IBM

IBM System z10 Capacity on Demand

- Describes the concepts and structure of Capacity on Demand
- Explains the offerings and how they can be applied
- Provides step-by-step examples

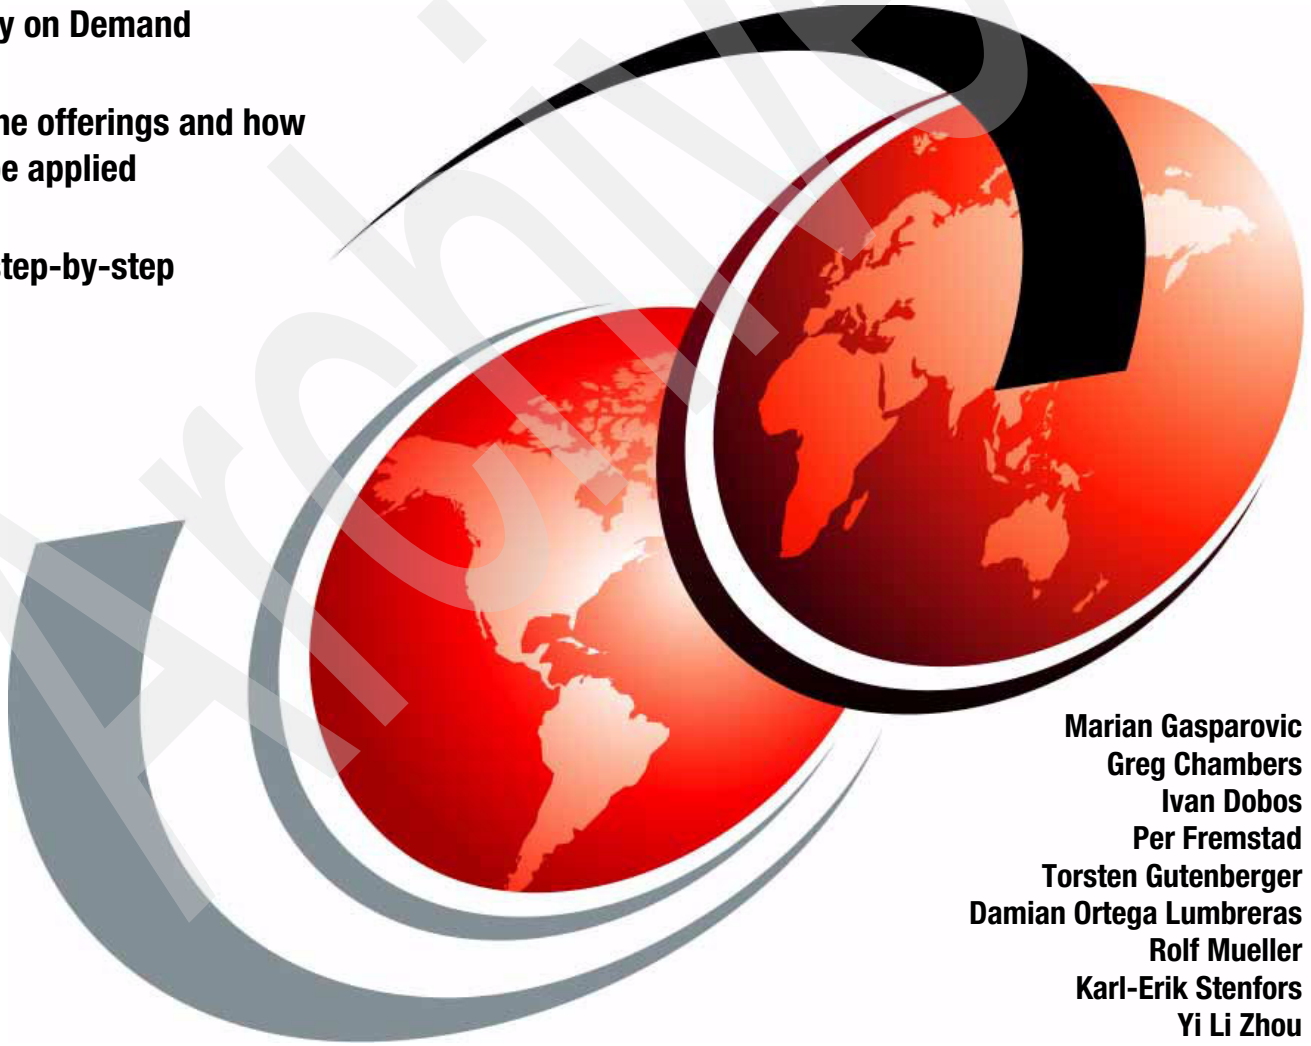

Marian Gasparovic
Greg Chambers
Ivan Dobos
Per Fremstad
Torsten Gutenberger
Damian Ortega Lumbreras
Rolf Mueller
Karl-Erik Stenfors
Yi Li Zhou

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM System z10 Capacity on Demand

January 2009

SG24-7504-01

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

Second Edition (January 2009)

This edition applies to the IBM System z10 servers.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described 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:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

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

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

CICS®
eServer™
IBM®
Parallel Sysplex®
PR/SM™
Processor Resource/Systems Manager™
RACF®
Redbooks®
Redbooks (logo) ®
Resource Link™
S/390®
System z10™
System z9®
System z®
z/Architecture®
z/OS®
z/VM®
z/VSE™
z9®
zSeries®

The following terms are trademarks of other companies:

Novell, SUSE, the Novell logo, and the N logo are registered trademarks of Novell, Inc. in the United States and other countries.

SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

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

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

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

Preface

The IBM® System z10™ servers provide a base for major server consolidation by further removing memory, processor, and channel constraints. Capacity on Demand offerings provide the flexibility to absorb temporary or permanent growth, maintenance, and disaster recovery situations in an enterprise.

The Capacity on Demand solutions offer permanent or temporary increases in processor capacity and memory. This IBM Redbooks® publication discusses the following topics:

- Understanding the common design of the various offerings
- Hardware and software areas relevant to technical planning
- Managing concurrent use of multiple Capacity on Demand offerings
- User-controlled and autonomic management of z/OS® images using z/OS Capacity Provisioning

This book is intended for systems engineers, infrastructure architects, and anyone wanting to understand Capacity on Demand functions. Readers are expected to be generally familiar with existing System z® technology and terminology.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Poughkeepsie Center.

Marian Gasparovic is an IT Specialist working for the IBM Server and Technology Group in IBM Slovakia. He worked as an Administrator for z/OS at Business Partner for 6 years. He joined IBM in 2004 as a Storage Specialist. Currently, he holds dual roles: one role is Field Technical Sales Support for System z in the CEMAAS region as a member of a team that handles new workloads; another role is for ITSO in Poughkeepsie, NY.

Greg Chambers is a Senior Certified FTSS in the Western Region of the United States. During his 30 years at IBM, he worked in various roles in service, design, and support of IBM large systems. He is a Regional Designated Specialist in multiple areas within the System z products, including performance/capacity planning, Crypto, Parallel Sysplex®, installation and migration, and Disaster Recovery. He is a member of the zChampions team. Currently, he provides pre-sales and post-sales technical support for the System z products.

Ivan Dobos is a System z IT Specialist with 10 years of experience with System z. During the past 5 years, he has worked with a large number of clients and spent most of his time supporting new workloads, z/VM® and Linux® on System z projects. Since 2005, he is based at the Products and Solutions Support Center (PSSC) Montpellier, France, where he spent 2 years as Technical Leader for Linux on System z projects in the System z Benchmark Center. He is currently working in System z New Technology Center where he specializes in IT Optimization and TCO studies. Ivan is a regular speaker at the Montpellier Executive Briefing Centre.

Per Fremstad is an IBM Certified Senior IT Specialist from the IBM Systems and Technology Group at IBM Norway. He has worked for IBM since 1982 and has extensive experience with mainframes and z/OS. Per also works extensively with Linux on System z and z/VM. During the past 25 years he has worked in various roles within IBM and with a large number of customers. He frequently teaches about z/OS and z/Architecture® subjects, and has been actively teaching at Oslo University College for the last 5 years. Per holds a BSc from the University of Oslo, Norway

Torsten Gutenberger is an IT specialist, working in the System z Front End Hardware Support in Mainz, Germany. He holds a degree in IT Engineering and has more than 12 years of experience in the s/390 and System z environment, including a three-year assignment to the zSeries® EMEA Product Support Group (EPSG) in Montpellier, France.

Damian Ortega Lumbreras is an IT Specialist, working as Field Technical Sales Support for System z in Spain. He worked for Learning Services for five years, teaching z/OS courses focusing mainly on DFSMS. He has participated in several mainframe courses in Spanish universities. Damian holds a degree in Physics from the University of Granada, Spain.

Rolf Mueller is an IT Specialist, working for IBM Techline, Germany. He works in pre-sales support, often dealing with sizing and capacity planning questions. He holds degrees in Business Administration and Informatics. Rolf has 20 years of information technology experience.

Karl-Erik Stenfors is a Senior IT Specialist in the Product and Solutions Support Centre (PSSC) in Montpellier, France. He has more than 38 years of experience in the large systems field, as a Systems Programmer, as a Consultant with IBM customers, and since 1986, with IBM. His areas of expertise include IBM System z hardware and operating systems, including z/VM, z/OS, and Linux. He teaches at numerous IBM user group and IBM internal conferences. He is currently working with the System z lab in Poughkeepsie, providing customer requirement input to create an IBM System vision for the future — the zAtlas and zChampions workgroups.

Yi Li Zhou is a Senior IT Architect in IBM China. He has 20 years of experience in the S/390® mainframe z/OS field. He holds a Bachelor degree in Computer Science from Beijing Information Engineering Institute. His areas of expertise include IT infrastructure solution design and implementation services. He has eight years of experience in assisting customers in architecting, installing, and sizing technology solutions to address business needs.

Thanks to the following people for their contributions to this project:

Robert Haimowitz, Franck Injey, Bill White
International Technical Support Organization, Poughkeepsie Center

Christine Axnix, Michael Groetzner, Horst Sinram
Systems &Technology Group, Boeblingen, Germany

Michael Allen, Brent Boisvert, Patrick Clas, Victor Lourenco, Brian Valentine
Systems &Technology Group, Endicott, NY

Michael Gerhart, Glen Poulsen
System z Product Engineering, Poughkeepsie, NY

Frank Wisnewski, Systems &Technology Group, Poughkeepsie, NY

Gregory Hutchison, IBM Sales & Distribution, Gaithersburg, MD

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

Overview

Unexpected events, such as unpredictable increases in workload, or even disasters can cause the need for IT capacity to grow tremendously. In these situations, IBM can supply the necessary capacity, quickly and simply. Capacity on Demand is the key feature that enables you to adjust your processing and memory capacity to your specific needs without having to shut down or restart the server, and without needing to re-IPL the operating system. This chapter introduces Capacity on Demand on the IBM System z10 server.

This chapter discusses the following topics:

1.1 System upgrades

The IBM System z10 servers have the capability of *concurrent* upgrades, providing additional capacity with no server outage. In most cases, with prior planning and operating system support, a concurrent upgrade can also be nondisruptive to the operating system.

In general, concurrency addresses the continuity of operations of the hardware part of an upgrade, for instance, whether a server (as a box) is required to be switched off during the upgrade. Disruptive versus nondisruptive refers to whether the running software or operating system has to be restarted for the upgrade to take an effect. Thus, even concurrent upgrades can be disruptive to those operating systems or programs that do not support them, while at the same time being nondisruptive to others.

The capabilities are based on the flexibility of the design and structure, which allows concurrent hardware installation and Licensed Internal Code (LIC) control over the configuration.

Upgrades can be *ordered* in two ways:

- MES process upgrade

  Traditional miscellaneous equipment specification (MES) process upgrade order is always performed by IBM. The result of the MES order can be either real hardware added to the server or installation of LIC Configuration Control (LICCC) to the server. In either case, installation is performed by IBM personnel.

- CIU facility upgrade

  The Customer Initiated Upgrade (CIU) facility is the IBM online infrastructure through which the customer can order upgrades for a System z server. Access to and use of the CIU facility requires the customer and IBM to sign a contract through which terms and conditions for use of the CIU facility are accepted. The use of the CIU facility for a given server requires that an online On-line CoD Buying contract is signed and that the feature is installed on the server.

  After all the prerequisites are in place, the entire process, from ordering to activation of the upgrade, is performed by the customer and does not require any on-site presence of IBM service personnel.

  The CIU facility supports LICCC upgrades only. It does not support I/O upgrades. All additional capacity required for an upgrade must be previously installed. Additional I/O drawers or I/O cards cannot be installed as part of an order placed through the CIU facility.

Upgrades can be *delivered* in two ways:

- MES upgrade

  An upgrade can be ordered with the miscellaneous equipment specification (MES) process only through an IBM representative. It delivers new hardware to be installed. MES installations require an IBM service personnel to perform the upgrade.

- CoD upgrade

  Capacity on Demand (CoD) upgrades are delivered in the LIC Configuration Control (LICCC) record. LICCC provides for server upgrade without hardware changes by activation of additional, previously installed, unused capacity.

  CoD upgrade can be *ordered* either through MES process (with limitations) or through CIU. If ordered through the MES process, IBM personnel is required to install it. If ordered through CIU, the customer may install the LICCC record without any assistance.

For different situations, different types of upgrades are necessary. After some time, you might realize that more memory is required, I/O cards have to be added, or processor more capacity is required for constantly growing workload. In other situations, only a short-term upgrade is needed, to handle a peak workload or to replace temporarily a server that is down during a data center maintenance, or a disaster situation. System z10 offers the following solutions for such situations:

- Permanent upgrade

  Permanent upgrades are requested when more channels, more memory, or more processing capacity for growing workloads are needed. A permanent upgrade can be performed either by IBM or by the customer without IBM personnel on site, as follows:

  – MES (IBM performs the upgrade)

    IBM performed upgrade is a result of an upgrade ordered through MES upgrade process. It can either add real hardware to the configuration or enable unused but present capacity through LICCC

  – On-line Permanent Upgrade (customer performs the upgrade)

    CIU permanent upgrade always results in an LICCC record. It allows you to activate unused processing units (PUs) or change the central processor (CP) capacity level. It also allows you to activate installed but unused memory. Upgrade is activated from the HMC.

- Temporary upgrade

  All temporary upgrades are LICCC-based. They can be installed either by downloading the code from IBM using CIU facility or by IBM installing the code on site, but activation is always done by the customer by using the HMC. Temporary upgrade is used in two situations: when peak in a workload has to be handled and when a capacity of another server in an enterprise has to be replaced temporarily. Therefore, two types of temporary offerings are available:

  – Billable

    To handle a peak workload, processors can be rented temporarily on a daily basis. Customer can activate up to double the purchased capacity of any PU type. The one billable capacity offering available is On/Off Capacity on Demand.

  – Replacement

    When a processing capacity is lost in another part of an enterprise, replacement capacity can be activated. It allows to activate any PU type up to authorized limit. The two replacement capacity offerings available are Capacity Backup and Capacity for Planned Events.

The remainder of this book discusses CoD upgrades; the book covers all upgrades that can be installed by a customer without IBM personnel on site.

1.2 Capacity on Demand on the System z10

The System z10 server, including z10 EC and z10 BC, introduces just-in-time deployment of additional computing capacity, known as Capacity on Demand (CoD). The functions are designed to provide more flexibility and to provide an easier way to dynamically change capacity when business requirements dictate.

Note: There is no difference in CoD functionality between z10 EC and z10 BC servers.

Additional capacity resources can be dynamically activated, either partially or in totality, by using granular activation controls directly from the management console of the z10, without having to interact with IBM Support. CoD delivers permanent or temporary capacity upgrades.

Online permanent upgrade

Up to the limits of the installed capacity on an existing server, a permanent upgrade can concurrently:

- Add processors:
  - Central processor (CP)
  - Integrated Facility for Linux (IFL) processor
  - Internal Coupling Facility (ICF) processor
  - IBM System z10 Application Assist Processor (zAAP)
  - IBM System z10 Integrated Information Processor (zIIP)
  - System assist processor (SAP)
- Add memory
- Change the subcapacity setting

A permanent upgrade is initiated by the customer using the CIU facility on IBM Resource Link™, and requires a contract to be signed. This contract covers the permanent upgrade buying capability.

Temporary upgrades

Temporary upgrades can be done by Capacity Backup (CBU), Capacity for Planned Events (CPE), or On/Off Capacity on Demand (On/Off CoD), all by using the CIU facility on IBM Resource Link. All three offerings are optional and require unused capacity to be available on the server, either as unused PUs or as a possibility to increase the CP capacity level, or both:

- Capacity Backup (CBU) is a *concurrent* and *temporary* activation of additional CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, an increase of the CP capacity level, or both. CBU is intended to replace capacity lost in the enterprise as a result of a disaster. CBU cannot be used for peak workload management. A CBU can last up to 90 days when a disaster situation occurs.

  A replacement capacity-offering contract and a contract specific to this offering must be signed. The standard CBU contract provides for five 10-day tests and one 90-day disaster activation over a five-year period. The contract period may be optionally set for one to five years.
- Capacity for Planned Events (CPE) is a *concurrent* and *temporary* activation of additional CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, an increase of the CP capacity level, or both. CPE is used to replace temporary lost capacity in an enterprise for planned downtime events, for example for data center changes. CPE is not meant to be used for peak workload management, or for a disaster situation.

  A replacement capacity-offering contract and a contract specific to this offering must be signed. The CPE provides for one 72-hour activation.
- On/Off Capacity on Demand (On/Off CoD) enables *concurrent* and *temporary* additional capacity on the server. On/Off CoD can be used for customer peak workload requirements. It has a daily hardware charge. The software charges vary according to the license agreement for the individual products; contact the IBM Software Group representative for details.

On/Off CoD can concurrently add processors (CPs, IFLs, ICFs, zAAPs. zIIPs, and SAPs), increase capacity level, or both, up to the limit of the installed capacity of an existing server, and is restricted to twice the currently purchased capacity.

A billable (online CoD buying) capacity offering contract has to be signed, and so does a contract specific to this offering. The On/Off Capacity on Demand provides for unlimited activation.

A set of contract documents is necessary to support the different Capacity on Demand offerings. The contracts are structured in a modular, hierarchical approach. For all Capacity on Demand offerings, a base contract specifies the basic Capacity on Demand terms. In addition to the base contract, each specific offering has a contract.

Table 1-1 lists the CoD offerings. For more information about using CoD features, refer to *IBM System z10 Capacity on Demand User's Guide,* SC28-6871.

Table 1-1 Capacity on Demand summary

Type	Offering	CPC resources	Process
Permanent	Permanent Upgrade	CPs, IFLs, ICFs, zAAPs, zIIPs, SAPs, and memory	LICCC, performed through CIU facility
Temporary	On/Off CoD	CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs	LICCC, performed through CIU facility
	CBU	CPs, IFLs, ICFs, zAAPs, zIIP, and SAPs	LICCC, performed through CIU facility
	CPE	CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs	LICCC, performed through CIU facility

MES upgrades do not use the CIU facility, they require the presence of IBM service personnel on site. MES upgrades may be LICCC only or LICCC with the addition of hardware, such as books (for z10 EC), I/O drawers (for z10 BC), memory, and I/O cards. Most MES upgrades are also concurrent.

For the z10 EC, adding a book to the central processor complex (CPC) is done concurrently by using Concurrent Book Add (CBA) task. However, addition of an I/O cage to the z10 EC, upgrade to a model E64, or physical memory upgrade on a single book server are not concurrent to server operations.

For the z10 BC MES upgrades for physical memory are not concurrent. However, processor capacity upgrades are concurrent, and so are the addition of I/O drawers and I/O cards.

Value and benefits

Capacity on Demand on the z10 server brings many benefits:

- Capacity on Demand helps keep investment in hardware low. Only the hardware that is required must be ordered and installed. Additional resources can be ordered and paid for when they are needed.
- Capacity on Demand gives a predictable structure of costs in case of an upgrade.
- Activating dormant capacity nondisruptively as the workload increases is possible, minimizing the risks with unexpected growth.
- Different Capacity on Demand offerings help to fit different requirements.
- A combination of several Capacity on Demand offerings is possible, providing for a high level of flexibility.

- Existing capacities can be used for production and activated as backup in parallel. This gives a high return of investment and a high level of protection in case of a disaster.
- Preplanning and preconfiguring for a number of different scenarios are possible, enabling an easier response to a changing environment.
- Capacity on Demand controls are enhanced to help customers stay within their expense budgets. The System z10 introduces capacity tokens. Capacity tokens are designed to provide the customers with the means to limit their financial liability when utilizing On/Off Capacity on Demand. Customers are assigned capacity tokens through the order process and tokens are consumed on a 24 hours billing window basis.
- z/OS Capacity Provisioning allows to manage processing capacity more reliably, more easily, and faster. Replacing manual monitoring with autonomic management, or supporting manual operation with recommendations can ensure that sufficient processing power is available with the least possible delay.
- The z/OS Capacity Provisioning function provides management of processing capacity through policy-based automation.

Customer Initiated Upgrade (CIU) facility

The CIU facility is the IBM online infrastructure that enables you to order permanent and temporary upgrades for a System z server.

The CIU facility supports LICCC upgrades only; it does not support I/O upgrades. All additional hardware (such as memory or I/O) required for an upgrade must be previously installed. Additional books (for the z10 EC server) or I/O cards cannot be installed as part of an order placed through the CIU facility.

Ordering is accomplished by logging on to the IBM Resource Link Web site and invoking the CIU application to upgrade the server's processors or memory. Defining additional IDs to be authorized to access CIU is possible. Additional IDs can be authorized to enter or approve new orders, or to view existing orders.

Figure 1-1 on page 7 shows the CIU process. An upgrade order is placed through the CIU facility by a customer. When it is approved by an alternative customer account, the order processed. The result of ordering is an LICCC record, which is placed on the IBM Support site. The customer receives a note that the order is ready. The customer retrieves the record from IBM to the Support Element (SE). When necessary, the customer installs the retrieved records. During installation and activation, being connected to the IBM Support site is not required.

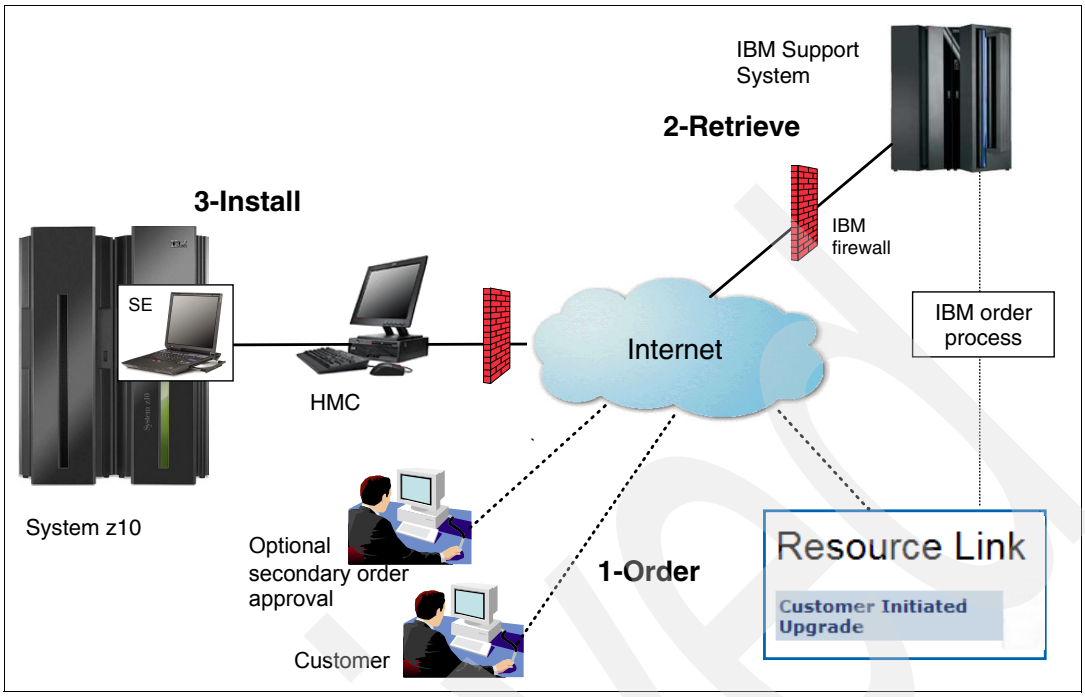

Figure 1-1 Using the Customer Initiated Upgrade facility

The actual upgrade process is fully automated and does not require on-site presence of IBM service personnel.

Access to, and use of the CIU facility requires the customer and IBM to agree and sign a contract through which terms and conditions for use of the CIU facility are accepted. The use of the CIU facility for a given server requires that the On-Line CoD Buying feature (FC 9900) is installed on the server.

To be able to use the CIU facility, the customer must have a registered user ID on Resource Link and the ID must be specified on the contract. After the contract is signed, customer support creates the machine profile and the customer is notified by e-mail when the CIU machine profile is ready. Optionally, a second ID can be specified as an approver.

1.3 Provisioning architecture

The configuration upgrades are contained in encrypted configuration *records.* These Licensed Internal Code Configuration Control (LICCC) records contain all necessary information about the upgrade. LICCC records are used every time a Capacity on Demand change to the configuration is performed, for temporary and permanent upgrades. The offerings are all built from a common LICCC record structure.

These records, referred to as Permanent or Temporary Entitlement Records (PER or TER), contain the information necessary to control which type of resource can be accessed and to what extent. They also contain time elements, such as test duration (measured in number of days) or real activation (measured in number of days). In addition, the TER may contain tokens that can represent the number of tests or the number of real activations. The TER may also contain capacity tokens, which may be used to control capacity that the customer already paid for (prepaid) or to limit the customer's liability for postpay capacity. The tokens

are reduced (consumed) in billing window intervals. More details are in “Capacity consumption and billing” on page 29.

A significant enhancement made to the CoD implementation is the ability to add permanent capacity to the server when there is temporary capacity active already. The CoD implementation includes support for:

- Multiple (up to 200) TERs can be concurrently staged on the SE.

  Staged means that the TER has been retrieved from either the Remote Support Facility or through an external media, and is loaded on the Support Element (SE) hard drive.
- Multiple (up to eight) TERs can be *installed* on the CPC and *active,* at any time.

  Installed means that the TER is promoted from the SE hard drive to one of the locations in system memory reserved for temporary records.

  Active means that resources of the TER have been enabled and are available for use.
- Variability in the number of resources that can be activated for each TER.
- The ability to control and update each TER independently of each other.
- Improved query function to monitor the state of TERs.
- Replenishment, which means that the contents of the temporary entitlement record (capacity and expiration date) can be dynamically updated.
- Permanent upgrades can be performed while a temporary upgrade is active.

Implementation allows access to additional capacity with increased flexibility.

The System z10 provides concurrent, on demand upgrades to the server hardware components. With operating system support, and appropriate planning, concurrent upgrades can be nondisruptive to the operating system.

The System z10 EC has five hardware models (E12, E26, E40, E56, E64) ranging from 12 - 64 PUs for customer use:

- Full capacity models have a capacity setting of 7nn, where nn ranges from 1 - 64.
- Subcapacity models have subcapacity settings with ranges of 401 - 412, 501 - 512, 601 - 612, which provide flexibility in capacity settings and configurations. A maximum of 12 subcapacity CPs are available.

The System z10 BC has one hardware model (E10) with 10 PUs. It can be configured with a maximum of five CPs. The configured CPs can be set to one of 26 different capacity levels resulting in a total of 130 (26 x 5) different capacity settings. The specialty engines are always at full capacity.

The following resources for the System z10 can be activated by Capacity on Demand in any amount up to the defined limit:

- Processor capacity by adding central processors, specialty engines, or both
- Processor capacity by increasing the capacity level of the CPs
- Processor capacity by adding CPs and increasing the capacity level of all CPs to the same level
- Memory capacity by enabling installed but not active memory (only a permanent upgrade)

Activation can be done manually or it can be automated through an application programming interface (API).

The Capacity on Demand implementation on the System z10 provides more flexibility, granularity, and responsiveness than previous implementations both for the customer and for IBM. In addition, an enhanced set of Capacity on Demand APIs is available for use by system management and automation software.

Figure 1-2 shows an overview of CoD. Processor capacity consists of purchased and dormant capacity. Purchased capacity consists of active and inactive processors. CoD offerings can activate dormant and purchased processors, but inactivate processors and (in a case of CPs) change the capacity level of purchased and active processors.

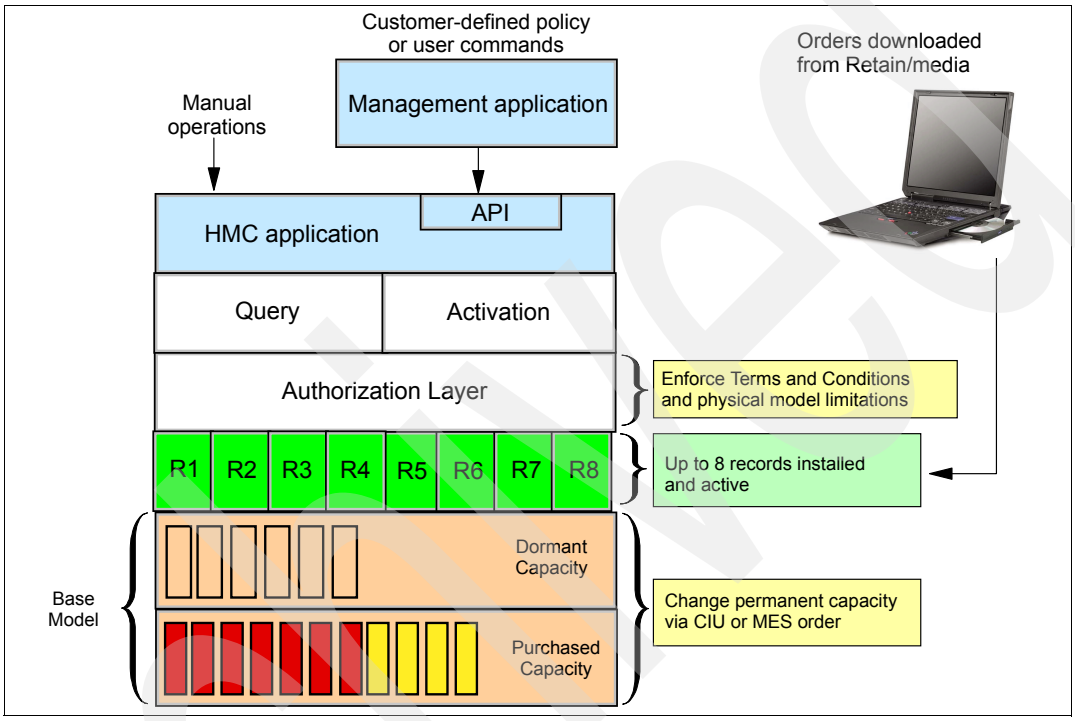

Figure 1-2 Capacity on Demand

The enhancements on System z10 include:

- Granular activation controls are available for temporary upgrades.

  The System z10 server provides the ability to select one or more processors in the temporary record for activation or deactivation. The activation level can be dynamically changed any time. It allows either increased or decreased temporary capacity, as needed.

  Rules are available for the allowable increase and decrease of processing power. The maximum levels are defined when you order the temporary records. The allowed activation and deactivation choices for each record are presented at the HMC/SE panels or as part of the Simple Network Management Protocol (SNMP) Capacity on Demand API query information.

- Temporary records are ready for self-contained activation on the system.

  Temporary capacity entitlement records are first staged on the SE and then installed into the system, available for activation at the time of need.

- Multiple temporary records can be active at the same time.

  Multiple (up to eight) temporary records can be active at the same time. The multiple record types include On/Off CoD, CBU, and CPE. Only one On/Off CoD record can be active at any given time.

- Support is available for permanent upgrades while temporary upgrade is active.

  The System z10 supports permanent upgrades while temporary capacity is active. It is possible to convert active temporary processors to permanent processors if they are the same processor type. This is done for On/Off CoD temporary processors.

  If not enough similar active On/Off CoD processors are available to match the newly requested permanent number of processors, dormant processors are made permanent. For example, if the permanent upgrade was to add three CPs with no capacity level change, and the temporary active capacity activated by the On/Off CoD record is two CPs (the same capacity level as permanent), those two temporary CPs are converted to permanent, and one dormant CP is used for the third new permanent CP.

- Temporary records staged on the SE (installed or active) remain valid after permanent upgrades and after repair actions.

  Temporary records store the allowed relative-capacity increase, they do not contain target capacity. When a temporary record is being activated, it offers possible target configurations based on a permanent capacity. For this reason, reordering temporary records after a permanent upgrade is not necessary because they are not invalidated.

- Automatic deactivate

  When an active temporary record expires (either record validity expires or activation limits expire), an automatic deactivation is performed. Automatic deactivation requires that the active temporary processor resources be defined as shared. System z10 does not deactivate dedicated processors or the last active processor of any type.

  If the automatic deactivation is not successful, deactivation will be tried every 24 hours. On/Off CoD billing will continue.

  Temporary record expiration warning messages are provided on the HMC. Resource Link sends warning e-mails to the user. The user can then choose to manually deactivate the temporary record or let automatic deactivation proceed. IBM has the right to revoke entitlement if consistent and repeated incidents occur in which resources cannot be automatically deactivated.

- Capacity on Demand automation

  Capacity on Demand automation is provided through SNMP APIs. z/OS Capacity Provisioning is a z/OS feature that utilizes these APIs.

The IBM System z10 has the capability of *concurrent* upgrades, providing additional capacity with no server outage. When operating system support is available, a concurrent Capacity on Demand upgrade can also be *nondisruptive* to the operating system.

Installation and activation of Capacity on Demand upgrades on a z10 is done through the Perform Model Conversion task on SE. The sequence is illustrated in Figure 1-3 on page 11. The Perform Model Conversion task has two paths, one for permanent and one for temporary upgrades.

From the permanent upgrade path, you may:

- Retrieve and apply processor or memory upgrade data from the IBM Support System.
- Retrieve processor or memory upgrade data, but not apply it.
- Apply processor or memory upgrade data that was previously retrieved.

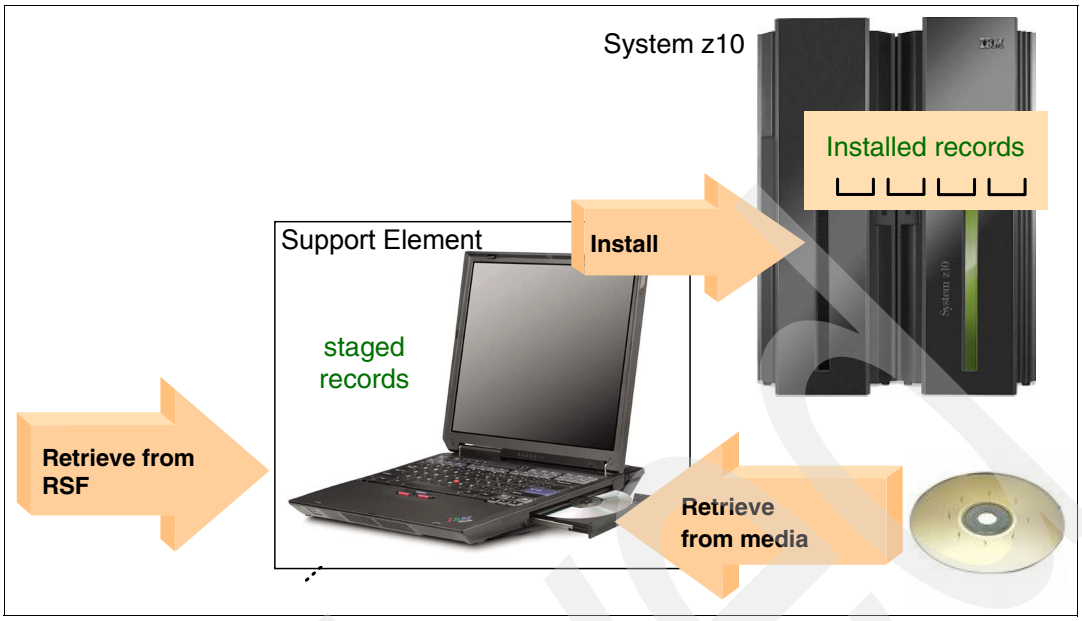

Figure 1-3 Perform Model Conversion task, temporary upgrades

From the temporary upgrade path, you have the option to retrieve or manage:

- Retrieve
  - Processor upgrade data from the IBM Service Support System
  - Processor upgrade data from media
- Manage
  - Temporary records, by installing or deleting them from the Staged Records tab
  - Processors, by adding or removing them from the Installed Records tab

1.4 z/OS Capacity Provisioning

z/OS Capacity Provisioning helps to manage the general purpose (CP), zAAP and zIIP capacity of a System z10, which is running one or more instances of z/OS. Based on the On/Off CoD offering, temporary capacity may be activated and deactivated under control of a defined policy. Combined with functions in z/OS, the System z10 provisioning capability gives the customer a flexible and automated process to control the configuration and activation of On/Off CoD records.

For detailed information, see the z/OS Capacity provisioning Web page:

`http://www.ibm.com/servers/eserver/zseries/zos/wlm/cp/`

The provisioning architecture enhances the already rich on-demand environment by opening up interfaces to the z/OS operating system through the SNMP APIs. The z/OS operating system can query which resources are in the On/Off CoD records and the status of these resources.

The provisioning of resources can be activated based on schedules or under control of automation functions by using the Workload Manager (WLM) performance indicators. Automation processes can activate resources in the On/Off CoD records as dictated by z/OS Capacity Provisioning automation policies, and deactivate the resources when they are no

longer needed, also according to the automation policies. Manual and supervised modes (where suggested changes need to be confirmed by the operator) are available.

z/OS Capacity Provisioning simplifies the monitoring of critical workloads, and its automation features can help activate additional resources faster than manual operation. When using Capacity Provisioning, different levels of automation provide an appropriate level of control.

With Capacity Provisioning, you may:

- Activate and deactivate temporary capacity through operator commands (manual mode)
- Activate and deactivate temporary capacity based on a defined *schedule*, without considering workload performance
- Have the Capacity Provisioning Manager *suggest* changes to the capacity of the System z10 server based on the observation of defined workloads. In this case, the operator has to confirm the suggested changes.
- Have the Capacity Provisioning Manager *automatically* implement changes to the capacity of the System z10 server based on the observation of defined workloads.
- Run Capacity Provisioning in analysis mode. In this mode the operator is informed when an action might occur, according to the defined rules. However, no action is taken unless the operator enters the necessary commands.

The Capacity Provisioning function included in the z/OS operating system is part of z/OS MVS Base Control Program (BCP) Release 9. The Capacity Provisioning Manager (CPM), one of the z/OS Capacity Provisioning components, monitors the workload on a set of z/OS instances. Within each z/OS instance WLM assigns resources to the workloads according to the importance defined by the administrator. Key performance indicators are derived from WLM, and if they indicate that a workload misses its goal, and the lack of processor resources is identified as the main reason, then CPM can recommend appropriate actions, or activate temporary capacity automatically. The automation and activation processes use the SNMP APIs on the HMC or SE. See Figure 1-4 on page 13.

For the observed workload, the administrator defines, in a CPM policy, which workload is critical and which criteria indicates it is suffering. Workload definition is based on WLM service classes. Each service class can be associated with a goal. Whether a workload reaches its goal is indicated by WLM through a performance index (PI) for that service class.

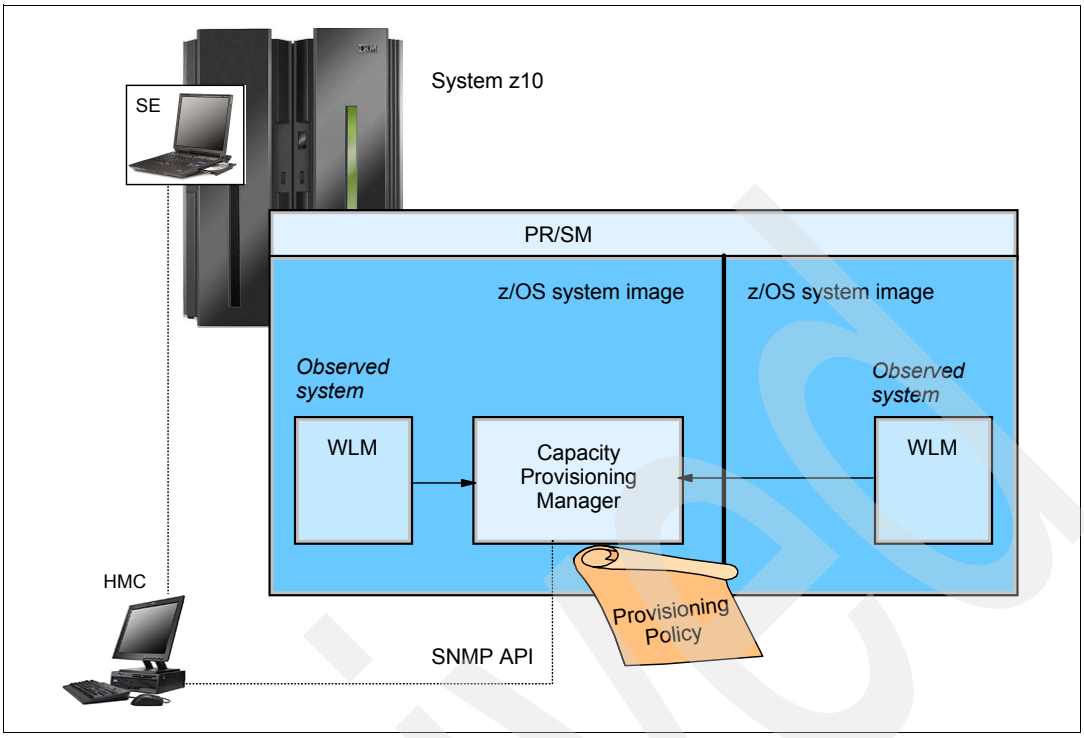

Figure 1-4 z/OS Capacity Provisioning Manager (CPM) overview

When CPM detects a workload that is suffering according to the rules in its policy, additional performance information is used to check whether it would help to increase the processing capacity. Based on the resources available to the system image, z/OS Capacity Provisioning can detect a need for general purpose (CP) or special purpose (zAAP or zIIP) capacity. General purpose capacity is also chosen depending on whether a higher capacity level would help or whether additional CPs would eliminate the resource shortage.

In the CPM policy, administrator defines which resources can be provisioned automatically, thereby limiting the amount of temporary resources made available. In the policy, the maximum number of zAAPs and zIIPs to be activated can be specified. For general purpose capacity, the maximum is specified in terms of Million Service Units (MSUs). z/OS Capacity Provisioning then automatically decides which target configuration best fits the subject workload and policy, and tries to activate it.

The z/OS CPM is one example of how system management components can exploit the Capacity on Demand SNMP APIs. Other IBM products, independent software vendor (ISV) products, or customer automation software can also benefit from the new functionality.

In previous z/OS environments, even with an ideal WLM service definition, it was not possible to achieve all specified goals when the total workload increased beyond the capacity of the server. z/OS Capacity Provisioning in concert with the On/Off Capacity on Demand provisioning capability gives a new flexible and automated process to control the configuration and activation of defined resources. CPM provides policy-based automation.

See 2.6, “Planning for z/OS Provisioning” on page 70 and for more information.

1.5 Terminology

Table 1-2 briefly describes the most frequently used terms related to Capacity on Demand.

Table 1-2 CoD terminology

Term	Description
Activated capacity	Capacity that is purchased and activated. Purchased capacity can be greater than activated capacity.
Billable capacity	Billable capacity is meant to help you to handle workload peaks, either expected or unexpected. There is one billable offering available, On/Off Capacity on Demand.
Book	A physical package on z10 EC that contains memory, a Multi-Chip Module (MCM), Memory Bus Adapter (MBA), and Host Channel Adapter (HCA) cards, and FSP cards. A book plugs into one of four slots in the central processor complex (CPC) cage of the z10 EC.
Capacity	Refers to hardware resources (processor and memory) able to process workload. Can be added to the system through various capacity offerings.
Capacity Backup (CBU)	The CBU function allows using spare capacity in a CPC to replace capacity from another CPC within an enterprise, for a limited time. Typically CBU is used when another CPC of the enterprise has failed or is unavailable because of a disaster event. The CPC using CBU replaces the missing CPC's capacity.
Capacity for Planned Events (CPE)	One of Capacity on Demand offerings, CPE is intended to be used when temporary replacement capacity is needed for a short term event.
Capacity levels	For the z10 EC server, the capacity levels for the CP engine are 7, 6, 5, and 4. Full capacity CP engine is indicated by 7; subcapacity CP engines are indicated by 6, 5, and 4. For the z10 BC server, the capacity levels for the CP engine are A - Z. Full capacity CP engine is indicated by the letter Z. Any other letter indicates the subcapacity CP engine.
Capacity setting	The capacity setting is derived from the capacity level and the number of CPs. For the z10 EC server, the capacity settings are 4xx, 5xx, 6xx, 7nn. The number of processors can have a range of: ► 0 - 64 for capacity level 7nn ► 1 - 12 for capacity settings 6xx, 5xx, 4xx For the z10 BC server, the capacity levels are Axx - Zxx. The number of CPs can be 1 - 5.
Concurrent book add (CBA)	CBA concurrently adds book hardware, including processors, physical memory, and I/O connectivity
Customer Initiated Upgrade (CIU)	An IBM Web-based facility that you use to request processor and memory upgrade through the Web, using IBM Resource Link and the system's RSF connection.
Capacity on Demand (CoD)	The ability of a computing system to increase or decrease its performance capacity as needed to meet fluctuations in demand.

Term	Description
CPC drawer	A physical package on z10 BC that contains memory, SCMs, HCA and MBA cards, FSP cards, OSC/ETR cards, and DCA cards.
Capacity Provisioning Manager (CPM)	CPM is a component of z/OS Capacity Provisioning.
Customer profile	This information resides on Resource Link and contains customer and machine information. A customer profile may contain information about more than one machine.
Dormant capacity	PUs present but not purchased. Same as uncharacterized capacity.
Enhanced book availability	In a multibook configuration (z10 EC only), the ability to have a book concurrently removed from the server and reinstalled during an upgrade or repair action.
Full capacity CP feature	For z10 EC feature (CP7) provides full capacity. Capacity settings 7xx are full capacity settings. For z10 BC feature (CPZ) provides full capacity. Capacity settings Z01, Z02, Z03, Z04, and Z05 are all full capacity settings.
High watermark (HWM)	HWM is capacity purchased and owned by the customer.
Installed record	The LICCC record has been downloaded, staged to the SE, and is now installed on the CPC. A maximum of eight different records can be concurrently installed and active.
Licensed Internal Code (LIC)	LIC is microcode, basic I/O system code, utility programs, device drivers, diagnostics, and any other code delivered with an IBM machine for the purpose of enabling the machine's specified functions.
LIC Configuration Control (LICCC)	Configuration control by the LIC provides for server upgrade without hardware changes by enabling the activation of additional previously installed capacity
Model-Capacity Identifier (MCI)	MCI shows the current active capacity on the server, including all replacement and billable capacity. For the z10 EC the Model-Capacity Identifier is in the form of 7nn, 6xx, 5xx or 4xx, where xx or nn indicates the number of active CPs. ► nn can range from 00-64. ► xx can range from 01-12. For the z10 BC the Model-Capacity Identifier is in the form of Axx - Zxx where xx indicates the number of active CPs and can range from 01 to 05.
(Model-Permanent-Capacity Identifier MPCI)	MPCI keeps information about capacity setting active before any temporary capacity was activated.
Model-Temporary-Capacity Identifier (MTCI)	MTCI reflects the permanent capacity with billable capacity only, without replacement capacity. If no billable temporary capacity is active, Model-Temporary-Capacity Identifier equals Model-Permanent-Capacity Identifier
Multi-Chip Module (MCM)	MCM is a compact, densely packed piece of technology with processing units (PUs), system controllers (SCs), memory bus adaptors (MBAs) and the clock chips.

Term	Description
On/Off Capacity on Demand (On/Off CoD)	On/Off CoD represents a function that allows a spare capacity in a CPC to be made available to increase the total capacity of a CPC. An On/Off CoD may be used to acquire additional capacity to handle a workload peak.
Permanent capacity	This is the capacity purchased and activated by the customer. It can be less capacity than the total capacity purchased by the customer.
Permanent upgrade	LIC licensed by IBM to enable the activation of applicable computing resources, such as processors or memory, for a specific CIU-eligible machine on a permanent basis.
Purchased capacity	Capacity delivered and owned by the customer. It can be higher than permanent capacity.
Permanent entitlement record (PER) Temporary entitlement record (TER)	PER or TER refers to the internal representation of a temporary or permanent capacity upgrade processed by the CIU facility. An entitlement record contains the encrypted representation of the upgrade configuration with the associated time limit conditions.
Replacement capacity	Replacement capacity is a temporary capacity used for situations when a customer has lost processing capacity in other parts of the enterprise, either in a planned event or during an unexpected disaster. The two replacement offerings available are Capacity for Planned Events and Capacity Backup.
Resource Link	IBM Resource Link is a technical support Web site included in the comprehensive set of tools and resources available from the IBM Systems technical support site: http://www.ibm.com/servers/resourcelink/
Single chip module (SCM)	In contrast to MCM, single chip contains only one chip with several cores. It is used in z10 BC.
Secondary approval	Customers may optionally define that each Capacity on Demand order be controlled by a second approver. When a secondary approval is required, the request is sent for approval or cancellation to the Resource Link secondary user ID.
Staged record	The point when a record representing a capacity upgrade, either temporary or permanent, has been retrieved and loaded on the Support Element (SE) disk.
Subcapacity	For the z10 EC, CP features (CP4, CP5 and CP6) provide reduced capacity relative to the full capacity CP feature (CP7). For the z10 BC, CP features (CPA - CPY) provide reduced capacity relative to the full capacity CP feature (CPZ).
Temporary capacity	This represents an optional capacity that is added to the current server capacity for a limited amount of time. It can be capacity owned or not owned by the customer.
Unassigned capacity	This is CP or IFL capacity that has been purchased but is not currently active.
Uncharacterized capacity	PUs present but not purchased. Same as dormant capacity.
Vital product data (VPD)	VPD is information that uniquely defines system, hardware, software, and microcode elements of a processing system.

2

Understanding the Capacity on Demand environment

The Capacity on Demand offerings are important features of the System z10 servers. They are developed in response to customer requirements for more flexibility, granularity, and better business control over the infrastructure, both operationally and financially.

Up to eight upgrade configurations can be concurrently installed and activated fully, partially, and in any sequence. These upgrades can be controlled manually through interfaces on the HMC, or programmatically through an application programming interface (API). This chapter discusses the environments and tools used to order, download, install and activate temporary or permanent upgrades.

This chapter discusses the following topics:

2.1 Capacity on Demand upgrades

The System z10 servers offer flexible permanent and temporary upgrade possibilities. With proper planning, all upgrades can be performed concurrently. This means that the server hardware does not have to be stopped and restarted to take advantage of the added capacity, and no power-on reset (POR), initial machine load (IML), or logical partition deactivation and reactivation is necessary.

If the operating systems used on the z10 support (and are configured to include) additional resources without requiring an IPL, all upgrades can be nondisruptive, that is, there is no need for the operating systems to be stopped and restarted to take advantage of the additional resources.

The CoD offerings can be characterized as follows:

- They can be temporary or permanent upgrades.
- A temporary capacity upgrade is removed when no longer needed, while a permanent upgrade remains installed on the server.
- A temporary upgrade can be replacement or billable capacity.
- Replacement capacity is always prepaid.
- Billable capacity can be prepaid or postpaid.

Replacement capacity is used for situations in which a customer loses processing capacity on a server, either because of a planned event or because of a disaster. Replacement capacity is designed to assure access to the same capacity on another server while the original server is unavailable. Replacement capacity is prepaid and usually does not involve additional hardware and software charges from IBM. The two replacement offerings available are Capacity Backup (CBU) and Capacity for Planned Events (CPE).

Billable capacity is designed to help a customer handle workload peaks, either planned or unplanned. Billable capacity can be prepaid or postpaid and the customer is charged based on the amount of capacity and the period that the billable capacity was active. Usage of billable capacity can generate additional software charges. The one billable offering available is On/Off Capacity on Demand.

Permanent upgrades can concurrently add processors (CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs) and memory, or change the CP capacity level, up to the limits of the installed capacity of the server.

Temporary upgrades can concurrently add processors (CPs, IFLs, ICFs, zAAPs, zIIPs, and SAPs) or change the CP capacity level. The set of capacity settings that can be reached by using a specific temporary upgrade configuration record is called the authorization space. Temporary memory upgrades are not available.

Note: All CoD permanent and temporary upgrades are limited by the physical resources currently installed in the server.

Each ordered and retrieved configuration is represented by an LICCC record. A record contains all necessary information about the upgrade. LICCC records are used every time a Capacity on Demand change to the configuration is performed, for temporary and permanent upgrades. The offerings are all built from a common LICCC record structure.

Each record contains the information necessary to control the type of resource that can be accessed and to what extent. Each type of offering uses its own interpretation of record content, but each record can contain three types of information regardless of whether they are used by that offering.

LICCC record

An LICCC record contains:

- Resources

  Resources define the type and amount of resources that can be activated by a record. The limits are set during ordering and can be further limited by the LIC.

- Time elements

  Time elements contain information such as test duration or real activation, each measured in number of days.

- Tokens.

  Tokens represent the number of tests or the number of real activations. Capacity tokens might be used to control capacity, which the customer already paid for (prepaid) or to limit the customer's liability for post-pay capacity.

Each Capacity on Demand offering requires the associated enablement feature codes ordered and installed on the server, and the associated contracts and administrative procedures.

Contracts

The structure of the contracts consists of multiple parts, depending on the individual offering. The contracts are built from a master agreement, the IBM Customer Master Agreement. The administrative system drives the necessary contract addendum. The addendum contains the server and serial number, the pricing, special terms and conditions associated with the offering, and the type and amount of resources provided. The structure of the contracts is:

- Base Capacity on Demand Agreement
  - Common Terms and Conditions
- Depending on the offering type to be used:
  - Replacement Capacity offering Attachment
  - On-Line CoD Buying Attachment
- Depending on the offering to be used:
  - CBU Offering Attachment
  - Capacity for Planned Events Offering Attachment
  - On/Off Capacity on Demand Offering Attachment

After the necessary contracts have been approved by both parties, the procedure installs or updates the customer's machine profiles on Resource Link. See Figure 2-22 on page 48 for an example of a machine profile.

For more information, see *IBM System z10 Capacity on Demand User's Guide,* SC28-6871. Certain conditions can vary by geography. Consult your local IBM sales representative for conditions specific to your country.

2.1.1 On-line Permanent Upgrade

This section describes the concepts of the On-line Permanent Upgrade offering, its ordering process, and activation.

Concepts

On-line Permanent Upgrades are ordered through the Customer Initiated Upgrade (CIU) facility on Resource Link. Permanent upgrades can be applied even if temporary records are already active on the server. Note the following points about the On-line Permanent Upgrade:

- It is enabled through the On-Line CoD buying Attachment. No additional contracts are required.
- It requires the Permanent Upgrade Enablement feature (FC 9898) to be installed.
- It allows ordering permanent capacity upgrades for processors and memory. All processor types are eligible for ordering.
- It cannot exceed the physical resources currently installed on the server.
- If one is being processed, another On-line Permanent Upgrade cannot be created for the same server until the VPD is received and that shows that the first upgrade was successfully installed.
- If an On-line Permanent Upgrade is activated when a temporary offering is already active:
  - The activation of the permanent upgrade will convert processors from billable (On/Off CoD) capacity if the upgrade is for the same processor type. If a sufficient number of processors are not in the activated temporary offering, additional resources will be permanently activated from unused processors.
  - The activation of the permanent upgrade never converts processors from active replacement (CBU or CPE) capacity, it always converts unused processors to permanent.
  - If the upgrade cannot be completely satisfied, no subset of the upgrade is activated.

Ordering

On-line Permanent Upgrade can be ordered through Resource Link only. See 2.3.2, “Resource Link” on page 48 for more details.

Activation

After the Resource Link order has been processed by IBM, the customer can retrieve the offering record and install it on the machine for activation at the customer’s discretion.

2.1.2 Capacity Backup

The following sections describe the concepts of the CBU offering, the ordering process, and activation and deactivation of temporary upgrades.

Concepts

CBU is a prepaid, replacement offering that can be used for disaster situations where a customer has lost processing capacity on a System z server. Use the CBU offering to temporarily add sufficient capacity to another server in order to continue executing an important workload even after an unplanned loss of computer capacity. CBU cannot be used for peak workload management.

CBU requires the CBU Enablement feature (FC 9910) to be installed on the server.

A CBU order specifies processor features for each type of processor. For specialty engines, the number of features decides how many processors of that type can be temporarily added. For CPs, the number of features specifies how many CPs are added or how many permanent CPs increase their capacity level or a combination of both. Details are in “Authorization space for CBU” on page 22. The number and type of temporary PUs that are allowed to be activated are stored in the CBU record.

CBU can be active for up to 90 days. The CBU record can be activated partially and its activation levels can be changed dynamically later as necessary.

A CBU record provides for five free tests over the contract period. Each test can be used for 10 days. However, more tests in groups of 5 up to a total of 15 can be ordered. The number of tests available for each CBU record is specified in tokens inside the record itself.

More than one CBU record may be installed and active on the same server, allowing various backup scenarios on one System z10. Another possibility is to have one large CBU record installed and then, at the time of activation, decide to use only a subset of the specified resources depending on the disaster situation. In general, IBM software charges are not affected by CBU or the test CBU activation.

Ordering and limits for CBU

Capacity Backup can be ordered through eConfig or Resource Link. For details, see:

- 2.3.1, “eConfig” on page 46
- 2.3.2, “Resource Link” on page 48
- 3.1.2, “Order, retrieve, and install a record” on page 83

The ordering limits are:

- The total number of zAAPs (temporary and permanent) cannot be higher than the total number of CPs (temporary and permanent).
- The total number of zIIPs (temporary and permanent) cannot be higher than the total number of CPs (temporary and permanent).
- Total number of CP CBU features equals the number of added CPs plus the number of permanent CPs that change their capacity level.

Replenishment

CBU allows the following replenishments:

- The replenishment can include additional CBU features for CPs or for specialty engines.
- The number of tests can be increased by multiples of 5 up to 15 tests per CBU offering.
- A real activation can be replenished only after it was used. There can be no more than one real activation per CBU record.
- The CBU expiration date cannot be replenished.

Record content

Each CBU record consists of resources, time elements, and tokens, as follows:

- Resources

  This record consists of definitions describing the maximum number of resources that can be added as part of CBU. The record contains the number of purchased features for each type of processor. All possible target configurations are determined at the activation time, and they depend on the current capacity setting of the server and available PUs.

- Time elements
  - Test duration time is set to 10 days.
  - Real activation time is set to 90 days.
  - The expiration date of the complete record can be set at 1 - 5 years.
  - A replenishment of the record cannot change the expiration date, which means that after the expiration date, a new CBU record must be ordered.
- Tokens
  - Controls the number of allowed tests. The default is five, but the number can be set to 5, 10, or 15. Test activations are possible only if a real activation still remains.
  - Controls the amount of real activation, which is set to one. Real activation is possible only while no test is active.

Each record is also subject to LIC rules and ordering limits.

Licensed Internal Code rules for CBU

The Licensed Internal Code (LIC) ensures that activation and deactivation result in a valid configuration for the model.

The LIC rules are:

- Increasing the number of processors is allowed.
- Removal of processors is not allowed.
- Increasing the CP capacity level is allowed.

Authorization space for CBU

CBU resources are purchased as CBU features for CPs or specialty engines, as follows:

- The number of purchased CBU features for specialty engines (IFLs, ICFs, zAAPs, zIIPs, and SAPs) represents a maximum number of temporary specialty engines that can be turned on.
- The number of purchased CP CBU features represents the number of CPs that can be added, the number of permanent CPs that change capacity level, or combination of both.

In other words, the number of CP CBU features required to get from the current capacity setting to a target capacity setting is calculated as the number of added CPs plus the number of permanent CPs that change capacity level. During the activation, with a given current configuration and number of CP CBU features in a purchased CBU record, all possible target configurations are shown and any of these can be activated. This can be changed dynamically at any time.

All capacity settings that can be reached by using a specific record are collectively called the *authorization space* for this record.

Figure 2-1 on page 23 shows all possible capacity settings for a CBU upgrade on a z10 EC server with a permanent capacity setting 401 and two CP CBU features.

The two CP CBU features can be used to:

- Add processors of the same capacity level 4xx, which allows reaching capacity settings 402 and 403.
- Convert capacity level of the permanent processor, which allows reaching capacity settings 501, 601, and 701

- Convert the capacity level of the permanent processor and add one processor, which allows reaching capacity settings 502, 602, and 702.

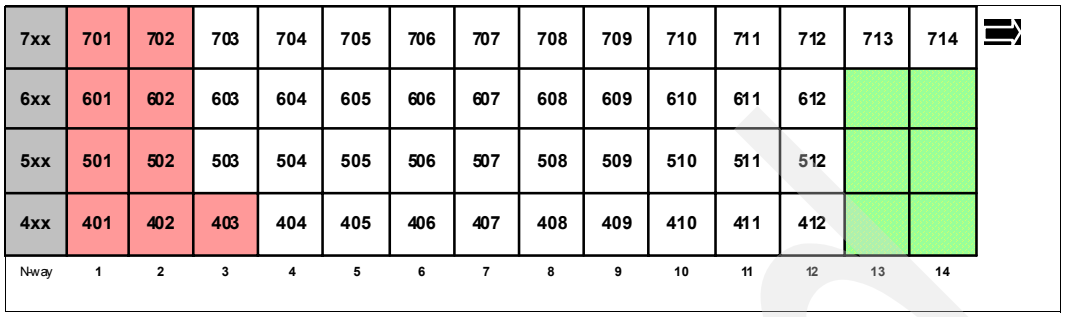

Figure 2-1 Authorization space for 401 and two CP CBU features

An example based on a permanent capacity setting of 504 and five CP CBU features is similar, as shown in Figure 2-2. The five CP CBU features can be used to:

- Add processors of the same capacity level 5xx, which allows capacity settings 505 - 509
- Convert the capacity level of the permanent processors, which allows capacity settings 604 and 704. Each permanent CP increasing its capacity level requires a CP CBU feature. In this case four CP CBU features are required.
- Convert the capacity level of the permanent processors and add one processor, which allows capacity settings 605 and 705
- Because decreasing the capacity level or the number of CPs is not allowed, operating on capacity level 4xx or with fewer than four CPs is not possible.

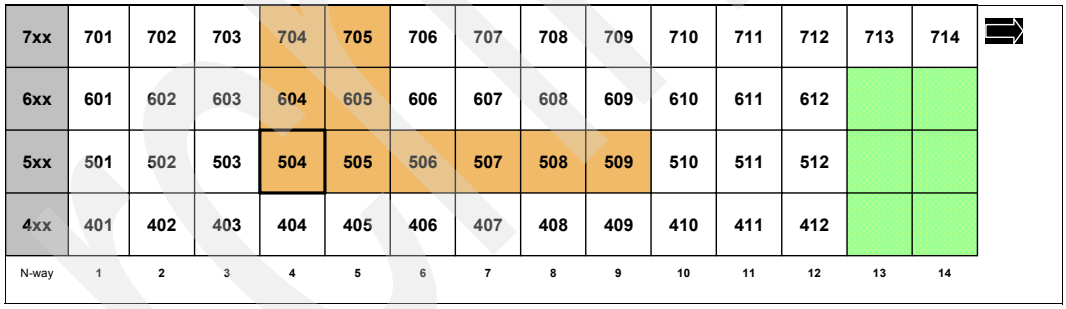

Figure 2-2 Authorization space for 504 and five CP CBU features

Figure 2-3 on page 24 shows an example with a permanent capacity setting of 612 and 14 CP CBU features. Adding more processors to a 6xx capacity level is not possible because only 12 CPs can be on subcapacity setting. Using 12 CP CBU features enables the conversion of all 12 CPs to a higher capacity level and upgrade the server to 712. The CP CBU features 13 and 14 can upgrade the server to 713 or 714, respectively.

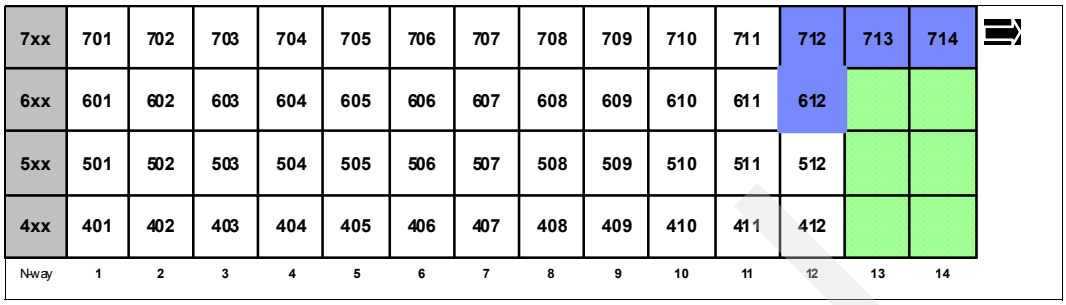

Figure 2-3 Authorization space for 612 and 14 CP CBU features

If fewer CP CBU features are in a CBU record than permanent CPs are in the current capacity setting, a CBU activation can only add CPs of the same capacity level. The reason is that not enough CP CBU features are available to convert CPs to a higher capacity level. For example, if permanent capacity is 504 and there are three CP CBU features, only 505 - 507 capacity settings can be activated. No capacity setting from 6xx or 7xx capacity level can be activated, because there are not enough CP CBU features to change the capacity level of four permanent CPs.

CBU Force activation

The Change Activation Level panel for CBU contains the option Force activation. If the Force activation option is selected, the server tries to activate the entire CBU record. If it could exceed the physical available resources, only a subset is activated. In this case, the remaining CBU processor resources that cannot be activated are put into a pending status and a message, as shown in Figure 2-4, is issued.

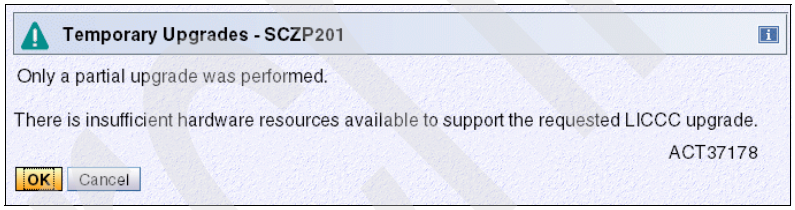

Figure 2-4 Temporary Upgrades - message ACT37178

When a PU becomes available (that is, another temporary CoD record is deactivated or physical resources are added), it is automatically used to activate the pending CBU resources.

If the server has subcapacity CPs and CBU can either add CPs or increase their capacity level or a combination of both, LICCC will choose the option that adds the most MSUs to the current configuration.

The Force activation option is available only for CBU.The CBU Record Details panel shows the amount of pending processors; see Figure 2-5 on page 25.

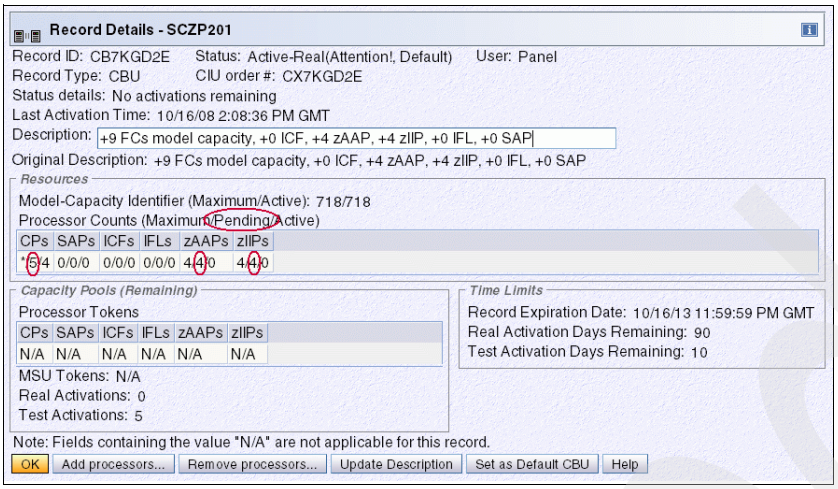

Figure 2-5 Record Details - showing the pending PUs

Force CBU is also used by APIs prior to System z10. Those APIs always activate the entire CBU record. Force CBU preserves the backward compatibility with those APIs, which are not aware of System z10 CoD flexibility and always expect to activate the full CBU record.

Force activation is only allowed for the default CBU record. Customer has to set CBU record as default before using the Force activation option manually. If no CBU record is set as the default and an API from before System z10 is used, the oldest CBU record is used.

APIs and the Set as default option

An API from before System z10 can be used to activate or deactivate a CBU record. The API expects that only one CBU record is available, but System z10 allows more than one CBU record installed. As the design of APIs stays unchanged due to compatibility reasons, the *Set as Default CBU* option is available on System z10 servers. It flags the selected CBU record as default and all the API calls (before System z10) will be able to see and use only this particular CBU record. Setting another CBU record as the default can release the flag on the former one. There is no reset option.

Code that uses an API from before System z10 does not have to change. It can work with only one CBU record, which is set as the default. Because the code is not aware of Capacity on Demand flexibility, which allows activating a record partially, the code will use, by default, the Force activation task to activate the whole CBU record, even if there are not enough resources.

To allow changes of capacity through API, the SE needs *Allow capacity change API requests* to be allowed, as shown in Figure 2-6 on page 26. It can be reached by selecting:

Console Actions →Support Element Settings →Customize API Settings

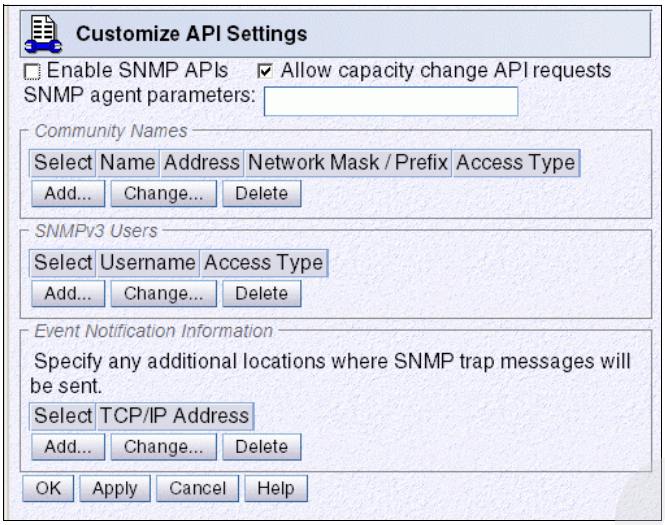

Figure 2-6 SE API settings

Chapter 3, “Capacity Backup” on page 81 provides a CBU example from ordering through activation, and deactivation to replenishment.

2.1.3 Capacity for Planned Events

This section describes the concepts of the CPE offering, the ordering process, and activation and deactivation of temporary upgrades.

Concepts

The CPE offering is a prepaid, replacement capacity offering. It is designed to help customers maintain their processing capacity during a planned event in the computing environment, such as relocation of servers, re-cabling, or general work on the physical infrastructure of the data processing environment. CPE cannot be used for peak workload management or for disaster recovery.

Capacity for Planned Events (CPE) requires the CPE Enablement feature (FC 9912) to be installed. The activation period for CPE is limited to three days. An activated CPE record can use all available resources installed on the server. It offers the ability to concurrently and temporarily activate additional CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, to increase the CP capacity level, or a combination of these. With the CPE record, all available PUs and all capacity-level increases are made available. The type and quantity of processors are defined as part of the activation of the CPE record and can be changed dynamically.

The CPE record can be activated only once. It contains an activation token, which will be consumed upon the record activation. CPE can be active for a three-day period.

There is no test available for CPE.

Ordering and limits for CPE

Capacity for Planned Events can be ordered through eConfig or Resource Link. See:

- 2.3.1, “eConfig” on page 46
- 2.3.2, “Resource Link” on page 48
- “Order, retrieve, and install a record” on page 106

There are no ordering limits on resources for a CPE record.

Replenishment

CPE records cannot be replenished.

Record content

Each CPE record consists of resources, time elements, and tokens, as follows:

- Resources

  Within the limits of resources installed and dormant, there is no limitation in the number of CPs, IFLs, ICFs, zIIPs, zAAPs, and SAPs that can be activated.

- Time elements

  Real activation duration is limited to 72 hours.

- Tokens
  - Number of tests is set to zero.
  - Number of real activations is set to one.

Each record is also subject to LIC rules.

Licensed Internal Code rules for CPE

Licensed Internal Code (LIC) ensures that activation and deactivation result in a valid configuration for the related model.

The LIC rules are:

- Increasing the number of processors is allowed
- Removal of permanent processors is not allowed.
- No restrictions exist regarding the number of IFLs, ICFs, zIIPs, or zAAPs.
- Incrementing the CP capacity level is possible.

Authorization space for CPE

CPE allows temporary activation of any available PU and to increase the capacity level of CPs. Any capacity setting that does not decrease the number of processors nor decrease the CP capacity level is allowed.

Chapter 4, “Capacity for Planned Events” on page 105 provides a CPE example, which describes ordering, activation, and deactivation of the CPE record.

2.1.4 On/Off Capacity on Demand

This section describes the concepts of the On/Off CoD offering, the ordering process, and activation and deactivation of temporary upgrades.

Concepts

On/Off Capacity on Demand can help customers handle workload peaks, either planned or unplanned, while keeping the related costs under control. It is a billable capacity offering that can have an associated hardware charge based on a 24-hour billing window, if the customer activates non-purchased resources. However, software charges can vary according to the license agreement for the individual software product.

On/Off CoD offers the ability to concurrently activate additional temporary CPs, ICFs, IFLs, zAAPs, zIIPs, and SAPs, to increase CP capacity level, or a combination of both. It allows increasing the capacity of any engine type up to twice the currently purchased capacity, which is the high watermark (HWM).

On/Off CoD can be used for customer peak workload requirements, for any length of time.

The On/Off Capacity on Demand contract allows one full function test. The test is delivered through a unique On/Off CoD test record, that can be used for 24 hours to validate the process and procedures.

On/Off CoD requires the On/Off CoD Enablement feature (FC 9896) to be installed.

Available PUs that are currently unassigned and non-purchased can be temporarily and concurrently activated as any processor type using an LICCC upgrade. On/Off CoD cannot change the server model, because additional book installation is not supported. However, On/Off CoD can change the capacity setting temporarily if the CP configuration is changed.

For specialty engines, the On/Off CoD limit is the number of purchased engines for each type. An On/Off CoD order can specify up to 100% of the MSU value of the purchased CP capacity. The Resource Link profile for each server enforces the rule. Capacity is computed based on processing capacity gained by adding the engines. It is based on the published Large Systems Performance Reference (LSPR) values for the configuration. For details, see “On/Off CoD authorization space” on page 33.

To achieve the highest flexibility, we recommend allowing maximum resources for the record as offered by default. This way, deciding which resources to activate can be decided at activation time.

More than one On/Off CoD record can be installed but only one can be active at a time.

Resources in the record define the upper limit of what resources can be activated. They can be active for unlimited time. To limit the time, capacity tokens can be used.

Capacity tokens

The capacity token is a representation of resources available for a given period of time. The measurement units used are *MSU days* for CP capacity and *specialty engine days* for specialty engines. More precisely, six types of capacity tokens are defined and managed for each On/Off CoD record:

- MSU tokens: One MSU token is worth one MSU day, a capacity of one MSU for 24 hours.
- ICF tokens: One ICF token is worth one ICF day, a capacity of one ICF for 24 hours.
- IFL tokens: One IFL token is worth one IFL day, a capacity of one IFL for 24 hours.
- SAP tokens: One SAP token is worth one SAP day, a capacity of one SAP for 24 hours.
- zAAP tokens: One zAAP token is worth one zAAP day, a capacity of one zAAP for 24 hours.
- zIIP tokens: One zIIP token is worth one zIIP day, a capacity of one zIIP for 24 hours.

The customer decides how and when the capacity tokens are consumed. For example, 100 MSU tokens can be spent in two days by using 50 MSUs each day, or in five days by using 20 MSUs each day, or any other combination. The resource limits set up in the record are based on the rules described in “Ordering and limits for On/Of CoD” on page 31.

Note: The actual balance of capacity tokens can be checked anytime on HMC.

Temporary capacity delivered by On/Off CoD can be either prepaid or postpaid, decided by the customer.

Prepaid capacity

The customer can order, pay for, and keep the prepaid processing capacity in reserve. The customer decides when and how the capacity will be consumed. Prepaid On/Off CoD capacity is always in the form of capacity tokens.

Prepaid capacity helps to control the costs because the budget related to the On/Off CoD is allocated and spent in advance with no additional hardware costs at the time of activation. Prepaid On/Off CoD allows you to limit the total amount of temporary capacity ordered.

Note: The On/Off CoD records that contain prepaid capacity tokens never expire.

Postpaid capacity

The On/Off CoD postpaid offering has no processing capacity prepaid in reserve. Resources can be activated by the customer anytime and remain active until the customer deactivates them or the On/Off CoD record expires.

Usage of capacity tokens is not mandatory, but possible for postpaid On/Off CoD records. The customer can use capacity tokens and limit the amount of MSU days or specialty engine days included in the record and thus limit the cost related to the activation.

Capacity consumption and billing

The hardware charge for temporary On/Off CoD resources is based on the consumed processing capacity during the 24-hour billing window.

When the On/Off CoD record is activated, a new 24-hour billing window starts. The hardware charge is based on the peak usage of the temporary activated resources during the billing window. The peak usage simply represents the highest capacity level activated during a particular billing window, it does not matter how long this capacity level was activated. Figure 2-7 on page 30 demonstrates how the peak usage for the 24-hour billing window is calculated.

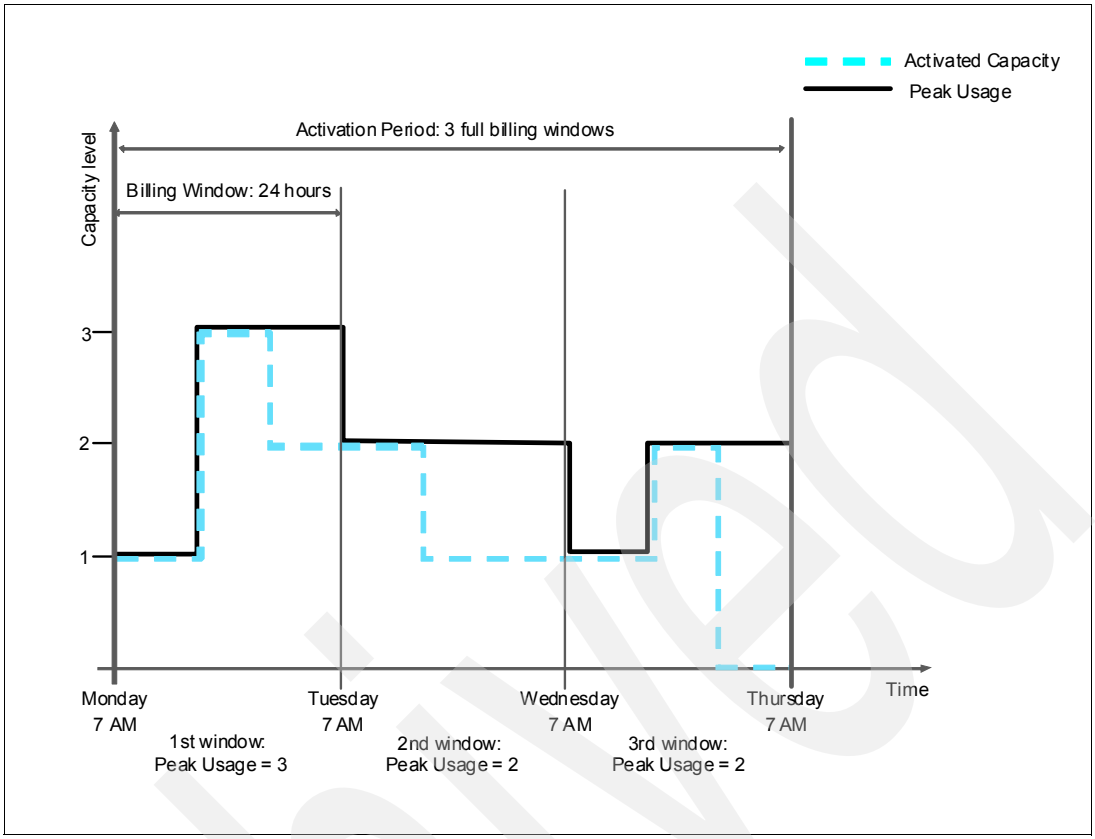

Figure 2-7 Peak usage calculation

Note: A one-hour grace period is available at the end of each billing window to protect customers from being charged for a new billing window if they deactivate the resources a little too late.

If the same On/Off CoD record is deactivated and later reactivated within the same billing window, the peak usage calculation happens as though the record was activated continuously and new billing window is not started. See Figure 2-8 on page 31.

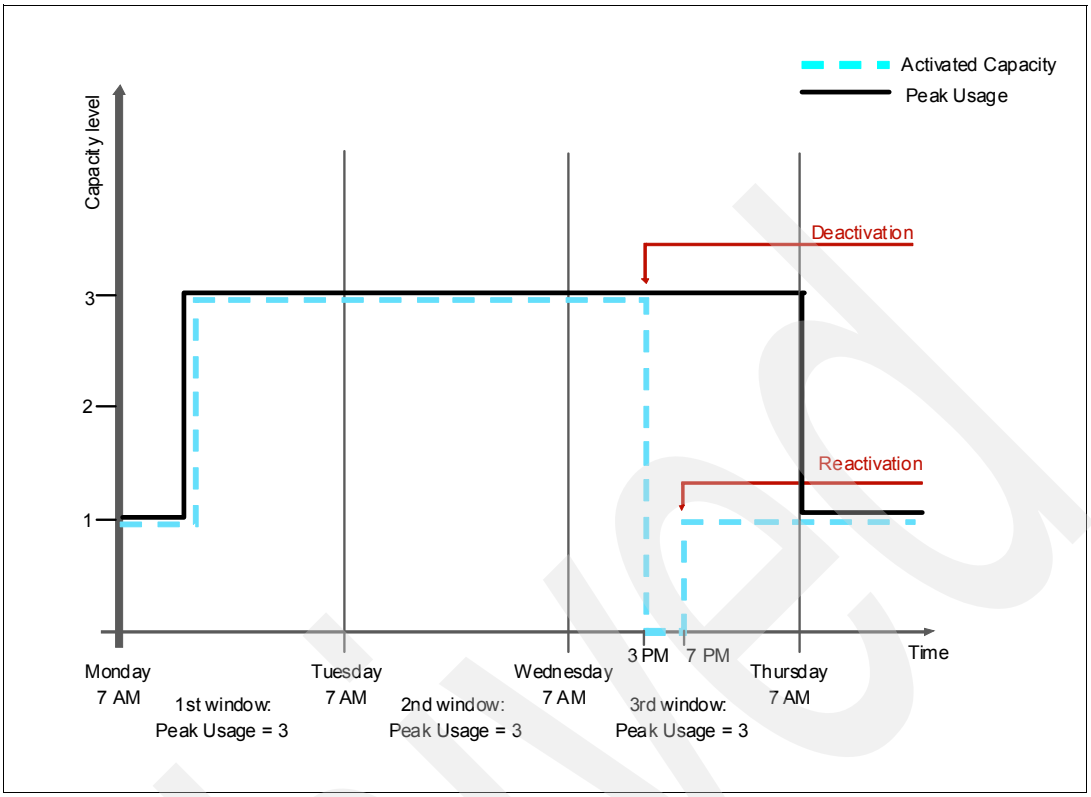

Figure 2-8 Deactivation and reactivation within the same billing window

Note: The record deactivation during the one-hour grace period ends the billing window definitely. Any subsequent record activation starts the new billing window.

The capacity tokens, prepaid or not, are decremented every 24 hours after the activation, based on the peak usage for that billing window.

For example, the record contains 120 MSU tokens and 10 IFL tokens at the beginning of the billing window when temporary capacity of 100 MSUs and two temporary IFLs are activated. At the end of the billing window 100 MSU tokens and two IFL tokens are consumed, thus 20 MSU tokens and eight IFL tokens remain in the record. The record now is deactivated by LIC, because although enough IFL tokens remain, not enough MSU tokens are available to support the current activated capacity setting for the next billing window.

Ordering and limits for On/Of CoD

On/Off Capacity on Demand record can be ordered through Resource Link only. See:

- 2.3.2, “Resource Link” on page 48
- “Order, retrieve, and install a record” on page 124

The ordering limits are:

- The order process limits resources for each record based on permanent configuration at the time of the order.
- Temporary CP capacity is limited to twice the MSU value of the purchased capacity. It is specified as a percentage of the purchased capacity.

- Temporary specialty engines are limited by an absolute number of purchased engines of the same type.
- The total number of zAAPs (temporary and permanent) cannot be higher than the total number of CPs (temporary and permanent). This is a record-based calculation, where the amount of CPs purchased is defined by the HWM.
- The total number of zIIPs (temporary and permanent) cannot be higher than the total number of CPs (temporary and permanent). This is a record-based calculation, where the amount of CPs purchased is defined by the HWM.
- The number of temporary IFLs cannot be higher than the number of purchased IFLs.
- The number of temporary ICFs cannot be higher than the number of purchased ICFs. The total number of ICFs cannot exceed 16.
- There is no limit on capacity tokens.

Replenishment

The customer has the opportunity to replenish the On/Off CoD record at any time, whether the record is active or not. Replenishments are allowed for an On/Off CoD record, as follows:

- The amount of On/Off CoD resources can be increased.
- The expiration date can be adjusted by not more than 180 days to the future.
- For records containing capacity tokens, the amount of capacity tokens can be increased.
- There are no replenishments available for a test On/Off CoD record.

Note: Replenish On/Off CoD records on a regular basis, at least every 180 days.

Record content

Each On/Off CoD record consists of resources, time elements, and tokens, as follows.

- Resources
  - On/Off CoD record activation can increase the MSU value up to 100% of the purchased capacity (HWM).
  - The number of specialty engines (IFLs, ICFs, zIIPs, zAAPs, or SAPs) can be doubled by Capacity on Demand record activation. The number of ICFs cannot exceed 16.
- Time elements
  - Automatic deactivation of activated On/Off CoD resources when the time limit is reached are as follows:
    - All resources stay active, if any resource cannot be removed.
    - The ability to remove resources is checked every 24 hours.
  - Record expiration defaults to 180 days, with the exception of prepaid temporary capacity (using prepaid MSU tokens or specialty engine tokens). No expiration is defined for prepaid temporary capacity.
  - A one-hour billing grace period allows an easy transition from upgraded capacity back to a permanent capacity without an extra charge for the next day.
- Tokens
  - MSU tokens represent the CP capacity available for a given period of time in terms of MSU days. One MSU token is worth one MSU day (capacity of one MSU for 24 hours).

– Specialty engine tokens represent the capacity of the specialty engine available for a given period of time in terms of engine days. One specialty engine token is worth one engine day (capacity of one specialty engine for 24 hours).
– Six independent specialty engine token pools are managed in the record:
  - MSU token pool
  - ICF token pool
  - IFL token pool
  - SAP token pool
  - zAAP token pool
  - zIIP token pool

Each record is also subject to LIC rules.

Licensed Internal Code rules for On/Off CoD

Licensed Internal Code (LIC) ensures that activation and deactivation result in a valid configuration for the related model.

The LIC rules are:

- Record activation must result in positive MSUs relative to the current capacity setting.
- Increasing the number of processors is allowed.
- Removal of permanent processors is not allowed.
- Incrementing the capacity level is possible without capacity level degradation.

On/Off CoD authorization space

Capacity on Demand offers the ability to activate a subset of the resources defined in a temporary record. For specialty engines a subset represents a part of the total number of specialty engines that can be activated. For CPs, the subset represents a combination of the number of CPs and the CP capacity settings. All the capacity settings which can be reached using a specific record are called the authorization space for this record. Each capacity setting corresponds to an MSU value.

Note: The MSU values of IBM servers are published in the *IBM System z10, z9®, eServer™ zSeries and S/390 Software Pricing Information* document available at:

`http://www.ibm.com/servers/resources/servers_eserver_zseries_library_swpriceinfo_pdf_swp_ibm.pdf`

Only capacity settings which do not exceed twice the MSU value of the purchased capacity (HWM) can be reached by activating On/Off Capacity on Demand records. Figure 2-9 on page 34 shows an example of how the authorization space is calculated. Note that the table shows an increase in the number of engines from left to right and an increase in capacity level from bottom to top.

Note: high watermark influences only On/Off CoD authorization area, it has no influence on CBU or CPE.

The authorization space goes right, and in subcapacity cases up, from the starting point. On/Off CoD cannot decrease the number of CPs, or decrease the capacity level. For these reasons the authorization space does not move left (decreasing the number of engines) or down (decreasing the capacity level). The authorization space can also be limited by physical resources of the system and configuration-related conditions.

Figure 2-9 shows an authorization space for capacity setting 504. The MSU value for this capacity setting is 207. All capacity settings in its authorization space have an MSU value lower than 414 MSU (2 x 207).

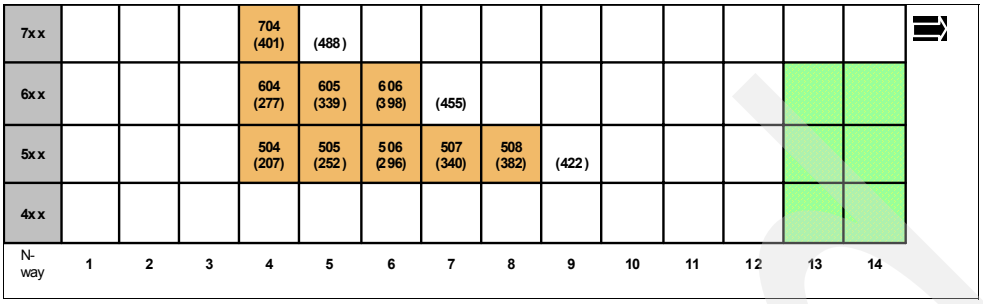

Figure 2-9 Examples of authorization space for On/Off CoD

The previous example showed a case where HWM was the same as the permanent capacity setting. Figure 2-10 shows an example where purchased capacity (owned by customer) is 504 (207 MSUs) and the permanent (activated) capacity setting is 402. The On/Off CoD record can activate any capacity setting below or equal to 414 MSUs (2 x 207), because the MSU value for capacity setting 504 is 207 MSUs. The authorization space also includes all the capacity settings highlighted in pink. This means that the authorization space includes all capacity settings from 402 up to twice the MSU value of the HWM setting.

When a temporary capacity activated by an On/Off CoD record is not higher than MSU value of HWM setting, there are no hardware charges applied. See the area marked by thick solid line on Figure 2-10.

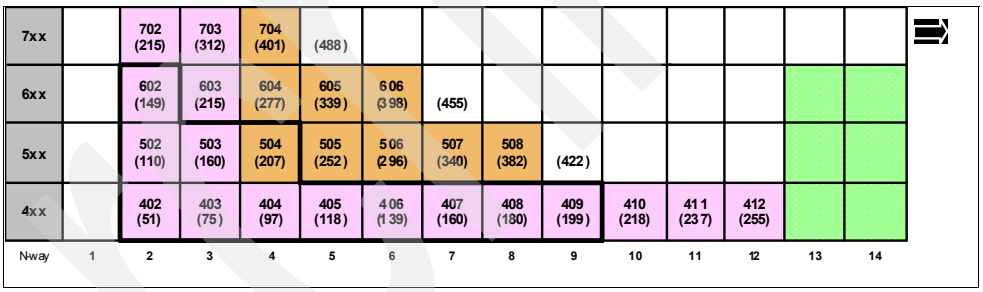

Figure 2-10 Authorization space - HWM higher than the permanent capacity

- If a server is permanently upgraded to a capacity setting with a higher MSU value than the HWM, the On/Off CoD authorization space will be adjusted automatically. After the upgrade, the licensed internal code calculates the new authorization space based on the new HWM. There is no need to order a new On/Off Capacity on Demand record. However, specialty engines limits remain the same, because the record is not updated when a permanent upgrade is performed.

Chapter 5, “On/Off Capacity on Demand” on page 123 provides an example of On/Off CoD record from ordering, through activation to deactivation.

Chapter 6, “On/Off Capacity on Demand with prepaid upgrades” on page 147 provides an example of On/Off CoD record with capacity tokens usage.

2.2 Multiple activations

Capacity on Demand design allows more than one offering to be active. It also allows you to perform a permanent upgrade while temporary offerings are active. This section provides information and rules that apply during multiple activations.

Capacity on Demand allows the following multiple activations:

- Multiple temporary upgrades can be active at the same time
- Permanent upgrades are allowed while temporary upgrades are active

The Licensed Internal Code (LIC) defines and controls the activation level change and the combination of records that can be activated at any given time. Each record has its own set of allowed resources, a specified duration for which resources can be used, an indication of whether tests are allowed, and a time frame for which the record is valid. A record can be activated by using only part of the defined resources and it can have defined resources added or removed dynamically while the record is active.

A Capacity on Demand offering can be characterized in two ways (see Table 2-1):

- Permanent or temporary capacity

  Permanent capacity stays with the server. The one permanent capacity offering is the On-line Permanent Upgrade.

  Temporary capacity is removed when it is no longer needed. The three temporary capacity offerings are Capacity Backup (CBU), Capacity for Planned Events (CPE), and On/Off Capacity on Demand (On/Off CoD).

- Billable or replacement temporary capacity
  - Billable capacity helps a customer handle workload peaks, either expected or unexpected. Billable capacity can be prepaid or postpaid; the customer is charged based on the amount of capacity and the length of time the capacity was used. The one billable capacity offering is On/Off Capacity on Demand (On/Off CoD).
  - Replacement capacity is when a customer has lost processing capacity in another part of the enterprise, either in a planned event or unexpected disaster. Replacement capacity helps a customer achieve the same level of capacity by using a different server while the replaced server is unavailable. Typically, replacement capacity is prepaid and does not involve additional charges from IBM. The two replacement capacity offerings are Capacity Backup (CBU) and Capacity for Planned Events (CPE).

Table 2-1 Capacity on Demand upgrades summary

Name	Permanent upgrade	Temporary upgrade	
		Billable	Replacement
On-line Permanent Upgrade	Yes	-	-
On/Off CoD	No	Yes	No
CBU	No	No	Yes
CPE	No	No	Yes

2.2.1 Multiple active temporary offerings

Approximately two hundred temporary CoD records can be staged on the Support Element. A maximum of eight records can be installed. These records are managed independently from

each other and independently from permanent system changes. Although any combination of installed records can be active at the same time, only one On/Off CoD record can be active at any given time. The order of activation determines the allocation of resources.

When a temporary record is activated, several rules apply:

- The record can use only those resources that are available.
- If a record allows activation of more resources than are needed, the resources can be activated partially and changed dynamically in the future as needed.
- If a record allows usage of more resources than are available, the resources can be activated only partially. Resources can be added dynamically later if they become available. The two ways that resources can become available are when they are deactivated from another temporary record or when a physical upgrade is performed, which adds new resources. Activation of additional resources is not automatic, except for the CBU record with the `Force` option enabled.
- Temporary records can be activated in any order and their activation level can be changed later as needed. Also, records can be deactivated in any order, independently of the order in which they were activated.

At any given time, the following three identifiers reflect the capacity of the server:

- Model-Capacity Identifier (MCI) shows the current active capacity, including all replacement and billable capacity.
- Model-Temporary-Capacity Identifier (MTCI) reflects the permanent capacity with billable capacity only, without replacement capacity. If no billable temporary capacity is active, Model-Temporary-Capacity Identifier equals Model-Permanent-Capacity Identifier.
- Model-Permanent-Capacity Identifier (MPCI) keeps information about capacity setting active before any temporary capacity was activated.

Capacity identifiers are displayed in several places, such as in the Temporary Upgrades panel on the SE, CPC details on HMC, or in SMF records. Figure 2-11 shows capacity identifiers on the CPC details panel. In this case, the permanent capacity of the server is 709, no billable upgrade is active, and a replacement upgrade is active to capacity 714.

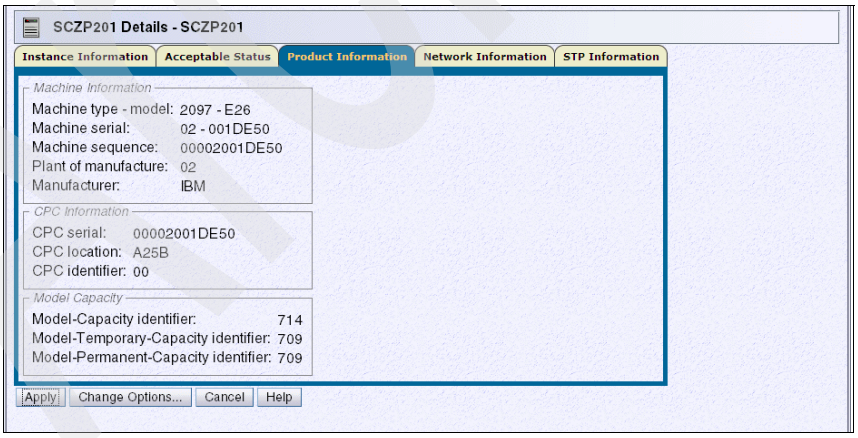

Figure 2-11 CPC details

Authorization space for On/Off CoD

When another temporary record is already active, On/Off Capacity on Demand authorization space is still calculated based on high watermark (HWM) MSU value as described in “On/Off CoD authorization space” on page 33. In other words, the shape of an authorization space is exactly the same as though no other temporary record is active.

Scenario one

Figure 2-12 shows a z10 BC with the permanent capacity setting C02. The solid, light gray box represents an authorization space for an On/Off CoD record that can activate up to 100% of the purchased MSU capacity. For example, suppose a CBU record is activated and temporarily upgrades the server to R03. Then, when an On/Off CoD record is to be activated, the On/Off CoD record authorization space shape stays exactly the same as it was before the CBU activation, as shown in the dark gray box.

C02 is the high watermark (HWM) for this configuration. The hardware cost for activating the On/Off CoD is always calculated from the HWM base. When CBU upgrades the server to an R03 capacity setting and then the On/Off CoD record upgrades R03 to V03, the cost is calculated the same as for a machine being upgraded from a C02 to a G02.

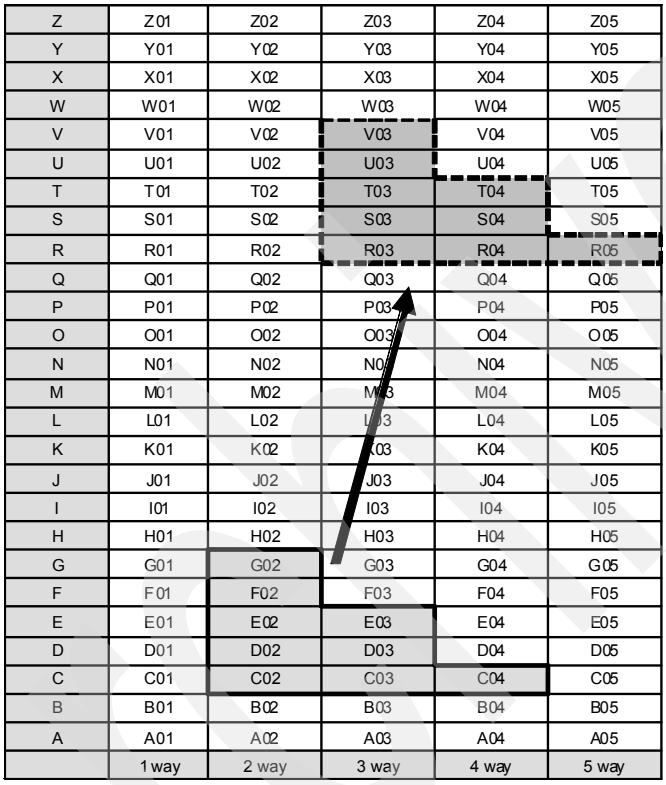

Figure 2-12 Scenario one: Permanent capacity is equal to the high watermark

Scenario two

Another scenario is shown in Figure 2-13 on page 38. Permanent setting is C02, but the high watermark is now F03, which means the customer has purchased an F03, but activated only C02 capacity setting.

In this case, any On/Off CoD from the C02 up to the capacity of the F03 can be activated without any hardware cost being incurred by IBM. In the figure, the light gray area represents capacity settings that can be activated by On/Off CoD in this way.

The authorization space for the On/Off CoD offerings are calculated based on the F03. If the On/Off CoD activates more than the equivalent of the F03, the incurred hardware cost will be based on the capacity difference between the F03 and whichever capacity is activated by the On/Off CoD.

The lower of the two dark gray areas represents the capacity settings that can be activated by On/Off CoD starting from the C02.

The upper dark gray area represents the On/Off CoD authorization space if a CBU activation had taken the C02 to a R03. The On/Off CoD authorization area shape stays the same, it is just cut by the physical limits. In other words, although the dashed area should have the same shape as the gray shape below it, its shape boundaries are *cut out* by the server's physical limitations (there are no rows over row Z, so the shape's height is only 9 rows, not 11).

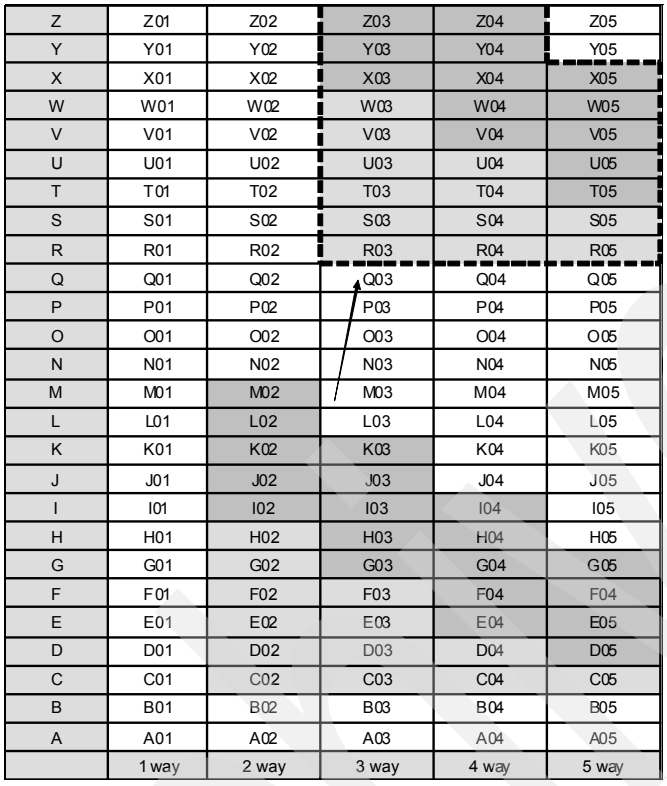

Figure 2-13 Scenario two: Permanent capacity is lower than the high watermark

CBU authorization space

For any CBU upgrade, the number of CP CBU features necessary is calculated as the number of permanent CPs that are changing capacity level, plus the number of CPs added to the current configuration, as described in “Authorization space for CBU” on page 22.

Figure 2-14 on page 39 shows a combined scenario. The z10 BC is configured with a permanent capacity of C02.

On/Off CoD upgrades the server to E04. If CBU is to activate the capacity setting F04, it needs two CP CBU features, because the capacity level of two permanent CPs will be changed and no CP is added to the current configuration. If On/Off CoD is deactivated the current capacity setting will become D02, because CBU only changed capacity level of two CPs by one step.

Similarly, when On/Off CoD upgrades C02 to E04, three CP CBU features are required to activate Z05, because two permanent CPs change capacity level and one CP is added to the current configuration. If On/Off CoD is deactivated, current capacity becomes X03, because CBU increased the capacity level of two permanent CPs by 21 capacity levels and added one CP.

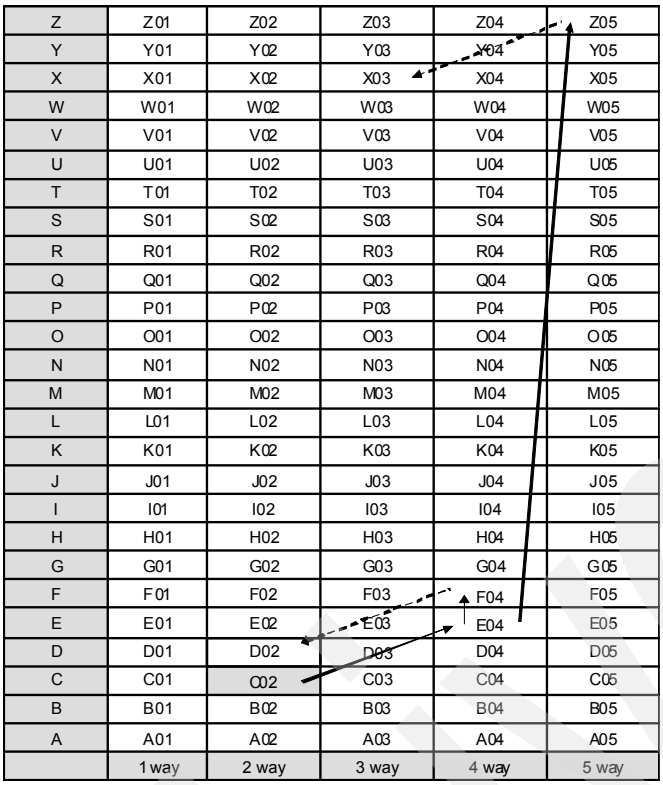

Figure 2-14 Combined scenario

2.2.2 Permanent upgrade with temporary offerings active

When a permanent upgrade is concurrently installed, it might affect active temporary records in different ways, depending on the record type and the physical resources available.

Upgrade rules

Resources must be available on the server to complete an installation of a permanent upgrade, because such an upgrade cannot be partially performed. When a permanent upgrade is performed with temporary upgrade concurrently active, the following rules apply:

- If a billable capacity record is active and if processor types match, temporary resources activated by this billable record are converted to permanent resources. If additional processor capacity is needed to satisfy the permanent upgrade, that capacity will be activated from non-active PUs.

  If a temporary billable record is active and a permanent upgrade is performed, the permanent upgrade converts any active temporary PU of the same type to a permanent PU. For example, a temporary zAAP cannot be converted to a permanent CP, only to a permanent zAAP. PUs of different types are activated from the pool of available PUs. Even if all activated resources from a billable record are converted, the record is not deactivated.

- If a replacement capacity record is active, resources activated by this record are never converted to permanent resources.

  If a temporary replacement record is active and a permanent upgrade is performed, the permanent upgrade converts only PUs that are in a pool of available PUs. Replacement capacity is applied in addition to the new permanent capacity setting.

These rules are described in more detail in the following sections.

Permanent upgrade with a billable capacity record active

When a permanent upgrade is performed and a billable record is active, some or all of the temporary PUs associated with the billable record are converted to permanent. Only PUs of the same type activated by billable temporary upgrade can be converted to permanent.

Figure 2-15 shows an example. A server with permanent capacity setting 708 is upgraded by On/Off CoD, which added two temporary CPs. Then, a permanent upgrade is installed, changing the Model-Permanent-Capacity Identifier from 708 to 709, which means an increase of one permanent CP. When the permanent upgrade is installed, one of the On/Off CoD temporary CPs is converted to a permanent CP and the count of the temporary CPs is reduced by one.

In the case of the permanent upgrade with On/Off CoD active, the only limit of the On-line Permanent Upgrade would be the physical resources available on the server.

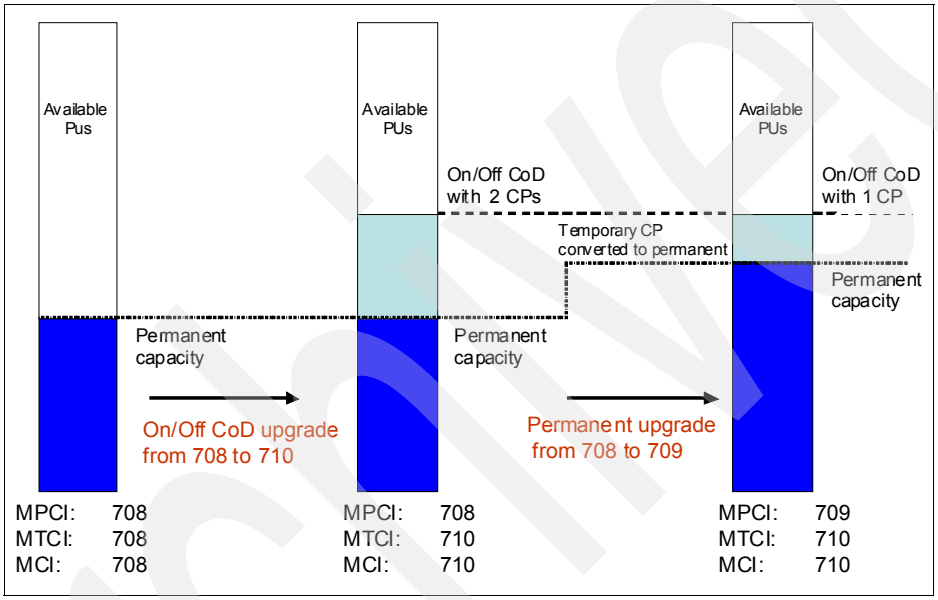

Figure 2-15 Permanent Upgrade with On/Off CoD active

Permanent upgrade with a replacement capacity record active

When a permanent upgrade is performed while temporary replacement capacity is active, the permanent upgrade adds resources to the permanent server capacity and the temporary replacement capacity resources are applied on top of the new permanent capacity.

PUs activated by a replacement capacity record are not converted to permanent. If not enough dormant resources are available to perform the permanent upgrade, resources associated with a temporary replacement capacity record can be manually changed to free resources for the permanent upgrade.

Figure 2-16 shows a server upgraded with CBU first, followed by a permanent upgrade. The permanent capacity of the server is 708. CBU activates another two CPs on top of the permanent configuration, and the current capacity identifier of the server becomes 710. When the On-line Permanent Upgrade adds one more CP to the permanent configuration the permanent capacity changes from 708 to 709. CBU remains active with two additional CPs on top of the permanent capacity, bringing the current capacity to 711.

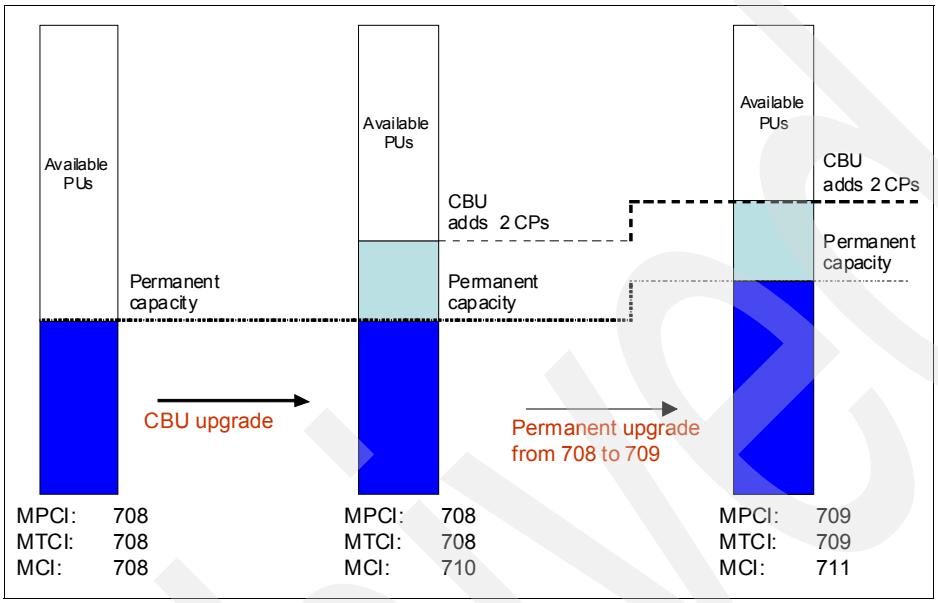

Figure 2-16 Permanent upgrade with CBU active

Considerations for permanent upgrade with active temporary upgrades

The ability to add permanent capacity while temporary capacity is active was added to the System z10 servers to help customers better manage their capacity.

General assumptions

The difference between how billable and replacement capacity is handled during a permanent upgrade is based on the following general assumptions:

- The most important assumption is that a permanent upgrade cannot be partially completed. Temporary billable capacity is converted to permanent capacity, assuming the customer has determined an economic advantage to having this capacity as part of the permanent configuration rather than having continuously or near continuously active temporary capacity.
- Keeping the replacement capacity intact, on top of the permanent upgrade, is based on the assumption that because temporary capacity is replacing lost capacity within the enterprise, any additional permanent capacity is considered a net addition to the enterprise.

Note: Not all customer conditions, preferences, and occurrences of temporary capacity activations can be taken into consideration to ensure that the permanent upgrade can be completed. Because not all customer situations can be considered here, only the customer can best determine the proper course of action.

Situations that can cause conflicts while a temporary upgrade is active

Therefore, certain situations exist when a permanent upgrade can cause conflicts while a temporary upgrade is active, as follows:

- When not enough PUs are available to fulfill the request

  A permanent upgrade can fail if not enough available PUs exist to fulfill the upgrade request. Permanent upgrades cannot finish with partial success.

  Note: This situation can only happen if too many or if all available PUs are used for a temporary capacity. If the permanent upgrade can convert temporary resources activated by a billable record (On/Off CoD) to permanent resources, depending on actual configuration, the upgrade might succeed.

- When a temporary upgrade is active, which increases the capacity level, and the resulting configuration after the permanent upgrade would result in the capacity settings area of 713 - 764.

  To better understand what could happen if such requests are allowed, see 2.2.3, “Subcapacity activation” on page 43.

  If the only temporary upgrade that has increased the capacity level is a billable (On/Off CoD) temporary upgrade, and the permanent upgrade can convert those temporary capacity level increases to permanent increases, then the permanent upgrade is accepted. Otherwise the permanent upgrade is rejected.

- When pending resources are in the system

  A permanent upgrade cannot succeed if pending resources are in the system as a result of partial completion of a CBU record forced activation.

  This means that no available PUs exist in the system and any attempt for further upgrades will fail.

  Note: This situation can only happen if a forced CBU record with pending resources is active in the system. It cannot happen for On/Off CoD or CPE records, or for a non-forced CBU record.

  Forced CBU can only be active if invoked through the APIs prior to System z10 or by selecting the `Force` option on the Add Processors panel.

- When an upgrade can bring a system to a non-existing target configuration

  If a permanent upgrade is requested while temporary records are active, the permanent upgrade can fail if the resulting configuration can result in a non-existent capacity setting. This applies to CPs and subcapacity settings only.

  If the permanent upgrade can convert temporary resources activated by a billable record (On/Off CoD) to permanent resources, depending on actual configuration, the upgrade might succeed.

- When an active temporary upgrade can cause a feature-code authorization error

  A permanent upgrade can fail if it leaves an active temporary record in an unauthorized state. This can happen with an active CBU record only *when both of the following statements are true*:

  – This record increased the CP capacity level.

  – This CBU record contains fewer CP CBU feature codes than the new number of permanent CPs plus the number of active CBU CPs.

- When a permanent upgrade can remove a dedicated processor or the last of a processor type that is in use by a logical partition

  A permanent upgrade cannot be performed if it involves removing the last of a processor type that is still in use by a logical partition. Permanent upgrade is also rejected if it removes a dedicated processor from an active partition.

  This situation is independent of concurrently active temporary resources and is mentioned here only for completeness.

- When a permanent upgrade is performed as part of an MES upgrade in an incomplete power-on reset state and a temporary upgrade was active before the server was brought down

  Important: Although this situation does not involve an On-line Permanent Upgrade but rather an MES upgrade, it might result in unexpected behavior.

  During power-on reset, the new permanent configuration is activated first, before the temporary capacity records are reactivated. Depending on the new permanent configuration, the activation of a previously active temporary upgrade record can fail because to any one of the conditions described previously (not enough available PUs, invalid target configuration, or feature code authorization error).

  If more that one temporary capacity record was active, after power-on reset, the server tries to activate all of them even if some fail.

  The order of activation is determined by the slot number in which the record is installed, with the exception that a forced CBU record is activated last.

  When a permanent upgrade is performed in a power-on reset incomplete state, we recommend verifying that the permanent upgrades work successfully with the currently active temporary records. If in doubt, we recommend deactivating all temporary records first before taking the system down for the MES upgrade.

2.2.3 Subcapacity activation

Certain combinations of multiple active records can cause the authorization space for subcapacity configuration to move into the capacity settings 713 - 764 area. Depending on the sequence of activation and deactivation of the records, the result could lead to an invalid configuration. To prevent an invalid configuration, the LIC calculates the subset of the record that can be activated without the risk of ending at an invalid configuration. Ending at an invalid configuration can happen only when subcapacity settings are used.

Figure 2-17 on page 44 shows an example how invalid configuration could be entered. The state is as follows:

- The permanent capacity is 605.
- The On/Off CoD record is installed.
- The CBU record is installed, with 18 CP CBU features.

If On/Off CoD upgrades the server to 707, and subsequently a CBU for additional ten CPs is activated, final capacity setting becomes 717.

If On/Off CoD is deactivated while CBU is still active, it can result in a capacity setting of 615, which is not a valid configuration.

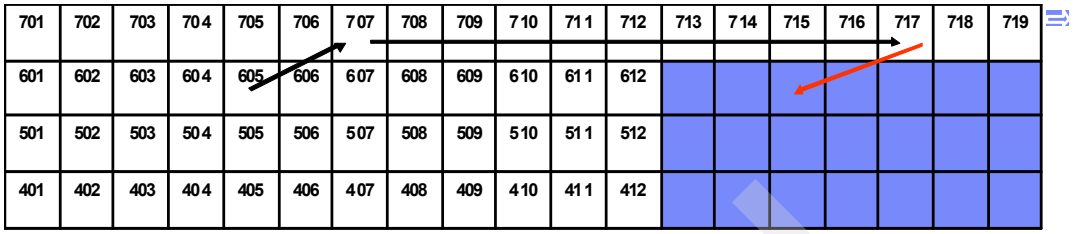

Figure 2-17 On/Off CoD and CBU entering the capacity settings 713-764 area

LIC rules for subcapacity settings

To prevent invalid configurations, three rules are enforced by the LIC for subcapacity settings.

- Only a record that increases the capacity level is allowed to enter the capacity settings 713 - 764.
- When more than one record increases the capacity level, none of them can enter the capacity settings 713 - 764 area.
- If the activation of any record that changes capacity level moves the capacity setting into the 713 - 764 area, no other temporary upgrade record is allowed to activate any additional CP. More CPs can only be added by the record that changes the capacity level and enters the capacity settings 713 - 764.

The rules, plus limits included in the LIC, guarantee that when the record is deactivated, the server returns to a valid configuration.

Examples of the rules

The rules are illustrated in the following four examples.

Example one

When a record increases CP capacity level but does not enter the capacity settings 713 - 764 area, any subsequent record can be activated only up to capacity setting 712.

In Figure 2-18, the first activated record changes the capacity setting from 605 to 707. The LIC allows activation of additional CPs from any subsequent temporary record to a maximum capacity setting 712, as shown.

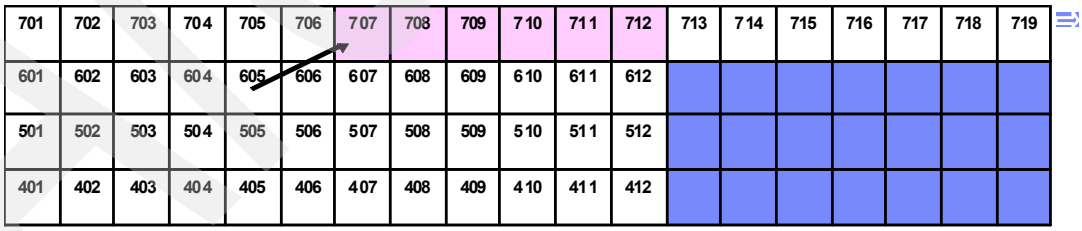

Figure 2-18 First record changes capacity level without entering the capacity settings 713-764 area

Example two

When multiple records increase the capacity level, but none of them enters the capacity settings 713 - 764 area, the LIC does not allow any record to enter the capacity settings 713 - 764 area.

Figure 2-19 on page 45 shows two records that change CP capacity levels. Entering the capacity settings 713 - 764 area is not possible. Only capacity settings in light gray are available by following the second rule.

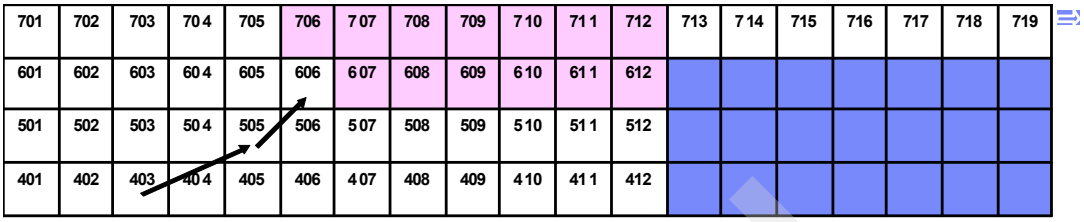

Figure 2-19 Two records change capacity levels

Example three

When a record increases the CP capacity level and enters the capacity settings 713 - 764 area, only this record can activate additional CPs, and other CoD records can only add specialty engines.

In Figure 2-20, the first record being activated changes the capacity setting from 504 to 716. While this record remains active, only this record is able to activate the remaining CPs (717 to 726 in light gray). At this time, all other temporary records can only add specialty engines, otherwise it would violate the third rule.

Figure 2-20 First record changes capacity level and enters the capacity settings 713-764 area

Example four

When the first record does not increase the CP capacity level, the second record is allowed to increase the CP capacity level and to enter the capacity settings 713 - 764 area. Figure 2-21 illustrates the following items:

- The first record upgrades a server from 607 to 611, and another record upgrades it from 611 to 714. This is a valid configuration, because the first record can be deactivated entirely without resulting in an invalid configuration.
- Adding more CPs with the first record is not possible because only a record that changed the capacity level can do it. This is in regard to the third rule. The second record can still add CPs (in light gray) if the CPs are in its authorization space.
- Both records can still remove activated CPs. That is, even the first record can *decrease* the number of activated CPs.

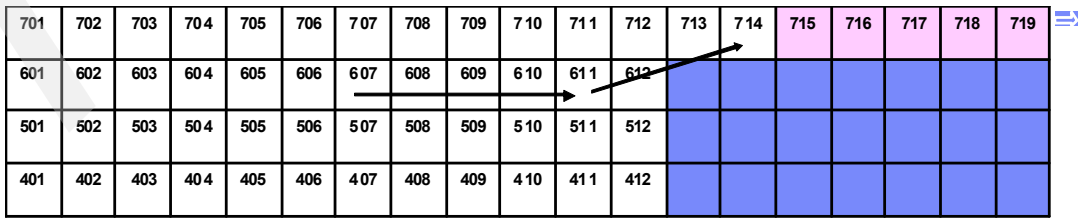

Figure 2-21 Second record changes capacity level and enters the capacity settings 713 - 764 area

2.3 Tools

The flexibility of the CoD offerings on the System z10 servers makes it important to understand how to best use the various tools for configuring, ordering, installing, using, and monitoring the offerings. This section describes the eConfig, Resource Link, HMC panels, and RMF used to perform these tasks.

2.3.1 eConfig

The eConfig tool is PC-based and is used by IBM or Business Partner personnel to configure hardware and software for a variety of IBM products. The eConfig tool provides support for the initial order configuration and upgrades involving LIC only or LIC along with hardware. Customers can see eConfig output and for this reason we decided to describe relevant CoD information here for better understanding.

To use the CoD offerings, the System z10 server must be conditioned with the Capacity on Demand enabling feature codes, which can be part of the initial server order, part of an upgrade from an earlier processor family, or be an upgrade containing only the enablement feature codes. The enablement feature codes can be installed concurrently. See Table 2-2 for the enabling feature codes.

Table 2-2 The System z10 Capacity on Demand enabling feature codes

Feature code	Description	Purpose
9900	Online Capacity on Demand Buying Enablement	Enables CIU machine profile for online ordering
9910	Capacity Backup Enablement	Enables Capacity Backup
9912	Capacity for Planned Events Enablement	Enables Capacity for Planned Events
9896	On/Off Capacity Enablement	Enables On/Off Capacity on Demand
9898	Permanent Upgrade enablement	Enables permanent upgrade ordering in CIU

The output from the configurator can be used to verify the server configuration and the current state of the conditioning. Example 2-1 shows extracts from an eConfig report.

Example 2-1 Configurator output

```
.....................
.....................

On Demand Capacity Selections:
NEW00001 - CBU - CP(2) - IFL(1) - ICF(1) - zAAP(1)
         zIIP(1) - SAP(1) - Tests(10) - Years(5)
NEW00002 - CBU - CP(2) - IFL(1) - ICF(1) - zAAP(1)
         zIIP(1) - SAP(1) - Tests(10) - Years(5)
NEW00003 - Capacity for Planned Event

5 tests are included with the initial CBU record.

System 1 - Comments

  1. Version : 20081020
```

```
20081112                    1 New Initial Order
18:17:10
 Warning: This configuration has errors and cannot be ordered. This output
 is provided for reference only and should not be used for quoting purposes
       Please view the errors in the messages tab for more information

Initial Hardware

Product             Description             Qty
_______  ________________________________  _____

2097-E26 IBM System z10 Enterprise Class      1
    0114  I/O Cage Full Card Airflow          30
.........................................
.........................................
    6702  2-Way Processor CP7                  1
    6810  CP7                                  2
    6811  IFL                                  1
    6812  ICF                                  1
    6814  zAAP                                 1
    6815  zIIP                                 1
    6817  1 CBU Year                          10
    6818  CBU                                  2
    6819  5 Additional CBU Tests               4
    6820  1 CBU CP                            20
    6822  1 CBU IFL                           10
    6824  1 CBU ICF                           10
    6826  1 CBU zAAP                          10
    6828  1 CBU zIIP                          10
    6830  1 CBU SAP                           10
    6833  Capacity for Planned Event           1
    7139  702 Capacity Marker                  1
    9896  On/Off COD authorization             1
    9898  Perm Upgr authorization              1
    9900  On Line CoD Buying (Flag)            1
    9910  CBU authorization                    1
    9912  CPE authorization                    1
    9968  z10 EC Site Tool Kit                 1
    9970  Weight Distribution Kit              1
```

The enablement feature codes 9896, 9898, 9900, 9910, and 9912 are all included in the order in Example 2-1 on page 46; all of the offerings can thus be ordered for this server.

Included in the eConfig report, shown in the example, is a section called `On Demand Capacity Selections` that shows the configured offerings. The example shows two CBU offerings, additional tests, and the term of the contract selected for each record. CPE is also configured for this order. In the configuration information are the totals for the resources ordered in the CBU offerings, including the physical resources, the total of additional CBU tests, and the total of the CBU-years that were ordered. When multiple records are ordered, the breakdown of the offerings is useful for verifying the configuration. Example 2-1 on page 46 shows `20` of `1 CBU CP` features. It is calculated as (2 x 5) + (2 x 5) = 20. It comes from two CP CBUs for five years from the first record and two CP CBUs for five years from the second record. When the

result is bigger than 25, the 25 CBU CP feature is used. For this reason, checking the `On Demand Capacity Selections` section is easier.

Note: There is no notion of On/Off CoD offerings in the order itself, only the enabling feature code. The On/Off CoD offerings can only be ordered by the customer through the Resource Link CIU application.

2.3.2 Resource Link

Resource Link allows the customer to configure, order, and download the various offerings. Using the Resource Link applications, you can create, replenish, cancel, and view upgrade orders, and also consult ordering history. The Resource Link Web site is:

`http://www.ibm.com/servers/resourcelink`

All offerings ordered through CIU Express may be staged in Resource Link for an extended period of time, unless one of the following conditions exists:

- The order is canceled by the customer.
- The system is no longer under warranty or under an IBM maintenance service agreement.
- The permanent PU or memory configurations are changed.

CIU Express is the default option when ordering on Resource Link. With CIU Express, the order can be ready for retrieval in a significantly shorter time.

Concurrent upgrades are not restricted to the physical resources currently installed in the System z10 server. Although the Concurrent Book Add (CBA) option on z10 EC allows concurrent installation of additional physical processors and memory, it requires a traditional MES ordering and IBM service personnel to install the required hardware. The CBA option is excluded from the Capacity on Demand features.

To use Resource Link, first register a user ID in Resource Link. This user ID must be included in CIU contract supplements. IBM then links the user ID to the CIU machine profile, enabling you to start using Resource Link for configuring and ordering Capacity on Demand upgrades. Figure 2-22 shows an example of a machine profile on Resource Link.

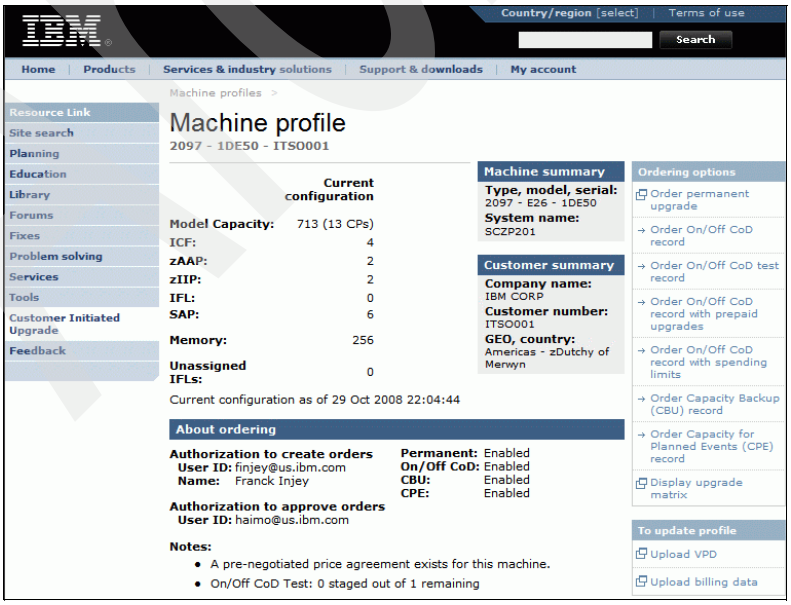

Figure 2-22 Example of a machine profile

With Resource Link, you may order only those configurations that are deemed valid by the order process, within the limits of already installed hardware resources.

Note: Temporary memory upgrades are not available. Memory upgrades are only offered as permanent upgrades.

In Figure 2-22 on page 48, the ordering options shown on the right list the order possibilities. One of the following options can be selected:

- Order an On-line Permanent Upgrade
- Order an On/Off CoD record
- Order an On/Off CoD test record
- Order an On/Off CoD record with prepaid upgrades
- Order an On/Off CoD record with spending limits
- Order a Capacity Backup (CBU) record
- Order a Capacity for Planned Events (CPE) record

For details and guidance, see *IBM System z10 Capacity on Demand User's Guide,* SC28-6871.

On-line Permanent Upgrade

On-line Permanent Upgrades are ordered through Resource Link. In the ordering panel, you specify the target configuration. See Figure 2-23 for an example. In this case, the permanent upgrade adds one CP, one zAAP, one zIIP and two IFLs.

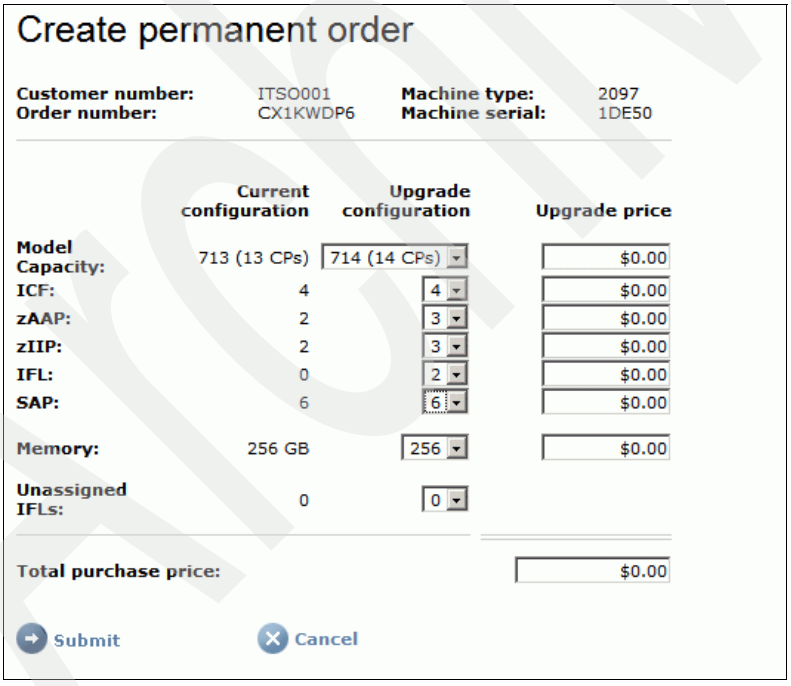

Figure 2-23 Resource Link panel for an On-line Permanent Upgrade

On/Off Capacity on Demand

Resource Link is the only place where the customer can order the On/Off CoD temporary upgrades records. It is also the only place for ordering a record replenishment.

Note: Resource Link limits the amount of resources customer can order based on the rules described in “Ordering and limits for On/Of CoD” on page 31.

Four types of On/Off CoD records can be ordered, each uses slightly different windows. The first window (Step 1), which is similar for each of the four types, is where to specify resource limits that can be activated at activation time, as shown in Figure 2-24. To achieve the highest flexibility, we recommend allowing maximum resources for the record as offered by default. This way, at activation time, a decision can made regarding which resources to activate.

No replenishment date is specified for On/Off CoD test record and On/Off CoD record with prepaid upgrades.

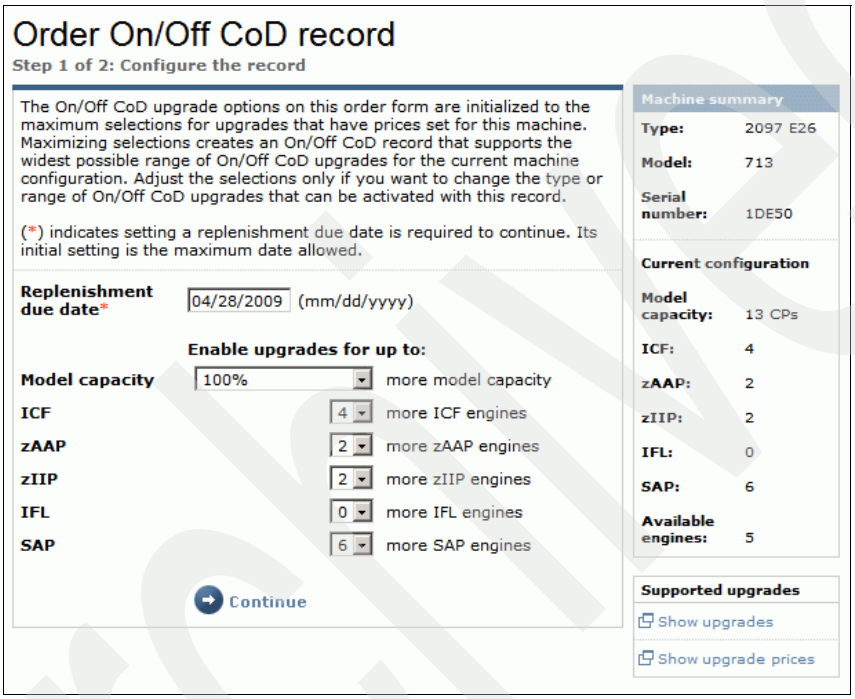

Figure 2-24 Order On/Off CoD record without capacity tokens

On/Off CoD record with prepaid upgrades and On/Off CoD record with spending limits orders require you to define the number of capacity tokens that the record will contain. You do this on a different window.

Figure 2-25 on page 51 shows the Step 2 window for On/Off CoD record with prepaid upgrades. This window helps you to define capacity tokens. You define the proposed upgrade and number of days for which the upgrade should be active. The proposed upgrade is used only for calculation, it is not stored anywhere in the record and, at activation time, you may choose any valid target configuration. Use the window to select and add multiple upgrade configurations of one or more engines types in a single prepaid order. The example in Figure 2-25 on page 51 shows 630 MSU tokens, which represents an upgrade from 713 to 715 (126 MSU difference) for five days.

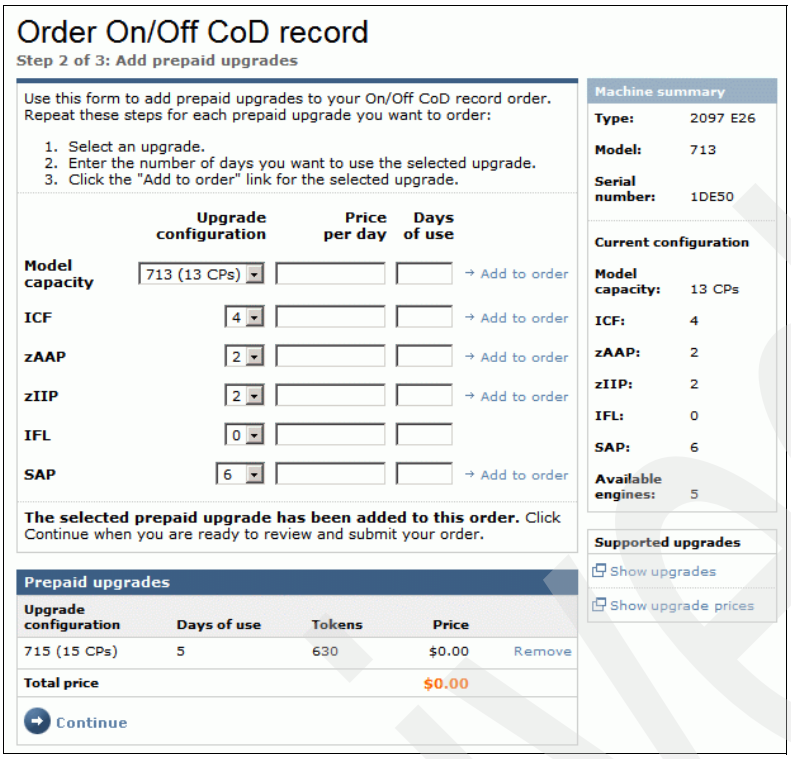

Figure 2-25 Order On/Off CoD record with prepaid upgrades

Figure 2-26 on page 52 shows the Step 2 window for On/Off CoD record with spending limits. This window helps you define capacity tokens. For each resource type, you specify the spending limit that must not be exceeded. The number of tokens is calculated. You determine when and how these tokens are used. They are paid for only after they are used and only those actually used are billed. Use the window to help control financial liability for possible upgrade without spending money before the upgrade is necessary. Example in Figure 2-26 on page 52 shows spending limits in fictional currency named Bucks. Spending limits are set to 50 000 bucks for CPs, 10 000 bucks for ICFs and zAAPs.

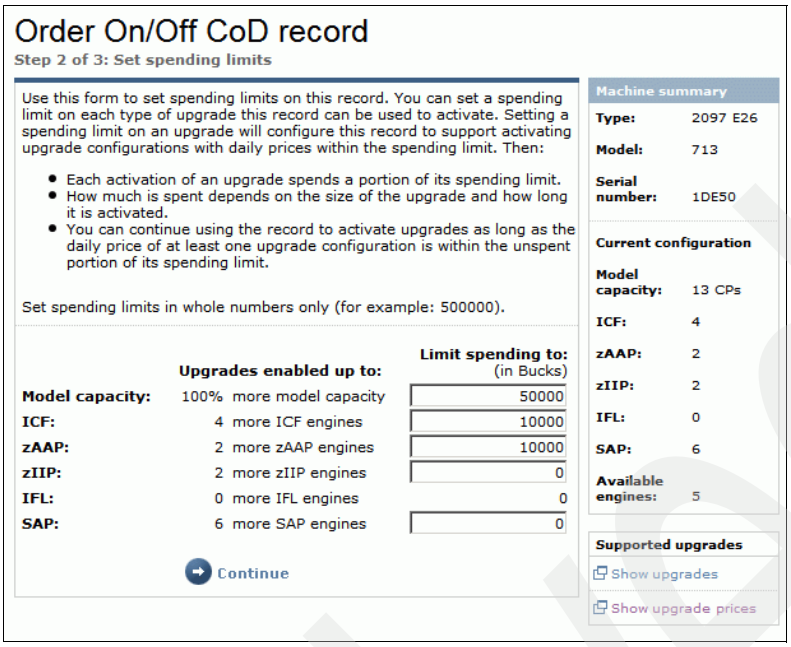

Figure 2-26 Order On/Off CoD record with spending limits

The last step in each case allows to review and finally submit the order. Figure 2-27 shows the Step 3 window for On/Off CoD record with spending limits order, where spending limits from Figure 2-26 were translated to 282 MSU tokens, two ICF tokens, and five zAAP tokens.

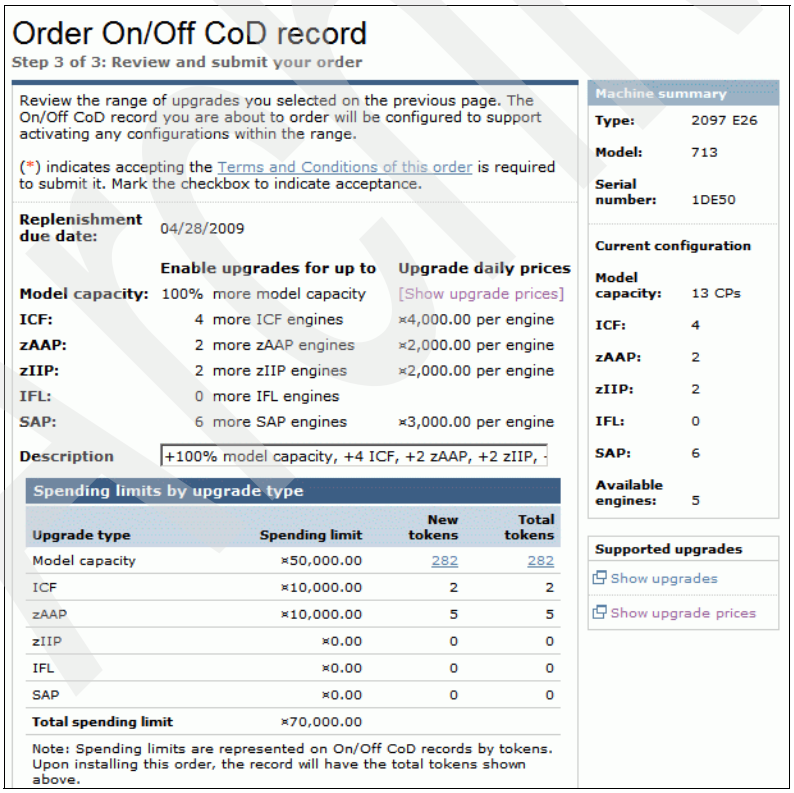

Figure 2-27 Review and submit On/Off CoD order

Capacity Backup

Ordering CBU on Resource Link consists of two steps:

1. The resources for CBU are defined.
2. The order is reviewed and submitted.

Figure 2-28 shows a CBU ordering window. The figure shows an order of a record that allows activation of up to 3 more CPs, 2 zAAPs, and 4 IFLs. The figure also shows that the contract will be valid for five years and will allow 15 CBU tests.

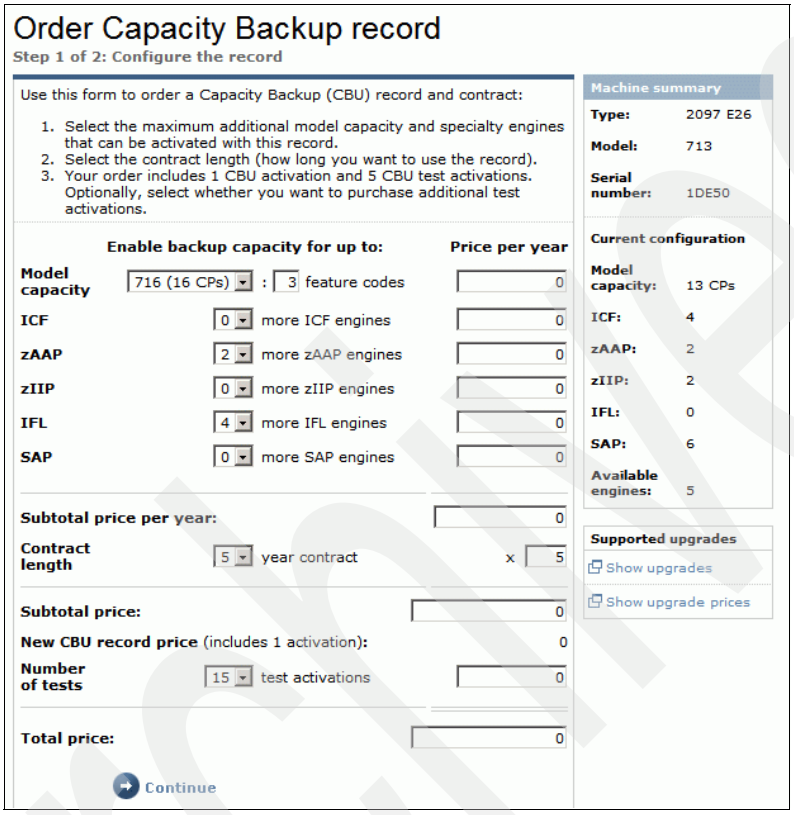

Figure 2-28 Order Capacity Backup record

When the CBU record is downloaded, Resource Link starts the billing.

Capacity for Planned Events

Figure 2-29 on page 54 shows a window for CPE ordering. There is nothing to specify for CPE because it makes available all dormant resources on the server.

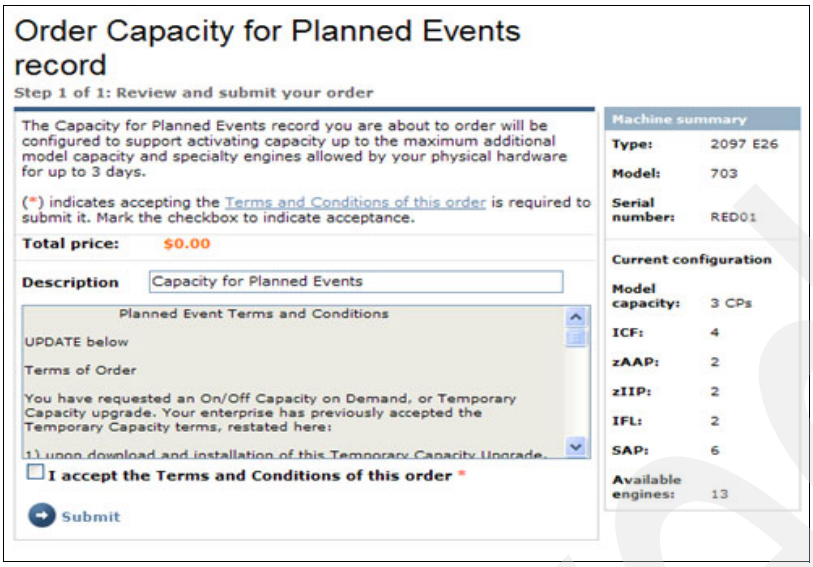

Figure 2-29 Order Capacity for Planned Events record

2.3.3 Perform Model Conversion task

After ordering a permanent or temporary upgrade, the order status will be *New Order* on the Resource Link profile. When the upgrade is ready to be retrieved, Resource Link generates an e-mail notification to the customer and changes the status of the order to *Download ready*.

The upgrade record can be retrieved and applied by using the Perform Model Conversion task on the Support Element (SE).

1. On the HMC, select **Systems Management** → **Target server** → **Recovery** → **Single Object Operations** to reach the SE console application.
2. From the SE console, select **CPC Configuration** → **Perform Model Conversion**, as shown in Figure 2-30.

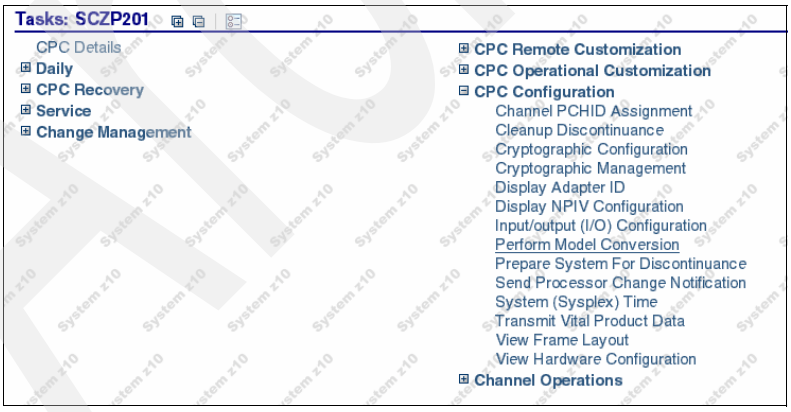

Figure 2-30 CPC configuration

The Perform Model Conversion panel opens, as shown in Figure 2-31 on page 55.

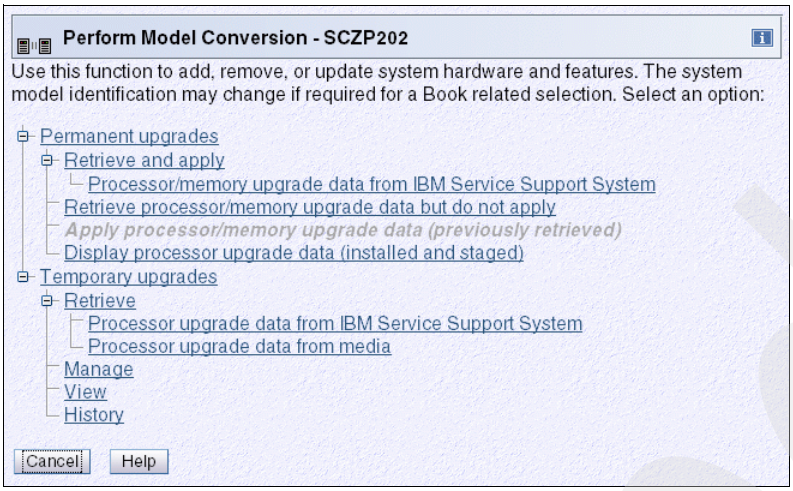

Figure 2-31 Perform Model Conversion panel

For further details, refer to *IBM System z10 Capacity on Demand User's Guide,* SC28-6871.

2.3.4 Remote Support Facility

To install a record, a connection to the IBM Remote Support Facility (RSF) is required. Therefore, at least one HMC, which has the related System z10 server defined, must have RSF enabled, and correctly defined. RSF connection and setup are described in *Hardware Management Console Operations Guide Version 2.10.0,* SC28-6867.

An easy way to verify a connection to RSF is by initiating a *Test automatic problem reporting* problem type, as follows:

1. Log on to the HMC by using the sysprog role.
2. Select **CPC** → **Report a Problem** → **Test automatic problem reporting**. See Figure 2-32 for an example.

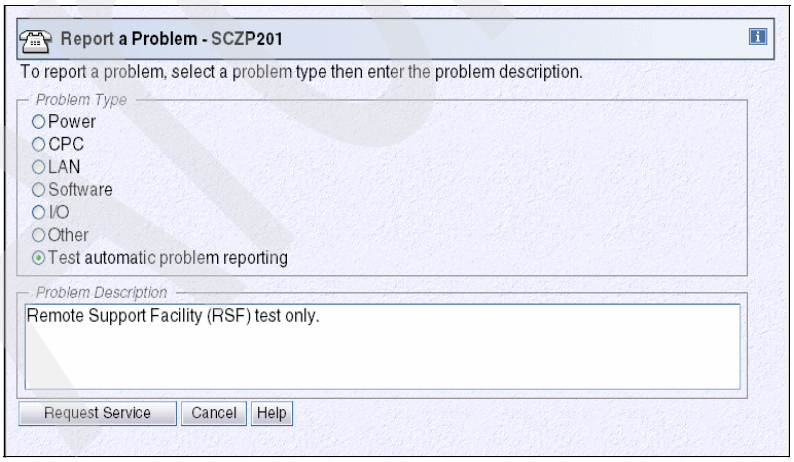

Figure 2-32 Panel to test RSF connection

Note: For secure accounts, where RSF access is prohibited, RPQ 8P2305 must be used to retrieve the offerings on a portable media, such as a CD or a USB key.

2.3.5 Scheduled operation task

For On/Off CoD, a scheduled operation can be configured on the SE to automatically deactivate the processors prior to expiration. Scheduled operation on SE activates or deactivates all resources in a record at once. See Figure 2-33, Figure 2-34, and Figure 2-35 for examples.

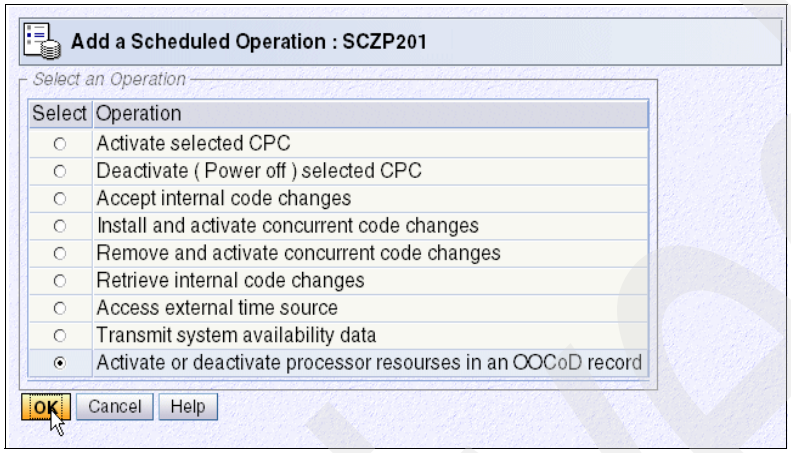

Figure 2-33 Set up a Scheduled Operation - 1 of 3

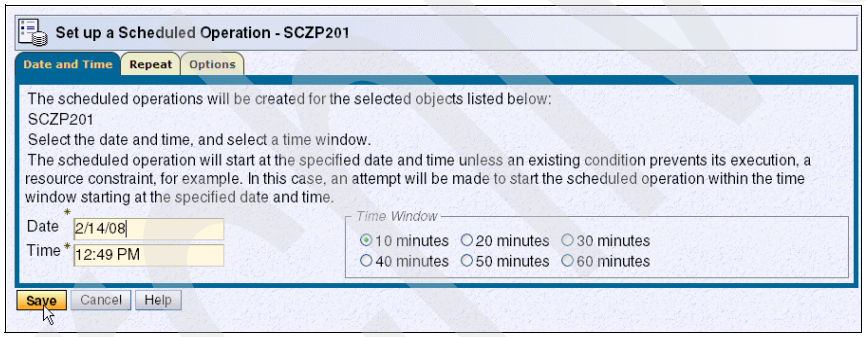

Figure 2-34 Set up a Scheduled Operation - 2 of 3

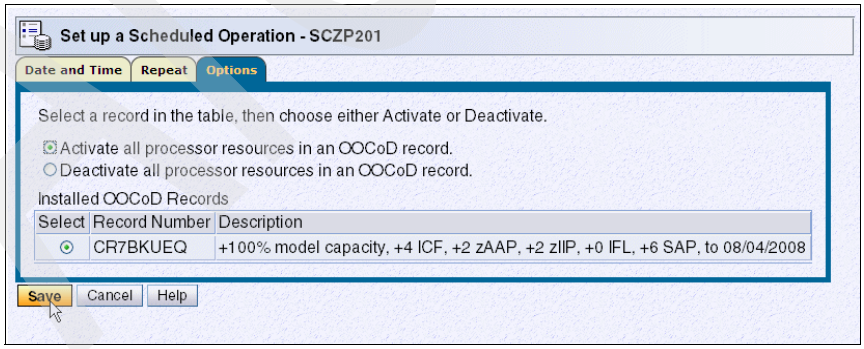

Figure 2-35 Set up a Scheduled Operation - 3 of 3

2.3.6 z/OS Resource Measurement Facility

z/OS Resource Measurement Facility (RMF) and System Measurement Facility (SMF), or equivalent products, can be used to collect data for deeper analysis.

RMF Monitor I is a very useful tool for long-term studies, to analyze trends and help with capacity planning. RMF Spreadsheet Reporter was created as the workstation front-end to the RMF postprocessor. Both extract information from SMF data to create different kinds of reports, including CPU reports.

In this book, we focus on RMF tools for short-term analysis. RMF Monitor III is better suited to get an immediate view of the situation, not to analyze long-term history of the system. The CPC Capacity report shows relevant information about partition activity and processor utilization. A sample CPC Capacity report is shown in Figure 2-36.

```
                        RMF V1R9   CPC Capacity                  Line 1 of 24
Command ===>                                                   Scroll ===> CSR
Samples: 100     System: SC80  Date: 02/15/08  Time: 17.00.00  Range: 100   Sec

Partition:   A01        2097 Model 713
CPC Capacity:    1076   Weight % of Max: ****   4h Avg:    1   Group:   N/A
Image Capacity:   828   WLM Capping %:    0.0   4h Max:    3   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
            Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total
*CP                      38.0                        0.2    31.1   31.3
A0E           0   166  NO   2.0        100     100    0.0    15.4   15.4
A0F           0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A01           0     1  NO   9.0        0.2     0.2    0.0     0.1    0.2
A02           0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E           0   165  NO   2.0       99.9    99.9    0.0    15.4   15.4
A1F           0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11           0     1  NO   9.0        0.2     0.2    0.0     0.1    0.1
A12           0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17           0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                             0.2            0.2

*AAP                     4.0                         0.0     0.0    0.0
A01                   NO   2.0        0.0     0.0    0.0     0.0    0.0
A11                   NO   2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                             0.0            0.0

*ICF                     0.0                         0.3     0.0    0.3
PHYSICAL                                             0.3            0.3

*IIP                     4.0                         0.0     0.0    0.0
A01                   NO   2.0        0.0     0.0    0.0     0.0    0.0
A11                   NO   2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                             0.0            0.0
```

Figure 2-36 Sample RMF CPC Capacity report

The top of the report shows information about the z/OS system such as where the report was generated, and the date and time.

The next three lines contain information about the partition where this z/OS system is running, the server (CPC) model, and capacity and Image Capacity (both in MSUs). Image Capacity is reported as the capacity limit for the partition if it has been defined, or it is calculated based on the logical CP configuration of the partition and capping state.

The bottom part of the report lists the active partitions in the CPC that utilize CPs, zAAPs, ICFs, and zAAPs. Listed are the actual capacity in MSUs used (during this 100-second interval), the number of processors online, and the logical and physical utilization of the processors.

The RMF CPC Capacity report shows information at the partition and CPC-level. To determine which workloads are using these processors, a more detailed report is necessary. The RMF System Information report is very useful to display the current workloads and information about their performance. A sample System Information report is shown in Figure 2-37.

```
                        RMF V1R9   System Information            Line 1 of 33
Command ===>                                                  Scroll ===> CSR

Samples: 200     System: SC80  Date: 02/16/08  Time: 17.00.00  Range: 200   Sec

Partition:    A01     2097 Model 713          Appl%:      54  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   55     EAppl%:     54  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   87     Appl% AAP:  60  Time:   23.22.20
IIPs Online:  2.0                            Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER
*SYSTEM      35  155   20        0.00  7.0  0.0   13.2  0.0  0.0  0.0  0.0
*TSO          1    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       35   20   20        0.00  7.0  0.0   13.1  0.0  0.0  0.0  0.0
*STC         31  131    0        0.00  0.0  0.0    0.1  0.0  0.0  0.0  0.0
*ASCH         0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  31  135    0  .000  0.00  0.0  0.0    0.1  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S  20  109    0  .000  0.00  0.0  0.0    0.1  0.0  0.0  0.0  0.0
          1  20  109    0  .000  0.00  0.0  0.0    0.1  0.0  0.0  0.0  0.0  N/A
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  35   20   20  .000  0.00  7.0  0.0   13.1  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  35   20   20  .000  0.00  7.0  0.0   13.1  0.0  0.0  0.0  0.0
          1  35   20   20  .000  0.00  7.0  0.0   13.1  0.0  0.0  0.0  0.0 2.21
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 2-37 Sample RMF Sysinfo report

This example was customized to display the local performance index in the last column. We use the local performance index (PI) because this metric can be used to trigger z/OS capacity provisioning as explained in 2.6.2, “CPM definitions” on page 72.

RMF Monitor III Data Portal for z/OS provides browser-based interface for viewing the RMF Monitor III data.

For more information about RMF reports, see the following publications:

- *z/OS Resource Measurement Faclity (RMF) Programmer’s Guide,* SC33-7994
- *z/OS Resource Measurement Facility (RMF) Report Analysis,* SC33-7991

RMF Performance Monitoring (PM) is a workstation tool that offers an alternative to Monitor III. It is a short-term study tool that shows similar reports to Monitor III, but in a graphical view. RMF PM monitors sysplex-wide information. Data views can be easily defined to display different metrics.

Figure 2-38 shows an example of an RMF PM window. The example shows three data views. The first is a list metric view, which offers a list of values for each interval (a value for each service class in this case). The other two data views are single metric; they offer a value every time interval (execution velocity for JAVAWKLD service class in one case, and PI for the same service class in the other).

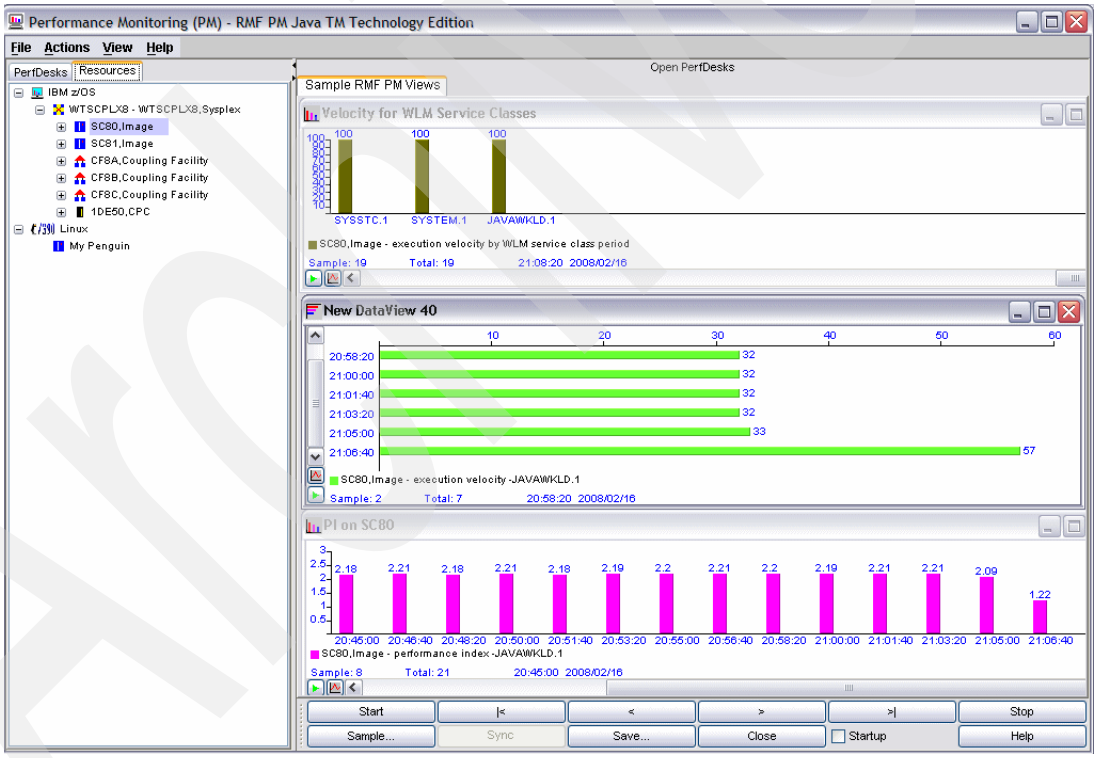

Figure 2-38 Sample RMF PM

RMF PM eases analysis tasks by providing point-and-shoot fields. More information about RMF and available tools can be found on the z/OS RMF Web site:

http://www.ibm.com/servers/eserver/zseries/zos/rmf

2.3.7 Additional tools

Other tools are available to help you understand how the offerings can be used to align server resources to meet the requirements of the business workloads.

For example, the Processor Capacity Reference for System z (zPCR) can be used to model possible configurations. The zPCR tool can be downloaded from:

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

2.4 Hardware planning

This section discusses the activation, deactivation, expiration, auto deactivation, and replenishment of temporary resources.

2.4.1 Activation of temporary resources

The System z10 servers can have up to eight temporary records concurrently active, with only one On/Off CoD active at any one time. The remaining records can be any combination of CBU or CPE. In addition, the records can be activated in any sequence.

The System z10 also offers the possibility to activate subsets of each record, so that the customer can decide, at the point of activation, how much capacity is needed for any given situation and activate only what is necessary to meet the requirement. The activation level can be changed later, and temporary resources can be added or removed as necessary without deactivating the entire record.

CoD provides enhanced flexibility, which means less technical constraint to adjust to diverse requirements. On the other hand, operations should be prepared to use the new functions to react to business situations.

Although this document describes activation or deactivation of a record, the controls provided by the Perform Model Conversion task on the Support Element do not show an activate or deactivate function. To control temporary upgrades, select **Add processors** or **Remove processors**. A temporary capacity record is active when adding processors results in a positive count, or inactive when removing processors brings the count back to zero.

The activation and deactivation can be a manual process, or automated through the APIs.

Note: Activating CBU does no longer requires a password to be exchanged with RSF.

Predefinition of resources

All CoD upgrades can be installed and activated concurrently. If the resources are added to expand an active logical partition, and if the operating system supports it, the upgrade is also nondisruptive.

With appropriate planning, the activation of a temporary upgrade is most often nondisruptive. If, however, the resources are used to activate new partitions, or if the existing partitions have not been prepared for additional resources, the upgrade is disruptive, because it necessitates creation, or modification, of logical partition activation profiles.

If CBU or CPE resources are activated on System z10 servers that are not part of a Parallel Sysplex with the servers for which they are replacement capacities, the activation is

disruptive. If the System z10 replacement capacity and the servers they are designed to replace are part of the same Parallel Sysplex, activation of the permanent or temporary record can be nondisruptive.

On System z10 servers, the maximum number of logical partitions are reserved:

- 60 for the z10 EC
- 30 for the z10 BC

Processor resources, initial and reserved, are defined in the image profile for each logical partition. The ability to configure additional hardware resources nondisruptively depends on functionality in the operating system executing in the logical partition, as follows.

- z/VM V5R3 with APAR VM64249 supports adding processors nondisruptively, but does not support nondisruptive addition of memory.
- z/VM V5R4 supports adding processors and memory nondisruptively. A logical partition that executes this level of z/VM must have both initial and reserved memory defined if dynamic addition of memory is needed. The dynamic addition of memory is performed by using the SET STORAGE command with the plus sign (+) parameter.
- z/OS supports nondisruptive addition of memory and nondisruptive addition of processor resources. Therefore, defining reserved processors, if partitions have to increase the number of processors nondisruptively, is mandatory. z/OS V1R10 supports the Dynamic Processor Add task and therefore does not require reserved that processors be defined in an image profile.
- If logical partitions executing z/OS require nondisruptive addition of memory, they must be defined with initial and reserved memory.

  In z/OS, the appropriate value must be set in the RSU parameter in IEAOPTxx in SYS1.PARMLIB. For details, see: *z/OS MVS Initialization and Tuning Reference,* SA22-7592.

The maximum number of processors, initial and reserved, that can be defined to a logical partition, depends on the System z10 model. For z10 EC models up to model E56, the maximum number of logical processors that can be specified is 56. For the E64 up to 64 logical processors can be specified. z10 BC supports up to 10 processors.

For maximum flexibility in logical partitions running z/OS prior to V1R10, defining as many reserved processors as possible is advisable. The total number of initial and reserved defined processors should not exceed the number of processors supported by the operating system. z/OS V1R9 and later supports up to 64 processors, z/OS V1R8 and z/OS V1R7 support 32 processors, including CPs, zAAPs, and zIIPs.

z/VM V5R3 and later supports up to 32 processors. Defining reserved processors to z/VM logical partitions is not required because dynamic addition of processors through Logical Processor Add task is supported in z/VM V5R3 and later releases.

Consider the following general recommendations:

- To take advantage of all physical processors in the server, the total number of initial and reserved logical processors in a logical partition should be equal to or higher than the number of available physical processors.
- Adding processors to the shared pool provides additional capacity to all logical partitions that have defined shared processors. The gain for each logical partition is proportional to the number of logical processors and weights that are defined for each logical partition.

  Every time a temporary resource is added or removed, the processing weights for the considered resource pool change.

The processing weight values for use when specific logical partitions are activated or deactivated should be calculated in advance, and be readily available after resources in a CoD record are activated or deactivated. Make sure you have processing weights, or rules to recalculate the weights, to match each business situation.

- If HiperDispatch is enabled, logical processors have high, medium, or low vertical polarity depending on the configuration of the logical partition.

  The relative processing weight that is defined for the logical partition effectively defines the amount of physical processor cycles the logical partition is entitled to have. The logical partition's processing weight has a direct effect on the number of high-polarity processors that the logical partition can have. An important recommendation is to set processing weights correctly to ensure that the workload is optimally allocated to the desired set of high-polarity processors.

  Note that when HiperDispatch is enabled and running on a logical partition, the partition's specialty processors, such as zAAPs and zIIPs, will also be running with HiperDispatch enabled. Plan for processing weights for all of the processor types to achieve the desired level of vertical processor allocations for each type.

  HiperDispatch has low or no influence on z10 BC servers.

- The actual utilization reported by monitors such as RMF can be different from the weights. Certain logical partitions can receive more than their share if other partitions are not using their share. The number of CPs configured online also affects the maximum resource allocated to the logical partition. For example, a partition configured with two CPs online on a ten-way CPC can never be allocated more than 20% of the server CP resources, no matter what its processing weight is.

For more information about processing weights, see *IBM System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide,* SB10-7153.

Limiting temporary resources

The Change Activation Levels panel is shown in Figure 2-39. You may activate only a subset of the record by choosing from among several Model-Capacity Identifiers, and then either activating all or some of them, or a subset of specialty engines. You may change the activation level any time later as necessary without deactivating the record.

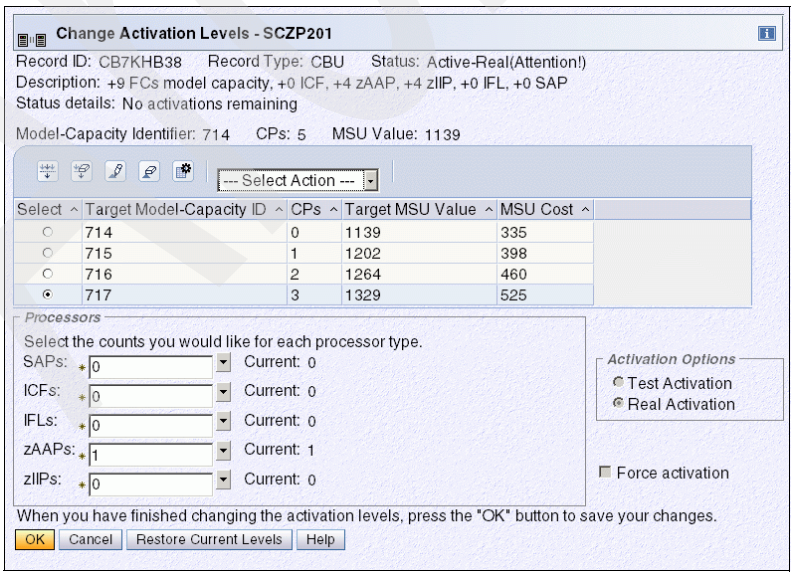

Figure 2-39 Change Activation Levels

2.4.2 Deactivation of temporary resources

Consider the following information when you prepare to deactivate a record:

- If the processors were made available into the shared pool without configuring any processor online to any logical partition, no preparatory actions are required. Deactivating the active record is sufficient because all the logical partitions defined with shared processors fall back to exactly the same situation as before the record was activated.
- If the processors were made available into the shared pool and reserved processors were configured online in logical partitions (or a new logical partition with shared processors was activated), they should be configured offline (or logical partition deactivated) before record deactivation. Otherwise, the logical-to-physical ratio would become different than before the activation of the record, which could lead to performance issues. Even if fewer physical processors are available in the shared pool than logical shared processors assigned to any logical partition, the record can be deactivated.
- If processors were made available and new dedicated processors were brought online, they should be configured offline before record deactivation because record deactivation removes processors only from the shared pool. If the processors are not configured offline and there are active partitions using shared processors:
  - Deactivation can fail if it would bring the number of processors in the shared pool to zero.
  - Although deactivation can proceed if at least one processor remains in the shared pool, the ratio of shared logical processors to physical processor will be changed.

For example, suppose the permanent capacity setting is 705, and a logical partition A has five online CPs. If the CPE is activated with eight CPs, then a logical partition B with eight dedicated CPs is activated. While logical partition B is active, deactivating temporary resources is not possible because capacity setting 705 cannot support eight dedicated CPs. In this case, a message is displayed, as shown in Figure 2-40.

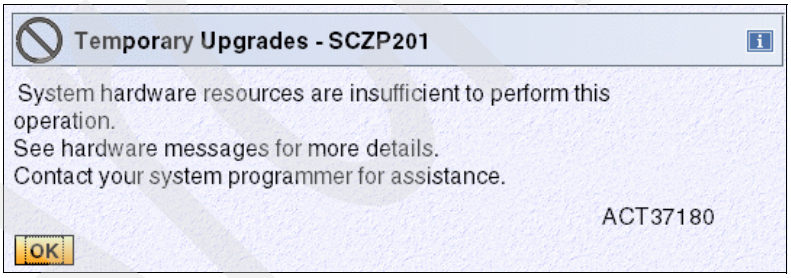

Figure 2-40 Temporary Upgrades - deactivation failure message ACT37180

2.4.3 Automatic deactivation of temporary resources

The Licensed Internal Code (LIC) deactivates all the resources in an active temporary record automatically if any one of three conditions happens:

- A temporary record expires. LIC deactivates the record within 0-24 hours after 23:59 GMT on the day of record expiration. Expiration warning messages provide the date when a record will expire. Consider the following information:
  - On/Off CoD record expires after 180 days from ordering (by default) or sooner if an earlier date was specified during ordering.
  - On/Off CoD record with prepaid tokens never expires.

  - CBU record expires when the contract period expires.
  - CPE record does not expire.
- Activation days exceed the set number. Warnings and expiration for these happen in exact 24-hour intervals from the time of activation. Warning messages for these provide the exact number of days (24 hours periods) till expiration. Consider the following information:
  - CBU test activation expires after 10 days.
  - CBU real activation expires after 90 days.
  - CPE activation expires after three days.
  - On/Off CoD test record activation expires after one day.
- Not enough capacity tokens are available for a next billing window for any of the activated resources in On/Off CoD record.

A Hardware message is generated on the SE and HMC if it is five or fewer days to record automatic deactivation. The message can occur when:

- A temporary record will expire in five or fewer days.
- Five or fewer activations days remain for a record.
- A capacity token from any token pool is five or fewer days away from being completely consumed, based on the current consumption rate.

Generating an event to signal the expiration through the SNMP APIs is possible.

Table 2-3 shows messages generated for a CPE record from its activation to automatic deactivation.

Table 2-3 Hardware Messages for CPE

Times and events	Hardware message
After CPE record activation ACT00056I Hardware problem.	WARNING! The allowed activation period of active temporary capacity record: CP7KHB6Y will expire in 3 days. These temporary resources will automatically be deactivated after expiration.
	Please deactivate these resources before expiration occurs
After 24 hours ACT00056I Hardware problem.	WARNING! The Allowed activation period of active temporary capacity record: CP78KQ93 will expire in 2 days. These temporary resources will automatically be deactivated after expiration.
	Please deactivate these resources before expiration occurs.
After 48 hours ACT00056I Hardware problem.	WARNING! The Allowed activation period of active temporary capacity record: CP78KQ93 will expire in 1 days. These temporary resources will automatically be deactivated after expiration.
	Please deactivate these resources before expiration occurs.
After 72 hours ACT00056I Hardware problem.	Active temporary capacity record: CP78KQ93 has exceeded its allowed activation duration. These temporary resources will automatically be deactivated in 0 days.
	Please deactivate these temporary resources now.

Times and events	Hardware message
Automatic deactivation	Temporary capacity record: CP78KQ93 has expired and all temporary resources for this record have been deactivated.
ACT00056I Hardware problem.	Please confirm that the dedicated/shared processor counts still fulfill your requirements.

All considerations discussed in 2.4.2, “Deactivation of temporary resources” on page 63 apply also to automatic deactivation. In addition:

- Automatic deactivation can fail if it would bring the number of processors in any used shared pool to zero.
- Automatic deactivation can proceed if at least one processor remains in the shared pool. Note that in this case the ratio of shared logical processors to physical processor will be changed.
- Automatic deactivation never removes dedicated processors.
- All resources stay active if any resource activated by the record cannot be deactivated. In other words, automatic deactivation deactivates the whole record or nothing.
- The ability to remove resources is checked every 24 hours.

Table 2-4 shows a hardware message generated when automatic deactivation is rejected.

Table 2-4 Hardware messages for rejected automatic deactivation

	Hardware Message
Automatic deactivation fails.	The SE licensed internal code has detected a problem. A synchronization problem was detected while concurrently updating the configuration, or the system was temporarily in a state that caused the operation to end without success.
ACT00060I Licensed internal code has detected a problem.	Retry the operation. If the failure continues, contact the next level of support.

2.4.4 Replenishment

Replenishment is the function used to update an already-installed temporary record. Replenishment can update installed and active temporary records.

CBU can be replenished through Resource Link or by an IBM Sales representative. On/Off CoD can be replenished only through Resource Link.

If deemed necessary by the customer, a replenishment can update an offering with resources, time limits, or tokens. After an offering is replenished and successfully processed by Resource Link, the associated offering can be retrieved and activated by using the HMC as discussed 2.3.3, “Perform Model Conversion task” on page 54.

Because a replenishment increases the limits for currently installed offerings without having to replace the record, no deactivation of resources is required.

See Figure 2-41 on page 66 for an example of a replenishment panel on Resource Link.

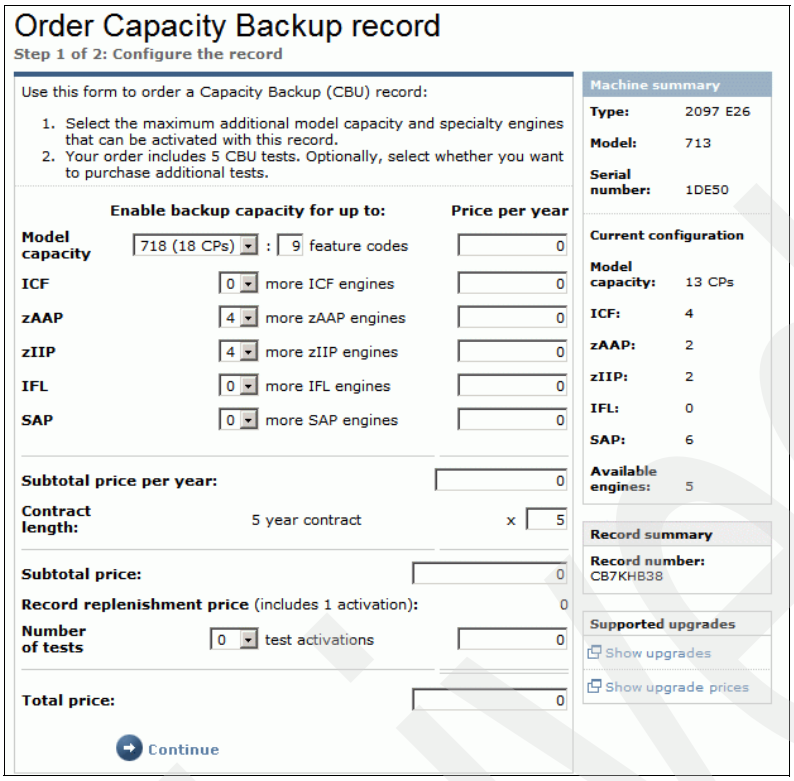

Figure 2-41 CBU record replenishment ordering

2.4.5 Carry forward of CBU features from z9, z890 or z990

In a situation where a customer's System z10 is a result of an MES upgrade from a z890, z990 or from a System z9®, the existing CBU offerings can be carried forward to the z10 server.

To take advantage of the enhanced CBU offering on the System z10, an existing CBU offering can be carried forward as it is or be split into multiple offerings on the System z10, at the customer's discretion.

The remaining time of the existing CBU offering contract is carried forward unchanged.

Splitting affects only the resources; the termination limit will be the same for all multiple records originated by splitting. However, all split records contain five test activations. Resources added during the carry-forward are charged accordingly and converting CBU engine types is possible. There is no charge for the anchor feature code of split records.

Note: The total CBU features count does not change. It is adjusted only when you create an order that requires the CBU feature count to be decremented.

For example, if we assume that a CBU offering for five CPs with a remaining entitlement time of three years was installed on a System z9, this offering could be carried forward as follows:

- No split of the CBU offering, which would mean:
  - 15 x FC 6820 in combination with 3 x FC 6817
  - FC 6820 defines a Single CBU CP per year; FC 6817 specifies the total CBU years. Five CPU CPs will be available within in the next three years in one CBU offering.

- Split of the CBU offering. The CBU offering is split into two CBU temporary records, where the first one offers two CPs and the latter offers three additional CPs in the next three years, which would mean:
  - 6 x FC 6820 in combination with 3 x FC 6817
  - 9 x FC 6820 in combination with 3 x FC 6817

2.5 Operating system considerations

This section describes the general software planning for the System z10 Capacity on Demand environment. The objective of the section is to help system programmers and system planners understand how the operating systems executing on the System z10 can be tailored to benefit from the enhanced offerings made available on the System z10.

The System z10 server hardware has the capability of making dormant hardware resources available to the workloads on the server depending on the requirements of those workloads. The utilization of these capabilities depends on functions in the operating systems. Each of the operating systems will have varying levels of support, from toleration of a new function to an exploitation of a new function, depending on version and release levels of the operating systems.

Operating systems supported on the System z10 are:

- z/OS V1R7 and later for the z10 EC (and z10 BC but under an extension of the support)
- z/OS V1R8 and later for the z10 BC
- z/VM V5R2 and later
- z/VSE™ V3R1 and later
- TPF V4R1 and z/TPF V1R1
- Linux on System z
  - Novell® SUSE® SLES9 and later
  - Red Hat RHEL4 and later

With appropriate planning, it is possible to enable server capacity growth on the System z10 servers concurrently and nondisruptively in most situations. However:

- The new capacity provisioning support is available only for z/OS V1 R9 and later releases.
- VSE, z/VSE, TPF, z/TPF, Linux for System z, and CFCC do not provide Dynamic I/O configuration support. The installation of the new hardware is performed concurrently, but defining the new hardware to these operating systems requires an IPL.
- If the concurrent upgrade is intended to satisfy an image upgrade to a logical partition, the operating system running in this partition must also have the capability to concurrently configure more capacity online. z/OS operating systems have this capability. z/VM V5R4 can concurrently configure new processors, memory and I/O devices online. Earlier releases of z/VM cannot concurrently configure memory online.
- If the concurrent upgrade is intended to satisfy the necessity for more operating system images, additional logical partitions can be created concurrently on the z10 servers, including all resources needed by such logical partitions. These additional logical partitions can be activated concurrently.
- If an operating system that is running under z/VM does not have the capability of adding more resources concurrently, it can still benefit from the z/VM capability to nondisruptively configure additional processors and memory online.

Nondisruptive upgrade planning

The primary reasons for disruptive upgrades on the System z10 servers are:

- z/OS logical partition processor upgrades when reserved processors were not previously defined.
- Logical partition memory upgrades when reserved storage was not previously defined.
- Installation of I/O cages is disruptive on the z10 EC. (Installation of I/O drawers on the z10 BC is nondisruptive)
- Linux, z/VSE, TPF, z/TPF, and CFCC do not support Dynamic I/O configuration.

Recommendations for avoiding disruptive upgrades:

- For z/OS logical partitions, configure as many reserved processors (CPs, IFLs, ICFs, zAAPs and zIIPs) as possible.
- Configure reserved storage to logical partitions that execute the z/OS operating system, and to logical partitions that execute z/VM V5R4 or later.
- Use a convenient entry point memory capacity and consider the flexible memory option, or the plan-ahead memory option to allow future upgrades using pre-installed memory cards.
- Use the plan-ahead concurrent condition for I/O for the z10 EC.

Processor identification

Two instructions are used to obtain processor information:

- Store System Information (STSI)

  STSI reports the processor model and Model-Capacity Identifier for the permanent configuration and for any additional configuration changes through temporary upgrade actions. It fully supports the concurrent upgrade functions and is the preferred way to request processor information.

- Store CPU ID (STIDP)

  STIDP is provided for purposes of backward compatibility.

Figure 2-42 on page 69 shows the relevant output from the STSI instruction on System z10 server. The STSI instruction returns the capacity setting for the permanent configuration, and the capacity setting for any temporary capacity. It allows you to distinguish between temporary billable capacity and temporary replacement capacity. The ability to distinguish is the key that makes it possible to combine several active temporary records at the same time and to upgrade a permanent capacity while temporary capacity is active. Consider the following information:

- The Model-Capacity Identifier contains the permanent capacity, the On/Off CoD capacity, the CBU capacity, and the CPE capacity.
- The Model-Permanent-Capacity Identifier contains the permanent capacity of the system.
- The Model-Temporary-Capacity Identifier contains the permanent capacity and the On/Off CoD capacity.

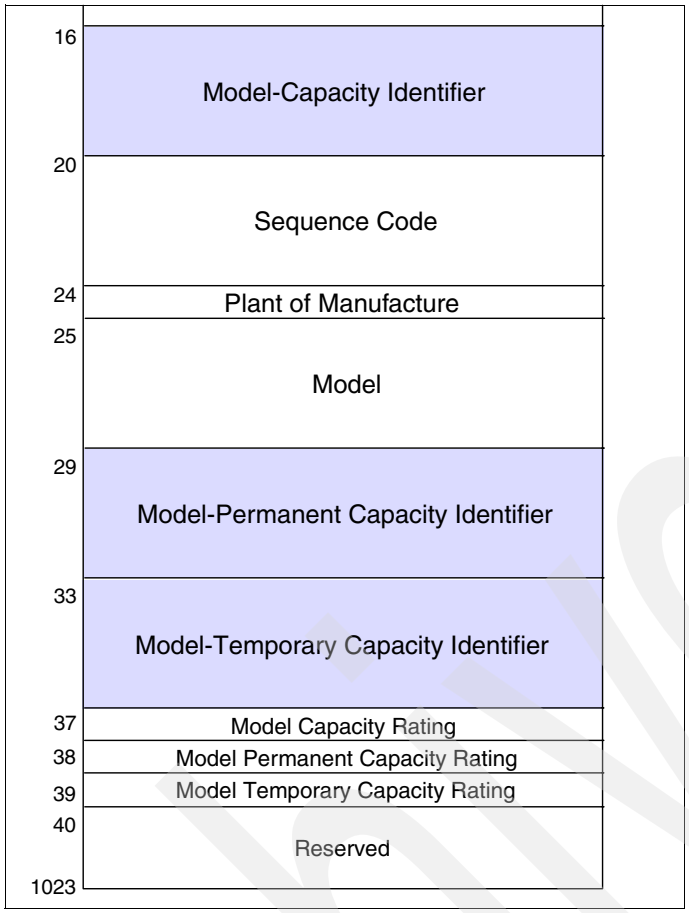

Figure 2-42 STSI output on a z10 server

The permanent and temporary capacity information is also available in SMF record type 70 subtype 1 (CPU activity); see Table 2-5. The permanent and temporary capacity identifiers and factors have been added to the record also.

Table 2-5 Model capacity fields in SMF record type 70

Field name	Offset	Length	Format	Description
SMF70MPC	x'74'	16	EBCDIC	When non-zero, the model permanent capacity identifier of the configuration. The identifier is left justified with trailing blanks if necessary.
SMF70MTC	x'84'	16	EBCDIC	When non-zero, model temporary capacity identifier of the configuration. The identifier is left justified with trailing blanks if necessary.
SMF70MCR	x'94'	4	Binary	Model capacity rating. When non-zero, this value is associated with the model capacity as identified in field SMF70MDL.
SMF70MPR	x'98	4	Binary	Model permanent capacity rating. When non-zero, this value is associated with the model permanent capacity as identified in field SMF70MPC.
SMF70MTR	x'9C'	4	Binary	Model temporary capacity rating. When non-zero, this value is associated with the model permanent capacity as identified in field SMF70MTC.

2.6 Planning for z/OS Provisioning

This section is intended for users who want to exploit the z/OS Capacity Provisioning functions of z/OS. It provides an overview and planning considerations of z/OS Capacity Provisioning Manager (CPM). For details, see *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

2.6.1 Overview of the z/OS Capacity Provisioning functions of z/OS

The On/Off Capacity on Demand offering has been available on System z servers for many years. The System Administrators have always been responsible for monitoring the capacity available and activating additional capacity when necessary. Although this approach works very well, it is error prone, inflexible, and requires human intervention.

System z10 servers, with support in z/OS V1R9 or later releases, facilitates monitoring and activation of additional resources by the system itself, automating the process. The functionality is known as z/OS Capacity Provisioning. It consists of two components:

- Capacity Provisioning Manager (CPM) - The code, executed in z/OS, that is responsible for monitoring, reporting, activating, and deactivating On/Off Capacity on Demand resources. It is part of the z/OS BCP component.
- Capacity Provisioning Control Center (CPCC) - A program that runs on a workstation. CPCC provides a GUI for defining CPM configurations and policies.

CPM modes

CPM monitors defined z/OS images (observed systems) and acts according to rules defined by the administrator. It can work in four modes:

- Manual

  In this mode, CPM does not act on its own but provides command interfaces for manual activation and deactivation of resources.

- Analysis

  In this mode, CPM monitors workloads and, according to defined rules, informs the system administrator through console messages whether additional resources are necessary. In this mode, CPM does not check whether additional resources are available.

- Confirmation

  In this mode, CPM proposes (to the system administrator) which resources should be activated and waits for command confirmation. If the command is not confirmed or is denied, CPM quiesces the proposals.

- Autonomic

  In this mode, CPM acts on its own, according to policy definitions. It activates and deactivates resources as needed and writes a message on the console after each action.

CPM resources

z/OS Capacity Provisioning, using CPM, works with On/Off CoD. CPM works with CPs, zAAPs, and zIIPs defined in On/Off CoD. The ICFs, SAPs, and IFLs are not considered by CPM.

CPM functionality

CPM monitors workloads on observed systems. It retrieves various metrics from RMF or equivalent products and determines, based on the metrics, whether to assist (if possible)

suffering workloads by adding On/Off CoD resources. System administrators define criteria to decide when a workload is considered suffering. These criteria include service class periods, performance indices, and time controls.

If a workload is considered to be suffering and additional resources might help, CPM activates those resources. Consider the following important information:

- CPM activates additional resources one at a time.

  After a resource is activated by CPM, a specified amount of time passes before the next possible activation, enabling the workload to stabilize. This time period is 15 minutes by default, but can be changed by the administrator. When this period expires, CPM determines whether additional resources are required to help the suffering workload.

- CPM can only activate installed On/Off Capacity on Demand records.

  CPM does not install staged records. If more than one On/Off CoD record is installed, the administrator can define which one should be used by CPM.

- CPM does not perform any actions inside observed z/OS systems.

  Resources are activated by using HMC API. Because CPM cannot provide any action on observed systems, it does not bring any additional processors online in the z/OS images. CPM activates a resource only if this action can help a workload without any other manual intervention. This leads to several consequences:

  – If a logical partition with a suffering workload is configured with dedicated processors, no new processors are activated. If those dedicated processors are subcapacity processors, CPM can change the capacity level of all processors if the On/Off Capacity on Demand offering allows it.

  – If a logical partition with a suffering workload is configured with shared processors, new processors are activated only if the sum of logical processors (in all logical partitions configured with shared processors) is greater than the number of physical processors in the respective processor pool.

  – Additional zAAPs are activated only if at least one zAAP is already in use by a logical partition with a suffering workload.

  – Additional zIIPs are activated only if at least one zIIP is already in use by a logical partition with a suffering workload.

  – Obviously, activating On/Off Capacity on Demand resources can influence more than simply the logical partition with a suffering workload. This point is important to have in mind; use soft or hard capping to control other logical partitions if necessary.

Figure 2-43 illustrates a sample scenario with shared processors.

```
Hardware - 2 CPs
LPAR A - 2 shared CPs
LPAR B - 2 shared CPs
LPAR A is observed. CPM will consider adding new CP because the sum of logical
CPs (four) is greater than sum of physical CPs (two).
```

Figure 2-43 Shared processors

Figure 2-44 shows a configuration with dedicated and shared processors.

```
Hardware - 6 CPs, no subcapacity
LPAR A - 2 dedicated CPs
LPAR B - 2 shared CPs
LPAR C - 2 shared CPs
LPAR A is observed. CPM will not consider adding new CPs because it will not
help LPAR A.
LPAR B is observed. CPM will not consider adding new CPs because the sum of
logical CPs (four CPs in the shared pool) is equal the sum of physical CPs
(four CPs in shared pool).
```

Figure 2-44 Dedicated and shared processors

Figure 2-45 contains dedicated and shared processors on a subcapacity level.

```
Hardware - 4 CPs, subcapacity level
LPAR A - 2 dedicated CPs
LPAR B - 2 shared CPs
LPAR C - 2 shared CPs
LPAR A is observed. CPM will consider increasing capacity level, because it is
the only way to help a logical partition configured with dedicated processors.
LPAR B is observed. CPM will consider adding a new CP because the sum of
logical CPs (four CPs in a shared pool) is greater than the sum of physical CPs
(two CPs in shared pool). It will also consider changing the capacity level.
```

Figure 2-45 Dedicated and shared processors with subcapacity

- CPM adds specialty engines and CPs.

  If either CPs or specialty engines (zAAPs or zIIPs) can help a suffering workload (for example a Java™ workload when IFAHONORPRIORITY=YES), CPM considers adding specialty engines first.

- CPM takes the smaller step when determining to add CP or change capacity level.

  If CPM can add processors or change capacity levels of processors, it generally chooses the possibility that adds less MSUs (smaller steps). However, because of the many metrics that can be observed, the decision depends on a real workload.

- CPM deactivates resources to adjust to the definitions.

  Resources stay active for a minimum of four hours (the administrator can change this). Resources are deactivated one at a time, after the workload criteria are met. If the policy no longer allows resources to be active, CPM deactivates as many resources as necessary to adjust to the definitions. Resources are also deactivated one at a time.

  If capacity tokens are present in an On/Off CoD record, for each resource that CPM should provision, enough capacity tokens must be available for that resource to be active for one billing window (24 hours). If that criteria is not met, no resources from that record will be provisioned.

2.6.2 CPM definitions

All CPM definitions are done from the CPCC. Configurations and policies are defined by using the CPCC interface and are interrogated by CPM.

Note the following information:

- Observed systems

  These are z/OS images defined to CPM by the administrator. They must run z/OS V1R9 or later and they must be on z10 EC CPCs. If these requirements are not met, CPM issues warning messages.

- Managed CPC

  Only z10 ECs that contain valid On/Off Capacity on Demand offerings can be managed. If the administrator defines CPCs that do not meet these criteria, CPM issues warning messages. The CPM itself does not have to execute on an observed system but it can reside on any z/OS V1R9 system, even on a server that is not a System z10 server.

- Capacity Provisioning Domain Configuration

  This contains definitions for which z/OS images are to be observed and for which CPCs will be managed. CPM requires this information to be able to match the z/OS image location and to communicate with the correct CPCs when On/Off Capacity on Demand offering resources are needed. z/OS images can be located on many CPCs. They can be part of a sysplex, in which case the recommendation is to have all observed z/OS images and all managed CPCs from the sysplex defined in the same domain. CPM can contain multiple domain configurations, but only one domain configuration can be active at a time.

Capacity Provisioning Policy

The policy defines conditions under which additional capacity may be provisioned, by determining when, which, and how much, as follows:

- *When* provisioning is allowed is a time condition:
  - Start time - Indicates when provisioning can begin.
  - Deadline - Provisioning of additional capacity is no longer allowed after this time.
  - End time - Deactivation of additional capacity can begin at this time.
- *Which* work qualifies for provisioning is a workload condition:
  - The z/OS systems that may execute eligible work
  - Importance filter - Eligible service class periods are identified by WLM importance.
  - Performance Indicator (PI) criteria are:
    - Activation threshold

      PI of service class periods must exceed the activation threshold for a specified duration before the work is considered to be suffering.

    - Deactivation threshold

      PI of service class periods must fall below the deactivation threshold for a specified duration before the work is considered to no longer be suffering.
  - Included service classes are the eligible service class periods.
  - Excluded service classes are the service class periods that should not be considered.
- *How much* additional capacity may be activated, expressed in MSUs is a provisioning scope. The number of zAAPs and number of zIIPs are allowed one specification per CPC.

Definitions in the Provisioning Policy are organized hierarchically. See Figure 2-46 on page 74.

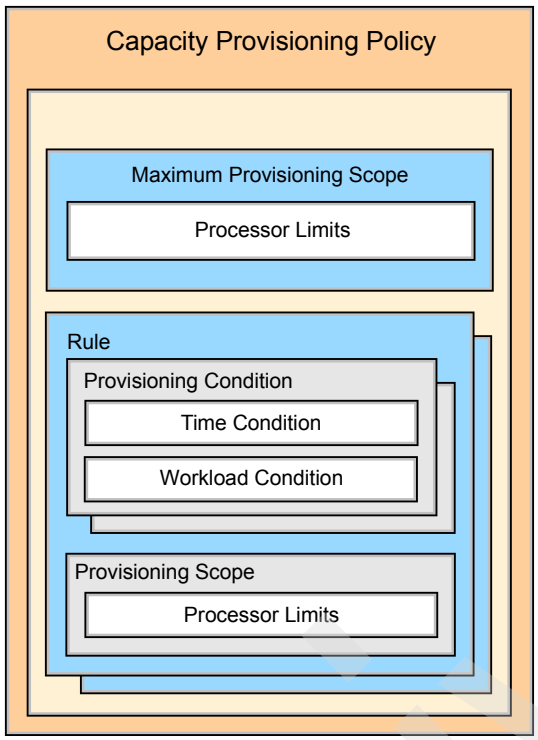

Figure 2-46 Capacity Provisioning Policy

The policy itself defines the maximum capacity that can be provisioned for each CPC in the domain. The maximum capacity is defined in terms of MSUs for CPs and in engine number for zAAPs and zIIPs.

The policy can contain one or more rules. Each rule has its own provisioning scope, which defines maximum capacity for each CPC provisioned by this rule. Each rule contains one or more conditions. The conditions determine when the rule is triggered. It can be triggered either by time condition or by workload condition.

A time condition defines time periods when a workload condition can be used, and is not triggered outside those time frames. If a time condition is specified without a workload condition, it is used for a scheduled capacity provisioning.

Workload conditions address WLM service classes. They specify a PI for a service class and a time period for how long a workload must be suffering before the condition is triggered. The time period provides protection against provisioning additional capacity caused by short spikes in a workload.

Note: At this time, CICS® and IMS transaction classes cannot be used to trigger provisioning actions.

CPM and capacity tokens

CPM, through the APIs, can interrogate the On/Off CoD records, including the information in the capacity tokens. Capacity tokens are always present in prepaid On/Off CoD records, and may be present in postpaid On/Off CoD records.

If CPM decides to activate resources from a record containing capacity tokens, a check is made for each resource in question to ensure that enough capacity tokens exists for that resource to be active for the next 24 hours (the billing window). If that is not the case, no

resources from that On/Off CoD record will be activated. An exception message is generated by CPM.

If CPM has successfully activated resources from an On/Off CoD record, and after a decrement at the end of a billing window results in any resource not having enough tokens for the next billing window, all the activated resources from that record are deactivated. An exception message is generated by CPM.

The system itself monitors capacity token consumption. If, by the decrement at the end of a billing window, any resource is approaching total consumption, the system issues exception messages. Those messages start five days before the total consumption of any resource, based on the current usage level. An important recommendation is that you implement procedures to handle such a situation. You may select to let the deactivation take place, then have CPM reactivate parts of the record that has remaining capacity tokens, or to replenish the capacity token in question by using Resource Link, to avoid the deactivation of the entire record.

2.6.3 Capacity Provisioning infrastructure

Several components make up the infrastructure for Capacity Provisioning. CPCC on a workstation is used to define domain configurations and policies to CPM. To connect from the CPCC to the CPM, a Common Information Model (CIM) server is required on the system where the Provisioning Manager runs.

The CIM server is provided as a function in the z/OS BCP. The CIM is the result of an open source project in which IBM has taken a lot of interest, known as the Distributed Management Task Force (DMTF), also known by its pet name, Pegasus.

CPM retrieves WLM statistics from observed systems through the RMF distributed data server (RMF DDS), or an equivalent product. The connection from CPM to RMF DDS is done by using the CIM server.

If the observed z/OS images are part of a sysplex, RMF DDS can run on any image in the sysplex. CIM servers have to know the address of the RMF DDS. A recommendation is to configure the RMF DDS in high availability mode with the same virtual IP address, regardless of the z/OS image on which it runs. See 7.6, “CPM hints and tips” on page 203 for details.

When CPM detects the need to provision additional capacity (either because of a suffering workload or as a scheduled activity), it connects to the HMC that controls the managed CPC. Connection is made using SNMP. The HMC API is used to activate the On/Off Capacity on Demand offering.

Figure 2-47 on page 76 shows the whole infrastructure required for z/OS Capacity Provisioning.

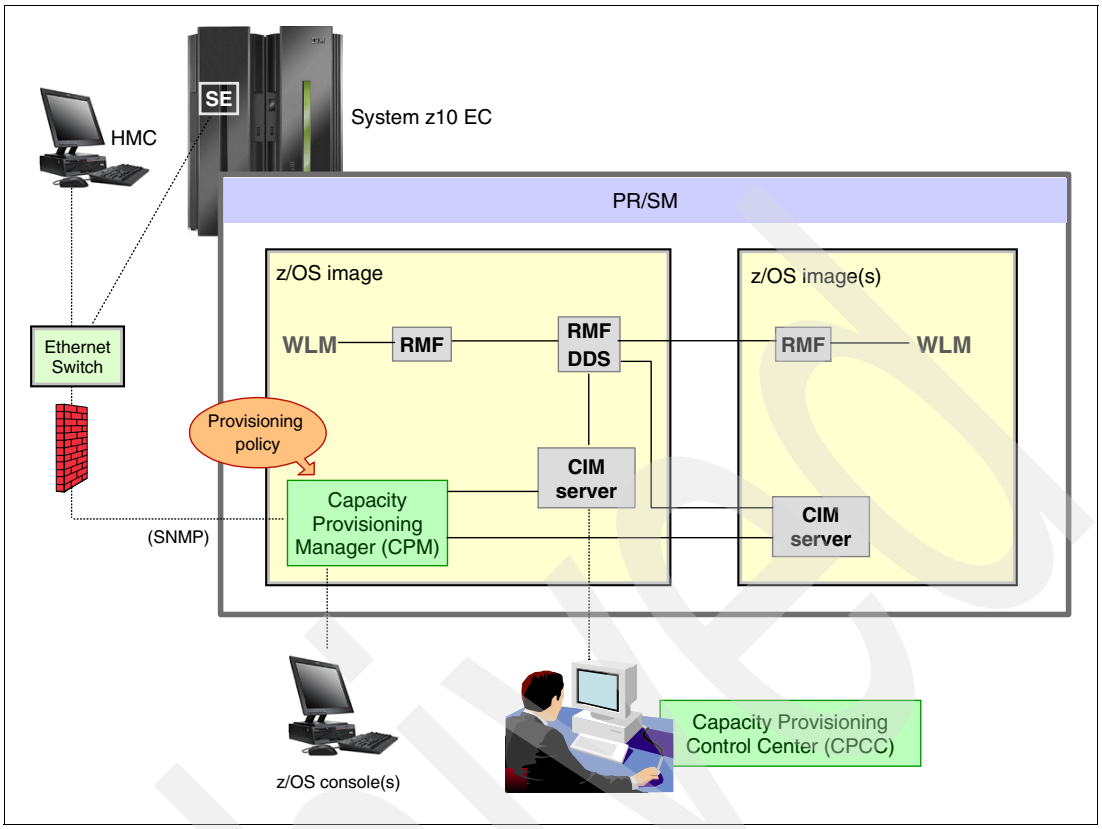

Figure 2-47 z/OS Capacity Provisioning infrastructure

2.6.4 Capacity Provisioning planning

Before planning the details and implementing the solution, you should understand the technology behind the offering. The basics are provided in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

Consider the following aspects when you plan for z/OS Capacity Provisioning:

- A thorough understanding of the current conditions is mandatory. Understand workload variations and needs, Service Level Agreements (SLAs), WLM service class definitions, z/OS environment and configuration, and sysplex structure
- Validate requirements and expectations. For example, determine the answers to the following questions:
  - What type of provisioning is needed and is it scheduled or workload-based?
  - Should autonomic mode be the ultimate goal?
  - Can CPM help in the current configuration (dedicated versus shared processors, logical versus physical processors)?
- Consider financial and legal consequences, such as additional software and hardware charges and licenses, also for ISV products.
- Determine the CPM mode. An easy starting place is with a CPM in analysis mode. This mode does not require any On/Off CoD record to be installed on the server, and it does not perform any action. CPM in this mode only displays proposals on the console. The next step would be CPM in confirmation mode, which requires a record to be installed. CPM in a confirmation mode proposes actions and can perform actions following operator

confirmation. When all definitions are tuned to satisfaction, CPM can be switched to an autonomic mode.

- If capacity tokens are used in an offering record, ensure that enough tokens are always available to cover at least one billing window (24 hours) for each of the resources in the record. If that criteria is not met, no resources from that record can be provisioned by CPM.

A number of software components are involved in z/OS Capacity Provisioning, as indicated in Figure 2-48.

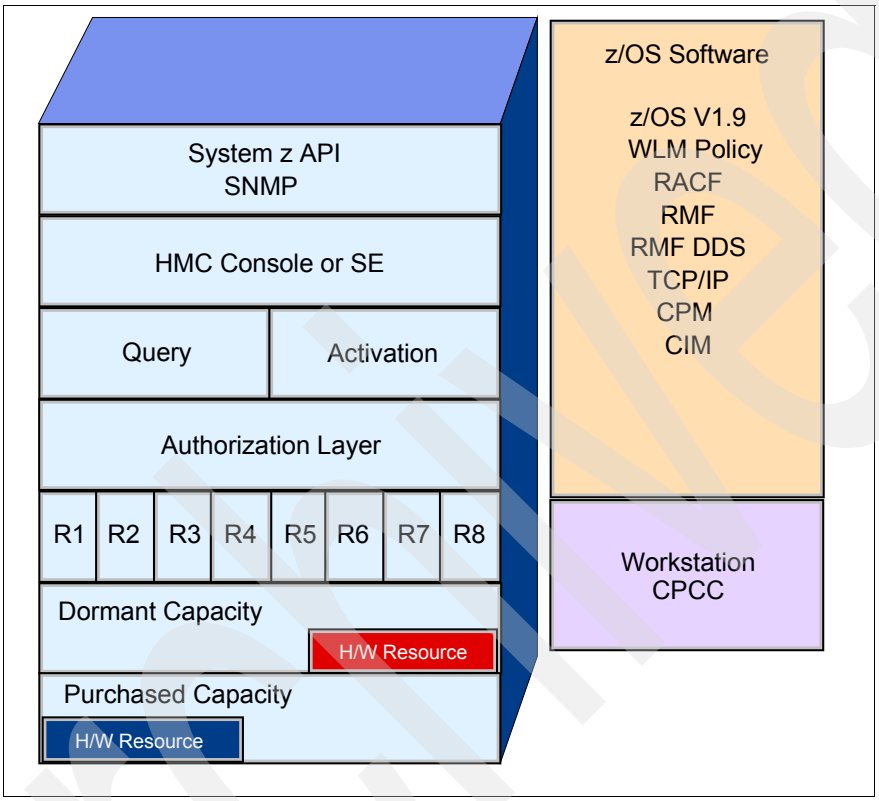

Figure 2-48 z/OS Capacity Provisioning infrastructure

The components are:

- CPCC

  This is used only to define policies and transfer them to CPM. CPCC is not required for CPM operation. The CPCC code can be downloaded from z/OS V1R9.

- CIM client for CPCC

  This can be downloaded from z/OS V1R9. Its location is specified during CPCC installation.

- Provisioning policies

  A provisioning policy, or policies, is defined by using the CPCC.

- CPM

  This is the controller of what will be provisioned, and when. It runs only on z/OS V1R9 and later.

- WLM

  This must be running in a goal mode. To be able to use workload conditions to trigger provisioning action, appropriate service classes must be defined.

- RMF

  This must be enabled to provide performance metrics to CPM. A product equivalent to RMF (including equivalent CIM RMF provider capability) can be used.

- RMF DDS

  This provides RMF metrics to CPM.

- CIM server

  This manages communication between CPCC and CPM and between CPM and RMF DDS.

- z/OS security server

  This has to support creation of pass tickets (R_GenSec) and evaluation through the SAF interfaces.

- Java 5 SDK

  This is the Java runtime for CPM.

- Naming conventions and standards

  A strong recommendation is to establish naming conventions and standards in the planning phase. *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299 contains naming suggestions for CPM. We highly recommend that those suggestions be used.

- Software levels

  All software levels required by the CPM infrastructure are defined in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

- TCP/IP

  All communication among CPM, CPCC, CIM and RMF DDS is based on TCP/IP, and SNMP connectivity from CPM to the Hardware Management Console or Support Elements.

Although many ways exist to start the z/OS Capacity Provisioning implementation, here are the suggested steps:

1. Verify the prerequisites.
2. Define the objectives.
   - Define managed CPCs and z/OS images.
   - Define provisioning and de-provisioning conditions, such as:
     - Which service classes
     - Scheduled provisioning
     - Workload-based provisioning
3. Set up prerequisite products.
4. Install recommended code levels.

   Consult *IBM System z10 Capacity on Demand User's Guide,* SC28-6871.
5. Install CPCC.
6. Install CPM.

7. Perform customization steps. Note the following guidelines:
   - A strong recommendation is to use the default naming conventions.
   - Develop a policy in CPCC. Start with a simple policy.
   - Start with CPM analysis mode.
8. Check for possible constraints where Capacity Provisioning cannot help, such as:
   - WLM definitions
   - Networking delays
   - Dedicated processors
9. Order required On/Off CoD records.
10. Test and re-evaluate, as follows:
    a. Verify and eventually modify the WLM parameters.
    b. Verify and eventually modify CPM rules. Check CPM parameters regarding utilization percentage.
11. If you use capacity tokens, a strong recommendation is to:
    - Closely monitor the capacity tokens in activated On/Off CoD records to avoid pre-mature deactivation of the resources activated from those records.
    - Set up procedures to handle situation where premature deactivation takes place, because not enough capacity tokens remain after the decrement at the end of a billing window.

3

Capacity Backup

Capacity Backup (CBU) provides replacement capacity for any type of processor in the event of a disaster or other unplanned outage. CBU is intended for use for disaster recovery purposes only and cannot be used to cover peaks in production workloads. Details about CBU can be found in 2.1.2, “Capacity Backup” on page 20.

This chapter discusses the Capacity Backup (CBU) offering usage, and the tasks to consider and plan.

This chapter discusses the following topics:

3.1 Preparation

To recover from a disaster, a CBU feature can be activated for up to 90 days. When a disaster strikes, making correct decisions could be too late. Therefore, strong recommendations include:

- Do all planning and preparations in advance.
- Test the CBU function before any real disaster.

3.1.1 Determine logical partition requirements

To take advantage of the resources made available by CBU, the server must be prepared accordingly. Preparation should start with analyzing how much replacement capacity is required to support a given workload and what types of PUs are necessary. Depending on your configuration, two scenarios (or a combination of them) are available for using the extra resources:

- The lost capacity can be recovered by providing additional resources to already active logical partitions.
- Additional resources are used by a new logical partition, which is activated after resources become available.

System z10 provides a fixed HSA, where the maximum number of logical partitions and I/O devices is already defined. All changes can be dynamic and they do not require server POR. A lot of planning can be saved if a new partition is needed during CBU activation.

CBU can only add processor resources. Additional memory, I/O devices, network connectivity and so forth must be determined and planned ahead of time. Required I/O, networking, and storage resources must be available at CBU activation time so that it is possible to take advantage of the extra resources enabled by CBU.

When resources are added to a server, the relative capacity of all partitions can change. The logical partition weights should be reviewed to insure that the needed capacity would be delivered to each logical partition. Refer to *IBM System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide,* SB10-7153.

The time spent planning is time well spent. Performing required tasks ahead of time helps avoid problems, stress, and mistakes when CBU is used in a disaster.

Prepare to activate an additional logical partition

If an extra logical partition has to be activated during a disaster situation, it must be activated as fast as possible. Planning ahead helps to avoid mistakes under stress. Plans should be checked and tested regularly. Tests show whether a plan works and helps you to get more experience in the process.

Each logical partition must have processor and memory resources defined in the image activation profile. The number of processors (CP, IFL, ICF, zAAP, zIIP, and SAP) and their operation mode (dedicated or shared) must be provided in this profile. To be able to add additional processors dynamically to this logical partition, while it is already active, define *reserved* processors or use the function Logical Processor Add if supported by the operating system. If processors are added to a logical partition, z/VM V5R3 recognizes them nondisruptively without having them previously defined as reserved. z/OS requires that these resources be defined as reserved.

Figure 3-1 shows the definition panel for processors in the image activation profile. It defines shared processors to a partition, ten CPs, out of which six are reserved. It also defines four zAAPs and four zIIPs; in each case two processors are reserved.

In addition to processor and memory definitions, calculating how logical partition weights should be set after temporary resources are added is also important.

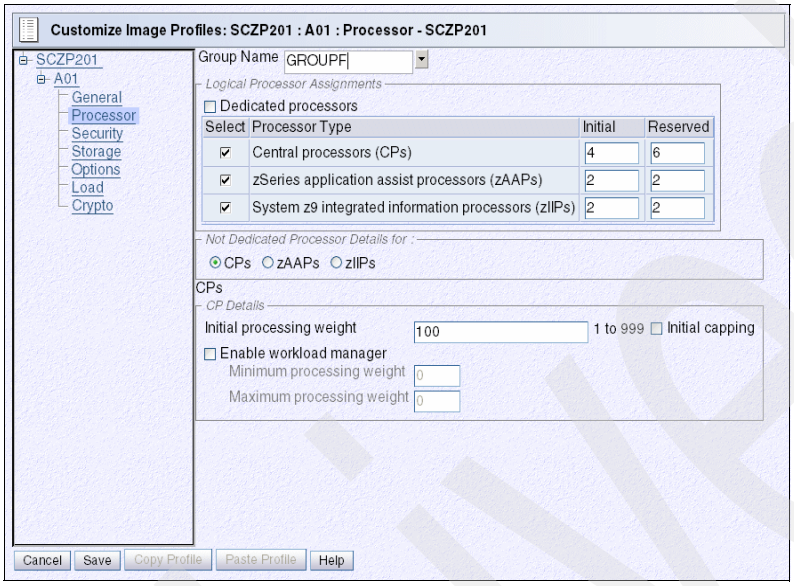

Figure 3-1 Image Activation profile - processor tab

Prepare to increase the resources for already active logical partitions

If all temporary resources will be given to existing partitions, the additional resources have to be added nondisruptively.

The extra resources have to be already defined in the image activation profiles as reserved processors.

You must determine how logical partition weights have to be adjusted. Temporary resources influence the whole environment and if weights do not reflect the change, serious performance implications can result. The same analysis is necessary for the logical-to-physical processor ratio for shared processors. Keep in mind that the System z10 allows you to activate a subset of the temporary record if necessary.

Note: If processors are added to a logical partition, z/VM V5R3 recognizes this nondisruptively without having them previously defined as reserved. z/OS requires that these resources be defined as reserved.

3.1.2 Order, retrieve, and install a record

This section contains a brief description of steps performed when a CBU record is ordered from Resource Link. You may also order a record from your IBM Sales Representative.

Order a record

Records can be ordered through Resource Link. Figure 3-2 on page 84 shows the Order Capacity Backup record panel on Resource Link, where the desired model capacity, the number of specialty engines, the entitlement length, and number of tests have to be defined.

In the figure, the order here contains nine CP CBU features, and four zIIP and four zAAP CBU features. The order is valid for five years; it contains five tests by default.

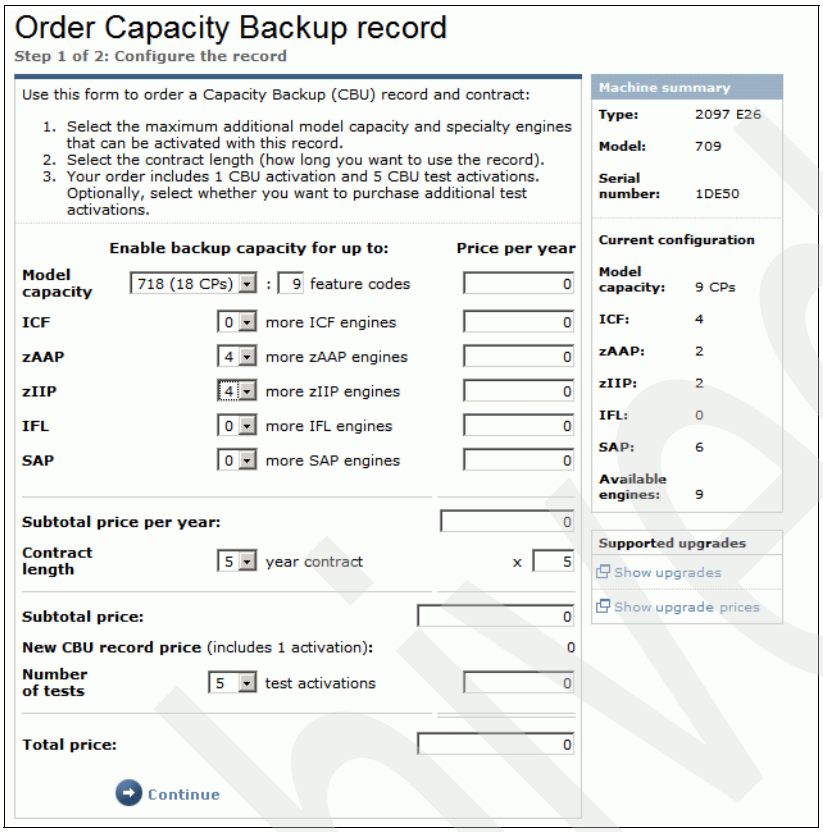

Figure 3-2 Order Capacity Backup record

When an order is submitted, it must be approved by an approver if one is defined. After the order is approved and is ready to be retrieved, a notification e-mail states that the record is ready for retrieval.

Example 3-1 shows an excerpt from a notification e-mail.

Example 3-1 Notification e-mail

```
Subject: Download ready: CBU order CX7KHB38 for machine (2097 -- 1DE50) is ready
to be downloaded and installed


Customer number: ITSO001

CBU order CX7KHB38 for machine 2097 -- 1DE50 has been processed and is
ready to be downloaded to the customer's machine. Please contact the
appropriate person to retrieve this order from the Hardware Management
Console (HMC).


Order Create Date: 10/17/2008
Record number: CB7KHB38
Order number: CX7KHB38

Order Confirmation Information
```

```
Expiration date:  10/17/2013

Model capacity:   718 (18 CPs) : 9 feature codes          $xxxxx.xx
ICF:              0 more ICF engines                     $xxxxx.xx
zAAP:             4 more zAAP engines                    $xxxxx.xx
zIIP:             4 more zIIP engines                    $xxxxx.xx
IFL:              0 more IFL engines                     $xxxxx.xx
SAP:              0 more SAP engines                     $xxxxx.xx
Subtotal price per year:                                 $xxxxx.xx

Contract length:        5 year contract                  x 5
Subtotal price:                                          $xxxxx.xx

CBU record price:                                        $xxxxx.xx
Number of tests:        5                                $xxxxx.xx

Total price:                                             $xxxxx.xx

Description:      +9 FCs model capacity, +0 ICF, +4 zAAP, +4 zIIP, +0 IFL,
+0 SAP
```

Retrieve a record

When Resource Link has sent the notification, the record is ready to be retrieved and the status in Resource Link is changed to “Download ready.”

To retrieve a record, a connection to the IBM Remote Support System is needed. This requires the Remote Support Facility (RSF) correctly defined and enabled. RSF connection and setup information is described in *Hardware Management Console Operations Guide Version 2.10.0,* SC28-6867.

Note: For situations where no RSF access is allowed, a request for price quotation (RPQ 8P2305) must be used to retrieve the Capacity on Demand records from portable media.

Figure 3-3 shows a Perform Model Conversion panel, which is used for record retrieval.

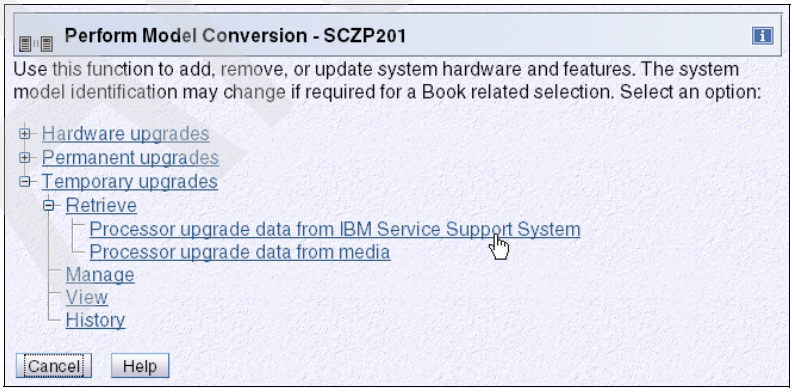

Figure 3-3 Perform Model Conversion panel

All temporary upgrade records are retrieved and staged on the Support Element (SE). The installation is done after, in a separate step.

Note: Any records ordered with the server and delivered from manufacturing will be staged on the Support Element, and they must be installed before they can be activated.

Install a record

When the record is retrieved, it is placed in the staged records area on the SE. The records in the staged area can be viewed, installed, or deleted. Figure 3-4 shows four staged records, one of which is CBU.

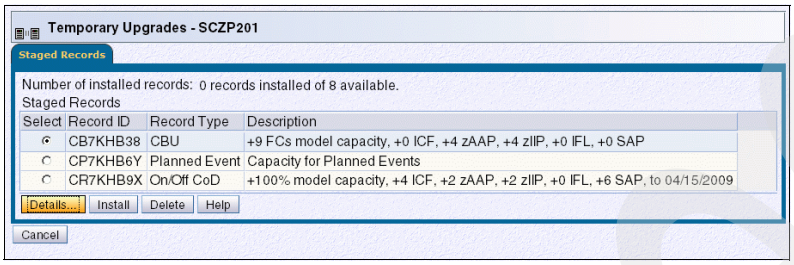

Figure 3-4 Temporary Upgrades - Staged Records

Figure 3-5 shows CBU record details. The panel provides the record ID, order number, the number of activations, and the duration information.

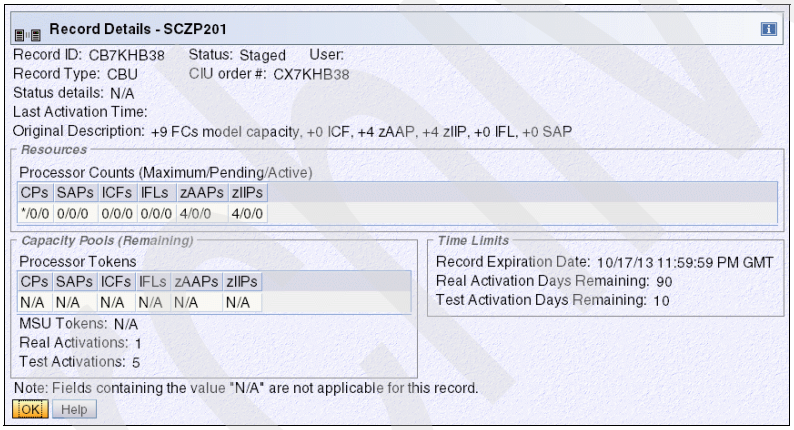

Figure 3-5 Record Details

To move the record into one of eight slots available for records that can be activated, it must be installed first by clicking **Install**, as shown in Figure 3-6.

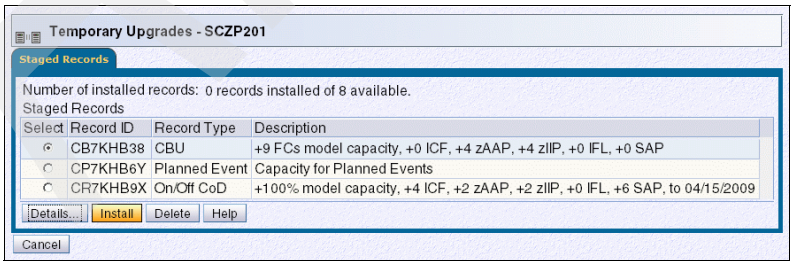

Figure 3-6 Temporary Upgrades, Staged Records

Moving an installed record back to a staged record is not possible, so one final confirmation is necessary. See “Temporary Upgrades - message ACT37466 after CBU record install” on page 87.

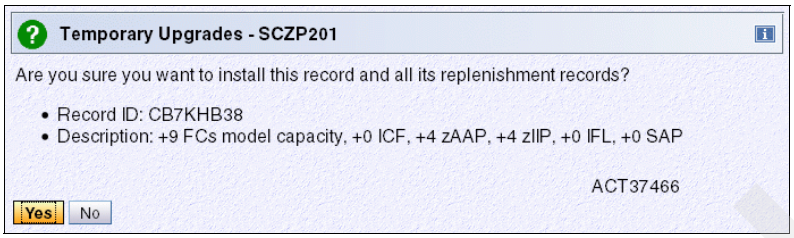

Figure 3-7 Temporary Upgrades - message ACT37466 after CBU record install

Figure 3-8 shows the installed records area. The CBU record is installed and allows activation of up to four zAAPs and up to four zIIPs. The number of CPs that can be activated is determined at the activation time. In this case, it would be up to nine additional CPs because no capacity level changes will be allowed.

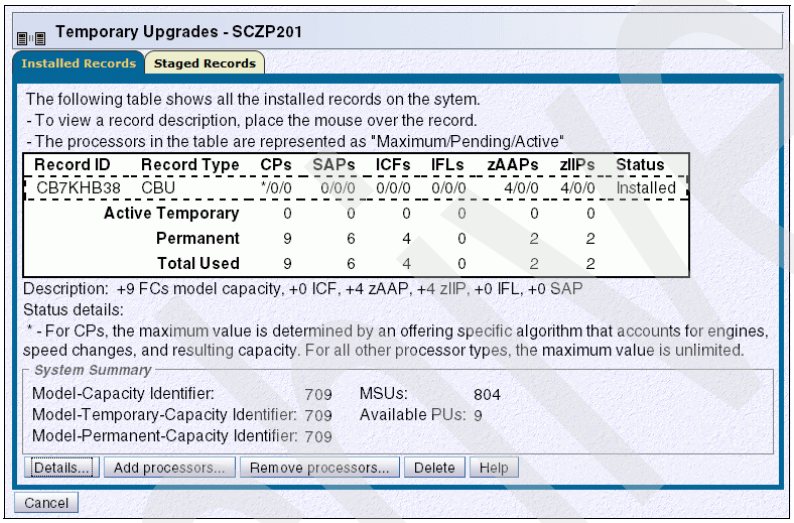

Figure 3-8 Temporary Upgrades

Additional details about ordering, retrieving, and installing the CBU record is in *IBM System z10 Capacity on Demand User's Guide,* SC28-6871.

3.2 Activation

After the CBU record is retrieved, external connectivity is unnecessary. An installed record can be activated at any time without access to IBM. For this reason, we strongly recommend to have the CBU records already retrieved to the SE when necessary to ensure the availability of the record independently of any modem and Internet connection in a disaster. Approximately 200 records can be placed in a staged-records area on the SE but only eight Capacity on Demand records can be installed at the same time. A record cannot be moved back from installed to staged. In other words, if more than eight Capacity on Demand records are available and because changing the status from an *installed record* to a *staged record* is not possible, ensure to size a CBU record properly. If it is installed, and more resources are necessary, a record can be replenished to contain additional resources.

Only installed CBU records can be activated through the APIs. Therefore, if you are relying on GPDS to activate CBU records, a recommendation is to have a CBU record installed at all times.

Figure 3-9 shows an output from the z/OS D M=CPU command. It shows nine online CPs, two online zAAPs, and two online zIIPs.

```
D M=CPU
IEE174I 11.15.34 DISPLAY M 002
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 3-9 D M=CPU command output before CBU activation

Figure 3-10 shows an RMF report from the same z/OS. It also shows nine CPs online for logical partition A01 and a server capacity setting of 709.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:   11   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  254   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                          33.0                      0.4    45.0   45.3
A0E            0   178  NO   2.0      99.8    99.8    0.0    22.2   22.2
A0F            0     1  NO   2.0       0.5     0.5    0.0     0.1    0.1
A01            0     2  NO   9.0       0.2     0.2    0.0     0.2    0.2
A02            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A1E            0   178  NO   2.0      99.7    99.7    0.0    22.1   22.1
A1F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A11            0     1  NO   4.0       0.5     0.5    0.0     0.2    0.2
A12            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
PHYSICAL                                              0.3            0.3
```

Figure 3-10 RMF report before CBU activation

Activation process

The Temporary Upgrades panel is used to activate a CBU record that is already installed. Figure 3-11 shows that a CBU record is selected and that **Add processors** has been clicked.

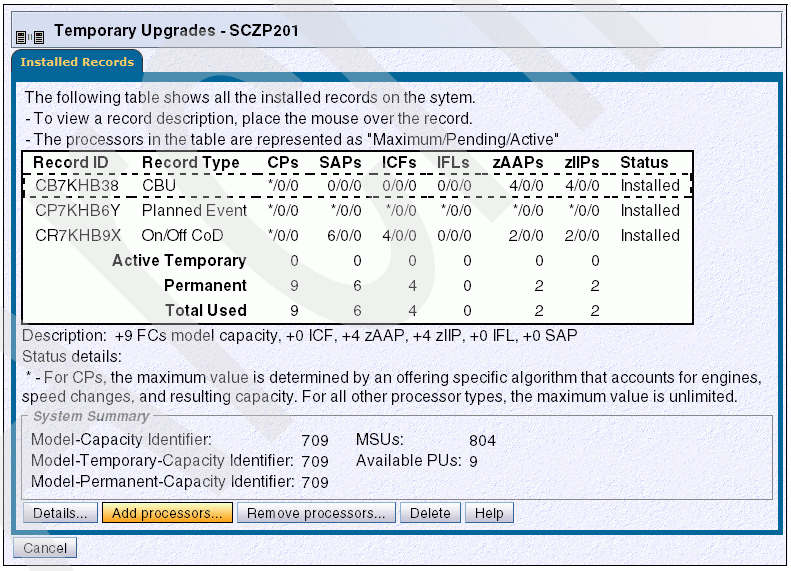

Figure 3-11 Temporary upgrades

Figure 3-12 shows the Change Activation Levels panel, which is where you define how many resources should be activated. In this case, five CPs, two zAAPs, and one zIIP will be added. The **Force activation** button is disabled because this record is not the default CBU record.

Note: Make sure that the correct Activation Option button (Test Activation or Real Activation) is selected in order to activate the desired CBU option.

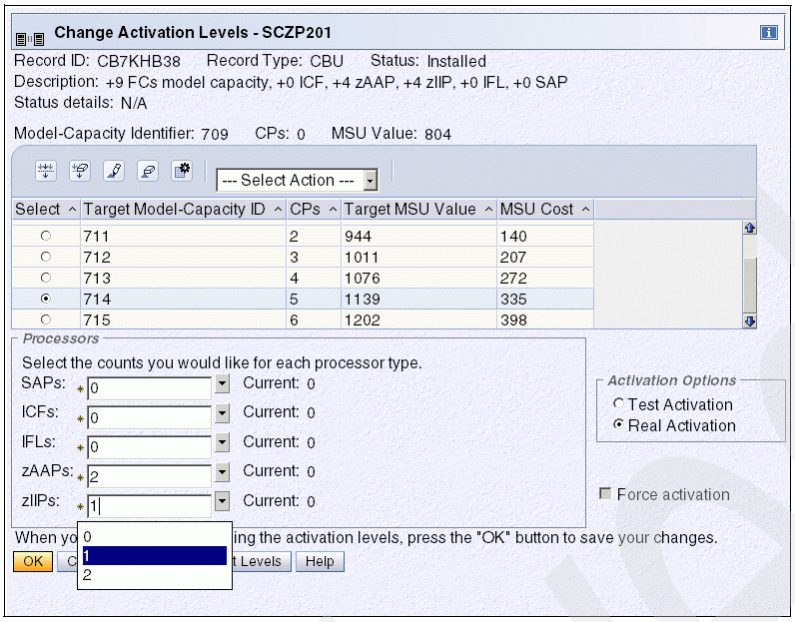

Figure 3-12 Change Activation Levels

After you click **OK**, a message is displayed. Figure 3-13 shows message ACT37464. It shows the original and the new configuration. After you click **Yes**, the CBU record is activated.

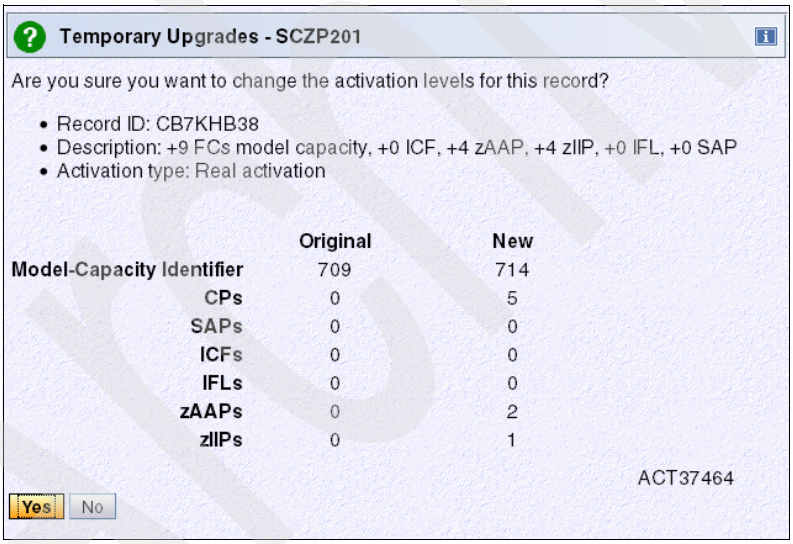

Figure 3-13 Temporary Upgrades, message ACT37464

The activation level can also be changed after the CBU is already activated.

After a record is active, the Temporary Upgrades panel reflects the following hardware changes, shown in Figure 3-14:

- 14 CPs active, 9 are permanent and 5 are temporary
- 4 zAAPs, 2 are permanent and 2 are temporary
- 3 zIIPs, 2 are permanent and 1 is temporary

The figure also reflects the new MSU value and capacity identifiers. One PU is still available. Because no billable capacity is active, Model-Temporary-Capacity Identifier has not changed.

Current capacity is represented by Model-Capacity Identifier, which has changed to 714. Five temporary CPs were added to a server, which has Model-Permanent-Capacity Identifier of 709.

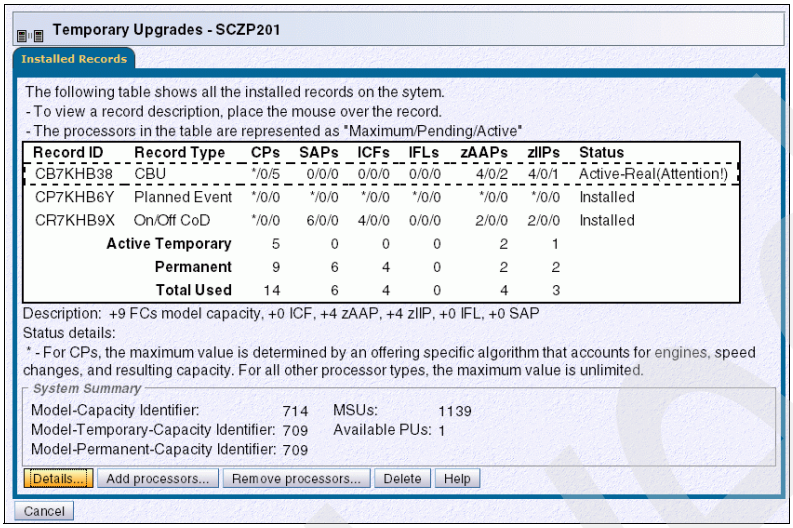

Figure 3-14 Temporary Upgrades - server configuration after CBU activation

The Record Details panel provides more information (such as number of activations, expiration dates, and so on). See Figure 3-15. In this example, because the real activation was just started, the number of real activations is already set to zero. The status shows that there are no activations remaining.

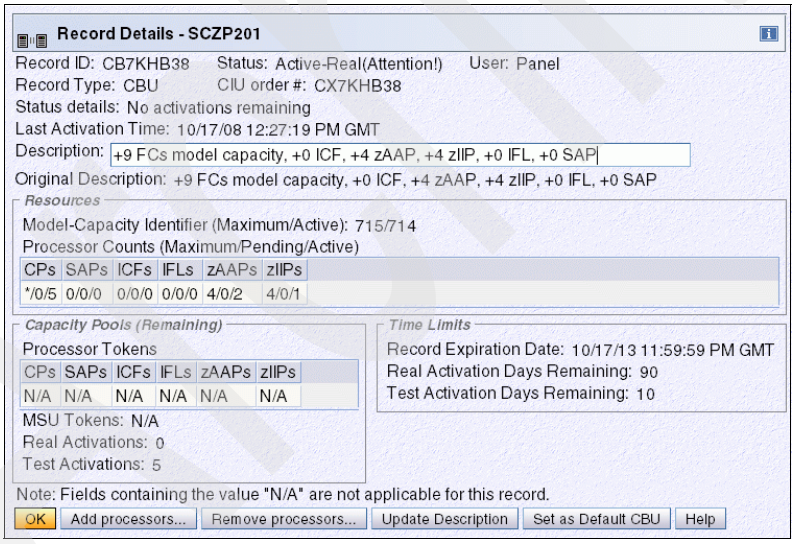

Figure 3-15 Record Details

Figure 3-16 shows the output of the D M=CPU command after activation. Four additional processors (one CP, two zAAPs, and one zIIP) which were not available before, are available now. Although CBU added five temporary CPs, the output shows only one additional CP known by z/OS. The reason is that this z/OS has ten CPs defined in the image profile, and nine of them had been online before the CBU activation. The figure also shows the new Model-Capacity Identifier 714.

```
D M=CPU
IEE174I 12.06.48 DISPLAY M 139
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  -
0E  -A
0F  -A
10  -I
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.714.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 3-16 D M=CPU output after CBU activation

The RMF output in Figure 3-17 also shows that the Model-Capacity Identifier and the CPC Capacity value have changed. Also, Image Capacity was recalculated as explained in 2.3.6, “z/OS Resource Measurement Facility” on page 57.

```
 Partition:   A01        2097 Model 714
 CPC Capacity:   1139   Weight % of Max: ****   4h Avg:   11   Group:   N/A
 Image Capacity:  814   WLM Capping %:    0.0   4h Max:  254   Limit:   N/A
```

Figure 3-17 RMF data after CBU activation

To adjust the performance, the weights of the logical partitions should be modified by using the Change Logical Partition Controls task.

Adding resources dynamically to an active LPAR

This section discusses the use of extra processors to dynamically increase the resources of an already active logical partition.

Note: If the CBU only increased the capacity level of the already-active processors (that is, Model-Capacity Identifier from 408 to 708), without increasing the number of processors, no further actions are required on the operating system.

If the number of logical processors (shared or dedicated) in a logical partition has to be increased, the processors must be configured online by using the appropriate operating system command. An extra processor can be used as either shared or dedicated, as follows:

- If this processor is configured online to a logical partition as dedicated, it is no longer available in the shared pool and therefore not usable by any other logical partition (either as shared or dedicated). The whole processor is used only by one partition.
- If this processor is configured online as shared, it is still available in a shared pool in order to be used as shared logical processor for other logical partitions. The benefit for each logical partition depends on the number of defined shared processors and partition weights. If required, review and adjust the logical partition weights accordingly to achieve the projected objectives.

Figure 3-18 on page 93 shows the z/OS view of its logical processor after upgrade and putting additional processors online. In this example, the shared processors 0D, 0E, 0F, and 10 have been configured online to logical partitions to drive the additional workload.

```
IEE174I 12.46.55 DISPLAY M 214
PROCESSOR STATUS
ID  CPU                SERIAL
00  +                   01DE502097
01  +                   01DE502097
02  +                   01DE502097
03  +                   01DE502097
04  +A                  01DE502097
05  +A                  01DE502097
06  +I                  01DE502097
07  +I                  01DE502097
08  +                   01DE502097
09  +                   01DE502097
0A  +                   01DE502097
0B  +                   01DE502097
0C  +                   01DE502097
0D  +                   01DE502097
0E  +A                  01DE502097
0F  +A                  01DE502097
10  +I                  01DE502097
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.714.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 3-18 PUs 0D, 0E, 0F,10 online

Change the activation level

Capacity on Demand design allows to change CBU activation levels dynamically. The Change Activation Levels task offers the possibility to add and remove processors. Adding and removing processors might make it necessary to configure processors offline or online in logical partitions, as described above.

Note: Always check the resulting MSU value. Decreasing the CPC MSU value by adding CPs when the capacity level is decreased in the authorization space is possible. Also, removing CPs in the authorization space can increase the CPC MSU value when the capacity level is increased.

Activating an additional LP after dynamic change of resources

We now describe a case when a new logical partition that requires four dedicated CPs has to be activated. Figure 3-19 on page 94 shows the processor definition tab of the related image profile. It needs four physical CPs to drive the four required dedicated CPs.

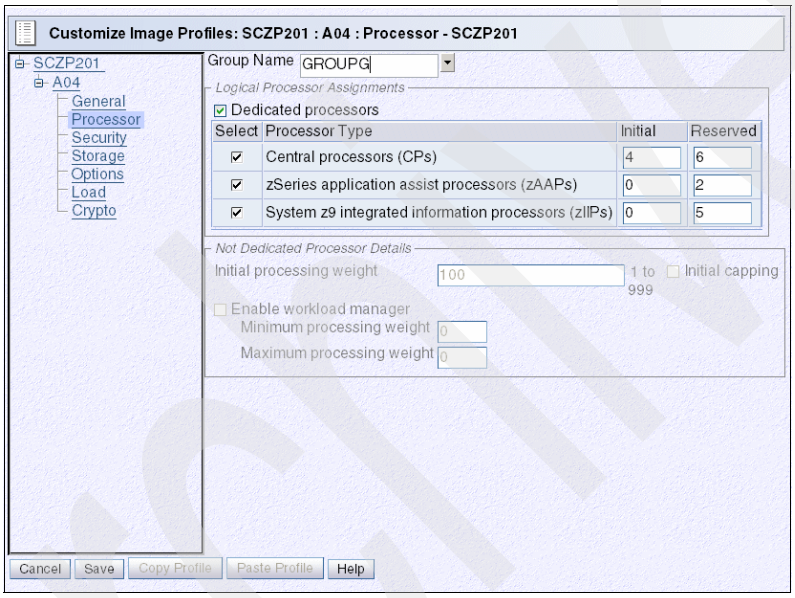

Figure 3-19 A04 Image Profile, Processor view

Because the current server does not have four unused CPs, they are provided as follows:

- One additional PU is still available on the server, as displayed in Figure 3-14 on page 91.
- Three temporary processors (CP 0D, zAAP 0F, and zIIP 10) will be configured offline from partition A01, as shown in Figure 3-20 on page 95.

```
CF CPU(D,F,10),OFFLINE
IEE505I CPU(D),OFFLINE
IEE505I CPU(F),OFFLINE
IEE505I CPU(10),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE

D M=CPU
IEE174I 14.51.51 DISPLAY M 291
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  -
0E  +A                   01DE502097
0F  -A
10  -I
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.714.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-
N NOT AVAILABLE
```

Figure 3-20 Configure processors offline

After the PUs have been configured offline from the logical partition, the CBU activation levels can be changed. Changing the PU type is required in our example. To change the PU type:

1. From a CBU record, deactivate the PU that is not needed.
2. Activate the PU of the desired type.

Note: Be careful not to remove the last processor from the CBU record. If you remove it, the record will be deactivated.

Figure 3-21 on page 96 shows removing one zIIP and one zAAP from activated resources.

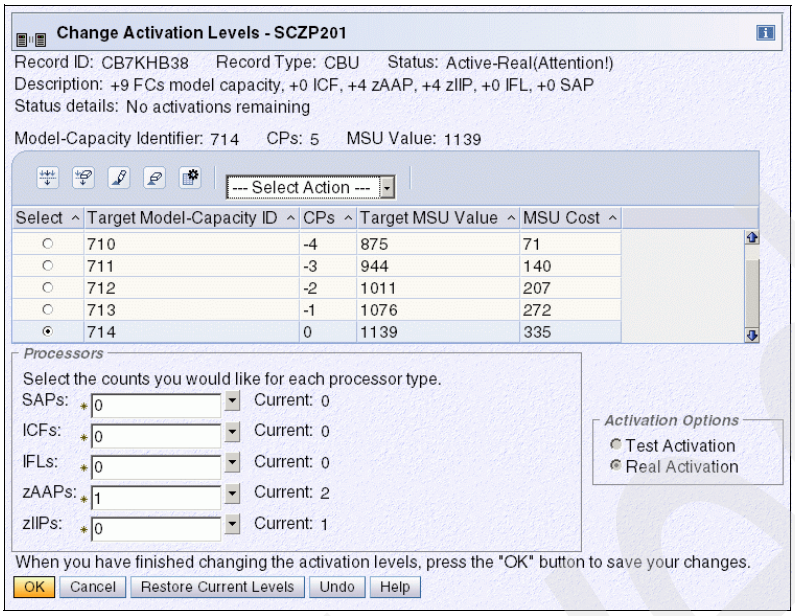

Figure 3-21 Change Activation Levels - removing one zAAP and one zIIP

After the zIIP and zAAP are deactivated, there are three available PUs on the server. Figure 3-22 shows adding three temporary CPs to the configuration.

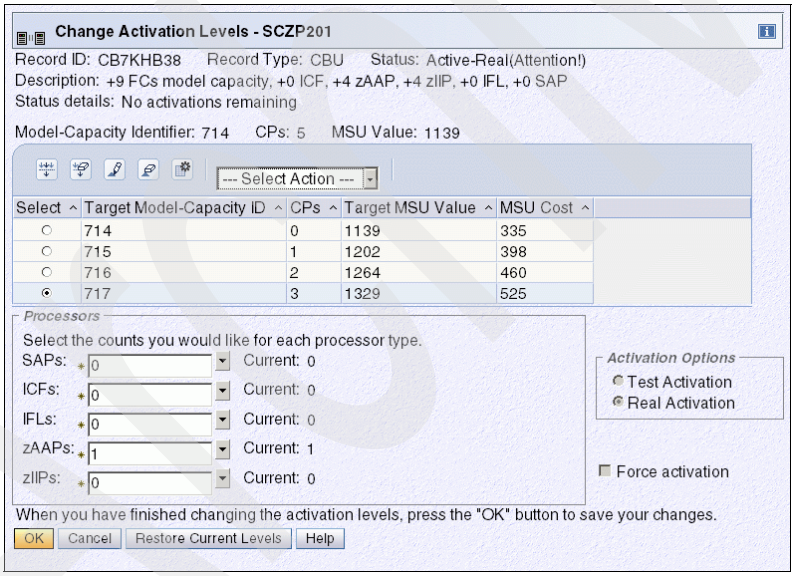

Figure 3-22 Change Activation Levels - add three CPs dynamically

After the dynamic change, review the new CoD configuration to ensure that the required resources are available.

Figure 3-23 shows the eight CPs activated by the CBU, the new Model-Capacity Identifier 717, and the related MSU value. One temporary zAAP is still activated by the CBU record.

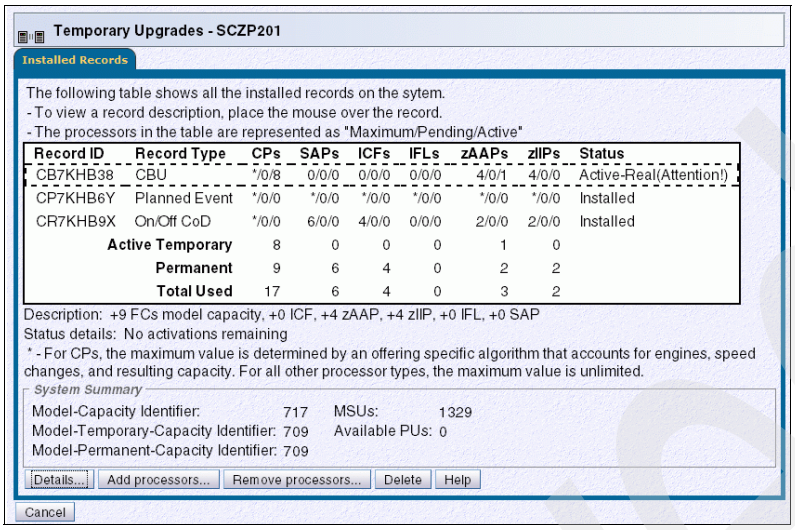

Figure 3-23 Temporary Upgrades - server configuration after the dynamic change

After the availability of the processors is verified, activate the logical partitions. When this has been done, the resource assignment to the logical partitions can be verified accordingly.

At this time, the RMF report also displays new logical partition A04 and the four associated CPs. See Figure 3-24. It also shows the new CPC Capacity value and the recalculated Image Capacity value.

```
Partition:   A01       2097 Model 717
CPC Capacity:   1329   Weight % of Max: ****   4h Avg:   97   Group:   N/A
Image Capacity:  782   WLM Capping %:   0.0   4h Max:  254   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        37.0                       0.1    35.6   35.7
A0E            0   156  NO   2.0       100     100    0.0    11.8   11.8
A0F            0     1  NO   2.0       0.4     0.4    0.0     0.0    0.0
A01            0   158  NO   9.0      22.4    22.4    0.0    11.8   11.8
A02            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A04            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A1E            0   156  NO   2.0      99.8    99.8    0.0    11.7   11.7
A1F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A11            0     1  NO   4.0       0.5     0.5    0.0     0.1    0.1
A12            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1
```

Figure 3-24 RMF output with A04 activated

3.3 Deactivation

Preparing the system to deactivate the CBU depends on the steps performed in 3.2, “Activation” on page 87.

Preparing the deactivation

To avoid future problems, revisit all definitions and activations that were made after CBU record activation. Because there is more flexibility for a user in Capacity on Demand design, more responsibility is required. Consider and follow certain rules, when appropriate:

- If CBU resources were used as dedicated processors to activate an additional logical partition, that partition should be deactivated. If it uses dedicated processors, deactivating CBU takes processors of the same type from the shared pool to satisfy the demand for dedicated processors while at least one processor is available in a shared pool. This can have severe performance implications for partitions that use shared processors.

  In other words, while at least one physical processor is available in a shared pool, regardless of any performance issue, System z10 allows the deactivation of temporary resources.

  If not enough resources are available to satisfy the request for all dedicated CPs and at least one CP for the shared pool, the server keeps the additional CoD resources active and posts an error message, as shown in Figure 3-25.

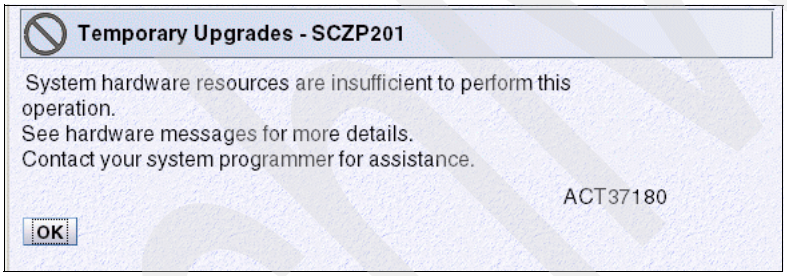

Figure 3-25 Temporary upgrades - message ACT37180

- If any dedicated processor has been configured online by the operating system, this processor should be configured offline in order to deactivate the CBU record. If the processor is not configured offline and the CBU record is deactivated, the dedicated processor will be used from the pool of shared processors, which changes the logical-to-physical processor ratio and can cause performance problems.
- If any shared processor was configured online, it should be configured offline in order to ensure the same logical-to-physical ratio that existed before the activation of the CBU record. As possibility is to have fewer physical processors available in the shared pool as logical shared processors assigned to a single logical partition. The server will allow the deactivation of CBU resources while at least one physical processor is available in any used pool.

Although in certain cases, you may ignore these rules, you should understand all implications.

Deactivation process

If the previously mentioned prerequisites are met, the CBU record can now be deactivated. After you click Remove Processors on the Temporary Upgrades panel, the Change Activation Levels panel is displayed. See Figure 3-26 on page 99.

The **Undo** button in Figure 3-26 does not deactivate the CBU at once. It updates all fields in the panels to their original values before record activation. Although you may also manually choose the original values, clicking **Undo** is more convenient. After the values are set and you click **OK**, the confirmation window opens and then deactivation is performed.

The Perform Model Conversion task on the Support Element does not show an activate or deactivate function. You control temporary upgrades by selecting **Add processors** or **Remove processors**. A temporary capacity record is active when **Add processors** results in a positive count, or inactive when **Remove processors** has brought the count back to zero.

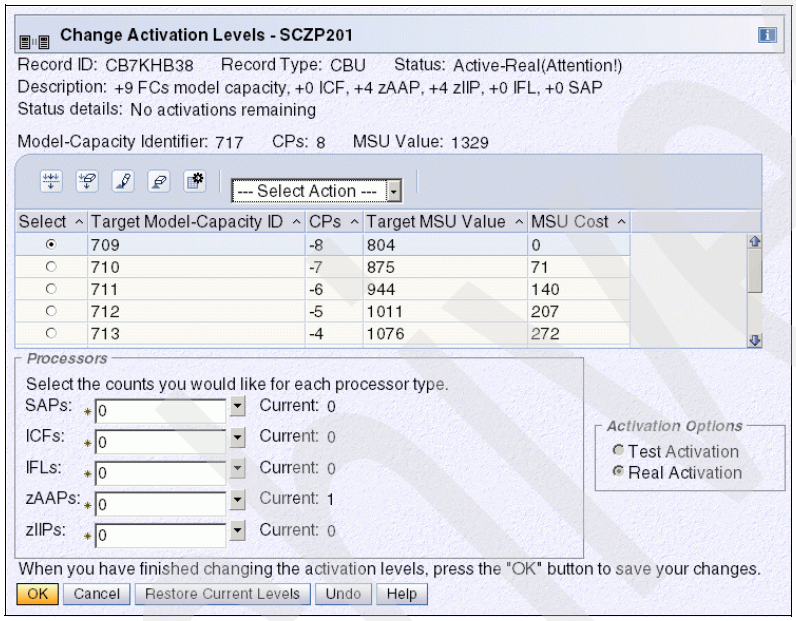

Figure 3-26 Change Activation Levels

Note: If a real CBU activation was performed (even for a short period), the whole CBU record (including all test activations) cannot be used until the activation value is replenished.

The Temporary Upgrades panel in Figure 3-27 on page 100 shows that the CBU was deactivated and the server configuration is back to its original. No temporary processor is active, the MSU value is set back to original. The Model-Capacity Identifier, the Model-Permanent-Capacity Identifier, and the Model-Temporary-Capacity Identifier are the same again. Also, the number of available PUs is restored to the original value.

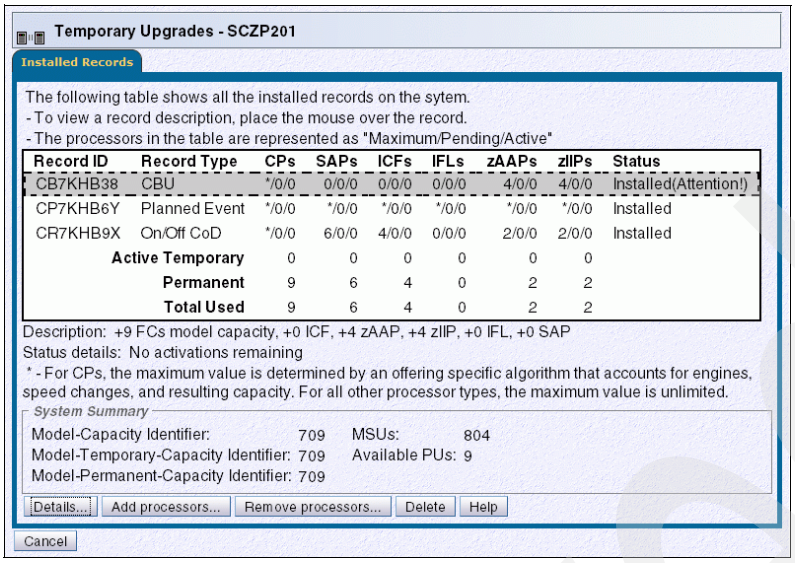

Figure 3-27 Temporary Upgrades - CBU record deactivated

Figure 3-28 shows an output from the D M=CPU command. As shown, processors 0D - 10 are no longer available and the capacity setting of a server is back to 709.

```
D M=CPU
IEE174I 14.34.17 DISPLAY M 002
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1
```

Figure 3-28 D M=CPU output after CBU deactivation

The RMF Output, shown in Figure 3-29, confirms a capacity setting of 709. CPC Capacity is back to its original value and Image Capacity is recalculated to its original value. Because there are nine physical CPs in the server and nine logical CPs in z/OS, the Image Capacity equals the CPC Capacity.

Keep in mind that the temporary record deactivation might require checking and adjusting the logical partition weights, especially if they had been modified previously.

```
Partition:   A01        2097 Model 709
CPC Capacity:    804    Weight % of Max: ****   4h Avg:  145   Group:   N/A
Image Capacity:  804    WLM Capping %:    0.0   4h Max:  448   Limit:   N/A
```

Figure 3-29 RMF Output after CBU deactivation

3.4 Replenishment

The Record Details panel in Figure 3-30 shows that the CBU record has no remaining real activations.

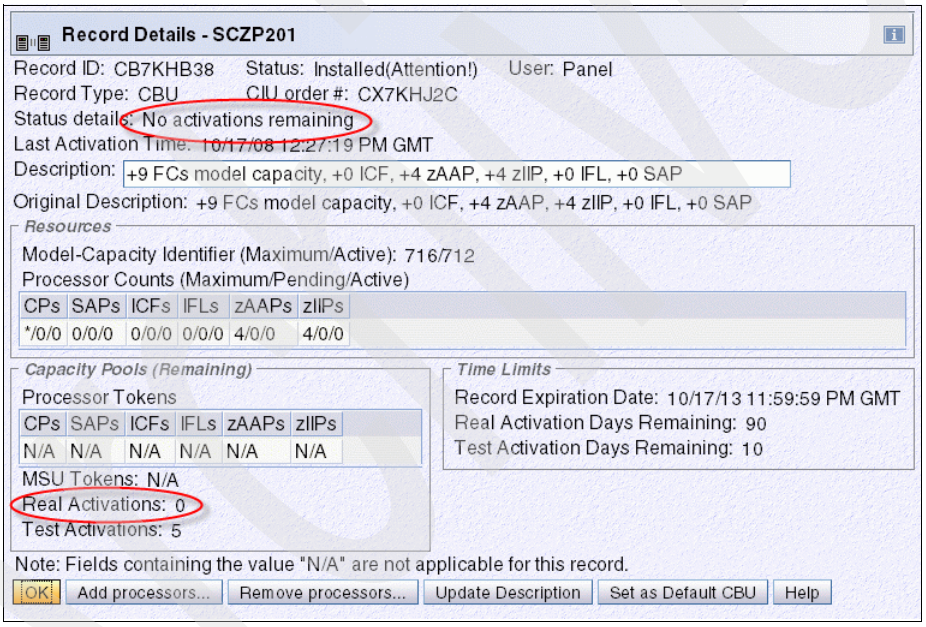

Figure 3-30 Record Details panel - no real activation remaining

CIU application on Resource Link is used to replenish an activation. It can be done only after Resource Link is aware that the activation took place.

Use one of the following methods to determine whether Resource Link is aware of the real activation are:

- Verify that e-mail was sent from Resource Link. A sample e-mail is shown in Example 3-2.

Example 3-2 Message from resource Link

```
Subject: Real Activation: CBU order CX7KHJ2C for machine 2097 -- 1DE50 has used
its only real activation

Customer number: ITSO001

CBU order CX7KHJ2C for machine 2097 -- 1DE50 has used its real activation.
You must replenish the real activation for this Capacity Backup to use it
for a future emergency.

 You can order a replenishment record from the CBU record page:
https://www.ibm.com/servers/resourcelink/ciu03010.nsf/cbuRecord/CB7KHB38?OpenDo
cument
```

- Open the Replenishment window in CIU applications to review the activation replenishment, as shown in Figure 3-31.

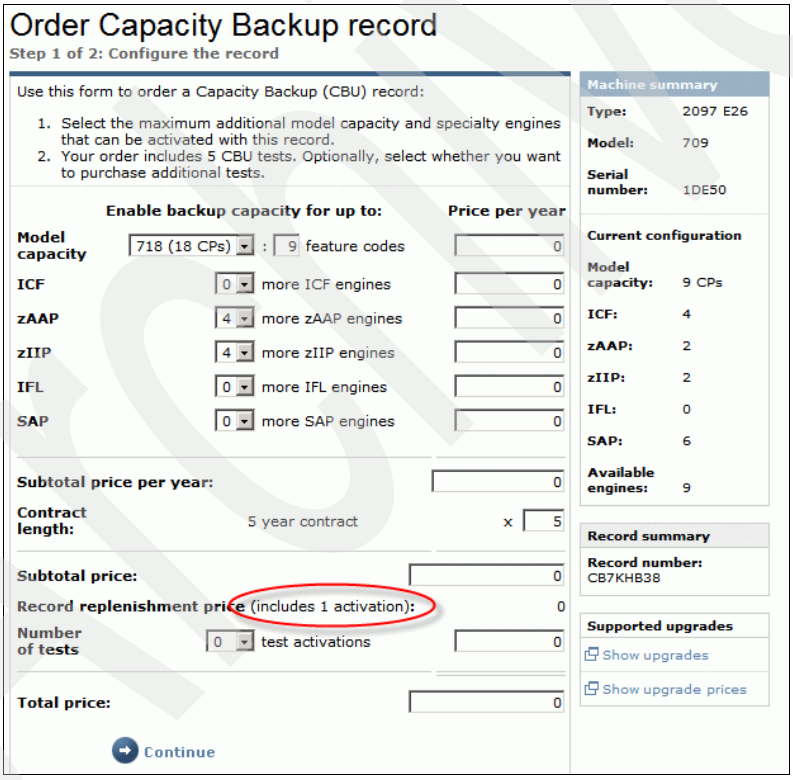

Figure 3-31 Order Capacity backup record - replenishment

When a record is retrieved by using the Perform Model Conversion panel, the installed record is automatically replenished. To verify replenishment, review the Record Details panel, as shown in Figure 3-32 on page 103.

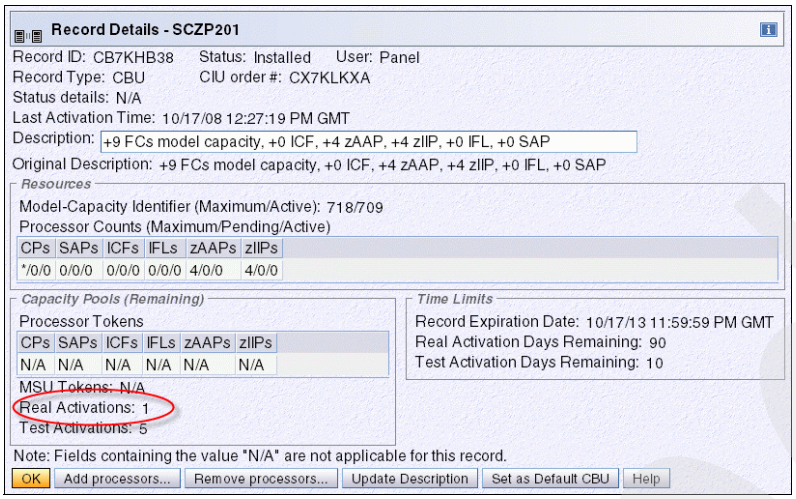

Figure 3-32 Record Details panel - replenished record

4

Capacity for Planned Events

The Capacity for Planned Events (CPE) provides replacement capacity for planned events. It is not intended for peak demand or disaster recovery usage. All unused PUs on the server are available for use for 72 hours as a replacement capacity. Examples of a planned event are a processor upgrade or computer room maintenance. Details about CPE are in 2.1.3, “Capacity for Planned Events” on page 26.

This chapter discusses the Capacity for Planned Events (CPE) offering usage, and the tasks to consider and plan.

This chapter discusses the following topics:

- “Preparation” on page 106
- “Activation” on page 109
- “Deactivation” on page 120

4.1 Preparation

The CPE has a 72-hour time limit and can be activated once. The CPE record makes all the unused processing units (PUs) available upon its activation. The PUs can be characterized as a CP, IFL, ICF, zAAP, zIIP, or SAP when the CPE record is activated.

4.1.1 Determine the requirements

Preparation should start with analyzing how much replacement capacity is necessary to support a given workload, and what types of PUs are necessary. If enough available PUs are on the server, and replacement capacity will be needed for less than 72 hours, CPE is a good solution.

All preparation steps depend on how the temporary capacity is intended to be used. It can be used for a new partition, it can be given to an existing partition, or as a combination of both approaches. Temporary capacity usage depends on the particular situation and should be determined during planning. In the example used in this chapter, a new logical partition is activated and processors are added to an existing partition. The example configuration consists of a z10 EC model E26, with nine CPs, four ICFs, two zAAPs, and two zIIPs. The CPE record can activate nine additional PUs.

During preparation, several areas must be addressed. The I/O device and network connectivity have to be in place to support the new workload. Before activation, determine the target partition and prepare its profile image so that activation can occur with minimal delay.

When resources are added to a server, the relative capacity of all partitions can change. Review the logical partition weights to ensure that the required capacity can be delivered to each logical partition. The zPCR tool is helpful in modeling the changes of adding more resources.

Performing required tasks in advance helps you to avoid problems, stress, and mistakes, when the CPE is actually used.

4.1.2 Order, retrieve, and install a record

This section contains a brief description of the steps performed when a CPE record is ordered from Resource Link. You may also order a record from an IBM Sales Representative.

Order the record

The CPE record is ordered through the Resource link, as shown in Figure 4-1 on page 107. The server configuration information is on the right side of the panel. The server (2097-E26) has nine CPs, four ICFs, two zAAPs, two zIIPs, and nine available engines. There are no customizable fields for CPE except the description, because CPE allows usage of any available PU. The order is submitted from this panel.

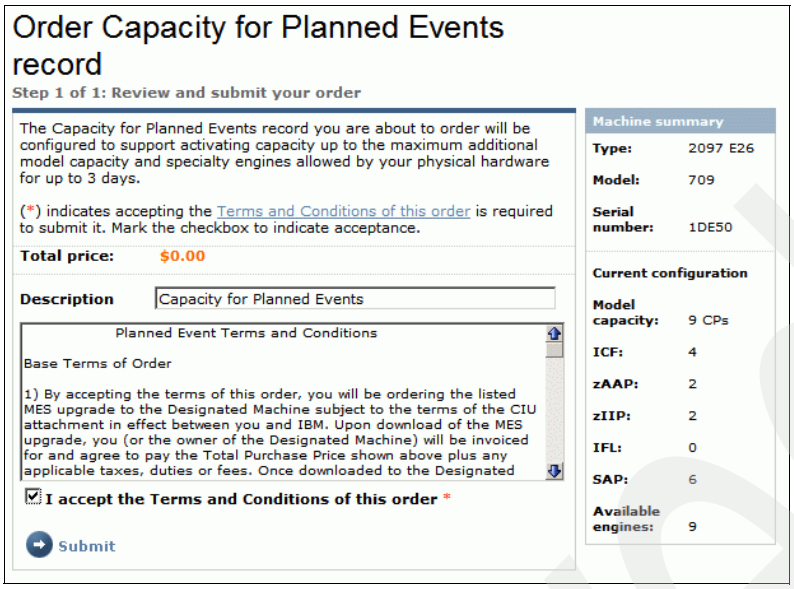

Figure 4-1 Resource link ordering

When an order is submitted, it must be approved by an approver if one is defined. After the order is approved and ready to be retrieved, a notification e-mail is sent stating that a record is ready for a retrieval. Example 4-1 shows an excerpt from a notification e-mail.

Example 4-1 Notification e-mail

```
Subject: Download ready: Capacity for Planned Events order CX7KHB6Y for machine
(2097 -- 1DE50) is ready to be downloaded and installed

Customer number: ITSO001

Capacity for Planned Events order CX7KHB6Y for machine 2097 -- 1DE50 has been
processed and is ready to be downloaded to the customer's machine.
Please contact the appropriate person to retrieve this order from the Hardware
Management Console (HMC).



Order Create Date: 10/17/2008
Record number: CP7KHB6Y
Order number: CX7KHB6Y

Order Confirmation Information


This Capacity for Planned Events order is limited to a single activation for 3
days
```

Retrieve the record

When the record is ready, use the Perform Model Conversion task from the IBM Service Support System, as shown in Figure 4-2 on page 108.

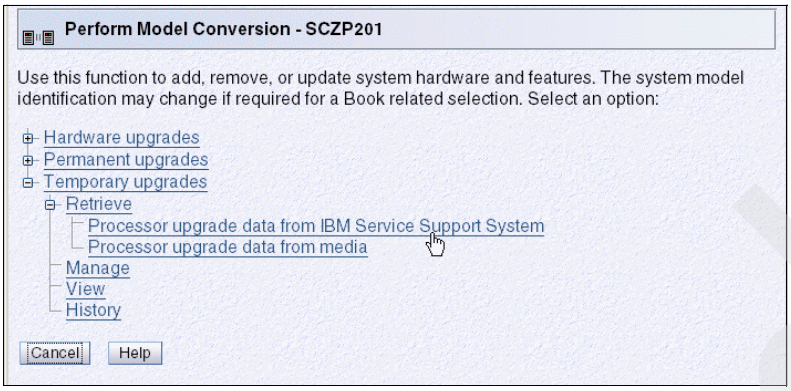

Figure 4-2 Perform Model Conversion

The retrieved record is placed in the staged records area on the Support Element (SE).

Note: Any records ordered with the server and delivered from manufacturing will be staged on the SE and they must be installed before they can be activated.

The records in the staged area can be viewed, installed, or deleted; see Figure 4-3.

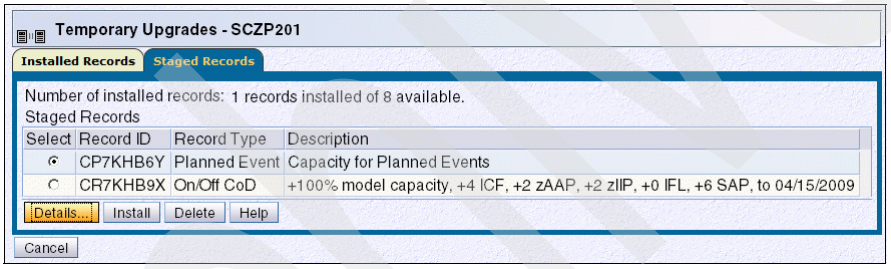

Figure 4-3 Temporary upgrades - CPE record staged to SE

Click **Details** to show the Record Details panel. The panel, shown in Figure 4-4, displays the record ID, order number, the number of activations, and the duration information. No processor assignment is shown because it is determined by the user at the time of activation.

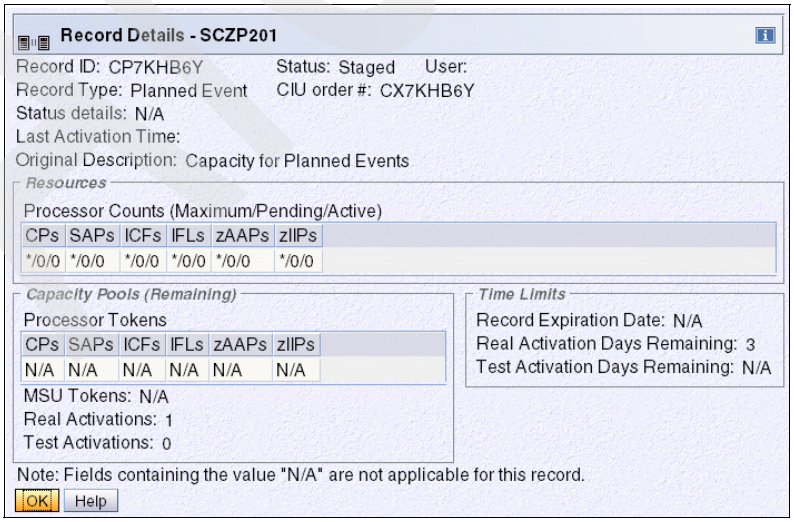

Figure 4-4 Record details

Install the record

To install the record, click **Install** on the Temporary Upgrades panel (see Figure 4-3 on page 108). The record is moved into one of the eight slots for installed records. Because moving an installed record back to a staged record later is not possible, a confirmation message ACT37466 is displayed, as shown in Figure 4-5.

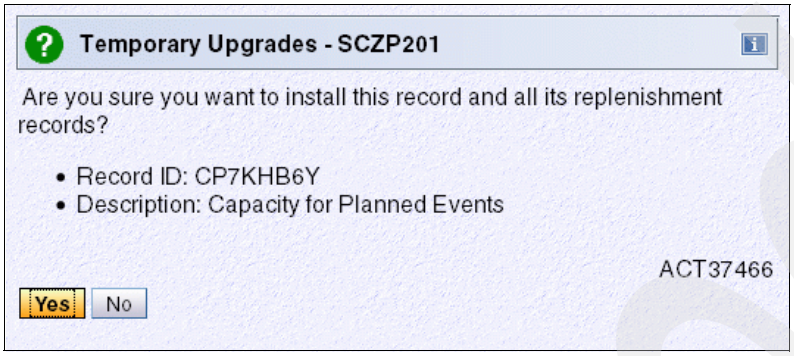

Figure 4-5 Temporary upgrades, message ACT37466

4.2 Activation

After everything is prepared, actual CPE activation can be performed. Review the current configuration before activating the record. Reviewing can performed in several places, depending on the tools and practices available. In our example, configuration information available from HMC and from RMF.

Activate CPE

The Temporary Upgrades panel shows the Model-Capacity Identifiers, the types and counts of PUs, and the number of available PUs.

Figure 4-6 shows three temporary records installed, none of them active.

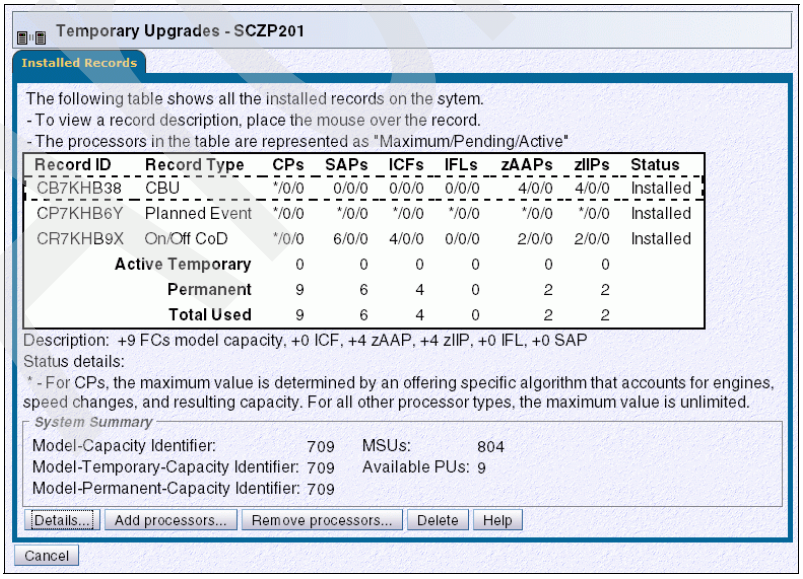

Figure 4-6 Temporary upgrades

As shown in Figure 4-7, RMF reports the Model number and CPC Capacity in MSUs. They match the values seen on SE.

```
 Partition:   A01        2097 Model 709
 CPC Capacity:     804  Weight % of Max: ****   4h Avg:  151   Group:   N/A
 Image Capacity:  804  WLM Capping %:    0.0  4h Max:  448   Limit:  N/A
```

Figure 4-7 RMF Report without CPE

With the base configuration confirmed, the next action is to activate the CPE record. Figure 4-8 shows that a CPE record is selected. All Capacity Identifiers show 709, which is the permanent capacity. The panel also displays the number of available PUs, in this case nine are available. These are the resources that are available and that can be assigned to the different engine types as part of the activation.

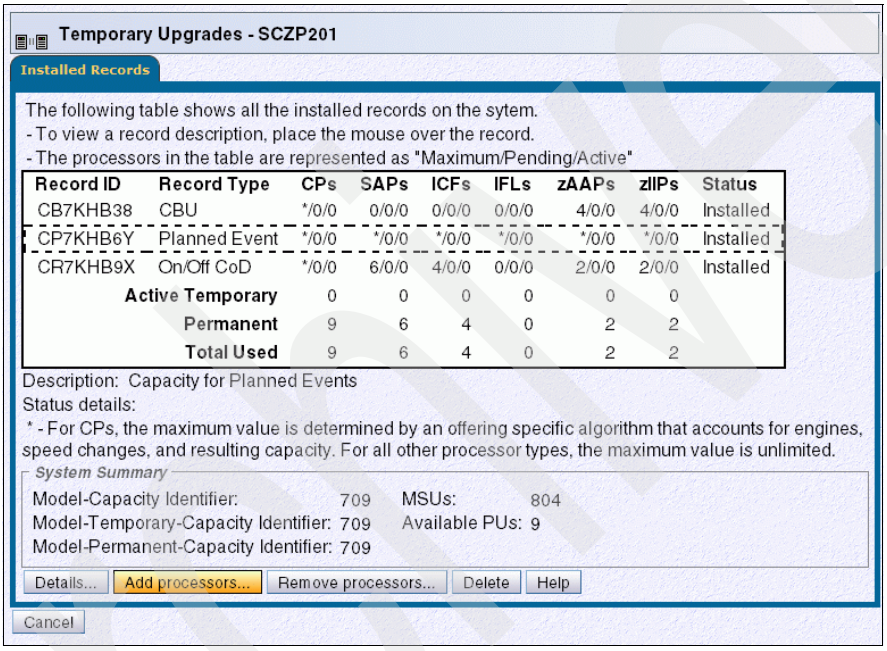

Figure 4-8 Temporary upgrades

After you click **Add processors**, the Change Activation Levels panel is displayed; see Figure 4-9 on page 111. Use this panel to select the target capacity and specialty engines. The selected target model here is 716, which adds seven of the available PUs as CPs. In addition, one zAAP and one zIIP are selected. These selections used all the available PUs.

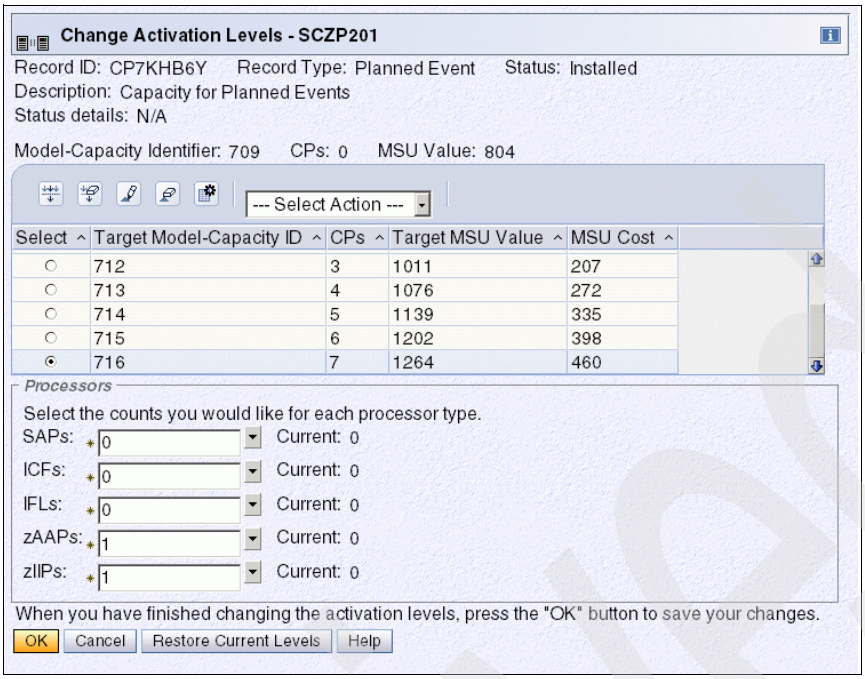

Figure 4-9 Change activation levels

After the definitions are complete, click **OK** to proceed with the activation.

The confirmation message, ACT37464, is displayed. The original configuration is shown beside the proposed configuration; see Figure 4-10. Click **Yes** to perform the configuration change.

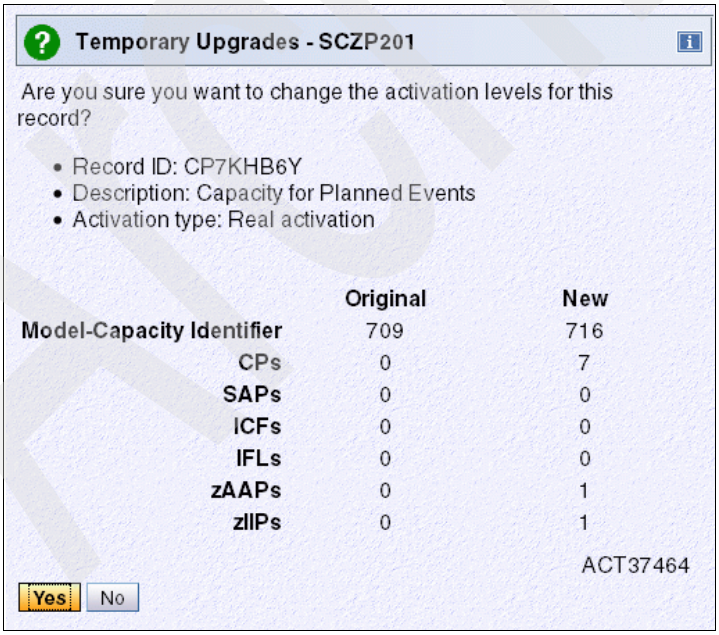

Figure 4-10 Temporary Upgrades, message ACT37464

After the resources in the CPE record are activated, the new configuration is displayed on the Temporary Upgrade panel, as shown in Figure 4-11. The Model-Capacity Identifier has changed to 716, which represents nine permanent CPs plus seven temporary CPs. Both Model-Temporary-Capacity Identifier and the Model-Permanent-Capacity Identifier remain at 709, because permanent capacity was not changed and billable capacity is not used.

Also, temporary zAAPs and zIIPs are displayed. The server can now use 16 CPs, three zIIPs, three zAAPs, and four ICFs. The status of the CPE record is now *Active Real (Attention!)*. The Status details field explains the Attention flag. It says that no activations remain in the record and only a few activation days remain.

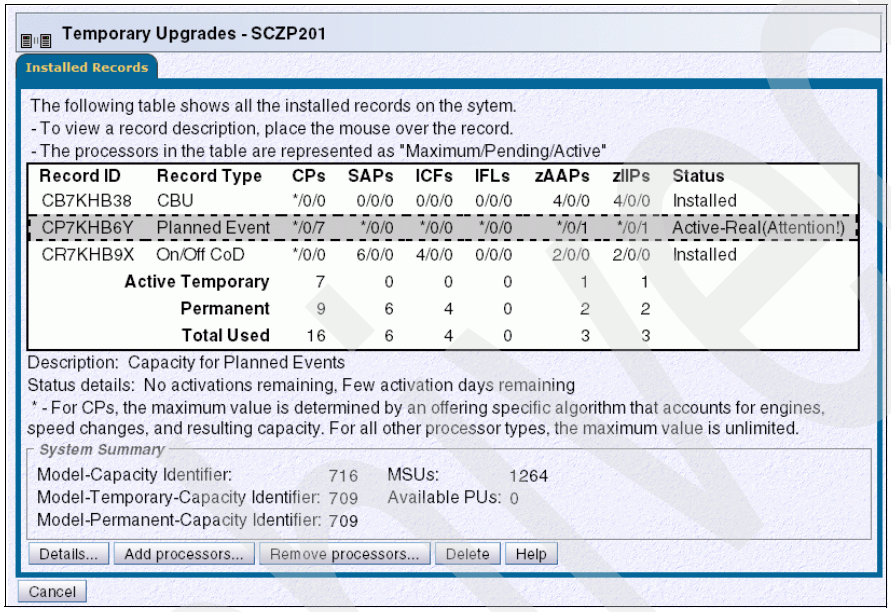

Figure 4-11 Temporary Upgrades, CPE active

The capacity changes are also detected by RMF; see Figure 4-12. The CPC Capacity Report shows new Model 716 and new CPC capacity in MSUs. Image Capacity was recalculated also, as described in 2.3.6, “z/OS Resource Measurement Facility” on page 57. The RMF report is useful for checking how z/OS sees additional capacity.

```
Partition:   A01        2097 Model 716
CPC Capacity:    1264   Weight % of Max: ****   4h Avg:  581   Group:   N/A
Image Capacity:   790   WLM Capping %:    0.0   4h Max:  617   Limit:   N/A
```

Figure 4-12 RMF Report showing additional resources

All three capacity identifiers can also be checked in the Product Information tab in the CPC Details panel for a server. See Figure 4-13 on page 113. The specialty engines are not reflected in the capacity identifier fields.

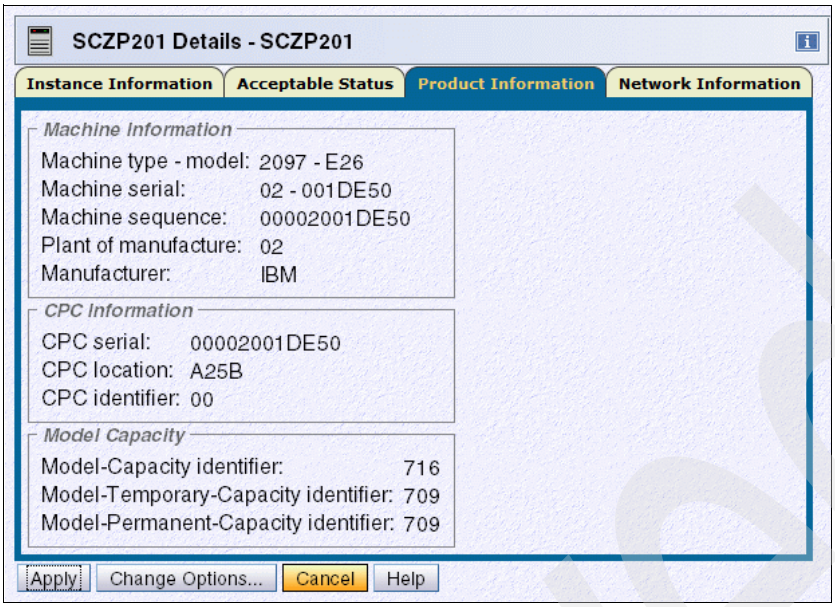

Figure 4-13 HMC Product Information display

Depending on particular needs, two scenarios, or their combination, will happen:

- If another logical partition is needed, it can be activated now with the replacement capacity. If planning was done correctly and the logical partition was defined and set up with the required CPs, reserved CPs, memory, and the partition weights, the server will now be able to handle the additional workload.
- If a partition requires an additional processor, depending on operating system functionality, it must adhere to either of the following items or a processor cannot be added without partition deactivation:
  - Have reserved processors defined in its image profile
  - Support dynamic processor add

In both scenarios, careful planning of partition weights is strongly advised. If needed, partition weights should be reevaluated and adjusted to match the new requirements. Partition processing weights can be changed dynamically through LPAR controls.

Figure 4-14 on page 114 displays the D M=CPU command output before the CPE record was activated. There are nine CPs active in this z/OS, with IDs 00-03 and 08-0C. Processor 0D has an `N` next to it, which means the physical processor is not available. This is correct because there were only nine CPs on the server before the CPE record was activated. This z/OS image is defined with four initial CPs and six reserved CPs. The `CPC SI` output shows capacity setting 709.

```
D M=CPU
IEE174I 17.09.54 DISPLAY M 922
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
```

Figure 4-14 Display CPU before CPE was active

The output of D M=CPU after CPE activation is shown in Figure 4-15 on page 115. Processor 0D is now in the offline state. This means that the physical processor is available in a pool of CPs, and it is offline in z/OS. The same can be observed for zAAP processor 0E and zIIP processor 10. The `CPC SI` line reflects a new capacity setting of 716.

```
D M=CPU
IEE174I 17.11.46 DISPLAY M 922
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  -
0E  -A
0F  NA
10  -I
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.716.IBM.02.000000000001DE50
```

Figure 4-15 Display CPU after CPE was active

Additional processors can now be varied online in z/OS with the CONFIG CPU command.

Figure 4-16 on page 116 shows the status of processors after processors 0D and 0E were varied online from z/OS. Both now have a plus (+) sign, which indicates they are online for z/OS.

```
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DE502097
01  +                     01DE502097
02  +                     01DE502097
03  +                     01DE502097
04  +A                    01DE502097
05  +A                    01DE502097
06  +I                    01DE502097
07  +I                    01DE502097
08  +                     01DE502097
09  +                     01DE502097
0A  +                     01DE502097
0B  +                     01DE502097
0C  +                     01DE502097
0D  +                     01DE502097
0E  +A                    01DE502097
0F  NA
10  -I
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.716.IBM.02.000000000001DE50
```

Figure 4-16 Display CPU after the processor was configured on line

Change the activation level

Capacity on Demand design provides great flexibility, including the ability to reconfigure resources after activation. If a different record has to be activated, a permanent upgrade is performed, or even if a change of type and amount of a given temporary resource is needed, the resources activated by the CPE record can be changed dynamically.

To change a capacity setting from 716 to 713, use the Temporary Upgrade task. Select the correct record and then click **Remove processors**; see Figure 4-17 on page 117.

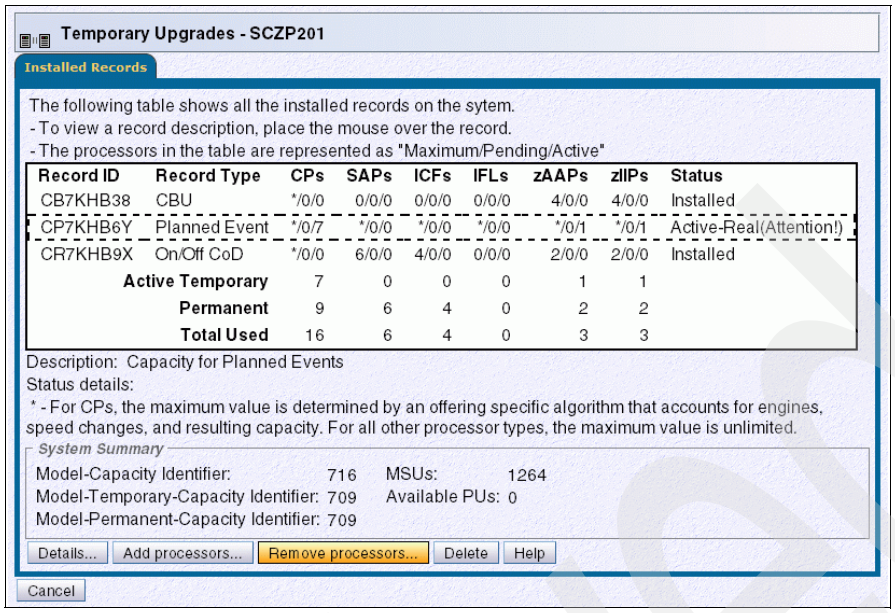

Figure 4-17 Temporary Upgrades

The Change Activation Levels panel opens, where the target model can be changed from a 716 to a 713, as shown in Figure 4-18 on page 118.

Caution: The Perform Model Conversion task on the Support Element does not display an activate or deactivate function. Use the **Add processors** or **Remove processors** buttons to control temporary upgrades. A temporary capacity record is active when the **Add processors** function results in a positive count, or is inactive when the **Remove processors** function brings the count back to zero.

When changing the configuration of the CPE record, do *not* reduce the processors back to the permanent capacity configuration. This causes the CPE Record to deactivate and a new record will be required.

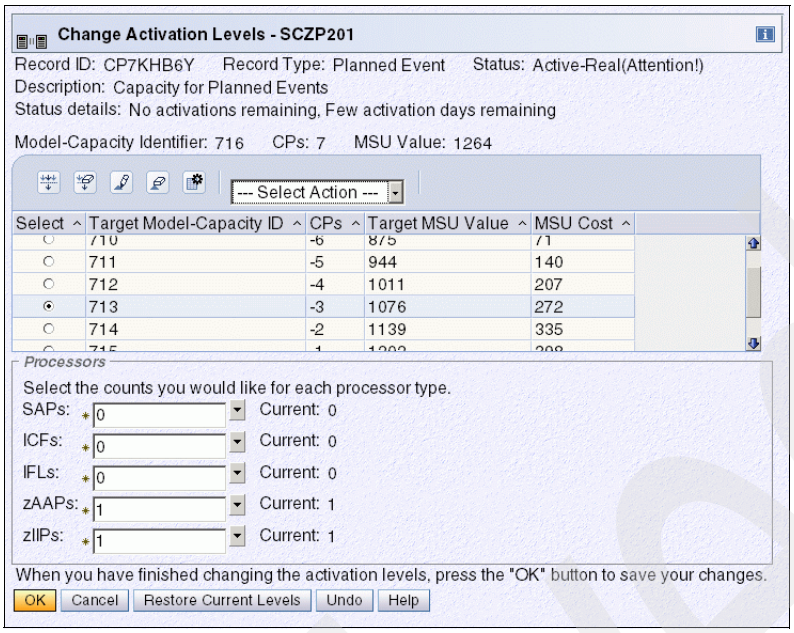

Figure 4-18 Change Activation Levels

The message ACT37464, in Figure 4-19, shows the requested change. In this case, the number of CPs is reduced from seven to four, although all other resources stay unchanged. Clicking Yes, confirms the action. After confirmation, the temporary resource reconfiguration is performed.

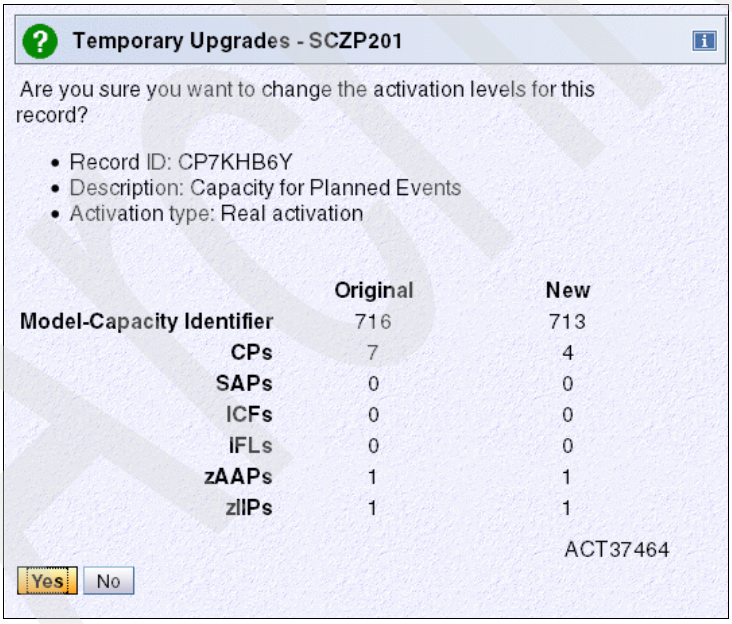

Figure 4-19 Temporary Upgrades - message ACT37464

Confirmation message ACT37155 is displayed, shown in Figure 4-20 on page 119. As mentioned previously, reviewing and understanding the actions being performed are critical. Reduction of resources could put the server in a resource shortage and there is a possibility of a performance problem. The warning message (ACT37155) is a reminder to verify the actions before performing them.

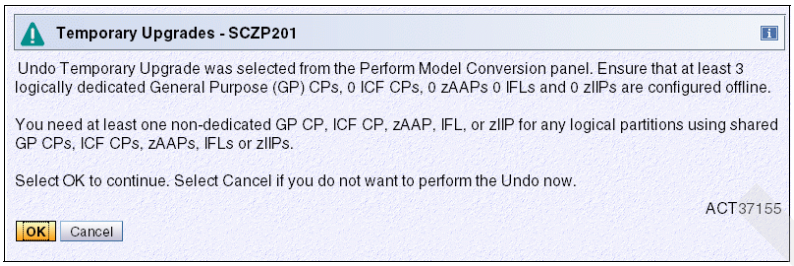

Figure 4-20 Temporary Upgrades - message ACT37155

The Temporary Upgrade Panel in Figure 4-21 shows a new configuration. The Model-Capacity Identifier is changed to 713, but the Model-Temporary-Capacity Identifier and Model-Permanent-Capacity Identifier remain at 709. Three PUs are now available.

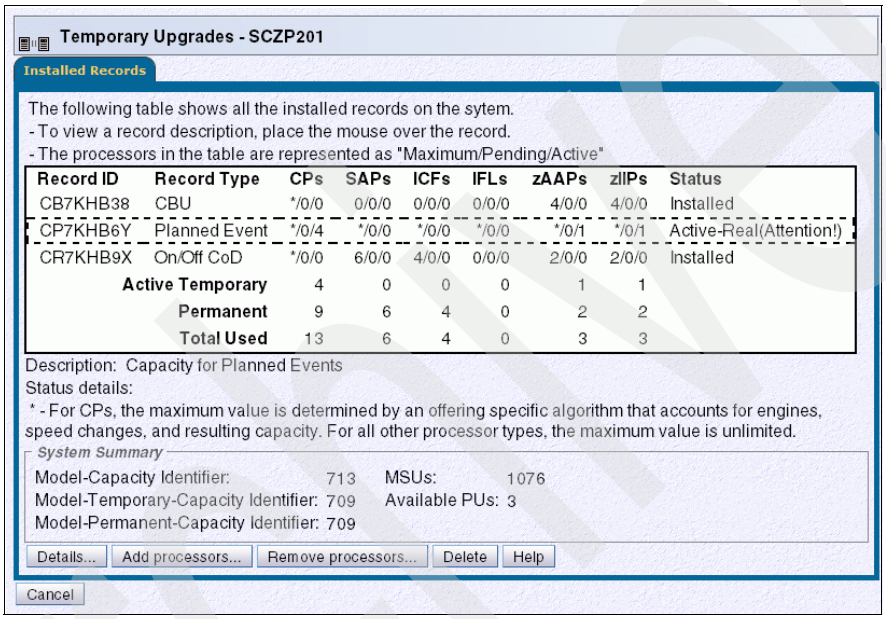

Figure 4-21 Temporary Upgrades

The RMF CPC Report (Figure 4-22) confirms the configuration change from the z/OS perspective.

```
Partition:   A01        2097 Model 713
CPC Capacity:    1076  Weight % of Max: ****   4h Avg:  581   Group:   N/A
Image Capacity:   828  WLM Capping %:    0.0   4h Max:  617   Limit:   N/A
```

Figure 4-22 RMF CPC Capacity Report with the 713 Model

4.3 Deactivation

The active CPE record can be deactivated as follows:

- Automatically at the end of the 72-hour period
- Manually by the operator before the 72-hour period expires

If portions of the resources activated by the record cannot be freed, all of the resources remain active. A check is made every 24 hours and notification is sent to the IBM Support Systems.

Note: Regardless of the deactivation method, removing temporary resources from active images before deactivation is important. Otherwise, the danger is that either resources could not be deactivated or resource deactivation can have a big performance impact.

Because the CPE record has only one activation allowed, the record can no longer be activated. An attempt to reactivate the record through the Temporary Upgrade task can result in error message ACT37469, as shown in Figure 4-23.

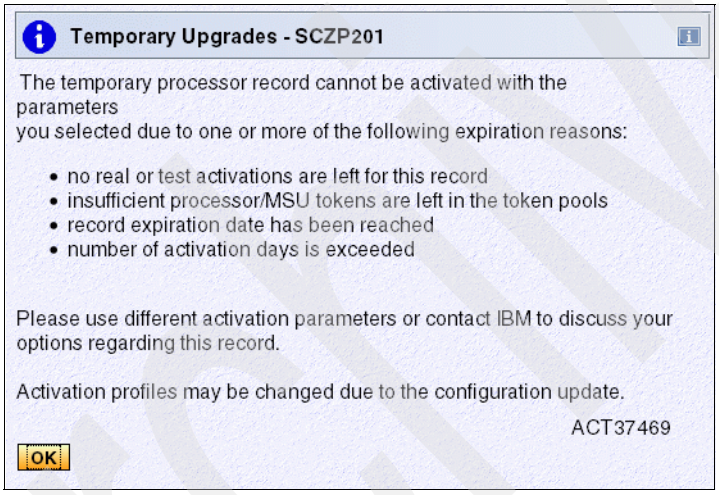

Figure 4-23 Temporary Upgrades - expired record activation failed, message ACT37469

The record's status changes from Active to Installed when it is deactivated. After the record is deactivated, it is no longer usable and should be deleted to avoid confusion.

To delete the CPE record from the installed records, select the correct record and click **Delete** on the Temporary Upgrade panel, as shown in Figure 4-24 on page 121. The deleted record is removed from the Installed records.

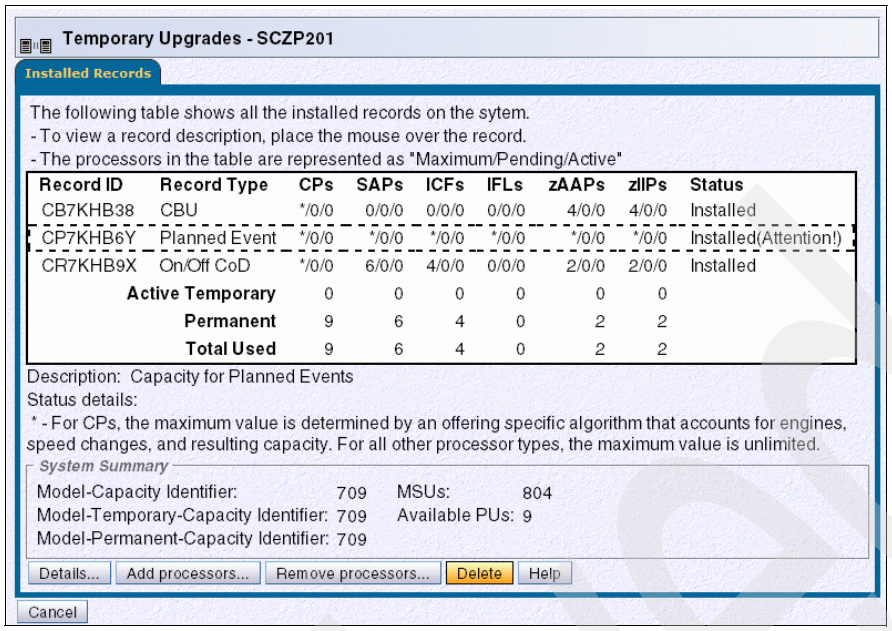

Figure 4-24 Temporary Upgrades

5

On/Off Capacity on Demand

On/Off Capacity on Demand (On/Off CoD) provides a billable temporary capacity to handle peak workload situations. Details about On/Off CoD are in 2.1.4, “On/Off Capacity on Demand” on page 27.

This chapter describes On/Off Capacity on Demand (On/Off CoD) offering usage, and the tasks to consider and plan.

This chapter discusses the following topics:

- “Preparation” on page 124
- “Activation” on page 128
- “Deactivation” on page 141

5.1 Preparation

This section describes the steps to use when an On/Off CoD record is prepared to handle a peak workload, when critical workloads suffer from low capacity. The current configuration is a server 2097 E26 with a capacity setting of 709 (9-way), with two zAAPs, two zIIPs, and four ICFs. Multiple logical partitions are running on the server. Two partitions consume the majority of the resources and are weighted to receive about 90% of the total resources.

Typical usage of On/Off CoD consists of adding temporary processors to existing logical partitions, configuring an additional logical processor in the operating system and reviewing the impact, and repeating the steps until the workload is under control. This is a common approach and is driven by the need to resolve the problem while keeping cost under control.

5.1.1 Review planned use of offering

With On/Off CoD, you may add up to twice the capacity of the CPC purchased capacity. The CPC is a 2097 E26 with a permanent capacity setting 709 (804 MSUs), thus the maximum temporary capacity setting available using On/Off CoD is 721 (1571 MSUs).

The capacity are added to the shared pool of CPs for use across all active partitions. Partition definitions are reviewed to ensure that they have reserved CPs, because the number of logical CPs might be increased to allow more resources to be used by the partition. The partition weights also have to be reviewed to determine how the resources should be distributed across partitions.

5.1.2 Check availability of additional resources

According to the configuration of the CPC, there are nine CPs, four ICFs, two zAAPs, and two zIIPS, leaving nine PUs available. Although doubling the current capacity MSU would allow a temporary capacity setting 721, the physical hardware has only nine PUs available for use. The On/Off CoD record can activate capacity settings up to 718.

5.1.3 Order, retrieve, and install a record

This section contains a brief description of steps performed when an On/Off CoD record is ordered from Resource Link.

Order a record

The On/Off CoD order is placed through Resource Link, Customer Initiated Upgrade (CIU) order. Figure 5-1 on page 125 shows the first step in the ordering process. The On/Off CoD configuration is entered. Note the Replenishment due date field entry in the panel. This field is required and is defaulted to 180 days, the maximum value. The date may be changed, but only to an earlier date, not beyond 180 days. Limits defined in this panel are the upper limits of how many resources can be activated by this record at its activation time. At the activation time, the amount of resources to activate can be chosen, and can be changed later within the record authorization space, as described in “On/Off CoD authorization space” on page 33.

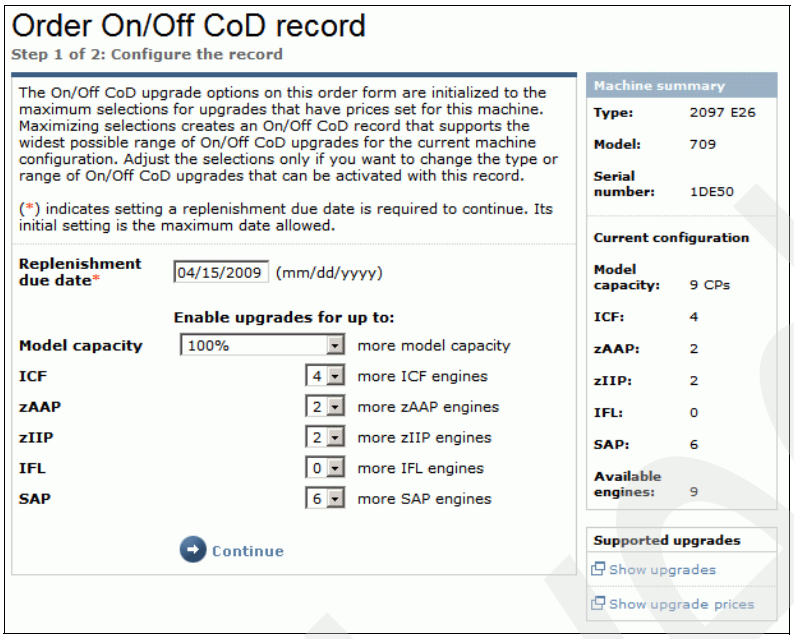

Figure 5-1 Configuration of the On/Off CoD order

After reviewing and submitting the order, it must be approved by an approver if one is defined. After the order is approved and ready to be retrieved, a notification e-mail is sent and indicates that a record is ready to be downloaded. Example 5-1 shows an excerpt from a notification e-mail.

Example 5-1 Notification e-mail

```
Subject: Download ready: On/Off CoD order CX7KHB9X for machine (2097 -- 1DE50) is ready to
be downloaded and installed

Customer number: ITSO001

On/Off CoD order CX7KHB9X for machine 2097 -- 1DE50 has been proessed and is ready to be
downloaded to the customer's machine. Please contact the appropriate person to retrieve
this order from the Hardware Management Console (HMC).

Order Create Date: 10/17/2008
Record number: CR7KHB9X
Order number: CX7KHB9X

Order Confirmation Information


Replenishment
due date:          04/15/2009

Model capacity:    100% more model capacity
ICF:               4 more ICF engines
zAAP:              2 more zAAP engines
zIIP:              2 more zIIP engines
IFL:               0 more IFL engines
SAP:               6 more SAP engines

Description:       +100% model capacity, +4 ICF, +2 zAAP, +2 zIIP, +0 IFL,
+6 SAP, to 04/15/2009
```

Retrieve a record

A record is retrieved by using the SE Perform Model Conversion task shown in Figure 5-2. This is the only time when a connection to an IBM site is required.

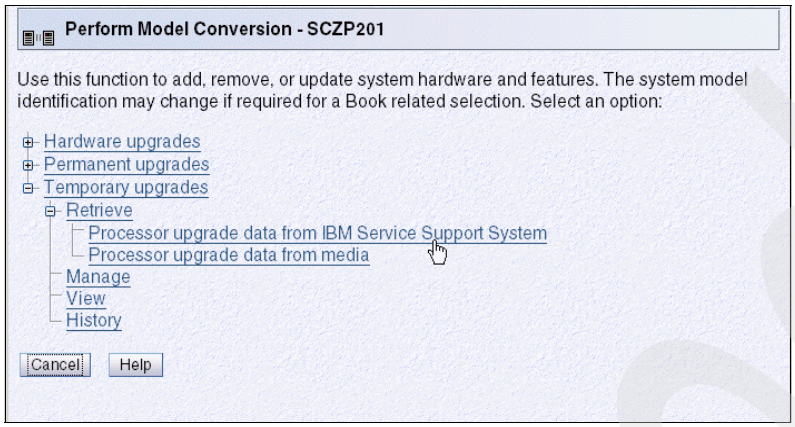

Figure 5-2 Perform Model Conversion

To begin the retrieval process, select the option **Processor upgrade data from the IBM Service Support System**. Retrieved records are staged on the Support Element (SE) hard drive.

Note: Any record ordered with the server and delivered from IBM is staged on the SE. The record must be installed before it can be activated.

To manage or view the staged records, select the **View** or **Manage** option on the Perform Model Conversion panel, as shown in Figure 5-3.

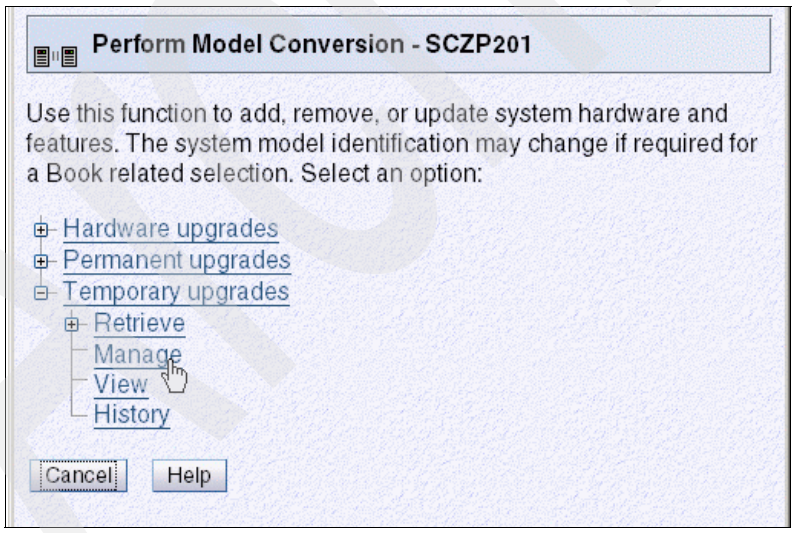

Figure 5-3 Perform Model Conversion

Select the Staged Records tab to display the list of staged records, as shown in Figure 5-4 on page 127. This tab includes all records that have not been installed or deleted yet. The On/Off CoD record that was just retrieved is shown in the Record ID list as CR7KHB9X.

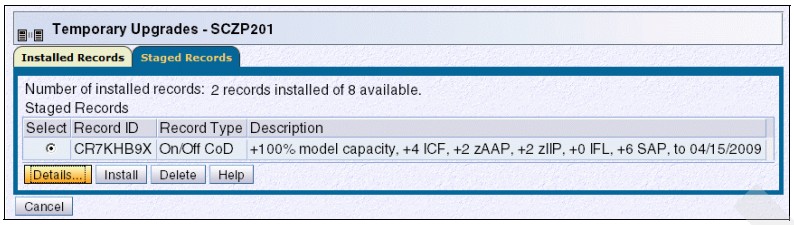

Figure 5-4 Temporary Upgrades - Staged Records

Click **Details.** The Record Details information is displayed, as shown in Figure 5-5.

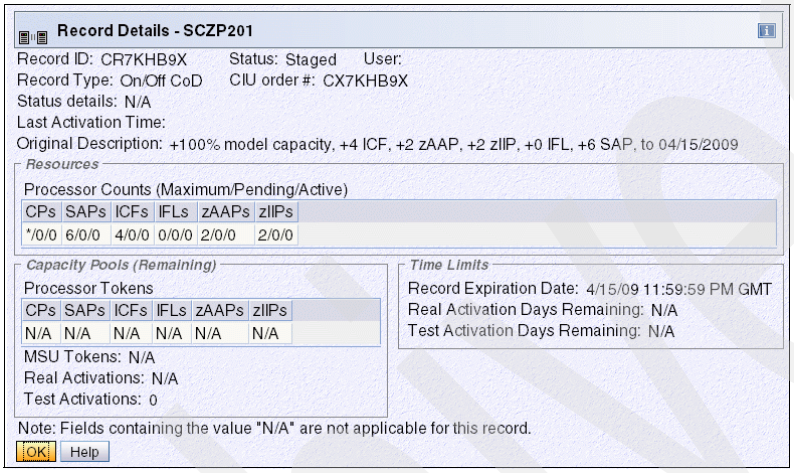

Figure 5-5 Record Details

Install a record

The On/Off CoD record is installed with HMC/SE by using the same path as was used to review the record after it was retrieved. The record is highlighted and the install option selected to perform the installation, as shown in Figure 5-6.

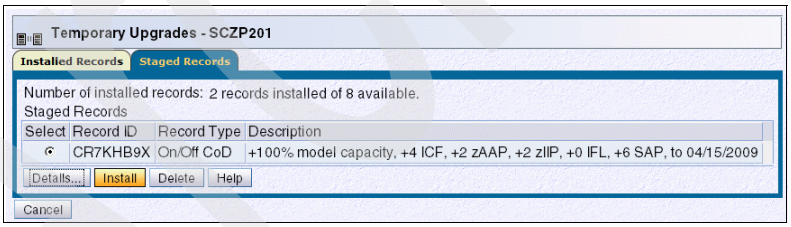

Figure 5-6 Temporary Upgrades - Install On/Off CoD record

The record is now installed and ready for use, as shown in the Installed Records tab in Figure 5-7 on page 128.

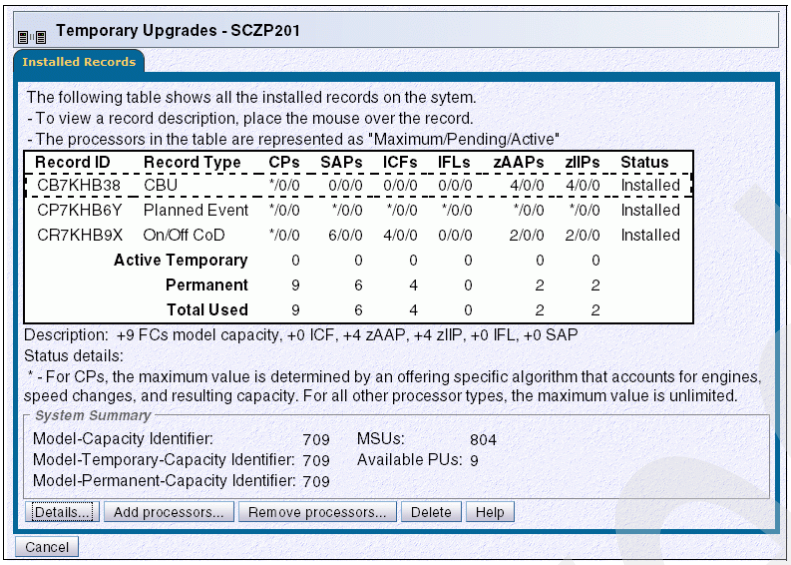

Figure 5-7 Temporary Upgrades

5.2 Activation

After the record is installed, it can be activated when the server requires more capacity for a peak workload. The record can be activated with all the resources available or a subset of them. Changes to the resources can be accomplished dynamically after the record is activated.

Reviewing reports before activation

Figure 5-8 on page 129 shows the status of the processors, from the z/OS image perspective, before activating On/Off CoD. The permanent capacity setting is 709 and the z/OS image has nine CPs online.

```
D M=CPU
IEE174I 10.16.37 DISPLAY M 019
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  NA
0F  NA
10  -I
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 5-8 Display command before activation of On/Off CoD

The RMF CPC capacity report in Figure 5-9 on page 130 shows the MSU value of capacity setting 709. The physical utilization percentage for CPs is 100%. This means that all nine CPs of the CPC are used all the time during this sampling period. The report also shows the number of processors available for logical partition A01, which is nine.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  342   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  675   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.9    100
A0E            0    30  NO   2.0       17.0    17.0    0.0     3.8    3.8
A0F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A01            0   371  NO   9.0       46.1    46.1    0.0    46.1   46.1
A02            0     0  NO   4.0        0.0     0.1    0.0     0.0    0.0
A1E            0    30  NO   2.0       17.0    17.0    0.0     3.8    3.8
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   371  NO   9.0       46.1    46.1    0.0    46.1   46.1
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
PHYSICAL                                                 0.1            0.1
```

Figure 5-9 CPC Capacity report from RMF before activation of On/Off CoD

Using the RMF Monitor III

This section discusses the manual activation of an On/Off CoD record. Resource Measurement Facility (RMF) is used to:

- Identify the workload that is suffering.
- Check the gain of the additional resources.

The System Information report of RMF Monitor III, shown in Figure 5-10 on page 131, indicates service classes defined in Workload Manager (WLM) and their actual performance.

```
Partition:    A01     2097 Model 709          Appl%:      46  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   46     EAppl%:     46  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:  100     Appl% AAP:   0  Time:   11.47.50
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      36  157   14        0.12  5.0  0.1    9.1  0.0  0.0  0.0  0.0
*TSO        100    3    0        0.12  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       36   12   12        0.00  4.4  0.0    7.7  0.0  0.0  0.0  0.0
*STC         34  139    2        0.00  0.6  0.1    1.4  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  35  145    2  .031  0.12  0.7  0.1    1.4  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S  18  119    1  .031  0.12  0.2  0.0    1.0  0.0  0.0  0.0  0.0
          1  18  119    1  .031  0.12  0.2  0.0    1.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S  56   23    1  .000  0.00  0.5  0.1    0.4  0.0  0.0  0.0  0.0
          1  56   23    1  .000  0.00  0.5  0.1    0.4  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  36   12   12  .000  0.00  4.4  0.0    7.7  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S  36   12   12  .000  0.00  4.4  0.0    7.7  0.0  0.0  0.0  0.0
          1  36   12   12  .000  0.00  4.4  0.0    7.7  0.0  0.0  0.0  0.0 2.07
JAVAWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 5-10 RMF Sysinfo report before activation of On/Off CoD

Based on the RMF System Information (Sysinfo) report in Figure 5-10, the critical workload is service class HIPRWKLD, showing a Performance Index (PI) of 2.07, which means its performance is less than half of the intended performance. Because the CPC Capacity report (in Figure 5-9 on page 130) shows that the CPC is 100% busy, and weights already favor partition A01, the determination is that additional capacity is required to address this situation.

To determine if PI can be lowered:

- Review the necessity for resources to address the performance problem versus the cost of the temporary capacity.
- Test the plan to add a small amount of capacity and another logical processor.

To activate the On/Off CoD record, select it from the list of the installed records, as shown in Figure 5-11 on page 132, and select the **Add Processors** option.

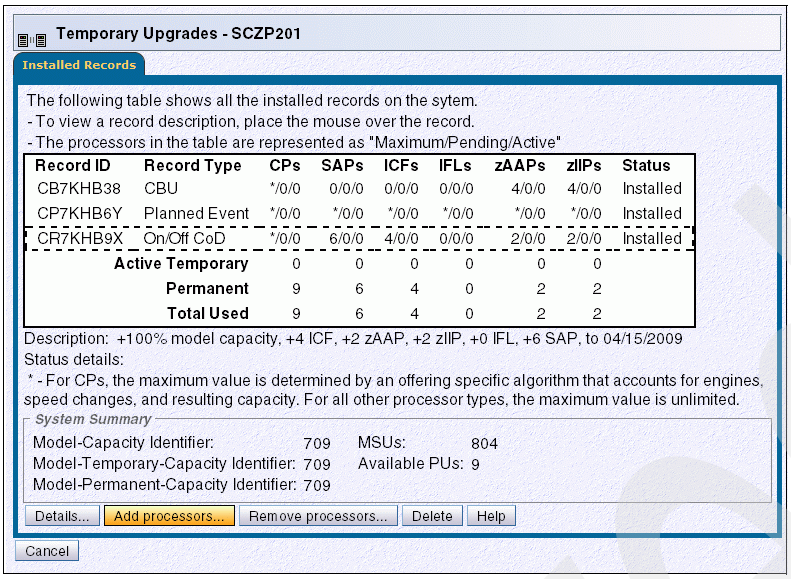

Figure 5-11 Temporary Upgrades - activating On/Off CoD

Normally, resources are added in small amounts to try to correct a problem. If the problem is not corrected, more resources are added. This helps to correct the problem and to control costs.

Because the workload uses only CPs, a single CP is added and the workload must be monitored. With On/Off CoD, the amount of resources to be activated is selectable and can be changed dynamically after activation. The target capacity setting 710 is selected, and no zAAPs or zIIPs are added; see Figure 5-12.

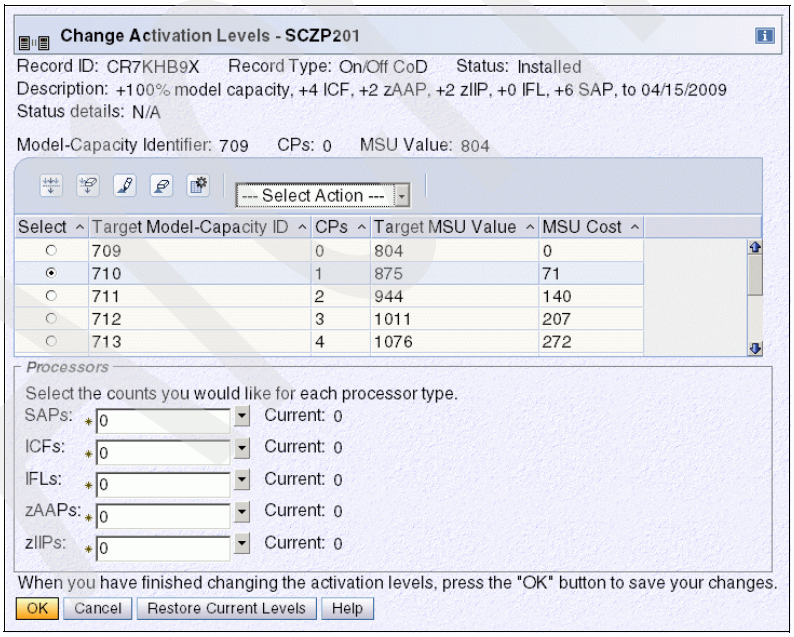

Figure 5-12 Change Activation Levels

Click **OK**. The confirmation panel opens. It lists the current and new configuration with the upgrade details, and the addition of one CP. See Figure 5-13 on page 133.

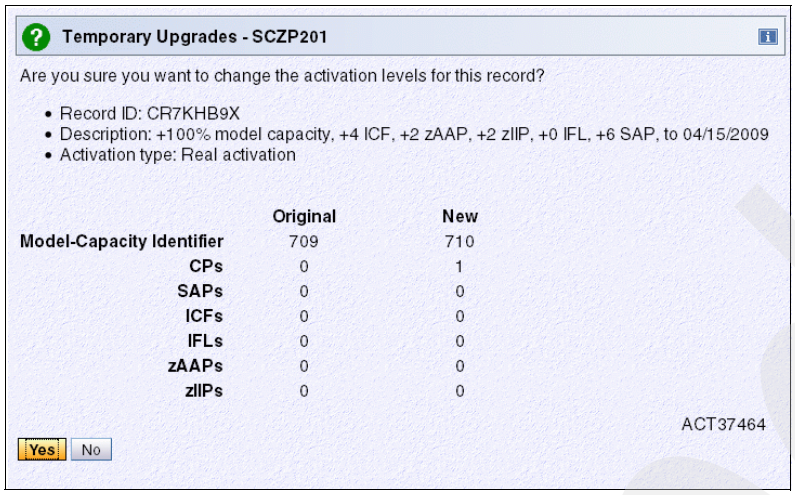

Figure 5-13 Temporary Upgrades - message ACT37464

The Temporary Upgrade panel is used to verify that the additional capacity was added. Checking the status of the record along with the Model-Capacity Identifiers shows the additional CP is now part of the configuration. The Model-Temporary-Capacity Identifier and Model-Capacity Identifier changed to 710 from 709 and the Active Temporary row shows one additional CP. The On/Off CoD record shows Active-Real and that one temporary CP is included in the active configuration; see Figure 5-14.

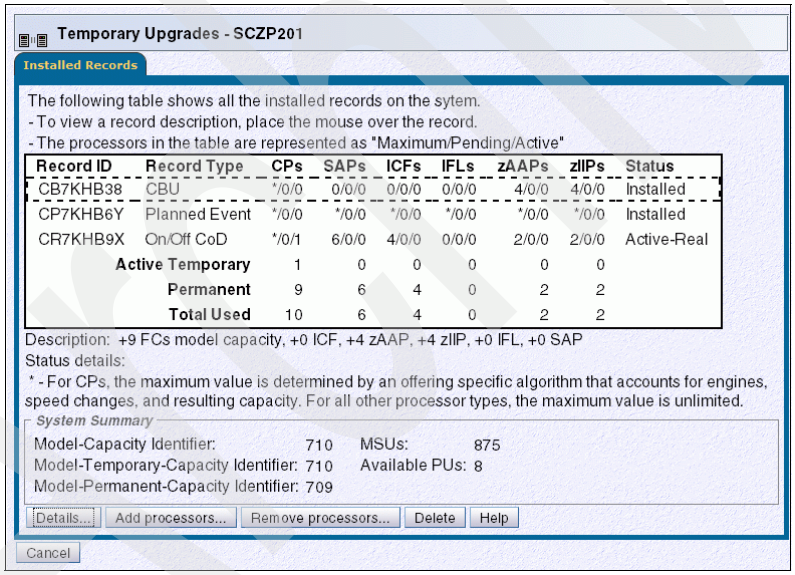

Figure 5-14 Temporary Upgrades - after On/Off CoD activation

The results of issuing the D M=CPU command from z/OS are shown in Figure 5-15 on page 134. When compared to Figure 5-8 on page 129, the processor 0D status is changed from N (not available) to a (offline), after the On/Off CoD was activated. In addition, the model number changed from 2097 709 to 2097 710.

```
D M=CPU
<< Omitted Lines >>
0C  +                    01DE502097
0D  -
0E  NA
<< Omitted Lines >>
CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.710.IBM.02.000000000001DE50
CPC ID = 00
...............................
```

Figure 5-15 Extract of Display command after activating On/Off CoD

Additional changes can be seen in the RMF CPC Capacity report in Figure 5-16. The MSU value for CPC activity changed. The report shows that the physical utilization percentage of the CPs is still 100, and the additional CP is fully consumed. This information indicates that the additional capacity is being used by the logical partition of the CPC. The logical partition is using 45.4% of the 10 CPs, although before it used 46.1% of nine CPs. Making a rough calculation, it is 4.54 CPs versus 4.15 CPs before. The capacity was added to the shared CP pool without changes to the logical partitions. By doing this, the capacity is added but not directly to a given logical partition; it is shared across all active partitions that use shared processors.

```
Partition:   A01       2097 Model 710
CPC Capacity:     875 Weight % of Max: ****   4h Avg:  369   Group:   N/A
Image Capacity:   875 WLM Capping %:    0.0   4h Max:  675   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.1    99.9    100
A0E            0    38  NO   2.0      21.4    21.4    0.0     4.3    4.3
A0F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A01            0   398  NO   9.0      50.5    50.5    0.0    45.4   45.4
A02            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A1E            0    38  NO   2.0      21.4    21.4    0.0     4.3    4.3
A1F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A11            0   401  NO   9.0      50.9    50.9    0.0    45.8   45.8
A12            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
PHYSICAL                                             0.1            0.1
```

Figure 5-16 CPC Capacity report from RMF after activating On/Off CoD

The CPC Capacity report shows how the additional resource has affected the logical partitions. Note that the number of logical processors in partitions has not changed. To determine how this has affected the critical workload, see the RMF Sysinfo report in Figure 5-17 on page 135. It shows that the workflow (WFL column) for service class HIPRWKLD has not changed, but the PI has improved from 2.07 to 1.92. Adding the resource to the shared CP pool allowed additional resources to be consumed across all the logical partitions that needed them.

```
Partition:    A01     2097 Model 710          Appl%:      50  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   50     EAppl%:     50  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:    0     Appl% AAP:   0  Time:   11.47.50
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      35  157   15        0.12  5.0  0.1    9.5  0.0  0.0  0.0  0.0
*TSO          3    0        0.12  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       36   12   12        0.00  4.3  0.0    7.7  0.0  0.0  0.0  0.0
*STC         32  139    3        0.00  0.7  0.1    1.7  0.0  0.0  0.0  0.0
*ASCH         0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  32  145    3  .041  0.13  0.7  0.1    1.7  0.0  0.0  0.0  0.0
SYSOTHER  S       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S  14  119    2  .041  0.13  0.2  0.0    1.3  0.0  0.0  0.0  0.0
          1  14  119    2  .041  0.13  0.2  0.0    1.3  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S  57   23    1  .000  0.00  0.6  0.0    0.4  0.0  0.0  0.0  0.0
          1  57   23    1  .000  0.00  0.6  0.0    0.4  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  36   12   12  .000  0.00  4.3  0.0    7.7  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S  36   12   12  .000  0.00  4.3  0.0    7.7  0.0  0.0  0.0  0.0
          1  36   12   12  .000  0.00  4.3  0.0    7.7  0.0  0.0  0.0  0.0 1.92
JAVAWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 5-17 RMF Sysinfo report after activation of On/Off CoD

Although adding resources to the shared pool showed an improvement to the critical workload, it was not enough because the PI is still high.

As part of the preparation, the logical partition was defined with reserved CPs. The number of logical CPs for the LPAR was changed from nine to ten to determine the impact of adding an additional logical processor to the partition. Logical partition weights are not changed. Processor 0D is configured online, as shown in Figure 5-18 on page 136.

```
CONFIG CPU(D),ONLINE
IEE504I CPU(D),ONLINE
IEE712I CONFIG   PROCESSING COMPLETE
D M=CPU
IEE174I 15.13.57 DISPLAY M 918
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DE502097
01  +                     01DE502097
02  +                     01DE502097
03  +                     01DE502097
04  +A                    01DE502097
05  +A                    01DE502097
06  +I                    01DE502097
07  +I                    01DE502097
08  +                     01DE502097
09  +                     01DE502097
0A  +                     01DE502097
0B  +                     01DE502097
0C  +                     01DE502097
0D  +                     01DE502097
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.710.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 5-18 Config CPU online and display command

The result of adding an additional logical processor on the critical workload, while keeping the weights the same, can be seen in the RMF Sysinfo report in Figure 5-19 on page 137. There was a slight improvement in workflow (WFL) and performance index (PI) for the monitored service class (37 and 1.91, respectively).

```
Partition:    A01     2097 Model 710          Appl%:      46  Policy: POLICY1
CPs Online:  10.0     Avg CPU Util%:   46     EAppl%:     46  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:    0     Appl% AAP:   0  Time:   11.47.50
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      34  157   15        0.29  5.0  0.1   10.0  0.0  0.0  0.0  0.0
*TSO         33    3    0        0.29  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       37   12   12        0.00  4.4  0.0    7.6  0.0  0.0  0.0  0.0
*STC         23  139    3        0.00  0.6  0.1    2.3  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  23  145    3  .204  0.30  0.6  0.1    2.4  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S   8  119    2  .204  0.30  0.1  0.1    1.9  0.0  0.0  0.0  0.0
          1   8  119    2  .204  0.30  0.1  0.1    1.9  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S  52   23    1  .000  0.00  0.5  0.1    0.5  0.0  0.0  0.0  0.0
          1  52   23    1  .000  0.00  0.5  0.1    0.5  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  37   12   12  .000  0.00  4.4  0.0    7.6  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S  37   12   12  .000  0.00  4.4  0.0    7.6  0.0  0.0  0.0  0.0
          1  37   12   12  .000  0.00  4.4  0.0    7.6  0.0  0.0  0.0  0.0 1.91
JAVAWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 5-19 RMF Sysinfo report after configuring an additional CP online

The next step is to add additional CPs and check the effect on performance. Resources in the record can be dynamically changed while the record is active. To activate more of the resources, the Temporary Upgrades task is used again. The active record is highlighted and the **Add processors** option is selected; see Figure 5-20 on page 138.

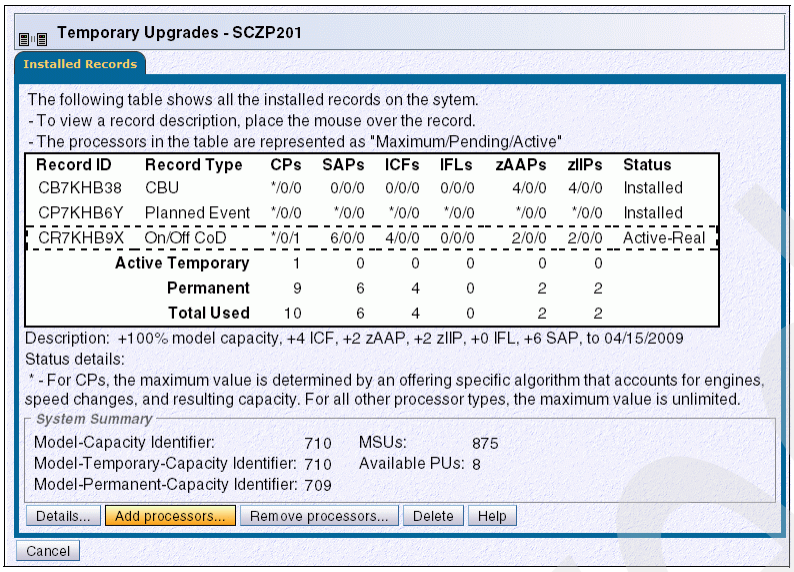

Figure 5-20 Temporary Upgrades

The Change Activation Levels panel is displayed and the target model of 714 is selected, which adds four additional CPs; see Figure 5-21.

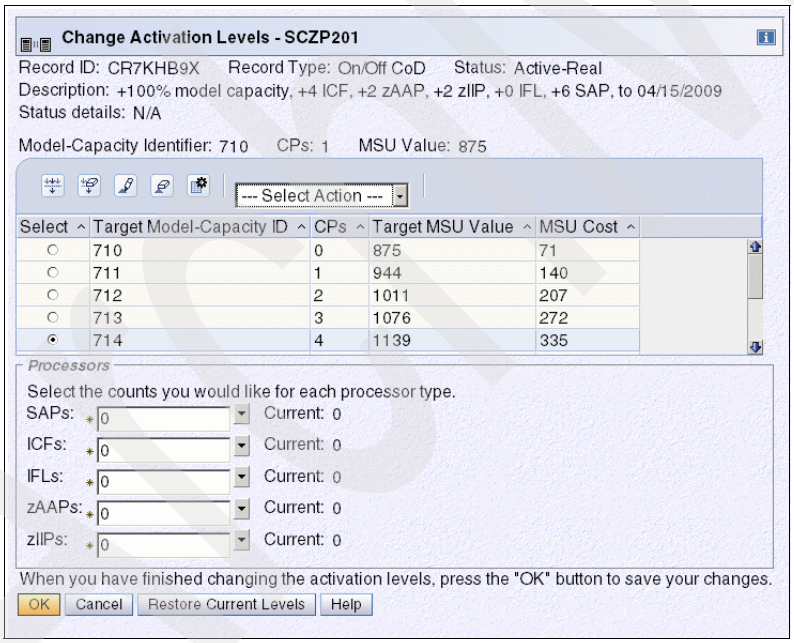

Figure 5-21 Change Activation Levels - selection of additional On/Off CoD capacity

An additional temporary capacity was added to the shared pool to be used by all active logical partitions. Monitored partition A01 can have only ten CPs online and they are already online, so no action can be performed in the monitored z/OS image.

Checking the RMF CPC Capacity in Figure 5-22 on page 139 shows that the additional capacity was added. Whole CPC Capacity increased, while Image capacity decreased. The reason is that partition A01 was able to use ten out of ten CPs before, now it can use only ten out of 14 CPs. CPs continue being used at 100% and logical partition A01 uses them 50.0%

of the time, which translates roughly to seven CPs. It is a considerable improvement and should be reflected in other reports.

```
Partition:   A01        2097 Model 714
CPC Capacity:    1139 Weight % of Max: ****   4h Avg:  408   Group:   N/A
Image Capacity:   814 WLM Capping %:    0.0   4h Max:  675   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        39.0                       0.1    99.9    100
A0E            0    80  NO   2.0      49.1    49.1    0.0     7.0    7.0
A0F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A01            0   570  NO  10.0      70.0    70.0    0.0    50.0   50.0
A02            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A1E            0    80  NO   2.0      49.1    49.1    0.0     7.0    7.0
A1F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A11            0   408  NO   9.0      55.7    55.7    0.0    35.8   35.8
A12            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
```

Figure 5-22 RMF CPC capacity report after adding four CPs

The RMF Sysinfo report shown in Figure 5-23 on page 140 confirms an improvement for the HIPRWRKLD service class. Workflow (WFL) increased to 51% and PI dropped to 1.31.

```
Partition:    A01     2097 Model 714          Appl%:      70  Policy: POLICY1
CPs Online:  10.0     Avg CPU Util%:   70     EAppl%:     70  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:  100     Appl% AAP:   0  Time:   11.47.50
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      45  157   16        0.13  7.0  0.2    8.9  0.0  0.0  0.0  0.0
*TSO         60    3    0        0.13  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       51   12   12        0.00  6.1  0.0    5.9  0.0  0.0  0.0  0.0
*STC         25  139    4        0.00  0.8  0.2    3.0  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  25  145    4  .249  0.13  0.8  0.2    3.1  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S  18  119    3  .249  0.13  0.4  0.1    2.5  0.0  0.0  0.0  0.0
          1  18  119    3  .249  0.13  0.4  0.1    2.5  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S  46   23    1  .000  0.00  0.4  0.1    0.6  0.0  0.0  0.0  0.0
          1  46   23    1  .000  0.00  0.4  0.1    0.6  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  51   12   12  .000  0.00  6.1  0.0    5.9  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S  51   12   12  .000  0.00  6.1  0.0    5.9  0.0  0.0  0.0  0.0
          1  51   12   12  .000  0.00  6.1  0.0    5.9  0.0  0.0  0.0  0.0 1.31
JAVAWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 5-23 RMF Sysinfo report after adding four CPs

The performance has improved after the additional capacity was provided. PI is still high, and the service class is not meeting its WLM goal. Adding more processors is possible, which requires a decision of whether to consider costs versus business importance; in this case, more processors will not be provided.

Using the RMF Performance Monitor

An alternative to RMF Monitor III panels is the RMF Performance Monitor. As shown in Figure 5-24 on page 141, you use the RMF Performance Monitor to check the performance metrics of a system in a graphical view.

Relevant Data views can be chosen and monitored the same way as in RMF Monitor III panels. In this case, there are processor-relative metrics (the three data views on the left) and a performance index for the critical workload HIPRWKLD in both systems separately and in the whole sysplex (the three data views on the right).

Figure 5-24 RMF PR Data views

RMF PM can be downloaded from RFM at:

`http://www.ibm.com/servers/eserver/zseries/zos/rmf`

5.3 Deactivation

After the peak workload has finished and the temporary capacity is no longer needed, it is time to deactivate a record.

Before removing temporary capacity, always check what it means for a running system. In this case, changing the capacity back to 709 would mean leaving the z/OS image in partition A01 with ten online processors that are dispatched on nine physical processors. Therefore, an important step is for those logical resources (that used temporary physical resources) to be put offline. One logical processor should be configured offline from the z/OS running in the A01 partition. This is shown in Figure 5-25 on page 142.

```
CONFIG CPU(D),OFFLINE
IEE505I CPU(D),OFFLINE
IEE712I CONFIG   PROCESSING COMPLETE
D M=CPU
IEE174I 15.48.07 DISPLAY M 946
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DE502097
01  +                     01DE502097
02  +                     01DE502097
03  +                     01DE502097
04  +A                    01DE502097
05  +A                    01DE502097
06  +I                    01DE502097
07  +I                    01DE502097
08  +                     01DE502097
09  +                     01DE502097
0A  +                     01DE502097
0B  +                     01DE502097
0C  +                     01DE502097
0D  -
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.714.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-MANAGED
N NOT AVAILABLE
```

Figure 5-25 Configure offline an excessive processor

Deactivation is done by using the same Temporary Upgrade panel as was used for an activation. Select the On/Off CoD active record and then click **Remove Processor**; see Figure 5-26.

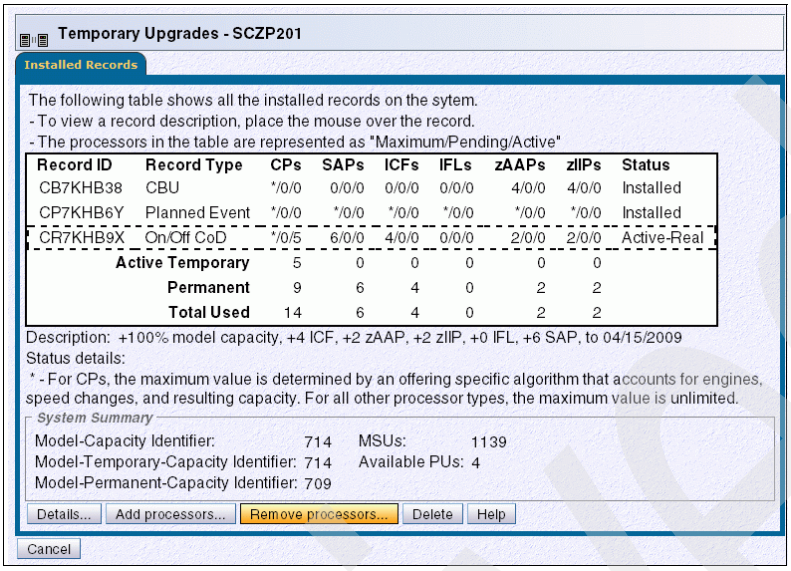

Figure 5-26 Temporary Upgrades - deactivate temporary capacity

The Change Activation Levels panel is displayed. Enter the values either manually or by clicking **Undo**, which returns the configuration back to the level it was before the record was activated. Figure 5-27 shows values that deactivate the record. No specialty engines were used and the capacity setting is set back to 709 from 714.

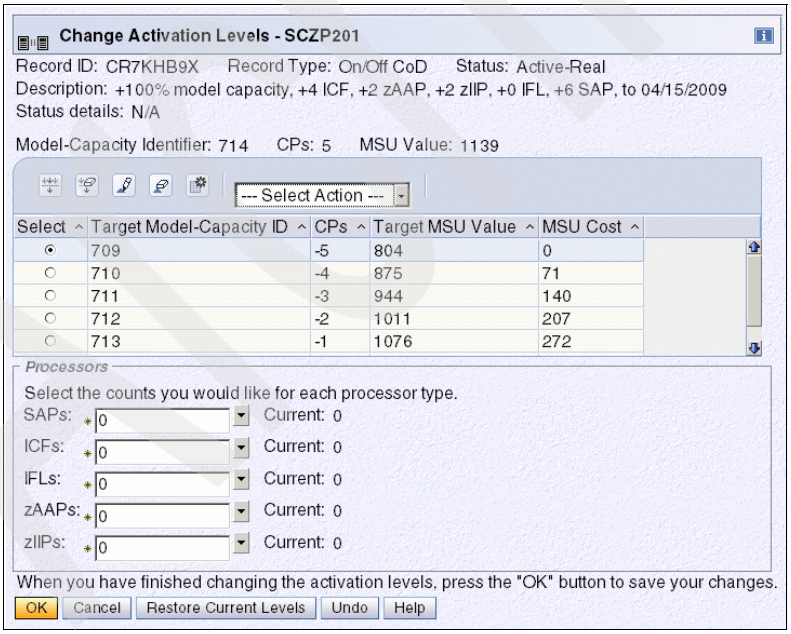

Figure 5-27 Change Activation Levels

Figure 5-28 on page 144 shows current and new configurations. There were originally five temporary CPs, now there will be none. Select **OK** to confirm. The temporary upgrade is deactivated.

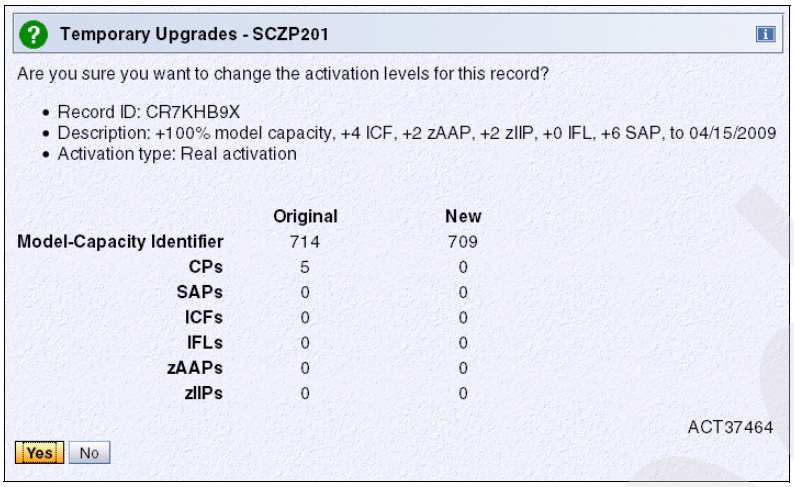

Figure 5-28 Temporary Upgrades - message ACT37464

Every time a change in activation levels is performed, a warning message is issued. See Figure 5-29.

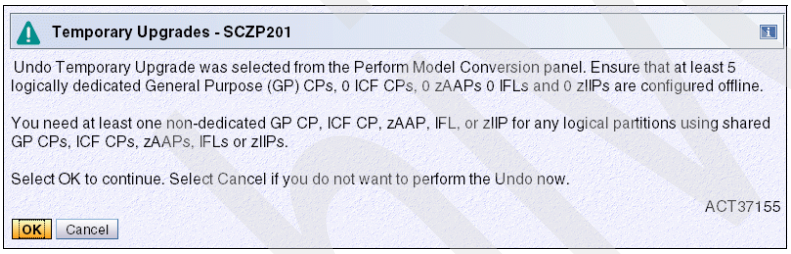

Figure 5-29 Temporary Upgrades - message ACT37155

The message serves as a reminder to check whether:

- Logical processors or specialty engines do not have to be configured offline.
- Logical partitions can be deactivated.
- Possible weight changes can be done.

Pay attention to this information and configure offline resources to avoid having more resources online than are physically available after the deactivation, which can cause a performance problem.

Perform a final check to ensure that the temporary resources are deactivated by reviewing the Temporary Upgrade panel. The Model Identifiers and MSU value are back to their original, and the On/Off CoD record is not active; see Figure 5-30 on page 145.

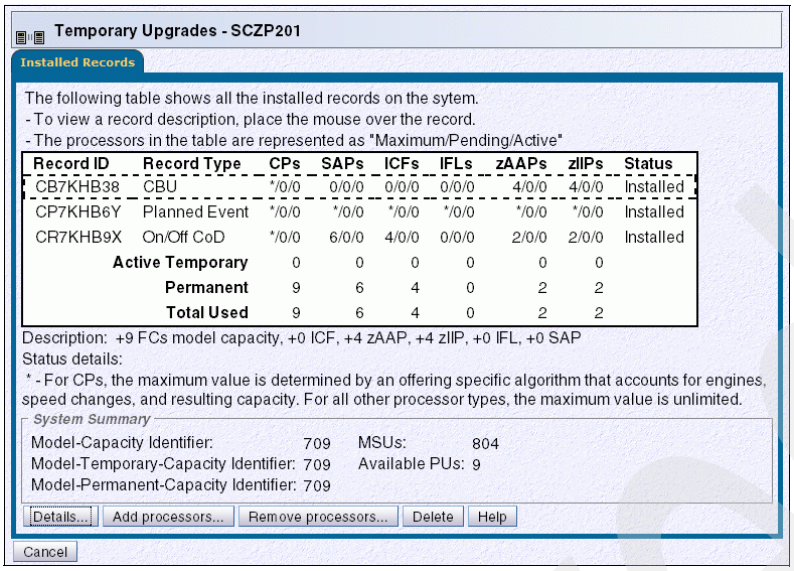

Figure 5-30 Temporary Upgrades

A final check also shows that z/OS recognizes only nine CPs, and that 709 is the current capacity setting; see Figure 5-31.

```
D M=CPU
IEE174I 15.53.39 DISPLAY M 948
PROCESSOR STATUS
ID  CPU                SERIAL
00  +                   01DE502097
01  +                   01DE502097
02  +                   01DE502097
03  +                   01DE502097
04  +A                  01DE502097
05  +A                  01DE502097
06  +I                  01DE502097
07  +I                  01DE502097
08  +                   01DE502097
09  +                   01DE502097
0A  +                   01DE502097
0B  +                   01DE502097
0C  +                   01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1
```

Figure 5-31 Display command after the deactivation of temporary capacity

6

On/Off Capacity on Demand with prepaid upgrades

This chapter shows an example On/Off Capacity on Demand with prepaid capacity tokens usage. All planning and monitoring tasks as described in Chapter 5, “On/Off Capacity on Demand” on page 123 apply also to this chapter. The sole purpose of this chapter is to show how capacity tokens work in reality because they provide a whole new dimension to the On/Off CoD offering.

This chapter discusses the following topics:

6.1 Preparation

This section describes the steps necessary for when an On/Off CoD record is prepared to handle a planned peak workload, such end-of-month processing. The current configuration is a z10 BC with a capacity setting of R02 active and V03 high watermark, with one ICF and two IFLs. During a peak workload, it is expected to need V04 for two days, one more IFL for two days and two more IFLs for one day.

6.1.1 Review planned use of offering

With On/Off CoD it is possible to add up to twice the capacity of the CPC purchased capacity. The CPC is a 2098 E10 with a permanent capacity setting R02, but purchased capacity is V03. Thus the maximum temporary capacity setting available using On/Off CoD is Z03, which is the highest capacity setting lower than double the capacity of V03.

6.1.2 Order, retrieve, and install a record

This section contains a brief description of steps performed when an On/Off CoD record with prepaid capacity tokens is ordered from a Resource Link.

Order the record

The On/Off CoD order is placed through Resource Link, Customer Initiated Upgrade (CIU). Figure 6-1 shows the first step in the ordering process. The On/Off CoD configuration is entered. Note that there is no Replenishment due date field in the panel compared to postpaid On/Off CoD record (shown in Figure 5-1 on page 125). Limits defined in this panel are the upper limits of how many resources can be activated by this record at its activation time. At activation time, you choose the amount of resources to activate; you may change it later within the record authorization space as described in “On/Off CoD authorization space” on page 33. We recommend using the defaults in this panel, which allows activation of the maximum amount of resources.

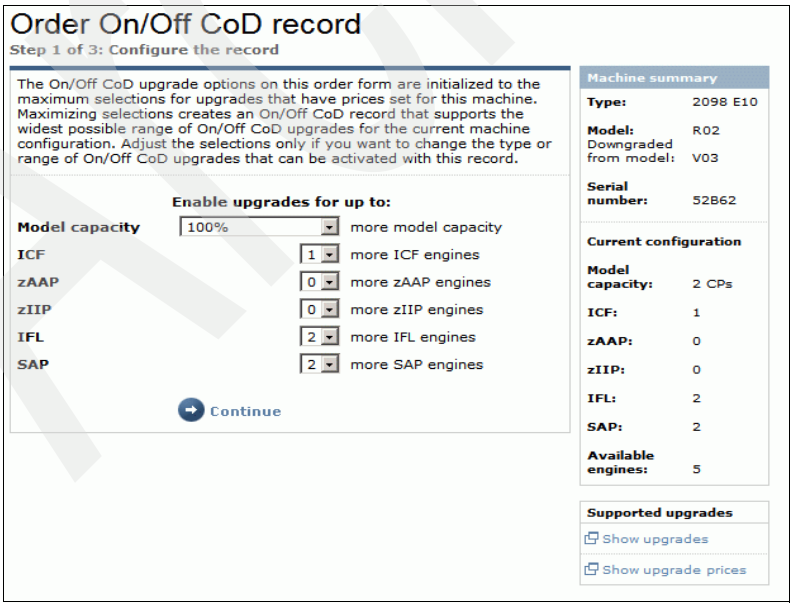

Figure 6-1 Configuration of the On/Off CoD order

The second step in the order is to define the number of tokens to be purchased. Figure 6-2 shows the following order:

- Upgrade to V04 for two days, which results in 78 MSU tokens because the difference between V03 and V04 is 39 MSUs
- Two more IFLs for two days, which results in four IFL tokens

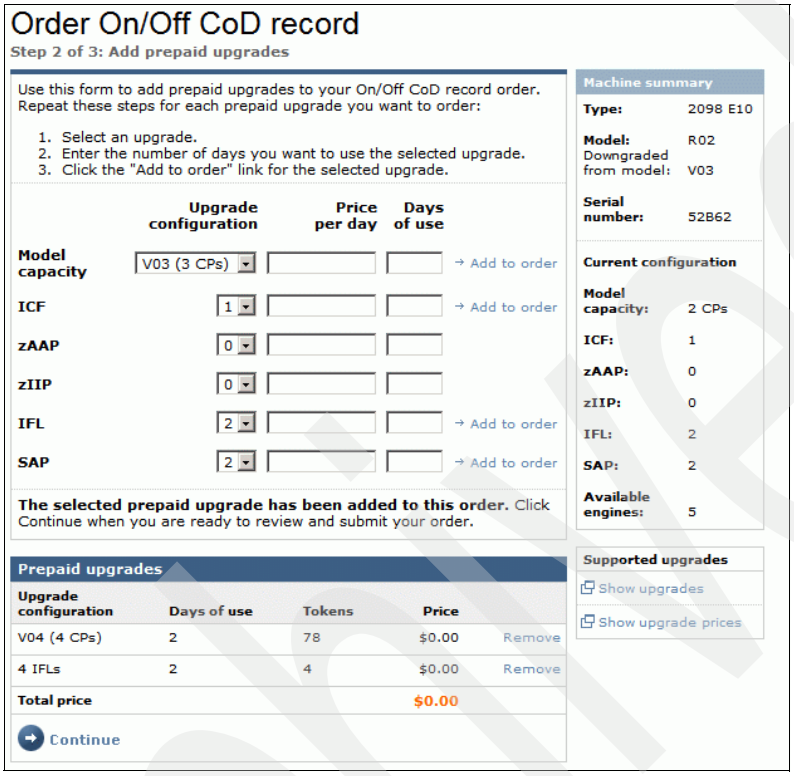

Figure 6-2 Add prepaid upgrades to On/Off CoD order

After the order is reviewed and submitted, it must be approved by an approver if one is defined. After the order is approved and ready to be retrieved, a notification e-mail is sent, indicating that a record is ready for a download. Example 6-1 shows an excerpt from a notification e-mail.

Example 6-1 Notification e-mail

```
Subject: Download ready: Prepaid On/Off CoD order CX7KVRZF for machine (2098 --
52B62) is ready to be downloaded and installed

Customer number: ITSO001

Prepaid On/Off CoD order CX7KVRZF for machine 2098 -- 52B62 has been processed and
is ready to be downloaded to the customer's machine. Please contact the
appropriate person to retrieve this order from the Hardware Management Console
(HMC).

Order Create Date: 10/29/2008
Record number: CR7KVRZF
Order number: CX7KVRZF

Order Confirmation Information
```

```
Model capacity:   100% more model capacity
ICF:              1 more ICF engines
zAAP:             0 more zAAP engines
zIIP:             0 more zIIP engines
IFL:              2 more IFL engines
SAP:              2 more SAP engines

Description:      Prepaid +100% model capacity, +1 ICF, +0 zAAP, +0 zIIP, +2 IFL,
+2 SAP

Note: Prepaid upgrades are represented on On/Off CoD records by tokens.
Use the link at the bottom of this email to view the tokens associated with this
order.
```

Retrieve and install the record

A record retrieval and installation does not differ from a postpaid On/Off CoD record, which is described in 5.1.3, “Order, retrieve, and install a record” on page 124.

6.2 Usage

This section describes the usage of an ordered record in a span of four days. It discusses how tokens are consumed after each day.

6.2.1 Day one

Figure 6-3 shows that a record is installed. Server capacity setting is R02 and five PUs are available.

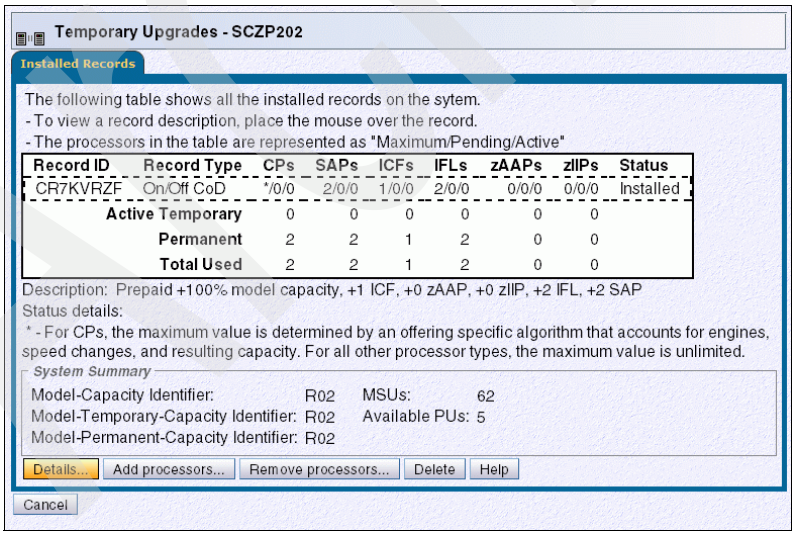

Figure 6-3 Temporary Upgrades

The Record Details panel in Figure 6-4 shows the following information:

- Defined resources for a record. You may add two SAPs, one ICF, and two IFLs using this record.
- Capacity tokens. There are 78 MSU tokens and four IFL tokens. Because no other processor type tokens are defined, activating processors, other than CP and IFL, is not possible.
- No expiration date. A prepaid On/Off CoD record does not expire.
- Record description. This shows the record is prepaid and what resources are defined.

Note: Prepaid On/Off CoD record description says, by default, that it is prepaid. We recommend keeping this information in a record description, because the only other way to distinguish between a prepaid and postpaid record is by checking the expiration date.

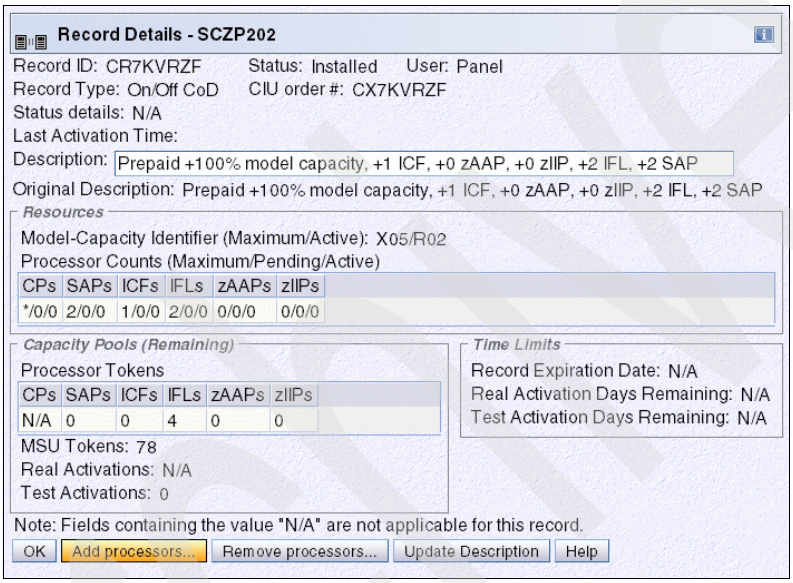

Figure 6-4 Record Details

The Change Activation Levels panel, shown in Figure 6-5 on page 152, is where to activate resources. All resources defined in the record are shown. The presence of tokens is checked only when an attempt to activate the resources is made.

The MSU Cost column shows how many MSU tokens per day will be used for each capacity setting, if activated. If the value is zero, the capacity setting is below or equal to high watermark capacity setting and does not consume tokens. In this case, U03 was chosen and one IFL was added to the configuration.

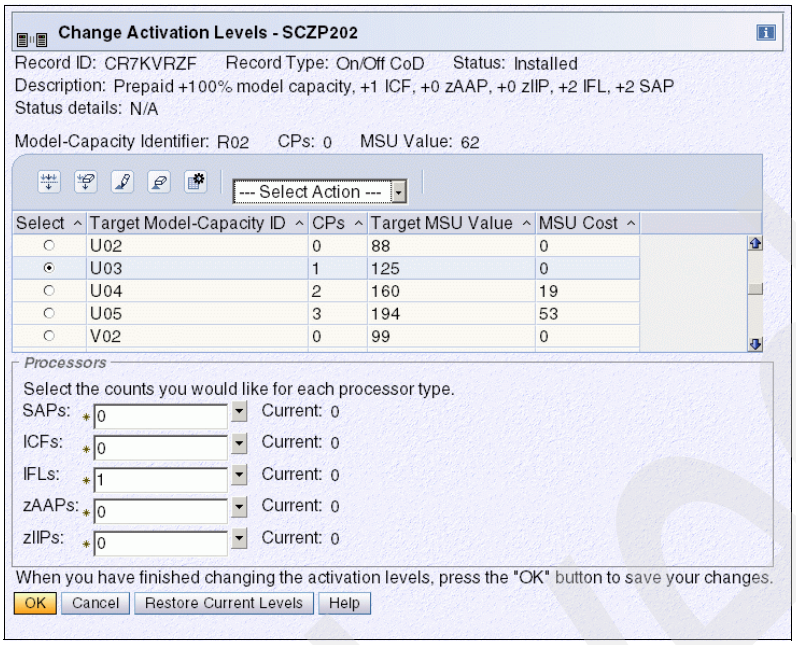

Figure 6-5 Change Activation Levels panel

Figure 6-6 shows that a record is active, Model-Capacity Identifier is U03, one temporary IFL is active, and a number of available PUs was decreased to three.

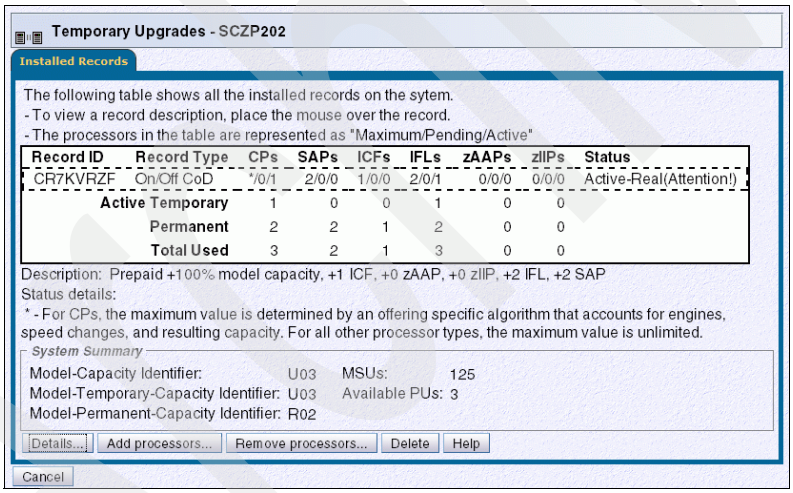

Figure 6-6 Temporary Upgrades panel

Figure 6-7 shows a hardware message, which appears after a record is activated. Based on current IFL tokens consumption, the message warns that tokens will be depleted in four days.

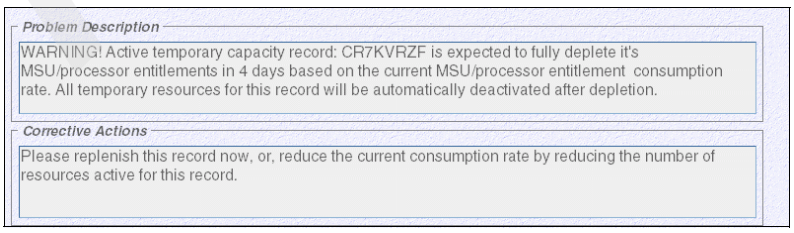

Figure 6-7 Hardware message

6.2.2 Day two

Hardware message appears after 24 hours, as shown in Figure 6-8. As expected, it indicates the record will be depleted in three days.

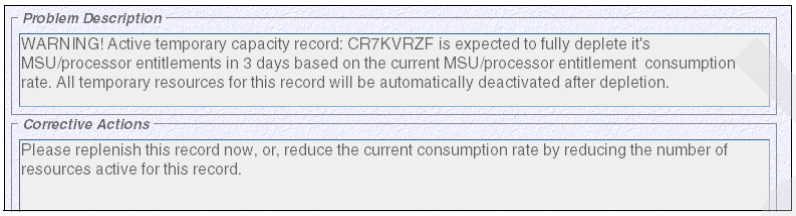

Figure 6-8 Hardware message

The Record Details panel, shown in Figure 6-9, confirms that after 24 hours one IFL token was consumed and three IFL tokens remain. It also confirms that no MSU token was consumed and all 78 MSU tokens are still available.

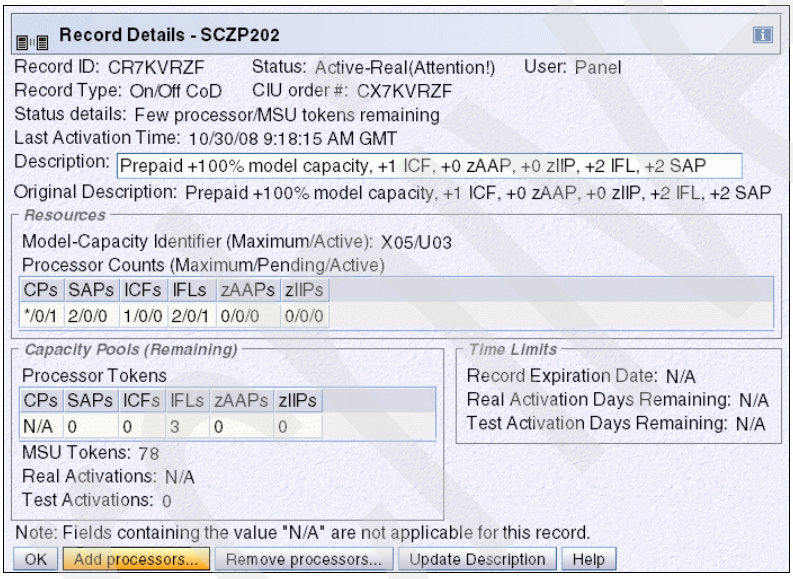

Figure 6-9 Record details panel

Because it was determined that V04 would not be enough for the workload, W04 has to be activated, as shown in Figure 6-10 on page 154. Select anything that adds less than the 79 MSU Cost.

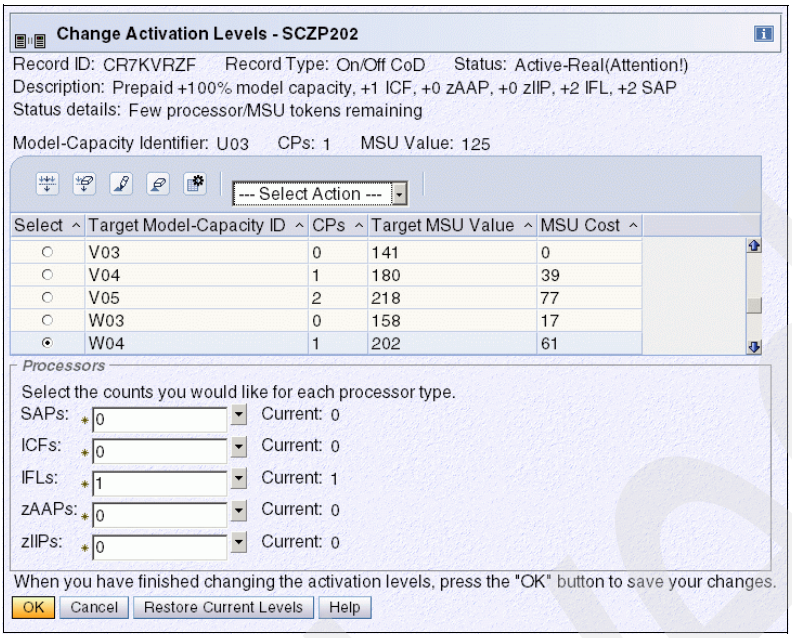

Figure 6-10 Change Activation Levels panel

After the change was activated, the hardware message shown in Figure 6-11 is issued. It indicates that the record will be depleted in one day rather than in three days. The reason is that W04 will consume 61 out of 78 MSU tokens and not enough tokens remain for the next day. Replenishing the record or deactivating W04 just before end of the billing window is possible. Otherwise the entire record will be deactivated automatically as shown in the next section.

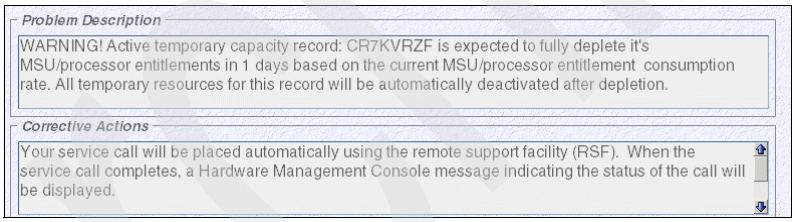

Figure 6-11 Hardware message

6.2.3 Day three

At the beginning of the next billing window, two hardware messages are issued.

The first hardware message, shown in Figure 6-12, indicates that the record tokens were depleted.

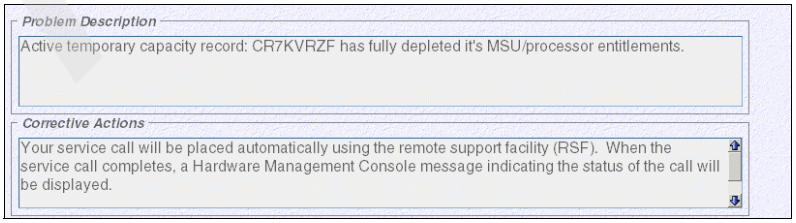

Figure 6-12 Hardware message

The second hardware message, in Figure 6-13, indicates that resources for the record were automatically deactivated.

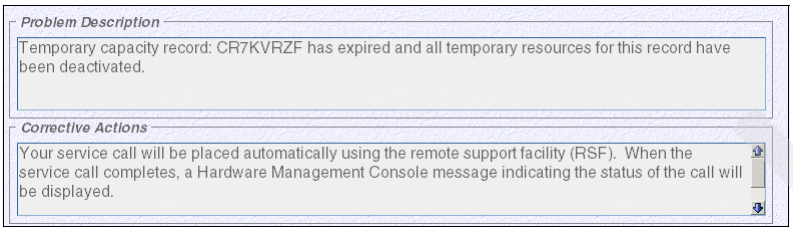

Figure 6-13 Hardware message

Checking the Temporary Upgrades panel confirms that the On/Off CoD record was fully deactivated, as shown in Figure 6-14.

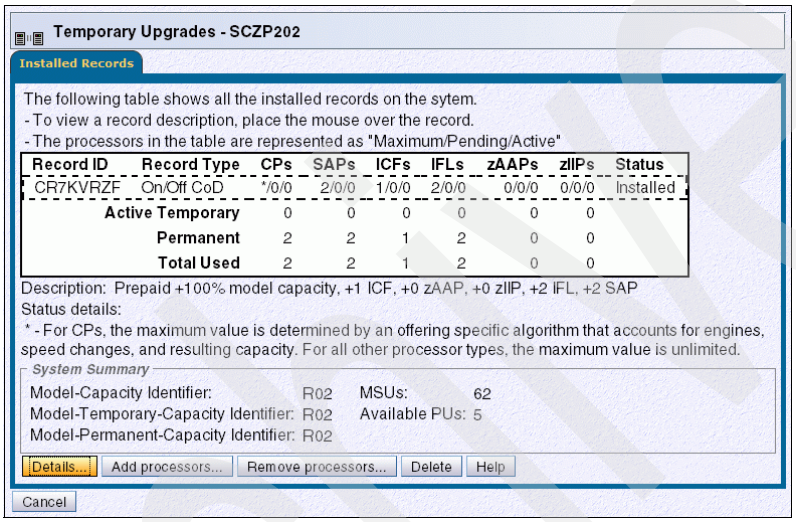

Figure 6-14 Temporary Upgrades panel

Figure 6-16 on page 156 shows two IFL and 17 MSU tokens remain in the record.

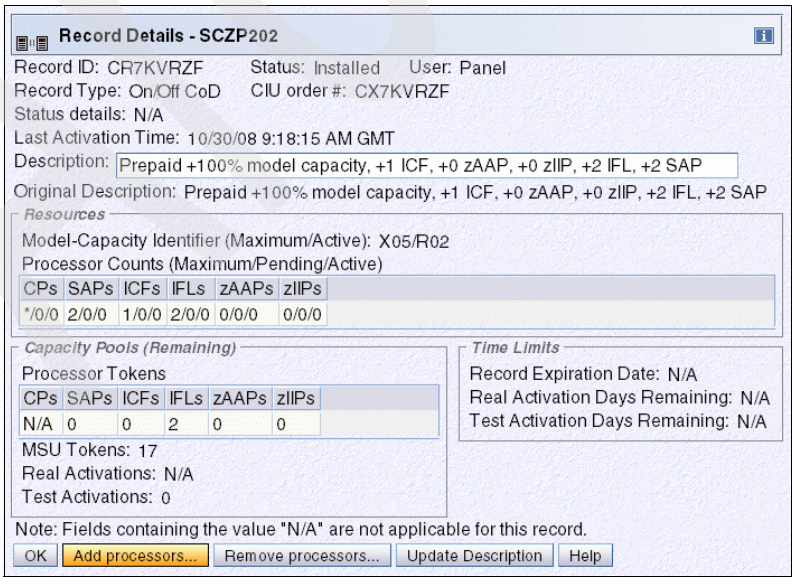

Figure 6-15 Record details panel

For this last day of active resources (day three), two IFLs will be activated and will consume the last two IFL tokens. V04 cannot be activated, because it requires 39 MSU tokens and only 17 remain. W03 will be activated together with IFLs, as shown in Figure 6-16.

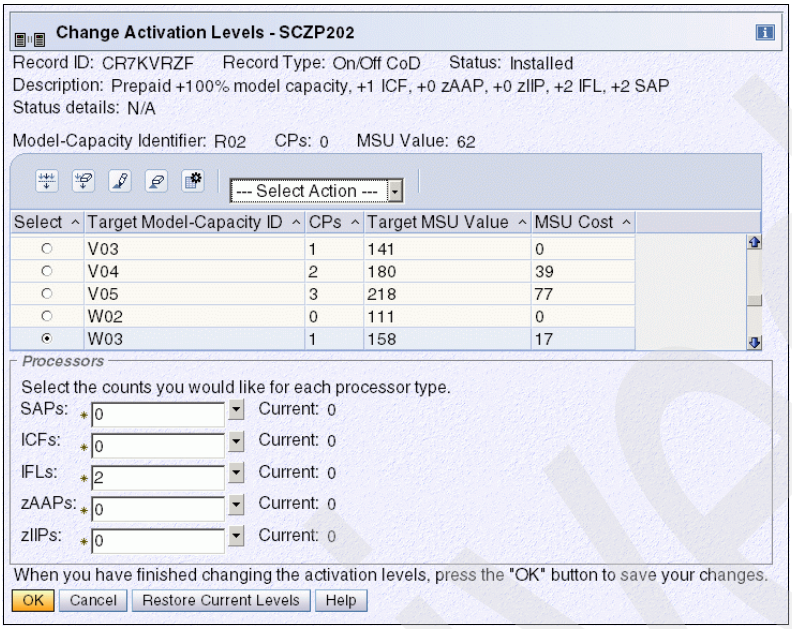

Figure 6-16 Change Activation Levels

Figure 6-17 shows that W03 and two temporary IFLs are active.

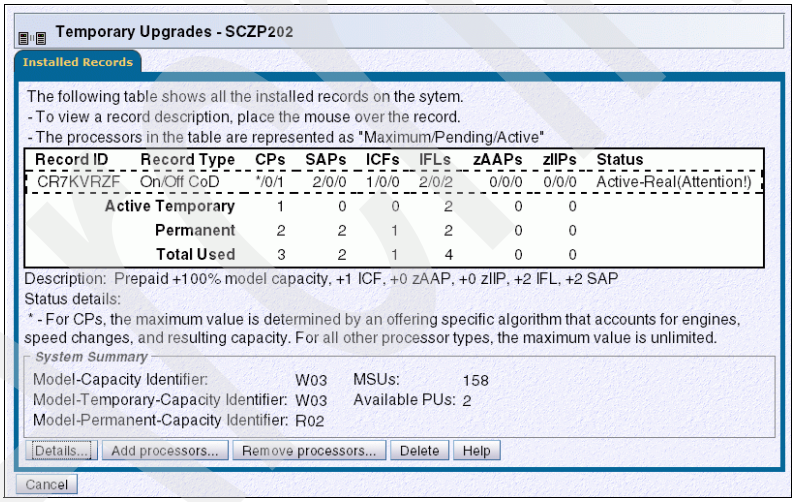

Figure 6-17 Temporary Upgrades panel

After activation, a hardware message appears, as shown in Figure 6-18 on page 157. The message warns that a record will be depleted in one day.

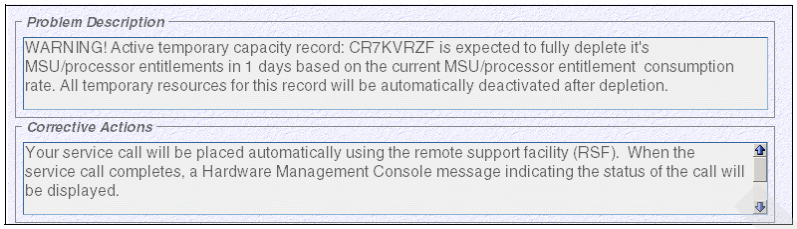

Figure 6-18 Hardware message

6.2.4 Day four

At the beginning of the next billing window, two hardware messages are issued.

The first hardware message, shown in Figure 6-19, indicates that the record tokens were depleted.

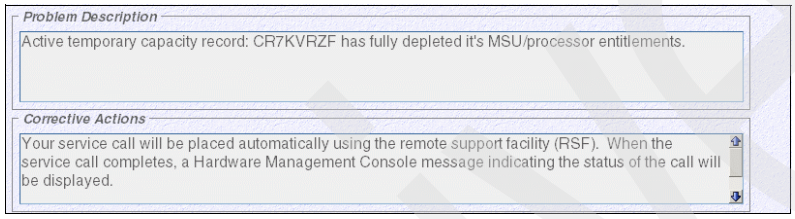

Figure 6-19 Hardware message

The second hardware message, in Figure 6-20, indicates that resources for the record were automatically deactivated.

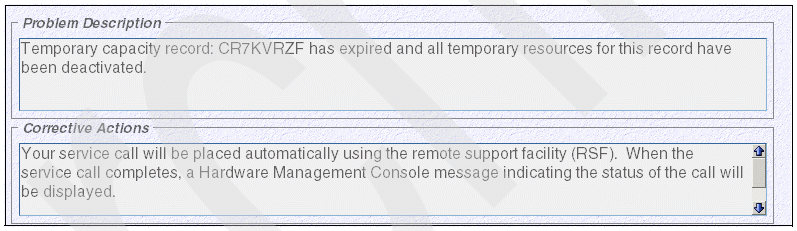

Figure 6-20 Hardware message

Checking the Temporary Upgrades panel confirms that the On/Off CoD record was fully deactivated, as shown in Figure 6-21 on page 158.

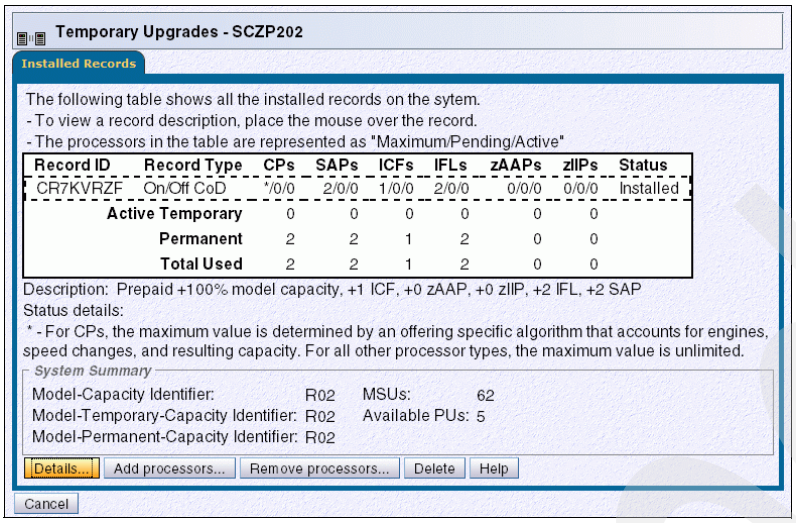

Figure 6-21 Temporary Upgrades panel

Figure 6-22 shows no tokens remain in any token pool. Only capacity setting lower than High Water Mark (HWM) can be activated; the remainder of the record cannot be activated until it is replenished.

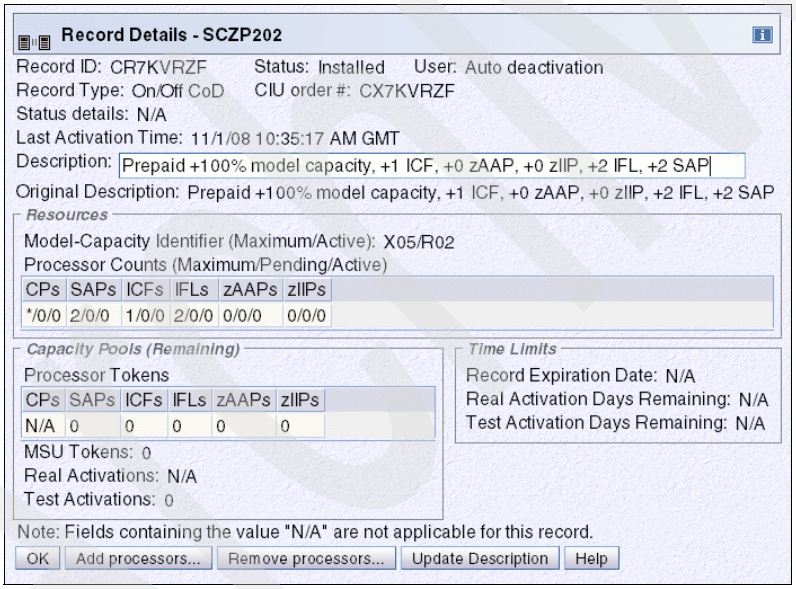

Figure 6-22 Record details panel

7

z/OS Capacity Provisioning

This chapter describes On/Off Capacity on Demand activation by using z/OS Capacity Provisioning Manager (CPM) in confirmation mode. CPM can run in one of four modes:

- Manual
- Analysis
- Confirmation
- Autonomic

We show confirmation mode because it exploits CPM functions, but the last decision is made by the system administrator before adding temporary capacity. A recommendation is that you start by using confirmation mode, and later switch to autonomic mode.

In this chapter, we discuss the components that are required to set up the environment, and explain the specific setup required for CPM. Explaining how to install and customize these components is beyond the scope of this book; that information is documented in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

This chapter discusses the following topics:

7.1 Environment

The scope of z/OS Capacity Provisioning is to manage server capacity delivered by general purpose CPs, and by zAAPs and zIIPs. This scenario shows capacity provisioning of CPs and zAAPs, provided by CPM in confirmation mode.

The CPM can interrogate the On/Off CoD records installed on the server, and understand which resources are available for provisioning, and the content of capacity tokens if they are present in the On/Off CoD records. If CPM is provisioning resources from offering records containing capacity tokens, CPM checks for each resource being provisioned, that the capacity token for each resource contains enough tokens to cover the next billing window (24 hours). If any capacity token does not meet this criteria, no resources from the On/Off CoD record will be activated.

In addition, care must be taken to avoid that the system itself deactivates resources that CPM has provisioned. When a capacity token is being fully consumed, the system will deactivate all resources provisioned from that offering record, even if the resources were provisioned by CPM, and even if other capacity tokens have resources remaining. When a capacity token approaches complete consumption, the system issues error messages, beginning 5 days before the token is completely consumed. A strong recommendation is to capture those messages and put routines in place to handle the situation. Either prepare for the deactivation of all the resources from the record, or use Resource Link to replenish the resource token. After an eventual deactivation, CPM can be used to reactivate resources from the deactivated offering record, provided the resources have enough tokens for a new billing window.

Infrastructure for CPM scenario

Figure 7-1 on page 161 shows the hardware and software environment that we used for our scenario.

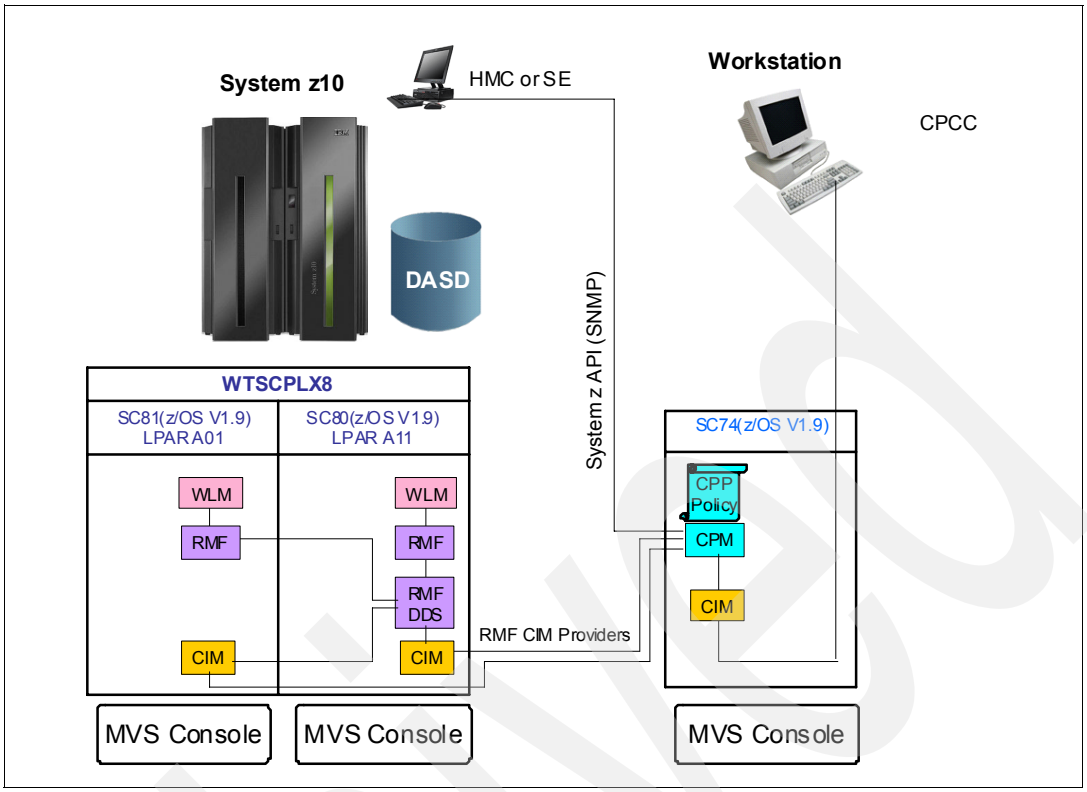

Figure 7-1 On/Off CoD scenario environment

The environment consisted of the following items:

- WTSCPLX8 consisted of two z/OS V1R9 images SC80 and SC81, both running on a System z10 EC server.
- CPM was located in SC74, which was a z/OS V1R9 image running on a System z9 EC.
- Both SC80 and SC81 were observed by CPM.
- System z10 was monitored by CPM.
- WLM was active sysplex-wide for workload management in WTSCPLX8.
- RMF was active in all z/OS images.
- RMF DDS server was active in high availability mode, and ran in SC80 or SC81.
- RMF was (and must be) configured to record appropriate data. RMF DDS must be configured so that CPM can retrieve data. Our environment had no specific setting.
- CIM servers were active in SC80 and SC81. CPM retrieves RMF data from RMFDDS through the CIM server. The CIM server contains the IP address of RMFDDS, so it can run anywhere. We decided to run it on the observed system.
- Another CIM server was active in SC74, which was in the same z/OS image as CPM. It was used for the CPCC-to-CPM connection.
- The security server was active to provide sysplex-wide security and access authorization control. There were many requirements from different components for security settings, as described in their respective documentation. In our environment, we used RACF®.
- Capacity Provisioning Control Center (CPCC) is the front-end tool to administer the Capacity Provisioning Domain, which includes CP domain configuration and CP policies. CPCC connects to CPM through the CIM server.

The procedure for CPCC installation is documented in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

HMC record definition and activation is discussed in 5.1, "Preparation" on page 124. We do not discuss On/Off CoD record installation in this chapter; the record must be installed and ready to use.

z/OS WLM Policy

CPM is intended to help critical workloads when they suffer. For this reason, choosing the correct service class period to monitor is critical. In this scenario, the performance of service class JAVAWKLD, with an execution velocity 70 goal, is monitored. It is a Java workload, and may be dispatched to zAAPs and CPs.

Having realistic goals is important. If the goal is too aggressive, its performance index (PI) might be high even when a significant amount of resources are available, and overall WLM management can be adversely affected.

7.2 Preparation

This section describes connections, and provisioning configuration and policies.

Set up the required connections

The following connections are required for the z/OS Capacity Provisioning On/Off CoD scenario by using CPM in confirmation mode.

- CPM connection to RMF DDS through the CIM server
- CPCC connection to CPM through the CIM server
- HMC and CPM connection

Be aware that these connections might require firewall changes in certain environments.

CPM connection to RMF DDS through the CIM server

A CIM server must be on each observed system, and the server must recognize an IP address of RMF DDS. In addition to the usual customization, as documented in *CIM Users's Guide,* SA33-7998, there is only one further customization necessary for CPM in CIM configuration, which is that RMF_CIM_HOST must point to the RMF DDS IP address, and RMF_CIM_PORT specifies port on which RMF DDS listens.

In our environment, we configured RMF DDS in high availability mode; for details, see "Setting up RMF DDS in high availability mode" on page 205. The dynamic virtual IP address for RMF DDS is set as:

```
RMF_CIM_HOST=9.12.5.26
RMF_CIM_PORT=8803
```

CPCC connection to CPM through the CIM server

A CIM server must be on the same z/OS image where CPM runs so that CPCC can connect and install definitions to CPM. Here we shown part of the NETSTAT output on the SC74 system:

```
CFZCIM   00001411 0.0.0.0..5988          0.0.0.0..0             Listen
CFZCIM   0000140D 127.0.0.1..1407        0.0.0.0..0             Listen
```

HMC and CPM connection

CPM communicates to HMCs through System z APIs (SNMP). The definitions in CPM, in HMC, and in the Support Element are described in the following list:

- Definitions in CPM

  CPM must be installed and configured as described in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299. CPCC is used to set up CPM operation. To adapt to our environment, the following three parameters are changed:

  – Domain name in started task JCL

    ```
    //DOMAIN   SET DOMAIN=DOMAIN1
    ```

  – HMC IP address in the CPM configuration file

    ```
    Topology.Address=9.12.6.46
    ```

  – SNMP community name in the CPM configuration file

    ```
    Topology.Community=public
    ```

- Definitions in HMC

  CPM communicates with HMC over IP by using SNMP. For this reason, SNMP must be configured and enabled on HMC.

  Figure 7-2 shows the Customize API Settings dialog, which is used to add SNMP information.

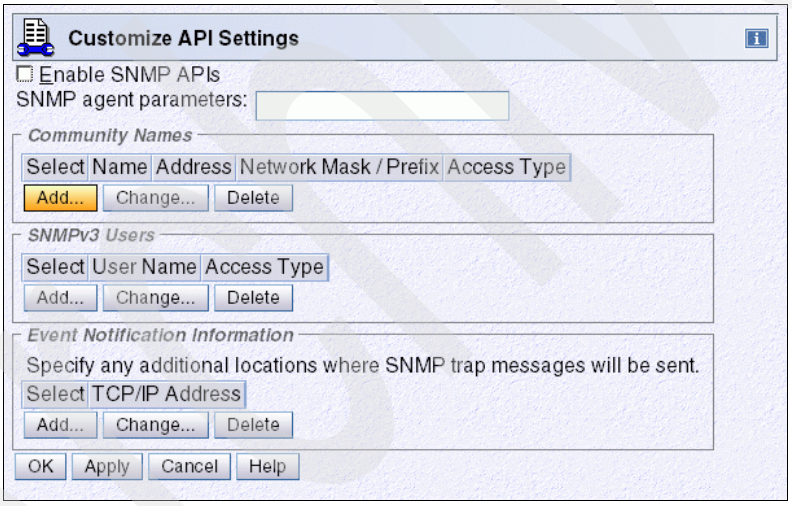

Figure 7-2 Customize API Settings dialog

Figure 7-3 on page 164 shows the dialog where SNMP information is entered. Community name is set to `public` and an IP address of `SC74`, where CPM is located. Network mask is set to `255.255.255.255` to indicate a host, not a network. CPM requires read/write SNMP access to be able to activate On/Off Capacity on Demand records.

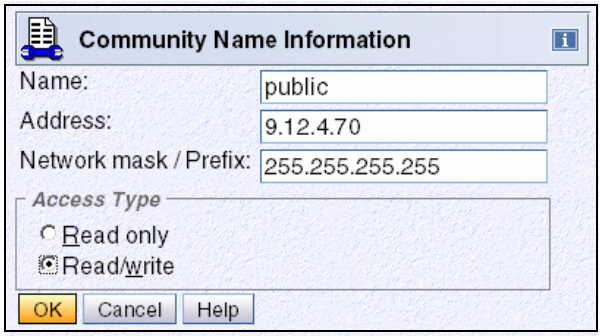

Figure 7-3 SNMP community name

Click **OK**. Finally, SNMP API is enabled, as shown in Figure 7-4.

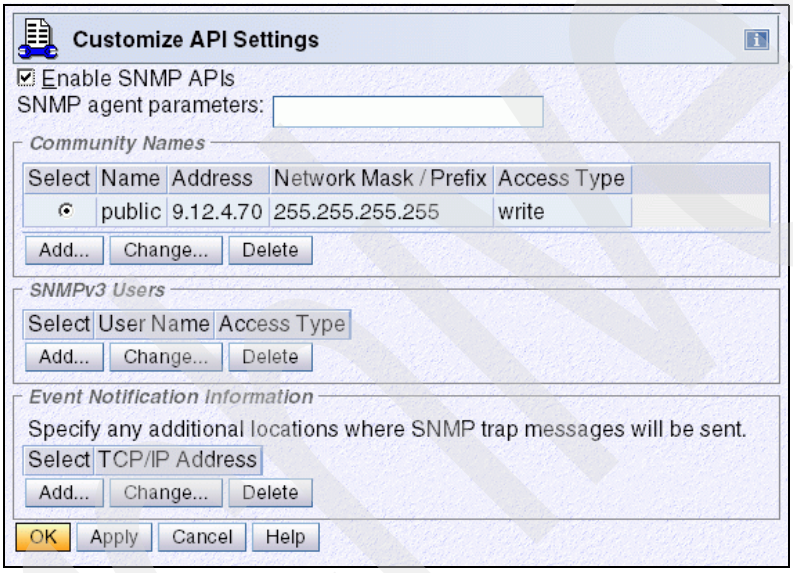

Figure 7-4 SNMP API enabled

- Definitions in the Support Element

  Note the following settings, as shown in Figure 7-5 on page 165:

  – Allow capacity change API requests

    To allow changes of capacity through API, the SE must allow capacity change API requests. It is not necessary to set SNMP values in the SE, because SNMP requests go through HMC in our environment.

  – Enable SNMP API and Community Names

    If the SE is used directly from CPM, then enabling SNMP API and Community Names would be required on the SE.

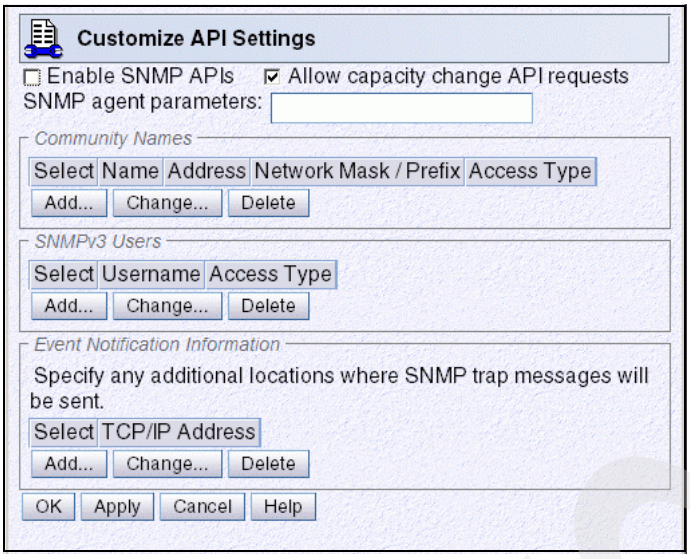

Figure 7-5 SE API settings

CPCC connection to CPM

CPCC is a workstation application that is used to define configurations and policies to CPM. It is necessary to connect to CPM to install any of them. To establish a connection between CPCC and CPM, CPCC needs an IP address and the listening port of the CIM server on the z/OS image, where CPM is running.

For a connection to CPM, user and password have to be entered in a login window. User is a TSO user ID with RACF USE access to groups CPOQUERY and CPOCTRL, to be able to display and modify CPM data.

For a detailed description of using the CPCC, see *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

Define provisioning configuration and policy

The following tasks summarize the configuration and policy used in this scenario to set up the provisioning domain:

- Define the provisioning configuration:
  - Servers (CPCs) managed
  - z/OS systems observed
- Define the provisioning policy:
  - Policy timeline
  - Maximum provisioning scope
  - Rules

Define provisioning configuration

The configuration keeps the definition of servers that CPM manages and a list of the z/OS images that are monitored.

To define a z/OS image to the configuration, enter its image name and sysplex name. These names must match system symbols in the z/OS image. In addition, the IP address and port number of the CIM server that is running on the observed system are required. The CIM server must have a connection to RMF DDS, which provides workload measurement data. z/OS images SC80 and SC81 are observed, as shown in Figure 7-6 on page 166.

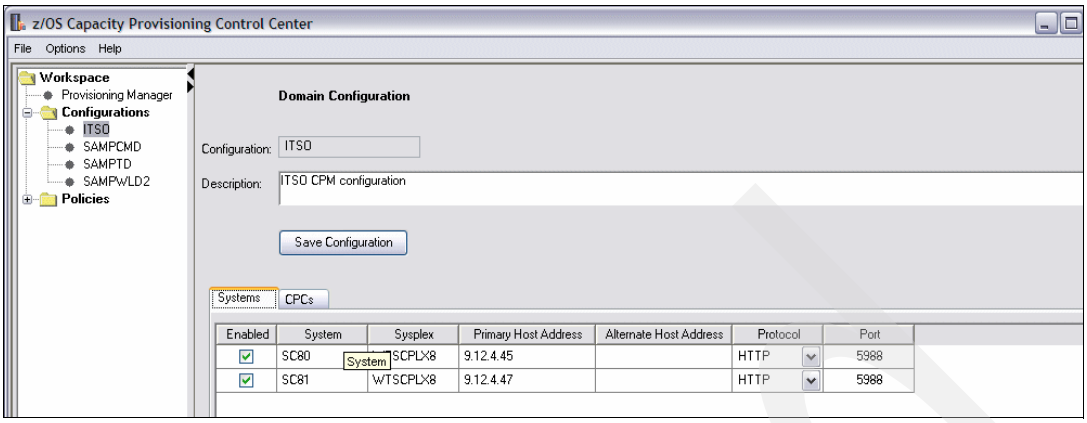

Figure 7-6 Systems in the provisioning domain

You must also enter the CPC name of the managed server, and it must match its name in HMC. Figure 7-7 shows the definition of SCZP201 in CPCC.

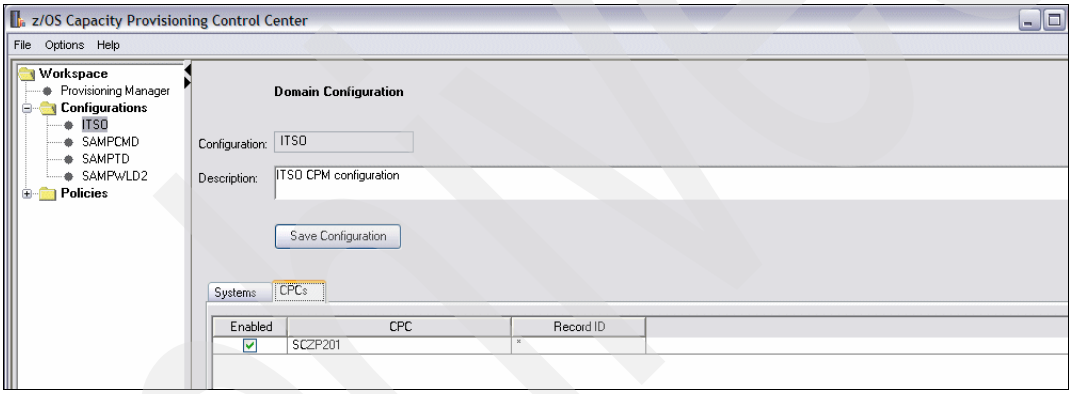

Figure 7-7 Server managed by CPM

After you save the configuration, install it in CPM. Right-click the name of the configuration (in this case, ITSO) and select Install. Before installing the domain configuration, a connection has to be established, which you do using the `Provisioning Manager` branch from the tree panel on the left side of the window.

After installing, the configuration can be activated in CPM from the z/OS console, as shown in Figure 7-8.

```
F CPOSERV,APPL=SET DOMAIN CFG=ITSO
CPO1041I Domain configuration ITSO successfully activated
CPO3806I The system at address 9.12.4.47 is SC81 in sysplex WTSCPLX8
CPO3813I The system at address 9.12.4.47 is running on CPC SCZP201
CPO3806I The system at address 9.12.4.45 is SC80 in sysplex WTSCPLX8
CPO3813I The system at address 9.12.4.45 is running on CPC SCZP201
CPO3880I System SC81 in sysplex WTSCPLX8 is now monitored
CPO3880I System SC80 in sysplex WTSCPLX8 is now monitored
```

Figure 7-8 Change of CPM configuration

Although more than one configuration may be defined in CPM, only one configuration can be active. To obtain a list of all defined configurations, or to receive a report of active

configuration details, use F CPOSERV,APPL=L C and F CPOSERV,APPL=R C, respectively, from the z/OS console. An example is shown in Figure 7-9.

```
F CPOSERV,APPL=R C
CPO1010I Configuration report generated at 02/27/2008 13:37:06
Domain configuration ITSO for domain DOMAIN1 is enabled
CPC SCZP201 with record * is enabled (default enabled)
 CPC is matched with serial 00002001DE50 since 02/22/2008 20:10:51
 Hardware is of type 2097 with model E26
 Current model is 711 with 944 MSUs, 2 zAAPs, and 4 zIIPs
 Permanent model is 710 with 875 MSUs
 Active record ID is CR7C3TP9
 Hardware has 5 spare processors
 Activation limits are 2 zAAPs, and 2 zIIPs
 Active resources GP/zAAP/zIIP 0(0/0)/0/0
System SC80 in sysplex WTSCPLX8 is enabled (default enabled)
 Primary host address: 9.12.4.45
 Protocol: HTTP, port: 5988
 The system at primary host address is observed
  This system is available since 02/25/2008 17:11:31
  This system is running on the CPC SCZP201
  WLM service definition: wlmdef1, active policy: POLICY1
System SC81 in sysplex WTSCPLX8 is enabled (default enabled)
 Primary host address: 9.12.4.47
 Protocol: HTTP, port: 5988
 The system at primary host address is observed
  This system is available since 02/25/2008 17:06:33
  This system is running on the CPC SCZP201
  WLM service definition: wlmdef1, active policy: POLICY1
```

Figure 7-9 Sample Report Configuration output

Define provisioning policy

A Capacity Provisioning Policy consists of the maximum provisioning scope and one or more rules. In this example, we show a policy with one rule.

The maximum provisioning scope limits the capacity that CPM is allowed to provision. Figure 7-10 shows a provisioning scope of 500 MSU and five zAAP processors.

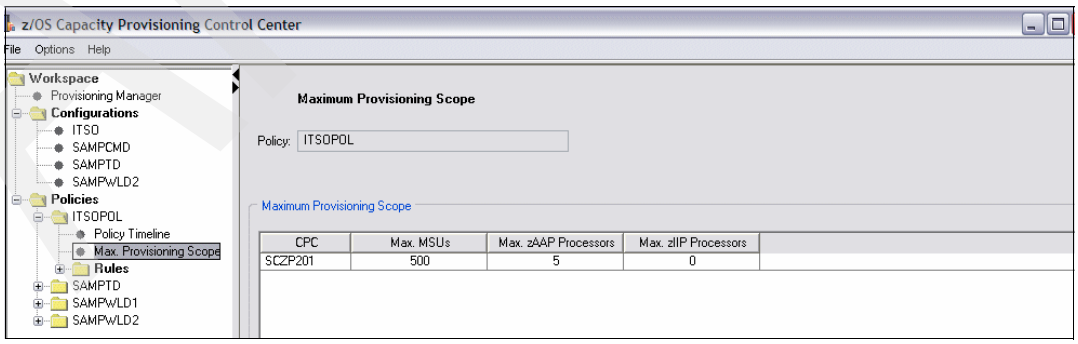

Figure 7-10 Maximum provisioning scope

Rules specify conditions under which additional capacity can be provisioned. Each rule can specify limits in time and maximum capacity.

Figure 7-11 shows a capacity limit for a rule, in this case 500 MSUs and two zAAPs.

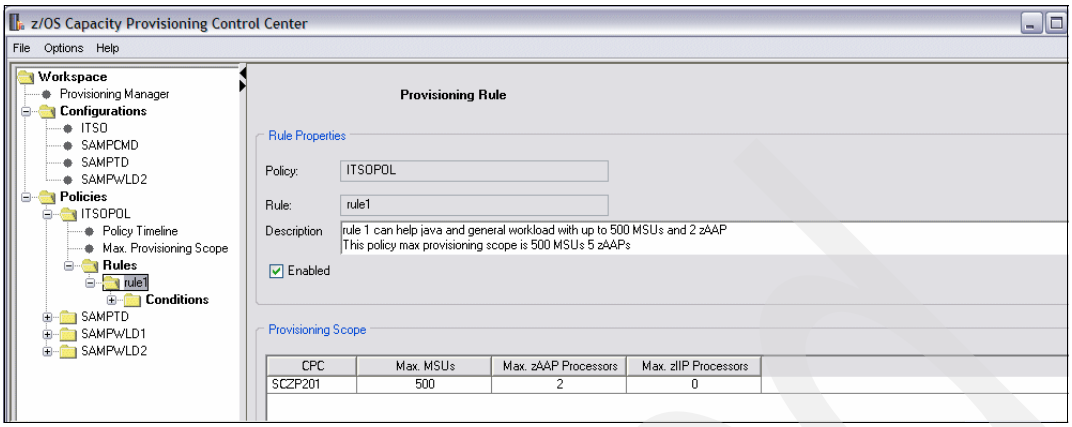

Figure 7-11 Provisioning scope for rule

Each rules has two types of conditions, which are time conditions and workload conditions.

- When only time conditions are specified, the full capacity as specified in the rule is provisioned at specified dates and times.
- When only workload conditions are specified, capacity is provisioned only if specified workload conditions are met.

Using both types of conditions can ensure that capacity is provisioned only during eligible time periods, and only if workload suffers and requires more capacity.

Figure 7-12 shows a time condition definition.

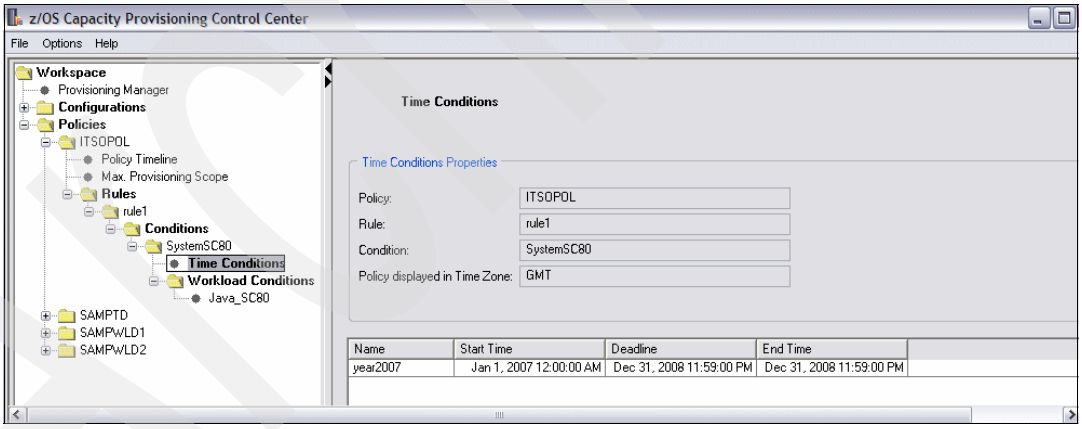

Figure 7-12 Time condition for the rule

Finally, the workload condition is specified, where the periods for the service classes that are considered critical can be used to trigger the provisioning and deprovisioning of a temporary capacity. The metric used to specify when a service class in a particular period requires additional capacity is the performance index (PI).

Figure 7-13 on page 169 shows a definition in which the critical workload is JAVAWKLD. A PI value larger than 2 for a duration of more than 5 minutes is considered unacceptable. This concept is the trigger for capacity provisioning.

The provisioned capacity (if any) should be deprovisioned when the PI for the same workload and period is under 1.5 for more than 5 minutes. Deprovisioning time is also influenced by parameters global to CPM, such as *Planner.MinimumActivationTime*.

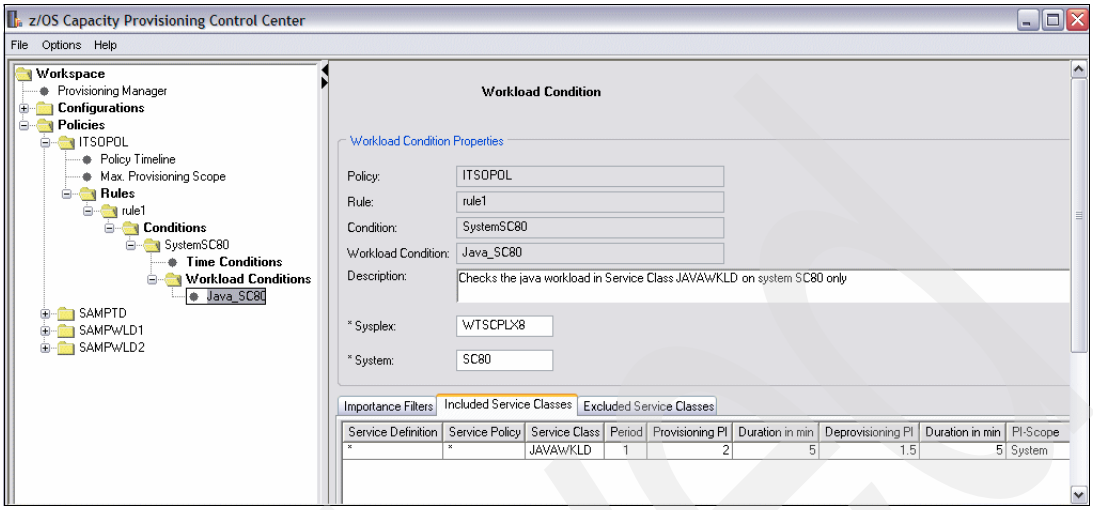

Figure 7-13 Workload condition for the rule

Only the tab of Included Service Classes is shown in Figure 7-13 because it is the only one used in this scenario.

After you complete and save the policy definition, install it in CPM by right-clicking the name of the policy (in this case, ITSOPOL) and clicking **Install**.

At this point, you can activate the policy from the z/OS console as shown in Figure 7-14.

```
F CPOSERV,APPL=SET DOMAIN POL=ITSOPOL
CPO1020I Policy successfully changed to ITSOPOL
```

Figure 7-14 Change of CPM policy

Several policies can be defined, but only one policy can be active. To list policies or show a report, use the following commands from the z/OS console:

`F CPOSERV,APPL=L P` List all available policies.
`F CPOSERV,APPL=R P` Show a report about an active policy.

Figure 7-15 on page 170 shows an example.

```
F CPOSERV,APPL=R P
CPO1005I Policy report generated at 02/27/2008 14:59:01
Policy ITSOPOL is enabled
Maximum provisioning scope:
 Limit for CPC SCZP201 is 500 MSUs, 5 zAAPs, 0 zIIPs
Rule rule1 is enabled (default enabled)
 Provisioning scope:
  Limit for CPC SCZP201 is 500 MSUs, 2 zAAPs, 0 zIIPs
 Provisioning condition SystemSC80 is enabled (default enabled)
  Workload condition Java_SC80 for sysplex WTSCPLX8
   System name "SC80"
   Included: */*/JAVAWKLD.1
    PL/PD/DL/DD/S 2.0 5 1.5 5 System
  Time condition year2007 is active and enabled
   s/d/e: 01/01/2007 00:00 / 12/31/2008 23:59 / 12/31/2008 23:59
```

Figure 7-15 Report active policy

After you finish setting up the configuration and the policy, you may set the mode of CPM processing. CPM is able to work in the following modes:

- Manual
- Analysis
- Confirmation
- Autonomic

When CPM is started the first time with an empty policy (with no rules), manual mode is the default. The mode can be set either in a CPM configuration, or it can be changed dynamically while CPM is active.

The recommendation is to start it in manual mode until a complete policy is defined. Then switch to confirmation mode, as shown in Figure 7-16.

```
F CPOSERV,APPL=SET DOMAIN MODE=CONF
CPO1009I Processing mode successfully changed to CONF
```

Figure 7-16 Change of CPM processing mode

A complete reference of CPM commands is documented in *z/OS MVS Capacity Provisioning User's Guide,* SA33-8299.

7.3 CPM provisioning

This scenario uses Java workload in both SC80 and SC81. They both run on z10 EC with a permanent capacity 709, with two zAAPs, two zIIPs, and four ICFs. The On/Off CoD record is installed. The z/OS image that triggers additional capacity provisioning if workload suffers (SC80) has nine logical CPs, two zAAPs, and two zIIPs online, as shown in Figure 7-17 on page 171.

The image profile was activated with the following resources:

- Four initial and six reserved CPs
- Two initial and two reserved zAAPs
- Two initial and two reserved zIIPs

```
D M=CPU
IEE174I 23.55.50 DISPLAY M 499
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DE502097
01  +                     01DE502097
02  +                     01DE502097
03  +                     01DE502097
04  +A                    01DE502097
05  +A                    01DE502097
06  +I                    01DE502097
07  +I                    01DE502097
08  +                     01DE502097
09  +                     01DE502097
0A  +                     01DE502097
0B  +                     01DE502097
0C  +                     01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1
```

Figure 7-17 Display command during normal activity

Note: All RMF reports shown in this chapter are from system SC80.

The RMF Monitor III System Information (Sysinfo) report, in Figure 7-18 on page 172, shows a high workflow (WFL) of 94 for the monitored service class JAVAWKLD. Its performance index (PI) on this z/OS system is 0.76.

```
Partition:    A01     2097 Model 709          Appl%:      41  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   42     EAppl%:     41  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   96     Appl% AAP:  57  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      94  163    6        0.00  5.5  0.0    0.4  0.0  0.0  0.0  0.0
*TSO              2    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       94   6    6        0.00  5.5  0.0    0.4  0.0  0.0  0.0  0.0
*STC             152    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH              0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS              3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE           0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W      157    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S      131    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1      131    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  94    6    6  .000  0.00  5.5  0.0    0.4  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  94    6    6  .000  0.00  5.5  0.0    0.4  0.0  0.0  0.0  0.0
          1  94    6    6  .000  0.00  5.5  0.0    0.4  0.0  0.0  0.0  0.0 0.76
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-18 RMF Sysinfo report during normal activity

SC80 is running in logical partition A01. As shown in Figure 7-19 on page 173, the physical utilization percentage for CPs and zAAPs is 100, so the entire server capacity is being used by the logical partitions. This is not a problem for the monitored workload at this point.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  304   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    55  NO   2.0       30.8    30.8   0.0     6.8    6.8
A0F            0     0  NO   2.0        0.2     0.2   0.0     0.0    0.0
A01            0   238  NO   9.0       42.0    42.0   0.0    42.0   42.0
A02            0     0  NO   4.0        0.0     0.1   0.0     0.0    0.0
A1E            0    55  NO   2.0       30.8    30.8   0.0     6.8    6.8
A1F            0     0  NO   2.0        0.2     0.2   0.0     0.0    0.0
A11            0   353  NO   9.0       43.9    43.9   0.0    43.9   43.9
A12            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.1
PHYSICAL                                              0.2            0.2

*AAP                       4.0                       0.0    99.9    100
A01                    NO  2.0       57.2    57.2   0.0    57.2   57.2
A11                    NO  2.0       42.8    42.8   0.0    42.8   42.8
PHYSICAL                                              0.0            0.0
```

Figure 7-19 RMF CPC Capacity report during normal activity

To understand how the partitions are sharing this 100% utilization of CPs and zAAPs, you should know the values of the partition weights, which is described in the following section.

Weights for CPs

Weights of the partitions are defined separately for CPs and each type of specialty engine. In the following section, we focus on weights for CPs.

The weights of the logical partitions sharing CPs are not changed by CPM. However, weights have to be taken into consideration and, in some cases, may have to be readjusted when processors are added or removed. When weights are changed while CPM is in action, you should to understand all performance consequences.

Table 7-1 on page 174 shows the number of logical CPs (LCPs) online and the weights of all the logical partitions. The resulting weight is calculated by using the weight of the corresponding logical partition and the sum of the weights of all of them. A partition can use less of the resulting weight if there is no activity on it to fulfill the capacity. A partition can use more of the resulting weight if it is not capped and other partitions are not consuming their resulting capacity. Table 7-1 on page 174 contains only the resulting weight. Be aware that the real weight assigned by Processor Resource/Systems Manager™ (PR/SM™) can vary within a range around this value. Refer to *IBM System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide,* SB10-7153, for the calculation of resulting weight range.

The column titled Maximum in 709 indicates the percentage of server capacity that can be used by a partition according to the online CPs. For example, A0E has only 2 logical CPs online. This means that its maximum capacity is limited to 2/9ths of the capacity of the server. A01 has 9 logical CPs online, so it could be able to use up to the 9 physical CPs of the server.

As mentioned, CPM does not change partition weight. The maximum percentage of utilization of the server shown in the Maximum in 709 column changes as the server capacity changes.

Table 7-1 Weights of logical partitions

LP	LCPs online	Weight	Weight per LCP	Resulting weight	Maximum in 709
A0E	2	10	5.0	3.7%	22.2%
A0F	2	10	5.0	3.7%	22.2%
A01	9	100	11.1	37.0%	100%
A02	4	10	2.5	3.7%	44.2%
A1E	2	10	5.0	3.7%	22.2%
A1F	2	10	5.0	3.7%	22.2%
A11	9	100	11.9	37.0%	100%
A12	4	10	2.5	3.7%	44.4%
A17	4	10	2.5	3.7%	44.4%
Total	38	270	-	100%	-

Four four partitions have high activity: A0E, A01, A1E and A11. These four partitions together reach a resulting weight of 81.4%. Assuming that the four of them have enough activity to fulfill their corresponding capacity and that they are not capped, the four partitions consume 81.4% of the server capacity, plus the free capacity that the remaining partitions do not consume.

Table 7-2 shows the situation during normal activity of these four partitions.

Table 7-2 Weights and utilization of f4 partitions

LP	LCPs online	Weight	Weight per LCP	Res. weight	Max. in 709	Actual use
A0E	2	10	5.0	3.7%	22.2%	6.8%
A01	9	100	11.1	37.0%	100%	42.0%
A1E	2	10	5.0	3.7%	22.2%	6.8%
A11	9	100	11.9	37.0%	100%	43.9%
Total	38	270	-	100%	-	100%

Comparing the Resulting weight (Res. weight) and Actual use columns shows that all of them are consuming more than their resulting weight. They take advantage of the capacity that is not consumed by the remaining partitions.

When the workload grows, with the same capacity as before, JAVAWKLD shows lower WFL (35) and a higher PI (2.23) than before. The Sysinfo report is shown in Figure 7-20 on page 175. If the PI stays the same for more than five minutes, then according to the CPM policy, more capacity should be provided.

```
 Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   99     Appl% AAP:  60  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      35  159   17        0.00  6.0  0.0   11.1  0.0  0.0  0.0  0.0
*TSO              2    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       35   17   17        0.00  6.0  0.0   11.1  0.0  0.0  0.0  0.0
*STC             137    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH              0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS              3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE           0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W      142    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S      116    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1      116    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  35   17   17  .000  0.00  6.0  0.0   11.1  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  35   17   17  .000  0.00  6.0  0.0   11.1  0.0  0.0  0.0  0.0
          1  35   17   17  .000  0.00  6.0  0.0   11.1  0.0  0.0  0.0  0.0 2.23
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-20 RMF Sysinfo report during peak activity

The RMF monitor III CPC capacity report shown in Figure 7-21 on page 176 confirms that the actual MSUs consumption has increased for the logical partition with the peak workload (A01). The physical utilization percentage of CPs is now 44.1% for this partition. The physical utilization percentage for zAAPs has increased, too. It was 57.2% during normal activity, and is 60.1% now.

The physical utilization of CPs and zAAPs is 100% in the entire server. This means that the peak activity makes logical partition A01 take processing capacity that was used by other partitions before.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  303   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      38.0                        0.2    99.8    100
A0E            0    49  NO   2.0       27.1    27.1    0.0     6.0    6.0
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   355  NO   9.0       44.1    44.1    0.0    44.1   44.1
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    49  NO   2.0       27.1    27.1    0.0     6.0    6.0
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   350  NO   9.0       43.5    43.5    0.0    43.5   43.5
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     1  NO   4.0        0.2     0.2    0.0     0.1    0.1
PHYSICAL                                           0.1             0.1

*AAP                     4.0                        0.0     100    100
A01                  NO  2.0       60.1    60.1    0.0    60.1   60.1
A11                  NO  2.0       39.9    39.9    0.0    39.9   39.9
PHYSICAL                                           0.0             0.0
```

Figure 7-21 RMF CPC Capacity report during peak activity

There are 38 logical CPs defined across partitions, and nine physical CPs in the server. CPM will consider CP provisioning.

There are four logical zAAPs defined across partitions, and two physical zAAPs in the server. CPM will consider zAAP provisioning.

In this case, both CP and zAAP can help a suffering workload because the workload is dispatched to both CPs and zAAPs. If CPM can choose between provisioning a new CP or a new specialty engine, it will choose a specialty engine first.

CPM proposes adding temporary capacity according to the rule of the active policy. It proposes provisioning one zAAP in this case, as shown in Figure 7-22. After operator confirmation, the new zAAP is provisioned from the On/Off CoD record.

```
*164 CPO4201I Proposed upgrade for CPC SCZP201 is model 709 with 1
zAAPs and 0 zIIPs. Enter 1 to activate or 2 reject
R 164,1
IEE600I REPLY TO 164 IS;1
CPO4108I Activation of resources for CPC SCZP201 successfully
initiated: model 709 (0/0) with 1 zAAPs and 0 zIIPs
CPO3030I Command completed successfully for CPC SCZP201
```

Figure 7-22 CPM message to provision a zAAP

After, a new zAAP is available in the server (as shown in Figure 7-23 on page 177) but not online in z/OS on SC80 (CPU 0E is available, but is offline).

Note: CPM provisions additional capacity for a server, but does not change the online configuration in the operating system.

```
D M=CPU
IEE174I 00.08.53 DISPLAY M 505
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  -A
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1
```

Figure 7-23 Display command after provisioning a zAAP

The Temporary Upgrades panel in SE can be used to check what has been activated. Figure 7-24 on page 178 shows the On/Off CoD record active. It can provide two zAAPs, and one of them is active.

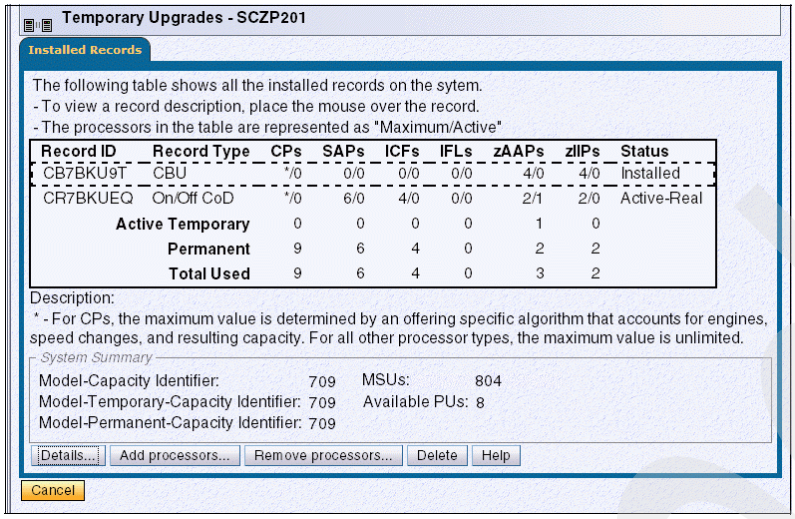

Figure 7-24 Temporary Upgrades - On/Off CoD activated

The additional capacity effect is that service class JAVAWKLD performs slightly better, but the PI is still high and observed workflow is unchanged (it remains at 35). See Figure 7-25.

```
Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:    0     Appl% AAP:  75  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      35  159   17        0.04  6.0  0.0   10.9  0.0  0.0  0.0  0.0
*TSO          2    0        0.04  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       35   17   17        0.00  6.0  0.0   10.9  0.0  0.0  0.0  0.0
*STC        100  137    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH         0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  142    0  .031  0.04  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S      3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1      3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S    116    0  .031  0.04  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1    116    0  .031  0.04  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  35   17   17  .000  0.00  6.0  0.0   10.9  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  35   17   17  .000  0.00  6.0  0.0   10.9  0.0  0.0  0.0  0.0
          1  35   17   17  .000  0.00  6.0  0.0   10.9  0.0  0.0  0.0  0.0 2.14
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-25 RMF Sysinfo report after provisioning a zAAP

Also, the MSU utilization is unchanged, as shown in Figure 7-26. The utilization of CPs by A01 has not changed. The change is in the zAAPs utilization where A01 now uses 50.3% and A11 uses 49.6% of the three zAAPs. So, both partitions benefit from the additional zAAP, not only the observed one.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  303   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    54  NO   2.0       30.3    30.3    0.0     6.7    6.7
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   355  NO   9.0       44.1    44.1    0.0    44.1   44.1
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    54  NO   2.0       30.3    30.3    0.0     6.7    6.7
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   338  NO   9.0       42.1    42.1    0.0    42.1   42.1
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1

*AAP                       4.0                       0.1    99.9    100
A01                    NO  2.0       75.4    75.4    0.0    50.3   50.3
A11                    NO  2.0       74.4    74.4    0.0    49.6   49.6
PHYSICAL                                                 0.1            0.1
```

Figure 7-26 RMF CPC Capacity report after provisioning a zAAP

You may use RMF PM to monitor similar information as you do with RMF Monitor III. Figure 7-27 on page 180 shows the data views selected for this case. The three data views show the delay for processor, the PI for the service class JAVAWKLD in the system SC80, and the list of service classes with higher delays.

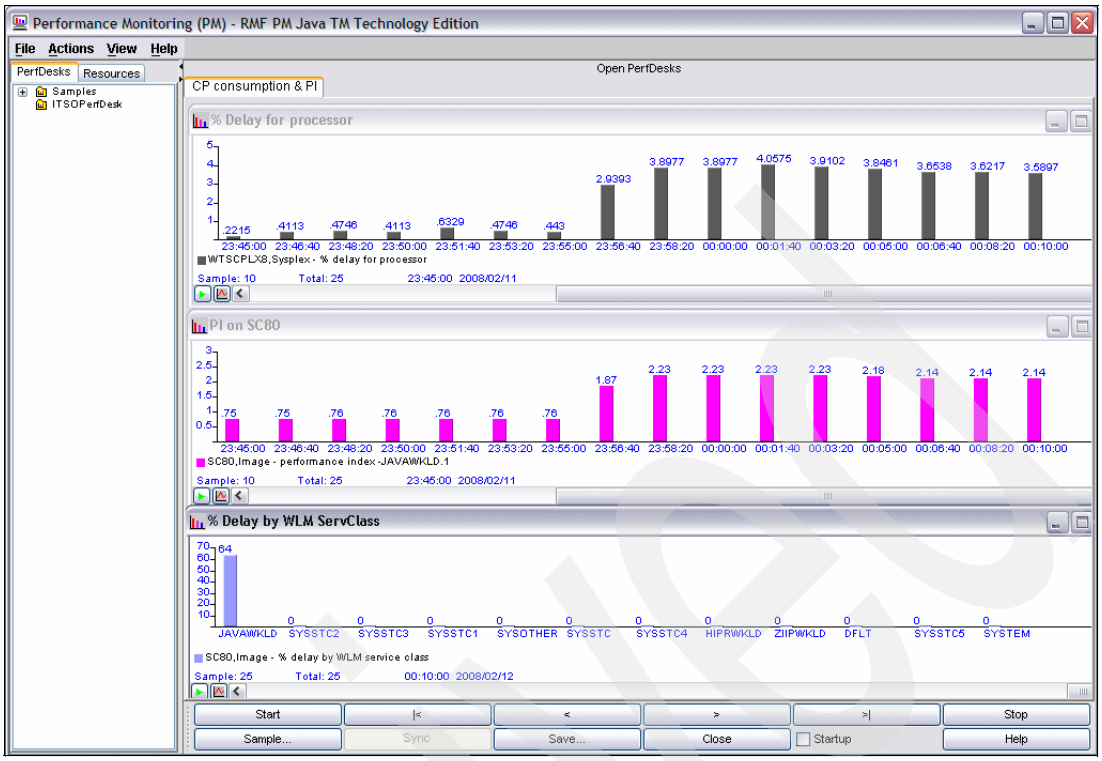

Figure 7-27 RMF PM with selected data views

A value greater than 2 for the PI indicates to CPM that more temporary capacity is needed.

There are 38 logical CPs defined across partitions, and nine physical CPs in the server. CPM will consider CP provisioning.

There are four logical zAAPs defined across partitions, and three physical zAAPs in the server. CPM will consider zAAP provisioning.

CPM proposes provisioning a second zAAP, and the request is confirmed by an operator, as shown in Figure 7-28.

```
*165 CP04201I Proposed upgrade for CPC SCZP201 is model 709 with 2
zAAPs and 0 zIIPs. Enter 1 to activate or 2 reject
R 165,1
IEE600I REPLY TO 165 IS;1
CP04108I Activation of resources for CPC SCZP201 successfully
initiated: model 709 (0/0) with 2 zAAPs and 0 zIIPs
CP03030I Command completed successfully for CPC SCZP201
```

Figure 7-28 CPM message to provision a second zAAP

Figure 7-29 on page 181 shows the activity report. Both activations are displayed, with the latest displayed at the top of the report.

```
F CPOSERV,APPL=R A
CPO1042I Activity report generated at 02/12/2008 05:11:15
Number of activities between 12/11/2007 and 02/12/2008 was 2
Activation for CPC SCZP201 at 02/12/2008 05:10:43
 Activation of model 709, 2 zAAPs and 0 zIIPs
 Active resources before activation: model 709, 1 zAAPs, 0 zIIPs
 Inducing policy element is policy ITSOPOL, rule rule1,
  provisioning condition SystemSC80, time condition year2007
 Inducing system is SC80 in sysplex WTSCPLX8
 Inducing workload is WLM service definition wlmdef1,
  policy POLICY1, service class period JAVAWKLD.1
 Inducing policy element is policy ITSOPOL, rule rule1,
  provisioning condition SystemSC80, time condition year2007
 Inducing system is SC80 in sysplex WTSCPLX8
 Inducing workload is WLM service definition wlmdef1,
  policy POLICY1, service class period JAVAWKLD.1
Activation for CPC SCZP201 at 02/12/2008 05:04:30
 Activation of model 709, 1 zAAPs and 0 zIIPs
 Active resources before activation: model 709, 0 zAAPs, 0 zIIPs
 Inducing policy element is policy ITSOPOL, rule rule1,
  provisioning condition SystemSC80, time condition year2007
 Inducing system is SC80 in sysplex WTSCPLX8
 Inducing workload is WLM service definition wlmdef1,
  policy POLICY1, service class period JAVAWKLD.1
 Inducing policy element is policy ITSOPOL, rule rule1,
  provisioning condition SystemSC80, time condition year2007
 Inducing system is SC80 in sysplex WTSCPLX8
 Inducing workload is WLM service definition wlmdef1,
  policy POLICY1, service class period JAVAWKLD.1
```

Figure 7-29 Report activity from CPM

As a consequence of having a new zAAP temporarily active, the PI for JAVAWKLD in SC80 has decreased to 2.02, as seen in Figure 7-30 on page 182.

```
Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   99     Appl% AAP:  99  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      35  159   17        0.00  6.0  0.0   11.0  0.0  0.0  0.0  0.0
*TSO          2    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       35   17   17        0.00  6.0  0.0   11.0  0.0  0.0  0.0  0.0
*STC        100  137    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH         0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  142    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S 100  116    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100  116    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  35   17   17  .000  0.00  6.0  0.0   11.0  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  35   17   17  .000  0.00  6.0  0.0   11.0  0.0  0.0  0.0  0.0
          1  35   17   17  .000  0.00  6.0  0.0   11.0  0.0  0.0  0.0  0.0 2.02
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-30 RMF Sysinfo report after provisioning a second zAAP

MSU utilization has not changed, as seen in Figure 7-31 on page 183. Again, the change is only in the zAAP utilization. Of the fours zAAPs, A01 now uses 50.0% and A11 uses 49.0%.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  303   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    56  NO   2.0       31.1    31.1    0.0     6.9    6.9
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   355  NO   9.0       44.1    44.1    0.0    44.1   44.1
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    56  NO   2.0       31.1    31.2    0.0     6.9    6.9
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   335  NO   9.0       41.7    41.8    0.0    41.7   41.8
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                              0.1            0.1

*AAP                       4.0                       0.1    99.0   99.1
A01                    NO  2.0        100     100    0.0    50.0   50.0
A11                    NO  2.0       98.0    98.0    0.0    49.0   49.0
PHYSICAL                                              0.1            0.1
```

Figure 7-31 RMF CPC Capacity report after provisioning a second zAAP

There are 38 logical CPs defined across partitions, and nine physical CPs in the server. CPM will consider CP provisioning.

There are four logical zAAPs defined across partitions, and four physical zAAPs in the server. CPM will *not* consider zAAP provisioning.

In this case, even if CPM could consider zAAP provisioning it would not be possible because the On/Off CoD record is limited to two zAAPs, and also because provisioning limits in the policy rule prevent it.

The PI remains larger than 2, so CPM proposes provisioning a CP. This implies a change from capacity setting 709 to 710, with the corresponding billing consequences for software. After operator confirmation, a new CP is activated; see Figure 7-32 on page 184.

```
*166 CPO4201I Proposed upgrade for CPC SCZP201 is model 710 with 2
 zAAPs and 0 zIIPs. Enter 1 to activate or 2 reject
 R 166,1
 IEE600I REPLY TO 166 IS;1
 CPO4108I Activation of resources for CPC SCZP201 successfully
 initiated: model 710 (1/0) with 2 zAAPs and 0 zIIPs
 CPO3030I Command completed successfully for CPC SCZP201
 F CPOSERV,APPL=R A
 CPO1042I Activity report generated at 02/12/2008 05:19:29
 Number of activities between 12/11/2007 and 02/12/2008 was 3
 Activation for CPC SCZP201 at 02/12/2008 05:18:46
  Activation of model 710, 2 zAAPs and 0 zIIPs
  Active resources before activation: model 709, 2 zAAPs, 0 zIIPs
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
.............................
```

Figure 7-32 CPM message to upgrade to model 710

CPs are shared by more partitions than zAAPs and the net effect of adding a CP is very small to an observed partition. The RMF Monitor III Sysinfo report, shown in Figure 7-33 on page 185, indicates that observed workload is not benefitting from the new CP in terms of workflow (WFL) and PI.

```
Partition:    A01     2097 Model 710          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   85     Appl% AAP:  99  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      35  156   17        0.00  5.9  0.0   11.1  0.0  0.0  0.0  0.0
*TSO          2    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       35   17   17        0.00  5.9  0.0   11.1  0.0  0.0  0.0  0.0
*STC        100  134    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH         0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  139    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S 100  113    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100  113    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  35   17   17  .000  0.00  5.9  0.0   11.1  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  35   17   17  .000  0.00  5.9  0.0   11.1  0.0  0.0  0.0  0.0
          1  35   17   17  .000  0.00  5.9  0.0   11.1  0.0  0.0  0.0  0.0 2.02
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-33 RMF Sysinfo report after upgrading to model 710

A closer look at Figure 7-34 on page 186 reveals details of why MSU utilization is now lower for an observed logical partition, although server capacity has increased.

```
Partition:   A01        2097 Model 710
CPC Capacity:     875 Weight % of Max: ****   4h Avg:  303   Group:   N/A
Image Capacity:   875 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    97  NO   2.0       55.3    55.3   0.0    11.1   11.1
A0F            0     0  NO   2.0        0.1     0.2   0.0     0.0    0.0
A01            0   348  NO   9.0       44.1    44.1   0.0    39.7   39.7
A02            0     0  NO   4.0        0.0     0.0   0.0     0.0    0.0
A1E            0    97  NO   2.0       55.3    55.3   0.0    11.1   11.1
A1F            0     0  NO   2.0        0.1     0.1   0.0     0.0    0.0
A11            0   331  NO   9.0       42.0    42.0   0.0    37.8   37.8
A12            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.1
PHYSICAL                                              0.1            0.1

*AAP                       4.0                       0.1    99.0   99.1
A01                    NO  2.0       100     100    0.0    50.0   50.0
A11                    NO  2.0      98.0    98.1    0.0    49.0   49.0
PHYSICAL                                              0.1            0.1

*ICF                       4.0                       0.0     100    100
A33                        4.0       100     100    0.0     100    100
PHYSICAL                                              0.0            0.0

*IIP                       4.0                       0.2     0.0    0.2
A01                    NO  2.0       0.0     0.0    0.0     0.0    0.0
A11                    NO  2.0       0.0     0.0    0.0     0.0    0.0
PHYSICAL                                              0.2            0.2
```

Figure 7-34 RMF CPC Capacity report after upgrading to 710

The new added capacity has been absorbed by logical partitions A0E and A1E. A01 is not using more processing capacity than before. The maximum utilization percentage has changed now, as shown in Table 7-3. A01 cannot reach more than 90% utilization with nine CPs online.

Table 7-3 Weights and utilization for model 710

LP	LCPs online	Weight	Weight per LCP	Res. weight	Max in 710	Actual use
A0E	2	10	5.0	3.7%	20.0%	11.1%
A01	9	100	11.1	37.0%	90.0%	39.7%
A1E	2	10	5.0	3.7%	20.0%	11.1%
A11	9	100	11.9	37.0%	90.0%	37.8%
Total	38	270	-	100%	-	100%

The PI stays over 2 and CPM proposes adding temporary capacity, as seen in Figure 7-35. This time, the provisioning of a second CP means a change from a model 710 to a model 711.

```
*167 CPO4201I Proposed upgrade for CPC SCZP201 is model 711 with 2
 zAAPs and 0 zIIPs. Enter 1 to activate or 2 reject
 R 167,1
 IEE600I REPLY TO 167 IS;1
 CPO4108I Activation of resources for CPC SCZP201 successfully
 initiated: model 711 (2/0) with 2 zAAPs and 0 zIIPs
 CPO3030I Command completed successfully for CPC SCZP201
 F CPOSERV,APPL=R A
 CPO1042I Activity report generated at 02/12/2008 05:50:40
 Number of activities between 12/11/2007 and 02/12/2008 was 4
 Activation for CPC SCZP201 at 02/12/2008 05:48:01
  Activation of model 711, 2 zAAPs and 0 zIIPs
  Active resources before activation: model 710, 2 zAAPs, 0 zIIPs
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
.............................................
```

Figure 7-35 CPM message to upgrade to model 711

After confirming the change, the RMF Monitor III Sysinfo report, shown in Figure 7-36 on page 188, indicates that an observed workload has improved in terms of workflow and PI. The workflow for JAVAWKLD is now 41 and the PI is 1.73.

```
Partition:    A01     2097 Model 711          Appl%:      54  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   55     EAppl%:     54  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   99     Appl% AAP:  99  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      41  156   17        0.15  7.0  0.0   10.0  0.0  0.0  0.0  0.0
*TSO        100    2    0        0.15  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       41   17   17        0.00  7.0  0.0   10.0  0.0  0.0  0.0  0.0
*STC         67  134    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W  75  139    0  .026  0.15  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S 100  113    0  .026  0.15  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100  113    0  .026  0.15  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S   0   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1   0   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  41   17   17  .000  0.00  7.0  0.0   10.0  0.0  0.0  0.0  0.0
DFLT      S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  41   17   17  .000  0.00  7.0  0.0   10.0  0.0  0.0  0.0  0.0
          1  41   17   17  .000  0.00  7.0  0.0   10.0  0.0  0.0  0.0  0.0 1.73
ZIIPWKLD  S        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-36 RMF Sysinfo message after upgrading to model 711

As shown in Figure 7-37 on page 189, the server capacity has increased, and so has MSU utilization by logical partition A01. MSU utilization is directly related to the percentage of physical CP utilization. This time, partition A01 gained most of the added capacity.

```
Partition:   A01        2097 Model 711
CPC Capacity:     944 Weight % of Max: ****   4h Avg:  304   Group:   N/A
Image Capacity:   858 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      38.0                        0.2    99.8    100
A0E            0    97  NO   2.0       56.3    56.3    0.0    10.2   10.2
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   361  NO   9.0       55.1    55.1    0.0    45.0   45.1
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    97  NO   2.0       56.3    56.3    0.0    10.2   10.2
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   322 NO   9.0       41.7    41.7    0.0    34.1   34.2
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                              0.1            0.1

*AAP                     4.0                         0.1    99.1   99.2
A01                 NO   2.0       100     100     0.0    50.0   50.0
A11                 NO   2.0      98.2    98.2     0.0    49.1   49.1
PHYSICAL                                              0.0            0.0

*ICF                     4.0                         0.0     100    100
A33                      4.0       100     100     0.0     100    100
PHYSICAL                                              0.0            0.0

*IIP                     4.0                         0.2     0.0    0.2
A01                 NO   2.0       0.0     0.0     0.0     0.0    0.0
A11                 NO   2.0       0.0     0.0     0.0     0.0    0.0
PHYSICAL                                              0.2            0.2
```

Figure 7-37 RMF CPC Capacity report after upgrading to model 711

CPM continues observing this workload on system SC80, but now the policy determines that no additional capacity is necessary because the PI is under 2.

7.4 CPM deprovisioning

Before discussing deprovisioning by CPM, an important point to remember is that the system itself can deprovision resources that were provisioned by CPM. If a customer is provisioning from offering records that contain capacity tokens, deprovisioning by the system itself can take place if capacity tokens for a provisioned resource are totally consumed. This section assumes that all resources have enough capacity tokens, so that CPM can perform the discussed functions. See “CPM and On/Off CoD offering records with capacity tokens” on page 209 for more details.

After some time, the workload decreases, as seen in Figure 7-38 on page 190.

```
Partition:    A01     2097 Model 711          Appl%:      49  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   50     EAppl%:     49  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   92     Appl% AAP:  98  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      92  156    7        0.25  6.3  0.0    0.6  0.0  0.0  0.0  0.0
*TSO        100    2    0        0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       92    7    7        0.00  6.2  0.0    0.6  0.0  0.0  0.0  0.0
*STC        100  144    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH         0    0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS         3    0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE      0  N/A    0         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  149    0  .028  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S     3    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1     3    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S 100  123    0  .028  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100  123    0  .028  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  92    7    7  .000  0.00  6.2  0.0    0.6  0.0  0.0  0.0  0.0
DFLT      S     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  92    7    7  .000  0.00  6.2  0.0    0.6  0.0  0.0  0.0  0.0
          1  92    7    7  .000  0.00  6.2  0.0    0.6  0.0  0.0  0.0  0.0 0.75
ZIIPWKLD  S     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1     0    0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-38 RMF Sysinfo report after peak activity

The same information is shown in the RMF PM in Figure 7-39 on page 191.

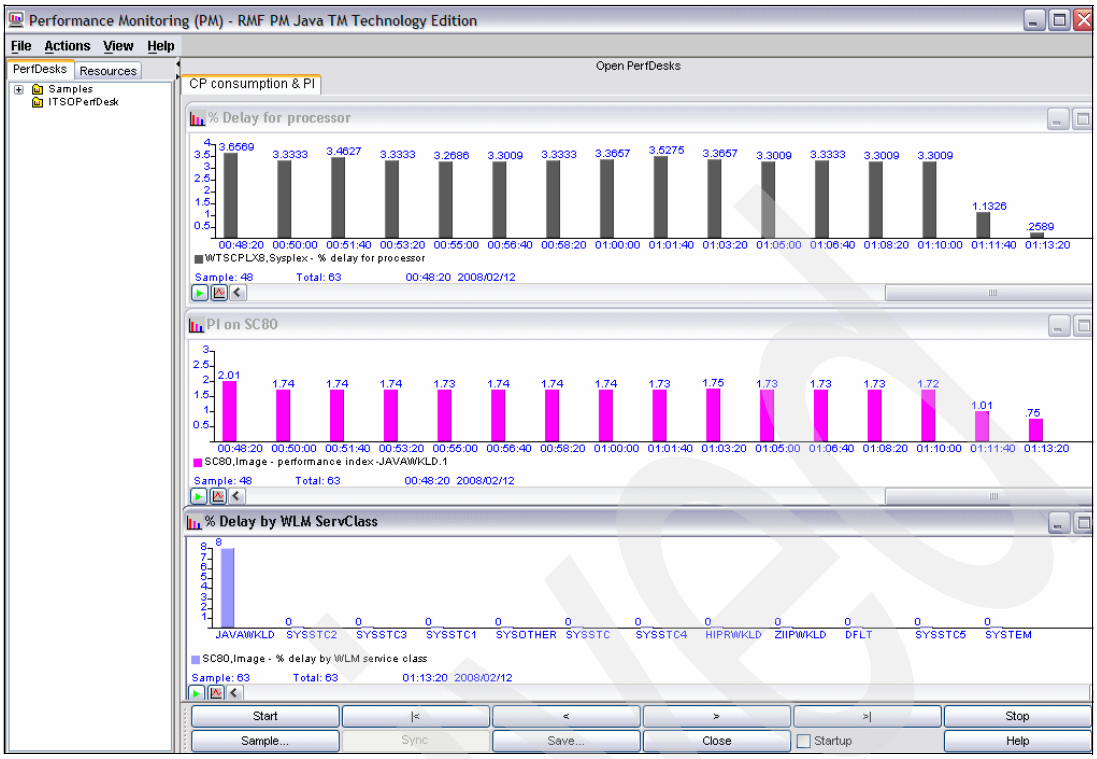

Figure 7-39 RMF PM with selected views after peak activity

MSU utilization is lower than in the peak moment, as shown in Figure 7-40 on page 192. The zAAPs have also decreased their physical utilization percentage to 98.3%.

```
 Partition:   A01        2097 Model 711
CPC Capacity:     944 Weight % of Max: ****   4h Avg:  309   Group:   N/A
Image Capacity:   858 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      38.0                        0.4    99.2   99.6
A0E            0   119  NO   2.0       69.5    69.6   0.0    12.6   12.6
A0F            0     0  NO   2.0        0.2     0.2   0.0     0.0    0.0
A01            0   382  NO   9.0       49.5    49.5   0.0    40.5   40.5
A02            0     0  NO   4.0        0.0     0.0   0.0     0.0    0.0
A1E            0   119  NO   2.0       69.5    69.6   0.0    12.6   12.6
A1F            0     0  NO   2.0        0.1     0.1   0.0     0.0    0.0
A11            0   314  NO   9.0       40.6    40.7   0.0    33.2   33.3
A12            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.0
PHYSICAL                                             0.3            0.3

*AAP                     4.0                         0.2    98.2   98.3
A01                  NO  2.0       98.2    98.2   0.0    49.1   49.1
A11                  NO  2.0       98.1    98.2   0.0    49.1   49.1
PHYSICAL                                             0.1            0.1
```

Figure 7-40 RMF CPC Capacity information after peak activity

When CPM observes that the PI is below 1.5, it proposes deprovisioning the additional temporary capacity that was previously provisioned. It first deprovisions CPs and then specialty engines, if possible.

CPM proposes a downgrade from model 711 to model 710, as shown in Figure 7-41.

```
*168 CPO4202I Proposed downgrade for CPC SCZP201 is model 710 with 2
zAAPs and 0 zIIPs. Enter 1 to deactivate or 2 to reject
R 168,1
IEE600I REPLY TO 168 IS;1
CPO4109I Deactivation of resources for CPC SCZP201 successfully
initiated: model 710 (1/0) with 2 zAAPs and 0 zIIPs
```

Figure 7-41 CPM message to downgrade to model 710

The PI shown in Figure 7-42 on page 193 has increased, but it is still less than 1.5, so deprovisioning continues.

```
Partition:    A01     2097 Model 710          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   83     Appl% AAP:  99  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      83  156    7        0.16  5.9  0.0    1.3  0.0  0.0  0.0  0.0
*TSO              2    0        0.16  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       85    7    7        0.00  5.9  0.0    1.1  0.0  0.0  0.0  0.0
*STC         32  144    0        0.00  0.1  0.0    0.2  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W   32  149    0  .003  0.16  0.1  0.0    0.2  0.0  0.0  0.0  0.0
SYSOTHER  S       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S   21  123    0  .003  0.16  0.0  0.0    0.2  0.0  0.0  0.0  0.0
          1   21  123    0  .003  0.16  0.0  0.0    0.2  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S  100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1  100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W   85    7    7  .000  0.00  5.9  0.0    1.1  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S   85    7    7  .000  0.00  5.9  0.0    1.1  0.0  0.0  0.0  0.0
          1   85    7    7  .000  0.00  5.9  0.0    1.1  0.0  0.0  0.0  0.0 0.82
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-42 RMF Sysinfo report after downgrading to model 710

CPC capacity has decreased as shown in Figure 7-43 on page 194. Also, the MSUs used by A01 have decreased to 290.

```
Partition:   A01        2097 Model 710
CPC Capacity:     875 Weight % of Max: ****   4h Avg:  309   Group:   N/A
Image Capacity:   875 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                       38.0                       0.2    99.8    100
A0E            0   100  NO   2.0       56.7    56.7    0.0    11.3   11.3
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   290  NO   9.0       43.7    43.7    0.0    39.3   39.3
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0   100  NO   2.0       56.7    56.7    0.0    11.3   11.3
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   329  NO   9.0       41.8    41.8    0.0    37.6   37.7
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1

*AAP                      4.0                        0.1    98.8   98.9
A01                   NO  2.0       99.7    99.7    0.0    49.8   49.8
A11                   NO  2.0       98.0    98.0    0.0    49.0   49.0
PHYSICAL                                                 0.1            0.1
```

Figure 7-43 RMF CPC Capacity report after downgrading to model 710

CPM proposed deprovisioning the only remaining temporary active CP, shown in Figure 7-44.

```
*169 CPO4202I Proposed downgrade for CPC SCZP201 is model 709 with 2
 zAAPs and 0 zIIPs. Enter 1 to deactivate or 2 to reject
 R 169,1
 IEE600I REPLY TO 169 IS;1
 CPO4109I Deactivation of resources for CPC SCZP201 successfully
 initiated: model 709 (0/0) with 2 zAAPs and 0 zIIPs
```

Figure 7-44 CPM message to downgrade to model 709

The current capacity setting is 709, but the effect in workflow and PI for an observed workload is minimum, as shown in Figure 7-45 on page 195.

```
Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:  100     Appl% AAP:  99  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      84  156    7        0.06  5.8  0.0    1.1  0.0  0.0  0.0  0.0
*TSO              2    0        0.06  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       84    7    7        0.00  5.8  0.0    1.1  0.0  0.0  0.0  0.0
*STC         50  144    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W   50  149    0  .007  0.06  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S     123    0  .007  0.06  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1     123    0  .007  0.06  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S   50   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1   50   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W   84    7    7  .000  0.00  5.8  0.0    1.1  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S   84    7    7  .000  0.00  5.8  0.0    1.1  0.0  0.0  0.0  0.0
          1   84    7    7  .000  0.00  5.8  0.0    1.1  0.0  0.0  0.0  0.0 0.82
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-45 RMF Sysinfo report after downgrading to model 710

MSU utilization for A01 as shown in Figure 7-46 on page 196 has not decreased, but increased. Logical partitions A0E and A1E have lost part of their processing capacity.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  309   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    57  NO   2.0       31.8    31.8    0.0     7.1    7.1
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   351  NO   9.0       43.7    43.7    0.0    43.7   43.7
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    57  NO   2.0       31.8    31.8    0.0     7.1    7.1
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   336  NO   9.0       41.8    41.8    0.0    41.8   41.8
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
PHYSICAL                                                 0.1            0.1

*AAP                       4.0                       0.1    98.8   98.9
A01                   NO   2.0       99.7    99.7    0.0    49.8   49.8
A11                   NO   2.0       98.0    98.0    0.0    49.0   49.0
PHYSICAL                                                 0.1            0.1
```

Figure 7-46 RMF CPC Capacity report after downgrading to model 710

Because no temporary CPs are active, CPM can only downgrade by decreasing the number of zAAPs that are active. Because the PI for the selected service class is still less than 1.5 (it is 0.82 now), CPM proposes deprovisioning a zAAP, as seen in Figure 7-47.

```
*170 CPO4202I Proposed downgrade for CPC SCZP201 is model 709 with 1
 zAAPs and 0 zIIPs. Enter 1 to deactivate or 2 to reject
 R 170,1
 IEE600I REPLY TO 170 IS;1
 CPO4109I Deactivation of resources for CPC SCZP201 successfully
 initiated: model 709 (0/0) with 1 zAAPs and 0 zIIPs
 CPO3032I Command completed successfully for CPC SCZP201
```

Figure 7-47 CPM message to deprovision a zAAP

After the deprovisioning of the zAAP, the RMF report in Figure 7-48 on page 197 shows that PI has increased to 0.83, but the conditions for deprovisioning remain.

```
Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:   85     Appl% AAP:  75  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      86  156    7        0.14  5.9  0.0    1.0  0.0  0.0  0.0  0.0
*TSO              2    0        0.14  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       86    7    7        0.00  5.9  0.0    1.0  0.0  0.0  0.0  0.0
*STC        100  144    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  149    0  .011  0.14  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S      123    0  .011  0.14  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1      123    0  .011  0.14  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  86    7    7  .000  0.00  5.9  0.0    1.0  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  86    7    7  .000  0.00  5.9  0.0    1.0  0.0  0.0  0.0  0.0
          1  86    7    7  .000  0.00  5.9  0.0    1.0  0.0  0.0  0.0  0.0 0.83
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-48 RMF Sysinfo report after deprovisioning a zAAP

With one zAAP less than before, MSU utilization for A01 has increased from 296 to 297, as shown in Figure 7-49 on page 198. Physical utilization of zAAPs has increased now to again reach 100% in the entire server. Of the three zAAPs, A01 is using 50% and A11 is using 49.9%.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  309   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      38.0                        0.2    99.8    100
A0E            0    55  NO   2.0       30.9    30.9    0.0     6.9    6.9
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   352  NO   9.0       43.8    43.8    0.0    43.8   43.8
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    55  NO   2.0       30.9    30.9    0.0     6.9    6.9
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   339  NO   9.0       42.1    42.1    0.0    42.1   42.1
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                0.1            0.1

*AAP                     4.0                        0.1    99.9    100
A01                  NO  2.0       75.0    75.0    0.0    50.0   50.0
A11                  NO  2.0       74.8    74.8    0.0    49.8   49.9
PHYSICAL                                                0.1            0.1
```

Figure 7-49 RMF CPC Capacity report after deprovisioning a zAAP

CPM continues deprovisioning and now the only remaining temporary capacity is one zAAP. CPM proposes deprovisioning the last zAAP, as shown in Figure 7-50.

```
*171 CPO4202I Proposed downgrade for CPC SCZP201 is model 709 with 0
 zAAPs and 0 zIIPs. Enter 1 to deactivate or 2 to reject
 R 171,1
 IEE600I REPLY TO 171 IS;1
 CPO4109I Deactivation of resources for CPC SCZP201 successfully
 initiated: model 709 (0/0) with 0 zAAPs and 0 zIIPs
```

Figure 7-50 CPM message to deprovision the second zAAP

After deactivating the last temporary zAAP, the service class continues with a desirable PI (0.83), as shown in Figure 7-51 on page 199. The peak activity is over and the permanent capacity is enough for the current situation.

```
Partition:    A01     2097 Model 709          Appl%:      43  Policy: POLICY1
CPs Online:   9.0     Avg CPU Util%:   44     EAppl%:     43  Date:   02/11/08
AAPs Online:  2.0     Avg MVS Util%:  100     Appl% AAP:  87  Time:   23.22.20
IIPs Online:  2.0                             Appl% IIP:   0
Group     T WFL --Users--  RESP TRANS -AVG USG-   -Average Number Delayed - PI
             %   TOT  ACT  Time  /SEC PROC  DEV   PROC  DEV STOR SUBS OPER

*SYSTEM      87  156    7        0.25  5.9  0.0    0.9  0.0  0.0  0.0  0.0
*TSO              2    0        0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*BATCH       87    7    7        0.00  5.9  0.0    0.9  0.0  0.0  0.0  0.0
*STC        100  144    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ASCH             0    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*OMVS             3    0        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A    0.0  N/A  0.0  N/A  N/A
SYSTEM    W 100  149    0  .347  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
SYSOTHER  S        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1        3    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
SYSSTC    S      123    0  .347  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1      123    0  .347  0.25  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
...
SYSTEM    S 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1 100   23    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
WRKLD1    W  87    7    7  .000  0.00  5.9  0.0    0.9  0.0  0.0  0.0  0.0
DFLT      S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
HIPRWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
JAVAWKLD  S  87    7    7  .000  0.00  5.9  0.0    0.9  0.0  0.0  0.0  0.0
          1  87    7    7  .000  0.00  5.9  0.0    0.9  0.0  0.0  0.0  0.0 0.83
ZIIPWKLD  S       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0
          1       0    0  .000  0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  N/A
```

Figure 7-51 RMF Sysinfo report after deprovisioning the last zAAP

The RMF CPC Capacity report, in Figure 7-52 on page 200, shows that MSU utilization has decreased for logical partition A01. Conversely, A01 has gained zAAP capacity, taking it from A11. A01 now uses 87.6% of two zAAPs; previously, it used 50.0% of three zAAPs.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804 Weight % of Max: ****   4h Avg:  309   Group:   N/A
Image Capacity:   804 WLM Capping %:    0.0   4h Max:  364   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      38.0                        0.2    91.1   91.3
A0E            0   152  NO   2.0       85.0    85.0   0.0    18.9   18.9
A0F            0     0  NO   2.0        0.3     0.3   0.0     0.1    0.1
A01            0   350  NO   9.0       43.5    43.5   0.0    43.5   43.5
A02            0     0  NO   4.0        0.0     0.0   0.0     0.0    0.0
A1E            0   152  NO   2.0       85.0    85.0   0.0    18.9   18.9
A1F            0     0  NO   2.0        0.1     0.1   0.0     0.0    0.0
A11            0    78  NO   9.0        9.7     9.7   0.0     9.7    9.7
A12            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1   0.0     0.0    0.1
PHYSICAL                                              0.1            0.1

*AAP                     4.0                         0.1    99.7   99.7
A01                  NO  2.0       87.6    87.6   0.0    87.6   87.6
A11                  NO  2.0       12.1    12.1   0.0    12.1   12.1
PHYSICAL                                              0.0            0.0
```

Figure 7-52 RMF CPC Capacity after deprovisioning the last zAAP

In this scenario, no processor was configured online or offline in a z/OS image, and no partition weights were changed. The scenario demonstrated the functionality and value of CPM.

Be aware that the capacity assigned to a logical partition depends on the workloads that run on it and the weights of all partitions in the server. Adding capacity to a server does not always mean that every partition gets more processing capacity. Also, adding more capacity for a partition does not always imply better performance for every workload on the partition.

The scenario described in this section was based on the confirmation mode of processing of CPM. The system administrator only replied to the WTOR messages with the proposed changes of CPM. At any time, the system administrator can also issue CPM commands to change the temporary capacity of the servers in the domain. Figure 7-53 shows an example of manual provisioning.

```
F CPOSERV,APPL=A R CPC=SCZP201 MODEL=711
CPO1023I Temporary upgrade for CPC SCZP201 to model 711 successfully initiated
CPO3030I Command completed successfully for CPC SCZP201
CPO4103I A change of the manually activated resources has been
detected for CPC SCZP201. The base levels for provisioning management
are now 1 CP, capacity level 0, 0 zAAP and 0 zIIP
```

Figure 7-53 Manual provisioning using CPM command

CPM keeps track of which resources were activated by its policy and which were activated manually. It does not try to deactivate resources that were provisioned manually.

7.5 Useful CPM reports

This section shows examples of the useful information that an administrator can obtain from CPM through the z/OS console. Note that these examples have been taken in the same domain, but at different times and dates, so they do not refer to the same situation in all cases.

Domain report

Figure 7-54 shows the report of the domain configuration. It displays the domain configuration (ITSO, in this case) and processing mode (confirmation mode, at that time), and when they were last changed.

```
F CPOSERV,APPL=R D
CPO1008I Domain report generated at 02/07/2008 16:22:59
Provisioning Manager for domain DOMAIN1 started at 02/04/2008 23:49:47
Active processing mode is CONFIRMATION since 02/05/2008 22:30:15
Active domain configuration is ITSO since 02/07/2008 14:26:31
Active policy is ITSOPOL since 02/07/2008 14:59:21
Code level is 2007122414
```

Figure 7-54 CPM Domain report

Configuration report

The configuration report shown in Figure 7-55 on page 202 has information about the current configuration for the servers in the domain and also about the observed z/OS systems.

```
F CPOSERV,APPL=R C

CPO1010I Configuration report generated at 02/07/2008 15:55:30
Domain configuration ITSO for domain DOMAIN1 is enabled
CPC SCZP201 with record * is enabled (default enabled)
 CPC is matched with serial 00002001DE50 since 02/07/2008 14:26:40
 Hardware is of type 2097 with model E26
 Current model is 709 with 804 MSUs, 2 zAAPs, and 2 zIIPs
 Active record ID is CR7BKUEQ
 Hardware has 9 spare processors
 Activation limits are 2 zAAP, and 2 zIIP
 Active resources GP/zAAP/zIIP 0(0/0)/0/0
System SC80 in sysplex WTSCPLX8 is enabled (default enabled)
 Primary host address: 9.12.4.45
 Protocol: HTTP, port: 5988
 The system at primary host address is observed
  This system is available since 02/07/2008 14:26:33
  This system is running on the CPC SCZP201
  WLM service definition: wlmdef1, active policy: POLICY1
System SC81 in sysplex WTSCPLX8 is enabled (default enabled)
 Primary host address: 9.12.4.47
 Protocol: HTTP, port: 5988
 The system at primary host address is observed
  This system is available since 02/07/2008 14:26:31
  This system is running on the CPC SCZP201
  WLM service definition: wlmdef1, active policy: POLICY1
```

Figure 7-55 CPM Configuration report

Policy report

The policy report is shown in Figure 7-56. Remember that installing a policy from CPCC does not make it active. You have to install it from CPCC and later activate it from the z/OS console.

```
F CPOSERV,APPL=R P

CPO1005I Policy report generated at 02/13/2008 12:57:57
Policy ITSOPOL is enabled
Maximum provisioning scope:
 Limit for CPC SCZP201 is 500 MSUs, 5 zAAPs, 0 zIIPs
Rule rule1 is enabled (default enabled)
 Provisioning scope:
  Limit for CPC SCZP201 is 500 MSUs, 2 zAAPs, 0 zIIPs
 Provisioning condition SystemSC80 is enabled (default enabled)
  Workload condition Java_SC80 for sysplex WTSCPLX8
   System name "SC80"
   Included: */*/JAVAWKLD.1
    PL/PD/DL/DD/S 2.0 5 1.5 5 System
  Time condition year2007 is active and enabled
   s/d/e: 01/01/2007 00:00 / 12/31/2008 23:59 / 12/31/2008 23:59
```

Figure 7-56 CPM Policy report

Policy reports show the enabled and disabled rules that are defined in the policy, but disabled rules are not taken into consideration by CPM.

Workload report

Workloads observed by CPM appear in the workload report, as seen in Figure 7-57. System SC81 is an observed system in the domain, but in this case none of the rules include workloads on system SC81 as a possible trigger for provisioning.

```
F CPOSERV,APPL=R W

CPO1047I Workload report generated at 02/13/2008 14:06:39
Workload is analyzed for 2 system(s)
No workload for system SC81 of sysplex WTSCPLX8 on CPC SCZP201 is
observed
Workload for system SC80 of sysplex WTSCPLX8 on CPC SCZP201
 JAVAWKLD.1 PL/PD/DL/DD/S 2.0 5 1.5 5 System
  PI from 02/13/2008 14:03 is 2.26. Last limit crossing was 02/13/2008
 12:53
```

Figure 7-57 CPM workload report

The workload report is a very useful report after the policy is installed. If the output is empty, it indicates that there are no valid workload conditions that would trigger capacity provisioning.

7.6 CPM hints and tips

This section provides helpful information regarding CPM:

Downloading the CPM configuration to CPCC

You may download CPM definitions to CPCC by using FTP. Definitions are stored in the following data sets:

- hlq.domainname.DOMCFG
- hlq.domainname.POLICIES

The hlq is a high level qualifier for CPM (the default is CPO). The domainname is name of the domain that CPM takes care of.

Both data sets are partitioned data sets, where each member represents a different configuration or a policy. In our scenario, domain configuration is named ITSO and it is stored in CPO.DOMAIN1.DOMCFG(ITSO). The policy is named ITSOPOL, and it is saved in CPO.DOMAIN.POLICIES(ITSOPOL).

Definitions for CPCC are stored in a directory specified in CPCC configuration. In our scenario, it is in \Documents and Settings\RESIDENT\My Documents\IBM\IBM Capacity

Provisioning Control Center\workspace. Domain configurations are saved as separate files in the DomainConfigurations\ subdirectory. Policies are saved as separate files in the ProvisioningPolicies\ subdirectory.

Members from PDS data sets can be copied to the appropriate directories by using FTP. Figure 7-58 shows an FTP client. On the left are the directories that CPCC uses. On the right are z/OS data sets containing CPM definitions.

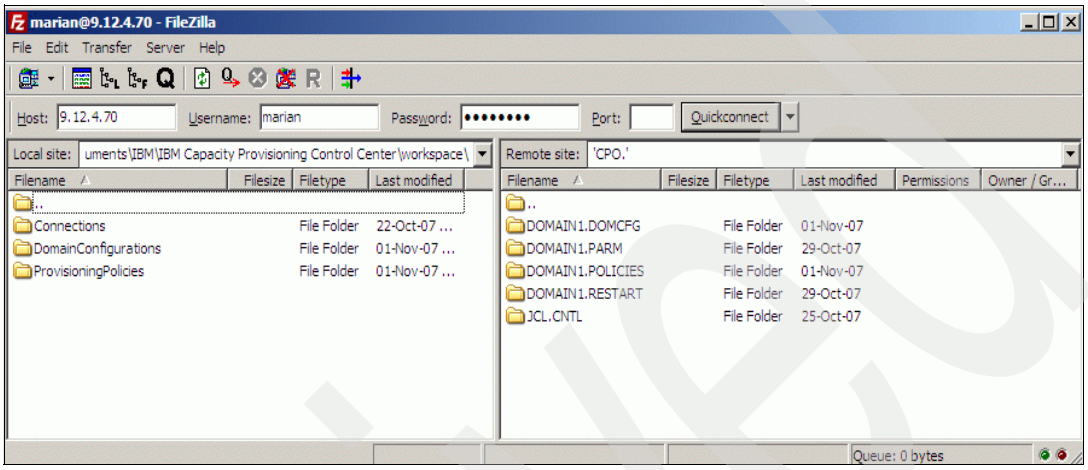

Figure 7-58 Directories with configurations

Figure 7-59 shows the ITSO domain configuration after FTP transfer. The transfer must be done in text (ASCII) mode, and the file must be renamed to end with an .xml extension.

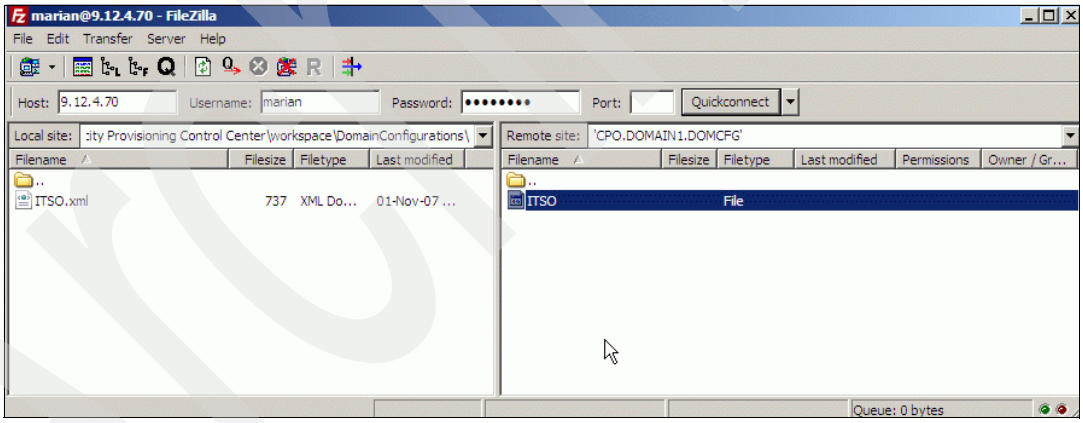

Figure 7-59 Domain configuration transfer

Figure 7-60 on page 205 shows the ITSO domain configuration after FTP transfer. The transfer must be done in text (ASCII) mode, and the file must be renamed to xml.

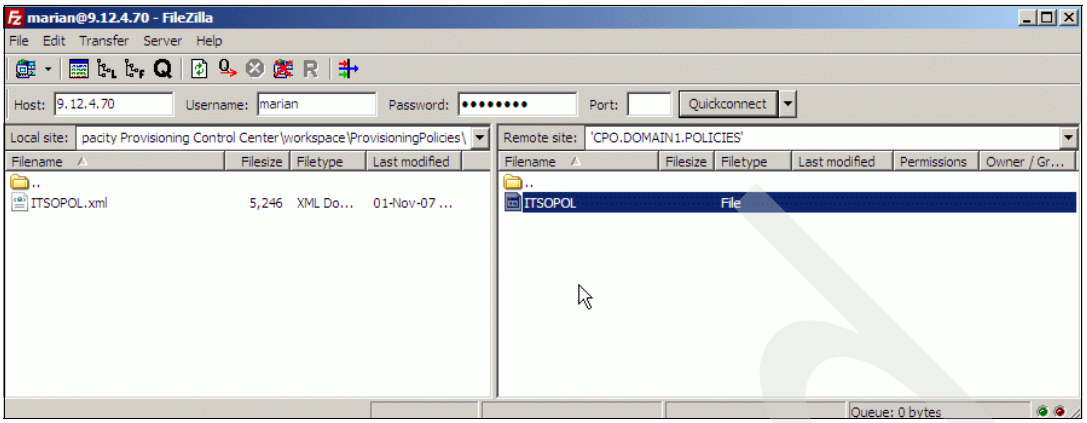

Figure 7-60 Policy transfer

Now CPCC contains the definitions from CPM.

Setting up RMF DDS in high availability mode

Each Parallel Sysplex requires one RMF DDS. The CIM server has to know where RMF DDS is running. However, if the image in the sysplex with RMF DDS is down, RMF DDS is automatically restarted in another sysplex image. To enable CIM so it can find the RMF DDS when the RMF DDS is restarted, the recommendation is that you run RMF DDS in high availability mode.

High availability mode for RMF DDS requires the use of Dynamically Activated IP Address (DVIPA). This requires RACF changes, and changes in the TCP/IP profile and task started with RMF DDS, as the examples in this section demonstrate.

RACF Authorizations

RACF Authorization are shown in Example 7-1.

Example 7-1 RACF Authorizations

```
RDEFINE SERVAUTH (EZB.MODDVIPA.sysname.tcpname) UACC(NONE)
   <<<Repeat for each image (sysname) in the sysplex>>>
PERMIT EZB.MODDVIPA.sysname.tcpname ACCESS(READ) CLASS(SERVAUTH) ID(USER1)
   <<<Repeat for each image (sysname) in the sysplex>>>
   PERMIT EZB.MODDVIPA.TPN.TCPIP ACCESS(READ) CLASS(SERVAUTH) ID(GPMSERVE)
   PERMIT EZB.MODDVIPA.JCO.TCPIP ACCESS(READ) CLASS(SERVAUTH) ID(GPMSERVE)
   PERMIT EZB.MODDVIPA.JBO.TCPIP ACCESS(READ) CLASS(SERVAUTH) ID(GPMSERVE)
   PERMIT EZB.MODDVIPA.ZO.TCPIP ACCESS(READ) CLASS(SERVAUTH) ID(GPMSERVE)
   :
   :
PERMIT EZB.MODDVIPA.TJ8O.TCPIP ACCESS(READ) CLASS(SERVAUTH) ID(GPMSERVE)
SETROPTS RACLIST(SERVAUTH) REFRESH
```

TCP/IP profile definitions

Reserve the Dynamic VIPA address in each TCP/IP profile so that the GPMSERVE procedure can activate it when the procedure is started in any z/OS image, either manually, or by ARM. This has to be conducted for each LPAR's TCP/IP profile.

The VIPARANGE statement must be contained within the TCP/IP profile's VIPADAYNAMIC block. See Example 7-2 on page 206.

Example 7-2 TCP/IP profile definitions

```
VIPADYNAMIC
;
;------------ RMF DDS (gpmserve) HA dvipa
   VIPARANGE DEFINE 255.255.255.255 172.31.1.1
ENDVIPADYNAMIC
```

GPMSERVE procedure modification

The GPMSERVE procedure has to invoke the MODDVIPA utility and create the Dynamic VIPA. Example 7-3

Example 7-3 GPMSERVE procedure modification

```
//GPMSERVE PROC MEMBER=HS
//* PARM='TRAP(ON),ENVAR(ICLUI_TRACETO=STDERR)/&MEMBER'
//*****************************************************************
//* Cleanup: *
//* this step will delete the application activated DVIPA prior *
//* to creating it for the case where the GPMSERVE ASID *
//* had not previously ended normally/cleanly. *
//* RC=8 is expected if the DVIPA was not in use *
//* *
//DELDVP EXEC PGM=MODDVIPA,REGION=0K,TIME=1440,
// PARM='POSIX(ON) ALL31(ON)/-p TCPIP -d 172.31.1.1'
//* *
//**------------ create the DVIPA --------------- *
//* *
//TCPDVP EXEC PGM=MODDVIPA,REGION=0K,TIME=1440,
// PARM='POSIX(ON) ALL31(ON)/-p TCPIP -c 172.31.1.1'
//* *
//*-----------------------------------------------
//* *
//STEP1 EXEC PGM=GPMDDSRV,REGION=0M,TIME=1440,
// PARM='TRAP(ON)/&MEMBER'
//GPMINI DD DISP=SHR,DSN=SYS1.SERBPWSV(GPMINI)
//GPMHTC DD DISP=SHR,DSN=SYS1.SERBPWSV(GPMHTC)
//CEEDUMP DD DUMMY
//SYSPRINT DD DUMMY
//SYSOUT DD DUMMY
//*-------------------------------------------------------
//* delete the DVIPA upon exit
//*
//TCPDVP EXEC PGM=MODDVIPA,REGION=0K,TIME=1440,
// PARM='POSIX(ON) ALL31(ON)/-p TCPIP -d 172.31.1.1'
//*-------------------------------------------------------
//*
// PEND
```

Note: If the TCP/IP stack is restarted and the GPMSERVE remains up, the GPMSERVE will not automatically reconnect to the restarted TCP/IP stack. GPMSERVE will need to be restarted in order to reserve the DVIPA in the restarted TCP/IP instance.

Monitor disk space usage

If tracing or logging is on, CPM can run out of disk space if there are many events to log. It is recommended that you check the disk space usage of these files. Trace data and log data are stored in /tmp by default, but their location can be changed in the CPM PARM member as Trace.Path and Log.Path.

Cleaning the history

Only a limited amount of space is available for console output, so after a time, output can look like the output in Example 7-4.

Example 7-4 Output before cleaning the history

```
F CPOSERV,APPL=R A
CPO2070W Following message may be truncated or missing:
CPO1042I Activity report generated at 02/08/2008 20:58:36
Number of ac
```

In this case, you have three possible actions:

- Redirect output to a file instead of console using DEST= parameter.
- Use FROM= and TO= parameters to limit the output.
- Clean the history, as follows:

  a. Provision down to the base.
  b. Stop CPM.
  c. Delete the PLANNER member in the *hlq.domain*.RESTART data set.
  d. Start CPM.

Note: Stop CPM only when no active resources are provisioned by CPM.

Avoid locking tasks while CPM is provisioning

CPM cannot provision additional capacity while a user is using a connection to the SE and managing temporary capacity. When a need for provisioning is detected, CPM can initiate the provisioning command, but the command cannot complete and the message shown in Figure 7-61 is received.

```
CPO3031W Unsuccessful command completion for CPC SCZP201. Reason 24
```

Figure 7-61 CPM message - unable to complete provisioning

This message indicates that CPM has initiated the command (it is considered as a CPM activity), but the server configuration has not changed.

Figure 7-62 on page 208 shows an example of this situation. Activity history has been cleared just prior to this. According to the CPM policy, additional capacity is needed and CPM proposes an upgrade. The system administrator replies to the WTOR to confirm the upgrade but another user has established a connection to SE and is managing temporary upgrades.

```
*146 CPO4201I Proposed upgrade for CPC SCZP201 is model 709 with 1
 zAAPs and 0 zIIPs. Enter 1 to activate or 2 reject
 R 146,1
 IEE600I REPLY TO 146 IS;1
 CPO4108I Activation of resources for CPC SCZP201 successfully
 initiated: model 709 (0/0) with 1 zAAPs and 0 zIIPs
 CPO3031W Unsuccessful command completion for CPC SCZP201. Reason 24
 F CPOSERV,APPL=R A
 CPO1042I Activity report generated at 02/11/2008 22:45:09
 Number of activities between 12/10/2007 and 02/11/2008 was 1
 Activation for CPC SCZP201 at 02/11/2008 22:43:52
  Activation of model 709, 1 zAAPs and 0 zIIPs
  Active resources before activation: model 709, 0 zAAPs, 0 zIIPs
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
```

Figure 7-62 Proposed upgrade by CPM and report activity

The initiated command fails, but the activity report shows the activation of the zAAP. This activation was initiated but not completed. A few minutes later, the message shown in Figure 7-63 appears in the z/OS console.

```
 CPO4203I Unexpected capacity setting for CPC SCZP201. Action may be
 pending. Expected 709(0/0) 1/0. Actual: 709(0/0) 0/0
*147 CPO4205I CPC SCZP201: Enter 1 to keep waiting for pending
 activation or 2 to accept current capacity setting
 R 147,2
 IEE600I REPLY TO 147 IS;2
 F CPOSERV,APPL=R A
 CPO1042I Activity report generated at 02/11/2008 22:56:09
 Number of activities between 12/10/2007 and 02/11/2008 was 1
 Activation for CPC SCZP201 at 02/11/2008 22:43:52
  Activation of model 709, 1 zAAPs and 0 zIIPs
  Active resources before activation: model 709, 0 zAAPs, 0 zIIPs
  Inducing policy element is policy ITSOPOL, rule rule1,
   provisioning condition SystemSC80, time condition year2007
  Inducing system is SC80 in sysplex WTSCPLX8
  Inducing workload is WLM service definition wlmdef1,
   policy POLICY1, service class period JAVAWKLD.1
```

Figure 7-63 Unexpected capacity setting found by CPM

In this case, the system administrator informs CPM that the actual configuration must be accepted and CPM can continue its provisioning work. The activity report keeps the information of the activation that was initiated but never completed.

The same situation occurs when CPM tries to deprovision resources, however, the warning message is different: message CPO3033W is sent when CPM cannot complete a deactivation.

Specialty engines provisioning with CPM

CPM can provision part of or all CPs, zIIPs and zAAPs included in the On/Off CoD record. These resources are limited by the maximum provisioning scope of the policy and the provisioning scope of the rules. If a requirement exists for CPM to provide only specialty engines, it can be achieved by using the policy which contains zero (0) MSU.

Even if CPs and specialty engines are in the provisioning scope, CPM takes into consideration what kind of additional capacity is needed. Running Java workload with IFAHONORPRIORITY set to NO causes the zAAP-eligible workload to run on zAAPs. This can require additional zAAP capacity but no general purpose CP capacity.

When a service class that is related to Java workload triggers CPM provisioning and CPM detects that only zAAP capacity is required, CPM proposes (if working in confirmation mode) only zAAP upgrades. When the maximum number of zAAPs in the scope is reached, or the PI of the triggering condition falls below the provisioning PI value, CPM stops provisioning. No CP capacity is provisioned when it is not the required type of capacity.

If IFAHONORPRIORITY is changed dynamically to YES, a Java workload can run on CPs and zAAPs. If, after changing IFAHONORPRIORITY, the PI of the triggering service class remains above the provisioning PI value, CPM can provide CP capacity.

CPM and On/Off CoD offering records with capacity tokens

If CPM is provisioning resources from an On/Off CoD offering record that contains capacity tokens, be especially careful to prevent the server from deactivating offering records prematurely. If a capacity token in an offering record is completely consumed, the system itself will deactivate all resources that were activated by that record, even if the activation was performed by CPM, and even if other capacity tokens in the same record have remaining resources available. However, when a capacity token approaches being consumed, the system issues error messages. Such error messages start appearing give days before the capacity token is fully consumed. A strong recommendation is to capture those error messages, and implement procedures to handle the situation. You may either prepare for the deactivation of the entire record, or use Resource Link to replenish the capacity token in question to avoid the premature deactivation of the offering record.

8

Multiple concurrent activations

Capacity on Demand (CoD) design allows you to have more than one offering active and also perform a permanent upgrade while temporary offerings are active. This chapter provides examples of the most commonly expected usages of multiple activations.

This chapter discusses the following topics:

The scenario in this chapter demonstrates the functionality provided by Capacity on Demand (CoD), which brings significant flexibility. The possibility to have several CoD records active at the same time and being able to dynamically change each of them independently allows formidable enhancements in planning and operating complex IT environments.

8.1 On/Off CoD and CBU

This section shows an example of two temporary records that are active at the same time. It shows a possible scenario in which an On/Off CoD record is active as a disaster strikes on a remote site. The CBU record has to be activated to cover the lost capacity. Later, when demand for peak workload goes down, some resources from On/Off CoD can be deactivated and used by CBU. The sequence of activities is shown in Figure 8-1.

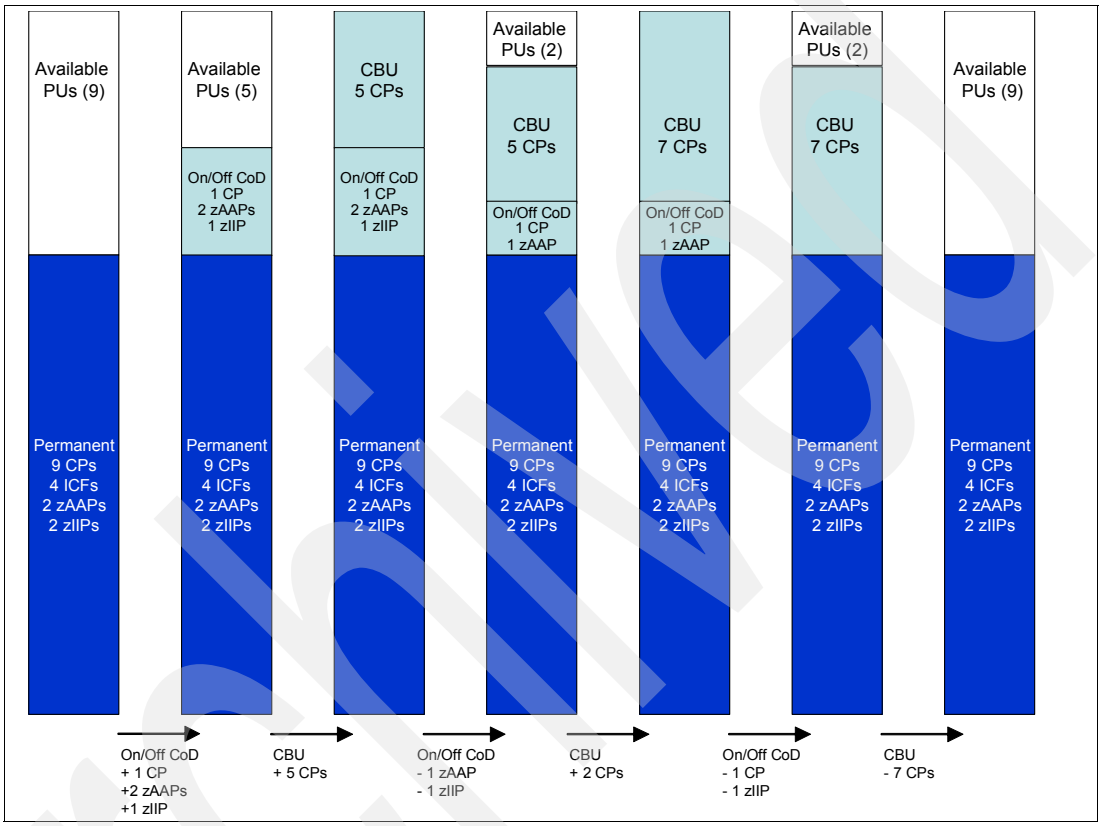

Figure 8-1 Change activation levels

To check the amount of available resources before activating any Capacity on Demand record, review the Installed Records tab on the Temporary Upgrades panel. For example, Figure 8-2 on page 213 shows that no Capacity on Demand record is currently active.

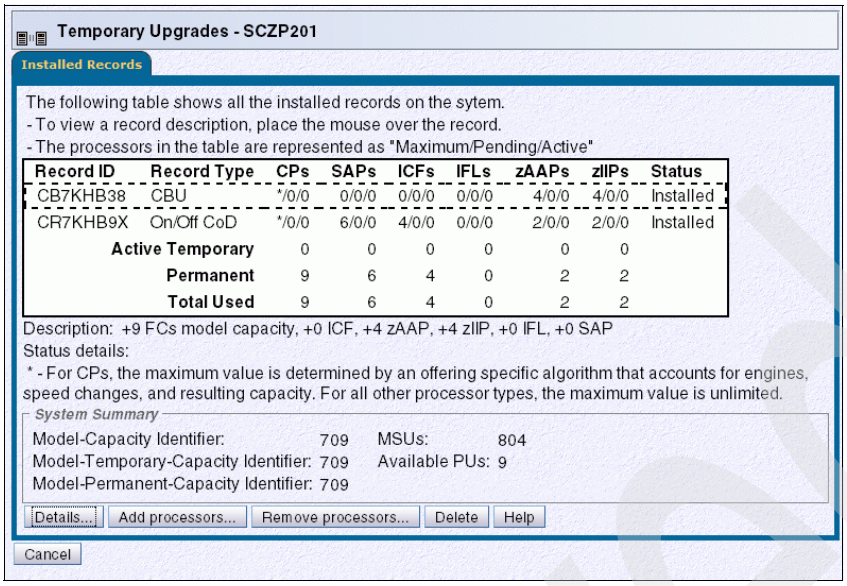

Figure 8-2 Temporary Upgrades panel for reviewing installed records

In our example, logical partition A04 is activated after the required resources are provided by the CBU record. Figure 8-3 shows the processor definitions of the corresponding image profile and that logical partition A04 uses an initial five dedicated central processors (CPs).

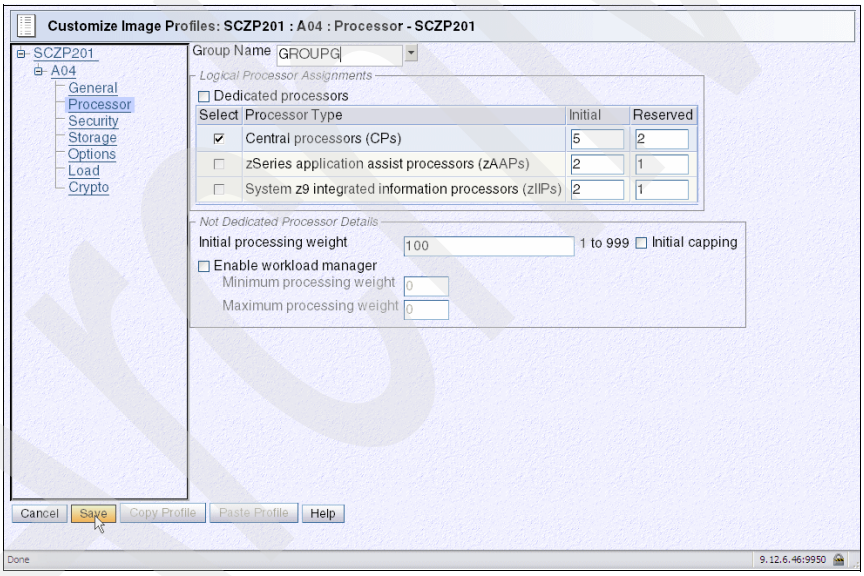

Figure 8-3 A04 image profile for the Processor selection

Although logical partition A04 is not active yet, logical partition A01 is already up and running. As shown in Figure 8-4, logical partition (LP) A01 is currently running with nine CPs, two zAAPs, and two zIIPs.

```
D M=CPU
IEE174I 13.29.47 DISPLAY M 871
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.709.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01        LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-
N NOT AVAILABLE
```

Figure 8-4 Processor assignment in A01

8.1.1 Activating On/Off CoD

The On/Off CoD record activates one CP, two zAAPs, and one zIIP. All resources are selected in the Change Activation Levels panel, as shown in Figure 8-5.

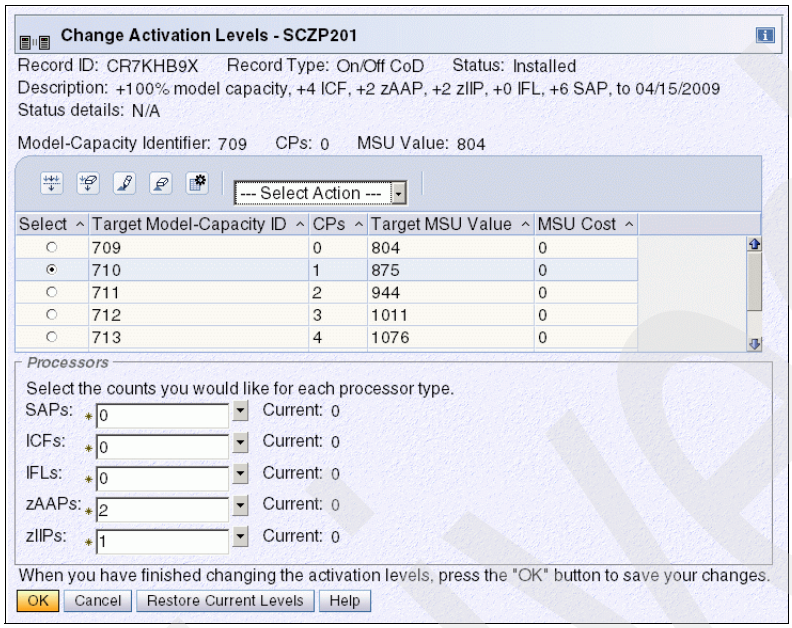

Figure 8-5 Change Activation Levels panel shows selected resources

The Temporary Upgrades panel in Figure 8-6 indicates the current server configuration. The Model-Capacity Identifier and the Model-Temporary-Capacity Identifier reflect the additional billable temporary capacity, while the Model-Permanent-Capacity Identifier remains unchanged. This panel also shows the number of active processors of all types.

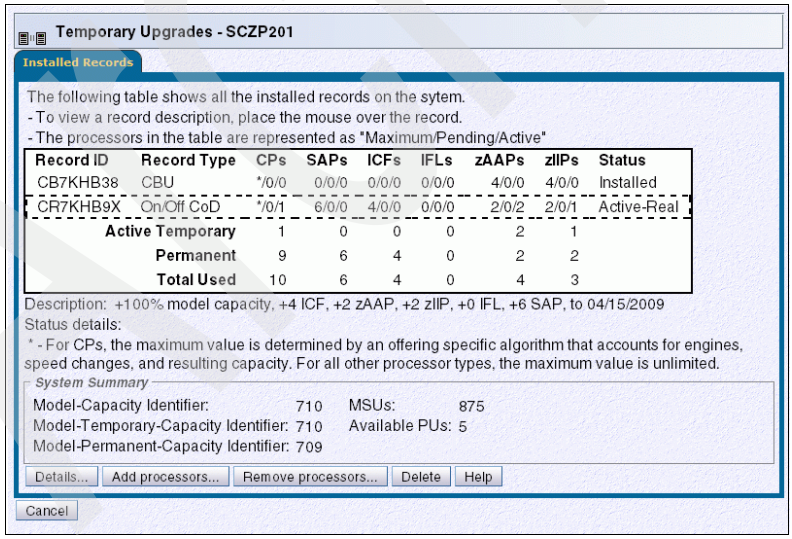

Figure 8-6 Temporary Upgrades panel shows the current server configuration

After the On/Off CoD activation has completed successfully, the additional processors (0D,0E,0F,10) are available to logical partition A01. To use them, configure them online manually by invoking the following z/OS command:

```
CF CPU(0D,0E,0F,10),ONLINE
```

Figure 8-7 shows that the four processors are configured online to logical partition A01. CPC SI indicates the new server capacity setting of 710.

```
D M=CPU
IEE174I 13.40.22 DISPLAY M 889
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  +                    01DE502097
0E  +A                   01DE502097
0F  +A                   01DE502097
10  +I                   01DE502097
11  NI

 CPC ND = 002097.E26.IBM.02.00000001DE50
 CPC SI = 2097.710.IBM.02.000000000001DE50
 CPC ID = 00
 CPC NAME = SCZP201
 LP NAME = A01       LP ID =  1
 CSS ID  = 0
 MIF ID  = 1

 + ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-
 N NOT AVAILABLE
```

Figure 8-7 Processors 0D, 0E, 0F and 10 now online in LP A01

RMF output on Figure 8-8 on page 217 shows the new MSU value for a CPC Capacity and for Image Capacity. Logical partition A01 is configured with ten CPs online.

```
Partition:   A01        2097 Model 710
CPC Capacity:     875   Weight % of Max: ****   4h Avg:  200   Group:   N/A
Image Capacity:   875   WLM Capping %:    0.0   4h Max:  489   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.8                      0.2    99.8    100
A0E            0    35  NO   2.0       20.0    20.0    0.0     4.0    4.0
A0F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A01            0   422  NO  10.0       48.5    48.5    0.0    47.7   47.7
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    35  NO   2.0       20.0    20.0    0.0     4.0    4.0
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   385  NO   9.0       48.9    48.9    0.0    44.0   44.0
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                             0.1            0.1

*AAP                       5.7                       0.0     0.0    0.0
A01                    NO  4.0        0.0     0.0    0.0     0.0    0.0
A11                    NO  2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                             0.0            0.0

*ICF                       4.0                       0.0     100    100
A33                        4.0        100     100    0.0     100    100
PHYSICAL                                             0.0            0.0

*IIP                       4.8                       0.0     0.0    0.0
A01                    NO  3.0        0.0     0.0    0.0     0.0    0.0
A11                    NO  2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                             0.0            0.0
```

Figure 8-8 RMF output after On/Off CoD resources have been made available

After changing the processor configuration, such as adding new physical or logical processors or changing the ratio of physical to logical processors, verify and adjust the logical partition weights as necessary by using the Change Logical Partition Controls task.

8.1.2 Activating CBU

If disaster strikes in another part of an enterprise, while the On/Off CoD record is active and providing additional resources to logical partition A01, a new partition on this server must be started. Partition A04 is prepared to handle replaced workload. However, at this time in our example, not enough resources are available to satisfy the extra demand of A04 to activate it. Partition A04 requires five dedicated CPs.

The CBU record is activated, while the On/Off CoD stays active. No changes are required to the On/Off CoD record or to partition A01. The CBU record in Figure 8-9 on page 218 activates five additional CPs, increasing the Model-Capacity Identifier to 715.

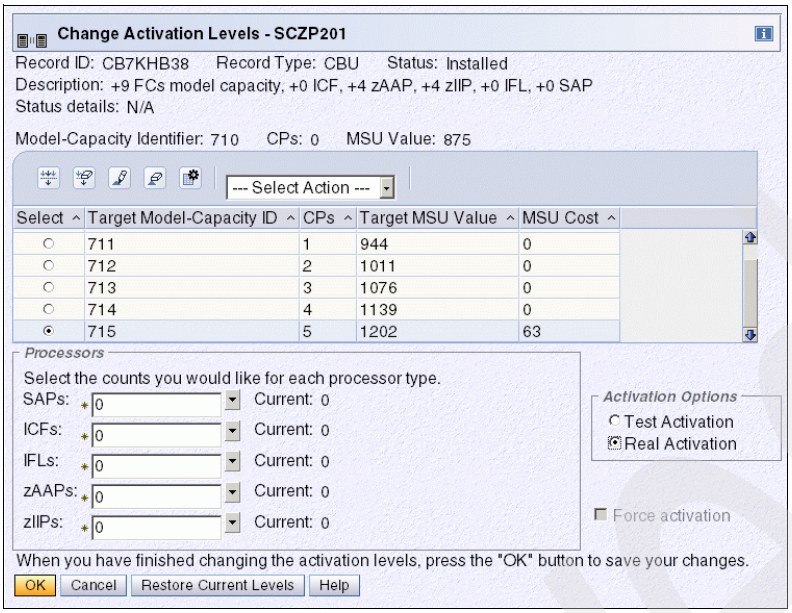

Figure 8-9 Change Activation Levels panel for activating additional CPs

The CBU activation completed successfully. Figure 8-10 shows the new processor configuration and Model-Capacity Identifier. The Model-Capacity Identifier has changed to 715 and the Model-Temporary-Capacity Identifier remains unchanged because billable capacity has not changed. The MSU value has changed. No remaining PUs are available.

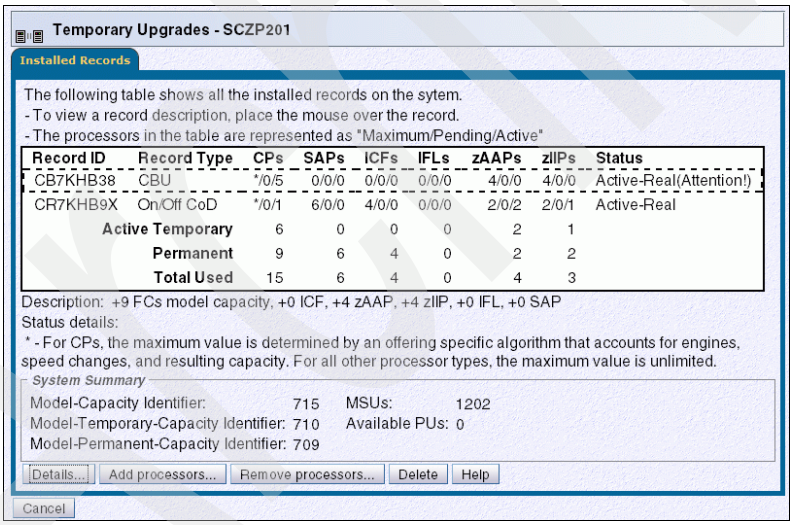

Figure 8-10 Temporary Upgrades panel shows the new configuration

Logical partition A04 can be activated now. Changing partition weights is unnecessary because all added processors are dedicated to a partition and they do not the influence the processor shared pool. The RMF output in Figure 8-11 on page 219 shows five CPs online in logical partition A04, and the Model-Capacity Identifier 715 and the related CPC Capacity MSU value. Notice the change of the Image Capacity value, which was recalculated after new processors were added.

```
Partition:   A01        2097 Model 715
CPC Capacity:    1202   Weight % of Max: ****   4h Avg:  225   Group:   N/A
Image Capacity:   801   WLM Capping %:    0.0   4h Max:  489   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        44.0                      0.1    99.9    100
A0E            0    79  NO   2.0       49.0    49.0    0.0     6.5    6.5
A0F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A01            0   640  NO  10.0       80.0    80.0    0.0    53.3   53.3
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A04            0     0  NO   5.0        0.0     0.0    0.0     0.0    0.0
A1E            0    79  NO   2.0       49.0    49.0    0.0     6.5    6.5
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   402  NO   9.0       55.7    55.7    0.0    33.4   33.4
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1
```

Figure 8-11 LPAR A04 on RMF report

8.1.3 Changing CBU and On/Off CoD activation levels

While the On/Off CoD and CBU records are active, their activation levels can be changed dynamically as needed. Resources can be moved between the different Capacity on Demand records indirectly. They have to be deactivated from one record so they can be used by another record. In our example, logical partition A04 requires two more dedicated CPs. Logical partition A01 can release one zAAP and one zIIP.

Because all physical resources are in use, we have to remove the required resources from the On/Off CoD record first, making them available as additional CBU resources. Before removing the resources from the On/Off CoD record, configure them offline in the associated logical partition.

After successfully configuring the zIIP (processor identifier 10) and the zAAP (processor identifier 0E), the On/Off CoD record can be modified, as shown in Figure 8-12. One zAAP and one zIIP are deactivated from the On/Off CoD record, and only one temporary zAAP stays active.

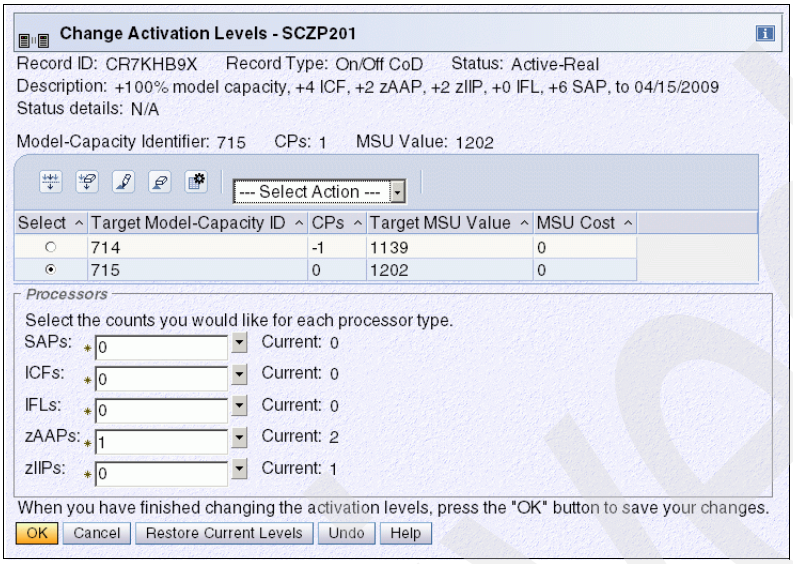

Figure 8-12 Change Activation Levels for modifying the record

One zAAP and one zIIP are removed, while the On/Off CoD and the CBU remain active, as shown in the Temporary Upgrades panel in Figure 8-13. Because CP settings have not changed, all Model-Capacity Identifiers and the MSU value remain unchanged. The number of available PUs is now two.

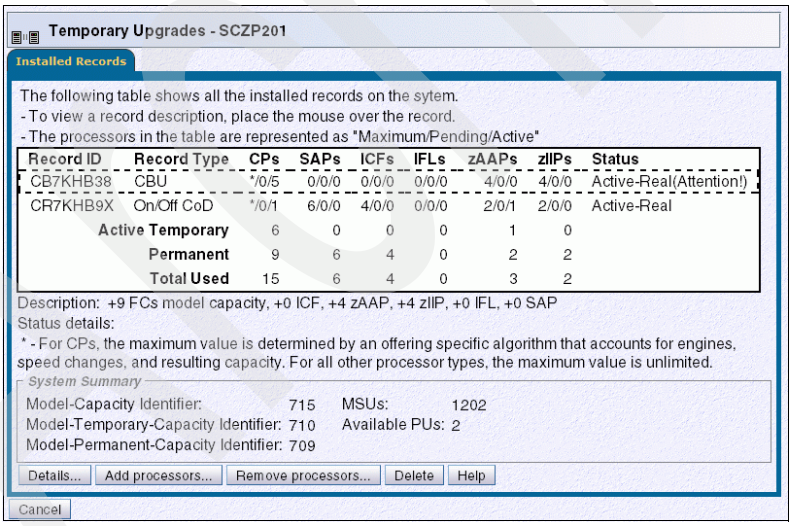

Figure 8-13 Temporary Upgrades panel shows active On/Off CoD and CBU

The two available PUs are activated as CPs in the CBU record, as shown in Figure 8-14.

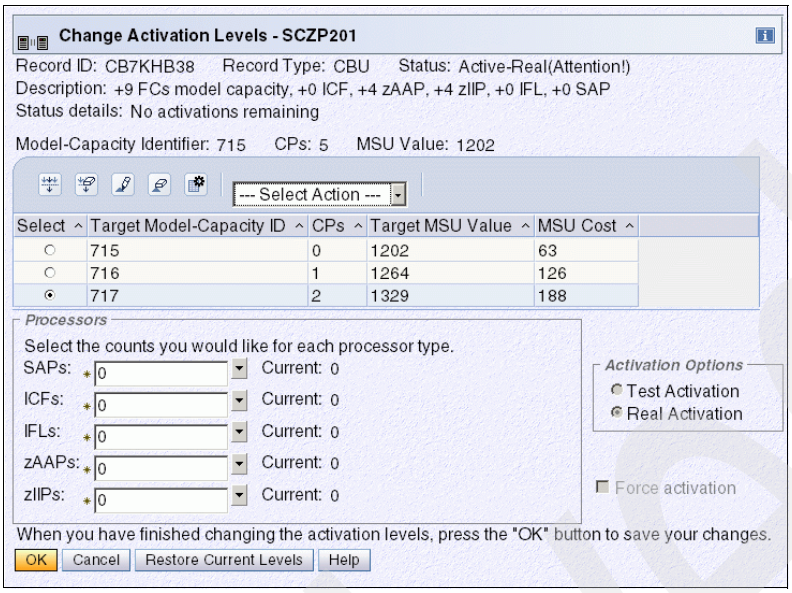

Figure 8-14 Change Activation Levels panel shows activated CPs in the CBU record

Figure 8-15 displays the processor configuration, after the two additional CPs are successfully added to the CBU. The Model-Capacity Identifier changes to 717 and the MSU value increases accordingly. No more PUs are available. The Model-Temporary-Capacity Identifier remains unchanged.

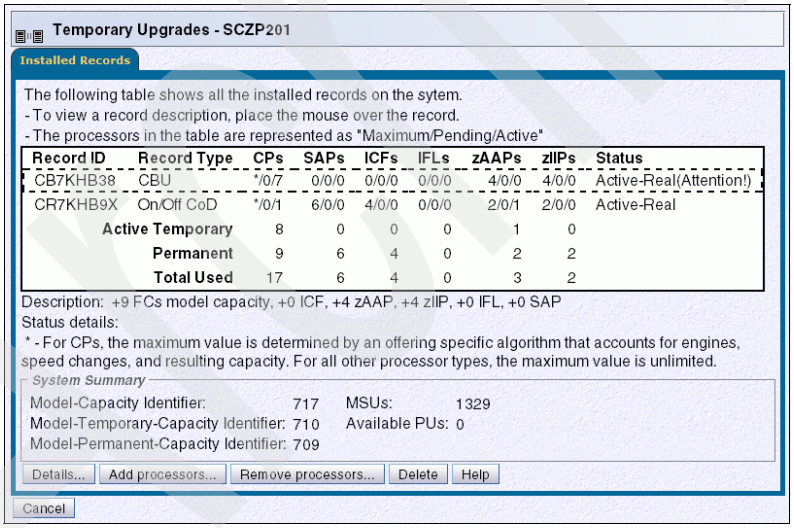

Figure 8-15 Temporary Upgrades panel shows two additional CPs for CBU

Two more CPs in partition A04 can now be brought online. The RMF report in Figure 8-16 on page 222 shows seven CPs online in partition A04.

```
Partition:   A01        2097 Model 717
CPC Capacity:    1329   Weight % of Max: ****   4h Avg:  270   Group:   N/A
Image Capacity:   781   WLM Capping %:    0.0   4h Max:  591   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        46.0                      0.1    97.4   97.5
A0E            0    74  NO   2.0      47.1    47.1    0.0     5.5    5.5
A0F            0     0  NO   2.0       0.1     0.2    0.0     0.0    0.0
A01            0   755  NO  10.0      96.6    96.6    0.0    56.8   56.8
A02            0     0  NO   4.0       0.0     0.0    0.0     0.0    0.0
A04            0     0  NO   7.0       0.0     0.0    0.0     0.0    0.0
A1E            0    74  NO   2.0      47.1    47.1    0.0     5.5    5.5
A1F            0     0  NO   2.0       0.1     0.1    0.0     0.0    0.0
A11            0   392  NO   9.0      55.7    55.7    0.0    29.5   29.5
A12            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0       0.1     0.1    0.0     0.0    0.0
PHYSICAL                                              0.1            0.1
```

Figure 8-16 RMF output after the additional resources had been online in LPAR A04.

8.1.4 Deactivating On/Off CoD

When resources are not needed anymore, the On/Off CoD record can be deactivated. The record remains installed for future use. Before deactivating physical resources, configure the related logical processors offline from the logical partitions, as shown in Figure 8-17 on page 223.

The zAAP and zIIP processors, which were configured offline and removed from On/Off CoD (in 8.1.3, “Changing CBU and On/Off CoD activation levels” on page 219), are not available in logical partition A01. Use the CONFIG command to configure the remaining On/Off CoD resources, which are one CP(0D) and one zAAP (0F), offline from z/OS.

```
CF CPU(0D,0F),OFFLINE

D M=CPU
IEE174I 14.19.14 DISPLAY M 907
PROCESSOR STATUS
ID  CPU                 SERIAL
00  +                    01DE502097
01  +                    01DE502097
02  +                    01DE502097
03  +                    01DE502097
04  +A                   01DE502097
05  +A                   01DE502097
06  +I                   01DE502097
07  +I                   01DE502097
08  +                    01DE502097
09  +                    01DE502097
0A  +                    01DE502097
0B  +                    01DE502097
0C  +                    01DE502097
0D  -
0E  -A
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.717.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-
N NOT AVAILABLE
```

Figure 8-17 Processors in partition taken offline

You can now deactivate the On/Off CoD record. Refer to Figure 8-18 on page 224. All specialty engines are set to zero, CP is set to 716. You can do this manually or by clicking **Undo**, which sets all values for record deactivation.

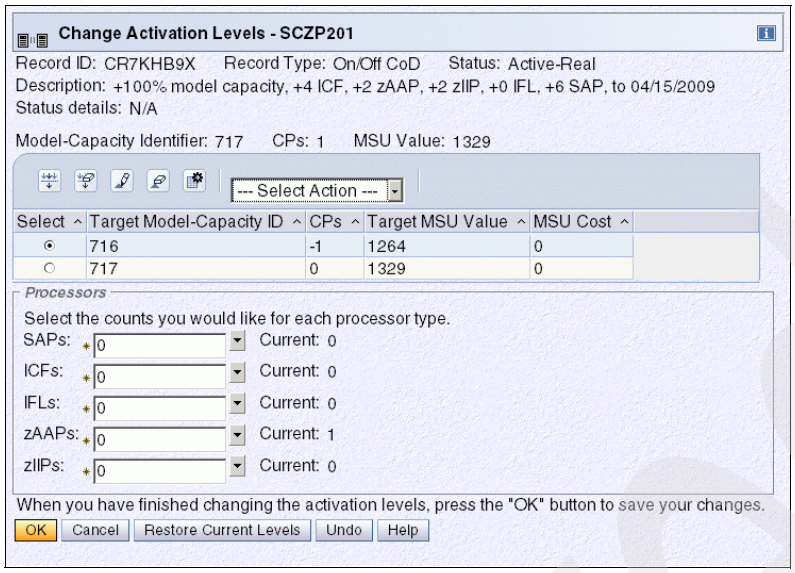

Figure 8-18 Change Activation Levels panel allows deactivation of the record

After a successful deactivation of the On/Off CoD record, the panel displays the results of the configuration, as shown Figure 8-19. The Model-Temporary-Capacity Identifier is the same as the Model-Permanent-Capacity Identifier because no billing capacity is active. The replacement capacity from CBU is still active.

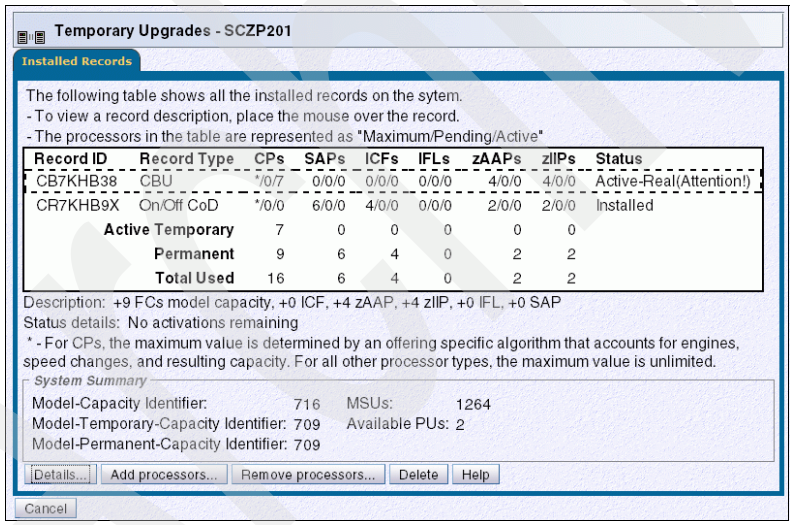

Figure 8-19 Temporary Upgrades panel shows the configuration after deactivation

The D M=CPU command also shows that no additional On/Off CoD resources are available in partition A01 after the On/Off CoD was deactivated; see Figure 8-20.

```
D M=CPU
IEE174I 14.24.12 DISPLAY M 909
PROCESSOR STATUS
ID  CPU                  SERIAL
00  +                     01DE502097
01  +                     01DE502097
02  +                     01DE502097
03  +                     01DE502097
04  +A                    01DE502097
05  +A                    01DE502097
06  +I                    01DE502097
07  +I                    01DE502097
08  +                     01DE502097
09  +                     01DE502097
0A  +                     01DE502097
0B  +                     01DE502097
0C  +                     01DE502097
0D  N
0E  NA
0F  NA
10  NI
11  NI

CPC ND = 002097.E26.IBM.02.00000001DE50
CPC SI = 2097.716.IBM.02.000000000001DE50
CPC ID = 00
CPC NAME = SCZP201
LP NAME = A01       LP ID =  1
CSS ID  = 0
MIF ID  = 1

+ ONLINE    - OFFLINE    . DOES NOT EXIST    W WLM-
N NOT AVAILABLE
```

Figure 8-20 No more additional Capacity on Demand resources available in A01

The RMF output shown in Figure 8-21 on page 226 confirms the same. The Model-Capacity Identifier is changed to 716, corresponding to the MSU values in CPC, and Image Capacity fields show the current capacity available. Partition A01 uses nine CPs, two zAAPs, and two zIIPs as before the On/Off CoD record activation.

```
Partition:   A01        2097 Model 716
CPC Capacity:    1264   Weight % of Max: ****   4h Avg:  298   Group:   N/A
Image Capacity:   790   WLM Capping %:    0.0   4h Max:  675   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                       45.0                      0.1     102    102
A0E            0    89  NO   2.0       56.4    56.4    0.0     7.1    7.1
A0F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A01            0   711  NO   9.0        100     100    0.0    56.2   56.2
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A04            0     0  NO   7.0        0.0     0.0    0.0     0.0    0.0
A1E            0    89  NO   2.0       56.4    56.4    0.0     7.1    7.1
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   396  NO   9.0       55.7    55.7    0.0    31.3   31.3
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                              0.1            0.1

*AAP                      4.0                       0.0     0.0    0.0
A01                  NO   2.0        0.0     0.0    0.0     0.0    0.0
A11                  NO   2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                         0.0            0.0

*ICF                      4.0                       0.0     100    100
A33                       4.0        100     100    0.0     100    100
PHYSICAL                                              0.0            0.0

*IIP                      4.0                       0.2     0.0    0.2
A01                  NO   2.0        0.0     0.0    0.0     0.0    0.0
A11                  NO   2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                              0.2            0.2
```

Figure 8-21 RMF output after On/Off CoD deactivation

When CBU resources are not needed anymore, the CBU record can be deactivated. As explained in 8.1.3, “Changing CBU and On/Off CoD activation levels” on page 219, you should configure the additional resources offline before record deactivation.

In this case, the operating system running in logical partition A04 can be brought down in order to deactivate the partition. Then, you can deactivate the CBU record because no other partition uses its resources; see Figure 8-22 on page 227. Clicking **Undo** sets all values for a record deactivation, or you can enter the values manually.

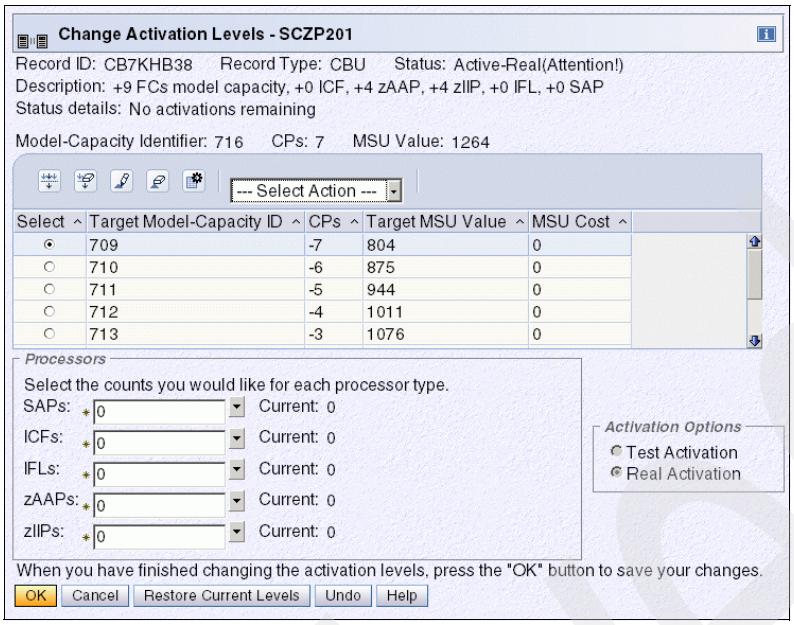

Figure 8-22 Change Activation Levels panel allows deactivation of records

Figure 8-23 shows that all identifiers are back to their original values and nine available PUs.

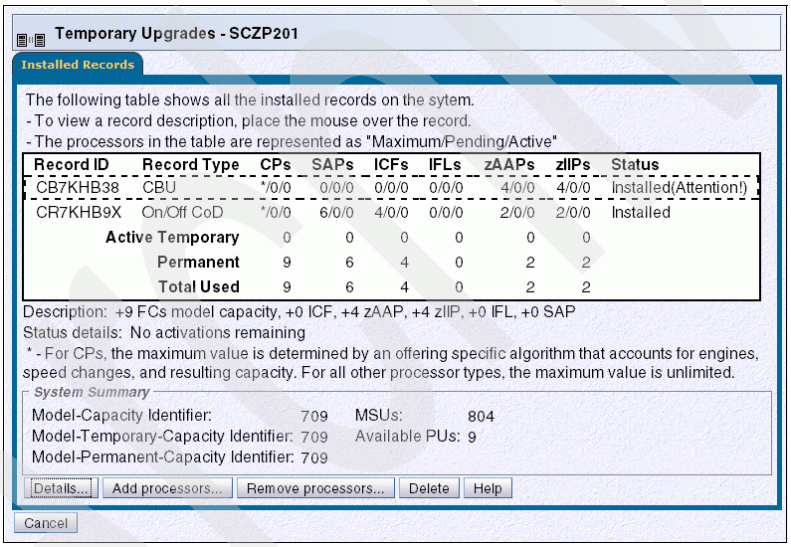

Figure 8-23 Temporary Upgrades panel shows a return to original values

Important: Remember to put resources that are used in logical partitions offline for those partitions, when physical resources are deactivated, otherwise there may be performance implications. Also, evaluate and adjust partition weights after every change.

Both CoD records are now deactivated. Although the On/Off CoD record is available for a future use, the CBU record is not usable because the *Real Activation* was used. To have a valid CBU available for future use, you have to either replenish the existing CBU record or purchase a new one.

8.2 On/Off CoD and CPE

This section describes an example similar to the one described in 8.1, “On/Off CoD and CBU” on page 212. Two temporary records are active at the same time. The On/Off CoD record is active and a replacement capacity for a planned event is needed.

CBU and CPE are both replacement offerings that restore capacities lost in another part of an enterprise. Because they are very similar offerings, the information in this section explains only the significant steps. The details are described in 8.1, “On/Off CoD and CBU” on page 212.

Planning is always essential. Refer to Chapter 2, “Understanding the Capacity on Demand environment” on page 17 for planning information.

Figure 8-24 shows the sequence of events performed in the example. The On/Off CoD record is activated first and then the CPE. Because not enough resources are available to cover lost capacity, some resources from On/Off CoD are deactivated and CPE resources are added. When replacement capacity is not needed, CPE is deactivated.

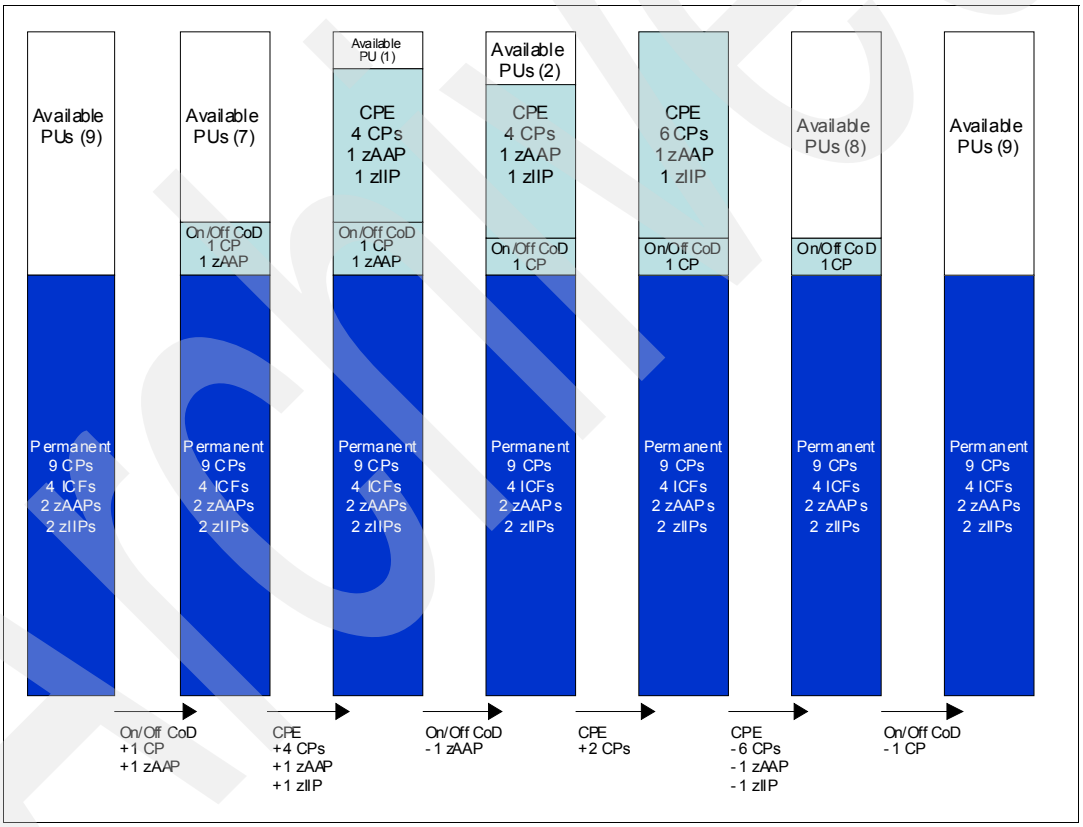

Figure 8-24 Activation - Change activation levels - Deactivation sequence

8.2.1 Activation sequence

On/Off CoD and CPE can be activated in any sequence because each is a separate record. However, for this scenario, the On/Off CoD is activated before the CPE record. Activation levels of each can be changed dynamically. Although many reasons exist for changing an activation level, changing it for On/Off CoD is more common. On/Off CoD is a net additional capacity and has an ongoing cost associated with it. Both records (On/Off CoD and CPE) are activated through the same path, which is the Perform Model Upgrade task on the SE.

8.2.2 Activating On/Off CoD

In this example, the peak workload period requires one additional CP and one additional zAAP. These processor resources are activated as Figure 8-25 shows. Current capacity setting is 710. Because On/Off CoD is a billable capacity, the Model-Capacity Identifier and the Model-Temporary-Capacity Identifier are modified accordingly.

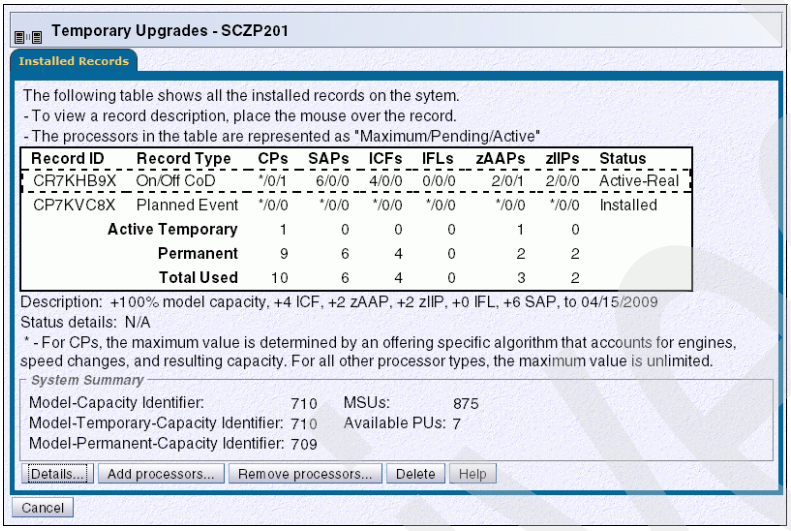

Figure 8-25 Temporary Upgrades panel shows activated processor resources

Additional processors are brought online in the z/OS image. The RMF output in Figure 8-26 on page 230 displays the new number of processors available in the associated logical partition A01, which has ten logical CPs, three zAAPs, and two zIIPs online.

```
Partition:   A01        2097 Model 710
CPC Capacity:     875 Weight % of Max: ****   4h Avg:  292   Group:   N/A
Image Capacity:   875 WLM Capping %:    0.0   4h Max:  362   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        39.0                      0.2    99.8   99.9
A0E            0    61  NO   2.0       34.8    34.8    0.0     7.0    7.0
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   262  NO  10.0       29.9    29.9    0.0    29.9   29.9
A02            0     0  NO   4.0        0.0     0.1    0.0     0.0    0.0
A1E            0    61  NO   2.0       34.8    34.8    0.0     7.0    7.0
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   489  NO   9.0       62.0    62.0    0.0    55.8   55.8
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1

*AAP                       5.0                       0.0    99.8   99.8
A01                    NO  3.0       59.5    59.5    0.0    49.8   49.8
A11                    NO  2.0        100     100    0.0    50.0   50.0
PHYSICAL                                                 0.0            0.0
*ICF                       4.0                       0.0     100    100
A33                        4.0        100     100    0.0     100    100
PHYSICAL                                                 0.0            0.0

*IIP                       4.0                       0.1     0.0    0.1
A01                    NO  2.0        0.0     0.0    0.0     0.0    0.0
A11                    NO  2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                                 0.1            0.1
```

Figure 8-26 RMF output after On/Off CoD processors had been configured online

While the On/Off CoD is active, the planned event also demands additional resources, which are provided by CPE. When CPE is activated, both On/Off CoD and CPE active at the same time, as shown in Figure 8-27 on page 231. It indicates that almost all resources have been activated. A CPE record allows you to activate all available PUs. However, activating only a subset of these resources is also possible, as follows:

- CPE activates four more CPs, one zAAP, and one zIIP.
- Permanent capacity is 709.
- On/Off CoD activated one CP, indicated by Model-Temporary-Capacity Identifier 710.
- CPE added four CPs, so the Model-Capacity Identifier is 714.
- One more PU is still available.

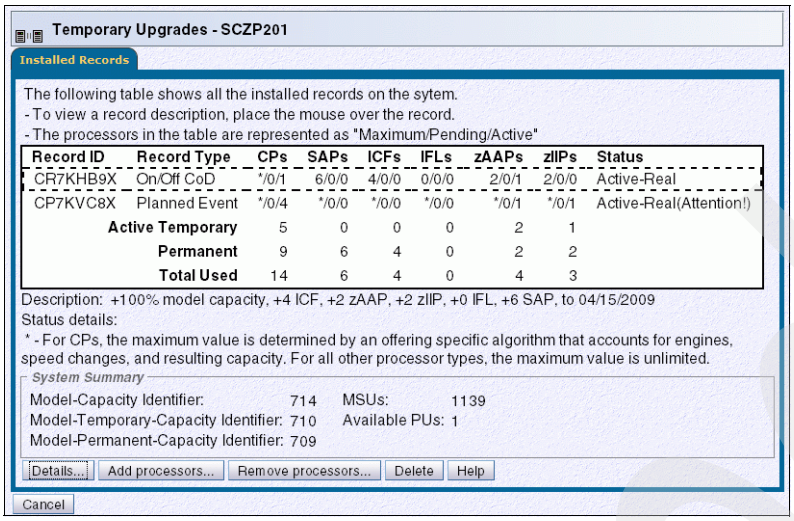

Figure 8-27 Two temporary upgrades active

Partition A04 can be activated now, which will serve the replaced workload. It uses dedicated processors, which means partition weights do not have to be adjusted. The RMF report in Figure 8-28 on page 232 presents all active partitions and their usage of processors. A04 is already active with four CPs, one zAAP, and one zIIP. The Capacity setting is shown as 714.

```
Partition:   A01        2097 Model 714
CPC Capacity:    1139   Weight % of Max: ****   4h Avg:  286   Group:   N/A
Image Capacity:   814   WLM Capping %:    0.0   4h Max:  362   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                      43.0                        0.1    97.3   97.5
A0E            0   112  NO   2.0       68.6    68.6    0.0     9.8    9.8
A0F            0     0  NO   2.0        0.2     0.2    0.0     0.0    0.0
A01            0   302  NO  10.0       37.1    37.1    0.0    26.5   26.5
A02            0     0  NO   4.0        0.0     0.1    0.0     0.0    0.0
A04            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0   112  NO   2.0       68.6    68.6    0.0     9.8    9.8
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   582  NO   9.0       79.5    79.5    0.0    51.1   51.1
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                0.1            0.1

*AAP                     6.0                         0.1    99.0   99.1
A01                  NO  3.0       65.4    65.4    0.0    49.0   49.0
A04                  NO  1.0        0.0     0.0    0.0     0.0    0.0
A11                  NO  2.0        100     100    0.0    50.0   50.0
PHYSICAL                                                0.1            0.1

*ICF                     4.0                         0.0     100    100
A33                      4.0        100     100    0.0     100    100
PHYSICAL                                                0.0            0.0

*IIP                     5.0                         0.1     0.0    0.1
A01                  NO  2.0        0.0     0.0    0.0     0.0    0.0
A04                  NO  1.0        0.0     0.0    0.0     0.0    0.0
A11                  NO  2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                                0.1            0.1
```

Figure 8-28 RMF output, after LPAR A04 was activated

Now a need arises to add resources to A04, for example two more CPs. Only one more PU is available. The second PU must be obtained from On/Off CoD. One zAAP is deactivated from the On/Off CoD record dynamically.

First, before removing it from the On/Off CoD, configure the zAAP offline from z/OS by using the following z/OS command:

```
CF CPU(xx),OFFLINE
```

For reasons why this is recommended, see 8.1.3, “Changing CBU and On/Off CoD activation levels” on page 219.

Figure 8-29 on page 233 shows the new processor configuration after the change. Note that the Model-Temporary-Capacity Identifier remained unchanged, although a processor was removed from the On/Off CoD. The removed processor was a zAAP, which is not reflected by any model identifier. The Model-Capacity Identifier is changed to 716 because of additional CPs that are made available by the CPE.

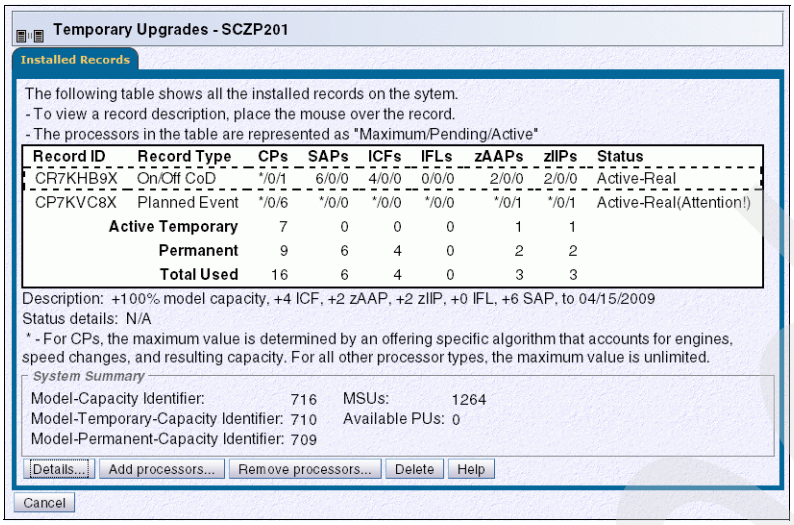

Figure 8-29 Temporary Upgrades panel shows modified configuration

After the additional CPs are available, they can be configured online in the A04 partition. They are dedicated to partition A04. Figure 8-30 on page 234 shows six CPs, one zAAP, and one zIIP active in logical partition A04.

```
Partition:   A01        2097 Model 716
CPC Capacity:    1264   Weight % of Max: ****   4h Avg:  282   Group:   N/A
Image Capacity:   790   WLM Capping %:    0.0   4h Max:  332   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                          45.0                     0.1    95.8   95.9
A0E            0   134  NO   2.0       84.7    84.7    0.0    10.6   10.6
A0F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A01            0   305  NO  10.0       38.6    38.6    0.0    24.1   24.1
A02            0     0  NO   4.0        0.0     0.1    0.0     0.0    0.0
A04            0     0  NO   6.0        0.0     0.0    0.0     0.0    0.0
A1E            0   134  NO   2.0       84.7    84.8    0.0    10.6   10.6
A1F            0     0  NO   2.0        0.1     0.1    0.0     0.0    0.0
A11            0   637  NO   9.0       89.6    89.6    0.0    50.4   50.4
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.0
PHYSICAL                                                0.1            0.1

*AAP                         6.0                      0.1    99.8    100
A01                     NO   3.0       48.7    48.7    0.0    48.7   48.7
A04                     NO   1.0        0.0     0.0    0.0     0.0    0.0
A11                     NO   2.0       76.8    76.8    0.0    51.2   51.2
PHYSICAL                                                0.1            0.1

*ICF                         4.0                      0.0     100    100
A33                          4.0        100     100    0.0     100    100
PHYSICAL                                                0.0            0.0

*IIP                         5.0                      0.1     0.0    0.1
A01                     NO   2.0        0.0     0.0    0.0     0.0    0.0
A04                     NO   1.0        0.0     0.0    0.0     0.0    0.0
A11                     NO   2.0        0.0     0.0    0.0     0.0    0.0
PHYSICAL                                                0.1            0.1
```

Figure 8-30 RMF output

After the additional workload related to the planned event is moved back to the original server, resources activated by CPE can be deactivated. Following the deactivation of partition A04, which used CPE resources, the entire CPE feature is deactivated. Figure 8-31 on page 235 shows Temporary Upgrade panel after CPE is deactivated.

Note: The CPE feature is only valid for a single activation, which can last up to 72 hours. After being deactivated, the CPE feature cannot be used anymore. Changing the CPE activation level (adding or removing processor resources) is possible if at least one processor stays active. Removing the last processor means the deactivation of the entire CPE feature. For details, see Chapter 4, “Capacity for Planned Events” on page 105.

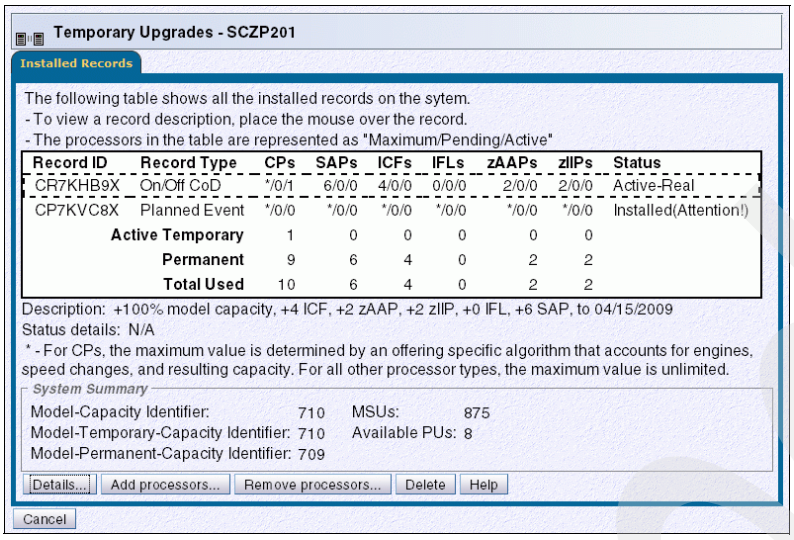

Figure 8-31 Temporary Upgrades panel shows configuration after CPE is deactivated

If resources are needed again from On/Off CoD, they can be activated again. When there is no more demand for temporary resources, the whole On/Off CoD record can be deactivated. To deactivate it, configure all associated processors offline from the operating systems that use them, and then deactivate the record.

Figure 8-32 shows the result after all temporary resources are deactivated. Again, nine PUs are available, and all three Capacity Identifiers show the same value 709, which is a permanent capacity setting.

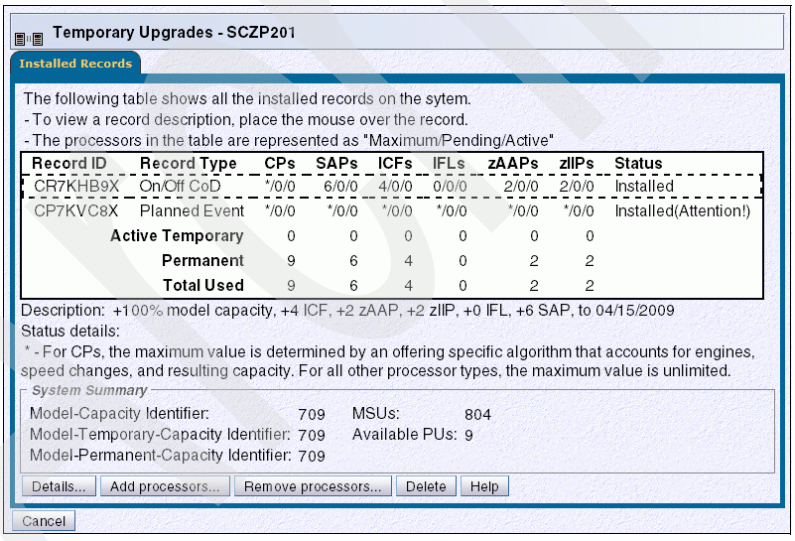

Figure 8-32 Temporary Upgrades panel indicates all temporary resources have been deactivated

The RMF report in Figure 8-33 on page 236 confirms that all CoD-associated processor resources were deactivated successfully. Only the original processors drive logical partition A01 again, and nine CPs, two zAAPs, and two zIIPs.

```
Partition:   A01        2097 Model 709
CPC Capacity:     804   Weight % of Max: ****   4h Avg:  280   Group:   N/A
Image Capacity:   804   WLM Capping %:    0.0   4h Max:  332   Limit:   N/A

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total

*CP                        38.0                      0.2    99.8    100
A0E            0    34  NO   2.0       18.9    18.9    0.0     4.2    4.2
A0F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A01            0   333  NO   9.0       41.4    41.4    0.0    41.4   41.4
A02            0     0  NO   4.0        0.0     0.0    0.0     0.0    0.0
A1E            0    34  NO   2.0       18.9    18.9    0.0     4.2    4.2
A1F            0     0  NO   2.0        0.1     0.2    0.0     0.0    0.0
A11            0   401  NO   9.0       49.9    49.9    0.0    49.9   49.9
A12            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
A17            0     0  NO   4.0        0.1     0.1    0.0     0.0    0.1
PHYSICAL                                                 0.2            0.2

*AAP                       4.0                       0.1    99.8    100
A01                    NO  2.0       74.7    74.8    0.0    49.8   49.8
A11                    NO  2.0       75.1    75.1    0.0    50.0   50.0
PHYSICAL                                                 0.1            0.1
*ICF                       4.0                       0.0     100    100
A33                        4.0        100     100     0.0     100    100
PHYSICAL                                                 0.0            0.0

*IIP                       4.0                       0.2     0.0    0.2
A01                    NO  2.0        0.0     0.0     0.0     0.0    0.0
A11                    NO  2.0        0.0     0.0     0.0     0.0    0.0
PHYSICAL                                                 0.2            0.2
```

Figure 8-33 RMF output showing no more CoD resources available

8.3 CBU and On-line Permanent Upgrade

With the flexibility of the Capacity on Demand design, it is possible to have temporary capacity active and perform a permanent upgrade at the same time. This section provides an example of performing a permanent upgrade when a CBU is active.

When a permanent upgrade is performed while a CBU upgrade is active, the permanent upgrade is added to the permanent capacity and the CBU resources are applied on top of the new permanent capacity. The replacement capacity resources are not converted to permanent resources. If the permanent upgrade is attempted and the needed resources are not available, the upgrade is blocked from being installing until the condition is corrected. For more detail, see 2.2.2, “Permanent upgrade with temporary offerings active” on page 39.

Figure 8-34 shows the steps performed in this scenario. First, the CBU record activates two temporary CPs. Then, a permanent upgrade with one CP is installed. The figure also shows the capacity identifiers associated with each step performed and how they change.

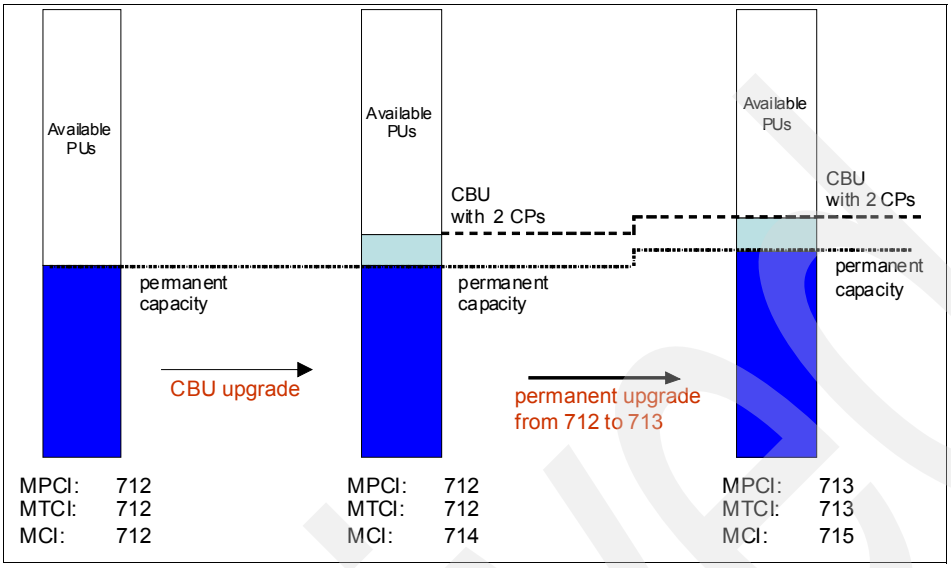

Figure 8-34 Permanent Upgrade with CBU Active

Activation consists of two parts. First, a CBU record that contains two CPs is activated. Next an On-line Permanent Upgrade that contains a single CP of permanent capacity is installed.

8.3.1 Activating CBU

The first part of the scenario is to activate a CBU record with two CPs. For details on the CBU activation, see Chapter 3, “Capacity Backup” on page 81

The permanent capacity setting of the server is 712 and all capacity identifiers have the same value. When the CBU with two CPs is activated, and only the Model-Capacity Identifier is changed, it reflects new current capacity setting 714. Four PUs are still available; see Figure 8-35.

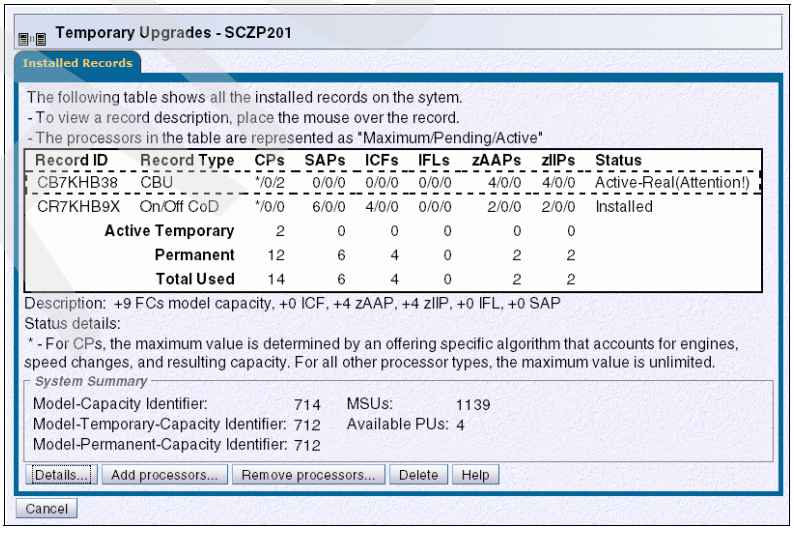

Figure 8-35 Temporary Upgrades panel shows a change in capacity setting

8.3.2 Installing Permanent upgrade

The On-line Permanent Upgrade is retrieved from the IBM Support System by using the Perform Model Conversion task on the SE. The **Permanent upgrades** option is selected, as shown in Figure 8-36.

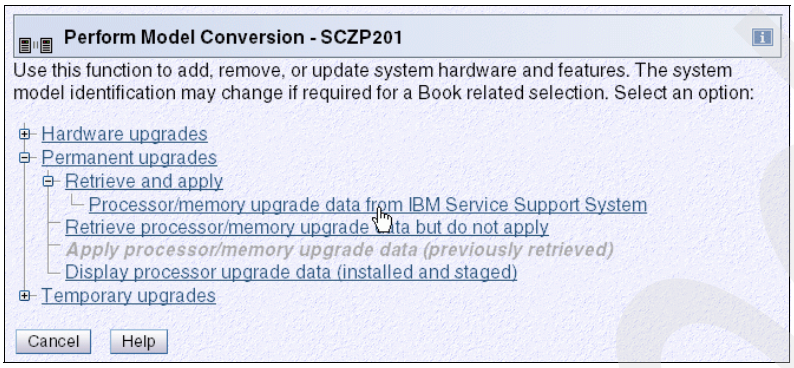

Figure 8-36 Perform Model Conversion

Retrieving a permanent upgrade and choosing to apply it or not is possible. If a permanent upgrade has already been retrieved, it may be applied now.

Selecting the **Retrieve and apply** option opens the Customer Initiated Upgrade Order Activation Panel, which prompts for the order activation number, as shown in Figure 8-37. The order activation number is provided by Resource Link when the order is created.

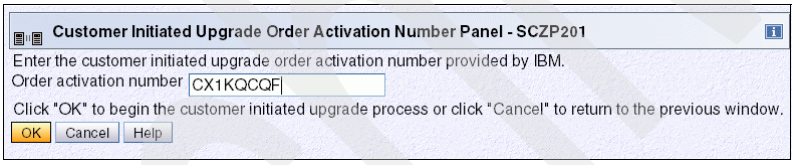

Figure 8-37 Customer Initiated Upgrade Order Activation Number

By clicking **OK**, the Permanent upgrades panel opens to confirm the original and new configurations. Figure 8-38 on page 239 shows only one difference in configuration, which is a change from capacity setting 712 to 713. It means that one CP will be added to the permanent configuration.

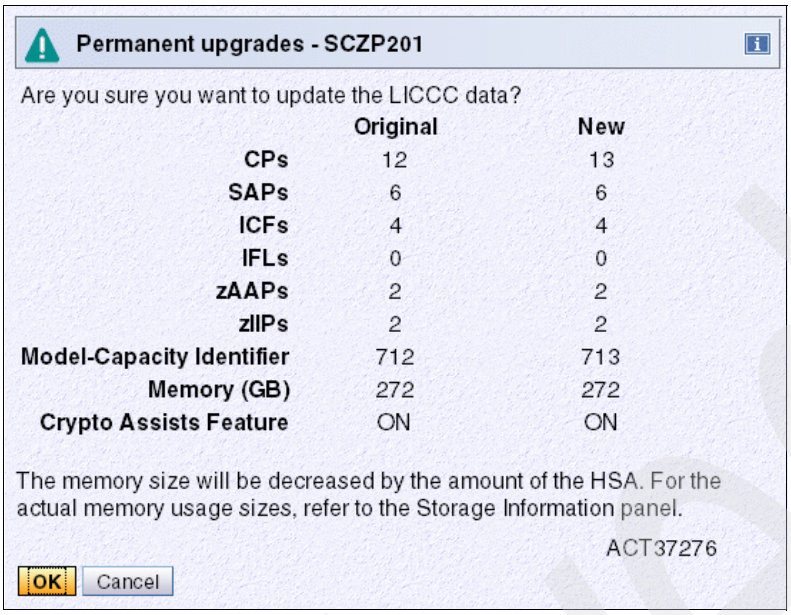

Figure 8-38 Permanent Upgrades, message ACT37276

After the permanent upgrade to 713 is performed, the Temporary Upgrades panel shows the updated number; see Figure 8-39. The Model-Permanent-Capacity Identifier reflects the new permanent capacity setting of 713. The Model-Temporary-Capacity Identifier is also 713, because no On/Off CoD record is active. The Model-Capacity Identifier is 715, indicating that the temporary capacity activated by the CBU record is still active with two CPs active. The number of available PUs is decreased by one, which is the one PU that was converted to permanent capacity.

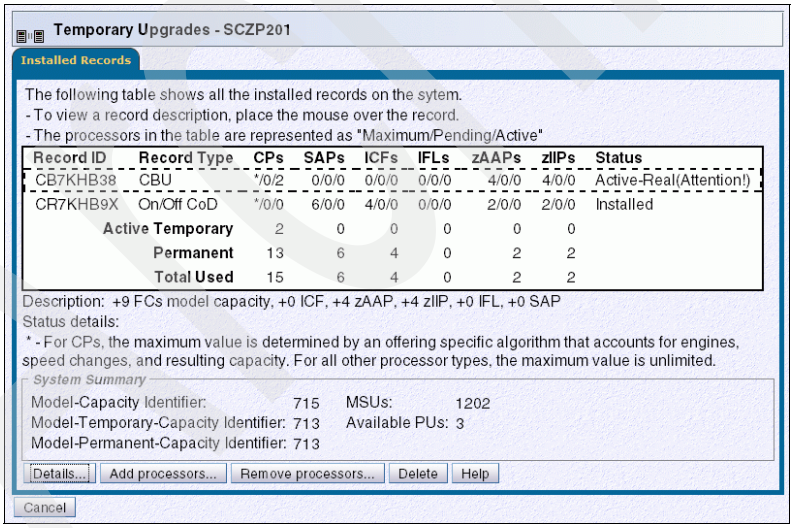

Figure 8-39 Temporary Upgrades panel shows changes after a permanent upgrade

As mentioned in 2.2.2, “Permanent upgrade with temporary offerings active” on page 39, a permanent upgrade never converts a temporary replacement capacity to a permanent capacity. Enough PUs must be available to perform an upgrade, otherwise an error message is issued as shown in Figure 8-40 on page 240.

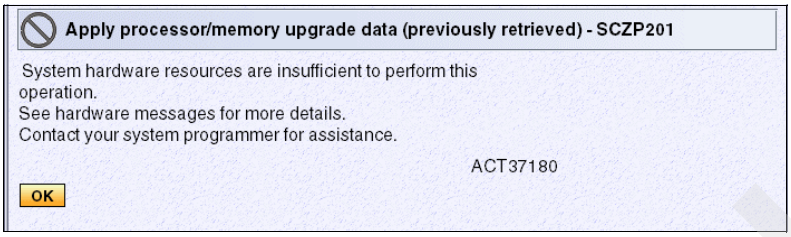

Figure 8-40 Permanent upgrade failed

8.4 On/Off CoD and On-line Permanent Upgrade

With the flexibility of the Capacity on Demand design, it is possible to have temporary capacity active and perform a permanent upgrade at the same time. This section provides a scenario of performing a permanent upgrade when On/Off CoD is active.

When a permanent upgrade is performed while an On/Off CoD upgrade is active, the permanent upgrade first converts billable capacity to permanent and then uses available PUs. More details are in 2.2.2, “Permanent upgrade with temporary offerings active” on page 39.

Figure 8-41 shows the steps performed in this scenario. An On/Off CoD temporary record with two CPs is activated. Next, the On-line Permanent Upgrade, which increases the permanent capacity from 709 to 712, is installed and the processors that are activated by the On/Off CoD are converted to permanent processors. When all temporary capacity provided by On/Off CoD is converted, available PUs are used for permanent upgrade. The figure also shows capacity identifiers associated with each step performed and how they change

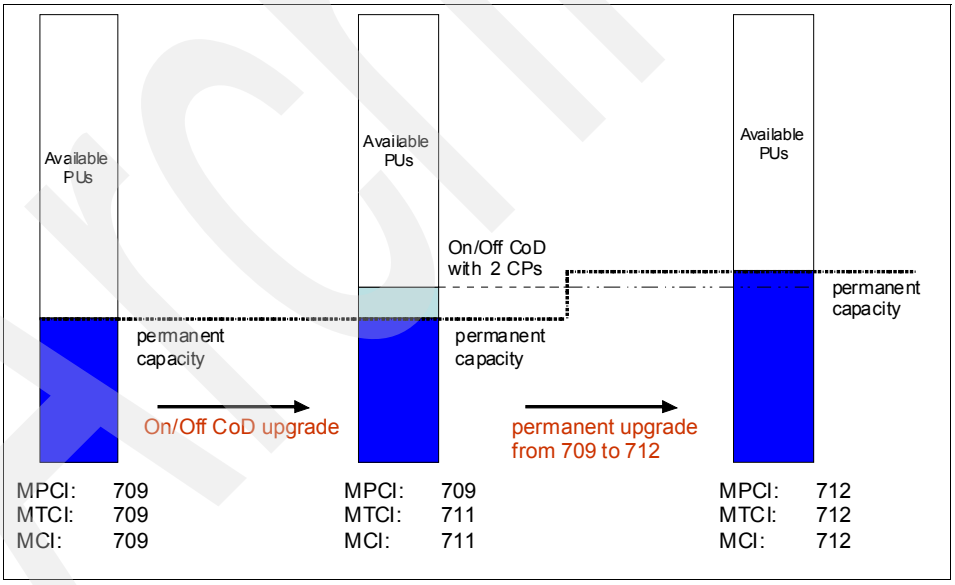

Figure 8-41 On-line Permanent Upgrade with On/Off CoD for two CPs active

Activation consists of two parts. First, an On/Off CoD record that contains two CPs of temporary capacity is activated. Next, an On-line Permanent Upgrade that contains three CPs of a permanent capacity is installed.

8.4.1 Activating On/Off CoD

The first part of the scenario activates an On/Off CoD record with two CPs. For details about activation of the On/Off CoD see Chapter 5, “On/Off Capacity on Demand” on page 123.

Figure 8-35 on page 237 shows a Temporary Upgrades panel with the On/Off CoD record active. Model-Permanent-Capacity Identifier is 709. Model-Temporary-Capacity Identifier is 711, because it reflects active billable capacity delivered by On/Off CoD. Model-Capacity Identifier is also 711 because no temporary replacement capacity is active. Seven PUs are available.

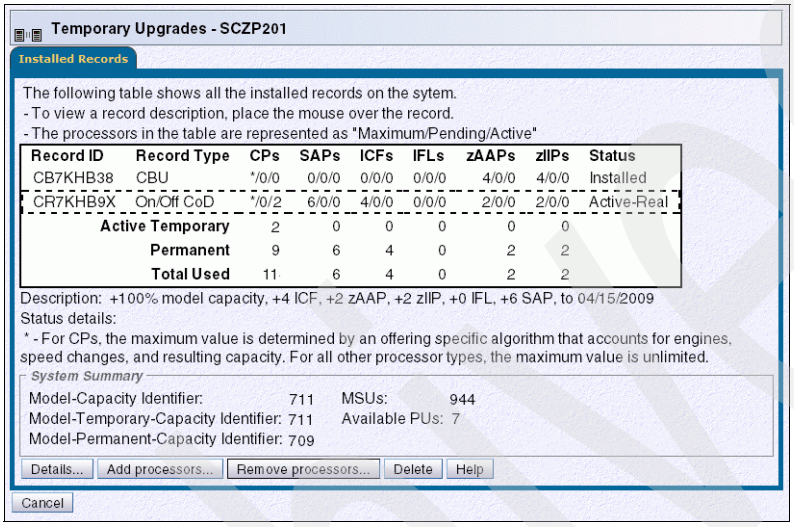

Figure 8-42 Temporary Upgrades shows active record

8.4.2 Activating On-line Permanent Upgrade

The On-line Permanent Upgrade is retrieved from the IBM Support Systems by using the Perform Model Upgrade task on SE. Because this is a permanent upgrade, the **Permanent upgrades** option is selected and expanded, as shown in Figure 8-43.

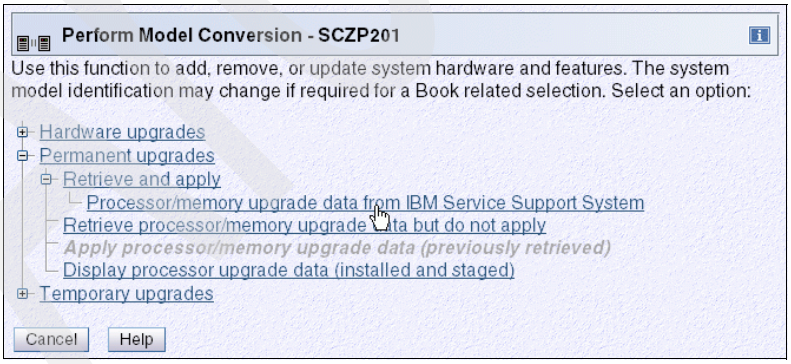

Figure 8-43 Perform Model Conversion

It is possible to retrieve a permanent upgrade and choose to apply it or not. If a permanent upgrade has already been retrieved, it may be applied now.

By selecting the **Retrieve and apply** option, the Customer Initiated Upgrade Order Activation panel opens and prompts for the order activation number as shown in Figure 8-44 on page 242. The order number is provided by Resource Link when the order is created.

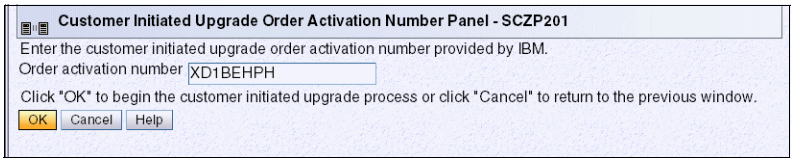

Figure 8-44 Customer Initiated Upgrade Order Activation Number

By clicking **OK**, the Permanent upgrades panel opens to confirm the original and new configurations. Figure 8-45 shows only one difference in the configuration, which is a change from 709 to 712. It means that three CPs will be added to a permanent configuration.

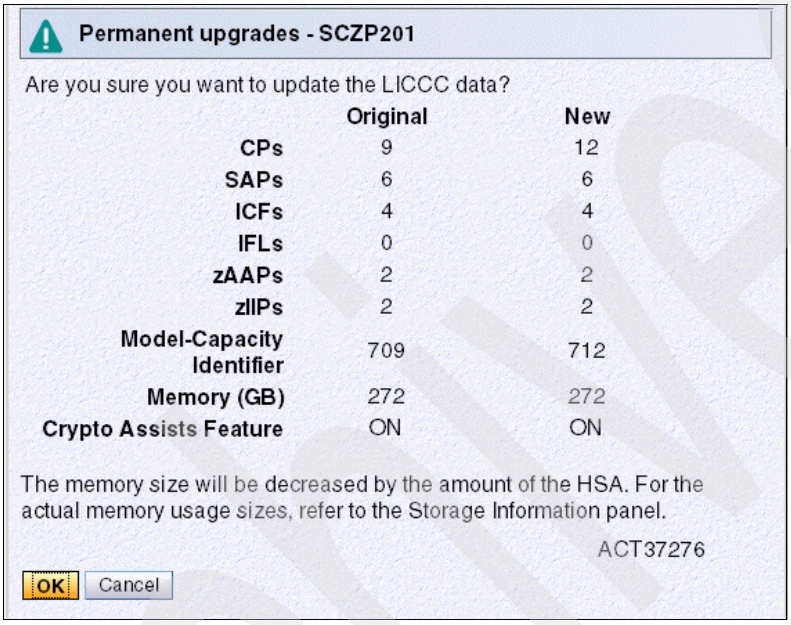

Figure 8-45 Permanent Upgrades, message ACT37276

When the upgrade completes successfully, a confirmation message indicates that the permanent upgrade was successful and that resources from an active temporary billable record were converted to fulfill the request. Because a permanent upgrade added CPs to a permanent configuration, active CPs from a temporary billable record were converted to permanent. See Figure 8-46.

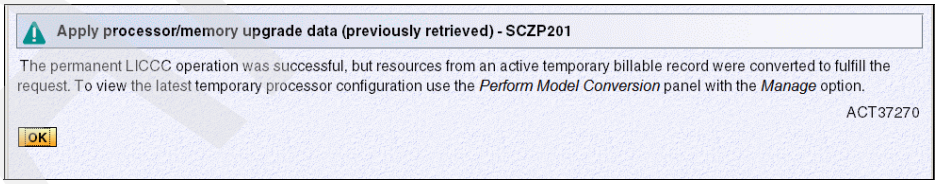

Figure 8-46 Apply processor/memory upgrade data, message ACT37270

The Temporary Upgrades panel in Figure 8-47 on page 243 shows all changes that were made. All capacity identifiers are now 712, which means no temporary billable and no temporary replacement capacity are active. However, the On/Off CoD record indicates the status is Active-Real. Close inspection of the On/Off CoD record reveals it is active, but no resources are active in it. It had two CPs active before and they were both converted to permanent CPs. Permanent upgrade was for three CPs, so one of the available PUs was converted too, and the number of available PUs is now 6. The On/Off CoD record should be manually deactivated if activating additional On/Off CoD resources is not needed at this time.

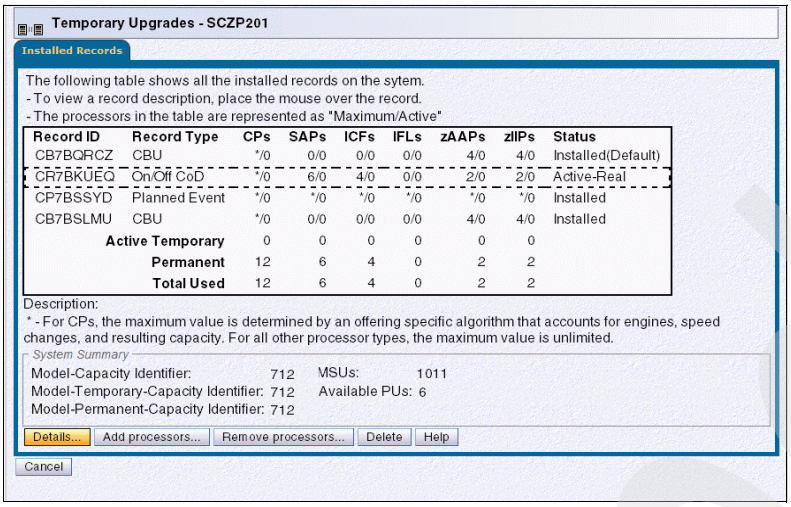

Figure 8-47 Temporary Upgrades panel

8.4.3 Deactivating On/Off CoD

The result of activities in the previous section is an active On/Off CoD record and upgraded permanent capacity. Although On/Off CoD is active, it uses no resources and can be deactivated.

8.5 CBU + On/Off CoD + CPE + CBU + On-line Permanent Upgrade

This section presents a scenario with four temporary records active at the same time, and a permanent upgrade activation while several temporary records are active. Although the order of the steps we perform in the scenario is unlikely in real situations, we show them to demonstrate the flexibility of Capacity on Demand.

All upgrades, temporary and permanent, are done concurrently and are not disruptive to the operating systems and applications. Furthermore, Capacity on Demand allows you to modify each record dynamically and independently while multiple records are active. This section briefly demonstrates the activation, modification, and deactivation of multiple records.

The hardware used for this example is a System z10 EC model E26, with the following permanent configuration (see Figure 8-48 on page 244):

- Capacity setting 705, which means 5 CPs
- 6 SAPs
- 4 ICFs
- 2 zAAPs
- 2 zIIPs
- 13 available PUs

All records are already retrieved to the Support Element and installed into four out of eight slots on System z10.

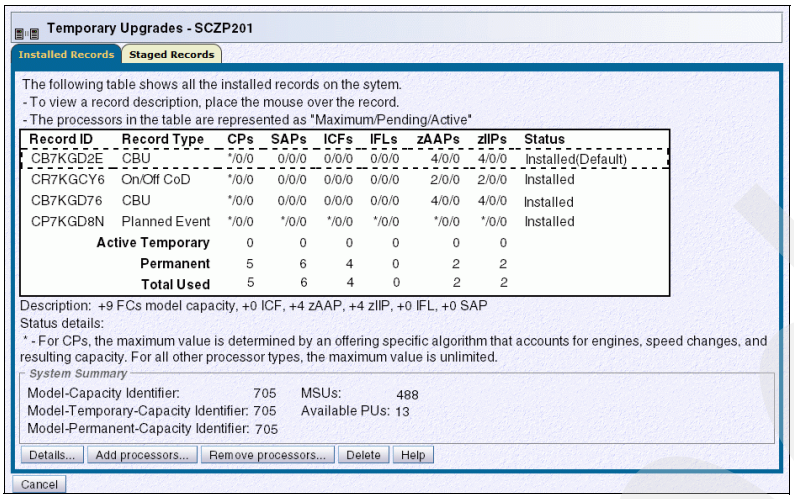

Figure 8-48 Temporary Upgrades panel shows installed records

Figure 8-49 shows the sequence of events. The multiple CoD records can be activated in any sequence, but obviously the order of activation determines the allocation of available resources.

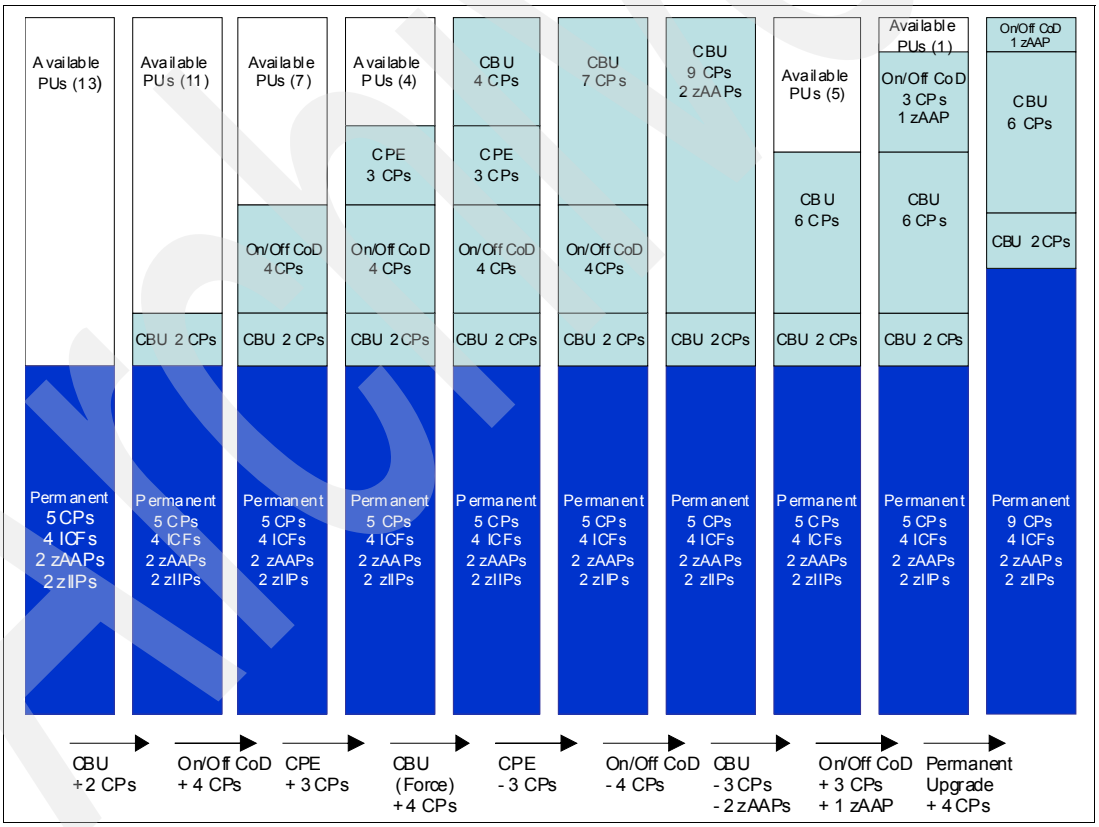

Figure 8-49 Sequence of activations

8.5.1 Activating the first CBU record

The first CBU activates two temporary CPs. The following list is the new configuration, also shown in Figure 8-50:

- The first CBU is active, all other records are inactive.
- Two temporary CPs are active.
- The number of available PUs decreased from 13 to 11.
- Model-Permanent-Capacity Identifier stays at 705.
- Model-Temporary-Capacity Identifier is also 705, because no billable capacity is active.
- Model-Capacity Identifier changed to 707 and reflects two temporary replacement CPs.

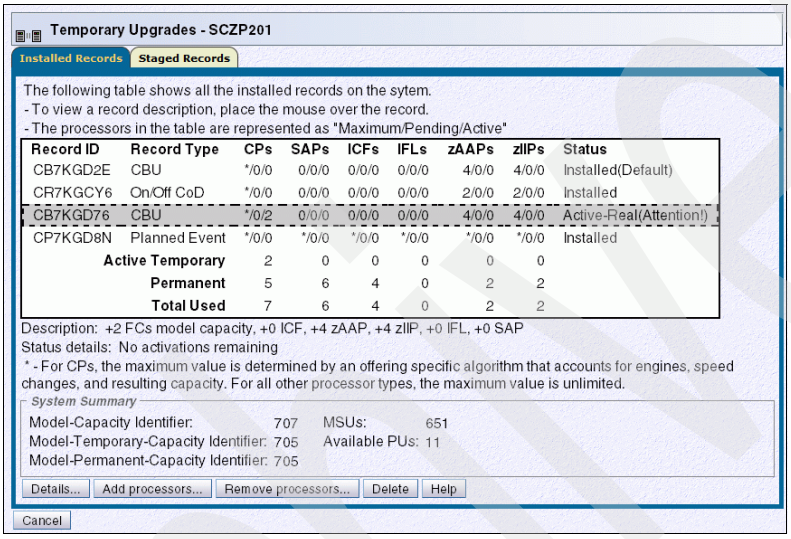

Figure 8-50 Additional CPs available after CBU activation

8.5.2 Activating the On/Off CoD record

The On/Off CoD record is now used to add four CPs to the configuration. Figure 8-51 on page 246 shows the following information:

- Six temporary CPs are active now.
- Two records are active; one provides four CPs and another provides two CPs.
- Seven PUs are available.
- Model-Permanent-Capacity Identifier is not changed, because the permanent capacity has not changed.
- Model-Temporary-Capacity Identifier is modified to 709. It indicates four CPs are added through billable temporary record (On/Off CoD).
- Model-Capacity Identifier is 711 because two CPs were also added by an active replacement capacity (CBU).

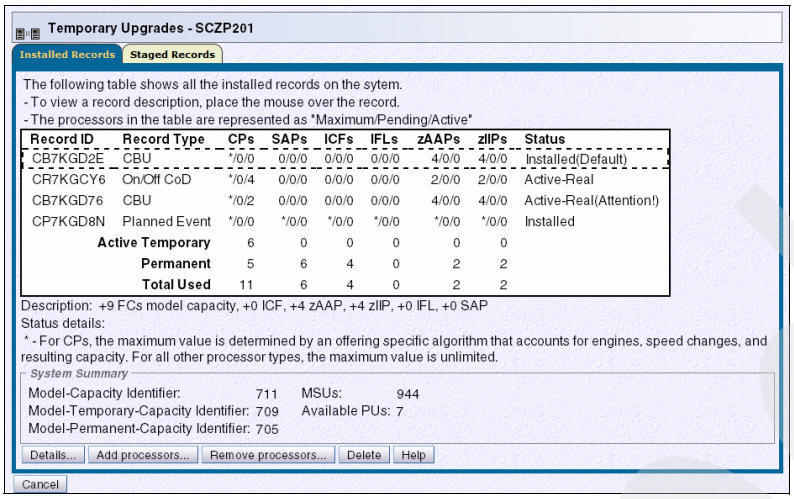

Figure 8-51 On/Off CoD record activated

8.5.3 Activating the CPE record

The CPE record is activated. It adds three additional CPs to the configuration. Figure 8-52 shows the updated processor configuration, including the extra CPs, as follows:

- Three records are active, providing 9 temporary CPs, and the number of available PUs has decreased to 4.
- Model-Permanent-Capacity Identifier stays at 705.
- Model-Temporary-Capacity Identifier stays at 709.
- Model-Capacity Identifier is increased to 714 because five replacement CPs are active (two activated by CBU and three activated by CPE).

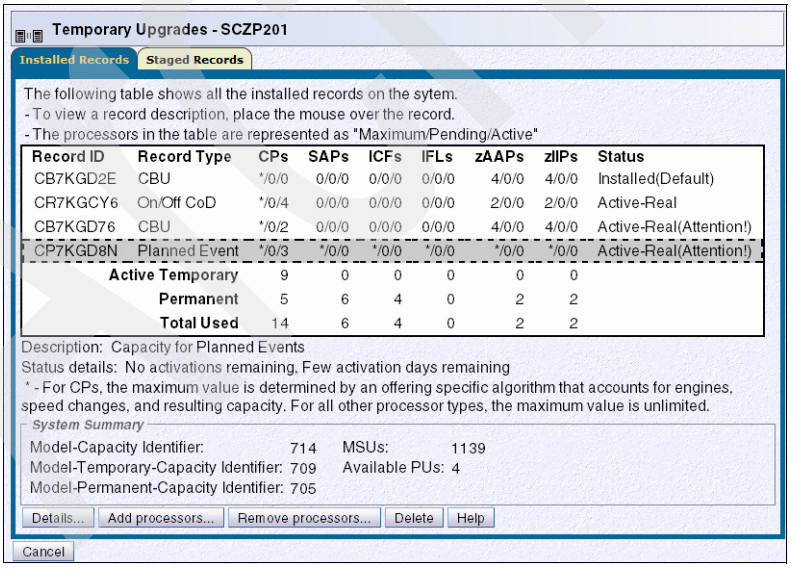

Figure 8-52 CPE record activated

8.5.4 Activating the second CBU record

The CBU with the Force option is activated. It is the same as though APIs used on systems prior to System z10 were used to activate a CBU. Because the API does not allow partial record activation, Force CBU is always used. In this case, it attempts to activate nine CPs, four zAAPs and four zIIPs. Because only four PUs are available, the record is activated partially and PUs are pending.

When using the Force option manually, a warning message (Figure 8-53) indicates that not enough PUs are available to activate the whole record.

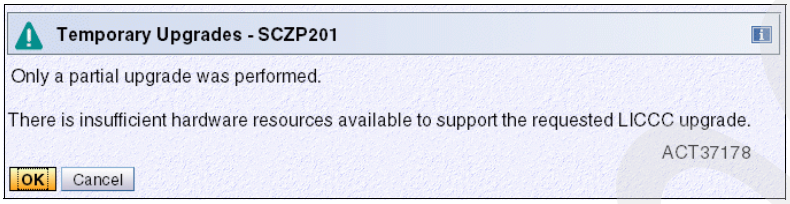

Figure 8-53 Not enough resource for forced CBU

Figure 8-54 shows an updated configuration, which indicates:

- No PUs are available.
- The second CBU activated four temporary CPs. Five CPs are pending activation and four zAAPs and four zIIPs.
- Model-Permanent-Capacity Identifier is still 705 because permanent capacity did not change.
- Model-Temporary-Capacity Identifier is still 709 because four billable capacity CPs are active from the On/Off CoD record.
- Model-Capacity Identifier is 718 because nine temporary CPs are activated by replacement capacity records. Two CPs are activated by the first CBU, three CPs by a CPE record and four CPs by the second CBU record.

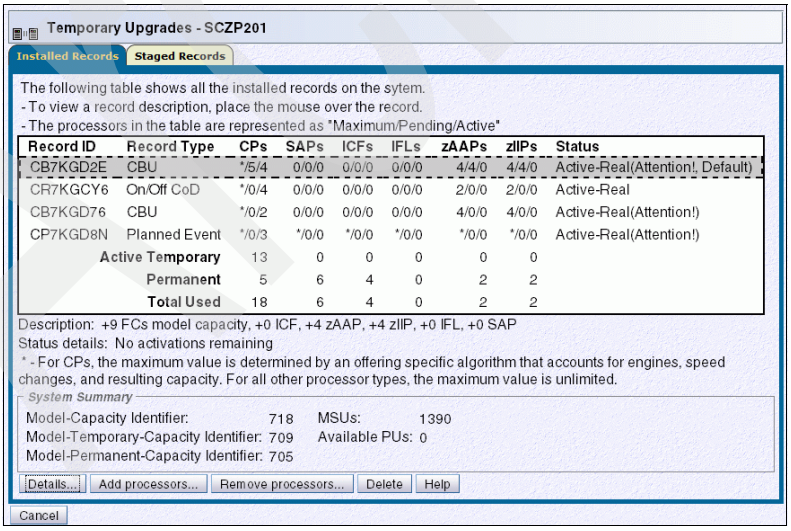

Figure 8-54 Temporary Upgrades panel shows an updated configuration

Note: PUs are pending activation, so performing a permanent upgrade is not possible now.

8.5.5 Deactivating CPE

Figure 8-55 shows a server configuration after the CPE record is deactivated. Three PUs released from the CPE were immediately taken by the CBU that has pending PUs. It now uses seven CPs instead of four, two CPs, four zAAPs and four zIIPs are still pending activation. The capacity identifiers are not changed.

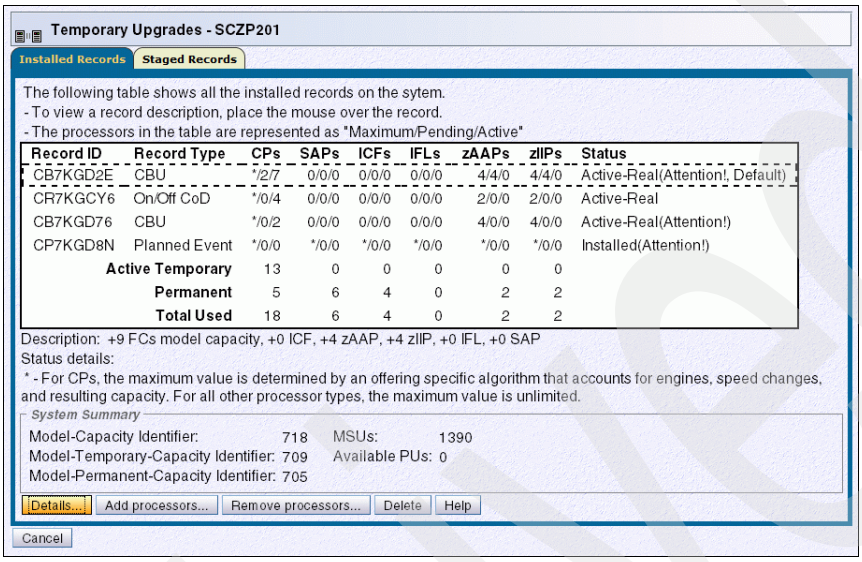

Figure 8-55 CPE record deactivated

8.5.6 Deactivating On/Off CoD

When an On/Off CoD record is deactivated, it releases its four CPs into the pool of available PUs. Again, all available PUs are immediately taken by pending PU activation from the second CBU record.

Figure 8-56 on page 249 shows the current configuration, as follows:

- Because all nine CPs from CBU record are active, two zAAPs can be activated now. If enough PUs were available, zAAPs would be activated and then zIIPs.
- Model-Temporary-Capacity Identifier was modified because no billable capacity is active now, therefore the identifier has the same value as the Model-Permanent-Capacity Identifier.
- The Model-Capacity Identifier shows a value of 716 because 11 CPs are activated by replacement capacity records.

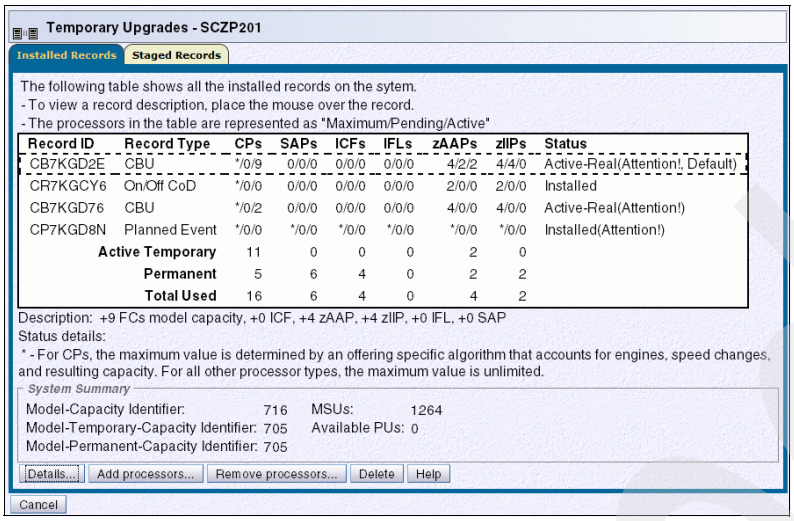

Figure 8-56 On/Off CoD deactivated

8.5.7 Changing the CBU activation level

The CBU record can be changed manually, as shown in Figure 8-57. Comparing the original and new configuration, three CPs and two zAAPs have been deactivated.

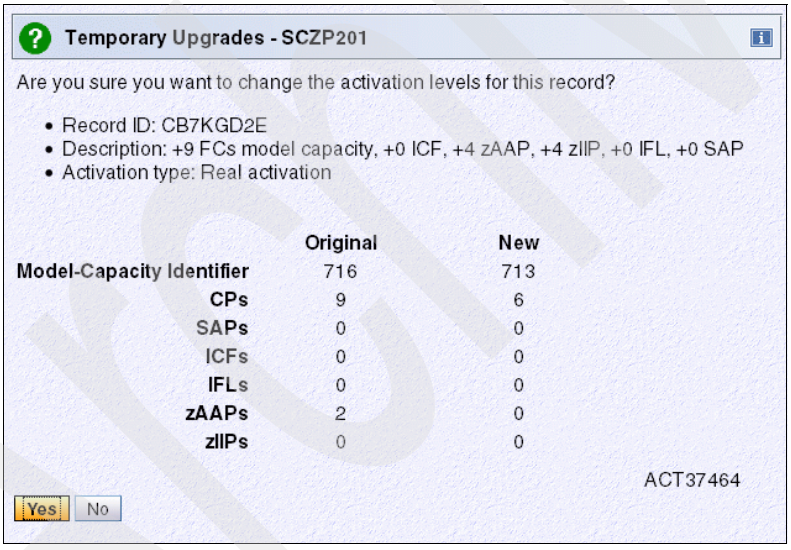

Figure 8-57 CBU record activation level change

Because the record with Force activation was modified, no more PUs are pending for this record. The current configuration in Figure 8-58 on page 250 shows five available PUs, which were released from the CBU record. The Model-Capacity Identifier indicates that only replacement capacity is active, bringing eight temporary CPs to the configuration.

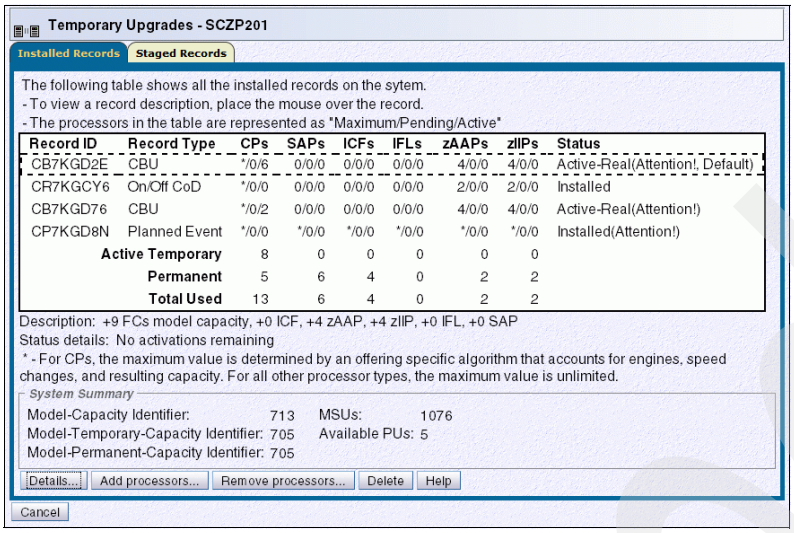

Figure 8-58 CBU activation levels changed

8.5.8 Activating the On/Off CoD

Now, when PUs are available, the On/Off CoD record can be activated again. This time, as shown in Figure 8-59, the configuration is as follows:

- The three CPs and one zAAP are activated.
- Model-Temporary-Capacity Identifier is set to 708. It represents three CPs activated by a billable capacity.
- Model-Capacity Identifier is 716 because eight replacement CPs are also active.
- One PU is available.

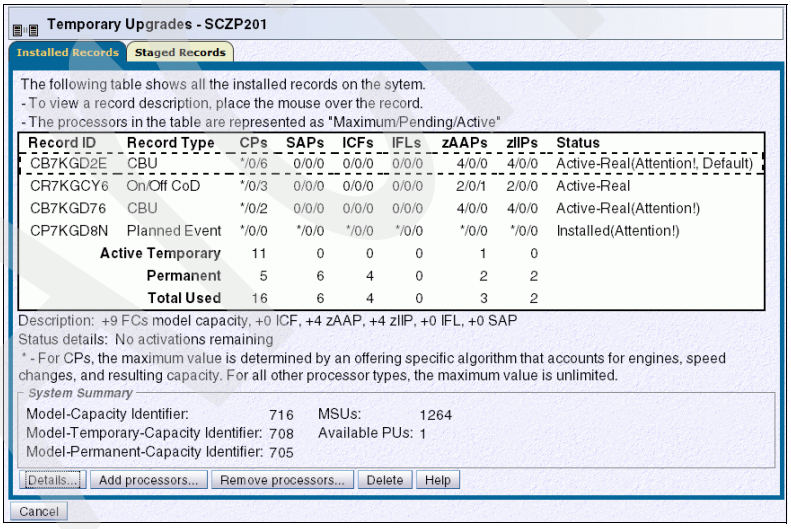

Figure 8-59 On/Off CoD activated

8.5.9 Performing a permanent upgrade

By performing a permanent upgrade from 705 to 709, four CPs have been added to the permanent configuration, as shown in Figure 8-60.

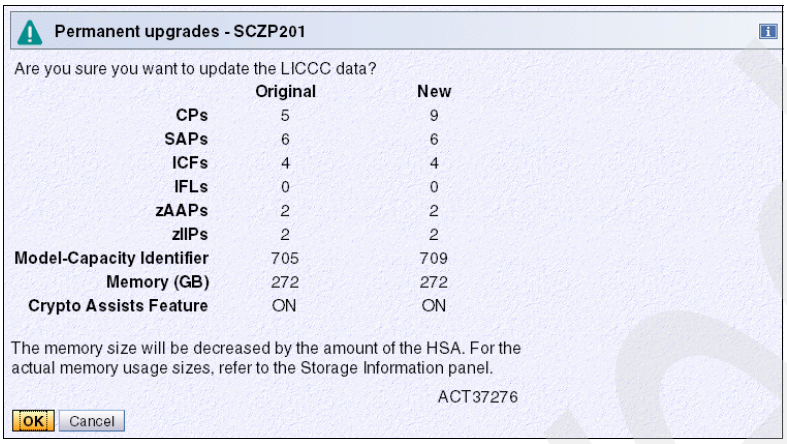

Figure 8-60 Permanent upgrade

After the permanent upgrade is performed, a message (shown in Figure 8-61) indicates temporary PUs from a billable capacity record were converted to permanent PUs.

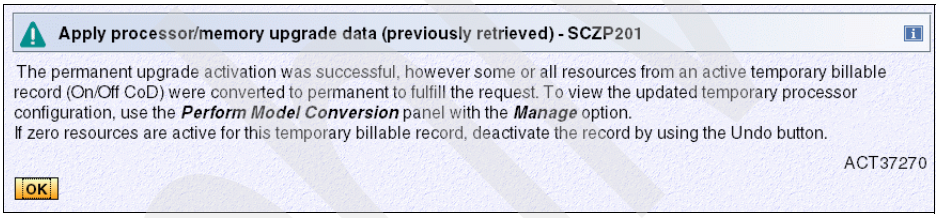

Figure 8-61 Permanent upgrade message

Figure 8-62 on page 252 shows the new configuration:

- The permanent upgrade uses all three CPs activated by the On/Off CoD record. Because the permanent upgrade added CPs only, it converted only CPs, not a zAAP activated by On/Off CoD. One more PU was needed for a permanent upgrade, so the last available PU was used.
- The On/Off CoD is still active with 1 zAAP.
- The Model-Permanent-Capacity Identifier has been modified to 709 as four CPs were added to the permanent configuration.
- All billable CPs from the On/Off CoD record were converted to permanent, so the Model-Temporary-Capacity Identifier is also 709.
- Model-Capacity Identifier is unchanged, eight CPs are activated by replacement capacity records.
- The only available PU was converted to CP, there are no available PUs now.

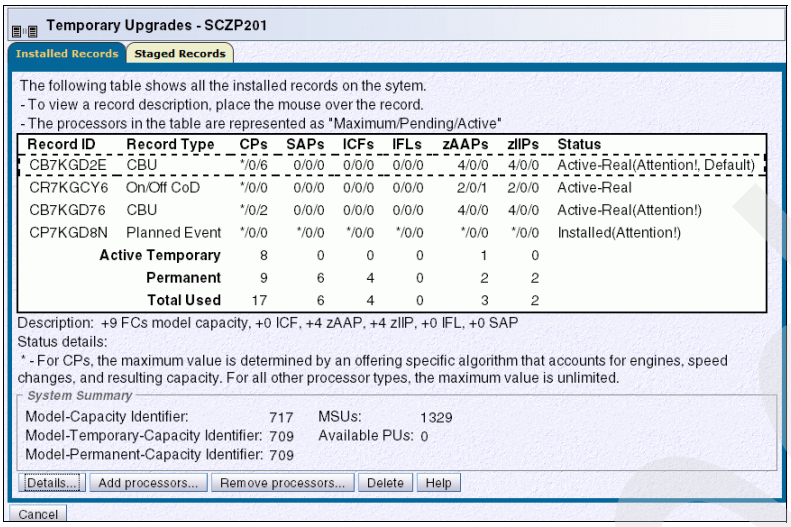

Figure 8-62 Temporary Upgrades shows a new configuration after upgrade

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 254. Note that some of the documents referenced here might be available in softcopy only.

- *IBM System z10 Enterprise Class Technical Introduction,* SG24-7515
- *IBM System z10 Enterprise Class Technical Guide,* SG24-7516
- *IBM System z10 Business Class Technical Overview,* SG24-7632
- *Getting Started with InfiniBand on System z10 and System z9,* SG24-7539
- *IBM System z Connectivity Handbook*, SG24-5444
- *Server Time Protocol Planning Guide*, SG24-7280
- *Server Time Protocol Implementation Guide*, SG24-7281
- *IBM System z10 Enterprise Class Configuration Setup,* SG24-7571

Other publications

These publications are also relevant as further information sources:

- *Hardware Management Console Operations Guide Version 2.10.0,* SC28-6867
- *Support Element Operations Guide V2.10.0,* SC28-6868
- *IOCP User’s Guide,* SB10-7037
- *Stand-Alone Input/Output Configuration Program User’s Guide,* SB10-7152
- *Planning for Fiber Optic Links,* GA23-0367
- *IBM System z10 Capacity on Demand User’s Guide,* SC28-6871
- *CHPID Mapping Tool User’s Guide,* GC28-6825
- *Common Information Model (CIM) Management Interfaces,* SB10-7154
- *IBM System z10 Enterprise Class Installation Manual, GC28-6865*
- *IBM System z10 Enterprise Class Installation Manual for Physical Planning,* GC28-6864
- *IBM System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide,* SB10-7153
- *IBM System z10 Enterprise Class System Overview,* SA22-1084
- *IBM System z Functional Matrix,* ZSW0-1335
- *z/Architecture Principles of Operation,* SA22-7832

Online resources

These Web sites are also relevant as further information sources:

- z/OS Capacity Provisioning

  http://www.ibm.com/servers/eserver/zseries/zos/wlm/cp/

- Software pricing information

  http://www.ibm.com/servers/resources/servers_eserver_zseries_library_swpriceinfo_pdf_swp_ibm.pdf

- z/OS RMF

  http://www.ibm.com/servers/eserver/zseries/zos/rmf

- Getting Started with zPCR (IBM's Processor Capacity Reference)

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

- IBM Resource Link Web site

  http://www.ibm.com/servers/resourcelink

  The Web site offers Capacity on Demand education courses. The courses include ordering a permanent upgrade, ordering On/Off CoD, CBU, and CPE records, approving orders, activating a record, and authorizing users to order upgrades, approve upgrades, and view orders. .

  https://www.ibm.com/servers/resourcelink/hom03010.nsf/pages/usingCIUandOOCODz10?OpenDocument&login

  After logging in, click **Education**. In the Customer Initiated Upgrade section, click **z10 EC**.

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

D

E

F

G

H

P

R

S

T

U

V

W

Z

IBM System z10 Capacity on Demand

(0.5” spine)
0.475”<->0.873”
250 <-> 459 pages

IBM System z10 Capacity on Demand

Describes the concepts and structure of Capacity on Demand

Explains the offerings and how they can be applied

Provides step-by-step examples

The IBM System z10 servers provide a base for major server consolidation by further removing memory, processor, and channel constraints. Capacity on Demand offerings provide the flexibility to absorb temporary or permanent growth, maintenance, and disaster recovery situations in an enterprise.

The Capacity on Demand solutions offer permanent or temporary increases in processor capacity and memory. This IBM Redbooks publication discusses the following topics:

- Understanding the common design of the various offerings
- Hardware and software areas relevant to technical planning
- Managing concurrent use of multiple Capacity on Demand offerings
- User-controlled and autonomic management of z/OS images using z/OS Capacity Provisioning

This book is intended for systems engineers, infrastructure architects, and anyone wanting to understand Capacity on Demand functions. Readers are expected to be generally familiar with existing System z technology and terminology.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7504-01 ISBN 0738432172