

IBM Z Performance and Capacity Analytics
Version 3 Release 1

*Capacity Planning
Guide and Reference*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 99.](#)

This edition applies to version 3, release 1 of IBM Z Performance and Capacity Analytics (program number 5698-AS3) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: March 2020

© **Teracloud S.A. 2018, 2020.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures.....	vii
Preface.....	ix
Who should read this book.....	ix
Accessing publications online.....	ix
Accessibility.....	ix
Support information.....	ix
Conventions used in this book.....	ix
Typeface conventions.....	x
What's new in this edition.....	x
Chapter 1. Introduction.....	1
Architecture.....	1
Subcomponents.....	2
Data tables.....	4
Lookup tables and control tables.....	5
Reports.....	5
Chapter 2. Implementing IBM Z Performance and Capacity Analytics.....	7
Step 1: Installing and configuring the Capacity Planning for z/OS component.....	7
Step 1.1: Deciding which subcomponents to install.....	7
Step 1.2: Making input data available.....	7
Step 1.3: Configuring lookup tables.....	8
Step 1.4: Configuring control tables.....	11
Step 2: Installing and configuring the Forecaster module.....	12
Step 2.1: Unpacking the Forecaster module.....	12
Step 2.2: Updating the table mapping files.....	13
Step 2.3: Configuring the Forecaster module.....	13
Chapter 3. Implementing Cognos reporting (optional).....	15
Chapter 4. Implementing ELK reporting (optional).....	17
Step 1: Setting up data streaming through IBM Common Data Provider for z.....	17
Step 2: Installing and configuring IBM Common Data Provider for z.....	18
Step 3: Configuring IBM Z Performance and Capacity Analytics for data streaming.....	18
Step 4: Configuring the Forecaster for data streaming.....	19
Step 5: Configuring IBM Common Data Provider for z for data streaming.....	22
Step 6: Configuring ELK.....	23
Step 6.1: Download and install the dashboard files.....	23
Step 6.2: Configuring Logstash.....	24
Step 6.3: Configuring Kibana.....	24
Step 7: Viewing the dashboard.....	26
Step 8: Ad hoc reporting.....	26
Chapter 5. Capacity planning.....	27
Introduction to capacity planning.....	27
Getting started with IBM Z Performance and Capacity Analytics.....	28
Chapter 6. Forecasting.....	31

Supplied forecasts.....	33
Tailoring a supplied forecast.....	35

Chapter 7. Data table reference..... 37

CPU subcomponent data tables.....	37
CP_CPU_LPAR_T.....	37
CP_CPU_LPAR_H.....	39
CP_CPU_LPAR_D.....	40
CP_CPU_LPAR_W.....	42
CP_CPU_LPAR_F.....	43
CP_CPU_LPAR_FH.....	44
CP_CPU_LPAR_PH.....	45
CP_LSPR_WKLD_H.....	47
CP_LSPR_WKLD_D.....	48
CP_LSPR_WKLD_W.....	48
CP_CPU_WKLD_H.....	49
CP_CPU_WKLD_D.....	50
CP_CPU_WKLD_W.....	52
CP_CPU_WKLD_F.....	53
Storage subcomponent data tables.....	54
CP_STORAGE_H.....	54
CP_STORAGE_D.....	55
CP_STORAGE_W.....	56
CP_STORAGE_F.....	57
Channel subcomponent data tables.....	58
CP_CHANNEL_H.....	58
CP_CHANNEL_D.....	59
CP_CHANNEL_W.....	60
Disk subcomponent data tables.....	61
DFSMS_VOLUME_D.....	61
DFSMS_VOLUME_M.....	62
CP_DISK_F.....	64
Tape subcomponent data tables.....	65
CP_TAPE_D.....	65
CP_TAPE_W.....	65
CP_TAPE_F.....	66

Chapter 8. Report reference..... 67

CPU subcomponent reports.....	67
CPU usage.....	68
CPU forecasting.....	76
LSPR workload analysis.....	81
Software cost analysis.....	82
Storage subcomponent reports.....	87
Storage usage.....	87
Storage forecasting.....	89
Channel subcomponent reports.....	89
Channel usage.....	89
Channel forecasting.....	90
Disk subcomponent reports.....	91
Disk usage.....	91
Disk forecasting.....	92
Tape subcomponent reports.....	93
Tape usage.....	94
Tape forecasting.....	95

Appendix A. Support information..... 97

Contacting IBM Support.....	97
Notices.....	99
Trademarks.....	100
Index.....	103

Figures

1. IBM Z Performance and Capacity Analytics Architecture.....	2
2. Example: CP_THRESHOLDS lookup table.....	9
3. Example: CP_SHIFT lookup table.....	10
4. Example: CP_BUSINESS_APPL lookup table.....	11
5. Streaming data to ELK by the Collector.....	17
6. Streaming data to ELK by the Forecaster.....	17
7. Defining an index pattern for retrieving data.....	25
8. CP CEC MIPS Weekly report.....	68
9. CP LPAR MIPS Weekly report.....	69
10. CP LPAR MIPS by CEC Average Daily report.....	70
11. CP LPAR MIPS by CEC Peak Weekly report.....	71
12. CP Workload MIPS Usage Daily report.....	72
13. CP Serv Class MIPS Daily report.....	73
14. CP Serv Class MIPS by Workload Avg Daily report.....	74
15. CP Serv Class MIPS by Workload Peak Weekly report.....	75
16. CP Application MIPS Daily report.....	76
17. CP CEC MIPS Forecast Weekly report.....	77
18. CP LPAR MIPS Forecast Weekly report.....	78
19. CP Workload MIPS Forecast Weekly report.....	79
20. CP Serv Class MIPS Forecast Weekly report.....	80
21. CP Application MIPS Forecast Weekly reports.....	81
22. CP LSPR Workload Analysis Hourly report.....	82
23. CP LPAR 4 Hour MSU Utilization Hourly report.....	83

24. CP LPAR MSU Statistics Hourly report.....	84
25. CP LPAR Product Matrix Daily report.....	85
26. CP Product LPAR Matrix Daily report.....	86
27. CP LPAR Product MSU Hourly report.....	87
28. CP Central Storage Available Hourly report.....	88
29. CP Channel Peak Busy Hourly report.....	90
30. CP Channel Peak Busy Forecast Weekly report.....	91
31. CP Disk Storage Group Allocated Space Monthly report.....	92
32. CP Disk Storage Group Space Forecast Monthly report.....	93
33. CP Tape Storage Group Space Weekly report.....	94
34. CP Tape Storage Group Space Forecast Weekly report.....	95

Preface

This book provides information about the capacity planning and forecasting feature of IBM Z Performance and Capacity Analytics.

You should use this book in conjunction with the documentation for IBM Z Performance and Capacity Analytics .

Who should read this book

IBM Z Performance and Capacity Analytics, Capacity Planning Guide and Reference is intended for those who are responsible for monitoring system capacity and ensuring that sufficient capacity is available.

Accessing publications online

Publications for this and all other IBM products, as they become available and whenever they are updated, can be viewed on the IBM Knowledge Center website where you can also download the associated pdf.

<https://www.ibm.com/support/knowledgecenter/SSPNK7>

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

For additional information, see the Accessibility Appendix in the *Administration Guide and Reference*.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides the following ways for you to obtain the support you need:

- Searching knowledge bases: You can search across a large collection of known problems and workarounds, Technotes, and other information.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Software Support.

For more information about these three ways of resolving problems, see [Appendix A, “Support information,”](#) on page 97.

Conventions used in this book

This guide uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

The following terms are used interchangeably throughout this book:

- MVS, OS/390, and z/OS.

- OPC, OPC/ESA, Tivoli Workload Scheduler for z/OS, and TWS.
- TCP/IP and TCP/IP for z/OS.
- VM and z/VM.
- WebSphere MQ for z/OS and MQSeries.

Typeface conventions

This guide uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip**, and **Operating system considerations**)
- Column headings in a table
- Keywords and parameters in text

Italic

- Citations (titles of books, diskettes, and CDs)
- Words defined in text
- Emphasis of words (words as words)
- Letters as letters
- New terms in text (except in a definition list)
- Variables and values you must provide

Monospace

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Except for editorial changes, updates to this edition are marked with a vertical bar to the left of the change.

What's new in this edition

The changes in this edition relate to IBM Z Performance and Capacity Analytics V3.1.0 new function and enhancements.

This edition includes new and changed information on the following:

Profiler

New: CPU component data table [“CP_CPU_LPAR_PH”](#) on page 45

LPAR Hourly Forecast

New: CPU component data table [“CP_CPU_LPAR_FH”](#) on page 44

Forecasting

- [“Supplied forecasts”](#) on page 33
- [“Tailoring a supplied forecast”](#) on page 35

- Removed: WKLD configurations WS and WR

Implementing Cognos for IBM Z Performance and Capacity Analytics reporting

Refer to the *Guide to Reporting*

PDF only: Except for editorial changes, updates to this edition are marked with a vertical bar [|] to the left of the change.

Chapter 1. Introduction

IBM Z Performance and Capacity Analytics is a bundle that includes IBM Z Performance and Capacity Analytics and all its components together with the Capacity Planning for z/OS component.

IBM Z Performance and Capacity Analytics is a reporting system that curates system and performance data generated by various systems and stores the curated data in a Db2 database. The Capacity Planning for z/OS component provides predictive analysis to manage future resource needs, to help understand when capacity or resource shortfalls might occur. IBM Z Performance and Capacity Analytics uses historic data in the IBM Z Performance and Capacity Analytics database to forecast expected future resource usage.

IBM Z Performance and Capacity Analytics includes the following:

- Capacity Planning for z/OS component
- Forecaster module
- Db2 data tables
- Capacity Planning reports and dashboards

Capacity Planning reports are provided on the following platforms:

- IBM Z Performance and Capacity Analytics Reporting dialog (with optional QMF and GDDM), from the IBM Z Performance and Capacity Analytics database
- Cognos for web-based reporting from the IBM Z Performance and Capacity Analytics database
- Splunk and ELK dashboards for off-platform reporting enabled by the integration of IBM Z Performance and Capacity Analytics with IBM Common Data Provider for z

IBM Z Performance and Capacity Analytics includes all the IBM Z Performance and Capacity Analytics features by default. IBM Z Performance and Capacity Analytics provides a complete Capacity Planning, Performance Management and System Health Check tool for z/OS systems.

IBM Z Performance and Capacity Analytics provides capacity planning and forecasting for the following resources:

- CPU
- Storage
- Channel
- Disk
- Tape

Architecture

The IBM Z Performance and Capacity Analytics collection engine transforms input log records into aggregated data records that consolidate data from multiple SMF intervals.

These metrics are stored in Db2 data tables. Forecaster modules then analyze the data tables, apply forecasting algorithms and store the results in Db2 Forecast data tables. Both Db2 tables provide input for capacity planning and forecasting reports.




Figure 1. IBM Z Performance and Capacity Analytics Architecture

IBM Z Performance and Capacity Analytics extends the IBM Z Performance and Capacity Analytics data and reporting infrastructure to add:

- Reports that are specifically designed for Capacity Planning
- Code to generate forecasts for future resource usage
- Reports to integrate the forecasts and the reports

The Forecaster module is written in Java and is zIIP eligible. A number of configurations are provided to drive forecasts for CPU consumption, storage, media, and channels.

The Forecast data table primary key mirrors the data table structure of the data table on which the forecast was run, but the Forecast data table structure also allows multiple forecasts to be stored, over time, against a single set of data.

Table 1. IBM Z Performance and Capacity Analytics Db2 tables, primary key structure	
Data Table primary key columns	Forecast Data Table primary key columns
COLUMNNA COLUMNB COLUMNNC	COLUMNNA COLUMNB COLUMNNC F_ALGORITHM F_TIMESTAMP

Subcomponents

Although IBM Z Performance and Capacity Analytics can analyze data and produce reports for all the supported resources, you might not need all that data. The Capacity Planning for z/OS component has multiple subcomponents that you can configure to meet your needs.

Subcomponents are groups of objects (for example, predefined update definitions, data tables, and reports). You can install one, some, or all of the subcomponents. When deciding which subcomponents to install, be aware that if you find that you need reports from a subcomponent that you have not installed,

you must install that subcomponent and then wait several days or weeks until enough data has been collected to create reports. However, if you install more subcomponents than you need, IBM Z Performance and Capacity Analytics collects needless data, which takes up disk space and uses processor time.

The Capacity Planning for z/OS component is made up of the following subcomponents:

- CPU
- Storage
- Channel
- Disk
- Tape

CPU

This subcomponent provides key capacity planning metrics on processor activity, collected on SMF timestamp, hourly, daily, and weekly level. It aggregates the data by capacity planning shift as well as CEC level, LPAR level, Workload level, Service or Report Class level, and Business Application level.

A Business Application is a grouping of Service or Report Classes, as defined in lookup table CP_BUSINESS_APPL.

A shift is a short period of time each day or particular times each week when the resource usage is at peak level and where planning of sufficient capacity is important. MIPS capacity planning is most effectively done on predefined shift levels. A more accurate forecasting result is obtained by restricting the scope of the data analyzed to the narrow window of a specific shift, because the input data is more consistent. Forecasting by shift would also result in less CPU and elapsed time used by the Forecaster to predict future usage.

This subcomponent requires the collection of data from the following SMF records:

- SMF type 70
- SMF type 71
- SMF type 72 (subtype 3)
- SMF type 73
- SMF type 89 (subtype 1)
- SMF type 113 (subtype 2) - These records are used to calculate MIPS usage and capacity. IBM Z Performance and Capacity Analytics calculates MIPS based on the LSPR Workload.

Storage

This subcomponent provides key capacity planning metrics on storage availability, collected on hourly, daily, and weekly level. It aggregates the data by capacity planning shift as well as MVS System ID.

This subcomponent requires the collection of data from the following SMF records:

- SMF type 71

Channel

This subcomponent provides metrics on channel availability, collected on hourly, daily, and weekly level. It aggregates the data by capacity planning shift as well as MVS System ID.

This subcomponent requires the collection of data from the following SMF records:

- SMF type 73

Disk

This subcomponent provides metrics on disk usage collected on daily and monthly level and by MVS System ID.

This subcomponent requires the collection of data from the following DCOLLECT records:

- DCOLLECT_V

Tape

This subcomponent provides metrics on tape usage collected on daily and weekly level and by MVS System ID.

This subcomponent requires the collection of data from the following RMM records:

- DFRMM_DATASET

Data tables

IBM Z Performance and Capacity Analytics has various Db2 data tables within each subcomponent. The data tables contain capacity planning and forecasting data collected, generated and used for analysis and reporting by the Capacity Planning for z/OS component.

For descriptions of common data tables used by IBM Z Performance and Capacity Analytics, refer to the *IBM Z Performance and Capacity Analytics Administration Guide and Reference*.

The following tables show the default configuration of the data tables within each subcomponent.

CPU

Table 2. CPU subcomponent data tables	
Populated by Collector	Populated by Forecaster
CP_CPU_LPAR_T	
CP_CPU_LPAR_H	CP_CPU_LPAR_FH, CP_CPU_LPAR_PH
CP_CPU_LPAR_D	
CP_CPU_LPAR_W	CP_CPU_LPAR_F
CP_CPU_WKLD_H	
CP_CPU_WKLD_D	
CP_CPU_WKLD_W	CP_CPU_WKLD_F
CP_LSPR_WKLD_H	N/A
CP_LSPR_WKLD_D	N/A
CP_LSPR_WKLD_W	N/A

Storage

Table 3. Storage subcomponent data tables	
Populated by Collector	Populated by Forecaster
CP_STORAGE_H	
CP_STORAGE_D	
CP_STORAGE_W	CP_STORAGE_F

Channel

<i>Table 4. Channel subcomponent data tables</i>	
Populated by Collector	Populated by Forecaster
CP_CHANNEL_H	
CP_CHANNEL_D	
CP_CHANNEL_W	CP_CHANNEL_F

Disk

<i>Table 5. Disk subcomponent data tables</i>	
Populated by Collector	Populated by Forecaster
DFSMS_VOLUME_D	
DFSMS_VOLUME_M	CP_DISK_F

Tape

<i>Table 6. Tape subcomponent data tables</i>	
Populated by Collector	Populated by Forecaster
CP_TAPE_D	
CP_TAPE_W	CP_TAPE_F

Lookup tables and control tables

IBM Z Performance and Capacity Analytics uses lookup tables and control tables to contain user-specified parameters specific to your installation. These tables are used when storing data in the IBM Z Performance and Capacity Analytics database and must be configured before data collection and reporting.

Lookup tables

The lookup tables contain installation-specific parameters including Business Application groups, Capacity Planning periods and threshold values for MIPS and Channel capacity.

Control tables

The Capacity Planning for z/OS component uses one control table which must be initialized with the MIPS ratings of your installation. It is a pre-requisite to reporting MIPS values.

Reports

IBM Z Performance and Capacity Analytics provides many types of reports that use curated historical data to provide predictive analysis both for managing future resource needs and to ensure that capacity shortfalls do not occur.

The forecasting features in the reports enable you to use data-driven practices to more accurately determine when the environment will reach a capacity threshold and require hardware upgrades. This can help you plan upgrades in a timely way that can enable you to save money and maintain business service levels.

The reports are grouped by Capacity Planning for z/OS component subcomponent.

CPU subcomponent reports

Reports provided in the CPU subcomponent provide information on the utilization of your environment in MIPS, allowing you to determine the portion of usage of a CPU that is attributable to certain workloads. Reports are available at the CEC, LPAR, Workload, Application and Workload Manager Service Class or Report Class level. CEC and LPAR views of the data show the utilization at an aggregated level to understand the cumulative utilization of the available capacity, while Application and Workload reports offer a more granular view.

Four categories of reports are provided in this subcomponent:

CPU usage and forecasting

CPU usage, as well as forecasted CPU usage, is reported as MIPS. The reports report on the utilization of your environment in MIPS, allowing you to determine the portion of usage of a CPU that is attributable to certain workloads. Reports are available at the CEC, LPAR, Workload, Application and Workload Manager Service Class or Report Class level. CEC and LPAR views of the data show the utilization at an aggregated level to understand the cumulative utilization of the available capacity, while Application and Workload reports offer a more granular view. Reports are also provided which show the actual MIPS usage versus the predicted MIPS usage, over time.

LSPR workload analysis

These reports show the LSPR workload calculated on every SMF interval during the reporting interval. You can use these reports to provide insight into how the system LSPR workload is calculated by IBM Z Performance and Capacity Analytics. IBM Z Performance and Capacity Analytics determines the most popular LSPR Workload (Low/Average/High) during the reporting interval and this is used as your system LSPR Workload for MIPS calculations.

Software cost analysis

Software cost analysis reports are provided to help you understand how to manage your software costs. Certain software product costs are charged based on the peak four-hour rolling average (4HRA) MSU usage of an environment during the month. Available reporting shows the utilization of the workloads identified from SMF type 89 records. The ability to report on the workloads that contribute to the peak periods of MSU usage can pinpoint opportunities for rescheduling work away from those peak periods to lower the overall utilization. For example, a report may show that Db2 is a significant contributor to the 4HRA which is attributed to very large, non-production queries running during a heavy processing period. These queries could be scheduled at a less busy time to level out the peaks.

Storage subcomponent reports

Storage (memory) reports provide information on CPU storage and help you monitor the overall usage, and the highs and lows of storage available for use. Available reports show the minimum, average and maximum central storage that is available over time as well as a definable threshold of acceptable levels. Reports are also provided which show the actual usage versus the predicted usage, over time.

Channel subcomponent reports

Channel reports provide information on the speed of the I/O operations in a system. Reports are also provided that show the channel busy percentage versus the predicted channel busy percentage, over time.

Disk subcomponent reports

Disk reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

Tape subcomponent reports

Tape reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

Chapter 2. Implementing IBM Z Performance and Capacity Analytics

IBM Z Performance and Capacity Analytics must be installed and in operation before IBM Z Performance and Capacity Analytics can be installed. Refer to the IBM Z Performance and Capacity Analytics publications for installation information.

Implementing IBM Z Performance and Capacity Analytics consists of two steps:

- Step 1: Install and configure the Capacity Planning for z/OS component.
- Step 2: Install and configure the Forecaster module.

Step 1: Installing and configuring the Capacity Planning for z/OS component

IBM Z Performance and Capacity Analytics must be installed and in operation before IBM Z Performance and Capacity Analytics can be installed. Refer to the IBM Z Performance and Capacity Analytics publications for installation information.

About this task

Before you install the Capacity Planning for z/OS component, you must do the following steps.

Procedure

1. Decide which subcomponents to install.
2. Make input data available.
3. Customize.

Step 1.1: Deciding which subcomponents to install

You can install one, multiple, or all subcomponents.

Procedure

Decide which subcomponents to install.

Installing only the component parts that are needed to meet your requirements ensures that the collected data is useful while it minimizes the performance impact caused by data collection and interpretation activities.

Step 1.2: Making input data available

In order to perform capacity planning and forecasting, you must make the relevant input records available for processing by IBM Z Performance and Capacity Analytics.

Procedure

1. Ensure that the following SMF record types are available for processing by IBM Z Performance and Capacity Analytics :
 - 70
 - 71
 - 72 (subtype 3)
 - 73
 - 89 (subtype 1)

- 113 (subtype 2) - This is mandatory for CPU reporting. These records are used to calculate and forecast MIPS usage and capacity because IBM Z Performance and Capacity Analytics calculates MIPS based on the LSPR Workload.

Refer to *MVS System Management Facilities (SMF)* for information on how to make SMF records available.

2. Ensure that the interval at which the SMF records are written must be the same for the following SMF records:

- 70
- 72 (subtype 3)
- 113 (subtype 2)

3. Ensure that the following RMM and DFSMS records are available:

- DFRMM_DATASET
- DCOLLECT_V

Refer to *System Performance Feature Reference 1* for information on these record types and how to generate the records.

Step 1.3: Configuring lookup tables

The Capacity Planning for z/OS component uses several lookup tables when storing data in the IBM Z Performance and Capacity Analytics database. Before collecting any data, you must update the lookup tables to include parameters that are specific to your installation.

Procedure

To update each of the following lookup tables, copy the corresponding member from *HLQ.SDRLEDFS* to your *HLQ.LOCAL.DEFS* data set, then edit the sample INSERT statements to suit your requirements.

Table 7. Lookup tables			
Lookup table name	Lookup table function	Used by Capacity Planning for z/OS component	Member in HLQ.SDRLEDFS
CP_BUSINESS_APPL	To define Business Applications for reporting and forecasting purposes. A business application is a combination of either Service Classes or Report Classes.	CPU	DRLTCPL3
CP_SHIFT	To define capacity planning shifts. For example: Online or Batch.	CPU Storage Channel MSU	DRLTCPL2
CP_THRESHOLDS	To specify threshold values for capacity planning reporting.	CPU	DRLTCPL1

CP_THRESHOLDS lookup table

The CP_THRESHOLDS lookup table is used to calculate resource capacity threshold.

Table 8. CP_THRESHOLDS lookup table			
Column name	Primary key	Data type	Description
SYSPLEX_NAME	K	CHAR(8)	Sysplex name. This can contain global search characters.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. This can contain global search characters.
LPAR_NAME	K	CHAR(8)	LPAR name. This can contain global search characters.
SHIFT	K	CHAR(10)	Capacity Planning shift. This can contain global search characters.
EXCEPTION_ID	K	CHAR(10)	The exception identifier. Default supplied values are: CEC_BUSY, LPAR_BUSY
THRESHOLD		DECIMAL(10,1)	The threshold value.
EXCEPTION_DS		CHAR(40)	The threshold exception description.

The following is an example of how the CP_THRESHOLDS lookup table might look.

SYSPLEX_NAME	MVS_SYSTEM_ID	LPAR_NAME	SHIFT	EXCEPTION_ID	THRESHOLD	EXCEPTION_DS
%	%	%	ONLINE	CEC_BUSY	90.0	CEC Busy > 90%
%	%	%	BATCH	CEC_BUSY	70.0	CEC Busy > 90%
%	%	%	%	LPAR_BUSY	90.0	LPAR Busy > 90%

Figure 2. Example: CP_THRESHOLDS lookup table

In this example a threshold of 90% busy is specified for all CECs during the Online shift, and a threshold of 70% busy is specified for all CECs during the batch shift. All LPARs have a threshold of 90% busy, regardless of the shift.

The following default values are supplied in this lookup table:

CEC Busy 90%
LPAR Busy 90%

CP_SHIFT lookup table

The CP_SHIFT lookup provides the ability to define the capacity planning shifts for which you need to plan the continued availability and capacity of resources.

Capacity planning should not be done for a 24 hour period. Instead, it should be done for identified capacity planning shifts only, where resource usage is high and availability must be guaranteed. You should define the capacity planning shifts that reflect the time periods that you need to measure for capacity planning purposes.

The CP_SHIFT lookup table contains these columns:

Table 9. CP_SHIFT lookup table			
Column name	Primary key	Data type	Description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. This can contain global search characters.

Table 9. CP_SHIFT lookup table (continued)			
Column name	Primary key	Data type	Description
CP_SHIFT_NM	K	CHAR(10)	Name of the capacity planning shift.
DAY_TYPE	K	CHAR(8)	Day type the period applies to. This can be any of the day types specified in the DAY_OF_WEEK and SPECIAL_DAY control tables.
START_TIME	K	TIME	Time when the shift starts.
END_TIME	K	TIME	Time when the period ends.
SHIFT_DESC		CHAR(40)	Description of the shift.

Refer to the *Administration Guide and Reference* for more information on the DAY_OF_WEEK and SPECIAL_DAY control tables.

This lookup table is used to populate column SHIFT, which is present in all of the Capacity Planning for z/OS component data tables of the following subcomponents: CPU, Storage, and Channel.

The following is an example of how the CP_SHIFT lookup table might look.

MVS_SYSTEM_ID	CP_SHIFT_NM	DAY_TYPE	START_TIME	END_TIME	SHIFT_DESC
%	ONLINE	MON	10.00.00	14.00.00	Online
%	ONLINE	TUE	10.00.00	14.00.00	Online
%	ONLINE	WED	10.00.00	14.00.00	Online
%	ONLINE	THU	10.00.00	14.00.00	Online
%	ONLINE	FRI	10.00.00	14.00.00	Online
%	BATCH	MON	20.00.00	23.00.00	Overnight Batch
%	BATCH	TUE	20.00.00	23.00.00	Overnight Batch
%	BATCH	WED	20.00.00	23.00.00	Overnight Batch
%	BATCH	THU	20.00.00	23.00.00	Overnight Batch
%	BATCH	FRI	20.00.00	23.00.00	Overnight Batch
%	OTHER	SAT	00.00.00	24.00.00	Not used
%	OTHER	SUN	00.00.00	24.00.00	Not used
%	OTHER	HOLIDAY	00.00.00	24.00.00	Not used

Figure 3. Example: CP_SHIFT lookup table

In this example, two capacity planning shifts are defined:

Online

Includes every week day from 10am to 2pm

Batch

Includes every week day from 8pm to 11pm

There are no default values provided for this lookup table. If there is no capacity planning shift defined for a particular hour of a particular day, IBM Z Performance and Capacity Analytics will populate column SHIFT with the value OTHER for all rows in all Capacity Planning for z/OS component data tables, for that particular day/hour.

CP_BUSINESS_APPL lookup table

Reporting and forecasting can be done on Service Class, Report Class, and Business Application level. A Business Application is a grouping of specified Service Classes or Report Classes. In order to report at the Business Application level, you must define your Business Applications in the CP_BUSINESS_APPL lookup table.

The CP_BUSINESS_APPL lookup table provides the ability to define your Business Application groups. If Business Applications are defined in this lookup table, the IBM Z Performance and Capacity Analytics Collector will use this information to aggregate curated SMF data by the specified Business Application groups, before it is stored in the Db2 data tables. Similarly, the Forecaster will forecast on Service Class,

Report Class, as well as Business Application level, if Business Applications are defined in this lookup table.

You do not need to specify Business Applications in the lookup table, if you do not want to report or forecast on Business Application level.

<i>Table 10. CP_BUSINESS_APPL lookup table</i>			
Column name	Primary key	Data type	Description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. This can contain global search characters.
WORKLOAD	K	CHAR(8)	MVS workload name. This can contain global search characters.
SERVICE_CLASS	K	CHAR(8)	WLM service class or report class name. This can contain global search characters.
SERV_CLASS_PERIOD	K	CHAR(1)	WLM service class or report class period. This can contain global search characters.
BUSINESS_APPL		CHAR(10)	Business application name

This lookup table is used to populate column BUSINESS_APPL of the following data tables:

CP_CPU_WKLD_H
CP_CPU_WKLD_D
CP_CPU_WKLD_W
CP_CPU_WKLD_F

The following is an example of how the CP_BUSINESS_APPL lookup table might look.

MVS_SYSTEM_ID	WORKLOAD	SERVICE_CLASS	SERV_CLASS_PERIOD	BUSINESS_APPL
-----	-----	-----	-----	-----
SYSA	BATCH	WEBPAY1	1	PAYMENTS
SYSA	BATCH	WEBPAY2	%	PAYMENTS
SYSB	SYSTEM	SYSC%	%	CREDITCARD

Figure 4. Example: CP_BUSINESS_APPL lookup table

The first two rows are used to create a Business Application called PAYMENTS. Business Application PAYMENTS consists of Service or Report Class WEBPAY1 (Period 1 only), as well as Service Class or Report Class WEBPAY2 (all Periods). The third row is used to create a Business Application called CREDITCARD. CREDITCARD consists of all Service or Report Classes whose names start with SYSC (all Periods).

There are no default values provided for this lookup table.

Step 1.4: Configuring control tables

The Capacity Planning for z/OS component uses one control table when storing data in the IBM Z Performance and Capacity Analytics database. This is the LSPR_MIPS control table and must be initialized before use with parameters specific to your installation.

About this task

Before collecting any data and reporting on MIPS values, the LSPR_MIPS control table must be initialized.

The LSPR_MIPS control table contains the multi image LSPR MIPS ratings for all processor and submodel types. The table has the following columns.

Table 11. Control table

Column name	Primary key	Data type	Description
CPU_MODEL_NO	K	SMALLINT	Machine type
CPU_SUB_MODEL	K	CHAR(16)	CPU submodel
PROCESSOR_TYPE	K	CHAR(4)	Processor pool type
LSPR_WORKLOAD	K	CHAR(4)	LSPR workload
MIPS_RATING		SMALLINT	Processor pool MIPS rating

Procedure

1. Initialize the MIPS ratings in the LSPR_MIPS table using sample JCL provided in the *HLQ.SDRLCNTL* data set member DRLJMIPS. Before you submit the job, review the instructions in the sample JCL and change the parameter settings to match your configuration.
2. When your changes are complete, submit the job.
Running this sample job DRLJMIPS will ensure that the LSPR_MIPS control table is initialized with the LSPR provided MIPS ratings for all processor model and submodel combinations.

Step 2: Installing and configuring the Forecaster module

You must install and configure the Forecaster module before you can use the Capacity Planning for z/OS component.

Step 2.1: Unpacking the Forecaster module

The Forecaster is a USS application and it is shipped as a renamed .tgz file.

Procedure

1. Locate the DRLPJFC member.
2. Extract the .tar file by using the following command where *destination* is the location where you want the contents of the .tar file to reside and the path name specified is the default path name.

```
tar -xvof /usr/lpp/IDSz/v2r1m0/IBM/DRLPJFC -C /destination
```

This creates the following subdirectories:

```
IDSz
IDSz/forecaster
IDSz/forecaster/config
IDSz/forecaster/secure
IDSz/forecaster/mappings
IDSz/forecaster/java
```

The directories are as follows:

forecaster

The working directory for the forecaster instance.

config

Stores the .properties files that the forecaster uses as configurations. These tell it what to forecast and how to make the forecast.

secure

Holds your Db2 access credentials.

mappings

Requires a copy of the JSON mapping files.

java

Holds the code.

3. Create a working copy of the forecaster directory by copying it and its subdirectories to another location.

This means you will not overwrite any changes to make to it when you unpack a service update to the forecaster and its configuration.

Step 2.2: Updating the table mapping files

You must update the table mapping files.

Procedure

1. Run the table mapping utility.
2. Copy the .json mapping files that are produced by the table mapping utility into the Forecaster module's mapping subdirectory `IDSz/forecaster/mappings`.

What to do next

If you subsequently make changes to the structure of any tables that are involved in the forecasting process or apply any updates to IBM Z Performance and Capacity Analytics or IBM Z Performance and Capacity Analytics, you must repeat this step.

Step 2.3: Configuring the Forecaster module

You must configure the Forecaster module.

About this task

Perform these steps in your working directory.

Procedure

1. Edit the `Forecaster.sh` shell script and update the values for the following variables to match your site:

rundir

The directory that contains the `Forecaster.sh` script.

logfile

The log file that sets the message output properties.

config

This maps the input parameter into a config file. By default the config files are stored in a subdirectory called `config`. You only need to change this if you want to store the configurations elsewhere.

db2jdbc

The directory that contains the JDBC .jar files for the Db2 that it is running against.

JAVA_HOME

The directory where you have an instance of the IBM Java SDK 8 installed.

2. Use the `DRLJFCP` sample to produce a JCL procedure to run the Forecaster module as a job.
You will use this JCL to run the Forecaster module.
3. Create a user ID that can only access the Db2 database that is used by IBM Z Performance and Capacity Analytics. Ensure that the ID has read/write access.
4. Edit the `IDSz/forecaster/secure/db2access.properties` sample to specify the userid and password that you created.
5. Restrict access to the `IDSz/forecaster/secure` subdirectory to only the following users:
 - The administrator who manages the user ID and password for the Db2 access must have read/write access.

- The Batch user ID that runs the forecasting must have read only access.

When the password is changed, only the `IDSz/forecaster/secure/db2access.properties` file needs to be updated.

Chapter 3. Implementing Cognos reporting (optional)

This is an optional feature. You can use Cognos for reporting from the IBM Z Performance and Capacity Analytics Db2 tables. If you are using this web reporting method for the first time, follow the implementation steps described in the *Guide to Reporting*.

Chapter 4. Implementing ELK reporting (optional)

This is an optional feature. You can use ELK for off-platform reporting.

IBM Z Performance and Capacity Analytics dashboards are also available as ELK reports. The IBM Z Performance and Capacity Analytics Collect engine can be configured to stream the rows collected to the Db2 data tables to ELK, by making use of the IBM Common Data Provider for z.




Figure 5. Streaming data to ELK by the Collector

The Forecaster can also be configured to stream the forecasted data to ELK for reporting. A new Shadowing process must be configured to allow the streaming of forecasted data to ELK via the Common Data Provider for z.




Figure 6. Streaming data to ELK by the Forecaster

The IBM Common Data Provider for z is an IBM offering that enables the gathering of data from many different data sources on a Z system and provides a single framework to forward that data to one or more remote receivers, including ELK.

Step 1: Setting up data streaming through IBM Common Data Provider for z

To implement ELK reporting, first set up data streaming from Db2 to ELK through IBM Common Data Provider for z.

About this task

IBM Z Performance and Capacity Analytics provides sample JCL member DRL.SDRLCNTL(DRLJCDPS) containing three utilities for generating stream definitions and table mappings from the Db2 database as required for the streaming of the Db2 data table content to ELK.

Procedure

Customize the sample JCL member DRL.SDRLCNTL(DRLJCDPS) to suit your environment.

DRLJCDPS runs three utility programs:

DRLELSTT

The first utility (DRLELSTT) creates a list of all the data tables in the IBM Z Performance and Capacity Analytics database. This is used as input to the other two utilities to ensure they produce consistent output.

DRLEMCDS

The second utility (DRLEMCDS) creates a `izds.streams.json` file, which contains an IBM Common Data Provider for z stream definition for each database table, as well as for each view based on each database table. The `izds.streams.json` file is required by IBM Common Data Provider for z. It needs to be copied into the configuration folder for the IBM Common Data Provider for z user interface. For more information, refer to [“Step 5: Configuring IBM Common Data Provider for z for data streaming” on page 22.](#)

DRLEMTJS

The third utility (DRLEMTJS) creates a `table.json` file for every table in the database, as well as for every view based on these tables. Each file maps the contents of the table. The `table.json` files are required by the Publication Forecaster. For more information, refer to [“Step 4: Configuring the Forecaster for data streaming” on page 19, *Customize the Publication Forecaster configuration.*](#)

What to do next

If you subsequently apply maintenance or customizations to your IBM Z Performance and Capacity Analytics Db2 database which result in database schema changes, you must rerun the sample JCL DRLJCDPS to create new data tables mappings and stream definitions.

Step 2: Installing and configuring IBM Common Data Provider for z

IBM Z Performance and Capacity Analytics makes use of the IBM Common Data Provider for z to enable the capacity planning and forecast data gathered in Db2 tables to be shadowed and forwarded to ELK.

Procedure

Install and configure IBM Common Data Provider for z.

Refer to the *IBM Common Data Provider for z User Guide* for information.

Step 3: Configuring IBM Z Performance and Capacity Analytics for data streaming

Complete the following steps to configure IBM Z Performance and Capacity Analytics to stream curated data to a publication log stream.

Procedure

1. Allocate a publication log stream.

Refer to the *Administration Guide and Reference* for instructions on how to set up a log stream.

A publication log stream is required for every system that is running an instance of the IBM Z Performance and Capacity Analytics Collector.

2. Specify the DISTRIBUTE parameter on the COLLECT statement.

A new parameter DISTRIBUTE on the COLLECT statement results in the streaming of data to the specified publication log stream.

```
COLLECT SMF
DISTRIBUTE TO IZDS.MY.NEW.LOGSTREAM
COMMIT AFTER 60 MINUTES
```

Refer to the *IBM Z Performance and Capacity Analytics Language Guide and Reference* for the syntax of the DISTRIBUTE parameter on the COLLECT statement.

Note:

- All data tables are streamed when the DISTRIBUTE parameter is specified.
- Streaming occurs at the time when the inserted and updated rows are committed to the database.
- All streamed rows will be converted to JSON and presented to the IBM Common Data Provider for z Data Streamer. The IBM Common Data Provider for z Data Streamer will discard any records for which it does not have a subscriber.
- The DISTRIBUTE parameter can only be used once all configuration steps described in this chapter are completed.

Step 4: Configuring the Forecaster for data streaming

Complete the following steps to configure the IBM Z Performance and Capacity Analytics Forecaster to move the data from the publication log stream to the IBM Common Data Provider for z.

About this task

This needs to be done by the Forecaster, and not by the Data Mover (as per IZDS Collect data streaming), because the Capacity Planning metrics are contained in Db2 views, rather than Db2 data tables. The Forecaster needs to read the contents of these views, and send it to ELK, as if it were actual data table content. This process is referred to as "Shadowing".

Procedure

1. Create a new Publication Forecaster directory.

Take a copy of the SMP/E installed IDSz/Forecaster directory. This new directory will be the working directory for the Publication Forecaster:

```
cp -R IDSz/Forecaster/* IZDS/publisher
```

2. Customize the Publication Forecaster configuration.

Customize the Publication Forecaster configuration used by the publisher to set up a pair of data routes. The first one is the classic publication data route, reading data from the publication logstream, transforming it into JSON and sending it to IBM Common Data Provider for z. The second one is the Shadow process, scanning Db2 to find new data added to views, transforming it into JSON and sending it to IBM Common Data Provider for z.

Edit the PubShadow.properties file in the IZDS/publisher directory as follows:

- Change input.1.logstream to the newly created publication log stream name.
- Change process.1.1.encoding if your system uses a JSON code page different to the default 1047.
- Change output.1.1.port to the output port on the TCPIP stage that the IBM Common Data Provider for z's Data Streamer is listening on.
- Change input.2.connection to the localhost and the port number that the Db2 holding the IBM Z Performance and Capacity Analytics scheme you want to forecast against has open for JDBC connections.
- Change input.2.location to the location value that you specified for the Db2 Subsystem when it was configured.
- Change input.2.schema to the schema that is holding the IBM Z Performance and Capacity Analytics tables you want to use for forecasting.
- Change output.2.1.port to the output port on the TCPIP stage that the IBM Common Data Provider for z's Data Streamer is listening on.

```
# Two routes
#
```

```

routes = 2
#
# Route 1 - Publisher
#
route.1.name = Publisher
#
# The input arrives in a log stream, put there by the COLLECTOR
#.
# This will take the data from the log stream and then erase it once it has been published.
#
input.1.type = LOGSTREAM
input.1.logstream = IFASMF.CF.LS01
input.1.wipe = YES
input.1.strip_header = NO
input.1.check_LRGH = YES
input.1.sourcename = COLLECTOR
input.1.sourcetype = AGGREGATE
#
process.1 = 2
#
# Filter out unwanted tables
# Only data for the tables listed gets past this stage.
#
process.1.1.type = FILTER
process.1.1.as = TDSz
process.1.1.tables = 19
process.1.1.table.1 = CP_CPU_LPAR_H
process.1.1.table.2 = CP_PROD_MSU_H
process.1.1.table.3 = CP_LSPR_WKLD_H
process.1.1.table.4 = CP_LSPR_WKLD_D
process.1.1.table.5 = CP_LSPR_WKLD_W
process.1.1.table.6 = CP_TAPE_D
process.1.1.table.7 = CP_TAPE_W
process.1.1.table.8 = CP_STORAGE_H
process.1.1.table.9 = CP_STORAGE_D
process.1.1.table.10 = CP_STORAGE_W
process.1.1.table.11 = DFSMS_VOLUME_D
process.1.1.table.12 = DFSMS_VOLUME_M
process.1.1.table.13 = CP_CHANNEL_H
process.1.1.table.14 = CP_CHANNEL_D
process.1.1.table.15 = CP_CHANNEL_W
process.1.1.table.16 = CP_TAPE_F
process.1.1.table.17 = CP_STORAGE_F
process.1.1.table.18 = CP_DISK_F
process.1.1.table.19 = CP_CHANNEL_F
#
# Convert to JSON
#
process.1.2.type = JSON
process.1.2.encoding = IBM-1047
process.1.2.float_format = IBM
#
# Send to CDPz
#
outputs.1 = 1
#
# Over TCPIP
#
output.1.1.type = TCPIP
output.1.1.host = localhost
output.1.1.port = 45020
output.1.1.directory = bufferR1
#
#
# Route 2 - Shadower
#
route.2.name = Shadower
#
# The SHADOW stage reads data from Views in Db2 and emits only new records.
#.
# This takes data from the listed tables and views. A High water mark is kept
# in the /shadow subdirectory.
#
# The JDBC details need updating, as do the security credentials.
#
# The JDBC connection details
# The Db2 location strings
# The Schema name
# The list of views to be shadowed
#
# Note that the shadow mechanism is only used for views. Updates to tables
# populated by the collector (continuous or batch) get published as they

```

```

# happen through the publisher.
#
input.2.type = SHADOW
input.2.interval = 15
input.2.protocol = db2
input.2.connection = localhost:5012
input.2.location = QXPBMR1DEC1
input.2.access = secure/db2access.properties
input.2.schema = DRL
#
#
input.2.tables = 8
#
input.2.table.1.name = CP_CPU_LPAR_HV
input.2.table.1.frequency = HOURLY
input.2.table.1.date = DATE
input.2.table.1.time = TIME
#
input.2.table.2.name = CP_CPU_LPAR_DV
input.2.table.2.frequency = DAILY
input.2.table.2.date = DATE
#
input.2.table.3.name = CP_CPU_LPAR_WV
input.2.table.3.frequency = WEEKLY
input.2.table.3.date = DATE
#
input.2.table.4.name = CP_CPU_LPAR_FV
input.2.table.4.frequency = WEEKLY
input.2.table.4.timestamp = F_TIMESTAMP
#
input.2.table.5.name = CP_CPU_WKLD_HV
input.2.table.5.frequency = HOURLY
input.2.table.5.date = DATE
input.2.table.5.time = TIME
#
input.2.table.6.name = CP_CPU_WKLD_DV
input.2.table.6.frequency = DAILY
input.2.table.6.date = DATE
#
input.2.table.7.name = CP_CPU_WKLD_WV
input.2.table.7.frequency = WEEKLY
input.2.table.7.date = DATE
#
input.2.table.8.name = CP_CPU_WKLD_FV
input.2.table.8.frequency = WEEKLY
input.2.table.8.timestamp = F_TIMESTAMP
#
input.2.check_dir = shadow
input.2.float_format = IBM
input.2.sourcename = COLLECTOR
input.2.sourcetype = AGGREGATE
#
process.2 = 1
#
# Convert to JSON
#
process.2.1.type = JSON
process.2.1.encoding = IBM-1047
process.2.1.float_format = IBM
#
# Send to CDPz
#
outputs.2 = 1
#
# Over TCPIP
#
output.2.1.type = TCPIP
output.2.1.host = localhost
output.2.1.port = 45020
output.2.1.directory = bufferR2
#

```

3. Copy the table.json files into the Publication Forecaster directory.

Use a normal copy command to move the data. If you have changed the output path in the job SDRLCNTL(DRLJCDPS) that generates the mapping, you will need to use the copy command.

```
cp /u/drl/* IDSZ/publisher/mappings
```

4. Create a new Started Task for the Publication Forecaster.

Refer to the *Administration Guide and Reference* for instructions to set up a started task.

Once started, the Publication Forecaster will be set up to read data from the publication log stream, convert it into JSON and send it to the IBM Common Data Provider for z for distribution.

Step 5: Configuring IBM Common Data Provider for z for data streaming

Complete the following steps to configure IBM Common Data Provider for z for data streaming to send data off-platform to ELK.

Procedure

1. Make IBM Z Performance and Capacity Analytics table definitions available to IBM Common Data Provider for z.

Copy the `izds.streams.json` file into the IBM Common Data Provider for z configuration UI folder.

The `izds.streams.json` file was created earlier with utility DRLEMCDs in [“Step 1: Setting up data streaming through IBM Common Data Provider for z”](#) on page 17.

2. Update your IBM Common Data Provider for z Policy to stream the data.
 - a) Find the stream definitions that correspond to each set of reports you want to be streamed from IBM Z Performance and Capacity Analytics and add it to your configuration. Based on the dashboard sets that you require, identify which of the following streams to subscribe to:

Streams for Capacity Planning CPU reports

`izds_cp_cpu_lpar_h`
`izds_cp_prod_msu_h`
`izds_cp_lspr_wkld_h`
`izds_cp_lspr_wkld_d`
`izds_cp_lspr_wkld_w`
`izds_cp_cpu_lpar_hv`
`izds_cp_cpu_lpar_dv`
`izds_cp_cpu_lpar_wv`
`izds_cp_lpar_fh`
`izds_cp_cpu_lpar_fv`
`izds_cp_cpu_lpar_fhv`
`izds_cp_cpu_wkld_hv`
`izds_cp_cpu_wkld_dv`
`izds_cp_cpu_wkld_wv`
`izds_cp_cpu_wkld_fv`

Streams for Capacity Planning Tape reports

`izds_cp_tape_d`
`izds_cp_tape_w`
`izds_cp_tape_f`

Streams for Capacity Planning Storage reports

`izds_cp_storage_h`
`izds_cp_storage_d`
`izds_cp_storage_w`
`izds_cp_storage_f`

Streams for Capacity Planning Disk reports

`izds_dfsms_volume_d`
`izds_dfsms_volume_m`

izds_cp_disk_f

Streams for Capacity Planning Channel reports

izds_cp_channel_h
izds_cp_channel_d
izds_cp_channel_w
izds_cp_channel_f

- b) Data will arrive encoded with the character set specified in `process.1.1.encoding` during configuration of the Publication Forecaster. For more information, refer to [“Step 4: Configuring the Forecaster for data streaming”](#) on page 19.

Configure IBM Common Data Provider for z to transcribe the data stream to UTF-8, a requirement for sending the data off-platform.

- c) For each recipient of the data, define a subscriber to IBM Common Data Provider for z using the correct IP address for the receiver.

Configure ELK receivers to use the Logstash protocol.

- d) Build, distribute, and activate the policy by following the instructions in the *IBM Common Data Provider for z User Guide*.

Step 6: Configuring ELK

Complete the following steps to install and configure the IBM Z Performance and Capacity Analytics dashboard components for off-platform ELK reporting.

Procedure

1. Download and install the Logstash and Kibana components.
2. Configure Logstash.
3. Configure Kibana.

Step 6.1: Download and install the dashboard files

IBM Z Performance and Capacity Analytics dashboards are provided by IBM Z Performance and Capacity Analytics as textual files. These must be downloaded (in ASCII) and installed.

About this task

The files are SMP/E installed as parts into the SDRLIDM target library. The default location is:

`/usr/lpp/IBM/IDSz/v1r9m0/IBM`

Table 12. Dashboard files for ELK reporting		
SDRLIDM File Name	Original File Name	Description
DRLWEING	IZDS_ELK_Ingest.tar	Ingestion package for ELK
DRLWECPZ	IZDS_ELK_CPZOS.tar	Capacity Planning for z Package for ELK

As they are initially installed as .tar files, you must unpack them.

Procedure

Extract the .tar files by using the following command where *filename* is given in [Table 12 on page 23](#) and *destination* is the location where you want the contents of the .tar file to reside.

```
tar -xvfo filename -C destination
```

This creates the following files:

DRLWEING (IZDS_ELK_Ingest.tar)

B_IZDS_Input.lsh
E_IZDS_Filter.lsh
Q_IZDS_Output.lsh

DRLWECpz (IZDS_ELK_CPZOS.tar)

IZDS_CP_CPU.json
IZDS_CP_Tape.json
IZDS_CP_Disk.json
IZDS_CP_Channel.json
IZDS_CP_Storage.json

Step 6.2: Configuring Logstash

Complete these steps to configure the Logstash protocol for ELK reporting.

Procedure

1. Take the three Logstash configuration files and place them in your logstash configuration directory.
Their naming convention is designed to work with the IBM Common Data Provider for z naming convention for members of the directory.
 - Edit B_IZDS_Input.lsh to set the correct IP port that you want it to listen on for the connection from the IBM Common Data Provider for z Data Streamer. The default is '40404'.
 - Edit Q_IZDS_Output.lsh to set the correct hostname (and port) for the elastic cluster. The default is simply 'localhost'.
2. Stop and restart Logstash to pick up the changes.
Logstash should now be able to ingest IBM Z Performance and Capacity Analytics data sent from the IBM Common Data Provider for z.

The data will be an index in Elastic, using the stream name as the index name (izds_table_name) and a monthly index rollover (yyyy.mm).
3. You can modify the output stage to produce tables with a shorter roll over period, but be aware that if you are running a minimal ELK configuration (as you might have on a stand-alone test system), that this will place a lot of stress on the shard and possibly cause Elastic to fail.

What to do next

Wait for data

Ensure the IBM Common Data Provider for z is feeding Logstash with data. Wait for a couple of hours to give it long enough to have downloaded at least one record in each of the data tables.

Step 6.3: Configuring Kibana

Complete these steps to configure Kibana for ELK reporting.

Procedure

1. Sign into Kibana.
2. Navigate to **Management**, then select **Saved Objects**.
3. Import the following IBM Z Performance and Capacity Analytics .json files:
 1. IZDS_CP_CPU.json
 2. IZDS_CP_Storage.json
 3. IZDS_CP_Channel.json
 4. IZDS_CP_Disk.json

5. IZDS_CP_Tape.json

Note that references between the objects are governed by identity strings that are encoded into the files. If you feel the need to recreate one of the supporting entities, you need to ensure that your recreation has the same identity as the IBM Z Performance and Capacity Analytics supplied version. If you don't do this, then you'll get errors about things not being found. You may be able to reimport the object from the supplied .json files. The Capacity Planning files can be installed in any sequence you wish, provided some data in the tables they use is present in ELK. You can select which subcomponents of Capacity Planning for z/OS component you want to install. You shouldn't need to restart ELK to pick up the changes.

4. To avoid the following message, select a default index pattern before continuing. You can select any index pattern as the default.




Figure 7. Defining an index pattern for retrieving data

What to do next

Configure IBM Common Data Provider for z to stream each of the following tables, forecasts and views to ELK.

Table 13. Data to configure for streaming to ELK			
File	Published Tables	Published Forecasts	Shadowed Views
IZDS_CP_CPU	CP_CPU_LPAR_H CP_PROD_MSU_H CP_LSPR_WKLD_H CP_LSPR_WKLD_D CP_LSPR_WKLD_W	(The forecast data is picked up from the shadowed *_FV views of the forecast tables.)	CP_CPU_LPAR_HV CP_CPU_LPAR_DV CP_CPU_LPAR_WV CP_CPU_LPAR_FV CP_CPU_LPAR_FHV CP_CPU_LPAR_PHV CP_CPU_WKLD_HV CP_CPU_WKLD_DV CP_CPU_WKLD_WV CP_CPU_WKLD_FV
IZDS_CP_Tape	CP_TAPE_D CP_TAPE_W	CP_TAPE_F CP_CPU_LPAR_FH	CP_CPU_LPAR_FHV CP_CPU_LPAR_PHV
IZDS_CP_Storage	CP_STORAGE_H CP_STORAGE_D CP_STORAGE_W	CP_STORAGE_F CP_CPU_LPAR_FH	CP_CPU_LPAR_FHV CP_CPU_LPAR_PHV
IZDS_CP_Disk	DFSMS_VOLUME_D DFSMS_VOLUME_M	CP_DISK_F CP_CPU_LPAR_FH	CP_CPU_LPAR_FHV CP_CPU_LPAR_PHV

Table 13. Data to configure for streaming to ELK (continued)

IZDS_CP_Channels	CP_CHANNEL_H	CP_CHANNEL_F	CP_CPU_LPAR_FHV
	CP_CHANNEL_D	CP_CPU_LPAR_FH	CP_CPU_LPAR_PHV
	CP_CHANNEL_W		

Refer to “[Step 5: Configuring IBM Common Data Provider for z for data streaming](#)” on page 22 for more information.

Step 7: Viewing the dashboard

The dashboards are set up like normal Kibana dashboards.

Search and Filter parameters

Time selector (top right) for the time range. This is usually an absolute range.

Filters (across the top, starting on the right) are used to frame the report.

Note:

1. The dashboards are typically designed to work with a single day and with the bulk of the filters active. Significantly expanding the search range and/or deactivating most of the filters may result in the queries behind the dashboard failing.
2. The visualizations are not designed to be viewed directly. They have search and filter parameters set on them that are critical to the correct functioning of the dashboards.

Do not change or remove the search and filter parameters. Changing or removing them may change the nature of the output and possibly corrupt the data values displayed.
3. In addition, the visualizations return mostly unfiltered data. If you run them without the additional filters the dashboards provide, this may result in the query failing because it returned too much data.

Step 8: Ad hoc reporting

When ad hoc reports are created on an off-platform base, for example ELK, be aware of potential data aggregation differences and duplicate records.

Data aggregation differences

Due to differences in rounding and computational algorithms, it is likely that data aggregated on the receiving platform, in customer written reports, will not exactly match the higher level aggregates that IBM Z Performance and Capacity Analytics creates in its Db2 tables.

Duplicate records

The IBM Z Performance and Capacity Analytics Collect engine inserts as well as updates rows in the Db2 database, and both the inserted and updated rows are streamed off-platform using the DISTRIBUTE parameter on the COLLECT statement. This could result in duplicate rows being streamed off-platform.

However, there is an ID field in each streamed record that uniquely identifies the row that it represents. ELK deals with this automatically, creating a new version of the updated document.

Chapter 5. Capacity planning

With IBM Z Performance and Capacity Analytics you can plan the future capacity needs of multiple resource types including CPU, storage, channel, disk, and tape.

Introduction to capacity planning

With IBM Z Performance and Capacity Analytics you can plan the future capacity needs of multiple resource types.

CPU

Evaluating the utilization of your environment in MIPS allows you to determine the portion of usage of the CPU that is attributable to certain workloads. Reports are available at the CEC, LPAR, Workload, Business Application and Workload Manager Service Class or Report Class level. CEC and LPAR views of the data show the utilization at an aggregated level to understand the cumulative utilization of the available capacity, while Application and Workload reports offer a more granular view.

CEC and LPAR reports are typically used for machine capacity planning purposes and to understand the nature and behavior of workload trends across the environment. For example, an hourly CEC MIPS usage report can identify specific times of day when the entire CEC may be saturated because of heavy utilization across all LPARs. This information can identify the opportunity for a workload shift to different times during the day to flatten the utilization, or the need for a capacity upgrade. Similarly, peak utilization reports can identify certain busy days, such as end of month, or seasonal spikes that allow for proactive planning to defer or reschedule workloads, or adjust capacity by LPAR re-weighting or activating Capacity on Demand.

Application and Workload Manager Service Class and Report Class reports provide a more granular view of utilization allowing for more fine-tuning of the environment. Using the CEC and LPAR reports to zero in on specific LPARs or peak periods in conjunction with the more granular reports allows for a drill-down approach to focus on specific areas to address.

Identifying the MIPS utilization of certain types of work can pinpoint specific stress points. Showing growth trends across time, or behavior tied to application changes or business growth allows you to extrapolate capacity needs and impact for the future. For example, reporting that an application's MIPS utilization doubled after a code change can detect the need to investigate the application's efficiency or to plan for more capacity if the new behavior is expected to continue.

MIPS capacity planning can be done using a focus shift which covers a short period of time each day or particular times each week when usage of the machine is heaviest and you are most concerned about having sufficient capacity. By restricting the scope of the data analyzed to a narrow window, you can get better forecast results (because the input data is more consistent) and reduce the CPU and elapsed time used to perform forecasting.

The default settings allow for forecasting for average CPU usage.

Business Application focused reporting provides a means of aggregating the reports and forecasts for a number of service and/or reporting classes to give a total value for the specific line of business.

Storage (memory)

Having sufficient central storage is vital to the overall performance and throughput of your mainframe. The ability to use storage effectively can mean the difference between meeting business service levels and suffering workload degradation. Monitoring the overall usage, and the highs and lows of storage available for use, ensures that you can plan for sufficient resources for the workload. Supplied reports to show the minimum, average and maximum central storage available over time.

The default settings allow for forecasting for minimum available storage.

Channel

I/O operations contribute significantly to the overall performance of a system's workloads. The faster the round trip to a peripheral device such as disk or tape, the better the performance of a workload and the faster it completes. Channels are the highway from the CEC to the device. Having the information to report on channel utilization and plan for capacity allows you to keep I/O traffic flowing as quickly as possible.

Available reports show the peak busy utilization of channels over various time periods. The forecasting capability can pinpoint when channels will reach a saturation point that will begin to degrade performance. For example, a report showing peak utilization of disk channels at a certain day and time during the week can reveal that an overload of weekly backups are taking place and degrading production work. Rescheduling the backups or adding channel capacity can avoid a missed service level for the business.

The default settings allow for channel peak busy loading.

Disk

Disk usage trends, as well as forecasted future disk storage usage, must be assessed in order to plan for future needs, allowing you to plan for enough space to meet the storage usage requirements of your systems.

Default configurations are provided to allow for forecasting on allocated space per storage group.

Tape

Tape usage trends, as well as forecasted future tape storage usage, must be assessed in order to plan for future needs, allowing you to plan for enough space to meet the storage usage requirements of your systems.

Default configurations are provided to allow for forecasting on allocated space per storage group.

Getting started with IBM Z Performance and Capacity Analytics

You can use IBM Z Performance and Capacity Analytics with existing IBM Z Performance and Capacity Analytics reports and metrics, to provide a complete system capacity planning and performance management solution.

About this task

Follow these steps to start using IBM Z Performance and Capacity Analytics for capacity planning and forecasting.

Procedure

1. Identify your capacity planning shifts.

Identify the days and hours for which you need to plan capacity (those shifts where resource usage is high and availability must be guaranteed). You can use IBM Z Performance and Capacity Analytics key performance measures (KPM) reports to identify peak shifts.

- KPMZLP01 (KPMZOS CEC CPU Busy, Hourly)
- KPMZLP02 (KPMZOS LPAR CPU Busy, Hourly)

2. Set up your capacity planning shifts.

Refer to section [“CP_SHIFT lookup table”](#) on page 9 for information about updating the CP_SHIFT lookup table to define your capacity planning shifts.

3. Set up your Business Applications.

If you want to report and forecast at the Business Application level, you must define your Business Applications. Refer to [“CP_BUSINESS_APPL lookup table”](#) on page 10 for information about updating the CP_BUSINESS_APPL lookup table to define your business applications.

4. Determine and set your capacity thresholds.

Review the default thresholds in the CP_THRESHOLDS lookup table and adjust them if necessary. You can set different thresholds for different shifts, and for different LPARs. Refer to [“CP_THRESHOLDS lookup table” on page 9](#) for information about updating the CP_THRESHOLDS lookup table with your thresholds.

5. Make your input data available.

Ensure that you are collecting the SMF and other record types that are required for capacity planning. Refer to [“Step 1.2: Making input data available” on page 7](#) for information about making your data available.

6. Install the Capacity Planning for z/OS component.

Refer to [“Step 1: Installing and configuring the Capacity Planning for z/OS component” on page 7](#) for information.

7. Collect data.

Collect data using the IBM Z Performance and Capacity Analytics collect engine.

8. Configure the Forecaster module.

Refer to [“Step 2: Installing and configuring the Forecaster module” on page 12](#) for information.

9. Do forecasting.

When you have sufficient data collected, use this data to forecast the required capacity during the shifts you are monitoring. The Forecaster will write the forecast values back to the IBM Z Performance and Capacity Analytics database.

Results

You can now use the capacity planning reports to monitor your usage compared with the forecast, and to make resource planning decisions.

Chapter 6. Forecasting

You can use IBM Z Performance and Capacity Analytics to forecast necessary capacity and to compare actual capacity usage with the forecast.

- IBM Z Performance and Capacity Analytics calculates CPU capacity in MIPS, on the following levels:
 - CEC
 - LPAR
 - Service/Report class
 - Business Application (this is a grouping of service or report classes)
- MIPS usage is forecast based on those levels
- MIPS capacity and usage dynamically is calculated, based on the actual workloads that are running on the machine at the time. IBM Z Performance and Capacity Analytics calculates the LSPR workload for this purpose, from the SMF type 113 subtype 2 records. The machines behavior is mapped to the best matching LSPR benchmark value, either low, medium or high.
- IBM Z Performance and Capacity Analytics includes the multi image LSPR MIPS values. This means that you do not need to manually determine the MIPS rating of your processors and then manually insert the required MIPS values into a lookup table, in order for MIPS values to be calculated.
- You can specify capacity thresholds for reporting purposes on CEC, LPAR and shift level.

In order to forecast you must first configure IBM Z Performance and Capacity Analytics to collect data, then collect the data. After the data is collected IBM Z Performance and Capacity Analytics can use it to generate reports that can be used for forecasting.

Configuration

The forecasting process is controlled by a configuration file that is read from the forecaster config subdirectory. The configuration specifies:

- Which table in Db2 to pull the input data from
- How far back to go to fetch historical data
- How to partition that data for forecasting
- The forecasting algorithm to use
- How far into the future to make the forecast for
- The time intervals to use for the forecast
- The table to write the results of the forecast to

A single forecasting run is normally executed as a batch job and will make a number of forecasts for distinct data sets that are held within the same table. The flow of a forecasting run is as follows:

1. A JDBC stage extracts the data from Db2 and partitions it for forecasting
2. Each partition is passed to the forecasting stage and a forecast is made for it
3. The forecast is passed through a SQL stage to prepare it for insertion into Db2
4. A JDBC stage then takes the SQL and executes it against Db2

The process runs mostly in parallel, although there is a pacing parameter on the JDBC input stage that you can use to slow it down.

System resource usage

The Forecaster is written in Java. This means it will be moved onto a zIIP processor if one is available. Note that the forecaster probably has a significantly longer runtime than most tasks moved onto a zIIP. A weekly forecasting run can easily consumer 100+ CPU seconds and run for 2 to 3 minutes.

The forecasting method used is an iterative least squares approach, meaning it's around $O(n^3)$ - twice the number of data points means 8x the CPU used. You will see the CPU cost for forecasting continue to rise until you've gathered a full set of forecast input data - 6 month's worth for the weekly forecasts. The CPU cost of running a weekly forecast on 6 months data is almost 3 orders of magnitude higher than the cost of running the same forecast on only 1 month's data. The cost of running a daily forecast on a 90 day's worth of data is yet another order of magnitude higher (as it's related to the number of data points, not the time span that those data points cover).

The cost of forecasting is related to the number of data points used to make the forecast, not the capacity of the system that the forecast is being made for. A 600 MSU system is no more expensive to forecast for than a 20 MSU system. The cost of forecasting a Sysplex depends on the number of systems in the Sysplex. Twice as many systems is twice the cost. This technique can be used to reduce the cost of forecasting by limiting the forecasting functions to only certain systems. This technique can be used to spread the cost of forecasting out over multiple days by splitting the forecasting up by Sysplex.

There are several ways to reduce the system resources that are required to forecast. You can start by using the supplied forecasts to make your own variants. Doing so can significantly reduce the cost of the forecasting process while still delivering the results that are important to you.

Scope

If you modify the query field in the JDBC Input step to add additional conditional clauses you can reduce the scope of a forecast. The provided samples do forecasting for all systems in a single forecasting run. This means you only need to set up a single batch job to schedule it once a month. By restricting the query you can restrict forecasting to only some of your systems or split the forecasting up into multiple batch jobs that you can run on different nights.

If, for example, you had four SYSPLEXs - (PLEXA, PLEXB, PLEXC and PLEXD) - you could split the forecasting into two batch jobs by taking a copy of the forecasting configuration and adding an AND SYSPLEX IN ('PLEXA', 'PLEXB') clause to one query and an AND SYSPLEX IN ('PLEXC', 'PLEXD') to the other. You would then need to create and schedule a batch job to run each forecasting configuration.

Focus

You can set up the Capacity Planning SHIFT definitions in a look up table. This allows you to focus on the specific time of day and day of week when you are most concerned about your capacity. Refer to ["CP_SHIFT lookup table"](#) on page 9.

After you set this up, you can add an AND SHIFT = 'FOCUS' (where FOCUS is the name you used) to the query and forecasts will only be made for your focus SHIFT. Depending on how small your focus shift is, this can drastically reduce the number of forecasts you need to make.

When selecting a focus shift, select a time range that usually has high usage for the workloads you are concerned about and within which the usage is fairly consistent. IBM Z Performance and Capacity Analytics projects the Average CPU usage, so if you focus shift has an equal mixture of 40% and 80% usage hours, then 60% will be the projected value. If you reduce the size of the focus shift so it just contains the 80% usage hours, 80% will be projected when you're using the machine the most heavily and it's capacity is the most critical. The reduction in the variation amongst the input values also reduces the uncertainty of the forecast output.

This approach can also be used to eliminate weekend data from forecasts. Adding a simple AND SHIFT IN ('BATCH', 'ON LINE') to the query will do this.

Data

By modifying the query in the JDBC query step you can increase or decrease the amount of source data that is used to make each forecast. The default for the weekly forecast is , which will generally pull in 26 weeks worth of data after you have gathered it. Changing it to CURRENT_DATE - 12 MONTHS will pull in 52 weeks worth of data.

When deciding how much data to use, consider the following:

- Doubling the number of input data points gives about an order of magnitude increase in the CPU cost of running the forecast.

- More input data does not, necessarily, make for a better forecast. It is more important to get an input range that is representative of how you expect the system to behave for the forecast period. The longer the time range you specify, the more likely you are to include data from before some transformative change on the system. This will make forecasts less accurate as the data from before the change is no longer representative of the system's current behavior.

Projecting into the future

The length of time a forecast goes into the future is controlled by the "future" parameter on the forecast step. This is the percentage of the input time range that the forecast is extrapolated for. The default value is 50% (entered without the % sign), so input data covering a 6 month time span will end up producing a forecast for those 6 months plus an additional 3 months into the future. By increasing this number you can increase the time the forecast goes into the future.

- There's no significant CPU cost to increasing the number.
- The accuracy of all forecasts should be considered to drop the further it gets into the future as the system is more likely to deviate from the behavioral patterns that are observed in the input data.
- A single transformative change can render all forecasts inaccurate and the further you look into the future, the more likely you are to encounter one.
- Once the projection exceeds 100% it starts to look less and less tenable. Although there are circumstances where it will be right, for a processor running general workload the accuracy of any forecast beyond 100 - 200% of the input time range becomes more and more questionable, especially if the input time range is short.
- Forecasts based on relatively short input ranges are very sensitive to minor fluctuations in the input data and projecting it over longer periods of time amplifies those inaccuracies.

CPU_LPAR Forecasts

When RMF runs on a z/OS system and collects LPAR performance data, it gathers it for all of the LPARs on the CPC that it is running on. So, if you have 3 z/OS LPARs running on the same CPC, then each one will gather a full set of LPAR data for all of the LPARs on the CPC. The MVS_SYSTEM_ID value you specify in the reports is to select the system whose gathered data will be used for the various LPARs on the CPC.

- The LPAR data gathered by each system is separately forecast. For example, if you have 3 systems on the same CPC and the forecast is restricted to only 1 of the systems, then you get a 60% reduction in the cost of running the CPU_LPAR forecasts for that CPC. When you do that, however, you must use that MVS_SYSTEM_ID to view reports including forecast data.
- If you have some LPARs that are always running on the same CPC and which have continuous availability, you could choose to do the forecasting only for them (and not for any of the other LPARs on the same CPC).
- If a system is down, it cannot gather data. this can skew some aggregated weekly values. Ideally you want to use data gathered by a system that's up all the time for forecasting.
- If a system moves to an LPAR on a different CPC, then the LPAR data it's gathered is going to be for a mixture of both CPC's. Forecasts based on this data are likely to be highly inaccurate as the data is both incomplete and muddled up. If you have some systems that are expected to be moved, you might want to simply not use their data for forecasting.

Supplied forecasts

IBM Z Performance and Capacity Analytics includes some supplied forecasts that you can use by configuring the associated properties files. These will help you get started with forecasting.

These supplied forecasts are designed to be used with the supplied capacity planning reports.

Forecasting configurations are supplied for both weekly and daily forecasting. As forecasting is a CPU intensive activity, it is recommended that you initially just do weekly forecasting on a monthly basis.

Forecast configurations

FC_CHANNEL_W_6M

This forecasts peak channel usage by Channel using weekly data gathered over the last six months. The forecast goes about 3 months into the future.

FC_CICS_REGN_W_6M

This forecasts takes the data from the CP_CICS_REGN_W table which contain aggregated weekly data for CICS Regions.

It uses the last 6 months' worth of data to generate a linear forecast for transaction volume and CPU usage out for another 3 months.

FC_CPU_LPAR_H_3M_PROFILE_WH

This forecasts takes data from the CP_CPU_LPAR_H table which contains aggregated Hourly usage data for LPARS.

It uses the last 3 months worth of data to generate a week by hour Profile for average CPU usage.

FC_CPU_LPAR_H_3M_WH

This forecasts takes data from the CP_CPU_LPAR_H table which contains aggregated Hourly usage data for LPARS.

It uses the last 3 months worth of data to generate a week by hour forecast for average CPU usage.

FC_CPU_LPAR_W_6M

This forecasts averages CPU consumption by LPAR using weekly data gathered over the last six months. The forecast goes about 3 months into the future.

FC_CPU_WKLD_W_6M

This forecasts takes data from the CP_CPU_WKLD_W table which contains aggregated weekly usage data for WORKLOADs.

It uses the last 6 months' worth of data to generate a linear forecast for average CPU usage out for another 3 months.

FC_DISK_M_1Y

This forecasts takes data from the DFSMS_VOLUME_M table which contains aggregated monthly disk usage data.

It uses the last year's worth of data to generate a linear forecast for channel usage out for another 6 months.

FC_MSU_CON_LPAR_M_1Y

This forecasts takes data from the CP_CPU_LPAR_M table which contains aggregated monthly usage data for LPARS.

It uses the last year's worth of data to generate a linear forecast for average MSU Consumption out for another 6 months.

FC_STORAGE_W_6M

This forecasts takes data from the CP_STORAGE_W table which contains aggregated weekly storage usage data.

It uses the last 6 months' worth of data to generate a linear forecast for storage usage out for another 12-13 weeks.

FC_TAPE_W_6M

This forecasts takes data from the CP_TAPE_W table which contains aggregated weekly disk usage data.

It uses the last 6 months' worth of data to generate a linear forecast for channel usage out for another 12 to 13 weeks.

The naming scheme for these is as follows: *FC_forecast_unit_qualifier_timespan* where:

- *forecast* is what is being forecast, such as CPU_LPAR, CPU_WKLD, STORAGE, CHANNEL
- *unit* is the unit and data used for forecasting such as **H**ourly, **D**aily, **W**eekly, **M**onthly, **Y**early

- *qualifier* is the (optional) special qualifier - S for Service Classes, R for Report Classes

Tailoring a supplied forecast

To use the supplied forecasts that come with IBM Z Performance and Capacity Analytics you must configure the associated .properties file for each forecast that you want to use. Forecasts are run as a batch job on your IBM Z Performance and Capacity Analytics Hub system.

Procedure

1. From the `IDSz/forecaster/config` directory, open the .properties file for the forecast that you want to configure.
2. Edit the .properties file to match your environment.

input.1.connection

The localhost and the port number that the Db2 holding the IBM Z Performance and Capacity Analytics scheme you want to forecast against has open for JDBC connections.

input.1.location

The location value that you specified for the Db2 Subsystem when it was configured.

input.1.schema

The schema that is holding the IBM Z Performance and Capacity Analytics tables you want to use for forecasting.

process.1.1.algorithm

The type of forecasting algorithm you wish to use, current algorithm options are MEAN and LINEAR.

process.1.1.inherits

The number of columns you wish to inherit from your base table.

process.1.1.inherit.1

The first column to inherit.

process.1.1.field

The field to forecast on.

process.1.1.span

The span (into the future) the forecaster should go to in weeks.

process.1.1.cf

Confidence factor of the forecast. Default is 2 (95%).

process.1.2.schema

The name of the schema that is holding the IBM Z Performance and Capacity Analytics tables you want to use for forecasting.

process.1.2.update

A value of 'yes' will generate a Db2 MERGE statement. A value of 'no' (the default) will generate an INSERT instead. This is used in the PROFILE forecast.

output.1.1.connection

The localhost and the port number that the Db2 holding the IBM Z Performance and Capacity Analytics scheme you wish to forecast against has open for JDBC connections.

output.1.1.location

The localhost and the port number that the Db2 holding the IBM Z Performance and Capacity Analytics scheme you wish to forecast against has open for JDBC connections.

Chapter 7. Data table reference

The Capacity Planning for z/OS component has various data tables within each subcomponent.

Naming conventions for data tables

The names of the data tables use this format:

CP_tablename_suffix

Where:

- **suffix** indicates the summarization level of the data in the table, or whether the table is populated by the Forecaster. A table name can have these suffixes:

A table name can have these summarization-level suffixes:

- _F**
The table is populated by the Forecaster, and holds forecasting data
- _T**
The table holds nonsummarized data (timestamped data)
- _H**
The table holds data summarized by hour (hourly data)
- _D**
The table holds data summarized by day (daily data)
- _W**
The table holds data summarized by week (weekly data)

Table descriptions

Each table description includes information about the table, a description of each of the key columns, and a description of each of the data columns:

Key columns

Marked K and form the primary key. They are sorted in the sequence they appear in the table.

Data columns

They follow the last key column and are sorted in alphabetic order with the underscore ignored.

Related concepts

[“Data tables” on page 4](#)

IBM Z Performance and Capacity Analytics has various Db2 data tables within each subcomponent. The data tables contain capacity planning and forecasting data collected, generated and used for analysis and reporting by the Capacity Planning for z/OS component.

CPU subcomponent data tables

The CPU subcomponent data tables contain CPU usage information.

CP_CPU_LPAR_T

This table provides timestamp-level capacity planning metrics on logical partitions and processor activity. It contains data from SMF type 70 records processed through a record procedure (DRL2S070).

This table is updated by the CP_SHIFT_NM lookup table.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started. From SMF70DAT.
TIME	K	TIME	Time that the RMF measurement interval started. From SMF70IST.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF70INT
CPU_USED		FLOAT	Total processor dispatch time during the RMF interval, in seconds. Calculated as the sum of SMF70PDT/1000000.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.

Column name	Primary key	Data type	Column description
LPAR_PHYS_PROC		FLOAT	Total number of physical processors, of this processor type, assigned to this logical partition. This number is based on the number of dedicated processors assigned to the partition, or on the weight of the shared processors allocated to the partition.

CP_CPU_LPAR_H

This table provides capacity planning metrics on logical partitions and processor activity, aggregated to an hourly level. It contains data from SMF type 70 records processed through a record procedure (DRL2S070).

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started. From SMF70DAT.
TIME	K	TIME	Time that the RMF measurement interval started, rounded down to the nearest hour. From SMF70IST.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.

Column name	Primary key	Data type	Column description
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF70INT
CPU_USED_AVG		FLOAT	Average processor dispatch time during the hour, in seconds.
CPU_PEAK		FLOAT	Peak processor dispatch time during the hour, in seconds.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
LPAR_PHYS_PROC		FLOAT	Total number of physical processors, of this processor type, assigned to this logical partition. This number is based on the number of dedicated processors assigned to the partition, or on the weight of the shared processors allocated to the partition.
NUM_RECORDS		FLOAT	Number of records aggregated to this interval.

CP_CPU_LPAR_D

This table provides capacity planning metrics on logical partitions and processor activity, aggregated to a daily level. It contains data from SMF type 70 records processed through a record procedure (DRL2S070).

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started. From SMF70DAT.

Column name	Primary key	Data type	Column description
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF70INT
CPU_USED_AVG		FLOAT	Average processor dispatch time during the day, in seconds.
CPU_PEAK_AVG		FLOAT	Average peak processor dispatch time during the day, in seconds.
CPU_PEAK		FLOAT	Peak processor dispatch time during the day, in seconds.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
LPAR_PHYS_PROC		FLOAT	Total number of physical processors, of this processor type, assigned to this logical partition. This number is based on the number of dedicated processors assigned to the partition, or on the weight of the shared processors allocated to the partition.

Column name	Primary key	Data type	Column description
NUM_RECORDS		FLOAT	Number of records aggregated to this interval.

CP_CPU_LPAR_W

This table provides capacity planning metrics on logical partitions and processor activity, aggregated to a weekly level. It contains data from SMF type 70 records processed through a record procedure (DRL2S070).

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started, grouped by first day of the week. From SMF70DAT.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF70INT
NUM_RECORDS		FLOAT	Number of records aggregated to this interval.
CPU_USED_AVG		FLOAT	Average processor dispatch time during the week, in seconds.

Column name	Primary key	Data type	Column description
CPU_PEAK_AVG		FLOAT	Average peak processor dispatch time during the week, in seconds.
CPU_PEAK		FLOAT	Peak processor dispatch time during the week, in seconds.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
LPAR_PHYS_PROC		FLOAT	Total number of physical processors, of this processor type, assigned to this logical partition. This number is based on the number of dedicated processors assigned to the partition, or on the weight of the shared processors allocated to the partition.

CP_CPU_LPAR_F

This table contains forecasted CPU metrics, generated by the Forecaster, on logical partition and processor level.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started.
AGGR_LEVEL	K	CHAR(1)	Aggregation level of the data contained in this table. Possible values are: D - Daily and W - Weekly.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.

Column name	Primary key	Data type	Column description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
F_CPU_USED_AVG		FLOAT	Forecasted average CPU seconds.
FV_CPU_USED_AVG		FLOAT	Standard variance calculated by the Forecaster.

CP_CPU_LPAR_FH

This table contains forecasted CPU metrics, generated by the Forecaster, on logical partition and processor level.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started.
TIME	K	INTEGER	Time when the RMF measurement interval started.
AGGR_LEVEL	K	CHAR(1)	Aggregation level of the data contained in this table. Possible values are: D - Daily and W - Weekly.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.

Column name	Primary key	Data type	Column description
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	CHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	CHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast. (When it was run)
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF72INT.
F_CPU_USED_AVG		FLOAT	Forecasted average CPU seconds.
FV_CPU_USED_AVG		FLOAT	Standard variance calculated by the Forecaster.
C_FACTOR		FLOAT	Confidence factor
FC_FACTOR_HIGH		FLOAT	High boundary for the field at the specified confidence factor.
FC_FACTOR_LOW		FLOAT	Low boundary for the field at the specified confidence factor.

CP_CPU_LPAR_PH

This table contains forecasted CPU metrics, generated by the Profiler, on logical partition and processor level.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date to which the profile applies.
TIME	K	INTEGER	Time interval to which the profile applies. (ie. Start Hour)

Column name	Primary key	Data type	Column description
PROFILE_NAME	K	VARCHAR(20)	Name of the profile which has been generated.
SHIFT	K	VARCHAR(10)	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
SYSPLEX_NAME	K	VARCHAR(8)	Name of the sysplex. From SMF70XNM.
MVS_SYSTEM_ID	K	VARCHAR(4)	MVS system ID. From SMF70SID.
LPAR_NAME	K	VARCHAR(8)	Name of the logical partition. From SMF70LPM.
PROCESSOR_TYPE	K	VARCHAR(4)	Logical processor type. Possible values are: CP,ICF+(includes IFA,IFL,ICF),IFA,IFL,ICF,II P. From SMF70CIX.
TOT_PHYS_PROC		FLOAT	Number of physical processors, used to determine MIPS.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF72INT.
WEEKDAY		CHAR(3)	Abbreviated day of the week (MON, TUE, WED, THU, FRI, SAT, SUN)
P_ALGORITHM	K	VARCHAR(10)	Algorithm used by the Profiler.
P_CPU_USED_AVG		FLOAT	Forecasted profile CPU seconds for the given hour.
P_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast. (When it was run)
PV_CPU_USED_AVG		FLOAT	Standard variance calculated by the Profiler.
P_C_FACTOR		FLOAT	Confidence factor
PC_FACTOR_HIGH		FLOAT	High boundary for the field at the specified confidence factor.

Column name	Primary key	Data type	Column description
PC_FACTOR_LOW		FLOAT	Low boundary for the field at the specified confidence factor.

CP_LSPR_WKLD_H

This table provides CPU Measurement Facility metrics on hourly level, calculated from table KPMZ_CPUMFM_PT_T, which is sourced from SMF type 113 subtype 2 records.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the SMF record was created. From SMF113DTE.
TIME	K	TIME	Time when the SMF record was created, rounded down to the nearest hour. From SMF113TME.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF113ID.
PROCESSOR_TYPE	K	CHAR(4)	he processor type. Possible values are CP, IFA or IIP. From SMF113_2_CPU_CLS.
LSPR_LOW_CNT		FLOAT	The number of occurrences when the LSPR Workload was LOW.
LSPR_AVG_CNT		FLOAT	The number of occurrences when the LSPR Workload was AVG.
LSPR_HIGH_CNT		FLOAT	The number of occurrences when the LSPR Workload was HIGH.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
CPU_MODEL_NO		SMALLINT	CPU model number.
CPU_SUB_MODEL		CHAR(16)	CPU submodel number.

CP_LSPR_WKLD_D

This table provides CPU Measurement Facility metrics on daily level, calculated from table KPMZ_CPUMFM_PT_T, which is sourced from SMF type 113 subtype 2 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the SMF record was created. From SMF113DTE.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF113ID.
PROCESSOR_TYPE	K	CHAR(4)	The processor type. Possible values are CP, IFA or IIP. From SMF113_2_CPU_CLS.
LSPR_LOW_CNT		FLOAT	The number of occurrences when the LSPR Workload was LOW.
LSPR_AVG_CNT		FLOAT	The number of occurrences when the LSPR Workload was AVG.
LSPR_HIGH_CNT		FLOAT	The number of occurrences when the LSPR Workload was HIGH.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
CPU_MODEL_NO		SMALLINT	CPU model number.
CPU_SUB_MODEL		CHAR(16)	CPU submodel number.

CP_LSPR_WKLD_W

This table provides CPU Measurement Facility metrics on weekly level, calculated from table KPMZ_CPUMFM_PT_T, which is sourced from SMF type 113 subtype 2 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the SMF record was created, grouped by the first day of the week. From SMF113DTE.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF70DAT and SMF70IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF113ID.
PROCESSOR_TYPE	K	CHAR(4)	The processor type. Possible values are CP, IFA or IIP. From SMF113_2_CPU_CLS.
LSPR_LOW_CNT		FLOAT	The number of occurrences when the LSPR Workload was LOW.
LSPR_AVG_CNT		FLOAT	The number of occurrences when the LSPR Workload was AVG.
LSPR_HIGH_CNT		FLOAT	The number of occurrences when the LSPR Workload was HIGH.
TOT_PHYS_PROC		FLOAT	Total number of physical processors available, of this processor type.
CPU_MODEL_NO		SMALLINT	CPU model number.
CPU_SUB_MODEL		CHAR(16)	CPU submodel number.

CP_CPU_WKLD_H

This table provides capacity planning metrics on service and report classes, aggregated by hourly level.

It contains data from SMF type 72 subtype 3 records.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF72DTE.
TIME	K	TIME	Time when the record is written, rounded down to the nearest hour. From SMF72IST.

Column name	Primary key	Data type	Column description
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF72DAT and SMF72IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF72SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF72XNM.
POLICY	K	CHAR(8)	Service policy name. From R723MNSP.
BUSINESS_APPL	K	CHAR(10)	Business application name, as derived from lookup table CP_BUSINESS_APPL.
WORKLOAD_NAME	K	CHAR(8)	Workload name. From R723MWNM.
SERVICE_CLASS	K	CHAR(8)	Service or report class name. From R723MCNM.
SERV_CLASS_PERIOD	K	CHAR(1)	Service or report class period number. From R723CPER.
SERV_OR_RPRT_FL	K	CHAR(1)	Service or Report Class. Possible values: S - Service Class, R - Report class.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF72INT.
CPU_PEAK		FLOAT	Peak processor dispatch time during the hour, in seconds.
CPU_USED_AVG		FLOAT	Average processor dispatch time during the hour, in seconds.
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_CPU_WKLD_D

This table provides capacity planning metrics on service and report classes, aggregated by daily level.

It contains data from SMF type 72 subtype 3 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF72DTE.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF72DAT and SMF72IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF72SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF72XNM.
POLICY	K	CHAR(8)	Service policy name. From R723MNSP.
BUSINESS_APPL	K	CHAR(10)	Business application name, as derived from lookup table CP_BUSINESS_APPL.
WORKLOAD_NAME	K	CHAR(8)	Workload name. From R723MWNM.
SERVICE_CLASS	K	CHAR(8)	Service or report class name. From R723MCNM.
SERV_CLASS_PERIOD	K	CHAR(1)	Service or report class period number. From R723CPER.
SERV_OR_RPRT_FL	K	CHAR(1)	Service or Report Class. Possible values: S - Service Class, R - Report class.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF72INT.
CPU_PEAK		FLOAT	Peak processor dispatch time during the day, in seconds.
CPU_USED_AVG		FLOAT	Average processor dispatch time during the day, in seconds.
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_CPU_WKLD_W

This table provides capacity planning metrics on service and report classes, aggregated by weekly level.

It contains data from SMF type 72 subtype 3 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written, grouped by the first day of the week. From SMF72DTE.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF72DAT and SMF72IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF72SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF72XNM.
POLICY	K	CHAR(8)	Service policy name. From R723MNSP.
BUSINESS_APPL	K	CHAR(10)	Business application name, as derived fromlookup table CP_BUSINESS_APPL.
WORKLOAD_NAME	K	CHAR(8)	Workload name. From R723MWNM.
SERVICE_CLASS	K	CHAR(8)	Service or report class name. From R723MCNM.
SERV_CLASS_PERIOD	K	CHAR(1)	Service or report class period number. From R723CPER.
SERV_OR_RPRT_FL	K	CHAR(1)	Service or Report Class. Possible values: S - Service Class, R - Report class.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF72INT.
CPU_PEAK		FLOAT	Peak processor dispatch time during the week, in seconds.

Column name	Primary key	Data type	Column description
CPU_USED_AVG		FLOAT	Average processor dispatch time during the week, in seconds.
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_CPU_WKLD_F

This table contains forecast CPU metrics, generated by the Forecaster, on workload and service or report class level.

It contains data from SMF type 72 subtype 3 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the RMF measurement interval started.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF72DAT and SMF72IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF72SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF72XNM.
POLICY	K	CHAR(8)	Service policy name. From R723MNSP.
BUSINESS_APPL	K	CHAR(10)	Business application name, as derived from lookup table CP_BUSINESS_APPL.
WORKLOAD_NAME	K	CHAR(8)	Workload name. From R723MWNM.
SERVICE_CLASS	K	CHAR(8)	Service or report class name. From R723MCNM.
SERV_CLASS_PERIOD	K	CHAR(1)	Service or report class period number. From R723CPER.

Column name	Primary key	Data type	Column description
SERV_OR_RPRT_FL	K	CHAR(1)	Service or Report Class. Possible values: S - Service Class, R - Report class.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
F_CPU_USED_AVG		FLOAT	Forecasted average CPU seconds.
FV_CPU_USED_AVG		FLOAT	Standard variance calculated by the Forecaster.

Storage subcomponent data tables

The Storage subcomponent data tables contain storage availability information.

CP_STORAGE_H

This table provides capacity planning metrics on storage, aggregated by hourly level.

It contains data from SMF type 71 records.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF71DAT.
TIME	K	TIME	Time when the record is written, rounded down to the nearest hour. From SMF71IST.
SYSPLEX_NAME	K	CHAR(8)	SYSPLEX_NAME. From SMF71XNM.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF71DAT and SMF71IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF71SID.

Column name	Primary key	Data type	Column description
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF71INT.
CSTOR_AVLBL_AVG		FLOAT	Central storage average available frames. This is the average of SMF71AVF.
CSTOR_AVLBL_MIN		FLOAT	Central storage minimum available frames. This is the minimum of SMF71MNF.
CSTOR_TOTAL_AVG		FLOAT	Central storage average total frames. This is the average of (SMF71AVT +SMF71FIN).
CSTOR_OK_THRESH		FLOAT	Central storage OK Threshold. This value indicates the number of frames on the available frame queue when stealing ends. This is the average of (SMF71AVF-SMF71CAA).
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_STORAGE_D

This table provides capacity planning metrics on storage, aggregated by daily level.

It contains data from SMF type 71 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF71DAT.
SYSPLEX_NAME	K	CHAR(8)	SYSPLEX_NAME. From SMF71XNM.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF71DAT and SMF71IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF71SID.

Column name	Primary key	Data type	Column description
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF71INT.
CSTOR_AVLBL_AVG		FLOAT	Central storage average available frames. This is the average of SMF71AVF.
CSTOR_AVLBL_MIN		FLOAT	Central storage minimum available frames. This is the minimum of SMF71MNF.
CSTOR_TOTAL_AVG		FLOAT	Central storage average total frames. This is the average of (SMF71AVT +SMF71FIN).
CSTOR_OK_THRESH		FLOAT	Central storage OK Threshold. This value indicates the number of frames on the available frame queue when stealing ends. This is the average of (SMF71AVF-SMF71CAA).
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_STORAGE_W

This table provides capacity planning metrics on storage, aggregated by weekly level.

It contains data from SMF type 71 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written, and grouped by the first day of the week. From SMF71DAT.
SYSPLEX_NAME	K	CHAR(8)	SYSPLEX_NAME. From SMF71XNM.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF71DAT and SMF71IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.

Column name	Primary key	Data type	Column description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF71SID.
MEASURED_SEC		FLOAT	Length of the RMF interval, in seconds. From SMF71INT.
CSTOR_AVLBL_AVG		FLOAT	Central storage average available frames. This is the average of SMF71AVF.
CSTOR_AVLBL_MIN		FLOAT	Central storage minimum available frames. This is the minimum of SMF71MNF.
CSTOR_TOTAL_AVG		FLOAT	Central storage average total frames. This is the average of (SMF71AVT +SMF71FIN).
CSTOR_OK_THRESH		FLOAT	Central storage OK Threshold. This value indicates the number of frames on the available frame queue when stealing ends. This is the average of (SMF71AVF-SMF71CAA).
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_STORAGE_F

This table contains forecasted storage metrics, generated by the Forecaster.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written.
SYSPLEX_NAME	K	CHAR(8)	Sysplex name. From SMF71XNM.
AGGR_LEVEL	K	CHAR(1)	Aggregation level of the data contained in this table. Possible values are: D - Daily and W - Weekly.

Column name	Primary key	Data type	Column description
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF71DAT and SMF71IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF71SID.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
F_CSTOR_AVLBL_AVG		FLOAT	Forecasted average available frames.
FV_CSTOR_AVLBL_AVG		FLOAT	Standard variance calculated by the Forecaster.

Channel subcomponent data tables

The Channel subcomponent data tables contain channel utilization information.

CP_CHANNEL_H

This table provides capacity planning metrics on channels, aggregated by hourly level.

It contains data from SMF type 73 records.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF73DTE.
TIME	K	TIME	Time when the record is written, rounded down to the nearest hour. From SMF73IST.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF73DAT and SMF73IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.

Column name	Primary key	Data type	Column description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF73SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF73XNM.
CHANNEL	K	CHAR(7)	Channel path ID and channel subsystem number.
CHANNEL_TYPE		CHAR(8)	Type of channel. From SMF73ACR.
CHAN_BUSY_AVG		FLOAT	Channel average total busy, calculated as the average of SMF73BSY*100/SMF73SMP.
CHAN_BUSY_PEAK		FLOAT	Channel peak total busy, calculated as the maximum of SMF73BSY*100/SMF73SMP.
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_CHANNEL_D

This table provides capacity planning metrics on channels, aggregated by Daily level.

It contains data from SMF type 73 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written. From SMF73DTE.
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF73DAT and SMF73IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF73SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF73XNM.

Column name	Primary key	Data type	Column description
CHANNEL	K	CHAR(7)	Channel path ID and channel subsystem number.
CHANNEL_TYPE		CHAR(8)	Type of channel. From SMF73ACR.
CHAN_BUSY_AVG		FLOAT	Channel average total busy, calculated as the average of SMF73BSY*100/ SMF73SMP.
CHAN_BUSY_PEAK		FLOAT	Channel peak total busy, calculated as the maximum of SMF73BSY*100/ SMF73SMP.
NUM_RECORDS		INTEGER	Number of records aggregated to this interval.

CP_CHANNEL_W

This table provides capacity planning metrics on channels, aggregated by weekly level.

It contains data from SMF type 73 records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when the record was written, grouped by the first day of the week
SHIFT	K	CHAR	Name of the capacity planning shift. This is derived using fields SMF73DAT and SMF73IST from the record as parameters to look up a defined SHIFT in lookup table CP_SHIFT_NM.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. From SMF73SID.
SYSPLEX_NAME	K	CHAR(8)	Name of the sysplex. From SMF73XNM.
CHANNEL	K	CHAR(7)	Channel path ID and channel subsystem number.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.

Column name	Primary key	Data type	Column description
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
CHANNEL_TYPE		CHAR(8)	
F_CHAN_BUSY_PEAK		FLOAT	Forecasted channel peak total busy.
FV_CHAN_BUSY_PEAKS		FLOAT	Standard variance calculated by the Forecaster.

Disk subcomponent data tables

The Disk subcomponent data tables contain disk storage information.

DFSMS_VOLUME_D

This table provides daily statistics on DASD volumes. It contains data written by the DFSMS DCOLLECT facility as type V records. Use this table for trend reporting on the overall utilization of DASD space.

The default retention periods for this table is 30 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when DCOLLECT was run against the volume. From DCUDATE.
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. This is the ID of the system running DCOLLECT. From DCUSYSID.
STORAGE_GROUP	K	VARCHAR(30)	Storage group name for the volume. From DCVSTGGP or set to the character string NON-SMS.
VOLSER	K	CHAR(6)	Volume serial number. From DCVVOLSR.
CAPACITY_TOTAL		FLOAT	Total capacity of the volume, in kilobytes. From DCVVLCAP.
DS_CTRL_BLOCKS		FLOAT	Total number of free data set control blocks in the VTOC. From DCVFDSCB.
EXTENT_MAX		FLOAT	Largest extent in the volume, in kilobytes. From DCVLGEXT.
EXTENTS_FREE		FLOAT	Total number of free extents in the volume. From DCVFREXT.

Column name	Primary key	Data type	Column description
FRAGMENT_INDEX		FLOAT	Fragmentation index of the volume. From DCVFRAGI.
SPACE_ALLOC_TOTAL		FLOAT	Total amount of space allocated to the data sets, in KB. For non-VSAM data sets, this is the sum of DCDALLSP (or DCDALLSX if the DCDFLAG5 bit is set). For VSAM data sets, this is instead calculated as the sum of DCAHARB/1024, where DCAHARB is from DCAHARBA or DCAHARBC For VSAM data sets, this is instead calculated as the sum of DCAHARB/1024, where DCAHARB is from DCAHARBA or DCAHARBC in extended VSAM data sets.
SPACE_FREE_PCT		FLOAT	Total amount of free space in the volume, as a percentage of the total volume capacity. From DCVPERCT.
SPACE_FREE_TOTAL		FLOAT	Total amount of free space in the volume, in kilobytes. From DCVFRESP.
VTOC_INDEX_RECORDS		FLOAT	Total number of free VTOC index records in the volume. From DCVFVIRS.

DFSMS_VOLUME_M

This table provides monthly statistics for DASD volumes. It contains data written by the DFSMS DCOLLECT facility as type V records. Use this table for trend reporting on the overall utilization of DASD space.

The default retention period for this table is 765 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when DCOLLECT was run against the volume. This is the date of the first day of the month. From DCUDATE.

Column name	Primary key	Data type	Column description
MVS_SYSTEM_ID	K	CHAR(4)	MVS system ID. This is the ID of the system running DCOLLECT. From DCUSYSID.
STORAGE_GROUP	K	VARCHAR(30)	Storage group name for the volume. From DCVSTGGP or set to the character string NON-SMS.
VOLSER	K	CHAR(6)	Volume serial number. From DCVVOLSR.
CAPACITY_TOTAL		FLOAT	Total capacity of the volume, in kilobytes. From DCVLCAP
DS_CTRL_BLOCKS_AVG		FLOAT	Average number of free data set control blocks in the VTOC.
EXTENT_MAX		INTEGER	Largest extent in the volume, in kilobytes. This is the maximum of EXTENT_MAX in the DFSMS_VOLUME_D table.
EXTENTS_FREE_AVG		FLOAT	Average number of free extents in the volume.
FRAGMENT_INDEX_AVG		FLOAT	Average fragmentation index of the volume.
SPACE_ALLOC_AVG		FLOAT	Average amount of allocated space in the volume, in kilobytes.
SPACE_ALLOC_MAX		INTEGER	Maximum amount of allocated space in the volume, in kilobytes. This is the maximum of SPACE_ALLOC_TOTAL in the DFSMS_VOLUME_D table.
SPACE_ALLOC_MIN		INTEGER	Minimum amount of space allocated to a data set, in KB. This is the minimum of DCDALLSP, or DCDALLSX if the DCDFLAG5 bit is set.
SPACE_FREE_AVG		FLOAT	Average amount of free space in the volume, in kilobytes.

Column name	Primary key	Data type	Column description
SPACE_FREE_AVG_PCT		FLOAT	Average amount of free space in the volume, as a percentage of the total volume capacity.
SPACE_FREE_MAX_PCT		INTEGER	Maximum amount of free space in the volume, as a percentage of the total volume capacity. Calculated as the maximum of SPACE_FREE_PCT in the DFSMS_VOLUME_D table.
SPACE_FREE_MIN_PCT		INTEGER	Minimum amount of free space in the volume, as a percentage of the total volume capacity. Calculated as the minimum of SPACE_FREE_PCT in the DFSMS_VOLUME_D table.
VOL_COUNT		INTEGER	Number of times the volume has been referenced. This is the count of DCVVOLSR. This is the counter used for calculating averages in this table.
VTOC_INDEX_REC_AVG		FLOAT	Average number of free VTOC index records in the volume.

CP_DISK_F

This table provides monthly forecast metrics on DASD volumes.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date when DCOLLECT was run against the volume. From DCUDATE.
AGGR_LEVEL	K	CHAR(1)	Aggregation level of the data contained in this table. Possible values are: D - Daily and W - Weekly.
MVS_SYSTEM_ID	K	CHAR(8)	MVS system ID. This is the ID of the system running DCOLLECT. From DCUSYSID.

Column name	Primary key	Data type	Column description
STORAGE_GROUP	K	CHAR(8)	Storage group name for the volume. From DCVSTGGP or set to the character string NON-SMS.
VOLSER	K	CHAR(6)	Volume serial number. From DCVVOLSR.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
F_SPACE_ALLOC_MAX		FLOAT	Forecasted amount of allocated space in the volume, in kilobytes.
FV_SPACE_ALLOC_MAX		FLOAT	Standard variance allocated space calculated by Forecaster.

Tape subcomponent data tables

The Tape subcomponent data tables contain tape storage information.

CP_TAPE_D

This table provides daily capacity planning metrics from RRM extract file dataset records.

The default retention period for this table is 60 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date set by pre-processor.
MVS_SYSTEM_ID	K	CHAR(8)	MVS system ID from RVCRSID.
STORAGE_GROUP	K	CHAR(8)	SMS storage group name from RDSGNAME.
KB_USED		FLOAT	Total storage used, in kilobytes. Calculated as the sum of RDDSSIZE.

CP_TAPE_W

This table provides weekly capacity planning metrics from RRM extract file dataset records.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date rounded to the week.

Column name	Primary key	Data type	Column description
MVS_SYSTEM_ID	K	CHAR(8)	MVS system ID from RVCRSID.
STORAGE_GROUP	K	CHAR(8)	SMS storage group name from RDSGNAME.
KB_USED		FLOAT	Total storage used, in kilobytes. Calculated as the sum of RDDSSIZE.

CP_TAPE_F

This table contains forecasted tape metrics, generated by the Forecaster.

The default retention period for this table is 366 days.

Column name	Primary key	Data type	Column description
DATE	K	DATE	Date set by pre-processor.
AGGR_LEVEL	K	CHAR(1)	Aggregation level of the data contained in this table. Possible values are: D - Daily and W - Monthly.
MVS_SYSTEM_ID	K	CHAR(8)	MVS system ID from RVCRSID.
STORAGE_GROUP	K	CHAR(8)	SMS storage group name from RDSGNAME.
F_ALGORITHM	K	CHAR(10)	Algorithm used by the forecaster.
F_TIMESTAMP	K	TIMESTAMP	Timestamp of the forecast.
F_KB_USED		FLOAT	Forecasted storage used, in kilobytes.
FV_KB_USED		FLOAT	Standard variance calculated by the Forecaster.

Chapter 8. Report reference

IBM Z Performance and Capacity Analytics provides many types of reports that use curated historical data to provide predictive analysis both for managing future resource needs and to ensure that capacity shortfalls do not occur.

The forecasting features in the reports enable you to use data-driven practices to more accurately determine when the environment will reach a capacity threshold and require hardware upgrades. This can help you plan upgrades in a timely way that can enable you to save money and maintain business service levels.

The reports are grouped by subcomponent of the Capacity Planning for z/OS component:

- CPU
- Storage
- Channel
- Disk
- Tape

CPU subcomponent reports

The CPU reports provide key capacity planning metrics on processor activity, collected on SMF timestamp, hourly, daily, and weekly level. Data is aggregated by capacity planning shift as well as CEC level, LPAR level, Workload level, Service or Report Class level, and Business Application level.

The CPU reports are grouped by report type:

- CPU usage
- CPU forecasting
- LSPR workload
- Software cost analysis

Understanding MIPS calculations in the CPU subcomponent reports

A control table is used by the CPU subcomponent for MIPS calculations. This is the LSPR_MIPS control table and must be initialized before use with parameters specific to your installation. Refer to [“Step 1.4: Configuring control tables” on page 11](#) for more information.

MIPS Capacity

The CEC MIPS capacity is calculated using the MIPS rating from the LSPR_MIPS control table. The LSPR_MIPS control table contains MIPS ratings for the following LSPR Workload categories: LOW, AVG, and HIGH. The LSPR category used for MIPS calculations is based on the most popular LSPR Workload during the reporting period. For example, if most of the SMF intervals in the reporting period reported an LSPR Workload of HIGH, then the MIPS rating for the LSPR HIGH Workload will be used when calculating MIPS capacity on the report.

The LPAR MIPS capacity is calculated as a percentage of the CEC MIPS capacity, and this percentage is calculated using the number of physical processors in the LPAR as well as the weight distribution of the logical processors in the LPAR.

MIPS Usage

MIPS usage is calculated on various levels, for example, CEC, LPAR, Workload, Service/Report Class or Business Application. MIPS usage is calculated as a percentage of the MIPS Capacity. This is the calculation used to calculate MIPS:

*CPU Seconds Used / Available Seconds * MIPS Capacity*

The Available Seconds is calculated as:

*Interval * Number of Physical Processors*

The MIPS Capacity is calculated using the MIPS rating from the LSPR_MIPS control table.

The LSPR_MIPS control table contains MIPS ratings for the following LSPR Workload categories: LOW, AVG and HIGH. The LSPR category used for MIPS calculations is based on the most popular LSPR Workload during the reporting period. For example, if most of the SMF intervals in the reporting period reported an LSPR Workload of HIGH, then the MIPS rating for the LSPR HIGH Workload will be used when calculating MIPS usage on the report.

CPU usage

CPU usage reports.

CP CEC MIPS Hourly/Daily/Weekly report

CP CEC MIPS usage reports show MIPS statistics on CEC level on an hourly, daily or weekly basis.

This information identifies the report:

Report ID

CPCPU01 / CPCPU05 / CPCPU06

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_HV / CP_CPU_LPAR_DV / CP_CPU_LPAR_WV

Attributes

MVS, CAPACITY, CEC, MIPS, HOURLY/DAILY/WEEKLY




Figure 8. CP CEC MIPS Weekly report

The report contains this information:

Time of day / Date

Time of day / date of the measurement.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

Threshold

The threshold of the MIPS capacity. The threshold is calculated as a percentage of the LSPR MIPS Capacity. The threshold percentage is obtained from lookup table CP_THRESHOLDS.

LSPR Capacity

This is the CEC MIPS capacity, based on the most popular LSPR Workload during the reporting period.

CP LPAR MIPS Hourly/Daily/Weekly report

This report shows MIPS statistics on LPAR level.

This information identifies the report:

Report ID

CPCPU02 / CPCPU07 / CPCPU08

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_HV / CP_CPU_LPAR_DV / CP_CPU_LPAR_WV

Attributes

MVS, CAPACITY, LPAR, MIPS, HOURLY/DAILY/WEEKLY




Figure 9. CP LPAR MIPS Weekly report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Average MIPS used by the LPAR during the interval.

Peak Used

Peak MIPS used by the LPAR during the interval.

Threshold

The threshold of the MIPS capacity. The threshold is calculated as a percentage of the LSPR MIPS Capacity. The threshold percentage is obtained from lookup table CP_THRESHOLDS.

LSPR Capacity

This is the LPAR MIPS capacity, which is calculated as a percentage of the CEC MIPS capacity, and this percentage is calculated using the number of physical processor in the LPAR as well as the weight distribution of the logical processors in the LPAR. The CEC MIPS capacity is based on the most popular LSPR Workload during the reporting period.

CP LPAR MIPS by CEC Average Hourly/Daily/Weekly report

This report shows MIPS statistics for all LPARs within a CEC.

This information identifies the report:

Report ID

CPCPU03 / CPCPU30 / CPCPU32

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_HV / CP_CPU_LPAR_DV / CP_CPU_LPAR_WV

Attributes

MVS, AVERAGE, CEC, MIPS, HOURLY/DAILY/WEEKLY




Figure 10. CP LPAR MIPS by CEC Average Daily report

The report contains this information:

Time of day / Date

Time / date of the measurement.

LPAR

Logical partition name.

Average used

Average MIPS used by each LPAR during the interval.

CP LPAR MIPS by CEC Peak Hourly/Daily/Weekly report

This report shows MIPS statistics for all LPARs within a CEC.

This information identifies the report:

Report ID

CPCPU04 / CPCPU31 / CPCPU33

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_HV / CP_CPU_LPAR_DV / CP_CPU_LPAR_WV

Attributes

MVS, PEAK, CEC, MIPS, HOURLY/DAILY/WEEKLY




Figure 11. CP LPAR MIPS by CEC Peak Weekly report

The report contains this information:

Time of day / Date

Time / date of the measurement.

LPAR

Logical partition name.

Peak Used

Peak MIPS used by each LPAR during the interval.

CP Workload MIPS Hourly/Daily/Weekly report

This report shows WLM Workload MIPS statistics.

This information identifies the report:

Report ID

CPCPU13 / CPCPU18 / CPCPU19

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_HV / CP_CPU_WKLD_DV / CP_CPU_WKLD_WV

Attributes

MVS, CAPACITY, WORKLOAD, MIPS, HOURLY/DAILY/WEEKLY




Figure 12. CP Workload MIPS Usage Daily report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

CP Serv Class MIPS Hourly/Daily/Weekly report

This report shows WLM Service or Report Class MIPS statistics.

This information identifies the report:

Report ID

CPCPU14 / CPCPU20 / CPCPU22

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_HV / CP_CPU_WKLD_DV / CP_CPU_WKLD_WV

Attributes

MVS, CAPACITY, WORKLOAD, MIPS, HOURLY/DAILY/WEEKLY




Figure 13. CP Serv Class MIPS Daily report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

CP Serv Class MIPS by Workload Avg Hourly/Daily/Weekly report

This report shows MIPS usage of all Service Classes within a WLM Workload.

This information identifies the report:

Report ID

CPCPU16 / CPCPU34 / CPCPU36

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_HV / CP_CPU_WKLD_DV / CP_CPU_WKLD_WV

Attributes

MVS, WORKLOAD, MIPS, HOURLY/DAILY/WEEKLY

CPCPU34 - CP SERV CLASS MIPS BY WORKLOAD AVG DAILY

MVS System ID: * M50 Shift: * OTHER Workload Name: * BATCH Date From: * Jun 14, 2018 Date To: * Jun 24, 2018

SERVICE CLASS MIPS WORKLOAD AVERAGE BY DAY - LSPR WKLD:AVG




Figure 14. CP Serv Class MIPS by Workload Avg Daily report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Name of the WLM Service Class.

Peak Used

Average MIPS used each WLM Service Class during the interval.

CP Serv Class MIPS by Workload Peak Hourly/Daily/Weekly report

This report shows MIPS usage of all Service Classes within a WLM Workload.

This information identifies the report:

Report ID

CPCPU17 / CPCPU35 / CPCPU37

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_HV / CP_CPU_WKLD_DV / CP_CPU_WKLD_WV

Attributes

MVS, WORKLOAD, MIPS, HOURLY/DAILY/WEEKLY

CPCPU37 - CP SERV CLASS MIPS BY WORKLOAD PEAK WEEKLY

MVS System ID: * M50 Shift: * OTHER Workload Name: * BATCH Date From: * Jun 11, 2018 Date To: * Aug 20, 2018

SERVICE CLASS MIPS WORKLOAD PEAK BY WEEK - LSPR WKLD:AVG




Figure 15. CP Serv Class MIPS by Workload Peak Weekly report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Name of the WLM Service Class.

Peak Used

Peak MIPS used each WLM Service Class during the interval.

CP Application MIPS Hourly/Daily/Weekly report

This report shows Business Application MIPS statistics. A Business Application is a grouping of Service or Report classes, as defined in lookup table CP_BUSINESS_APPL.

This information identifies the report:

Report ID

CPCPU15 / CPCPU21 / CPCPU23

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_HV / CP_CPU_WKLD_DV / CP_CPU_WKLD_WV

Attributes

MVS, CAPACITY, APPLICATION, MIPS, HOURLY/DAILY/WEEKLY

CPCPU21 - CP APPLICATION MIPS USAGE DAILY

MVS System ID: * M50
 Shift: * OTHER
 Business Application: * MOBILE
 Date From: * Jun 14, 2018
 Date To: * Jun 24, 2018

APPLICATION MIPS USAGE BY DAY - LSPR WKLD:AVG




Figure 16. CP Application MIPS Daily report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

CPU forecasting

CPU forecasting reports.

CP CEC MIPS Forecast Weekly report

This report shows MIPS statistics and forecast values on CEC level.

This information identifies the report:

Report ID

CPCPU10

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_FV

Attributes

MVS, CAPACITY, CEC, MIPS, WEEKLY, FORECAST




Figure 17. CP CEC MIPS Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Forecast

Forecast average MIPS for the CEC during the interval.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

Threshold

The threshold of the MIPS capacity. The threshold is calculated as a percentage of the LSPR MIPS Capacity. The threshold percentage is obtained from lookup table CP_THRESHOLDS.

LSPR Capacity

This is the CEC MIPS capacity, based on the most popular LSPR Workload during the reporting period.

CP LPAR MIPS Forecast Weekly report

This report shows MIPS statistics and forecast values on LPAR level.

This information identifies the report:

Report ID

CPCPU12

Report group

Capacity Planning - CPU

Source

CP_CPU_LPAR_FV

Attributes

MVS, CAPACITY, LPAR, MIPS, WEEKLY, FORECAST




Figure 18. CP LPAR MIPS Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Forecast

Forecasted average MIPS for the CEC during the interval.

Average Used

Average MIPS used by the LPAR during the interval.

Peak Used

Peak MIPS used by the LPAR during the interval.

Threshold

The threshold of the MIPS capacity. The threshold is calculated as a percentage of the LSPR MIPS Capacity. The threshold percentage is obtained from lookup table CP_THRESHOLDS.

LSPR Capacity

This is the LPAR MIPS capacity, which is calculated as a percentage of the CEC MIPS capacity, and this percentage is calculated using the number of physical processor in the LPAR as well as the weight distribution of the logical processors in the LPAR. The CEC MIPS capacity is based on the most popular LSPR Workload during the reporting period.

CP Workload MIPS Forecast Weekly report

This report shows WLM Workload MIPS usage and forecast statistics.

This information identifies the report:

Report ID

CPCPU25

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_FV

Attributes

MVS, CAPACITY, WORKLOAD, MIPS, WEEKLY, FORECAST




Figure 19. CP Workload MIPS Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Forecast

Forecast average MIPS for the CEC during the interval.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

CP Serv Class MIPS Forecast Weekly report

This report shows WLM Service or Report Class MIPS usage and forecast statistics.

This information identifies the report:

Report ID

CPCPU28

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_FV

Attributes

MVS, CAPACITY, WORKLOAD, MIPS, WEEKLY, FORECAST




Figure 20. CP Serv Class MIPS Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Forecast

Forecast average MIPS for the CEC during the interval.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

CP Application MIPS Forecast Weekly report

This report shows Business Application MIPS usage and forecast statistics.

This information identifies the report:

Report ID

CPCPU29

Report group

Capacity Planning - CPU

Source

CP_CPU_WKLD_FV

Attributes

MVS, CAPACITY, APPLICATION, MIPS, WEEKLY, FORECAST




Figure 21. CP Application MIPS Forecast Weekly reports

The report contains this information:

Date

Date of the measurement.

Forecast

Forecasted average MIPS for the CEC during the interval.

Average Used

Average MIPS used by the CEC during the interval.

Peak Used

Peak MIPS used by the CEC during the interval.

LSPR workload analysis

LSPR workload analysis reports.

CP LSPR Workload Analysis Hourly/Daily/Weekly report

These reports show the LSPR Workload of your system, as calculated from SMF type 113 records.

This information identifies the report:

Report ID

CPCPU38 / CPCPU39 / CPCPU40

Report group

Capacity Planning - CPU

Source

CP_LSPR_WKLD_H / CP_LSPR_WKLD_D / CP_LSPR_WKLD_W

Attributes

MVS, CAPACITY, WORKLOAD, LSPR, HOURLY/DAILY/WEEKLY

CPCPU38 - CP LSPR WORKLOAD ANALYSIS HOURLY




Figure 22. CP LSPR Workload Analysis Hourly report

The report contains this information:

LSPR Workload High Count

The total number of SMF intervals where the LSPR Workload is HIGH.

LSPR Workload Average Count

The total number of SMF intervals where the LSPR Workload is AVERAGE.

LSPR Workload Low Count

The total number of SMF intervals where the LSPR Workload is LOW.

Software cost analysis

CPU software cost analysis reports.

CP LPAR 4 Hour MSU Utilization Hourly report

This report shows LPAR 4 hour rolling average MSU usage.

This information identifies the report:

Report ID

CPMSU01

Report group

Software Cost Analysis

Source

CP_CPU_LPAR_H

Attributes

CP, LPAR, MSU, 4HRA, HOURLY

CPMSU01 - CP LPAR 4 HOUR MSU UTILIZATION BY HOUR




Figure 23. CP LPAR 4 Hour MSU Utilization Hourly report

The report contains this information:

Time of day

Time of the measurement.

LPAR Name

Name of the logical partition.

4 hour average

Four hour rolling average MSU.

CP LPAR MSU Statistics Hourly report

This report shows MSU usage statistics for a given LPAR per hour.

This information identifies the report:

Report ID

CPMSU02

Report group

Software cost analysis

Source

CP_CPU_LPAR_H

Attributes

CP, LPAR, MSU, HOURLY

CPMSU02 - CP LPAR MSU STATISTICS HOURLY

LPAR Name:
 Shift:
 Date:

LPAR MSU STATISTICS BY HOUR




Figure 24. CP LPAR MSU Statistics Hourly report

The report contains this information:

Time of day

Time of the measurement.

4 hour average

Four hour rolling average MSU.

Actual consumed

Actual MSU consumed by the partition.

CP LPAR Product Matrix Daily report

This report shows the distribution of products on partitions. It lists all partitions, and for each partition, it lists all the products which ran on that partition for a given day.

This information identifies the report:

Report ID

CPMSU03

Report group

Software cost analysis

Source

CP_PROD_MSU_H

Attributes

CP, LPAR, PRODUCT, DAILY

CPMSU03 - CP LPAR PRODUCT MATRIX DAILY

LPAR Name	Shift	Date
LPAR Name	Shift	Dec 11, 2013

LPAR PRODUCT MATRIX BY DAY

LPAR Name	Product ID
M50	5605-DB2 DB1R
	5645-001 TSO/E
	5655-W65 WAS Z/OS
	5655-Y26 COGNOS
	5698-AA6 COGNOS
	5698-AA6 USINGWAS

Figure 25. CP LPAR Product Matrix Daily report

The report contains this information:

Product

Product ID and name.

LPAR

Name of the logical partition.

CP Product LPAR Matrix Daily report

This report shows the distribution of products on partitions. It lists all partitions, and for each partition, it lists all the products which ran on that partition for a given day.

This information identifies the report:

Report ID

CPMSU04

Report group

Software cost analysis

Source

CP_PROD_MSU_H

Attributes

CP, LPAR, PRODUCT, DAILY

CPMSU04 - CP PRODUCT LPAR MATRIX DAILY

Shift	Product ID	Product Feature Name	Date
Shift	Product ID	Product Feature Name	* Dec 11, 2013

PRODUCT LPAR MATRIX BY DAY

Product	LPAR
5605-DB2 DB1R	M50
5645-001 TSO/E	M50
5655-W65 WAS Z/OS	M50
5655-Y26 COGNOS	M50
5698-AA6 COGNOS	M50
5698-AA6 USINGWAS	M50

Figure 26. CP Product LPAR Matrix Daily report

The report contains this information:

Product

Product ID and name.

LPAR

Name of the logical partition.

CP LPAR Product MSU Hourly report

This report shows the MSU usage of products which ran on a specific LPAR during a given day. It lists the Products and MSU consumption per hour.

This information identifies the report:

Report ID

CPMSU05

Report group

Software cost analysis

Source

CP_PROD_MSU_H

Attributes

CP, MSU, PRODUCT, HOURLY

CPMSU05 - CP LPAR PRODUCT MSU HOURLY

LPAR Name:
 Shift:
 Date:

LPAR PRODUCT MSU BY HOUR




Figure 27. CP LPAR Product MSU Hourly report

The report contains this information:

Time of day

Time of the measurement.

Product

Product ID and name.

MSU

MSU consumed by the product during the interval.

Storage subcomponent reports

Storage reports provide information on CPU storage and help you monitor the overall usage, and the highs and lows of storage available for use.

Available reports show the minimum, average and maximum central storage that is available over time as well as a definable threshold of acceptable levels. Reports are also provided which show the actual usage versus the predicted usage, over time.

The Storage reports are grouped by report type:

- Storage usage
- Storage forecasting

Storage usage

Storage usage reports.

CP Central Storage Available Hourly/Daily/Weekly report

This report shows memory statistics on system level, per hour, day or week.

This information identifies the report:

Report ID

CPSTO01 / CPSTO02 / CPSTO03

Report group

Capacity Planning - Storage

Source

CP_STORAGE_H / CP_STORAGE_D / CP_STORAGE_W

Attributes

MVS, STORAGE, AVAILABLE, FRAMES, HOURLY/DAILY/WEEKLY

CPSTO01 - CP CENTRAL STORAGE AVAILABLE HOURLY

MVS System ID Shift Date

* M50 Shift * Jun 17, 2018  [Submit](#)

CENTRAL STORAGE BY HOUR




Figure 28. CP Central Storage Available Hourly report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Available AVG

Average memory available during the interval, in megabytes.

Available MIN

Minimum memory available during the interval, in megabytes.

OK Threshold

This threshold indicates the number of frames on the available frame queue when stealing ends.

Storage forecasting

Storage forecasting reports.

CP Central Storage Forecast Weekly report

This report shows memory statistics, included forecast values, on system level, by week.

This information identifies the report:

Report ID

CPSTO05

Report group

Capacity Planning - Storage

Source

CP_STORAGE_F

Attributes

MVS, STORAGE, AVAILABLE, FRAMES, WEEKLY

The report contains this information:

Date

Date of the measurement.

Forecast available MIN

Forecast minimum memory available during the interval, in megabytes.

Available AVG

Average memory available during the interval, in megabytes.

Available MIN

Minimum memory available during the interval, in megabytes.

OK Threshold

This threshold indicates the number of frames on the available frame queue when stealing ends.

Channel subcomponent reports

Channel reports provide information on the speed of the I/O operations in a system. Reports are also provided that show the channel busy percentage versus the predicted channel busy percentage, over time.

The Channel reports are grouped by report type:

- Channel usage
- Channel forecasting

Channel usage

Channel usage reports.

CP Channel Peak Busy Hourly/Daily/Weekly report

This report shows channel utilization statistics on system level, by hour, day or week.

This information identifies the report:

Report ID

CPCHAN01 / CPCHAN02 / CPCHAN03

Report group

Capacity Planning - Channel

Source

CP_CHANNEL_H / CP_CHANNEL_D / CP_CHANNEL_W

Attributes

MVS, CHANNEL, PEAK, BUSY, HOURLY / DAILY / WEEKLY

CPCHAN01 - CP CHANNEL PEAK BUSY HOURLY

MVS System ID	Shift	Channel Type	Peak %	Date
* M50	* OTHER	* FC_S	* 10	* Jun 14, 2018

CHANNEL PEAK BUSY BY HOUR




Figure 29. CP Channel Peak Busy Hourly report

The report contains this information:

Time of day / Date

Time / date of the measurement.

Channel

Channel name.

Peak busy

Highest percentage busy during the interval.

Channel forecasting

Channel forecasting reports.

CP Channel Peak Busy Forecast Weekly report

This report shows channel utilization and forecasting statistics on system level.

This information identifies the report:

Report ID

CPCHAN05

Report group

Capacity Planning - Channel

Source

CP_CHANNEL_F

Attributes

MVS, CHANNEL, PEAK, BUSY, FORECAST, WEEKLY




Figure 30. CP Channel Peak Busy Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Channel

Channel name.

Forecast peak busy

Forecast highest percentage busy during the interval.

Peak busy

Highest percentage busy during the interval.

Disk subcomponent reports

Disk reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

The Disk reports are grouped by report type:

- Disk usage
- Disk forecasting

Disk usage

Disk reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

CP Disk Storage Group Allocated Space Daily/Monthly report

This report shows disk statistics on storage group level, by day or by month.

This information identifies the report:

Report ID

CPDISK01 / CPDISK02

Report group

Capacity Planning - Disk

Source

DFSMS_VOLUME_D / DFSMS_VOLUME_M

Attributes

MVS, DISK, STORAGE, DAILY/MONTHLY

CPDISK02 - CP DISK STORAGE GROUP ALLOCATED SPACE MONTHLY




Figure 31. CP Disk Storage Group Allocated Space Monthly report

The report contains this information:

Date

Date of the measurement.

Storage group

Storage group name.

Allocated space

Space allocated in the storage group, in megabytes.

Disk forecasting

Disk reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

CP Disk Storage Group Space Forecast Monthly report

This report shows disk statistics on storage group level.

This information identifies the report:

Report ID

CPDISK03

Report group

Capacity Planning - Disk

Source

CP_DISK_F

Attributes

MVS, DISK, STORAGE, MONTHLY, FORECAST




Figure 32. CP Disk Storage Group Space Forecast Monthly report

The report contains this information:

Date

Date of the measurement.

Storage group

Storage group name.

Allocated space

Space allocated in the storage group, in megabytes.

Tape subcomponent reports

Tape reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

The Tape reports are grouped by report type:

- Tape usage
- Tape forecasting

Tape usage

Tape reports provide information on the actual storage usage, as well as predicted usage, per storage group over time.

CP Tape Storage Group Space Daily/Weekly report

This report shows tape statistics on storage group level, by day or by week.

This information identifies the report:

Report ID

CPTAPE01 / CPTAPE02

Report group

Capacity Planning - Tape

Source

CP_TAPE_D / CP_TAPE_W

Attributes

MVS, TAPE, STORAGE, DAILY/WEEKLY

CPTAPE02 - CP TAPE STORAGE GROUP SPACE WEEKLY




Figure 33. CP Tape Storage Group Space Weekly report

The report contains this information:

Date

Date of the measurement.

Storage group

Storage group name.

Allocated space

Space allocated in the storage group, in megabytes.

Tape forecasting

Tape forecasting reports.

CP Tape Storage Group Space Forecast Weekly report

This report shows tape usage and forecast statistics on storage group level.

This information identifies the report:

Report ID

CPTAPE03

Report group

Capacity Planning - Tape

Source

CP_TAPE_F

Attributes

MVS, TAPE, STORAGE, WEEKLY, FORECAST




Figure 34. CP Tape Storage Group Space Forecast Weekly report

The report contains this information:

Date

Date of the measurement.

Storage group

Storage group name.

Forecast space

Forecast allocated space in the storage group, in megabytes.

Allocated space

Space allocated in the storage group, in megabytes.

Appendix A. Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides a number of ways for you to obtain the support you need.

- Searching knowledge bases: You can search across a large collection of known problems and workarounds, Technotes, and other information.
- Obtaining fixes: You can locate the latest fixes that are already available for your product.
- Contacting IBM Software Support: If you still cannot solve your problem, and you need to work with someone from IBM, you can use a variety of ways to contact IBM Support. See [“Contacting IBM Support”](#) on page 97 for more information.

Contacting IBM Support

This topic describes how to contact IBM Support if you have been unable to resolve a problem with IBM Z Performance and Capacity Analytics.

Before contacting IBM Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have. For more information, refer to the IBM Support website at the following links:

IBM Support

<https://www.ibm.com/mysupport/s/>

IBM Z Support

<https://www.ibm.com/support/pages/ibm-enterprise-support-and-preferred-care-options-ibm-z>

To contact IBM Support to report a problem (*open a case*), follow these steps:

1. Determine the business impact.
2. Describe the problem and gather information.
3. Submit the problem report.

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem that you are reporting. Use the following criteria:

Severity 1

The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2

The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3

The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4

The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing the problem and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Support specialists can help you solve the problem efficiently. To save time, know the answers to the following questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, product-specific customization, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting the problem

You can submit your problem to IBM Support in either of the following ways:

Online

Go to <https://www.ibm.com/mysupport/s/>, click on **Open a case**, and enter the relevant details into the online form.

By email or phone

For the contact details in your country, go to the IBM Support website at <https://www.ibm.com/support/>. Look for the tab on the right and click **Contact and feedback > Directory of worldwide contacts** for a list of countries by geographic region. Select your country to find the contact details for general inquiries, technical support, and customer support.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support website, so that other users who experience the same problem can benefit from the same resolution.

Notices

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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 might 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 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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 Corporation
224A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

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 measurement 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.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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

COPYRIGHT LICENSE:

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

If you are viewing this information in softcopy form, the photographs and color illustrations might not display.

Trademarks

Listed below are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

IBM
IBM logo
IBM Z
ibm.com
AIX
AS/400
CICS
Db2
DB2
developerWorks
Domino
eServer
GDDM
IBMLink
IMS
iSeries
Lotus
Lotus Notes
MQSeries

MVS
NetView
OS/390
Passport Advantage
PR/SM
pSeries
RACF
RAMAC
Rational
RMF
S/390
Tivoli
VSE/ESA
VTAM
WebSphere
z/OS
z/VM
zSeries

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.



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

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

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

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

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

Index

A

APAR (Authorized Program Analysis Report) [97](#)
applications, business
 CP_BUSINESS_APPL [10](#)
 lookup tables [8](#)
architecture [1](#)
Authorized Program Analysis Report (APAR) [97](#)

B

business applications
 CP_BUSINESS_APPL [10](#)
 lookup tables [8](#)

C

capacity
 CPU [28](#)
 installing and configuring [7](#)
capacity planning
 channel [27](#)
 CPU [27](#)
 disk [27](#)
 input data [7](#)
 SMF records required [7](#)
 storage [27](#)
 tape [27](#)
Capacity Planning for z/OS component
 installing and configuring [7](#)
 introduction [1](#)
central storage
 CP Central Storage Available Hourly/Daily/Weekly report [88](#)
 CP Central Storage Forecast Weekly report [89](#)
 reporting [5](#)
 reports [87](#)
 See also storage
central storage usage
 reports [87](#)
changes in this edition [x](#)
channel
 CP Channel Peak Busy Forecast Weekly report [90](#)
 CP Channel Peak Busy Hourly/Daily/Weekly report [89](#)
Channel
 tables [58](#)
channel forecasting
 reports [90](#)
channel reporting [5](#)
channel subcomponent
 reports [89](#)
Cognos reporting
 implementing [15](#)
Collector
 architecture [1](#)
Common Data Provider for z
 install and configure [18](#)

configure
 ELK [23](#)
configuring
 control tables [11](#)
 Forecaster module [13](#)
control tables
 configuring [11](#)
 described [5](#)
 updating [11](#)
CP Application MIPS Forecast Weekly report [80](#)
CP Application MIPS Hourly/Daily/Weekly report [75](#)
CP CEC MIPS Forecast Weekly report [76](#)
CP CEC MIPS Usage Hourly/Daily/Weekly report [68](#)
CP Central Storage Available Hourly/Daily/Weekly report [88](#)
CP Central Storage Forecast Weekly report [89](#)
CP Channel Peak Busy Forecast Weekly report [90](#)
CP Channel Peak Busy Hourly/Daily/Weekly report [89](#)
CP Disk Storage Group Allocated Space Daily report [92](#)
CP Disk Storage Group Space Forecast Monthly report [93](#)
CP LPAR 4 Hour MSU Utilization Hourly report [82](#)
CP LPAR MIPS by CEC Average Hourly/Daily/Weekly report [70](#)
CP LPAR MIPS by CEC Peak Hourly/Daily/Weekly report [71](#)
CP LPAR MIPS Forecast Weekly report [77](#)
CP LPAR MIPS Hourly/Daily/Weekly report [69](#)
CP LPAR MSU Statistics Hourly report [83](#)
CP LPAR Product Matrix Daily report [84](#)
CP LPAR Product MSU Hourly report [86](#)
CP LSPR Workload Analysis Hourly/Daily/Weekly report [81](#)
CP Product LPAR Matrix Daily report [85](#)
CP Serv Class MIPS by Workload Avg Hourly/Daily/Weekly report [73](#)
CP Serv Class MIPS by Workload Peak Hourly/Daily/Weekly report [74](#)
CP Serv Class MIPS Forecast Weekly report [79](#)
CP Serv Class MIPS Hourly/Daily/Weekly report [72](#)
CP Tape Storage Group Space Daily/Weekly report [94](#)
CP Tape Storage Group Space Forecast Weekly report [95](#)
CP Workload MIPS Forecast Weekly report [78](#)
CP Workload MIPS Hourly/Daily/Weekly report [71](#)
CP_BUSINESS_APPL lookup table [10](#)
CP_CHANNEL_D [59](#)
CP_CHANNEL_H [58](#)
CP_CHANNEL_W [60](#)
CP_CPU_LPAR_D [40](#)
CP_CPU_LPAR_F [43](#)
CP_CPU_LPAR_FH [44](#)
CP_CPU_LPAR_H [39](#)
CP_CPU_LPAR_pH [45](#)
CP_CPU_LPAR_T [37](#)
CP_CPU_LPAR_W [42](#)
CP_CPU_WKLD_D [50](#)
CP_CPU_WKLD_F [53](#)
CP_CPU_WKLD_H [49](#)
CP_CPU_WKLD_W [52](#)
CP_DISK_F [64](#)

- [CP_LSPR_WKLD_D 48](#)
- [CP_LSPR_WKLD_H 47](#)
- [CP_LSPR_WKLD_W 48](#)
- [CP_SHIFT lookup table 9](#)
- [CP_STORAGE_D 55](#)
- [CP_STORAGE_F 57](#)
- [CP_STORAGE_H 54](#)
- [CP_STORAGE_W 56](#)
- [CP_TAPE_D 65](#)
- [CP_TAPE_F 66](#)
- [CP_TAPE_W 65](#)
- [CP_THRESHOLDS 9](#)
- [CPCHAN01 CP Channel Peak Busy Hourly report 89](#)
- [CPCHAN02 CP Channel Peak Busy Daily report 89](#)
- [CPCHAN03 CP Channel Peak Busy Weekly report 89](#)
- [CPCHAN05 CP Channel Peak Busy Forecast Weekly report 90](#)
- [CPCPU01 CP CEC MIPS Hourly report 68](#)
- [CPCPU02 CP LPAR MIPS Hourly report 69](#)
- [CPCPU03 CP LPAR MIPS by CEC Average Hourly report 70](#)
- [CPCPU04 CP LPAR MIPS by CEC Peak Hourly report 71](#)
- [CPCPU05 CP CEC MIPS Daily report 68](#)
- [CPCPU06 CP CEC MIPS Weekly report 68](#)
- [CPCPU07 CP LPAR MIPS Daily report 69](#)
- [CPCPU08 CP LPAR MIPS Weekly report 69](#)
- [CPCPU10 CP CEC MIPS Forecast Weekly report 76](#)
- [CPCPU12 CP LPAR MIPS Forecast Weekly report 77](#)
- [CPCPU13 CP Workload MIPS Hourly report 71](#)
- [CPCPU14 CP Serv Class MIPS Hourly report 72](#)
- [CPCPU15 CP Application MIPS Hourly report 75](#)
- [CPCPU16 CP Serv Class MIPS by Workload Avg Hourly report 73](#)
- [CPCPU17 CP Serv Class MIPS by Workload Peak Hourly report 74](#)
- [CPCPU18 CP Workload MIPS Daily report 71](#)
- [CPCPU19 CP Workload MIPS Weekly report 71](#)
- [CPCPU20 CP Serv Class MIPS Daily report 72](#)
- [CPCPU21 CP Application MIPS Daily report 75](#)
- [CPCPU22 CP Serv Class MIPS Weekly report 72](#)
- [CPCPU23 CP Application MIPS Weekly report 75](#)
- [CPCPU25 CP Workload MIPS Forecast Weekly report 78](#)
- [CPCPU28 CP Serv Class MIPS Forecast Weekly report 79](#)
- [CPCPU29 CP Application MIPS Forecast Weekly report 80](#)
- [CPCPU30 CP LPAR MIPS by CEC Average Daily report 70](#)
- [CPCPU31 CP LPAR MIPS by CEC Peak Daily report 71](#)
- [CPCPU32 CP LPAR MIPS by CEC Average Weekly report 70](#)
- [CPCPU33 CP LPAR MIPS by CEC Peak Weekly report 71](#)
- [CPCPU34 CP Serv Class MIPS by Workload Avg Daily report 73](#)
- [CPCPU35 CP Serv Class MIPS by Workload Peak Daily report 74](#)
- [CPCPU36 CP Serv Class MIPS by Workload Avg Weekly report 73](#)
- [CPCPU37 CP Serv Class MIPS by Workload Peak Weekly report 74](#)
- [CPCPU38 CP LSPR Workload Analysis Hourly report 81](#)
- [CPCPU39 CP LSPR Workload Analysis Daily report 81](#)
- [CPCPU40 CP LSPR Workload Analysis Weekly report 81](#)
- [CPDISK01 CP Disk Storage Group Allocated Space Daily report 92](#)
- [CPDISK02 CP Disk Storage Group Allocated Space Monthly report 92](#)
- [CPDISK03 CP Disk Storage Group Space Forecast Monthly report 93](#)
- [CPMSU01 CP LPAR 4 Hour MSU Utilization Hourly report 82](#)
- [CPMSU02 CP LPAR MSU Statistics Hourly report 83](#)
- [CPMSU03 CP LPAR Product Matrix Daily report 84](#)
- [CPMSU04 CP Product LPAR Matrix Daily report 85](#)
- [CPMSU05 CP LPAR Product MSU Hourly report 86](#)
- [CPSTO01 CP Central Storage Available Hourly report 88](#)
- [CPSTO02 CP Central Storage Available Daily report 88](#)
- [CPSTO03 CP Central Storage Available Weekly report 88](#)
- [CPSTO05 CP Central Storage Forecast Weekly report 89](#)
- [CPTAPE01 CP Tape Storage Group Space Daily report 94](#)
- [CPTAPE02 CP Tape Storage Group Space Weekly report 94](#)
- [CPTAPE03 CP Tape Storage Group Space Forecast Weekly report 95](#)
- CPU
 - [capacity 28](#)
 - [forecasting 31](#)
 - [LSPR workload analysis 81](#)
 - [tables 37](#)
- CPU forecasting
 - [reports 76](#)
- CPU reporting [5](#)
- CPU software cost analysis
 - [reports 82](#)
- CPU subcomponent
 - [reports 67](#)
- CPU usage
 - [CP Application MIPS Forecast Weekly report 80](#)
 - [CP Application MIPS Hourly/Daily/Weekly report 75](#)
 - [CP CEC MIPS Forecast Weekly report 76](#)
 - [CP CEC MIPS Hourly/Daily/Weekly report 68](#)
 - [CP LPAR MIPS by CEC Average Hourly/Daily/Weekly report 70](#)
 - [CP LPAR MIPS by CEC Peak Hourly/Daily/Weekly report 71](#)
 - [CP LPAR MIPS Forecast Weekly report 77](#)
 - [CP LPAR MIPS Hourly/Daily/Weekly report 69](#)
 - [CP Serv Class MIPS by Workload Avg Hourly/Daily/Weekly report 73](#)
 - [CP Serv Class MIPS by Workload Peak Hourly/Daily/Weekly report 74](#)
 - [CP Serv Class MIPS Forecast Weekly report 79](#)
 - [CP Serv Class MIPS Hourly/Daily/Weekly report 72](#)
 - [CP Workload MIPS Forecast Weekly report 78](#)
 - [CP Workload MIPS Hourly/Daily/Weekly report 71](#)
- customer support
 - [contacting IBM Support 97](#)

D

- dashboard
 - [configuring Kibana 24](#)
 - [configuring Logstash 24](#)
 - [download and install for ELK reporting 23](#)
 - [ELK reporting 23, 26](#)
 - [Kibana 26](#)
- data tables
 - [architecture 1](#)
 - [by subcomponent 37](#)
- Db2 tables
 - [architecture 1](#)
 - [data tables 1](#)
 - [forecast data tables 1](#)
 - [lookup tables, updating 8](#)

DFSMS_VOLUME_D [61](#)

DFSMS_VOLUME_M [62](#)

disk

CP Disk Storage Group Allocated Space Daily/Monthly report [92](#)

CP Disk Storage Group Space Forecast Monthly report [93](#)

Disk

tables [61](#)

disk forecasting

reports [92](#)

disk subcomponent

reports [91](#)

disk usage

reports [91](#)

E

ELK reporting

ad hoc reporting [26](#)

Common Data Provider for z, install and configure [18](#)

configuring IBM Z Performance and Capacity Analytics [18](#)

configuring Kibana index pattern [24](#)

configuring Logstash protocol [24](#)

configuring the Forecaster [19](#)

dashboard [26](#)

install and configure [17](#), [23](#)

installing Logstash and Kibana components [23](#)

Logstash protocol [22](#)

onfiguring IBM Common Data Provider for z for

streaming [22](#)

setting up data streaming [17](#)

F

forecast data tables [1](#)

forecast reports

channel [90](#)

CPU [76](#)

disk [92](#)

introduction [31](#)

storage [89](#)

tape [95](#)

Forecaster

architecture [1](#)

forecaster module

subdirectories [12](#)

unpacking [12](#)

Forecaster module

configuring [13](#)

install and configure [12](#)

forecasting

CP Application MIPS Forecast Weekly report [80](#)

CP CEC MIPS Forecast Weekly report [76](#)

CP Central Storage Forecast Weekly report [89](#)

CP Channel Peak Busy Forecast Weekly report [90](#)

CP Disk Storage Group Space Forecast Monthly report [93](#)

CP LPAR MIPS Forecast Weekly report [77](#)

CP Serv Class MIPS Forecast Weekly report [79](#)

CP Tape Storage Group Space Forecast Weekly report [95](#)

forecasting (*continued*)

CP Workload MIPS Forecast Weekly report [78](#)

CPU [31](#)

data [31](#)

future [31](#)

input data [7](#)

MIPS [31](#)

resource usage [31](#)

scope [31](#)

SMF records required [7](#)

supplied forecasts [33](#)

tailoring a supplied forecast [35](#)

I

IBM Common Data Provider for z [17](#)

IBM Support [97](#)

implementing

Cognos reporting [15](#)

IBM Z Performance and Capacity Analytics [7](#)

initialize

control tables [5](#)

lookup tables [5](#)

input data

SMF records required [7](#)

install and configure

Capacity Planning for z/OS component [7](#)

ELK [23](#)

for ELK reporting [17](#)

Forecaster module [12](#)

installing [7](#)

introduction

Capacity Planning for z/OS component [1](#), [7](#)

J

Java [1](#), [31](#)

JCL

lookup tables [8](#)

setting up data streaming through IBM Common Data Provider for z [17](#)

K

Kibana

dashboard [26](#)

KPM reports [28](#)

L

log record input [1](#)

Logstash [22–24](#)

lookup tables

configuring [8](#)

CP_BUSINESS_APPL [10](#)

CP_SHIFT [9](#)

CP_THRESHOLDS [9](#)

described [5](#)

LSPR MIPS ratings [11](#)

LSPR workload analysis

reports [81](#)

LSPR_MIPS table [11](#)

M

MIPS

- calculating in CPU subcomponent reports [67](#)
- forecasting [31](#)
- LSPR_MIPS control table [11](#)

MSU reporting [5](#)

N

naming standard tables [37](#)

O

off-platform reporting ELK [17](#)

P

pacing parameter [31](#) parameters

- lookup tables [8](#)

problem determination, IBM Support determining business impact [97](#)

publication log stream architecture, off-platform reporting [1](#) ELK reporting [18](#)

R

remote receiver ELK [17](#), [22](#)

report

- CP Application MIPS Forecast Weekly report [80](#)
- CP Application MIPS Hourly/Daily/Weekly report [75](#)
- CP CEC MIPS Forecast Weekly report [76](#)
- CP CEC MIPS Usage Hourly/Daily/Weekly report [68](#)
- CP Central Storage Available Hourly/Daily/Weekly [88](#)
- CP Central Storage Forecast Weekly report [89](#)
- CP Channel Peak Busy Forecast Weekly report [90](#)
- CP Channel Peak Busy Hourly/Daily/Weekly report [89](#)
- CP Disk Storage Group Allocated Space Daily report [92](#)
- CP Disk Storage Group Space Forecast Monthly [93](#)
- CP LPAR 4 Hour MSU Utilization Hourly report [82](#)
- CP LPAR MIPS by CEC Average Hourly/Daily/Weekly [70](#)
- CP LPAR MIPS by CEC Peak Hourly/Daily/Weekly report [71](#)
- CP LPAR MIPS Forecast Weekly report [77](#)
- CP LPAR MIPS Usage Hourly/Daily/Weekly report [69](#)
- CP LPAR MSU Statistics Hourly report [83](#)
- CP LPAR Product Matrix Daily report [84](#)
- CP LPAR Product MSU Hourly report [86](#)
- CP LSPR Workload Analysis Hourly/Daily/Weekly report [81](#)
- CP Product LPAR Matrix Daily report [85](#)
- CP Serv Class MIPS by Workload Avg Hourly/Daily/Weekly report [73](#)
- CP Serv Class MIPS by Workload Peak Hourly/Daily/Weekly report [74](#)
- CP Serv Class MIPS Forecast Weekly report [79](#)
- CP Serv Class MIPS Hourly/Daily/Weekly report [72](#)
- CP Tape Storage Group Space Daily/Weekly report [94](#)

report (*continued*)

- CP Tape Storage Group Space Forecast Weekly report [95](#)
- CP Workload MIPS Forecast Weekly report [78](#)
- CP Workload MIPS Hourly/Daily/Weekly report [71](#)

reporting

- ad hoc, ELK [26](#)

reports

- central storage usage [87](#)
- channel forecasting [90](#)
- channel subcomponent [89](#)
- CPU forecasting [76](#)
- CPU software cost analysis [82](#)
- CPU subcomponent [67](#), [87](#)
- CPU usage [68](#)
- disk forecasting [92](#)
- disk subcomponent [91](#)
- disk usage [91](#)
- LSPR workload analysis [81](#)
- storage forecasting [89](#)
- tape forecasting [95](#)
- tape subcomponent [93](#)
- tape usage [94](#)

resource usage

- forecasting [31](#)

S

severity

- contacting IBM Support [97](#)
- determining business impact [97](#)

shadowing, for ELK reporting [19](#)

shifts

- CP_SHIFT [9](#)
- lookup tables [8](#)

SMF records [7](#)

software cost analysis

- CP LPAR 4 Hour MSU Utilization Hourly report [82](#)
- CP LPAR MSU Statistics Hourly report [83](#)
- CP LPAR Product Matrix Daily report [84](#)
- CP LPAR Product MSU Hourly report [86](#)
- CP Product LPAR Matrix Daily report [85](#)

software cost analysis reports

- CPU [82](#)

storage

- CP Central Storage Available Hourly/Daily/Weekly report [88](#)
- CP Central Storage Forecast Weekly report [89](#)

Storage

- tables [54](#)

storage (memory) reporting [5](#)

storage forecasting

- reports [89](#)

storage subcomponent

- reports [87](#)

storage usage

- reports [87](#)

subcomponents

- data tables [4](#)

subdirectories

- forecaster module [12](#)

Support

- contacting IBM [97](#)
- describing problems [97](#)

Support (*continued*)
 determining business impact [97](#)
 submitting problems [97](#)
support information [97](#)

T

table mapping files
 updating [13](#)
tables
 by subcomponent [37](#)
 Channel [58](#)
 control [5](#)
 control tables, updating [11](#)
 CP_CHANNEL_D [59](#)
 CP_CHANNEL_H [58](#)
 CP_CHANNEL_W [60](#)
 CP_CPU_LPAR_D [40](#)
 CP_CPU_LPAR_F [43](#)
 CP_CPU_LPAR_FH [44](#)
 CP_CPU_LPAR_H [39](#)
 CP_CPU_LPAR_PH [45](#)
 CP_CPU_LPAR_T [37](#)
 CP_CPU_LPAR_W [42](#)
 CP_CPU_WKLD_D [50](#)
 CP_CPU_WKLD_F [53](#)
 CP_CPU_WKLD_H [49](#)
 CP_CPU_WKLD_W [52](#)
 CP_DISK_F [64](#)
 CP_LSPR_WKLD_D [48](#)
 CP_LSPR_WKLD_H [47](#)
 CP_LSPR_WKLD_W [48](#)
 CP_STORAGE_D [55](#)
 CP_STORAGE_F [57](#)
 CP_STORAGE_H [54](#)
 CP_STORAGE_W [56](#)
 CP_TAPE_D [65](#)
 CP_TAPE_F [66](#)
 CP_TAPE_W [65](#)
 CPU [37](#)
 Db2 [8](#)
 descriptions [37](#)
 DFSMS_VOLUME_D [61](#)
 DFSMS_VOLUME_M [62](#)
 Disk [61](#)
 lookup [5](#)
 lookup tables, updating [8](#)
 LSPR_MIPS [11](#)
 naming standard [37](#)
 Storage [54](#)
 Tape [65](#)
tailoring
 forecast [35](#)
tape
 CP Tape Storage Group Space Daily/Weekly report [94](#)
 CP Tape Storage Group Space Forecast Weekly report [95](#)
Tape
 tables [65](#)
tape forecasting
 reports [95](#)
tape subcomponent
 reports [93](#)
tape usage

tape usage (*continued*)
 reports [94](#)
thresholds
 CP_THRESHOLDS [9](#)
 lookup tables [8](#)

U

unpacking
 ELK reporting dashboard components [23](#)
 forecaster module [12](#)
updating
 control tables [11](#)
 table mapping files [13](#)
usage
 CCP Channel Peak Busy Hourly/Daily/Weekly report [89](#)
 CP Central Storage Available Hourly/Daily/Weekly report [88](#)
 CP Disk Storage Group Allocated Space Daily/Monthly report [92](#)
 CP Tape Storage Group Space Daily/Weekly report [94](#)
usage reports
 central storage [87](#)
 channel [89](#)
 CPU [68](#)
 disk [91](#)
 tape [94](#)

W

what's new in this edition [x](#)
workload analysis reports
 LSPR [81](#)

Z

zIIP [1](#), [31](#)



SC28-3213-00

