IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning Version 7.2 Fix Pack 2

User's Guide





IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning Version 7.2 Fix Pack 2

User's Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 393.

This edition applies to Version 7.2 Fix Pack 2 of IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning (product number 5724-L92).

© Copyright IBM Corporation 2011, 2013. US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures
Tables
Chapter 1. Tivoli Monitoring for Virtual Environments structure
Chapter 2. PowerVM: Installing components

· · · · · · · · · · · · · · · · · ·
components
PowerVM: Using the IBM Installation Manager GUI
to install components
PowerVM: Using the IBM Installation Manager in
silent mode to install the Capacity Planner 11
PowerVM: Using the IBM Installation Manager
Console to install the Capacity Planner
PowerVM: Installing the Capacity Planner 16
PowerVM: Using the IBM Installation Manager GUI
to uninstall the Capacity Planner
PowerVM: Using IBM Installation Manager Console
to uninstall the Capacity Planner
PowerVM: Using the IBM Installation Manager in
silent mode to uninstall the Capacity Planner 26
PowerVM: Types of report installation
PowerVM: Installing Capacity Planner Reports 27
Installing reports by using the command line
Installing reports by using the silent mode 32
PowerVM: Uninstalling Capacity Planning Reports 33

Chapter 3. VMware: Installing

components
VMware: Installing prerequisites
VMware: Types of installation
VMware: Using the Launchpad to install
components
VMware: Using the IBM Installation Manager GUI
to install components
VMware: Using the IBM Installation Manager silent
mode to install the Capacity Planner and the
Dashboard
VMware: Using the IBM Installation Manager
Console to install the Capacity Planner and the
Dashboard
VMware: Installing the Capacity Planner and the
Dashboard
VMware: Using the IBM Installation Manager GUI
to uninstall the Capacity Planner and the Dashboard 53

VMware: Using the IBM Installation Manager in		
silent mode to uninstall the Capacity Planner and		
the Dashboard		54
VMware: Using IBM Installation Manager Console		
to uninstall the Capacity Planner and the Dashboard	d	55
VMware: Types of report installation		57
VMware: Installing Capacity Planner Reports		58
Installing reports by using the command line .		63
Installing reports by using the silent mode		64
VMWare: Uninstalling Capacity Planner Reports .		64

Chapter 4. Updating the Capacity Planners and Dashboard for VMware

from an earlier version 6	37
Updating components by using the IBM Installation	
Manager GUI	67
Updating components by using the IBM Installation	
Manager console mode	77
Updating components by using the IBM Installation	
Manager silent mode	83

Chapter 5. Rolling back Capacity Planners and Dashboard for VMware to

an earlier version	•	•	ОЭ
Rolling back components by using the IBM			
Installation Manager GUI.			. 85
Rolling back components by using the IBM			
Installation Manager console mode			. 94

Chapter 6. Linux KVM and VMware Performance and Capacity

Management Reports: Installation 101
Linux KVM and VMware: Prerequisites for
installing the Performance and Capacity
Management Reports
Linux KVM and VMware: Cataloging databases
and servers
Linux KVM and VMware: Types of report
installation
Linux KVM: Installing Performance and Capacity
Management Reports
Installing reports by using the command line 110
Installing reports by using the silent mode 111
VMware: Installing Performance and Capacity
Management Reports
Installing reports by using the command line 119
Installing reports by using the silent mode 119
Linux KVM and VMware: Uninstalling
Performance and Capacity Management Reports . 119

Chapter 7. PowerVM and VMware: Configuring the Capacity Planners

and the Dashboard for VMware	121
Configuring a connection to the IBM Tivoli	
Monitoring dashboard data provider	. 121
Configuring a connection for the Capacity Planner	
for VMware and the Capacity Planner for	
PowerVM	. 121
Verifying the Capacity Planner database	
connection details	. 121
Configuring historical data collection for the	
Capacity Planner for PowerVM on Linux and	
UNIX	. 122
Configuring historical data collection for the	
Capacity Planner for PowerVM in Windows	. 123
Configuring homogeneous federation for the	4.00
Capacity Planner for PowerVM	. 123
Configuring homogeneous federation for the	405
Capacity Planner for VMware	. 127
Verifying the Capacity Planner federation	. 131
Verifying the Capacity Planner federation by	
using the summary of created federation views	. 131
Verifying the Capacity Planner federation by	
using the log files	. 133
Configuring connection for the Dashboard for	400
VMware	. 133
Configuring Dashboard for VMware to	
communicate with Tivoli Application Dependency	104
Discovery Manager	. 134
Configuring historical data collection for the	
Performance and Capacity Management Reports	405
for VMware.	. 135
Heterogeneous federation configuration for Oracle	10(
and MSSQL Server	. 136
Configuring heterogeneous federation	10/
prerequisites for Oracle and MSSQL Server	136
Registering the service or node for the MSSQL	10(
server on Windows	136
Registering the service or node for the MSSQL	107
Server on UNIX	. 137
Registering the rederation server definition for	107
MSSQL	. 137
Registering the service or node for the Oracle	120
Begistering the service of node for the Oregle	. 139
Compared by LINIX	140
Server on UNIX	. 140
Registering the server definition for Oracle .	140
Ludating and execution recovered for Dobb and	. 142
Departing configuration passwords for Dashboard,	144
Undating password for Capacity Planner	. 144
database	144
Undefine recovered for the dechlosered date	. 144
provider connection	145
Undating password for connection from	143
Dashboard for VMware to Tiveli Application	
Dependency Discovery Manager	145
Dependency Discovery Manager	140
	143

Chapter 8. Dashboard for VMware	147
Dashboard for VMware Clusters	. 147
Accessing the Dashboard for VMware Clusters	147
Clusters	. 147
Servers	. 153
Virtual Machines	. 159
Datastores	. 164
Dashboard for VMware ESX Servers	. 170
Accessing the Dashboard for VMware ESX	
Servers	. 170
Servers	. 170
Launching to Tivoli Common Reporting	. 173
Launching to Tivoli Enterprise Portal	. 173
Setting a trace	. 174
Sharing the current dashboard page link	. 175
Dashboard Health Checks for VMware	. 175
Accessing the Dashboard Health Checks for	
VMware	. 175
Health Checks for VMware	. 175

Chapter 9. VMware VI agent: Using Performance and Capacity

Management Reports	179
Report packages	. 179
Performance and Capacity Management report	
descriptions	. 179
Prerequisites for running the Performance and	
Capacity Management reports	. 182
Accounting	. 185
Performance trends and resource forecasts	. 185
Prerequisites checking	. 203
What-if analysis for capacity estimation:	
estimate the capacity of your environment .	. 205
Workload right sizing and balancing:	
determining the performance of your	
environment	. 217
Using the Cognos data model	. 230

Chapter 10. Linux Kernel-based Virtual Machines agent: Using Performance and Capacity

Management Reports
More information about Tivoli Common Reporting 243
Linux KVM: Performance and Capacity
Management Reports descriptions
Prerequisites Checking report
Accounting report
Performance trends and resource forecasts 246
What-if analysis for workload placement reports 250
Workload right-sizing and balancing report 253

Chapter 11. Capacity Planner for

VMware	255
Planning Center for VMware wizard	. 255
Loading configuration data for physical servers	256
Scenario: Consolidate and plan for growth	. 258
Scenario: Adding more servers	. 265
Scenario: Adding new virtual machines	. 267
Cloning physical servers and virtual machines .	. 268

Cloning physical servers.	268
Cloning virtual machines	270
Detailed reference: Edit Current Environment .	273
Detailed reference: Edit Recommended	
Environment settings.	286
Defining resource sizing by using application	
sizing rules	297
Rule sizing	297
Detailed reference: Capacity efficiency and	
performance risk indices	299
Uploading benchmark and server model catalog	
files in Capacity Planner.	300
Uploading custom benchmark file in Capacity	
Planner	301
Use cases for network and storage aware capacity	
planning	301
Cognos based report packages	305
Capacity Planner Reports for VMware	
description	305
-	

Chapter 12. Capacity Planner for

enapter in eapaerry i lamer ier	
PowerVM	315
PowerVM PlanningCenter wizard	. 315
Configuring data for managed systems	. 317
Scenario: Providing a sizing recommendation of	
LPARs on systems.	. 319
PowerVM LPAR Sizing wizard: Selected	
workload page	. 326
PowerVM LPAR Sizing wizard: Workload	
definition page	. 328
PowerVM LPAR Sizing wizard: Selected system	
page	. 328
PowerVM Recommended Environment report	329
Scenario: Adding more servers	. 331
Scenario: Adding new logical partitions	. 333

Detailed reference: Edit Current Environment	
window	333
Detailed reference: Edit recommended environment	
settings	348
Defining resource sizing by using application	
sizing rules	351
Supported database paths for writing rules	353
Writing rules for custom attributes	354
Importing rules in Capacity Planner	355
Chapter 13. Troubleshooting	357
Installation and configuration problems	357
Timeout and memory problems	362
Performance and Capacity Management Reports	
problems	364
Dashboard for VMware problems	369
Capacity Planner problems	375
Appendix A. CSV format for data	381
Appendix B. Editing knowledge base	385
Appendix C. Capacity Planner	
database recommendations	389
Support Information	391
Notices	
The design of a	393
Irademarks	393 395
Privacy policy considerations	393 395 395
Privacy policy considerations	393 395 395 397

vi IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Figures

1.	ITM for Virtual Environments page	. 9 47
2.	Install Packages page	10
3.	Install Packages page	16 48
4.	Install Packages page: Common Configurations	17
5.	Install Packages page: PowerVM Capacity	49
	Planner Configuration	18
6.	Install Packages page: Summary Information	19
7.	Install Packages page	20 50
8.	Installed packages details	21
9.	Planning Center for PowerVM	22
10.	Choose the reports for the installation page	28 51
11.	Cognos Engine Configuration page	29 52
12.	Cognos Data Source	
	TADFDCP_REPORTING_DS Configuration	53
	page	30
13.	Report Installer Summary page	31 54
14.	Installation results page	32 55
15.	Delete icon	33 56
16.	ITM for Virtual Environments page	38 57
17.	Install Packages page	39 58
18.	Install Packages page	46 59
19.	Install Packages page	47 60
20.	Install Packages page: Common Configurations	47 61
21.	Install Packages page: Common Configurations	48 62
22.	Install Packages page: VMware Capacity	63
	Planner Configuration	49 64
23.	Install Packages page: Summary Information	50 65
24.	Install Packages page	51 66
25.	Installed packages details	52 67
26.	Planning Center for VMware	53 68
27.	Choose the reports for the installation page	59
28.	Cognos Engine Configuration page	60 69
29.	Cognos Data Source	70
	TADFDC_REPORTING_DS Configuration page	61 71
30.	Report Installer Summary page	62 72
31.	Installation results page	63 73
32.	Reports packages list	65 74
33.	Delete icon	65 75
34.	ITM for Virtual Environments page	67 76
35.	IBM Installation Manager page	68 77
36.	Update Packages page for existing package	78
	group	69 79
37.	Update Packages page for updates to install	70 80
38.	Update packages page for features to install	71 81
39.	Update Packages page to validate WebSphere	82
	Application Server Credentials	72 83
40.	Update Packages page to validate the Capacity	84
	Planner Database Schema details for VMware .	73 85
41.	Update Packages page to validate the Capacity	86
	Planner Database Schema details for PowerVM	74 87
42.	Update Packages page to show the summary	75 88
43.	Update Packages to show that the components	
	are updated	76
44.	Installed Packages and fixes page	77 89
45.	IBM Installation Manager	86
46.	Roll Back Packages page	87

47.	Roll back Packages page to select the	
	components to roll back	88
48	Roll Back Packages page to validate	
10.	Webenham Application Service Credentials	00
10	Websphere Application Server Credentials	. 09
49.	Roll Back Packages page to validate the	
	Capacity Planner Database Schema details for	
	PowerVM	90
50.	Roll Back Packages page to validate the	
	Capacity Planner Database Schema details for	
	VMware	91
E 1	Poll Pool: Dool: and poor to show the summary	02
51.	Kon back rackages page to show the summary	92
52.	Roll Back Packages to show that the	
	components are rolled back	. 93
53.	Install Packages page to show the currently	
	installed packaged	94
54.	Choose the reports for the installation page	104
55.	Cognos Engine Configuration page	105
56	Cognos Data Source TDW Configuration page	106
57	IDBC User Credentials tab	107
57.	IDBC Oser Credentials tab	107
50.		100
59.	Report Installer Summary page	109
60.	Installation results page	110
61.	Choose the reports for the installation page	113
62.	Cognos Engine Configuration page	114
63.	Cognos Data Source TDW configuration	115
64.	IDBC User Credentials tab	116
65.	IDBC Database Credentials tab	117
66	Report installer summary page	118
67	Installation results page	110
607.	Work with reports Calact the reports page	110
00.	work with reports: Select the reports package	100
60	that you want to delete	120
69.		120
70.	Data navigation tree	231
71.	Daily folder	232
72.	Sample dimensions	233
73.	Server filter screen	234
74.	Missing values filter screen	235
75.	Sample report	236
76	Cross tab	237
77	Chart wizard	237
79	Chart Wizard	201
70.		200
79.		238
80.	Prompt page	239
81.	TCR date range prompt	239
82.	Query Explorer	240
83.	Expression editor	240
84.	Play button	241
85.	Report values	241
86	Report results	242
87	Planning Center for VMware wizard	255
QQ	Spanshot config data with closed I and data	200
00.	for selected above entry local data	
	for selected physical servers and Clean	<u> </u>
	database before loading check boxes	257
89.	Snapshot config data with selected the Load	
	data for selected physical servers and cleared	
	Clean database before loading check boxes .	257

90.	Snapshot config data with selected the Load	
	data for selected physical servers and Clean database before loading check boxes	258
91.	Snapshot config data with cleared the Load	
	data for selected physical servers and	
	selected Clean database before loading check	
	hoves	258
92	Current Environment Physical Servers view	259
03	Add Custom Tag window	260
97 97	Physical Servers Inventory view Reports	200
71.	menu	261
05	Virtual Machines Utilization view	261
96	Virtual Machines Utilization view	201
<i>J</i> 0.	menu	262
97	Edit Usage window	262
08	Virtual Machines Utilization view Reports	202
<i>J</i> 0.	menu	263
99	Edit Candidate Tag window	263
100	Cenerate Optimization Plan section of	200
100.	Planning Center page	264
101	Table showing recommended resources	264
101.	Individual bar charts comparing current	201
102.	resources with recommended resources	265
103	Individual tabular results comparing current	200
100.	resources with recommended resources	265
104	Physical Servers Actions menu	266
105	Create Server Instances window	266
106	Optimization plan: adding servers	267
107.	Optimization plan including new virtual	207
	machine	268
108.	Edit VMware Current Environment page	269
109.	Clone Physical Servers window	269
110.	Clone Physical Servers second page	270
111.	Edit VMware Current Environment page for	
	virtual machines	271
112.	Edit VMware Current Environment Actions	
	tab	271
113.	Clone Virtual Machines window	272
114.	Clone Virtual Machines second page	272
115.	Current Environment Physical Servers view	274
116.	Benchmark Search Results window	275
117.	Benchmark Search Results window, single	
	match	275
118.	Capacity Planner Physical Server Inventory	
	report	277
119.	Current Environment Virtual Machines view	278
120.	Virtual Machine Inventory report	279
121.	Virtual Machines Utilization view	280
122.	Compute Usage window	281
123.	Generate Stability Characteristic Parameters	
	window	282
124.	Edit Usage window	282
125.	Utilization Aggregated Timeseries report	283
126.	Capacity Planner Utilization Detailed	
	Timeseries report	284
127.	Current Environment report	285
128.	Current Environment report (continued)	285

129.	Edit Recommended Environment Settings	
	window	286
130.	Refresh Knowledge Base Content window	294
131.	Minimize systems optimization without rules	295
132.	Minimize systems optimization with rules	296
133.	Refresh Knowledge Base Content window for	
	benchmark and server model catalog files	300
134.	Refresh Knowledge Base Content window for	
10 11	custom benchmark file	301
135	PlanningCenter for PowerVM interface	316
136	Planning Contor for PowerVM	317
130.	Soloctivo data load window	318
107.		210
130.	Success message	210
139.	Planning Center for PowerVM.	319
140.	Edit PowerVM Current Environment window	320
141.	Actions menu: Edit PowerVM Current	
	Environment window.	321
142.	Selected workload page	322
143.	Workload definition page	323
144.	Selected system page	324
145.	Selected system	324
146.	Selected system tab	325
147.	PowerVM LPAR Sizing	326
148.	Selected workload page	327
149.	Workload definition page	328
150.	Selected system page	329
151.	Summary of the PowerVM Recommended	
	Environment report	330
152.	Capacity Planner PowerVM Recommended	
	Environment report	331
153	Actions menu	332
154	Create Managed System Instances window	332
155	Managed System Inventory View	334
156	Create Managed System Instances window	335
150.	Edit PowerVM Current Environment page	336
157.	Clone Managed Systems window	226
150.	Clone Managed Systems willdow	227
160	Managad System Javantary View	221
160.	Managed System Inventory view	338
161.	Logical Partition Inventory view	338
162.	Edit PowerVM Current Environment page for	0.40
1 (0	logical partitions	340
163.	Edit PowerVM current Environment Actions	
	tab	340
164.	Clone Logical Partitions window	341
165.	Clone Logical Partitions second page	342
166.	LPAR Inventory report	343
167.	VIO Partition Inventory View	343
168.	Logical Partition Utilization View	344
169.	Compute Usage window	345
170.	Generate Stability Characteristic Parameters	
	window	346
171.	Edit Usage window	347
172.	Utilization Aggregated Timeseries report	348
173.	Edit PowerVM Recommended Environment	
	Settings window	349
174.	Refresh Knowledge Base Content window	355
	0	-

Tables

1.	Setup files for Tivoli Common Reporting 2.1.1	27
2.	Setup files for Tivoli Common Reporting 3.1 or	r
	later	. 27
3.	Setup files for Tivoli Common Reporting 2.1.1	58
4.	Setup files for Tivoli Common Reporting 3.1 or	r
	later	. 58
5.	Setup files for Tivoli Common Reporting 2.1.1	103
6.	Setup files for Tivoli Common Reporting 3.1	
	or later.	103
7.	Setup files for Tivoli Common Reporting 2.1.1	112
8.	Setup files for Tivoli Common Reporting 3.1	
0.	or later	112
9	Cluster scorecard	148
10	Situation Events table	151
11	Change history table	153
12	Server scorecard	154
12.	Situation Events table	157
13. 14	Change history table	157
15	Notwork table	150
15.	Virtual machine scorecard	160
10.	Situation Events table	160
17.	Change history table	163
10.	Change history table	164
19. 20	Situation Events table	160
20.	Situation Events table	100
21.		100
22.		169
23.	Server scorecard table.	171
24.	Health checks for liveli Monitoring	176
25.	Health checks for Virtual Infrastructure	1
•	Monitoring	176
26.	Attribute groups where historical collection	
	must be enabled for preconfigured reports to	100
~ 7	Work	183
27.	Attribute groups where historical collection	104
20	must be enabled for custom reporting	184
28.	Attribute groups and collection intervals	185
29.	VMware VI Number of DCs Clusters Hosts	105
20	and VMs Monitored report	185
30.	VMware VI Cluster Forecast Alerts report	186
31.	VMware VI Cluster Workload Irend and	107
22	Forecast report	187
32.	VMware VI Host Server Forecast Alerts	107
22	report	187
33.	VMware VI Host Server Workload Irend and	100
~ .	Forecast report	188
34.	VMware VI Host Server CPU Ready and	100
	Utilization report	189
35.	VMware VI VM CPU Ready and Utilization	
	report	190
36.	VMware VI Cluster Performance Trends	
	report	191
37.	VMware VI Host Server Performance Trends	
	report	192
38.	VMware VI VM Performance Trends report	193
39.	VMware VI Host Server Heat Chart report	195
40.	VMware VI VM Heat Chart report	196

	/1	VMware VI Host Server Utilization Details	
	т1.	report	107
	10		197
	42.	VMware VI Host Server Memory Allocation	
		Details report	198
	43.	VMware VI VM Utilization Details report	199
	44.	VMware VI Data Store Capacity and	
		Performance Trend report	200
	45.	VMware VI Cluster Weekly Comparison	
		report	201
	46	VMware VI Host Servers Weekly Comparison	
	10.	report	202
	17	What a Promodulate Scanner report	202
	47. 10	White Prefequisite Scalifier report	204
	40.	Charles and a second se	205
	10	Clusters report	205
	49.	VMware VI Number of Workloads for	
		Clusters or Host Servers report	208
,	50.	VMware VI Resources Needed for Additional	
		Workloads on Clusters report	211
,	51.	VMware VI Resources Needed for Additional	
		Workloads on Host Servers report	214
	52.	VMware VI Top or Bottom Workload	
		Consumers Clusters report	217
	53	VMware VI Top or Bottom Host Servers	-17
	00.	Weekly Comparison report	210
	E 1	Wayara VI Tan an Pattern Warkland	219
	94.	Viviware vi top or bottom workload	010
		Consumers Host Servers report	219
	55.	VMware VI Top N VMs by Host Server	
		report	221
	56.	VMware VI Bottom N VMs by Host Server	
		report	221
,	57.	VMware VI Top N VMs by Resource Pool	
		report	222
,	58.	VMware VI Bottom N VMs by Resource Pool	
		report	223
,	59.	VMware VI Balanced and Unbalanced	
		Clusters report	224
	60	VMware VI Balanced and Unbalanced Host	
	00.	Sarvars report	225
	61	VMware VI Top or Bottom Virtual Switches	225
	01.	by Network L/O	226
	\sim	VM www.VII Ten and Pattern Disciplination NIICo	220
	62.	VMware VI Top or Bottom Physical NICS	227
	63.	VMware VI Network Usage by VMs	229
	64.	VMware VI Host Server Memory Allocation	
		report	229
	65.	Linux KVM Report Prerequisite Scanner	245
	66.	Linux KVM Number of Hosts and VMs	
		monitored	246
	67.	Linux KVM Physical CPU Utilization by VM	247
	68.	Linux KVM Storage Pool Capacity and	
		Performance Trend.	249
	69.	Linux KVM Number of VMs for Host	251
	70	Linux KVM Resources Needed for Additional	
	, 0.	VMS on a Host	252
	71	Linux KVM Top or Bottom VMc by Physical	292
	/ 1.	CPU Utilization report	254
	70	Dhysical component the last	204
	12.	r nysical servers attributes	291

73.	Virtual machine attributes	1
74.	CSV format details	3
75.	Supported values in <path> tag of dependent</path>	
	literal	3
76.	Supported dynamic path in <term> tag of</term>	
	dependent literal	3
77.	Host level network capacity check 30	1
78.	Host level storage access check	2
79.	Host level storage capacity check	3
80.	Localized placement for a group of virtual	
	machines where groups are specified through	
	user input	1
81.	Capacity Planner Physical Server Inventory	
	report	5
82.	Capacity Planner Virtual Server Inventory	
	report	5
83.	Capacity Planner Utilization Detailed	
	Timeseries report	7
84.	Capacity Planner Utilization Aggregated	
	Timeseries report	7
85.	Capacity Planner Current Environment report 30	3
86.	Capacity Planner Optimized Environment	
	Plan report	3
87.	Capacity Planner Additional Workloads	_
	Analysis report)
88.	VMware Expense Reduction report 310)
89.	Capacity Planner Optimized Environment	_
	Plan Cluster Details report	2

90.	. Target fields required for application sizing	
	rule	352
91.	Logical partition attributes	354
92.	CSV format details	355
93.	Capacity Planner installation and	
	configuration problems	357
94.	Linux Kernel-based Virtual Machines agent	
	installation and configuration problems	357
95.	Reports installation and configuration	
	problems	358
96.	VMware agent reports problems	359
97.	Timeout and memory problems	362
98.	Performance and Capacity Management	
	Reports problems	365
99.	Dashboard for VMware problems	369
100.	General Capacity Planner problems and	
	solutions	376
101.	Capacity Planner for PowerVM problems and	
	solutions	377
102.	Capacity Planner for VMware problems and	
	solutions	378
103.	CFG_PHYSICAL_SERVER for PowerVM	381
104.	CFG_PHYSICAL_SERVER for VMware	382
105.	CFG_VIRTUAL_MACHINE for PowerVM	382
106.	CFG_VIRTUAL_MACHINE for VMware	383
107.	Server Model Catalog.	385
108.	User-defined benchmark	386
109.	Virtualization overhead	387

Chapter 1. Tivoli Monitoring for Virtual Environments structure

Tivoli[®] Monitoring for Virtual Environments structure consists of the Data collection, Server, Data analysis, and presentation layers.

Data collection layer

The *data collection layer* consists of Tivoli Monitoring agents, including Tivoli Monitoring for Virtual Environments agents that collect data about virtual environments, and other agents that collect data about storage, network, and guest operating systems.

For storage, Tivoli Monitoring for Virtual Environments is used to complement information about virtual storage, with minimal information about associated physical storage, and make this information available in the Dashboard for VMware:

- Information that provides a view of physical storage that is associated with the virtual environment
- The link between virtual and physical storage, which facilitates communication between the operator or administrator of the virtual environment and the storage administrator for problem resolution

Tivoli Monitoring for Virtual Environments provides predefined integration with IBM[®] Tivoli Storage Productivity Center (through the Tivoli Storage Productivity Center agent) and with NetApp DataFabric Manager (through the NetApp Storage agent).

For networks, Tivoli Monitoring for Virtual Environments is used to complement information about virtual networks, with minimal information about the physical network that is connected to the virtual environment, and make this information available in the Dashboard for VMware:

- Information that provides a basic view of the physical network that is associated with the virtual environment. With this information, a customer can understand the relationship between the virtual network and the physical switches (ports) connected to the hypervisor hosts, and can monitor the switches and ports that are being used by virtual machines of interest.
- The link between the virtual and physical network, to facilitate communication between the operator or administrator of the virtual environment and the network administrator for problem resolution

For guest operating systems, Tivoli Monitoring for Virtual Environments is used to complement information about the host, the hypervisor, and the virtual machines, with minimal information about the guest operating system and its processes and services that run in the virtual machines:

- Information that provides a basic view of how resources are being used by the various processes that run in the guest operating system
- The link between the virtual machine and the guest operating system, to facilitate communication between the operator or administrator of the virtual environment and the application owner for problem resolution

Server layer

Tivoli Monitoring for Virtual Environments provides a server layer that contains the following IBM Tivoli Monitoring server components:

- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Portal Server
- Tivoli Data Warehouse

For more information about these Tivoli Monitoring components, see the IBM Tivoli Monitoring Installation and Setup Guide (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/overview_components.htm).

Data analysis and presentation layer

Tivoli Monitoring for Virtual Environments provides a *data analysis and presentation layer* where most of the new values are displayed.

Dashboard for VMware

The Dashboard for VMware provides a summary view of the health of the whole environment (including the associated physical environment). You can then quickly assess problems and fix them.

The Dashboard for VMware brings together important information in context about components of the virtual environment from various sources. By using the Dashboard for VMware, you can view key components of the whole environment, including server, storage, and network (both virtual and physical), how they are related, and key attributes (configuration, metrics) in the context of the problem. Instead of overwhelming the operator or IT administrator with information, the Dashboard for VMware aggregates, associates, and summarizes data by showing only the key required information in context.

The Dashboard for VMware relies on situations from Tivoli Monitoring (including trending or forecasting) to determine the status of the environment that is monitored. Tivoli Monitoring for Virtual Environments provides a configuration file with a list of situations that are used for a specific hypervisor. You can add or remove situations from the list when you configure the Dashboard for VMware. When you apply the list, all situations in the list can be received by the Dashboard for VMware.

The Dashboard for VMware destination is configured as one of the EIF destinations for the Tivoli Monitoring situations, so that situations can be forwarded by using the **EIF sender** tab in the Tivoli Monitoring situation editor. All preselected situations that are listed in the configuration of Tivoli Monitoring for Virtual Environments are forwarded. To forward any situation that is not preselected, you must set the EIF destination for the situation by using the situation editor in Tivoli Monitoring, which is a standard behavior in that product.

Status in the Dashboard for VMware is organized in three columns: Server, Storage, and Network. In the data centers of most enterprises, operations and responsibilities are organized around these three areas. This determination is done by association of the attribute group that is related to the situation with the status column. This association is provided by Tivoli Monitoring for Virtual Environments and does not require any action from you.

From the Dashboards, you have access to other user interfaces. Two levels of UI integration exist in Dashboard for VMware:

- The Tivoli Application Dependency Discovery Manager (TADDM) change and configuration details view can be displayed in context in the Dashboard for VMware. Tivoli Monitoring for Virtual Environments uses the TADDM console view so that you can visualize change and configuration details about the resources that are selected in the Dashboard for VMware. The TADDM view gets data directly from the TADDM server (configured in Tivoli Monitoring for Virtual Environments).
- You can launch-in-context from the Dashboard for VMware or panels into Tivoli Monitoring and the IBM Dashboard Application Services Hub.

Capacity management

Tivoli Monitoring for Virtual Environments capacity management enables you to assess your current capacity usage, identify bottlenecks, and forecast usage constraints. You can then support your planning activities to address any issues that are identified, or to better consolidate and optimize environment and plan for growth in terms of capacity.

Performance and Capacity Management Reports

Tivoli Monitoring for Virtual Environments includes predefined reports that help you to understand current capacity. The Performance and Capacity Management Reports are based on IBM Cognos[®] software and work with Tivoli Common Reporting. The reports are based on historical data that is collected in the Tivoli Data Warehouse and user inputs. In addition to the predefined reports, you are also provided with a data model and tools for creating ad hoc reports. The reports are available from the Tivoli Integrated Portal tasks. The following three sets of reports are defined:

- Workload Estimation: Simple reports to gauge how much more workload (virtual machines) the environment can handle, and where to place them based on demand. These reports are useful for IT engineers who deal with everyday capacity management tasks.
- Performance Trends and Resource Forecasts: Reports that show historical trends and future forecasts of key metrics. These reports are useful for IT engineers for comparing historical resource usage trends and identifying short-term future bottlenecks.
- Workload Right Sizing and Balancing: Reports that provide a holistic view of the environment that is geared towards IT managers and IT engineers. These reports provide a high-level understanding of the top or bottom resource consumers, and show how well-balanced the environment is.

With these reports provided by Tivoli Monitoring for Virtual Environments, you can answer questions such as:

- How is my environment performing overall?
- Which are my most used servers for a specific resource type?
- Do bottlenecks exist in my current environment, and where?
- Am I reaching capacity on resources, and on which resource? When will I exhaust capacity?
- Which are my top or bottom virtual machine resource consumers for a specific resource type?
- Which are my least used servers for a specific resource type?
- Has any abnormal behavior been detected this week compared with last week (or other period)?
- Are my systems or workloads balanced or unbalanced?

Capacity Planner for PowerVM[®] and Capacity Planner for VMware

With the Capacity Planner tool, you can answer the following questions:

- How many more virtual machines can I add to a cluster or server, which is based on usage history?
- How much more resource do I need to add more virtual machines to the environment?
- How and where do I add capacity if existing systems are not enough for future growth for optimized capacity usage?
- Where do I place new workloads? Must I add more resources?
- How can I optimize the virtual machine placement to maximize usage and minimize costs?
- How can I optimize the application placement to maximize usage and minimize costs?

Capacity planning involves a five-step process that is guided by a wizard. During these steps, you import data for analysis, select the scope of analysis, time, servers that you want to work with, understand the characteristics of the servers or virtual machines, observe trends and patterns. In the *recommendation generation* step, you define the objectives or results to achieve with planning activity, the policies (rules) to be applied during analysis, and the analysis output characteristics (reports), and generate the recommendation. The reports can be visualized in the Dashboard for VMware.

The IBM Infrastructure Management Capacity Planner for PowerVM and IBM Infrastructure Management Capacity Planner for VMware tool uses usage data available in Tivoli Data Warehouse. IBM Infrastructure Management Capacity Planner for PowerVM and IBM Infrastructure Management Capacity Planner for VMware creates a data mart in the DB2[®] database (the database server can be shared with Tivoli Data Warehouse). Capacity planning requires DB2 federation capability, because it is not copying utilization data from Tivoli Data Warehouse is using data that ise already captured in Tivoli Data Warehouse (summarized data).

Integration with other products

Tivoli Monitoring for Virtual Environments also functions in a broader context and integrates with various products.

IBM Tivoli Netcool/OMNIbus

Tivoli Monitoring for Virtual Environments forwards situations to OMNIbus by using the Tivoli Monitoring base mechanism. You can use Tivoli Monitoring tooling to define which situations to forward. Good examples include situations about virtual machine changes, including deletion, creation, and movements of virtual machines.

IBM Tivoli Application Dependency Discovery Manager

Tivoli Monitoring for Virtual Environments discovery library adapters can be imported into TADDM and reconciled within TADDM with other sources, including the sensors for virtual platforms or hypervisors. A second integration with TADDM is the ability to dynamically update TADDM about virtual machine changes.

IBM Tivoli Business Service Manager

Tivoli Monitoring for Virtual Environments discovery library adapters can also be imported into Tivoli Business Service Manager for services management. Situations that are generated by Tivoli Monitoring for Virtual Environments can also be used in Tivoli Business Service Manager for services impact determination.

IBM System Director

From IBM Tivoli Monitoring for Virtual Environments, you can launch-in-context to the administration UIs of IBM Systems Director to administer and configure the VMware hypervisor. The launch-in-context is available from the VMware cluster, the ESX servers workspaces in the Tivoli Enterprise Portal, or both.

New in this release

The dashboard, reporting, and capacity planning components contain enhancements that were made since Version 7.2, including Fix Pack 1.

For the Version 7.2 Fix Pack 2 dashboard, reporting, and capacity planning components, the following enhancements were made:

- Capacity Planner for PowerVM
 - Ability to copy or clone managed systems and logical partitions (LPARs)

You can copy existing managed systems and LPARs and create new workloads with similar configuration and utilization profiles for what-if analysis.

 Ability to define policies for setting resource usage that is based on application parameters by using sizing rules

With utilization rules available in a previous release, you were able to update computed usage data for LPARs that matched the rule criteria according to business requirements. With the updated sizing rules, you can set the size of resource usage to match configuration or application parameters. The parameters can be discovered attributes or custom tags to represent application characteristics.

Persistent session lock to prevent other users from modifying data in a planning session
 The persistent session lock prevents other users from logging in to the Capacity Planner for a hypervisor and modifying or overwriting data in your planning scenario. The lock ensures that other users cannot use the Capacity Planner functions until the lock is released by the user who acquired it or forcibly released by an administrator.

- Federation script enhancements

The federation script output is improved to generate a report on the status of federation configuration. The updated report shows the views for which the configuration was successful and the missing summarization levels for resources in case of failures.

- Capacity Planner for VMware
 - Ability to copy or clone physical servers and virtual machines

It is now possible to copy existing physical servers and virtual machines, and create new workloads with similar configuration and utilization profiles for what-if analysis.

 Ability to define policies for setting resource usage that is based on application parameters by using the Sizing Rules

With previously available sizing rules, you were able to update computed usage data for virtual machines that matched the rule criteria per the business requirement. With the updated application sizing rules, you can set the size of resource usage that is based on matching configuration or application parameters. The parameters can be the discovered attributes or custom tags to represent application characteristics.

- Persistent session lock to prevent other users from modifying data in a planning session

The persistent session lock prevents other users from logging in to the Capacity Planner for a hypervisor and modifying or overwriting data in your planning scenario. The lock ensures that the other users cannot use the Capacity Planner functions until the lock is released by the user who acquired it or the lock can be released by an administrator.

• Enhanced the federation script

The federation script output is improved to generate a report on the status of the federation configuration. The updated report shows the views for which the configuration was successful and the missing summarization levels for resources in case of failures.

- Capacity Planner Reports for VMware
 - Added the Capacity Planner Optimized Environment Plan Cluster Detail report.
 - Enhanced the Capacity Planner Optimized Environment Plan report to group the actionable recommendation at cluster level that enables the recommendations to be more actionable.
- Virtual Environments VMware Reports
 - Added the VMware VI Top or Bottom Host Servers Weekly Comparison report.
 - Enhanced the following reports:
 - VMware VI Host Server Heat Chart
 - VMware VI VM Heat Chart
 - VMware VI Top or Bottom Workload Consumers Clusters
 - VMware VI Cluster Performance Trends
 - VMware VI Cluster Weekly Comparison report
 - VMware VI Cluster Workload Trend and Forecast
 - VMware VI Balanced and Unbalanced Clusters
 - VMware VI Bottom N VMs by Resource Pool
 - VMware VI Top N VMs By Resource Pool
 - VMware VI Top or Bottom Physical NICs
 - VMware VI Top or Bottom Workload Consumers Clusters
 - VMware VI Top or Bottom Workload Consumers Host Servers
- Dashboard for VMware

Enhanced the Dashboard for VMware with the following features:

- Added a Copy URL feature.
- Added a launch to Tivoli Enterprise Portal feature that can be used from all component details pages.

- Added more graphical data and two sets of charts where one chart shows absolute values and other chart shows percentage values.
- Added a data stores situation to the Cluster, Server, and Virtual Machines group tables.
- Updated some names and labels to be more intuitive.
- Updated the Situations Events table to add a source event link to open the component (source) details page.
- Updated the Server details page to separate the Link chart into Link Speed (MBps) and Link Utilization (%) charts for the Network tab.
- Updated all components details pages such that loading of all data is avoided on details pages that have multiple tabs until a particular tab is selected.
- Updated scorecard tables to hide Sort Order Number (1, 2, 3) to avoid confusion with the aggregate number of situations.
- Updated the Resource Relationships view such that when only one component (for example, a virtual machine) is related to the selected component (for example, server), the name of that component is displayed. By clicking the name, you can view the related component details page.

Chapter 2. PowerVM: Installing components

To install the PowerVM components, you must install the prerequisites, select a type of installation depending on your environment, complete the selected procedure to install the Capacity Planner for PowerVM and install the capacity planner reports.

PowerVM: Installing prerequisites

Before you complete your installation, you must install the required prerequisites.

Procedure

- 1. Install Jazz[™] for Service Management V1.1 or later. You must select the following Jazz for Service Management components that are required for the Capacity Planner for PowerVM:
 - WebSphere[®] Application Server V8.5 or later
 - Jazz for Service Management Extension to WebSphere Application Server V8.5 or later
 - Dashboard Application Services Hub V3.1 or later
- Create a new database for both the local and remote installations of the Capacity Planner database by using the following command: db2 create database TADFDCDB USING CODESET UTF-8 TERRITORY US COLLATE USING SYSTEM PAGESIZE 16384.

Important: This database is shared between the Capacity Planner for PowerVM and the Capacity Planner for VMware. If you install both of these capacity planners, you must create the database once only. You can create the capacity planner database as the db2 instance owner on one of the following systems:

- On Windows, the instance owner is **db2admin**.
- On Linux and UNIX, the instance owner is **db2inst1**.

The Post Install Federation step might fail if the Capacity Planner database is created as a user with non compliant authorization levels.

3. To install and configure the base IBM Tivoli Monitoring monitoring server, portal server, and portal client components, see the: IBM Tivoli Monitoring Installation and Setup Guide (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/itm_install.htm).

Remember: When you configure the portal server, you must enable the dashboard data provider.

4. Install and configure the monitoring agents so that the monitoring data displays in the Capacity Planner application. If you want to monitor the virtual environment remotely, install IBM Tivoli Monitoring Linux or UNIX OS Agents V6.3 or later on the AIX[®] LPARs that are monitored. Then, you install the VIOS agents on the VIOS LPARs, and you install the HMC agent to monitor the HMC remotely. For more information about installing the application support in the monitoring servers, portal server, and desktop portal client, see the IBM Tivoli Monitoring Installation and Setup Guide.

PowerVM: Types of installation

You can use either the Launchpad or the IBM Installation Manager to install the PowerVM components:

The PowerVM components include:

- IBM Infrastructure Management Capacity Planner for PowerVM
- IBM Infrastructure Management Capacity Planner Reports for PowerVM

Select the type of installation that suits your environment:

• Launchpad

If you have a browser and a graphics environment, select the launchpad to complete your installation. See Using the launchpad to install PowerVM components.

• Installation Manager GUI

If you do not have a browser but you have a graphics environment, use the Installation Manager GUI to complete your installation. See Using the IBM Installation Manager to install Power[®] VM components.

• Installation Manager Silent

If you do not have a browser or a graphics environment, or you want to complete a silent installation, use the Installation Manager silent type of installation to install the PowerVM components. See "PowerVM: Using the IBM Installation Manager in silent mode to install the Capacity Planner " on page 11.

• Installation Manager Console

If you want to use IBM Installation Manager in console mode to install packages from a command line interface, use the Installation Manager console type of installation. See "PowerVM: Using the IBM Installation Manager Console to install the Capacity Planner " on page 11.

Note: Both the Jazz for Service Management components and the IBM Tivoli Monitoring for Virtual Environments components for a dashboard environment must be installed in the same user mode; install both as a root user or install both as a non-root user. Information cannot be exchanged between these components when there is a combination of root and non-root installations.

PowerVM: Using the Launchpad to install components

You can use the launchpad to start the IBM Installation Manager to install the Power VM components.

Before you begin

Ensure that you have a supporting browser that is installed in your target environment, because the Launchpad requires a browser to run. The following browsers are supported:

- Firefox 2.0 or later, Firefox 3.0 or later and Firefox 3.5 or later
- Internet Explorer 6.0 or later
- Mozilla 1.7 or later
- SeaMonkey 1.1.4 or later and SeaMonkey 2.0 or later

Attention: If your target system does not have a browser that is installed, use Installation Manager in GUI mode.

Ensure that you do not have multiple instances of the Launchpad open, because this might cause a conflict to occur.

When you first download the Launchpad image and extract to a directory on the local system, ensure that the path to the extracted Launchpad package does not contain any spaces; otherwise the Launchpad will not run.

Procedure

- Download and extract the dashboard image file for the launchpad that contains the IBM Installation Manager repository, and Reports Installer. You can download this image to any directory on your local system.
- 2. Start the Launchpad.

Important: The best practice is to have only one instance of the Launchpad open at a time.

Option	Description
On Windows systems	Browse to the local \Launchpad directory, and run launchpad64 for 64-bit Windows computers.
On Linux or AIX systems	In a command window, open the local \Launchpad directory and run ./launchpad.sh .

These commands start the IBM Installation Manager GUI.

- **3**. Choose one of the following options:
 - To install the capacity planners and reports as an administrative user, click **Full Installation as administrative user**.
 - To install the capacity planners and the reports as a non-administrative user, click **Full Installation** as non-administrative user

ITM for Virtual E	nvironments	Select a language: English	• OK
	Installation Options This product contains installation con and reports. You can use this launch perform a full installation or install co Full Installation as administrativ Full Installation as non-administ You will be able to choose which cap Install Capacity Planners and Da	pponents for capacity planners, dashboards bad to install them. Click the links below to imponents individually. re user irative user acity planners and dashboards to be installed. shboards as administrative user	
	Install Capacity Planners and Da Install Reports View Release Information Readme Installation Guide	shboards as non-administrative user	
IBM.			

Figure 1. ITM for Virtual Environments page

What to do next

To complete the installation, see the step to select a component to install in the "PowerVM: Using the IBM Installation Manager GUI to install components" procedure. Also see the "PowerVM: Installing Capacity Planner Reports " on page 27 procedure.

PowerVM: Using the IBM Installation Manager GUI to install components

You start the Installation Manager to install the PowerVM components.

Procedure

- 1. To start the Installation Manager, do one of the following steps:
 - On Windows systems, select Start > IBM Installation Manager > IBM Installation Manager.
 - On Linux or AIX systems, in a command window, open the /opt/IBM/InstallationManager/ eclipse directory and enter ./IBMIM or ./groupinst for group mode.
 - Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where the IBM Installation Manager is installed. The default path for IBM Installation Manager on Windows is C:\Program Files\IBM\Installation Manager or C:\Program Files(86)\IBM\Installation Manager and on UNIX it is /opt/IBM/Installation Manager.
- 2. Click **File** > **Preferences**. Then click **Add**.
- 3. Enter the path or click **Browse**.
- 4. In the **Repository** field, enter the file path by clicking **Browse** to point to the directory where you extracted the installation image and select the diskTag.inf file, and then click **OK**.
- 5. Select a component to install. For more information about the components to install, see "PowerVM: Installing the Capacity Planner" on page 16. Then, click **Next**.

Results

>	IBM Installatio	on Manager		<u>+</u>
Install Packages				~
Select packages to install:				4
Installation Packages		Status	Vendor	License Ke
🗢 🗹 😫 IBM Infrastructure	e Management Capacity Planner for	PowerVM		
🗹 🙀 Version 7.2.0.2		Will be installed	d IBM	
🗢 🗌 🔖 IBM Infrastructure	e Management Capacity Planner for	VMware		
🗌 🙀 Version 7.2.0.2			IBM	
🗢 🗌 🔖 IBM Infrastructure	e Management Dashboard for VMwa	are		
🗌 🕌 Version 7.2.0.2			IBM	
<u>.</u>	m			
 ✓ C Show <u>all versions</u> 	10	Check for O	ther Versions, Fixes	s, and Extensions
 ✓ Show <u>all versions</u> 	10	Check for O	ther Versions, Fixes	s, and Extensions
 Show <u>all versions</u> Details IBM Infrastructure Mar 	magement Capacity Planner fo	Check for O	ther Versions, Fixes	i, and Extensions
Show <u>a</u> ll versions Show <u>a</u> ll versions Details IBM Infrastructure Manage	nagement Capacity Planner fo ement Capacity Planner for PowerVM	⊆heck for O Check for O r PowerVM 7.2.0.2 4 installer <u>More info</u>	ther Versions, Fixes	and Extensions
Show <u>a</u> ll versions Details IBM Infrastructure Manage Repository: /tmp/cicvolo	magement Capacity Planner fo ement Capacity Planner for PowerVM eache_root/com.ibm.tivoli.monitorir	<u>C</u> heck for O r PowerVM 7.2.0.2 A installer <u>More info</u> ng.powervmcp_7.2.100.20130	ther Versions, Fixes 723_0241/md	s, and Extensions
Show <u>a</u> ll versions Details IBM Infrastructure Manage Repository: /tmp/cicvolc	magement Capacity Planner fo ement Capacity Planner for PowerVM cache_root/com.ibm.tivoli.monitorir	<u>C</u> heck for O r PowerVM 7.2.0.2 4 installer <u>More info</u> ig.powervmcp_7.2.100.20130	ther Versions, Fixes 723_0241/md	s, and Extensions
Show <u>a</u> ll versions Details IBM Infrastructure Manage Repository: /tmp/cicvolo	nagement Capacity Planner fo ement Capacity Planner for PowerVN cache_root/com.ibm.tivoli.monitorir	Check for O r PowerVM 7.2.0.2 M installer <u>More info</u> Ig.powervmcp_7.2.100.20130	ther Versions, Fixes 723_0241/md	, and Extensions
Show <u>a</u> ll versions Details IBM Infrastructure Manage Repository: /tmp/cicvolc	magement Capacity Planner fo ement Capacity Planner for PowerVM eache_root/com.ibm.tivoli.monitorir	<u>Check for O</u> r PowerVM 7.2.0.2 A installer <u>More info</u> ng.powervmcp_7.2.100.20130	ther Versions, Fixes 723_0241/md	s, and Extensions

Figure 2. Install Packages page

Note: The IBM Installation Manager checks for license, languages, and the location of installation. You are prompted for configuration parameters. For more information about configuration parameters, see "Configuring a connection for the Capacity Planner for VMware and the Capacity Planner for PowerVM" on page 121.

PowerVM: Using the IBM Installation Manager in silent mode to install the Capacity Planner

There are specific installation steps to install the Capacity Planner for PowerVM by using the silent mode.

Before you begin

- Ensure that the Jazz for Service Management V1.1 or later is installed with the following three required components:
 - IBM Websphere Application Server V8.5 or later
 - Jazz for Service Management extension for IBM WebSphere V8.5 or later
 - Dashboard Application Services Hub V3.1 or later
- If you want to install Capacity Planner for PowerVM as a non-root user then Jazz for Service Management must be installed as a non-root user.
- Ensure that no other Installation Manager process is running before you install the offering silently.

Procedure

- 1. Download and extract the dashboard image. You can download this image to any directory on your local system.
- 2. Go to the silentInstall directory.
- 3. In the directory, modify the following silent response file: install_PowervM_CP_sample.xml
- 4. Run the following silent installation scripts:

Installation as Administrator

On Windows systems: silentInstall_PowerVM_CP_win.bat

On UNIX systems: silentInstall_PowerVM_CP_unix.sh

Installation as Non-Administrator

For Window systems: silentInstall_powerVM_CP_win.bat user

For UNIX systems: silentInstall_PowerVM_CP_unix.sh user

Note: If you want to run the script for a non-root user then, you must pass one variable user to the script. The user does not mean the actual non-root user name. It must be only the user.

PowerVM: Using the IBM Installation Manager Console to install the Capacity Planner

There are specific installation steps to install the Capacity Planner for PowerVM by using the Console mode.

Procedure

- 1. Start the console mode.
 - a. Open the command line.
 - b. From the command line, navigate to the IBM Installation Manager tools directory.
 - On Windows systems, the tools directory is in the following location: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

On operating systems other than Windows, the tools directory is in the following location: opt/IBM/InstallationManager/eclipse/tools

c. Run the following command: On Windows systems

imcl.exe -c

For example,

C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools>imcl.exe -c

On operating systems other than Windows, the tools directory is in the following location: ./imcl -c

2. Add the repository. The command line displays the following options:

====> IBM Installation Manager

Select:

- Install Install software packages
 Update Find and install updates and fixes to installed software packages
 Modify Change installed software packages
 Roll Back Revert to an earlier version of installed software packages
 Uninstall Remove installed software packages

 Other Options:
 - L. View Logs
 - S. View Installation History
 - V. View Installed Packages
 - P. Preferences
 - -----
 - E. Export Data for Problem Analysis
 - A. About IBM Installation Manager
 - X. Exit Installation Manager
- a. On the command line, type **P** to set the preferences. The command line displays the preferences.

====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- 2. Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
 6. Passport Advantage
- 7. Updates
- R. Return to Main Menu
- b. Type **1** to select the repository. The command line displays other options.

====> IBM Installation Manager> Preferences> Repositories

Repositories:

Other Options:

D. Add Repository

- S. [X] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- c. Type **D** to add the repository.
- d. Type the repository location. For example, C:\v7.2.0.2\diskTag.inf.

====> IBM Installation Manager> Preferences> Repositories> Add repository

Enter a new repository location.

To skip, press Enter: ----> C:\v7.2.0.2\diskTag.inf Checking repositories...

====> IBM Installation Manager> Preferences> Repositories

Repositories:

1. [X] C:\v7.2.0.2\diskTag.inf

Other Options:

- D. Add Repository
- S. [X] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- e. Type S to clear the option "Search service repositories during installation and updates".

====> IBM Installation Manager> Preferences> Repositories

Repositories:

1. [X] C:\v7.2.0.2\diskTag.inf

Other Options:

D. Add Repository

- S. [] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- f. Type **A** to apply changes and return to preferences menu.

====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- 2. Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
- 6. Passport Advantage
- 7. Updates
- R. Return to Main Menu
- g. Type **R** to return to main menu.
- 3. Selecting the package to install. The command line displays the following option:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- -----
- P. Preferences
- _____
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- X. Exit Installation Manager
- a. Type **1** to install the packages.

```
Checking repositories...
   Loading repositories...
   Checking availability of packages...
   Loading fixes...
   ====> IBM Installation Manager> Install
   Select packages to install:
        1. [] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
        0. Check for Other Versions, Fixes, and Extensions.
        C. Click Cancel.
b. Type 1 to install IBM Infrastructure Management Capacity Planner for PowerVM.
   ====> IBM Installation Manager> Install> Select
   IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
   Options:
        1. Choose version 7.2.0.2 for installation.
        2. Show all available versions of the package.
        C. Cancel
c. Type 1 to select version 7.2.0.2.
   ====> IBM Installation Manager> Install
   Select packages to install:
        1. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
        0. Check for Other Versions, Fixes, and Extensions
        N. Next,
                      C. Cancel
d. Type n to find compatible package groups. The command line displays the details of the
   compatible package groups.
   Finding compatible package groups...
   ====> IBM Installation Manager> Install> Location
   Existing package groups:
        1. [X] Core services in Jazz for Service Management
   Incompatible existing package groups:
        2. IBM WebSphere Application Server V8.5.0.1
   Selected group id: "Core services in Jazz for Service Management"
   Selected location: "C:\Program Files\ibm\JazzSM"
   Selected architecture: 64-bit
        B. Back,
                      N. Next,
                                    C. Cancel
e. Type n. The command line displays the features of the packages.
   ====> IBM Installation Manager> Install> Location> Features
   IBM Infrastructure Management Capacity Planner for powerVM
        1. [X] Installation
        2. [X] Configuration
        B. Back,
                      N. Next,
                                    C. Cancel
```

- f. Type **n** to continue the installation.
- 4. Validate Dashboard Application Services Hub credentials. The command line displays the following options:

====> IBM Installation Manager> Install> Location> Features> Custom panels
---- Common configuration:
User name for WebSphere Application Server:

- a. Type the user name for the WebSphere Application Server. For example, **tipadmin**. Password for WebSphere Application Server:
- b. Type the password for the WebSphere Application Server.

---- Configuration validation:

Validating WebSphere Application Server credentials... Provided credentials are valid.

5. Validate the Capacity Planner database credentials.

---- Configuration for IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2 Capacity Planner Database Server Host name:

- a. Type the Capacity Planner Database Server Host name. For example, **localhost**. Capacity Planner Database Server Port:
- b. Type the Capacity Planner Database Server Port. For example 50000.
 Capacity Planner Database Admin name:
- c. Type the Capacity Planner Database Admin name. For example, db2admin. Capacity Planner Database Admin Password:
- d. Type the Capacity Planner Database Admin Password.Capacity Planner Database Name:
- e. Type the Capacity Planner Database Name. For example, **TADFDCDB**. The system automatically validates the data that you entered. If the validation fails, enter the correct details again.
- 6. Check the summary of the installation. The command line displays the summary of the packages that you are going to install.

```
====> IBM Installation Manager> Install> Location> Features> Custom panels> Summary
```

```
Target Location:

Package Group Name : Core services in Jazz for Service Management

Installation Directory : C:\Program Files\ibm\JazzSM

Shared Resources Directory : C:\Program Files (x86)\IBM\IMShared

Packages to be installed:

IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
```

Options:

G. Generate an Installation Response File.

B. Back, I. Install, C. Cancel

a. Type **I** to install the packages.

----> [I]

25% 50% 75% 100%

====> IBM Installation Manager> Install> Location> Features> Custom panels> Summary> Completion

There were problems during the installation. WARNING: Administration Service Task Bundle Registration Failed to register resource. Returned status: Failed to register resource. Returned status:. Options:

- F. Finish
- b. Type **F** to finish the installation.

PowerVM: Installing the Capacity Planner

Installing the Capacity Planner for PowerVM requires specific installation steps.

Procedure

- 1. Select the type of installation. For more information on selecting the type of installation, see "PowerVM: Types of installation" on page 7.
- 2. Start the launchpad or Installation Manager GUI. For more information, see "PowerVM: Using the Launchpad to install components" on page 8.
- 3. Select the following check box: IBM Infrastructure Management Capacity Planner for PowerVM.

Note: When you select a component, the Installation Manager automatically selects the version of the respective component that you selected.

- 4. Click Next.
- 5. The Installation Manager displays the existing package groups. Verify the details, and then click **Next**.

IBM Ins	tallation Manager	
I nstall Packages A package group is a location that contains one or mo group only and will share a common user interface. Se	ore packages. Extensions can be installed into elect an existing package group for the exten	o a common package
Install Location Features Sum	mary	
<u>U</u> se the existing package group <u>C</u> reate a new package group		
Package Group Name	Installation Directory	Architecture
🐛 Core services in Jazz for Service Management	/opt/IBM/JazzSM	64-bit
Package Group Name: Core services in Jazz for Servic	ce Management	
Installation Directory: /opt/IBM/JazzSM		Browse
Architecture Selection: O 32-bit 64-bit		
Details • Shared Resources Directory: /opt/IBM/IBMIMShared • Eclipse IDE: /opt/IBM/JazzSM Installed Packages	Disk Space Inform Volume Available / 23.36 GB	ation Space
 IBM Dashboard Application Services Hub 3.1.0.1 		

Figure 3. Install Packages page

6. Enter the WebSphere Application Server Username and Password for Dashboard Application Services Hub, and then click **Validate**.

7. After the Username and Password are validated, click Next.

tinue.
Features Summary
Common Configurations WebSphere Application Server Credentials
User name: tipadmin Password: •••••• Validate
< <u>B</u> ack <u>N</u> ext > <u>Install</u> Cancel

Figure 4. Install Packages page: Common Configurations

- 8. Enter the Capacity Planner Database Schema Creation details, and then click Validate.
- 9. After the details are validated, click Next

٥	IBM Installation Manager	_ = ×
Install Packages Press the Validate button to con Install Location	tinue. Features Summary	
 ♥ Common Configurations ♥ WebSphere Application ♥ IBM Infrastructure Manage ♥ Capacity Planner Datab 	Configuration for IBM Infrastructure Mana PowerVM 7.2.0.2 Capacity Planner Database Schema Creation Capacity Planner Database Server Host name: Capacity Planner Database Server Port: Capacity Planner Database Admin name: Capacity Planner Database Admin Password: Capacity Planner Database Name: Validate	Iocalhost 50000 db2inst1 ••••••• TADFDCDB
< <u> </u>	< <u>B</u> ack	Next > Install Cancel

Figure 5. Install Packages page: PowerVM Capacity Planner Configuration

10. Review the information in the Install Packages Summary window, and then click Install.

	IBM Installation Manager		
Install Packages			
Review the summary info	rmation.		7
Install Location	Features Summary		
Target Location			
Package Group Name:	Core services in Jazz for Service Management		
Installation Directory:	/opt/IBM/JazzSM		
Shared Resources Directo	ory: /opt/IBM/IBMIMShared		
Packages	Sur Barra Barra Barra Contra en Contra en Contra en Contra de Contra de Contra de Contra de Contra de Contra de		
Packages			
🕼 Installation			
Configuration	Disk Space Information		
Configuration Configuration	Disk Space Information	Total Availa	ble Space
Configuration Configuration	Disk Space Information	Total Availa	ble Space 23.32 GB
Configuration Configuration	Disk Space Information / Total Download Size: 168.53 MB Total Installation Size: 343.99 MB	Total Availa	ble Space 23.32 GB
Configuration Environment English Repository Information	Disk Space Information / Total Download Size: 168.53 MB Total Installation Size: 343.99 MB	Total Availa	ble Space 23.32 GB

Figure 6. Install Packages page: Summary Information

11. The Results window displays the package installation status. Click **Finish** to complete the installation.



Figure 7. Install Packages page

12. To verify the installation, click File > View Installed Packages.

1 Installed Packages **Installed Packages and Fixes** Installed Packages Vendor License ▼ %, IBM WebSphere Application Server V8.5 івм IBM WebSphere Application Server_8.5.0.1 Jazz for Service Management extension for IBM WebSphere 8.5 1.1 IBM Core services in Jazz for Service Management IBM Dashboard Application Services Hub_3.1.0.1 IBM IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2 IBM 💫 IBM Infrastructure Management Capacity Planner for VMware_7.2.0.2 IBM IBM Infrastructure Management Dashboard for VMware_7.2.0.2 IBM > Details **IBM WebSphere Application Server V8.5** Shared Resources Directory: /opt/IBM/IBMIMShared Installation Directory: /opt/IBM/WebSphere/AppServer Eclipse IDE: /opt/IBM/WebSphere/AppServer Translations: English

Figure 8. Installed packages details

13. Log in to the Dashboard Application Services Hub to verify that you can start the **Planning Center for PowerVM** by using the following syntax: https://dash-server-hostname:16311/ibm/console. You replace dash-server-hostname with the host name of the server in your environment.

Close

Flamming Center for Fower VM	×		-
Planning Center for	PowerVM		
You have acquired the P	anning Center for PowerVM Io	ck at 25 Jul 2013 06:46:28 AM EDT. You need to explicitly release the lock for o	other
to use the Planning Cen	er for PowerVM. Release Lo	ock	
Any change in steps 1 to	needs a re-generation of the	plan in step 5 to view the latest recommendation.	
Step 1: Snapshot conf	g data.		
Load the latest configuration da	a for managed systems and I	logical partitions for analysis. You can change this data for what-if analysis.	
Advanced options: Select the data load options:			
Load data for selected Mar Clean database before loa	aged Systems ding		
Load Config			
Step 2: Set analysis ti	ne period.		
Set the time period for which the measurement data is federated	measurement data in the wa from the warehouse.	rehouse can be analyzed corresponding to the logical partitions loaded in Ste	ep 1.
Set Time			
Stop 2: Scope the infr	structure for analysis	e	
While scoping you can also edi systems and logical partitions, Define Scope	the current configuration on th f required.	1e <u>Edit Current Environment</u> page to add new attributes or clean the data for n	nana
Step 4: Analyze logica	partition characteris	stics.	
Analyze the warehouse data wit weekly utilization).	nin the time limit set in Step 2	to compute the LPAR-level sizing characteristics using default settings (maxi	mum
Advanced options: Experts can customize sizing o logical partitions based on mea	the <u>Edit Current Environment</u> surement data.	t page. While on this edit page, you can trigger several custom actions to char	acter
Analyze LPARs			
Current Environment Report			
Step 5: Generate reco	nmended sizing plan.		
Any change in steps 1 to	needs a re-generation of the	plan in step 5 to view the latest recommendation.	
Generate a recommended sizir	g plan based on recommende	ed environment settings.	
Advanced options: Custom settings can be made	n the <u>Edit Recommended En</u>	vironment Settings page where experts can select sizing rules.	
Generate Plan			
Recommended Sizing Plan Re	port		

Figure 9. Planning Center for PowerVM

PowerVM: Using the IBM Installation Manager GUI to uninstall the Capacity Planner

You can uninstall the Capacity Planner for PowerVM by using the IBM Installation Manager GUI.

Procedure

- 1. Start the IBM Installation Manager, and click Uninstall.
- 2. On the Uninstall Packages page, select the following PowerVM component, and then click Next.IBM Infrastructure Management Capacity Planner for PowerVM
- 3. Enter the credentials for WebSphere Application Server, and then click Validate.
- 4. After the credentials are verified, click Next.
- 5. Review the summary of the components that you want to uninstall.
- 6. Click Uninstall.

PowerVM: Using IBM Installation Manager Console to uninstall the Capacity Planner

There are specific steps to uninstall the Capacity Planner for PowerVM by using the Console mode.

Procedure

- 1. Start the console mode.
 - a. Open a command line.
 - b. From the command line, navigate to the IBM Installation Manager tools directory.

Default path of the tools directory for administrator

Windows systems: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

Linux, UNIX, IBM i, and z/OS[®] systems: opt/IBM/InstallationManager/eclipse/tools

Default path of the tools directory for non-administrator

Windows systems: C:\Users\username\IBM\Installation Manager\eclipse\tools

Linux, UNIX, IBM i, and z/OS systems: /home/username/IBM/InstallationManager/ eclipse/tools

For more information on the default paths for IBM Installation Manager installation directory, see Install as an administrator, nonadministrator, or group (http://pic.dhe.ibm.com/infocenter/install/v1r5/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_admin_nonadmin.html).

c. Run the following command:

On Windows systems

imcl.exe -c

```
For \ example, \ \texttt{C:Program Files (x86)} ibm \ \texttt{Installation Manager} eclipse \ \texttt{tools>imcl.exe -c}
```

On operating systems other than Windows, the tools directory is in the following location: ./imcl -c

2. Uninstall the packages. The command line displays the following options:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

```
L. View Logs
S. View Installation History
V. View Installed Packages
  -------
P. Preferences
  ------
E. Export Data for Problem Analysis
A. About IBM Installation Manager
  ------
X. Exit Installation Manager
```

a. On the command line, type 5 to uninstall the packages. The command line displays the following options:

```
====> IBM Installation Manager> Uninstall
```

Select one package group to uninstall from:

- 1. [X] Core services in Jazz for Service Management
- 2. [] IBM WebSphere Application Server V8.5.0.1

C. Cancel

b. Type **1**. The command line displays other options.

====> IBM Installation Manager> Uninstall

```
Select one package group to uninstall from:
     1. [X] Core services in Jazz for Service Management
    2. [] IBM WebSphere Application Server V8.5.0.1
Details of package group Core services in Jazz for Service Management:
```

```
Package Group Name : Core services in Jazz for Service Management
Shared Resources Directory : C:\Program Files (x86)\IBM\IMShared
Installation Directory : C:\Program Files\ibm\JazzSM
InstallationE.t.Eclipse IDE: L:\riog.Translations: English: 64-bit
                          : C:\Program Files\ibm\JazzSM
```

N. Next, C. Cancel

c. Type **N** to continue the uninstallation. The command line displays the following options: ====> IBM Installation Manager> Uninstall> Packages

Package group: Core services in Jazz for Service Management

Installed Packages:

- 1. [] IBM Dashboard Application Services Hub 3.1.0.1
- 2. [] Registry Services 1.1
- 3. [] Administration Services 1.1
 4. [] Security Services 1.1
- 5. [] Administration Services UI 1.1
- 6. [] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2

Other Options:

A. Select All Packages

B. Back, C. Cancel

d. Type 6 to uninstall the IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2. The command line displays the following options:

Preparing and resolving the selected packages...

====> IBM Installation Manager> Uninstall> Packages

Package group: Core services in Jazz for Service Management

Installed Packages: 1. [] IBM Dashboard Application Services Hub 3.1.0.1
```
2. [] Registry Services 1.1
3. [] Administration Services 1.1
4. [] Security Services 1.1
5. [] Administration Services UI 1.1
6. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
Other Options:
    U. Unselect All Packages
    A. Select All Packages
```

B. Back, N. Next, C. Cancel

- e. Type N to continue the uninstallation.
- f. Type the user name for WebSphere Application Server. For example, tipadmin. Password for WebSphere Application Server:
- g. Type the password for WebSphere Application Server.

---- Configuration validation:

```
Validating WebSphere Application Server credentials...
Provided credentials are valid.
B. Back, N. Next, C. Cancel
```

- h. Type **N** to continue the uninstallation.
- Validate Dashboard Application Services Hub credentials. The command line displays the following options:

```
====> IBM Installation Manager> Uninstall> Packages> Custom panels ---- Common configuration:
```

a. Type the user name for the WebSphere Application Server. For example, tipadmin.

Password for WebSphere Application Server:

b. Type the password for the WebSphere Application Server.

---- Configuration validation:

```
Validating WebSphere Application Server credentials...
Provided credentials are valid.
B. Back, N. Next, C. Cancel
```

```
c. Type N to continue.
```

4. Verify the summary of the uninstallation. The command line displays the summary of the packages that you are going to uninstall. The command line displays the following options:

```
====> IBM Installation Manager> Uninstall> Packages> Custom panels> Summary
```

Package group: Core services in Jazz for Service Management

====> IBM Installation Manager> Uninstall> Packages> Custom panels> Summary> Completion

```
The uninstall completed successfully.
Options:
F. Finish
b. Type F to finish the uninstallation.
```

PowerVM: Using the IBM Installation Manager in silent mode to uninstall the Capacity Planner

There are specific uninstallation steps to uninstall the Capacity Planner for PowerVM by using the silent mode.

Procedure

- 1. Download and extract the dashboard image. You can download this image to any directory on your local system.
- 2. Go to the silentInstall directory.
- 3. In the directory, modify the following silent response file: uninstall PowerVM CP sample.xml
- 4. Run the following commands:

For Windows systems: *IIM location\eclipseIBMIMc -s -input path of uninstall_PowerVM_CP.xml* -nosplash -acceptLicense -log *logfile* -showProgress

For Linux systems: *IIM location\eclipse.*/IBMIMc -s -input *path of uninstall_PowerVM_CP.xml* -nosplash -acceptLicense -log *logfile* -showProgress

For example, IIM location\eclipseIBMIMc -s -input C:\PowerVM7.2.0.2\silentInstall \uninstall_PowerVM_CP.xml -nosplash -acceptLicense -log C:\PowerVM7.2.0.2\uninstall_VMD.xml -showProgress

Note: The log file must be an XML file.

When the silent uninstallation is successful, a status is displayed as zero. Unsuccessful operation returns a non-zero number. When you run the Installation Manager installer, the Installation Manager reads the response file and (optionally) generates a log file in the specified directory. If you specified a log file and a directory, the log file is empty when the operation is successful.

For example,

```
<?xml version="1.0" encoding="UTF-8"?>
<result>
</result>
```

The log file is available at the following locations:

- On Windows systems:
 - Administrator : C:\ProgramData\IBM\Installation Manager\logs
 - Non-administrator: C:\Documents and Settings\username\Application Data\IBM\Installation Manager\logs
- On Linux systems: /var/ibm/InstallationManager/logs

PowerVM: Types of report installation

You can use either the Launchpad or the Reports IBM Installation Manager to install the Capacity Planner Reports for PowerVM .

Select the type of installation that suits your environment:

• Launchpad

If you have a browser and a graphics environment, select the launchpad to complete your reports installation. See Using the launchpad to install PowerVM components.

• Reports Installation Manager GUI

If you do not have a browser but you have a graphics environment, run the installer in GUI mode by using the following syntax: setup_platform.exe/.bin. For more information about using the Reports Installation Manager GUI to install reports, see "PowerVM: Installing Capacity Planner Reports ."

PowerVM: Installing Capacity Planner Reports

You can use the Launchpad or the Reports installer to install Capacity Planner Reports for PowerVM.

Procedure

- 1. Select the type of installation. For more information on selecting the type of installation, see "PowerVM: Types of report installation" on page 26.
- 2. You might need to point to Java[™] 1.5+ through your system PATH. Make sure that your system PATH contains a valid path to a Java virtual machine, for example: # PATH=\$PATH:/ibmjre50/ibm-java-i386-50/jre/bin
- **3**. From the directory from which you extracted the reports package, run the file in the following table depending on your operating system and version of Tivoli Common Reporting that is installed.

Operating system	File
AIX	setup_aix.bin
HP-UX	setup_hpux.bin
Linux	setup_linux.bin
Solaris	setup_solaris.bin
Windows	setup_windows.exe

Table 1. Setup files for Tivoli Common Reporting 2.1.1

Table 2. Setup	files for	Tivoli	Common	Reporting	3.1	or later
----------------	-----------	--------	--------	-----------	-----	----------

Operating system	File
AIX	setup_aix.bin
Linux	setup_linux.bin
Windows	setup_windows.exe

- 4. Select a language, and then click OK.
- 5. Accept the license agreement.
- 6. Select the location where the Tivoli Common Reporting server is installed (not the location where the reports are to be installed).
 - For Tivoli Common Reporting V2.1, the default path is C:\IBM\tivoli\tipv2Components\ TCRComponent or /opt/IBM/tivoli/tipv2Components/TCRComponent. The path must end with the /TCRComponent folder.
 - For Tivoli Common Reporting V3.1 or later, the default path is C:\Program Files\IBM\JazzSM\ reporting or /opt/IBM/JazzSM/reporting. The path must end with the /reporting folder.

Note: If Tivoli Common Reporting installation is distributed, reports must be installed on the dispatcher site only.

- 7. On the Welcome page, click Next
- 8. On the Choose the Installation Folder page, select the location where the Tivoli Common Reporting server is installed, and then click **Next**.
- 9. On the Choose the reports for the installation page, select the type of report that you want to install, and then click **Next**:

• IBM Infrastructure Management Capacity Planner Reports for PowerVM

CReport Installer	Choose the reports for the installation
	Choose the reports for the installation Cognos reports IBM Infrastructure Management Capacity Planner Reports for VM IBM Infrastructure Management Capacity Planner Reports for Pov IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports v IBM Tivoli Monitoring for Virtual Environments VMware Reports v7
	Description
InstallAnywhere Cancel <u>H</u> elp	ty Planner dashboard component to be instal

Figure 10. Choose the reports for the installation page

- **10**. Configure Cognos data sources to connect to Tivoli Data Warehouse by using one of the following choices:
 - If you have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting (from a previous installation of reports), go to the next step. To test whether you have Tivoli Data Warehouse defined, go to TCR > Launch Administration > Configuration > Data Source Connections and see whether there is an entry called TDW.
 - If you do not have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting, you must configure the data source in Tivoli Common Reporting. Configure the data source through the administration panel as described in Configuring database connection (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ttcr_config_db.html). You must enter the database alias name or the ODBC name for the database name input field.
 - a. On the **Cognos Engine Configuration** page, enter the user name and password for Tivoli Common Reporting, and then click **Next**.

🛎 Report Installer	
	Cognos Engine Configuration
	Enter the Tivoli Common Reporting user name jsmadmin Enter the Tivoli Common Reporting user password ******
InstallAnywhere Cancel <u>H</u> elp	<u>P</u> revious <u>N</u> ext

Figure 11. Cognos Engine Configuration page

- 11. Define common dimensions on the Tivoli Data Warehouse by using one of the following choices:
 - If you have common dimensions (Time Dimension, Weekday Lookup, Month Lookup, and Computer System under IBM_TRAM schema) in your Tivoli Data Warehouse from a previous installation and you want to modify those dimensions to define a different time granularity, you can run the scripts manually as described in *Creating shared dimension tables and populating the time dimensions table* in the *IBM Tivoli Monitoring Administrator's Guide* V6.2.2 Fix Pack 2.
 - If you do not have common dimensions in your Tivoli Data Warehouse, in the next panel, enter the JDBC credentials. Provide the database admin (db2admin, system, and so on) user name and password in the Configure data script window for JDBC User Credentials. You use the JDBC connection to run the Common Dimensions scripts on Tivoli Data Warehouse. Admin privileges are required in this step to create the IBM_TRAM schema and required tables. If you are using an Oracle database and you do not have the USERS and TEMP tablespaces in your database, you must create them in your Tivoli Data Warehouse before you can run these scripts.
 - a. If you have selected any cognos reports, enter the following details, and then click Next:
 - In the Enter the database user name field, type the DB2 administrator user name.
 - In the **Enter the user password** field, type the password.
 - In the **Choose the database type** list, select DB2.
 - In the Enter the database name field, type TADFDCDB.

🛎 Report Installer		- 🗆 X
Cognos Dat	ta Source TADFDCP_REPORTING_DS Configur	ation
	Enter the database user name administrator	
	Enter the user password	_
	Choose the database type	
	DB2	-
	TADFDCDB	
	Skip this panel (not recommended)	
InstallAnywhere Cancel Help	Previous	xt

Figure 12. Cognos Data Source TADFDCP_REPORTING_DS Configuration page

12. On the Report Installer Summary page, verify the reports to be installed, click **Install**, and wait for the installer to finish.

CReport Installer		
	The following report sets will be installed: Cognos reports IBM Infrastructure Management Capacity Planner Reports for PowerVM v7.2.0.2	
InstallAnywhere Cancel <u>H</u> elp	<u>P</u> revious <u>I</u> r	istall

Figure 13. Report Installer Summary page

Note: The **Installation results** page shows the status of all installation actions for every item or report. One log file and one trace file are included. Both files are in the user home directory, with the following names: Report_Installer_Install_mm_dd_yyyy_hh_mm_ss.log (Log) and Report_Installer_For_TCR_Output.txt (Trace)



Figure 14. Installation results page

13. Click Done.

Installing reports by using the command line

You can install the reports by using the command line.

Procedure

- 1. Run the **setup_***platform*.**exe** -**i** or **setup_***platform*.**bin** -**i** command.
- 2. Choose your installation language.
- 3. Enter the location of the TCRComponent directory.
- 4. Choose the type of reports to be installed.
- 5. Enter your Tivoli Common Reporting user name and password.
- 6. Configure your datasource and data scripts. Some report packages might not have data scripts.
- 7. An installation summary is provided, then press Enter to begin installation.

Installing reports by using the silent mode

You can install the reports by using the silent mode.

Procedure

- 1. Create the silent installer response file, and name the file as silent_installer.properties.
- 2. Run the **setup_***platform*.**exe**/.**bin** -**i silent** -**f** *path_to_response_file* command.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

PowerVM: Uninstalling Capacity Planning Reports

Uninstallation of reports is not supported by the reports installer. You can however manually delete the reports package on the Dashboard Application Services Hub for Tivoli Common Reporting version 2.1 or later, or on the Dashboard Application Services Hub for Tivoli Common Reporting version 3.1 or later.

Procedure

- 1. Log in to the Tivoli Common Reporting interface and go to Common Reporting.
- 2. In the **Public Folders** tab of the Work with reports window, select the reports package that you want to delete.
- 3. Click the **Delete** icon on the toolbar.



Figure 15. Delete icon

Results

The selected reports package is uninstalled.

34 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 3. VMware: Installing components

To install the VMware capacity planning components, and the Dashboard for VMware, you must install the prerequisites, select a type of installation depending on your environment, complete the selected procedure to install the Capacity Planner, and the Dashboard for VMware. Then, you must install the Capacity Planner Reports.

VMware: Installing prerequisites

Before you complete your installation, you must install the required prerequisites.

Procedure

- 1. Install Jazz for Service Management 1.1 or later. You must select the following Jazz for Service Management components that are required for the Capacity Planner for VMware and the Dashboard for VMware:
 - WebSphere Application Server V8.5 or later
 - Jazz for Service Management Extension to WebSphere Application Server V8.5 or later
 - Dashboard Application Services Hub V3.1 or later

For information about database recommendations, see "Capacity Planner database recommendations" in the *IBM Tivoli Monitoring for Virtual Environments Dashboard, Reporting, and Capacity Planning Troubleshooting* documentation.

 Create a new database for both the local and remote installations of the Capacity Planner database by using the following command: db2 create database TADFDCDB USING CODESET UTF-8 TERRITORY US COLLATE USING SYSTEM PAGESIZE 16384.

Important: This database is shared between the Capacity Planner for PowerVM and the Capacity Planner for VMware. If you install both of these capacity planners, you must create the database once only. You can create the capacity planner database as the db2 instance owner on one of the following systems:

- On Windows systems, the instance owner is **db2admin**.
- On operating systems other than Windows, the instance owner is db2inst1.

The Post Install Federation step might fail if the Capacity Planner database is created as a user with non compliant authorization levels.

Note: The database instance owner must have the following DB2 permissions for the Capacity Planner database:

- Connect to database
- Create tables
- Create schema
- Create packages
- **3.** To install and configure the base IBM Tivoli Monitoring monitoring server, portal server, and portal client components, see the: IBM Tivoli Monitoring Installation and Setup Guide (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/itm_install.htm).

Note: When you configure the portal server, you must enable the dashboard data provider.

4. Install and configure the monitoring agents so that the monitoring data displays in the monitoring dashboards or the Capacity Planner application. If you want to monitor the virtual environment remotely, for the Dashboard for VMware, you must install the VMware VI, NetApp Storage, and

Tivoli Storage Productivity agents. For the Capacity Planner for VMware, you must install the VMware VI agent. For more information about installing the application support in the monitoring servers, portal server, and desktop portal client, see the IBM Tivoli Monitoring Installation and Setup Guide.

VMware: Types of installation

You can use either the Launchpad or the IBM Installation Manager to install the VMware components:

The VMware components include:

- IBM Infrastructure Management Capacity Planner for VMware
- IBM Infrastructure Management Capacity Planner Reports for VMware
- IBM Infrastructure Management Dashboard for VMware

Select the type of installation that suits your environment:

• Launchpad

If you have a browser and a graphics environment, select the launchpad to complete your installation. See Using the launchpad to install VMware components.

• Installation Manager GUI

If you do not have a browser but you have a graphics environment, use the Installation Manager GUI to complete your installation. See Using the IBM Installation Manager to install VMware components.

• Installation Manager Silent

If you do not have a browser or a graphics environment, or you want to complete a silent installation, use the Installation Manager silent type of installation to install the VMware components. See "VMware: Using the IBM Installation Manager silent mode to install the Capacity Planner and the Dashboard" on page 39.

• Installation Manager Console

If you want to use IBM Installation Manager in console mode to install packages from a command line interface, use the Installation Manager console type of installation. See "VMware: Using the IBM Installation Manager Console to install the Capacity Planner and the Dashboard " on page 40.

Note: Both the Jazz for Service Management components and the IBM Tivoli Monitoring for Virtual Environments components for a dashboard environment must be installed in the same user mode; install both as a root user or install both as a non-root user. Information cannot be exchanged between these components when there is a combination of root and non-root installations.

VMware: Using the Launchpad to install components

You can use the launchpad to start the IBM Installation Manager to install the VMware components.

Before you begin

Ensure that you have a supporting browser that is installed in your target environment, because the Launchpad requires a browser to run. The following browsers are supported:

- Firefox 2.0 or later, Firefox 3.0 or later and Firefox 3.5 or later
- Internet Explorer 6.0 or later
- Mozilla 1.7 or later
- SeaMonkey 1.1.4 or later and SeaMonkey 2.0 or later

Attention: If your target system does not have a browser that is installed, use Installation Manager in GUI mode.

Ensure that you do not have multiple instances of the Launchpad open, because this might cause a conflict to occur.

When you first download the Launchpad image and extract to a directory on the local system, ensure that the path to the extracted Launchpad package does not contain any spaces; otherwise the Launchpad will not run.

Procedure

- 1. Download and extract the dashboard image file for the launchpad that contains the IBM Installation Manager repository, and Reports Installer. You can download this image to any directory on your local system.
- 2. Start the Launchpad.

Important: The best practice is to have only one instance of the Launchpad open at a time.

Option	Description
On Windows systems	Browse to the local \Launchpad directory, and run launchpad64 for 64-bit Windows computers.
On Linux or AIX systems	In a command window, open the local \Launchpad directory and run ./launchpad.sh .

These commands start the IBM Installation Manager GUI.

- 3. Choose one of the following options:
 - To install the capacity planners and reports as an administrative user, click **Full Installation as administrative user**.
 - To install the capacity planners and the reports as a non-administrative user, click **Full Installation** as non-administrative user
 - To install the capacity planners and the dashboards as an administrative user, click **Install Capacity Planners and Dashboards as administrative user**. When you select this option, the Installation Manager GUI starts.
 - To install the capacity planners and the dashboards as a non-administrative user, click **Install Capacity Planners and Dashboards as non-administrative user**. When you select this option, the Installation Manager GUI starts.
 - To install Performance and Capacity Management reports, click **Install Reports**. When you select this option, the Reports Installer starts. For more information about installing reports, see "VMware: Installing Capacity Planner Reports" on page 58.

ITM for Virtual Er	vironments	Select a language: English	▼ OK
	Installation Options This product contains installation comp and reports. You can use this launchpa perform a full installation or install cor Full Installation as administrative Full Installation as non-administr You will be able to choose which capae Install Capacity Planners and Das Install Capacity Planners and Das Install Reports	oonents for capacity planners, dashboards ad to install them. Click the links below to nponents individually. • user ative user city planners and dashboards to be installed. hboards as administrative user hboards as non-administrative user	
	View Release Information Readme Installation Guide		
IBM.			

Figure 16. ITM for Virtual Environments page

What to do next

To complete the installation, see the step to select a component to install in the "VMware: Using the IBM Installation Manager GUI to install components" procedure. Also see the "VMware: Installing Capacity Planner Reports" on page 58 procedure.

VMware: Using the IBM Installation Manager GUI to install components

You start the Installation Manager to install the VMware components.

Procedure

1. To start the Installation Manager, do one of the following steps:

- On Windows systems, select Start > IBM Installation Manager > IBM Installation Manager.
- On Linux or AIX systems, in a command window, open the /opt/IBM/InstallationManager/ eclipse directory and enter ./IBMIM or ./groupinst for group mode.
- Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where the IBM Installation Manager is installed. The default path for IBM Installation Manager on Windows systems is C:\Program Files\IBM\Installation Manager or C:\Program Files(86)\IBM\ Installation Manager and on UNIX systems, the path is /opt/IBM/Installation Manager.
- 2. Click **File** > **Preferences**. Then click **Add**.
- 3. Enter the path or click **Browse**.
- 4. In the **Repository** field, enter the file path by clicking **Browse** to point to the directory where you extracted the installation image and select the diskTag.inf file, and then click **OK**.

5. Select a component to install. For more information about the components to install, see "VMware: Installing the Capacity Planner and the Dashboard " on page 45. Then, click **Next**.

Results

	Status	Vendor	License Key Type
↓ ↓ IBM Infrastructure Management Capacity Planner for PowerVM 		IBM	
Version 7.2.0.2	Will be installed	IBM	
	Will be installed	IBM	
Show <u>all</u> versions		Check	for Other Versions, Fixes, and Extension
tails			
M Infrastructure Management Dashboard for VMware 7.2.0.2			
Repository: D:			
Repository: D:			

Figure 17. Install Packages page

Note: The IBM Installation Manager checks for license, languages, and the location of installation. You are prompted for configuration parameters. For more information about configuration parameters, see "Configuring a connection for the Capacity Planner for VMware and the Capacity Planner for PowerVM" on page 121.

VMware: Using the IBM Installation Manager silent mode to install the Capacity Planner and the Dashboard

There are specific installation steps to install the Capacity Planner for VMware and the Dashboard for VMware by using the silent mode.

Before you begin

- Ensure that the Jazz for Service Management 1.1 or later is installed with the following three required components:
 - IBM WebSphere Application Server V8.5 or later
 - Jazz for Service Management Extension to WebSphere Application Server V8.5 or later

- Dashboard Application Services Hub V3.1 or later
- If you want to install Capacity Planner for VMware and Dashboard for VMware as a non-root user then there must be Jazz for Service Management that is installed as a non-root user.
- Ensure that no other Installation Manager process is running before you install the offering silently.

Procedure

- 1. Download and extract the dashboard image. You can download this image to any directory on your local system.
- 2. Go to the silentInstall directory.
- In the directory, modify the following silent response file install_VMware_CP_sample.xml
- 4. Run the following silent installation scripts:

Installation as Administrator

On Windows systems: silentInstall_VMware_CP_win.bat

On UNIX systems: silentInstall_VMware_CP_unix.sh

Installation as Non-Administrator

For Windows systems: silentInstall_VMware_CP_win.bat user

For UNIX systems: silentInstall_VMware_CP_unix.sh user

Note: If you want to run script for non-root user, you must pass one variable user to the script. The user does not mean actual non-root user name. It must be only user.

VMware: Using the IBM Installation Manager Console to install the Capacity Planner and the Dashboard

There are specific installation steps to install the Capacity Planner for VMware and the Dashboard for VMware by using the Console mode.

Procedure

- 1. Start the console mode.
 - a. On a command line, navigate to the IBM Installation Manager tools directory.

Default path of the tools directory for administrator

On Windows systems: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

On operating systems other than Windows: opt/IBM/InstallationManager/eclipse/tools

Default path of the tools directory for non-administrator On Windows systems: C:\Users\username\IBM\Installation Manager\eclipse\tools

On operating systems other than Windows: /home/username/IBM/InstallationManager/ eclipse/tools

For more information on the default paths for IBM Installation Manager installation directory, see Install as an administrator, nonadministrator, or group (http://pic.dhe.ibm.com/infocenter/install/v1r5/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_admin_nonadmin.html).

b. Run the following command:

On Windows systems

imcl.exe -c

For example, C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools>imcl.exe -c On operating systems other than Windows:

./imcl -c

For example, opt/IBM/InstallationManager/eclipse/tools>./imcl -c

2. Add the repository. The command line displays the following options:

====> IBM Installation Manager

Select:

- Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- P. Preferences
- -----
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- X. Exit Installation Manager
- a. At the command line, type P to set the preferences. The command line displays the preferences.

====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- 2. Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
- 6. Passport Advantage
- 7. Updates
- R. Return to Main Menu
- b. Type **1** to select the repository. The command line displays other options.

====> IBM Installation Manager> Preferences> Repositories

Repositories:

Other Options:

- D. Add Repository
- S. [X] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- c. Type **D** to add the repository.
- d. Type the repository location where you extracted the installation image, and select the diskTag.inf file. For example, C:\v7.2.0.2\diskTag.inf.

====> IBM Installation Manager> Preferences> Repositories> Add repository

```
Enter a new repository location.
To skip, press Enter:
----> C:\v7.2.0.2\diskTag.inf
Checking repositories...
```

====> IBM Installation Manager> Preferences> Repositories

Repositories: 1. [X] C:\v7.2.0.2\diskTag.inf Other Options:

D. Add Repository

- S. [X] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- e. Type S to clear the option "Search service repositories during installation and updates".

====> IBM Installation Manager> Preferences> Repositories Repositories:

1. [X] C:\v7.2.0.2\diskTag.inf

- Other Options:
- D. Add Repository
- S. [] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu

f. Type **A** to apply changes and return to the Preferences menu.

====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
- 6. Passport Advantage
- 7. Updates
- R. Return to Main Menu
- g. Type **R** to return to the main menu.
- **3**. Select the package to install. The command line displays the following option:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- P. Preferences
 - _____
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- X. Exit Installation Manager
- a. Type **1** to install the packages.

```
Checking repositories...
Loading repositories...
Checking availability of packages...
Loading fixes...
```

```
====> IBM Installation Manager> Install
```

Select packages to install: 1. [] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2

- 2. [] IBM Infrastructure Management Dashboard for VMware 7.2.0.2
- O. Check for Other Versions, Fixes, and Extensions.
- C. Click Cancel.
- b. Type 1 to install the IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2.
 ====> IBM Installation Manager> Install> Select

IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2

Options:

- 1. Choose version 7.2.0.2 for installation.
- 2. Show all available versions of the package.

C. Cancel

c. Type **1** to select version 7.2.0.2.

====> IBM Installation Manager> Install

- Select packages to install:
 - [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
 [] IBM Infrastructure Management Dashboard for VMware 7.2.0.2
 - 0. Check for Other Versions, Fixes, and Extensions

N. Next, C. Cancel

d. Type 2 to install the IBM Infrastructure Management Dashboard for VMware. ====> IBM Installation Manager> Install> Select

IBM Infrastructure Management Dashboard for VMware 7.2.0.2

Options:

- 1. Choose version 7.2.0.2 for the installation.
- 2. Show all available versions of the package.
- C. Cancel
- e. Type **1** to select version 7.2.0.2.

Preparing and resolving the selected packages...

====> IBM Installation Manager> Install

Select packages to install:

```
1. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
```

2. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.2

O. Check for Other Versions, Fixes, and Extensions

N. Next, C. Cancel

f. Type N. The command line displays the details of the compatible package groups. Finding compatible package groups...

====> IBM Installation Manager> Install> Location
Existing package groups:
 1. [X] Core services in Jazz for Service Management
Incompatible existing package groups:
 2. IBM WebSphere Application Server V8.5

Selected group id: "Core services in Jazz for Service Management"

Selected location: "C:\Program Files\IBM\JazzSM" Selected architecture: 64-bit

B. Back, N. Next, C. Cancel

g. Type N. The command line displays the features of the packages.
====> IBM Installation Manager> Install> Location> Features

```
IBM Infrastructure Management Capacity Planner for VMware
1. [X] Installation
2. [X] Configuration
IBM Infrastructure Management Dashboard for VMware
3. [X] Installation
4. [X] Configuration
```

B. Back, N. Next, C. Cancel

h. Type N to continue the installation.

4. Validate the Dashboard Application Services Hub credentials. The command line displays the following options:

```
====> IBM Installation Manager> Install> Location> Features> Custom panels
---- Common configuration:
User name for WebSphere Application Server:
----> [tipadmin] tipadmin
```

- a. Type the user name for WebSphere Application Server. For example, **tipadmin**. Password for WebSphere Application Server:
- b. Type the password for WebSphere Application Server.

---- Configuration validation:

Validating WebSphere Application Server credentials... Provided credentials are valid.

5. Validate the Capacity Planner database credentials.

---- Configuration for IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 Capacity Planner Database Server Host name:

- a. Type the Capacity Planner Database Server host name. For example, **localhost**. Capacity Planner Database Server Port:
- b. Type the Capacity Planner Database Server Port. For example, 50000.
 Capacity Planner Database Admin name:
- C. Type the Capacity Planner Database Admin name. For example, db2admin.
 Capacity Planner Database Admin Password:
- d. Type the Capacity Planner Database Admin password.Capacity Planner Database Name:
- e. Type the Capacity Planner Database Name. For example, **TADFDCDB**. The system automatically validates the data that you entered. If the validation fails, the command line displays the following options:

Validation in progress. Please wait...

Validation failed.

[jcc][t4][2013][11249][4.7.85] Connection authorization failure occurred. Reaso n: User ID or Password invalid. ERRORCODE=-4214, SQLSTATE=28000

Validation of your parameters failed, do you want to enter parameters again?

Y. Yes N. No

----> [Y]

f. Enter the values for the parameters. The command line displays the following options: Validation in progress. Please wait...

The parameters were verified successfully.

B. Back, N. Next, C. Cancel

6. Check the summary of the installation. The command line displays the summary of the packages that you are going to install.

```
====> IBM Installation Manager> Install> Location> Features> Custom panels>
  Summary
Target Location:
 Package Group Name : Core services in Jazz for Service Management
Installation Directory : C:\Program Files\ibm\JazzSM
  Shared Resources Directory : C:\Program Files (x86)\IBM\IMShared
Packages to be installed:
        IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
        IBM Infrastructure Management Dashboard for VMware 7.2.0.2
Options:
    G. Generate an Installation Response File.
     B. Back,
                  I. Install,
                                   C. Cancel
a. Type I to install the packages.
   ----> [I]
                    25%
                                                         75%
                                                                           100\%
                                       50%
   -----|-----|-----|
   ====> IBM Installation Manager> Install> Location> Features> Custom panels> Summary> Completion
   The install completed successfully.
   Options:
        F. Finish
```

b. Type **F** to finish the installation.

VMware: Installing the Capacity Planner and the Dashboard

Installing the Capacity Planner for VMware and the Dashboard for VMware requires specific installation steps.

Procedure

- 1. Select the type of installation. For more information about selecting the type of installation, see "VMware: Types of installation" on page 36.
- 2. Start the launchpad or Installation Manager GUI. For more information, see "VMware: Using the Launchpad to install components" on page 36.
- 3. Select the following check boxes:
 - IBM Infrastructure Management Capacity Planner for VMware
 - IBM Infrastructure Management Dashboard for VMware

Note: When you select a component, the Installation Manager automatically selects the version of the respective component that you selected.

4. Click Next.

Installation Packages	Status	Vendor	License Key Type
IBM Infrastructure Management Capacity Planner for PowerVM Wersion 7.2.0.2 With Infrastructure Management Capacity Planner for VMware Wersion 7.2.0.2 With Infrastructure Management Dashboard for VMware Wersion 7.2.0.2	Will be installed Will be installed	IBM IBM IBM	
Show <u>a</u> ll versions etails BM Infrastructure Management Dashboard for VMware 7.2.0.2 3M Infrastructure Management Dashboard for VMware Installer <u>More info</u> Repository: D:		Check	for Other Versions, Fixes, and Extension

Figure 18. Install Packages page

5. The Installation Manager displays the existing package groups. Verify the details, and then click **Next**.

	Server 2008 energing system, you must install the application on diversities a directory	that is not uich aligned
The Program Files directory on Windows Server 2008 is virtualized.	s Server 2008 operating system, you must install the application package into a directory	that is not virtualized.
To run an application with administrative privileges after it is installed, r	right-click the package, and click "Run as administrator".	
Install Location Features Summary		
Use the existing package group		
Create a new package group		
Bardinan Court Mana	Televille Deve Principal	A
Package Group Name	C:\Program Files\ibm\]azzSM	64-bit
	C:\Program Files (x86)\IBM\WebSphere\AppServer	
ackage Group Name: Core services in Jazz for Service Management		
Package Group Name: Core services in Jazz for Service Management		Browce
Package Group Name: Core services in Jazz for Service Management nstallation Directory: C:\Program Files\bm\JazzSM		Browse
Vackage Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\bm\JazzSM Architecture Selection: 32-bit 64-bit		Browse
Package Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\bm\JazzSM Architecture Selection: O 32-bit O 64-bit		Browse
Vackage Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\bm\JazzSM Architecture Selection: Image: Selection Selection Image: Selection Selection: Image: Selection Selection	Dick Space Information	Browse
Vackage Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\bm\JazzSM Architecture Selection: Image: Selection Selection Image: Starsed Resources Directory: C:\Program Files (v86\\TBM\TMShared	Disk Space Information	Browse
Package Group Name: Core services in Jazz for Service Management Installation Directory: C: \Program Files \Jbm \JazzSM Architecture Selection: Image: Selection Selection: Image: Selection: Image: Selection Select	Disk Space Information Volume Available Space	Browse
Package Group Name: Core services in Jazz for Service Management Installation Directory: C: \Program Files \Jbm \JazzSM Architecture Selection: Image: Selection Selection: Image: Selection: Image: Selection Selection: Image: Selection Sel	Disk Space Information Volume Available Space C: 7,46 GB	Browse
Package Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\Jbm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\JBM\JMShared • Eclipse IDE: C:\Program Files\Jbm\JazzSM Installed Packages • Administration Services 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browse
Package Group Name: Core services in Jazz for Service Management installation Directory: C:\Program Files\Jbm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\JBM\JMShared • Eclipse IDE: C:\Program Files\Jbm\JazzSM Installed Packages • Administration Services 1.1 • JEM Dackbared Amplicing Services Hub 2.1	Disk Space Information Volume Available Space C: 7.46 GB	Browsenn
Package Group Name: Core services in Jazz for Service Management installation Directory: C:\Program Files\Jbm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared • Eclipse IDE: C:\Program Files\Jbm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Descint, Comment 1.	Disk Space Information Volume Available Space C: 7.46 GB	Browsen
Package Group Name: Core services in Jazz for Service Management installation Directory: C: Program Files \bm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C: \Program Files (x86)\IBM\IMShared • Eclipse IDE: C: \Program Files \bm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browse
Ackage Group Name: Core services in Jazz for Service Management Installation Directory: C: \Program Files \Jbm \JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C: \Program Files (x86) \JBM \JMShared • Edipse IDE: C: \Program Files \Jbm \JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1 • Security Services 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browse
Package Group Name: Core services in Jazz for Service Management Installation Directory: C:Program Files\bm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C:Program Files (x86)\LBM\LMShared • Edipse IDE: C:Program Files\bm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1 • Security Services 1.1 • Administration Services UI 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browsen
Package Group Name: Core services in Jazz for Service Management Installation Directory: C:\Program Files\Jbm\JazzSM Architecture Selection: ● 32-bit ● 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\JBM\JMShared • Eclipse IDE: C:\Program Files\Jbm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1 • Security Services 1.1 • Administration Services UI 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browse
Package Group Name: Core services in Jazz for Service Management installation Directory: C:\Program Files\bm\JazzSM Architecture Selection: O 32-bit O 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\BM\JMShared • Eclipse IDE: C:\Program Files\bm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1 • Security Services 1.1 • Administration Services UI 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browsen
Package Group Name: Core services in Jazz for Service Management installation Directory: C:\Program Files\Jbm\JazzSM Architecture Selection: ① 32-bit ② 64-bit Details • Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared • Eclipse IDE: C:\Program Files\Jbm\JazzSM Installed Packages • Administration Services 1.1 • IBM Dashboard Application Services Hub 3.1 • Registry Services 1.1 • Security Services 1.1 • Administration Services UI 1.1	Disk Space Information Volume Available Space C: 7.46 GB	Browsen

Figure 19. Install Packages page

- 6. Enter the WebSphere Application Server user name and password for Dashboard Application Services Hub, and then click **Validate**.
- 7. After the user name and password are validated, click Next.

Figure 20. Install Packages page: Common Configurations

Install Deskages						_
Click the Validate button to continue.						
Testal Learning Ex	the second					Feature
Common Configurations	Common Cor WebSphere Ap User name: Password: Validate	figurations slication Server Credentials smadmin 				
			[[
			< <u>B</u> ack	Next >	Install	Cancel

Figure 21. Install Packages page: Common Configurations

- 8. Enter the Capacity Planner Database Schema Creation details, and then click Validate.
- 9. After the details are validated, click **Next**



Figure 22. Install Packages page: VMware Capacity Planner Configuration

10. Review the information in the Install Packages summary window, and then click Install.

Installation Manager			
nstall Packages			
Review the summary information.			4
cever ac sennery mondant			
Install Location Features Summary			
Target Location			
Package Group Name: Core services in Jazz for Service Management			
Installation Directory: C:\Program Files\ibm\JazzSM			
Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared			
Packages			
Packages			
🖃 🕼 IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2			
🕼 Installation			
Configuration			
IBM Infrastructure Management Dashboard for VMware 7.2.0.2			
Installation			
Configuration			
énvironment	Disk Space Information		
English	Total Available Space		
	C. 7.00 GD		
	Total Download Size: 35.00 MB		
	Total Installation Size: 77.31 MB		
Repository Information			
, repository and finding			
		< Back Nevt >	Install Can
		Line Line A	Listaii Calit

Figure 23. Install Packages page: Summary Information

11. The Results window displays the package installation status. Click **Finish** to complete the installation.



Figure 24. Install Packages page

12. To verify the installation, click **File** > **View Installed Packages**.

🖄 Installed Packages

Installed Packages and Fixes

Istalieu Packages	Vendor	License	
Real IBM WebSphere Application Server V8.5			
IBM WebSphere Application Server_8.5.0.1	IBM		
💫 IBM WebSphere SDK Java Technology Edition (Optional)_7.0.2.0	IBM		
🙀 Jazz for Service Management extension for IBM WebSphere 8.5_1.1	IBM		
Rea Core services in Jazz for Service Management			
Registry Services_1.1	IBM		
Administration Services_1.1	IBM		
Security Services_1.1	IBM		
IBM Dashboard Application Services Hub_3.1	IBM		
Administration Services UI_1.1	IBM		
IBM Infrastructure Management Capacity Planner for PowerVM_7.2.0.0	IBM		
IBM Infrastructure Management Capacity Planner for VMware_7.2.0.2	IBM		
IBM Infrastructure Management Dashboard for VMware_7.2.0.2	IBM		
taile			
etails IM WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86) \IBM\IMShared			
etails IM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86) \IBM\IMShared Installation Directory: C: \Program Files (x86) \IBM\WebSphere \AppServer Eclipse IDE: C: \Program Files (x86) \IBM\WebSphere \AppServer			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86) \IBM\IMShared Installation Directory: C: \Program Files (x86) \IBM\WebSphere \AppServer Eclipse IDE: C: \Program Files (x86) \IBM\WebSphere \AppServer Translations: English			
etails IM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails IM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\UMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86) \IBM\\IMShared Installation Directory: C: \Program Files (x86) \IBM\\WebSphere \AppServer Eclipse IDE: C: \Program Files (x86) \IBM\\WebSphere \AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86)\IBM\\IMShared Installation Directory: C: \Program Files (x86)\IBM\\WebSphere\AppServer Eclipse IDE: C: \Program Files (x86)\IBM\\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C: \Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English			
etails M WebSphere Application Server V8.5 Shared Resources Directory: C: \Program Files (x86)\IBM\\Imshared Installation Directory: C: \Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C: \Program Files (x86)\IBM\WebSphere\AppServer Translations: English			

Figure 25. Installed packages details

13. Log in to the Dashboard Application Services Hub to verify that you can start the **Planning Center for VMware** wizard.

Planning Center for VMware ×

Planning Center for VMware

You have acquired the Planning Center for VMware lock at 25 Jul 2013 03:14:51 PM IST. You need to explicitly release the lock for other users to use the Planning Center for VMware.	Release Lock
O Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommendation.	
Step 1: Snapshot config data.	
Load the latest configuration data for physical servers and virtual machines for analysis. You can change this data for what-if analysis. Advanced options: Select the data load options:	
Load data for selected Physical Servers Clean database before loading	
Load Config	
Step 2: Set analysis time period.	
Set the time period for which the measurement data in the warehouse can be analyzed corresponding to the virtual machines loaded in Step 1. The measurement data is federated from the	warehouse.
Set Time	
Step 3: Scope the infrastructure for analysis.	
Default scope includes all physical servers loaded in Step 1. Click Define Scope to go to an expert mode page where you can select the subset of physical servers that you want to be part of	fthis analysis.
Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to add new attributes or clean the data for physical servers and virtual machines, if required.	
Define Scope	
Step 4: Analyze virtual machine characteristics.	
Analyze the warehouse data within the time limit set in Step 2, to compute the VM-level sizing estimates using default settings (average daily utilization).	
Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you can trigger several custom actions to characterize virtual machines based on measurement	data.
Size VMs	
Current Environment Report	
Step 5: Generate optimization plan.	
O Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommendation.	
Generate an optimization plan based on recommended environment settings.	
Advanced options: Custom settings can be made on the Edit Recommended Environment Settings page where experts can select optimization strategies, such as applicable business and technical policies,	optimization goal and so o
Generate Plan	
Optimized Plan Report	
Additional Workload Analysis Report	
VMware Expense Reduction Report	

Figure 26. Planning Center for VMware

14. Verify that the Dashboard for VMware is installed.

What to do next

If an HTTP server is being used with Dashboard Application Services Hub for load balancing, regenerate the web server plug-in. For more information about regenerating the web server plug-in, see "Generating the plugin-cfg.xml file" in the Jazz for Service Management Configuration Guide.

Note: If you do not complete the steps, a blank page is displayed when you access the Dashboard for VMware or Capacity Planner for VMware pages.

VMware: Using the IBM Installation Manager GUI to uninstall the Capacity Planner and the Dashboard

You can uninstall the Capacity Planner for VMware and Dashboard for VMware by using the IBM Installation Manager GUI.

Procedure

- 1. Start the IBM Installation Manager, and click Uninstall.
- 2. On the Uninstall Packages page, select the following VMware components, and then click Next.
 - IBM Infrastructure Management Capacity Planner for VMware
 - IBM Infrastructure Management Dashboard for VMware
- 3. Enter the credentials for WebSphere Application Server, and then click Validate.
- 4. After the credentials are verified, click **Next**.
- 5. Review the summary of the components that you want to uninstall.
- 6. Click Uninstall.

VMware: Using the IBM Installation Manager in silent mode to uninstall the Capacity Planner and the Dashboard

There are specific uninstallation steps to uninstall the Capacity Planner for VMware and the Dashboard for VMware by using the Silent mode.

Procedure

- 1. Download and extract the dashboard image. You can download this image to any directory on your local system.
- 2. Go to the silentInstall directory.
- In the directory, modify the following silent response file: uninstall VMware CP.xml
- 4. Run the following commands:

For Windows systems: *IIM location\eclipse*>IBMIMc -s -input *path of uninstall_VMware_CP.xml* -nosplash -acceptLicense -log *logfile* -showProgress

For Linux systems: *IIM location**eclipse*>./IBMIM -s -input *path of uninstall_VMware_CP.xml* -nosplash -acceptLicense -log *logfile* -showProgress

For example, IIM location\eclipse>IBMIMc -s -input C:\VMware7.2.0.2\silentInstall \uninstall_VMware_CP.xml -nosplash -acceptLicense -log C:\VMware7.2.0.2\uninstall_VMD.xml -showProgress

Note: the log file must be an XML file.

When the silent uninstallation is successful, a status is displayed as zero. Unsuccessful operation returns a non-zero number. When you run the Installation Manager installer, the Installation Manager reads the response file and (optionally) generates a log file in the specified directory. If you specified a log file and a directory, the log file is empty when the operation is successful.

```
For example,
```

```
<?xml version="1.0" encoding="UTF-8"?>
<result>
</result>
```

The log file is available at the following locations:

- On Windows systems:
 - Administrator : C:\ProgramData\IBM\Installation Manager\logs
 - Non-administrator: C:\Documents and Settings\username\Application Data\IBM\Installation Manager\logs
- On Linux systems: /var/ibm/InstallationManager/logs

VMware: Using IBM Installation Manager Console to uninstall the Capacity Planner and the Dashboard

There are specific steps to uninstall the Capacity Planner for VMware and Dashboard for VMware.

Procedure

- 1. Start the console mode.
 - a. On a command line, navigate to the IBM Installation Manager tools directory.

```
Default path of the tools directory for administrator
```

On Windows systems: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

On operating systems other than Windows: opt/IBM/InstallationManager/eclipse/tools

```
Default path of the tools directory for non-administrator
```

On Windows systems: C:\Users\username\IBM\Installation Manager\eclipse\tools

On operating systems other than Windows: /home/username/IBM/InstallationManager/ eclipse/tools

For more information on the default paths for IBM Installation Manager installation directory, see Install as an administrator, nonadministrator, or group (http://pic.dhe.ibm.com/infocenter/ install/v1r5/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_admin_nonadmin.html).

b. Run the following command:

On Windows systems:

imcl.exe -c

```
For example, C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools>imcl.exe -c
On operating systems other than Windows:
```

./imcl -c

```
For example, opt/IBM/InstallationManager/eclipse/tools>./imcl -c
```

2. Uninstall the packages. The command line displays the following options:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages

```
P. Preferences
```

- -----
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- X. Exit Installation Manager
- a. At the command line, type **5** to uninstall the packages. The command line displays the following options:

```
====> IBM Installation Manager> Uninstall
```

```
Select one package group to uninstall from:
```

- 1. [X] Core services in Jazz for Service Management
- 2. [] IBM WebSphere Application Server V8.5
- C. Cancel
- b. Type1. The command line displays other options.

====> IBM Installation Manager> Uninstall

N. Next, C. Cancel

c. Type **N** to continue the uninstallation. The command line displays the following options: ====> IBM Installation Manager> Uninstall> Packages

Package group: Core services in Jazz for Service Management

```
Installed Packages:
```

- 1. [] IBM Dashboard Application Services Hub 3.1
- 2. [] Registry Services 1.1
- 3. [] Administration Services 1.1
- 4. [] Security Services 1.1
- 5. [] Administration Services UI 1.1
- 6. [] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
- 7. [] IBM Infrastructure Management Dashboard for VMware 7.2.0.2 $\,$

```
Other Options:
```

A. Select All Packages

B. Back, C. Cancel

d. Type 6 to uninstall IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 and type 7 to uninstall IBM Infrastructure Management Dashboard for VMware 7.2.0.2. The command line displays the following options:

Preparing and resolving the selected packages...

====> IBM Installation Manager> Uninstall> Packages

Package group: Core services in Jazz for Service Management

Installed Packages:

- 1. [] IBM Dashboard Application Services Hub 3.1
- 2. [] Registry Services 1.1
- 3. [] Administration Services 1.1
- 4. [] Security Services 1.1
- 5. [] Administration Services UI 1.1
- 6. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
- 7. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.2

Other Options:

- U. Unselect All Packages
- A. Select All Packages
- B. Back, N. Next, C. Cancel

e. Type N to continue the uninstallation.

3. Validate the Dashboard Application Services Hub credentials. The command line displays the following options:

```
====> IBM Installation Manager> Uninstall> Packages> Custom panels ---- Common configuration:
```

- a. Type the user name for the WebSphere Application Server. For example, **tipadmin**. Password for WebSphere Application Server:
- b. Type the password for the WebSphere Application Server.
 - ---- Configuration validation:

Validating WebSphere Application Server credentials... Provided credentials are valid. B. Back, N. Next, C. Cancel

- c. Type N to continue.
- 4. Verify the summary of the uninstallation. The command line displays the summary of the packages that you are going to uninstall. The command line displays the following options:

```
Package group:
      Core services in Jazz for Service Management
Packages to be uninstalled:
  1. IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
  2. IBM Infrastructure Management Dashboard for VMware 7.2.0.2
Options:
   G. Generate an Uninstall Response File
   B. Back.
               U. Uninstall,
                              C. Cancel
a. Type U to uninstall the packages. The command line displays the following options:
                                             75%
                               50%
                25%
          ====> IBM Installation Manager> Uninstall> Packages> Custom panels> Summary>
    Completion
  The uninstall completed successfully.
  Options:
      F. Finish
b. Type F to finish the uninstallation.
```

====> IBM Installation Manager> Uninstall> Packages> Custom panels> Summary

VMware: Types of report installation

You can use either the Launchpad or the Reports IBM Installation Manager to install the Capacity Planner Reports for VMware .

Select the type of installation that suits your environment:

• Launchpad

If you have a browser and a graphics environment, select the launchpad to complete your reports installation. See Using the launchpad to install VMware components.

• Reports Installation Manager GUI

If you do not have a browser but you have a graphics environment, run the installer in GUI mode by using the following syntax: setup_platform.exe/.bin. For more information about using the Reports Installation Manager GUI to install reports, see "VMware: Installing Capacity Planner Reports" on page 58.

VMware: Installing Capacity Planner Reports

You can use the Launchpad or the Reports installer to install Capacity Planner Reports for VMware .

Procedure

- 1. Select the type of installation. For more information on selecting the type of installation, see "VMware: Types of report installation" on page 57.
- 2. You might need to point to Java 1.5+ through your system PATH. Make sure that your system PATH contains a valid path to a Java virtual machine, for example: # PATH=\$PATH:/ibmjre50/ibm-java-i386-50/jre/bin
- **3**. From the directory from which you extracted the reports package, run the file in the following table depending on your operating system and version of Tivoli Common Reporting that is installed.

Operating system	File
AIX	setup_aix.bin
HP-UX	setup_hpux.bin
Linux	<pre>setup_linux.bin</pre>
Solaris	setup_solaris.bin
Windows	setup_windows.exe

Table 3. Setup files for Tivoli Common Reporting 2.1.1

Table 4. Setup files for Tivoli Common Reporting 3.1 or later

Operating system	File
AIX	setup_aix.bin
Linux	setup_linux.bin
Windows	setup_windows.exe

- 4. Select a language, and then click **OK**.
- 5. Accept the license agreement.
- 6. Select the location where the Tivoli Common Reporting server is installed (not the location where the reports are to be installed).
 - For Tivoli Common Reporting V2.1, the default path is C:\IBM\tivoli\tipv2Components\ TCRComponent or /opt/IBM/tivoli/tipv2Components/TCRComponent. The path must end with the /TCRComponent folder.
 - For Tivoli Common Reporting V3.1 or later, the default path is C:\Program Files\IBM\JazzSM\ reporting or /opt/IBM/JazzSM/reporting. The path must end with the /reporting folder.

Note: If Tivoli Common Reporting installation is distributed, reports must be installed on the dispatcher site only.

- 7. On the Welcome page, click Next
- 8. On the Choose the Installation Folder page, select the location where the Tivoli Common Reporting server is installed, and then click **Next**.
- 9. On the Choose the reports for the installation page, select the type of report that you want to install, and then click **Next**:
 - IBM Infrastructure Management Capacity Planner Reports for VMware

🛎 Report Installer	
	Choose the reports for the installation
	Choose the reports for the installation Cognos reports IBM Infrastructure Management Capacity Planner Reports for VM IBM Infrastructure Management Capacity Planner Reports for Pou IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports v IBM Tivoli Monitoring for Virtual Environments VMware Reports v7
InstallAnywhere	Description Capacity Planner reports require the Capaci
Cancel <u>H</u> elp	<u>Previous</u> <u>N</u> ext

Figure 27. Choose the reports for the installation page

10. Configure Cognos data sources to connect to Tivoli Data Warehouse.

a. On the Cognos Engine Configuration page, enter the user name and password for Tivoli Common Reporting, and then click **Next**.

📲 Report Installer	
	Cognos Engine Configuration
	Enter the Tivoli Common Reporting user name
	tipadmin
	Enter the Tivoli Common Reporting user password

InstallAnywhere	
Cancel Help	Previous Next

Figure 28. Cognos Engine Configuration page

- 11. Define common dimensions on the Tivoli Data Warehouse.
 - a. Enter the following details, and then click **Next**:
 - In the Enter the database user name field, type the DB2 administrator user name.
 - In the Enter the user password field, type the password.
 - In the **Choose the database type** list, select DB2.
 - In the Enter the database name field, type TADFDCDB.
| CREPORT Installer | |
|--|--|
| Cognos D | ata Source TADFDC_REPORTING_DS Configuration |
| | Enter the database user name
db2admin |
| | Enter the user password |
| | ***** |
| | Choose the database type |
| | DB2 |
| | Enter the database name |
| | TADFDCDB |
| | Skip this panel (not recommended) |
| InstallAnywhere
Cancel <u>H</u> elp | <u>Previous</u> <u>N</u> ext |

Figure 29. Cognos Data Source TADFDC_REPORTING_DS Configuration page

12. On the Report Installer Summary page, verify the reports to be installed, click **Install**, and wait for the installer to finish.

CREPORT Installer		<u>- </u>
	The following report sets will be installed: Cognos reports IBM Infrastructure Management Capacity Planner Reports for VMware v7.2.0.2	
InstallAnywhere Cancel <u>H</u> elp	<u>Previous</u>	stall

Figure 30. Report Installer Summary page

Note: The Installation results panel shows the status of all installation actions for every item or report. One log file and one trace file are included. Both files are in the user home directory, with the following names: Report_Installer_Install_mm_dd_yyyy_hh_mm_ss.log (Log) and the Report_Installer_For_TCR_Output.txt (Trace)



Figure 31. Installation results page

13. Click Done.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

Installing reports by using the command line

You can install the reports by using the command line.

Procedure

- 1. Run the **setup_***platform*.**exe** -**i** or **setup_***platform*.**bin** -**i** command.
- 2. Choose your installation language.
- 3. Enter the location of the TCRComponent directory.
- 4. Choose the type of reports to be installed.
- 5. Enter your Tivoli Common Reporting user name and password.
- 6. Configure your datasource and data scripts. Some report packages might not have data scripts.
- 7. An installation summary is provided, then press Enter to begin installation.

Installing reports by using the silent mode

You can install the reports by using the silent mode.

Procedure

- 1. Create the silent installer response file, and name the file as silent_installer.properties.
- 2. Run the setup_platform.exe/.bin -i silent -f path_to_response_file command.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

VMWare: Uninstalling Capacity Planner Reports

The reports installer does not support uninstalling reports. However, you can manually delete the reports packages by using the Dashboard Application Services Hub.

About this task

To delete the reports manually, use the Dashboard Application Services Hub for Tivoli Common Reporting 3.1 or later.

Procedure

Use the following procedure to uninstall reports manually:

- 1. Log in to the Tivoli Common Reporting interface, and go to Common Reporting.
- 2. Select the check box for the reports package that you want to delete.

Note: You can only view the reports that you installed.

ly Folders

 Image: Name €

 Image: Common Reporting

 Image: IBM Infrastructure Management Capacity Planner Reports for PowerVM

 Image: IBM Infrastructure Management Capacity Planner Reports for VMware

 Image: IBM Tivoli Monitoring for Linux KVM Reports

 Image: IBM Tivoli Monitoring for Virtual Environments Reports

Figure 32. Reports packages list

3. Click the **Delete** icon.

tipadmin	\$ª [_		Q	~	1	*	8.	L		in 🗸	3	*
											1			
		:8:		B	88	*	9	68	of	Þ	13	×		Į a ⊮≋

Figure 33. Delete icon

66 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 4. Updating the Capacity Planners and Dashboard for VMware from an earlier version

If you installed Capacity Planner for PowerVM, Capacity Planner for VMware, or Dashboard for VMware version 7.2 or V7.2 with a fix pack, you can update the components to 7.2 Fix Pack 2.

Updating components by using the IBM Installation Manager GUI

You can use Installation Manager to update the Capacity Planner for PowerVM, Capacity Planner for VMware, and Dashboard for VMware.

Procedure

1. Start the Launchpad, and select an appropriate installation option.

For more information about selecting the options, see "VMware: Using the Launchpad to install components" on page 36 or For more information about selecting the options, see "PowerVM: Using the Launchpad to install components" on page 8.

🗐 ITMfVE			
ITM for Virtual E	nvironments	Select a language: English	• OK
	Installation Options This product contains installation of and reports. You can use this laum- perform a full installation or install Full Installation as administra Full Installation as non-admin You will be able to choose which or Install Capacity Planners and Install Capacity Planners and Install Reports	omponents for capacity planners, dashboards chpad to install them. Click the links below to components individually. tive user istrative user apacity planners and dashboards to be installed. Dashboards as administrative user Dashboards as non-administrative user	
	View Release Information Readme Installation Guide		
IBM.			

Figure 34. ITM for Virtual Environments page

2. Click Update.

🖄 IBM Installation Manager Eile Help			×
IBM Installation Manager			
	Install Install software packages.		
	Update Discover and install updates and fixes to installed software packages.	Manage Licenses	
	Change installed software packages by adding or removing features and functions.	Uninstall	
IBM.			

Figure 35. IBM Installation Manager page

3. The Installation Manager displays the existing package groups. Verify the details, and then click **Next**.

51 22 23 33.0 V H2		
ackage Group Name	Directory	
Teo IBM WebSphere Application Server V8.5	C:\Program Files (x86)\IBM\WebSphere\AppServer	
Core services in Jazz for Service Management	C: Program Files (Dm Qazz5M	
0		
Update <u>a</u> ll		
etails		
ore services in Jazz for Service Management		
Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared		
Installation Directory: C:\Program Files\jbm\JazzSM		
Eclipse IDE: C:\Program Files\ibm\JazzSM		
Translations: English		
Architecture: 64-bit		
stalled Packages and Fixes		
Registry Services 1.1		
Administration Services 1.1		
Security Services 1.1		
IBM Dashboard Application Services Hub 3.1		
Administration Services UI 1.1		
IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0		
IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0		1
	-	 _
In this deale Hundgenen especty hannel for while 7,200	-	

Figure 36. Update Packages page for existing package group

- 4. Select the following check boxes depending on the components you want to install:
 - IBM Infrastructure Management Capacity Planner for PowerVM
 - IBM Infrastructure Management Capacity Planner for VMware
 - IBM Infrastructure Management Dashboard for VMware

Also, select the check box for the Version.

elect the updates to install.			CL.
Jodate	Recommended	Vendor	
⊟… 🗹 🗞 Core services in Jazz for Service Management			_
🛱 🔽 🔀 IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0 (Installed)	83		
	~	IBM	
Unitrastructure Management Capacity Planner for VMWare 7.2.0.0 (Installed)	1	TBM	
Version 7.2.0.2 IBM Infrastructure Management Dashboard for VMware 7.2.0.0 (Installed)		1011	
└── 🗹 🕼 Version 7.2.0.2	\checkmark	IBM	
Show recommended only		Select Recomme	ende
).	

Figure 37. Update Packages page for updates to install

5. Review the information on the Update Packages page, and then click Next.

Update Package	es								F
Select the features to insi	tall.								
Update Packages	Updates	Features	Summary						
Features									
E 🗹 🚺 IBM Infrastruc	ture Managemer	nt Capacity Planner 1	for PowerVM 7.2.0.	.2					
	n tion								
E 🖸 🚺 IBM Infrastruc	ture Managemer	nt Capacity Planner f	for VMware 7.2.0.2	2					
Installatio	n								
Configura	tion ture Managemer	at Dashboard for VM	ware 7 2 0 2						
Installatio	n n	IC Dashboard for VM	Ware 7.2.0.2						
Configura	tion								
Show dependencies							Expand All	Collapse All	Restore
Show dependencies ≶- Selected by Installatic	on Manager beca	use of dependencies	s				Expand All	Collapse All	Restore
Show dependencies	on Manager beca	use of dependencies	s			 	Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installatio Details	on Manager beca	use of dependencies	s			 	Expand All	Collapse All	Restore
Show dependencies Selected by Installatio Details IBM Infrastructure Ma	on Manager beca	use of dependencies	s • PowerVM 7.2.0.	2		 	Expand All	Collapse All	<u>R</u> estore
Show dependencies Selected by Installation Details IBM Infrastructure Manag	n Manager beca nagement Cap	use of dependencies Pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. 1 installer	2		 	Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installation Details IBM Infrastructure Manag	on Manager beca nagement Cap ement Capacity P	use of dependencies pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. 1 installer	2		 	Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installatio Details IBM Infrastructure Manag	on Manager beca nagement Cap ement Capacity F	use of dependencies pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. 1 installer	2			Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installatio Details IBM Infrastructure Manag	on Manager beca nagement Cap ement Capacity F	use of dependencies pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. I installer	2		 	Expand All	<u>C</u> ollapse All	Restore
Show dependencies - Selected by Installation Details IBM Infrastructure Manag Disk Space Information	n Manager beca nagement Car ement Capacity F	use of dependencies pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. I installer	2		 	Expand All	<u>C</u> ollapse All	<u>R</u> estore
Show dependencies Selected by Installation Details IBM Infrastructure Manag IBM Infrastructure Manag Disk Space Information	n Manager beca nagement Caj ement Capacity F	use of dependencies pacity Planner for Planner for PowerVM	s • PowerVM 7.2.0. installer	2			Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installation Setails IBM Infrastructure Manag Disk Space Information	on Manager beca nagement Caj ement Capacity F n Volume	use of dependencies pacity Planner for Planner for PowerVM Required	s • PowerVM 7.2.0. 1 installer Temporary	2 Total	Available		Expand All	<u>C</u> ollapse All	Restore
Show dependencies Selected by Installatio Details IBM Infrastructure Manag Disk Space Informatio Shared Resources Area	on Manager beca nagement Ca ement Capacity F n Volume C:	use of dependencies pacity Planner for Planner for PowerVM Required 184.72 MB	s PowerVM 7.2.0. I installer Temporary 24.57 MB	2 Total 209.29 MB	Available 13.60 GB		Expand All	<u>C</u> ollapse All	Restore

Figure 38. Update packages page for features to install

- 6. Enter the WebSphere Application Server user name and password for Dashboard Application Services Hub, and then click **Validate**.
- 7. After the user name and password are validated, click Next.

Update Padages Updates Summary Common Configurations WebSphere Application Server Credentals WebSphere Application Server Credentals WebSphere Application Server Credentals Web	Update Packages ① Provided credentials are valid.	
	Update Packages Updates Features Summary Common Configurations Some c	Common Configurations WebSphere Application Server Credentials User name: smadmin Password: ••••••••• Walidate

Figure 39. Update Packages page to validate WebSphere Application Server Credentials

- 8. Enter the Capacity Planner Database Schema details, and then click Validate.
- 9. After the details are validated, click Next.



Figure 40. Update Packages page to validate the Capacity Planner Database Schema details for VMware

Update Packages		
Press the Validate button to continue.		
Common Configurations WebSphere Application Server Credentials Given Management Capacity Planner for VMware 7.2.0.2 Gapacity Planner Database Schema Creation Gapacity Planner Database Schema Creation Gapacity Planner Database Schema Creation	Capacity Planner Database Schema Creation Capacity Planner Database Schema Creation Capacity Planner Database Server Host name: Capacity Planner Database Server Port: Capacity Planner Database Admin Password: Capacity Planner Database Name: Validate	Igement Capacity Planner for PowerVM 7.2.0.2
)	< Back Next > Update Cancel

Figure 41. Update Packages page to validate the Capacity Planner Database Schema details for PowerVM

10. Review the information in the Update Packages summary page, and then click **Update**.



Figure 42. Update Packages page to show the summary

11. The Results window displays the package installation status. Click **Finish** to complete the installation.



Figure 43. Update Packages to show that the components are updated

12. To verify the updates, click **File** > **View Installed Packages**.

🖄 Installed Packages

Installed Packages and Fixes

	Vendor	License
E R IBM WebSphere Application Server V8.5		
IBM WebSphere Application Server_8.5.0.1	IBM	
🕞 IBM WebSphere SDK Java Technology Edition (Optional)_7.0.2.0	IBM	
🙀 Jazz for Service Management extension for IBM WebSphere 8.5_1.1.0.1	IBM	
😑 🔩 Core services in Jazz for Service Management		
Registry Services_1.1	IBM	
Administration Services 1.1.0.1	IBM	
Security Services 1.1.0.1	IBM	
IBM Dashboard Application Services Hub 3.1.0.1	IBM	
Administration Services UI_1.1.0.1	IBM	
IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2	IBM	
IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2	IBM	
IBM Infrastructure Management Dashboard for VMware 7.2.0.2	IBM	
Details		
IBM WebSphere Application Server V8.5 • Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared • Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer • Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer		
IBM WebSphere Application Server V8.5 • Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared • Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer • Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer • Translations: English		
IBM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English		
IBM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English		
IBM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English		
IBM WebSphere Application Server V8.5 Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer Translations: English		

Figure 44. Installed Packages and fixes page

Updating components by using the IBM Installation Manager console mode

There are specific steps to update the Capacity Planner for VMware, Capacity Planner for PowerVM, and the Dashboard for VMware by using the Console mode.

Procedure

- 1. Start the console mode.
 - a. Open a command line.

×

b. From the command line, navigate to the IBM Installation Manager tools directory.

Default path of the tools directory for administrator

On Windows systems: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

On operating systems other than Windows: opt/IBM/InstallationManager/eclipse/tools

Default path of the tools directory for non-administrator

On Windows systems: C:\Users\username\IBM\Installation Manager\eclipse\tools

On operating systems other than Windows: /home/username/IBM/InstallationManager/ eclipse/tools

For more information on the default paths for IBM Installation Manager installation directory, see Install as an administrator, nonadministrator, or group (http://pic.dhe.ibm.com/infocenter/install/v1r5/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_admin_nonadmin.html).

c. Run the following command:

On Windows systems:

imcl.exe -c

For example, C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools>imcl.exe -c

On operating systems other than Windows:

./imcl -c

For example, opt/IBM/InstallationManager/eclipse/tools>./imcl -c

2. Adding the repository. The command line displays the following options:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- P. Preferences
- ------
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- -----
- X. Exit Installation Manager
- a. At the command line, type **P** to set the preferences. The command line displays the preferences. ====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- 2. Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
- 6. Passport Advantage
- 7. Updates
- R. Return to Main Menu
- b. Type **1** to select the repository. The command line displays other options.

====> IBM Installation Manager> Preferences> Repositories

```
Repositories:
```

- Other Options:
 - D. Add Repository
 - S. [X] Search service repositories during installation and updates

- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- c. Type **D** to add the repository.
- d. Type the repository location where you extracted the installation image, and select the diskTag.inf file. For example, D:\diskTag.inf.

====> IBM Installation Manager> Preferences> Repositories> Add repository
Enter a new repository location. To skip, press Enter:
----> D:\diskTag.inf
Checking repositories...

====> IBM Installation Manager> Preferences> Repositories

Repositories:

1. [X] D:\diskTag.inf

Other Options:

- D. Add Repository
- S. [X] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- e. Type S to unselect the option Search service repositories during installation and updates.
 - ====> IBM Installation Manager> Preferences> Repositories

Repositories:

1. [X] D:\diskTag.inf

Other Options:

- D. Add Repository
- S. [] Search service repositories during installation and updates
- R. Restore Defaults
- A. Apply Changes and Return to Preferences Menu
- P. Temporarily Keep Changes and Return to Preferences Menu
- f. Type **A** to apply changes and return to the Preferences menu.

====> IBM Installation Manager> Preferences

Select:

- 1. Repositories
- 2. Appearance
- 3. Files for Rollback
- 4. SSL/TLS
- 5. HTTP/FTP Proxy
- 6. Passport Advantage
- 7. Updates

R. Return to Main Menu

- g. Type R to return to the main menu.
- 3. Selecting the package to update. The command line displays the following option:

====> IBM Installation Manager

Select:

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

```
Other Options:
```

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- P. Preferences

- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- -----
- X. Exit Installation Manager

```
a. Type 2 to install packages.
```

```
Checking repositories...

Loading repositories...

====> IBM Installation Manager> Update

Select a package group to update:

1. [] IBM WebSphere Application Server V8.5

2. [] Core services in Jazz for Service Management

Other Options:

U. Update All

C. Cancel
```

b. Type 2 to select package group Core services in Jazz for Service Management.

```
====> IBM Installation Manager> Update
```

- Select a package group to update:
 - 1. [] IBM WebSphere Application Server V8.5
 - 2. [X] Core services in Jazz for Service Management

Details of package group Core services in Jazz for Service Management:

```
Package Group Name:Core services in Jazz for Service ManagementShared Resources Directory:C:\Program Files (x86)\IBM\IMSharedInstallation Directory:C:\Program Files\ibm\JazzSMEclipse IDE:C:\Program Files\ibm\JazzSMTranslations:EnglishArchitecture:64-bit
```

Other Options:

```
U. Update All
A. Unselect All
```

N. Next, C. Cancel

c. Type **N** to continue the update.

```
Preparing installed packages in the selected package groups...
Validating package group locations...
Searching updates...
Preparing and resolving the selected packages...
Preparing and resolving the selected packages...
Preparing and resolving the selected packages...
Validating the selected updates...
====> IBM Installation Manager> Update> Packages
Package group: Core services in Jazz for Service Management
Update packages:
     1. [ ] Registry Services 1.1

    [] Administration Services 1.1.0.1
    3. [] Security Services 1.1.0.1

     4. [ ] IBM Dashboard Application Services Hub 3.1.0.1
     5. [] Administration Services UI 1.1.0.1
     6-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0
       7. [X] Version 7.2.0.2
     8-. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0
       9. [X] Version 7.2.0.2
    10-. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.0
      11. [X] Version 7.2.0.2
Other Options:
     A. Show All
     R. Select Recommended
```

B. Back, N. Next, C. Cancel

d. Type N to continue with by default selected all three components.

```
Unpreparing com.ibm.tivoli.monitoring.powervmcp 7.2.100.20130725 0625.
   Validating the selected updates...
   ====> IBM Installation Manager> Update> Packages
   Package group: Core services in Jazz for Service Management
   Update packages:
        1. [] Registry Services 1.1
        2. [] Administration Services 1.1.0.1
        3. [ ] Security Services 1.1.0.1
        4. [ ] IBM Dashboard Application Services Hub 3.1.0.1
        5. [] Administration Services UI 1.1.0.1
        6-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0
          7. [X] Version 7.2.0.2
        8-. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0
          9. [X] Version 7.2.0.2
       10-. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.0
         11. [X] Version 7.2.0.2
   Other Options:
        A. Show All
        R. Select Recommended
        B. Back,
                      N. Next,
                                    C. Cancel
e. Type N to continue the update.
   Preparing and resolving the selected packages...
   Validating the selected updates...
   ====> IBM Installation Manager> Update> Packages
   Package group: Core services in Jazz for Service Management
   Update packages:
        1. [] Registry Services 1.1
        2. [] Administration Services 1.1.0.1
        3. [] Security Services 1.1.0.1
        4. [ ] IBM Dashboard Application Services Hub 3.1.0.1
        5. [] Administration Services UI 1.1.0.1
        6-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0
          7. [X] Version 7.2.0.2
        8-. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0
          9. [X] Version 7.2.0.2
       10-. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.0
         11. [X] Version 7.2.0.2
   Other Options:
        A. Show All
        R. Select Recommended
                                    C. Cancel
        B. Back,
                      N. Next,
f. Type N. The command line displays the features of the packages.
   ====> IBM Installation Manager> Update> Packages> Features
   IBM Infrastructure Management Capacity Planner for PowerVM
        1. [X] Installation
       2. [X] Configuration
   IBM Infrastructure Management Capacity Planner for VMware
       3. [X] Installation
       4. [X] Configuration
   IBM Infrastructure Management Dashboard for VMware
       5. [X] Installation
       6. [X] Configuration
                                    C. Cancel
       B. Back,
                      N. Next,
```

g. Type **N** to continue the update. The command line displays the features of the packages.

====> IBM Installation Manager> Install> Location> Features
IBM Infrastructure Management Capacity Planner for VMware
1. [X] Installation
2. [X] Configuration
IBM Infrastructure Management Dashboard for VMware
3. [X] Installation
4. [X] Configuration
B. Back, N. Next, C. Cancel
4. Validating Dashboard Application Services Hub credentials. The command line displays the following
options:
====> IBM Installation Manager> Update> Packages> Features> Custom panels
---- Common configuration:

> WebSphere Application Server Credentials
User name for WebSphere Application Server:
----> [smadmin] smadmin

- a. Type the user name for WebSphere Application Server. For example, smadmin. Password for WebSphere Application Server:
- b. Type the password for WebSphere Application Server.

---- Configuration validation: Validating WebSphere Application Server credentials... Provided credentials are valid.

- 5. Validating the Capacity Planner database credentials for Capacity Planner for VMware. ---- Configuration for IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 Capacity Planner Database Server Host name:
 - a. Type the Capacity Planner Database Server Host name. For example, localhost. Capacity Planner Database Server Port:
 - b. Type the Capacity Planner Database Server Port. For example 50000. Capacity Planner Database Admin name:
 - C. Type the Capacity Planner Database Admin name. For example, db2admin.
 Capacity Planner Database Admin Password:
 - d. Type the Capacity Planner Database Admin Password. Capacity Planner Database Name:
 - **e.** Type the Capacity Planner Database Name. For example, **TADFDCDB**. If you do not select any option for the validation, the system automatically validates the data that you entered. If the validation fails, enter the correct details again.

Validation in progress. Please wait... The parameters were verified successfully. B. Back, N. Next, C. Cancel

6. Validating the Capacity Planner database credentials Capacity Planner for PowerVM.

---- Configuration for IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2 Capacity Planner Database Server Host name:

- a. Type the Capacity Planner Database Server Host name. For example, localhost.
 Capacity Planner Database Server Port:
- b. Type the Capacity Planner Database Server Port. For example 50000.
 Capacity Planner Database Admin name:
- C. Type the Capacity Planner Database Admin name. For example, db2admin.
 Capacity Planner Database Admin Password:
- d. Type the Capacity Planner Database Admin Password.Capacity Planner Database Name:

e. Type the Capacity Planner Database Name. For example, **TADFDCDB**. If you do not select any option for the validation, the system automatically validates the data that you entered. If the validation fails, enter the correct details again.

```
Validation in progress. Please wait...
The parameters were verified successfully.
B. Back, N. Next, C. Cancel
```

7. Type **N** to continue the update. Check the summary of the update. The command line displays the summary of the packages that you are going to update.

```
====> IBM Installation Manager> Update> Packages> Features> Custom panels> Summary
Target Location:
 Shared Resources Directory : C:\Program Files (x86)\IBM\IMShared
Update packages:
    1-. Core services in Jazz for Service Management (C:\Program Files\ibm\JazzSM)
     2-. IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0
       3-. Features to install:
         4. Installation
         5. Configuration
     6-. IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0
       7-. Features to install:
         8. Installation
         9. Configuration
     10-. IBM Infrastructure Management Dashboard for VMware 7.2.0.0
      11-. Features to install:
        12. Installation
        13. Configuration
Options:
    G. Generate an Update Response File
    B. Back, U. Update, C. Cancel
a. Type U to update the packages.
                 25%
                               50% 75%
                                                                 100%
   -----|-----|-----|
   ====> IBM Installation Manager> Update> Packages> Features> Custom panels>Summary> completion
  The update completed successfully.
  Options:
       F. Finish
a. Type F to finish the updates.
```

Updating components by using the IBM Installation Manager silent mode

You can update the packages by using the silent mode. For more information about installing Capacity Planner for PowerVM in silent mode, see "PowerVM: Using the IBM Installation Manager in silent mode to install the Capacity Planner " on page 11, and Capacity Planner for VMware and Dashboard for VMware, see "VMware: Using the IBM Installation Manager silent mode to install the Capacity Planner and the Dashboard" on page 39.

84 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 5. Rolling back Capacity Planners and Dashboard for VMware to an earlier version

If you installed Capacity Planner for PowerVM, Capacity Planner for VMware, and Dashboard for VMware with a version 7.2 Fix Pack 2, you can roll back to an earlier installed version.

Rolling back components by using the IBM Installation Manager GUI

You can use the IBM Installation Manager GUI to roll back the following components: Capacity Planner for PowerVM, Capacity Planner for VMware, and Dashboard for VMware

Procedure

- 1. Start the Installation Manager by using any one of the following steps:
 - On Windows systems, select Start > IBM Installation Manager > IBM Installation Manager.
 - On Linux or AIX systems, in a command window, open the /opt/IBM/InstallationManager/ eclipse directory, and enter ./IBMIM or ./groupinst for group mode.
 - Double-click the IBMIM.exe file that is located in the eclipse subdirectory in the directory where the IBM Installation Manager is installed. The default path for IBM Installation Manager on Windows systems is C:/Program Files/IBM Installation Manager or C:/Program Files(86)/IBM Installation Manager and on UNIX systems, the path is /opt/IBM/Installation Manager.
- 2. Click Roll Back.

🖄 IBM Installation Manager <u>Fi</u> le <u>H</u> elp		_ _ ×
IBM Installation Manager		
6	Install Install software packages.	
	Update Discover and install updates and fixes to installed software packages.	Manage Licenses Roll Back
	Modify Change installed software packages by adding or removing features and functions.	Uninstall
IBM.		

Figure 45. IBM Installation Manager

3. On the Roll Back Packages page, select a package group to find a roll back package for, and then click **Next**.

Select a package group to find rollback packages for.		the second
Package Group Name	Directory	
$= {}^{\rm e_{\rm Rs}}$ IBM WebSphere Application Server V8.5 $= {}^{\rm e_{\rm Rs}}$ Core services in Jazz for Service Management	C:\Program Files (x86)\IBM\WebSphere\AppServer C:\Program Files\bm\JazzSM	
etails		
ore services in Jazz for Service Management		
Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared		
Installation Directory: C:\Program Files\jbm\JazzSM		
Eclipse IDE: C:\Program Files\jbm\JazzSM		
Translations: English		
Architecture: 64-bit		
installed Packages and Fixes		
Registry Services 1.1		
Administration Services 1.1.0.1		
Security Services 1.1.0.1		
IBM Dashboard Application Services Hub 3.1.0.1		
Administration Services UI 1.1.0.1		
IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2		
IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2		
	_	

Figure 46. Roll Back Packages page

4. Select the check boxes for the components that you want to roll back to, and then click Next.

Rollback Packages	Dependency	Vendor	
El- 🕑 🗞 Core services in Jazz for Service Management			
Em U Administration Services 1.1.0.1 (Installed)		IBM	
🛱 🔲 🚺 Security Services 1.1.0.1 (Installed)			
Im T 🖗 Version 1.1		IBM	
		IBM	
🛱 🗖 🕼 Administration Services UI 1.1.0.1 (Installed)			
		IBM	
Version 7.2.0.0		IBM	
🗄 🖓 🕼 IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 (Installed)		12200	
E IBM Infrastructure Management Dashboard for VMware 7.2.0.2 (Installed)		IBM	
		IBM	
etails 3M Infrastructure Management Dashboard for VMware 7.2.0.0 3M Infrastructure Management Dashboard for VMware Installer <u>More info</u> Repository: C: Users (Administrator Desktop \tmfve_dashboard_201302260740			

Figure 47. Roll back Packages page to select the components to roll back

- 5. Enter the WebSphere Application Server user name and password for Dashboard Application Services Hub, and then click **Validate**.
- 6. After the user name and password are validated, click Next.

Roll Back Packages	
Roll Back Packages Features Summary	dra dra
Common Configurations WebSphere Application Server Credentials Capacity Planner Database Schema Creation Capacity Planner Database Schema Creation Capacity Planner Database Schema Creation Capacity Planner Database Schema Creation	Common Configurations WebSphere Application Server Credentials User name: smadmin Password: •••••••• Walidate: ••••••••
	<back next=""> Roll Back Cancel</back>

Figure 48. Roll Back Packages page to validate WebSphere Application Server Credentials

- 7. Enter the Capacity Planner Database Schema details, and then click Validate.
- 8. After the details are validated, click Next.

T Albumation Homoget		
Roll Back Packages Press the Validate button to continue. Roll Back Packages Features Summary		
Connon Configurations Construction Constr	Configuration for IBH Infrastructure Mana Capacity Planner Database Schema Creation Capacity Planner Database Server Host name: Capacity Planner Database Server Port: Capacity Planner Database Admin name: Capacity Planner Database Admin Password: Capacity Planner Database Name: Validate	agement Capacity Planner for PowerVII 7.2.0.0
		< <u>Back Next</u> > <u>Roll Back</u> Cancel

Figure 49. Roll Back Packages page to validate the Capacity Planner Database Schema details for PowerVM



Figure 50. Roll Back Packages page to validate the Capacity Planner Database Schema details for VMware

9. Review the information in the Roll Back Packages summary page, and then click Roll Back.

Roll Back Packages		
Review the summary information.		The second se
Roll Back Packages Features	Summary	
Target Location		
Package Group Name: Core services in Jaz	z for Service Management	
Installation Directory: C:\Program Files\jbr	n\JazzSM	
Shared Resources Directory: C:\Program Files (xi	86)\IBM\IMShared	
Rollback Packages		
Version	Package	Package Group Name
Version 7.2.0.0	IBM Infrastructure Management Capacity P	Core services in Jazz for Service Management
0. Version 7.2.0.0	IBM Infrastructure Management Capacity P	Core services in Jazz for Service Management
10. Version 7.2.0.0	IBM Infrastructure Management Dashboard	Core services in Jazz for Service Management
Disk Space Information	Details	
Disk Space Information	Details IBM Infrastruct	ure Management Capacity Planner for PowerVM 7.2.0.0
Disk Space Information Total Available Space C: 12.91 GB	Details IBM Infrastruct	ure Management Capacity Planner for PowerVM 7.2.0.0
Disk Space Information Total Available Space C: 12.91 GB Fotal Download Size: 0.0 KB Fotal Installation Size: 212.68 MB	Details IBM Infrastruct	ure Management Capacity Planner for PowerVM 7.2.0.0
Disk Space Information Total Available Space C: 12.91 GB Fotal Download Size: 0.0 KB Fotal Installation Size: 212.68 MB ▶ Repository Information	Details IBM Infrastruct	ure Management Capacity Planner for PowerVM 7.2.0.0

Figure 51. Roll Back Packages page to show the summary

10. The Roll Back Packages results window displays the package that is rolled back. Click **Finish** to complete the roll back.



Figure 52. Roll Back Packages to show that the components are rolled back

11. To verify the rolled back components, click File > View Installed Packages.

🖄 Installed Packages

Installed Packages and Fixes

Installed Packages	Vendor	License	
E Rea IBM WebSphere Application Server V8.5			
IBM WebSphere Application Server_8.5.0.1	IBM		
💫 IBM WebSphere SDK Java Technology Edition (Optional)_7.0.2.0	IBM		
🙀 Jazz for Service Management extension for IBM WebSphere 8.5_1.1.0.1	IBM		
E ^P ₂₀ Core services in Jazz for Service Management			
Registry Services_1.1	IBM		
Administration Services_1.1.0.1	IBM		
Security Services_1.1.0.1	IBM		
IBM Dashboard Application Services Hub_3.1.0.1	IBM		
Administration Services UI_1.1.0.1	IBM		
BM Infrastructure Management Capacity Planner for PowerVM_7.2.0.0	IBM		
BM Infrastructure Management Capacity Planner for VMware_7.2.0.0	IBM		
IBM Infrastructure Management Dashboard for VMware 7.2.0.0	IBM		
petails			
BM WebSphere Application Server V8.5			
Shared Resources Directory: C:\Program Files (x86)\IBM\IMShared			
Installation Directory: C:\Program Files (x86)\IBM\WebSphere\AppServer			
Eclipse IDE: C:\Program Files (x86)\IBM\WebSphere\AppServer			
Translations: English			
• Translations: English			
Translations: English			Close

Figure 53. Install Packages page to show the currently installed packaged

Rolling back components by using the IBM Installation Manager console mode

There are specific steps to roll back the Capacity Planner for VMware, Capacity Planner for PowerVM, and the Dashboard for VMware by using the console mode.

Procedure

- 1. Start the console mode.
 - a. Open a command line.

b. From the command line, navigate to the IBM Installation Manager tools directory.

Default path of the tools directory for administrator

On Windows systems: C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools

On operating systems other than Windows: opt/IBM/InstallationManager/eclipse/tools

Default path of the tools directory for non-administrator

On Windows systems: C:\Users\username\IBM\Installation Manager\eclipse\tools

On operating systems other than Windows: /home/username/IBM/InstallationManager/ eclipse/tools

For more information on the default paths for IBM Installation Manager installation directory, see Install as an administrator, nonadministrator, or group (http://pic.dhe.ibm.com/infocenter/ install/v1r5/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_admin_nonadmin.html).

c. Run the following command:

On Windows systems:

imcl.exe -c

```
For example, C:\Program Files (x86)\ibm\Installation Manager\eclipse\tools>imcl.exe -c
```

On operating systems other than Windows:

./imcl -c

For example, opt/IBM/InstallationManager/eclipse/tools>./imcl -c

2. Type 4 when the command line displays the following options:

```
====> IBM Installation Manager
```

```
Select:
```

- 1. Install Install software packages
- 2. Update Find and install updates and fixes to installed software packages
- 3. Modify Change installed software packages
- 4. Roll Back Revert to an earlier version of installed software packages
- 5. Uninstall Remove installed software packages

Other Options:

- L. View Logs
- S. View Installation History
- V. View Installed Packages
- -----
- P. Preferences
- E. Export Data for Problem Analysis
- A. About IBM Installation Manager
- -----
- X. Exit Installation Manager

```
----> 4
```

3. Select a package group to roll back when the command line displays the following option:

====> IBM Installation Manager> Roll Back

Select a package group to roll back:

- 1. [] IBM WebSphere Application Server V8.5
- 2. [] Core services in Jazz for Service Management

C. Cancel

- a. Type **2** to select a package group to roll back.
 - ====> IBM Installation Manager> Roll Back

Select a package group to roll back:

- [] IBM WebSphere Application Server V8.5
- 2. [X] Core services in Jazz for Service Management

Details of package group Core services in Jazz for Service Management: Package Group Name : Core services in Jazz for Service Management

hared

N. Next, C. Cancel

b. Type N to continue.

Preparing installed packages in the selected package groups... Validating package group locations... Searching rollback...

4. The command line displays the following list of roll back packages.

====> IBM Installation Manager> Roll Back> Packages Package group: Core services in Jazz for Service Management Roll Back packages: 1-. [] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2 2. [] Version 7.2.0.0 3-. [] IBM Dashboard Application Services Hub 3.1.0.1 4. [] Version 3.1 5-. [] Administration Services UI 1.1.0.1 6. [] Version 1.1 7-. [] Administration Services 1.1.0.1 8. [] Version 1.19-. [] IBM Infrastructure Management Dashboard for VMware 7.2.0.2 10. [] Version 7.2.0.0 11-. [] Security Services 1.1.0.1 12. [] Version 1.1 13-. [] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 14. [] Version 7.2.0.0 15. [] Registry Services 1.1 B. Back, C. Cancel a. Type 2 to select IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2 to roll back. Preparing and resolving the selected packages... Validating packages... ====> IBM Installation Manager> Roll Back> Packages Package group: Core services in Jazz for Service Management Roll Back packages:

```
1-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
  2. [X] Version 7.2.0.0
 3-. [] IBM Dashboard Application Services Hub 3.1.0.1
  4. [] Version 3.1
 5-. [] Administration Services UI 1.1.0.1
  6. [] Version 1.1
 7-. [] Administration Services 1.1.0.1
  8. [] Version 1.1
 9-. [] IBM Infrastructure Management Dashboard for VMware 7.2.0.2
10. [] Version 7.2.0.0
11-. [ ] Security Services 1.1.0.1
  12. [ ] Version 1.1
13-. [] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
  14. [] Version 7.2.0.0
15. [] Registry Services 1.1
 B. Back,
                             C. Cancel
               N. Next.
```

b. Type 10 to select IBM Infrastructure Management Dashboard for VMware 7.2.0.2 to roll back.
```
Preparing and resolving the selected packages...
Validating packages...
====> IBM Installation Manager> Roll Back> Packages
Package group: Core services in Jazz for Service Management
Roll Back packages:
     1-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
      2. [X] Version 7.2.0.0
     3-. [ ] IBM Dashboard Application Services Hub 3.1.0.1
      4. [] Version 3.1
    5-. [ ] Administration Services UI 1.1.0.1
6. [] Version 1.1
    7-. [] Administration Services 1.1.0.1
       8. [] Version 1.1
    9-. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.2
      10. [X] Version 7.2.0.0
    11-. [ ] Security Services 1.1.0.1
      12. [ ] Version 1.1
    13-. [ ] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
      14. [] Version 7.2.0.0
    15. [] Registry Services 1.1
    B. Back,
                   N. Next,
                                 C. Cancel
```

c. Type **14** to select IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2 to roll back.

```
Preparing and resolving the selected packages...
Validating packages...
====> IBM Installation Manager> Roll Back> Packages
Package group: Core services in Jazz for Service Management
```

d. Review the list of packages selected to roll back.

```
Roll Back packages:
     1-. [X] IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.2
       2. [X] Version 7.2.0.0
     3-. [ ] IBM Dashboard Application Services Hub 3.1.0.1
       4. [ ] Version 3.1
     5-. []_Administration Services UI 1.1.0.1
     6. [] Version 1.1
7-. [] Administration Services 1.1.0.1
       8. [] Version 1.1
     9-. [X] IBM Infrastructure Management Dashboard for VMware 7.2.0.2
      10. [X] Version 7.2.0.0
    11-. [] Security Services 1.1.0.1
      12. [] Version 1.1
    13-. [X] IBM Infrastructure Management Capacity Planner for VMware 7.2.0.2
      14. [X] Version 7.2.0.0
    15. [] Registry Services 1.1
                                  C. Cancel
     B. Back,
                   N. Next,
```

- e. Type N to continue the roll back.
- 5. Validating Dashboard Application Services Hubcredentials. The command line displays the following options:

```
====> IBM Installation Manager> Update> Packages> Features> Custom panels
---- Common configuration:
> WebSphere Application Server Credentials
User name for WebSphere Application Server:
-----> [smadmin] smadmin
a. Type the user name for WebSphere Application Server. For example, smadmin.
```

- Password for WebSphere Application Server:
- b. Type the password for WebSphere Application Server.

```
---- Configuration validation:
Validating WebSphere Application Server credentials...
Provided credentials are valid.
```

- Validating the Capacity Planner database credentials for Capacity Planner for VMware.
 ---- Configuration for IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0 Capacity Planner Database Server Host name:
 - a. Type the Capacity Planner Database Server Host name. For example, localhost. Capacity Planner Database Server Port:
 - b. Type the Capacity Planner Database Server Port. For example 50000.
 Capacity Planner Database Admin name:
 - c. Type the Capacity Planner Database Admin name. For example, **db2admin**. Capacity Planner Database Admin Password:
 - d. Type the Capacity Planner Database Admin Password. Capacity Planner Database Name:
 - e. Type the Capacity Planner Database Name. For example, **TADFDCDB**. If you do not select any option for the validation, the system automatically validates the data that you entered. If the validation fails, enter the correct details again.

```
Validation in progress. Please wait...
The parameters were verified successfully.
B. Back, N. Next, C. Cancel
```

7. Validating the Capacity Planner database credentials Capacity Planner for PowerVM.

---- Configuration for IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0 Capacity Planner Database Server Host name:

- a. Type the Capacity Planner Database Server Host name. For example, localhost.
 Capacity Planner Database Server Port:
- b. Type the Capacity Planner Database Server Port. For example 50000.
 Capacity Planner Database Admin name:
- c. Type the Capacity Planner Database Admin name. For example, db2admin. Capacity Planner Database Admin Password:
- d. Type the Capacity Planner Database Admin Password.Capacity Planner Database Name:
- e. Type the Capacity Planner Database Name. For example, **TADFDCDB**. If you do not select any option for the validation, the system automatically validates the data that you entered. If the validation fails, enter the correct details again.

```
Validation in progress. Please wait...
The parameters were verified successfully.
B. Back, N. Next, C. Cancel
```

8. Type N to continue the roll back. Check the summary of the roll back. The command line displays the summary of the packages that you want to roll back.

```
====> IBM Installation Manager> Roll Back> Packages> Custom panels> Summary
Target Location:
   Shared Resources Directory : C:\Program Files (x86)\IBM\IMShared
Roll Back packages:
    1-. Core services in Jazz for Service Management (C:\Program Files\ibm\JazzSM)
      2. IBM Infrastructure Management Capacity Planner for PowerVM 7.2.0.0
      3. IBM Infrastructure Management Capacity Planner for VMware 7.2.0.0
      4. IBM Infrastructure Management Dashboard for VMware 7.2.0.0
Options:
    G. Generate a Rollback Response File
             R. Roll Back,
    B. Back,
                               C. Cancel
a. Type R to continue the roll back.
                  25%
                                    50%
                                                    75%
                                                                       100%
                               --|
```

```
====> IBM Installation Manager> Roll Back> Packages> Custom panels> Summary> Completion
The roll back completed successfully.
Options:
F. Finish
```

b. Type ${\bf F}$ to finish the roll back.

100 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 6. Linux KVM and VMware Performance and Capacity Management Reports: Installation

Before you install the Performance and Capacity Management Reports for Linux KVM and Performance and Capacity Management Reports for VMware, you must catalog the Tivoli Data Warehouse database and database server. Also, you must catalog the Capacity Planner database and database server if the Capacity Planner database is not on the same server as the Tivoli Data Warehouse database.

Linux KVM and VMware: Prerequisites for installing the Performance and Capacity Management Reports

The Cognos reports require the completion of prerequisite steps for the reports to run.

All of the following prerequisite steps must be completed before you run the Performance and Capacity Management Reports:

- 1. Install Tivoli Common Reporting V3.1 or later.
- 2. If you have Tivoli Common Reporting V2.1.1, update the OS and virtualization agent reports and schedules from an earlier version of Tivoli Common Reporting to V3.1.

You can complete the migration by exporting the packages from the earlier version of Tivoli Common Reporting, and then importing the packages into the version 3.1.

For more information about upgrading the Tivoli Common Reporting V3.1, see "Upgrading to Tivoli Common Reporting Version 3.1" in the *Jazz for Service Management Installation Guide*.

- 3. Catalog the Tivoli Data Warehouse database and database server.
- 4. Obtain the reports from the product media.
- 5. Configure historical collection for the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines and the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI. For more information about configuring this historical collection, see "Configuring historical data collection for the Performance and Capacity Management Reports for VMware" on page 135.
- 6. Connect to Tivoli Data Warehouse by using the database client over ODBC.
- 7. Run the following commands on your Tivoli Data Warehouse database, where *DBNAME* is the name of the database and *itmuser* is the database administrator ID:

```
db2 connect to $DBNAME
db2 grant dbadm on database to user itmuser
db2 disconnect $DBNAME
```

For more information on the DB2 V9.7 requirement, see "Schemas" in the DB2 Database Administration information (Schemas).

Linux KVM and VMware: Cataloging databases and servers Procedure

 To catalog the Tivoli Data Warehouse database and database server, use the following syntax: db2 "catalog tcpip node node_id remote remote_host server remote_port" For example:

db2 "catalog tcpip node WHNODE remote testwarehouse1 server 50000"

Where:

node_id

The ID that is associated with the node, for example, WHNODE.

remote_host

The host name of the remote Tivoli Data Warehouse database server, for example, testwarehouse1.

remote_port

The port number of the remote Tivoli Data Warehouse database server, for example, 50000.

2. To catalog the Tivoli Data Warehouse database, use the following syntax: db2 "catalog database remote_db as local_db at node node_id authentication server" For example,

db2 "catalog database warehouse as WHDB at node WHNODE authentication server"

Where:

remote_db

The name of the remote Tivoli Data Warehouse database, for example, warehouse.

local_db

The alias is used to identify the remote warehouse database on the local computer, for example, WHDB.

node_id

The ID that is associated with the node, for example, WHDB.

Linux KVM and VMware: Types of report installation

You can use either the Launchpad or the Reports IBM Installation Manager to install the Performance and Capacity Management Reports for Linux KVM and Performance and Capacity Management Reports for VMware.

Select the type of installation that suits your environment:

• Launchpad

If you have a browser and a graphics environment, select the launchpad to complete your reports installation. See Using the launchpad to install VMware components.

Reports Installation Manager GUI

If you do not have a browser but you have a graphics environment, run the installer in GUI mode by using the following syntax: setup_platform.exe/.bin. For more information about using the Reports Installation Manager GUI to install reports, see "Linux KVM: Installing Performance and Capacity Management Reports" and "VMware: Installing Performance and Capacity Management Reports" on page 111.

Linux KVM: Installing Performance and Capacity Management Reports

You can use the Launchpad or the Reports installer to install Performance and Capacity Management Reports for Linux KVM . You must import the IBM Tivoli Monitoring for Linux KVM Cognos reports to run any report from the Linux KVM Reports package or the VMware Reports package.

Before you begin

Note: With this release, the reports package name does not contain the release version, such as 7.2, and so on. When the package name does not have a version number, the name is the same across releases, so different versions of the packages cannot co-exist in the Dashboard Application Services Hub. Before you install the reports, back up the existing reports package if the name is *IBM Tivoli Monitoring for Linux KVM Reports* or *IBM Tivoli Monitoring for Virtual Environments Reports*.

About this task

The IBM Tivoli Monitoring for Linux KVM Reports and VMware package contains an installer that performs the following tasks:

- · Importing the reports and data model into Tivoli Common Reporting
- Configuring a data source to connect to Tivoli Data Warehouse
- Running scripts to create and populate the common dimensions in Tivoli Data Warehouse

After you complete the steps for importing and running Cognos reports, you can run any report from the IBM Tivoli Monitoring for Linux KVM Reports package.

Procedure

- 1. Install the prerequisites. See "Linux KVM and VMware: Prerequisites for installing the Performance and Capacity Management Reports" on page 101
- 2. Select the type of installation. For more information on selecting the type of installation, see "Linux KVM and VMware: Types of report installation" on page 102.
- 3. You might need to point to Java 1.5+ through your system PATH. Make sure that your system PATH contains a valid path to a Java virtual machine, for example: # PATH=\$PATH:/ibmjre50/ibm-java-i386-50/jre/bin
- 4. From the directory from which you extracted the reports package, run the file in the following table depending on your operating system and version of Tivoli Common Reporting that is installed.

Operating system	File
AIX	setup_aix.bin
HP-UX	setup_hpux.bin
Linux	setup_linux.bin
Solaris	setup_solaris.bin
Windows	setup_windows.exe

Table 5. Setup files for Tivoli Common Reporting 2.1.1

Table 6. Setup files for Tivoli Common Reporting 3.1 or later

Operating system	File
AIX	setup_aix.bin
Linux	setup_linux.bin
Windows	setup_windows.exe

- 5. Select a language, and then click **OK**.
- 6. Accept the license agreement.
- 7. Select the location where the Tivoli Common Reporting server is installed (not the location where the reports are to be installed).
 - For Tivoli Common Reporting V2.1, the default path is C:\IBM\tivoli\tipv2Components\ TCRComponent or /opt/IBM/tivoli/tipv2Components/TCRComponent. The path must end with the /TCRComponent folder.
 - For Tivoli Common Reporting V3.1 or later, the default path is C:\Program Files\IBM\JazzSM\ reporting or /opt/IBM/JazzSM/reporting. The path must end with the /reporting folder.

Note: If Tivoli Common Reporting installation is distributed, reports must be installed on the dispatcher site only.

- 8. On the Welcome page, click Next
- 9. On the Choose the Installation Folder page, select the location where the Tivoli Common Reporting server is installed, and then click **Next**.
- **10**. On the Choose the reports for the installation page, select the type of report that you want to install, and then click **Next**:

• IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports

Choose the reports for the installation
 Cognos reports IBM Infrastructure Management Capacity Planner Reports for VM IBM Infrastructure Management Capacity Planner Reports for Pow IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports v IBM Tivoli Monitoring for Virtual Environments VMware Reports v7

Figure 54. Choose the reports for the installation page

- 11. Configure Cognos data sources to connect to Tivoli Data Warehouse by using one of the following choices:
 - If you have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting (from a previous installation of reports), go to the next step. To test whether you have Tivoli Data Warehouse defined, go to TCR > Launch Administration > Configuration > Data Source Connections and see whether there is an entry called TDW.
 - If you do not have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting, you must configure the data source in Tivoli Common Reporting. Configure the data source through the administration panel as described in Configuring database connection (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ttcr_config_db.html). You must enter the database alias name or the ODBC name for the database name input field.
 - a. On the Cognos Engine Configuration page, enter the user name and password for Tivoli Common Reporting, and then click **Next**.

📲 Report Installer	
	Cognos Engine Configuration
	Enter the Tivoli Common Reporting user name
	tipadmin
	Enter the Tivoli Common Reporting user password

InstallAnywhere	
Cancel Help	Previous

Figure 55. Cognos Engine Configuration page

12. Define common dimensions on the Tivoli Data Warehouse by using one of the following choices:

- If you have common dimensions (Time Dimension, Weekday Lookup, Month Lookup, and Computer System under IBM_TRAM schema) in your Tivoli Data Warehouse from a previous installation and you want to modify those dimensions to define a different time granularity, you can run the scripts manually as described in *Creating shared dimension tables and populating the time dimensions table* in the *IBM Tivoli Monitoring Administrator's Guide* V6.2.2 Fix Pack 2.
- If you do not have common dimensions in your Tivoli Data Warehouse, in the next panel, enter the JDBC credentials. Provide the database admin (db2admin, system, and so on) user name and password in the Configure data script window for JDBC User Credentials. You use the JDBC connection to run the Common Dimensions scripts on Tivoli Data Warehouse. Admin privileges are required in this step to create the IBM_TRAM schema and required tables. If you are using an Oracle database and you do not have the USERS and TEMP tablespaces in your database, you must create them in your Tivoli Data Warehouse before you can run these scripts.
- a. If you have selected any cognos reports, enter the following details, and then click Next:
 - In the Enter the database user name field, enter the administrator user name.
 - In the Enter the user password field, enter the password.
 - In the **Choose the database type** list, select the database type.
 - In the Enter the database name field, enter the warehouse database name.

📲 Report Installer	
	Cognos Data Source TDW Configuration
	Enter the database user name itmuser
	Enter the user password

	Choose the database type
	DB2
	Enter the database name
	WAREHOUS
	Skip this panel (not recommended)
InstallAnywhere Cancel Help	Previous Next

Figure 56. Cognos Data Source TDW Configuration page

13. On the Data Script runDbScript Configuration page, click the **JDBC User Credentials** tab, and enter the required details.

🚾 Report Installer	Data Script rupDbScript Configuration
	JDBC User Credentials JDBC Database Credentials Enter the database user name
InstallAnywhere Cancel Help	Skip this panel (not recommended) Previous Next

Figure 57. JDBC User Credentials tab

14. Click the **JDBC Database Credentials** tab, and select **database type**. Edit the JDBC URL, JDBC driver file names, and JDBC driver class for the selected database type and then click **Next**.

Database	Required driver file name
DB2	db2jcc.jar and db2jcc_license_cu.jar Note: JDBC credentials must have db2admin privileges.
Oracle	oraclethin.jar
SQL Server	sqljdbc.jar

On the pre-installation summary panel, all reports that are selected for installation are displayed.

Report Installer	Data Script runDbScript Configuratio
	JDBC User Credentials JDBC Database Credentials
	Choose the database type
	DB2
	Enter the database JDBC URL
	jdbc:db2://10.44.185.108:50000/WAREHOUS
	Provide the JDBC driver files separated by the semicolon
	C:\Program Files\IBM\SQLLIB\java\db2jcc_license_cu.jar
	Enter the JDBC driver class
	com.ibm.db2.jcc.DB2Driver
	Skip this panel (not recommended)
InstallAnywhere	
Cancel Help	Previous Next

Figure 58. JDBC Database Credentials tab

15. On the Report Installer Summary page, verify the reports to be installed, click **Install**, and wait for the installer to finish.

🐙 Report Installer	
InstallAnywhere	The following report sets will be installed: Cognos reports IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports v7.2 The following database scripts will be run: runDbScript
Cancel Help	Previous

Figure 59. Report Installer Summary page

Note: The **Installation results** page shows the status of all installation actions for every item or report. One log file and one trace file are included. Both files are in the user home directory, with the following names: Report_Installer_Install_mm_dd_yyyy_hh_mm_ss.log (Log) and Report_Installer_For_TCR_Output.txt (Trace)

उ Report Installer	
	The installation results: SUCCESS The database scripts: runDbScript SUCCESS Install the reports: IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports v7.2 SUCCESS Create Cognos data source: TDW
InstallAnywhere Cancel <u>H</u> elp	Previous Done

Figure 60. Installation results page

16. Click Done.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

Use the following steps to make sure that your installation was successful:

- 1. Go to Tivoli Common Reporting and see whether IBM Tivoli Monitoring for Linux KVM Reports is displayed in the Public Folders.
- 2. Go to TCR> Launch Administration> Configuration> Data Source Connections and see whether Tivoli Data Warehouse was defined. Click Tivoli Data Warehouse.
- **3**. On the next page, Tivoli Data Warehouse has a Test Connection icon next to it. Click **Test connection** to make sure that you are connected to the database.
- 4. Go to **TCR**> **Launch Query Studio**. Select **IBM Tivoli Monitoring for Linux KVM Reports**. In the navigation area, all the data is displayed.

Installing reports by using the command line

You can install the reports by using the command line.

Procedure

- 1. Run the **setup_***platform*.**exe** -**i** or **setup_***platform*.**bin** -**i** command.
- 2. Choose your installation language.
- 3. Enter the location of the TCRComponent directory.
- 4. Choose the type of reports to be installed.
- 5. Enter your Tivoli Common Reporting user name and password.
- 6. Configure your datasource and data scripts. Some report packages might not have data scripts.
- 7. An installation summary is provided, then press Enter to begin installation.

Installing reports by using the silent mode

You can install the reports by using the silent mode.

Procedure

- 1. Create the silent installer response file, and name the file as silent_installer.properties.
- 2. Run the setup_platform.exe/.bin -i silent -f path_to_response_file command.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

Use the following steps to make sure that your installation was successful:

- 1. Go to Tivoli Common Reporting and see whether IBM Tivoli Monitoring for Linux KVM Reports is displayed in the Public Folders.
- 2. Go to TCR> Launch Administration> Configuration> Data Source Connections and see whether Tivoli Data Warehouse was defined. Click Tivoli Data Warehouse.
- **3**. On the next page, Tivoli Data Warehouse has a Test Connection icon next to it. Click **Test connection** to make sure that you are connected to the database.
- 4. Go to **TCR**> **Launch Query Studio**. Select **IBM Tivoli Monitoring for Linux KVM Reports**. In the navigation area, all the data is displayed.

VMware: Installing Performance and Capacity Management Reports

You can use the Launchpad or the Reports installer to install Performance and Capacity Management Reports for VMware.

Procedure

- 1. Install the prerequisites. See "Linux KVM and VMware: Prerequisites for installing the Performance and Capacity Management Reports" on page 101
- 2. Select the type of installation. For more information on selecting the type of installation, see "Linux KVM and VMware: Types of report installation" on page 102.
- 3. You might need to point to Java 1.5+ through your system PATH. Make sure that your system PATH contains a valid path to a Java virtual machine, for example: # PATH=\$PATH:/ibmjre50/ibm-java-i386-50/jre/bin
- 4. From the directory from which you extracted the reports package, run the file in the following table depending on your operating system and version of Tivoli Common Reporting that is installed.

Operating system	File
AIX	setup_aix.bin
HP-UX	setup_hpux.bin
Linux	setup_linux.bin
Solaris	setup_solaris.bin
Windows	setup_windows.exe

Table 7. Setup files for Tivoli Common Reporting 2.1.1

Table 8. Setup files for Tivoli Common Reporting 3.1 or later

Operating system	File
AIX	setup_aix.bin
Linux	setup_linux.bin
Windows	setup_windows.exe

- 5. Select a language, and then click **OK**.
- 6. Accept the license agreement.
- 7. Select the location where the Tivoli Common Reporting server is installed (not the location where the reports are to be installed).
 - For Tivoli Common Reporting V2.1, the default path is C:\IBM\tivoli\tipv2Components\ TCRComponentor /opt/IBM/tivoli/tipv2Components/TCRComponent. The path must end with the /TCRComponent folder.
 - For Tivoli Common Reporting V3.1 or later, the default path is C:\Program Files\IBM\JazzSM\ reporting or /opt/IBM/JazzSM/reporting. The path must end with the /reporting folder.

Note: If Tivoli Common Reporting installation is distributed, reports must be installed on the dispatcher site only.

- 8. On the Choose the Installation Folder page, select the location where the Tivoli Common Reporting server is installed, and then click **Next**.
- 9. On the Choose the reports for the installation page, select the type of report that you want to install, and then click **Next**:
 - IBM Tivoli Monitoring for Virtual Environments VMware Reports

🛎 Report Installer	
	Choose the reports for the installation
	Choose the reports for the installation
	Cognos reports IBM Infrastructure Management Capacity Planner Reports for VM IBM Infrastructure Management Capacity Planner Reports for Pou IBM Tivoli Monitoring for Virtual Environments Linux KVM Reports IBM Tivoli Monitoring for Virtual Environments VMware Reports v7
	IBM Tivoli Monitoring for Virtual Environme
InstallAnywhere Cancel <u>H</u> elp	<u>Previous</u> <u>N</u> ext

Figure 61. Choose the reports for the installation page

- **10**. Configure Cognos data sources to connect to Tivoli Data Warehouse by using one of the following choices:
 - If you have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting (from a previous installation of reports), go to the next step. To test whether you have Tivoli Data Warehouse defined, go to TCR > Launch Administration > Configuration > Data Source Connections and see whether there is an entry called TDW.
 - If you do not have a Tivoli Data Warehouse connection defined in Tivoli Common Reporting, you must configure the data source in Tivoli Common Reporting. Configure the data source through the administration panel as described in Configuring database connection (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ttcr_config_db.html). You must enter the database alias name or the ODBC name for the database name input field.
 - a. On the Cognos Engine Configuration page, enter the user name and password for Tivoli Common Reporting, and then click **Next**.

📲 Report Installer	
	Cognos Engine Configuration
	Enter the Tivoli Common Reporting user name
	tipadmin
	Enter the Tivoli Common Reporting user password

InstallAnywhere	
Cancel Help	Previous Next

Figure 62. Cognos Engine Configuration page

11. Define common dimensions on the Tivoli Data Warehouse by using one of the following choices:

- If you have common dimensions (Time Dimension, Weekday Lookup, Month Lookup, and Computer System under IBM_TRAM schema) in your Tivoli Data Warehouse from a previous installation and you want to modify those dimensions to define a different time granularity, you can run the scripts manually as described in *Creating shared dimension tables and populating the time dimensions table* in the *IBM Tivoli Monitoring Administrator's Guide* V6.2.2 Fix Pack 2.
- If you do not have common dimensions in your Tivoli Data Warehouse, in the next panel, enter the JDBC credentials. Provide the database admin (db2admin, system, and so on) user name and password in the Configure data script window for JDBC User Credentials. You use the JDBC connection to run the Common Dimensions scripts on Tivoli Data Warehouse. Admin privileges are required in this step to create the IBM_TRAM schema and required tables. If you are using an Oracle database and you do not have the USERS and TEMP table spaces in your database, you must create them in your Tivoli Data Warehouse before you can run these scripts.
- a. If you have selected any cognos reports, enter the following details, and then click Next:
 - In the Enter the database user name field, type the administrator user name.
 - In the Enter the user password field, type the password.
 - In the **Choose the database type** list, select the database type.
 - In the Enter the database name field, type the warehouse database name.

📲 Report Installer	
	Cognos Data Source TDW Configuration
	Enter the database user name itmuser
	Enter the user password

	Choose the database type
	DB2
	Enter the database name
	WAREHOUS
	Skip this panel (not recommended)
InstallAnywhere Cancel Help	Previous Next

Figure 63. Cognos Data Source TDW configuration

12. On the Data Script runDbScript Configuration page, click the **JDBC User Credentials** tab, and enter the database user name and user password.

📲 Report Installer	
	Data Script runDbScript Configuration
	JDBC User Credentials JDBC Database Credentials
	Enter the database user name
	Enter the user password *******
	Skip this panel (not recommended)
InstallAnywhere Cancel Help	Previous Next

Figure 64. JDBC User Credentials tab

13. Click the **JDBC Database Credentials** tab, enter the following details: database type, database JDBC URL, JDBC driver files, and JDBC driver class. Then click **Next**.

JDBC User Credentials JDBC Database Credentials
Choose the database type
DB2
Enter the database JDBC URL
jdbc:db2://10.44.184.157:50000/WAREHOUS
Provide the JDBC driver files separated by the semicolon
C:\Program Files\IBM\SQLLIB\java\db2jcc_license_cu.jar Browse
Enter the JDBC driver class
com.ibm.db2.jcc.DB2Driver
Skip this panel (not recommended)

Figure 65. JDBC Database Credentials tab

14. On the Report Installer Summary page, verify the reports to be installed, click **Install**, and wait for the installer to finish.

CREPORT Installer		
	The following report sets will be installed: Cognos reports IBM Tivoli Monitoring for Virtual Environments VMware Reports v7.2.0.2 The following database scripts will be run: runDbScript	
Cancel Help	<u>Previous</u>	istall

Figure 66. Report installer summary page

Note: The **Installation results** page shows the status of all installation actions for every item or report. One log file and one trace file are included. Both files are in the user home directory, with the following names: Report_Installer_Install_*mm_dd_yyyy_hh_mm_ss.*log (Log) and Report_Installer_For_TCR_Output.txt (Trace)



Figure 67. Installation results page

15. Click Done.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

Installing reports by using the command line

You can install the reports by using the command line.

Procedure

- 1. Run the **setup**_*platform*.**exe** -**i** or **setup**_*platform*.**bin** -**i** command.
- 2. Choose your installation language.
- 3. Enter the location of the TCRComponent directory.
- 4. Choose the type of reports to be installed.
- 5. Enter your Tivoli Common Reporting user name and password.
- 6. Configure your datasource and data scripts. Some report packages might not have data scripts.
- 7. An installation summary is provided, then press Enter to begin installation.

Installing reports by using the silent mode

You can install the reports by using the silent mode.

Procedure

- 1. Create the silent installer response file, and name the file as silent_installer.properties.
- 2. Run the setup_platform.exe/.bin -i silent -f path_to_response_file command.

What to do next

On Windows systems in the Run window, type %USERPROFILE% to open the file explorer to the directory where the log and trace files are created. If you skipped running the database scripts or a script failed, you can run the script manually by using the instructions in Creating shared dimension tables and populating the time dimensions table in the *IBM Tivoli Monitoring Administrator's Guide V6.2.2 Fix Pack 2* or later.

Linux KVM and VMware: Uninstalling Performance and Capacity Management Reports

Uninstallation of reports is not supported by the reports installer. You can however manually delete the reports package on the Dashboard Application Services Hub for Tivoli Common Reporting version 2.1 or later, or on the Dashboard Application Services Hub for Tivoli Common Reporting version 3.1.

Procedure

- 1. Log in to the Tivoli Common Reporting interface and go to Common Reporting.
- 2. In the **Public Folders** tab of the Work with reports window, select the reports package that you want to delete, for example, IBM Tivoli Monitoring for Linux KVM Reports.

Vork	with reports	
Conne	ction	
~	Public Folders	My Folders
Public	Folders	
ublic	Folders	
Public	Folders Name ≑	
Public	Folders Name ≑ Common Reportir	ng

Figure 68. Work with reports: Select the reports package that you want to delete.

3. Click the **Delete** icon on the toolbar.

smadmin 🔗 📃		-	Q	~	6	•	2.	auni	sh 🚩	3
									1	. je
	=	₽		*	W	R	of	Ē	×	<u>□</u> / ↓ ^a _∞

Figure 69. Delete icon

Results

The selected reports package is uninstalled.

Chapter 7. PowerVM and VMware: Configuring the Capacity Planners and the Dashboard for VMware

You can configure the Capacity Planner for PowerVM, the Capacity Planner for VMware and the Dashboard for VMware. In addition, you can configure federation for the capacity planners.

Configuring a connection to the IBM Tivoli Monitoring dashboard data provider

Data providers retrieve requested information from the data sources. You must create a connection to a remote data provider to retrieve the information. For more information about configuring a connection to the IBM Tivoli Monitoring dashboard data provider, see "Creating a connection to the IBM Tivoli Monitoring dashboard data provider" in the *IBM Tivoli Monitoring Administrator's Guide*. In addition, you can review the configuration of your dashboard environment. For more information, see Chapter 3, "Preparing your dashboard environment" in the *IBM Tivoli Monitoring Administrator's Guide*.

Configuring a connection for the Capacity Planner for VMware and the Capacity Planner for PowerVM

Data providers retrieve requested information from the data sources. You must create a connection to a remote data provider to retrieve the information. For information on configuring connection to the data provider for the Capacity Planner for VMware and Capacity Planner for PowerVM, see "Creating a connection to the IBM Tivoli Monitoring dashboard data provider" in the *IBM Tivoli Monitoring Administrator's Guide*.

In addition, the data provider connection must exist with IBM Tivoli Monitoring for Virtual Environments Capacity Planner Database.

Verifying the Capacity Planner database connection details

You must verify the Capacity Planner database connection details.

About this task

If you have the Capacity Planner for PowerVM or the Capacity Planner for VMware installed, you must complete the following steps to verify the connection details for the Capacity Planner.

Procedure

- 1. On the Connections page, select the Capacity Planner database connection for which you want to verify the connection details, and then click *selection* (for verifying existing connection details).
- 2. Verify the values in the following fields:
 - Hostname:
 - Port:
 - Database Name
 - Schema
 - Username
 - Password
- 3. After you have verified the fields, click OK.

Results

The connection details for the Capacity Planners are configured.

Configuring historical data collection for the Capacity Planner for PowerVM on Linux and UNIX

Provided with the Capacity Planner installation media, are scripts to configure historical data collection and summarization and pruning for all the attribute groups that are referenced by the Capacity Planner. These scripts include information about data collection that is used for capacity planning.

About this task

The powervm_hist_config script uses the command-line interface tacmd commands to set up historical collections. When you run the script, you use parameters and commands that are shown in the procedure. The historical data collection configuration scripts are saved to the following path: \installedDashboards\com.ibm.tivoli.cppowervm\AnalyticsDatabaseInstaller\hdc_config_scripts\.

Procedure

- 1. Copy the powervm_check_hist_config.shscript to a temporary directory on the computer where the Tivoli Enterprise Monitoring Agent, Tivoli Enterprise Monitoring Server, or the Tivoli Enterprise Portal Server is installed.
- 2. Change to the directory where you copied the scripts and enter the following command: ./powervm_hist_config.sh *TEPS_Hostname username password* [ITM PATH]

Note: By default, the script uses the default Tivoli Monitoring path of /opt/IBM/ITM. To specify a different path, type the relative path after the password. If you want to modify any of the settings, edit the script. Comments at the beginning of the script describe the parameters. Each attribute group has a parameter for configuring the historical collection interval and upload times.

Example

In the following example, the -i parameter specifies that data should be uploaded to the data warehouse every hour. The -c parameter specifies that historical data should be collected every 15 minutes:

\$CANDLEHOME/bin/tacmd histcreatecollection -a "KPH SERVER DETAILS" -t kph -o "KPH SERVER DETAILS" -i 1h -l TEMA -c 15m -e "historical collection for HMC Base:KPH SERVER DETAILS"

A statement for each attribute group defines the summarization and pruning intervals. The **-d** parameter specifies the summarization intervals to configure.

In this example, D specifies Daily Summarization: \$CANDLEHOME/bin/tacmd histconfiguregroups -t kph -o "KPH SERVER DETAILS" -m -d D -p D=30d,R=3d

To configure Hourly and Daily summarization, you must specify **DH**. The **-p** flag specifies the pruning intervals.

In this example, D=30d,R=3d indicates that the Daily Summarization data is pruned after 30 days and the Raw or detailed data is not pruned after 3 days. Each pruning interval must be specified in a comma separated list with no spaces between the pruning intervals.

Configuring historical data collection for the Capacity Planner for PowerVM in Windows

You configure historical data collection for the Capacity Planner for PowerVM in Windows by using specific scripts.

Procedure

- The historical data collection configuration scripts are saved to the following path: \installedDashboards\com.ibm.tivoli.cppowervm\AnalyticsDatabaseInstaller\hdc_config_scripts\.
- 2. Copy the following scripts: powervm_hist_config.cmd to a temporary directory on the computer where the Tivoli Enterprise Monitoring Agent, Tivoli Enterprise Monitoring Server, or the Tivoli Enterprise Portal Server is installed.
- 3. Change to the directory where you copied the scripts and enter the following command: ./powervm_hist_config.cmd *TEPS_hostname username password* [ITM installation path]

Note: By default, the script uses the C:\IBM\ITM path for Tivoli Monitoring. To specify a different path, type the relative path after the password. If you want to modify any of the settings, edit the script. Comments at the beginning of the script describe the parameters. Each attribute group has one entry for configuring the historical collection interval and upload times and the summarization and pruning intervals.

Example

In the following example, the first parameter is the attribute group: //"KPH SERVER DETAILS":[" -m -d D -p D=30d,R=3d"//

The **-d** parameter specifies the summarization intervals to configure and **D** specifies Daily Summarization. To configure Hourly and Daily summarization, you specify **DH**. The **-p** flag specifies the pruning intervals. In this example, **D=30d**, **R=3d** indicates that the Daily Summarization data is pruned after 30 days and the Raw or detailed data is pruned after 3 days. Each pruning interval must be specified in a comma separated list with no spaces between the pruning intervals.

Configuring homogeneous federation for the Capacity Planner for PowerVM

To configure the homogeneous federation, you must create federation between the IBM Tivoli Data Warehouse database and the Capacity Planner database.

Note: The federation server configuration for Tivoli Data Warehouse is supported on the DB2, MSSQL Server, and Oracle databases.

Before you begin

- The Java Runtime Environment (JRE) version 1.6.0 or later must be available on the database server where you want to complete the federation steps.
- The federation steps must be completed by using a user account that has instance owner authority, permission on the database instance for which the Capacity Planner database was created, or both.

Note: Typically, the instance owner is "db2inst1" on the UNIX operating system and "db2admin" on the Windows operating system. You might observe failures if you are completing the federation steps by using a user account with non-compliant authorization levels.

• You must enable the federation.

Complete the following steps to enable the federation:

1. Keep the federated mode as On for the database manager configuration.

2. Check the parameter by connecting to the database and running the following command on the command line:

db2 get dbm cfg | grep -i federated

Note: Use this command on the Linux operating system.

If the parameter is set to no, set the parameter to yes by using the following command:

```
# db2 update dbm cfg using federated yes
```

- 3. Restart the database manager by using the following command:
 - # db2stop
 - # db2start
- For federation of Tivoli Data Warehouse on an Oracle or MSSQL server, you must install and configure IBM InfoSphere[®] Federation Server.

About this task

Use this task to create federation between the IBM Tivoli Data Warehouse database and Capacity Planner database.

Procedure

- 1. Log on to the computer where the database is installed.
- 2. Open a command line.
- **3**. Copy the AnalyticsDatabaseInstaller directory to a temporary directory. It contains the following scripts for Windows, Linux, and UNIX operating systems:
 - federation_power\fed_admin.bat
 - federation_power\fed_admin.sh

For example, copy the AnalyticsDatabaseInstaller directory from the following path: C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cppowervm\AnalyticsDatabaseInstaller to C:\temp. Update the fed_config.cfg script to include the Capacity Planner Database details and the Tivoli Data Warehouse database details.

4. Edit the fed_config.cfg file in the temporary directory, and set the following properties with appropriate values:

TDW_REMOTE_HOST_NAME

The IP address or complete host name of the database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_PORT

The port number of the DB2 database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_DB_NAME

The name of the remote Tivoli Data Warehouse database.

TDW_REMOTE_DB_SCHEMA_NAME

The schema name in the Tivoli Data Warehouse database.

TDW_REMOTE_DB_INSTANCE_OWNER

The instance owner of the remote Tivoli Data Warehouse database. The default owner is ITMUSER, who created the Tivoli Data Warehouse database on the remote Tivoli Data Warehouse database server.

TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD

The password of the remote Tivoli Data Warehouse database instance owner. This password is required when you register the federation server definition.

NODE_ID

The node name that is used to create DB2 node on database server. This NODE_ID must not be more than 8 characters in length, and must not contain special characters.

TDW_LOCAL_ALIAS

The alias name that is used to catalog the remote Tivoli Data Warehouse database on local DB2 database server.

FED_DB_SERVER_TYPE

The DB2 server data source, for example, DB2/CS.

FED_DB_VERSION

The version of the DB2 Server, for example 10.1 or 9.7.

USE_EXISTING_CATALOGED_TDW

This property must be set to Yes if you manually created an alias for remote Tivoli Data Warehouse database, and want to use the alias for federation. If the catalog is not created, the federation script attempts to create an alias for remote Tivoli Data Warehouse by using the value that is provided for TDW_LOCAL_ALIAS property. Set this property to NO if you do not want to use existing alias to be used in federation.

Note: Set **USE_EXISTING_CATALOGED_TDW** to NO only if you are sure that no alias is available on the capacity planner database with the name provided for the TDW_LOCAL_ALIAS property. Otherwise, the scripts might remove this alias after it fails to register the federation server.

CP_DB_NAME

The name of the capacity planner database, for example, TADFDCDB.

CP_DB_INSTANCE_OWNER

The instance owner of the capacity planner database. This owner is the same user that was used to create the capacity planner database.

CP_DB_INSTANCE_OWNER_PASSWORD

The password of the capacity planner database instance owner. This password is required to connect to the capacity planner database.

CP_DB_SCHEMA_NAME

The schema name in the Capacity planner database. For PowerVM, the schema name is TADFDCP. For VMware, the schema name is TADFDC.

FED_SERVER_NAME

The name of the federation server that is configured. This name might be any unique name without special characters and space characters in it. You use this federation server when you federate the Tivoli Data Warehouse tables or views.

- 5. Navigate to the temporary directory, and run any of the following commands that match your operating system:
 - fed_admin.bat: Use this script on the Windows operating system.
 - fed_admin.sh: Use this script on the UNIX and AIX operating systems.

You can run the **fed_admin.sh** or **fed_admin.bat** script in any one of the following modes:

• Non-interactive mode

In this mode, use one of the following scripts:

- fed_admin.sh
- fed_admin.bat
- fed_admin.sh JAVA_HOME -f custom_config_file
- fed_admin.bat JAVA_HOME -f custom_config_file

where *custom_config_file_name* is the file with non-default configuration parameters.

For example, the following script configures the federation server, creates all views for the physical server and the virtual machine, and prints the summary of views that were created in the console: fed_admin.sh JAVA_HOME -f fed_config.cfg

• Interactive mode

In this mode, use the one of the following scripts, which have the **-i** argument, to select the task that you want to complete:

- fed_admin.sh JAVA_HOME -i -f custom_config_file
- fed_admin.bat JAVA_HOME -i -f custom_config_file

Where:

- -i Runs the script in interactive mode. For example, you can configure a new server; update user mapping; add views; and delete views. If you choose to add all views, a summary of the views that were created is displayed.
- **-f** *custom_config_file_name*

custom_config_file_name is the file with non-default configuration parameters.

Note: When this file is not used, the script uses the default fed_config.cfg file. To get the correct result from the scripts in any mode, you must specify correct input parameters in the fed_config.cfg file.

• DBA verification mode

In this mode, use the **fed_admin.sh** or **fed_admin.bat** script with the **-v** argument to print the federation commands for the task in a file instead of running the commands against the database.

- Change the directory to the directory that contains the federation.
- On the command line, run the **fed_admin** script in interactive mode and DBA verification mode:
 - ./fed_admin.sh JAVA_HOME -i -v -f custom_config_file
 - When you are prompted for input, select one of the following options:
 - Configure new server
 - Update user mapping
 - Add views
 - Delete views
 - If you select Add views or Delete views, type the view names to be added or deleted separated by a space, for example, VM_CPU_D VM_CPU_M. For all views, type the argument all.

Note: View names must be typed in uppercase.

- On the command line, run the fed_admin script only in DBA verification mode to print the queries that are to be run against the database for the following options:
 - Configure the federation server
 - Add all views for the physical server and the virtual machine: ./fed_admin.sh JAVA_HOME -v -f custom_config_file

The federation commands are printed in the federationqueries.ddl file.

6. Configure a new federation server by running the fed_admin.sh or fed_admin.bat script.

Important: Ensure that you set all input parameters in the fed_config.cfg file.

- If you run the **fed_admin.sh** or **fed_admin.bat**, or the **fed_admin.sh** -**f** or **fed_admin.bat** -**f** *custom_config_file* script, the script configures the federation server and creates all views for the physical server and virtual machines.
- If you run the **fed_admin.sh -i -f** or **fed_admin.bat -i -f** *custom_config_file* script, type 1. This script configures a federation server that is based on the configuration parameters that are specified in the fed_config.cfg file.

- 7. If the credentials of the Tivoli Data Warehouse instance owner are changed on the remote database server after the federation server is configured, you must update the existing federation server details.
 - a. Set the **TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD** property in the configuration file to a new password. The default configuration file is fed_config.cfg.
 - b. Run the **fed_admin.sh** or **fed_admin.bat** script.
 - If you run the **fed_admin.sh** or **fed_admin.bat**, or the **fed_admin.sh** -**f** or **fed_admin.bat** -**f** *custom_config_file* script, the script updates the federation server details.
 - If you run the **fed_admin.sh** -i or **fed_admin.bat** -i script, type 2. This script uses the configuration parameters that are specified in the fed_config.cfg file to update the federation server details.
- 8. Add the selected views.
 - a. Run the fed_admin.sh -i or fed_admin.bat -i script.
 - b. Type 3.
 - c. Type the names of views that you want to add separated by spaces, for example, VM_CPU_H VM_NET_M.

Note: You must type the view names in uppercase. If you want to add all views, type the argument all.

- 9. Delete the selected views.
 - a. Run the fed_admin.sh -i or fed_admin.bat -i script.
 - b. Type 4.
 - c. Type the names of views you want to delete separated by spaces, for example, VM_CPU_H VM_NET_M.

Note: You must type the view names in uppercase. If you want to delete all views, type the argument all.

Configuring homogeneous federation for the Capacity Planner for VMware

To configure the homogeneous federation, you must create federation between the IBM Tivoli Data Warehouse database and the Capacity Planner database.

Note: The federation server configuration for Tivoli Data Warehouse is supported on the DB2, MSSQL Server, and Oracle databases.

Before you begin

- The Java Runtime Environment (JRE) version 1.6.0 or later must be available on the database server where you want to complete the federation steps.
- The federation steps must be completed by using a user account that has instance owner authority and/or permission on the database instance for which the Capacity Planner database was created.

Note: Typically, the instance owner is "db2inst1" on the UNIX systems and "db2admin" on the Windows systems. You might observe failures if you are completing the federation steps by using a user account with non-compliant authorization levels.

• You must enable the federation.

Complete the following steps to enable the federation:

- 1. Keep the federated mode as On for the database manager configuration.
- **2**. Check the parameter by connecting to the database, running the following command at the command line:

```
# db2 get dbm cfg | grep -i federated
```

Note: Use this command on the Linux systems.

If the parameter is set to no, set the parameter to yes by using the following command:

db2 update dbm cfg using federated yes

- 3. Restart the database manager by using the following command:
 - # db2stop
 - # db2start
- For federation of Tivoli Data Warehouse on Oracle or MSSQL server, you must install and configure IBM InfoSphere Federation Server.

About this task

This task provides information about creating federation between the IBM Tivoli Data Warehouse database and Capacity Planner database.

Procedure

- 1. Log on to the computer where the database is installed.
- 2. Open a command line.
- 3. Copy the federation directory from *dash_home*\installedDashboards\com.ibm.tivoli.cpdash\ AnalyticsDatabaseInstaller to a temporary directory. The federation scripts are available in the federation directory.

After installation, the scripts are placed in the following directory for the Capacity Planner for VMware: dash_home\installedDashboards\com.ibm.tivoli.cpdash\AnalyticsDatabaseInstaller where dash_home has the following path on the Windows systems: \Program Files\IBM\JazzSM or this path: /opt/IBM/JazzSM on the UNIX systems.

Note: The directory where you copy these scripts must have write permissions for the database instance owner. For the Windows systems, the target directory should not have any spaces in the path, otherwise the federation scripts fail.

4. Edit the fed_config.cfg file in the temporary directory, and set the following properties with appropriate values:

TDW_REMOTE_HOST_NAME

The IP address or complete host name of the database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_PORT

The port number of the DB2 database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_DB_NAME

The name of the remote Tivoli Data Warehouse database.

TDW_REMOTE_DB_SCHEMA_NAME

The schema name in the Tivoli Data Warehouse database.

TDW_REMOTE_DB_INSTANCE_OWNER

The instance owner of the remote Tivoli Data Warehouse database. The default owner is ITMUSER, who created the Tivoli Data Warehouse database on the remote Tivoli Data Warehouse database server.

TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD

The password of the remote Tivoli Data Warehouse database instance owner. This password is required when you register the federation server definition.

NODE_ID

The node name that is used to create the DB2 node on the database server. This NODE_ID must not be more than 8 characters in length, and must not contain special characters.

TDW_LOCAL_ALIAS

The alias name that is used to catalog the remote Tivoli Data Warehouse database on the local DB2 database server.

FED_DB_SERVER_TYPE

The DB2 server data source, for example, DB2/CS.

FED_DB_VERSION

The version of the DB2 Server, for example 10.1 or 9.7.

USE_EXISTING_CATALOGED_TDW

This property must be set to Yes if you manually created an alias for the remote Tivoli Data Warehouse database and want to use the alias for federation. If the catalog is not created, the federation script attempts to create an alias for the remote Tivoli Data Warehouse by using the value that is provided for the TDW_LOCAL_ALIAS property. Set this property to NO if you do not want to use existing alias to be used in federation.

Note: Set **USE_EXISTING_CATALOGED_TDW** to NO only if you are sure that no alias is available on the capacity planner database with the name provided for the TDW_LOCAL_ALIAS property. Otherwise, the scripts might remove this alias after it fails to register the federation server.

CP_DB_NAME

The name of the capacity planner database, for example, TADFDCDB.

CP_DB_INSTANCE_OWNER

The instance owner of the capacity planner database. This owner is the same user that was used to create the capacity planner database.

CP_DB_INSTANCE_OWNER_PASSWORD

The password of the capacity planner database instance owner. This password is required to connect to the capacity planner database.

CP_DB_SCHEMA_NAME

The schema name in the Capacity planner database. For PowerVM, the schema name is TADFDCP. For VMware, the schema name is TADFDC.

FED_SERVER_NAME

The name of the federation server that is configured. This name might be any unique name without special characters and space characters in it. You use this federation server when you federate the Tivoli Data Warehouse tables or views.

- 5. Navigate to the temporary directory, and run any of the following commands that matches your operating system:
 - **fed_admin.bat**: Use this script on the Windows systems.

• **fed_admin.sh**: Use this script on the UNIX and AIX systems.

You can run the **fed_admin.sh** or **fed_admin.bat** script in any one of the following two modes:

- Non-interactive mode
 - In this mode, use one of the following scripts:
 - fed_admin.sh
 - fed_admin.bat
 - fed_admin.sh JAVA_HOME -f custom_config_file
 - fed_admin.bat JAVA_HOME -f custom_config_file

where *custom_config_file_name* is the file with non-default configuration parameters.

For example, the following script configures the federation server and creates all views for the physical server and the virtual machine and prints the summary of views created in the console:

fed_admin.sh JAVA_HOME -f fed_config.cfg

Interactive mode

In this mode, use the one of the following scripts, which have the **-i** argument to select the task that you want to complete:

- fed_admin.sh JAVA_HOME -i -f custom_config_file

- fed_admin.bat JAVA_HOME -i -f custom_config_file

Where:

- -i Runs the script in interactive mode. For example, you can configure a new server; update user mapping; add views; and delete views. If you choose to add all views, a summary of the views that were created is displayed.
- **-f** custom_config_file_name

custom_config_file_name is the file with non-default configuration parameters.

Note: When this file is not used, the script uses the default fed_config.cfg file. To get the correct result from the scripts in any mode, you must specify correct input parameters in the fed_config.cfg file.

• DBA verification mode

In this mode, use the **fed_admin.sh** or **fed_admin.bat** script with the **-v** argument to print the federation commands for the task in a file instead of running the commands against the database.

- Change the directory to the directory that contains the federation.
- At the command line, run the **fed_admin** script in interactive mode and DBA verification mode:
 - ./fed_admin.sh JAVA_HOME -i -v -f custom_config_file
 - When you are prompted for input, select one of the following options:
 - Configure new server
 - Update user mapping
 - Add views
 - Delete views
 - If you select Add views or Delete views, type the view names to be added or deleted separated by a space, for example, VM_CPU_D VM_CPU_M. For all views, type the argument all.

Note: View names must be typed in uppercase.

- At the command line, run the **fed_admin** script only in DBA verification mode to print the queries that are to be run against the database for the following options:
 - Configure federation server
 - Add all views for physical server and virtual machine: ./fed_admin.sh JAVA_HOME -v -f custom_config_file

The federation commands are printed in the federationqueries.ddl file.

6. Configure new federation server by running the fed_admin.sh OR fed_admin.bat script.

Note: Ensure that you have set all the input parameters in the fed_config.cfg file.

- If you run the **fed_admin.sh** or **fed_admin.bat**, or the **fed_admin.sh** -**f** or **fed_admin.bat** -**f** *custom_config_file* script, the script configures the federation server and creates all views for the physical server and virtual machines.
- If you run the **fed_admin.sh -i -f** or **fed_admin.bat -i -f** *custom_config_file* script, type 1. This script configures a federation server that is based on the configuration parameters that are specified in the fed_config.cfg file.
- 7. If the credentials of Tivoli Data Warehouse instance owner are changed on the remote database server after the federation server is configured, you must update the existing federation server details.
 - a. Set TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD property in configuration file to new password. The default configuration file is fed_config.cfg.

- b. Run the **fed_admin.sh/bat** script.
 - If you run the fed_admin.sh or fed_admin.bat, or the fed_admin.sh -f or fed_admin.bat -f *custom_config_file* script, the script updates the federation server details.
 - If you run the fed_admin.sh -i or fed_admin.bat -i script, type 2.
 This script uses the configuration parameters that are specified in the fed_config.cfg file to update the federation server details.
- 8. Add the selected views.
 - a. Run the fed_admin.sh -i or fed_admin.bat -i script.
 - b. Type 3.
 - c. Type space separated view names that you want to add. For example, VM_CPU_H VM_NET_M.

Note: You must type the view names in uppercase. If you want to add all the views, type the argument **all**.

- 9. Delete the selected views.
 - a. Run the fed_admin.sh -i or fed_admin.bat -i script.
 - b. Type 4.
 - c. Type space separated view names that you want to delete. For example, VM_CPU_H VM_NET_M.

Note: You must type the view names in uppercase. If you want to delete all the views, type the argument **all**.

Verifying the Capacity Planner federation

After you configure the Capacity Planner federation, you can verify the federation by using the summary of created federated views and log files.

Verifying the Capacity Planner federation by using the summary of created federation views

Summary for Capacity Planner for PowerVM

After you configure the Capacity Planner federation, the script prints the summary of created federated views. If all the views are successfully created, the following summary is displayed:

[INFO] [main] 2013/07/05 07:13:13 - For CPU [INFO] [main] 2013/07/05 07:13:13 - Success: You are configured to use Hourly, Daily, Weekly, Monthly, Quarterly, Yearly data in the Capacity planning tool [INFO] [main] 2013/07/05 07:13:13 [INFO] [main] 2013/07/05 07:13:13 - For MEMORY - Success: You are configured to use Hourly, Daily, Weekly, Monthly, Quarterly, Yearly data in the Capacity planning tool [INFO] [main] 2013/07/05 07:13:13 - For STORAGE [INF0] [main] 2013/07/05 07:13:13 - Success: You are configured to use Hourly, Daily, Weekly, Monthly, Quarterly, Yearly data in the Capacity planning tool [INFO] [main] 2013/07/05 07:13:13 - For NETWORK - Success: You are configured to use Hourly, Daily, Weekly, Monthly, [INFO] [main] 2013/07/05 07:13:13 Quarterly, Yearly data in the Capacity planning tool [INF0] [main] 2013/07/05 07:13:13 - ANA01009I : Successfully added view all. [INFO] [main] 2013/07/05 07:13:13 - ANA01006I : Successfully created views.

This summary confirms that the views are successfully created for CPU, MEMORY, STORAGE, and NETWORK metrics for Hourly, Daily, Weekly, Monthly, Quarterly, and Yearly summarized utilization data.

If all or any of the federation views are not created, the script displays the attributes groups that must be enabled for Historical Data Collection in IBM Tivoli Monitoring:

[INFO] [main] [INFO] [main] in history col [INFO] [main] [INFO] [main]	2013/07/24 06:09:57 2013/07/24 06:09:57 lection configuration 2013/07/24 06:09:57 2013/07/24 06:09:57	- - for - -	For CPU Failure : Please enable at least one summarization interval the following attribute groups in ITM KPH_SERVER_LPARS
[INFO] [main] [INFO] [main] in history col [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main]	2013/07/24 06:09:57 2013/07/24 06:09:57 lection configuration 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57	- for - - - - - -	For MEMORY Failure : Please enable at least one summarization interval the following attribute groups in ITM Unix_Memory Machine_Information AIX_LPAR KVH_SERVER_LPARS KVA_PHYSICAL_MEMORY KVA_STORAGE_MAPPINGS KVA_TADDM
[INFO] [main] [INFO] [main] in history col [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main]	2013/07/24 06:09:57 2013/07/24 06:09:57 lection configuration 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57	- for - - - -	For STORAGE Failure : Please enable at least one summarization interval the following attribute groups in ITM Disk_Performance Machine_Information AIX_LPAR KPH_SERVER_LPARS
[INFO] [main] [INFO] [main] in history col [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main] [INFO] [main]	2013/07/24 06:09:57 2013/07/24 06:09:57 lection configuration 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57 2013/07/24 06:09:57	- for - - -	For NETWORK Failure : Please enable at least one summarization interval the following attribute groups in ITM Network Machine_Information AIX_LPAR KPH_SERVER_LPARS

In this case, you must enable the summarization, and run the federation again. For more information about enabling the summarization, see "Configuring historical data collection for the Capacity Planner for PowerVM in Windows" on page 123.

Summary for Capacity Planner for VMware

<pre>[INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34 in history collection configuration [INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34</pre>	 For CPU Failure : Please enable at least one summarization interval for the following attribute groups in ITM KVM_VIRTUAL_MACHINES -
[INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34 in history collection configuration [INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34	 For MEMORY Failure : Please enable at least one summarization interval for the following attribute groups in ITM KVM_VM_MEMORY -
[INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34 in history collection configuration [INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34	 For STORAGE Failure : Please enable at least one summarization interval for the following attribute groups in ITM KVM_VM_DATASTORE_UTILIZATION
[INFO] [main] 2013/07/26 06:31:34 [INFO] [main] 2013/07/26 06:31:34 in history collection configuration [INFO] [main] 2013/07/26 06:31:34	 For NETWORK Failure : Please enable at least one summarization interval for the following attribute groups in ITM KVM_VM_NETWORK
[INFO] [main] 2013/07/26 06:31:34 [ERROR] [main] 2013/07/26 06:31:34 - CTGCP1016E : Failed to create view . See
federation.log file for details.

Note: If the **fed_admin.sh** or **fed_admin.bat** script is used in interactive mode, the print summary works only for the option to add all views.

Verifying the Capacity Planner federation by using the log files Procedure

- 1. Log in to the computer system with the database server as the database instance owner, and then complete the federation. The federation script attempts to federate all possible warehouse tables. If the summarization interval is not defined, the federation script reports an error.
- 2. Review the federation.log file, and confirm that at least one summarization interval was successful for each attribute group.

The federation.log file is in the following location:

- On Windows systems: C:\Program Files\IBM\JazzSM\installedDashboards\ com.ibm.tivoli.cpdash\AnalyticsDatabaseInstaller\federation\federation.log
- On operating systems other than Windows: /opt/IBM/JazzSM/installedDashboards/ com.ibm.tivoli.cpdash/AnalyticsDatabaseInstaller/federation/federation.log

If the Tivoli Data Warehouse database alias and catalog entry were manually created before the fed_admin.bat or fed_admin.sh script is run, the federation.log file contains error messages that you can ignore.

- **3.** Optional: If the fed_admin.bat or fed_admin.sh script fails, you must drop the "server" and "wrapper" for the database before you run the script for federation again. For example, run the following DB2 commands:
 - db2 connect to tadfdcdb
 - db2 drop server whserver
 - db2 drop wrapper drda

Where

```
tadfdcdb
```

The name of the Capacity Planner database.

```
whserver
```

The server.

```
drda The wrapper.
```

Note: If the fed_admin.bat or fed_admin.sh script indicates that federation was successful but the last set of messages in the federation.log file indicate that the federation failed, check the log entries for each federation step. If each step was successful, you can ignore the error message that is at the end of the federation.log file.

Configuring connection for the Dashboard for VMware

Data providers retrieve requested information from the data sources. You must create a connection to a remote data provider to retrieve the information. For information about configuring connection to the data provider for the Dashboard for VMware, see "Creating a connection to the IBM Tivoli Monitoring dashboard data provider" in the *IBM Tivoli Monitoring Administrator's Guide*.

Configuring Dashboard for VMware to communicate with Tivoli Application Dependency Discovery Manager

If your environment contains a Tivoli Application Dependency Discovery Manager server, the Dashboard for VMware uses information from the Tivoli Application Dependency Discovery Manager server to display the change history details for selected resources. To configure the Dashboard for VMware to communicate with Tivoli Application Dependency Discovery Manager, complete the following steps:

Procedure

- 1. Log on to IBM Dashboard Application Services Hub by using an administrator user ID.
- 2. Select Console Settings > Connections.
- 3. In the Connection table, select the VMware Dashboard row, click Z Edit existing provider, and then enter the Tivoli Application Dependency Discovery Manager connection details in the following fields:
 - TADDM Hostname
 - TADDM Port
 - TADDM Username
 - TADDM Password
- 4. Click **OK** to save the connection configuration.

You can configure Tivoli Application Dependency Discovery Manager to run the VMware sensors. In addition, you can configure the VMware workflow policies to forward VMware virtual machine updates to Tivoli Application Dependency Discovery Manager so that Tivoli Application Dependency Discovery Manager is up-to-date.

Note: When you configure a connection to Tivoli Application Dependency Discovery Manager, specify the host name or IP address of the Tivoli Application Dependency Discovery Manager server that is listed for each of the following deployment types:

- Domain server deployment: Domain server
- Streaming server deployment: Primary storage server
- Synchronization server deployment: Synchronization server (also known as the enterprise domain server)

To configure the VMware workflow policy to forward updates to Tivoli Application Dependency Discovery Manager, see the following topics:

- "Predefined policies" in the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI User's Guide
- "Tivoli Application Dependency Discovery Manager policies" in the *Tivoli Enterprise Portal User's Guide*

For more information about generating and loading IdML book on Tivoli Application Dependency Discovery Manager, see Using TADDM to Determine Monitoring Coverage for IBM Tivoli Monitoring (https://www.ibm.com/developerworks/mydeveloperworks/wikis/home/wiki/Tivoli Application Dependency Discovery Manager/page/Using TADDM to Determine Monitoring Coverage for IBM Tivoli Monitoring?lang=en).

For more information about enabling the Tivoli Application Dependency Discovery Manager policies, see Tivoli Application Dependency Discovery Manager policies (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=%2Fcom.ibm.itm.doc_6.2.3%2Fpolicy_taddm_intro_tep.htm).

Configuring historical data collection for the Performance and Capacity Management Reports for VMware

Provided with the Dashboard for VMware, Capacity Reporting, and Capacity Planner installation media are scripts to configure historical data collection and summarization and pruning for all the attribute groups that are referenced by the Tivoli Common Reporting Performance and Capacity Management Reports.

These scripts include best practice recommendations for data collection that is used for reporting and capacity planning.

About this task

The vmware_hist_config script uses the command-line interface **tacmd** commands to set up the historical collections. When executing the script, you pass in some parameters. The commands and parameters are listed in the procedure.

Procedure

- On operating systems other than Windows
 - Locate the following scripts on the installation media and copy to a temporary directory on the computer where a Tivoli Enterprise Monitoring Agent, a Tivoli Enterprise Monitoring Server, or Tivoli Enterprise Portal Server is installed: vmware_hist_config.sh and vmware_check_hist_config.sh.
 - 2. Change to the directory where you copied the scripts and enter the following command. ./vmware_hist_config.sh *TEPS_Hostname username password* [ITM path]

By default, the script uses the default Tivoli Monitoring path of /opt/IBM/ITM. To specify a different path, type the relative path after the password. If you want to modify any of the settings, edit the script. Comments at the beginning of the script describe the parameters. Each attribute group has a parameter for configuring the historical collection interval and upload times.

In the following example, the -i parameter specifies that data must be uploaded to the data warehouse every hour. The -c parameter specifies that historical data must be collected every 15 minutes.

```
$CANDLEHOME/bin/tacmd histcreatecollection -a "KVM_DISTRIBUTED_VIRTUAL_SWITCHES"
-t "vm" -o "KVM DISTRIBUTED VIRTUAL SWITCHES" -i 1h -l TEMA -c 15m -e
"historical collection for Distributed Virtual Switches"
```

A statement for each attribute group defines the summarization and pruning intervals. The -d parameter specifies which summarization intervals to configure. In the following example, D specifies Daily Summarization. (To configure Hourly and Daily summarization, specify DH.) The -p flag specifies the pruning intervals. In the example, D=30d,R=3d indicates that the Daily Summarization data will be pruned after 30 days and the Raw (detailed) data will be pruned after 3 days. Each pruning interval must be specified in a comma-separated list with no spaces between the pruning intervals.

\$CANDLEHOME/bin/tacmd histconfiguregroups -t "vm" -o "KVM DISTRIBUTED VIRTUAL SWITCHES" -m -d D -p D=30d,R=3d

- On Windows systems
 - Locate the following script on the installation media and copy to a temporary directory on the computer where a Tivoli Enterprise Monitoring Agent, a Tivoli Enterprise Monitoring Server, or Tivoli Enterprise Portal Server is installed: vmware_hist_config.cmd
 - 2. Change to the directory where you copied the script and enter the following command: vmware_hist_config.cmd *TEPS_Hostname username password* [ITM path]

By default, the script uses the C:\IBM\ITM path for Tivoli Monitoring. To specify a different path, type the relative path after the password. If you want to modify any of the settings, edit the script. Comments at the beginning of the script describe the parameters. Each attribute group has one entry for configuring the historical collection interval and upload times and the summarization and pruning intervals.

In the following example, the first parameter is the attribute group. The -d parameter specifies which summarization intervals to configure. In the following example, D specifies Daily Summarization. (To configure Hourly and Daily summarization, specify DH.) The -p flag specifies the pruning intervals. In the example, D=30d,R=3d indicates that the Daily Summarization data will be pruned after 30 days and the Raw (detailed) data will be pruned after 3 days. Each pruning interval must be specified in a comma-separated list with no spaces between the pruning intervals.

//"KVM DISTRIBUTED VIRTUAL PORTGROUPS":[" -m -d D -p D=30d,R=3d", "*VMWARE_VI_AGENT"],//

Heterogeneous federation configuration for Oracle and MSSQL Server

You use the IBM InfoSphere Federation Server to federate databases to run a database query that is called a Federated Query. You use the federated query to work with objects, for example, tables or views in different relational database systems.

To create a federated system, you must install the federated engine. Then, you must configure the federated engine to communicate with the data sources. The basic federated objects include:

- The federated server that communicates with the data sources by using software modules that are called wrappers.
 - Each data source must be identified to the system as a server. If the data source require authentication, the remote authentication information must be registered with the federated system as user mappings.
- Remote data sets that you access by using nicknames to the federated system.
 - You must reference the nickname as a local table in your application.

Configuring heterogeneous federation prerequisites for Oracle and MSSQL Server

You must configure the prerequisites for heterogeneous federation for the Oracle and MSSQL Server.

Procedure

- 1. Install the Database Client software for the corresponding Oracle or MSSQL Server database that is installed on the Capacity planner database server.
- 2. Install IBM InfoSphere Federation Server, version 9.7 Fix Pack 7 or later. You use an existing copy of a compatible DB2 database system or you use the IBM InfoSphere Federation Server installation wizard to install a new copy of the bundled DB2 Enterprise Server Edition. IBM InfoSphere Federation Server requires a DB2 database for the Linux, UNIX, and Windows systems.
- **3.** If the data source client software is required, install, configure, and verify the client software for each federated data source that you want to access before you install the wrappers. The ODBC driver must be installed and configured on the federated server.
- 4. Ensure that the user account has authorization levels that are compliant or you might observe failures when you complete the federation steps. You can use one of the following instance owners:
 - On Windows systems, use db2admin
 - On UNIX systems, use db2inst1

Registering the service or node for the MSSQL server on Windows

You must register the service or node for the MSSQL server on the Windows systems so that the data source information is updated in Windows systems.

Procedure

1. Set ODBC Connections.

Note: Ensure that ODBC connections are already set on the computer where the SQL Server is installed.

- 2. Select Administrative Tools > Data Source (ODBC) to open ODBC Data Source Administrator.
- 3. Click the System DSN tab.
- 4. Click Add, select SQL Server, and then click Finish.
- 5. Specify the Data source Name and SQL Server IP or Host name, and then click Next.

Note: Data source name must not contain spaces or special characters.

- 6. Select the With SQL Server authentication using login id and password entered by user check box.
- 7. Select the Connection to SQL Server to obtain default settings for the additional options check box.
- 8. Enter the credentials of SQL Server, and then click Next.
- 9. Click **Finish** and verify the connection.

Registering the service or node for the MSSQL Server on UNIX

You must register the service or node for the MSSQL Server by using database client software on UNIX systems so that the data source information is updated in UNIX systems.

Procedure

- 1. Install and configure the Data Direct ODBC drivers.
- 2. After you install the ODBC driver, configure the odbc.ini system information file as shown in the following example:

```
[ODBC Data Sources]
WAREHOUS=MS SQL Server 2000
[WAREHOUS]
Driver=/opt/odbc/lib/ddmsss20.so
Description=MS SQL Server Driver for AIX
Address=9.112.98.123,1433
```

- 3. Test the connection with SQL server.
- 4. Edit the db2dj.ini file and set the environment variables as shown in the following examples:
 - ODBCINI

This environment variable provides the path to the odbc.ini file. ODBCINI=/opt/IBM/odbcdriver/branded odbc/IBM Tools/odbc.ini

• DJX_ODBC_LIBRARY_PATH

DJX_ODBC_LIBRARY_PATH = /opt/IBM/odbcdriver/branded_odbc/lib/

- 5. Stop and start DB2 by using the following commands:
 - db2stop;
 - db2start;

Registering the federation server definition for MSSQL

You must register the federation server definition for MSSQL and add links for all the required Tivoli Data Warehouse tables. Then, the **fed_config.cfg** script and the **fed_admin.sh** or **fed_admin.bat** script that is running are updated.

Procedure

- 1. After you register the service or node by using the client software, open a command line. For more information about registering the service or node on Windows and UNIX systems, see "Registering the service or node for the MSSQL server on Windows" on page 136 and "Registering the service or node for the MSSQL Server on UNIX."
- 2. Copy the AnalyticsDatabaseInstaller folder from the installed location to a temporary directory. The path of the AnalyticsDatabaseInstaller folder for PowerVM:

- On Windows systems, copy the AnalyticsDatabaseInstaller directory from the C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cppowervm\AnalyticsDatabaseInstaller directory to the c:\tmp directory.
- On operating systems other than Windows, copy the AnalyticsDatabaseInstaller directory from the opt/IBM/JazzSM/InstalledDashboards/com.ibm.tivoli.cppowervm/AnalyticsDatabaseInstaller directory to the /tmp directory.

The path of the AnalyticsDatabaseInstaller folder for VMware:

- On Windows systems: DASH_HOME\installedDashboards\com.ibm.tivoli.cpdash For example, C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cpdash
- On operating systems other than Windows: DASH_HOME/installedDashboards/ com.ibm.tivoli.cpdash

For example, /opt/IBM/JazzSM/installedDashboards/com.ibm.tivoli.cpdash/

The temporary directory must not include space characters in the file name.

Note: After you copied the AnalyticsDatabaseInstaller directory to a temp directory, the fed_admin.sh or fed_admin.bat and fed_config.cfg files are available in the AnalyticsDatabaseInstaller/federation directory.

- 3. Change the working directory to the directory where the fed_admin.sh or fed_admin.bat federation scripts files are copied.
 - The PowerVM federation scripts are stored in the federation_power directory.
 - The VMware federation scripts are stored in the federation directory.
- 4. Edit the fed_config.cfg file in the same directory and set the following properties with appropriate values:

TDW_REMOTE_HOST_NAME

The IP address or complete host name of the database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_PORT

The port number of the DB2 database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_DB_NAME

The name of the remote Tivoli Data Warehouse database, for example, CP_DB_NAME. The name of the capacity planner database, for example, TADFDCDB.

TDW_REMOTE_DB_SCHEMA_NAME

The schema name in the Tivoli Data Warehouse database.

TDW_REMOTE_DB_INSTANCE_OWNER

The instance owner of the remote Tivoli Data Warehouse database. The default owner is ITMUSER, who created the Tivoli Data Warehouse database on the remote Tivoli Data Warehouse database server.

TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD

The password of the remote Tivoli Data Warehouse database instance owner. This password is required when you register the federation server definition.

NODE_ID

The node name that is configured in the odbc.ini file or the data source name that you configured. The NODE_ID must not be more than eight characters in length, and must not contain special characters.

TDW_LOCAL_ALIAS

The alias name that is used to catalog the remote Tivoli Data Warehouse database on local DB2 database server.

FED_DB_SERVER_TYPE

The MSSQL server data source, for example, MSSQL.

FED_DB_VERSION

The version of the MSSQL Server, for example 2000 or 2008.

USE_EXISTING_CATALOGED_TDW

This property must be set to Yes if you manually created an alias for remote Tivoli Data Warehouse database, and want to use the alias for federation. If the catalog is not created, the federation script attempts to create an alias for remote Tivoli Data Warehouse by using the value that is provided for TDW_LOCAL_ALIAS property. Set this property to NO if you do not want to use existing alias to be used in federation.

Note: Set **USE_EXISTING_CATALOGED_TDW** to NO only if you are sure that no alias is available on the capacity planner database with the name provided for the TDW_LOCAL_ALIAS property. Otherwise, the scripts might remove this alias after it fails to register the federation server.

CP_DB_NAME

The name of the capacity planner database, for example, TADFDCDB.

CP_DB_INSTANCE_OWNER

The instance owner of the capacity planner database. This owner is the same user that was used to create the capacity planner database.

CP_DB_INSTANCE_OWNER_PASSWORD

The password of the capacity planner database instance owner. This password is required to connect to the capacity planner database.

CP_DB_SCHEMA_NAME

The schema name in the Capacity planner database. For PowerVM, the schema name is TADFDCP.

FED_SERVER_NAME

The name of the federation server that is configured. This name might be any unique name without special characters and space characters in it. You use this federation server when you federate the Tivoli Data Warehouse tables or views.

5. After you update the fed_config.cfg file, run the **fed_amdin.sh** or **fed_amdin.bat** script from the same directory with the appropriate parameter.

Registering the service or node for the Oracle Server on Windows

You must register the service or node for the Oracle Server on Windows systems by using Oracle software. Then, by using Net Manager, the tnsnames.ora file is updated to include the Oracle Tivoli Data Warehouse server information.

Procedure

- 1. Using the oracle client software, run Net Manager.
- 2. Right-click Service-Naming and click new service.
- **3**. On the Service page of the **Net Service Name** wizard, type the name of the service to connect to the remote Tivoli Data Warehouse, and then click **Next**.
- 4. On the Protocol page of the Net Service Name wizard, from the network protocol list, select **TCP/IP** (Internet Protocol), and then click Next.
- 5. On the Protocol Settings page of the Net Service Name wizard, type the host name or the IP address of the database server where the Tivoli Data warehouse is installed on the Oracle database. Also, type the port number. The default port number is 1521.
- 6. On the Service page of the Net Service Name wizard, type TDW as the Service Name and from the **Connection Type** list, select **Database Default**.

7. In the Connection Test window click **Test**. In the Change Login window, type the user name and the password of the person who connects to the Tivoli Data Warehouse on the Oracle database. Test the service connection. Then, save the settings and exit.

Registering the service or node for the Oracle Server on UNIX

You must register the service or node for the Oracle Server on UNIX systems by using Oracle software. Then, by using Net Manager, the tnsnames.ora file is updated to include the Oracle Tivoli Data Warehouse server information in UNIX systems.

Procedure

- 1. Install and configure the Oracle client software on the server where you installed the Capacity Planner database.
- 2. Install the IBM federation server.
- **3**. Edit the tnsnames.ora file with the **Tivoli Data Warehouse** connection details as shown in the following example:

```
tdw_node =
  (DESCRIPTION =
  (ADDRESS_LIST =
  (ADDRESS = (PROTOCOL = TCP)(HOST = ip_address)(PORT = 1521)))
  (CONNECT_DATA = (SERVICE_NAME = warehous)))
```

4. Edit the db2dj.ini file and set the environment variables as shown in the following example:

ORACLE_HOME

Path to the oracle client software installation directory, for example, **ORACLE_HOME**=\usr\oracle\ 8.1.7

TNS_ADMIN

Path to where the tnsnames.ora file is stored, for example, **TNS_ADMIN**=\$ORACLE_HOME/NETWORK/ ADMIN

DB2LIBPATH

```
ORACLE_HOME/lib folder, for example, DB2LIBPATH=$ORACLE_HOME/lib
```

5. Add the following environment variables to the bash_profile file of the db2inst1 user:

```
ORACLE_HOME=/home/db2inst1/app/db2inst1/product/11.2.0/client_1/
export ORACLE_HOME
DB2LIBPATH=/home/db2inst1/app/db2inst1/product/11.2.0/client_1/lib
export DB2LIBPATH
PATH=$PATH:$HOME/bin
export PATH=$ORACLE_HOME/bin:$PATH
LD_LIBRARY_PATH=/home/db2inst1/app/db2inst1/product/11.2.0/client_1/lib:$LD_LIBRARY_PATH
export LD_LIBRARY_PATH
```

6. Stop and start DB2 by using the following commands: db2stop; and db2start;

Registering the server definition for Oracle

You must register the federation server definition for Oracle and add links for all the required Tivoli Data Warehouse tables. Then, the **fed_config.cfg** script and the **fed_admin.sh** or **fed_admin.bat** script that is running is updated.

Procedure

- 1. After you register the service or node by using the client software, open a command line.
- 2. Copy the AnalyticsDatabaseInstaller folder from the installed location to a temporary directory. The path of the AnalyticsDatabaseInstaller folder for PowerVM:
 - On Windows systems, copy the AnalyticsDatabaseInstaller directory from the C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cppowervm\ AnalyticsDatabaseInstaller directory to the c:\tmp directory.

• On operating systems other than Windows, copy the AnalyticsDatabaseInstaller directory from the opt/IBM/JazzSM/InstalledDashboards/com.ibm.tivoli.cppowervm/AnalyticsDatabaseInstaller directory to the /tmp directory.

The path of the AnalyticsDatabaseInstaller folder for VMware:

- On Windows systems: DASH_HOME\installedDashboards\com.ibm.tivoli.cpdash For example, C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cpdash
- On operating systems other than Windows: DASH_HOME/installedDashboards/ com.ibm.tivoli.cpdash

For example, /opt/IBM/JazzSM/installedDashboards/com.ibm.tivoli.cpdash/

The temporary directory must not include space characters in the file name.

Note: After you copied the AnalyticsDatabaseInstaller directory to a temp directory, the fed_admin.sh or fed_admin.bat and fed_config.cfg files are available in the AnalyticsDatabaseInstaller/federation directory.

- 3. Change the working directory to the directory where the fed_admin.sh or fed_admin.bat federation scripts files are copied.
 - The PowerVM federation scripts are stored in the federation_power directory.
 - The VMware federation scripts are stored in the federation directory.
- 4. Edit the fed_config.cfg file in the same directory and set the following properties with appropriate values:

TDW_REMOTE_HOST_NAME

The IP address or complete host name of the database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_PORT

The port number of the DB2 database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_DB_NAME

The name of the remote Tivoli Data Warehouse database, for example, CP_DB_NAME. The name of the capacity planner database, for example, TADFDCDB.

TDW_REMOTE_DB_SCHEMA_NAME

The schema name in the Tivoli Data Warehouse database.

TDW_REMOTE_DB_INSTANCE_OWNER

The instance owner of the remote Tivoli Data Warehouse database. The default owner is ITMUSER, who created the Tivoli Data Warehouse database on the remote Tivoli Data Warehouse database server.

TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD

The password of the remote Tivoli Data Warehouse database instance owner. This password is required when you register the federation server definition.

NODE_ID

The node name that is configured in the tnsnames.ora file on a UNIX system or the net service name that is configured by using Net8 connection manager on a Windows systems. The NODE_ID must not be more than eight characters in length, and must not contain special characters.

TDW_LOCAL_ALIAS

The alias name that is used to catalog the remote Tivoli Data Warehouse database on local DB2 database server.

FED_DB_SERVER_TYPE

The Oracle server data source, for example, Oracle.

FED_DB_VERSION

The version of the Oracle Server, for example 10 or 11.

Note: If a character is used in the **FED_DB_VERSION** field, this may cause problems. Therefore, you must use 11 instead of 11g for the **DB_VERSION** field.

USE_EXISTING_CATALOGED_TDW

This property must be set to Yes if you manually created an alias for remote Tivoli Data Warehouse database, and want to use the alias for federation. If the catalog is not created, the federation script attempts to create an alias for remote Tivoli Data Warehouse by using the value that is provided for TDW_LOCAL_ALIAS property. Set this property to NO if you do not want to use existing alias to be used in federation.

Note: Set **USE_EXISTING_CATALOGED_TDW** to NO only if you are sure that no alias is available on the capacity planner database with the name provided for the TDW_LOCAL_ALIAS property. Otherwise, the scripts might remove this alias after it fails to register the federation server.

CP_DB_NAME

The name of the capacity planner database, for example, TADFDCDB.

CP_DB_INSTANCE_OWNER

The instance owner of the capacity planner database. This owner is the same user that was used to create the capacity planner database.

CP_DB_INSTANCE_OWNER_PASSWORD

The password of the capacity planner database instance owner. This password is required to connect to the capacity planner database.

CP_DB_SCHEMA_NAME

The schema name in the Capacity planner database.

FED_SERVER_NAME

The name of the federation server that is configured. This name might be any unique name without special characters and space characters in it. You use this federation server when you federate the Tivoli Data Warehouse tables or views.

Note: Ensure that the parameter names in the fed_config.cfg file must not contain spaces.

5. Run the fed_admin.sh or fed_admin.bat script.

- On Windows systems, run the fed_admin.bat script.
- On operating systems other than Windows, complete the following steps:
 - a. Open a command line.
 - b. Run the fed_admin.sh \$JAVA_HOME -f fed_config.cfg script.

For more information about running the scripts, see "Configuring homogeneous federation for the Capacity Planner for VMware" on page 127.

Registering the server definition for DB2

You must register the federation server definition for DB2 and add links for all the required Tivoli Data Warehouse tables.

Procedure

1. Log on to the database server where the capacity planner database is installed.

Note: You must use the same user name that you used to create capacity planner database.

2. Copy the AnalyticsDatabaseInstaller directory from the installed location to a temporary directory.

- On Windows systems, copy the AnalyticsDatabaseInstaller directory from the C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cppowervm\AnalyticsDatabaseInstaller directory to the c:\tmp directory.
- On operating systems other than Windows, copy the AnalyticsDatabaseInstaller directory from the opt/IBM/JazzSM/InstalledDashboards/com.ibm.tivoli.cppowervm/AnalyticsDatabaseInstaller directory to the /tmp directory.

The temporary directory must not include space characters in the file name.

- 3. Change the working directory to the directory where the fed_admin.sh or fed_admin.bat federation script files are copied.
 - The PowerVM federation scripts are stored in the federation_power directory.
 - The VMware federation scripts are stored in the federation directory.
- 4. Edit the fed_config.cfg file in the same directory and set the following properties with appropriate values:

TDW_REMOTE_HOST_NAME

The IP address or complete host name of the database server where the Tivoli Data Warehouse is installed.

TDW_REMOTE_PORT

The port number of DB2 database server where the Tivoli Data Warehouse is installed

TDW_REMOTE_DB_NAME

The name of the remote Tivoli Data Warehouse database.

TDW_REMOTE_DB_SCHEMA_NAME

The schema name in the Tivoli Data Warehouse database.

TDW_REMOTE_DB_INSTANCE_OWNER

The instance owner of the remote Tivoli Data Warehouse database. The default owner is ITMUSER, who created the Tivoli Data Warehouse database on the remote Tivoli Data Warehouse database server.

TDW_REMOTE_DB_INSTANCE_OWNER_PASSWORD

The password of the remote Tivoli Data Warehouse database instance owner. This password is required when you register the federation server definition.

NODE_ID

The node name that is used to create DB2 node on database server. The NODE_ID must not be more than eight characters in length, and must not contain special characters.

TDW_LOCAL_ALIAS

The alias name that is used to catalog the remote Tivoli Data Warehouse database on local DB2 database server.

FED_DB_SERVER_TYPE

The DB2 server data source, for example, DB2/CS.

FED_DB_VERSION

The version of the DB2 Server, for example 10.1 or 9.7.

USE_EXISTING_CATALOGED_TDW

This must be set to Yes if you have manually created an alias for remote Tivoli Data Warehouse database, and want to use the alias for federation. If catalog is not created, then federation script will attempt to create an alias for remote Tivoli Data Warehouse by using the value provided for TDW_LOCAL_ALIAS property. Set this to NO if you do not want to use existing alias to be used in federation. **Note:** Set USE_EXISTING_CATALOGED_TDW to NO only if you are sure that no alias is available on capacity planner database with the name provided for property TDW_LOCAL_ALIAS. Otherwise, the scripts might remove such alias after it fails to register federation server.

CP_DB_NAME

The name of capacity planner database, for example, TADFDCDB.

CP_DB_INSTANCE_OWNER

The instance owner of the capacity planner database. This owner is the same user that was used to create the capacity planner database.

CP_DB_INSTANCE_OWNER_PASSWORD

The password of the capacity planner database instance owner. This password is required to connect to the capacity planner database.

CP_DB_SCHEMA_NAME

The schema name in the Capacity planner database. For PowerVM, the schema name is TADFDCP; and for VMware the schema name is TADFDC.

FED_SERVER_NAME

The name of the federation server that is configured. This name might be any unique name without special characters and space characters in it. You use this federation server when you federate the Tivoli Data Warehouse tables or views.

5. Run the fed_admin.sh/bat script.

- On Windows systems, run the fed_admin.bat %JAVA_HOME% -f fed_config.cfg script.
- On operating systems other than Windows, run the fed_admin.sh \$JAVA_HOME -f fed_config.cfg script.

You can run the **fed_admin.sh/bat** script in the following ways:

- Non-Interactive mode
 - a. fed_admin.sh %JAVA_HOME% -f fed_config.cfg
 - b. fed_admin.sh
 - c. ./fed_admin.sh /opt/-f fed_config.cfg
- Interactive Mode
 - a. fed_admin.sh -i
 - b. fed_admin.sh -i -f fed_config.cfg
 - c. ./fed_admin.sh %JAVA_HOME% -i -f fed_config.cfg

Updating configuration passwords for Dashboard, Reports, and Capacity Planners

You can update the passwords for dashboard, reports, and capacity planners regularly by completing the following steps:

Updating password for Capacity Planner database Procedure

Use the following steps to update the password for capacity planner applications and Tivoli Common Reporting:

- Configure the password again, which is used by the capacity planner applications. For more information about configuring a password, see "Verifying the Capacity Planner federation" on page 131.
- **2**. Configure the password again, which is used by Tivoli Common Reporting by completing the following steps:

- a. Log in to Dashboard Application Services Hub, and select Common Reporting.
- b. From the Launch list, select Administration.
- c. Click the **Configuration** tab.
- d. Select Data Source Connections.
- e. Select any one of the following data sources for Capacity Planner database:
 - TADFDC_Reporting_DS
 - TADFDCP_Reporting_DS
- f. Select the user for the data source (typically dbadmin or db2inst1), and select the **Set properties** option.
- g. Click the Signon tab, and then click the Edit the signon link.
- h. Type an updated password, and click OK.

Updating password for the dashboard data provider connection Procedure

Use the following steps to update the password for the dashboard data provider connection:

1. Log in to the Dashboard Application Services Hub.

Note: You must log in to the Dashboard Application Services Hub as a user assigned with the "administrator" and "iscadmins" roles. If the dashboard data provider connection is configured for single sign-on, you must also be logged in as a user in the LDAP user registry that is shared by the Dashboard Application Services Hub and the portal server.

- 2. Select Console Settings > Connections.
- 3. Delete the details for the connection with the ITMSD provider ID.

Important: Keep a backup of the details for the connection with the ITMSD provider ID.

4. Click **Create new remote provider**, and re-create the connection with the same settings as before except for the password.

For more details about creating a dashboard data provider connection, see "Creating a connection to the IBM Tivoli Monitoring dashboard data provider" in the *IBM Tivoli Monitoring Administrator's Guide*.

Updating password for connection from Dashboard for VMware to Tivoli Application Dependency Discovery Manager

You can again configure the password for connection from Dashboard for VMware to Tivoli Application Dependency Discovery Manager. For more information, see "Configuring Dashboard for VMware to communicate with Tivoli Application Dependency Discovery Manager" on page 134

Database tuning

For best performance of the Dashboard for VMware and Performance and Capacity Management Reports, keep the Tivoli Data Warehouse database and Capacity Planner database tuned.

- Ensure that buffer pools are at a minimum of 1 GB.
- For information about tuning the Tivoli Data Warehouse, see *Database Tuning* in the *IBM Tivoli Monitoring Installation and Setup Guide*.

146 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 8. Dashboard for VMware

You can determine the health of your clusters and all the components that are associated with the clusters by using the Dashboard for VMware.

You can investigate issues that the cluster might have by using the dashboard pages for servers, virtual machines, and data store. In addition, you can determine the health of an individual server that does not belong to a cluster.

To use Dashboard for VMware, ensure that you have installed the following prerequisites:

- IBM Tivoli Monitoring Version 6.3 or later
- IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI Version 7.1 or later
- VMware VI Agent Application Support 7.2 Fix Pack 2 or later
- Tivoli Application Dependency Discovery Manager 7.2.1 Fix Pack 4 or later
- NetApp agent
- NetApp Agent Support
- TPC agent
- TPC agent support

Dashboard for VMware Clusters

Use the Dashboard for VMware Clusters to determine the health of the clusters and all the associated components.

Accessing the Dashboard for VMware Clusters

Log on to the Tivoli Common Reporting or Dashboard Application Services Hub in your web browser to access the Dashboard for VMware Clusters.

Before you begin

Ensure that you created a connection to the data provider that you want to use to retrieve the information. For more information about creating a connection, see "Configuring connection for the Dashboard for VMware" on page 133.

Procedure

In the navigation tree, select **System Status and Health** > **Dashboard for VMware Clusters** to open the Dashboard for VMware Clusters page.

Clusters

The Clusters group page contains a cluster scorecard that displays the VMware-defined clusters that are available in the environment.

The Clusters group page displays a table that contains an entry for each VMware-defined cluster in the environment. All clusters are listed in the scorecard. The Clusters group page also displays the bar charts for the CPU, memory, and storage that are used by the top five clusters.

Note: For rows to display in the cluster scorecard, the VMware VI agent must be configured to monitor a VMware Virtual Center.

The upper-right corner of the Clusters group page displays the **Actions** menu. You can use the **Copy URL** option from the **Actions** menu to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

To filter rows, click inside the Filter field, and type the partial or full text to filter by. As you type, the

filter results are displayed. You can also click Apply Filter to view the filter results. You can sort columns in the scorecard by moving the mouse pointer over a column heading, and clicking the sort direction.

The cluster scorecard displays the following columns:

Table 9	Cluster	scorecard
rabic 0.	olusiol	Scorcoura

Column Name	Description
Cluster Name	Displays the name of the cluster.
	Click the name of the cluster to view the "Cluster details" on page 149.
Fatal event	Displays the number of fatal events.
Critical event	Displays the number of critical events.
Warning event	Displays the number of warnings.
Harmless event	Displays the number of harmless events.
Informational event	Displays the number of information messages.
Unknown event	Displays the number of unknown events.
Datacenter	Displays the name of the data center that the cluster belongs to.
Servers	Displays the number of servers that belong to the cluster.
Storage Used (%)	Displays the amount of storage capacity (in percentage) that is used by the cluster.
CPU Used (%)	Displays the amount of CPU (in percentage) that is used by the cluster.
Memory Used (%)	Displays the amount of memory (in percentage) that is used by the cluster.

Top 5 Clusters by % CPU Used

The **Top 5 Clusters by % CPU Used** chart shows the top five clusters for which the CPU usage (in percentage) is high. The chart is sorted by the descending order of the usage. You can view two sets of the bar charts. One chart displays the value of the CPU usage in percentage, and the second chart displays the total value of the CPU usage in GHz. Both types of charts are sorted with percentage data.

By default, the **Top 5 Clusters by % CPU Used** chart shows a chart of the CPU usage for the clusters in percentage. You can click the **Next** arrow to view the total value of the CPU usage in GHz.

- The Used CPU (%) bar shows the amount of effective CPU (in percentage) that is being used.
- The **Free CPU** (%) bar shows the amount of CPU (in percentage) that is available to run virtual machines in the clusters.
- The Used CPU (GHz) bar shows the amount of effective CPU (in GHz) that is being used.
- The **Free CPU (GHz)** bar shows the amount of CPU (in GHz) that is available to run virtual machines in the clusters.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Clusters by % Memory Used** and **Top 5 Clusters by % Storage Used** charts.

Top 5 Clusters by % Memory Used

The **Top 5 Clusters by % Memory Used** chart shows the top five clusters for which the memory usage (in percentage) is high. The chart is sorted by the descending order of the usage. You can view two sets

of bar charts. One chart displays the value of the memory usage in percentage, and the second chart displays the total value of the memory usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Clusters by** % **Memory Used** chart shows a bar chart of the memory usage for the clusters in percentage. You can click the **Next** arrow to view the total value of the memory usage in GB.

- The Used Memory (%) bar shows the amount of effective memory (in percentage) that is being used.
- The **Free Memory (%)** bar shows the amount of memory that is available (in percentage) to run virtual machines in the clusters.
- The Used Memory (GB) bar shows the amount of effective memory (in GB) that is being used.
- The **Free Memory (GB)** bar shows the amount of memory that is available (in GB) to run virtual machines in the clusters.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Clusters by % CPU Used** and **Top 5 Clusters by % Storage Used** charts.

Top 5 Clusters by % Storage Used

The **Top 5 Clusters by % Storage Used** chart shows the top five clusters for which the storage usage (in percentage) is high. The chart is sorted by the descending order of usage. You can view the two sets of bar charts. One chart displays the storage usage in percentage, and the second chart displays total value of the storage usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Clusters by** % **Storage Used** chart shows a bar chart of the storage usage for the clusters in percentage. You can click the **Next** arrow to view the total value of the storage usage in GB.

- The **Used Storage (%)** bar shows the amount of effective storage (in percentage) that is being used.
- The **Free Storage (%)** bar shows the amount of storage that is available (in percentage) to run virtual machines in the clusters.
- The Used Storage (GB) bar shows the amount of effective storage (in GB) that is being used.
- The **Free Storage (GB)** bar shows the amount of storage that is available (in GB) to run virtual machines in the clusters.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Clusters by** % **CPU Used** and **Top 5 Clusters by** % **Memory Used** charts.

Cluster details

The Cluster details page displays different tabs that contain information about the selected cluster and its components.

Each tab displays Resource Relationships table at the end. The Resource Relationships table displays similar information about each tab for that group page.

Resource Relationships: The Resource Relationships table displays a representative image and data from the components of the VMware infrastructure for the selected cluster. For example, data stores, virtual machines, servers, and the selected cluster.

- The Cluster component displays the name of the selected cluster.
- The Servers component displays the number of effective servers that are defined for the selected cluster. Click the number to view the "Servers" on page 153 group page. If only one effective server is defined for the selected cluster, the Servers component displays the name of that server. By clicking the name of the server, you can view the related Server details page.
- The Virtual Machines component displays the number of virtual machines that belong to the selected cluster. Click the number to view the "Virtual Machines" on page 159 group page. If only one virtual machine belongs to the selected cluster, the Virtual Machines component displays the name of that virtual machine. By clicking the name of the virtual machine, you can view the related Virtual Machine details page.

• The Datastores component displays the number of data stores that are related to the selected cluster. Click the number to view the "Datastores" on page 164 group page. If only one data store is related to the selected cluster, the Datastores component displays the name of the data store. By clicking the name of the data store, you can view the related Datastore details page.

Actions: The upper-right corner of the details page displays the **Actions** menu. You can complete the following tasks by using the **Actions** menu:

- Launch the Tivoli Enterprise Portal from the Actions menu. For more information about launching the Tivoli Enterprise Portal, see "Launching to Tivoli Enterprise Portal" on page 173.
- Launch the Tivoli Common Reporting from the Actions menu. For more information about launching the Tivoli Common Reporting, see "Launching to Tivoli Common Reporting" on page 173
- Use the **Copy URL** option to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

Overview:

The Overview tab provides information about a selected cluster in the environment.

The Overview tab displays the bar charts for the CPU, memory, storage, and servers that are used by the selected cluster. The Overview tab also displays the list of situation events for the selected cluster.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

Cluster CPU (*GHz*): The Cluster CPU (GHz) chart shows a bar chart of the total, effective, and used CPU for the selected cluster.

- The Total bar shows the total CPU resources that are defined for the cluster.
- The **Effective** bar shows the amount of CPU that is available to run virtual machines in the selected cluster.
- The **Used** bar shows the amount of CPU that is being used.

You can move the mouse pointer over a bar to view the exact CPU usage.

Cluster Memory (*GB*): The Cluster Memory (GB) chart shows a bar chart of the total, effective, used, and allocated memory for the selected cluster.

- The Total bar shows the total memory resources that are defined for the cluster.
- The Allocated bar shows the memory resources that have been allocated to the cluster.
- The **Effective** bar shows the amount of memory available to run virtual machines in the cluster. The effective memory resources do not include memory resources that are associated with hosts that are not responding or hosts in VMware maintenance mode.
- The **Used** bar shows the amount of memory that is being used.

You can move the mouse pointer over a bar to see the exact memory usage.

Cluster Storage (*GB*): The Cluster Storage (GB) chart shows a bar chart of the allocated, used, free, and total data store capacity for servers in the selected cluster.

- The Total bar shows the total amount of storage capacity available for the cluster.
- The Allocated bar is the data store capacity that is allocated to virtual machines in the cluster.
- The Free bar is the available data store capacity for servers in the cluster.
- The Used bar is the data store capacity currently being used by the servers in the cluster.

You can move the mouse pointer over a bar to view the exact storage usage.

of Servers: The *#* of Servers chart shows a bar chart that displays the total number of servers in the cluster. In addition, the bar chart displays the number of the servers that are effective, unavailable, and in maintenance mode.

- The Total bar represents the total number of servers that are defined for the cluster.
- The **Effective** bar is the number of servers that are defined for the cluster, and that are available to run virtual machines.
- The **Maintenance** bar is the number of unavailable servers that are in maintenance mode.
- The **Unavailable** bar is the number of servers that are defined for the cluster, and that are not available to run virtual machines.

Situation Events: The Situation Events table shows a table of situation events for the selected cluster. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Severity	Displays the severity of the situation event: Unknown, Informational, Harmless, Warning, Minor, Critical, or Fatal.
Status	Displays the status of the situation event.
Situation Name	Displays the short name that is assigned to the situation.
Display Item	Displays the attribute name if the situation is for a multiple-row attribute group and if a display item was selected.
Source	Displays the component source that originated the event. You can click the link to view the source of the event. Note: If the source of the event is available on the same page, the following message is displayed when you click the link: "The source of selected component situation is available on the same page!"
Global Timestamp	Displays the date and time (in the MMM DD, YYYY HH:MM:SS AM/PM format) when the event occurred or the condition was observed by the originating managed system.
Туре	Displays the type of event.

Table 10. Situation Events table

To filter rows, click inside the **Filter** field, type the partial or full text to filter by, and then click **Apply Filter**. You can sort columns in the table by moving the mouse pointer over a column heading, and then clicking the sort direction.

Note: The following are the prerequisites to view a situation in the Dashboard for VMware:

- The situation must sample a VMware attribute group that contains NodeID attribute.
- The situation must be associated with a Navigator item. For more information about associating a situation, see *IBM Tivoli Monitoring Information Center*.

To view the list of situations and VMware attribute groups that contains NodeID attribute, see "Dashboard for VMware problems" on page 369.

Properties:

The Properties tab provides information about the properties of the selected cluster in the environment.

The Properties tab displays the bar charts for the guests, virtual machines, servers, data stores, physical storage, and network that are used by the selected cluster.

of Guests: The *#* of Guests chart shows a bar chart of the guest operating systems for the virtual machines that are configured within the selected cluster.

- The **Unknown** bar shows the unknown guest operating system for the virtual machines that are configured within the selected cluster.
- The **Other** bar shows the guest operating systems other than Windows or Linux.
- The **Linux** bar shows the Linux guest operating system for the virtual machines that are configured within the selected cluster.
- The **Windows** bar shows the Windows guest operating system for the virtual machines that are configured within the selected cluster.
- The **Total Guests** bar shows the total guest operating system for the virtual machines that are configured within the selected cluster.

You can move the mouse pointer over a bar to view the exact number of guests.

of Virtual Machines: The *#* of Virtual Machines chart shows a bar chart of the running, powered on, and total virtual machines that are configured within the selected cluster.

- The **Running** bar shows the virtual machines in the selected cluster that are running.
- The **Powered On** bar shows the virtual machines in the selected cluster that are powered on.
- The Total bar shows the total virtual machines in the selected cluster.

You can move the mouse pointer over a bar to view the exact number of virtual machines.

of Servers: The *#* of Servers chart shows a bar chart of the effective and total servers along with the servers that are under maintenance for the selected cluster.

- The **Maintenance** bar shows the unavailable servers that are in the maintenance mode.
- The **Effective** bar is the number of servers that are defined for the cluster, and that are available to run the virtual machines.
- The Total bar is the total number of servers that are defined for the cluster.

You can move the mouse pointer over a bar to view the exact number of servers.

of Datastores: The *#* of Datastores chart shows a bar chart of the number of data stores for a particular file system like virtual machine file system (VMFS), network file system (NFS), and total file systems in the selected cluster.

- The VMFS bar shows the number of data store of VMFS file system in the selected cluster.
- The NFS bar the number of data store of NFS file system in the selected cluster.
- The Total bar is the total number of data stores of all file system in the selected cluster.

You can move the mouse pointer over a bar to view the exact number of data stores for a particular file system.

of Physical Storage: The *#* of Physical Storage chart shows a bar chart of the number of SAN volumes, NAS volumes, and total volumes that are associated with the selected cluster.

- The SAN bar shows the number of SAN volumes that are associated with the selected cluster.
- The NAS bar is the number of NAS volumes that are associated with the selected cluster.
- The Total bar is the total number of volumes that are associated with the selected cluster.

You can move the mouse pointer over a bar to view the exact type of volume.

of Physical NICs: The *#* of Physical NICs chart shows a bar chart of the number of physical network interface cards and physical network interface cards with a link status of down in the selected cluster.

- The **Physical NICs Down** bar shows the number of physical network interface cards (NIC) in the cluster with a link status of down.
- The Physical NIC bar is the number of physical NIC in the cluster.

You can move the mouse pointer over a bar to view the exact number of physical NIC.

Change History:

The Change History tab shows a table that lists the changes that are associated with the selected cluster as provided by the Tivoli Application Dependency Discovery Manager.

You can sort columns in the table by moving the mouse pointer over a column heading, and clicking the sort direction. The table contains the following columns:

Column name	Description
Туре	Displays the name of the cluster that changed.
Component	Displays the identifier for the specific component that changed.
Change	Displays the change action from one of the following actions:
	• Created
	• Updated
	• Deleted
Date	Displays the date and time of the change.
Attribute	Displays the component attribute changed.
Old	Displays the value before the change.
New	Displays the value following the change.

Table 11. Change history table

Note:

- The data on the Change History tab is displayed only when Tivoli Application Dependency Discovery Manager is configured.
- The change history data is fetched from Tivoli Application Dependency Discovery Manager for last two weeks only.

For detailed information, see the Tivoli Application Dependency Discovery Manager Information Center.

Servers

The Servers group page contains a server scorecard that displays the VMware-defined servers that are available in the environment.

The Servers group page displays a table that contains an entry for each VMware-defined server in the environment. All servers are listed in the scorecard. The Servers group page also displays the bar charts for the CPU, memory, and storage that are used by the top five servers.

Note: For rows to display in the server scorecard, the VMware VI agent must be configured to monitor a VMware Virtual Center.

The upper-right corner of the Servers group page displays the **Actions** menu. You can use the **Copy URL** option from the **Actions** menu to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

To filter rows, click inside the Filter field, and type the partial or full text to filter by. As you type, the

filter results are displayed. You can also click Apply Filter to view the filter results. You can sort columns in the scorecard by moving the mouse pointer over a column heading, and clicking the sort direction.

The server scorecard displays the following columns:

Column Name	Description
Server Name	Displays the name of the server.
	Click the name of the server to view the "Server details" on page 155.
Fatal event	Displays the number of fatal events.
Critical event	Displays the number of critical events.
Warning event	Displays the number of warnings.
Harmless event	Displays the number of harmless events.
Informational event	Displays the number of information messages.
Unknown event	Displays the number of unknown events.
Cluster	Displays the name of the cluster that the server belongs to.
Virtual Machines	Displays the number of virtual machines that belong to the server.
Average VMs CPU % Ready	Displays the average of all CPU percent ready values for all the virtual machines on the server.
CPU Used (%)	Displays the amount of CPU (in percentage) used by the server.
Memory Used (%)	Displays the amount of memory (in percentage) used by the server.
Overall Status	Displays the overall status of the server.
Number of CPUs	Displays the number of CPUs that belong to the server.
Memory Used (MB)	Displays the amount of memory (in MB) used by the server.
Connection Status	Displays the connection status of the server.

Table 12. Server scorecard

When you select a server in the scorecard, all other views are displayed in the context of the selected server.

Top 5 Servers by % CPU Used

The **Top 5 Servers by** % **CPU Used** chart shows the top five servers for which the CPU usage (in percentage) is high. The chart is sorted by the descending order of usage. You can view two sets of the bar charts. One chart displays the value of the CPU usage in percentage, and the second chart displays the total value of the CPU usage in MHz. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by** % **CPU Used** chart shows a chart of the CPU usage for the servers in percentage. You can click the **Next** arrow to view the total value of the CPU usage in MHz.

- The Used CPU (%) bar shows the amount of effective CPU (in percentage) that is being used.
- The Free CPU (%) bar shows the amount of the CPU (in percentage) that is available for the servers.
- The Used CPU (MHz) bar shows the amount of effective CPU (in MHz) that is being used.
- The Free CPU (MHz) bar shows the amount of the CPU (in MHz) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % Memory Used** and **Top 5 Servers by % Storage Used** charts.

Top 5 Servers by % Memory Used

The **Top 5 Servers by** % **Memory Used** chart shows the top five servers for which the memory usage (in percentage) is high. The chart is sorted by the descending order of usage. You can view two sets of bar charts. One chart displays the value of the memory usage in percentage, and the second chart displays the total value of the memory usage in MB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by % Memory Used** chart shows a bar chart of the memory usage for the servers in percentage. You can click the **Next** arrow to view the total value of the memory usage in MB.

- The Used Memory (%) bar shows the amount of effective memory (in percentage) that is being used.
- The **Free Memory (%)** bar shows the amount of memory (in percentage) that is available for the servers.
- The Used Memory (MB) bar shows the amount of effective memory (MB) that is being used.
- The Free Memory (MB) bar shows the amount of memory (MB) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % CPU Used** and **Top 5 Servers by % Storage Used** charts.

Top 5 Servers by % Storage Used

The **Top 5 Servers by % Storage Used** chart shows the top five servers for which the storage usage (in percentage) is high. The chart is sorted by the descending order of usage. You can view the two sets of bar charts. One chart displays the storage usage in percentage, and the second chart displays total value of the storage usage in MB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by % Storage Used** chart shows a bar chart of the storage usage for the servers in percentage. You can click the **Next** arrow to view the total value of the storage usage in MB.

- The Used Storage (%) bar shows the amount of effective storage (in percentage) that is being used.
- The **Free Storage** (%) bar shows the amount of storage (in percentage) that is available for the servers.
- The Used Storage (MB) bar shows the amount of effective storage (in MB) that is being used.
- The Free Storage (MB) bar shows the amount of storage (in MB) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % CPU Used** and **Top 5 Servers by % Memory Used** charts.

Server details

The Server details page displays different tabs that contain information about the selected server and its components.

Each tab displays Resource Relationships table at the end. The Resource Relationships table displays similar information about each tab for that group page.

Resource Relationships: The Resource Relationships table displays a representative image and data from the components of the VMware infrastructure for the selected server. For example, data stores, virtual machines, cluster, and the selected server.

- The Cluster component display the name of the cluster that is related to the selected server. Click the name to view the "Cluster details" on page 149.
- The Servers component displays the name of the selected server.
- The Virtual Machines component displays the number of virtual machines that belong to the selected server. Click the number to view the "Virtual Machines" on page 159 group page. If only one virtual machine belongs to the selected server, the Virtual Machines component displays the name of that virtual machine. By clicking the name of the virtual machine, you can view the related Virtual Machine details page.
- The Datastores component displays the number of data stores that are related to the selected server. Click the number to view the "Datastores" on page 164 group page. If only one data store is related to the selected server, the Datastores component displays the name of that data store. By clicking the name of the data store, you can view the related Datastore details page.

Actions: The upper-right corner of the details page displays the Actions menu.

The following actions are available on the **Actions** menu:

- Launch to Tivoli Common Reporting. For more information about launching to Tivoli Common Reporting, see "Launching to Tivoli Common Reporting" on page 173.
- Launch to Tivoli Enterprise Portal. For more information about launching to Tivoli Enterprise Portal, see "Launching to Tivoli Enterprise Portal" on page 173.
- Use the **Copy URL** option to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

Overview:

The Overview tab provides information about a selected server in the environment.

The Overview tab displays the bar charts for the CPU, memory, storage, and CPU and memory that is used by the selected server. The page also displays the list of situation events for the selected server.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

Server CPU (MHz): The Server CPU (MHz) chart shows a bar chart of the total and used CPU for the selected server.

- The Total bar shows the total CPU resources that are defined for the server.
- The **Used** bar shows the amount of the CPU that is being used by the server.

You can move the mouse pointer over a bar to view the exact CPU usage.

Server Memory (MB): The Server Memory (MB) chart shows a bar chart of the total and used memory for the selected server.

- The Used bar shows the amount of effective memory that is being used.
- The **Total** bar shows the total memory resources that are defined for the server.

You can move the mouse pointer over a bar to view the exact memory usage.

Server Storage (MB): The Server Storage (MB) chart shows a bar chart of the used and total data store capacity for the selected server.

- The Used bar is the data store capacity currently being used by the server.
- The **Total** bar shows the total amount of storage capacity of the server.

You can move the mouse pointer over a bar to view the exact storage usage.

CPU and Memory (%): The CPU and Memory (%) chart shows a bar chart of the memory, CPU, and average of all CPU percent ready values for the selected server.

- The **Memory** bar shows the amount of the effective memory that is being used.
- The **CPU** bar shows the amount of the effective CPU that is being used.
- The **Average VMs CPU % Ready** bar shows the average of all CPU percent ready values for all the virtual machines on the selected server.

You can move the mouse pointer over a bar to view the exact values.

Situation Events: The Situation Events table shows a table of situation events for the selected server. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Table 13. Situation Events table

Column name	Description
Severity	Displays the severity of the situation event: Unknown, Informational, Harmless, Warning, Minor, Critical, or Fatal.
Status	Displays the status of the situation event.
Situation Name	Displays the short name that is assigned to the situation.
Display Item	Displays the attribute name if the situation is for a multiple-row attribute group and if a display item was selected.
Source	Displays the component source that originated the event. You can click the link to view the source of the event. Note: If the source of the event is available on the same page, the following message is displayed when you click the link: "The source of selected component situation is available on the same page!"
Global Timestamp	Displays the date and time (in the MMM DD, YYYY HH:MM:SS AM/PM format) when the event occurred or the condition was observed by the originating managed system.
Туре	Displays the type of event.

To filter the rows, click inside the **Filter** field, and type the partial or full text to filter by, and then click **Apply Filter**. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction.

Note: The following are the prerequisites to view a situation in the Dashboard for VMware:

- The situation must sample a VMware attribute group that contains NodeID attribute.
- The situation must be associated with a Navigator item. For more information about associating a situation, see *IBM Tivoli Monitoring Information Center*.

To view the list of situations and VMware attribute groups that contains NodeID attribute, see "Dashboard for VMware problems" on page 369.

Configuration:

The Configuration tab provides configuration details about the server.

You can click the required tab to view the configuration details. For detailed information, see the Tivoli Application Dependency Discovery Manager Information Center.

Note: The data on the Configuration tab is displayed only when Tivoli Application Dependency Discovery Manager is configured.

Change History:

The Change History tab shows a table that lists the changes that are associated with the selected server as provided by the Tivoli Application Dependency Discovery Manager.

You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Туре	Displays the name of the server that changed.
Component	Displays the identifier for the specific component that changed.

Table 14. Change history table

Column name	Description
Change	Displays the change action from one of the following actions:
	• Created
	• Updated
	• Deleted
Date	Displays the date and time of the change.
Attribute	Displays the component attribute changed.
Old	Displays the value before the change.
New	Displays the value following the change.

Table 14. Change history table (continued)

Note:

- The data on the Change History tab is displayed only when Tivoli Application Dependency Discovery Manager is configured.
- The change history data is fetched from Tivoli Application Dependency Discovery Manager for last two weeks only.

For detailed information, see the Tivoli Application Dependency Discovery Manager Information Center.

Network:

The Network tab provides information about the network usage for a server in the environment.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

The Network tab displays the bar charts for the transmit and receive rate, link, and packets that are sent and received in the performance interval.

Transmit and Receive Rate (KBps): The Transmit and Receive Rate (KBps) chart shows a bar chart of the receive rate and transmit rate for the selected server.

- The **Receive Rate** bar shows the receive rate of the host on the virtual switch.
- The Transmit Rate bar shows the transmitted rate of the host on the virtual switch.
- The **Transmit and Receive Rate** bar shows the transmitted and received rate of the host on the virtual switch.

You can move the mouse pointer over a bar to view the exact values.

Link Speed (MBps): The Link Speed (MBps) chart shows a bar chart of the link speed for the selected server. You can move the mouse pointer over a bar to view the exact values.

By default, the Network tab displays the Link Speed (MBps) chart. You can click the **Next** or **Previous** arrow to view the Link Utilization (%) chart.

Link Utilization (%): The Link Utilization bar shows the usage (in percentage) of the NIC relative to the capacity of the link. You can move the mouse pointer over a bar to view the exact values.

You can click the Next or Previous arrow to view the Link Speed (MBps) chart.

By default, the Network tab displays the Link Speed (MBps) chart.

of Packets Sent and Received: The *#* of Packets Sent and Received chart shows a bar chart of the number of packets that received and packets that are sent in the performance interval for the selected server.

- The **Received** bar shows the number of packets that are received in the performance interval.
- The **Sent** bar shows the number of packets that are sent in the performance interval.

You can move the mouse pointer over a bar to view the exact values.

Network table:

Table 15. Network table

Column name	Description
Virtual Switch	Displays the name of the virtual switch.
NIC Name	Displays the name or label of the network interface.
Status	Displays the current status of the NIC.
Link Utilization	Displays the percent usage of the NIC relative to the capacity of the link (including duplex).
Link Speed (MBps)	Displays the current operating speed of the NIC (in MBps).
Transmission and Receive (KBps)	Displays the total rate (in KBps) that the host is transmitting and receiving data on the virtual switch.
Transmit (KBps)	Displays the total transmission rate (in KBps) of the host on the virtual switch.
Receive (KBps)	Displays the total reception rate (in KBps) of the host on the virtual switch.
Sent Packets	Displays the number of packets that are sent in the performance interval.
Received Packets	Displays the number of packets that are received in the performance interval.

Virtual Machines

The Virtual Machines group page contains a virtual machine scorecard that displays the VMware-defined virtual machines that are available in the environment.

The Virtual Machines group page displays a table that contains an entry for each VMware-defined virtual machine in the environment. All the virtual machines are listed in the scorecard.

Note: The virtual machine templates that are available at Tivoli Enterprise Portal are not displayed in the list.

The Virtual Machines group page also displays the bar charts for the CPU, memory, and storage that are used by the top five virtual machines.

Note: For rows to display in the virtual machine scorecard, the VMware VI agent must be configured to monitor a VMware Virtual Center.

The upper-right corner of the Virtual Machine group page displays the **Actions** menu. You can use the **Copy URL** option from the **Actions** menu to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

To filter rows, click inside the Filter field, and type the partial or full text to filter by. As you type, the

filter results are displayed. You can also click Apply Filter to view the filter results. You can sort columns in the scorecard by moving the mouse pointer over a column heading, and clicking the sort direction.

The virtual machine scorecard displays the following columns:

Column Name	Description
Virtual Machine	Displays the name of the virtual machine.
Name	Click the name of the virtual machine to view the "Virtual Machine details" on page 162.
Fatal event	Displays the number of fatal events.
Critical event	Displays the number of critical events.
Warning event	Displays the number of warnings.
Harmless event	Displays the number of harmless events.
Informational event	Displays the number of information messages.
Unknown event	Displays the number of unknown events.
Cluster	Displays the name of the cluster that the server belongs to.
Server	Displays the name of the server that the virtual machine belongs to.
Datastores	Displays the number of data store that the virtual machine is using.
CPU Utilization (%)	Displays the amount of CPU (in percentage) for virtual machines that are powered on.
Memory Utilization (%)	Displays the amount of memory (in percentage) used by the virtual machine.
Storage Utilization (%)	Displays the amount of storage (in percentage) used by the virtual machine.
Host Name	Displays the host name of the server that runs the virtual machine.
Power Status	Displays the current power status of the virtual machine.
Overall Status	Displays the overall status of the virtual machine.
CPU Used (MHz)	Displays the amount of CPU (in MHz) used by the server.
Memory Used (MB)	Displays the amount of memory (in MB) used by the server.
VM % Ready	Displays the CPU percent ready metric across all the virtual machine CPUs.
Storage Provisioned (GB)	Displays the total amount of space (in GB) that has been provisioned for use by VMs on the server.
Storage Committed (GB)	Displays the amount of space (in GB) on the data store, that this virtual machine is using.
Storage Uncommitted (GB)	Displays the reserved but unused amount of space (in GB) on the data store, that this virtual machine can use in the future.
Swap to File (MB)	Displays the total amount of virtual machine memory that has been swapped out to the swap file (in MB).
Number of Snapshots	Displays the number of snapshots for the virtual machine.

Table 16. Virtual machine scorecard

When you select a virtual machine in the scorecard, all other views are displayed in the context of the selected virtual machine.

Top 5 VMs by % CPU Used

The **Top 5 VMs by % CPU Used** chart shows the top five virtual machines for which the CPU usage (in percentage) is high. This chart shows a bar chart of the used and available CPU (MHz) for the virtual

machines. You can view two sets of the bar charts. One chart displays the value of the CPU usage in percentage, and the second chart displays the total value of the CPU usage in MHz. Both types of charts are sorted with percentage data.

By default, the **Top 5 VMs by % CPU Used** chart shows a chart of the CPU usage for the clusters in percentage. You can click the **Next** arrow to view the total value of the CPU usage in MHz.

- The Used CPU (%) bar shows the amount of effective CPU (in percentage) that is being used.
- The **Free CPU** (%) bar shows the amount of the CPU (in percentage) that is available for the virtual machines.
- The Used CPU (MHz) bar shows the amount of effective CPU (in MHz) that is being used.
- The **Free CPU (MHz)** bar shows the amount of the CPU (in MHz) that is available for the virtual machines.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 VMs by % Memory Used** and **Top 5 VMs by % Storage Used** charts.

Top 5 VMs by % Memory Used

The **Top 5 VMs by** % **Memory Used** chart shows the top five virtual machines for which the memory usage (in percentage) is high. This chart shows a bar chart of the used and available memory for the virtual machines. One chart displays the value of the memory usage in percentage, and the second chart displays the total value of the memory usage in MB. Both types of charts are sorted with percentage data.

By default, the **Top 5 VMs by % Memory Used** chart shows a bar chart of the memory usage for the virtual machines in percentage. You can click the **Next** arrow to view the total value of the memory usage in MB.

- The Used Memory (%) bar shows the amount of effective memory (in percentage) that is being used.
- The **Free Memory (%)** bar shows the amount of total memory (in percentage) that is available for the virtual machines.
- The Used Memory (MB) bar shows the amount of effective memory (in MB) that is being used.
- The **Free Memory (MB)** bar shows the amount of total memory (in MB) that is available for the virtual machines.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 VMs by % CPU Used** and **Top 5 VMs by % Storage Used** charts.

Top 5 VMs by % Storage Used

The **Top 5 VMs by** % **Storage Used** chart shows the top five virtual machines for which the storage usage (in percentage) is high. This chart shows a bar chart of the used and available memory for the virtual machines. You can view the two sets of bar charts. One chart displays the storage usage in percentage, and the second chart displays total value of the storage usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 VMs by** % **Storage Used** chart shows a bar chart of the storage usage for the clusters in percentage. You can click the **Next** arrow to view the total value of the storage usage in GB.

- The **Used Storage (%)** bar shows the average amount of effective storage (in percentage) that is being used.
- The **Free Storage (%)** bar shows the average amount of total storage (in percentage) that is available for the virtual machines.
- The Used Storage (GB) bar shows the average amount of effective storage (in GB) that is being used.
- The **Free Storage (GB)** bar shows the average amount of total storage (in GB) that is available for the virtual machines.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 VMs by % CPU Used** and **Top 5 VMs by % Memory Used** charts.

Virtual Machine details

The Virtual Machine details page displays different tabs that contain information about the selected virtual machine and the components that are related to the selected virtual machine.

Each tab displays Resource Relationships table at the end. The Resource Relationships table displays similar information about each tab for that group page.

Resource Relationships: The Resource Relationships table displays a representative image and data from the components of the VMware infrastructure for the selected virtual machine. For example, data stores, server, cluster, and the selected virtual machine.

- The Cluster component display the name of the cluster that is related to the selected virtual machine. Click the name to view the "Cluster details" on page 149 page.
- The Server component displays the name of the server that is related to the selected virtual machine. Click the name to view the "Server details" on page 155 page.
- The Virtual Machine component displays the name of selected virtual machine.
- The Datastores component displays the number of data stores that are related to the selected virtual machine. Click the number to view the "Datastores" on page 164 group page. If only one data store is related to the selected virtual machine, the Datastores component displays the name of that data store. By clicking the name of the data store, you can view the related Datastore details page.
- The Guest OS component displays the name of the guest operating system for the selected virtual machine. Click the link to view the respective Server Dashboards page.

Note: The Guest OS link is enabled only when OS agent is installed and online on the managed system. The dashboard data provider verifies that the OS agent is connected to the same Tivoli Enterprise Monitoring Server and is online.

Actions: The upper-right corner of the details page displays the **Actions** menu. You can complete the following tasks by using the **Actions** menu:

- Launch the Tivoli Enterprise Portal from the Actions menu. For more information about launching the Tivoli Enterprise Portal, see "Launching to Tivoli Enterprise Portal" on page 173.
- Launch the Tivoli Common Reporting from the Actions menu. For more information about launching the Tivoli Common Reporting, see "Launching to Tivoli Common Reporting" on page 173
- Use the **Copy URL** option to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

Overview:

The Overview tab provides information about the selected virtual machine in the environment.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

The Overview tab displays the bar charts for the CPU, memory, storage, and CPU and memory that is used by the selected virtual machine. The page also displays situation events for the selected virtual machine.

VM CPU (MHz): The VM CPU (MHz) chart shows a bar chart of the free, total, and used CPU for the selected virtual machine.

- The **Free** bar shows the amount of available CPU for the virtual machine.
- The **Used** bar shows the amount of CPU that is being used by the virtual machine.
- The Total bar shows the total CPU resources that are defined for the virtual machine.

You can move the mouse pointer over a bar to view the exact CPU usage.

VM Memory (*MB*): The VM Memory (MB) chart shows a bar chart of the free, total, and used memory for the selected virtual machine.

- The Free bar shows the amount of memory that is currently free.
- The Used bar shows the amount of memory that is being used.
- The **Total** bar shows the total memory resources that are defined for the virtual machine.

You can move the mouse pointer over a bar to see the exact memory usage.

VM Storage (*GB*): The VM Storage (GB) chart shows a bar chart of the free, total, and used data store capacity for the selected virtual machine.

- The Free bar is the data store capacity that is available for the virtual machine.
- The Used bar is the data store capacity currently being used by the virtual machine.
- The Total bar shows the total amount of storage capacity available.

You can move the mouse pointer over a bar to view the exact storage usage.

Situation Events: The Situation Events table shows a table of situation events for the selected virtual machine. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Severity	Displays the severity of the situation event: Unknown, Informational, Harmless, Warning, Minor, Critical, or Fatal.
Status	Displays the status of the situation event.
Situation Name	Displays the short name that is assigned to the situation.
Display Item	Displays the attribute name if the situation is for a multiple-row attribute group and if a display item was selected.
Source	Displays the component source that originated the event. You can click the link to view the source of the event. Note: If the source of the event is available on the same page, the following message is displayed when you click the link: "The source of selected component situation is available on the same page!"
Global Timestamp	Displays the date and time (in the MMM DD, YYYY HH:MM:SS AM/PM format) when the event occurred or the condition was observed by the originating managed system.
Туре	Displays the type of event.

Table 17. Situation Events table

To filter the rows, click inside the **Filter** field, and type the partial or full text to filter by, and then click **Apply Filter**. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction.

Note: The following are the prerequisites to view a situation in the Dashboard for VMware:

- The situation must sample a VMware attribute group that contains NodeID attribute.
- The situation must be associated with a Navigator item. For more information about associating a situation, see *IBM Tivoli Monitoring Information Center*.

To view the list of situations and VMware attribute groups that contains NodeID attribute, see "Dashboard for VMware problems" on page 369.

Change History:

The Change History tab shows a table that lists the changes that are associated with the selected virtual machine as provided by the Tivoli Application Dependency Discovery Manager.

You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Туре	Displays the name of the virtual machine that changed.
Component	Displays the identifier for the specific component that changed.
Change	Displays the change action from one of the following actions:
	• Created
	• Updated
	• Deleted
Date	Displays the date and time of the change.
Attribute	Displays the component attribute changed.
Old	Displays the value before the change.
New	Displays the value following the change.

Table 18. Change history table

Note:

- The data on the Change History tab is displayed only when Tivoli Application Dependency Discovery Manager is configured.
- The change history data is fetched from Tivoli Application Dependency Discovery Manager for last two weeks only.

For detailed information, see the Tivoli Application Dependency Discovery Manager Information Center.

Datastores

The Datastores group page contains a data store scorecard that displays the VMware-defined data stores that are available in the environment. You can control the context of the charts and views on the Datastores group page by selecting a row in the data store scorecard. When you select a data store, the information is updated for the data store that you selected.

The Datastores group page displays a table that contains an entry for each VMware-defined data store in the environment. All servers are listed in the scorecard. The Datastores group page also displays the bar charts for the CPU, memory, and storage that are used by the top five data stores.

Note: For rows to display in the data store scorecard, the VMware VI agent must be configured to monitor a VMware Virtual Center.

The upper-right corner of the Datastores group page displays the **Actions** menu. You can use the **Copy URL** option from the **Actions** menu to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

To filter rows, click inside the Filter field, and then type the partial or full text to filter by. As you type,

the filter results are displayed. You can also click Apply Filter to view the filter results. You can sort columns in the scorecard by moving the mouse pointer over a column heading, and clicking the sort direction.

The data stores scorecard displays the following columns:

Column Name	Description
Datastore Name	Displays the name of the data store.
	Click the name of the data store to view the "Datastore details" on page 166 page.
Fatal event	Displays the number of fatal events.
Critical event	Displays the number of critical events.
Warning event	Displays the number of warnings.
Harmless event	Displays the number of harmless events.
Informational event	Displays the number of information messages.
Unknown event	Displays the number of unknown events.
Connected Clusters	Displays the number of clusters that are connected to the data store.
Connected Servers	Displays the number of servers that are connected to the data store.
Storage Used (GB)	Displays the total amount of data store storage (in GB) that is actually in use by the selected server.
Storage Used (%)	Displays the total amount of data store storage (in percentage) that is actually in use by the selected server.
Snapshots	Displays the number of snapshots for a virtual machine on the selected data store.
Status	Displays the overall status of the data store.
Total Capacity (GB)	Displays the storage capacity (in GB) of the data store.
Percent Overcommitted (%)	Displays the percentage of the total capacity of the data store, which is overcommitted.

Table 19. Data store scorecard

When you select a data store in the scorecard, all other views are displayed in the context of the selected data store.

Top 5 Datastores by % Space Used

The **Top 5 Datastores by % Space Used** chart shows the top five data stores for which the amount of allocated storage (in percentage) is high. The chart is sorted by the descending order of usage. You can view two sets of the bar charts. One chart displays the value of the CPU usage in percentage, and the second chart displays the total value of the CPU usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Datastores by % Space Used** chart shows a chart of the space usage for the data stores in percentage. You can click the **Next** arrow to view the total value of the space usage in GB.

- The **Used Storage (%)** bar shows the amount of allocated storage (in percentage) that is being used.
- The Free Storage (%) bar shows the amount of storage (in percentage) that is available on the data store.
- The Used Storage (GB) bar shows the amount of allocated storage (in GB) that is being used.
- The Free Storage (GB) bar shows the amount of storage (in GB) that is available on the data store.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Datastores by % Snapshot Space Used** and **Top 5 Datastores by Number of VMs** charts.

Top 5 Datastores by % Snapshot Space Used

The **Top 5 Datastores by % Snapshot Space Used** chart shows the top five data stores for which the amount of disk space that is used by the snapshot (in percentage) is high. The chart is sorted by the descending order of the usage. You can view two sets of the bar charts. One chart displays the value of the disk space usage in percentage, and the second chart displays the total value of the disk space usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Datastores by % Snapshot Space Used** chart shows a chart of the disk space usage for the data stores in percentage. You can click the **Next** arrow to view the total value of the disk space usage in GB.

- The Used Storage (%) bar shows the amount of storage (in percentage) that is used by snapshot.
- The Free Storage (%) bar shows the amount of storage (in percentage) that is available on the data store.
- The Used Storage (GB) bar shows the amount of storage (in GB) that is used by snapshot.
- The Free Storage (GB) bar shows the amount of storage (in GB) that is available on the data store.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Datastores by % Space Used** and **Top 5 Datastores by Number of VMs** charts.

Top 5 Datastores by Number of VMs

The **Top 5 Datastores by Number of VMs** chart shows the top five data stores for which the amount of space that is used by the virtual machine is high. The chart is sorted by the descending order of usage. You can view two sets of the bar charts. One chart displays the value of the space usage in percentage, and the second chart displays the total value of the space usage in GB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Datastores by Number of VMs** chart shows a chart of the space usage for the data stores in percentage. You can click the **Next** arrow to view the total value of the space usage in GB.

- The **Used Storage (%)** bar shows the amount of space (in percentage) on the data store used by the virtual machine.
- The Free Storage (%) bar shows the amount of space (in percentage) that is available on the data store.
- The **Used Storage (GB)** bar shows the amount of space (in GB) on the data store used by the virtual machine.
- The Free Storage (GB) bar shows the amount of space (in GB) that is available on the data store.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Datastores by % Space Used** and **Top 5 Datastores by % Snapshot Space Used** charts.

Datastore details

The Datastore details page displays different tabs that contain information about the selected data store and the components that are related to the selected data store.

Each tab displays Resource Relationships table at the end. The Resource Relationships table displays similar information about each tab for that group page.

Resource Relationships: The Resource Relationships table displays a representative image and data from the components of the VMware infrastructure for the selected data store. For example, data stores, server, cluster, and the selected virtual machine.

• The Connected Clusters component display the number of the clusters that are related to the selected data store. Click the number to view the "Clusters" on page 147 group page. If only one cluster is related to the selected data store, the Connected Clusters component displays the name of that connected cluster. By clicking the name of the connected cluster, you can view the related Clusters details page.

Note: This link is disabled when you access the Datastore details page from Dashboard for VMware ESX Servers.

- The Connected Servers component displays the number of the servers that are related to the selected data store. Click the number to view the "Servers" on page 153 group page. If only one server is related to the selected data store, the Connected Servers component displays the name of that connected server. By clicking the name of the connected server, you can view the related Servers details page.
- The Connected Virtual Machines component displays the number of virtual machines that are related to the selected data store. Click the number to view the "Virtual Machines" on page 159 group page. If only one virtual machine is related to the selected data store, the Connected Virtual Machines component displays the name of that connected virtual machine. By clicking the name of the connected virtual machine, you can view the related Virtual Machine details page.
- The Datastore component displays the name of selected data store.

Actions: The upper-right corner of the details page displays the **Actions** menu. You can complete the following tasks by using the **Actions** menu:

- Launch the Tivoli Enterprise Portal from the **Actions** menu. For more information about launching the Tivoli Enterprise Portal, see "Launching to Tivoli Enterprise Portal" on page 173.
- Launch the Tivoli Common Reporting from the Actions menu. For more information about launching the Tivoli Common Reporting, see "Launching to Tivoli Common Reporting" on page 173
- Use the **Copy URL** option to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

Overview:

The Overview tab provides information about a selected data store in the environment.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

The Overview tab displays the bar charts for the space that is used, data store I/O, and virtual machines and snapshots that is used by the selected data store. The page also displays situation events for the selected data store.

Datastore Space Used (GB): The Datastore Space Used (GB) chart shows a bar chart of the total, and used space for the selected data store.

- The Used Space bar shows the amount of space that is being used by the selected data store.
- The **Total Capacity** bar shows the total space on the selected data store.

You can move the mouse pointer over a bar to view the exact values.

of VMs and Snapshots: The *#* of VMs and Snapshots chart shows a bar chart of the virtual machines and snapshots for the selected data store.

- The **Virtual Machines** bar shows the number of virtual machines that are connected to the selected data store.
- The **Snapshots** bar shows the total number of snapshots on a selected data store.

You can move the mouse pointer over a bar to view the exact values.

Datastore I/O (kBps): The Datastore I/O (kBps) chart shows a bar chart for the amount of data being read and written (kBps) by the virtual machines on the selected data store.

• The **Write** bar is the KB (kilobytes) written per second by all virtual machines that are configured for the selected data store.

- The **Read** bar is the KB read per second by all virtual machines that are configured for the selected data store.
- The **Total** bar shows the sum of total KB read and written per second by all virtual machines that are configured for the selected data store.

You can move the mouse pointer over a bar to view the exact values.

Situation Events: The Situation Events table shows a table of situation events for the selected data store. You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Severity	Displays the severity of the situation event: Unknown, Informational, Harmless, Warning, Minor, Critical, or Fatal.
Status	Displays the status of the situation event.
Situation Name	Displays the short name that is assigned to the situation.
Display Item	Displays the attribute name if the situation is for a multiple-row attribute group and if a display item was selected.
Source	Displays the component source that originated the event.
Global Timestamp	Displays the date and time (in the MMM DD, YYYY HH:MM:SS AM/PM format) when the event occurred or the condition was observed by the originating managed system.
Туре	Displays the type of event.

Table 20. Situation Events table

To filter the rows, click inside the **Filter** field, and type the partial or full text to filter by, and then click **Apply Filter**.

Note: Following are the prerequisites to view a situation in the Dashboard for VMware:

- The situation must sample a VMware attribute group that contains NodeID attribute.
- The situation must be associated with a Navigator item. For more information about associating a situation, see "Dashboard for VMware problems" on page 369.

Change History:

The Change History tab shows a table that lists the changes that are associated with the selected data store as provided by the Tivoli Application Dependency Discovery Manager.

You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction. The table contains the following columns:

Column name	Description
Туре	Displays the name of the data store that changed.
Component	Displays the identifier for the specific component that changed.
Change	Displays the change action from one of the following actions:
	• Created
	• Updated
	• Deleted
Date	Displays the date and time of the change.
Attribute	Displays the component attribute changed.

Table 21. Change history table
Table 21. Change history table (continued)

Column name	Description	
Old	Displays the value before the change.	
New	Displays the value following the change.	

Note:

- The data on the Change History tab is displayed only when Tivoli Application Dependency Discovery Manager is configured.
- The change history data is fetched from Tivoli Application Dependency Discovery Manager for last two weeks only.

For detailed information, see the Tivoli Application Dependency Discovery Manager Information Center.

Volume:

The Volume tab provides information about the volume usage for a data store in the environment.

Select the required time range from the time selector list. You can select real time or a historic time period for the data to plot in the chart. The time selector feature is applicable only for charts.

The Volume tab displays the bar charts for the latency and space that is used by the selected data store.

Latency (millisecond): The Latency (millisecond) chart shows a bar chart of the latency of the data store on a physical volume.

- The Other bar shows the amount of time that is taken for any other operation from the data store.
- The Write bar shows the amount of time that is taken for a write operation from the data store.
- The **Read** bar shows the amount of time that is taken for a read operation from the data store.
- The **Total** bar shows the total time that is taken to complete read, write, and other operations from the data store.

You can move the mouse pointer over a bar to view the exact values.

Volume Space Used (GB): The Volume Space Used (GB) chart shows a bar chart of the used and total capacity for the selected data store.

- The **Total Capacity** bar shows the total capacity for the selected data store.
- The **Used** bar shows the used capacity for the selected data store.

You can move the mouse pointer over a bar to view the exact values.

Volume table: You can sort columns in the table by moving the mouse pointer over a column heading and clicking the sort direction.

Column name	Description
Volume Name	Displays the name of the volume.
Volume Type	Displays the type of the volume.
Total Capacity (GB)	Displays the total space on the selected data store.
Space Used (GB)	Displays the amount of space (in GB) that is being used by a particular volume.
Total Operations/sec	Displays the number of operations per second on a particular volume.

Table 22. Volume table

Table 22. Volume table (continued)

Column name	Description
Read Latency	Displays the average amount of time that is taken for a read operation by a particular volume.
Write Latency	Displays the average amount of time that is taken for a write operation by a particular volume.
Other Latency	Displays the average amount of time that is taken for any other operation by a particular volume.

Dashboard for VMware ESX Servers

Use the Dashboard for VMware ESX Servers to determine the health of an individual server that might or might not belong to a cluster.

Accessing the Dashboard for VMware ESX Servers

Log on to the Tivoli Common Reporting or Dashboard Application Services Hub in your web browser to access the Dashboard for VMware ESX Servers.

Before you begin

Ensure that you created a connection to the data provider that you want to use to retrieve the information. For more information about creating a connection, see "Configuring connection for the Dashboard for VMware" on page 133.

Procedure

In the navigation tree, select **System Status and Health** > **Dashboard for VMware ESX Servers** to open the Dashboard page.

Servers

The Servers group page contains a server scorecard that displays all the VMware-defined servers that are available in the environment. A server can be a stand-alone server or might belong to a cluster.

The Servers group page displays a table that contains an entry for each VMware-defined server in the environment. All servers are listed in the scorecard table. The Servers group page also displays the bar charts for the CPU, memory, and storage that is used by the top five servers.

Note: For rows to display in the server scorecard table, the VMware VI agent must be configured to monitor a VMware Virtual Center.

The upper-right corner of the Servers group page displays the **Actions** menu. You can use the **Copy URL** option from the **Actions** menu to share the current dashboard page. For more information about sharing the current dashboard page, see "Sharing the current dashboard page link" on page 175.

To filter rows, click inside the Filter field and type the partial or full text to filter by. As you type, the

filter results are displayed. You can also click Apply Filter to view the filter results. You can sort columns in the scorecard by moving the mouse pointer over a column heading and clicking the sort direction.

The server scorecard table displays the following columns:

Column Name	Description
Server Name	Displays the name of the server.
	Click the name of the server to view the "Server details" on page 155.
Fatal event	Displays the number of fatal events.
Critical event	Displays the number of critical events.
Warning event	Displays the number of warnings.
Harmless event	Displays the number of harmless events.
Informational event	Displays the number of information messages.
Unknown event	Displays the number of unknown events.
Cluster	Displays the name of the cluster if the cluster belongs to the server.
Virtual Machines	Displays the number of virtual machines that belong to the server.
Average VMs CPU % Ready	Displays the average of all CPU percent ready values for all the virtual machines on the server.
CPU Used (%)	Displays the amount of CPU (in percentage) used by the server.
Memory Used (%)	Displays the amount of memory (in percentage) used by the server.
Overall Status	Displays the overall status of the server.
Number of CPUs	Displays the number of CPUs that belong to the server.
Memory Used (MB)	Displays the amount of memory (in MB) used by the server.
Connection Status	Displays the connection status of the server.

Table 23. Server scorecard table

When you select a server in the scorecard, all other views are displayed in the context of the selected server.

Top 5 Servers by % CPU Used

The Top 5 Servers by % CPU Used chart shows the top five servers for which the CPU usage (in percentage) is high. The chart is sorted by the descending order of utilization. You can view two sets of the bar charts. One chart displays the value of the CPU usage in percentage, and the second chart displays the total value of the CPU usage in MHz. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by % CPU Used** chart shows a chart of the CPU usage for the servers in percentage. You can click the **Next** arrow to view the total value of the CPU usage in MHz.

- The Used CPU (%) bar shows the amount of effective CPU (in percentage) that is being used.
- The Free CPU (%) bar shows the amount of the CPU (in percentage) that is available for the servers.
- The Used CPU (MHz) bar shows the amount of effective CPU (MHz) that is being used.
- The Free CPU (MHz) bar shows the amount of the CPU (MHz) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % Memory Used** and **Top 5 Servers by % Storage Used** charts.

Top 5 Servers by % Memory Used

The **Top 5 Servers by** % **Memory Used** chart shows the top five servers for which the memory usage (in percentage) is high. The chart is sorted by the descending order of utilization. You can view two sets of bar charts. One chart displays the value of the memory usage in percentage, and the second chart displays the total value of the memory usage in MB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by % Memory Used** chart shows a bar chart of the memory usage for the servers in percentage. You can click the **Next** arrow to view the total value of the memory usage in MB.

- The Used Memory (%) bar shows the amount of effective memory (in percentage) that is being used.
- The **Free Memory (%)** bar shows the amount of memory (in percentage) that is available for the servers.
- The Used Memory (MB) bar shows the amount of effective memory (in MB) that is being used.
- The Free Memory (MB) bar shows the amount of memory (in MB) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % CPU Used** and **Top 5 Servers by % Storage Used** charts.

Top 5 Servers by % Storage Used

The **Top 5 Servers by** % **Storage Used** chart shows the top five servers for which the storage usage (in percentage) is high. The chart is sorted by the descending order of utilization. You can view the two sets of bar charts. One chart displays the storage usage in percentage, and the second chart displays total value of the storage usage in MB. Both types of charts are sorted with percentage data.

By default, the **Top 5 Servers by % Storage Used** chart shows a bar chart of the storage usage for the servers in percentage. You can click the **Next** arrow to view the total value of the storage usage in MB.

- The Used Storage (%) bar shows the amount of effective storage (in percentage) that is being used.
- The **Free Storage (%)** bar shows the amount of storage (in percentage) that is available for the servers.
- The Used Storage (MB) bar shows the amount of effective storage (in MB) that is being used.
- The Free Storage (MB) bar shows the amount of storage (in MB) that is available for the servers.

By clicking the **Previous** or **Next** arrow, you can view the **Top 5 Servers by % CPU Used** and **Top 5 Servers by % Memory Used** charts.

Server, Cluster, Virtual Machine, and Datastore details

The pages that display details of the server, cluster, virtual machine, and data store are same whether accessed from the Dashboard for VMware Clusters or the Dashboard for VMware ESX Servers. The difference is the context: When accessed from the Dashboard for VMware ESX Servers, the details are displayed for all servers. A server can be a stand-alone server or might belong to a cluster.

Related concepts:

"Server details" on page 155

The Server details page displays different tabs that contain information about the selected server and its components.

"Clusters" on page 147

The Clusters group page contains a cluster scorecard that displays the VMware-defined clusters that are available in the environment.

"Cluster details" on page 149

The Cluster details page displays different tabs that contain information about the selected cluster and its components.

"Virtual Machines" on page 159

The Virtual Machines group page contains a virtual machine scorecard that displays the VMware-defined virtual machines that are available in the environment.

"Virtual Machine details" on page 162

The Virtual Machine details page displays different tabs that contain information about the selected virtual machine and the components that are related to the selected virtual machine.

"Datastores" on page 164

The Datastores group page contains a data store scorecard that displays the VMware-defined data stores that are available in the environment. You can control the context of the charts and views on the Datastores group page by selecting a row in the data store scorecard. When you select a data store, the

information is updated for the data store that you selected.

"Datastore details" on page 166

The Datastore details page displays different tabs that contain information about the selected data store and the components that are related to the selected data store.

Launching to Tivoli Common Reporting

You can launch to the Tivoli Common Reporting interface to view and manage reports. You can select the report that you want to view from the details page of the dashboard.

Before you begin

You must have Tivoli Common Reporting and VMware VI reports installed and configured in your environment.

About this task

Complete these steps to open a new browser window and launch to the Tivoli Common Reporting to view the selected report.

Procedure

- 1. From the Dashboard for VMware Clusters or Dashboard for VMware ESX Servers, go to the details page. For example, Cluster details page.
- 2. On the upper-right corner, click **Actions** > **Launch to TCR**, and select the required report from the available list.

Results

The selected report is displayed in Tivoli Common Reporting in your browser. For more information about Tivoli Common Reporting, see the Tivoli Common Reporting Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ic-home.html).

Launching to Tivoli Enterprise Portal

You can launch to the Tivoli Enterprise Portal browser client to complete tasks such as editing situations and acknowledging events.

Before you begin

Your user ID must be registered with the Tivoli Enterprise Portal Server. If single sign-on (SSO) has been established, you must not authenticate again when you launch to the Tivoli Enterprise Portal. For more information, see "Enabling user authentication" in the *IBM Tivoli Monitoring Administrator's Guide*.

About this task

Complete these steps to open a new browser window, and launch to the Tivoli Enterprise Portal to the default workspace. For example, if you launch from the Server details page of the Dashboard for VMware ESX Servers, the workspace for that particular ESX server is displayed.

Note: You can launch to Tivoli Enterprise Portal from any of the following dashboard pages:

- Cluster details
- Server details
- Virtual Machine details
- Datastore details

Procedure

- 1. From the Dashboard for VMware Clusters or Dashboard for VMware ESX Servers, go to any of the following dashboard pages that you want to launch Tivoli Enterprise Portal from:
 - Cluster details
 - Server details
 - Virtual Machine details
 - Datastore details
- 2. On the upper-right corner, click **Actions** > **Launch to TEP**.

Results

The Tivoli Enterprise Portal opens in your browser. Context information is passed from the dashboard to the Tivoli Enterprise Portal, and an associated default workspace is displayed.

Note: If the Tivoli Enterprise Portal does not start, there might be a problem running the browser client for such reasons as the Java Plug-in not being installed or configured correctly. You can review the browser client start up information in the *IBM Tivoli Monitoring Installation and Setup Guide* and browser client troubleshooting information in the *IBM Tivoli Monitoring Troubleshooting Guide*.

Setting a trace

IBM Infrastructure Management Dashboard for VMware has several levels of tracing that you can set while using the dashboards. You can start a higher level of tracing exactly at the point in the user interface where you are having a problem, then set tracing back to a lower level when you have captured the necessary log data.

Adjust the trace settings to help your administrator or IBM Support to diagnose the cause of problems with the Dashboard for VMware.

For example, if a particular dashboard is behaving unexpectedly, you can raise the trace level before opening the dashboard to log the activity and then return trace logging to the normal level.

About this task

Take the following steps to set the trace level when you want to increase or reduce the amount of trace logging.

Procedure

- 1. If Dashboard for VMware is not already open, select **System Status and Health** > **Dashboard for VMware** and navigate to the dashboard where you want to change the trace level.
- 2. Click Actions > Trace level and select one of the following levels:
 - Verbose to have all activity logged. Verbose trace level includes Moderate, Light, and Minimal trace logging.
 - **Moderate** to have variable changes logged, such as what parameters were passed in and what calculations were made. Moderate trace level includes Light and Minimal trace logging.
 - **Light** to log error and variable activity. You might want to set the trace to this level if you have a problem such as no data being returned but the dashboard continues to function. Light trace level includes Minimal trace logging.
 - **Minimal** is the default setting and records only fatal errors. You can set the trace level back to minimal after collecting a specific activity sequence. Even if a different trace level was set before logout, the trace is always reset to the lowest level the next time you log in.

Results

The trace is adjusted to the level chosen for this and all subsequent dashboards selected. To keep communications traffic to a minimum, the log messages are transferred in batches to the Dashboard Application Services Hub and, finally, after you log out, whether manually or after a time-out period. (If the browser fails, no final logging is sent.)

The logs are saved on the server computer and named *userid*.log.0 where *userid* is the ID used to log in to the Dashboard Application Services Hub and "0" is the first log. Three log files of 750 KB total are used to record trace data in a cyclical manner: After the *userid*.log.0 reaches 250 KB, log entries are saved to *userid*.log.1; after *userid*.log.1 reaches 250 KB, log entries go to *userid*.log.2 until it reaches the maximum, at which time *userid*.log.0 is cleared and new entries are saved there.

Sharing the current dashboard page link

You can use the **Copy URL** feature to share the current dashboard page link and access the content of the dashboard page in another browser.

Procedure

- 1. From the Dashboard for VMware Clusters or Dashboard for VMware ESX Servers, go to a dashboard page, for example, the Clusters group page.
- 2. On the upper-right corner, click Actions > Copy URL. A Share this page window opens that displays the Link to the current page link.
- **3**. Right-click the **Link to the current page** link, and copy the link location or select another option to copy the link.
- 4. Paste the link where you want to share it.

Dashboard Health Checks for VMware

You can use the Dashboard Health Checks for VMware to view the overall health of the data connections in your system.

Accessing the Dashboard Health Checks for VMware

Log on to the Tivoli Common Reporting or Dashboard Application Services Hub in your web browser to access the Dashboard Health Checks for VMware.

Procedure

In the navigation tree, select **System Status and Health** > **Dashboard Health Checks** to open the Dashboard page.

Health Checks for VMware

You can check the status of the dashboard components by using the Dashboard Health Checks page for VMware.

The Tivoli Monitoring components show health checks for the installed dashboards. You can check the status of the dashboard components by using the Dashboard Health Checks page. For each component, the table displays the name of the check, its status, and its description. Green, Red, Yellow, or Gray status is displayed against each check.

Table 24. Health checks for Tivoli Monitoring

Component	Check	Description
<itm dashboard="" data<br="">Provider ID> Note: The Provider ID value ITMSD must be configured in the IBM Infrastructure Management Dashboard.</itm>	Summary	Displays the overall status of this component.
	Dashboard Data Provider Available	Displays Green status if the dashboard data provider is available at version 6.3 or later.Displays Gray status if the required information cannot be retrieved.
	Dashboard Data Provider Enabled	Displays Green status if the dashboard data provider is enabled.Displays Gray status if the required information cannot be retrieved.
	TEPS Ping	Displays green status if the host specified for the Tivoli Enterprise Portal Server responds to a network ping.Displays Gray status if the required information cannot be retrieved.

For each component, the table displays the name of the check, its status, and its description. Green, Red, Yellow, or Grav status is displayed against each check

Table 25. Health checks for Virtual Infrastructure Monitoring		
Component	Check	Description
Dashboard for VMware	Summary	Displays the overall status of this component.
	VMware VI Agent Status	 Displays Green status if at least one VMware VI agent of supported version is online at this moment. Displays Gray status if the required information cannot be retrieved.
	VMware VI Agent Support	• Displays Green status if VMware VI agent support is installed on Tivoli Enterprise Portal Server and its available version is supported by Dashboard for VMware.
		• Displays Gray status if the required information cannot be retrieved.
	NetApp Agent Support	 Displays Green status if NetApp agent support is installed on Tivoli Enterprise Portal Server. Displays Grav status if the required information cannot be retrieved.
	NetApp Agent Status	 Displays Green status if at least one NetApp agent is online at this moment.
		• Displays Gray status if the required information cannot be retrieved.
	TPC Agent Support	Displays Green status if TPC agent support is installed on Tivoli Enterprise Portal Server.

Enterprise Portal Server.

Displays Gray status if the required information cannot be retrieved.

Displays Gray status if the required information cannot be retrieved.

Displays Green status if at least one TPC agent is online at this

•

•

•

moment.

TPC Agent

Status

Tab

Table 25. Health checks for Virtual Infrastructure Monitoring (continued)

Component	Check	Description
	TADDM Hostname Ping	 Displays Green status if the host specified for theTivoli Application Dependency Discovery Manager responds to a network ping. Displays Gray status if the required information cannot be retrieved.

178 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 9. VMware VI agent: Using Performance and Capacity Management Reports

Depending on the performance and capacity management report, there are specific prerequisites that apply.

The Performance and Capacity Management Reports component includes model and reports for the following agents:

- IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI
- IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines

The Performance and Capacity Management Reports component includes model for IBM Tivoli Monitoring for Virtual Environments Agent for NetApp Storage.

Other IBM Tivoli Monitoring for Virtual Environments agents include Cognos-based reports designed for use with Tivoli Common Reporting. For more information, See the chapter named "Tivoli Common Reporting" in the following agent user's guides for these reports:

- IBM Tivoli Monitoring for Virtual Environment Agent for Cisco UCS
- IBM Tivoli Monitoring for Virtual Environments Agent for Citrix XenApp
- IBM Tivoli Monitoring for Virtual Environments Agent for Citrix XenServer

For complete documentation of the Tivoli Common Reporting tool, see Tivoli Common Reporting Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ ic-home.html).

Report packages

The Tivoli Monitoring for Virtual Environments V7.2 reports are historical reports, reporting against summarized data that is collected in Tivoli Data Warehouse V6.2.2 or later.

The following databases for Tivoli Data Warehouse are supported for all reports:

- DB2
- Oracle
- SQL Server

The reports can be administered, run, and edited by Tivoli Common Reporting V2.1.1 software included with IBM Tivoli Monitoring V6.2.3.

Tivoli Common Reporting V3.1 is included with Tivoli Monitoring for Virtual Environments V7.2. This version of Tivoli Common Reporting includes Cognos Business Intelligence and Reporting V10.2.

The predefined reports that are run against VMware VI agent V7.2 only. The package contains 36 predefined reports and Cognos data models for the VMware VI agent and the NetApp Storage agent. Custom reports can be built for both agents.

Performance and Capacity Management report descriptions

The reports are organized into five groups. Each report contains parameters, uses tables or views, and provides output that is related to the purpose of the report.

• Accounting

The VMware VI Number of DCs Clusters Hosts and VMs Monitored report displays information, such as the data centers that are monitored.

• Performance trends and resource forecasts

This set of reports shows resource usage over time for various levels of the virtualized environment. The reports show trends at the cluster, host server, and virtual machine levels. Some reports show linear forecasts in addition to historical trends. The reports can be run with different date ranges, different summarization types (daily or hourly) and with varied thresholds. These reports are useful in determining trends, patterns, and forecasts, and in doing comparisons across the environment. Some of the reports are interactive and allow drill-through to other reports for more problem determination.

- Cluster Forecast Alerts

Calculates the forecast for resources that are used by all the clusters and alerts you if any of the clusters reach the specified thresholds in the next 30 days

- Cluster Performance Trends

Shows a historical trend of all the key performance metrics in percentages for one or more clusters

- Cluster Weekly Comparison

Compares the key attributes for clusters between two weeks and shows if any significant changes in the values occurred between the week

- Cluster Workload Trend and Forecast

Displays trend charts with historical and forecasted values of the key metrics for one or more clusters

- Data Store Capacity and Performance Trend

Shows the capacity, performance trends, and forecasting for one or more data stores in the environment

- Host Server CPU Ready and Utilization

Displays two line charts, which show CPU utilization of the host server, the number of virtual machines on the host and the percentage of this number that are ready to be allocated CPU

- Host Server Forecast Alerts

Calculates the forecast for resources that are used by all the host servers in one or more clusters and alerts you if any of the servers reach the specified thresholds in the next 30 days

- Host Server Heat Chart

Shows patterns of CPU, memory utilization, or both CPU and memory utilization over a period of time for the selected servers

- Host Server Memory Allocation Details

Shows the clusters, host servers, and virtual machines that consume the most or least resources (CPU, memory, storage)

- Host Server Performance Trends

Shows a historical trend of all the key performance metrics in percentages for one or more host servers

- Host Server Utilization Details

Shows a historical trend of the average (or maximum or minimum) daily (or hourly or weekly or monthly) CPU utilization in percentages for one or more host servers

- Host Server Workload Trend and Forecast

Displays trend charts with historical and forecasted values of the key resources for one or more host servers

- Host Servers Weekly Comparison

Compares the key attributes for host servers between two weeks and shows if any significant changes in the values occurred between the weeks

- VM CPU Ready and Utilization

Shows the CPU performance for one virtual machine

- VM Heat Chart

Shows patterns of CPU, memory utilization, or both CPU and memory utilization over a period of time for the selected virtual machines

- VM Performance Trends

Shows a historical trend of all the key performance metrics in percentages for one or more virtual machines

- VM Utilization Details

Shows a historical trend of the average (or maximum or minimum) daily (or hourly or weekly or monthly) CPU utilization in percentages for one or more virtual machines

· Prerequisites checking

This report provides a prerequisite scanner that checks whether the Tivoli Data Warehouse contains the list of all tables and views that are needed to run the predefined reports, and those needed to support custom reporting.

• What-if analysis for capacity estimation

This set of reports provides you with a simple what-if analysis tool to determine an approximate figure of how many more workloads your environment can accommodate and what more resources you would need. You can modify several parameters to obtain the wanted output. The capacity estimation reports consider historical usage data that is collected by the VMware VI agent and user inputs. The calculations are also based on forecasted data.

- Number of Workloads for Clusters

Performs an analysis to determine the number of additional virtual machines that can be placed on a cluster, which is based on the historical usage or user inputs

- Number of Workloads for Clusters or Host Servers

Performs an analysis to determine the number of additional virtual machines that can be placed on a cluster or a group of host servers, which are based on the historical usage or user inputs

- Resources Needed for Additional Workloads on Clusters

Performs an analysis to determine the resources (CPU, memory, storage, and so on) that you need, to add virtual machines to a cluster

- Resources Needed for Additional Workloads on Host Servers
 Performs an analysis to determine the resources (CPU, memory, storage, and so on) that you need, to add virtual machines to a server or a group of servers
- Workload right sizing and balancing

This set of reports is useful for determining the overall performance of the environment.

- Balanced and Unbalanced Clusters

Shows the balanced and unbalanced clusters in the environment

- Balanced and Unbalanced Host Servers

Shows the balanced and unbalanced host servers in one or more clusters

- Bottom N VMs by Host Server
 Shows the bottom N virtual machines by CPU or memory utilization on one or more selected host servers on a cluster
- Bottom N VMs by Resource Pool
 Shows the bottom N virtual machines by CPU and memory utilization on one or more resource pools
- Host Server Memory Allocation

Shows the average memory that is allocated to virtual machines on the host servers in comparison with the actual physical memory on the host

- Network Usage by VMs

Shows, which virtual machines have physical network interface cards that consume the most, or least, network I/O, and in addition, the amount they consume

- Top N VMs by Host Server
 Shows the top N virtual machines by CPU or memory utilization on one or more selected host servers on a cluster
- Top N VMs by Resource Pool
 Shows the top N virtual machines by CPU and memory utilization on one or more resource pools on a cluster
- Top or Bottom Host Servers Weekly Comparison

Compares the key attributes for top or bottom host servers between two weeks and shows if any significant changes in the values occurred between the weeks

- Top or Bottom Physical NICs
 - Shows the top or bottom N physical network interface cards for selected virtual machines
- Top or Bottom Virtual Switches by Network IO
 Shows the switches that consume the most, or least, network I/O, and in addition, the amount they consume
- Top or Bottom Workload Consumers Clusters
 Shows the top or bottom clusters in the environment by CPU, memory, or storage utilization
- Top or Bottom Workload Consumers Host Servers
 Shows the top or bottom host servers in the environment by CPU, memory, or storage utilization

The reports use the following attribute groups:

- KVM_SERVER (Daily, Hourly, Weekly and Monthly Views)
- KVM_SERVER_DATASTORE_DV
- KVM_SERVER_DISK (Daily, Hourly, Weekly and Monthly Views)
- KVM_SERVER_NETWORK (Daily, Hourly, Weekly and Monthly Views)
- KVM_VM_MEMORY (Daily, Hourly, Weekly and Monthly Views)
- KVM_VM_NETWORK (Daily, Hourly, Weekly and Monthly Views)
- KVM_VIRTUAL_MACHINES (Daily, Hourly, Weekly and Monthly Views)
- KVM_CLUSTERS (Daily, Hourly, Weekly and Monthly Views)
- KVMVMDSUTL (Daily, Hourly, Weekly and Monthly Views)
- KVMCLTRRPS_DV
- KVM_DATASTORES (Daily, Hourly, Weekly and Monthly Views)
- KVM_SERVER_VIRTUAL_SWITCHES_DV
- KVM_SERVER_MEMORY (Daily, Hourly, Weekly, and Monthly Views)

Prerequisites for running the Performance and Capacity Management reports

This topic outlines prerequisites for the Performance and Capacity Management reports.

The reports are developed on a data model that is based on IBM Tivoli Monitoring for Virtual Environments V7.2 and later only. The reports do not run with earlier versions of IBM Tivoli Monitoring for Virtual Servers.

Historical Data collection and aggregation

Historical data collection and aggregation is a key component of the infrastructure of IBM Tivoli Monitoring.

See the information center for managing historical data (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/adminuse/history_manage_intro.htm) for more information about how history can be collected and summarized over time.

The data that is collected by the VMware agent can be aggregated at the following levels of hierarchy:

- Data center
- Cluster
- Host server
- Virtual machine
- Data store
- Network information center

A huge amount of performance metrics is collected. These metrics are pulled into a data model to make it more usable. Most of the preconfigured capacity reports are focused on the following key metrics:

- CPU usage
- · Percent ready time
- Memory consumption
- Data store space usage
- Disk I/O
- Network I/O
- Number of virtual machines on the host

For the preconfigured reports to work, you must enable historical collection on the following attribute groups.

Table 26. Attribute groups where historical collection must be enabled for preconfigured reports to work

Attribute groups	Summarization type
KVM CLUSTERED RESOURCE POOLS	DAILY
KVM CLUSTERS	DAILY, HOURLY, WEEKLY, MONTHLY
KVM DATASTORES	DAILY, HOURLY, WEEKLY, MONTHLY
KVM SERVER	DAILY, HOURLY, WEEKLY, MONTHLY
KVM SERVER DATASTORE	DAILY
KVM SERVER DISK	DAILY, HOURLY, WEEKLY, MONTHLY
KVM SERVER NETWORK	DAILY, HOURLY, WEEKLY, MONTHLY
KVM VIRTUAL MACHINES	DAILY, HOURLY, WEEKLY, MONTHLY
KVM SERVER VIRTUAL SWITCHES	DAILY
KVM VM DATASTORE UTILIZATION	DAILY, HOURLY, WEEKLY, MONTHLY
KVM VM MEMORY	DAILY, HOURLY, WEEKLY, MONTHLY
KVM VM NETWORK	DAILY, HOURLY, WEEKLY, MONTHLY
KVM SERVER MEMORY	DAILY, HOURLY, WEEKLY, MONTHLY

Note: Hourly and daily summarization intervals must be enabled. Although reports can be run against weekly and monthly historical data, you are not required to enable those summarization intervals unless you plan to run reports using those intervals.

To use the data model for custom reporting, you must enable historical collection on the following attribute groups for the daily, hourly, weekly, and monthly summarization types.

Table 27. Attribute groups where historical collection must be enabled for custom reporting

Attribute groups
KVM SERVER
KVM SERVER DATASTORE
KVM SERVER DISK
KVM SERVER NETWORK
KVM VM MEMORY
KVM VM NETWORK
KVM VIRTUAL MACHINES
KVM CLUSTERS
KVM VM DATASTORE UTILIZATION
KVM CLUSTERED RESOURCE POOLS
KVM DATASTORES
KVM CLUSTERED VIRTUAL MACHINES
KVM DISTRIBUTED VIRTUAL PORTGROUPS
KVM DISTRIBUTED VIRTUAL SWITCH
KVM NETWORKED VIRTUAL MACHINES
KVM CLUSTERED DATASTORES
KVM CLUSTERED SERVERS
KVM NETWORKED SERVERS
KVM NETWORKS
KVM RESOURCE POOL CPU
KVM RESOURCE POOL GENERAL
KVM RESOURCE POOL MEMORY
KVM SERVER CPU
KVM SERVER SAN
KVM SERVER VIRTUAL SWITCHES
KVM VIRTUAL MACHINES
KVM VM CPU
KVM VM DISK
KVM VM PARTITION
KVM SERVER MEMORY

For the definition of each attribute, see the Attributes reference chapter of the VMware VI agent User's Guide that is located at IBM Tivoli Monitoring for Virtual Environments Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.tivoli.itmvs.doc_7.2.0.2/welcome_ve72fp2.htm).

Historical data is collected and aggregated at various levels. The reports make use of hourly and daily data. Some reports include the option of using weekly and monthly data if required. For best practices on how to set up the warehouse, see Historical Collections Best Practices in Tivoli Monitoring 6.2.2 (https://www.ibm.com/developerworks/wikis/display/tivolimonitoring/ Historical+Collections+Best+Practices+in+Tivoli+Monitoring+6.2.2).

The following table contains some examples of the types of attributes that can be collected and the intervals that can be used for collection.

Table 28. Attribute groups and collection intervals

Attribute group type	Collection and warehouse interval
Processor and memory utilization	• 5- or 15-minute interval; retain for 2 weeks
	Hourly summarization; retain for 3 months
	• Daily summarization; retain for 1 year
Disk I/O	5-minute interval; retain for 2 weeksHourly summarization; retain for 3 months
Volume usage	• Daily summarization; retain for 1 year
Network adapter rates	 5-minute interval; retain for 2 weeks Daily summarization; retain for 3 months
Process data	• 15-minute interval; retain for 1 to 2 weeks

Accounting

The VMware VI Number of DCs Clusters Hosts and VMs Monitored report displays information, such as the Data centers that are monitored.

Report structure	Description
Name	VMware Hypervisor: Number of DCs, Clusters, Hosts and VMs monitored
Description	This report displays the number of Data centers, Clusters, Hosts, Virtual Machines, and Processor Cores that are monitored by the VMware VI agent for the VMware Hypervisor.
Purpose	This report gives the count of the number of Data centers, Clusters, Host servers, and Virtual machines that the VMware agent is monitoring. In addition, this report gives number of cores monitored per server. As Smart Cloud Monitoring is priced on a Virtual Server or Virtual Machine basis, instead of the traditional Resource Value Unit (RVU), this report helps user to know what to pay or charge for.
Tables or views used	KVM_SERVER KVM_VIRTUAL_MACHINES
Output	 This report displays a table that contains the following data: Number of Datacenters being monitored Number of Clusters being monitored Number of Host Servers being monitored Number of Virtual Machines being monitored This table is followed by another table that displays the servers being monitored and the number of processor cores on each server.

Table 29. VMware VI Number of DCs Clusters Hosts and VMs Monitored report

Performance trends and resource forecasts

These reports show resource usage over time for various levels of the virtualized environment. The reports show trends at the cluster, host server, and virtual machine level. Some reports show linear forecasts in addition to historical trends. You can run the reports with different date ranges, with different summarization types (daily or hourly), and with varied thresholds. You can use these reports to determine trends, patterns, and forecasts, and to compare across the environment. Some of the reports are interactive and allow drill-through to other reports for more problem determination.

Use case 1:

When do my clusters or host servers reach specified thresholds?

Table 30. VMware VI Cluster Forecast Alerts report

Report structure	Description		
Name	VMware VI: Cluster Forecast Alerts		
Description	This report calculates the forecast for resources that are used by all the clusters, and alerts you if any of the clusters reach the specified thresholds in the next 30 days. This report can be run on a schedule regularly with user-defined thresholds. If the thresholds are exceeded, the corresponding cluster row is highlighted. You can click a cluster to see more details about the trend. The values are given in percentage.		
Purpose	Determines whether any of the clusters reach the specified thresholds in the next 30 days.		
Parameters	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from the list of clusters that belong to the selected data center. Date Range for computing forecast You can choose a date range for computing the virtual machine profile. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days Forecast Period (Days) This period is the number of days in the future that the forecast will be calculated for. The default is 30 days. Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1. Vacation Days If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled. Metric You can select one of the key metrics, or all, to forecast for. The key metrics are CPU, Memory, Storage, or Number of VMs per Host. Thresholds (%) You can choose the thresholds for CPU, Memory, and Storage to mark the clusters for which the usage is above this threshold level. The CPU, Memory, and Storage fields displays the following default values in percentage: CPU: 80 		
	Memory: 80 Storage: 80		
Tables or views	KVM_CLUSTERS_DV		
used	KVM_SERVER_DV		
Output	The report displays Average CPU, Data Store, and Memory usage for each of the selected clusters. Note: The rows marked in red show the clusters that are estimated to exceed the user-defined threshold.		

Table 31. V	'Mware VI	Cluster	Workload	Trend a	and	Forecast	report
-------------	-----------	---------	----------	---------	-----	----------	--------

Report structure	Description		
Name	VMware VI: Cluster Workload Trend and Forecast		
Description	This report displays trend charts with historical and forecasted values of the key metrics for one or more clusters. You can specify thresholds to see the trends in reference. The values are given in percentages. The forecasting is purely linear and based on least square regression.		
Purpose	Determines historical and forecasted values of key metrics.		
Parameters	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from the list of clusters that belong to the selected data center. Date Range for computing forecast You can choose a date range for computing the virtual machine profile. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days Forecast Period (Days) This period is the number of days in the future that the forecast will be calculated for. The default is 30 days. Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1. 		
	 Vacation Days If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled. Metric You can select one of the key metrics, or all, to forecast for. The key metrics are CPU, Memory, Storage, or Number of VMs per Host. Thresholds (%) You can choose the thresholds for CPU, Memory, and Storage. 		
Tables or views used	KVM_CLUSTERS_DV KVM_SERVER_DV KVM_SERVER_DISK_DV KVM_SERVER_NETWORK_DV		
Output	The report displays a daily historical trend and forecast for resource (CPU, Memory, Storage, Network, Disk, and Number of VMs per Host) usage for a cluster or a group of clusters.		

Table 32. VMware VI Host Server Forecast Alerts report

Report structure	Description
Name	VMware VI: Host Server Forecast Alerts
Description	This report calculates the forecast for resources that are used for all host servers in one or more clusters, and alerts you if any of the servers reach the specified thresholds in the next 30 days. This report can be run on a schedule regularly with user-defined thresholds. You can click a host server to see more details about the trend. The values are given in percentage.
Purpose	Determines whether any of the servers reach the specified thresholds in the next 30 days.

Table 32. VMware VI Host Server Forecast Alerts report (continued)

Report structure	Description				
Parameters	Data Center				
	You can choose one data center from a list of data centers.				
	Cluster You can choose one cluster from a list of clusters that belong to the selected data				
	center.				
	Host Servers				
	You can choose one or more nost servers from a list of nost servers that belong to the				
	Date Range for computing forecast				
	You can choose a date range for computing the virtual machine profile. The following are the options for Date Range selection: • Last 7 Days • Last 30 Days				
	• Last 90 Days				
	Forecast Period (Days)				
	This period is the number of days in the future that the forecast will be calculated for. The default is 30 days.				
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.				
	Vacation Days				
	or 2 if the vacation period is enabled				
	Metric You can select one of the key metrics, or all, to forecast for. The key metrics are CPU, Memory, Storage, or Number of VMs per Host.				
	Thresholds (%)				
	You can choose the thresholds for CPU, Memory, and Storage to mark the clusters for which the usage is above this threshold level. The CPU , Memory , and Storage fields displays the following default values in percentage:				
	• CPU: 80				
	• Memory: 80				
	• Storage: 80				
Tables or views used	KVM_SERVER_DV				
Output	This report alerts you on when a server or group of selected servers are expected to reach capacity limitations. The report calculates a linear trend, and determines whether any of the servers are expected to exceed the user-defined threshold for the CPU, Memory, and Storage of the server.				

Table 33. VMware VI Host Server Workload Trend and Forecast report

Report structure	Description
Name	VMware VI: Host Server Workload Trend and Forecast
Description	This report displays trend charts with historical and forecasted values of the key resources for one or more host servers. You can specify thresholds to see the trends in reference. The values are in percentages. The forecasting is purely linear and is based on least square regression.
Purpose	Determines whether any of the servers reach the specified thresholds in the next 30 days.

Report structure	Description			
Parameters	Data Center You can choose one data center from a list of data centers.			
	Cluster You can choose one cluster from a list of clusters that belong to the selected data center.			
	Host Servers			
	You can choose one or more host servers from a list of host servers that belong to selected Cluster and data center.			
	Date Range for computing forecast			
	You can choose a date range for computing the virtual machine profile. The following are the options for date range selection: • Last 7 Days			
	Last 30 Days Last 90 Days			
	Forecast Period (Days) This period is the number of days in the future that the forecast will be calculated for. The default is 30 days.			
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.			
	Vacation Days If the vacation period is not enabled, the default value is -1. You can enter the value 1			
	or 2 if the vacation period is enabled.Metric You can select one of the key metrics, or all, to forecast for. The key metrics are CPU, Memory, Storage, or Number of VMs per Host.			
	Thresholds (%) You can choose the thresholds for CPU, Memory, and Storage.			
Tables or views used	KVM_SERVER_DV			
	KVM_SERVER_NETWORK_DV			
Output	This report shows a daily historical trend and forecast for resources (CPU, Memory, Storage, and Network) usage for a host server or a group of servers.			

Table 33. VMware VI Host Server Workload Trend and Forecast report (continued)

Use case 2:

Which host servers are overloaded by processor? That is, the virtual machines are either using more processor capacity or are waiting for a long time to get processor capacity from the host server.

Report structure	Description
Name	Host Server CPU Performance
Description	This report shows two line charts. One chart shows the processor usage of the host server, overall virtual machine percent ready for all the virtual machines on the host, and the number of virtual machines on that host. The second chart shows the overlaid line chart for virtual machine percent ready of all the virtual machines on the host.
Purpose	Shows the host servers that are overloaded by processor, that is, where the virtual machines are either using more processor capacity or are waiting for a long time to get processor capacity from the host server.

Table 34. VMware VI Host Server CPU Ready and Utilization report

Report structure	Description
Parameters	Date Range You can choose a date range to display data for. Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Host Servers You can choose one host server from the list of host servers that belong to the selected cluster and data center.
	Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly from the list.
Tables or views used	 KVM_SERVER (Daily, Hourly, Weekly, and Monthly views) KVM_VIRTUAL_MACHINES (Daily, Hourly, Weekly, and Monthly views) Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	 This report shows the following two line charts: The first chart shows the processor usage of the host server, overall virtual machine percent ready for all the virtual machines on the host, and the number of virtual machines on that host. All lines are overlaid on a single chart with multiple y-axes. Percent ready indicates how long the virtual machines must wait to be allocated processor capacity. Acceptable percent ready range is 0 - 5%. High virtual machine percent ready and high overall processor usage on the host server indicates that there is not enough processor capacity. If the percent ready is high and the number of virtual machines is high, then too many virtual machines are competing for processor capacity on that host. The second chart shows the overlaid line chart for virtual machine percent ready of all the virtual machines on the host. By using this chart, it is easy to identify which virtual machines must wait longest to obtain processor capacity from the host. Where spikes occur on the graphs, you can click the line to see a weekly comparison report of the server. This shows whether any significant changes in the environment might have caused the spike.

Table 34. VMware VI Host Server CPU Ready and Utilization report (continued)

Table 35. VMware VI VM CPU Ready and Utilization report

Report structure	Description
Name	VM CPU Performance
Description	This report shows the processor performance for one virtual machine. The report compares the virtual machine percent ready, processor usage of the virtual machine, overall processor usage on the host and number of virtual machines on the host.
Purpose	Shows the processor performance for one virtual machine.

Report structure	Description
Parameters	 Date Range You can choose a date range for which to display data. Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Host Servers You can choose one host server from a list of host servers that belong to the selected cluster and data center.
	Virtual Machines You can choose one Virtual Machine from a list of virtual machines that belong to the selected data center, cluster, and host server.
	Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly from the list.
Tables or views used	 KVM_SERVER (Daily, Hourly, Weekly, and Monthly views) KVM_VIRTUAL_MACHINES (Daily, Hourly, Weekly, and Monthly views) Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	This report shows a line chart. The chart compares the VM percent ready, CPU usage of the VM, overall CPU usage on the host and number of VMs on the host. Acceptable percent ready range is 0-5%. High VM percent ready, low host CPU utilization, and high number of VMs on the host indicate that too many virtual machines are competing for CPU on that host. Where spikes occur on the graph, you can click the line to see a weekly comparison report of the server to observe any significant changes in the environment might have caused the spike.

Table 35. VMware VI VM CPU Ready and Utilization report (continued)

Use case 3:

What are the historical performance trends for processor, memory, storage, and network usage for one or more clusters or servers? Help me to compare the trends.

Table 36.	VMware	VI Cluste	r Performance	Trends report
-----------	--------	-----------	---------------	---------------

Report structure	Description		
Name	VMware VI: Cluster Performance Trends		
Description	This report shows a dashboard view of historical trends of all the key performance metrics for one or more clusters. The report displays the following values:		
	Average and maximum processor usage in percent		
	 Average and maximum memory usage in percent 		
	Average data store space usage against capacity in GB		
Network I/O in Mbps			
	• disk I/O in KBps		
	Average number of virtual machines per host		
Average number of effective host servers against number of hosts			
	Average number of virtual machines on against total number of virtual machines		
	The trends can be seen in an hourly or a daily format.		
Purpose	Shows historical trends of all key performance metrics for clusters.		

Table 36. VMware VI Cluster Performance Trends report (continued)

Report structure	Description		
Parameters	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from the list of clusters that belong to the selected data center. Summarization Type You can choose daily or hourly granularity of the trend. Display Option You can choose whether you want to see the trend chart or the summary table. 		
Tables or views used	 KVM_CLUSTERS (Daily, Hourly, Weekly, and Monthly views) KVM_SERVER_DV KVM_SERVER_DISK (Daily, Hourly, Weekly, and Monthly views) KVM_SERVER_NETWORK (Daily, Hourly, Weekly, and Monthly views) Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server. 		
Output	The report shows trends for multiple performance metrics of multiple clusters in a matrix. For percentages (CPU, Memory), the scales are 0 - 100. The yellow line indicates the maximum usage and the blue line indicates the average usage. Data store usage (shown as a solid line) is plotted against the total capacity (shown as a dashed line). The number of virtual machines that are in the ON state (shown as a solid line) is plotted against the total number of virtual machines on the cluster (shown as a dashed line). Network I/O is plotted in Mbps, and Disk I/O is plotted in KBps. Scale for both is relative to the maximum value for all the clusters. The other metrics that are shown are the number of virtual machines per host, and the number of effective host servers (shown as a solid line). The red dots show the peaks and the green dots show the minimum values.		

Table 37. VMware VI Host Server Performance Trends report

Report structure	Description
Name	VMware VI: Host Server Performance Trends

Report structure	Description
Description	This report shows a dashboard view of historical trends of all the key performance metrics for one or more host servers. The report displays the following values:
	Average and maximum processor usage in percent
	Average and maximum overall virtual machine percent ready
	Average and maximum memory usage in percent
	Data store space usage against capacity in GB
	Network I/O in mbps
	Disk I/O in kbps
	Number of virtual machines against total number of virtual machines
	The trends can be seen in an hourly or a daily format.
Purpose	Shows historical trends of all key performance metrics for host servers.
Parameters	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days
	Data Center
	You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center.
	You can choose one or more host servers that belong to the selected cluster.
	You can choose daily or hourly granularity of the trend.
	Display Option You can choose whether you want to view the trend chart or the summary table.
Tables or views	KVM_SERVER (Daily, Hourly, Weekly, and Monthly views)
used	KVM_SERVER_DISK (Daily, Hourly, Weekly, and Monthly views)
	KVM_SERVER_NETWORK (Daily, Hourly, Weekly, and Monthly views)
	Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	The report shows trends for multiple performance metrics of multiple host servers in a matrix. For percentages (CPU, Memory, VM Percent Ready), the scales are 0 - 100.
	The yellow line indicates the maximum usage and the blue line indicates the average usage. Data store usage (shown as a solid line) is plotted against the total capacity (shown as a dashed line).
	The number of virtual machines that are in the ON state (shown as a solid line) is plotted against the total number of virtual machines on the cluster (shown as a dashed line). Network I/O is plotted in mbps and its scale is relative to the maximum value for all the servers. The red dots show the peaks and the green dots show the minimum values.

Table 37. VMware VI Host Server Performance Trends report (continued)

Table 38. V	'Mware VI	' VM	Performance	Trends	report
-------------	-----------	------	-------------	--------	--------

Report structure	Description
Name	VMware VI: VM Performance Trends

 Table 38. VMware VI VM Performance Trends report (continued)

Report structure	Description
Description	This report shows a dashboard view of historical trends of all the key performance metrics for one or more virtual machines on one or more hosts. The report displays the following values:
	Average and maximum processor usage in percent
	Average and maximum virtual machine percent ready
	Average and maximum memory usage in percent
	Data store space usage against capacity in GB
	Network I/O in mbps
	Disk I/O in kbps
	The trends can be seen in an hourly or a daily format.
Purpose	Shows historical trends of all key performance metrics for virtual machines.
Parameters	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers that belong to the selected cluster. Virtual Machines You can choose one or more virtual machines that belong to the selected host servers. Summarization Type You can choose daily or hourly granularity of the trend.
	Display Options
Tables or views	KVM SERVER (Daily, House Workly, and Monthly views)
used	KVM_SERVER (Daily, Houriy, Weekly, and Monthly Views)
	KVM_SERVER_DISK (Daily, Hourly, Weekly, and Monthly views)
	KVM_SERVER_NETWORK (Daily, Hourly, Weekly, and Monthly views)
	KVM_VIRTUAL_MACHINES (Daily, Hourly, Weekly, and Monthly views
	KVMVMDSUTL (Daily, Hourly, Weekly, and Monthly views)
	KVM_VM_MEMORY (Daily, Hourly, Weekly, and Monthly views)
	KVM_VM_NEWORK (Daily, Hourly, Weekly, and Monthly views)
	Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.

Table 38. VMware VI VM Performance Trends report (continued)

Report structure	Description
Output	The report shows trends for multiple performance metrics of multiple host servers in a matrix. For percentages (CPU, Memory, VM Percent Ready), the scales are 0 - 100.
	Memory Utilization (%) can be greater than 100%, if Virtual Machine utilizes shared memory of Host Server.
	The yellow line indicates the maximum usage and the blue line indicates the average usage. Data store usage (shown as a solid line) is plotted against the total capacity (shown as a dashed line). The number of virtual machines that are in the ON state (shown as a solid line) is plotted against the total number of virtual machines on the cluster (shown as a dashed line).
	Network I/O is plotted in mbps and its scale is relative to the maximum value for all the servers. The red dots show the peaks and the green dots show the minimum values.

Use case 4:

What are the historical performance trends for processor, memory, storage, and network usage for one or more clusters, host servers, or virtual machines? Help me to compare the trends and identify patterns in usage and performance.

VMware VI Cluster Performance Trends report

See table for VMware VI Cluster Performance Trends report.

VMware VI Host Server Performance Trends report

See table for VMware VI Host Server Performance Trends report.

Table 39.	VMware	VI Host	Server	Heat	Chart	report
-----------	--------	---------	--------	------	-------	--------

Report structure	Description	
Name	Virtual Machine Host Server - Utilization Heat Chart	
Description	This report shows a historical trend of the average hourly processor usage and memory usage in percentage for one or more host servers.	
Purpose	Shows historical trend of average hourly processor and memory usage for one or more host servers.	
Parameters	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers that belong to the selected cluster. Date Range You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1. Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. 	

Table 39. VMware VI Host Server Heat Chart report (continued)

Report structure	Description
Tables or views used	KVM_SERVER (Daily and Hourly views)
Output	This report presents a different visualization for observing patterns in hourly processor or memory usage. This chart is called a heat chart.
	In a heat chart, the X-axis shows the hours during the day and the Y-axis shows the dates. For each host server, hourly averages for the metric are shown. This chart helps in identifying patterns, such as during which times of day the server gets busy. The chart is useful for determining maintenance schedules or observing whether the pattern of the virtual machine matches the pattern of its target host during virtual machine placement exercises. Different colors on the heat chart represent different percentage bands. You can modify the threshold values and colors for these bands.

Table 40. VMware VI VM Heat Chart report

Report structure	Description		
Name	Virtual Machine Utilization Heat Chart		
Description	This report shows a historical trend of the average hourly processor usage and memory usage in percentage for one or more virtual machines.		
Purpose	Shows historical trend of average hourly processor and memory usage for one or more virtual machines.		
Parameters	 Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers that belong to the selected cluster. VM You can choose one or more virtual machines that belong to the selected host servers. Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days 		
	 If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1. Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. 		
Tables or views used	KVM_VIRTUAL_MACHINES (Daily and Hourly views) KVM_VM_MEMORY_HV		

Table 40. VMware VI VM Heat Chart report (continued)

Report structure	Description
Output	This report presents a different visualization for observing patterns in hourly processor or memory usage. This chart is called a heat chart.
	In a heat chart, the X-axis shows the hours during the day and the Y-axis shows the dates. For each host server, hourly averages for the metric are shown. This chart helps in identifying patterns, such as during which times of day the server gets busy. The chart is useful for determining maintenance schedules or observing whether the pattern of the virtual machine matches the pattern of its target host during virtual machine placement exercises. Different colors on the heat chart represent different percentage bands. You can modify the threshold values and colors for these bands.

Use case 5:

What are the average, maximum, or minimum processor, memory, data store space usage, network I/O, and disk I/O, for a list of host servers over a period of time?

Table 41. VMware VI Host Server Utilization Details r	report
---	--------

Report structure	Description
Name	VMware VI: Host Server Utilization Details
Description	This report shows the average, maximum, or minimum processor usage, memory usage, data store space usage, network I/O, and disk I/O over a period of time for one or more host servers across a cluster in a data center. You can view the data in hourly, daily, weekly, or monthly format. You can drill down by clicking a data point on the chart to see more granular data. You can scroll to the next page of the report to see tabular data.
Purpose	Shows average, maximum, or minimum, values for a range of metrics for one or more host servers over time.

Report structure	Description
Parameters	Data Center
	You can choose one or more data centers from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data centers.
	Host Server
	You can choose one or more host servers that belong to the selected clusters.
	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days
	• Last 30 Days
	• Last 90 Days
	Metric You can choose the metric or resource to run the report on.
	You can choose average, maximum, or minimum, to aggregate the resources over a period of time.
	Summarization
	You can choose how the data is summarized for the resources. Data can be summarized on an hourly, or daily, or weekly, or monthly basis.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
Tables or views used	KVM_SERVER (Daily, Hourly, Monthly, and Weekly views)
	KVM_SERVER_DISK (Daily, Hourly, Monthly, and Weekly views)
	KVM_SERVER_NETWORK (Daily, Hourly, Weekly, and Monthly views)
	Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	The report shows a line chart, plotting the average percentage CPU or host memory usage for the selected host servers, on all the servers on which they are located, over the selected date range. The report also shows the results in table form.

Table 41. VMware VI Host Server Utilization Details report (continued)

Table 42. VMware VI Host Server Memory Allocation Details report

Report structure	Description
Name	VMware VI: Host Server Memory Allocation Detail
Description	This report shows the average, maximum, or minimum memory usage over a period of time for one or more host servers across a cluster in a data center. You can view the data in hourly format. You can scroll to the next page of the report to see tabular data.
Purpose	Shows average, maximum, or minimum, values for memory usage for one or more host servers over time.

Table 42. VMware VI Host Server Memory Allocation Details report (continued)

Report structure	Description
Parameters	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data centers. Host Server You can choose one host server that belong to the selected clusters.
Tables or views used	KVM_SERVER_MEMORY_DV KVM_VIRTUAL_MACHINES_DV
Output	The report shows a line chart, plotting the average percentage host memory usage for the selected host servers, on all the servers on which they are located, over the selected date range. The report also shows the results in table form.

Table 43. VMware VI VM Utilization Details report

Report structure	Description
Name	Virtual Machine Utilization Details
Description	This report shows the average, maximum, or minimum processor usage, memory usage, data store space usage, network I/O, and disk I/O over a period of time for one or more virtual machines on host servers across a cluster in a data center. You can view the data in hourly, daily, weekly, or monthly format. You can drill down by clicking a data point on the chart to see more granular data. You can scroll to the next page of the report to see tabular data.
Purpose	Shows average, maximum, or minimum, values for a range of metrics for one or more host servers over time.

Table 43. VMware VI VM Utilization Details report (continued)

Report structure	Description
Parameters	Data Center
	You can choose one or more data centers from a list of data centers.
	Cluster You can choose one or more clusters from a list of clusters that belong to the selected
	data centers.
	You can choose one or more host servers that belong to the selected clusters
	VM Name
	You can choose one or more virtual machines from a list of virtual machines that
	belong to the selected data centers, clusters, and host servers.
	Date Range
	You can choose a date range for which the trend is to be generated. The following are the options for date range selection:
	• Last 7 Days
	• Last 30 Days
	• Last 90 Days
	Metric You can choose the metric or resource to run the report on.
	Aggregation
	neriod of time
	Summarization Type
	You can choose how the data is summarized for the resources. Data can be summarized on an hourly, or daily, or weekly, or monthly basis.
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.
	Vacation Days
	If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled.
Tables or views	KVMVMDSUTL (Daily, Hourly, Weekly, and Monthly views)
useu	KVM_VIRTUAL_MACHINES (Daily, Hourly, Weekly, and Monthly views)
	KVM_VM_MEMORY (Daily, Hourly, Weekly, and Monthly views)
	KVM_VM_NETWORK (Daily, Hourly, Weekly, and Monthly views)
	Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	The report shows a line chart, plotting the average percentage CPU or host memory usage for the selected virtual machines, on all the servers on which they are located, over the selected date range. The report also shows the results in table form.

Use Case 6:

What are the historical performance trends and capacity for all data stores in the environment?

Table 44. VMware VI Data Store Capacity and Performance Trend report

Report structure	Description
Name	VMware VI: Data Store Trend and Forecast Report
Description	This report shows the capacity and performance trends for one or more data stores in the environment.

Report structure	Description
Purpose	Shows historical trends of all key performance metrics for data stores.
Parameters	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for date range selection: Last 7 Days Last 30 Days Last 90 Days Data Center You can choose one data center from a list of data centers. Data Store You can choose one or more data stores.
	Summarization Type Select the summarization type, such as Hourly, Daily, Weekly, Monthly from the list.
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.
	 Vacation Days If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled. For display options, you can choose Trend Charts or Forecast Charts to see either trending or trending with forecasting. Forecasting is added only for Total I/O and % Used.
Tables or views	KVM_DATASTORES (Daily, Hourly, Weekly, and Monthly views)
used	Note: The report supports Weekly and Monthly Summarization types. If you do not plan to run the report for these summarizations, do not configure these summarizations for the attribute groups on Tivoli Enterprise Portal Server.
Output	The report displays the following charts:
	• The first chart is a stacked bar chart. When the display option is Trend Charts, used space and free space are combined to show the total space in GB.
	• The second chart shows the space that is used in percent trend lines over time.
	• The third chart shows the total I/O in kbps trend over time.
	• The fourth chart shows the number of connected virtual machines over time.
	When the display option is Forecast Charts, the first chart shows the total I/O in KB per second trending and forecasting over time. The second chart shows the space that is used in percent trending and forecasting over time.
	In addition, this report displays the following tables in the list format:
	• The first table shows the number of connected clusters over time.
	• The second table shows the number of connected host servers over time.
	• The third table shows the space that is provisioned for virtual machines over time in GB.

Table 44. VMware VI Data Store Capacity and Performance Trend report (continued)

Use case 7:

Were there any significant changes in my environment between two weeks?

Table 45. VMware VI Cluster Weekly Comparison report

	2
Report structure	Description
Name	VMware VI: Clusters Weekly Comparison

Report structure	Description
Description	This report compares the key attributes for clusters between two weeks and shows whether significant changes occurred in the values between the weeks. You can choose the percent change in values. The default value is 20, which indicates that if a change occurred of more than 20% in the value of any attribute between the weeks, then that field is highlighted.
Purpose	Compares key attributes for clusters between weeks.
Parameters	Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center. Percent Change You can choose the Percent Change value, which indicates that if change occurred more than this for any attribute between the weeks, then the field is highlighted. Date Range You can choose a date range for which the trend is to be generated. You can either choose a historical period for which the trend is to be generated, for example: Last 7 Days Last 30 Days Last 90 Days or you can enter the start week and end week for which you want to compare attributes. Known issue: You might need to enter the default value of start week or end week, even if you choose a predefined date range such as Last 30 days. If you use the date range, then choose 1 for the start week and 52 for the end week. If you want to choose specific weeks, select the start and end week and year that you want the report for.
Tables or views used	KVM_CLUSTERS_DV KVM_SERVER_DV KVM_SERVER_DISK_DV KVM_SERVER_NETWORK_DV
Output	The report displays a table, with comparisons between weeks for each selected cluster. A percent change in values is highlighted if it is higher than the value specified. For configuration changes such as a change in the number of servers or vMotions, any change in value is highlighted.

Table 45. VMware VI Cluster Weekly Comparison report (continued)

Table 46. VMware VI Host Servers Weekly Com	parison rep	ort
---	-------------	-----

Report structure	Description
Name	VMware VI: Weekly Comparison of Key Metrics for Host Servers
Description	This report compares the key attributes for host servers between two weeks and shows whether significant changes occurred in the values between the weeks. You can choose the percent change in values. The default value is 20, which indicates that if a change occurred of more than 20% in the value of any attribute between the weeks, then that field is highlighted.
Purpose	Compares key attributes for host servers between weeks.

Table 46. VMware VI Host Servers Weekly Comparison report (continued)

Report structure	Description
Parameters	Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers that belong to the selected clusters.
	 Percent Change You can choose the Percent Change value, which indicates that if a change occurred more than this value for any attribute between the weeks, then the field is highlighted. Date Range You can choose a date range for which the trend is to be generated. You can either choose a historical period for which the trend is to be generated, for example: Last 7 Days Last 30 Days Last 90 Days
	or you can enter the start week and end week for which you want to compare attributes. Known issue: You might need to enter the default value of start week or end week, even if you choose a predefined date range such as Last 30 days . If you use the date range, then choose 1 for the start week and 52 for the end week. If you want to choose specific weeks, select the start and end week and year that you want the report for.
Tables or views used	KVM_SERVER_DV KVM_SERVER_DATASTORE_DV KVM_SERVER_DISK_DV KVM_SERVER_NETWORK_DV
Output	The report displays a table, with comparisons between weeks for each selected host server. A percent change in values is highlighted if it is higher than the value specified. For configuration changes such as a change in the number of virtual machines, any change in value is highlighted.

Prerequisites checking

These reports provide a prerequisite scanner that checks whether the Tivoli Data Warehouse contains the list of all tables and views that are needed to run the predefined reports, and those needed to support custom reporting.

The reports also direct you to appropriate documentation that can be helpful in the following ways:

- Enabling historical collection and summarization and pruning
- Creating IBM_TRAM schema, the Time Dimension, and other shared dimensions, such as WEEKDAY_LOOKUP, MONTH_LOOKUP, and ComputerSystem
- Populating the Time dimension

The reports also provide a list of attribute groups for the agent to guide you while enabling historical collection.

Note: Despite having all the prerequisite tables and views in the Tivoli Data Warehouse, you might not be able to run the reports because of insufficient data in the warehouse. For example, this can happen if adequate time stamps are not generated for the Time Dimension.

The documentation that is provided helps you to run the appropriate database scripts to populate the tables. However, before you run the prerequisite scanner reports, ensure that you have defined and tested the appropriate database connection (DB2, Oracle, or MS SQL Server) to the Tivoli Data Warehouse.

Use case:

- (a) While running a predefined report, an error message similar to the following is displayed: RQP-DEF-0177 An error occurred while performing operation 'sqlPrepareWithOptions' status='-201'. UDA-SQL-0196 The table or view "ITMUSER.KVM_SERVER_NETWORK_DV" was not found in the dictionary.
- (b) While creating an ad hoc query (custom report) in Query Studio, I see a blank screen while trying to drag a column from a table.

How do I find out whether my warehouse has all the tables and views needed to run a report or create a custom report?

Review the following reports to help answer the questions in the use case:

You can verify whether the reporting function is installed and configured correctly by running the sample report for the monitoring agent for VMware. The monitoring agent for VMware reporting package that you imported into Tivoli Common Reporting includes three sample reports. By using these sample reports, you can monitor the reporting activity and see what a typical Cognos report includes. These sample reports are available in the Common Reporting panel in Tivoli Common Reporting. The following tables describe these reports.

Report structure	Description
Name	VMware VI Prerequisite Scanner report
Description	This report runs on the DB2 or MS SQL Server or Oracle database to check whether all prerequisite tables and views are available on the Tivoli Data Warehouse.
Purpose	You can use this report to verify whether all the prerequisite tables and views that are required to run a predefined report or to create a custom report that is available on the Tivoli Data Warehouse.
Parameters	Database Type Select the required database type, such as DB2 or MS SQL Server or Oracle. Display Options Select the display option, such as Check all reports or Check a specific report. If you select the Check a specific report option, select the required report category. Reports Category Select the report category, such as Accounting or Performance Trends and Resource Forecasts. Reports Select the required report name.

Table 47. VMware Prerequisite Scanner report
Table 47. VMware Prerequisite Scanner report (continued)

Report structure	Description
Output	The report lists all the prerequisite tables and views and displays a status against each of them.
	A red cross (x) indicates that the table or view requires to run some of the predefined reports and is missing on the Tivoli Data Warehouse.
	A yellow exclamation mark (!) indicates that the table or view is missing on the Tivoli Data Warehouse. The missing view does not affect the running of the predefined reports. However, ad hoc reporting on the attribute group is not possible. You are directed to the documentation to enable historical collection and summarization and pruning in either case.
	This report displays a table with the following two columns:
	• The first column lists the prerequisite tables and views from VMware VI agent in the Tivoli Data Warehouse.
	• The second column lists the prerequisite tables for Tivoli Common Reporting shared dimensions.

What-if analysis for capacity estimation: estimate the capacity of your environment

This set of reports provides you with a simple what-if analysis tool. You can determine approximately how many more workloads your environment can accommodate and what additional resources you might need based on demand. You can modify several parameters to obtain the output you want. The capacity estimation reports take into consideration historical usage data collected by the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI and user inputs. The calculations are based on current environment characteristics and the forecasted data.

Use case 1:

How many more virtual machines can I place on a cluster, a host server, or a group of servers?

One of the most important questions that VMware VI administrators or capacity planners ask is how many more virtual machines they can place, based on the demand created by the virtual machines in their environment. This analysis can be done by using the following reports:

Report Structure	Description
Name	VMware VI Number of Workloads for Clusters
Description	Use this report to determine the number of additional virtual machines that can be placed on a cluster, based on the forecasted period and historical usage or user inputs.
Purpose	Determine the number of additional virtual machines that can be placed on a cluster based on the historical data and forecasted period for the next 30 days.

Table 48. VMware VI Number of Workloads for Clusters report

Table 48. VMware VI Number of Workloads for Clusters report (continued)

Report Structure	Description
Parameters	Data Center
	You can choose one data center from a list of data centers.
	Clusters
	You can choose one cluster from a list of clusters that belong to the selected data
	VM Profile
	A profile for virtual machines is computed, which is the amount of resources (CPU in
	GHz, memory in MB, and storage space in GB) that the virtual machines in the cluster
	consume. You can choose from the following three VM Profiles:
	Average
	machines in the cluster consume, averaged for the selected cluster over the
	range of time specified by the Date Range parameter.
	Peak In this profile, you use the maximum amount of resources that the virtual
	machines in the selected cluster consume over the range of time specified by
	the Date Kange parameter.
	In this profile, you use resource characteristics of the virtual machine that you
	enter.
	Buffer The buffer is the amount of resources that are not allocated based on user preferences.
	For example, if you do not want to allocate all your storage, you can specify how
	much space you want to leave unallocated. The analysis takes that figure into account.
	CPU: 2 GHz
	Data store Space: 30 GB
	• Memory: 256 MB
	• Disk I/O: 50%
	• Network 1/O: 50%
	User-defined Resource Usage for VM Profile
	You can use this parameter to define the VM profile in case you selected the VM
	parameter shows the following default values:
	• CPU (GHz): 2
	• Data Store Space (GB): 30
	• Mamory (MB): 256
	$\frac{1}{2} \frac{1}{2} \frac{1}$
	• Disk $1/O$ (Kbps): 10
	• Network 1/O (Mbps): 0.01 Date Range for computing VM Profile
	You can choose a date range for computing the virtual machine profile. The options for
	Date Range selection are as follows:
	Last 7 Days
	Last 30 Days
	• Last 90 Days
	The default value is Last 7 Days.
	Forecast Period (Days)
	davs.
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending
	on the peak and off-peak hours that are configured in the data warehouse. The daily
	table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1
	default value is -1
	Vacation Days
	If the vacation period is not enabled, the default value is -1. You can enter the value 1
	or 2 if the vacation period is enabled.

Table 48. VMware VI Number of Workloads for Clusters report (continued)

Report Structure	Description
Tables or views used	KVM_CLUSTERS_DV
	KVM_VIRTUAL_MACHINES_DV
	KVM_VM_MEMORY_DV
	KVMVMDSUTL_DV (virtual machine Data Store utilization)
	KVM_SERVER_DISK_DV
	KVM_SERVER_NETWORK_DV
	KVM_VM_NETWORK_DV
Output	The report is in table form. This reports displays two tables. The first table shows the data based on historical usage or user inputs and the second table shows data based on the forecasted period (default 30 days).
	• The first column shows the type of resource: namely CPU in GHz, Data Store Space in GB, Memory in MB, Network I/O in Mbps, and Disk I/O in Kbps.
	• The second column shows the VM profile based on average resource used by all virtual machines on this cluster. This column displays the resources consumed by all the deployed virtual machines on that cluster over the selected date range. The default date range is the last 7 days.
	• The third column shows the Available Cluster Capacity (before applying Buffer) . The Available Cluster Capacity is the amount of resources available as a whole for the cluster (the cumulative sum of the available capacities on all servers in the cluster) before applying the user-defined buffer. The available capacity is recorded for the last 48 hours.
	• The fourth column shows the Buffer that you enter. The buffer is to indicate the resources that you do not want to allocate.
	• The fifth column shows the Available Cluster Capacity after applying Buffer.
	This figure is calculated using the formula Available Cluster Capacity (before applying Buffer) – Buffer.
	• The sixth column shows the Number of VMs that can be placed on the cluster for each resource.
	This figure is calculated using the formula Available Cluster Capacity (after applying Buffer) / VM Profile.
	The constrained resources are highlighted in red.

VMware VI Number of Workloads for Clusters or Host Servers report

The key difference between this report and the VMware VI Number of Workloads for Clusters report is that in this report the calculations are done at a finer level of granularity, that is, at the host server level instead of looking at the aggregated values at the cluster level.

For example, the cluster might have 10 servers, each server having 5 MB of memory. Cumulatively the cluster has 50 MB of memory. Assume that a typical virtual machine consumes 6 MB of memory. Using the VMware VI Number of Workloads for Clusters report, the calculation yields 50 / 6 = 8 more virtual machines can be added to the cluster. However, if you look at each server individually, no more virtual machines can be added, because each server has only 5 MB of memory remaining.

From this report, you can see how many virtual machines can be accommodated by each host server individually in a cluster. Use the VMware VI Number of Workloads for Clusters report to assess how many virtual machines can be placed on the cluster if resources are moved and reallocated, using all available capacity. This report, in contrast, assesses how many virtual machines can fit on a cluster without changing the clusters. A combination of these two reports shows whether, by moving virtual machines, you can fit more virtual machines on the same cluster.

Report Structure	Description
Name	VMware VI Number of Workloads for Clusters or Host Servers
Description	You can use this report to determine the number of additional virtual machines that can be placed on a cluster or a group of host servers, based on the forecasted period and historical usage or user inputs. You can run this report for all the servers in a cluster, for selected host servers in the cluster, or for stand-alone servers.
Purpose	Determine the number of additional virtual machines that can be placed on a cluster or a group of host servers, based on the historical usage or user inputs and the forecasted period for the next 30 days.

Table 49. VMware VI Number of Workloads for Clusters or Host Servers report

Report Structure Description **Parameters** Data Center You can choose one data center from a list of data centers. **Cluster** You can choose one cluster from a list of clusters that belong to the selected data center. Host Servers You can choose one or more host servers from a list of host servers that belong to the selected Cluster and Data Center. **VM** Profile A profile for virtual machines is computed, which is the amount of resources (CPU in GHz, memory in MB, and storage space in GB) that the virtual machines in the cluster consume. You can choose from the following three VM Profiles: Average In this profile, you use the average amount of resources that the virtual machines in the cluster consume, averaged for the selected cluster over the range of time specified by the Date Range parameter. Peak (average) In this profile, you use the maximum amount of resources that the virtual machines in the selected cluster consume averaged over the range of time specified by the Date Range parameter. Peak (maximum) In this profile, you use the maximum amount of resources that the virtual machines in the selected cluster consume maximized over the range of time specified by the Date Range parameter. **User-defined** In this profile, you use resource characteristics of the virtual machine that you enter. Buffer The buffer is the amount of resources that are not allocated based on user preferences. For example, if you do not want to allocate all your storage, you can specify how much space you want to leave unallocated. The analysis takes that figure into account. The default buffer values are: • CPU: 2 GHz • Data Store Space: 30 GB Memory: 256 MB • Disk I/O: 50% Network I/O: 50% Note: This buffer is applied to each individual server, so to calculate the buffer, divide the total resources you do not want to allocate by the number of servers in the cluster. **User-defined Factor for VM Profile** You can use this parameter to define the VM profile in case the you selected the VM profile parameter as "User-defined". The User-defined Factor for VM Profile parameter shows the following default values: • CPU (GHz): 2 Data Store Space (GB): 30 Memory (MB): 256 • Disk I/O (KBps): 10 • Network I/O (Mbps): 0.01 **Include Peak Hours** You can choose the peak hours range.

Table 49. VMware VI Number of Workloads for Clusters or Host Servers report (continued)

Report Structure	Description
Parameters	Date Range for computing VM Profile You can choose a date range for computing the virtual machine profile. The options for Date Range selection are as follows: • Last 7 Days • Last 30 Days • Last 90 Days
	The default value is Last 7 Days. Forecast Period (Days) The number of days that the forecast will be calculated for. The default value is 30 days.
	Shifts If shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.
	Vacation Days If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled.
Tables or views used	KVM_SERVER_DV
	KVM_SERVER_DISK_DV
	KVM_SERVER_MEMORY_DV
	KVM_SERVER_NETWORK_DV
	KVM_VIRTUAL_MACHINES (Daily or Hourly View)
	KVM_VM_MEMORY (Daily or Hourly View)
	KVM_VM_NETWORK (Daily or Hourly View)
	KVMVMDSUTL (virtual machine data store utilization) (Daily or Hourly View)

Table 49. VMware VI Number of Workloads for Clusters or Host Servers report (continued)

Table 49. VMware VI Number of Workloads for Clusters or Host Servers report (continued)

Report Structure	Description
Output	Although the layout and functionality of this report are similar to the VMware VI Number of Workloads for Clusters report, this report provides a more granular capacity calculation at the host server level instead of the high-level cluster view.
	Each row in the table shows a host server. The servers are grouped by cluster.
	• The first column shows the type of resource: namely CPU in GHz, Data Store Space in GB, Memory in MB, Network I/O in Mbps and Disk I/O in Kbps.
	• The second column lists the cluster name.
	• The third column lists the host server name.
	• The fourth column shows the virtual machine profile, that is, the resources consumed by all the deployed virtual machines on the selected host server over the selected date range.
	• The fifth column shows the Available Capacity of the host server (before applying buffer). The Available Capacity is the amount of resources available on the single host server before applying the user-defined buffer. The available capacity is recorded for the last 48 hours.
	• The sixth column shows the Buffer that you enter. The buffer is to indicate the resources that you do not want to allocate on a host server.
	• The seventh column shows the Available Cluster Capacity after applying Buffer.
	This figure is calculated using the formula Available Cluster Capacity (before applying Buffer) – Buffer.
	• The eighth column shows the Number of VMs that can be placed on the cluster for each resource.
	This figure is calculated using the formula Available Cluster Capacity (after applying Buffer) / VM Profile.
	The constrained resource is highlighted in red.
	Note: A value of 999,999,999 for Number of VMs indicates that no limit exists on the number of virtual machines that can be added to the cluster. This value appears if the VM profile is 0.

Use case 2:

I want to place a number of workloads or virtual machines to a cluster or group of servers. How much additional resource, such as processor, memory, and storage, do I need?

This analysis is derived from the use case about the number of virtual machines needed. In this case, you know how many more virtual machines you want to add to the cluster or host server, but are not sure whether the available capacity is sufficient, or whether you will need additional resources. This analysis can be done by using the following reports:

Report Structure	Description
Name	VMware VI Resources Needed for Additional Workloads on Clusters
Description	You can use this report to determine the resources (CPU, memory, storage, and so on) that you need in order to add virtual machines to a cluster.
Purpose	Determine the amount of resources that you need in order to add virtual machines to a cluster based on the historical usage or user inputs and forecasted period (default value is 30 days).

Table 50. VMware VI Resources Needed for Additional Workloads on Clusters report

Report Structure	Description	
Parameters	Data Center	
	You can choose one data center from a list of data centers.	
	Cluster You can choose one cluster from a list of clusters that belong to the selected data	
	center.	
	A profile for virtual machines is computed which is the amount of resources (CPL) i	n
	GHz, memory in MB, and storage space in GB) that the virtual machines in the clust	er
	consume. You can choose from the following three VM Profiles:	
	Average	
	In this profile, you use the average amount of resources that the virtual machines in the cluster consume, averaged for the selected cluster over the range of time specified by the Date Range parameter.	
	Peak In this profile, you use the maximum amount of resources that the virtual	
	machines in the selected cluster consume over the range of time specified b	y
	the Date Range parameter.	
	User-defined	
	In this profile, you use resource characteristics of the virtual machine that y	ou
	• CPU: 2 GHz	
	• Data store space: 30 GB	
	• Memory: 256 MB	
	• Disk I/O: 50%	
	• Network I/O: 50%	
	Number of VMs to add to the cluster	
	how much additional resources you need to accommodate those virtual machines	
	Buffer The buffer is the amount of resources that are not allocated based on user preference	s.
	For example, if you do not want to allocate all your storage, you can specify how much space you want to leave unallocated. The analysis takes that figure into account	nt.
	User-defined Resource Usage for VM Profile	
	You can use this parameter to define the VM profile in case the you selected the VM profile parameter as "User-defined". The User-defined Resource Usage for VM Profil parameter shows the following default values:	e
	• CPU (GHz): 2	
	• Data Store Space (GB): 30	
	• Memory (MB): 256	
	• Disk I/O (KBps): 10	
	• Network I/O (Mbps) : 0.01	
	Date Range for computing VM Profile	
	You can choose a date range for computing the virtual machine profile. The options date range selection are as follows:	for
	Last 7 Days	
	• Last 30 Days	
	• Last 90 Days	
	The default value is Last 7 Days.	

Table 50. VMware VI Resources Needed for Additional Workloads on Clusters report (continued)

Table 50. VMware VI Resources Needed for Additional Workloads on Clusters report (continued)

Report Structure	Description
Parameters	 Forecast Period (Days) The number of days that the forecast will be calculated for. The default value is 30 days. Shifts If shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1. Vacation Days
	If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled.
Tables or views used	KVM_CLUSTERS_DV
	KVMVMDSUTL_DV (virtual machine data store utilization)
	KVM_VIRTUAL_MACHINES_DV
	KVM_VM_MEMORY
	KVM_SERVER_DISK_DV
	KVM_SERVER_NETWORK_DV
	KVM_VM_NETWORK_DV
Output	The report is in table form. This reports displays two tables. The first table shows the data based on historical usage or user inputs and the second table shows data based on the forecasted period (default 30 days).
	Memory in MB, Network I/O in Mbps, and Disk I/O in Kbps.
	• The second column shows the VM profile based on average resource used by all VMs on this cluster. This column displays the resources consumed by all the deployed virtual machines on that cluster over the selected date range. The default date range is the last 7 days.
	• The third column shows the Resources Needed by Additional VMs , that is, the amount of resources that are needed to accommodate the virtual machines you want to add.
	This figure is calculated using the formula Number of VMs to add to the cluster* VM Profile.
	• The fourth column shows the Available Cluster Capacity (before applying Buffer) . The Available Cluster Capacity is the amount of resources available as a whole for the cluster before applying the user-defined buffer. The available capacity is recorded for the last 48 hours.
	• The fifth column shows the Buffer that you enter.
	• The sixth column shows the Available Cluster Capacity after applying Buffer.
	This figure is calculated using the formula Available Cluster Capacity (before applying Buffer) – Buffer.
	The seventh column shows the Capacity Needed.
	This figure is calculated using the formula Resources Needed by Additional VMs - Available Cluster Capacity (after applying Buffer). This column shows a value of 0 if the value for Available Cluster Capacity (after applying Buffer) is greater than the value for Resources Needed by Additional VMs.
	Note: The rows marked in red show the clusters that need additional capacity in order to accommodate additional virtual machines.

Report Structure	Description
Name	VMware VI Resources Needed for Additional Workloads on Clusters
Description	You can use this report to determine the resources (CPU, memory, storage, and so on) that you need in order to add virtual machines to a server, or group of servers. This report is similar to the VMware VI Resources Needed for Additional Workloads on Clusters report, except the calculation is done at a more granular level, that is, by looking at individual host servers instead of clusters.
Purpose	Determine the amount of resources that you need in order to add virtual machines to a server, or group of servers based on the historical usage or user inputs and forecasted period (default value is 30 days).

Table 51. VMware VI Resources Needed for Additional Workloads on Host Servers report

Report Structure	Description
Parameters	Data Center
	You can choose one data center from a list of data centers.
	Cluster You can choose one cluster from a list of clusters that belong to the selected data
	center.
	You can choose one or more host servers from a list of host servers that belong to the
	selected Cluster and Data Center.
	VM Profile
	A profile for virtual machines is computed, which is the amount of resources (CPU in GHz, memory in MB, and storage space in GB) that the virtual machines in the cluster consume. You can choose from the following three VM Profiles:
	In this profile, you use the average amount of resources that the virtual machines in the cluster consume, averaged for the selected cluster over the range of time specified by the Date Range parameter.
	In this profile, you use the maximum amount of resources that the virtual machines in the selected cluster consume, averaged over the range of time specified by the Date Range parameter.
	Peak (maximum) In this profile, you use the maximum amount of resources that the virtual machines in the selected cluster consume maximized over the range of time specified by the Date Range parameter.
	User-defined
	In this profile, you use resource characteristics of the virtual machine that you
	enter.
	Number of VMs to add to each nost server This parameter is the number of virtual machines that you add to the cluster to see
	Buffer The buffer is the number of virtual intermites that you add to the cluster to see how much additional resources are needed to accommodate those virtual machines.Buffer The buffer is the amount of resources that are not allocated based on user preferences. For example, if you do not want to allocate all your storage, you can specify how much space you want to leave unallocated and the analysis takes that into account. The default buffer values are:
	 CPU: 2GHz Data Store Space: 30GB Memory: 256 MB Disk I/O: 50% Network I/O: 50%
	User-defined Resource Usage for VM Profile You can use this parameter to define the VM profile in case the you selected the VM profile parameter as "User-defined". The User-defined Resource Usage for VM Profile parameter shows the following default values:
	• CPU (GHz): 2
	• Data Store Space (GB): 30
	• Memory (MB): 256
	• Disk I/O (KBps): 10
	• Network I/O (Mbps): 0.01
	Include Peak Hours You can choose the peak hours range.

Table 51. VMware VI Resources Needed for Additional Workloads on Host Servers report (continued)

Report Structure	Description
Parameters	Date Range for computing VM Profile
	You can choose a date range for computing the virtual machine profile. The options for
	Date Range selection are as follows:
	• Last 7 Days
	• Last 90 Days
	Forecast Period (Days)
	The number of days that the forecast will be calculated for. The default value is 30
	days.
	 Shifts If shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1. Vacation Days
	If the vacation period is not enabled, the default value is -1. You can enter the value 1 or 2 if the vacation period is enabled.
Tables or views used	KVM_SERVER_DV
	KVM_SERVER_MEMORY_DV
	KVMVMDSUTL (virtual machine data store utilization) (Daily and Hourly View)
	KVM_VIRTUAL_MACHINES (Daily and Hourly View)
	KVM_VM_MEMORY (Daily and Hourly View)
	KVM_SERVER_DISK_DV
	KVM_SERVER_NETWORK_DV
	KVM_VM_NETWORK (Daily and Hourly View)

Table 51. VMware VI Resources Needed for Additional Workloads on Host Servers report (continued)

Table 51. VMware VI Resources Needed for Additional Workloads on Host Servers report (continued)

Report Structure	Description
Output	The report is in table form. This reports displays two tables. The first table shows the data based on historical usage or user inputs and the second table shows data based on the forecasted period (default 30 days).
	• The first column shows the type of resource: namely CPU in GHz, Data Store Space in GB, Memory in MB, Network I/O in Mbps, and Disk I/O in Kbps.
	• The second column lists the cluster name.
	• The third column lists the server host name.
	• The fourth column shows the VM profile , that is, the resources consumed by the virtual machines on that host server over the selected date range. The default date range is the last 7 days.
	• The fifth column shows the Resources Needed by Additional VMs , that is, the amount of resources that are needed to accommodate the virtual machines you want to add.
	This figure is calculated using the formula Number of VMs to add to the host server*VM Profile.
	• The sixth column shows the Available Capacity (before applying Buffer) . The Available Capacity is the amount of resources available as a whole for the cluster before applying the user-defined buffer. The available capacity is recorded for the last 48 hours.
	• The seventh column shows the Buffer that you enter.
	• The eighth column shows the Available Capacity after applying Buffer.
	This figure is calculated using the formula Available Capacity (before applying Buffer) – Buffer.
	The ninth column shows the Capacity Needed.
	This figure is calculated using the formula Resources Needed by Additional VMs - Available Capacity (after applying Buffer). This column shows a value of 0 if the value for Available Capacity (after applying Buffer) is greater than the value for Resources Needed by Additional VMs.
	Note: The rows marked in red show the clusters that need additional capacity in order to accommodate additional virtual machines.

Workload right sizing and balancing: determining the performance of your environment

You can use these reports to determine the overall performance of the environment.

Use case 1:

Which clusters, host servers, or virtual machines use the most or least resources (processor, memory, storage, network I/O, or disk I/O)? For example, what are the top 10 servers by processor usage, or bottom 10 virtual machines by memory usage less than 10% memory, or the top 10 clusters by processor usage in GHz?

Report Structure	Description
Name	VMware VI: Top or Bottom Workload Consumers - Clusters
Description	This report shows the clusters that use the most, or least, resources in the environment by processor, memory, storage, network, and disk usage. You can view the values in units such as processor capacity in GHz, memory in MB, data store space in GB, Network I/O in Mbps, Disk I/O in KBps, or in percentage.
Purpose	Shows the clusters that use the most, or least, resources (processor, memory, storage, network I/O, and disk I/O).

Table 52. VMware VI Top or Bottom Workload Consumers Clusters report

Report Structure Description **Parameters Data Center** You can choose one data center from a list of data centers. **Top or Bottom** You can choose whether you want to see most used or least used resources. Display You can choose whether the charts display real values (CPU in GHz, memory in MB, data store space in GB, network I/O in Mbps, and disk I/O in KBps) or percentages. Number of Resources You can choose the number of clusters to be displayed, for example the top 10 or bottom 5. **Date Range** You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: • Last 7 Days Last 30 Days · Last 90 Days Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1. Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. **Forecast Range** This period is the number of days in the future that the forecast will be calculated for. The default is 7 days. Use Threshold for Ranking You can choose whether the Threshold should be used or not. Thresholds for Rank (%) You can choose the upper and lower thresholds that are used to filter the percentage values. For example, you can see the top 10 clusters that use more than 80% CPU or the bottom 5 clusters that use less than 10% memory. KVM_SERVER_DV Tables or views used KVM_CLUSTERS_DV KVM_SERVER_DISK_DV KVM_SERVER_NETWORK_DV Output The report shows bar charts that display the clusters that use the most, or least, resources, and also the amount that they use. One bar chart is shown for each of CPU, memory, data store space, network, and disk usage. If you specify that the charts display real values, you can see the total processor capacity or memory used, the effective processor capacity, or memory available. This processor capacity or memory available s the processor capacity or memory available for allocation, and total memory available, which includes the processor capacity and memory unavailable for allocation, network I/O in Mbps and disk I/O in KBps. If you specify that the charts display percentages, only the used and free processor capacity, memory, data store space and network usage are shown. You can click a cluster and see the top or bottom host servers on that cluster. You can then click a host server to see the top or bottom virtual machines on the host. Alternatively you can click the cluster to see the historical trend of that attribute of the selected cluster over time.

Table 52. VMware VI Top or Bottom Workload Consumers Clusters report (continued)

Report Structure	Description
Name	VMware VI Top or Bottom Host Servers Weekly Comparison
Description	This report compares the key attributes for Top or Bottom N host servers for the past few weeks, and shows whether significant changes occurred in the values between the weeks. You can choose the percent change in values. The default value is 20, which indicates that if a change occurred of more than 20% in the value of any attribute between the weeks, then that field is highlighted.
Purpose	Compares key attributes for Top or Bottom N host servers between weeks.
Parameters	Data Center You can choose one data center from the list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. Uset Server
	 Host Server You can choose one or more host servers that belong to the selected clusters. Percent Change You can choose the Percent Change value, which indicates that if change occurred more than the selected percent value for any attribute between the weeks, then the field is highlighted. Top or Bottom You can choose whether you want to see the most used or least used host servers. Metric You can select any one metric to view the Top or Bottom host servers. The key metrics are CPU, Memory, Storage, Network, and Disk. Number of Resources You can choose the number of host servers to be displayed. Date Range You can choose a date range for which the trend is to be generated. You can either choose a historical period for which the trend is to be generated, for example: Last 30 Days Last 90 Days
Tables or views used	KVM_SERVER_DV KVM_SERVER_DATASTORE_DV KVM_SERVER_DISK_DV KVM_SERVER_NETWORK_DV
Output	The report displays a table, with comparisons between key metrics of Top or Bottom N host servers for the specified weeks. A percent change in values is highlighted if it is higher than the value specified. For configuration changes such as a change in the number of virtual machines, any change in value is highlighted.

Table 53. VMware VI Top or Bottom Host Servers Weekly Comparison report

Table 54. VMware VI Top or Bottom Workload Consumers Host Servers report

Report Structure	Description
Name	VMware VI: Top or Bottom Workload Consumers - Host Servers
Description	This report shows the host servers that use the most, or least, resources in the environment by processor, memory, and storage usage.
Purpose	Show the host servers that use the most, or least, resources (processor, memory, storage).

Report Structure	Description
Parameters	Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center
	Top or Bottom You can choose whether you want to see most used or least used resources.
	You can choose whether the charts display real values (processor capacity in GHz, memory in MB, or data store space in GB) or percentages. Number of Resources
	You can choose the number of host servers to be displayed, for example the top 10 or bottom 5.
	 You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days
	Shifts If the shifts are enabled, the hourly table displays the shift period as 1 or 2, depending on the peak and off-peak hours that are configured in the data warehouse. The daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
	Forecast Range This period is the number of days in the future that the forecast will be calculated for. The default is 7 days.
	Use Threshold for Ranking You can choose whether the Threshold must be used or not. Thresholds for Rank (%)
	You can choose the upper and lower thresholds that are used to filter the percentage values. For example, you can see the top 10 host servers that use more than 80% CPU or the bottom 5 host servers that use less than 10% memory.
Tables or views used	KVM_SERVER_DV
	KVM_SERVER_NETWORK_DV
Output	The report shows bar charts that display the host servers that use the most, or least, resources, and also the amount that they use. One bar chart is shown for each of CPU, memory, and data store space usage.
	If you specify that the charts display real values, you can see used and free processor capacity in GHz, memory in MB, data store space in GB, and network I/O in mbps. If you specify that the charts display percentages, only the used and free processor capacity, memory, and data store space are shown.
	You can click a host server and see the top or bottom virtual machines on the host. Alternatively you can also click to see the historical trend of that attribute of the selected host server over time.

Table 54. VMware VI Top or Bottom Workload Consumers Host Servers report (continued)

Table 55. VMware VI Top N VMs by Host Server report

Report Structure	Description
Name	Top N Virtual Machines by Host Server
Description	This report shows the top N virtual machines on one or more selected host servers by processor or memory usage.
Purpose	Show the virtual machines that use the most processor capacity or memory.
Parameters Tables or views	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center. Host Server You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days Number of Resources You can choose the number of virtual machines to be displayed, for example the top 10. Threshold You can choose the upper threshold that is used to filter the percentage values. For example, you can see the top 10 virtual machines that use more than 80% processor capacity. Display You can choose whether the charts display real values (processor capacity in GHz) or percentages. Metric You can select one or all the key metrics. The key metrics are CPU, Memory, Storage, Network I/O, and Disk I/O. Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1. Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
used	KVMVMDSUTL_DV
	KVM_VM_MEMORY_DV
	KVM_VM_NETWORK_DV
Output	The report shows a bar chart that displays the virtual machines that use the most resources, and also the amount they use.
	By clicking a virtual machine, you can see a historical trend of the virtual machine over time.

Table 56. VMware VI Bottom N VMs by Host Server report

Report Structure	Description
Name	Bottom N Virtual Machines by Host

Report Structure	Description
Description	This report shows the bottom N virtual machines on one or more selected host servers by processor or memory usage.
Purpose	Show the virtual machines that use the least processor capacity or memory.
Parameters	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers from the selected cluster. Date Range You can choose a data range for which the trend is to be generated. The following are You can choose a data range for which the trend is to be generated. The following are You can choose a data range for which the trend is to be generated.
	 You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days
	Number of Resources You can choose the number of virtual machines to be displayed, for example the bottom 5.
	You can choose the lower thresholds that are used to filter the percentage values. For example, you can see the bottom 5 virtual machines that use less than 20% processor capacity.
	Display You can choose whether the charts display real values (processor in GHz) or percentages.
	Metric You can select any one or all the key metrics. The key metrics are CPU, Memory, Storage, Network I/O, and Disk I/O.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
Tables or views	KVM_VIRTUAL_MACHINES_DV
used	KVMVMDSUTL_DV
	KVM_VM_MEMORY_DV
	KVM_VM_NETWORK_DV
Output	The report shows a bar chart that displays the virtual machines that use the least resources, and also the amount that they use.
	By clicking a virtual machine, you can see a historical trend of the virtual machine over time.

Table 56. VMware VI Bottom N VMs by Host Server report (continued)

Table 57. VMware VI Top N VMs by Resource Pool report

Report Structure	Description
Name	Top N Virtual Machines by Resource Pool
Description	This report shows the top N virtual machines by processor or memory usage on one or more resource pools on a cluster.

Report Structure	Description
Purpose	Show the virtual machines that use the most processor capacity or memory.
Parameters	Cluster You can choose one cluster or All from a list of clusters that belong to the selected data center.
	Resource Pool You can choose one or more resource pools from a list of resource pools that belong to the selected cluster.
	 You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days
	Number of Resources You can choose the number of virtual machines to be displayed, for example the top 10.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
	Use Threshold You can choose whether the Threshold must be used or not. Threshold You can choose the upper threshold that is used to filter the percentage values. For example, you can see the top 10 virtual machines that use more than 80% processor capacity.
Tables or views used	KVMCLTRRPS_DV
	KVM_VIKTUAL_MACHINES_DV KVM_VM_MEMORY_DV
Output	The report shows a bar chart that displays the virtual machines that use the most resources, and also the amount they use.
	By clicking a virtual machine, you can see a historical trend of the virtual machine over time.

Table 57. VMware VI Top N VMs by Resource Pool report (continued)

Table 58. VMware VI Bottom N VMs by Resource Pool report

Report Structure	Description
Name	VMware VI Bottom N VMs by Resource Pool
Description	This report shows the bottom N virtual machines by processor capacity or memory usage on one or more resource pools on a cluster.
Purpose	Show the virtual machines that use the least processor capacity or memory.

Table 58. VMwar	e VI Bottom	N VMs by Resc	ource Pool report	(continued)
-----------------	-------------	---------------	-------------------	-------------

Report Structure	Description		
Parameters	Cluster You can choose one cluster or All from a list of clusters that belong to the selected data center.		
	Resource Pool You can choose one or more resource pools from a list of resource pools that belong to the selected cluster.		
	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days Number of Resources You can choose the number of virtual machines to be displayed, for example the bottom 5 		
	 Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1. 		
	 Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. Use Threshold You can choose whether the Threshold must be used or not. Threshold You can choose the lower thresholds that are used to filter the percentage values. For example, you can see the bottom 5 virtual machines that use less than 20% processor 		
Tables or views	capacity. KVMCLTRRPS DV		
used	KVM_VIRTUAL_MACHINES_DV		
	KVM_VM_MEMORY_DV		
Output	The report shows a bar chart that displays the virtual machines that use the least resources, and also the amount that they use.		
	By clicking a virtual machine, you can see a historical trend of the virtual machine over time.		

Use case 2:

What clusters or host servers in the environment are balanced or unbalanced compared with the rest of the clusters or host servers in the environment?

Table 59. VMware VI Balanced and Unbalanced Clusters report

Report Structure	Description
Name	VMware VI Balanced and Unbalanced Clusters
Description	This report shows the balanced and unbalanced clusters in the environment.
Purpose	Identify the balanced and unbalanced clusters in the environment.

Report Structure	Description
Parameters	Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	 Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. Date Range You can choose a date range for which the report is to be generated.
Tables or views used	KVM_CLUSTERS_DV KVM_SERVER_DV KVM_SERVER_DISK_DV
	KVM_SERVER_NETWORK_DV
Output	The report shows three bar charts that display usage of processor capacity, memory, data store, Network I/O, and Disk I/O for the specified clusters in the specified time period. The clusters that are close to the average and are far from the average (either below or above the average) are shown.
	The following three reference lines show how the cluster is balanced:
	• The first reference line is the mean, which is the average of all the clusters in the environment.
	• The second reference line is the statistical maximum, which is determined by the following expression: 75th percentile value + 1.5 * (75th percentile value - 25th percentile value). The statistical maximum uses percentiles to determine values and does not appear in the chart if the values are off the axis.
	• The third reference line is the statistical minimum, which is determined by the following expression: 25th percentile value - 1.5 * (75th percentile value - 25th percentile value). The statistical minimum uses percentiles to determine values and does not appear in the chart if the values are off the axis

Table 59. VMware VI Balanced and Unbalanced Clusters report (continued)

Report Structure	Description
Name	VMware VI: Balanced and Unbalanced Host Servers
Description	This report shows the balanced and unbalanced host servers in the environment. The report can be used to view the trend charts with historical and forecasted values of the key resources for respective host server.
Purpose	Identify the balanced and unbalanced host servers in the environment.

Table 60. VMware VI Balanced and Unbalanced Host Servers report (continued)

Report Structure	Description
Parameters	Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected data center. Host Server You can choose one or more host servers from the selected cluster.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	 Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. Metric You can select any one or all the key metrics to view the balanced and unbalanced host servers. The key metrics are CPU, Memory, Storage, Network, and Disk. Date Range You can choose a date range for which the report is to be generated.
Tables or views used	KVM_SERVER_DV
	KVM_SERVER_DISK_DV
	KVM_SERVER_NETWORK_DV
Output	The report shows three bar charts that display usage of processor capacity, memory, data store, link, Network I/O, and Disk I/O for the specified host servers in the specified time period. The host servers that are close to the average and are far from the average (either below or above the average) are shown.
	The following three reference lines show how balanced the cluster is:
	• The first reference line is the mean, which is the average of all the clusters in the environment.
	• The second reference line is the statistical maximum, which is determined by the following expression: 75th percentile value + 1.5 * (75th percentile value - 25th percentile value). The statistical maximum uses percentiles to determine values and does not appear in the chart if the values are off the axis.
	• The third reference line is the statistical minimum, which is determined by the following expression: 25th percentile value - 1.5 * (75th percentile value - 25th percentile value). The statistical minimum uses percentiles to determine values and does not appear in the chart if the values are off the axis.

Use case 3:

Which are the top or bottom N virtual switches used by a list of host servers, which are based on network I/O?

Table 61. VMware VI Top or Bottom Virtual Switches by Network I/O

Report Structure	Description
Name	VMware VI Top or Bottom Virtual Switches by Network I/O
Description	This report shows the top or bottom N virtual switches by network I/O. The report can be used to drill down to see the top or bottom physical network interface cards that are connected to a virtual switch.

Report Structure	Description		
Purpose	Show the top or bottom N virtual switches by network I/O.		
Parameters	 Data Center You can choose one data center from a list of data centers. Cluster You can choose one cluster from a list of clusters that belong to the selected data center. 		
	 Host Servers You can choose one or more host servers from a list of host servers that belong to the selected Cluster and data center. Date Range You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days 		
	 Last 90 Days Top or Bottom You can choose whether you want to see the most used or least used virtual switches. Display You can choose whether the charts display real values (network I/O in Mbps) or percentages. Number of Resources You can choose the number of switches to be displayed, for example the top 10. 		
	 Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1. Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, then the default value is -1. 		
Tables or views used	KVM_SERVER_DV KVM_SERVER_NETWORK_DV		
Output	KVM_SERVER_VIRTUAL SWITCHES_DV The report shows a bar chart that displays the switches that use the most, or least, network I/O, and also the amount they use. The results are also shown in table form.		

Table 61. VMware VI Top or Bottom Virtual Switches by Network I/O (continued)

Use case 4:

Which are the top or bottom N physical network interface cards that are used by a list of host servers, which are based on network I/O?

Report Structure	Description
Name	VMware VI Top or Bottom Physical NICs
Description	This report shows the top or bottom physical network interface cards in one or more host servers. The output of this report provides a view into the network I/O. The report can be used to drill down to see the network usage by virtual machines that are hosted on a server.
Purpose	Show the top or bottom N physical network interface cards.

Table 62. VMware VI Top or Bottom Physical NICs

Table 62. VMware VI Top or Bottom Physical NICs (continued)

Description
Data Center
You can choose one data center from a list of data centers.
Cluster You can choose one cluster from a list of clusters that belong to the selected data center.
Host Servers
You can choose one or more host servers from a list of host servers that belong to the selected Cluster and data center.
Virtual Switches You can choose one or more virtual switches from a list of virtual switches that belong to the selected data center, cluster and host server.
 You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days
Top or Bottom You can choose whether you want to see the most used or least used network interface cards.
Display You can choose whether the charts display real values (network I/O in Mbps) or percentages. Number of Resources
You can choose the number of physical network interface cards to be displayed, for example the top 10.
Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1. Use Threshold
You can choose whether the Threshold must be used or not.
Threshold You can choose the upper or lower thresholds that are used to filter the percentage values. For example, you can see either top N physical NICs having Network usage more than 100% or bottom N physical NICs having Network usage less than 100%.
KVM_SERVER_DV
KVM_SERVER_NETWORK_DV
The report shows a bar chart that displays the switches that use the most, or least, network I/O, and also the amount they use. The results are also shown in table form.

Use case 5:

Which are the top or bottom N physical network interface cards that are used by a list of virtual machines, which are based on network I/O?

Report Structure	Description
Name	VMware VI Network Usage by VMs
Description	This report shows the top or bottom N physical network interface cards for selected virtual machines. The output of this report provides a view into the network I/O.
Purpose	Shows that virtual machines have physical network interface cards that use the most, or least, network I/O.
Parameters	Data Center You can choose one data center from a list of data centers.
	Cluster You can choose one cluster from a list of clusters that belong to the selected data center.
	Host Servers You can choose one host server from a list of host servers that belong to the selected cluster and data center.
	Physical NICs You can choose one or more physical NICs from a list of physical NICs that belong to the selected data center, cluster, and host server.
	 Date Range You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: Last 7 Days Last 30 Days Last 90 Days
	You can choose whether you want to see the most used or least used virtual machines. Number of Resources You can choose the number of physical network interface cards to be displayed, for example the top 10.
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
Tables or views used	KVM_VIRTUAL_MACHINES_DV
	KVM_VM_NETWORK_DV
Output	The report shows a bar chart that displays the virtual machines corresponding to the physical network interface cards that use the most, or least, network I/O, and also the amount they use. The results are also shown in table form.

Table 63. VMware VI Network Usage by VMs

Use case 6:

What is the memory allocation on all the host servers? On which servers does the allocation exceed the physical memory present?

Table 64. VMware VI Host Server Memory Allocation report

Report Structure	Description
Name	VMware VI Host Server Memory Allocation

Report Structure	re Description This report shows the average memory that is allocated to virtual machines on the host servers in comparison with the actual physical memory on the host.				
Description					
Purpose	Show the memory allocation on all the host servers and highlight the servers where the allocation exceeds the physical memory present.				
Parameters	Date Range You can choose a date range for which the trend is to be generated. The following are the options for Date Range selection: • Last 7 Days • Last 30 Days • Last 90 Days • Last 90 Days Data Center You can choose one data center from a list of data centers. Cluster You can choose one or more clusters from a list of clusters that belong to the selected Data Center. Host Server You can choose one or more host servers that belong to the selected clusters.				
Tables or views used	KVM_VIRTUAL_MACHINES_DV KVM_SERVER_MEMORY_DV				
Output	The report displays a bar chart, showing the total physical memory on a host server that is compared with the total memory allocated to virtual machines on that server. The results are also shown in table form. Total memory that is allocated to virtual machines on a host must not exceed the total physical memory of the host server. In cases where the allocated memory might exceed the total physical memory on the host, the servers are highlighted. By clicking a particular host server, you can see the memory allocation trend over time.				

Table 64. VMware VI Host Server Memory Allocation report (continued)

Using the Cognos data model

The historical data collected by the VMware VI agent and the NetApp Storage agent can be used in Tivoli Common Reporting for building ad hoc reports and queries. A Cognos data model that is ready to use is provided to enable this function.

The VMware VI agent data model is part of the IBM Tivoli Monitoring for Virtual Environments Reports V7.2 package. The data model is located under IBM Tivoli Monitoring for Virtual Environments Reports in the **Data Navigation** tree of the Query and Report Studio. The data model is a layer built on top of the Tivoli Data Warehouse to make the data more usable. The data model contains predefined relationships so that when you drag elements across different tables and views, they are joined so you do not have to manually write SQL code.

After importing the package in Tivoli Common Reporting, click **Launch** > **Query Studio** from the **TCR** menu and select **IBM Tivoli Monitoring for Virtual Environments Reports**. The Query Studio is a web-based ad hoc query tool that you can use to build quick tables and charts by dragging.

The structure of the VMware VI and NetApp Storage agent data models is shown in Figure 70 on page 231:

🖶 IBM Tivoli Monitoring for Virtual Environments Reports	
🖻 🏭 ITM for Virtual Environments (Query)	
🛨 🏭 TRAM Shared Dimensions (Query)	
😑 🟭 VMware VI Agent (Query)	
😠 🤠 🏭 ITM for Virtual Environments Shared Dimensions (Query)	
🕀 🧰 Detailed	
🕀 🧰 Hourly	
🕀 🧰 Daily	
🕀 🧰 Weekly	
🗉 🧰 Monthly	
🛱 🖶 NetApp Storage Agent (Query)	
🗄 🏭 ITM for NetApp Storage Shared Dimensions (Query)	
🗉 🧰 Hourly	
🕀 🧰 Daily	
🖅 🔀 Configuration (Query)	

Figure 70. Data navigation tree

Cognos data models are virtual star schema models separated into facts and dimensions. Facts are measurable quantities that can be aggregated, such as CPU utilization and number of processors. Dimensions are the main identifiers by which facts can be grouped, aggregated, and organized. For example, time and server are dimensions by which the fact CPU utilization can be grouped.

The data model contains two types of dimensions:

- **TRAM Shared Dimensions (TRAM** stands for Tivoli Reporting and Analytics Model), which are shared across Tivoli by products such as **Time**
- ITM for Virtual Environments Shared Dimensions, which are dimensions shared across the IBM Tivoli Monitoring for Virtual Environments agents such as Datacenter, Clusters, Environments, VMs, Data Stores, and so on

The facts in the data model are organized under folders by their summarization type, such as **Daily** and **Hourly**.

When you expand **Daily** and **Hourly**, you can see the attribute groups (see Figure 71 on page 232).



Figure 71. Daily folder

Each attribute group corresponds to a table or view in the data warehouse. Each attribute group contains a group of facts or measures, such as **MIN_CPU_Utilization** and **AVG_CPU_Utilization**, and some identifiers, such as **CPU_Number**, and shift and vacation periods.

The data model has relationships defined between the different tables and with the NetApp Storage agent data model, so that you can correlate VMware data stores to NetApp volumes.

By using the various controls in Query Studio, you can build a report in minutes.

The following procedure shows an example for building a report:

 Drag one of the IBM Tivoli Monitoring for Virtual Environments Shared Dimensions, such as the Datacenter/Clusters/Environments/VMs.VM_Name dimension. When looking at data at the cluster and server level, drag the dimension from the Datacenter/Clusters/Environments dimension. When looking at the virtual machine level, use the Datacenter/Clusters/Environments/VMs dimension. Followed by the VM_Name dimension, drag the Datacenter/Clusters/Environments/ VMs.Server_Hostname dimension. An example is shown in Figure 72 on page 233.

VM_Name	Server_Hostname
3vm	itm64vm5.tivlab.raleigh.ibm.com
IBM-EC2960CA51B	itm64vm14.tivlab.raleigh.ibm.com
ITMESX1	itm64vm16.tivlab.raleigh.ibm.com
ITMESX1	itm64vm16.tivlab.raleigh.ibm.com
ITMWINVMЗ	itm64vm4.tivlab.raleigh.ibm.com
ITMWinVM	itm64vm4.tivlab.raleigh.ibm.com
ITMWinVM2	itm64vm4.tivlab.raleigh.ibm.com
VMware Studio	itm64vm11.tivlab.raleigh.ibm.com
demoVM	itm64vm1.tivlab.raleigh.ibm.com
esx01t	itm64vm4.tivlab.raleigh.ibm.com
esx02t	itm64vm4.tivlab.raleigh.ibm.com
fac2w2k3b	itm64vm16.tivlab.raleigh.ibm.com
fac2w2k3b	itm64vm16.tivlab.raleigh.ibm.com
fred	itm64vm2.tivlab.raleigh.ibm.com
itm64rh1	itm64vm4.tivlab.raleigh.ibm.com
itm64rh3	itm64vm2.tivlab.raleigh.ibm.com
itm64rh4	itm64vm2.tivlab.raleigh.ibm.com
itm64suse1	itm64vm2.tivlab.raleigh.ibm.com

Figure 72. Sample dimensions

2. Filter the data by server in order to narrow down your results. Click the **Server_Hostname** column and click the **Filter** icon in the toolbar at the top of the screen. This action displays a list of host servers. Select the server you want to view and click **OK**. Alternatively, you can search for servers by clicking **Search for values**. Select the **Prompt every time the report runs** check box to make this a report parameter. This specification reduces the number of entries automatically included on the report.

These actions are illustrated in Figure 73 on page 234.

Font Size Aar B I	
Filter (Pick values from a list)	
Reduce the amount of data in the report. With the Prompt of immediate the report runs.	option selected, the filter can be changed each Search for values Type in value
Filter on: Server_Hostname	Prompt every time the report runs
Condition: Show only the following	
itm64vm1.tivlab.raleigh.ibm.com itm64vm10.tivlab.raleigh.ibm.com itm64vm11.tivlab.raleigh.ibm.com	
itm64vm13.tivlab.raleigh.ibm.com	
itm64vm14.tivlab.raleigh.ibm.com itm64vm15.tivlab.raleigh.ibm.com itm64vm16.tivlab.raleigh.ibm.com	
🔲 itm64vm2.tivlab.raleigh.ibm.com 🔽	
Select all Deselect all	

Apply the filter to individual values in the data source

Figure 73. Server filter screen

- 3. Drag any metric from a related attribute group, for example, if you want to look at server-related metrics, drag any metric in the attribute groups with names that start with Server. For virtual machine-related metrics, drag data from the attribute groups with names that start with VM. For this example, drag the Virtual Machines Daily.AVG_Used_CPU_MHz metric.
- 4. Filter the metric to eliminate values of -1, which indicate that no data was collected. Click on the AVG_Used_CPU_MHz column, click the Filter icon and enter the values as shown in Figure 74 on page 235.

ilter on: VG_CPU	_Utilization	Prompt every time the report runs
Condition		Apply the filter to:
Show or	nly the following	Ovalues in the report
		Individual values in the data source
rom: 💿	0	
0) Lowest value	
īo: O		
۲) Highest value	

Figure 74. Missing values filter screen

5. Drag a time element from the **TRAM Shared Dimensions.Time** Dimension. For daily data you can use **Date**, and for hourly you can use **Standard Timestamp**. You can use any of the other time metrics to organize the data. The data now appear as in Figure 75 on page 236.

VM_Name	Server_Hostname	Date	AVG_Used_CPU_MHz
vsvdash1	itm64vm13.tivlab.raleigh.ibm.com	Feb 14, 2011	182.73
vsvdash2	itm64vm13.tivlab.raleigh.ibm.com	Feb 14, 2011	124.45
vsvdash3	itm64vm13.tivlab.raleigh.ibm.com	Feb 14, 2011	122.98
vsvtaddm	itm64vm13.tivlab.raleigh.ibm.com	Feb 14, 2011	123.69
win2008vm1	itm64vm13.tivlab.raleigh.ibm.com	Feb 14, 2011	2,654.675
vsvdash1	itm64vm13.tivlab.raleigh.ibm.com	Feb 15, 2011	169.17
vsvdash2	itm64vm13.tivlab.raleigh.ibm.com	Feb 15, 2011	109.04333333
vsvdash3	itm64vm13.tivlab.raleigh.ibm.com	Feb 15, 2011	111.82
vsvtaddm	itm64vm13.tivlab.raleigh.ibm.com	Feb 15, 2011	123.195
win2008vm1	itm64vm13.tivlab.raleigh.ibm.com	Feb 15, 2011	2,653.92666667
vsvdash1	itm64vm13.tivlab.raleigh.ibm.com	Feb 16, 2011	166.52666667
vsvdash2	itm64vm13.tivlab.raleigh.ibm.com	Feb 16, 2011	155.74333333
vsvdash3	itm64vm13.tivlab.raleigh.ibm.com	Feb 16, 2011	226.14333333
vsvtaddm	itm64vm13.tivlab.raleigh.ibm.com	Feb 16, 2011	122.2
win2008vm1	itm64vm13.tivlab.raleigh.ibm.com	Feb 16, 2011	2,655
vsvdash1	itm64vm13.tivlab.raleigh.ibm.com	Feb 17, 2011	163.02
vsvdash2	itm64vm13.tivlab.raleigh.ibm.com	Feb 17, 2011	117.35
21.20%	a 64 46 0 1 1 1 1 1 1	- 1 - 1 - 0 - 0 - 1	100.105

Figure 75. Sample report

6. Create a cross tab by pivoting on the date. Click the **Date** column and select **Pivot** (**Create a Cross tab**) from the menu. This action creates a cross tab as shown in Figure 76 on page 237, with dates along the x-axis and the virtual machine names along the y-axis.

8 C

<u>Title</u>

Summary		641.705	669.87642857	703.90285714	635.08	633.07	630.373	666.49428571	671.54
	win2008vm1	2,654.675	2,653.92666667	2,655	2,653.35	2,654.66	2,654.465	2,655.15666667	2,655.0
win2008vm1	itm64vm13.tivlab.raleigh.ibm.com	2,654.675	2,653.92666667	2,655	2,653.35	2,654.66	2,654.465	2,655.15666667	2,655.0
	vsvtaddm	123.69	123.195	122.2	121.255	121.965	125.1475	123.55	
vsvtaddm	itm64vm13.tivlab.raleigh.ibm.com	123.69	123.195	122.2	121.255	121.965	125.1475	123.55	
	vsvdash3	122.98	111.82	226.14333333	120.425	112.77	107.215	110.56666667	
vsvdash3	itm64vm13.tivlab.raleigh.ibm.com	122.98	111.82	226.14333333	120.425	112.77	107.215	110.56666667	
	vsvdash2	124.45	109.04333333	155.74333333	117.35	114.705	110.3375	109.87666667	
vsvdash2	itm64vm13.tivlab.raleigh.ibm.com	124.45	109.04333333	155.74333333	117.35	114.705	110.3375	109.87666667	
	vsvdash1	182.73	169.17	166.52666667	163.02	161.25	154.7	152.34	
vsvdash1	itm64vm13.tivlab.raleigh.ibm.com	182.73	169.17	166.52666667	163.02	161.25	154.7	152.34	
A¥G_Used_	_CPU_MHz	Feb 14, 2011	Feb 15, 2011	Feb 16, 2011	Feb 17, 2011	Feb 18, 2011	Feb 19, 2011	Feb 20, 2011	Feb 21

Figure 76. Cross tab

7. Create a chart. Click the **chart** icon and select **Area** and the **Stacked** option from the Chart wizard, as in Figure 77. Click **OK**.

Server Hostname: itm64vm13.	tivlab.raleigh.ibm.com	<u>Title</u>
hart		
pecify how the data is to be g	raphed. Choose None to r	emove the chart.
Chart type:		Show the values on the chart
O None		Show the following in the report:
💿 Area 🛛 💌		Ohart and table
o M Standard	Stacked	Chart only
100 Percent Stacked	Standard with 3-D Axis	

Figure 77. Chart wizard

This selection produces a stacked area chart for all the virtual machines on the filtered host server. As shown in Figure 78, five virtual machines exist on this host and their average processor capacity used in MHz is stacked and plotted against time.



Figure 78. Stacked chart

- 8. Save the report. Click the Save icon on the toolbar at the top of the screen. Enter a name for this report, such as "VM CPU Stacked Area Chart" and click OK. When you go back by clicking the Back button at the top of the screen, you see that the report has been saved under the main ITM for Virtual Environments Reports v6.2.3 folder.
- 9. You can further enhance this report. In order to filter by date and time, use the Report Studio. In the TCR navigation, click **More**, next to the report actions, as in Figure 79.

Public Folders > IBM Tivoli Monitoring for Virtual Environments Reports	🔲 🖩 🔛 📽 💖 🖬 😽 🗈 🗙 🗔 🛱		
	Entries: 1	- 6	
□ Name ≑	Modified	Actions	
🗖 🗎 Accounting	January 18, 2013 1:21:20 PM	More	
🗖 🗎 Performance Trends and Resource Forecasts	January 18, 2013 1:21:36 PM	More	
🗖 🗎 Prerequisites Checking	January 18, 2013 1:21:42 PM	More	
🔲 🛅 What-If Analysis for Capacity Estimation	January 18, 2013 1:22:08 PM	More	
🔲 🛅 Workload Right-Sizing and Balancing	January 18, 2013 1:22:27 PM	More	
🔲 💽 VM CPU Stacked Area Chart	January 18, 2013 4:44:13 PM	🔲 🕨 🔪 🔛 🏢 More	

Figure 79. TCR navigation

- 10. From the next screen, select **Open with Report Studio**. The Report Studio is a more advanced report editing tool where you can manipulate various parts of the reports. When the Report Studio opens, click the title in the report header to edit the report.
- 11. Add the time range filtering. Hover over **Page Explorer** in the center of the page to browse to the prompt page, as in Figure 80. Parameters that filter the report are called **Prompts** in Cognos. The prompt page is visible every time you run the report.

	Proce Fundamentary	
Page Explorer	Page Explorer > Page I Page I Prompt Page I Prompt Page I Prompt Classes Classes	VM CPU Stacked Area Chart

Figure 80. Prompt page

- 12. When the prompt page is open, notice that the **Server_Hostname** prompt already exists because you filtered the data by the server host name.
- **13**. Click the **Toolbox** tab in the navigation. Drag **TCR Date Range prompt** from the toolbox onto the prompt page, as in Figure 81.

Insertable Objects			Drop items here to add them to the page header
TCR Collapse Table Below Item	A Page E	Server_Hostname Provide a value:	
Re TCR Start Date Display Value	plore	Server_Hostname?	
[法] TCR End Date Display Value	pt: drag and drop control	into panel Filter	
ab Text Item			
E Block			
III Table	***		
[ab] Field Set			
Calculated Member		Date Range	
Calculated Measure			
Totersection (Tunle)			
Query Calculation			
Cayout Calculation			
😰 Image		HINT: in order to use "TCR Date	Range" you need to add following filter to your query(ies): #prompt('TCRDateFilter', 'token', '1=1', '[YOUR-COLUMN-NAME]')#
Crosstab Space	-1		
		L	

Figure 81. TCR date range prompt

A hint appears. Copy the text #prompt('TCRDateFilter', 'token', '1=1', '[YOUR-COLUMN-NAME]')#.
14. Browse to the query to add this filter. Hover over the center of the page and select Query under Query Explorer, as in Figure 82 on page 240.



Figure 82. Query Explorer

You see next the query that was created when you dragged data from the data model into the report. As you can see on the right of the screen, a filter on the **Server_Hostname** already exists.

15. Filter by time. Browse to the Toolbox tab and drag Filter onto the Detail Filters section. This action opens an Expression editor. In this box, paste the value you copied from the prompt page, that is, #prompt('TCRDateFilter','token','1=1', '[YOUR-COLUMN-NAME]')#, and replace YOUR-COLUMN-NAME with [TRAM Shared Dimensions (Query)].[Time].[Standard Timestamp] dragged from the Time dimension under Available Components, as shown in Figure 83. Click OK to save.



Figure 83. Expression editor
16. Save the report and run it by clicking **Play** on the toolbar at the top of the page, as shown in Figure 84.



Figure 84. Play button

17. When you run the report, you see two prompts, one for host server name and the other for time range. Select a host, select **Last 7 Days** from the date filter, as in Figure 85, and click **Finish**.

st 7 days	•
Feb 23, 2011	
12 : 00 AM	
Feb 23, 2011	-
11 · 50 PM	
	Feb 23, 2011 12 : 00 AM Feb 23, 2011 11 : 59 PM

Figure 85. Report values

The report runs and displays values only for the last seven days and the server selected, as in Figure 86 on page 242.

VM CPU Stacked Area Chart

-1

-1

-1

-1

-1

-1

Server_Hostname: itm64vm15.tivlab.raleigh.ibm.com



Figure 86. Report results

After you have the desired output, you can save the report and use it like any other report. You can run the report later and view it in different formats such as PDF and Excel, schedule it, email it, and so on.

Chapter 10. Linux Kernel-based Virtual Machines agent: Using Performance and Capacity Management Reports

You can find complete information about prerequisites, importing reports, and running reports by using the agent-specific information with the Tivoli Common Reporting information in the IBM Tivoli Monitoring Administrator's Guide.

Cognos data model and reports to be used in Tivoli Common Reporting were introduced with IBM Tivoli Monitoring V6.2.2 Fix Pack 2.

The reports in this package are historical reports, that use summarized data that is collected in Tivoli Data Warehouse V6.2.2. These reports are built to run only with the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines.

The DB2, Oracle, and SQL Server databases are supported for running all reports.

The Cognos reports can be administered, run, and edited by using Tivoli Common Reporting V2.1 or V3.1 or later software that is included with IBM Tivoli Monitoring V6.2.2 Fix Pack 2 or later. For more information about Tivoli Common Reporting, see the Tivoli Common Reporting Community (https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/communityview?communityUuid=9caf63c9-15a1-4a03-96b3-8fc700f3a364).

This version of Tivoli Common Reporting includes Cognos Business Intelligence and Reporting V8.4 or later.

More information about Tivoli Common Reporting

You can find information about Tivoli Common Reporting at the Tivoli Common Reporting documentation Information Center and the Tivoli Common Reporting website.

For more information about the Tivoli Common Reporting V3.1 or later requirements, see the JAZZ for SM Information Center.

For complete documentation for the Tivoli Common Reporting tool, see the Tivoli Common Reporting documentation Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ic-home.html).

The Tivoli Common Reporting website contains information and how-to videos about subjects such as how to create IBM Tivoli Monitoring reports by dragging, import Tivoli Common Reporting and Cognos reports, and set up Cognos and Tivoli Common Reporting data connections. You can find a report catalog and information about reporting across Tivoli products at the Tivoli Common Reporting Community (https://www.ibm.com/developerworks/mydeveloperworks/groups/service/html/ communityView?communityUuid=9caf63c9-15a1-4a03-96b3-8fc700f3a364). Training videos can be downloaded at https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/ wiki/IBM%20SmartCloud%20Monitoring/page/Training%20on%20ITMfVE%20vNext%20Linux%20KVM %20reports to help you learn, install, and troubleshoot Cognos reports for the IBM Tivoli Monitoring for Virtual Environments Agent for Linux Kernel-based Virtual Machines and the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI.

Linux KVM: Performance and Capacity Management Reports descriptions

The Linux Kernel-based Virtual Machines agent provide five categories of Performance and Capacity Management Reports: Prerequisites Checking, Accounting, Performance trends and resource forecasts, What if analysis for workload placement, and Workload right-sizing and balancing.

The following Performance and Capacity Management reports are available:

- Prerequisites Checking
 - Linux KVM Report Prerequisite Scanner
- Accounting
 - Linux KVM Number of Hosts and VMs monitored
- Performance trends and resource forecasts
 - Linux KVM Physical CPU Utilization by VM
 - Linux KVM Storage Pool Capacity and Performance Trend
- What if analysis for workload placement
 - Linux KVM Number of VMs for Host
 - Linux KVM Resources Needed for additional VMs on a Host
- Workload right-sizing and balancing
 - Linux KVM Top or Bottom VMs by Physical CPU Utilization

Attribute groups

The Performance and Capacity Management reports use the following attribute groups:

- KV1_HOSTS
- KV1_HOSTS_HV
- KV1_HOSTS_DV
- KV1_HOSTS_WV
- KV1_HOSTS_MV
- KV1_VIRTUAL_MACHINES_HV
- KV1_VIRTUAL_MACHINES_DV
- KV1_VIRTUAL_MACHINES_WV
- KV1_VIRTUAL_MACHINES_MV
- KV1_STORAGE_POOLS_HV
- KV1_STORAGE_POOLS_DV
- KV1_STORAGE_POOLS_WV
- KV1_STORAGE_POOLS_MV
- KV1_HOST_CPU_HV
- KV1_HOST_CPU_DV

Prerequisites Checking report

You can use prerequisites checking to create a prerequisite scanner report for the Linux KVM agent by using predefined reports.

The following report is available for prerequisites checking:

• Linux KVM Report Prerequisite Scanner

Linux KVM Report Prerequisite Scanner

This report runs on DB2, Oracle, and MS SQL Server databases. The report shows you if all the prerequisite tables and views are present to successfully run Linux KVM reports, either all or on a per-report basis.

Report element	Details
Parameters	To run the prerequisite scanner, ensure that you defined and tested a database (DB2/MS SQL Server/Oracle) connection to the Tivoli Data Warehouse. Also ensure that you choose the appropriate connection to generate Prerequisite Scanner Report.
	Database Type DB2/Oracle/MS SQL Server
	Display Options Check all reports or a specific report by choosing from a category within the reports package.
Tables or views used	DB2:
	SYSCAT.VIEWS SYSCAT.TABLES
	Oracle:
	SYS.ALL VIEWS
	SYS.ALL_TABLES
	MS SQL Server:
	INFORMATION_SCHEMA.VIEWS INFORMATION_SCHEMA.TABLES
Output	A legend is displayed at the beginning of the report that shows the meaning of the symbols displayed under the Status column. A red x and a yellow exclamation point (!) indicate error conditions. When an error is indicated, a corrective action is suggested that includes links to the appropriate documentation. The table contains two columns:
	 Missing Tables/Views from IBM Tivoli Monitoring for Linux KVM agent
	In the first column, missing tables and views are listed in order, showing status with a red x or a yellow exclamation point (!). Available tables and views are shown with a green check mark.
	Missing Table/Views for Shared Dimensions
	In the second column, the IBM_TRAM, TIME_DIMENSION, WEEKDAY_LOOKUP, MONTH_LOOKUP, and ComputerSystem under the IBM_TRAM schema are checked for availability.
	Even if all the tables and views are available in the warehouse, a report might still not run due to inadequate time stamps generated. In this case, you must run the appropriate database scripts to populate the TIME_DIMENSION table. When you run the prerequisite scanner to check a specific report, the IBM Tivoli Monitoring tables used for its implementation are checked for availability and a status is displayed for each of the tables used. Since the TIME_DIMENSION table is used by almost all predefined reports, you can check its availability by clicking a hyperlink provided in the report.

Table 65. Linux KVM Report Prerequisite Scanner

Accounting report

You can create an Accounting report for the Linux KVM agent using predefined reports. The report shows information about the numbers of hosts and VMs that are monitored.

The following report is available for Accounting:

· Linux KVM Number of Hosts and VMs monitored

Linux KVM Number of Hosts and VMs monitored

This report displays the number of Hosts, defined virtual machines, active virtual machines and processor cores monitored by the Linux KVM Agent for the KVM Hypervisor.

Report element	Details
Parameters	None
Tables or views used	KV1_HOSTS
Output	This report contains a table that shows the number of defined and active virtual machines and processor cores against each host in the environment. The report also shows the total number of hosts that are being monitored for the KVM hypervisor.

Table 66. Linux KVM Number of Hosts and VMs monitored

Performance trends and resource forecasts

You can create performance trends and resource forecasts reports using predefined reports. These reports shows information about VM processor utilization and storage pool capacity and performance.

The following reports are available for performance trends and resource forecasts:

- Linux KVM Physical CPU Utilization by VM
- Linux KVM Storage Pool Capacity and Performance Trend

Linux KVM Physical CPU Utilization by VM

This report shows the CPU usage of all VMs stacked up in a host over time.

Table 67. Linux KVM Physical CPU Utilization by VM

Report element	Details
Parameters	Date Range
	Report Period You can choose from a predefined date range such as Last Week, Current Month, Last 30 Days. Alternatively, you can enter a start and end date and time for the reporting period by choosing the Date Range (below) option.
	Start Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	End Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	Resource Selection
	Host Select the required host from the environment either by using a menu list or by applying search and filter.
	Display Options
	Summarization Type Choose the summarization type from the menu list. The options are Hourly, Daily (the default value), Weekly, and Monthly.
	Shifts If shifts are enabled, the hourly tables have a value for SHIFTPERIOD of 1 or 2, based on off-peak and peak hours that are configured in the data warehouse. The daily tables have values of 1 and 2. 1 corresponds to the off-peak hours and 2 corresponds to peak hours. A value of -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.
	Vacation days If the Vacation Period option is not enabled, the default value is -1. Otherwise, enter 0 for work days or 1 for vacation days.
Tables or views used	KV1_VIRTUAL_MACHINES_HV KV1_VIRTUAL_MACHINES_DV KV1_VIRTUAL_MACHINES_WV KV1_VIRTUAL_MACHINES_WV KV1_VIRTUAL_MACHINES_MV KV1_HOSTS_HV KV1_HOSTS_DV KV1_HOSTS_WV KV1_HOSTS_WV KV1_HOSTS_WV Tip: Although the report supports Weekly and Monthly Summarization types, if you do not plan to run the report for these summarizations, do not configure these summarizations for the

Table 67. Linux KVM Physical CPU Utilization by VM (continued)

Report element	Details
Output	This report contains an area chart that shows the total CPUs used by virtual machines over a selected time. The crosstab shows Average CPUs used by each VM (%) for the time stamps within the selected date range.

Linux KVM Storage Pool Capacity and Performance Trend

This report shows various performance and capacity trends for one or more selected storage pools.

Report element	Details
Parameters	Date Range
	Report Period You can choose from a predefined date range such as Last Week, Current Month, Last 30 Days. Alternatively, you can enter a start and end date and time for the reporting period by choosing the Date Range (below) option.
	Start Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	End Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	Resource
	Host Select the required host from the environment.
	Storage Pools Select either all or a subset of the storage pools connected to the selected host.
	Display Options
	Summarization Type Choose the summarization type from the menu list. The options are Hourly, Daily (the default value), Weekly, and Monthly.
	Shifts If shifts are enabled, the hourly tables have a value for SHIFTPERIOD of 1 or 2, based on off-peak and peak hours that are configured in the data warehouse. The daily tables have values of 1 and 2. 1 corresponds to the off-peak hours and 2 corresponds to peak hours. A value of -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.
	Vacation days If the Vacation Period option is not enabled, the default value is -1. Otherwise, enter 0 for work days or 1 for vacation days.

Report element	Details
Tables or views used	KV1_HOSTS _ DVKV1_STORAGE_POOLS_HVKV1_STORAGE_POOLS_DVKV1_STORAGE_POOLS_WVKV1_STORAGE_POOLS_MVTip: Although the report supports Weekly and MonthlySummarization types, if you do not plan to run thereport for these summarizations, do not configure thesesummarizations for the attribute groups on TivoliEnterprise Portal Server.
Output	This report displays an area chart to compare the storage pool used capacity against available capacity. It also displays a line chart for storage pool utilization over the selected time. A default threshold of 80% is applied to the storage pool utilization.

Table 68. Linux KVM Storage Pool Capacity and Performance Trend (continued)

What-if analysis for workload placement reports

You can use what-if analysis to create a workload placement report for the agent for Linux Kernel-based Virtual Machines by using predefined reports.

The following reports are available for what-if analysis for workload placement:

- Linux KVM Number of VMs for Host
- · Linux KVM Resources Needed for Additional VMS on a Host

Linux KVM Number of VMs for Host

Use this report to do what-if analysis to determine the number of additional virtual machines that can be placed on a host based on the average historical usage and other user inputs.

VM Profile is the resource used by all the deployed VMs on each host over time, average, peak or can be user specified. Available Capacity, before applying the buffer, is the amount of resource available for allocation on each host before applying the buffer value. The buffer is the amount of resource that the user does not want to allocate. Available Capacity (after applying Buffer) = Available Capacity (before applying Buffer) - Buffer and Number of VMs = Available Capacity(after applying Buffer) / VM Profile. The number of VMs that can be added to this host is the minimum of the values calculated for each resource. A value of 999,999,999 for the number of VMs indicates that the number of VMs that can be added to the host is unlimited. This value is displayed if the VM Profile is 0. The rows marked in red show the constrained resources.

Table 69. Linux KVM Number of VMs for Host

Report element	Details
Parameters	Date Range
	Report Period You can choose from a predefined date range such as Last Week, Current Month, Last 30 Days. Alternatively, you can enter a start and end date and time for the reporting period by choosing the Date Range (below) option.
	Start Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	End Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	Resource Selection
	Host Select the required host from the environment either by using a menu list or by applying search and filter.
	User Inputs for Analysis
	VM Profile Select the required profile to run the report against. The options are Average (default), Peak, and User-defined.
	CPU Buffer (GHz) The buffer indicates the resources that the user does not want to allocate.
	User-defined Factor for VM Profile – CPU (GHz) Enter user-defined values to be used alongside the user-defined profile.
	Include Peak Hours
	Start You can choose a start time from a menu list
	End You can choose an end time from a menu list.
Tables or views used	KV1_HOST_CPU_HV KV1_HOST_CPU_DV
Output	This report contains a table that shows the number of VMs that can be added to a host, based on the resource usage of the monitored VMs and the available resource capacity on the host after allowing for user-defined buffers. The table shows information related to CPU on the selected host and how this resource affects the total number of VMs that can be added. The average resources used is the historical average of all the deployed VMs on the host. The available resource capacity is the current deallocated resources. The number of VMs that can be deployed is the Available Resource Capacity/Average Resource Usage per VM. The row in the table is highlighted in green if additional VMs can be deployed on the host and red, if no more VMs can be deployed on the selected host.

Linux KVM Resources Needed for Additional VMS on a Host

This report provides an estimate of how many more resources (CPU) are needed to add additional VMs to the host.

Report element	Details
Parameters	Date Range
	Report Period You can choose from a predefined date range such as Last Week, Current Month, Last 30 Days. Alternatively, you can enter a start and end date and time for the reporting period by choosing the Date Range (below) option.
	Start Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	End Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	Resource Selection
	Host Select the required host from the environment either by using a menu list or by applying search and filter.
	User Inputs for Analysis
	VM Profile Select the required profile to run the report on. The options are Average (default), Peak, and User-defined.
	Number of VMs to add to the Host Enter the number of VMs you want to add to the Host.
	CPU Buffer (GHz) The buffer is to indicate the resources that the user does not want to allocate.
	User-defined Factor for VM Profile – CPU (GHz) Enter user-defined values to be used alongside the User-defined profile.
	Include Peak Hours
	Start You can choose a start time from a menu list End You can choose an end time from a menu list.
Tables or views used	KV1_HOST_CPU_HV KV1_HOST_CPU_DV

Table 70. Linux KVM Resources Needed for Additional VMS on a Host (continued)	
---	--

Report element	Details
Output	This report contains a table that shows the resources required to successfully add VMs to the selected host, based on the current resource usage of the monitored VMs. The table shows information related to CPU on the selected host and how much of these resources is required to add the required number of VMs. A value of 0 for a particular resource means that no additional capacity is required for this resource to accommodate the new VMs. The average resource usage per VM is the historical average of all the deployed VMs on the host. Resources required by additional VMs to be added to the host is the Average Resource Usage per VM * Number of VMs to be added. The available resource capacity is the current deallocated resources. Additional capacity required for new VMs is the Available Resource Capacity - Resources Required by Additional VMs. If additional resources are required, the row is highlighted in red.

Workload right-sizing and balancing report

You can use the predefined workload right-sizing and balancing report to determine the overall performance of the environment for the agent for Linux Kernel-based Virtual Machines.

The following report is available for workload right-sizing and balancing:

• Linux KVM Top or Bottom VMs by Physical CPU Utilization

Linux KVM Top or Bottom VMs by Physical CPU Utilization report

This report displays average CPU Utilization for all VMs deployed on the selected host during the report period, with bar charts that show the top and bottom n VMs based on CPU utilization.

Table 71. Linux KVM Top or Bottom VMs by Physical CPU Utilization report

Report element	Details
Parameters	Date Range
	Report Period You can choose from a predefined date range such as Last Week, Current Month, Last 30 Days. Alternatively, you can enter a start and end date and time for the reporting period by choosing the Date Range (below) option.
	Start Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	End Date You can choose a start date from a calendar and start time from the time widget. Both date and time must be selected.
	Resource Selection
	Host Select the required host from the environment either by using a menu list or by applying search and filter.
	Display Options
	Number of VMs (N) You can choose any integer to filter the number of top/bottom VMs visible in the bar charts.
	Top/Bottom You can choose to run the report to show the top N VMs or the bottom N VMs.
	Units You can choose to run the report to show the top/bottom charts with either real values or percentages.
	Shifts If shifts are enabled, the hourly tables have a value for SHIFTPERIOD of 1 or 2, based on off-peak and peak hours that are configured in the data warehouse. The daily tables have values of 1 and 2. 1 corresponds to the off-peak hours and 2 corresponds to peak hours. A value of -1 corresponds to the summarized value for that day. If shifts are not enabled, the default value is -1.
	Vacation days If the Vacation Period option is not enabled, the default value is -1. Otherwise, enter 0 for work days or 1 for vacation days.
Tables or views used	KV1_HOSTS_DV KV1_VIRTUAL_MACHINES_ DV
Output	This report shows a bar chart: the top n VMs based on average CPU Utilization (number of CPUs consumed or % CPU consumption) or the bottom n VMs. A table below these charts displays CPU attributes for all VMs deployed on the selected host during the report period.

Chapter 11. Capacity Planner for VMware

Capacity Planner works on the configuration and historical usage data collected by IBM Tivoli Monitoring to provide recommendations for optimizing the virtual environment.

Planning Center for VMware wizard

The Planning Center for VMware provides a simple five-step wizard to generate an optimized plan for the Virtual Environment.

Important: Capacity Planner for VMware does not support multiple concurrent users.

Planning Center for VMware ×	₽ +
Planning Center for VMware	
You have acquired the Planning Center for VMware lock at 25 Jul 2013 03:14:51 PM IST. You need to explicitly release the lock for other users to use the Planning Center for VMware.	:k
O Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommendation.	
Step 1: Snapshot config data.	
Load the latest configuration data for physical servers and virtual machines for analysis. You can change this data for what-if analysis. Advanced options: Select the data load options:	
Load data for selected Physical Servers	
Clean database before loading	
Load Config	
Step 2: Set analysis time period.	
Set the time period for which the measurement data in the warehouse can be analyzed corresponding to the virtual machines loaded in Step 1. The measurement data is federated from the warehouse.	
Set Time	
Step 3: Scope the infrastructure for analysis.	
Default scope includes all physical servers loaded in Step 1. Click Define Scope to go to an expert mode page where you can select the subset of physical servers that you want to be part of this analysis.	
Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to add new attributes or clean the data for physical servers and virtual machines, if required.	
Define Scope	
Step 4: Analyze virtual machine characteristics.	
Analyze the warehouse data within the time limit set in Step 2, to compute the VM-level sizing estimates using default settings (average daily utilization).	
Advanced options:	
Expension and external standy of the <u>Cancernet Content of the</u> standy and the standy of the the standy of the tent of	
Current Environment Report	
Step 5: Generate optimization plan.	
O Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommendation.	
Generate an optimization plan based on recommended environment settings.	
Advanced options: Custom settings can be made on the Edit Recommended Environment Settings page where experts can select optimization strategies, such as applicable business and technical policies, optimization	joal and so on
Generate Plan	
Optimized Plan Report	
Additional Workload Analysis Report	
V/Ware Expense Reduction Report	

Figure 87. Planning Center for VMware wizard

1. In **Step 1: Snapshot config data**, click **Load Config** to load the latest configuration data for the physical servers and the virtual machines that are related to the physical servers.

The Capacity Planner for VMware uses the IBM Tivoli Monitoring data provider that provides the configuration data for capacity planning. By using the incremental load for Capacity Planner for VMware, you can refresh the configuration data for physical servers and virtual machines. The incremental data load ensures that the configuration data and the custom modifications (such as custom tags, edited usage data, and so on) remains unchanged after the incremental load is completed.

By default, the configuration data load is incremental. You can refresh the configuration data for physical servers by clicking **Load Config**. You can use the following options while loading the configuration data:

• Load data for selected Physical Servers

By selecting the **Load data for selected Physical Servers** check box for Capacity Planner for VMware, you can load the configuration data for certain physical servers and the related virtual machines for analysis. If the data for physical server is already loaded, then this option refreshes the configuration data for selected physical server. If this option is selected with the **Clean database before loading** check box, the Capacity Planner database becomes empty before loading configuration data of selected server.

• Clean database before loading

By selecting the **Clean database before loading** you can clear the Capacity Planner database for Capacity Planner. All the configuration data that was loaded previously becomes empty before loading the latest configuration data for physical servers and related virtual machine.

For more information about loading the configuration data, see ""Loading configuration data for physical servers.""

- 2. In **Step 2: Set analysis time period**, click **Set Time** to set the time period for which the measurement data in the warehouse can be analyzed corresponding to the virtual machines that are loaded in the **step 1**.
- **3**. In **Step 3**: **Scope the infrastucture for analysis**, click **Define Scope** to open a window where you can choose physical servers to analyze. The default scope includes all the physical servers that are loaded in **Step 1**: **Snapshot config data**. You can edit the current configuration in the Edit Current Environment window, which is shown in Figure 92 on page 259.
- 4. In **Step 4: Analyze virtual machine characteristics**, click **Size VMs** to calculate the virtual machine sizing estimates based on defaults. You can customize and edit the sizing estimates in the Edit Current Environment window, which is shown in Figure 92 on page 259. The Current Environment report opens a new window that displays bar charts showing Capacity, Resource Usage, and Spare.
- 5. In **Step 5: Generate optimization plan**, click **Generate Plan** to generate a recommended environment. You must view this report after you complete the sizing step to see the computed usage and estimated spare. You can edit optimization strategies in the Edit Recommended Environment Settings window.

Capacity Planner has a multi-user access conflict removal feature. When a user starts working on a particular Capacity Planner, the system acquires a persistent lock on the Capacity Planner for this user automatically. No other user can work on the Capacity Planner while this user has the lock. The user retains this lock, even across sessions, until the lock is released by the user or until the Dashboard Application Services Hub administrator releases the lock. The administrator (with the iscadmin role) can release the lock for any user.

Loading configuration data for physical servers

You can load the configuration data for physical servers and virtual machines that are related to the physical servers to the Capacity Planner database for analysis.

Before you begin

• Install and configure IBM Tivoli Monitoring, Version 6.3 so that you can view a list of physical servers that are available.

- Install and configure the VMware VI agent, Version 7.1 or later.
- Install an agent support for VMware VI agent, Version 7.2 or later, and configure the historical data collection of the agent.
- Ensure that the Representational State Transfer (REST) service is enabled on the Tivoli Enterprise Portal Server.

Procedure

- 1. Log on to the IBM Dashboard Application Services Hub.
- 2. Open Planning Center for VMware.
- 3. In the Step 1: Snapshot config data section, select any one of the following options:
 - Clear the Load data for selected physical servers and Clean database before loading check boxes. The configuration data is updated for all the physical servers and virtual machines. In addition, the earlier configuration data remains unchanged.

Step 1: Snapshot config data.

Load the latest configuration data for physical servers and virtual machines for analysis. You can change this data for what-if analysis. Advanced options: Select the data load options:

Load data for selected Physical Servers

Clean database before loading

Load Config

Figure 88. Snapshot config data with cleared Load data for selected physical servers and Clean database before loading check boxes

• Select the Load data for selected Physical Servers check box, and clear the Clean database before loading check box.

The configuration data for selected physical servers and the virtual machines that are related to the physical servers are updated. In addition, the earlier configuration data remains unchanged.

Step 1: Snapshot config data. Load the latest configuration data for physical servers and virtual machines for analysis. You can change this data for what-if analysis. Advanced options: Select the data load options: Image: Load data for selected Physical Servers

Clean data for selected Physical Selected Clean database before loading

Figure 89. Snapshot config data with selected the Load data for selected physical servers and cleared Clean database before loading check boxes

• Select the **Load data for selected Physical Servers** and **Clean database before loading** check boxes. The configuration data in the Capacity Planner database is overwritten with the latest configuration data. The configuration data for selected physical servers and the virtual machines that are related to the physical servers are loaded. Load data for selected Physical Servers

Clean database before loading

Load Config

Figure 90. Snapshot config data with selected the Load data for selected physical servers and Clean database before loading check boxes

• Select the **Clean database before loading** check box, and clear the **Load data for selected Physical Servers** check box.

The data in the Capacity Planner database is overwritten. The complete configuration data is loaded for all the physical servers and virtual machines.

Step 1: Snapshot config data.

Load the latest configuration data for physical servers and virtual machines for analysis. You can change this data for what-if analysis. Advanced options: Select the data load options:

Load data for selected Physical Servers

Clean database before loading

Load Config

Figure 91. Snapshot config data with cleared the Load data for selected physical servers and selected Clean database before loading check boxes

- 4. If you have selected the **Load Data for selected Physical Server** check box, click **Load config**. The Selective Data Load window opens that shows the available physical servers.
- 5. In the Selective Data Load window, select one or more available physical servers for which you want to update the configuration data, and then either click **Add** or **Add All**.

Note: You can search a physical server by adding a server name in the Filter field.

6. After you have added the physical servers, click OK.

Results

The configuration data for the selected physical servers and the virtual machines that are related to the physical servers is updated.

Scenario: Consolidate and plan for growth

A capacity planner completes "what-if" planning to accommodate capacity growth while the capacity planner optimizes the environment. Consider a cloud environment that has multiple customer virtual machines that are distributed on two clusters: Production and Development. The capacity analyst can periodically analyze the cloud usage data to do capacity planning regularly. The two clusters currently have three hosts and 37 virtual machines.

About this task

An administrator or an analyst can use this scenario:

- An IT administrator wants to balance workloads to avoid a capacity bottleneck.
- A capacity analyst wants to plan for an IT infrastructure that anticipates business growth.

Procedure

- 1. Start the Planning Center, which is shown in "Planning Center for VMware wizard" on page 255.
- 2. Select the overall data set you want to analyze:
 - a. Click Load Config to load the latest configuration data for the what-if analysis.
 - b. Click **Set Time** to set the time period to analyze measurement data in Tivoli Data Warehouse.
- 3. Click Define Scope to open the Edit Current Environment window, as shown in Figure 92.

iews >> Inven	tory >> Physic	cal Servers				
Views 🔹	Actions		• 🔜 😡			Filter 🗙
Select Server	Status	Datacenter*	Cluster*	Server Hostname*	Data Source	Architecture*
	۲	HJ-IBMITMFVS	Unavailable	ibmitmfvsesx3.itmfvs1.c	Discovery	Intel(R) Xeon(R) CP X5675 @ 3.07GHz
	8	HJ-IBMITMFVS	Unavailable	ibmitmfvsesx1.itmfvs1.c	Discovery	Intel(R) Xeon(R) CP X5675 @ 3.07GHz
		HJ-IBMITMFVS	Unavailable	ibmitmfvsesx- 7.itmfvs1.com	Discovery	Intel(R) Xeon(R) CP E5506 @ 2.13GHz
	8	TestDataCenter	Unavailable	10.44.184.234	Discovery	Intel(R) Xeon(R) CP X5650 @ 2.67GHz
>	0	HJ-IBMITMFVS	Unavailable	10.44.184.233	Discovery	Intel(R) Xeon(R) CP X5650 @ 2.67GHz

Figure 92. Current Environment Physical Servers view

- 4. Select the check boxes for the servers you want to analyze. The default scope includes all physical servers that were loaded when you clicked **Load Config**.
 - You can add servers; see "Scenario: Adding more servers" on page 265 for more details.
 - You can add attribute columns by clicking Actions > Add Custom Tag. Enter a name for the new column in the Column Header field, which is shown in Figure 93 on page 260. Attribute columns are added to augment the discovered data with more attributes. You can use these attributes for writing policies as mentioned in "Detailed reference: Edit Recommended Environment settings" on page 286.

Add a New	Tag	
Column	Header:	
Tag Nam	ne:	

Figure 93. Add Custom Tag window

Note: When you add or edit a value for any tag column with tag type of MULTIPLE_VALUE, and add a semicolon (;) in the value, the semicolon is considered as a separator, and the values are saved as separate values for the tag column.

- You can double-click any cell in the table to edit configuration data. However, for bulk editing or global search and replace, use the **Export Data** and **Import Data** options on the **Actions** menu. These options can be used to complete the following functions:
 - Export or import configuration data for bulk editing
 - Global search and replace or add new attributes

Click Actions > Export Data to export data in csv format. After you edit the data, click Actions > Import Data to import the csv file.

Note: The Export and Import data supports the .zip file. The .CSV file is available in the .zip file. A required format for csv files is given in Appendix A, "CSV format for data import and export," on page 381.

• You can generate context-sensitive reports from the **Reports** menu, as shown in Figure 94 on page 261.

Important: In Capacity Planner, reports are in the context of the current view and might not give the correct result if run separately.



Figure 94. Physical Servers Inventory view Reports menu

Each report opens in a new window.

See Physical Server Inventory view reports for more details about the Capacity Planner Physical Server Inventory report.

5. Click **Size VMs** to compute the default sizes of the virtual machines, which are based on historical usage. The sizes can be viewed in the view that is shown in Figure 95. A custom sizing profile can be applied in the same view, as shown in the next step. Other detailed sizing options are mentioned in Compute Usage.

ws >	> Virtual Machi	nes Utilization						
/iews	- Act	ions 🔻	Reports	• 📑 I			Filter	×
	VM Hostname	VM Name	Server Hostnam e	CPU Usage [MHz]	Memory Usage [MB]	Storage Usage [Kbps]	Network Usage [Kbps]	CPU Deviati
~	Unavailable	ibmtest08	ibmucsesx1	589	256	756	1,024	
~	Unavailable	ibmtest07	ibmucsesx1	589	256	756	1,024	
~	Unavailable	ibmtest06	ibmucsesx1	589	256	756	1,024	
~	Unavailable	ibmtest05	ibmucsesx1	589	256	756	1,024	
~	Unavailable	ibmtest04	ibmucsesx1	589	256	756	1,024	
~	Unavailable	ibmtest03	ibmucsesx1	589	256	756	1,024	
~	win2k332bit	ibmtest02	ibmucsesx1	589	256	756	1,024	

Figure 95. Virtual Machines Utilization view

• You can apply a growth profile by using the Actions menu, as shown in Figure 96 on page 262.

Edit VMv	ware Cu	rrent Envi	ronment		
Views >:	> Virtual Ma	chines Utilizatio	n		
Views	•	Actions 💌	Reports	Ŧ	
		Compute Usa	ge		
	VM Hostn	Generate Worl Edit Usage	doad Stability	Гуре	Server Hostname
	Unavailab	le ik	mtest09		ibmucsesx1.itmfvs1.co

Figure 96. Virtual Machines Utilization view Actions menu

Click Edit Usage on this menu to apply a growth profile, as shown in Figure 97.

🔵 Absolute Value 🖲 Gr	owth	
CPU Usage:		9
Memory Usage:		9
Storage Usage:		9
Network Usage:		9

Figure 97. Edit Usage window

• You can generate context-sensitive reports from the **Reports** menu. An example is shown in Figure 98 on page 263.

Important: In Capacity Planner, reports are in the context of the current view and might not give the correct result if run separately.

Edit VMv	ware Current E	nviro	nment
Views >:	> Virtual Machines Uti	lization	
Views		•	Reports 🔹 🛛 🔤
		1	Current Environment
	VM Hostname	VMP	Utilization Aggregated Timeseries Utilization Detailed Timeseries
	Unavailable	ibmte	est09 ibmucsesx1.itmfvs1.co

Figure 98. Virtual Machines Utilization view Reports menu

Each report opens in a new window.

For details of these reports, see the Virtual Machine Utilization reports.

• You can select a value for the **Candidate** attribute as **Target** only. An example is shown in Figure 99. If you select the **Target** candidate for the server that has some virtual machines that already exist on the server. The CPU and memory demand of the virtual machine that already exists on the target only server is also considered while computing available capacity of the server. If there is no memory or CPU demand, by default, the reservation value is used. If you want to use the CPU and memory limit values for the virtual machine while calculating server capacity, you can make the following edit in the analytics.properties file: Set **PLACEMENT_RESERVATION_IF_NO_USAGE_DATA** to MAXIMUM

Note: After you edit the file, restart the server.

Edit Canc	lidate	Tag
Source		
	ок	Cancel

Figure 99. Edit Candidate Tag window

If the CPU Usage and Memory Usage values are null (the usage is not edited or computed), the placement engine must consider Minimum or Maximum Memory, or Minimum or Maximum CPU from the CFG_VIRTUAL_MACHINE table as a suggested reservation based on the configuration that is provided in the analytics.properties file.

The analytics.properties file is in the following location:

- On Windows systems: DASH_HOME\profile\installedApps\Node_Name\isc.ear\ AnalyticsWebUI.war\WEB-INF\classes
- On operating systems other than Windows: DASH_HOME/profile/installedApps/Node_Name/ isc.ear/AnalyticsWebUI.war/WEB-INF/classes

- 6. Click Edit Recommended Environment Settings to open a new window, where you can edit optimization strategies. Choose rules to apply in optimization. For more information about optimization rules, see "Detailed reference: Edit Recommended Environment settings" on page 286.
- 7. Click Generate Plan to generate a recommended environment, as shown in Figure 100.



Figure 100. Generate Optimization Plan section of Planning Center page

Results

The optimization plan opens in a new window. For this example, the optimization plan shows that, based on historical usage data and rules, virtual machines can be consolidated on two physical servers instead of three physical servers in the current environment. The optimization plan contains different types of output:

• A table that compares overall current resources with recommended resources, as shown in Figure 101

	Current		Recommendation			
Physical Servers	4	4				
Virtual Machines	8			8		
	CPU (GHz) Men	nory (GB)	CPU (GHz)M	emory (GB)		
Total Capacity	15.62	19.53	15.62	19.53		
Total Reservation	7.81	7.81	7.81	7.81		
Total Unused Capacity (excluding headroom)	6.25	9.77	6.25	9.77		
Capacity Efficiency Index	50.00		<mark>50.0</mark>	00		

Report As Of : Jan 28, 2013 3:14:43 PM

Figure 101. Table showing recommended resources

Note: See "Detailed reference: Capacity efficiency and performance risk indices" on page 299 for an explanation of *capacity efficiency index* and *performance risk index*. By consolidating the number of servers, we can improve the capacity efficiency index.

• Detailed output for individual data centers, clusters, physical servers, and virtual machines, which are displayed in bar charts, as shown in Figure 102 on page 265



Figure 102. Individual bar charts comparing current resources with recommended resources

• Individual results in table form. Figure 103 shows the placement recommendation, that is, physical server to virtual machine mapping, along with the recommended reservation for CPU and memory. The reservation that is recommended can be implemented by the administrator in the case of virtual machines where no risk of performance degradation must be taken. The placement recommendation can be implemented independently of the reservation recommendation.

Virtual Ma	firtual Machines :										
Virtual Machine	Number Of vCPUs	Current Reservation CPU (GHz)	Recommended Reservation CPU (GHz)	Current Reservation Memory (GB)	Recommended Reservation Memory (GB)	Performance Risk Index	OS	Middleware Name			
S2V1	1	0.977	0.977	0.977	0.977						
83V1	1	0.977	0.977	0.977	0.977						
84V1	1	0.977	0.977	0.977	0.977						

Figure 103. Individual tabular results comparing current resources with recommended resources

Any virtual machines that cannot be placed are listed at the end of the report. This situation usually occurs either when there is no CPU or memory usage data for those virtual machines, or no capacity is remained on the target servers.

Scenario: Adding more servers

A capacity planner adds more servers in what-if analysis to remove a capacity shortfall with existing servers.

Procedure

- 1. On the Planning Center for VMware page, which is shown in "Planning Center for VMware wizard" on page 255, click **Define Scope**, which opens the Edit Current Environment page, which is shown in Figure 92 on page 259.
- 2. Click Actions > Add Server, as shown in Figure 104,



Figure 104. Physical Servers Actions menu

which opens the Create Server Instances window, as shown in Figure 105.

and the second second					-
Belect Model:	BladeCenter HS12, Intel Core 2 Duo E6305, 2 Cores				
Number of server instances:	1				

Figure 105. Create Server Instances window

In this window, you can select a model from the list and enter the number of servers that you want to add.

The new servers are displayed at the bottom of the list of servers on the Edit Current Environment page. Ensure that the check boxes are selected that correspond to the new servers to be added in scope.

3. On the Planning Center for VMware page, which is shown in "Planning Center for VMware wizard" on page 255, click **Generate Plan** to regenerate the optimization plan with the increased scope.

Results

An example optimization plan is shown in Figure 106.

	Curren	nt	Recom	nendation
Physical Servers	4			4
Virtual Machines	8			8
	CPU (GHz) Mer	mory (GB)	CPU (GHz)	Memory (GB)
Total Capacity	15.62	19.53	15.62	19.53
Total Reservation	7.81	7.81	7.81	7.81
Total Unused Capacity (excluding headroom)	6.25	9.77	6.25	9.77
Capacity Efficiency Index	<mark>50.00</mark>		5	0.00

Figure 106. Optimization plan: adding servers

Note that this optimization plan shows four servers in the recommended environment, which implies that an additional server is required to address a shortfall in capacity.

Scenario: Adding new virtual machines

A capacity planner wants to allocate new virtual machines in a cloud.

About this task

Consider a cloud environment that has multiple customer virtual machines that are distributed on two clusters: Production and Development. The capacity analyst can periodically analyze the cloud usage data to do capacity planning regularly. The two clusters currently have three hosts and 37 virtual machines.

Consider a scenario where three new virtual machines must be added to the cloud for a specific customer. The planner can assess where to allocate the resources: at a cluster level or optionally at a host level. The request can be actual or a predicted future request for what-if analysis.

Procedure

- 1. On the PlanningCenter page, which is shown in "Planning Center for VMware wizard" on page 255, click **Define Scope** to open the Edit Current Environment page, which is shown in Figure 92 on page 259.
- 2. Select the checkboxes corresponding to the Production and Development servers.
- **3**. Click **Views** > **Virtual Machines** to open the virtual machines view. This view shows the virtual machines from the selected servers. A column specifying the customers that the virtual machines belong to can be added as described in adding a custom tag.
- 4. Click **Actions** > **Add Virtual Machine** to add a virtual machine for the customer. The new virtual machine is displayed in the scope.
- 5. On the PlanningCenter page, click **Edit Recommended Environment Settings**, which is shown in Figure 100 on page 264, to open a new window, which is shown in Figure 129 on page 286. In this window, you can detail customer-specific rules. For more information about rules, see "Detailed reference: Edit Recommended Environment settings" on page 286.
- 6. On the PlanningCenter page, click Generate Plan.

Results

A new optimization plan is generated, including the virtual machine that you added to the cloud. An example of the output is shown in Figure 107 on page 268:

	Currer	ıt	Recom	nendation
Physical Servers	4			4
Virtual Machines	8			8
	CPU (GHz) Mei	nory (GB)	CPU (GHz)	Memory (GB)
Total Capacity	15.62	19.53	15.62	19.53
Total Reservation	7.81	7.81	7.81	7.81
Total Unused Capacity (excluding headroom)	6.25	9.77	6.25	9.77
Capacity Efficiency Index	50.00		5	0.00

Figure 107. Optimization plan including new virtual machine

Note: Information about the mapping of the added virtual machines to the recommended physical server is contained in the detailed section of the optimization plan.

Cloning physical servers and virtual machines

With this feature, you can create clones of existing physical servers and virtual machines in the Capacity Planner. The cloned entries have the same configuration and utilization profile as the original entries. This feature can be helpful in what-if analyses where you want to analyze the impact of having new workloads that are similar to existing workloads in the environment.

About this task

The Cloning physical servers and virtual machines feature has the following benefits:

- You are not required to add physical servers or virtual machines and complete various required fields.
- All field values from the physical servers or virtual machines entry that is being cloned is automatically copied to cloned entries.
- You can associate all clone virtual machines to a physical server that is available in the working set.
- You can choose whether to clone sizing data for virtual machines that are being cloned.
- You can choose whether to copy tags for the clone entries that are being created.

Cloning physical servers Procedure

Use the following steps to clone a physical server:

- 1. Start the Planning Center page for VMware.
- 2. Go to the Edit Current Environment page.
- 3. On the menu, click Actions > Clone Physical Servers (Figure 108 on page 269).



Figure 108. Edit VMware Current Environment page

The Clone Physical Servers window is displayed (Figure 109).

wailable Physical Servers		Selected Physical Servers	
Filter	×	Filter	×
Clone_10.44.184.150_1375422889047_ Clone_10.44.184.155_1375422889197_		10.44.184.150 10.44.184.155 10.44.184.161 10.44.184.175 10.44.232.199	

Figure 109. Clone Physical Servers window

- 4. Move the physical servers from the **Available Physical Servers** list to the **Selected Physical Servers** list (Figure 109).
- 5. Click **Next**. The second pane of the Clone Physical Servers window is displayed (Figure 110 on page 270).

	Filter	X			
Physical Server Name	Number of Instances*		Clone Name Prefix*		Copy Server Tags
ibmitmfvsesx3.itmfvs1.cor	1	* *	Clone_ibmitmfvsesx3.itm		
ib <mark>m</mark> itmfvsesx4.itmfvs1.cor	1	*	Clone_ibmitmfvsesx4.itm	nfvs1.com	
Total: 2 Selected: o			<1→	5 10 25 5	0 100 All 1

Figure 110. Clone Physical Servers second page

- 6. In the Number of Instances field, select a value.
 - The number indicates the number of clone entries for the physical servers that you want to create.
 - The default value is 1.
- 7. In the **Clone Name Prefix** field, enter a value.
 - Each clone entry is prefixed with the string value provided in this field and with a unique time stamp value.
 - The default value is Clone_PHYSICAL_SERVER_NAME.
- 8. Select the Copy Server Tags check box to enable the tags.
 - If enabled, all tags for the selected physical server are copied to clone entries.
 - By default, this check box is selected.
- 9. Click Clone.

Cloning virtual machines

Procedure

- 1. Start the Planning Center page for VMware.
- 2. Go to the Edit Current Environment page.
- 3. Click Views > Inventory > Virtual Machines (Figure 111 on page 271).

Edit	VMware	Current E	nvironme	nt
Vi	iews >> Inven	tory >> Virtual N	lachines	
	Views *	Actions -	Reports	-
	No filte	r applied		
	Action	Server Hostname*	VM Hostname*	VM Name*
		ibmitmfvsesx	IBMESX2V13_	Clone_IBMESX2-V13-RHEL5- 32Bit_1373374508815_0

Figure 111. Edit VMware Current Environment page for virtual machines

 Click Actions > Clone Virtual Machines (Figure 112). The Clone Virtual Machines window is displayed (Figure 113 on page 272).



Figure 112. Edit VMware Current Environment Actions tab

Clone Virtual Machines Available Virtual Machines Selected Virtual Machines Filter × Filter × * IBMESX2-V12-RHEL5-64Bit IBMESX2-V11-SLES10-32Bit IBMESX2-V13-RHEL5-32Bit IBMESX2-V15-RHEL5-32Bit >> IBMESX2-V16-SUSE11-64 Bit < IBMESX2-V19-Win2K8-R2 \ll IBMESX2-V21-RHEL5-64Bit IBMESX2-V22-Win2k3-32Bit IBMESX2-V23-Win2k8-32Bit Previous Next Clone Cancel

Figure 113. Clone Virtual Machines window

- 5. Move the virtual machine from the **Available Virtual Machines** list to the **Selected Virtual Machine** list.
- 6. Click Next. The second pane of the Clone Virtual Machines window is displayed (Figure 114).

			Filter X	
Virtual Machine Name	Number of Instances*	Clone Name Prefix*	Copy Server Tags	Phy:
IBMESX2-V11-SLES10-32Bit	1	Clone_IBMESX2-V11-SLES10-32Bit		ibi
IBMESX2-V13-RHEL5-32Bit	1	Clone IBMESX2-V13-RHEI 5-32Bit		ih.
Total: 2 Selected: o		4 1 b 5	10 25 50 100	All 4

Figure 114. Clone Virtual Machines second page

- 7. In the Number of Instances field, select a value.
 - The number indicates the number of clone entries for the selected virtual machine that you want to create.
 - The default value is 1.
- 8. In the **Clone Name Prefix** field, select a value.

- Each clone entry is prefixed with the string value that is provided in this field and with a unique time stamp value.
- The default value is Clone_VIRTUAL_MACHINE.
- 9. Select the Copy Server Tags check box to enable the tags.
 - If enabled, all tags for the selected physical servers are copied to clone entries.
 - By default, this check box is selected.
- 10. In the Physical Server Name field, select a physical server.
 - By default, the physical server that is displayed is the one to which the selected virtual machine is pointing.
 - All physical servers that are in the working set are populated in the Physical Server Name field.
- 11. Select the Copy Sizing Data check box to enable the copying of sizing data.
 - If enabled, sizing data values for selected virtual machines are copied to clone entries.
 - By default, this check box is selected.
 - Special Case: The **Copy Sizing Data** check box is selected, and the selected virtual machine is a discovered virtual machine with no sizing data values. The following are the sizing data values for the CPU Usage [MHz] and Memory Usage [MB] columns in the **Virtual Machines Utilization** view from the virtual machine inventory view:
 - The value for CPU Usage [MHz] is Minimum CPU [MHz]
 - The value for Memory Usage [MB] is **Minimum Memory** [**MB**]
- 12. Click Clone.

Detailed reference: Edit Current Environment

The Edit Current Environment window presents a spreadsheet-like view within a browser window. In this view, you can view and edit the configuration and utilization profile, in addition, define scope.

Data that is loaded in this view is a snapshot of the current environment that is loaded by using **Load config**, or by adding servers, as described in "Scenario: Adding more servers" on page 265. You can then add to the data.

Physical server inventory view

ews >> Invent	ory >> Physic	cal Servers				
/iews 🔻	Actions	▼ Reports	- 🔤 🐼			Filter 🗙
Select Server	Status	Datacenter*	Cluster*	Server Hostname*	Data Source	Architecture*
	۲	TestDataCenter01	Unavailable	10.44.184.241	Discovery	Intel(R) Xeon(R) CPL
•	0	HJ-IBMITMFVS	Unavailable	ibmucsesx1.itmfvs1.co	Discovery	Intel(R) Xeon(R) CPL E7520 @ 1.87GHz
		HJ-IBMITMFVS	Unavailable	ibmitmfvssap3.itmfvs1.	Discovery	Intel(R) Xeon(R) CPL E5506 @ 2.13GHz
~	0	HJ-IBMITMFVS	Unavailable	ibmitmfvssap2.itmfvs1.	Discovery	Intel(R) Xeon(R) CPL X5675 @ 3.07GHz
~	0	HJ-IBMITMFVS	Unavailable	ibmitmfvsesx9.itmfvs1.c	Discovery	Intel(R) Xeon(R) CPL

Figure 115. Current Environment Physical Servers view

The Physical Servers view presents the physical server configuration data in the current environment. To define the scope, you can select a subset of physical servers to include in the analysis session. The candidate column in the view has *Source* and *Target* tags. These tags determine how the server is used in optimization generation.

- The *Source* tag ensures that during optimization, virtual machines from this physical server can only be moved to another physical server.
- The *Target* tag ensures that during optimization, this physical server can only receive virtual machines from another physical server.
- *Source* and *Target* tags together ensure that during optimization, virtual machines from the server can be moved for optimization and the server can also be used to place virtual machines.

The Status column indicates whether any problems exist with matching performance benchmark data for the particular server. Additionally, you can click the Status column to see benchmark search results. A

green check mark \blacksquare icon indicates that a single matching benchmark entry exists. A yellow triangle \triangle icon indicates that multiple benchmark matches to the server architecture exist, but an approximate

benchmark value is assigned to the server. An X in a red circle 🔇 icon indicates that no benchmark match exists.

You can verify the search result by clicking the Status column. If multiple matches or no match, you can use the information from the search results to modify the architecture column to narrow down the match to the correct benchmark.

For example, for a server with architecture Intel Xeon CPU E5506 @2.13GHz, the default result is a single match:



Figure 116. Benchmark Search Results window

You can enter the correct string from the Benchmark Search Results window to the **Architecture** field in the physical server inventory view to narrow the match. In this example, if the architecture is updated to IBM BladeCenter HS21 (Intel Xeon CPU E5506 2.13GHz) from the search results, a single match is displayed:

ws >> Inven	tory >> Physical Server	s				
ïews 💌	Actions 👻	Reports 🔹 🕴 🖥			Filter	×
Status	Datacenter*	Cluster*	Server Hostname*	Data Source	Architecture*	Numbe
۲	TestDataCenter01	Unavailable	10.44.184.241	Discovery	Intel(R) Xeon(R) CPU X5650 @ 2.67GHz	-
۲	HJ-IBMITMFVS	Unavailable	ibmucsesx1.itmfvs1.coi	Discovery	Intel(R) Xeon(R) CPU E7520 @ 1.87GHz	
	HJ-IBMITMFVS	Unavailable	ibmitmfvssap3.itmfvs1.	Discovery	Intel(R) Xeon(R) CPU E5506 @ 2.13GHz	
0	HJ-IBMITMFVS	Unavailable	ibmitmfvssap2.itmfvs1.	Discovery	Intel(R) Xeon(R) CPU X5675 @ 3.07GHz	
8	HJ-IBMITMFVS	Unavailable	ibmitmfvsesx9.itmfvs1.(Discovery	Intel(R) Xeon(R) CPU	

Figure 117. Benchmark Search Results window, single match

The built-in benchmark database of the Capacity Planner might not be complete with all the types and configuration variations (for example, vendor, model, architecture, number of cores, and so on) of

physical servers that are available. Therefore, the Capacity Planner might not find an accurate benchmark value for the missing types and configuration variations.

In this case, you can feed in your environment-specific server data into the Capacity Planner database. This supplementary data is searched first before the built-in database by the Capacity Planner to find the benchmark values. For more information about the custom benchmark data, see the CUSTOM_USER_DEFINED_BENCHMARK.csv section in Appendix B, "Editing knowledge base," on page 385.

If the environment is homogeneous, use the raw CPU capacity option for analysis. Enable this option by updating the normalization benchmark setting as described in "Normalization benchmark setting" on page 296.

The following actions are available on the Actions menu:

Add Server

You can add more physical servers to this view. These servers can be used to provide additional hardware, or to try what-if scenarios. When you choose this action, a window opens with a list of available models, as shown in Figure 105 on page 266. You can select the appropriate model and create multiple instances as needed.

When you click **Create**, a new row is added to the inventory grid with architecture details populated from the knowledge base data.

Add Custom Tag

You can extend and augment the discovered data with user-defined attributes (tags). When you select this action, you are prompted to provide the tag name and a new column is added to the inventory grid. You can then add values for this column.

These tags can be used to formulate and apply rules during optimization. For more information, see "Detailed reference: Edit Recommended Environment settings" on page 286.

Export Data

You can download the data in the Physical Servers inventory view to a CSV file, which can be edited offline to add missing information.

Import Data

You can import a CSV file that was downloaded using Export Data.

Important: Note the following when you edit this file:

- The first column, Physical_Server_PK, must not be edited. If an extra physical server is added, you must keep the Physical_Server_PK column empty.
- Only previously NULL or blank data is updated. If data exists in a specified column, the data cannot be updated if edited.

Note: The Export and Import data supports the .zip file. The .CSV file is available in the .zip file.

See Appendix A, "CSV format for data import and export," on page 381 for more information about editing CSV files.

Reports

The following reports are available:

Capacity Planner Physical Server Inventory

The Capacity Planner Physical Server Inventory report presents an overall view of the physical environment in the current Capacity Planner session. The report contains a summary table of the inventory, and bar charts that are organized by hypervisor name and version, or data center and
cluster, as shown in Figure 118.

v7.2 Capacity Planner Physical Server Inventory

About this report	Report As Of : Jan 28, 2013	3 3:30:04 PM
Physical Server Summary		
	Number of Physical servers Number of Physical CPUs CPU Speed (GHz) Memory Installed (GB)	4 16 15.62 19.53
Server Distribution: by Hypervisor type and v	ersion	
45		
4		
3.5		
3		
2.5		
2		
15		
1		
0.5		
٥	4.0	
	VM	

Figure 118. Capacity Planner Physical Server Inventory report

Virtual machine inventory view

ews >> Inventory >> Vir	tual Machines					
Views 👻 Action	ns 🔻 Report	s • 📑		F	ilter ×	-0
Server Hostname*	VM Hostname*	VM Name*	Data Source	Number of CPU Cores*	Minimum CPU [MHz]*	
ibmucsesx1.itmfvs1.coi	Unavailable	ibmtest09	Discovery	1	0	
ibmucsesx1.itmfvs1.coi	Unavailable	ibmtest08	Discovery	1	0	
ibmucsesx1.itmfvs1.coi	Unavailable	ibmtest07	Discovery	1	0	
ibmucsesx1.itmfvs1.coi	Unavailable	ibmtest06	Discovery	1	0	

Figure 119. Current Environment Virtual Machines view

The Virtual Machines view presents the virtual machine configuration data from the selected set of physical servers. You can make corrections, add missing data or add new attributes (tags) that cannot be discovered in the current environment (such as middleware details, or tags such as criticality, SLA factors, and so on).

The following actions are available on the Actions menu:

Add Virtual Machine

You can add a virtual machine to provide for future workloads. A new row is added to the inventory grid where you can populate details of the new virtual machine. Virtual Machines can be added to any server in the working set, if spare capacity exists for CPU and memory reservations.

Add Custom Tag

You can extend and augment the discovered data with user-defined attributes (tags). When you select this action, you are prompted to provide the tag name and a new column is added to the inventory grid. You can then add values for this column.

These tags can be used to formulate and apply rules during optimization. For more information, see "Detailed reference: Edit Recommended Environment settings" on page 286.

Export Data

You can download the data in the Virtual Machines inventory view to a CSV file, which can be edited offline to add missing information.

Import Data

You can import a CSV file that was downloaded by using Export Data.

Important: Note the following when you edit this file:

• The first column, VIRTUAL_MACHINE_PK, must not be edited. If an extra virtual machine is added, you must keep the VIRTUAL_MACHINE_PK column empty.

• Only previously NULL or blank data is updated. If data exists in a specified column, it cannot be updated if edited.

Reports

Virtual Machine Inventory

The Virtual Machine Inventory report presents an overall view of the virtual environment in the current capacity planning session. The report contains a summary table of the inventory and overall organizational graphical representations that are organized by datacenter and cluster, operating system name and version, and middleware name and version. The report also contains bar charts that are organized by hypervisor name and version, or data center and cluster, as shown in Figure 120.



Figure 120. Virtual Machine Inventory report

Virtual Machines Utilization View

ews	>> Virtual Machines	Hilization					
/iew:	s 🔹 Actions	▼ Reports	- -			Filter	40 4
	VM Hostname	VM Name	Physical Server Key*	Server Hostname	CPU Usage [MHz]	Memory Usage [MB]	Memory Dei
	Ser1V1	Ser1V1	33	server1	1,000	1,000	-
	Ser1V2	Ser1V2	33	server1	1,000	1,000	
	S2V1	S2V1	35	server2	1,000	1,000	
	S2V2	S2V2	35	server2	1,000	1,000	
	S3V1	S3V1	37	server3	1,000	1,000	
	S3V2	S3V2	37	server3	1,000	1,000	
C	S4V1	S4V1	39	server4	1,000	1,000	
100	S4V2	S4V2	39	server4	1.000	1.000	

Figure 121. Virtual Machines Utilization view

A capacity analyst can use this view to view and analyze the utilization profiles of the virtual machines on the servers currently in the working set. The virtual machine utilization data is federated from the Tivoli Data Warehouse tables. The capacity analyst can define the time interval for analysis of data on the PlanningCenter (**Step 2: Set analysis time period**).

The following actions are available on the **Actions** menu:

Compute Usage

You can compute the usage requirement of virtual machines by using different parameters, as shown in Figure 122 on page 281:

			Summarizatio	on Selected Su	ummarizatio
CPU:	Maximum	•	Weekly 📘	All First Wee Second V	k Veek 👻
Memory:	Maximum	•	Weekly 📘	All First Wee Second V	k Veek 👻
Storage:	Maximum	•	Weekly 📘	All First Wee Second V	k Veek 👻
Network:	Maximum	•	Weekly 📘	All First Wee Second V	k Veek 👻

Figure 122. Compute Usage window

Compute Usage calculates the sizing of the virtual machine for the parameters: CPU, memory, network bandwidth, and disk I/O usage. This sizing is done by analyzing the utilization data available in Tivoli Data Warehouse based on the summarization and aggregation levels specified. Aggregation levels available are **Average**, **Minimum**, **Maximum**, and **90th Percentile**. Summarization levels available are **Hourly**, **Daily**, **Weekly**, **Monthly**, and **Yearly**. The available values in the **Selected Summarization** field depend on which value was selected in the **Summarization** field.

Important: Usage numbers are generated only for virtual machines that have utilization data that is collected in the Tivoli Data Warehouse.

Generate Workload Stability Type

The Generate Stability Characteristic Parameters window is shown in Figure 123 on page 282:

	Measurement Type	
CPU:	Average 💌	
Memory:	Average 💌	
Storage:	Average 💌	
Network:	Average	

Figure 123. Generate Stability Characteristic Parameters window

Generate Workload Stability Type analyzes the hourly utilization data for a virtual machine and determines whether the resource utilization is stable or unstable, depending on the variation in usage.

Edit Usage

The Edit Usage window is shown in Figure 124:

Absolute Value O Growth	1	
CPU Usage:		MHz
Memory Usage:		МВ
Storage Usage:		Kbps
Network Usage:		Kbps

Figure 124. Edit Usage window

You can manually edit or *Adjust-for-growth* the resource usage. You can apply different growth profiles to adjust as needed. The usage parameters can be specified in absolute units, for example, 1024 MHz CPU, or a growth percentage can be applied, for example, add 10% growth to memory.

Reports

The following reports are available:

Utilization Aggregated Timeseries report

This report can be used to identify utilization patterns of virtual machines. Because you can view aggregations of multiple virtual machines at a time, you can also identify correlations in the resource utilizations. You can use these observations to determine the usage sizing summarization level. An example graph is shown in Figure 125.



Figure 125. Utilization Aggregated Timeseries report

Utilization Detailed Timeseries report

This report helps you identify any data gaps in the utilization data that is collected for the virtual machines Data points come directly from aggregated measurement tables in utilization schema. An example graph is shown in Figure 126 on page 284.

Network Utilization



Figure 126. Capacity Planner Utilization Detailed Timeseries report

Current Environment report

The Capacity Planner Current Environment report shows the recommended usage and sizing for each virtual machine that is based on time series utilization data and sizing process: summarization, granularity type, and selections. The report shows the capacity gaps, if any, in the current environment before any optimization. Example output is shown in Figure 127 on page 285 and Figure 128 on page 285.



Figure 127. Current Environment report



Figure 128. Current Environment report (continued)

Detailed reference: Edit Recommended Environment settings

Capacity Planner provides a framework to capture various business and technical rules and apply them during optimization. The policies ensure that when an optimization is done for an environment, the generated recommendation complies with the stated business requirements, and follows technical best practices.

Rule framework overview

Figure 129 shows the window where you can edit rules:

Colocation//	Anti-colocation
Active	Rule Instances
	Do not colocate VMs with DB2 and WAS
	Place DB2 workloads on servers with hypervisor as VMWare ESXi
Boundary	
Active	Rule Instances
	Create a Boundary for Critical VMs
	Create a Boundary for Win2003 32-bit VMs
Utilization	
Activo	Rule Instances

Figure 129. Edit Recommended Environment Settings window

This window displays all rule instances from the database. Select the check box that is displayed corresponding to a rule instance to enable that rule. Clear the check box beside a rule instance to disable that rule.

Capacity Planner provides the following rule templates that can be used to encode rule instances that are catering to business, and in addition, technical requirements:

Colocation and anti-colocation

This rule template is the most commonly used. You choose two sets of virtual machines and specify whether these two sets must share the same host or must not share the same host. In the anti-colocation case, no two virtual machines from different sets share the same host. In the

colocation case, you want to put all of the set on the same host. If the virtual machines do not all fit on the same host, place the excess virtual machines on the next available host.

A variant of the colocation template specifies a set of virtual machines to be *pinned* to a set of hosts. In this case, instead of both sets being virtual machines, a set of virtual machines is placed on a specific set of container hosts.

The following examples of colocation and anti-colocation rule instances illustrate the nature of the templates:

Colocate Windows 2003 VMs

This rule consolidates Windows virtual machines on a group of hosts and reduces the number of physical processors that the Windows system is deployed on, thus reducing the license cost.

Anti-colocate customer 1 and customer 2 VMs

This rule imposes a business rule restriction to share hosts among competitor virtual machines in a shared cloud.

Use x3950 high memory configuration hosts in Net1 network for all VMs running databases This rule ensures a technical best practice for database consolidation.

In all these examples, it is assumed that the configuration or usage attributes that are used to create these sets of virtual machines and hosts are known. You discover or create these attributes in the Edit Current Environment page.

Boundary

In this template, you can pick a set of virtual machines and specify them to be placed on an isolated set of containers, which is not shared by any other virtual machine that does not belong to that set. The container is a host or set of hosts. This rule creates an isolation zone for the specified set of virtual machines. A boundary rule can also be created by several anti-colocation policies.

The following example shows a boundary rule:

Create a boundary for all WebSphere TEST VMs

This rule prevents these virtual machines from sharing a host with any other type of virtual machine.

Utilization

This template specifies a growth factor for the processor resource demand of the virtual machines in excess of the current values of processor use that are displayed in the Virtual Machine Utilization view of the Edit Current Environment page.

The following example shows a utilization rule:

Use 20% CPU growth for all Critical VMs

This rule specifies a 20% increase in processor demand for all virtual machines that have a *Critical* tag. This rule ensures that enough safety margin is kept to maintain the performance of the applications that are running in these virtual machines.

Rule syntax

Colocation and anti-colocation

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
<group1 operator="GROUPING_OPERATOR">
<literal>
<path>DATABASE_PATH</path>
<operator>"PATH_OPERATOR"</operator>
<term>PATH_VALUE</term>
</literal>
<literal>
<path>DATABASE_PATH</path>
```

```
<operator>"PATH OPERATOR"</operator>
  <term>PATH VALUE</term>
 </literal>
</group1>
<constraint>CONSTRAINT TYPE</constraint>
<group2 operator="GROUPING_OPERATOR">
 <literal>
  <path>DATABASE PATH</path>
  <operator>"PATH_OPERATOR"</operator>
  <term>PATH VALUE</term>
 </literal>
<literal>
  <path>DATABASE PATH</path>
  <operator>"PATH OPERATOR"</operator>
  <term>PATH_VALUE</term>
 </literal>
</group2>
<type>COLOCATION_TYPE</type>
```

</rule>

where

GROUPING_OPERATOR

AND or OR

PATH_OPERATOR

EQ or NEQ

DATABASE PATH

See "Supported database paths for writing rules" on page 291

PATH_VALUE

All possible values that are defined for DATABASE_PATH

CONSTRAINT TYPE

Inclusion or Exclusion

COLOCATION_TYPE

- **S2S** Specifies a constraint for two groups of virtual machines
- **S2T** Specifies a constraint for a group of virtual machines and a group of physical servers

Example colocation rule: Colocate all Windows 2003 workloads

```
CFG_VIRTUAL_MACHINE.OPERATING_SYSTEM_NAME</path>
<operator>EQ</operator>
<term>Windows 2003</term>
</literal>
</group2>
<type>S2S</type>
</rule>
```

Boundary

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
 <if>
  <antecedent operator="GROUPING_OPERATOR">
  <literal>
   <path>DATABASE PATH</path>
   <operator>"PATH_OPERATOR"</operator>
   <term>PATH_VALUE</term>
   </literal>
<literal>
   <path>DATABASE_PATH</path>
   <operator>"PATH_OPERATOR"</operator>
   <term>PATH VALUE</term>
  </literal>
 </antecedent>
 </if>
 <then>
 <dependent>
  <literal>
   <path>Boundary</path>
   <operator>EQ</operator>
   <term>BOUNDARY NAME</term>
  </literal>
 </dependent>
</then>
</rule>
```

where

GROUPING_OPERATOR AND or OR

PATH_OPERATOR

EQ or NEQ

DATABASE PATH

See "Supported database paths for writing rules" on page 291

PATH_VALUE

All possible values that are defined for DATABASE_PATH

BOUNDARY NAME

Unique name for boundary

Note: If multiple boundary rules are specified with same name, then they are all treated as one boundary.

Example boundary rule: Do not mix workloads that include a production environment with other workloads

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
<if>
<antecedent operator="AND">
```

```
<literal>
        <path>Virtual Machine.ANL VM PLACEMENT SET V*
              ANL VIRTUAL MACHINE TAG MAP.SERVER TAGS PK.TAG TYPE</path>
            <operator>EQ</operator>
            <term>ENVIRONMENT</term>
           </literal>
           teral>
        <path>Virtual Machine.ANL VM PLACEMENT SET V*
              ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_NAME</path>
            <operator>EQ</operator>
            <term>Production</term>
           </literal>
          </antecedent>
         </if>
         <then>
          <dependent>
           <literal>
           <path>Boundary</path>
            <operator>EQ</operator>
            <term>Environment.Production</term>
           </literal>
          </dependent>
        </then>
        </rule>
Utilization
        <?xml version="1.0" encoding="UTF-8"?>
```

```
<rule>
<if>
  <antecedent operator="GROUPING_OPERATOR">
  <literal>
   <path>DATABASE PATH</path>
   <operator>"PATH_OPERATOR"</operator>
   <term>PATH_VALUE</term>
   </literal>
<literal>
    <path>DATABASE PATH</path>
   <operator>"PATH OPERATOR"</operator>
   <term>PATH_VALUE</term>
   </literal>
 </antecedent>
 </if>
 <then>
  <dependent>
  <literal>
   <path>target.capacity</path>
   <operator>Add</operator>
   <term>NUMBER</term>
  </literal>
 </dependent>
</then>
</rule>
```

where

GROUPING_OPERATOR

AND or OR

PATH_OPERATOR

EQ or NEQ

DATABASE PATH

See "Supported database paths for writing rules" on page 291

PATH_VALUE

All possible values that are defined for DATABASE_PATH

NUMBER Valid number in percentages (positive or negative)

Example utilization rule: apply 20% buffer on DB2 workloads

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
<if>
 <antecedent operator="AND">
  <literal>
   <path>Virtual Machine.ANL VM PLACEMENT SET V*
             CFG_VIRTUAL_MACHINE.MIDDLEWARE_NAME</path>
   <operator>EQ
   <term>IBM Universal Database</term>
  </literal>
 </antecedent>
</if>
<then>
 <dependent>
  <literal>
   <path>target.capacity</path>
   <operator>Add</operator>
   <term>20</term>
  </literal>
 </dependent>
</then>
</rule>
```

Supported database paths for writing rules

Commonly used discovered attributes and their corresponding database paths:

Attribute	Database path
Host name	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER.HOST_NAME
Cluster name	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER. SERVER_POOL_NAME
Data center name	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER. DATA_CENTER_NAME
Model	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER.MODEL
Architecture	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER. ARCHITECTURE
Hypervisor type	Physical_Server.ANL_WORKING_SET_V*CFG_PHYSICAL_SERVER. HYPERVISOR_TYPE

Table 72. Physical servers attributes

Table 73. Virtual machine attributes

Attribute	Database path
Virtual machine ID	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.VMID
Virtual machine name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.VM_NAME
Operating system name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.OPERATING_SYSTEM_NAME
Operating system version	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.OPERATING_SYSTEM_VERSION
Application name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.MIDDLEWARE_NAME

Table 73. Virtual machine attributes (continued)

Attribute	Database path
Application version	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.MIDDLEWARE_VERSION
Physical server host name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK.HOST_NAME
Cluster name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK. SERVER_POOL_NAME
Data center name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK. DATA_CENTER_NAME
Physical servers model	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK.MODEL
Physical servers architecture	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK. HYPERVISOR_TYPE
CPU variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.CPU_VAR
Memory variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.MEM_VAR
Network variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.NETWORK_VAR
Disk variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.DISK_VAR

Note:

- Database paths are used to specify specific attributes from the database. A database path consists of table and view names, column names, and operators such as asterisks (*) and dots (.). An asterisk is used where tables and views have a one-to-many relationship. A dot is used where tables and views have a one-to-many relationship.
- Ensure that database paths do not include spaces or new line characters.

Writing rules for custom attributes

Use the following code in rules for custom attributes for virtual machines:

Use the following code in rules for custom attributes for physical servers:

```
<literal>
<path>Physical_Server.ANL_WORKING_SET_V*ANL_SERVER_TAG_MAP.SERVER_TAGS_PK.TAG_TYPE
</path><operator>EQ</operator>
<term>TAG_TYPE</term>
</literal>
<literal>
<path>Physical_Server.ANL_WORKING_SET_V*ANL_SERVER_TAG_MAP.SERVER_TAGS_PK.TAG_NAME
</path><operator>EQ</operator>
<term>TAG_VALUE</term>
</literal>
```

In each case, replace **TAG_TYPE** and **TAG_VALUE** with appropriate values, for example, to write a rule for virtual machines with production environment, then set **TAG_TYPE** to Environment and **TAG_VALUE** to Production. These values must exactly match the database values. Use AND as the grouping operator.

Importing rules in Capacity Planner

The XML snippets that are used to author policies must be kept in a single file in CSV format and loaded into the Planner Knowledge Base by clicking **Load Knowledge Data**. Any existing policies in the Knowledge Base are deleted when a new set is loaded.

Column name	Maximum column length	Details
RuleDimKey	8	Primary key of this rule
RuleName	50	Readable identifier of this rule
Priority	4	Rule Priority. 1 is the highest priority. All rules that cannot be broken must have priority 1.
Source	50	Metadata, for example, General, Customer, and BestPractice
IsActive	1	1 for yes, θ for no
RuleType	50	Type of rule, such as App Selection, OS selection, and so on
Notes	255	Additional information
RuleXML	Binary large object	XML format of rule

Table 74. CSV format details

A sample CSV file can be found in the installer, in the *dash_home* \installedDashboards\ com.ibm.tivoli.cpdash\AnalyticsDatabaseInstaller\samples directory.

Complete the following steps to import rules in Capacity Planner:

- 1. Open the Edit Current Environment window, as shown in Figure 92 on page 259.
- 2. Click the **Load Knowledge Data** icon is , which is in the upper-right corner of the window. The Refresh Knowledge Base Content window opens, as shown in Figure 130 on page 294:

1	Server Model Catalog		Browse
~	Rules	C:\samples\RuleDim.cs	Browse
	Benchmark		Browse
1	Custom Benchmark		Browse
77	Virtualization Overhead		Browse

Figure 130. Refresh Knowledge Base Content window

- **3**. Select the **Rules** check box. In the corresponding field, enter the path of the CSV file that contains the rules.
- 4. Click Upload.

Optimization goal

Currently, Capacity Planner supports *Minimize Systems* as an optimization goal. This goal consists of packing virtual machines together to minimize the number of physical servers that are used in the recommended environment.

Example: Optimization without rules

P1, P2, P3, and P4 are physical servers. V1 to V9 are virtual machines. Each physical server is marked as source, target, or both.

Source and Target

During optimization, virtual machines from this physical server can be moved to another physical server. This physical server can also receive virtual machines from another physical server.

Source

During optimization, virtual machines from this physical server can only be moved to another physical server.

Target During optimization, this physical server can only receive virtual machines from another physical server.

In this example, assume that the size of the physical server indicates the total CPU capacity, and for virtual machines CPU requirement is in descending order for V1 through V9. Place virtual machines from physical servers that marked as target on their corresponding physical servers, that is, V3 and V4 are placed on P2. During minimize systems optimization, we start by placing the virtual machine with the highest CPU requirement on the physical server with the highest CPU capacity and continue until the physical server does not have enough capacity. We then move to the next highest physical server.

In this example, four physical servers were consolidated to two physical servers, as shown in Figure 131.



Figure 131. Minimize systems optimization without rules

Example: Optimization with rules

Rules for optimization for Minimize Systems with rules:

- Run the optimization for each boundary rule
- Apply growth on virtual machine CPU usage that is based on utilization policies
- Try to pack workloads together that comply with colocation and anti-colocation polices

Example:

P1, P2, P3, and P4 are physical servers. V1 to V9 are virtual machines. Each physical server is marked as source, target, or both.

Source and Target

During optimization, virtual machines from this physical server can be moved to another physical server. This physical server can also receive virtual machines from another physical server.

Source

During optimization, virtual machines from this physical server can only be moved to another physical server.

Target During optimization, this physical server can only receive virtual machines from another physical server.

Policies that are applied in this example:

- Colocate virtual machines for Windows systems
- · Do not mix production virtual machines with non-production virtual machines
- Do not mix virtual machines for the Windows and Linux systems

In this example, all physical servers have the same total CPU capacity. For virtual machines the CPU requirement is in descending order for V1 through V9. Place virtual machines from physical servers that marked as target on their corresponding physical servers, that is, V4 is placed on P3.

During *Minimize Systems* optimization with policies, we start with production virtual machines. To minimize the system, V2 and V6 are placed with V4. From the remaining virtual machines, V5, V7, V9 are placed next on P1 because these virtual machines are from the colocation group. V1 and V3 are placed on same physical server because P1 still has enough capacity. V8 is not placed on P1 or P2 so as not to violate rules. Hence, V8 is placed on P3.

After optimization with rules, four physical servers were consolidated to three physical servers and placement ensures compliance with all policies, as shown in Figure 132.



Figure 132. Minimize systems optimization with rules

Optimization options

Keep existing recommendations

You can do incremental capacity planning.

Overwrite existing recommendations

You can ignore existing recommendations and optimize a selected set of physical servers.

Normalization benchmark setting

You can define settings in the analytics.properties file, which is in the following directories:

- For Tivoli Integrated Portal 2.1: *TIP_HOME*\tipv2\profiles\TIPProfile\installedApps\TIPCell\ isc.ear\AnalyticsWebUI.war\WEB-INF\classes
- For Dashboard Application Services Hub
 - On Windows systems: C:\Program Files\IBM\JazzSM\profile\installedApps\node-name\isc.ear\ AnalyticsWebUI.war\WEB-INF\classes
 - On operating systems other than Windows: /opt/IBM/JazzSM/profile/installedApps/node-name/ isc.ear/AnalyticsWebUI.war/WEB-INF/classes

Capacity normalization setting

Set BENCHMARK as SPEC or NONE.

SPEC Use the Standard Performance Evaluation Corporation (SPEC) benchmark for capacity

normalization. Standard Performance Evaluation Corporation benchmarks are published by the Standard Performance Evaluation Corporation organization. For details, see the Standard Performance Evaluation Corporation home page (http://www.spec.org).

NONE Use the raw capacity for capacity normalization. The raw capacity of the physical server is calculated based on the CPU speed and the number of cores.

Defining resource sizing by using application sizing rules

You can use the application sizing rules policy template to help define resource sizing that is based on application parameters. Use the sizing rules to adjust computed usage data for virtual machines that match given criteria. By using application sizing rules, you can override computed usage sizing data for resource parameters for virtual machines that match the rule criteria.

About this task

Use of the sizing rules has the following benefits:

- You can define policies to set resource allocation based on discovered attributes or custom tags with application parameters information.
- You can set target allocation values for the following metrics:
 - CPU Usage
 - Memory Usage
 - Storage Usage
 - Network Usage

Procedure

Complete these steps to define and use application sizing rules:

- 1. Identify existing attributes for virtual machines such as Middleware Name, or define custom attributes such as Transactions per second by clicking **Actions** > **Add Custom Tag** in the virtual machine inventory.
- 2. Set the values for these parameters in the Inventory view and click **Save**.
- 3. Define the rule (see samples below) and save the RuleDim.csv file.
- 4. Import the rules to Capacity Planner by using the **Load Knowledge Data** icon that is displayed in the upper-right corner of the **Edit VMware Current Environment** window.
- 5. Navigate to the Planning Center Edit Recommended Environment settings page and enable or disable the rule as required.
- 6. Complete the Generate plan step to generate recommendations. The enabled rules guide the sizing of the virtual machines that match the rule criteria in the recommendation generation.

Rule sizing

When you are using the application sizing rules, use the following Rule XML format:

```
<?xml version=1.0" encoding="UTF-8"?>
<rule>
<if>
<antecedent operator="AND">
<literal>
<path>Source.ANL_VM_PLACEMENT_SET_V<ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_TYPE</path>
<operator>EQ</operator>
<term>WEBSPHERE_USERS</term>
</literal>
<literal>
<path>Source.ANL_VM_PLACEMENT_SET_V<ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_NAME</path>
<operator>EQ</operator>
<term>1000</term>
```

</literal> </antecedent> </if> <!-- set demand in target intermediate target table --> <then> <dependent> <literal> <path>target.capacity</path> <operator>set</operator> <term>3000</term> </literal> <literal> <path>target.memory</path> <operator>set</operator> <term>8000</term> </literal> </dependent> </then> </rule>

"Defining resource sizing by using application sizing rules" on page 297 provides descriptions of the supported operators in the <operator> tag of the dependent literal: SET.

Supported values in <path> tag of dependent literal</path>	Description
target.capacity	Set CPU capacity (in MHz) required for a virtual machine
target.memory	Set memory capacity (in MB) required for a virtual machine
target.network	Set network bandwidth demand that is required for a virtual machine (in Mbps)
target.storage	Set storage I/O demand that is required for a virtual machine (in Mbps)

Table 75. Supported values in <path> tag of dependent literal

Table 76. Supported dynamic path in <term> tag of dependent literal

Supported dynamic path in <term> tag of dependent literal</term>	Description
CFG_VIRTUAL_MACHINE {any_column_in_cfg_virtual_machine}	Values from any of the numeric columns in the CFG_VIRTUAL_MACHINE table

A sample RuleDim.csv entry is provided here for the following example rule:

Sharepoint 2010 needs CPU reservation as 1000 MHz and network bandwidth demand as 1024 Kbps.

Custom tags: Virtual Machine Inventory view > Transactions per minute

```
Sharepoint 2010 needs CPU reservation as 1000 MHz and network bandwidth demand as 1024 Kbps MIDDLEWARE_VERSION
```

```
,"Sharepoint 2010 needs CPU reservation as 1000 MHz and network bandwidth demand as 1024 Kbps",1
```

,"General",0,"0","Application Sizing","No Notes",

"<?xml version="1.0" encoding="UTF-8"?>

<rule><if><antecedent operator="AND">

<liiteral><path>Source.ANL_VM_PLACEMENT_SET_V*CFG_VIRTUAL_MACHINE.MIDDLEWARE_NAME

</path><operator>EQ</operator><term>Sharepoint</term></literal>

<literal><path>Source.ANL_VM_PLACEMENT_SET_V*CFG_VIRTUAL_MACHINE.MIDDLEWARE_VERSION

</path><operator>EQ</operator><term>2010</term></literal></antecedent></if>

<then><dependent><literal><path>target.capacity</path><operator>set</operator>

<term>1000</term></literal><literal><path>target.network</path><operator>set</operator>

<term>1024</term></literal></dependent></then></rule>"

Detailed reference: Capacity efficiency and performance risk indices

The plan generated contains two indices to explain the quality of the plan.

The *capacity efficiency index* explains how efficiently we are using resources at the containers, such as hosts, clusters, and data centers, and overall across data centers. The measure is based on spare resources available at these multiple levels and the indices are provided at all levels. The spare is computed from usage analysis, not allocations.

- A virtual environment container level can have a negative spare, hence a negative capacity efficiency index, if resources are overcommitted at that level.
- A low positive capacity efficiency index means that the resources are used efficiently with not much wastage and low risk because no overcommittal occurs, according to observed usage.
- A high positive capacity efficiency index means that wastage occurs in the environment and efficiency is low.

These three states are color-coded as red (negative capacity efficiency index), green (low capacity efficiency index), and yellow (high capacity efficiency index).

The *performance risk index* is computed for each virtual machine. The index is computed differently for the current and the optimized environments.

In the current environment, the performance risk index of a virtual machine is based on the comparison of usage versus the reservation of the virtual machine. The performance risk index is color-coded as follows:

- Red if the virtual machine uses more than the reservation,
- Green if the virtual machine is using close to the reservation,
- Yellow if the virtual machine is using much less than the reservation

For the recommended environment, the performance risk index for a virtual machine is computed as the likelihood of a virtual machine violating the recommended reservation. The performance risk index is color-coded as follows:

- Red if the likelihood is high
- Green if the likelihood is low
- Yellow if the likelihood is moderate

•

The capacity efficiency index and performance risk index are computed for CPU and memory. The more conservative estimate determines the capacity efficiency index for the host, cluster or datacenter and the performance risk index for the virtual machine. All values are normalized to the total CPU and memory capacity. Red takes precedence over yellow, which in turn takes precedence over green. for example, if for a host, the capacity efficiency index for CPU is red and for memory is green, then the capacity efficiency index for the host is taken as red.

Some default settings for performance risk index calculation in the recommended environment report can be modified through the analytics.properties file, which is located in the *TIP_HOME/profiles/TIPProfile/installedApps/TIPCell/isc.ear/AnalyticsWebUI.war/WEB-INF/classes* directory, according to the following details.

Important: The performance risk index of virtual machines is calculated using GHz and GB values. After any change in this file, restart Dashboard Application Services Hub.

#RISK_SCORE_CALCULATION

This value enables risk scoring for the recommended environment report. Valid values are DEFAULT and UTILIZATION. In DEFAULT mode, the performance risk index is essentially disabled

and all virtual machines are shown as green for the recommended environment. UTILIZATION mode enables the calculation to estimate the performance risk index based on historical data.

Syntax: RISK_SCORE_CALCULATION=DEFAULT

#RISK_SCORE_GRANULARITY

This value gives the time scale granularity, for example, HOURLY, DAILY, or WEEKLY. The risk is calculated based on this value. The default value is HOURLY. This value is only used if RISK_SCRORE_CALCULATION=UTILIZATION.

Syntax: RISK SCORE GRANULARITY=HOURLY

#RISK_SCORE_PARAMETERS

This value denotes the resource types. You can choose from CPU, MEM, and ALL. The default value is ALL.

Syntax: RISK_SCORE_PARAMETERS=ALL

Uploading benchmark and server model catalog files in Capacity Planner

You can upload the latest benchmark and server model catalog files in Capacity Planner.

Procedure

- 1. Open the Edit Current Environment window.
- 2. Click the **Load Knowledge Data** icon *icon*, which is in the upper-right corner of the window. The Refresh Knowledge Base Content window opens.

~	Server Model Catalog	C:\Documents and Settin	Browse
	Rules		Browse
~	Benchmark	C:\Documents and Settil	Browse
	Custom Benchmark	[Browse
	Virtualization Overhead		Browse

Figure 133. Refresh Knowledge Base Content window for benchmark and server model catalog files

- 3. Select the **Server Model Catalog** check box, and in the corresponding field, enter the path by clicking **Browse** to point to the SERVER_MODEL_CATALOG.csv file.
- 4. Select the **Benchmark** check box, and in the corresponding field, enter the path by clicking **Browse** to point to the USER_DEFINED_BENCHMARK.csv file.
- 5. Click Upload.

Uploading custom benchmark file in Capacity Planner

You can upload the custom benchmark file in Capacity Planner.

Procedure

- 1. Open the Edit Current Environment window.
- 2. Click the **Load Knowledge Data** icon *icon*, which is in the upper-right corner of the window. The Refresh Knowledge Base Content window opens.

	Server Model Catalog		Browse
	Rules		Browse
	Benchmark		Browse
/	Custom Benchmark	C:\Documents and Settin	Browse
	Virtualization Overhead		Browse

Figure 134. Refresh Knowledge Base Content window for custom benchmark file

- 3. Select the **Custom Benchmark** check box, and in the corresponding field, enter the path by clicking **Browse** to point to the CUSTOM_USER_DEFINED_BENCHMARK.csv file.
- 4. Click Upload.

Use cases for network and storage aware capacity planning

The following are the network and storage aware capacity planning use cases:

- Host level network capacity check
- Host level storage access check
- Host level storage capacity check
- Localized placement for virtual machines that communicate frequently and are specified through the user input

Use case: Host level network capacity check

Table 77. Host level network capacity check

Category	Description
Description	In this use case, the combined network I/O demand of all virtual machines that is placed on a host is verified against the combined network capacity of all NICs on a host. In addition, it is ensured that the combined network capacity is not violated.

Table 77. Host level network capacity check (continued)

Category	Description
Output	When the plan is generated, the host level storage capacity checks are considered. In addition, the plan records any constraint check fails. Note: This use case must take into account the percentage of a network I/O demand of a virtual machine to another virtual machine that is collocated on the same host. This percentage must be decreased from the total network I/O demand of virtual machine since the network traffic does not go out of the host (unless VEPA mode is being used, which is rare). However, this requires VM to VM flow level information, and this use case assumes that such flow level information is not available.
How to enable or disable this check	 The host level network capacity check can be enabled or disabled using the flag "DISABLE_NETWORK_CAPACITY_CHECK" in the analytics.properties file. The location of the analytics.properties file: On Windows systems: DASH_HOME\profile\installedApps\Node_Name\isc.ear\ AnalyticsWebUI.war\ WEB-INF\classes On operating systems other than Windows: DASH_HOME/profile/installedApps/Node_Name/ isc.ear/AnalyticsWebUI.war/ WEB-INF/classes This check is disabled by default (DISABLE_NETWORK_CAPACITY_CHECK=TRUE). To enable the network capacity check, modify the following property in the analytics.properties file: DISABLE_NETWORK_CAPACITY_CHECK=FALSE Headroom for network capacity is specified through a property name (GLOBAL_NETWORK_HEADROOM_PERCENTAGE) in the analytics.properties file. User can specify any value 0 - 100 as a network IO headroom. The default headroom value is 10%. After you enable this check, you must restart the Dashboard Application Services Hub.

Use case: Host level storage access check

Table 78. Host level storage access check

Category	Description
Description	In this use case, the Capacity Planner must ensure that the virtual machine to data store connectivity or access exists, and the virtual machine must be placed only on those hosts that have access to data stores that are used by the virtual machines.
Output	The generation of plan ensures that a virtual machine is placed only on those hosts that have access to the data stores used by that virtual machine, and the plan logs the failures.

Table 78. Host level storage access check (continued)

Category	Description
How to enable or disable this check	The host level storage checks can be enabled or disabled by using the flag "DISABLE_STORAGE_CHECKS" in the analytics.properties file.
check	The location of the analytics.properties file:
	 On Windows systems: DASH_HOME\profile\installedApps\Node_Name\isc.ear\ AnalyticsWebUI.war\ WEB-INF\classes
	 On operating systems other than Windows: DASH_HOME/profile/installedApps/Node_Name/ isc.ear/AnalyticsWebUI.war/ WEB-INF/classes
	The check is disabled by default (DISABLE_STORAGE_CHECKS=TRUE).
	Note: You can enable this check if the CFG_HBA table contains positive values for the SPEED column in the capacity planner database.
	To enable the storage checks, modify the following property in the analytics.properties file:
	DISABLE_STORAGE_CHECKS=FALSE
	After you enable this check, you must restart the Dashboard Application Services Hub.

Use case: Host level storage capacity check

Table 79. Host level storage capacity check

Category	Description
Description	In this use case, the combined storage (disk) I/O demand of all virtual machines that are placed on a host is checked against the combined storage capacity of all HBAs or storage adapters on a host. In addition, it is ensured that the combined storage capacity is not violated.
Output	When the plan is generated, the host level storage capacity checks are considered. In addition, the plan records any constraint check fails.
How to enable or disable this check	The host level storage checks can be enabled or disabled by using the flag "DISABLE_STORAGE_CHECKS" in the analytics.properties file.
	 On Windows systems: DASH_HOME\profile\installedApps\Node_Name\isc.ear\ AnalyticsWebUI.war\ WEB-INF\classes
	 On operating systems other than Windows: DASH_HOME/profile/installedApps/Node_Name/ isc.ear/AnalyticsWebUI.war/ WEB-INF/classes
	This check is disabled by default (DISABLE_STORAGE_CHECKS=TRUE). To enable storage checks, modify the following property in the analytics.properties file:
	DISABLE_STORAGE_CHECKS=FALSE
	After you enable this check, you must restart the Dashboard Application Services Hub.
Assumption	The checks are completed against the overall disk IO demand of all virtual machines on the server and the entire spare disk IO capacity of the physical server (which is a sum of individual HBA capacities on the host) rather than against the individual host bus adapter (HBA) disk IO capacity.
	Headroom for HBA disk IO capacity is specified through a property name "GLOBAL_DISKIO_HEADROOM_PERCENTAGE" in the analytics.property file. User can specify any value 0 - 100 as disk IO headroom.
	The default headroom value is 10%.

Use case: Localized placement for a group of virtual machines where groups are specified through user input

Category	Description
Description	In this use case, a set of virtual machines that is specified by a user as group (through a user tag) are placed on the servers that belong to the same physical server cluster.
	For example, you can specify a group or a pair of virtual machines that communicate frequently and have high traffic rate between the virtual machines. The capacity planner tries to ensure that the virtual machines with high traffic rate between them are placed on the servers in the same cluster.
	Each virtual machine is associated with a VMCluster, and a target server is associated with a TargetCluster. The use case introduces the following clustering constraint on the placement scheme:
	All virtual machines that belong to the same VMCluster are placed on target servers that belong to the same TargetCluster. However, this constraint is soft: the constraint must be ensured only if the constraint does not result in violation of other constraints (for example, collocation and utilization), or the constraint does not result in some remaining virtual machine not placed. Note the localized placement of a set of virtual machines that communicates frequently can be ensured by associating the virtual machines in the set with the same VMCluster.
	The VM Clustering algorithm follows a first fit algorithm in the descending order of free CPU on target cluster. Target clusters with already used target servers are given higher priority, followed by target clusters with larger free CPUs to minimize number of systems.
How to enable or disable this check	To disable or enable VM Clustering, you can use the flag (DISABLE_VM_CLUSTERING_CONSTRAINT) in the analytics.properties file.
	The location of the analytics.properties file:
	 On Windows systems: DASH_HOME\profile\installedApps\Node_Name\isc.ear\ AnalyticsWebUI.war\ WEB-INF\classes
	 On operating systems other than Windows: DASH_HOME/profile/installedApps/ Node_Name/isc.ear/AnalyticsWebUI.war/ WEB-INF/classes
	The flag is set to TRUE (disable) by default. It must be changed to FALSE to enable VM Clustering.
	The VM CLUSTER ID is an arbitrary String. Identical Cluster ID strings are internally mapped to a unique number (starting from id value of 100), and this mapping is printed in the log. The log uses the mapped long value to identify the VM Clusters.
	After you enable this check, you must restart the Dashboard Application Services Hub.
Assumption	• The VM_CLUSTER_ID of a virtual machine, which specifies the VMCluster of the virtual machine, must be populated by the user. The VM CLUSTER ID can be any arbitrary String.
	• The current implementation assumes that the VisianPlacement is the algorithm that is used for Placement, and modified that algorithm.

Table 80. Localized placement for a group of virtual machines where groups are specified through user input

Table 80. Localized placement for a group of virtual machines where groups are specified through user input (continued)

Category	Description
Required input data	• A user specifies a tag for each virtual machine that is a part of a group that must be placed in the same physical server cluster. This is done through the capacity planner user interface. The tag is referred to as the VMCluster ID of a virtual machine. VMCluster of a virtual machine is specified in the table "ANL_SERVER_TAGS" where the TAG_TYPE column is "VM_CLUSTER_ID" and the TAG_NAME column specifies the value. The value must be specified by the user. The VM CLUSTER ID can be an arbitrary String. If the value is unspecified for VM_CLUSTER_ID, the VM_CLUSTER_ID is set to the default value of -1 , indicating that the virtual machine has no associated clustering constraint.
	• For a target server, the poolName of the target server (column SERVER_POOL_NAME in TargetView) is used as TargetClusterID of the target server. (TargetClusterID of a target server specifies its TargetCluster.)
Output	The generation of plan considers the user input on a group or a pair of virtual machines that must be placed together in the same physical cluster (for example, virtual machines with high network traffic between them), and the plan tries to allocate those virtual machines on the servers within the same cluster. If no use input is specified, the placement can be either within the same cluster or across different clusters. In addition, the plan logs any constraint check fails.
Failure or Violation event logging	Information or Debug logging in following cases:Clustering constraint is violated because two virtual machines in the same VM
	 Capsule have different VMClusterID. Unable to place a virtual machine in presence of the Clustering Constraint such that no target in the selected TargetCluster satisfies the colocation and capacity rules. (In this case, after raising a warning, placement is revised for the virtual machine without the clustering constraint.)

Cognos based report packages

Capacity Planner uses Tivoli Common Reporting.

Complete documentation for the Tivoli Common Reporting tool is located at the Tivoli Common Reporting Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.tivoli.tcr.doc_211/ic-home.html).

The Cognos based reports can be administered, run, and edited by Tivoli Common Reporting Version 2.1.1 software included with IBM Tivoli Monitoring 6.2.3.

For more information about Tivoli Common Reporting, see the developerWorks[®] page on Tivoli Common Reporting (http://www.ibm.com/developerworks/spaces/tcr).

Capacity Planner Reports for VMware description

All the Capacity Planner reports work on the data set that is selected in the Capacity Planner Edit Current Environment view. These reports can be started from the Capacity Planner web pages and are started in the context of the current data set.

Capacity Planner contains the following predefined Cognos based reports:

• Capacity Planner Physical Server Inventory

Shows configuration summary details for physical servers that are present in the working set.

• Capacity Planner Virtual Server Inventory

Shows virtual machine configuration details. The virtual machines that are shown are hosted on servers in the working set.

- Capacity Planner Utilization Detailed Timeseries Shows a detailed, unaggregated plot for various servers in a provided time interval.
- Capacity Planner Utilization Aggregated Timeseries Shows multiple time series plots for servers.
- Capacity Planner Current Environment Shows the resource demand numbers that are stored for each server present.
- Capacity Planner Optimized Environment Plan

Compares current resources that are used in the environment with optimized resources.

• Capacity Planner Additional Workloads Analysis

Shows the number of virtual machines that can be placed in the available CPU and Memory capacity (excluding the headroom).

- VMware Expense Reduction Report Identifies problems in the virtualized environment of the customer, suggests optimizations, and demonstrates Return on Investment savings on using Smart Cloud Monitoring.
- Capacity Planner Optimized Environment Plan Cluster Details

Shows number of servers to be retired from cluster, number of VMs moved in and out of the cluster, VMs with no change in the Cluster and number of VMs not placed with reason.

Report structure	Description
Name	Capacity Planner Physical Server Inventory
Description	This report shows configuration summary details for physical servers that are present in the working set. The values that are displayed are Number of Physical servers, Number of Physical CPUs, CPU Speed (GHz), and Memory Installed (GB). The distribution of servers is displayed for different data centers, clusters, and hypervisors.
Purpose	Show configuration summary details for physical servers in the working set.
Output	 The report displays bar charts of the distribution of physical servers. The first chart displays the distribution for different clusters and data centers. The second chart displays the distribution for different hypervisors.
1	

Table 81. Capacity Planner Physical Server Inventory report

Table 82. Capacity Planner Virtual Server Inventory report

Report structure	Description
Name	Capacity Planner Virtual Server Inventory
Description	This report shows virtual machine configuration details. The virtual machines that are shown are hosted on servers in the working set. The values that are displayed are Number of Virtual Machines, Number of Virtual CPUs, Current CPU Reservation (GHz), and Current Memory Reservation (GB). The distribution of virtual machines is displayed for different data centers, clusters, operating systems, and applications.
Purpose	Show configuration summary details for virtual machines that are hosted on servers in the working set.
Output	The report displays bar charts of the distribution of virtual machines.
	• The first chart displays the distribution for different clusters and data centers.
	• The second chart displays the distribution for different operating systems.
	• The third chart displays the distribution for different applications.

Report structure	Description		
Name	Capacity Planner Utilization Detailed Timeseries		
Description	This report shows a detailed, unaggregated plot for various servers in a provided time interval. Data points come directly from aggregated measurement tables in utilization schema. Utilization is displayed for CPU, memory, network, and disk.		
Purpose	Shows utilization for various servers in a provided time interval.		
Parameters	Aggregation You can choose minimum, maximum, or average aggregation. Summarization You can choose the granularity of how the data are generated. You can choose Hourly, Daily, Weekly, Monthly, or Yearly. Data Center You can choose one data center, or all. Cluster You can choose one cluster, or all, among clusters that belong to the selected data centers. Host Server You can choose one server, or all, among servers that belong to the selected clusters. Top/Bottom VMs You can choose whether to view the top used or bottom used virtual machines. Value of N You can enter the number of virtual machines or what percentage of virtual machines to view at the top or bottom.		
Output	The report displays separate line graphs for CPU, memory, network, and disk utilization.		

Table 83. Capacity Planner Utilization Detailed Timeseries report

Table 84.	Capacit	v Planner	Utilization	Aggregated	Timeseries	report
10010 01.	oupuon	y i iaiiii0i	ounzation	, iggi ogulou	111100001100	ropon

Report structure	Description		
Name	Capacity Planner Utilization Aggregated Timeseries		
Description	This report shows multiple time series plots for servers. The report compares aggregated utilization values on the y axis for various virtual machines that are selected or hosted on servers present in the working set. Utilization is displayed for CPU, memory, network, and storage.		
Purpose	Show utilization for various servers in a provided time interval.		
Parameters	 Aggregation You can choose minimum, maximum, or average aggregation. Summarization You can choose the granularity of how the data are generated. You can choose Hourly, Daily, Weekly, Monthly, or Yearly. Data Center You can choose one data center, or all. Cluster You can choose one cluster, or all, among clusters that belong to the selected data centers. Host Servers You can choose one server, or all, among servers that belong to the selected clusters. Top/Bottom VMs You can choose whether to view the top used or bottom used virtual machines. Value of N You can enter the number of virtual machines or what percentage of virtual machines to view at the top or bottom. 		
Output	The report displays separate line graphs for CPU, memory, network, and storage utilization.		

Table 85. Capacity Planner Current Environment report

Report structure	Description
Name	Capacity Planner Current Environment
Description	This report shows the resource demand numbers that are stored for each server present. CPU and memory details are displayed for selected data centers, clusters, and host servers. Recommended reservation is displayed with current reservation of virtual machine CPU and memory details. Expected network usage and storage are displayed. Threshold values of CPU, memory, network, and disk are given.
Purpose	Show resource demand numbers for each server. Compare recommended reservation with current reservation. Show expected usage and storage, and threshold values.
Parameters	Data Center You can choose one data center, or all. Cluster You can choose one cluster, or all, among clusters that belong to the selected data centers. Host Servers You can choose one server, or all, among servers that belong to the selected clusters.
Output	The report displays two bar charts for each selected data center, cluster, and host server. The first bar chart displays CPU details. The second bar chart displays memory details. Both bar charts display capacity, recommended reservation, buffer, and space. These bar charts are followed by two bar charts that compare current reservation with recommended reservation for each individual virtual machine in the selected data center, cluster, and host server. The first of these bar charts display data for CPU. The second bar chart displays data for memory. The final two bar charts display expected usage for the individual virtual machines. The first bar chart displays expected network usage. The second bar chart displays expected storage usage. At the end of the report, a table displays Thresholds Used for CPU, Memory, Network, and Disk. In each of these four categories, thresholds are displayed for Utilization (High (Madium (Law)) and Verinese (Stable (Unstable))

Table 86. Capacity Planner Optimized Environment Plan report

Report structure	Description
Name	Capacity Planner Optimized Environment Plan
Description	This report compares current resources that are used in the environment with optimized resources. Current, and recommended, values are displayed for total capacity, reservation, and spare CPU (GHZ) and memory (GB). Expected utilization is displayed before and after optimization.
Purpose	Compares recommended values with current values for capacity, reservation, and spare CPU and memory. Show expected utilization before and after optimization.

Table 86. Capacity Planner Optimized Environment Plan report (continued)

Report structure	Description
Output	The report displays a table with the current, and recommended, number of physical servers and virtual machines. For both CPU and memory, the table also displays the current, and recommended, values of total capacity, reservation, and spare, which is the difference between capacity and reservation.
	This table is followed by summary of cluster level actionable recommendations, which shows the following statistics:
	Current number of servers and VMs
	Current CPU utilization (in GHz) and Memory utilization (in GB)
	Recommended number of servers and VMs
	• Recommended CPU utilization (in GHz) and Memory utilization (in GB)
	Number of VMs moved out and moved in
	• Number of servers recommended for retirement and number of VMs not placed
	In this table, when you click the Cluster name, the "Capacity Planner Optimized Environment Plan Cluster Details" report opens, which displays details of the actionable recommendations for the selected cluster.
	This table is followed by two bar charts - one for CPU and one for memory. Each bar chart displays Expected Utilization before Optimization, Expected Utilization after Optimization, Capacity, Buffer, and Spare.
	The bar charts are followed by a table that contains details of individual virtual machines: Name, Number of vCPUs, Current Reservation CPU, Current Reservation Memory, Recommended Reservation CPU, Recommended Reservation Memory, Operating System, and Application.
	This sequence of charts and tables is repeated for each physical server specified.
	This report also displays the following actionable recommendation scenarios:
	• Virtual Machines that should be moved from one cluster to another
	• Virtual Machines that could not be placed and associated reason(s)
	Data Centers that have significant change in Capacity Efficiency Index (CEI)
	Clusters that have significant change in Capacity Efficiency Index (CEI)
	Physical Servers that have significant change in Capacity Efficiency Index (CEI)
	• Virtual Machines that have significant change in Performance Risk Index (PRI)
	Summary usage of fictitious servers

Table 87	Capacity	Planner	Additional	Workloads	Analysis	report
----------	----------	---------	------------	-----------	----------	--------

Report structure	Description
Name	Capacity Planner Additional Workloads Analysis
Description	This report shows the number of virtual machines that can be placed in the available CPU and Memory capacity (excluding the headroom). In addition, this report shows the distribution of virtual machines. Based on this distribution of virtual machines, the report takes the values that a user entered to show how many virtual machines can be placed in the available capacity.
Purpose	This report helps capacity planner to complete what-if analysis on the additional workloads that are based on Average, Minimum, Maximum, or Custom virtual machine profile. The Average profile is the default profile. A profile of the virtual machine is computed based on the resource recommendations that are displayed in the Capacity Planner Optimized Environment Plan report.

Table 87. Capacity Planner Additional Workloads Analysis report (continued)

Report structure	Description
Parameter	 VM Profile You can select a virtual machine profile for which you want to view the report. The following are the virtual machine profiles: Average Minimum Maximum Custom
Output	 The report displays the following tables and graphs for the additional workloads: The first table displays information about Recommended CPU (in GHz) and Recommended Memory (in GB). The second table displays the values for the following parameters: Physical Server Name Total CPU GHz Spare CPU GHz (Excluding Overhead) Total Memory GB Spare Memory GB (Excluding Overhead) Constraint No. of Additional Workloads The first graph displays the virtual machine distribution for CPU usage (in GHz).

Table 88. VMware Expense Reduction report

Report Structure	Description	
Name	VMware Expense Reduction Report	
Description	This report helps the IBM technical sales team to explain Return on Investment (ROI) for the Smart Cloud Monitoring offering that is based on the historical data and capacity analytics for the VMware agent. In addition, this report can be used by customers to explain Return on Investment savings.	
Purpose	The VMware Expense Reduction report:	
	• identifies problems in the virtualized environment of the customer	
	suggests optimizations	
	• demonstrates Return on Investment savings because of the usage of Smart Cloud Monitoring	
Parameter	Date Range Report Period Select the report period from a predefined date range, such as Last Week, Current Month, Last 30 Days. You can also enter a start date, an end date, and the time for the reporting period.	
	Start Date Select a start date from the calendar, and the start time from the time widget. You must select both date and time.	
	End Date Select an end date from the calendar and an end time from the time widget. You must select both date and time.	
	Resources	
	Data Center(s) Select the data center that you want to evaluate from the environment. You can select multiple data centers.	
	Select the cluster that you want to evaluate within the selected data center from the environment. You can also select the multiple clusters for one data center.	

Table 88. VMware Expense Reduction report (continued)

Report Structure	Description
Parameter	User Inputs for Analysis
	Shift Periods If the shifts are enabled, the Hourly table displays the shift period as 1 or 2 based on the peak and off-peak hours that are configured in the data warehouse. The Daily table consists of 1 and 2 corresponding to the peak and off-peak hours, and -1 corresponding to the summarized value for that day. If the shifts are not enabled, then the default value is -1.
	Vacation Periods Enter 0 for work days or 1 for the vacation days. If the vacation period is not enabled, the default value is -1.
	Number of Clusters Type the number of clusters. This is the number of clusters across the enterprise.
	Number of vSphere Servers Type the number of vSphere servers. This is the number of vSphere servers across the enterprise.
	Number of Virtual Machines Type the number of virtual machines for which you want the report to be displayed. This is the number of virtual machines across the enterprise.
	Number of Data Stores Type the number of data stores for which you want the report to be displayed. This is the number of data stores across the enterprise.
	Memory (GB) Type the value for the memory. This is the memory capacity that is measured in GB across the enterprise.
	Number of CPU cores Type the number of CPU cores. This is the number of CPU cores across the enterprise.
	CPU (GHz) Type the value for CPU. This is the capacity that is measured in GHz across the enterprise.
	DASD (TB) Type the value for the Direct Access Storage Device. This is the storage capacity that is measured in Terabyte (TB) across the enterprise.
	Currency Select the currency. This is the local currency measure. The default currency is US Dollars.
	Hardware cost per server Type a value for the cost per server. This is the cost of each server that comprises the VMware environment.
	Administrative cost per server Type a value for the cost per server. This is the average cost of administering a vSphere Server. This cost is based on the location of the data center.

Table 88. VMware Expense Reduction report (continued)

Report Structure	Description
Parameter	Virtualization Licensing per CPU socket Type a value for virtualization licensing per CPU socket. This is an average cost for the CPU sockets per server in the environment.
	Energy cost per server per year Type a value for the energy cost per server per year. This is an annual cost for power for an average server in the VMware environment.
	Flooring space per server Type a value for the flooring space that is required per server. This is the average cost for floor space for each of the vSphere servers.
	Storage cost per TB Type a value for storage cost per server. This is the average cost to purchase and maintain a Terabyte of storage for the VMware environment.
	Smart Cloud Monitoring Software Cost per VM Type the cost of IBM SmartCloud [®] Monitoring per virtual machine. Consider the following points when you provide this information:
	 What is the software price as per the currency in your country? The list price for IBM SmartCloud Monitoring in US dollars is \$225/Virtual Machine.
	What discount will you offer your customer?
	• If the customer is already paying for VMware monitoring, you might consider entering zero for this value to represent the fact that there is no difference in the price the customer is currently paying for management software.
Output	This report displays the following sections on a page:
	Evaluated Environment
	This section displays charts that are based on the key performance indicators, such as CPU, Memory, and Storage for Clusters and Virtual Machines. In addition, the Evaluated Environment identifies potential problems in the environment of the customer.
	Optimization
	This section suggests optimizations that are based on the parameters that are selected in the Capacity Planning tool from step 1 through step 5.
	• Savings
	This section displays the Return on Investment (ROI) savings that are based on various cost factors, such as Administration, Licensing, Flooring, Storage, Energy, Hardware, and Smart Cloud Monitoring software.
	Note: The VMware Expense Reduction Report provides an estimated projected savings if the VMware environment is optimized based on the recommendations in the report. This report depends on input data from the customer to identify their costs for power, cooling, and licenses.

Table 89. Capacity Planner Optimized Environment Plan Cluster Details report

Report Structure	Description
Name	Capacity Planner Optimized Environment Plan Cluster Details
Description	This report shows details of actionable recommendations for the selected cluster.
Purpose	Shows number of servers to be retired from cluster, number of VMs moved in and out of the cluster, VMs with no change in the Cluster and number of VMs not placed with reason.
Parameters	Data Center You can choose one data center. Cluster You can choose one cluster from the clusters that belong to the selected data center.
Table 89. Capacity Planner Optimized Environment Plan Cluster Details report (continued)

Report Structure	Description
Output	This report displays the following actionable recommendation scenarios for the selected cluster:
	• The "Servers recommended for retirement" chart displays number of servers to be retired from the selected cluster.
	• The "Virtual Machines that should be moved out of the Cluster" chart displays number of virtual machines that must be moved out of the selected cluster.
	• The "Virtual Machines that should be moved into the cluster" chart displays number of virtual machines that must be moved into the selected cluster.
	• The "Virtual Machines with no change in the Cluster and have recommended configuration change" chart displays number of virtual machines with no change in the selected cluster, and have recommended configuration change for CPU or Memory demand.
	• In the "Virtual Machines that were not placed" chart, the reason "No utilization data" indicates that no recommended CPU and memory reservation exist for the associated virtual machine and for all other cases the reason is "Not determined".
	• The links at the end of this report display the following IBM Tivoli Monitoring for Virtual Environments Reports:
	- VMware VI Top N VMs By Resource Pool
	– VMware VI Bottom N VMs by Resource Pool

314 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 12. Capacity Planner for PowerVM

The capacity planner for PowerVM works on sizing plans for the virtual environment.

PowerVM PlanningCenter wizard

The PlanningCenter provides a simple five-step wizard to generate a sizing recommendation plan for the Power systems virtual environment.

Important: Capacity Planner does not support multiple concurrent users. A system acquires a lock for the user that prevents other users from using the hypervisor-specific Capacity Planner for working on their planning scenarios.

Planning Center for PowerVM ×	
Planning Center for PowerVM	
You have acquired the Planning Center for PowerVM lock at 25 Jul 2013 06:46:28 AM EDT	You need to explicitly release the lock for othe
to use the Planning Center for PowerVM. Release Lock	
In the second	ommendation.
Step 1: Snapshot config data.	
Load the latest configuration data for managed systems and logical partitions for analysis. You c	an change this data for what-if analysis.
Advanced options: Select the data load options:	
Load data for selected Managed Systems	
Clean database before loading	
Load Config	
Step 2: Set analysis time period.	
Set the time period for which the measurement data in the warehouse can be analyzed correspondence of the standard from the warehouse	nding to the logical partitions loaded in Step 1.
Sof Time	
Set time	
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an elemanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to the stope of the stope	opert mode page where you can select the sub o add new attributes or clean the data for mana
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope	xpert mode page where you can select the sub o add new attributes or clean the data for mana
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics.	xpert mode page where you can select the sub o add new attributes or clean the data for mana
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization).	xpert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you	xpert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data.	xpert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs	xpert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report	o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan.	o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan. Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the plan in step 5 to view the latest recommended size of the	opert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan.	o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan. Image: Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommended sizing plan based on recommended environment settings. Advanced options: Custom settings can be made on the Edit Recommended Environment Settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings and based on the settings page where explored settings page where explored settings page where explored settings page where explored sett	opert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte ommendation.
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan. Image: Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommended sizing plan based on recommended environment settings. Advanced options: Custom settings can be made on the Edit Recommended Environment Settings page where expression of the plan	opert mode page where you can select the sub o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte ommendation.
Step 3: Scope the infrastructure for analysis. Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an emanaged systems that you want to be part of this analysis. Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page t systems and logical partitions, if required. Define Scope Step 4: Analyze logical partition characteristics. Analyze the warehouse data within the time limit set in Step 2, to compute the LPAR-level sizing of weekly utilization). Advanced options: Experts can customize sizing on the Edit Current Environment page. While on this edit page, you logical partitions based on measurement data. Analyze LPARs Current Environment Report Step 5: Generate recommended sizing plan. @ Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommended sizing plan based on recommended environment settings. Advanced options: Custom settings can be made on the Edit Recommended Environment Settings page where explored and the state plan Report	o add new attributes or clean the data for mana characteristics using default settings (maximur can trigger several custom actions to characte ommendation.

1. In **Step 1: Snapshot config data**, click **Load Config** to load the latest configuration data for managed systems and logical partitions. For more information on loading the configuration data to the capacity planner database, see "Configuring data for managed systems" on page 317.

- 2. In **Step 2: Set analysis time period**, click **Set Time** to set the time period for which the measurement data in the warehouse is analyzed corresponding to the logical partitions that are loaded in step 1.
- 3. In **Step 3: Scope the infrastucture for analysis**, click **Define Scope** to select a subset of managed systems for analysis. You can edit the current configuration, add new attributes, and clean the data for managed systems in the Edit Current Environment window here: "Detailed reference: Edit Current Environment window" on page 333.
- 4. In **Step 4: Analyze logical partition characteristics:**, click **Analyze LPARs** to invoke Compute Usage to analyze the warehouse data within the time limit that you set in step 2. You can customize and edit the sizing estimates in the Edit Current Environment window. For more information on sizing LPARs, see "Scenario: Providing a sizing recommendation of LPARs on systems" on page 319.
- 5. In **Step 5: Generate recommended sizing plan**, click **Generate Plan** to generate a sizing plan that is based on recommended environment settings. This opens the PowerVM LPAR Sizing window where you complete the sizing steps to complete your sizing. To view the sizing report after you complete the sizing step, you click the **Recommended Sizing Plan Report**. You can also edit the sizing rules on the Edit Recommended Environment Settings page.

Capacity Planner has a multi-user access conflict removal feature. When a user starts working on a particular Capacity Planner, the system acquires a persistent lock on the Capacity Planner for this user automatically. No other user can work on the Capacity Planner while this user has the lock. The user retains this lock, even across sessions, until the lock is released by the user or until the Dashboard Application Services Hub administrator releases the lock. The administrator (with the iscadmin role) can release the lock for any user.

Configuring data for managed systems

You can load the configuration data for physical systems and LPARS into the Capacity Planner Database for analysis purposes.

Before you begin

To view the list of managed systems that are available, you must configure IBM Tivoli Monitoring to collect data from the HMC that manages these systems. You must ensure that the correct configuration is defined for your IBM Tivoli Monitoring Data Provider in the Connections page. For more information, see "PowerVM: Installing the Capacity Planner" on page 16.

Procedure

1. Log in and open the **Planning Center for PowerVM** wizard.



Figure 136. Planning Center for PowerVM

- 2. In step 1 of the **Planning Center for PowerVM** wizard, select one of the following options:
 - Select the Load data for selected Managed Systems check box to load or update configuration information for selected managed systems from IBM Tivoli Monitoring. The configuration

information of the Managed Systems and LPAR, for example, the host name or the entitlement information is overwritten. The tag information and LPAR utilization information is not deleted or overwritten.

- Select the **Clean database before loading** check box to clear the Capacity planner database, and then load fresh configuration information from Tivoli Monitoring. All the previously stored information about Managed Systems and LPARs is deleted. The deleted information includes tag information, working set scope, source and target tags of managed systems, and computed or edited usage for LPARs.
- 3. Next, click Load Config.

Load Config updates all the information for the Managed Systems or LPARs that was loaded into the Capacity Planner database previously, except for the following items:

- Base Configuration
- Scope selected for analysis
- User added Custom Tags
- Computed or edited usage

The Selective data load page opens that shows the managed systems that are available. You filter the managed systems that are available by using the filter input box.

otal Managed System : 49 wailable Managed System			Selected Manage	d System
SN064B	×		Filter	×
Server-8203-E4A-SN064B	ED2	Add >	Server-8203-E4A	-SN064BEC2
Server-8203-E4A-SN064B	EE2	Add All >		
Server-8203-E4A-SN064BEF2 Server-8203-E4A-SN064BF02		< Remove		
	f and a	< Remove All		

Figure 137. Selective data load window

- 4. Select one or more available managed systems and click either Add> or Add All>.
- 5. Then, click **OK**. The configuration data and the LPARs for the selected managed system are loaded to the **Capacity Planner** database.
- 6. To trigger the loader to load the configuration data and LPARs for the selected managed systems, click **OK**. When the configuration data and LPARs for the selected systems are loaded, the following message is shown:



Figure 138. Success message

- 7. Click OK.
- 8. To view or to edit the managed systems that are loaded to the **Capacity Planner** database, in step 3 of the **Planning Center for PowerVM** wizard, click **Edit Current Environment**.

Planning Center for PowerVM ×
Planning Center for PowerVM
 Any change in steps 1 to 4 needs a re-generation of the plan in step 5 to view the latest recommendation.
Step 1: Snapshot config data.
Load the latest configuration data for managed systems and logical partitions for analysis. You can change this data for what-if analysis.
Advanced options: Select the data load options:
Load data for selected Managed Systems Clean database before loading
Load Config
Step 2: Set analysis time period.
Set the time period for which the measurement data in the warehouse can be analyzed corresponding to the logical partitions loaded in Step 1. The measurement data is federated from the warehouse. Set Time
Step 3: Scope the infrastructure for analysis.
Default scope includes all managed systems loaded in Step 1. Click Define Scope to go to an expert mode page where you can select the subset of managed systems that you want to be part of this analysis.
Advanced options: While scoping you can also edit the current configuration on the Edit Current Environment page to add new attributes or clean the data for managed systems and logical partitions, if required.
Define Scope

Figure 139. Planning Center for PowerVM

Scenario: Providing a sizing recommendation of LPARs on systems

An IT administrator wants to ensure that the existing infrastructure is optimally used. Also, they want to plan ahead for any contingencies that arise in the current environment.

About this task

An administrator or an analyst can use this scenario.

- An IT Administrator wants to evaluate the spare capacity that is available on selected existing Power Systems[™].
 - Check that the systems that run the workloads have sufficient resources that are based on the historical utilization data for LPARs.
 - Complete the sizing by keeping all LPARS on the same system.
 - Ensure that the work loads are not consolidated across LPARs.

Ultimately, the IT administrator wants to provide a sizing recommendation of LPARs by keeping the workloads on the same system.

- An IT administrator wants to transfer the workloads on a new target power system.
 - The IT Administrator knows that the existing systems that run workloads are not sufficient and wants to know the required LPAR sizes on a new target system given the historical utilization.
 - Size the specific workloads with recommendations for moving the LPARs to potentially different target systems.
 - Ensure that the workloads are not consolidated across LPARs. This means that a specific workload
 must run in the existing LPAR. The number of LPARs remain the same but the configuration might
 vary on the target environment.

Ultimately, the IT Administrator wants to provide a sizing recommendation of LPARs for a different target system that is automatically taken from the sizing catalog.

Procedure

Edit PowerVM Current Environment

Selecting a physical system to size

1. In the Edit PowerVM Current Environment window, select one or more systems to size.

Views >> Invent Views •	ory >> Mana	ged Systems - Reports	- 💀 🔯	C G								tivoli	× *
Select Managed System	Status	Machine Serial Number	HMC Host Name	Managed System Name*	Data Source	Total Processors*	CPU Cores Activated*	CPU Speed [MHz]*	Total Memory [MB]*	Model	Model Name		
V	V	064664H	Tivoli-HMC	Server-9133-55A- SN064664H	Discovery	8	8	1,648	16,384	9133-55A	p5-550Q-9133		
V	Ø	100F84P	Tivoli-HMC	Server-8233-E8B- SN100F84P	Discovery	32	32	3,300	131,072	8233-E8B	750-8233-E8B		

Figure 140. Edit PowerVM Current[®] Environment window

All the selected systems and the LPARs that are associated with them are included in the sizing. However, if you select an LPAR and no information is available, the LPAR data is not included in the sizing.

- 2. If all the relevant information is available to uniquely identify a model, the status column shows a green indicator. The model name column is also populated. If the status column shows a red indicator, you must try to enter the values manually. For example, you enter the processing speed of the discovered system.
- 3. Click Views > Logical Partition Utilization.
- 4. This view shows the utilization for your LPARs. To provide the utilizations manually, complete the next step.
- 5. In the Actions list, select Edit Usage.

Edit PowerVM Current Environment

Views >> Logical Partition Utilization

	Compute U	sage	matrices 1 wantan							Pritter X	4
3	Logical P Generate W HostNam Edit Usage	orkload Stability Type	Managed System Name	CPU Usage [Share]	Memory Usage [MB]	Read Operations	Write Operations	Read Usage[KB]	Write Usage[KB]	Network Operations	ľ
3	lcolumbus.in.ibm.com	Icolumbus	Server-9133-55A- SN064664H	70	4,063						
	smecaix05	smec05	Server-8233-E8B- SN100F84P	16.4	6,749	1.0	187.0	862.0	4280.0	16.0	0
0	smecaix06.in.ibm.com	smec06	Server-8233-E8B- SN100F84P	84.8	8,027	1.0	679.0	36757.0	3162.71	20.0	0
	smecaix07.in.ibm.com	smec07	Server-8233-E8B- SN100F84P	14	6,291	37.0	746.0	5835.78	2878.66	12.0	13
0	smecaix08.in.ibm.com	smec08	Server-8233-E8B- SN100F84P	13.2	6,859	1.0	894.0	538.0	2402.11	12.0	0
	smecaix11.in.ibm.com	smecaix11	Server-8233-E8B- SN100F84P	129.7	2,035	159.0	350.0	7553.96	3427.86	18.0	0
	ljackson.in.ibm.com	ljackson	Server-9133-55A- SN064664H	70	4,059	1464.19	963.18	48.05	71.33	17481.0	2
0	smecaix02	smecaix02-vio-server	Server-8233-E8B- SN100F84P	180.3	2,011						
	NewLPAR1.in.ibm.com	NewLPAR1	Server-8233-E8B- SN100F84P	70	1,434						
10	of 10 items		5	10 25 50 100 All						H + 1 ►	31.5

Save Cancel

Figure 141. Actions menu: Edit PowerVM Current Environment window

- 6. In step 4 of the **Planning Center for PowerVM** wizard, click **Analyze LPARs**. This step computes the usage characteristics for the LPAR utilization values. The default setting is to take the maximum weekly utilization.
- 7. Then, in step 5 of the Planning Center for PowerVM wizard, click Generate Plan.

Selected workload

Workload definition Selected workload Selected system The target structure of your estimation is listed below. Click on a workload link to see the information about the workload. Note: You can also use the navigation tabs above, to move to a particular screen. SCM Power System Solution 🔤 🔄 Smart Cloud Monitoring IBM Power System System: Server-9133-55A-SN064664H (1 interval defined) 🗧 🔄 Partition: Icolumbus 🛄 🖌 Icolumbus Bartition: Ijackson 🛄 🖌 ljackson - 🔄 Partition: NewLPAR2 MewLPAR2 System: Server-8233-E8B-SN100F84P (1 interval defined) 🗝 🛅 Partition: smec05 🛄 🖌 smec05 Partition: smec06 🚞 ✔ <u>smec06</u> 🔄 🔄 Partition: smec07 🛄 🖌 smec07 Partition: smec08 🛄 🖌 <u>smec08</u> Partition: smecaix11 🛄 🗸 smecaix11 Partition: NewLPAR1 E V NewLPAR1 Partition: smecaix02-vio-server i smecaix02-vio-server



Figure 142. Selected workload page

For more information about the **Selected workload** page, see "PowerVM LPAR Sizing wizard: Selected workload page" on page 326.

8. To view a workload that is associated with a partition, click the **Workload** link that relates to the partition. For more information about the **Workload definition** page, see "PowerVM LPAR Sizing wizard: Workload definition page" on page 328.

Icolumbus

SCM Workload Definition

Selected workload	Workload definition Selec	ted system	10		
er: Smart Cloud Mor	nitoring System: Server-9133-55	A-SN064664	H Icolumbus /	AIX - 6.1, LPA	R Dedicated F
<u>Model</u> : p5-550Q-9133 <u>Feature</u> : 8285 <u>Clock Speed</u> : 1650 M⊦	łz				
1. Total CPU Utilization	i.				70
2. Processor cores act	ivated				8
3. Assigned Processo	r Cores				2
4. Memory					4063
5. Disk Configuration					
Group Name	Storage Used(GB) Consumed	Read Ops	Read IOSize (bytes)	Write Ops	Write IOSize (bytes)
	54	0	0.0	0	0.0

Figure 143. Workload definition page

9. To select a system for sizing, complete one of the following steps:

- Select the **Selected system** page. For more information about the **Selected system** page, see "PowerVM LPAR Sizing wizard: Selected system page" on page 328.
- On the **Workload definition** page, click **Continue**.

Selected system		
Selected workload Workload c	lefinition Selected system	
System Information		
Tier	System	Immediate Solution
Smart Cloud Monitoring IBM Power System	Server-9133-55A-SN064664H	Choose base for Server-9133-65A-SN064664H
	Server-8233-E8B-SN100F84P	Choose base for Server-8233-E8B-SN100F84P
Back		

Figure 144. Selected system page

10. To choose an option for the selected system, click the **base for <Server>** link.

PowerVM LPAR Sizing

Selected system

Choose B	ase System
Make one	selection below, for your existing system. The sizing engine will recommend a solution based on your selection.
Select the	existing system option to give preference to this base model.
For consid	deration of typical migration paths, select the migration path option.
The final s	selection in the list below allows the system selection algorithms to choose the best system without consideration of any migration or upgrade paths.
0	Size to Existing System Model p5-5500-9133 8285 (1650 MHz, 4-8 cores, 20.25-38.34 rPerfs)
œ	Choose a new model based on typical migration paths if possible (default selection)
0	Allow the system selection algorithms to choose the best system without giving any preference based on existing system migration or upgrade paths
S Ba	ck 🕑 Continue

Figure 145. Selected system

11. Select a model or an algorithm to size your existing system and click **Continue**. You must repeat this step for all selected systems.

Selected workload Workload defin	inition Selected system	
ystem Information		
Tier	System	Immediate Solution
Smart Cloud Monitoring IBM Power System	Server-9133-55A-SN064664H	720-8202-E4C 6 cores, 3000 MHz 27% utilized ♥ Select for planning ♥ Change base for Server-9133-55A-SN064664H
	Server-8233-E8B-SN100F84P	750-8233-E8B 8 of 8-32 cores, 3300 MHz 48% utilized IV Select for planning Ochange base for Server-8233-E8B-SN100F84P

Figure 146. Selected system tab

- 12. To select the systems for planning and to return to the Planning Center wizard, click **Complete Sizing**.
- 13. To view the sizing details, click the **System** link.

Immediate Solution



Figure 147. PowerVM LPAR Sizing

- 14. To select a different model for sizing purposes, open the Model/Feature list and select a model.
- 15. To compare the model that is recommended by the sizing engine with another model, click the (+) icon.

PowerVM LPAR Sizing wizard: Selected workload page

You use the PowerVM LPAR Sizing wizard: Selected workload page to view the LPARs on a particular system and the workloads that are associated with them.

The **Selected workload** page shows a tree structure for your solution.

Selected workload

Selected workload Workload definition Selected system

The target structure of your estimation is listed below. Click on a workload link to see the information about the workload.

Note: You can also use the navigation tabs above, to move to a particular screen.



Continue

Figure 148. Selected workload page

Each workload is preceded by a green check mark icon. These check mark icons show that the data selected for sizing was validated correctly against the sizing engine.

The tree structure shows the Default tier that contains the selected systems, the partitions or LPARs that are on the system, and the workload that is associated with each LPAR. At a minimum, each partition has one workload that is associated with it.

If you included incorrect data for your LPAR sizing, error messages are shown on this page.

For more information about a workload, click the specific Workload link.

PowerVM LPAR Sizing wizard: Workload definition page

You use the PowerVM LPAR Sizing wizard: Workload definition page to view a workload that is associated with a partition.

The Workload definition page provides the following information about a workload:

- · System configuration on which the workload runs
- · CPU, memory, and storage utilizations

PowerVM LPAR Sizing

COIUMDUS SCM Workload Definition	n				
Selected workload	Workload definition Select	ted system	L		
Fier: Smart Cloud Moni	itoring System: Server-9133-5	5A-SN064664	H Icolumbus	AIX - 6.1, LPA	AR Dedicate
Model: p5-550Q-9133 Feature: 8285 Clock Speed: 1650 MH:	z				
1. Total CPU Utilization					7
2. Processor cores acti	vated				3
3. Assigned Processor	Cores				
4. Memory					406
5. Disk Configuration					
Group	Storage	Read	Read	Write	Write
	Used(GB) Consumed	Ops	lOSize (bytes)	Ops	IOSize (bytes)
Name					

Figure 149. Workload definition page

PowerVM LPAR Sizing wizard: Selected system page

You use the PowerVM LPAR Sizing wizard: Selected system page to select a system for sizing.

The Selected System page shows the selected systems and solutions that are proposed for sizing.

Selected system

Choose Base System

Make one selection below, for your existing system. The sizing engine will recommend a solution based on your selection.

Select the existing system option to give preference to this base model.

For consideration of typical migration paths, select the migration path option.

The final selection in the list below allows the system selection algorithms to choose the best system without consideration of any migration or upgrade paths.

•	Back	Continue
	С	Allow the system selection algorithms to choose the best system without giving any preference based on existing system migration or upgrade paths
	œ	Choose a new model based on typical migration paths if possible (default selection)
	0	Size to Existing System Model p5-550Q-9133 8285 (1650 MHz, 4-8 cores, 20.25-38.34 rPerfs)

Figure 150. Selected system page

PowerVM Recommended Environment report

The **PowerVM Recommended Environment** report shows the current and recommended processing and memory requirements for each of the LPARs on the selected physical systems.

The PowerVM Recommended Environment report is a new report.

After you size the LPARs in step 5 in the **PlanningCenter** wizard, to view a summary of the recommendation report, click the **Capacity Planner PowerVM Optimized Environment Plan Report** link.

Power VM sizing complete.	zed Environmen	t Plan Report	Report	link
Physical Server Name	Curren	t Model	Selecte	d Model
Server-8203-E4A-SN060A7F4	520-82 4.0 Max. Core	03-E4A 25,4200.0 MHz	720-82 4.0 Max. Core	02-E4C 25,3000.0 MHz
LPAR Name	Cur	rent	Recom	mended
	Entitlement	Memory (MB)	Entitlement	Memory (MB)
Tivoli-HA1	1.0	4096	1.0	1338.0
Tivoli-HA2	1.0	4096	0.8	4762.0

Physical Server Name	Curren	t Model	Selecte	d Model
Server-8233-E8B-SN100F84P	750-82 32.0 Max. Cor	33-E8B es,3300.0 MHz	750-82 32.0 Max. Cor	33-E8B es,3600.0 MHz
LPAR Name	Cur	rent	Recom	mended
	Entitlement	Memory (MB)	Entitlement	Memory (MB)
smec05	2.0	8192	1.1	1338.0
smec08	2.0	8192	0.7	1359.0
tpmapserver116_aakanksha	1.5	1024	1.4	4771.0

Figure 151. Summary of the PowerVM Recommended Environment report

By default, the PowerVM Recommended Environment report shows recommended environment details for all of the selected physical systems. To filter this report to show the recommended environment for one or more physical systems only, in the **Managed System** list, select one or more of the servers.

v7.2 Capacity Planner PowerVM Recommended Environment

Report As Of : Oct 5, 2012 4:44:02 PM

The recommended environment report presents the recommended environment after sizing. The report shows summary of environment before and after sizing along with detait capacity, utilization for each Managed System and LPAR before and after sizing. The recommended optimization plan can be used for implementing changes, comparison with different plans to evaluate different optimization strategies.



Figure 152. Capacity Planner PowerVM Recommended Environment report

Scenario: Adding more servers

A capacity planner might add one or more servers from the PowerVM Model catalog to a what-if analysis.

Procedure

- 1. On the PlanningCenter page, which is shown in the "PowerVM PlanningCenter wizard" on page 315, click **Define Scope**, which opens the Edit Current Environment page.
- 2. Click Actions > Add Managed System, as shown here:

Edit PowerVM Current Environment

	Add Ma	anaged System				
 No filter a 	Add C	uatam Tag				
Select Managed System	Export Import Clone	Data Data Data Managed System	IMC Host lame*	Managed System Name*	Data Source	Total Processors*
		064664H	laudi	Server- 9133-55A- SN064664H	Discovery	
		656EB7A	laudi	Server- 9111-520- SN656EB7A	Discovery	

Views >> Inventory >> Managed Systems

Figure 153. Actions menu

The Create Managed System Instances window is displayed, as shown here:

Machine Serial Number	HMC Host Name	Managed System Name*	Data Source	Total Processors*	CPU Cores Activated*	CPU Speed [MHz]*
064664H	Create Mana	aged System Insta	nces		8	1,648
100F84P					32	3,300
	Select Model:		B, 720-8202-E4B, 30	00 MHz, 6 Cores [Proc Code:83	151] -	
	Number of Mar	aged System instances:	8202-E4B, 720-8202- 8202-E4B, 720-8202- 8202-E4C, 720-8202- 8202-E4C, 720-8202- 8202-E4C, 720-8202- 8202-E4C, 720-8202-	E4B, 3000 MHz, 6 Cores [Proc E4B, 3000 MHz, 8 Cores [Proc E4C, 3000 MHz, 4 Cores [Proc E4C, 3000 MHz, 6 Cores [Proc E4C, 3000 MHz, 8 Cores [Proc	Code:8351] Code:8352] Code:EPC5] Code:EPC6] Code:EPC7]	
			8205-E6B, 740-8205-	E6B, 3700 MHz, 4-8 Cores [Pro	oc Code:8347]	-

Figure 154. Create Managed System Instances window

In this window, you can select a model from the list and enter the number of servers that you want to add.

The new servers are displayed after the list of servers on the Edit Current Environment page. Ensure that the check boxes are selected that correspond to the new servers to be added in scope.

3. On the PlanningCenter page, which is shown in the "PowerVM PlanningCenter wizard" on page 315, click **Generate Plan** to regenerate the optimization plan with the increased scope.

Scenario: Adding new logical partitions

A capacity planner wants to add one or more virtual machines to the PowerVM managed systems.

About this task

Consider a PowerVM environment that is running multiple managed systems. The capacity analyst can periodically analyze the usage data to do capacity planning regularly. The capacity planner can add new fictitious LPARs to existing managed systems to do a what-if analysis. The planner can asses on which managed system to allocate the resources. This can be an actual or a predicted future request for what-if analysis.

Procedure

- 1. On the PlanningCenter page, which is shown in "PowerVM PlanningCenter wizard" on page 315, click **Define Scope** to open the Edit Current Environment page.
- 2. Select the managed systems to which you want to add the LPARs.
- **3**. To open the virtual machines view, click **Views** > **Inventory** > **Logical Partitions**. This view shows the LPARs from the selected servers.
- 4. Click Actions > Add Logical Partitions.
- 5. On the PlanningCenter page, click Edit Recommended Environment Settings.

Detailed reference: Edit Current Environment window

The Edit Current Environment window presents a spreadsheet-like view within a browser window. In this view, you can view and edit the configuration and utilization profile, as well as define scope.

Data loaded in this view is a snapshot of the current environment loaded by using **Load config**, or by adding servers, as described in "Scenario: Adding more servers" on page 331. You can subsequently add to the data.

Managed Systems Inventory View



Figure 155. Managed System Inventory View

The Managed System view presents the physical server configuration data in the current environment. To define the scope, you can select a subset of physical servers to include in the analysis session. The Status column indicates if any problem exists with matching the particular server data to a unique model in the

PowerVM catalog. A green check mark icon \bowtie indicates that a unique model entity was mapped to the server. The Model Name column is populated with the name of the model in this case. An X in a red

circle icon ⁶⁰ indicates that the managed system data could not be matched to a unique model in the PowerVM catalog. This problem of matching the server data to a unique model occurs when the LPARs on the managed system are not monitored by IBM Tivoli Monitoring OS agents. To use this server in further analysis, you update the CPU Speed [MHz] of this model manually and save the data.

The following actions are available on the Actions menu:

Add Managed System

You can add additional managed systems to this view. These systems can be used to provide for additional hardware, or to try what-if scenarios. When you choose this action, a window opens with a list of available models as shown in the Create Managed System Instances window. You can select the appropriate model and create multiple instances as needed.

When you click **Create**, a new row is added to the inventory grid with architecture details populated from the knowledge base data.

Add Custom Tag

You can extend and augment the discovered data with user-defined attributes (tags). When you select this action, you are prompted to provide the tag name and a new column is added to the inventory grid. You can then add values for this column.

These tags can be used to formulate and apply rules during optimization. See "Detailed reference: Edit recommended environment settings" on page 348 for more information.

Select Model:	S703-7891-73X, :	2400 MHz, 16	Cores [Proc Code	:52CC] *
Number of Managed System instances:	1			07.50
20 20				52

Figure 156. Create Managed System Instances window

Export Data

You can download the data in the Managed System inventory view to a CSV file, which can be edited offline to add missing information.

See Appendix A, "CSV format for data import and export," on page 381 for more information about editing CSV files.

Import Data

You can import a CSV file that was downloaded using Export Data.

Important: Note the following while editing this file:

- The first column, PHYSICAL_SERVER_PK, should NOT be edited.
- Only previously NULL or blank data is updated. If data already exists in a specified column, it will not be updated if edited.

See Appendix A, "CSV format for data import and export," on page 381 for more information about editing CSV files.

Clone managed systems

With this action, you can create clones of existing managed systems in the Capacity Planner. The cloned entries have the same configuration and utilization profile as the original entries. This feature can be helpful in what-if analyses where you want to analyze the impact of having new workloads that are similar to existing workloads in the environment.

The Cloning managed systems and logical partitions feature has the following benefits:

- You are not required to add managed systems and complete various required fields.
- All field values from the managed system entry that is being cloned are automatically copied to cloned entries.
- You can choose whether to copy tags for the clone entries that are being created.

Use the following steps to clone a managed system:

- 1. Start the Planning Center page for PowerVM.
- 2. Go to the Edit Current Environment page.
- 3. From the menu, click Actions > Clone Managed Systems (Figure 157 on page 336).



Figure 157. Edit PowerVM Current Environment page

The Clone Managed Systems window is displayed (Figure 158).

vailable Managed Systems		Selected Mar	naged System	S	
Filter	×	Filter			×
		Deview	Nort	Claus	0

Figure 158. Clone Managed Systems window

- 4. Move the managed system from the **Available Managed Systems** list to the **Selected Managed Systems** list (Figure 158).
- 5. Click Next.

The second pane of the Clone Managed Systems window is displayed (Figure 159).

			Filter	×
Managed System Name	Number of Instances*	Clone Name Prefix*		Copy Server Tags
Server-8203-E4A- SN060A7E4	1	Clone_Server-8203-E4A	SN060A7F4	

Figure 159. Clone Managed System second page

- 6. In the Number of Instances field, select a value.
 - The number indicates the number of clone entries for the managed system that you want to create.
 - The default value is 1.
- 7. In the **Clone Name Prefix** field, enter a value.
 - Each clone entry must be prefixed with the string value provided in this field and with a unique time stamp value.
 - The default value is Clone_MANAGED_SYSTEM_NAME.
- 8. Select the Copy Server Tags check box to enable the tags.
 - If enabled, all tags for the selected managed system are copied to clone entries.
 - By default, this check box is selected.
- 9. Click Clone.

Reports

Capacity Planner Managed System Inventory

The Capacity Planner Managed System Inventory report presents an overall view of the servers

that are selected for analysis. The report contains a summary table that includes the managed system resources and distribution charts from the HMC, Data Center, and Power Family.



Figure 160. Managed System Inventory View

Logical Partition Inventory view

ews	 Actions 	Reports •								Filter	
tion	Managed System Name*	Logical Partition HostName*	LPAR Name*	Data Source	Entitlement*	Memory Assigned [MB]*	OS*	OS Version*	Shared Mode	Capped Mode	Midd
	Server-8233-E8B- SN100F84P	smecaix05	smec05	Discovery	2.0	8,192	AIX	6.1	shared	uncapped	
	Server-8233-E8B- SN100F84P	smecaix06.in.ibm.com	smec06	Discovery	2.0	8,192	AIX	6.1	shared	uncapped	
	Server-8233-E8B- SN100F84P	smecaix07.in.ibm.com	smec07	Discovery	2.0	8,192	AIX	6.1	shared	uncapped	
	Server-8233-E8B- SN100F84P	smecaix08.in.ibm.com	smec08	Discovery	2.0	8,192	AIX	6.1	shared	uncapped	
	Server-8233-E8B- SN100F84P	smecaix11.in.ibm.com	smecaix11	Discovery	0.2	2,048	AIX	7.1	shared	uncapped	
	Server-8203-E4A- SN060A7F4	tivoliha1	Tivoli-HA1	Discovery	1.0	4,096	AIX	7.1	dedicated	donating	
	Server-8203-E4A- SN060A7F4	tivoliha2.in.ibm.com	Tivoli-HA2	Discovery	1.0	4,096	AIX	7.1	shared	uncapped	



The Logical Partition Inventory view presents the logical partition configuration data from the selected set of physical servers. You can make corrections, add missing data or add new attributes (tags) that cannot be discovered in the current environment (such as middleware details, or tags such as criticality, SLA factors, and so on).

The following actions are available on the Actions menu:

Add Logical Partition

You can add a logical Partition to provide for future workloads. A new row is added to the

inventory grid where you can populate details of the new logical partition. Logical Partitions can be added to any server in the working set. These added LPARs are used to perform What-if analyses to plan for new workloads.

Add Custom Tag

You can extend and augment the discovered data with user-defined attributes (tags). When you select this action, you are prompted to provide the tag name and a new column is added to the inventory grid. You can then add values for this column.

These tags can be used to formulate and apply rules during optimization. See "Detailed reference: Edit recommended environment settings" on page 348 for more information.

Export Data

You can download the data in the Logical Partition inventory view to a CSV file, which can be edited offline to add missing information.

Import Data

You can import a CSV file that was downloaded using Export Data.

Important: Note the following while editing this file:

- The first column, VIRTUAL_MACHINE_PK, should NOT be edited.
- Only previously NULL or blank data is updated. If data already exists in a specified column, it will not be updated if edited.

Clone logical partitions

With this feature, you can create clones of existing logical partitions in the Capacity Planner. The cloned entries have the same configuration and utilization profile as the original entries. This feature can be helpful in what-if analyses where you want to analyze the impact of having new workloads that are similar to existing workloads in the environment.

The cloning logical partitions feature has the following benefits:

- You are not required to add logical partitions and complete various required fields.
- All field values from the logical partition entry that is being cloned are automatically copied to cloned entries.
- You can associate all clone logical partitions to a managed system that is available in the working set.
- You can choose whether to clone sizing data for logical partitions that are being cloned.
- You can choose whether to copy tags for the clone entries that are being created.

Use the following steps to clone a logical partition:

- 1. Start the Planning Center page for PowerVM.
- 2. Go to the Edit Current Environment page.
- 3. Click Views > Inventory > Logical Partitions (Figure 162 on page 340).



Figure 162. Edit PowerVM Current Environment page for logical partitions

4. Click Actions > Clone Logical Partitions (Figure 163).

Views 👻	Actions 🕶 Re	ports	- -	(909)
Action	Add Logical Partition Add Custom Tag Export Data		al Partition lame*	LPAR Name*
	Import Data		aix10.in.ibm.com	smecaix10
	Clone Logical Partitio	ns		
	Server-8203-E4A- SN060A7E4	ha1.i	n.ibm.com	HA1

Figure 163. Edit PowerVM current Environment Actions tab

The Clone Logical Partitions window is displayed (Figure 164 on page 341).

vailable Logical Partitions			Selected Logica	I Partitions	
Filter	×		Filter		×
history.in.ibm.com jackson topeka		> >> «			

Figure 164. Clone Logical Partitions window

- 5. Move the logical partition from the **Available Logical Partitions** list to the **Selected Logical Partitions** list (Figure 164).
- 6. Click Next.

The second pane of the Clone Logical Partitions window is displayed (Figure 165 on page 342).

				Filter	×
LPAR Name	Number of Instances*	Clone Name Prefix*	Copy Server Tags	Managed System Name	Copy Sizing Data
latlanta	1	Clone_latlanta		Server-9133-55A-SN064665F *	
ljackson	1	Clone_ljackson		Server-9133-55A-SN064664F -	
Cotate o Calantade				5 10 25 5	0 1 100 1 44 +

Figure 165. Clone Logical Partitions second page

- 7. In the Number of Instances field, select a value.
 - The number indicates the number of clone entries for the selected logical partition that you want to create.
 - The default value is 1.
- 8. In the Clone Name Prefix field, select a value.
 - Each clone entry is prefixed with the string value that is provided in this field and with a unique time stamp value.
 - The default value is Clone_LOGICAL_PARTITION.
- 9. Select the Copy Server Tags check box to enable the tags.
 - If enabled, all tags for the selected managed system are copied to clone entries.
 - By default, this check box is selected.
- 10. In the Managed System Name field, select a managed system.
 - By default, the managed system that is displayed is the one to which the selected logical partition is pointing.
 - All managed systems that are in the working set are populated in this drop-down list.
- 11. Select the Copy Sizing Data check box to enable the copying of sizing data.
 - If enabled, sizing data values for selected logical partitions are copied to clone entries.
 - By default, this check box is selected
 - Special Case: The **Copy Sizing Data** check box is enabled and the selected logical partition is a discovered partition with no sizing data values.
- 12. Click Clone.

Reports

Logical Partition Inventory

The Logical Partition Inventory report shows an overall view of the LPARs and VIO partitions on

the selected Managed Systems. The report has a summary table of the resources that are allocated to the LPARs and distribution charts by Data Center and HMC, OS and Middleware that are installed.



Figure 166. LPAR Inventory report

VIO Partition Inventory View

uncapped	shared					Data Source	LPAR Name*	HostName*	Name*	
		2.2.1.1	VIOS	2,048	0.2	Discovery	Power-vio2	power-vio2	Server-8203-E4A- SN060A7F4	
uncapped	shared	2.2.1.1	VIOS	2,048	0.2	Discovery	Power-Vio	powervio1	Server-8203-E4A- SN060A7F4	
uncapped	shared	2.2.1.4	VIOS	2,048	0.2	Discovery	smecaix02-vio-server	smecaix02	Server-8233-E8B- SN100F84P	
	shared	2.2.1.4	VIOS	2,048	0.2	Discovery	smecaix02-vio-server	smecaix02	er-8233-E8B- 00F84P	Serv SN1



The VIO Partition Inventory view shows details of the VIO partitions on the selected Managed Systems. You can edit and update the data in this view with custom tags.

The following actions are available on the Actions menu:

Add VIO Partition

You can add a VIO partition to perform What-if analysis. A new row is added to the current grid where you populate the details of the VIO server.

Important: In the VIO Inventory view, if you change the Managed System for a VIO partition manually, you must also update the Logical Partition VIO associations manually in the Logical Partition Inventory view.

Logical Partition Utilization View

View Action Report Report CPU Manged System Removed Manged Mange	Views >	> Logical Par tion Utiliza	tion									
Image: Partition Image: Partition Image: Partition Image: Partition Read Operation Read Operation Read Use operation	Views • Actions • Reports									Filter	₽. ₽	
Image: Sense 2005 Sense 2003-EBB- Sense 2003-EBB- 1250 7650 22.0 1 Image: Sense 2005 in Libration Sense 2003-EBB- Sense 2003-EBB- 1660 17.0 657.0 17.40 876.0 24.40 2 Image: Sense 2005 in Libration Sense 2003-EBB- Sense 2003-EBB- 1660 657.0 17.40 876.0 246.0 1 Image: Sense 2003 EBB- Sense 2003-EBB- Sense 2003-EBB- 1660 657.0 546.0 870.0 2456.0 1 Image: Sense 2003 EBB- Sense 2003-EBB- Sense 2003-EBB- 160.0 650.0 540.0 870.0 2456.0 1 Image: Sense 2003 EBB- Sense 2003-EBA- Sense 2003-EBA- 15.0 100.0 760.0 540.0 870.0 2456.0 1 Image: Sense 2003 EBB- Sense 2003-EBA- Sense 2003-EBA- 100.0 1		Logical Partition HostName	LPAR Name	Managed System Name	CPU Usage [Share]	Memory Usage [MB]	Read Operations	Write Operations	Read Usage[Bytes]	Write Usage[Bytes]	Network Operations	Netv Usaj
Image:		smecaix05	smec05	Server-8233-E8B- SN100F84P	28	1,265	1223.0	4326.0	125.0	765.0	222.0	1
Image: Sinter		smecaix06.in.ibm.com	smec06	Server-8233-E8B- SN100F84P	65	1,124	1767.0	657.0	174.0	876.0	244.0	2
Image: Sense Sinder		smecaix07.in.ibm.com	smec07	Server-8233-E8B- SN100F84P	45	876	176.0	657.0	546.0	879.0	2456.0	1
Smecaix11 Smecaix11 <t< td=""><td></td><td>smecaix08.in.ibm.com</td><td>smec08</td><td>Server-8233-E8B- SN100F84P</td><td>37.9</td><td>532</td><td>176.0</td><td>65.0</td><td>54.0</td><td>87.0</td><td>2456.0</td><td>1</td></t<>		smecaix08.in.ibm.com	smec08	Server-8233-E8B- SN100F84P	37.9	532	176.0	65.0	54.0	87.0	2456.0	1
Involima1 Twoli-HA1 Server-8203-E4A- SN060A7F4 15.1 129.373 614.0 1800.0 6 Involima1 Twoli-HA2 Server-8203-E4A- SN060A7F4 52.5 129.373 514.0 1800.0 6	~	smecaix11.in.ibm.com	smecaix11	Server-8233-E8B- SN100F84P	59.8	2,345	980.0	760.0	544.0	8723.0	2456.0	3
Introduction Twoli-HA2 Server-8203-E4A- SN080A7F4 52.5 129.373 614.0 1800.0 6		tivoliha1	Tivoli-HA1	Server-8203-E4A- SN060A7F4	15.1	129,373				614.0	1800.0	6
		tivoliha2.in.ibm.com	Tivoli-HA2	Server-8203-E4A- SN060A7F4	52.5	129,373				614.0	1800.0	6
	_	f e e literare			10 25 50 100 40							

Figure 168. Logical Partition Utilization View

A Capacity Planner uses this view to analyze the usage and performance profiles of Logical Partitions on the selected Managed Systems. The usage computation is done by using the federated performance data in the Tivoli Data Warehouse that is configured to the Tivoli Monitoring server. As per the analysis time period that is set in Step 2 on the Planning center, the calculation is performed on the specified interval. For more information about setting the analysis time period, see step 2 of the PowerVM Planning Center wizard: "PowerVM PlanningCenter wizard" on page 315.

The following actions are available on the Actions menu:

Compute Usage

You can compute the usage requirement of logical partitions by using different parameters, as shown in the Compute Usage window:

		Summarizati	on Selected Summarization
CPU:	Maximum	Veekly 🗸	All First Week Second Week
Memory:	Maximum	▼ Weekly ▼	All First Week Second Week
Storage:	Maximum	▼ Weekly ▼	All First Week Second Week Third Week
Network:	Maximum	Veekly -	All First Week Second Week

Figure 169. Compute Usage window

Compute Usage calculates the resource usage on the logical partitions: CPU, memory, network bandwidth and disk I/O. This sizing is done by analyzing the utilization data available in Tivoli Data Warehouse based on the summarization and aggregation levels specified. Aggregation levels available are **Average**, **Minimum**, **Maximum**, and **90th Percentile**. Summarization levels available are **Hourly**, **Daily**, **Weekly**, and **Monthly**. The available values in the **Selected Summarization** field depend on which value was selected in the **Summarization** field.

Important: Usage numbers are generated only for logical partitions that have Tivoli Monitoring OS agents configured for monitoring purposes.

Generate Workload Stability Type

The **Generate Workload Stability Type** analyzes the hourly utilization data for a logical partition and determines whether the resource utilization is stable or unstable, depending on the variation in usage.

	Measurement Type	
CPU:	Average 🔻	
Memory:	Average 💌	
Storage:	Average 💌	
Network:	Average 💌	

Figure 170. Generate Stability Characteristic Parameters window

Edit Usage

You can manually edit or *Adjust-for-growth* the resource usage. You can apply different growth profiles to adjust as needed. The usage parameters can be specified in absolute units, for example, 60% CPU, or a growth percentage can be applied, for example, add 10% growth to memory.

Absolute Value O Growth	
CPU Usage:	%
Memory Usage:	MB
Read Usage	KBytes
Write Usage	KBytes
Network Usage	MBytes
Read Ops	
Write Ops	
Network Ops	

Figure 171. Edit Usage window

Reports

The following reports are available:

Utilization Aggregated Timeseries report

This report can be used to identify utilization patterns of logical partitions. Because you can view aggregations of multiple virtual machines at a time, you can also identify correlations in the resource utilizations. You can use these observations to determine the usage sizing summarization level. An example graph is shown in the Utilization Aggregated Timeseries report:



Figure 172. Utilization Aggregated Timeseries report

Utilization Detailed Timeseries report

This report helps you identify any data gaps in the utilization data collected for the logical partitions. Data points come directly from aggregated measurement tables in utilization schema. An example graph is shown in the Utilization Aggregated Timeseries report.

Detailed reference: Edit recommended environment settings

A capacity planner guides the Sizing wizard by adding and enabling the Sizing rules. These rules are derived from technical best practices, environment experience, or performance data trends and forecasts. The rules act upon the usage values that are calculated by the Compute Usage action in the Logical Partition Utilization view for the matching logical partitions.

Sizing Rules overview

Use the Edit PowerVM Recommended Environment Settings window to edit rules (Figure 173 on page 349).
×

Edit PowerVM Recommended Environment Settings

×

Select sizing rules.

ACUAC	Rule Instances	
	Add 30% more CPU to DB2 LPARs	
	Add 25% more CPU to AIX 7.1 LPARs	
	Provide 10% more CPU for Critical VMs	
	Add 35% more CPU to AIX 6.1 with DB2	
Active	Rule Instances	
	For DD2 workloads with 4 million transaction act antitlement to 2.5 and memory to 0.00	
	For DB2 workloads with 1 million transaction set entitlement to 3.5 and memory to 8GE	
	For WAS workloads with 1000 users set entitlement to 2.3	
	For Tivoli Group Workloads set Memory to 6.5GB	

Figure 173. Edit PowerVM Recommended Environment Settings window

The Edit PowerVM Recommended Environment Settings window shows the Sizing rules that are available in the Capacity Planner database. You select or clear the check box that relates to a rule to enable or disable the rule and to guide the sizing procedure. The sizing rule specifies a growth factor that is applied to the CPU entitlement of a logical partition that matches a required condition. The Sizing wizard uses this updated CPU entitlement value to perform further sizing analysis. These rules might be based on discovered attributes, such as operating system name or user-defined custom tags.

Rule syntax

```
Utilization
        <?xml version="1.0" encoding="UTF-8"?>
        <rule>
         <if>
          <antecedent operator="GROUPING_OPERATOR">
           <literal>
            <path>DATABASE_PATH</path>
            <operator>"PATH OPERATOR"</operator>
            <term>PATH_VALUE</term>
           </literal>
        <literal>
            <path>DATABASE PATH</path>
            <operator>"PATH_OPERATOR"</operator>
            <term>PATH_VALUE</term>
           </literal>
          </antecedent>
         </if>
         <then>
          <dependent>
           <literal>
            <path>target.capacity</path>
            <operator>Add</operator>
            <term>NUMBER</term>
           </literal>
          </dependent>
         </then>
        </rule>
```

Where:

GROUPING_OPERATOR

AND or OR

PATH_OPERATOR

EQ or NEQ

DATABASE PATH

See "Supported database paths for writing rules" on page 353.

PATH_VALUE

All possible values defined for DATABASE_PATH

NUMBER Valid number in percentages (positive or negative)

Example sizing rule: Increase CPU entitlement by 50% for LPARs tagged Critical

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
<if>
  <antecedent operator="AND">
  <literal>
    <path>Source.ANL VM PLACEMENT SET V*
              ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_TYPE</path>
   <operator>EQ/operator>
   <term>CRITICALITY_SLA</term>
   </literal>
   </literal>
   <path>Source.ANL VM PLACEMENT SET V*
     ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_NAME</path>
   <operator>EQ</operator>
   <term>High</term>
   </literal>
```

```
</antecedant>
</if>
<then>
<dependent>
<literal>
<path>target.capacity</path>
<operator>Add</operator>
<term>50</term>
</literal>
</dependent>
</then>
</rule>
```

Defining resource sizing by using application sizing rules

You can use the application sizing rules policy template to help define resource sizing that is based on application parameters. Use the sizing rules to adjust computed usage data for LPARs that match given criteria. By using application sizing rules, you can override computed usage sizing data for resource parameters for LPARs that match the rule criteria.

About this task

Use of the sizing rules has the following benefits:

- You can define policies to set resource allocation based on discovered attributes or custom tags with application parameters information.
- You can set target allocation and entitlement values for the following metrics:
 - CPU Entitlement
 - Memory Allocation
 - Storage read/write operations
 - Network operations

For example:

- For DB2 workloads with one million transactions, set entitlement to 3.5 and memory to 8GB.
- For WebSphere Application Server workloads with 1,000 users, set entitlement to 2.3.

Procedure

Use these steps to define and use application sizing rules:

- 1. Identify existing attributes for LPARs such as Middleware Name, or define custom attributes such as Transactions per second by clicking **Actions** > **Add Custom tag** in the LPAR inventory.
- 2. Set the values for these parameters in the Inventory view and click Save.
- 3. Define the rule (see the sample in "Rule syntax") and save the RuleDim.csv file.
- 4. Import the rules to Capacity Planner by using the Load Knowledge base action in the Inventory view in the upper-right corner.
- 5. Navigate to the Planning Center Edit Recommended Environment settings page and enable or disable the rule as required.
- 6. Complete the Generate plan step to generate recommendations. The enabled rules guide the sizing of the LPARs that match the rule criteria in the recommendation generation.

Rule syntax

When you are using the application sizing rules, use the following Rule XML format:

```
<?xml version="1.0" encoding="UTF-8"?>
<rule>
<if>
```

```
<antecedent operator="GROUPING OPERATOR">
  <literal>
    <path>DATABASE PATH</path>
    <operator>"PATH_OPERATOR"</operator>
   <term>PATH VALUE</term>
   </literal>
   <literal>
    <path>DATABASE PATH</path>
    <operator>"PATH_OPERATOR"</operator>
   <term>PATH_VALUE</term>
  </literal>
 </antecedent>
 </if>
 <then>
  <dependent>
  <literal>
   <path>TARGET FIELD</path>
   <operator>SET</operator>
   <term>NUMBER</term>
   </literal>
   <literal>
    <path>TARGET FIELD</path>
    <operator>SET</operator>
   <term>NUMBER</term>
  </literal>
 </dependent>
</then>
</rule>
```

Where:

GROUPING_OPERATOR

AND or OR

PATH_OPERATOR

EQ or NEQ

DATABASE_PATH

See "Supported database paths for writing rules" on page 353.

PATH_VALUE

All possible values defined for DATABASE_PATH

TARGET_FIELD

All possible values defined for TARGET_FIELD

NUMBER

Valid number in percentages (positive or negative)

SET is supported as the valid value in the operator tag in the dependent literal of rule definitions.

Table 90 provides descriptions of the target fields that can be used in the dependent literal path of a rule.

Table 90. Target fields required for application sizing rule

Target field name	Description
Target.capacity	Set CPU capacity required for LPAR or VM
Target.Entitlement	Set CPU entitlement for an LPAR
Target.memory	Set memory capacity required for LPAR or VM
Target.read_ops	Set number of read operations per second required for LPAR or VM
Target.read_size	Set read bytes per read operation seconds required for LPAR or VM

Table 90. Target fields required for application sizing rule (continued)

Target field name	Description
Target.write_ops	Set number of write operations per second required for LPAR or VM
Target.write_size	Set write bytes for LPAR or VM
Target.net_ops	Set number of network operations per second required for LPAR or VM
Target.net_ops_size	Set network transfers in bytes per network operation required for LPAR or VM

A sample RuleDim.csv entry is provided here for two example rules:

• For DB2 workloads with 1 million transactions, set entitlement to 3.5 and memory to 8 GB.

Attributes used: LPAR Inventory view > Middleware name

Custom tags: LPAR Inventory view > Transactions Per Minute

```
,"For DB2 workloads with 1 million transaction set entitlement to 3.5 and memory to
8GB",1,"General",0,"Application Sizing","No Notes","<?xml version="1.0" encoding="UTF-8"?
><rule><if><antecedent
operator="AND"><literal><path>Source.ANL_VM_PLACEMENT_SET_V<CFG_VIRTUAL_MACHINE.MIDDLEWARE_NAME</
path><operator>EQ</operator><term>DB2</term></literal><literal><path>Source.ANL_VM_PLACEMENT_SET_
V*ANL_VIRTUAL_MACHINE_TAG_MAP.SERVER_TAGS_PK.TAG_TYPE</path><operator>EQ</operator><term>TRANSACT
```

IONSPERMINUTE</term></literal><literal><path>Source.ANL_VM_PLACEMENT_SET_V*ANL_VIRTUAL_MACHINE_TA G_MAP.SERVER_TAGS_PK.TAG_NAME</path><operator>EQ</operator><term>1000000</term></literal></antece dent></if><then><dependent><literal><path>target.memory</path><operator>set</operator><term>8192< /term></literal><literal><path>target.entitlement</path><operator>set</operator><term>3.5</term></literal></if>/literal><path>target.capacity</path><operator>set</operator><term>85</term></literal></term></literal></term></literal></term></literal></term></literal></term></literal></term></literal></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term></term>

• For WebSphere Application Server workloads with 1,000 users, set entitlement to 2.3.

Attributes used: LPAR Inventory view > Middleware name

Custom tags: LPAR Inventory view > Num Of Users

```
,"For WAS workloads with 1000 users set entitlement to 2.3",1,"General",0,"Application
Sizing","No Notes","<?xml version="1.0" encoding="UTF-8"?><rule><if><antecedent
operator="AND"><literal><path>Source.ANL_VM_PLACEMENT_SET_V<CFG_VIRTUAL_MACHINE.MIDDLEW
ARE_NAME</path><operator>EQ</operator><term>WebSphere Application
Server</term></literal><literal><path>Source.ANL_VM_PLACEMENT_SET_V*ANL_VIRTUAL_MACHINE
_TAG_MAP.SERVER_TAGS_PK.TAG_TYPE</path><operator>EQ</operator><term>NUMOFUSERS</term></literal><literal><path>Source.ANL_VM_PLACEMENT_SET_V*ANL_VIRTUAL_MACHINE
_TAGS_PK.TAG_NAME</path><operator>EQ</operator><term>1000</term></literal><literal><path>target.entitlement</path><operator>set</operator><
term>2.3</term></literal><literal></path>target.capacity</path><operator><term>100</term></literal></path><operator><term>100</term></literal></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path></path>
```

Supported database paths for writing rules

Database paths are used to specify specific attributes from the database. A database path consists of table and view names, column names, and operators such as asterisks (*) and dots (.). An asterisk is used where tables and views have a one-to-many relationship. A dot is used where tables and views have a one-to-one relationship.

Ensure that database paths do not include spaces or new line characters.

Table 91 on page 354 contains a list of commonly used discovered attributes and their corresponding database paths.

Table 91. Logical partition attributes

Attribute	Database path
Virtual machine ID	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.VMID
Virtual machine name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.VM_NAME
Operating system name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.OPERATING_SYSTEM_NAME
Operating system version	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.OPERATING_SYSTEM_VERSION
Application name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.MIDDLEWARE_NAME
Application version	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.MIDDLEWARE_VERSION
Physical server host name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK.HOST_NAME
Data center name	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK. DATA_CENTER_NAME
Physical servers model	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK.MODEL
Physical servers architecture	Virtual_Machine.ANL_VM_PLACEMENT_SET_V* CFG_VIRTUAL_MACHINE.PHYSICAL_SERVER_PK. HYPERVISOR_TYPE
CPU variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.CPU_VAR
Memory variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.MEM_VAR
Network variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.NETWORK_VAR
Disk variance	Virtual_Machine.ANL_PLACEMENT_SET_V* ANL_RESOURCE_DEMAND.DISK_VAR

Writing rules for custom attributes

Use the following code in rules for custom attributes for virtual machines:

```
</path><operator>EQ</operator>
<term>TAG_TYPE</term>
</literal>
<literal>
<path>Physical_Server.ANL_WORKING_SET_V*ANL_SERVER_TAG_MAP.SERVER_TAGS_PK.TAG_NAME
</path><operator>EQ</operator>
<term>TAG_VALUE</term>
</literal>
```

In each case, replace **TAG_TYPE** and **TAG_VALUE** with appropriate values, for example, to write a rule for virtual machines with production environment, then set **TAG_TYPE** to Environment and **TAG_VALUE** to Production. These values should exactly match the database values. Use AND as the grouping operator.

Importing rules in Capacity Planner

The XML snippets used to author policies must be kept in a single file in CSV format and loaded into the Planner Knowledge Base by clicking **Load Knowledge Data**. Any existing policies in the Knowledge Base are deleted when a new set is loaded.

Column name	Maximum column length	Details
RuleDimKey	8	Primary key of this rule
RuleName	50	Readable identifier of this rule
Priority	4	Rule Priority. 1 is the highest priority. All rules that cannot be broken should have priority 1.
Source	50	Metadata, for example, General, Customer, and BestPractice
IsActive	1	1 for yes, θ for no
RuleType	50	Type of rule, such as App Selection, OS selection, and so on
Notes	255	Additional information
RuleXML	Binary large object	XML format of rule

Table 92. CSV format details

A sample CSV file is in the dash_home\installedDashboards\com.ibm.tivoli.cppowervm\ AnalyticsDatabaseInstaller\samples directory.

Complete the following steps in order to import rules in Capacity Planner:

- 1. Open the Edit Current Environment window.
- 2. Click the Load Knowledge Data icon 🖾 .
- **3**. Select the **Rules** check box. In the corresponding field, enter the path of the CSV file that contains the rules.

Rules	Browse	e_
		F -

Figure 174. Refresh Knowledge Base Content window

4. Click Upload.

356 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Chapter 13. Troubleshooting

Some problems and solutions for those problems are identified in the following areas: Installation and configuration, timeout and memory, Performance and Capacity Management Reports, Dashboard for VMware, and the Capacity Planners.

Installation and configuration problems

Use the installation problems table to troubleshoot errors you might encounter when you install and configure the Dashboard for VMware, Performance and Capacity Management Reports, Capacity Planner for PowerVM, and Capacity Planner for VMware.

Problem	Solution
Capacity Planner for VMware A problem occurred when you run the heterogeneous	Set the following Environment Variables or Permissions for heterogeneous federation with Oracle on the Windows operating system.
federation (with Oracle Server) scripts for Capacity	• In the System Variables, set a value for ORACLE_HOME.
The log displays the following error messages:	• In the User Variables, set the values for DB2LIBPATH and ORACLE_HOME.
SQL1822N Unexpected error code -1 received from data source "WHSERV".	• Provide all permissions to the db2 instance owner for the lib directory in \$ORACLE_HOME.
Associated text and tokens are "Do not have permission to read Oracle Client Libraries".	• Add the details for the following parameters in the db2dj.ini file:
SQLSTATE=560BD.	- ORACLE_HOME
	– DB2LIBPATH
The DBinstaller is looking for bash and failing on AIX.	Bash is a prerequisite on AIX, but is not available by default on AIX. You can download it from IBM AIX Toolbox download (http://www-03.ibm.com/systems/ power/software/aix/linux/toolbox/download.html). See also IBM developerWorks System Administration Toolkit: Get the most out of bash (http://www.ibm.com/ developerworks/aix/library/au-satbash.html).

Table 93. Capacity Planner installation and configuration problems.

Table 94. Linux Kernel-based Virtual Machines agent installation and configuration problems

Problem	Solution
Report installation fails on Linux systems with the following error found in the logs: The <i>fileset</i> type doesn't support the "erroronmissingdir" attribute.	The Linux environment has another version of ANT earlier than 1.8.1 configured as the default. Set the following environment variables and run the installer again: export PATH=Install_dir\reports\lib\ apache-ant-1.8.1\bin:\$PATH export ANT_HOME=Install_dir\reports\lib\ apache-ant-1.8.1 export CLASSPATH=Install_dir\reports\lib\ apache-ant-1.8.1\bin:\$CLASSPATH Before you start the installer again, run ant -version to check that the returned results show ant 1.8.1 compiled.

Problem	Solution	
You get errors when you run the reports installer.	The report installer completes the following 3 steps:	
	1. Imports reports	
	2. Defines the Tivoli Data Warehouse data source in Cognos	
	3. Makes schema updates by running scripts against the Tivoli Data Warehouse (adding and populating Time Dimension tables, creating indexes if history is enabled on the attribute groups).	
	Check the following logs on Windows systems at C:\Documents and Settings\Administrator or on Linux or UNIX systems at \$HOME:	
	 Report_Installer_InstallLog.log 	
	 Report_Installer_For_TCR_Output.txt 	
When you install reports, the installer fails with error messages similar to the following error message: JVMDUMP006I Processing dump event "systhrow", detail "java/lang/OutOfMemoryError" - please wait.	This problem is a known problem with Tivoli Common Reporting. For more information, see the Tivoli Common Reporting technote (OutOfMemoryError exception occurs when issuing trcmd command).	
 JVMDUMP013I Processed dump event "systhrow", detail "java/lang/OutOfMemoryError". CTGTRQ010E Processing has ended because of an unexpected error.		

Table 94. Linux Kernel-based Virtual Machines agent installation and configuration problems (continued)

Problem		Solution
In the Reports Installer, when you complete the following steps to install reports, you have some problems:		Restart the Reports Installer.
1.	On the Choose the reports for the installation page, select all reports that you want to install, and then click Next . The Cognos Engine Configuration page is displayed.	
2.	On the Cognos Engine Configuration page, click Previous to return to the Choose the Installation Folder page.	
3.	On the Choose the Installation Folder page, click Next	
	Problem: The Choose the reports for the installation page that contains selected check boxes for all reports is displayed, and Next is disabled.	
4.	Clear all check boxes for the reports, and select the reports to install again. The Next button becomes available.	
5.	Click Next.	
	Problem: The following message is displayed on the next page: THERE ARE NO ITEMS SELECTED FOR THE INSTALLATION.	

Table 95. Reports installation and configuration problems (continued)

Problem	Solution
Labels are not displayed on the installation panels or in the dialog boxes on an AIX operating system with Turkish locales when the Report Installer was run on	This problem occurs because both the background color and the font color are white. Use one of the following workarounds:
Java 6.	• Change the style palette to defaultmono when the Report Installer is running (if it is not set to defaultmono by default). This solution works for Java 5 and Java 6 and is the solution that is preferred.
	 Run the Report Installer by using Java 5. You can specify Java by using the following command:
	setup_aix.bin lax_vm /opt/ibm/java5/jre/bin/java

Table 96.	VMware	agent	reports	problems
10010 00.	v mnulo	ugoin	1000110	problotitio

Problem	Solution
You get errors when you run the reports installer.	The report installer completes the following 3 steps:
	1. Imports reports
	 Defines the Tivoli Data Warehouse data source in Cognos
	3 . Makes schema updates by running scripts against the Tivoli Data Warehouse (adding and populating Time Dimension tables, creating indexes if history is enabled on the attribute groups).
	Check the following logs on Windows systems at C:\Documents and Settings\Administrator or on Linux or UNIX systems at \$HOME:
	 Report_Installer_InstallLog.log
	 Report_Installer_For_TCR_Output.txt

Analyzing the report installer log

Review the Report_Installer_For_TCR_Output.txt file (on Windows systems under C:\Documents and Settings\Administrator; on Linux and UNIX systems under \$HOME.) to identify the step on which the installer failed.

Sample log output

```
INSTALLATION COMPLETED.
The status of installation steps:
TCRRunDBScripts(runDbScript): FAILED
INFORMATION: /tmp/450480.tmp/reports/itmfvs/build.xml:31:
The <fileset> type doesn't support the "erroronmissingdir" attribute.
InstallReportsAction(IBM Tivoli Monitoring for
Virtual Environments Reports v7.1): SUCCESS
CognosDataSource(TDW): SUCCESS
```

Analysis

In the sample log, the success or failure of each step is evident:

1. InstallReportsAction (Step 1 - Importing Reports) succeeded.

- 2. CognosDataSource(TDW) (Step 2 Defining the Tivoli Data Warehouse data source in Cognos) succeeded.
- **3**. RunDBScripts (Step 3 Updating schema by running scripts against the Tivoli Data Warehouse) failed.

Step 2: Define the Tivoli Data Warehouse data source in Cognos.

Possible causes of the failure:

- The database alias that is specified during installation did not match the cataloged DB2 database alias, the Oracle local TNS service name, or the MS SQL Server ODBC data source name.
- The credentials are incorrect for connecting to the Tivoli Data Warehouse.

Solution:

• Ensure that you installed the database client on the same server as Tivoli Common Reporting and cataloged the database. If you are using Oracle, the TNS service name in the tnsnames.ora file must be defined. If you are using MS SQL server, an ODBC data source must be defined. See Connecting to the Tivoli Data Warehouse using the database client over ODBC (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/ adminuse/tcr_tdwconnect.htm) in the *IBM Tivoli Monitoring Administrator's Guide*. If you already have a Tivoli Data Warehouse data source defined, adding another one overwrites the existing data source.

Step 3: Make schema updates

Possible causes of failure:

- Database administrative privileges (such as db2admin or sys) are required for this step; if user ITMUSER is specified, the schema cannot be updated.
- Database issues like connectivity problems, full logs, space issues, or any other performance problems that prevents writing to the database.

Solution:

- An error at Step 3 is accompanied by an informational message that contains SQL errors with SQL codes. You can search on the SQL code to determine the problem.
- If Time Dimension tables are present in the database, you can choose to skip the schema update (JDBC) step when you run the dashboard installer. If you want to create time dimension with a different granularity, you must edit the following sql file:
 - 1. Go to reports package\reports\cognos_reports\itmfvs\db_scripts.
 - 2. Open call_proc_DB2.sql , call_proc_MSSQL.sql, or call_proc_ORACLE.sql depending on the database used.
 - **3**. Edit the last parameter in the call to IBM_TRAM.CREATE_TIME_DIMENSION.

Notes

• The Database scripts for creating indexes are provided for enhanced reporting performance in the Tivoli Data Warehouse. If your data warehouse is not prepared with history before installation, this step is skipped by the installer. You can manually run one the following scripts, depending on your database type:

create_index_DB2.sql
create_index_MSSQL.sql
create_index_ORACLE.sql

For more information, see Creating shared dimensions tables and populating the time dimensions table (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/adminuse/tcr_reports_dimensionsshared.htm).

• Although indexes help enhance report performance, some limitations apply: Use indexes only on large tables with thousands of rows; because indexes degrade the performance of insert, update, and delete operations on a table.

You can run a script to drop these indexes if you run into either of these performance issues: drop_index_DB2.sql drop_index_MSSQL.sql drop index ORACLE.sql

- Connections under the Tivoli Data Warehouse are overwritten by the report installer. This is a limitation of the current installer.
- The privileges that are required when you run the installer are ITMUSER (database user) for the Tivoli Data Warehouse creation step and ADMIN (database administrator) for the schema update step. The Database Test connection for the schema update panel does not check for privileges of the database user. Installation fails at the schema update step if the database user does not have administrative privileges.

Dashboard Health Checks logs

After you install and configure the Dashboard for VMware, the Dashboard Health Checks is installed and configured with the Dashboard for VMware.

• You can view the runtime log for Dashboard Health Checks in the ITMDHC_LOGYYYYMMDDHHMMSS.log file at the following location:

TIP_INSTALL_DIR\IBM\JazzSM\profile\logs\

• You can view the configuration for Dashboard Health Checks in the ITMDHCHealthChecker.cfg file at the following location:

 $\label{eq:listalledApps} TIP_INSTALL_DIR \IBM \JazzSM \profile \installedApps \JazzSM \Node01Cell \isc.ear \ITMDHC.war \WEB-INF \config$

Dashboard Health Checks configuration parameters and their descriptions:

ITMcDP_PROVIDER_TYPE

The type of IBM Tivoli Monitoring Dashboard data provider. The value of the ITMcDP_PROVIDER_TYPE parameter is IBMTivoliMonitoringServices for IBM Tivoli Monitoring V6.3 or later.

COMMON_TEPS_PING_TIME_OUT

The Tivoli Enterprise Portal Server ping timeout value, in seconds. This value indicates the maximum amount of time the Tivoli Enterprise Portal Server Ping check must take. If the operation times out before a response, the Tivoli Enterprise Portal Server is deemed unreachable. The default value is 15 seconds.

Timeout and memory problems

If the default timeout values are short for the Tivoli Common Reporting or Cognos console login, you can change the settings. If your Java virtual machine runs out of memory, you can increase the heap size.

Table 97. Timeout and memory problems

Problem	Solution
You must log in to the Cognos server frequently because of automatic timeout (the default is 60 minutes).	 Edit the following file: For Tivoli Common Reporting V2.1.1 On operating systems other than Windows, the file is in the following location: /opt/IBM/tivoli/ tipv2Components/TCRComponent/cognos/ configuration/cogstartup.xml
	 On Windows systems: C:\IBM\tivoli\ tipv2Components\TCRComponent\cognos\ configuration\cogstartup.xml
	For Tivoli Common Reporting V3.1 or later
	 On operating systems other than Windows, the file is in the following location: /opt/IBM/JazzSM/ reporting/cognos/configuration/cogstartup.xml
	 On Windows systems: C:\Program Files\IBM\JazzSM\reporting\cognos\ configuration\cogstartup.xml
	2. Locate the mdmSessionTimeout parameter.
	 Change the value from 60 minutes to a longer timeout interval. Do not set it to -1, which gives an unlimited timeout period, because connections to the Cognos server might be kept open.
	4. Save the file.
	5. Restart the Tivoli Integrated Portal Server and the Cognos Report Server by using the startTCRserver script (On operating systems other than Windows, use startTCRserver.sh; on Windows systems, use startTCRserver.bat). If you are also editing the Dashboard Application Services Hub timeout value, you can restart the server when you have made the changes that are described in the pext row

Problem	Solution	
You are prompted to log in again to the Dashboard Application Services Hub after being logged in for a while, so you would like to increase the session timeout.	Use the WebSphere Administrative Console to set the session timeout and LTPA timeout values to larger values. To open the administrative console from within the Dashboard Application Services Hub, select Settings > WebSphere Administrative Console in the navigation tree and click Launch WebSphere administrative console . 1 Set the session timeout for Dashboard Application	
	Services Hub:	
	 a. In the Navigation frame on the left side of the WebSphere Administrative Console, expand Applications > Application Types, and click WebSphere enterprise applications. 	
	b. In the list of Enterprise Applications, click isc.	
	 On the Configuration tab, click Session management. 	
	d. For the Session timeout setting, specify the number of minutes for the timeout or specify No timeout.	
	e. Click OK.	
	2. Set the LTPA timeout:	
	 a. In the Navigation frame on the left side of the WebSphere Administrative Console, expand Security, and click Global security. 	
	b. In the Global security panel, click LTPA in the Authentication section.	
	c. Enter the LTPA timeout value in minutes.	
	d. Click OK.	
	Important: Use high or unlimited values for timeouts with caution, because such values can lead to poor server performance or out of memory conditions. These timeouts are used by the server to release storage that is associated with sessions that are no longer active. Such sessions can occur when users close their browsers without logging off or are disconnected from the server because of network disruptions.	

Table 97. Timeout and memory problems (continued)

Table 97.	Timeout and	memory problems	(continued)
-----------	-------------	-----------------	-------------

Problem	Solution
You get Java core dumps when you are running out of Java virtual machine (JVM) memory.	If your JVM runs out of memory within the Tivoli Integrated Portal, you receive a Java core dump. The Java core files are displayed in the /opt/IBM/tivoli/ tipv2/profiles/TIPProfile directory on operating systems other than Windows, and in C:\ibm\tivoli\tipv2\profiles\TIPProfile on Windows systems. This problem is common on AIX systems because the default maximum heap size for Java 1.6 on AIX systems is low. Set the maximum heap size to 512 MB or higher.
	Take the following steps to update the maximum heap size for the Tivoli Integrated Portal JVM:
	 On operating systems other than Windows, the file is in the following location: /data/IBM/tivoli/ tipv2/profiles/TIPProfile/config/cells/ TIPCell/nodes/TIPNode/servers/server1
	 On Windows systems: C:\IBM\tivoli\tipv2\ profiles\TIPProfile\config\cells\TIPCell\ nodes\TIPNode\servers\server1
	 Locate the genericJvmArguments=""> parameter towards the end of the file. Notice that there are no default JVM heap settings.
	 Add -Xmx512m or -Xmx1024m to the genericJvmArguments, depending on how much system memory you have on your server. For example, genericJvmArguments="-Xmx1024m">
	 Restart the Tivoli Integrated Portal Server and the Cognos Report Server using the startTCRserver script (On operating systems other than Windows, use startTCRserver.sh; on Windows systems, use startTCRserver.bat).
	Note: . The values that are given here are appropriate for typical cases, but be aware that the heap size might be already set for other applications. The heap size adjustment must be finely tuned. If set too high, the Java process can use too much memory and slow the system. for more information, see IBM developerWorks Roadmap for WebSphere Application Server - Tuning.

Performance and Capacity Management Reports problems

Use the Performance and Capacity Management Reports table to troubleshoot any errors you encounter when you view the Cognos reports.

The reports are accessed from the Dashboard Application Services Hub navigation tree: **Reporting** > **Virtualization**.

Problem	Solution
When you run Tivoli Common Reporting reports or test the database connection in Cognos, you get errors (such as The logon failed) that reference libdb2.a.	Ensure that you set up the correct library path environment variables. The example uses LD_LIBRARY_PATH, which might be LIBPATH on some operating systems.
	Complete the following steps for Tivoli Common Reporting V2.1.1:
	 Stop Tivoli Common Reporting: /opt/IBM/tivoli/ tipv2Components/TCRComponent/bin/ stopTCRserver.sh tipadmin tippass
	2. Open /opt/IBM/tivoli/tipv2Components/ TCRComponent/bin/startTCRserver.sh in a text editor and add the following two lines at line 26 of the script (prior to WebSphere being started): export LD_LIBRARY_PATH=/opt/ibm/db2/V9.7/ lib32:\$LD_LIBRARY_PATH . /home/db2inst1/sqllib/ db2profile
	 Start Tivoli Common Reporting: /opt/IBM/tivoli/ tipv2Components/TCRComponent/bin/ startTCRserver.sh tipadmin tippass
	Another option is to add these environment variables to your .bashrc/.profile so that the variables are set up every time you log into the system.
	Complete the following steps for Tivoli Common Reporting V3.1 or later:
	1. Open the ld.so.conf file.
	 Paste the following path in the ld.so.conf file: /opt/ibm/db2/V10.1/lib32
	3. Open a command line, and run the following command:
	ldconfig
	4. Restart the Dashboard Application Services Hub server.
The reports schema update fails. An SQL error message is displayed in the Report_Installer_For_TCR_Output.txt log file.	The error message indicates that the transaction log file has run out of space. Set the following database parameters to increase the size of the transaction log file: • LOGBUFSIZE: 1024
	• BUFFPAGE: 2000 or 3000
	Number of primary log files: 20
	• Log File size: 8192
	Number of secondary log files: 10

Table 98. Performance and Capacity Management Reports problems

Problem	Solution
You have trouble with the database connection.	The Tivoli Monitoring reports must connect to the Tivoli Data Warehouse to run. The data source, <i>TDW</i> , is defined in Cognos.
	Test the Database Connection to the Tivoli Data Warehouse:
	1. Depending on the database type, make sure the Tivoli Data Warehouse is cataloged locally in the database client (DB2), the local TNS service name is defined in tnsnames.ora (Oracle), or the ODBC data source is created (Microsoft SQL Server). The alias/tns service name/odbc data source name is used in the data source connection.
	2. Test the connection to the data warehouse:
	a. Select Reporting > Common Reporting > Launch > Administration > Configuration > Data Source Connections > TDW.
	b. Click TDW to discover the data source connections (also called TDW).
	c. Click the <i>test</i> icon next to TDW to test the connection.
	3 . If TDW is not defined, manually define the data source in Cognos.
	a. Using the database client, catalog the Tivoli Data Warehouse database.
	b. Log in to the Dashboard Application Services Hub.
	 c. In the navigation tree, select Reporting > Common Reporting.
	 d. Follow the instructions under Configuring database connection (http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=%2Fcom.ibm.tivoli.tcr.doc_211%2Fttcr_config_db.html) in the Tivoli Common Reporting information center to create the data source in Cognos. Ensure that you call this data source "TDW".
You are able to connect to Tivoli Data Warehouse using the database client, but when you run a report, you are asked to enter the database credentials again.	If you get this prompt when you attempt to run a report, the database connection under TDW is not configured properly.
In the Work with reports page, you are prompted to Type a user ID and password and you get the following message: An attempt to connect to the data source failed.	Some common causes might be that the connection is to an invalid DSN, or an incorrect user ID was given, or the database server might be down. Confirm that you created the correct database connection and configure the Tivoli Data Warehouse correctly before running the report again. See the Configuring database connection (http://publib.boulder.ibm.com/infocenter/tivihelp/ v3r1/topic/com.ibm.tivoli.tcr.doc_211/ ttcr_config_db.html) for details.

Table 98. Performance and Capacity Management Reports problems (continued)

Problem	Solution
You open a report, but the report does not populate	Check for one or more of these possible causes:
with data; the report is empty.	• The Tivoli Data Warehouse required historical tables for attributes but does not have the dimension tables. Follow the instructions in Creating shared dimensions tables and populating the time dimensions table (http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/ topic/com.ibm.itm.doc_6.2.3fp1/adminuse/ tcr_reports_dimensionsshared.htm).
	• The historical data in the warehouse is not summarized and pruned. If you know that summarization and pruning is configured and enabled, wait until the process is completed at least one time before opening a report. If you are not sure, see "Configuring historical collection for the Performance and Capacity Management Reports for VMWare" in the <i>IBM Tivoli</i> <i>Monitoring for Virtual Environments Dashboard, Reporting,</i> <i>and Capacity Planning Installation and Configuration</i> <i>Guide</i> .
	• The Tivoli Data Warehouse does not have data. Set summarization and pruning for at least hourly summarization and daily pruning for the required attribute groups Check the "Prerequisite" topic in the Tivoli Common Reporting chapter of your agent user's guide and look for the queries to run to validate the required attribute groups. With Tivoli Monitoring for Virtual Environments V7.1 (and later), Prereq Scanner Reports are provided for these checks. If historical collection is configured and started, in addition to summarization and pruning, and you see seeing missing tables or data for a monitoring agent, contact customer support.

Table 98. Performance and Capacity Management Reports problems (continued)

Problem	Solution
When you run a report, you encounter the following error or similar: RQP-DEF-0177 An error occurred while performing operation 'sqlPrepareWithOptions' status='-201'. UDA-SQL-0196 The table or view "ITMUSER.KVM_SERVER_NETWORK_DV"was not found in the dictionary.	This means that some of the tables or views (or both) are missing in the database, or indicates that the shared dimensions are not created.
	For missing tables or views Check whether historical collection and summarization and pruning are enabled for all the prerequisite attribute groups for the VMware agent. If any of them are missing, enable historical collection and Summarization and Pruning for that particular attribute group. After the historical data collection is configured and historical data is saved, run the report again. Follow the instructions in "Configuring historical data collection for the Performance and Capacity Management Reports for VMWare" in the <i>IBM</i> <i>Tivoli Monitoring for Virtual Environments</i> <i>Dashboard, Reporting, and Capacity Planning</i> <i>Troubleshooting</i> documentation.
	For missing dimensions The shared dimensions are not created and the appropriate database scripts must be run to create shared dimensions and populate them. For details, see Creating shared dimensions tables and populating the time dimensions table (http://pic.dhe.ibm.com/infocenter/tivihelp/ v15r1/topic/com.ibm.itm.doc_6.2.3fp1/ adminuse/tcr_reports_dimensionsshared.htm).
	ITMUSER You can also check whether you are using ITMUSER as your schema name in the Tivoli Data Warehouse. If you are using any schema other than ITMUSER, see the ITMUSER schema entry.
In the VMware VI Balanced and Unbalanced Host Servers report, when you move the mouse pointer over any bar chart, the tooltip displays the Host Server static value in the English language on all the non-English-language interfaces of Tivoli Common Reporting.	No solution is available for this problem.
You choose to view the reports in Portuguese (Brazilian), but the change in locale is not reflected in the report prompt page or the output. You still see English strings instead of Portuguese (Brazilian).	For this release, when you choose to view the reports in Portuguese (Brazilian) with Tivoli Common Reporting 3.1, the text is displayed in English. This problem is a known problem with Cognos 10.2. However, the reports can be viewed in Portuguese (Brazilian) using Tivoli Common Reporting 2.1.1.

Table 98. Performance and Capacity Management Reports problems (continued)

Problem	Solution
 In the following reports, the Month column in the report table might show corrupted values on the non-English-language interfaces of Tivoli Common Reporting: VMware VI Cluster Weekly Comparison VMware VI Host Servers Weekly Comparison VMware VI Top or Bottom Host Servers Weekly Comparison 	 This problem occurs because of garbage values that are inserted in the MONTH_LOOKUP table of Tivoli Data Warehouse database by the reports installer. To resolve this problem, run the database scripts manually to populate the MONTH_LOOKUP table with the correct values. For more information about running the scripts manually, see "Creating shared dimensions tables and populating the time dimensions table" in the <i>IBM Tivoli Monitoring Administrator's Guide V6.2.3 Fix Pack 1.</i>
 In the following reports, the Month column in the report table shows values in English language on the Simplified Chinese, Traditional Chinese, German, Korean, and Portuguese (Brazil) interfaces of Tivoli Common Reporting: VMware VI Cluster Weekly Comparison VMware VI Host Servers Weekly Comparison VMware VI Top or Bottom Host Servers Weekly Comparison 	No solution is available for this problem.

Table 98. Performance and Capacity Management Reports problems (continued)

Dashboard for VMware problems

Use the Dashboard for VMware problems table to troubleshoot errors you might encounter when you are using the Dashboard for VMware.

Note: The Dashboard Health Checks page provides information on the health of the data connections in your system. See "Dashboard Health Checks for VMWare" in the *IBM Tivoli Monitoring for Virtual Environments Dashboard, Reporting, and Capacity Planning User's Guide.*

Table 99. Dashboard for VMware problems

Problem	Solution
Data does not populate the chart or table, does not refresh, or the chart or table only partially displays.	 Log out of Dashboard Application Services Hub. Clear the browser cache. (For example, in the Firefox browser, click Tools > Options > Advanced > Network.)
	3. In the Cached Web Content section, click Clear Now .
	4. Recycle the browser by restarting the browser.
	5. Log on to Dashboard Application Services Hub.

Table 99. Dashboard for VMware problems (continued)

Problem	Solution
You are getting an error message	This message indicates that the micro broker service that communicates asynchronous updates from the server to
ATKRST123E Java Message Service initialization error	the tables in the browser client is not running or unable to start. A common cause of this situation is the port that is being used by another process. The port number that is
What is wrong?	required is shown in the URL in the message. In this example, it is the default port 16324. This problem can
The message that is mentioned in the following example is displayed in the Charts:	sometimes be resolved by restarting the Dashboard Application Services Hub server. If that does not solve the problem investigate what other processes on the host
ATKRST123 ATKRST123 Java Message Service	might be using the port.
initialization error. The error message is as follows: 'Create connection to local://ibm.tivoli.rest_16324 failed: Unexpected exception.'.	The WebSphere Appliance Management Center product published another solution to this problem that you can try if one of the mentioned solutions does not resolve the issue. The alternative solution is described in the
or ATKRST123 Java Message Service initialization error	WebSphere Appliance Management Center Information Center topic, Fixing the display problem with the Device window, Firmware window and other windows
Could not connect to broker URL: tcp:// <hostname>:16324. Reason:java.net.ConnectException:Connection refused: connect.</hostname>	(http://publib.boulder.ibm.com/infocenter/wamcinfo/ v4r0m0/topic/com.ibm.wamc.doc/ tr_widgets_display_issue.html).

Problem	Solution
On some of the Dashboard for VMware pages, the Internet Explorer 8 displays the following message: Do you want to view only the web page content that was delivered securely?	The browser displays this message to alert you that some of the content on the page is being delivered by a secure (https) connection and some content is being delivered by a nonsecure (http) connection. ForDashboard for VMware, this situation occurs because the Tivoli Application Dependency Discovery Manager (TADDM) server is accessed by using the http protocol and the content hosted on the local Dashboard Application Services Hub is accessed by using https. You can respond to the prompt in several ways:
	• Respond "Yes" to the message prompt to suppress the nonsecure content from the Tivoli Application Dependency Discovery Manager server for the remainder of the current browser session.
	• Respond "No" to the message prompt. This response allows the nonsecure content for the remainder of the current browser session.
	 Change your settings to suppress the prompt. If you want to suppress this prompt, you can edit the Security Settings for the zone that applies to the Dashboard Application Services Hub and Tivoli Application Dependency Discovery Manager servers. If you set Enable mixed content to true, this setting allows http and https content to be mixed on a single page without prompting. This setting increases the exposure to certain security issues, therefore, use the setting with caution in zones that apply to untrusted sites. Conversely, setting Enable mixed content to false always suppresses the http content on a secure page. To change the setting in the Internet Explorer 8, select Tools > Internet Options > Security, and then click Custom level. Scroll down to Display mixed content, and select Disable or Enable, depending on your preference. Click OK2 times to save your changes.

Table 99. Dashboard for VMware problems (continued)

Problem	Solution
On some of the Dashboard for VMware pages, the Internet Explorer 9 displays the following message: Only secure content is displayed.	The browser displays this message to alert you that some of the content on the page is being delivered by a secure (https) connection, and some content is being delivered by a nonsecure (http) connection. For Dashboard for VMware, this situation occurs because the Tivoli Application Dependency Discovery Manager server is accessed by using the http protocol and the content hosted on the local Dashboard Application Services Hub is accessed by using https. You can respond to the prompt:
	• Respond "Show all content" to the message prompt, which allows the nonsecure content for the remainder of the current browser session.
	• Close the message prompt without any action to suppress the nonsecure content from the Tivoli Application Dependency Discovery Manager server for the remainder of the current browser session.
	 Change your settings to suppress the prompt. If you want to suppress this prompt, you can edit the Security Settings for the zone that applies to the Dashboard Application Services Hub and Tivoli Application Dependency Discovery Manager servers. If you set Enable mixed content to true, this setting always allows http and https content to be mixed on a single page without prompting. This setting increases the exposure to certain security issues, therefore, use the setting with caution in zones that apply to untrusted sites. Conversely, setting Enable mixed content to false always suppresses the http content on a secure page. To change the setting in the Internet Explorer 9, select Tools > Internet Options > Security, and click Custom level. Scroll down to Display mixed content and select Disable or Enable, depending on your preference. Click OK 2 times to save your changes.

Table 99. Dashboard for VMware problems (continued)

Problem	Solution		
A situation is not displayed in the Situation Events list on the Dashboard for VMware, V7.2.	Following are the prerequisites to view a situation in the "Situation Events" list on the Dashboard for VMware, V7.2:		
	• The situation must sample a VMware attribute group that contains NodeID attribute.		
	• The situation must be associated with a Navigator item.		
	The following is the list of VMware attribute groups that contain a NodeID attribute:		
	KVM_DATACENTERS		
	KVM_DATASTORES		
	KVM_VM_DATASTORE_UTILIZATION		
	KVM_DATASTORE_HOST_DISKS		
	KVM_CLUSTERS		
	KVM_CLUSTERED_SERVERS		
	KVM_CLUSTERED_DATASTORES		
	KVM_CLUSTERED_RESOURCE_POOLS		
	KVM_CLUSTERED_VIRTUAL_APPS		
	KVM_CLUSTERED_VIRTUAL_MACHINES		
	KVM_SERVER		
	KVM_SERVER_DATASTORE		
	• KVM_SERVER_HEALTH		
	• KVM_SERVER_HBA		
	• KVM_SERVER_CPU		
	• KVM_VM_CPU		
	KVM_SERVER_DISK		
	• KVM_SERVER_SAN		
	• KVM_VM_DISK		
	KVM_VM_PARTITION		
	KVM_SERVER_VM_DATASTORE_UTILIZATION		
	KVM_SERVER_MEMORY		
	KVM_VM_MEMORY		
	KVM_SERVER_NETWORK		
	• KVM_VM_NETWORK		
	• KVM_RESOURCE_POOL_GENERAL		
	• KVM_RESOURCE_POOL_CPU		
	• KVM_RESOURCE_POOL_MEMORY		
	KVM_VIRTUAL_MACHINES		

Table 99. Dashboard for VMware problems (continued)

Problem	Solution		
Continued from previous row.	The following is the list of the VMware VI agent Situations that are provided by the product:		
	KVM_Cluster_CPU_Util_High		
	KVM_Cluster_Effective_Svrs_Low		
	KVM_Cluster_Effective_CPU_Low		
	KVM_Cluster_Effective_Mem_Low		
	KVM_Cluster_Memory_Util_High		
	KVM_Datastore_Bad_Status		
	KVM_Datastore_Inaccessible		
	KVM_Datastore_Usage_High		
	KVM_Server_CPU_Util_High		
	KVM_Server_Datastore_Free_Low		
	 KVM_Server_Disk_Reads_High 		
	KVM_Server_Disk_Writes_High		
	KVM_Server_HBA_Fault		
	KVM_Server_Memory_Util_High		
	KVM_Server_NIC_Down		
	KVM_Server_Receive_Rate_High		
	KVM_Server_Transmit_Rate_High		
	KVM_VM_CPU_Ready_High		
	KVM_VM_CPU_Util_High		
	KVM_VM_Disk_Free_Low		
	 KVM_VM_Guest_Memory_Util_High 		
	KVM_VM_Host_Memory_Util_High		
	KVM_VM_Powered_Off		
	KVM_VM_Receive_Rate_High		
	KVM_VM_Transmit_Rate_High		
The Virtual Machine details page displays the Guest OS link. When you click the Guest OS link, the link does not open the respective Server Dashboards page.	To resolve the problem, you must install the Server Dashboards.		
When you click the Guest OS link, the page loading icon is displayed for a longer period to open the Server	If the performance of the computer is slow, the page loading icon is displayed.		
	To view the Server Dashboards page contents, either wait for a few more minutes or close the Server Dashboards page, and again click the Guest OS link.		
In the Select Time Range window, after you select the Set Time Range option, sometimes the slider does not work correctly and the slider changes the time in an hourly interval only.	When the time range slider does not work, set the time by selecting a required option from the Time list.		
For example, 11:38AM, 4:38PM, 9:38PM, and so on.			
In addition, the slider does not display the time between the time range that is displayed on the slider.			

Table 99. Dashboard for VMware problems (continued)

Table 99. Dashboard for VMware problems (continued)

Problem	Solution
On the Dashboard Health Checks page, the components	To expand the component, press Ctrl+Right Arrow Key.
tables do not specify any instructions for expanding or	To collapse the component, press Ctrl+Left Arrow Key.
collapsing the component. This is not possible for the	
blind users to know that they can expand or collapse the	
component.	

Capacity Planner problems

Consult the Capacity Planner problems table to troubleshoot errors you might encounter while you use the Capacity Planner for VMware and Capacity Planner for PowerVM.

Enabling debugging information in the Capacity Planner logs

For the Load Configuration logs, complete the following steps:

- Update the .properties file that is saved to the following path: For Capacity Planner for PowerVM: Jazz_SM_Home\profile\installedApps\JazzSMNode01Cell\ isc.ear\AnalyticsWebUIPower.war\WEB-INF\classes\SYSTEMPLOADER_log4j.properties For Capacity Planner for VMware: Jazz_SM_Home\profile\installedApps\JazzSMNode01Cell\isc.ear\ AnalyticsWebUI.war\WEB-INF\classes\DMLOADER_log4j.properties
- 2. Set the SystemPLoader to DEBUG as shown in the following example: For Capacity Planner for PowerVM: log4j.logger.SystemPLoader=DEBUG,SystemPLoader_FILE,SystemPLoader_CONSOLE For Capacity Planner for VMware: log4j.logger.DMLoader=DEBUG,DMLoader FILE,DMLoader CONSOLE
- 3. Restart the Jazz for Service Management server.

For other server logs, complete the following steps:

- Update the .properties file that is saved to the following path: For Capacity Planner for PowerVM: Jazz_SM_Home\profile\installedApps\JazzSMNode01Cell\ isc.ear\AnalyticsWebUIPower.war\WEB-INF\classes\log4j.properties
 For Capacity Planner for VMware: Jazz_SM_Home\profile\installedApps\JazzSMNode01Cell\ isc.ear\AnalyticsWebUI.war\WEB-INF\classes\log4j.properties
- Set Analytics to DEBUG as shown in the following example: log4j.logger.Analytics=DEBUG,ANALYTICS_FILE
- 3. Restart the Jazz for Service Management server.

Problem	Solution	
During Capacity Planner workflow, the following error is displayed in the Analytics.log file: For DB2 exception SQLCODE=-805 SQLSTATE=51002.	 Log in as a DB2 administrator. Run the following command: db2inst1@tfam02:/opt/ ibm/db2/V9.7/bnd]db2 bind /opt/ibm/db2/V9.7/bnd/ @db2cli.lst CLIPKG 30 Note: If DB2 is installed in a different location, change the path accordingly. A success message similar to the following message is displayed after completion of the command: LINE MESSAGES FOR db2cli.lst SQL0061W The binder is in progress. SQL0091N Binding was ended with "0" errors and "0" warnings. db2inst1@tfam02:/opt/ibm/db2/V9.7/bnd] 	
The DBinstaller is looking for bash and failing on AIX.	Bash is a prerequisite on AIX, but is not available by default on AIX. You can download it from IBM AIX Toolbox download (http://www-03.ibm.com/systems/ power/software/aix/linux/toolbox/download.html). See also IBM developerWorks System Administration Toolkit: Get the most out of bash (http://www.ibm.com/ developerworks/aix/library/au-satbash.html).	
The Capacity Planner fails randomly in loading data for the configuration objects (Physical Server, VM, and Datastores) from the Tivoli Enterprise Portal Server.	In rare cases, because of environmental issues, Capacity Planner data loaders might face problems in sourcing data from the Tivoli Enterprise Portal Server. Click Load Config to rerun the data load.	
In the Edit Current Environment page, you click the Load knowledge data icon and load the server model catalog. In the Physical Servers view, you then click Actions > Add server to add a server. Newly added physical server models are not visible in the list of available models.	Log out of the Capacity Planner user interface and log back in. The newly added physical server models are displayed in the list of available models.	
In the Utilization view for Capacity Planner for VMware or Capacity Planner for PowerVM, when you clear a selected check box for the available virtual machines (for VMware) or logical partitions (for PowerVM), the JAWS Screen Reading Software cannot read out this action of clearing the check box.	No solution is available for this problem.	

Table 100. General Capacity Planner problems and solutions

Table 101. Capacity Planner for	or PowerVM problems	and solutions
---------------------------------	---------------------	---------------

Problem	Solution
After you complete the Load Config procedure, you receive the following warning message: ANA000000W : Load Config is successful. However, one or more of the mandatory parameters could not be loaded. Check SystemPLoader.log for more details.	 For more information about the error message, check the log file that is saved to the following path: <jazz-sm_home>\profile\logs\SystemPLoader.log.</jazz-sm_home> Note: It is possible that certain Managed Systems or LPARs are not loaded to the Capacity Planner database.
ANA00001I : Data initialized	If you see the following message in the log:
successfully.	Skipping The managed system with hostname
	This message means that the Managed System is not in an operating state but that it is in an offline, or not connected state.
	If you see the following message in the log:
	Skipped Lpar Details are as below
	This message means that the LPARs might be hosted on an IVM based managed system that can not be loaded.
After you complete the Load Config procedure, you receive the following error: ANA00135E: SystemPLoader failed. ADL00032E: Failed to get response from URI http://hostname:port/ ibm/tivoli/rest/providers? param_providerVersion=06.30.00.00	 Check the log file that is saved to the following path:<jazz_sm_home>\profile\logs\ SystemPLoader.log for detailed error messages</jazz_sm_home> Check the network connection between the ITM TEPS data provider server and port and the Jazz SM server.
You complete the step to resize the LPAR in the Capacity Planner for PowerVM tool. Then, when you run the Capacity Planner PowerVM Recommended Environment Report, the following pop up message displays in Internet Explorer 9: "IE has blocked cross-site scripting".	 Add the Cognos URL to the Trusted Sites list. Modify the options of the Trusted Sites zone and set the "Enable XSS Filter" property to "Disable". Restart the web browser. Run the Capacity Planner PowerVM Recommended Environment report.
	For more information, see the following: Technote and Troubleshooting steps.
A Red Cross error icon displays in the Status column on the Edit Current Environment → Inventory Managed Systems View. This error occurs when the CPU Speed for a specific	 Configure a Tivoli Monitoring OS monitoring agent on an LPAR on the managed system and refresh the configuration by completing the Load Config procedure. Update the CPU Speed column with the correct value
Managed System is not discovered. Also, the error occurs when the LPARs on the Managed System are not monitored by using IBM Tivoli Monitoring Unix or Linux OS agents.	as per the Managed System model. Note: To determine the correct value, search for the model in the Add Managed System UI.
When you run the Analyze LPARs or Compute usage action, values are not populated for CPU, Memory, Network, or Storage parameters for all or some of the LPARs.This error occurs when there is not any Performance or utilization data available in Tivoli Monitoring Warehouse for the LPARs.	 Check if federation configuration is performed as part of the Post-install configuration of the capacity planner. Ensure that the correct Historic data collection is configured on Tivoli Monitoring as per the steps shown in the <i>Configuring historical data collection for</i> <i>Capacity Planner for PowerVM</i> section of this guide. Update the Analysis time period in step 2 of the Planning Center wizard to the time period when performance data is available.

Problem	Solution
When you click Generate plan, you see the following error on the PowerVM LPAR Sizing page: ANA00147E : Data Validation Error SERVER:: <server-name> LPAR::<lpar-name> ANA00143E : Total CPU Utilization value 0 is invalid. The value must be greater than or equal to 0.01% ANA00146E : Memory value 0 is invalid. The value must be greater than or equal to 1.0</lpar-name></server-name>	 Run the Analyze LPARs step to compute the usage data for CPU, Memory, Network and Storage bandwidths for the selected LPARs on the Managed System. If the performance metrics of the LPARs are not available, you can navigate to Edit Current Environment > LPAR Utilization view and add usage data manually by using Actions > Edit Usage.
 When you click Generate Plan, you see the following error on the PowerVM LPAR Sizing page: ANA00147E : Data Validation Error: SERVER::<server-name>ANA00160E : Total CPU Entitlement for all LPAR and VIO Partition cannot exceed activated Cores on the Managed System.</server-name> ANA00165E : Total Memory Allocation for all LPAR and VIO Partition cannot exceed Total Memory Installed on the Managed System. Note: This error occurs when configured entitlements or allocations for CPU or Memory for LPARs errorially 	 Open the Edit current environment > Inventory View for Managed System. Check the Total Activated CPU cores and the Total Memory Installed for the selected Managed System. The Allocation of Total CPU entitlement for all LPARs or VIOs must not be greater than the Total Activated CPU cores of the selected Managed System. The Allocation of Total Memory Assigned for all LPARs or VIOs must not be greater than the Total Memory Installed for the selected Managed System.
allocations for CPU or Memory for LPARs, especially fictitious LPARs, exceed the resources that are available on the Managed System.	

Table 101. Capacity Planner for PowerVM problems and solutions (continued)

Table 102.	Capacity	Planner for	VMware	problems	and	solutions
------------	----------	-------------	--------	----------	-----	-----------

Problem	Solution
For Physical Servers, the status is displayed as red in the Status column.	 Create the CUSTOMER_USER_DEFINED_BENCHMARK.csv file with the following content:
	,Manufacturer,,,Manufacturer Model (CPU Architecture),,Processor chips,Processor cores,CPU Benchmark Value,,,,,,Model (CPU Architecture)
	Note: Ensure that the values for Architecture, Number of Core CPU, Model, and Manufacturer match with the values for the Server for which the status is displayed as red.
	 After you have created the CUSTOMER_USER_DEFINED_BENCHMARK.csv file, upload the file from the Knowledge Base.
You choose to view the reports in Portuguese (Brazilian), but the change in locale is not reflected in the report prompt page or the output. You still see English strings instead of Portuguese (Brazilian).	For this release, when you choose to view the reports in Portuguese (Brazilian) with Tivoli Common Reporting 3.1, the text is displayed in English. This problem is a known problem with Cognos 10.2. However, the reports can be viewed in Portuguese (Brazilian) using Tivoli Common Reporting 2.1.1.
When you open the Current Environment Report in the PDF format, the bar charts and table for the report in the PDF are not clearly visible.	If you open the report in a PDF file, zoom in the PDF to 300% - 400% to view the bar charts and table clearly in the PDF.

Table 102. Capacity Planner for VMware problems and solutions (continued)

Problem	Solution
In the Capacity Planner Optimized Environment Plan report, the bar charts display the value for every second bar when there is large number of values on the ordinal axis.	To see the missing value for each bar, move the mouse pointer over the bar that is related to the ordinal axis. The value for the ordinal axis is displayed in a tooltip.
In the VMware Expense Reduction Report, a table is displayed on the upper-left corner of the window. The table summarizes the user inputs that are entered on prompt page for enterprise-wide number of clusters, servers, processor cores, and so on. In the table some of the text overlaps in the PDF format, while the text is displayed properly in the HTML format.	View the VMware Expense Reduction Report in the HTML format.
<pre>When you load the configuration data, the DMloader.log file displays the following error message: Failed to get response from URI: http:// 10.44.184.157:15210/ibm/tivoli/rest//providers/ itm.HUB_IBMESX7V5/datasources/TMSAgent. %25IBM.STATIC016/datasets/MetricGroup.KVMSERVERG/ items?properties=all&param_SourceToken=*VMWARE_VI &param_refId=lynx_1374225142804 &param_providerVersion=06.30.00.00 org.apache.wink.client.ClientRuntimeException: java.lang.RuntimeException: java.net.SocketTimeoutException: Read timed out</pre>	 To resolve this problem, increase the value for the following parameters in the dmloader.cfg file: ITMDP.READ_TIMEOUT ITMDP.READ_TIMEOUT The read timeout (in seconds). If you have slow responding agents, you can increase the value of this parameter. The default value is 180 seconds. ITMDP.CONNECT_TIMEOUT The connection timeout (in seconds). The default

380 IBM Tivoli Monitoring for Virtual Environments: Dashboard, Reporting, and Capacity Planning: User's Guide

Appendix A. CSV format for data import and export

Data to be imported and exported must be kept in CSV format.

Table 103. CFG_PHYSICAL_SERVER for PowerVM

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
PHYSICAL_ SERVER_PK	BIGINT		Primary key
SERVER_POOL _NAME	VARCHAR	256	Currently not used
DATA_ CENTER_ NAME	VARCHAR	256	Data center to which physical server belongs, default value "Default"
FQDN	VARCHAR	256	Currently not used
HOST_NAME	VARCHAR	256	Host name of Managed System
IP_ADDRESS	VARCHAR	256	IP address of Managed System
MODEL	VARCHAR	256	Model name, for example, POWER [®] 570, Power 750
MODEL NAME	VARCHAR	256	Human_Desired_Model for example 750-8233-E8B
MANUFACTURER	VARCHAR	256	Name of manufacturing company, for example, IBM
ARCHITECTURE	VARCHAR	256	Currently not used
OPERATING_ SYSTEM_NAME	VARCHAR	64	Operating system name, for example, Power Linux or AIX
OPERATING_ SYSTEM_ VERSION	VARCHAR	32	Operating system version, for example, 6.1 or 7.1
CPU_SPEED	BIGINT		Processor speed in MHZ
CPU_CORES_INSTALLED	BIGINT		Total CPU Installed
NUM_CPUS	INTEGER		Number of cores Enabled
MEMORY_LIMIT	BIGINT		Physical memory capacity in MB
MEMORY_ INSTALLED	BIGINT		Currently not used
STATUS	VARCHAR	256	Do not edit this column
HYPERVISOR_ VERSION	VARCHAR	32	Hypervisor version, Currently not used
HYPERVISOR_ TYPE	VARCHAR	64	Hypervisor Name, for example, Power
SERIAL_NUMBER	VARCHAR	128	Serial number of Managed System
DATA_SOURCE	VARCHAR	64	Do not edit this column
HMC_HOST_NAME	VARCHAR	256	HOST NAME OF THE HMC SERVER
TAG TYPES			
TAG TYPE	DATA TYPE	STRING LENGTH	DETAILS
CUST_PRIMARY_ BU	VARCHAR	50	Business unit with most significant dependency on this server
CUST_PRIMARY_ BUSINESS_ APP	VARCHAR	50	Business application with most significant dependency on this server
SERVER_ FUNCTION	VARCHAR	50	Application, database, print server, or web server
ENVIRONMENT	VARCHAR	50	Production, development, test, data management zone, or eminent security zone
CRITICALITY_ SLA	VARCHAR	50	Business criticality or service level agreement info
CITY_CENTER	VARCHAR	50	If multiple locations, city or data center name location of server or host
WORKINGSET	VARCHAR	50	Flag indicating physical server is part of working set

Instructions:

· Do not add any value to the MODEL_NAME column when you import the CSV for any new managed system.

• Do not alter the content of the first column of the csv file, namely PHYSICAL_SERVER_PK.

• Update the values in the rest of the columns only if the existing values are null or empty.

• To add new servers, keep the values in the PHYSICAL_SERVER_PK column blank and add values in the rest of the columns.

· The length of the input strings must not exceed the string length as shown.

• You can add tags as new columns in the table.

• You can add multiple tags of the same tag type with a ; separator. For example, you can add two tag names, PostPaid and PrePaid, of the same tag type, CUST_PRIMARY_BU, to a server by combining the tag names as PostPaid; PrePaid.

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
PHYSICAL SERVER PK	BIGINT		Primary key
SERVER_POOL_NAME	VARCHAR	256	Name of server pool if a hierarchy exists
DATA_ CENTER_ NAME	VARCHAR	256	Data center to which physical server belongs
FQDN	VARCHAR	256	Fully qualified host name
HOST_NAME	VARCHAR	256	Host name of physical server
IP_ADDRESS	VARCHAR	256	IP address of physical server
MODEL	VARCHAR	256	Model name, for example, SUN-FIRE-V240 or POWER 570
MANUFACTURER	VARCHAR	256	Name of manufacturing company, for example, IBM, HP, or DELL
ARCHITECTURE	VARCHAR	256	System Architecture, for example, SUN4 or POWER
OPERATING_ SYSTEM_NAME	VARCHAR	64	Operating system name, for example, Linux or Windows
OPERATING_ SYSTEM_ VERSION	VARCHAR	32	Operating system version
CPU_SPEED	BIGINT		Processor speed in MHZ
NUM_CPUS	INTEGER		Number of cores
MEMORY_LIMIT	BIGINT		Physical memory capacity in MB
MEMORY_ INSTALLED	BIGINT		Currently not used
STATUS	VARCHAR	256	Do not edit this column
HYPERVISOR_ VERSION	VARCHAR	32	Hypervisor version, for example, 3.5 U3 for VMware
HYPERVISOR_ TYPE	VARCHAR	64	Hypervisor Name, for example, pHyp or VMware ESX
BIOS_DATE	DATE		BIOS date of server
DATA_SOURCE	VARCHAR	64	Do not edit this column
TAG TYPES			
TAG TYPE	DATA TYPE	STRING LENGTH	DETAILS
CUST_PRIMARY_ BU	VARCHAR	50	Business unit with most significant dependency on this server
CUST_PRIMARY_ BUSINESS_ APP	VARCHAR	50	Business application with most significant dependency on this server
SERVER_ FUNCTION	VARCHAR	50	Application, database, print server, or web server
ENVIRONMENT	VARCHAR	50	Production, development, test, data management zone, or eminent security zone
CRITICALITY_ SLA	VARCHAR	50	Business criticality or service level agreement info
CITY_CENTER	VARCHAR	50	If multiple locations, city or data center name location of server or host
WORKINGSET	VARCHAR	50	Flag indicating physical server is part of working set
TYPE	VARCHAR	50	Candidate type indicating whether a physical server is Source or Target

Table 104. CFG_PHYSICAL_SERVER for VMware

Instructions:

- Do not alter the content of the first column of the csv file, namely <code>PHYSICAL_SERVER_PK</code>.

• Update the values in the rest of the columns only if the existing values are null or empty.

• To add new servers, keep the values in the PHYSICAL_SERVER_PK column blank and add values in the rest of the columns.

• The length of the input strings must not exceed the string length as shown.

• You can add tags as new columns in the table.

• You can add multiple tags of the same tag type with a ; separator. For example, you can add two tag names, PostPaid and PrePaid, of the same tag type, CUST_PRIMARY_BU, to a server by combining the tag names as PostPaid; PrePaid.

Table 105. CFG_	VIRTUAL_	MACHIN	IE for Powe	rVM

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
VIRTUAL_ MACHINE_PK	BIGINT		Primary key
SHARED_MODE	VARCHAR	32	Logical Partition CPU sharing Mode
CAPPED_MODE	VARCHAR	32	Logical Partition CPU Capping Mode
PS_HOST_ NAME	VARCHAR	256	Host name of Managed System

Table 105. CFG	_VIRTUAL	MACHINE	for PowerVM	(continued)
----------------	----------	---------	-------------	-------------

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
VMID	VARCHAR	256	Currently Not Used
HOST_NAME	VARCHAR	256	Host name of Logical Partition
IP_ADDRESS	VARCHAR	256	IP address of Logical Partition
FQDN	VARCHAR	256	Currently Not Used
MEMORY_ SIZE	BIGINT		Memory currently allcoated to LPAR
OPERATING_ SYSTEM_NAME	VARCHAR	256	Name of operating system
OPERATING_ SYSTEM_ VERSION	VARCHAR	256	Operating system version
MIDDLEWARE _NAME	VARCHAR	256	Middleware name
MIDDLEWARE _VERSION	VARCHAR	256	Middleware version
DATA_SOURCE	VARCHAR	64	Source of configuration data of virtual machine. Do not edit this information.
PHYSICAL_ SERVER_PK	BIGINT		Physical server primary key
VM_NAME	VARCHAR	256	Display name of Logical Partition
WORKINGSET	VARCHAR	50	Flag indicating that virtual machine is part of working set
VIOS_STORAGE_NAME y STORAGE_PROTECTION VARCHAR 128	VARCHAR	256	VIO Partition name providing storage to the LPAR
VIOS_NETWORK_NAME	VARCHAR	256	VIO Partition name providing network to the LPAR
VIOS_STORAGE_VIRT_MC_PK	BIGINT		VIO Partition for Storage primary key
VIOS_NETWORK_VIRT_MC_PK	BIGINT		VIO Partition for Network primary key
STORAGE_ASSIGNED DOUBLE			Storage space assigned in MB to LPAR
STORAGE_CONSUMED DOUBLE			Storage space consumed in MB to LPAR
STORAGE_PROTECTION	VARCHAR	128	Currenly Not used
CITY_CENTER	VARCHAR	50	If multiple locations, city or data center name location of server or host
CRITICALITY_SLA	VARCHAR	50	Business criticality or service level agreement info
CUST_PRIMARY_BU	VARCHAR	50	Business unit with most significant dependency on this server
CUST_PRIMARY_BUSINESS_ APP	VARCHAR	50	Business application with most significant dependency on this server
ENVIRONMENT	VARCHAR	50	Production, development, test, data management zone, or eminent security zone
SERVER_FUNCTION	VARCHAR	50	Application, database, print server, or web server

Instructions:

• Do not alter the content of the first column of the csv file, namely VIRTUAL_MACHINE_PK.

• Update the values in the rest of the columns only if the existing values are null or empty.

• To add new servers, keep the values in the VIRTUAL_MACHINE_PK column blank and add values in the rest of the columns.

• The length of the input strings must not exceed the string length as shown.

1. When you import any LPAR, if VIO is not added on the Managed System in the database, you must not use any VIO name or VIOS_STORAGE_VIRT_MC_PK / VIOS_NETWORK_VIRT_MC_PK

2. You can import VIO with the Logical Partition import or export.

3. To import VIO, set the OS Type to VIOS.

4. Before you complete the association of the VIO and the LPAR, ensure that the VIO Partition details imported successfully.

Table 106. CFG_VIRTUAL_MACHINE for VMware

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
VIRTUAL_ MACHINE_PK	BIGINT		Primary key
PS_HOST_ NAME	VARCHAR	256	Host name of physical server
VMID	VARCHAR	256	ID of the virtual machines
HOST_NAME	VARCHAR	256	Host name of virtual machine
IP_ADDRESS	VARCHAR	256	IP address of virtual machine
FQDN	VARCHAR	256	Fully qualified domain name
CPU_MIN	BIGINT		CPU reservation

Table 106. (CFG_V	IRTUAL	MACHINE	for \	/Mware	(continued)
--------------	-------	--------	---------	-------	--------	-------------

COLUMN NAME	DATA TYPE	COLUMN LENGTH	DETAILS
CPU_MAX	BIGINT		CPU limit
CPU_DESIRED	BIGINT		Currently not used
MEMORY_MIN	BIGINT		Memory reservation
MEMORY_MAX	BIGINT		Currently not used
MEMORY_ DESIRED	BIGINT		Currently not used
OPERATING_ SYSTEM_NAME	VARCHAR	256	Name of operating system
OPERATING_ SYSTEM_ VERSION	VARCHAR	256	Operating system version
MIDDLEWARE _NAME	VARCHAR	256	Middleware name
MIDDLEWARE _VERSION	VARCHAR	256	Middleware version
DATA_SOURCE	VARCHAR	64	Source of configuration data of virtual machine. Do not edit this information.
PHYSICAL_ SERVER_PK	BIGINT		Physical server primary key
VM_NAME	VARCHAR	256	Display name of virtual machine
WORKINGSET	VARCHAR	50	Flag indicating that virtual machine is part of working set
CITY_CENTER	VARCHAR	50	If multiple locations, city or data center name location of server or host
CRITICALITY_SLA	VARCHAR	50	Business criticality or service level agreement info
CUST_PRIMARY_BU	VARCHAR	50	Business unit with most significant dependency on this server
CUST_PRIMARY_BUSINESS_ APP	VARCHAR	50	Business application with most significant dependency on this server
ENVIRONMENT	VARCHAR	50	Production, development, test, data management zone, or eminent security zone
SERVER_FUNCTION	VARCHAR	50	Application, database, print server, or web server

Instructions:

• Do not alter the content of the first column of the csv file, namely VIRTUAL_MACHINE_PK.

• Update the values in the rest of the columns only if the existing values are null or empty.

• To add new servers, keep the values in the VIRTUAL_MACHINE_PK column blank and add values in the rest of the columns.

• The length of the input strings must not exceed the string length as shown.
Appendix B. Editing knowledge base

Knowledge base data consists of rules, the server model catalog, user-defined benchmark values, and virtualization overhead values. This data that is loaded into the Capacity Planner can be adjusted as the environment requires.

The sample files that are loaded by default during installation are available in CSV format in the following directories:

- For Tivoli Integrated Portal V2.1: TIP_HOME/ITMFVSDash/dbinstaller/samples/
- For Dashboard Application Services Hub:
 - Windows systems: C:\Program Files\IBM\JazzSM\installedDashboards\com.ibm.tivoli.cpdash\ AnalyticsDatabaseInstaller\samples
 - Operating systems other than Windows: /opt/IBM/JazzSM/installedDashboards/ com.ibm.tivoli.cpdash/AnalyticsDatabaseInstaller/samples

You can edit these files and update them in the tool using the Refresh Knowledge Base Content window, shown in Figure 130 on page 294.

Rules See "Importing rules in Capacity Planner" on page 293.

Server Model Catalog

This file contains the server models that are listed in the Add Physical Server window. If you must make changes to the models available for adding new servers, this file can be modified. Use the file format in Table 107.

Column name	Maximum column length	Details
SERVER_MODEL_CATALOG_PK	19	This column must be blank.
MODEL_DESCRIPTION	50	Description of the model. This column can be blank.
VENDOR	20	Model vendor. This column can be blank.
MODEL	50	Model name, for example, System X3550.
PROCESSOR	50	Processor Family, for example, Intel Xeon E5335.
PROCESSOR_SPEED_MHZ	10	Processor speed in MHz.
PROCESSOR_CHIPS	10	Number of processor chips. This column can be blank.
PROCESSOR_CORES	10	Number of processor cores.

Table 107. Server Model Catalog

The following entry is a sample:

,,IBM,System X3200,Intel Xeon X3210,,,4

USER_DEFINED_BENCHMARK.csv

This file contains the benchmark values for the server models. These values are used to match comparative benchmarks for the server models and the result of the matching is displayed in the **Status** column in the Physical Server Inventory view. Use the file format in Table 108 on page 386.

Column name	Maximum column length	Details
USER_DEFINED_BENCHMARK_PK	19	This column must be blank.
VENDOR	20	Server vendor. This column can be blank.
SERVER_DESCRIPTION	50	Server description. This column can be blank.
FAMILY	20	Server family. This column can be blank.
PROCESSOR	128	System name with Processor architecture, for example, IBM System x3550 M3 (Intel Xeon X5650).
PROCESSOR_SPEED_MHZ	10 (Integer)	Processor speed. This column can be blank.
PROCESSOR_CHIPS	10 (Integer)	Number of processor chips.
PROCESSOR_CORES	10 (Integer)	Number of processor cores
CPU_BENCHMARK_CAPACITY	19 (Integer)	Benchmark value. The value can be 0 - 100000.
CPU_OLTP_BENCHMARK_CAPACITY	19	Currently not used.
0S	50	Operating system. This column can be blank.
DATE	10	Currently not used. This column must be blank.
UVALUES	10 (Integer)	Currently not used. This column must be blank.
HEIGHT_MM	10	Currently not used. This column must be blank.
WIDTH_MM	10	Currently not used. This column must be blank.
DEPTH_MM	10	Currently not used. This column must be blank.
MAX_HEAT_BTU	19	Currently not used. This column must be blank.
MAX_POWER_WATTS	19	Currently not used. This column must be blank.
NOTES	50	Currently not used. This column must be blank.
MODEL	50	Currently not used. This column must be blank.

The following entry is a sample:

,,,,IBM Corporation IBM BladeCenter HS21 XM (Intel Xeon E5430),,1,8,21200

, , , , , , , , , , , ,

CUSTOM_USER_DEFINED_BENCHMARK.csv

You can use the CUSTOM_USER_DEFINED_BENCHMARK.csv file to supplement the benchmark data of the Capacity Planner with your environment specific physical server models and the benchmark data. The data loaded using the CUSTOM_USER_DEFINED_BENCHMARK.csv file is searched on precedence to the built-in benchmark data. The search result is displayed in the **Status** column in the Physical Server Inventory view.

The built-in benchmark database of Capacity Planner contains some benchmark values for some types of physical servers. While loading knowledge base data for benchmark if you choose **Benchmark** to upload the knowledge base, the built-in benchmark database is overwritten by the values that are uploaded. If this is not wanted, upload the created benchmark csv file using **Custom Benchmark**. The benchmarks uploaded through Custom Benchmark are giving preference over the built-in benchmark database entries while finding a matching performance benchmark for the physical server.

For more information about uploading a custom benchmark file, see "Uploading custom benchmark file in Capacity Planner" on page 301.

Note: The CUSTOM_USER_DEFINED_BENCHMARK.csv file and the USER_DEFINED_BENCHMARK.csv file have similar file format.

VIRTUALIZATION_OVERHEAD.csv

This file contains the virtualization overhead values that are used in optimization generation. These values are by default set to 10% of CPU, memory, network bandwidth, and storage I/O resources. These values can be modified and loaded to the Capacity Planner as needed by the environment. Use the file format inTable 109.

Column name	Maximum column length	Details
VIRTUALIZATION_OVERHEAD_PK	8	This column must be blank.
HYPERVISOR_TYPE	64	Hypervisor type, for example, VMware ESX.
HYPERVISOR_VERSION	64	Hypervisor version.
CPU_CORES_RANGE_LOWER	4 (Integer)	Lower limit for CPU cores range for which the overhead applies.
CPU_CORES_RANGE_UPPER	4 (Integer)	Upper limit for CPU cores range for which the overhead applies.
PERCENTAGE_CPU_OVERHEAD	4 (Integer)	CPU overhead value (percentage).
PERCENTAGE_MEMORY_OVERHEAD	4 (Integer)	Memory overhead value (percentage).
PERCENTAGE_NETWORK_OVERHEAD	4 (Integer)	Currently not used.
PERCENTAGE_INTERNAL_DISK_ OVERHEAD	4 (Integer)	Currently not used.
PERCENTAGE_DISK_I0_OVERHEAD	4 (Integer)	Currently not used.
APPLICATION_TYPE	30	Currently not used.

Table 109. Virtualization overhead

Appendix C. Capacity Planner database recommendations

Recommended minimum disk space required is 300 MB.

Tuning parameters

- update db cfg for itmdw2 using LOGFILSIZ 8192
- update db cfg for itmdw2 using LOGPRIMARY 20
- update db cfg for itmdw2 using LOGBUFSZ 1024

See "Capacity Planner problems" in the *IBM Tivoli Monitoring for Virtual Environments Dashboard, Reporting, and Capacity Planning Troubleshooting Guide* for additional database settings.

- If Tivoli Data Warehouse is on DB2, use the Tivoli Data Warehouse database server. Otherwise, either install DB2 on the server where the Dashboard Application Services Hub and Tivoli Common Reporting V3.1 or later are installed or install DB2 on a separate server, and catalog the server on the system for the Dashboard Application Services Hub and Tivoli Common Reporting.
- Increase the transaction logs and buffer pool sizes.

Support Information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides different ways for you to obtain the support you need, such as online or IBM Support Assistant.

Online

The following sites contain troubleshooting information:

- Go to the IBM Software Support site at http://www.ibm.com/software/support/probsub.html and follow the instructions.
- Go to the IBM Tivoli Distributed Monitoring and Application Management Wiki at http://www.ibm.com/developerworks/wikis/display/tivolimonitoring/Home. Feel free to contribute to this wiki.

IBM Support Assistant

The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to http://www.ibm.com/software/support/isa

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:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement 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/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation 2Z4A/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.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© IBM 2009. Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2009. All rights reserved.

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

Trademarks

IBM, the IBM logo, and ibm.com[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information (www.ibm.com/legal/copytrade.shtml).

Intel, Intel logo, and Intel Xeon, are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Microsoft and Windows 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, or service names may be trademarks or service marks of others.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Index

Α

accounting report 246 adding 266, 331, 333 adjust usage data 351 analytics Cognos reports overview 305 application sizing rules 351

С

capacity efficiency index 299 capacity planner 8, 11, 16, 23, 26, 36, 39, 40, 45, 54, 55, 77, 94 Cognos reports report descriptions 305 graphical user interface See GUI GUI 258 Capacity planner 317 Capacity Planner 255 capacity efficiency index 299 Optimization rules 286, 348 performance risk index 299 Planning Center 255, 315 troubleshooting 375 Capacity Planner configuration post migrate 300, 301 Capacity Planner database recommendations 389 cataloging 101 change history 157, 164 Cluster 150, 151 change history 153 cluster dashboard 147 Cluster details 149 clusters 147 Cognos errors 365 Cognos-based report packages prerequisites 101 uninstalling 33, 119 Cognos-based reports attribute groups 244 Linux KVM Number of Hosts and VMs monitored 246 Linux KVM Number of VMs for Host 250 Linux KVM Physical CPU Utilization by VM 247 Linux KVM Report Prerequisite Scanner 245 Linux KVM Resources Needed for Additional VMS on a Host 252 Linux KVM Storage Pool Capacity and Performance Trend 249 Linux KVM Top or Bottom VMs by Physical CPU Utilization report 254 predefined 244 communication with servers 121 configuration 157 Capacity Planner migrate 300, 301 connections 121 configure 317 Configuring 122, 123, 136

connections 121 cookies 395 Create Connection 134 CSV formatting 381

D

Dashboard 147, 164, 175 Virtual Machines 159 VMware Clusters 147 Dashboard for VMware overview 147 troubleshooting 369 dashboards setting a trace 174 data collection 122, 123 data store 164 Data store 167, 169 change history 168 Data store details 166 database 317 Capacity Planner database tuning 145 tuning 145 databases 101

Ε

edit current environment 333 enhancements 4 entitlement values 351 ESX Server details 172 ESX Server 170 ESX Server Dashboard 170 ESX Servers 170 export CSV formatting 381

F

federation 137, 140

Η

health check 175 Health Check Dashboard for VMware 175 heterogeneous federation 136 historical 122, 123 history collection configuration 135

IBM Support Assistant 391IBM Tivoli Distributed Monitoring and Application Management Wiki 391IIM 10, 38 import CSV formatting 381 install 11, 16, 23, 26, 39, 40, 45, 54, 55, 77, 94 installation troubleshooting 357 installing 7, 27, 35, 58, 102, 111 Installing 7, 35 integration 4 ISA 391 Iterative Data Load 317

Κ

Knowledge base editing 385 KVM Linux 7, 36

Launch TCR 173 TEP 173 Launchpad 8, 36 LinuxKVM 101 Load Config 317 logging on 170, 175 LPAR Sizing 326 LPARs 317, 326

Μ

managed systems 317 MSSQL 137

Ν

Network 158 new in this release 4 node 140 Number of Hosts and VMs monitored 246 Number of VMs for Host 250

0

OMNIbus 4 oracle 140 Oracle 140 Overview 150, 156, 162, 167

Ρ

Performance and capacity management reports 101
Performance and Capacity Management Reports configuring historical data collection 135
performance risk index 299
performance trends and resource forecasts 246
Physical CPU Utilization by VM 247
planning 319
Planning Center 255, 315
Planning Center for PowerVM 317
PlanningCenter 329
Power 7, 36, 121
PowerVM 326
PowerVM LPAR Sizing 328

PowerVM LPAR Sizing wizard 328 PowerVM Sizing 329 Prerequisite Scanner 245 prerequisites 7, 35, 136 Cognos-based report packages 101 prerequisites checking report 244 privacy policy 395 products 4 Properties 151

R

registering 140 Registering 137, 140 report 329 reports 26, 27, 57, 58, 102, 111 accounting 246 performance trends and resource forecasts 246 prerequisites checking 244 uninstalling 64 what if analysis for workload placement 250 workload right-sizing and balancing 253 Resources Needed for Additional VMS on a Host 252

S

Selected system page 328 Server 156, 157, 158 server definition 137, 140 Server details 155 servers 101, 266, 331 Servers 153 service 140 sizing rules 351 Software Support 391 starting 10, 38 Storage Pool Capacity and Performance Trend 249 support assistant 391

Т

TADDM 4 target allocation 351 TBSM See Tivoli Business Service Manager timeout problems 362 Tivoli Application Dependency Discovery Manager See TADDM Tivoli Business Service Manager 4 Tivoli common reporting Cognos reports data collection 182 performance trends 186 report descriptions 179 resource forecasts 186 What-if analysis for capacity estimation 205 Workload right sizing and balancing 217 reports data models 230 overview 179 Tivoli Common Reporting documentation 243 Prerequisites 182 uninstalling reports 64 Tivoli Common Reporting Community 243 Tivoli Common Reporting documentation 243

Tivoli Data Warehouse database tuning 145 Tivoli Enterprise Portal Server communications 121 Tivoli Monitoring 1 top clusters by percentage usage 147 Top or Bottom VMs by Physical CPU Utilization report 254 trace Dashboard for VMware 174 troubleshooting 357 Capacity Planner 375 Dashboard for VMware 369 installation 357 timeout problems 362 Tivoli Common Reporting 365 types of installation 26, 57, 102

U

uninstalling reports 64 UNIX 140

V

virtual environments Tivoli common reporting overview 179 Virtual Environments 1 Virtual machine 162, 164 Virtual machine details 162 virtual machines 333 Virtualization Reporting troubleshooting 365 VMware 35, 101 VMWare 7, 36 VMware Health Checks 175 Volume 169

W

what if analysis reports 250 workload 326 Workload definition 328 workload right-sizing and balancing reports 253

IBW ®

Product Number: 5724-L92

Printed in USA

SC14-7493-03

