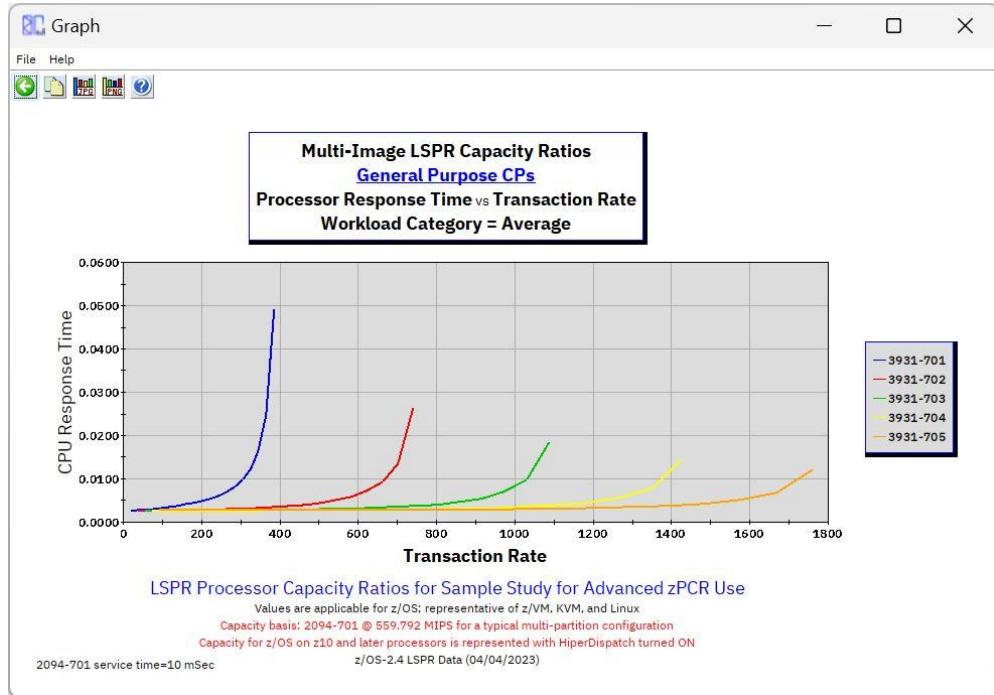
The Single-Image table is intended to show only reasonable Single-Image partition configurations. It is therefore limited to displaying a maximum of 30 CPs.

Once the Single-Image table is closed, the table display default will be restored to Multi-Image.

Charts and Graphs

Several charts are available when viewing any of the **LSPR Processor Capacity Ratio Table** windows by clicking **Graph** on the menu bar. Some examples are shown below.





Graphs are available in the following functions as follows:

- **LSPR Capacity Ratios** tables for General Purpose or IFL CPs
 - Bar graphs showing processor capacity for workload categories
 - Bar graph showing processor capacity for all workload categories
 - Line graphs showing CPU response time for workload categories
- **LPAR Configuration Capacity Planning** function
 - Pie charts showing distribution of capacity by CP pool
 - Bar chart showing distribution of CPs across books or drawers
 - Bar graphs showing processor capacity by CP pool
- **zAAP Capacity** estimator
 - Bar graph showing General Purpose and zAAP capacity
 - Line graph showing General Purpose utilization as zAAP utilization diminishes

Titles on each graph will relate the intended purpose. Either the version of LSPR data or the **zPCR** version used to generate the graph, and the date generated, is provided at the bottom of the graph. All graphs relating capacity results will also include the **Reference-CPU** assumption that was used.

If a **Study Identification** is provided on the **Control Panel** window, that information will be included in the X-axis title.

Most bar graphs will reveal a capacity value at the top of each bar. However, for some bar graphs capacity values are not shown, due to the number of bars involved. In these cases, capacity values can be momentarily displayed by letting the mouse hover on the desired bar (known as fly-over display of text).

Charts and graphs can be captured for documentation purposes in three different formats:

- **Bitmap** Write the chart to the Windows clipboard as a bitmap, which can subsequently be pasted into a document.
- **JPEG** Save the chart as a named **JPG** file; this file can be viewed with various PC applications.
- **PNG** Save the chart as a named **PNG** file; this file can be viewed with various PC applications.

Captured file **JPG** or **PNG** files can be saved in a common directory as a collection of documentation relating to a study.

An alternative method to capture any **zPCR** window, including graphs can also be used. With focus on the desired window, press **Alt-PrintScreen**. The entire active window is copied as a bitmap to the windows clipboard. The contents of the windows clipboard can subsequently be pasted into a document.

A currently captured bitmap must be pulled from the Windows clipboard before another bitmap is captured, or it will be lost. The Windows clipboard can be pasted into the Windows **Paint** utility, and the picture subsequently saved in any of several graphic formats.

Click the **Copy to Clipboard** toolbar icon to create a copy of the chart in the Windows clipboard.

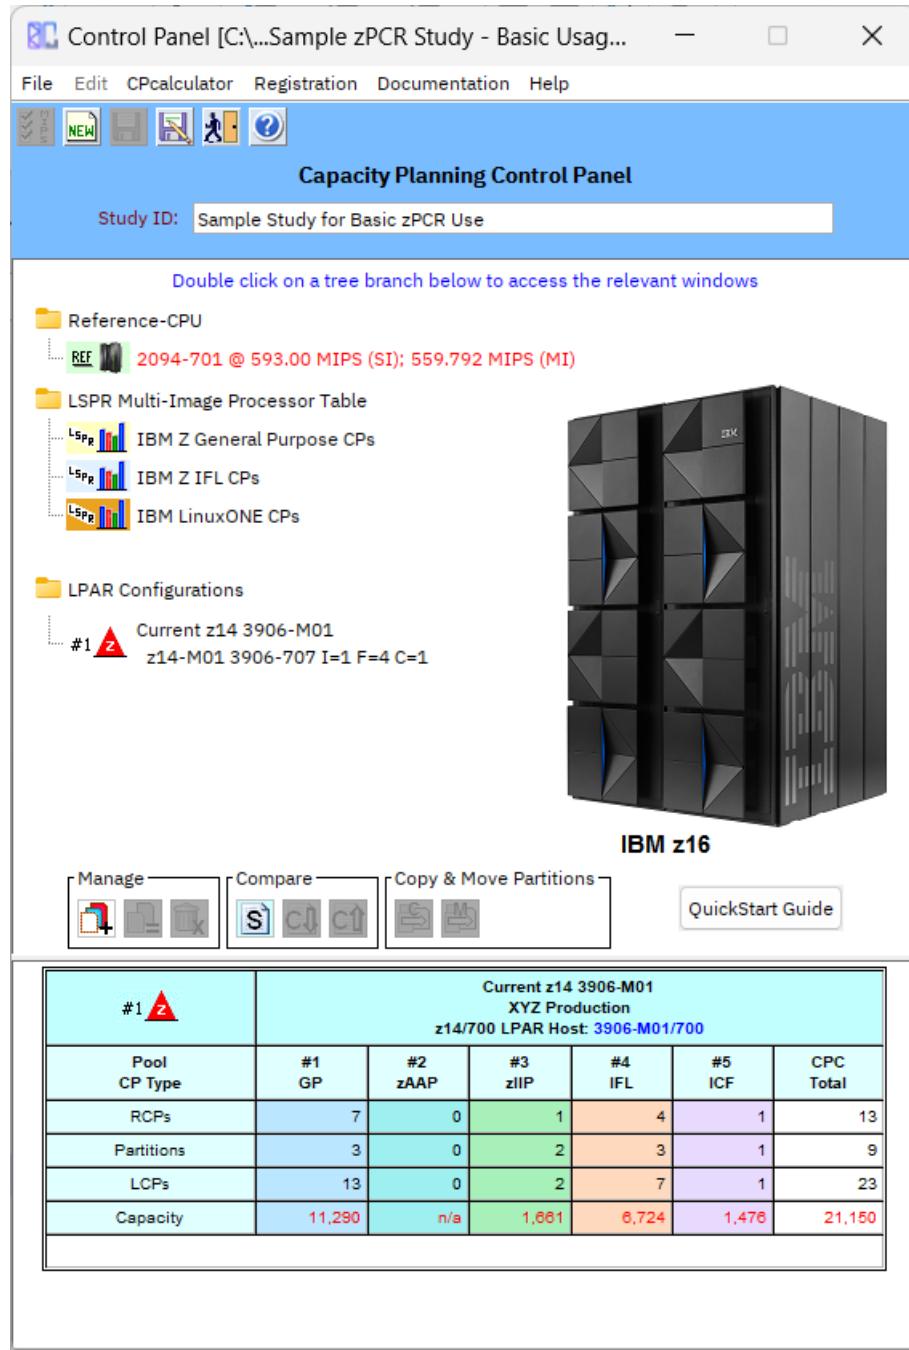
Click the **Create JPG** toolbar icon to save a **JPG** file of the chart.

Click the **Create PNG** toolbar icon to save a **PNG** file of the chart.

Click the **Return** toolbar icon to close the **Graph** window.

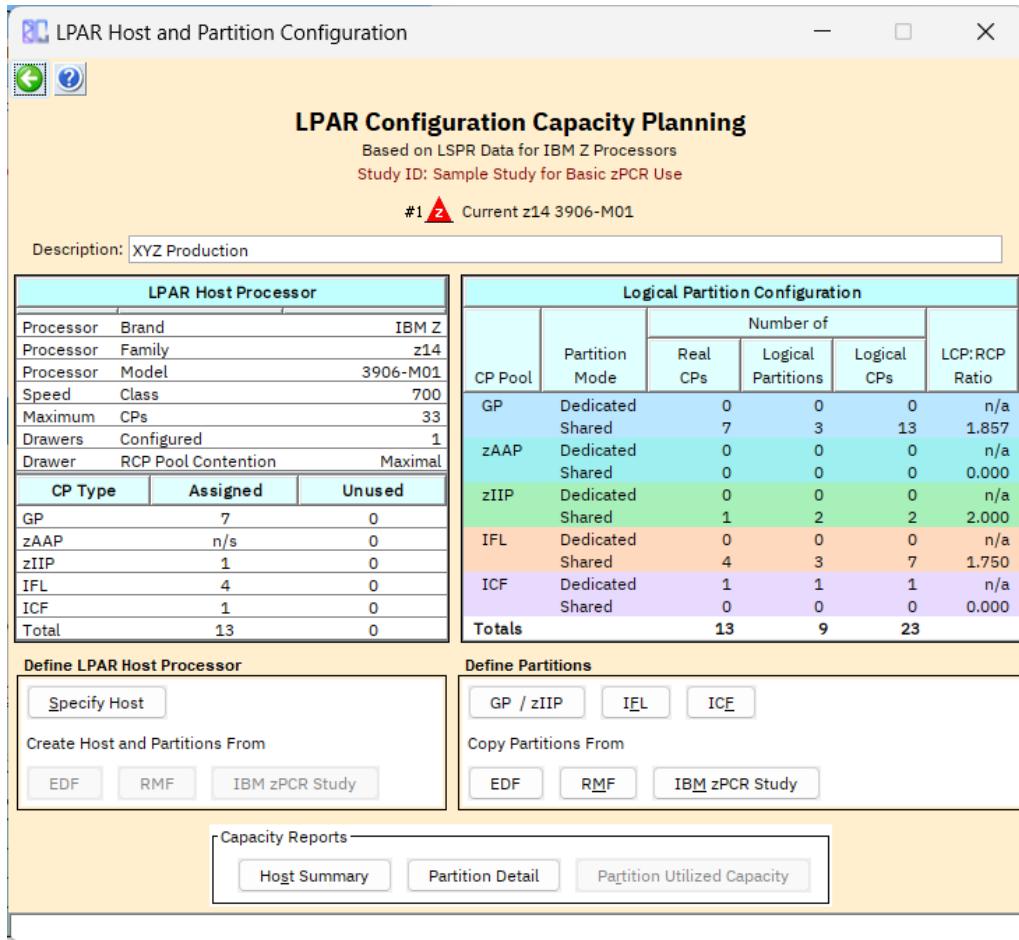
LPAR Configuration Capacity Planning

Note: The **Sample zPCR Study - Basic Usage.zpcr**, included with the **zPCR** package, will be the primary source of the examples shown in this section. Once this study file is loaded the **Control Panel** window will appear as shown below.



LPAR Host and Partition Configuration

LPAR Configuration Capacity Planning



The **LPAR Configuration Capacity Planning** function provides capacity projections for any IBM mainframe LPAR host processor with any specific partition configuration running specific SCP/workload environments in each individual partition. Capacity projections are generated for each partition as well as for the LPAR host as a whole. All capacity results for this function are based on the currently defined **Reference-CPU**.

The **LPAR Host and Partition Configuration** window is the control point for one fully defined LPAR configuration. A **Description** entry field is available at the top with which to provide any informative information concerning the configuration.

When the LPAR configuration is obtained from EDF or RMF, the **Description** field is automatically filled in with the CPC ID, interval #, date, and time. If that description is replaced, the EDF or RMF source identification will be lost. However, this description is saved in a zPCR study, but can only be retrieved by reviewing the study file with a text editor.

The **LPAR Host and Partition Configuration** window is accessed by double-clicking on any defined LPAR icon under the **LPAR configuration** tree. Each icon represents a single LPAR configuration.

An LPAR configuration is created manually using the following steps:

1. **Define the LPAR host:** The LPAR host processor must be specified before any partitions can be defined. In the **Define LPAR Host Processor** group box, click the **Specify Host** button to display the **LPAR Host Processor** window (see [LPAR Host Processor](#)). Upon return, the LPAR host configuration is summarized on the left side of the **LPAR Host and Partition Configuration** window.
2. **Define one or more partitions:** Once the LPAR host has been specified, individual partitions can be defined using the **GP**, **IFL**, and **ICF** buttons in the **Define Partitions** group box. zAAP and zIIP partitions are always associated with a parent GP partition running z/OS. IFL partitions may be associated with a Parent GP partition running z/VM (hardware definition for the partition must be **Mode=z/VM**).

The **LSPR Single-Image Capacity Ratios** table is used exclusively to support this function (the multi-image table plays no role). Since all **zPCR** function is based on the same **Reference-CPU** setting, one can jump between LPAR configurations and the LSPR tables and continue to observe capacity values that are all on the same scale.

Several generations of IBM mainframe families are supported by **zPCR** as LPAR hosts. The SCP assigned to a partition will dictate the workload categories assignable.

- **z/OS** 5 z/OS workload categories
- **z/VM** 5 z/VM workload categories
- **z/VSE** 5 z/VSE workload categories
- **KVM** 5 KVM workload categories
- **Linux** 5 Linux workload categories
- **zAware** 1 zAware workload (not shown in LSPR tables)
- **zACI** 1 zACI workload (not shown in LSPR tables)
- **SSC** 1 SSC workload (not shown in LSPR tables; intended to replace zAware and zACI)
- **CFCC** 1 CFCC workload (not shown in LSPR tables)
- **Not IPL'd** Special designation for partitions with no SCP running.

Multiple SCP versions are supported for both z/OS and z/M. The SCP version information is used to enforce configuration rules (i.e., the number of LCPs that can be defined, or whether associated LCPs are supported (zAAP or zIIP for z/OS or IFL for z/VM). Capacity results will not be affected by changing the version.

z/OS versions supported in a GP partition for each IBM Z family

z/OS Version	IBM Processor Family								
	z17	z16	z15	z14	z13	z12	z11	z10	z9
z/OS-3.2	S	S	S						
z/OS-3.1	S	S	S	S					
z/OS-2.5	S	S	S	S	S				
z/OS-2.4	S	S	S	S	S	S			
z/OS-2.3		S	S	S	S	S			
z/OS-2.2		S	S	S	S	S	S	S	
z/OS-2.1			S	S	S	S	S	S	
z/OS-1.13				S	S	S	S	S	S
z/OS-1.12					S	S	S	S	S
z/OS-1.9							S	S	S
z/OS-1.8							S	S	S
z/OS-1.7							S	S	S
z/OS-1.6									S
z/OS-1.5									S
z/OS-1.4									S

S = Supported

z/VM versions supported in a GP or IFL partition for each IBM Z family

z/VM Versions	IBM Processor Family								
	z17	z16	z15	z14	z13	z12	z11	z10	z9
z/VM-7.4	S	S	S						
z/VM-7.3	S	S	S	S					
z/VM-7.2		S	S	S	S				
z/VM-7.1		S	S	S	S	S			
z/VM-6.4			S	S	S	S	S		
z/VM-6.3 *			S	S	S	S	S	S	
z/VM-5.4						S	S	S	S

S = Supported

* Note: To represent z/VM-6.3 assign z/VM-6.4 For **zPCR** purposes, z/VM-6.4 support was extended to include the z10 family.

Any SCP/workload category shown on the **Workloads** window, including CFCC, can be assigned to a General Purpose partition. Only z/VM, Linux, and KVM can be assigned to an IFL partition. And only CFCC can be assigned to an ICF partition. zAAP and zIIP LCPs must be associated with a parent z/OS GP partition, and are assumed to be running the same z/OS version and workload category. IFL LCPs can be associated with a parent z/VM GP partition (hardware definition for the partition must be **Mode=z/VM**), and are assumed to be running the same z/VM workload category.

For z/VSE, a maximum of four LCPs is supported in **zPCR**.

The currently defined **Reference-CPU** (processor model, scaling-factor and scaling-metric), is used as the basis for all capacity projections. **Capacity projections for various partition configurations on the same LPAR host, or on different LPAR hosts are comparable as long as the Reference-CPU and its scaling-factor remain unchanged.**

Click the **Return** toolbar icon to close the **LPAR Host and Partition Configuration** window and return to **Control Panel** window.

Features supported by **zPCR** for each IBM Z family

Feature	IBM Processor Family								
	z17	z16	z15	z14	z13	z12	z11	z10	z9
zAAP CPs						S	S	s	S
zIIP CPs	S	S	S	S	S	S	S	S	S
Power Save Mode				S	S	S	S		
Absolute Capping	S	S	S	S	S	S			
Confirmation when Partition LCPs Exceed Drawer Size	S	S	S	S	S				
zAware Partition					S	S			
zACI Partition					S				
SSC Partition	S	S	S	S	S				
Simultaneous Multi Threading (SMT) (zIIP/IFL partitions)	S	S	S	S	S				
System Recovery Boost (SRB)	S	S	S						
Not IPL'd partition	S	S	S	S	S	S	S	S	S

S = Supported

Define Partitions Group Box

Once the LPAR host has been selected and configured, buttons are enabled for the purpose of defining partitions. Only the buttons that are valid for the LPAR host's defined CP types are enabled.

- Click the **GP** button to define partitions with workloads to be run on General Purpose CPs.
zAAP and/or zIIP partitions are defined (associated) in conjunction with a parent z/OS GP partition. When such specialty CPs are configured on the host, the **GP** button name will be expanded to include the specialty CPs (i.e., **GP / zAAP / zIIP**).

An IFL partition may be defined in conjunction (associated) with a GP partition running z/VM.

- Click the **IFL** button to define partitions with workloads to be run on IFLs.
- Click the **ICF** button to define partitions with the CFCC workload to be run on ICFs.

Clicking any of these buttons will open the **Partition Definition window**, where individual partition definitions are created (see [Partition Definition](#)). Upon return, the logical partition configuration is summarized on the right side of the **LPAR Host and Partition Configuration** window.

Reports Group Box

With an LPAR host processor specified, and a legitimate partition configuration defined, three reports are available via buttons in the **Capacity Reports** group box:

1. Click the **Host Summary** button to display the **Host Summary Report** window, which includes a description of the LPAR host, a summary of the partition configuration, and the projected capacity available (see [Host Summary Report](#)).
2. Click the **Partition Detail** button to display the **Partition Detail Report** window, which shows each partition as entered, with its capacity expectation (see [Partition Detail Capacity Report](#)).
3. Click the **Partition Utilized Capacity** button to display the **Utilized Capacity Report** window (see [Utilized Capacity Report](#)). This button is enabled only if the entire LPAR host and its entire partition configuration was created from EDF or RMF.

Note that there is no output capability for the **LPAR Host and Partition Configuration** window. A complete review of the LPAR host and its partition configuration, including capacity results, can be obtained as output from the **Host Summary Report** window and the **Partition Detail Report** window.

Creating LPAR Configuration from EDF

An **EDF** (Enterprise Data File) can be used to create the entire LPAR host and partition configuration or to copy individual partition definitions into a currently active LPAR configuration. A single EDF interval must be chosen, from which the requested information will be drawn. For detailed information on EDFs, see [EDF Overview](#).

From **zPCR**'s **Host and Partition Configuration** window, there are two ways that configuration information can be obtained from an EDF. The EDF file extension must be ".edf" or ".txt".

1. In the **Define LPAR Host Processor** group box, under **Create Host and Partitions From**, click the **EDF** button to create the entire LPAR configuration (LPAR host processor and some or all of its partitions). This button is enabled only when the LPAR host has yet to be defined (i.e., a new LPAR configuration). This capability allows you to quickly build the LPAR host and its entire partition configuration, representing the current operating environment.

Note: The **LPAR Configuration Name** shown on the **Control Panel** has 2 levels, The 1st level has default text which can be user modified from the **Control Panel** (see [Renaming LPAR Configurations](#)). This level is displayed on all report windows. The 2nd level is generated automatically when the LPAR host processor from the EDF input becomes defined. This name is based on the processor hardware model with the real GP CP model information. Appended is the number of other real CP types defined ("A=" for zAAPs; "I=" for zIIPs; "F=" for IFLs; "C=" for ICFs). This level is displayed only on the **Control Panel**. When a duplicate name would result, a repetition index is appended. Whenever the LPAR host is changed, a new name is generated.

2. In the **Define Partitions** group box, under **Copy Partitions From**, click the **EDF** button to copy some or all partition definitions into a current LPAR configuration. This button is enabled only when the LPAR host processor has already been defined. This capability allows the addition of existing partitions to an LPAR configuration.

As an alternative to clicking the buttons, an EDF file can be dragged and dropped onto either of these **EDF** buttons when it is active.

Use of either button will display the **EDF Interval Selection** window (see [EDF Interval Selection](#)) from which an EDF report can be loaded. Once an EDF interval is selected, click the **Show Partitions** button to display the **Create LPAR Configuration from EDF** or **Get Partitions from EDF** window (see [Get Partitions from EDF](#)).

System Recovery Boost Considerations

If any partition in the selected EDF interval has **System Recovery Boost** activity, the boosted partitions will be flagged in the **Create LPAR Configuration from EDF** window. Such flags will be perpetuated to ensuing **zPCR** windows.

For detail concerning SRB support see [System Recovery Boost](#).

Creating LPAR Configuration from RMF

A z/OS **RMF Report** can be used to create the entire LPAR host and partition configuration or to copy individual partition definitions into a currently active LPAR configuration. A single RMF interval must be chosen, from which the requested information will be drawn. For detailed information on RMFs, see [RMF Overview](#).

From **zPCR**'s **Host and Partition Configuration** window, there are two ways that configuration information can be obtained from RMF. The RMF file extension must be ".rmf" or ".txt".

1. In the **Define LPAR Host Processor** group box, under **Create Host and Partitions From**, click the **RMF** button to create the entire LPAR configuration (LPAR host processor and some or all of its partitions). This button is enabled only when the LPAR host has yet to be defined (i.e., a new LPAR configuration). This capability allows you to quickly build the entire LPAR host and partition configuration model that represents a current operating environment.

Note: The **LPAR Configuration Name** shown on the **Control Panel** has 2 levels,

The 1st level has default text which can be user modified from the **Control Panel** (see [Renaming LPAR Configurations](#)). This level is displayed on all report windows.

The 2nd level is generated automatically when the LPAR host processor from the RMF input becomes defined. This name is based on the processor hardware model with the real GP CP model information. Appended is the number of other real CP types defined ("A=" for zAAPs; "I=" for zIIPs; "F=" for IFLs; "C=" for ICFs). This level is displayed only on the **Control Panel**.

When a duplicate name would result, a repetition index is appended. Whenever the LPAR host is changed, a new name is generated.

2. In the **Define Partitions** group box, under **Copy Partitions From**, click the **RMF** button to copy some or all partition definitions into a current LPAR configuration. This button is enabled only when the LPAR host processor has already been defined. This capability allows the addition of existing partitions to a current study.

As an alternative to clicking the buttons, an RMF file can be dragged and dropped onto either of the **RMF** buttons when the button is active.

Use of either button will display the **RMF Interval Selection** window (see [RMF Interval Selection](#)) from which an RMF report can be loaded. Once an RMF interval is selected, click the **Show Partitions** button to display the **Create LPAR Configuration from RMF** or **Get Partitions from RMF** window (see [Get Partitions from RMF](#)).

System Recovery Boost Considerations

If any partition in the selected RMF interval has **System Recovery Boost** activity, the boosted partitions will be flagged in the **Create LPAR Configuration from RMF** window. Such flags will be perpetuated to ensuing **zPCR** windows.

For detail concerning SRB support see [System Recovery Boost](#).

Creating LPAR Configuration from a zPCR Study

From **zPCR**'s **Host and Partition Configuration** window, there are two ways that configuration information can be obtained from a previous **zPCR** study. The study file extension must be **.zpcr**.

1. In the **Define LPAR Host Processor** group box, under **Create Host and Partitions From**, click the **zPCR Study** button to create the entire LPAR configuration (LPAR host processor and all of its partitions). This button is enabled only when the LPAR host has yet to be defined (i.e., a new LPAR configuration). This capability allows you to quickly build the LPAR host and its entire partition configuration, using information from a previous study.

Note: The **LPAR Configuration Name** shown on the **Control Panel** has 2 levels, The 1st level has default text which can be user modified from the **Control Panel** (see [Renaming LPAR Configurations](#)). This level is displayed on all report windows. The 2nd level is generated automatically when the LPAR host processor from the **zPCR** study input becomes defined. This name is based on the processor hardware model with the real GP CP model information. Appended is the number of other real CP types defined ("A=" for zAAPs; "I=" for zIIPs; "F=" for IFLs; "C=" for ICFs). This level is displayed only on the **Control Panel**. When a duplicate name would result, a repetition index is appended. Whenever the LPAR host is changed, a new name is generated.

2. In the **Define Partitions** group box, under **Copy Partitions From**, click the **zPCR Study** button to copy some or all partition definitions into a current LPAR configuration. This button is enabled only when the LPAR host processor has already been defined. This capability allows the addition of existing partitions to an LPAR configuration.

As an alternative to clicking the buttons, a **zPCR** study file can be dragged and dropped onto either of the **zPCR Study** buttons when it is active. For detail concerning these transfers, see [Create Host and Partitions from Study](#) or [Copy Partitions from Study](#).

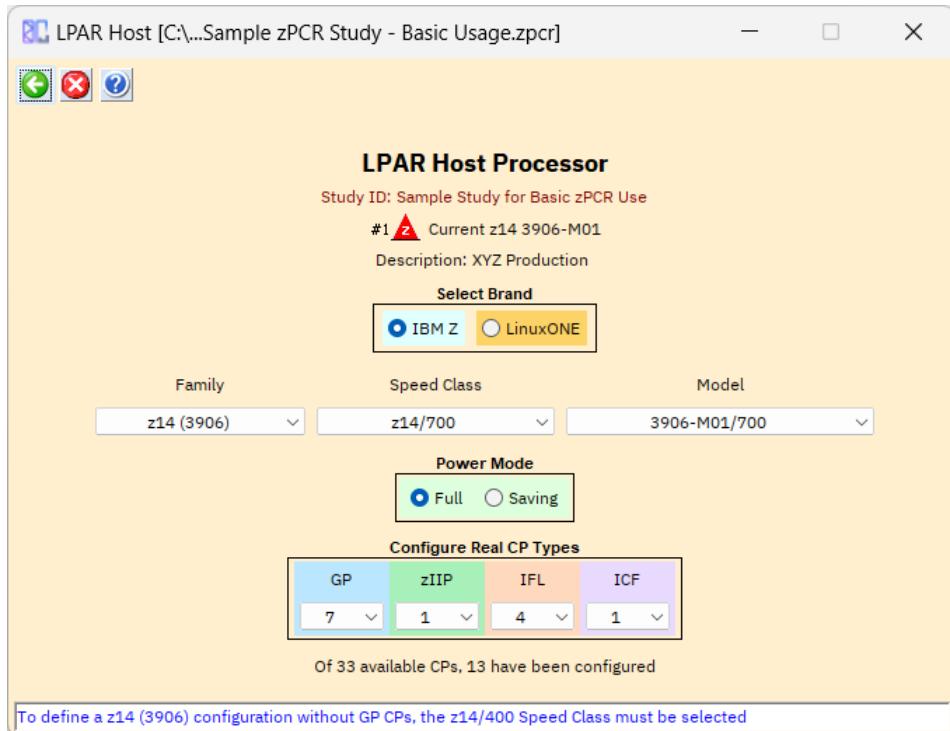
System Recovery Boost Considerations

If any partition in the selected study has **System Recovery Boost** activity, the boosted partitions will be flagged in the **Copy Partitions from Previous Study** window. Such flags will be perpetuated to ensuing zPCR windows.

For detail concerning SRB support see [System Recovery Boost](#).

LPAR Host Processor

LPAR Configuration Capacity Planning



The **LPAR Host** window is accessed from the **LPAR Host and Partition Configuration** window by clicking the **Specify Host** button. From this window, any IBM mainframe processor model may be defined as the LPAR host for the configuration.

Note: The **LPAR Configuration Name** shown on the **Control Panel** has 2 levels, The 1st level has default text which can be user modified from the **Control Panel** (see [Renaming LPAR Configurations](#)). This level is displayed on all report windows. The 2nd level is generated automatically when the LPAR host processor from the **LPAR Host** window becomes defined. This name is based on the processor hardware model with the real GP CP model information. Appended is the number of other real CP types defined ("A=" for zAAPs; "I=" for zIIPs; "F=" for IFLs; "C=" for ICFs). This level is displayed only on the **Control Panel**. When a duplicate name would result, a repetition index is appended. Whenever the LPAR host is changed, a new name is generated.

To define the host make selections as follows:

1. Choose the brand as **IBM Z** or **LinuxONE**.
2. Choose a processor family using the **Family** dropdown list.
3. Choose the processor speed for GP CPs using the **Speed Class** dropdown list.
4. Choose the processor model using the **Model** dropdown list.

Set the number of CPs for each of the engine types that are to be configured using the dropdown lists provided. Every processor family and model has a limit on the total number of CPs that can be configured. As the various types of CPs are added to the host configuration, the count of potential CPs remaining is reduced. Text below the **Configure Real CP Types** group box indicates the configured CP status.

For EC processor families with a **Speed Class** setting less than the highest speed, the number of General Purpose CPs that can be configured is limited, and may be less than the total CPs available. For these cases, the GP count is limited to that maximum.

For BC processor families, the number of General Purpose CPs that can be configured is limited, and may be less than the total CPs available. For these cases, the GP count is limited to that maximum.

Certain specialty CP types are not supported on specific processor families. In these cases, that CP type dropdown will not appear.

Notes:

1. IBM z17, z16, z15, z14, and z13 processors do not support zAAP CPs. The zAAP CP type will be omitted from the **Configure Real CP Types** group box. If converting a LPAR host with zAAP CPs configured, the associated zAAP partitions will need to be converted to zIIP associated partitions manually (assumes that **zAAP on zIIP** has been enabled).
2. **LinuxONE** processors can only be configured with IFL CPs, and optionally, 1 GP CP. The zAAP, zIIP, and ICF CP types will be omitted from the **Configure Real CP Types** group box.
3. The **LinuxONE Rockhopper 2828** models are actually known as **2828-H06** and **2828-H13**. However, in **zPCR** they are identified as **2828-L06** and **2828-L13**. This change was necessary in order to differentiate the **LinuxONE** models from the equally named **IBM z12** models.

The specific partition configuration that can be defined will be dependent on the processor family, model, and CP configuration selected, based on the following items:

- Maximum number of partitions that may be configured
- Number of General Purpose RCPs configured
- Number of zAAP RCPs configured (there are rules about how many CPs relative to General Purpose CPs can be defined while remaining a standard zAAP configuration).
- Number of zIIP RCPs configured (for z15 and prior families there are rules about how many zIIP CPs relative to General Purpose CPs can be defined while remaining a standard zIIP configuration).
- Number of IFL RCPs configured
- Number of ICF RCPs configured

Switching the LPAR host from LinuxONE to IBM Z

The LPAR host can be changed from **LinuxONE** to **IBM Z** since all **LinuxONE** CP types are supported. All defined RCPs and defined partitions will be transferred.

The LPAR host cannot be changed from **IBM Z** to **LinuxONE**. This is because most **IBM Z** CP types are not supported on **LinuxONE**. In addition, there are limitations concerning the number of GP CPs allowed on **LinuxONE**.

SCPs that can be defined to Partitions

Note: The IBM z17, z16, z15, and z14 can only be IPL'd in z/Architecture mode. Therefore, older versions of SCPs cannot be run. **zPCR** supports specific SCP versions for both z/OS and z/VM. When migrating to a IBM z16, z15, and z14 host, old versions of these SCPs will be converted to a default version that will IPL. For all other SCPs, **zPCR** assumes that a supported version will be used.

z/OS can be defined to GP partitions on any **IBM Z** processor model. **LinuxONE** allows a single GP CP, which must be assigned with z/OS.

- On IBM z17 processors, z/OS-2.4 or later must be specified.
- On IBM z16 processors, z/OS-2.2 or later must be specified.
- On z15 processors, z/OS-2.1 or later must be specified.
- On z14 processors, z/OS-1.13 or later must be specified.
- On z13 and z12 processors, z/OS-1.10 or later must be specified.
- On z11, z10, and z9 processors, z/OS must be between z/OS-2.2 and z/OS-1.7.

If changing the LPAR host to a newer family, any partitions defined with older versions of z/OS may need to be converted to a newer version in order for partition definitions to be valid. Note that, changing the z/OS version will have no effect on capacity results, since all capacity data is drawn from the same z/OS LSPR table.

When supported by the specified z/OS version, zAAP LCPs and zIIP LCPs can also be associated with a parent z/OS GP partition.

If changing an older LPAR host to an IBM z17, z16, z15, z14, or z13, zAAP CPs are no longer supported. GP partitions with associated zAAP LCPs must be converted to use zIIP LCPs. The zAAP workload must run on zIIP LCPs using “zAAP on zIIP” enabled.

z/VM can be defined to **IBM Z** GP or IFL partitions, or to **LinuxONE** IFL partitions. The number of LCPs supported is limited based on the z/VM version.

- On IBM z17, z16, z15, and z14, with z/VM-7.4, -7.3, -7.2, or -7.1, the LCP limit is 80 when SMT is not enabled and 40 when SMT is enabled.
- On z13, with z/VM-7.2, -7.1 or -6.4, the limit is 64 when SMT is not enabled and 32 when SMT is enabled.
- When z/VM is run in an IFL partition, Linux guests must be assumed.

z/VSE can only be assigned to **IBM Z** GP partitions. A maximum of four LCPs is supported by **zPCR**.

KVM can be defined to GP or IFL partitions on **IBM Z** or **LinuxONE** IBM z17, z16, z15, z14, z13, z12 processor models. KVM can support any number of LCPs.

Linux can be defined to GP or IFL partitions on any **IBM Z** or **LinuxONE** processor model. Linux can support any number of LCPs.

zAware can be defined to GP or IFL partitions on z13 and z12 processors only. For zAware, a limit of 16 LCPs is currently enforced by **zPCR**. SSC is intended to replace zAware.

zACI can be defined to General Purpose or IFL partitions on **IBM** z12 and z13 processors only. For a single zACI image, any number of LCPs can be defined. SSC is intended to replace zACI.

SSC (Secure Service Container) can be defined to General Purpose or IFL partitions on **IBM** z13 and later processors. SSC is intended to replace zAware and zACI.

CFCC can be defined to General Purpose or ICF partitions on any **IBM Z** processor model. A CFCC image is limited to supporting 16 LCPs on all processor models.

Not IPL'd should be assigned to any partition that is HMC activated and have no SCP running.

Note: If the LPAR host is changed to a processor model that does not support one of the z/OS or z/VM versions currently defined to a partition, a dialog will offer the chance to convert to a supported version or to cancel the change.

Window Controls

Click the **Return** toolbar icon to accept your host specification. At least 1 General Purpose, IFL, or ICF must be configured. If a nonstandard zAAP or zIIP CP configuration is detected, a dialog will be posted providing a checkbox to justify the reason for being nonstandard. If not justified, zAAP and zIIP RCPs will be restored to be the maximum allowed for a standard configuration.

Click the **Cancel** toolbar icon to discard your changes.

In either case you will return to the **LPAR Host and Partition Configuration** window, where the LPAR host specification is summarized.

Power-Save Mode

Applies to **IBM z14, z13, z12 and z11** and **LinuxONE** Emperor I and Emperor II only. It is not available on IBM z17, z16, z15, LinuxONE III, or LinuxONE 4.

While normally run at full power, these families can also be run in **Power-Save** mode. **Power-Save** mode reduces the processor capacity (all CP types are affected equally) and the MSU rating, while also reducing power requirements. All processor models default to full-power. When defined as a model in one of these families, the **LPAR Host** window provides a means to set it so that capacity results will represent its operation in **Power-Save** mode.

To change to **Power-Save** mode, click the **Saving** radio button in the **Power Mode** group box. If the LPAR host is configured for **Power-Save** mode, all capacity reporting windows will indicate so in the title area.

Absolute Capping

Applies to **IBM Z** (z12 and later) and all **LinuxONE** processors

An **Absolute Capping** value may be specified for a partition. The value is expressed as a fractional number of the partition's LCPs. The value has the effect of reducing the partition's **Maximum Capacity** value to that of the number of LCPs specified.

For cases where the LPAR host supports **Absolute Capping**, metrics can be entered on the **Partition Detail Report** window.

Nonstandard zAAP or zIIP Configurations

Standard configurations for zAAPs and for zIIPs limit the number of CPs of either type in relation to the number of General Purpose CPs configured.

- IBM z17 and z16: any number of zIIPs can be defined.
- z15, z14, z13, and z12 models, zIIPs exceeding 2 for each GP are considered nonstandard.
- For all prior models, zAAPs or zIIPs RCPs cannot exceed 1 for each GP RCP.

With certain processor upgrade scenarios, the number of zAAP and zIIP CPs allowed is based on the number previously installed, and therefore the new LPAR host could become a nonstandard zAAP/zIIP configuration.

When defining the LPAR host, if a nonstandard zAAP or zIIP configuration is detected, a **Notice** dialog is offered with a checkbox to suggest the reason for being nonstandard. If checked, the nonstandard configuration will be allowed. If not checked, the zAAP and zIIP CPs will revert to those of a standard configuration.

Whenever a nonstandard zAAP/zIIP configuration has been defined, all subsequent windows will indicate such. Should the nonstandard configuration be changed to become standard, the nonstandard status is removed.

zAAP on zIIP Capability

On all **IBM** z9 and later processors, GP partitions can be configured to run zAAP eligible work on zIIP engines (**zAAP on zIIP**). On z13 and later processor families, zAAP eligible work must be run on zIIP engines.

When **zAAP on zIIP** is enabled for a GP partition, the following rules apply:

- zAAP LCPs cannot be used by the GP partition, and, therefore, should not be associated with the GP partition.
- z/OS-1.11 is enhanced to allow use of this capability. z/OS-1.10 and z/OS-1.9 require a PTF to make use of this capability. The SCP for GP partitions with zAAP eligible work should be specified as z/OS-1.9 or later.

In **zPCR**, if **zAAP on zIIP** is to be assumed for a z/OS GP partition, then only zIIP LCPs should be associated with it.

LPAR Host Hardware Designations

zPCR enforces the configuration rules applicable for each processor family

For older and for lesser capacity families, the processor model designation indicates the specific number of General Purpose CPs configured. The number of specialty engines (zAAPs, zIIPs, IFLs, and ICFs) that can be configured is unique to each model, depending on whether supported by the family, and depending on the number of CPs available to be configured as such. zAAP CPs are only supported on **IBM Z** families prior to the z13. zIIP, IFL, and ICF CPs are supported on all **IBM Z** families.

Note: z900, z800, z990 and z890 processors are no longer supported by **zPCR** as an LPAR host. However, they may be viewed in the LSPR Multi-Image ITRR table.

The IBM z17, z16, z15, z14 and z13 processor families (drawer-based) and z12, z11, z10, and z9 processor families (book-based), have a designation method where the model number indicates the number of books/drawers and the maximum number of CPs that can be configured. The actual number of GP CPs (limitations apply for some models) and each of the other CP types are configured independently.

The z10 through z13 processor families have a designation method where the model number indicates the number of books or drawers assumed and the maximum number of CPs that can be configured. The number of General Purpose CPs is limited to a subset of the total, and are based on the specific z/OS model designated. Each of the other CP types are configured independently.

In addition to the **Family**, **Speed Class**, and **Model** metrics, newer processor models have an additional **MaxN** designation specifying the maximum number of RCPs that can be defined. The **MaxN** setting also determines the number of drawers installed.

- **z17-ME1 (9175):** **Max43, Max90, Max136, Max183, and Max208.**
- **z16-A01 and -LA1 (3931):** **Max39, Max82, Max125, Max168, and Max200.**
- **z16-A02/AGZ and -LA2/AGL (3932):** **Max5, Max16, Max32, and Max68.**
- **z15-T01 and -LT1 (8561):** **Max34, Max71, Max108, Max145, and Max190.**
- **z15-T02 and -LT2 (8562):** **Max4, Max13, Max21, Max31, and Max65.**
- **z14-ZR1 and -LR1 (3907):** **Max4, Max12, Max24, and Max30.**

When reading EDF or RMF, if the **MaxN** setting cannot be determined, the largest **MaxN** will be assumed.

IBM z17 Processor Models

IBM z17 (9175)

When configuring, there is one model with 5 CP sizes

9175-ME1(Max43)	1 drawer; up to 43 CPs may be configured
9175-ME1(Max90)	2 drawers; up to 90 CPs may be configured
9175-ME1(Max136)	3 drawers; up to 136 CPs may be configured
9175-ME1(Max183)	4 drawers; up to 183 CPs may be configured
9175-ME1(Max208)	4 drawers; up to 208 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 43 General Purpose CPs can be configured for the /400, /500, and /600, while up to 208 can be configured for the /700. All General Purpose CPs must be the same speed class. The z17/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 43 for the /400, /500, and /600), as a zIIP, as an IFL, or as an ICF. Note that zAAP CPs are not available on the IBM z17. Former zAAP workload must be run on zIIP CPs using “zAAP on zIIP” capability.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z16 Processor Models

IBM z16 (3931)

When configuring, there is one model with 5 CP sizes

3931-A01(Max39)	1 drawer; up to 39 CPs may be configured
3931-A01(Max82)	2 drawers; up to 82 CPs may be configured
3931-A01(Max125)	3 drawers; up to 125 CPs may be configured
3931-A01(Max168)	4 drawers; up to 168 CPs may be configured
3931-A01(Max200)	4 drawers; up to 200 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 39 General Purpose CPs can be configured for the /400, /500, and /600, while up to 200 can be configured for the /700. All General Purpose CPs must be the same speed class. The z16/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 39 for the /400, /500, and /600), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF. Note that zAAP CPs are not available on the IBM z16. Former zAAP workload must be run on zIIP CPs using "zAAP on zIIP" capability.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z16 (3932)

When configuring, there are two models, each with 4 CP sizes

3932-A02(Max5)	1 drawer; up to 5 CPs may be configured
3932-AGZ(Max5)	
3932-A02(Max16)	1 drawer; up to 16 CPs may be configured
3932-AGZ(Max16)	
3932-A02(Max32)	1 drawer; up to 32 CPs may be configured
3932-AGZ(Max32)	
3932-A02(Max68)	2 drawers; up to 68 CPs may be configured
3932-AGZ(Max68)	

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 6 General Purpose CPs can be designated (for example 3932-Z01, Z02, Z03, Z04, Z05, or Z06). All General Purpose CPs must be the same speed class. The z16/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z15 Processor Models

IBM z15 (8561)

When configuring, there is one model with 5 CP sizes

8561-T01(Max34)	1 drawer; up to 34 CPs may be configured
8561-T01(Max71)	2 drawers; up to 71 CPs may be configured
8561-T01(Max108)	3 drawers; up to 108 CPs may be configured
8561-T01(Max145)	4 drawers; up to 145 CPs may be configured
8561-T01(Max190)	5 drawers; up to 190 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 34 General Purpose CPs can be configured for the /400, /500, and /600, while up to 190 can be configured for the /700. All General Purpose CPs must be the same speed class. The z15/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 34 for the /400, /500, and /600), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF. Note that zAAP CPs are not available on the z15. Former zAAP workload must be run on zIIP CPs using “zAAP on zIIP” capability.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z15 (8562)

When configuring, there is one model with 5 CP sizes

8562-T02(Max4)	1 drawer; up to 4 CPs may be configured
8562-T02(Max13)	1 drawer; up to 13 CPs may be configured
8562-T02(Max21)	1 drawer; up to 21 CPs may be configured
8562-T02(Max31)	1 drawer; up to 31 CPs may be configured
8562-T02(Max65)	2 drawers; up to 65 CPs may be configured

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 6 General Purpose CPs can be designated (for example 8562-Z01, Z02, Z03, Z04, Z05, or Z06). All General Purpose CPs must be the same speed class. The z15/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z14 Processor Models

IBM z14 (3906)

When configuring, there is a choice of 5 models

3906-M01	1 drawer; up to 33 CPs may be configured
3906-M02	2 drawers; up to 69 CPs may be configured
3906-M03	3 drawers; up to 105 CPs may be configured
3906-M04	4 drawers; up to 141 CPs may be configured
3906-M05	4 drawers; up to 170 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 33 General Purpose CPs can be configured for the /400, /500, and /600, while up to 170 can be configured for the /700. All General Purpose CPs must be the same speed class. The z14/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 33 for the /400, /500, and /600), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF. Note that zAAP CPs are not available on the z14. Former zAAP workload must be run on zIIP CPs using “zAAP on zIIP” capability.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z14 (3907-ZR1)

When configuring, there is one model with 4 CP sizes

3907-ZR1(Max4)	1 drawer; up to 4 CPs may be configured
3907-ZR1(Max12)	1 drawer; up to 12 CPs may be configured
3907-ZR1(Max24)	1 drawer; up to 24 CPs may be configured
3907-ZR1(Max30)	1 drawer; up to 30 CPs may be configured

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 6 General Purpose CPs can be designated (for example 3907-Z01, Z02, Z03, Z04, Z05, or Z06). All General Purpose CPs must be the same speed class. The z14/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z13 Processor Models

IBM z13 (2964)

When configuring, there is a choice of 5 models

2964-N30	1 drawer; up to 30 CPs may be configured
2964-N63	2 drawers; up to 63 CPs may be configured
2964-N96	3 drawers; up to 96 CPs may be configured
2964-NC9	4 drawers; up to 129 CPs may be configured
2964-NE1	4 drawers; up to 141 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 30 General Purpose CPs can be configured for the /400, /500, and /600, while up to 141 can be configured for the /700. All General Purpose CPs must be the same speed class. The z13/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 30 for the /400, /500, and /600), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF. Note that zAAP CPs are not available on the z13. Former zAAP workload must be run on zIIP CPs using “zAAP on zIIP” capability.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z13s (2965)

When configuring, there is a choice of 2 models

2965-N10	1 drawer; up to 10 CPs may be configured
2965-N20	2 drawers; up to 20 CPs may be configured

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 6 General Purpose CPs can be designated (for example 2965-Z01, Z02, Z03, Z04, Z05, or Z06). All General Purpose CPs must be the same speed class. The z13s/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z12 Processor Models

IBM zEnterprise EC12 (zEC12)

When configuring, there is a choice of 5 models

2827-H20	1 book; up to 20 CPs may be configured
2827-H43	2 books; up to 43 CPs may be configured
2827-H66	3 books; up to 66 CPs may be configured
2827-H89	4 books; up to 89 CPs may be configured
2827-HA1	4 books; up to 101 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 20 General Purpose CPs can be configured for the /400, /500, and /600, while up to 101 can be configured for the /700. All General Purpose CPs must be the same speed class. The zEC12/400 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 20 for the /400, /500, and /600), as a zAAP (standard zAAP configurations are limited to 2X the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM zEnterprise BC12 (zBC12)

When configuring, there is a choice of 2 models

2828-H06	1 drawer; up to 6 CPs may be configured
2828-H13	2 drawers; up to 13 CPs may be configured

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 6 General Purpose CPs can be designated (for example 2828-Z01, Z02, Z03, Z04, Z05, or Z06). All General Purpose CPs must be the same speed class. The z12BC/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zAAP (standard zAAP configurations are limited to 2X the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to 2X the number of General Purpose CPs), as an IFL, or as an ICF.

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z11 Processor Models

IBM zEnterprise 196 (z196)

When configuring, there is a choice of 5 models

2817-M15	1 book; up to 15 CPs may be configured
-----------------	--

2817-M32	2 books; up to 32 CPs may be configured
-----------------	---

2817-M49	3 books; up to 49 CPs may be configured
-----------------	---

2817-M66	4 books, up to 66 CPs may be configured
-----------------	---

2817-M80	4 books; up to 80 CPs may be configured
-----------------	---

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 15 General Purpose CPs can be configured for the /400, /500, and /600, while up to 80 can be configured for the /700. All General Purpose CPs must be the same speed class. The z196/700 speed class is the only z196 host that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 15 for the /400, /500, and /600), as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF (limited to 16).

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM zEnterprise 114 (z114)

When configuring, there is a choice of 2 models

2818-M05	1 drawer; up to 5 CPs may be configured
-----------------	---

2818-M10	2 drawers; up to 10 CPs may be configured
-----------------	---

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 5 General Purpose CPs can be designated (for example 2818-Z01, Z02, Z03, Z04, or Z05). All General Purpose CPs must be the same speed class. The z114/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF.

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z10 Processor Models

IBM Z System z10 EC

When configuring, there is a choice of 5 models

2097-E12	1 book; up to 12 CPs may be configured
2097-E26	2 books; up to 26 CPs may be configured
2097-E40	3 books; up to 40 CPs may be configured
2097-E56	4 books; up to 56 CPs may be configured
2097-E64	4 books; up to 64 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 12 General Purpose CPs can be configured for the /400, /500, and /600, while up to 64 can be configured for the /700. All General Purpose CPs configured must be the same speed class. The z10 EC/700 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 12 for the /400, /500, and /600), as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zAAP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF (limited to 16).

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM Z System z10 BC

When configuring, there is a single model

2098-E10	1 drawer; up to 10 CPs may be configured
-----------------	--

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). A maximum of 5 General Purpose CPs can be designated (for example 2098-Z01, Z02, Z03, Z04, or Z05). All General Purpose CPs must be the same speed class. The z10 BC/A00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF.

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM z9 Processor Models

IBM Z System z9 EC

When configuring, there is a choice of 5 models

2094-S08	1 book; up to 8 CPs may be configured
2094-S18	2 books; up to 18 CPs may be configured
2094-S28	3 books; up to 28 CPs may be configured
2094-S38	4 books; up to 38 CPs may be configured
2094-S54	4 books; up to 54 CPs may be configured

There are 4 General Purpose CP **Speed Classes**, the /400, /500, /600, and /700 (the /700 is the full speed model). No more than 8 General Purpose CPs can be configured for the /400, /500, and /600, while up to 54 can be configured for the /700. All General Purpose CPs configured must be the same speed class. The z9 EC/700 is the only speed class that can be configured with no General Purpose CPs.

Any available CP can be configured as a General Purpose CP (limited to 8 for the /400, /500, and /600), as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF (limited to 16).

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM Z System z9 BC

When configuring, there is a choice of 2 models

2096-R07	1 drawer; up to 6 CPs may be configured
2096-S07	1 drawer; up to 7 CPs may be configured

There are 26 General Purpose CP **Speed Classes**, /A00, /B00, /C00, /D00, .../Z00 (/Z00 is the full speed model). Speeds /A00 through /J00 can be configured on the 2096-R07. Speeds /K00 through /Z00 can be configured on the 2096-S07.

Depending on the speed, a maximum of 4 General Purpose CPs can be designated (for example 2096-Z01, Z02, Z03, or Z04). All General Purpose CPs must be the same speed class. The z9 BC/Z00 is the only speed class that can be configured with no General Purpose CPs.

After the General Purpose CPs are designated, any available CP can be configured as a zAAP (standard zAAP configurations are limited to the number of General Purpose CPs), as a zIIP (standard zIIP configurations are limited to the number of General Purpose CPs), as an IFL, or as an ICF.

zAAPs, zIIPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

IBM zSeries Processor Models

Note: Processors on this page may no longer be defined in **zPCR** as an LPAR host. They do remain visible in the **LSPR ITRR Table** windows when “**All Families**” is selected.

IBM zSeries 990

When configuring, there is a choice of 4 models

2084-A08	1 book; up to 8 CPs may be configured
2084-B16	2 books; up to 16 CPs may be configured
2084-C24	3 books; up to 24 CPs may be configured
2084-D32	4 books, up to 32 CPs may be configured

Any available CP can be configured as a General Purpose CP, a zAAP (limited to the number of General Purpose CPs), an IFL, or an ICF (limited to 16). There is only 1 **Speed Class** for the z990.

All z990 CPs are the same speed regardless of the type assigned.

The z990 operates with 2 CP pools:

- 1st pool is for General Purpose CPs
- 2nd pool is for zAAP/IFL/ICF CPs

IBM zSeries 890

When configuring, there is a choice of 4 models

2086-070	No General Purpose CPs; IFL and ICF only
2086-1x0	1 General Purpose CP
2086-2x0	2 General Purpose CPs
2086-3x0	3 General Purpose CPs
2086-4x0	4 General Purpose CPs

Up to 4 CPs can be configured on any model above. The number of General Purpose CPs is based on the model selection above. The remaining CPs can be assigned as zAAP (limited to the number of General Purpose CPs), as IFL, or as ICF. There are 7 General Purpose **Speed Classes**, designated as 1 (slowest) through 7 (full speed). For the General Purpose models above, substitute the “**x**” with the desired speed.

zAAPs, IFLs, and ICFs are always full speed CPs regardless of the General Purpose CP speed class installed.

The z890 operates with 2 CP pools:

- 1st pool is for General Purpose CPs
- 2nd pool is for zAAP/IFL/ICF CPs

LinuxONE z17 Processor Models

LinuxONE Emperor 5 (9175)

When configuring, there is one model with 5 CP sizes

9175-ML1(Max43)	1 drawer; up to 43 CPs may be configured
9175-ML1(Max90)	2 drawers; up to 90 CPs may be configured
9175-ML1(Max136)	3 drawers; up to 136 CPs may be configured
9175-ML1(Max183)	4 drawers; up to 183 CPs may be configured
9175-ML1(Max208)	4 drawers; up to 208 CPs may be configured

The GP **Speed Class** must be designated as /400.

A maximum of one General Purpose CP may be configured (seen by z/OS as 9175-401). Remaining CPs may be only be configured as full speed IFLs.

LinuxONE z16 Processor Models

LinuxONE Emperor 4 (3931)

When configuring, there is one model with 5 CP sizes

3931-LA1(Max39)	1 drawer; up to 39 CPs may be configured
3931-LA1(Max82)	2 drawers; up to 82 CPs may be configured
3931-LA1(Max125)	3 drawers; up to 125 CPs may be configured
3931-LA1(Max168)	4 drawers; up to 168 CPs may be configured
3931-LA1(Max200)	4 drawers; up to 200 CPs may be configured

The GP **Speed Class** must be designated as /400.

A maximum of one General Purpose CP may be configured (seen by z/OS as 3931-401). Remaining CPs may be only be configured as full speed IFLs.

LinuxONE Rockhopper 4 (3932)

When configuring, there are two models, each with 4 CP sizes

3932-LA2(Max5)	1 drawer; up to 5 CPs may be configured
3932-AGL(Max5)	
3932-LA2(Max16)	1 drawer; up to 16 CPs may be configured
3932-AGL(Max16)	
3932-LA2(Max33)	1 drawer; up to 32 CPs may be configured
3932-AGL(Max33)	
3932-LA2(Max68)	2 drawers; up to 68 CPs may be configured
3932-AGL(Max68)	

The GP **Speed Class** may be designated as /A00 (no GP CPs) or /C00 (1 GP CP seen by z/OS as 3932-C01).

A maximum of one General Purpose CP may be configured. Remaining CPs may be only be configured as full speed IFLs.

LinuxONE z15 Processor Models

LinuxONE III (8561)

When configuring, there is one model with 5 CP sizes

8561-LT1(Max34)	1 drawer; up to 34 CPs may be configured
8561-LT1(Max71)	2 drawers; up to 71 CPs may be configured
8561-LT1(Max108)	3 drawers; up to 108 CPs may be configured
8561-LT1(Max145)	4 drawers; up to 145 CPs may be configured
8561-LT1(Max190)	5 drawers; up to 190 CPs may be configured

The GP **Speed Class** must be designated as /400.

A maximum of one General Purpose CP may be configured (seen by z/OS as 8561-401). Remaining CPs may be only be configured as full speed IFLs.

LinuxONE III (8562)

When configuring, there is one model with 5 CP sizes

8562-LT2(Max4)	1 drawer; up to 4 CPs may be configured
8562-LT2(Max13)	1 drawer; up to 13 CPs may be configured
8562-LT2(Max21)	1 drawer; up to 21 CPs may be configured
8562-LT2(Max31)	1 drawer; up to 31 CPs may be configured
8562-LT2(Max65)	2 drawers; up to 65 CPs may be configured

The GP **Speed Class** may be designated as /A00 (no GP CPs) or /C00 (1 GP CP seen by z/OS as 8562-C01).

A maximum of one General Purpose CP may be configured. Remaining CPs may be only be configured as full speed IFLs.

LinuxONE z14 Processor Models

LinuxONE Emperor II (3906)

When configuring, there is a choice of 5 models

3906-LM1	1 drawer; up to 33 CPs may be configured
3906-LM2	2 drawers; up to 69 CPs may be configured
3906-LM3	3 drawers; up to 105 CPs may be configured
3906-LM4	4 drawers; up to 141 CPs may be configured
3906-LM5	4 drawers; up to 170 CPs may be configured

The GP **Speed Class** must be designated as /400.

A maximum of one General Purpose CP may be configured (seen by z/OS as 3906-401). Remaining CPs may be only be configured as full speed IFLs.

LinuxONE Rockhopper II (3907)

When configuring, there is one model with 4 CP sizes

3907-LR1(Max4)	1 drawer; up to 4 CPs may be configured
3907-LR1(Max12)	1 drawer; up to 12 CPs may be configured
3907-LR1(Max24)	1 drawer; up to 24 CPs may be configured
3907-LR1(Max30)	1 drawer; up to 30 CPs may be configured

The GP **Speed Class** may be designated as /A00 (no GP CPs) or /C00 (1 GP CP seen by z/OS as 3907-C01).

A maximum of one General Purpose CP may be configured. Remaining CPs may be only be configured as full speed IFLs.

LinuxONE z13 Processor Models

LinuxONE Emperor (2964)

When configuring, there is a choice of 5 models

2964-L30	1 drawer; up to 30 CPs may be configured
2964-L63	2 drawers; up to 63 CPs may be configured
2964-L96	3 drawers; up to 96 CPs may be configured
2964-LC9	4 drawers; up to 129 CPs may be configured
2964-LE1	4 drawers; up to 141 CPs may be configured

The GP **Speed Class** must be designated as /400.
A maximum of one General Purpose CP may be configured (seen by z/OS as 2964-401). Remaining CPs may be only be configured as full speed IFLs.

LinuxONE Rockhopper (2965)

When configuring, there is a choice of 2 models

2965-L10	1 drawer; up to 10 CPs may be configured
2965-L20	2 drawers; up to 20 CPs may be configured

The GP **Speed Class** may be designated as /A00 (no GP CPs) or /C00 (1 GP CP seen by z/OS as 2965-C01).
A maximum of one General Purpose CP may be configured. Remaining CPs may be only be configured as full speed IFLs.

LinuxONE z12 Processor Models

LinuxONE Rockhopper (2828)

When configuring, there is a choice of 2 models

2828-H06	1 drawer; up to 6 CPs may be configured
2828-H13	2 drawers; up to 13 CPs may be configured

The GP **Speed Class** must be designated as /A00 (no GP CPs) or /F00 (1 GP CP seen by z/OS as 2828-F01).
A maximum of one General Purpose CP may be configured. Remaining CPs may be only be configured as full speed IFLs.

Notice: In zPCR, the references to these models is as follows:

2828-H06 will appear as **2828-L06**

2828-H13 will appear as **2828-L13**

This change was necessary in order to differentiate the **LinuxONE** models from the equally named **zBC12** models.

Specify the LPAR host's RCP configuration using the dropdown list provided for each of the following categories:

- **General Purpose** On **IBM z17-ME1** the range allowed is based on the MaxN setting.
On **IBM z16-A01** the range allowed is based on the MaxN setting.
On **IBM z16-A02/AGZ** the range is limited to 6 for GPs.
On **IBM z15-T01** the range allowed is based on the MaxN setting.
On **IBM z15-T02** the range is limited to 6 for GPs.
On **IBM z14-ZR1** the range is limited to 6 for GPs.
On the remaining **IBM Z** high-end z9 and later models, the range allowed is implied by the model specification, and you must provide a number.
On the remaining **IBM Z** mid-range and later models, this number is limited, based on the model specification and the maximum GPs allowed.
On **LinuxONE** processors, a single GP RCP (special purpose) may be defined (optional).
- **zAAP** May only be configured on z9 through z12 processor models. The maximum number of standard zAAPs is generally restricted, based on the total number of General Purpose CPs configured.
- **zIIP** May only be configured on z10 and later processor models. The maximum number of zIIPs is generally restricted, based on the total number of General Purpose CPs configured.
- **IFL** May be configured on any **IBM Z** or **LinuxONE** processor.
- **ICF** May be configured on any **IBM Z** processor. For z12 and later processor models, any number of RCPs may be configured as ICF. For z9 through z11 processor models, the number of ICF CPs is limited to 16.

Notes:

1. Engine assignment is validated by CP type as you make your choices. Therefore, you cannot specify more of any type than are available to be configured on the host processor model you have chosen.
2. zAAP and zIIP LCPs assigned to a z/OS partition are considered by **zPCR** as separate partitions in order to report capacity. However, these zAAP and zIIP partitions do not count towards the partition limit of the family being configured.
3. GPs, zAAP, zIIPs, IFLs, and ICF partitions each operate as a separate CP pool.
4. Any SCP may be assigned to a GP partition.
 - o z/OS, z/VM, and z/VSE can only be assigned to a GP partition.
 - o z/VM (running Linux guests), KVM, Linux, zAware, zACI, and SSC can be assigned to an IFL partition.
 - o CFCC can only be assigned to an ICF partition.
5. The workload category assigned to a zAAP or zIIP partition will be identical to that assigned to the parent z/OS GP partition.
6. If you reconfigure the LPAR host for which a partition configuration has already been defined, and the reconfiguration causes one or more partitions to become invalid, a dialog box appears, giving you a chance to cancel the change. If you accept the change, any partitions that become invalid will have their **Include** box unchecked; you will need to correct each excluded partition before it can be restored to the LPAR configuration. Any time that a partition has been automatically excluded, must be manually re-included.

For **IBM Z** processors, RCPs are considered in two groups:

1. **GP, zAAP, and zIIP Real CPs**
2. **IFL and ICF Real CPs**

When an LPAR configuration is established, these groups will be isolated from each other to the extent possible.

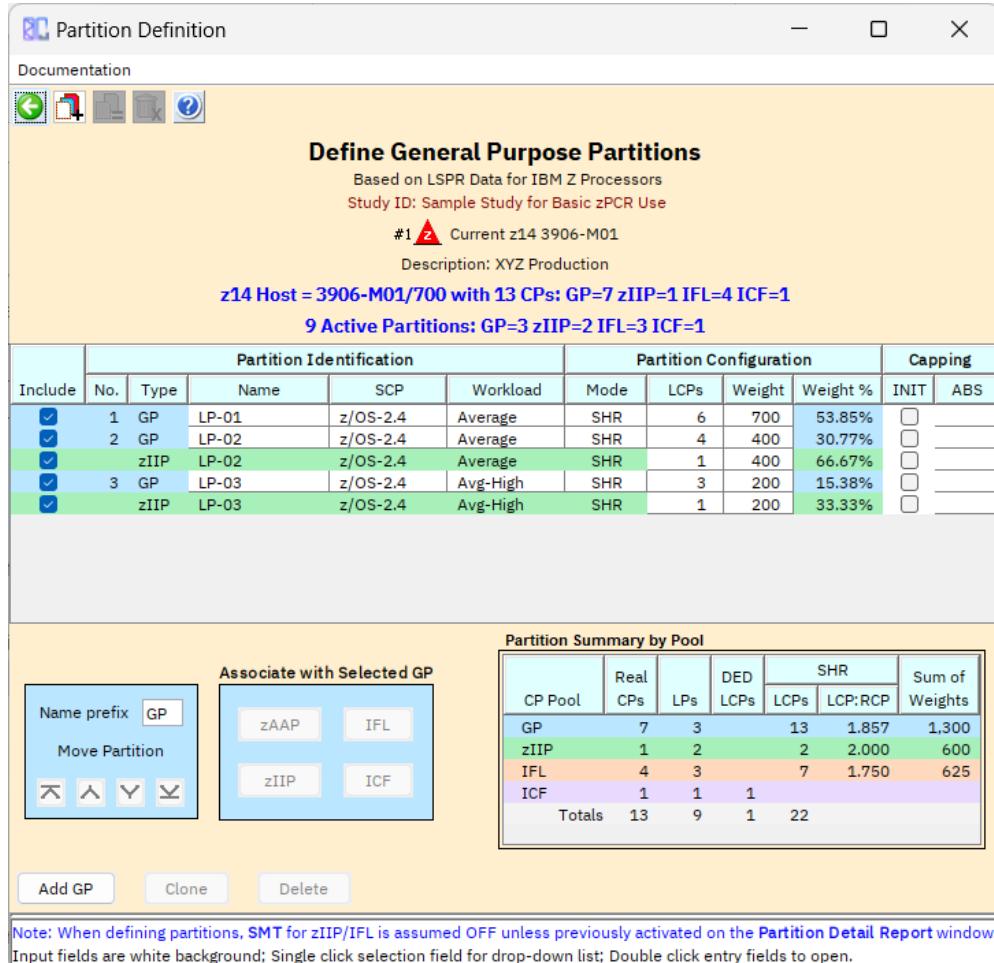
GP, zAAP and zIIP RCPs are allocated starting at one end of storage, while **IFL and ICF RCPs** are allocated starting at the other end of storage. On multi-book configurations an intersection of **GP/zAAP/zIIP RCPs** with **IFL/ICF RCPs** can occur in only one book/drawer. When such an intersection exists, algorithms apply a partitioning cost that reflects level of contention between the two CP groups in that book/drawer only. The partitioning cost for the RCPs in the remaining books will be less. A bar graph is available from the **Partition Detail Report** window showing how these RCPs would be allocated and where any contention between them might occur.

- On single-book/drawer systems, these two groups will always be in contention with each other (i.e., 100% contention).
- On multi-book/drawer systems, these two groups will be isolated from each other to the extent possible. There can only be one book/drawer where contention exists (contention can vary from 0% to 100%).

For **zPCR** chart examples depicting contention, see [Book and Drawer Considerations](#).

Partition Definition

LPAR Configuration Capacity Planning



The **Partition Definition** window is accessed from the **LPAR Host and Partition Configuration** window by clicking the **GP** button to define GP partitions, the **IFL** button to define IFL partitions, or the **ICF** button to define ICF partitions.

The **Partition Definition** table (at the top) provides an area where individual partitions are defined to the LPAR host, whose configuration is stated above. The **Partition Summary by Pool** table is presented below, summarizing the current partition definitions for each CP pool. This table is updated dynamically as partitions are added or modified.

When each of the 3 possible **Partition Definition** windows is initially opened, a single default partition of the requested type is automatically defined. This partition definition entry should be modified to match your requirements. Click the **Add** button to create additional partitions of the same type.

When defining additional GP partitions, click the **Add GP** button.

When defining a GP partition running z/OS, associated zAAP or zIIP partitions may also be defined.

When defining a GP partition running z/VM, zAAP, zIIP, IFL and ICF associated partitions may also be defined. Note that the only way that zAAP/zIIP and ICF LCPs could be exploited is to attach them to z/OS guest running in the z/VM partition. It is beyond **zPCR**'s capability to characterize capacity for z/VM guests.

The **Partition Definition** window serves to enter data defining each logical partition. The following fields comprise each partition's definition:

- **Include** Include/exclude this partition in the capacity assessment. Exclude assumes the partition is not HMC activated.

Partition Identification

- **No.** The relative partition number (determined by **zPCR**)
- **Type** The RCP pool to which the partition is assigned
- **Name** Partition name; assigned default can be modified
- **SCP** z/OS, z/VM, z/VSE, KVM, Linux, SSC, CFCC may be assigned to a GP partition. z/VM, KVM, Linux, zAware, zACI, and SSC can be assigned to an IFL partition. CFCC can be assigned to an ICF partition.

When the partition is HMC activated, but no SCP is running, the SCP should be set to “**Not IPL'd**”. logical CPs will remain in order to accurately represent the distribution of LCPs across the RCPs, and the assigned weights will be applied to the CP pool. The partition's minimum and maximum capacity values will set to zero.

- **Workload** For z/OS, z/VM, z/VSE, KVM, and Linux, 5 workload categories are available.
For zAware*, the workload must be zAware.
For zACI*, the workload must be zACI.
For SSC, the workload must be SSC.
For CFCC, the workload must be CFCC.

* Note: On z13 and later, use SSC in lieu of zACI and zAware.

Partition Configuration

- **Mode** DED (dedicated) or SHR (shared)
- **LCPs** Number of active Logical CPs
- **Weight** SHR partitions only; relative weight assigned

Capping metrics (SHR partitions only)

- **INIT** Checkbox indicates that partition is Initial Capped.
- **ABS** An Absolute Capping value may be set for **IBM z12** and later, and all **LinuxONE**. For detail see [Absolute Capping](#).

General Purpose Partitions with Associated Specialty LCPs

- **z/OS with zAAPs** Number of zAAP LCPs to be associated with a z/OS-1.6 (or later) partition. A switching cost is applied to both the GP and zAAP LCPs depending on zAAP utilization.
- **z/OS with zIIPs** Number of zIIP LCPs to be associated with a z/OS-1.6 (or later) partition.. A switching cost is applied to both the GP and zIIP LCPs depending on zIIP utilization.
- **z/VM with IFLs** Number of IFL LCPs to be associated with a z/VM partition. The partition must be on a z10 or later processor and must be designated to the hardware as **Mode=zVM**. No switching cost is applied for this association.

The following associations could only be exploited when assigned to a z/OS guest running under the parent z/VM. No switching cost is applied for these associations.

- **z/VM with zAAPs** Number of zAAP LCPs to be associated with the z/VM partition.
- **z/VM with zIIPs** Number of zIIP LCPs to be associated with the z/VM partition.
- **z/VM with ICFs** Number of ICF LCPs to be associated with the z/VM partition.

Note: **Associated LCP Partitions** must be defined with the same **LP Mode** (DED or SHR) as the parent GP partition.

Specifying Partitions

Upon initial access to this window, a single partition is automatically defined with default values supplied for the input fields. Input fields always have a white background; click on the field to obtain a dropdown list; for fields where manual entry is required, double click to open the field. Fixed and computed fields have a background color other than white. Modify the partition setting entries as desired. The codes below identify the status for each of the partition defining fields above:

- **P** Predefined field displaying generated information
- **E** Entry field; Double click field to open, key in text, and press **Enter**.
- **▼** Dropdown list; Click field to access dropdown list and make selection.
- **✓** Checkbox; Click field to check or un-check.

As partition data is entered, it is validated against the currently specified LPAR host, and the RCP resources remaining. Should an entry make the partition configuration invalid, you will be required to correct it before continuing.

Up to 130 partitions (sum of General Purpose, IFL, and ICF partitions) may be defined in **zPCR**. The maximum number of partitions (General Purpose, IFL, and ICF) that can be active depends on the host processor model selected (zAAP, zIIP, IFL, and ICF partitions associated with a parent GP partition do not count against these limits). Once that limit has been exceeded, a subsequently defined partition will have its **Include** un-checked. The partition **Include** checkboxes can be used to test various combinations of partitions on a given host when planning scenarios include many potential partitions and more than one CPC.

Partition Defining Controls

Click the **Add** button to add a new partition defined with default values; then modify the entries as desired.

Click the **Clone** button to make a copy of any single partition. In order to clone a partition, a previously defined partition entry must be selected. Click any non-entry field of the desired partition (e.g., the partition number) to select it. The original partition remains selected to accommodate multiple clone actions.

Any GP, IFL, or ICF partition may be cloned. If a GP partition has associated zAAPs, zIIPs, IFLs, or ICFs, they will be cloned along with the GP partition.

Note: Each of the above items will result in new partitions. When the real CPs are insufficient to accommodate the partition, its **Include** will be unchecked and an explanation dialog will be presented.

Click the **Delete** button to erase an existing partition definition. In order to delete a partition, it must be selected. Click any non-entry field of the desired partition (e.g., the partition number) to select it.

When defining GP partitions, any individual associated zAAP, zIIP, IFL, or ICF partition may be deleted. If deleting the parent GP partition all associated partitions will also be deleted.

Partition Specification Fields

A number of input fields are provided to configure a partition and define its workload activity. Default values are assumed when new partitions are defined, which should be modified as needed. These input fields are:

Include: Individual partitions that have been defined may be excluded from the capacity analysis by un-checking the partition's **Include** checkbox. When there are specialty LCPs associated with a GP partition, this setting also applies to them. This action can also be done from the **Partition Detail Report** window for each individually identified partition.

Partition Identification

E **Partition Name:** A default name is generated for each new partition as it is added. This name will be a short prefix (default = **LP**) followed by a dash and a sequential 3 digit number, indicating the order in which the partition was defined. This method assures unique names for each partition. The prefix to be assigned can be specified in the **Name prefix** entry field. Click on the input field, and enter a new prefix string.

A partition name can be changed by double clicking the field, keying in a new name, and pressing **Enter**. Names must be unique between partitions within an LPAR configuration. The one exception is that an associated zAAP, zIIP, and IFL partition name must be identical to that of its parent partition.

SCP: Must be specified for each partition. When defining GP partitions, z/OS, z/VM, z/VSE, KVM, Linux, zAware, zACI, SSC, or CFCC, where supported, can be selected. For IFL partitions, z/VM, KVM, Linux, zAware, zACI, SSC can be selected. For an ICF partition, CFCC (Coupling Facility Control Code) must be selected. If the SCP for a partition is changed, the workload field is changed to a default value for that SCP. Make your workload selection from the dropdown list provided.

1. **z/OS**: Various versions may be specified, based on the processor families listed below. The version is used only to enforce version-specific partitioning rules such as LCP limitations, associated zAAP/zIIP usage, and the use of SMT. Changing the z/OS version does not affect capacity results.

- (z17) **z17-ME1** **z/OS-3.2 through z/OS-2.4**
- (z16) **z16-A01, -A02, -AGZ** **z/OS-3.2 through z/OS-2.2**
- (z15) **z15-T01, -T02** **z/OS-3.2 through z/OS-2.2**
- (z14) **z14 and z14-ZR1** **z/OS-3.1 through z/OS-1.13**
- (z13) **z13 and z13s** **z/OS-2.5 through z/OS-1.10**
- (z12) **zEC12 and zBC12** **z/OS-2.4 through z/OS-1.10**
- (z11) **z196 and z114** **z/OS-2.2 through z/OS 1.7**
- (z10) **z10 EC and z10 BC** **z/OS-2.2 through z/OS-1.7**
- (z9) **z9 EC and z9 BC** **z/OS-2.2 through z/OS-1.4**

The z/OS version serves to enforce support rules shown in the table below. If in doubt about the actual version, specify the most recent LSPR measured version (z/OS-3.1).

z/OS Feature Support by Version				
z/OS Version	GP+zAAP/zIIP Maximum LCPs	Associated zAAP/zIIP LCPs	z13 and Later SMT zIIP LCPs	z10 and Later HiperDispatch
z/OS-3.2	208 *	Yes	Yes	Yes
z/OS-3.1	208 *	Yes	Yes	Yes
z/OS-2.5	208 *	Yes	Yes	Yes
z/OS-2.4	208 *	Yes	Yes	Yes
z/OS-2.3	200 **	Yes	Yes	Yes
z/OS-2.2	200 **	Yes	Yes	Yes
z/OS-2.1	190 ***	Yes	Yes	Yes
z/OS-1.13	100	Yes	No	Yes
z/OS-1.12	100	Yes	No	Yes
z/OS-1.11	80	Yes	No	Yes
z/OS-1.10	64	Yes	No	Yes
z/OS-1.9	64	Yes	No	No
z/OS-1.8	32	Yes	No	No
z/OS-1.7	32	Yes	No	No
z/OS-1.6	32	Yes	No	No

* 208 LCPs are supported only on IBM z17, and only if zIIP SMT is not enabled.
 ** 200 LCPs are supported only on IBM z16, and only if zIIP SMT is not enabled.
 *** 190 LCPs are supported only on z15, and only if zIIP SMT is not enabled.

Actual LCP limit is based on the LPAR host processor configured.

2. **z/VM**: Various versions may be specified based on the processor families listed below. The version is used only to enforce version-specific partitioning rules such as LCP limitations, associated IFL usage, and the use of SMT. Changing the z/VM version will not affect capacity results.

• (z17)	z17-ME1	z/VM-7.4 through z/VM-7.3
• (z16)	z16-A01, -A02, -AGZ	z/VM-7.4 through z/VM-7.1
• (z15)	z15-T01, -T02	z/VM-7.4 through z/VM-6.4
• (z14)	z14 and z14-ZR1	z/VM-7.3 through z/VM-6.4
• (z13)	z13 and z13s	z/VM-7.2 through z/VM-6.4
• (z12)	zEC12 and zBC12	z/VM-7.1 through z/VM-5.4
• (z11)	z196 and z114	z/VM-6.4 through z/VM 5.4
• (z10)	z10 EC and z10 BC	z/VM-6.4 through z/VM-5.4
• (z9)	z9 EC and z9 BC	z/VM-6.4 through z/VM-5.4

Logical CP considerations for z/VM IFLs:

- **z/VM-7.4, -7.3, -7.2, -7.1, and -6.4** support SMT for IFL CPs.
 - On IBM z17, z16, z15, and z14 processors a maximum of 80 IFL LCPs is supported; however, when SMT is enabled, IFL LCPs are limited to 40.
 - On z13 processors a maximum of 64 IFL LCPs is supported; however when SMT is enabled, IFL LCPs are limited to 32.
 - On z12 and prior processors, SMT is not supported; IFL LCPs are limited to 32.
- **z/VM-5.4** implies it and any prior z/VM version. These versions are out of support, but may be specified to partitions on z12 and prior LPAR hosts. IFL LCPs are limited to 32 (SMT is not supported).
- For **LinuxONE** processors, **z/VM-7.4, z/VM-7.3, z/VM-7.2, z/VM-7.1, or z/VM-6.4** is required. The same LCP configuration rules specified above apply.

The table below differentiates the features supported by the various z/VM versions.

z/VM Feature Support by Version				
z/VM Version	Maximum IFL CPs	Associated IFL LCPs	z13 and Later SMT on IFLs	z10 and Later HiperDispatch
z/VM-7.4	80 **	Yes	Yes	Yes
z/VM-7.3	80 **	Yes	Yes	Yes
z/VM-7.2	80 **	Yes	Yes	Yes
z/VM-7.1	80 **	Yes	Yes	Yes
z/VM-6.4	64 *	Yes	Yes	Yes
z/VM-6.3	64 *	Yes	Yes	Yes
z/VM-6.2	32	Yes	No	Yes
z/VM-6.1	32	Yes	No	Yes
z/VM-5.4	32	Yes	No	No
z/VM-5.3	32	No	No	No
Prior to 5.3	16	No	No	No

** 80 IFL LCPs are supported on IBM z17, z16, z15, and z14. When SMT is enabled the limit is 40.

* 64 LCPs are supported on z13. If IFL when SMT is enabled the limit is 32.

The remaining possible SCP assignments are considered generically (i.e., **zPCR** does not support specific partitioning rules for versions of these SCPs).

3. **z/VSE** can be defined on any IBM Z processor GP LCPs (limit = 4).
4. **KVM** can be defined on **IBM z12** and later processor GP or IFL LCPs. This also applies for **LinuxONE** IFL CPs.
5. **Linux** can be defined on any IBM Z processor GP or IFL LCPs.
6. **zAware** can be defined on **IBM z12** and **z13** processor GP or IFL LCPs only. A single workload (zAware) is available for assignment. zAware capacity data is derived from the z/OS Average workload category. **SSC** is intended to replace zAware on **IBM z13** and later processor models.
7. **zACI** (Appliance Container Infrastructure) can be defined on **IBM z13** processors only. A single workload (zACI) is available for assignment. zACI capacity data is derived from the Average workload category. **SSC** is intended to replace zACI on **IBM z13** and later processor models.
8. **SSC** (Secure Service Container) can be defined on **IBM z13** and later processor GP or IFL LCPs. A single workload (SSC) is available for assignment. **SSC** capacity data is derived from the Average workload category. **SSC** is intended to replace zAware and zACI; on the IBM z17, z16, z15, z14, **SSC** must be specified.
9. **CFCC** can be defined on all **IBM Z** processor GP or ICF LCPs. A single workload (CFCC) is available for assignment. Capacity is based on measurements for the unique CFCC level available for each processor family.

Note concerning optimal performance: Despite the maximum number of LCPs supported by an SCP, the count should normally be limited to the number of real CPs available in the largest drawer or book. For z/OS partitions this applies to the sum of its GP and zIIP LCPs. For z/VM partitions this applies individually to its GP and to its IFL LCPs.

Note that for **IBM Z** (z14, z13, zEC12, z196), and **LinuxONE Emperor**, **Power-Save** mode can be turned on regardless of the SCP version that is running. However, the ability to report on its status is only available from a partition running z/OS-1.10 or later.

The z/OS or z/VM version selected for a partition will not affect capacity results since they are driven off a single LSPR table derived from z/OS-2.4 measurements. While actual capacity relationships may vary somewhat for different SCP versions, the difference is considered insignificant enough to allow the use of a single LSPR table. For z/VM-5.4 and later, LCPs can be General Purpose or IFL CPs. On all z10 and later processors, IFLs can be associated with a parent z/VM GP partition if the partition is defined as **Mode=zVM**.

Defined partitions that are HMC activated, but have no SCP IPL'd (i.e., not running) the SCP should be set to "**Not IPL'd**". Logical CPs will remain in order to correctly represent the distribution of LCPs across the RCPs, and the assigned weights will be applied across the CP pool. The partition's minimum and maximum capacity values will be assumed as zero.

Defined partitions that are not HMC activated, should not be included, since no LPAR overhead is generated on behalf of them. The **Include** checkbox should be made unchecked.

Workload: Must be specified for each partition. Selection is limited to those workloads that are valid for the SCP specified. For **z/OS**, **z/VM**, **z/VSE**, **KVM**, and **Linux**, the workload can be any of the 5 supported workload categories. Make your selection from the dropdown list provided. For **zAware**, **zACI**, **SSC**, and **CFCC**, a single workload selection available.

Partition Configuration

Mode: Must be specified for each partition as DED (for a dedicated partition) or SHR (for a shared partition). Make your selection from the dropdown list provided.

LCPs: Set the number of active Logical CPs to be assigned to the partition. The selection list only includes counts that are valid for the currently defined LPAR host configuration. Adequate RCPs must be available in the pool to which the partition is assigned. Dedicated partitions will remove RCPs from the shared pool.

Note: There are limitations on the number of LCPs that can be specified for z/OS, z/VM, z/VSE, and CFCC partitions.

E **Weight:** Specified for each shared partition. Double click the field, keying in a whole number and press Enter. A default value of 100 is used if no value is provided. The individual weights for partitions in each pool are summed to determine each partition's weight percentage.

Initial Capping

: A checkbox used to indicate that a shared partition's Maximum Capacity value will be limited to its Minimum Capacity value. The capacity values are based on its Logical CPs and Weight % relative to its CP pool.

Absolute Capping

ABS: Supported on all z12 and later processors. The **ABS** value is entered as a fractional N.nn, and must be between zero and the number of LCPs defined. The minimum value allowed will be dependent on the overall CP pool configuration.

Partition capacity results will be based on the least restrictive of the capping settings.

Partitions must initially be defined from the **Partition Definition** window. Once defined, any partition's definition metrics (of those shown above) can be modified from the **Partition Detail Report** window, discussed later.

LCP associations can only be changed from the **Partition Definition** window.

HiperDispatch

HiperDispatch is supported on all z10 and later processor families. Both z/OS and z/VM support HiperDispatch.

z/OS LSPR data for all z10 and later processor families represents HiperDispatch as active for defined partitions. HiperDispatch achieves improved performance by attempting to keep workload elements on the same logical (and physical) CP. This aspect of HiperDispatch is inherent in the LSPR data used by [zPCR](#).

HiperDispatch can also improve z/OS performance by dynamically parking partition LCPs that it considers to be excessive for the workloads competing for the shared CP resource. This aspect of HiperDispatch is not included in the LSPR data or algorithms used by [zPCR](#). To fairly represent capacity for these configurations, **parked LCPs should not be included when defining the number of shared LCPs for a partition**. When **Total LCPs** is reduced by **Parked LCPs**, the result is termed **Active LCPs**. **Parked LCPs** are determined from EDF (z/OS or z/VM) or RMF data (z/OS only). Note that a separate input file is required for each partition to obtain parked LCPs.

With EDF input, parked LCPs that are not represented by an EDF (GP, zAAP, zIIP, and IFL LCPs) can be estimated for z/OS and z/VM partitions. This capability is available when the EDF interval has been selected and before the configuration is to be transferred into [zPCR](#).

Special consideration when defining a current operating environment

For shared partition configurations, [zPCR](#) uses the weight percent to allocate processor resource, which is quantified as the **Minimum Capacity** available to each partition. This assumes that every partition is competing for CPU resource, and the partition's weight percent determines the capacity realized. The sum of the partition **Minimum Capacity** values is taken to be the effective capacity of the LPAR host.

For production workloads, actual partition weights often may not align with the actual distribution of the processor resource. When a shared partition with a significant weight assignment, tends to idle, its presumed processor resource is available to be used by other active partitions. In such cases, it may be better to use the actual partition utilization (from RMF's **Partition Data Report**) as the weight value in [zPCR](#). In this way, the actual distribution of the processor resource to each partition is represented, resulting in a more accurate capacity projection. When modeling the planned LPAR host replacement, a decision will be necessary concerning the weights to be assigned.

Entry Fields

The **Default Prefix** entry field provides a way to customize partition names as they are defined. When the prefix is changed, it will affect partitions that are subsequently added. The prefix will have a sequential number appended, indicating the order in which the partition was defined (assuring a unique default name for each partition). The default name can be changed to any desired character string.

To change the **Default Prefix**, click on the entry field to set focus, key in your change, and press Enter.

Repositioning Partitions

The sequence of defined partitions can be modified using the **Move** buttons provided. There are move controls for **Top**, **Bottom**, **Up**, and **Down**.

Partitions can only be repositioned within their LP pool boundaries. All reports showing multiple LP pools will have the General Purpose pool first, followed by the zAAP, zIIP, IFL, and ICF pools.

Any repositioning done with these controls is retained with a saved study.

Specialty LCP Associations with a Parent GP Partition

zAAP and **zIIP** LCPs may be associated with a parent GP partition running **z/OS**. zAAP associations are supported starting with z9. zIIP associations are supported starting with z10.

zAAP, **zIIP**, **IFL**, and **ICF** LCPs may be associated with a parent GP partition running **z/VM**. These associations are supported on z10 and later families.

To create an LCP association, select the GP partition. Then, from the **Associate with Selected GP** group box, click the **zAAP**, **zIIP**, **IFL** or **ICF** button. Buttons are enabled only if the LCP type is supported and configured on the LPAR host.

Note: It is unlikely that one would manually define a z/VM partition with zAAP, zIIP and/or ICF associated LCPs. However, such configurations could exist when loading from an EDF or RMF.

Associated LCPs are considered as separate partitions for capacity projection purposes. However, they do not count against the number of partitions allowed on the LPAR host. A single central storage is shared by associated partitions, and each of their parent GP partitions.

When created, an associated partition inherits the parent GP partition **Number**, **Name**, **SCP**, and **Workload** assignment, which can only be changed by changing that of the parent GP partition. Two exceptions are note below:

- IFL LCPs associated with z/VM: the workload will be initially set to that of the parent z/VM partition. However, the workload may be changed to any of those valid for z/VM.
- ICF LCPs associated with z/VM, the workload will be CFCC.

If the parent partition is z/OS, the **SCP** can only be changed to another version of z/OS. If the parent partition is z/VM, the **SCP** can only be changed to another version of z/VM.

The associated partition's **Mode**, (DED or SHR) must be identical to that of the parent GP partition.

The associated partition's **Weight** and **Capping** assignments are initialized to default settings, and may be changed to any legitimate value. These metrics are totally independent of those of the parent GP partition.

Any change to the LPAR host or a partition's definition that causes it to become invalid will trigger a dialog box explaining the problem, and the partition's **Include** box will be unchecked; you will need to correct the excluded partition before it can be restored to the LPAR configuration. Any time that partition has been automatically excluded, it must be manually re-included.

Should the SCP assigned to a partition be changed, and the workload becomes an invalid name for that SCP, "**unknown**" will appear as the partition's workload, and the partition will be excluded from the study until a proper workload is selected.

Partition Summary by Pool

A table beneath the partition definitions labeled ***Partition Summary by Pool*** provides a summary of the partition resources for each pool. The table is dynamically updated as partitions are added or modified.

System Recovery Boost Considerations

System Recovery Boost metrics cannot be specified when manually defining a partition, SRB metrics can only be input to **zPCR** via EDF or RMF.

When the LPAR configuration was created from EDF or RMF, and includes SRB activity, a **SRB** column will be added to the ***Partition Definition*** window, window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

Note: Initial Capping and Absolute Capping are ignored for zIIP partitions using zIIP Boost.

For detail concerning SRB support see [System Recovery Boost](#).

Window Controls

Click the **Return** toolbar icon to return to the ***LPAR Host and Partition Configuration*** window. All partition definitions, as entered, will participate in the capacity assessment.

Click the **Add (+)** toolbar icon to add another partition, initialized with default settings.

Click the **Clone (=)** toolbar icon to clone a selected partition

Click the **Delete (x)** toolbar icon to delete the selected partition.

Click the **Help (?)** toolbar icon for help concerning this window.

Menu-bar

Documentation

Provides access to the various **Notice** windows associated with this window

EDF Overview

LPAR Configuration Capacity Planning

An **EDF** (Enterprise Data File) can be used to create the entire LPAR host and partition configuration or to copy individual partition definitions into a currently active LPAR configuration. EDFs can be created from partitions running z/OS or z/VM. Each EDF represents a single partition. EDFs generally represent multiple time intervals. EDFs from multiple partitions can be read into **zPCR** at the same time. All EDFs should include the same intervals, and all must include the same interval selected for **zPCR**.

zPCR can read EDFs generated by reasonably current versions of z/OS or z/VM, using the applicable extract program (**CP3KEXTR** for z/OS; **CP3KVMXT** for z/VM). If the EDF was generated from an unsupported z/OS or z/VM version, or an unsupported version of the extract program that creates the EDF, a dialog will be presented concerning the supported versions.

EDFs must be downloaded to your PC (preferred file extension = **.edf**). Only one EDF is necessary to obtain the LPAR host, partition definition, and partition utilization information (referred to as the primary EDF).

The first EDF, considered the master, contains all the necessary information to define the LPAR host processor and its entire partition configuration. It also includes information specific to the partition represented. The remaining EDFs will be used to obtain partition specific information. The following EDF information is used by **zPCR**.

Note: The filename assigned to an EDF should not contain any special characters. Otherwise, problems may occur when transferring the configuration into **zPCR**.

The table below documents the EDF information used by **zPCR**.

LPAR Configuration Metrics Determined from EDF	
An EDF created from any single partition reveals the following information for the entire CPC. This information is always obtained from the 1 st EDF read.	
LPAR host	Selected Interval (Date, Time, and Duration)
	CPC ID
	Specific processor family and model
	IBM z17, z16, z15, z14-ZR1 MaxN definition
	RCP count for GP, zAAP, zIIP, IFL, ICF
Partition detail	Type: GP, zAAP, zIIP, IFL, ICF
	Name
	Mode: Dedicated or Shared
	LCP count
	Weight
	Capping
	Absolute Capping
	Utilization
	HMC activated but not IPL'd
	System Recovery Boost activity
Individual EDFs can reveal additional information concerning each partition	
SCP: Specific z/OS or z/VM version running	
Workload: RNI (CPU MF required) for appropriate LSPR workload assignment	
HiperDispatch: Active and number of parked LCPs	
Measured SMT Benefit: IBM z17, z16, z15, z14, z13 only; zIIP/IFL partitions only	
Topology: IBM z17 and z16 only	
Where indeterminate, the following assumptions apply. If an actual assignment differs, it can be changed after the configuration is loaded into zPCR.	
SCP	GP/zAAP/zIIP: z/OS
	IFL: z/VM
	ICF: CFCC
Workload	Average category

Additional information specific to a partition (listed below) may also be captured (note that a separate EDF must be generated for each partition).

1. If **CPU MF** is enabled (z10 and later) for the partition, counters will be captured. This information is used to compute an RNI value, and subsequently the appropriate LSPR workload category to be assigned for that partition.
2. If **HiperDispatch** is enabled for the partition (z10 and later running z/OS or z/VM), the EDF will provide information concerning the number of parked LCPs for that partition. General Purpose LCPs (and any associated ZAAP, zIIP, or IFL LCPs) may be parked. For these EDFs, **zPCR** will identify both the number of assigned LCPs and the number of parked LCPs for the interval selected for transfer. When partitions are transferred into **zPCR**, all parked LCPs that were identified will be removed from the total partition LCP count by default (a checkbox is provided to override this default). Any parked LCPs not removed when the partition is transferred, can subsequently be removed in **zPCR** itself.
3. If **SMT** is enabled (z13 and later; zIIP and IFL LCPs only), the EDF will provide information concerning the actual benefit over running without SMT enabled.

In order for multiple EDFs to be processed, the following conditions must exist or the EDF processing will be skipped.

- Each EDF must represent the same CPC ID as the master. In cases where multiple CPC IDs exist, each must be defined independently as a separate LPAR host.
- Each EDF must represent the same period of time, within a defined tolerance. Note: For any EDF describing a partition with its time offset different than the master EDF partition, its interval times will be adjusted to be represented based on the master EDF time zone. Doing so assures that the interval picked will represent the same physical time for all partitions.
- The start time for each EDF interval is adjusted to the nearest 5 minutes. This allows the intervals across multiple EDFs to be aligned more consistently.

On z/OS partitions where IRD is actively managing resources, partition LCPs are reported by EDF as fractional values. For **zPCR** purposes, fractional LCP numbers are rounded up to the nearest whole number. Since weights can also be dynamically modified by IRD, it is important that the interval selected for analysis be one that is representative for a load period of particular interest.

EDFs are read in via the **EDF Interval Selection** window (see [EDF Interval Selection](#)). Once processed, each measurement interval found is listed. A single measurement interval must be selected from which to obtain the LPAR configuration information. Using the master EDF, the LPAR configuration is determined. Partition specific information (i.e., **CPU MF**, **HiperDispatch** and **SMT Benefit**) is obtained from each individual EDF.

EDFs created from z/OS

When running z/OS in a partition, SMF data is normally captured for the purpose of monitoring performance and capacity. SMF can capture a wide variety of record types, resulting in data which is voluminous and binary. This data can be reduced to an **EDF (Enterprise Data File)** that can be used by **zPCR** to automatically generate the LPAR configuration (same as can be done from an RMF report). With EDFs, **zPCR** usability can be further enhanced when certain SMF record types are captured for partitions. The record types used are listed below.

SMF Record Types used by zPCR

- **SMF Type 70** – CPU Activity Record (required)
- **SMF Type 113** – CPU Measurement Facility (**CPU MF**) counters (optional but highly desired; only available on z10 and later processors). This information is used by **zPCR** to choose the LSPR workload category that will best represent characteristics of the partition's production workload.

Information on capturing **CPU Measurement Facility** data can be found at:

www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TC000066

Creating z/OS EDFs

The z/OS extract program, **CP3KEXTR**, is available from your **zPCR** download website along with instructions for its use. It is distributed as a z/OS “Load and Go” job, to be run on any **IBM Z** processor. **CP3KEXTR** reads SMF data, creating an EDF. Each EDF represents a single z/OS partition, while also including the information necessary to determine the LPAR host processor and its entire partition configuration. When there are multiple partitions for which SMF type 113 (**CPU MF** hardware counters) are to be used, a separate EDF must be created for each. To be useful, all the EDFs for a single study must represent the same time period and should normally include the same number of intervals.

In order for EDF data to be valid for **zPCR** use, a recent version of **CP3KEXTR** is necessary. If the version used to generate the EDF is not currently supported, **zPCR** will present a dialog with information concerning the version(s) required.

EDFs created from z/VM

When running z/VM in a partition, VM Monitor data is normally captured for the purpose of understanding performance and capacity. VM Monitor can capture a wide variety of data. This data can be reduced to an **EDF (Enterprise Data File)** that can be used by **zPCR** to automatically generate LPAR configuration information.

When the LPAR host is a z10 or later processor, z/VM-5.4 and later can capture CPU MF data in an EDF. This information is used by **zPCR** to choose the LSPR workload category that will best represent characteristics of the partition's production workload. When z/VM is running in a GP partition while also supporting an associated IFL partition, separate CPU MF data is captured for each. Hence, the workload category chosen for each type could be different.

Creating z/VM EDFs

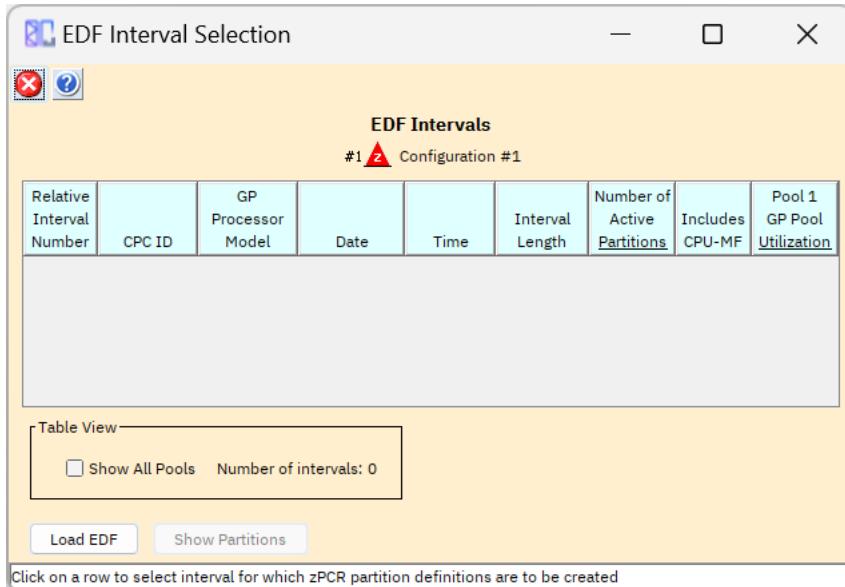
The z/VM extract program, **CP3KVMXT**, is available on your **zPCR** download website along with instructions for its use. It must be run under z/VM in a CMS guest. **CP3KVMXT** reads VM Monitor data to generate an EDF. The EDF represents the overall workload of a single z/VM partition, while also including the information necessary to determine the LPAR host processor and its entire partition configuration. To be useful, all the EDFs for a single study must represent the same time period and should normally include the same number of intervals.

In order for EDF data to be valid for **zPCR** use, a recent version of **CP3KVMXT** is necessary. If the version used to generate the EDF is not current enough, **zPCR** will present a dialog with information concerning the version(s) required.

EDF Interval Selection

LPAR Configuration Capacity Planning

From the **LPAR Host and Partition Configuration** window, click the **EDF** button in either the **Define LPAR Host Processor** group box or the **Define Partitions** group box to open the **EDF Interval Selection** window.



The **EDF Interval Selection** window provides the means to load one or more EDFs. For general information concerning EDFs, see [EDF Overview](#).

Click the **Load EDF** button. A file dialog box will appear from which you can specify the drive:\directory\filename of the desired EDF (multiple files can be selected). While the default EDF extension is **".edf"**, **".txt"** may also be used.

When read, the EDFs are searched for all report intervals. If the file does not appear to be a standard EDF or appears to be invalid, a message is issued and no EDF intervals are shown.

Loading Multiple EDFs

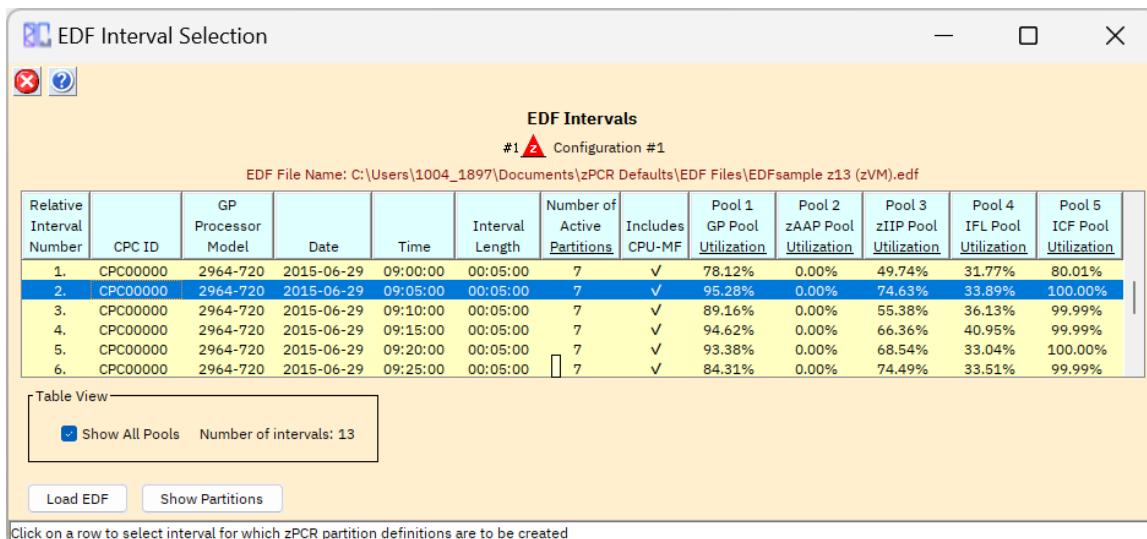
zPCR can process multiple EDFs in a single pass, each representing an individual partition. The 1st file in the selection list is considered the master, serving as the basis for the LPAR host processor and its entire partition configuration. Each individual EDF is used to gather partition specific information.

To load multiple EDFs, hold down the **Ctrl** (or **Shift**) key while selecting files. The selected files can be dragged to any undefined LPAR icon on the **Control Panel** window, and they will all be processed. On the **EDF Interval Selection** window, when multiple EDFs are selected, they are placed in the selection list in the same order as displayed in the **Load EDF** dialog. With this order, the 1st entry may not be desired as the master EDF. The order of the selection list can be controlled somewhat by going to the **Details** view and sorting (ascending or descending) on **Name**, **Size**, or **Date**, and then selecting the multiple files. Alternatively, before clicking the **Open** button, cut

(**Ctrl-X**) the desired master EDF from the list, and then paste (**Ctrl-V**) it back at the beginning of the list.

zPCR assumes that the report intervals of multiple EDFs, when loaded, will be equal in number. If they are not, a dialog will appear showing the total number of intervals, and the number of intervals that do not match. Two choices are provided: 1) **Proceed** using only the intervals that do match; or 2) **Cancel**. In addition, **zPCR** assumes that the interval time stamps between the EDFs will be closely aligned. If they are not, no EDF intervals will be shown.

Once validated, the EDF has loaded and a title line reveals the EDF filename and location. The EDF will remain available until a different EDF is loaded.



The screenshot shows the 'EDF Interval Selection' dialog box. At the top, there are buttons for Close (X) and Help (question mark). The main title is 'EDF Intervals' with a note '#1 2 Configuration #1'. Below that, it says 'EDF File Name: C:\Users\1004_1897\Documents\zPCR Defaults\EDF Files\EDFsample z13 (zVM).edf'. The main area is a table with the following columns: Relative Interval Number, CPC ID, GP Processor Model, Date, Time, Interval Length, Number of Active Partitions, Includes CPU-MF, Pool 1 GP Pool Utilization, Pool 2 zAAP Pool Utilization, Pool 3 zIP Pool Utilization, Pool 4 IFL Pool Utilization, and Pool 5 ICF Pool Utilization. The table contains 6 rows of data. At the bottom left of the table area, there is a 'Table View' section with a checkbox labeled 'Show All Pools' and a note 'Number of intervals: 13'. Below the table, there are buttons for 'Load EDF' and 'Show Partitions'. A note at the bottom says 'Click on a row to select interval for which zPCR partition definitions are to be created'.

Each EDF interval found is displayed in a sequentially numbered row. Interval descriptions include **CPC ID**, **GP Processor Model**, **Date**, **Time**, **Interval Length**, **Number of Active Partitions**, **Includes CPU MF**, and **GP Pool Utilization** (generally the CP pool of most interest). Utilization for all CP pools can be displayed by checking the **Show All Pools** checkbox.

In the **Includes CPU MF** column, a checkmark will appear if **CPU MF** is included for at least one partition for the report interval.

For z17 and z16 LPAR hosts with topology information, a **Topology Changed** column will be included. The presence of this column indicates that topology information is available. A check mark will appear for each interval where the topology changed.

To facilitate identification of report intervals where the greatest amount of CPU resource is consumed, the columns with underlined headings are enabled for sorting. Click on the heading once to sort in descending sequence, again to sort in ascending sequence, and once again to restore the original order.

From the table, select the report interval for which LPAR configuration information is to be obtained by clicking on the desired interval. The selected row will be highlighted.

Then Click the **Show Partitions** button to open the **Create LPAR Configuration from EDF** or **Get Partitions from EDF** window where the partition definitions can be reviewed.

System Recovery Boost Considerations

Intervals that include any partition running with **System Recovery Boost** activity will be highlighted with a red background.

Note: To recognize SRB activity, the latest version of CP3KEXTR (v4.06 or higher) must be used to generate the EDF.

For detail concerning SRB support, see [System Recovery Boost](#).

Window Controls

Click the **Load EDF File** button to select the specific EDF to be loaded.

Click the **Show Partitions** button to review the partitions that were active during the selected interval. The **Create LPAR Configuration** or **Get Partitions from EDF** window will appear.

Click the **Cancel** toolbar icon to close the window and return to the **LPAR Host and Partition Configuration** window.

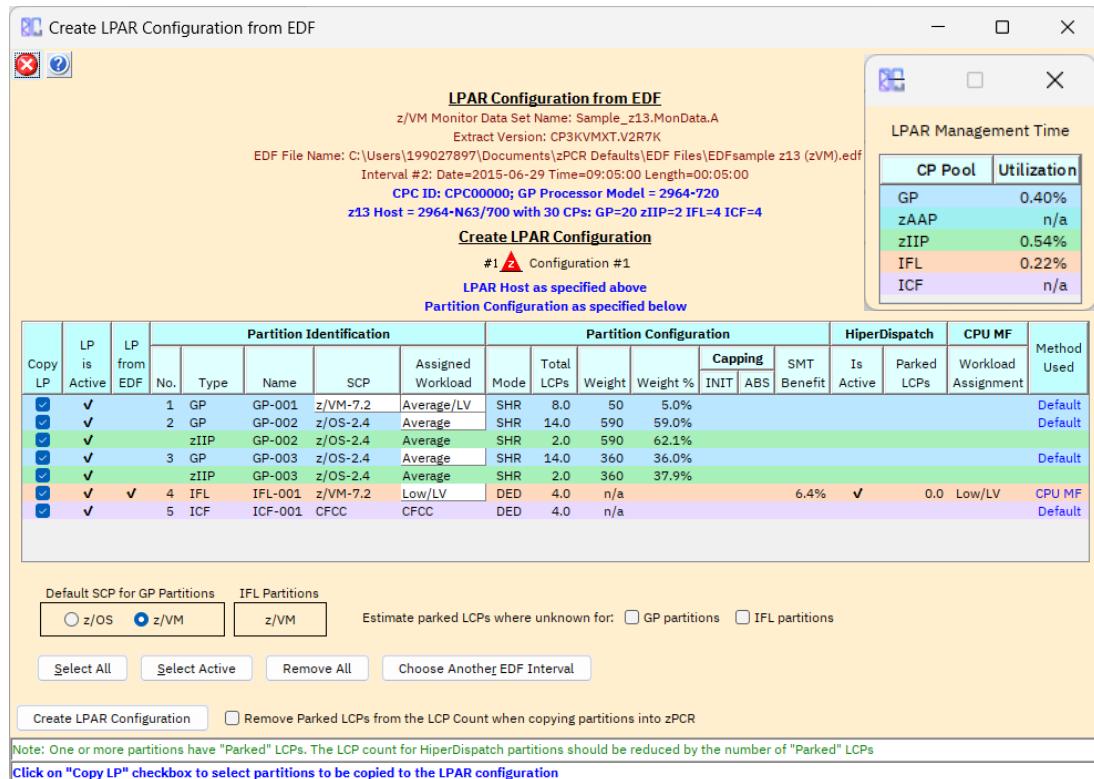
Get Partitions from EDF

LPAR Configuration Capacity Planning

From the **EDF Interval Selection** window, click on an interval row to enable the **Show Partitions** button. Clicking the **Show Partitions** button will:

- When entered via [Create Host and Partitions From EDF](#), the **Create LPAR Configuration from EDF** window appears.
- When entered via [Copy Partitions From EDF](#), the **Copy Partitions from EDF** window appears.

Except for the window titles and some button names, these two windows are identical.



In addition to the primary window, a small **LPMT** window appears showing the LPAR management time for each CP pool as reported by EDF (informational purposes only).

At the top of the window, title lines reveal the location and filename of the currently loaded EDF report, and the interval selected. Just below is listed the LPAR host system ID and its General Purpose processor model as noted by EDF. The next line identifies the overall LPAR host configuration as would be shown in [zPCR](#).

When the window title reads **Create LPAR Configuration from EDF**, the lines below reveal, what will become the host processor for the LPAR configuration. When the window title bar reads **Copy Partitions from EDF**, the currently defined LPAR host processor is simply identified. In either case, all the partitions that are selected in the table below will be transferred into the LPAR configuration.

Partition Utilization Considerations

Partition utilizations, while known, are not displayed on the ***LPAR Configuration from EDF*** window. However, they will be transferred into **zPCR** and will be included in a saved study. When the study for the original configuration is reloaded, the partition utilization values can be viewed in the ***Utilized Capacity Report*** window. In addition, the actual zAAP and zIIP partition utilizations will be available for assignment on the ***zAAP/zIIP Loading*** window.

SMT Benefit

For z13 and later processors, measured ***SMT Benefit*** values can be obtained from RMF for SMT enabled zIIP and IFL partitions. An additional column, ***SMT Benefit***, is displayed immediately following the ***Capping*** columns with this information. When SMT information is found, it will be transferred to **zPCR** as a ***Measured SMT Benefit***.

When a zIIP or IFL partition's utilization is less than 20% or when the measured SMT Benefit is greater than 90%, the SMT Benefit value is considered as being unreliable. While transferring the configuration into **zPCR**, a dialog will identify such cases and the default ***Estimated SMT*** value will be assigned instead. While viewing the ***SMT Benefit*** window, the ability to use the ***Measured SMT*** benefit is provided.

Partition Identification and Partition Configuration

The primary table contains a list of all the partitions found in the selected EDF interval with their then defined metrics. From this table you select the specific partitions to be copied.

All active partitions found in the selected EDF interval are listed in the table in the following order:

1. GP (General Purpose)
2. zAAP
3. zIIP
4. IFL
5. ICF

Note: zAAP, zIIP, IFL, and ICF LCPs that are associated are listed just below the parent GP partition.

The format of each row is similar to that used in **zPCR**, with background colors differentiating each of the partition pools. Partition weights are summed by pool to determine each partition's weight percent.

A checkmark in the column identified as ***From EDF*** indicates that an EDF was produced for that partition.

SCP Assignment

The **SCP** field is populated with the SCP version found in the partition's EDF or z/VM EDF, when available. Otherwise, the field is populated with a default SCP for the partition type as follows:

- All remaining GP partitions will be assigned the default **z/OS** version for the LPAR host. When another SCP is actually in control, the SCP will need to be changed manually.
- Any remaining IFL partitions will be assigned the default **z/VM** version for the LPAR host. When another SCP is actually in control, the SCP will need to be changed manually.
- For **zAAP** or **zIIP** partitions associated with a parent z/OS GP partition, the same **z/OS** is assigned for each.
- For **zAAP**, **zIIP**, **IFL**, and **ICF** partitions associated with a parent z/VM GP partition, the GP **z/VM** version is assigned as the SCP for each.
- If the SCP settings for GP and IFL partitions are corrected before the configuration is transferred into **zPCR**, **Parked LCPs** will automatically be corrected to only represent those partitions running z/OS or z/VM. If these changes are to be made in **zPCR**, GP **Parked LCPs** should not be estimated.

Partitions that are HMC activated, but have no SCP IPL'd (i.e., not running) will have the SCP set to “**Not IPL'd**”. Logical CPs will remain in order to accurately represent the distribution of LCPs across the RCPs, and the assigned weights will be applied across the CP pool. In **zPCR**, the partition's minimum and maximum capacity values are assumed as zero (no value is displayed).

Note: **Every partition detected in an EDF is HMC activated.**

All partitions found in an EDF should be included for reliable capacity results.

Workload

The **Method Used** column (on the right side of the window) reveals the method applied for **Assigned Workload**. When the partition is represented by an EDF which includes CPU MF information, the workload will be determined by the RNI and L1MP.

If an **Assigned Workload** is subsequently altered, the **Method Used** will change accordingly.

A zAAP or zIIP partition associated with a parent z/OS or z/VM partition will be assigned the same workload as the parent. An IFL partition associated with a parent z/VM partition will be assigned the same workload as the parent (it can be modified in **zPCR**). An ICF partition associated with a parent z/VM partition will have CFCC assigned as the workload.

Mode (Dedicated or Shared)

Total LCPs

Shows the total number of LCPs assigned to the partition. A fractional LCP count indicates that the partition had LCP(s) varied off or on during the measurement interval. When the configuration is transferred into **zPCR**, fractional values are rounded up to the next integer.

Weight and Weight %

Shows the partition weight and calculated weight percentage relative to the other partitions in the CP pool.

Initial Capping

If a partition was initial capped, a checkmark will appear in the INIT column.

Absolute Capping (z12 and later processors only)

If **Absolute Capping** is defined to a partition, a value will appear in the **ABS** column.

SMT Benefit for zIIP and IFL Partitions

On z13 and later processors, **SMT Benefit** is obtained from EDF or RMF for SMT enabled zIIP partitions running z/OS and for IFL partitions running z/VM, KVM, or Linux. An additional column, **SMT Benefit**, is displayed with this information immediately following the **Capping** columns.

The SMT benefit is considered reliable when the partition's utilization is 20% or greater. In this case the benefit will be assigned as a **Measured SMT Benefit**. When the partition's utilization is less than 20%, the default **Estimated SMT Benefit** value will be assigned instead.

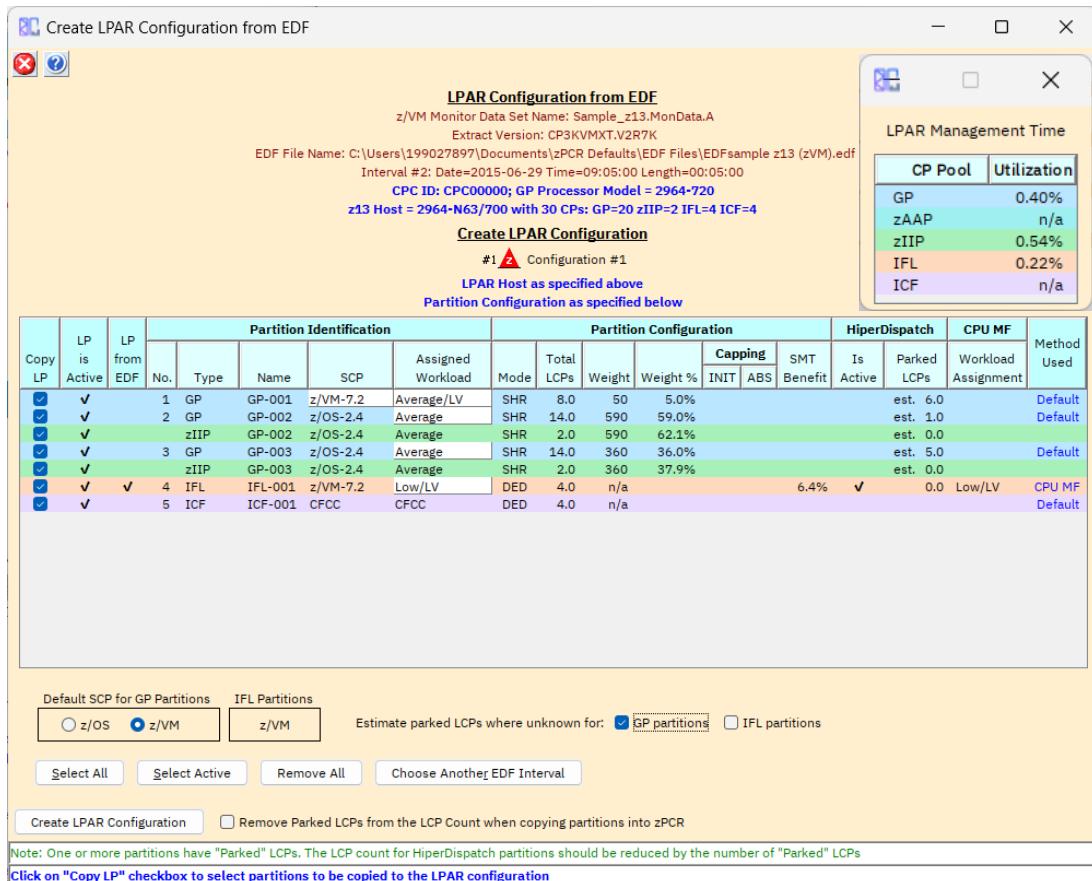
There are other scenarios where the measured SMT Benefit value is limited to a maximum value. In this case the value will be assigned as a **Measured SMT Benefit**. When the value is zero or less, the **Measured SMT Benefit** will be set to 0%. SMT can be manually disabled if desired.

HiperDispatch (z10 and later; running z/OS or z/VM)

- **HiperDispatch Is Active:** A checkmark indicates whether or not HiperDispatch was active for the partition.
- **Parked LCPs:** Indicates the actual number of the partitions LCPs that were parked during the measurement interval. The Parked LCP count will display with 1 decimal place, since parking activity is dynamic across the selected interval.

Estimating Parked LCPs

If the LPAR host supports Hiperdispatch (z10 and later), parked LCPs can be estimated for z/OS and z/VM shared partitions that are not represented by an EDF. This also applies to any associated zAAP or zIIP partitions. Parked LCPs can be estimated for IFL shared partitions not represented by an EDF and running z/VM. The Parked LCP estimate will display with one decimal place.



When the **Create/Copy from EDF** window is opened, a dialog is displayed to help emphasize the ability to estimate parked LCPs. The dialog can remain open while making changes to the window.

Parked LCP estimates are controlled using the check boxes displayed below the J-table. When GP or IFL is checked, a parked LCP estimate of will be made (1 decimal precision) for the non-EDF shared partitions in that pool. Estimated values are prefixed with "Est." so as to differentiate them from actual parked values. Unchecking the box will remove the parked LCP estimates.

Notes:

- Each partition's **Copy LP** checkbox defaults to its **LP is Active** checkbox. If **Copy LP** is changed, all **Parked LCP** estimates for the CP pool will be recomputed.

- Changes to any partition's SCP other than z/OS or z/VM will cause **Parked LCPs** estimates for the CP pool will be recomputed. If such changes are to be made after transfer to **zPCR**, **Parked LCPs** should not be estimated, since the estimates would no longer be relevant.

To most accurately reflect partitioned capacity, **Parked LCPs** should be removed from each partition's **Total LCP** count. Use the line at the bottom of the window for this purpose (disabled if there are no parked LCPs).

When partitions are transferred into **zPCR**, parked LCP values are rounded down to the next integer before being subtracted from **Total LCPs**.

When partitions are transferred into **zPCR**, parked LCP values are rounded down to the next integer before being subtracted from **Total LCPs**.

Once loaded into **zPCR**, the values will be identified as **Active LCPs**. Active LCP values can always be changed on the **Partition Detail Report** window. Unparked LCP information is carried forward into **zPCR** during the invocation (this information is not retained in a study file). Parked LCPs can subsequently be removed by clicking the **LCP Alternatives** button. For more information, see [LCP Alternatives](#).

Note that the original **Total LCPs** partition information is not retained in a study file. To obtain that information, the EDF would have to be input again.

CPU MF and Workload Category Determination

When **CPU MF** data is available for a partition, information for the following items are displayed in the **CPU MF** column.

- **Workload Assignment:** Displays the z/OS or z/VM workload category that would be selected based on **CPU MF** data.

The **Assigned Workload** for each partition (under **Partition Identification**) is initialized to one of the **Workload Choices**, in the priority listed below:

1. **CPU MF** workload choice (z/OS and z/VM partitions only). **When available, this workload category should be considered as the most reliable choice for the partition.** For the **CPU MF** method, the **RNI** and internal algorithms specific to each processor family are used to determine the **Workload Choice**.
2. **Default** workload choice (by SCP)

• z/OS	(General Purpose, zAAP, & zIIP)	Average
• z/VM	(General Purpose and/or IFL)	Average/LV
• z/VSE	(General Purpose only)	Average/VS
• KVM	(General Purpose or IFL)	Average/K
• Linux	(General Purpose or IFL)	Average/L
• CFCC	(General Purpose or ICF)	CFCC (cannot be changed)

When CPU MF data is not available, the default workload assignment for the SCP should generally be used.

3. **User Modified** from one of the above assignment.

The **SCP** for IFL partitions may be altered to either **z/VM**, **KVM**, **Linux**, **zAware**, **zACI**, or **SSC**. If the **SCP** for a GP partition needs to be changed, it must be done after the configuration has been transferred into **zPCR**.

The **SCP** and **Assigned Workload** for zAAP and zIIP partitions are always derived from the parent GP partition.

The **Assigned Workload** field for a **z/OS**, **z/VM**, **z/VSE**, **KVM**, or **Linux** partition may be altered as desired. The **Assigned Workload** categories available will depend on the SCP specified.

CPU MF workload assignment considerations

A single EDF interval is used to define partitioning metrics to **zPCR**. The workload selected is also based on that same interval (only **Low**, **Average**, or **High** will ever be selected). Since the actual workload characteristics may vary between intervals, it might be useful to load the configuration from several intervals before and after the desired one to see if the selected workload does vary. If so, a judgement could be made about the actual workload that should be used (the **Low-Avg** or **Avg-High** workloads could also be considered).

Workload Category Determination

The **Method Used** column (on the right side of the window) reveals the method applied for **Assigned Workload (CPU MF or Default)**. If the workload is manually changed, this indicator will be set to “**Other**”.

Transferring Configuration Information into zPCR

To enable partitions for transfer to **zPCR**, click each of the desired partition's **Copy LP** checkbox. Note that zAAP and zIIP partitions are always transferred with their parent GP partition. Use the **Select All** button to check all partitions and the **Remove All** button to uncheck all partitions. Use the **Select Active** button to check only those partitions that were active in the interval.

Once all the desired partitions are checked, when this function was entered via:

- [**Create Host and Partitions From EDF**](#), click the **Create zPCR Configuration** button to create the entire LPAR configuration, including the LPAR host and the selected partitions. Partition utilization values will also be transferred, thus making the **Utilized Capacity Report** available in **zPCR**. The description field on the **LPAR Host and Partition Configuration** window will be initialized to indicate the source.
When an entire LPAR configuration is created from EDF, the **LPAR configuration name** is automatically generated. It may be renamed if desired.
- [**Copy Partitions From EDF**](#), click the **Copy Partitions** button to add the selected partition definitions to the currently active LPAR configuration.

Note. If any partition's LCPs exceed the size of the host's largest drawer, a dialog will appear with the partition names. An exception check box is provided to allow those partitions to be automatically included when the configuration is ported into **zPCR**. If not checked, each of these partitions need to be manually included later.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

Note: To recognize SRB activity, the latest version of CP3KEXTR (v4.06 or higher) must be used to generate the EDF.

For detail concerning SRB support, see [System Recovery Boost](#).

If **Estimate Parked LCPs where unknown for** **GP** or **IFL** partitions, the estimated parked LCPs are added to the partition's total LCP count (fractional parked LCPs are always rounded up). Otherwise, the defined LCP count is transferred to the LPAR configuration.

In **zPCR**, transferred partitions can be viewed in the **Partition Detail Report** window. If a partition definition conflicts with the overall configuration, its **Include** checkbox will be unchecked. When you attempt to check **Include**, the partition's definition is validated; if invalid, **Include** will remain unchecked. In this case you must make corrections to the partition definition or to the LPAR host definition before the partition can be included. Most partition definition metrics can be modified from these windows (also the **Partition Definition** window).

Window Controls

Click the **Select All** button to include all of the partitions.

Click the **Select Active** button to include only the active partitions.

Click the **Remove All** button to exclude all of the partitions.

Click the **Choose Another EDF Interval** button to select another EDF interval or to load another EDF report.

To estimate parked LCPs for GP (including any associated zAAP or zIIP) partitions or IFL partitions use the checkboxes in the line: **Estimate parked LCPs where unknown for:** **GP partitions** **IFL partitions**. Estimated parked values will appear in red to differentiate from actual parked.

Use the **When copying partitions into zPCR, remove Parked LCPs from the LCP Count** checkbox to have all active and estimated parked LCPs when the LPAR configuration is transferred to **zPCR**.

Depending on how this window was entered:

- Click the **Create LPAR Configuration** button to transfer the LPAR host and all the selected partitions to **zPCR**.
- Click the **Copy Partitions**, to transfer all the selected partitions to the currently defined LPAR host in **zPCR**.

Note that, if the selected interval contains System Recover Boost (SRB) activity, these buttons are disabled, and the LPAR configuration cannot be transferred into **zPCR**.

Click the **Cancel** toolbar icon to return to the **EDF Interval Selection** window. A 2nd **Cancel** will return to **zPCR**.

Partition Utilization Values

When an LPAR configuration is obtained from EDF, utilization values of each partition for the selected interval are also transferred into **zPCR**. These can be used to determine "utilized" capacity, which can then be compared to the capacity that would be allocated based on the partition weights.

To make this comparison, the ***Utilized Capacity Report*** window can be accessed from the ***LPAR Host and Partition Configuration*** window, using the **(Partition Utilized Capacity)** button. Note that this button is enabled only when the LPAR host and its entire partition configuration were transferred and only if no LPAR host or partition configuration changes have been made. For more information, see [Utilized Capacity Report](#).

The EDF data represents a close approximation of utilized partition capacity. Some minor deviation from the capacity actually consumed may exist because:

1. What EDF views as ***LPAR Management Time*** (i.e., Physical) is not exactly equivalent to what **zPCR** views as ***LPAR Management Time***.
2. **zPCR** algorithms assume the ***Capacity Perspective*** (i.e., LPAR costs are based on full utilization of the CPC's CP resources). Since actual utilization for the CP pools may be less than full utilization, actual LPAR costs may be less than that reported by **zPCR**.

Note: CFCC in a dedicated partition always runs at 100% utilization. For such partitions, the actual capacity being consumed cannot be determined as for the other partition types.

RMF Overview

LPAR Configuration Capacity Planning

An **RMF Report** can be used as input to **zPCR** for the purpose of creating the entire LPAR host and partition configuration or copying individual partition definitions into a currently active LPAR configuration. An RMF report is created from a partition running z/OS. It must include the following:

1. **Partition Data Report (PDR)**
2. At least one **CPU Activity Report**, where the LPAR host processor model is identified. When a **CPU Activity Report** is included for more than one partition, each will be processed, providing additional detail for those partitions.

All other RMF reports are ignored.

RMF data generally represents multiple time intervals. One specific interval must be selected as input to **zPCR**.

The following **RMF Report** information is used by **zPCR**.

LPAR Configuration Metrics Determined from RMF	
The CPU Activity Report provides the following information	
LPAR Host	Selected Interval (Date, Time, and Duration)
	System ID
	Specific processor family and model
The Partition Data Report provides the following information	
LPAR Host	RCP count for: GP, zAAP, zIIP, IFL, ICF
	Name
	Type: GP, zAAP, zIIP, IFL, ICF
	Mode: Dedicated or Shared
	LCP count
	Weight
	Capping
	Absolute Capping (Indicator without a value)
	Utilization
	HMC activated but not IPL'd
Individual partition CPU Activity Reports provide the following information	
Specific z/OS version running	
HiperDispatch: Active and number of parked LCPs	
Measured SMT Benefit: IBM z17, z16, z15, z14, z13 only; zIIP/IFL partitions only	
Where indeterminate, the following assumptions apply. If an actual assignment differs, it can be changed after the configuration is loaded into zPCR.	
SCP	GP/zAAP/zIIP: z/OS
	IFL: z/VM
	ICF: CFCC
Workload	Average category

RMF is transferred via a traditional RMF report in text format. The RMF report must be available as a flat file, downloaded to your PC (preferred file extension = **.rmf** or **.txt**). The file may either contain or exclude carriage-control characters in the 1st byte position of each record.

Note that the RMF report in the newer XML format is not supported.

RMF reports are read in via the **RMF Interval Selection** window (see [RMF Interval Selection](#)). Once processed, each measurement interval found is listed. A single measurement interval must be selected from which to obtain the LPAR configuration information. Partition specific information (i.e., **HiperDispatch** and **SMT Benefit**) is obtained from each individual partition **CPU Activity Report** that was made available.

For z/OS partitions where HiperDispatch is active (z10 and later), some General Purpose LCPs (and any associated zAAP, zIIP, or IFL LCPs) may be parked. The number of parked LCPs is computed and thus noted before the partition is transferred into **zPCR**. Parked LCPs should be removed from the partition LCP definition to more accurately relate capacity. Parked LCPs can automatically be removed from partitions when the configuration is transferred into **zPCR**.

On systems where IRD is actively managing the partition resources, partition LCPs are reported by RMF as fractional values. For **zPCR** purposes, fractional LCP numbers are rounded up to the nearest whole number. Since weights can also be dynamically modified by IRD, it is important that the interval selected for analysis be one that is representative for a load period of particular interest.

Capping information from RMF

In 2016, a PTF for RMF became available which identifies 3 types of capping, any of which may be assigned to a partition:

1. Initial Capping	Supported by zPCR
2. Absolute Capping	Supported by zPCR; requires manual input
3. Group Absolute Capping	<u>Not supported</u> by zPCR

In **zPCR**, partition capping will be set for only the 1st two types.

When RMF is generated with the PTF applied, **zPCR** will identify a partition's capping type as one of the above. At the time the RMF is transferred into **zPCR**, a dialog will list each capped partition and its capping type. Use this dialog to identify the "Absolute Capped" partitions. Absolute capping values are not available from RMF; they must be obtained externally and the capping value must be manually entered into **zPCR** on the **Partition Detail Report** window.

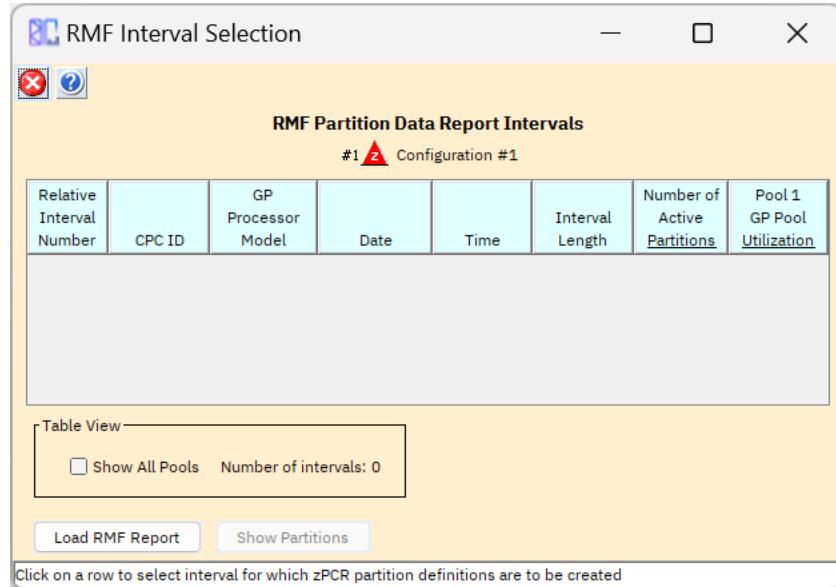
When RMF is generated without the PTF, **zPCR** can only identify partitions as being "Initial capped". At the time the RMF is transferred into **zPCR**, a dialog will list each capped partition. Any information concerning "Absolute Capped" partitions must be obtained externally and the capping value must be manually entered into **zPCR** on the **Partition Detail Report** window.

With z/OS V2R4, some RMF fields are formatted differently. Starting with **zPCR** version 9.3a, this change is supported.

Note: The filename assigned to an RMF should not contain any special characters. Otherwise, problems may occur when transferring the configuration into **zPCR**.

RMF Interval Selection

LPAR Configuration Capacity Planning



From the **LPAR Host and Partition Configuration** window, click the **RMF** button in either the **Define LPAR Host Processor** group box or the **Define Partitions** group box to open the **RMF Interval Selection** window.

The **RMF Interval Selection** window provides the means to load an RMF report file (for general information concerning RMFs, see [RMF Overview](#)).

Click the **Load RMF Report** button. A file dialog box will appear from which you can specify the drive:\directory\filename of the desired RMF report. The input must be a traditional RMF report (text format, file extension assumed is “**rmf**” or “**txt**”). Traditional RMF reports must have a record length of 133 bytes, or if Carriage Control information is included, 134 bytes.

When read, the RMF file is searched for all report intervals. No RMF intervals are shown and a message is issued if the RMF report file:

- Does not appear to be a standard RMF report
- Was generated on a z/OS version not supported by **zPCR**
- Does not include a **Partition Data Report (PDR)** and a **CPU Activity Report**
- Appears to be invalid

Once validated, the RMF report is loaded. It will remain available until a different RMF report is loaded.

RMF Interval Selection

RMF Partition Data Report Intervals
#1  Configuration #1
RMF Report File: C:\Users\1004_1897\Documents\zPCR Defaults\RMF Files\RMFSample z10 E56.rmf

Relative Interval Number	CPC ID	GP Processor Model	Date	Time	Interval Length	Number of Active Partitions	Pool 1 GP Pool Utilization	Pool 2 zAAP Pool Utilization	Pool 3 zIIP Pool Utilization	Pool 4 IFL Pool Utilization	Pool 5 ICF Pool Utilization
1.	SYSA	2097-715	11/10/2008	07.59.00	000.59.59	9	53.97%	18.70%	0.00%	0.00%	0.00%
2.	SYSA	2097-715	11/11/2008	07.59.00	001.00.00	9	15.32%	10.04%	0.00%	0.00%	0.00%
3.	SYSA	2097-715	11/12/2008	07.59.00	001.00.00	9	60.03%	19.69%	0.00%	0.00%	0.00%
4.	SYSA	2097-715	11/13/2008	07.59.00	001.00.00	9	67.49%	19.82%	0.00%	0.00%	0.00%
5.	SYSA	2097-715	11/14/2008	07.59.00	<u>001.00.00</u>	9	84.45%	19.81%	0.00%	0.00%	0.00%
6.	SYSA	2097-715	11/17/2008	07.59.00	000.59.59	9	72.03%	20.06%	0.00%	0.00%	0.00%
7.	SYSA	2097-715	11/18/2008	07.59.00	001.00.00	9	67.00%	19.42%	0.00%	0.00%	0.00%
8.	SYSA	2097-715	11/19/2008	07.59.00	001.00.00	9	68.72%	19.53%	0.00%	0.00%	0.00%

Table View
 Show All Pools Number of intervals: 13

Load RMF Report Show Partitions

Click on a row to select interval for which zPCR partition definitions are to be created

Once an RMF report is successfully loaded, a title line reveals the location and filename of that report.

Each RMF interval found in the report is displayed in a sequentially numbered row. Interval descriptions include **System ID**, **GP Processor Model**, **Date**, **Time**, **Interval Length**, **Number of Active Partitions**, and **Utilization** within each CP pool. By default, only the utilization column for the General Purpose pool appears (normally the pool of greatest interest). The **Show All Pools** checkbox can be used to reveal utilization values for the remaining CP Pools.

To facilitate identification of report intervals where the greatest amount of CPU resource is consumed, several columns (with underlined headings) are enabled for sorting. Click on the heading once to sort in descending sequence, again to sort in ascending sequence, and once again to restore the original order.

From this table, select the report interval for which LPAR information is to be captured by clicking on the desired interval. The selected row will be highlighted. Click the **Show Partitions** button to go to the **Create LPAR Configuration from RMF** or **Copy Partitions from RMF** window where the individual partition definitions can be reviewed and selected for transfer to the **zPCR** LPAR configuration.

RMF cannot identify the SCP/workload assigned to a partition in terms that **zPCR** uses. Therefore, SCP/workload assignments are made using default settings. These SCP/workload assignments can be subsequently changed in **zPCR**, on the **Define Partitions** window or the **Partition Detail Report** window.

System Recovery Boost Considerations

Intervals that include any partition running with **System Recovery Boost** activity will be highlighted with a red background.

Note: To recognize SRB activity, the RMF CPU Activity Report must be generated by z/OS V2R3 or later.

For detail concerning SRB support, see [System Recovery Boost](#)

Window Controls

Click the **Load RMF Report** button to select the specific RMF file to be loaded.

Click the **Show Partitions** button to reveal the partitions that were active during the selected interval. The window discussed in [Get Partitions from RMF](#) will appear.

Click the **Cancel** toolbar icon to close the window, return to the **LPAR Host and Partition Configuration** window with no changes applied.

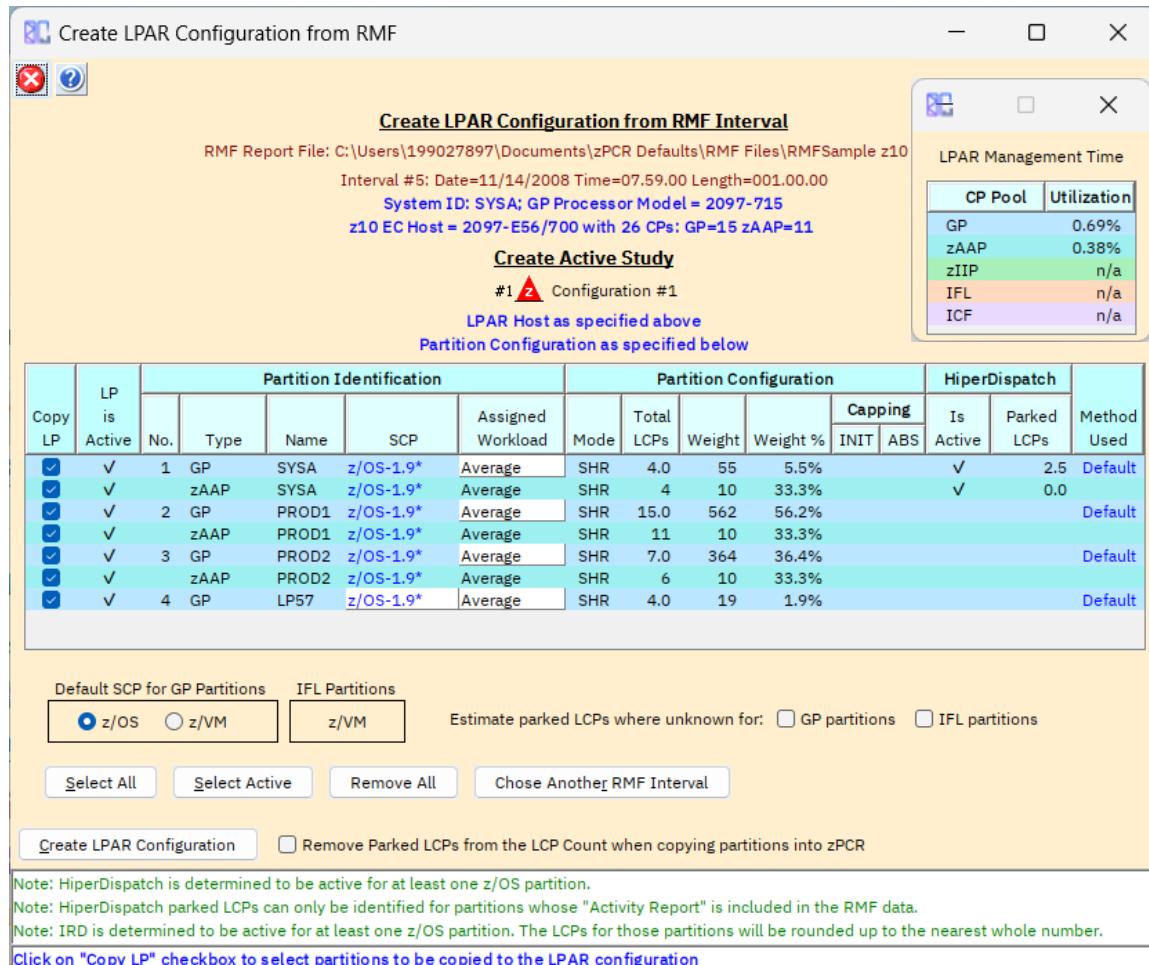
Get Partitions from RMF

LPAR Configuration Capacity Planning

From the **RMF Interval Selection** window, click on an interval row to enable the **Show Partitions** button. Clicking the **Show Partitions** button will:

- When entered via [Create Host and Partitions From RMF](#), the **Create LPAR Configuration from RMF** window appears.
- When entered via [Copy Partitions From RMF](#), the **Copy Partitions from RMF** window appears.

Except for the window titles and some button names, these two windows are identical.



In addition to the primary window, a small **LPMT** window appears showing the LPAR management time for each CP pool as reported by RMF (informational purposes only).

At the top of the window, title lines reveal the location and filename of the currently loaded RMF report, and the interval selected. Just below is listed the LPAR host system ID and its General Purpose processor model as noted by RMF. The next line identifies the overall LPAR host configuration as would be shown in **zPCR**.

When the window title reads **Create LPAR Configuration from RMF**, the lines below reveal, what will become the host processor for the LPAR configuration. When the window title bar reads **Copy Partitions from RMF**, the currently defined LPAR host is simply identified. In either case, all the partitions that are selected in the table below will be transferred into the LPAR configuration.

Partition Utilization Considerations

Partition utilizations, while known, are not displayed on the **LPAR Configuration from RMF** window. However, they will be transferred into **zPCR** and will be included in a saved study. When the study for the original configuration is reloaded, the partition utilization values can be viewed in the **Utilized Capacity Report** window. In addition, the actual zAAP and zIIP partition utilizations will be available for assignment on the **zAAP/zIIP Loading** window.

SMT Benefit

For z13 and later processors, measured **SMT Benefit** values can be obtained from RMF for SMT enabled zIIP and IFL partitions. An additional column, **SMT Benefit**, is displayed immediately following the **Capping** columns with this information. When SMT information is found, it will be transferred to **zPCR** as a **Measured SMT Benefit**.

When a zIIP or IFL partition's utilization is less than 20% or when the measured SMT Benefit is greater than 90%, the SMT Benefit value is considered as being unreliable. While transferring the configuration into **zPCR**, a dialog will identify such cases and the default **Estimated SMT** value will be assigned instead. While viewing the **SMT Benefit** window, the ability to use the **Measured SMT** benefit is provided.

Partition Identification and Partition Configuration

The primary table contains a list of all the partitions found in the selected RMF interval with their definition metrics. From this table you select the specific partitions to be copied.

All active partitions found in the selected RMF interval are listed in the table, by pool, in the following order:

1. GP (General Purpose)
2. zAAP
3. zIIP
4. IFL
5. ICF

Note: zAAP, zIIP, IFL, and ICF LCPs that are associated are listed just below the parent GP partition.

The format of each row is similar to that used in **zPCR**, with background colors differentiating each of the partition pools. Partition weights are summed by pool to determine the weight percent.

SCP Assignment

The **SCP** field is populated with the z/OS SCP version found in the partition's **RMF CPU Activity Report**, when available. Otherwise, the field is populated with a default SCP for the partition type as follows:

- All remaining GP partitions will be assigned the default **z/OS** version for the LPAR host. When another SCP is actually in control, the SCP will need to be changed manually.
- Any remaining IFL partitions will be assigned the default **z/VM** version for the LPAR host. When another SCP is actually in control, the SCP will need to be changed manually.
- For **zAAP**, **zIIP**, **IFL**, and **ICF** partitions that are associated with a parent z/OS or z/VM GP partition, the SCP of the parent is assigned for each.
- If the SCP settings for GP and IFL partitions are corrected before the configuration is transferred into **zPCR**, **Parked LCPs** will automatically be corrected to only represent those partitions running z/OS or z/VM. If these changes are to be made in **zPCR**, GP **Parked LCPs** should not be estimated.

Partitions that are HMC activated, but have no SCP IPL'd (i.e., not running) will have the SCP set to "**Not IPL'd**". Logical CPs will remain in order to accurately represent the distribution of LCPs across the RCPs, and the assigned weights will be applied across the CP pool. In **zPCR**, the partition's minimum and maximum capacity values are assumed as zero (no value is displayed).

Note: **Every partition detected with RMF is HMC activated.**

All partitions found in RMF should be included for reliable capacity results.

Workload

For RMF input a partition's workload assignment will always default to the **Average** category for the SCP. The workload may be changed to any of the SCP's available categories. For **ICF** partitions, **CFCC** is always assigned.

Mode (Dedicated or Shared)

Total LCPs

Shows the total number of Logical CPs assigned to the partition.

Weight and Weight %

Shows the partition weight and calculated weight percentage relative to the other partitions in the CP pool.

Initial Capping

If a partition was capped, a checkmark will appear in the **INIT** column.

Absolute Capping (z12 and later processors only)

If **Absolute Capping** is defined to a partition, in addition to the checkmark, a value will appear in the **ABS** column.

SMT Benefit for zIIP and IFL Partitions

On z13 and later processors, ***SMT Benefit*** is obtained from EDF or RMF for SMT enabled zIIP partitions running z/OS and for IFL partitions running z/VM, KVM, or Linux. An additional column, ***SMT Benefit***, is displayed with this information immediately following the ***Capping*** columns.

The SMT benefit is considered reliable when the partition's utilization is 20% or greater. In this case the benefit will be assigned as a ***Measured SMT Benefit***. When the partition's utilization is less than 20%, the default ***Estimated SMT Benefit*** value will be assigned instead.

There are other scenarios where the measured SMT Benefit value is limited to a maximum value. In this case the value will be assigned as a ***Measured SMT Benefit***. When the value is zero or less, the ***Measured SMT Benefit*** will be set to 0%. SMT can be manually disabled if desired.

HiperDispatch (z/OS and z/VM partitions only)

The following are reported only for partitions where a ***CPU Activity Report*** is included.

- **HD Active:** A checkmark indicates whether or not HiperDispatch was active for the partition.
- **Parked LCPs:** Indicates the number of the GP partitions total LCPs that were parked during the measurement interval. Parked LCPs should be removed from the partition LCP count to more accurately reflect capacity. Use the ***Remove Parked LCPs from Partition LCP Count*** checkbox to have all parked LCPs removed from partitions when the LPAR configuration is transferred to **[zPCR](#)**.

Note that parked LCP information is carried forward into **[zPCR](#)** (this information is not retained in a study file). Parked LCPs can subsequently be removed from these partitions from the ***Partition Detail Report*** window, by clicking the **LCP Alternatives** button. For more information, see [LCP Alternatives](#).

HiperDispatch (z10 and later; running z/OS)

The following are reported only for partitions where a ***CPU Activity Report*** is included.

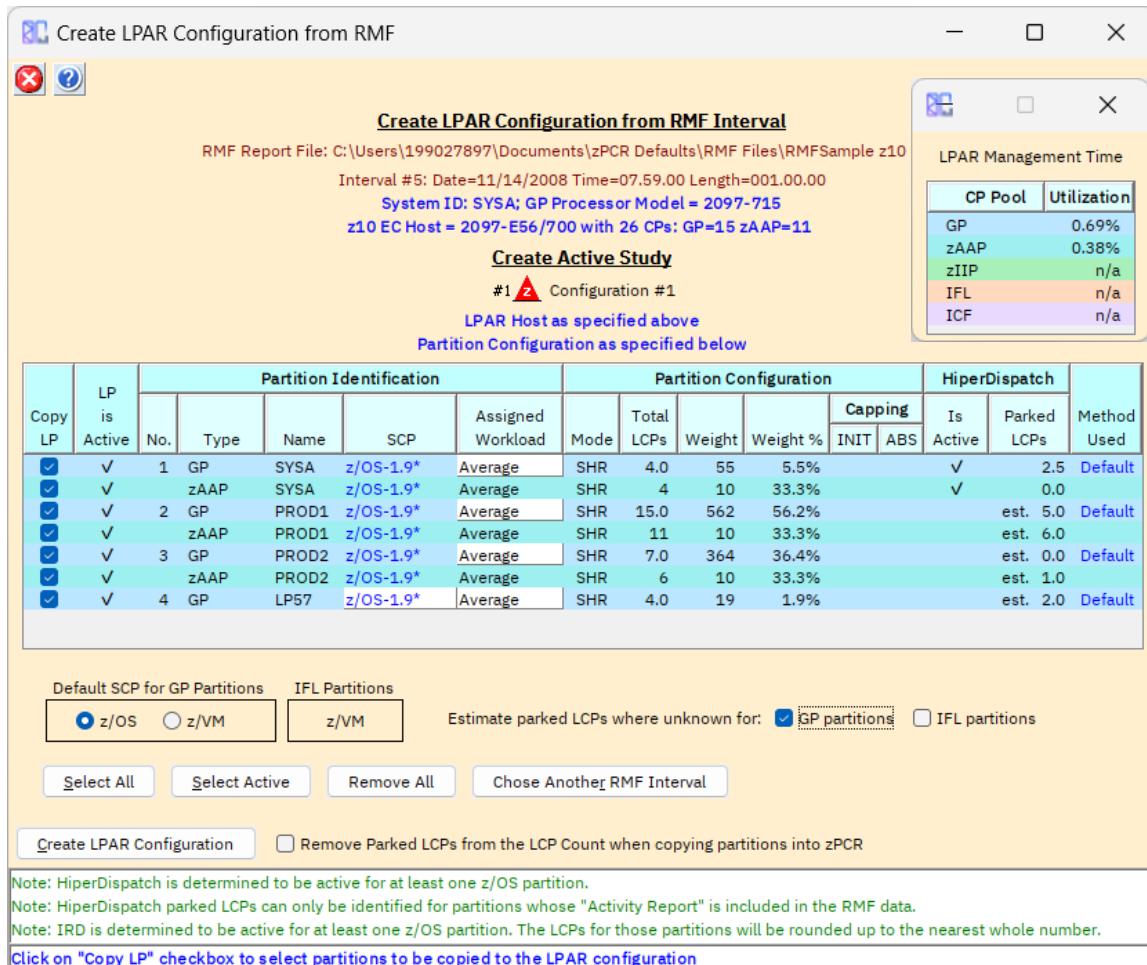
- **HiperDispatch Active:** A checkmark indicates whether or not HiperDispatch was active for the GP partition.
- **Parked LCPs:** Indicates the actual number of the GP partition's LCPs that were parked during the measurement interval. The Parked LCP count will display with 1 decimal place, since parking activity is dynamic across the selected interval.

Parked LCPs may be estimated for the remaining z/OS and z/VM partitions.

Estimating Parked LCPs

If the LPAR host supports Hiperdispatch (z10 and later), parked LCPs can be estimated for z/OS and z/VM shared partitions that are not represented by the RMF input. This also applies to associated zAAP or zIIP partitions. Parked LCPs can be estimated for IFL shared partitions not represented by an RMF and running z/VM. The Parked LCP estimate will display with one decimal place.

When the **Create/Copy from RMF** window is opened, a dialog is displayed to help emphasize the ability to estimate parked LCPs. The dialog can remain open while making changes to the window.



Parked LCP estimates are controlled using the check boxes displayed below the J-table. When GP or IFL is checked, a parked LCP estimate of will be made (1 decimal precision) for the non-RMF shared partitions in that pool. Estimated values are prefixed with "Est." to differentiate them from actual parked values. Unchecking the box will remove the parked LCP estimates.

Notes:

- Each partition's **Copy LP** checkbox defaults to its **LP is Active** checkbox. If **Copy LP** is changed, all **Parked LCP** estimates for the CP pool will be recomputed.

- Changes to any partition's SCP other than z/OS or z/VM will cause **Parked LCPs** estimates for the CP pool will be recomputed. If such changes are to be made after transfer to **zPCR**, **Parked LCPs** should not be estimated, since the estimates would no longer be relevant.

To accurately reflect partitioned capacity, **Parked LCPs** should be removed from each partition's **Total LCP** count. Use the line at the bottom of the window for this purpose (disabled if there are no parked LCPs).

When partitions are transferred into **zPCR**, parked LCP values are rounded down to the next integer before being subtracted from **Total LCPs**.

When partitions are transferred into **zPCR**, parked LCP values are rounded down to the next integer before being subtracted from **Total LCPs**.

Once loaded into **zPCR**, the values will be identified as **Active LCPs**. Active LCP values can always be changed on the **Partition Detail Report** window. Unparked LCP information is carried forward into **zPCR** during the invocation (this information is not retained in a study file). Parked LCPs can subsequently be removed by clicking the **LCP Alternatives** button. For more information, see [LCP Alternatives](#).

Note that the original **Total LCPs** partition information is not retained in a study file. To obtain that information, the RMF would have to be input again.

Workload Category Determination

The **Method Used** column (on the right side of the window) reveals the method applied for **Assigned Workload (Default or Other)**. If the workload is manually changed, this indicator will change accordingly.

zAAP and zIIP partitions will always be assigned the same workload as the parent GP partition. IFL partitions associated with a parent GP partition may have independent workload category assignments. ICF partitions associated with a parent GP partition will have a CFCC as the assigned workload.

Transferring Configuration Information into zPCR

To enable partitions for transfer to **zPCR**, click each of the desired partition's **Copy LP** checkbox. Note that zAAP and zIIP partitions are always transferred with their parent GP partition. Use the **Select All** button to check all partitions and the **Remove All** button to uncheck all partitions. Use the **Select Active** button to check only those partitions that were active in the interval.

Once all the desired partitions are checked, when this function was entered via:

- [Create Host and Partitions From RMF](#), click the **Create zPCR Configuration** button to create the entire LPAR configuration, including the LPAR host and the selected partitions. Partition utilization values will also be transferred, thus making the **Utilized Capacity Report** available in **zPCR**. The description field on the **LPAR Host and Partition Configuration** window will be initialized to indicate the source.
When an entire LPAR configuration is created from RMF, the **Partition Name** is automatically generated. It may be manually renamed if desired.
- [Copy Partitions From RMF](#), click the **Copy Partitions** button to add the selected partition definitions to the currently active LPAR configuration.

Note. If any partition's LCPs exceed the size of the host's largest drawer, a dialog will appear with the partition names. An exception check box is provided to allow those partitions to be automatically included when the configuration is ported into **zPCR**. If not checked, each of these partitions need to be manually included later.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

Note: To recognize SRB activity, the RMF CPU Activity Report must be generated by z/OS V2R3 or later.

For detail concerning SRB support, see [System Recovery Boost](#).

In **zPCR**, partitions can be viewed in the **Partition Detail Report** window. If a partition definition conflicts with the overall configuration, its **Include** checkbox will be unchecked. When you attempt to check **Include**, the partition's definition is validated; if invalid, **Include** will remain unchecked. In this case you must make corrections to the partition definition or to the LPAR host definition before the partition can be included. Most partition definition metrics can be modified from these windows (also the **Partition Definition** window).

Window Controls

Click the **Select All** button to include all of the partitions.

Click the **Select Active** button to include only the active partitions.

Click the **Remove All** button to exclude all of the partitions.

Click the **Choose a Different RMF Interval** button to select another RMF interval or to load another RMF report.

Use the **When copying partitions into zPCR, remove Parked LCPs from the LCP Count** checkbox to have all parked LCPs that have been identified removed from partitions when the LPAR configuration is transferred to **zPCR**.

Depending on how this window was entered:

- Click the **Create LPAR Configuration** button to transfer the LPAR host and all the selected partitions to **zPCR**.
- Click the **Copy Partitions**, to transfer all the selected partitions to the currently defined LPAR host in **zPCR**.

Click the **Cancel** toolbar icon to return to the **RMF Interval Selection** window. A 2nd **Cancel** will return to **zPCR**.

Partition Utilization Values

When both the LPAR host and its entire partition configuration are obtained from RMF, utilization values of each partition for the selected interval are also transferred to **zPCR**. These can be used to determine “utilized” capacity, which can then be compared to the capacity that would be allocated based on the partition weights.

To make this comparison, the ***Utilized Capacity Report*** window can be accessed from the ***LPAR Host and Partition Configuration*** window, using the **(Partition Utilized Capacity)** button. Note that this button is enabled only when the LPAR host and its entire partition configuration were transferred and only if no LPAR host or partition configuration changes have been made. For more information, see [Utilized Capacity Report](#).

The RMF data represents a close approximation of utilized partition capacity. Some minor deviation from the capacity actually consumed may exist because:

1. What RMF views as ***LPAR Management Time*** (i.e., Physical) is not exactly equivalent to what **zPCR** views as ***LPAR Management Time***.
2. **zPCR** algorithms assume the ***Capacity Perspective*** (i.e., LPAR costs are based on full utilization of the CPC's CP resources). Since actual utilization for the CP pools may be less than full utilization, actual LPAR costs may be less than that reported by **zPCR**.

Note: CFCC in a dedicated partition always runs at 100% utilization. For such partitions, the actual capacity being consumed cannot be determined as done for the other partition types.

Create Host and Partitions from Study

Copy Partitions from Study

LPAR Configuration Capacity Planning

This capability is provided via the **LPAR Host and Partition Configuration** window.

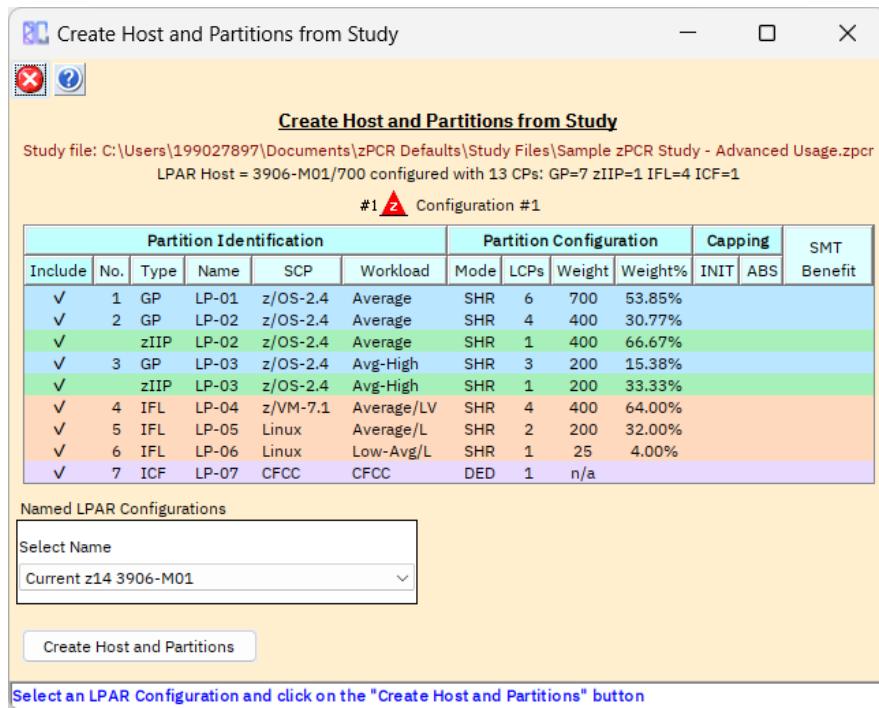
1. To define an entire LPAR configuration from a study: In the **Define LPAR Host Processor** group box, click on the **zPCR Study** button. This button is disabled if the LPAR host has already been defined.

The **Create Host and Partitions from Study** window will open.

2. To copy Partitions from a study: In the **Define Partitions** group box, click on the **zPCR Study** button.

The **Copy Partitions from Study** window will open.

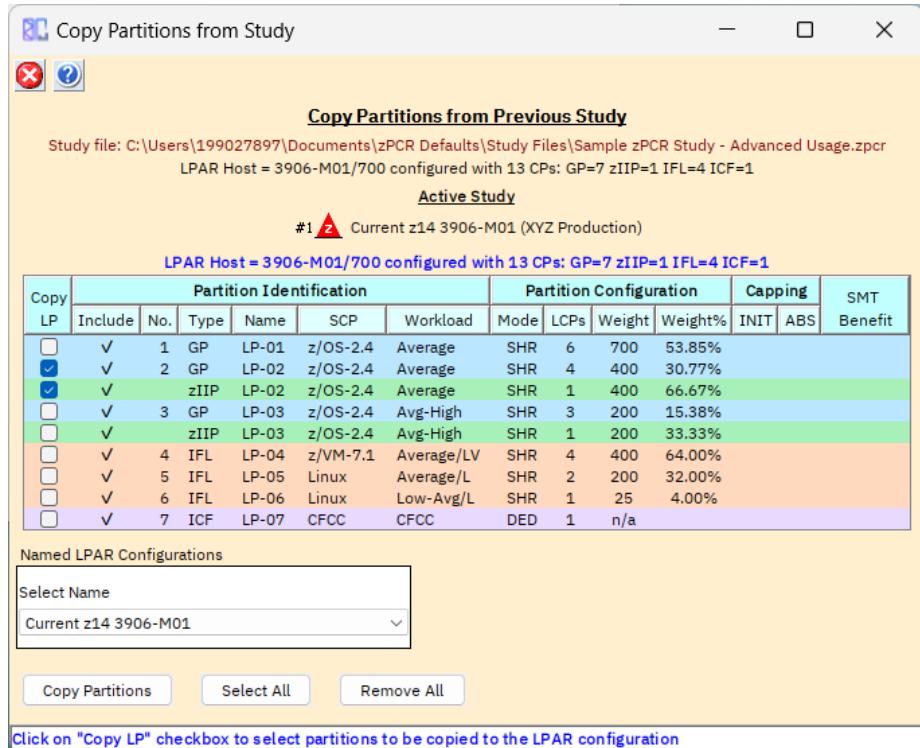
The **Load Study** dialog will appear from which you can select the appropriate drive:\directory\filename. Note that only studies made with recent **zPCR** versions are able to be loaded (a dialog will appear if the version is no longer supported). Once a valid study has been loaded, the **Copy Partitions from Study** window will appear.



The **Create Host and Partitions from Study** window can only be used to transfer an entire LPAR configuration into **zPCR**.

The target study file may contain multiple LPAR configurations. In this case, one of those configurations must be selected from a drop down list. That configuration will become the object from which information is transferred.

The **Copy Partitions from Study** window provides the capability to select the specific partitions to be copied using the check boxes provided.



The file ID of the previous study is displayed followed its hardware configuration. For the active study, the hardware configuration and **Reference-CPU** capacity basis are provided.

Transferring Configuration Information to zPCR

Partitions are presented in a table format similar to the **Partition Detail Report** window. Partition definitions cannot be modified from this window.

A **Copy LP** checkbox is presented for each partition in the first column, defaulting to unchecked status. To copy one or more partitions to the active study, check each of those desired. A **Select All** button and a **Remove All** button are available to assist in controlling partition selection.

Once a partition has been checked, the **Copy Partitions** button is enabled. Clicking this button will copy each of the selected partitions into zPCR's current LPAR configuration, and the **Copy Partitions from Study** window is closed. A dialog box will appear with information about the number of partitions copied and their status. Since some of the imported partition's definition metrics may be invalid for the currently defined LPAR host (i.e., number of LCPs defined or references to hardware features that are not configured such as zAAPs, zIIPs, IFLs or ICFs), all partition **Include** checkboxes will initially be unchecked.

For GP partitions with associated zAAP, zIIP, IFL, or ICF partitions, the associated partitions will be included in the copy. However, in zPCR, if either the required hardware or the required software is not specified, neither the General Purpose nor the associated partition's **Include** checkbox will be able to be checked.

Partitions copied from a previous study will appear in the **Partition Definition** window (**GP**, **IFL**, and **ICF**) and the **Partition Detail Report** window with the **Include** checkbox unchecked. When you attempt to check **Include**, the partition's definition is validated; if invalid, **Include** will remain unchecked. In this case you must make corrections to the partition definition or to the LPAR host definition before the partition can be included. Most partition definition metrics can be modified from either of these windows. However, changes to the number of zAAP or zIIP LCPs, and IFLs associated with z/VM, must be made from the **Define Partitions** window.

Capping

If a partition was **Initial Capped**, a checkmark will appear in the **INIT** column.

If **Absolute Capping** is defined to a partition, a value will appear in the **ABS** column (z12 and later processors only). If copying a partition with **Absolute Capping** specified to an LPAR host that does not support the feature, the values will be discarded.

SMT Benefit

If a z13 or later study was saved with zIIP or IFL **SMT Benefit** values applied, an **SMT Benefit** column will appear with all measured and estimated values that have been applied. If copying a partition that has **SMT Benefit** values specified to an LPAR host that does not support SMT, these values will be discarded.

Click the **Cancel** toolbar icon to close the **Copy Partitions from Study** window without transferring any partitions to the active study.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** partition was active, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

zAAP/zIIP Capacity Considerations

LPAR Configuration Capacity Planning

zAAP CPs and zIIP CPs, purchased as a feature code, are configured to the LPAR host like IFLs and ICFs. They are used in the form of zAAP or zIIP partitions, associated with a parent z/OS GP partition. zAAP/zIIP associations are considered as unique partitions in **zPCR** for the purposes of projecting capacity. Each of these partition types will be defined identically to its parent z/OS partition for the following items:

- Name
- SCP
- Workload
- Mode

The following items can be uniquely set.

- Weight
- Initial CAP
- ABS CAP

Several questions concerning potential zAAP or zIIP capacity can be answered using **zPCR**

1. What is the capacity expectation for the parent GP partition when also managing its associated zAAP or zIIP partition?
2. What capacity expectation for the associated zAAP or zIIP partition?
3. If the capacity requirement for a parent z/OS GP partition is reduced due to the addition of an associated zAAP or zIIP partition, how many LCPs might still be required?
4. How would configuration changes to a parent GP partition and/or its associated zAAP or zIIP partition affect the overall capacity perspective of the LPAR host?

zAAP (Application Assist Processor)

zAAP CPs may be configured on z9 through z12 processor models (they are not available on z13 and later). When running Java applications under z/OS-1.6 or later, zAAP LCPs may be used in lieu of GP LCPs to run certain Java code. This has the effect of relieving the GP LCPs assigned to the z/OS partition of some of the load, potentially reducing software licensing costs.

For zAAP capacity planning purposes, the workload of a z/OS image can be considered as having three components.

- **Normal content** is the percent of workload that has no Java content (i.e., no zAAP eligible work). When no zAAP LCPs are defined, the entire workload is considered to be normal content.
- **Java Content** is the percent of workload that contains zAAP eligible work (e.g., WebSphere).
- **zAAP Eligible Work** is the portion of Java Content that can actually run on zAAP LCPs. This is the portion of the workload to consider when determining how much zAAP capacity could actually be consumed.

zAAP partition capacity is always associated with that of the Java eligible content.

zIIP (Integrated Information Processor)

zIIP CPs may be configured on all z10 and later processor models. When running certain workloads under z/OS-1.6 or later (Db2 is the prime example), zIIP LCPs may be used in lieu of GP LCPs to run portions of the work. This has the effect of relieving the GP LCPs assigned to the z/OS partition of some of the load, potentially reducing software licensing costs.

For zIIP capacity planning purposes, the workload of a z/OS image can be considered as having three components.

- **Normal content** is the percent of workload that has no zIIP eligible work). When there are no zIIP LCPs defined, the entire workload is considered normal content.
- **Db2 Content** is the total combination of those applications that contain zIIP eligible work.
- **zIIP Eligible Work** is the portion of Db2 Content that can actually run on zIIP LCPs. This is the portion of the workload to consider when determining how much zIIP capacity could actually be consumed.

zIIP partition capacity is always associated with that of the Db2 content workload.

zAAP on zIIP Capability

zAAP on zIIP capability may be enabled for a zIIP partition on z9 and later processor models running z/OS-1.9 or later. This capability allows zAAP eligible work to run on zIIP LCPs (this is the only way that zAAP work can be run on z13 and later since zAAP LCPs cannot be defined). When **zAAP on zIIP** is enabled, all zAAP and zIIP eligible work will be directed to the zIIP LCPs.

Note that **zPCR** does not know that a zIIP partition is enabled for **zAAP on zIIP** ; therefore you need to ensure that the partition is defined without any zAAP LCPs.

Notice Concerning Specialty Engines

Neither **zPCR** nor this document provides descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zAAP, zIIP, and IFL). IBM authorizes clients to use IBM Specialty Engines only to process Eligible Workloads of specific Programs expressly authorized by IBM. These programs are specified in the "Authorized Use Table for IBM Machines", found at:

www.ibm.com/systems/support/machine_warranties/machine_code/aut.html

No other workload processing is authorized for execution on an SE.

IBM offers Specialty Engines at a lower price than General Processors/Central Processors because clients are authorized to use Specialty Engines only to process certain types and/or amounts of workloads as specified by IBM in the Authorized Use Table.

General Comments

zPCR supports zAAPs used as zAAP LCPs, and zIIPs used as zIIP LCPs. For a GP partition defined with SCP as z/OS-1.6 or later, and zAAP or zIIP RCPs exist in the hardware, additional controls are available to define an associated specialty partition. For capacity reporting purposes, each partition's associated LCPs are considered as a separated partition. While GP CPs may be slugged model, zAAP and zIIP CPs are always full speed.

When adding an associated partition to a parent GP partition, eligible work that formerly ran on GP LCPs can now run on the associated LCPs. As a result, it may be possible to reduce the number of parent GP LCPs, and therefore reduce the number of GP RCPs required.

If the amount of zAAP or zIIP eligible workload exceeds the specialty LCP capacity available, some of that work can be run on the GP LCPs.

To derive capacity projections for an associated zAAP or zIIP partition, the number of LCPs being controlled is considered as the sum of the GP LCPs and the associated LCPs. A small switching cost is factored in, based on the processor family and the number of drawers (or books) involved. These costs apply to both the parent GP partition and its associated partition. Once these costs are factored in, the partitions are considered identically to other partition types to determine their capacity projections.

Capacity projections are shown for the parent GP partition and its associated zAAP or zIIP partition as separate values on the **Partition Detail Report** window. You can relate the capacity of an associated partition to its parent GP partition by the matching partition name and/or number. **zPCR** projects the capacity available to accommodate work. The consumable capacity needs to be considered separately for the GP LCPs and the associated LCPs, based on the characteristics of the application. How the combined capacity of parent GP partition and its associated partition actually gets used will depend on the ability of the workload to exploit the specialty LCPs.

If zAAP or zIIP partitions are defined to a parent GP partition, they will have some effect on the capacity of partitions already defined.

zPCR assumes the capacity perspective for all partitions. Therefore zAAP and zIIP partitions are assumed to be 100% busy. If the actual utilization were less than 100%, switching costs would be less, thus allowing the GP LCP capacity to increase somewhat. A window is available where these utilization values may be adjusted (see [zAAP/zIIP Loading](#)).

A **CPcalculator** function, **zAAP Capacity Estimator**, is available in **zPCR** to assess scenarios with specific Java content on various zAAP configurations (see [zAAP Capacity Estimator](#)).

White Paper documentation available

A white paper entitled "z/OS Performance: Capacity Planning Considerations for zAAP Processors" is available on the ATS Techdocs web site:

www.ibm.com/support/techdocs/atstrmstr.nsf/WebIndex/WP100417

Host Summary Report

LPAR Configuration Capacity Planning

Host Summary Report

LPAR Host Summary

Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Advanced zPCR Use

#1 A Current z14 3906-M01
Description: XYZ Production

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Host Processor		
Processor	Brand	IBM Z
Processor	Family	z14
Processor	Model	3906-M01
Speed	Class	700
Maximum	CPs	33
Drawers	Configured	1
Drawer	RCP Pool Contention	Maximal
CP Type	Assigned	Unused
GP	7	0
zAAP	n/s	0
zIIP	1	0
IFL	4	0
ICF	1	0
Total	13	0

CP Pool	Number of			LCP:RCP Ratio	Effective Capacity
	Real CPs	Logical Partitions	Logical CPs		
GP	7	0	0	n/a	11,290
Dedicated	0	0	0	n/a	
Shared	7	3	13	1.857	
zIIP	1	0	0	n/a	1,661
Dedicated	0	0	0	n/a	
Shared	1	2	2	2.000	
IFL	4	0	0	n/a	6,724
Dedicated	0	0	0	n/a	
Shared	4	3	7	1.750	
ICF	1	0	0	n/a	1,476
Dedicated	1	1	1	n/a	
Shared	0	0	0	0.000	
Totals	13	9	23		21,150

Estimate of LPAR Management Time

CP Pool	Utilization level	
	High	Low
GP	0.79%	2.22%
zIIP	0.54%	1.50%
IFL	0.73%	2.04%
ICF	0.32%	0.32%

Calibrate Capacity

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error

The **Host Summary Report** window is accessed from the **LPAR Host and Partition Configuration** window by clicking the **Host Summary** button in the **Capacity Reports** group box. This report is intended to provide an overview of the LPAR host configuration and characterize its partition configuration and capacity for each of the CP pools. Also included is a totals line for the overall LPAR host. Estimates for LPAR management time by pool (as would be reported by RMF) are also provided.

For information concerning capacity projections, see [Accuracy of LPAR Capacity Projections](#).

When there are unused real CPs in the configuration, **zPCR** will provide an estimate for Reserve Capacity representing those CPs. Reserve Capacity information will appear on the **Partition Detail Report** window and in the capacity summary table on the **Control Panel** window. For details [see Reserve Capacity](#).

The fields describing the LPAR host processor (shown on the left side of the **Host Summary Report** window) are as follows:

LPAR Host Processor

- Processor Brand
- Processor Family
- Processor Model
- Speed Class
- Maximum CPs
- Books/Drawers Configured
- Book/Drawer RCP Pool Contention (None → Minimal → Nominal → Maximal)

Real CP Configuration

• GP:	Number Assigned	Number Unused
• zAAP:	Number Assigned	Number Unused
• zIIP:	Number Assigned	Number Unused
• IFL:	Number Assigned	Number Unused
• ICF:	Number Assigned	Number Unused

The fields summarizing the partition configuration (shown on the right side of the **Host Summary Report** window) are as follows:

Rows displayed for each CP pool as follows

Pool Name	RCPs			Capacity
Dedicated	RCPs	Partitions	LCPs	
Shared	RCPs	Partitions	LCPs	LCP:RCP

Regardless of the processor family, General Purpose CPs are always managed as a single pool. The first set of rows will reveal the General Purpose pool metrics.

On **z9 and later processor families**, zAAPs, zIIPs, IFLs, and ICFs are each managed as a separate pool. In these cases, four additional sets of rows will be displayed, one set for each of the CP pools.

A **Totals** line summarizes the allocations and capacity for the entire CPC.

Key to acronyms used

- **RCP** Real CPs
- **Partitions** Logical Partitions
- **LCP** LCPs
- **LCP/RCP** Ratio of Shared LCPs to RCPs
- **Capacity** Effective capacity of the CP pool or CPC

In general, shared partitions should be configured such that the LCP:RCP ratio for the pool remains relatively low, especially when there are many CPs in the pool. Doing so will minimize capacity lost due to LPAR overhead related to excessive sharing of CP resource.

When a CP pool has only one shared RCP, the LCP:RCP ratio will be dictated by the number of logical partitions (i.e., 10 partitions = LCP:RCP of 10.00). As the number of shared RCPs in the CP pool increases, the tolerance for a large LCP:RCP ratio diminishes. As a general rule, when the number of shared RCPs in the CP pool is six, the LCP:RCP ratio should be kept below 3.00. And when the number of shared RCPs in the CP pool is 12 or greater, the LCP:RCP ratio should be kept below 2.00. Of course, if the various partition workloads have complimentary peak utilizations, or tend to remain idle, this rule-of-thumb becomes less important. **zPCR** computes two LCP:RCP limits for each CP pool, based on the number of shared RCPs in the pool:

- **Point at which the LCP:RCP ratio is considered unreasonable.** If the LCP:RCP for a CP pool exceeds this value, capacity results continue to be generated and a note is issued in the message box.
- **Point at which the LCP:RCP ratio is considered excessive.** If the LCP:RCP for a CP pool exceeds this value, capacity results are not generated and an error message is issued in the message box.

Effective capacity values for the LPAR host are expressed as the sum of the **Minimum Capacity** projection that each partition could see. Results represent the “Capacity Perspective”, i.e., maximum sharing contention, where every partition is seeking all the capacity it can get and the partition weights are deciding who gets it. For multiple partitions that tend to have complimentary peak utilizations, or tend to remain idle, the effective capacity could vary somewhat from the projection, and the capacity available to any individual partition could vary significantly, up to its **Maximum Capacity** projection.

Messages Displayed

The following messages may be issued in the message area at the bottom of the **Host Summary Report** window:

Displayed when the LCP:RCP ratio is considered excessive for the number of CPs in the pool. Capacity results continue to be provided. See note below.

Note: GP shared LCP:RCP ratio appears excessive, exceeding threshold value of NN.NN for X shared RCPs.

Displayed when the LCP:RCP ratio is unreasonable for the number of CPs in the pool. In this case, capacity results are not provided. See note below.

Error: GP shared LCP:RCP ratio exceeds cutoff value of NN.NN for X shared RCPs; No results will be generated.

Note: Often the number of LCPs defined to a partition is in excess of what is normally needed, and **HiperDispatch** (z10 or later processors only) or **IRD** is used to manage them down to a reasonable level.

Whenever HiperDispatch is active, parked LCPs should be removed from all partitions in the study that are being run that way.

Whenever IRD is active, LCPs that are varied offline should be removed from all partitions in the study that are being run that way.

Displayed when the number of defined/active partitions is excessive:

Note: Capacity results for LPAR configurations with a large number of partitions have not been validated by benchmarks or client experience.

Displayed when some of the CP pool capacity cannot be consumed due to capping:

Note: xx.xx% of XX (pool) capacity cannot be used by partitions due to capping assignments.

Displayed when the number of LCPs assigned to a partition is inadequate to support the assigned weight percent:

Note: One or more partitions weights indicate more capacity than can be provided with the LCPs defined.

Displayed when the **Include** checkbox for one or more partitions is unchecked:

Note: N defined partitions are excluded from the capacity assessments.

LPAR Management Time

The ***Estimate of LPAR Management Time*** group box is presented below the main table. Estimates of the LPAR management time are shown for each CP pool. These estimates are intended to reflect what would be observed in an RMF ***Partition Data Report***.

The results represent the percent of the CP resource used by LPAR to manage the dispatching of LCPs on RCPs. Both a low percent and high percent are reported, representing the expected range for LPAR management time.

- Utilization level **High** – represents the condition when overall utilization tends to be high and/or small partitions are less likely to interrupt the execution of large ones on a frequent basis.
- Utilization level **Low** – represents the condition when overall utilization tends to be low and/or small partitions are likely to interrupt large ones on a frequent basis.

LPAR management time is generally a small percent of total capacity. The LPAR management time results from **zPCR** are merely estimates, which should not be expected to carry a high degree of accuracy.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, the following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

Window Controls

Click the **Calibrate Reference-CPU** button to calibrate the current LPAR host and partition configuration such that the LPAR host capacity result comes out at a desired value. See [Calibrate Capacity to LPAR host](#) for a discussion of this capability.

Click the **Return** toolbar icon to return to the **LPAR Host and Partition Configuration** window.

Click the **HTM** toolbar icon to create a filename.HTML with the information from this report.

Click the **Help** toolbar icon to access context sensitive help for this window

Partition Detail Report

LPAR Configuration Capacity Planning

Partition Detail Report

Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Advanced zPCR Use
#1 A Current z14 3906-M01
Description: XYZ Production

z14 Host = 3906-M01/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include V	Partition Identification				Partition Configuration										
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Capping		SMT		Capacity		
								Percent	INIT	ABS	V	Benefit	Minimum	Maximum	
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	6,151	9,791
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	3,507	6,513
<input checked="" type="checkbox"/>	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,123	1,684	
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,632	4,547
<input checked="" type="checkbox"/>	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	538	1,613	
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	4,296	6,713
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2,148	3,357
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	279	1,744
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a					1,476	1,476	

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions							
<input checked="" type="radio"/>	With Parent GP	<input type="radio"/>	Separate by Pool				
Show GP Pool Specialty Pools							
<input checked="" type="radio"/>	All Partitions	<input checked="" type="checkbox"/>	GP	<input type="checkbox"/>	zAAP	<input checked="" type="checkbox"/>	zIIP
<input type="radio"/>	Includes Only	<input checked="" type="checkbox"/>	IFL	<input type="checkbox"/>	ICF		

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR		Sum of Weights	SMT Benefit	Capacity Totals
				LCPs	LCP:RCP			
GP	7	3	13	1.857	1,300			11,290
zIIP	1	2	2	2.000	600			1,661
IFL	4	3	7	1.750	625			6,724
ICF	1	1	1					1,476
Totals	13	9	22					21,150

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

The **Partition Detail Report** window is accessed from the **LPAR Host and Partition Configuration** window by clicking the **Partition Detail** button in the **Capacity Reports** group box. This section discusses all the standard items on this window.

Note: For z13 and later hosts, this window also includes SMT benefit values (see [SMT Benefit](#)).

This window includes the necessary metrics concerning each partition participating in the analysis and the range of capacity expectation for each. For information concerning capacity projection accuracy, see [Accuracy of LPAR Capacity Projections](#).

The title area includes the description field set on the **LPAR Host and Partition Configuration** window, followed by the number of active partitions with their distribution across RCP pools. The basis for all capacity results is also provided (i.e., the current **Reference-CPU** and its scaling-factor/metric).

Detail for the partition configuration is displayed below in a J-table. Fields with a white background are alternate input; changing them will modify results accordingly. See [User Controls](#) under "Getting Started" on how to enter data into J-table input fields.

Partition Control

- **Include** Include/exclude this partition in the capacity assessment. Exclude assumes the partition is not HMC activated.
- **Partition Identification metrics**
- **No.** Automatically assigned sequential partition number
- **Type** CP pool to which the partition is assigned
- **Name** Partition name
- **SCP** Operating system
- **Workload** Workload category

Partition Configuration metrics

- **Mode** Partition is dedicated or shared
- **Logical CPs** Number of active (not parked) LCPs
- **Weight** Weight value assigned (SHR only)
- **Weight %** Relative weight within the partition's CP pool (SHR only)

Capping metrics

- **INIT** A checkbox indicates that Partition is Initial Capped (SHR only)
- **ABS** Absolute Capping value (z12 and later processors only). For detail see [Absolute Capping](#).

Partition Capacity metrics

- **SMT Benefit** [z13 and later processors only] Displays the current percent SMT benefit settings (measured or estimated) for zIIP and for IFL partitions. For detail see [SMT Benefit](#).
- **Minimum** Minimum capacity realizable when sharing contention is 100%
- **Maximum** Maximum capacity realizable when there is no sharing contention for the RCPs being used by this partition. The value can seem misleading in cases where different workload categories are assigned to one or more of the other partitions. This value is affected when Absolute Capping is specified for the partition. Generally, Maximum capacity is not useful for capacity planning purposes.

Defined partitions that are HMC activated, but have no SCP IPL'd (i.e., not running the SCP) should be set to “**Not IPL'd**”. Active LCPs will remain to accurately represent the distribution of LCPs across the RCPs, and the assigned weights will be applied across the CP pool. The partition's minimum and maximum capacity values will be assumed as zero (no value is displayed).

Partitions may be excluded in **zPCR**. No LPAR overhead is generated on behalf of excluded partitions. To exclude, the **Include** checkbox should be unchecked.

Note-1: The partition number assigned for associated zAAP, zIIP, IFL, or ICF partitions is the same as the parent GP partition. These partitions do not count towards the maximum number of partitions allowed. The partition name assigned to an associated zAAP, zIIP, IFL, or ICF partition must be identical to the parent GP partition.

Note-2: (IBM z17, z16, z15, z14, z13 only). When a partition's LCP count is excessive, **Logical CPs** is displayed with **this background color**. Best performance is obtained when a partition's LCP count does not exceed the number of RCPs on a single drawer. The partition is automatically excluded from the LPAR configuration, and a dialog is presented allowing an exception for the partition to be included. The warning can be eliminated by reducing the LCP count such that it fits in the largest drawer. Allowing excessive LCPs is not necessarily a problem for **zPCR**; it is merely intended to warn that performance will be less than optimal.

Note-3: When a partition's LCP count falls within 10% of the total RCPs of a drawer, its **Logical CPs** is displayed with **this background color**; a note will appear in the message area.

Note-4: A partition's **Logical CPs** and **Weight** are displayed with a **violet background** when its **Logical CPs** is greater than that which can be consumed, based on its weight percent relative to that of the other partitions in the same CP pool. A note in the message box at the bottom of the window indicates that the setting is less than optimal. The remedy is to reduce the partition's **Logical CPs** or increase the partition's **Weight**. This condition is not necessarily a problem for **zPCR**, as any capacity that cannot be used by such a partition (i.e., **Unusable Capacity**) is redistributed to the other non-capped partitions within the same CP pool.

Note-5: A partition's **Weight** is displayed with an **orange background** when its weight results in a weight percent that cannot be accommodated due to the limited number of LCPs assigned. A note in the message box at the bottom of the window indicates the nature of the problem. The remedy is to increase the partition's LCPs or lower its weight assignment. This condition is not necessarily a problem for **zPCR**, as any capacity that cannot be used by such a partition (i.e., **Unusable Capacity**) is redistributed to the other partitions within the same CP pool.

Note-6: A partition's **Absolute Capping** is displayed with a **pink background** when it is specified such that the partition's **Minimum Capacity** must be reduced. A note in the message box at the bottom of the window indicates that the setting is outside the suggested range. The remedy is to increase the **Absolute Capping** value so that the **Maximum Capacity** exceeds the **Minimum Capacity**. This condition is not necessarily a problem for **zPCR**, as any capacity that cannot be used by such a partition (i.e., **Unusable Capacity**) is redistributed to the other non-capped partitions within the same CP pool.

Note-7: Contention between two classes of CPs will have an effect on capacity results. On multi-drawer CPCs, the GP/zAAP/zIIP class CPs and the IFL/ICF class CPs are each isolated to drawers to the extent possible. There can be only one drawer where the two CP classes might intersect. The drawer where they intersect will have some additional contention cost over that of the other drawers. Maximum contention will always exist on single drawer CPCs. A bar chart is available to show how these 2 CP classes are distributed across drawers and reveals where contention would exist (see [Book and Drawer Considerations](#)).

Explain Minimum and Maximum Capacity

Dedicated Partitions

Minimum Capacity will be equal to its **Maximum Capacity**. Capacity is based solely on the number of LCPs assigned.

Shared Partitions - Minimum Capacity

Minimum Capacity is that which can be realized when all shared partitions within a CP pool are actively competing for the available RCP resource. The determinants for **Minimum Capacity** are the partition's LCP count and its relative weight (or weight percent).

In cases where a partition's weight percent is greater than its LCP count would warrant ($LCP \div RCPs$), the **Minimum Capacity** will be determined by the number of LCPs, and the excess capacity, which cannot be consumed, is redistributed to the other shared partitions within the same CP pool.

In cases where **Absolute Capping** is defined such that the partition's **Minimum Capacity** must be reduced, the excess capacity, which cannot be consumed, is redistributed to the other non-capped shared partitions within the same CP pool.

Shared Partitions - Maximum Capacity

A shared partition's **Maximum Capacity** can only be realized when no other partitions are competing for the same RCP resource within the CP pool. The primary determinant for **Maximum Capacity** is the number of LCPs assigned, since a partition can consume no more RCP resource than it has LCPs.

LPAR cost is assumed to be the same when computing both the **Minimum Capacity** and **Maximum Capacity** values, even though a case could be made that LPAR cost might be less when only 1 partition is competing.

A partition that is **Initial Capped** will have its **Maximum Capacity** set equal to its **Minimum Capacity** (which is controlled by weight %). In actual production, a capped partition capacity may see $\pm 3.6\%$ of its **Minimum Capacity**. Capping a partition has no effect on the capacity of other partitions.

A partition that is **Absolute Capped** will normally have its **Maximum Capacity** set to a value somewhere between its **Minimum Capacity** value and its uncapped **Maximum Capacity** value. The value is established based on the **Absolute Capping** value provided, expressed as the fractional number up to its total LCPs. Should the partition's **Absolute Capping** value result with **Maximum Capacity** less than **Minimum Capacity**, its **Minimum Capacity** is reduced to equal **Maximum Capacity**.

The **Capacity Summary by Pool** group box below the table reports effective capacity values for the LPAR host's processor pools as follows:

All processor families (General Purpose CPs are always managed as a single pool)

- **GP** Sum of GP partition minimum capacity values

z9 and later (separate CP pools exist for specialty engines)

- **zAAP** Sum of zAAP partitions' **Minimum Capacity** values
Note: z13 and later processor models do not support zAAP CPs
- **zIIP** Sum of zIIP partitions' **Minimum Capacity** values
- **IFL** Sum of IFL partitions' **Minimum Capacity** values
- **ICF** Sum of ICF partitions' **Minimum Capacity** values

Overall LPAR Host

- **Totals** Sum of all partitions' **Minimum Capacity** values

Capacity results are determined based on the number of RCPs in each pool that can be exploited. If the total number of LCPs assigned to a pool is less than the number of RCPs defined to partitions, a note appears in the message box, and only the usable RCP capacity is considered.

Partition Pools

A unique background color is assigned to distinguish each of the partition types.



This color key is applied consistently throughout **zPCR** windows.

By default, GP partitions will be positioned first, followed immediately by any associated zAAP, zIIP, IFL, and ICF partitions.

An alternate order can be chosen, such that all partitions in each pool are grouped together, GP first, followed by zAAP, zIIP, IFL, and ICF.

When opening the **Partition Detail Report** window, partitions are initially displayed in the order most recently chosen from this window.

Report Controls

From the ***Partition Detail Report*** window, the following columns can be changed for any partition (various restrictions apply):

- **Include** a partition from being considered in the capacity assessment by un-checking its **Include** checkbox. zAAP, zIIP, IFL, and ICF partitions can be excluded without excluding their parent GP partition. However, if a parent GP partition is excluded, any associated zAAP, zIIP, IFL, and ICF partition will also be excluded.
Note: If a partition is HMC activated, it should not be excluded since its LCPs and weights still participate in producing the capacity results.
- **Name**. Any name up to 8 characters may be specified. If the name of a GP partition is changed, any associated zAAP, zIIP, IFL or ICF partition's name will also be changed.
- **SCP**. Displays the SCP assumed to be controlling the partition. The SCP assigned to an associated zAAP, zIIP, IFL, or ICF partition will be the same as that of the parent GP partition and cannot be changed. The SCP for a parent z/OS partition can only be changed to another version of z/OS. The SCP for a parent z/VM partition can only be changed to another version of z/VM. Any other changes to the SCP must be made via the ***Partition Definition*** window for General Purpose CPs.
- **Assigned Workload**. A zAAP or zIIP partition workload will always reflect the workload of the parent GP partition, and cannot be changed. The workload classification for an associated IFL partition may be set differently than its parent GP partition.
- **Mode**. Partitions may be dedicated and shared. The Mode for associated partitions must be identical to that of the parent partition.

Changing a partition's Mode to dedicated will remove RCPs from the shared pool, which may cause other partitions to become invalid (**Include** will become unchecked).

- **LCPs**. The LCP count is limited to the number of sharable RCPs in the pool.
- **Weight**. Changing a partition's weight value will rebalance the ***Minimum Capacity*** results between partitions within the CP pool.
- **Capping**. Limit the capacity available to a partition.
 - **Initial Capping** (any SHR partition) - Simply click the checkbox for the partition. This will limit its ***Maximum Capacity*** to its ***Minimum Capacity***.
 - **Absolute Capping** (z12 and later processors only). The value set for **ABS** has the effect of reducing the partition's ***Maximum Capacity*** value to that of the number of LCPs specified. For a detailed discussion, see [Absolute Capping](#).

Neither capping case will affect the capacity value results of other partitions.

Any changes to partition metrics made on the ***Partition Detail Report*** window are reflected throughout **zPCR** for this LPAR configuration.

[z/OS Partition with associated zAAP/zIIP LCPs](#)

These associations must be defined to parent GP partition from the **Partition Definition** window. They can only be defined for a parent GP partition running z/OS-1.6 or later. An overhead switching cost applies to both the GP partition and zAAP/zIIP partition. Capacity results consider switching cost and zAAP/zIIP utilization.

- **Name, SCP, Mode** and **Workload** can only be changed via the parent GP partition. All must be identical to that of the parent GP partition.
- **LCPs, Weight**, and **Capping** can be set unique from the parent GP partition.

[z/VM Partition with associated IFL LCPs](#)

These associations must be defined to a parent GP partition from the **Partition Definition** window. They are only supported on z10 and later processors. No overhead cost is considered (i.e., z/VM is provided the full capacity of the GP CPs and z/VM Linux guests are provided the full capacity of the IFLs)

- **Name, SCP**, and **Mode** can only be changed via the parent GP partition. All must be identical to that of the parent GP partition.
- **Workload** can be changed from that of the parent GP partition.
- **LCPs, Weight**, and **Capping** can be set unique from the parent GP partition.

[z/VM Partition with associated zAAP/zIIP/ICF LCPs](#)

These associations must be defined to a parent GP partition from the **Partition Definition** window. They are only supported on z10 and later processors. They can only be exploited by attaching them to the z/VM partition assuming that a z/OS guest will use them. The value of including these is to accurately represent the LPAR Host's resources being used to **zPCR**. No overhead cost is considered (i.e., the parent GP partition sees the full capacity of each partition type). The associated zAAP, /zIIP, and ICF capacity would be utilized by the z/OS guest (**zPCR** has no capability to understand z/VM guest activity).

- **Name, SCP, Workload**, and **Mode** can only be changed via the parent GP partition (ICF partition workload must be **CFCC**).
- **LCPs, Weight**, and **Capping** can be set unique from the parent GP partition.

To delete an associated partition, on the **Partition Definition** window, select it and click the [Delete](#) button.

Any currently defined partition can be excluded by un-checking its **Include** checkbox. When a partition is excluded, it no longer participates in the capacity scenario, and all capacity values are adjusted accordingly. No LPAR costs are assumed an excluded partition. Even though a partition is excluded, it remains in the study's partition list.

Changes to a partition's definition that causes it or other partitions to become invalid will cause their **Include** to become unchecked. In many cases, a dialog box will explain the problem. Any such excluded partition will need to be corrected if it is to be restored to the LPAR configuration. Whenever a partition has been automatically excluded, must be manually re-included.

Table View Controls

- Under ***Display zAAP/zIIP/IFL/ICF Partitions***

Ⓐ **With Parent GP** is the default order. This means that zAAP, zIIP, IFL, and ICF partitions that are associated with a GP partition will be displayed immediately beneath the owning GP partition. Standalone IFL and ICF partitions will follow the GP partitions.

Ⓐ **Separate by Pool** is an alternative order. All partitions of the same type will be grouped together, General Purpose first, followed by zAAP, zIIP, IFL, and ICF.

Note: The most recent ***Display zAAP/zIIP/IFL/ICF Partitions*** setting made from any open ***Partition Detail Report*** window will serve as the default order for every window that displays the individual partitions of an LPAR configuration. The default order can only be changed from a ***Partition Detail Report*** window.

If a table sort is requested, the table is temporarily placed in ***Separate by Pool*** order. The default order is not affected.

- To limit the partitions shown in the window to those that are active (for documentation purposes), click the Ⓢ ***Includes Only*** radio button under **Show** in the ***Table View Controls*** group box. Click the Ⓢ ***All Partitions*** radio button to restore excluded partitions to the view. To re-activate an excluded partition, the Ⓢ ***All Partitions*** radio button must be set so that the partition's ***Include*** checkbox can be reset.

You can limit the report to show only partitions residing in specific pools by using checkboxes associated with each. All partitions continue to participate in the capacity scenario; capacity results are not affected. The un-checked pool(s) simply are not revealed in the table, while their capacity values continue to be reflected in the ***Capacity Summary by Pool*** group box.

To add partitions, to delete a partition, or to change the sequence of partitions within a pool, the ***Partition Definition*** window must be used.

When generating output, the partitions included in the report will always reflect what is currently being presented in this window.

Making Global Changes to the SCP/Workload of Partitions

zPCR provides the ability to make global changes for a defined LPAR configuration to the SCPs and/or workloads assigned to partitions. Requested changes are applied temporarily, such that updated capacity results can be viewed in the ***Partition Detail Report*** window. At any point in the change process, the temporary LPAR configuration can then be committed, becoming permanent for the active study. Use the ***Modify SCP/Workload*** button to display the ***Modify Partitions*** window, providing access to this capability (see [Modify SCP/Workload](#)).

Controlling the Order of Displayed Partitions

By default, GP partitions are always displayed first, followed by zAAP, zIIP, IFL, and ICF partitions. Within each pool partitions are displayed in the order in which they were defined.

zAAP and zIIP partitions must always be associated with a parent GP partition. IFL and ICF partitions can be independent or associated with a GP partition. Any associated zAAP, zIIP, IFL, and ICF partition may either be displayed immediately following its parent GP partition or displayed grouped by pool (the order is determined by the setting in the **Table View Controls** group box). Independent IFL and ICF partitions are displayed last.

Associated zAAP, zIIP, IFL, and ICF partitions can be determined by partition's number and name matching that of the parent GP partition.

The current table view can be temporarily sorted across all the partitions within a pool, based on **LP Name**, **LP Weight Percent**, **LP Min Capacity**, or **LP Max Capacity** (click on the related underlined title). The first sort request will sort ascending, the second will sort descending, and the third will restore the original order. While a sort is active, the table is temporarily placed in **Separate by Pool** order (the default order is not affected). Whenever output is requested, the current view will be used.

With the exception of associated partitions, permanent changes to the sequence of partitions within a pool can be made using the Move icons on the **Partition Definition** window.

Summary by Pool

The **Capacity Summary by Pool** group box shows the **Minimum Capacity** totals for each CP pool. A totals line shows the **Minimum Capacity** total for the entire CPC.

When there are unusable real CPs, two additional **Reserve Capacity** columns will appear in the table. For details see [Reserve Capacity](#).

Messages Concerning Generated Results

Various indicators and text may be issued in the message area at the bottom of the **zPCR** window.

Results Pertaining to the LPAR Host Definition

Text appears in the message area at the bottom of the window.

Displayed when the LCP:RCP ratio is unreasonably high for the number of CPs in the pool. In this case, capacity results are not provided.

Error: GP shared LCP:RCP ratio exceeds cutoff value of NN.NN for X shared RCPs; No results will be generated.

Displayed when the LCP:RCP ratio is considered excessive for the number of CPs in the pool. Capacity results continue to be provided.

Note: GP shared LCP:RCP ratio appears excessive, exceeding threshold value of NN.NN for X shared RCPs.

Note: Often the number of LCPs defined to a partition is in excess of what is normally needed, and **HiperDispatch** (z10 and later processors) or **IRD** is used to manage them down to a reasonable level.

When HiperDispatch is active, parked LCPs should be removed from all partitions that are run that way.

When IRD is active, LCPs that are varied offline should be removed from all partitions that are run that way.

Displayed when the number of defined/active partitions is excessive:

Note: Capacity results for LPAR configurations with an excessively large number of partitions have not been validated by benchmarks or client experience.

Displayed when a portion of the CP pool capacity cannot be consumed due to capping settings:

Note: xx.xx% of XX (pool) capacity cannot be used by partitions due to capping assignments.

Displayed when the number of LCPs assigned to a partition is inadequate to support the assigned weight percent:

Note: One or more partitions weights indicate more capacity than can be provided with the LCPs defined.

Displayed when the **Include** checkbox for one or more partitions is unchecked:

Note: N defined partitions are excluded from the capacity assessments.

Results Pertaining to a Partition Definition

Specific metrics of a partition will be flagged with a background color with elaborating text, using the same background color, appearing in the message area.

Displayed with this background color for partition **Logical CPs** and **Weight**.

Note: The partition's LCP count exceeds the recommended number of LCPs based on its weight relative to that of the other shared partitions in the pool.

Displayed with this background color for the partition **Weight**.

Note: The partition's weight indicates more capacity than its LCPs can provide; Unusable capacity is redistributed to other partitions in the same pool.

Displayed with this background color for the partition **ABS Capping**.

Warning: The "ABS Capping" value will reduce "Minimum Capacity"; Unusable capacity is redistributed to other uncapped partitions within the pool.

IBM z17, z16, z15, z14, and z13 models only. Displayed with this background color for partition **Logical CPs**. The partition's **Include** will be unchecked and a dialog will appear that can be used to override the exception condition.

Warning: On zXX (xxxx) models, best performance is obtained when the partition's Logical CP count doesn't exceed the number of RCPs on a single drawer.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, the **SRB** column will be added to the **Partition Detail Report** window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

Note: **Initial Capping** and **Absolute Capping** are ignored for zIIP partitions using zIIP Boost.

For detail concerning SRB support see [System Recovery Boost](#).

Menu-bar

Edit

Modify SCP/Workload

Opens window with editing capability for globally modifying the SCP and/or Workload of z/OS partitions. See [Modify SCP/Workload](#) for a description of this capability.

Calibrate Capacity

Adjust the LPAR configuration total capacity result to a specific user-specified value. This capability is limited to the 1st LPAR configuration. The reference-CPU setting is so adjusted thus affecting all defined LPAR configurations. See [Calibrate Capacity to LPAR host](#).

Graph

Distribution of RCPs	On multi-drawer or multi-book processors, the GP/zAAP/zIIP CPs and the IFL/ICF CPs are isolated to individual drawers to the extent possible. A bar chart shows how these 2 CP classes are likely distributed across the drawers and reveals the one where contention may exist. For IBM z17, z16, z15, z14, and z13 processors, the location of these 2 CP classes is managed dynamically by PR/SM; therefore, the precise drawer allocation layout cannot be portrayed. However, the total number of RCPs in each drawer and an estimate of the number of GP/zAAP/zIIP and IFL/ICF CPs in contention is provided.
Distribution of capacity	Pie chart showing all partitions or partitions by CP pool.
Capacity available	Bar chart showing all partitions or partitions by CP pool.
Documentation	Links to various informational dialogs.

Window Controls

Click the **Summary** button to display the **Host Summary Report** window, which includes a description of the LPAR host, a summary of the partition configuration, and the projected capacity available (see [Host Summary Report](#)).

Click the **SMT Benefit** button (only available for z13 and later) to add or change SMT benefit percentages for zIIP and IFL partitions running an eligible SCP. See [SMT Benefit](#) for a description of this capability.

Click the **LCP Alternatives** button to test the effect on capacity when changing the SCP and or workloads assigned to partitions. See [LCP Alternatives](#) for a description of this capability.

Click the **zAAP/zIIP Loading** button to assess the effect on General Purpose capacity when the load on associated zAAP or zIIP LCPs is changed from the default. See [zAAP/zIIP Loading](#) for a description of this capability.

Click the **HiperDispatch** button to present a window showing the logical CP type assignments assumed for each partition. See [HiperDispatch](#) for details.

[z17 and z16 only] Click the **Topology** button for a window showing the allocation of partition logical CPs across drawers. See [Topology](#) for details.

Click the **Return** toolbar icon to return to the **LPAR Host and Partition Configuration** window. Any alterations to a partition's **Include** checkbox, **Weight** or **CAP** assignment will be retained.

Click the **HTM** toolbar icon to create an HTML file with the report tables.

Click the **CSV** toolbar icon to create a CSV file with the report tables. (Note that CSV output was enabled for a particular project, and is not being considered for universal implementation in [zPCR](#).)

Click the **Help** toolbar icon to access context sensitive help for this window.

Reserve Capacity

When existent, creates a special case for the Partition Detail Report

When defined real CPs are greater than can be consumed by the LPAR configuration, **zPCR** will provide an estimate of the unusable CP capacity (each CP pool is considered independently). The estimate is referred to as "Reserve Capacity". Such capacity can exist when:

1. Defined partitions cannot make use of all the real CPs available.
2. Partitions defined with SCP = **Not-IPL'd** may cause real CPs to become unused.

To demonstrate Reserve Capacity, the **Sample zPCR Study - Basic Usage** (used throughout this section) was modified as follows:

1. Clone the 1st LPAR configuration.
2. Increase GP real CPs to 10 and zIIP real CPs to 3.
3. Change partition LP-01 SCP to Not-IPL'd.

Partition Detail Report
Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Advanced zPCR Use

#2 **Configuration #2**
Description: Cloned from Current z14 3906-M01

z14 Host = 3906-M01/700 with 18 CPs: GP=10 zIIP=3 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include V	Partition Identification				Partition Configuration								
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight Percent	Capping		SMT	Configured Capacity	
								INIT	ABS	V	Benefit	Minimum	Maximum
✓	1	GP	LP-01	Not_IPL'd	None	SHR	6	700	53.85%	□		6,629	6,629
✓	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	□		1,611	1,611
✓	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	□		4,631	4,631
✓	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	□		1,528	1,528
✓	3	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	□		4,279	6,686
✓	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	□		2,140	3,343
✓	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	□		278	1,737
✓	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	□		1,467	1,467
✓	7	ICF	LP-07	CFCC	DED		1	n/a					

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions			Capacity Summary by Pool									
<input checked="" type="radio"/> With Parent GP	<input type="radio"/> Separate by Pool		CP Pool	Real CPs	LPs	DED LCPs	SHR	Sum of Weights	SMT Benefit	Configured	Estimated Reserve	Totals
<input checked="" type="radio"/> All Partitions	<input checked="" type="radio"/> GP	<input type="radio"/> zAAP	<input checked="" type="radio"/> zIIP	GP	10	3	13	1,300	1,300	11,260	4,326	15,587
<input type="radio"/> Includes Only	<input checked="" type="radio"/> IFL	<input checked="" type="radio"/> ICF		zIIP	3	2	2	0.667	600	3,138	1,719	4,858
				IFL	4	3	7	1.750	625	6,696		6,696
				ICF	1	1	1			1,467		1,467
				Totals	18	9	22			22,562	6,046	28,608

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch Explain Reserve Capacity

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Note: Estimated Reserve Capacity is shown due to under-configured hardware. Once the hardware is completely configured to IPLed partitions the capacity estimate will change.
Note: Unused RCPs consist of GP = 3 of 10, zIIP = 1 of 3
Note: This partition's weight indicates more capacity than its LCPs can provide; Unusable capacity is redistributed to other partitions within the pool

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

When Reserve Capacity exists, additional capacity columns will appear on the **Partition Detail Report** window (as shown above) in the **Capacity Summary by Pool** group box. These columns are labeled:

1. **Configured Capacity** (for the real CPs that can execute workload). This is the traditional **zPCR** capacity result, representing only the real CPs being used.
2. **Reserve Capacity** (for the real CPs that have no workload to execute). This must be considered an estimate because it is unknown know what workload classification would be appropriate. The Average workload classification is used for the estimate.
3. **Total Capacity** is the sum of Configured Capacity and Reserve Capacity.

The **Explain Reserve Capacity** button appears beneath the summary table. It can be used to display the online help for this section of the User's Guide.

The summary table at the bottom of the **Control Panel** window will display additional rows when Reserve Capacity is present, as shown below.

Configuration #2 Cloned from Current z14 3906-M01 z14/700 LPAR Host: 3906-M01/700						
Pool CP Type	#1 GP	#2 zAAP	#3 zIIP	#4 IFL	#5 ICF	CPC Total
RCPs	10	0	3	4	1	18
Partitions	3	0	2	3	1	9
LCPs	13	0	2	7	1	23
Configured Capacity	11,260	n/a	3,138	6,696	1,467	22,562
Est. Reserve Capacity	4,326	0	1,719	0	0	6,046
Total Capacity	15,587	n/a	4,858	6,696	1,467	28,608

Reserve Capacity information is not provided with any other **zPCR** windows.

If the LPAR host configuration is subsequently scaled down such that all real CPs are used, Reserve Capacity information will no longer appear.

On the **Partition Detail Report** window certain actions can trigger a Reserve Capacity situation, such as:

1. Excluding a partition.
2. Changing the number of logical CPs defined to a partition.

When configuration changes cause real CPs to become usable, Reserve Capacity information will no longer be displayed.

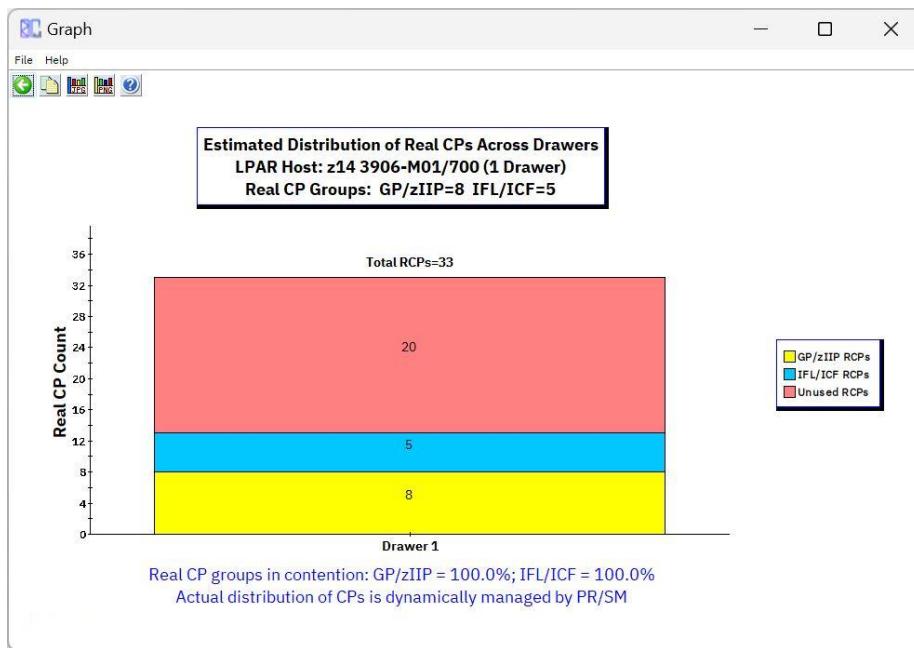
Book and Drawer Considerations

On z12 and later processors, partitions can realize a capacity benefit when spread across multiple drawers (or books). LPAR attempts to minimize contention between **GP/zAAP/zIIP** real CPs and **IFL/ICF** real CPs by isolating the two groups from each other. This isolation minimizes HSB cache interference.

GP/zAAP/zIIP real CPs are generally allocated starting at one end of the book/drawer spectrum while **IFL/ICF** real CPs are allocated starting at the opposite end. Only one drawer (or book) can exist where there would be contention between these RCP groups. A capacity cost for the entire configuration will be generated depending on the number of RCPs in contention.

A chart showing the **Distribution of Real CPs across Books or Drawers** is available from the **Partition Detail Report** window.

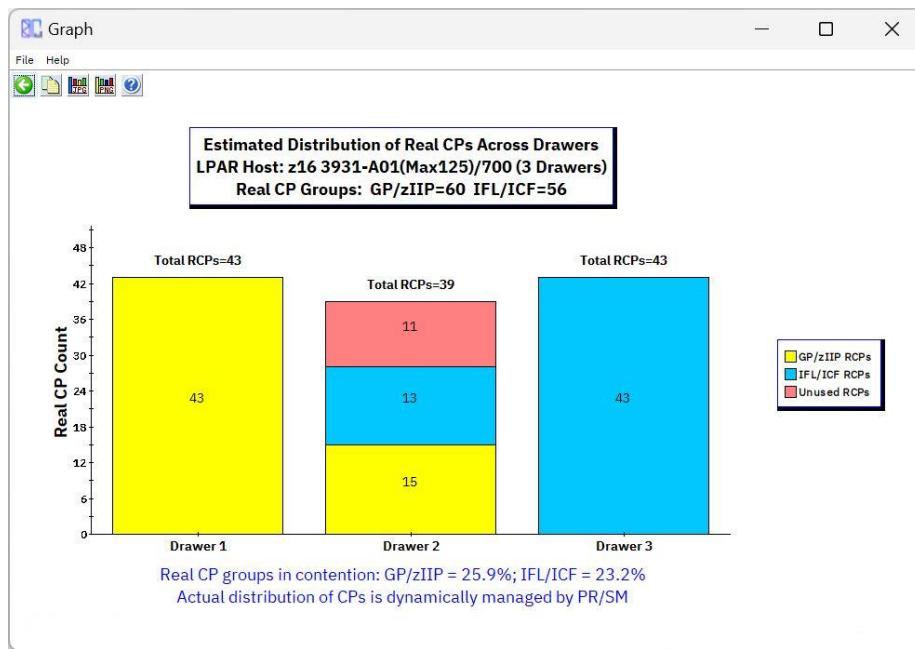
Click the Show Distribution of RCPs Graph tool bar icon  to display a chart showing **zPCR**'s estimation for the distribution of real CPs across books (or drawers). Alternatively, from the menu bar click **Graph→Distribution of RCPs**.



For configurations with a single drawer (or book), each RCP group will be in full contention with the other. The chart below is an example.

Of the 33 RCPs on the drawer, 8 are GP or zIIP and 5 are IFL or ICF. The remaining 20 RCPs are unused.

When the number of drawers increases to two or more, contention between the RCP groups will diminish. There can be only 1 drawer where the RCPs of a group will be in contention. The chart below is an example of this case.



Using three drawers, this example shows:

43 GP/zIIP RCPs are assigned to drawer #1 (max RCPs = 43).

43 IFL.ICF RCPs are assigned to drawer #3 (max RCPs = 43).

The remaining RCPs are assigned to drawer #2 (max RCPs = 39) including:

17 GP zIIP RCPs.

13 IFL/ICF RCPs.

9 Unused RCPs.

Result: The following RCPs are considered to be in contention, thus realizing slightly degraded capacity.

GP/zIIP – 28.3% of all the GP/zIIP RCPs.

IFL/ICF – 23.2% of all the IFL/ICF RCPs.

Capacity cost factors are determined based on the RCP contention percentages. The cost factors are then applied across the entire LPAR configuration (note that **zPCR** does not consider which drawer (or drawers) any given partition is running on). The cost factors applied do consider the effect of slugged GP RCPs where capacity is less than that of the full speed zIIP, IFL and ICF RCPs.

Note: This chart provides an estimation of how RCPs are likely to be distributed across drawers. With z17 and z16 models, the actual distribution of RCPs across the drawers can be seen using the **Topology Report** window (click the **Topology** button on the **Partition Detail Report** window).

Menu-bar**Edit****Modify SCP/Workload**

Opens window with editing capability for globally modifying the SCP and/or Workload of z/OS partitions. See [Modify SCP/Workload](#) for a description of this capability.

Graph**Distribution of RCPs**

On multi-drawer (book) processors, the GP/zAAP/zIIP class CPs and the IFL/ICF class CPs are isolated to drawers to the extent possible. A bar chart shows how these 2 CP classes are likely distributed across the drawers and reveals the one where contention may exist. For IBM z17, z16, z15, z14, and z13 processors, the location of these 2 CP classes is managed dynamically by PR/SM; therefore the precise drawer allocation layout cannot be portrayed. However, the maximum number of RCPs in each drawer and an estimate of the number of GP/zAAP/zIIP and of IFL/ICF CPs in contention are revealed.

Distribution of capacity

Pie chart showing all partitions or partitions by pool

Capacity available

Bar chart showing all partitions or partitions by pool

Documentation

Links to various informational dialogs

System Recovery Boost

LPAR Configuration Capacity Planning

The **System Recovery Boost (SRB)** feature is available only on IBM z17, z16, and z15 processors. **zPCR** supports several aspects of the **SRB** feature. It is used to accelerate partition startup (IPL) and shutdown or to accelerate specific Sysplex recovery events. **SRB** requires operating system support for exploitation.

During the boost period, GP capacity may be temporarily increased by:

- **Speed Boost:** Enable sub-capacity GP engines (/400, /500, and /600) to run at full speed (/700) and/or
- **zIIP Boost:** Increases parallelism by enabling zIIP engines to run any work of the associated GP partition.

The length of the boost period is limited based on the purpose (class) of the boost. Optionally, additional zIIP engines can be activated for up to 6 hours to provide additional zIIP capacity using the SRB Upgrade priced feature.

SRB support applies only to **zPCR**'s **LPAR configuration Capacity Planning** function. If SRB activity does not exist in the current study, tables are unaltered. Only when SRB activity exists, the table format changes discussed below will occur.

Source of SRB data

SRB information cannot be manually defined for a partition. It can only be transferred into **zPCR** via EDF or RMF input. When selecting a performance interval for analysis, intervals that include SRB activity on behalf of any partition are flagged with a red background.

If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the various partition description tables, and the actual SRB partitions are highlighted in the **No.**, **Type**, and **SRB** columns. The SRB column denotes the type of boost applied. Both boost types may be assigned to a GP partition.

- **S = GP speed boost.**
The speed class of sub-capacity LCPs is converted to full speed for the duration of boost period, resulting in temporarily enhanced capacity.
- **I = zIIP boost assist.**
Any z/OS GP workload content can be dispatched to its associated zIIP LCPs during the boost period.
With an SRB Upgrade activated, unused zIIP CPs can be added to the zIIP pool. Any z/OS workload content can then make use of these as LCPs.
Both of these boost types will result in temporarily enhanced capacity for these partitions.

Note: Increased SRB capacities can only be realized while System Recovery Boost is active on behalf of one or more partitions. Such capacity increases cannot be redistributed to other non-boosted partitions. Therefore, when viewing an interval where SRB is active, making adjustments to the overall partition configuration is not valid, and should not be done. Capacity results are only valid for configuration that was active during the loaded SRB interval.

Include ✓	Partition Identification					
	No.	Type	SRB	Name	SCP	Assigned Workload
<input checked="" type="checkbox"/>	1	GP	SI	LP-01	z/OS-2.4	Average
<input checked="" type="checkbox"/>		zIIP	I	LP-01	z/OS-2.4	Average

Whenever SRB was active, each **zPCR** window related to the study will include the SRB indicators and metrics where applicable. In addition, a message line will indicate that SRB was active for one or more partitions.

Configuration has one or more partitions in System Recovery Boost. Don't use this configuration for a processor sizing.

GP speed boost is supported by:

- z/OS-3.2, 3.1, 2.5, 2.4, and 2.3
- z.VM-7.4, 7.3, 7.2, and 7.1

zIIP boost is supported by:

- z/OS-3.2, 3.1, 2.5, 2.4, and 2.3

zPCR can only represent the SRB status that existed at the time of the captured interval. SRB activity is likely to vary over time.

The following operations will include any partition SRB status when present:

- **Clone** LPAR configuration
- **Copy** partitions from EDF, RMF, or a previous zPCR study.
- **Copy or Move** partitions from another LPAR configuration.

Changing the LPAR host, addition or deletion of partitions, or modification of any partition's definition metrics will not affect the status of SRB partitions. However, there is no certainty that SRB behavior would remain identical to that of the original captured EDF/RMF interval.

Note: Performance intervals that include boost activity should not be used for capacity planning or sizing purposes due to the temporary nature of the adjustments made to the LPAR configuration, which are not representative of a typical production workload interval.

Absolute Capping

LPAR Configuration Capacity Planning

Absolute Capping is a feature available on z12 and later processor models. An **Absolute Capping** column will always appear on the **Partition Detail Report** window. **Absolute Capping** can be specified for any shared-mode partition in any CP pool. The **Absolute Capping** value has the effect of decreasing the partition's **Maximum Capacity** value to that of the number of LCPs specified.

Partition Detail Report
Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Basic zPCR Use
#1 A Current z14 3906-M01
Description: XYZ Production

z14 Host = 3906-M01/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include V	Partition Identification				Partition Configuration										
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Weight Percent	Capping		SMT	Capacity		
									INIT	ABS	V	Benefit	Minimum	Maximum	
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	6,151	9,791
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	3,507	5,699
<input checked="" type="checkbox"/>		zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,123	1,684
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,632	4,547
<input checked="" type="checkbox"/>		zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	538	1,613
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	4,296	6,713
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2,148	3,357
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	279	1,744
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	DED		1	n/a						1,476	1,476

Table View Controls
Display zAAP/zIIP/IFL/ICF Associated Partitions
 With Parent GP Separate by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR LCPs	SHR LCP:RCP	Sum of Weights	SMT Benefit	Capacity Totals
GP	7	3	13	1.857	1,300			11,290
zIIP	1	2	2	2.000	600			1,661
IFL	4	3	7	1.750	625			6,724
ICF	1	1	1					1,476
Totals	13	9	22					21,150

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

Absolute Capping values may be specified for any z13 or later shared partition. The **ABS** fields for such partitions will be enabled, and a value may be entered. Values are expressed as a fractional value (N.nn) greater than zero and less or equal to than the number of LCPs defined to the partition. Normally, only the partition's **Maximum Capacity** value is affected by this setting.

Capping Conditions

- Without **Capping**, a partition's **Maximum Capacity** represents the capacity it could get if no other partitions are competing for the shared CP resource.
- With **Initial Capping**, a partition's **Maximum Capacity** is limited to its **Minimum Capacity**.
- An **Absolute Capping** value will result in reducing a partition's **Maximum Capacity** to that of the number of LCPs specified for Absolute Capping

Absolute Capping values are entered manually. If the value entered is greater than the partition's LCPs, a note appears and the value is assumed to equal the number of LCPs defined to the partition. If the value entered generates a **Maximum Capacity** value greater than its **Minimum Capacity**, its **Minimum Capacity** is reduced to match, and a note appears indicating such. This results in unusable capacity for the partition, which is then redistributed to the other non-capped partitions within the same CP pool.

Exception Considerations

Absolute Capping values may become note-worthy or invalid if a partition's LCPs or relative weight is changed. In these cases, the partition's **Absolute Capping** field will be highlighted with a red background. The conditions are:

1. Note: A partition's **Absolute Capping** value exceeds its LCP count; The number of LCPs defined to the partition is assumed.
2. Note: **Absolute Capping** values specified are such that the partition's **Maximum Capacity** would be less than its **Minimum Capacity** (hardware will actually allow this, but it is not recommended). In this case, the partition's **Minimum Capacity** is lowered to equal its **Maximum Capacity**.

EDF and RMF Input

Absolute Capping values can be obtained from EDF or RMF for a z13 or later partition. Such values will become part of the LPAR configuration study.

Note: RMF can only identify if a partition is absolute capped; it does not provide the value assigned. The value must be determined externally and manually entered on the **Partition Detail Report** window.

When copying z12 or later partitions from an EDF, **Absolute Capping** values will remain only if the currently specified LPAR host supports it.

zPCR Study File

Absolute Capping values will be saved as part of the LPAR configuration in a study file.

When copying partitions from a study file where the LPAR host was a z12 or later processor model, **Absolute Capping** values will be included only if the LPAR host is a model that supports it.

Changing the LPAR Host

If the LPAR host is changed to a model that does not support **Absolute Capping**, any **Absolute Capping** values will be purged. If the model is subsequently changed to one that does support Absolute Capping, the values must be reentered.

SMT Benefit

LPAR Configuration Capacity Planning

z13 and later processor models can be run with **Simultaneous Multi-Threading (SMT)** enabled for zIIP and IFL partitions. SMT can provide a capacity benefit for these CP types when the partition is managed by an SCP that supports it.

Partition Detail Report

Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Basic zPCR Use
#1 ! Current z14 3906-M01
Description: XYZ Production

z14 Host = 3906-M01/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include V	Partition Identification				Partition Configuration										
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Capping		SMT	Capacity			
								Percent	INIT	ABS	✓	Benefit	Minimum	Maximum	
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/>	<input type="checkbox"/>			6,151	9,791
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>		3,507	5,699
<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/>	<input type="checkbox"/>			1,123	1,684
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>	<input type="checkbox"/>			1,632	4,547
<input checked="" type="checkbox"/>	3	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/>	<input type="checkbox"/>			538	1,613
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>	<input type="checkbox"/>			4,296	6,713
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>	<input type="checkbox"/>			2,148	3,357
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>	<input type="checkbox"/>			279	1,744
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a					1,476	1,476	

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions
 With Parent GP Separate by Pool

Show GP Pool Specialty Pools
 All Partitions GP zAAP zIIP
 Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR		Sum of Weights	SMT Benefit	Capacity Totals
				LCPs	LCP:RCP			
GP	7	3	13	1.857	1.300			11,290
zIIP	1	2	2	2.000	600			1,661
IFL	4	3	7	1.750	625			6,724
ICF	1	1	1					1,476
Totals	13	9	1	22				21,150

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

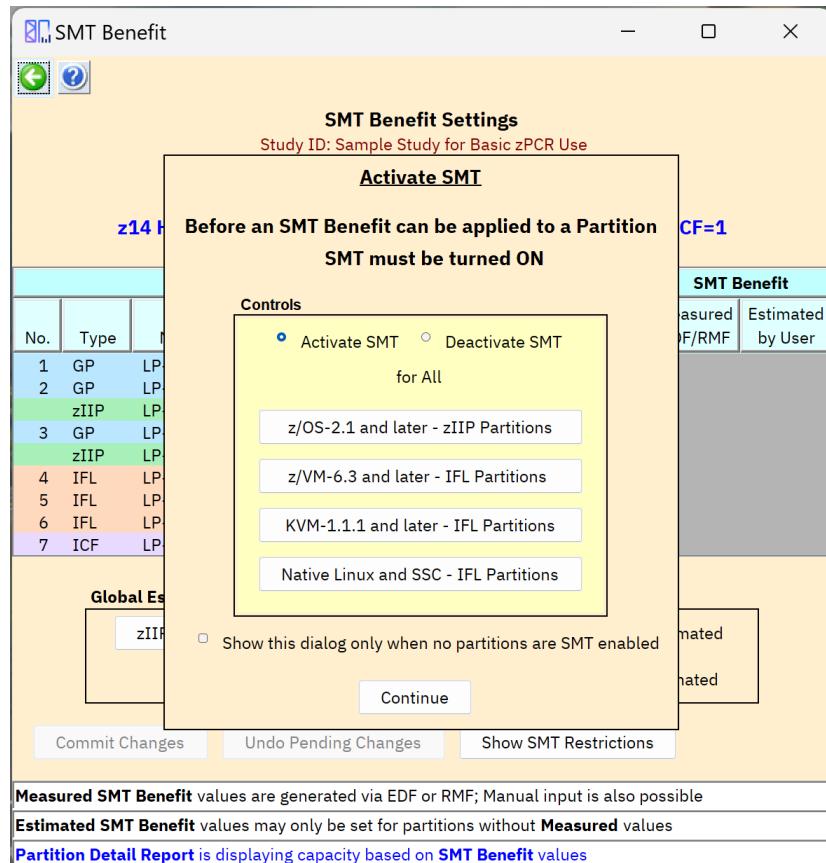
Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

Whenever a z13 or later processor is defined as the LPAR host, two **SMT Benefit** columns are displayed on the **Partition Detail Report** window immediately after the **Capping** columns. An **SMT Benefit** column is also displayed in the **Capacity Summary by Pool** group box. An **SMT Benefit** button will also appear at the bottom of the window.

When **SMT** is checked for a partition, SMT is assumed to be enabled for that partition. The initial SMT benefit of 0% will be refined later via the **SMT Benefit** window.

The SMT benefit checkboxes may also be activated (or deactivated) by group at the time the **SMT Benefit** window is displayed. See below.

SMT benefit values for zIIP and/or IFL partitions are established via the **SMT Benefit** window. Click the **SMT Benefit** button to present the window; it will be displayed adjacent to the **Partition Detail Report** window.



Initially a dialog overlays the **SMT Benefit** window, providing the ability to activate SMT for groups of zIIP and IFL partitions. When requested, SMT activation occurs for all partitions within a group, based on having an SCP assignment that supports SMT.

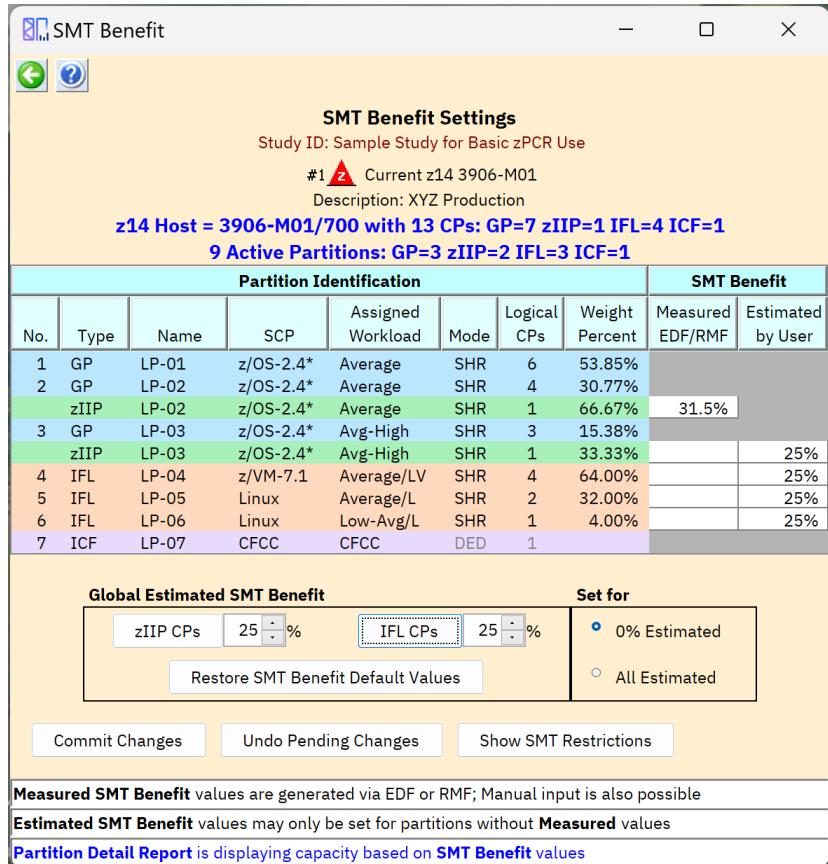
Set the **Activate SMT** radio button to activate or the **Deactivate SMT** radio button to deactivate. The buttons below are available for the requested action:

- **z/OS-2.1 and Later – zIIP Partitions** (SCP version enforced by zPCR)
- **z/VM-6.3 and Later – IFL Partitions** (SCP version enforced by zPCR)
- **KVM-1.1.1 and Later – IFL Partitions** (SCP version not enforced by zPCR)
- **Native Linux and SSC – IFL Partitions** (SCP version not enforced by zPCR)

The overlay dialog may be skipped for subsequent accesses to the **SMT Benefit** window by checking the **Show this dialog only ...** checkbox.

Click **Continue** to proceed to the **SMT Benefit** window.

SMT benefit values to be applied are entered on the **SMT Benefit** window.



SMT Benefit Values

There are two types of SMT benefit values. Only 1 type can be defined for each zIIP or IFL partition where SMT has been enabled.

- **Measured:** Measured values are typically obtained from performance data produced by z/OS or z/VM. When obtaining partition information via EDF or RMF, these measured values are loaded into the LPAR configuration. Measured values may also be entered manually (double click on the entry field, enter a value between $\pm 90.0\%$, and press **Enter**). Negative values indicate that there is no benefit for enabling SMT; such values will be displayed with **brown text**. A measured value precludes the assignment of an estimated value for that partition.
- **Estimated:** Estimated values are intended to approximate an expected benefit due to SMT. Estimated entry fields appear as spin buttons. Values may be set between 0% and 60% in 1% increments. For the estimated setting to take effect, simply move the mouse off of the estimated entry field.

Note that a **Measured SMT Benefit** may be converted to an **Estimated SMT Benefit**, right click the entry field, then click **Make this an Estimated Value**. If an estimated value is not desired, set it to zero.

Global Estimated SMT Benefit

Estimated values may be set globally for all zIIP partitions and for all z/VM IFL partitions that have SMT enabled, using the controls in the ***Global Estimated SMT Benefit*** group box. The default values are provided as follows: zIIP is 25%; IBM z17, z16, z15, z14 IFL is 25% and z13 IFL is 20%. These may be modified using the spin buttons. Click the **zIIP CPs** or the **IFL CPs** button to apply the estimated values. For IFL partitions, estimated SMT benefit values can only be set for z/VM and KVM partitions. Linux partitions running on IFLs may be set manually (recognize that there is currently no SMT support available for Linux).

Radio buttons are provided on the right to control how the estimated values for either zIIP or IFL partitions are to be applied.

- 0% Estimated Only update estimated values that are currently 0%
- All Estimated Update all estimated values

The **Restore SMT Benefit Default Values** button can be used to reset the values to the supplied defaults, in case they have been modified.

Changes made on the **SMT Benefit** window are considered temporary. To commit pending changes, click the **Commit Changes** button. To remove pending changes, click the **Undo Pending Changes** button.

Click the **Show SMT Restrictions** button to display considerations for applying **SMT Benefit** values.

Click the **Return** tool bar icon to close the **SMT Benefit** window. If changes are pending, a dialog will ask whether the changes are to be committed or discarded. The full function of the **Partition Detail Report** window will be restored, reflecting the SMT benefit values that have been set.

Clicking the **Return** tool bar icon on the **Partition Detail Report** window will close both windows, returning to the **LPAR Host and Partition Configuration** window.

As **SMT Benefit** values are set, the **SMT Benefit**, **Minimum Capacity**, and **Maximum Capacity** columns on the (simultaneously displayed) **Partition Detail Report** window are updated to reflect the changes.

Partition Detail Report

Based on LSPR Data for IBM Z Processors
Study ID: Sample Study for Basic zPCR Use
#1 ! Current z14 3906-M01
Description: XYZ Production

z14 Host = 3906-M01/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include	Partition Identification				Partition Configuration										
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Weight Percent	Capping INIT	Capping ABS	SMT	Capacity		
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	6,151	9,791	
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	3,502	5,690	
<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	32%	1,476	2,215
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	1,629	4,538	
<input checked="" type="checkbox"/>	3	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	est. 25%	672	2,017
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	est. 25%	5,371	8,392
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	est. 25%	2,685	4,196
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	est. 25%	349	2,179
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a				<input checked="" type="checkbox"/>	1,476	1,476	

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions
 With Parent GP Separate by Pool

Show GP Pool Specialty Pools
 All Partitions GP zAAP zIIP
 Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR		Sum of Weights	SMT Benefit	Capacity Totals
				LCPs	LCP:RCP			
GP	7	3	13	1.857	1.300			11,282
zIIP	1	2	2	2.000	600	est. 29%	2,149	
IFL	4	3	7	1.750	625	est. 25%	8,405	
ICF	1	1	1				1,476	
Totals	13	9	22					23,311

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

In the **SMT Benefit** column, **Estimated** values are prefixed with “est”, while **Measured** values are not.

Measured SMT benefit values captured via EDF or RMF that are negative indicate that there is no benefit realized by enabling SMT for the partition; such negative values will be displayed with **brown text**.

The **Capacity Summary by Pool** group box values are also updated. The **SMT Benefit** column shows the combined benefit applied for all the zIIP partitions and for all the IFL partitions. If any of the contributing partitions have estimated SMT benefit values, “est” will prefix the value.

Whenever SMT is deactivated for a partition, its SMT benefit is retained behind the scene, but not applied. When SMT is subsequently activated for the partition, its SMT benefit reappears and is applied. Use of the SMT checkbox facilitates the comparison of a partition’s capacity without and with SMT enabled.

Clicking the **Return** tool bar icon on the **Partition Detail Report** window will close both windows, returning to the **LPAR Host and Partition Configuration** window.

SCP Considerations

With the current **zPCR** implementation **be certain that each zIIP and IFL partition will be controlled by an appropriate SCP version before applying **SMT Benefit** values to it.**

SMT is currently supported only by the SCPs listed below.

- To represent **z/OS** using SMT on zIIP LCPs, **zPCR** requires that the SCP version be **z/OS-2.1** or later.
- To represent **z/VM** using SMT on IFL LCPs, **zPCR** requires that the SCP version be **z/VM-6.4** or later.
- To represent **KVM** using SMT on IFL LCPs, **KVM for IBM Z v1.1.1** or later is required. The KVM SCP is not qualified by a version number in **zPCR**. Be certain that the appropriate KVM version will be run before applying SMT benefit estimates.
- To represent **Linux** using SMT on IFL LCPs, a version supporting SMT is required. The Linux SCP is not qualified by a version number in **zPCR**.
- To represent **Secure Service Container (SSC)** using SMT on IFL LCPs, a version supporting SMT is required. The SCP is not qualified by a version number in **zPCR**.

Where **Estimated SMT Benefit** values have been set for zIIP and IFL partitions, the expected margin-of-error for capacity results must be considered greater than that normally assumed due to uncertainty of the estimate. An additional line of text will appear at the bottom of each capacity result window to indicate such. If the default **Estimated SMT Benefit** settings are used, consider the margin-of-error to be $\pm 10\%$.

For any partition where an SMT value has been set, its SMT checkbox on the **Partition Detail Report** window may be used to disable or enable SMT, adjusting capacity results accordingly.

With EDF or RMF input, when a zIIP or IFL partition's utilization is less than 20% or when the **Measured SMT** benefit is greater than 90%, the SMFT Benefit is considered as being unreliable, and the default **Estimated SMT** value is assigned instead. While viewing the **SMT Benefit** window the ability to use the actual **SMT Measured** value is provided. An **SMT Benefit** window example is presented below.

Partition Identification								SMT Benefit	
No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight Percent	Measured EDF/RMF	Estimated by User
1	GP	VGIA	z/OS-2.3*	Average	SHR	17	42.55%	25%	25%
	zIIP	VGIA	z/OS-2.3*	Average	SHR	6	40.00%		
2	GP	GD1P	z/OS-2.3*	Average	SHR	2	1.01%	13.88%	13.88%
3	GP	VGIB	z/OS-2.3*	Average	SHR	3	13.88%		
	zIIP	VGIB	z/OS-2.3*	Average	SHR	5	20.00%	25%	25%

Global Estimated SMT Benefit

zIIP CPs
25
IFL CPs
25

0% Estimated
 All Estimated

Commit Changes
Undo Pending Changes
Show SMT Restrictions

Note: An unreliable measured SMT Benefit value has been captured from the EDF/RMF data for the selected interval.

The SMT Benefit cells of the subject partition(s) are displayed with a pink background and a related note is displayed in the message area at the bottom of the window.

Right click on the pink background area of a subject partition to reveal a pop-up and select it: **Evaluate unreliable measured value**.

A dialog documenting the unreliable measured value is presented with selection buttons to make a choice:

- Click **Yes** to apply the **Measured SMT Benefit** value.
- Click **No** to retain the **Estimated SMT Benefit** value.

The selection made is considered final after committing the changes and the pink background color will be removed if the unreliable measured value was applied. Before committing the changes, the **Undo Pending Changes** button can be used to restore the previous benefit value. If every subject partition is using its unreliable measured benefit value, the **SMT Benefit** window will return to its normal format.

Changing the LPAR Host

If the LPAR host is changed to a processor model that does not support SMT, the SMT benefit column will be excluded from the **Partition Detail Report** window and all **SMT Benefit** values will be purged (purged values cannot be recovered).

zPCR Study File

SMT Benefit values will be saved as part of the LPAR configuration in a study file.

When copying partitions from a study file where the LPAR host is a z13 or later model, **SMT Benefit** values will be included only if the currently specified LPAR host is a model that supports SMT.

Study files that include **SMT Benefit** values cannot be loaded into previous **zPCR** versions.

LSPR Table: Displaying IFL CPs

When displaying IFL CPs in the ***LSPR Table*** window, a ***SMT Benefit*** value may be specified, uplifting the z13 or later processor model capacity values to represent that benefit (see **Capacity Exceptions** under **Settings** on the ***LSPR Table Control*** window menu-bar.

System Recovery Boost Considerations

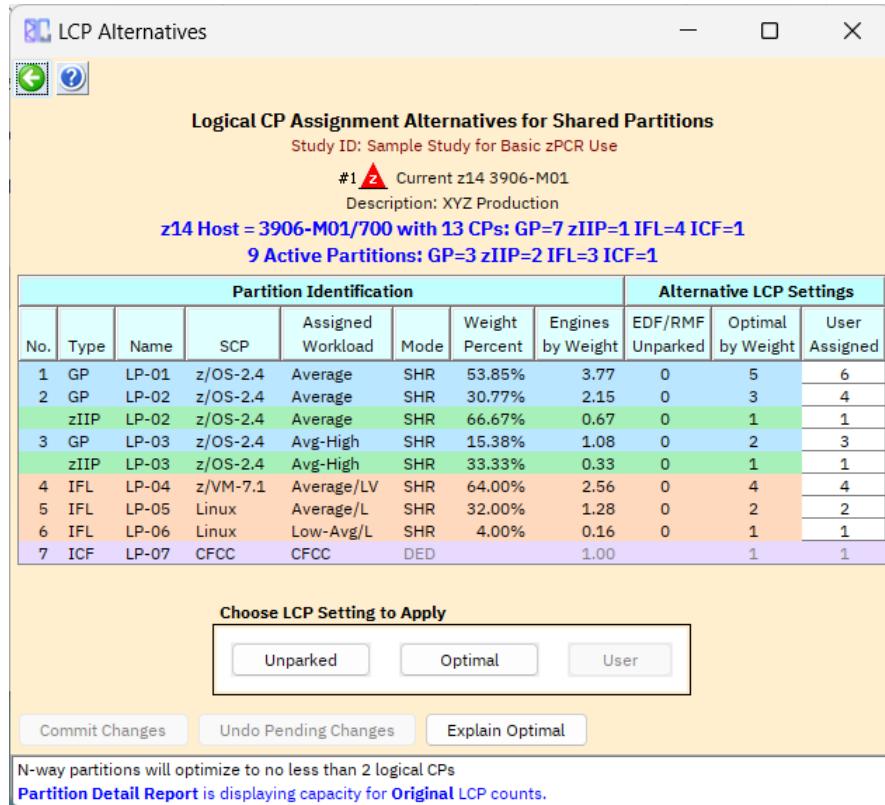
If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

LCP Alternatives

LPAR Configuration Capacity Planning



The **LCP Alternatives** window is accessed from the **Partition Detail Report** window by clicking the **LCP Alternatives** button. This window provides the ability to test the effect of various sets of LCP counts on the overall capacity of the configuration. Only shared partition LCP counts can be changed.

While the **LCP Alternatives** window is active, the **Partition Detail Report** window remains visible, with its capacity results reflecting the effects of the alternative LCP setting applied. While most of the function on the **Partition Detail Report** window is disabled, **HTML** output and **Graph** creation are allowed.

Four variations of LCP settings may be applied, temporarily replacing the current LPAR host configuration's LCP settings:

1. **EDF/RMF Unparked:** If any of the original partitions were generated from EDF or RMF data, and include **Parked LCPs** for 1 or more partitions, this column is enabled, and the LCPs for those partitions are set to the unparked LCP count (**Total LCPs** minus **Parked LCPs**) when the configuration was transferred into **zPCR**. The LCP count for all the other partitions remains at the **Total LCPs** value. If partition LCPs have been modified, these values can be restored by clicking the **Unparked** button in the **LCP Setting Alternatives** group box. Note that the LCP counts in this column are only available during the **zPCR** invocation when the EDF or RMF was read. They are not retained in a saved study.

Note: To allow **zPCR** to represent capacity more accurately, parked LCPs should always be removed from partitions running under the control of HiperDispatch. This includes z/OS partitions and any associated zAAP or zIIP LCPs. It also includes z/VM and any associated IFL LCPs. The **EDF/RMF Unparked** count will only reflect information for the measurement intervals that were chosen. The number of parked LCPs could vary considerably across intervals. Therefore, additional analysis outside of **zPCR** may be necessary to more accurately assess the average number of unparked LCPs. In this case, the **User Assigned** column should be used to enter and test results for these LCP counts.

2. **Weight Based:** LCP counts are computed using a new algorithm that calculates the **Optimal** number of LCPs for each partition with shared LCPs.
 - The Optimal shared LCP count setting utilizes the number of weight-based engines and no more than 1 or 2 additional LCPs, depending on the n-way, to provide the appropriate amount of white space, adhering to established best practices.

N-way partitions will optimize to no less than 2 LCPs for SCP availability reasons.

To update the **Partition Detail Report** window to reflect capacity for these values, click the **Optimal** button in the **LCP Setting Alternatives** group box.

3. **User Assigned:** When the **Alternative LCPs** window is opened, partition LCPs are set to the values on the **Partition Detail Report** window. The values can be user modified to any legitimate count. To update the **Partition Detail Report** window to reflect capacity for these values, click the **User** button in the **LCP Setting Alternatives** group box. The **User** button will remain disabled until at least one of the individual user LCP counts has been changed, or one of the other buttons has been used.

Partitions that have been excluded (**Include** is unchecked) will not have their LCPs adjusted. In this case, their relative weight no longer contributes to the weight percent total used to distribute LCP counts for the other partitions.

Suggested LCP counts are capped at the maximum LCPs supported by the partition's assigned SCP. When zAAP or zIIP LCPs are associated with a GP partition, the total LCPs for the combination cannot exceed the maximum supported by the SCP. For such associations, the LCPs determined for the specialty partition will be limited if necessary. If the GP partition is determined to have the maximum LCPs assigned, that number will be reduced by 1 so that a single specialty LCP can be assigned. This same logic applies for IFL LCPs associated with a z/VM partition.

As **LCP** counts are changed, the **LCPs**, **Minimum Capacity**, and **Maximum Capacity** columns on the **Partition Detail Report** window are dynamically updated to reflect the changes. The **Capacity Summary by Pool** group box values are also updated.

Changes made to LCP settings are considered temporary. To commit pending changes to the LPAR configuration, click the **Commit Changes** button. Click the **Undo Pending Changes** button to restore settings back to the previous commit.

Use the **Explain Optimal** button for a high-level description of how Optimal LCPs are calculated.

Click the **Return** tool bar icon to close the **LCP Alternatives** window. If any changes are pending, a dialog will ask whether the changes are to be committed or discarded. The full function of the **Partition Detail Report** window will be restored, reflecting the last LCP assignment setting made. Note that, once the study has been saved and subsequently reloaded, the original **EDF/RMF Unparked** counts are lost.

Clicking the **Return** tool bar icon on the **Partition Detail Report** window will close both windows, returning to the **LPAR Host and Partition Configuration** window.

System Recovery Boost Considerations

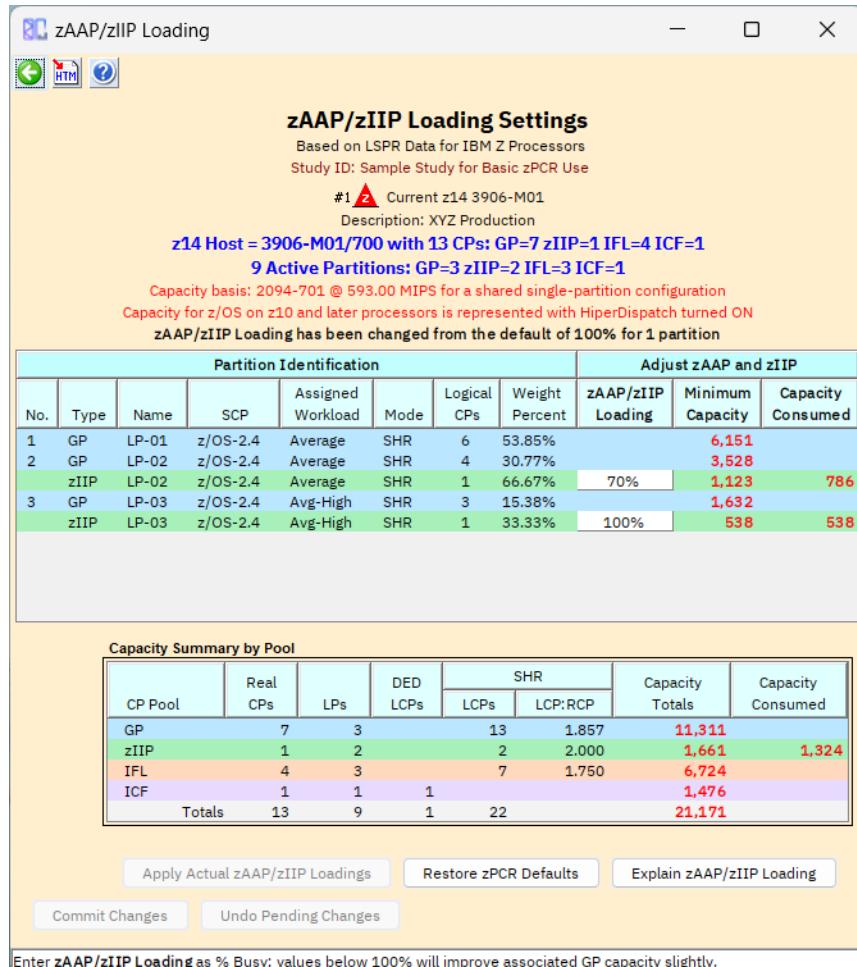
If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

zAAP/zIIP Loading

LPAR Configuration Capacity Planning



The **zAAP/zIIP Loading** window is accessed from the **Partition Detail Report** window by clicking the **zAAP/zIIP Loading** button. This window provides the ability to assess the effect on General Purpose capacity when an associated zAAP or zIIP partition is running at less than full loading. The window can only be accessed if associated zAAP or zIIP partitions are defined in the configuration.

Only parent GP partitions, and their associated zAAP and zIIP partitions are presented on this window. zAAP and zIIP partitions are displayed immediately below the parent GP partition.

With the following two exceptions, columns displayed are the same as those shown on the **Partition Detail Report** window.

- The **zAAP/zIIP Loading** column provides a spin button input field (white background) for each zAAP and zIIP partition. These fields are used to set the zAAP or zIIP loading to any value between default 100% and 0%. Move the mouse onto the field to access the spin button. Select any value from 0 to 100 representing the percent loading and press **Enter**. Specifying a loading of less than 100% will result in increasing the parent GP partition's capacity, reflecting a reduction in management and switching cost.
- zAAP and zIIP partition capacity available is not affected by the loading value specified; however, the capacity considered as being consumed will be reduced to the value shown in the **Capacity Consumed** column.

When an LPAR configuration is generated from EDF or RMF, actual GP and zAAP/zIIP utilizations are known. In such cases the **Apply Actual zAAP/zIIP Loadings** button will be enabled. To apply those values, click the button. Since GP LCPs in this window are portrayed at full capacity, the measured GP utilization is normalized to 100% utilization, using the formula.

$$\text{zAAP/zIIP Loading} = \text{zAAP/zIIP Utilization} \div \text{GP Utilization} \times 100\%$$

This formula is a reasonable approximation of the loading factor to be applied using partition physical utilization values. In reality, partition logical utilization values should be used, which is the case when EDF or RMF input is available.

For cases where an actual zAAP/zIIP loading value is unknown, either estimate it or leave it at 100%. Some guidelines for estimating **zAAP/zIIP Loading** follow:

Note that, for z17, z16, z15, z14, and z13 LPAR host processors, the use of SMT on zIIP logical CPs is considered when determining the influence on capacity for the parent GP partition. The above interpolation is based on the capacity values determined after any SMT benefit is applied.

zPCR algorithms assume 50% offload from GP to associated zAAP or zIIP as the basis for overhead to manage the CP combination. This means that the zAAP/zIIP capacity requirement would be equal to that of the parent GP partition.

- When GP capacity is equal to or greater than the associated zAAP or zIIP capacity, one can assume the zAAP or zIIP LCPs would be fully utilized and **zAAP/zIIP Loading** should remain at 100%.
- When GP capacity is less than zAAP or zIIP capacity, it is likely that the zAAP or zIIP LCPs would be less than fully utilized and **zAAP/zIIP Loading** should be reduced. This case becomes more likely when GP CPs are less than full speed engines (zAAP and zIIP CPs are always full speed).
- If the expected offload from GP to zAAP or zIIP is known, compute the zAAP/zIIP capacity requirement by multiplying GP partition capacity by the expected offload percent times 2.0. Then adjust the **zAAP/zIIP Loading** so that amount of zAAP or zIIP capacity is perceived as being used (**Capacity Consumed** column).

In the **Capacity Summary by Pool** group box, the total capacity for the General Purpose pool reflects the sum of GP partition's **Minimum Capacity** values using the currently specified **zAAP/zIIP Loading** values. **Capacity Consumed** is computed for both the zAAP and zIIP CP pools, also based on the currently specified **zAAP/zIIP Loading** values. This information can be useful for the purpose of increasing or decreasing zAAP and zIIP capacity by altering the number of RCPs available (such modifications must be made on the **LPAR Host** window).

Changes made to **zAAP/zIIP Loading** values are considered temporary. To commit pending changes to the LPAR configuration, click the **Commit Changes** button. To remove pending changes, click the **Undo Pending Changes** button.

Click the **Restore zPCR Defaults** button to set all of the **zAAP/zIIP Loading** percentages to the default 100%.

If enabled, click the **Apply Actual zAAP/zIIP Loadings** button to reset the **zAAP/zIIP Loading** values to represent measured utilizations. Note that this can only be done for an LPAR configuration created from EDF or RMF input.

Click the **Explain zAAP/zIIP Loading** button for a short dialog discussing the benefit of providing **zAAP/zIIP Loading** values less than 100%.

Explain zAAP/zIIP Loading

By default, **zPCR** capacity results assume 100% utilization, regardless of CP type. In cases where zAAP/zIIP LCP utilization is less than 100%, the associated GP LCP capacity available will increase due to smaller N-way effect and lower switching cost. This effect can be demonstrated by specifying a **zAAP/zIIP Loading** as the value of the zAAP/zIIP utilization.

The N-way of a GP + zAAP/zIIP association is considered to be the sum of the logical CPs involved. Capacity values are derived for:

1. GP only (0% zAAP/zIIP utilization; i.e., no switching cost)
2. GP + zAAP/zIIP (100% zAAP/zIIP utilization; i.e., maximum switching cost)

Actual GP partition capacity is determined by interpolating between these 2 capacity results based on the **zAAP/zIIP Loading** value specified.

When an LPAR configuration is generated from EDF or RMF, actual zAAP/zIIP utilization values are known. In such cases, the **Apply Actual zAAP/zIIP Loadings** button is available to make assignments representing those utilizations. For zPCR purposes, the measured zAAP/zIIP utilizations are applied; thus a **zAAP/zIIP Loading** value is generated.

Click the **Return** toolbar icon to close the **zAAP/zIIP Loading** window and return to the **Partition Detail Report** window with the **Minimum Capacity** of the affected GP partitions adjusted. If changes are pending, a dialog will ask whether the changes are to be committed or discarded.

To create an HTML file reflecting the window, click the **HTM** toolbar icon.

To access context sensitive help for this window, click the **Help** toolbar icon.

If any user specified **zAAP/zIIP Loading** values (other than the default 100%) are committed to the configuration, they will be carried forward in a saved study file.

Cloned configurations will include the **zAAP/zIIP Loading** values of the original LPAR configuration. When comparing LPAR configurations, matching partitions (Name, SCP, Workload, and DED/SHR), should have matching **zAAP/zIIP Loading** values in order for the comparison to be fair. When making such comparisons, if **zAAP/zIIP Loading** values for any partition in either configuration is less than 100%, a note will appear in the message area at the bottom of each comparison window.

System Recovery Boost Considerations

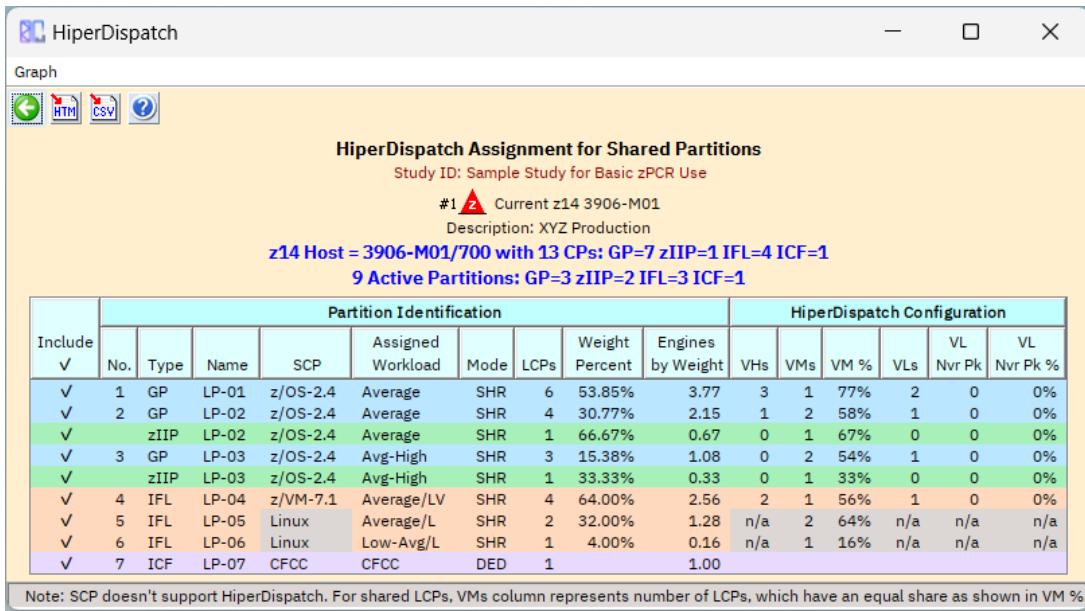
If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

HiperDispatch

LPAR Configuration Capacity Planning



The **HiperDispatch** window is accessed from the **Partition Detail Report** window by clicking the **HiperDispatch** button. When the **HiperDispatch** window is active, the **Partition Detail Report** window remains visible with most of its function enabled. Any changes will be applied to the **HiperDispatch** window.

HiperDispatch supports shared logical CPs running:

- z/OS (v1.7 and later) in a GP partition. zAAP and zIIP shared logical CPs associated with the z/OS partition are similarly affected.
- z/VM (v6.3 and later) in a GP or IFL partition. IFL shared logical CPs associated with the z/VM partition are similarly affected.

HiperDispatch support does not apply to partitions running any other SCP. When a partition SCP is defined as other than z/OS or z/VM, the background color of the **SCP** field and the **HiperDispatch** fields will be set to silver. In addition the following note will appear in the message box at the bottom of the window.

Note: SCP doesn't support HiperDispatch. For shared LCPs, VMs represents number of LCPs, which have an equal share shown as VM %.

The **HiperDispatch** function projects the way logical CPs of these partitions will be distributed in three categories:

1. **Vertical High (VH)** LCPs are essentially dedicated to the partition. They do not service other partitions. They are fully assigned to that partition's workload demand.
2. **Vertical Medium (VM)** LCPs are shared among partitions. The percent of time that a partition is entitled to have their services is depicted on the charts..
3. **Vertical Low (VL)** LCPs have 0% share but are available to a partition in the event other partitions do not require the level of service specified by their weight. Vertical Lows are provided so a partition may use GP LCP capacity above the amount guaranteed by its LPAR weight. If there is no available GP LCP capacity above the weight, then Vertical Lows are parked.

There is a special kind of Vertical Low processor which is never parked. These are typically used to assure that small partitions always have a second logical engine available. The following note will apply.

Note: Partition has a **VL** that is never parked. The **VM %** will be dispatched across the **VM** and the **VL** engines.

The **HiperDispatch** window reiterates most of the **Partition Identification** information from the **Partition Detail Report** window. The following columns are unique.

- **Engines by Weight:** Partition Weight % times the number of real CPs in the pool.
- **VHs:** Number of LCPs categorized as **Vertical High**.
- **VMs:** Number of LCPs categorized as **Vertical Medium**.
- **VM %:** Percent of time the partition's **Vertical Medium** LCPs are committed.
- **VLs:** Number of LCPs categorized as **Vertical Low**.
- **VL Nvr Pk** Number of LCPs categorized as **Vertical Low Never Parked**.
- **VL Nvr Pk%** Percent of time the partition's **Vertical Low Never Parked** LCPs are committed.

As input fields are modified on the **Partition Detail Report** window, results shown on the **HiperDispatch** window will be updated accordingly. Note that when exiting the HiperDispatch window, any changes made to the **Partition Detail Report** window are not automatically reset.

For GP or IFL partitions where HiperDispatch is not supported, only the **VMs** and **VM %** columns apply. For ICF partitions, none of the HiperDispatch columns apply.

Menu-bar

Graph

For the selected CP pool, a bar for each partition shows the number of **Vertical High**, **Vertical Medium**, and **Vertical Low** LCPs. On each of the Vertical Medium LCPs the value shows the percent of time the logical CP is active with that partition.

GP LCP HD Assignments

zAAP LCP HD Assignments

zIIP LCP HD Assignments

IFL LCP HD Assignments



On each chart, the Y-axis shows the partition name along with its count of **VH**, **VM**, and **VL** LCPs. Each color coded rectangle represents a logical CP and its HiperDispatch type that is defined to that partition. For GP and IFL partitions where HiperDispatch is not supported, the LCPs will be depicted with a silver color.

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, a **SRB** column will be added to the window, and each partition with SRB activity will have its **No.**, **Type**, and **SRB** fields highlighted. The following message will appear at the bottom of the window.

Partition was in System Recovery Boost during the selected interval. Don't use this interval for a processor sizing.

For detail concerning SRB support, see [System Recovery Boost](#).

Topology Report

LPAR Configuration Capacity Planning

Partition topology information is available for z17 and z16 processors only. z/OS-2.3 (or later) SMF data or z/VM-7.2 (or later) Monitor data is required to capture topology information which can only be passed to zPCR via EDF input.

Topology Report
Study ID: Not specified

#1 z16 with Topology data
z16 Host = 3931-A01(Max82)/700 with 55 CPs: GP=29 zIIP=10 IFL=9 ICF=7
16 Active Partitions: GP=5 zIIP=5 IFL=3 ICF=3
Note: Topology configuration changed during the measurement interval.

No.	Name	Type	Drawer 1								Drawer 2							
			DCM 1		DCM 2		DCM 3		DCM 4		DCM 2		DCM 3		DCM 4			
			Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 1	Chip 2	Chip 2	Chip 1	Chip 1	Chip 2	Chip 2	Chip 1	Chip 1	Chip 2
1	ZOSPRD1	GP	(6)H		(6)H	(6)H	(6)H	(2)H (1)M		(2)L								
	ZOSPRD1	zIIP			(2)M					(5)H	(1)L (2)-							
2	ZOSPRD2	GP									(3)M (1)L							
	ZOSPRD2	zIIP									(1)M (1) (8)-							
3	ZOSPRD3	GP						(1)L			(1)M (1)L							
	ZOSPRD3	zIIP			(1)-					(1)L	(2)M (1)L (5)-							
4	ZOSTST1	GP									(2)M							
	ZOSTST1	zIIP									(2)M (8)-							
5	ZOSTST2	GP						(1)M (1)L										
	ZOSTST2	zIIP			(1)M (1)L					(8)-								
6	VMPRD1	IFL										(4)H (1)M						
7	VMPRD2	IFL									(1)M (2)L							
8	VMSTST1	IFL									(1)L							
9	CFPRD1	ICF											(5)-					
10	CFTST1	ICF										(2)-						
11	CFTST2	ICF										(2)-						

Key to Above Table
- : Not Polarized
H: Vertical High
L: Vertical Low
M: Vertical Medium

Partition View Controls

<input type="checkbox"/> Select All	<input checked="" type="checkbox"/> Drawer 1
<input type="checkbox"/> Remove All	<input checked="" type="checkbox"/> Drawer 2

Partition Summary by Pool

View	Partition		GCP Pool		zIIP Pool		IFL Pool		ICF Pool	
	No.	Name	Topology	Weight Percent	Topology	Weight Percent	Topology	Weight Percent	Topology	Weight Percent
1	ZOSPRD1	(26)H (1)M (2)L	91.74%	(5)H (2)M (1)L (2)-	62.50%					
2	ZOSPRD2	(1)M (1)L	1.84%	(1)M (1)L (8)-	6.94%					
3	ZOSPRD3	(1)M (2)L	0.92%	(2)M (2)L (6)-	13.89%					
4	ZOSTST1	(2)M	4.59%	(2)M (8)-	11.11%					
5	ZOSTST2	(1)M (1)L	0.92%	(1)M (1)L (8)-	5.56%					
6	VMPRD1					(4)H (1)M	50.00%			
7	VMPRD2					(1)M (2)L	10.00%			
8	VMSTST1					(3)H (1)M (1)L	40.00%			
9	CFPRD1							(5)-	n/a	
10	CFTST1							(2)-	50.00%	

The **Topology** window is accessed from the **Partition Detail Report** window by clicking the **Topology** button. Note: This button only appears if the LPAR host is a z17 or z16 model and topology information is present. The example below represents a z16 with 2 drawers and 11 active partitions.

When the selected EDF interval indicates 'Topology Changed', the following will appear above the table. Note: Topology configuration changed during measurement interval.

Caution: If the partition configuration has been altered, the **Topology Report** is no longer applicable and therefore, cannot be displayed.

While the **Topology** window is open, the **Partition Detail Report** window remains visible but with most of its function disabled (configuration changes are not allowed). The order in which partition associated logical CPs are displayed can be changed from the **Partition Detail Report** window using the **Table View Controls** radio buttons:

- **With Parent GP:** Each GP partition is immediately followed by any associated zAAP, zIIP, IFL, and ICF partitions.
- **Separate by Pool:** All GP partitions are listed together, followed by all zIIP partitions, all IFL partitions, and finally all ICF partitions.

The **Topology Report** displays logical CP information for each active partition. The Logical CP classification keys are displayed below the table, as follows:

- Not Polarized (—)
- Vertical High (H)
- Vertical Low (L)
- Vertical Medium (M)

The table portrays how the partition's logical CPs and their classification are distributed:

- Across the installed drawers (maximum of 4 on z16 and z17)
- Across the 4 Dual Chip Modules (DCMs) on each drawer
- Across the 2 chips on each DCM

Note that the sample output shown in this section is unrelated to any EDF or sample study file shipped with zPCR.

Logical CP classifications used in the table are preceded by the count of LCPs of that classification. For example, “(8)H” is interpreted as eight Vertical High LCPs for the partition on that Drawer, DCP and Chip.

Partition View Controls group box

- The **Select All** button causes all drawers to be selected and all partitions to be selected (default scenario). In this state, individual drawers and/or partitions may be removed as desired.
- The **Remove All** button causes all drawers to be unselected and all partitions to be unselected. In this state individual drawers and/or partitions can be added as desired.

When a specific drawer is checked, all the partitions that have one or more LCPs on that drawer are displayed. If one of those partitions becomes unselected, the drawer itself becomes unchecked indicating that the drawer configuration is incomplete. The Drawer and View check boxes determine which partitions are included on the **Topology Report**.

Partition Summary by Pool group box

This area summarizes each partition's logical CP classification and relative weight for each CP pool. The View check boxes to the left determine which partitions are included on the ***Topology Report***.

Output is available by clicking either the **HTM** or the **CSV** tool bar icons.

Focus will be restored to the ***Partition Detail Report*** by clicking **Return**.

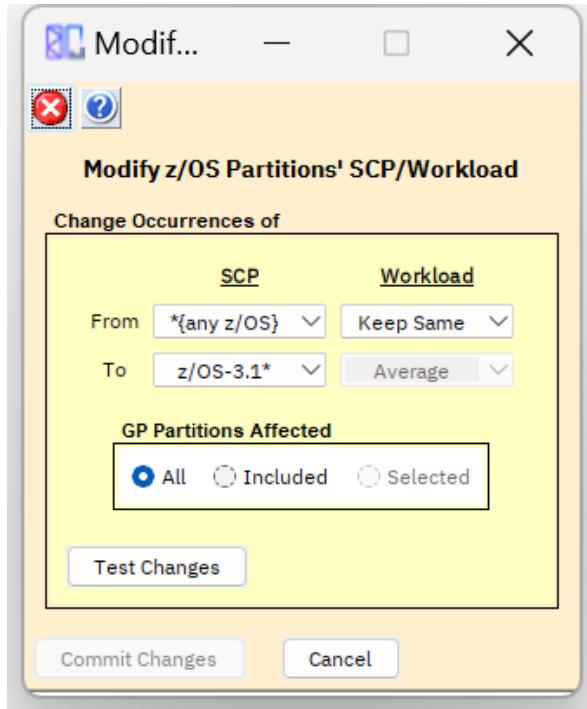
Note: The ***Topology Report*** window reveals the actual distribution of RCPs across drawers for the EDF interval selected. However, **zPCR** capacity results are based on the estimated distribution of RCPs across drawers accessed via the ***Partition Detail Report*** window (see [Book and Drawer Considerations](#)).

System Recovery Boost Considerations

If a **System Recovery Boost (SRB)** interval was selected, the ***Topology Report*** window will appear. Generally, it is unwise to use such data for planning purposes.

Modify SCP/Workload

LPAR Configuration Capacity Planning



The **Modify Partitions** window is accessed via the **Partition Detail Report** window by clicking **Modify SCP/Workload** under **Edit** on the menu-bar. The intent of this window is to provide an easy means to test the effect on capacity for the currently defined LPAR configuration by changing the SCP and/or workload in one or more partitions. Changes are limited to z/OS GP partitions and any associated zAAP or zIIP partitions. The target SCP can be any z/OS version supported by **zPCR** for the LPAR host processor. The target workload is limited to one of the z/OS workload categories.

While the **Modify Partitions** window is active, the **Partition Detail Report** window will remain visible, with the capacity results reflecting the effects of any changes that have been applied. While most of the function on the **Partition Detail Report** window is disabled, **HTML** output and **Graph** creation are allowed.

At any time the changes being tested can be committed to the current study using the **Commit Changes** button. If the changes being tested are not intended to be committed, **Cancel** can be used to restore the original LPAR configuration.

Defining Changes

Changes to the SCP and/or workload can be applied to all qualifying GP partitions in the current LPAR configuration using the dropdown lists as follows:

- **From SCP** change can be set to:
 1. **Keep Same** – no change is to be made
 2. *** {Any z/OS}** (default)
 3. **A specific z/OS version** that is supported on the current LPAR host processor.
- **From Workload** change can be set to:
 1. **Keep Same** (default) – no change is to be made
 2. *** {Any} workload**
 3. **Any specific z/OS workload category**
- **To SCP** change can be set to **Any specific z/OS version** supported by **zPCR** for the current LPAR host processor.
- **To Workload** change can be set to **Any of the z/OS workload categories**.

Only GP partitions with z/OS assigned are affected by the ***Modify Partitions*** function (IFL and ICF partitions will not be affected). When a GP partition has an associated zAAP or zIIP partition, a change will only be applied if the **To SCP** is a z/OS version that supports zAAPs and zIIPs (i.e., z/OS-1.6 or later). Otherwise, these partitions will not be modified.

In the ***GP Partitions Affected*** group box changes can be directed to the following, using the associated radio buttons,:

1. All GP partitions
2. Only GP partitions that are currently included in the study
3. GP partitions that have been user selected before opening the ***Modify Partitions*** window.

Why the “Modify SCP/Workload” function can be helpful

1. When capacity planning for:
 - IBM z17, z/OS partitions must be specified with z/OS-2.4 or later.
 - IBM z16, z15, or z14, z/OS partitions must be specified with z/OS-2.1 or later.
 - z13, z/OS partitions must be specified with z/OS-1.13 or later.
 - z12, z/OS partitions must be specified with z/OS-1.10 or later.
 - z11 or z10, z/OS partitions must be specified with z/OS-1.7 or later (older releases are not supported).

When the current processor LPAR configuration includes older z/OS releases, this function provides the ability to easily change all the z/OS partitions to one of the newer releases, thus allowing newer processor models to be assigned as the new LPAR host without incurring errors.

2. Regardless of the workload(s) used for capacity planning, it is often useful to test the effect of alternate workloads to determine their effect on capacity. This function provides the ability to change all the z/OS partitions to another workload. The capacity result can be reviewed without committing the changes to the study.

Such changes should be made to both the current LPAR configuration and the planned new LPAR configuration. The effect on absolute capacity results due to these changes should be ignored. Rather, the It is the effect on the capacity relationship between the two configurations that is important to note.

Window Controls

Click the **Test Changes** button to temporarily apply the intended SCP/workload changes to the current LPAR configuration and update the capacity results accordingly in the **Partition Detail Report** window.

Click the **Commit Changes** button to make the current (test) LPAR configuration the current study. A dialog will be offered allowing the current study to be saved under a new name. The **Modify Partitions** window will be closed and the **Partition Detail Report** window's full function will be restored, reflecting the last modified state for the LPAR configuration.

Click the **Cancel** button to restore the original LPAR configuration. The **Modify Partitions** window will be closed and the **Partition Detail Report** window's full function will be restored, reflecting the original LPAR configuration.

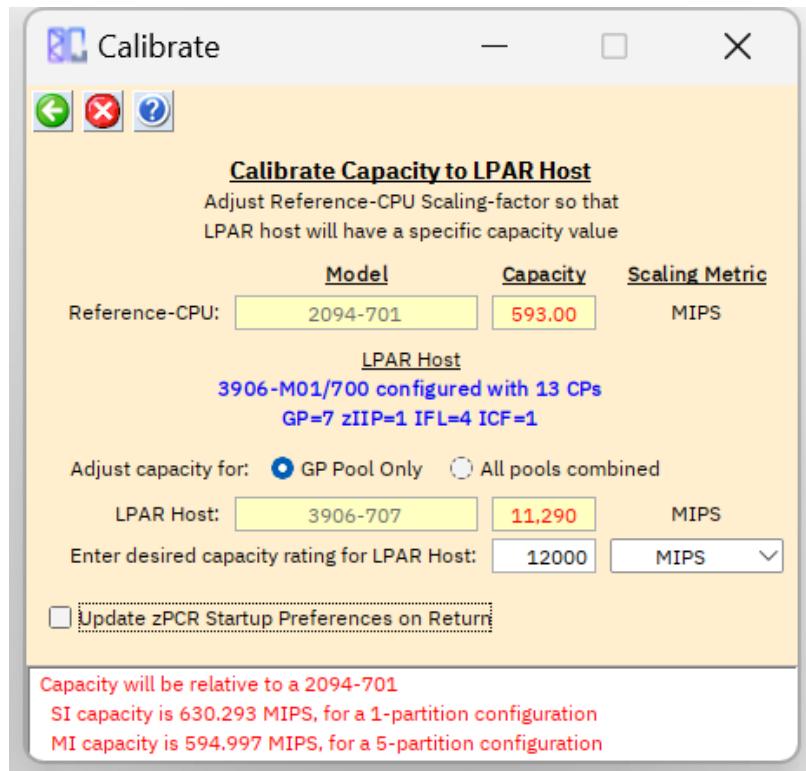
Calibrate Capacity to LPAR Host

LPAR Configuration Capacity Planning

For capacity planning purposes, it is often desirable to define the LPAR host and partition configuration for a currently installed system such that its capacity result comes out to be some specific value. While one would like to assign the LPAR host as the **Reference-CPU**, doing so is not always possible since a 1-way processor is required. A technique adjusting the 1-way **Reference-CPU** scaling-factor can be used to resolve this dilemma.

Using the **LPAR Configuration Capacity Planning** function, once the LPAR host has been defined with a valid partition configuration, calibration is available. Calibration simply adjusts the **Reference-CPU**'s current scaling factor such that the LPAR host capacity projection is the desired value. The scaling-metric may also be changed.

The **Calibrate** function can only be accessed for the 1st configuration on the LPAR configuration tree. To open the **Calibrate** window, from the **Host Summary Report** window click the **Calibrate Capacity** button, or from the **Partition Detail Report** widow menu bar click **Edit → Calibrate Capacity**.



The **Calibrate** window reflects the current **Reference-CPU** settings, and a summary of the currently defined LPAR host's CP configuration. Also displayed is the current capacity result for the LPAR host. Below, a single entry field (white background) provides the means to enter the desired capacity value for this LPAR host. Click on the field, key in a value, and press **Enter**. If a capacity value is entered, the scaling-metric may also be changed. The message box below will reflect what the new **Reference-CPU** metrics will be.

The desired capacity value can represent either that of the General Purpose CP pool only, or that of the entire CPC (i.e., combination of all CP pools). Radio buttons are available for this selection (default is General Purpose pool only). Upon Return, the new **Reference-CPU** metrics will be applied and the current LPAR host configuration capacity result will be recomputed to reflect this setting.

Calibrate does not require the **Reference-CPU** to be preset with any particular scaling-factor or scaling-metric. Even if default values are currently assigned (i.e., **1.00 ITRR**), the LPAR host capacity can be designated with a much larger MIPS value. The only limit applied is the maximum value that is allowed for the **Reference-CPU** scaling-factor.

Once a valid capacity number has been entered, the Return toolbar icon is enabled. By clicking Return, the ratio between the projected LPAR capacity and the desired LPAR capacity is applied to the **Reference-CPU** scaling-factor, and control is returned to the calling window. Now the adjusted **Reference-CPU** scaling-factor appears at the top of the calling window, and the LPAR host configuration capacity result will reflect the desired value.

The new **Reference-CPU** metrics can be saved as the **Preferences** setting by checking the **Update zPCR Startup Preferences on Return** checkbox. The Preferences setting will be updated when Return is clicked.

Click the Cancel toolbar icon to exit without making any **Reference-CPU** scaling-factor adjustment.

Once the new **Reference-CPU** setting have been established, you should continue to use it unmodified, to model potential upgrades or new processors with the same or modified partition definitions. By retaining the same **Reference-CPU** settings, capacity results between each modeled configuration can be directly compared, since they will all be based on the same capacity assumption. In addition, since a 1-way must be used as the **Reference-CPU**, capacity results will be presented with a perspective that can be more easily understood.

Utilized Capacity Report

LPAR Configuration Capacity Planning

The **Utilized Capacity Report** window is accessed from the **LPAR Host and Partition Configuration** window by clicking the **Partition Utilized Capacity** button in the **Capacity Reports** group box.

Utilized Capacity Report

Utilized Capacity vs Weighted Capacity

Based on LSPR Data for IBM Z Processors
Study ID: Not specified
#1 Configuration #1

Description: Created from EDF XYZ 3906.edf for CPC00001 interval #1: Date=2020-02-03 Time=08:00:00 Length=00:15:00
z14 Host = 3906-M01/700 with 7 CPs: GP=7
8 Active Partitions: GP=8

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration							Weight Based			Utilization Based		Ratio
No.	Type	Name	SCP	Workload	Mode	LCPs	Weight%	Minimum Capacity	Pool Capacity	Busy	Capacity	
							Weight%	Minimum Capacity	Pool Capacity	Busy	Capacity	
1	GP	CICSA	z/OS-2.4	Average	SHR	7	43.00%	4,607	10,745	44.82%	4,815	1.045
2	GP	BATCHA	z/OS-2.4	Average	SHR	7	19.50%	2,089	10,745	20.72%	2,227	1.066
3	GP	BATCHB	z/OS-2.4	Average	SHR	2	3.20%	352	10,745	2.12%	228	0.649
4	GP	TESTB	z/OS-2.4	Average	SHR	2	1.20%	132	10,745	0.67%	72	0.543
5	GP	TESTIMS	z/OS-2.4	Average	SHR	5	3.60%	392	10,745	2.88%	310	0.790
6	GP	CICSB	z/OS-2.4	Average	SHR	7	23.10%	2,475	10,745	22.93%	2,464	0.996
7	GP	IMSA	z/OS-2.4	Average	SHR	5	4.90%	534	10,745	4.19%	451	0.844
8	GP	TESTCICS	z/OS-2.4	Average	SHR	2	1.50%	165	10,745	0.84%	90	0.547

Summary by Pool

CP Pool	Number of			Pool Capacity	Busy	Utilized Capacity	Ratio
	RCPs	Partitions	LCPs				
GP	7	8	37	10,745	99.18%	10,656	0.992
zAAP	0	0	0	0	0.00%	0	0.000
zIP	0	0	0	0	0.00%	0	0.000
IFL	0	0	0	0	0.00%	0	0.000
ICF	0	0	0	0	0.00%	0	0.000
Totals	7	8	37	10,745		10,656	0.992

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

This report is only enabled when the entire LPAR configuration was created using EDF or RMF input since these are the only sources where actual utilization values are known.

The intent of this report is to provide a comparison of the actual capacity used (for the EDF or RMF interval loaded) to the minimum capacity that is available to the partition based on each partition's relative weight assignment. In addition, comparisons are available for each CP pool, as well as for the overall LPAR host.

The title area includes the description set on the **LPAR Host and Partition Configuration** window, followed by the number of active partitions with their distribution across RCP pools. The basis for all capacity results is also provided (i.e., the current **Reference-CPU** and its scaling-factor/metric).

Partition Reporting Section

By default, the GP partitions are listed first, followed by zAAP, zIIP, IFL, and ICF partitions. A unique background color is assigned to distinguish each of the partition types. This color key is applied consistently throughout **zPCR**.

GP partition CPs always operate from a single pool of RCPs. On z9 and later processor models, zAAP partitions, zIIP partitions, IFL partitions, and ICF partitions each operate from a separate pool of RCPs.

The partition table columns include:

LPAR Configuration metrics

- **No.** Automatically assigned sequential partition number
- **Type** CP pool to which the partition is assigned
- **Name** Partition name
- **SCP** Operating system
- **Workload** Workload category
- **Mode** Partition is dedicated or shared
- **LCPs** Number of active Logical CPs

Weight Based capacity metrics

- **Weight %** Relative weight within the partition's CP pool (SHR only)
- **Minimum Capacity** Capacity realizable at 100% sharing contention (**Minimum Capacity** value from the **Partition Detail Report** window)
- **Pool Capacity** Capacity determined for the entire pool, based on the LPAR configuration defined. This is the capacity to which EDF or RMF utilization values will be applied.

Utilization Based capacity metrics

- **Busy** The partition utilization determined from EDF or RMF for the measurement interval that was captured.
- **Capacity** The capacity actually utilized by the partition (busy times the pool capacity).

Ratio The computed ratio between the actual capacity used and the **Minimum Capacity** that should always be available to the partition, based on its weight percent.

Summary by Pool

The **Summary by Pool** group box reports by CP pool the actual capacity used and the **Minimum Capacity** based on the partition weight percentages. The ratio between these is also provided. A totals line draws a similar comparison for the entire LPAR host as a whole.

System Recovery Boost Considerations

If any partition in the selected interval has **System Recovery Boost** activity, the **Utilized Capacity vs Weighted Capacity** window is disabled.

For detail concerning SRB support see [System Recovery Boost](#).

Window Controls

No change capability is provided on the ***Utilized Capacity Report*** window itself. Changes can only be made via the ***LPAR Host*** window the ***Partition Definition*** window, or the ***Partition Detail Report*** window.

The only change allowed without disabling this report is the modification of the ***Reference-CPU*** setting, in which case all capacity results will be adjusted accordingly.

If any modifications are made to the LPAR host or to any individual partition, the report is no longer accurate and the ***Utilized Capacity Report*** is disabled. Utilization values can only be representative for the original partition definitions. Any changes to the LPAR host or to partition definitions would affect the ***Weight-Based*** metrics thus affecting the comparison to the ***Utilization-Based*** metrics. Partitions that have been excluded on the ***Partition Detail Report*** window would also compromise the accuracy of the comparison.

Partition utilization values from EDF of RMF are saved in a ***zPCR*** study file. When a study is reloaded, and no configuration changes have been made, the ***Utilized Capacity Report*** remains available.

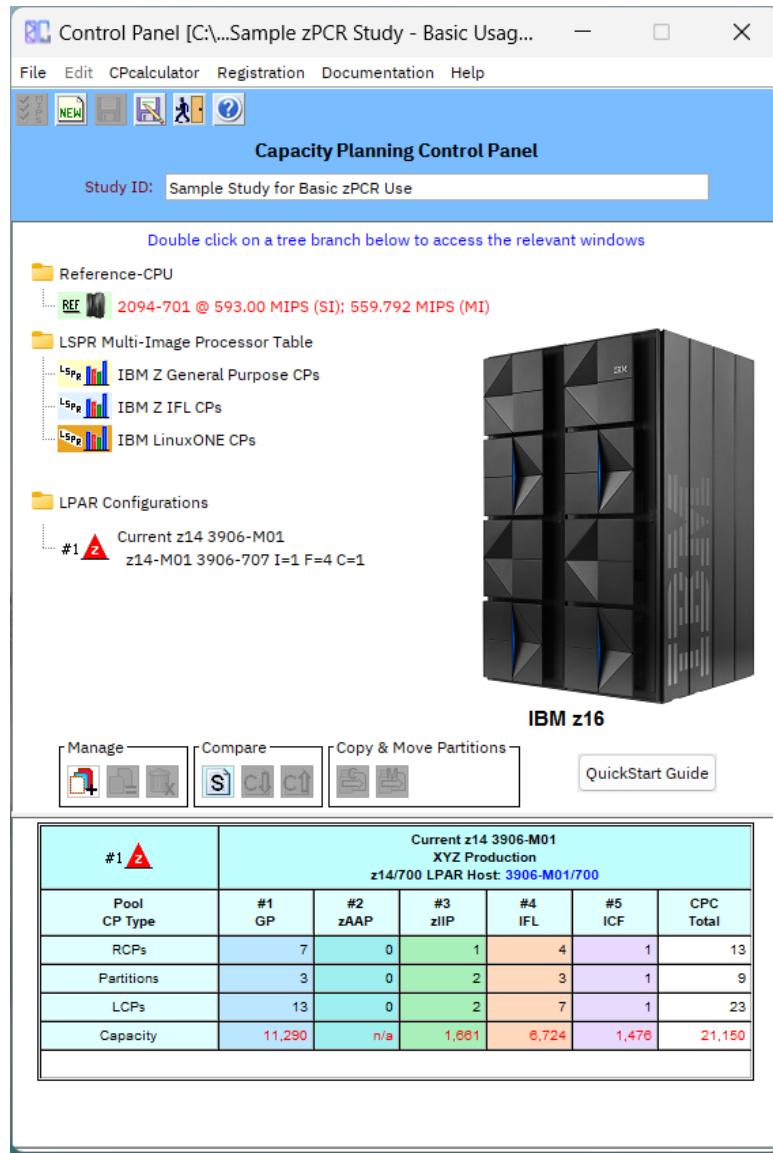
Click the **Return** toolbar icon to return to the ***LPAR Host and Partition Configuration*** window.

Click the **HTM** toolbar icon to create an HTML file with the report tables.

Click the **CSV** toolbar icon to create a CSV file with the report tables.

Click the **Help** toolbar icon to access context sensitive help for this window.

Control Panel - Advanced Usage



The **Control Panel** window is displayed immediately following the **Logo** window (once the Java code is compiled). This is the primary window to access all the function of **zPCR**, including

1. Provide a description for the Study ID.
2. Set the **Reference-CPU** metrics  ITRR Scaling-Factor and Metric are used throughout all **zPCR** function.
3. Review **LSPR Multi-Image capacity** for **IBM Z** and **LinuxONE** processors.
 - IBM Z GP CPs  (z/OS, z/VM, z/VSE, KVM, or Linux assumed).
 - IBM Z IFL CPs  (z/VM, KVM or Linux assumed).
 - IBM LinuxONE CPs  (z/VM, KVM or Linux assumed).
4. Define or review an **LPAR Partition Configuration** 

For detail about defining an LPAR configuration, see [Control Panel](#). Note: The **Sample zPCR Study - Basic Usage.zpcr** study file included with the **zPCR** package, is the source used for the examples shown in that and the ensuing chapters.

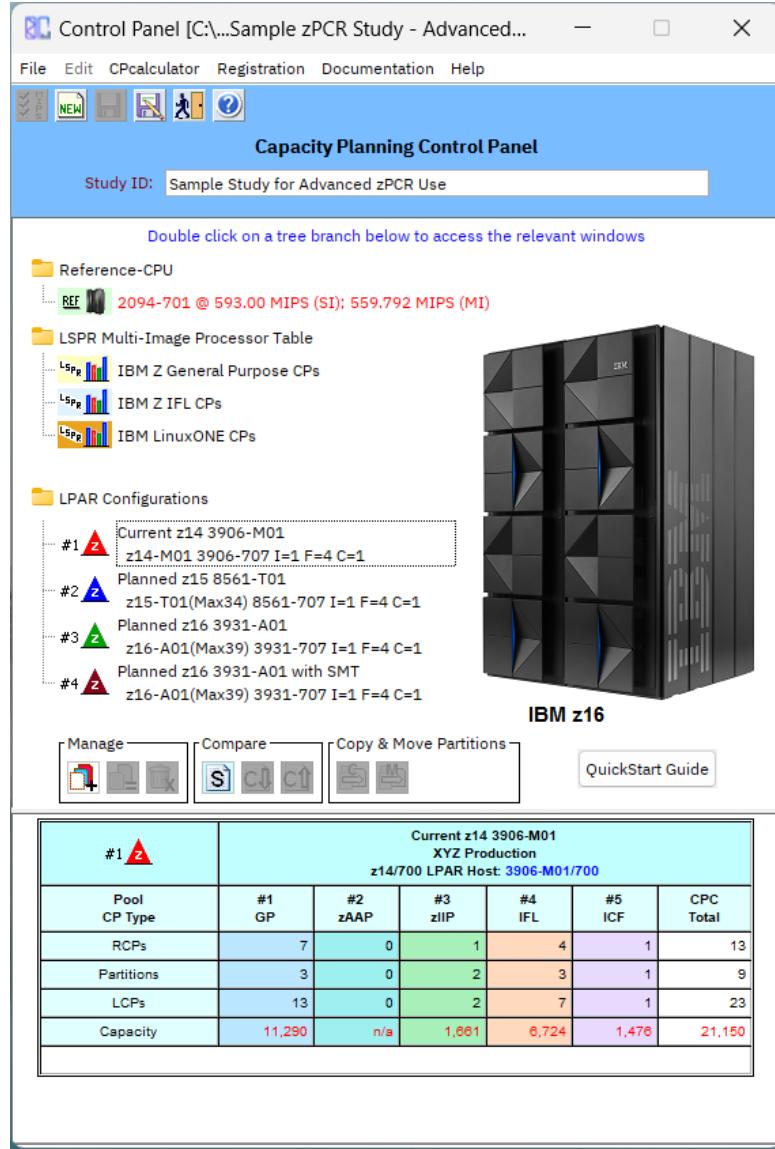
When multiple LPAR Configurations are defined, additional **zPCR** function is available. Shown below is a sample study with 4 LPAR configurations defined. This study will be used to demonstrate the advanced function available. Note: The **Sample zPCR Study - Advanced Usage.zpcr** study file, included with the **zPCR** package, is the source used for the examples shown for this and the ensuing chapters.

zPCR study files can be loaded by dragging them onto the **Control Panel** window.

This section discusses the following **zPCR** capabilities:

- Add, Clone, or Delete LPAR Configurations
- Compare capacity results between LPAR Configurations.
- Copy or Moving partitions between LPAR configurations.

IBM zPCR User's Guide



For capacity comparisons to be useful, the partition configurations being compared should both contain some or all the same partitions (i.e., in terms of partition type, name, SCP, and workload). Generally, the capacity differences for a new LPAR configuration, over that of the current LPAR configuration will be a result of one or more of the following configuration changes:

1. Change of the LPAR host processor family and/or model
2. Change of the LPAR host processor CP configuration
3. Change of the way that one or more partitions are defined
4. Add one or more new partitions
5. Delete one or more current partitions.

The **QuickStart Guide** button reveals a short paper describing the process using **zPCR** to define a current and an alternate LPAR configuration and make capacity comparisons between them.

Summary of Control Panel Function

1. The **Reference-CPU** metrics can be set.
2. Of the LSPR tables, the **Multi-Image LSPR Capacity Ratios** table is available for General Purpose CPs or for IFL CPs. The Multi-Image table is the preferred table with which to generalize capacity. Alternatively, a **Single-Image LSPR Capacity Ratio** tables can be temporarily displayed.
3. Multiple LPAR configurations (limited to 10) with a 1st level name such as **Configuration #1** and **Configuration #2** can be defined (these can be renamed, if desired; see [Renaming LPAR Configurations](#)). Several report windows are available with which to compare capacity results between various configurations. A 2nd level name generated by zPCR appears below.
4. The **Reference-CPU** scaling-factor can be calibrated for the 1st LPAR configuration only, for the purpose of producing a desired total capacity result for that LPAR configuration. Once calibration has been done, the resulting **Reference-CPU** scaling-factor will be used for all the LPAR configurations that are defined. Note that this calibration is applied from windows other than the **Control Panel**.
5. The **Workloads** window must be accessed from the **LSPR Processor Table**.

The **Control Panel** window provides the capability to define multiple LPAR configurations and to make direct comparisons between those configurations. The control panel is intended to remain visible while setting the **Reference-CPU**, viewing **LSPR Tables** and defining **LPAR Configurations**. Windows for any of these functions can be displayed simultaneously.

Notes concerning window size and scroll bars:

- If the Control Panel does not fit on a monitor's screen, a scroll bar for the entire window will appear on the right.
- When more than 6 LPAR configurations are defined, a scroll bar will appear in that area, providing access to the 7th and above.
- When text appears at the bottom of the window, a scroll bar may appear in that area.

The only **Reference-CPU** modification allowed from windows other than the **Control Panel** is by using the **Calibrate Reference-CPU** button from either the **Host Summary Report** window or the **Partition Detail Report** window. This button is available for the first LPAR configuration only (i.e., the LPAR configuration named by default **Configuration #1**).

Several windows are available, with which to compare capacity results between LPAR configurations (for detail see [Host Capacity Comparison](#) and [Partition Capacity Comparison](#)). Also, reports are available showing the effect on capacity when ±5% margin-of-error is applied (see [Host Margin-of-Error](#) and [Partition Margin-of-Error](#)). The **Control Panel** window will be locked while comparison windows are open.

The **Study Identification** text can be entered from **Control Panel** window. In addition, to further describe scenarios, an **LPAR configuration description** entry field is available on the **LPAR Host and Partition Configuration** window for each LPAR configuration defined.

The **Control Panel** window is composed of two sections:

1. The top panel is presented in the form of tree structures, with branches representing the controls that are available. Each branch provides access to function previously described in this document.
 - **Reference-CPU** tree: This branch provides access to the **Reference-CPU** settings which control how capacity will be represented throughout. The setting status is shown in the text attached to the branch. The scaling-factor used for the **LSPR Multi-Image Capacity Ratios** table is adjusted down from the **Reference-CPU**. The adjusted value is revealed on the **LSPR Multi-Image Capacity Ratios** table. Double-click the branch icon to open the **Reference-CPU** window.
 - **LSPR Processor Table** tree: Double-click the branch icon to open an **LSPR Multi-Image Processor Capacity Ratios** table.
 - **IBM Z** General Purpose CPs
 - **IBM Z** IFL CPs
 - **LinuxONE** CPs.
 - **LPAR Configurations** tree: Branches are displayed for each LPAR configuration; the first will be named **Configuration #1**, the second, when created, will be named **Configuration #2**, and the third, **Configuration #3**, etc. Each of these branches provide access to separate invocations of the **LPAR Configuration Capacity Planning** function, including the **LPAR host**, the **Partition Definition**, the **Host Summary Report**, and **Partition Detail Report** windows. The discussion below pertains to these branches.
Whenever the mouse is passed over an LPAR configuration branch, its description (if one has been provided) will appear as flyover text.
2. The bottom panel alternates between presenting user instructions and presenting a summary of the currently selected LPAR configuration.

Current z14 3906-M01 XYZ Production z14/700 LPAR Host: 3906-M01/700						
Pool CP Type	#1 GP	#2 zAAP	#3 zIIP	#4 IFL	#5 ICF	CPC Total
RCPs	7	0	1	4	1	13
Partitions	3	0	2	3	1	9
LCPs	13	0	2	7	1	23
Capacity	11,290	n/a	1,661	6,724	1,476	21,150

When the **Control Panel** window is initially presented, a single icon representing the branch for LPAR Configuration #1 is shown under the **LPAR Configurations** tree. This icon has a name of **Configuration #1**, and its configuration is initialized as "undefined".

There are several ways that an undefined configuration may be initialized:

1. To manually define a configuration, double click the branch icon. The normal **LPAR Configuration Capacity Planning** function windows will appear with which to define the LPAR host and its partition configuration.
 - a. The LPAR configuration may be defined by dragging a previous **zPCR** study file on top of the LPAR configuration branch icon. All of the defined LPAR configurations will be loaded (i.e., **Configuration #1**, **Configuration #2**, **Configuration #3**, etc).

In lieu of dragging a study file, it can be loaded by using **File → Load** from the menu-bar.

2. The LPAR configuration may be defined by dragging an EDF on top of the **LPAR configuration** branch icon. The LPAR host and the partition configuration will be taken from the EDF interval selected.
3. The LPAR configuration may be defined by dragging an RMF report (text file) on top of the **LPAR configuration** branch icon. The LPAR host and the partition configuration will be taken from the RMF interval selected.

Notes:

Dragging an EDF, RMF, or a **zPCR** study file to an undefined LPAR configuration icon will cause that configuration to become defined.

If the 1st LPAR icon ("Configuration #1") is undefined, it can be defined by dragging an EDF, RMF, or study file to that icon or to the area beneath it.

If more than one LPAR icon exists (defined or undefined), an additional LPAR configuration will be created when dragging an EDF, RMF, or study file to the area immediately below; the existing LPAR icons will be unaffected.

If an EDF, RMF, or study file is dragged to an LPAR configuration that is already defined, it is then considered to be a source for copying partitions into the existing configuration.

When dragging a **zPCR** study file to the area below existing LPAR icons, new icons will be created for each LPAR configuration in the study up to the maximum LPAR configurations allowed. Should the number of LPAR configurations in the study cause the maximum of LPAR configurations that can be defined, loading any of the LPAR configurations is prevented.

Once an LPAR configuration icon has become defined, it may be refined by double clicking its icon to access to the normal **LPAR Configuration Capacity Planning** windows. The associated icon and configuration name will appear on each of these windows to help differentiate which LPAR configuration the window belongs to.

Managing LPAR configurations

The **LPAR Configuration** branches are managed using the toolbar icons in the **Manage** group box:



Add toolbar icon: Creates an icon representing the next LPAR configuration, starting with the name **Configuration #1**, **Configuration #2**, etc. Each added icon is initialized as “undefined”. It can be defined in any of the ways that were described for the first LPAR configuration.



Clone toolbar icon: If a single defined LPAR configuration is selected, and room remains for additional LPAR configurations, the Clone icon is enabled. An icon representing the next available LPAR configuration is created with the same name appended with a repetition indicator. Cloning allows you to use one LPAR configuration as a starting basis for another LPAR configuration. Changes can then be made to the LPAR host and/or the partition configuration as needed.



Delete toolbar icon: An LPAR configuration can be deleted by selecting its icon and then clicking the delete button. When the icon contains a defined configuration, a dialog will be presented for confirmation of the delete action. All LPAR configurations below a deleted one will move up. Note that when the last LPAR configuration icon is deleted, the first icon (**Configuration #1**) will appear with its LPAR configuration undefined (equivalent to starting a new study).

Renaming LPAR Configurations

Each LPAR configuration icon has a 2 level name. For the 1st level, a default is assigned at the time the configuration icon is added. This level can be renamed with up to 35 characters. Steps to rename the 1st level:

1. Left-click the LPAR configuration icon to select it
2. Right-click the LPAR configuration to display a pop-up menu
3. Click **Rename Configuration** on the pop-up menu.

The 2nd level name is automatically generated by zPCR at the time the LPAR host is defined. The name is based on the processor model and the real CPs defined. This name cannot be user modified.

Comparing Capacity between LPAR Configurations

Direct comparisons between **LPAR Configuration** branches can be made using the following **Compare** toolbar icons:

S **Summary** toolbar icon: Generates the **Host Capacity Comparison Summary** window, showing the hardware configuration and the capacity results for each defined LPAR configuration. Capacity projections are provided for each CP pool, with a total for the overall CPC, which can be viewed as **Full** capacity (default) or **Single-CP** capacity. For detail concerning this summary report, see [Host Capacity Comparison Summary](#).

C↓ and **C↑** **Compare** toolbar icons: Generates various reports comparing capacity between any two LPAR configurations. Two LPAR configuration icons must be present, and both must have LPAR configurations defined. The arrow on the toolbar icon indicates the order in which the LPAR configuration branches will be compared. Select either one or two LPAR configuration icons and click one of the compare icons to view the **Host Capacity Comparison** window and subsequently, the various **Partition Capacity Comparison** windows. If only one LPAR configuration was selected, the one adjacent to it will become the focus of the comparison. For detail concerning these capacity comparison reports, see [Host Capacity Comparison](#) and [Partition Capacity Comparison](#).

Migrating Partitions between LPAR Configurations

Partitions can be copied from one LPAR configuration to another using the toolbar icons in the **Migrate & Analyze** group box.

C→ **Copy** and **M→** **Move** toolbar icons: Provide the ability to modify an LPAR configuration with partitions from another LPAR configuration, and assess the overall effect on capacity. Select a single defined LPAR configuration that is to be the target of the migration and click **Copy** or **Move**. Partitions can be migrated to the target from any other defined LPAR configuration. In addition, partitions existing on the target can be modified. As each change is made, capacity values are recomputed and compared back to the original partition capacity values. Changes to the target configuration can either be retained (**Commit**) or discarded (**Undo**). For detail concerning this capability, see [Copy Partitions to an LPAR Configuration](#) or [Move Partitions between Configurations](#).

Note

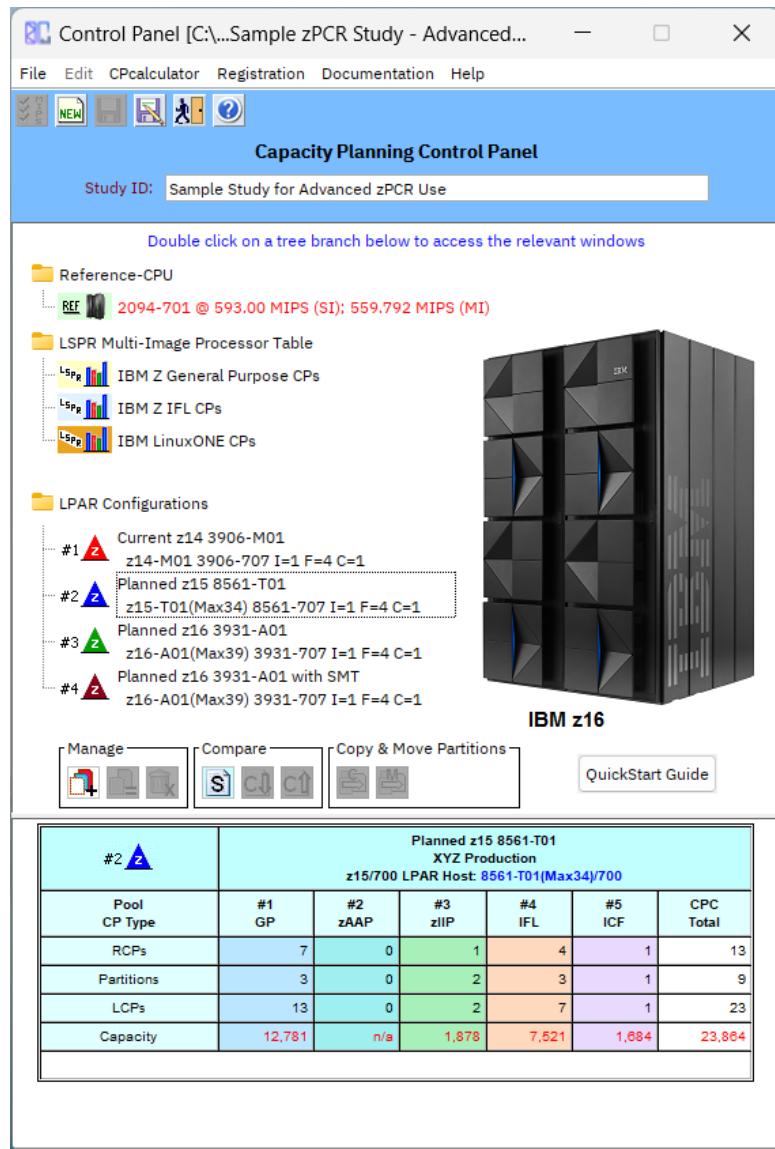
The **Sample zPCR Study – Advanced Usage** file included with the **zPCR** package, is the source used for the examples shown in the **Control Panel** chapter and following sections.

Menu-bar

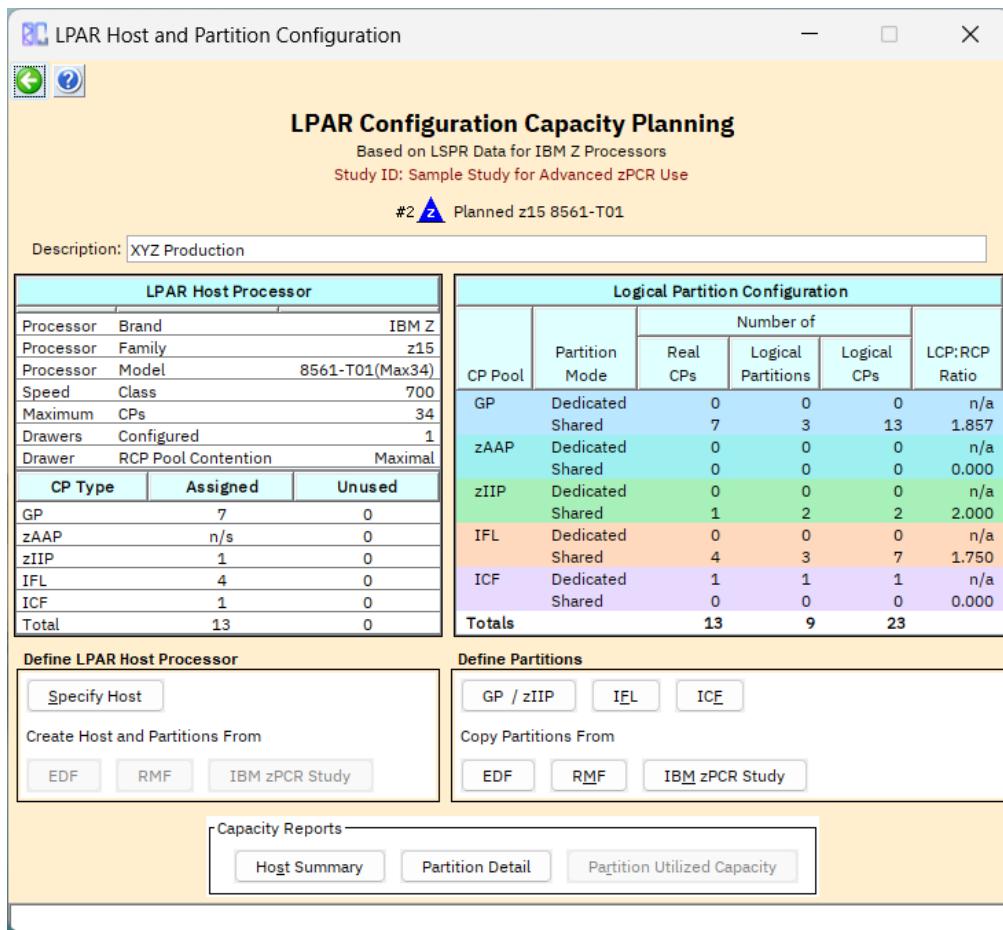
File	
New	Start a new (untitled) study.
Load	Open a previously saved study.
Save	Save this study.
Save as...	Save this study with a new name.
Up to 10 filenames	Recent study files are listed for possible loading.
Exit (<i>Ctrl+E</i>)	Terminate zPCR execution.
Fast Exit (<i>Ctrl+Q</i>)	Terminate zPCR execution immediately
Edit	
Preferences	Personalize zPCR startup settings (see Preferences).
CPcalculator	Capacity planning calculator function (see CP Calculator)
zAAP Capacity	Capacity estimator for migration to a zAAP or zAAP on zIIP configuration (see zAAP Capacity Estimator Input)
Registration	
Remove	Delete your CPS tool registration (disables zPCR and any other CPS tools installed). A new registration will be required for continued tool usage.
Documentation	Various supporting documentation, including:
	zPCR NEWS file
	zPCR User's Guide (if copied to the zPCR Documentation directory as zPCRUG.PDF)
	LSPR FAQ
	LSPR Workloads
	LSPR Document
	HiperDispatch Consideration
	zAAP/zIIP Considerations
	zAAP White Paper
	Concerning Accuracy
	Obtain CP3KEXTR
	Obtain CP3KVMXT
Help	
Context Help (F1)	
QuickStart Guide	Provides guidance on using zPCR .
Check for updates	Verifies that the latest version of zPCR is being used.
About zPCR	Display Logo window information including the of zPCR version.

Defining and Managing LPAR Configurations

On the **Control Panel** window, for each LPAR configuration defined, a unique icon and name is displayed. The first icon default name will be **Configuration #1** and the subsequent icons will be named **Configuration #2**, **Configuration #3**, etc. (the icons may be renamed). Double clicking any of these icons will provide access to the various **LPAR Configuration Capacity Planning** windows represented by that icon.



While defining or altering any single LPAR configuration, each window will display that LPAR configuration's icon and name in the title area to assist in associating the various windows presented.



An LPAR configuration description input field is available on the **LPAR Host and Partition Configuration** window, serving as additional documentation detail. When a configuration is loaded from EDF or RMF, its file name and information concerning the CPC ID and the chosen interval will be placed in the description field. Each LPAR configuration's description field may be altered at any time to the text desired.

The **Utilized Capacity Report** window can be displayed for any LPAR configuration when the entire configuration was obtained from EDF or RMF (click the **Partition Utilized Capacity** button on the **LPAR Host and Partition Configuration** window). Once any LPAR host or partition configuration changes have been made, this report is no longer valid, and therefore, will no longer be available.

The **LPAR Configuration Capacity Planning** windows of separate LPAR configurations can be open at the same time, allowing changes to be made that are appropriate to the specific window. Use the LPAR configuration icon and description for guidance to assure that changes are being made to the LPAR configuration that was intended.

IBM zPCR User's Guide

Partition Detail Report

Based on LSPR Data for IBM Z Processors

Study ID: Sample Study for Advanced zPCR Use

#2 Planned z15 8561-T01

Description: XYZ Production

z15 Host = 8561-T01(Max34)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1

9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration

Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Include V	Partition Identification				Partition Configuration										
	No.	Type	Name	SCP	Assigned Workload	Mode	Logical CPs	Weight	Weight Percent	Capping		SMT		Capacity	
									INIT	ABS	V	Benefit	Minimum	Maximum	
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	6,950	11,063
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	3,964	7,361
<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,267	1,900
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1,867	5,201
<input checked="" type="checkbox"/>	4	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	612	1,835
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	4,808	7,513
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	2,404	3,756
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	309	1,932
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a					1,684	1,684	

Table View Controls

Display zAAP/zIIP/IFL/ICF Associated Partitions

With Parent GP Separate by Pool

Show GP Pool Specialty Pools

All Partitions GP zAAP zIIP

Includes Only IFL ICF

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED LCPs	SHR		Sum of Weights	SMT Benefit	Capacity Totals
				LCPs	LCP:RCP			
GP	7	3	13	1,857	1,300			12,781
zIIP	1	2	2	2,000	600			1,878
IFL	4	3	7	1,750	625			7,521
ICF	1	1	1					1,684
Totals	13	9	22					23,864

Host Summary SMT Benefit LCP Alternatives zAAP/zIIP Loading HiperDispatch

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

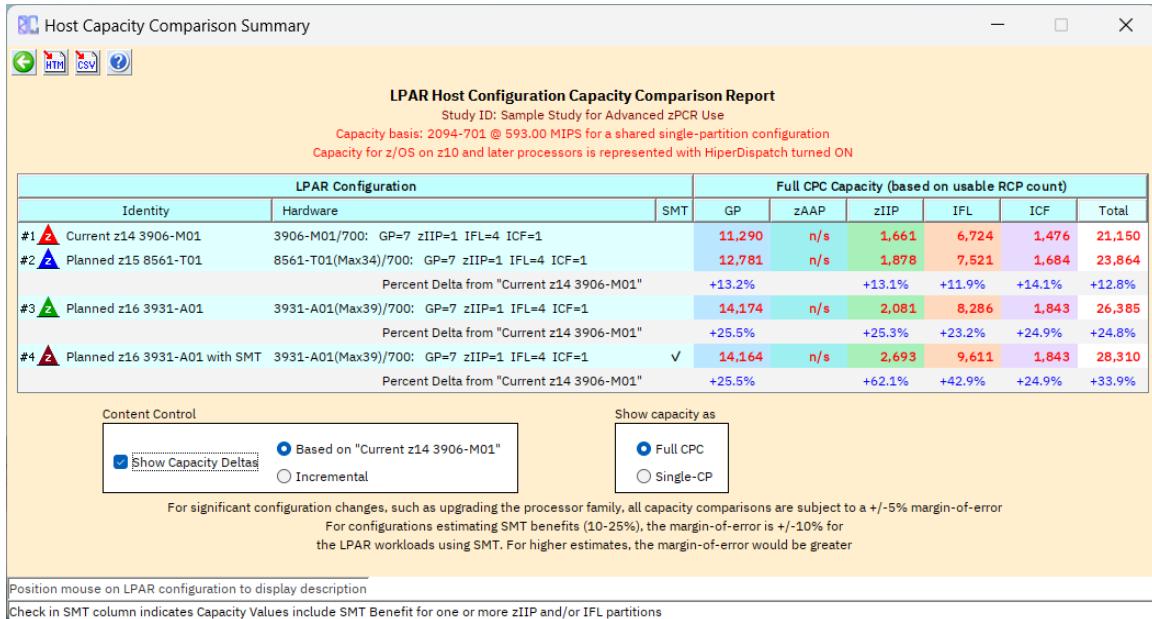
Whenever the **Host Capacity Comparison Summary** window or any of the **Capacity Comparison** windows are open, any open **LPAR Configuration Capacity Planning** windows are hidden. Once these capacity windows are closed, any open **LPAR Configuration Capacity Planning** windows are restored.

Clicking the **Double Return** toolbar icon will return directly to the **Control Panel** window, closing all windows relative to this particular LPAR configuration.

When the **Exit** toolbar icon on the **Control Panel** window is clicked, all open windows, including the **LPAR Configuration Capacity Planning** windows, will automatically be closed.

Host Capacity Comparison Summary

Once any LPAR configurations have been defined, the **S** **Summary** toolbar icon on the **Control Panel** window is activated. The **Host Capacity Comparison Summary** window is displayed by clicking this icon. While the **Host Capacity Comparison Summary** window is being displayed, the **Control Panel** window remains locked. Any open **LPAR Configuration Capacity Planning** windows are hidden to prevent LPAR configuration changes being made while the window is open.



This report presents the capacity projections for each LPAR configuration that has been defined. The LPAR configurations displayed can be controlled; whenever two or more LPAR configurations are selected, only those selected will appear (select multiple LPAR configurations using the **Ctrl** key). Otherwise, all the LPAR configurations appear. LPAR configurations always appear in the order defined.

For each LPAR configuration, its icon and name are provided, along with the processor model information and the number of RCPs configured in each pool. To display the description set on the **LPAR Host and Partition Configuration** window, place the mouse pointer anywhere on that row. The description will appear in the message line at the bottom of the window.

Capacity values are provided for each CP pool. A total capacity value is also provided on the right as the sum of the individual CP pool capacity values.

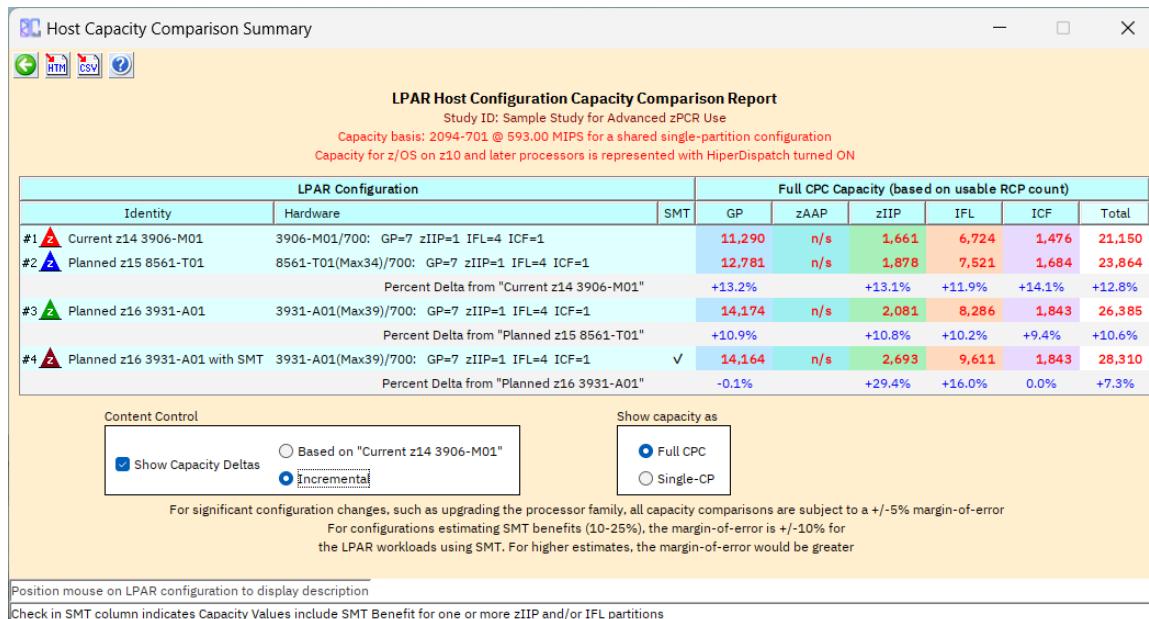
Whenever an IBM z17, z16, z15, z14, or z13 processor has been defined as the LPAR configuration's host, a checkmark appearing in the SMT column indicates that an SMT benefit has been established, resulting in increased capacity for zIIP and/or IFL CPs. Refer to the LPAR configuration's **Partition Detail Report** window to determine what SMT benefits have been defined.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for detail see [zAAP/zIIP Loading](#)). The **GP** row heading will appear as “**GP****” and the General Purpose capacity number will be presented in **brown** rather than the normal **red**. The reason is that General Purpose capacity for the configuration is improved somewhat over that which would be shown using the default zAAP/zIIP loading values. The concern is that, when comparing General Purpose capacity between such configurations, user specified zAAP/zIIP loading values should be identical for matching partitions in both configurations. An explanation will be provided in the message box at the bottom of the window.

Capacity Deltas

Additional rows showing percent deltas for capacity values between LPAR configurations can be added to this window by checking in **Show Capacity Deltas** in the **Content Control** group box. The percent deltas can be displayed in either of two ways.

- Each percent delta is relative to the 1st LPAR configuration.
- Each percent delta is relative to the immediately previous LPAR configuration in the list.

The screenshot shows the 'Host Capacity Comparison Summary' window. The title bar includes buttons for 'HTML' and 'CSV'. The main content is a 'LPAR Host Configuration Capacity Comparison Report' with the following details:
 Study ID: Sample Study for Advanced zPCR Use
 Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
 Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration		Full CPC Capacity (based on usable RCP count)						
Identity	Hardware	SMT	GP	zAAP	zIIP	IFL	ICF	Total
#1  Current z14 3906-M01	3906-M01/700: GP=7 zIIP=1 IFL=4 ICF=1		11,290	n/s	1,661	6,724	1,476	21,150
#2  Planned z15 8561-T01	8561-T01(Max34)/700: GP=7 zIIP=1 IFL=4 ICF=1		12,781	n/s	1,878	7,521	1,684	23,864
		Percent Delta from "Current z14 3906-M01"	+13.2%	+13.1%	+11.9%	+14.1%		+12.8%
#3  Planned z16 3931-A01	3931-A01(Max39)/700: GP=7 zIIP=1 IFL=4 ICF=1		14,174	n/s	2,081	8,286	1,843	26,385
		Percent Delta from "Planned z15 8561-T01"	+10.9%	+10.8%	+10.2%	+9.4%		+10.6%
#4  Planned z16 3931-A01 with SMT	3931-A01(Max39)/700: GP=7 zIIP=1 IFL=4 ICF=1	✓	14,164	n/s	2,693	9,611	1,843	28,310
		Percent Delta from "Planned z16 3931-A01"	-0.1%	+29.4%	+16.0%	0.0%		+7.3%

Content Control
 Show Capacity Deltas Based on "Current z14 3906-M01" Incremental

Show capacity as
 Full CPC Single-CP

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
 For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Position mouse on LPAR configuration to display description
 Check in SMT column indicates Capacity Values include SMT Benefit for one or more zIIP and/or IFL partitions

Full CPC / Single-CP Capacity

Capacity projections on this window can be cycled between **Full CPC** capacity and **Single-CP** capacity, using the radio buttons in the **Show Capacity as** group box. By default, projections are presented as **Full CPC**, representing the total capacity for all CPs involved. **Single-CP** capacity represents the average capacity of each CP (determined by dividing the full capacity by the number of CPs involved). **Single-CP** capacity can be useful for revealing relative engine speed when comparing LPAR configurations where the host processor family is changed

Host Capacity Comparison Summary

LPAR Host Configuration Capacity Comparison Report
Study ID: Sample Study for Advanced zPCR Use
Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

LPAR Configuration		Single-CP Capacity (based on usable RCP count)						
Identity	Hardware	SMT	GP	zAAP	zIIP	IFL	ICF	Total
#1 A Current z14 3906-M01	3906-M01/700: GP=7 zIIP=1 IFL=4 ICF=1		1,613	n/s	1,661	1,681	1,476	1,627
#2 A Planned z15 8561-T01	8561-T01(Max34)/700: GP=7 zIIP=1 IFL=4 ICF=1		1,826	n/s	1,878	1,880	1,684	1,836
#3 A Planned z16 3931-A01	3931-A01(Max39)/700: GP=7 zIIP=1 IFL=4 ICF=1		2,025	n/s	2,081	2,072	1,843	2,030
#4 A Planned z16 3931-A01 with SMT	3931-A01(Max39)/700: GP=7 zIIP=1 IFL=4 ICF=1	✓	2,023	n/s	2,693	2,403	1,843	2,178
	Percent Delta from "Planned z15 8561-T01"		+13.2%		+13.1%	+11.9%	+14.1%	+12.8%
	Percent Delta from "Planned z16 3931-A01"		+10.9%		+10.8%	+10.2%	+9.4%	+10.6%
	Percent Delta from "Planned z16 3931-A01"		-0.1%		+29.4%	+16.0%	0.0%	+7.3%

Content Control

Show capacity as

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error
For configurations estimating SMT benefits (10-25%), the margin-of-error is +/-10% for the LPAR workloads using SMT. For higher estimates, the margin-of-error would be greater

Position mouse on LPAR configuration to display description
Check in SMT column indicates Capacity Values include SMT Benefit for one or more zIIP and/or IFL partitions

Formatted Output

Click the **HTM** toolbar icon to create an HTML file with the report table data.

Click the **CSV** toolbar icon to create a CSV file with the report tables data.

Closing the Window

When the **Return** toolbar icon is clicked, the **Host Capacity Comparison Summary** window will be closed, and control will be restored to the **Control Panel** window. Any open **LPAR Configuration Capacity Planning** windows will be unhidden.

System Recovery Boost Considerations

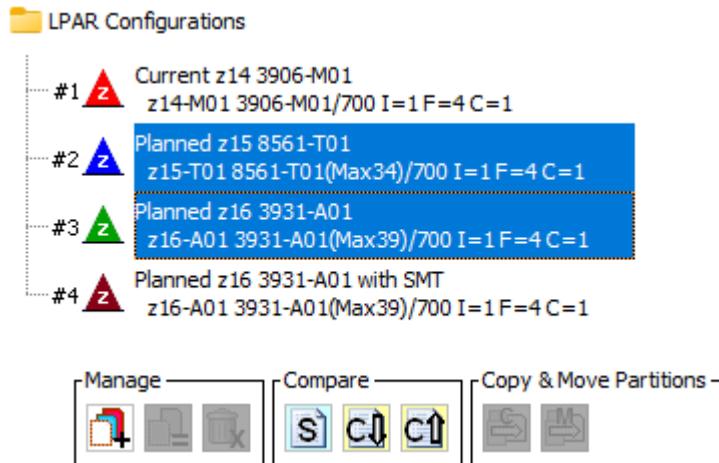
If any partition in either of the selected LPAR configurations Has **System Recovery Boost** activity, the following message will display at the bottom of the window.

Configuration #1 has one or more partitions in **System Recovery Boost**. Don't use this configuration for a processor sizing.

For detail concerning SRB support see [System Recovery Boost](#).

Host Capacity Comparison

The **Host Capacity Comparison** window is accessed by clicking either the the **C↓ Compare Down** or the **C↑ Compare Up** toolbar icon on the **Control Panel** window. Exactly two LPAR configurations must be selected to activate these icons (use the **Ctrl** key to select the 2nd LPAR configuration).



The arrow on the **Compare** toolbar icon indicates the order in which the 2 LPAR configuration branches will be compared. The base of the arrow indicates which LPAR configuration branch to be used as the basis for the comparison, and the head of the arrow indicates which LPAR configuration branch will be used as the focus of the comparison.

Partition Type	Planned z15 8561-T01					Planned z16 3931-A01					Capacity Net Change		
	# of LPs	Usable RCPs	Logical CPs	SHR	LCP:RCP	Full Capacity	# of LPs	Usable RCPs	Logical CPs	SHR	LCP:RCP	Full Capacity	MIPS
GP	3	7	13	1.857	12,781	3	7	13	1.857	14,174	+1,393	+10.9%	
zAAP	0	0	0			0	0	0		2,081	+203	+10.8%	
zIIP	2	1	2	2.000	1,878	2	1	2	2.000	8,286	+765	+10.2%	
IFL	3	4	7	1.750	7,521	3	4	7	1.750	1,843	+159	+9.4%	
ICF	1	1	1		1,684	1	1	1		26,385	+2,521	+10.6%	
Total	9	13	23		23,864	9	13	23					

While the **Host Capacity Comparison** window is being displayed, the **Control Panel** window remains locked. Also, any open **LPAR Configuration Capacity Planning** windows are hidden to prevent LPAR configuration changes from being made while the window is open.

A processor-oriented summary of the two LPAR host configurations is presented in this window. The first LPAR configuration is shown on the left, and the second is shown to the right. Each partition type is listed in a separate row, with a total line at the bottom. The metrics represent the combined values for each partition type, showing the number of partitions defined, the number of RCPs, the number of partition LCPs, the shared LCP:RCP ratio, and the capacity result. Note that, for partition types with only dedicated LCPs, the shared LCP:RCP ratio is omitted.

In the rightmost columns, the capacity **Net Change** (increase or decrease) and the **% Delta** for each is shown. The total row provides the capacity perspective between each of the LPAR configurations as a whole.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for detail see [zAAP/zIIP Loading](#)). The **GP** column heading will appear as "**GP***" and the GP capacity number will be presented in **brown** rather than the normal **red**. The reason is that GP capacity for the configuration is improved somewhat over that which would be shown using the default zAAP/zIIP loading values. The concern is that, when comparing General Purpose capacity between such configurations, user specified zAAP/zIIP loading values should be identical for matching partitions in both configurations. An explanation will be provided in the message box at the bottom of the window.

Full / Single-CP Capacity

Capacity projections on this window can be cycled between **Full** capacity and **Single-CP** capacity, using the radio buttons in the **Show Capacity as** group box. By default, projections are presented as **Full**, representing the total capacity for all CPs involved. **Single-CP** capacity represents the average capacity of each CP (determined by dividing the full capacity by the number of CPs involved). Single-CP capacity can be useful for revealing relative engine speed when comparing LPAR configurations where the host processor family is changed

Partition Capacity Comparisons

The **Partition Capacity Comparison** window is displayed by clicking either the **Minimum Capacity** or the **Maximum Capacity** button, found in the **Comparison Report by Partition** group box. The **Host Capacity Comparison** window remains open while the new window appears, and will reveal any capacity changes due to adjustments made from it. Note that the **Partition Capacity Comparison** window is not available when viewing **Single-CP** capacity. For detail see [Partition Capacity Comparison](#).

Margin-of-Error Consideration

Click the **Consider Margin-of-Error** button to display the **Host Margin-of-Error** window. This window will reveal the capacity results with the full 5% margin-of-error applied. For detail see [Host Margin-of-Error](#).

Output

Output for this window can be obtained in HTML format by clicking the ***Output to HTML file*** toolbar icon.

Closing the Comparison Windows

When the **Return** toolbar icon is clicked, the ***Host Capacity Comparison*** window and any additional windows chained off it will automatically be closed, and control will be restored to the ***Control Panel*** window. Any open ***LPAR Configuration Capacity Planning*** windows will be unhidden.

System Recovery Boost Considerations

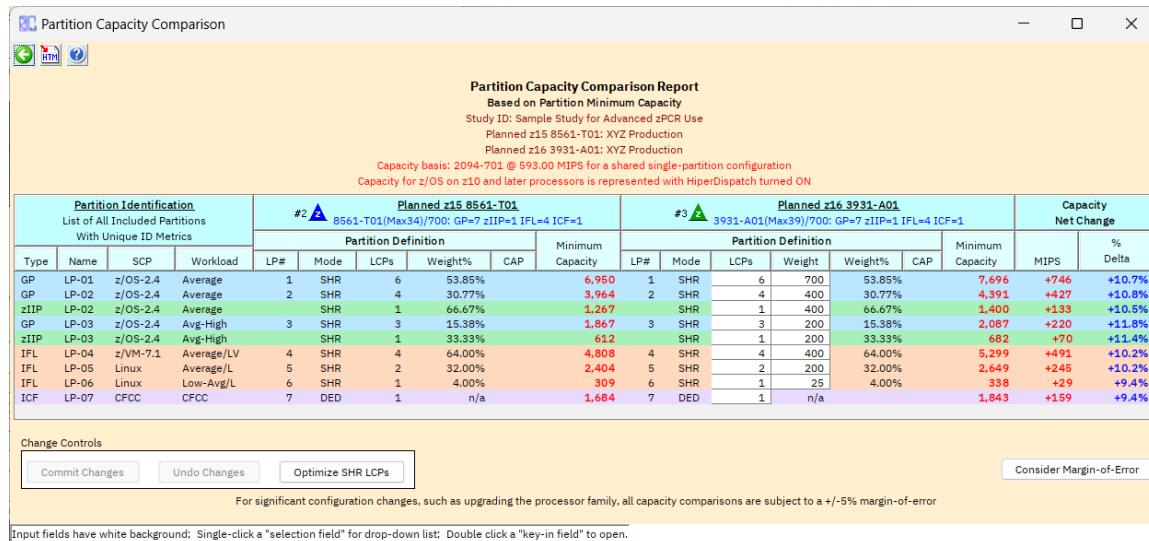
If any partition in either of the selected LPAR configurations Has ***System Recovery Boost*** activity, the following message will display at the bottom of the window.

Configuration #N has one or more partitions in ***System Recovery Boost***. Don't use this configuration for a processor sizing.

For detail concerning SRB support see [System Recovery Boost](#).

Partition Capacity Comparison

The **Partition Capacity Comparison** window is presented by clicking either the **Minimum Capacity** or the **Maximum Capacity** button on the **Host Capacity Comparison** window. The button used determines whether the **Minimum Capacity** values or the **Maximum Capacity** values will be the subject of the comparison. The **Host Capacity Comparison** window remains open, and will reflect capacity results of configuration changes made via the **Partition Capacity Comparison** window.



The screenshot shows the 'Partition Capacity Comparison Report' window. The report is based on 'Partition Minimum Capacity' for Study ID: Sample Study for Advanced zPCR Use. It compares two configurations: #2 (8561-T01/Max34/700: GP=7 zIIP=1 IFL=4 ICF=1) and #3 (3931-A01/Max39/700: GP=7 zIIP=1 IFL=4 ICF=1). The capacity basis is 2094-701 @ 593.00 MIPS for a shared single-partition configuration. The report table lists partitions by type, name, and workload, showing their definition metrics (LP#, Mode, LCPs, Weight%, CAP) and capacity results (Minimum Capacity, MIPS, % Delta). The table highlights changes in green for configuration #3. A note at the bottom states: 'For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error'.

Partition Identification				Planned z15 8561-T01					Planned z16 3931-A01					Capacity Net Change		
List of All Included Partitions With Unique ID Metrics				#2 8561-T01/Max34/700: GP=7 zIIP=1 IFL=4 ICF=1					#3 3931-A01/Max39/700: GP=7 zIIP=1 IFL=4 ICF=1							
Type	Name	SCP	Workload	Partition Definition					Partition Definition					Minimum Capacity	MIPS	% Delta
				LP#	Mode	LCPs	Weight%	CAP	LP#	Mode	LCPs	Weight	Weight%			
GP	LP-01	z/OS-2.4	Average	1	SHR	6	53.85%	6,950	1	SHR	6	700	53.85%	7,696	+746	+10.7%
GP	LP-02	z/OS-2.4	Average	2	SHR	4	30.77%	3,964	2	SHR	4	400	30.77%	4,391	+427	+10.8%
zIIP	LP-02	z/OS-2.4	Average		SHR	1	66.67%	1,267		SHR	1	400	66.67%	1,400	+133	+10.5%
GP	LP-03	z/OS-2.4	Avg-High	3	SHR	3	15.38%	1,867	3	SHR	3	200	15.38%	2,087	+220	+11.8%
zIIP	LP-03	z/OS-2.4	Avg-High		SHR	1	33.33%	612		SHR	1	200	33.33%	682	+70	+11.4%
IFL	LP-04	z/VM-7.1	Average/LV	4	SHR	4	64.00%	4,808	4	SHR	4	400	64.00%	5,299	+491	+10.2%
IFL	LP-05	Linux	Average/L	5	SHR	2	32.00%	2,404	5	SHR	2	200	32.00%	2,649	+245	+10.2%
IFL	LP-06	Linux	Low-Avg/L	6	SHR	1	4.00%	309	6	SHR	1	25	4.00%	338	+29	+9.4%
ICF	LP-07	CFCC	CFCC	7	DED	1	n/a	1,684	7	DED	1	n/a		1,843	+159	+9.4%

Change Controls: Commit Changes, Undo Changes, Optimize SHR LCPs. Consider Margin-of-Error.

Input fields have white background: Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.

The **Partition Capacity Comparison** window compares the capacity results between similarly defined partitions of the two LPAR configurations. Partitions are considered similar when they have the same partition type, name, SCP**, and workload. The partitions' definition metrics and capacity result are shown for both configurations. The capacity **Net Change** (increase or decrease) and **% Delta** for each matched partition is shown on the right.

SCP **: z/OS and z/VM are always specified with a version, used solely for the purpose of enforcing configuration rules. Capacity values are unaffected. In cases where the versions differ, the SCP will display with “***” in lieu of the version.

zAAP and zIIP partitions will be displayed in the order most recently chosen on the **Partition Detail Report** window (any defined LPAR configuration). The default order is to display zAAP and zIIP partitions immediately following their parent GP partition.

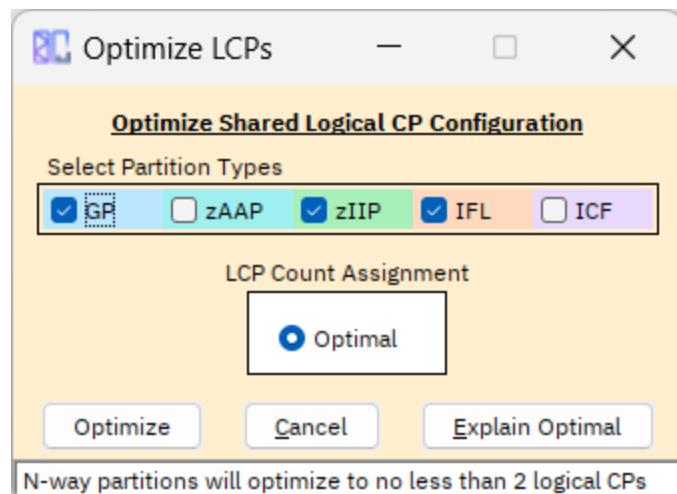
While displayed, the **Partition Capacity Comparison** window provides the capability to modify several of the partition defining metrics for the 2nd LPAR configuration. These input fields, shown with a white background, include number of LCPs, weight, and capping. Changes to these metrics can be made to quickly assess the effect on the capacity of the second LPAR configuration, and, hence, on the capacity **Net Change** (increase or decrease) and the **% Delta** for each partition. Each time a change is made, the LPAR configuration is recalculated, and the capacity results are updated. Changes are cumulative, and are considered to be temporary. Modifications to these metrics can be restored to their initial setting in the **Change Controls** group box, by clicking **Undo**, or by simply closing the window. To change the LPAR configuration to permanently include the modified metrics, click **Commit** button before returning from the window.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for detail see [zAAP/zIIP Loading](#)). The **GP** row heading for affected partitions will appear as “**GP***” and the General Purpose capacity number will be presented in **brown** rather than the normal **red**. The reason is that General Purpose capacity for the configuration is improved somewhat over that which would be shown using the default zAAP/zIIP loading values. The concern is that, when comparing General Purpose capacity between such configurations, user specified zAAP/zIIP loading values should be identical for matching partitions in both configurations. An explanation will be provided in the message box at the bottom of the window. In cases where actual zAAP/zIIP loading values differ between matching partitions, a **Notice** dialog window will appear.

Optimizing LPAR Configuration Capacity

Often, the number of LCPs assigned to a shared partition is more than adequate to satisfy the capacity requirement indicated by its relative weight. Such overstatements can easily happen when planning for a newer processor family, with faster but fewer RCPs. The number of LCPs will be reduced only to satisfy the partitioning rules without consideration of the relative weight assigned.

On the **Partition Capacity Comparison** window (**Minimum Capacity** or **Maximum Capacity** version), an **Optimize SHR LCPs** button is provided in the **Change Controls** group box. Click this button to open the **Optimize LCPs** window, which helps in streamlining the LCP counts for the LPAR configuration to achieve more efficient use of the available RCPs.



Any or all the partition types may be selected for optimization, using the checkboxes provided in the **Select Partition Types** group box. By default, all partition types with shared LCPs available are checked.

The **Optimal** level of optimization for **LCP Count Assignment** uses a new algorithm that adheres to established best practices.

- The Optimal shared LCP count setting utilizes the number of weight-based engines and no more than 1 or 2 additional LCPs, depending on the n-way, to provide the appropriate amount of white space, adhering to established best practices.

N-way partitions will optimize to no less than 2 LCPs for SCP availability reasons.

The resulting LCP assignments are based strictly on the relative weight of the partition. The number of LCPs assigned will determine the **Maximum Capacity** that could ever be seen by a partition. In cases where partition workload peaks might require more capacity than is indicated by **Maximum Capacity**, the optimized LCP assignment may need to be increased.

Use the **Explain Optimal** button for a high-level description of how Optimal LCPs are calculated.

Click **Optimize** to apply the new LCP assignments. This action will return to the **Partition Capacity Comparison** window, with the new LCP assignments temporarily applied and capacity values updated to reflect the change. Note that you must click the **Commit** button for the new LCP counts to be assigned permanently to the LPAR configuration. Clicking the **Undo** button will restore the LCP counts to what they were before optimizing them.

Click **Cancel** to return to the **Partition Capacity Comparison** window without changing any LCP assignments.

Special Considerations Concerning Partition LCP Assignments

HiperDispatch: On z10 and later processors running z/OS or z/VM with HiperDispatch enabled, excessively configured LCPs will tend to be parked, based on the overall demand for capacity. Parked LCPs of partitions should not be included in an LPAR configuration, since only active LCPs contribute to cost of partitioning.

IRD: On **IBM Z** platforms running z/OS with the **Intelligent Resource Director (IRD)**, shared partition LCPs that are varied offline should not be included in an LPAR configuration, since only the active LCPs contribute to cost of partitioning.

Margin-of-Error Reports

Click the **Consider Margin-of-Error** button to display the **Partition Margin-of-Error** window. This window will reveal the capacity results with the full 5% margin-of-error applied. For detail see [Partition Margin-of-Error](#).

Note: This capability is provided only on the **Minimum Capacity** version of the **Partition Capacity Comparison** window.

Output

Output for this window can be obtained in HTML format by clicking the ***Output to HTML file*** toolbar icon.

Closing the Partition Comparison Window

When the **Return** toolbar icon is clicked, the ***Partition Capacity Comparison*** window and any additional windows chained off it will automatically be closed, and control will be restored to the ***Control Panel*** window. Any open ***LPAR Configuration Capacity Planning*** windows will be unhidden.

System Recovery Boost Considerations

If any partition in either of the selected LPAR configurations has ***System Recovery Boost*** activity, that configuration will be **highlighted in the title area** at the top of the table, and the following message will display at the bottom of the window.

Configuration #N has one or more partitions in **System Recovery Boost**. Don't use this configuration for a processor sizing.

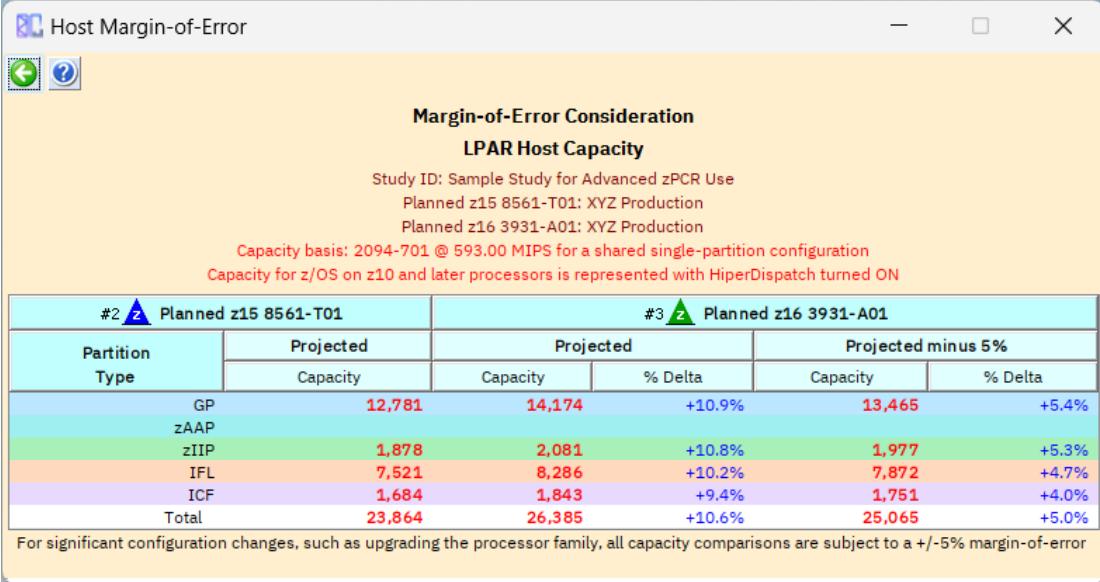
In addition, the individual boosted partition's **Type** and **SRB** columns are highlighted.

For detail concerning SRB support see [System Recovery Boost](#).

Host Margin-of-Error

The capacity expectation derived from **zPCR** for a new processor should normally be considered as having a margin-of-error of up to 5% margin-of-error. The full $\pm 5\%$ margin of error should be considered whenever the LPAR host processor family is changed, or when very significant changes are made to either the LPAR host CP configuration or to the partition configuration itself. When changes are minor, the margin-of-error should be less.

The **Host Margin-of-Error** window is available from the **Host Capacity Comparison** window by clicking the **Consider Margin-of-Error** button.



The screenshot shows the 'Margin-of-Error Consideration' window for LPAR Host Capacity. It compares two configurations: #2 Planned z15 8561-T01 and #3 Planned z16 3931-A01. The window includes a message bar at the top with study details and capacity basis information. Below is a table comparing projected capacities and their 5% margin of error for various partition types (GP, zAAP, zIIP, IFL, ICF) and totals.

Partition Type	#2 Planned z15 8561-T01		#3 Planned z16 3931-A01		Projected minus 5%	
	Projected Capacity	Projected Capacity	Projected % Delta	Projected Capacity	Projected % Delta	
GP	12,781	14,174	+10.9%	13,465	+5.4%	
zAAP	1,878	2,081	+10.8%	1,977	+5.3%	
zIIP	7,521	8,286	+10.2%	7,872	+4.7%	
IFL	1,684	1,843	+9.4%	1,751	+4.0%	
ICF						
Total	23,864	26,385	+10.6%	25,065	+5.0%	

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a $\pm 5\%$ margin-of-error

The **Projected Capacity**, summed by partition type, is shown for both LPAR configurations and the capacity **% Delta** is computed for the 2nd LPAR configuration. The **Projected Capacity** and **% Delta** columns for each partition type of the 2nd LPAR configuration are then repeated on the right, with the -5% margin-of-error applied.

Note that the capacity projection for the 2nd LPAR configuration, the margin-of-error consideration is shown for all partition types regardless of whether the partition type was defined on the 1st LPAR configuration.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for detail see [zAAP/zIIP Loading](#)). The **GP** row heading will appear as "**GP***" and the General Purpose capacity number will be presented in **brown** rather than the normal **red**. The reason is that General Purpose capacity for the configuration is improved somewhat over that which would be shown using the default zAAP/zIIP loading values. The concern is that, when comparing General Purpose capacity between such configurations, user specified zAAP/zIIP loading values should be identical for matching partitions in both configurations. An explanation will be provided in the message box at the bottom of the window.

System Recovery Boost Considerations

If any partition in either of the selected LPAR configurations has **System Recovery Boost** activity, that configuration will be **highlighted in the title area** at the top of the table, and the following message will display at the bottom of the window.

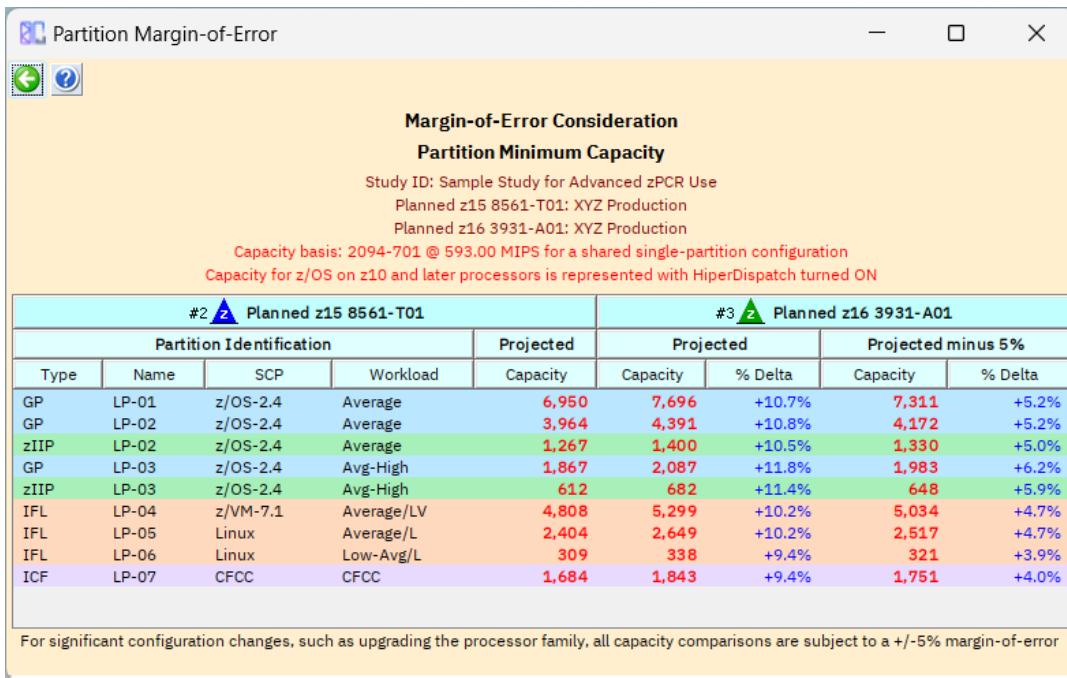
Configuration #N has one or more partitions in **System Recovery Boost**. Don't use this configuration for a processor sizing.

For detail concerning SRB support see [System Recovery Boost](#).

Partition Margin-of-Error

The capacity expectation derived from **zPCR** for a new processor should normally be considered as having a margin-of-error of up to 5% margin-of-error. The full $\pm 5\%$ margin of error should be considered whenever the LPAR host processor family is changed, or when very significant changes are made to either the LPAR host CP configuration or to the partition configuration itself. When changes are minor, the margin-of-error should be less.

The **Partition Margin-of-Error** window is available from the **Partition Capacity Comparison** window by clicking the **Consider Margin-of-Error** button. Note that the **Partition Margin-of-Error** window is available only when **Minimum Capacity** values are displayed.



The **Projected Capacity** for similarly defined partitions is shown for both LPAR configurations and the capacity **% Delta** is computed for each partition of the 2nd LPAR configuration. The **Projected Capacity** and **% Delta** columns for each partition of the second LPAR configuration are repeated on the right, with the 5% margin-of-error applied.

zAAP and zIIP partitions will be displayed in the order most recently chosen on the **Partition Detail Report** window (any defined LPAR configuration). The default order is to display zAAP and zIIP partitions immediately following their parent GP partition.

Note that the capacity projection for the 2nd LPAR configuration that includes the margin-of-error consideration is shown for all defined partitions regardless of whether the partition was defined on the 1st LPAR configuration.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for detail see [zAAP/zIIP Loading](#)). The **GP** row heading will appear as “**GP***” and the General Purpose capacity number will be presented in **brown** rather than the normal **red**. The reason is that General Purpose capacity for the configuration is improved somewhat over that which would be shown using the default zAAP/zIIP loading values. The concern is that, when comparing General Purpose capacity between such configurations, user specified zAAP/zIIP loading values should be identical for matching partitions in both configurations. An explanation will be provided in the message box at the bottom of the window.

System Recovery Boost Considerations

If any partition in either of the selected LPAR configurations has **System Recovery Boost** activity, that configuration will be **highlighted in the title area** at the top of the table, and the following message will display at the bottom of the window.

Configuration #N has one or more partitions in **System Recovery Boost**. Don't use this configuration for a processor sizing.

For detail concerning SRB support see [System Recovery Boost](#).

Copy Partitions to an LPAR Configuration

When multiple LPAR configurations are defined, one of them can be selected to be the target for partitions to be copied from any of the other LPAR configurations. Select the

LPAR configuration that is to be the target for migrated partitions and click  in the **Migrate & Analyze** group box to open the **Copy Partitions: Receiving LPAR Configuration** window. The target configuration can be an existing LPAR configuration or may be an LPAR host definition with no partitions defined.

 Copy Partitions: Receiving LPAR Configuration

   #1  #2  #3  #4 

Copy Partitions: Receiving LPAR Configuration

#3  Planned z16 3931-A01
Description: XYZ Production

z16 Host = 3931-A01(Max39)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Orig LPAR Config	Partition Identification						Partition Configuration						Partition Capacity			
	Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT	Original Minimum	Modified Minimum	Net Change	% Delta
INIT	ABS	V	Benefit													
<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%				7,696	7,696		
<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%				4,391	4,391		
<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%				1,400	1,400		
<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%				2,087	2,087		
<input checked="" type="checkbox"/>	3	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%				682	682		
<input checked="" type="checkbox"/>	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%				5,299	5,299		
<input checked="" type="checkbox"/>	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%				2,649	2,649		
<input checked="" type="checkbox"/>	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%				338	338		
<input checked="" type="checkbox"/>	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a					1,843	1,843		

Capacity Summary by Pool

CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta					
		Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP
GP	7	3	13	1.857	14,174	3	13	1.857	14,174						
zIIP	1	2	2	2.000	2,081	2	2	2.000	2,081						
IFL	4	3	7	1.750	8,286	3	7	1.750	8,286						
ICF	1	1	1	All DED	1,843	1	1	All DED	1,843						
Totals	13	9	23	26,385		9	23	26,385							

Commit Changes **Undo Changes**

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.
Using the LPAR configuration icons at the top of this window, select the configuration from which the partitions are to be migrated.

This window is similar to the **Partition Detail Report** window, except for the **Partition Capacity** columns which here, always relate to the partition **Minimum Capacity** values.

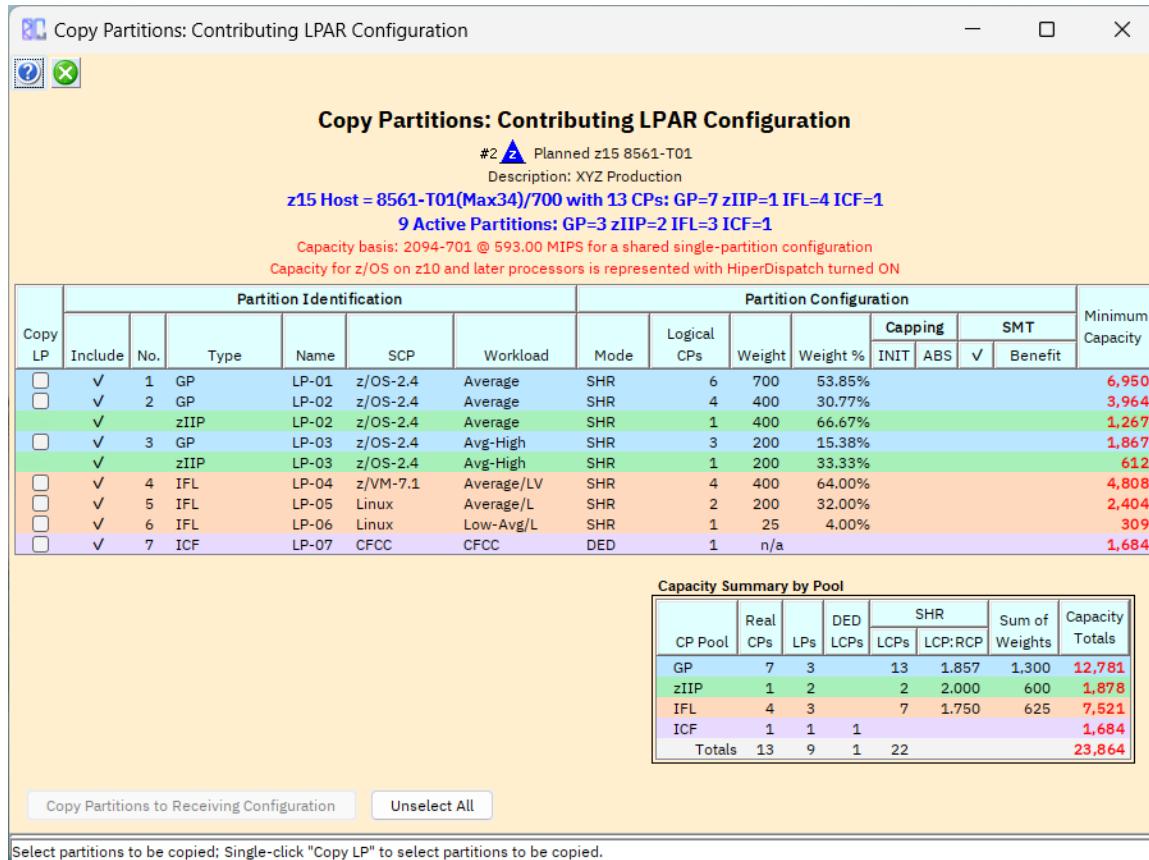
The **Original Minimum** column reveals partition capacity when the window was initially opened (before any partitions are copied or modified). For partitions that have been copied into the LPAR configuration, the **Original Capacity** value will be that from the contributing LPAR configuration.

The **Modified Minimum** column shows capacity that includes the effect of any partition changes as they are made. Possible changes include additional partitions that are 1) copied from another LPAR configuration, 2) changing partition **LCPs**, **Weight**, or **CAP** setting, or 3) unchecking a partition's **Include** checkbox.

All capacity values are recomputed with each change to the LPAR configuration. **Net Change** and **% Delta** metrics are computed expressing the difference between the **Modified Minimum** capacity column and the **Original Minimum** capacity column.

The **Capacity Summary by Pool** table reveals CP pool metrics for the Original Configuration (state when the window was opened) and for the Modified Configuration (state after partitions have been copied or modified). In addition, **% Delta** values are provided for each of the metrics except **SHR LCP:RCP**.

Icons representing each of the remaining LPAR configurations appear enabled at the top of the window. Click one of them to select the LPAR configuration from which partitions will be copied. The **Copy Partitions: Contributing LPAR Configuration** window will open to the left of the **Copy Partitions: Receiving LPAR Configuration** window.

The screenshot shows the 'Copy Partitions: Contributing LPAR Configuration' window. At the top, it displays the host information: 'z15 Host = 8561-T01(Max34)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1' and '9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1'. Below this is a note about capacity basis: 'Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration' and 'Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON'. The main table, 'Partition Identification', lists 13 partitions (LP-01 to LP-07) with columns for Copy LP, Include, No., Type, Name, SCP, Workload, Mode, Logical CPs, Weight, Weight %, and Capping/SMT values. The 'Minimum Capacity' column shows values like 6,950, 3,964, 1,267, etc. The 'Capacity Summary by Pool' table at the bottom shows metrics for GP, zIIP, IFL, and ICF pools. Buttons at the bottom include 'Copy Partitions to Receiving Configuration' and 'Unselect All'.

Copy LP	Partition Identification						Partition Configuration							Minimum Capacity
	Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT		
<input type="checkbox"/>	✓	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%			6,950	
<input type="checkbox"/>	✓	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%			3,964	
	✓		zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%			1,267	
<input type="checkbox"/>	✓	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%			1,867	
	✓		zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%			612	
<input type="checkbox"/>	✓	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%			4,808	
<input type="checkbox"/>	✓	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%			2,404	
<input type="checkbox"/>	✓	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%			309	
<input type="checkbox"/>	✓	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a				1,684	

CP Pool	Real CPs		DED		SHR		Sum of Weights	Capacity Totals
	CPs	LPs	LPs	LCPs	LCP:RCP			
GP	7	3	13	1.857	1,300	12,781		
zIIP	1	2	2	2.000	600	1,878		
IFL	4	3	7	1.750	625	7,521		
ICF	1	1	1			1,684		
Totals	13	9	1	22		23,864		

Select partitions to be copied: Single-click "Copy LP" to select partitions to be copied.

This window is similar to the **Partition Detail Report** window, except for the **Capacity** column which, here, reveals only partition **Minimum Capacity**.

To copy a partition into the **Receiving LPAR Configuration**, click its **Copy LP** checkbox. Once all desired partitions are selected, click the **Copy Partitions to Receiving Configuration** button and the partition definition for each will be copied into the **Receiving LPAR Configuration**. Partitions that have already been copied to the **Receiving LPAR Configuration** will display a dark background for its **Copy LP** checkbox and cannot be re-selected.

Copy Partitions: Contributing LPAR Configuration

Copy Partitions: Contributing LPAR Configuration

#2 Planned z15 0561-T01
Description: XYZ Production

z15 Host = 8561-T01 (Max34)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Copy LP	Partition Identification					Partition Configuration							Minimum Capacity
	Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT	
INIT	ABS	V	Benefit										
<input type="checkbox"/>	✓	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%			6,950
<input type="checkbox"/>	✓	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%			3,964
<input type="checkbox"/>	✓		zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%			1,267
<input checked="" type="checkbox"/>	✓	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%			1,867
<input type="checkbox"/>	✓		zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%			612
<input type="checkbox"/>	✓	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%			4,808
<input type="checkbox"/>	✓	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%			2,404
<input type="checkbox"/>	✓	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%			309
<input type="checkbox"/>	✓	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a				1,684

Capacity Summary by Pool

CP Pool	Real CPs	LPs	DED	SHR		Sum of Weights	Capacity Totals
				LCPs	LCP:RCP		
GP	7	3		13	1.857	1,300	12,781
zIIP	1	2		2	2.000	600	1,878
IFL	4	3		7	1.750	625	7,521
ICF	1	1		1			1,684
Totals	13	9	1	22			23,864

Select partitions to be copied; Single-click "Copy LP" to select partitions to be copied.

Partitions that would be invalid on the **Receiving LPAR Configuration** will not have their **Copy LP** checkbox enabled. This would normally be due to insufficient RCPs needed to support the partition. In cases where this problem condition is undetected, the partition will be copied, but its **Include** on the **Receiving LPAR Configuration** will be unchecked.

Parent GP partitions with associated zAAP, zIIP, IFL, and ICF LCPs will always be copied together (if the associated partition is not desired in the **Receiving Configuration**, its **Include** can be unchecked). Note that associated zAAP, zIIP, IFL, and ICF partitions cannot be copied without its parent GP partition.

Click the **Unselect All** button to clear all **Copy LP** checkboxes.

Close the **Contributing LPAR Configuration** window by clicking **Cancel** in the upper right corner. Another LPAR configuration can then be selected from the **Receiving LPAR Configuration** to become the **Contributing LPAR Configuration** window.

Partitions that have been copied into the **Receiving LPAR Configuration** have their **Contributing LPAR Configuration** number and icon indicated in the **Original LPAR Config** column. The **Modified Capacity** values are computed for the new partition configuration and compared back to the **Original Capacity** values as revealed in the **Net Change** and **% Delta** columns.

Copy Partitions: Receiving LPAR Configuration

#1 #2 #3 #4

Copy Partitions: Receiving LPAR Configuration

#3 Planned z16 3931-A01
Description: XYZ Production

z16 Host = 3931-A01(Max39)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
11 Active Partitions: GP=4 zIP=3 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Orig LPAR Config	Partition Identification					Partition Configuration							Partition Capacity				
	Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping INIT	Capping ABS	SMT V	Original Minimum	Modified Minimum	Net Change	% Delta
	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	46.67%					7,696	6,588	-1,108	-14.4%
	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	26.67%					4,391	3,759	-632	-14.4%
	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	50.00%					1,400	1,036	-363	-26.0%
	4	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	13.33%					2,087	1,786	-301	-14.4%
#2	5	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	25.00%					682	504	-177	-26.0%
#2	6	GP	LP-03+	z/OS-2.4	Avg-High	SHR	3	200	13.33%					1,867	1,786	-81	-4.3%
	7	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	25.00%					612	504	-108	-17.6%
	8	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%					5,299	5,292	-7	-0.1%
	9	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%					2,649	2,646	-3	-0.1%
	10	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%					338	338	-0	-0.1%
	11	ICF	LP-07	CFCC	CFCC	DED	1	n/a						1,843	1,843	-0	-0.0%

Capacity Summary by Pool

CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta					
		Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP
GP	7	3	13	1.857	14,174	4	16	2.286	13,919	+33.3%	+23.1%			+23.1%	-2%
zIIP	1	2	2	2.000	2,081	3	3	3.000	2,045	+50.0%	+50.0%			+50.0%	-2%
IFL	4	3	7	1.750	8,286	3	7	1.750	8,276	0.0%	0.0%			0.0%	-0%
ICF	1	1	1	All DED	1,843	1	1	All DED	1,843	0.0%	0.0%			0.0%	-0%
Totals	13	9	23		26,385	11	27		26,082	+22.2%	+17.4%			-1%	

Commit Changes Undo Changes

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.
Using the LPAR configuration icons at the top of this window, select the configuration from which the partitions are to be migrated.

Whenever a copied partition's name matches one already in the configuration, a plus will be appended to make the name unique.

On this window, partition **LCPs** and **Weights** may be adjusted to rebalance CPU resource between partitions in each CP pool. Any change to a partition will likely have some effect on the capacity values of other partitions. All partitions are eligible to be excluded or modified to help understand the overall effect on capacity values.

Once the partition migration plan is complete, click **Commit Changes** to make the copied partitions permanent to the **Receiving LPAR Configuration**.

The **Original Capacity** values will remain fixed as long as the **Receiving LPAR Configuration** window is open. Additional partition migrations can still be applied, but another **Commit Changes** will be required to make those changes permanent.

Click the **Undo Changes** button to remove all copied partitions and partition changes since the last **Commit Changes**.

Note that there is no output capability for this function. A bitmap can be captured by pressing **Alt-PrintScreen** while it is the in-focus window.

System Recovery Boost Considerations

If any partition in the LPAR configuration has **System Recovery Boost** activity, the **SRB** column is included in the table and the individual boosted partition's **No.** **Type** and **SRB** columns are highlighted.

In addition, the following message is displayed at the bottom of the window.

Note: Partition was in System Recovery Boost during the selected interval. Don't use this configuration for a processor sizing.

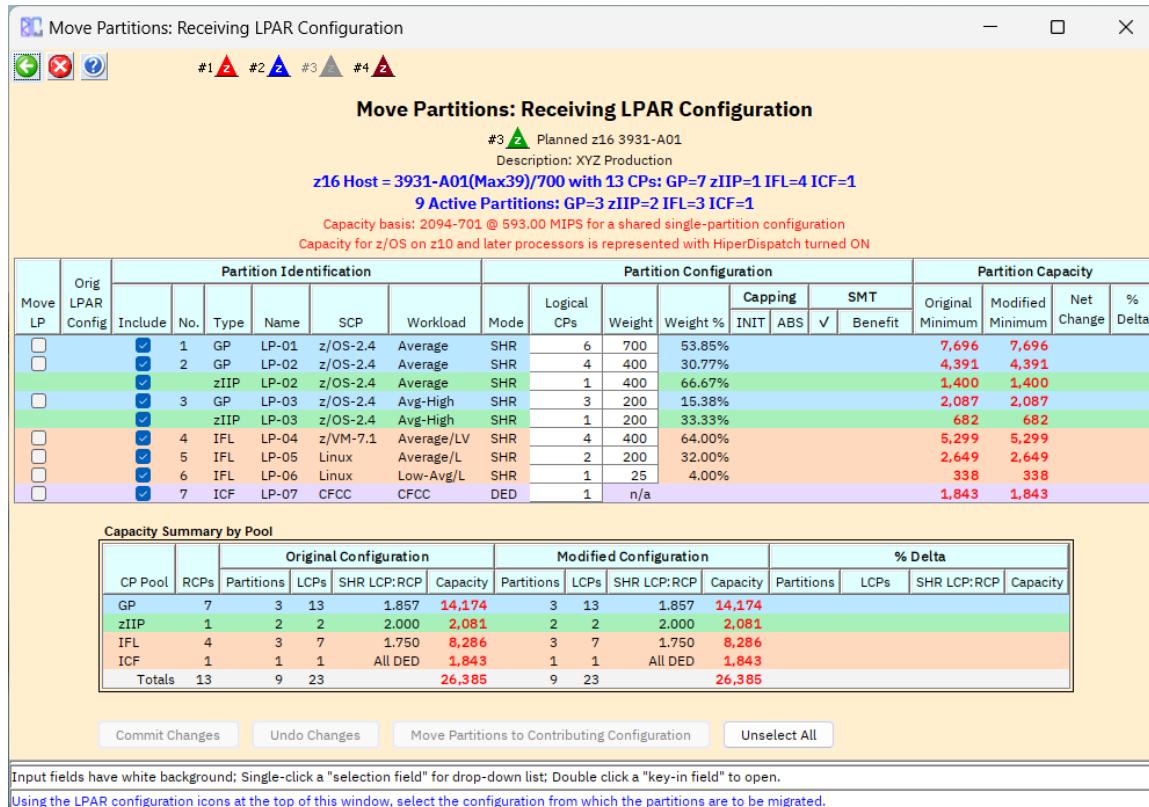
Note: **Initial Capping** and **Absolute Capping** are ignored for zIIP partitions using zIIP Boost.

For detail concerning SRB support see [System Recovery Boost](#).

Move Partitions between Configurations

When multiple LPAR configurations are defined, one of them can be selected to be the target for partitions to be moved from any of the other LPAR configurations. Select the

LPAR configuration that is to be the target for migrated partitions and click  in the **Migrate & Analyze** group box to open the **Move Partitions: Receiving LPAR Configuration** window. The target configuration can be an existing LPAR configuration or may be an LPAR host definition with no partitions defined.

Move Partitions: Receiving LPAR Configuration

#3  Planned z16 3931-A01
Description: XYZ Production
z16 Host = 3931-A01(Max39)/700 with 13 CPUs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1
Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Move LP	Orig LPAR Config	Partition Identification						Partition Configuration						Partition Capacity			
		Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT	Original Minimum	Modified Minimum	Net Change	% Delta
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	7,696	7,696	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	4,391	4,391	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	1,400	1,400	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	4	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	2,087	2,087	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	5	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	682	682	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	6	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	5,299	5,299	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	7	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	2,649	2,649	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	8	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%	<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	338	338	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	9	ICF	LP-07	CFCC	CFCC	DED	1	n/a		<input type="checkbox"/> INIT	<input type="checkbox"/> ABS	<input type="checkbox"/> V	<input type="checkbox"/> Benefit	1,843	1,843	

Capacity Summary by Pool

CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta					
		Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP
GP	7	3	13	1.857	14,174	3	13	1.857	14,174						
zIIP	1	2	2	2.000	2,081	2	2	2.000	2,081						
IFL	4	3	7	1.750	8,286	3	7	1.750	8,286						
ICF	1	1	1	All DED	1,843	1	1	All DED	1,843						
Totals	13	9	23	26,385		9	23	26,385							

Commit Changes Undo Changes Move Partitions to Contributing Configuration Unselect All

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.
Using the LPAR configuration icons at the top of this window, select the configuration from which the partitions are to be migrated.

This window is similar to the **Partition Detail Report** window, except for the **Partition Capacity** columns which, here, always relate to the partition **Minimum Capacity** values.

The **Original Minimum** column reveals partition capacity when the window was initially opened (before any partitions are moved or modified). For partitions that have been moved to the **Receiving LPAR Configuration**, the **Original Capacity** value will be that from the **Contributing LPAR Configuration**.

The **Modified Minimum** column shows capacity that includes the effect of any partition changes as they are made. Possible changes include additional partitions that are 1) moved from another LPAR configuration, 2) changing partition **LCPs**, **Weight**, or **CAP** setting, or 3) unchecking a partition's **Include** checkbox.

All capacity values are recomputed with each change to the LPAR configuration. **Net Change** and **% Delta** metrics are computed expressing the difference between the **Modified Minimum** capacity column and the **Original Minimum** capacity column.

The **Capacity Summary by Pool** table reveals CP pool metrics for the Original Configuration (state when the window was opened) and for the Modified Configuration (state after partitions have been moved or modified). In addition, **% Delta** values are provided for each of the metrics except **SHR LCP:RCP**.

Icons representing each of the remaining LPAR configurations appear enabled at the top of the window. Click one of them to select the LPAR configuration from which partitions will be moved. The **Move Partitions: Contributing LPAR Configuration** window will open to the left of the **Move Partitions: Receiving Configuration** window.

Move Partitions: Contributing LPAR Configuration

Move Partitions: Contributing LPAR Configuration

#2 Planned z15 8561-T01
Description: XYZ Production

z15 Host = 8561-T01(Max34)/700 with 13 CPs: GP=7 zIIP=2 IFL=4 ICF=1
9 Active Partitions: GP=3 zIIP=2 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Move LP	Orig LPAR Config	Partition Identification						Partition Configuration						Partition Capacity			
		Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT	Original Minimum	Modified Minimum	Net Change	% Delta
<input type="checkbox"/>		✓	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	53.85%		6,950	6,950			
<input type="checkbox"/>		✓	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	30.77%		3,964	3,964			
		✓	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	66.67%		1,267	1,267				
<input checked="" type="checkbox"/>		✓	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	15.38%		1,867	1,867			
<input checked="" type="checkbox"/>		✓	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	33.33%		612	612				
<input type="checkbox"/>		✓	4	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%		4,808	4,808			
<input type="checkbox"/>		✓	5	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%		2,404	2,404			
<input type="checkbox"/>		✓	6	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%		309	309			
<input type="checkbox"/>		✓	7	ICF	LP-07	CFCC	CFCC	DED	1	n/a			1,684	1,684			

Capacity Summary by Pool

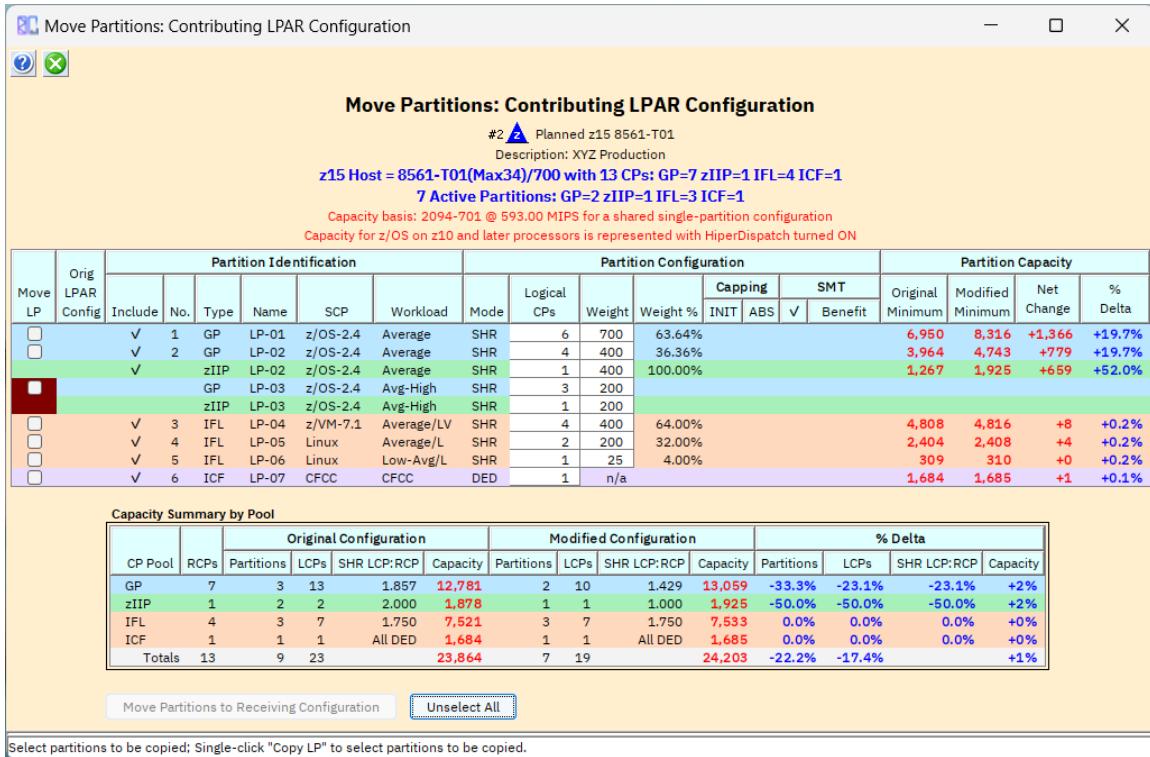
CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta			
		Partitions	LCPs	SHR LCP:RCP	Capacity	Partitions	LCPs	SHR LCP:RCP	Capacity	Partitions	LCPs	SHR LCP:RCP	Capacity
GP	7	3	13	1.857	12,781	3	13	1.857	12,781				
zIIP	1	2	2	2.000	1,878	2	2	2.000	1,878				
IFL	4	3	7	1.750	7,521	3	7	1.750	7,521				
ICF	1	1	1	All DED	1,684	1	1	All DED	1,684				
Totals	13	9	23		23,864	9	23		23,864				

Move Partitions to Receiving Configuration Unselect All

Select partitions to be copied; Single-click "Copy LP" to select partitions to be copied.

This window is like the **Partition Detail Report** window, except for the **Capacity** column which, here, reveals partition **Minimum Capacity**.

To move a partition into the **Receiving LPAR Configuration**, click its **Move LP** checkbox. Once all desired partitions are selected, click the **Move Partitions to Receiving Configuration** button and the partition definition for each will be moved into the **Receiving LPAR Configuration**. On the **Contributing LPAR Configuration** window, moved partitions will display a dark background for its **Copy LP** checkbox partition definition is removed. All capacity values for the both LPAR configurations are recomputed. **Net Change** and **% Delta** metrics are computed expressing the difference between the **Modified Minimum** capacity column and the **Original Minimum** capacity column.



Parent GP partitions with associated zAAP, zIIP, IFL, and ICF LCPs will always be copied together (if the associated partition is not desired in the **Receiving Configuration**, its **Include** can be unchecked). Note that zAAP, zIIP, IFL, and ICF partitions, when associated with a GP partition), cannot be copied without the parent GP partition.

On the **Contributing LPAR Configuration** window, click the **Unselect All** button to clear any **Move LP** checkboxes.

Close the **Contributing LPAR Configuration** window by clicking **Cancel** in the upper right corner. Another LPAR configuration may then be selected from the **Receiving LPAR Configuration** to become the **Contributing LPAR Configuration** window.

Partitions that have been moved into the **Receiving LPAR Configuration** have their **Contributing LPAR Configuration** indicated in the **Original LPAR Config** column. The **Modified Capacity** values are computed for the new partition configuration and compared back to the **Original Capacity** values as revealed in the **Net Change** and **% Delta** columns.

Whenever a moved partition's name matches one already in the configuration, a plus will be appended to make the name unique.

Note that a change to any partition will likely have some effect on the capacity values of every partition. This means that, in addition to partitions moved, partitions can also be excluded or modified to understand the overall effect on capacity values.

Once the partition migration plan is complete, click **Commit Changes** to make the copied partitions permanent to the **Receiving LPAR Configuration**.

The **Original Capacity** values will remain fixed as long as the **Receiving LPAR Configuration** window is open. Additional partition migrations can still be applied, but another **Commit Changes** will be required to make those changes permanent.

On the **Contributing LPAR Configuration** window, partitions that were moved will be flagged with a dark background in the **Move LP** column, indicating that this partition is no longer participating in this configuration.

Move Partitions: Receiving LPAR Configuration

#3 Planned z16 3931-A01
Description: XYZ Production

z16 Host = 3931-A01(Max39)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
11 Active Partitions: GP=4 zIIP=3 IFL=3 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HiperDispatch turned ON

Move LP	Orig LPAR Config	Partition Identification						Partition Configuration						Partition Capacity			
		Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	Capping	SMT	Original Minimum	Modified Minimum	Net Change	% Delta
<input type="checkbox"/>	<input checked="" type="checkbox"/>	1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	46.67%			7,696	6,588	-1,108	-14.4%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	26.67%			4,391	3,759	-632	-14.4%	
	<input checked="" type="checkbox"/>	3	zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	50.00%			1,400	1,036	-363	-26.0%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	3	GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200	13.33%			2,087	1,786	-301	-14.4%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	4	zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200	25.00%			682	504	-177	-26.0%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	4	GP	LP-03+	z/OS-2.4	Avg-High	SHR	3	200	13.33%			1,867	1,786	-81	-4.3%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	4	zIIP	LP-03+	z/OS-2.4	Avg-High	SHR	1	200	25.00%			612	504	-108	-17.6%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	5	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	64.00%			5,299	5,292	-7	-0.1%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	6	IFL	LP-05	Linux	Average/L	SHR	2	200	32.00%			2,649	2,646	-3	-0.1%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	7	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	4.00%			338	338	-0	-0.1%	
<input type="checkbox"/>	<input checked="" type="checkbox"/>	8	ICF	LP-07	CFCC	CFCC	DED	1	n/a				1,843	1,843	-0	-0.0%	

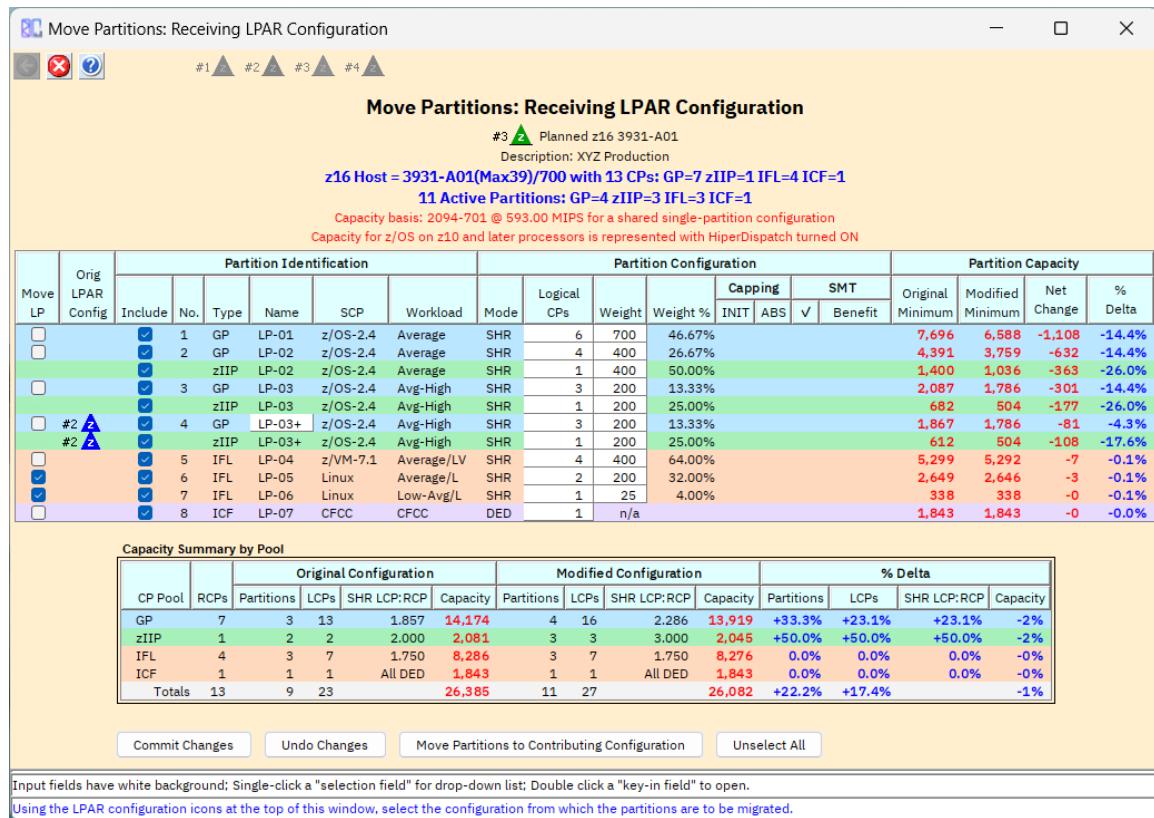
Capacity Summary by Pool

CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta					
		Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP
GP	7	3	13	1.857	14,174	4	16	2.286	13,919	+33.3%	+23.1%	+23.1%	-2%		
zIIP	1	2	2	2.000	2,081	3	3	3.000	2,045	+50.0%	+50.0%	+50.0%	-2%		
IFL	4	3	7	1.750	8,286	3	7	1.750	8,276	0.0%	0.0%	0.0%	-0%		
ICF	1	1	1	All DED	1,843	1	1	All DED	1,843	0.0%	0.0%	0.0%	-0%		
Totals	13	9	23	26,385		11	27	26,082	+22.2%	+17.4%	-1%				

Buttons: Commit Changes | Undo Changes | Move Partitions to Contributing Configuration | Unselect All

Input fields have white background; Single-click a "selection field" for drop-down list; Double click a "key-in field" to open.
Using the LPAR configuration icons at the top of this window, select the configuration from which the partitions are to be migrated.

Original partitions can also be moved from the **Receiving LPAR Configuration** to the **Contributing LPAR Configuration**. Such partitions are moved by checking the selections on the **Receiving LPAR Configuration** window, and clicking the **Move Partitions to Contributing Configuration** button.



IBM zPCR User's Guide

Move Partitions: Contributing LPAR Configuration

#2 Planned z15 8561-T01
Description: XYZ Production

z15 Host = 8561-T01(Max34)/700 with 13 CPs: GP=7 zIIP=1 IFL=4 ICF=1
9 Active Partitions: GP=2 zIIP=1 IFL=5 ICF=1

Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Capacity for z/OS on z10 and later processors is represented with HyperDispatch turned ON

Move LP	Orig LPAR Config	Partition Identification						Partition Configuration						Partition Capacity				
		Include	No.	Type	Name	SCP	Workload	Mode	Logical CPs	Weight	Weight %	INIT	ABS	✓	Benefit	Original Minimum	Modified Minimum	Net Change
<input type="checkbox"/>	✓ 1	GP	LP-01	z/OS-2.4	Average	SHR	6	700	63.64%						6,950	8,316	+1,366	+19.7%
<input type="checkbox"/>	✓ 2	GP	LP-02	z/OS-2.4	Average	SHR	4	400	36.36%						3,964	4,743	+779	+19.7%
<input type="checkbox"/>	✓ zIIP	LP-02	z/OS-2.4	Average	SHR	1	400	100.00%						1,267	1,925	+658	+52.0%	
<input type="checkbox"/>	✓ GP	LP-03	z/OS-2.4	Avg-High	SHR	3	200											
<input type="checkbox"/>	✓ zIIP	LP-03	z/OS-2.4	Avg-High	SHR	1	200											
<input type="checkbox"/>	✓ 3	IFL	LP-04	z/VM-7.1	Average/LV	SHR	4	400	47.06%						4,808	3,471	-1,337	-27.8%
<input type="checkbox"/>	✓ 4	IFL	LP-05	Linux	Average/L	SHR	2	200	23.53%						2,404	1,735	-669	-27.8%
<input type="checkbox"/>	✓ 5	IFL	LP-06	Linux	Low-Avg/L	SHR	1	25	2.94%						309	223	-86	-27.8%
<input type="checkbox"/>	#3	IFL	LP-05+	Linux	Average/L	SHR	2	200	23.53%						2,649	1,735	-914	-34.5%
<input type="checkbox"/>	#3	IFL	LP-06+	Linux	Low-Avg/L	SHR	1	25	2.94%						338	223	-115	-34.0%
<input type="checkbox"/>	✓ 8	ICF	LP-07	CFCC	CFCC	DED	1	n/a							1,684	1,685	+2	+0.1%

Capacity Summary by Pool

CP Pool	RCPs	Original Configuration				Modified Configuration				% Delta				
		Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR	LCP:RCP	Capacity	Partitions	LCPs	SHR
GP	7	3	13	1.857	12,781	2	10	1.429	13,059	-33.3%	-23.1%	-23.1%	+2%	
zIIP	1	2	2	2.000	1,878	1	1	1.000	1,925	-50.0%	-50.0%	-50.0%	+2%	
IFL	4	3	7	1.750	7,521	5	10	2.500	7,388	+66.7%	+42.9%	+42.9%	-2%	
ICF	1	1	1	All DED	1,684	1	1	All DED	1,685	0.0%	0.0%	0.0%	+0%	
Totals	13	9	23	23,864		9	22	24,057		0.0%	-4.3%		+1%	

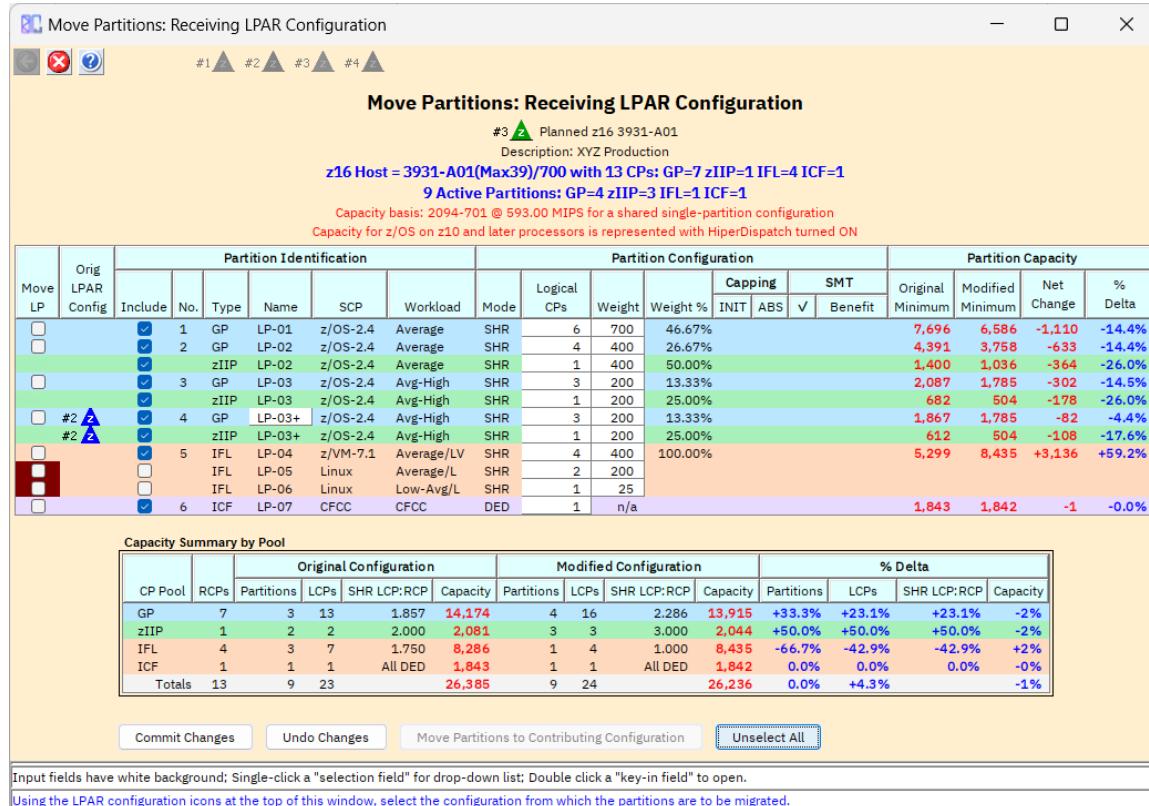
Move Partitions to Receiving Configuration Unselect All

Select partitions to be copied; Single-click "Copy LP" to select partitions to be copied.

The moved partitions now appear on the **Contributing LPAR Configuration** window flagged with the original LPAR configuration icon.

The **Modified Capacity** values are computed for both of the partition configurations and compared back to the **Original Capacity** values as revealed in the **Net Change** and **% Delta** columns.

On the **Receiving LPAR Configuration** window, partitions that were moved will be flagged with a dark background in the **Move LP** column, indicating that the partition is no longer participating in this configuration



On the **Receiving LPAR Configuration** window, click the **Undo Changes** button to remove all moved partitions and partition definition changes since the last **Commit Changes**.

Note that there is no output capability for this function. A bitmap can be captured by pressing **Alt-PrintScreen** while it is the in-focus window.

System Recovery Boost Considerations

If any partition in the LPAR configuration has **System Recovery Boost** activity, the **SRB** column is added to the table and the individual boosted partition's **No.** **Type** and **SRB** columns are highlighted.

In addition, the following message is displayed at the bottom of the window.

Note: Partition was in System Recovery Boost during the selected interval. Don't use this configuration for a processor sizing.

Note: Initial Capping and **Absolute Capping** are ignored for zIIP partitions using zIIP Boost.

For detail concerning SRB support see [System Recovery Boost](#).

Accuracy of LPAR Capacity Projections

zPCR projections represent the **Capacity Perspective**, i.e., all partitions are competing for CPU resource, and partition weights (shared partitions) decide how the CPU resource will be distributed between them. Taken into account is the potential capacity available to each image, less the capacity costs due to partitioning. Results represent capacity that is available for the operating system (i.e., z/OS, Linux, etc.) and its workload.

With normal production workloads, some partitions will tend to dominate, while others will idle to various degrees. When this is the case, actual partitioning costs will be less than those projected by **zPCR**. The same is true when processor utilization averages less than 100%. In addition, partitions that tend to idle release capacity to those who dominate, thereby allowing the dominate partitions to achieve more than their projected minimum capacity.

Since every LPAR configuration defined to **zPCR** provides the **Capacity Perspective**, making capacity comparisons between them is reasonable, assuming that the current and new production systems are run in essentially the same manner.

zPCR accuracy lies partly in the selection of appropriate workload categories to represent the production environment. For both z/OS and z/VM, when CPU MF data is available for a partition (via EDF), an algorithm automatically chooses the appropriate LSPR workload category. Otherwise, the default workload category **Average** will be assigned. Some judgment will be required to change from the default. If the workload selections are appropriate, when comparing the capacity of one processor's LPAR configuration to another, zPCR results should be considered with up to $\pm 5\%$ margin-of-error. This is particularly true when the LPAR host processor's family is being changed, or when the host CP or LPAR configuration is significantly changed. When configuration changes are minor, the margin-of-error will likely be less than $\pm 5\%$.

Because there is a potential $\pm 5\%$ margin-of-error when sizing the capacity of a replacement processor, one should always consider how the new processor's utilization might look if its actual capacity happened to be 5% less than that projected. The safe approach to plan for a new processor would be to always consider at least 5% growth. This concept is particularly important when the current processor is running at high utilization levels and all of the work is considered as high to medium priority.

SCP (System Control Program) considerations

- **z/OS on General Purpose CPs**: Capacity results in **zPCR** should be considered as reliable, given the number of LSPR workload categories that are defined (based on measured LSPR workload primitives) and number of processor models that are measured (see accuracy statement above). The degree to which capacity results will be accurate lies primarily in the decision concerning which LSPR workload category to use for representing a given production workload. Even when uncertain of the z/OS production environment, one can look at the LSPR-based worst case capacity relationship across all the z/OS workloads and assume that his capacity relationship will likely be no worse.

- **zAAP and zIIP LCPs associated with a z/OS partition:** Only a few representative workloads have been tested for such configurations (such measurements are beyond the scope of IBM's LSPR). Considerable variation can exist on how zAAP/zIIP CPs might actually be exploited over those benchmarked. Therefore, **zPCR** capacity results for these specialty engines should only be considered as **reasonable approximations of realizable capacity**.

zPCR must assume the same workload is running on specialty engines as on the parent z/OS partition. Since **zPCR** results represent the **Capacity Perspective**, consider, that, when a single zAAP/zIIP LCP is associated with a single z/OS GP LCP, both are considered to be 100% busy, and maximum switching costs apply. If, in reality, the zAAP were only 1% busy, General Purpose capacity would be more in line with that when no zAAP is configured, because the switching cost would be very low. To consider General Purpose capacity for zAAP/zIIP configurations that will be less than fully utilized, one might interpolate GP capacity based on expected zAAP/zIIP loading values between the two extremes projected by **zPCR**.

For **IBM Z** processors, when large z/OS partitions (many LCPs) are defined with associated zAAPs or zIIPs, capacity results for the zAAPs or zIIPs should be considered as rough approximations. This is because the only benchmarks available, with which to validate **zPCR** algorithms, involve much smaller configurations. If anything, capacity projections for such zAAP/zIIP configurations are likely conservative (low).

zPCR does provide the ability to define the amount of work that might run on zAAP or zIIP partitions. This definition is made in terms of the minimum capacity of the parent GP partition. By default, **zPCR** assumes the zAAP or zIIP capacity requirement to be equal to that of the parent GP partition (i.e., 100% zAAP/zIIP loading). By specifying zAAP/zIIP loading less than 100%, GP capacity will increase slightly due to reduced switching cost. zAAP and zIIP capacity is not affected.

- **z/VM, KVM, and Linux on IFL or General Purpose CPs:** Only one LSPR workload is run for each of these environments. z/VM has been measured on models up to the 40-way, which is the maximum that is supported. Linux has been measured on CPU models up to 16-way only, while more than 16 CPs can actually be supported.

In actuality, production workloads on z/VM, KVM, or Linux are likely to have capacity relationships that vary as much (or more) as those for the z/OS workloads. With z/VM, there is an extremely wide variety of ways the SCP might be exploited. Therefore, **zPCR** capacity results for z/VM, KVM, and Linux, while reasonably represented by the z/OS LSPR workloads measured, must be considered with a wider margin of error than those for z/OS when used to represent a production workload. This accuracy consideration is true regardless of whether General Purpose CPs or IFLs are used.

For both z/VM, KVM, and Linux, workload categories with names similar to those of z/OS are available for assignment to a partition. The capacity values displayed will be identical to those if z/OS (with the same workload category) was defined as the partition's SCP. These workload categories provide additional flexibility when considering z/VM and Linux, particularly when defined on an IFL.

A z/VM GP partition can have associated zAAP, zIIP, IFL, and ICF LCPs. **zPCR** supports all combinations. The partition number and name of the associated LCPs will be the same as the GP partition. The workload assignment for:

- An associated zAAP or zIIP partition will be that of the parent GP partition.
- An associated IFL partition may differ from that of the parent GP partition.
- An associated ICF partition must be CFCC.

No attempt is made to factor in a cost for z/VM to manage associated LCPs.

Therefore, the capacity values generated for zAAP, zIIP, IFL, and ICF partitions will be identical to that if the associated partitions were independent of the parent z/VM GP partition.

- **z/VSE on General Purpose CPs**: z/VSE environments are not measured for LSPR purposes. As a convenience, 5 z/VSE workload environments have been defined in **zPCR**, based on the z/OS LSPR measurements. Capacity results for z/VSE should be considered as **reasonable approximations of realizable capacity**.
- **CFCC on ICF of General Purpose CPs**: CFCC capacity data, carried internally, is used when configuring partitions as coupling facilities. This CFCC data is representative of the CFCC level that became available at the time of the processor family announce, and, therefore may not be totally consistent across all processor families. In addition, CFCC numbers are based on only a single operating environment. Therefore, CFCC capacity results in **zPCR** should only be considered as **reasonable approximations of realizable capacity**.

The potential accuracy of **zPCR** capacity results varies considerably across the various SCPs supported, as discussed above, with z/OS results being the most reliable. When various partition types are configured on the same LPAR host, **zPCR** also considers their impact on each of the other partitions. Such impacts are generally small; there should be little concern that potential accuracy is reduced for any given partition running on such configurations.

CP Calculator

A side function provided in **zPCR**, identified as **Capacity Planning Calculator**, can be accessed from the **CPcalculator** menu-bar item on the **Control Panel** window.

The only function currently implemented is the **zAAP Capacity Estimator**, used to estimate the capacity available to and needed by a z/OS partition when migrating from a purely GP CP environment to one with associated zAAP (or zIIP) CPs.

While **CPcalculator** function can be used in parallel with any normal **zPCR** function.

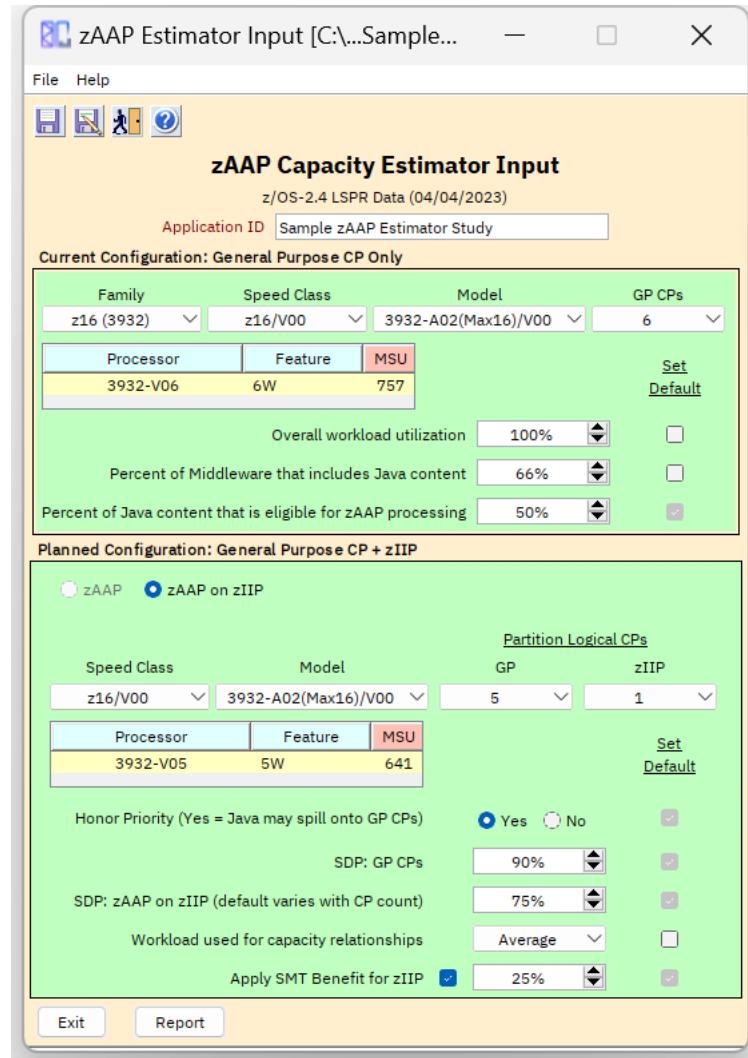
Inputs to **CPcalculator** functions are not retained as part of a **zPCR** saved study. Rather, you can save your inputs as a special file type, unique to the function. In this way, any previously saved **CPcalculator** function input can be loaded in conjunction with any **zPCR** study.

Note that **Reference-CPU** settings, when needed to provide capacity results, are derived from **zPCR**'s current **Reference-CPU**. Any change to the **Reference-CPU** will immediately be reflected in capacity values revealed on **CPcalculator** windows. The **Reference-CPU** settings are not retained as part of a saved **CPcalculator** function.

zAAP Capacity Estimator Input

Capacity Planning Calculator

The **zAAP Capacity Estimator** is accessed from the **Control Panel** window by clicking on the **CPcalculator → zAAP Capacity** menu item. The **zAAP Capacity Estimator Input** window appears.



This function is intended to preview the effects on capacity when moving a z/OS image from a purely General Purpose CP environment to one that includes and can exploit zAAP CPs or zIIP CPs using **zAAP on zIIP** capability. Capacity comparisons should be considered as being appropriate for the current LSPR measurement data (z/OS-2.4).

Capacity values presented are based on **zPCR**'s currently defined **Reference-CPU** and its scaling metrics. If you change the **Reference-CPU** or its scaling-factor while the **zAAP Capacity Estimator** is active, capacity results will immediately reflect the change.

Capacity values presented with this function assume a single z/OS partition and its associated zAAP (or zIIP) partition, without regard to any other partitions that may be active. To evaluate the effects on capacity when defining multiple partitions sharing the same RCP resources, you must use the **LPAR Configuration Capacity Planning** function of **zPCR**.

All inputs are defined on the **zAAP Capacity Estimator Input** window. You may enter an **Application ID** by clicking on the entry area and keying in the name. If entered, the application name will appear on the **Report** window and the associated graphs and table output.

The codes below identify the type of entry for the input fields detailed below:

-  **Entry field**; Double click field to open, key in text, and press **Enter**.
-  **Dropdown list**; Click field to access dropdown list and make selection.
-  **Checkbox**; Click field to check or un-check.
-  **Spin button**; Click top to increase value and bottom to lower value.
-  **Radio Button**; Make selection by clicking on desired button.

The controlling input fields are detailed below.

Current Configuration: General Purpose CP Only

Describe the z/OS image (a single partition) which is being considered for exploiting zAAP CPs or zIIP CPs using "zAAP on zIIP" capability. The capacity requirement must be characterized in terms of the current **General Purpose CP Only** configuration. You must provide a processor family, speed class, and model. In addition, supply the overall utilization represented by the workload, an estimate of the percent of the workload that has Java content, and an estimate of the percent of the Java content workload that will be eligible for zAAP processing.

-  **(Processor) Family** - This field is initially displayed as a preset value. Click on the dropdown icon for a list of processor families, and make a selection. All **IBM Z** families, z9 and later are available for selection.
-  **(Processor) Speed Class** - This field is initially displayed as the full speed setting for the processor family chosen. Click on the dropdown icon for a list of processor speed settings for the family, and make a selection.
-  **(Processor) Model** - This field is initially displayed as the default value for the processor family chosen. Click on the dropdown icon for a list of processor models in the family, and make a selection. The model determines the maximum number of CPs that can be defined.

- **GP CPs** - This field is initially set to 1. Click on the dropdown icon to change the number of General Purpose LCPs active for the z/OS image. The value can be in the range of 1 to 128; only numbers valid for the selected processor model are shown in the list.
The **Processor** (as identified by z/OS), **Feature** (N-way), and **MSU** rating are redisplayed in the table below.
- **Overall workload utilization** - Click on the spin button to select a utilization value (or click on the field and key in the value desired). Click the **Set Default** checkbox to restore the default setting of 100%.
- **Percent of Middleware that includes Java Content** - Click on the spin button to select a percent (or click on the field and key in the value desired). Click the **Set Default** checkbox to restore the default setting of 50%. This value represents the percent of processor's utilization represented by all the applications that have Java content (e.g., WebSphere).
- **Percent of Java content that is eligible for zAAP processing** - Click on the spin button to select a percent (or click on the field and key in the value desired). Click the **Set Default** checkbox to restore the default setting of 50%. This value represents the percent of the Java content activity expected to run under the Java Virtual Machine (JVM).

The workload's total capacity requirement is computed as the capacity rating of the selected processor times the **Overall Workload Utilization**. The percent of Java content determines the amount of the total workload that will incur zAAP processing costs. The **Percent of Java content that is eligible for zAAP processing** defines the amount of workload that can actually be run on zAAP LCPs.

Planned Configuration: General Purpose CP+zAAP (or zAAP on zIIP)

The intended new General Purpose plus zAAP processing environment is described, using the following fields:

- **zAAP** or **zAAP on zIIP** - Set the appropriate radio button for the intended environment. For an LPAR host that does not support zAAP (z13 and later processors), **zAAP on zIIP** will be preset.
When **zAAP on zIIP** is selected, **zIIP** will be used as the hardware designation, while the workload remains zAAP content.
- The **Processor Family** is assumed to be the same as the **Current Configuration**. Click on the dropdown icon for a list of processor families. Only processors within the currently designated family may be selected.
- **(Processor) Speed Class** - This field is initially displayed as the full speed setting for the processor family chosen. Click on the dropdown icon for a list of processor speed settings for the family, and make a selection.
- **(Processor) Model** - This field is initially displayed as the default value for the processor family chosen. Click on the dropdown icon for a list of processor models in the family, and make a selection. The model determines the maximum number of CPs that can be defined.

The **Processor Family** cannot be changed as it must be the same as that of the **Current Configuration**. The **Speed Class** and/or **Model** may be changed.

- **GP (LCPs)** - This field is initially set to 1. Click on the dropdown icon to change the number of GP LCPs active for this z/OS image. Only numbers valid for the selected processor model are shown in the list. At least 1 zAAP (or zIIP) LCP must be assumed; the number of GP+zAAP LCPs cannot exceed the CPs available with the designated processor model.

The **Processor** (as identified by z/OS), **Feature** (N-way), and **MSU** rating are redisplayed in the table below.

For the GP+zAAP environment, the speed class and number of GP LCPs may both be changed; the goal is to lower the MSU rating (by reducing the GP CP count, meaning software cost are lower) while providing adequate zAAP CPs capacity to support the zAAP eligible work as defined.

- **zAAP or zIIP (LCPs)** - This field is initially set to 1. Click on the dropdown icon to change the number of zAAP (or zIIP) LCPs active for this z/OS image. The values allowed are in the range of **1** to **n**, where **n** cannot exceed the total CPs available minus number of GP CPs defined. Note that capacity results for configurations with a large number of zAAP (or zIIP) LCPs may be less reliable, due to the absence of supporting measurement data.
- **Honor Priority (Yes = Java may spill onto GP CPs)** - Click on the **Yes** or **No** radio button for the desired setting. Click the **Set Default** checkbox to restore the default setting of **Yes**.
- **E SDP: GP CPs** - Click on the spin button to set the upper limit for utilization on the General Purpose CPs. Click the **Set Default** checkbox to restore the default setting of 90%.
- **E SDP: zAAPs (default varies with CP count)** - Click on the spin button to set the upper limit for utilization on the zAAPs. Click the **Set Default** checkbox to restore the default setting. Note that the default setting varies with the number of zAAP (or zIIP) LCPs assigned.
- **Workload used for capacity relationships** - Click on the dropdown to set the z/OS workload category to be used for capacity relationships between models and for zAAP (or zIIP) LCPs. Click the **Set Default** checkbox to restore the default setting of **Low-Avg**.
- **Apply SMT Benefit for zIIP** (z13 and later processors only) - Check the box to activate SMT. A default value of 25% will be initially applied to the zIIP capacity values. Use the spin button to set the **SMT Benefit** to any value between 0% and 60%. Click the **Set Default** checkbox to restore the default setting of 25%.

Capacity available to z/OS is determined by the General Purpose processor family, its speed class, and the number of zAAP (or zIIP) LCPs defined. The combination of GP and zAAP (or zIIP) LCPs is viewed as a single z/OS N-way image for capacity purposes. That capacity is then apportioned to the actual number of CPs of each type configured. Nominal overheads associated with zAAP dispatching decisions are also taken into account.

For capacity planning purposes, the workload of a z/OS image is considered to have two separate components.

- **Normal content** is the total combination of those applications that have no zAAP eligible work. When there are no zAAP (or zIIP) LCPs defined, the entire workload is considered normal content.
- **Java Content** is the total combination of those applications that contain zAAP eligible work (e.g., WebSphere).

Capacity for the zAAP (or zIIP) LCPs and for the parent GP partition is determined as follows:

- **Java content** work sees capacity relating to the sum of General Purpose CPs and zAAPs (i.e., N-way of the combination) plus overhead related to switching work between the two.
- **Normal content** work sees capacity relating to the General Purpose CPs (N-way of the General Purpose CPs), as if no zAAP (or zIIP) LCPs were defined.

zAAP (or zIIP) capacity is always associated with that of Java content workload.

Within a single invocation of **zPCR**, all **zAAP Capacity Estimator** inputs are retained as the most recent value entered. To save all of the inputs for future use, click on **Save** or **Save as** under **File** on the menu-bar (inputs are not captured as part of a **zPCR** study file, but rather as a unique file extension of **zPCRZAP**). To restore a previously saved set of inputs, click on **Load** under **File** on the menu-bar. To restore all of the inputs to their initialization state, click on **New** under **File** on the menu-bar.

Click the **Report** to open the **zAAP Capacity Estimator Report** window. For detail concerning the report, see [zAAP Capacity Estimator Report](#).

Click the **Exit** toolbar icon or the **Exit** button to terminate this **CPcalculator** function.

Menu-bar

File

New	Refresh inputs to initialization state
Load	Open a previously saved set of zAAP Capacity Estimator inputs
Save	Save this set of inputs (must already be a titled study)
Save as...	Save this set of inputs with a new name
Exit function (Ctrl+E)	Terminate zAAP Capacity Estimator . Exit can also be invoked using the icon on the tool bar.
Exit tool (Ctrl+Q)	Terminate zPCR execution. If a zPCR study has been started, you will be asked for confirmation concerning the saving of those inputs.

Help

zAAP Capacity Estimator Report

Capacity Planning Calculator

The **zAAP Capacity Estimator Report** window will normally appear alongside the input window. When not in view, Click the **Report** button on the input window. Changes made on the input window will be immediately reflected in the results on the report window. For detail concerning inputs, see [zAAP Capacity Estimator Input](#).

zAAP Estimator Output [C:\...Sample zAAP Estimator Study.zpcrzap]

zAAP Capacity Estimator Report
z/OS-2.4 LSPR Data (04/04/2023)

Study ID: Sample Study for Advanced zPCR Use - Application ID: Sample zAAP Estimator Study
Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration
Workload used to establish capacity relationships: Average

Current Configuration: z16/V00, 3932-A02(Max16)/V00 with 6 GP CPs (3932-V06)

Configured CPs	Capacity		Percent Used	Comments
	Available	Used		
06 General Purpose CPs	6,622	6,622	100%	Total workload capacity requirement
	2,252	2,252	34%	Portion of workload that has no Java content
	4,371	4,371	66%	Portion of workload containing Java content
	2,185	2,185	50%	Portion of Java content that is zAAP eligible

Planned Configuration: z16/V00, 3932-A02(Max16)/V00 with 5 GP CPs (3932-V05) and 1 zIIP

Configured CPs	Capacity		Percent Needed	Comments
	Available	Needed		
01 zIIP	2,369	2,185	92%	zIIP CP projected utilization is 92%
Not possible to fulfill		0		
Not fulfilled at 75% SDP		408		zAAP eligible work likely to spill at 75% SDP
05 General Purpose CPs	5,521	4,437	80%	GP CP projected utilization for GP work is 80%
Spill from zIIP		408		zAAP eligible work likely to spill to GP CPs
Total		4,845	88%	GP CP projected utilization = 88%

An SMT Benefit of 25% has been applied to zIIP CP capacity

Compare Solutions	Current			Delta
	GP-only	GP	zIIP	
CPs Configured	6	5	1	0%
Capacity(MIPS)	6,622	5,521	2,369	19%
MSU Rating	757	641	N/A	-15%

Graphs

[Capacity vs MSU Rating](#) [Utilization](#)

Close

For significant configuration changes, such as upgrading the processor family, all capacity comparisons are subject to a +/-5% margin-of-error

All capacity values presented on the **zAAP Capacity Estimator Report** window and the related graphs are based on the currently defined **Reference-CPU**. If you change the **Reference-CPU** or its scaling-factor while the **zAAP Capacity Estimator** is active, capacity results will immediately reflect the change.

Current Configuration Group Box

A summary of the workload as defined to the GP only configuration is presented in the upper group box. The precise configuration defined is summarized in the group box title.

The **Overall workload utilization** is applied to the selected processor's capacity rating to reveal the capacity being consumed by the entire workload. The **Portion of**

workload with Java content is applied to the capacity being consumed to reveal the capacity consumed by work containing Java content. And the **Java content eligible for zAAP processing** is applied to the Java content capacity to reveal the capacity that is eligible to run on zAAPs.

Planned Configuration Group Box

An analysis of capacity for the GP+zAAP (or zIIP) configuration is presented in the lower group box. The precise configuration defined is summarized in the group box title.

Since only a portion of the total workload can be eligible to be executed on zAAP (or zIIP) LCPs, capacity for such a configuration must be viewed as two distinct entities, one for the zAAPs (or zIIPs) and one of the General Purpose CPs.

For zAAPs, the zAAP capacity needed is compared to the zAAP capacity available. If **Honor Priority** is set to **No**, then the zAAP capacity is deemed either adequate or inadequate. If **Honor Priority** is set to **Yes**, then any zAAP capacity needed that would exceed the zAAP SDP is assumed to spill over to the General Purpose CPs.

For General Purpose CPs, the General Purpose capacity needed is compared to the General Purpose capacity available, and General Purpose capacity is deemed either adequate or inadequate. If **Honor Priority** is set to **Yes**, then any unfulfilled zAAP capacity (at zAAP SDP) is added to the General Purpose capacity needed, and that result is compared to the General Purpose capacity available.

For both the zAAPs and the General Purpose CPs, the comments will indicate that their respective SDP settings are exceeded.

When the host processor is defined as a z13 or later processor, an indication of zIIP SMT benefit is displayed at the bottom of the **Planned Configuration** group box.

Compare Solutions Group Box

Below the configuration results is a comparison of the solutions, based on total CPs, Capacity, and MSU Rating. A percent delta is computed for each.

Summary

Capacity results are generated directly from [zPCR's LPAR Capacity Planning Function](#). The **Current** and **Planned Configurations** are each modeled as a single partition configuration. Capacity results are then extracted and presented in the **zAAP Capacity Estimator Report** window. zAAP (or zIIP) utilization is used to determine zAAP/zIIP loading value used in the model. As zAAP (or zIIP) utilization decreases, the switching cost diminishes, resulting in somewhat increased GP capacity.

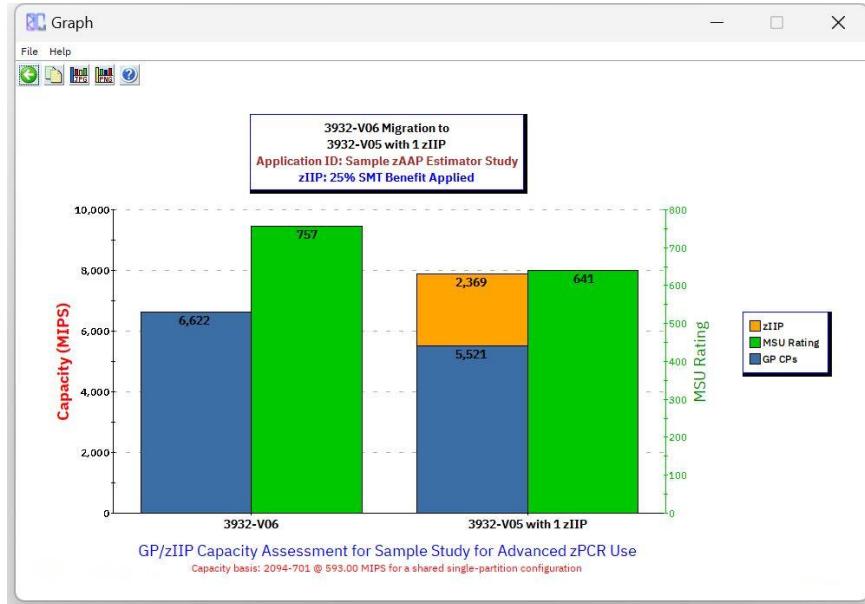
Keep in mind that this function considers the LCPs for only a single GP+zAAP (or GP+zIIP) pairing. If in reality, multiple partitions would compete for the same RCP resources, capacity results will be different. In such cases, [zPCR's LPAR Capacity Planning Function](#) should be used to model the entire LPAR configuration.

The report can be captured as an HTML file by clicking the **Output to HTML** toolbar icon. As an alternative form of output, you can capture either of the two windows as a bitmap by keying **Alt+PrintScreen** when the desired window is the active window. A copy of the window is written to the windows clipboard, which may subsequently be pasted into a document.

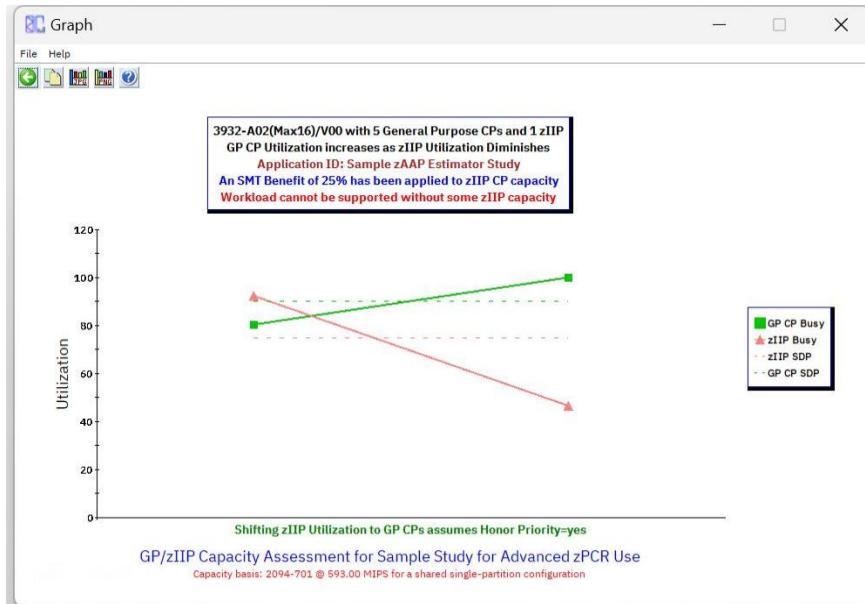
Graphs

Two graphs are available with this function.

1. Click the **Capacity vs MSU Rating** button for a bar graph showing the previous General Purpose capacity and the new GP+zAAP environment capacity. The zAAP environment capacity is shown as a stacked bar, with separate values for the General Purpose CPs and for the zAAPs.



2. Click on **Utilization** for a line graph depicting utilization of the GP CPs as utilization of the zAAP CPs declines. Dotted lines showing the currently set SDP values will appear if they fall within the utilization range plotted.



Important Considerations

1. Capacity values presented here assume a single z/OS partition and its associated zAAP (or zAAP on zIIP) partition, without regard to other partitions that may be active. To evaluate the effects on capacity when defining multiple partitions sharing the same RCP resources, you must use the ***LPAR Configuration Capacity Planning*** function of ***zPCR***.
2. The ***Capacity Available*** result for GP CPs of the ***Planned Configuration*** is based on the computed zAAP utilization. As zAAP utilization diminishes below 100%, the GP CP cost associated with managing them also diminishes, resulting in slightly higher GP capacity. If the zAAP utilization were 0%, the GP CP available capacity would align with that of the GP CPs when no zAAPs are configured. For configurations where the GP CPs are less than full speed, the difference in zAAP speed is considered when determining the GP CP capacity results.

Window Controls

Click the ***Return*** tool bar icon to close the report window.

Click the ***Exit*** button on the ***zAAP Capacity Estimator Input*** window to terminate this ***CPcalculator*** function.

Support

Problems and Suggestions

Every effort has been made to make **zPCR** a useful and intuitive capacity planning application. Should you detect problems or desire to make suggestions for future enhancements, please send a note to the following Capacity Planning Support (CPS) address:

E-mail: zpcr@us.ibm.com

Specify "Subject: **zPCR** (Processor Capacity Reference)." Describe the problem or suggestion as thoroughly as possible, and suggest any possible solutions. Please include your name, your location, and phone number should additional information be required. If appropriate, send the saved study file (*.zpcr) representing the particular **zPCR** inputs with which you are working so that your scenario can be easily recreated.

All comments and suggestions will be considered.

Maintaining Currency

As new or additional LSPR data becomes available, **zPCR** will be updated. In addition, as further experience is gained, algorithms related to the **LPAR Configuration Capacity Planning** function may be revised.

It is your responsibility to ascertain that you are always working with the most current LSPR data and **zPCR** algorithms. From the menu-bar on the **Control Panel** window, click **Help** → **Check for updates**. If the version is not current, go to the Web site from which you obtained **zPCR** and download the newer version.

External Study File

zPCR provides the ability to input LPAR configurations as follows:

1. Manual input
2. A saved **zPCR** study
3. EDF input created with the CP3KEXTR program from z/OS SMF data, or created with the CP3KVMXT program from z/VM Monitor data.
4. RMF report (from z/OS SMF) using the **Partition Data Report** and **CPU Activity Report**.

There is one additional method to define LPAR configurations to **zPCR**, using a source referred to as an **External Study File** (ESF). This source would typically be created by a vendor's software product, designed to read hardware/software performance data such as z/OS SMF. Information extracted by that software would then create a file containing the specific XML tags required by **zPCR**. The resulting file can then be loaded directly into **zPCR** similar to a study file, resulting with the LPAR host and its partition configuration entirely defined.

If you are interested in obtaining documentation concerning the format of a **zPCR External Study File**, send an E-mail request to:

zpcr@us.ibm.com

Definition of Terms

Processor Terminology

IBM Z	IBM Z processor models, including: z17: 9175-ME1 z16: 3931-A01 and 3932-A02/AGZ z15: 8561-T01 and 8562-T02 z14: 3906 and 3907 z13: 2964 and 2965 zEnterprise EC12 (2827) and BC12 (2828) zEnterprise 196 (2817) and 114 (2818) System z10 EC (2097) and z10 BC (2098) System z9 EC (2094) and z9 BC (2096) zSeries 990 (2084) and 890 (2086) zSeries 900 (2064) and 800 (2066)
LinuxONE	LinuxONE processor models, including z17: Emperor 5 9175-ML1 z16: Emperor 4 3931-LA1 and Rockhopper 4 (3932-LA2/AGL) z15: LinuxONE III 8561-LT1 and 8562-LT2 z14: Emperor II 3906 and Rockhopper II 3907 z13: Emperor 2964 and Rockhopper 2965 z12: Rockhopper 2828
IBM Z	IBM z17, z16, z15, z14, z13, z13s models; also includes models below.
zEnterprise EC12	IBM zEC12 models
zEnterprise BC12	IBM zBC12 models
zEnterprise 196	IBM z196 models
zEnterprise 114	IBM z114 models
System z10	IBM System z10 EC and z10 BC models
System z10 EC	IBM System z10 Enterprise Class models
System z10 BC	IBM System z10 Business Class models
System z9	IBM System z9 EC and z9 BC models
System z9 EC	IBM System z9 Enterprise Class models
System z9 BC	IBM System z9 Business Class models
zSeries	IBM eServer z990, x900, z890, or z800 processor models
S/390	IBM System/390 processor models (G6, G5, G4, G3.G2, MP3000, MP2000, 9021, and 9121). G6 and G5 processors are also considered to be zSeries processors. S/390 processors are no longer supported in zPCR .

Processor Family Naming Convention

z9	IBM Z: z9 EC (2094) and z9 BC (2096)
z10	IBM Z: z10 EC (2097) and z10 BC (2098)
z11	IBM Z: z196 (2817) and z114 (2818)
z12	IBM Z: zEC12 (2827) and zBC12 (2828) LinuxONE: Rockhopper (2828)
z13	IBM Z: z13 (2964) and z13s (2965) LinuxONE: Emperor (2964) and Rockhopper (2965)
z14	IBM Z: z14 (3906) and z14-ZR1 (3907) LinuxONE: Emperor II (3906) and Rockhopper II (3907)
z15	IBM Z: z15-T01 (8561) and z15-T012 (8562) LinuxONE: Emperor III (8561) and Rockhopper III (8562)
z16	IBM Z: z16-A01 (3931) and z16-A02/AGZ (3932) LinuxONE: z16-LA1 Emperor 4 (3931) and z16-LA2/AGL Rockhopper 4 (3932)
z17	IBM Z: z17-ME1 (9175) LinuxONE: z17-ML1 Emperor 5 (9175)

Hardware Terminology

CEC	Central Electronic Complex (the computer)
CP	A PU that has been configured as a General Purpose CP, a zAAP, a zIIP, an IFL, or an ICF.
CPC	Central Processing Complex (the computer; same as CEC)
CPU MF	CPU Measurement Facility (hardware counter data captured in z/OS SMF type 113 records or z/VM Monitor records). This information is used by zPCR to choose the LSPR workload category that most reliably represents the production workload of a partition.
General Purpose CP	Real CPs (also referred to as GP CPs) that are eligible to run z/OS workloads (software pricing is based on MSU rating). A General Purpose CP can run any SCP that supports z/Architecture, including z/OS, z/VM, z/VSE, KVM, Linux, SSC, or CFCC.
GP	An abbreviation used for “General Purpose” (i.e., GP CPs or GP partitions). Other CP types include zAAP, zIIP, IFL, and ICF.
GP CP Pool	All of the active General Purpose CPs on a CPC. They are managed by LPAR as a single pool of Real CPs, in support of all defined General Purpose partitions. General Purpose partitions can only be dispatched on the CPs in this pool.
ICF	Integrated Coupling Facility (a special purpose CP). Each ICF is a Real CP on the CPC. An ICF can only run the CFCC control program, at the CFCC level appropriate for the hardware family.
ICF Pool	All of the CPs in the ICF pool.
IFL	Integrated Facility for Linux (a special purpose CP). Each IFL is a Real CP on the CPC. An IFL can only run z/VM (running Linux guests), KVM, Linux, or SSC.
IFL Pool	All of the CPs in the IFL pool. (GP, zAAP, zIIP, IFL, and ICF. CPs each operate as a separate pool).
LCP	Logical CP (engines defined to a logical partition)
LCP:RCP Ratio	The ratio of all the partitions' LCPs to all of the processor's Real CPs within a single CP pool. Dedicated CPs are generally removed from the Real CP count of the pool before computing the ratio.
LP	Logical partition
RCP	Real CP (an engine on the real hardware)

Multi-Image	Capacity represents the combination of multiple partitions, each running a copy of the same SCP and workload. Partition configurations assumed are intended to be typical for the processor model represented. The value of multi-image tables lies in the ability to represent capacity on models that cannot be supported with a single copy of an SCP.
PU	Processor Unit which can be client configured as General Purpose CPs, zAAP CPs, zIIP CPs, IFL CPs, and/or ICF CPs.
SCP	System Control Program (i.e., z/OS, z/VM, z/VSE, KVM, Linux, SSC, and CFCC)
Single-Image	Capacity represents a single shared partition running the SCP and workload, with all the processor CPs assigned. SCPs are limited in the number of CPs supported; therefore, capacity can be represented only on models with CPs in the supported range.
zAAP	z Application Assist Processor (a special purpose CP). Each zAAP is a Real CP on the CPC for all z9 through z12 processors only. zAAP LCPs, using zAAP Real CPs, can only run certain Java code on behalf of its parent z/OS-1.6 (or higher) partition. On z13 and later processors, zAAP workload must be run on zIIP LCPs (zAAP on zIIP must be enabled).
zAAP Pool	All of the CPs in the zAAP pool. zAAPs can be configured only on z9 through z12 processor models, and always operate as a separate CP pool
zIIP	z Information Integration Processor (a special purpose CP). Each zIIP is a Real CP on the CPC for all z9 and later processor models). zIIP LCPs, using zIIP Real CPs, can only run certain redirected Db2 code on behalf of its parent z/OS-1.6 (or later) partition. Note: When zAAP on zIIP is enabled for a partition, the zAAP workload will run on the zIIP LCPs.
zIIP Pool	All of the CPs in the zIIP pool. zIIPs can be configured on all z9 and later processor models, and always operate as a separate CP pool.

Hardware Features Terminology

Initial Capping	When this capping is set on, a partition's Maximum Capacity value is set equal to its Minimum Capacity value. Sometimes referred to as Hard Capping or Normal Capping.
Absolute Capping	(A feature on z12 and later processor only). An Absolute Capping value is specified as a fractional value between zero and the number of LCPs defined to the partition (format N.nn). This value has the effect of reducing the partition's Maximum Capacity value to that of the number of LCPs specified for Absolute Capping.
Group Absolute Capping	Applies capping for a group of partitions. zPCR does not support this form of capping.
SMT	Simultaneous Multi-Threading is a feature on z13 and later processors and is applicable only for zIIP and IFL CPs. LSPR data is published based on single-threaded workload. To evaluate expected zIIP and IFL capacity, an SMT Benefit value can be applied manually for each zIIP or IFL partition. Suggested default SMT Benefit value for zIIP CPs is 25%; for IFL CPs on z17, z16, z15, z14 is 25% and on z13 is 20%. For workloads actually measured on z13 and later processors, the actual SMT Benefit obtained can be determined from SMF data. Such measured SMT Benefit values can be transferred into zPCR via EDF or RMF.
SRB	System Recovery Boost (SRB) is a feature available on the IBM z17, z16, and z15. It is used to accelerate partition startup (IPL) and shutdown or to accelerate specific Sysplex recovery events. During the boost period, GP capacity may be increased by temporarily enabling sub-capacity GP engines to run at full speed (Speed Boost) and/or increasing parallelism by enabling zIIP engines to run any work for the partition (zIIP Boost). The length of the boost period is limited based on the purpose (class) of the boost. Optionally, additional zIIP engines can be activated for up to 6 hours to provide additional zIIP capacity using the SRB Upgrade priced feature. SRB information can only be transferred into zPCR via EDF or RMF input. Performance intervals that include boost activity should not be used for capacity planning or sizing purposes due to the temporary nature of the adjustments made to the LPAR configuration, which are not representative of a typical production workload interval.

Capacity Rating Terminology (Scaling-Metric)

ITRR	Internal Throughput Rate Ratio The capacity of a processor expressed relative to the capacity the processor assigned as the Reference-CPU for a specific workload. The scaling-factor of the Reference-CPU may be set to any reasonable value, but is typically set to 1.00.
MIPS	Millions of Instructions per Second (this literal translation is no longer applicable but the acronym remains commonly used). The capacity of a processor expressed relative to the capacity the processor assigned as the Reference-CPU for a specific workload. The scaling-factor of the Reference-CPU may be set to any reasonable value, but is normally set to a commonly accepted value for that processor.
PCI	Processor Capacity Index The capacity of a processor expressed relative to the capacity the processor assigned as the Reference-CPU for a specific workload. The scaling-factor of the Reference-CPU may be set to any reasonable value, but is normally set to a commonly accepted value for that processor. This term is useful for those who oppose the MIPS terminology.
Capacity Index	A wordy version of PCI .

Note that the scaling-metric assigned to the **Reference-CPU** is not limited to these 4 items. Any text up to 12 characters can be assigned as the scaling-metric.

System Control Program (SCP) Terminology

z/OS	Any version of z/OS. zPCR always uses the most current z/OS LSPR data to generate capacity tables. However, algorithms related to partitioning are sensitive to the version of z/OS specified.
OS/390	Any version of OS/390. zPCR no longer includes OS/390 LSPR data, which was the sole source for making comparisons that included S/390 processors.
MVS	Includes any version of z/OS or OS/390.
z/VM	Any version of z/VM. zPCR always uses the most current z/OS LSPR data to generate capacity tables. However, algorithms related to partitioning are sensitive to the version of z/VM specified.
z/VSE	Any version of VSE. zPCR considers VSE generically, i.e., not specific to any particular release.
KVM	Any version of KVM. zPCR considers KVM generically, i.e., not specific to any particular release.
Linux	Any version of Linux. zPCR considers Linux generically, i.e., not specific to any particular release.
SSC	Secure Service Container.
CFCC	Coupling Facility Control Program , which normally runs in an ICF partition. It may also be run in a General Purpose partition.
<u>Not-IPL'd</u>	Partition is defined but has no active SCP.

Workload Terminology Related to z/OS and zAAP/zIIP CPs

Normal content	The total combination of those applications that include no zAAP eligible work. When there are no zAAP LCPs defined, the entire workload is considered normal content.
Java content	The percent of the overall workload that represents Java exploiting applications (e.g., WebSphere), and therefore have code that can exploit zAAP CPs.
zAAP eligible	The percent of Java work that can execute on zAAP CPs.
Db2 content	The percent of the overall workload that represents applications (in this case, Db2) that can exploit zIIP CPs.
zIIP eligible	The percent of Db2 work that can execute on zIIP CPs.

zAAP/zIIP Loading	A percentage representing the amount of zAAP or zIIP capacity that is being used. zPCR capacity algorithms assume full loading (i.e., 100% utilization) for all CP types. Specifying zAAP or zIIP utilization values less than 100%, will produce a slight increase for the parent GP partition's capacity, due to a reduction in switching cost.
--------------------------	--

Files Recognized by zPCR

EDF	<p>An EDF (Enterprise Data File) is generated by running the CP3KEXTR program against z/OS SMF data on a IBM Z processor. The SMF data (binary and voluminous) is reduced to a size that can be downloaded to a PC (preferred file extension = edf), while remaining useful for capacity planning purposes.</p> <p>For z/VM, create an EDF by running the CP3KVMXT program against CP Monitor Data on a IBM Z processor.</p> <p>An EDF can be used by zPCR to define the LPAR host and its entire partition configuration or to copy specific partition definitions into a currently active LPAR configuration. An advantage of the EDF is that CPU MF hardware counter data, if captured, can be used by zPCR to choose the LSPR workload category that best represents a z/OS partition's production workload. In addition, for each partition represented by an EDF, information about the use of HiperDispatch, the number of parked LCPs, and SMT benefit (z13 and later only) is available.</p>
RMF	<p>A Resource Measurement Facility report is generated using standard z/OS software against z/OS SMF data on a IBM Z processor. The report generally consists of several selectable sections. For zPCR purposes, the CPU Activity Report and the Partition Data Report are necessary. When downloaded to a PC (preferred file extension = rmf), these reports can be used by zPCR to define the LPAR host and its entire partition configuration or to copy specific partition definitions into a currently defined LPAR configuration. In addition, for each partition represented by a CPU Activity Report, information about the use of HiperDispatch, the number of parked LCPs, and SMT benefit (z13 and later only) is available.</p>

Study File

A **Study file** may include up to 10 sets of LPAR host and partition configuration definitions. The file extension is **zPCR**. Study files can be reloaded into **zPCR** for future reference or used for the purpose of copying previously defined partitions into a currently active **zPCR** study.

External Study File

A file created by an OEM software product that reads hardware/software performance data such as z/OS SMF. Information extracted is used to generate specific XML tags required by **zPCR**. The resulting file can then be loaded directly into **zPCR** similar to a study file, to define the LPAR host and its partition configuration. For more information, see [External Study File](#).

