IBM Capacity Management Analytics
Version 2.1.1

Solution Guide

Note

Before using this information and the product it supports, read the information in "Notices" on page 365.

Product Information

This document applies to IBM Capacity Management Analytics Version 2.1.1 and may also apply to subsequent releases.

Contents

Introduction

IBM® Capacity Management Analytics provides IT infrastructure system managers, system administrators, and data center executives with access to the information they need to optimize their IT performance, availability, and scalability.

This information is used for cost effective analysis of service objectives, usage and resource utilization, system tuning, accounting, and cost recovery. Using the measurement and management of available and planned capacity, IT can move quickly from insight to action to improve services for the business.

Two products are available for IBM Capacity Management Analytics:

IBM Capacity Management Analytics on distributed platforms

This product includes:

- IBM Cognos® Analytics
- IBM SPSS® Modeler Gold
- IBM ILOG® CPLEX® Optimization Studio
- IBM Capacity Management Analytics Solutions Kit
- IBM Capacity Management Analytics Solutions Installer

Important: IBM Tivoli® Decision Support for z/OS® is not included with this product. You will need to purchase it separately.

IBM Capacity Management Analytics on z/OS

This product includes:

- IBM Capacity Management Analytics on distributed platforms
- IBM Capacity Management Analytics Product Registration
- IBM Tivoli Decision Support for z/OS

Note: To run Capacity Management Analytics, you must have IBM Tivoli Decision Support for z/OS already installed. The following options are recommended to be able to manage most systems:

- System
- CICS®
- IMS™
- Network
- Distributed
- AS400

The Capacity Management Analytics Solutions Kit provides built-in expertise on capacity management, which is designed to support the core capabilities associated with capacity management for IBM z Systems® and distributed-based infrastructures. It consists of a predefined data model to warehouse machine data from across the data center, a set of predefined reports to visualize what is defined and how it is being used, a set of predictive models to understand what is needed in the future and likely usage scenarios, and a set of prescriptive models to optimize the use of IT resources.

The Capacity Management Analytics Solution Kit leverages the operational intelligence platform using industry best practices and prebuilt integration between the components to get you up and running quickly and effectively, focusing on the following core capabilities:

System management and optimization

Sophisticated reporting capabilities to allow analysts, systems managers and executives to view, interact and personalize their reporting that is based on their unique requirements.

Problem identification and resolution

A top-down view of zEnterprise® and distributed system workloads with the ability to drill into further detail, perform ad hoc analysis, create system alerts, or monitor performance to predict potential issues before they affect the business.

Software costs analysis

Ability for your customer to help manage z/OS® software costs and enable capacity management planners to identify when and where workloads need to be adjusted and when more capacity is required.

Capacity forecasting

Ability to forecast your future requirements to ensure that capacity is available when the business needs it.

Real-time anomaly detection

Ability to improve systems management response time with a tool that can detect CICS transaction anomalies in real time.

Audience

This guide is intended to provide users with an understanding of how to install, maintain, and use the Capacity Management Analytics Solution Kit. It describes the reports that are included with the solution kit and the models and streams that are used to provide the data for the reports. It is intended for Capacity Management Analytics administrators, system administrators, database administrators, analysts, report authors, content administrators, and modelers.

Finding information

To find Capacity Management Analytics product documentation, go to Capacity Management Analytics (http://www.ibm.com/support/knowledgecenter/SSUFR9_2.1.1) on IBM Knowledge Center.

Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products. Some of the components included in the Capacity Management Analytics solution have accessibility features. For information about these features, see the Appendix A, "Accessibility features," on page 291 section in this document.

Capacity Management Analytics HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

What's new?

The following new features are available in Capacity Management Analytics 2.1.1.

New features in Capacity Management Analytics 2.1.1

A new WEB UI to manage and drive automated installation and configuration of the Capacity Management Analytics 2.1.1 operational intelligence platform components that install on distributed systems as well as the Solution Kit expert-based application modules. The following sub-hills provide more details:

Installer WEB UI

Installation is made easier with a WEB user interface that leverages NodeJS thus eliminating the need for a fat client that existed in prior versions of Capacity Management Analytics. The UI provides a recommended configuration and recommended defaults for the required attributes.

The WEB UI drives a solution installer to automatically install and configure the platform components and solution kit modules.

Install verification via use of sample database

A sample database with TDSz data is provided that can be used to verify that all of Capacity Management Analytics is properly installed and configured. It is an optional selection for installation that when selected will create the sample TDSz database and load the data. Subsequent Solution Kit modules will use this data allowing for the verification of the streams and reports for those modules.

Automatic installation and configuration of platform components

A solution installer is leveraged to install and configure Cognos, SPSS Modeler, CPLEX, and DB2® LUW on a single node. The installer WEB UI is used to override the minimal set of attributes required if so desired.

Automatic installation and configuration of Solution Kit modules

A solution installer is leveraged to install and configure the Solution Kit modules and sample database on a feature basis. Only the BASE module is required to install the first time. The rest of modules can be installed on an as needed basis.

Solution Kit Command line installation, configuration, and management

The installation, configuration, and management of the Solution Kit can also be done via command line instead of through the WEB UI and Solution Installer. This is useful for continued control of solution kit features after the initial installation and configuration. It is also useful in operating systems where the solution installer is not supported.

Commands are available for CMAHOME and CMAINSTANCE which drives the individual solution kit modules. Each solution kit module provides its own command line with additional levels of control.

Preview Solution Kit installation, configuration, management

A preview mode is provided for the installation, configuration, and management of the Solution Kit. This allows you to see what will be executed before running it so you can review and make sure you organization’s standards are being followed.

Solution Kit Instance

The Solution Kit installation is separated into a CMAHOME directory and a CMAINSTANCE directory. The CMAHOME is for executable code, configuration information, and data that should not be changed after the initial installation. The CMAINSTANCE directory is for configuration and data specific to an individual instance of Capacity Management Analytics. The CMAHOME can be used by multiple CMAINSTANCE(s) each with its own directory and Capacity Management Analytics metadata database.

Chapter 1. IBM Capacity Management Analytics

IBM Capacity Management Analytics architecture can combine multiple data sources and other IBM database products to provide extensive capacity management reporting.

The following diagram shows a typical core architecture for Capacity Management Analytics where a single warehouse is used to store all data. The Capacity Management Analytics solution runs on an IBM zEnterprise system and uses an IBM DB2 for z/OS database server. The core software components include the following:

- IBM Tivoli Decision Support for z/OS
- IBM SPSS Modeler Professional
- IBM Cognos Analytics
- IBM ILOG CPLEX Optimization Studio
- IBM Capacity Management Analytics Solution Installer
- IBM Capacity Management Analytics Solution Kit

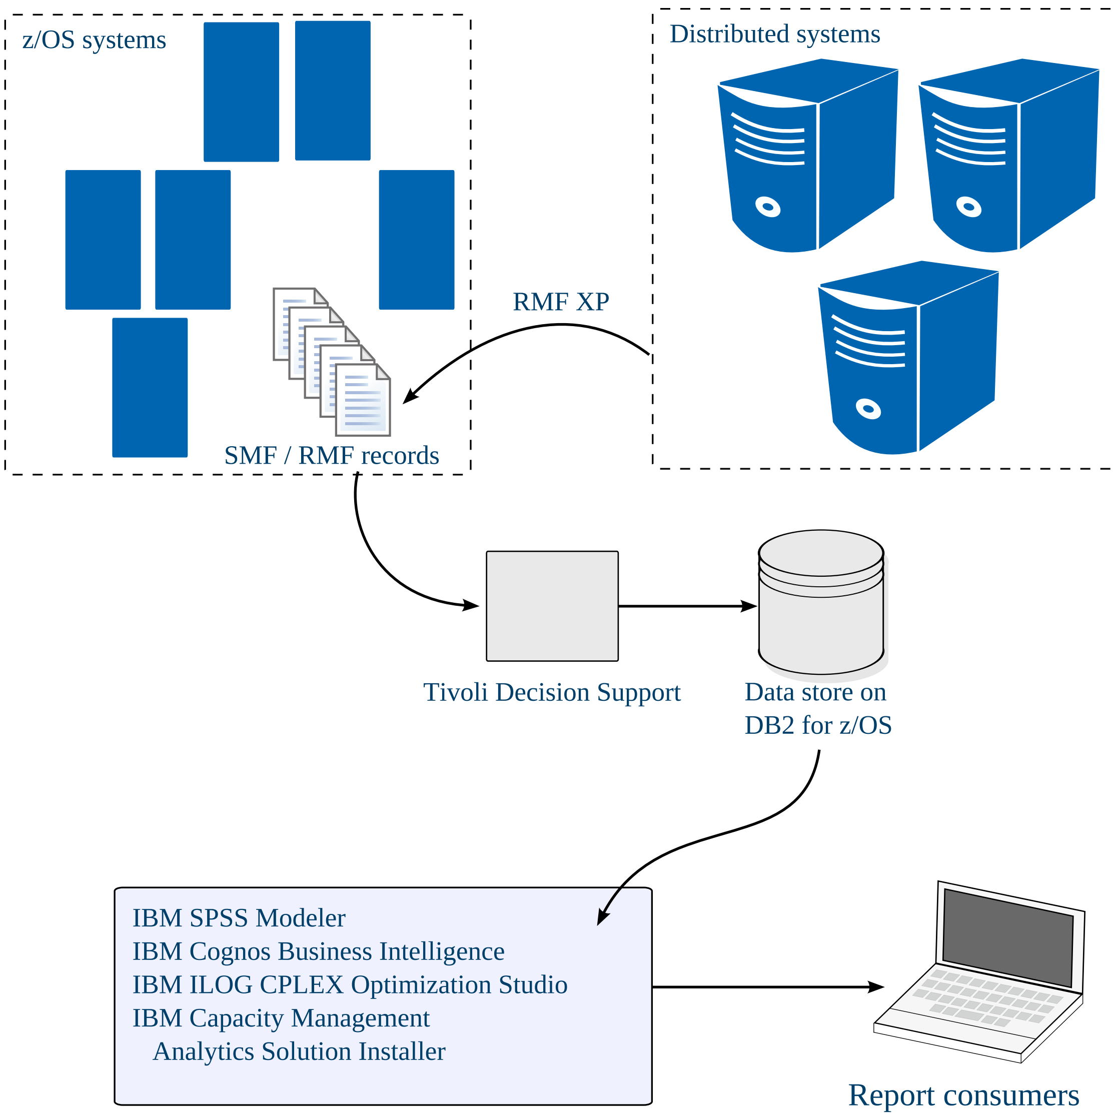

Figure 1: A typical Capacity Management Analytics architecture on z systems

Server-side architecture and operations

The server-side components in the Capacity Management Analytics solution run on the IBM zEnterprise system. All related data for the system is stored in an IBM DB2 for z/OS database server.

System Management Facility (SMF) and Resource Measurement Facility™ (RMF™) data is generated by zEnterprise systems and loaded into a DB2 for z/OS database through Tivoli Decision Support for z/OS. Some SMF and RMF data is further processed by SPSS Modeler streams and loaded into the Capacity Management Analytics database. After the data is loaded, users can view the Capacity Management Analytics reports from the IBM Cognos Analytics portal.

The following diagram shows the flow of data for the Capacity Management Analytics solution.

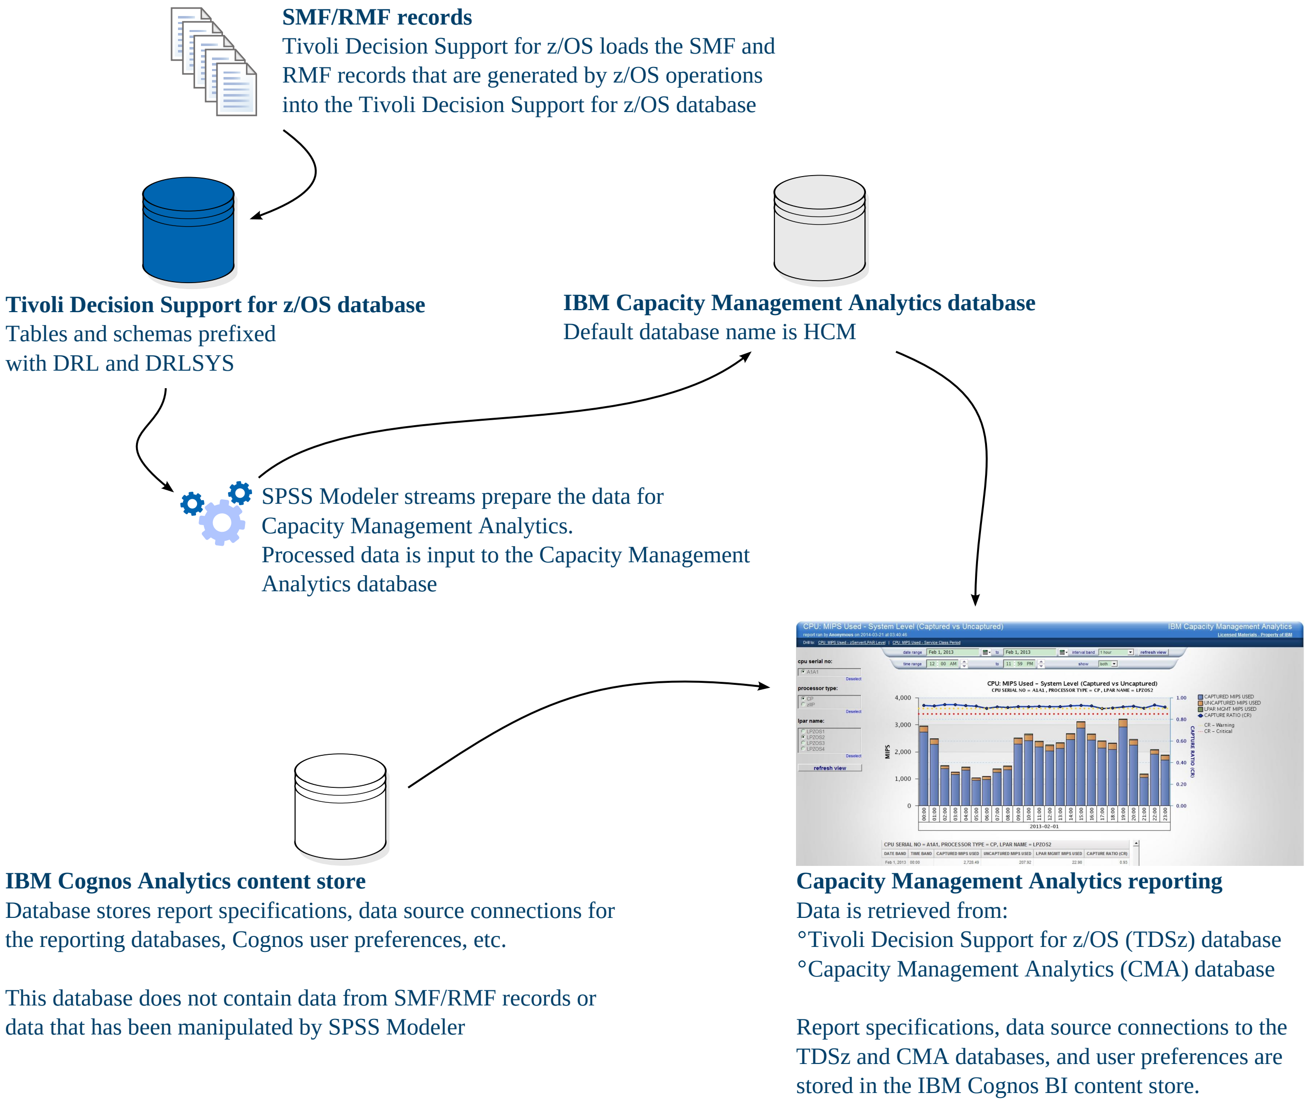

Figure 2: Capacity Management Analytics data flow

The Capacity Management Analytics Solution Kit provides the reports that are deployed into IBM Cognos Analytics and the SPSS Modeler streams.

Client-side applications

The client side of the Capacity Management Analytics solution can run different types of components, such as IBM Cognos Framework Manager, the SPSS Modeler Client, and web browsers for consuming reports.

The Capacity Management Analytics Solution Kit provides the Framework Manager model and the SPSS streams.

Assumptions for successfully running Capacity Management Analytics

- Naming conventions are used within your datacenter to keep LPARs and SYSIDs are unique.
- Best practice installations, configurations, and performance tuning have been applied to the core software products.

Chapter 2. Downloading IBM Capacity Management Analytics 2.1.1

After you have ensured your environment is set up correctly, you can download Capacity Management Analytics 2.1.1.

First, decide whether you will purchase IBM Capacity Management Analytics for distributed platforms, or IBM Capacity Management Analytics on z/OS, which includes IBM Tivoli Decision Support for z/OS.

Next, decide whether you will use the IBM Capacity Management Analytics Solution Installer to automatically install all distributed components or whether you will install each individual component manually.

For more information about the software that you must download, see the IBM Capacity Management Analytics 2.1.1 download document (www.ibm.com/support/docview.wss?uid=swg24042285).

Procedure

1. Order Capacity Management Analytics 2.1.1 on Passport Advantage® or Capacity Management Analytics 2.1.1 on z/OS on ShopZ .
2. If you have ordered IBM Capacity Management Analytics 2.1.1 distributed systems, then follow these steps.
   a) Access the IBM Passport Advantage web site.

      Tip: If you receive an error, try using a different web browser to access Passport Advantage.
   b) Sign in and navigate to the software downloads page.
   c) Find the eImages for Capacity Management Analytics 2.1.1 and select the check box beside the name of the desired eImages to download
   d) If you are installing using the solution installer, download the following:
      - IBM Capacity Management Analytics Solution Installer 2.1.1 for the required platform (Linux on System z or x86-64).
      - IBM Capacity Management Analytics Server Bundle 2.1.1 for the required platform (Linux on System z or x86-64).
      - IBM Capacity Management Analytics Solution Kit 2.1.1 Multiplatform English.
   e) If you are installing manually, you must download each individual server component for the required system and platform.
   f) If you intend to use client side components, download each individual client component. The following are recommended:
      - IBM Cognos for Microsoft Office Microsoft Windows Multilingual
      - IBM Cognos Framework Manager Microsoft Windows Multilingual
      - IBM SPSS Modeler Gold Client Keyless 64-bit Microsoft Windows Multilingual
      - IBM SPSS Data Access Pack Multiplatform English
   g) When the download is complete, a **Download Complete** message will be displayed. The location of the downloaded files will be displayed in the message window.
   h) Save the packages to the corresponding platforms.
   i) Follow the instructions of the Chapter 4, "Installing IBM Capacity Management Analytics," on page 33 section to install and configure the Capacity Management Analytics Solution Kit.
3. If you have ordered IBM Capacity Management Analytics on z/OS 2.1.1, follow these steps:
   a) Sign in to the ShopZ website and navigate to the software download page.

b) Search using product ID 5698-CMA or description IBM Capacity Management Analytics under Package type z/OS – CBPDO and Group MVS™ – Application Development.

c) Download from the z/OS (Z038) SREL: IBM Capacity Management Analytics Product Registration (HHCN211) and the required IBM Tivoli Decision Support for z/OS FMIDs and follow instructions in their respective program directories to complete their installation.

d) Download Accessing IBM Capacity Management Analytics on z/OS DVDs (LCD7704900). Read and follow the instructions it provides to download the bundled products that install on distributed systems from the IBM Web Membership (IWM) website.

e) Log in to the IBM Web Membership website with your IBM ID and the security key provided.

f) If you will install using the solution installer, download the appropriate server package bundle for your platform from the following:

- IBM Capacity Management Analytics Server Package Bundle with Solution Installer for Linux on z Systems (recommended)
- IBM Capacity Management Analytics Server Package Bundle with Solution Installer for Linux x86-64.

g) If you are installing manually, you must download each individual server component for the required system and platform for the product bundles in the IBM Web Membership website.

h) If you intend to use client side components, download the IBM Capacity Management Analytics Client Package Bundle (recommended) or each individual client component.

i) Save the packages to the corresponding platforms.

j) Follow the instructions to install the software.

Chapter 3. Preparation of your environment

Before you can install IBM Capacity Management Analytics, you must ensure that your environment is set up. Preparing your environment includes applying settings for Tivoli Decision Support for z/OS, identifying the Tivoli Decision Support for z/OS schemas that Capacity Management Analytics requires, creating users and groups, and configuring IBM SPSS Modeler for Capacity Management Analytics.

Software, hardware, and operating system requirements

Ensure your environment is prepared with the correct software, hardware, and operating systems.

To review an up-to-date list of environments that are supported by Capacity Management Analytics, see the IBM Software Product Compatibility Reports (www.ibm.com/support/docview.wss?uid=swg27045028).

For the list of required fix packs and APARs for the products that are used with Capacity Management Analytics, see the Required fix packs and APARs for Capacity Management Analytics page (www.ibm.com/support/docview.wss?uid=swg27046400).

Recommended

RedHat 7 or later on Linux System z is recommended for all distributed components, including IBM Cognos Analytics and IBM SPSS Modeler Premium. It is further recommended that the solution installer be used to install all of the server components and the solution kit.

Operating system requirements

The Capacity Management Analytics Solution Installer must be on 64bit servers that are running either Red Hat Enterprise Linux Server Edition version 6.7 or version 7.2 operating system for either z Systems or x86-64 hardware.

Note: If you use the solution installer, the operating system for the computer where you run the solution installer must be set to use English as the language. The solution installer does not run on other languages.

The Capacity Management Analytics Solution Kit can be installed on 64bit servers that are running either Red Hat Enterprise Linux Server Edition version 6.7 or version 7.2 operating system or SUSE Linux Enterprise Server 12.0 for either z Systems or x86-64 hardware.

The Capacity Management Analytics client components must be on 64bit computers that are running either Microsoft Windows 7 or Microsoft Windows 8 operating systems.

Hardware requirements

The server must have at minimum:

- 2 processors, with 4 recommended
- 8GB memory, with 12GB recommended
- 60 GB of free hard disk space

Prerequisite software for installing Capacity Management Analytics

To run IBM Capacity Management Analytics, you must have certain IBM software already installed.

It is recommended to use the solution installer for all prerequisite software that is supported by the solution installer including Cognos, SPSS, CPLEX, and DB2 LUW. Use the documentation for the respective products to install the software that you wish to install yourself without the solution installer or is not supported by the solution installer.

Note: The solution installer does not support the installation of any z/OS software.

The prerequisite software that is listed here are the minimum versions that are required for Capacity Management Analytics. It is recommended to always use the most recent version or fix pack of the prerequisite software.

Java

Java is included. Do not use external or unapproved Java versions.

Tivoli Decision Support for z/OS 1.8.2

Tivoli Decision Support for z/OS version 1.8.2 is required for Capacity Management Analytics.

Note: You must apply fixes PI37340 and PI41870 for Tivoli Decision Support for z/OS version 1.8.2 for Capacity Management Analytics.

For information about upgrading to Tivoli Decision Support for z/OS version 1.8.2, see Migration from TDSz 1.8.0 and 1.8.1 to 1.8.2 (www.ibm.com/support/docview.wss?uid=swg21699114).

The need for more Tivoli Decision Support for z/OS components is determined by your environment and reporting requirements. The following Tivoli Decision Support for z/OS components are required for all installations:

- z/OS Performance Management (MVSPM)
  - System statistics
  - Workload statistics
  - Global storage statistics
  - Virtual storage statistics
  - Device statistics
- TCP/IP for z/OS
- RMF XP Distributed

The following Tivoli Decision Support for z/OS components are required for the Capacity Management Analytics Application Analytics feature:

- Key Performance Metrics - z/OS
  - Address Space Statistics (SMF type 30)

If Customer Information and Control System (CICS) is installed, the following Tivoli Decision Support for z/OS components are also recommended:

- CICS Monitoring Partitioned Component
- CICS OMEGAMON® Monitoring Component
- CICS Statistics Partitioned Component
- CICS Transaction and UOW Analysis Component (required by the Anomaly Detection feature)
- CICS Monitoring (required by the Application Analytics feature)
- CICS Transaction Gateway Component

Additionally, the following Tivoli Decision Support for z/OS components are recommended for all installations:

- z/OS System (MVS) component
- z/OS Interval Job/Step Accounting component
- z/OS Availability component

- Data Set component

For DB2, the Tivoli Decision Support for z/OS DB2 components are also recommended.

For more information about Tivoli Decision Support for z/OS, see IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2).

IBM DB2 for z/OS 11

This database is required for IBM DB2 Analytics Accelerator (IDAA). For more information about IBM DB2 for z/OS, see IBM Knowledge Center DB2 for z/OS welcome page (www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/home/src/tpc/db2z_11_prodhome.html).

IBM DB2 for z/OS 10

For more information about IBM DB2 for z/OS, see IBM Knowledge Center DB2 for z/OS welcome page (www.ibm.com/support/knowledgecenter/SSEPEK_10.0.0/com.ibm.db2z10.doc/src/alltoc/db2z_10_prodhome.html).

IBM System Director Platform Agent and RMF XP

For the distributed platform monitoring capability of Capacity Management Analytics, you must use RMF XP on z/OS and install the latest version (6.3.3 or later) of IBM System Director Platform Agent on distributed operating systems.

For the Redhat xLinux 64-bit platform, install Platform Agent for RHEL 6 for Linux on IBM X86-64 on KVM.

For more information about IBM System Director Platform Agent, see the IBM System Director introduction (www.ibm.com/systems/director/).

For more information about RMF XP, see the Cross platform monitoring with RMF XP (www.ibm.com/support/knowledgecenter/SSLTBW_2.1.0/com.ibm.zos.v2r1.erbb200/gpm4cim.htm) topic on IBM Knowledge Center.

Note: RMF APAR OA45331 is required for the distributed platform monitoring capability of Capacity Management Analytics.

Note: Important: Additional prerequisite software is required if the Capacity Management Analytics solution is installed by not using the Solution Installer.

IBM Cognos Analytics 11

For information about installing IBM Cognos Analytics, see the Installation and Configuration Guide at IBM Cognos Analytics 11.0 documentation (http://www.ibm.com/support/knowledgecenter/SSEP7J_11.0.0/com.ibm.swg.ba.cognos.cbi.doc/welcome.html).

Note: You must also install the IBM Cognos Analytics Software Development Kit component before you install Capacity Management Analytics.

IBM SPSS Modeler Server 17.1

SPSS Modeler Server must be configured to use the DB2 ODBC driver connection. You cannot use Data Direct. For more information, see Configuring IBM SPSS Modeler Server with DB2 CLI/ODBC on Linux on System z® operating systems (Configuring IBM SPSS Modeler Server with DB2 CLI/ODBC on Linux operating systems)

The SPSS Modeler Server installation must include Modeler Batch. For more information, see IBM SPSS Modeler V17.1.0 documentation

For information about the required ports for SPSS Modeler, see the maximum and minimum server port information IBM SPSS Modeler V17.1.0 documentation

IBM ILOG CPLEX Optimization Studio 12.6.3

For more information about IBM ILOG CPLEX Optimization Studio, see IBM ILOG CPLEX Optimization Studio V12.6.3 documentation

IBM DB2 Enterprise Server 10.5

For more information about IBM Data Server Client, see IBM DB2 10.5 for Linux, Unix and Windows documentation

Note: Important: Your license for DB2 must include connectivity to DB2 for z/OS.

IBM DB2 Analytics Accelerator (IDAA)

IDAA with Capacity Management Analytics is described here: "Use IBM DB2 Analytics Accelerator (IDAA) with IBM Capacity Management Analytics" on page 294

DB2 package requirements for SLES and RHEL distributions

The library libstdc++.so.6 is required for DB2 database servers and clients. For more information, see IBM DB2 10.5 for Linux, Unix and Windows documentation

System Management Facility and Resource Measurement Facility records

IBM Capacity Management Analytics uses the following System Management Facility (SMF) and Resource Measurement Facility (RMF) records.

The System Management Facility (SMF) is a z/OS component that provides the means for gathering and recording information for evaluating system usage, applications, security, error conditions, and hardware resources.

The Resource Measurement Facility (RMF) provides performance and usage instrumentation of hardware resources such as processor, memory, disk, cache, workload, and virtual storage. RMF records are a subset of SMF records, in the range from SMF 70 to SMF 79.

The SMF and RMF records are captured and stored in the Tivoli Decision Support for z/OS (TDSz) datastore in a DB2 for z/OS database. Capacity Management Analytics uses the Tivoli Decision Support for z/OS datastore as the source for the report data.

Tip: A sample file is provided to help you load the SMF data into the TDSz datastore. The sample file is named `HCMCOLL`, and it is in the *CMA_ROOT*`/scripts` directory where you install the Capacity Management Analytics Solution Kit. On z/OS systems, the sample file is found in the samples directory in the SMP/E installation folder.

SMF and RMF records by feature

The following information shows the SMF and RMF records that are used by the different Capacity Management Analytics features and the Capacity Management Analytics objects and database tables that the SMF and RMF records correlate to.

SMF 30

Table 1: SMF 30 record ID

TDS Tables	TDS Tablespaces	Description and Volume	CMA features
KPMZ_JOB_INT_H	DRLSKZJ2	Key Performance Indicators on address space High volume	Application Analytics (AA)

SMF 70

Table 2: SMF 70 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_LPAR_H MVSPM_SYSTEM_H MVSPM_LPAR_MSU_T MVSPM_CLUSTER_H	DRLSMP14 DRLSMP26 DRLSMP4A DRLSMP40	RMF processor activity Low volume	Base System Management & Optimization (SMO) Capacity Planning & Forecasts (CPF) Application Analytics (AA) Software Cost Analysis (SCA)

SMF 71

Table 3: SMF 71 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_PAGING_H	DRLSMP22	RMF paging activity Low volume	System Management & Optimization (SMO)

SMF 72

Table 4: SMF 72 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_GOAL_ACT_H MVSPM_WORKLOAD2_H	DRLSMP13 DRLSMP35	RMF workload activity and storage data High volume	System Management & Optimization (SMO) Capacity Planning & Forecasts (CPF) Application Analytics (AA)

SMF 73

Table 5: SMF 73 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_CHANNEL_H	DRLSMP07	RMF channel path activity High volume	Base System Management & Optimization (SMO)

SMF 74

Table 6: SMF 74 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_DEVICE_H	DRLSMP11	RMF activity of resources High volume	Base System Management & Optimization (SMO)

SMF 78

Table 7: SMF 78 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_VS_CSASQA_H	DRLSMP29	RMF virtual storage and I/O queuing activity Low volume	System Management & Optimization (SMO)

SMF 89

Table 8: SMF 89 record ID			
TDS Tables	**TDS Tablespaces**	**Description and Volume**	**CMA features**
MVSPM_PROD_T MVSPM_PROT_INT_T	DRLSMP4B DRLSMP4C	Product MSU utilization MVSPM_ PROD_T: low to medium volume depending on how many Sub Capacity products are running MVSPM_ PROD_INT_T: medium to high volume depending on how many Sub Capacity products are running and the frequency of the SMF interval records	Software Cost Analysis (SCA)

SMF 104

Table 9: SMF 104 record ID

TDS Tables	TDS Tablespaces	Description and Volume	CMA features
Z_OPERATING_SYS_T Z_OPERATING_SYS_H X_OPERATING_SYS_T X_OPERATING_SYS_H A_COMPUTER_SYS_T A_COMPUTER_SYS_H A_OPERATING_SYS_T A_MEMORY_T A_MEM_EXPANSION_T W_OPERATING_SYS_T W_OPERATING_SYS_H W_PROCESSOR_T	S104*	Distributed systems – CIM agents Low volume	Application Analytics (AA) System Management & Optimization (SMO)

SMF 110

Table 10: SMF 110 record ID

TDS Tables	TDS Tablespaces	Description and Volume	CMA features
CICS_T_TRAN_T CICS_TRANSACTION_H	DRLSCU01 DRLSC801	CICS statistics High volume	Application Analytics (AA) Problem Identification (PI)

SMF 119

Table 11: SMF 119 record ID

TDS Tables	TDS Tablespaces	Description and Volume	CMA features
TCP_TCPSERV_PORT_H	DRLSTCPO	Hourly TCP/IP statistics about port connection activity High volume	System Management & Optimization (SMO)

SMF and RMF records by component

The following information shows the SMF and RMF records that are used by the different Capacity Management Analytics components and the Capacity Management Analytics objects and database tables that the SMF and RMF records correlate to.

Base components

Tivoli Decision Support control and lookup tables

MVS_MIPS_T

MVSPM_TIME_RES

DAY_OF_WEEK

PERIOD_PLAN

SPECIAL_DAY

For more information about the control and lookup tables, see "Update Tivoli Decision Support for z/OS control and lookup tables" on page 23.

CMA streams

distribution_procedure.str

mvspm_lpar_clone.str

mvspm_lpar_insert.str

mvspm_lpar_refresh_full.str

mvspm_lpar_refresh_roll.str

mvspm_lpar_update.str

Sources tables

DRL.MVSPM_LPAR_H

DRL.MVSPM_SYSTEM_H

HCM.MVSPM_LPAR

SMF source tables

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_CHANNEL_H

MVSPM_DEVICE_H

SMF records

SMF 70

SMF 73

SMF 74

System management and optimization components

The following list shows the Capacity Management Analytics streams with the source tables and SMF records that they use.

CMA streams

channel_top_etl.str

channel_top_etl_refresh_full.str

channel_top_etl_refresh_roll.str

device_lpar_agg_etl.str

device_lpar_agg_etl_refresh_full.str

device_lpar_agg_etl_refresh_roll.str

device_top_etl.str

device_top_etl_refresh_full.str

device_top_etl_refresh_roll.str

lwo.jar

Sources tables

DRL.MVSPM_LPAR_H

DRL.MVSPM_SYSTEM_H

DRL.MVSPM_CHANNEL_H

DRL.MVSPM_DEVICE_H

DRL.MVSPM_WORKLOAD2_H

HCM.MVSPM_LPAR

HCM.MVSPM_CHANNEL

HCM.MVSPM_DEVICE

SMF source tables

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_WORKLOAD2_H

MVSPM_CHANNEL_H

MVSPM_DEVICE_H

SMF records

SMF 70

SMF 72

SMF 73

SMF 74

The following list shows the Capacity Management Analytics reports with the source tables and SMF records that they use.

CMA reports

CPU: AIX - 10 Highest/Lowest CPU Usage Peaks

CPU: AIX - CPU Usage

CPU: Latent Demand

CPU: Linux for System X - 10 Highest/Lowest CPU Usage Peaks

CPU: Linux for System X - CPU Usage

CPU: Linux for System Z - 10 Highest/Lowest CPU Usage Peaks

CPU: Linux for System Z - CPU Usage

CPU: MIPS Used - LPAR Level by WLM Importance

CPU: MIPS Used - Service Class Period Level

CPU: MIPS Used - System Level (Captured vs Uncaptured)

CPU: MIPS Used - zServer/LPAR Level

CPU: Over/Under Shared Weight - CPC by LPAR

CPU: LPAR Weight Optimization Run Result

CPU: Percentage of Servers with High/Low Usage

CPU: Servers Reporting by Operating System - Enterprise Level

CPU: Windows - 10 Highest/Lowest CPU Usage Peaks

CPU: Windows - CPU Usage

CPU: z/OS - 10 Highest/Lowest CPU Usage Peaks

Memory: AIX - Memory Usages

Memory: CSA/ECSA/SQA/ESQA Utilization

Memory: Linux for System X - Memory Usage

Memory: Linux for System Z - Memory Usage

Memory: Windows - Memory Usage

WLM: Delays by Importance Level

WLM: Delays by Service Class Period

Workspace: Enterprise Dashboard

Workspace: zServer Monitoring Dashboard

Sources tables

DRL.MVSPM_LPAR_H

DRL.MVSPM_SYSTEM_H

DRL.MVSPM_PAGING_H

DRL.MVSPM_WORKLOAD2_H

DRL.MVSPM_GOAL_ACT_H

DRL.MVSPM_CHANNEL_H

DRL.MVSPM_DEVICE_H

DRL.MVSPM_VS_CSASQA_H

DRL.X_OPERATING_SYS_H

DRL.X_OPERATING_SYS_T

DRL.Z_OPERATING_SYS_H

DRL.Z_OPERATING_SYS_T

DRL.W_OPERATING_SYS_H

DRL.W_OPERATING_SYS_T

DRL.W_PROCESSOR_T

DRL.A_COMPUTER_SYS_H

DRL.A_COMPUTER_SYS_T

DRL.A_COMPUTER_SYS_T

DRL.A_OPERATING_SYS_T

DRL.A_MEMORY_T

DRL.A_MEM_EXPANSION_T

DRL.TCP_TCPSERV_PORT_H

HCM.MVSPM_LPAR

HCM.LPAR_WEIGHT_OPTIMIZATION

HCM.OPTIMIZATION_METADATA

HCM.LPAR_WEIGHT_OPTIMIZATION_PARAMETERS

HCM.LPAR_WEIGHT_OPTIMIZATION_DETAIL

HCM.TIMEZONES

HCM.DIM_DATE

HCM.MVSPM_LPAR_CPU_BUSY_MIPS_PRED

HCM.MVSPM_CPC_CPU_BUSY_MIPS_PRED

HCM.FORECAST_METEDATA

HCM.MVSMP_DEVICE

SMF source tables

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_PAGING_H

MVSPM_WORKLOAD2_H

MVSPM_GOAL_ACT_H

MVSPM_CHANNEL_H

MVSPM_DEVICE_H

X_OPERATING_SYS_H

X_OPERATING_SYS_T

Z_OPERATING_SYS_H

Z_OPERATING_SYS_T

W_OPERATING_SYS_H

W_OPERATING_SYS_T

W_PROCESSOR_T

A_COMPUTER_SYS_H

A_COMPUTER_SYS_T

TCP_TCPSERV_PORT_H

SMF records

SMF 70

SMF 71

SMF 72

SMF 73

SMF 74

SMF 78

SMF 104

SMF 119

Capacity planning and forecast component

The following list shows the Capacity Management Analytics streams with the source tables and SMF records that they use.

CMA streams

mvspm-lpar-daily-forecast-timeseries.str

mvspm-lpar-hourly-forecast-timeseries.str

mvspm-lpar-monthly-forecast-timeseries.str

mvspm-lpar-daily-peak-forecast-timeseries.str

mvspm-cpc-daily-peak-forecast-timeseries.str

Sources tables

HCM.MVSPM_LPAR

SMF source tables

MVSPM_LPAR_H

SMF records

SMF 70

The following list shows the Capacity Management Analytics reports with the source tables and SMF records that they use.

CMA reports

CPU: MIPS Used - LPAR Level by WLM Importance

CPU: MIPS Used - zServer/LPAR Level w/Forecast

Workspace: zIIP/zAAP What ifs - LPAR Level

Sources tables

DRL.MVSPM_LPAR_H

DRL MVSPM_SYSTEM_H

DRL MVSPM_WORKLOAD2_H

HCM.MVSPM_LPAR

HCM.MVSPM_LPAR_CPU_BUSY_MIPS_PRED

HCM.FORECAST_METADATA

HCM.MVSPM_CPC_CPU_BUSY_MIPS_PRED

HCM.LPAR_WEIGHT_OPTIMIZATION

SMF source tables

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_WORKLOAD2_H

SMF records

SMF 70

SMF 72

Application analytics component

The following list shows the Capacity Management Analytics streams with the source tables and SMF records that they use.

CMA streams

appl_capture_ratio.str

appl_lob.strappl_lob_z.str

appl_avg_daily_forecast_timeseries.str

appl_avg_hourly_forecast_timeseries.str

appl_peak_daily_forecast_timeseries.str

Sources tables

DRL.X_OPERATING_SYS_H

DRL.A_OPERATING_SYS_H

DRL.Z_OPERATING_SYS_H

DRL.W_OPERATING_SYS_H

DRL.CICS_TRANSACTION_H

DRL.KPMZ_JOB_INT_H

HCM.APPL_MAPPING

HCM.APPL_UTIL_ZOS_VIEW

HCM.APPL_DIST_MAPPING

HCM.MIPS_CAPACITY

SMF source tables

KPMZ_JOB_INT_H

MVSPM_WORKLOAD2_H

X_OPERATING_SYS_H

A_OPERATING_SYS_H

Z_OPERATING_SYS_H

W_OPERATING_SYS_H

CICS_TRANSACTION_H

SMF records

SMF 30

SMF 72

SMF 104

SMF 110

The following list shows the Capacity Management Analytics reports with the source tables and SMF records that they use.

CMA reports
Application Analytics: Distributed CPU Usage - Applications

Application Analytics: Distributed CPU Usage - Application Summary

Application Analytics: MIPS Used - Application

Application Analytics: MIPS Used - Application by Function

Application Analytics: MIPS Used - Application by LPAR

Application Analytics: MIPS Used - Report Class by Job Names

Sources tables
DRL.KPMZ_JOB_INT_H

DRL. MVSPM_LPAR_H

DRL. MVSPM_SYSTEM_H

DRL.MVSPM_WORKLOAD2_H

DRL. X_OPERATING_SYS_H

DRL. Z_OPERATING_SYS_H

DRL.W_OPERATING_SYS_H

DRL. A_COMPUTER_SYS_H

HCM.APPL_DIST_MAPPING

HCM.TIMEZONES

HCM.MVSPM_LPAR

HCM.DIM_TIME

HCM.DIM_DATE

HCM.APPL_MAPPING

HCM.MIPS_CAPACITY

SMF source tables
KPMZ_JOB_INT_H

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_WORKLOAD2_H

X_OPERATING_SYS_H

Z_OPERATING_SYS_H

W_OPERATING_SYS_H

A_COMPUTER_SYS_H

SMF records
SMF 30

SMF 72

SMF 104

SMF 110

Software cost analysis component

The following list shows the Capacity Management Analytics streams with the source tables and SMF records that they use.

CMA streams

sca_no89.str

sca_pricing.str

sca_subcap.str

sca_tiers.str

msu_gssp_forecast_timeseries.str

msu_lpar_forecast_timeseries.str

msu_prod_forecast_timeseries.str

prod_bill_msu_calculate.str

prod_bill_msu_optimization.str

Sources tables

DRL.MVSPM_CLUSTER_H

HCM.PARENT_PROGRAMS

HCM.SUBCAP_PROGRAMS

HCM.MSU_TIERS

HCM.VUE_TIERS

HCM.CUSTOMER_PRICE

HCM.MVSPM_LPAR

HCM.NO89_PRODUCTS

HCM.MVSPM_PROD_VIEW

HCM.LPAR_MSU_FORECAST

HCM.PROD_MSU_FORECAST

SMF source tables

MVSPM_LPAR_MSU_T

MVSPM_CLUSTER_H

MVSPM_LPAR_H

MVSPM_PROD_INT_T

MVSPM_PROD_T

SMF records

SMF 70

SMF 89

The following list shows the Capacity Management Analytics reports with the source tables and SMF records that they use.

CMA reports

SCA: LPAR MSU Utilization

SCA: NO89 Product MSU and Price

SCA: NO89 Products Matrix

SCA: Product MSU and Prices

SCA: Registered Products Matrix

Workspace: SCA: Software Cost Analysis Summary

Sources tables

DRL.MVSPM_LPAR_H

DRL.MVSPM_SYSTEM_H

DRL.MVSPM_LPAR_MSU_T

DRL.MVSPM_PROD_T

DRL.MVSPM_PROD_INT_T

HCM.MVSPM_LPAR

HCM.PARENT_PROGRAMS

HCM.SUBCAP_PROGRAMS

HCM.PROD_MSU_FORECAST

HCM.PROD_MSU_OPTIMIZATION

HCM.LPAR_MSU_FORECAST

HCM.LPAR_MSU_OPTIMIZATION

HCM.FORECAST_BILLABLE

HCM.OPTIMIZATION_BILLABLE

HCM.OPTIMIZATION_MOVEMENT

HCM.NO89_PRODUCTS

HCM.CUSTOMER_PRICE

HCM.TIMEZONES

SMF source tables

MVSPM_LPAR_H

MVSPM_SYSTEM_H

MVSPM_LPAR_MSU_T

MVSPM_PROD_T

MVSPM_PROD_INT_T

SMF records

SMF 70

SMF 89

Problem identification component

The following list shows the Capacity Management Analytics streams with the source tables and SMF records that they use.

CMA streams

anomaly_classification.str

anomaly_clean_up.str

anomaly_detect_building_c.str

anomaly_detect_building_r.str

anomaly_detect_scoring.str

anomaly_detect_treebased.str

Sources tables

DRL.CICS_T_TRAN_T

HCM.ANOMALIES_CICS

SMF source tables

CICS_T_TRAN_T

SMF records

SMF 110

The following list shows the Capacity Management Analytics reports with the source tables and SMF records that they use.

CMA reports

Anomaly Detection Analytics: CICS Anomaly Details

Sources tables
DRL.CICS_T_TRAN_T

HCM.ANOMALIES_CICS

HCM.DIM_MINUTE

SMF source tables
CICS_T_TRAN_T

SMF records
SMF 110

Settings for Tivoli Decision Support for z/OS

For IBM Capacity Management Analytics reports, you must apply certain settings for your IBM Tivoli Decision Support for z/OS data store.

Note: The default tablespaces DRLSMP07 and DRLSMP11 are segments. If there are some DASD and CHANNEL devices in your z/OS environment, tablespace DRLSMP07 and DRLSMP11 can be updated to partitioned tablespaces by running the Tivoli Decision Support for z/OS component installation. For more information, see Introducing Tivoli Decision Support for z/OS on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSH53X/welcome).

The Tivoli Decision Support for z/OS tables that IBM Capacity Management Analytics uses usually have hourly, daily, and monthly tables. IBM Capacity Management Analytics uses the data from the hourly table. Often the table name ends with "_H".

The settings for your Tivoli Decision Support for z/OS data store are as follows:

- Enable DB2 New Function mode.

  For more information, see IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSH53X_1.8.1/com.ibm.tivoli.dszos.doc_1.8.1/drl5ba1337.htm%23a02pres).
- DRL and DRLSYS are the default schema names that are used by Capacity Management Analytics. If you use different schema names, you must update references to those objects in the Capacity Management Analytics Solution Kit with the appropriate schemas.

  If you use different schema names, you must update the references in the Framework Manager models (`CMA.cpf` and `CMA DW\CMA DW.cpf`), and then republish the packages to IBM Cognos Connection.

  For more information, see "Changing the default schema names in Framework Manager models" on page 293.
- Configure your Tivoli Decision Support for z/OS data store with the appropriate buffer pool.

  Ensure that you use separate buffer pools for the Capacity Management Analytics objects and the Tivoli Decision Support for z/OS objects.

  For your buffer pool settings, for page sizes of 4 and 8 K, you can start with 2000 pages. For page sizes of 16 and 32 K, start with 5000 pages. Then, you can use Tivoli OMEGAMON to review the pages in use and adjust the sizes as needed.
- The forecasting that is done with predictive analytics is most reliable with bigger sets of data. Try to keep at least two years worth of data in the Tivoli Decision Support for z/OS data store. However, if that volume of data is not allowed in your organization, you must check the prediction results to ensure that they are within your tolerated accuracy range.
- Capacity Management Analytics provides flexible self-service reporting that allows for different information intervals.

  Important: Collect and store System Management Facility (SMF) data at an interval that allows the highest range of use. Ideally, collect data at the smallest possible interval, which is 1 minute. However, that interval might not be realistic for certain types of data, such as storage devices. In some cases, intervals of 15 or 30 minutes might be more appropriate. You must select an interval that is most appropriate for your organization and requirements.

Setting the interval must be done in SMF and in Tivoli Decision Support for z/OS.

The default interval for Tivoli Decision Support for z/OS is 1 hour. You can change the interval in the MVSPM_TIME_RES lookup table. Ensure that the Tivoli Decision Support for z/OS interval is not smaller than the SMF interval.

Update Tivoli Decision Support for z/OS control and lookup tables

IBM Tivoli Decision Support for z/OS control and lookup tables provide information about your environment.

The following tables are the core control and lookup tables that need to be updated in the Tivoli Decision Support for z/OS ISPF Administration screens to reflect your specific environment. The updates must be applied before you load data into the Tivoli Decision Support for z/OS data store.

For more information about these tables, see the Tivoli Decision Support for z/OS System Performance Feature Reference guide (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2/com.ibm.tivoli.dszos.doc.1.8.2/SysRef1/DRL5FT16.pdf).

MVS_MIPS_T table

This lookup table must contain the MIPS ratings for the systems that you want to analyze in IBM Capacity Management Analytics, whether currently running or historical.

For example, if you are collecting data into the IBM Tivoli Decision Support for z/OS database for the past two years, then you need entries in the MVS_MIPS_T table for all systems that were installed over the past two years, both active and historical data.

MVSPM_TIME_RES table

This lookup table defines the time resolution to use for each row of data that is stored in a set of MVSPM tables. This lets you specify that data should be recorded for a time period other than the 1-hour default for MVSPM tables. The default TIME_RESOLUTION value must be equal to your RMF recording interval to analyze and report at the most granular level.

For Software Cost Analysis, SMF type 70 records are collected to the MVSPM_LPAR_MSU_T table, and SMF type 89 records are collected to the MVSPM_PROD_T and MVSPM_PROD_INT_T tables. Data is collected by SMF interval to these tables. The aggregation interval is not determined by the MVSPM_TIME_RES lookup table, as is the case with other MVSPM_* tables.

The following tables tend to be high volume, so a 60-minute time resolution is appropriate in the interest of keeping table sizes manageable:

- MVSPM_CACHE_ESS_H
- MVSPM_CACHE_H
- MVSPM_DEVICE_H
- MVSPM_DEVICE_AP_H
- MVSPM_CHANNEL_H
- MVSPM_LCU_IO_H

If your RMF recording interval is 15 minutes for all systems, then the following settings store data in the MVSPM tables at the most granular level of 15 minutes. The exceptions to this are the six tables that are previously identified as requiring a 60-minute time resolution.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
sssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssss
EDIT       SYS13036.T104206.RA000.XXXXXXXX.R0100033         Columns 00001 00046
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
=NOTE=                      TABLE : xxxxxxxx.MVSPM_TIME_RES
====== HO|PERIOD_N|SYSTEM_I|TABLE_SET_NAME    |TIME_R
====== UR|AME     |D       |                  |ESOLUT
======   |        |        |                  |ION
====== ---------------------------------------------
000001 %  %        %        %                  15
000002 %  %        %        MVSPM_CACHE_ESS_H  60
000003 %  %        %        MVSPM_CACHE_H      60
000004 %  %        %        MVSPM_DEVICE_H     60
000005 %  %        %        MVSPM_DEVICE_AP_H  60
000006 %  %        %        MVSPM_CHANNEL_H    60
000007 %  %        %        MVSPM_LCU_IO_H     60
****** **************************** Bottom of Data ****************************
```

Figure 3: MVSPM_TIME_RES table with 15-minute resolution

If the RMF recording interval is not the same on all systems, then the following settings take into account these differences. In this example, systems MVS1 and MVS2 have a 30-minute recording interval while all other systems have a 15-minute recording interval.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
sssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssss
EDIT       SYS13036.T104206.RA000.xxxxxxxx.R0100033         Columns 00001 00046
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
=NOTE=                      TABLE : xxxxxxxx.MVSPM_TIME_RES
====== HO|PERIOD_N|SYSTEM_I|TABLE_SET_NAME    |TIME_R
====== UR|AME     |D       |                  |ESOLUT
======   |        |        |                  |ION
====== ---------------------------------------------
000001 %  %        MVS1     %                  30
000002 %  %        MVS2     %                  30
000003 %  %        %        %                  15
000004 %  %        %        MVSPM_CACHE_ESS_H  60
000005 %  %        %        MVSPM_CACHE_H      60
000006 %  %        %        MVSPM_DEVICE_H     60
000007 %  %        %        MVSPM_DEVICE_AP_H  60
000008 %  %        %        MVSPM_CHANNEL_H    60
000009 %  %        %        MVSPM_LCU_IO_H     60
****** **************************** Bottom of Data ****************************
```

Figure 4: MVSPM_TIME_RES table with varied resolution

Important: If you are using data from tables that are at different time intervals, you must normalize or roll up the lowest granularity to the highest before you can join the data. For example, Table 1 has an interval of 30 minutes and Table 2 has an interval of 60 minutes. Before you join the data, you must ensure that Table 1 aggregates two 30-minute intervals to equal one 60-minute interval so that it matches the rows in Table 2.

DAY_OF_WEEK table

This control table defines the day type to be returned by the DAYTYPE function for each day of the week.

The day type is used as a key in the PERIOD_PLAN and SCHEDULE control tables.

The settings in this example assign Sunday to be Day 1 of the week.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
ssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssss
EDIT       SYS13036.T095831.RA000.XXXXXXXX.R0100030         Columns 00001 00015
Command ===>                                                  Scroll ===> PAGE
****** ***************************** Top of Data ******************************
=NOTE=                      TABLE : xxxxxxxx.DAY_OF_WEEK
====== DAY_OF|DAY_TYPE
====== _WEEK |
====== ---------------
000001 1      SUN
000002 2      MON
000003 3      TUE
000004 4      WED
000005 5      THU
000006 6      FRI
000007 7      SAT
****** **************************** Bottom of Data ****************************
```

Figure 5: DAY_OF_WEEK table

PERIOD_PLAN table

This control table defines the periods that are returned by the PERIOD function.

A period plan defines periods (such as shifts) for each day type that is defined by the DAY_OF_WEEK and SPECIAL_DAY control tables.

In this example, the PRIME period is defined as 6 a.m. to 9 p.m Monday through Friday, and the OFFPRIME period is defined as 9 p.m. to 6:00 a.m Monday through Friday. Both WEEKEND and HOLIDAY periods are also defined. The ranges should be >=START_TIME and <END_TIME. For example, in row 2, `MON 06.00.00 21.00.00 PRIME` means that all records on Monday from 6 am until 8:59:59 pm belong to the PRIME period.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
ssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssss
EDIT       SYS13036.T100502.RA000.XXXXXXXX.R0100031         Columns 00001 00044
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
=NOTE=                      TABLE : xxxxxxxx.PERIOD_PLAN
====== PERIOD_P|DAY_TYPE|START_TI|END_TIME|PERIOD_N
====== LAN_ID  |        |ME      |        |AME
====== --------------------------------------------
000001 %        MON      00.00.00 06.00.00 OFFPRIME
000002 %        MON      06.00.00 21.00.00 PRIME
000003 %        MON      21.00.00 24.00.00 OFFPRIME
000004 %        TUE      00.00.00 06.00.00 OFFPRIME
000005 %        TUE      06.00.00 21.00.00 PRIME
000006 %        TUE      21.00.00 24.00.00 OFFPRIME
000007 %        WED      00.00.00 06.00.00 OFFPRIME
000008 %        WED      06.00.00 21.00.00 PRIME
000009 %        WED      21.00.00 24.00.00 OFFPRIME
000010 %        THU      00.00.00 06.00.00 OFFPRIME
000011 %        THU      06.00.00 21.00.00 PRIME
000012 %        THU      21.00.00 24.00.00 OFFPRIME
000013 %        FRI      00.00.00 06.00.00 OFFPRIME
000014 %        FRI      06.00.00 21.00.00 PRIME
000015 %        FRI      21.00.00 24.00.00 OFFPRIME
000016 %        SAT      00.00.00 24.00.00 WEEKEND
000017 %        SUN      00.00.00 24.00.00 WEEKEND
000018 %        HOLIDAY  00.00.00 24.00.00 HOLIDAY
****** **************************** Bottom of Data ****************************
```

Figure 6: PERIOD_PLAN table

SPECIAL_DAY table

This control table defines the day type to be returned by the DAYTYPE function for special dates such as holidays.

The day type is used as a key in the PERIOD_PLAN and SCHEDULE control tables.

This example defines company holidays for 2012 and 2013.

```
  File  Edit  Edit_Settings  Menu  Utilities  Compilers  Test  Help
ssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssssss
EDIT       SYS13036.T102640.RA000.XXXXXXXX.R0100032         Columns 00001 00019
Command ===>                                                  Scroll ===> CSR
****** ***************************** Top of Data ******************************
=NOTE=                         TABLE : xxxxxxxx.SPECIAL_DAY
====== DATE      |DAY_TYPE
====== -------------------
000010 2012-01-02 HOLIDAY
000011 2012-01-16 HOLIDAY
000012 2012-05-28 HOLIDAY
000013 2012-07-04 HOLIDAY
000014 2012-09-03 HOLIDAY
000015 2012-11-22 HOLIDAY
000016 2012-11-23 HOLIDAY
000017 2012-12-24 HOLIDAY
000018 2012-12-25 HOLIDAY
000019 2013-01-01 HOLIDAY
000020 2013-01-21 HOLIDAY
000021 2013-05-27 HOLIDAY
000022 2013-07-04 HOLIDAY
000023 2013-07-05 HOLIDAY
000024 2013-09-02 HOLIDAY
000025 2013-11-28 HOLIDAY
000026 2013-11-29 HOLIDAY
000027 2013-12-24 HOLIDAY
000028 2013-12-25 HOLIDAY
****** **************************** Bottom of Data ****************************
```

Figure 7: SPECIAL_DAY table

Users and groups for installation and configuration

The IBM Capacity Management Analytics solution is designed to be used and shared among several users.

Groups for installation and configuration

The IBM Capacity Management Analytics solution is designed to be used and shared among several users.

To configure your group for installation and configuration, a group on Linux needs to be created. All Capacity Management Analytics administrators and users will need to be part of this group.

This group will have read, write, and execute access on the Solution Installer directory, CMAHOME directories, and the CMAINSTANCE directories.

To create a group, run these commands: `groupadd <group name> => groupadd cmagrp`.

Users for installation and configuration

The IBM Capacity Management Analytics solution is designed to be used and shared among several users.

User IDs will need to be set up for the Capacity Management Analytics administrator and users. All Capacity Management Analytics work will be done by these user IDs on both Linux and DB2 z/OS.

The Capacity Management Analytics administrator ID is used to install, uninstall, instantiate a Capacity Management Analytics instance, execute features of the Capacity Management Analytics solution, start and stop the Capacity Management Analytics component products (Cognos, SPSS, and so on). The administrator ID needs to have sudo with no password set up on Linux.

To create the Capacity Management Analytics administrator ID, issue this sample command:

```
useradd -m -g <groupname> <username> => useradd -m -g cmagrp cmaadm
```

Note: This command needs to be issued by sudo or root.

A generic Capacity Management Analytics user ID is used to instantiate, execute features, run SPSS streams and connect to DB2 zOS. The same user ID must also be created in z/OS to be used for DB2 connections and the password must be maintained the same on both Linux and z/OS. This ID does not need sudo:

```
useradd -m -g <groupname> <username> => useradd -m -g cmagrp cmauser
```

Subsquent other Capacity Management Analytics user IDs can be added in the same manner. If IDs already exist, they can be added to the Capacity Management Analytics group with the usermod –g cmagrp or usermod –G cmagrp commands.

Modifying the sudoers file for the user who runs the installation

To install IBM IBM Capacity Management Analytics by using the solution installer, you must be either the Capacity Management Analytics administrator, root user or be a user with sudo permissions. To run the installation as a user with sudo permissions, you must add that user to the sudoers file.

Procedure

1. Log in as the Capacity Management Analytics administrator or and ID with sudo or root.
2. Open the `sudoers` file for editing.
   For example, enter `visudo -f /etc/sudoers`.
3. Add the following line for your user:

   ```
   username ALL=(ALL) ALL
   ```

4. Comment out the TTY requirement.
5. #Defaults requiretty.
6. Save and close the file.

Configuring the administrator and user profiles

Before you can use the IBM Capacity Management Analytics IDs, you must set them up correctly.

Note: The CMAHOME and CMAINSTANCE directory values will be used in the Solution Installer and various other Capacity Management Analytics programs. It must be consistent everywhere.

Configuring the administrator and user profiles

Follow this procedure to configure your IBM Capacity Management Analytics administrator and user profiles.

About this task

What follows is a sample configuration.

Procedure

1. Copy and paste the contents of the following example into a file in your home directory called: hcm.profile using your preferred method.
2. Open your profile file.
   a) For RHEL, this is usually .bash_profile
   b) For SLES, this is usually .profile
3. Append this line at the very bottom of the profile: “`. hcm.profile`”

Example

```
#!/bin/sh
###############################################################################
#  Licensed Materials - Property of IBM
#  IBM Products: IBM Capacity Management Analytics
#  (C) Copyright IBM Corp. 2013, 2015
#  US Government Users Restricted Rights
#  Use, duplication or disclosure restricted by
#  GSA ADP Schedule Contract with IBM Corp.
################################################################################
# NAME: hcm.profile
#
# DESCRIPTION: IBM Capacity Management Analytics
#              Sample Unix/Linux profile. This profile should be
#              used by both the CMA_ADM and HCM_USER defined user IDs.
#
# Notes:
#    - Either rename this file to .profile or source it from your
 existing .profile
#    - Replace the values for the environment variables as needed
#    - ulimits: verify with your admin what settings are appropriate for your
 system
#      #unlimited may not always be available
#
# Set:
#    - ulimit settings
#    - DISPLAY to the IP of your desktop
#    - JAVAHOME to the JAVA home directory
#    - CMAHOME to the CMA home directory
#    - CMAINSTANCE to the CMA instance directory
#
###############################################################################

TRUE=1
FALSE=0IS_LINUX=${FALSE}

OS=`echo \`uname\` | tr '[A-Z]' '[a-z]'`     #Operating System in lowercase
echo "OS=${OS}"

case "$OS" in
    "linux")
        IS_LINUX=${TRUE}
    ;;
    *)
        "ERROR: Cannot determine this Operating System"
        exit 1
    ;;
esac

#Execute DB2 profile
if [ -f /home/db2inst1/sqllib/db2profile ] ; then
  . /home/db2inst1/sqllib/db2profile
  export IBM_DB_HOME=$DB2_HOME
fi

#vars common to all platforms
#The CMA_ADM and HCM_USER users must be in the same group
#The instance directories created by CMA_ADM must have rwx for group
umask 002      #Give owner rwx, group rwx and give other rx

export EDITOR=vi
export TMPDIR=${HOME}/temp
export TERM=vt100
export LANG=En_US
```

```
set -o vi

#export PATH="${PATH}:/usr/sbin:/usr/local/bin:/usr/bin:/bin:"

export CMAHOME=/opt/ibm/Capacity_Management_Analytics/v2.1.1
export CMAINSTANCE=/opt/ibm/Capacity_Management_Analytics/v2.1.1_instance

export PYTHONPATH=$CMAHOME/bin:$CMAHOME/lib:$PYTHONPATH

#export JAVA_HOME=/opt/IBM/SPSS/ModelerServer/17.1/jre
#export PATH=$JAVA_HOME/bin:$PYTHONPATH:$PATH
export PATH=$PYTHONPATH:$PATH

#set common ulimits
ulimit -t unlimited
ulimit -d unlimited
ulimit -f unlimited

if [ -f .envfile ] ; then
    . .envfile
fi
```

Granting permissions to the data store

DB2 permissions are needed for theIBM Capacity Management Analytics users to access the IBM Tivoli Decision Support for z/OS data store for Capacity Management Analytics and to create and use the Capacity Management Analytics data store.

You do not have to manually apply these grants. Capacity Management Analytics 2.1.1 separates the access to DB2 into privileged and user role based security. It requires a DB2 privileged user to create the Capacity Management Analytics user role, trusted context, databases, and to grant the requires access to resources for the Capacity Management Analytics role created.

All SQL files for the Capacity Management Analytics instance are in the `$CMAINSTANCE/sql` directory. The SQL files beginning with admin* are the ones that are run with the DB2 privileged user. Look for the SQL files with `grant` in the name to see the specific grants that are done.

Setting ulimit values on the node computers

Before you install IBM Capacity Management Analytics on Linux operating systems, you must ensure that you have appropriate ulimit values on each node.

Procedure

Log in to the computer as the root user or as a user with sudo permissions.

a) Go to the `/etc/security` directory.

b) Open the `limits.conf` file for editing.

c) Add the following lines to the file:

```
@root soft nofile 10240
@root hard nofile 16384
* soft nofile 10240
* hard nofile 16384
```

d) Save and close the file.

e) Restart the computer for the changes to take effect.

Server side network configuration

Server side network configuration includes host file configuration and firewall configuration.

Configuring the host file

If you are installing on a Linux operating system by using the solution installer, you must ensure proper communication.

Procedure

1. On the node computer, open the `/etc/hosts` file.
2. Ensure that the node computer is listed in the file. For example, ensure that your `hosts` files contain entries in the following pattern:
   - 127.0.0.1 localhost ##.##.##.## yourhost.domain.com yourhost
3. Save and close the file.

Configuring the firewall

Configuring your firewall is a necessary part of your network configuration.

Various ports for input and output need to be opened on the server for proper communication. The default ports to be opened are:

- Cognos: 9300
- SPSS Modeler: 28055
- Solution Installer UI: 8080
- Chef: 9683
- SSL: 443
- SSH: 22

If the default ports are not available then the ports will need to be changed after installation.

For information about the required ports for SPSS Modeler, see the Maximum and minimum server port information in the Performance and Optimization topic on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/ clementine/server/adminguidesource/admin_pem_performance.dita).

Configuring RedHat6

You must configure the ports on RedHat 6 as part of your network configuration.

To open the ports on RedHat 6, utilize the `iptables` command for your network zone.

Run `iptables -I INPUT -p tcp --dport 8080 -j ACCEPT iptables -I OUTPUT -p tcp --sport 8080 -m state --state RELATED,ESTABLISHED -j ACCEPT`. For further information, contact your Linux administrator.

Configuring RedHat7

You must configure the ports on RedHat 7 as part of your network configuration.

To open the ports on RedHat 7 utilize the `firewalld` command for your network zone.

For instance, `firewall-cmd --permanent --zone=<your zone> --add-port=9683/tcp`. For further information, contact your Linux administrator.

RedHat Linux (RHEL) utilities

You must install these utilities to use RedHat.

About this task

Procedure

Install the following utilities by running `sudo yum install wget zip unzip file`.

a) Install Wget. Wget, meaning web get, is a command-line utility that downloads files over a network.

b) Install Zip. Zip is a compression and file packaging utility for Linux and Unix.

c) Install UnZip. UnZip will list, test, or extract files from Zip archive files.

d) Install File. File determines the file type.

e) Verify these utilities were installed successfully by running `sudo yum whatprovides /usr/bin/wget`

Chapter 4. Installing IBM Capacity Management Analytics

Use the solution installer to deploy the IBM Capacity Management Analytics component programs on computers or virtual machines that are running 64-bit Red Hat Enterprise Linux Server Edition version 7 operating systems. For other operating systems, you must install the product manually.

The solution installer runs on a Red Hat Enterprise Linux version 7 operating system. It provides a web interface that you use to define, validate, and install your environment.

You define your environment by identifying node computers, and dragging the component programs onto the nodes. The solution installer provides guidance for where you can add components and you can use it to validate the environment that you define.

If you want to deploy the IBM Capacity Management Analytics component programs on other operating systems, you must use a manual installation. For more information, see Chapter 6, “Manually installing IBM Capacity Management Analytics 2.1.1,” on page 59.

Note: If you want to customize the installation of certain components, such as distributing IBM Cognos Analytics on different computers, you can use the manual installation instructions. For more information, see the installation documentation for the component.

Also, if you want to configure multiple nodes of IBM SPSS Collaboration and Deployment Services, you must use a manual installation of IBM Capacity Management Analytics. The solution installer allows only one instance of IBM SPSS Collaboration and Deployment Services in a configuration.

For more information about manually installing IBM Predictive Customer Intelligence, see Manually install IBM Predictive Customer Intelligence.

Important: If you use the solution installer, the operating system for the computer where you run the solution installer must be set to use English as the language. The solution installer does not run on other languages. Also, the client computers must also be using English as the language.

Migration considerations

Migrating IBM Capacity Management Analytics involves upgrading software components and deploying new Capacity Management Analytics report content.

Due to the software components changing and the changes to the underlying architecture, it is recommended that a new Capacity Management Analytics installation be made regardless of the prior version. However, if TDSz Version 1.8.2 is already in use then no migration needs to be performed for Capacity Management Analytics. If you wish to keep the same database and schema names from the prior Capacity Management Analytics datastore, then the old one will need to be purged before doing so.

The prior version of Capacity Management Analytics can exist in parallel with the new installation as long as the new Capacity Management Analytics database is not sharing the prior database and schema names and TDSz is Version 1.8.2. It is recommended that you install and validate the new Capacity Management Analytics installation before removing the prior version.

When you migrate, install the new software and deploy the report content to a test environment where you can verify the report outputs. Then, after you are satisfied that the new product is working as expected, you can migrate your production environment.

When you migrate from an existing version, all of the report content that is deployed to IBM Cognos Analytics is overwritten. If you created new reports or modified any reports in the Capacity Management Analytics Solution Kit, you must take extra steps to preserve your new or modified reports.

IBM Cognos Analytics reports

In IBM Cognos Business Intelligence, you should back up your new and modified reports before you deploy Capacity Management Analytics 2.1.1. You can redeploy the reports, but you should do so to a new location in IBM Cognos Business Intelligence.

For more information about upgrading or migrating IBM Cognos BI, see the upgrading documentation on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J_10.2.2/com.ibm.swg.ba.cognos.inst_cr_winux.10.2.2.doc/c_upgradingtocognos8.html%23UpgradingtoCognos8)

Framework Manager models

Also, if you modified the IBM Cognos Framework Manager models, you must apply your modifications to the Capacity Management Analytics 2.1.1 version of the models.

If you want to continue to use the previous version of the model, you should rename the models and store them in a new location. You must also remap the reports in IBM Cognos Business Intelligence to use the previous version package as the path will have changed.

Report templates

If you want to keep the previous version report templates, you can copy the entire template folder into Capacity Management Analytics 2.1.1, but ensure that you use a new name for the folder in IBM Cognos Business Intelligence. Additionally, you must remap objects as the path will have changed due to the new name.

To migrate any report objects that you created, you should rename the model, report, template, or workspace folders in the previous version before you copy them to the new version.

SPSS Modeler streams

If you have created any IBM SPSS Modeler streams in your previous version, you will have to modify the streams to use in the new version as the database tables and column names have changed. Also, if you have modified any streams, the table and column names have changed in the new version, so you must manually update any streams that you have modified.

Capacity Management Analytics database

Before you migrate your Capacity Management Analytics ensure that you back up your Capacity Management Analytics database. Specifically, ensure that you back up all tables that have HCM as the prefix and all of the tablespaces that have IDHTS as the prefix.

If you use any Capacity Management Analytics database tables as source data for any applications, you must change the names in the new version as the table names will have changed.

Tivoli Decision Support for z/OS database

It is also recommended that you back up the Tivoli Decision Support for z/OS database. For more information, see IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2/com.ibm.tivoli.dszos.doc.1.8.2/Admin/abackup.dita).

Installing and configuring the Capacity Management Analytics solution

Capacity Management Analytics 2.1.1 has a new WEB UI to manage and drive automated installation and configuration of the Capacity Management Analytics 2.1.1 operational intelligence platform components as well as the Solution Kit expert-based application modules. It provides a web interface that you use to define, validate, and install your environment.

If you want to deploy the Capacity Management Analytics component programs on a platform other than 64-bit Red Hat Enterprise Linux Server Edition version 6 and 7, you must install the product manually.

For more information, see Chapter 6, "Manually installing IBM Capacity Management Analytics 2.1.1," on page 59.

Note: Manual installation also allows you to customize the installation of certain components, such as distributing IBM Cognos Business Intelligence on different computers.

Note: If you use the Solution Installer, the operating system where you run the installer must be set to use English as the language. The Solution Installer does not run on other languages. Client computers must also be set to English.

Uploading Capacity Management Analytics

After you have installed all of your prerequisite software and downloaded IBM Capacity Management Analytics, you are ready to upload the solution.

Procedure

1. Navigate to the location where you downloaded the eImages for Capacity Management Analytics.
2. Upload the eImages to your designated Capacity Management Analytics server in the Capacity Management Analytics administrator ID home directory.
3. Extract the `cma_sol_ins_2.1.1.0_lxz_en.tar.gz` for Linux on System z or `cma_sol_ins_2.1.1.0_lx86_en.tar.gz`, Linux x86, file depending on your server with command `tar xvf filename.tar.gz` If you are using RedHat 6, use the command `tar --no-same-owner --no-same-permissions -xvf filesname.tar.gz` with command `tar xvf filename.tar.gz`. If you are using RedHat6, use the command `tar --no-same-owner --no-same-permissions -xvf filesname.tar.gz`. The file will be extracted into a directory named `SolutionInstaller`.
4. Copy the `cma_srv_bndl_2.1.1.0_lxz_en.tar.gz` or `cma_srv_bndl_2.1.1.0_lx86_en.tar.gz` file to the `SolutionInstaller/ NodeRoot/ Downloads/Software` directory.
5. Go back to your `SolutionInstaller` folder, run command `./cmaupdate.py`. You can also run the command `./cmaupdate.py -d SolutionInstallerDirectory`. You can find the log for `cmaupdate.py` in the `SolutionInstaller/log` folder.
6. Copy the `cma_sol_kit_2.1.1.0_mp_en.tar.gz` to the `SolutionInstaller/NodeRoot/ Downloads/Software/CMA` folder.

Starting the solution installer

After the solution installer is running, you can access the installer interface from another computer by using a web browser and start defining the environment.

Note: Ensure that you copy the solution installer files to a directory in which you have permissions to execute files.

Important: Ensure that the operating system on which you run the solution installer is set to use English as the language. The solution installer does not run on other languages. Also, the client computers must also be using English as the language.

Procedure

1. Log on to the computer where you decompressed the installation files as the administrator or as a user with sudo or root.
2. Go to the `SolutionInstaller` directory where you decompressed the solution installer files.
3. Enter the following command: `setup.sh`.

   After the solution installer starts, open a web browser on another computer, and go to the solution installer URL: `https://`*`servername`*`:8080/UI/index.html`.

Important: Do not use a web browser on the computer on which you are running the solution installer or any computer on which a IBM Capacity Management Analytics component is being installed. You must use a web browser on another computer.

The computer on which you are using the browser to access the solution installer must have a screen resolution that is greater than 1024 by 760.

The solution installer runs on Google Chrome 44, or later, or Mozilla Firefox 38 or later, web browsers. It does not run on Internet Explorer.

Ensure that you are using an English language web browser to access the solution installer URL.

4. Validation:

   a) Node:

   - Run command `ps -ef | grep node`
   - You should see the node is running as follows:

```
root 5811 1 0 13:35 ? 00:00:00 /opt/ibm/nodejs1.2.0.1/bin/node
 /opt/ibm/nodejs1.2.0.1/lib/node_modules/forever/bin/monitor NodeJS/
server.js
```

```
root      5816  5811  0 13:35 ?        00:00:00 /opt/ibm/
nodejs1.2.0.1/bin/node
/home/nthuy/sandbox-local/si0003/Capacity_Management_Analytics/v2.1.1/
SolutionInstaller/
NodeJS/server.js
```

   b) Chef client:

   - run command `rpm -qa | grep chef`. You should see the chef client rpm `chef-12.1.2-1.el6.s390x`.

   c) Cookbook. Open cookbook in `cd <your_home_directory>/sandbox-local/si0003/Capacity_Management_Analytics/v2.1.1/SolutionInstaller/chef-repo/cookbooks` to ensure all cookbooks have been copied.

Using the solution installer

You use the solution installer to define your IBM Capacity Management Analytics environment you want to install.

You can save your configurations. For more information, see "Saving and opening configurations" on page 42.

Important: The solution installer runs on Google Chrome 44 or later or Mozilla Firefox 38 or later. It does not run on Internet Explorer.

For more information about supported software, see the IBM Software Product Compatibility Reports (http://www-01.ibm.com/support/docview.wss?uid=swg24042285).

Procedure

1. Open the solution installer in a web browser.

   After the solution installer is running, you can access the URL from any computer from a Firefox or Chrome web browser.

   The URL is `https://`*`servername`*`:8080/UI/index.html`, where *`servername`* is the name of the computer where you ran the solution installer.

   Important: Do not use a web browser on the computer on which you are running the solution installer or any computer on which a IBM Capacity Management Analytics component is being installed. You must use a web browser on another computer.

Ensure that you are using an English language web browser to access the solution installer URL.

2. On the license information page, click **Accept**.

   CAUTION: Declining the license allows you to uninstall the solution installer and remove all of the files from the solution installer directory.

3. Select **Predefined Configuration**, select the **CMA Single Node Configuration** and click **OK**.

   Note: RedHat 6 users must delete the CMA Node and install the solution using a two step process. This process entails the following:

   The configuration is displayed in the **Configuration Editor** pane. Icons are displayed on each component where configuration information is required. For instance, where you must provide user names or passwords.

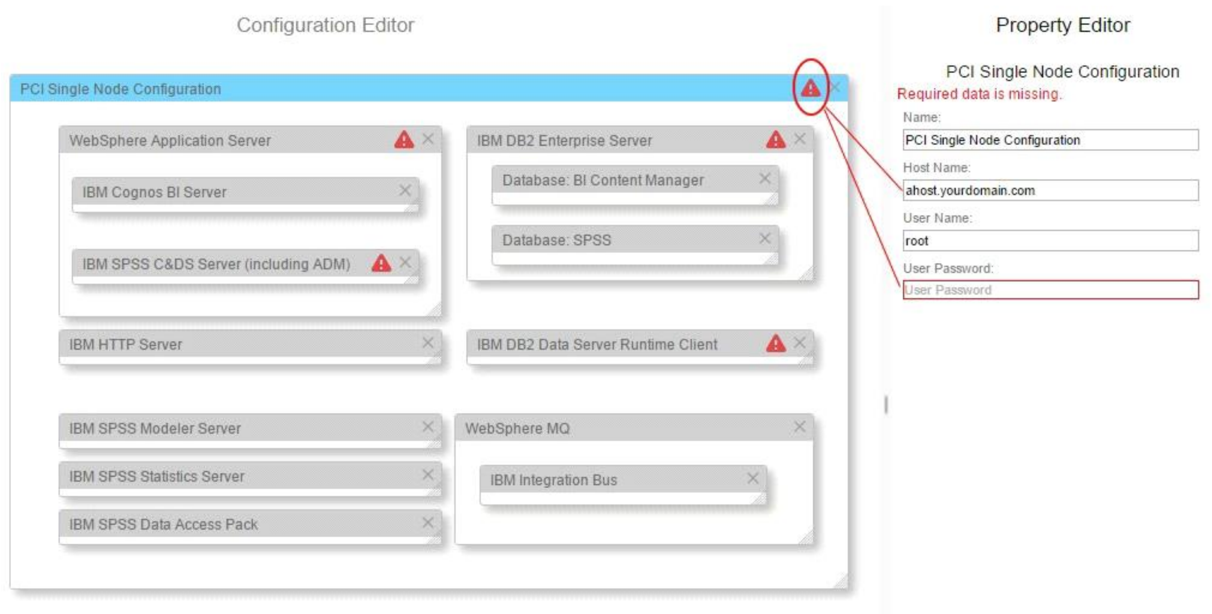

   Note: Ensure that you change the server names to an actual server in your environment. Default values are provided.

   Ensure all of the properties are valid for your environment for each of the components and the nodes in the configuration.

4. In the **Property Editor** pane, select the node object, and enter the following information. For instance, in the OS Node:

   a) Enter a name for the node in **Name**, and select **Enter**.
      For example, enter `BI node`.

   b) Enter the server name in **Host Name**, and select **Enter**.

      Important: Unless you are using a single node configuration, ensure that the host names or IP addresses are for different computers. Do not use the same computer for more than one **Node** object.

   c) Enter the user who has access to install the components in **User Name**, and press the Enter key.
      For example, enter `root` or a user with sudo permissions.

   d) Enter the user's password in **User Password**, and select **Enter**.

   e) Repeat these steps for each node object that you add.
      For example, if you are installing IBM Capacity Management Analytics on four nodes, you must repeat these steps for each of the four node objects.

   Important: You must enter any mandatory information for nodes and other objects as indicated. Ensure that you review all of the properties to ensure that the values, both mandatory and optional, are valid.

Ensure that you select **Enter** after you type a value for a property.

5. From the **Optional Software List**, drag the objects to the node on which you want the installation files copied.
6. As you build your environment, click **Validate** to identify errors or incomplete information.
7. After all nodes are defined and each node is validated, click **Run**.

   Important: After you click the **Run** button, **do not close** the browser window until the installation completes. If you close the browser, you cannot re-initiate the session.

   However, the installation does proceed even with the browser closed. But you cannot monitor the progress.

   If the browser window closes, do the following steps:

   a. Do nothing until enough time elapses to allow the installation to complete. Allow the installation at least 2 hours to complete.
   b. Open a new browser window, and open the solution installer again.
   c. Re-create the configuration that you used or open the saved session that you used. You must ensure that all of the values that you entered in the fields are the same as the configuration that you used when the browser window closed.
   d. Rerun the installation. The solution installer verifies that all of the components are installed.

      After the solution installer completes, verify that the components are working. For more information, see "Testing and validating the installation" on page 43.

   Tip: If you click **Validate** before all the mandatory content is added, a message indicates that the node is not valid. For example, if you validate the BI Node before you add the **IBM DB2** component, then a message indicates that the node is not yet valid.

Configuration properties for the solution installer

You must provide some configuration properties for each of the mandatory software components.

If you select any components from the **Optional Software List**, you must provide a deployment location for the installation files. You must manually install any item that you select from the **Optional Software List**.

Note: For RedHat 6 systems, delete the CMA node.

IBM Capacity Management Analytics Solution Kit properties

Sample Database
A sample TDSz and CMA database is installed that can be used to verify all components of the CMA solution are working properly. Note: Once you decide to use your production data then the sample database must be dropped. Default: Yes.

CMA Feature – Base
Installs and configures the base of CMA providing the database setup and reports. For first time installations it must be installed. Default: Yes.

CMA Feature – Application Analytics
Installs and configures this feature to provide insights to application usage. Default: Yes.

CMA Feature – Capacity Planning & Forecasting
Installs and configures this feature to provide reports and models to plan and forecast for servers and systems usage. Default: Yes.

CMA Feature – Problem Identification
Installs and configures this feature to provide a top down view into systems to identify and resolve problems. Default: Yes.

CMA Feature – Software Cost Analysis
Installs and configures this feature to manage z/OS software costs. Default: Yes.

CMA Feature – Systems Management & Optimization
Installs and configures this feature to provide insights and optimization recommendations to servers and systems. Default: Yes.

Run CMA Instance
A CMA Instance will be created and configured for all the previously listed features. If set to No, an instance will then need to be created manually. However, the instance folder and data files will still be copied over but no stream execution or database work will be done. Refer to cmainstance.py for details.

Note: If you want to use your production data first, instead of samples, then set to No. Default: Yes

Note: Your DB2 administrator may want to review the SQL that is executed before installation occurs. Set this value to No to generate a preview of the SQL files. These files can be reviewed in the CMAINSTANCE/sql folder.

Verbose
Provides more information onto the console and log files. It can also be used to provide trace information. To run a trace, use numeric values 1 to 4. All verbose levels are inclusive.

- 1 = Verbose mode is on
- 2 = Provides informational messages about the function call stack
- 3 = Provides some variable and informational messages
- 4 = Deep tracing, more variables, and informational messages

Default: Yes

CMA Solution Kit Install Directory
This is the CMAHOME directory where all executable and configuration files reside.

Note: If you change this default you must update your CMA administrator and user Linux profiles to reflect this. Default: `/opt/ibm/Capacity_Management_Analytics_/v2.1.1`.

CMA Instance Directory
This is the CMAINSTANCE directory where all data & configuration files for a specific instance resides.

Note: If you change this default you must update your CMA administrator and user Linux profiles to reflect this. Default: `/opt/ibm/Capacity_Management_Analytics_/v2.1.1_instance`.

Cognos Namespace
Optional Cognos Namespace. Value unused.

DB2 Privileged User Name
Credentials for the DB2 z/OS user account with privilege to create databases and tables. This ID is not created for you and must exist before going further in the installation process. See Chapter 3, “Preparation of your environment,” on page 7.

DB2 Privileged User Password
Password for the DB2 privileged user.

DB2 Host Name
DB2 zOS Server hostname.

DB2 Database Port
DB2 zOS Port number.

CMA Database Sub-System ID
The database sub-system ID.

DB2 Database Location Name
The DB2 z/OS Location Name

CMA Database Name
CMA Database Name. Default: HCMDB.

CMA Database User Name
CMA database user. This ID is not created for you and must exist before going further in the installation process. See Chapter 3, “Preparation of your environment,” on page 7.

CMA Database User Password
CMA database user password

CMA Schema Name
CMA schema name. Default: HCM.

CMA Database Storage Group Name
Default: SYSDEFLT

CMA Database CCSID
CMA Database encoding. Default: EBCDIC.

Note: Important: if this database is to be loaded into IDAA then the value must be EBCIDC.

CMA Database Bufferpool
Database bufferpool. It is recommended you use a 8K size bufferpool. Default: BP0.

CMA Database Index Bufferpool
Index bufferpool. It is recommended you use a 8K size bufferpool. Default: BP1.

DB2 for z/OS General WLM Environment Name
WLM environment of DB2 for stored procedures. You must assign the Capacity Management Analytics stored procedure to one predefined WLM application environment to route the stored procedure work to the appropriate WLM-managed address space. For example, you can look for your WLM in SYS1.PROCLIB, otherwise consult with your DB2 sysadmin. Default: CMAWLM_GENERAL.

TDSz Database Name
The Tivoli Decision Support for z/OS database name. TDSz must be installed before going any further in the installation process. Default: DRLDB.

TDSz Schema Name
TDSz schema name. Default: DRL.

TDSz Database Storage Group Name
TDSz database storage group name. Default: SYSDEFLT.

IBM Cognos Analytics Server properties

IBM Cognos BI Analytics Install Directory
The location where IBM Cognos BI Analytics Server is to be installed.

Default: `/opt/ibm/cognos/analytics`

IBM DB2 Enterprise Server properties

DB2 Server Install Directory
The location where IBM DB2 is to be installed on the data node computer.

Default: `/opt/ibm/db2/V10.5`

DB2 Database Port Number
The port number that is used by IBM DB2.

Default: `50000`

DB2 FCM Port Number
The port number that is used by IBM DB2.

Default: `60000`

DB2 Server Admin User Name
The user ID that runs the administration server.

Default: `dasusr1`

This user must not exist on the computer.

DB2 Server Admin Password
The password for the user ID that runs the administration server.

The password cannot contain more than 8 characters and it must be lowercase.

DB2 Server Admin Home Directory
The home directory for the user ID that runs the administration server. This directory is on the data node computer's file system.

Default: `/home/dasusr1`

DB2 Instance Owner User Name
The user ID that controls the DB2 processes and owns the directories that are used by the database instance.

Default: `db2inst1`

This user must not exist on the computer.

DB2 Instance Owner Password
The password for the DB2 instance owner user.

The password cannot contain more than 8 characters and it must be lowercase.

DB2 Instance Owner Home Directory
The home directory for the DB2 instance owner user. This directory is on the data node computer's file system.

Default: `/home/db2inst1`

DB2 Fenced User Name
The user ID that can run user-defined functions and store procedures.

Default: `db2fenc1`

This user must not exist on the computer.

DB2 Fenced User Password
The password for the DB2 fenced user.

The password cannot contain more than 8 characters and it must be lowercase.

DB2 Fenced User Home Directory
The home directory for the DB2 fenced user.

Default: `/home/db2fenc1`

Database: Cognos Content Store

Content Manager Database Name
The IBM Cognos content store database.

This component must be contained within **IBM DB2 Enterprise Server.**

Default: `COG11CS`

IBM SPSS Modeler Server properties

SPSS Modeler Server Install Directory
The location where IBM SPSS Modeler Server is to be installed.

Default: `/opt/IBM/SPSS/ModelerServer/17.1`

IBM SPSS Modeler Batch properties

SPSS Modeler Batch Install Directory
The location where IBM SPSS Modeler Batch is to be installed.

Default: `/opt/ibm/SPSS/ModelerBatch/17.1`

IBM SPSS Data Access Pack properties

SPSS Data Access Pack Install Directory
The location where IBM SPSS Data Access Pack is to be installed.

Default: /opt/ibm/SPSS/SDAP711

IBM ILOG CPLEX Optimization Studio properties

ILOG CPLEX Optimization Studio Install Directory
The location where IBM ILOG CPLEX Optimization Studio is to be installed.

Default: /opt/ibm/ILOG/CPLEX_Studio1263

Saving and opening configurations

You can save IBM Capacity Management Analytics configurations and open existing configurations in the solution installer.

Procedure

1. To save a configuration:
   a) In the solution installer, define your configuration.
   b) Enter a name in the **Save Configuration** box, and click the **Save Configuration** icon.
2. To open a saved configuration:
   a) Click **Open**, and browse to the saved configuration file, and click **Open**.

Starting the installation

The IBM Capacity Management Analytics solution installer installs the component products. Before you start the installation, you can validate the configuration.

Procedure

1. In the solution installer, open the configuration that you want to install.
2. Click **Validate**.

   Any errors or missing information is displayed. You must correct the error or provide the information before you can run the installation.
3. Click **Run**.

   Note: Running times may vary depending on the amount of data loaded within your TDSz datastore and system loads. If just sample data is being created then running times should be within an hour.

RedHat 6 systems

For RedHat 6 systems, you must manually catalog your database connection to DB2 z/OS once you have installed DB2.

Procedure

Run the following commands:

- db2 catalog tcpip node <nodename> remote <dbhostname> server <dbport>
- db2 catalog database <db2locname> as <dbalias> at node <nodename> authentication dcs"

What to do next

Proceed to Chapter 6, "Manually installing IBM Capacity Management Analytics 2.1.1," on page 59

Testing and validating the installation

After the installation of the IBM Capacity Management Analytics is complete, you should test the installation.

Verify that the products exist:

- /opt/ibm/Capacity_Management_Analytics/v2.1.1
- /opt/ibm/Capacity_Management_Analytics/v2.1.1_instance
- /opt/ibm/ILOG
- /opt/ibm/SPSS/ModelerServer
- /opt/ibm/SPSS/ModelerBatch
- /opt/ibm/SPSS/SDAP711

If there are any issues, refer to the *Troubleshooting* section.

Testing the database installation and configuration

You can test the database installation and configuration by logging in to the IBM Capacity Management Analytics data node computer and listing the databases on the server. You can also list the databases available on each node computer where the IBM DB2 client was installed.

Procedure

Log on to the data node computer.

a) Open a terminal window, and change to the database instance owner user.
For example, `su - db2inst1`.

b) Enter the following command to list the databases on the computer:

```
db2 list database directory
```

For an installation on a Red Hat Linux operating systems that used the solution installer, the following databases are listed:

- `COG112CS`
- `CMA_<DBSID>` or the DB2 alias you set using the command line interface.

Testing the installation of the server components on a single node installation

You can test the installation of the IBM Capacity Management Analytics server components by accessing some of the component portals.

For example, you can open the IBM Cognos BI portal. If the portal opens, IBM Cognos BI is installed and running. You can also access the SPSS configuration tool on the Analytics node computer.

Procedure

1. Open a web browser.
2. In the address bar, type the address for the IBM Cognos Content Manager status page.

   The address is `http://`*`servername`*`:9300/p2pd/servlet`.

   The Content Manager status says **Running**.
3. In the address bar, type the address for the IBM Cognos Analytics portal. You can access the IBM Cognos BI portal from the dispatcher URL: `http://`*`servername`*`:9300/bi`
4. Navigate to **Team Content > IBM Capacity Management Analytics > Models** and verify that the date is your install date.
5. To verify that SPSS Modeler Server is running: `/opt/ibm/opt/ibm/SPSS/ModelerServer/17.1/modelersrv.sh list`.

Post installation

After a successful installation of IBM Capacity Management Analytics, restart Cognos.

See Appendix C, “Stop and start solution software services,” on page 305

Using production data

When using IBM Capacity Management Analytics with your production data for the first time, further input is required for proper execution. Various data files for each feature that is pending installation will need to be filled out accordingly.

Refer to the corresponding feature chapter for details about filling out the data files correctly.

To ensure and protect the data files from being overridden, take write capabilities off all the data files in the `CMAINSTANCE/data` folder. Run “`chmod w-ugo *`”.

Note: If you need to make changes to the file then write access will need to be granted again.

Run the cmainstance.py program manually. See Chapter 6, “Manually installing IBM Capacity Management Analytics 2.1.1,” on page 59.

Using the RUNSTATS utility to improve performance

The RUNSTATS online utility of DB2 gathers summary information about the characteristics of data in tablespaces, indexes, and partitions. These JCL jobs can be executed under z/OS TSO.

The samples JCL jobs are in the *CMAHOME*/`scripts` directory.

We recommend that you schedule this utility to run periodically.

Procedure

1. On a Linux on System z operating system, upload HCMSTAT1, HCMSTAT2, and HCMSTAT3 JCL jobs to the z/OS machine and copy into one PDS dataset with RECFM=FB and LRECL=80.
2. Within the JCL jobs, change parameters as described in the following table.

   JCL jobs and descriptions

Table 12:

JCL job name	Description	Parameters to change
HCMSTAT1	Executes RUNSTATS utility for all tables and indexes in Tivoli Decision Support for z/OS DRL and Capacity Management Analytics HCM databases.	Change #DB2SID# to appropriate DB2 SYSTEM ID. Change #DB2LIB# to appropriate DB2 SDSNLOAD LIBRARY. Change #DRLDB# to appropriate Tivoli Decision Support for z/OS database name. Change #HCMDB# to appropriate Capacity Management Analytics HCM database name.

Table 12: (continued)		
JCL job name	**Description**	**Parameters to change**
HCMSTAT2	Executes RUNSTATS utility for special tables and indexes in Tivoli Decision Support for z/OS DRL and Capacity Management Analytics HCM databases.	Change #DB2SID# to appropriate DB2 SYSTEM ID. Change #DB2LIB# to appropriate DB2 SDSNLOAD LIBRARY. Change #DRLDB# to appropriate Tivoli Decision Support for z/OS database name. Change #HCMDB# to appropriate Capacity Management Analytics HCM database name. Change #DRL# to appropriate Tivoli Decision Support for z/OS schema name. Change #HCM# to appropriate Capacity Management Analytics HCM schema name. Change #HCM_ID# to the unique ID you used when creating the HCM tablespace.
HCMSTAT3	Executes RUNSTATS utility for special tables and indexes in Tivoli Decision Support for z/OS DRL and Capacity Management Analytics HCM databases.	Change #DB2SID# to appropriate DB2 SYSTEM ID. Change #DB2LIB# to appropriate DB2 SDSNLOAD LIBRARY. Change #DRLDB# to appropriate Tivoli Decision Support for z/OS database name. Change #HCMDB# to appropriate Capacity Management Analytics HCM database name. Change #DRL# to appropriate Tivoli Decision Support for z/OS schema name. Change #HCM# to appropriate Capacity Management Analytics HCM schema name. Change #HCM_ID# to the unique ID you used when creating the HCM tablespace.

3. Submit these JCL jobs using HCMADM or other database user with STATS privilege for Tivoli Decision Support for z/OS and Capacity Management Analytics databases. Alternatively, you can submit the jobs as a user with DBADM, DBCTRL, or DBMAINT authority for these databases.

Maintaining Capacity Management Analytics

After you install and configure the Capacity Management Analytics Solution Kit, data is loaded into the Capacity Management Analytics data store that you can run reports from. However, as your systems run, SMF and RMF data is continually being collected in the Tivoli Decision Support for z/OS database. To ensure that the Capacity Management Analytics reports continue to reflect your environment and provide useful insights, you must continue to load data into the Capacity Management Analytics datastore and run forecasts on that data.

Tip: If you want to extend or customize Capacity Management Analytics, you can use the Solution Kit reports, models, or streams as examples. However, you should not change any reports, models, or streams. You can make a copy of the report, model, or stream, and modify the copy. Do not modify any of the reports, models, or streams that are provided with the Capacity Management Analytics Solution Kit.

The following are some general guidelines for TDSz loads:

- Data can be loaded into DB2 from TDSz daily, as frequently as your systems allow, or for your business needs.

The RUNSTATS online utility of DB2 gathers summary information about the characteristics of data in tablespaces, indexes, and partitions of Capacity Management Analytics tables and Tivoli Decision Support for z/OS tables. For more information, see "Using the RUNSTATS utility to improve performance" on page 44.

Tip: Sample files are provided to help you do RUNSTATS. The sample files are named `HCMSTAT1`, `HCMSTAT2`, and `HCMSTATS`, and they are in the `CMAHOME/samples` directory where you install the Capacity Management Analytics Solution Kit. On z/OS systems, the sample files are found in the samples directory in the SMP/E installation folder.

The following are some general guidelines for running DB2 RUNSTATS:

- RUNSTATS can be done weekly or if there is a noticeable performance issue or large amount of data loaded.

Running the streams to load data and run forecasts is done by manually running the Capacity Management Analytics feature program. For more information, see Chapter 7, "Manually running features for IBM Capacity Management Analytics 2.1.1," on page 71.

The following are some general guidelines for loading data and running the Capacity Management Analytics feature program:

Important: The Capacity Management Analytics streams can modify the Capacity Management Analytics tables in the database. Always use the Capacity Management Analytics feature program to modify the tables unless it is necessary to use the DB2 SQL statement directly.

The following are some general guidelines for loading data and running the Capacity Management Analytics feature program:

- Loading data can be done on a daily, weekly, or monthly basis, depending on your business needs and system toleration.
- Running the forecast streams can be done every 4 - 6 weeks or whenever there is enough volatility or fluctuation within the historical data or deviations between the forecasts and the historical data.

  Note: All of the preceding steps can be automated through your own tooling. You can also use standard UNIX cronjobs to schedule the work.
- Cleaning up log and temporary files should be done every month. Log and temporary files are in the `*CMAHOME*/log` and `*CMAINSTANCE*/log` and `*CMAINSTANCE*/temp` directories.

- If you want to use any new features after installation, load new feature data with using cmahome.py to install new features, cmainstance.py to create an instance and the Capacity Management Analytics feature programs for further granular control.

The following is a general guideline for Lpar Weight Optimization:

- Run the Lpar Weight Optimizations program for busy Lpars to determine optimized weight value. Chapter 11, “LPAR weight optimization,” on page 111

The following is a general guideline for Cognos reports:

- Batch repeat scheduling after loads & forecasts schedule in Cognos.

For more Cognos scheduling information, see IBM Cognos Business Intelligence documentation.

Chapter 5. Creating a sample database

You can install a sample TDSz and IBM Capacity Management Analytics database that can be used to verify that Capacity Management Analytics has been properly installed and configured.

The sample database is an optional selection for installation. When selected, it will create the sample TDSz database and load the data. Subsequent Solution Kit modules will use this data to allow for the verification of the streams and reports for those modules.

Obtaining a sample database

The sample database can be created automatically from the **Property Editor** pane in the Solution Installer. The default value is **Yes**.

A sample database can also be created manually by running cmainstance.py or cmasampledb.py. See Chapter 6, "Manually installing IBM Capacity Management Analytics 2.1.1," on page 59.

Setting up the sample database

After you have selected the sample database to be installed, the following procedure is performed automatically:

- TDSz database is created with a chosen schema name. The default is DRL.
- All TDSz tables are created. The sample DDLs can be found in `<cma_instance_home>/sql`.
- Sample data is inserted into these TDSz tables.
- The Capacity Management Analytics database is created with the chosen schema name. The default is HCM.
- All Capacity Management Analytics tables are created. The sample DDLs can be found in `<cma_instance_home>/sql`.
- SPSS streams are run based on the selected features to insert sample data into the Capacity Management Analytics tables. To verify the results, run the sql `hcm_dates_rows.sql` and `drl_dates_rows.sql` in `<cma_instance_home>/sql`. The command to run is: `db2 -tf hcm_dates_rows.sql > hcm_dates_rows.log`. This log will show the number of rows that were inserted into each table and the date ranges, minimum and maximum dates, that are used to run the sample reports.
- Best practices are separating schema, storegroup, bufferpool and database from production.

Note: Some features are not available for sample data or date ranges. Refer to the table in this section.

Running reports with the sample database

You can run reports with the sample database in IBM Capacity Management Analytics.

The table in this section contain the Capacity Management Analytics tables, the number of rows within each table, the features and the reports that are mapped with each table. Date ranges are also provided for the sample reports.

Before running the reports, ensure all the packages of Framework Manager models are re-published with the correct schema. See "Changing the default schema names in Framework Manager models" on page 293

Table 13: IBM Capacity Management Analytics Tables (HCM schema).

CMA TABLES (HCM schema)	Number of rows	Feature	Report	Date Ranges (MIN/MAX dates) YYYY-MM-DD
ANOMALIES_CICS	0	PI	CICS Anomaly Details	N/A
APPL_DIST_MAPPING	0	AA	Application Analytics: Distributed CPU usage - applications	N/A
APPL_MAPPING	92	AA	Application Analytics: MIPS Used - Application	2016-09-30 creation date of map
APPL_ZOS_FORECAST	54032	SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used - zServer/LPAR Level	2015-05-21 2015-05-23
CUSTOMER_PRICE	159	SCA	SCA: NO89 Product MSU and Price	1900-01-01 1900-01-01
DIM_DATE	13878	SCA	CPU: AIX - 10 Highest/ Lowest CPU Usage Peaks	2015-01-01 2015-01-01
DIM_MINUTE	1440	SMO	Anomaly Detection Analytics: CICS Anomaly Details	N/A
DIM_TIME	24	PI	Application Analytics: MIPS Used - Application	N/A
Distributed Server data from SMF 104		AA	CPU:AIX - CPU Usage CPU:Linux for System X - CPU Usage CPU:Linux for System Z - CPU Usage	2015-06-02 -2015-07-04 2015-07-31 – 2015-08-03

Table 13: IBM Capacity Management Analytics Tables (HCM schema).(continued)				
CMA TABLES (HCM schema)	**Number of rows**	**Feature**	**Report**	**Date Ranges (MIN/MAX dates) YYYY-MM-DD**
FORECAST_BILLABLE	247350	SCA	SCA: LPAR MSU Utilization	N/A
FORECAST_METADATA	4	SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used - zServer/LPAR Level	No forecast data
LPAR_MSU_FORECAST	9373	SCA	SCA: LPAR MSU Utilization SCA: NO89 Product MSU and Price SCA: Product MSU and Price SCA: Software Cost Analysis Summary	2015-04-01 - 2015-05-03 Observed data 2015-04-01 – 2015-09-10 Forecast data
LPAR_MSU_ OPTIMIZATION	1488	SCA	SCA: LPAR MSU Utilization SCA: NO89 Product MSU and Price SCA: Product MSU and Price SCA: Software Cost Analysis Summary	2015-05-02 - 2015-07-10

Table 13: IBM Capacity Management Analytics Tables (HCM schema). (continued)				
CMA TABLES (HCM schema)	**Number of rows**	**Feature**	**Report**	**Date Ranges (MIN/MAX dates) YYYY-MM-DD**
LPAR_WEIGHT_ OPTIMIZATION	2	SMO	CPU: LPAR Weight Optimization Run Result CPU: MIPS Used - LPAR Level by WLM Importance CPU: Over/ Under Shared Weight - CPC by LPAR	2015-05-21 – 2015-05-23
LPAR_WEIGHT_ OPTIMIZATION_DETAIL	144	SMO	CPU: LPAR Weight Optimization Run Result	2015-05-21 - 2015-05-23
LPAR_WEIGHT_ OPTIMIZATION_ PARAMETERS	4	SMO	CPU: LPAR Weight Optimization Run Result	N/A
MIPS_CAPACITY	46	AA	Application Analytics Reports	N/A
MVSPM_CHANNEL	240	SMO	zServer Monitoring Dashboards	2015-05-21 - 2015-05-21
MVSPM_DEVICE	264	SMO	zServer Monitoring Dashboards	2015-05-21 - 2015-05-21

CMA TABLES (HCM schema)	Number of rows	Feature	Report	Date Ranges (MIN/MAX dates) YYYY-MM-DD
MVSPM_LPAR	12432	AA, CPF, SCA, SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used -Service Class Period Level CPU: MIPS Used - System Level (Captured vs Uncaptured) CPU: MIPS Used - zServer/LPAR Level	2015-04-01 -2015-05-28
MVSPM_LPAR continued	12432	AA, CPF, SCA, SMO	CPU: Over/ Under Shared Weight - CPC by LPAR Application Analytics: MIPS Used – Application Application Analytics: MIPS Used - Application by Function	
MVSPM_LPAR continued	12432	AA, CPF, SCA, SMO	SCA: LPAR MSU Utilization SCA: NO89 Product MSU and Price SCA: Product MSU and Price	

Table 13: IBM Capacity Management Analytics Tables (HCM schema). (continued)

CMA TABLES (HCM schema)	Number of rows	Feature	Report	Date Ranges (MIN/MAX dates) YYYY-MM-DD
MVSPM_LPAR continued	12432	AA, CPF, SCA, SMO	SCA: Software Cost Analysis Summary zIIP/zAAP What ifs - LPAR Level zServer Monitoring Dashboards	2015-04-01 -2015-05-28
KPMZ_JOB_INT_H - SMF 30 data	N/A		Application Analytics: MIPS Used - Report Class by Job Names	No SMF 30 data available, only list of Report Classes 2015-05-21 – 2015-05-23
MVSPM_WORKLOAD2_H from TDS/z, type 72 data			CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used -Service Class Period Level Application Analytics: MIPS Used – Application Application Analytics: MIPS Used - Application by Function Application Analytics: MIPS Used - Application by LPAR	2015-05-21 – 2015-05-23 for LPAR JB0
MVSPM_LPAR_CPU_BUSY_ MIPS_PRED (hourly)	28320	SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used - zServer/LPAR Level	No forecast data for WLM reporting 2015-04-01 – 2015-07-27 is available for zServer/ LPAR level reporting of forecast

Table 13: IBM Capacity Management Analytics Tables (HCM schema). (continued)

Table 13: IBM Capacity Management Analytics Tables (HCM schema). (continued)				
CMA TABLES (HCM schema)	**Number of rows**	**Feature**	**Report**	**Date Ranges (MIN/MAX dates) YYYY-MM-DD**
MVSPM_LPAR_CPU_BUSY_ MIPS_PRED (daily)	0	SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used - zServer/LPAR Level	N/A
MVSPM_LPAR_CPU_BUSY _MIPS_PRED (monthly)	0	SMO	CPU: MIPS Used - LPAR Level by WLM Importance CPU: MIPS Used - zServer/LPAR Level	N/A
NO89_PRODUCTS	23	SCA	SCA: NO89 Product MSU and Price SCA: NO89 Products Matrix	1900-01-01 1900-01-01
OPTIMIZATION_BILLABLE	39312	SCA	SCA: NO89 Product MSU and Price SCA: Product MSU and Price SCA: Registered Products Matrix SCA: Software Cost Analysis Summary	2015-05-02 - 2015-07-02
OPTIMIZATION_METADATA	1	SMO	CPU: LPAR Weight Optimization Run Result CPU: Over/ Under Shared Weight - CPC by LPAR	2015-05-21 -2015-05-23

CMA TABLES (HCM schema)	Number of rows	Feature	Report	Date Ranges (MIN/MAX dates) YYYY-MM-DD
OPTIMIZATION_ MOVEMENT	0	SCA	SCA Reports	N/A
PARENT_PROGRAMS	29	SCA	SCA Reports	N/A
PROD_MSU_FORECAST	257074	SCA	SCA: LPAR MSU Utilization SCA: Registered Products Matrix	2015-04-01 - 2016-04-30 The MSU forecast can go to 2016-04-30
PROD_MSU_OPTIMIZATION	40802	SCA	SCA: LPAR MSU Utilization SCA: Product MSU and Price SCA: Registered Products Matrix SCA: Software Cost Analysis Summary	2015-05-02 - 2015-07-02
SUBCAB_PROGRAMS	467	SCA	SCA: Product MSU and Price SCA: Registered Products Matrix SCA: Software Cost Analysis Summary SCA: NO89 Product MSU and Price	2013-01-01 2013-01-01
TIMEZONES	31	AA, SCA, SMO	CPU Reports, Application Analytics Reports, scar EPORTS, Memory Reports	N/A

Table 13: IBM Capacity Management Analytics Tables (HCM schema). (continued)

Dropping the sample database

Once you have used your production data, then the sample database must be dropped.

Procedure

1. Drop the feature base first by running the command `cmabase.py --dropdb --run`.
2. Drop the sample database by running the command `cmasampledb.py –dropdb --run`.

Chapter 6. Manually installing IBM Capacity Management Analytics 2.1.1

If you want to run IBM Capacity Management Analyticson a platform other than 64-bit Red Hat Enterprise Linux Server Edition version 6 and 7, you must install the product manually.

Before you begin

Ensure you have all the prerequisite software installed before beginning your manual installation of Capacity Management Analytics.

About this task

For operating systems other than 64-bit Red Hat Enterprise Linux Server Edition version 6 and 7, you cannot use the Solution Installer to run Capacity Management Analytics and must install the product manually.

Procedure

1. Extract the Capacity Management Analytics Solution Kit in your `CMAHOME` directory.
2. Update your profile. You can use the sample profile hcm.profile in `CMAHOME/samples`.
3. Run `cmaconfig.py`. `cmaconfig.py` configures Capacity Management Analytics to update the SPSS Modeler setup, Cognos setup, SDAP setup, and so on. This program must be run using your Capacity Management Analytics administrator ID and sudo.
   - `sudo CMAHOME=$CMAHOME PYTHONPATH=$PYTHONPATH ${CMAHOME}/bin/cmaconfig.py --COGNOSHOME=<COGNOSHOME> --SDAPHOME=<SDAPHOME> --SPSSHOME=<SPSSHOME> --CLEMBHOME=<CLEMBHOME> --JAVAHOME=<JAVA_HOME> --DB2_HOME=<DB2_HOME> --CPLEXHOME=<CPLEXHOME> --CMA_ADM=<CMA_ADM> --install`

     --COGNOSHOME=COGNOSHOME
     Cognos Home Directory

     --SDAPHOME=SDAPHOME
     SDAP Home Directory

     --SPSSHOME=SPSSHOME
     SPSS Modeler Server Home Directory

     --CLEMBHOME=CLEMBHOME
     SPSS Modeler Batch Home Directory

     --JAVAHOME=JAVAHOME
     JAVA Home Directory

     --DB2_HOME=DB2_HOME
     DB2 Home Directory

     --CPLEXHOME=CPLEXHOME
     CPLEX Home Directory

     --CMA_ADM=CMA_ADM
     CMA Admin ID

     --install
     Install CMA Configuration

     [--verbose]
     Verbose mode

[--trace=tracelevel]
Level of trace, 1-4

Note: To learn more about cmaconfig.py, run `CMAHOME/bin/cmaconfig.py -h`.

4. Restart Cognos to apply your settings.
5. Run `cmahome.py`. `cmahome.py` helps create your own configuration and installs Capacity Management Analytics features. It will create a cmainstance-config.js in the `CMAHOME/configuration` folder with your inputs. It will also install the Capacity Management Analytics features you have specified. This program must be run using your Capacity Management Analytics administrator ID.

- ```
  cmahome.py --CMADBHOST=<CMADBHOST> --HCM_INDEXBP=<HCM_INDEXBP>
  --HCM_BP=<HCM_BP> --CMADBPORT=<CMADBPORT> --CMADBSID=<CMADBSID>
  --CMADBLOCNAME=<CMADBLOCNAME> --HCM_CCSID=<HCM_CCSID> --
  HCM_STOGROUP=<HCM_STOGROUP> --WLM_GENERAL=<WLM_GENERAL> --
  HCM_DATABASE=<HCM_DATABASE> --HCM_SCHEMA=<HCM_SCHEMA> --
  DRL_DATABASE=<DRL_DATABASE> --DRL_SCHEMA=<DRL_SCHEMA> --
  DRL_STOGROUP=<DRL_STOGROUP> --HOST=<HOST> --DB2_HOME=<DB2_HOME> --
  CPLEXHOME=<CPLEXHOME> --SPSSHOME=<SPSSHOME> --CLEMBHOME=<CLEMBHOME>
  --SDAPHOME=<SDAPHOME> --COGNOSHOME=<COGNOSHOME> --
  COGNOSNAMESPACE=<COGNOSNAMESPACE> --JAVAHOME=<JAVAHOME> --CMA_ADM=<CMA_ADM>
  --CMA_PWD=<CMA_PWD> --CMADBAUSR=<CMADBAUSR> --CMADBAPWD=<CMADBAPWD> --
  DB2LUW_ADM=<DB2LUW_ADM> --DB2LUW_PWD=<DB2LUW_PWD> --HCM_USER=<HCM_USER> --
  HCM_PWD=<HCM_PWD> --CMADBALIAS=<CMADBALIAS> --feature=<FEATURE> --sample --
  install –run
  ```

--CMADBHOST=CMADBHOST
CMA Database Host Name

--HCM_INDEXBP=HCM_INDEXBP
CMA Database Index Buffer Pool

--HCM_BP=HCM_BP
CMA Database Buffer Pool

--CMADBPORT=CMADBPORT
CMA Database Port

--CMADBSID=CMADBSID
CMA Database Sub System ID

--CMADBLOCNAME=CMADBLOCNAME
CMA Database Location Name

--HCM_CCSID=HCM_CCSID
CMA Database CCSID Encoding

--HCM_STOGROUP=HCM_STOGROUP
CMA Database Storage Group Name

--WLM_GENERAL=WLM_GENERAL
DB2 for z/OS General WLM Environment Name

--HCM_DATABASE=HCM_DATABASE
CMA Database Name

--HCM_SCHEMA=HCM_SCHEMA
CMA Database Schema Name

--DRL_DATABASE=DRL_DATABASE
TDSz Database Name

--DRL_SCHEMA=DRL_SCHEMA
TDSz Database Schema Name
  ```

--DRL_STOGROUP=DRL_STOGROUP
TDSz Database Storage Group

--HOST=HOST
Host Name where Cognos / SPSS / CPLEX is installed

--DB2_HOME=DB2_HOME
DB2 Database Path

--CPLEXHOME=CPLEXHOME
CPLEX Home Directory

--SPSSHOME=SPSSHOME
SPSS Modeler Server Home Directory

--CLEMBHOME=CLEMBHOME
SPSS Modeler Batch Home Directory

--SDAPHOME=SDAPHOME
SDAP Home Directory

--COGNOSHOME=COGNOSHOME
Cognos Home Directory

--COGNOSNAMESPACE=COGNOSNAMESPACE
Cognos Security Namespace

--JAVAHOME=JAVAHOME
Java Path

--CMA_ADM=CMA_ADM
CMA Admin User

--CMA_PWD=CMA_PWD
CMA Admin Password

--CMADBAUSR=CMADBAUSR
CMA Database Admin User

--CMADBAPWD=CMADBAPWD
CMA Database Admin Password

--DB2LUW_ADM=DB2LUW_ADM
DB2LUW Admin User

--DB2LUW_PWD=DB2LUW_PWD
DB2LUW Admin Password

--HCM_USER=HCM_USER
CMA User Name

--HCM_PWD=HCM_PWD
CMA User Password

Note: HCM_USER and HCM_PWD use the zlinux machine and z/OS database. Their usernames and passwords should match. Passwords are case sensitive.

--CMADBALIAS=CMADBALIAS
Name of CMA Database Alias

--feature=FEATURE
Install CMA OR CMA specific features into your home. BASE (CMA base), AA (Application Analytics),PI (Anomaly Detection), CPF (Capacity Planning and Forecasting), SCA (Software Cost Analysis), SMO (Systems Management and Optimizations) Example: AA, CPF, SMO. Default is ALL

--install
Install CMA

--run
Run/Execute the generated commands for a feature

--sample
Will install the sample database feature

[--trace=tracelevel]
Trace Level, 1-4

[--verbose]
Verbose mode

Note: To learn more about cmahome.py, run `CMAHOME/bin/cmahome.py -h`.

6. Run cmainstance. cmainstance.py creates your own instance and installs Capacity Management Analytics features. You can also use cmainstance to create a database, set up Cognos, run SPSS Modeler streams and the CPLEX optimization program. All Capacity Management Analytics users, including the administrator, can use cmainstance to create and manage their instance. cmainstance will use the cmainstance-config.js in the `CMAHOME/configuration` folder and change the parameters according to what you specify on the command line and will save the file in your `CMAINSTANCE/configuration` directory. You can always update your configuration file by executing `cmainstance.py`. For instance, `cmainstance.py --HCM_SCHEMA=<HCM_SCHEMA>` .

   - `cmainstance.py --CMAINSTANCE=<CMAINSTANCE> --instance --run --sample --feature=<FEATURE>`

     --feature=FEATURE
     Install CMA OR CMA specific features into your instance. BASE (CMA base), AA (Application Analytics),PI (Anomaly Detection), CPF (Capacity Planning and Forecasting), SCA (Software Cost Analysis), SMO (Systems Management and Optimizations) Example: AA, CPF, SMO. Default is ALL.

     --instance
     Create CMA instance

     --run
     Run and execute the generated commands for a feature

     --sample
     Will create the sample database

     [--trace=tracelevel]
     Trace Level, 1-4

     [--verbose]
     Verbose mode

   The command will create your instance, install Capacity Management Analytics features into your instance, set up Cognos reports, create a sample database and a Capacity Management Analytics database. It will also execute streams with the sample data. You will the results in the Capacity Management Analyticsdatabase and Cognos reports.

   If you want to use your own data instead of the sample data, you can execute `cmainstance.py --instance` first to create your instance and install Capacity Management Analytics features in your instance. Then review and make changes corresponding to your data and run the individual `feature.py` program to do feature work.

   Note: To learn more about cmahome.py, run `CMAHOME/bin/cmainstance.py -h`.

Completing Capacity Management Analytics input files

This section will provide you with details on how to fill in your IBM Capacity Management Analytics input files.

cmabase.parms

-PBLOCKSIZE
You can decrease the number of times data is inserted into the database by increasing BLOCKSIZE, and the stream can run faster. However, larger BLOCKSIZE increases the duration of each insertion. BLOCKSIZE is measured in days. When you enter the BLOCKSIZE value, enter the number of days of data that you want to treat as a block. For example, 30.

-PSDATE
The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2015-01-01"

-PEDATE
The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2016-12-31"

-PSTART_DATE_DIS
The start date you want use for analyzing data on a distributed system. For example, "2015-01-01"

-PEND_DATE_DIS
The end date you want use for analyzing data on a distributed system. For example, "2015-12-31"

cmasca.parms

-PUSERSELECTION
“Define” or “weight”. “Define” indicates that only the define capacity is considered in the optimization algorithm. “Weight” indicates that both weighted capacity and define capacity are considered in the optimization algorithm.

-PDATEMIN
The SCA product forecast stream uses the date from DATEMIN to the latest date to forecast all of the product and LPAR combinations. For example, "2015-01-01"

-PODD
Obsolete detection days. If a product no longer exists or does not run for recent ODD days in one LPAR, we assume that this combination is obsolete and the stream does not forecast this combination. All of the obsoleted combinations appear in the product forecast log file. For example, 20.

-PREFRESH
“True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PIWP
“True” or “false”. IWP is Integrated Workload Pricing. For more information on IWP and the appropriate settings, see IWP Defining Programs (http://www.ibm.com/systems/z/resources/swprice/reference/exhibits/iwp.html)

-PGSSPFILE
The full path of the server location for the `gssp_scaling_factors.csv` file. For example, `CMAINSTANCEDATA/gssp_scaling_factors.csv`

-PMLCFILE
The full path of the server location for the `mlc_prod.csv` file. For example, `CMAINSTANCEDATA/mlc_prod.csv`.

-PIPLAFILE
The full path of the server location for the `ipla_prod.csv` file. For example, `CMAINSTANCEDATA/ipla_prod.csv`.

-PENTERPRISEFILE
The full path of the server location for the `enterprise.csv` file. For example, `CMAINSTANCEDATA/enterprise.csv`.

-PCURRENCYFILE
The full path of the server location for the `currency.csv` file. For example, `CMAINSTANCEDATA/currency.csv`.

-PNO89_FILE
The full path of the server location for the no89 file. For example, `CMAINSTANCEDATA/no89`.

-PCAREFILE
The full path of the server location for the `opt_care_prod.csv` optimization file. For example, `CMAINSTANCEDATA/opt_care_prod.csv`.

-PRESERVEDFILE
The full path of the server location for the `opt_reserved_prod.csv` optimization file. For example, `CMAINSTANCEDATA/opt_reserved_prod.csv`.

-PBOUNDFILE
The full path of the server location for the `opt_bound_prod.csv` optimization file. For example,`CMAINSTANCEDATA/opt_bound_prod.csv`.

cmasmo.parms

-PTOP
The amount of TOP records you want to get. For example, 10.

-PSTART_DATE_DIS
The start date you want use for analyzing data on a distributed system. For example, "2015-01-01"

-PEND_DATE_DIS
The end date you want use for analyzing data on a distributed system. For example, "2015-12-31"

-PREFRESH
“True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PBLOCKSIZE
You can decrease the number of times data is inserted into the database by increasing BLOCKSIZE, and the stream can run faster. However, larger BLOCKSIZE increases the duration of each insertion. BLOCKSIZE is measured in days. When you enter the BLOCKSIZE value, enter the number of days of data that you want to treat as a block. For example, 30.

-PSDATE
The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2015-01-01"

-PEDATE
The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2016-12-31"

cmapi.parms

-PANOMALY_RATIO
Defines the ratio that identifies abnormal transactions. A smaller ratio means that a smaller deviation identifies an anomaly, resulting in more anomalies being detected. A larger ratio means that a larger deviation is required to identify an anomaly, resulting in fewer anomalies being detected. The default value is 1.

-PMEMORY
The amount of memory on your modeler server machine.

-PCPU_RESPONSE_RATIO
Identifies what should be considered the abnormal ratio between response seconds and CPU seconds of a single transaction. If your normal CICS transactions have a small differential between response seconds and CPU seconds, the ratio you set should be small. When the differential between response

seconds and CPU seconds is larger than the CPU_RESPONSE_RATIO you specify, the transaction is considered abnormal. The default value is 50.

-PSCORE_RATIO
Defines the ratio that the tree-based anomaly model uses to judge whether the transaction is abnormal or not. The default value is 0.5. The recommended range of this parameter value is [0.5,0.6].

-PTRANSACTIONFILE
The full path of server location where the `transaction.csv` file is located. For example, `CMAINSTANCEDATA/transaction.csv`.

-PDATEFILE
The full path of server location where the `data_info.csv` file is located. For example `,CMAINSTANCEDATA/date_info.csv`.

cmacpf.parms

-PDATE_START
The start date you want to use for MIPS forecast. For example, “2015-01-01”

-PDATE_END
The start date you want to use for MIPS forecast. For example, “2015-12-31”

-PHOUR_START
The start hour of the observed data that you are using. For example, 0.

-PHOUR_END
The end hour of the observed data that you are using. For example, 23".

-PPEAKRANK
The Nth highest peak rank values. Valid values are in the range [1,5]. You can also input multiple rank values. For example, 1, 2, 3, and 4 represent highest peak, 2nd highest peak, 3rd highest peak, and 4th highest peak. For example, "1,2"

cmaaa.parms

-PDATE_START_APPL
The start date that you want use for capture ratio data. For example, 2015-01-01

-PDATE_END_APPL
The end date that you want use for capture ratio data. For example, 2016-01-01

-PTIME_START_APPL
The start time that you want use for capture ratio data. The default value is 00:00:00.

-PTIME_END_APPL
The end time that you want use for capture ratio data. The default value is 23:59:59.

-PREFRESH
“True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PDATE_START
The start date of the observed data that you are using. For example, "2015-01-01".

-PDATE_END
The end date of the observed data that you are using. For example, "2015-06-30"

-PHOUR_START
The start hour of the observed data that you are using. For example, 0.

-PHOUR_END
The end hour of the observed data that you are using. For example, 23"

-PPEAKRANK
The Nth highest peak rank values. Valid values are in the range [1,5]. You can also input multiple rank values. For example, “1,2”

-PAPPFILE

The full path of server location where the `appl_lob_z.csv` file is located. For example, `CMAINSTANCEDATA/appl_lob_z.csv`.

-PAPPFILEDIS

The full path of server location where the `appl_lob.csv` file is located. For example `CMAINSTANCEDATA/appl_lob.csv`.

mlc_prod.csv

This file defines the licensing, pricing, and system information for each MLC product for which you want to view monetary values in the Software Cost Analysis reports and workspace. The monetary values are expressed in the currency that is defined for an enterprise in the `currency.csv` file.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. The CPU serial numbers that are listed here must match with the serial numbers that are listed in the `enterprise.csv` file.

PID is an eight-character field that must follow the IBM format (XXXX-XXX) for a product ID.

LICENSE is a six-character field. It must be one of the acronyms that are listed here. To view details on each license, refer to IBM System z Software Pricing (www.ibm.com/systems/z/resources/swprice/mlc/index.html).

- AWLC – Advanced Workload License Charge
- AEWLC – Advanced Entry Workload License Charge
- EWLC – Entry Workload License Charge
- VWLC – Variable Workload License Charge
- MWLC – Midrange Workload License Charge
- ZNALC – System z New Application Workload License Charge

TIER is a five-character field that specifies the tier/level for an associated LICENSE. To view details on the tiers/levels for each license, refer to IBM System z Software Pricing (www.ibm.com/systems/z/resources/swprice/mlc/index.html).

PRICE_PER_MSU is a decimal value that goes up to two decimal places.

ipla_prod.csv

This file defines the VUE, pricing, and system information for each IPLA product for which you want to view monetary values in the Software Cost Analysis reports and workspace. The entitlement for each IPLA is by CPU Serial Number, not the entire logical enterprise.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. The CPU serial numbers that are listed here must match with the serial numbers that are listed in the `enterprise.csv` file.

PID is an eight-character field that must follow the IBM format (XXXX-XXX) for a product ID.

SS_PID is an eight character field that must follow the IBM format (XXXX-XXX) for the Subscription and Support product ID. This value can be left blank if Subscription and Support were not purchased for a product.

VUE is the six-character field for the Value Unit Exhibit (VUE) for a product. It must be one of the VUEs listed here. Refer to Overview of Value Unit Exhibits (www.ibm.com/systems/z/resources/swprice/zipla/vue.html) for more details.

- VUE001
- VUE007
- VUE020

ASSIGNED_ENTITLED_MSU is a numeric field. The entitled MSUs are specifically for the corresponding CPU_SERIAL_NO, not the entire enterprise entitlement. Adding up each entitlement totals the enterprise entitlement.

PRICE_PER_VU is a decimal value that goes up to two decimal places.

ASSIGNED_SS_ENTITLED_MSU is a numeric field. The entitled MSUs are specifically for the corresponding CPU_SERIAL_NO, not the entire enterprise entitlement. This value can be left blank if Subscription and Support were not purchased for a product.

SS_PRICE_PER_MTH specifies the monthly price of Subscription and Support. It is a decimal value that goes up to two decimal places.

enterprise.csv

This file sets the CPU serial number for each enterprise.

ENTERPRISE is a 20-character enterprise name. The enterprise names in the `enterprise.csv` file must match with the names defined in the `currency.csv` file.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. If a CPU serial number is missing in this table, then the **CPU Serial No** is missing in "Software Cost Analysis Summary workspace" on page 277.

currency.csv

This file defines the currency of each logically named enterprise. An enterprise is defined as a grouping of mainframe systems that all have the same currency.

ENTERPRISE is a 20-character enterprise name.

CURRENCY is the three character international symbol for the currency that is used to display monetary values for the associated ENTERPRISE in reports.

Note: The default decimal separator for values in these files is the period, which is the ISO format. If a non-ISO format is used, then the SPSS Modeler stream properties must be modified to use the same decimal.

`gssp_scaling_factors.csv` file:

- PROD_ID: the IBM product ID registered on the System z machine.
- GSSP_SCALING_FACTOR: the GSSP scaling factor for the product.

no89

Prepare the registration file of NO89 products. The file should follow the conventions of NO89 DD statements that are described in the topic "Modifying the NO89 DD statement" in the *SCRT Users Guide* (www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=ZSL03022USEN). The definition section of the file should begin with `//NO89` and end with `//SYSLIN DD`

For example:

```
//no89
 * LOTUS ACTIVE INSIGHT FOR Z/OS V6
 5655-U15=2827-BAA6:JB0,2817-4E15:JE0,2097-99FF:JC0
 * COBOL FOR OS/390 & VM V2
 5648-A25=CF3,CF2,J90,JB0
 * VA PL/I FOR OS/390 V2
 5655-B22=*NONE
 //SYSLIN  DD
```

holiday.csv

1. Date—The date of the holiday or the day on which the MIPS volume is notable.
2. Isholiday—Whether this date is a holiday. `1` indicates that the date is a holiday, `0` indicates that it is not.
3. D_name—The name of the holiday. Note that this column is for reference and is not required.
4. Ishigh—Whether the MIPS volume on a date is high over the data range. `1` indicates MIPS_USED is high for this date, `0` means MIPS_USED is not high for this date.
5. Islow—Whether the MIPS volume on a date is low below the data range. `1` indicates MIPS_USED is low for this date, `0` means MIPS_USED is not low for this date.

Note: Date format in `holiday.csv` should be ISO format (YYYY-MM-DD). If using a non-ISO format, the stream properties of forecast streams must be updated to match the date format in `holiday.csv`. For more information about modifying the `holiday.csv` file, see "Time period definition" on page 119.

Tip: If you use Microsoft Excel to edit csv files and your computer uses a different locale than what is used in the SPSS Modeler streams, you must ensure that the data format in the csv remains consistent with date format that is used in the streams. Microsoft Excel can change the locale that is used in the file to match your computer's locale.

opt_care_prod.csv

This file is used to identify the cared products. If you do not need to identify any products as cared, do not enter anything in this file.

1. For any product that is considered a cared product, enter the product ID in the PROD_ID column.
2. For each cared product, enter 1 in the CARED column. If a PROD_ID is entered, but CARED is not set or is set to 0, this product is considered as not cared.

opt_bound_prod.csv

This file is used to identify bounded products. If you do not need to identify any products as bounded, do not enter anything in this file.

1. If two or more products need to be bounded, enter the product IDs individually in the PROD_ID column.
2. Enter the identical numeric identifier for the products to be bounded in the BOUND_ID column. The numeric identifier must be unique for each group of bounded products. For example, if product IDs 5635-A03 and 5655-R36 should be bounded, enter each product ID in separate cells in the PROD_ID column, and assign the numeric identifier 1 in the BOUND_ID column for each product.

PROD_ID	BOUND_ID
5635-A03	1
5655-R36	1
5655-S97	2
5605-DB2	2

3. Repeat steps 1 and 2 for each group of products that should be bounded.

opt_reserved_prod.csv

This file is used to identify reserved products. If you do not need to identify any products as reserved, do not enter anything in this file.

1. For each product to be reserved, enter the product ID in the **PROD_ID** column.
2. Enter the name of the LPAR reserved for the product in the **LPAR_NAME** column.
3. Enter the CPC number of the LPAR in the **CPC_SERIAL_NO** column.
4. Set the reserved flag in the **RESERVED** column.
5. Enter the reserved amount by MSU in the **RESERVED_MSU** column.

PROD_ID	LPAR_NAME	CPC_SERIAL_NO	RESERVED	RESERVED_MSU
5635-A03	LP1	AAAA	1	10
5655-R36	LP2	BBBB	1	10

If **RESERVED** is 1 and **RESERVED_MSU** is *n*, the product is arranged in the specified LPAR with an amount equal or above *n*. If **RESERVED** is 1 and **RESERVED_MSU** is 0, the product is arranged in the specified LPAR with an amount above 0. If **RESERVED** is 1 and **RESERVED_MSU** is *n*, but *n* is beyond the MSU the product can produce, then the product is kept in its original LPAR after optimization. If the product is reserved in an LPAR that can not host it for any reason, then the reservation is ignored and the stream seeks any suitable LPAR for the product.

transaction.csv

Specify the transactions for which you want to perform anomaly detection. This file contains a single column, TRANSACTION_ID.

Note: You must enter the full value for the TRANSACTION_ID. You cannot use wildcard characters.

date_info.csv

This file contains three columns. Do not modify the value in the first column. In the second column, specify the date at which anomaly detection scoring should begin for the specified TRANSACTION_ID. Leave the third column value empty. The third column can optionally specify the end date for anomaly detection scoring, but Capacity Management Analytics generates scoring data for all dates after the specified start date.

For information on the following files, refer to "Configuring input files for LPAR weight optimization" on page 114.

- lwo_parameters.csv
- lwo_parameters_model.csv
- lwo_parameters_date_list.csv
- lwo_parameters_per_lpar.csv

For information on the following files, refer to "Mapping table definitions for application analytics reporting" on page 146.

- appl_lob_z.csv
- appl_lob.csv

Chapter 7. Manually running features for IBM Capacity Management Analytics 2.1.1

Before you begin

Prior to running any IBM Capacity Management Analytics feature programs, ensure you have the following environment variables set in your profile:

- CMAINSTANCE
- CMAHOME
- PYTHONPATH=$CMAHOME/bin:$CMAHOME/lib:$PYTHONPATH

About this task

Capacity Management Analytics feature programs help you to manage your features. For instance, they enable you to create instances, cleaning up instances, and uninstalling features. They also help perform all feature functions such as creating and removing databases, running SPSS streams, and running CPLEX optimization programs.

Procedure

1. After you install the Capacity Management Analytics Solution Kit, run `cmahome.py` and create your own cmainstance by running `cmainstance.py`.

   Note: You can run the feature program to invoke other feature functions, for instance, `cmabase.py --createdb –run` to create a database. For more information, see each Capacity Management Analytics feature program.
2. Before you begin invoking feature functionality to execute SPSS streams or the CPLEX optimization program, ensure you have the required input files in your `CMAINSTANCE/data` folder and provide the correct information.

   Note: If you make any customized change to any file in the folder, update the file to read only.

Table 14: Summary of stream and input files by feature

Feature name-mode	Streams	Required input files
smo -- etl	channel_top_etl.str device_top_etl.str device_lpar_agg_etl.str	cmasmo.parms
smo -- refresh_roll	channel_top_etl_refresh_roll.str device_top_etl_refresh_roll.str device_lpar_agg_etl_refresh_roll.str	cmasmo.parms
smo – refresh_full	channel_top_etl_refresh_full.str device_top_etl_refresh_full.str device_lpar_agg_etl_refresh_full.str	cmasmo.parms

Table 14: Summary of stream and input files by feature(continued)		
Feature name-mode	**Streams**	**Required input files**
smo -- optimization	Lwo.jar	lwo_parameters.csv lwo_parameters_model.csv lwo_parameters_date_list.csv lwo_parameters_per_lpar.csv
sca -- etl	sca_subcap.str sca_tiers.str sca_no89.str	cmasca.parms gssp_scaling_factors.csv no89
sca -- pricing	sca_pricing.str	cmasca.parms mlc_prod.csv ipla_prod.csv enterprise.csv currency.csv
sca -- forecast	msu_gssp_forecast_timeseries.str msu_prod_forecast_timeseries.str msu_lpar_forecast_timeseries.str	cmasca.parms holiday.csv
sca -- optmization	prod_bill_msu_calculate.str prod_bill_msu_optimization.str	cmasca.parms opt_bound_prod.csv opt_care_prod.csv opt_reserved_prod.csv
aa -- etl	appl_lob_z.str appl_lob.str appl_capture_ratio.str appl_lob_z.csv appl_lob.csv	cmaaa.parms
aa -- forecast	appl_avg_daily_forecast_timeseries.str appl_avg_hourly_forecast_timeseries.str appl_peak_daily_forecast_timeseries.str	cmaaa.parms holiday.csv
cpf -- forecast	mvspm-lpar-hourly-forecast-timeseries.str mvspm-lpar-daily-forecast-timeseries.str mvspm-lpar-monthly-forecast-timeseries.str mvspm-lpar-daily-peak-forecast-timeseries.str mvspm-cpc-daily-peak-forecast-timeseries.str	cmacpf.parms holiday.csv

Table 14: Summary of stream and input files by feature(continued)		
Feature name-mode	**Streams**	**Required input files**
pi -- anomaly	anomaly_clean_up.str anomaly_detect_building_r.str anomaly_detect_scoring.str anomaly_detect_building_c.str anomaly_detect_scoring.str anomaly_detect_treebased.str anomaly_classification.str	cmapi.parms date_info.csv transaction.csv

3. Optional. To run a sample database, CMASAMPLEDB, run `cmasamledb.py`.

   --instance
   Manage CMA sample database instance. If run is specified, invoke `createdb` and import. Import should be the same format as `createdb`.

   --createdb
   Generate the sql and commands to create the CMA sample database

   --import
   Generate the sql and commands to import data into the CMA sample database

   --run
   Run and execute the generated SQL and/or commands

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

   Note: To learn more about cmasampledb.py, run `cmasampledb.py -h`.

4. To create the feature directory structure for this CMA instance, CMABASE, run `cmabase.py`.

   Note: CMABASE is the prerequisite of CMASMO, CMAAA, CMASCA, CMAPI, and CMACPF detailed in the following steps.

   --instance
   Create feature directory structure for this CMA instance. If run is specified, invoke creatdb, defdata, setup, and etl

   --createdb
   Generate the sql and commands to create the CMA data warehouse

   --defdata
   Generate the commands to create the default data for this feature

   --setup
   Generate the commands to set up Cognos: deploy images, import reports and update FM

   --etl
   Generate the commands etl data warehouse data for this feature

   --refresh_roll
   Run BASE etl (refresh-roll) streams

   --refresh_full
   Run BASE etl (refresh-full) streams

   --run
   Run and execute the generated SQL commands

[--verbose]
Verbose

[--trace=TRACE]
Trace Level

Note: To learn more about cmabase.py, run `cmabase.py -h`.

5. Optional. To create a SMO instance, cmasmo, and to create directories and prepare files and data, run cmasmo.py.

   --instance
   Create SMO instance - create directories and prepare files and data. If --run is specified, invokeetl and lwo optimization. Lwo optimization should be the same format as etl.

   --etl
   Run SMO etl streams

   --optimization=OPTIMIZATION
   Optimization model to run. lwo--Lpar Weight Optimization

   --refresh_roll
   Run SMO etl (refresh-roll) streams

   --refresh_full
   Run SMO etl (refresh-full) streams

   --run
   Run and execute the generated commands

   [--plwo=PLWO]
   Parameter file for optimization. Default: `CMAINSTANCE/data/lwo_parameters.csv`

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

   Note: To learn more about cmasmo.py, run `cmasmo.py -h`.

6. Optional. To create a SCA instance, cmasca, and to create directories and prepare files and data, run cmasca.py.

   --instance
   Create SCA instance: create directories and prepare files and data. If --run is specified, invoke etl, defdata, pricing, forecast and optimization

   --etl
   Run SCA etl streams

   --defdata
   Generate SCA default data

   --pricing
   Run SCA pricing stream

   --forecast
   Run SCA forecast streams

   --optimization
   Run SCA optimization streams

   --run
   Run and execute the generated commands

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

Note: To learn more about cmasca.py, run `cmasca.py -h`.

7. Optional. To create an AA instance, CMAAA, and to create directories and prepare files and data, run `cmaaa.py`.

   --instance
   Create AA instance: create directories and prepare files and data. If --run is specified, invoke etl and forecast

   --etl
   Run AA etl streams

   --forecast
   Run AA forecast streams

   --run
   Run and execute the generated commands

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

   Note: To learn more about cmaaa.py, run `cmaaa.py -h`.

8. Optional. To create a CPF instance, cmacpf, and to create directories and prepare files and data, runcmacpf.py.

   --instance
   Create CPF instance: create directories and prepare files and data. If --run is specified, invoke forecast

   --forecast
   Run forecast streams

   --run
   Run and execute the generated commands

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

   Note: To learn more about cmacpf.py, run `cmacpf.py -h`.

9. Optional. To create a PI instance, cmapi, and to create directories and prepare files and data, run cmapi.py.

   --instance
   Create PI instance: create directories and prepare files and data. If --run is specified, invoke anomaly

   --anomaly
   Run anomaly detection streams

   --run
   Run and execute the generated commands

   [--verbose]
   Verbose

   [--trace=TRACE]
   Trace Level

   Note: To learn more about cmapi.py, run `cmapi.py -h`.

Chapter 8. Completing Capacity Management Analytics input files

This section will provide you with details on how to fill in your IBM Capacity Management Analytics input files.

cmabase.parms

-PBLOCKSIZE
: You can decrease the number of times data is inserted into the database by increasing BLOCKSIZE, and the stream can run faster. However, larger BLOCKSIZE increases the duration of each insertion. BLOCKSIZE is measured in days. When you enter the BLOCKSIZE value, enter the number of days of data that you want to treat as a block. For example, 30.

-PSDATE
: The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2015-01-01"

-PEDATE
: The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2016-12-31"

-PSTART_DATE_DIS
: The start date you want use for analyzing data on a distributed system. For example, "2015-01-01"

-PEND_DATE_DIS
: The end date you want use for analyzing data on a distributed system. For example, "2015-12-31"

cmasca.parms

-PUSERSELECTION
: “Define” or “weight”. “Define” indicates that only the define capacity is considered in the optimization algorithm. “Weight” indicates that both weighted capacity and define capacity are considered in the optimization algorithm.

-PDATEMIN
: The SCA product forecast stream uses the date from DATEMIN to the latest date to forecast all of the product and LPAR combinations. For example, "2015-01-01"

-PODD
: Obsolete detection days. If a product no longer exists or does not run for recent ODD days in one LPAR, we assume that this combination is obsolete and the stream does not forecast this combination. All of the obsoleted combinations appear in the product forecast log file. For example, 20.

-PREFRESH
: “True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PIWP
: “True” or “false”. IWP is Integrated Workload Pricing. For more information on IWP and the appropriate settings, see IWP Defining Programs (http://www.ibm.com/systems/z/resources/swprice/reference/exhibits/iwp.html)

-PGSSPFILE
: The full path of the server location for the `gssp_scaling_factors.csv` file. For example, `CMAINSTANCEDATA/gssp_scaling_factors.csv`

-PMLCFILE
: The full path of the server location for the `mlc_prod.csv` file. For example, `CMAINSTANCEDATA/mlc_prod.csv`.

-PIPLAFILE
: The full path of the server location for the `ipla_prod.csv` file. For example, `CMAINSTANCEDATA/ipla_prod.csv`.

-PENTERPRISEFILE
: The full path of the server location for the `enterprise.csv` file. For example, `CMAINSTANCEDATA/enterprise.csv`.

-PCURRENCYFILE
: The full path of the server location for the `currency.csv` file. For example, `CMAINSTANCEDATA/currency.csv`.

-PNO89_FILE
: The full path of the server location for the no89 file. For example, `CMAINSTANCEDATA/no89`.

-PCAREFILE
: The full path of the server location for the `opt_care_prod.csv` optimization file. For example, `CMAINSTANCEDATA/opt_care_prod.csv`.

-PRESERVEDFILE
: The full path of the server location for the `opt_reserved_prod.csv` optimization file. For example, `CMAINSTANCEDATA/opt_reserved_prod.csv`.

-PBOUNDFILE
: The full path of the server location for the `opt_bound_prod.csv` optimization file. For example,`CMAINSTANCEDATA/opt_bound_prod.csv`.

cmasmo.parms

-PTOP
: The amount of TOP records you want to get. For example, 10.

-PSTART_DATE_DIS
: The start date you want use for analyzing data on a distributed system. For example, "2015-01-01"

-PEND_DATE_DIS
: The end date you want use for analyzing data on a distributed system. For example, "2015-12-31"

-PREFRESH
: “True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PBLOCKSIZE
: You can decrease the number of times data is inserted into the database by increasing BLOCKSIZE, and the stream can run faster. However, larger BLOCKSIZE increases the duration of each insertion. BLOCKSIZE is measured in days. When you enter the BLOCKSIZE value, enter the number of days of data that you want to treat as a block. For example, 30.

-PSDATE
: The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2015-01-01"

-PEDATE
: The start date that you want use for MIPS, CHANNEL, and DEVICE related data. For example, "2016-12-31"

cmapi.parms

-PANOMALY_RATIO
: Defines the ratio that identifies abnormal transactions. A smaller ratio means that a smaller deviation identifies an anomaly, resulting in more anomalies being detected. A larger ratio means that a larger deviation is required to identify an anomaly, resulting in fewer anomalies being detected. The default value is 1.

-PMEMORY
: The amount of memory on your modeler server machine.

-PCPU_RESPONSE_RATIO
Identifies what should be considered the abnormal ratio between response seconds and CPU seconds of a single transaction. If your normal CICS transactions have a small differential between response seconds and CPU seconds, the ratio you set should be small. When the differential between response seconds and CPU seconds is larger than the CPU_RESPONSE_RATIO you specify, the transaction is considered abnormal. The default value is 50.

-PSCORE_RATIO
Defines the ratio that the tree-based anomaly model uses to judge whether the transaction is abnormal or not. The default value is 0.5. The recommended range of this parameter value is [0.5,0.6].

-PTRANSACTIONFILE
The full path of server location where the `transaction.csv` file is located. For example, `CMAINSTANCEDATA/transaction.csv`.

-PDATEFILE
The full path of server location where the `data_info.csv` file is located. For example `,CMAINSTANCEDATA/date_info.csv`.

cmacpf.parms

-PDATE_START
The start date you want to use for MIPS forecast. For example, “2015-01-01”

-PDATE_END
The start date you want to use for MIPS forecast. For example, “2015-12-31”

-PHOUR_START
The start hour of the observed data that you are using. For example, 0.

-PHOUR_END
The end hour of the observed data that you are using. For example, 23".

-PPEAKRANK
The Nth highest peak rank values. Valid values are in the range [1,5]. You can also input multiple rank values. For example, 1, 2, 3, and 4 represent highest peak, 2nd highest peak, 3rd highest peak, and 4th highest peak. For example, "1,2"

cmaaa.parms

-PDATE_START_APPL
The start date that you want use for capture ratio data. For example, 2015-01-01

-PDATE_END_APPL
The end date that you want use for capture ratio data. For example, 2016-01-01

-PTIME_START_APPL
The start time that you want use for capture ratio data. The default value is 00:00:00.

-PTIME_END_APPL
The end time that you want use for capture ratio data. The default value is 23:59:59.

-PREFRESH
“True” or “false”. If you want to replace current data, then set it to true, otherwise set it to false.

-PDATE_START
The start date of the observed data that you are using. For example, "2015-01-01".

-PDATE_END
The end date of the observed data that you are using. For example, "2015-06-30"

-PHOUR_START
The start hour of the observed data that you are using. For example, 0.

-PHOUR_END
The end hour of the observed data that you are using. For example, 23"

-PPEAKRANK
: The Nth highest peak rank values. Valid values are in the range [1,5]. You can also input multiple rank values. For example, “1,2”

-PAPPFILE
: The full path of server location where the `appl_lob_z.csv` file is located. For example, `CMAINSTANCEDATA/appl_lob_z.csv`.

-PAPPFILEDIS
: The full path of server location where the `appl_lob.csv` file is located. For example `CMAINSTANCEDATA/appl_lob.csv`.

mlc_prod.csv

This file defines the licensing, pricing, and system information for each MLC product for which you want to view monetary values in the Software Cost Analysis reports and workspace. The monetary values are expressed in the currency that is defined for an enterprise in the `currency.csv` file.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. The CPU serial numbers that are listed here must match with the serial numbers that are listed in the `enterprise.csv` file.

PID is an eight-character field that must follow the IBM format (XXXX-XXX) for a product ID.

LICENSE is a six-character field. It must be one of the acronyms that are listed here. To view details on each license, refer to IBM System z Software Pricing (www.ibm.com/systems/z/resources/swprice/mlc/index.html).

- AWLC – Advanced Workload License Charge
- AEWLC – Advanced Entry Workload License Charge
- EWLC – Entry Workload License Charge
- VWLC – Variable Workload License Charge
- MWLC – Midrange Workload License Charge
- ZNALC – System z New Application Workload License Charge

TIER is a five-character field that specifies the tier/level for an associated LICENSE. To view details on the tiers/levels for each license, refer to IBM System z Software Pricing (www.ibm.com/systems/z/resources/swprice/mlc/index.html).

PRICE_PER_MSU is a decimal value that goes up to two decimal places.

ipla_prod.csv

This file defines the VUE, pricing, and system information for each IPLA product for which you want to view monetary values in the Software Cost Analysis reports and workspace. The entitlement for each IPLA is by CPU Serial Number, not the entire logical enterprise.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. The CPU serial numbers that are listed here must match with the serial numbers that are listed in the `enterprise.csv` file.

PID is an eight-character field that must follow the IBM format (XXXX-XXX) for a product ID.

SS_PID is an eight character field that must follow the IBM format (XXXX-XXX) for the Subscription and Support product ID. This value can be left blank if Subscription and Support were not purchased for a product.

VUE is the six-character field for the Value Unit Exhibit (VUE) for a product. It must be one of the VUEs listed here. Refer to Overview of Value Unit Exhibits (www.ibm.com/systems/z/resources/swprice/zipla/vue.html) for more details.

- VUE001
- VUE007
- VUE020

ASSIGNED_ENTITLED_MSU is a numeric field. The entitled MSUs are specifically for the corresponding CPU_SERIAL_NO, not the entire enterprise entitlement. Adding up each entitlement totals the enterprise entitlement.

PRICE_PER_VU is a decimal value that goes up to two decimal places.

ASSIGNED_SS_ENTITLED_MSU is a numeric field. The entitled MSUs are specifically for the corresponding CPU_SERIAL_NO, not the entire enterprise entitlement. This value can be left blank if Subscription and Support were not purchased for a product.

SS_PRICE_PER_MTH specifies the monthly price of Subscription and Support. It is a decimal value that goes up to two decimal places.

enterprise.csv

This file sets the CPU serial number for each enterprise.

ENTERPRISE is a 20-character enterprise name. The enterprise names in the `enterprise.csv` file must match with the names defined in the `currency.csv` file.

CPU_SERIAL_NO is the last 4 characters of the machine serial number. If a CPU serial number is missing in this table, then the **CPU Serial No** is missing in "Software Cost Analysis Summary workspace" on page 277.

currency.csv

This file defines the currency of each logically named enterprise. An enterprise is defined as a grouping of mainframe systems that all have the same currency.

ENTERPRISE is a 20-character enterprise name.

CURRENCY is the three character international symbol for the currency that is used to display monetary values for the associated ENTERPRISE in reports.

Note: The default decimal separator for values in these files is the period, which is the ISO format. If a non-ISO format is used, then the SPSS Modeler stream properties must be modified to use the same decimal.

`gssp_scaling_factors.csv` file:

- PROD_ID: the IBM product ID registered on the System z machine.
- GSSP_SCALING_FACTOR: the GSSP scaling factor for the product.

no89

Prepare the registration file of NO89 products. The file should follow the conventions of NO89 DD statements that are described in the topic "Modifying the NO89 DD statement" in the *SCRT Users Guide* (www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=ZSL03022USEN). The definition section of the file should begin with `//NO89` and end with `//SYSLIN DD`

For example:

```
//no89
 * LOTUS ACTIVE INSIGHT FOR Z/OS V6
 5655-U15=2827-BAA6:JB0,2817-4E15:JE0,2097-99FF:JC0
 * COBOL FOR OS/390 & VM V2
 5648-A25=CF3,CF2,J90,JB0
 * VA PL/I FOR OS/390 V2
 5655-B22=*NONE
 //SYSLIN  DD
```

holiday.csv

1. Date—The date of the holiday or the day on which the MIPS volume is notable.
2. Isholiday—Whether this date is a holiday. 1 indicates that the date is a holiday, 0 indicates that it is not.
3. D_name—The name of the holiday. Note that this column is for reference and is not required.
4. Ishigh—Whether the MIPS volume on a date is high over the data range. 1 indicates MIPS_USED is high for this date, 0 means MIPS_USED is not high for this date.
5. Islow—Whether the MIPS volume on a date is low below the data range. 1 indicates MIPS_USED is low for this date, 0 means MIPS_USED is not low for this date.

Note: Date format in `holiday.csv` should be ISO format (YYYY-MM-DD). If using a non-ISO format, the stream properties of forecast streams must be updated to match the date format in `holiday.csv`. For more information about modifying the `holiday.csv` file, see "Time period definition" on page 119.

Tip: If you use Microsoft Excel to edit csv files and your computer uses a different locale than what is used in the SPSS Modeler streams, you must ensure that the data format in the csv remains consistent with date format that is used in the streams. Microsoft Excel can change the locale that is used in the file to match your computer's locale.

opt_care_prod.csv

This file is used to identify the cared products. If you do not need to identify any products as cared, do not enter anything in this file.

1. For any product that is considered a cared product, enter the product ID in the PROD_ID column.
2. For each cared product, enter 1 in the CARED column. If a PROD_ID is entered, but CARED is not set or is set to 0, this product is considered as not cared.

opt_bound_prod.csv

This file is used to identify bounded products. If you do not need to identify any products as bounded, do not enter anything in this file.

1. If two or more products need to be bounded, enter the product IDs individually in the PROD_ID column.
2. Enter the identical numeric identifier for the products to be bounded in the BOUND_ID column. The numeric identifier must be unique for each group of bounded products. For example, if product IDs 5635-A03 and 5655-R36 should be bounded, enter each product ID in separate cells in the PROD_ID column, and assign the numeric identifier 1 in the BOUND_ID column for each product.

PROD_ID	BOUND_ID
5635-A03	1
5655-R36	1
5655-S97	2
5605-DB2	2

3. Repeat steps 1 and 2 for each group of products that should be bounded.

opt_reserved_prod.csv

This file is used to identify reserved products. If you do not need to identify any products as reserved, do not enter anything in this file.

1. For each product to be reserved, enter the product ID in the **PROD_ID** column.
2. Enter the name of the LPAR reserved for the product in the **LPAR_NAME** column.
3. Enter the CPC number of the LPAR in the **CPC_SERIAL_NO** column.
4. Set the reserved flag in the **RESERVED** column.
5. Enter the reserved amount by MSU in the **RESERVED_MSU** column.

PROD_ID	LPAR_NAME	CPC_SERIAL_NO	RESERVED	RESERVED_MSU
5635-A03	LP1	AAAA	1	10
5655-R36	LP2	BBBB	1	10

If **RESERVED** is 1 and **RESERVED_MSU** is *n*, the product is arranged in the specified LPAR with an amount equal or above *n*. If **RESERVED** is 1 and **RESERVED_MSU** is 0, the product is arranged in the specified LPAR with an amount above 0. If **RESERVED** is 1 and **RESERVED_MSU** is *n*, but *n* is beyond the MSU the product can produce, then the product is kept in its original LPAR after optimization. If the product is reserved in an LPAR that can not host it for any reason, then the reservation is ignored and the stream seeks any suitable LPAR for the product.

transaction.csv

Specify the transactions for which you want to perform anomaly detection. This file contains a single column, TRANSACTION_ID.

Note: You must enter the full value for the TRANSACTION_ID. You cannot use wildcard characters.

date_info.csv

This file contains three columns. Do not modify the value in the first column. In the second column, specify the date at which anomaly detection scoring should begin for the specified TRANSACTION_ID. Leave the third column value empty. The third column can optionally specify the end date for anomaly detection scoring, but Capacity Management Analytics generates scoring data for all dates after the specified start date.

For information on the following files, refer to "Configuring input files for LPAR weight optimization" on page 114.

- lwo_parameters.csv
- lwo_parameters_model.csv
- lwo_parameters_date_list.csv
- lwo_parameters_per_lpar.csv

For information on the following files, refer to "Mapping table definitions for application analytics reporting" on page 146.

- appl_lob_z.csv
- appl_lob.csv

Chapter 9. Report, workspace, and folder customization

IBM Capacity Management Analytics lets you create custom reports, workspaces, and folders that satisfy your unique analytic requirements. You can modify the reports, workspaces, and folders by using IBM Cognos Business Intelligence tools.

When you create custom reports or workspaces, save to a location outside of the **IBM Capacity Management Analytics Solution Kit** folder. Saving to a location outside of the **IBM Capacity Management Analytics Solution Kit** folder ensures that your reports and workspaces are not overwritten when updates occur to Capacity Management Analytics.

When you create a new folder, you can move any report, object, or template into the folder. If you cut and paste reports, but you want to keep the report shortcut in the **Reports** > **All** folder in IBM Cognos Connection, then you should create a new shortcut in the **Public Folders** > **IBM Capacity Management Analytics Solution Kit** > **Reports** > **All** folder.

Customizing a Capacity Management Analytics report

There are four options to build a custom report with IBM Capacity Management Analytics, listed in order of preference:

- Building a custom report from an existing report
- Building a custom report from a category template
- Building a custom report from the base template
- Building a custom report without a template

To use any of the first three options, you must copy a report or template from the IBM Capacity Management Analytics Solution Kit folder to your own folder. For detailed instructions on copying reports, see IBM Knowledge Center Cognos Business Intelligence welcome page (http://www-01.ibm.com/support/knowledgecenter/SSEP7J/welcome).

Important: If you are using data from tables that are at different time intervals, you must normalize or roll up the lowest granularity to the highest before you can join the data. For example, Table 1 has an interval of 30 minutes and Table 2 has an interval of 60 minutes. Before you join the data, you must ensure that Table 1 aggregates two 30 minute intervals to equal one 60 minute interval so that it matches the rows in Table 2.

There are several options for creating a custom workspace, as described in "Customizing a Capacity Management Analytics workspace" on page 89.

Building a custom report from an existing report

You can edit an existing IBM Capacity Management Analytics report to create a custom report. For example, you can add a chart to an existing report.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Reports` directory.
2. Open the directory that contains the report that you want to customize.

   Do not copy from the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Reports\All` directory. This directory contains report shortcuts, with each shortcut pointing to a Capacity Management Analytics report. Any modification of a shortcut automatically updates the source report. For more information about shortcuts, see the topic "Create a shortcut" in the *IBM Cognos Connection User Guide*.

3. Copy the report to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory.

   Saving to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory ensures that your custom reports are not overwritten when updates occur to Capacity Management Analytics.
4. Open the copied report in IBM Cognos Report Studio.
5. In Cognos Report Studio, edit the report as necessary. For detailed instructions on editing reports, refer to the *IBM Cognos Report Studio User Guide*.

   You can also insert the Base or Prompt objects that are provided in the IBM Capacity Management Analytics Solution Kit to your report. For more information on Base and Prompt objects, see "Customizing reports with Base and Prompt objects" on page 88.
6. In Cognos Report Studio, click **File**, **Save As** and save the report with a new name.

Building a custom report from a specific template

You can use a specific template to build a new report in one of several categories. For example, the CPU Template is a specific template that holds commonly used prompts for a CPU metric based report, and you can use this template to build your own CPU report.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Templates\Specific Templates` directory.
2. Copy a specific template to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory.

   Saving to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory ensures that your custom reports are not overwritten when updates occur to IBM Capacity Management Analytics.
3. Open the copied template in IBM Cognos Report Studio.
4. In Cognos Report Studio, edit the template as necessary. For detailed instructions on editing reports, refer to the *IBM Cognos Report Studio User Guide*.

   You can also insert the Base and Prompt objects that are provided in the Capacity Management Analytics Solution Kit to your report. For more information on Base and Prompt objects, see "Customizing reports with Base and Prompt objects" on page 88.
5. In Cognos Report Studio, click **File**, **Save As** and save the report with a new name.

Building a custom report from the base template

You can use the Base template as a starting point to create a custom IBM Capacity Management Analytics report.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Templates` directory.
2. Copy the Base template to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory.

   Saving to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory ensures that your custom reports are not overwritten when updates occur to Capacity Management Analytics.
3. Open the copied Base template in IBM Cognos Report Studio.
4. In Report Studio, edit the template as necessary. For detailed instructions on editing reports, refer to the *IBM Cognos Report Studio User Guide*.

You can also insert the Base and Prompt objects that are provided in the Capacity Management Analytics Solution Kit to your report. For more information on Base and Prompt objects, see "Customizing reports with Base and Prompt objects" on page 88.

5. In Report Studio, click **File**, **Save As** and save the report with a new name.

Building a custom report without a template

You can use IBM Cognos Report Studio to create a new custom IBM Capacity Management Analytics report that is not based on any other report or template.

Procedure

1. In IBM Cognos Connection, open the directory where you want to save the new report.
2. Click **Launch**, **Report Studio**.
3. If the **Select Package** dialog box appears, navigate to the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Models` directory in the **List of all packages**, then select a Capacity Management Analytics package for your report.
4. Use the Cognos Report Studio features to create a report that is useful for your analysis. For detailed instructions on using Cognos Report Studio, refer to the *IBM Cognos Report Studio User Guide*.

   You can also insert the Base/Prompt objects that are provided in the IBM Capacity Management Analytics Solution Kit to your report. For more information on Base/Prompt objects, see "Customizing reports with Base and Prompt objects" on page 88.
5. Click **File**, **Save**.

What to do next

Pre-built query data items DATE_BAND & TIME_BAND for the interval band

You run `hcm_udf_create_01.sql` to create user-defined functions such as YEAR_BAND and TIME_BAND. These user-defined functions can be used in Capacity Management Analytics reporting, for query data items [DATE_BAND] and [TIME_BAND].

In a CMA pre-built report, query data item [DATE_BAND], invoking user-defined functions YEAR_BAND, QUARTER_BAND, MONTH_BAND and DATE_BAND, creates interval bands for yearly, quarterly, monthly, and daily level. The query data item [TIME_BAND], invoking user-defined function TIME_BAND, creates interval bands for hourly and minute level.

These two data items, associated with Layout Component Reference `common_prompt_set_2` or `common_prompt_set_3`, provide the user flexibility to report on different user-selected interval bands. Each Capacity Management Analytics pre-built report or template includes data items [DATE_BAND] and [TIME_BAND] in its main query. When you build your own report and want to enable reporting on different interval bands, you need to have these two data items in the main query as well. You can copy them from any Capacity Management Analytics pre-built report or template to your report.

Each Capacity Management Analytics pre-built report or template also includes filters for [DATE] and [TIME], which can be used in conjunction with the [DATE_BAND] and [TIME_BAND] data items. You can copy them from any Capacity Management Analytics pre-built report or template to your report.

Customizing reports with Base and Prompt objects

IBM Cognos Capacity Management Analytics provides a number of Base and Prompt objects that you can use when creating a new report or modifying an existing report.

Inserting a Base object into a report

Base objects are common objects that can be used in any IBM Capacity Management Analytics report or template.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Templates\Objects` directory.
2. Run the Base report.
3. Identify the object that you want to insert into your report.

   You can review the object properties, which appear in yellow, to help identify the object that you want.
4. In IBM Cognos Report Studio, open the Capacity Management Analytics report into which you want to insert the Base object.
5. From the Cognos Report Studio **Toolbox** tab, drag the **Layout Component Reference** object to the location in your report at which you want the Base object to appear.
6. In the **Component Location** box, click **Another report**.
7. Click the **More** button.
8. Navigate to the `\Public Folders\IBM Capacity Management Analytics Solution Kit \Templates\Objects` directory and open the Base report.
9. In the **Available components to reference** box, select the name of the object you want to insert into your report, then click **OK**.

Results

A copy of the Base object appears in your Capacity Management Analytics report at the location where you placed the Layout Component Reference object.

Modifying a child object of a Base object

If you insert a Base object that contains children into your report, you can modify the child objects as required. For example, if you insert the app_header1 object into your report, you can modify the text item in the title bar to display your report name.

Procedure

1. In your report, select the copy of the parent object you created with the **Layout Component Reference** object.
2. In the **Properties** pane, double-click the **Overrides** property.
3. In the **Overrides** dialog box, select the child object you want to modify, then click **OK**.
   The child object is replaced by the text **Drop item to override component child**.
4. Drag any object from the **Toolbox** list to replace the child object.

Inserting a Prompt object into a report

A Prompt object asks a user to provide parameter values before a report is run. You can insert a Prompt object into any IBM Capacity Management Analytics report

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Templates\Objects` directory.
2. Open the Prompt report in IBM Cognos Report Studio.
3. Identify the Prompt object that you want to insert into your report.

You can review the object properties, which appear in yellow, to help identify the object that you want.

Note: The Prompt report is a multi-page report. Use Page Explorer to help you locate the Prompt object that you want to insert into your report.

4. In Cognos Report Studio, open the Capacity Management Analytics report into which you want to insert the Prompt object.
5. From the Cognos Report Studio **Toolbox** tab, drag the **Layout Component Reference** object to the location in your report at which you want the Prompt object to appear.
6. In the **Component Location** box, click **Another report**.
7. Click the **More** button.
8. Navigate to the `\Public Folders\IBM Capacity Management Analytics Solution Kit \Templates\Objects` directory and open the Prompt report.
9. In the **Available components to reference** box, select the name of the object you want to insert into your report, then click **OK**.
10. In Cognos Report Studio, open the Query Explorer and copy the query that the identified Prompt object uses.

    The yellow document for each Prompt object provides the related information for the object, such as Object Name, Query, and Parameter.
11. In the report with the Layout Component Reference object, open the Query Explorer and paste the query for the Prompt object.

Results

A copy of the Prompt object appears in your Capacity Management Analytics report at the location where you placed the Layout Component Reference object.

Customizing a Capacity Management Analytics workspace

IBM Capacity Management Analytics lets you create custom workspaces that satisfy your unique analytic requirement using standard IBM Cognos Business Intelligence authoring tools.

When you create custom workspaces, save your workspaces to a location outside of the IBM Capacity Management Analytics Solution kit folder. Saving to a location outside of the IBM Capacity Management Analytics Solution Kit folder ensures that your workspaces are not overwritten when updates occur to Capacity Management Analytics.

There are two options to build a custom workspace with IBM Capacity Management Analytics:

- “Building a custom workspace from an existing workspace” on page 89
- “Creating a new custom workspace ” on page 90

Building a custom workspace from an existing workspace

You can edit an existing IBM Capacity Management Analytics workspace to create a custom workspace. For example, you can add a chart to an existing workspace.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces` directory.
2. Copy the workspace you want to customize to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory. Saving to a location outside of the `\Public Folders\IBM Capacity Management Analytics Solution Kit` directory ensures that your custom workspaces are not overwritten when updates occur to Capacity Management Analytics.

3. Open the copied workspace in IBM Cognos Workspace.
4. In IBM Cognos Workspace, edit the workspace as necessary.

   For detailed instructions on editing workspaces, refer to the *IBM Cognos Workspace User Guide*, available on the IBM Cognos Business Intelligence welcome page of IBM Knowledge Center.

   You can also insert the workspace objects that are provided with the IBM Capacity Management Analytics Solution Kit into your workspace. For more information on workspace objects, see "Customizing a workspace with widgets" on page 90.
5. In Cognos Workspace, click **Save As** on the application bar and save the workspace with a new name.

Creating a new custom workspace

You can use IBM Cognos Workspace to create a new custom IBM Capaticy Management Analytics workspace without modifying an existing workspace.

Procedure

1. In IBM Cognos Connection, open the directory where you want to save the new workspace.
2. Click **Launch** > **Cognos Workspace**.
3. Click the **Create New** icon.
4. Use the Cognos Workspace features to create a workspace that is useful for your analysis.

   For detailed instructions on using workspaces, refer to the *IBM Cognos Workspace User Guide*, available on the IBM Cognos Business Intelligence welcome page of IBM Knowledge Center.

   You can also insert the workspace objects that are provided with the IBM Capacity Management Analytics Solution Kit into your workspace. For more information on workspace objects, see "Customizing a workspace with widgets" on page 90.
5. Click **Save** on the application bar.

Customizing a workspace with widgets

IBM Capacity Management Analytics provides a number of pre-built widgets, which you can use when creating a new workspace or modifying an existing workspace.

About this task

Review "IBM Capacity Management Analytics Workspace pre-built widgets" on page 91 for a list of available pre-built widgets.

Procedure

1. In the **Workspace Content** pane, navigate to `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces\Objects\Prompts & Titlebars` directory.
2. Open an existing **Prompts & Titlebars** report, and then drag the **Prompt** page to the canvas, or to the global area for a workspace with multiple tabs. For more information on Prompts & Titlebars objects, see "Creating your own Prompts & Titlebars report" on page 96.

   Note: You can use the keyboard shortcut **Ctrl+Shift+Enter** to add a selected widget, or right-click the widget in the **Content** pane and select **Insert in Current Tab**.
3. In the **Workspace Content** pane, navigate to `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces\Objects\Common` directory.
4. Open the **Titlebar Template** report.
5. Drag the Titlebar page from the **Content** pane and edit the titlebar as required for your purposes. For details on editing the titlebar text, see "Editing titlebars in Cognos Workspace Advanced" on page 91.

   Some titlebars also incorporate drill-through capabilities. For more information on editing such tilebars, see "Creating your own Prompts & Titlebars report" on page 96.
6. Perform a search for the widget in the **Workspace Content** pane.

For more information about workspace search, see "Performing a search" in the *IBM Cognos Workspace User Guide*.

7. Click the **Show Information Card** icon to enable displaying high-level information about workspace content such as type, owner and description.
8. Drag a report part (chart or table) from the **Result** pane to the canvas. For more information on widgets used in Canvas, see "Creating your own widgets report used in Canvas" on page 97.
9. To align widgets, drag one widget towards another widget until a dotted line appears, and then drop the widget.
10. To remove a widget from the workspace, click the **Widget Actions** icon, then click **Remove from Workspace**.
11. Click the **Save** button on the Application Bar to save the workspace.

Editing titlebars in Cognos Workspace Advanced

The Titlebar Template is a base template you can use in your workspace. You can edit this template to satisfy your particular requirements in Cognos Workspace Advanced.

Procedure

1. Open your workspace.
2. Click the **Actions** icon for the titlebar widget, then click **Do More**.
   The titlebar widget opens in Cognos Workspace Advanced.
3. Edit the titlebar according to your needs.
4. For detailed instructions on using workspace advanced, refer to the *IBM Cognos Workspace Advanced User Guide*, available on the IBM Cognos Business Intelligence welcome page of IBM Knowledge Center.
5. Click **Done** on the tool bar to save your changes and return to your workspace

IBM Capacity Management Analytics Workspace pre-built widgets

The following widgets are available for inclusion in a new or existing workspace.

Table 15: Capacity Management Analytics Workspace pre-built widgets

Object Type	Object Name	Location
Chart	As if all on zIIP Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/zAAP What-ifs – LPAR Level
Chart	As if all on zIIP Engines Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/zAAP What-ifs – LPAR Level
Chart	As if all on zIIP/zAAP Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/zAAP What-ifs – LPAR Level
Chart	As if all on zIIP/zAAP Engines Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/zAAP What-ifs – LPAR Level

Table 15: Capacity Management Analytics Workspace pre-built widgets (continued)		
Object Type	**Object Name**	**Location**
Chart	DASD I/O Response Time Components & I/O Rate Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > DASD I/O Performance
Chart	Latent Demand Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > Latent Demand
Chart	LCU I/O Queuing Delay Time - Top N at Machine Level Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > DASD I/O Performance
Chart	Memory Usage Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Memory > Memory Usage
Chart	MIPS Used at LPAR Level Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > MIPS Used – LPAR Level
Chart	Observed Engines Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/ zAAP What-ifs – LPAR Level
Chart	Observed MIPS Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/ zAAP What-ifs – LPAR Level
Chart	Percentage of Servers with High Usage Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Percentage of Servers with High/Low Usage
Chart	Percentage of Servers with Low Usage Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Percentage of Servers with High/Low Usage
Chart	Servers reporting by Operating System Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Servers Reporting by Operating System - Enterprise Level

Table 15: Capacity Management Analytics Workspace pre-built widgets (continued)		
Object Type	**Object Name**	**Location**
Chart	Top N DASD I/O Channel Path Utilizations Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > Channel Utilization
Chart	Top N Max TCP/IP Server Ports Connections Chart	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > TCP/IP Server Ports Connection Statistics
Page	FOR MACHINE LEVEL SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Page	FORE PRODUCT SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Page	OBS MACHINE LEVEL SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Page	OBS PRODUCT SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Page	OPT MACHINE LEVEL SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Page	OPT PRODUCT SUMMARY	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary

Object Type	Object Name	Location
Page	OPT SUGGESTED ACTION	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > Software Cost Analysis Summary
Table	AIX - 10 Highest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > AIX - 10 Highest/Lowest CPU Usage Peaks
Table	AIX - 10 Lowest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > AIX - 10 Highest/Lowest CPU Usage Peaks
Table	DASD I/O Channel Utilization - Top N at Machine Level Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > Channel Utilization
Table	DASD I/O Queuing Delay Intensity - Top N at Machine Level Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > DASD I/O Performance
Table	Linux for System X –10 Lowest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System X - 10 Highest/Lowest CPU Usage Peaks
Table	Linux for System X –10 Highest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System X - 10 Highest/Lowest CPU Usage Peaks
Table	Linux for System Z –10 Lowest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System Z - 10 Highest/Lowest CPU Usage Peaks
Table	Linux for System Z –10 Highest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System Z - 10 Highest/Lowest CPU Usage Peaks
Table	LPAR Configuration Metrics Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > MIPS Used – Machine Level

Table 15: Capacity Management Analytics Workspace pre-built widgets (continued)

Table 15: Capacity Management Analytics Workspace pre-built widgets (continued)		
Object Type	**Object Name**	**Location**
Table	MIPS Used at Machine Level Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > MIPS Used – Machine Level
Table	SCA Summary - NO89 Products Matrix Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > Software Cost Analysis > NO89 Products Matrix
Table	Windows - 10 Highest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Windows - 10 Highest/Lowest CPU Usage Peaks
Table	Windows - 10 Lowest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Windows - 10 Highest/Lowest CPU Usage Peaks
Table	zIIP zAAP What-ifs - LPAR Level Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > CPU > zIIP/zAAP What-ifs – LPAR Level
Table	z/OS - 10 Lowest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > z/OS - 10 Highest/Lowest CPU Usage Peaks
Table	z/OS - 10 Highest CPU Usage Peaks Table	Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > z/OS - 10 Highest/Lowest CPU Usage Peaks
Common Objects	Titlebar Template	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Common > Titlebar Template
Common Objects	Date & Time Prompt Group	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Common > Date & Time Prompt Group
Common Objects	License	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Common > License

Table 15: Capacity Management Analytics Workspace pre-built widgets (continued)

Object Type	**Object Name**	**Location**
Specific Prompt Group and Titlebar	Software Cost Analysis Prompt & Titlebar • Prompts • Titlebar	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Prompts & Titlebars > Software Cost Analysis Prompt & Titlebar
Specific Prompt Group and Titlebar	zIIP/zAAP What-ifs - LPAR Level Prompts & Titlebars • Prompts • Titlebar	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Prompts & Titlebars
Specific Prompt Group and Titlebar	Distributed Prompt & Titlebar • Prompts • Titlebar	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Prompts & Titlebars > Distributed Prompt & Titlebar
Specific Prompt Group and Titlebar	zServer Monitoring Dashboard Prompts & Titlebars • Prompts • Machine Level Titlebar • CPU & Storage Titlebar • I/O & Network Titlebar	Public Folders > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Prompts & Titlebars

Creating your own Prompts & Titlebars report

A Prompts & Titlebars report contains prompts and titlebars which are used in your own workspace. You can insert Prompts & Titlebars object into any IBM Capacity Management Analytics workspace.

Procedure

1. In IBM Cognos Connection, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces\Objects\Prompts & Titlebars` directory.
2. Create your own Prompts & Titlebars report and name the report as appropriate according to your workspace usage.

   You can create your own Prompts & Titlebars report from either one of the Prompts & Titlebars reports that Capacity Management Analytics provides. You can also create Prompts & Titlebars report using the common objects Capacity Management Analytics provides.

   Note: The common objects that you can use to customize your report are located in `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces\Objects\Common`.

   See Chapter 9, “Report, workspace, and folder customization,” on page 85 on creating reports.
3. Edit your own Prompts & Titlebar report as necessary.
4. In **Cognos Report Studio**, click **File** > **Save As** and save the report.

Creating your own widgets report used in Canvas

A widgets report used in Canvas contains charts and tables you want to show in the workspace. You can insert any chart or table widget into any IBM Capacity Management Analytics workspace.

Procedure

1. In **IBM Cognos Connection**, open the `\Public Folders\IBM Capacity Management Analytics Solution Kit\Workspaces\Objects\Charts & Tables` directory.
2. Create your own widgets report, and name the report as appropriate for your report usage. You can create your own widgets report from one of the widgets reports Capacity Management Analytics provides or from scratch.

   See Chapter 9, “Report, workspace, and folder customization,” on page 85 on creating reports.

   Note: Reports should be saved in different folders according to the usage of each report.
3. Edit your own widgets report as necessary.
4. In **Cognos Report Studio**, click **File** > **Save As** and save the report.

Customizing top N display for the zServer Monitoring Dashboard

Many tables and charts on the **zServer Monitoring Dashboard** are configured to display the top five or top ten items from a query. You can modify the widget associated with a given table or chart to return any top number (top N) of items.

About this task

Generally, to customize the top N value for a chart or table, you open the associated widget source report in IBM Cognos Report Studio, and in the widget source query, set the filter for [RANK] to the desired number. For example, changing [RANK] <= 10 to [RANK] <= 12 returns the top 12 items instead of the top 10 items.

The following table lists the widgets on the **zServer Monitoring Dashboard** for which top N can be modified.

Widgets on the **zServer Monitoring Dashboard** for which top N can be modified

Table 16:

Widget Name	Widget Source Report	Widget Source Query
Top 5 Max TCP/IP Server Ports Connections Chart	Public > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > TCP/IP Server Ports Connection Statistics	qry_main_chart
DASD I/O Channel Utilization - Top N at Machine Level Table	Public > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > Channel Utilization	qry_channel_machine
LCU I/O Queuing Delay Time - Top N 10 at Machine Level Chart	Public > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > DASD I/O Performance	qry_main_chart_machine_level

Table 16: (continued)

Widget Name	**Widget Source Report**	**Widget Source Query**
DASD I/O Queuing Delay Intensity - Top N at Machine Level Table	Public > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > DASD I/O Performance	qry_dasd_qi_by_machine
Top 5 DASD I/O Channel Path Utilization Chart	Public > IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > Channel Utilization	qry_channel

Procedure

1. Open the **zServer Monitoring Dashboard** workspace.
2. Right-click the widget for which you want to change the top N, then click **Widget Actions** > **Properties**.
3. Click the **Report** tab on the **Properties** dialog box.
4. Note the location of the source report for the widget as displayed on the dialog box. You will have to go to this location to edit the source report.

   The report location will look similar to this:`Public Folders >IBM Capacity Management Analytics Solution Kit > Workspaces > Objects > Charts & Tables > I/O > Channel Utilization`
5. In **IBM Cognos Connection**, navigate to the location of the source report for the widget you want to modify.
6. Click the **Open with Report Studio** icon for the report you want to edit.
7. Click the **Select Object** icon to select the table or chart widget you want to modify.
8. Right-click the table or chart widget, then click **Go to Query**.
9. In the **Detail Filters** pane of **Report Studio**, double-click the **[Rank]** filter.
10. In the **Detail Filter Expression** dialog box, modify the **Expression Definition** to use the top N filter you want to see in your chart or table. For example, set the Expression Definition to `[RANK] <= 7` if you want to see the top seven items in a chart or table.
11. Click **OK** to save the Expression Definition.
12. For the **LCU I/O Queuing Delay Time - Top 10 at Machine Level** chart, **Top 5 DASD I/O Channel Path Utilizations** chart, and **Top 5 Max TCP/IP Server Ports Connections** chart, edit the chart title to reflect the new top N value you set in Step 10.

    All other chart and table titles use variables to automatically display the top N value you set.
13. Click **Save** on the **Report Studio** toolbar to save the report.
14. Reopen the **zServer Monitoring Dashboard** workspace.
15. Right-click the chart or table for which you updated the top N value, then click **Widget Actions** > **Reset**.
16. Click **Yes** to confirm that you want to reset the widget.
17. Click the **Refresh View** icon in the Global Area of the **zServer Monitoring Dashboard** workspace to view the updated chart or table.
18. Click **Save** on the **zServer Monitoring Dashboard** toolbar to save the workspace.

Customizing top N display for the Enterprise Dashboard

Many tables and charts on the **Enterprise Dashboard** are configured to display the top 10 items from a query. You can modify the widget that is associated with a given table or chart to return any top number (top N) of items.

Generally, to customize the top N value for a chart or table, you open the associated widget source report in IBM Cognos Report Studio, and in the widget source query, set the filter for [RANK] to the wanted number. For example, changing [RANK] <= 10 to [RANK] <= 12 returns the top 12 items instead of the top 10 items.

The following table lists the widgets on the **Enterprise Dashboard** for which top N can be modified.

*Table 17: Widgets on the **Enterprise Dashboard** for which top N can be modified*

Widget Name	Widget Source Report	Widget Source Query
AIX® - 10 Highest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > AIX - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_top_list_aix`
AIX - 10 Lowest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > AIX - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_bottom_list_aix`
Linux for System X - 10 Highest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System X - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_top_list_xLinux`
Linux for System X –10 Lowest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System X - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_bottom_list_xLinux`
Linux for System Z –10 Highest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System Z - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_top_list_zLinux`
Linux for System Z –10 Lowest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Linux for System Z - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_bottom_list_zLinux`
Windows - 10 Highest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Windows - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_top_list_win`
Windows - 10 Lowest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Windows - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_bottom_list_win`
z/OS - 10 Highest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > z/OS - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_top_list_zos`
z/OS - 10 Lowest CPU Usage Peaks Table	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > z/OS - 10 Highest/Lowest CPU Usage Peaks	`qry_main_highest _usage_bottom_list_zos`

Procedure

1. Open the **Enterprise Dashboard** workspace.
2. Right-click the widget for which you want to change the top N, then click **Widget Actions** > **Properties**.
3. Click the **Report** tab on the **Properties** dialog box.
4. Note the location of the source report for the widget as displayed on the dialog box. You must go to this location to edit the source report.

   The report location looks similar to this: `Public Folders > IBM Capacity Management Analytics Solution Kit > Reports > CPU > z/OS - 10 Highest/Lowest CPU Usage Peaks`.
5. In **IBM Cognos Connection**, navigate to the location of the source report for the widget you want to modify.
6. Click the **Open with Report Studio** icon  for the report you want to edit.
7. Click the **Select Object** icon  to select the table or chart widget you want to modify.
8. Right-click the table or chart widget, then click **Go to Query**.
9. In the **Detail Filters** pane of **Report Studio**, double-click the **[Rank]** filter.
10. In the **Detail Filter Expression** dialog box, modify the **Expression Definition** to use the top N filter you want to see in your chart or table. For example, set the Expression Definition to `[RANK] <= 7` if you want to see the top seven items in a chart or table.
11. Click **OK** to save the Expression Definition.
12. For the **AIX - 10 Highest/Lowest CPU Usage Peaks**, **Linux for System X - 10 Highest/Lowest CPU Usage Peaks**, **Linux for System Z - 10 Highest/Lowest CPU Usage Peaks**, **Windows - 10 Highest/ Lowest CPU Usage Peaks**, and **z/OS - 10 Highest/Lowest CPU Usage Peaks** tables, edit the table title to reflect the new top N value you set in Step 10.

    All other table titles use variables to automatically display the top N value you set.
13. Click **Save** on the **Report Studio** toolbar to save the report.
14. Reopen the **Enterprise Dashboard** workspace.
15. Right-click the table for which you updated the top N value, then click **Widget Actions** > **Reset**.
16. Click **Yes** to confirm that you want to reset the widget.
17. Click the **Refresh View** icon in the Global Area of the **zServer Monitoring Dashboard** workspace to view the updated chart or table.
18. Click **Save** on the **Enterprise Dashboard** toolbar to save the workspace.

Customizing high usage and low usage thresholds for the Enterprise Dashboard

The Percentage of Servers with High/Low Usage charts on the **Enterprise Dashboard** are configured with the default high usage threshold as 90% busy and the default low usage threshold as 10% busy. You can modify the widget that is associated with a given chart to use any high or low usage threshold.

About this task

Generally, to customize the usage threshold value for the Percentage of Servers with High Usage chart or the Percentage of Servers with Low Usage chart, you open the associated widget source report in IBM Cognos Report Studio, and in the widget source query, set the filter for [CPU_BUSY_PCNT] to the wanted threshold value. For example, changing [CPU_BUSY_PCNT] > 90 to [CPU_BUSY_PCNT] > 80 returns the percentage of the total number of servers having a peak utilization greater than 80%.

The following table lists the widgets on the **Enterprise Dashboard** for which usage threshold can be modified.

Widgets on the **Enterprise Dashboard** for which usage threshold can be modified

Table 18:

Widget Name	**Widget Source Report**	**Widget Source Queries**
Percentage of Servers with High Usage Chart	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Percentage of Servers with High/Low Usage	qry_main_ninety_percent_zLinux qry_main_ninety_percent_xLinux qry_main_ninety_percent_win qry_main_ninety_percent_aix qry_main_ninety_percent_zos
Percentage of Servers with Low Usage Chart	Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Percentage of Servers with High/Low Usage	qry_main_ten_percent_zLinux qry_main_ten_percent_xLinux qry_main_ten_percent_win qry_main_ten_percent_aix qry_main_ten_percent_zos

Procedure

1. Open the **Enterprise Dashboard** workspace.
2. Right-click the widget for which you want to change the usage threshold, then click **Widget Actions** > **Properties**.
3. Click the **Report** tab on the **Properties** dialog box.
4. Note the location of the source report for the widget as displayed on the dialog box. You must go to this location to edit the source report.

   The report location looks similar to this: `Public > IBM Capacity Management Analytics Solution Kit > Reports > CPU > Percentage of Servers with High/Low Usage`
5. In **IBM Cognos Connection**, navigate to the location of the source report for the widget you want to modify.
6. Click the **Open with Report Studio** icon for the report you want to edit.
7. Click the **Query Explorer** icon, then click the **Select Queries** icon Queries to select the query you want to modify.
8. In the **Detail Filters** pane of **Report Studio**, double-click the usage threshold filter.
9. In the **Detail Filter Expression** dialog box, modify the **Expression Definition** to use the usage threshold filter you want to see in your chart or table. For example, set the Expression Definition to `[CPU_BUSY_PCNT] > 80` if you want to see the percentage of the total number of servers having a peak utilization greater than 80% in a chart or table.
10. Click **OK** to save the Expression Definition.
11. Click **Save** on the **Report Studio** toolbar to save the report.

Chapter 10. Software cost analysis

Software cost analysis allows you to better manage z/OS software costs and identify where and when workloads need to be adjusted and when more capacity is required.

Software cost analysis in IBM Capacity Management Analytics includes one workspace and four reports.

- The "Software Cost Analysis Summary workspace" on page 277 shows the machine level summary and registered IBM product MSU and monetary value summary.
- The "SCA: LPAR MSU Utilization " on page 249 shows the MSU utilization metrics from an LPAR perspective.
- The "SCA: Product MSU and Price " on page 257 shows an IBM product's billable MSUs and monetary cost on a CPC from product and CPC perspective.
- The "SCA: NO89 Product MSU and Price " on page 256 is similar to the Product MSU and Price report, but shows the IBM NO89, unregistered products.
- The "SCA: NO89 Products Matrix " on page 255 shows the running matrix of LPARs and IBM NO89 products.

Important: The Sub-Capacity Reporting Tool (SCRT) is the official pricing tool for IBM. The Software Cost Analysis capability of Capacity Management Analytics is for your reference and analysis. If there is any discrepancy between SCRT and Capacity Management Analytics, SCRT values are used for billing purposes. The monetary value shown in Capacity Management Analytics reports is based on the highest MSU value calculated by Capacity Management Analytics and the price value customer supply in the HCM.CUSTOMER_PRICE table, which may differ from actual IBM charges.

The Sub-Capacity Reporting Tool (SCRT) is the official pricing tool for IBM. Capacity Management Analytics is to be used only for reference and planning. Any discrepancies in Capacity Management Analytics will not be considered valid against SCRT. The monetary values shown in Capacity Management Analytics reports are based on the highest calculated MSU value, and the monetary values given by the customer may be different with the final bill that IBM provides.

Capacity Management Analytics does not enforce any of the Terms & Conditions set forth by the IBM products. Anything not listed is therefore not supported by Capacity Management Analytics. Among the unsupported items in Capacity Management Analytics are:

- Capacity On Demand
- Technology Update Pricing
- Mobile Pricing
- Entitlement IDs / Features
- Single Version Charging
- SYSPLEX / Parallel SYSPLEX Pricing

Supported scenarios

Capacity Management Analytics supports three scenarios for Software Cost Analysis: observed, forecasted, and optimized.

Observed scenarios

The observed scenario shows what has been observed in the past. The data is from Tivoli Decision Support for z/OS with SMF70 and SMF89 records.

Note: You must ensure that data is aligned between SMF70 and SMF89. If they are not both collected at the same time interval, report values might be inaccurate.

Forecasted scenarios

The forecasted scenario shows what is predicted to happen in the future. The data is from tables that store the result from forecast streams in the HCM schema.

Although the forecast stream is for predicting the future, the stream also provides referenced forecasted value for previous periods, so you can sometimes see both observed and forecasted data in reports.

For comparison purposes, some reports and workspaces show observed values as well as forecasted values. In those places, observed values are retrieved directly from observed data, and not from forecasted result tables, so you can see the latest actual observed value.

For more information about forecast streams, such as the default date ranges that the stream is supporting, see "Software Cost Analysis forecasting streams" on page 128.

Optimized scenarios

The optimized scenario shows suggested actions based on forecasted MSU values. The data is from tables that store the result from optimization streams in the HCM schema.

As the suggested actions are generated based on forecasted data, you should run optimization streams immediately after forecast stream. This ensures that you can see aligned forecasted and optimized data in reports. If there are new forecasted results, but the optimization stream has not been run, then the values that are shown in forecasted and optimized scenarios might not be aligned. When this occurs, reports provide a warning message to remind you to run the optimization stream: **`CMA-RW002 Scenario Synchronization Warning Message`** `Please note current optimization result is based on the forecast result running at <`*`timestamp`*`>. There are new updates, please run optimization again.`

For more information about optimization streams, such as the default date ranges that the stream is supporting, see "Billable MSU optimization stream" on page 137.

Supported pricing structures

For registered IBM products, Software Cost Analysis supports the following pricing structures.

Sub-capacity pricing of Monthly License Charge (MLC) pricing

Software Cost Analysis supports these pricing structures for Monthly License Charges.

- Workload License Charges (WLC) - Variable WLC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/wlc.html.
- Advanced Workload License Charges (AWLC), including Integrated Workload Pricing (IWP), when appropriate - AWLC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/awlc.html.
- Entry Workload License Charges (EWLC) - EWLC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/ewlc.html.
- Advanced Entry Workload License Charges (AEWLC), including Integrated Workload Pricing (IWP), when appropriate - AEWLC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/aewlc.html.
- Midrange Workload License Charges (MWLC) - MWLC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/mwlc.html.
- System z New Application License Charges (zNALC) - zNALC Structure is considered in monetary value calculations. For details, see http://www.ibm.com/systems/z/resources/swprice/mlc/znalc.html.

System z International Product License Agreement (IPLA) Software

Software Cost Analysis supports these pricing structures for System z International Product License Agreement (IPLA) Software.

- Execution-based licensing terms
- Getting Started Sub-capacity Pricing (GSSP)
- Reference-based licensing terms
- z/OS based licensing terms

Value Unit Pricing is applied. The following Value Unit Exhibits (VUE) which convert MSU to VU are supported:

- VUE001
- VUE007
- VUE020

For details on Value Unit Exhibits, see http://www.ibm.com/systems/z/resources/swprice/zipla/vue.html.

IBM NO89 products

For IBM NO89 products, Software Cost Analysis supports the same pricing structures as registered products.

z/OS V1 (5694-A01) and z/OS V2 (5655-ZOS)

z/OS products can use either traditional pricing structures or System z New Application License Charges (zNALC). Capacity Management Analytics calculates MSU and monetary value according to the pricing structure in effect: zNALC or traditional pricing.

WebSphere Transformation Extender V8 (5655-R95)

Software Cost Analysis in Capacity Management Analytics considers product ID as the unique identifier of a product and uses the same pricing structure, except for WebSphere® Transformation Extender V8 (5655-R95).

WebSphere Transformation Extender V8 (5655-R95) differs from all other Getting Started Sub-Capacity Pricing products because only parts of this product (just two of the four priced features) are eligible for Getting Started Sub-Capacity Pricing . Capacity Management Analytics considers WebSphere Transformation Extender V8 as one product, and use the GSSP values in the HCM.SUBCAP_PROGRAMS table to identify the pricing structure that is applied in Software Cost Analysis.

MSU calculations

In Capacity Management Analytics reports and workspaces, MSU is calculated according to the date and time range you select

When observing MSU calculations you should be aware that:

- The date and time range used in calculations is based on the time zone you select in the report. If the LPARs in a CPC are in different time zones, or the CPCs in an enterprise are in different time zones, all of them are converted to the time zone you specify in the report. Unlike SCRT, which calculates billable MSU within one month only, Software Cost Analysis calculates all values according to the date and time range you select in prompts. In the four-hour rolling average calculation of the first three hours of the date and time range that you select in report or workspace, the values of the previous hours are included in the calculation (if there is any data).

The following terms are important to understand and are referenced in the descriptions of the MSU calculations performed by Capacity Management Analytics.

LPAR 4HRA MSU
The four-hour rolling average of LPAR Hourly MSU Utilization of a LPAR at each hour. In the implementation of Observed scenario, this value is not calculated by Capacity Management Analytics, but by DRL.MVSPM_LPAR_MSU_T.MSU_4HRA, which is from SMF70LAC directly.

LPAR BILLABLE MSU
The minimum value of LPAR 4HRA MSU and defined capacity.

Defined capacity is from DRL. MVSPM_LPAR_MSU_T.CAPACITY_LIMIT_MSU. If defined capacity is not set or the value is 0, then use the value from DRL.MVSPM_LPAR_MSU_T.PROC_CAPACITY_MSU to compare with LPAR 4HRA MSU.

If the minimum value is not an integer, rounding is performed:

- If value is < 1, then use 1. For example, 0.2 is rounded up to 1.
- If value >= 1, then round the value down to the nearest integer. For example, 2.6 is rounded down to 2.
- If the value is very near to an integer (less than 10^-4), then use the integer value. For example, 3.9999 is rounded up to 4.

GROUP CAPACITY LIMIT
The capacity limit of the group; an LPAR can be set to a group. The value is from DRL.MVSPM_LPAR_MSU_T.CAPACITY_GRP_LMT.

PRODUCT HOURLY MSU UTILIZATION
The MSU that the product utilized at each hour.

Each hour's value is the sum of each SMF collection interval's value within an hour.

PRODUCT 4HRA MSU
The four hour rolling average of PRODUCT HOURLY MSU UTILIZATION of the product on the LPAR at each hour.

MLC & IPLA Execution-based calculation

MLC & IPLA Execution-based is calculated as follows:

1. Calculate TOTAL BILLABLE MSU PER CPC, which is the sum of LPAR BILLABLE MSU of each LPAR the product is running on at each hour.

   GROUP CAPACITY LIMIT is considered. If one or more LPARs on which the product runs belong to a group, then the sum of the LPAR BILLABLE MSU in the group must be equal to or less than the GROUP CAPACITY LIMIT.
2. Find the peak TOTAL BILLABLE MSU PER CPC during the specified date and time range.

MLC-IWP Adjusted calculation

MLC-IWP Adjusted is calculated as follows:

1. For each LPAR on which the product runs, calculate the TOTAL MSU UTILIZATION OF IWP DEFINING PRODUCTS. This value is the sum of PRODUCT 4HRA MSU of IWP defining products which are identified in the HCM.SUBCAP_PROGRAMS table and are running on the same LPAR.
2. For each LPAR on which the product runs, calculate the LPAR ADJUSTED BILLABLE MSU. This value equals the rounded value of (LPAR BILLABLE MSU) – (TOTAL MSU UTILIZATION OF IWP DEFINING PRODUCTS). In this formula, LPAR BILLABLE MSU and TOTAL MSU UTILIZATION OF IWP DEFINING PRODUCTS are not rounded prior to calculation.
3. Calculate TOTAL ADJUSTED BILLABLE MSU PER CPC, which is the sum of LPAR ADJUSTED BILLABLE MSU of each LPAR on which the product is running at each hour.

   GROUP CAPACITY LIMIT is considered. If one or more LPARs on which the product runs belong to a group, then the sum of the LPAR ADJUSTED BILLABLE MSU in the group must be equal to or less than the GROUP CAPACITY LIMIT.
4. Find the peak TOTAL BILLABLE MSU PER CPC during the specified date and time range.

IPLA Execution-based GSSP calculation

IPLA Execution-based GSSP is calculated as follows:

1. For each LPAR on which the product runs, get the PRODUCT 4HRA MSU of this product, and PRODUCT INT 4HRA MSU.

   PRODUCT INT 4HRA MSU is the four hour rolling average utilization of the DB2 time, which is used by this product via the DB2 JDBC type 2 connector. The one hour utilization value is calculated by the same formula of PRODUCT HOURLY MSU UTILIZATION, but the parameters are from the DRL.MVSPM_PROD_INT_T table instead of the DRL.MVSPM_PROD_T table.
2. For each LPAR on which the product runs, calculate LPAR GSSP MSU1, which is equal to (PRODUCT 4HRA MSU) – (PRODUCT INT 4HRA MSU).
3. For each LPAR on which the product runs, calculate LPAR GSSP BILLABLE MSU, which is the smaller of LPAR BILLABLE MSU and (LPAR GSSP MSU1) * (GSSP SCALING FACTOR). GSSP SCALING FACTOR is from the HCM.SUBCAP_PROGRAMS table.
4. Calculate TOTAL GSSP BILLABLE MSU PER CPC, which is the sum of LPAR GSSP BILLABLE MSU of each LPAR on which the product is running at each hour.

   GROUP CAPACITY LIMIT is considered. If one or more LPARs on which the product runs belong to a group, then the sum of the LPAR GSSP BILLABLE MSU in the group must be equal to or less than the GROUP CAPACITY LIMIT.
5. Find the peak TOTAL GSSP BILLABLE MSU PER CPC during the specified date and time range.

IPLA z/OS-based calculation

IPLA z/OS-based is calculated as follows:

1. Calculate TOTAL ZOS BILLABLE MSU PER CPC, which is the sum of LPAR BILLABLE MSU of all LPARs that runs z/OS at each hour. Because only z/OS system can collect SMF70 and SMF89 records, all the LPARs for which data is collected are included.

   GROUP CAPACITY LIMIT is considered. If one or more LPARs belong to a group, then sum of the LPAR BILLABLE MSU in the group must be equal to or less than the GROUP CAPACITY LIMIT.
2. Find the peak TOTAL ZOS BILLABLE MSU PER CPC as ZOS MSU PER CPC during the specified date and time range.

IPLA Reference-based calculation

IPLA Reference-based is calculated as follows:

1. Find the parent products of this product. This product's parent names are from HCM.SUBCAP_PROGRAMS.PARENT_NAME. The parent products can be identified by mapping PROD_ID to PARENT_NAME in the HCM.PARENT_PROGRAMS table.
2. Calculate the HIGHEST MSU PER CPC of all the parents.
3. Calculate TOTAL PARENT MSU, which is the sum of HIGHEST MSU PER CPC of all this product’s parents.
4. Calculate the HIGHEST MSU PER CPC of this child product by getting the minimum value of TOTAL PARENT MSU PER CPC and ZOS MSU PER CPC.

Product display names

To help you select products easily in report prompts, Capacity Management Analytics shows products using the product display name.

A product display name uses this syntax:
`[PRODUCT_ID]_[PRODUCT_NAME]_[PRODUCT_PRICING_STRUCTURE]`.

The components of a product display name are described in the following table.

Table 19: Product display name components	
Display name component	**Description**
PRODUCT_ID	• Is unique per product. • For registered products, the ID is retrieved from DRL. MVSPM_PROD_T.PROD_ID. • For NO89 or unregistered products, the ID is retrieved from HCM.NO89_PRODUCTS.PRODUCT_ID.
PRODUCT_NAME	If DRL.MVSPM_PROD_T.PROD_OWNER contains "IBM", then get product name from HCM.SUBCAP_PROGRAMS. If this product ID is not found or the value is null in HCM.SUBCAP_PROGRAMS table, then use "Not Defined". If DRL.MVSPM_PROD_T.PROD_OWNER does not contain "IBM", then get product name from DRL.MVSPM_PROD_T.PROD_NAME.
PRODUCT_PRICING_STRUCTURE	For registered products, if DRL.MVSPM_PROD_T.PROD_OWNER contains "IBM", then check the product ID in HCM.SUBCAP_PROGRAMS. For IBM NO89 or unregistered products, check the product ID in HCM.SUBCAP_PROGRAMS. If the product ID exists in HCM.SUBCAP_PROGRAMS, then the possible values are: • MLC • MLC-IWP Adjusted • Execution-based • Execution-based GSSP • z/OS-based • Reference-based If the product ID does not exist in the HCM.SUBCAP_PROGRAMS table, then use Execution-based. If DRL.MVSPM_PROD_T.PROD_OWNER does not contain "IBM", then use the DRL.MVSPM_PROD_T.PROD_OWNER value. If the value is null, then set to "Non-IBM".

Examples of product display names include:

- 5605-DB2_DB2 10 for z/OS_MLC
- 5635-A02_IMS V11_MLC-IWP Adjusted
- 5655-W65_WebSphere Application Server for z/OS V8_Execution-based GSSP
- 5655-DGS_InfoSphere Optim™ Data Growth Solution for z/OS with Classic Fed V7_z/OS-based
- 5655-H91_CICS VSAM Recovery V3_Reference-based
- 5655-XXX_Not Defined_Execution-based

- XX_XXXX_XX CORP

z/OS product identification

z/OS products are identified by the product ID in SMF89 which has the "z/OS" feature.

Only 5650-ZOS and 5694-A01 are z/OS products. If you register another product ID and name it as "z/OS", Capacity Management Analytics does not consider it as z/OS and it is treated as a normal product in either the MLC or IPLA categories. Accordingly, only 5650-ZOS and 5694-A01 will have zNALC consideration.

Historical data support for pricing information and NO89 Product Matrix information

Historical information is **not** considered for pricing-related calculations or in the NO89 Product Matrix report.

Chapter 11. LPAR weight optimization

IBM Capacity Management Analytics uses IBM ILOG CPLEX Optimization Studio to provide optimal LPAR weight values for your environment.

The objective of the optimization is to minimize the maximum or sum of the over weight amounts, in terms of % or number of engines, for all LPARs on a CPC while also considering the constraints set out by the relationships among the workload and weight for each LPAR and the total weights for all of the LPARs on the CPC.

Weight values for LPARs on a CPC are important when the CPC gets very busy; for example, when the LPARs are over 80% busy overall. The optimization model can help you determine more effective weights for the LPARs.

The LPAR weight optimization model uses the data from the SMF records for all the LPARs of the same processor type in the shared pool. The model expects a consistent list of LPARs during all intervals of the day for each date in the date list. The objective is to be able to use days that have the higher peak usage and represent the amount of workload on each LPAR that is expected to be seen repeatedly on those peak days. With 2 to 5 dates of similar peak usage of the same LPARs, the optimizer can better create a best fit set of weights that support the peak usage of all of the LPARs on the CPC.

It is important to note that it is not the consistency of the total workload of the LPAR that is important. It is the target workload amount that must fit within the guaranteed weight amount. As you select the dates, ensure that you look at the usage of each LPAR in the MIPS Used - Service Class Period Level report. Each service class is stacked on the chart by its importance level. You can also see the current guaranteed weight. By understanding what business work runs in each service class and what importance level it is running at, you can determine what portion of the total workload usage on the LPAR that you want to ensure has enough guaranteed capacity to complete when it wants to (this is also known as the demand workload). The other amount of workload above the target amount is considered deferrable.

Note: To use the LPAR weight optimization program, you must use Java version 7. Java version 7 is provided with SPSS Modeler 17. If you are using SPSS Modeler 16, you must install Java version 7.

For more information about the LPAR weight optimization report, see "CPU: LPAR Weight Optimization Run Result" on page 218.

Optimization information is also available in the "CPU: MIPS Used - LPAR Level by WLM Importance" on page 229 and "CPU: Over/Under Share Weight - CPC by LPAR" on page 234 reports.

Creating optimization programs

IBM Capacity Management Analytics provides an optimization program for LPAR weights. You can create your own optimization programs by using IBM ILOG CPLEX Optimization Studio. For more information about using Optimization Studio, see the IBM ILOG CPLEX Optimization Studio Getting Started with CPLEX guide (www.ibm.com/support/knowledgecenter/SSSA5P_12.6.1/ilog.odms.studio.help/pdf/gscplex.pdf).

Important: When you create an optimization program in the version of ILOG CPLEX Optimization Studio that is provided with Capacity Management Analytics, you can use up to 1000 variables and 1000 constraints. You must ensure that you are within this limitation before you run the program to solve the problem. You can include a check and an audit in the program, and if the number of variables or constraints is over 1000, then the programs must stop.

If you must use more than 1000 variables or constraints, you can purchase the full version of IBM ILOG CPLEX Optimization Studio.

You can use the following formula to calculate the number of constraints and variables:

```
# of constraints = IloCplex.getNcols + IloCplex.getNrows + IloCplex.getNQCs +
IloCplex.getNSOSs
```

Where:

- `IloCplex.getNrows` queries the number of rows of A (that is, the number of linear constraints)
- `IloCplex.getNQCs` queries the number of quadratic constraints
- `IloCplex.getNcols` queries the number of special ordered sets (SOSs)

`# of variables = IloCplex.getNcols`

Where:

- `IloCplex.getNcols` queries the number of columns of A (that is, the number of variables)

For more information about these classes, see the ILOG CPLEX documentation (www.ibm.com/support/knowledgecenter/SSSA5P_12.5.1/ilog.odms.cplex.help/refjavacplex/html/ilog/cplex/IloCplex.html).

More documentation about classes is provided with the product installation. This documentation is available in *CMA_ROOT*`/documents/optimizations/index.html`.

Choosing dates for the optimization

Selecting the dates to use as input involves looking at the target amount of workload for each LPAR on the CPC and ensuring that the dates that you select best represent the peak and the competition among LPARs that you want optimized.

For example, Monday is the peak day of the week for having multiple applications that are running on the CPC. And a holiday falls on a Monday so the applications are not nearly as busy as they would be on a regular workday. Therefore, choosing four recent Mondays but avoiding any Monday that was a holiday could be a good set of dates to use for the optimization. Or, say that your application peak is even higher during the first of each quarter and Monday is the busiest day of the week for your applications. If your last quarter did not perform as desired during the peak Mondays and you expect the next quarter to be similar and want to have the optimizer provide a different set of weight values for the competing LPARs. You could choose one or 2 of the Mondays that occurred during that last quarter.

Also, keep in mind that any day that had an abnormally high amount of target workload, which is not expected to be normal, should not be included in the date list. The optimizer attempts to support that full amount of target workload under the weight value and this impacts the other LPARs ability to fit their target workload under the weight value.

The reverse is also true. If you have weight values set such that an abnormally large amount of total workload was able to run on one LPAR and thus forcing another LPAR's target workload to not be able to fit within its weight. This means that the target workload amount that was accomplished is lower than what is expected to occur. You don’t want to include that date in the date list as it is not representative of the amount of work you want to optimize for.

Objectives and constraints

The LPAR weight optimization model produces optimized weights based on your objectives and within the bounds of your constraints.

Objectives

The objective of the optimization is to find the best weight value for each LPAR within a CPC. For example:

```
Minimize (max_over_target_weight * max_over_target_lpar_weight +
      max_over_total_weight * max_over_total_lpar_weight +
      sum_over_target_weight * sum_over_target_lpar_weight +
      sum_over_total_weight * sum_over_total_lpar_weight)
```

Where:

max_over_target_weight
: The weight or punishment for the sub-objective max_over_target_lpar_weight. The higher the value, the more important the sub-objective is.

max_over_target_lpar_weight
The total maximum amount in % of the total physical shared CPUs that the target workload is over the guaranteed capacity that is set by LPAR weight on all LPARs.

max_over_total_weight
The weight or punishment for the sub-objective max_over_total_lpar_weight. The higher the value, the more important the sub-objective is.

max_over_total_lpar_weight
The total maximum amount in % of the total physical shared CPUs that the total workload is over the guaranteed capacity that is set by LPAR weight on all LPARs.

sum_over_target_weight
The weight or punishment for the sub-objective sum_over_target_lpar_weight. The higher the value is, the more important the sub-objective is.

sum_over_target_lpar_weight
The total amount in % of the total physical shared CPUs that the target workload is over the guaranteed capacity that is set by LPAR weight on all LPARs at all hours and on all days.

sum_over_total_weight
The weight or punishment for the sub-objective sum_over_total_lpar_weight. The higher the value, the more important the sub-objective is.

sum_over_total_lpar_weight
The total amount in % of the total physical shared CPUs that the total workload is over the guaranteed capacity that is set by LPAR weight on all LPARs at all hours and on all days.

Constraints

The basic constraints that are used in the optimization are the relationships between the workload and weight, and the 100% of total weights.

Objective related constraints

ct_sum_target_wl_over_weight presents the relationship between the decision expression sum_over_target_lpar_weight (used in objective) and the decision variable over_target_lpar_weight. sum_over_target_lpar_weight is the sum of over_target_lpar_weight for all LPARs at all hours and on all days.

ct_sum_total_wl_over_weight presents the relationship between the decision expression sum_over_total_lpar_weight (used in objective) and the decision variable over_total_lpar_weight. sum_over_total_lpar_weight is the sum of over_total_lpar_weight for all LPARs at all hours and on all days.

ct_max_target_wl_over_weight presents the relationship among the decision expression max_over_target_lpar_weight (used in objective), the decision variable max_over_target_lpar_weight, and the input data lpar_priority. max_over_target_lpar_weight is the sum of max_over_target_lpar_weight times lpar_priority for all LPARs.

ct_max_total_wl_over_weight presents the relationship among the decision expression max_over_total_lpar_weight (used in objective), the decision variable max_over_total_lpar_weight, and the input data lpar_priority. max_over_total_lpar_weight is the sum of max_over_total_lpar_weight times lpar_priority for all LPARs.

Relationship constraints

ct_target_wl_over_under_weight presents the relationship between the target workload and the LPAR weight on each LPAR at each hour of each day. When the target workload is higher than the LPAR weight, then the decision variable over_target_lpar_weight gets the over value. When the workload is less than the weight, then the decision variable under_target_lpar_weight gets the under value. The over values are used in objective.

ct_total_wl_over_under_weight presents the relationship between the total workload and the LPAR weight on each LPAR at each hour of each day. When the total workload is higher than the LPAR weight, then the decision variable over_total_lpar_weight gets the over value. When the workload is less than the weight, then the decision variable under_total_lpar_weight variable gets the under value. The over values are used in objective.

ct_max_target_wl_over_weight_per_lpar presents the relationship between the decision variable max_over_target_lpar_weight and the decision variable over_target_lpar_weight. max_over_target_lpar_weight is the maximum value of over_target_lpar_weight for each LPAR across all hours and days.

ct_max_total_wl_over_weight_per_lpar presents the relationship between the decision variable max_over_total_lpar_weight and the decision variable over_total_lpar_weight. max_over_total_lpar_weight is the maximum value of over_total_lpar_weight for each LPAR across all hours and days.

Total weight constraints

ct_sum_weight presents the relationship among the weight of all LPARs. The total weight of all LPARs should be 100%.

Logical engine number constraint

ct_weight_lessthan_logic_proc presents the relationship between the guaranteed capacity set by LPAR weight and logical processor capacity. The LPAR weight in engine numbers should not be more than the logical process numbers set for the LPAR.

Fixed weight constraint

ct_fixed_weight fixes the LPAR weight for an LPAR if the user provides a fixed_weight to this LPAR in the input file `lwo_parameters_per_lpar`.

Configuring input files for LPAR weight optimization

To run the IBM Capacity Management Analytics LPAR weight optimization scripts, you must provide input data in separate CSV files.

The files that you must provide inputs in are:

lwo_parameters.csv

This file contains the following fields:

- EXECUTION_NAME—A name for the optimization run. The default value is *modelname_timestamp*. The optimization records are recognized by *execution_name*, *processor_type*, and *cpu_serial_no*.
- HOUR_START—The start time for the run. An integer number between 0-23.
- HOUR_END—The end time for the run. An integer number between 0-23. This value should be greater than or equal to the HOUR_START value.

  HOUR_END means another full hour. For example, if HOUR_START = 1 and HOUR_END = 13, then the time range is 01:00:00 - 13:59:59.
- CPU_SERIAL_NO—An identifier for the CPC. The last 4 digits of the CPU serial number taken from SMF70SER. This value can contain only English and numeric characters.
- PROCESSOR_TYPE—The name of the logical processor type. This value can be CP, zIIP, zAAP, or IFL.
- FILE_NAME_MODEL_PARA—The path and file name for the model parameter file. For example, *`CMAINSTANCE`*`/data/lwo_parameters_model.csv`. If the model parameter file is in the default location, you do not need the path.
- FILE_NAME_DATE_LIST—The path and file name for the date parameter file. For example, *`CMAINSTANCE`*`/data/lwo_parameters_date_list.csv`. If the model parameter file is in the default location, you do not need the path.
- FILE_NAME_PARA_PER_LPAR—The path and file name for the model parameter file. For example, *`CMAINSTANCE`*`/data/lwo_parameters_per_lpar.csv`. If the model parameter file is in the default location, you do not need the path.

lwo_parameters_model.csv

This file contains the sub-objective names and weight values for the optimization run. The sub-objective names include MAX_OVER_TARGET_WEIGHT, MAX_OVER_TOTAL_WEIGHT, SUM_OVER_TARGET_WEIGHT, and SUM_OVER_TOTAL_WEIGHT. The sub-objective weight values can be any positive number, where the larger the number means a higher priority.

lwo_parameters_date_list.csv

This file contains the dates to use as inputs for the optimization run. The dates should be representative of the usage you are trying to optimize. For more information, see "Choosing dates for the optimization" on page 112.

lwo_parameters_per_lpar.csv

This file contains the following fields:

- LPAR_NAME

  The name for each LPAR to be optimized.
- LPAR_PRIORITY

  The current priority for the LPAR. This value should be a positive integer number. The larger the number, the higher the priority. The default value is 1.
- FIXED_WEIGHT

  The fixed weight assigned to the LPAR, if any. This value is the percentage, such as 0.25 (25%), of the shared processor pool that you want the LPAR to have. This value can be a number between 0-1.
- TARGET_WORKLOAD_IMPORTANCE_LEVEL

  This value can be an integer number between 0-6 if the target workload of this LPAR is set by the WLM importance level from the SMF72 records.

  This value is null if the target workload of this LPAR is not set by the total LPAR CPU usage from the SMF70 records.

  This value is used to calculate the target workload.
- TARGET_WORKLOAD_PERCENT

  A number between 0-1. The default value is 1. For example, 0.25 is 25%.

  This value is used to calculate the target workload.

The target workload is calculated by the value in the TARGET_WORKLOAD_IMPORTANCE_LEVEL field with the value in the TARGET_WORKLOAD_PERCENT field.

If TARGET_WORKLOAD_IMPORTANCE_LEVEL is an integer number between 0-6, then the target workload = workload under the TARGET_WORKLOAD_IMPORTANCE_LEVEL + workload of the TARGET_WORKLOAD_IMPORTANCE_LEVEL * TARGET_WORKLOAD_PERCENT.

For example, if the TARGET_WORKLOAD_IMPORTANCE_LEVEL is 5, and TARGET_WORKLOAD_PERCENT is 80%, then the target workload = (workload of WLM importance level: 0 - 4) + ((workload of WLM importance level: 5) * 80%).

If TARGET_WORKLOAD_IMPORTANCE_LEVEL is null, then the target workload = LPAR CPU usage * TARGET_WORKLOAD_PERCENT.

Table 20: Target workload calculation examples

TARGET_WORKLOAD_ IMPORTANCE_LEVEL	**TARGET_WORKLOAD_ PERCENT**	**Target workload**
5	0.2	= sum of (imp level 0-4)'s workload + (20% imp level 5)'s workload
6	0.2	= sum of (imp level 0-5)'s workload + (20% imp level 6)'s workload
6	0	= sum of (imp level 0-5)'s workload

Table 20: Target workload calculation examples (continued)		
TARGET_WORKLOAD_ IMPORTANCE_LEVEL	**TARGET_WORKLOAD_ PERCENT**	**Target workload**
6	1	= sum of (imp level 0-6)'s workload
null	0.8	= total workload * 80%

Procedure

1. Go to the `CMAINSTANCE/data` directory.
2. Open and edit the contents of the `lwo_parameters.csv`, `lwo_parameters_model.csv`, `lwo_parameters_date_list.csv`, and `lwo_parameters_per_lpar.csv` files.
3. Save and close the files.
4. Update the files you have changed to read only.

Chapter 12. SPSS Modeler stream reference

Much of the data in IBM Capacity Management Analytics reports and workspaces is processed using IBM SPSS Modeler. This software uses programs that are referred to as "streams," which contain the instructions required to retrieve, prepare, model, and forecast the data.

MIPS forecasting stream

The predicted scores for CPU_BUSY_MIPS shown in the reports that include forecasts are produced using the MIPS forecasting stream.

Creating forecasts

The process of creating the forecast values that are provided as part of the IBM Capacity Management Analytics solution follows a standard approach used to forecast future values of things like stock prices, interest rates, and population. The process does not need to be understood at a technical level to make use of the predictions provided. The description of this process focuses on the how the data is handled and provides some basic information on evaluating the accuracy and utility of the predictions.

There are four basic tasks involved in building the forecasts for each LPAR/Processor type combination found on your system.

1. Preparing the data by aggregating it up to the appropriate level
2. Defining the time period for the forecast, including the number of rows to be used for validation of the predictions
3. Building the actual forecast model
4. Saving the predicted scores for use in reporting

For more information about IBM SPSS Modeler, including procedural information on modifying streams, see the product documentation (pic.dhe.ibm.com/infocenter/spssmodl/v15r0m0/index.jsp).

Data aggregation

The forecasts provided as part of the IBM Capacity Management Analytics solution are produced using the IBM SPSS Modeler application. The program code is contained in "streams", including the stream used to produce the Capacity Management Analytics daily forecast. There are corresponding streams for the hourly and monthly forecasts which are similar except for the level of aggregation used in the data preparation and a few other details.

The Daily Forecast stream

The IBM SPSS Modeler stream used to produce the daily forecast is named `mvspm-lpar-daily-forecast-timeseries.str`. The flow of the program is essentially left to right, following the arrows that connect the nodes that make up the stream, as shown in the following figure. This stream contains a script that drives the process and adds connections dynamically.

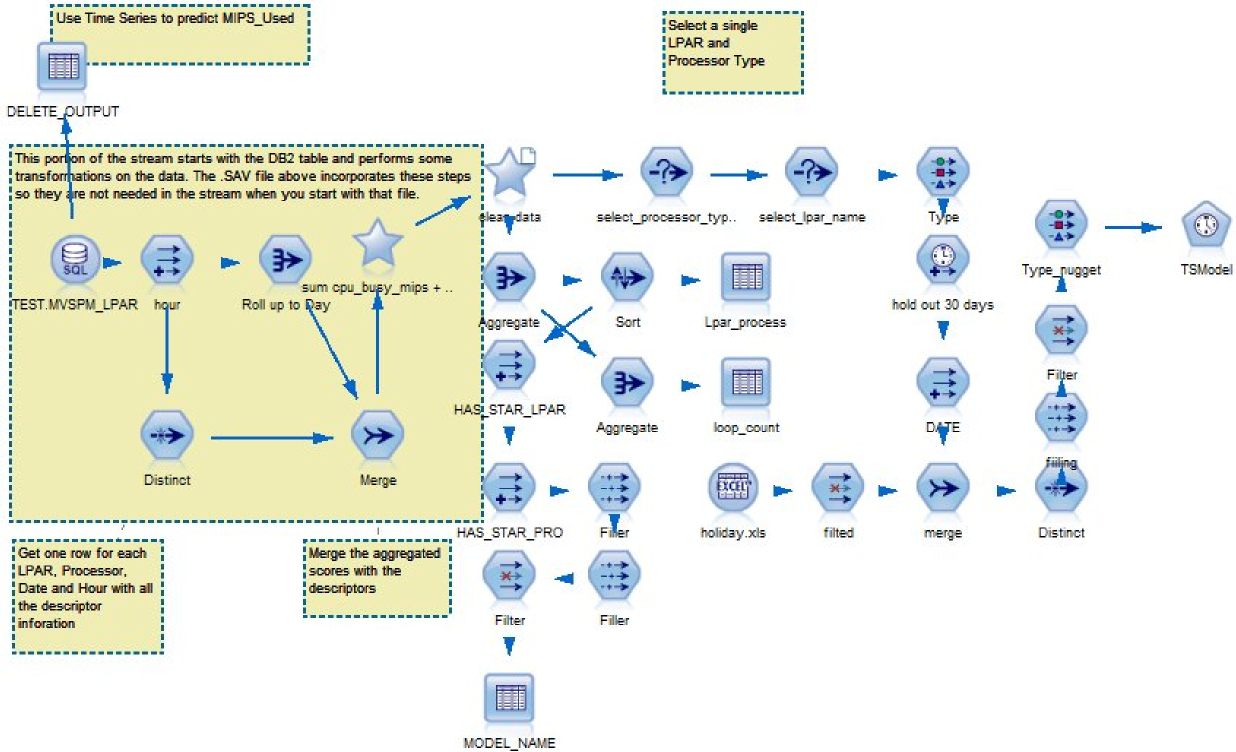

Figure 8: SPSS Modeler daily forecast stream

The first step in creating a forecast is to retrieve the data from the database and aggregate it up to the appropriate level. System Management Facilities data is often captured at 10 or 15 minute intervals and needs to be aggregated to the daily level for use in IBM Capacity Management Analytics. The following figure illustrates how this aggregation is configured in SPSS Modeler.

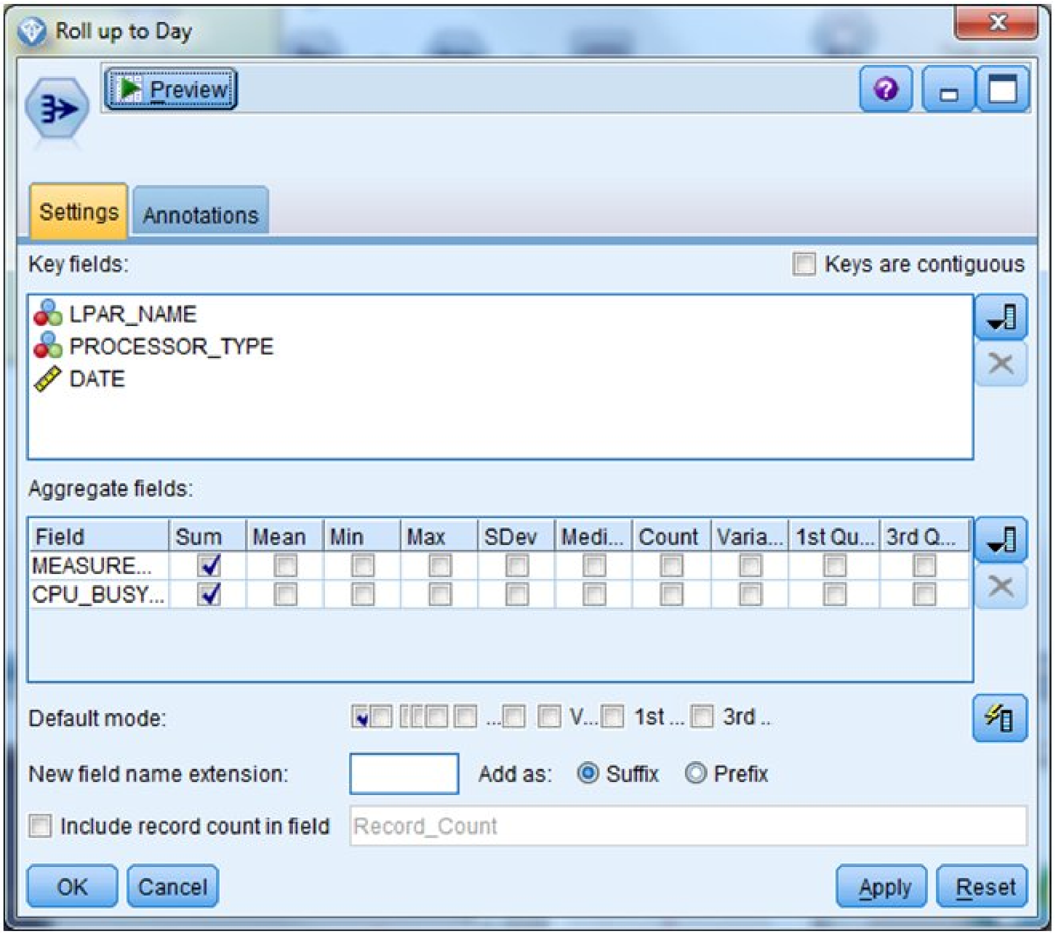

Figure 9: Aggregation to the daily level in SPSS Modeler

Time period definition
Defining the time period encompasses several distinct tasks, including modifying the `holiday.csv` file, modifying the Time Interval node, and setting the number of time periods.

Modifying the `holiday.csv` file

The `holiday.csv` file is where you identify, by date, important information that is related to demand on your system. This includes information such as holidays when load is typically lower or days of the year when you expect load to be unusually high or low. These high-demand and low-demand dates are days that you can anticipate based on your experience. A low-demand day might be the day before a major holiday. A high-demand day might be a day following a major holiday or a day at the beginning of a month (or quarter) when you do end-of-month reporting for the previous month. Identifying atypical days gives the forecasting algorithm additional information to use in accounting for fluctuations in CPU usage, which can help you to build more accurate models at the hourly and daily level.

Note: Date format in `holiday.csv` should be ISO format. If using a non-ISO format, the stream properties of forecast streams must be updated to match the date format in `holiday.csv`.

Tip: If you use Microsoft Excel to edit csv files and your computer uses a different locale than what is used in the SPSS Modeler streams, you must ensure that the data format in the csv remains consistent with date format that is used in the streams. Microsoft Excel can change the locale that is used in the file to match your computer's locale.

Modifying the Time Interval node

The Time Interval node, which is labeled **hold out 30 days** in the flow chart of the Daily Forecast stream, sets up the time-based parameters for your models to use. The time interval for modeling is defined on the **Intervals** tab in this node. Typically this is 7 days a week, as in the following figure. If your data does not include the full 24x7 period, the settings in this node must be modified. For example, for data that includes only weekdays, the **Number of days per week** field would be 5 and **Week begins on** option would be set to Monday.

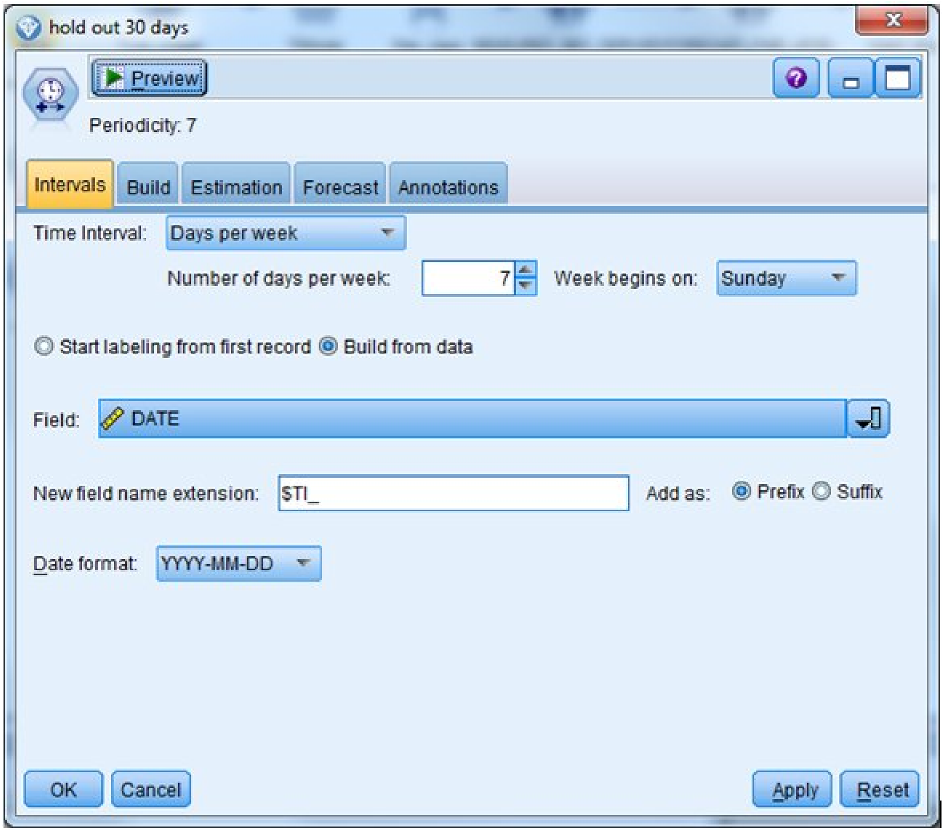

Figure 10: Time interval settings in the Time Interval node

A hold out period for validation purposes can also be defined in the Time Interval node. In a daily forecast situation, the last 30 days of actual data are "held out", meaning that the forecast model does not use this data to build the model. This allows the last 30 days of known CPU usage to be compared with the predicted values from the forecast model. This comparison serves as a means of evaluating how accurately the model predicted results for the last 30 days of CPU usage. This is the best way to determine how well the model will perform in predicting future events.

You can modify the hold out period on the **Estimation** tab of the Time Interval node.

Setting the number of time periods

You set the number of time periods to be forecast in the future on the **Forecast** tab of the Time Interval node. This instructs the program to generate rows of data for dates in the future beyond the end of the actual data available in your database. In the following image, the time period is days and a full year (365 days) is requested for the forecast. In the hourly and monthly streams, this value is set in hours and months, respectively. The values of the identifier fields such as LPAR_NAME and PROCESSOR_TYPE are taken from the most recent value in the original data.

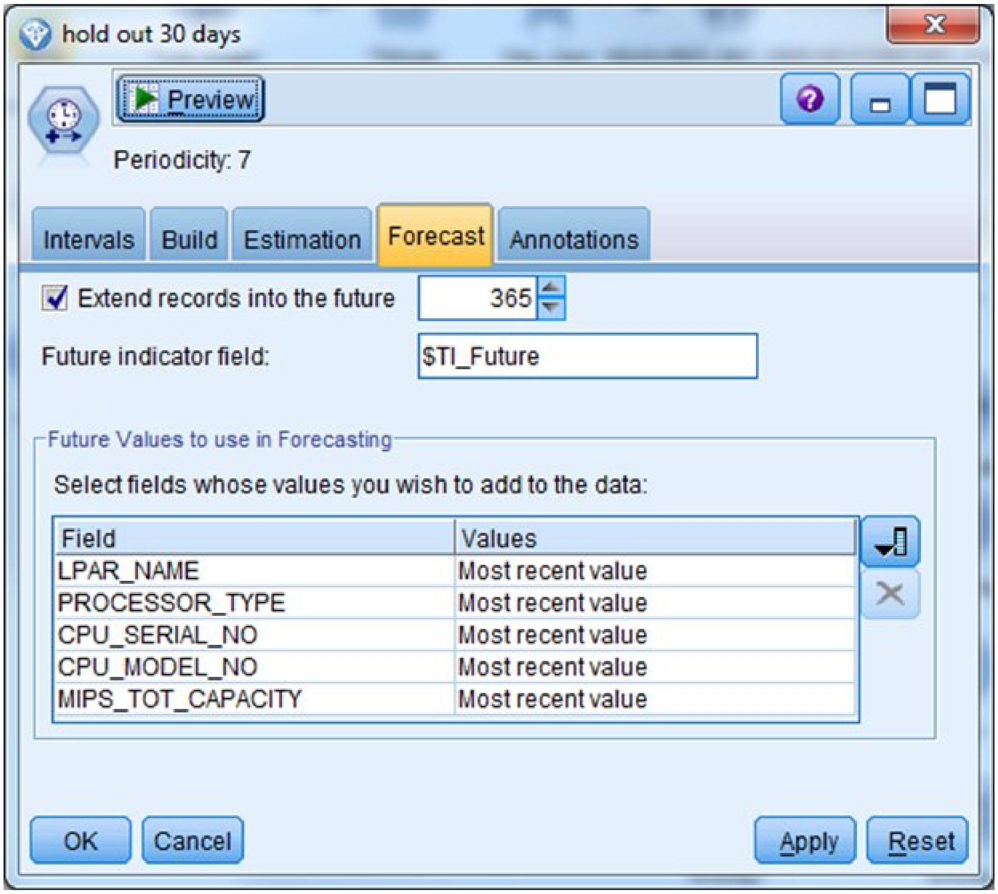

Figure 11: Forecast settings in the Time Interval node

It is important to note that the ability to forecast into the future is based in part on the amount of data that is used to construct a model. A year or two of data is desirable if the goal is to project daily and monthly usage for a year into the future. If only a few months of data are used to build a model, it is possible that for some LPAR and PROCESSOR TYPE combinations it is impossible to build a usable model.

Also, to prevent the forecast processing from stopping due to no usable model output, a check is made to filter out data for certain combinations. If the number of records in input data of a combination is not above a specific threshold (three records for a monthly forecast, 60 records for a daily forecast, and 720 records for an hourly forecast), forecast processing is not performed for that combination to build the model. If there is no model for certain combinations after forecast processing, you should review the input data for the combinations to ensure that the number of records exceeds the threshold at which forecast processing occurs.

Forecast model creation

The forecast model is built based on the specifications in the Time Series node in each stream.

The Time Series node is labelled **TSModel** in the flow chart of the Daily Forecast stream. This node is where the field to be predicted is identified, along with the fields to be used as inputs to the model. In IBM Capacity Management Analytics, CPU_BUSY_MIPS (summed across the relevant rows to the daily, hourly or monthly level) is the field that is predicted. In addition to the Holiday, high, and low fields from the `holiday.csv` file, the day of week is included as a predictor to be evaluated by the modeling algorithm.

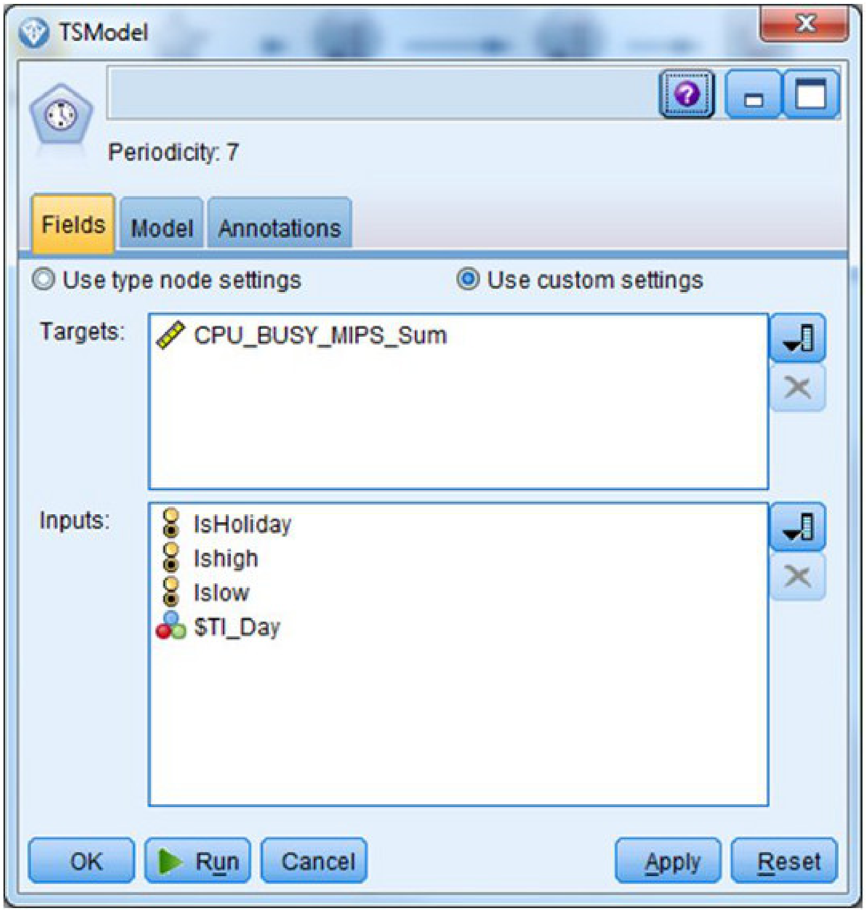

Figure 12: The Time Series node

The modeling itself is set on the **Model** tab of Time Series node. Capacity Management Analytics uses the **Expert Modeler** facility for time series models, set to evaluate Auto Regressive Integrated Moving Average (ARIMA) models and select the one that yields the most accurate result.

Running the stream

IBM SPSS Modeler streams for the hourly, daily, and monthly program are designed to run in an automated manner. When they run, all of the LPAR and PROCESSOR TYPE combinations found in the data are identified and a separate model is produced for each. These models are used to calculate the forecast for the original data and produce predicted values for the future. Each stream contains a script that manages this process. When the script is started it prompts for parameters that can be modified as necessary to point to the database source and the output directory for storing the models and the output directory for logs. Also, for hourly and daily forecast, an input PATH for `holiday.csv` is required. The output directories and the input path must exist before the script runs.

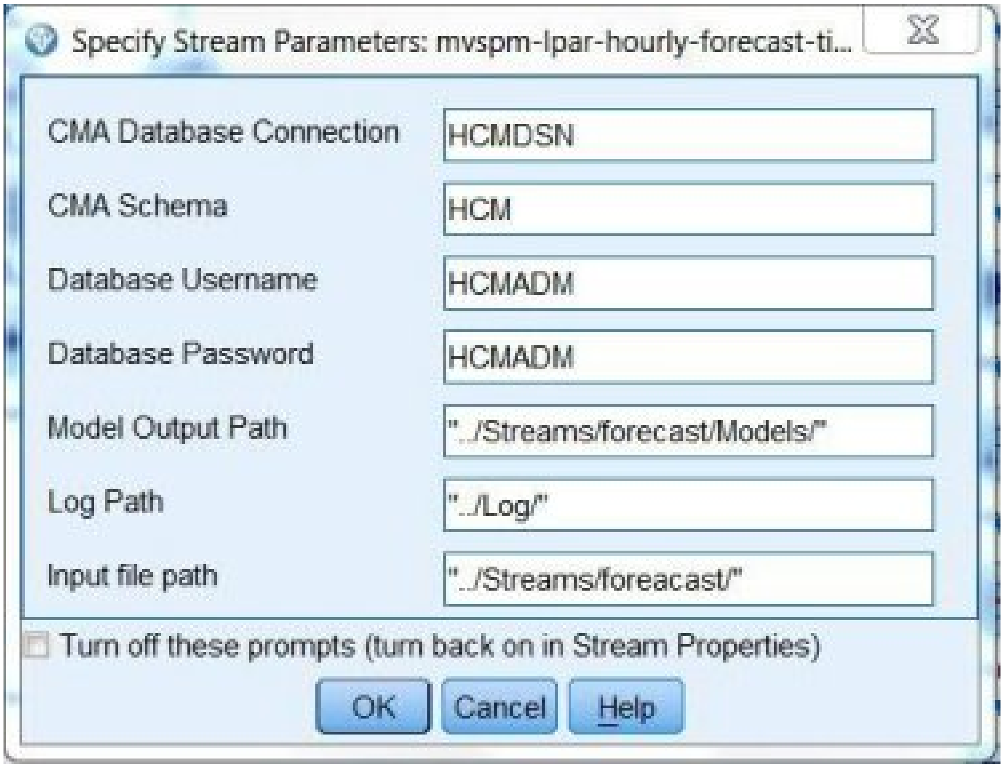

Figure 13: Prompts for stream parameters

The script itself is unlikely to ever require modification, but you can click **Tools**, **Stream Properties**, then **Script** in SPSS Modeler to view the script.

Predicted scores
The predictions that are produced by the models are saved to the database and stored in the `HCM.MVSPM_LPAR_CPU_BUSY_MIPS_PRED` table.

The predicted score field is named `CPU_BUSY_MIPS_SUM_PRED`, with its confidence interval named `CPU_BUSY_MIPS_SUM_PRED_L` as the lower bound and `CPU_BUSY_MIPS_SUM_PRED_U` as the upper bound.

There is also a `FORECAST_TIME_LEVEL` that indicates whether the forecast is for the hourly, daily, or monthly level. This table can be accessed with SPSS Modeler using an SQL source node similar to the node used to retrieve the original data from the same database.

The evaluation of models

As models are produced independently for each LPAR/PROCESSOR TYPE combination, they need to be evaluated separately to validate their utility.

The accuracy of a forecasting model can be assessed using a number of standard criteria. The usefulness of the forecasts is usually related to their ability to predict peak CPU usage with reasonable accuracy. Off-peak usage is generally less critical to predict accurately, which means that the overall accuracy of the model may not be the best gauge of its usefulness. The reports showing predicted values provide you with the best means of evaluating the forecasts.

The best way to evaluate the model results is to examine the forecast reports. These reports are available through IBM Cognos Connection. For example, the "CPU: MIPS Used - zServer/LPAR Level" on page 224 report contains both actual values and forecasted values.

The following image is an example of the CPU: MIPS Used -zServer/LPAR Level report with forecasts included. The gap between the dots on the line and the top of the bar indicate the degree to which the predicted value matched actual CPU usage.

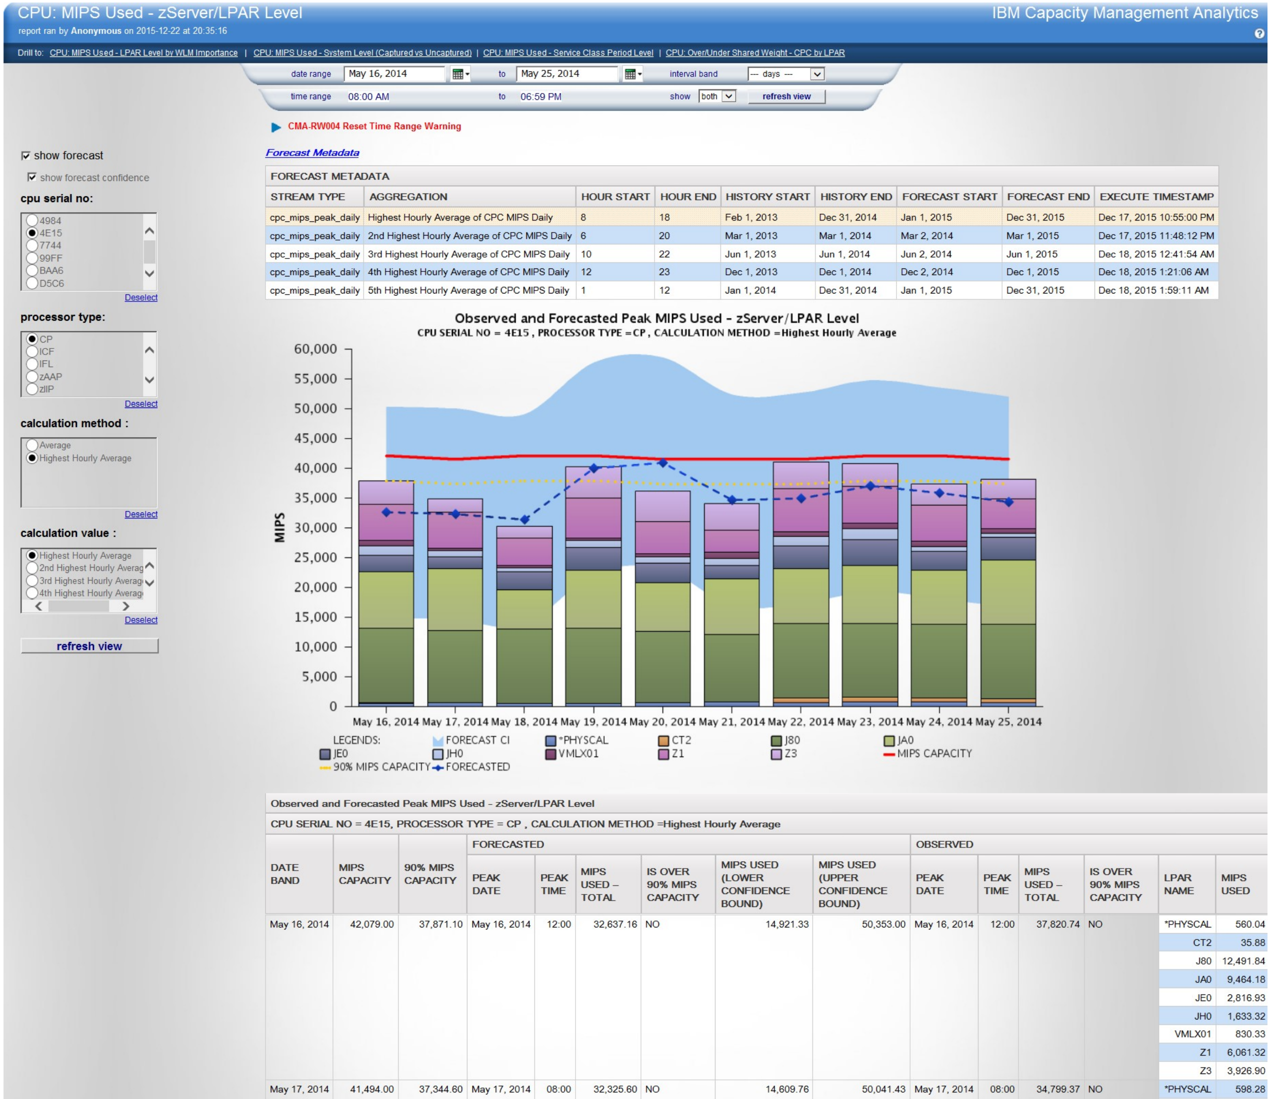

Figure 14: CPU: MIPS Used -zServer/LPAR Level report with forecasting

You can use the table that accompanies the graph to make precise comparisons. If your main objective is to accurately identify peak loads, you should confirm that in most cases the predicted value is equal to or greater than actual value. The closer the actual and forecasted values are over the last 30 days, the more confidence you can have in the accuracy of future forecast values. Around any estimate there is a confidence interval and the width of this interval (or band around the predicted value) widens as time extends into the future. This means forecasts closer to the last data point are likely to be more accurate than those farther out in time.

Details for hourly forecasts

Time Interval node settings for the hourly forecast are based on a 24x7 time period. Note that the Time Interval is set to **Hours per day** and the timestamp field provides the information necessary to identify the date, day of the week, and hour of the day.

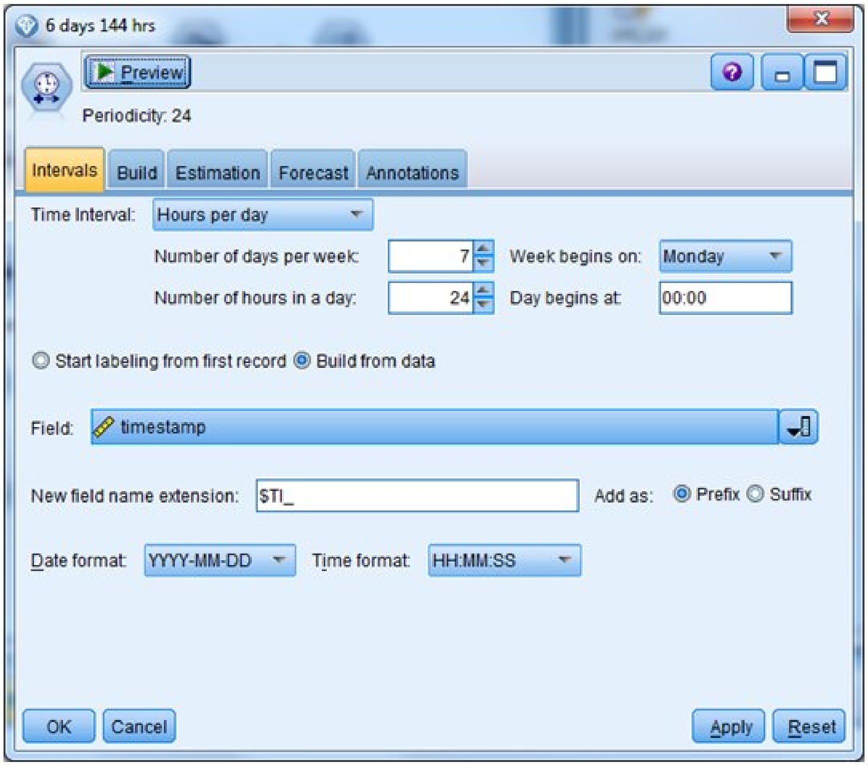

Figure 15: Time Interval node for the hourly forecast

The hold out period for the hourly forecast is 144 hours (6 days).

Details for monthly forecasts

Time Interval node settings for the monthly forecast is set to months and the date field provides the information necessary to assign each row of data to the appropriate month.

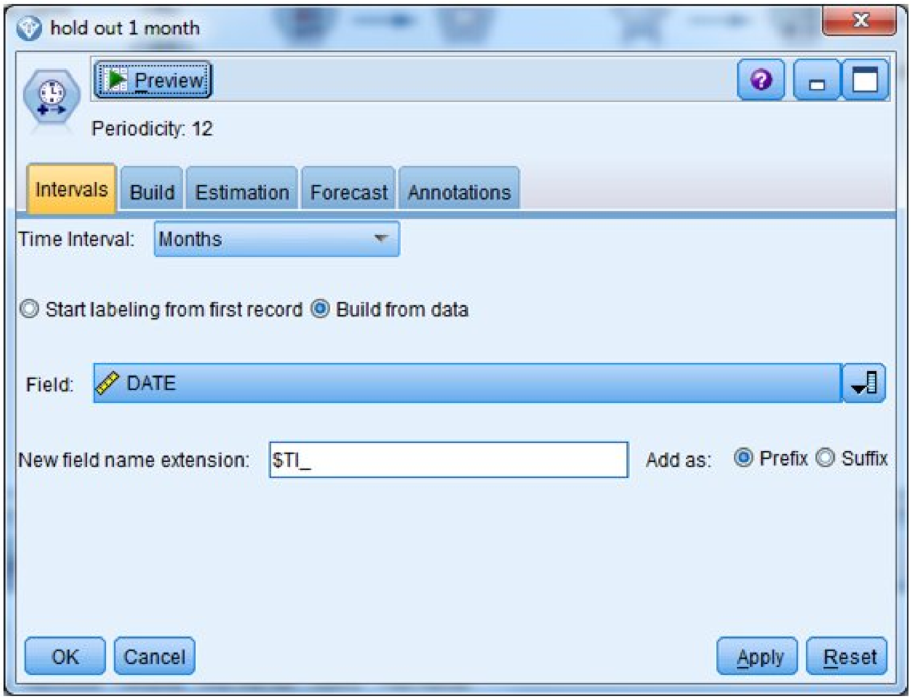

Figure 16: Time Interval node for the monthly forecast

In the monthly forecast, one month of data is held out for validation purposes. Forecasts are provided for twelve months into the future.

Details for the CPC daily peak forecast stream

The CPC daily peak forecast stream (`mvspm-cpc-daily-peak-forecast-timeseries.str`) is used to produce daily CPU_BUSY_MIPS :)peak forecasts at the CPC level. The stream analyzes a defined time range of the historical data, calculates its peak value of MIPS, and forecasts its one year peak value at the CPC level for the future.

CPC level

The CPC level is defined by the combination of the CPU_SERIAL_NO and PROCESSOR_TYPE fields. A different combination of the fields result in a different CPC level.

Peak values

The peak values are defined as a serial that is formed by the Top-N MIPS values in each day, where N is specified as the Peak Highest Rank. You can also input multiple N values to forecast multiple peak value serials in a single run. The maximum value of N is 5.

Input table

- HCM.MVSPM_LPAR: the source of history data for MIPS.

Output table

- HCM.MVSPM_CPC_CPU_BUSY_MIPS_PRED: the forecasted data and the related history data for MIPS of each day by each CPC level for each Peak Highest Rank is stored in this table.
- HCM.FORECAST_METADATA: the execution information of peak forecasting stream for each Peak Highest Rank is stored in this table.

Input parameters

- CMDW: the CMA database connection.
- CMASCHEMA: the schema for CMA.
- UID: the database user.
- PWD: the database user's password.
- ENCODED: the encoded database user's password flag. (true/false)
- OUTPUTMODELPATH: the model output path.
- TMPPATH: temporary data saved path.
- INPUTHOLIDAYFILE: the input holiday file path.
- DATE_START: the start date for the data that is used.
- DATE_END: the end date for the data that is used.
- HOUR_START: the start hour for the data that is used. An integer number between 0-23.
- HOUR_END: the end hour for the data that is used. An integer number between 0-23.

  This value should be greater than or equal to the HOUR_START value. HOUR_END means another full hour. For example, if HOUR_START = 1 and HOUR_END = 13, then the time range is 01:00:00 - 13:59:59.
- PEAKRANK: the Nth highest peak rank values. Valid values are integers in the range [1,5]. Use a comma (,) to separate rank values if you are using multiple rank values.

Validating inputs

The stream verifies that input values are valid. For example, the stream verifies whether the path and database are writable. If not, it stops and an error is written to the log file. The log file is stored in LOGPATH.

Validating historical data

The stream verifies that the historical data in the HCM.MVSPM_LPAR table is valid to do a forecast. For example, the stream skips the series if the data points of a CPC for a peak are less than 60.

Forecast data generation and output

There is a Python script in the peak forecast stream that drives the process. The stream finds the peak values by each Peak Highest Rank N by each CPC level within each day. The peak MIPS values are selected from the MIPS values by each hour within a day. After a serial of peak values is produced, the stream analyzes the data by using the Time Serials model, outputs and stores each of the models that are trained in OUTPUTMODELPATH, and uses them to forecast one year of peak values in the future by each CPC level. The Peak Highest Rank N of the forecasted data is labeled in the field FORECAST_LEVEL such as Highest Hourly Average of CPC MIPS Daily for N=1. The saved models are named `<CPU_SERIAL_NO>_<PROCESSOR_TYPE>_Peak_<N>_daily_<timestamp>.gm`.

After the forecast values are produced, they are gathered in a temporary file that is stored in TMPPATH named as CPC_MIPS_DAILY_PEAK_FORECAST_TEMP_<*peak_n*> and then output to the database. If the database already contains MIPS forecasted data with the same peak level, the new values overwrite the existing values. Values that do not have the same peak level as the new values are kept in the database.

Estimation for missing data

If there are any missing data points in a series, the stream generates an estimator by the average of the other data points in the series.

Adjusting outliers

If any prediction is greater than 2 times the max value of the historical data in a series, then that prediction value, the upper bound and the lower bound of the confidence level, are adjusted to 2 times the max value of the historical data. If any prediction (including the 2 bounds of the confidence level) is less than 0, then the value is adjusted to 0.

Execution data generation and output

The information of the latest stream execution time and the data time range for history data and forecast data along with the related Peak Highest Rank N (represented by filed AGGREGATION in the table, for example, Highest Hourly Average of CPC MIPS Daily) are gathered when the forecasting runs. This information is written to the database after the forecasting finishes. Each running for each Peak Highest Rank makes one record in the table. If there are old records in the database, any values with the same Peak Highest Rank are overwritten if new values are produced. Any existing values that do not have the same Peak Highest Rank are kept in the database.

Log file and error handling

You can find the IBM SPSS Modeler stream log files in the log folder that you specify in the LOGPATH parameter. The log file is named `mvspm_cpc_daily_peak_forecast_timeseries_<timestamp>.log`. For information about error and warning messages, see Appendix F, “Error, warning, and installation messages,” on page 315.

Details for LPAR peak daily forecast stream

The LPAR peak daily forecast stream uses the historical daily peak MIPS value for each CPC, LPAR, and processor type combination and forecasts the daily peak MIPS for one year into the future. The stream requires at least 60 days of continuous data, without any missing values, to do the forecasting.

You can specify the date and time range that you want to use as the historical data when you run the stream. This stream also supports different ranks of peak calculations; for example, highest peak, 2nd highest peak.

The stream takes external factors as input values; for example, holiday information is used as an external factor. You can input your own external factors to replace the holiday files. When you use customized external factors, these factors must contain a value for all future records. You can specify customized external factors in a text file. The text file must use the same format as the sample holiday file.

To use the customized external factors, you must specify the path to the text file in the INPUTHOLIDAYFILE parameter when you run the stream.

Error handling for model building

If an error is encountered while building a model for a specific LPAR/PROCESSOR TYPE combination, tables are generated that contain details related to the error.

These tables are saved into the `Log` directory that was created as part of the IBM Capacity Management Analytics configuration process. The name of a given table contains the LPAR/PROCESSOR TYPE that produced the error. The information in these tables is useful for diagnosing underlying issues with your model. In most cases, the problem will be related to insufficient data on which to build a model. This can occur if a LPAR/PROCESSOR TYPE combination has only a few months of data available.

Software Cost Analysis forecasting streams

Product forecast

There are three forecasting streams in the `sca` folder:

- `msu_prod_forecast_timeseries.str`

- msu_lpar_forecast_timeseries.str
- msu_gssp_forecast_timeseries.str.

These forecast streams are used to predict the future MSU usage of a machine for a certain product or certain LPAR, which is the basis for billable MSU calculation and optimization.

Approach for forecasting

The product forecast streams and the GSSP forecast stream use similar approaches to forecasting MSU. The LPAR stream aggregates MSU from product level to LPAR level.

MSU (Millions of Service Units) refers to software pricing capacity units. Product MSU, which refers to product consumed in millions of service units hourly, is the target used to forecast. The IBM SPSS Modeler stream used to forecast MSU for each product on each LPAR is msu_prod_forecast_timeseries.str. A product named ‘9999-OTH’, also called Others, represents the remaining MSUs that are unaccounted for, such as NO89 products. This Other category is also forecast. The first step in creating a forecast is to retrieve the data from the database and aggregate it up to the hourly level. PROD_MSU is the field name of the forecasting target, while PROD_MSU_FORECAST is what is forecast from the built model. PROD_MSU_4HRA_FORECAST is the average of PROD_MSU_FORECAST over the past four hours and PROD_MSU_4HRA_UTILRA_FORECAST is the minimum of PROD_MSU_4HRA_FORECAST and defines capacity.

SCA forecast is different from the MIPS forecast stream in structure. To improve the accuracy of the model, the data is split to the daily level. This means the data is put together for Monday and then forecast. So there are seven models in each product on each LPAR level in general.

All models are stored in a specified location. This location is specified in the prompt window when starting running the stream in client mode or in the CMAINSTANCE/data/cmasca.parms file if running the stream in shell mode. The default value for the location is the CMAINSTANCE/model folder.

The time interval nodes in SCA don’t need to be modified manually, but can be managed by scripts automatically instead. Product MSU is forecast for the next year by default, that is, the forecast is performed for the next 52 weeks multiplied by 24 hours, which equals 1248 records. You can optimize this but you need to forecast at least one month's (4x24=96 records) data.

Inputs for model

The holiday file (CMAINSTANCE/data/holiday.csv) is used to input for the build model for MIPS forecasting. For data that has a large enough time range, you can use data from the same time the previous year as another input for forecasting. Scripts can manage this automatically. For data consisting of one day of the week on each product and an LPAR combination that has more than an 18 month time range, holiday and PROD_MSU_LAST_YEAR are used as input fields. For data that is less than 18 months, only holiday is used as an input field.

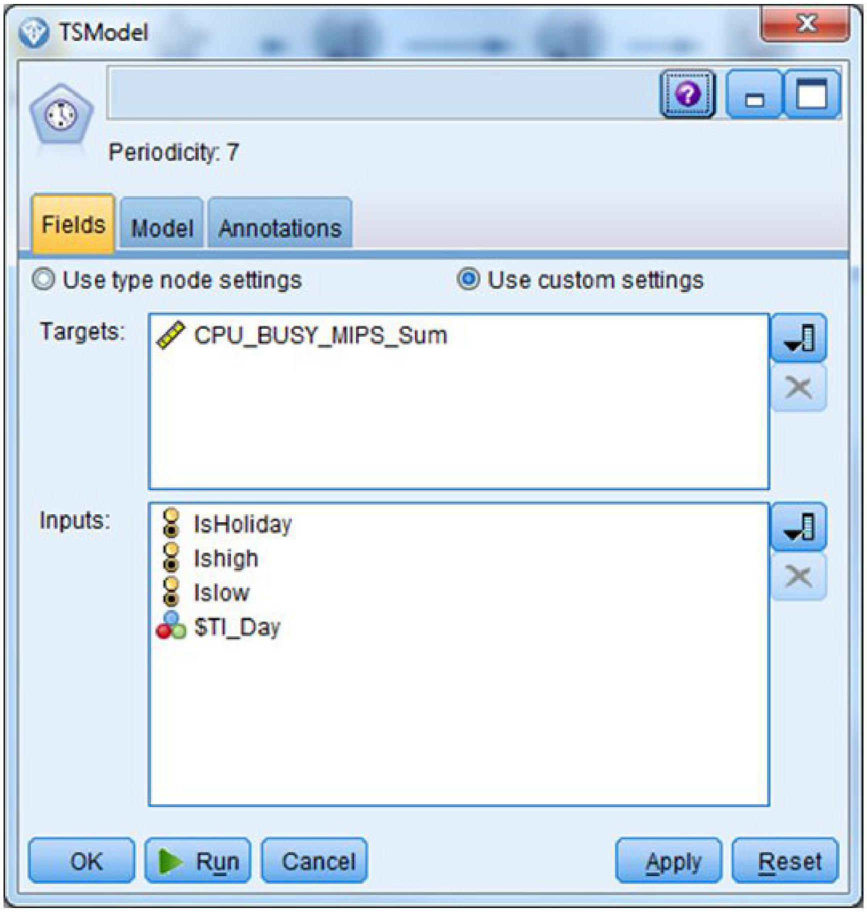

Forecast protection

An interruption can occur if a time series is incorrect for the model used or when Modeler stops because there is not enough data for building or has too many null values for the target field. In SCA, forecasting streams address some of these types of interruptions.

For models that are built for hours of one day of the week, the time series used builds the model if the data used for the last day of building has 24 records and there is no null value in target field. So data previous to the 24 hour period that has no null value is used to build a model for each day of the week on each product and LPAR combination.

In MSU data, some of the LPARs or products run only for a short time and then stop. The ODD parameter is used to handle these obsoleted products. If a product no longer exists or does not run for recent ODD days in one LPAR, then we assume that this combination is obsolete and the stream does not forecast this combination.

You can view a prompt table that includes all the combinations that are used to forecast which meet the previously mentioned conditions if running the stream in the client.

The SCA product forecast stream uses the date from DATEMIN to the latest date to forecast all of the product and LPAR combinations. The forecast stream uses the latest operation system version that is in the observed data to perform a forecast for each LPAR, and ignores any other operation system versions previously run by the LPAR.

The combinations that cannot meet the previously mentioned conditions are written into the log file where you entered the parameter of `Log Path` in the prompt window. The default is `CMAINSTANCE/log`, at the beginning of running the stream.

Note: there is no combination message if there is no record of Sunday in a product on one LPAR. You can fill in the corresponding missing values in the database if you want to forecast the MSU for a combination in the log file.

Monthly level

Aggregate the MSU hourly forecast data to a monthly level using the max method and insert it into the same table where hourly forecast data exists in the database for the SCA forecast procedure. You can distinguish the hourly data from the monthly data with the `FORECAST_TIME_LEVEL` field.

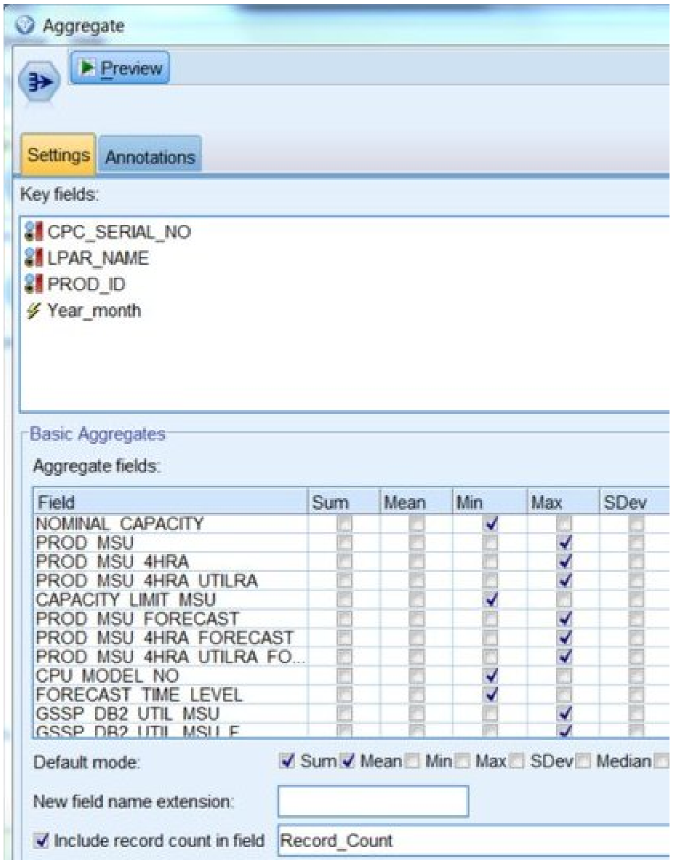

The LPAR Forecast Stream

The IBM SPSS Modeler stream used to forecast MSU for each LPAR is `msu_lpar_forecast_timeseries.str`. Aggregate all product MSU and Others MSU to the LPAR level and get the MSU usage of each LPAR in one CPC (CEC).

The GSSP Forecast Stream

The IBM SPSS Modeler stream used to forecast the GSSP MSU for each product on each LPAR is `msu_gssp_forecast_timeseries.str`. As the use of the GSSP forecast stream and the product forecast stream are almost the same, refer to the "Product forecast" on page 128 section for information.

The GSSP forecast data isn't inserted into a separate table in the database. Instead, it is saved in a temporary file and merged with product forecast data. It is then inserted into a table in `PROD_MSU_FORECAST` in the product forecast stream database. So you should run GSSP forecast stream before running product forecast stream.

CICS anomaly detection streams

Customer Information Control System (CICS) anomaly detection helps to improve system management by proactively detecting anomalies in forecast capacity so that you can identify issues before they affect the business.

You can identify items that differ from other items or do not fit into the expected pattern in a data set. Items that are found to be anomalous typically indicate a problem such as a structural defect, bank fraud, finding errors in text, or medical problems. Anomalies can also be referred to as outliers, novelties, noise, deviations, and exceptions. The faster you can identify unanticipated transactions, the faster you can take corrective action to circumvent any effects on the performance of the business application.

Results of the stream are shown in the “Anomaly Detection Analytics: CICS Anomaly Details” on page 194 report.

IBM Capacity Management Analytics provides a scoring stream that reads the model building information, calculates the prediction results for each transaction type, determines which transactions are abnormal, and then exports the results to the database. Logging the history of abnormal transactions lets you report on trends and patterns of the anomalies. It is also a source for any specific alerts that are based on different levels or scope of anomalies detected.

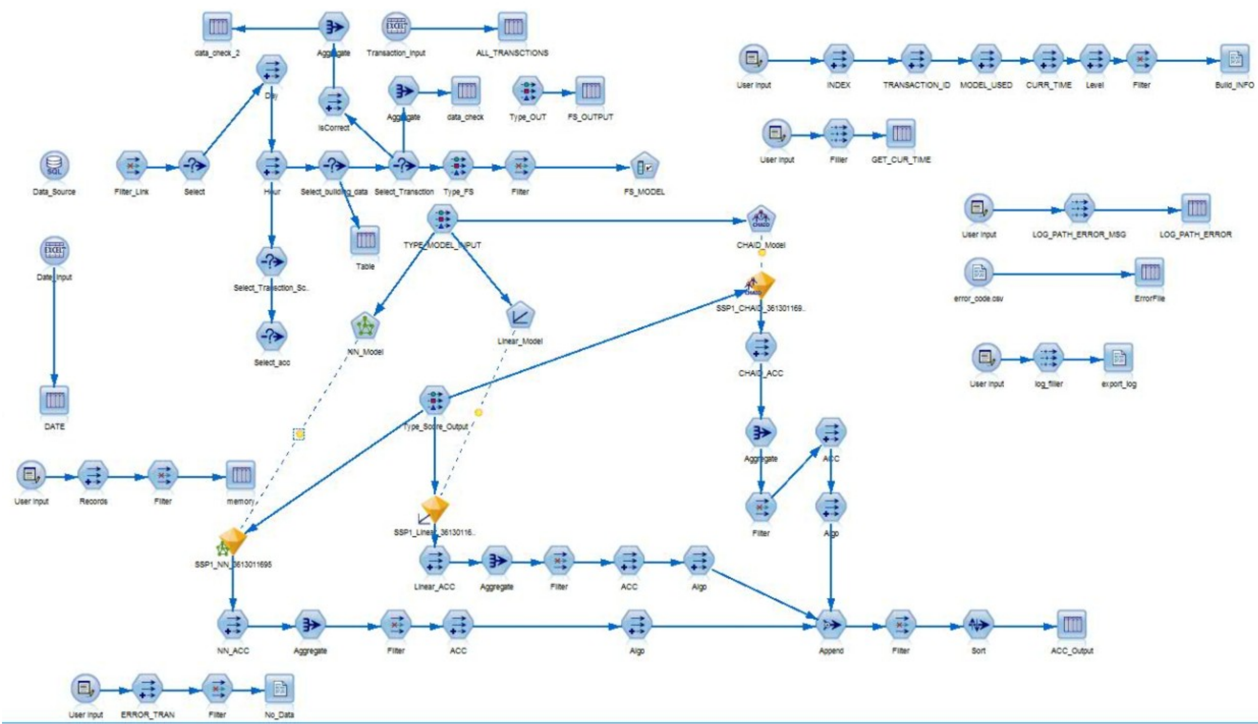

Transaction

A transaction is a set of operations that perform a task together. Typically, a transaction is a simple task such as entering an account debit or credit or requesting an inventory list. A transaction should be primarily atomic.

Abnormal transaction

A transaction is abnormal if its execution characteristics differ by some ratio from its transaction model, based on normal transaction processing. Characteristics that can change only by logic in the program for a transaction would cause the transaction to be detected as abnormal and its logic be investigated. However, a number of characteristics can be influenced in varying degrees by conditions outside the scope of the transaction itself, such as the environment in which the transaction is running (in effect, shared resources such as CPU, memory, storage, and queuing because of different concurrency controls). The characteristics to focus on can be determined from the SMF 110 records that are written by CICS. For more information, see the CICS_T_TRAN_T_table documentation on IBM Knowledge Center.

Anomaly model building

Capacity Management Analytics has two different model building streams: `anomaly_detect_building_c.str` and `anomaly_detect_building_r.str`.

The streams contain these steps:

- Data preparation
- Model building
- Model evaluation

The models are built with default options. You can change the options to customize the models for your environment.

Trained models are stored locally in the `CMAINSTANCE/model` folder and model information is found in a log file in the `CMAINSTANCE/log` folder.

Data source

The CICS_T_TRAN_T table provides statistics on CICS transactions. It contains information from CICS performance class monitoring records for CICS/ESA (SMF 110, subtype 1) and CICS/MVS (SMF 110, subtype 0). Related table definitions can be found in the CICS_T_TRAN_T table documentation.

Data preparation

CICS transaction data creates records for each transaction with more than 300 columns. In a production environment, records with more than 10 million transactions per day is not uncommon. Most of the columns are not necessary to analyze abnormal transactions, so Capacity Management Analytics includes an iterative process that uses the modeling tool to help you build a detection model that fits your transaction profile. The process helps to speed the process of learning which of those columns are the most important ones for your transactions.

- Record reduction for an input table
  - you can specify which transaction needs to be examined by using the `transaction.csv` file.

	A
1	TRANSACTION_ID
2	SSC1
3	

  - you can specify the date range for model building and evaluation by using the `date_info.csv` file. The stream will use data from the time 00:00:00 in the start date to 23:59:59 in the end date.

	A	B	C
1	Type	Start_date	end_date
2	building	2014-05-29	2014-05-29
3	evaluation	2014-05-30	2014-05-30
4	Scoring	2014-06-10	
5			

- Dimension reduction for the input table
  - Reduction dimension of input data by using predefined filter node
  - Reduction dimension of input data by using feature selection model in model.
- Adding external factors for model building

  The day of the week and the time of day need to be introduced because transactions in a production environment have some periodical patterns. For example, many specific types of transactions happen on Mondays.

Note: IBM SPSS Modeler has a data limitation for model building. The limitation depends on the size of input data. You can find detailed information from the Ram Requirements section of the SPSS Modeler documentation on IBM Knowledge Center. The stream is protected such as to avoid a memory exceeded error. If an error occurs, you can verify it by viewing the error message in the error log. If the error is because of the size of the data set used, you need to limit the data set to just a sample size. This prevents errors that are caused by data size and memory limits.

Model building

Different streams build a model by using different targets.

- The `anomaly_detect_building_c.str` stream builds a model on the transactions that detect an anomaly that occurs in CICS.
- The `anomaly_detect_building_r.str` stream builds a model on the transactions that detect anomalies in both CICS and in an z/OS environment.

Both streams use different data mining algorithms for predicting normal transactions. Models are built with default options. Input field models use results from the feature selection model that is used in the data preparation steps.

Model evaluation and selection

The mean of relative error and standard deviation of relative error is used to evaluate and predict results accuracy. Rank models by accuracy and select the one with best accuracy for scoring.

$$\frac{\sum_1^n |\text{target} - \text{predict})/\text{target}|}{n}$$

Anomaly model scoring

CMA provides a scoring stream to score models built by model building streams. Streams need to read model-building information, calculate prediction results for each transaction type, determine which transaction is abnormal, and then export the result to the database. Scoring streams support SQL push-back in Modeler, and can use the scoring adapter for zEnterprise V16 to increase scoring performance. Refer to the zEnterprise V16 documentation for more information.

Define abnormal transactions
Abnormal transactions can run for a long time or for a short time.

Store scoring result
A scoring stream stores the scoring result to the database for future reference. The **Abnormal** field indicates whether the transaction is abnormal or not.

Tree-based anomaly detection model

There is one tree-based anomaly detection stream: `anomaly_detect_treebased.str`.

Data source
The `CICS_T_TRAN_T` table provides statistics on CICS transactions. It contains information from CICS performance class monitoring records for CICS/ESA (SMF 110, subtype 1) and CICS/MVS (SMF 110, subtype 0). Related table definitions can be found in the CICS_T_TRAN_T_table documentation.

Data preparation

CICS transaction data has records for each transaction with more than 300 columns. In a production environment, more than 10 million transactions occur in a day and, to analyze abnormal transactions, most of the columns are not required. This tree-based model requires that all input columns are continuous values. Therefore, a data preparation process must be run before the model is built.

There is a process of record reduction for the input table:

- You can specify which transactions need to be examined by using the `transaction.csv` file.

	A
1	TRANSACTION_ID
2	SSC1
3	

- You can specify the date range for model building and scoring in the `date_info.csv` file. The stream uses data from a time of 00:00:00 on the start date to 23:59:59 on the end date.

	A	B	C
1	Type	Start_date	end_date
2	building	2014-05-29	2014-05-29
3	evaluation	2014-05-30	2014-05-30
4	Scoring	2014-06-10	
5			

- The tree-based model requires only a sample of the data to build the trees. It uses no more than 2048 records, so a sample function is applied to select the model building data.

There is a process of dimension reduction for the input table:

- Reduction dimension of the input data by using a predefined filter node.

Model building

This tree-based model is implemented by an IBM SPSS Modeler embedded Jython script. The model builds 10 different trees.

Model scoring

Scoring data feeds this model and it outputs an anomaly score based on the average path length of the input data in the trees. A conversion is applied to the average path length to get the final scoring result. When the score result is greater than 0.5, it means that this transaction is probably an abnormal transaction. In the stream, this ratio is set with 0.5 as the default. The value for this ratio is recommended to be no greater than 0.6. and no less than 0.5 according to this algorithm.

The scoring result is saved to the database with anomaly_level = "TREE".

Anomaly scoring classification

This stream takes the scoring results classified result to a different level of abnormality. Mark all transactions with red, yellow, or green flags, then store the result to the database and output a distribution graph.

Note: Before you use this classification stream, the other streams must run for the same transaction in the database or else an error is recorded in the error log and you are prompted for data.

CICS Anomaly detection customization stream

The anomaly detection stream is used as a general approach for anomaly detection problems. The advantage of the stream is that it can run on a CICS transaction record from a different industry. However, its generic approach means that model accuracy might not be good for different scenarios. This section explains how to customize the stream to fulfill a specific requirement.

Model building streams customization

Data source customization

The anomaly detection stream uses the CICS statistic table that is provided by TDSz. You can use your own SMF parser to parse SMF 110 data, and then use the data for anomaly detection. If you choose to do this, you need to change the data source node to the type of data source you are using. Also, in Jython script you need to change the data source definition ("`Data_Source_Node=s.findByType("database","Data_Source")`" to the type of data source you are using. Keep the node name as `"Data_Source"` and do not change it.

Data preparation customization

- Change the filter node at the beginning to include or exclude fields that are related with the target on the data set.
- Add other external factors by using derive nodes. You can insert derive nodes between the filter and day node.

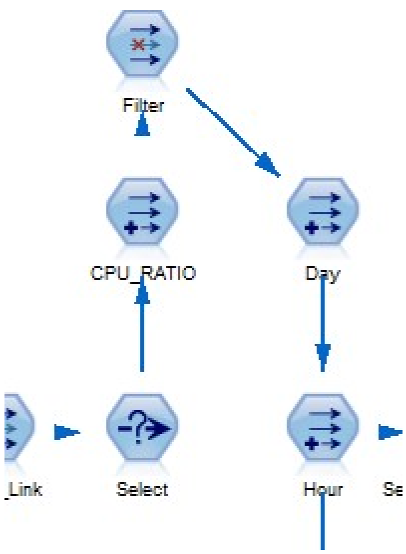

Algorithm customization

The anomaly detection model building stream uses different algorithms to build model candidates for evaluation. According to your specified data, you can select a different modeling node in SPSS Modeler. If a new modeling node is introduced into the stream, the stream script needs to be changed. Use the existing script as an example in developing a new script. Related scripts start from line 137.

Model evaluation customization

Model evaluation uses related errors as evaluation criteria. You can specify your own evaluation criteria. To do this, you change the evaluation formula node for each model you build.

Model scoring streams customization

The model scoring stream is related to the model building stream. If you have customized the model building stream, you need to also customize the model scoring stream.

Data source customization

- Use the data source with the same field that is used in the model building stream
- Use the same filter node that is used in the model building stream. `Filter_Link_C` should be the same as what is used in the `anomaly_detect_building_c.str` stream and `Filter_Link_R` should be the same as what is used in the `anomaly_detect_building_r.str` stream.

In the `date_info.csv` file, you can specify the start date for model scoring. The scoring stream uses data according to `START_TIMESTAMP >= start_date`.

Anomaly transaction criteria customization

The scoring stream uses a formula for its criteria. You can change the script (line 220 and line 233) to apply your own criteria.

Anomaly classification stream customization

Classification logic is found in the combine super node. You can open this super node and edit the logic inside to apply your own classification.

Billable MSU optimization stream

This stream is used to determine how to rearrange qualified products, or parts of products, among LPARs to lower the total cost of the billable MSUs. A qualified product is an IBM product that can produce billable MSUs in the forecasted data. If the forecasted MSU usage for a product is zero or if it is not an IBM product, then it is not included in the optimization.

This optimization stream is named `prod_bill_msu_optimization.str` and is in the `CMAHOME/features/cmasca/analytics` folder.

The main objective of the optimization is to move each qualified product, or part of the product, to one or more suitable LPARs to reduce the billable MSUs and total cost for the product. The report includes suggestions to help you determine where you can move the product. If the suggestion is to move part of the product to an LPAR, the suggestion also includes a recommendation for the amount of the product to move. The amount is measured in MSUs, with a granularity of one MSU, and it is recommended that you move the product's workload, equal to this MSU amount, to the suggested LPAR.

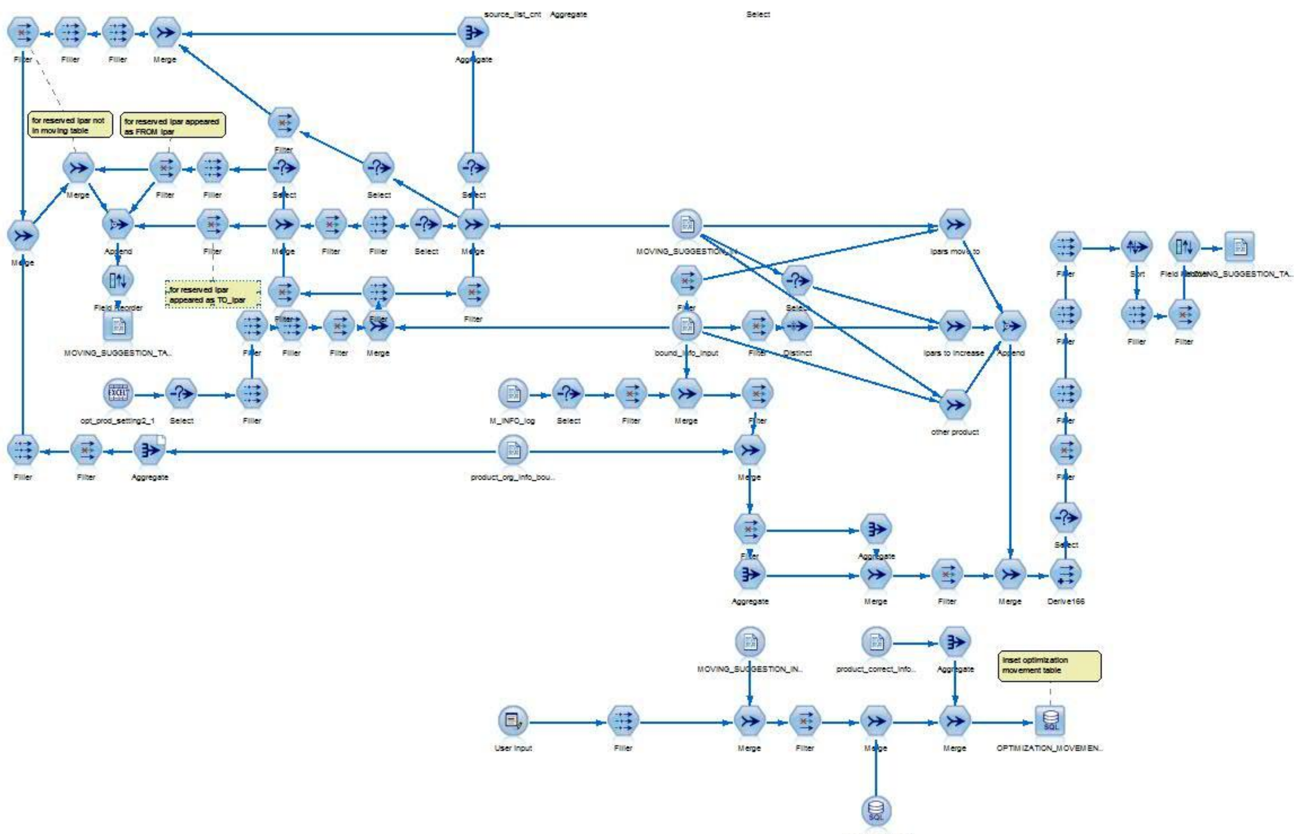

Figure 17: Billable MSU optimization stream

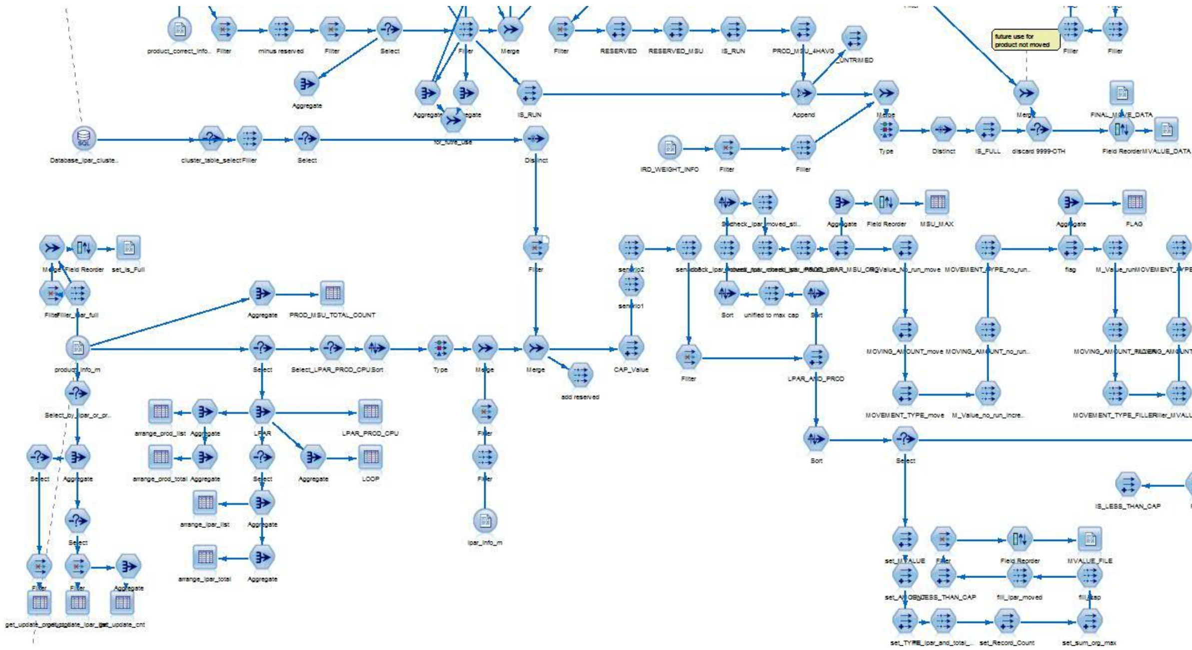

Figure 18: Billable MSU optimization stream

The results of this stream are output to the database and stored in four tables individually:

- PROD_MSU_OPTIMIZATION—for new distribution of the products among the LPARs, at the product level.
- LPAR_MSU_OPTIMIZATION—for new distribution of the products among the LPARs, at the LPAR level.
- OPTIMIZATION_MOVEMENT—for moving suggestions that are used for the new distribution.
- OPTIMIZATION_BILLABLE—for billable MSU calculated from the new distribution.

This stream performs the optimization based on the forecasted MSU usage of the products and LPARs.

Note: there could be one year for the forecasted data, but the optimization stream uses the first month of the forecasted data to determine the new distribution. You can assume that the configuration for this new distribution would be applied only for the next two months from the beginning of the forecasted data. It is better to have a short range for the optimization suggestion as the performance of the computer that is running the optimization might change significantly during a longer run and the forecasted data needs to be updated continually to give a more accurate prediction.

The second day in each month is when IBM SCRT billing occurs, therefore the starting day of the forecasted data that is input into the optimization stream is also the second day in each month. The stream prolongs the starting day to the second day in the next month if the actual beginning of the forecasted day is later than second day in the current month. The second day in the current month is used by the stream if the actual beginning of the forecasted day is the first or second day in the current month. Any forecasted data earlier than the starting day is ignored.

Capping strategy

During the optimization of the arrangement for one product to an LPAR, the usage of MSU of the product should not exceed the cap value of the LPAR. The cap value is calculated from the defined capacity of the LPAR and the capacity that is limited by the shared weight of the LPAR. It is assumed that these two values are near to each other in most cases, and the capping value would be the minimum one of these two values. However, if the defined capacity is significantly bigger than the other (for example, five percent for the difference part), there is a parameter to control the capping value calculation strategy. There are two choices:

- "weight": means to use the capacity that is limited by the shared weight as the capacity capping of the LPAR, but this selection might waste available capability that the LPAR has from the defined capacity and, therefore, decreases the possibility of an optimized result.

- "define": means to use the define capacity as the capacity capping of the LPAR, but this selection might cause certain products to not get the capability expected from the define capacity when CEC is busy and the product's performance will be degraded.

The parameter "PUSERSELECTION" is used to specify the choice.

The cared product

By default, the stream is trying to reduce the billable MSU for all qualified products. You can also designate important products as cared products to notify the stream to lower the cost of the products in priority. This is because there could be a case for an optimization result where the total cost of the summed billable MSU of all the products is reduced with no cost reduction or increase for some of its products. Therefore, this feature can identify these products to guarantee cost reduction for cared products with no guarantee for cost reduction for a non-cared product. You can specify a cared product by inputting the product ID and set the cared flag in the input file.

PROD_ID	CARED
5635-A03	1
5655-R36	1
5655-S97	1
5605-DB2	1

The bounded product

In some cases, products can be related and should appear together in the LPARs arranged for them. For example, product A relies on product B, therefore, in each LPAR where product A is located, product B must exist. The optimization stream supports this scenario by using a bounding feature. For each product group that contains products that are related, you can input them in the input file and allocate a bound ID for them. The file is read during optimization and the products are located together after the arrangement of the optimization.

PROD_ID	BOUND_ID
5635-A03	1
5655-R36	1
5655-S97	2
5605-DB2	2

The reserved product

In some cases, you might want certain products to be arranged in specific LPARs after optimization. The optimization stream supports this scenario by using a reserving feature. For each product expected to be arranged in each LPAR, in the input file you can input the product ID, set the reserved flag, and specify the reserved amount. The new distribution of the products is then contained in the product-LPAR combinations that are specified in the reserved information in the input file. Also, the MSU amount of the products in a reserved LPAR is the amount that is specified in the reserved information.

PROD_ID	LPAR_NAME	CPC_SERIAL_NO	RESERVED	RESERVED_MSU
5635-A03	LP1	AAAA	1	10
5655-R36	LP2	BBBB	1	10

Logging and error handling

Whether the optimization is successful or not, a log is provided to view the running status. The log information is found in the file that is named `prod_bill_msu_optimization_xxx.log`, where *xxx* is the timestamp. The location of the log is specified by the LOGPATH parameter. If the optimization is completed successfully, the final log shows an entry similar to:

```
2014-07-11 19:28:31, "Stream execution completed successfully."
```

To provide better product performance, in most cases, the optimization doesn't include any products that have an MSU that is above the defined capacity of the LPAR or the capacity that is limited by the shared weight of the LPAR. This is to avoid delays in processing.

However, before optimization, it is possible that the MSU usage of any product is already above the defined capacity of the LPAR or the capacity that is limited by the shared weight of the LPAR, so there might be a case that not all products can be arranged when the optimization is performed. Currently, the stream tries to increase the capacity of the LPAR and host the rest of the products in the LPAR. If the attempt fails, the stream produces a log entry and no result is output in the table.

A log entry for increasing capacity successfully is similar to:

```
2014-07-11 17:20:40, "CMA-MW012: The total capacity of the LPARs is not
enough for Product ID 5655-Y04, the optimization is now trying to increase
the capacity of certain LPAR... "

2014-07-11 17:20:57,"CMA-MW009: The total capacity of the LPAR: JF0 is
increased 17.86% percent to host the remaining of the Product ID 5655-Y04."
```

In this case, the optimization is still considered as being successful, but certain LPARs mentioned in the log might have an MSU above the define capacity or weight limited capacity in peak hour. The percentage of the exceeded part is also given in the log.

A log entry for increasing capacity unsuccessfully is similar to:

```
2014-07-09 04:35:07, "CMA-MW012: The total capacity of the LPARs is not
enough for Product ID 5655-Y04, the optimization is now trying to increase
the capacity of certain LPAR..."

2014-07-09 04:35:23, "CMA-ME013: The total capacity of the LPARs is not
enough. No optimization suggestion could be given and the optimization
related reports could not be generated."
```

In this case, the optimization is unsuccessful.

For more information about the error messages, see Appendix F, “Error, warning, and installation messages,” on page 315.

Input parameters

There are two types of inputs for the optimization stream:

- Input file: provides the stream the product information that is related to features used for optimization. Its name and location is `CMAINSTANCE/data/opt_care_prod.csv`, `CMAINSTANCE/data/opt_reserved_prod.csv`, and `CMAINSTANCE/data/data/opt_bound_prod.csv`. You can change the name of the file but the tab name and its format in each tab must be kept.
  - **opt_care_prod.csv**: if there is no requirement for setting products as cared, do not input anything in this tab.

    PROD_ID: the IBM product ID of the cared product.

    CARED: the flag that indicates that the product is to be cared.

    Input any product ID that needs to lower cost in priority in **PROD_ID** field and set the **CARED** field as 1. If **PROD_ID** is entered but the corresponding **CARED** field is blank or is set to `0`, the product is considered to be not cared.
  - **opt_bound_prod.csv**: if there is no requirement for setting products as bounded, do not input anything in this tab.

    PROD_ID: the IBM product ID of the product that needs to be bounded in a group.

    BOUND_ID: the unique number that is allocated to the group.

If two or more products need to be bounded, input the ID in the **PROD_ID** field and allocate a number for them. For each bounded group, the number should be the same for each product in the group.

- **opt_reserved_prod.csv**: if there is no requirement for reserving any products, do not input anything in this tab. This tab is used to input any product that needs to be reserved in any LPAR.

  PROD_ID: the IBM product ID of the reserved product.

  LPAR_NAME: the name of the LPAR that is reserved for the product.

  CPC_SERIAL_NO: the CPC number of the LPAR.

  RESERVED: the reserved flag.

  RESERVED_MSU: the reserved amount by MSU.

  If **RESERVED** is set as 1 and **RESERVED_MSU** is X MSU, the product is arranged in the LPAR with an amount equal or above X.

  If **RESERVED** is 1 and **RESERVED_MSU** is `0`, the product is arranged in the LPAR with an amount above `0`.

  If **RESERVED** is 1 and **RESERVED_MSU** is X MSU, but X is beyond the MSU the product can produce, then the product is kept in its original LPARs location after optimization is completed.

  If the product is reserved in an LPAR that cannot host it (for example, no capacity for the LPAR), then the reserve setting is ignored. However, the stream looks for any LPAR that is suitable for the product.

Runtime parameters

- If you are using the scripts, you can specify these runtime parameters in the `cmasca.parms` file that is located in the `CMAINSTANCE/data` folder:
  - `PUSERSELECTION`: the capping calculation option, set from two strategies that are used to apply capping: “`weight`” or “`define`”. The default setting is "weight".
  - `PCAREFILE`: the location of the `opt_care_prod.csv` input file.
  - `PRESERVEDFILE`: the location of the `opt_reserved_prod.csv` input file.
  - `PBOUNDFILE`: the location of the `opt_bound_prod.csv` input file.
  - `PODD`: obsolete detection days.

  Note: The `opt_care_prod.csv`, `opt_reserved_prod.csv`, or `opt_bound_prod.csv` file names can be changed if required.
- if you are using Modeler Client, you can specify the same runtime parameters as in a shell:

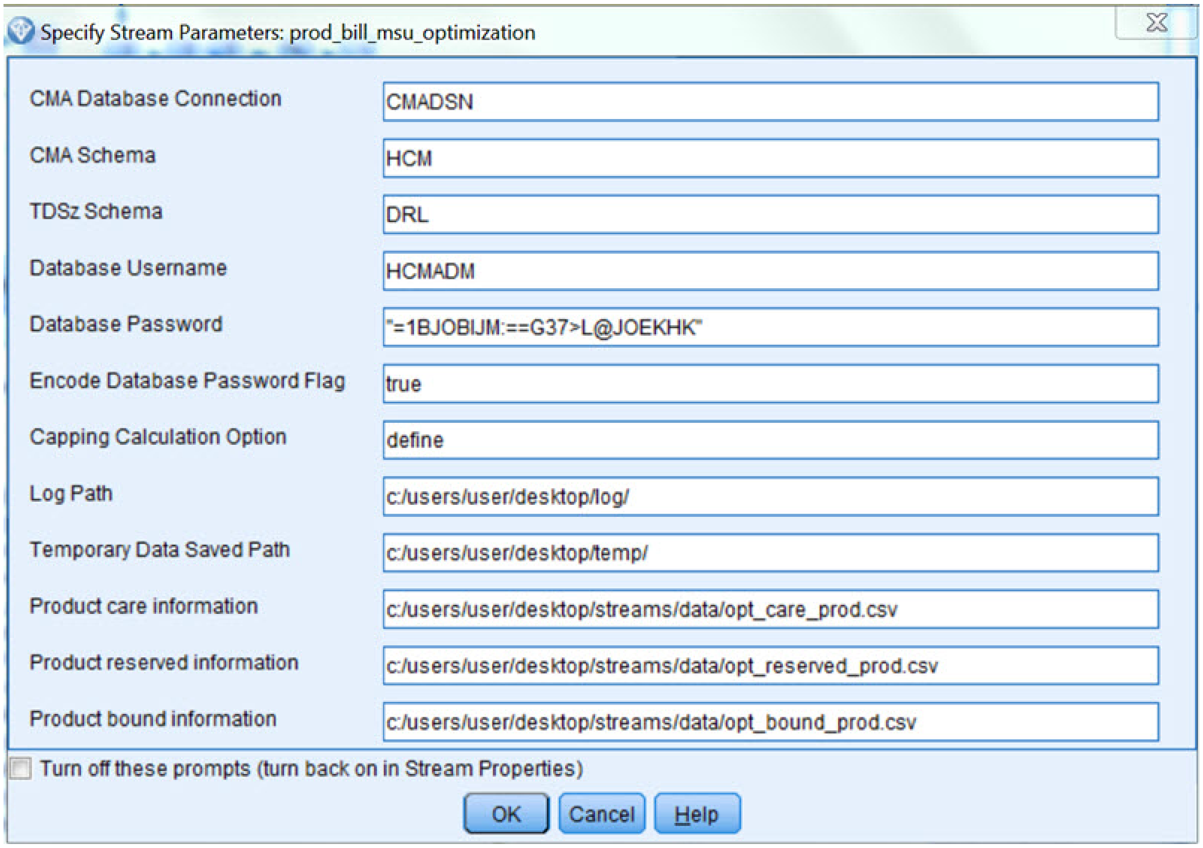

Moving suggestion

After the optimization stream is executed successfully, the moving suggestion of each product in the optimization can be viewed in this report: SCA Workspace, SCA LPAR MSU Utilization & SCA Product MSU Utilization.

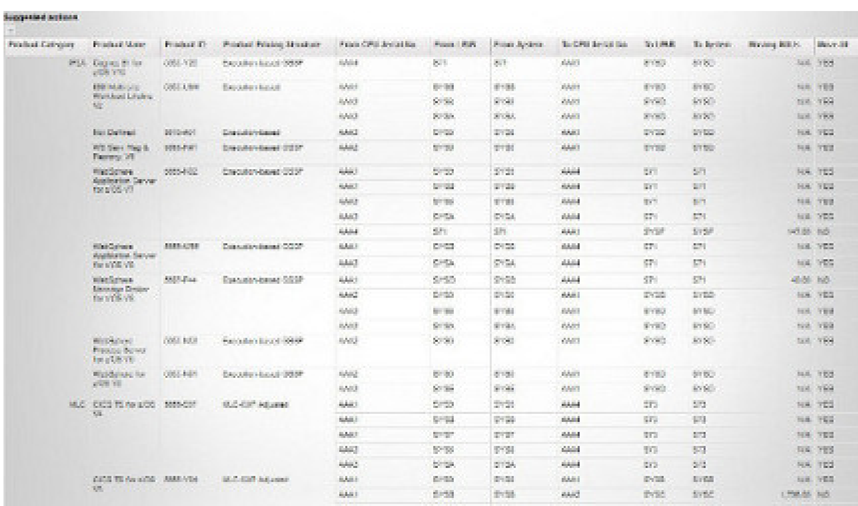

Billable MSU

Once the optimization stream is executed successfully, the billable MSU of each product after the optimization can be viewed this report:

Product Summary

Product Category	Product Name	Product ID	All Machines: Optimized Total Billable MSUs	All Machines: Forecasted Total Billable MSUs	All Machines: Total Difference
IPLA	CICS VSAM Recovery	5655-P30	1,933	8,457	-6,524
	CICS VSAM Recovery V3	5655-H91	1,933	8,457	-6,524
	Cognos BI for z/OS V10	5655-Y26	711		
	IBM Multi-site Workload Lifeline V2	5655-UM4	711	4,882	-4,171
	IBM Security zSecure Admin	5655-N16	5,692	8,457	-2,765
	Not Defined	5610-A01	711	1,510	-799
		5655-B17	5,692	8,457	-2,765
		5655-JAV	5,692	8,457	-2,765
		5655-M23	5,692	8,457	-2,765
		5695-045	5,692	8,457	-2,765
		5697-F51	5,692	8,457	-2,765
		5698-AAB	4,452	7,413	-2,961
		5698-ARA	5,163	7,816	-2,653
		5748-XX9	5,692	8,457	-2,765
	WebSphere Application Server for z/OS V7	5655-N02	301	2,283	-1,982
	WebSphere Application Server for z/OS V8	5655-W65		20	
	WebSphere for z/OS V6	5655-N01	1	1	0
	WebSphere Message Broker for z/OS V6	5697-P44	32	111	-79
	WebSphere Process Server for z/OS V6	5655-N53	1		
	WS Serv Reg & Repstory, V6	5655-R41	1		
	zSecure Admin	5655-T01	5,692	8,457	-2,765
	zSecure Audit	5655-T02	5,692	8,457	-2,765
MLC	CICS TS for z/OS V4	5655-S97		7,126	
	CICS TS for z/OS V5	5655-Y04	1,933	7,987	-6,054
	DB2 10 for z/OS	5605-DB2		8,457	

Application average forecast stream

The `appl_avg_hourly_forecast_timeseries.str` and `appl_avg_daily_forecast_timeseries.str` streams are used to predict the application future values of CP_SECONDS, IFA_ SECONDS, and IIP_ SECONDS for hourly and daily levels. The Application Analytics reports use the forecast data that is provided by these streams to calculate MIPS values.

Forecasting approach

The first step in creating a forecast is to retrieve the data from the database and aggregate it up to the appropriate level. In the hourly forecast stream, the data is aggregated by function key, environment key, APPL key, CPC, MVS_SYSTEM_ID, DATE, HOUR to hourly level to sum the MEASURED_SEC, CP_SECONDS, IIP_SECONDS, and IFA_SECONDS, which are marked as MEASURED_SEC_Sum, CP_SECONDS_Sum, IIP_SECONDS_Sum, and IFA_SECONDS_Sum accordingly. Then, the stream uses time series to forecast CP_SECONDS_Sum, IIP_SECONDS_Sum and IFA_SECONDS_Sum. The daily forecast stream aggregates the data to the daily level.

Prompt parameters

- CMDW: the CMA database connection.
- CMASCHEMA: the schema for CMA.
- UID: the database user.
- PWD: the database user's password.
- OUTPUTMODELPATH: model output path.
- LOGPATH: the log path.
- INPUTHOLIDAYFILE: the input holiday file.
- TMPPATH: temporary data saved path.
- DATE_START: the start date for the data that is used.
- DATE_END: the end date for the data that is used.
- HOUR_START: the start hour for the data that is used.

- HOUR_END: the end hour for the data that is used.

If you do not specify all of the parameters or if you specify incorrect parameters, the stream stops and an error is written to the logs specified in LOGPATH. If you have not specifies the LOGPATH or the path is no valid, then the stream stops and a message is displayed. The stream verifies the database connection and the values for the INPUTHOLIDAYFILE and TMPPATH parameters before the stream is run into the main branch.

For the special parameters DATE_START, DATE_END, HOUR_START, and HOUR_END, the stream verifies that DATE_START < DATE_END, HOUR_START <= HOUR_END, and HOUR_START and HOUR_END are in range[0,23]. If the values that you specified for parameters DATE_START and DATE_END are out of range between the actual history start date and the history end date, then the stream forecasts using the actual history start or actual history end date and a warning message shows the date range that is used by the stream.

Source table

CMASCHEMA.APPL_MIPS_UTIL_ZOS_VIEW

Export table

CMASCHEMA.APPL_ZOS_FORECAST

CMASCHEMA.FORECAST_METADATA

These two streams export forecast seconds data to the CMASCHEMA.APPL_ZOS_FORECAST table with FORECAST_LEVEL= "avg_hourly" or FORECAST_LEVEL="avg_daily".

These two streams each insert one record of forecast stream metadata properties into the CMASCHEMA.FORECAST_METADATA table. For hourly forecast stream, STREAM_TYPE="appl_sec_avg_hourly" and AGG_METHOD="avg_hourly", and for the daily forecast stream, STREAM_TYPE="appl_sec_avg_daily" and FORECAST_LEVEL="avg_daily".

When the streams insert new data into the tables, the data from the last time the stream was run for the FORECAST_LEVEL is deleted.

Time interval

For the hourly forecast stream, the time interval nodes on “Number of hours in a day” and “Day begins at” can be automatically set so that they are consistent with the vales that are specified in the HOUR_START and HOUR_END parameters by using scripts. The number of hours in a day = HOUR_END - HOUR_START + 1, and Day begins at = HOUR_START.

For the daily forecast stream, because we forecast seconds in daily level, we use the observed data from HOUR_START to HOUR_END to summarize to daily level and then predict the daily seconds. We do not have to specify the time interval by using scripts.

The hourly stream forecast seconds for next 60 days. If the specified time range is 0 to 23 for HOUR_START and HOUR_END, then the stream predicts for a future 24 multiplied by 60 records, which is equal to 1440 records. The stream adjusts the predicted records number by using scripts automatically. The daily forecast stream predicts seconds for the next 1 year, which is always 365 records.

Inputs for time series

The holiday file that you specified in the INPUTHOLIDAYFILE parameter. For these two streams, the holiday date file is the only default input for the time series.

Log file and error handling

You can find the IBM SPSS Modeler stream log files in the log folder that you specify in the LOGPATH parameter. The log file is named `appl_avg_hourly_forecast_timeseries_<timestamp>.log` or `appl_avg_daily_forecast_timeseries_<timestamp>.log`.

For information about error and warning messages, see Appendix F, "Error, warning, and installation messages," on page 315.

Application peak value forecast stream

The application peak value forecast stream (`appl_peak_daily_forecast_timeseries.str`) analyzes a defined time range of the history data for CPU seconds, calculates its peak value of MIPS, and forecast its one year peak value at an application level for the future.

CPU seconds and MIPS

The CPU seconds refer to three types of processor seconds: CP_SECONDS, IFA_ SECONDS, and IIP_ SECONDS. MIPS is calculated based on the CPU seconds, so there are three types of MIPS derived from three types of the CPU seconds.

Application level

The application level is defined by the field APPL_NAME. A different value in the field results in a different application level.

Peak values

The peak values are defined as a serial formed by the Top-N MIPS values in each day, where N is specified as the Peak Highest Rank within the range [1,24]. You can also input multiple N values to forecast multiple peak value serials in a single run. The maximum value of N is 5.

Input table

- APPL_MIPS_UTIL_ZOS_VIEW: the source of history data for CPU seconds in each measurement interval.
- MIPS_CAPACITY: this table is used to look up the MIPS_TOTAL_CAPACITY to calculate the MIPS.
- APPL_MAPPING: this table is used to derive the value of MAPPING_TIME.

Output table

- APPL_ZOS_FORECAST: the forecasted data and the related history data for MIPS of each day by each application level for each Peak Highest Rank is stored in this table.
- FORECAST_METADATA: the execution information of peak forecasting stream for each Peak Highest Rank is stored in this table.

Input parameters

- CMDW: the CMA database connection.
- UID: the database user.
- PWD: the database user's password.
- CMASCHEMA: the schema for CMA.
- OUTPUTMODELPATH: the model output path
- LOGPATH: the logging path.
- INPUTHOLIDAYFILE: the input holiday file.
- TMPPATH: temporary data saved path.
- DATE_START: the start date for the data that is used.
- DATE_END: the end date for the data that is used.
- HOUR_START: the start hour for the data that is used.

- HOUR_END: the end hour for the data that is used.
- PEAKRANK: the Nth highest peak rank values. Valid values are in the range [1,5]. You can also input multiple rank values.

  Note: For the DATE_END and HOUR_END parameter, the inputted value must be greater than the DATE_START and HOUR_START values, respectively.

  For PEAKRANK, use a comma (,) to separate rank values if you are using multiple rank values.

Validating inputs

The stream verifies that input values are valid. For example, the stream verifies whether the path and database are writable. If not, it stops and an error is written to the log file. The log file is stored in LOGPATH.

Forecast data generation and output

There is a Python script in the peak forecast stream that drives the process. The stream finds the peak values for each type of MIPS by each Peak Highest Rank N by each application level within each day. The peak MIPS values are selected from the MIPS values by each hour within a day. After a serial of peak values is produced, the stream analyzes the data by using the Time Serials model, outputs and stores each of the models that are trained in OUTPUTMODELPATH, and uses them to forecast one year of peak values in the future for each type of MIPS by each application level. For each related Peak Highest Rank N of the forecasted data, peak_n is labeled in the field FORECAST_LEVEL. The saved models are named `APPL_NAME_<timestamp>.gm`.

After the forecast values are produced, they are gathered in a temporary file that is stored in TMPPATH and output to the database. If the database already contains MIPS forecasted data with the same “peak_n”, the new values overwrite the existing values. Values that do not have the same “peak_n” as the new values are kept in the database.

Execution data generation and output

The information of the latest stream execution time and the data time range for history data and forecast data along with related Peak Highest Rank “peak_n” (represented by filed AGGREGATION in the table) are gathered when the forecasting runs. This information is written to the database after the forecasting finishes. Each running for each Peak Highest Rank makes one record in the table. If there are old records in the database, values with the same Peak Highest Rank are overwritten if new values are produced. Any existing values that do not have the same Peak Highest Rank are kept in the database.

Log file and error handling

You can find the IBM SPSS Modeler stream log files in the log folder that you specify in the LOGPATH parameter. The log file is named `appl_peak_daily_forecast_timeseries_<<timestamp>.log`.

For information about error and warning messages, see Appendix F, “Error, warning, and installation messages,” on page 315.

Mapping table definitions for application analytics reporting

The application analytics reports categorize the application information by using mapping tables. The application mapping is defined in two spreadsheet files that are provided with the Capacity Management Analytics Solution Kit installation. one file (`appl_lob_z.csv`) is used for z/OS systems. The other (`appl_lob.csv`) is used for distributed systems.

The `appl_lob_z.csv` and `appl_lob.csv` files are in the *CMAINSTANCE*`/data` directory.

You must modify the files to suit your environment before you run the application analytics streams.

Application mapping for z/OS systems

The `appl_lob_z.csv` file is used for z/OS systems and allows you to map workloads based on their report class and job name into three categories that describe the application. The categories are application, environment, and function.

The input is determined by either the REPORT_CLASS or the JOB_NAME from the SMF records. The user-defined values for APPL_NAME and FUNCTION determine a hierarchy where the APPL_NAME is the high-level name and FUNCTION is a sub-category of that application. The ENVIRONMENT value allows you to separate out different environments for the same application with multiple functions; for example, Prod, QA, Test, Devl, or whatever you define.

The mapping file uses the following columns for classification:

APPL_NAME

The name of the application or collection of applications. This can be a single application or a collection of applications, depending on how you want to roll up the applications. This functional grouping allows you to subset large applications into smaller logical groups that map to business or application functions.

ENVIRONMENT

For example, use Prod for production, Test or QA for quality assurance, Devl for development, or Sand for sandbox environments. You can use any environment classification that supports your separation of the workload of a single application.

FUNCTION

This is a second-level categorization for an application that allows you to have a group of applications with a single name but still have separate sub-application names in the FUNCTION field.

REPORT_CLASS

The report class definition that is used in IBM Workload Manager for z/OS (WLM). You can define the report class for a function of an application. The default value is *.

If you used * in the REPORT_CLASS column of the mapping table, use a specific value in the JOB_NAME column.

Note: The report class is specified in SMF 72 records, and the job name is specified in SMF 30 records. Do not use both a REPORT_CLASS and a JOB_NAME value for the same entry.

JOB_NAME

The detail address space name. The default value is *.

If you used * in the JOB_NAME column of the mapping table, use a specific value in the REPORT_CLASS column.

Tip: You can use wild-card characters for the JOB_NAME column. For example, you can enter values such as ABC*, ABC%, or AB%C. * represents any characters, whereas % represents only a single character. You cannot use both * and % in the same JOB_NAME value, and they cannot be used as the first character in a JOB_NAME value. * should be used only at the end of a string of characters, rather than in the middle.

If you want to use multiple % characters in a string, you must use them only in adjacent positions in the string. You cannot separate % characters within the string. For example, you can use A%%D. But, you cannot use A%C%.

If you specify a JOB_NAME of ABC* or ABC% and you also specify a JOB_NAME of ABCD, that is rolled up into the same FUNCTION or APPL_NAME, it results in ABCD being counted twice.

TRANS_CAPTURE_RATIO

The ratio of the CPU time that is accounted-for in the SMF 30 records for the CICS address space to the total processor time that is measured in the SMF 110 records for the transactions that map to the REPORT_CLASS.

If a report class is tracking a subset of transactions within a CICS region, both the SMF 30 records for that CICS region and the SMF 110 records are used to identify the transactions that use that report class. Capacity Management Analytics develops the TRANS_CAPTURE_RATIO and the Trans_Rate value by using a separate utility. Alternatively, you can provide the values in the mapping file.

In this case, it is possible to see both REPORT_CLASS and JOB_NAME values used together. However, it applies only to CICS or IMS applications where you want to breakout some transactions into separate applications or functions.

The default value for TRANS_CAPTURE_RATIO is 1.0.

PROGRAM_TYPE
The report class or address space program type. For example, CICS, IMS, WAS, DB2. The default is blank.

TRANS_RATE
CICS or IMS transaction counts per CPU second. The default is 0.

FLAG
Indicates that record is to be added, updated, or deleted from the database. The values can be ADD, UPDATE, or DELETE.

Application mapping for distributed systems

For distributed systems, a similar mapping is done, except that you can use only server names as input. A server is assumed to be running only one application or function. However, many servers can map to the same application or function.

Use the `appl_lob.csv` mapping file to provide the server names, and to specify the values for the application name, environment, and function.

APPL_NAME

The name of the application or collection of applications. This can be a single application or a collection of applications, depending on how you want to roll up the applications. This functional grouping allows you to subset large applications into smaller logical groups that map to business or application functions.

ENVIRONMENT
For example, use Prod for production, Test or QA for quality assurance, Devl for development, or Sand for sandbox environments. You can use any environment classification that supports your separation of the workload of a single application.

FUNCTION
This is a second-level categorization for an application that allows you to have a group of applications with a single name but still have separate sub-application names in the FUNCTION field.

SERVER
The name of the server on which the application is running.

Chapter 13. Framework Manager model reference

The IBM Cognos Framework Manager models used in IBM Capacity Management Analytics consists of several namespace/folder hierarchies that group tables and objects into a logically identifiable way within IBM Cognos Framework Manager.

There are four data source connections available in the Capacity Management Analytics model:

- HCM—References the Capacity Management Analytics data store tables. The default schema is HCM.
- DRL—References the Tivoli Decision Support for z/OS (TDSz) data store tables. The default schema is DRL.
- DRLSYS—References the Tivoli Decision Support for z/OS metadata tables. The default schema is DRLSYS.

There are two packages available in the Capacity Management Analytics model:

- Capacity Management Analytics—The Capacity Management Analytics package consists of only the tables, filters, and other components that are used in the final reports.
- Capacity Management Analytics Development—The Capacity Management Analytics Development package provides full access to all the tables and columns used to develop reports. When you verify the Development package, you may see some warning messages. Framework Manager repairs the errors automatically, so no action is required.

The Capacity Management Analytics model is segmented into four layers: Physical View, Logical View, Presentation View, and Dimensional View. The four layers are built upon one another, with the Physical View being the foundation of the model.

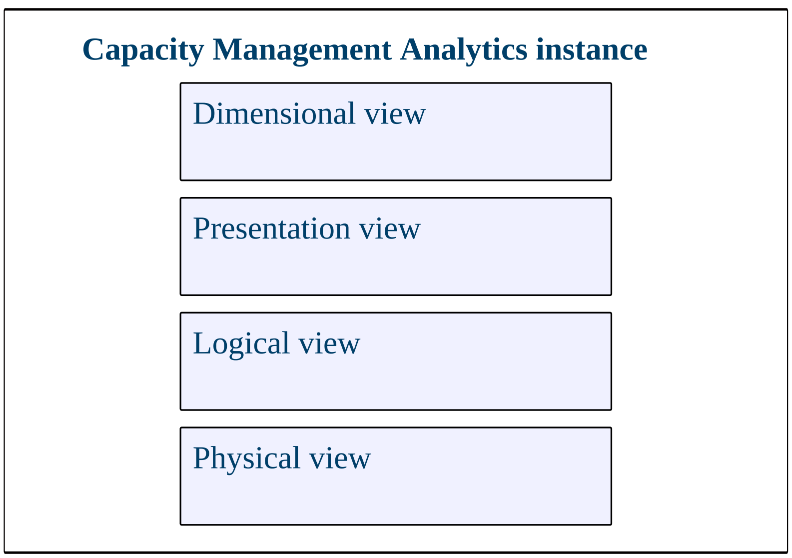

Figure 19: Layers within the Capacity Management Analytics model

The Capacity Management Analytics Instance is the high-level container for Capacity Management Analytics which holds the four model layers.

Capacity Management Analytics Instance

The IBM Capacity Management Analytics Instance is the high-level container for Capacity Management Analytics which holds the four model layers.

Dimensional view

The Dimensional View contains regular dimensions, measure dimensions, and scope relationships that are created in Framework Manager.

The dimensional reporting style is recommended for dimensionally-modeled relational (DMR) and Online Analytical Processing (OLAP) data sources. Dimensional data is best represented by crosstabs, maps, and charts. This data is shown in IBM Cognos Report Studio in dimensions, hierarchies, levels, and members.

In dimensional reporting, you summarize data by using member summaries and within set aggregation. You focus data in dimensional reporting by adding only the relevant members to the edge of a crosstab or to the context filter. You can also enable drilling up and drilling down in dimensional reports.

Table 21: Shortcuts in the Dimensional View

Shortcut Name	Description	Target Object Reference
Linux for System Z DIST Date	This table provides yearly, monthly, weekly and daily dimension from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[DIST_DATE]`
Linux for System Z DIST Hour	This table provides hourly dimension from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[DIST_HOUR]`
Linux for System Z DIST Time	This table provides minutely dimension from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[DIST_Z_TIME]`
Linux for System Z Image	This table provides image dimension from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[Linux for System Z IMAGE]`
Linux for System Z CPU Fact	This table provides CPU measure from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[Linux for System Z CPU]`
Linux for System Z Memory FACT	This table provides Memory measure from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records.	`[Dimensional view].[Linux for System Z Memory]`
Linux for System Z CPU Hour FACT	This table provides hourly statistics on CPU measure from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Dimensional view].[Linux for System Z CPU Hour]`
Linux for System Z Image Name	This table provides image metrics from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Logical View].[Z_IMAGE]`
Linux for System Z CPU Raw Metrics	This table provides CPU metrics from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Logical View].[Z_OPERATING_SYS_T_CPU]`
Linux for System Z Memory Raw Metrics	This table provides Memory metrics from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Logical View].[Z_OPERATING_SYS_T_MEM]`
Linux for System Z CPU Hour Metrics	This table provides hourly statistics on CPU metrics from the ZLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 43.	`[Logical View].[Z_OPERATING_SYS_H_CPU]`
Linux for System X DIST Date	This table provides yearly, monthly, weekly and daily dimension from the XLINUX_IMAGE resource used by RMF XP for Linux on System z. Data is taken from SMF 104 records subtype 23.	`[Dimensional view].[DIST_DATE]`

Presentation View

The Presentation View includes shortcuts to the existing items in the Logical View.

The shortcuts available in the Presentation View are described in the following table. The Target Object References are shown with line breaks for publication formatting purposes only.

Table 22: Shortcuts in the Presentation View

Shortcut Name	Description	Target Object Reference
Channel Path Activity Metrics	This table provides statistics on channel path activity on MVS systems. It contains data from SMF type 73 records. This table is updated by Tivoli Decision Support for z/OS MVSPM_CHANNEL_H table	`[Logical View].[HCM_MVSPM_CHANNEL]`
CSA/ECSA/SQA/ESQA Usage Metrics	This table provides hourly statistics on the usage of virtual storage in the CSA and SQA. It contains data from SMF type 78, subtype 2 records. For more information, see MVSPM_VS_CSASQA_H in the Tivoli Decision Support for z/OSdocumentation (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2/com.ibm.tivoli.dszos.doc.1.8.2/SysRef1/DRL5FT16.pdf).	`[Logical View].[MVSPM_VS_CSASQA_H]`
Customer Pricing Information	This table contains customer-provided pricing information for the all running products.	`[Logical View].[CUSTOMER_PRICE]`
Device Activity Metrics	This table provides statistics on the I/O activity of MVS devices. The data comes from SMF type 74, subtype 1 records. This table is updated by Tivoli Decision Support for z/OS MVSPM_DEVICE_H table.	`[Logical View].[HCM_MVSPM_DEVICE]`
LPAR Level Forecast Metrics	This table provides forecast results. All the forecast results come from three forecast streams: hourly, daily and monthly. For details of this table, see “LPAR Level Forecast Metrics table” on page 161.	`[Logical View].[MVSPM_LPAR_CPU_BUSY_MIPS_PRED]`
LPAR MSU Forecast Metrics	This table contains hourly MSU utilization forecasts for each LPAR.	`[Logical View].[LPAR_MSU_FORECAST]`
LPAR Level Metrics	This table is a consolidated version of DRL.MVSPM_LPAR_H, which eliminates any duplicate entries from DRL.MVSPM_LPAR_H. For details of this table, see “LPAR Level Metrics table” on page 162.	`[Logical View].[HCM_MVSPM_LPAR]`
LPAR MSU Optimization Metrics	This table contains LPAR MSUs based on the optimized product placement.	`[Logical View].[LPAR_MSU_OPTIMIZATION]`
LPAR MSU Metrics	This view returns LPAR MSU metrics between specific dates.	`[Logical View].[MVSPM_LPAR_MSU_VIEW]`
MIPS Capacity Ratings	This table provides the CPU_SERIAL_MIPS rating for your processors. For details of this table, see “MIPS Capacity Ratings table” on page 169.	`[Logical View].[MIPS_CAPACITY]`
Movement Suggestion by Optimization Metrics	This table contains suggested product placement after an optimization.	`[Logical View].[OPTIMIZATION_MOVEMENT]`
MVS Paging Activity Metrics	This table provides statistics on paging activity on MVS systems. It contains data from SMF	`[Logical View].[MVSPM_PAGING_H]`

Channel Path Activity Metrics table

This table provides statistics on channel path activity on MVS systems. It contains data from SMF type 73 records. This table is updated by the Tivoli Decision Support for z/OS MVSPM_CHANNEL_H table Channel Path Activity Metrics table

Table 23:

Column	Description
DATE	Date of the mid measurement. From SMF73DAT, SMF73IST, and SMF73INT.
TIME	Time (rounded down as specified in the MVSPM_TIME_RES lookup table) of the mid measurement. From SMF73DAT, SMF73IST, and SMF73INT.
PERIOD_NAME	Name of the period. This is derived using fields SMF73DAT, SMF73IST, and SMF73INT from the record as parameters in the PERIOD function.
MVS_SYSTEM_ID	MVS system ID. This is the SMF system ID. From SMF73SID.
CHANNEL_PATH_ID	Channel path ID. From SMF73PID.
CHANNEL_TYPE	Type of channel. From SMF73ACR.
CHANNEL_FLAGS	Channel flag bits. From SMF73FG3 and SMF73FG4. The following list shows bits and meaning when set. • 0 - ES connection channel • 1 - ES connection director attached to the channel path • 2 - ES connection converter attached to this channel • 3 - Channel path modified • 4 - Channel path deleted • 5 - Channel path added • 6 - Valid channel path • 7 - Channel path is shared between logical partitions • 8 - CPMB entry is not valid • 9 - Channel path is channel-to-channel
RECORD_LEVEL	SMF record level change number. This column enables processing of SMF record level changes in an existing release. From SMF73SRL.
CHANNEL_BUSY_PCT	Average percentage of time that the channel path was busy. Calculated as the average of SMF73BSY*100/SMF73SMP. Please note this value may exceed 100% due to data from micro code.

Column	Description
PART_CHAN_BUSY_PCT	Average percentage of time that the channel path was busy for the partition. Calculated as the average of SMF73PBY*100/SMF73PTI. Please note this value may exceed 100% due to data from micro code.
CHANNEL_BUSY_MAX	Maximum percentage of time that the channel path was busy during an RMF interval. Calculated as the maximum of SMF73BSY*100/SMF73SMP. Please note this value may exceed 100% due to data from micro code.
PART_CHAN_BUSY_MAX	Maximum percentage of time that the channel path was busy for the partition. Calculated as the maximum of SMF73PBY*100/SMF73PTI. Please note this value may exceed 100% due to data from micro code.
TIME_RESOLUTION	Time resolution used for this table, in minutes. The default is 60 minutes. From TIME_RESOLUTION in the MVSPM_TIME_RES lookup table.
MEASURED_SEC	Measured time period, in seconds. This is the duration of the RMF measurement intervals. Calculated as the sum of SMF73INT/1000.
SAMPLES	Number of store channel path status (STCP) samples taken by SRM. This is the sum of SMF73SMP.
SAMPLES_BUSY	Number of store channel path status (STCP) samples taken by SRM in which the channel path related to this entry was found busy. This is the sum of SMF73BSY.
CPMF_BUSY_SEC	Logical partitions channel path busy time as measured by the channel path measurement facility (CPMF), in seconds. Calculated as the sum of SMF73PBY*0.001024.
CPMF_INTERVAL_SEC	Sum of the interval times as measured by the channel path measurement facility (CPMF), in seconds. Calculated as the sum of SMF73PTI*0.001024.
CPMF_GROUP	CPMF channel measurement group. From SMF73CMG
CHANNEL_BUSY_TOT	Total channel path busy time in seconds when the measurement group number is 1. Calculated as the sum of SMF73TUT/7.8125.
PART_CHAN_BUSY_TOT	LPAR channel path busy time in seconds when the measurement group number is 1. Calculated as the sum of SMF73PUT/7.8125.
BUSCYCLE_MAX_SEC	Maximum bus cycles per second when the measurement group number is 2. Calculated as the maximum value of SMF73MBC.

Table 23: (continued)	
Column	**Description**
CHAN_UNIT_MAX_SEC	Maximum channel work units per second when the measurement group number is 2. Calculated as the maximum value of SMF73MCU.
WRITE_MB_MAX_SEC	Maximum write data units per second when the measurement group number is 2. Calculated as the maximum value of SMF73MWU.
READ_MB_MAX_SEC	Maximum read data units per second when the measurement group number is 2. Calculated as the maximum value of SMF73MRU.
BUSCYCLE	Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit. Calculated as the average of (SMF73TBC*100.)/(SMF73MBC*SMF73PTI*1024E-6). This value may exceed 100% due to data from micro code.
CHAN_UNIT_TOT	Total channel work unit count when the measurement group Number is 2. Calculated as the sum of SMF73TUC.
LPAR_CHAN_UNIT_TOT	LPAR channel work unit count when the measurement group number is 2. Calculated as the sum of SMF73PUC.
CHAN_WRITE_MB_TOT	Channel write in MB when the measurement group number is 2. Calculated as the sum of SMF73TWU*(SMF73US/1024E3).
LPAR_WRITE_MB_TOT	LPAR write in MB when the measurement group number is 2. Calculated as the sum of SMF73PWU*(SMF73US/1024E6).
CHAN_READ_MB_TOT	Channel read in MB when the measurement group number is 2. Calculated as the sum of SMF73TRU*(SMF73US/1024E3).
LPAR_READ_MB_TOT	LPAR read in MB when the measurement group number is 2. Calculated as the sum of SMF73PRU*(SMF73US/1024E3).
SHARED_IND	Channel Path is shared between LPAR. This is Y (yes) if bit 8 of SMF73FG3 is on; otherwise, N (no).
MULTI_SUBSYS_IND	Channel Path is shared between LPAR. This is Y (yes) if bit 8 of SMF73FG3 is on; otherwise, N (no).
CHANNEL_SUBSYS_ID	Channel Subsystem ID. Valid only for hardware that supports Multiple Channel Subsystem. From SMF73CSS.
DATA_TYPE	Rank Flag id. A flag which distinguishes the Top N CHANNEL BUSY MAX record in this table.
LPAR_NAME	Name of the logical partition.
CPU_SERIAL_NO	Last 4 digits of the CPU serial number.

Table 23: (continued)	
Column	**Description**
LPAR_SYSTEM_ID	System ID of the logical partition.

Customer Pricing Information table

The Customer Pricing Information table contains the following columns.

STARTDATE
Start date of the mid measurement. The default is 1900-01-01.

ENDDATE
End date of the mid measurement. The default is 2900-01-01

ENTERPRISE
Logically named enterpise. Enterprise is defined as a grouping of mainframe systems that will all have the same currency.

CPU_SERIAL_NO
Last 4 digits of the CPU serial number. From SMF70SER.

CURRENCY
The three character international symbol for the currency.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

PID
Product ID number.

CATEGORY
Software Licensing category at IBM: MLC or IPLA.

LICENSE
Monthly License Charge metrics name of MLC products or Value Unit Exhibit name of IPLA products.

TIER
Monthly License Charge metrics tier of MLC products or Value Unit Exhibit tier of IPLA products

TIER_MIN_MSU
Minimum MSU of this tier.

TIER_MAX_MSU
Maximum MSU of this tier.

PRICE_PER_MSU
Price per MSU of MLC products. Decimal value that goes up to two decimal places.

C_TIER_PRICE
Interim price calculation result for MLC products for reporting usage.

VU_PER_MSU
VU per MSU.

C_TIER_VU
Interim VU calculation result for IPLA products for reporting usage.

PRICE_PER_VU
Price per VU. Decimal value that goes up to two decimal places.

VU_ENTITLED_MSU
Entitled MSUs for this IPLA product on this CPU Serial No. It is not the entire enterprise entitlement. Adding up each entitlement totals the enterprise entitlement.

SS_PID
Subscription and Support product ID. This value may be left blank if Subscription and Support was not purchased.

SS_VU_ENTITLED_MSU
Entitled MSUs for this IPLA product on this CPU Serial No. It is not the entire enterprise entitlement. This value may be left blank if Subscription and Support was not purchased.

SS_PRICE_PER_MTH
Subscription and Support product price per month. Decimal value that goes up to two decimal places.

Device Activity Metrics table

This table provides statistics on the I/O activity of MVS devices. The data comes from SMF type 74, subtype 1 records. This table is updated by Tivoli Decision Support for z/OS MVSPM_DEVICE_H table. Device Activity Metrics table

Table 24:

Column	Description
DATE	Date of the mid measurement. From SMF74DAT, SMF74IST, and SMF74INT.
TIME	Time of the mid measurement. From SMF74DAT, SMF74IST, and SMF74INT.
DATA_TYPE	Rank Flag id. A flag which distinguishes DASD top record and DASD aggregate record in this table. If the flag equals 0, it represents DASD top record. If the flag equals -1, it represents DASD aggregate record.
LPAR_NAME	Name of the logical partition.
CPU_SERIAL_NO	Last 4 digits of the CPU serial number.
PERIOD_NAME	Name of the period. This is derived using fields SMF74DAT, SMF74IST, and SMF74INT from the record as parameters in the PERIOD function.
MVS_SYSTEM_ID	MVS system ID. From SMF74SID.
LPAR_SYSTEM_ID	System ID of the logical partition.
IO_RATE	IO Rate, MEASUREMENT_EVENTS/RMF_INTERVAL.
IO_INTENSITY	IO Intensity, IO_RATE/Response Time. Response time is sum of average Connection, disconnection, pending and IOSQ.
QI_INTENSITY	Queuing Intensity, IO_RATE/Queue Delay Time. Queue Delay Time is sum of average disconnection, pending and IOSQ.
DEVICE_NUMBER	Device number. From MVS_ADDR in the MVSPM_DEVICE_ADDR lookup table, or from SMF74NUM if no match is found in the lookup table.
CACHE_SUBSYSTEM_ID	3990 control unit subsystem ID. From CACHE_SUBSYSTEM_ID in the MVSPM_DEVICE_ADDR lookup table using SMF74SID and SMF74LCU as lookup argument.
DEVICE_CLASS	Device class code. This is derived using field SMF74TYP from the record as key.

Table 24: (continued)	
Column	**Description**
UNIT_TYPE	Device model name or unit type. From SMF74DEV or from UNIT_TYPE in the MVSPM_UNIT_TYPE lookup table if SMF74DEV is missing.
LCU_NUMBER	Logical control unit number. From SMF74LCU.
VOLSER	Volume serial of the volume mounted on the device (tape or direct access device only). From VOLSER in the MVSPM_DEVICE_ADDR lookup table, or from SMF74SER if VOLSER is missing or no match is found in the lookup table.
STORAGE_GROUP	Storage group name. From SMF74SGN.
CU_NAME	Control unit name. This column is blank if control unit name cannot be determined. From SMF74CU.
RECORD_LEVEL	SMF record level change number. This column enables processing of SMF record level changes in an existing release. From SMF74SRL.
TIME_RESOLUTION	Time resolution used for this table, in minutes. The default is 60 minutes.
RMF_INTERVAL	Duration of the RMF measurement interval, in minutes. Calculated as SUM(SMF74INT/60000).
CYCLE_LENGTH_MSEC	RMF sampling cycle length, in milliseconds. From SMF74CYC.
INPUT_RECORDS	Number of input records collected from the input log.
SAMPLES	Number of RMF samples taken. This is calculated as the sum of SMF74SAM.
SSCH_COUNT	Start subchannel count. This is the number of SSCH instructions. This is the sum of SMF74SSC.
MEASUREMENT_EVENTS	Measurement event count. This is the number of SSCH instructions for which connect, pending, and active times were stored. This is the sum of SMF74MEC.
REQUESTS_QUEUED	Number of requests queued in the IOS for the device. This is the sum of SMF74QUE.
SAMPLES_UTL	Number of samples taken when the device was reserved but an SSCH instruction had not been issued to the device. This is the sum of SMF74UTL.
SAMPLES_RSV	Number of samples taken when the device was reserved. This is the sum of SMF74RSV.
ALLOCATIONS	For RMF 3.5.1 releases and later, this is the total number of allocations in effect for the device. For RMF releases earlier than 3.5.1, this is the total number of data sets open on the device. This is the sum of SMF74NDA and SMF74DSO.

Table 24: (continued)	
Column	**Description**
CONNECT_MSEC	Device connect time, in milliseconds. This is the sum of SMF74CNN*0.128.
PENDING_MSEC	Device pending time, in milliseconds. This is the sum of SMF74PEN*0.128.
ACTIVE_MSEC	Device active time, in milliseconds. This is sum of SMF74ATV*0.128.
DISCONN_MSEC	Device disconnect time, in milliseconds. This is the sum of SMF74DIS*0.128.
DEVICE_BUSY_DELAY	Device busy delay time, in milliseconds. This is from subchannel information block (SCHIB). This is the sum of SMF74DVB*0.128.
CU_BUSY_DELAY	Control unit busy delay time, in milliseconds. This is the sum of SMF74CUB*0.128. If you installed the RMF APAR OW54347, this field is no longer used.
DPB_MSEC	Amount of time, in milliseconds, during the measurement interval that I/O requests to a device were delayed because a director port was busy. This is the sum of SMF74DPB*0.128. If you installed the RMF APAR OW54347, this field is no longer used.
SAMPLES_ALLOCATED	Number of samples that indicated that the device was allocated. This is the sum of SMF74ALC.
SAMPLES_MOUNT_PEND	Number of samples that indicated a mount pending condition. This is the sum of SMF74MTP.
SAMPLES_NOT_READY	Number of samples that indicated that the device was not ready. This is the sum of SMF74NRT.
MOUNTS_DETECTED	Number of mounts detected. This is the sum of SMF74MTC.
ALLOCATED_TAPE_MAX	Maximum number of mounts detected. This is the maximum of SMF74MTC.
FLAGS	Flags for the device class; 80 = number option active, 40 = SG option active, 20 = SG name changed during the interval, 10 = Open mount at interval start, 08 = Open mount at interval end. From SMF74CLF.
SYSPLEX_NAME	This is the name of the sysplex. From SMF74XNM.
SYSTEM_NAME	MVS system name. From SMF74SNM.
PAV_SUCC_SAMPLES	Number of unsuccessful PAV sample counts. From SMF74PCT.
NUM_UCB_PAV	Average number of HyperPAV devices. Calculated as the sum of (SMF74NUX/SMF74PSM).

Table 24: (continued)

Column	Description
INIT_CMD_RSP_MSEC	Initial command response time, in milliseconds. For a start or resume function of the subchannel, it is the time required until the device indicates that has accepted the command. Calculated as SUM of SMF74CMR*0.128.
DEV_HPAV_ALIAS	Number of configured HyperPAV aliases for the LSS of the device. Calculated as the sum of SMF74HPC.
NUMBER_OF_CYL	DASD volume capacity (specified by the number of available cylinders)
IOSQ_MESC	IOS Queue time, in milliseconds, REQUESTS_QUEUED * CYCLE_LENGTH_MSEC.

LPAR Level Forecast Metrics table
This table provides forecast results. All the forecast results come from three forecast streams: hourly, daily and monthly.

Table 25: The LPAR Level Forecast Metrics table

Column	Description
CPU_BUSY_MIPS_SUM	Sum of CPU_BUSY_MIPS in MVSPM_LPAR table by different time intervals (hourly, daily, monthly).
CPU_BUSY_MIPS_SUM_PRED	Predicted value of CPU_BUSY_MIPS_SUM based on different time intervals (hourly, daily, monthly).
CPU_MODEL_NO	CPU model number. From SMF70MOD.
CPU_SERIAL_NO	Last 4 digits of the CPU serial number. From SMF70SER.
DATE	Date of the mid measurement.
FORECAST_TIME_LEVEL	Aggregated time interval for the record (hourly, daily, monthly).
LPAR_NAME	Name of the logical partition. From SMF70LPM.
MEASURED_SEC_DERIVED	Fixed MEASURED_SEC_SUM. For hourly it is 3,600, for daily is 86,400, and for monthly is 2,592,000.
MEASURED_SEC_SUM	Sum of MEASURED_SEC in MVSPM_LPAR table by different time intervals (hourly, daily, monthly).
MIPS_TOT_CAPACITY	Total MIPS capacity for specified process type and LPAR.
MVS_SYSTEM_ID	MVS system ID. This is the SMF system ID. From SMF70SID.

Table 25: The LPAR Level Forecast Metrics table (continued)	
Column	**Description**
PROCESSOR_TYPE	Name of the logical processor type. Possible values are: **Name** Description **CP** General purpose processor **ICF+** ICF pool (includes IFA, IFL and ICF) **IFA** Integrated Facility for Application processor **IFL** Integrated Facility for Linux processor **ICF** Internal Coupling Facility processor **IIP** Integrated Information Processor
TIME	Time (rounded down as specified in the MVSPM_TIME_RES lookup table) of the mid measurement.
GUARANTEED_MIPS	Predicted guaranteed MIPS value of processors based on different time intervals (hourly, daily, monthly).
LOGIC_PROC_MIPS	Logical processor busy MIPS time.
GMT_OFFSET	Offset Greenwich Mean Time (GMT) to local time, in seconds, and calculated as SMF70LGO/4096E6.
ADJUSTED_VALUE	Adjusted flag for confidence interval.
CPU_BUSY_MIPS_SUM_PRED_L	Predicted lower bound value of CPU_BUSY_MIPS_SUM based on different time intervals (hourly, daily, monthly).
CPU_BUSY_MIPS_SUM_PRED_U	Predicted upper bound value of CPU_BUSY_MIPS_SUM based on different time intervals (hourly, daily, monthly).

LPAR Level Metrics table
This table is a consolidated version of DRL.MVSPM_LPAR_H, which eliminates any duplicate entries from DRL.MVSPM_LPAR_H.

Table 26: LPAR Level Metrics table	
Column	**Description**
CAPACITY_GRP_LMT	Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.
CAPACITY_GRP_NM	Name of the capacity group this partition belongs to. From SMF70GNM.

Table 26: LPAR Level Metrics table (continued)	
Column	**Description**
CAPACITY_LIMIT_MSU	Defined capacity limit of the partition in MSUs per hour. A value of 0 indicates that the partition is not under control of the License Manager. Valid only for CP processors. From SMF70MSU.
CENTRAL_STOR_FRAME	Central storage frames assigned to this partition. Calculated as MAX of SMF70CSF.
COMPLEX_SEC	Total processor complex time, in seconds. Each processor type is considered separately.
COUNT_SHR	Number of occurrences of LPROC shared within the reported interval. Each processor type is considered separately.
CPC_MODEL_ID	CPC model identification. From the SMF70MDL.
CPU_BUSY_MIPS	Virtual processor busy MIPS time, in seconds. Calculated as the value of (SMF70PDT/1000000.) * (MIPS/TOT_PROC). MIPS is a column in lookup table MVS_MIPS_T. TOT_PROC is the physical processor total.
CPU_DISPATCH_SEC	Logical processor dispatch time, in seconds.
CPU_EFF_DISP_SEC	Logical processor effective dispatch time (excluding LPAR management time), in seconds.
CPU_MODEL_NO	CPU model number. From SMF70MOD.
CPU_SERIAL_NO	Last 4 digits of the CPU serial number. From SMF70SER.
CPU_WAIT_SEC	Processor wait time, in seconds. This is set to zero if the processor was run in LPAR mode and wait completion was not enabled. Calculated as the sum of SMF70WAT/4 096 000 000.
DATE	Date of the mid measurement.
DEDICAT_PROC_LPAR	Total number of dedicated processors for this LPAR. From LPDEDICF + LPDEDCP + LP_SPP_IFA_DED + LP_SPP_IFL_DED + LP_SPP_ICF_DED.
DEDICAT_PROCESSORS	Number of dedicated processors for this LPAR. Each processor type is considered separately.
DEDICATED_PROC_IND	Indicates whether the partition has been assigned one or more dedicated processors: 1 = yes, 0 = no. From SMF70BPS.
DISPATCH_INTERVAL	Dispatch interval time, in seconds. A zero value indicates that the interval was dynamically determined. Calculated as the value of SMF70GTS/1000.
EXPAND_STOR_FRAME	Expanded storage frames assigned to this partition. Calculated as maximum of SMF70CSF.
FROM_LPAR_NO	Number of the logical partition that wrote the record. From SMF70PTN.
LOGICAL_PROC_TOT	Total number of logical processor assigned to the partition for all occurrences of LPROC shared. Each processor type is considered separately.
LOGICAL_PROCESSORS	Number of logical processors assigned to the partition. From SMF70BDN.

Column	Description
LP_ONLINE_SEC	Total logical processor online seconds for this LPAR. From SMF70INT/1000. For z/OS system is SMF70ONT/1000000.
LPAR_FLAGS	LPAR flags. 80 = partition deactivated during interval, 40 = number of logical processors changed. From SMF70PFG.
LPAR_NAME	Name of the logical partition. From SMF70LPM.
LPAR_NO	Number of the logical partition. From SMF70LPN.
LPAR_SYSTEM_ID	System ID of the logical partition.
LPAR_WEIGHT	Logical partition processor resource weight factor. This column is set to zero if the weight factor is X'FFFF'. From SMF70BPS.
LPAR_WEIGHT_TOT	Sum of all partition processor resource weight factors. Each processor type is considered separately.
LPROC_SEC	Available logical processor time, in seconds. This is the time available for all nondedicated processors in the complex. Each processor type is considered separately.
MEASURED_SEC	Measured time period, in seconds. This is the duration of the RMF measurement intervals. From SMF70INT/1000.
MIPS_PER_LOG	Total number of MIPS per logical processor type. The number of MIPS is taken from a specific Look Up table, MVS_MIPS_T.
MIPS_PER_LOG_CP	Number of MIPS per single logical CPU.
MIPS_TOT_CAPACITY	Total MIPS capacity for specified process type and LPAR.
MVS_LEVEL	MVS software level. This is the acronym, version number, release number, and modification level. From SMF70MVS.
MVS_SYSTEM_ID	MVS system ID
NONDED_PROCESSORS	Number of processors that were not assigned as dedicated in the processor complex. Each processor type is considered separately.
NUM_CONSUMED_MSU	Actual number of consumed MSUs. Valid only for CP processors. Calculated as Sum(SMF70PDT)*3600*16/(SMF70CPA*SMF70INT*1000).
PERIOD_NAME	Name of the period.
PHY_PROC_CNT	The number of physical processors currently recorded.
PHY_PROC_SPP_ICF	Total number of ICF physical processors assigned for use by PR/SM. From TOT_SPP_ICF_PRO.
PHY_PROC_SPP_IFA	Total number of IFA physical processors assigned for use by PR/SM. From TOT_SPP_IFA_PRO.
PHY_PROC_SPP_IFL	Total number of IFL physical processors assigned for use by PR/SM. From TOT_SPP_IFL_PRO.
PHY_PROC_SPP_IIP	Total number of IIP physical processors assigned for use by PR/SM. From TOT_SPP_IIP_PRO.
PHY_PROCESSORS_CP	Total number of CP physical processors assigned for use by PR/SM. From TOTCPPRO.

Table 26: LPAR Level Metrics table (continued)

Table 26: LPAR Level Metrics table (continued)	
Column	**Description**
PHYS_PROC_LPAR	Total number of physical CPUs utilized in LPAR.
PHYS_PROCESSORS	Total number of physical processors assigned for use by PR/SM. From TOTCPPRO+TOTICFPRO+TOT_SPP_IFA_PRO+ TOT_SPP_IFL_PRO+ TOT_SPP_ICF_PRO.
PROCESSOR_FLAGS	Logical processor flags. 80 = wait state assist enabled; 40 = wait state assist changed; 20 = relative share changed; 10 = partition capping is enabled; and 08 = partition capping status changed. From SMF70VPF.
PROCESSOR_TYPE	Name of the logical processor type. Possible values are: **Name** Description **CP** General purpose processor **ICF+** ICF pool (includes IFA, IFL and ICF) **IFA** Integrated Facility for Application processor **IFL** Integrated Facility for Linux processor **ICF** Internal Coupling Facility processor **IIP** Integrated Information Processor
PRSM_FLAGS	PR/SM flags. 80 = diagnose X'204' failure; 40 = number of physical processors changed; 20 = global time slice value changed; and 10 = physical partition included. From SMF70INB.
RECORD_LEVEL	SMF record level change number. This column enables processing of SMF record level changes in an existing release. From SMF70SRL.
RMF_VERSION	RMF version number. From SMF72MFV.
SYSPLEX_NAME	The sysplex name to be associated with the MVS system ID or the MVS system name.
SYSTEM_NAME	MVS system name
TIME	Time (rounded down as specified in the MVSPM_TIME_RES lookup table) of the mid measurement.
TIME_RESOLUTION	Time resolution for the set of tables, in minutes. This defines the time period for which data is to be recorded. Valid values are 1, 2, 3, 5, 6, 10, 12, 15, 20, and 30 or any multiple of 60.
WEIGHT_ACTUAL	Actual weighting of the shared processor resources. Valid only for z/OS system with WLM active. Calculated as the average of SMF70ACS / SMF70DSA.
WLM_CAPPING_PCT	Percentage of WLM capping of the partition. Valid only for CP processors. Calculated as the maximum of (SMF70NSW/ SMF70DSA) * 100.

Table 26: LPAR Level Metrics table (continued)

Column	Description
GUARANTEED_MIPS	Guaranteed busy MIPS time.
LOGIC_PROC_MIPS	Logical processor busy MIPS time.
GMT_OFFSET	Offset Greenwich Mean Time (GMT) to local time, in seconds, and calculated as SMF70LGO/4096E6.

LPAR MSU Forecast Metrics table

The LPAR MSU Forecast Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

LPAR_MSU
LPAR consumed MSU hourly

LPAR_MSU_4HRA
LPAR four hour rolling average MSU.

LPAR_MSU_4HRA_UTILRA
LPAR four hour rolling average MSU after capping.

CPC_MSU_4HRA_UTILRA_FORECAST
CPC four hour rolling average MSU after capping.

LPAR_MSU_FORECAST
Forecasted LPAR consumed MSU hourly.

LPAR_MSU_4HRA_FORECAST
Forecasted LPAR four hour rolling average MSU.

LPAR_MSU_4HRA_UTILRA_FORECAST
Forecasted LPAR four hour rolling average MSU after capping.

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

CPU_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

NOMINAL_CAPACITY
Nominal processor capacity available to the central processor complex, measured in MSU/h. Calculated as the average of SMF70NCR.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group, but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

MACHINE_SERIAL_NO
z/OS machine Serial No.

HOUR_NO
Hour No.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

FORECAST_TIME_LEVEL
Forecast time level. Possible values are hourly, monthly.

LPAR MSU Optimization Metrics table

The LPAR MSU Optimization Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

LPAR_MSU_FORECAST
Forecasted LPAR consumed MSU hourly

LPAR_MSU_OPTIMIZED
Optimized LPAR consumed MSU hourly

LPAR_MSU_4HRA_OPTIMIZED
Optimized LPAR four hour rolling average MSU

LPAR_MSU_4HRA_UTILRA_OPTIMIZED
Optimized LPAR four hour rolling average MSU after capping.

CPC_MSU_4HRA_UTILRA_OPTIMIZED
Optimized CPC four hour rolling average MSU after capping.

LPAR_MSU_4HRA_FORECAST
Forecasted LPAR four hour rolling average MSU.

LPAR_MSU_4HRA_UTILRA_FORECAST
Forecasted LPAR four hour rolling average MSU after capping.

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

CPU_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

NOMINAL_CAPACITY
Nominal processor capacity available to the central processor complex, measured in MSU/h. Calculated as the average of SMF70NCR.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group, but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

MACHINE_SERIAL_NO
z/OS machine Serial No.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

HOUR_NO
Hour No.

LPAR MSU Metrics table

The LPAR MSU Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

CPC_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

MSU_4HRA
Long-term average of CPU service (millions of service units). From SMF70LAC.

LPAR_MSU_4HRA
LPAR 4 hour rolling average MSU

LPAR_MSU_4HUTILRA
Four hour rolling average LPAR MSU after capping.

LPAR_CONSUMED_MSU
LPAR consumed MSU

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

LPAR_MSU_1HTOT
LPAR consumed MSU hourly.

HOUR_NO
Hour No.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group, but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

CPC_MSU_4HUTILRA
CPC four hour rolling average MSU after capping.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

MIPS Capacity Ratings table

This table provides the CPU_SERIAL_MIPS rating for your processors.

Table 27: MIPS Capacity Ratings table

Column	Description
CPC_MODEL_ID	CPC model identification. From the field SMF70MDL.
CPU_MODEL_NO	CPU model number. From SMF70MOD.
CPU_SERIAL_NO	Last 4 digits of the CPU serial number. From SMF70SER.
MIPS_TOT_CAPACITY	Total MIPS capacity for specified process type and LPAR.
PHY_PROC_CNT	The number of physical processors. For example, when processor type is CP, phy_proc_cnt = phy_processors_cp.

Column	Description
PROCESSOR_TYPE	Name of the logical processor type. Possible values are: **Name** Description **CP** General purpose processor **ICF+** ICF pool (includes IFA, IFL and ICF) **IFA** Integrated Facility for Application processor **IFL** Integrated Facility for Linux processor **ICF** Internal Coupling Facility processor **IIP** Integrated Information Processor

Table 27: MIPS Capacity Ratings table (continued)

Movement Suggestion by Optimization Metrics table

The Movement Suggestion by Optimization Metrics table contains the following columns.

MOVE_STEP_ID
The sequence to do the movement.

TIMESTAMP
Timestamp of the mid measurement.

FROM_LPAR_NAME
Optimization movement source LPAR name.

FROM_CPC_SERIAL_NO
Source CPC serial number.

FROM_MVS_SYSTEM_ID
Source MVS system ID.

PROD_ID
Product ID number.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PRICING_STRUCTURE
Product sub-capacity program.

TO_LPAR_NAME
Optimization movement target LPAR name.

TO_CPC_SERIAL_NO
Movement target CPC serial number.

TO_MVS_SYSTEM_ID
Movement source MVS system ID.

MOVING_AMOUNT
The amount of the workload (MSU) for this product that should be moved in this step. (-1 mean the remainder of the workload.)

MOVE_ALL
1 - move all of the remaining workload (MSU) of this product in this step. 0 - move the amount of the workload indicated by MOVING_AMOUNT.

TIMEZONE
The time zone of the timestamp.

NO89 Products table

The NO89 Products table contains the following columns.

STARTDATE
Start date of the mid measurement. The default is 1900-01-01.

ENDDATE
End date of the mid measurement. The default is 2900-01-01

PRODUCT_ID
Product ID number.

CPC_MODEL_NO
CPC model number. From SMF70MOD.

CPC_SERIAL_NO
Last 4 digits of the CPC serial number. From SMF70SER.

PRODUCT_NAME
Product name.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

Parents of Sub-Capacity Reference-based Programs table

The Parents of Sub-Capacity Reference-based Programs table contains the following columns.

PROD_ID
Product ID number.

PARENT_NAME
Product Parent name.

DESCRIPTION
Product standard name from IBM website.

Product Billable MSU Forecast Metrics table

The Product Billable MSU Forecast Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

TIMESTAMP
Timestamp of the mid measurement.

PROD_ID
Product ID number.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PRICING_STRUCTURE
Product sub-capacity program.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

IWP
Product IWP program. For IWP program, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

IWP_ADJUSTED
Product IWP adjusted program. For IWP adjusted product, this is 1.

ZOS_BASED
Product z/OS based program. For z/OS based product, this is 1.

REFERENCE_BASED
Product reference based program. For reference based product, this is 1.

EXECUTION_BASED
Product IPLA execution based program. For execution based product, this is 1.

ZNALC_LICENSE
Indicates that LICENSE=zNALC was specified in EASYSxx. Possible values are Y or N. From Bit 1 of SMF89PFL.

IWP_DEFINING
Product IWP defining program. For IWP defining product, this is 1.

PROD_PARENT
Product Parent name.

PROD_LPAR_BILLABLE_MSU
LPAR billable MSU.

PROD_LPAR_BILLABLE_MSU_F
Forecasted LPAR billable MSU.

PROD_BILLABLE_MSU
Product billable MSU.

PROD_BILLABLE_MSU_F
Forecasted product billable MSU.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

HOUR_NO
Hour No.

PROD_SHORT_VERSION
Product short version.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

Product Billable MSU Optimization Metrics table

The Product Billable MSU Optimization Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

TIMESTAMP
Timestamp of the mid measurement.

PROD_ID
Product ID number.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PRICING_STRUCTURE
Product sub-capacity program.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

IWP
Product IWP program. For IWP program, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

IWP_ADJUSTED
Product IWP adjusted program. For IWP adjusted product, this is 1.

ZOS_BASED
Product z/OS based program. For z/OS based product, this is 1.

REFERENCE_BASED
Product reference based program. For reference based product, this is 1.

EXECUTION_BASED
Product IPLA execution based program. For execution based product, this is 1.

ZNALC_LICENSE
Indicates that LICENSE=zNALC was specified in EASYSxx. Possible values are Y or N. From Bit 1 of SMF89PFL.

IWP_DEFINING
Product IWP defining program. For IWP defining product, this is 1.

PROD_PARENT
Product parent name.

PROD_LPAR_BILLABLE_MSU_O
Optimized LPAR billable MSU.

PROD_LPAR_BILLABLE_MSU_O
Optimized LPAR billable MSU.

PROD_LPAR_BILLABLE_MSU_F
Forecasted LPAR billable MSU.

PROD_BILLABLE_MSU_O
Optimized product billable MSU.

PROD_BILLABLE_MSU_F
Forecasted product billable MSU.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

HOUR_NO
Hour No.

PROD_SHORT_VERSION
Product short version.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

Product MSU Forecast Metrics table

The Product MSU Forecast Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_ID
Product ID number.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PRICING_STRUCTURE
Product sub-capacity program.

ZOS_BASED
Product z/OS based program. For z/OS based product, this is 1.

REFERENCE_BASED
Product reference based program. For reference based product, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

EXECUTION_BASED
Product IPLA execution based program. For execution based product, this is 1.

GSSP_DB2_UTIL_MSU
GSSP DB2 MSU value.

GSSP_DB2_UTIL_MSU_F
Forecasted GSSP DB2 MSU.

PROD_MSU
Product consumed MSU hourly.

PROD_MSU_4HRA
Four hour rolling average product MSU.

PROD_MSU_4HRA_UTILRA
Four hour rolling average product MSU after capping.

PROD_MSU_FORECAST
Forecasted product MSU.

PROD_MSU_4HRA_FORECAST
Forecasted product four hour rolling average MSU.

PROD_MSU_4HRA_UTILRA_FORECAST
Forecasted product four hour rolling average MSU after capping.

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

IWP
Product IWP program. For IWP program, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

IWP_ADJUSTED
Product IWP adjusted program. For IWP adjusted product, this is 1.

CPU_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

NOMINAL_CAPACITY
Nominal processor capacity available to the central processor complex, measured in MSU/h. Calculated as the average of SMF70NCR.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group, but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

ZNALC_LICENSE
Indicates that LICENSE=zNALC was specified in EASYSxx. Possible values are Y or N. From Bit 1 of SMF89PFL.

MACHINE_SERIAL_NO
z/OS machine Serial No.

HOUR_NO
Hour No.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

PROD_PARENT
Product Parent name.

GSSP_SCALING_FACTOR
GSSP Product scaling factor. Default is 2.

IWP_DEFINING
Product IWP defining program. For IWP defining product, this is 1.

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

FORECAST_TIME_LEVEL
Forecast time level. Possible values are hourly, monthly.

PROD_SHORT_VERSION
Product short version.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

Product MSU Optimization Metrics table

The Product MSU Optimization Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

MVS_SYSTEM_ID
MVS system ID.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_ID
Product ID number.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PRICING_STRUCTURE
Product sub-capacity program.

ZOS_BASED
Product z/OS based program. For z/OS based product, this is 1.

REFERENCE_BASED
Product reference based program. For reference based product, this is 1.

EXECUTION_BASED
Product IPLA execution based program. For execution based product, this is 1.

PROD_MSU_FORECAST
Forecasted product MSU.

PROD_MSU_OPTIMIZED
Optimized product MSU.

PROD_MSU_4HRA_OPTIMIZED
Optimized product four hour rolling average MSU.

PROD_MSU_4HRA_UTILRA_OPTIMIZED
Optimized product four hour rolling average MSU after capping.

PROD_MSU_4HRA_FORECAST
Forecasted product four hour rolling average MSU.

PROD_MSU_4HRA_UTILRA_FORECAST
Forecasted product four hour rolling average MSU after capping.

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

IWP
Product IWP program. For IWP program, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

IWP_ADJUSTED
Product IWP adjusted program. For IWP adjusted product, this is 1.

CPU_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

NOMINAL_CAPACITY
Nominal processor capacity available to the central processor complex, measured in MSU/h. Calculated as the average of SMF70NCR.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long-term average of CPU service in millions of service units which would be allowed by the limit of the capacity group, but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

ZNALC_LICENSE
Indicates that LICENSE=zNALC was specified in EASYSxx. Possible values are Y or N. From Bit 1 of SMF89PFL.

MACHINE_SERIAL_NO
z/OS machine Serial No.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

PROD_PARENT
Product Parent name.

GSSP_SCALING_FACTOR
GSSP Product scaling factor. Default is 2.

IWP_DEFINING
Product IWP defining program. For IWP defining product, this is 1.

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

HOUR_NO
Hour No.

PROD_SHORT_VERSION
Product short version.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

Product MSU Utilization Metrics tables

The Product MSU Utilization Metrics table, z/OS Product MSU Utilization Metrics table, and Product MSU Utilization Metrics without z/OS table all use the same structure.

The Product MSU Utilization Metrics tables contains the following columns.

DATE
: Date of the mid measurement.

TIME
: Time of the mid measurement.

CPC_MNFCTR
: V1-CPC manufacturer. From SMF89MNF.

CPC_SERIAL_NO
: Last four digits of the CPC serial number. From SMF70SER.

CPC_MODEL_NO
: CPU model number. From SMF70MOD.

CPC_MODEL_ID
: CPC model ID. From SMF70MDL.

MACHINE_SERIAL_NO
: zOS machine Serial No.

LPAR_NAME
: Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
: MVS system ID.

PROD_NAME
: Product name. From SMF89UPN or SMF89T2ProdName.

PROD_ID
: Product ID number.

PROD_PARENT
: Product Parent name.

SMF89_REGISTERED
: Product register in SMF89 records. For SMF89 registered, this is 1.

ZOS_BASED
: Product zOS Based program. For zOS Based product, this is 1.

REFERENCE_BASED
: Product reference based program. For reference Based product, this is 1.

GSSP
: Product GSSP program. For GSSP product, this is 1.

GSSP_SCALING_FACTOR
: GSSP Product scaling factor. Default is 2.

GSSP_DB2_UTIL_MSU
: GSSP DB2 MSU.

IWP
: Product IWP program. For IWP program, this is 1.

IWP_ADJUSTED
: Product IWP adjusted program. For IWP adjusted product, this is 1.

IWP_DEFINING
: Product IWP defining program. For IWP defining product, this is 1.

EXECUTION_BASED
: Product IPLA execution based program. For execution based product, this is 1.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_SHORT_VERSION
Product short version.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

PRODUCT_DISPLAY_VALUE
Combined product ID, product name and pricing structure.

PRICING_STRUCTURE
Product sub-capacity program.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PROD_CP_SECONDS
Product CPU time.

PROD_MSU_CONSUMED
Product consumed MSU.

PROD_MSU_1HAVG
Product consumed MSU hourly.

PROD_MSU_4HUTILRA
Product four hour rolling average MSU after capping.

LPAR_MSU_4HUTILRA
LPAR four hour rolling average MSU after capping.

HOUR_NO
Hour No.

UNIQUE_HOUR_NO
Unique hour No. from 1980-01-01 00:00:00

PROC_CAPACITY_MSU
Processor capacity available to MVS image, measured in MSUs. Calculated as the maximum of SMF70WLA.

LPAR_CONSUMED_MSU
LPAR consumed MSU.

CAPACITY_LIMIT_MSU
Defined capacity limit of the partition in MSUs per hour.

LPAR_MSU_4HRA
LPAR four hour rolling average MSU.

LPAR_MSU_1HTOT
LPAR consumed MSU hourly.

OTHER_PRODS_MSU_1HTOT
Other products consumed MSU hourly.

CPC_MSU_4HUTILRA
CPC four hour rolling average MSU after capping.

CAPACITY_GRP_NM
CAPACITY_GRP_NMacity group this partition belongs to. From SMF70GNM.

CAPACITY_GRP_LMT
Group maximum licensing units which may be consumed per unit of time, on average. From SMF70GMU.

CAPACITY_GRP_AVL
Long term average of CPU service in millions of service units which would be allowed by the limit of the capacity group but is not used by its members. If the value is negative, this capacity group is subject to capping. From SMF70GAU.

ZNALC_LICENSE
Indicates that LICENSE=zNALC was specified in EASYSxx. Possible values are Y or N. From Bit 1 of SMF89PFL.

NOMINAL_CAPACITY
Nominal processor capacity available to the central processor complex, measured in MSU/h. Calculated as the average of SMF70NCR.

IWP_DEFINING_PROD_MSU_4HUTILRA
IWP defining product four hour rolling average MSU after capping.

IWP_ADJUSTED_PROD_MSU_4HRA
IWP adjusted product four hour rolling average MSU.

GSSP_PROD_MSU_4HRA
GSSP product db2 four hour rolling average MSU after capping.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

IWP_DEFINED_MSU
IWP defining product consumed MSU.

CAPACITY_GRP_MSU_4HUTILRA
Total LPAR MSU by group.

DESCRIPTION
Product standard name from IBM websie.

MLC
Product MLC program. For MLC product, this is 1.

IPLA
Product IPLA program. For IPLA program, this is 1.

PRODUCT_CATEGORY
Product is MLC program or IPLA program, others are blank.

Prompts/CPC Prompt View table

The Prompts/CPC Prompt View table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

Prompts/LPAR MSU Prompt View table

The Prompts/LPAR MSU Prompt View table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

CPC_SERIAL_NO
Last 4 digits of the CPC serial number. From SMF70SER.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

Prompts/Product Billable MSU Prompt View table

The Prompts/Product Billable MSU Prompt View table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PROD_ID
Product ID number.

PROD_PARENT
Product Parent name.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PROD_FEATURE_NM
Product qualifier. From SMF89UPQ or SMF89T2FeatureName.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

Prompts/Product Prompt for LPAR MSU Report View table

The Prompts/Product Prompt for LPAR MSU Report View table contains the following columns.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

Raw LPAR Product MSU Utilization Metrics table

The Raw LPAR Product MSU Utilization Metrics table contains the following columns.

DATE
Date of the mid measurement.

TIME
Time of the mid measurement.

CPC_SERIAL_NO
Last four digits of the CPC serial number. From SMF70SER.

CPC_MODEL_NO
CPU model number. From SMF70MOD.

CPC_MODEL_ID
CPC model ID. From SMF70MDL.

LPAR_NAME
Name of the logical partition. From SMF70LPM.

MVS_SYSTEM_ID
MVS system ID.

PROD_NAME
Product name. From SMF89UPN or SMF89T2ProdName.

PROD_ID
Product ID number.

PROD_VERSION
Product version. From SMF89UPV or SMF89T2ProdVers.

PRODUCT_DISPLAY_VALUE
Combined product ID, product name and pricing structure.

PRODUCT_USE_VALUE
Combined product ID, product name and pricing structure.

PROD_MSU_CONSUMED
Product consumed MSU.

HOUR_NO
Hour No.

LPAR_CONSUMED_MSU
LPAR consumed MSU.

LPAR_MSU_1HTOT
LPAR consumed MSU hourly.

GMT_OFFSET
Offset GMT to local time, in seconds. Calculated as SMF70LGO/4096E6.

PROD_OWNER
Product owner or vendor name. From SMF89UPO or SMF89T2ProdOwner.

Sub-Capacity Program Pricing Structures table

The Sub-Capacity Program Pricing Structures table contains the following columns.

DATE
Date of the mid measurement.

PROD_ID
Product ID number.

PARENT_NAME
Product Parent name.

SMF89_REGTERED
Product register in SMF89 records. For SMF89 registered, this is 1.

MLC
Product MLC program. For MLC product, this is 1.

IPLA
Product IPLA program. For IPLA program, this is 1.

ZOS_BASED
Product z/OS Based program. For z/OS based product, this is 1.

REFERENCE_BASED
Product reference based program. For reference based product, this is 1.

GSSP
Product GSSP program. For GSSP product, this is 1.

GSSP_SCALING_FACTOR
GSSP Product scaling factor. Default is 2.

IWP_ADJUSTED
Product IWP adjusted program. For IWP adjusted product, this is 1.

IWP_DEFINING
Product IWP defining program. For IWP defining product, this is 1.

EXECUTION_BASED
Product IPLA execution based program. For execution based product, this is 1.

DESCRIPTION
Product standard name from IBM website.

IWP
Product IWP program. For IWP program, this is 1.

PRICING_STRUCTURE
Product sub-capacity program.

Timezone Metrics table

The Timezone Metrics table contains the following columns.

TIMEZONE_NAME
Timezone name.

Description
Timezone description.

GMT_OFFSET
Offset GMT to local time, in seconds.

FLAG
Frequently-used timezone, default is 1.

Logical View

The Logical View organizes elements into logical query subjects.

The Logical View contains some Query Subjects.

All relationships should be defined in the Logical View.

Physical View

The physical view, using `SQL (“select * from TableName”)` to retrieve table information, reflects the underlying tables and join relationships in the data source.

CMA DW

The CMA DW namespace holds folders for the Tivoli Decision Support for z/OS (TDSz) schema and the IBM Capacity Management Analytics schema.

DRL

DRL is the folder for all the TDSz components. Within this folder all the TDSz component tables are grouped by their prefix. For example, all DB2_ tables are found in the DB2 folder.

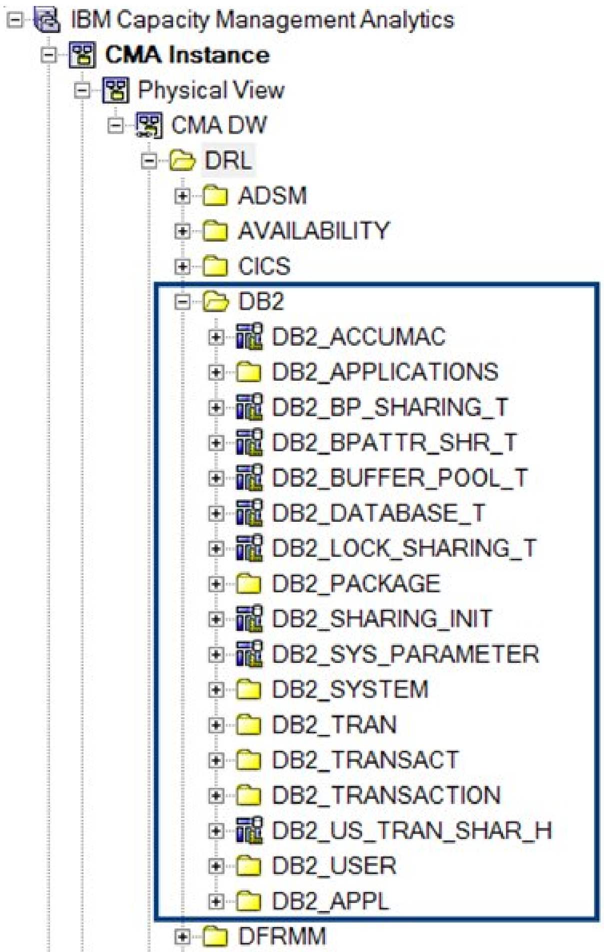

Figure 20: Contents of the DRL folder

Tables within a folder are further grouped by their name. For example, the DB2_APPLICATIONS folder is within the DB2 folder, and it holds all DB2_APPLICATIONS tables.

Refer to the TDSz documentation to find the table you need for your specific reporting needs. Especially, the SMF_113 folder that is provided with TDSz with the z/OS System (MVS) Component. This folder supports the z13 SMF type 113 records.

DRLSYS

The DRLSYS folder contains the catalog tables specifically for TDSz. These tables are used to do reporting on TDSz itself.

Refer to the TDSz documentation to find the table you need for your specific reporting needs.

HCM

The HCM folder holds all the Capacity Management Analytics specific tables.

FORECASTS
: Holds all forecast results. These results are used to do MIPS forecast reporting. All the forecast results come from the following forecast streams: hourly, daily and monthly.

HCM.MVSPM_LPAR
: A consolidated version of the DRL.MVSPM_LPAR_H table, which eliminates any duplicate entries from the DRL.MVSPM_LPAR_H table.

HCM.MIPS_CAPACITY
: Holds the CPU_SERIAL_MIPS rating for your processors

HCM.DATE_BAND
Based on interval band, this table returns the first day of every date period.

HCM.TIME_BAND
Based on interval band, this table returns the first minute of every time period.

HCM.QUARTER_BAND
Returns the first day of every season.

HCM.MONTH_BAND
Returns the first day of every month.

HCM.YEAR_BAND
Returns the first day of every year.

HCM.GMT_TIME
Returns the time , depending on the format of time expression, that results from adding integer expression seconds to time expression.

HCM.GMT_DATE
Returns the DATE, depending on the format of date and time expression, that results from adding integer expression seconds to date expression.

HCM.GMT_TIMESTAMP
Returns the TIMESTAMP, depending on the format of timestamp expression, that results from adding integer expression seconds to timestamp expression.

HCM.ROUNDOFF
Round off MSU value. For example, ROUNDOFF (5.6) = 5, ROUNDOFF (0.5)=1, ROUNDOFF (4.99999999) = 5.

HCM.UTILMSU
For LPAR capping, compare LPAR MSU with defined capacity or processor capacity, and return minimum value.

HCM.UTIL_GRP_MSU
For sysplex group capping, compare total LPAR MSU values with group capacity limit, and return minimum value.

HCM.PRICE_PER_MSU_DETERMINE
Returns pricing per MSU by CPC Serial No, Product ID, and Product MSU.

HCM.TIER_DETERMINE
Returns pricing tier by CPC Serial No, Product ID, and Product MSU.

HCM.TIER_MAX_MSU
Returns maximum MSU of this pricing tier by CPC Serial No, Product ID, and Product MSU.

HCM.TIER_MIN_MS
Returns minimum MSU of this pricing tier by CPC Serial No, Product ID, and Product MSU.

HCM.TIER_PRICE
Returns price of this pricing tier by CPC Serial No, Product ID, and Product MSU

HCM.TIER_VU
Returns Interim VU calculation result by CPC Serial No, Product ID, and Product MSU.

HCM.TIER_VU_PER_MSU
Returns entitled MSUs by CPC SERIAL NO, Product ID, and Product MSU.

HCM.GSSP_MSU
Using product 4 hours rolling average MSU value, and LPAR four hours rolling average MSU value, returns GSSP product billable MSU.

HCM.IWP_ADJUSTED_MSU
Returns IWP adjusted product billable of one LPAR.

Opening the Capacity Management Analytics Model in Framework Manager

You can view the IBM Capacity Management Analytics model in IBM Cognos Framework Manager. The Framework Manager project file (.cpf) for the Capacity Management Analytics model is included in the Capacity Management Analytics installation.

Before you begin

Copy the entire contents of the `C:\Program Files (x86)\IBM\Capacity Management Analytics 1.2\FM_model` directory from the computer where you installed Capacity Management Analytics to the computer where you have IBM Cognos Framework Manager installed.

Procedure

1. Launch IBM Cognos Framework Manager.
2. From the **Welcome** page, click **Open a project**.
3. Navigate to the project folder and click the `CMA.cpf` file.
4. Click **Open**.

Enabling a Capacity Management Analytics package with Dynamic Query Mode

To optimize performance, configure your IBM Capacity Management Analytics package to use Dynamic Query Mode (DQM).

Testing query subjects against Dynamic Query Mode

Test your query subjects to confirm that they run correctly using Dynamic Query Mode (DQM). If you modify any reports in your model, you should confirm that all queries work correctly with DQM.

Procedure

1. In the Project Viewer, expand the IBM Capacity Management Analytics model.
2. Right-click any shortcut under **Presentation View**, then click **Go to target**.
3. Right-click the target query subject, then click **Test**.
4. Click **Test Sample**.
5. Click **Close**.

Publishing a package with Dynamic Query Mode support

Once a package has been modified in any way, you must explicitly publish the package to make changes available to users.

Procedure

1. In the Project Viewer, right-click **Packages**, then click **Create**, **Package**.
2. In the Create Package dialog box, enter a **Name** for the package you are publishing.
3. Optionally, enter a **Description** and **Screen Tip** for the package.
4. Click **Next**.
5. Select the metadata you want to include in your published package.
6. Click **Finish**.
7. Click **Yes** when prompted to open the Publish Package wizard.
8. In the Publish Wizard – Select Location Type window, click **IBM Cognos 10 Content Store**.

9. Specify the folder to which you want to publish the package. For example `Public Folders\Customization\Models`.
10. Click **Next**.
11. Click **Next** again.
12. Click **Publish**.
13. Click **Finish**.

Verifying the query mode used by a published package

After you publish a package, you should verify that the package uses Dynamic Query Mode.

Procedure

1. In IBM Cognos Connection, open the folder where your published package resides.
2. On the **Actions** column, click **Set Properties**.
3. On the **General** properties tab, confirm that the **Query Mode** is set to Dynamic.

Chapter 14. Report and workspace reference

IBM Capacity Management Analytics provides several predefined reports and workspaces that let you monitor and analyze the capacity of your IT services and the IT infrastructure.

Opening a report or workspace

You can access all IBM Capacity Management Analytics reports and workspaces through IBM Cognos Connection.

Procedure

1. Open IBM Cognos Connection.
2. Click the **Public Folders** tab.
3. Click **IBM Capacity Management Analytics Solution Kit**.
4. Click **Reports**.
5. Click **All** to view links to all available reports, or **Workspaces** to view links to all available workspaces.
6. Optional: Click **Anomaly Detection Analytics**, **Application Analytics**, **CPU**, **Memory**, **Software Cost Analysis**, or **WLM** to view links to the reports available in any of those categories.
7. Click the title of the report or workspace that you want to view.

   The report or workspace opens in IBM Cognos Viewer. For details on using Cognos Viewer, see IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J/welcome).

Common report and workspace features

The IBM Capacity Management Analytics reporting application provides consistent features across all reports and workspaces.

Use this sample report to become familiar with the elements that are common to all of the Capacity Management Analytics reports and workspaces.

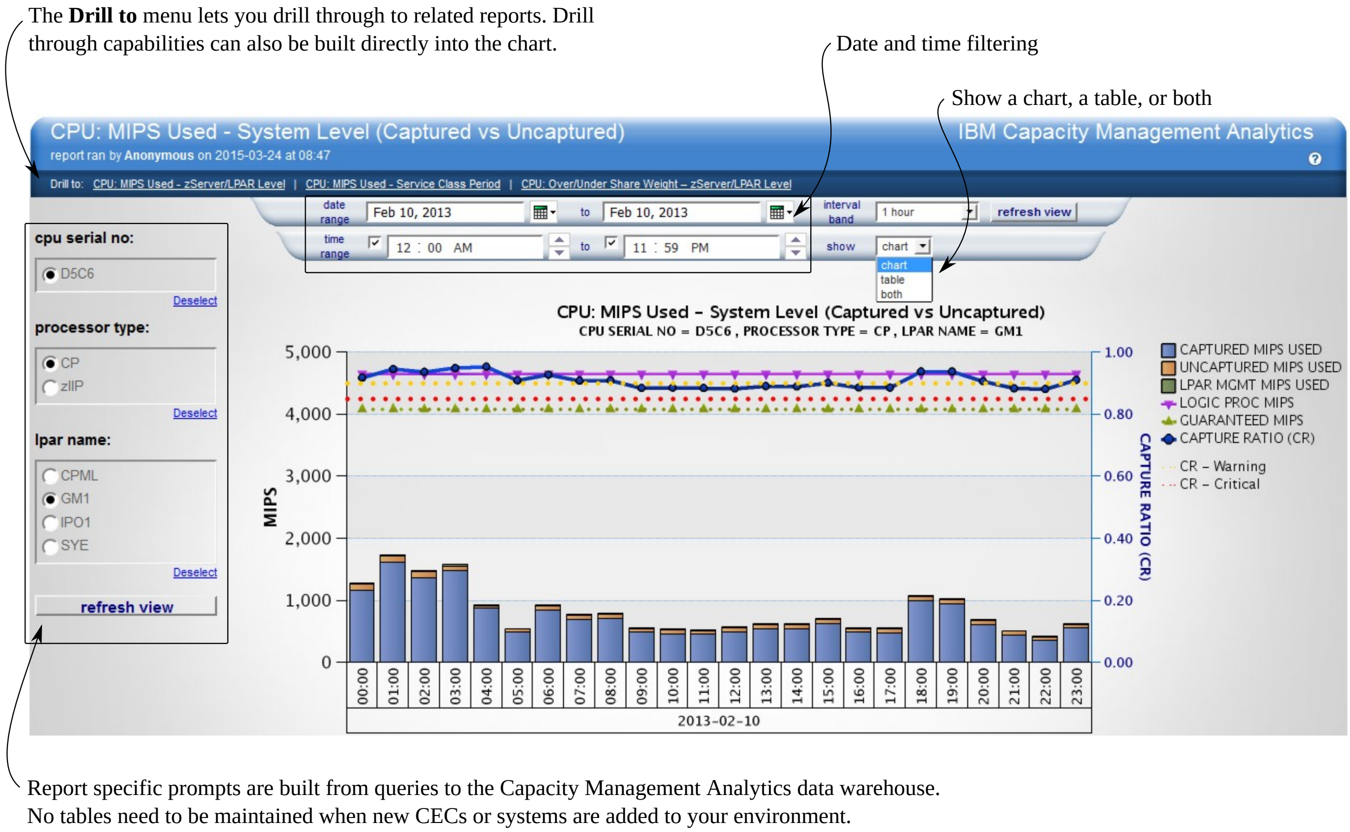

Figure 21: Sample Capacity Management Analytics report

Drill to

Contains a list of related reports that a user can drill through to for additional information. Drill to capabilities are also provided in some reports by clicking a section in the report body.

Tip: Note: When you drill through to related reports, the prompt criteria is automatically populated.

Standard prompts

date range

The date range of the data to be analyzed.

time range

The time range within each day of the selected date range to be analyzed.

interval band

The level to which you want to aggregate.

By changing the interval band, a single report can aggregate to different levels, such as 30 minutes or 30 days. You must view the data at an interval that is a multiple of the recording interval. For example, if your recording interval is 10 minutes, you can view data at intervals of 10, 20, or 30 minutes, but not at an interval of 5 minutes.

show

The report can be displayed in different formats.

chart

Information is displayed in a graphical format.

table

Information is displayed in a table.

both

Information is displayed in both a chart and table.

Report specific prompts

Prompt criteria that is specific to the displayed report. In the sample report, the report specific prompts are **cpu serial no**, **processor type**, and **lpar name**.

The values that are populated in the report-specific prompt boxes are generated dynamically by queries to the IBM Tivoli Decision Support for z/OS or Capacity Management Analytics data stores. For example, if a new CEC (cpu serial no), processor type, or LPAR is added to the environment, there are no additional tables or lists need to be updated by the user to reflect the new items.

refresh view

The **refresh view** buttons populate the next prompt box and generate the report.

Workspace prompt and refresh buttons

To avoid refreshing charts unnecessarily, **next >** buttons are available for each prompt value in the workspace prompt area. After you select a prompt value, you must click **next >** to refresh the next prompt. Other prompts or workspace charts and tables are not refreshed until the last prompt is selected.

Important: The **next >** buttons are only for the workspace, and not for the report. You must set each prompt for them to take effect. You must also set them in sequence and click **next >** to submit the setting.

On the last prompt, click **refresh workspace** to refresh all of the charts and tables in the workspace.

Each workspace can vary according to their business needs.

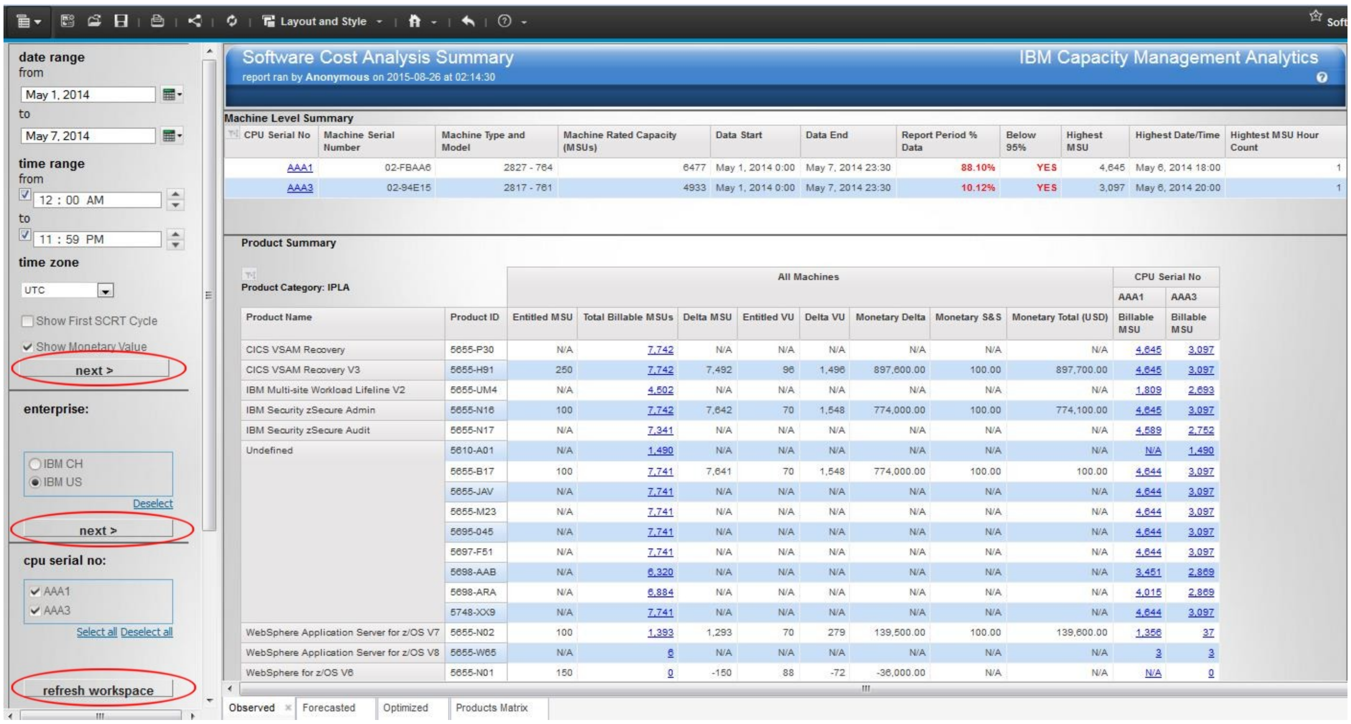

Figure 22: Workspace prompt and refresh buttons

Setting default values for report prompts

You can set default values for report prompts. When the report is run, the data is automatically filtered based on the prompt values that you specify. The user does not have to specify prompt values when the report is run.

For detailed instruction on setting default prompt values, refer to the *IBM Cognos Connection User Guide*, available on the IBM Cognos Business Intelligence welcome page (http://www.ibm.com/support/knowledgecenter/SSEP7J/welcome) of IBM Knowledge Center. For most prompts, you can select and set the default value directly in the report.

In some circumstances, such as when running a report with **Run with options**, you must manually input the value of some prompts. In these cases, refer to this table for the information you need to set prompt values.

Table 28: Report prompt values

Prompt name	**Prompt internal name**	**Prompt value for setting default**	**Prompt display value**
process type	p_processor_type	IIP IFA	zIIP zAAP
importance level	p_implvl	0 1 2 3 4 5 6	System Highest High Medium Low Lowest Discretionary
timezone	p_timezone	The integer number of GMT offset seconds. For example, 0 for UTC/GMT, -25200 for PDT(UTC-7:00).	The timezone name and the GMT offset hours. For example, UTC, GMT, PDT(UTC-7:00).
show	p_What_To_Show	0 1 2	chart table both

Table 28: Report prompt values (continued)			
Prompt name	**Prompt internal name**	**Prompt value for setting default**	**Prompt display value**
interval band (for all reports except distributed reports)	p_Interval_Type	yy:01 qq:01 mo:01 dd:01 dd:07 dd:14 dd:30 mm:60 mm:120 mm:180 mm:240 mm:360 mm:480 mm:720 mm:01 mm:05 mm:10 mm:15 mm:30	-yearly- -quarterly- -monthly- 1 day 7 days 14 days 30 days 1 hour 2 hours 3 hours 4 hours 6 hours 8 hours 12 hours 1 minute 5 minutes 10 minutes 15 minutes 30 minutes
interval band (for distributed reports)	p_Interval_Type	1 2 3 4 5	-yearly- -monthly- -weekly- -daily- -hourly-
show first SCRT cycle	p_show_first_cycle	SHOW_CYCLE	N/A
SMF collection interval	p_SMF_Interval	true	N/A
scenario	p_scenarios	observed forecast what-if	Observed Forecasted Optimized
Show Monetary Value	p_show_monetary	1	N/A

Anomaly Detection Analytics: CICS Anomaly Details

This report shows the abnormal transactions that were detected by the CICS anomaly detection streams. This report shows a summary of the abnormal transaction counts and the detail transactions that are determined to be abnormal over a short period.

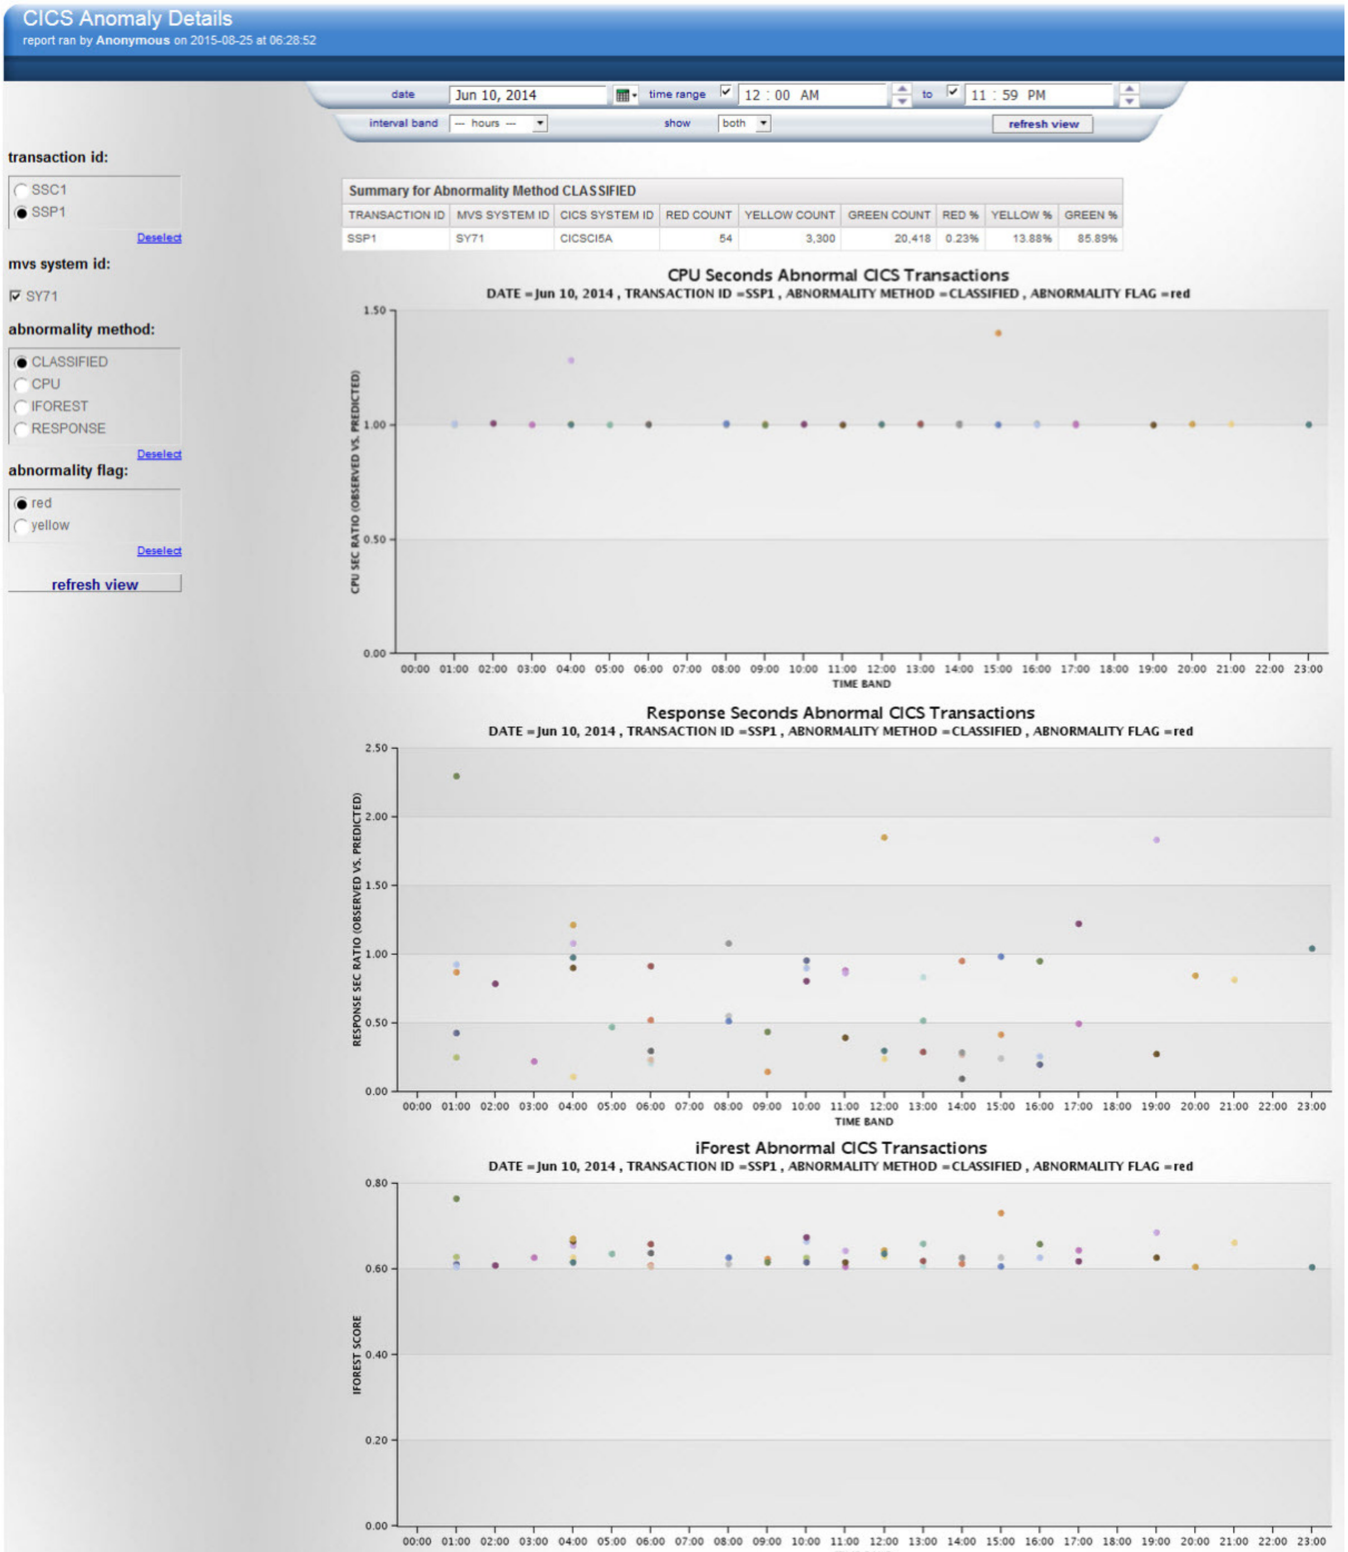

Figure 23: CICS Anomaly Details

Source tables
[Presentation View].[CICS Anomaly Detection]

[Presentation View].[DIM MINUTE]

Standard prompts:
date

time range

interval band: hours and minutes

show

Report-specific prompts
transaction id

mvs system id

abnormality method

You can select CLASSIFIED, CPU, RESPONSE, or IFOREST for the abnormality method.

abnormality flag

This prompt appears if you select CLASSIFIED in the **abnormality method**. You can select red or yellow for the flag.

The Summary for Abnormality Method CLASSIFIED table contains the following information:

TRANSACTION ID

MVS SYSTEM ID

CICS SYSTEM ID

RED COUNT
The number of transactions that are classified as red.

YELLOW COUNT
The number of transactions that are classified as yellow.

GREEN COUNT
The number of transactions that are classified as green.

RED %
The percentage of transactions that are classified as red.

YELLOW %
The percentage of transactions that are classified as yellow.

GREEN %
The percentage of transactions that are classified as green.

The Summary for Abnormality Method CPU, Summary for Abnormality Method RESPONSE, and Summary for Abnormality Method IFOREST tables contain the following information:

TRANSACTION ID

MVS SYSTEM ID

CICS SYSTEM ID

ABNORMAL COUNT
Number of transactions that are marked as abnormal by the selected abnormality method.

NORMAL COUNT
Number of transactions that are marked as normal by the selected abnormality method.

ABNORMAL %
Percentage of transactions that are marked as abnormal by the selected method.

NORMAL %
Percentage of transactions that are marked as normal by the selected method.

The CICS Anomaly Details for Abnormality Method CLASSIFIED table contains the following information:

TIME BAND
The time band is not shown if there are no abnormal transactions in that band.

START TIMESTAMP

MVS SYSTEM ID

CICS SYSTEM ID

TRANSACTION TYPE

STOP TIMESTAMP

CPU SEC (OBSERVED)
The observed CPU seconds of the transaction.

CPU SEC (PREDICTED)
The predicted CPU seconds of the transaction.

CPU SEC RATIO
The ratio between the observed CPU seconds and the predicted CPU seconds of the transaction.

RESPONSE SEC (OBSERVED)
The observed response seconds of the transaction.

RESPONSE SEC (PREDICTED)
The predicted response seconds of the transaction.

RESPONSE SEC RATIO
The ratio between the observed response seconds and the predicted response seconds of the transaction.

IFOREST SCORE
The score of the transaction using the IFOREST method.

The CICS Anomaly Details for Abnormality Method CPU table contains the following information:

TIME BAND
The time band is not shown if there are no abnormal transactions in that band.

START TIMESTAMP

MVS SYSTEM ID

CICS SYSTEM ID

TRANSACTION TYPE

STOP TIMESTAMP

CPU SEC (OBSERVED)
The observed CPU seconds of the transaction.

CPU SEC (PREDICTED)
The predicted CPU seconds of the transaction.

CPU SEC RATIO
The ratio between the observed CPU seconds and the predicted CPU seconds of the transaction.

The CICS Anomaly Details for Abnormality Method RESPONSE table contains the following information:

TIME BAND
The time band is not shown if there are no abnormal transactions in that band.

START TIMESTAMP

MVS SYSTEM ID

CICS SYSTEM ID

TRANSACTION TYPE

STOP TIMESTAMP

RESPONSE SEC (OBSERVED)
The observed response seconds of the transaction.

RESPONSE SEC (PREDICTED)
The predicted response seconds of the transaction.

RESPONSE SEC RATIO
The ratio between the observed response seconds and the predicted response seconds of the transaction.

The CICS Anomaly Details for Abnormality Method IFOREST table contains the following information:

TIME BAND
The time band is not shown if there are no abnormal transactions in that band.

START TIMESTAMP

MVS SYSTEM ID

CICS SYSTEM ID

TRANSACTION TYPE

STOP TIMESTAMP

IFOREST SCORE
The score of the transaction using the IFOREST method.

Application Analytics: Distributed CPU Usage - Application Summary Report

This report shows the 90% CPU busy line for applications. It can be displayed as hourly, daily, weekly, monthly, or yearly.

This report provides a summary of the CPU usage of all the servers associated with an application by showing the count of the servers that have usage within the different thresholds. The report is intended to help provide a quick view of all applications and their CPU usage of the servers in which they reside, by showing the overall results as well as day by day. This daily view can pinpoint when a change occurs and how many servers of an application are affected at the same time.

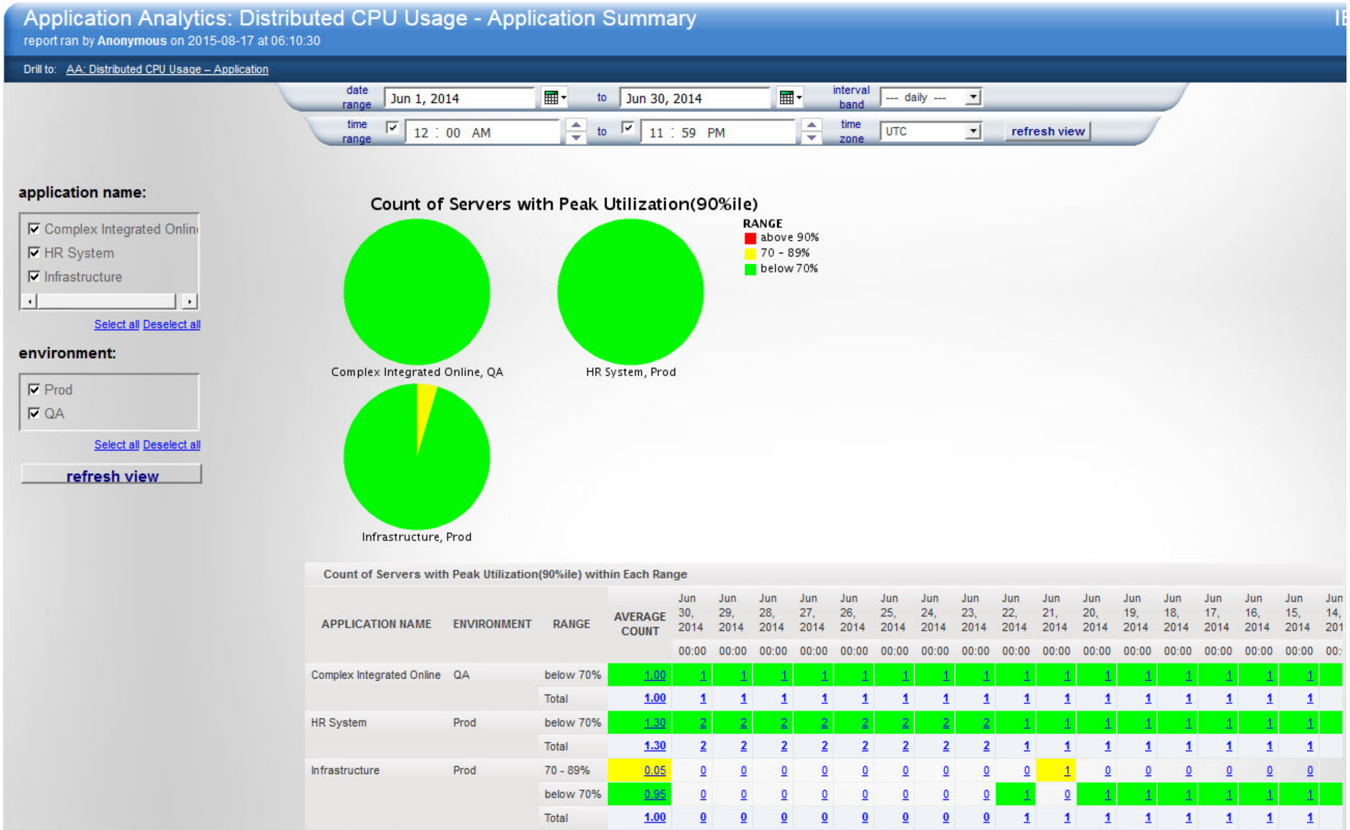

Figure 24: Application Analytics: Distributed CPU Usage - Application Summary

Source tables

[Presentation View].[Distributed Application Mapping]

[Presentation View].[AIX CPU Raw Hour Metrics]

[Presentation View].[Linux for System X CPU Hour Metrics]

[Presentation View].[Linux for System Z CPU Hour Metrics]

[Presentation View].[WIN OPERATING CPU Hour Metrics]

Drill through to the following reports

"Application Analytics: Distributed CPU Usage - Application Report" on page 199

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

Report-specific prompts

Application name

Environment

The Analytics Application: Distributed CPU usage - Application report table contains the following information:

Application Name

Environment

Server

Detail CPU Utilization by different interval bands (hourly, daily, weekly, monthly, yearly)

Green is less than 70%, yellow is between 70% and less than 90%, and red is 90% or greater.

Application Analytics: Distributed CPU Usage - Application Report

This report shows the 90% CPU busy line for applications. It can be displayed as hourly, daily, weekly, monthly, or yearly.

This report provides a view of all the servers associated with an application and shows the CPU usage for each of them with respect to thresholds of overall usage. The report is intended to help provide a quick view of which servers that support an application are running at higher utilizations than others. This report is normally used after viewing the Application Summary report.

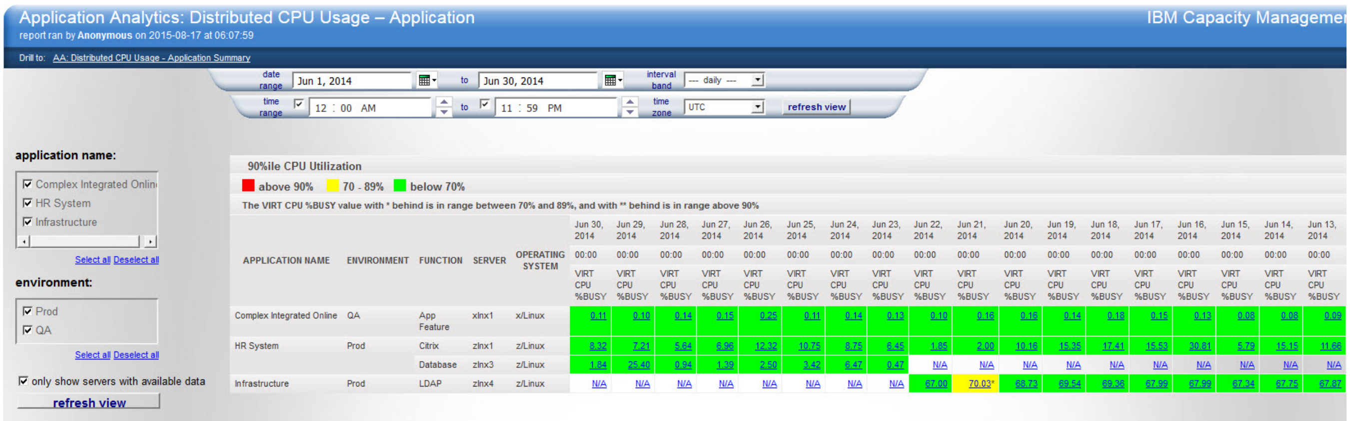

Application Analytics: Distributed CPU Usage – Application
report ran by **Anonymous** on 2015-08-17 at 06:07:59
IBM Capacity Manageme
Drill to: AA: Distributed CPU Usage - Application Summary

date range: Jun 1, 2014 to Jun 30, 2014 interval band: --- daily ---
time range: 12 : 00 AM to 11 : 59 PM time zone: UTC refresh view

application name:
- Complex Integrated Onlin
- HR System
- Infrastructure

Select all Deselect all

environment:
- Prod
- QA

Select all Deselect all

only show servers with available data

refresh view

90%ile CPU Utilization

above 90% 70 - 89% below 70%

The VIRT CPU %BUSY value with * behind is in range between 70% and 89%, and with ** behind is in range above 90%

APPLICATION NAME	ENVIRONMENT	FUNCTION	SERVER	OPERATING SYSTEM	Jun 30, 2014 00:00 VIRT CPU %BUSY	Jun 29, 2014 00:00 VIRT CPU %BUSY	Jun 28, 2014 00:00 VIRT CPU %BUSY	Jun 27, 2014 00:00 VIRT CPU %BUSY	Jun 26, 2014 00:00 VIRT CPU %BUSY	Jun 25, 2014 00:00 VIRT CPU %BUSY	Jun 24, 2014 00:00 VIRT CPU %BUSY	Jun 23, 2014 00:00 VIRT CPU %BUSY	Jun 22, 2014 00:00 VIRT CPU %BUSY
Complex Integrated Online	QA	App Feature	xlnx1	x/Linux	0.11	0.10	0.14	0.15	0.25	0.11	0.14	0.13	0.10
HR System	Prod	Citrix	zlnx1	z/Linux	8.32	7.21	5.64	6.96	12.32	10.75	8.75	6.45	1.85
		Database	zlnx3	z/Linux	1.84	25.40	0.94	1.39	2.50	3.42	6.47	0.47	N/A
Infrastructure	Prod	LDAP	zlnx4	z/Linux	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	67.00

APPLICATION NAME	SERVER	Jun 21, 2014 00:00 VIRT CPU %BUSY	Jun 20, 2014 00:00 VIRT CPU %BUSY	Jun 19, 2014 00:00 VIRT CPU %BUSY	Jun 18, 2014 00:00 VIRT CPU %BUSY	Jun 17, 2014 00:00 VIRT CPU %BUSY	Jun 16, 2014 00:00 VIRT CPU %BUSY	Jun 15, 2014 00:00 VIRT CPU %BUSY	Jun 14, 2014 00:00 VIRT CPU %BUSY	Jun 13, 2014 00:00 VIRT CPU %BUSY
Complex Integrated Online	xlnx1	0.16	0.16	0.14	0.18	0.15	0.13	0.08	0.08	0.09
HR System	zlnx1	2.00	10.16	15.35	17.41	15.53	30.81	5.79	15.15	11.66
	zlnx3	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Infrastructure	zlnx4	70.03*	68.73	69.54	69.36	67.99	67.99	67.34	67.75	67.87

Figure 25: Application Analytics: Distributed CPU Usage - Application Report

Source tables

[Presentation View].[Distributed Application Mapping]

[Presentation View].[AIX CPU Raw Hour Metrics]

[Presentation View].[Linux for System X CPU Hour Metrics]

[Presentation View].[Linux for System Z CPU Hour Metrics]

[Presentation View].[WIN OPERATING CPU Hour Metrics]

Drill through to the following reports

"Application Analytics: Distributed CPU Usage - Application Summary Report" on page 197

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

Report-specific prompts

Application name

Environment

Only show servers with available data

The Analytics Application: Distributed CPU usage - Application report table contains the following information:

Application Name
Environment
Function
Server
Detail CPU Utilization by different interval bands (hourly, daily, weekly, monthly, yearly)

Green is less than 70%, yellow is between 70% and less than 90%, and red is 90% or greater.

Application Analytics: MIPS Used - Application

This report shows the average MIPS used by each application or line-of-business grouping across all CPCs and LPARs. The average can be displayed as hourly, daily, weekly, monthly, or yearly.

The purpose of this report is to be able to compare the MIPS usage of all applications which are mapped in the application mapping table. This also shows you where the MIPS are being consumed by each application so you can focus your attention on which applications use the most MIPS and where they use those MIPS within the environment. You can drill down to see the functions or select fewer applications to compare. You can also see if there is a particular CPC or LPAR that has a large amount of MIPS being consumed by one or more application.

This report shows the CPU average busy MIPS for the z/OS application. You can define z/OS workload information in the mapping table.

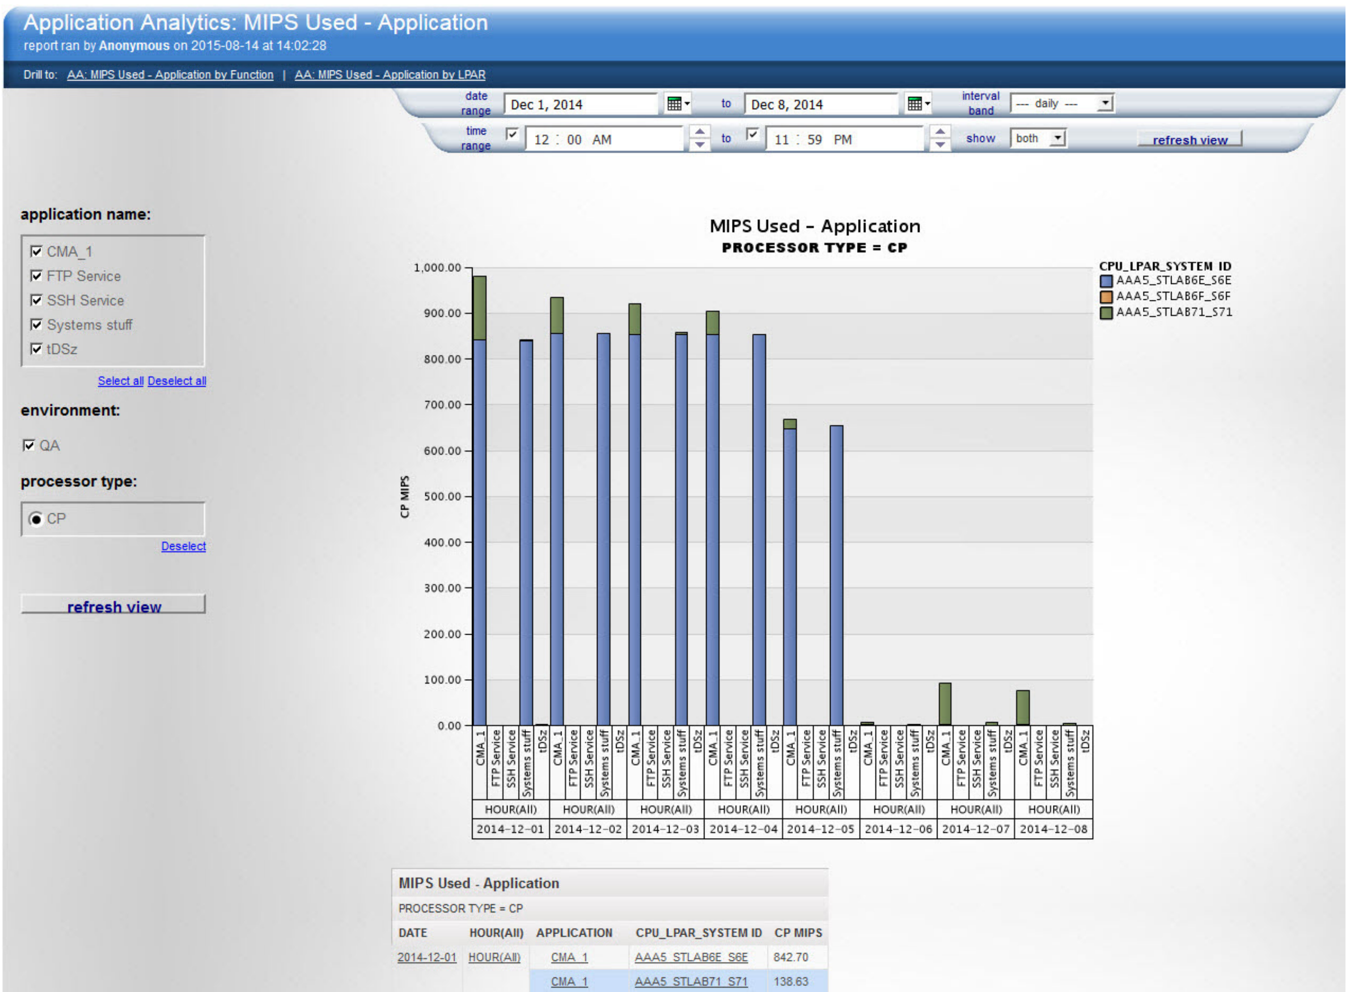

Figure 26: Application Analytics: MIPS Used - Application

Important: If you change the mapping after the forecast stream is executed, there may be discrepancies between the application MIPS on which the forecast is based and the observed application MIPS.

Source tables

[Appl zOS CPU MIPS Usage]

[Application zOS Mapping]

[Application LPAR Prompt]

[Application Util View]

[Application KPMZ Address Space]

[Application Prompt View]

Drill through to the following reports

"Application Analytics: MIPS Used - Application by Function" on page 201

"Application Analytics: MIPS Used - Application by LPAR" on page 206

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

show

Report-specific prompts

application name

environment

processor type

The Application Analytics: MIPS Used - Application table contains the following information:

DATE
Displays on all DMR levels

TIME
Displays on all DMR levels

APPLICATION NAME
Displays on application level

ENV-FUNCTION
Displays environment-function level

CPU_LPAR_SYSTEM ID
Displays **CPU Serial No**, **LPAR Name** and **System ID**

DATA TYPE
Display data type is the report class or job name

REPORT CLASS
Displays on the report class level

JOB NAME
Displays on job name level

CP MIPS/IIP MIPS/AAP MIPS
Displays MIPS values by different processor type: CP/IIP/AAP

CAPTURE RATIO
Displays on the report class or job name level

PROGRAM TYPE
Displays on the report class/job name level

TRANSACTION RATE
Displays transaction rate of the report class or job name level

TRANSACTIONS
Displays on the report class or job name level

Application Analytics: MIPS Used - Application by Function

This report shows the average or peak MIPS used by the selected application or line-of-business grouping across all CPCs and LPARs, according to function. The value can be displayed as hourly, daily, weekly, monthly, or yearly. This report also includes a forecast option.

In the application mapping table, you have the flexibility of defining functions within an application. The functions that are defined can represent sub-entities of the application that have their own business growth. This report enables you to view the MIPS usage by application or specific functions within an application. Additionally, you can also see where the MIPS are being consumed. This report provides a forecast capability for the application and where you can include only the functions that make the most sense for your forecast.

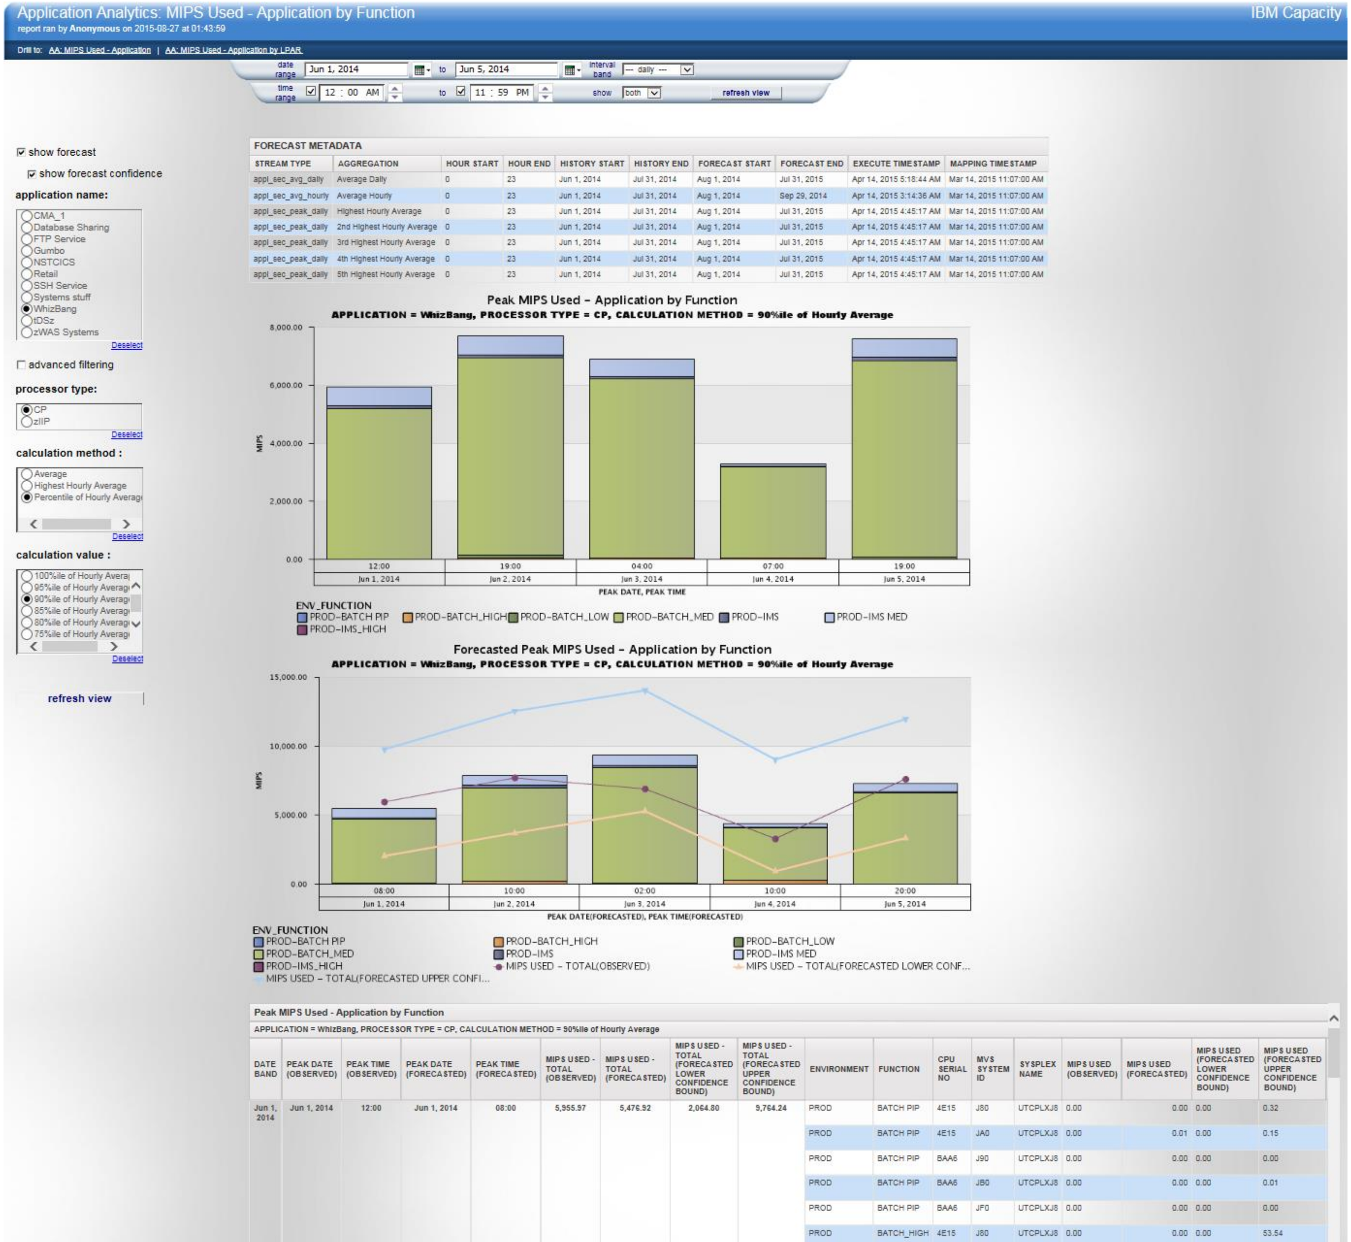

Figure 27: Application Analytics: MIPS Used - Application by Function

Important: If you change the mapping after the forecast stream is executed, there can be discrepancies between the application MIPS on which the forecast is based and the observed application MIPS.

Source tables

[Appl zOS CPU MIPS Usage]

[Application zOS Forecast]

[Application zOS Mapping]

[Application LPAR Prompt]

[Application Util View]

[Application KPMZ Address Space]

[Application Prompt View]

[Forecast Metadata]

Drill through to the following reports

“Application Analytics: MIPS Used - Application” on page 200

“Application Analytics: MIPS Used - Application by LPAR” on page 206

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

show

Report-specific prompts

show forecast

application name

advanced filtering, if selected, the following prompts are available:

- environment
- function
- cpu serial no
- lpar_system id

processor type

calculation method

There are 3 calculation methods available: Average, Highest Hourly Average, and Percentile of Hourly Average.

calculation value

This prompt is shown only if you select Highest Hourly Average or Percentile of Hourly Average as the calculation method.

If you select Highest Hourly Average, the values are:

- Highest Hourly Average
- 2nd Highest Hourly Average
- 3rd Highest Hourly Average
- 4th Highest Hourly Average
- 5th Highest Hourly Average

If you select Percentile of Hourly Average, the values are:

- 100%ile of Hourly Average
- 95%ile of Hourly Average
- 90%ile of Hourly Average
- 85%ile of Hourly Average
- 80%ile of Hourly Average
- 75%ile of Hourly Average
- 70%ile of Hourly Average

The Observed MIPS Used - Application by Function chart and table contains the following information. When the **show forecast** prompt is not selected, only observed values are displayed.

DATE BAND

The start date of the time interval, display only when calculation method is average.

TIME BAND
The start time of the time interval, display only when calculation method is average.

PEAK DATE

The date when the peak of the time interval happens.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

PEAK TIME

The time when the peak of the time interval happens.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

MIPS USED - TOTAL (OBSERVED)
The sum of the observed MIPS utilization of the application for the selected environment, function, cpu serial no, and lpar_system id, filtering for the selected processor type at each time interval or at the observed peak hour.

ENVIRONMENT

FUNCTION

CPU_SERIAL_NO
The zServer/CEC serial number.

SYSPLEX_NAME

MIPS USED (OBSERVED)
The detail value of the observed MIPS utilization of the application for each environment, function, cpu serial no, and lpar_system id, filtering for the selected processor type at each time interval or at the observed peak hour.

The Forecasted MIPS Used – Application by Function chart and table contains the information that is shown in the Observed MIPS Used - Application by Function chart and table as well as the following information. When the **show forecast** prompt is selected, both forecasted and observed values are displayed.

PEAK DATE (FORECASTED)

The date when the forecasted peak of the time interval happens.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

PEAK TIME (FORECASTED)

The time when the forecasted peak of the time interval happens.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

MIPS USED – TOTAL (FORECASTED)
The sum of the forecasted MIPS utilization of the application for the selected environment, function, cpu serial no, and lpar_system id, filtering for the selected processor type at each time interval or at the forecasted peak hour.

MIPS USED – TOTAL (FORECASTED LOWER CONFIDENCE BOUND)
The lower confidence bound of forecasted MIPS.

MIPS USED – TOTAL (FORECASTED UPPER CONFIDENCE BOUND)
The upper confidence bound of forecasted MIPS.

MIPS USED (FORECASTED)

The detail value of the forecasted MIPS utilization of the application for each environment, function, cpu serial no, and lpar_system id, filtering for the selected processor type at each time interval or at the forecasted peak hour.

If the application peak value forecast stream is used in the report, then this data item is not displayed.

MIPS USED (FORECASTED LOWER CONFIDENCE BOUND)
The lower confidence bound of forecasted MIPS. If the application peak value forecast stream is used in the report, then this data item is not displayed.

MIPS USED (FORECASTED UPPER CONFIDENCE BOUND)
The upper confidence bound of forecasted MIPS. If the application peak value forecast stream is used in the report, then this data item is not displayed.

FORECAST LEVEL
The aggregation level of the forecast stream that is used in the forecast MIPS calculation.

The Forecast Metadata table contains the following information. This table is displayed when the **show forecast** prompt is selected.

STREAM TYPE
The stream type that is used by the available forecast results set, for example appl_sec_avg_hourly. This means that the stream is for the application forecast, and it forecasts the CPU seconds at the hourly average aggregation level.

AGGREGATION
The aggregation level of the available forecast results set. The values are: Average Hourly, Average Daily, Highest Hourly Average, 2nd Highest Hourly Average, 3rd Highest Hourly Average, 4th Highest Hourly Average, and 5th Highest Hourly Average.

HOUR START
The start hour of the forecasted hour range of the available forecast results set.

HOUR END
The end hour of the forecasted hour range of the available forecast results set.

HISTORY START
The start date of the based historical data range of the available forecast results set.

HISTORY END
The end date of the based historical data range of the available forecast results set.

FORECAST START
The start date of the available forecast results set.

FORECAST END
The end date of the available forecast results set.

EXECUTE TIMESTAMP
The forecast execution timestamp of the available forecast results set.

MAPPING TIMESTAMP
The last updated timestamp of the mapping table on which the forecast is based.

The rules to choose the forecast stream aggregation level when **show forecast** is selected are:

If you select an hourly interval band:

- If you select the Average calculation method:
  1. If there is an hourly average stream available, then it is used.
  2. Otherwise, **no data available** is displayed.
- If you select the Highest Hourly Average or Percentile of Hourly Average calculation method:
  – The daily interval band is used and the stream, the logic for displaying the report, and the interval band prompts are set to the daily interval band as well.

If you select daily, weekly, monthly, or yearly:

- If you select the Average calculation method:
  1. If there is a suitable daily average stream available, then it is used. A suitable daily average stream means that the time range that you select is the same as the time range of the daily average stream.

2. Otherwise, if there is a suitable hourly average stream available then it is used. A suitable hourly average stream means that the time range that you select is within the time range of the hourly average stream.
3. Otherwise, **no data available** is displayed.

- If you select the Highest Hourly Average calculation method:
  - If you select **advanced filtering**, then the hourly average stream is used. A suitable hourly average stream means that the time range that you select is within the time range of the hourly average stream. Otherwise, **no data available** is displayed.
  - If you do not select **advanced filtering** and there is a suitable daily peak stream available, then the daily peak stream is used. Otherwise, if there is a suitable hourly average stream available, then it is used.

    A suitable daily peak stream means that the calculation value (highest, 2nd highest, etc.) that you select has the same aggregation number stream available, and that the time range that you select is the same as the time range of the daily peak stream.
- If you select the Percentile of Hourly Average calculation method:
  1. If there is a suitable hourly average stream available, then it is used.
  2. Otherwise, **no data available** is displayed.

Application Analytics: MIPS Used - Application by LPAR

This report shows the average MIPS used by all applications or line-of-business groupings on the selected LPAR. The average can be displayed as hourly, weekly, monthly, or yearly.

This report focuses on a LPAR and shows the MIPS usage of all applications in which it is running. If there is a planned outage change to a LPAR, then this report can help show the usage by application across time. You can plan ahead by finding out which applications may be impacted and what days and times show the least impact.

If there is a performance issue or a capacity issue on a LPAR, this report will help identify the usage patterns of each application.

This report shows the CPU average busy MIPS for the z/OS application. You can define z/OS workload information to the mapping table.

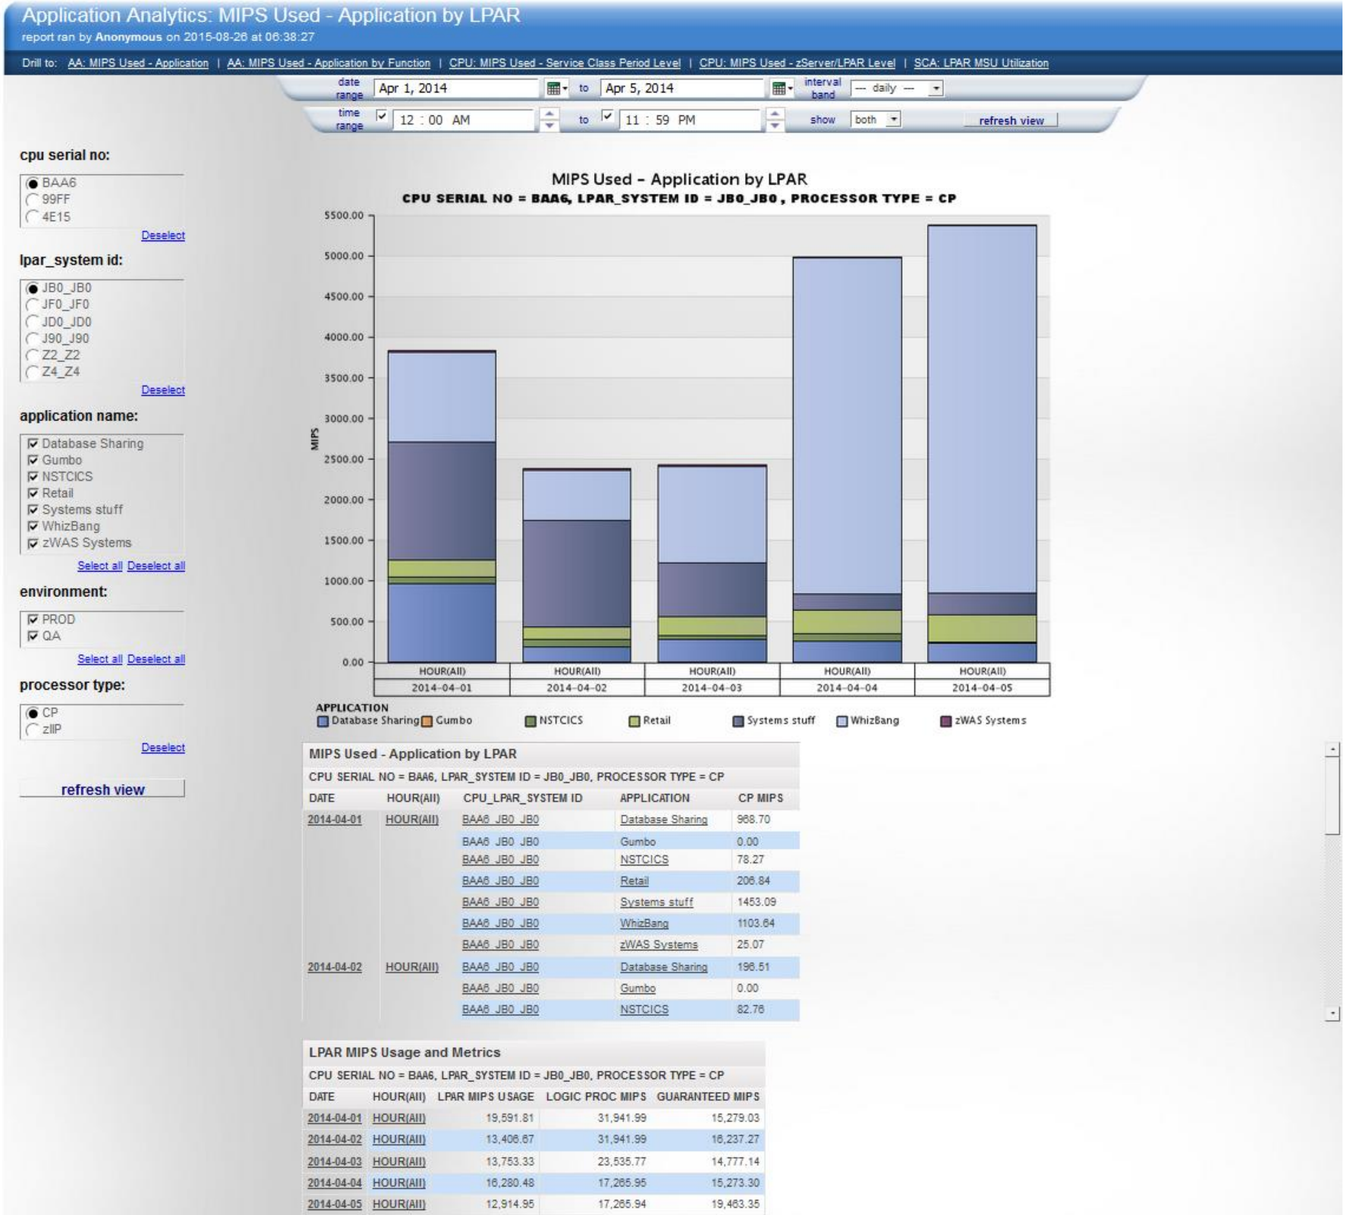

Figure 28: Application Analytics: MIPS Used - Application by LPAR

Source tables

[Appl zOS CPU MIPS Usage]

[Application zOS Mapping]

[Application LPAR Prompt]

[Application Util View]

[Application KPMZ Address Space]

[Application Prompt View]

Drill through to the following reports

"Application Analytics: MIPS Used - Application" on page 200

"Application Analytics: MIPS Used - Application by LPAR" on page 206

"CPU: MIPS Used - Service Class Period Level" on page 220

"CPU: MIPS Used - zServer/LPAR Level" on page 224

"SCA: LPAR MSU Utilization " on page 249

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

show

Report-specific prompts

cpu serial no

lpar_system id

application name

environment

processor type: CP is displayed in the **processor type** prompt as a default in some cases. Ensure that you click **refresh view** after you have completed all relevant fields.

The Application Analytics: MIPS used - Application by LPAR table contains the following information:

DATE
Displays on all DMR levels

TIME
Displays on all DMR levels

APPLICATION
Displays on application level

ENV-FUNCTION
Displays environment-function level

DATA TYPE
Display data type is report class or job name

REPORT CLASS OR JOB NAME
Displays on report class or job name level

CP MIPS/IIP MIPS/AAP MIPS
Displays MIPS values by different processor type: CP/IIP/AAP

CAPTURE RATIO
Display on report class or job name level

TRANSACTION COUNTS
Displays transaction counts on report class or job name level

TRANSACTION RATE
Displays transaction rate report class or job name level

Application Analytics: MIPS Used - Report Class by Job Name

This report shows the average MIPS used for a detailed list of the job names that were mapped with Workload Manager to a Report Class. The average can be displayed as hourly, daily, weekly, monthly, or yearly.

Report classes are used to map to functions within an application. For instance, you can look at an application usage over several days, then select a single day to view the hourly usage and drill into it to determine if something has changed.

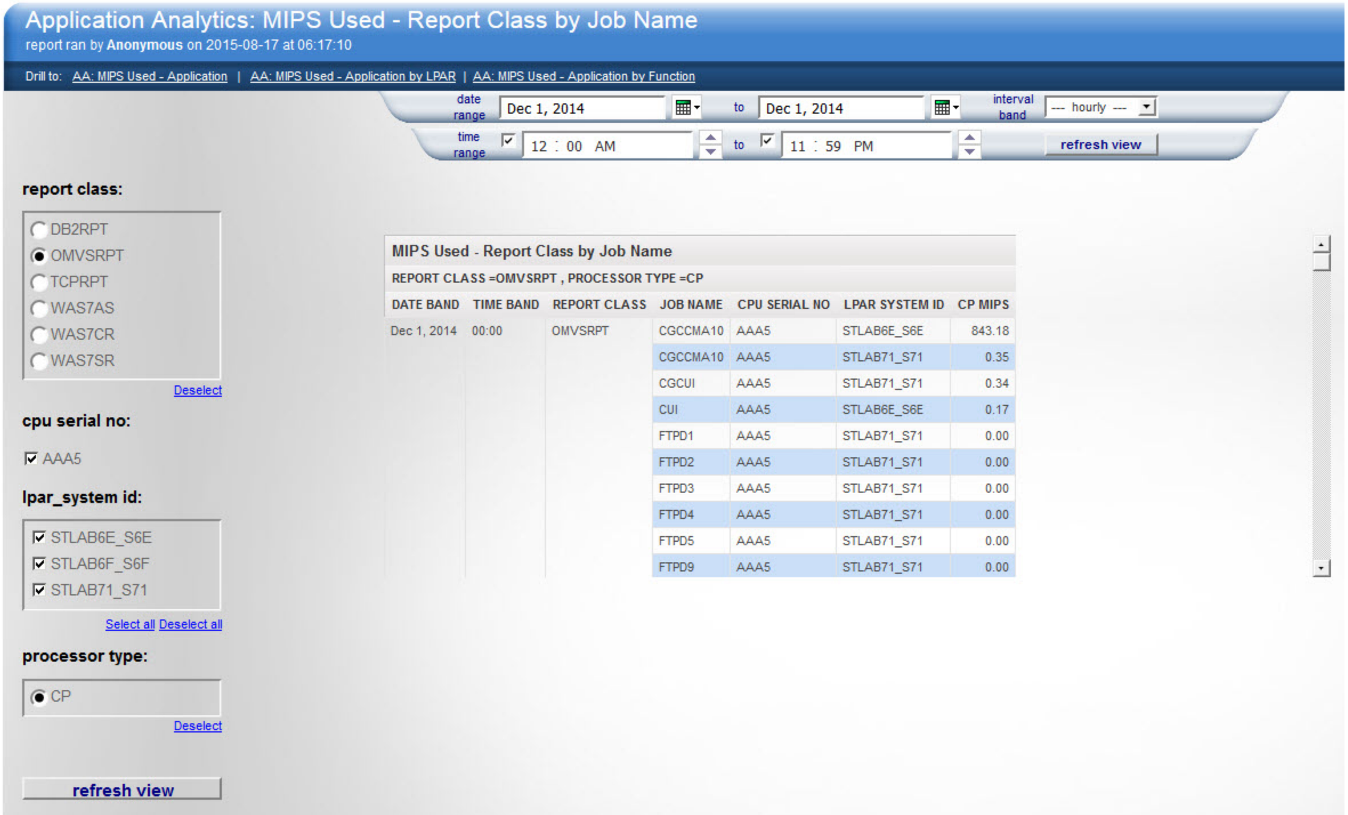

Figure 29: Application Analytics: MIPS Used - Application by LPAR

Source tables

[Appl zOS CPU MIPS Usage]

[Application zOS Mapping]

[Application LPAR Prompt]

[Application Util View]

[Application KPMZ Address Space]

[Application Prompt View]

Drill through to the following reports

"Application Analytics: MIPS Used - Application" on page 200

"Application Analytics: MIPS Used - Application by Function" on page 201

Standard prompts:

date range

time range

interval band: hourly, daily, weekly, monthly, yearly

Report-specific prompts

report class

cpu serial no

lpar_system id

processor type

The Analytics Application: MIPS Used by Report Class by Job Name chart or table contains the following information:

DATE

Displays on all DMR levels

TIME
: Displays on all DMR levels

REPORT CLASS
: Displays report class the user selected

JOB NAME
: Displays job name of one report class the user selected

CPU SERIAL NO
: cpu serial no

LPAR SYSTEM ID
: lpar name_system id

CP MIPS/IIP MIPS/AAP MIPS
: Displays MIPS values by different processor type: CP/IIP/AAP

CPU: AIX - CPU Usage

This report shows the percentage of CPU used from both a virtual and physical perspective.

The server might be virtualized or it might be a dedicated physical server. Showing both virtual and physical servers lets you know what type of server it is and how it is using its resources.

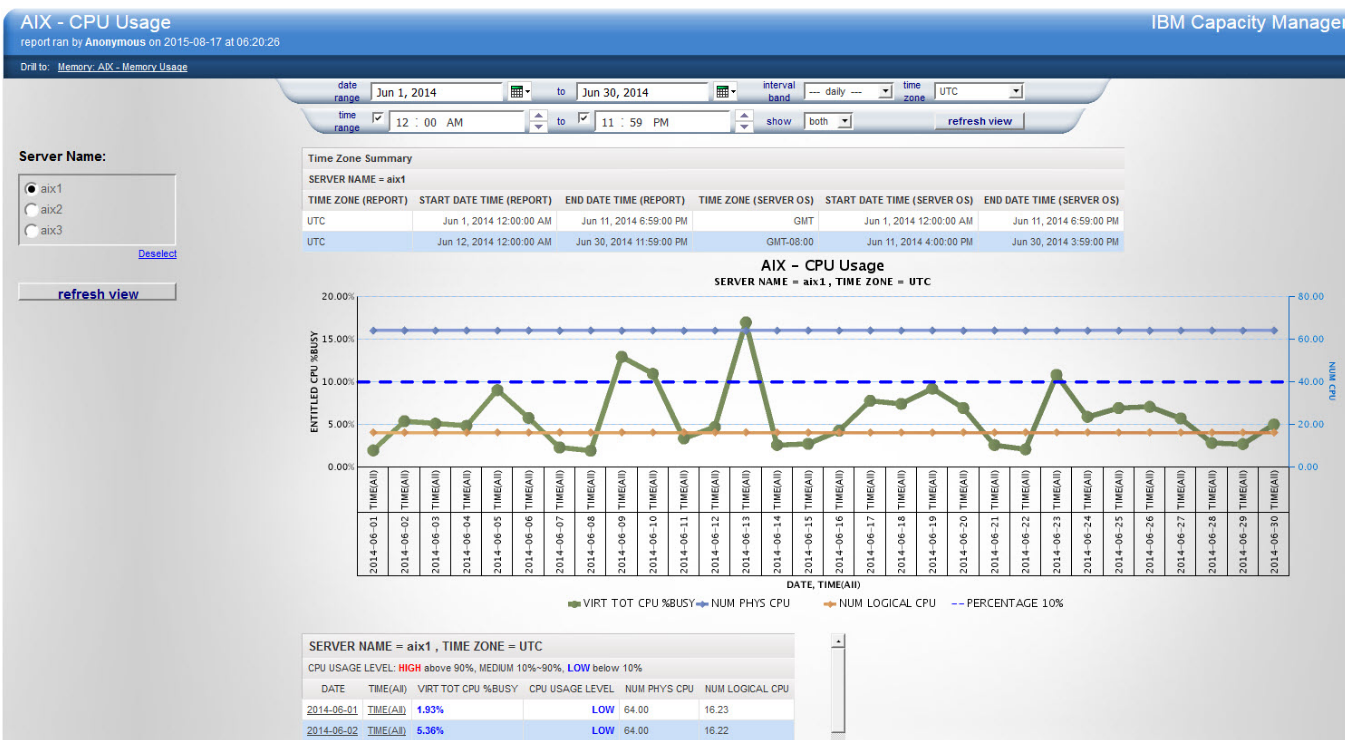

Figure 30: AIX - CPU Usage example

Source table
: [Presentation View].[AIX DIST_Date]

[Presentation View].[AIX DIST_Hour]

[Presentation View].[AIX DIST TIME]

[Presentation View].[AIX CPU Fact]

[Presentation View].[AIX IMAGE]

Drill through to the following report
: "Memory: AIX - Memory Usage" on page 240

Standard prompts:
date range

time range

interval band

time zone

show

Report-specific prompts
Server Name

The AIX - CPU Usage report contains the following information:

server name
The name of the server against which the report is run.

NUM PHYS CPU
The number of active virtual processors.

NUM LOGICAL CPU
The logical number of CPUs that AIX can dispatch on.

VIRT TOT CPU %BUSY
Total percent CPU busy of the Logical CPU count. This is the sum of **Virt SYS CPU % Busy** and **Virt USR CPU % Busy**.

PERCENTAGE 10%
10% CPU busy line.

PERCENTAGE 90%
90% CPU busy line. If the maximum CPU Usage percentage is less than 90%, then only 10% CPU busy line is displayed.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

CPU: Latent Demand

Latent demand represents the work that is waiting to be dispatched on a CP where all of the other CPs are currently busy. This report shows you which systems are overloaded and, as a result, could potentially slow down your high-priority applications. The report allows you to be proactive and reroute this work to less-utilized systems for faster processing.

Using this report, you can see the In-and-Ready Count and the Capture Ratio for a system. When the In-and-Ready Count exceeds the number of logical processors that are assigned to a system, then you have latent demand because work is waiting to be dispatched to a CP that is not already busy.

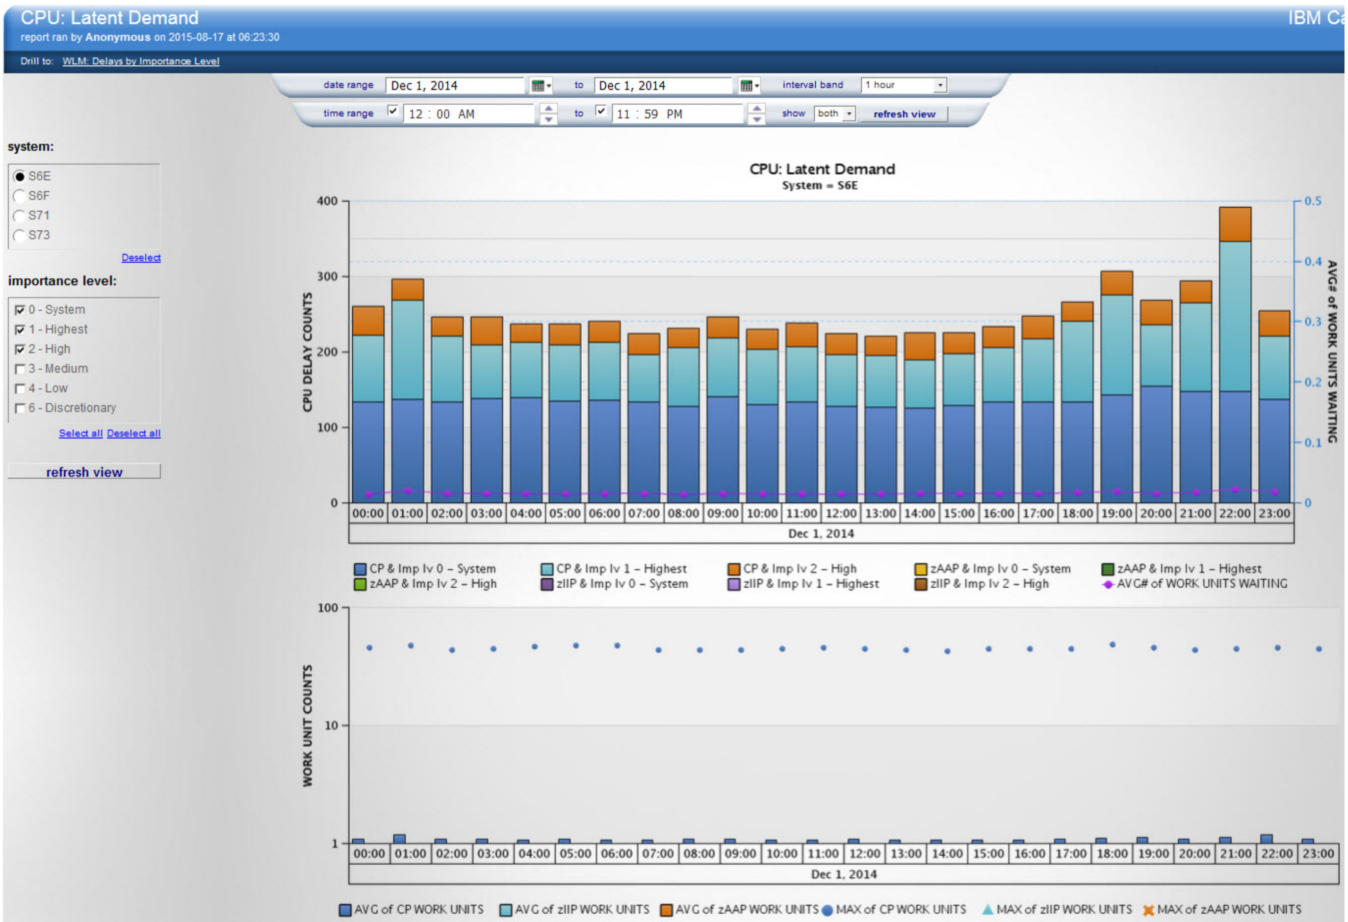

Figure 31: Latent Demand example

Source tables

[Presentation View].[z/OS System Level Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports

WLM: Delays by Importance Level

Standard prompts:

date range

time range

interval band

show

Report-specific prompts

system

importance level

This report contains the following information:

system
The z/OS system id.

importance level
The WLM importance level that is assigned to a service class period.

processor type
The engine type: CP, zIIP, zAAP.

AVG # OF ADDRESS SPACES WAITING
The weighted average number of address spaces that are swapped in and ready to be dispatched.

CPU DELAY COUNTS
Number of CPU delay samples. Calculated as the sum of R723CCDE.

AVG of CP WORK UNITS
Average number of work units for general-purpose processors over the interval. Calculated as the average of SMF70CTT.

AVG of zIIP WORK UNITS
Average number of work units for zIIP processors over the interval. Calculated as the average of SMF70ETT.

AVG of zAAP WORK UNITS
Average number of work units for zAAP processors over the interval. Calculated as the average of SMF70DTT.

MAX of CP WORK UNITS
Maximum number of work units for general-purpose processors over the interval. Calculated as the minimum of SMF70CMM.

MAX of zIIP WORK UNITS
Maximum number of work units for zIIP processors over the interval. Calculated as the minimum of SMF70EMM.

MAX of zAAP WORK UNITS
Maximum number of work units for zAAP processors over the interval. Calculated as the minimum of SMF70DMM.

AVG# of WORK UNITS WAITING
The weighted average number of any work is queued and waiting to be processed.

ONLINE PROCESSORS
The number of engines online to a system on which work can be dispatched.

CPU: Linux for System x - CPU Usage

This report shows the percentage of CPU used from both a virtual and physical perspective.

The server might be virtualized or it might be a dedicated physical server. Showing both virtual and physical servers lets you know what type of server it is and how it is using its resources.

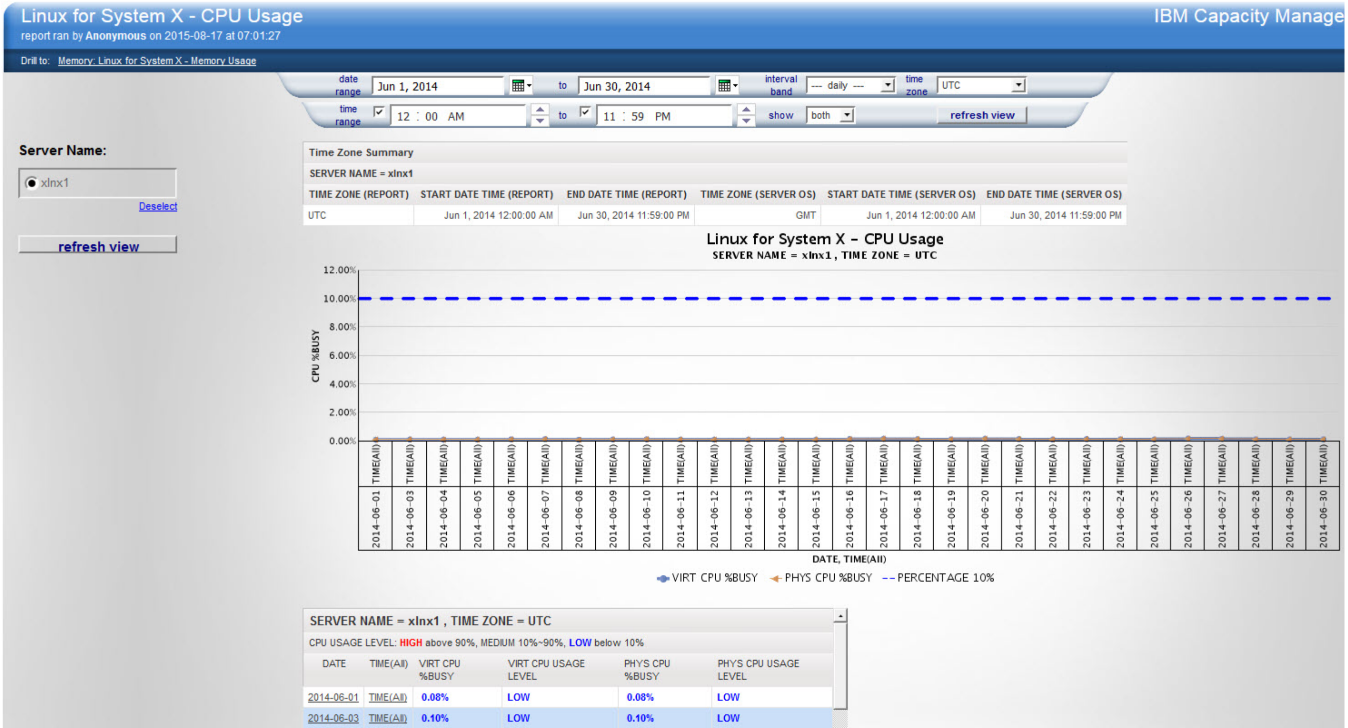

Figure 32: Linux for System X - CPU Usage example

Source table

[Presentation View].[Linux for System X DIST Date]

[Presentation View].[Linux for System X DIST Hour]

[Presentation View].[Linux for System X DIST Time]

[Presentation View].[Linux for System X CPU Fact]

[Presentation View].[Linux for System X Image]

Drill through to the following report

"Memory: Linux for System x - Memory Usage" on page 244

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Linux for System x - CPU Usage report contains the following information:

server name

The name of the server against which the report is run.

VIRT CPU %BUSY

The percentage of CPU resources that are used by this operating system in relation to the CPU capacity of the virtual server, without taking capping effects or competition with other partitions into account.

PHYS CPU %BUSY

The percentage of physical CPU resources that are used by this operating system

PERCENTAGE 10%

10% CPU busy line.

PERCENTAGE 90%
: 90% CPU busy line. If the maximum CPU Usage percentage is less than 90%, then only 10% CPU busy line is displayed.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
: The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
: The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
: The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
: The local time of the selected server. The format is GMT -5:00

  Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
: The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
: The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

CPU: Linux for System z - CPU Usage

This report shows the percentage of CPU used from both a virtual and physical perspective.

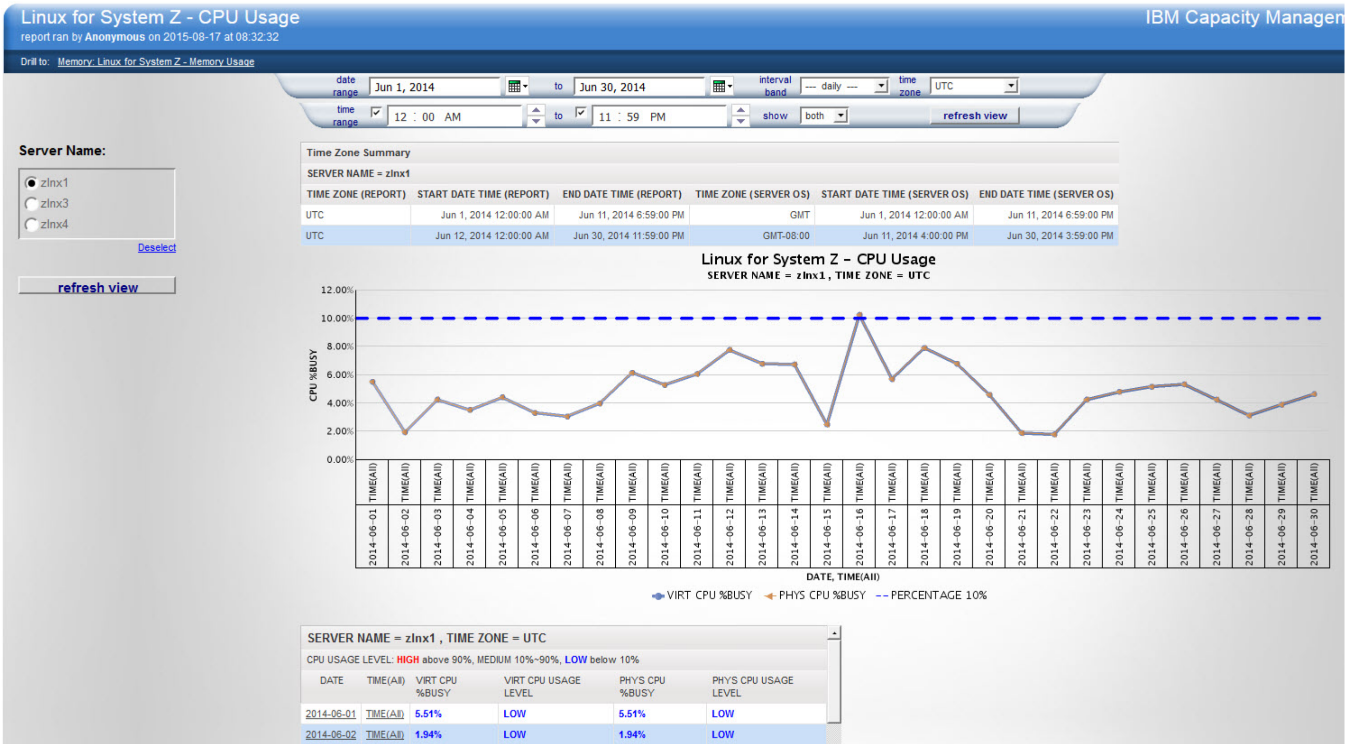

Figure 33: Linux for System z - CPU Usage example

The server might be virtualized or it might be a dedicated physical server. The view that shows both virtual and physical lets you know what type of server it is and how its resources are being used.

Source table

[Presentation View].[Linux for System z DIST Date]

[Presentation View].[Linux for System z DIST Hour]

[Presentation View].[Linux for System z DIST Time]

[Presentation View].[Linux for System z CPU Fact]

[Presentation View].[Linux for System z Image]

Drill through to the following report

"Memory: Linux for System z - Memory Usage" on page 246

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Linux for System z - CPU Usage report contains the following information:

server name

The name of the server against which the report is run.

VIRT CPU % BUSY
The percentage of CPU resources that are used by this operating system in relation to the CPU capacity of the virtual server, without taking capping effects or competition with other partitions into account.

PHYS CPU % BUSY
The percentage of physical CPU resources that are used by this operating system.

PERCENTAGE 10%
10% CPU busy line.

PERCENTAGE 90%
90% CPU busy line. If the maximum CPU Usage percentage is less than 90%, then only 10% CPU busy line is displayed.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

CPU: LPAR Weight Optimization Run Result

This report shows the available LPAR Weight Optimization runs. It also shows the input settings and optimized weight results for a run selected for each CPC and processor type.

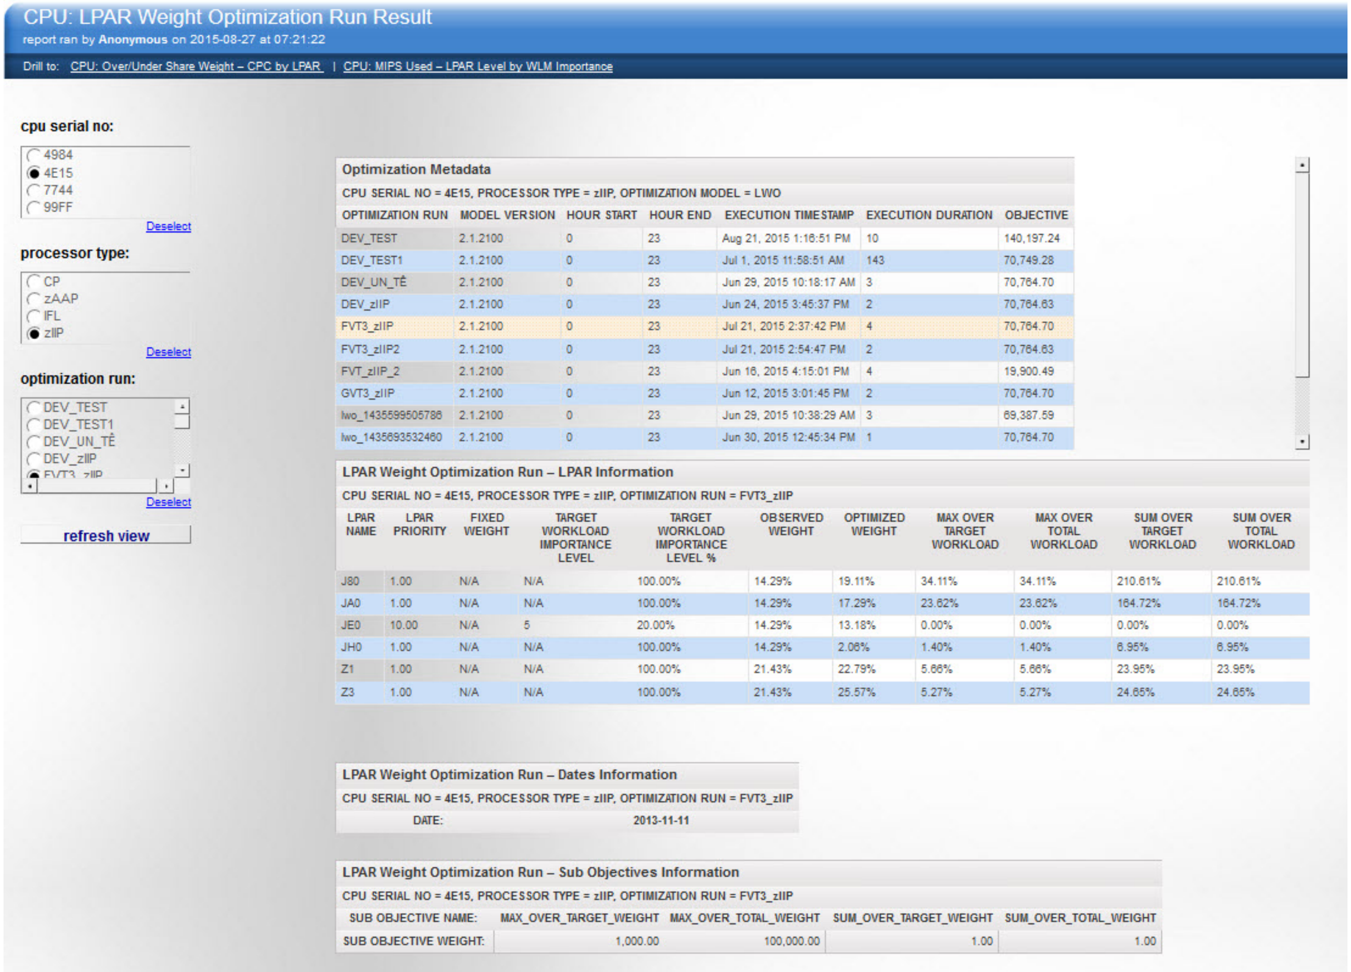

Figure 34: CPU: LPAR Weight Optimization Run Result example

Source tables

[Presentation View].[Optimization Metadata]

[Presentation View].[LPAR Weight Optimization Metrics]

[Presentation View].[LPAR Weight Optimization Parameters]

[Presentation View].[LPAR Weight Optimization Detail]

Drill through to the following reports

"CPU: MIPS Used - LPAR Level by WLM Importance" on page 229

"CPU: Over/Under Share Weight - CPC by LPAR" on page 234

Report-specific prompts

cpu serial no

processor type

optimization run

The Optimization Metadata table contains the following information:

OPTIMIZATION RUN

The name that is used for the execution of the LPAR weight optimization.

MODEL VERSION
The version of the optimization model.

HOUR START
The start hour of the dates that are used in the LPAR weight optimization run.

HOUR END
The end hour of the dates that are used in the LPAR weight optimization run.

EXECUTION TIMESTAMP
The execution start time of this optimization run.

EXECUTION DURATION
The length of time that the optimization run took to run.

OBJECTIVE
The total objective value of this optimization model.

The LPAR Weight Optimization Run - LPAR Information table contains the following information:

LPAR NAME
The name of the LPAR.

LPAR PRIORITY
The priority of the LPAR. Higher numbers mean that an LPAR has a higher priority.

FIXED WEIGHT
The fixed weight value for the LPAR, if specified. If the value is "N/A", it means that the weight is not fixed and it is optimized by the LPAR weight optimization run.

TARGET WORKLOAD IMPORTANCE LEVEL

The importance level of the target workload's boundary of this LPAR.

A NULL value means that the WLM importance level does not need to be considered. The target workload is considered to be the total workload multiplied by the value of the TARGET WORKLOAD IMPORTANCE LEVEL % column.

A value of 0 to 6 means that the workloads under the importance level number are considered as the target workload for this LPAR. Also, it means that a certain percent (the value of the TARGET WORKLOAD IMPORTANCE LEVEL % column) of the importance level is considered as the target workload for this LPAR.

TARGET WORKLOAD IMPORTANCE LEVEL %

The percentage that is used in the target workload calculation.

When TARGET WORKLOAD IMPORTANCE LEVEL is null, then it is the percentage of the total workload that should be considered as the target workload.

When TARGET WORKLOAD IMPORTANCE LEVEL is a number from 0 to 6, then it is the percentage of the TARGET WORKLOAD IMPORTANCE LEVEL workload that should be considered as the target workload.

OBSERVED WEIGHT
The average observed weight of the dates and hour range considered in the optimization run.

OPTIMIZED WEIGHT
The optimized weight that is the result from the LPAR Weight Optimization run.

MAX OVER TARGET WORKLOAD
The maximum value that the target workload is over the optimized weight for this LPAR. The value is a percentage of the capacity of the shared pool of this processor type.

MAX OVER TOTAL WORKLOAD
The maximum value that the total workload is over the optimized weight for this LPAR. The value is a percentage of the capacity of the shared pool of this processor type.

SUM OVER TARGET WORKLOAD
The sum of the amount that the target workload is over the optimized weight for each time interval for this LPAR. The value is a percentage of the capacity of the shared pool of this processor type. Because it is a sum value for each time interval, it is possible that the value can be larger than 100%.

SUM OVER TOTAL WORKLOAD
The sum of the amount that the total workload is over the optimized weight for each time interval for this LPAR. The value is a percentage of the capacity of the shared pool of this processor type. Because it is a sum value for each time interval, it is possible that the value can be larger than 100%.

The LPAR Weight Optimization Run – Dates Information table contains the following information.

DATE
The dates used for the optimization run.

The LPAR Weight Optimization Run – Sub Objectives Information table contains the following information.

SUB OBJECTIVE NAME
The name of each sub objective for the LPAR Weight Optimization model.

SUB OBJECTIVE WEIGHT
The weight or punishment value for each sub objective for the optimization run.

CPU: MIPS Used - Service Class Period Level

This report provides a comprehensive view of all of the applications that are running across the entire systems landscape. This helps you quickly rebalance system allocation and to focus on the most critical applications.

This report helps you assess whether you might need to purchase more capacity, and when that extra capacity is needed. However, what might appear as running out of capacity, might not necessarily mean that more capacity is needed—it does mean that you must review your usage and ensure that you are optimizing your usage effectively. The objective is to run 100% busy 24/7. The key to achieving that objective is ensuring that you can prioritize your workloads and move applications around efficiently. For example, scheduling lower priority jobs for off-peak hours. This ensures that the priority applications have the capacity they need when they need it.

This report analyzes the workloads (service classes) that are driving CPU usage on a system by showing MIPS used at the Workload Manager service class period level. The total CPU time that is used by all WLM service class periods on a system is known as captured CPU time.

For CPs, the zIIP / zAAP eligible MIPS are also shown. For zIIP / zAAP, the MIPS that were eligible are also shown.

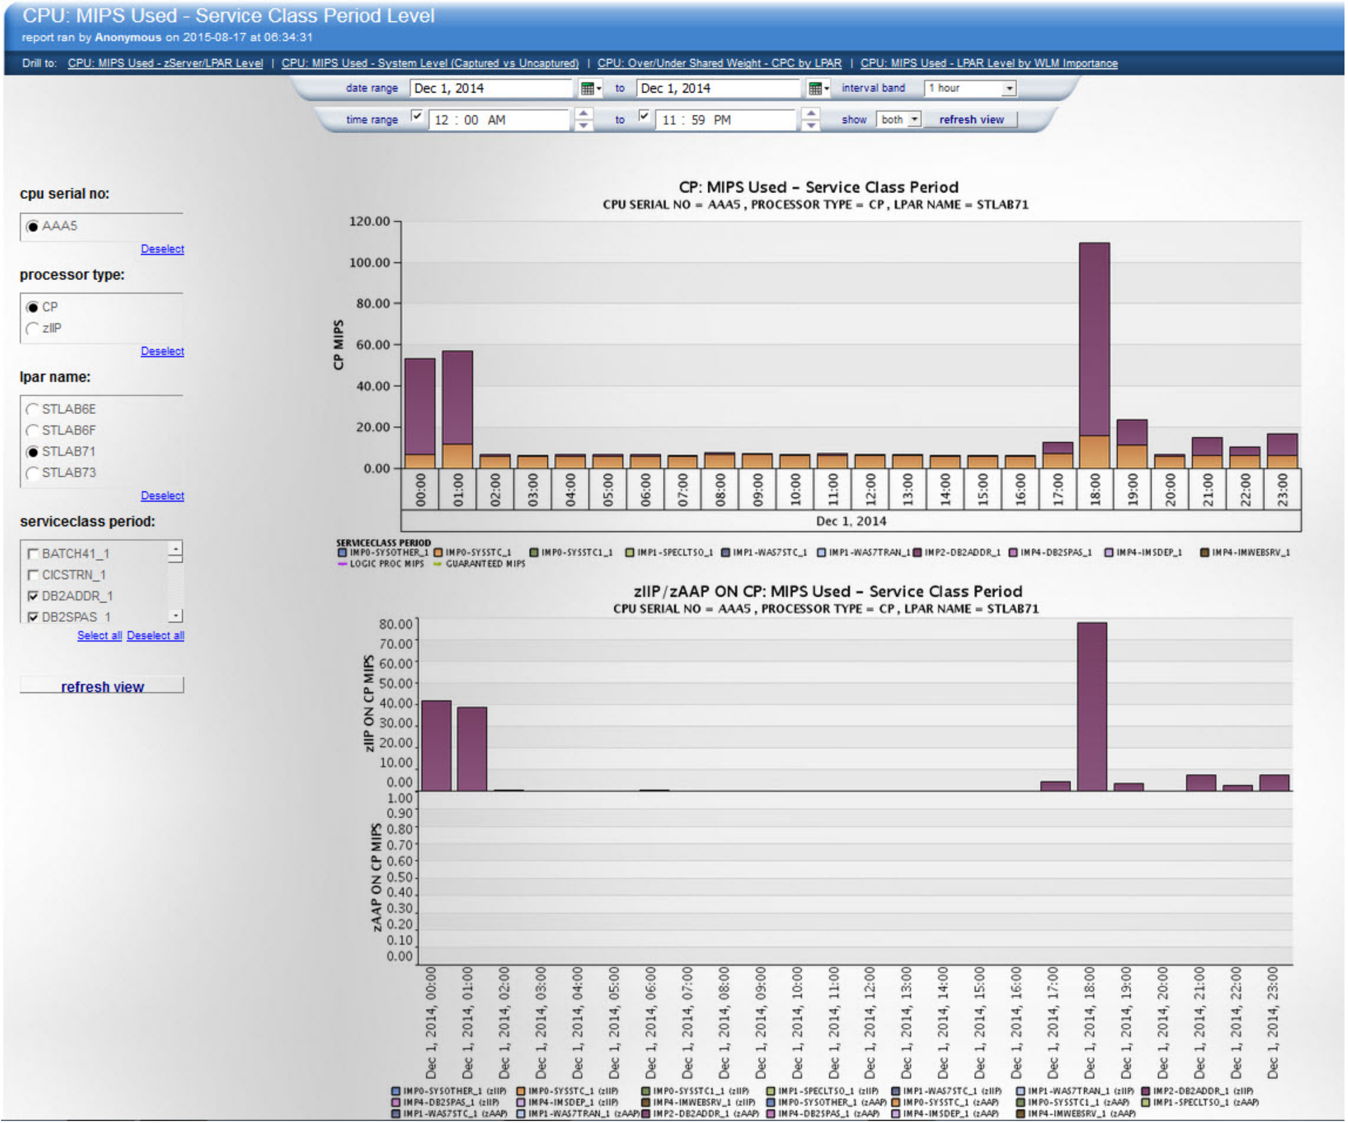

Figure 35: MIPS Used - Service Class Period Level example

Source table

[Presentation View].[LPAR Level Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports

"CPU: MIPS Used - zServer/LPAR Level" on page 224

"CPU: MIPS Used - System Level (Captured vs Uncaptured)" on page 222

"CPU: MIPS Used - LPAR Level by WLM Importance" on page 229

"CPU: Over/Under Share Weight - CPC by LPAR" on page 234

Standard prompts:

date range

time range

interval band

show

Report-specific prompts

cpu serial no

processor type

lpar name

serviceclass period

This report contains the following information:

CPU SERIAL NO
The zServer/CEC serial number.

processor type
The engine type: CP, zIIP, or zAAP. The list contains only installed processor types.

lpar name
The name of the logical partition to which the specified processor type is assigned.

LOGIC PROC MIPS

The logical processors MIPS value is the MIPS capacity of the logical processors that are assigned to this LPAR

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

WEIGHT %
The per cent of the shared pool that is guaranteed to the LPAR, if the LPAR uses shared processors. If the LPAR uses dedicated processors, then the value is shown as N/A.

GUARANTEED MIPS

The guaranteed capacity in MIPS on this LPAR. The value is calculated according to the weight assigned to this LPAR.

If dedicated processors are assigned to this LPAR, then this line does not appear in the chart and the value is shown as N/A in the table.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

serviceclass period
The WLM service class name and period. The concatenation of columns SERVICE_CLASS and SERV_CLASS_PERIOD.

CAPTURED MIPS USED
The CPU time that is consumed by serviceclass_period, converted to MIPS used.

zIIP ON CP MIPS
The CPU time that is used for IIP on CP processing, converted to MIPS used.

zAAP ON CP MIPS
The CPU time that is used for IFA on CP processing, converted to MIPS used.

zIIP ELIGIBLE MIPS
The CPU time that is used for IIP on CP processing, converted to MIPS used.

zAAP ELIGIBLE MIPS
The CPU time that is used for IFA on CP processing, converted to MIPS used.

CPU: MIPS Used - System Level (Captured vs Uncaptured)

This report provides insights into potential MIPS overages. This can help you to determine the need to tune and adjust your system processes and avoid overcapacity by allocating MIPS back to business applications. Aside from the potential cost savings, a regular assessment, based on this report, can help you maintain your SLA by keeping up the performance and availability level.

This report analyzes a system’s capture ratio to determine whether the CPU time that is used by system-related processes (uncaptured CPU time) is too high. It shows both the captured and uncaptured CPU MIPS used at the system level. It also shows the system's capture ratio. A notable increase or decrease in a system's capture ratio can identify a system problem. It is recommended that you closely monitor a system's capture ratio after you apply any CPU hardware changes, major operating system upgrades, or software upgrades.

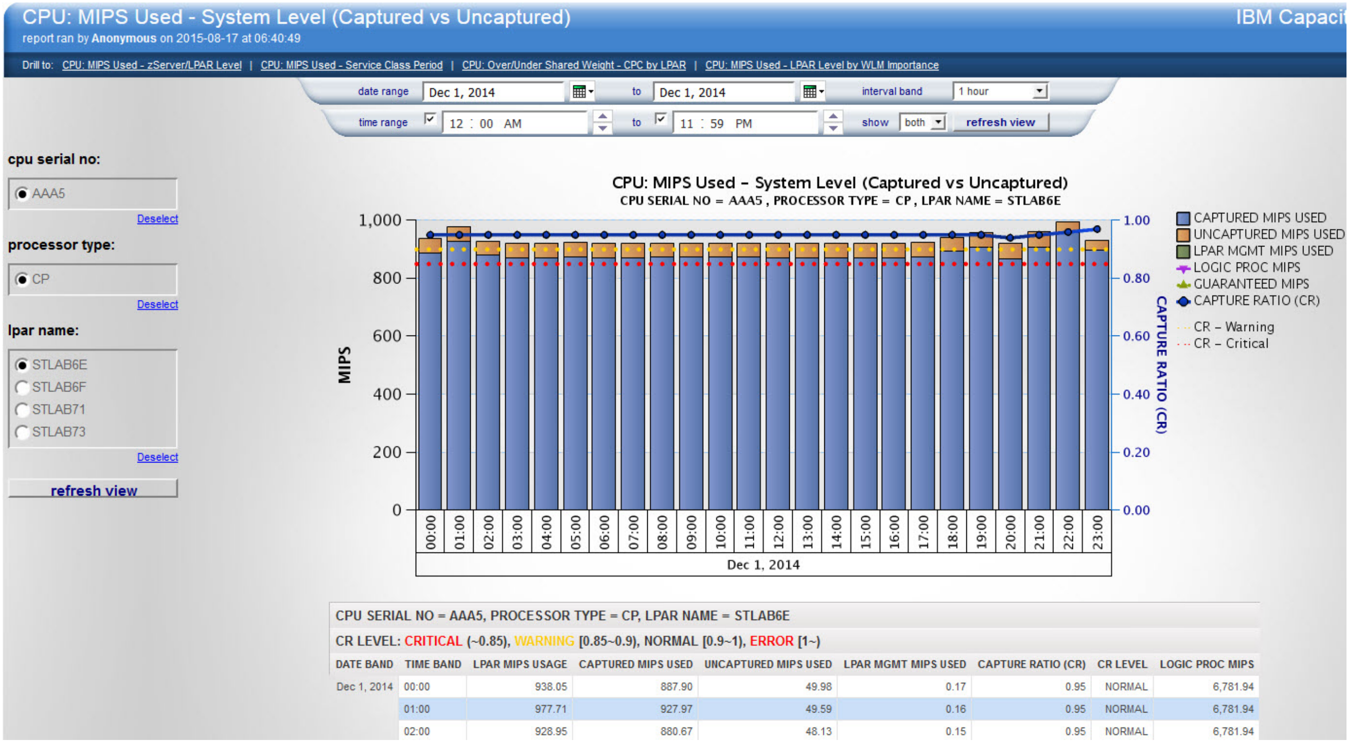

Figure 36: MIPS Used - System Level (Captured vs Uncaptured) example

Source tables

[Presentation View].[LPAR Level Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports

"CPU: MIPS Used - zServer/LPAR Level" on page 224

"CPU: MIPS Used - LPAR Level by WLM Importance" on page 229

"CPU: MIPS Used - Service Class Period Level" on page 220

"CPU: Over/Under Share Weight - CPC by LPAR" on page 234

Standard prompts:

date range

time range

interval band

show

Report-specific prompts

cpu serial no

processor type

lpar name

The CPU: MIPS Used - System Level (Captured vs Uncaptured) report contains the following information:

cpu serial no

The zServer/CEC serial number.

processor type

The engine type: CP, zIIP, or zAAP. The list contains only installed processor types.

system

The z/OS system id.

LPAR MGMT MIPS USED
The difference between LPAR Dispatch CPU Time and LPAR Effective Dispatch CPU Time, converted to MIPS used.

CAPTURED MIPS USED
The total CPU time that is captured in the Type72 workload data, converted to MIPS used.

UNCAPTURED MIPS USED
The difference between LPAR Effective Dispatch CPU Time and the total CPU time that is captured in the Type72 workload data, converted to MIPS used.

CAPTURE RATIO (CR)
The ratio of total CPU time that is captured in the Type72 workload data to LPAR Dispatch CPU Time.

CR - Warning
The capture ratio warning threshold line = 0.90.

CR- Critical
The capture ratio critical threshold line = 0.85.

LOGIC PROC MIPS

The logical processors MIPS value is the MIPS capacity of the logical processors that are assigned to this LPAR

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

WEIGHT %
The per cent of the shared pool that is guaranteed to the LPAR, if the LPAR uses shared processors. If the LPAR uses dedicated processors, then the value is shown as N/A.

GUARANTEED MIPS

The guaranteed capacity in MIPS on this LPAR. The value is calculated according to the weight assigned to this LPAR.

If dedicated processors are assigned to this LPAR, then this line does not appear in the chart and the value is shown as N/A in the table.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

CPU: MIPS Used - zServer/LPAR Level

This report shows the observed average or peak CPU MIPS used at the zServer/CPC level and the forecasted average or peak CPU MIPS used along with their confidence interval. The average CPU MIPS used (observed and forecasted) and peak CPU MIPS used (observed and forecasted) can be displayed for each processor type (CP, zIIP, zAAP, IFL, or ICF) at each interval band and is broken down to the LPAR level. The forecasted peak CPU MIPS used can be displayed for each processor type at each interval band on the CPC level.

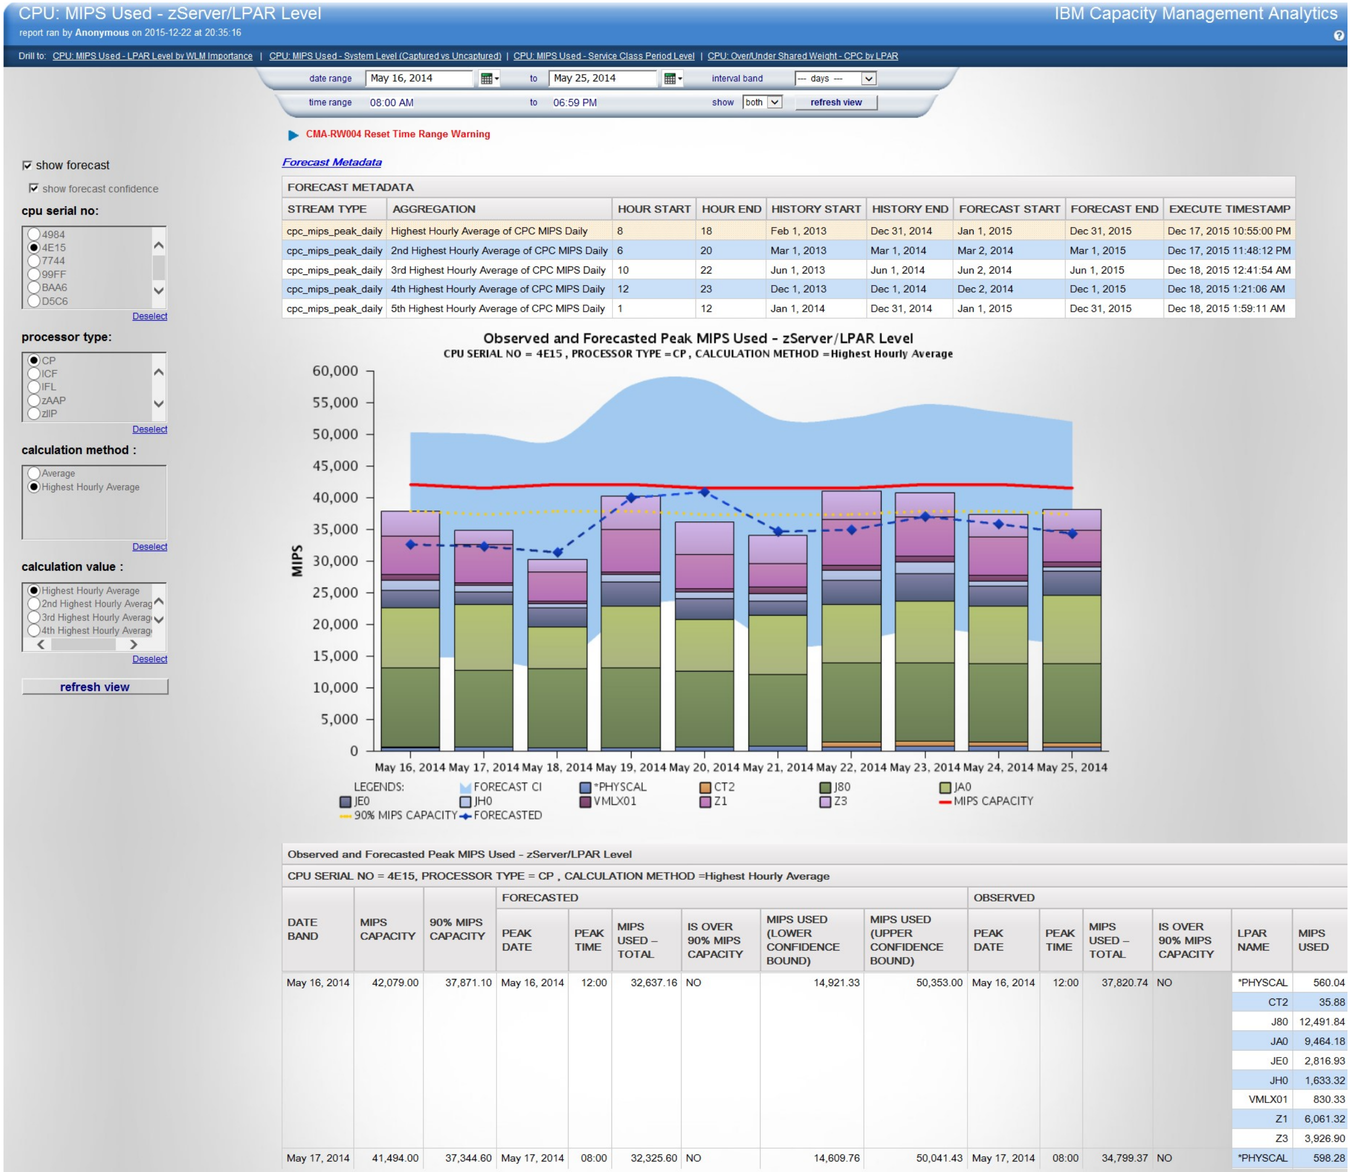

Figure 37: MIPS Used - zServer/LPAR Level example

Source table

[Presentation View].[LPAR Level Metrics]

[Presentation View].[LPAR Level Forecast Metrics]

[Presentation View].[CPC Level Forecast Metrics]

Drill through to the following reports

"CPU: MIPS Used - LPAR Level by WLM Importance" on page 229

"CPU: MIPS Used - System Level (Captured vs Uncaptured)" on page 222

"CPU: MIPS Used - Service Class Period Level" on page 220

"CPU: Over/Under Share Weight - CPC by LPAR" on page 234

Standard prompts:

date range

time range

interval band (If an interval band of less than one day is selected, and peak usage is shown, then the calculation uses an interval band of one day for the calculation.)

show

Report-specific prompts

show forecast

cpu serial no

processor type

calculation method

There are three calculation methods available: Average, Highest Hourly Average, and Percentile of SMF Interval.

If **show forecast** is selected, then Percentile of SMF Interval is not available.

calculation value

This prompt is shown only if you select Highest Hourly Average or Percentile of SMF Interval as the calculation method.

If you select Highest Hourly Average, the values are:

- Highest Hourly Average
- 2nd Highest Hourly Average
- 3rd Highest Hourly Average
- 4th Highest Hourly Average
- 5th Highest Hourly Average

If you select Percentile of SMF Interval, the values are:

- 100%ile of SMF Interval
- 95%ile of SMF Interval
- 90%ile of SMF Interval
- 85%ile of SMF Interval
- 80%ile of SMF Interval
- 75%ile of SMF Interval
- 70%ile of SMF Interval

lpar name

If Highest Hourly Average is selected as the calculation method, then **lpar name** is not shown.

The Observed Average MIPS Used - zServer/LPAR Level and Observed Peak MIPS Used -zServer/LPAR Level charts or tables contain the following information. If **show forecast** is cleared, then only observed values are displayed.

cpu serial no
The zServer/CEC serial number.

processor type
The engine type: CP, zIIP, zAAP, IFL, or ICF. The list contains only installed processor types.

lpar name
The name of the logical partition to which the specified processor type is assigned.

DATE BAND
The start date of the time interval.

TIME BAND
The start time of the time interval.

This value displays only when the calculation method is Average.

PEAK DATE
The date when the peak period occurs.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of SMF Interval.

PEAK TIME
The start time of the peak period.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of SMF Interval.

MIPS Capacity
The total MIPS capacity available for the selected processor type.

If MIPS capacity has changed during the time band, then the maximum value is displayed. In this case, there is a "(ii)" marker behind the MIPS Capacity value and a footnote is displayed on the top of this table stating: "(ii) the value is the maximum value during the time band".

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

90% MIPS Capacity
The 90% MIPS capacity line (warning) for the selected processor type.

90% MIPS Capacity is displayed only if MIPS Capacity is also displayed.

MIPS Used
The amount of available relative processor capacity consumed for the selected processor type.

If an LPAR is assigned with dedicated processors, and this LPAR does not collect SMF data, then CPU_DISPATCH_SEC is used which shows the dedicated engine's MIPS capacity but not the real utilization of this LPAR. In this case, there is a "(i)" marker behind the MIPS Used value and a footnote is displayed on the top of this table stating: (i) the value is calculated by CPU_DISPATCH_SEC collected from other LPAR because this LPAR uses dedicated processors and does not collect SMF data.

FORECASTED
The total forecasted MIPS used for the selected processor type. This metric is derived from the CPU forecast model (SPSS).

The Forecasted Average MIPS Used - zServer/LPAR Level and the Observed and Forecasted Peak MIPS Used – zServer/LPAR Level charts and tables contain the information that is shown in the Observed Average MIPS Used - zServer/LPAR Level and Observed Peak MIPS Used - zServer/LPAR Level charts and tables as well as the following information. If **show forecast** is cleared, then only observed values are displayed.

PEAK DATE (FORECASTED)
The date when the forecasted peak of the time interval occurs.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

PEAK TIME (FORECASTED)
The time when the forecasted peak of the time interval occurs.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of Hourly Average.

MIPS USED - TOTAL (OBSERVED)
The total MIPS used for the selected processor type.

If an LPAR is assigned with dedicated processors, and this LPAR does not collect SMF data, then CPU_DISPATCH_SEC is used which shows the dedicated engine's MIPS capacity but not the real utilization of this LPAR. In this case, there is a "(i)" marker by the MIPS Used value and a footnote is displayed on the top of this table stating: "(i) the value is calculated by CPU_DISPATCH_SEC collected from other LPAR because this LPAR uses dedicated processors and does not collect SMF data."

MIPS USED (OBSERVED)
The total MIPS used for the selected processor type.

If an LPAR is assigned with dedicated processors, and this LPAR does not collect SMF data, then CPU_DISPATCH_SEC is used which shows the dedicated engine's MIPS capacity but not the real utilization of this LPAR. In this case, there is a "(i)" marker by the MIPS Used value and a footnote is displayed on the top of this table stating: "(i) the value is calculated by CPU_DISPATCH_SEC collected from other LPAR because this LPAR uses dedicated processors and does not collect SMF data."

MIPS USED (FORECASTED)
The total forecasted MIPS used for the selected processor type. This metric is derived from the CPU forecast model (SPSS).

MIPS USED (FORECAST LOWER CONFIDENCE BOUND)
The lower bound of the confidence interval for the MIPS USED forecasting.

MIPS USED (FORECAST UPPER CONFIDENCE BOUND)
The upper bound of the confidence interval for the MIPS USED forecasting.

Values that are marked with an asterisk (iii) are adjusted values. The values were originally negative in the forecasting result and are adjusted to 0.

Values that are marked with an asterisk (iv) are adjusted values. The values were originally extremes in the forecasted result and are adjusted to maximum forecasted value.

The Forecast Metadata table contains the following information. This table is displayed when the **show forecast** prompt is selected and Highest Hourly Average is selected as calculation method.

STREAM TYPE
The stream type that is used by the available forecast results set, for example lpar_mips_peak_daily. This means that the stream is for the LPAR MIPS forecast, and it forecasts the CPU seconds at the daily peak aggregation level.

AGGREGATION
The aggregation level of the available forecast results set. The values are: Highest Hourly Average of LPAR MIPS Daily, 2nd Highest Hourly Average of LPAR MIPS Daily, 3rd Highest Hourly Average of LPAR MIPS Daily, 4th Highest Hourly Average of LPAR MIPS Daily, and 5th Highest Hourly Average of LPAR MIPS Daily.

HOUR START
The start hour of the forecasted hour range of the available forecast results set.

HOUR END
The end hour of the forecasted hour range of the available forecast results set.

HISTORY START
The start date of the based historical data range of the available forecast results set.

HISTORY END
The end date of the based historical data range of the available forecast results set.

FORECAST START
The start date of the available forecast results set.

FORECAST END
The end date of the available forecast results set.

EXECUTE TIMESTAMP
The forecast execution timestamp of the available forecast results set.

MAPPING TIMESTAMP
The last updated timestamp of the mapping table on which the forecast is based.

The rules to choose the forecast stream aggregation level when **show forecast** is selected are:

If you select the Average calculation method:

- If you select daily, quarterly, or yearly, and a daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select monthly and a monthly stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select hourly and an hourly stream is available, then it is used. Otherwise, **no data available** is displayed.

If you select the Highest Hourly Average calculation method:

- If you select Highest Hourly Average calculation value and a Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 2nd Highest Hourly Average calculation value and a 2nd Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 3rd Highest Hourly Average calculation value and a 3rd Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 4th Highest Hourly Average calculation value and a 4th Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 5th Highest Hourly Average calculation value and a 5th Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.

CPU: MIPS Used - LPAR Level by WLM Importance

This report shows the observed average or peak CPU MIPS used at the LPAR level and the forecasted average or peak with their confidence intervals. The CPU MIPS used (both observed and forecasted) can be displayed for each processor type at each interval band. The observed CPU MIPS used is also broken down by the WLM importance level. The optimized LPAR weight and the target workload, according to the settings in the selected optimization run, can be also displayed.

The overall usage is shown in MIPS of the LPAR by stacking the usage that is attributed to the workloads that are running in WLM importance levels. The chart shows a reference line that represents the amount of MIPS that the hypervisor attempts to guarantee as available for this LPAR. The configured Weight value for the LPAR is used to determine the guarantee line. If other LPARs are using all of their guaranteed amount of capacity, then the LPAR being viewed must fit its workload demand within the guaranteed amount of capacity. WLM manages the prioritization of work based on the goals and priorities that are established by the customer. WLM defers the lower priority work and attempts to achieve the goals of the higher priority work with the capacity that is available.

By stacking the workload by Importance level order and showing the guarantee line, the report shows you what priority work is at risk by running over the guarantee line. If it shows any importance level work running over the guarantee line, then that work is at risk of not having capacity available if other LPARs use more of their capacity.

By setting up the WLM rules, you can define what work runs at what Importance level. You can use this chart to ensure that the Weight values for each LPAR are satisfying the needs of the important work as you have defined it. If your higher priority workload is consistently running above the guarantee line, then it is time to evaluate the LPAR’s weight value setting.

The peak value for the selected LPAR is shown for each interval band, and the values of each WLM importance level are stacked. The peak time (by hour or SMF interval) is shown if the user chooses to show the peak usage.

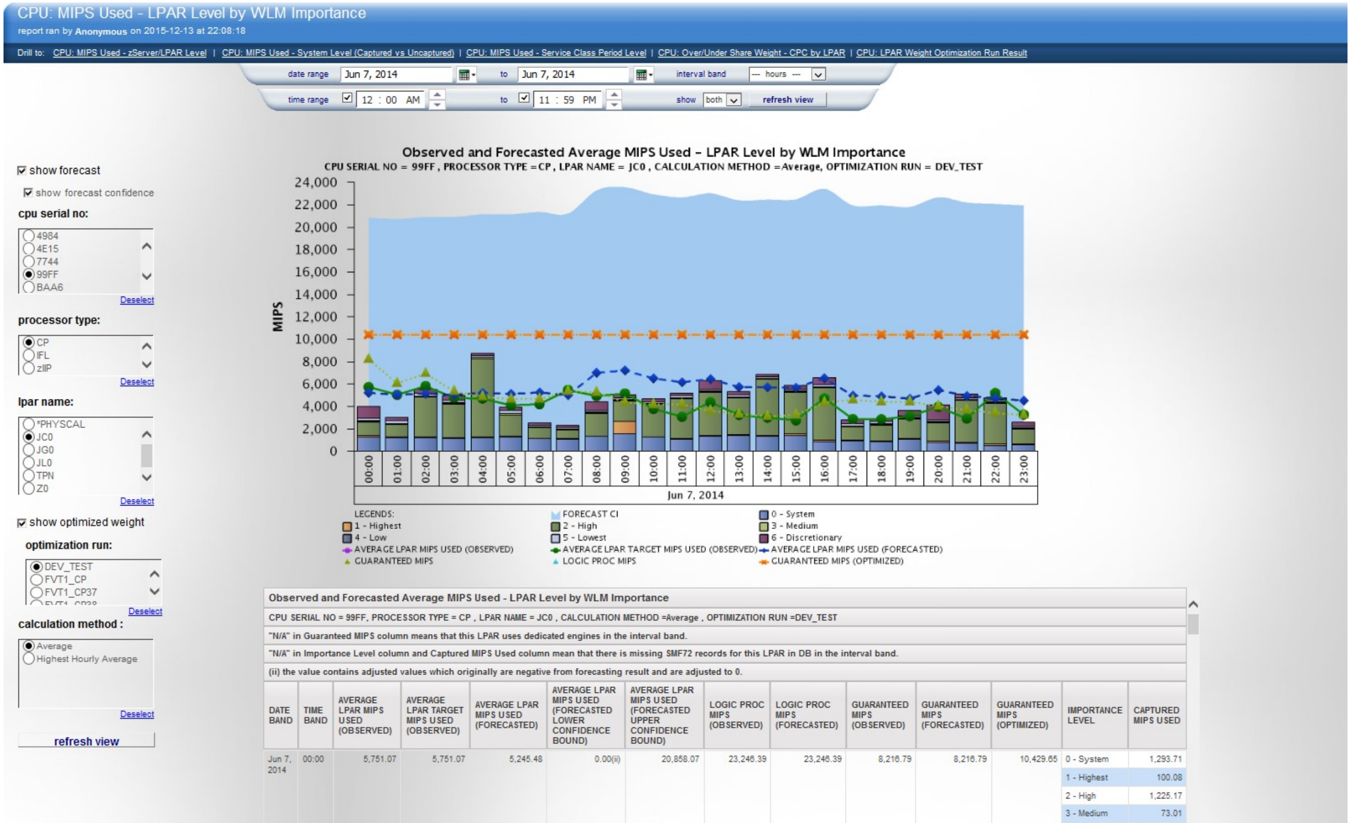

Figure 38: CPU: MIPS Used - LPAR Level by WLM Importance

Source table

[Presentation View].[LPAR Level Metrics]

[Presentation View].[LPAR Level Forecast Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

[Presentation View].[LPAR Weight Optimization Metrics]

Drill through to the following reports

"CPU: MIPS Used - zServer/LPAR Level" on page 224

"CPU: MIPS Used - System Level (Captured vs Uncaptured)" on page 222

"CPU: MIPS Used - Service Class Period Level" on page 220

"CPU: Over/Under Share Weight - CPC by LPAR" on page 234

"CPU: LPAR Weight Optimization Run Result" on page 218

Standard prompts:

date range

time range

interval band (If an interval band of less than one day is selected, and peak usage is shown, then the calculation uses an interval band of one day for the calculation.)

show

Report-specific prompts

show forecast

cpu serial no

processor type

lpar name

show optimized weight

optimization run

calculation method

Three calculation methods are available: Average, Highest Hourly Average, and Percentile of SMF Interval. Percentile of SMF Interval is not available when **show forecast** is selected.

calculation value

This prompt is shown only if you select Highest Hourly Average or Percentile of SMF Interval as the calculation method.

If you select Highest Hourly Average, the values are:

- Highest Hourly Average
- 2nd Highest Hourly Average
- 3rd Highest Hourly Average
- 4th Highest Hourly Average
- 5th Highest Hourly Average

If you select Percentile of SMF Interval, the values are:

- 100%ile of SMF Interval
- 95%ile of SMF Interval
- 90%ile of SMF Interval
- 85%ile of SMF Interval
- 80%ile of SMF Interval
- 75%ile of SMF Interval
- 70%ile of SMF Interval

The Observed Average MIPS Used - LPAR Level by WLM Importance and Observed Peak MIPS Used - LPAR Level by WLM Importance charts or tables contains the following information:

Note: If the **show forecast** prompt is cleared, only observed values are displayed.

cpu serial no
The zServer/CEC serial number.

processor type
The engine type: CP, zIIP, zAAP, IFL, or ICF. The list contains only installed processor types.

lpar name
The name of the logical partition to which the specified processor type is assigned.

DATE BAND
The start date of the time interval.

PEAK DATE
The date when the peak period occurs.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of SMF Interval.

PEAK TIME
The start time of the peak period.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of SMF Interval.

TIME BAND
The start time of the time interval.

This value displays only when the calculation method is Average.

AVERAGE LPAR MIPS USED
The average MIPS usage for the selected LPAR.

This value displays only when the calculation method is Average.

If the SMF72 records for this LPAR in the interval band are missing from the database, a "(i)" marker is shown beside the AVERAGE LPAR MIPS USED value and a footnote is displayed on the top of the table stating: "(i) there is missing SMF72 records for this LPAR in DB in the interval band, and the MIPS value is normalized to the whole interval."

PEAK LPAR MIPS USED

The peak MIPS usage for the selected LPAR.

This value is displayed only when the calculation method is Highest Hourly Average or Percentile of SMF Interval.

If the SMF72 records for this LPAR in the interval band are missing from the database, a "(i)" marker is shown beside the PEAK LPAR MIPS USED value and a footnote is displayed on the top of the table stating: "(i) there is missing SMF72 records for this LPAR in DB in the interval band, and the MIPS value is normalized to the whole interval."

LOGIC PROC MIPS

The logical processors MIPS value is the MIPS capacity of the logical processors that are assigned to this LPAR

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

GUARANTEED MIPS (OBSERVED)

The guaranteed capacity in MIPS on this LPAR. The value is calculated according to the weight assigned to this LPAR.

If dedicated processors are assigned to this LPAR, then this line does not appear in the chart and the value is shown as N/A in the table.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

GUARANTEED MIPS (OPTIMIZED)

The guaranteed capacity in MIPS on this LPAR. The value is calculated according to the optimized weight assigned to this LPAR.

This value is displayed only when **show optimized weight** is selected.

AVERAGE LPAR TARGET MIPS USED

The average LPAR target MIPS usage at each time interval. The target setting is in the setting of the selected optimization run.

This value is displayed only when **show optimized weight** is selected and the calculation method is Average.

If the SMF72 records for this LPAR in the interval band are missing from the database, a "(i)" marker is shown beside the AVERAGE LPAR MIPS USED value and a footnote is displayed on the top of the table stating: "(i) there is missing SMF72 records for this LPAR in DB in the interval band, and the MIPS value is normalized to the whole interval."

PEAK LPAR TARGET MIPS USED

The peak LPAR target MIPS usage at each time interval. The target setting is in the setting of the selected optimization run.

This value is displayed only when **show optimized weight** is selected and the calculation method is Highest Hourly Average or Percentile of SMF Interval.

If the SMF72 records for this LPAR in the interval band are missing from the database, a "(i)" marker is shown beside the PEAK LPAR MIPS USED value and a footnote is displayed on the top of the table stating: "(i) there is missing SMF72 records for this LPAR in DB in the interval band, and the MIPS value is normalized to the whole interval."

IMPORTANCE LEVEL

The WLM importance level assigned to a service class period.

CAPTURED MIPS USED
The CPU time that is consumed by serviceclass_period, converted to MIPS used.

The Observed and Forecasted Average MIPS Used – LPAR Level by WLM Importance and the Observed and Forecasted Peak MIPS Used – LPAR Level by WLM Importance charts and tables contain the information that is shown in the Observed Average MIPS Used - LPAR Level by WLM Importance and Observed Peak MIPS Used - LPAR Level by WLM Importance charts and table, as well as the following information:

PEAK DATE (FORECASTED)
The date when the forecasted peak of the time interval occurs.

This value is displayed only when the calculation method is Highest Hourly Average.

PEAK TIME (FORECASTED)
The time when the forecasted peak of the time interval occurs.

This value is displayed only when the calculation method is Highest Hourly Average.

AVERAGE LPAR MIPS USED (FORECASTED)
The forecasted average MIPS usage for the selected LPAR.

This value is displayed only when the calculation method is Average.

AVERAGE LPAR MIPS USED (FORECASTED LOWER CONFIDENCE BOUND)
The lower confidence bound of the forecasted MIPS.

AVERAGE LPAR MIPS USED (FORECASTED UPPER CONFIDENCE BOUND)
The upper confidence bound of the forecasted MIPS.

Values that are marked with an asterisk (ii) are adjusted values. These values were originally negative in the forecasting result and are adjusted to 0.

Values that are marked with an asterisk (iii) are adjusted values. These values were originally extremes in the forecasted result and are adjusted to the maximum forecasted value.

LOGIC PROC MIPS (FORECASTED)
The forecasted MIPS capacity of the logical processors that are assigned to this LPAR.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

GUARANTEED MIPS (FORECASTED)
The forecasted guaranteed capacity in MIPS on this LPAR.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

The Forecast Metadata table contains the following information. This table is displayed when the **show forecast** prompt is selected and Highest Hourly Average is selected as calculation method.

STREAM TYPE
The stream type that is used by the available forecast results set, for example lpar_mips_peak_daily. This means that the stream is for the lpar mips forecast, and it forecasts the CPU seconds at the daily peak aggregation level.

AGGREGATION
The aggregation level of the available forecast results set. The values are: Highest Hourly Average of LPAR MIPS Daily, 2nd Highest Hourly Average of LPAR MIPS Daily, 3rd Highest Hourly Average of LPAR MIPS Daily, 4th Highest Hourly Average of LPAR MIPS Daily, and 5th Highest Hourly Average of LPAR MIPS Daily.

HOUR START
The start hour of the forecasted hour range of the available forecast results set.

HOUR END
The end hour of the forecasted hour range of the available forecast results set.

HISTORY START
The start date of the based historical data range of the available forecast results set.

HISTORY END
The end date of the based historical data range of the available forecast results set.

FORECAST START
The start date of the available forecast results set.

FORECAST END
The end date of the available forecast results set.

EXECUTE TIMESTAMP
The forecast execution timestamp of the available forecast results set.

MAPPING TIMESTAMP
The last updated timestamp of the mapping table on which the forecast is based.

The rules to choose the forecast stream aggregation level when **show forecast** is selected are:

If you select the Average calculation method:

- If you select daily, quarterly, or yearly, and a daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select monthly and a monthly stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select hourly and an hourly stream is available, then it is used. Otherwise, **no data available** is displayed.

If you select the Highest Hourly Average calculation method:

- If you select Highest Hourly Average calculation value and a Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 2nd Highest Hourly Average calculation value and a 2nd Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 3rd Highest Hourly Average calculation value and a 3rd Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 4th Highest Hourly Average calculation value and a 4th Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.
- If you select 5th Highest Hourly Average calculation value and a 5th Highest Hourly Average of LPAR MIPS Daily stream is available, then it is used. Otherwise, **no data available** is displayed.

CPU: Over/Under Share Weight - CPC by LPAR

This report shows the stacked usage of each LPAR against the total number of engines of each type in the shared pool of a CPC. It also shows the usage percentage over or under its guaranteed share of each LPAR, according to both the observed and optimized weights, due to the competition among LPARs. The target workload, according to the selected optimization run, can also be displayed.

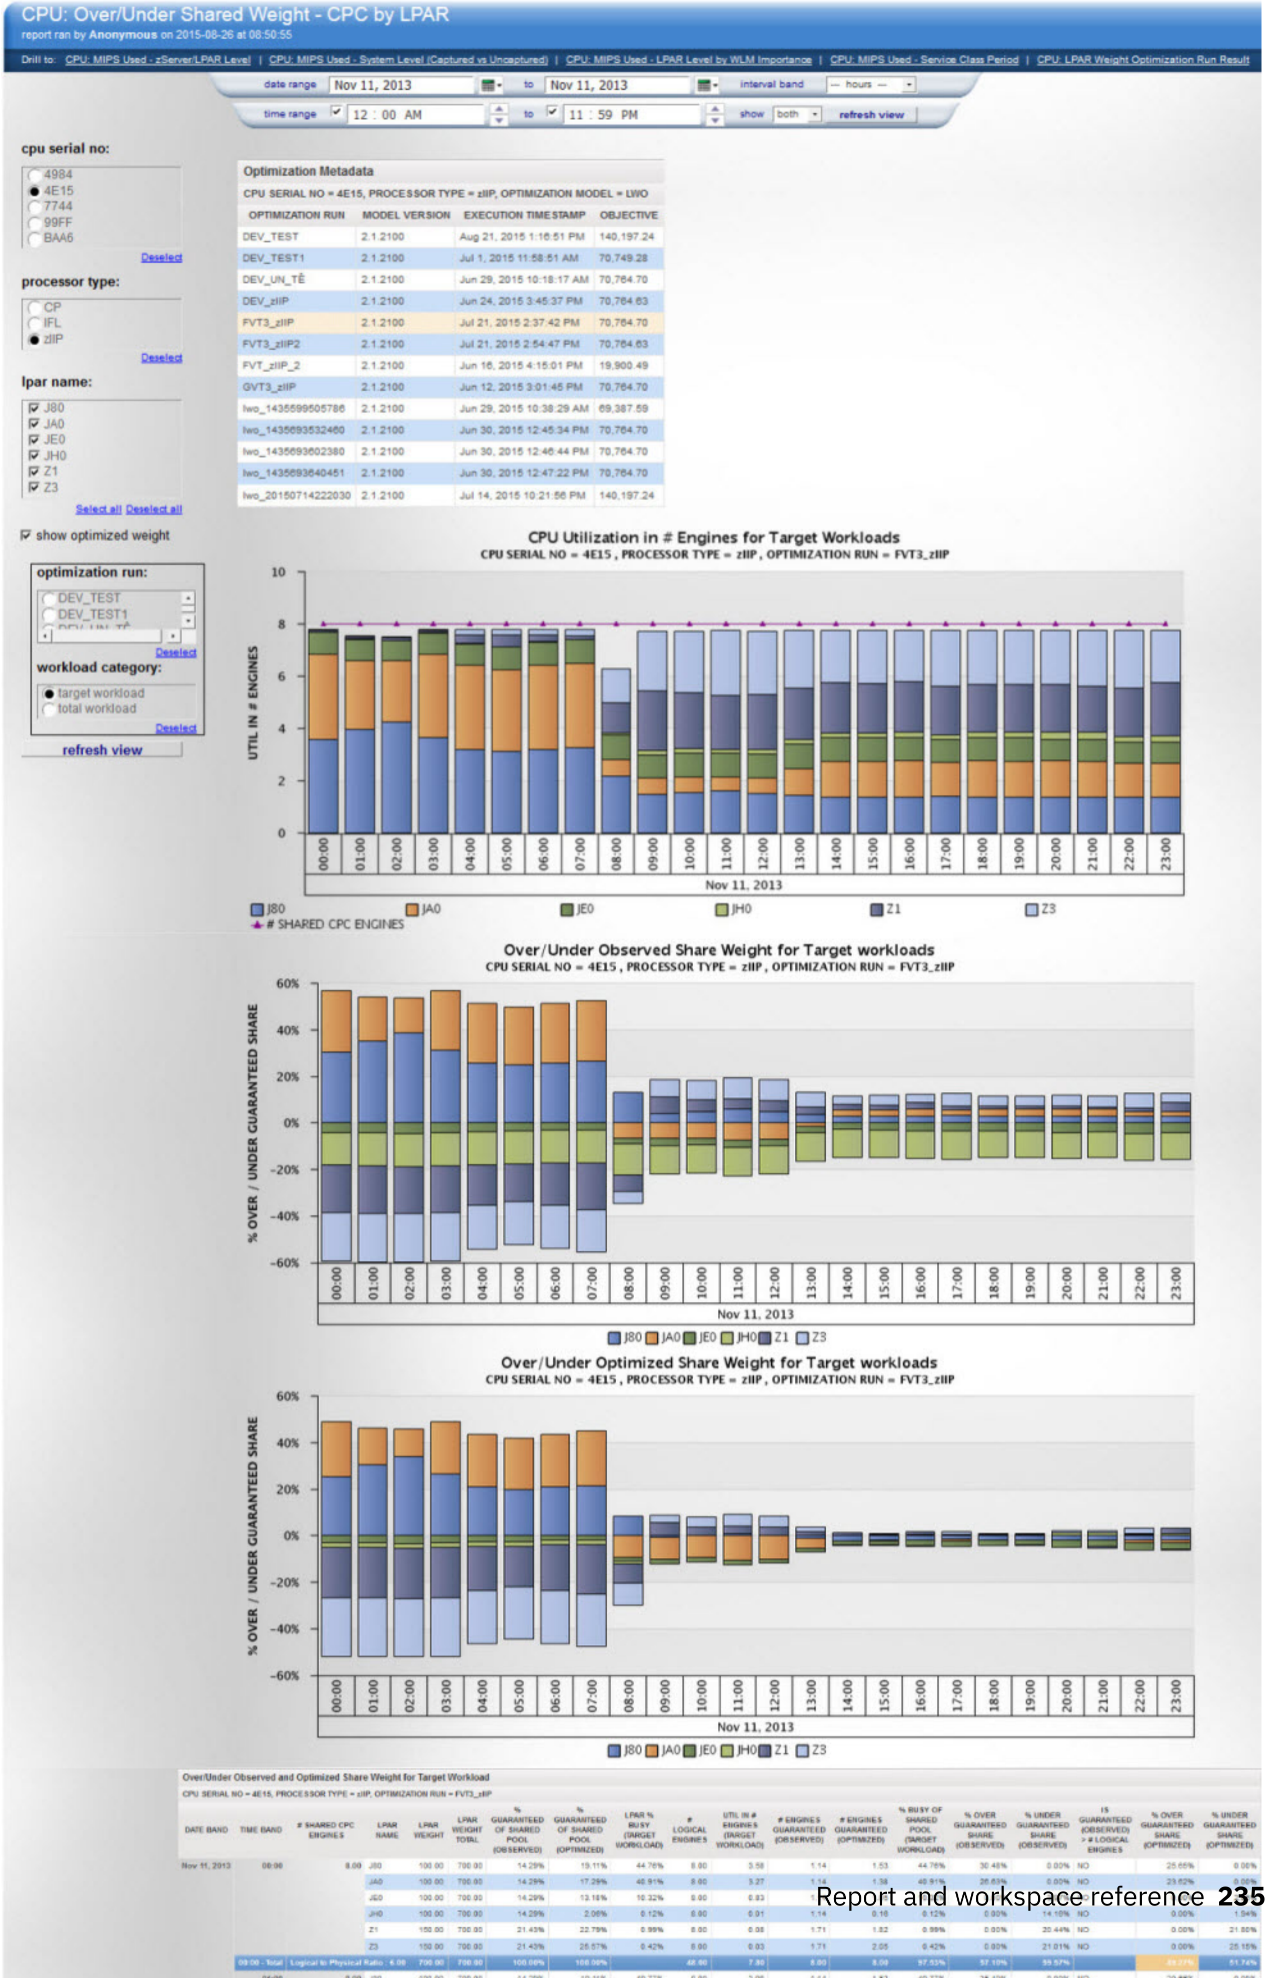

CPU: Over/Under Shared Weight - CPC by LPAR
report ran by Anonymous on 2015-08-26 at 08:50:55
Drill to: CPU: MIPS Used - zServer/LPAR Level | CPU: MIPS Used - System Level (Captured vs Uncaptured) | CPU: MIPS Used - LPAR Level by WLM Importance | CPU: MIPS Used - Service Class Period | CPU: LPAR Weight Optimization Run Result
date range Nov 11, 2013 to Nov 11, 2013 interval band -- hours --
time range 12 : 00 AM to 11 : 59 PM show both refresh view
cpu serial no:
4984
4E15
7744
99FF
BAA6
Deselect
processor type:
CP
IFL
zIIP
Deselect
lpar name:
J80
JA0
JE0
JH0
Z1
Z3
Select all Deselect all
show optimized weight
optimization run:
DEV_TEST
DEV_TEST1
Deselect
workload category:
target workload
total workload
Deselect
refresh view
Optimization Metadata
CPU SERIAL NO = 4E15, PROCESSOR TYPE = zIIP, OPTIMIZATION MODEL = LWO
OPTIMIZATION RUN
MODEL VERSION
EXECUTION TIMESTAMP
OBJECTIVE
CPU Utilization in # Engines for Target Workloads
CPU SERIAL NO = 4E15 , PROCESSOR TYPE = zIIP , OPTIMIZATION RUN = FVT3_zIIP
UTIL IN # ENGINES
Nov 11, 2013
J80
JA0
JE0
JH0
Z1
Z3
SHARED CPC ENGINES
Over/Under Observed Share Weight for Target workloads
CPU SERIAL NO = 4E15 , PROCESSOR TYPE = zIIP , OPTIMIZATION RUN = FVT3_zIIP
% OVER / UNDER GUARANTEED SHARE
Nov 11, 2013
Over/Under Optimized Share Weight for Target workloads
CPU SERIAL NO = 4E15 , PROCESSOR TYPE = zIIP , OPTIMIZATION RUN = FVT3_zIIP
% OVER / UNDER GUARANTEED SHARE
Nov 11, 2013
Over/Under Observed and Optimized Share Weight for Target Workload

On a CPC, there are multiple LPARs that have a configuration of CPUs that they share among other LPARs that are defined on the same CPC. And each LPAR has a defined maximum number of CPUs that it can consume. It is also given a weight value that is calculated as a percentage of the entire set of CPUs that are shared on the CPC. The percentage of the physical CPUs is the target, or guaranteed, amount of CPU that the hypervisor ensures that this LPAR is able to use. If there is unused capacity on the CPC above this target and the LPAR wants more, then it is allowed to use it up to the maximum number of CPUs that are defined for that LPAR.

There is a shared pool for each type of processor; General CP, zIIP, and zAAP. Each shared pool is managed separately and each LPAR has a target or guaranteed amount for each processor type that is based on weight. An LPAR participates in the sharing of a processor pool only if it has at least one logical processor of that processor type defined.

Important: When there are several LPARs on a CPC that are competing for more than their guaranteed amount during different portions of the day, it can be difficult to assess just what those guaranteed amounts should be to satisfy the workloads in the proper priority. When the CPC is running at 100% busy, then the hypervisor is trying to satisfy each LPAR's demand up to the guaranteed amount.

This report uses two charts to show the guaranteed weight. One chart shows the weight from the CPC perspective, and the other shows it from the perspective of each individual LPAR against the shared pool for the CPC.

- From the CPC perspective, the chart shows the total number of CPUs in the shared pool and stacks the usage of each LPAR that is sharing the pool. This provides a view of how much of the shared pool is being utilized. This might not be the entire CPC capacity, as there can be other LPARs with dedicated CPs.
- The competition among LPARs is shown by the usage of each LPAR against it’s own guaranteed share as a percent of the CPC shared processor pool.

  For example, if an LPAR is guaranteed 30% of the CPC shared pool and it uses only 25%, then it shows as using 5% under its target. If there are 3 LPARs that each have a guarantee of 30% and each use 25%, then the chart shows a stacked bar of 3 LPARs each at 5% under, showing a total of 15% under target. The 15% of the CPC that is not being consumed is available for other LPARs.

  In the example, there is 15% capacity available above the planned target capacity to support the workload according to business priorities.

  However, at the same time interval there might be another LPAR with a guarantee of 10% and it might be using 25% of the CPC shared pool. This would be 15% over its guarantee and would be shown as 15% over the target usage. In this case, the CPC is actually running at 100% busy of the shared processor pool.

  The chart shows which LPARs are using more than their share. They are allowed to because the others are not using all of their guarantee.

This report, with both of its charts, can be used to quickly assess if the time period of the CPC being 100% is good or bad. If this usage matches the business priority as expected, then it is good.

The share weight reports should be viewed by using a fine interval band, preferably no higher than hourly. If all of the LPARs are using the same SMF interval duration, then you can go to that finest level. If a different SMF interval duration is used by different LPARs, then the selected interval duration is displayed.

Source table

[Presentation View].[LPAR Level Metrics]

[Presentation View].[LPAR Weight Optimization Metrics]

[Presentation View].[Optimization Metadata]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports

“CPU: MIPS Used - Service Class Period Level” on page 220

“CPU: MIPS Used - zServer/LPAR Level” on page 224

“CPU: MIPS Used - System Level (Captured vs Uncaptured)” on page 222

“CPU: MIPS Used - LPAR Level by WLM Importance” on page 229

“CPU: LPAR Weight Optimization Run Result” on page 218

Standard prompts:

date range

time range

interval band

show

Report-specific prompts

cpu serial no

processor type

lpar name

show optimized weight

optimization run

workload category

This report contains the following information:

cpu serial no

The zServer/CEC serial number.

processor type

The engine type: CP, zIIP, zAAP, IFL, or ICF. The list contains only installed processor types.

lpar name

The name of the logical partition to which the specified processor type is assigned.

Only LPARs with shared processors are shown.

 # SHARED CPC ENGINES

The number of shared engines on the CPC.

LPAR WEIGHT

The weight that is assigned to this LPAR.

LPAR WEIGHT TOTAL

The sum of the weights for all active LPARs on the CPC for the selected processor type.

% GUARANTEED OF SHARED POOL (OBSERVED)

The percent of the shared pool that is guaranteed to the LPAR. The value is calculated according to the weight assigned to this LPAR.

% GUARANTEED OF SHARED POOL (OPTIMIZED)

The percent of the shared pool that is guaranteed to the LPAR. The value is calculated according to the optimized weight assigned to this LPAR.

This value is displayed only when **show optimized weight** is selected.

LPAR % BUSY

The percentage of time during which the logical engines for this LPAR are busy.

This information is displayed when **show optimized weight** is cleared, or when **show optimized weight** is selected and **total workload** in **workload category** is also selected.

LPAR % BUSY (TARGET WORKLOAD)

The percentage of time during which the logical engines for this LPAR are busy for the target workload.

This information is displayed only when **show optimized weight** and **target workload** in **workload category** are selected.

 # LOGICAL ENGINES

The number of logical engines for this LPAR.

UTIL IN # ENGINES

The utilization in number of the logical engines on this LPAR.

This information is displayed when **show optimized weight** is cleared, or when **show optimized weight** is selected and **total workload** in **workload category** is also selected.

UTIL IN # ENGINES (TARGET WORKLOAD)
The utilization of target workload in number of the logical engines on this LPAR.

This information is displayed only when **show optimized weight** and **target workload** in **workload category** are selected.

 # ENGINES GUARANTEED (OBSERVED)
The guaranteed capacity in number of engines on this LPAR. This value is calculated according to the weight assigned to the LPAR.

 # ENGINES GUARANTEED (OPTIMIZED)
The guaranteed capacity in number of engines on this LPAR. This value is calculated according to the optimized weight assigned to the LPAR.

This value is displayed only when **show optimized weight** is selected.

% BUSY OF SHARED POOL
The utilization of this LPAR as a percentage of the shared pool of processors on the CPC.

This information is displayed when **show optimized weight** is cleared, or when **show optimized weight** is selected and **total workload** in **workload category** is also selected.

% BUSY OF SHARED POOL (TARGET WORKLOAD)
The utilization of target workload of this LPAR as a percentage of the shared pool of processors on the CPC.

This information is displayed only when **show optimized weight** and **target workload** in **workload category** are selected.

% OVER GUARANTEED SHARE (OBSERVED)
The usage percentage over the guaranteed share for each LPAR due to the competition among LPARs.

% UNDER GUARANTEED SHARE (OBSERVED)
The usage percentage under the guaranteed share for each LPAR due to the competition among LPARs.

IS GUARANTEED (OBSERVED) > # LOGICAL ENGINES
If the guaranteed number of engines is larger than the number of logical engines, then the LPAR cannot get its guaranteed number of engines forever. In this case, the value is set to YES as a reminder of this configuration conflict on the LPAR.

% OVER GUARANTEED SHARE (OPTIMIZED)
The usage percentage for the total or target workload over the optimized guaranteed share for each LPAR due to the competition among LPARs.

This information is displayed only when **show optimized weight** is selected.

% UNDER GUARANTEED SHARE (OPTIMIZED)
The usage percentage for the total or target workload under the optimized guaranteed share for each LPAR due to the competition among LPARs.

This information is displayed only when **show optimized weight** is selected.

LOGICAL TO PHYSICAL RATIO
The ratio that is presented is for all LPARs selected against the total physical shared pool.

The Optimization Metadata table contains the following information. The table is shown when **show optimized weight** is selected.

OPTIMIZATION RUN
The name that is used for the execution of the LPAR weight optimization.

MODEL VERSION
The version of the optimization model.

EXECUTION TIMESTAMP
The execution start time of this optimization run.

OBJECTIVE

The total objective value of this optimization model.

CPU: Windows - CPU Usage

This report shows the percentage of CPU used from both a virtual and physical perspective.

The server might be virtualized or it might be a dedicated physical server. Showing both virtual and physical servers lets you know what type of server it is and how it is using its resources.

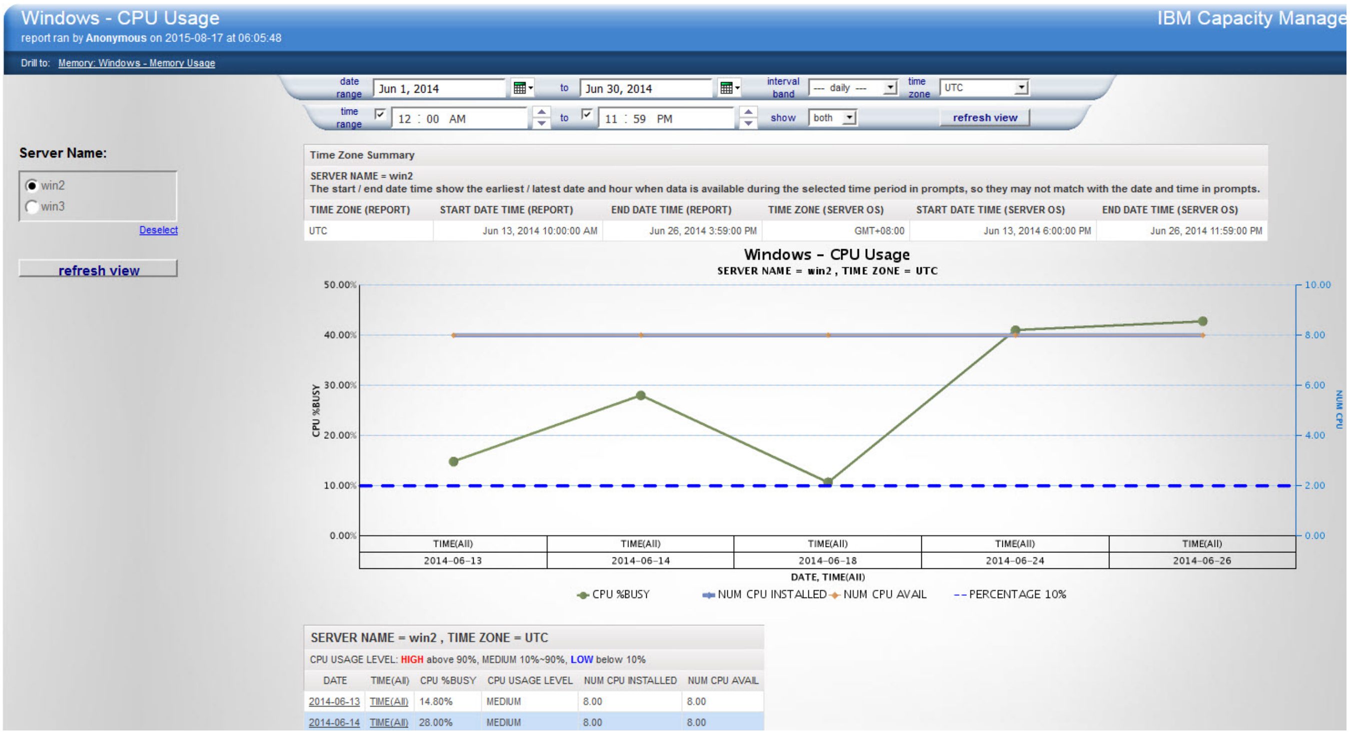

Figure 40: Windows - CPU Usage example

Source tables

[Presentation View].[WIN DIST DATE]

[Presentation View].[WIN DIST HOUR]

[Presentation View].[WIN DIST TIME]

[Presentation View].[WIN CPU Fact]

[Presentation View].[WIN IMAGE]

Drill through to the following report

"Memory: Windows - Memory Usage" on page 248

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Windows - CPU Usage report contains the following information:

server name
The name of the server against which the report is run.

CPU % BUSY
CPU utilization in percent. For systems with multiple cores this value represents the average utilization over all cores

NUM CPU AVAIL
The number of processors currently available on this system.

NUM CPU INSTALLED
The number of processors installed on this system.

PERCENTAGE 10%
10% CPU busy line.

PERCENTAGE 90%
90% CPU busy line. If the maximum CPU Usage percentage is less than 90%, then only 10% CPU busy line is displayed.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the LPAR time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Memory: AIX - Memory Usage

This report shows the free physical memory, free virtual memory, and active virtual memory on the system. The report also displays metric showing the turnover of physical memory, such as Page-In/sec and Page-Out/sec (file system related pages), as well as PageSpace-In/sec and PageSpace-Out/sec.

Memory usage cannot be fully understood by using just one metric. For example, if you see a spike in Page-in/sec, maybe there is also page out activity. But if the active virtual memory has a minimal increase and there is a large amount of free memory on the system, then there is not a memory shortage on the server. However, there is some memory turnover occurring. This might be the nature of the application or it might suggest that there is an opportunity to tune the application on the server.

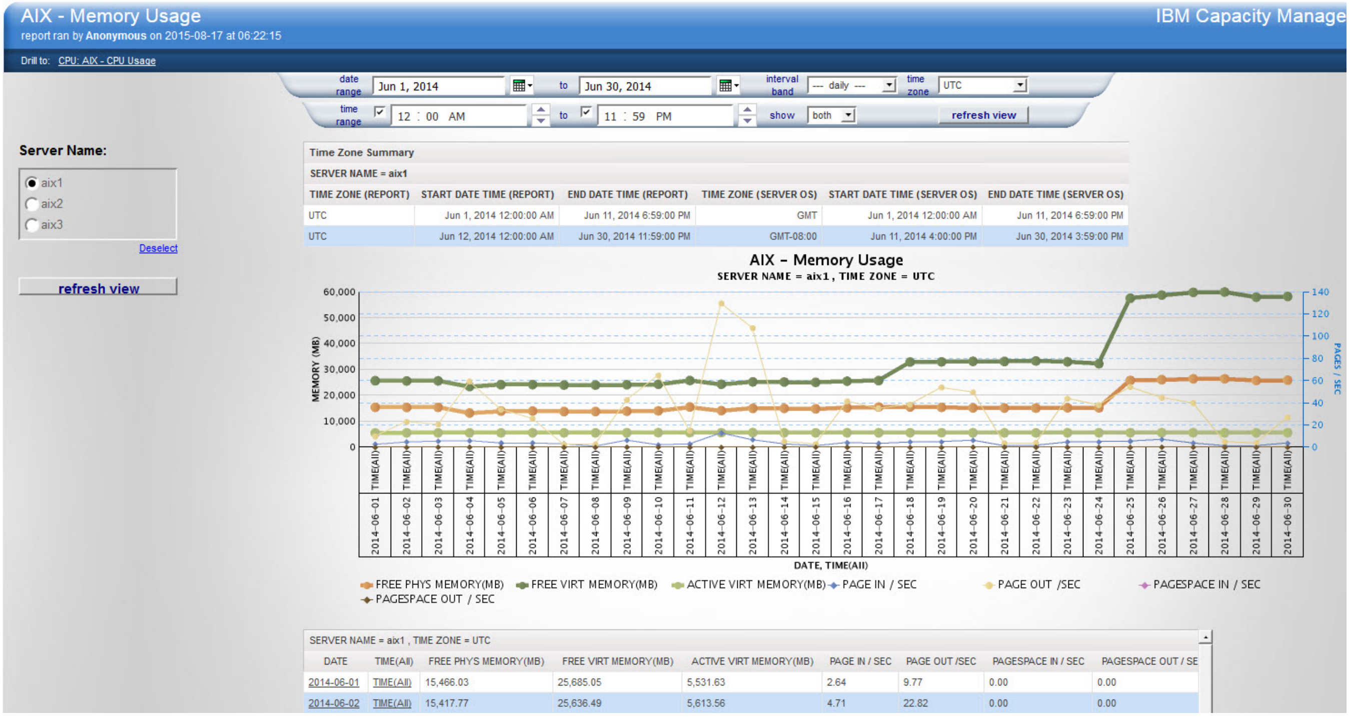

Figure 41: AIX - Memory Usage example

Source tables

[Presentation View].[AIX DIST Date]

[Presentation View].[AIX DIST Hour]

[Presentation View].[AIX DIST TIME]

[Presentation View].[AIX Memory Fact]

[Presentation View].[AIX IMAGE]

Drill through to the following report

"CPU: AIX - CPU Usage" on page 210

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The AIX - Memory Usage report contains the following information:

server name

The name of the server against which the report is run.

FREE PHYS MEMORY (MB)

The amount of physical memory currently unused and available. In a virtualized environment this is the free memory available to the operating system

FREE VIRT MEMORY (MB)

The amount of virtual memory currently unused and available.

ACTIVE VIRT MEMORY (MB)

The current number of active pages in virtual memory. Note that this metric is expressed in pages. Pages are 4K in size, so the number of pages are multiplied by 4, then divided by 1024 to get MB.

Page In / Sec
The number of page-in operations per second from the file system.

Page Out / Sec
The number of page-out operations per second from the file system.

Pagespace In / Sec
Number of page-in operations per second from the paging space.

Pagespace Out / Sec
Number of page-out operations per second from the paging space.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Memory: CSA/ECSA/SQA/ESQA Utilization

This report provides information on detailed memory utilization to help you decide whether you need to adjust and rebalance memory allocations to prevent costly system outages.

The report shows the peak (maximum) utilization for common virtual storage areas.

- Common service area (CSA)
- Extended common service area (ECSA)
- System queue area (SQA)
- Extended system queue area (ESQA)

System outages can occur when CSA or ECSA storage is exhausted, so it is critical to monitor the CSA, ECSA, SQA, and ESQA utilization.

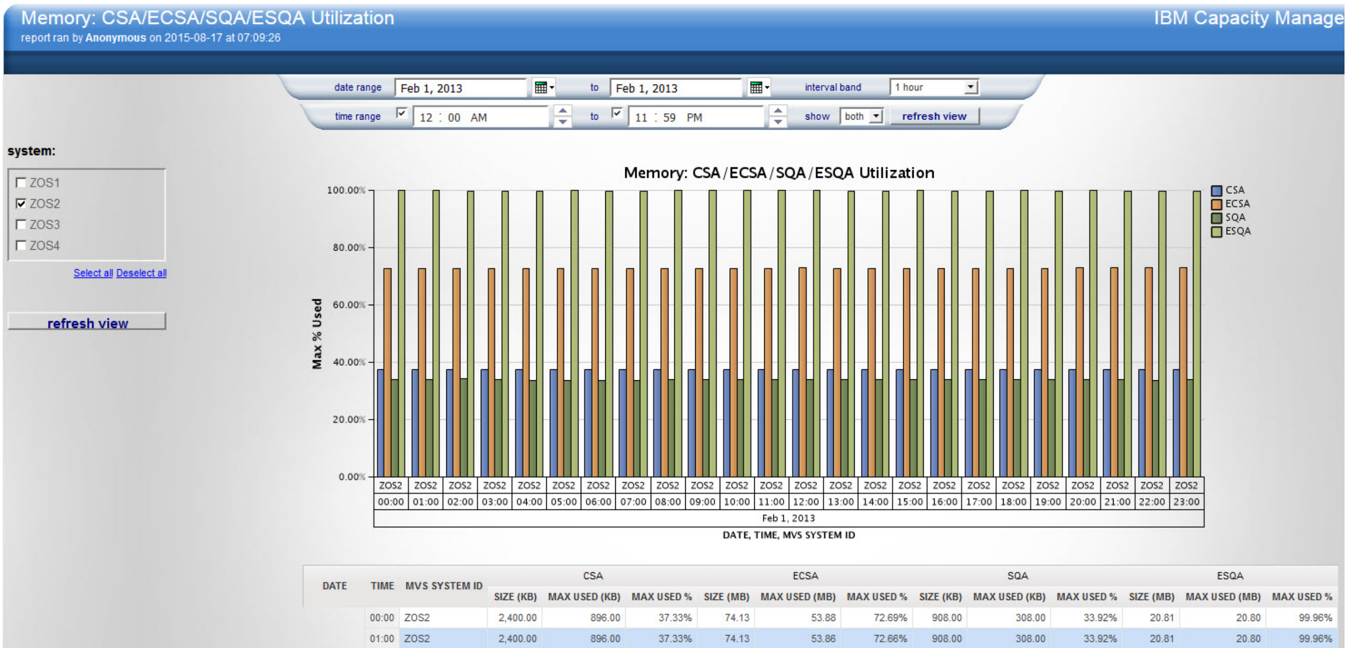

Figure 42: Memory: CSA/ECSA/SQA/ESQA Utilization sample report

Source tables
[Presentation View].[CSA/ECSA/SQA/ESQA Usage Metrics]

Drill through to the following reports
No drill through available.

Standard prompts:
date range

time range

interval band

show

Report-specific prompts
system

The Memory: CSA/ECSA/SQA/ESQA Utilization report contains the following information:

system
The z/OS system ID.

CSA
The maximum percentage of CSA used below the 16MB line

ECSA
The maximum percentage of CSA used above the 16MB line.

SQA
The maximum percentage of SQA used below the 16MB line.

ESQA
The maximum percentage of SQA used above the 16MB line.

Memory: Linux for System x - Memory Usage

This report shows the installed (visible) memory, free physical memory, and free virtual memory on the server.

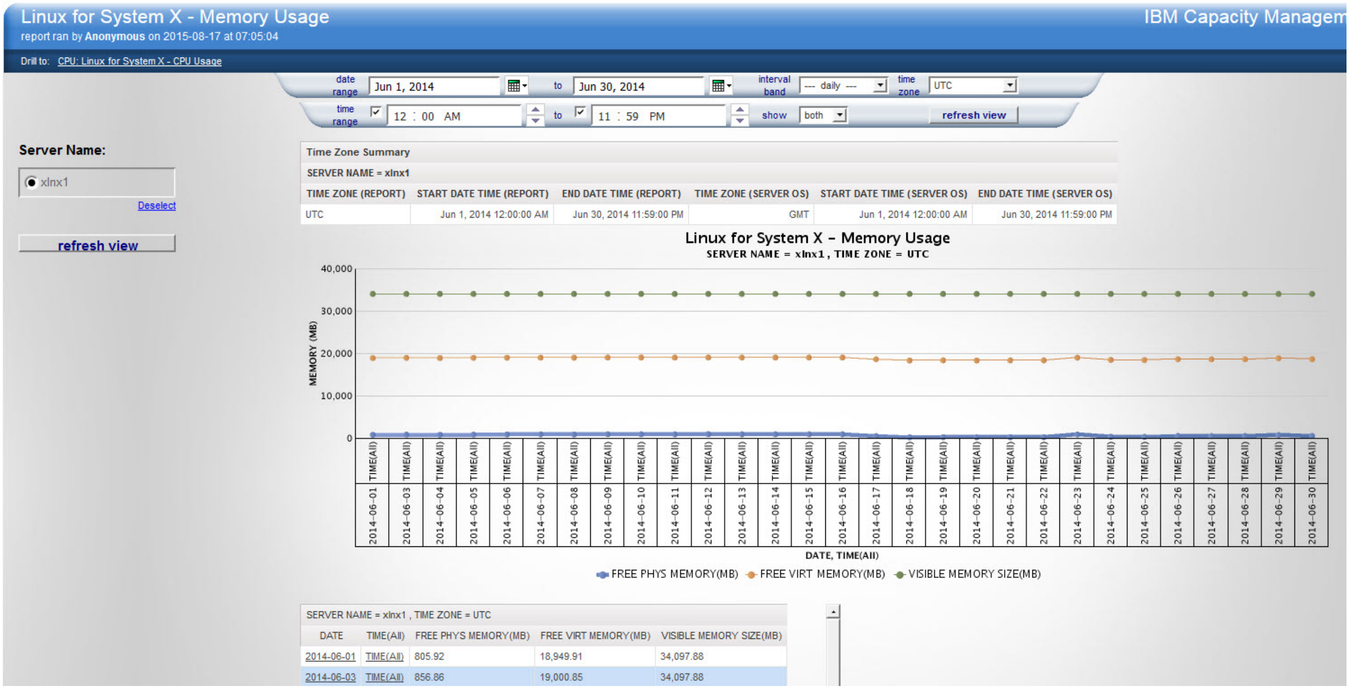

Figure 43: Linux on System x - Memory Usage example

Source table

[Presentation View].[Linux for System X DIST Date]

[Presentation View].[Linux for System X DIST Hour]

[Presentation View].[Linux for System X DIST Time]

[Presentation View].[Linux for System X Memory Fact]

[Presentation View].[Linux for System X Image]

Drill through to the following report

"CPU: Linux for System x - CPU Usage" on page 213

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Linux for System x - Memory Usage report contains the following information:

server name

The name of the server against which the report is run.

FREE PHYS MEMORY(MB)

The amount of physical memory currently unused and available. In a virtualized environment, this is the free memory available to the operating system.

FREE VIRT MEMORY(MB)
The amount of virtual memory currently unused and available.

VISIBLE MEMORY SIZE(MB)
The total amount of physical memory available to the operating system. This value does not necessarily indicate the true amount of physical memory, but indicates the amount that is reported as being available to the operating system.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Memory: Linux for System z - Memory Usage

This report shows the installed (visible) memory, free physical memory, and free virtual memory on the server.

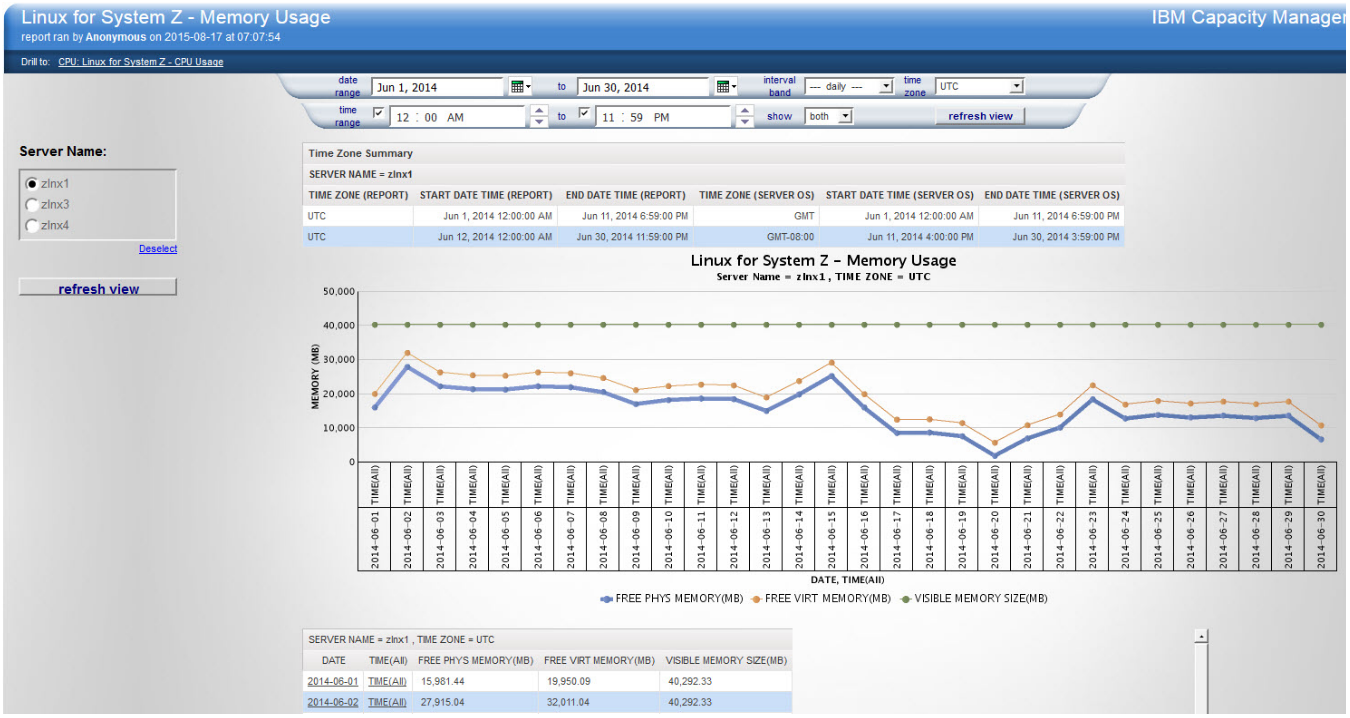

Figure 44: Memory: Linux for System z - Memory Usage sample report

Source table

[Presentation View].[Linux for System Z DIST Date]

[Presentation View].[Linux for System Z DIST Hour]

[Presentation View].[Linux for System Z DIST Time]

[Presentation View].[Linux for System Z Memory FACT]

[Presentation View].[Linux for System Z Image]

Drill through to the following report

"CPU: Linux for System z - CPU Usage " on page 216

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Linux for System z - Memory Usage report contains the following information:

server name

The name of the server against which the report is run.

FREE PHYS MEMORY(MB)

The amount of physical memory currently unused and available. In a virtualized environment, this value is the free memory that is available to the operating system.

FREE VIRT MEMORY(MB)
The amount of virtual memory currently unused and available.

VISIBLE MEMORY SIZE(MB)
The total amount of physical memory available to the operating system. This value does not necessarily indicate the true amount of physical memory, but indicates the amount that is reported as being available to the operating system.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Memory: Windows - Memory Usage

This report shows how much memory is installed and the amount of free memory that is available for the server.

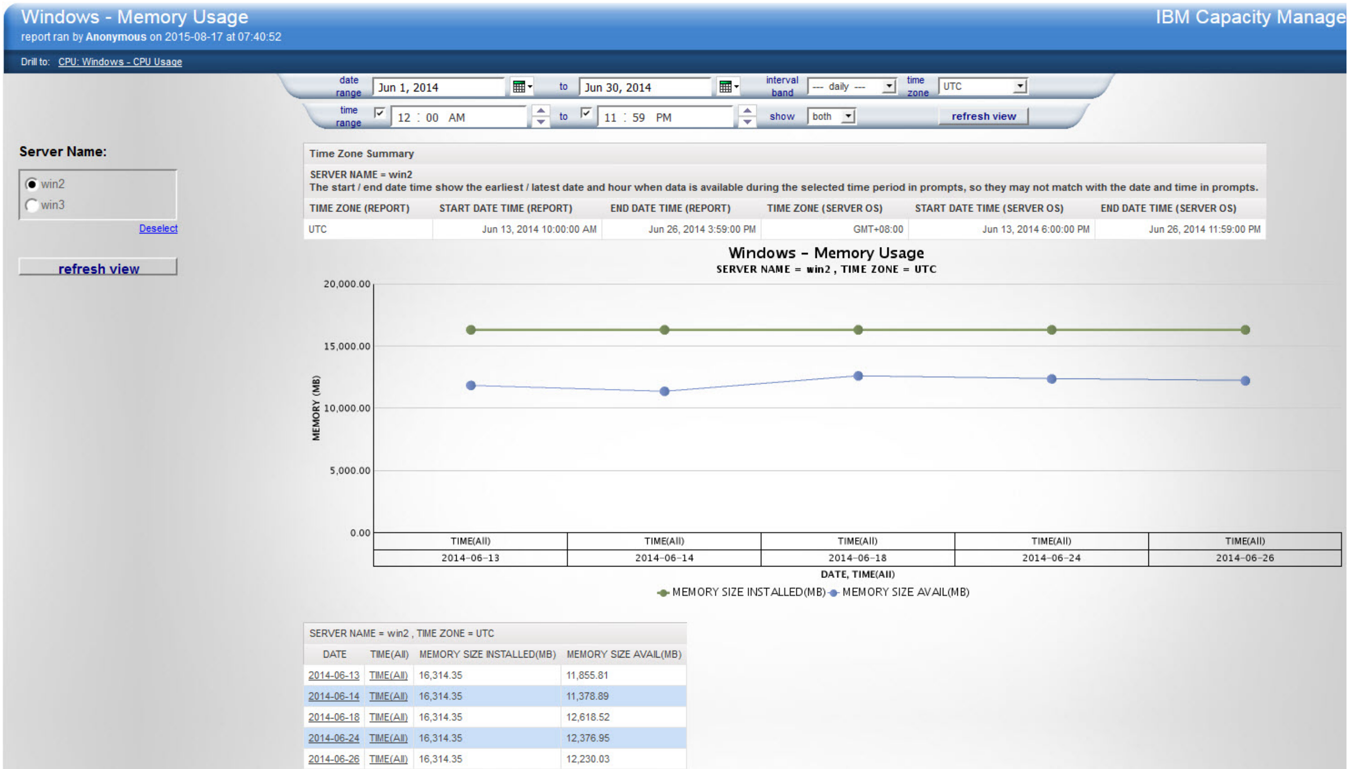

Figure 45: Windows - Memory Usage example

Source table

[Presentation View].[WIN DIST DATE]

[Presentation View].[WIN DIST HOUR]

[Presentation View].[WIN DIST TIME]

[Presentation View].[WIN Memory Fact]

[Presentation View].[WIN IMAGE]

Drill through to the following report

"CPU: Windows - CPU Usage" on page 239

Standard prompts:

date range

time range

interval band

time zone

show

Report-specific prompts

Server Name

The Windows - Memory Usage report contains the following information:

server name

The name of the server against which the report is run.

MEMORY SIZE AVAIL(MB)

Amount of free and available memory on this system.

MEMORY SIZE INSTALLED(MB)
The amount of memory that is installed on this system.

Time Zone Summary table

This table displays the time zone summary information for a selected server and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (SERVER OS)
The local time of the selected server. The format is GMT -5:00

Important: the server time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (SERVER OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (SERVER OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

SCA: LPAR MSU Utilization

This report helps you understand and forecast your costs based on your MSU usage. You can compare your historical MSU usage with the forecasted usage to see upcoming trends and take proactive action to adjust your workload and capacity limits to help control your bill.

The report shows MSU utilization metrics from an LPAR perspective. It includes the following metrics:

- LPAR utilization MSU
- LPAR 4 Hour Rolling Average (4HRA) utilization MSU—The 4 hour rolling average of the MSU that is used by the LPAR each hour compared with the forecast and the recommended optimization.
- LPAR billable MSU
- Product MSU utilization

The report displays data for each Software Cost Analysis scenario:

- Observed (can also display data at the SMF collection interval)
- Forecasted
- Optimized

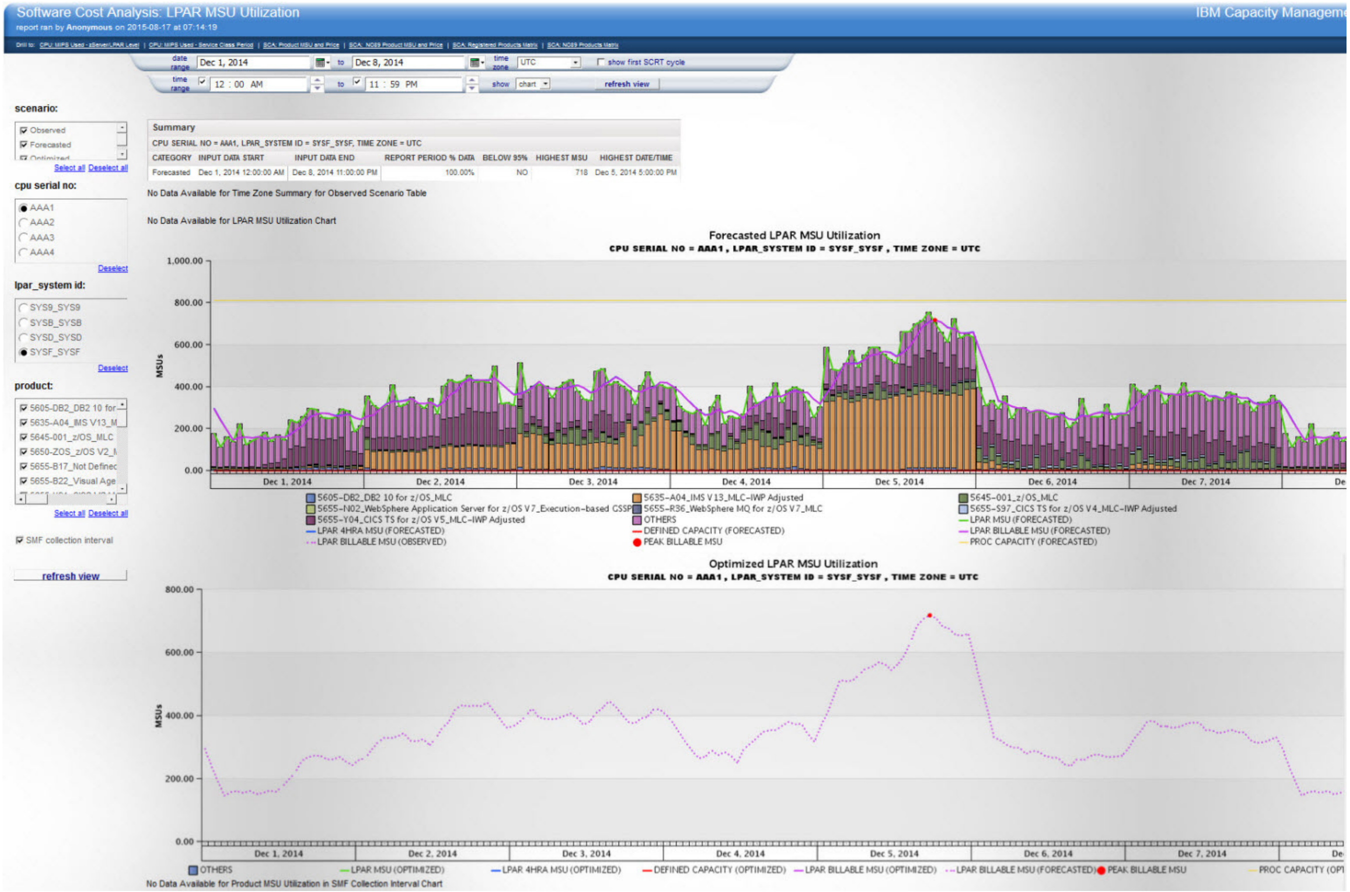

Figure 46: SCA: LPAR MSU Utilization example

Sub-capacity pricing has several pricing structures. One of the more popular structures is WLC which prices the product based on the MSU usage of the entire LPAR. This pricing model presents the quesion of what is driving the LPAR's usage and thus driving the cost of the WLC-based products on this LPAR. The report provides the LPAR 4HRA peak for the selected date range. It also provides a breakdown of all of the registered product's MSU usage during each interval. The report totals the MSUs that are used by all products and subtracts that value from the LPAR MSU usage to arrive at the MSU usage of all other things on the LPAR, which is called the "Other" bucket. The Observed view shows which products are driving the usage of the LPAR. Comparing the overall LPAR MSU usage and individual product MSU usage provides some guidance as to the cost effectiveness of running each product on this LPAR.

The forecast view provides a forecast for each individual product, including the "Other" bucket. You can select all products and see the breakdown of Observed MSU and the aggregate forecast of all products. Based on the Observed MSU usage of a select set of products, you might want to see the forecast for just those products. The report provides the aggregate MSU forecast for the group of products that you select.

If the "Other" bucket represents a large portion of the MSU usage on this LPAR, then you might want to move some of that workload to another LPAR to reduce the WLC costs for this LPAR.

Source tables and views

[Presentation View].[Timezone Metrics]

[Presentation View].[Prompts/Product Prompt for LPAR MSU Report View]

[Presentation View].[Prompts/LPAR MSU Prompt View]

[Presentation View].[Product MSU Forecast Metrics]

[Presentation View].[Product MSU Optimization Metrics]

[Presentation View].[Raw LPAR MSU Utilization Metrics]

[Presentation View].[Raw Product MSU Utilization Metrics]

[Presentation View].[Raw LPAR Product MSU Utilization Metrics]

[Presentation View].[LPAR MSU Metrics]

[Presentation View].[LPAR MSU Optimization Metrics]

[Presentation View].[LPAR MSU Forecast Metrics]

[Presentation View].[Movement Suggestion by Optimization Metrics]

[Presentation View].[LPAR Level Metrics]

Drill through to the following reports

“CPU: MIPS Used - zServer/LPAR Level” on page 224

“CPU: MIPS Used - Service Class Period Level” on page 220

“SCA: NO89 Product MSU and Price ” on page 256

“SCA: Product MSU and Price ” on page 257

Standard prompts

date range

time range

show first SCRT cycle

time zone

show

Report-specific prompts

scenario

cpu serial no

lpar_system_id

product

SMF collection interval

The scenario prompt is independent of other prompts. This means that the values in other prompts, such as cpu serial no and lpar system id, will include all the possible values of the three scenarios. This also means you might sometimes see no data available even when there are cpu serial no, lpar system id, and product values are available in prompts.

All the charts and tables that are shown in this report are in the hourly interval band except when you select the SMF collection interval. When you select the SMF collection interval, another chart (Observed Product MSU Utilization in SMF Collection Interval) reports the product MSU utilization in SMF collection interval, and more columns are added to the Observed LPAR MSU Utilization table to show the observed LPAR and product MSU Utilization in SMF Collection Interval. Only the observed scenario has data in the SMF collection interval; forecasted and optimized scenarios are in the hourly interval band.

The LPAR MSU Utilization report includes the following charts and tables.

Summary table

The Summary table shows the summary information about the percent of data that is collected and highest MSU. This table is always displayed if any scenario is selected. The table includes the following information:

CATEGORY

The scenario name: Observed, Forecasted or Optimized.

INPUT DATA START

The earliest date with data in the database during the selected date range.

INPUT DATA END

The latest date with data in the database during the selected date range.

REPORT PERIOD % DATA
The time period with data in the database for this LPAR vs. the selected period. In the observed scenario, the data is from the DRL.MVSPM_LPAR_MSU_T table, which is from SMF70 records. In the forecasted scenario, the data is from the HCM.LPAR_MSU_FORECAST table, which is the result of forecast streams. In the optimized scenario, the data is from the HCM.LPAR_MSU_OPTIMIZATION table, which is the result of optimization streams.

BELOW 95%
If the value of REPORT PERIOD % DATA is below 95%, then "YES", otherwise "NO". This value and the related REPORT PERIOD % DATA are highlighted by red font color.

HIGHEST MSU
The highest billable MSU on this LPAR during this period.

HIGHEST DATE/TIME
The date and time at which the highest billable MSU first occurred.

Time Zone Summary table

This table displays the time zone summary information for a selected LPAR and contains the following elements:

TIME ZONE (REPORT)
The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)
The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (LPAR OS)
The local time of the selected server. The format is GMT -5:00

Important: the LPAR time zone is the time zone of the server OS and not the time of where the server is physically located.

START DATE TIME (LPAR OS)
The start date and time in the selected server local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (LPAR OS)
The end date and time in the selected server local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Observed/Forecasted/Optimized LPAR MSU Utilization charts and tables

These charts and tables show MSU utilization metrics in different scenarios from an LPAR perspective. Metrics include LPAR utilization MSU, LPAR 4HRA utilization MSU, LPAR billable MSU, and product MSU utilization for each scenario. Clicking any of these charts or tables drills through to the Product MSU and Price report.

The charts and tables include the following information:

PRODUCT UTILIZED MSU
The MSU utilized by the selected product at each hour.

LPAR UTILIZED MSU/LPAR MSU (FORECASTED)/LPAR MSU (OPTIMIZED)

The MSU LPAR utilized at each hour, which can be compared with the forecast and the recommended optimization.

If the values of a scenario are all NULL, then in the chart, the line is not shown; in the table, this column is not shown for this scenario.

LPAR 4HRA MSU/ LPAR 4HRA MSU (FORECASTED) / LPAR 4HRA MSU(OPTIMIZED)

The 4 hour rolling average of the MSU being used by the LPAR each hour, which can be compared with the forecast and recommended optimization.

If the values of a scenario are all NULL, then in the chart, the line is not shown; in the table, this column is not shown for this scenario.

DEFINED CAPACITY/ DEFINED CAPACITY (FORECASTED)/ DEFINED CAPACITY (OPTIMIZED)

The defined capacity of this LPAR at each hour of the related scenario.

If some or all values of a scenario are NULL, then in the chart, the line is not shown; in the table, this value is shown as 0 for this scenario.

LPAR BILLABLE MSU / LPAR BILLABLE MSU (FORECASTED) / LPAR BILLABLE MSU (OPTIMIZED)

The billable MSU of LPAR that is the minimum value of LPAR 4HRA MSU and defined capacity at each hour compared with the forecast and the recommended optimization.

If the values of a scenario are all NULL, then in the chart, the line is not shown; in the table, this column is not shown for this scenario.

PEAK BILLABLE MSU

The peak billable MSU point during the specified date and time range of the selected scenario.

This metric appears only in charts. If LPAR BILLABLE MSU values of a scenario are all NULL, then the point is not shown.

GROUP CAPACITY NAME

The name of the group the LPAR is in at each hour.

This metric appears only in tables. If some or all values of a scenario are NULL, then the value is shown as `N/A` for this scenario.

GROUP CAPACITY LIMIT

The upper capacity limit of the group the LPAR is in at each hour.

This metric appears only in tables. If some or all values of a scenario are NULL, the value is shown as `0` for this scenario.

IS HIGHEST BILLABLE MSU

If the LPAR BILLABLE MSU is the highest MSU during the specified date and time range, then the value is "YES", and the font color of LPAR BILLABLE MSU and this field is red. Otherwise, the value is "NO".

This metric appears only in tables. If LPAR BILLABLE MSU values of a scenario are all NULL, all values are "NO" for this scenario.

PROC CAPACITY / PROC CAPACITY (FORECASTED) / PROC CAPACITY (OPTIMIZED)

The MSU rating of the LPAR using the logical CP count.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

If the values of a scenario are all NULL, then the line is not shown in the chart and this column is not shown in the table.

GUARANTEED CAPACITY

The guaranteed capacity in MSU on this LPAR. This value is calculated according to the weight assigned to the LPAR.

If the LPAR uses dedicated processors, then the value is shown as N/A.

If this value is double the maximum of the total MIPS usage, then this line does not appear in the chart. Showing the line would make the chart appear small.

The product utilized MSU value is retrieved from the DRL.MVSPM_PROD_T table, which is from SMF89 records. The product utilized MSU for some products might always be 0. In this case, the report is not displayed these products in charts or tables.

In charts, the product utilized MSU of each scenario is shown in the chart and table of the related scenario. So, observed product utilized MSU is shown in the chart and table for the observed scenario; forecasted product utilized MSU is shown in the chart and table for the forecasted scenario; optimized product utilized MSU is shown in the chart and table for the optimized scenario.

In charts, there is always a legend for product utilized MSU. If the product selected has no data or 0 MSU, PRODUCT UTILIZED MSU still appears in the legend.

In charts, if defined capacity is not set or is 0, then the line is not shown in the chart but DEFINED CAPACITY still appears in the legend.

In the Observed LPAR MSU Utilization table, products are grouped by IBM CORP or other vendor’s name.

There is a product that is named "OTHERS" whose MSU is the MSU utilized by NO89 products, non-IBM products, and LPAR management on the specified LPAR.

In the chart and table for the Observed scenario, only observed data is shown.

For the Forecasted scenario, in addition to forecasted data, OBSERVED LPAR BILLABLE MSU is also shown in the chart as a dotted line. OBSERVED LPAR UTILIZED MSU, LPAR 4HRA MSU, LPAR BILLABLE MSU, and IS HIGHEST BILLABLE MSU are also shown in the table.

For the Optimized scenario, in addition to optimized data, FORECASTED LPAR BILLABLE MSU is also shown in the chart as a dotted line. FORECASTED LPAR UTILIZED MSU, LPAR 4HRA MSU, LPAR BILLABLE MSU, and IS HIGHEST BILLABLE MSU are also shown in the table.

Observed Product MSU Utilization in SMF Collection Interval chart and table

This chart and table show product MSU utilization at SMF collection interval in the Observed scenario from an LPAR perspective. Clicking any of these charts or tables drills through to the Product MSU and Price report.

The chart displays the PRODUCT UTILIZED MSU metric, which is the MSU utilized by the selected product at each SMF collection interval.

The table displays all the same metrics in the Observed LPAR MSU Utilization table, as well as the product utilized MSU for each product and LPAR MSU utilization at each SMF collection interval.

This chart and table apply only to Observed scenario.

When you select the **SMF Collection Interval** option, this chart appears in addition to the Observed LPAR MSU Utilization chart.

When you select the **SMF Collection Interval** option, the Observed LPAR MSU Utilization table is replaced by this table.

Suggested Actions in Optimized Scenario table

This table shows the suggested actions that are related to the selected LPAR from the optimization stream's result.

FROM CPU SERIAL NO
: The CPU Serial No from which it is suggested to move the product's MSU.

FROM LPAR
: The LPAR name from which it is suggested to move the product's MSU.

FROM SYSTEM
: The system ID from which it is suggested to move the product's MSU.

TO CPU SERIAL NO
The CPU Serial No to which it is suggested to move the product's MSU.

TO LPAR
The LPAR to which it is suggested to move the product's MSU.

TO SYSTEM
The system ID to which it is suggested to move the product's MSU.

PRODUCT NAME
The product name of the product suggested to move.

PRODUCT ID
The product ID of the product suggested to move.

PRICING STRUCTURE
The pricing structure of the product suggested to move.

MOVE ALL
If all MSUs are suggested to be moved, then "YES"; otherwise, "NO".

MOVING MSU
If all MSUs are suggested to be moved, then "N/A"; otherwise, it's the number of MSU to be moved.

SCA: NO89 Products Matrix

This report provides a single place to review all of the NO89 products that use sub-capacity pricing. It also shows where these products are designated to be run (both the CPCs and the LPARs). It shows all of the locations where a product is running along with all of the other NO89 products to provide a planning resource for LPAR and CPC consolidations. From this report, you can also drill to the NO89 Product MSU and Price report for any product.

This report provides a listing of LPARs and CPCs where the unregistered IBM NO89 products run.

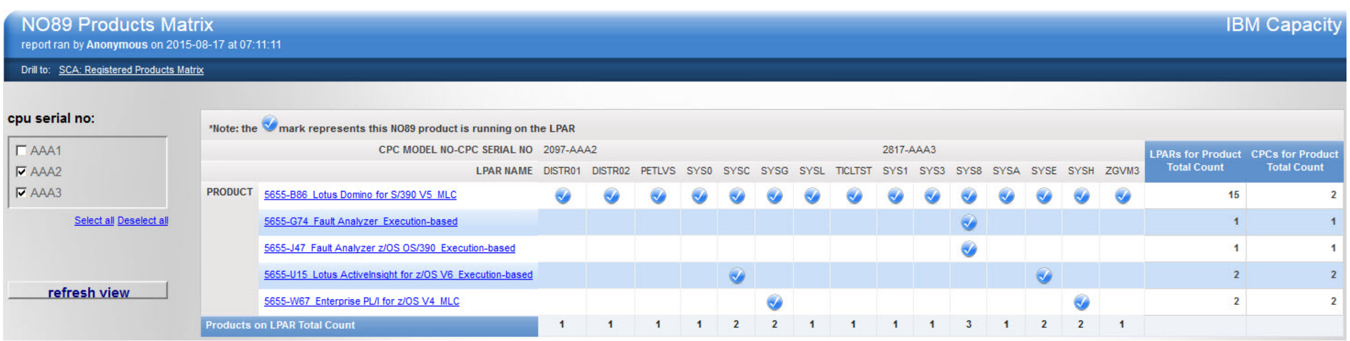

Figure 47: SCA: NO89 Products Matrix example

Source tables
[Presentation View].[NO89 Products]

[Presentation View].[Sub-Capacity Program Pricing Structures]

Drill through to the following reports
"SCA: Registered Products Matrix " on page 265

Report-specific prompts
cpu serial no

This report contains a single table, the NO89 Products Matrix table, which shows the LPARs and CPCs where the unregistered IBM NO89 products run. The table contains the following information:

CPC MODEL NO-CPC SERIAL NO
The zServer/CEC model number and serial number.

LPAR NAME
The LPAR name.

PRODUCT

The product display name, which uses the syntax [PRODUCT_ID]_[PRODUCT_NAME]_[PRODUCT_PRICING_STRUCTURE].

For more information about naming rules, see "Product display names" on page 107.

Products on LPAR Total Count
: The total number of NO89 products that are running on the LPAR.

LPARs for Product Total Count
: The total number of LPARs that the product is running.

CPCs for Product Total Count
: The total number of CPCs that the product is running.

In this table, if the product is running on this LPAR, then the cross point of LPAR NAME and PRODUCT is marked by a check mark .

SCA: NO89 Product MSU and Price

This report helps you understand and forecast your costs based on your MSU usage. It also shows you the monetary cost of your MSU for each NO89 product. You can compare the historical MSU usage with the forecasted usage, see upcoming trends, and take proactive measures to adjust your workload and capacity limits to help control your costs.

This report shows the billable MSUs for the IBM NO89 products, along with their prices.

This report can identify which LPARs are driving the cost of each NO89 product. You can then view the LPAR MSU report to see which other products are driving the MSU usage up on the LPARs where the NO89 product runs. By using a variety of reports, you can see what is driving a product's cost up. You can also determine where you can move a product to reduce costs.

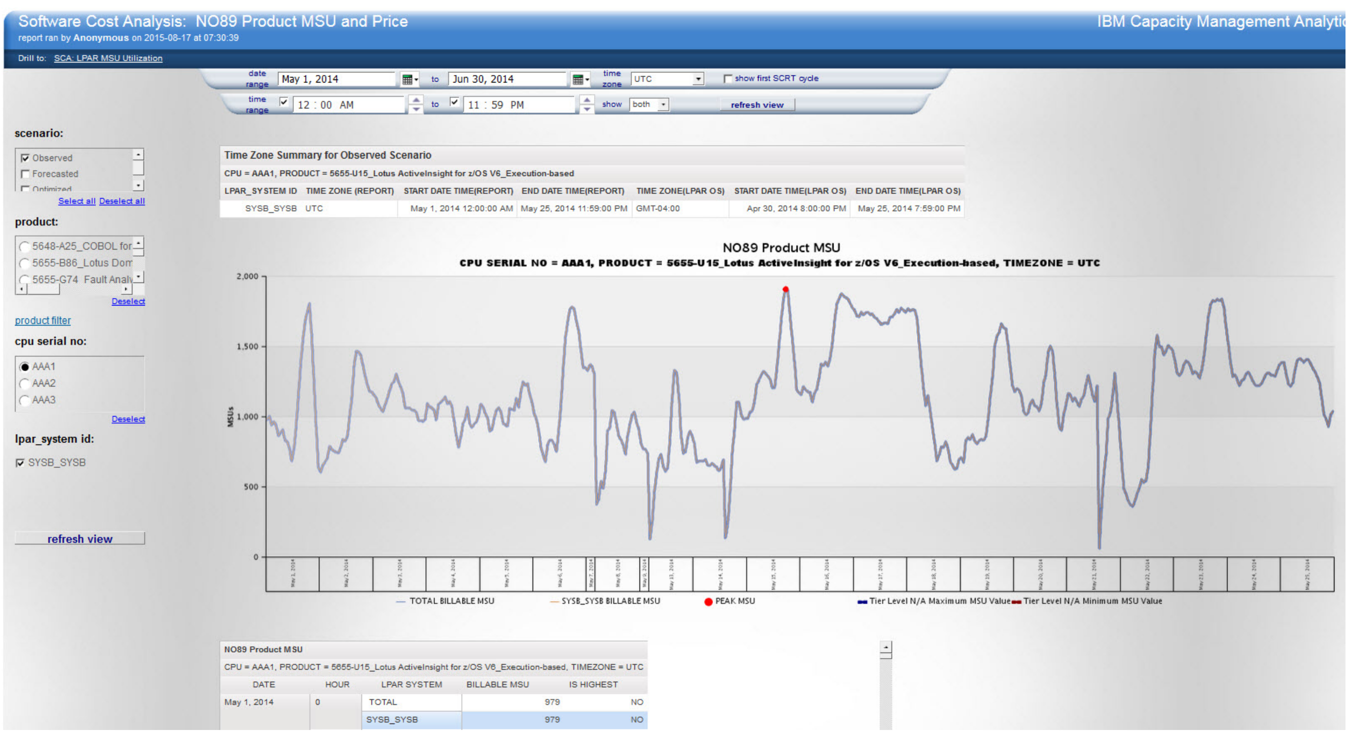

Figure 48: SCA: NO89 Product MSU and Price example

Source tables
: [Presentation View].[Customer Pricing Information]

[Presentation View].[LPAR MSU Forecast Metrics]

[Presentation View].[LPAR MSU Metrics]

[Presentation View].[LPAR MSU Optimization Metrics]

[Presentation View].[NO89 Products]

[Presentation View].[Parents of Sub-Capacity Reference-based programs]

[Presentation View].[Product Billable MSU Forecast Metrics]

[Presentation View].[Product Billable MSU Optimization Metrics]

[Presentation View].[Product MSU Utilization Metrics Without Z/OS]

[Presentation View].[Product MSU Utilization Metrics]

[Presentation View].[Prompts/LPAR MSU Prompt View]

[Presentation View].[Raw Product MSU Utilization Metrics]

[Presentation View].[Sub-Capacity Program Pricing Structures]

[Presentation View].[Timezone Metrics]

Drill through to the following report

"SCA: LPAR MSU Utilization " on page 249

"SCA: Registered Products Matrix " on page 265

Standard prompts

date range

time range

show first SCRT cycle

timezone

show

Report-specific prompts

scenario

product

product filter (search and select product from product lists)

cpu serial no

lpar_system id

The scenario prompt is independent of other prompts. This means that the values in other prompts, such as cpu serial no and lpar system id, will include all the possible values of the three scenarios. This also means you may sometimes see no data available even when cpu serial no, lpar system id, and product values are available in prompts.

All the charts and tables shown in this report are in hourly interval bands, except when the product is in IPLA Reference-based pricing structure. When the product is in IPLA reference-based pricing structure, the **lpar_system id** prompt is not available to select. All LPARs are counted in calculations.

The NO89 Product MSU and Price report contains the same tables and charts available in the "SCA: Product MSU and Price " on page 257 except tables and charts for z/OS and the optimized product suggested actions table.

SCA: Product MSU and Price

This report shows a registered IBM product's billable MSUs and monetary cost from a single CPC perspective for each scenario: Observed, Forecasted, and Optimized.

This report not only helps you understand and forecast your costs based on your MSU usage, but it also shows you the monetary cost of your MSU by product. You can compare the historical MSU usage with the forecasted usage and see upcoming trends. You can then take proactive action to adjust your workloads and capacity limits to help manage your bill. You can also identify which products drive the LPAR's usage and increase the cost of other products based on LPAR MSUs as a result. You can also see how a product runs on the other LPARs across your CPC to help determine consolidation plans.

This report allows you to focus on one product so that you can see its MSU usage across all LPARs on a CPC. You can quickly select another CPC and see the usage across those LPARs as well. You can also see whether the product is being used as expected on each LPAR and CPC. It is common for a product that is licensed for an LPAR to provide the occasional but necessary use of the product from that LPAR. It is also common for workload usage to change over time. This report can show you where there are opportunities to reduce your license requirements based on your product usage. This can be based on currently observed usage or historical usage.

You can also use the Forecast view to see what the pattern of usage will be in the future. The Forecast view provides a separate forecast for each LPAR that you select. The aggregate forecast for the product is shown for all selected LPARs. You can also view any number of LPARs you want and see the aggregate product forecast for that set of LPARs.

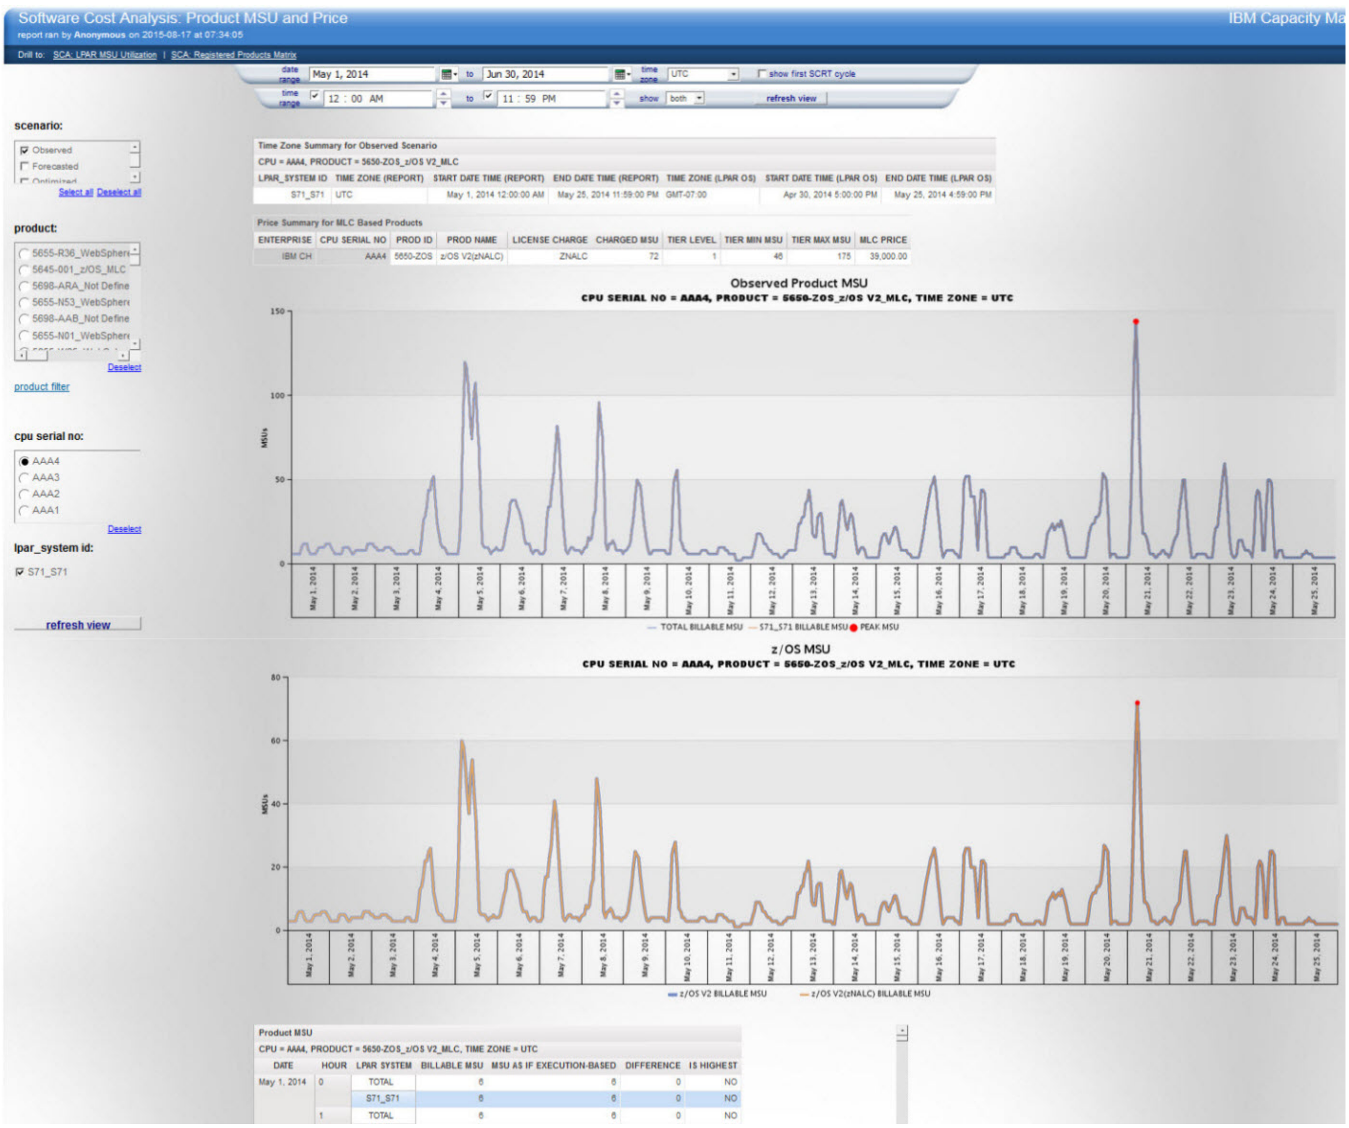

Figure 49: SCA: Product MSU and Price example

Source tables and views

[Presentation View].[Customer Pricing Information]

[Presentation View].[LPAR MSU Forecast Metrics]

[Presentation View].[LPAR MSU Metrics]

[Presentation View].[LPAR MSU Optimization Metrics]

[Presentation View].[Movement Suggestion by Optimization Metrics]

[Presentation View].[Product Billable MSU Forecast Metrics]

[Presentation View].[Product Billable MSU Optimization Metrics]

[Presentation View].[Product MSU Optimization Metrics]

[Presentation View].[Product MSU Utilization Metrics Without Z/OS]

[Presentation View].[Product MSU Utilization Metrics]

[Presentation View].[Prompts/Product Billable MSU Prompt View]

[Presentation View].[Raw Product MSU Utilization Metrics]

[Presentation View].[Sub-Capacity Program Pricing Structures]

[Presentation View].[Parents of Sub-Capacity Reference-based programs]

[Presentation View].[Timezone Metrics]

[Presentation View].[Z/OS Product MSU Utilization Metrics]

Drill through to the following report

"SCA: LPAR MSU Utilization " on page 249

Standard prompts

date range

time range

show first SCRT cycle

timezone

show

Report-specific prompts

scenario

product

product filter (search and select product from product lists)

cpu serial no

lpar_system id

The scenario prompt is independent of other prompts. This means that the values in other prompts, such as cpu serial no and lpar system id, includes all the possible values of the three scenarios. This also means you might sometimes see no data available even when cpu serial no, lpar system id, and product values are available in prompts.

All the charts and tables that are shown in this report are in hourly interval bands, except when the product is in IPLA Reference-based pricing structure. When the product is in IPLA reference-based pricing structure, the lpar_system id prompt is not available to select. All LPARs are counted in calculations.

The Product MSU and Price report includes the following charts and tables.

Time Zone Summary table

This table displays the time zone summary information for a selected LPAR and contains the following elements:

LPAR_SYSTEM ID

The LPAR name and its system ID.

TIME ZONE (REPORT)

The time zone used when displaying the report. The format is UTC -5:00.

START DATE TIME (REPORT)

The start date and time in reporting the time zone. The start date and time is the latest time between the start date and time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (REPORT)
The end date and time in reporting the time zone. The start date and time is the earliest time between the end date and time in the prompt and the latest date and time of available data in the database.

TIME ZONE (LPAR OS)
The local time of the selected LPAR. The format is GMT -5:00

Important: the LPAR time zone is the time zone of the LPAR OS and not the time of where the server is physically located.

START DATE TIME (LPAR OS)
The start date and time in the selected LPAR local time zone. The start date and time is the latest time between the start date and the time in the prompt and the earliest date and time of available data in the database.

END DATE TIME (LPAR OS)
The end date and time in the selected LPAR local time zone. The end date and time is the earliest time between the end date and the time in the prompt and the latest date and time of available data in the database.

Observed/Forecasted/Optimized Product MSU charts and tables

These charts and tables show a registered IBM product's billable MSUs for each scenario from the perspective of a single CPC and the LPARs it contains. You can click a chart to drill through to the LPAR MSU Utilization report.

The Optimized Product MSU charts and tables are not shown if the product is in IPLA Reference-based pricing structure.

In the chart and table for the Observed scenario, only observed data is shown.

For the Forecasted scenario, in addition to forecasted data, OBSERVED BILLABLE MSU, and Tier Level N Maximum/Minimum MSU Value for Observed are also shown in the chart as a dotted line, while OBSERVED BILLABLE MSU and IS HIGHEST are also shown in the table.

For the Optimized scenario, in addition to optimized data, FORECASTED BILLABLE MSU, and Tier Level N Maximum/Minimum MSU Value for Forecasted are also shown in the chart as a dotted line, while FORECASTED BILLABLE MSU and IS HIGHEST are also shown in the table.

For z/OS products, there are no Tier Level N Maximum/Minimum MSU Value lines in the chart.

The Product MSU charts and tables include the following information:

CPU SERIAL NO
The zServer/CEC serial number.

LPAR SYSTEM
The LPAR name and system ID.

PRODUCT
The product display name. For details on display name rules, see "Product display names" on page 107.

BILLABLE MSU
TOTAL BILLABLE MSU: the product billable MSU on a CPC at each hour.

LPAR BILLABLE MSU: the product billable MSU on each LPAR at each hour. Only selected LPARs are shown in the chart and table.

The calculation formula varies according to each pricing structure. If there are no records for a product at some hour, then in the table, the cell is empty. If the values of a scenario are all NULL, then in the chart, the line is not shown; in the table, this column is not shown for this scenario.

PEAK MSU
The peak billable MSU point during the date and time range that is selected for the related scenario. This metric appears only in charts. If BILLABLE MSU values of a scenario are all NULL, the point is not shown.

MSU AS IF EXECUTION-BASED
The billable MSU value by assuming the product uses execution-based pricing structure. This metric appears only in the table for the Observed scenario.

DIFFERENCE
The difference between BILLABLE MSU and MSU AS IF EXECUTION-BASED, which shows the amount of MSU you save from using the current pricing structure. This metric appears only in tables.

IS HIGHEST
If the product BILLABLE MSU is the highest MSU during the selected date and time range, then the value is "YES", and the product BILLABLE MSU and this field appear in red. Otherwise, the value is "NO". This metric exists only in tables. If product BILLABLE MSU values of a scenario are all NULL, the value is shown as "NO" for this scenario.

Tier Level N Maximum MSU Value

The maximum MSU value of the tier to which the products' highest MSU belongs. This metric appears only in charts. If BILLABLE MSU values of a scenario are all NULL, or the value of Maximum MSU is not defined, then the line is not shown in the chart, and the legend label for this line is displayed as "Tier Level N/A Maximum MSU Value ([SCENARIO_NAME])",

Tier Level N Minimum MSU Value

The minimum MSU value of the tier to which the products' highest MSU belongs. This metric appears only in charts. If BILLABLE MSU values of a scenario are all NULL, then the line is not shown in the chart, and the legend label for this line is displayed as "Tier Level N/A Minimum MSU Value ([SCENARIO_NAME])".

Observed/Forecasted/Optimized Price Summary tables

These tables show an IBM product's billable MSUs and the monetary cost from a single CPC perspective for each of the observed, forecasted, and optimized scenarios.

The displayed IPLA MSUs and price are only for the chosen CPC and not for the entire enterprise. To view IPLA MSUs and price for the entire enterprise, see the "Software Cost Analysis Summary workspace" on page 277.

The Price Summary tables require that price values are correctly loaded into the HCM.CUSTOMER_PRICE table to display related monetary values. If price values are not loaded in the HCM.CUSTOMER_PRICE table, a `No data available` message appears in the Price Summary tables.

The Price Summary tables contain the following information:

ENTERPRISE
The enterprise name.

CPU SERIAL NO
The zServer/CEC serial number.

PROD NAME:
The product name.

PROD ID:
The product ID.

CATEGORY
Sub-capacity pricing of Monthly License Charge (MLC) or System z International Product License Agreement (IPLA).

LICENSE CHARGE
Monthly License Charge metrics name of MLC products or Value Unit Exhibit name of IPLA products.

HIGHEST MSU
The peak billable MSU point during the specified date and time range of the related scenario. This metric exists only in the table for MLC products.

TIER LEVEL
The tier to which the highest MSU belongs. For MLC products, the tier is the one within a License Charge; for IPLA products, the tier is the one within a Value Unit Exhibit.

TIER MIN MSU
The minimum MSU of the TIER LEVEL.

TIER MAX MSU
The maximum MSU of the TIER LEVEL. If the value is not defined, then this column is not shown.

ASSIGNED ENTITLED MSU
The assigned entitled MSU of the CPC. This metric exists only in the table for IPLA products.

ASSIGNED ENTITLED VU
The assigned entitled VU of the CPC. This metric exists only in the table for IPLA products.

TOTAL USAGE MSU
The peak billable MSU point during the specified date and time range of the related scenario. This metric exists only in the table for IPLA products.

TOTAL USAGE VU
The usage of VU that is converted from TOTAL USAGE MSU. This metric exists only in the table for IPLA products.

DELTA MSU
(TOTAL USAGE MSU) – (ASSIGNED ENTITLED MSU). This metric exists only in the table for IPLA products.

DELTA VU
(TOTAL USAGE VU) - (ASSIGNED ENTITLED VU). This metric exists only in the table for IPLA products.

MONETARY ENTITLED
The assigned entitled monetary value of the CPC. This metric exists only in the table for IPLA products.

MONETARY USAGE
The billable monetary value. If DELTA VU is negative, then (MONETARY USAGE) = (MONETARY ENTITLED); otherwise, (MONETARY USAGE) = (DELTA VU * PRICE PER VU). This metric exists only in the table for IPLA products.

MONETARY DELTA
(MONETARY USAGE) – (MONETARY ENTITLED). This metric exists only in the table for IPLA products.

MONETARY SS
S&S monetary value. MONETARY SS PER MONTH * the natural months of the specified date range. This metric exists only in the table for IPLA products.

TOTAL
Total monetary value. For IPLA products, if (MONETARY DELTA) is more than 0, then (MONETARY DELTA) + (MONETARY SS); otherwise,MONETARY SS. For MLC products, the total monetary value that is calculated by the method of its license charge.

Observed/Forecasted/Optimized Parent Billable MSU chart and Billable MSU Summary for Reference-Based Products table

These charts and tables show the billable MSUs for each parent product that has IBM Reference-Based charging from a single CPC perspective for each of the observed, forecasted, and optimized scenarios.

These charts and tables exist only for IPLA Reference-Based products. If there is no record for any parent of a child product, then the chart or table displays a `No data available` message.

You can click a parent product name in the Billable MSU Summary for Reference-Based Products table, or click a piece of the pie of the Parent Billable MSU chart, to drill through to the "SCA: Product MSU and Price " on page 257.

The charts and tables contain the following information:

PARENT PEAK MSU
: The peak MSU of each parent product during the specified date and time range.

CHILD BILLABLE MSU
: The total MSU of all the parents' peak MSU.

PEAK ZOS MSU
: The peak MSU of z/OS on the CPC during the specified date and time range.

BILLABLE MSU
: The final billable MSU that is the minimum value between CHILD BILLABLE MSU and PEAK ZOS MSU.

Observed/Forecasted/Optimized z/OS MSU chart and zNALC/Traditional Workloads tables

These charts and tables show MSU summary and details for zNALC and traditional workloads for a z/OS product from a single CPC.

There is one chart that is named z/OS MSU chart, and three tables: MSU Summary for zNALC and Traditional Workload table, Price Summary for zNALC and Traditional Workload table, and Billable MSU Details for zNALC and Traditional Workloads table.

These charts and tables exist only for z/OS products. Particularly, the MSU Summary for zNALC and Traditional Workload table and Price Summary for zNALC and Traditional Workload table exist only when there are both zNALC and Traditional workloads are running on the CPC during the specified date and time range.

In the z/OS MSU chart and Billable MSU Details for zNALC and Traditional Workloads table, for the Observed scenario, only observed data is shown. For the Forecasted scenario, in addition to forecasted data, OBSERVED BILLABLE MSU is also shown in the z/OS MSU chart as a dotted line, while OBSERVED BILLABLE MSU and IS HIGHEST appear in the Billable MSU Details for zNALC and Traditional Workloads table. For the Optimized scenario, in addition to optimized data, FORECASTED BILLABLE MSU is also shown in the z/OS MSU chart as a dotted line, while FORECASTED BILLABLE MSU and IS HIGHEST appear in the Billable MSU Details for zNALC and Traditional Workloads table.

The Price Summary for zNALC and Traditional Workload tables requires that price values are correctly loaded into the HCM.CUSTOMER_PRICE table to display related monetary values. If price values are not loaded in the HCM.CUSTOMER_PRICE table, a `No data available` message appears in the tables.

The MSU Summary for zNALC and Traditional Workload table contains the following information:

CPU SERIAL NO
: The zServer/CEC serial number.

PROD ID
: The product ID.

PROD NAME
: The product name.

FULL CAPACITY MSU
: Rated capacity of the CPC.

z/OS HIGHEST MSU
: The highest MSU of the z/OS product regardless of zNALC workloads or traditional workloads.

TRAD HIGHEST MSU
: The highest MSU of the z/OS product for traditional workloads.

zNALC HIGHEST MSU
: The highest MSU of the z/OS product for zNALC workloads.

CHARGED TRAD MSU
: The MSU used in monetary value calculations for traditional workloads.

CHARGED zNALC MSU
: The MSU used in monetary value calculations for zNALC workloads.

The Price Summary for zNALC and Traditional Workload table contains the following information:

LICENSE CHARGE
Monthly License Charge metrics name.

CHARGED MSU
The MSU used in monetary value calculations.

TIER LEVEL
The tier within a License Charge to which the CHARGED MSU belongs.

TIER MIN MSU
The minimum MSU of the TIER LEVEL.

TIER MAX MSU
The maximum MSU of the TIER LEVEL. If the value is not defined, then this column is not shown.

MLC PRICE
Total monetary value.

The z/OS MSU charts and Billable MSU Details for zNALC and Traditional Workloads tables contain the following information:

z/OS PROD NAME
The z/OS product name. The z/OS product name presents the total z/OS MSU, while the z/OS product name with (zNALC) or (Traditional) presents the MSU for each type of workload.

BILLABLE MSU
The sum of billable MSU for each LPAR the workload is running at each hour. If the values of a scenario are all NULL, then in the chart, the line is not shown; in the table, the column is not shown for this scenario.

PEAK MSU
The peak billable MSU point during the specified date and time range of the related scenario. This metric appears only in charts. If BILLABLE MSU values of a scenario are all NULL, the point is not shown.

IS HIGHEST
If the BILLABLE MSU is the highest MSU during the specified date and time range, then the value is "YES" and the font color of product BILLABLE MSU and this field is red. Otherwise, the value is "NO". If BILLABLE MSU values of a scenario are all NULL, then the value is shown as "NO" for the scenario. This metric appears only in tables.

Optimized Product Suggested Actions table

This table shows the suggested actions that are related to the selected product from the optimization stream's results that are related to the selected product.

The table contains the following information:

PRODUCT CATEGORY
Subcapacity pricing of Monthly License Charge (MLC) or System z International Product License Agreement (IPLA).

PRODUCT NAME
The product name of the product suggested to move.

PRODUCT ID
The product ID of the product suggested to move.

PRODUCT PRICING STRUCTURE
The pricing structure of the product suggested to move.

FROM CPU SERIAL NO
The CPU Serial No from which it is suggested to move the product’s MSU.

FROM LPAR
The LPAR name from which it is suggested to move the product’s MSU.

FROM SYSTEM
The system ID from which it is suggested to move the product’s MSU.

TO CPU SERIAL NO
: The CPU Serial No to which it is suggested to move the product's MSU.

TO LPAR
: The LPAR to which it is suggested to move the product's MSU.

TO SYSTEM
: The system ID to which it is suggested to move the product's MSU.

MOVING MSUs
: If all MSUs are suggested to be moved, then "N/A"; otherwise, it is the number of MSU to be moved.

MOVE ALL
: If all MSUs are suggested to be moved, then "YES"; otherwise, "NO".

SCA: Registered Products Matrix

This report provides a listing of LPARs and CPCs where the SCA registered products are running.

This matrix view of where the products are running allows a quick assessment of whether the products are running on the correct LPARs. The matrix also shows which products are co-located. This helps you plan your software costs and to determine where you might need to license a new product.

The report shows which products are forecasted in a forecast stream. It also shows a new matrix after a product has been moved because of an optimization stream suggestion. The report shows observed, forecasted, and optimized scenarios.

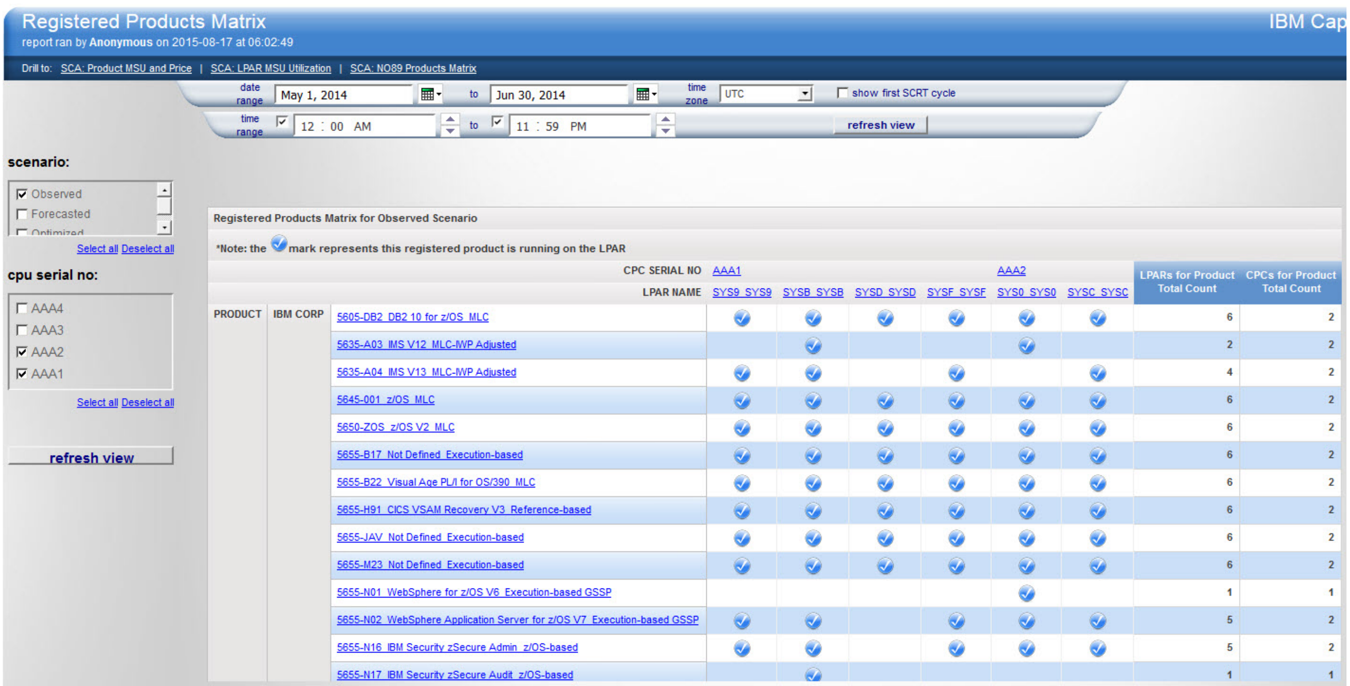

Figure 50: SCA: Registered Products Matrix example

Source tables
: [Presentation View].[Prompts/Product Billable MSU Prompt View]
: [Presentation View].[Raw Product MSU Utilization Metrics]
: [Presentation View].[Customer Pricing Information]
: [Presentation View].[Sub-Capacity Program Pricing Structures]
: [Presentation View].[Product MSU Forecast Metrics]
: [Presentation View].[Product MSU Optimization Metrics]

Drill through to the following reports

"SCA: Product MSU and Price " on page 257

"SCA: LPAR MSU Utilization " on page 249

"SCA: NO89 Products Matrix " on page 255

Standard prompts

date range

time range

show first SCRT cycle

time zone

Report-specific prompts

scenario

cpu serial no

This report shows a table for each scenario selected. The tables contain the following information:

CPC SERIAL NO
The zServer/CEC serial number.

LPAR NAME
The LPAR name.

PRODUCT

The product display name, which uses the syntax [PRODUCT_ID]_[PRODUCT_NAME]_[PRODUCT_PRICING_STRUCTURE].

See "Product display names" on page 107 for details on naming rules.

PRODUCT OWNER
The product's owner, for example, IBM CORP.

Products on LPAR Total Count
The total number of 89 products that are running on the LPAR.

LPARs for Product Total Count
The total number of LPARs that the product is running.

CPCs for Product Total Count
The total number of CPCs that the product is running.

In this table, if the product is running on this LPAR, then the cross point of LPAR NAME and PRODUCT is marked by a check mark .

WLM: Delays by Importance Level

This report shows you the effect of delays for each WLM importance level that is running on the system. Delaying high importance workloads can result in missed Service Level Agreements (SLAs) which has a higher cost than delaying less important workloads. It is important to monitor these delays so that you can take timely action to control your costs and maintain your SLAs.

This report shows the percentage of time when delays (CPU, IO, and paging) are negatively affecting work. Delays at the service class period level are aggregated and displayed for each WLM importance level.

You can use this report to ensure that your most important work is not negatively affected by delays.

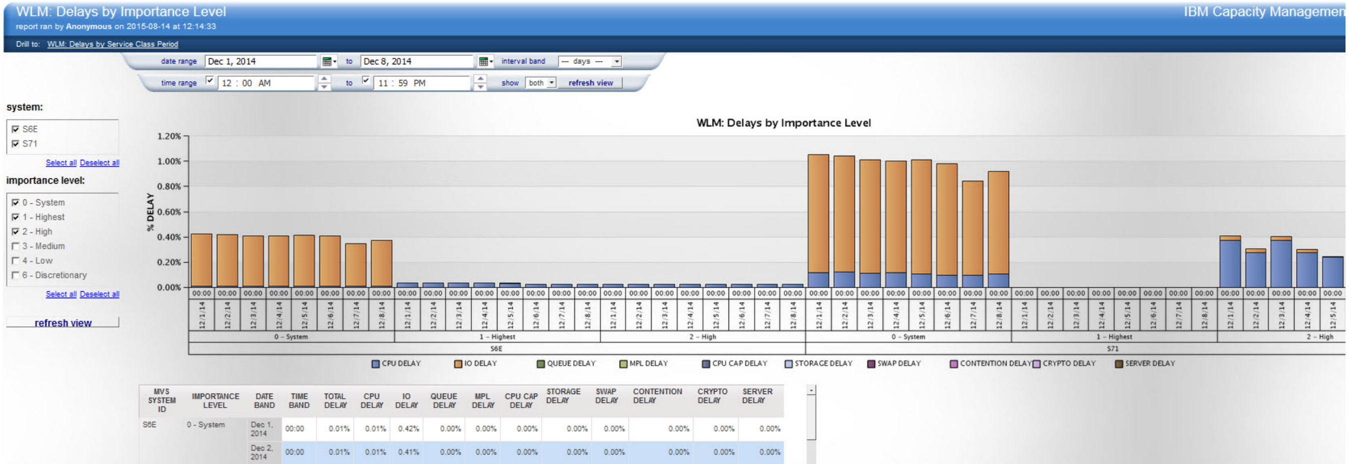

Figure 51: WLM: Delays by Importance Level example

Source tables
: [Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports
: "WLM: Delays by Service Class Period" on page 268

Standard prompts:
: date range
: time range
: interval band
: show

Report-specific prompts
: system
: importance level

The WLM: Delays by Importance Level report contains the following information:

system
: The z/OS system id.

importance level
: The WLM importance level that is assigned to a service class period.

TOTAL DELAY
: Total percentage of samples where the reported WLM service class period has been found delayed for any used type.

CPU DELAY
: The percentage of samples when a CPU delay existed (a TCB or SRB is waiting to be dispatched).

IO DELAY
: The percentage of samples when an IO delay existed (a TCB or SRB has initiated an I/O request that is delayed obtaining a path to the device). This includes IOSQ and Q+PEND components.

QUEUE DELAY
: The percentage of samples when a queue delay existed (work is waiting for a server address space or batch initiator).

MPL DELAY
: The percentage of samples when an MPL delay existed (ready to be dispatched but a swap-in has not started).

CPU CAP DELAY
: The percentage of samples when a CPU capping delay existed (a TCB or SRB is marked non-dispatchable because a resource group maximum is being enforced or because of discretionary

goal management). That is, if certain types of work are overachieving their goals, that work might be capped so that the resources can be diverted to run discretionary work.

STORAGE DELAY

The percentage of samples when a storage-related delay existed. Storage-related delays can be attributed to auxiliary paging from private, auxiliary paging from common, auxiliary paging from cross memory, auxiliary paging from VIO, auxiliary paging from standard hiperspaces, or auxiliary paging from ESO hiperspace (a page that is being read was not in the ESO hiperspace and must be read from DASD by the program that is managing the hiperspace).

SWAP DELAY

The percentage of samples when a swap delay existed (a Swap-In has started but not completed).

CONTENTION DELAY

The percentage of samples when a contention-related delay existed (work is waiting for resources).

CRYPTO DELAY

The percentage of samples when a Crypto-related delay existed (a TCB or SRB was found to be waiting for a Crypto Asynchronous Message Processor, an Adjunct Processor, or a processor feature queue).

SERVER DELAY

The percentage of samples when a server-related delay existed. Server-related delays can be attributed to a private area paging delay for a server address space, a VIO paging delay for a server address space, a hiperspace paging delay for a server address space, or a Swap-in delay for a server address space.

WLM: Delays by Service Class Period

This report shows where delays are having a negative impact on performance so that you can take corrective action. Monitoring delays is important because delays can cause you to miss SLAs and incur a penalty.

This report shows the percentage of time that delays (CPU, IO, and paging) are negatively impacting work. Delays are displayed at the service class period level. Use this report to ensure that your most important workloads (service classes) are not negatively impacted by delays.

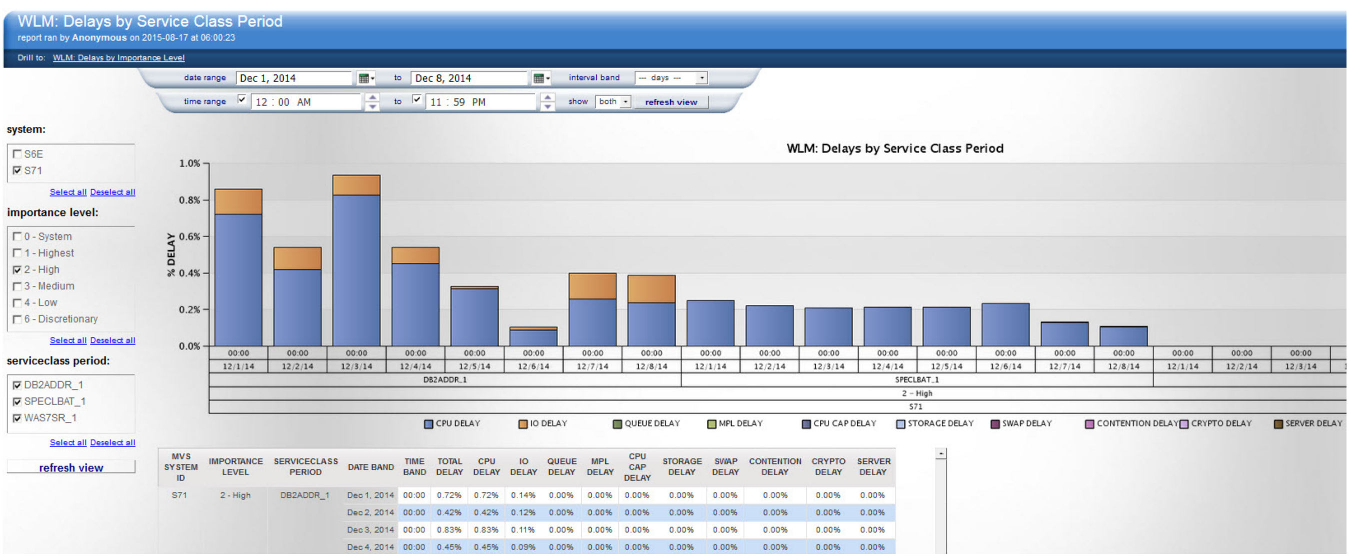

MVS SYSTEM ID	IMPORTANCE LEVEL	SERVICECLASS PERIOD	DATE BAND	TIME BAND	TOTAL DELAY	CPU DELAY	IO DELAY	QUEUE DELAY	MPL DELAY	CPU CAP DELAY	STORAGE DELAY	SWAP DELAY	CONTENTION DELAY	CRYPTO DELAY	SERVER DELAY
S71	2 - High	DB2ADDR_1	Dec 1, 2014	00:00	0.72%	0.72%	0.14%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
			Dec 2, 2014	00:00	0.42%	0.42%	0.12%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
			Dec 3, 2014	00:00	0.83%	0.83%	0.11%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
			Dec 4, 2014	00:00	0.45%	0.45%	0.09%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%

Figure 52: WLM: Delays by Service Class Period example

Source tables

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports

"WLM: Delays by Importance Level" on page 266

Standard prompts:
date range

time range

interval band

show

Report-specific prompts
system

importance level

serviceclass period

The WLM: Delays by Service Class Period report contains the following information:

system
The z/OS system id.

importance level
The WLM importance level that is assigned to a service class period.

serviceclass period
The WLM service class name and period. The concatenation of columns SERVICE_CLASS and SERV_CLASS_PERIOD.

importance level
The WLM importance level that is assigned to a service class period.

CPU DELAY
The percentage of samples when a CPU delay existed (a TCB or SRB is waiting to be dispatched).

IO DELAY
The percentage of samples when an IO delay existed (a TCB or SRB has initiated an I/O request that is delayed obtaining a path to the device). This includes IOSQ and Q+PEND components.

QUEUE DELAY
The percentage of samples when a queue delay existed (work is waiting for a server address space or batch initiator).

MPL DELAY
The percentage of samples when an MPL delay existed (ready to be dispatched but a swap-in has not started).

CPU CAP DELAY
The percentage of samples when a CPU capping delay existed (a TCB or SRB is marked non-dispatchable because a resource group maximum is being enforced or because of discretionary goal management). That is, if certain types of work are overachieving their goals, that work might be capped so that the resources can be diverted to run discretionary work.

STORAGE DELAY
The percentage of samples when a storage-related delay existed. Storage-related delays can be attributed to auxiliary paging from private, auxiliary paging from common, auxiliary paging from cross memory, auxiliary paging from VIO, auxiliary paging from standard hiperspaces, or auxiliary paging from ESO hiperspace (a page that is being read was not in the ESO hiperspace and must be read from DASD by the program that is managing the hiperspace).

SWAP DELAY
The percentage of samples when a swap delay existed (a Swap-In has started but not completed).

CONTENTION DELAY
The percentage of samples when a contention-related delay existed (work is waiting for resources).

CRYPTO DELAY
The percentage of samples when a Crypto-related delay existed (a TCB or SRB was found to be waiting for a Crypto Asynchronous Message Processor, an Adjunct Processor, or a processor feature queue).

SERVER DELAY
: The percentage of samples when a server-related delay existed. Server-related delays can be attributed to a private area paging delay for a server address space, a VIO paging delay for a server address space, a hiperspace paging delay for a server address space, or a Swap-in delay for a server address space.

WLM: Performance Indexes

This report helps you align goal achievement metrics with overall business performance measurements, such as the cost associated with system performance (for example, adequate response times).

This report analyzes how well Workload Manager (WLM) is doing with achieving its goals. The report shows the distribution of WLM performance indexes across eight performance index ranges. It provides a view of how often WLM goals are met (performance index equal to or less than 1) or missed (performance index greater than 1), and to what degree.

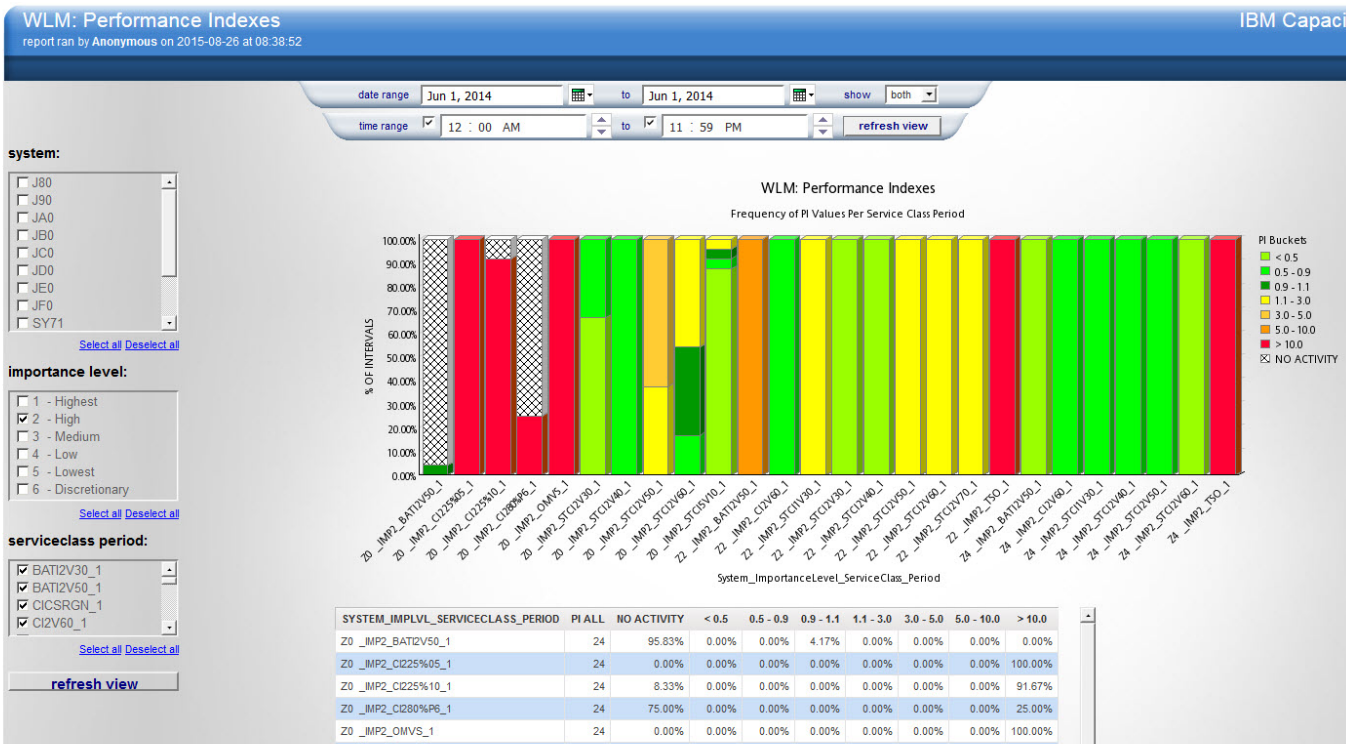

Figure 53: WLM: Performance Indexes example

Source tables
: [Presentation View].[WLM Service Class Period Performance Metrics]
: [Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following reports
: There is no drill through available on this report.

Standard prompts:
: date range
: time range
: show

Report-specific prompts
: system
: importance level
: serviceclass period

The WLM: Performance Indexes report contains the following information:

system
The z/OS system id.

importance level
The WLM importance level that is assigned to a service class period.

serviceclass period
The WLM service class name and period. The concatenation of columns SERVICE_CLASS and SERV_CLASS_PERIOD.

SYSTEM_IMPLVL_SERVICECLASS_PERIOD
The concatenation of system, importance_level and serviceclass_period.

PI Bucket (<0.5)
The percentage of intervals analyzed where the PI was less than 0.5.

PI Bucket (0.5 – 0.9)
The percentage of intervals analyzed where the PI was greater than or equal to 0.5 and less than 0.9.

PI Bucket (0.9 – 1.1)
The percentage of intervals analyzed where the PI was greater than or equal to 0.9 and less than 1.1.

PI Bucket (1.1 – 3.0)
The percentage of intervals analyzed where the PI was greater than or equal to 1.1 and less than 3.0.

PI Bucket (3.0 – 5.0)
The percentage of intervals analyzed where the PI was greater than or equal to 3.0 and less than 5.0.

PI Bucket (5.0 – 10.0)
The percentage of intervals analyzed where the PI was greater than or equal to 5.0 and less than 10.0.

PI Bucket (>= 10.0)
The percentage of intervals analyzed where the PI was greater than or equal to 10.0.

zIIP/zAAP What-ifs - LPAR Level workspace

This workspace puts analytical power at your fingertips. Scenario modeling lets you decide how to allocate processing costs to the zIIP/zAAP engines where it is most efficient and cost effective. Operational improvements that are based on the best scenario can drive significant cost reductions and savings.

The workspace displays the observed zIIP and zAAP usage along with eligible workloads in MIPS and engines. You can use this workspace to see the impact of moving around workloads. To help you get started, two what-if scenarios are included that show

- Increased usage from moving zIIP/zAAP-eligible MIPS
- Engines from CPs on zIIPs/zAAPs

You can also use this workspace to see the impact of adding zIIP/zAAPs to your environment or the impact of adding more zIIPs/zAAPs. All MIPS are normalized to a CP while engines are represented as actual engines. If zIIP/zAAP engines are not present, then the estimated zIIP/zAAP engines for the what-if scenarios that are based on the same speed as the installed CPs.

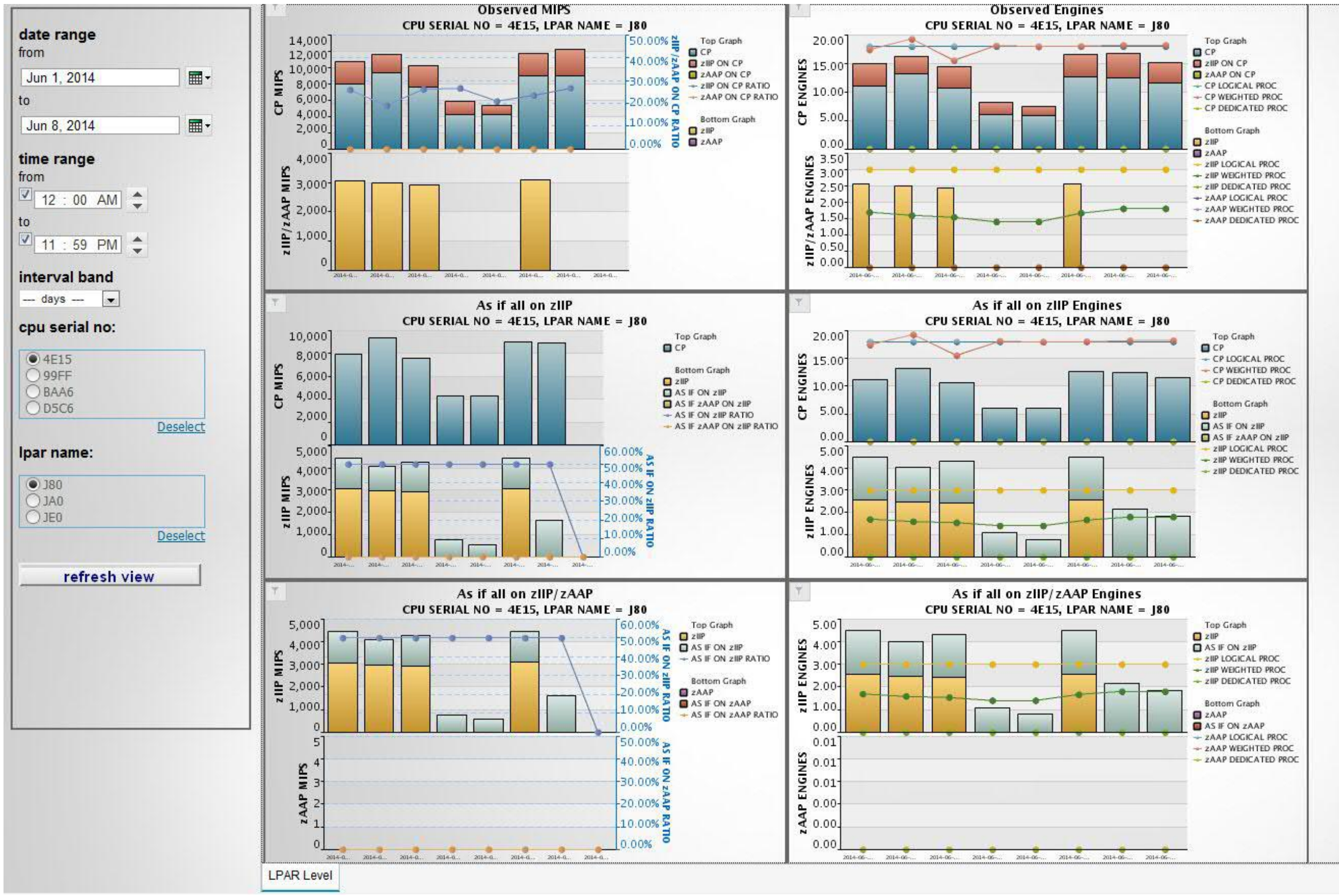

Figure 54: zIIP/zAAP What-ifs - LPAR Level workspace sample

Source tables

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

[Presentation View].[LPAR Level Metrics]

Drill through to the following report:

"CPU: MIPS Used - Service Class Period Level" on page 220

Workspace prompts

date range

time range

interval band

cpu serial no

lpar name

The **zIIP/zAAP What-ifs - LPAR Level workspace** includes the following charts and tables.

In some cases, a value can be too small to observe graphically on a chart. You can review all charted values in the **zIIP/zAAP What-ifs table**, which is included in the workspace.

The engine numbers that are shown in a chart or table is estimated from MIPS used, and is *not* the actual required engine numbers.

When you plan the configuration of your zIIP/zAAP, consider the following aspects. This list of considerations is not definitive, and you might have other considerations unique to your model or environment.

- With more engine numbers installed, LPAR management time is higher, so that more engines are required. The more processors that are installed, the fewer MIPS per processor are supported.

- Consider peak and average CPU utilization. Charts display average, but many systems need to support peak workload. (Peak-to-Average Ratio (PAR) and Saturation Design Point (SDP).)
- With greater CPU utilization, MIPS per transaction is higher. Generally, the CPU utilization that is required to maintain response time is different when engine numbers are different. With more engines, less utilization is required due to flexibility.
- Consider failover support requirements among all systems.
- Consider future workload increases or decreases.

Observed MIPS chart

This chart shows the average Observed MIPS utilization by CP and zIIP/zAAP. The top half of the chart shows the observed CP MIPS with the eligible zIIP/zAAP MIPS along with a ratio of zIIP/zAAP on CP. The bottom half shows just the observed zIIP/zAAP MIPS.

As if all on zIIP chart

This chart shows the what-if scenario for moving all the zIIP/zAAP eligible and all zAAP work onto zIIPs (zAAP on zIIP). The top half of the chart shows CP MIPS without the zIIP/zAAP eligible work. The bottom half of the chart shows all the zIIP/zAAP eligible work added to the zIIPs. Ratios of zIIP / zAAP eligible work to the total zIIP work is also shown.

As if all on zIIP/zAAP chart

This chart shows the what-if scenario that assumes all zIIP/zAAP eligible workloads that are run on zIIP/zAAP. The top and bottom half show the new total MIPS of the zIIP and zAAP eligible work added to the existing zIIP and zAAP MIPS separately. A ratio of the zIIP/zAAP eligible MIPS to total zIIP/zAAP work is also shown.

Observed Engines chart

This chart shows utilization based on the average number of observed engines for each processor type in respect to logical, dedicated, and weighted processor configuration for the LPAR. The top half of the chart shows engine counts for CPs including eligible zIIP/zAAP work. The bottom half shows the various engine counts for zIIP/zAAPs.

As if all on zIIP Engines chart

This chart illustrates the what-if scenario for moving all zIIP/zAAP eligible work onto zIIPs expressed in engine counts with respect to the LPAR configuration for logical, weighted, and dedicated processors. The top half of the chart shows various engine counts for CPs without zIIP/zAAP eligible work. The bottom half shows various engine counts of the zIIP/zAAP eligible work that was added to the existing zIIP counts.

If zIIP/zAAP engines are not present, then the estimated zIIP/zAAP engines are based on the same speed of the installed CPs.

As if all on zIIP/zAAP Engines chart

This chart illustrates the what-if scenario for moving all zIIP/zAAP eligible work onto zIIPs/zAAPs expressed in engine counts with respect to the LPAR configuration for logical, weighted, and dedicated processors. The top and bottom half show various engine counts of the zIIP and zAAP eligible work added to the existing zIIP and zAAP counts separately.

If zIIP/zAAP engines are not present, then the estimated zIIP/zAAP engines are based on the same speed of the installed CPs.

zIIP/zAAP What-ifs table

This table shows observed zIIP/zAAP usage along with eligible usage in MIPS and engines. In addition, two what-if scenarios are included showing the increased usage of zIIP/zAAP eligible MIPS from CPs on zIIPs/zAAPs. All MIPS are normalized to a CP while engines are represented as actual engines.

If zIIP/zAAP engines are not present, then the estimated zIIP/zAAP engines for the what-if scenarios are based on the same speed of the installed CPs.

Enterprise Dashboard workspace

The Enterprise Dashboard workspace brings together information for all of your servers across the enterprise to provide an overview of your inventory and utilization.

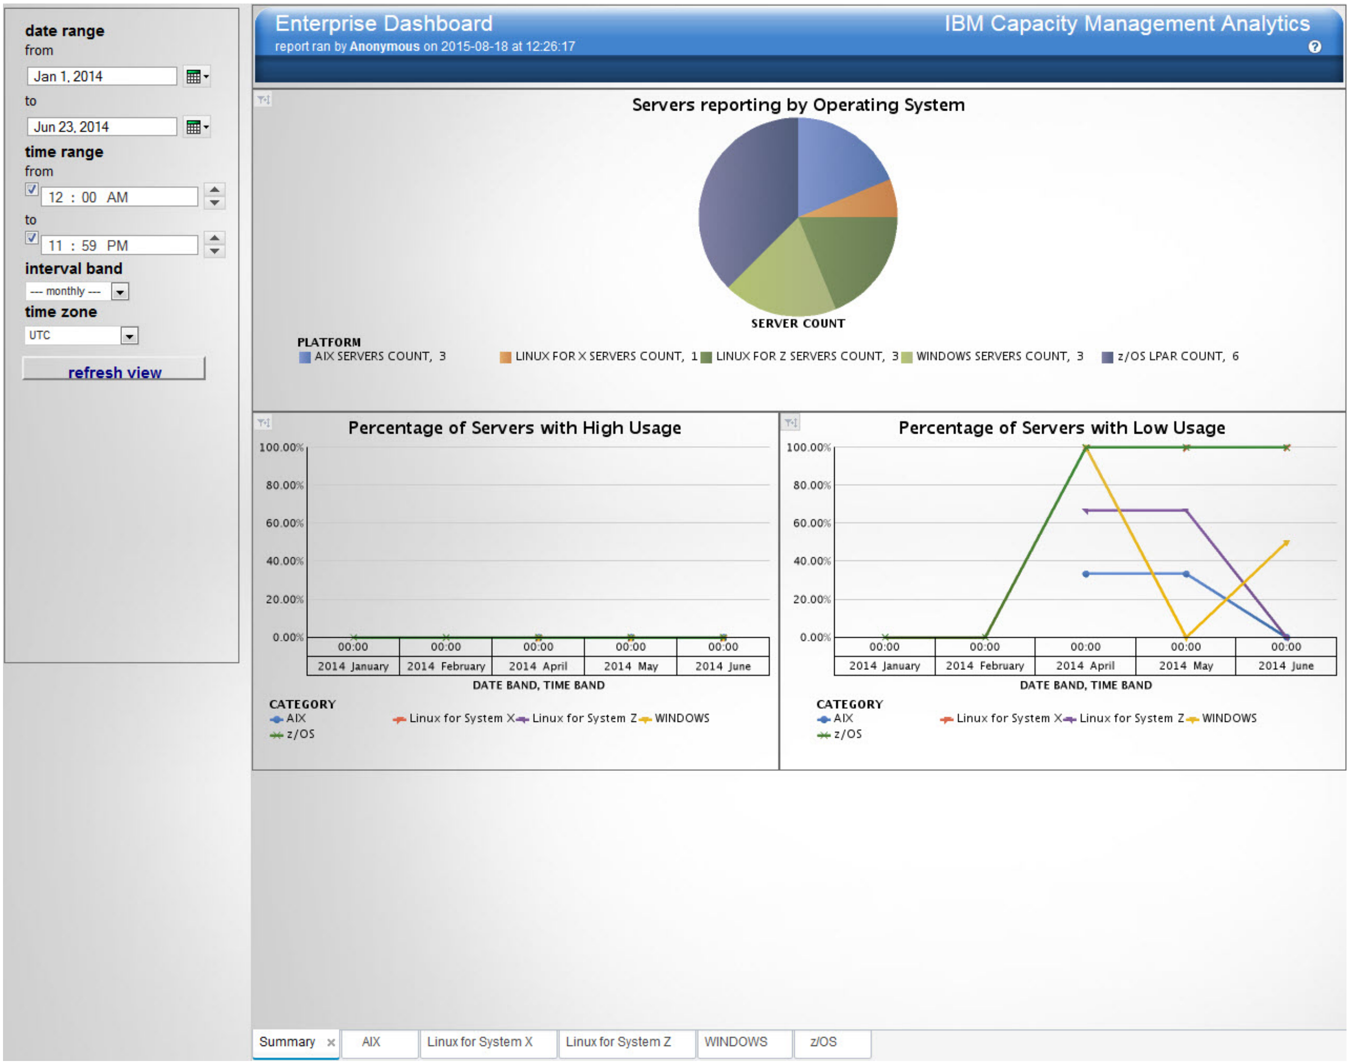

Figure 55: Summary tab example

Source tables

[Dimensional View].[Linux for System X CPU Hour Metrics]

[Dimensional View].[Linux for System Z CPU Hour Metrics]

[Dimensional View].[z/OS System Level Metrics]

[Dimensional View].[WIN OPERATING CPU Hour Metrics]

[Dimensional View].[AIX CPU Raw Hour Metrics]

Standard prompts

date range

time range

interval band

time zone

The charts and tables available on the Enterprise Dashboard are grouped on the following tabs:

- **Summary**
- **AIX**
- **Linux for System x**
- **Linux for System z**
- **WINDOWS**
- **z/OS**

Summary tab

This enterprise-wide view shows the inventory count of servers that are reporting data metrics for each platform. It also shows the data for each individual server. Its peak CPU usage is tracked over time. This allows you to see what percent of the servers in each platform are running hot (running with a peak usage over 90% busy). It also lets you track SLAs for what percent of servers can be running hot for each platform. You can also use that same Peak CPU usage tracking and show the percent of servers in each platform that are very low CPU usage. If the peak usage is less than 10% busy, that is considered to be low usage. The percent of servers that are running very low CPU usage can be tracked to monitor progress of right sizing the servers over time. These thresholds are customizable.

This report provides an overview of all the servers that have data that is being collected about its capacity. You can change the data range of the chart to see the changes that have occurred by platform. The two reports below the inventory pie chart are meant to provide a quick glance at each platform's CPU utilization. If you have a high percentage of servers that have very low utilization, that is typically a target for cost reduction. If you have a high percentage of servers that are very high in CPU usage, that shows that you might need to spread out some work or increase the capacity of some servers. These high-level views allow tracking of how you are doing overall towards the balance of cost and SLA adherence. Over time, you might want to modify these reports to filter out specific servers that are supposed to be low usage or high usage.

- Servers Reporting by Operating System chart

  This chart shows the number of servers on which data was found during the reporting date/time range, by operating system.

- Percentage of Servers with High Usage chart

  This chart shows the percentage of the total number of servers that have a peak utilization greater than the high usage threshold. The default high usage threshold is 90% busy. There is a separate high usage report for each operating system group of servers. Peak utilization is calculated by using the 90 percentile CPU Busy value over the interval band.

- Percentage of Servers with Light Usage chart

  This chart shows the percentage of the total number of servers that have a peak utilization lower than the low usage threshold. The default low usage threshold is 10% busy. There is a separate low usage report for each operation system group of servers. Peak utilization is calculated by using the 90 percentile CPU Busy value over the interval band.

From the Summary tab, you can see the percentage of servers at high usage and low usage. The platform tabs provide the list of the highest usage servers for the selected date range. It also shows their history of usage over time so you can see whether it is consistently at that level of usage or not. The lowest usage servers for the selected date are also shown. The server history of usage is shown over time so you can see whether it is consistently low usage or not. For both the high usage and low usage servers, this might be an action list of servers that require further review.

AIX tab

- 10 Highest CPU Usage Peaks table

  This table shows the top 10 AIX servers with the highest 90 percentile CPU busy values.
- 10 Lowest CPU Usage Peaks table

  This table shows the bottom 10 AIX servers with the lowest 90 percentile CPU busy values.

Linux for System x tab

- 10 Highest CPU Usage Peaks table

  This table shows the top 10 Linux for System x servers with the highest 90 percentile CPU busy values.
- 10 Lowest CPU Usage Peaks table

  This table shows the bottom 10 Linux for System x servers with the lowest 90 percentile CPU busy values.

Linux for System z tab

- 10 Highest CPU Usage Peaks table

  This table shows the top 10 Linux for System z servers with the highest 90 percentile CPU busy values.
- 10 Lowest CPU Usage Peaks table

  This table shows the bottom 10 Linux for System z servers with the lowest 90 percentile CPU busy values.

Windows tab

- 10 Highest CPU Usage Peaks table

  This table shows the top 10 Windows servers with the highest 90 percentile CPU busy values.
- 10 Lowest CPU Usage Peaks table

  This table shows the bottom 10 Windows servers with the lowest 90 percentile CPU busy values.

z/OS tab

Using high and low usage percentage is an effective method to determine which servers should be looked into. For consistency, the same report is included for the servers (LPARs) within the mainframe platform. Having a view of the high and low usage LPARs allows you to monitor changes in the behavior of your LPARs that can suggest further investigation is necessary.

- 10 Highest CPU Usage Peaks table

  This table shows the top 10 z/OS LPARs for System z with the highest 90 percentile CPU busy values.
- 10 Lowest CPU Usage Peaks table

  This table shows the bottom 10 z/OS LPARs for System z with the lowest 90 percentile CPU busy values.

Software Cost Analysis Summary workspace

Managing System z capacity can be costly and complex. This workspace allows you to better manage your software costs on System z by providing graphs and tables that show both the details and the summary of the software costs.

This workspace shows MSU utilization and pricing information for all the registered IBM products for CPCs within an enterprise. The workspace displays information for the observed, forecasted, and optimized scenarios. Lists of the registered products in three scenarios and IBM NO89 products are also provided.

The execution time of the Software Cost Analysis Summary workspace is directly proportional to the number of days, sub-capacity products, CPCs, and LPARs. Therefore, to minimize the waiting time for processing the workspace with a long date range over many products, CPC, and LPARs, you can schedule the source report of this workspace to run at regular intervals. The source report is named Software Cost Analysis Summary under the **Workspace** > **Objects** > **Charts & Tables** > **Software Cost Analysis** folder. When you do so, you can quickly view the most recently saved version of the contents without having to wait for the report to populate.

Source tables

[Presentation View].[Customer Pricing Information]

[Presentation View].[LPAR MSU Forecast Metrics]

[Presentation View].[LPAR MSU Metrics]

[Presentation View].[LPAR MSU Optimization Metrics]

[Presentation View].[Movement Suggestion by Optimization Metrics]

[Presentation View].[NO89 Products]

[Presentation View].[Product Billable MSU Forecast Metrics]

[Presentation View].[Product Billable MSU Optimization Metrics]

[Presentation View].[Product MSU Forecast Metrics]

[Presentation View].[Product MSU Optimization Metrics]

[Presentation View].[Product MSU Utilization Metrics Without Z/OS]

[Presentation View].[Product MSU Utilization Metrics]

[Presentation View].[Prompts/Product Billable MSU Prompt View]

[Presentation View].[Raw Product MSU Utilization Metrics]

[Presentation View].[Sub-Capacity Program Pricing Structures]

[Presentation View].[Parents of Sub-Capacity Reference-based programs]

[Presentation View].[Timezone Metrics]

[Presentation View].[Z/OS Product MSU Utilization Metrics]

Standard prompts

date range

time range

timezone

Show First SCRT Cycle

Show Monetary Value

enterprise

cpu serial no

The values for the **Enterprise** and **cpu serial no** prompts are retrieved from the HCM.CUSTOMER_PRICE table. You must ensure that the required values are loaded in the table. Refer to the correct Configuration section for loading.

Workspace objects

The widgets in the Software Cost Analysis Summary workspace are in the **Workspace** > **Objects** > **Charts & Tables** > **Software Cost Analysis** folder.

The tables available on the Software Cost Analysis Summary are grouped in the following tabs that are at the bottom of the workspace:

- **Observed**—shows the machine level summary and product summary information for the observed scenario.
- **Forecasted**—shows the summary information, for comparison purposes, of both the forecasted and observed scenarios.
- **Optimized**—shows the summary information, for comparison purposes, of both the optimized and forecasted scenarios.
- **Products Matrix**—shows a list of registered products in three scenarios (Observed, Forecasted and Optimized) and IBM NO89 products.

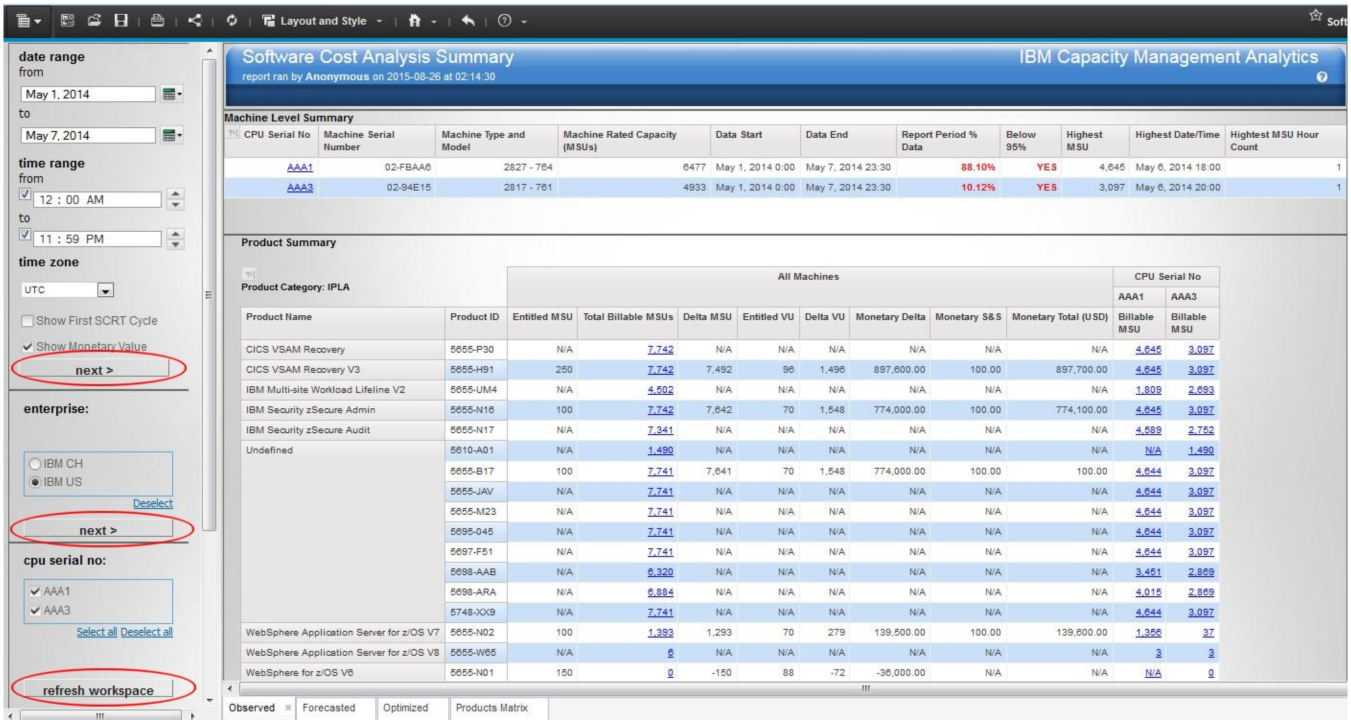

Figure 56: Software Cost Analysis Summary workspace Observed tab

Observed tab: Machine Level Summary table

This table shows the summary information for each machine that you select in the prompts.

If the machine model changes during the specified period for a machine, the table shows you the highest machine model observed in the specified period.

Data Start and Data End is the date and time range within the specified period and with observed data for the specific CPC. If the value of Report Period % Data is below 95%, the value displays in red for quick identification.

Observed tab: Product Summary table

This table shows the billable MSU and/or monetary value for each registered IBM product that is running on the selected enterprise.

When **Show Monetary Value** is not selected, the Product Summary table is split into two tables: one for all IPLA and MLC products except z/OS, and another for z/OS with zNALC consideration.

When **Show Monetary Value** is checked, the Product Summary table is split into three tables: one for IPLA, one for MLC except z/OS, and one for z/OS with zNALC consideration.

Important: When **Show Monetary Value** is checked, if the price of some products on some CPCs is not set in the HCM.CUSTOMER_PRICE table, then the billable MSUs of them is counted in the Total Billable MSUs column for All Machines, and not counted in the Monetary Total column for All Machines. The Monetary Total column for All Machines can calculate only the total of all CPCs with the price.

The Summary row is used to do the total aggregation for the Monetary Total column and the MSU columns.

If a product is running on some CPCs, but not on all the user-specified CPCs, then the related billable MSUs is shown as empty and the related monetary value appears as "N/A" or empty.

If all the values of a column are NULL, the column is not displayed in the Product Summary table.

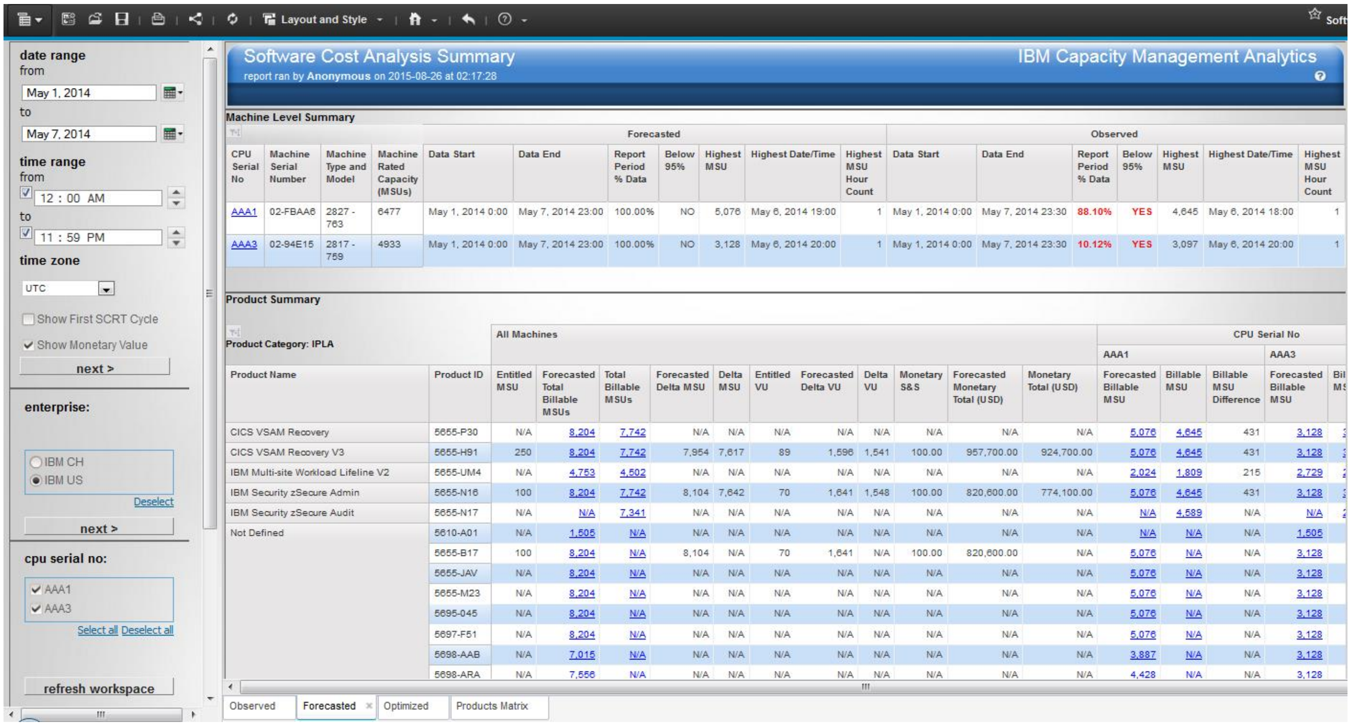

Figure 57: Software Cost Analysis Summary workspace Forecasted tab

Forecasted tab: Machine Level Summary table

This table shows the Forecasted and Observed summary information for each machine that you select in the prompts.

If the machine model changes during the specified period for a machine, the table shows you the highest machine model in forecast streams for the specified period.

If there is no data for the Observed or Forecasted scenario for all machines, only the **Below 95%** and **Highest MSU Hour Count** columns are displayed in the table, with "N/A" or 0 values.

Forecasted tab: Product Summary table

This table shows the product summary information for both Forecasted and Observed scenarios.

When **Show Monetary Value** is not selected, the Product Summary table is split into two tables: one for all IPLA and MLC products except z/OS, and another for z/OS with zNALC consideration.

When **Show Monetary Value** is checked, the Product Summary table is split into three tables: one for IPLA, one for MLC except z/OS, and one for z/OS with zNALC consideration.

Important: When **Show Monetary Value** is checked, if the price of some products on some CPCs is not set in the HCM.CUSTOMER_PRICE table, then the billable MSUs of them is counted in the Total Billable MSUs column for All Machines, and not counted in the Monetary Total column for All Machines. The Monetary Total column for All Machines can calculate only the total of all CPCs with the price.

The Summary row is used to do the total aggregation for the Monetary Total column and the MSU columns.

If a product is running on some CPCs, but not on all the user-specified CPCs, then the related billable MSUs is shown as empty and the related monetary value appears as "N/A" or empty.

If all the values of a column are NULL, the column is not displayed in the Product Summary table.

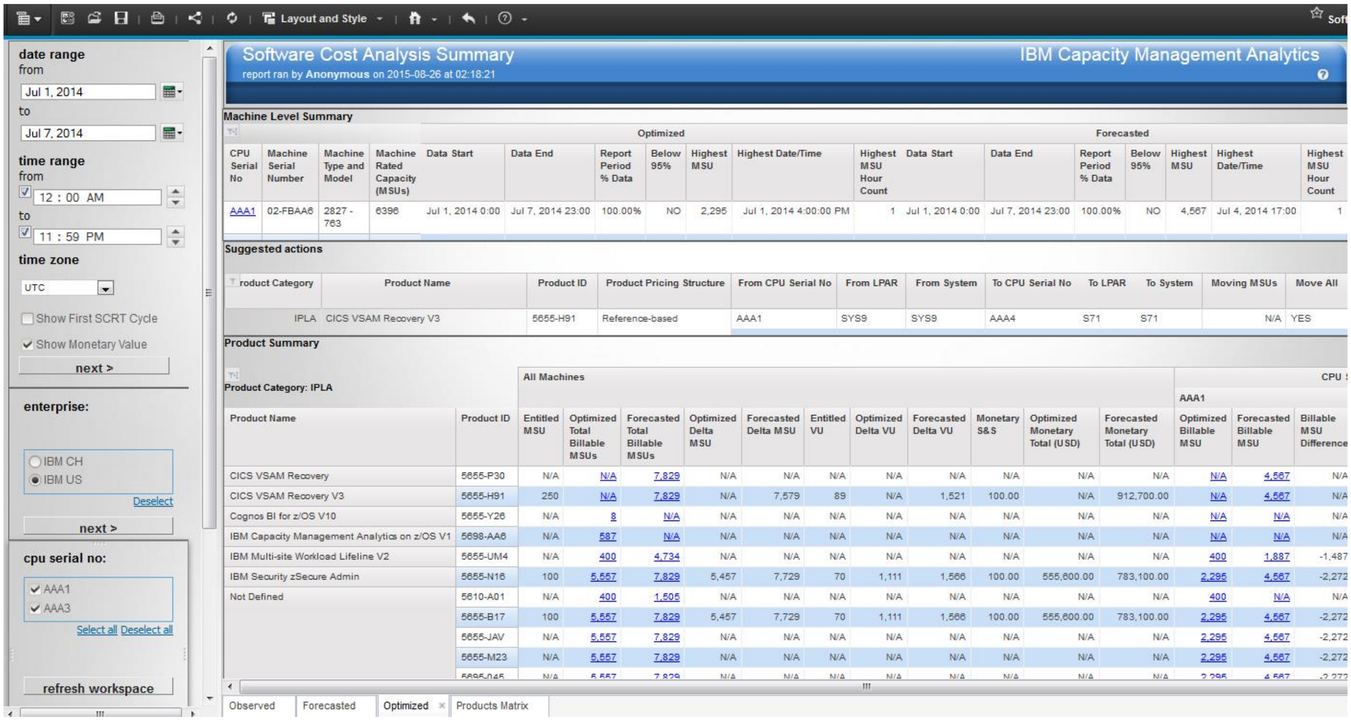

Figure 58: Software Cost Analysis Summary workspace Optimized tab

Optimized tab: Machine Level Summary table

This table shows the Optimized and Forecasted summary information for each machine that you select in the prompts.

If the machine model changes during the specified period for a machine, the table shows you the highest machine model in optimization streams for the specified period.

If there is no data for the Optimized or Forecasted scenario for all machines, only the **Below 95%** and **Highest MSU Hour Count** columns are displayed in the table, with "N/A" or 0 values.

Optimized tab: Product Summary table

This table shows the product summary information for both Optimized and Forecasted scenarios.

When **Show Monetary Value** is not selected, the Product Summary table is split into two tables: one for all IPLA and MLC products except z/OS, and another for z/OS with zNALC consideration.

When **Show Monetary Value** is checked, the Product Summary table is split into three tables: one for IPLA, one for MLC except z/OS, and one for z/OS with zNALC consideration.

The Summary row is used to do the total aggregation for the Monetary Total column and the MSU columns.

Important: When **Show Monetary Value** is checked, if the price of some products on some CPCs is not set in the HCM.CUSTOMER_PRICE table, then the billable MSUs of them is counted in the Total Billable MSUs column for All Machines, and not counted in the Monetary Total column for All Machines. The Monetary Total column for All Machines can calculate only the total of all CPCs with the price.

If a product is running on some CPCs, but not on all the user-specified CPCs, then the related billable MSUs is shown as empty and the related monetary value appears as "N/A" or empty.

If all the values of a column are NULL, the column is not displayed in the Product Summary table.

Optimized tab: Suggested Actions table

This table shows the suggested actions from the optimization stream results.

When **Show Monetary Value** is not selected, this table does not appear on the Optimized tab. Provide the overall suggested actions on all CPCs, so that the enterprise and cpu serial no prompts do not affect the result.

To provide the overall suggested actions for all CPCs, the result in the Suggested Actions table is not filtered by the enterprise and cpu serial number value prompts in the workspace global area.

Products Matrix tab

This tab shows the following tables:

- Registered Products Matrix for Observed Scenario
- Registered Products Matrix for Forecasted Scenario
- Registered Products Matrix for Optimized Scenario
- NO89 Products Matrix table

If the product is running on this LPAR, then the cross point of LPAR NAME and PRODUCT is marked by a check mark. For more information, see “SCA: Registered Products Matrix ” on page 265 and “SCA: NO89 Products Matrix ” on page 255.

zServer Monitoring Dashboard workspace

This workspace provides a 360-degree view of the zServer system environment. Through timely identification of capacity bottlenecks and potential performance issues, IT can take steps to eliminate deficiencies and optimize provisioning for critical business functions. The benefits of using this dashboard include growing revenue by meeting business goals and SLAs for constant availability and scalability. It also helps with cost savings and cost avoidance of unnecessary software licenses. There are also efficiency and productivity gains from the reduction of administrative efforts.

This workspace provides a quick visual status of a zServer, its LPARs, and disk subsystems. Key metrics are provided at the overall machine level. CPU and Storage metrics and I/O and Network metrics are displayed at the LPAR level.

Standard prompts

date range

time range

interval band

cpu serial no

lpar name

The charts and tables available on the zServer Monitoring Dashboard are grouped on tabs that are located at the bottom of the workspace:

- Machine Level

- CPU & Storage at LPAR Level
- I/O & Network at LPAR Level

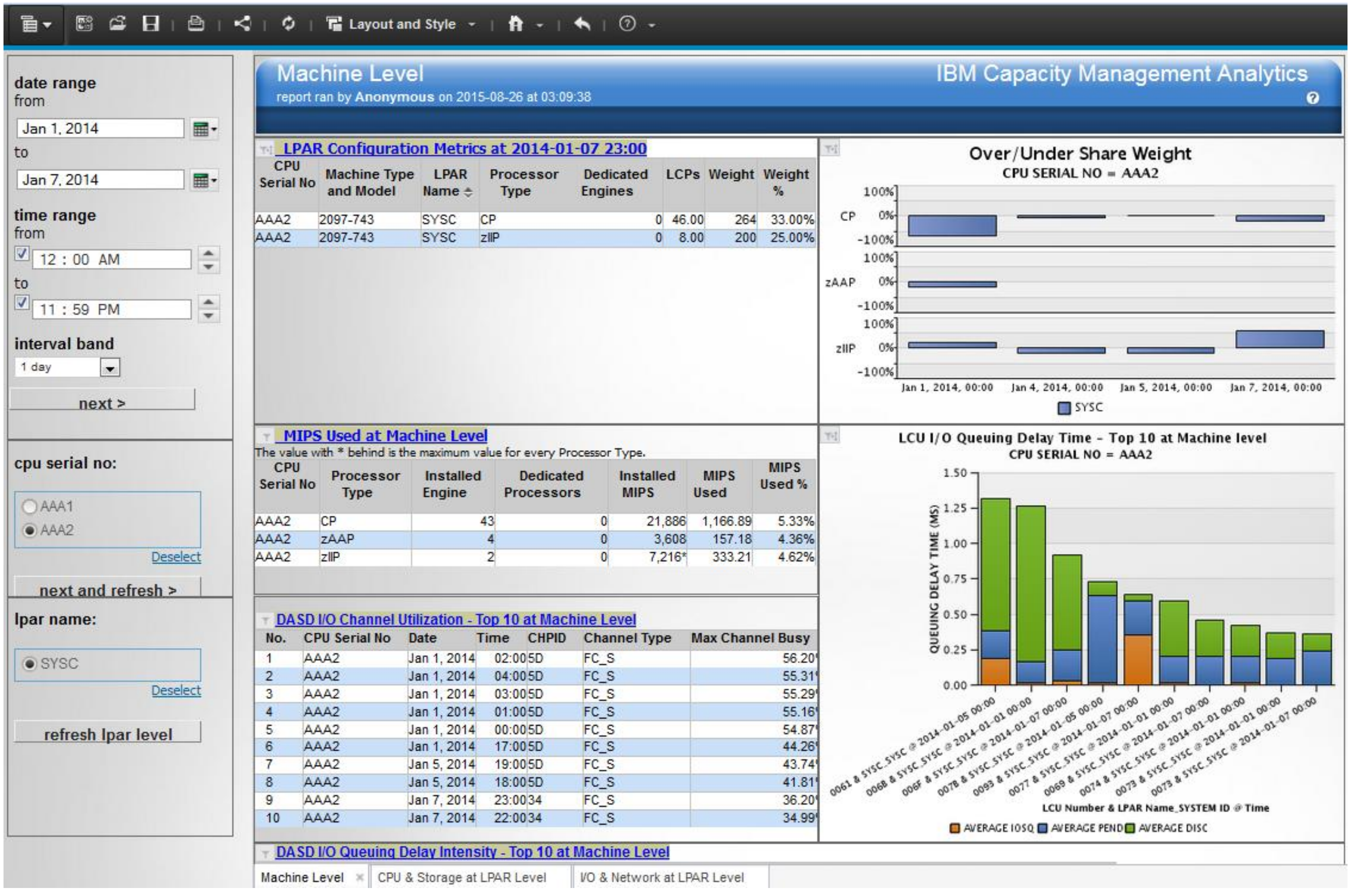

Figure 59: zServer Monitoring Dashboard Machine Level tab example

Machine Level tab: LPAR Configuration Metrics table

This table shows the latest LPAR configuration information of a zServer over the specified date and time range. Metrics include dedicated processor numbers, defined logical processor numbers, defined weight, and weight percentage over all LPARs for each processor type of each LPAR.

Source table
[Presentation View].[LPAR Level Metrics]

Machine Level tab: Over/Under Share Weight table

This table shows the usage of each LPAR against the total number of engines of each type in the shared pool of a CPC. It also shows the usage percentage over or under its guaranteed share of each LPAR due to the competition among LPARs.

Each processor type is shown separately.

Source table
[Presentation View].[LPAR Level Metrics]

Machine Level tab: MIPS Used at Machine Level table

This table shows summary level information for the zServer, including the number of installed engines, the number of dedicated engines, installed MIPS, MIPS used, and the percentage of installed MIPS used.

When there is value change during the selected date and time period, there is an asterisk marker (*) behind the value, and a footnote is displayed in top of the table that states "The value with * behind is maximum value for every Processor Type".

Source table
: [Presentation View].[LPAR Level Metrics]

Machine Level tab: DASD I/O Channel Utilization - Top 10 at Machine Level table

This table displays when and where the top 10 DASD I/O Channel Utilization occurs and helps you identify potential problem points in the I/O channel.

Although the default table configuration displays the top 10, you can modify the table to display any top number, as described in "Customizing top N display for the zServer Monitoring Dashboard" on page 97.

Source table
: [Presentation View].[Channel Path Activity Metrics]

Machine Level tab: LCU I/O Queuing Delay Time - Top 10 at Machine Level chart

This chart shows the top 10 longest delays by LCU, LPAR Name, and MVS System ID.

The amount of time that is required to generate this chart increases as the date range for the chart is increased. To minimize waiting time for chart generation, run live reports for weekly or fewer date ranges, and run scheduled reports for monthly or longer date ranges.

Although the default chart configuration displays the top 10, you can modify the table to display any top number, as described in "Customizing top N display for the zServer Monitoring Dashboard" on page 97.

Source tables
: [Presentation View].[Raw Device Activity Metrics]

  [Presentation View].[LPAR Level Metrics]

Machine Level tab: DASD I/O Queuing Delay Intensity - Top 10 at Machine Level table

This table identifies the top 10 largest queuing delay intensities by Volume Serial Number (VOLSER) and LPAR name.

Although the default table configuration displays the top 10, you can modify the table to display any top number, as described in "Customizing top N display for the zServer Monitoring Dashboard" on page 97.

Source table
: [Presentation View].[Device Activity Metrics]

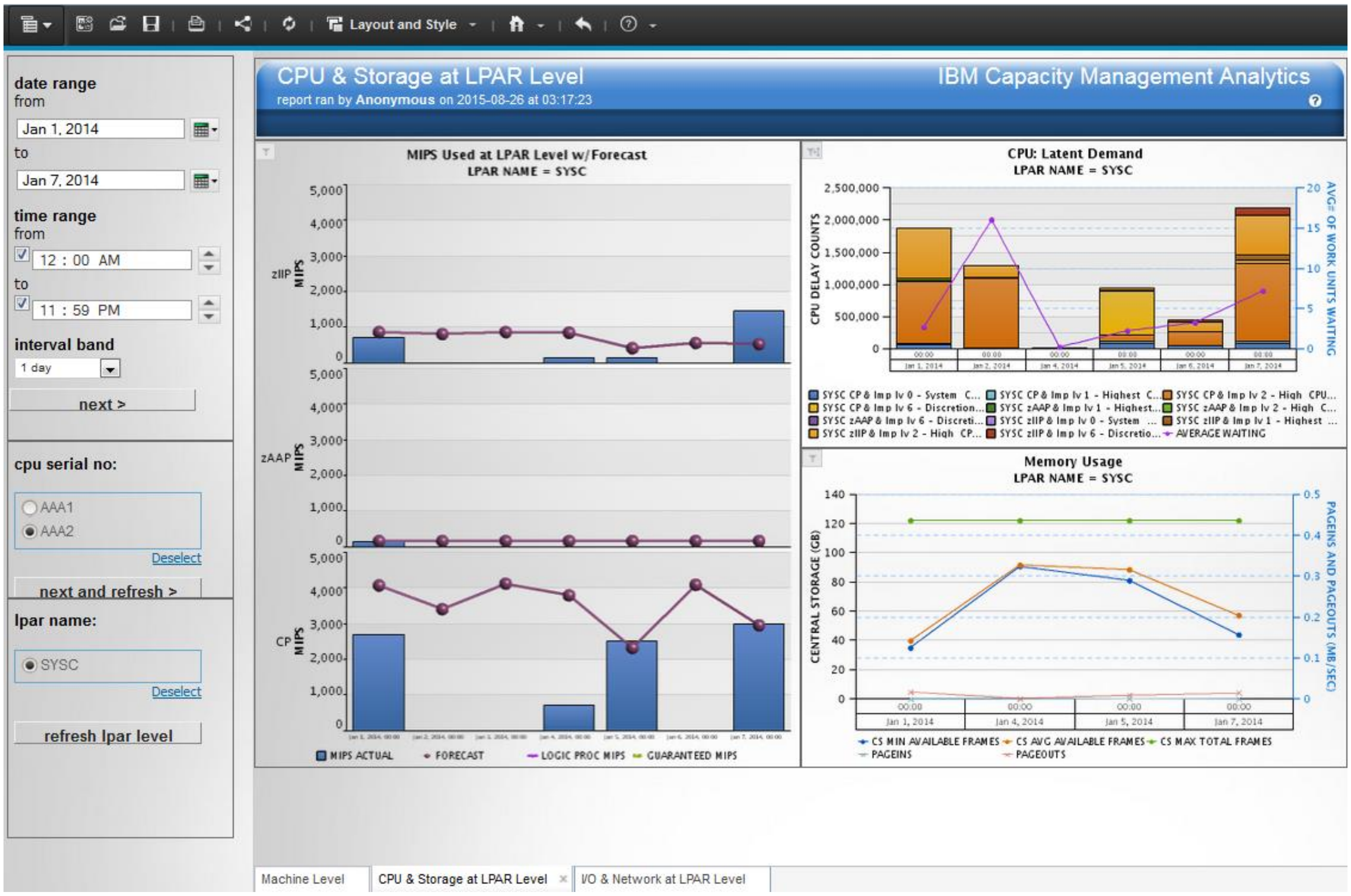

Figure 60: zServer Monitoring Dashboard CPU & Storage at LPAR Level tab example

CPU & Storage at LPAR Level tab: MIPS Used at LPAR Level chart

This chart shows the actual and forecasted CPU MIPS used at the LPAR level separated by processor type.

Source tables

[Presentation View].[LPAR Level Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

[Presentation View].LPAR Level Forecast Metrics]

Drill through to the following report:

CPU: MIPS Used - Service Class Period Level

CPU & Storage at LPAR Level tab: CPU - Latent Demand chart

This chart shows latent demand in terms of In & Ready Address Spaces or Work Units (for z/OS 1.12 or higher) in relation to the CPU Delay counts by WLM Importance Level.

Source tables

[Presentation View].[z/OS System Level Metrics]

[Presentation View].[LPAR Level Metrics]

[Presentation View].[WLM Service Class Period Usage/Delay Metrics]

Drill through to the following report:

"WLM: Delays by Importance Level" on page 266

CPU & Storage at LPAR Level tab: Memory Usage chart

This chart shows minimum, maximum, and average available central storage, along with page-in and page-out values.

Source tables

[Presentation View].[MVS Paging Activity Metrics]

[Presentation View].[LPAR Level Metrics]

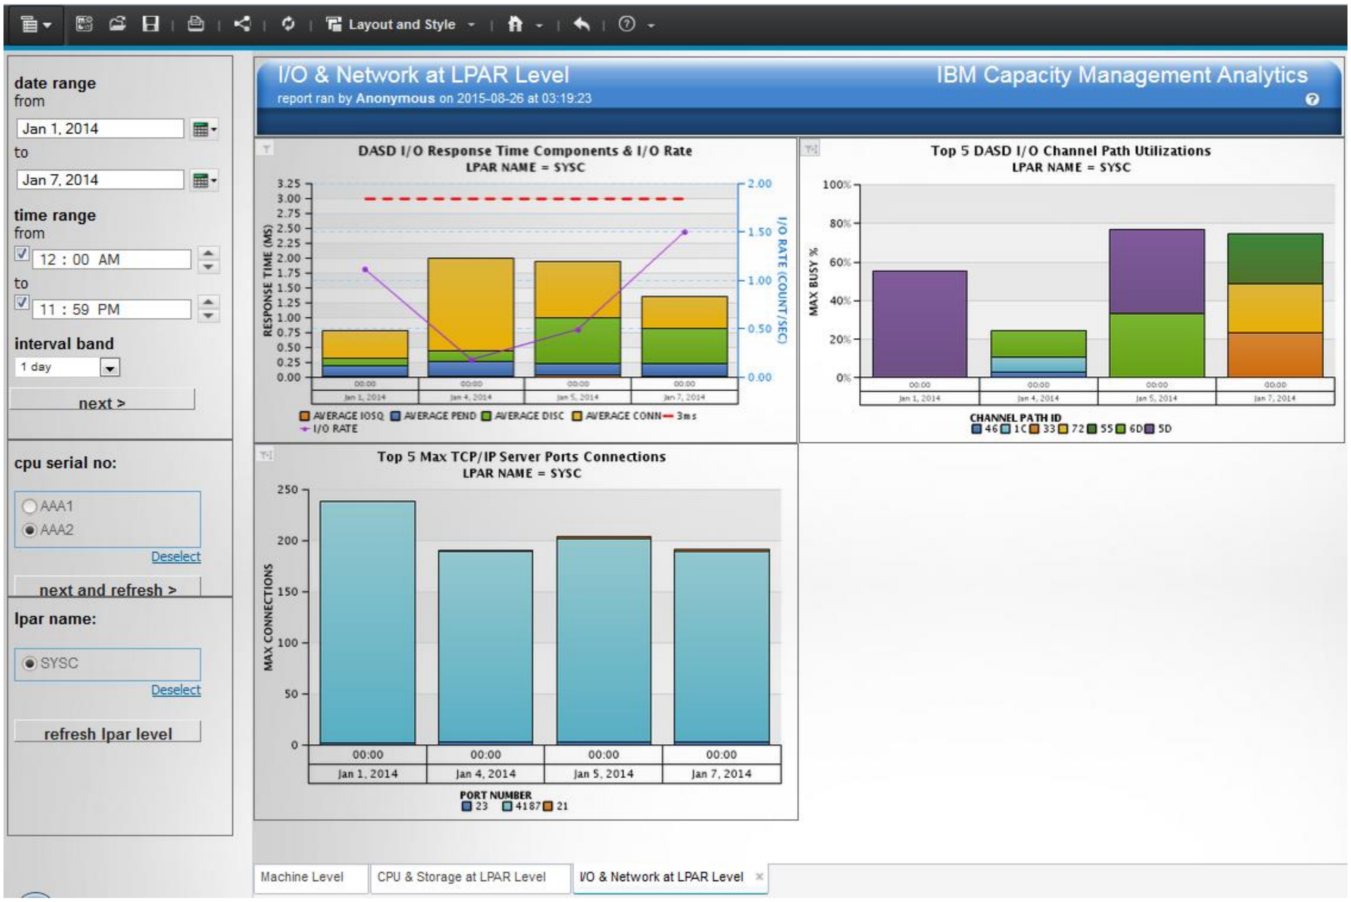

Figure 61: zServer Monitoring Dashboard I/O & Network at LPAR Level tab example

I/O & Network at LPAR Level tab: DASD I/O Response Time Components & I/O Rate chart

This chart shows DASD I/O Response Time components and I/O Rate by interval and LPAR, assuming LPAR names are unique across all zServers.

Source table

[Presentation View].[Device Activity Metrics]

I/O & Network at LPAR Level tab: Top 5 DASD I/O Channel Path Utilizations chart

This chart shows the top five DASD I/O Channel Path Utilizations, which occur at a System Management Facilities (SMF) time interval on a channel path.

When the interval band you select coincides with the SMF time interval, the chart displays five stacked bars. When the interval band is larger than the SMF time interval, there might be fewer than five stacked bars on the chart if the same channel path ID occurs in the top five at different time points.

Although the default chart configuration displays the top five, you can modify the table to display any top number, as described in "Customizing top N display for the zServer Monitoring Dashboard" on page 97.

Source table

[Presentation View].[Channel Path Activity Metrics]

I/O & Network at LPAR Level tab: Top 5 Max TCP/IP Server Ports Max Connections chart

This chart shows the top five maximum TCP/IP Server Port connections.

Although the default chart configuration displays the top five, you can modify the table to display any top number, as described in "Customizing top N display for the zServer Monitoring Dashboard" on page 97.

Source tables

[Presentation View].[TCP Server Port Connection Metrics]

[Presentation View].[LPAR Level Metrics]

Product Version

The Product Version report displays the IBM Capacity Management Analytics copyright, version, and build information.

The report is available in the **IBM Capacity Management Analytics Solution Kit** folder in IBM Cognos Connection.

The same information is available in other IBM Capacity Management Analytics reports by clicking the help icon (?), and clicking **About IBM CMA**.

Chapter 15. Uninstalling Capacity Management Analytics

Follow these procedures to uninstall IBM Capacity Management Analytics.

Uninstalling Capacity Management Analytics

Procedure

1. Drop the IBM Capacity Management Analytics database.

   Note: You must always drop the Capacity Management Analytics database first.

   a) Run `cmabase.py --dropdb [--run] [--verbose] [--trace=tracelevel]`

      Note: You can run `cmabase.py --dropdb` first to review the commands and run `cmabase.py --dropdb --run` to drop the Capacity Management Analytics database.

   b) Drop the sample database. Run `cmasampledb.py --dropdb [--run] [--verbose] [--trace=tracelevel]`.

      Note: You can run `cmasampledb.py --dropdb` first to review the commands and run `cmasampledb.py --dropdb --run` to drop the sample database.

2. Clean up CMAINSTANCE and the Capacity Management Analytics features. Clean up of the Capacity Management Analytics features removes all related files in your CMAINSTANCE directory, including logs and input files. Clean up of CMAINSTANCE deletes the entire CMAINSTANCE directory.

   a) Run `cmainstance.py --cleanup=<feature> [--verbose] [--trace=tracelevel]`

      --cleanup=feature
      Cleanup of the Capacity Management Analytics Instance. BASE (CMA base), AA (Application Analytics), PI (Anomaly Detection), CPF (Capacity Planning and Forecasting), SCA (Software Cost Analysis), SMO (Systems Management and Optimizations), SAMPLE (Sample Database), INSTANCE (Instance Directory) For instance, AA, SCA, PI.

      Note: You also can run the Capacity Management Analytics feature program to clean up individual features. For instance, `cmabase.py –cleanup`. Additionally, if you clean up any features by mistake, you can always create a new one by running `cmainstance.py --instance` or another feature program, for instance, `cmabase.py --instance`. For more information, see Chapter 6, “Manually installing IBM Capacity Management Analytics 2.1.1,” on page 59

3. Uninstall Capacity Management Analytics feature and the Capacity Management Analytics Solution Kit. You can uninstall one or more Capacity Management Analytics features or uninstall the entire Solution Kit.

   a) Run `cmahome.py –uninstall=<feature> [--verbose] [--trace=tracelevel]`

      --uninstall=feature
      Uninstall Capacity Management Analytics. BASE (CMA base), AA (Application Analytics), PI (Anomaly Detection), CPF (Capacity Planning and Forecasting), SCA (Software Cost Analysis), SMO (Systems Management and Optimizations), SAMPLE (Sample Database), SK (CMA Solution Kit) For instance, AA, SCA, PI.

      Note: You can also run the Capacity Management Analytics feature program to uninstall individual features For instance, `cmabase.py --uninstall`. Additionally, you can always reinstall Capacity Management Analytics features by running `cmahome.py --install --feature=<feature>`. If you want to reinstall the Capacity Management Analytics Solution Kit, you have to manually extract the tarball.

Uninstalling the Capacity Management Analytics Solution Kit content

You can delete the solution kit report content by deleting the **IBM Capacity Management Analytics Solution Kit** folder from the **Public folders** tab in **IBM Cognos Connection**, and then deleting the artifacts from **IBM Cognos Administration**.

Procedure

1. In **IBM Cognos Connection**, click **Public folders**.
2. Select the **IBM Capacity Management Analytics Solution Kit** folder, and click the **Delete** icon.
3. In **IBM Cognos Connection**, click **Launch** > **IBM Cognos Administration**.
4. Click the **Configuration** tab.
5. Select the **CMADW** entry, then click the **Delete** icon.
6. On the **Configuration** tab, click **Content Administration**.
7. Select the deployment entry and click the **Delete** icon.
8. Go to the `*Cognos_install_dir*/deployment` directory on the computer where IBM Cognos Business Intelligence is installed and delete `cma-2.1.zip`.

Uninstalling CPLEX components from the server

Once you have stopped all the services, you can proceed to uninstalling CPLEX components from the server.

About this task

To uninstall CPLEX, see the ILOG CPLEX Optimization Studio ILOG CPLEX Optimization Studio(http://www.ibm.com/support/knowledgecenter/SSSA5P)

Procedure

1. Log in as the Capacity Management Analytics administrator or an ID that has sudo or root.
2. Run the following commands:
   - `sudo /opt/ibm/ILOG/CPLEX_Studio1263/Uninstall/Uninstall`
   - `sudo rm -rf /opt/ibm/ILOG`

Uninstalling SPSS components from the server

Once you have stopped all the services, you can proceed to uninstalling SPSS components from the server.

About this task

To uninstall SPSS Modeler, see the SPSS Modeler documentation at SPSS Modeler V17.1.0 documentation(http://www.ibm.com/support/knowledgecenter/SS3RA7_17.1.0/clementine/knowledge_center/product_landing.html).

Procedure

1. Log in as the Capacity Management Analytics administrator or an ID that has sudo or root.
2. Run the following commands:
   - `sudo /opt/ibm/SPSS/ModelerBatch/17.1/Uninstall_IBM_SPSS_MODELER_BATCH/Uninstall_IBM_SPSS_MODELER_BATCH`

- `sudo /opt/ibm/SPSS/ModelerServer/17.1/Uninstall_IBM_SPSS_PREMIUM_SERVER/Uninstall_IBM_SPSS_Premium_SERVER`
- `sudo /opt/ibm/SPSS/ModelerServer/17.1/Uninstall_IBM_SPSS_MODELER_SERVER/Uninstall_IBM_SPSS_MODELER_SERVER`
- `sudo rm -rf /opt/ibm/SPSS`

Uninstalling Cognos Analytics from the server

Once you have stopped all the services, you can proceed to uninstalling Cognos Analytics components from the server.

About this task

To uninstall IBM Cognos Business Intelligence, see the IBM Cognos Analytics documentation at IBM Cognos Analytics 11.0 documentation(http://www.ibm.com/support/knowledgecenter/SSEP7J_11.0.0/com.ibm.swg.ba.cognos.cbi.doc/welcome.html).

Procedure

1. Log in as the Capacity Management Analytics administrator or an ID that has sudo or root.
2. Run the following commands:
   - `sudo /opt/ibm/cognos/analytics/Uninstall_IBM_Cognos_Analytics`
   - `sudo rm -rf /opt/ibm/cognos`

Uninstalling DB2 from the server

Once you have stopped all the services, you can proceed to uninstalling DB2 components from the server.

About this task

To uninstall IBM DB2, see the IBM DB2 documentation at IBM DB2 10.5 for Linux, Unix and Windows documentation(http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.kc.doc/welcome.html).

Procedure

Remove the licence and iso-swid directories.

Uninstalling the solution installer

If you used the solution installer to install IBM Capacity Management Analytics on Red Hat Enterprise Linux operating systems, you can uninstall the solution installer separately.

When you uninstall the solution installer, the processes that are used by the solution installer are stopped and removed from your computer. Also, all of the files in the solution installer directory are removed. This includes the Capacity Management Analytics server component installation files that you copied into the solution installer directory. If you want to preserve those files, you must copy them to another directory before you uninstall the solution installer.

Note: Uninstalling the solution installer does not uninstall the Capacity Management Analytics component applications. If you want to uninstall those components, you must use each component's uninstallation programs.

Procedure

1. Log in as the Capacity Management Analytics administrator or with an ID that has sudo or root.

   Note: Optional: go to the `SolutionInstaller/NodeRoot/Downloads/` software directory, and copy the *.tar.gz files outside of the SolutionInstaller directory.
2. Remove the Chef client.
   a) `sudo rpm -ev chef`
3. Remove NodeJS by running the following:
   a) `sudo - sh -c 'export PATH=${PATH}:/opt/ibm/nodejs1.2.0.1/bin; forever stop 0'`
   b) `sudo - sh -c 'export PATH=${PATH}:/opt/ibm/nodejs1.2.0.1/bin; npm uninstall express'`
   c) `sudo - sh -c 'export PATH=${PATH}:/opt/ibm/nodejs1.2.0.1/bin; npm uninstall async`
   d) `sudo - sh -c 'export PATH=${PATH}:/opt/ibm/nodejs1.2.0.1/bin; npm uninstall body-parser`
   e) `sudo - sh -c 'export PATH=${PATH}:/opt/ibm/nodejs1.2.0.1/bin; npm uninstall forver`
   f) `sudo rm -rf /opt/ibm/nodejs1.2.0.1/`
4. Remove the solution installer folder.
   a) Run `sudo rm -rf SolutionInstaller`

Appendix A. Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products.

Keyboard shortcuts

Standard Microsoft Windows navigation keys are used in addition to application-specific keys.

You can use keyboard shortcuts to navigate through the application and perform tasks. If you are using a screen reader, you might want to maximize your window so the keyboard shortcut table is completely expanded and accessible. You might want to turn high contrast on in your operating system so the lines in diagrams and charts in the application are more visible.

Note: The following keyboard shortcuts are based on U.S. standard keyboards.

Table 29: Keyboard shortcuts

Action	Shortcut keys
Enables or disables the accessibility feature - an IBM Cognos Report Studio restart is required. The default is **disabled.**	Ctrl+Alt+A
Opens the context menu for the selected item, if available.	Shift+F10
Switches focus from or to the main menu bar and the main worksheet.	Ctrl+F10
Opens the online help.	F1
Closes the Web browser window.	Alt+F4
Copies objects.	Ctrl+C
Pastes objects.	Ctrl+V
Closes objects, such as a dialog box, menu, drop-down list, or page.	Esc
Cycles through objects, such as the tabs in a dialog, the objects in a page, the cells in a list, the rows in a crosstab, or check boxes.	Tab, Shift+Tab, Arrow keys
Selects the first or last item in a list.	Home/Pg Up or End/Pg Dn keys
Selects or clicks one or multiple objects, such as a radio button, menu or toolbar item, or node in a tree widget. In the **Properties** pane, expands or collapses a property group.	Space or Enter, Ctrl+space or Shift+space, Ctrl +Enter or Shift+Enter

IBM and accessibility

See the IBM Accessibility Center (http://www.ibm.com/able) for more information about the commitment that IBM has to accessibility.

Appendix B. Optional: Manual configuration tasks

The tasks in this chapter are run automatically by the IBM Capacity Management Analytics Solution Kit installation. Use these tasks only if you want to perform any of the configuration tasks manually.

All of the tasks in this section are optional.

Changing the default schema names in Framework Manager models

The IBM Capacity Management Analytics Framework Manager models use default schema names for the data sources. If you use different schema names, you must update the references to the appropriate schema names in the Framework Manager models.

There are 2 models: `CMA.cpf` and `CMA DW.cpf`. Each model has references to 3 data sources that contain the schema names: DRL, DRLSYS, and HCM.

After you change the schema names in the models, you must republish the Framework Manager packages to IBM Cognos Business Intelligence.

For more information about publishing packages, see IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J_10.2.1/com.ibm.swg.ba.cognos.ug_fm.10.2.1.doc/t_publishapackage.html%23PublishaPackage)..

Procedure

1. Launch IBM Cognos Framework Manager.
2. From the **Welcome** page, click **Open a project**.
3. Navigate to the location where you installed Capacity Management Analytics.
4. In the `fm_model` folder, select `CMA.cpf`, and click **Open**.
   a) In the **Project Viewer** pane, expand **Data Sources**.
   b) Select **CMADW_DRL**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   c) Select **CMADW_DRLSYS**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   d) Select **CMADW_HCM**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   e) Click **File** > **Save**.
5. Click **File** > **Open**.
6. In the `fm_model\CMA DW` folder, select `CMA DW.cpf`, and click **Open**.
   a) In the **Project Viewer** pane, expand **Data Sources**.
   b) Select **CMADW_DRL**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   c) Select **CMADW_DRLSYS**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   d) Select **CMADW_HCM**, and in the **Properties** pane, enter the appropriate schema name in the **Schema** box.
   e) Click **File** > **Save**.
7. Republish the packages.
   a) Open the `CMA.cpf` model again.
   b) In the **Project Viewer** pane, expand **Packages**.
   c) Right-click **IBM Capacity Management Analytics**, and click **Publish Package**.

d) Right-click **IBM Capacity Management Analytics Development**, and click **Publish Package**.

Use IBM DB2 Analytics Accelerator (IDAA) with IBM Capacity Management Analytics

To use IBM DB2 Analytics Accelerator (IDAA) with IBM Capacity Management Analytics, you can either configure ZPARMs in your DB2 configuration or enable IDAA in IBM Cognos Business Intelligence and IBM SPSS Modeler.

- For information about configuring ZPARMs, see "Setting ZPARMs for IDAA in DB2 for z/OS" on page 294.
- For information about configuring Cognos and SPSS Modeler, see "Configuring IBM Capacity Management Analytics to use IBM DB2 Analytics Accelerator" on page 295.

Before you enable IDAA

To use IDAA with IBM Capacity Management Analytics, you must ensure that all database object's (the database, tables, tablespaces, UDFs, and stored procedures for both DRL and HCM databases) CCSID uses EBCDIC encoding.

For example, the following function is encoded in EBCDIC by using the PARAMETER CCSID option:

```
CREATE FUNCTION HCM.YEAR_BAND(d date)
RETURNS DATE
LANGUAGE SQL
PARAMETER CCSID EBCDIC
DETERMINISTIC
READS SQL DATA
NO EXTERNAL ACTION
RETURN
(
date ( varchar(year(d)) || '-01-01')
);
```

The Tivoli Decision Support for z/OS DRL database uses EBCDIC encoding. To use IDAA, the HCM database must also use EBCDIC encoding. If you use Unicode for the HCM database, you must create a separate EBCDIC encoded database with a different schema and use the PARAMETER CCSID option.

Setting ZPARMs for IDAA in DB2 for z/OS

You can enable IBM DB2 Analytics Accelerator (IDAA) acceleration for IBM Capacity Management Analytics by setting the appropriate ZPARMS in your DB2 configuration.

For more information, see the following topics in IBM Knowledge Center:

- Setting ZPARMs for IBM DB2 Analytics Accelerator in DB2 11 for z/OS (www.ibm.com/support/knowledgecenter/SS4LQ8_4.1.0/com.ibm.datatools.aqt.doc/installmanual/task/t_idaa_inst_zparms_db2_11.html)
- Enabling DB2 to work with IBM DB2 Analytics Accelerator for z/OS (www.ibm.com/support/knowledgecenter/SSEPEK_11.0.0/com.ibm.db2z11.doc.perf/src/tpc/db2z_enablingaccelerators.dita)

If you set QUERY_ACCELERATION to ALL, then all queries use IDAA and you do not have to do any additional configuration.

If you are using another setting, some queries might not use IDAA. In this case, you can enable IDAA in IBM Cognos Business Intelligence and IBM SPSS Modeler. For more information, see "Configuring IBM Capacity Management Analytics to use IBM DB2 Analytics Accelerator" on page 295.

Configuring IBM Capacity Management Analytics to use IBM DB2 Analytics Accelerator

To use IBM DB2 Analytics Accelerator (IDAA) with IBM Capacity Management Analytics, you can add a DB2 open session command block to the data source connection in IBM Cognos Business Intelligence and enable IDAA in SPSS Modeler.

Alternatively, you can use ZPARMs for IDAA. For more information, see "Setting ZPARMs for IDAA in DB2 for z/OS" on page 294.

Procedure

1. Enable IDAA in IBM Cognos Business Intelligence:
   a) In **IBM Cognos Connection**, click **Launch** > **IBM Cognos Administration**.
   b) On the **Configuration** tab, select **Data Source Connections**.
   c) Click the data source name.
   d) Click **Set properties** .
   e) On the **Connection** tab, click **Set** for **Open session commands**.
   f) Enter the following text in the **XML database commands** box:

```
<commandBlock>
  <commands>
    <sqlCommand>
      <sql>SET CURRENT QUERY ACCELERATION ALL</sql>
    </sqlCommand>
  </commands>
</Block>
```

   It is recommended that you set the value to ALL, but you can set it to the value that is more appropriate for your environment. For more information, see the "Setting the query acceleration register from IBM Cognos BI" section in Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS

   If this data source is shared among other Cognos applications, all queries would use IDAA. If data from those shared applications is not also loaded into IDAA, then those reports may not complete. If loading the data for the shared applications into IDAA is not appropriate for your environment, consider setting up a second data source.

   g) Click **OK**, and then click **OK** again.
2. Enable IDAA in IBM SPSS Modeler:
   a) Catalog the IDAA data source.
   b) Edit the `odbc.ini` file to connect to the IDAA data source. Ensure that the data source name is the same as the database alias name that you used when you cataloged the database. For more information, see "Configuring IBM SPSS Modeler Server with DB2 CLI/ODBC on Linux on System z operating systems" on page 296.
   c) For more information, see Enabling Integration with IBM DB2 Analytics Accelerator for z/OS in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/clementine/dbmining_zdb2_enabling.html).

   Ensure that you modify the `odbc-db2-accelerator-names.cfg` file in the Enabling the integration of IBM® DB2 for z/OS® in IBM SPSS Modeler task (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/clementine/dbmining_zdb2_enabling_helperapps.html).

   d) Restart SPSS Modeler.

Verify that IDAA is being used in IBM Capacity Management Analytics

You can verify that the query or stream was run with IDAA by using IBM Data Studio or IBM z®/OS System Display and Search Facility (SDSF) on the LPAR where the database resides.

Verify by using IBM Data Studio

You can verify that a query or stream was run with IDAA from IBM Data Studio.

From the **Administration Explorer**, expand the data source, and select **Accelerators**, and then click **Manage**. The executed queries or streams are displayed in the **Query Monitoring** pane.

Verify by using SDSF

If you use SDSF, you can run the following command to verify whether a Cognos query or an SPSS Modeler stream was run with IDAA: `-DatabaseName DISPLAY THD(*) ACCEL(*)`

For a Cognos query, the expected result would be something like:

```
DSNV401I -DatabaseName DISPLAY THREAD REPORT FOLLOWS -
DSNV402I -DatabaseName ACTIVE THREADS - 697
NAME ST A REQ ID AUTHID PLAN ASID TOKEN
SERVER AC * 23 db2jcc_appli COGADM DISTSERV 0071 4185
V437-WORKSTATION=name
USERID=cogadm
APPLICATION NAME=db2jcc_application
V442-CRTKN=::IP_address.36684.CFD341BE1A54
V445-G91E8B0B.OF4C.CFD341BE1A54=4185 ACCESSING DATA FOR ::IP_address
V444-G91E8B0B.OF4C.CFD341BE1A54=4185 ACCESSING DATA AT
 ZGRYPHO2-::IP_address..1400
DISPLAY ACTIVE REPORT COMPLETE
```

For an SPSS stream, the expected result would be something like:

```
-DatabaseName DISPLAY THD(*) ACCEL(*)
DSNV401I  -DatabaseName DISPLAY THREAD REPORT FOLLOWS -
DSNV402I  -DatabaseName ACTIVE THREADS - 606
NAME     ST A   REQ ID           AUTHID   PLAN     ASID TOKEN
SERVER   RA *   230 modelerserve COGADM   DISTSERV 0071 10446
V437-WORKSTATION=ADMINIB-7ILJPF1
USERID=cogadm
APPLICATION NAME=modelerserver.exe
V442-CRTKN=IP_address.61028.151111053635
V445-G9778346.EE64.151111053635=10446 ACCESSING DATA FOR ::IP_address

V444-G9778346.EE64.151111053635=10446 ACCESSING DATA AT
 ZGRYPHO2-::IP_address..1400
DISPLAY ACTIVE REPORT COMPLETE
DSN9022I  -DatabaseName DSNVDT '-DISPLAY THD' NORMAL COMPLETION
```

Configuring IBM SPSS Modeler Server with DB2 CLI/ODBC on Linux on System z operating systems

Complete this procedure to configure IBM SPSS Modeler Server with DB2 CLI/ODBC on Linux on System z operating systems.

Before you begin

Ensure that the following prerequisites are met before proceeding.

- Modeler Server is installed on Linux on System z

- SDAP 7.1 is installed on Linux on System z
- DB2 client is installed on Linux on System z and the db2inst user exists

Procedure

1. Edit `odbc.ini` in the SDAP installation directory to include the following lines:

```
[ODBC Data Sources]
<DataSourceName>=IBM DB2 CLI ODBC

[<DataSourceName>]
Driver=/home/db2inst1/sqllib/lib/libdb2o.so
Description=IBM DB2 CLI/ODBC Driver
DriverUnicodeType=1
```

   For example:

```
[ODBC Data Sources]
STLAB71=IBM DB2 CLI ODBC

[STLAB71]
Driver=/home/db2inst1/sqllib/lib/libdb2o.so
Description=IBM DB2 CLI/ODBC Driver
DriverUnicodeType=1
```

   Note: Ensure that the data source name (*DataSourceName* in the example) does not contain spaces.

2. Edit `odbc.sh` in the SDAP installation directory so that the *LD_LIBRARY_PATH* variable appears as follows:

```
if [ "$LD_LIBRARY_PATH" = "" ]; then
LD_LIBRARY_PATH=/opt/IBM/SPSS/SDAP71/lib:/home/db2inst1/sqllib/lib
else
LD_LIBRARY_PATH=/opt/IBM/SPSS/SDAP71/lib:$LD_LIBRARY_PATH:/home/db2inst1/
sqllib/lib
```

3. In the `db2inst/sqllib/cfg` directory, create or edit `db2cli.ini` to include following lines.

```
[<data source name>]
Hostname=<hostname>
Port=8040
Protocol=TCPIP
DriverUnicodeType=1
Database=DB1R
```

   For example:

```
[STLAB71]
Hostname=myserver.mydomain.com
Port=8040
Protocol=TCPIP
DriverUnicodeType=1
Database=DB1R
```

   Note: The data source name should be the same as the name specified in the SDAP `odbc.ini` file.

4. Add the following line to `modelersrv.sh` in the SPSS Modeler Server installation directory:

```
. /opt/IBM/SPSS/SDAP71/odbc.sh
```

5. In the SPSS Modeler Server <install_dir>/bin directory, link libspssodbc.so with libspssodbc_datadirect_utf16.so.

```
rm -fr libspssodbc.so
ln -s libspssodbc_datadirect_utf16.so libspssodbc.so
```

6. Restart SPSS Modeler Server.

Import Capacity Management Analytics content to IBM Cognos Business Intelligence

View IBM Capacity Management Analytics reports in the IBM Cognos Business Intelligence portal. The reports are imported to the portal from a deployment zip file that is included in the Capacity Management Analytics installation.

Ensure that images for the reports are displayed in IBM Cognos BI

Some configuration is required to ensure that the IBM Capacity Management Analytics images appear in the IBM Cognos Business Intelligence reports.

Enabling the reports images for WebSphere Application Server

If you use IBM Cognos BI for z/OS or Linux on System z, enable the IBM Capacity Management Analytics report images by deploying the cma_images.ear file to the IBM WebSphere Application Server.

Before you begin

You must copy the cma_images.ear file to a location where you can deploy the images to WebSphere Application Server. The cma_images.ear file is installed by the Capacity Management Analytics installer, and is in the C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0\Images directory on the computer where you installed Capacity Management Analytics.

If you are upgrading, you must uninstall the existing CMA_Images application, and then complete the following steps.

Procedure

1. In a web browser, enter the URL for WebSphere Application Server Integrated Solutions Console.

   For example, http://*host_name*:*port*/ibm/console. Where *host_name* is the name of the computer where WebSphere Application Server is installed and *port* is the WebSphere Application Server HTTP transport port number.
2. On the **Welcome** page, click **Application**, **New Application**.
3. Click **New Enterprise Application**,
4. Click **Local file system**, then browse to C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0\Images\cma_images.ear.
5. Click **Next**.
6. Click **Fast Path**, then click **Next**.
7. Enter CMA_Images in the **Application name box**, then click **Next**.
8. Click **Next** on the **Map modules to servers** screen.
9. Click **Next** on the **Metadata for modules** screen.
10. Review the **Summary** screen, then click **Finish**.
11. Click **Save** when notified that the **CMA_Images** application was successfully installed.
12. Click **Applications**, **Application Types**, **WebSphere.enterprise applications**.
13. Select the **CMA_Images** application, then click **Start**.

Deploying the report images to WebSphere Liberty or Apache Tomcat application server

You can manually deploy the IBM Capacity Management Analytics report images to the WebSphere Liberty Profile or Apache Tomcat application server instance where IBM Cognos Business Intelligence is running.

Procedure

1. Go to the *CMA_ROOT*`\Images` directory.
   For example, on a Microsoft Windows operating system, go to the `C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0\Images` directory.
2. Do one of the following steps:
   - If you are using WebSphere Liberty, copy the `cma.war` file to the *c10_location*`/wlpdropins` directory.
   - If you are using Apache Tomcat, copy the `cma.war` file to the *c10_location*`/webapps` directory.
3. Restart the WebSphere Liberty Profile or Apache Tomcat server instance.

Enabling the report images for IBM HTTP Server or Apache server for Linux on System z

On IBM HTTP Server or Apache server, you must add directives to your server configuration file to display the images for the IBM Capacity Management Analytics reports.

Procedure

1. Copy the report images to a location on your web server computer.

   The report images are installed by the Capacity Management Analytics installer, and are in the `C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0\Images` directory on the computer where you installed Capacity Management Analytics.

   Copy the images to `/opt/IBM/Capacity_Management_Analytics_2.1.0/Images` on your Linux on System z operating system.
2. In the *webserver_location*`/conf` directory, open the `httpd.conf` file in a text editor.
3. Create an alias that is named cma that points to the directory where you have the Capacity Management Analytics report images.
   For example, add or modify the following information:

```
Alias /cma "/opt/IBM/Capacity_Management_Analytics_2.1.0"
<Directory "/opt/IBM/Capacity_Management_Analytics_2.1.0">
Options Indexes FollowSymLinks MultiViews
AllowOverride None
Order allow,deny
Allow from all
</Directory>
```

4. Save the file.
5. Restart your web server.

Results

The Capacity Management Analytics images are displayed when the user runs a report.

Enabling the report images for Report Studio

To view the IBM Capacity Management Analytics images when you author reports in Report Studio, you must enable WebDAV in IBM HTTP Server or Apache server.

Procedure

1. Copy the report images to a location on your web server computer.

The report images are installed by the Capacity Management Analytics installer, and are in the `C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0\Images` directory on the computer where you installed Capacity Management Analytics.

Copy the images to `/opt/IBM/Capacity_Management_Analytics_2.1.0/Images` on your Linux on System z operating system.

2. In the *webserver_location*`/conf` directory, open the `httpd.conf` file in a text editor.
3. Uncomment the directives that load the `modules/mod_dav.so` and `modules/mod_dav_fs.so` components.

```
LoadModule dav_module modules/mod_dav.so
LoadModule dav_fs_module modules/mod_dav_fs.so
```

4. Provide a location for the `DAVLockDB` directive.

   For example,

   `DAVLockDB "`*webserver_location*`/var/DavLock"`

   Ensure that the directory exists.
5. Set a time-out for WebDAV.

   For example,

   `DAVMinTimeout 600`
6. Create an alias for the directory where your images are stored.
   For example, create an alias that is named `cma` that points to the directory where you have the Capacity Management Analytics report images. On Linux on System z, that directory might be `/opt/IBM/Capacity_Management_Analytics_2.1.0/Images`.
7. Add or edit the `Dav On` information for the alias.

   For example, add or modify the following information:

```
Alias /cma "/opt/IBM/Capacity_Management_Analytics_2.1.0"
<Directory "/opt/IBM/Capacity_Management_Analytics_2.1.0">
Dav On
Options Indexes FollowSymLinks MultiViews
AllowOverride None
Order allow,deny
Allow from all
</Directory>
```

8. Save the file.
9. Restart your web server.

Results

With WebDAV enabled, the Capacity Management Analytics images are displayed in Report Studio.

Importing the IBM Cognos report deployment file

Import the IBM Capacity Management Analytics report deployment file into IBM Cognos Business Intelligence by using IBM Cognos Administration.

Before you begin

If you ran the installer on a Microsoft Windows operating system, ensure that you upload the contents of the `C:\Program Files (x86)\IBM\Capacity_Management_Analytics_2.1.0` directory on the computer where you installed Capacity Management Analytics to the computer where the IBM Cognos Business Intelligence is installed.

Important: If you already have IBM Capacity Management Analytics installed and you have customized reports or created your own reports, you must back up your reports to another location before you delete the existing Solution Kit content.

For more information about deleting the solution kit content, see "Uninstalling the Capacity Management Analytics Solution Kit content" on page 288.

Procedure

1. On the computer where the IBM Cognos Content Manager component is installed, copy the `cma-2.1.zip` file from the Capacity Management Analytics `reports` directory to the *`c10_location`*`\deployment` directory.

   Note: If you are not using the default deployment directory, ensure that you copy the `cma-2.1.zip` file to the appropriate directory.
2. Open a web browser.
3. Enter the URL for the IBM Cognos Business Intelligence portal, and press Enter.

   For example, enter

   `http://`*`host_name`*`/ibmcognos/`

   Or, if you are using IBM Cognos BI on z/OS, enter

   `http://`*`host_name`*`:`*`port`*`/p2pd/servlet/dispatch/ext`

   Where *`host_name`* is the name of the computer where IBM Cognos BI is installed. For IBM Cognos BI on z/OS, *`port`* is the WebSphere Application Server HTTP transport port number.
4. On the **Welcome** page, click **Administer IBM Cognos** content. Or click **Launch** > **IBM Cognos Administration** on IBM Cognos content.
5. On the **Configuration** tab, click **Content Administration**.
6. On the toolbar, click the **New Import** icon .
7. On the **Select a deployment archive - New Import** wizard, click **Next**.
8. In the **Deployment Archive** pane, select **cma-2.1 file**, then click **Next**.
9. On the **Specify a name and description - New Import** wizard, enter a **Name**, then click **Next**.
10. In the **Public folders, directory, and library content** section, select **IBM Capacity Management Analytics Solution Kit**, and click **Next**.
11. On the **Specify the general options** pane, select whether to include access permissions and references to external namespaces, and who owns the entries after they are imported, and click **Next**.

    The summary information is displayed.
12. Review the summary information, and click **Next**.
13. Click **Save and run once**, then click **Finish**.
14. On the **Run with options-deployment** page, specify the time and date for the run, or select **NOW** if you want to run immediately.
15. Click **Upgrade all report specifications to the latest version**.
16. Click **Do not assign new IDs during import**, then click **Run**.

    Important: To ensure that the workspaces work correctly, you must select **Do not assign new IDs during import** when you import the content.
17. Review the run time, click **View the details of this import after closing this dialog**, then click **OK**.
18. In the **View run history details - deployment** pane, click **Refresh** and check import messages in the **View an import deployment record - deployment** pane. A message such as `CM-REQ-2300 Import is complete` indicates that the import action is successful.
19. Click **Close**.

The deployed content is available in **Public Folders**.

20. Click **Launch** > **IBM Cognos Connection**. The **Launch** icon is in the upper right of the outer frame.

Creating data source connections to the Capacity Management Analytics data store

You must create a data source connection to the Capacity Management Analytics data store to be able to run the IBM Capacity Management Analytics reports.

About this task

If you are performing a new installation of Capacity Management Analytics 2.1 or you are migrating from Capacity Management Analytics 1.2 Fix Pack 1 and the data configuration differs from the base installation (for example, a different schema name), then you must complete this procedure.

If you are migrating from Capacity Management Analytics 1.2 Fix Pack 1 and the data configuration **does not differ** from the base installation, then you can skip this procedure.

Procedure

1. Open a web browser.
2. Enter the URL for the IBM Cognos Business Intelligence portal, and press Enter.

   For example, enter

   `http://`*`host_name`*`/ibmcognos/`

   Or, if you are using IBM Cognos BI on z/OS, enter

   `http://`*`host_name`*`:`*`port`*`/p2pd/servlet/dispatch/ext`

   Where *`host_name`* is the name of the computer where IBM Cognos BI is installed. For IBM Cognos BI on z/OS, *`port`* is the WebSphere Application Server HTTP transport port number.
3. On the **Welcome** page, click **Administer IBM Cognos content**. Or click **Launch** > **IBM Cognos Administration**.
4. On the **Configuration** tab, and click **Data Source Connections**.
5. On the toolbar, click **New Data Source**.
6. In the **Name** box, type `CMADW`, and click **Next**.
7. On the **Specify the connection** page, do the following steps:
   a) In **Type**, select **IBM DB2**.
   b) In **Isolation level**, select **Specify a value**, and choose **Read uncommitted**.
   c) Select **Configure JDBC connection**.
   d) Click **Next**.
8. In the **DB2 database name** field, enter the name of your Capacity Management Analytics data store.
9. Under **Signons**, select both the **Password** and **Create a signon that the Everyone group can use** check boxes, type the user ID and password for the data store, and then click **Finish**.
   For example, enter `hcmadm` for the user ID and `hcmadm` for the password.

   Tip: To test whether the parameters are correct, click **Test the connection**.
10. On the **Specify the IBM DB2 (JDBC) connection string** panel, enter the name of the computer where your Capacity Management Analytics data store database is in the **Server name** box.
11. In the **Port number** box, enter the port number for the database server.
12. In the **Database name** box, enter the name of the database.

    Tip: To test whether the parameters are correct, click **Test the connection**.
13. Click **Finish**.

Calibrating streams and models for new hardware generations

If you are using IBM Capacity Management Analytics with new generations of hardware, you must ensure that the Capacity Management Analytics assets (the streams and models) are calibrated correctly.

To calibrate the assets, you must generate new models by rerunning all of the IBM SPSS predictive streams, such as the forecasting and anomaly detection streams.

Special consideration must be taken if you are using Capacity Management Analytics on IBM z13 due to the multithreading capabilities within specialty processors.

- You must rerun the following MIPS forecasting streams:
  - `mvspm-lpar-hourly-forecast-timeseries.str`
  - `mvspm-lpar-daily-forecast-timeseries.str`
  - `mvspm-lpar-monthly-forecast-timeseries.str`
- You must rerun the following anomaly forecasting streams:
  - `anomaly_detect_building_c.str`
  - `anomaly_detect_building_r.str`
  - `anomaly_classification.str`
  - `anomaly_detect_scoring.str`

Appendix C. Stop and start solution software services

IBM Capacity Management Analytics is an integrated solution that includes many products. If you must stop the services, you must do so in the correct order. The product services must also be started in the correct order.

Stop solution services

You must stop the IBM Capacity Management Analytics node services in a specific order.

Stop the node services in the following order:

1. Cognos Analytics
2. SPSS Modeler Server
3. DB2 Enterprise Server

Stopping services on IBM Cognos Analytics

To stop solution services, you must begin with stopping services on IBM Cognos Analytics.

About this task

You must stop the IBM Cognos Business Intelligence services and IBM HTTP Server on the IBM Capacity Management Analytics node computer.

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. Go to `/opt/ibm/cognos/analytics/bin64` and run `./cogconfig.sh -stop`.

Stopping services on the SPSS Modeler Server

You must stop the IBM SPSS services on the IBM Capacity Management Analytics node computer.

About this task

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. Go to the IBM SPSS Modeler Server directory. For instance, go to `/opt/ibm/SPSS/ModelerServer/17.1`. If you manually installed and accepted the defaults, the path is `/usr/ibm/SPSS/ModelerServer/17.1`.
3. Enter the following command: `./modelersrv.sh stop`.
4. To verify whether any services are still running, enter the following command: `./modelersrv.sh list`
5. To stop any services that are still running, enter the following command: `kill -9 'cat modelersrv.pid'`.

Stopping services on the IBM DB2 Enterprise Server

You must stop the IBM DB2 instance on the IBM Capacity Management Analytics node computer.

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. In a terminal window, type the following command to change the DB2 instance owner: `su - db2inst1`
3. Enter the following command to stop the DB2 administration server: `db2stop`.

Start solution services

You must start the IBM Capacity Management Analytics node services in a specific order.

Start the node services in the following order:

1. DB2 Enterprise Server
2. Cognos Analytics
3. SPSS Modeler Server

Starting services on the IBM DB2 Enterprise Server

You must start the IBM DB2 instance on the IBM Capacity Management Analytics node computer.

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. In a terminal window, type the following command to change the DB2 instance owner: `su - db2inst1`.
3. Enter the following command to start the DB2 administration server: `db2start`.

Starting services on IBM Cognos Analytics

You must start the IBM Cognos Analytics services on the IBM Capacity Management Analytics node computer.

About this task

On Linux operating systems, or if you installed IBM Cognos Business Intelligence to WebSphere Application Server on Microsoft Windows operating systems, you start the IBM Cognos Business Intelligence services by starting the WebSphere Application Server profile where IBM Cognos Business Intelligence is running.

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. Go to `/opt/ibm/cognos/analytics/bin64` and run `./cogconfig.sh start`.

Starting services on the IBM SPSS Modeler Server

You must start the IBM SPSS services on the IBM Capacity Management Analytics node computer.

Procedure

1. On Linux operating systems, log in to the node computer as the administrator or as a user with sudo or root.
2. Go to the IBM SPSS Modeler Server directory. For example, go to `/opt/ibm/SPSS/ModelerServer/17.1`. If you manually installed and accepted the defaults, the path is `/usr/ibm /SPSS/ModelerServer/17.1`.
3. Enter the following command: `./modelersrv.sh start`.

Appendix D. Capacity Management Analytics and the IBM Sub-Capacity Reporting Tool (SCRT)

The IBM Sub-Capacity Reporting Tool (SCRT) is a no-charge tool that IBM policy requires to be used to report license capacity for sub-capacity eligible products that run on the z/OS system. If you purchased Capacity Management Analytics on z/OS, which is a sub-capacity eligible product, you will be required to submit the SCRT reports to IBM.

For information about SCRT, see the IBM System z Software Pricing web site.

At startup, all Capacity Management Analytics on z/OS included products that run on z/OS must register themselves for sub-capacity pricing with the Capacity Management Analytics on z/OS product identifier (PID) and related information. System Management Facility (SMF) 89 records are written as part of the sub-capacity registration process. When the SCRT is run, the SMF 89 records are processed and IBM Capacity Management Analytics on z/OS is recognized as being installed and is eligible for sub-capacity pricing.

Each included product provides the support to properly register as Capacity Management Analytics on z/OS by the following APARs:

- PM94322 for Cognos Business Intelligence for z/OS
- PM94316 for SPSS Modeler with Scoring Adapter for zEnterprise
- PM94238 for IBM Tivoli Decision Support for z/OS

Before you continue, you must install the products and apply the APARs as soon as they become available. Otherwise, the included product's own PIDs shows up in the SCRT report instead of IBM Capacity Management Analytics on z/OS. If that happens and the APARs are not yet available, you must add an Exclude control statement for it when you run SCRT.

Registering IBM Cognos BI for z/OS

You must copy the IBM Capacity Management Analytics on z/OS registration file to you IBM Cognos Business Intelligence for z/OS installation.

The registration information file is provided with the IBM Capacity Management Analytics Solution Kit. The file is in the `CMA_ROOT\configuration` directory.

Procedure

1. On the computer where you installed the Capacity Management Analytics Solution Kit, go to the `CMA_ROOT\configuration` directory.
2. Use FTP or another means to copy `CMApid.registrar` as a binary file to the z/OS system where IBM Cognos BI is running.

   Important: Copying `CMApid.registrar` as a binary file avoids conversion to EBCDIC.
3. Do one of the following steps:
   - If IBM Cognos is running on WebSphere Application Server for z/OS, copy `CMApid.registrar` to the `WAS_instance/AppServer/properties/version` directory.
   - If IBM Cognos is running on Tomcat, copy `CMApid.registrar` to the `c10_location/configuration` directory.
4. Set the environment variable CGC_SCRT_REGISTRATION_FILE to be the fully qualified path and file name of the registration file for the user ID in which IBM Cognos is running. This is typically done in the `.profile` for that user.
   For example: `export CGC_SCRT_REGISTRATION_FILE=$WAS_instance/AppServer/properties/version/CMApid.registrar`

If IBM Cognos is running on WebSphere Application Server, you must also set the environment variable in the administrative console in **Application Servers** > ***server_name*** > **Process Definition** > **Environment Entries**.

5. Restart your IBM Cognos BI services.

Registering IBM SPSS Modeler with Scoring Adapter for zEnterprise

You must copy the IBM Capacity Management Analytics on z/OS registration file to your IBM SPSS Modeler with Scoring Adapter for zEnterprise installation.

The registration information file is provided with the IBM Capacity Management Analytics Solution Kit. The file is in the *CMA_ROOT*`\configuration` directory.

Procedure

1. On the computer where you installed the Capacity Management Analytics Solution Kit, go to the *CMA_ROOT*`\configuration` directory.
2. Use FTP or another means to copy `CMApid.registrar` as a binary file to the z/OS system where IBM SPSS Modeler with Scoring Adapter for zEnterprise is running.

   Important: Copying `CMApid.registrar` as a binary file avoids conversion to EBCDIC.
3. Copy `CMApid.registrar` to the directory where the scoring adapter is installed. The directory must also be defined in the LIBPATH environment variable in the IBM DB2 User Defined Function for the scoring adapter.

   The default path is `/usr/lpp/spss/cfscoring_16.0.0`.
4. Rename `CMApid.registrar` to `pid.registrar`.
5. Restart the Workload Manager Application Environment in which the scoring adapter is running by using a QUIESCE and START command.

Registering IBM Tivoli Decision Support for z/OS

You must copy the IBM Capacity Management Analytics on z/OS registration file to your IBM Tivoli Decision Support for z/OS installation.

The registration information file is provided with the Capacity Management Analytics Solution Kit. The file is in the *CMA_ROOT*`\configuration` directory.

A sample JCL job named `DRLJCCMA` is provided with Tivoli Decision Support for z/OS in APAR PM94238. The sample creates the registration load module for Tivoli Decision Support for z/OS from the Capacity Management Analytics supplied file containing the registration information.

Procedure

1. On the computer where you installed the Capacity Management Analytics Solution Kit, go to the *CMA_ROOT*`\configuration` directory.
2. Use FTP or another means to copy `CMApid.registrar` to the z/OS system where IBM Tivoli Decision Support for z/OS is running.

   Important: `CMApid.registrar` must be copied as EBCDIC for IBM Tivoli Decision Support for z/OS.
3. Copy `CMApid.registrar` to a data set of your choice.

   This should be a z/OS sequential data set or PDS member with RECFM=FB and LRECL=80.
4. Edit the JCL in `DRLJCCMA` in the `SDRLCNTL` data set.

   Details are provided in the comments in the JCL header. Ensure that the DRLPIN DD name of the READ step refers to the EBCDIC copy of `CMApid.registrar`.

5. Run the JCL.

   The JCL reads the IBM Capacity Management Analytics on z/OS registration file and builds a registration load module in the Tivoli Decision Support for z/OS load library SDRLLOAD.
6. Run Tivoli Decision Support for z/OS.

Verifying that products are registered correctly

You can use the JCL sample script HCMPROD to verify the registration of IBM Capacity Management Analytics on z/OS. The sample script is in the *CMA_ROOT*/scripts directory. On z/OS systems, the sample file is found in the samples directory in the SMP/E installation folder.

Before you begin

Verify that IBM Capacity Management Analytics components are running. You can verify that the registration was applied correctly only if the products are running.

Procedure

1. If you installed on a Microsoft Windows or a Linux on System z operating system, upload the HCMPROD script to the z/OS machine, and copy into one PDS dataset with RECFM=FB and LRECL=80.
2. Change #SMFDS# to the SMF record dataset.
3. Submit this job and check the SYSPRINT of the joblog in the SDSF panel.

   The results show IBM CMA with each product or feature.

```
 Totals                                      445       0    46      290
-Product       Product     Product                 Product     Product
 Owner         Number      Name                    Version     Function
 ---------     --------    ----------------        --------    --------
 IBM CORP      5694-A01    z/OS                    01.13.00    z/OS
 IBM CORP      5694-A01    z/OS                    01.13.00    JES2
 IBM CORP      5694-A01    z/OS                    .  .        RACF
 IBM CORP      5694-A01    z/OS                    .  .        RMF
 IBM CORP      5694-A01    z/OS                    .  .        Security Server
 IBM CORP      5694-A01    z/OS                    .  .        SDSF
 IBM CORP      5694-A01    z/OS                    .  .        TCP/IP BASE
 IBM CORP      5645-001    z/OS                    01.13.00    TSO/E
 IBM CORP      5655-Y26    Cognos BI 10.2.1        NOTUSAGE    COGNOS
 IBM CORP      5655-Y26    Cognos BI 10.2.1        NOTUSAGE    USINGWAS
 IBM CORP      5698-AA6    IBM CMA 2.1             NOTUSAGE    COGNOS
 IBM CORP      5698-AA6    IBM CMA 2.1             NOTUSAGE    USINGWAS
 IBM CORP      5655-W65    WAS FOR Z/OS            NOTUSAGE    WAS Z/OS
```

Appendix E. Security considerations for IBM Capacity Management Analytics

This section provides information about the users and roles that are used in IBM Capacity Management Analytics and about password usage in configuration and script files.

Users and roles

Multiple roles are required to install, configure, and use IBM Capacity Management Analytics. This section provides a brief introduction to these roles or users.

Tivoli Decision Support for z/OS database user

To initialize the IBM DB2 database when you install Tivoli Decision Support for z/OS, you must be a user that has DB2 SYSADM authority or a user that has the authority to create plans, storage groups, and databases, and who has access to the DB2 catalog.

For more information about the security and permission requirements, see the Tivoli Decision Support for z/OS documentation (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2/com.ibm.tivoli.dszos.doc.1.8.2/Admin/asecure.html).

For more information about initializing the database, see the Tivoli Decision Support for z/OS documentation (www.ibm.com/support/knowledgecenter/SSH53X_1.8.2/com.ibm.tivoli.dszos.doc.1.8.2/Admin/aininor.html).

Capacity Management Analytics database user

The Capacity Management Analytics data in the database is protected by DB2. Administrators and users must be granted DB2 privileges to be able to access the data.

The Capacity Management Analytics database administrator should have DB2 SYSADM authority (or be a user with the authority to create plans, storage groups, and databases, and have the authority to grant create tablespace, table, view, UDF, storage procedure, and index privileges to other users.

The Capacity Management Analytics database user should have createdba, createalias, bindadd, accessctrl, and dataaccess privileges. The user also needs packadm on all collections. The user should have full access to the HCM database (the Capacity Management Analytics database) and DRL database (the Tivoli Decision Support for z/OS database).

SPSS Modeler Server user

IBM SPSS Server normal need to be installed/start/stop by root user, but SPSS Modeler does provide no-root start/stop functions.

For more information, see the SPSS Modeler documentation (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/clementine/server/adminguidesource/non_root_unix_container.dita).

When you use SPSS Modeler Server with IBM Capacity Management Analytics, the user must have read and write access to the folder, data files, temporary files, and log files that are used in the Capacity Management Analytics streams.

IBM Cognos Business Intelligence user

For more information about user account information for IBM Cognos Business Intelligence for z/OS, see the User accounts for the installation and running the product in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J_10.2.2/com.ibm.swg.ba.cognos.c8bizos_ig.10.2.2.doc/c_cbizos_prepcrteusrgrp.html).

For more information about user account information for IBM Cognos Business Intelligence on distributed platforms, see Configure a User Account or Network Service Account for IBM Cognos Business Intelligence in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J_10.2.2/com.ibm.swg.ba.cognos.inst_cr_winux.10.2.2.doc/c_configureauseraccountforcognos8.html%23ConfigureAUserAccountForCognos8).

Password handling in IBM Capacity Management Analytics

IBM Capacity Management Analytics requires multiple user passwords during installation and daily usage. These passwords include an IBM Cognos Business Intelligence account password, a WebSphere Application Server administration password, a database administration password, a database user password, and an SPSS Modeler Server password.

During the installation of IBM Capacity Management Analytics

During the IBM Capacity Management Analytics installation, you are asked to enter all required passwords. After the installation, the passwords are encrypted and stored in the `*CMA_ROOT*/bin/setenv.sh` file. You can change the passwords by editing the `setenv.sh` file.

Note: If you must change passwords in the `setenv.sh` file, you must ensure that the passwords are encrypted.

All passwords are encrypted by using the AES 256 encryption algorithm. You can use `cmaencrypt.jar` to encrypt a password. For example, you can use the following command to encrypt a password: `java -jar cmaencrypt.jar *yourpassword*`.

Standard Java™ does not support 256-bit encryption. You must replace the security policy file in your Java `lib` directory before you use `cmaencrypt.jar`.

While configuring IBM Capacity Management Analytics

When you configure IBM Capacity Management Analytics, the script files use the password from the `setenv.sh` file and the passwords are decrypted when you are using them. The decryption program does not expose the password to the user. Also, all Capacity Management Analytics SQL and configuration step execution is initialized by a Java program to prevent passwords from being exposed during execution.

While using IBM Capacity Management Analytics

When you run the Capacity Management Analytics streams, the SPSS Modeler Server password is read from the `setenv.sh` file. The password is decrypted in a Java program and used to connect to SPSS Modeler Server.

The SPSS Modeler streams require input stream parameters, including database passwords. SPSS Modeler provides a password encoding function in SPSS Modeler Client. You can use the encoding function to encode your database password and then enter the encrypted password in the `parameter.txt` file.

For more information about encoding passwords in SPSS Modeler, see the SPSS Modeler documentation (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/clementine/epassword_tool.dita)

Appendix F. Error, warning, and installation messages

Error, warning, and installation messages are displayed for IBM Capacity Management Analytics.

Modeling error and warning messages

Running the IBM Capacity Management Analytics anomaly streams, SCA streams, and ETL streams generate stream logs in the log folder. The log folder is specified by the LOGPATH parameter.

Log file names are formed by the stream name and a time stamp. For example, a `prod_bill_msu_optimization.str` log would be named `prod_bill_msu_optimization_3629024907.log`.

If a stream failed to run, error or warning messages are written to the log file.

CMA-ME001

UID and/or PWD for database is incorrect.

Explanation

The account name for parameter UID or the password for parameter PWD is invalid.

The execution of the current stream has stopped and no valid result is produced.

User Response

Check and make sure the account name or the password is valid, then rerun the stream.

CMA-ME002

No database connection has been selected.

Explanation

No database connection is provided in the data source node in filed “data source.” This can happen when you run a single branch of a stream manually and the branch has a data source node.

User Response

Ensure that there is a non-empty and valid data connection in the data source node in filed “data source.” Then, rerun the stream branch.

CMA-ME003

CMDW (CMA database connection) is incorrect.

Explanation

The value for the parameter CMDW refers to an invalid database connection.

User Response

Ensure the value for the CMDW parameter is a valid database connection and that it exists on the computer on which the stream is running. Then, rerun the stream.

CMA-ME004

Excel input file type is incorrect.

Explanation

Capacity Management Analytics uses Microsoft Excel 2003 formatting. The input file that you specified might be another format.

User Response

Convert your file to Microsoft Excel 2003 format.

CMA-ME005

<transaction> transaction doesn't have enough data for model building.

Explanation

The data recorded in the database for the specified type of <transaction> is not enough, and the anomaly model cannot be built.

User Response

Change model building date in the `transaction.csv` file to ensure that the records for the specified <transaction> in the DRL.CICS_T_TRAN_T table are greater than 40.

CMA-ME006

<transaction> transaction has exceed the data limitation for model building, which is SIZE. Please sample your data before you build the model.

Explanation

The specified <transaction> contained too much data for model building. IBM SPSS Modeler has limits to the amount of data used in model building process. The limitation can be found in the Modeler documentation (http://www.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/server/adminguidesource/admin_ram.htm).

User Response

Change model building date in the `date_info.csv` file to ensure that the input data records are less than the limitation. Alternatively, you can sample the input data to reduce the data records.

CMA-ME007

<transaction> transaction incorrect data, which is response time is less than CPU second. Please check your input data.

Explanation

The specified <transaction> contained invalid records that indicate that the transaction response time is less than a CPU second.

User Response

Check the input data for the specified <transaction> and correct the records so that the response time is greater than a CPU second.

CMA-ME008

Do not have permission to read/write to `<file>` or this file path is incorrect.

Explanation

The file path you specified in stream parameter is incorrect or Capacity Management Analytics does not have the correct permission to read or write to the file.

User Response

Ensure the path and permissions are correct to read or write to the `<file>`.

CMA-ME009

No <parameter> specified!

Explanation

You have not specified the parameter <parameter>.

User Response

Specify the value for <parameter> in `parameter.txt` or in Modeler client.

CMA-ME010

CMA Schema or TDSz Schema is incorrect, or related table did not exist in database.

Explanation

The parameter you specified for CMASCHEMA or TDSz is incorrect or the related table does not exist in database.

User Response

Modify the parameter you specified for CMASCHEMA or TDSz and ensure the related table exists in the database.

CMA-ME011

Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> can't fit a model, please inspect your data.

Explanation

The modeler could not fit a model in time series for Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no>.

User Response

Inspect your data of Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> to ensure that it contains valid data or has enough data for building a model.

CMA-ME012

The storage of certain fields or the number of total fields does not comply to the expected value as defined in the source node.

Explanation

The storage of certain fields has changed or certain fields have been removed or added to the inputted data source for the current stream.

The execution of the current stream was stopped and no valid result is produced.

User Response

Get the source node from the node indicated in the log where the stream interruption occurred. Check whether the data definition (The storage of certain fields or the number of total fields) in the data source has changed. If so, try to restore it or use the original data source to rerun the stream.

CMA-ME013

The total capacity of the LPARs is not enough. No optimization suggestion could be given and the optimization-related reports could not be generated!

Explanation

There are still some products that are left and they cannot be arranged because each LPAR reached its maximum capacity.

The execution of the optimization stream was stopped and no optimization result was produced.

User Response

The defined capacity of some LPARs should be increased if USERSELECTION=define to give enough space to do the optimization, and, at the same time, the weight of these LPARs should also be increased if USERSELECTION=weight. Also, you can try to free the products if it is not necessary to be reserved in a certain LPAR in the setting file.

If these changes are not allowed or they do not increase the capacity, then the optimization cannot be conducted for these products against these LPARs.

CMA-ME014

No capacity in each LPARs before arrangement, please check your data or configuration.

Explanation

Each LPAR has a heavy load because the N89 program, ZOS program, or the system overhead has a large MSU load that is already beyond the capacity of each LPAR in the data.

The execution of the optimization stream was stopped and no optimization result is produced.

User Response

Ensure that the MSU of the LPARs is valid. Then try to increase the capacity for each LPAR by checking the setting of define capacity and the weight setting of each LPAR. It should not be too low. In this case, the define capacity of each LPAR should be increased to provide enough space for the optimization if USERSELECTION=define, and, at the same time, the weight of these LPARs should also be increased if USERSELECTION=weight.

CMA-ME015

No suitable LPAR in the data!

Explanation

The data inputted to CMA is invalid and contains no LPAR information.

The execution of the optimization stream stopped and no optimization suggestion is given.

User Response

Check and make sure the data is valid.

CMA-ME016

No product needs to be optimized!

Explanation

The data during optimized period contains null value for field PROD_ID or each of the optimizing product in field PROD_ID have zero MSU in the history data. So the data is invalid or has nothing to be optimized. The optimizing products refer to those that are waiting to be optimized.

The execution of the optimization stream has stopped and no optimization suggestion is given.

User Response

Check why the optimizing product has zero MSU. The possible reasons include the machine was halted recently in a long term, data inputted to CMA is invalid, or abnormal execution results in CMA forecast processing.

CMA-ME017

The product *<product>* is cared but not suitable for any LPAR.

Explanation

A product can be specified as a cared product which has priority and is arranged first. If it is not suitable for any of the LPARs, a dead loop may occur when the stream is run. To avoid this, the stream will detect this and interrupt.

The execution of the optimization stream stopped and no optimization suggestion is given.

User Response

Check if the capacity of all the LPARs are normal and check if *<product>* has an valid amount of MSU.

CMA-ME018

Invalid condition meet, optimization stopped.

Explanation

The optimization stream encountered an unknown problem and cannot continue to output the result.

The execution of the optimization stream stopped and no optimization suggestion is given.

User Response

Do not delete any logs or temporary files. Contact customer support to resolve the issue.

CMA-ME019

Do not have permission to read/write to the path or path is incorrect. Log *<file>* cannot be created!

Explanation

The log path you specified for LOGPATH is incorrect or we do not have permission to write to it.

User Response

Ensure that the LOGPATH parameter you specified in `parameter.txt` or Modeler client is correct, and ensure we can write a log file to it.

CMA-ME020

CPU and response scoring results didn't have intersection data for classify results.

Explanation

The same transaction does not have both CPU sec and response time scoring results in the database. This is needed for classification stream.

User Response

The scoring stream needs to be rerun based on both CPU sec model and response time for the same transaction.

CMA-ME021

Didn't find any CPU and response scoring results in database, please run CPU and response scoring before execute classification stream.

Explanation

No scoring results for anomaly detection in database.

User Response

The scoring stream needs to be rerun based on both CPU sec model and response time model for the same transaction.

CMA-ME022

Didn't find any response scoring results in database, please run response scoring before execute classification stream.

Explanation

Response time scoring results for anomaly detection didn't exist in database.

User Response

The scoring stream needs to be rerun based on response time model for the same transaction.

CMA-ME023

Didn't find any cpu scoring results in Database, please run cpu scoring before execute classification stream.

Explanation

CPU second scoring results for anomaly detection did not exist in the database.

User Response

The scoring stream need to be rerun based on CPU second model for the same transaction.

CMA-ME024

Can not access the database because password for database is expired.

Explanation

The password you specified in parameter PWD for database is expired. We cannot access the database.

User Response

Modify your password for the database, and then specify the new password for parameter PWD.

CMA-ME025

Unidentified error.

Explanation

We cannot identify the error and will show you the Modeler error.

User Response

Review the Modeler exception and take the appropriate measures.

CMA-ME026

<transaction> transaction doesn't have enough data for model scoring.

Explanation

The specified *<transaction>* does not have any scoring input data.

User Response

Change the model scoring date in the `date_info.csv` file to ensure that the specified *<transaction>* has records.

CMA-ME027

Wrong value for USERSELECTION specified! Input either 'define' or 'weight'.

Explanation

The value inputted for the USERSELECTION parameter for the optimization stream is invalid. It should be either 'define' or 'weight'.

The execution of the optimization stream stopped and no optimization suggestion is given.

User Response

Correct the value for the USERSELECTION parameter, and re-run the stream.

CMA-ME028

User does not have enough privilege. Please access the database with high privilege account.

Explanation

The database user specified in the Modeler stream parameters does not have enough privilege to operate the table.

User Response

Use another database user with higher privilege or grant specified privilege to the user used in stream.

CMA-ME029

Wrong value for <parameter> specified! Input either 'true' or 'false'.

Explanation

<parameter> input value is incorrect.

User Response

Check the `parameter.txt` file or parameter input in the Modeler stream and ensure the input value for *<parameter>* is true or false.

CMA-ME030

start_date is incorrect

Explanation

Parameter start_date is incorrect.

User Response

Check the `parameter.txt` file or parameter input in the Modeler stream and ensure the input value for start_date is in the correct format.

CMA-ME031

end_date is incorrect

Explanation

Parameter end_date is incorrect.

User Response

Check the `parameter.txt` file or parameter input in the Modeler stream and ensure the input value for end_date is in the correct format.

CMA-ME032

Record record is not valid, following case may meet both report class and job name filled with *, both of them have specified value, capture ratio is greater than 1.

Explanation

Parameter end_date is incorrect.

User Response

Check the `parameter.txt` file or parameter input in the Modeler stream and ensure the input value for end_date is in the correct format.

CMA-ME033

<job_name> has duplicate records.

Explanation

Parameter <job_name> has duplicate records.

User Response

Check your data and rerun the stream.

CMA-ME034

Report class <class_name> has duplicate records.

Explanation

Parameter <class_name> has duplicate records.

User Response

Check your data and rerun the stream.

CMA-ME035

Server <Server> has duplicate records.

Explanation

Parameter <Server> has duplicate records.

User Response

Check your data and rerun the stream.

CMA-ME036

Record<record> is invalid, no record existed in TDSz tables.

Explanation

There is no record that exists in the TDSz tables, so the record is invalid.

User Response

Check your data and ensure the record exists in the TDSz tables.

CMA-ME037

Application *app_name* for environment *environment* function *function* on CPC Series NO *cpc_series_no* MVS_SYSTEM_ID *mvs_system_id* cannot fit a model. Please check your data.

Explanation

The SPSS Statistics Expert Modeler could not fit a model.

User Response

Check your data for application *app_name* for environment *environment* function *function* on CPC Series NO *cpc_series_no* MVS_SYSTEM_ID *mvs_system_id* to ensure that it contains valid data or that it has enough data to build a model.

CMA-ME038

Application *app_name* for environment *environment* function *function* on CPC Series NO *cpc_series_no* MVS_SYSTEM_ID *mvs_system_id* is missing values. Please check your data.

Explanation

IBM SPSS Modeler could not fit the model in time series for application *app_name* for environment *environment* function *function* on CPC Series NO *cpc_series_no* MVS_SYSTEM_ID *mvs_system_id* due to missing values.

User Response

Check your data for application *app_name* for environment *environment* function *function* on CPC Series NO *cpc_series_no* MVS_SYSTEM_ID *mvs_system_id* to ensure that it has enough data to build a model.

CMA-ME039

Parameter HOUR_START and HOUR_END must be in range [0, 23], HOUR_START must not larger than HOUR_END.

Explanation

The values that you specified for parameter HOUR_START and HOUR_END must be an integer and between 0 and 23. HOUR_START must not larger than HOUR_END.

User Response

Check your data for the values that you specified for the HOUR_START and HOUR_END parameters.

CMA-ME040

DATE_START must be earlier than DATE_END.

Explanation

The value that you specified for the DATE_START parameter must be earlier than the value that you specified for the DATE_END parameter.

User Response

Check your data to ensure that the DATE_START value is earlier than the DATE_END value.

CMA-ME041

Wrong PEAKRANK input format. Valid format is 'level1,level2,...,leveln' and 0<level<'*<hour>*'.

Explanation

The input parameter PEAKRANK has an incorrect value. The value must be formatted as 'level1,level2,...,leveln' or it is not in the range of 0<level<*<hour>*, where *<hour>* is automatically calculated from the HOUR_START and HOUR_END parameters.

User Response

Check your data to ensure that the PEAKRANK value is in the correct format.

CMA-ME042

Wildcard character in Job_name <JOB_NAME> is incorrect, please check input application mapping file.

Explanation

A wildcard character was found in <JOB_NAME> and is incorrect. Full values must be entered.

User Response

Check the input application file and ensure that you enter the full value. Then, rerun the stream.

CMA-ME043

No needed data in the table *<Schema.TableName>* from date_start to date_end.

Explanation

The process failed because no data was not found in the *<Schema.TableName>* table that is within the time range from data_start to data_end.

User Response

Check your data to ensure that the date range is correct or that the data in the *<Schema.TableName>* table is correct.

CMA-ME044

No application data to process.

Explanation

The input data is invalid or empty.

User Response

Check to ensure that you have valid input data.

CMA-ME045

No Transaction specified in the anomaly input file.

Explanation

No transaction has been specified in the anomaly input file.

User Response

Specify a transaction to be examined in the `transaction.csv` file to ensure that your model can be built.

CMA-ME046

The start date specified in the file is greater than the end date.

Explanation

The value that you specified for the DATE_START must be earlier than the DATE_END parameter.

User Response

Check your data to ensure that the DATE_START value is earlier than the DATE_END value.

CMA-ME047

TIME_START should be earlier than TIME_END.

Explanation

The value that you specified for the TIME_START must be earlier than TIME_END.

User Response

Check your data to ensure that the TIME_START value is earlier than the TIME_END value.

CMA-ME048

The score ratio is greater than 1. The score ratio should be greater than 0 but less than or equal to 1.

Explanation

The SCORE_RATIO value that defines the ratio that the tree-based anomaly model uses to judge whether the transaction is abnormal or not is greater than 1. The default value is 0.5, with a recommended range of this parameter value of [0.5,0.6].

User Response

Edit the value of the parameter SCORE_RATIO and ensure it does not exceed 1. Use the recommended values.

CMA-ME049

There were no tree-based scoring results in the database. Run the tree-based scoring before you execute the classification system.

Explanation

The `anomaly_classification.str` can be run only after all other scoring streams have run. The HCM anomaly table should contain CPU, response, and tree scoring results for the same transaction.

User Response

Run all scoring streams, including the tree-based scoring stream, then run the classification system.

CMA-MW001

Product ID *<product>* on LPAR *<lpar>* and CPC Series NO *<cpc_series_no>* has been obsoleted from *<date>*, it will be skipped in forecast sequence.

Explanation

If the latest date in Product ID *<product>* on LPAR *<lpar>* and CPC Series NO *<cpc_series_no>* (DA) is earlier than the latest date of all the product (DB), and DB minus DA is greater than or equal to ODD (obsolete of days which client specified in stream), then it will be skipped in forecast sequence.

User Response

If you want the product to be included in the forecast sequence, specify a larger number for the parameter ODD or modify the records in the database so that DB minus DA is less than ODD.

CMA-MW002

Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> on <weekday> contained null value and does not have 24 hours data in one day, it cannot be forecasted.

Explanation

If the product in one LPAR of one CPC contains only null values in any day of the week and does not have 24 hours of data in any day of the week, it cannot be forecasted.

User Response

Fill in the null values and enter the missing values in the database to correct this warning.

CMA-MW003

Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> on <weekday> does not have 24 hours data in one day, it cannot be forecasted.

Explanation

If the product in one LPAR of one CPC does not have 24 hours of data in any day of the week, it cannot be forecasted.

User Response

Enter the missing values, according to your SMF data, in the database to correct this warning.

CMA-MW004

Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> on <weekday> contained a null value, it cannot be forecasted.

Explanation

If the product in one LPAR of one CPC contains only null values in any day of week, it cannot be forecasted.

User Response

Enter the missing values in the database to correct this warning.

CMA-MW005

Product ID <product> on LPAR <lpar> and CPC Series NO <cpc_series_no> contained bad data, stream automatically fix it and forecast.

Explanation

Product ID <product> on LPAR <lpar> and CPC Series NO<cpc_series_no> contained data which PROD_MSU is not less than PROC_CAPACITY_MSU, Stream will replace the data with last normal data in the same hour of the same week of day for the specified product.

User Response

Correct the data in database for the specified product so that the PROD_MSU is always less than PROC_CAPACITY_MSU.

CMA-MW006

The record count of LPAR <lpar> on processor type <processor type> and CPC Series NO <CPU series no> is less than the limitation of the algorithm, do not forecast.

Explanation

Historical specified combination is less than algorithm requirement.

User Response

Add enough historical data for this combination. Hourly forecast required at least 721 records, daily forecast required at least 61 records, monthly forecast required at least 4 records.

CMA-MW007

LPAR <lpar> on processor type <processor type> and CPC Series NO <cpc_series_no> have incorrect scoring result, the predict can't be processed.

Explanation

Scoring results for specified combination is incorrect, and forecasted result will not be ignored.

User Response

Check the historical data for this specified combination. If the variance of historical data is too big, then this combination can not be handled by the algorithm.

CMA-MW008

Product ID <product> cannot be put in LPAR <lpar> due to no information, forbidden in setting, or have no overlapped duration, ignored!

Explanation

The product referred by <product> cannot be put in LPAR referred by <lpar> during the optimization arrangement.

The execution of the optimization stream will not stop.

User Response

There is no user response. The message indicates that some products were not placed in LPARs.

CMA-MW009

The total capacity of the LPAR <lpar> is increased <xxx> percent to host the remaining of the Product ID <product>.

Explanation

This warning will follow after CMA-MW012.

This warning indicates that some products are left and there is no LPAR capacity for them to be arranged into. But the optimization algorithm finds a suitable LPAR referred by <lpar> to increased its capacity (the increasing amount is indicated by <xxx> percent) and hosted the product referred by <product>. The capacity is increased means that the actual overload for this LPAR after the optimization may be over the configured capacity, but the optimization stream will not change any configuration for the capacity of this LPAR. The increasing amount is the expected maximum amount above that configured capacity. The capacity is permitted to be increased to the limitation that not above the NOMINAL CAPACITY of the LPAR. The reason that some products are left to be arranged is because the total capacity of all the LPARs is not enough to host all the products.

The execution of the optimization stream will continue, it may have another attempt to increase other LPARs for other products which cannot be arranged. In that case, there will be multiple CMA-MW009 warnings.

User Response

There is no user response.

CMA-MW010

Do not have permission to save models to model path or model path is incorrect.

Explanation

The model path you specified in stream parameter is incorrect or we do not have the permission to save model to the path.

User Response

Modify the OUTPUTMODELPATH parameter you specified in `parameter.txt` or Modeler client to a valid path.

CMA-MW011

LPAR <lpar> is invalid and ignored!

Explanation

LPAR referred by <lpar> is not suitable for all products to be arranged in it, the stream will continue to run, but this LPAR will be ignored for the succeeding processing.

User Response

There is no user response.

CMA-MW012

The total capacity of the LPARs is not enough for Product ID <product>, the optimization is now trying to increase the capacity of certain LPAR.

Explanation

This warning indicates that some products are left and there is no LPAR capacity for them to be arranged into. The optimization algorithm is trying to locate a suitable LPAR to host the <product> product by increasing the capacity of the LPAR.

User Response

There is no user response.

CMA-MW013

Failed to reading `cma_message.csv`, detailed CMA information cannot be provided.

Explanation

The Modeler stream cannot read `cma_message.csv` so some of CMA error cannot be shown.

User Response

On Windows operating systems, ensure that you have read and write permissions for *$CMA_HOME*/Streams/*data folder*. On Unix or Linux operating systems, ensure that the `Streams` folder is copied to the correct location and *$CMA_HOME* environment variable is set correctly.

CMA-MW014

Job xxx and service xxx have intersections.

Explanation

Job xxx belongs to service class xxx. Both are mentioned in the input file **appl_lob_z.csv**, causing double counting in the calculation.

User Response

Remove either the job or the class in the input file **appl_lob_z.csv**.

CMA-MW015

The stream used actual history data from *<DATE_MIN>* to *<DATE_MAX>*.

Explanation

If the values that you specified for the DATE_START and DATE_END parameters are out of range between the actual history start date and the actual history end date. The stream forecast seconds using the actual history start date or actual history end date and a warning message is displayed to show the data range that was used in the stream.

User Response

There is no user response.

CMA-MW016

The input start date *<input_data_start>* is earlier than the actual observed start date *<observ_date_start>*. The later date was used.

Explanation

The date in the DATE_START parameter is earlier than the actual observed start date. The stream used the actual observed start date instead.

User Response

There is no user response.

CMA-MW017

The input end date *<input_data_end>* is later than the actual observed end date *<observ_date_end>*. The later date was used.

Explanation

The date in the DATE_END parameter is later than the actual observed end date. The stream used the actual observed end date instead.

User Response

There is no user response.

CMA-MW018

The selected period contains incorrect capture ratio for the job. The capture ratio must be less than or equal to 1.

Explanation

The specified period contains an incorrect capture ratio for the job.

User Response

Change the selected period so the capture ratio remains below or at a value of 1.

CMA-MW019

The selected period contains incorrect capture ratio for the report class. The capture ratio must be less than or equal to 1.

Explanation

The specified period contains an incorrect capture ratio for the report class.

User Response

Change the selected period so the capture ratio remains below or at a value of 1.

CMA-MW020

The mapping table does not contain any records and no data is exported to the mapping table. Run `appl_lob_z.str` to export the data.

Explanation

The mapping table contains no records and no data is exported to the mapping table.

User Response

Run `appl_lob_z.str` to export the data.

Reporting error and warning messages

Error or warning messages are displayed in reports and workspaces. In reports, the message summary is displayed in the right top of the subtitle bar. In workspaces, the message summary is displayed at the end of the global prompts area on the left. The message summary is in red text. You can click the message to show the details.

CMA-RW001

Partial Interval Band Warning.

Explanation

This warning message is displayed because partial interval band is detected for your selection on date/time range and interval band value. Please consider one of the following suggestions:

- Change date/time range to be a multiple of the interval band.
- Change interval band selection to a smaller value.

Condition to display

When you select a partial interval band with your date/rime range and your interval band choices.

CMA-RW002

Scenario Synchronization Warning.

Explanation

The current optimization result is based on the forecast result running at **<TIMESTAMP>**. There are new updates. Please run optimization again.

Condition to display

This warning message is for the Software Cost Analytics feature, and is issued when optimization records are later than forecast records.

Optimization error and warning messages

The following error or warnings may occur for the LPAR weight optimization.

CMA-OE000

The message file **cma_messages.csv** is not found in the default location.

Explanation

This message is displayed when the file that contains error and warning messages is not found in the default location `CMA_HOME/optimizations/data`.

User Response

Ensure the file **cma_messages.csv** is under the default location `CMA_HOME/optimizations/data` and rerun the optimization.

CMA-OE001

`<lparName>` does not have records for every hour.

Explanation

This message is displayed because the LPAR does not have records for each hour in the time range for the dates indicated.

User Response

Ensure the chosen dates and time ranges that have a consistent and representative set of LPAR records as input and rerun the optimization.

CMA-OE002

<lparName> does not have SMF72 records for any hour.

Explanation

This message is displayed because the LPAR <lparName> does not have SMF72 records for any hour in the LPAR parameter file <fileNameParaPerLpar>.

User Response

To resolve this issue, set the **TARGET_WORKLOAD_IMPORTANCE_LEVEL** as `null` and rerun the optimization. SMF70 data will be used for the LPAR.

CMA-OE003

CPU_SERIAL_NO is missing in the parameter file.

Explanation

CPU_SERIAL_NO is missing in the parameter file <fileNameParameter>.

User Response

To resolve this issue, enter the item and value in the file and rerun the optimization.

CMA-OE004

Invalid **HOUR_START** in the parameter file.

Explanation

This message is displayed because there is an invalid **HOUR_START**: <hourStart> in the parameter file <fileNameParameter>. **HOUR_START** is expected to be an integer between 0 and 23.

User Response

To resolve this error, correct the value in the parameter file and rerun the optimization.

CMA-OE005

HOUR_START is missing in the parameter file.

Explanation

This message is displayed when **HOUR_START** is missing in the parameter file <fileNameParameter>.

User Response

To resolve this error, enter the item and value in the parameter file and rerun the optimization.

CMA-OE006

Invalid **HOUR_END** in the parameter file.

Explanation

This message is displayed when there is an invalid **HOUR_END**: <hourEnd> in the parameter file <fileNameParameter>. **HOUR_END** is expected to be an integer between 0 and 23 and should be greater than **HOUR_START**.

User Response

To resolve this error, correct the value in the parameter file and rerun the optimization.

CMA-OE007

HOUR_START should be a lesser value than **HOUR_END** in the parameter file.

Explanation

This message is displayed when **HOUR_START** is a higher value than **HOUR_END** in the parameter file <fileNameParameter>. **HOUR_START** should always be a lesser value than **HOUR_END**.

User Response

To resolve this error, correct the value in the parameter file and rerun the optimization.

CMA-OE008

HOUR_END is missing in the parameter file.

Explanation

This message is displayed when **HOUR_END** is missing in the parameter file <fileNameParameter>.

User Response

To resolve this error, enter the item and value in the parameter file and rerun the optimization.

CMA-OE009

Invalid **PROCESSOR_TYPE** in the parameter file.

Explanation

This message is displayed when the **PROCESSOR_TYPE** is invalid in the parameter file <fileNameParameter>. The value can be CP, zIIP, zAAP, or IFL.

User Response

To resolve this error, correct the value in the parameter file and rerun the optimization.

CMA-OE010

PROCESSOR_TYPE is missing in the parameter file.

Explanation

This message is displayed when the **PROCESSOR_TYPE** is missing in the parameter file <fileNameParameter>. The value can be CP, zIIP, zAAP, or IFL.

User Response

To resolve this error, enter the item and value in the parameter file and rerun the optimization.

CMA-OE011

No input date is available or all of the dates specified are invalid in the parameter file.

Explanation

This message means that either no input date is available, or all of the dates specified in the date list parameter file <fileNameDateList> are invalid.

User Response

To resolve this error, enter valid dates in the date list parameter file and rerun the optimization.

CMA-OE012

The decimal format of the LOCALE: <locale> is <decimalFormat>. The **-OPTIM_DELIMITER**:<delimiter> is given as a parameter. This can result in incorrect values for the decimal format numbers from the input csv files.

Explanation

This message means that the given parameter **-OPTIM_DELIMITER** and the decimal format of the locale are the same. This can lead to interpreting incorrect values for the decimal format numbers.

User Response

To resolve this issue, use the correct delimiter that matches with the LOCALE as parameter and rerun the optimization.

CMA-OE013

CPU_SERIAL_NO in the parameter file cannot be found in the table <cmaSchema>.MVSPM_LPAR.

Explanation

This message means that the **CPU_SERIAL_NO** in the parameter file <fileNameParameter> cannot be found in the SMF70 table.

User Response

To resolve this issue, correct the value in the parameter file and rerun the optimization.

CMA-OE014

PROCESSOR_TYPE in the parameter file cannot be found in the table <cmaSchema>.MVSPM_LPAR.

Explanation

This message means that the **PROCESSOR_TYPE** in the parameter file <fileNameParameter> cannot be found in the SMF70 table.

User Response

To resolve this issue, correct the value in the parameter file and rerun the optimization.

CMA-OE015

TIME_RESOLUTION is not consistent between the tables MVSPM_LPAR and MVSPM_WORKLOAD_H.

Explanation

This message is displayed when **TIME_RESOLUTION** in the parameter file <fileNameParameter> is not consistent between the tables <cmaSchema>.MVSPM_LPAR and <cmaSchema>.MVSPM_WORKLOAD_H.

User Response

To resolve this error, use only <cmaSchema>.MVSPM_LPAR records by setting **TARGET_WORKLOAD_IMP_LEVEL** as null in the LPAR parameter file <fileameParaPerLpar>, or use dates that have matching **TIME_RESOLUTION** as input for the optimization, or reload the tables by using the same **TIME_RESOLUTION** and rerun the optimization.

CMA-OE016

The **TARGET_WORKLOAD_IMPORTANCE_LEVEL** of **PROCESSOR_TYPE** IFL is not null in the LPAR parameter file.

Explanation

This message is displayed when the **TARGET_WORKLOAD_IMPORTANCE_LEVEL** of **PROCESSOR_TYPE** IFL is not null in the LPAR parameter file <fileNameParaPerLpar>. The value is expected to be null.

User Response

To resolve this error, set **TARGET_WORKLOAD_IMPORTANCE_LEVEL** to `null` in the LPAR parameter file and rerun the optimization.

CMA-OE017

<lparName> does not have active logical processors for every interval.

Explanation

This message is displayed because the LPAR <lparName> is inactive in the table <cmaSchema>.MVSPM_LPAR.

User Response

To resolve this error, use dates and a time range that have a consistent and representative set of active LPARs as input and rerun the optimization

CMA-OE018

Invalid **FILE_NAME_MODEL_PARA** in the parameter file.

Explanation

This message is displayed when the file **FILE_NAME_MODEL_PARA** specified in the parameter file <fileNameParameter> does not exist.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_MODEL_PARA** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and the file name **lwo_parameters_model.csv** will be used.

CMA-OE019

Invalid **FILE_NAME_MODEL_PARA_PER_LPAR** in the parameter file.

Explanation

This message is displayed because the file **FILE_NAME_MODEL_PARA_PER_LPAR** specified in the parameter file <fileNameParameter> does not exist.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_MODEL_PARA_PER_LPAR** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and the file name **lwo_parameters_per_lpar.csv** will be used.

CMA-OE020

Invalid **FILE_NAME_DATE_LIST** in the parameter file.

Explanation

This message means that the file **FILE_NAME_DATE_LIST** in the parameter file <fileNameParameter> does not exist.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_DATE_LIST** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and file name **lwo_parameters_date_list.csv** will be used.

CMA-OE021

Invalid LPAR during the date and time range.

Explanation

This message is displayed when there is no valid LPAR during the specified date and time range.

User Response

Enter a valid LPAR that is available during the specified date and time range and rerun the optimization.

CMA-OW000

The folder for storing optimization model LP files does not exist.

Explanation

This message is displayed when the folder for storing optimization model LP files cannot be found. A new folder is created in <modelPath>.

CMA-OW001

Duplicate records may be loaded in the table for LPAR.

Explanation

This message is displayed when **MEASURED_SEC** is larger than the **TIME_RESOLUTION** in the table for the specified date and time intervals. As a result, duplicate records may be loaded in the table. To resolve this issue, double check your data.

User Response

Double check your data.

CMA-OW002

Invalid **LPAR_PRIORITY** of LPAR.

Explanation

This message is displayed when the **LPAR_PRIORITY** value is invalid. The value is expected to be a non-negative integer, with a default value of 1.

User Response

The warning can be ignored if the **LPAR_PRIORITY** for the LPAR is 1, or you must change the value in the LPAR parameter file.

CMA-OW003

Invalid **FIXED_WEIGHT** of LPAR.

Explanation

This message is displayed when the **FIXED_WEIGHT** value is invalid. The value should be between 0 and 1. The default value is null.

User Response

The warning can be ignored if the **FIXED_WEIGHT** for the LPAR is 1, or you must change the value in the LPAR parameter file.

CMA-OW004

The **TARGET_WORKLOAD_PERCENT** of LPAR is missing in the parameter file.

Explanation

This message is displayed when the **TARGET_WORKLOAD_PERCENT** value is missing in the LPAR parameter file. The default value is 1.

User Response

The warning can be ignored if the **TARGET_WORKLOAD_PERCENT** for the LPAR is 1, or you must change the value in the LPAR parameter file.

CMA-OW005

Invalid **TARGET_WORKLOAD_PERCENT** of LPAR in the parameter file.

Explanation

This message is displayed when the **TARGET_WORKLOAD_PERCENT** value is invalid in the LPAR parameter file. The value should be between 0 and 1, with a default value of 1.

User Response

The warning can be ignored if the **TARGET_WORKLOAD_PERCENT** for the LPAR is 1, or you must change the value in the LPAR parameter file.

CMA-OW006

<lparName> in the LPAR parameter file cannot be found in the table <cmaSchema>.MVSPM_LPAR.

Explanation

This message is displayed because the LPAR <lparName> in the LPAR parameter file cannot be found in the table <cmaSchema>.MVSPM_LPAR for the given CPU_SERIAL_NO, PROCESSOR_TYPE, dates and time ranges. The LPAR is excluded.

User Response

Double check your data in the LPAR parameter file.

CMA-OW007

FILE_NAME_MODEL_PARA is missing from the parameter file.

Explanation

This message is displayed when **FILE_NAME_MODEL_PARA** is missing from the parameter file. The default value <fileNameModelPara> is assigned.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_DATE_PARA** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and file name **lwo_parameters_model.csv** will be used.

CMA-OW008

FILE_NAME_PARA_PER_LPAR is missing from the parameter file.

Explanation

This message is displayed when **FILE_NAME_PARA_PER_LPAR** is missing from the parameter file. The default value <fileNameParaPerLpar> is assigned.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_PARA_PER_LPAR** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and file name **lwo_parameters_per_lpar.csv** will be used.

CMA-OW009

FILE_NAME_DATE_LIST is missing from the parameter file.

Explanation

This message is displayed when **FILE_NAME_DATE_LIST** is missing from the parameter file. The default value <fileNameDateList> is assigned.

User Response

To resolve this error, provide the absolute path of the file in **FILE_NAME_DATE_LIST** or if no value is provided, the <default> location `CMA_HOME/optimizations/data` and file name **lwo_parameters_date_list.csv** will be used.

CMA-OW010

EXECUTION_NAME is missing from the parameter file.

Explanation

This message is displayed when **EXECUTION_NAME** is missing from the parameter file. The default value <executionName> is assigned.

User Response

The warning can be ignored if you want to use the **EXECUTION_NAME** by default, or enter the value in the parameter file.

CMA-OW011

Invalid value in the model parameter file.

Explanation

This message is displayed when the **SUB_OBJECTIVE_WEIGHT** <subObjectiveWeightValue> of the **SUB_OBJECTIVE_NAME** <subObjectiveName> is invalid in the model parameter file <fileNameModelPara>. The value is expected to be non-negative and a default value of 1 is assigned.

User Response

The warning can be ignored if the **SUB_OBJECTIVE_WEIGHT** for the **SUB_OBJECTIVE_NAME** is 1 or you must enter the correct value in the model parameter file.

CMA-OW012

Missing **SUB_OBJECTIVE_WEIGHT** value in the model parameter file.

Explanation

This message is displayed when the **SUB_OBJECTIVE_WEIGHT** <subObjectiveWeightValue> of the **SUB_OBJECTIVE_NAME** <subObjectiveName> is missing in the model parameter file <fileNameModelPara>. The default value of 1 is assigned.

User Response

The warning can be ignored if the **SUB_OBJECTIVE_WEIGHT** for the **SUB_OBJECTIVE_NAME** is 1 or you must enter the correct value in the model parameter file.

CMA-OW013

Missing **SUB_OBJECTIVE_NAME** in the model parameter file.

Explanation

This message is displayed when the **SUB_OBJECTIVE_NAME** is missing in the model parameter file <fileNameModelPara>. The default value of 1 is assigned to the missing objective.

User Response

The warning can be ignored if the **SUB_OBJECTIVE_WEIGHT** for the **SUB_OBJECTIVE_NAME** is 1 or you must enter the correct value in the model parameter file.

CMA-OW015

Invalid format for **DATE** in the date list parameter file.

Explanation

This message is displayed when the format for **DATE** <date> in the date list parameter file <fileNameDateList> is invalid. The date format should be `YYYY-MM-DD` or the date format of the locale. The date will be excluded.

User Response

Check the file and the input value of the locale.

CMA-OW016

The given double value does not match the decimal format of the locale.

Explanation

This message occurs when the given double value in the file <filename> does not match the decimal format of the locale, which can lead to assigning the incorrect value.

User Response

Check the file and the input value of the locale.

CMA-OW017

Records already exist in the Optimization Metadata table.

Explanation

This message is displayed as follows: **EXECUTION_NAME**: <executionName> **CPU_SERIAL_NO**: <cpuSerialNo> AND **PROCESSOR_TYPE**: <processorType> already has records in the <cmaSchema> .**OPTIMIZATION_METADATA** table. The old records will be replaced with a new record.

User Response

Use a different execution name.

CMA-OW018

Records already exist in the LPAR Weight Optimization Parameters table.

Explanation

This message is displayed as follows: **EXECUTION_NAME**: <executionName> **CPU_SERIAL_NO**: <cpuSerialNo> AND **PROCESSOR_TYPE**: <processorType> already has records in the <cmaSchema> .**LPAR_WEIGHT_OPTIMIZATION_PARAMETERS** table. The old records will be replaced with a new record.

User Response

Use a different execution name.

CMA-OW019

The captured seconds of the target workload calculated from the <TDSzSchema> table is larger than the LPAR busy seconds calculated from the <cmaSchema> table.

Explanation

The captured seconds of the target workload calculated from the <TDSzSchema> table is larger than the LPAR busy seconds (total workload) calculated from the <cmaSchema> table. To resolve this issue, check the data in both tables.

User Response

To resolve this issue, check the data in both tables.

CMA-OW020

MEASURED_SEC does not match between records in the <TDSzSchema>.MVSPM_WORKLOAD_H table and the <cmaSchema>.MVSPM_LPAR table.

Explanation

MEASURED_SEC does not match between the records in the <TDSzSchema>.MVSPM_WORKLOAD_H table and the <cmaSchema>.MVSPM_LPAR table. To resolve this issue, check the data in both tables.

User Response

To resolve this issue, check the data in both tables.

CMA-OW021

Records already exist in the LPAR Weight Optimization table.

Explanation

This message is displayed as follows: **EXECUTION_NAME**: <executionName> **CPU_SERIAL_NO**: <cpuSerialNo> AND **PROCESSOR_TYPE**: <processorType> already has records in the <cmaSchema> .**LPAR_WEIGHT_OPTIMIZATION** table. The old records will be replaced with a new record.

User Response

Use a different execution name.

CMA-OW022

Records already exist in the LPAR Weight Optimization Detail table.

Explanation

This message is displayed as follows: **EXECUTION_NAME**: <executionName> **CPU_SERIAL_NO**: <cpuSerialNo> AND **PROCESSOR_TYPE**: <processorType> already has records in the <cmaSchema> .**LPAR_WEIGHT_OPTIMIZATION_DETAIL** table. The old records will be replaced with a new record.

User Response

Use a different execution name.

CMA-OW023

Invalid **TARGET_WORKLOAD_IMPORTANCE_LEVEL** of LPAR in the parameter file.

Explanation

This message displays when there is an invalid **TARGET_WORKLOAD_IMPORTANCE_LEVEL** of LPAR <lparName> in the LPAR parameter file<fileNameParaPerLpar>. Ensure the value is set to null.

User Response

This warning can be ignored if the **TARGET_WORKLOAD_IMPORTANCE_LEVEL** for the LPAR is set to 1, or you can change the value in the LPAR parameter file.

CMA-OW024

Optimized weight for LPAR is less than 0.1%.

Explanation

This message displays when the optimized weight for LPAR is less than 0.1%. Check the value and validate if any adjustment is required.

User Response

Check the data and input values.

CMA-OW025

<lparName> is missing in the LPAR parameter file.

Explanation

This message displays when <lparName> is missing in the LPAR parameter file <fileName>. The default setting is assigned to the LPAR.

User Response

If it is okay to assign default values to the LPAR, you can ignore this warning, or you can add the LPAR and parameters to the LPAR parameter file.

CMA-OW026

<lparName> does not have any active logical processors for any interval.

Explanation

This message displays when <lparName> does not have any active logical processors for any interval (**LP_ONLINE_SEC** is always 0). The LPAR will be excluded.

User Response

Check the data and input values.

Installation error, information, and warning messages

When installing IBM Capacity Management Analytics, it's important to understand what error messages, warnings, and information messages you may encounter.

Installation error messages

When installing IBM Capacity Management Analytics, it's important to understand what error messages, warnings, and information messages you may encounter.

CMA-E-002
Unable to load JSON file <jsonfile>.

Explanation

Installation error: cannot load the JSON <jsonfile> file.

User Response

Check if the file exists and isn't corrupted and the format is correct.

CMA-E-003
Cannot create Python object <object>.

Explanation

Installation error: cannot create Python object <object>.

User Response

Check if PYTHONPATH is set correctly in the user profile, and ensure the program file exists.

CMA-E-004
Unable to load parameter file <file>.

Explanation

Installation error: cannot load the parameter file <file>.

User Response

Check if the file exists and isn't corrupted and the format is correct.

CMA-E-005
Cannot find file <file>.

Explanation

Installation error: cannot find the file <file>.

User Response

Check if the file exists and isn't corrupted and the format is correct.

CMA-E-007
Environment variable CMAHOME not set.

Explanation

Installation error: the environment variable CMAHOME is not set.

User Response

Check the CMAHOME setup in the user profile.

CMA-E-008
More than one file, <file>, was found.

Explanation

Installation error: more than one file, <file>, was found. Cannot differentiate the files.

CMA-E-009
Cannot create directory: <dir>.

Explanation

Installation error: cannot create directory: <dir>.

User Response

Check if the parent folder exists and ensure the user has the correct permission.

CMA-E-010
There are no feature tarballs detected.

Explanation

Installation error: there are no feature tarballs detected.

User Response

Check if the feature tarballs are in the correct directory: `cmahome/zipfiles` and also ensure that the user has the correct permission to the tarballs.

CMA-E-011
Environment variable CMAHOME: <dir> is not a directory.

Explanation

Installation error: environment variable CMAHOME: <dir> is not a directory.

User Response

Check if CMAHOME is set correctly in the profile and ensure the director exists.

CMA-E-012
Cannot find a support version of CPLEX Optimization Studio.

Explanation

Installation error: cannot find a support version of CPLEX Optimization Studio.

User Response

Refer to the *Solution Guide* to ensure you have the correct version installed in the correct location.

CMA-E-013
Version of Python is not supported.

Explanation

Installation error: version of Python is not supported.

User Response

Refer to the *Solution Guide* for a supported version of Python.

CMA-E-014
Could not detect `cmainstance-config.js`. The CMAINSTANCE specified is not valid.

Explanation

Installation error: could not detect `cmainstance-config.js`. The CMAINSTANCE specified is not valid.

User Response

Check the CMAINSTANCE setting in the user profile. Check if the `<cmainstance>` directory exists and ensure the user has the correct permission. Check if `cmainstance-config.js` exists in the `<cmainstance>/configuration` directory, that it isn't corrupted, and the file format is correct.

CMA-E-016
JSON file: <file> content is invalid.

Explanation

Installation error: JSON file: <file> content is invalid.

User Response

Check if the file is corrupt, ensure the file is in the correct format, and double check the user's permissions.

CMA-E-018
No CMA features are installed. Cannot create Instance.

Explanation

Installation error: no CMA features are installed. Cannot create Instance.

User Response

Ensure that CMAHOME and CMAINSTANCE are set correctly. Also check if the files and folders exist. Lastly, check if the user has correct permissions.

CMA-E-019
Feature: <feature> was not activated in CMAINSTANCE: <instance> because <feature> was not detected as installed.

Explanation

Installation error: feature: feature: <feature> was not activated in CMAINSTANCE: <instance> because <feature> was not detected as installed.

User Response

Check if the feature tarballs are in the correct directory: `cmahome/zipfiles` and also ensure that the user has the correct permission to the tarballs.

CMA-E-020
Invalid CMAHOME detected at `<dir>`. You may need to recreate the CMAHOME directory.

Explanation

Installation error: invalid CMAHOME detected at `<dir>`. You may need to recreate the CMAHOME directory.

User Response

Create CMAHOME at `<dir>` and ensure that the user has the correct permission to access CMAHOME.

CMA-E-021
Invalid function: <function>.

Explanation

Installation error: invalid function: <function>.

CMA-E-022
Untarring file: <file> failed.

Explanation

Installation error: untarring file: <file> failed.

User Response

Check if the file exists and ensure the user has the correct permissions to the file.

CMA-E-023
Unable to load Python file: <file>.

Explanation

Installation error: unable to load Python file: <file>.

User Response

Check if the file exists and ensure the user has the correct permissions to the file.

CMA-E-024
Cognos Namespace was not specified. Set the Cognos Namespace before proceeding.

Explanation

Installation error: Cognos Namespace was not specified. Set the Cognos Namespace before proceeding.

User Response

Set the Cognos Namespace.

CMA-E-025
Environment variable <envar> not set.

Explanation

Installation error: Environment variable <envar> not set.

User Response

Set <envar> in the user profile.

CMA-E-026
Directory: <dir> is invalid or doesn't exist.

Explanation

Installation error: Directory: <dir> is invalid or doesn't exist.

User Response

Check if the <dir> exists and ensure that the user has the correct permissions to the <dir>.

CMA-E-027
Failed to build the CMA images application war file.

Explanation

Installation error: failed to build the CMA images application war file.

CMA-E-028
Failed to execute Python command: `<cmd>`.

Explanation

Installation error: failed to execute Python command: `<cmd>`.

CMA-E-029
Failed to create directory(ies): <directory>.

Explanation

Installation error: failed to create directory(ies): <directory>.

User Response

Check to see if the user has the correct permissions.

CMA-E-030
External command: <cmd> failed with return code <retcode>.

Explanation

Installation error: external command: <cmd> failed with return code <retcode>.

User Response

Check to see if the user has the correct permissions and check if the username and password are correct.

CMA-E-031
Logger has not been defined.

Explanation

Installation error: logger has not been defined.

User Response

Check if your environment variables are set and if you have the correct permission to write to the log.

CMA-E-032
Logger has not been defined.

Explanation

Installation error: logger has not been defined.

User Response

CMA-E-033
Logger has not been defined.

Explanation

Installation error: logger has not been defined.

User Response

CMA-E-034
HOST <host> is invalid: Name or service not known.

Explanation

Installation error: HOST <host> is invalid: Name or service not known.

User Response

Check if your input host is valid.

CMA-E-035
Unable to get group name of user <user>.

Explanation

Installation error: unable to get group name of user <user>

User Response

Ensure your ID has been set to cmagroup. Refer to the cmagroup section for more detailed information.

CMA-E-040
Parameter <parameter> doesn't exist or invalid.

Explanation

Installation error: parameter <parameter> doesn't exist or invalid.

User Response

Ensure your command line is valid. Use `program.py -h` to get more information about parameters.

CMA-E-042
Fail to <operation>.

Explanation

Installation error: fail to <operation>.

User Response

Check if your environment variables are set correctly and ensure you have valid input.

Installation information messages

When installing IBM Capacity Management Analytics, it's important to understand what error messages, warnings, and information messages you may encounter.

CMA-I-001
Untarring <file>.

Explanation

Installation information: untarring <file>.

CMA-I-002
Installing feature: <feature>.

Explanation

Installation information: installing feature: <feature>.

CMA-I-003
<operation> of CMAINSTANCE: <path> started.

Explanation

Installation information: <operation> of CMAINSTANCE: <path> started.

CMA-I-004
<operation> of CMAINSTANCE: <`path`> finished.

Explanation

Installation information: <operation> of CMAINSTANCE: <`path`> finished.

CMA-I-005
<operation> of CMAHOME: <path> started.

Explanation

Installation information: <operation> of CMAHOME: <path> started.

CMA-I-006
<operation> of CMAHOME: <path> finished.

Explanation

Installation information: <operation> of CMAHOME: <path> finished.

CMA-I-007
`--run` was not specified. No actions taken.

Explanation

Installation information: `--run` was not specified. No actions taken.

User Response

Rerun the program and append `--run` on command to activate actions.

CMA-I-008
Check log <log> for more information.

Explanation

Installation information: check log <log> for more information.

User Response

Check the log for more information.

CMA-I-009
<operation> of <target>: <path> started.

Explanation

Installation information: <operation> of <target>: <path> started.

CMA-I-010
<operation> of <target>: <path> finished.

Explanation

Installation information: <operation> of <target>: <path> finished.

CMA-I-011
<operation> started.

Explanation

Installation information: <operation> started.

CMA-I-012
<operation> finished.

Explanation

Installation information: <operation> finished.

CMA-I-013
No features were installed.

Explanation

Installation information: no features were installed.

CMA-I-014
No features were instantiated.

Explanation

Installation information: no features were instantiated.

Installation warning messages

When installing IBM Capacity Management Analytics, it's important to understand what error messages, warnings, and information messages you may encounter.

CMA-W-001
Environment variable CMAINSTANCE not set.

Explanation

Installation warning: environment variable CMAINSTANCE not set.

User Response

Check if the CMAINSTANCE is set in the user profile.

CMA-W-002
Environment variable CMAHOME not set.

Explanation

Installation warning: environment variable CMAHOME not set.

User Response

Check if the CMAHOME is set in the user profile.

CMA-W-003
Feature: <feature> was not activated in CMAINSTANCE: <instance> because <feature> was not detected as installed.

Explanation

Installation warning: feature: <feature> was not activated in CMAINSTANCE: <instance> because <feature> was not detected as installed.

User Response

1. Check if the feature tar file exists in `<home>/zipfiles` and isn't corrupted.

2. Check if the feature is installed in the `<home>/features` directory.
3. Check the user permissions for the `<home>` and `<instance>` directories.

CMA-W-004
Feature: <feature> was not installed in CMAHOME: <home> because <feature> was not detected in the `<home>/zipfiles` directory.

Explanation

Installation warning: feature: <feature> was not installed in CMAHOME: <home> because <feature> was not detected in the `<home>/zipfiles` directory.

User Response

Check if the feature tar file exists in the `<home>/zipfiles` directory and that it isn't corrupted.

CMA-W-005
Fail to get system locale, use default locale en_US.

Explanation

Installation warning: fail to get system locale, use default locale en_US.

User Response

Check the system locale setting.

CMA-W-006
Fail to get delimiter from csv file <csvfile> or no csv file to read, use default delimiter, .

Explanation

Installation warning: fail to get delimiter from csv file <csvfile> or no csv file to read, use default delimiter, .

User Response

Check if the file exists, that it isn't corrupted, and the format is correct.

CMA-W-007
External command: `<cmd>` finished with warning code <retcode> .

Explanation

Installation warning: external command: `<cmd>` finished with warning code <retcode>.

User Response

The command may have failed. Check the warning code for more detailed information.

CMA-W-008
External command: `<cmd>` only can run with sudo on Linux Red Hat 6.

Explanation

Installation warning: external command: `<cmd>` only can run with sudo on Linux Red Hat 6.

CMA-W-009
External command: Cognos configuration has been updated. Restart Cognos.

Explanation

Installation warning: external command: Cognos configuration has been updated. Please restart Cognos.

User Response

Restart Cognos.

CMA-V-001
No CMA features were instantiated.

Explanation

Installation verbose: no CMA features were instantiated.

User Response

Check if you have valid inputs.

CMA-V-002
Update/create file: <file>.

Explanation

Installation verbose: update/create file: <file>.

CMA-V-003
Copy file <file> to directory: `<dir>`.

Explanation

Installation verbose: copy file <file> to directory: `<dir>`.

CMA-V-004
Reinstate file <file>.

Explanation

Installation verbose: reinstate file <file>.

CMA-V-005
<operation> started.

Explanation

Installation verbose: <operation> started.

CMA-V-006
<operation> finished.

Explanation

Installation verbose: <operation> finished.

CMA-V-007
Backup original file <file>.

Explanation

Installation verbose: backup original file <file>.

CMA-V-008
Saving configuration file: <file>.

Explanation

Installation verbose: saving configuration file: <file>.

CMA-V-009
Updating configuration data.

Explanation

Installation verbose: updating configuration data.

CMA-V-010
Getting environment variable: <envar>.

Explanation

Installation verbose: getting environment variable: <envar>.

CMA-V-011
Name match: <match>.

Explanation

Installation verbose: name match: <match>.

CMA-V-012
Run external command: <cmd>.

Explanation

Installation verbose: run external command: <cmd>

CMA-V-013
Run python command: <cmd>.

Explanation

Installation verbose: name match: run python command: <cmd>.

CMA-V-014
Preview command: <cmd>.

Explanation

Installation verbose: name match: preview command: <cmd>.

Appendix G. IBM Capacity Management Analytics terminology

This glossary includes terms and definitions for IBM Capacity Management Analytics.

AA

Application Analytics

This feature uses a user supplied application mapping table to provide reports by application, function, and environment.

APPL

Application

CMA

Capacity Management Analytics

CPC

Processor and CPU can refer to either the complete system box, or to one of the processors (CPUs) within the system box. Although the meaning might be clear from the context, even mainframe professionals must clarify which processor or CPU meaning they are using. IBM uses the term central processor complex (CPC) to refer to the physical collection of hardware that includes main storage, one or more central processors, timers, and channels. (Some system programmers use the term central electronic complex (CEC) to refer to the mainframe box, but the preferred term is CPC.)

CPLEX

IBM ILOG CPLEX Optimization Studio is an optimization software package.

DRL

DRL is the default three letter identifier for Tivoli Decision Support for z/OS. It is used to reference various objects such as Framework Manager items and database objects.

GSSP

IBM Getting Started Sub-Capacity Pricing (GSSP) for z/OS offers entry pricing for customers with select z/OS International Program License Agreement (IPLA) programs and the applicable Subscription and Support annual maintenance charge.

Getting Started Sub-capacity Pricing helps you:

- Deploy small projects on z/OS with improved price/performance

- Provide greater software licensing flexibility by improving alignment between software usage and software charges
- Increase LPAR capacity without changing software charges for a constant workload
- Transition smoothly to traditional sub-capacity pricing as workload grows

HCM

HCM is the default three letter identifier for CMA. It is used to reference various objects such as Framework Manager items and database objects.

HIS

Hardware Instrumentation Services (HIS)—data collection of CPU metrics written to SMF 113 record.

IPLA

Related to software pricing—IPLA programs have a one-time-charge (OTC) and an (optional) annual maintenance charge, called Subscription & Support. For more information, see the IBM z Systems Software Pricing page (www.ibm.com/systems/z/resources/swprice/zipla/index.html).

IWP

IBM provides Integrated Workload Pricing (IWP) to enhance utilization reporting capabilities of systems that are running IBM z/OS, which can improve the price/performance for eligible Monthly License Charge (MLC) programs that are running within the same logical partition (LPAR) as select defining programs. These programs are listed in the Description section, Integrated Workload Pricing Defining Programs. The IWP enhancements to the reporting capabilities in z/OS enable the IBM Sub-capacity Reporting Tool (SCRT) to calculate the general-purpose processor time that is consumed by eligible defining programs and adjust the reported MSU values for eligible MLC programs that are running in the same LPAR for eligible Monthly License Charge IWP Adjusted Programs running within the same LPAR as select IWP Defining Programs.

LPAR

A Logical Partition on a physical server. Essentially LPAR is a virtual server.

LWO

LPAR Weight Optimization—this feature has reports as well as a CPLEX model that optimizes the LPAR weight value to achieve fitting Target workload within the weight amount of MIPS.

MIPS

We use MIPS to represent the capacity of the zSeries frame or LPAR. Early metrics tended to concentrate on the rate at which a processor executes instructions to represent capacity. One metric of this type is MIPS (millions of instructions per processor second). Although MIPS is used as the term, the measure is not actually millions of instructions per processor second. The value is derived from numbers published by IBM in the Large System Performance Reference (LSPR) table. The default is to use the single image

Processor Capacity Index (PCI) value with the average Relative Nest Intensity (RNI) for the hardware. For an LPAR, we divide the total MIPS capacity by the number of physical processors and then multiply by the number of processors defined to the LPAR. You can supply a MIPS table of your own, which might use the results from a zPCR study, which would be more accurate because it takes into account the workload characteristics.

zPCR is a PC-based productivity tool under Windows. It is designed to provide capacity planning insight for IBM System z processors that are running various z/OS, z/VM®, z/VSE®, Linux, IBM zAware, and CFCC workload environments on partitioned hardware. Capacity results are based on IBM's most recently published LSPR data for z/OS.

For more information, see the Large Systems Performance Reference for IBM z Systems page (https://www.ibm.com/servers/resourcelink/lib03060.nsf/pages/lsprindex?OpenDocument).

MLC

Monthly License Charge.

For more information, see the IBM z Systems Software Pricing page (www.ibm.com/systems/z/resources/swprice/mlc/index.html).

MSU

A million service units (MSU) are a measurement of the amount of processing work a computer can perform in one hour. MSU ratings are used for Software Pricing charges.

NO89

Refers to software products that are priced using a sub-capacity pricing method. NO89 products do not generate SMF type 89 records, which report on usage by the product over a time interval.

PEAKRANK

In the Application peak value forecast stream, PEAKRANK is an input parameter that specifies which peak values you want the forecast to support.

PWD

Password

SCA

Software Cost Analysis.

This feature reports and forecasts software cost based on Sub-Capacity Pricing for z/OS.

UID

User ID.

Appendix H. Troubleshooting

These topics can help you troubleshoot problems and errors you encounter in IBM Capacity Management Analytics.

General troubleshooting

This information can help you troubleshoot general problems you encounter while using IBM Capacity Management Analytics.

Handling duplicate records

Before executing a Tivoli Decision Support for z/OS collect job, you should sort Service Management Facility (SMF) records to avoid duplicate records. For more information, see Tivoli Decision Support for z/OS General best practices (https://www.ibm.com/developerworks/community/wikis/home/wiki/Tivoli+Decision+Support+for+zOS/page/General+best+practices).

If you see duplicate record in your software cost analysis reports, please remove duplicate records in the MVSPM_PROD_T, MVSPM_PROD_INT_T and MVSPM_LPAR_MSU_T tables, sort the SMF records, then again execute a Tivoli Decision Support for z/OS collect job.

Maintaining date format consistency

You must ensure that date formats are consistent between SPSS Modeler streams and input files, such as `holiday.csv`.

For example, if you use Microsoft Excel to edit csv files and your computer uses a different locale than what is used in the SPSS Modeler streams, you must ensure that the data format in the csv remains consistent with date format that is used in the streams. Microsoft Excel can change the locale that is used in the file to match your computer's locale.

Tune IBM Cognos Business Intelligence for Capacity Management Analytics

Use the following information to help you tune your IBM Cognos Business Intelligence report performance for IBM Capacity Management Analytics.

The following factors can influence IBM Cognos report performance:

- Peak system usage periods
- Number of concurrent users the system can support
- Number of concurrent users expected
- Acceptable response times for the system
- Size of the data that is being transferred and the processing capacity of the database server

You can also use cached data when you run report to performance improvements.

For more performance tuning information, see the IBM Cognos Business Intelligence Server documentation on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SSEP7J_10.2.2/com.ibm.swg.ba.cognos.crn_arch.10.2.2.doc/c_arch_performancetuning.html%23arch_PerformanceTuning)

Tune IBM SPSS Modeler for Capacity Management Analytics

IBM Capacity Management Analytics performance can be improved by ensuring that SPSS Modeler is performing optimally.

For more performance tuning information, see the IBM SPSS Modeler documentation on IBM Knowledge Center (www.ibm.com/support/knowledgecenter/SS3RA7_17.0.0/clementine/server/adminguidesource/admin_perf_optimize.dita).

Error: AEQMJ0100E: Script error (Array index out of range: 1) on line 252 column 1

This error occurs when you are running a stream but have a mismatch between the ENCODED parameter and whether you entered an encoded password for the PWD value.

For example, if you set the ENCODED value in a `parameter.txt` file to true, but you did not enter an encoded password in PWD, then this error occurs.

SMO stream troubleshooting

This information can help you troubleshoot SMO stream problems and errors you encounter while using IBM Capacity Management Analytics.

You can find error log files in the log folder that you specify with the LOGPATH parameter or CMAINSTANCE/data. The log file names consist of the stream name and a timestamp.

Troubleshooting distribution_procedure.str

- The ETL stream `distribution_procedure.str` reports the following error:

  `AEQMJ0142E: cannot concatenate 'str' and {1} objects`

  You can resolve this in IBM SPSS Modeler Client by clicking **Tools** > **Stream Properties** > **Parameters**, and changing the storage of stream parameters start_date and end_date from Date to String.

SCA stream troubleshooting

This information can help you troubleshoot SCA stream problems and errors you encounter while using IBM Capacity Management Analytics.

Troubleshooting SCA forecast streams

You can find the following error log files for the SCA forecast streams in the log folder that you specify with the LOGPATH parameter.

- `msu_prod_forecast_timeseries_<timestamp>.log` for the product forecast stream.
- `msu_gssp_forecast_timeseries_<timestamp>.log` for the GSSP forecast stream.

The log files may contain the following errors:

not enough data error
: If the product in one LPAR of one CPC doesn't have 24 hours of data in any day of the week, the log file will contain the message `Do not have 24 hours data in one day`. You can enter the missing values according to your SMF data in the database to correct this error.

null value error
: If the product in one LPAR of one CPC contains only null values in any day of week, the log file will contain the message `Contained null value`. You can fill in the null values in the database to correct this error.

not enough data and null value messag
: If the product in one LPAR of one CPC contains only null values in any day of week **and** doesn't have 24 hours of data in any day of the week, the log file will contain with message `Contained null`

value and do not have 24 hours data in one day. You can fill in the null values and enter the missing values in the database to correct this error.

Troubleshooting SCA optimization streams

For any of the following situations, there should be one or more error messages in the log file named prod_bill_msu_optimization_*xxx*.log, where *xxx* is the timestamp when the log is created:

- If you cannot get an optimized result.
- If the stream is interrupted.

The location of the log is specified by the LOGPATH parameter. If you run the stream in batch mode, LOGPATH is set in the file named parameter.txt. If you run the stream in client mode, LOGPATH is set in the parameter input window.

For more information about the error messages, see Appendix F, “Error, warning, and installation messages,” on page 315.

In addition to the previous errors, in some cases you might receive an optimized result containing the following errors:

Didn’t find the parent program in the LPAR where child program locates after optimization
: Confirm that the program is configured in the bounded_prod file in opt_bound_prod.csv. This can be found in the Configuration chapter.

The program is reserved in a certain LPAR, but it is located in other LPARs after the optimization
: There are two possible causes for this error:

  - The LPAR that the program reserved may not suitable, already out of capacity, or be configured with a very small capacity definition. You should try to increase the capacity for this LPAR or change to another LPAR to reserve this product.
  - The program is a child program and set to be in a bound group, and it also set to be reserved in an LPAR, but the parent program is not reserved. In this case you should also reserve the parent program in that LPAR.

The program is reserved in a certain LPAR, but it is still in the original LPAR after the optimization
: Try to decrease the reserved amount. If the reserved amount is greater than the actual MSU load of the program, the optimization ignores this reservation and keep the program in its original location.

Troubleshooting sca_no89.str

- If this stream returns a no data source error, ensure that the NO89 file is in the location specified by the NO89_FILE parameter. Additionally, ensure that the NO89 file contains the correct start and end strings. The file should start with //NO89 DD * NON-89 PRODUCT SECTION and end with // SYSLIN DD *.
- If the stream fails with a merge error, verify that the HCM.MVSPM_LPAR and HCM. SUBCAP_PROGRAMS HCM tables contain all installed product information and LPAR information. To insert correct data into these tables, execute the mvspm_lpar_insert.str, mvspm_lpar_update.str, and sca_subcap.str streams in sequence.
- If this stream returns a table not found error, verify that the HCM.NO89_PRODUCTS table exists in your HCM database.

Troubleshooting sca_pricing.str

- If this stream returns a no data source error, ensure that the mlc_prod.csv, ipla_prod.csv, enterprise.csv, and currency.csv files are in the CMAINSTANCE/data directory.
- If the stream fails with a merge error, verify that the SUBCAP_PROGRAMS, MSU_TIERS and VUE_TIERS HCM tables contain all installed product information, MSU tier and VU tier. To insert correct data into these tables, execute the sca_subcap.str and sca_tiers.str streams.

- If this stream returns a `table not found` error, verify that the `HCM.CUSTOMER_PRICE` table exists in your HCM database.

Troubleshooting sca_tiers.str

- If this stream returns a `no data source` error, ensure that the `u_tiers.csv` and `msu_tiers.csv` files are in the `CMAINSTANCE/data` directory.
- If this stream returns a `table not found` error, verify that the `HCM.VUE_TIERS` and `HCM.MSU_TIERS` tables exists in your HCM database.

Troubleshooting sca_subcap.str

- If this stream returns a `no data source` error, ensure that the `gssp_scaling_factors.csv`, `gssp.csv`, `iwpadjust.csv`, `iwpdefine.csv`, `mlc.csv`, `parent.csv`, , `referencebased.csv`, and `zos_based.csv` files are in the *CMA_ROOT*`/streams/data` directory.
- If this stream returns a `table not found` error, verify that the `HCM.SUBCAP_PROGRAMS` table exists in your HCM database.

Anomaly Detection troubleshooting

This information can help you troubleshoot Anomaly Detection stream problems and errors you encounter while using IBM Capacity Management Analytics.

Troubleshooting the model building stream

You can find the Modeler stream log files in the log folder that you specify with the LOGPATH parameter.

- `anomaly_detect_building_c_<timestamp>.log` for the response stream.
- `anomaly_detect_building_r_<timestamp>.log` for the CPU stream.

The log files may contain the following errors:

Not enough data error
Model building streams generate errors when there is not enough data. If you have a log entry such as `XXXX doesn't have enough data for model building`, where XXXX is the TRANSACTION_ID, you need to check the data for that specified transaction to determine how many records exist between the start and end date you specified for model building. Capacity Management Analytics requires at least 40 records for model building.

Large data exceed memory error
Model building streams generate errors when the amount of data used for model building is larger than the memory capacity of your machine. In this case, you will see a log entry such as `XXXX has exceed data limitation for model building, which is xxxxx. Please sample your data before you build the model`. In this log entry, XXXX is TRANSACTION_ID, and xxxxx is the record number limitation for model building. There are two ways of solving this problem: you can either sample your data before you build the model or reduce the range for your model building input data.

Troubleshooting the model scoring stream

This stream writes errors to the `anomaly_detect_scoring_<timestamp>.log`.

Not enough data error:
The model scoring stream generates errors when your input data for model scoring is empty. In this case, you will see a log entry such as `XXXX doesn't have enough data for model scoring`, where XXXX is TRANSACTION_ID. You should ensure that the date set for scoring in `date_info.csv` and `transaction.csv` are valid, and that your database contains the data want to score.

Troubleshooting the classification stream

This stream writes errors to the `anomaly_classification_<timestamp>.log`.

No scoring result for both CPU and RESPONSE model

The classification stream generates errors when you don't have any scoring result from both the CPU model and RESPONSE model in database. In this case, you will see a log entry such as `Didn't find any cpu and response scoring results in database, please run cpu and response scoring before executing the classification stream.` To resolve this issue, must score both the CPU and REPONSE models, and store the results in the database.

No scoring result for CPU model

The classification stream generates errors when you don't have any scoring result from the CPU model. In this case, you will see a log entry such as `Didn't find any cpu scoring results in database, please run cpu scoring before executing the classification stream.` To resolve this issue, you need to score the CPU model and store the result to the database before executing the Classification stream.

No scoring result for RESPONSE model

The classification stream generates errors when you don't have any scoring result from the RESPONSE model. In this case, you will see a log entry such as `Didn't find any response scoring results in database, please run response scoring before execute classification stream.` To resolve this issue, you need to score the RESPONSE model and store the result to the database before executing Classification stream

CPU result and RESPONSE result doesn't match each other

The classification stream generates errors when CPU and RESPONSE don't have any intersection data. In this case, you will see a log entry such as `CPU and RESPONSE scoring results don't have intersection data to classify results.` .To resolve this issue, you need to make sure the CPU and RESPONSE results in the database are consistent with each other. For any given transaction, both CPU and RESPONSE scoring results must exist.

Solution Installer troubleshooting

This information can help you troubleshoot the Solution Installer.

About this task

To update the database alias, you must first manually remove the existing one.

Procedure

1. Uncatalog DB2 alias:
   a. uncatalog db2 node <CMADBALIAS>
   b. uncatalog db2 database <CMADBALIAS>
2. Clean up the CMADBALIAS value in the `DB2HOME/cfg/db2cli.ini` file, remove the section for the CMADBALIAS.

   For example:

   ```
   [CMA_CMAD]

   Hostname=myserver.mydomain.com

   Port=1234

   Protocol=TCPIP

   DriverUnicodeType=1

   Database=CMAD
   ```
3. Create the alias with the correct information by executing cmainstance.py or cmahome.py.

For more information about cmainstance.py and cmahome.py, see Chapter 6, “Manually installing IBM Capacity Management Analytics 2.1.1,” on page 59.

Notices

This information was developed for products and services offered worldwide.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
Attention: Licensing

3755 Riverside Dr.
Ottawa, ON
K1V 1B7
Canada

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

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

Trademarks

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

The following terms are trademarks or registered trademarks of other companies:

- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.