IBM

International Technical Support Organization

Sub-Capacity Reporting Tool

September 2001

First Edition (September 2001)

This edition applies to z/OS Version 1 Release 1 of z/OS (5694-A01), and to subsequent releases and modifications until otherwise indicated in new editions.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. HYJ Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Preface

This document describes how to use the Sub-Capacity Reporting Tool (SCRT) that supports sub-capacity software pricing until the IBM License Manager is available. The SCRT analyzes a month's data from a z900 processor, which includes changes to LPAR sizes and movement of products across LPARs. The SCRT creates a report, called the Sub-Capacity Report. The report indicates the required license capacity for each Variable Workload License Charge product for a z900 processor.

IBM trademarks

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

Domino	SP
IBM ®	TME
Lotus 1-2-3	VM/ESA
PR/SM	VSE/ESA
QMF	z/OS
RMF	z/VM
S/390	zSeries
Sequent	

Sub-Capacity pricing introduction

This chapter describes the IBM's newest software pricing model called Workload License Charges (WLC), designed to support e-business reality.

The following topics are addressed:

- ► IBM pricing model for zSeries
- ► IBM Workload License Charge concepts and benefits
- ► Prerequisites for WLC
- ► Categories of licensing

The objective is to give you an overview of the newest pricing model made available by IBM as well as show the benefits that you can derive from this new model.

1.1 Pricing model for zSeries

Workload License Charges (WLC) is IBM's newest software pricing model, designed to support e-business reality. With Workload License Charges, customers pay only for the software capacity that they need. Because the new pricing structure is designed around leading edge technology and the concept of paying for what you need, customers can grow a single workload at a time, with ease and granularity.

The IBM eServer zSeries Workload License Charges provides:

► New S/390 software pricing model designed to support e-business reality.

► Changes the metric of software pricing from machine-capacity to sub-machine capacity, for qualified products.

► Allows customers to grow one workload at a time, in a simple and granular fashion.

► Provides enhanced price performance.

► Enables customers to utilize excess hardware capacity without incurring software charges.

► Permits customer pay for software based on a defined average capacity requirement instead of their peak requirement.

1.2 Workload License Charge benefits and concepts

IBM zSeries and IBM S/390 software pricing methodologies offer:

► Price-to-value

► Flexibility to run software where it is most efficient

► Ability to predict software charges

► Help with cost of new applications

► Flexibility to pay for software based on workload requirements

zSeries delivers a comprehensive set of software pricing initiatives that facilitate the growth of the industry's hottest applications (Enterprise Resource Planning, Business Intelligence, e-business and Linux) on the zSeries platform, as well as support of traditional workloads.

z/OS is the next generation of the OS/390 operating system. z/OS combined with a zSeries server offers flexibility in how your IBM zSeries software product licenses are managed and charged. When z/OS runs on a zSeries server, a new charge type is available: Workload License Charges.

Workload License Charges (WLC) is a new pricing model that provides flexibility, supports the variability and growth of e-business workloads, and improves the cost of computing. Customers now have the flexibility to configure their system to match workload requirements and be charged for the products used by these workloads at less than the full machine capacity, allowing customers to pay for what they need.

When running z/OS on a zSeries server (or equivalent), Workload License Charges allows you:

► To pay for software based on workload capacity requirements, not necessarily the full capacity of the machine

► The flexibility to add capacity for new workloads, and not have software charges increase for existing or stable workloads

- The flexibility to add capacity for an existing workload and have only the existing workload software charge rise
- To buy reserve hardware capacity (white space) for future growth with no increase in your software bill
- The flexibility to manage e-business spikes
- Enhanced price performance as your workloads grow

WLC offers customers new flexibility in how their IBM zSeries software product licenses are managed and charged. Through a combination of technologies provided in z/OS Workload Manager and zSeries, customers have the flexibility to configure their system to match their workload requirements and be charged for the products needed by these workloads at less than full machine capacity.

In addition, the capacity that a customer defines for a given workload will be used to manage the average capacity level over a rolling four-hour period versus the peak, or maximum capacity allowed. That is, a workload can use more than its defined capacity should a workload spike occur. Such spikes are common today as more workloads are subject to e-business demand. As long as the average of the capacity used by a workload during the four-hour period is not exceeded, the workload can continue to spike above its defined capacity.

IBM provides a two step approach to implementing sub-capacity WLC: presently, the Sub-Capacity Reporting Tool; and in the future, through a planned production-level version of IBM License Manager (ILM), a component of z/OS. Charges at less than full machine capacity apply when:

- z/OS is licensed and running on an IBM zSeries 900 or equivalent plug compatible server
- z/OS is the only MVS-based operating system licensed to that server
- The customer provides the required service data and all other information defined by IBM at the frequency set by IBM

Sub-capacity charging is one approach to better address today's explosive growth of new workloads, driven largely by new e-business applications. It is designed to support the variability of e-business workloads and makes the zSeries platform a more affordable platform for growth.

WLC pricing supports two categories of products:

- Variable Workload License Charge (VWLC) products
- Flat Workload License Charge (FWLC) products.

1.2.1 Variable Workload License Charge products

Variable WLC products include z/OS and a select set of IBM's recent subsystem products; for example, DB2, CICS TS, IMS, MQSeries, COBOL, and others (see Appendix A, "Products eligible for license charges" on page 39 for a complete list of products). All of the Variable WLC products give customers the opportunity to define a *product capacity* that is less than the total capacity of the zSeries server. The defined product capacity for a given VWLC product is the largest number of MSUs allocated to the LPARs in which a VWLC product runs concurrently during a reporting period. The number of MSUs can be defined explicitly by specifying the LPAR defined capacity on the Hardware Management Console, or implicitly by

the MSU-equivalent capacity of the LPAR. If a product is running in a z/OS image running under z/VM, the number of MSUs assigned to the z/VM LPAR is used. If applicable, this value must be manually added to the product defined capacity generated by the Sub-Capacity Reporting Tool.

> **Important:** z/OS, Linux and/or z/VM on every system image is required to define a product at less than full capacity of the machine. No other OS/390 or MVS workload can be running on that zSeries machine.

The LPARs defined capacity is managed by the z/OS Workload Manager to a defined value across a rolling four-hour period allowing for e-business spikes. The benefits that customers will derive from Variable WLC products include the following:

► Pay only for what you need

► Buy reserve hardware (white space) for instant capacity growth or e-business spikes without increasing your software bill

► Flexibility to manage e-business spikes

Sub-Capacity value unit products

IBM has extended sub-capacity pricing to database management and application development tools; value unit pricing is a one-time charge with annual support/subscription. Value units are derived from the MSU capability of the machine; an example of translation from MSUs to Value Units is shown in Table 1-1:

Table 1-1   Variable units tiers

	MSUs	Value units / MSU
Base	1-3	1.00
Tier A	4-45	0.45
Tier B	46-175	0.36
Tier C	176-315	0.27
Tier D	316+	0.20

The following example shows how the total number of Value Units is calculated.

If the customer has installed 1000 MSUs, the applicable Value Units would be: 3 x 1.00 + 42 x 0.45 + 130 x 0.36 + 140 x 0.27 + 685 x 0.20 = 244 Value Units. When calculating the total number of Value Units, the sum is rounded up to the next integer.

The sub-capacity value units will be based on the sub-capacity of the z/OS product on the machine on which the tool is running.

1.2.2 Flat Workload License Charge products

The second category of products supported under WLC is Flat Workload License Charge (FWLC) products. Flat WLC products include legacy S/390 software products, VM/ESA, VSE/ESA, and other middleware (see list in Appendix A, "Products eligible for license charges" on page 39). These products have one "flat" MLC charge, *per machine*, independent of the size of the zSeries server to which the product is licensed. Upgrading a zSeries will have no effect on the license charge for these products.

1.3 WLC price structure

Customers may define a capacity in MSUs for each Variable WLC product to run on a zSeries server. This defined capacity represents the sum of the defined capacities specified for each LPAR in which the product will run, and is the basis for determining the Variable WLC product license charge. Each WLC product will have a MLC price structure, as shown in Figure 1-1.
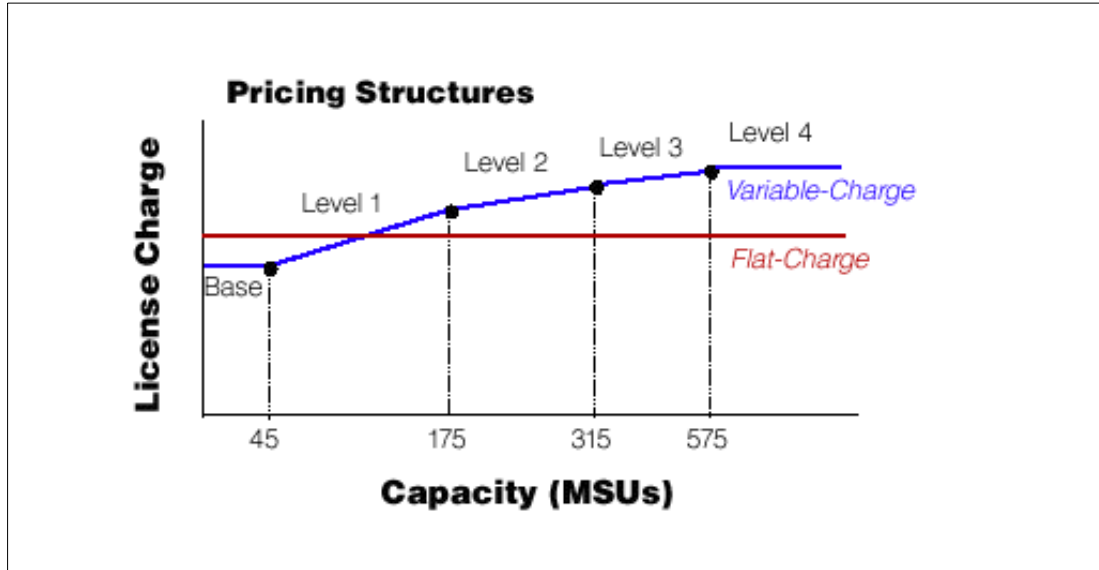
Figure 1-1   MLC pricing structure

The Table 1-2 gives you the variable workload license charge structure and the corresponding MSU levels.

Table 1-2   VWLC levels

Base WLC	45 MSUs
Level1 WLC	46-175 MSUs
Level 2 WLC	176-315 MSUs
Level 3 WLC	316-575 MSUs
Level 4 WLC	576+ MSUs

Flat WLC products have one simple "flat" MLC charge per machine, independent of zSeries server size.

1.3.1 Aggregation rules

Aggregation of Variable WLC product MSUs across multiple zSeries servers is allowed if the environment meets current aggregated WLC guidelines (as stated in the Attachment for zSeries Workload License charges), and only for zSeries servers running z/OS in that Parallel Sysplex. FWLC products will not be allowed to aggregate MSUs since their charge is independent of zSeries server size or how the products are used within a zSeries server.

1.3.2 WLC pricing on z900 machines

When z/OS is licensed to a zSeries server, the Workload License Charge (WLC) applies to all products licensed to that server that have WLC available. If OS/390 is licensed to a zSeries with z/OS, then all Variable WLC products will be priced at the full capacity of the server.

Restriction: If z/OS is not licensed to and running on the zSeries server then the standard PSLC or other charge type applies and WLC is not available.

2

Setting up the system for sub-capacity pricing

This chapter describes the steps necessary for setting-up a zSeries machine running z/OS for sub-capacity pricing. It discusses the optional hardware definitions that the customer can use as well as the necessary software definitions.

This chapter describes the following:

- ► How to define the capacity of your partition (hardware definitions)
- ► How to collect the SMF records (software definitions) necessary for the Sub-Capacity Report Tool (SCRT)
- ► How to extract and file the records for the tool

2.1 Controlling resources at the LPAR level

This section introduces you to some of the concepts and functions for the hardware definitions on a zSeries machine that allow you to optimize your system for sub-capacity pricing. Though some implementation details are provided, in many cases this section refers you to other books to find more detailed information.

The following are methods you can use to control the size of a logical partition:

- ► Using a defined capacity
- ► Defining a hard cap
- ► Defining no cap
- ► Using CPU management

> **Attention:** Refer to *z/OS Planning for Workload License Charges Release 2*, SA22-7506 for a description of the methods.

While several types of LPAR capping methods are discussed in this chapter, using a defined capacity, a new function available when you run z/OS on a zSeries 900, is the recommended method for enabling sub-capacity pricing. The ability to use a hard cap or no cap existed prior to z/OS. You can still use these capping methods when z/OS runs on a zSeries 900. However, these methods have specific implications for sub-capacity pricing which are discussed in this chapter.

2.1.1 Using a defined capacity

A defined capacity refers to the LPAR capacity you define for a workload. Using a defined capacity (also called a soft cap) can help you overcome the limitations of previous capping methods. It also allows for spikes in your workload while maintaining an overall 4-hour average, and is defined in MSUs. A defined capacity applies to all the products running in an LPAR, regardless of the number of products the LPAR contains. You specify a defined capacity for an LPAR using the appropriate LPAR controls.

Suppose you define a 40-MSU LPAR running only CICS. That is, your defined capacity for this LPAR is 40 MSUs. Your CICS application is allowed to spike above the 40 MSU level as long as its 4-hour rolling average stays under its defined capacity, as shown in Figure 2-1. If the LPAR's 4-hour average rises above its defined capacity value, the Workload Manager limits the LPAR to using no more CPU resource than its defined capacity value. Because of this conservative approach to capping the LPAR, you could at times see a 4-hour average that is greater than the defined capacity. The LPAR is capped until its 4-hour average drops below the defined capacity limit. When capping by the Workload Manager is not in effect, a product running in an LPAR with a defined capacity can spike up to the total capacity available to the LPAR.
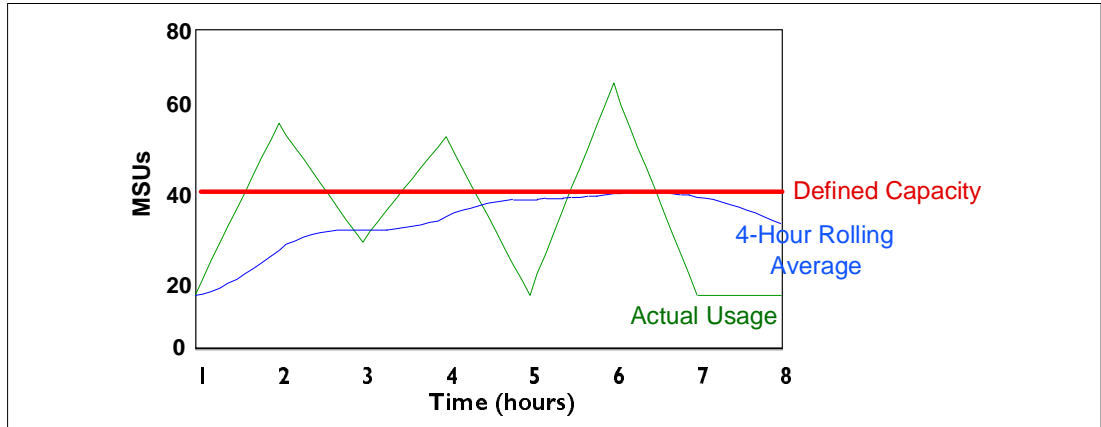
Figure 2-1   How the defined capacity works

You set up the defined capacity for an LPAR by specifying the number of MSUs you want for this partition in the Defined Capacity field on the Change Logical Partition Controls.

Setting up defining capacity

You can define the capacity you want to be assigned to a partition using the HMC profile definition sequence. The first thing you have to do is to log on as SYSPROG, for example, as shown in Figure 2-2.
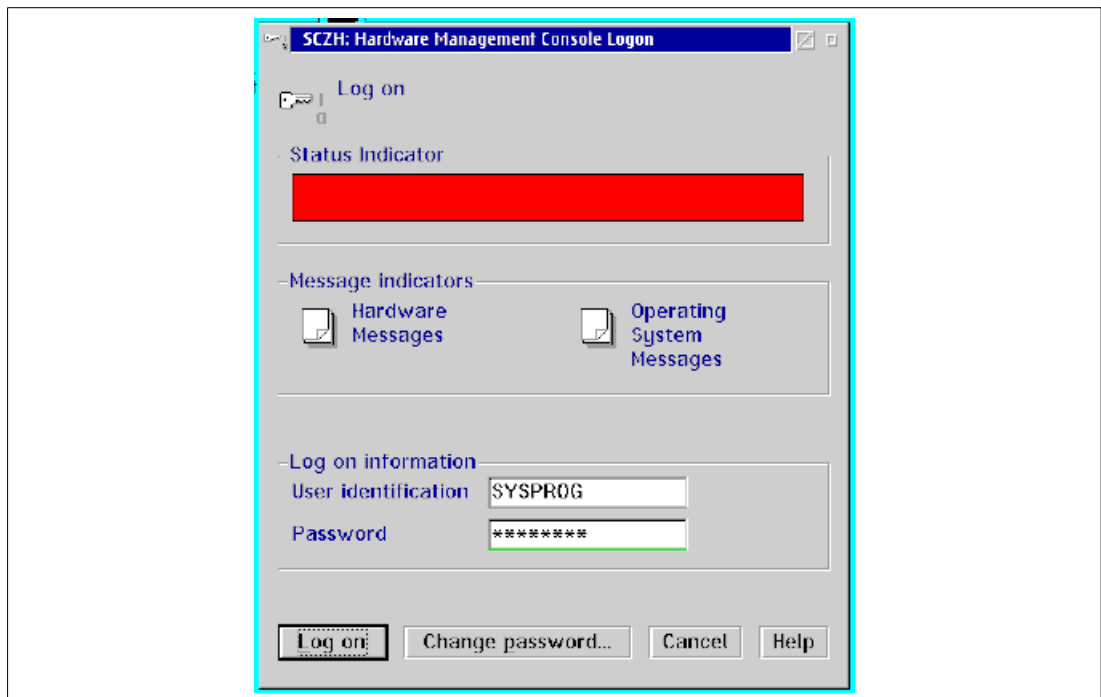
Figure 2-2   HMC logon panel

When you log on as SYSPROG, the panel shown in Figure 2-3 on page 10 is displayed; you have to select the CPC for which you want to customize the profile. In our example the processor is SCZP801.
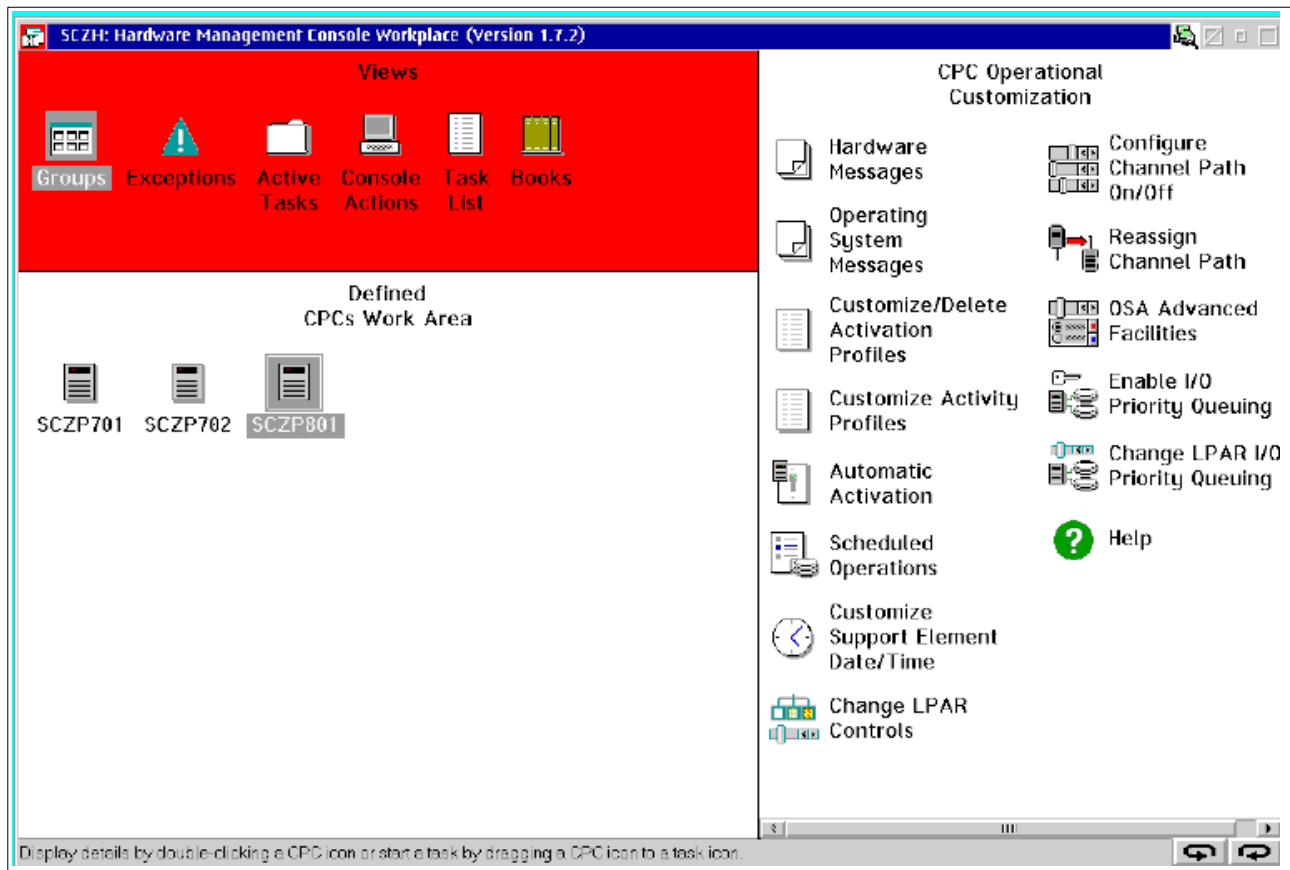
Figure 2-3   HMC Group Work Area Panel

You select the CPC icon and double-click the Customize Activity Profiles icon. This icon is part of the CPC Operational Customizing task. When you select it, the panel shown in Figure 2-4 on page 11 is displayed.
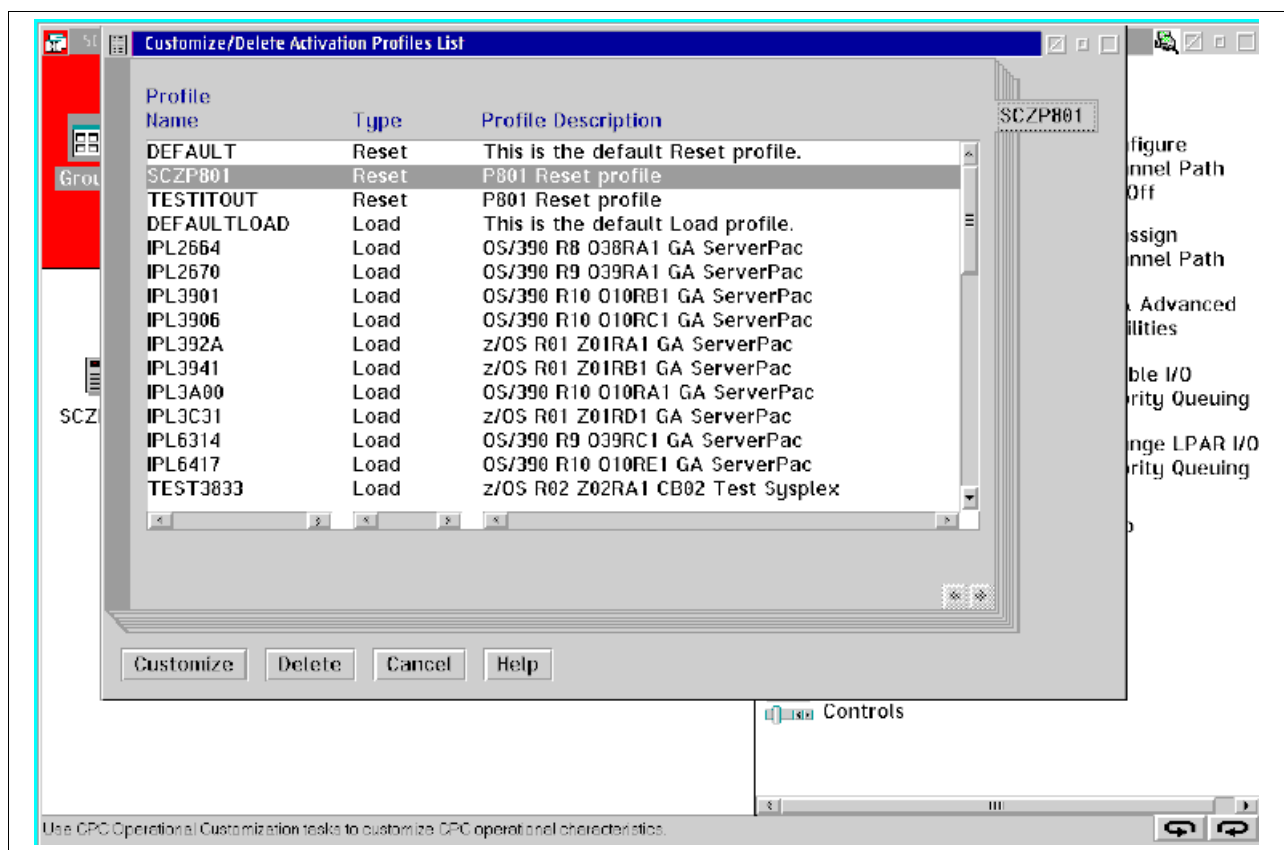
Figure 2-4   Customize/Delete Activation Profile List panel

In Figure 2-4, select the Reset profile for your CPC and click the Customize button at the bottom of the panel. The customization and activation panel for profile SCZP801 is displayed, as shown in Figure 2-5 on page 12.

Figure 2-5   Activation Profile Customizing panel

Listed in this panel are all the partitions defined for this CPC. In our example we have ESA/390 partitions as well as CF partitions. Select the ESA/390 partition for which you want to define a capacity value. It is recommended that you define a capacity value for all z/OS partitions in a z/Series processor.

To select a partition, Click on the tab (right side of the panel) containing the name of the partition; in our case we clicked at **SCZP801:A7**, our A7 partition. The partition activation profile panel is displayed; select the **Options** tab in this panel. When you do this the image options panel shown in Figure 2-6 on page 13 is displayed. In this panel you can define, among other things, the capacity you want to assign to this partition; in our example we assigned the value 45. Remember that the number you specify in this field represents the number of MSUs you want to associate to this partition.

You must save the profile after executing this procedure. The next time the partition is activated, the defined capacity of the partition will be used. So, in order to set up a defined capacity at system initialization, you must use the Customize Activation Profiles selection in the CPC Operational Customizing task list.
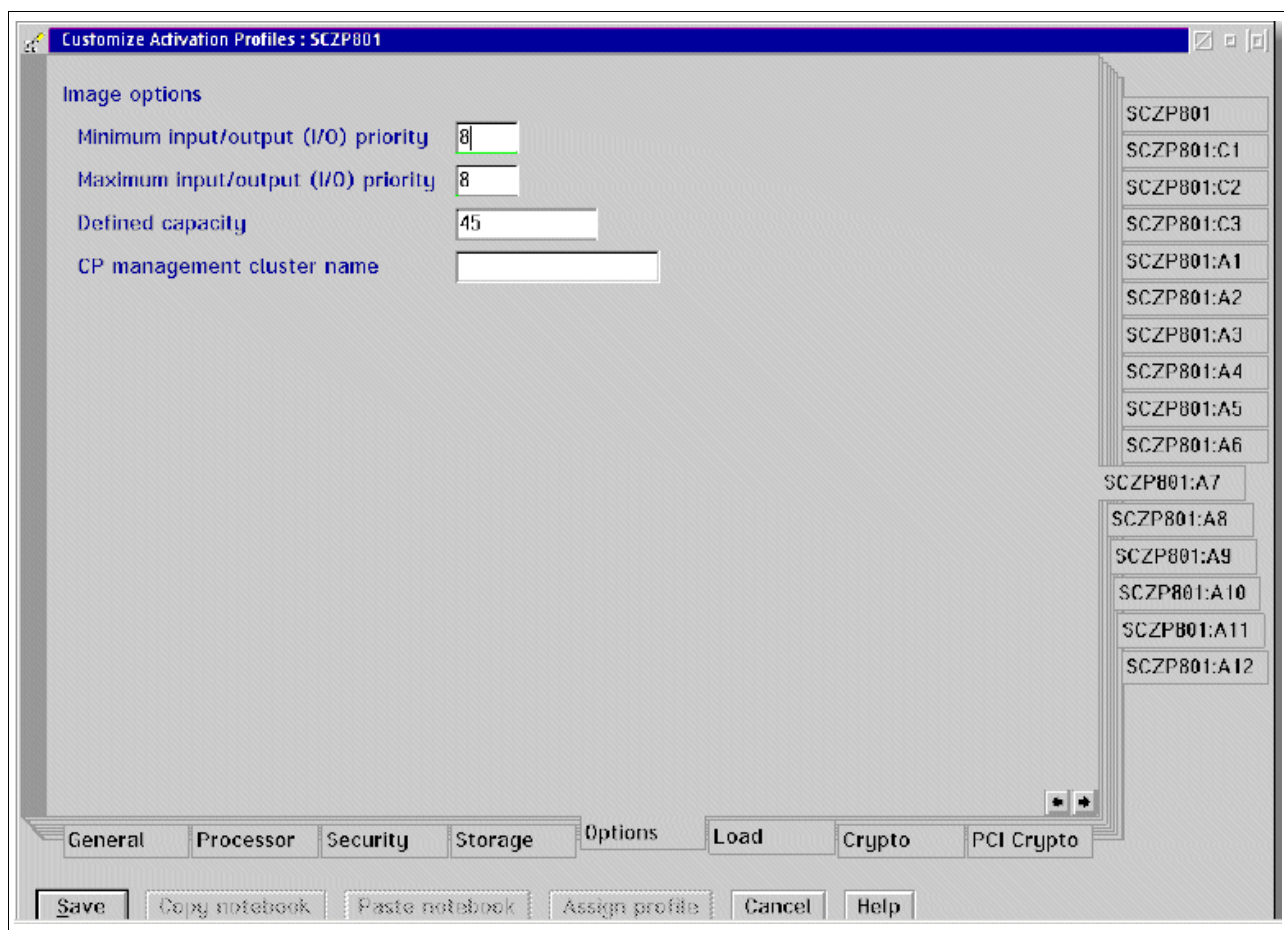
Figure 2-6   Image options panel

> **Important:** Once your system is up and running you can use the Change Logical Partition Controls panel to change a defined capacity.

Dynamically changing a defined capacity

You can increase or decrease the defined capacity for an LPAR without stopping z/OS or the other product (or products) running in the LPAR. To dynamically change the defined capacity for an LPAR go to the Defined Capacity field on the Change Logical Partition Controls panel and specify the new MSUs for this LPAR.

To perform this change you have to use the HMC; you first log on as a system programmer. When you do that, the panel shown in Figure 2-3 on page 10 is presented. In this panel you select the CPC that contains the partition you want to change the defined capacity of. After selecting the CPC, double-click the **Change LPAR Control** icon. When you do this, the Change Logical Partitions Control panel, shown in Figure 2-7 on page 14, pops up. In this panel you can overtype the defined capacity value specified in the image activation profile. After entering the new value, click the **Change Running System** tab (or Change and Save tab) to dynamically reset the defined capacity for the partition to the new value.
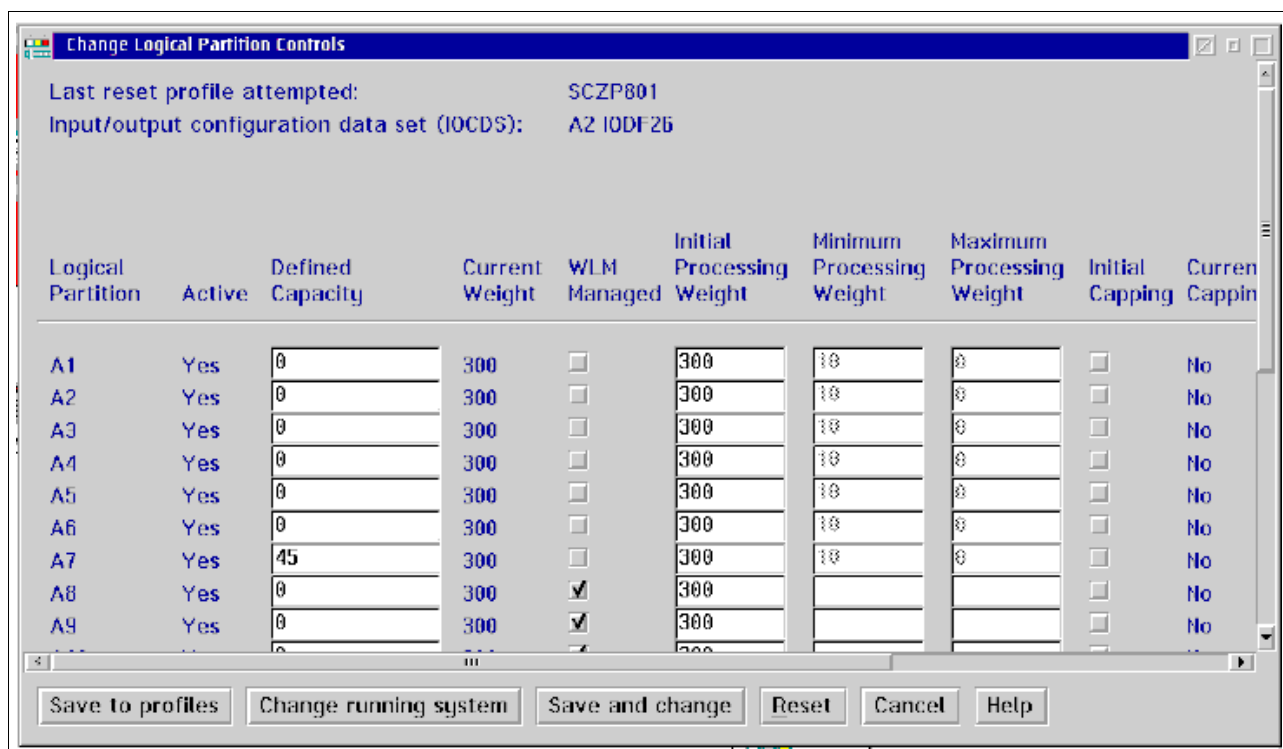
Figure 2-7   Change Logical Partition Controls panel

When you do this the panel shown in Figure 2-8 is displayed, confirming that the requested operation was completed. Whenever you dynamically change a partition's defined capacity, the operating system is notified and the new value will be recorded by SMF.
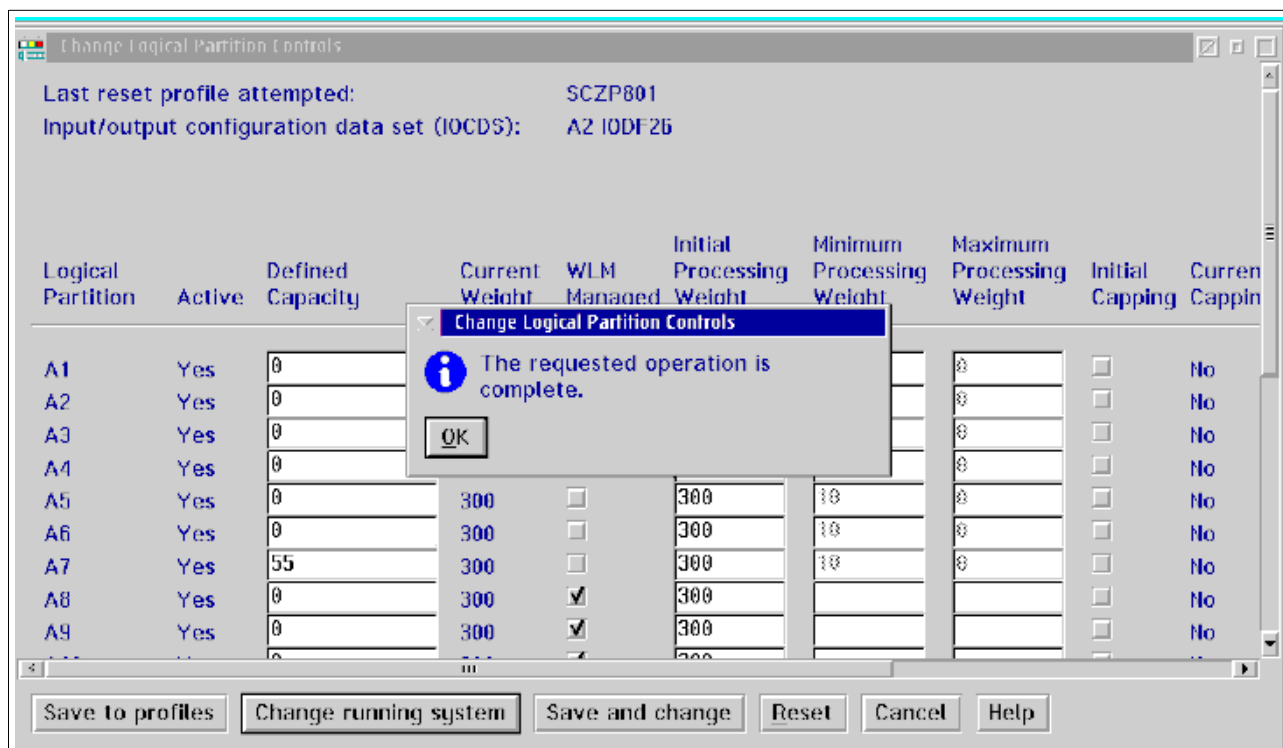
Figure 2-8   Change Logical Partition Controls confirmation

Implications for sub-capacity pricing

Using a defined capacity is the recommended method for setting up sub-capacity pricing for a product. Your software charges are based on the capacities you define. Spikes in your workload, up to the total size of the LPAR, are accommodated.

2.2 Monitoring your system

WLM in z/OS provides support to allow an LPAR to be managed to a specified defined capacity. The defined capacity value allows an LPAR to be capped when the LPAR's capacity exceeds this limit as judged by a rolling four hour average. This applies only to z/OS images running on a 2064 z900 CPC. When the actual CPU usage of the LPAR (calculated as the average over 4 hours) exceeds the defined capacity, WLM in conjunction with PR/SM will apply a processor cap to reduce the CPU usage of the LPAR. This cap will be applied to reduce the 4 hour average usage to no longer exceed defined capacity. This function is called soft capping an LPAR. As CPU demand is reduced, and the rolling four hour average no longer exceeds the defined capacity, the soft capping function will be removed dynamically by WLM and PR/SM.

The WLM support, introduced with z/OS V1R1, is provided in support of sub capacity Workload License Charges. Installations are warned the WLM soft capping function is active whenever a defined capacity is specified on the LPAR Change Logical Partition Controls panel, even if sub-capacity pricing is not being used. Installations are cautioned to only use the defined capacity definitions when they want to have the potential to soft cap a partition.

Performance problems and throughput problems may be seen if WLM soft caps are accidentally turned on.

Installations should also be aware of the WLM calculation of a rolling four hour average, and the application of a soft cap is active on any z/OS image on a z900 processor, even if there are OS/390 partitions active on the CEC. This is contrary to sub-capacity WLC which requires all active MVS partitions to be running z/OS to be eligible for sub-capacity WLC. For more information on Workload License Charges, see:

```
http://www.ibm.com/zseries/wlc_lm
```

In order to determine if a defined capacity has been specified, review either the LPAR Logical Partition Control panel, or review an RMF Partition Data Report. For information on how an installation would specify a defined capacity, see "Setting up defining capacity" on page 9.

Any specification of a defined capacity greater than zero, will cause WLM to start managing the partition to a soft cap. The WLM Managed check box, shown in Figure 2-7 on page 14, is not used to control the soft capping function. This check mark is used to indicate if WLM should be performing IRD functions. If an installation inadvertently specifies a defined capacity, the function can be turned off by dynamically updating the running system and specifying a zero, for the partition's defined capacity.

2.2.1 Using RMF reports

Using RMF, if WLM applies a soft cap to the LPAR's capacity exceeding the defined capacity you will see a value in the CAPPING WLM% field of the Partition Data Report. An example of a partition which is having a soft cap applied is Figure 2-9. Partition NP4, 87% of the time is being limited in it's ability to use CPU because of a soft cap. For more information on soft capping, review *MVS Planning: Workload Management*, SA22-7602. For more information on the PR/SM interfaces used in this support, review *PR/SM Planning Guide*, SB10-7033.

```
--------- PARTITION DATA ---------------- -- LOGICAL PARTITION PROCESSOR DATA --

                  ----MSU----   -CAPPING--   PROCESSOR-   ----DISPATCH TIME DATA----
NAME      S   WGT DEF    ACT    DEF   WLM%    NUM   TYPE    EFFECTIVE        TOTAL

NP1       A   400   0      5    NO    0.0      4    CP     00.02.27.649   00.02.30.388
NP2       A   200   0      5    NO    0.0      4    CP     00.02.28.330   00.02.30.449
NP3       A   200   0     65    NO    0.0      4    CP     00.31.47.073   00.31.47.881
NP4       A   400  50     54    NO   86.7      4    CP     00.26.46.161   00.26.45.617
NP5       A   850  50      5    NO    0.0      5    CP     00.02.14.117   00.02.16.743
CB88      A    10   0      0    NO    0.0      2    CP     00.00.00.000   00.00.00.000
CFC1      A   DED   0     30          0.0      1    CP     00.14.59.839   00.14.59.862
CFC2      A   DED   0     30          0.0      1    CP     00.14.59.945   00.14.59.963
*PHYSICAL*                                                               00.00.50.252
                                                          ------------   ------------
TOTAL                                                     01.35.43.117   01.36.41.158
```

Figure 2-9   RMF report

2.3 Capacity planning

The capacity planning life-cycle, as it relates to WLC, consists of the following tasks:

1. Measuring your use of system resources

2. Configuring defined capacities

3. Optimizing your resources

2.3.1 Measuring your use of system resources

This section describes the tools you can use to measure your existing LPARs in preparation for establishing defined capacities. These tools can be used in conjunction with your current operating system (OS/390 V2 R7, for example) prior to your upgrade to z/OS. You can even run these tools on your current 9672 G6 CPCs in preparation for installing a z/900. The first step is to measure the system resources being consumed by your existing operating systems (LPARs). To do this you can use the following tools:

► CP2000 capacity planning tool

This tool can be used by your IBM representative to display the rolling 4-hour average for each LPAR.

► WLC Tool

This tool, available to IBM customers, also reports the rolling 4-hour average for each LPAR. You can download this tool from the following Web site:

    http://www.ibm.com/servers/eserver/zseries/wlc_lm/tool.html

► RMF Partition Data report

RMF allows you to run reports that tell you how much system capacity your LPARs are actually using. You can use the RMF Partition Data Report to determine:

– The capacity of an LPAR in MSUs
– The CPU usage of an LPAR in MSUs
– The percentage of time an LPAR was capped to enforce a soft cap

> **Note:** The RMF Partition Data report with soft capping data is only available with z/OS.

For more information, see *z/OS Resource Measurement Facility User's Guide*, SC33-7990.

Configuring LPARs

Once you have measured the capacity your system is using, you next need to configure your system in terms of LPAR sizes.

Setting up your LPARs

To optimally implement WLC you should set a defined capacity for each LPAR. A typical value for a defined capacity is the highest 4-hour rolling average you observe during a given measurement period. A more conservative approach is to set the defined capacity equal to the maximum processing power of the LPAR (based on the number of logical CPUs associated with the LPAR).

Note that you can choose to divide a CPC into LPARs that have a total capacity less than the full capacity of the CPC. This allows you to be charged for less than the full capacity of the CPC. The portion of hardware that does not carry any software charges is called *white space*. You can use white space to support short term workload spikes and long term workload growth.

Additional considerations

The following are additional configuration recommendations:

► Make sure FWLC products run on appropriate CPCs in your configuration. For these types of products you are charged for each CPC running the application.

► Defined Capacity does not apply to z/VM LPARs. z/OS guests in VM LPAR will use the image capacity of the z/VM LPAR for their software pricing. That image capacity is determined as follows:

 – If the z/OS guests are dedicated then it is number of engines available to the LPAR times engine size.

 – If there is a hard cap, then it is the hard cap as a percentage of the shared pool.

 Otherwise it is the number of engines available to the LPAR times engine size.

► SCRT is not aware of the LPAR capacity used by z/OS when z/OS in 64-bit mode runs in a virtual machine under VM. If z/OS runs under a z/VM LPAR, then the defined capacity for the LPAR must satisfy all the CPU demands of the z/VM guest, and software prices for z/OS guest usage will be based on that defined capacity.

► If you order the Integrated Facility for Linux (IFL) with your zSeries 900, be aware that the capacity of an IFL is "invisible" to z/OS in that it is not considered to be part of the full CPC capacity. This means that WLC does not apply to any Linux LPARs running in an IFL. (Similarly, you can order Integrated Coupling Facilities to run coupling facility LPARs. These LPARs are also invisible to z/OS). If, however, you set up a non-IFL LPAR containing Linux, the Linux LPAR capacity is included in the total CPC capacity. This means that if for some reason full capacity pricing goes into effect, then the full capacity price you pay for a VWLC product will include the MSUs being used by the Linux LPAR.

Optimizing your resources

To optimize your system, you should attempt to maintain a balance between system performance and WLC pricing benefits. Once you have an operating system running in an LPAR:

► You should consider gradually lowering the LPAR size until the Workload Manager begins capping the LPAR. Examine the LPAR RMF report that shows the percentage of time the LPAR was capped. If necessary, raise the capacity of the LPAR so that the LPAR is capped only 1 to 5 percent of the time. This ensures you are getting your work done optimally at the lowest cost.

► You should consider tuning your system capacity from time to time based on RMF reports. For example:

  – If an RMF report shows that an LPAR is consistently running at 95% of capacity (MSUs), you should consider increasing the size of the LPAR.
  – If an RMF report shows that an LPAR is consistently running at 40% of capacity (MSUs), you should consider decreasing the size of the LPAR.

2.4 Software definitions

This section discusses the definitions you have to implement in your operating system to collect the data necessary for the Sub-Capacity Reporting Tool, in order to qualify for sub-capacity pricing.

2.4.1 SMF requirements

System management facilities (SMF) collects and records system- and job-related information that your installation can use for accounting and billing purposes. SMF formats the information that it gathers into system-related records (or job-related records). System-related SMF records include information about the configuration, paging activity, and workload. Job-related records include information on the CPU time, SYSOUT activity, and data set activity of each job step, job, APPC/MVS transaction program, and TSO/E session.

The volume and variety of information in the SMF records enables installations to produce many types of analysis reports and summary reports. The Sub-Capacity Reporting Tool requires the installation to collect SMF records 70 and 89. To accomplish this, the installation must tell SMF that it has to collect these records; this is done by customizing SMF using the SMFPRMxx Parmlib member. The SMF parameters allow you to select specific records to be collected, specify SMF data set names, specify the system identifier to be used in all SMF records, select specific record subtypes, collect SMF statistics, and so on.

SMF records are selected by specifying either the type desired (or the types not desired) with the TYPE or NOTYPE option of the SYS or SUBSYS Parmlib parameter.

```
ACTIVE                         /*ACTIVE SMF RECORDING*/
DSNAME(SYS1.&SYSNAME..MAN1,  /*SMF DATA SET NAMES*/
       SYS1.&SYSNAME..MAN2,  /*SMF DATA SET NAMES*/
       SYS1.&SYSNAME..MAN3)  /*SMF DATA SET NAMES*/
NOPROMPT                       /*DON'T PROMPT THE OPERATOR*/
REC(PERM)                      /*TYPE 17 PERM RECORDS ONLY*/
SID(&SYSNAME(1:4))          /* SYSTEM ID IS &SYSNAME  */
LISTDSN                     /* LIST DATA SET STATUS AT IPL*/
LASTDS(MSG)                    /*DEFAULT TO MESSAGE          */
NOBUFFS(MSG)                   /*DEFAULT TO MESSAGE          */
SYS(TYPE(0:255),EXITS(IEFU83,IEFU84,IEFU85,IEFACTRT,IEFUJV,IEFUSI,
                IEFUJP,IEFUSO,IEFUJI,IEFUTL,IEFU29),
    NOTYPE(99),NOINTERVAL,NODETAIL)

SUBSYS(STC,NOTYPE(99),
    EXITS(IEFU29,IEFU83,IEFU84,IEFU85,IEFUJP,IEFUSO))
```

Figure 2-10   SMFPRMxx member

In Figure 2-10, we have requested that all SMF records be collected; for STCs we do not want to collect SMF records 99. You can specifically request for records 70 and 89; to do so you have to code in your SMFPRMxx member the following statements:

```
SYS(TYPE(70,89))
SUBSYS(STC,TYPE(70,89))
```

2.4.2 Extracting records to be used by SCRT

The input for the Sub-Capacity Reporting Tool is a sequential data set containing SMF records 70 (CPU activity) and 89 (Product Usage). Refer to *z/OS MVS System Management Facilities (SMF)*, SA22-7630 for the description of the fields in these records. To obtain these records from your monthly SMF file, which contains the SMF data for all the LPARs on the machine, you can execute a job based on the JCL in Figure 2-11 on page 19, after making changes to the job card, data set names, SPACE parfameter, and the DATE.

> **Attention:** The next three JCL examples are samples and you should use your own JCL techniques to obtain the data required by the SCRT.

```
//LUIZ1    JOB
//STEP001  EXEC PGM=IFASMFDP
//SYSPRINT DD SYSOUT=*
//DUMPIN   DD DSN=SYS1.SYSA.MAN1,DISP=SHR
//SCRT     DD DSN=SCRT.SYSA.SMF.DATA,DISP=(MOD,CATLG),
//         UNIT=SYSDA,SPACE=(CYL,(40,10))
//SYSIN    DD *
 INDD(DUMPIN,OPTIONS(DUMP))
 OUTDD(SCRT,TYPE(70,89))
 DATE(2001213, 2001244)
```

Figure 2-11   SMF Reduction Program for records 70 and 89

The input for the tool will be in the data set defined by the SCRT DD card.

Another way of achieving this is: Collect the data into its own data sets by SYSID; to do this you can modify the job which dumps off the SMF data to add another DD to send the 70s and 89s to a separate data set, one per SYSID. For example take your existing job and add the SCRT DD and the OUTDD(SCRT...) card, as shown in Figure 2-12.

```
//SMFDMP   JOB
//         EXEC  PGM=IFASMFDP
//SYSPRINT DD  SYSOUT=*
//DUMPIN   DD  DISP=SHR,DSN=SYS1.MANB
//ALLSMF   DD  DSN=ALLSMF.DATA(+1),DISP=(,CATLG),
//             SPACE=(CYL,(40,10)),UNIT=SYSDA
//SCRT     DD  DSN=SCRT.SYSA.SMF.DATA,DISP=(MOD,CATLG),
//         UNIT=SYSDA,SPACE=(CYL,(40,10))
//SYSIN    DD  *
 INDD(DUMPIN,OPTIONS(DUMP))
 OUTDD(ALLSMF,TYPE(0:255))
 OUTDD(SCRT,TYPE(70,89))
/*
```

Figure 2-12   Another SMF reduction example

To collect it by month, determine the average frequency that the SMF data set switches, for example every 4 hours. Create a job for your automation package that runs at half that interval after midnight on the second of each month that renames this data set (see Figure 2-13).

```
//RENAME   JOB
//IDCAMS   EXEC  PGM=IDCAMS
//SYSPRINT DD SYSOUT=*
//SYSIN    DD  *
ALTER 'SCRT.SYSA.SMF.DATA' NEWNAME('SCRT.SYSA.SMF.AUG2001.DATA')
```

Figure 2-13   Monthly procedure

2.4.3 Data collection environment

With the process just described, you will be able to implement a process to collect the SMF data and assemble the data into the single data set for each z900, which will serve as the input to the SCRT. Remember that it is your responsibility to create an input file with data from the entire z900 machine, including the correct range of dates (2nd of month X through the 1st of month Y). In other words, you have to extract the SMF records 70 and 89 from all z/OS images in each z900 and accumulate the data on the same data set. The SMF records have to be writen at least every hour; in other words the maximum duration of an interval is one hour.

Sub-Capacity Reporting Tool - system programmer tasks

This chapter describes the tasks a system programmer has to perform to obtain, install, and run the IBM Sub-Capacity Reporting Tool.

This chapter contains the following:

- ▶ SCRT description and capabilities
- ▶ How to obtain the tool
- ▶ How to use the tool

3.1 SCRT description and capabilities

Sub-Capacity Reporting Tool (SCRT) is an as-is, no-charge IBM interim tool aimed to support sub-capacity software pricing until IBM License Manager is available. The SCRT analyzes a month of data for a particular z900 machine, including changes to LPAR sizes and movement/addition/deletion of products across LPARs. The SCRT creates a report, called a "Sub-Capacity Report"; this report indicates the required licensed capacity for each Variable Workload License Charge product for a particular z900 machine.

For each z900 wishing to qualify for sub-capacity pricing:

► An entire month of SMF data must be collected on that machine.

► The tool must be run against data from that machine.

► The Sub-Capacity Report must be produced for that machine.

► The Sub-Capacity Report must be submitted to IBM for that machine.

► The machine must be configured to send weekly Transmit System Availability Data (TSAD) to IBM via the IBM eServer zSeries 900 RSF. If the machine does not perform a weekly "call home" via the RSF, an alternate means to collect and transmit the data is documented in the z/OS publication *Planning for Workload License Charges*, SA22-7506 at:

```
http://www.ibm.com/zseries/wlc_lm
```

A small subset of the data already being transmitted via RSF today will enable IBM to verify that capacity allocation data submitted to IBM on Sub-Capacity Reports is consistent with the actual machine configuration. Failure to send weekly TSAD to IBM may result in charges at full machine capacity.

The SCRT post-processes SMF data, specifically SMF70s (CPU Activity) and SMF89s (Product Usage) records. The SCRT takes the specified data sets in the JCL as input, and has a single data set as output. The input data sets must contain one month's SMF70 and SMF89 records for all the images on a z900 machine. The output data set is a comma-separated-value (CSV) flat file, referred to as the Sub-Capacity Report. The recommended approach is to download the Sub-Capacity Report to a desktop machine and review it for accuracy (and edit it in some cases) using any spreadsheet package, prior to submission to IBM via e-mail.

> **Attention:** It is possible to view and edit the report on the host using a text editor or on the desktop using a text editor, but this is a less convenient approach.

3.1.1 SCRT capabilities

The SCRT capabilities are to:

► Provide simple, non-intrusive means for customers to gather and report sub-capacity information to IBM

► Allow IBM to receive, view and store sub-capacity information from customers

► Allow IBM to update customer bills according to sub-capacity information received from customers

► Enhance the pure trust-model by providing IBM some degree of audit ability

> **Important:** IBM intends to utilize the sub-capacity report in conjunction with Call Home data for audit purposes.

3.1.2 SCRT data collection

As mentioned in "Data collection environment" on page 20, the input for the SCRT tool is SMF data sets containing SMF70 and SMF89 records for the month being reported, from all images in the same z900 system. The SCRT tool, however, analyzes any SMF70 and SMF89 records in the input file, regardless of the range of dates. Other SMF record types are ignored. The tool is designed to support sub-capacity pricing; therefore, it *only* analyzes data if the data was generated by a z900 machine running z/OS.

If the tool is given an input data set containing data from a non-z900 machine or if z/OS is not present on the z900 machine being analyzed, the tool exits and a report is not generated. If multiple z900 machines are contained within the input data sets, only the first z900 machine is analyzed, or the machine specified in the parameter list for the tool. Figure 3-1 shows the possible data collection environments and whether or not what type of report is generated.

	z900	non z900 hardware	multiple mainframe machines
z/OS	analyze data, produce report		analyze data, produce report for 1st z900 in input dataset
mixed z/OS and OS/390	analyze data, produce report		analyze data, produce report for 1st z900 in input dataset
OS/390			
MVS			

Figure 3-1   Data collection environments

Attention: To produce SMF70 records, a systems management product is required. IBM's product in this space is the Resource Measurement Facility (RMF) feature of the z/OS operating system. For customers who do not have RMF, other vendor's products (like BMC's CMF Monitor) can produce compatible SMF70 records. No additional product, beyond the base z/OS operating system, is required to produce SMF89s records.

3.2 How to obtain and install the SCRT tool

The first thing you have to do to use the SCRT tool is to get it from the following WLC Web site which is the same site used for the planning tool:

```
http://www.ibm.com/servers/eserver/zseries/wlc_lm/tool.html
```

As described on this site, the steps to obtain the tool are:

1. Download the tool, which is contained in a self extracting zip file.

2. Execute the file you downloaded (*scrttool.exe*) on your PC to decompress the file. The decompressed file will be named *scrttool.bin*.

3. Upload the file (*scrttool.bin*) to your host system, using binary transfer, and move the file to an 80 byte fixed data set. This file will contain the JCL that executes a load-and-go job (see Example 3-2 on page 24) and the MVS loader is executed to link the object deck into memory and execute the tool for every iteration.

If you want to link the tool to a library you certainly can do that, but IBM does not provide the JCL to do this; you have to customize the job card to your installation standards.

3.3 How to use the SCRT tool

To run the SCRT Tool, the SMF data does not need to be sorted. You will need to make one invocation of the tool for each CEC eligible for sub-capacity pricing. The invocation for each CEC will need to have the data sets for all the included LPARs concatenated. The JCL that comes with the tool is shown in Figure 3-2.

```
//SCRTRPT  JOB
//DELETE   EXEC  PGM=IEFBR14
//DD1      DD  DSN=HLQ.CECA.SCRTTOOL.CSV,DISP=(MOD,DELETE),
//             UNIT=SYSDA,SPACE=(TRK,(1))
//EXTR     EXEC  PGM=LOADER,PARM='/CUSTOMER NAME'
//SYSPRINT DD    SYSOUT=*
//SYSLOUT  DD    SYSOUT=*
//SMF      DD    DISP=SHR,DSN=SCRT.SYSA.SMF.AUG2001.DATA
//         DD    DISP=SHR,DSN=SCRT.SYSB.SMF.AUG2001.DATA
//         DD    DISP=SHR,DSN=SCRT.SYSC.SMF.AUG2001.DATA
//         DD    DISP=SHR,DSN=SCRT.SYSD.SMF.AUG2001.DATA
//OUTPUT   DD    DISP=(,CATLG),DSN=HLQ.CECA.SCRTTOOL.CSV,
//               UNIT=SYSDA,SPACE=(TRK,(15,15))
//SYSLIN   DD    *
```

Figure 3-2   JCL to execute the Sub-Capacity Reporting Tool

The first step is used to delete the file created by a previous execution of the tool. The following sections describe the JCL statements and parameters specified in Figure 3-2.

3.3.1 EXEC card

The PARM field must be specified to contain at least one parameter (mandatory); the second parameter is optional, as follows:

Customer name The first parameter is the customer name which must be present; if this parameter is not specified a message is issued and the tool terminates with return code 16. The size of the first parameter, the customer name, can be up to 80-bytes; for example, see Example 4-1 on page 30 where the customer name is THEBIGBANK.

Note: The / following the PARM= is required because the parameters before the / are passed to the loader and the parameters after the / are passed to the program being loaded.

3.3.2 DD cards

The execution of the tool requires the definition of three DD cards: SMF, OUTPUT, and SYSPRINT. The SMF DD card identifies the SMF input data set. This data set contains all SMF records 70 and 89 for all z/OS systems executing in all partitions defined in the same z/900 machine (see 2.4.1, "SMF requirements" on page 18). Example 3-2 on page 24 shows the concatenation of data sets containing SMF records 70 and 89, one for each image running in that CEC (four images in the example).

The file defined by the SMF DD card is processed as follows:

► Open the SMF file and check return codes; if OK, proceed to read the file.

► Read the first record; if it is a record 70 then the record CPU type and serial number is checked against the ones specified in the parm field of the JCL EXEC statement; if it does not match then the record is skipped and a message is issued to the SYSPRINT output indicating that a different type-serial was found in the input data set. This message is only issued once for each unique type-serial pair. If the type-serial parameter is not specified in the EXEC parm field then it is taken from this record (first SMF 70 record) and all the SMF 70 records which do not match the type-serial of the first SMF 70 record will be ignored.

► The processor type of the SMF 70 record is checked; if it is not 2064, a message is issued and processing terminates with return code 16.

► The next check is to verify if the record was generated by an image running z/OS; if not, a message is issued, this record is ignored and the next record is processed; this message is only issued once for each SYSID.

► Next a check is made to see if the machine is in LPAR mode (not basic mode); if it is not, a message is issued and the program terminates with return code 16 since Sub-Capacity pricing is not available in basic mode.

► This concludes the validation for SMF records 70. The records are then processed to build the matrix of processor and partition utilization.

► If it is a record 89 and this is the first record, then all checks made for a record 70 (as described above) are also performed for record 89. Record 89s are processed to obtain information about products that have been executed as well as the earliest and latest time of their execution.

► Each product in the record is matched against an internal table of products. If the product is listed in the table, it is used; if it is not found, it is discarded.

► If it is not a record 70 or 89, read the next record; at end of file, the SMF file is closed.

The OUTPUT DD card defines the file that is created by the tool: the sub-capacity report. This is a sequential file; it contains a comma-separated-value (CSV) flat file that is validated by the software asset manager and sent to IBM. The processing of the CSV flat file is described in Chapter 4, "Software asset manager tasks" on page 27.

> **Note:** The record size for this DD is VB 1024.

The SYSPRINT DD card defines the message file. All messages issued by the tool are recorded in this file. The messages issued by the tool are described in Appendix B, "SCRT Messages" on page 43.

4

Software asset manager tasks

This chapter describes the tasks that the software asset manager (SAM), the person in your organization responsible for controlling software billing, will have to perform to analyze and, if necessary change and comment on the output produced by the SCRT. It also describes the procedure SAM has to follow to send the information to IBM.

In this chapter, the following topics are discussed:

- ► How to analyze the sub-capacity report
- ► How to change information in the sub-capacity report
- ► How to send the sub-capacity report to IBM
- ► How IBM handles the report

4.1 How to analyze the sub-capacity report

IBM's recommendation is that the sub-capacity report produced by the SCRT, a comma-separated-value file, should be downloaded to a directory in your desktop machine and its contents reviewed for accuracy (and edited if necessary). This review can be accomplished using any spreadsheet package, such as Lotus 1-2-3 or Microsoft EXCEL. To do this using Lotus 1-2-3:

1. Select File->Open
2. Then in the dialog box select text (csv) in the types of files drop-down box
3. Then select the file from the directory
4. Click open to get another dialog box for Text File Options, and select the radio button for "Parse as CSV file" and click OK
5. Rearrange the column widths, especially the first column to see it properly.

A similar procedure can be used with Microsoft EXCEL.

> **Note:** It is possible to view and edit the sub-capacity report on the host using a text editor, or on the desktop using a text editor, but this is a less convenient approach.

The sub-capacity report contains two sections: the first report section contains data that IBM Customer Support Organization (CSO) and IBM Techline will need to reconcile your company's sub-capacity bills, and is referred to as the *Billing-related section*. The second report section contains information that is useful to you when verifying the information in the Sub-Capacity Report and is referred to as the *Customer-verification section*.

4.1.1 Billing-related section

The key component of the Billing-related section, shown in Example 4-1 on page 30, is the tool's ability to report on the license requirements for each VWLC product of which it is aware (see Appendix , "Variable Workload License Charges (VWLC) products" on page 40 for the VWLC product list). The license requirements generated by the tool are derived by comparing hourly snapshots across all LPARs, across an entire month, for a particular z900 machine.

The fields presented in this section of the report, as shown in Example 4-1 on page 30, and their sources, are:

Customer information
► Customer Name - customer input (via parameter)

► Machine Serial Number - serial number for first machine found in input dataset, or customer input via parameter (optional); this is the machine to be analyzed.

► Machine Type/Model - machine type and model for machine selected above.

► Machine Rated Capacity (MSUs) - announced MSU capacity for machine selected above, derived from lookup table.

► Purchase Order Number - customer input (via spreadsheet edit), meant for customer to indicate associated product order (optional).

► Customer Comments - customer input (via spreadsheet edit), meant for general customer comments about report (optional).

Tool information
► Tool Release Date - indicator of the version of SCRT that you are utilizing to generate the report.

- Start Date/End Date - earliest and latest dates found in input dataset.
- % of Data Collected - indicator of amount of data analyzed across a month.

> **Note:** The tool takes the number of hours between the START DATE and END DATE. Then it determines how many hours are described by the SMF records in the input data set. For an hour to be counted as described, it must have both an SMF 70 and an SMF 89 for that hour. It does not matter if there is a 70 and 89 record from LPAR1, LPAR2 and LPAR3, or only from one of those systems.

- Justification for Low Data Collection - customer input via spreadsheet edit, this field shown if "% data collected", above, is less than 95% (required). The possible reasons this can happen are:
  - If you take down all of your LPARs on the same day one day a week for maintenance, that means you won't have SMF 70 records for 4 days. That means your month of SMF data will only have records for 26 days (26/30 = 86.6%)
  - Likewise, if you have a test LPAR that you only run occasionally, and have never activated SMF data collection on that test LPAR (because looking at RMF data isn't meaningful for you), the report will show 0% data collected for that LPAR. Keep in mind, however, that you are using SMF records from *all partitions in that CEC; the percentage you see is related to data from all of them together.*
- Any MVS or OS/390? - indicator of whether MVS or OS/390 was detected; value can be either YES or NO.
- Any z/OS running as z/VM Guest? - indicator of whether z/OS was detected as running under VM or z/VM; value can be either YES or NO
- Manual z/VM Calculation Completed? - customer input via spreadsheet, this field is shown if "z/OS running as VM Guest", is = "YES" (required). See "z/VM Considerations and manual calculations" on page 33.

Product summary information

- VWLC Product Name - listing of all VWLC products (by name), derived from lookup table.
- VWLC Product ID - listing of all VWLC products (by PID), derived from lookup table.
- Tool MSUs - the license requirements, in MSUs, for each VWLC product.
  - For VWLC products that produce SMF 89 records, the are MSUs calculated by the SCRT.
  - For VWLC products that do not produce SMF 89 records, the SCRT set the value to N/A.
- Customer MSUs - customer input via spreadsheet, you may override the value shown in Tool MSUs (optional).
  - For VWLM products that produce SMF 89 records, an example of an override is shown in Example 4-1 on page 30 for CICS TS for OS/390.
  - For VWLC products that do not produce SMF 89 records, you must specify here the MSU value of the LPARs the products run in.
- Customer Comment - your input via spreadsheet edit, to be filled in if you indicate a value in Customer MSUs field.
  - For VWLM products that produce SMF 89 records, you must input a comment if you override the Tool MSU value with a new value, as shown in Example 4-1 on page 30 for CICS TS for OS/390.
  - For VWLC products that do not produce SMF 89 records, you may specify here the reason for setting the MSU value of the LPARs the products run in.

Example 4-1   Sub-capacity report - Billing-related section

============== SUB-CAPACITY REPORT ==============			
Run Date/Time	05 Sept 2001 - 12:46		
Name of Person Submitting Report	XXXXX		
E-Mail Address of Report Submitter:	XXXXX		
Phone Number of Report Submitter:	XXXXX		
Customer Name	THEBIGBANK		
Machine Serial Number	12345		
Machine Type and Model	2064-111		
Machine Rated Capacity (MSUs)	350		
Purchase Order Number	XXXXX		
Customer Comments	XXXXX		
The data supplied in this sheet will be used to adjust your billing for all VWLC Programs			
on this machine. In accordance with our agreement, IBM will treat a change in product licensed			
capacity as an order. If the MSUs have increased since the last report, your billing will increase.			
For any VWLC Program, the total number of MSUs should not exceed the rated machine capacity.			
Note: This report is expected to provide a "% data collected" > 95% and data reporting period			
beginning on the 2nd of the previous month and ending on the 1st of the current month.			
===			
TOOL INFORMATION			
Tool Release Date	14 Sept 2001		
Data Start Date	02 Sept 2001		
Data End Date	01 Oct 2001		
% Data Collected	99% for 30 days		
Any MVS or OS/390?	NO	<- If yes, ineligible for sub-CEC	
Any zOS running as VM guest?	NO	<- If yes, manual calc req'd	
===			
PRODUCT SUMMARY INFORMATION			

VWLC Product Name	VWLC Product I	Tool MSUs	Customer MSUs	Customer Comments
z/OS V1	5694-A01	332	XXXXX	
DB2 for MVS/ESA V4	5695-DB2	0	XXXXX	
DB2 UDB for OS/390 V7	5675-DB2	0	XXXXX	
DB2 UDB for OS/390 V6	5645-DB2	132	XXXXX	
IBM DB2 for OS/390 Version 5	5655-DB2	0	XXXXX	
QMF MVS Version 3	5706-254	0	XXXXX	
CICS TS for OS/390 V2	5697-E93	0	XXXXX	
CICS TS for OS/390	5655-147	132	100	relocated CICS
CICS/ESA V4	5655-018	0	XXXXX	
MQSeries MVS/ESA	5695-137	0	XXXXX	
MQSeries for OS/390	5655-F10	0	XXXXX	
MQSERIES FOR OS/390 V.2.1	5655-A95	200	XXXXX	
IMS V7	5655-B01	0	XXXXX	
IMS/ESA V6	5655-158	200	XXXXX	
IMS/ESA V5	5695-176	0	XXXXX	
Netview Perf Mon	5655-043	N/A	XXXXX	
Tivoli NetView	5697-B82	N/A	XXXXX	
OPC	5697-OPC	N/A	XXXXX	
Lotus Domino for S/390	5655-B86	N/A	XXXXX	
System Automation for OS/390	5645-005	N/A	XXXXX	
System Automation OS/390 V2	5645-006	N/A	XXXXX	
COBOL for OS/390 & VM V2	5648-A25	N/A	XXXXX	
VA PL/I for OS/390	5655-B22	N/A	XXXXX	
IBM Enterprise PL/I for z/OS and OS/390	5655-G31	N/A	100	From LPAR A
IBM Enterprise Cobol for z/OS and OS/39	5655-G53	N/A	32	From LPAR B
===				

4.1.2 Customer-verification section

The customer verification section contains information, derived by the tool, on the amount of MSUs utilized by the products that generate SMF 89 records. This section is shown in Example 4-2 on page 32.

Detail data sections

The detail LPAR data section shows changes in the defined capacity of LPARs as follows:

- ► Interval - listing of all intervals in which a change in the defined capacity for a partition was reported.
- ► LPAR - all partitions encountered by the tool and their capacity.
- ► MAX - maximum capacity defined (or calculated) for the machine.
- ► MAX Interval - the first interval that reported the maximum capacity.

> **Note:** There must be an SMF 70 record sometime for an LPAR over the month. The MSU data for LPARx can only come from SMF records written by LPARx, not from data available to other images. Therefore, in order to qualify for VWLC, you must activate SMF and RMF (or equivalent) on all LPARs, for at least some of the time during the month.

Product summary drill-down section

This section shows the intervals that reported the highest values (sum of all defined capacity for all partitions executing the product during the reported interval) for the VWLC products as follows:

- ► VWLC Product Name - listing of all VWLC products (by name), derived from lookup table.
- ► VWLC Product ID - listing of all VWLC products (by PID), derived from lookup table.
- ► Highest - the highest value encountered for the product. If the product does not create a SMF 89 record, the value is zero.

> **Note:** If the input data set does not contain an SMF 89 record for a product, the value is zero.

- ► Date and Time - the date and time that value was reached.
- ► 2nd highest - the second highest value encountered for the product. If the product does not generate a SMF 89 record, the value is zero.

> **Note:** If the input data set does not contain an SMF 89 record for a product, the value is zero. If the MSU values do not change over the entire month, the 2nd Highest MSU value reported is 0, and the Date/Time field will be blank. If multiple intervals all contain the same MSU values, the first interval where that value occurs is reported, but not any others.

- ► Date and Time - the date and time that value was reached.

Product grid snapshot

This section shows in which LPAR the program was detected as being executed, at the beginning of the interval and at the end of the interval, as follows:

- ► VWLC Product Name - listing of all VWLC products (by name), derived from lookup table.
- ► VWLC Product ID - listing of all VWLC products (by PID), derived from lookup table.
- ► LPARx - LPAR name where the product was found to have been executed.

Example 4-2   Sub-capacity report - Customer verification sections

```
================================================================
DETAIL DATA SECTIONS - FOR CUSTOMER ANALYSIS PURPOSES ONLY
================================================================
DETAIL LPAR DATA SECTION (Only Intervals with Changes Are Shown)
```

Interval	LPAR A	LPAR B	LPAR C	
02 Sept 2001 - 00:00	75	32	175	
10 Sept 2001 - 00:00	100	32	175	
15 Sept 2001 - 00:00	100	32	200	
01 Oct 2001 - 23:45				
MAX ==>>	332			
MAX Interval ==>	15 Sept 2001 - 00:00			

```
================================================================
PRODUCT SUMMARY DRILL DOWN
```

WLC Product Name	WLC Product	HIGHEST	DATE/TIME	2nd HIGHEST	DATE/TIME
z/OS V1	5694-A01	332	15 Sept 2001 - 00:00	307	10 Sept 2001 - 00:00
DB2 for MVS/ESA V4	5695-DB2	0		0	
DB2 UDB for OS/390 V7	5675-DB2	0		0	
DB2 UDB for OS/390 V6	5645-DB2	132	10 Sept 2001 - 00:00	107	02 Sept 2001 - 00:00
IBM DB2 for OS/390 Version 5	5655-DB2	0		0	
QMF MVS Version 3	5706-254	0		0	
CICS TS for OS/390 V2	5697-E93	0		0	
CICS TS for OS/390	5655-147	132	12 Sept 2001 - 00:00	100	12 Sept 2001 - 00:00
CICS/ESA V4	5655-018	0		0	
MQSeries MVS/ESA	5695-137	0		0	
MQSeries for OS/390	5655-F10	0		0	
MQSERIES FOR OS/390 V.2.1	5655-A95	200	15 Sept 2001 - 00:00	175	02 Sept 2001 - 00:00
IMS V7	5655-B01	0		0	
IMS/ESA V6	5655-158	200	15 Sept 2001 - 00:00	175	02 Sept 2001 - 00:00
IMS/ESA V5	5695-176	0		0	

```
================================================================
PRODUCT GRID SNAPSHOT
Data Start -
02 Sept 2001
```

		LPAR A	LPAR B	LPAR C	
z/OS V1	5694-A01	x	x	x	
DB2 for MVS/ESA V4	5695-DB2				
DB2 UDB for OS/390 V7	5675-DB2				
DB2 UDB for OS/390 V6	5645-DB2	x	x		
IBM DB2 for OS/390 Version 5	5655-DB2				
QMF MVS Version 3	5706-254				
CICS TS for OS/390 V2	5697-E93				
CICS TS for OS/390	5655-147		x		
CICS/ESA V4	5655-018				
MQSeries MVS/ESA	5695-137				
MQSeries for OS/390	5655-F10				
MQSERIES FOR OS/390 V.2.1	5655-A95			x	
IMS V7	5655-B01				
IMS/ESA V6	5655-158			x	
IMS/ESA V5	5695-176				

Data End -
01 Oct 2001

		LPAR A	LPAR B	LPAR C	
z/OS V1	5694-A01	x	x	x	
DB2 for MVS/ESA V4	5695-DB2				
DB2 UDB for OS/390 V7	5675-DB2				
DB2 UDB for OS/390 V6	5645-DB2	x	x		
IBM DB2 for OS/390 Version 5	5655-DB2				
QMF MVS Version 3	5706-254				
CICS TS for OS/390 V2	5697-E93	x			
CICS TS for OS/390	5655-147				
CICS/ESA V4	5655-018				
MQSeries MVS/ESA	5695-137				
MQSeries for OS/390	5655-F10				
MQSERIES FOR OS/390 V.2.1	5655-A95			x	
IMS V7	5655-B01				
IMS/ESA V6	5655-158			x	
IMS/ESA V5	5695-176				

4.2 How to change information in the sub-capacity report

The main function of the software asset manager is to verify whether the data reported by the sub-capacity report is accurate. The SCRT tool is capable of accurately compiling all the data available, but there are some instances when the data is not available. These instances are:

► Products not generating SMF89 records

Certain VWLC products which qualify for sub-capacity pricing do NOT generate SMF89 records; if that is the case the tool is unable to track the usage of the product. VWLC products that are in this category are shown in Figure 4-1. See "Manual product calculations for certain VWLC PIDs" on page 36.

```
5655-043 - Netview PM V2 MVS
5645-005 - System Automation for MVS V1
5645-006 - System Automation for OS/390 V2
5655-B86 - Lotus Domino V5 for S/390
5655-G31 - IBM Enterprise PL/I for z/OS and OS/390 V3R1
3655-G53 - IBM Enterprise Cobol for z/OS and OS/390 V3R1
5697-B82 - Tivoli NetView for OS/390
5697-OPC - TME10 OPC V2
5648-A25 - Cobol for OS/390 & VM V2
5655-B22 - VA PL/1 for OS/390 V2
```

Figure 4-1   VWLC products that do not generate SMF89 records

► z/OS is running under z/VM

If the tool analyzes data from a z900 machine running z/OS as a guest of z/VM, the tool reports on each z/OS guest, but it is unable to determine if all the guests are running under a single instance of z/VM or under multiple instances of z/VM. Since customers are required to pay for MVS-based middleware on MVS-based guests running in z/VM LPARs, these guests need to be considered when making sub-capacity decisions.

The tool outputs information about MVS-based guests of z/VM in the Sub-Capacity Report but place these guests to the right of the other LPAR calculations. Manual calculations in this scenario are required. See "z/VM Considerations and manual calculations" on page 33.

It is the software asset manager's responsibility to edit the sub-capacity report and input the information he believes is correct. The software asset manager should not alter the values in the column titled *Tool MSUs*, instead he should only edit the columns titled *Customer MSUs* and *Customer Comments (see* Example 4-1 on page 30). Customer MSUs should override Tool MSUs during the billing process.

4.2.1 z/VM Considerations and manual calculations

If SCRT analyzes data from a z900 machine running z/OS as a guest of z/VM, it reports on each z/OS guest but it is unable to determine if all the guests are running under a single instance of z/VM or under multiple instances of z/VM. Since customers are required to pay for MVS-based middleware on MVS-based guests running in z/VM LPARs, these guests need to be considered when making sub-capacity decisions. SCRT outputs information about MVS-based guests of z/VM in the Sub-Capacity Report, but places these guests to the right of the other LPAR calculations. Manual calculations in this scenario are required.

Software Pricing in a z/VM and z/OS Environment on z900

The picture in Figure 4-2 shows a z900 server configured with 4 LPARs, as follows:

► LPAR A runs z/OS natively, with CICS and DB2.
► LPAR B runs z/VM with one z/OS guest.

- ► LPAR C runs z/VM with two z/OS guests and one Linux guest.
- ► LPAR D is not discussed in this example because sub-capacity pricing does not apply in a Linux environment.

Software is priced based on the size of the LPAR, even in a z/VM environment, regardless of the number of guests running under z/VM.
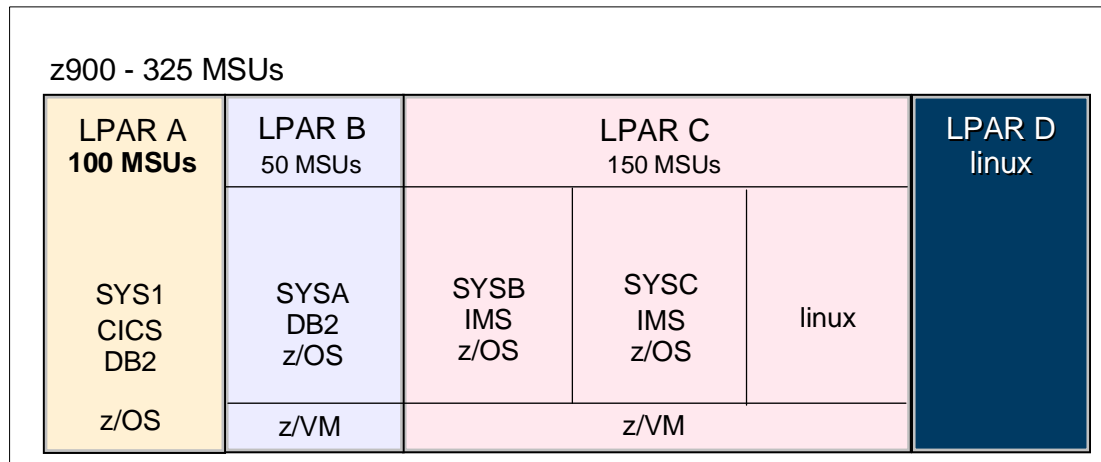

Figure 4-2   z/900 sub-capacity pricing and z/VM

In Figure 4-2, the LPAR-based view is:

z/OS   LPAR A + LPAR B + LPAR C

CICS   LPAR A

DB2   LPAR A + LPAR B

IMS   LPAR C

z/VM   Flat, regardless of placement on z900 (not covered by SCRT)

One other complicating factor is that z/VM does not support the concept of a defined capacity. Thus the mechanism that is used to determine the software capacity of z/VM guests is as follows:

- ► If they are dedicated then it is number of engines available to the LPAR times engine size.
- ► If there is a hard cap then it is the hard cap as a percentage of the shared pool.

Important: For each LPAR shown in Figure 4-2, the SCRT calculates the MSU value based on the number of engines assigned to the LPAR. See Figure 4-3 on page 35 for the MSU values calculated by the SCRT for each LPAR.

Using the MSU values calculated by the SCRT, the following licenses are required:

z/OS   LPAR A (100) + LPAR B (50) + LPAR C (150) = 300 MSUs

CICS   LPAR A (100) = 100 MSUs

DB2   LPAR A (100) + LPAR B (50) = 150 MSUs

IMS   LPAR C (150) = 150 MSUs

z/VM   Flat, regardless of placement on z900 (not covered by SCRT)

Sub-capacity report in a z/VM environment on a z900

The Sub-Capacity Report Tool *sees* each instance of z/OS running on the z900 server, regardless of whether that z/OS runs natively (as is the case with LPAR A) or whether z/OS runs as a guest of z/VM (as is the case with LPARs B and C). When computing summary data, the tool will **only** include native z/OS instances. z/OS instances running as guests of z/VM need to be manually added to the summary data provided by the SCRT.

Note: VM guest are identified by an asterisk (*) before the system name, as shown in Figure 4-3.

==								
PRODUCT SUMMARY INFORMATION								
VWLC Product Name	VWLC Product ID	Tool MSUs	Customer MSUs	Customer Comments		*SYSA	*SYSB	*SYSC
z/OS V1	5694-A01	100	XXXXX			50	150	150
DB2 for MVS/ESA V4	5695-DB2	0	XXXXX					
DB2 UDB for OS/390 V7	5675-DB2	0	XXXXX					
DB2 UDB for OS/390 V6	5645-DB2	100	XXXXX			50		
IBM DB2 for OS/390 Version 5	5655-DB2	0	XXXXX					
QMF MVS Version 3	5706-254	0	XXXXX					
CICS TS for OS/390 V2	5697-E93	0	XXXXX					
CICS TS for OS/390	5655-147	100	XXXXX					
CICS/ESA V4	5655-018	0	XXXXX					
MQSeries MVS/ESA	5695-137	0	XXXXX					
MQSeries for OS/390	5655-F10	0	XXXXX					
MQSERIES FOR OS/390 V.2.1	5655-A95	0	XXXXX					
IMS V7	5655-B01	0	XXXXX				150	150
IMS/ESA V6	5655-158	0	XXXXX					
IMS/ESA V5	5695-176	0	XXXXX					

Figure 4-3   SCRT product summary information - z/VM environment

Therefore, in the example shown in Figure 4-3, the SCRT summarizes what it *sees* in LPAR A and indicates these findings in the column titled *Tool MSUs*. The SCRT also *sees* the guests running in LPARs B and C. Each guest runs a copy of z/OS and is reported separately. Unfortunately, software pricing is driven by LPAR (not by guest) and the SCRT cannot determine which guests belong to which LPARs. Hence, you must determine that SYSA runs in LPAR B, and SYSB and SYSC run in LPAR C.

You can tell from the report that LPAR B (SYSA) runs z/OS and DB2. You can also tell from the report that LPAR C (SYSB and SYSC) runs z/OS and IMS. Products running under SYSA, in LPAR B, should be licensed at the capacity of LPAR B which is 50 MSUs. You can determine the capacity of LPAR B by examining the guests running in LPAR B. Products running under SYSB and SYSC, in LPAR C, should be licensed at the capacity of LPAR C which is 150 MSUs. You can determine the capacity of LPAR C by examining the guests running in LPAR C.

Updating the sub-capacity report in a z/VM environment on a z900

The Figure 4-4 on page 36 represents a correctly updated report:

PRODUCT SUMMARY INFORMATION							
VWLC Product Name	VWLC Product ID	Tool MSUs	Customer MSUs	Customer Comments	*SYSA	*SYSB	*SYSC
z/OS V1	5694-A01	100	300	added VM guests	50	150	150
DB2 for MVS/ESA V4	5695-DB2	0					
DB2 UDB for OS/390 V7	5675-DB2	0					
DB2 UDB for OS/390 V6	5645-DB2	100	150	added VM guests	50		
IBM DB2 for OS/390 Version 5	5655-DB2	0					
QMF MVS Version 3	5706-254	0					
CICS TS for OS/390 V2	5697-E93	0					
CICS TS for OS/390	5655-147	100					
CICS/ESA V4	5655-018	0					
MQSeries MVS/ESA	5695-137	0					
MQSeries for OS/390	5655-F10	0					
MQSERIES FOR OS/390 V.2.1	5655-A95	0					
IMS V7	5655-B01	0	150	added VM guests		150	150
IMS/ESA V6	5655-158	0					
IMS/ESA V5	5695-176	0					

Figure 4-4   Updating the sub-capacity report manually

You should not alter the values in the column titled "Tool MSUs", instead you should only edit the columns titled "Customer MSUs" and "Customer Comments". Customer MSUs should override Tool MSUs during the billing process.

4.2.2 Manual product calculations for certain VWLC PIDs

The SCRT only *sees* products which generate SMF89 records. Certain VWLC products which qualify for sub-capacity pricing do **not** generate SMF89 records. Hence, you must manually calculate the license requirement for these VWLC products.

The VWLC products which do not have SMF89 records are shown in Figure 4-1 on page 33.

> **Attention:** The process to do the manual calculation is the same as described in "z/VM Considerations and manual calculations" on page 33. For each product that does not create a SMF 89 record, you have to determine the MSU values for each LPAR that they concurrently executed on and add this sum to the report manually. If they only execute in one LPAR at a time, and over the month they execute in all three LPARs, you must add the highest MSU value from the three LPARs to the report

The SCRT will place a "N/A" in the column titled "Tool MSUs" for products which do not generate SMF89 records, as shown in Example 4-1 on page 30.

You should not alter the values in the column titled *Tool MSUs*, instead they should only edit the columns titled *Customer MSUs* and *Customer Comments*. Customer MSUs should override Tool MSUs during the billing process.

4.3 How to send the sub-capacity report to IBM

After reviewing all the information in the sub-capacity report, the software asset manager must generate a CSV flat file, like the one generated by the tool, with the changes and comments he added when performing the task described in the previous sections. This is easily done using the spreadsheet packages mentioned above, or by saving the edit file. To qualify for sub-capacity pricing, the software asset manager must submit the Sub-Capacity Report to IBM via e-mail. The IBM e-mail address, for each Country, can be found in the following URL address: *http://www.ibm.com/servers/eserver/zseries/wlc_lm/scrt.html*. IBM will acknowledge that the file has been received.

4.3.1 SCRT report checklist

Before sending your report to IBM, be sure of the following:

1. Was the report generated using the most current level of the SCRT?

2. Was the report generated using the minimum set of SMF data as input?

    a. Was the data collected from the 2nd of last month through the 1st of this month?

    b. Was at least 95% of the data from that period?

3. If the answer to #2 is no, did you provide a technical explanation for the missing data?

4. Did you supply the required information not generated by the tool?

    a. Name of person submitting report, email and phone.

    b. Product Order (PO) Number (if required by your business accounting rules).

5. Have you reviewed all PDCs generated by the SCRT, and provided alternate values where appropriate?

    a. If you are running z/OS as a guest under z/VM, have you provided PDCs for products that run in those guests to account for the size of the z/VM LPAR?

    b. To account for any periods during which the SMF data collected did not represent your normal or intended use of the products (e.g. for problem determination or disaster recovery testing purposes)?

6. Have you provided PDCs for any VWLC products that do not generate SMF type 89 records, for which you are licensed and which you have used during this reporting period?

7. Are you licensed for all the products for which a PDC is specified on this report?

8. Have you already submitted a report for this reporting period?

9. Have any conditions been detected by the SCRT that would indicate this machine is not eligible for sub-capacity charges?

    a. MVS or OS/390 are not running on this machine

    b. z/OS is not being run in 31-bit mode on this machine

4.4 How IBM handles the report

The IBM CSO will receive the file and acknowledge its receipt to the software asset manager. The sub-capacity report provides a percentage of data collected based upon number of missing records in a given timeframe. The percentage acceptable by IBM is 95%. If that is the case the file received will be reviewed by the CSO and sent for billing.

IBM will accept reports in Lotus 1-2-3, Microsoft Excel or plain text format.

A

Products eligible for license charges

This appendix provides information on which products are elegible for Variable Workload License Charge or eligible for Flat Workload License Charges.

The following topis are presented:

► Variable Workload License Charges (VWLC) Products

► Flat Workload License Charges (FWLC) Products

► Sub-Capacity Value Unit Products

Variable Workload License Charges (VWLC) products

Example 4-3 VWLC Products

```
5694-A01 - z/OS V1
5695-DB2 - DB2 V4
5655-DB2 - DB2 V5
5645-DB2 - DB2 V6
5675-DB2 - DB2 V7
5655-018 - CICS V4
5655-147 - CICS TS V1
5697-E93 - CICS TS V2
5655-B01 - IMS V7
5655-158 - IMS V6
5695-176 - IMS V5
5695-137 - MQSeries V1 MVS/ESA
5655-A95 - MQSERIES FOR OS/390 V2.1
5655-F10 - MQSERIES FOR OS/390 V5R2
5648-A25 - Cobol for OS/390 & VM V2
5655-B22 - VA PL/I FOR OS/390 V2
5706-254 - QMF V3 MVS
5697-B82 - TME10 NETVIEW FOR OS/390
5697-OPC - TME10 OPC V2
5655-043 - Netview PM V2 MVS
5645-005 - System Automation for MVS V1
5645-006 - System Automation for OS/390 V2
5655-B86 - Lotus Domino V5 for S/390
5655-G31 - IBM Enterprise PL/I for z/OS and OS/390 V3R1
3655-G53 - IBM Enterprize Cobol for z/OS and OS/390 V3R1
```

Flat Workload License Charges (FWLC) products

These are some examples of Flat Workload License Charge products:

Example 4-4 FWLC Products

```
5655-007 - NETVIEW V3R1
5688-197 - AD/C COBOL/370 V1
5668-958 - VS COBOL II V1
5685-111 - NETVIEW V2 MVS/ESA-CS
5685-083 - CICS V3 MVS/ESA
5695-007 - OPC/A MVS/ESA V1
5685-DB2 - DB2 V3
5685-016 - NETVIEW DM V1 MVS/XA
5688-235 - AD/C PL/I MVS/VM V1
5685-093 - IMS/VS DB TOOLS V2
5688-218 - CSP/370 AD V4
5665-366 - SDF/II
5648-092 - APPLICATION SYSTEM V
5685-108 - NETVIEW FTP V2 MVS
5668-806 - FORTRAN/390
5695-171 - INFO V6 MVS
5655-103 - DITTO/ESA FOR MVS V1
5688-008 - ESCON
5668-910 - OS PL/1 V2R1 COMP.&L
5668-909 - OS PL/I V2
```

```
5688-121 - TPNS V3
5695-013 - IBM COMPILER FOR REXX
5655-002 - ASF V3 FOR MVS
5695-101 - E.P.D.M.
5695-014 - IBM LIBRARY FOR REXX
5655-A17 - IBM SMARTBATCH FOR OS/390
5688-228 - APL2 MVS/VM V2
5688-216 - C/370 AD/C COMP.
5688-188 - C/370 LIBR.
5685-106 - OV/MVS V1
5695-010 - CICS VSAM RECOVERY /ESA V2
5688-206 - CSP/370RS V4
5648-142 - INFO V7 MVS
5685-013 - IMS/ESA TM V4
5688-087 - VS FORTRAN COMP & LIBR
5655-102 - DB2 PM V4
5695-041 - FAF V2
5665-290 - DISOSS V3 MVS
5685-151 - AOC V1 V1
5695-042 - IODM/2 V2
5648-A12 - IXFP SNAPSHOT
5655-A23 - DPROPR Capture for MVS V5.1
5685-101 - DW/370 MVS CICS V2
5695-081 - CICSPLEX SM/ESA V1
5688-132 - COMM SUBSYS FOR INTER
```

Sub-Capacity value unit pricing

These are the tools that are elegible for sub-capacity unit pricing:

Example 4-5   Value Unit Pricing Tools

IBM DB2 Tools
► IBM DB2 Administration for z/OS, V2
► IBM DB2 Archive Log Compression Tool for z/OS, V1
► IBM DB2 Automation Tool for z/OS, V1
► IBM DB2 Bind Manager for z/OS, V2
► IBM DB2 Change Accumulation Tool for z/OS, V1
► IBM DB2 DataPropagator for OS/390, V7
► IBM DB2 Diagnostic and Recovery Utilities for OS/390 and z/OS, V7.1
► IBM DB2 Forms for OS/390, V3
► IBM DB2 High Performance Unload for OS/390, V1.1
► IBM DB2 Log Analysis Tool for OS/390, V1.1
► IBM DB2 Object Comparison Tool for z/OS, V1.1
► IBM DB2 Object Restore Tool for OS/390, V1
► IBM DB2 Operational Utilities for OS/390 and z/OS, V7.1
► IBM DB2 Path Checker for z/OS, V1
► IBM DB2 Performance Monitor for OS/390, V7
► IBM DB2 Query Monitor for OS/390, V1
► IBM DB2 Recovery Manager for OS/390, V1
► IBM DB2 Row Archive Manager for OS/390, V1
► IBM DB2 SQL Performance Analyzer for OS/390, V1
► IBM DB2 Table Editor for z/OS, V4.1
► IBM DB2 Utilities Suite for OS/390 and z/OS, V7.1
► IBM DB2 Web Query Tool for OS/390, V1

IBM IMS Tools

- ► IBM IMS Advanced ACB Generator for OS/390, V1
- ► IBM IMS Batch Terminal Simulator for OS/390
- ► IBM IMS Command Control Facility for z/OS, V1
- ► IBM IMS Connect for OS/390, V1
- ► IBM IMS Database Control Suite for z/OS, V2
- ► IBM IMS Database Repair Facility for OS/390, V1.1
- ► IBM IMS DEDB Fast Recovery for OS/390, V2
- ► IBM IMS Dynamic Resource Control Facility V1.2
- ► IBM IMS ETO Support for OS/390, V2
- ► IBM IMS Fast Path Basic Tools for OS/390, V1
- ► IBM IMS Fast Path Online Tools for z/OS, V2
- ► IBM IMS Hardware Data Compression - Extended for z/OS, V2
- ► IBM IMS High Performance Load for OS/390, V1
- ► IBM IMS High Performance Pointer Checker for OS/390, V1
- ► IBM IMS High Performance Prefix Resolution for OS/390, V1
- ► IBM IMS High Performance Sysgen Tools for z/OS, V1
- ► IBM IMS High Performance Unload for OS/390, V1
- ► IBM IMS Image Copy Extensions for OS/390, V1
- ► IBM IMS Index Builder for OS/390, V2
- ► IBM IMS Library Management Utilities for OS/390, V1
- ► IBM IMS Message Format Services Reversal Utilities for z/OS, V1
- ► IBM IMS Online Recovery Service for z/OS, V1
- ► IBM IMS Parallel Reorganization for OS/390
- ► IBM IMS Performance Analyzer for OS/390, V3
- ► IBM IMS Program Restart Facility for OS/390, V2
- ► IBM IMS Queue Control Facility for OS/390, V1
- ► IBM IMS Recovery Saver for OS/390, V2
- ► IBM IMS Sequential Randomizer Generator for OS/390, V1
- ► IBM IMS Workload Router V2.2

IBM AD Tools
- ► IBM Fault Analyzer for z/OS and OS/390, V2
- ► IBM File Manager for z/OS and OS/390, V2
- ► IBM Breeze for SCLM for z/OS
- ► IBM Cloud 9 for SCLM for z/OS
- ► IBM ISPF's SCLM
- ► IBM z/OS SCLM Suite

SCRT Messages

This appendix contains all messages issued by SCRT. A brief explanation follows each message; some of them are self-explanatory.

****** SUB-CAPACITY REPORTING TOOL MM/DD/YY ******

Explanation: This is the first message written to SYSPRINT. MM/DD/YY is the date the report was generated.

SCRTTOOL ERROR: CUSTOMER NAME PARAMETER OMITTED

Explanation: User did not supply the customer name as a parameter in the EXEC card. This is a mandatory parameter. The tool will abort its execution and exit with return code 16. No file will be generated.

SCRTTOOL ERROR: UNABLE TO OPEN SMF DD STATEMENT

Explanation: The tool was not able to open the SMF input data set. Probable user error; the DD statement was not defined or the data set defined by the SMF DD statement is invalid.

SCRTTOOL: PROCESSING COMPLETE

Explanation: The tool has completed its processing. This message is always put out independent of the success of the execution.

SCRTTOOL ERROR: UNABLE TO OPEN OUTPUT DD STATEMENT

Explanation: The tool was unable to open the output data set (CSV file). Probable user error; the DD statement was not defined or the data set defined by the OUTPUT DD card is invalid.

SCRTTOOL: TYPE SERIAL FORMAT MUST BE TTTT-SSSSS.

Explanation: The user-supplied type and serial number is not valid. This is the second parameter (optional) entered in the PARM field of the EXEC statement. The format of this parameter must be TTTT (4-digit machine type, usually 2064) and SSSSS (5-digit machine serial number), separated by a -.

SCRTTOOL: SYSID @@@@ NOT RUNNING ZOS: tttt-sssss

Explanation: The system with SYSID @@@@ was not executing z/OS. Processing continues, but the machine identified by tttt-sssss is not eligible for sub-capacity pricing.

SCRTTOOL: PROCESSOR NOT FOUND IN TABLE @@@@ VC-@@ STSI-@@@@@@

Explanation: The processor was not in the internal table hard-coded in the tool. An update to the tool is required to include the processor.

SCRTTOOL: NON-ZOS TABLE EXCEEDED. CONTACT DEVELOPMENT

Explanation: The threshold of 50 non-z/OS operating systems was encountered. The number of non-z/OS SYSIDs encountered exceeded the hard-coded limit of 50. Please contact your local IBM support.

SCRTTOOL: STORAGE EXHAUSTED. INCREASE REGION SIZE.

Explanation: GETMAIN failure; increase the region size.

SCRTTOOL: CONTACT DEVELOPMENT.

Explanation: This message is issued with message PROCESSOR NOT FOUND IN TABLE @@@@ VC-@@ STSI-@@@@@@ to indicate that IBM should be contacted.

SCRTTOOL: SYSID @@@@ HAS MISSING OR ZERO SMF70WLA.

Explanation: The field SMF70WLA is required. It should always be present for z/OS. If it is zero, or if the RMF record is too short to have it, then this message is issued. Contact IBM for possible RMF maintenance.

SCRTTOOL: APPLY PTFS

Explanation: This is issued after message SYSID @@@@ HAS MISSING OR ZERO SMF70WLA to indicate that PTFs may be required*.*

SCRTTOOL: ERROR SYSID @@@@ DUPLICATE RECORD WITH DIFFERENT WLA

Explanation: All duplicate SMF records are eliminated; however, they were found to be duplicate. Two records for the same interval specify different WLA values; processing exits with return code 16.

SCRTTOOL: PROCESSOR MUST BE Z900

Explanation: SCRT tool only supports z900 processors; processing is terminated.

SCRTTOOL: @@@@@@@@ OF @@@@@@@@ SMF @@ RECORDS PROCESSED

Explanation: The message contains the number of SMF records that have been processed. The tool will issue two messages: one for SMF 70 records and one for SMF 89 records. The first @@@@@@@@ is the number of SMF records of that type processed by the tool; the second @@@@@@@@ is the number of SMF records of that type encountered in the data set defined by the SMF DD card.

SCRTTOOL: NO TYPE @ @ DATA FOUND TO PROCESS

Explanation: No SMF record @@ was found in the input data set. If @@ is 70, then no file is generated; if @@ is 89, then a flat file is generated with partition information only; manual updates are necessary.

SCRTTOOL: NO ZOS DATA FOUND TO PROCESS

Explanation: No data was found to process. Processing finishes. Nothing is generated.

SCRTTOOL: ADDITIONAL TYPE SERIAL FOUND IN INPUT FILE:

tttt-sssss.

Explanation: The input file contains data for type-serials other than the one been processed.

SCRTTOOL: PROCESSOR tttt-sssss NOT IN LPAR MODE.

Explanation: The processor is in basic mode; processors in basic mode are not supported for sub-capacity pricing.

SCRTTOOL: CEC NOT VALID FOR SUB CAPACITY PRICING.

Explanation: This message is issued after message *PROCESSOR MUST BE z900; only z900 CECs are valid for sub-capacity pricing*.

SCRTTOOL: DETECTED z/OS IN ESA/390 MODE IN SYSID @ @ @ @

SCRTTOOL: THIS VIOLATES THE TERMS AND CONDITIONS FOR

RUNNING z/OS ON THIS

SCRTTOOL: MACHINE; z/OS IS INTENDED TO BE IN z/ARCHITECTURE

MODE.

Explanation: The system @ @ @ @ is not running in 64 bit mode and is not valid for sub-capacity pricing.

Special notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere.,The Power To Manage., Anything. Anywhere.,TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

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

PC Direct is a trademark of Ziff Communications Company in the United States and/or other

countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.