zPCR User's Guide

Processor Capacity Reference
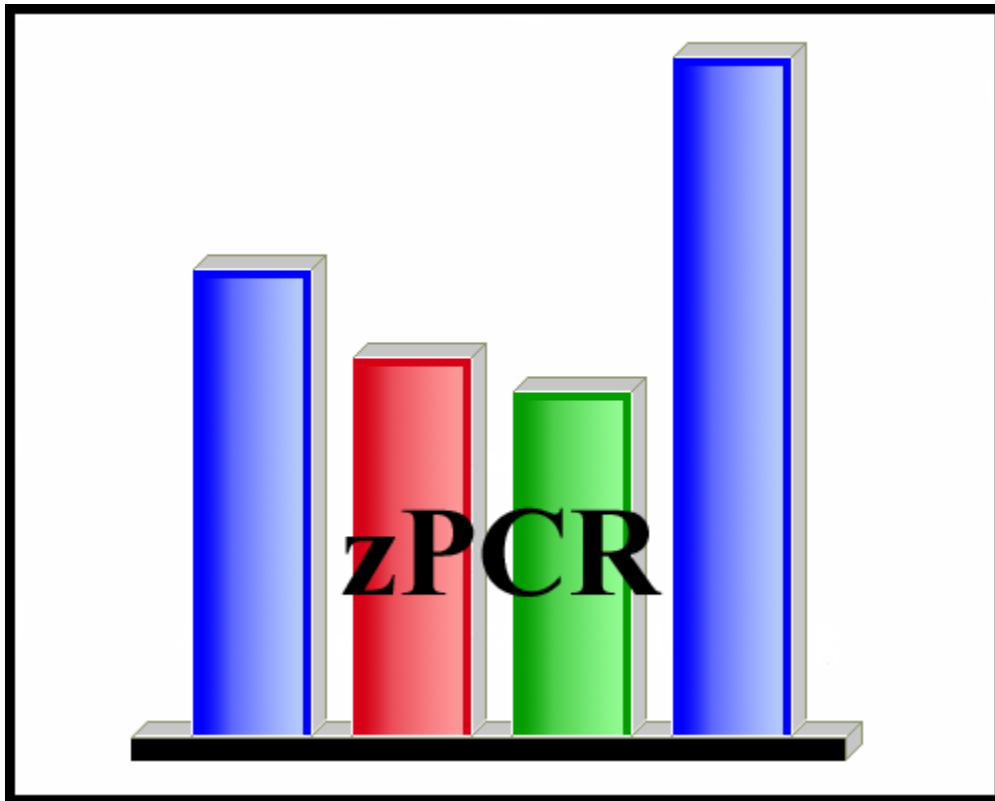
for

IBM System z

Version 8.5

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

CICS®	System z®
DB2®	System z9®
Enterprise Storage Server®	System z10®
eServer	VM/ESA®
IBM®	WebSphere®
IMS	z/Architecture®
MVS	z9®
OS/390®	z10
Parallel Sysplex®	zEnterprise
Processor Resource/Systems Manager	z/OS®
Resource Measurement Facility	z/VM®
RMF	z/VSE®
System/390®	zSeries®

The following are trademarks or registered trademarks of other companies.

Linux is a registered trademark of Linus Torvalds.

Java and all Java-related trademarks and logos are trademarks of Sun Microsystems, Inc., in the United States and other countries

Microsoft, Windows 7, Vista, Windows XP, and Windows 2000 are registered trademarks of Microsoft Corporation.

InstallShield 2011 Premier Edition is a registered trademark of Flexera Software, Incorporated.

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

Table of Contents

Introduction

zPCR

Processor Capacity Reference
for
IBM System z

zPCR is a PC-based productivity tool that runs under Windows 7 or Windows XP. It is designed to provide capacity planning insight for IBM System z processors running various SCP/workload environments using various LPAR partition configurations. Capacity results are based primarily on IBM's published LSPR data for System z. Generally, we have found the accuracy of using **zPCR** with an appropriate z/OS workload category to be well within the range of ±5% for reasonable LPAR configurations.

zPCR is based on a single version of LSPR data, normally published at the time of a new IBM System z processor announcement. The most recent LSPR data, identified as **z/OS-1.13**, adds the **IBM zEnterprise BC12** (**zBC12**) family to the entire set of System z processor families. The specific processor families included are:

- **zEC12** (2827)
- **z196** (2817)
- **z10**-**EC** (2097)
- **z9**-**EC** (2094)
- **z990** (2084)
- **z900** (2064)

- **zBC12** (2828)
- **z114** (2818)
- **z10-BC** (2098)
- **z9-BC** (2096)
- **z890** (2086)
- **z800** (2066)

LSPR capacity values used by **zPCR** are based entirely on the most recent set of LSPR data, which includes various workload environments run under z/OS-1.13 on up to 100 CPs. Additional LSPR workloads for z/VM (one multiple Linux guest workload, measured on up to 32 CPs) and for native Linux (one workload on up to 16 CPs) are also measured. For **zPCR** purposes however, the z/OS measurement data is considered to be representative for both z/VM and Linux.

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

- **Low**   Low relative nest intensity
- **Low-Avg**   Defined midway between *Low* and *Average* (**zPCR** only)
- **Average**   Average relative nest intensity
- **Avg-High**   Defined midway between *Average* and *High* (**zPCR** only)
- **High**   High relative next intensity

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.

Partition configurations can be created in **zPCR** several ways:
- Automatically from a z/OS or z/VM EDF
- Automatically from a z/OS RMF Partition Data Report (PDR)
- Manual entry
- By cloning a previously created configuration

z/OS and z/VM production workloads running on **zEC12**, **zBC12**, **z196**, **z114**, 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 workload categories mentioned above.

When CPU-MF is not available, a default workload category 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 primary 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, z/VM, or Linux.

   - The **Multi-Image LSPR** table provides capacity relationships for processors running with multiple partitions that are assumed to exploit all the available real CPs. 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 (no zAAP or zIIP specialty engines are considered). The same workload is assumed to be running in every General Purpose or IFL partition. A 1-way processor can be considered as running with five partitions. Specific detail concerning the specific partition configurations represented for other models, is not provided. But, as the number or real CPs increases, the number of partitions increases, the average number of partition logical CPs increases, and the logical CP vs real CP ratio diminishes.

Note: The ***Multi-Image LSPR Processor Capacity Ratios*** table serves only 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 any specific LPAR host, since it considers the exact partition configuration, including the individual SCP/workload environments being run in each partition and any specialty CPs that are configured.

- The **Single-Image LSPR** table relates capacity for a single copy of the SCP (1 partition) managing all the real CPs. The table can be presented with General Purpose CPs only. For z/OS-1.13, capacity relationships are available up to the maximum of 100 CPs that are supported. When considering z/VM capacity, note that 32 CPs is the maximum supported.

- 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 any specific partition configuration on any 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 they are 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 logical CPs are always associated with a General Purpose z/OS-1.6 (or later) partition. IFL logical CPs may be associated with a General Purpose z/VM partition.

The LPAR host can be assigned as any IBM System z processor model. z/OS (various versions), z/VM, z/VSE, Linux, or CFCC may be defined as the operating system to any appropriate partition. Limitations concerning the z/OS version supported are listed below.

- **zEC12** and **zBC12**   z/OS-1.10 or later must be specified
- **z196** and **z114**   z/OS-1.7 or later must be specified
- **z10**   z/OS-1.7 or later must be specified
- **z9** and prior   Any supported z/OS version may be specified

For z/OS, any of the 5 LSPR workload categories may be used.

For z/VM and Linux, 5 workload categories similar to those for z/OS may be used.

z/OS, z/VM, and Linux capacity data is taken from the currently published LSPR tables, based on a single version of the SCP (z/OS-1.13). Various versions of z/OS may be specified for a partition. However, changing versions will have no effect on capacity results. Versions are used solely to enforce support rules, such as logical CP limitations or zAAP/zIIP usage.

For z/VSE, three predefined workloads are available, representing batch, online, and mixed environments.

For zAware (runs on zEC12 processors only), a single workload (zAware) is available. zAware capacity data is derived from the z/OS Average workload category.

For CFCC, a single workload (CFCC) is available. Capacity is based on CFCC measurements made at the time each processor family was announced.

Two modes of operation are provided for the **LPAR Configuration Capacity Planning** function:

- **Basic-mode**- allows a single LPAR configuration to be defined. Comparisons between LPAR configurations must be made manually, using output from separate invocations of the tool.

- **Advanced-Mode** 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.

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.

The **zPCR** tool's original design was in support of the z990 processor announcement. Features unique to the z990 include the ability to configure more than 16 real CPs on a CPC, and the ability to define up to 30 logical partitions. In addition, **zPCR** supports workload offload to zAAP logical CPs that are associated with a z/OS-1.6 (or later) partition, and workloads on IFL and ICF partitions on processors where such partitions are supported. On the z990 and z890 processors, zAAP, IFL, and ICF exist in a single common real CP pool. For capacity planning purposes, zAAP logical CPs associated with a z/OS partition are considered as separate partitions, though they do not count against the maximum number of partitions supported.

zPCR has subsequently been extended to support all newer IBM System z processor announcements. Features unique to the newer models include zIIP engines, and separate real CP pools for each CP type (i.e., General Purpose, zAAP, zIIP, IFL, and ICF). Also, weight and CAP assignments for zAAP and zIIP partitions are no longer tied to the associated General Purpose partition. zEC12, z196, z114, z10, and z9 processor models can support up to 60 logical partitions. For capacity planning purposes, zAAP and zIIP logical CPs associated with a z/OS partition are considered as separate partitions, though they do not count against the maximum number of partitions supported.

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.

Users should always assure that the most recently available version of **zPCR** is being used. Your download site can be checked by clicking **Check for updates** under **Help** on the *Function Selection* window menu-bar.

Getting Started

zPCR is a PC-based productivity tool that runs under Microsoft's Windows. It is one of a family of tools produced and maintained by your Capacity Planning Support (CPS) team, part of IBM's Advanced Technical Sales (ATS) in Gaithersburg, Maryland.

PC System (minimum requirement)

Processor:	Intel Core Duo or better
Memory:	1GB is required to install; more is recommended
Graphics:	1280×1024 or higher recommended
OS:	Microsoft Windows 7

Note: The currently supported environment for **PCR** is Windows 7. **zPCR** should function under Windows 8, XP, and Vista. However, reported problems will be addressed only if they can be recreated under Windows 7.

Obtaining zPCR

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

www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/PRS1381

Questions and Feedback

Contact Capacity Planning Support via …

- E-mail **zpcr@us.ibm.com**

Installation

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.

Prompts are provided during installation to specify the target directory for the application code and the program folder for the application icons. An application icon can optionally be placed on your desktop.

Note: Starting with **zPCR v8.5**, there is no longer a need for Windows *Administrator* authority to install the tool. **zPCR** must be run from the same user ID from which it was installed.

In the directory where **zPCR** was installed, the following subdirectories have been created and are available for your use:

- ✓ **Documentation**   Contains all of the documentation files that are provided with **zPCR**. The *zPCR User's Guide* (this document) is not included with the installation package; rather, it is available as a separate download. If the *zPCR User's Guide PDF* is copied into this directory, it will be directly available from **zPCR** under <u>**Documentation**</u> on the menu-bar.

- ✓ **Study Files**   Default folder for Study files. Includes a sample basic-mode study file and a sample advanced-mode study file (same study that was used for examples in this document).

- ✓ **EDF Files**   Default folder for EDFs. Includes a sample EDF

- ✓ **RMF Files**   Default folder for RMF files. Includes several sample RMF files

- ✓ **CSV Files**   Default folder for CSV files.

Java Requirement

zPCR requires the *IBM 32-bit Java runtime environment*. Special versions of the install material for **zPCR** are available to IBM Customers 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:

<p align="center"><u>**http://www.adobe.com**</u>.</p>

The related online help consists of the user's guide text, implemented as Java Help. This help is context sensitive to the window currently being viewed.

Notes:

The *zPCR User's Guide* PDF can be made directly accessible from **zPCR** (under <u>**Documentation**</u> on the menu bar of the primary windows) if copied to the **zPCR** *Documentation* sub-directory using the PDF filename as downloaded.

The *LSPR Document* PDF can be made directly accessible from **zPCR** (under <u>**Documentation**</u> on the menu bar of the primary windows) if copied to the **zPCR** *Documentation* sub-directory using the PDF filename as downloaded.

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 must be connected to the internet 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 completing and successfully submitting 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 (CPS Tools) 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 firewall will prevent direct internet registration. In this situation you should click the Register (e-mail) button. This will 1) attempt to initiate a properly addressed e-mail for you, and 2) invoke a dialog box with instructions to copy the encoded registration information into the e-mail note. Send the e-mail as addressed with the encoded registration information and wait for an e-mail response (generally within a few hours). Once received, start **zPCR** again, click the Register (e-mail) button, and continue with the dialog box instructions, indicating that you are to copy the contents of the returned note and click the Complete Registration button.

Registration is only required once, the first time that **zPCR** is installed and attempted to be used. Once registered, you have unlimited access to the tool. Occasionally, as major versions of **zPCR** become available, your registration will be renewed (automatically, if direct internet registration is allowed).

Capability to modify or remove your registration is provided under **Registration** on the *Function Selection* 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, tabs, pushbuttons, and radio buttons are provided on the *Function Selection* window to select between the various LSPR tables and the *LPAR Configuration Capacity Planning* function.

Usage Tips

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

Low Resolution Displays: **zPCR** is designed for a display resolution of 1280 × 1024 pels or higher. While a display resolution of 1024 × 768 pels may function, some windows may be less appealing to the eye. 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 un-checking "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 P, 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 a row.

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 E : 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 has been implemented to respect the various national language standards. All output fields are truly separated by commas. However, if your national language setting inserts additional commas, the CSV output cannot be properly read. In this case, your national language should be temporarily set to English (United States) before generating the CSV output. Alternatively, HTML output can be used without changing the national language setting.

- **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*.

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 100 unique partition definitions is allowed, regardless of any limitation on how many can be simultaneously active.)

When in ☑ **Advanced-Mode**, a study being saved will be flagged as an *Advanced-Mode* study, meaning that it includes 1 or more defined LPAR configurations. If not running in ☐ **Advanced-Mode**, the study will be flagged as a *Basic-Mod*e study which can have only a single LPAR configuration defined.

Basic-Mode studies can be saved from the *Function Selection* window, by using the **Save** toolbar icons or clicking on **File** on the menu-bar. *Advanced-Mode* studies can also be saved from the *Advanced-Mode Control Panel* window. Basic-mode 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 *Function Selection* window, click **File**, **Load** on the menu-bar
- Drag and drop a study file III icon onto the *Function Selection* window
- Drag and drop a study file onto the **zPCR** desktop III icon
- Double click on a **zPCR** study file III icon.

When an *Advanced-Mode* study file is loaded, ☑ **Advanced-Mode** will be checked, and it cannot be changed. When loading a basic-mode study file, ☐ **Advanced-Mode** will be unchecked. In this case, ☑ **Advanced-Mode** may be changed before entering the *LPAR Configuration Capacity Planning* function if desired.

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 7.x, a study file from **zPCR** version 6.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 *Function Selection* window or the *Advanced-Mode 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

zPCR can be customized to initialize with your own preset settings: From the menu-bar on the *Function Selection* window, click on **Edit ➔ Preferences** or Click the *Preferences* toolbar icon. Customization settings include the *Reference-CPU*, the list of workloads displayed in LSPR tables, the default drive:\directory\ for saving and loading **zPCR** studies, and the *Advanced-Mode* checkbox setting.

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.

Note that *Advanced-Mode* allows multiple LPAR configurations to be defined within a single **zPCR** invocation, thus eliminating the usual requirement for multiple **zPCR** invocations.

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 better 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* - an Adobe PDF document, 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 downloaded and copied to the **zPCR** *Documentation* folder, this document can be accessed from the *Function Selection* window or the *Advanced-Mode Control Panel* window by clicking **zPCR User Guide** under **Documentation** on the 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 by pressing PF1, by clicking on **Help** ➔ **Context Help** on the menu-bar.

- *zPCR QuickStart Guide* - an Adobe PDF document available via a button on the *Function Selection* window or the *Advanced-Mode Control Panel* window. Use this for guidance to defining a current and an alternate LPAR configuration, and make capacity comparisons between them.

- *LSPR Document* - current version of IBM's LSPR document in PDF format. If downloaded and copied to the **zPCR** *Documentation* folder, this document can be accessed from the *Function Selection* window or the *Advanced-Mode Control Panel* window by clicking **LSPR Document** under **Documentation** on the menu bar.

- *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. This information can be accessed by clicking on the **LSPR FAQ** button found on the *Function Selection* window.

- *LSPR Workloads and Capacity Planning Considerations* - information and considerations concerning the LSPR workloads and their use with **zPCR**. This information can be accessed from the *Function Selection* window or the *Advanced-Mode Control Panel* window; click **LSPR Workloads** under **Documentation** on the 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. *Choosing a Workload Category* - provides rationale concerning capacity planning workloads and how to select a category to represent a given production workload. This section can be accessed by clicking **Choosing a Workload Category** button on the *Workloads* window.

  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 offering insight on zAAP and zIIP capacity. A web site pointer to a white paper concerning zAAPs is also provided. Click the
  zAAP/zIIP Capacity Considerations button on the ***LPAR Host and Partition Configuration*** window.
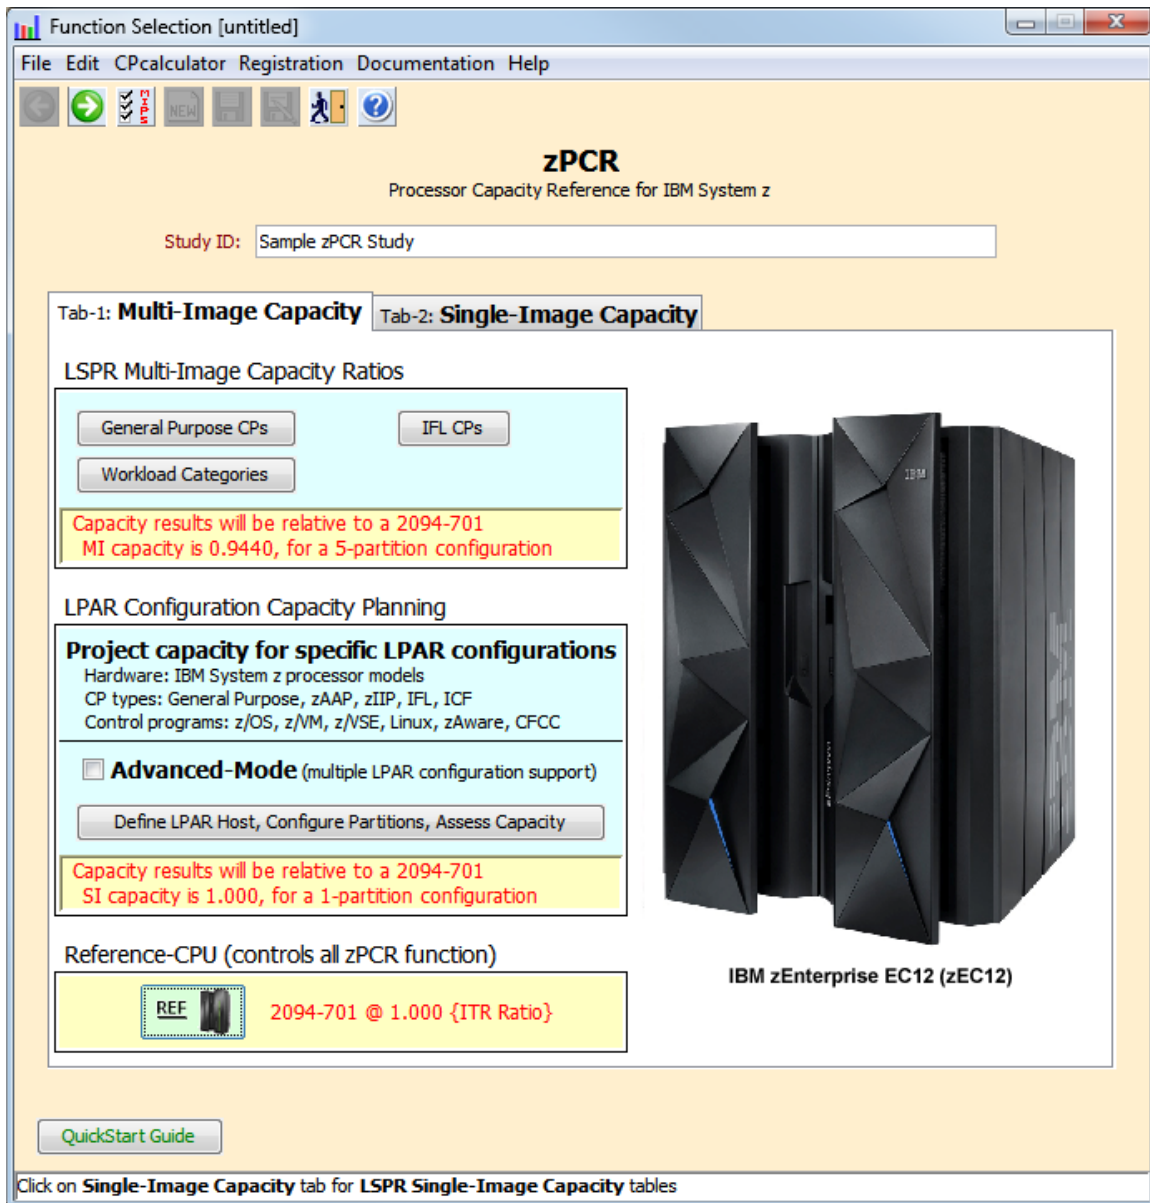
Logo Window
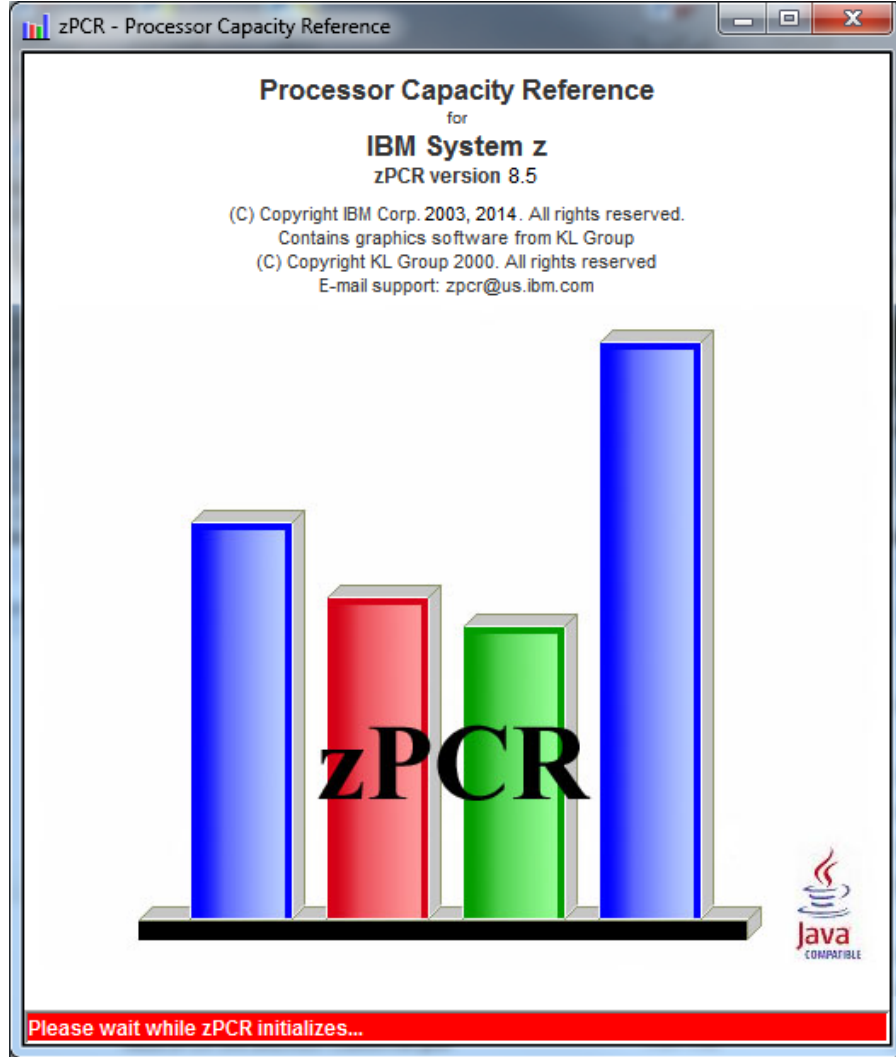
When **zPCR** execution is initiated, the *Logo* window is displayed. This window identifies the application, provides version information (also available under **Help ➔ About zPCR** on the menu-bar of the primary windows).

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 *Function Selection* window to appear. The "Please wait" message indicates a short delay while the Java code is dynamically compiled for execution.

Function Selection

Multi-Image Capacity Tab

The **Function Selection** window is displayed immediately following the **Logo** window (once the Java code is compiled). This is the primary window to access all the basic-mode function of the tool. In basic-mode, when you finish a **zPCR** function, control is returned to this window. From the **Function Selection** window a basic-mode study can be saved, retaining the **Reference-CPU** and the LPAR configuration information that has been entered. You will also terminate **zPCR** from this window.

There are two selectable tabs on the **Function Selection** window. Available activities include:

1. From the **Multi-Image Capacity** tab, you can:

   - View the **LSPR Multi-Image Capacity Ratios** table, based on current z/OS LSPR data. All IBM System z processor models are included. General Purpose CPs with up to 5 z/OS workloads can be viewed. Capacity ratios can also be considered to be representative of z/VM (maximum of 32 CPs) and Linux. IFL CPs can also be displayed, where capacity ratios can only represent z/VM or Linux. All IBM System z processor models through 101-way are included.

   - Perform **LPAR Configuration Capacity Planning** for a specific LPAR host processor running with a specific partition configuration. For z/OS, any recent version can be assigned to a General Purpose partition (the version is used only to enforce partitioning rules). z/VM and Linux can be assigned to General Purpose or IFL partitions. Any of the 5 workload categories can be assigned to z/OS, z/VM, and Linux partitions. z/VSE, with three workloads available, can be assigned to General Purpose partitions. CFCC can be assigned to General Purpose or ICF partitions. On zEC12 and zBC12 processors, zAware can be assigned to General Purpose or IFL partitions.

   - From the **Single-Image Capacity** tab, you can view General Purpose CPs in the **LSPR Single-Image Capacity Ratios** table, based on currently published z/OS LSPR data. All IBM System z processor models through 100-way are included.

User activity related to all **zPCR** function:
   1. Set the **Reference-CPU**, and its **scaling-factor** and **metric**

User activities related to the **LSPR Processor Capacity Ratios** tables (multi-image is the preferred table) include:
   2. Select SCP/workloads to be displayed in LSPR tables
   3. View **General Purpose CP** capacity (assumes z/OS, z/VM, or Linux)
   4. View **IFL CP** capacity (assumes z/VM and Linux only)
   5. Select specific processors for review, capture, and graphing

User activities related to the **LPAR Configuration Capacity Planning** function in **Basic-Mode**, include:
   2. From the **Function Selection** window, click the **Define LPAR Host, Configure Partitions, Assess Capacity** button.
   3. Define the current LPAR host processor with General Purpose CPs, zAAPs, zIIPs, IFLs, and ICFs.
   4. Define each General Purpose, IFL, and ICF partition. zAAP and zIIP logical CPs are always associated with z/OS General Purpose partitions. IFL logical CPs can be associated with a z/VM General Purpose partition.
   5. View the **Host Summary Report** for an overview of the LPAR configuration and its capacity results.
   6. View the **Partition Detail Report** to review individual partition capacity. Partition definition changes can also be made from this window.
   7. Capture table and graph output.
   8. Save basic-mode study inputs for future use.

9. Model a 2nd LPAR configuration to project its capacity. Then manually compare those results to the 1st (current configuration).

User activities related to the **LPAR Configuration Capacity Planning** function in **Advanced-Mode**, include:

2. From the **Function Selection** window, check **Advanced-Mode** and click the **Enter Advanced-Mode** button.

3. On the **Advanced-Mode Control Panel** window, double-click the **"Current" LPAR Configuration** icon.

4. Define the current LPAR host processor with General Purpose CPs, zAAPs, zIIPs, IFLs, and ICFs

5. Define each General Purpose, IFL, and ICF partition. zAAP and zIIP logical CPs are always associated with General Purpose partitions. IFL logical CPs can be associated with a z/VM General Purpose partition.

6. View the **Host Summary Report**, for an overview of the LPAR configuration and the capacity results.

7. View the **Partition Detail Report** to review individual partition capacity. Partition definition changes can also be made from this window.

8. On the **Advanced-Mode Control Panel** window, select the **Configuration #1** LPAR Configuration icon and click the **Clone** toolbar icon. Then double-click the **Configuration #2** LPAR configuration icon.

9. Modify the LPAR host processor to a different model or a new processor family.

10. Update each defined partition as needed.

11. From the **Advanced-Mode Control Panel** window, compare the **Configuration #2** LPAR configuration capacity to that of the **Configuration #1** LPAR configuration.

12. Save advanced-mode study inputs for future use

In **zPCR**, most System z 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.

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

To help to generalize capacity relationships representing processors running in basic-mode (where supported) or LPAR-mode, with a single partition exploiting the hardware, the **Single-Image Capacity** tab is provided. Single-Image capacity is a reasonable way to compare processors when the maximum number of CPs that could be configured is exactly what is to be supported by the SCP. Single-Image tables ignore any effect on capacity due to configurations that involve more than one partition.

LPAR-mode is required on all contemporary System z processors. To more fairly show capacity relationships for System z processors, the **Multi-Image Capacity** tab is provided. Both the **LSPR Multi-Image Capacity Ratios** table and the **LPAR**

Configuration Capacity Planning function can be used for comparing capacity projections of various configurations where multiple partitions are defined.

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 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.

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

Function Selection

Single-Image Capacity Tab
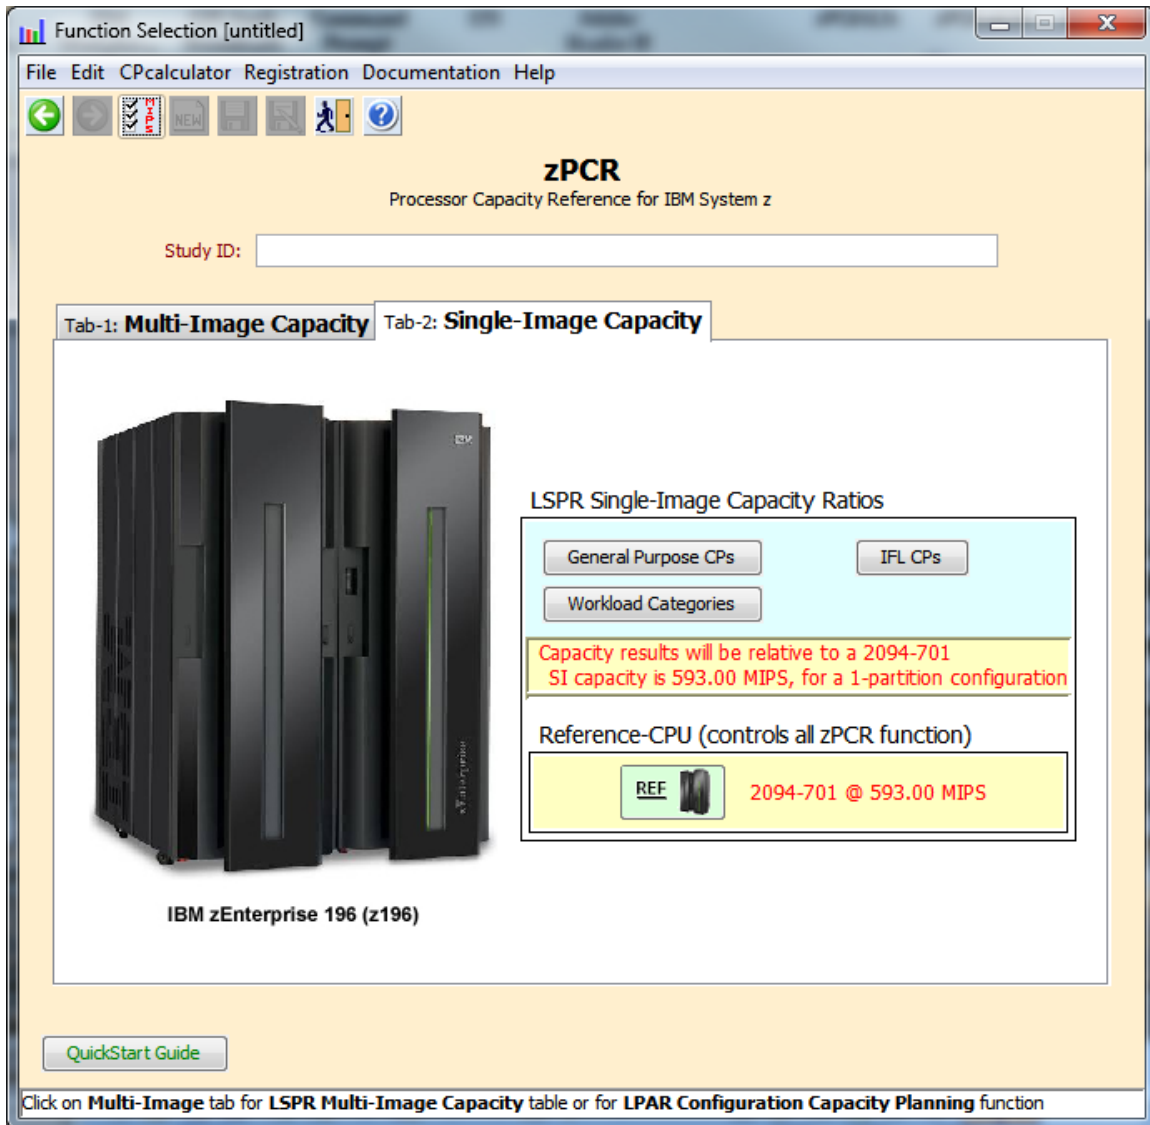
Click the ***LSPR Single-Image Capacity*** tab to reveal the 2nd page, providing access to the ***LSPR Single-Image Capacity Ratios*** tables.

The single-image table provides capacity relationships between processors running a single copy of the SCP, supporting all the available CPs. The single-image table can only be displayed from the ***Single-Image Capacity*** tab. To reveal Single-Image LSPR capacity ratios, click the **General Purpose CPs** button or the **IFL CPs** button.

z/OS: Only the General Purpose CP capacity ratios can be considered for z/OS. Ratios are shown for any of the 5 z/OS workload categories shown on the ***LSPR Workload Categories*** window). All possible IBM System z processor General Purpose configurations are available in this table, regardless of any N-way limitation that a specific version of z/OS might have.

z/VM and Linux: Either the IFL CP or the General Purpose CP capacity ratios may be considered for z/VM or Linux. The z/OS capacity relationships used can be assumed to be representative for z/VM (support limited to 32 CPs) and Linux. IFLs are always full speed engines.

Considerations for both the Multi-Image and Single-Image tabs

When viewing **General Purpose CPs** or **IFL CPs**, all capacity values are generated to be relative to the current **Reference-CPU** settings

Click the **Workload Categories** button to review the z/OS workload category names that are used for the LSPR Tables.

When initially loading **zPCR**, a new study is assumed; hence, the title bar will indicate an **[untitled]** study. When a study is saved, or a saved study is loaded, the name of the study will be indicated on the title bar.

Study Identification (**E** entry field)

The **Study Identification** entry field provides a means for identifying each unique study. Typically, this would be the name of a client or a site for which the study is being done. Character entry is accepted up to a count of 50. The name entered will appear on various tables and graphs that are produced. This name will be retained with a saved study.

Multi-Image Capacity (1st tab of 2)

This tab provides access to the **LSPR Multi-Image Capacity Ratios** tables. It also makes available the **LPAR Configuration Capacity Planning** function, providing the ability to configure partitions on System z processors. z/OS, z/VM, z/VSE, Linux, zAware, and CFCC workload environments are available to be defined to partitions.

1. **LSPR Multi-Image Capacity Ratios** tables

    Click the **General Purpose CPs** button to reveal LSPR capacity relationships for processors configured with General Purpose CPs. Each processor has an LPAR configuration considered representative for the N-way of the model. Every partition is assumed to be running z/OS with the same LSPR workload category. Up to 5 workload categories can be viewed in the table.

    Click the **IFL CPs** button to reveal LSPR capacity relationships for processors configured with IFL CPs. Each processor has an LPAR configuration considered representative for the N-way of the model. Every partition is assumed to be running z/VM or Linux with the same LSPR workload category. Up to 5 workload categories can be viewed in the table.

    Capacity values are relative to the current **Reference-CPU** settings (Click the **Reference-CPU** button). While a default is supplied, any 1-way General Purpose processor model contained in the table can be defined as the **Reference-CPU**. A scaling-factor and scaling-metric may also be set (see Reference-CPU).

    Use the **Workload Categories** button to select the specific workload categories for which capacity ratios are to be produced in the **LSPR Processor Capacity Ratios** table.

2. **LPAR Configuration Capacity Planning**

This function provides capacity results for a specific LPAR host processor defined with a specific partition configuration. Any IBM System z processor model is supported as an LPAR host. Partitions can be defined as running z/OS, z/VM, z/VSE, Linux, zAware, and CFCC workloads. Specific (recent) versions of z/OS can be specified for partitions where supported. When ☐ **Advanced-Mode** is not checked, click the ⌈ **Define LPAR Host, Configure Partitions, Assess Capacity** ⌋ button to start defining the partitioned environment in basic-mode. In basic-mode, the *LPAR Host and Partition Configuration* window will appear.

☑ **Advanced-Mode**: This checkbox provides the capability to define multiple LPAR configurations and make direct comparisons between them. When checked, the push button below becomes the ⌈ **Enter Advanced-Mode** ⌋ button. Click this button to open the *Advanced-Mode Control Panel* window, used to manage one or more LPAR configurations. Once in *Advanced-Mode*, the study cannot be reverted back to basic-mode. A discussion of *Advanced-Mode* features is provided towards the end of this document (see Advanced-Mode).

The basis for **zPCR** capacity projections is controlled by the *Reference-CPU* setting (Click the ⌈ **Reference-CPU** ⌋ button). Only 1-way General Purpose processors may be set as the *Reference-CPU*. A scaling-factor and scaling-metric for it may also be set (see Reference-CPU).

Partitions may be defined as running any one of the supported SCPs with one of the available workload categories. In some cases, such as z/VSE and CFCC, the workload choices are different.

Single-Image Capacity (2[nd] tab of 2)

Only the most recently published *LSPR Single-Image Capacity Ratios* table is available for viewing.

Reference-CPU

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

Click the ⌈ **Reference-CPU** ⌋ button to set the processor model and its metrics (see Reference-CPU).

Workloads

LSPR Capacity Ratios tables are displayed showing ITR ratios for multiple workload categories. Workloads that can be displayed include any of those included in the current LSPR tables. Click the ⌈ **Workload Categories** ⌋ button to review the available workload categories specify those to be included in the capacity tables (see Workloads).

Note that the *Workloads* window is common to the *LSPR Multi-Image Table*, the *LSPR Single-Image Table*, and the *LPAR Configuration Capacity Planning* function. While any of the SCP/workload categories shown can be used in the *LPAR Configuration Capacity Planning* function, only the published LSPR SCP/workload categories can be viewed in the LSPR tables.

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 ☑ **Advanced-Mode** is checked, save study capability is disabled on the *Function Selection* window. *Advanced-Mode* studies must be saved from the *Advanced-Mode Control Panel* window.

When not in *Advanced-Mode*, a basic-mode study will be saved, containing only the single LPAR configuration that was defined. In addition to the *Function Selection* window, basic-mode studies can also be saved from several of the *LPAR Configuration Capacity Planning* windows. When a basic-mode study is subsequently loaded, **zPCR** is placed in basic-mode and only the single LPAR configuration is loaded.
☑ **Advanced-Mode** can subsequently be checked to allow definition of additional LPAR configurations.

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:
1. *Reference-CPU Metrics*, processor model, scaling-factor and scaling-metric
2. *zPCR Operating Mode*, Advanced-Mode or Basic-Mode
3. *Workload category Display List* selections
4. *Directory Path for Saved Studies*, default folder for loading or saving study files
5. *Directory Path for RMF Reports*, default folder for loading RMF reports
6. *Directory Path for EDFs*, default folder for loading EDFs
7. *Directory Path for CSV files*, default folder for loading EDFs

To activate customization, click on **Edit ➔ Preferences** on the menu-bar; the *Preferences* window will open (see Preferences). Customization can only be invoked when a new study is started.

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** 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 zEC12, zBC12, z196, z114, and z10 processors.

✓ **LSPR Document** displays the *Large Systems Performance Reference* document (PDF format), published by IBM development and available via the internet at:

http://www.ibm.com/servers/eserver/zseries/lspr/lsprintr.html

The version included in **zPCR** is complete 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* and *Single-Image Capacity Ratios* tables, which always include the latest data. In addition, **zPCR** gives you the flexibility of setting the basis of the capacity values by setting the associated *Reference-CPU*s and their associated scaling-factor and 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.

Checking for zPCR Updates

You can access your CPS tools download web site directly by clicking on Check for updates under **Help** on the menu-bar. This function can help you to assure that you are using the latest version of **zPCR**.

Menu-bar

File

New	Start a new (untitled) study. A tool bar icon is also available.
Load	Open a previously saved study
Save	Save this study (must already be a titled study). A tool bar icon is also available.
Save as…	Save this study with a new name [1]. A tool bar icon is also available.
Up to 5 filenames	The last 5 study files referenced are listed for reference.
Exit (*Ctrl+E*)	Terminate **zPCR** execution. A tool bar icon is also available.
Fast Exit (*Ctrl+Q*)	Terminate **zPCR** execution immediately

Edit

Preferences	Personalize **zPCR** startup settings (see Preferences). Preferences can also be invoked using the tool bar icon.

CPcalculator	Capacity planning calculator functions (see CP Calculator)
zAAP Capacity	Capacity estimator for migration to a zAAP configuration (see zAAP Capacity Estimator Input)

Registration

Update	Change your CPS registration information
Remove	Delete your CPS registration (disables **zPCR**)
Documentation	Various supporting documentation, including:

> **zPCR NEWS file**
> **zPCR User's Guide** (if copied into **zPCR**'s *Documentation* folder)
> **LSPR FAQ**
> **LSPR Workloads**
> **LSPR Document** (if copied into **zPCR**'s *Documentation* folder)
> **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	Access your CPS tools download site
About zPCR	Displays information from the *Logo* window including the version of **zPCR**.

---

[1] In *Basic-Mode*, *Save as…* will create a basic-mode study.

Preferences
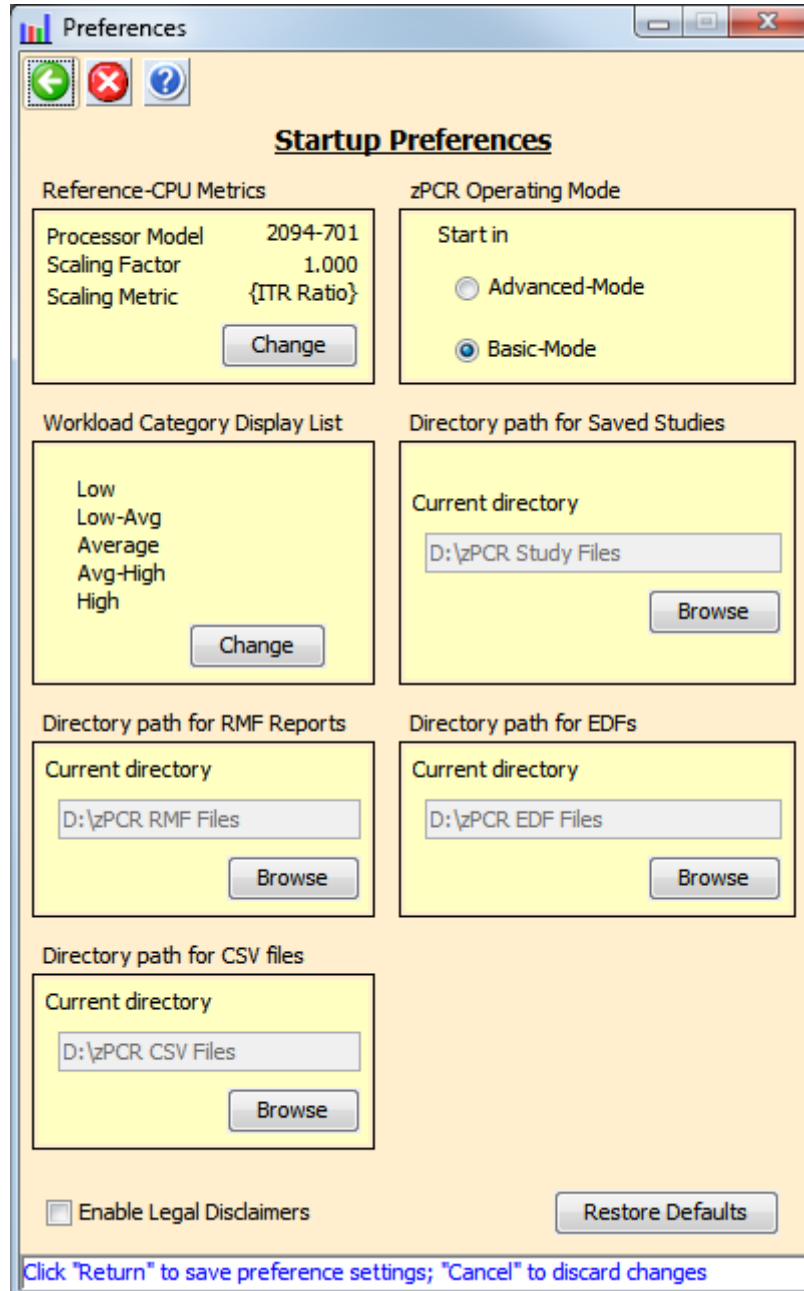
The **Preferences** window provides the capability to establish your own initialization conditions at **zPCR** startup. The following items can be set:

1. The **Reference-CPU** and its scaling-factor and metric can be set.

   To set the reference-CPU, click the **Change** button in the **Reference Processor** group box. The **Reference-CPU** window will be presented, from which to make your selection, setting the processor mode, and its scaling-factor and scaling-metric. These reference-CPU preferences can alternatively be made directly from the **Reference-CPU** window via a checkbox, when accessed (see Reference-CPU).

2. The workloads that will be displayed in the **LSPR Processor Capacity Ratios** tables can be set by clicking the `Change` button in the **Workload Category Display List** group box. The **Workloads** window will be presented, from which to make your selection. The workload list can alternatively be set directly from the **Workloads** window via a checkbox, when accessed (see Workloads).

3. In the **zPCR Operating Mode** group box, **Advanced-Mode** can be set as the startup default by clicking ⊙ **Advanced-Mode** (for details, see Advanced-Mode). The default setting is **Basic-Mode**.

4. The default location for loading and saving study files can be set by clicking `Browse` in the **Directory Path for Saved Studies** group box. A standard windows dialog is presented, from which to make your **drive:\directory** selection.

5. The default location for loading RMF reports can be set by clicking `Browse` in the **Directory Path for RMF Reports** group box. A standard windows dialog is presented, from which to make your **drive:\directory** selection.

6. The default location for loading EDFs can be set by clicking `Browse` in the **Directory Path for EDFs** group box. A standard windows dialog is presented, from which to make your **drive:\directory** selection.

7. The default location for saving CSV files can be set by clicking `Browse` in the **Directory Path for CSV files** group box. A standard windows dialog is presented, from which to make your **drive:\directory** selection.

The Preference settings for the **Reference-CPU** are used each time **zPCR** is started. Any of the customization settings can subsequently be changed during the execution instance of the tool. When changing a customized **Reference-CPU** setting, if you click `Default`, the values applied will be the **zPCR** defaults and not your customized values. You can click `Startup` to restore the customized value, if one has been established.

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

The **Preferences** window is only available before a study is started. Once any of these settings have been changed during the process of a study, access to this window is disabled. However, the **Reference-CPU** preference settings and the workloads list preference settings can still be altered from their respective windows.

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

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

Click the `Restore Defaults` button to return preference settings to their default values. In addition, any repeating dialogs that you have requested to stop will resume.

Reference-CPU

The ***Reference-CPU*** window is accessed primarily from the ***Function Selection*** window by clicking the ⬚**Reference-CPU**⬚ button. There are two of these buttons, one on each of the ***Function Selection*** window tabs. Both buttons control the same object.

The ***Reference-CPU*** can be set to any IBM System z processor 1-way model. Only models representing General Purpose CPs (as known to z/OS) can be selected. 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 set to the default value of **1.00** and the scaling-metric is set to the default value of **{ITRR}**.

Processor Model and Capacity Assumption group box

- **Family** Select the processor family from the drop down list (this field also relates the General Purpose CP speed setting where applicable)

- **Model** Select the specific processor model from the drop down list. The list is limited to 1-way models. (default is generally the full speed 1-way model of a System z family previous to that currently announced)

- **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-Saving** mode, the scaling-factor always represents the **Reference-CPU** at full power.

- **Scaling-Metric** Set the metric assumed for the capacity value, using the drop down list or by keying in up to 12 characters.
- (Default scaling-metric =. **{ITRR}**)

  Note: Whenever the **Reference-CPU** processor model is changed, the scaling-factor entry field will revert to "**1.00**" and the scaling-metric entry field will revert to "**{ITRR}**". These values must be changed explicitly by the user if other values are desired.

The **LSPR Multi-Image Capacity Ratios** table and the **LPAR Configuration Capacity Planning** function support all IBM System z processor 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 error 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 global **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 details, 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.

Provisional Reference-CPU

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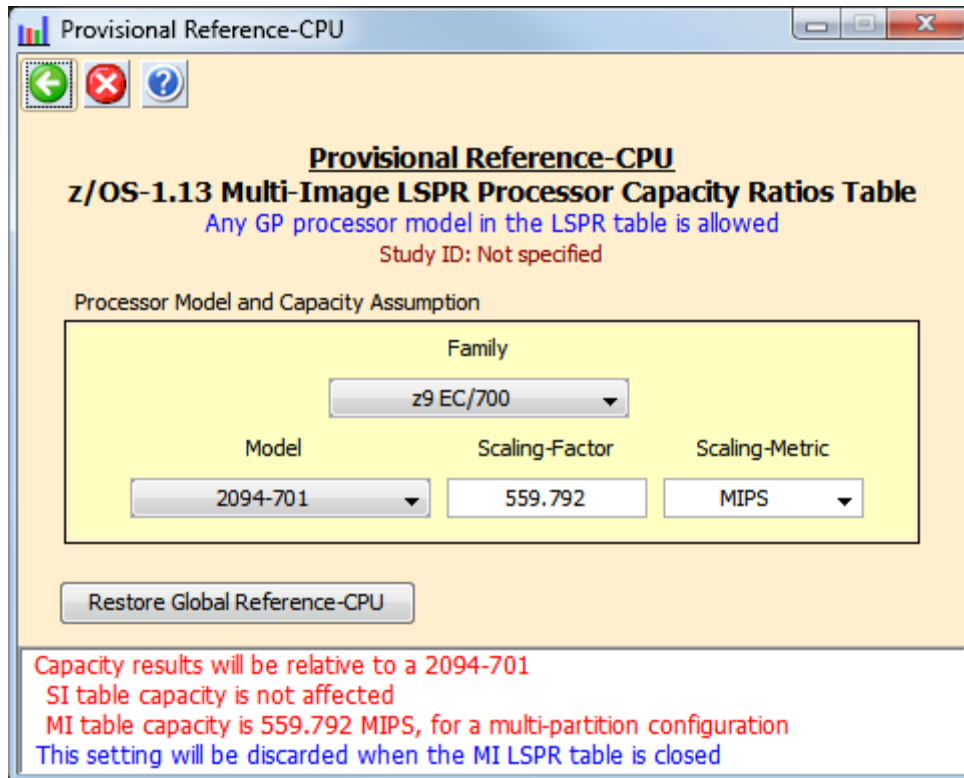
The *Provisional Reference-CPU* window is aware of which *LSPR Processor Capacity Ratios* table window is calling it; a title line will so indicate. Entry fields, similar to those on the *Reference-CPU* window, are described below.

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

Tip for quickly setting the *Provisional Reference-CPU*

From the *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 metric need be entered.

Processor Model and Capacity Assumption group box

- **Family**   Select the processor family from the drop down list (this field also relates the General Purpose CP speed setting where applicable)

- **Model**   Select the specific processor model from the drop down list. The list is limited to 1-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-Saving*** mode, the scaling-factor always represents the ***Reference-CPU*** at full power.

- **Scaling-Metric**   Set the metric assumed for the capacity value, using the drop down list or by keying it in (default =. **{ITRR}**)

Click the ⌷**Restore Global Reference-CPU**⌷ button to restore the global ***Reference-CPU*** values.

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.

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**. To help in understanding how the *Reference-CPU* settings are interpreted by the various **zPCR** functions, the following table is provided.

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) *Scaling-factor* (capacity of a shared 1-partition configuration) *Scaling-metric*	**2094-701** **1.00** **{ITRR}**	**2094-701** **593** **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 logical CP 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

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.

Workloads
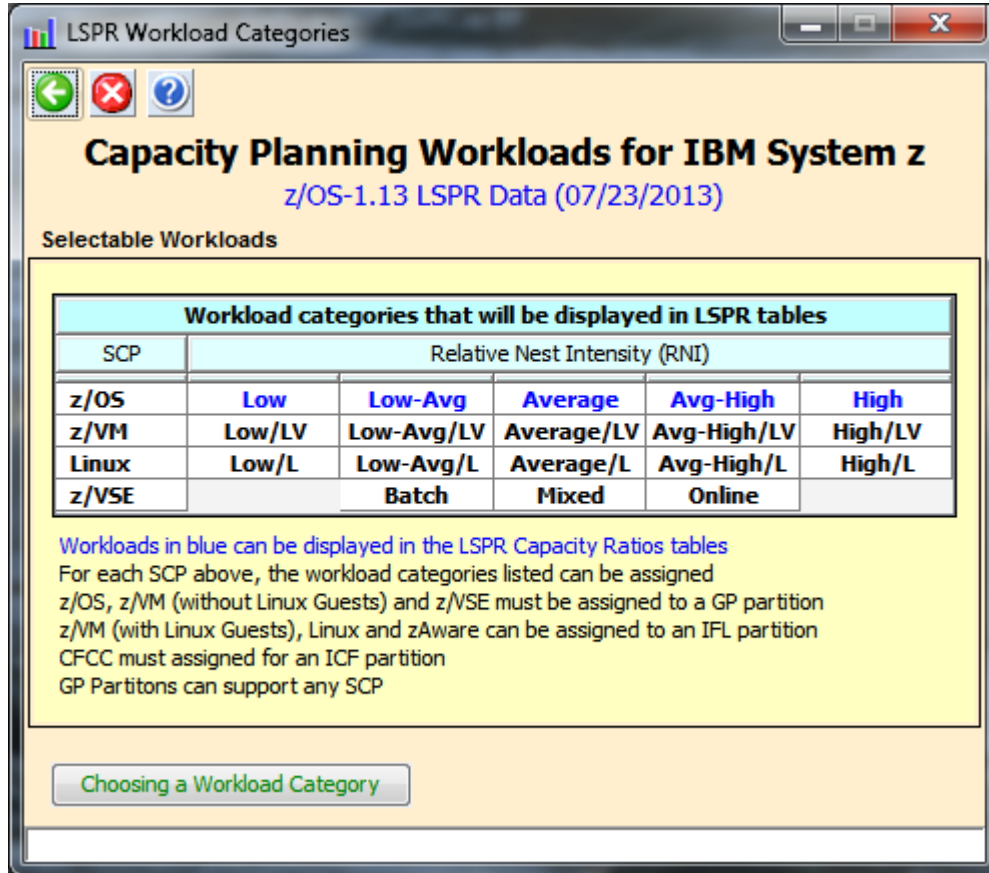
The **Workloads** window is accessed from the **Function Selection** window by clicking the **Workload Categories** button from either the **Multi-Image Capacity** tab or the **Single-Image Capacity** tab. It can also be accessed directly from either of the **LSPR Processor Capacity Ratios** tables via 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 System 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 and Linux. When General Purposed CPs are being displayed, the capacity values are considered applicable for z/OS, z/VM, or Linux. When IFL CPs are being displayed, the capacity values are considered applicable for z/VM or Linux (z/OS 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* will be the default assignment.
- **z/VM** can be defined to IFL or GP partitions; *Average/LV* will be the default assignment.
- **Linux** can be defined to IFL or GP partitions; *Average/L* will be the default assignment.
- **z/VSE** can only be defined to GP partitions; *Mixed* will be the default assignment.
- *zAware* can be assigned to IFL or GP partitions; *zAware* is also the workload assignment.
- *CFCC* can be defined to ICF or GP partitions; *CFCC* is also the workload assignment.

The latest versions of z/OS can support up to 100 logical CPs. z/VM can support a maximum of 32 logical CPs. Linux can support any number of logical CPs.

LSPR Workloads

In the past, the actual LSPR workload primitives were published. Such data could not be used reliably to represent capacity for a production workload. This is because some of the z/OS workload primitives tend to stress hardware more than would be typical of production work. As a result, the primitives (and suggested mixes of them) are no longer supported in **zPCR**. The z/OS workload primitives are useful to understand the total capacity envelope (or bounds) expected for each processor model

The 5 z/OS LSPR workload categories now carried in **zPCR** provide a more reliable way of characterizing production workloads. These categories allow more precise positioning in the potential capacity spectrum, and tend to smooth out any workload primitive anomalies that may exist.

The five z/OS LSPR workload categories are listed in the table below. 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 design will affect the absolute *Nest Intensity* of any given workload.

Name	Characterization
Low	Low *Relative Nest Intensity*
Low-Avg	Mid-way between **Low** and **Average**
Average	Average *Relative Nest Intensity*
Avg-High	Mid-way between **Average** and **High**
High	High *Relative Nest Intensity*

The 5 workload categories are expected to satisfy most z/OS capacity planning requirements. All of these categories are available in **zPCR**, in the *LSPR Multi-Image*, the *LSPR Single-Image* table, and for the *LPAR Configuration Capacity Planning* function. Similarly named workload categories are also available for z/VM and Linux when assigned to a partition.

While the 5 workload categories are derived from z/OS measurement data, they are also applicable for z/VM and Linux production workloads.

Controls

Click the **Return** toolbar icon to return to the calling window (usually the *Function Selection* window).

Click the **Choosing a Workload Category** 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** (General Purpose, and associated zAAP or zIIP partitions)

  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 zEC12, zBC12, z196, z114, and z10 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 recognize. 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*.

  When *CPU-MF* data is not available, the default LSPR workload category "Average" will be assigned.

  The workload used to represent a z/OS partition can be manually designated as any of the 5 available z/OS workload categories.

- **z/VM** (General Purpose or IFL partitions)

  CPU Measurement Facility (**CPU-MF**) data provides the most reliable method for assigning a workload category to a z/VM partition. **CPU-MF** data can be obtained from z/VM systems running on zEC12, zBC12, z196, z114, and z10 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 recognize. 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**..

  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.

  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 available z/VM workload categories.

- **Linux** (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 available Linux workload categories.

- **z/VSE** (General Purpose partitions only)

  The default workload category assignment for z/VSE is **Mixed**.

  The workload used to represent a z/VSE partition can be manually designated as any of the 3 available z/VSE workload categories.

- **zAware** (General Purpose or IFL partitions; zEC12 and zBC12 processors only)

  The only workload category assignment for CFCC is CFCC.

- **CFCC** (General Purpose or ICF partitions)

  The only workload category assignment for CFCC is CFCC.

LSPR Processor Capacity Ratios
Capacity by Workload Type

Multi-Image Table

The **LSPR Multi-Image Processor Capacity Ratios** window is displayed from the **Function Selection** window by clicking on the General Purpose CPs or the IFL CPs button in the **LSPR Multi-Image Capacity Ratios** group box on the **Multi-Image Capacity** tab.

LSPR Processor Capacity Ratios
Capacity by Workload Type

Single-Image Table

The **LSPR Single-Image Capacity Ratios** window is displayed from the **Function Selection** window (**Single-Image Capacity** tab) by clicking on the General Purpose CPs button in the **LSPR Single-Image LSPR Processor Capacity Ratios** group box.

These windows 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 **Workloads** window by completing the **Workload Display List**.

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 is 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 apply only to General Purpose CPs.

Capacity Rating by Workload

- **SCP/Workload 1**   Relative capacity for this workload
- **SCP/Workload 2**   Relative capacity for this workload
- etc.

Note: **For zEC12, zBC12, z196, z114, and z10 processors, z/OS capacity is represented with HiperDispatch turned ON**. A white paper is available concerning **HiperDispatch** at:

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

LSPR Multi-Image Capacity Ratios Table

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 logical CPs 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** 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 and Linux 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 may vary from that of the "Typical" LPAR configuration:
   - Number of partitions defined
   - Number of logical CPs defined to each partition
   - Relative weights assigned to each partition
   - Capping of partitions
   - The workload category assigned to each partition may vary
   - SCPs other than z/OS (z/VM, z/VSE, Linux, CFCC) may be assigned to a partition.

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. In addition, capacity projections for the zAAP CPs will be 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. In addition, capacity projections for the zIIP CPs will be provided.

4. When IFLs are configured to run z/VM or Linux, capacity results for other partitions will be impacted slightly. In addition, capacity projections for the IFLs will be provided.

5. When ICFs are configured to run CFCC, the capacity results for other partitions will be impacted slightly. In addition, capacity projections for the ICFs will be provided.

6. When books on zEC12, z196, z10-EC, z9-EC, and z990 processors are configured in excess of that needed for the real CPs being used, capacity results for all partitions will be impacted slightly.

7. For zEC12, zBC12, z196, z114, and z10 processors running z/OS, the ***HiperDispatch*** effect of keeping workload content on the same real CPs is inherent in the LSPR capacity data. However, any potential parking of logical CPs is ignored. If parked logical CPs 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 real CP type (i.e., CP pool), and for the LPAR host as a whole.

The **LSPR Multi-Image Capacity Ratios** table plays no roll 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

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 real CP resource. This table is a reflection of the z/OS workload environments that are considered measured for LSPR purposes.

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-1.13, which supports a maximum of 100 CPs. If considering capacity for previous z/OS versions, one must ignore capacity ratios for N-way processors that are not supported. Versions prior to z/OS-1.6 can support a maximum of 16 CPs. z/OS-1.6 through z/OS-1.8 can support a maximum of 32 CPs. z/OS-1.9 can support a maximum of 64 CPs. z/OS-1.10 and later can support up to 100 CPs.

While the single-image capacity ratios are based on z/OS measurement data, they can also be considered reasonable for z/VM and Linux. z/VM supports a maximum of 32 CPs; therefore, do not consider the z/OS single-image capacity ratios for models with more than 32 CPs. Linux can support any number of CPs.

For z900 and z800 processors only, one may assume that the capacity relationships represent either basic-mode or LPAR-mode. When a zEC12, zBC12, z196, z114, z10, z9, z990, or z890 processor is 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 (both tables) 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 or 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.

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.

By default the **Reference-CPU** is selected when the table is initially displayed. To select additional 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.

Table View Group Box

1. **Families & Models**

   The number of processor families that are displayed in the table can be controlled by clicking the **Processor Families** button (see Processor Families for a discussion of this capability). The ⊙**Subset** radio button will cause only the processor families that have been chosen to be included in the table. Setting the ⊙**All** radio button causes all processor families to be included in the table. By default, older processor families are excluded from the table and the setting is ⊙**Subset**.

   Note: The **Processor Families** feature offers the capability to limit the table to only certain specific families. For example, when specifying only the z114, z10-BC or z9-BC models, the table can then be sorted to assist in finding a model with a desired capacity or MSU rating while excluding all the other families.

   Individual processor models that have been selected can be viewed independently by clicking the ⊙**Selected** radio button. Restore the processor list by clicking the ⊙**Subset** or the ⊙**All** radio button.

2. **Capacity** (Single-Image table only)

   With the **LSPR Single-Image Capacity Ratios** table, the capacity values displayed can be changed to represent the capacity of a single engine by clicking the ⊙**1-CP** radio button. Information in the title area will reflect the status of the capacity data. To help emphasize that the table is being displayed in this state, capacity numbers are changed from red to brown. To restore the table to full CPC capacity, click the ⊙**All CPs** radio button.

3. **zEC12 & z196 Power**

    While zEC12 and z196 processors will normally be run at Full Power, they can be operated in *Power-Saving* mode, reducing capacity, reducing the MSU rating, and reducing energy requirements. The zEC12 and z196 processor models can be displayed in either mode using the radio buttons provided. Click ⊙*Saving* to display the *Power-Saving* mode MSU ratings and capacity values for all zEC12 and z196 processor models. A green background color will be displayed. Click ⊙*Full* to restore the zEC12 and z196 MSU ratings and capacity values back to Full Power. Note that the *Reference-CPU* scaling-factor is always assumed as being Full Power.

    The zEC12 and z196 processor families will default to displaying at Full Power when:
    - The LSPR table is opened.
    - The *Provisional Reference-CPU* is set or changed.

Controlling the Processor Sequence

The processor table may be sorted on the *Processor* column, the *MSU* column, or on any *SCP/workload* column by clicking its 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, and are not displayed.

In the *Table View* group box under *Families & Models*, with ⊙*Selected* set, sorting is disabled. By sorting the table with ⊙*Subset* or ⊙*ALL* set, and selecting processors of interest, they can then be displayed with ⊙*Selected* by first restoring the normal sort order.

In the *Table View* group box under *Capacity*, the processor table may be sorted with either ⊙*All CPs* or ⊙*1-CP* radio button set. However, switching the setting will change a sorted table back to its normal order, since the capacity numbers would no longer match the sort request.

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

Controls

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 window. Unlike the *Reference-CPU*, the *Provisional Reference-CPU* can be set to any processor N-way model in the table. When the window is closed, the global *Reference-CPU* settings are restored.

> **Note 1**: Alternatively, when viewing General Purpose Engine 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 the *Reference-CPU* processor model.

> **Note 2**: Setting the *Provisional Reference-CPU* will cause zEC12 and z196 processor models, if being displayed in *Power-Saving* mode, to revert to *Full Power*. If *Power-Saving* mode is desired, it can be restored by clicking ⊙ *Saving* in the *zEC12 & z196 Power* group box.

Click the ⎹ **Processor Families** ⎸ button to review and control the processor families that appear in the window.

Click the ⎹ **Workload Categories** ⎸ button to go to the *Workloads* window to alter which workloads are displayed or the order in which they are displayed.

Click the **Return** toolbar icon to return to the *Function Selection* window.

Click the **PRN** toolbar icon to create flat PRN file output

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

Menu-bar

Workload

Describe	Show workload description for LSPR workload categories.
Graph	Create charts for processors that have been selected (General Purpose CPs only; not available for IFLs).
Workload ITRR	Generate capacity bars showing any specific workload.
All categories	Generate capacity bars showing capacity for all of the published LSPR workload categories side-by-side.
Response	Generate line graph showing processor internal response time.

Help

Processor Families

LSPR Processor Capacity Ratios
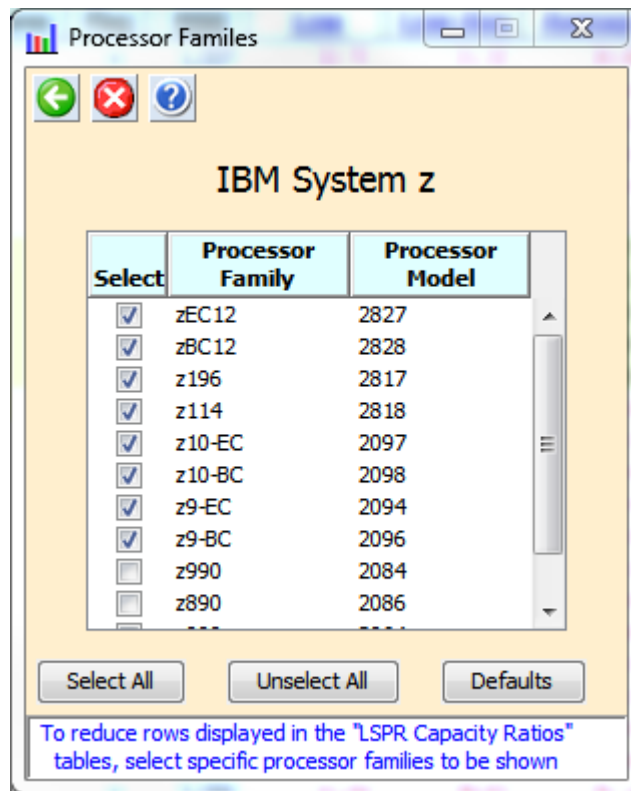
The ***Processor Families*** window is accessed by clicking the Processor Families button on either the ***LSPR Multi-Image*** or ***Single-Image Capacity Ratios*** table windows. This window provides a list of all the processor families included in the LSPR table. It also provides a means to limit the ***LSPR Processor Capacity Ratios*** tables to include only those processor families of interest.

Each processor family carried in **zPCR** is listed on the left following an ☑ ***Include*** checkbox. Use the ☑ ***Include*** checkboxes to indicate which families are to be displayed when viewing any given ***LSPR Processor Capacity Ratios*** table (single-image and multi-image). Processor families that remain checked are displayed in the ***LSPR Processor Capacity Ratios*** table when the ⊙***Subset*** radio button under ***Families*** is set. Note that several of the older System z processor families default to being excluded.

Click the **Return** toolbar icon invoke the family settings and return to the ***LSPR Processor Capacity Ratios*** table. If any processor families are unchecked, the ⊙***Subset*** radio button under ***Families*** is set, and only the selected families appear. The ⊙***Complete*** radio button restores the table to reveal all System z processor families.

Click the **Cancel** toolbar icon to exit without saving any checkbox settings that were made.

To assist in making changes, Select All and Unselect All buttons are provided. At least one processor family must be selected in order to return from this window.

Processor Family settings are retained only for the current **zPCR** invocation.

LPAR Configuration Capacity Planning
<u>LPAR Host and Partition Configuration</u>

The **LPAR Configuration Capacity Planning** function provides capacity projections for any System z 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. An entry field is available at the top with which to provide a verbal description for the configuration.

zPCR may be run in *Basic-Mode* or *Advanced-Mode*. How the *LPAR Host and Partition Configuration* window is accessed depends on the mode used.

- *Basic-Mode*: From the *Function Selection* window, click the
  Define LPAR Host, Configure Partitions, Assess Capacity button.

- *Advanced-Mode*: Double-click on any defined LPAR icon under the *LPAR configuration* tree.

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 logical CPs are always defined in conjunction with a General Purpose partition running z/OS. IFL logical CPs may be defined to a General Purpose 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 System z processor families are supported by **zPCR** as LPAR hosts. The SCP assigned to a partition will dictate the workload that can be assigned:

- **z/OS**   5 z/OS workload categories
- **z/VM**   5 z/VM workload categories
- **Linux**   5 Linux workload categories
- **z/VSE**   3 z/VSE workloads (not shown in LSPR tables)
- **zAware**   1 zAware workload (not shown in LSPR tables)
- **CFCC**   1 CFCC workload (not shown in LSPR tables)

Note: For z/OS partitions on zEC12 and zBC12 processors, z/OS-1.10 or later must be specified; and on z196, z114, and z10 processors, z/OS-1.7 or later must be specified. For all other processor families, z/OS can be defined as any version from 1.4 to 1.13. The SCP version information is used only to enforce configuration rules (i.e., the number of logical CPs that can be defined, or whether zAAP or zIIP logical CPs are supported). Capacity results will not be affected by changing the version.

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

z/VSE data is not carried in any LSPR table. However, with the **LPAR Configuration Capacity Planning** function, z/VSE may be defined to a General Purpose partition with one of three predefined workload environments, **Batch**, **Mixed**, and **Online**. z/VSE capacity for 1-way processors has been derived using similar z/OS LSPR workload categories. To support N-way processors the VSE Turbo Dispatcher is required; z/VSE N-way processor capacity is carried as a simple approximation of that realizable with the Turbo Dispatcher. A maximum of four logical CPs is supported for z/VSE 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 (in **Basic-Mode**) the **Function Selection** window or (in **Advanced-Mode**) the **Advanced-Mode Control Panel** window.

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 logical CPs are defined in conjunction with its associated General Purpose 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**). Additionally, IFL logical CPs may be defined in conjunction with a z/VM general purpose partition.

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

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

Clicking any of these buttons will open the **LPAR 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 **RMF Utilized Capacity vs Weighted Capacity** window (see Utilized Capacity Report). This button is enabled only if the LPAR host and its entire partition configuration (where partition utilizations are available) were created from RMF or EDF.

Note that there is no output capability for the **LPAR Host and Partition Configuratio**n 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 configuration or copy individual partition definitions into a currently active **zPCR** study. A single EDF interval must be chosen, from which the requested information will be drawn.

Only EDF from an IBM System z processor is supported. **zPCR** only supports EDF reports generated under reasonably current z/OS or z/VM versions. If the EDF was generated from the data of 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 explaining which versions are supported.

EDFs created from z/OS

When the LPAR host is a zEC12, zBC12, z196, z114, or z10, CPU-MF hardware counter data can be captured via SMF type 113 records. This information, transferred to **zPCR** in the form of an EDF, is used to determine a **Workload Choice for CPU-MF**. When this information is available, it should be considered as the most reliable method for workload category assignment to a partition.

Each EDF represents a single z/OS partition. Therefore, separate EDFs may be needed to represent all the partitions of interest. When multiple EDFs are being referenced, each should represent the same time period.

On systems where HiperDispatch is active (z10 and later), some LCPs may be parked for General Purpose partitions and any associated zAAP or zIIP logical CPs. 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 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 logical CPs 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.

Partitions that are defined, but inactive (i.e., no SCP is IPL'ed) are unchecked in the **Active** column. Normally, such partitions should not be included, since no LPAR overhead is generated on behalf of them.

EDFs created from z/VM

EDFs can also be captured from z/VM systems. Each EDF represents a single z/VM partition and all of its guests. If there are multiple z/VM partitions, separate EDFs may be needed to represent all the partitions of interest. When multiple EDFs are being referenced, each should represent the same time period.

When the LPAR host is a zEC12, zBC12, z196, z114, or z10, CPU-MF hardware counter data can be captured via VM Monitor. If z/VM is running in a General Purpose partition, while also supporting IFL logical CPs, CPU-MF data is captured for each type. This information, transferred to **zPCR** in the form of an EDF, is used to determine a **Workload Choice for CPU-MF**. When this information is available, it should be considered as the most reliable method for workload category assignment to a partition.

EDFs must be downloaded to your PC (preferred extension = **edf**). Only one EDF is necessary to obtain the LPAR host, partition definition, and partition utilization information (referred to as the primary EDF). Additional EDFs are useful when **CPU-MF** information is available for partitions (one EDF required for each such partition).

For 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:

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 entire LPAR host and partition configuration model that represents a current operating environment.

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 a current study.

As an alternative to clicking 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).

Creating LPAR Configuration from RMF

A z/OS **RMF Report** can be used to create the entire LPAR configuration or to copy individual partition definitions into a currently active **zPCR** study. A single RMF interval must be chosen, from which the requested information will be drawn.

Only RMF data from an IBM System z processor is supported. **zPCR** only supports RMF reports generated under reasonably current z/OS versions. If the RMF was generated under an unsupported z/OS version, a dialog will be presented explaining which versions of z/OS are supported.

Two types of RMF report data can be processed.

1. A traditional RMF report in text format. The RMF report must be available as a flat file, downloaded to your PC (preferred extension = **txt**). The file may either contain or exclude carriage-control characters in the 1st byte position of each record.

2. An RMF report in the newer XML format. The RMF report must be downloaded to your PC (preferred extension = **xml**).

The RMF report must include 1) the **Partition Data Report** (**PDR**), where the actual partition configuration is determined, and 2) at least 1 **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.

On systems where HiperDispatch is active (z10 and later), some LCPs may be parked for General Purpose partitions and any associated zAAP or zIIP logical CPs. 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 definition to more accurately relate capacity. It is up to the user to manually remove parked LCPs.

On systems where IRD is actively managing the partition resources, partition logical CPs 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.

Partitions that are defined, but inactive (i.e., no SCP is IPL'ed) are unchecked in the **Active** column. Normally, such partitions should not be included in a study, since no LPAR overhead is generated on behalf of them.

From **zPCR**'s **Host and Partition Configuration** window, there are two ways that configuration information can be obtained from RMF:

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.

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 buttons, an RMF file can be dragged and dropped onto either of the $\boxed{\textbf{RMF}}$ buttons when the button it is active.

Use of either button will display the **RMF Interval Selection** window (see <u>RMF Interval Selection</u>) from which an RMF report can be loaded. Once an RMF interval is selected, click the $\boxed{\textbf{Show Partitions}}$ button to display the **Create LPAR Configuration from RMF** or **Get Partitions from RMF** window (see <u>Get Partitions from RMF</u>).

Copying Partitions from a Previous Study

Partitions defined in a previous **zPCR** study can be copied into the currently active study. In the **Define Partitions** group box, under **Copy Partitions From**, click the $\boxed{\textbf{zPCR Study}}$ button (see <u>Copy Partitions from Study</u>).

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 System z processor model may be defined as the LPAR host.

Choose the intended processor family using the **Family** dropdown list. Then choose the specific processor model using the **Model** dropdown list. Finally, set the number of CPs for each engine type that are to be configured using the dropdown list 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.

Some processor families have a limit on the number of General Purpose CPs that can be configured that is less than the total CPs available. In these cases, the GP dropdown list will be limited to that count.

For some processor families, the model name indicates the actual number of General Purpose CPs configured. In these cases the GP count will be set and the dropdown will be locked.

Some CP types are not supported on certain families or models. In these cases, the CP count will be set to zero and the dropdown will be locked.

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 real CPs configured
- Number of zAAP real CPs 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 real CPs configured (there are rules about how many CPs relative to General Purpose CPs can be defined while remaining a standard zAAP configuration).
- Number of IFL real CPs configured
- Number of ICF real CPs configured

z/OS can be defined to General Purpose partitions on any System z processor model. Each version has its own limit on the number of CPs that an image can support. The maximum for z/OS-1.12 and later is 100 CPs.

For z/OS partitions on zEC12 and zBC12 processors, z/OS-1.10 or later must be specified; and on z196, z114, and z10 processors, z/OS-1.7 or later must be specified for General Purpose partitions (previous z/OS versions are not supported). If changing an older LPAR host to a zEC12, zBC12, z196, z114, or z10 processor model, partitions defined with older versions of z/OS must 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.

On z9 processors and prior, any z/OS version (1.6 and later) can be defined to a General Purpose partition. Where supported, zAAP LCPs and zIIP LCPs can also be configured to the LPAR host processor.

z/VM can be defined to General Purpose or IFL partitions on any System z processor model. A z/VM image is limited to supporting 32 CPs. When z/VM is run in an IFL partition, Linux guests must be assumed.

Linux can be defined to General Purpose or IFL partitions on any System z processor model. A Linux image can support any number of CPs.

z/VSE can only be assigned to General Purpose partitions. A maximum of four logical CPs is supported by **zPCR**.

zAware can be defined to General Purpose or IFL partitions on zEC12 and zBC12 processors only. For a single zAware image, a limit of 16 CPs is currently enforced by **zPCR**.

CFCC can be defined to General Purpose or ICF partitions on any System z processor model. A CFCC image is limited to supporting 16 CPs on all processor models.

<u>Controls</u>

Click the **<u>Return</u>** 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 CPs will be restored to be the maximum standard configuration.

Click the **<u>Cancel</u>** 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-Saving Mode (zEC12 and z196)

While normally run at full power, zEC12 and z196 processors can be run in **Power-Saving** mode. **Power-Saving** mode reduces the processor MSU rating and capacity (all CP types are affected equally), while also reducing power requirements. When the LPAR host is defined as a zEC12 or z196, the **LPAR Host** window provides a means to set it so that capacity results will represent its operation in **Power-Saving** mode.

zEC12 and z196 processors defined as the LPAR host will initially default to full power. To change to **Power-Saving** mode, click the **Saving** radio button in the **zEC12 & z196 Power** group box. If the LPAR host is configured for **Power-Saving** mode, all capacity reporting windows will indicate so in the title area.

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.

- zEC12 and zBC12 processors, zAAP or zIIP CPs cannot exceed 2 for each GP CP.
- All other System z processors, zAAP or zIIP CPs cannot exceed 1 for each GP CP.

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 ☑ check box 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 zEC12, zBC12, z196, z114, z10, and z9 processors, GP partitions can be configured to run zAAP eligible work on zIIP engines (**zAAP on zIIP**).

When **zAAP on zIIP** is enabled for a GP partition, the following rules apply:
- zAAP logical CPs 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 GP partition, then only zIIP logical CPs 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. zAAPs are not supported on z900 and z800 processors. zIIPs are not supported on z900, z800, z990 and z890 processors. IFLs and ICFs are supported on all System z families.

The zEC12, z196, z10-EC, z9-EC, and z990 processors, which are book-based models, have a method of processor model identification where the model number indicates the number of books and the maximum number of General Purpose CPs. The actual number of General Purpose CPs configured is designated independently.

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 capacity indicators (or speeds), 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. The 400 capacity indicator is the only zEC12 host that can be configured with no General Purpose CPs. Any available CP can be configured as a General Purpose CP (within the limits for the 400, 500, and 600 described above), as a zAAP (limited to 2X the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

IBM zEnterprise BC12 (zBC12)	
When configuring, there is a choice of 2 models	
2828-H06	Up to 6 CPs may be configured
2828-H13	Up to 16 CPs may be configured

There are 26 General Purpose CP capacity indicators (or speeds), *A*, *B*, *C*, *D*…*Z* (*Z* 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. The 2828-A00 is the only model 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 (limited to 2X the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

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 capacity indicators (or speeds), 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. The 700 capacity indicator 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 (within the limits for the 400, 500, and 600 described above), as a zAAP (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

IBM zEnterprise 114 (z114)	
When configuring, there is a choice of 2 models	
2818-M05	Up to 5 CPs may be configured
2818-M10	Up to 10 CPs may be configured

There are 26 General Purpose CP capacity indicators (or speeds), *A*, *B*, *C*, *D*…*Z* (*Z* 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. The 2818-A00 is the only model 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 (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

IBM 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 capacity indicators (or speeds), 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. The 700 capacity indicator is the only z10-EC host that can be configured with no General Purpose CPs. Any available CP can be configured as a General Purpose CP (within the limits for the 400, 500, and 600 described above), as a zAAP (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.	

IBM System z10-BC When configuring, there is a single model	
2098-E10	Up to 10 CPs may be configured
There are 26 General Purpose CP capacity indicators (or speeds), *A*, *B*, *C*, *D*…*Z* (*Z* 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. The 2098-A00 is the only model 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 (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.	

IBM 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 capacity indicators (or speeds), 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. The 700 capacity indicator is the only z9-EC host that can be configured with no General Purpose CPs. Any available CP can be configured as a General Purpose CP (within the limits for the 400, 500, and 600 described above), as a zAAP (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

IBM System z9-BC
  When configuring, there is a choice of 2 models

2096-R07	Up to 6 CPs may be configured
2096-S07	Up to 7 CPs may be configured

There are 26 General Purpose CP capacity indicators (or speeds), *A*, *B*, *C*, *D*…*Z* (*Z* is the full speed model). Speeds **A** through **J** can be configured on the 2096-R07. Speeds **K** through **Z** may 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. The 2096-Z00 is the only model 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 (limited to the number of General Purpose CPs), as a zIIP (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 capacity indicator installed.

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

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

- **General Purpose**   For zEC12, z196, z10-EC, z9-EC, and z990 processors, the range allowed is implied by the model specification, and you must provide a number. For all other processor families, this number is fixed based on the model specification.

- **zAAP**   May only be configured on zEC12, zBC12, z196, z114, z10, z9, z990, and z890 processors. The maximum number of zAAPs is generally restricted to the total number of General Purpose CPs configured.

- **zIIP**   May only be configured on zEC12, zBC12, z196, z114, z10, and z9 processors. The maximum number of zIIPs is generally restricted to the total number of General Purpose CPs configured.

- **IFL**   May be configured on any System z processor.

- **ICF**   May be configured on any System z processor. For zEC12 processors any number of the CPs configured may be ICF. For all other processor families, the maximum 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 logical CPs 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. zAAPs (zEC12, zBC12, z196, z114, z10, z9, z990, and z890 processors), zIIPs (zEC12, zBC12, z196, z114, z10, z9 processors), IFLs, and ICFs are configured separately. On zEC12, zBC12, z196, z114, z10 and z9 processors, each of these CP types operates as a separate CP pool. On z990 and z890 families, these 4 specialty engine types operate as a single pool, and any zAAP, IFL, or ICF partition may be configured to use any/all of the CPs in the pool, after DED/SHR rules are applied. This reflects the way the hardware actually works, but does not reflect any usage restrictions implied by the specific processor features that have been ordered and are installed.

4. z/OS, z/VM without Linux guests, and z/VSE can only be assigned to General Purpose partitions. z/VM with Linux guests or Native Linux may be assigned to an IFL partition or to a General Purpose partition. zAware can be assigned to and IFL partition or to a General Purpose partition. CFCC may be assigned to an ICF partition or to a General Purpose partition.

5. The workload category assigned to a zAAP or zIIP partition will be identical to that assigned to the associated General Purpose partition, which must be designated with z/OS-1.6 or later.

6. If you reconfigure the LPAR host for which an LPAR 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 partition has been automatically excluded, must be manually re-included.

On multi-book configurations for zEC12, z196, z10-EC, z9-EC, and z990, **zPCR** recognizes that General Purpose, zAAP, and zIIP CPs are allocated upward, starting in the first book and IFL and ICF CPs are allocated downward, starting in the last book. An intersection of GP/zAAP/zIIP CPs with IFL/ICF CPs can occur in only one book. When such an intersection exists, algorithms apply a partitioning cost that reflects contention between the two CP groups in that book only. The partitioning cost for the CPs in the remaining books will be less. A bar graph is available from the *Partition Detail Report* window showing how these CPs would be allocated and where any contention between them might occur.

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 General Purpose partitions (and any associated zAAP and/or zIIP logical CPs), the IFL button to define IFL partitions, or the ICF button to define ICF partitions.

The primary table (at the top) provides an area where individual partitions are defined to the host, whose configuration is stated above. The overall partition configuration, shown on the line below, is dynamically updated as partitions are added or modified.

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.

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

- **Include**   ☑   Include/exclude this partition in the capacity assessment

Partition Identification

- **No**.   P   The relative partition number (determined by **zPCR**)

- **Type**   P   The real CP pool to which the partition is assigned

- **Name**   E   Partition name; assigned default can be modified

- **SCP**   ▼   z/OS, z/VM, z/VSE, Linux, zAware, or CFCC may be assigned to a General Purpose partition. z/VM (if supporting Linux guests), Linux, and zAware can be assigned to and IFL partition. And CFCC can be assigned to an ICF partition.

- **Workload**   ▼   For z/OS, z/VM, and Linux: 5 LSPR workload categories.
  For z/VSE, three generic workloads are available.
  For zAware the workload must be zAware.
  For CFCC the workload must be CFCC.

Partition Configuration

- **Mode**   ▼   DED (dedicated) or SHR (shared)

- **LCPs**   ▼   Number of logical CPs assigned

- **Weight**   E   SHR partitions only; relative weight assigned

- **CAP**   ☑   SHR partitions only; checkbox if capped

General Purpose Partitions with Associated Specialty Logical CPs

- **z/OS zAAPs**   ▼   Number of zAAP logical CPs to be associated with a z/OS-1.6 (or later) partition. This column will appear only if real zAAP CPs have been configured.

- **z/OS zIIPs**   ▼   Number of zIIP logical CPs to be associated with a z/OS-1.6 (or later) partition. This column will appear only if real zIIP CPs have been configured.

- **z/VM IFLs**   ▼   Number of IFL logical CPs to be associated with a z/VM-5.4 (or later) partition. The partition must be on a zEC12, zBC12, z196, z114, or z10 processor and must be designated as *Mode=zVM*. This column will appear only when real IFLs have been configured on the host.

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 100 partitions (sum of General Purpose, zAAP, zIIP, 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 and zIIP partitions do not count against the maximum). 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 definition. 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.

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.

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 logical CPs associated with a General Purpose partition, this setting also applies to them. This action can also be done from the *Partition Detail Report* window for each individually identified partition.

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 to any character string by double clicking the field, keying in the name, and pressing **Enter**.

SCP: Must be specified for each partition. When defining General Purpose partitions, z/OS , z/VM, z/VSE, Linux, zAware, or CFCC can be selected. When the partition is not General Purpose, the SCP selection is limited. For an IFL partition, z/VM (Linux guests assumed), Linux, or zAware 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 selection from the dropdown list provided.

For z/OS, a specific version must be selected. 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-1.13).

z/OS Feature Support by Version			
Version	**Maximum Logical CPs**	**zAAP and zIIP Logical CPs**	**HiperDispatch**
z/OS-1.4	**16**	**No**	**No**
z/OS-1.5	**16**	**No**	**No**
z/OS-1.6	**32**	**Yes**	**No**
z/OS-1.7	**32**	**Yes**	**No**
z/OS-1.8	**32**	**Yes**	**Yes**
z/OS-1.9	**64**	**Yes**	**Yes**
z/OS-1.10	**64**	**Yes**	**Yes**
z/OS-1.11	**80**	**Yes**	**Yes**
z/OS-1.12	**100**	**Yes**	**Yes**
z/OS-1.13	**100**	**Yes**	**Yes**
z/OS-2.1	**100**	**Yes**	**Yes**

Note that zEC12 and z196 *Power-Saving* mode can be turned on regardless of the SCPs that are running. However, the ability to report on its status is only available from a partition running z/OS-1.10 or later.

Whenever the *LPAR host* processor is defined as a **zEC12** or **zBC12** processor, z/OS-1.10 or later must be specified; and whenever the *LPAR host* processor is defined as a **z196**, **z114**, or **z10** processor, z/OS-1.7 or later must be specified, since prior versions are not supported. The z/OS version selected will not affect capacity results since they are driven off a single LSPR table derived from z/OS-1.13 measurements. While actual capacity relationships may vary somewhat for different z/OS versions, the difference is considered insignificant enough to allow the use of a single LSPR table.

z/VM is considered generically by **zPCR**, rather than by specific versions.

- For z/VM version 5.4 and later, up to 32 real CPs are supported, which can be General Purpose or IFL CPs. On zEC12, zBC12, z196, z114 and z10 processors, IFLs can be associated with a General Purpose z/VM if the partition is defined as *Mode=zVM*.

- For z/VM versions prior to 5.4, up to 16 real CPs are supported, which can be either General Purpose CPs or IFL CPs.

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*, and *Linux*, the workload can be any of the 5 supported workload categories. For *z/VSE*, the workload can be specified as *Batch*, *Mixed*, or *Online*.   Make your selection from the dropdown list provided. For *zAware*, the workload will be *zAware*. For *CFCC*, the workload will be set to *CFCC*.

Type (LP-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 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 real CPs must be available in the pool to which the partition is assigned. Dedicated partitions will remove real CPs from the shared pool.

Note: There are limitations on the number of logical CPs that can be specified for some SCP/workload environments. For example, z/VM is limited to 32 logical CPs, since that is all it supports. Linux is limited to 32 logical CPs due to the absence of LSPR measurement data.

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.

CAP: A checkbox used to indicate that a shared partition is to be considered as being capped. The CAP specification can be changed from the *Partition Detail Report* window, discussed later.

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.

z/OS partitions with HiperDispatch (zEC12, zBC12, z196, z114, and z10 only)

LSPR data for zEC12, zBC12, z196, z114, and z10 processors assumes that HiperDispatch is active for z/OS 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 logical CPs that it considers to be excessive for the z/OS 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 logical CPs should not be considered when defining the number of shared LCPs for a z/OS partition***. Parked logical CPs can be determined from EDF or RMF data (note that a separate report is required for each such partition.)

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.

Specialty Engines Associated with a General Purpose Partition

zAAP and zIIP engines must be defined as logical CPs associated with a General Purpose partition running a z/OS version that supports them. IFL engines may be defined as logical CPs associated with a General Purpose partition running z/VM (hardware definition for the partition must be ***Mode=z/VM***).

▼ **z/OS zAAPs**: Define the number of zAAP logical CPs to be associated with the General Purpose partition when the SCP is z/OS-1.6 or later. This column will only appear if zAAP real CPs have been defined to the LPAR host. The selection list only includes logical CP counts that are valid for the currently defined LPAR host configuration (adequate CPs must be available in the pool to which the partition is assigned). Initially, the zAAP partition will assume the **Name**, **SCP**, **Workload**, **Mode**, **Weight**, and **CAP** assignment of the associated General Purpose partition. On z9 and later processors, **Mode**, **LCPs**, **Weight**, and **CAP** can subsequently be modified via the ***Partition Detail Report*** window. A zAAP partition can only be created or deleted from the ***LPAR Partition Definition*** window, via its associated General Purpose partition.

▼ **z/OS zIIPs**: Define the number of zIIP logical CPs to be associated with the General Purpose partition when the SCP is z/OS-1.6 or later. This column will only appear if zIIP

real CPs have been defined to the LPAR host. The selection list only includes logical CP counts that are valid for the currently defined LPAR host configuration. Initially, the zIIP partition will assume the **Name**, **SCP**, **Workload**, **Mode**, **Weight**, and **CAP** assignment of the associated General Purpose partition. **Mode**, **LCPs**, **Weight**, and **CAP** can subsequently be modified via the ***Partition Detail Report*** window. A zIIP partition can only be created or deleted from the ***LPAR Partition Definition*** window, via its associated General Purpose partition.

▼ **z/VM IFLs**: Define the number of IFL logical CPs to be associated with the General Purpose partition when the SCP is z/VM. This column will only appear for zEC12, zBC12, z196, z114, and z10 processors if IFL real CPs have been defined to the LPAR host. The selection list only includes logical CP counts that are valid for the currently defined LPAR host configuration. Initially, the IFL partition will assume the **Name**, **SCP**, **Workload**, **Mode**, **Weight**, and **CAP** assignment of the associated General Purpose partition. **Mode**, **LCPs**, **Weight**, and **CAP** can subsequently be modified via the ***Partition Detail Report*** window. An IFL partition associated with a General Purpose z/VM partition can only be created or deleted from the ***Partition Definition*** window, via its associated General Purpose partition.

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.

zAAP and zIIP Logical CP Considerations

zPCR considers zAAP logical CPs that are associated with a z/OS General Purpose partition as a separate zAAP partition, for capacity projection purposes. Likewise, zIIP logical CPs that are associated with a z/OS General Purpose partition are considered as a separate zIIP partition. zAAP partition logical CPs and zIIP partition logical CPs are managed by z/OS. A single central storage is shared by zAAP logical CPs, zIIP logical CPs, and the associated General Purpose logical CPs.

When initially created, a zAAP or zIIP partition inherits the associated General Purpose partition's ***Number*** (preceded by an asterisk), **Name**, **SCP**, **Workload**, **Mode**, **Weight**, and **CAP** assignments. On zEC12, zBC12, z196, z114, z10, and z9 processors, a zAAP or zIIP partition's **Mode**, **Weight** and **Capping** assignment can subsequently be changed (these changes can only be made from the ***Partition Detail Report*** window).

On z990 and z890 processors, the zAAP partition's **Mode**, **Weight**, and **CAP** assignment cannot be changed. If the LPAR host is changed from a zEC12, zBC12, z196, z114, z10 or z9 processor to a z990 or z890 processor, any zAAP partitions will automatically inherit the associated General Purpose partition's **Mode**, **Weight**, and **CAP** settings.

IFL Logical CPs Associated with a z/VM General Purpose Partition

zPCR considers IFL logical CPs that are associated with a z/VM General Purpose partition as a separate IFL partition, for capacity projection purposes. The IFL logical CPs are managed by z/VM. A single central storage is shared by the IFL logical CPs and the associated General Purpose logical CPs.

When initially created, an associated IFL partition inherits the General Purpose partition's **Number** (preceded by an asterisk), **Name**, **SCP**, **Workload**, **Mode**, **Weight**, and **CAP** assignments. The IFL partition's **Workload**, **Mode**, **Weight** and **CAP** assignment can subsequently be changed (these changes can only be made from the **Partition Detail Report** window). Associated IFL partitions are only allowed on zEC12, zBC12, z196, z114, and z10 processors.

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.

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.

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

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 zEC12, zBC12, z196, z114, and z10 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 **CP3KEXTR** program 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 System 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.

EDFs will provide information concerning the number of parked logical CPs for partitions that are running with HiperDispatch enabled. Note that a separate EDF must be generated for each such partition. For these EDFs, **zPCR** will identify both the number of assigned LCPs and the number of parked LCPs at the time the partitions are being prepared for transfer to the study. 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.

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.

z/VM v5.4 and later can capture CPU-MF data which will also be included in an EDF. CPU-MF data is useful to understand the characteristics of the overall z/VM workload. 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 General Purpose partition while also supporting IFL logical CPs, separate CPU-MF data is captured for each. Hence, the workload category chosen for each type could be different.

Creating z/VM EDFs

The **CP3KVMXT** program 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.

Reading EDFs

Each EDF contains all the necessary information to define the LPAR host processor and its entire partition configuration. **zPCR** can read EDFs for the purpose of defining the entire configuration or to simply copy partitions into an existing configuration.

For partitions where **CPU-MF** data is desired, a separate EDF is required for each. Within **zPCR**, multiple EDFs can be selected to be read at the same time. 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 **CPU-MF** information specific to each of the other partitions.

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.
- 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.

EDFs are read in via the **EDF Interval Selection** window. 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 concerning **CPU MF Counters** is obtained from each individual EDF.

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.
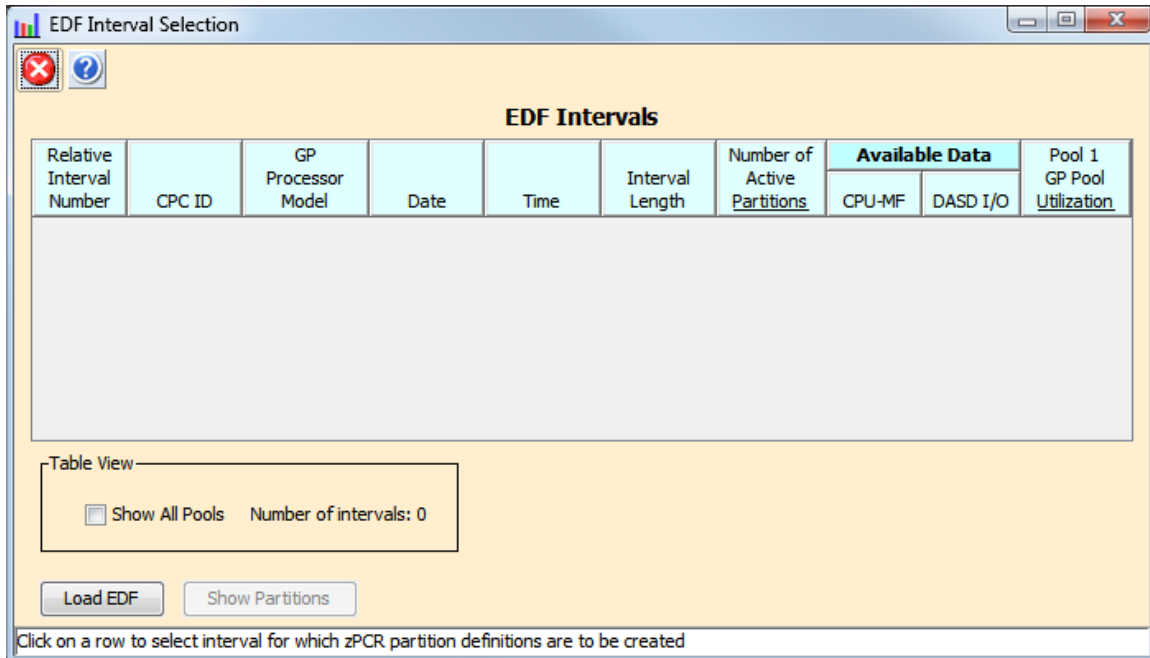
The ***EDF Interval Selection*** window provides the means to load one or more EDFs. For information concerning EDFs, see EDF Overview.

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 assumed is "***edf***", any file extension can be loaded.

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, an error message is issued and no EDF intervals are shown.

Loading Multiple EDFs

zPCR can process multiple EDFs (each representing an individual partition) in a single pass. The 1st file in the selection list is 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 CPU-MF information.

To load multiple EDFs, hold down the ***Ctrl*** (or ***Shift***) key while selecting files. When in ***Advanced-Mode***, the selected files can be dragged to any undefined LPAR icon on the ***Advanced-Mode*** 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 align closely. If they do not, no EDF intervals will be shown.

Once validated, the EDF is loaded and a title line reveals the EDF filename and location. The EDF will remain available until a different EDF is loaded, allowing the subsequent return to this window to select a different interval.

EDF Intervals
EDF File Name: D:\zPCR EDF Files\UGSample.edf

Relative Interval Number	CPC ID	GP Processor Model	Date	Time	Interval Length	Number of Active Partitions	Available Data		Pool 1 GP Pool Utilization
							CPU-MF	DASD I/O	
1.	CECAAAA	2097-713	2010-09-26	23:45:00	00:15:00	4		✓	18.80%
2.	CECAAAA	2097-713	2010-09-27	10:00:00	01:00:00	4	✓	✓	55.56%
3.	CECAAAA	2097-713	2010-09-27	11:00:00	01:00:00	4	✓	✓	58.32%
4.	CECAAAA	2097-713	2010-09-27	12:00:00	01:00:00	4	✓	✓	52.14%
5.	CECAAAA	2097-713	2010-09-27	13:00:00	01:00:00	4	✓	✓	49.95%
6.	CECAAAA	2097-713	2010-09-27	14:00:00	01:00:00	4	✓	✓	52.73%
7.	CECAAAA	2097-713	2010-09-27	15:00:00	01:00:00	4	✓	✓	50.64%
8.	CECAAAA	2097-713	2010-09-27	16:00:00	01:00:00	4	✓	✓	40.30%
9.	CECAAAA	2097-713	2010-09-27	17:00:00	01:00:00	4	✓	✓	29.83%
10.	CECAAAA	2097-713	2010-09-27	18:00:00	01:00:00	4	✓	✓	26.37%
11	CECAAAA	2097-713	2010-09-27	19:00:00	01:00:00	4	✓	✓	28.48%

Table View
☐ Show All Pools   Number of intervals: 15

Load EDF   Show Partitions

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.

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.

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, when an interval has been selected and the
⬛ **Show Partitions** ⬛ button is clicked:

- When entered via **Create Host and Partitions From**, the **Create LPAR Configuration from EDF** window appears.

- When entered via **Copy Partitions From**, 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 LPAR host in the **zPCR** study. When the window title bar
reads **Copy Partitions from EDF**, 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 **zPCR** configuration.

Partition Identification and Partition Configuration

The primary table contains a list of all the partitions found in the selected EDF interval
with their definition 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, by pool, in the following order:

1. General Purpose
2. zAAP (listed below the associated General Purpose partition)
3. zIIP (listed below the associated General Purpose partition)
4. IFL
5. ICF

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.

SCP Assignment

A checkmark in the column identified as *From EDF* indicates that EDF was produced for the partition. The *SCP* field is populated with the SCP version found in the partition's z/OS EDF or z/VM EDF, when available. Otherwise, the field is populated with a default SCP for the partition type as follows:

- *General Purpose* partitions:
  - **EDF from z/OS**: All General Purpose partitions are defaulted to *z/OS* (which also applies to any associated **zAAP** or **zIIP** partitions). Note that the specific version of z/OS has no effect on capacity results; rather it is used to enforce partitioning rules.

    z/VM, *z/VSE*, *Linux*, or *CFCC* could actually be in control of a General Purpose partition. When this is the case, the SCP will need to be changed after the LPAR configuration has been transferred into **zPCR**.

  - **EDF from z/VM**: All partitions characterized by the EDF are set to z/VM. All remaining General Purpose partitions are defaulted to the radio button setting in the *Change SCP for GP Partitions* group box (z/OS or z/VM). Note that the specific version of z/VM or z/OS has no effect on capacity results.

    z/OS, *z/VSE*, *Linux*, or *CFCC* could actually be in control of a General Purpose partition. When this is the case, the SCP will need to be changed after the LPAR configuration has been transferred into **zPCR**.

- *IFL* partitions are defaulted to *z/VM*. If *Linux* is actually in control of the partition, it can be changed on this window or after the LPAR configuration has been transferred into **zPCR**.

- *ICF* partitions are defaulted to *CFCC*.

Additional Information (z/OS Partitions Only)

- **HD Active**: A checkmark indicates whether or not HiperDispatch was active for the partition.

- **LCPs Parked**: Indicates the number of the partitions logical CPs that were parked during the subject 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* check box to have all parked LCPs remove from partitions when the LPAR configuration is transferred to **zPCR**.

Workload Category Determination

When **CPU-MF** data is available for a partition, information for the following items are displayed in the **CPU-MF** columns.

- **RNI** "Relative Nest Intensity" is computed for each z/OS partition where CPU-MF data is included in the EDF.

- **Workload Assignment**: Displays the z/OS or z/VM workload category that would be selected based on **CPU-MF** data.

The **Method Used** column (on the right side of the window) reveals the method applied for **Assigned Workload** (**CPU-MF**, **Default**, or **Other**).

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.

2. **Default** workload choice (by SCP)
   - **z/OS**   (General Purpose, zAAP, & zIIP) **Average**
   - **z/VM**   (General Purpose and/or IFL) **Average/LV**
   - **Linux**   (General Purpose or IFL) **Average/L**
   - **z/VSE**   (General Purpose only) **Mixed**
   - **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.

The **SCP** for IFL partitions may be altered to either **z/VM** or **Linux**. If the **SCP** for a General Purpose partition needs to be changed, it must be done after the configuration has been transferred into **zPCR**.

The **Assigned Workload** field for a z/OS, z/VM, Linux, or z/VSE partition may be altered as desired. The **Assigned Workload** categories available will depend on the SCP specified.

The **Method Used** column will reflect the highest priority assignment method that matches the **Assigned Workload** field (i.e., **CPU-MF** or **Default**). If there is no match, **Other** will appear in red.

The **SCP** and **Assigned Workload** for zAAP and zIIP partitions are always derived from the associated General Purpose partition.

For the **CPU-MF** method, the **RNI** and internal algorithms specific to each processor family are used to determine the **Workload Choice**.

Some exceptions to consider

- For z990 processors, the actual hardware model (i.e., 2084-A08, B16, C24, or D32) might not be precisely determined by EDF. In these cases the model assumed is the smallest that could accommodate all of the CPs defined. The actual processor model should be determined, and, if different, the LPAR host definition in **zPCR** should be changed accordingly. The model designation will have some minor effect on capacity results.

- For z990, z890, z900, and z800 processors, partitions that are not General Purpose are always identified in EDF as Type=ICF. If any such partitions are actually IFLs, they will need to be modified to Type=IFL. All ICF partitions will have the *Type* field enabled for input. Click on the *Type* field entry to reveal a dropdown list which provides a selection for IFL. When this change is made, the *SCP/workload* is changed to the default for an IFL (z/VM with Average/LV workload), which can be subsequently modified from within **zPCR**.

- For z990 and z890 processors, zAAP CPs are identified as ICFs. However, since a zAAP partition always has the same name as its associated General Purpose partition, type=zAAP is automatically assigned and the SCP/workload of the associated General Purpose partition is used.

Other than the items noted above, partition definition changes such as changing the SCP for a General Purpose partition, and adding or removing zAAP or zIIP logical CPs for a General Purpose partition must be made after the configuration has been transferred into **zPCR**.

HiperDispatch Considerations

For partitions that are running with HiperDispatch enabled, the partition's EDF provides information concerning the number of parked logical LCPs. A separate EDF should be included for each of the partitions where knowing the number or parked LCPs is important.

> **The number of LCPs assigned to a partition should always be reduced by the number of parked LCPs. This is necessary to obtain reasonable capacity projections for the configuration**.

If ☑ *Remove Parked LCPs from Partition LCP Count* is checked, any partition parked LCPs are subtracted from the defined LCP count (fractional LCPs are always rounded up). Otherwise, the defined LCP count is transferred to the **zPCR** study.

Note that parked LCP information is carried forward into **zPCR**. 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 Assignment Alternatives.

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 associated General Purpose 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 LPAR configuration description will be initialized to indicate the source.

- **Copy Partitions From EDF**, click the ⸢ **Copy Partitions** ⸥ button to add the selected partition definitions to the currently active 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 either of these windows.

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.

When copying partitions intp zPCR, remove Parked LCPs from the LCP Count check box to have all parked LCPs that have been identified removed from partitions when the LPAR configuration is transferred to **zPCR**. It is advisable that this box always be checked.

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 *EDF Interval Selection* window. A 2nd **Cancel** will return to **zPCR**.

Referencing Partition Utilizations from EDF in zPCR

When both the LPAR host and its <u>entire</u> partition configuration are obtained from EDF (or RMF), 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 100% 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). 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 concerning CFCC Partitions: CFCC, when 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. The correct utilization values can be independently determined via EDF.

RMF Interval Selection
LPAR Configuration Capacity Planning

From the *LPAR Host and Partition Configuration* window, click the $\boxed{\text{RMF}}$ button in either the *Define LPAR Host Processor* group box or the *Define Partitions* group box to open the *RMF Interval Selection* window.
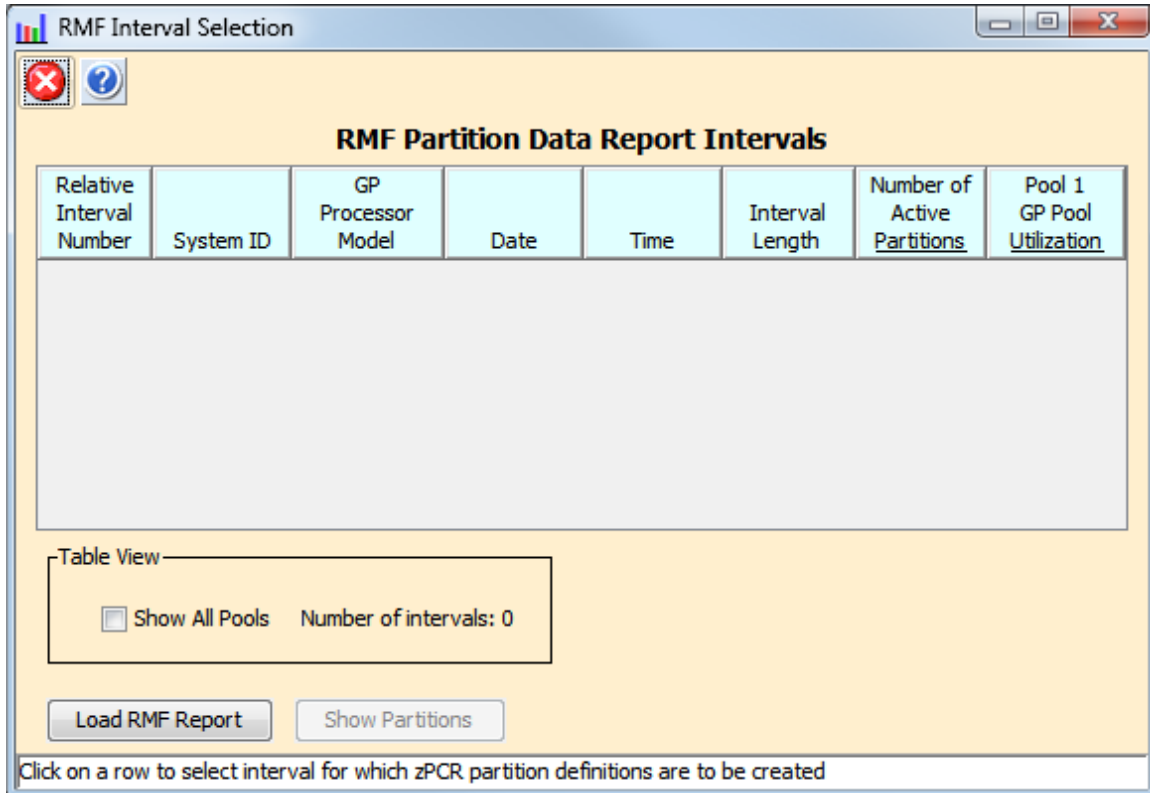
The *RMF Interval Selection* window provides the means to load an RMF report file. Click the $\boxed{\textbf{Load RMF Report}}$ button. A file dialog box will appear from which you can specify the drive:\directory\filename for the RMF report. The input may be a traditional RMF report (text format, default file extension assumed is "*txt*"), or the newer RMF XML format (default RMF file extension assumed is "*xml*"). However, for either of these formats, any file extension can be used. Traditional RMF reports must have a record length of 133 bytes, or if CC information is included, 134 bytes.

When read, the RMF file is searched for all report intervals. If the RMF report file 1) does not appear to be a standard RMF report, 2) was generated on a z/OS version not supported by **zPCR**, 3) does not include both a *Partition Data Report* (*PDR*) and *CPU Activity Report*, or 4) appears to be invalid, an error message is issued and no RMF intervals are shown.

Once validated, the RMF report is loaded. It will remain available until a different RMF report is loaded, allowing you to subsequently return to this window to select a different interval.

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 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 the default settings shown in the *Default SCP/Workload for Partitions* group box. The default SCP/workload assignments can be subsequently changed in **zPCR**, on the *Define Partitions* window or the *Partition Detail Report* 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 <u>Get Partitions from RMF</u> 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, when an interval has been selected and the
⌐Show Partitions⌐ button is clicked:

- When entered via **Create Host and Partitions From**, the ***Create LPAR
  Configuration from RMF*** window appears.

- When entered via **Copy Partitions From**, 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 LPAR host in the **zPCR** study. 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 **zPCR** configuration.

Note: On zEC12, zBC12, z196, z114, and z10 processors, *CPU-MF* (CPU Measurement Facility) data can be captured for z/OS or z/VM partitions. The use of CPU-MF data provides a reliable method for determining the appropriate LSPR workload category to represent a partition's actual production workload. CPU-MF information can only be transferred to **zPCR** via EDF. For details, see Get Partitions from EDF.

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 (listed below the associated General Purpose partition))
3. zIIP (listed below the associated General Purpose partition)
4. IFL
5. ICF

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 for each partition is populated with a default value for the partition type as follows:

- *General Purpose* partitions are assigned *z/OS* (which also applies to any associated *zAAP* or *zIIP* partitions). Note that the specific version of z/OS has no effect on capacity results; rather it is used to enforce partitioning rules.

  z/VM, *z/VSE*, *Linux*, or *CFCC* could actually be in control of a General Purpose partition. When this is the case, the SCP will need to be changed after the LPAR configuration has been transferred into **zPCR**.

- *IFL* partitions are defaulted to *z/VM*. If *Linux* is actually in control of the partition, it can be changed on this window or after the LPAR configuration has been transferred into **zPCR**.

- *ICF* partitions are defaulted to *CFCC*.

Additional Information (z/OS 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.

- **LCPs Parked**: Indicates the number of the partitions logical CPs that were parked during the subject 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* check box to have all parked LCPs remove from partitions when the LPAR configuration is transferred to **zPCR**.

Workload Category Determination

The **Method Used** column (on the right side of the window) reveals the method applied for **Assigned Workload** (**Default**, or **Other**).

With RMF input, the default LSPR workload category for the SCP is set as the **Assigned Workload** for each partition (under **Partition Identification**), as one of the following:

- **z/OS**   (General Purpose, zAAP, & zIIP)   **Average**
- **z/VM**   (General Purpose or IFL)   **Average/LV**
- **Linux**   (General Purpose or IFL)   **Average/L**
- **z/VSE**   (General Purpose only)   **Mixed**
- **CFCC**   (General Purpose or ICF)   **CFCC**

When CPU-MF data (via EDF) is not available, the default workload assignment for the SCP should generally be used.

The **SCP** for IFL partitions may be altered (either z/VM or Linux). If the **SCP** for General Purpose partitions needs to be altered, it must be done after the configuration has been transferred to **zPCR**.

The **Assigned Workload** field for a partition may be altered as desired. The **Assigned Workload** categories available will depend on the SCP specified (i.e., z/OS, z/VM/ Linux, z/VSE, or CFCC). The **Method Used** column will reflect the assignment method of **Default**, or, if there is no match, **Other** will appear in red.

The **SCP** and **Assigned Workload** for zAAP and zIIP partitions are always derived from the associated General Purpose partition.

Some exceptions to consider

- For z990 processors, the actual hardware model (i.e., 2084-A08, B16, C24, or D32) cannot be precisely determined by RMF versions prior to v1.7. In these cases the model assumed is the smallest that could accommodate all of the CPs defined. The actual processor model should be determined, and, if different, the LPAR host definition in **zPCR** should be changed accordingly. The model designated will have some effect on capacity results.

- For z990, z890, z900, and z800 processors, partitions that are not General Purpose are always identified in RMF as Type=ICF. If any such partitions are actually IFLs, they will need to be modified to Type=IFL. All ICF partitions will have the **Type** field enabled for input. Click on the **Type** field entry to reveal a dropdown list which provides a selection for IFL. When this change is made, the **SCP/workload** is changed to the default for an IFL (z/VM with Average/LV workload), which can be subsequently modified from within **zPCR**.

- For z990 and z890 processors, zAAP CPs are identified as ICFs. However, since a zAAP partition always has the same name as its associated General Purpose partition, type=zAAP is automatically assigned and the SCP/workload of the associated General Purpose partition is used.

Other than the items noted above, partition definition changes such as changing the SCP for a General Purpose partition, and adding or removing zAAP or zIIP logical CPs for a General Purpose partition must be made after the configuration has been transferred into **zPCR**.

HiperDispatch Considerations

For partitions that are running with HiperDispatch enabled, RMF's **CPU Activity Report** provides information concerning the number of parked logical LCPs. A separate **CPU Activity Report** should be included in the RMF for each of the partitions where knowing the number or parked LCPs is important.

> **The number of LCPs assigned to a partition should always be reduced by the number of parked LCPs. This is necessary to obtain reasonable capacity projections for the configuration**.

If ☑ **Remove Parked LCPs from Partition LCP Count** is checked, any partition parked LCPs are subtracted from the defined LCP count (fractional LCPs are always rounded up). Otherwise, the defined LCP count is transferred to the **zPCR** study.

Note that parked LCP information is carried forward into **zPCR**. 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 Assignment Alternatives.

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 associated General Purpose 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 LPAR configuration description will be initialized to indicate the source.

- **Copy Partitions From RMF**, click the ⌈ **Copy Partitions** ⌋ button to add the selected partition definitions to the currently active LPAR configuration.

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 either of these windows.

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 intp zPCR, remove Parked LCPs from the LCP Count* check box 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 2^nd **Cancel** will return to **zPCR**.

Referencing Partition Utilizations from RMF in zPCR

When both the LPAR host and its <u>entire</u> 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 100% 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). 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 concerning CFCC Partitions: CFCC, when in a dedicated partition, always runs at 100% utilization. For these partitions, the actual capacity being consumed cannot be determined as for the other partition types. The correct utilization values can be independently determined via RMF.

Copy Partitions from Study

LPAR Configuration Capacity Planning

When defining the LPAR configuration, you can copy partitions that were defined in a previous **zPCR** study using the **Copy Partitions from Previous Study** function, available on the **Host and Partition Configuration** window.

In the **Define Partitions** group box, under **Copy Partitions From**, click the **zPCR Study** button to open a previous **zPCR** study file. 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 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.

If the **zPCR** study is an *Advanced-Mode* study with more than one LPAR configuration defined, the *Advanced-Mode Configurations* group box will appear. From the drop down list, select the name of the specific LPAR configuration from which partitions are to be copied and that partition configuration will appear.

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 for controlling partition selection en-masse.

Once a partition has been checked, the **Copy Partitions** button is enabled. Clicking this button will copy each 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 logical CPs 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 General Purpose partitions with associated zAAP, zIIP, or IFL logical CPs, 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 LCP partition's ☐ *Include* checkbox will be able to be checked.

Partitions copied from a previous study will appear in the *Define Partitions* 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 logical CPs, and IFLs associated with z/VM, must be made from the *Define Partitions* window.

Click the **Cancel** toolbar icon to close the *Copy Partitions from Study* window without transferring any partitions to the active study.

zAAP/zIIP Capacity Considerations

LPAR Configuration Capacity Planning

On zEC12, zBC12, z196, z114, z10, z9, z990, and z890 processor models, zAAPs (*System z Application Assist Processor* may be configured. When running Java applications under z/OS-1.6 or later, zAAP processors may be used in lieu of General Purpose CPs to run certain Java code. This has the effect of relieving the General Purpose CPs assigned to that z/OS partition of some of the load and, hence, potentially reducing software licensing costs.

On zEC12, zBC12, z196, z114, z10, and z9 processor models, zIIPs (*System z Integrated Information Processor*) may be configured. When running certain workloads under z/OS-1.6 or later (DB2 is a prime example), zIIP processors may be used in lieu of General Purpose CPs to run portions of the work. This has the effect of relieving the General Purpose CPs assigned to that z/OS partition of some of the load and, hence, potentially reducing software licensing costs.

On zEC12, zBC12, z196, z114, z10, and z9 processor models running z/OS-1.9 or later, *zAAP on zIIP Capability* may be enabled for a partition, allowing zAAP eligible work to run on zIIP logical CPs. When enabled, zAAP and zIIP eligible work can only be directed to zIIP logical CPs. Therefore, zAAP logical CPs cannot be used by that partition (note that **zPCR** does not know that the partition is using this capability; therefore you need to ensure that the partition is defined with no zAAP logical CPs).

zAAPs and zIIPs, purchased as a feature code, are configured to the LPAR host similar to IFLs and ICFs. They are used in the form of zAAP or zIIP logical CPs, associated with a z/OS-1.6 (or later) partition. zAAP/zIIP associations are considered as unique partitions in **zPCR** for the purposes of projecting capacity. Each of these partitions will be defined identically with its associated z/OS partition for the following items:

- Name
- SCP and Workload
- Mode (may be modified on System z9 and later processors)
- Weight (may be modified on System z9 and later processors)
- CAP (may be modified on System z9 and later processors)

Several questions concerning potential zAAP or zIIP capacity can be answered using **zPCR**

1. What is the capacity expectation for the General Purpose partition's logical CPs when also managing associated zAAP or zIIP logical CPs?

2. What capacity expectation for the partition's zAAP or zIIP logical CPs?

3. With one or more zAAP partitions configured on an IBM System z processor (z990 or z890 only), how might the weights for IFL and ICF partitions be rebalanced to best assure that each partition gets its intended capacity? Note that this is not an issue for zEC12, zBC12, z196, z114, z10, and z9 processors, since zAAPs, zIIPs, IFLs, and ICFs run in separate real CP pools.

4. If the capacity requirement for a z/OS General Purpose partition is reduced due to the addition of an associated zAAP or zIIP partition, what number of logical CPs might still be required?

5. How would the use of zAAP or zIIPs, and any associated partition configuration changes affect the overall capacity perspective of the LPAR host processor?

zPCR supports zAAPs used as zAAP logical CPs, and zIIPs used as zIIP logical CPs. Whenever a General Purpose partition is defined with its SCP as running z/OS-1.6 or later, and zAAP or zIIP real CPs exist in the hardware, additional input fields will appear, used to designate the number of zAAP/zIIP logical CPs (if any) to be associated with the General Purpose partition. For **zPCR** capacity reporting purposes, each partition's zAAP logical CPs and each partition's zIIP logical CPs are considered as a separated partition.

zAAP/zIIP partitions will appear in **zPCR** reports as a separate RCP pools on zEC12, zBC12, z196, z114, z10, and z9 processors. zAAP partitions will appear as part of the combined zAAP/IFL/ICF pool on z990 and z890 processors. zAAP/zIIP partition logical CPs will be as specified in the definition of the associated General Purpose partition, while the other LPAR configuration designations will be initialized as equal to those of the associated General Purpose partition. For zAAP/zIIP logical CPs on a On zEC12, zBC12, z196, z114, z10, and z9 processors, weight and CAP assignments can subsequently be adjusted.

On z990 and z890 processors, once the zAAP partitions are configured, the IFL and ICF weights can be rebalanced to get the desired capacity effect for each within the shared pool of zAAP/IFL/ICF real CPs.

zPCR capacity projections take into account the costs relating to a single z/OS managing the General Purpose logical CPs and its associated zAAP/zIIP logical CPs. Algorithms consider the processor family (i.e., zEC12, zBC12, z196, z114, z10, z9, z990, z890), the total number of logical CPs seen by z/OS, and the number of books installed. The same SCP/workload defined to the General Purpose partition is also assumed for its associated zAAP/zIIP partition.

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 customers 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 customers are authorized to use Specialty Engines only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

zAAP Specific Considerations

To derive capacity projections for zAAP logical CPs associated with a z/OS partition, the number of logical CPs available is considered as the sum of the General Purpose logical CPs and the zAAP logical CPs. A small switching cost is factored in, based on the processor family and the number of books. These items apply to both the General Purpose partition and the associated zAAP partition. Once these items are factored in, the partitions are considered identically to other partition types to determine the final capacity projections.

When adding zAAP CPs to a configuration, zAAP eligible work that formerly ran on General Purpose logical CPs can now run on the zAAP logical CPs. As a result, it may be possible to reduce the number of General Purpose logical CPs, and therefore reduce the number of General Purpose RCPs.

If the amount of zAAP eligible workload exceeds the zAAP logical CP capacity available, some of that work can be run on the General Purpose logical CPs.

zPCR can only project the capacity that is available on zAAP logical CPs. Whether or not that capacity can be exploited can only be determined from sources outside of **zPCR**.

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 there are no zAAP logical CPs defined, the entire workload is considered 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 logical CPs. This is the portion of the workload to consider when determining how much zAAP capacity could actually be consumed.

For Java content workloads such as WebSphere, on the order of 50% or more is generally seen as being zAAP eligible.

Capacity projections are shown for the General Purpose partition and its associated zAAP logical CPs, as separate entries on the **Partition Detail Report** window. You can relate the capacity of the zAAP logical CPs (partition) to its associated General Purpose partition capacity by finding 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 General Purpose logical CPs and zAAP logical CPs, based on the characteristics of the application. How the combined capacity of General Purpose logical CPs and the associated zAAP logical CPs actually get used will depend on the ability of the workload to exploit zAAP processors.

Additional partitions defined will have an effect on the overall capacity expectation for the LPAR host, and for each individual partition.

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/atsmastr.nsf/WebIndex/WP100417

zIIP Specific Considerations

To derive capacity projections for zIIP logical CPs associated with a z/OS partition, the number of logical CPs available is considered as the sum of the General Purpose logical CPs and the zIIP logical CPs. A small switching cost is factored in, based on the processor family and the number of books. These items apply to both the General Purpose partition and the associated zAAP partition. Once these items are factored in, the partitions are considered identically to other partition types to determine the final capacity projections.

When adding zIIP CPs to a configuration, zIIP eligible work that formerly ran on General Purpose logical CPs can now run on the zAAP logical CPs. As a result, it may be possible to reduce the number of General Purpose logical CPs, and therefore reduce the number of General Purpose RCPs.

If the amount of zIIP eligible workload exceeds the zIIP logical CP capacity available, some of that work can be run on the General Purpose logical CPs.

zPCR can only project the capacity that is available on zIIP logical CPs. Whether or not that capacity can be exploited can only be determined from sources outside of **zPCR**.

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 Java content (i.e., no zAAP eligible work). When there are no zAAP logical CPs 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 logical CPs. This is the portion of the workload to consider when determining how much zIIP capacity could actually be consumed.

zIIP logical CP capacity is always associated with that of DB2 content workload.

Capacity projections are shown for the General Purpose partition and its associated zIIP partition, as separate entries on the **Partition Detail Report** window. You can relate the capacity of a zIIP partition to its associated General Purpose partition capacity by finding 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 General Purpose logical CPs and zIIP logical CPs, based on the characteristics of the application. How the combined capacity of General Purpose logical CPs and the associated zIIP logical CPs actually get used will depend on the ability of the workload to exploit zIIP processors.

Additional partitions defined will have an effect on the overall capacity expectation for the LPAR host, and for each individual partition.

Host Summary Report

LPAR Configuration Capacity Planning

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.

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 Family
- Processor Model
- Speed Setting
- Books Configured
- Books Unused
- Maximum CPs
- Maximum Partitions

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 **System 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.

On the **z990 and z890 processor families**, zAAPs, IFLs, and ICFs are all managed from a single pool. In these cases, one additional set of rows will be displayed as the zAAP/IFL/ICF pool.

On the **z900 and z800 processor families**, IFLs and ICFs are managed from a single pool. In these cases, one additional set of rows will be displayed as the IFL/ICF pool.

A **Totals** line summarizes the allocations and capacity for the entire CPC.

Key to acronyms used
- **RCP**   Real CPs
- **Partitions**   Logical Partitions
- **LCP**   Logical CPs
- **LCP/RCP**   Ratio of Shared Logical CPs to Shared Real CPs
- **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 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). As the number of shared RCPs in a pool increases, the tolerance for a large LCP:RCP ratio diminishes. As a general rule, when the number of shared RCPs in pool is six, the LCP:RCP ratio should be kept below 3.00. And when the number of shared RCPs in 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 warning message 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:

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 unreasonable for the number of CPs in the pool. In this case, capacity results are not 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 considered excessive for the number of CPs in the pool. Capacity results continue to be provided. See note below.

> **Note:** Often the number of logical CPs 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 logical CPs should be removed from all partitions in the study that are being run that way.**
>
> **Whenever IRD is active, logical CPs that are varied offline should be removed from all partitions in the study that are being run that way.**

Note: Capacity results for LPAR configurations with a large number of partitions have not been validated by benchmarks or customer experience.
Displayed when the number of defined/active partitions exceeds 30:

Note: xx.xx% of XX (pool) capacity cannot be used by partitions due to capping assignments.
Displayed when some of the CP pool capacity cannot be consumed due to capping assignments:

Note: One or more partitions weights indicate more capacity than can be provided with the LCPs defined.
Displayed when the number of logical CPs assigned to a partition is inadequate to support the assigned weight percent:

Note: N defined partitions are excluded from the capacity assessments.
Displayed when the ☑ **Include** checkbox for one or more partitions is unchecked:

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.

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 description 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

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 window includes the description metrics for 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 a description of the LPAR host and its CP configuration, followed by the number of active partitions with their distribution across real CP 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, including capacity values is displayed below in a J-table format. Many of the fields (white background) serve as an alternate input; changing them will instantly modify results accordingly. See User Controls under "Getting Started" for a discussion on how to enter data into J-table input fields.

Partition Control

- **Include**   Checkbox to control partition participation in study

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
- **LCPs**   Number of logical CPs
- **Weight**   Weight value assigned (SHR only)
- **Weight %**   Relative weight within the partition's CP pool (SHR only)
- **CAP**   Capping assignment (SHR only)

Partition Capacity metrics

- **Minimum**   Minimum capacity realizable if 100% sharing contention
- **Maximum**   Maximum capacity realizable if no sharing contention

Note-1: The partition number assigned for zAAP and zIIP partitions is the number of its associated General Purpose partition. These partitions do not count towards the maximum number of partitions allowed.

Note-2: If the assigned weight for a partition results in a weight percent that cannot be accommodated due to the limited number of logical CPs assigned, the background for this field displays red. A note in the message box at the bottom of the window indicates the nature of the problem. Correction consists of increasing the number of logical CPs assigned to the partition or lowering its weight assignment. This condition is not necessarily problem for **zPCR**, as any capacity that cannot be used by such a partition is redistributed to the other partitions.

Minimum Capacity

A dedicated partition's minimum capacity will be the same as its maximum capacity. Capacity is based solely on the number of logical CPs assigned.

A shared partition's minimum capacity is that which can be realized when all partitions are actively competing for the available real CP resource. The primary determinant for minimum capacity is the partition's relative weight (or weight percent). In cases where a partition's weight percent is greater than its logical CP count would warrant (LCP ÷RCPs), the minimum capacity will be determined by the number of LCPs, and the capacity in excess, which cannot be consumed, is redistributed to the other shared partitions.

Maximum Capacity

A shared partition's maximum capacity can only be realized when no other partition is competing for the same real CP resource. The primary determinant for maximum capacity is the number of logical CPs 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 and maximum capacity values, even though a case could be made that LPAR cost might be less when only 1 partition is competing.

A capped partition will have its maximum capacity set exactly equal to its minimum capacity (which is controlled by weight %). In actual production, a capped partition capacity may see ±3.6% of its minimum capacity. Capping a partition has no effect on the capacity of other partitions.

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 General Purpose partition minimum capacity values

System z9 and later (separate CP pools exist for specialty engines

- **zAAP**   Sum of zAAP partition minimum capacity values
- **zIIP**   Sum of zIIP partition minimum capacity values
- **IFL**   Sum of IFL partition minimum capacity values
- **ICF**   Sum of ICF partition minimum capacity values

z990 and z990 (common pool of CPs exists for specialty engines)

- **zAAP/IFL/ICF**   Sum of zAAP/IFL/ICF partition minimum capacity values

z900 and z800 (common pool of CPs exists for specialty engines)

- **IFL/ICF**   Sum of IFL/ICF partition minimum capacity values

Overall LPAR host

- **Totals**   Sum of all partitions' minimum capacity values

Capacity results are determined based on the number of real CPs in each pool that can actually be exploited. If the total number of logical CPs assigned to a pool is less than the number of real CPs defined, a warning appears in the message box, and only the usable real CP 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**.

By default, General Purpose partitions will be positioned first, followed immediately by any (z/OS) associated zAAP or zIIP partitions (also applies to IFL partitions that are associated with a General Purpose partition running z/VM with Linux guests).

An alternate order can be chosen, such that all partitions in a single pool are grouped together, General Purpose 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 this report window, for any partition you can:

- **Exclude** a partition from being considered in the capacity assessment by un-checking its ☑ ***Include*** checkbox. zAAP, zIIP, and IFL partitions can be excluded without excluding their associated General Purpose partition. However, if a General Purpose partition is excluded, any associated zAAP, zIIP, of IFL partition will also be excluded.

- **Change Name**.

- **Modify SCP**. A zAAP or zIIP partition SCP or an IFL partition SCP that is associated with a General Purpose partition cannot be changed; it will always reflect the SCP of the associated General Purpose partition. If a General Purpose partition has an associated zAAP, zIIP, or IFL partition, its SCP cannot be changed to a different SCP (such as changing z/OS to z/VM).

- **Modify Workload**. A zAAP or zIIP partition workload cannot be changed; it will always reflect the workload of the associated General Purpose partition.

- **Modify Mode**. Partitions may be switch between dedicated and shared. However, changing a partition to dedicated will remove real CPs from the shared pool, which may cause some partition to become invalid (☑ ***Include*** will become unchecked).

- **Modify LCPs**. The LCP count allowed will be limited to the number of LCPs in shared real CP pool of the partition.

- **Modify Weight**. Changing a partition's weight values will rebalance the ***Minimum Capacity*** results between partitions within the CP pool.

- **Modify CAP**. Capping a partition will limit its ***Maximum Capacity*** to its ***Minimum Capacity***. There is no effect on the capacity of other partitions.

Any changes to partition metrics made on the ***Partition Detail Report*** window are reflected throughout **zPCR** for this LPAR configuration.

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.

zAAPs / zIIPs associated with a z/OS General Purpose Partition

These partitions must be initially defined from the **Partition Definition** window. They can only be associated with a General Purpose partition running z/OS-1.6 or later. Except for the number of logical CPs, the zAAP/zIIP partition is initialized with partition definition metrics identical to its associated General Purpose partition. Once defined, the number of **LCPs** specified for the zAAP/zIIP partition can be modified from this window. In addition, on z9 and later processors, the **Mode**, **Weight**, and **CAP** assignment can also be changed from this window (not needing to be tied to the settings of the associated General Purpose partition). The **Name**, **SCP**, and **Workload** can only be changed via the associated General Purpose partition.

To delete a zAAP or zIIP partition, set the LCP count to zero on the **Partition Definition** window that includes the General Purpose partition.

IFLs associated with a z/VM General Purpose partition considerations

These partition configurations are only supported on z10 and later processors. They must be initially defined from the **Partition Definition** window, and can only be associated with a General Purpose partition running z/VM. Except for the number of logical CPs, the IFL partition is initialized with partition definition metrics identical to its associated General Purpose partition. Once defined, the number of logical CPs specified for the IFL partition can be modified from this window. The **Mode**, **Weight**, and **CAP** assignment can also be changed from this window. The **Name**, **SCP**, and **Workload** can only be changed via the associated General Purpose partition.

To delete an IFL partition associated with a General Purpose partition, set the LCP count to zero on the **Partition Definition** window that includes the General Purpose partition.

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. Even though a partition is excluded, it remains in the study's partition list.

Table View Controls

- Under **Display zAAP/zIIP/IFL Partitions**

  ⊙ **With Associated GP** is the default order. This means that zAAP, zIIP, and IFL partitions that are owned by a General Purpose partition will be displayed immediately beneath the owning General Purpose partition. Partitions in the IFL and ICF pools will follow.

  ⊙ **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 Partitions** setting made from any open **Partition Detail Report** window will serve globally as the default order for every window that displays the individual partitions of an LPAR configuration. The global 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 global 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 Partitions).

Controlling the Order of Displayed Partitions

By default, General Purpose partitions are always displayed first, followed by zAAP, zIIP, IFL, and ICF partitions. zAAP and zIIP partitions (and General Purpose associated IFL partitions) are either displayed immediately following their owning General Purpose partition or displayed grouped by pool (the order is determined by the setting in the **Table View Controls** group box). IFL partitions followed by ICF partitions are displayed last.

Within each pool (General Purpose, IFL when not associated with a General Purpose, and ICF), partitions are displayed in the order in which the partitions were defined. zAAP and zIIP partitions (and IFL partitions when associated with a General Purpose partition) will appear in the same order as their owning General Purpose partition. General Purpose partition ownership of zAAP, zIIP, and IFL partitions can be determined by the matching partition number and name.

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 global default order is not affected). Whenever output is requested, the current view will be used.

With the exception of zAAP, zIIP, and associated IFL partitions, permanent changes to the sequence of partitions within a pool can be made 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.

Messages Concerning Generated Results

The following messages may be issued in the message area at the bottom of the **Partition Detail Report** window:

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 unreasonable for the number of CPs in the pool. In this case, capacity results are not 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 considered excessive for the number of CPs in the pool. Capacity results continue to be provided. See note below.

> **Note:** Often the number of logical CPs 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 logical CPs should be removed from all partitions in the study that are being run that way.**
>
> **Whenever IRD is active, logical CPs that are varied offline should be removed from all partitions in the study that are being run that way.**

Note: Capacity results for LPAR configurations with a large number of partitions have not been validated by benchmarks or customer experience.
Displayed when the number of defined/active partitions exceeds 30:

Note: xx.xx% of XX (pool) capacity cannot be used by partitions due to capping assignments.
Displayed when some of the CP pool capacity cannot be consumed due to capping assignments:

Note: One or more partitions weights indicate more capacity than can be provided with the LCPs defined.
Displayed when the number of logical CPs assigned to a partition is inadequate to support the assigned weight percent:

Note: N defined partitions are excluded from the capacity assessments.
Displayed when the ☑ **Include** checkbox for one or more partitions is unchecked:

Controls

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).

Click the Modify SCP/Workload button to test the effect on capacity when changing the SCP and or workloads assigned to partitions. See Modify Partitions 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 Assignment 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 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 description of this capability.

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.

Menu-bar

Graph

Distribution of capacity   (Pie charts by the various CP pools)
Capacity available   (Bar graphs, by the various CP pools)

Modify SCP/Workload

LPAR Configuration Capacity Planning

The **Modify Partitions** window is accessed from the **Partition Detail Report** window by clicking the **Modify SCP/Workload** button. 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 General Purpose z/OS partitions and any associated zAAP or zIIP logical CPs. 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 General Purpose 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 General Purpose partitions with z/OS assigned are affected by the *Modify Partitions* function (IFL and ICF partitions will not be affected). When a General Purpose 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 General Purpose partitions
2. Only General Purpose partitions that are currently included in the study
3. General Purpose 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 zEC12 and zBC12 processors, z/OS partitions must be specified with z/OS-1.10 or later. When capacity planning for z196, z114, or z10 processors, 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.

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.

LCP Assignment Alternatives

LPAR Configuration Capacity Planning

The **LCP Assignment 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 Assignment 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. Closing the **Partition Detail Report** window will also close the **LCP Assignment Alternatives** window, restoring all of the original partition LCP counts.

Four different LCP setting variations can be applied, temporarily replacing the current LPAR configuration's LCP settings:

1. **EDF/RMF Unparked**: If any of the original partitions were generated from EDF or RMF data, and include detail for 1 or more partitions that were running with HiperDispatch enabled, this column is enabled and the LCPs for those partitions are set to the unparked LCP count (total count minus parked count). The LCP count for all the other partitions remains at the original value. To update the **Partition Detail Report** window to reflect capacity for these values, click the ⎴Unparked⎴ button in the **LCP Setting Alternatives** group box.

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. The ***EDF/RMF Unparked*** count will only reflect information for the single measurement interval that was 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*** column should be used to enter and test results for these LCP counts.

2. ***Weight Based Moderate***: LCP counts are computed to reflect the relative weight assigned to the partition. When the weight percent indicates that the number of logical CPs should be less than 2.6, the exact number of logical CPs (rounded up to the nearest whole number) will be assigned. When the weight percent indicates that the number of logical CPs is greater than or equal to 2.6, the exact number of logical CPs plus 1 (rounded up to the nearest whole number) will be assigned To update the ***Partition Detail Report*** window to reflect capacity for these values, click the Moderate button in the ***LCP Setting Alternatives*** group box.

3. ***Weight Based Minimum***: LCP counts are computed to reflect the relative weight assigned to the partition. Fractional LCPs are then rounded up. To update the ***Partition Detail Report*** window to reflect capacity for these values, click the Minimum button in the ***LCP Setting Alternatives*** group box.

4. ***User Assigned***: LCP counts are originally assigned to each partition with the LCP counts when this window was entered. 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.

Click the **Commit LCP Changes** button to make the current LCP assignments become the LPAR configuration current study. A dialog will be offered allowing the current study to be saved under a new name. The ***LCP Assignment Alternatives*** window will be closed and the full function of the ***Partition Detail Report*** window will be restored, reflecting the last LCP assignment setting used. Note that, once a set of LCP counts has been committed, any original parked LCP counts that may have been read from EDF or RMF data will be discarded.

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.

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 associated zAAP or zIIP logical CPs are less than full utilization. The window can only be accessed if zAAP and/or zIIP logical CPs are defined in the configuration.

Only General Purpose, zAAP, and zIIP partitions are presented on this window, with zAAP and zIIP partitions displayed immediately following their associated General Purpose partition.

With the following two exceptions, columns displayed are the same as those shown on the **Partition Detail Report** window.

- The **zAAP/zIIP Utilization** column provides an input field (white background) for each zAAP and zIIP partition. These fields are used to set the zAAP or zIIP utilization to a value less than 100%. Double click on the field to enter an integer from 1 to 100 representing the utilization. Specifying a utilization less than 100% will result in increasing the associated General Purpose partition's capacity, reflecting a reduced amount of management and switching cost.

- zAAP and zIIP partition capacity available is not affected by the utilization specified; however, the capacity considered as being consumed will be reduced to the value shown in the **Capacity Consumed** column.

Explain zAAP/zIIP Loading

zPCR capacity results assume full utilization regardless of CP type. In cases where zAAP or zIIP utilization is actually less than 100%, a smaller **zAAP/zIIP Utilization** value can be specified. This has the effect of reducing management and switching cost, thus increasing the productive capacity of the associated GP partition.

The N-way of a GP+zAAP or GP+zIIP association is considered to be the sum of all the logical CPs involved plus some switching cost. Capacity values are derived for:

1. The purely GP partition (0% **zAAP/zIIP Utilization**; no switching cost).
2. The GP+zAAP/zIIP partition association (100% **zAAP/zIIP Utilization**; full switching cost).

Actual GP partition capacity is determined by interpolating between these 2 capacity values based on the **zAAP/zIIP Utilization** specified (value between 1% and 100%).

For cases where the actual zAAP/zIIP utilization may be unknown, either estimate it or leave it at 100%. Some guidelines for estimating **zAAP/zIIP Utilization** follow:

zPCR algorithms assume 50% offload from GP to 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 associated GP partition.

- When GP capacity is equal to or greater than the associated zAAP or zIIP capacity, one can assume the zAAP or zIIP logical CPs would be fully utilized and **zAAP/zIIP Utilization** should remain at 100%.

- When GP capacity is less than zAAP or zIIP capacity, it is likely that the zAAP or zIIP logical CPs would be less than fully utilized and **zAAP/zIIP Utilization** 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 Utilization** 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 General Purpose partition's **Minimum Capacity** values using the currently specified **zAAP/zIIP Utilization** values. **Capacity Consumed** is computed for both the zAAP and zIIP CP pools, also based on the currently specified **zAAP/zIIP Utilization** values. .This information can be useful for the purpose of increasing or decreasing zAAP and zIIP capacity by altering the number of real CPs available. Note that such modifications must be made back on the **LPAR Host** window.

Click the **Commit Changes** button for the current **zAAP/zIIP Utilization** percentage settings to take effect for the study.

Click the **Restore zPCR Defaults** button to restore all of the **zAAP/zIIP Utilization** percentages back to 100%.

Click the **Restore Previous** button to set all of the **zAAP/zIIP Utilization** percentages to the values that were presented when the window was opened.

Click the **Explain** button for a short dialog discussing the benefit of providing **zAAP/zIIP Utilization** percentages less than 100%.

Click the **Return** tool bar icon to close the ***zAAP/zIIP Loading*** window and return to the ***Partition Detail Report*** window's. If any zAAP/zIIP loading percentage changes are pending, you will be prompted via a dialog to either ⬚ **Commit Changes** ⬚ or ⬚ **Restore Defaults** ⬚.

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

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

If any user specified ***zAAP/zIIP Utilization*** percentages (other than the default 100%) are committed to the configuration, they will be carried forward in a saved study file.

In ***Advanced-Mode***, cloned configurations will include the ***zAAP/zIIP Utilization*** percentages of the original configuration. When comparing LPAR configurations, matching partitions (Name, SCP, Workload, and DED/SHR), should have matching ***zAAP/zIIP Utilization*** percentages for the comparisons to be completely fair. When making such comparisons, if ***zAAP/zIIP Utilization*** percentages of matching partitions differ, a warning dialog will be presented.

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. From either the **Host Summary Report** window or the **Partition Detail Report** widow, click the Calibrate Reference-CPU button to open the **Calibrate** window.

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 *Reference-CPU* settings 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 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 vs Weighted Capacity
Based on LSPR Data for IBM System z Processors
Study ID: Not specified
Description: Created from RMF e:\...RMFSample z10 E26.txt interval #9
z10 EC/700 Host = 2097-E26/700 with 12 CPs: GP=7 zAAP=4 ICF=1
12 Active Partitions: GP=6 zAAP=4 ICF=2
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 | | |
| | | Partition | | | | | | Minimum | Pool | | | |
No.	Type	Name	SCP	Workload	Mode	LCPs	Weight%	Capacity	Capacity	Busy	Capacity	Ratio
1	GP	SYSB	z/OS-1.13	Average	SHR	3	36.50%	1,809	4,975	36.91%	1,836	1.015
	zAAP	SYSB	z/OS-1.13	Average	SHR	2	41.07%	1,187	2,817	19.32%	544	0.459
2	GP	TESTCICS	z/OS-1.13	Average	SHR	2	2.00%	97	4,975	1.27%	63	0.650
	zAAP	TESTCICS	z/OS-1.13	Average	SHR	2	2.05%	60	2,817	0.05%	1	0.023
3	GP	PROD1	z/OS-1.13	Average	SHR	5	12.20%	597	4,975	12.50%	622	1.041
	zAAP	PROD1	z/OS-1.13	Average	SHR	3	10.78%	306	2,817	3.44%	97	0.317
4	GP	TEST1	z/OS-1.13	Average	SHR	1	2.00%	108	4,975	0.85%	42	0.393
5	GP	TEST2	z/OS-1.13	Average	SHR	1	2.00%	108	4,975	0.75%	37	0.347
6	GP	PROD2	z/OS-1.13	Average	SHR	7	45.30%	2,256	4,975	46.95%	2,336	1.035
	zAAP	PROD2	z/OS-1.13	Average	SHR	2	46.10%	1,264	2,817	2.64%	74	0.059
7	ICF	CF0001	CFCC	CFCC	SHR	1	50.00%	429	858	0.17%	1	0.003
8	ICF	CF0002	CFCC	CFCC	SHR	1	50.00%	429	858	0.08%	1	0.002

Summary by Pool

| | Number of | | | Pool | | Utilized | |
CP Pool	RCPs	Partitions	LCPs	Capacity	Busy	Capacity	Ratio
GP	7	6	19	4,975	99.23%	4,937	0.992
zAAP	4	4	9	2,817	25.45%	717	0.254
zIIP	0	0	0	0	0.00%	0	0.000
IFL	0	0	0	0	0.00%	0	0.000
ICF	1	2	2	858	0.25%	2	0.002
Totals	12	12	30	8,650		5,656	0.654

For significant configuration changes, capacity comparisons should be considered to have a +/-5% margin-of-error
Upgrading the processor family is considered a significant configuration change

This report is only enabled in cases where the LPAR host and the entire partition configuration was loaded via EDF or RMF data, since these are the only cases where 100% of the utilization is accounted for.

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 a description of the LPAR host and its CP configuration, followed by the number of active partitions with their distribution across real CP 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 General Purpose 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**.

General Purpose partition logical CPs always operate from a single pool of real CPs. On z9 and later processor models, zAAP partitions, zIIP partitions, IFL partitions, and ICF partitions each operate from a separate pool of real CPs. On z990 and z890 models, logical CPs defined for these specialty engines operate from a single pool of real CPs.

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 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.

Controls

No change capability is provided on the *Utilized Capacity Report* window itself. Changes can only be made via prior windows.

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.

Capacity results on this report only make sense when the entire LPAR configuration is being considered. Therefore, partitions that may have been excluded on the *Partition Detail Report* window will still be shown on this report.

If modifications are made to the LPAR host processor or to any individual partition, the report can no longer be considered an accurate representation. Once any such modification is made, the **Partition Utilized Capacity** button will become inactive and the report will be unavailable.

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. (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.

Advanced-Mode

On the **Function Selection** window (**Multi-Image Capacity** tab), when
☑ **Advanced-Mode** is checked, clicking the ⎡**Enter Advanced-Mode**⎤ button initiates
the **LPAR Configuration Capacity Planning** function in advanced-mode. In advanced-
mode, the **Advanced-Mode Control Panel** widow provides access to all **zPCR** function,
adding the capability to define and compare multiple LPAR configurations.
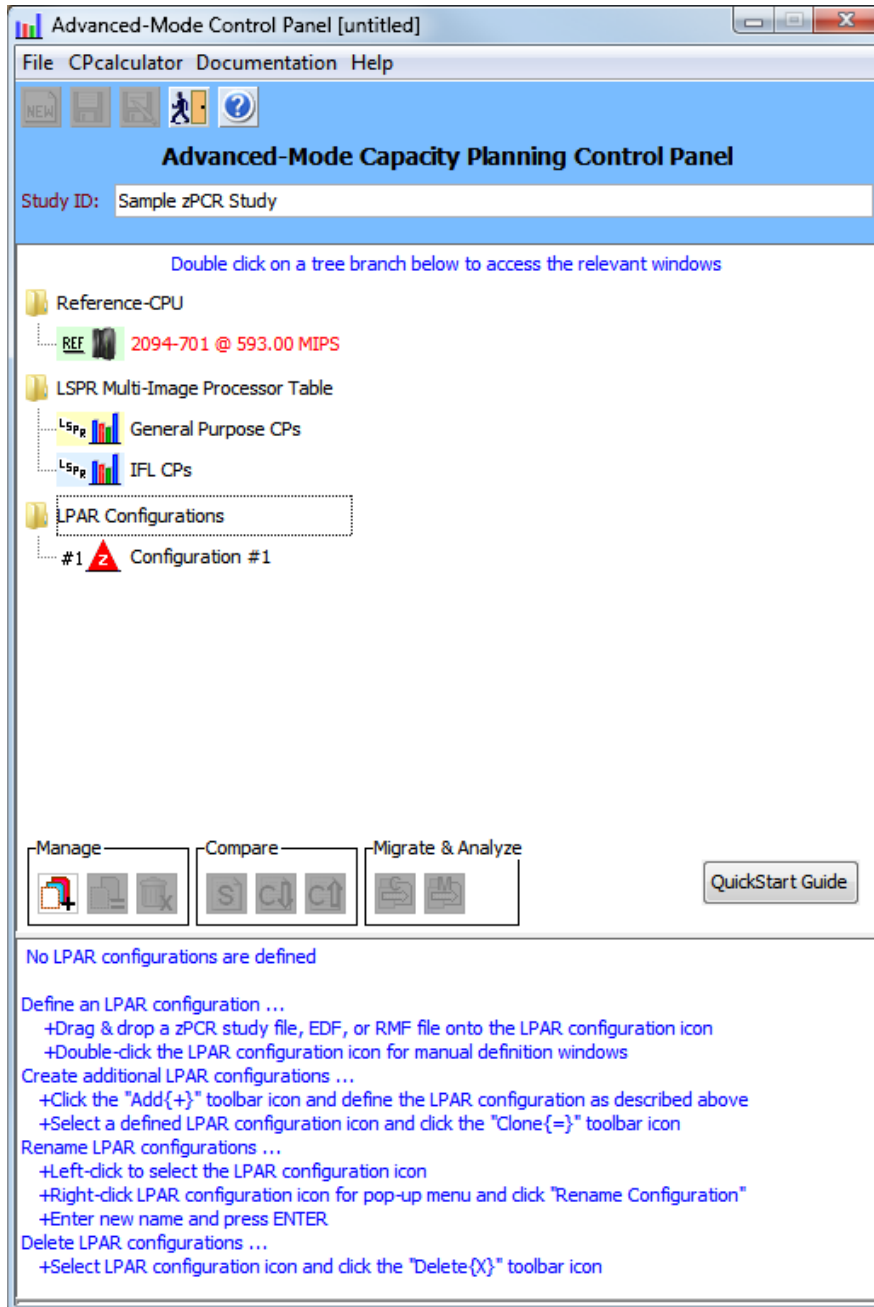
The primary advantage of **zPCR**'s *Advanced-Mode* lies in the ability to make direct comparisons of capacity results between LPAR configurations. For capacity comparisons to be useful, the partition configurations being compared should both contain some or all of the same partitions (i.e., in terms 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. Changing the LPAR host processor family and/or model
2. Changing the LPAR host processor's CP configuration
3. Changing the way that one or more partitions are defined
4. Adding one or more new partitions
5. Deleting one or more current partitions.

Note that, once a study is initiated in *Advanced-Mode*, it cannot revert to a basic-mode study. All control affecting the LPAR configuration(s) must be accessed via the *Advanced-Mode Control Panel* widow. Saved study files will be in the *Advanced-Mode* format.

The **QuickStart Guide** button reveals a short paper describing the process using **zPCR** in *Advanced-Mode* to define a current and an alternate LPAR configuration and make capacity comparisons between them.

Summary of Advanced-Mode function

☑ **Advanced-Mode** provides the following **zPCR** function:

1. The *Reference-CPU* metrics can be set.

2. Of the LSPR tables, only the *Multi-Image LSPR Capacity Ratios* table (General Purpose CPs or IFLs) is available. This is the preferred table with which to generalize capacity. (Note: To view the *Single-Image LSPR Capacity Ratios* tables, invoke a 2nd copy of **zPCR** and access the table from the *Function Selection* window.)

3. Multiple LPAR configurations (currently limited to six) with names such as *Configuration #1* and *Configuration #2* can be defined (these can be renamed, if desired). Several report windows are available with which to compare capacity results between various configurations.

4. The *Reference-CPU* scaling-factor can be calibrated for the 1st LPAR configuration only, for the purpose of producing a desired capacity result for that LPAR configuration. Once calibration has been done, the resulting *Reference-CPU* scaling-factor will be use for all the LPAR configurations that are defined.

5. The *Workloads* window must be accessed from the *LSPR Processor Table*, since the *Function Selection* window is no longer accessible in *Advanced-Mode*.

When in ☑ **Advanced-Mode**, the following **zPCR** function is not available:

1. The *LSPR Single-Image Processor Capacity Ratios* tables cannot be viewed (the *LSPR Multi-Image Processor Capacity Ratios* table is preferred for generalizing on partitioned capacity). To subsequently access the *LSPR Single-Image Processor Capacity Ratios* tables, start a second **zPCR** invocation in basic-mode (be sure the *Reference-CPU* settings are as desired).

2. Basic-mode study files cannot be created (studies will be saved in *Advanced-Mode* format), and will include all of the LPAR configurations that are defined.

The *Advanced-Mode 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.

The only *Reference-CPU* modification allowed from windows other than the *Advanced-Mode 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*).

In *Advanced-Mode*, several additional windows are available, with which to compare capacity results between LPAR configurations (for details, 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 *Advanced-Mode Control Panel* window will be locked while comparison windows are open.

The *Study Identification* text can be entered from either the *Advanced-Mode Control Panel* window or the *Function Selection* 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 *Advanced-Mode 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 either the General Purpose CPs or the IFLs *LSPR Multi-Image Processor Capacity Ratios* table. These are the preferred tables for generalizing on partitioned capacity.

- **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.

#1 ⚠z	Configuration #1 XYZ Production on z196 z196/700 LPAR Host: 2817-M15/700					
Pool CP Type	#1 GP	#2 zAAP	#3 zIIP	#4 IFL	#5 ICF	CPC Total
RCPs	7	1	1	4	1	14
Partitions	3	1	1	3	1	9
LCPs	14	1	1	7	1	24
Capacity	7,359	1,112	1,081	4,390	1,147	15,089
Capacity basis: 2094-701 @ 593.00 MIPS for a shared single-partition configuration						

Example of LPAR Configuration Summary

When the **Advanced-Mode Control Panel** 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.

2. The LPAR configuration may be defined by dragging a previous **zPCR** study file on top of the LPAR configuration branch icon.

   a. If a **Basic-Mode** study is used, its LPAR configuration is loaded.

   b. If an **Advanced-Mode** study is used, 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.

3. 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.

4. 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 an Advanced-Mode 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 name **Configuration #x**. 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.

X **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. When only two LPAR configuration icons exist, and the first (**Configuration #1**) is deleted, the second will become the first, and be renamed to **Configuration #1**. Note that when the last LPAR configuration icon is deleted, the first icon (**Configuration #1**) will reappear, but with no LPAR configuration defined. This state is equivalent to a new **Advanced-Mode** study.

Renaming LPAR Configurations

A default name is assigned to each LPAR configuration icon when it is created. Each icon can be renamed with up to 20 characters. To rename:

1. Left-click the LPAR configuration icon to select it
2. Right-click the LPAR configuration to generate a pop-up menu
3. Click **Rename Configuration** on the pop-up menu.

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 details 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 details 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 *Migrate & Analyze* toolbar icon:

➔ **Migrate** toolbar icon: Provides the ability to modify an LPAR configuration and assess the overall effect on capacity. Select a single defined LPAR configuration and click **Migrate** to make it the target for changes. Partitions can be copied to the target from any other defined LPAR configuration and 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 details concerning this capability, see Copy Partitions into LPAR configuration.

Saving Studies in Advanced-Mode

When in *Advanced-Mode*, saving a study from the *Advanced-Mode Control Panel* window will result in an *Advanced-Mode* study file format. When subsequently loaded, **zPCR** is automatically placed in ☑ **Advanced-Mode**, with the LPAR configuration(s) loaded.

In ☑ **Advanced-Mode**, save is no longer available from the *LPAR Host and Partition Configuration* window.

Menu-bar

File

New	Start a new study (untitled)
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 [1]
Exit (*Ctrl+E*)	Terminate **zPCR** execution. Exit can also be invoked using the icon on the tool bar.
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**
> **Reference-CPU Notice**

Help

Context Help (F1)	
QuickStart Guide	Provides guidance on using **zPCR**
Check for updates	Access your CPS tools download site
About zPCR	Displays information from the *Logo* window including the version of **zPCR**.

---

[1] In *Advanced-Mode*, *Save as…* will create an advanced-mode study.

Defining and Managing LPAR Configurations in Advanced-Mode

On the **Advanced-Mode Control Panel** window, for each LPAR configuration defined, a unique icon and name is displayed. The first icon will be named **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 a basic-mode study, the study's file identification will be placed in the description field as a reminder of its origin. When a configuration is loaded from RMF, the name of the RMF file will be placed in the description filed. Each LPAR configuration's description field can be altered at any time to the text desired.

In **Advanced-Mode**, the **Utilized Capacity Report** window can be displayed for any LPAR configuration when the entire configuration was obtained from 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, a 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.

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.

When the **Exit** toolbar icon on the **Advance-Mode Control Panel** window is clicked, all open windows, including the **LPAR Configuration Capacity Planning** windows, will automatically be closed.

Host Capacity Comparison Summary

Advanced-Mode: Capacity Comparison Reports

Once any LPAR configurations have been defined, the **S** **Summary** toolbar icon on the **Advanced-Mode 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 **Advanced-Mode Control Panel** window remains locked. Also, any open **LPAR Configuration Capacity Planning** windows are hidden to prevent LPAR configuration changes being made while the summary window is open.

This report presents the capacity projections for each defined LPAR configuration. For each configuration, its icon and name are provided, along with the processor model information and the number of real CPs configured in each pool. To display the description field of any LPAR configuration, 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 in the LPAR configuration. A total capacity value is also provided on the right as the sum of the individual CP pool capacity values.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for details, 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.

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

Formatted Output

Output for the **Host Capacity Comparison Summary** window can be obtained in HTML format by clicking the **Output to HTML file** toolbar icon.

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 **Advanced-Mode Control Panel** window. Any open **LPAR Configuration Capacity Planning** windows will be unhidden.

Host Capacity Comparison

Advanced-Mode: Capacity Comparison Reports

The **Host Capacity Comparison** window is accessed by clicking either the
C↓ Compare Down or the **C↑ Compare Up** toolbar icon on the **Advanced-Mode Control Panel** window. At least two LPAR configurations must be defined. If only a single LPAR configuration is selected, it is used as the basis for comparison. The **C↓ Compare Down** toolbar icon uses the LPAR configuration branch immediately below as the focus of the comparison. The **C↑ Compare Up** toolbar icon uses the LPAR configuration branch immediately above as the focus of the comparison. When exactly two LPAR configurations are selected, the arrow on the **C↓** or **C↑ Compare** toolbar icon indicates the order in which the 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.

While the **Host Capacity Comparison** window is being displayed, the **Advanced-Mode 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 summary 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 host configurations as a whole.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for details, 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 details, 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 details, 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 *Advanced-Mode Control Panel* window. Any open *LPAR Configuration Capacity Planning* windows will be unhidden.

Partition Capacity Comparison

Advanced-Mode: Capacity Comparison Reports

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 ***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.

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 associated General Purpose partition.
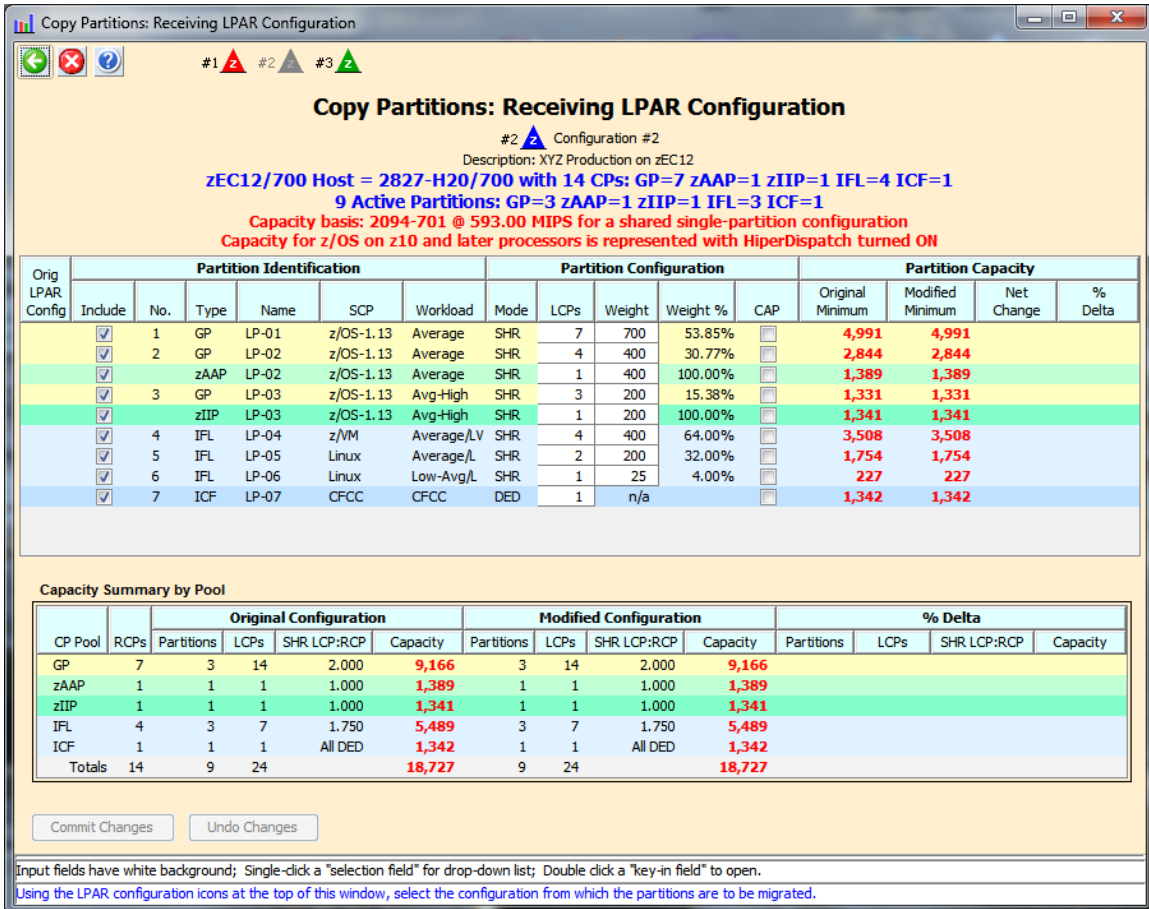
While displayed, the **Partition Capacity Comparison** window provides the capability to modify several of the partition defining metrics for the 2nd LPAR configuration. These input fields, shown with a white background, include number of LCPs, weight, and capping. Changes to these metrics can be made to quickly asses 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 details, see <u>zAAP/zIIP Loading</u>). 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 provides assistance in streamlining the LCP counts for the LPAR configuration to achieve more efficient use of the available RCPs.
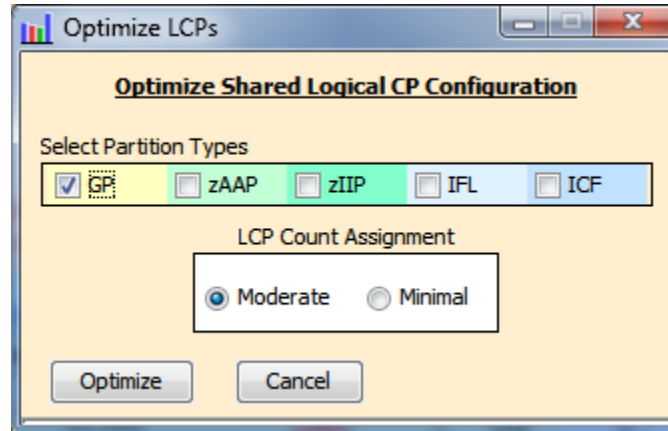
Any or all of the partition types may be selected for optimization, using the checkboxes provided in the **Select Partition Types** group box. By default, only the General Purpose partition type is checked.

The level of optimization for **LCP Count Assignment** can be chosen as:

1. **Moderate**: When the weight percent indicates that the number of logical CPs should be less than 2.6, the exact number of logical CPs (rounded up to the nearest whole number) will be assigned. When the weight percent indicates that the number of logical CPs is greater than or equal to 2.6, the exact number of logical CPs plus 1 (rounded up to the nearest whole number) will be assigned.

2. **Minimum**:   The weight percent is used to determine the exact number of logical CPs (rounded up to the nearest whole number) to be assigned.

The resulting logical CP assignments are based strictly on the relative weight of the partition. The number of logical CPs 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.

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 logical CP Assignments

HiperDispatch: On zEC12, zBC12, z196, z114, and z10 processors running z/OS with HiperDispatch enabled, excessively configured logical CPs will tend to be parked, based on the overall demand for capacity. Parked logical CPs should not be included in an LPAR configuration, since only active logical CPs contribute to cost of partitioning.

IRD: On System z processors running with the **Intelligent Resource Director** (**IRD**), shared partition logical CPs that are varied offline should not be included in an LPAR configuration, since only the active logical CPs 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 details, 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 **Advanced-Mode Control Panel** window. Any open **LPAR Configuration Capacity Planning** windows will be unhidden.

Host Margin-of-Error

Advanced-Mode: Capacity Comparison Reports

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 ±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 *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 details, 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.

Partition Margin-of-Error

Advanced-Mode: Capacity Comparison Reports

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 ±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. Not 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 2[nd] 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 associated General Purpose partition.

Note that the capacity projection for the 2[nd] LPAR configuration that includes the margin-of-error consideration is shown for all defined partitions regardless of whether the partition was defined on the 1[st] LPAR configuration.

LPAR configurations where zAAP/zIIP loading values less than the default 100% have been specified will be flagged (for details, 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.

Copy Partitions to an LPAR Configuration

Advanced-Mode: Migrate & Analyze

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 $\boxed{C→}$ 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.

This window is similar to the *Partition Detail Report* window, except for the *Partition Capacity* columns which here, always relate to 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.

This window is similar to the *Partition Detail Report* window, except for the *Capacity* column which, here, reveals 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 copied into the *Receiving LPAR Configuration*. Partitions that have already been copied to the *Receiving LPAR Configuration* cannot be reselected.

General Purpose partitions with associated zAAP or zIIP LCPs will always be copied together (if the zAAP or zIIP partition is not desired in the **Receiving Configuration**, its ☑ **Include** can be unchecked). Note that zAAP, zIIP, and IFL (when associated with a General Purpose partition) partitions cannot be copied without the associated General Purpose 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** 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 copied 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 copied, 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.

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 focus window.

Move Partitions to an LPAR Configuration

Advanced-Mode: Migrate & Analyze

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 **M➡** 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.

This window is similar to the *Partition Detail Report* window, except for the *Partition Capacity* columns which, here, always relate to 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 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) 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.

This window is similar to 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**. Partitions that are moved to the **Receiving LPAR Configuration** are removed from the **Contributing LPAR Configuration**. In addition, all capacity values for the **Contributing LPAR Configuration** are recomputed. **Net Change** and **% Delta** metrics are computed expressing the difference between the **Modified Minimum** capacity column and the **Original Minimum** capacity column.

General Purpose partitions with associated zAAP or zIIP LCPs will always be moved together (if the zAAP or zIIP partition is not desired in the **Receiving Configuration**, its ☑ **Include** can be unchecked). Note that zAAP, zIIP, and IFL (when associated with a

General Purpose partition) partitions cannot be moved without the associated General Purpose partition.

Click the **Unselect All** button to clear the ☑ **Move 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 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.

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. The **Modified Capacity** values are computed for both new partition configurations and compared back to the **Original Capacity** values as revealed in the **Net Change** and **% Delta** columns.

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 focus window.

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 z/OS, when CPU-MF (SMF type 113) data is available, an algorithm automatically chooses the appropriate LSPR workload category. For the remaining z/OS environments, rules-of-thumb have been implemented to assist in this selection. For other operating systems, some judgment will be required. Even if the workload selections are totally appropriate, **when comparing the capacity of one processor's LPAR configuration to another, zPCR results should be considered with up to ± 5% margin-of-error**. This is particularly true when the LPAR host processor's family is 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 **± 5%**.

Because there is a potential ±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 CPs 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 associated z/OS partition. Since **zPCR** results represent the *Capacity Perspective*, consider, that, when a single zAAP logical CP is associated with a single z/OS General Purpose logical CP, 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 General Purpose capacity based on expected zAAP/zIIP utilization between the two extremes projected by **zPCR**.

  For zEC12, zBC12, z196, z114, z10-EC, and z9-EC processors, when large z/OS partitions (many logical CPs) 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 logical CPs. This definition is made in terms of the minimum capacity of the associated General Purpose partition. By default, **zPCR** assumes the zAAP or zIIP capacity requirement to be equal to that of the associated General Purpose partition (i.e., 100% zAAP/zIIP loading). By specifying zAAP/zIIP loading less than 100%, General Purpose capacity will increase slightly due to reduced switching cost. zAAP and zIIP capacity is not affected.

- **z/VM and Linux on General Purpose CPs or IFLs**: Only one LSPR workload is run for each of these environments. z/VM has been measured on models up to the 32-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 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 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 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 General Purpose partition can have associated zAAP, zIIP, IFL, and ICF CPs. **zPCR** currently supports only the GP/IFL combination. The partition number and name of the associated IFLs will be the same as the General Purpose partition. The workload assignment for the associated IFLs may be different than that for the General Purpose partition. No attempt is made to factor in a cost for z/VM to manage the IFLs. Therefore, the capacity result generated for the IFLs will be identical to that if the IFLs were a partition independent of the z/VM General Purpose partition.

- **z/VSE on General Purpose CPs**: z/VSE environments are not measured for LSPR purposes. As a convenience, three z/VSE environments have been defined in **zPCR**, representing *Batch*, *Online*, and *Mixed*. To develop capacity relationships between 1-way processors, similar z/OS workload environments were used. For capacity relationships of 2, 3, and 4-ways, N-way relationships, considered representative of z/VSE's turbo dispatcher, were used. Capacity results for z/VSE should only be considered as **reasonable approximations of realizable capacity**.

- **CFCC on General Purpose CPs or IFLs**: CFCC capacity data, carried internally, is used when configuring partitions as coupling facilities. Generally, this CFCC data is representative of the version that was available at a processor announce, and, therefore may not be 100% 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 shown 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 *Function Selection* window.

The sole 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 General Purpose CP environment to one with associated zAAP CPs.

Note: Additional CPcalculator functions may be added as time and resources permit.

While *CPcalculator* functions are started from the *Function Selection* window, they can be used in parallel with the basic **zPCR** function. This means that one could access the *zAAP Capacity Estimator* function while using the **LPAR Configuration Capacity Planning** 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 each 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 *Function Selection* 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

zAAPs. Capacity comparisons should be considered as being most appropriate for the most recent LSPR measurement data (z/OS-1.13).

Capacity values presented are based on **zPCR**'s currently defined *Reference-CPU* and its scaling. 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 assume a single z/OS partition and its associated zAAP logical CPs, without regard to any other partitions that may be active. To evaluate the effects on capacity when defining multiple partitions sharing the same real CP 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's identification 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:

- **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.
- ⬍ **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.

Previous Configuration: General Purpose CPs Only

For a z/OS image which is being considered to start exploiting a zAAPs, its capacity requirement must be characterized in terms of the previous General Purpose processor. You must provide a processor model, 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 work 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 System z families are available for 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.

- ▼ **General Purpose CPs** -   This field is initially set to 1. Click on the dropdown icon to change the number of General Purpose CPs active for this z/OS image. The value can be in the range of 1 to 100 (what z/OS-1.13 and later can support); only numbers valid for the selected processor model are shown in the list. Note that, for z114, z10-BC, z9-BC, and z890 processors, the number of General Purpose CPs is implied by the processor model, and cannot be changed.

  The *Processor Model* and *General Purpose CPs* fields are redisplayed in a table, showing the processor name as identified by z/OS, feature, and MSU rating appropriate for the z/OS image.

- **⬍ 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%. This value represents the percent of the processor's capacity required to manage the **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). Use 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). Use 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 workload's utilization. The percent of Java content determines the amount of the total workload that will incur zAAP processing costs. The percent of zAAP eligible work determines that amount of zAAP capacity that might be used.

Planned Configuration: General Purpose CP + zAAPs

The intended new General Purpose plus zAAP processing environment is described, using the following fields:

- **▼ 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. Only families that can support zAAPs are available for selection (zEC12, zBC12, z196, z114, z10, z9, z990, and z890).

- **▼ 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.

- **▼ General Purpose CPs** - This field is initially set to 1. Click on the dropdown icon to change the number of General Purpose CPs active for this z/OS image. The value must be in the range of 1 to ½ the maximum number of CPs available (the number of zAAPs cannot exceed the number of General Purpose CPs defined); only numbers valid for the selected processor model are shown in the list. . Note that, for z114, z10-BC, z9-BC, and z890 processors, the number of General Purpose CPs is implied by the processor model, and cannot be changed.

  The *Processor Model* and *General Purpose CPs* fields are redisplayed in a table to show the processor name as identified by z/OS, feature, and MSU rating appropriate for the z/OS image. If migrating to a zAAP environment on the same hardware, and the number of General Purpose CPs can be reduced, the MSU rating change will be of particular interest, being an indicator of potential software cost savings.

⌄ **zAAPs** - This field is initially set to 1. Click on the dropdown icon to change the number of General Purpose CPs active for this z/OS image. The value must be in the range of 1 to *n*, where *n* cannot exceed the number of General Purpose CPs; only numbers valid for the selected processor model are shown in the list. Note that capacity results for configurations defined with a large number of zAAP (or zIIP) logical CPs 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** **DP: General Purpose 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** - 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 zAAPs 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 logical CPs. Click the ☑ **Set Default** checkbox to restore the default setting of *Low-Avg*.

Available capacity is determined by the General Purpose processor model and the number of zAAPs assigned. The combination of General Purpose CPs and zAAPs is viewed as a single z/OS image for capacity purposes. Nominal overheads associated with zAAP dispatching decisions are 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 logical CPs 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 logical CPs and for the associated General Purpose 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 zAAPs were defined.

zAAP capacity is always associated with that of Java content workload.

General Purpose CP capacity is a value interpolated between that of the normal content work and that for Java content work, based on the percent of represented by each. As the amount of normal content workload diminishes in favor of Java content workload, the General Purpose CP capacity value will move towards the capacity assumed for purely Java content work.

The resulting capacity requirement is then allocated to the zAAPs where needed, and then to the General Purpose CPs.

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 ⬚ **Reference-CPU** to open the *Reference-CPU* window, allowing its settings to be changed. Any changes made here will also affect **zPCR**'s *Single Image Capacity Ratios* tables and the *LPAR Configuration Capacity Planning* function.

Click the ⬚ **Report** to open the *zAAP Capacity Estimator Report* window. For details 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 details concerning inputs, see zAAP Capacity Estimator Input.

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.

Previous Configuration Group Box

A summary of the workload as defined to the all General Purpose CP 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 new 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 zAAPs, capacity for a zAAP configuration must be viewed as two distinct entities, one for the zAAPs 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, a warning is issued when their respective SDP settings are exceeded.

Capacity results for both the current and planned configurations are drawn directly from **zPCR**'s **LPAR Configuration Capacity Planning** function. Normally we would view General Purpose capacity as somewhat dependent on the zAAP utilization. For example, if the zAAP were not utilized at all, capacity would simply be that of the General Purpose logical CPs alone. The **zAAP Capacity Estimator** looks at the zAAP utilization result, and interpolates for the General Purpose logical CP capacity between that with no zAAPs defined, and that with the zAAPs defined. As a result the General Purpose utilization will also be affected.

Keep in mind that this function considers the logical CPs for only a single GP/zAAP pairing. If in reality, multiple partitions would compete for the same real CP resources, capacity results will be different.

You can capture the report 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. Click the ⬚ **Capacity** ⬚ button for a bar graph showing the previous General Purpose capacity and the new 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. Click on ⬚ **Utilization** ⬚ for a line graph depicting utilization of the General Purpose CPs as utilization of the zAAPs 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 logical CPs, without regard to any other partitions that may be active. To evaluate the effects on capacity when defining multiple partitions sharing the same real CP resources, you must use the *LPAR Configuration Capacity Planning* function of **zPCR**.

2. The planned configuration *Capacity Available* result for General Purpose CPs is based on fully utilized zAAP CPs (the zAAP projected utilization is not considered in the model as currently implemented). In reality, as the utilization of the zAAP CPs diminishes below 100%, the General Purpose costs associated with managing them would also diminish. If the zAAP utilization were 0%, the General Purpose available capacity would be somewhat greater, aligning more closely with that of the General Purpose CPs with no zAAPs configured. For configurations where the General Purpose CPs are less than full speed, the difference in zAAP CP speed must be considered when adjusting General Purpose CP capacity available results.

Controls

Click the ⬚ **Return** ⬚ button to close the report window.

Click the ⬚ **Exit** ⬚ button on the *zAAP Capacity Estimator Input* window to terminate this *CPcalculator* function.

Charts and Graphs

Throughout the various **zPCR** functions, a variety of pie charts and bar/line graphs are available. Each of these is presented in a common *Graph* window.
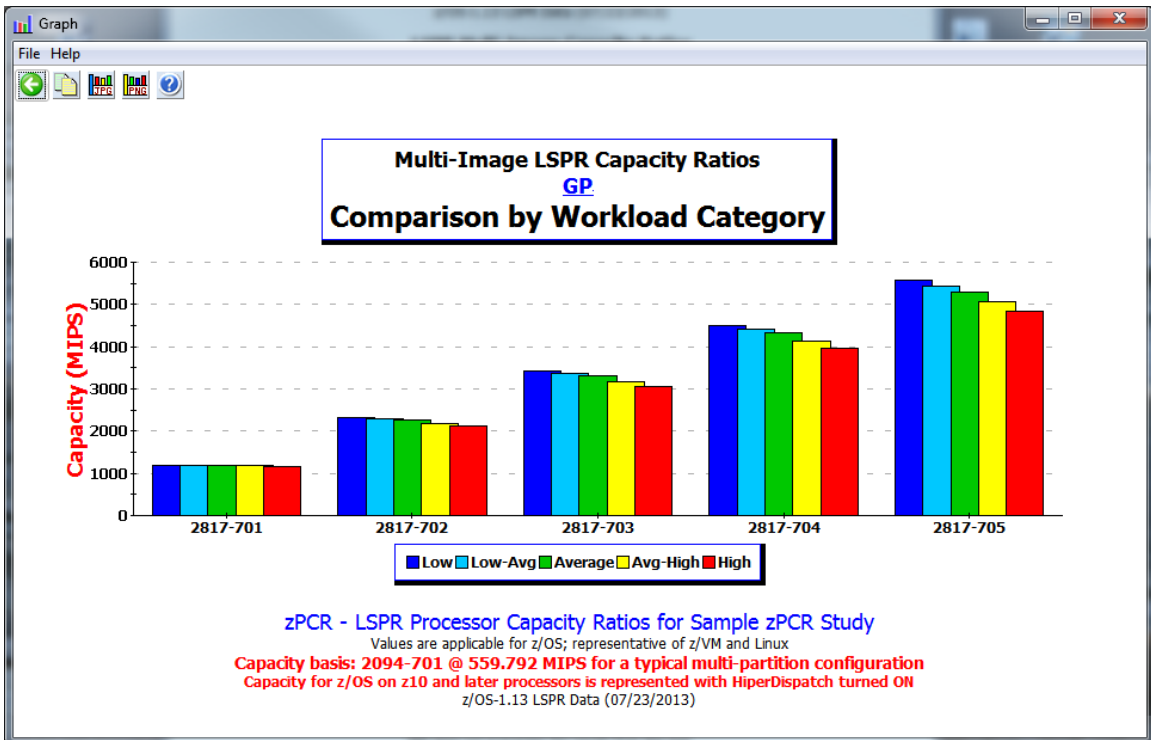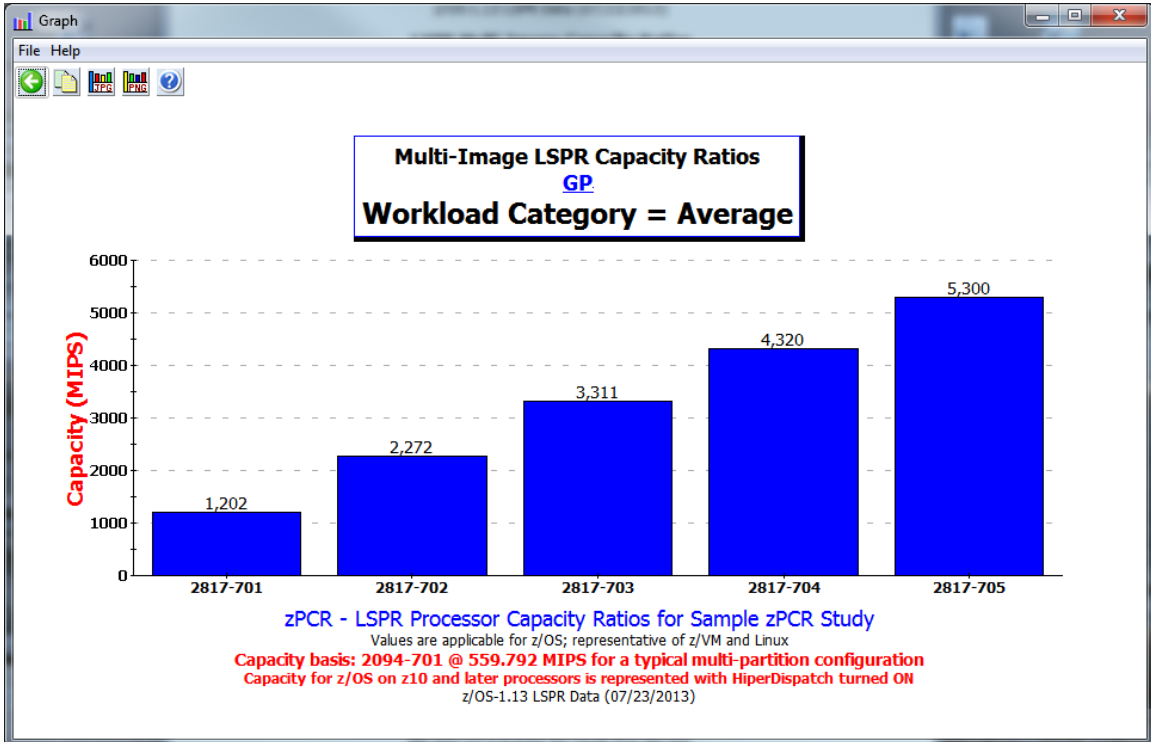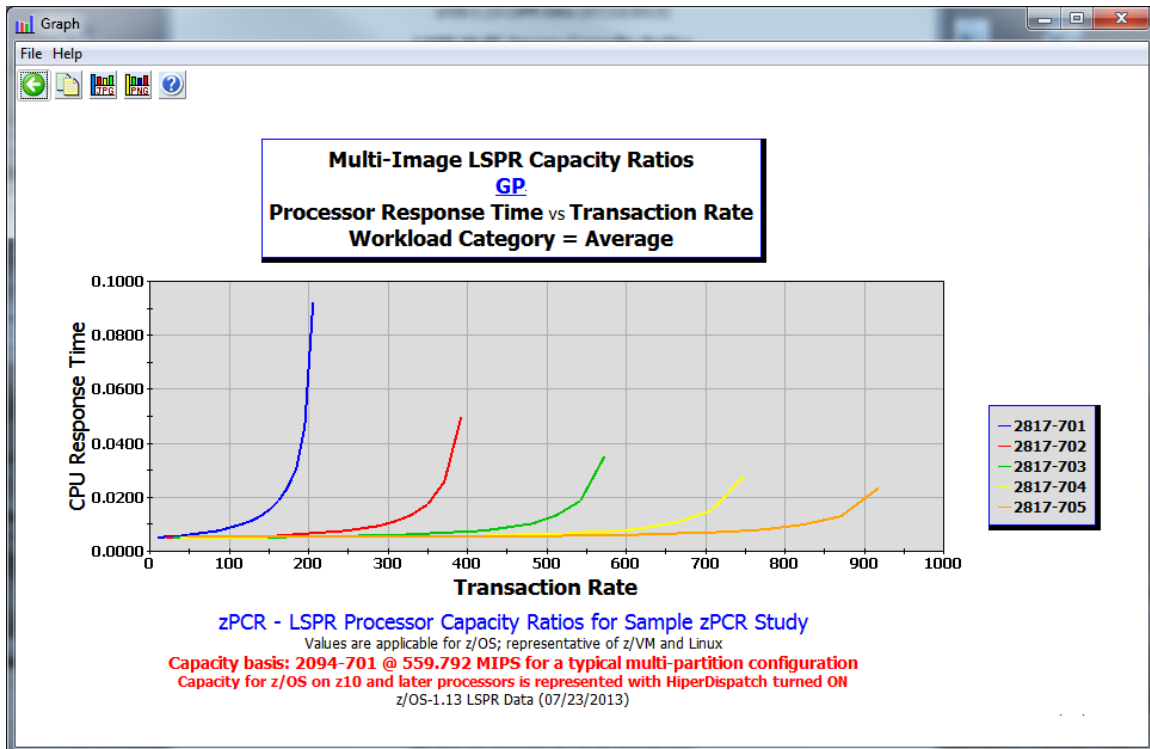
Graphs are available in the following functions as follows:

- **LSPR Multi-Image Capacity Ratios** table
  - Bar graphs showing processor capacity for z/OS workloads
  - Bar graph showing processor capacity for all z/OS workloads
  - Line graphs showing CPU response time for z/OS workloads

- **LSPR Single-Image Capacity Ratios** tables
  - Bar graphs showing processor capacity by SCP/workload
  - Bar graph showing processor capacity for all workloads of a SCP
  - Line graphs showing CPU response time by SCP/workload

- **LPAR Configuration Capacity Planning** function
  - Pie charts showing distribution of capacity by CP pool
  - 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 **Function Selection** 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 a number of PC applications.

- **PNG**   Save the chart as a named *PNG* file; this file can be viewed with a number of 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.

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 for IBM System z)." 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. You can check the Web site from which you obtained **zPCR** to determine if the version of the tool has changed. From the menu bar on the *Function Selection* window or the *Advanced-Mode Control Panel* window, click **Help** ➔ **Check for updates**.

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*. 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 Vintage Terminology

System z   Includes the following processor models:
          IBM zEnterprise EC12 (2827)
          IBM zEnterprise 196 (2817)
          IBM zEnterprise 114 (2818)
          IBM System z10-BC (2098) and z10-EC (2097)
          IBM System z9-BC (2096) and z9-EC (2094)
          IBM eServer z890 (2086) and z990 (2084)
          IBM eServer z800 (2066) and z900 (2064)

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 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 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**.

Mainframe   All of the above processors considered as a group

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, Linux, zAware, or CFCC.

GP   An abbreviation used for "General Purpose", used in conjunction with the acronyms zAAP, IFL, and ICF when referring to real CPs, logical CPs, or partitions.

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.

ICF Pool   All of the CPs in the ICF pool. For processor models subsequent to the z990 and z890, zAAP, IFL, and ICF CPs each operate as a separate pool.

IFL   Integrated Facility for Linux (a **Special Purpose CP**). Each IFL is a real CP on the CPC. An IFL can only run native Linux or z/VM with Linux guests.

IFL Pool   All of the CPs in the IFL pool. For processor models subsequent to the z990 and z890, zAAP, IFL, and ICF CPs each operate as a separate pool.

LCP   Logical CP (engines defined to a logical partition)

LCP:RCP Ratio   Ratio of all the partitions' Logical CPs to all of the processor's Real CPs (this ratio is only meaningful for LCPs and RCPs within a single CP pool)

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 customer 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, or Linux)

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 System z Application Assist Processor (a **Special Purpose CP**). Each zAAP is a real CP on the CPC (zEC12, zBC12, z196, z114, z10, z9, z990, and z890 only). zAAP LCPs, using zAAP real CPs, can only run certain Java code on behalf of its associated z/OS-1.6 (or higher) partition.

zAAP/IFL/ICF Pool For z990 and z890 processors, all of the zAAP, IFL, and ICF CPs are managed by LPAR as a single pool of real CPs, in support of all defined zAAP, IFL, and ICF partitions. zAAP, IFL, and ICF LCPs can only be dispatched on the real CPs in this pool. On processors prior to the z990 and z890, zAAPs are not available and the IFL and ICF CPs operate as a single pool.

zAAP Pool All of the CPs in the zAAP pool. For processor models subsequent to the z990 and z890, zAAP, IFL, and ICF CPs each operate as a separate pool.

zIIP System z Information Integration Processor (a **Special Purpose CP**). Each zIIP is a real CP on the CPC (zEC12, zBC12, z196, z114, z10, and z9 only). zIIP LCPs, using zIIP real CPs, can only run certain redirected DB2 code on behalf of its associated z/OS-1.6 (or higher) partition.

zIIP Pool All of the CPs in the zIIP pool. zIIPs can be configured only on zEC12, zBC12, z196, z114, z10, and z9 processor models, and always operate as a separate pool.

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 LSPR data (based on a single version of z/OS) to generate capacity tables and to support the algorithms related to partitioning.

OS/390
Any version of OS/390. **zPCR** no longer includes any OS/390 LSPR data, which was the sole source for making comparisons that included /390 processors.

MVS
Includes any version of z/OS or OS/390.

z/VM
Any version of VM. **zPCR** considers VM generically, i.e., not specific to any particular release.

z/VSE
Any version of VSE. **zPCR** considers VSE 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.

CFCC
Coupling Facility Control Program, which normally runs in an ICF partition, but may also be run in a General Purpose partition.

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 can be used in terms of the capacity of the associated General Purpose partition. **zPCR** algorithms assume 100% loading (i.e., the amount of zAAP or zIIP capacity required is equal to the capacity of the associated GF partition. Specifying zAAP/zIIP loading percentages less than 100%, will project an increase for the associated General Purpose partition's capacity, due to a reduction in switching cost.

Files Recognized by zPCR

Study File   A *Study file* includes LPAR host and partition configuration definitions saved by **zPCR** (file extension = *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.

RMF   A *Resource Measurement Facility* report is generated using standard z/OS software against z/OS SMF data on an IBM System 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 = *txt*), these reports can be used by **zPCR** to define the LPAR host and its partition configuration or to copy partition definitions into a currently active **zPCR** study. In addition, for each partition represented by a *CPU Activity Report*, information about the use of HiperDispatch and the number of parked LCPs is available.

EDF

An *EDF* (*Enterprise Data File*) is generated by running the *CP3KEXTR* program against z/OS SMF data on an IBM System 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.

For z/VM, create an *EDF* by running the *CP3KVMXT* program against CP Monitor Data on an IBM System z processor.

An EDF can be used by **zPCR** to define the LPAR host and its partition configuration or to copy partition definitions into a currently active **zPCR** study. 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 and the number of parked LCPs is available.

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.