



OnCommand® Insight 7.1

Reporting Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-10165_B0
June 2015

Contents

Welcome to OnCommand Insight reporting	7
Using predefined reports to answer common questions	7
Locating predefined reports	8
Variations due to installed licenses	8
OnCommand Insight licenses	8
Reporting user roles	9
OnCommand Insight Reporting Portal installation	10
OnCommand Insight data model descriptions	11
Accessing the Data Warehouse portal	12
Data Warehouse software components	13
Data Warehouse processes	13
How Data Warehouse extracts data	14
Where to find more information about OnCommand Insight	17
Accessing OnCommand Insight Reporting Connection Help	17
OnCommand Insight on the web	18
Product documentation	18
Reporting made easy	20
Accessing the OnCommand Insight Reporting Portal from Data Warehouse	20
Opening the OnCommand Insight Reporting Portal from a browser	21
Navigating to the pre-defined OnCommand Insight reports and dashboards	21
Locating IBM Cognos documentation	22
Running a report	23
Managing reports	24
Customizing a report's output format and delivery	24
Scheduling reports	25
Setting up email notification in Reporting Connection	25
Managing dashboards	26
Accessing dashboards and creating a dashboard tab	26
Adding a dashboard from User Preferences	27
Using dashboards to survey information and display details	28
Inventory reports	30
Inventory reports	30
Fabric Summary report	30
Host HBAs report	31
Host Summary report	32
Mapped But Not Masked Volumes report	33
Storage Summary report	34
VM Paths report	36
Capacity and chargeback dashboards and reports	38
Accessing capacity reports	38
Storage Summary report	39

VM Capacity Summary report	40
Creating custom ad hoc reports	42
Report authoring tools	42
What you need to do before you design your report	42
Process of creating reports	43
How to plan the results of your custom report	43
Tips for designing reports	44
Creating your first Query Studio report	44
Accessing Query Studio	45
Selecting the data model and data elements for your report	46
Filtering and sorting data	46
Pivoting the report	47
Saving your report	47
Selecting another data model for more reporting	48
Returning to OnCommand Insight Data Warehouse	48
Creating reports with Workspace Advanced	48
Creating a sample report: Raw Capacity by Data Center (table)	48
Creating a sample report: Raw Capacity by Data Center (chart)	49
Importing and exporting report definitions in Report Studio	50
Reporting data model	50
Data model overview	50
Capacity data models	53
Chargeback data model	77
Inventory data model	81
Performance data models	86
Storage Efficiency data model	92
Data marts for SQL queries	97
Data marts in Data Warehouse	97
Inventory data mart for applications in OnCommand Insight	
(dwh_inventory)	97
Capacity data mart (dwh_capacity)	98
Capacity Efficiency data mart (dwh_capacity_efficiency)	98
File System Utilization data mart (dwh_fs_util)	98
Performance data mart (dwh_performance)	98
Port Capacity data mart for switches (dwh_ports)	98
Storage and Storage Pool Capacity data mart	
(storage_and_storage_pool_capacity_fact)	98
Volume Capacity data mart (dwh_capacity)	99
Accessing MySQL using the command-line interface	99
Accessing the database schema	100
Inventory data mart for SQL queries	100
Questions you can answer with the Inventory data mart	101
Inventory data mart considerations	101
Inventory data mart sample SQL queries	107
Capacity data mart for SQL queries	109

Questions you can answer with the Capacity data mart	109
Capacity data mart considerations	110
Capacity data mart components	112
Capacity Efficiency data mart for SQL queries	117
Questions you can answer with the Capacity Efficiency data mart	117
Capacity producers and consumers	117
Storage efficiency calculation	118
Capacity Efficiency data mart considerations	118
Capacity Efficiency data mart sample SQL query	120
File System Utilization data mart for SQL queries	120
Questions you can answer with the File System Utilization data mart	120
FS Utilization data mart considerations	120
FS Utilization data mart sample SQL query	121
Performance data mart for SQL queries	121
Questions you can answer with the Performance data mart	122
Performance data mart considerations	122
Performance data mart components	125
Performance data mart sample SQL query	127
Port Capacity data mart for SQL queries	129
Questions you can answer with the Port Capacity data mart	129
Port Capacity data mart considerations	129
Port Capacity data mart sample SQL query	130
Storage and Storage Pool Capacity data mart for SQL queries	131
Questions you can answer with the Storage and Storage Pool Capacity data mart	131
Storage and Storage Pool Capacity data mart considerations	131
Storage and Storage Pool Capacity data mart sample SQL query	132
Volume Capacity data mart for SQL queries	133
Questions you can answer with the Volume Capacity data mart	133
Volume Capacity data mart sample SQL query	134
OnCommand Insight Reporting data model glossary	135
Data model dimension tables for Reporting Connection	135
Application Dimension	135
Application Group Dimension	136
Application Dimension	136
Business Entity Hierarchy	137
Business Unit Dimension	137
Capacity	138
Capacity Types full description	139
Connected to Dimension	141
Data Center	141
Date Dimension	141
Fabric Dimension	143
File System Dimension	143
Host Dimension	144

Host Group Dimension	144
Internal Volume Dimension	144
Port Dimension	145
Qtree Dimension	145
Service Level Dimension	146
Storage Dimension	146
Storage Pool Dimension	147
Switch Dimension	148
Switch Port Dimension	149
Switch Performance	149
Tape Dimension	151
Tier Dimension	151
Time Dimension	152
Time period fact tables	152
VM Dimension	152
Volume Dimension	153
Data mart dimension tables for SQL queries	154
Application dimension in data marts	154
Application group dimension in data marts	154
Business Entity dimension in data marts	154
Connected Device dimension in data marts	154
Date dimension in data marts	154
Data store dimension	155
Disk dimension in data marts	155
Fabric dimension in data marts	156
File System Dimension in data marts	156
Host dimension in data marts	156
Hosts and host group dimensions	156
Internal volume dimension in data marts	157
Port dimension in data marts	157
Qtree dimension in data marts	157
Service Level dimension in data marts	157
Storage dimension in data marts	157
Storage pool dimension in data marts	157
Switch dimension in data marts	158
Tape dimension in data marts	158
Tier dimension in data marts	158
Time dimension in data marts	158
VM dimension in data marts	158
Volume dimension in data marts	158
Copyright information	159
Trademark information	160
How to send comments about documentation and receive update notifications	161
Index	162

Welcome to OnCommand Insight reporting

The OnCommand Insight Reporting Connection is a business intelligence tool that enables you to view pre-defined reports or create your own custom reports. The OnCommand Insight Reporting Connection generates the reports from the Data Warehouse (DWH) data.

You can do the following tasks using OnCommand Insight Reporting Connection:

- Run a report.
- Create your own custom report.
- Customize the report format and delivery method.
- Schedule reports to run.
- Email reports.
- Use colors to represent thresholds on data.
- Access a collection of report results on a dashboard.
- Add tabs to views to show a dashboard collection of report results.

The pre-defined Inventory reports are the standard OnCommand Insight reports. Additional pre-defined reports are provided with the Assure, Perform, and Plan licenses. This guide describes the pre-defined reports available with all of the product licenses.

Note: The OnCommand Insight data can be viewed in any standard browser. However, the report generation engine requires the Microsoft Windows Internet Explorer (IE) browser. Reports cannot be dependably generated when using other web browsers.

Using predefined reports to answer common questions

The following table lists the OnCommand Insight predefined reports that answer common questions.

Question	Report	Locating the report
Where can I get an overview of storage utilization by host?	Host Summary report	Public Folders > Reports
How do I determine where a certain device is connected and find free switch ports?	Fabric Summary report	Public Folders > Reports
How can I get an overview of all of my storage?	Storage Summary report	Public Folders > Reports
Do I have any unmasked volumes that can be seen by any host?	Mapped But Not Masked Volumes report	Public Folders > Reports
What is the comparison of ESX to VM performance?	ESX vs VM Performance report	Public Folders > Reports
What data is occupying capacity but is unused?	Orphaned Storage Details	Public Folders > Storage Manager Dashboard

Question	Report	Locating the report
Where can I find storage pools reports?	Storage Pools reports	Public Folders > Storage Manager Dashboard
Where can I find vendor-specific reports (HDS, NetApp, Symmetrix, XIV)	Vendor reports	Public Folders > Vendor specific reports

Related concepts

[Inventory reports](#) on page 30

[Capacity and chargeback dashboards and reports](#) on page 38

Locating predefined reports

The predefined OnCommand Insight reports can be accessed from different locations in the OnCommand Insight Reporting Portal.

One report can be part of a dashboard and can also be generated from a list of individual reports; there can be several locations in the portal for any one report. The multiple access points for reports are a result of OnCommand Insight customer requests for report groupings to help them with different types of analysis.

This guide lists the shortest paths to locate the individual reports.

- **Public Folders > Reports**
- **Public Folders > Packages**
- **Public Folders > Storage Manager Dashboard**
- **Public Folders > Vendor Specific Reports**

However, the documentation lists only the shorter path to help you to find the report quickly. If an individual report is generated on a dashboard, that information is also provided.

As you use the reports, you can group favorite dashboards and reports together into your own folder in the portal.

Variations due to installed licenses

Data in the OnCommand Insight reports is based upon the OnCommand Insight licenses that you have purchased. For example, if you do not have OnCommand Insight Plan installed, then you will not receive any data in the capacity reports.

You might see differences between the available reports in your OnCommand Insight windows compared to the illustrations in the documentation. These variations are due to differences between the installed licenses on your system and the licenses on the system used to create the illustrations.

OnCommand Insight licenses

OnCommand Insight OCI operates with licenses that enable specific features on the Insight Server.

Discover

Basic Insight license that supports Inventory.

Assure

Supports Inventory and Assurance policy and violation functionality.

Perform

Supports Inventory and Performance and policy management.

Plan

Supports the Inventory and Planning functions.

Licenses are based on the following product characteristics:

- By terabyte of managed capacity
- By time (or perpetual)
- By product (Assure, Perform, and Plan), which can be purchased separately.

License keys are a set of unique strings that are generated for each customer. Obtain them from your OnCommand Insight representative.

Your installed licenses control the options available in the OCI software:

License	Assurance menu	Planning menu
Perform only	<ul style="list-style-type: none"> • Changes • Switch Port Performance Alerts • Port Balance Violations • Disk Utilization Violations 	Does not appear
Plan only	<ul style="list-style-type: none"> • Reservation Violations 	<ul style="list-style-type: none"> • Requests • Pending Tasks

Reporting user roles

Each user account is assigned a role with a set of permissions. The number of users is limited by the number of Reporting licenses attached to each role.

Each role can perform the following actions:

Recipient

Views OnCommand Insight Reporting portal dashboards and reports and sets personal preferences such as those for languages and time zones.

Note: Recipients cannot create reports, run reports, schedule reports, export reports, nor perform administrative tasks.

Business Consumer

Runs reports in Workspace and runs reports interactively in addition to performing all Recipient options.

Business Author

Views scheduled reports, runs reports interactively, and creates some reports in addition to performing all Business Consumer options.

Pro Author

Creates reports in addition to performing all Business Author options and has advanced reporting tool options that are available also with Report Studio.

Administrator

Performs reporting administrative tasks such as the import and export of report definitions, configuration of reports, configuration of data sources, and the shutdown and restart of reporting tasks.

The following table shows the privileges and the maximum number of users allowed for each role:

Feature	Recipient	Business Consumer	Business Author (Insight Plan only)	Pro Author (Insight Plan only)	Admin
View reports in the Public Folders and My Folders tabs	Yes	Yes	Yes	Yes	Yes
Run reports	No	Yes	Yes	Yes	Yes
Schedule reports	No	Yes	Yes	Yes	Yes
Create reports in Query Studio	No	No	Yes	Yes	No
Create reports in Workspace	No	Yes	Yes	Yes	No
Create reports in Workspace Advanced	No	No	Yes	Yes	No
Create reports in Report Studio	No	No	No	Yes	No
Perform administrative tasks	No	No	No	No	Yes
Number of users	Number of OnCommand Insight users	20	2	1	1

When you add a new Data Warehouse and Reporting user, if you exceed the limit in a role, the user is added as “deactivated,” and you need to deactivate or remove another user with that role to give a new user membership.

Note: Report authoring capabilities require Insight Plan license. You can add additional Business Author and Pro Author users by purchasing the ARAP (Additional Report Authoring Package). Contact your OnCommand Insight representative for assistance.

These reporting user roles do not affect direct database access. These reporting user roles do not impact your ability to create SQL queries using the data marts.

OnCommand Insight Reporting Portal installation

The Insight Reporting Portal requires the installation of other components to work as designed.

The following environment is necessary to provide optimum use of the OnCommand Insight reporting data model:

- The current version of OnCommand Insight
- The current version of OnCommand Insight Data Warehouse

See the following resources for additional information:

- For installation instructions for the OnCommand Insight Reporting Portal, see the *OnCommand Insight Installation and Administration Guide*.
- For more information about Cognos features, see the Cognos documentation.

OnCommand Insight data model descriptions

OnCommand Insight includes several data models from which you can select predefined reports or create your own custom report.

Capacity data model

Enables you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, and VM capacity. The Capacity data model is a container for several capacity data models:

Storage and Storage Pool Capacity data model

Enables you to answer questions about storage capacity resource planning, including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model, because it is a simpler, targeted data model.

File System Utilization data model

Enables you to answer questions about file system utilization. This data model provides visibility into capacity utilization by hosts on the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type.

Internal Volume Capacity data model

Enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time.

Port Capacity data model

Enables you to answer questions about switch port connectivity, port status, and speed over time. This data can be used when planning for purchases of new SAN switches.

Qtree Capacity data model

Enables you to trend qtree utilization (with data such as used versus allocated capacity) over time. You can also see the information by different dimensions, for example, by business entity, application, tier, and service level.

VM Capacity data model

Enables you to report your virtual machine environment and its capacity usage. This data model lets you report on changes in capacity usage over time for virtual machines and data stores. The data model also provides thin provisioning and virtual machine chargeback data.

Volume Capacity data model

Enables you to analyze all aspects of volumes in your environment and look at them by vendor, model, tier, service level, and data center. You can view capacity related to orphan volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technology (iSCSI or FC) and compare virtual to non-virtual volumes for array virtualization issues.

Chargeback data model

Enables you to answer questions about used capacity and allocated capacity on storage resources (volumes, internal volumes, and qtrees). This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data. Report data can be categorized by service level and storage tier.

You can use this data model to generate chargeback reports by finding the amount of capacity used by a business entity. This data model enables you to create unified reporting of multiple protocols (including NAS, SAN, FC, and iSCSI).

Inventory data model

Enables you to answer questions about inventory resources such as hosts, storage systems, switches, disks, tapes, qtrees, quotas, virtual machines and servers, and generic devices. The Inventory data model includes several submarts that enable you to see information about replications, FC paths, iSCSI paths, NFS paths, and violations. The Inventory data model does not include historical data.

Performance data model

Enables you to answer questions about performance for volumes, application volumes, internal volumes, switches, applications, VMs, ESX vs VM, hosts, and application nodes.

The Performance data model includes information that helps you determine appropriateness of tiers, storage misconfigurations for applications, and volume and internal volume last access times. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

Storage Efficiency data model

Enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount used or consumed (the physical measurement). For example, when thin provisioning is enabled, OnCommand Insight indicates how much capacity is actually taken from the device. You can also use this model to determine efficiency when deduplication is enabled.

Accessing the Data Warehouse portal

The OnCommand Insight Data Warehouse portal is a web-based user interface that you can use to update connector information, view job queues, schedule daily builds, select annotations, set up email notifications, view system information, build the database, reset Data Warehouse, perform a backup and restore of the database, troubleshoot issues, manage Data Warehouse and Reporting portal user accounts, and access documentation and schema diagrams.

About this task

The following is the default user name and password. It is a good idea to change these defaults after installation.

- User name: admin
- Password: admin123

Steps

1. Log in to the Data Warehouse portal at `https://hostname/dwh`, where *hostname* is the name of the system where OnCommand Insight Data Warehouse is installed.
2. Enter your user name and password.
3. Click **Login**.

The Data Warehouse portal opens.

The screenshot shows the 'DWH Admin' interface with a sidebar menu on the left containing options like Connectors, Jobs, Schedule, Annotations, Email Notification, System Information, Build from history, Reset DWH, Backup/Restore, Troubleshooting, and User Management. The main area displays a table of jobs with columns for ID, Name, Status, Start time, and End time. The jobs listed include Job Status Notification (Completed), Performance (Aborted), File System Utilization (Aborted), Storage Efficiency (Aborted), Ports (Aborted), Capacity (Aborted), Datamarts Preparation (Aborted), Dimensions (Aborted), Post Inventory (Aborted), and two failed inventory connector jobs. A pagination bar at the bottom of the table shows page 1 of 9.

#	Name	Status	Start time	End time
9476	Job Status Notification	COMPLETED	4/2/15 5:01 AM	4/2/15 5:02 AM
9475	Performance	ABORTED		4/2/15 5:01 AM
9474	File System Utilization	ABORTED		4/2/15 5:01 AM
9473	Storage Efficiency	ABORTED		4/2/15 5:01 AM
9472	Ports	ABORTED		4/2/15 5:01 AM
9471	Capacity	ABORTED		4/2/15 5:01 AM
9470	Datamarts Preparation	ABORTED		4/2/15 5:01 AM
9469	Dimensions	ABORTED		4/2/15 5:01 AM
9468	Post Inventory	ABORTED		4/2/15 5:01 AM
9467	Inventory: connector 'MPSC-STM-PVOCI1', 'Apr 2, 2015 5:00 AM'	FAILED	4/2/15 5:01 AM	4/2/15 5:01 AM
9466	Inventory: connector 'CHDC-STM-PROVCI1', 'Apr 2, 2015 5:00 AM'	FAILED	4/2/15 5:00 AM	4/2/15 5:01 AM
9465	Pre Inventory	COMPLETED	4/2/15 5:00 AM	4/2/15 5:00 AM

Data Warehouse software components

OnCommand Insight Data Warehouse includes several software components.

- MySQL database
The back-end repository for data mart tables
- IBM Cognos
The reporting engine for OnCommand Insight
- Apache Derby Database
Used for storing Cognos configuration and content
- JBoss
The Java Enterprise application server that hosts OnCommand Insight components

Data Warehouse processes

Data Warehouse performs many types of processes.

ETL process

The Extract Transform and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The Data Warehouse build process is an ETL process.

Jobs

Data Warehouse performs and reports on jobs such as these: inventory, dimensions, capacity, port capacity, VM capacity, file system utilization, performance, capacity efficiency, licenses, history build, dynamic annotations, connector removal, skipped build, AutoSupport option, and maintenance jobs.

Consolidation process


Data Warehouse supports the consolidation of multiple OnCommand Insight servers into the same Data Warehouse database. In many configurations it might happen that the same

object is reported from multiple connectors (that is, the same switch exists in two OnCommand Insight instances). In that case, Data Warehouse consolidates the multiple objects into one (a primary connector is chosen and the object's data is taken from that connector only).

How Data Warehouse extracts data

The Extract, Transform, and Load (ETL) process retrieves data from multiple OnCommand Insight databases, transforms the data, and saves it into the data marts.

OnCommand Insight connectors invoke a series of batch jobs to extract data from multiple OnCommand Insight MySQL databases and publish the data in various data marts, as shown in the following diagram.



The ETL process includes these individual processes:

Extract

This process takes data from multiple OnCommand Insight databases, transforms the data, and saves it into the data mart. The process is performed against each OnCommand Insight instance at the same time. To ensure that data cleansing and deduplication is performed, it is not possible to split the ETL process into multiple scheduled ETL operations.

Transform

This process applies business logic rules or functions to extract the data from the OnCommand Insight database.

Load

This process loads the transformed data into public data marts.

ETL frequency and date data

You should run the Extract, Transform, and Load (ETL) process at least once per day; however, you choose to run ETL numerous times if needed.

By default, the Cognos reporting engine treats all capacity and performance facts as additive. As a result, there is a risk of double counting capacity data if the ETL process is run multiple times per day without the proper time filters.

Two date data elements in the Date dimension are related to the daily ETL process. The Date dimension, which is used in several data models, includes the following data elements that are affected by the ETL:

Is Day Representative

The "Is Day Representative" data element is set to a value of 1 (true) during the first ETL process run during any given day. If the first ETL process is run at 1:00 a.m., Is Day Representative is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If a

second ETL is scheduled later (for example, 1:00 p.m.), Is Day Representative is set to 0 (false) for the data loaded during that ETL process.

Is Latest

The "Is Latest" member is set to a value of 1 (true) after each ETL process completes. If the first ETL process is run at 1:00 a.m., Is Latest is set to 1 for all of the data loaded during the 1:00 a.m. ETL process. If another ETL process is scheduled later (for example, 1:00 p.m.), Is Latest is set to 1 for data loaded during the 1 p.m. ETL process. The ETL process also sets the 1:00 a.m. ETL load's Is Latest entry to 0 (false).

How historical data is retained in Data Warehouse

Data is maintained in Data Warehouse according to a schedule.

Data Warehouse retains historical data based on the data marts and granularity of the data, as shown in the following summary:

Data mart	Measured object	Granularity	Retention period
Performance marts	Volumes and internal volumes	Hourly	14 days
Performance marts	Volumes and internal volumes	Daily	13 months
Performance marts	Application	Hourly	13 months
Performance marts	Host	Hourly	13 months
Performance marts	Switch performance for port	Hourly	5 weeks
Performance marts	Switch performance for host, storage and tape	Hourly	13 months
Performance marts	Storage node	Hourly	14 days
Performance marts	Storage node	Daily	13 months
Capacity marts	All (except individual volumes)	Daily	13 months
Capacity marts	All (except individual volumes)	Monthly representative	14 months and beyond
Inventory mart	Individual volumes	Current state	1 day (or until next ETL)

After 13 months, which is configurable, Data Warehouse retains only one record per month instead of one record per day for capacity, performance, and resource data in the following tables:

- Chargeback fact table (dwh_capacity.chargeback_fact)
- File System Utilization fact table (dwh_fs_util.fs_util_fact)
- Host fact table (dwh_sa.sa_host_fact)
- Internal Volume Capacity fact table (dwh_capacity.internal_volume_capacity_fact)
- Ports fact table (dwh_ports.ports_fact)
- Qtree Capacity fact table (dwh_capacity.qtree_capacity_fact)

- Storage and Storage Pool Capacity fact table (dwh_capacity.storage_and_storage_pool_capacity_fact)
- Volume capacity fact table (dwh_capacity.vm_capacity_fact)
- Storage Node Hourly Performance (storage_node_hourly_performance_fact) and Storage Node Daily Performance (storage_node_daily_performance_fact) fact tables

Data retention, ETL, and time periods

OnCommand Insight Data Warehouse retains data obtained from the Extract, Transform, and Load (ETL) process for different time periods based on the different data marts and time granularity of the data.

Performance Marts and hourly granularity for volumes and internal volumes

The OnCommand Insight Data Warehouse records the hourly averages, hourly maximums, and access bit for each hour of the day (24 data points) for 14 days. The access bit is a Boolean value that is true if the volume is accessed or false if the volume is not accessed during the hourly interval. All 24 data points for the preceding day are obtained during the first ETL process of the day.

You do not need to run one ETL process per hour to gather the hourly data points. Running additional ETL processes during the day does not obtain any performance information from the OnCommand Insight Servers.

Performance Marts and daily granularity for volumes and internal volumes

Each day when the ETL is processed, the daily averages for the preceding day are calculated and populated within Data Warehouse. The daily average is a summary of the 24 data points for the previous day. The performance data marts retain daily summaries for volumes and internal volumes for 13 months.

Capacity marts and daily granularity

The Capacity marts provide daily measurements for various capacity facts on a daily basis for a period of 13 months. The capacity facts in Data Warehouse are current as of the last data source acquisition for the device prior to the ETL.

Capacity marts and monthly granularity

Data Warehouse retains daily capacity data for 13 months. After the 13-month threshold is reached, the capacity data is summarized on a monthly basis. The monthly data is based on the values reflected by the date that is the month representative date.

The following table shows which monthly data is included in the monthly summary:

Date	Is Month Representative value	Allocated capacity
Jan 1	1 (True)	50 TB
Jan 2	0 (False)	52 TB
...
Jan 31	0 (False)	65 TB
Feb 1	1 (True)	65 TB

Based on the table, a monthly report would show 50 TB allocated for January and 65 TB allocated for February. All of the other capacity values for January would not be included in the monthly summary.

Inventory mart

The Inventory data mart is not historical. Each time an ETL process is run, the Inventory mart is erased and rebuilt. Therefore, any reports generated out of the Inventory mart do not reflect historical inventory configuration.

Where to find more information about OnCommand Insight

You can find more information about OnCommand Insight on the NetApp Support Site and in other OnCommand Insight documentation.


Accessing OnCommand Insight Reporting Connection Help

You can find descriptions of predefined reports, instructions on how to create custom reports, and descriptions of data models in the Help for OnCommand Insight Reporting Connection.

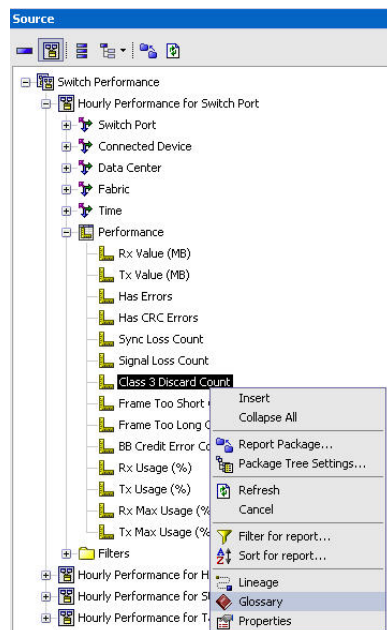
Step

1. Do one of the following:

- From OnCommand Insight Reporting Connection, click the **?** in the upper right corner. From the OnCommand Insight Reporting Connection Help menu, select **OnCommand Insight Reporting Help**.



- In Query Studio or Report Studio, after you select a data model from the Public Folders tab and begin to develop a custom report, expand the data element tree on the left, right-click on a data element, and select **Glossary**.



OnCommand Insight on the web

For comprehensive, up-to-date information about OnCommand Insight, use these NetApp web site resources.

- OnCommand Insight product web site at www.netapp.com/oncommandinsight
- The NetApp Support Site at: mysupport.netapp.com
- The OnCommand Insight data source Interoperability Matrix at mysupport.netapp.com/matrix.

Product documentation

You can access the OnCommand Insight guides at the NetApp Support Site to learn how to use the product.

You can access the following documents from the NetApp Support Site at mysupport.netapp.com/documentation/productsatoz/index.html:

OnCommand Insight Installation Guide for Microsoft Windows

Describes the installation prerequisites and procedures for the web-based interface version of OnCommand Insight.

OnCommand Insight Getting Started Guide

Helps new OnCommand Insight users set up and customize their installed system and begin using it for improved efficiency.

OnCommand Insight Configuration and Administration Guide

Provides an overview of suite architecture, including instructions for starting your system, discovering the logical and physical storage resources in your storage environment, and performing administrative tasks.

Also, describes the configuration parameters and some installation procedures for data sources used by OnCommand Insight and provides recommended methods of discovering your storage environment for the OnCommand Insight deployment.

OnCommand Insight Inventory User Guide

Provides information about the tools and features of your SAN or NAS inventory environment. The Inventory features are the foundation used by all other products in the OnCommand Insight suite. After looking at the *OnCommand Insight Getting Started Guide*, use this guide to learn about basic features common across all OnCommand Insight modules.

OnCommand Insight Assurance User Guide

Provides an overview of how to use OnCommand Insight to analyze and validate your storage network environment and to automate monitoring of operations, as well as procedures for making safe changes without disrupting availability.

OnCommand Insight Performance User Guide

Provides an overview of how to use OnCommand Insight to reclaim underutilized resources, manage tiers, identify multipath risks, and troubleshoot ongoing performance bottlenecks.

OnCommand Insight Planning User Guide

Provides an overview of how to use OnCommand Insight to enable educated capacity management decisions by managing the end-to-end resource order and allocation process.

Also, describes how system administrators can communicate storage capacity requirements to SAN managers using OnCommand Insight Connect Applications web access instead of the full OnCommand Insight feature set.

OnCommand Insight Reporting Guide

Describes reports from the centralized reports portal that support viewing critical inventory and capacity-related information from the data warehouse.

Also describes the metadata model upon which OnCommand Insight reports are based.

Information in this guide is also available from the OnCommand Insight Data Warehouse portal and in the OnCommand Insight Reporting Connection.

OnCommand Insight Data Warehouse Administration Guide

Describes a data warehouse repository that consolidates multiple OnCommand Insight operational databases in an easy-to-query format.

OnCommand Insight Connect API Reference

Provides an overview of how to use the API interface to enable integration with other applications, such as reporting and monitoring systems. Helps customers and independent software vendors (ISVs) to develop applications using the OnCommand Insight API interface.

Reporting made easy

You can generate pre-defined reports from the OnCommand Insight Reporting Portal, email them to other users, and even modify them. Several reports enable you to filter by device, business entity, or tier. The reporting tools use IBM Cognos as a foundation and give you many data presentation options.

- The OnCommand Insight pre-defined reports show your inventory, storage capacity, chargeback, performance, and storage efficiency data. You can modify the pre-defined reports and save your modifications.
- You can create custom reports using the OnCommand Insight Reporting Portal tools: Business Insight and Business Insight Advanced.
- OnCommand Insight provides quick access to report data through dashboards. Additionally, you can group your own report selections and, using the Business Insight tools, you can display them in your own dashboards.

The report data available to you is controlled by several things, including the following:

- Login access to the OnCommand Insight Reporting Portal, which is defined by roles.
- The setup of the OnCommand Insight Data Warehouse, which stores the data for the reports.

You can generate reports in various formats, including HTML, PDF, CSV and Excel.

OnCommand Insight accommodates multiple tenancy in reporting by enabling you to associate users with business units. With this feature, administrators can separate data or reports according to the attributes of a user or his affiliation.

For additional instructions about how to use the IBM Cognos features, see the Cognos online Help.

Related tasks

[Locating IBM Cognos documentation](#) on page 22

Accessing the OnCommand Insight Reporting Portal from Data Warehouse


You can access the OnCommand Insight Reporting Portal directly from a web browser or from Data Warehouse. Use the portal to access the predefined reports in the Query Studio or Report Studio (where you can create your own reports using Data Warehouse data). IBM Cognos software is used to create and generate the predefined reports.

About this task

This procedure describes how to access the portal from the Data Warehouse.

Steps

1. Open a web browser.
2. Enter the following URL:
`http://host_name/dwh`
3. Enter your user name and password, and then click **OK**.

4. In the Data Warehouse main menu, click .
5. From the login browser page that appears, enter your user name and password, and then click OK.

Result

The Storage Manager Dashboard is displayed.

Opening the OnCommand Insight Reporting Portal from a browser

You can access the OnCommand Insight Reporting Portal directly from a web browser.

Steps

1. Open a web browser.
2. Enter the following URL:
`http://<host-name>/reporting`
3. Enter your user name and password, and then click **OK**.

Navigating to the pre-defined OnCommand Insight reports and dashboards

When you open the Reporting Portal, the Public Folders tab is the starting point for you to select the type of information you need in the reports and dashboards.

Steps

1. Click the Public Folders link and select the information category that you want to use.
2. In the information category, you can select dashboards, folders, and individual reports.
For example, the Capacity > Storage Capacity information category has Dashboards and Reports links as well as custom folders for your own reports and dashboards.



3. Click **Dashboards** or **Reports** to navigate to specific OnCommand Insight reports.

Related concepts

[Locating predefined reports](#) on page 8

Locating IBM Cognos documentation

For basic information such as how to start and stop the Reporting portal software, see the IBM Cognos documentation installed with the product. You can search with a web browser for information about any of the IBM Cognos reporting products, such as Query Studio, Report Studio, Business Insight, or Business Insight Advanced on the IBM website in the Information Centers for those software products.

Steps

1. To locate the IBM Cognos documentation installed with OnCommand Insight, navigate to this directory.

```
<install_dir>\cognos\c10_64\webcontent\documentation\help_docs.html
```
2. You can also display topics describing individual IBM Cognos windows used in the OnCommand Insight Reporting Portal. Click the ? icon on the window toolbar.

Running a report

You can quickly run the pre-defined reports from the OnCommand Insight Reporting Portal. Some reports might require you to select parameters that filter the data to be displayed in the generated report.

Steps

1. Select the folder containing the types of reports you want.
2. Click the report link to generate the report or to open the parameter filtering page for that report.
3. If the report has a parameters page, select the report parameters you want to use to filter the data displayed.

In many reports, you can leave the parameter selections untouched to select all of the possible types.

4. Click **Finish** to generate the report.

Managing reports

For each report, you can select the **More** link in the Actions column and access all of the report operations, such as setting report properties, scheduling reports, or emailing reports. Administrators have more management options available than other users.

Administrators can set permissions for other report users according to their OnCommand Insight roles.

Related tasks

[Customizing a report's output format and delivery](#) on page 24

Customizing a report's output format and delivery

You can customize the format and delivery method of a pre-defined OnCommand Insight report.

Steps

1. Log into the OnCommand Insight Data Warehouse.
2. Open the OnCommand Insight Reporting Portal and select the report you want to customize.
3. To the right of the report name are several Actions icons and a More link. Click the **Run with options** triangle icon to open the output format options for this report.

Run with options - Storage Capacity Usage Trends by Timeframe

Select how you want to run and receive your report.

Format:

Accessibility:
☐ Enable accessibility support

Language:

Delivery:
☒ View the report now
☐ Save the report
☐ Print the report
Printer location: [Select a printer...](#)

Prompt values:
 No values saved
☒ Prompt for values

4. Set these options:
 - **Format** as HTML, PDF, Excel 2007, Excel 2003, Delimited Text (CSV), or XML
 - **Accessibility** check to enable accessibility support for this report
 - **Language** select from the list of available languages. The default is English (United States).
 - **Delivery** select to view the report, save it, or print the report on a selected printer
 - **Prompt values** check to prompt for values
5. Click **Run** to produce the report using the selections.

Scheduling reports

You can set up the report to run at a recurring date and time and set its runtime priority. You can run it using the default values or specify the options, and you can disable the schedule without losing any of its details.

Steps

1. From the OnCommand Insight Reporting Portal, select the report.
2. Click the **More** link and then click **New schedule**.
3. Enter the schedule parameters for the report.
For more information about these selections, click the **Help** link.
4. Click **OK**.

Setting up email notification in Reporting Connection

Setting up email notification is required to send out reports from Query Studio or Reports Studio via email.

Steps

1. From the Windows Start menu, select **All Programs > IBM Cognos > IBM Cognos configuration**.
2. Under the Explorer column at the end, select **Notification**.
3. On the right side of the **Notification** screen, enter the following information:
 - SMTP mail server: Enter the server name registered in DNS or SMTP server IP, for example, 10.10.10.10:25.
 - Account and password needed to log into the mail server.
 - Default sender: Enter the email account that was added to the generated emails as the sender, for example, user@domain.com.
4. Select **File > Save**.
5. In the Cognos Configuration window, right-click **Local Configuration**, and then select **Restart** to restart the Cognos service.
You could also select **Actions > Restart** or **Start** if the server is not running.
6. Test the email server setting:
 - a. From the Windows Start menu, select **All Programs > IBM Cognos > IBM Cognos configuration**.
 - b. Right-click **Notification**, and then select **Test**.

Managing dashboards

OnCommand Insight provides report groupings in the form of dynamic dashboards to show different views of data relating to a specific management concern. Clicking elements in the dashboards displays more detailed information.

About this task

The dashboards group information on a subject of interest and provide these dynamic features:

- Click on chart elements to regenerate charts showing more details for the selected items.
- Links to other reports and more detailed information.

If you have a favorite dashboard, you can add a tab for it using one of these methods:

- Add the dashboard tab from the public folders.
- Add the dashboard tab from User Preferences.

Related tasks

[Using dashboards to survey information and display details](#) on page 28

Accessing dashboards and creating a dashboard tab

You might want to display a dashboard group of reports in a tab for easy access. You can add the tab from the public folders list.

Steps

1. Open the OnCommand Insight Reporting Portal.
2. Select **Public Folders > Storage Manager Dashboard**.



Individual reports and dashboards are listed together in the folder.

Within the package, several options appear, including Dashboards and Reports.

3. For quick access to a dashboard in the future, you can create a separate tab for it: to the right of the dashboard, click the **More** link.
4. In the list that opens, click the **Add to my portal tabs** link.


The new tab then displays at the top of the window with the selected dashboard as the tab label.

Adding a dashboard from User Preferences

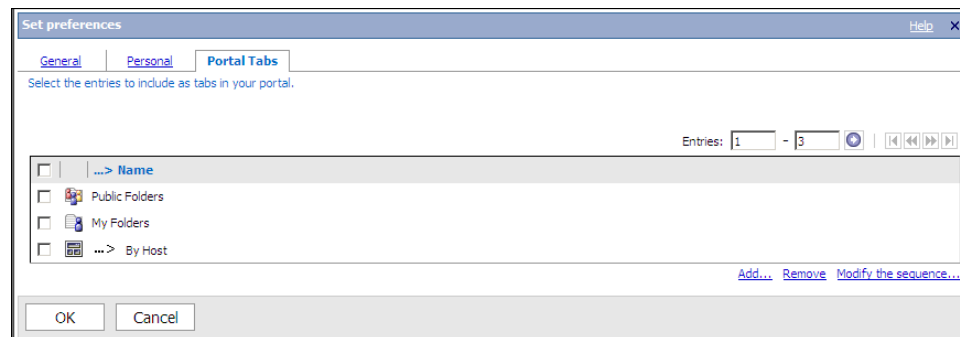
You might want to display a dashboard group of reports in a tab for easy access. You can add the tab from User Preferences.

Steps

1. Open the OnCommand Insight Reporting Connection.
2. At the upper right, click the Person icon to show a drop-down menu.



3. Select **My Preferences**.
4. In the Set preferences view, click the **Portal Tabs** tab.



5. In the Portal Tabs tab, click the **Add** link.
6. In the Public Folders and Dashboard folder, locate the dashboard that you want as a tab and check it.
7. Click the arrow to move it to the Selected Entries box, and click **OK**.


Using dashboards to survey information and display details

You can access report data through individual reports or by viewing that data in a group of report results on a dashboard.

Steps

1. To display dashboards, open the OnCommand Insight Reporting Portal.
2. Select **Public Folders** and one of the report categories.
3. Select **Dashboards**.

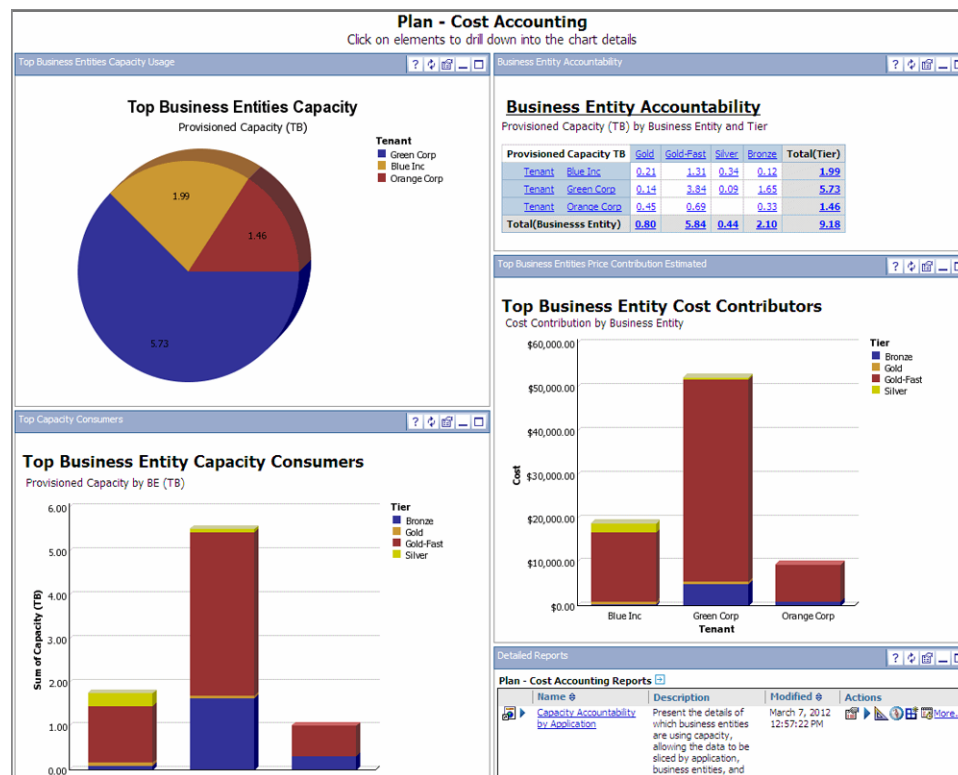
The list that opens may contain reports, folders, and dashboards.

4. To open a dashboard, select an item in the list with the dashboard icon  beside it.

The dashboard might display immediately or require you to identify environment elements for the reports and submit your selections.

5. In the dashboard, click items in graphs to open more detailed versions of the reports for those specific items.

In this example, you might want to know more about the Gold-Fast tier, so you would click portions of the graphs that depict that tier.



6. To display more information about data in report lists, click the links.

In the example, the Business Entity Accountability report in the Plan-Cost Accounting dashboard contains linked information.

7. At the bottom of most dashboards is a list of report links. Click any of these links to generate additional reports.

Related tasks

[*Managing dashboards*](#) on page 26

Inventory reports

The OnCommand Insight provides reports that display inventory information.

The Inventory reports supply information in these categories:

- Available FC and NFS paths
- Connected ports
- Data sources
- Device inventory
- Fabric summary
- Glossary
- Host data
- Internal volumes
- Mapped but not masked volumes
- Qtree information
- SAN path violation information
- Switch connectivity

Inventory reports

The detailed Inventory reports provide information about devices, FC and NSF paths, connected ports, devices, data sources, hosts, internal volumes, and qtrees.

You can access the Inventory reports from **Public Folders > Packages** or in the Inventory **Dashboards** location.

You can create Inventory reports using the report authoring tools described in this guide.

Related concepts

[*Creating custom ad hoc reports*](#) on page 42

Fabric Summary report

This report contains an overview of all the fabrics in the environment. For each fabric, the report indicates the switch ports that create the fabric.

Accessing this report

To display this report, select **Public Folders > Reports > Fabric Summary**.

Sample

Fabric Summary						
Switch Name	Switch Firmware	Serial Number	Port Count	Connected Port Count	Licensed Port Count	Licensed Connected Port Count
brocade_alpha			12	0	12	0
brocade_beta			13	0	13	0
brocade_gamma			16	3	16	3
F0brcd-a	v5.0.3		64	4	64	4
F1brcd-a	v5.0.3		64	4	64	4
F2brcd-a	v5.0.3		2	2	2	2
F2brcd-b	v5.0.3		2	2	2	2

Report columns

The report includes these columns:

Switch Name

Name of the switch.

Switch Firmware

Version of the firmware running on the switch.

Serial Number

Serial number of the switch, if available.

Port Count

Number of ports that the switch can accommodate.

Connected Port Count

Number of switch ports connected to the fabric.

Licensed Port Count

Number of licensed ports on the switch.

Licensed Connected Port Count

Number of switch ports logically connected to the fabric.

Host HBAs report

This report contains an overview of hosts in the environment including the vendor, model, and firmware version of HBAs, and firmware level of the switches to which they are connected.

Accessing this report

To display this report, select **Public Folders > Reports > Host HBAs**.

Report generation options

You can select one or more hosts for this report.

Sample

Host HBAs								
Hosts: Agassi, Safin, Samphire, Sampras, San Antonio-Host1, San Antonio-Host2, San Antonio-NAS, SE-Esx1, SE-Esx2, Spinach								
Host	Host OS	Host Port WWN	HBA Vendor	HBA Model	HBA Firmware Version	HBA Driver	Switch Serial Number	Switch Firmware Version
Agassi		10:B0:00:00:00:00:10:12	EMULEX	LP9002	FV3.93A0			v5.0.3
Agassi		10:B0:00:00:00:00:10:11	EMULEX	LP9002	FV3.93A0			v5.0.3
Safin		10:B0:00:00:00:00:10:16	EMULEX	LP9002	FV3.93A0			v5.0.3
Safin		10:B0:00:00:00:00:10:15	EMULEX	LP9002	FV3.93A0			v5.0.3
Sampras		10:B0:00:00:00:00:10:14	EMULEX	LP9002	FV3.93A0			v5.0.3
Sampras		10:B0:00:00:00:00:10:13	EMULEX	LP9002	FV3.93A0			v5.0.3
San Antonio-Host1		10:B0:00:00:00:00:08:39	EMULEX	LP9002	FV3.93A0			
San Antonio-Host2		10:B0:00:00:00:00:08:42	EMULEX	LP9002	FV3.93A0			

Interpretation

The primary use of this report is to analyze firmware compatibility when doing a firmware upgrade for a switch or an HBA.

Report columns

The report includes these columns:

Host

Selected host.

Host OS

Operating system running on the host.

Host Port WWN

World Wide Name of the HBA port.

HBA Vendor

Name of the vendor who sells the HBA.

HBA Model

Manufacturer's model number for the HBA.

HBA Firmware Version

Version of the firmware running on the HBA.

HBA Driver

Version of the driver running the HBA.

Switch Serial Number

The serial number for the switch.

Switch Firmware Version

Version of the firmware running on the switch that is connected to the HBA.

Host Summary report

This report contains an overview of storage utilization by each selected host with information for Fibre Channel and iSCSI hosts.

Accessing this report

To display this report, select **Public Folders > Reports > Host Summary**.

You can also generate this report by clicking a bar in the Host Capacity report.

Report generation options

You need to enter the names of one or more hosts and click **Search**. You can select some or all of the search results and click **Finish** to generate this report.

Sample

Host Summary							
Host Name	Port Count	FC Path Count	ISCSI Path Count	Accessed Capacity (TB)	Accessed Raw Capacity (TB)	FC Violation Count	ISCSI Violation Count
Safin	2	1	0	0.01	0.02	1	0
Samphire	0	0	4	0.03	0.06	0	4
Sampras	2	1	0	0.01	0.02	1	0
SE-Esx1	0	0	0	0.00	0.00	0	0
SE-Esx2	0	0	0	0.00	0.00	0	0
Spinach	0	0	2	0.02	0.03	0	4
SV-Host1	1	1	0	0.01	0.02	23	0
SV-Host2	1	1	0	0.01	0.02	23	0
Summary	6	4	6	0.08	0.16	48	8

Interpretation

This summary report allows you to compare ports and paths and the Fibre Channel and ISCSI capacity and violation counts.

Report columns

The report includes these columns:

Host Name

Selected host.

Port Count

Number of ports on that host.

FC Path Count

Number of paths connected to Fibre Channel hosts.

ISCSI Path Count

Number of paths connected to ISCSI hosts.

Accessed Capacity (TB)

Amount of storage that is accessible for host, in terabytes.

Accessed Raw Capacity (TB)

Actual raw disk capacity for the device, in terabytes.

FC Violation Count

Number of Fibre Channel host violations.

ISCSI Violation Count

Number of ISCSI host violations.

Mapped But Not Masked Volumes report

This report lists those volumes whose LUN (logical unit number) has been mapped for use by a particular host, but not masked to that host. Unmasked volumes can be accessed by any host, making them vulnerable to data corruption.

Accessing this report

To display this report, select **Public Folders > Reports > Mapped But Not Masked Volumes**.

Report generation options

You can select one or more data centers, storage vendors, or volumes with capacities larger than a specified amount. Leave the selection boxes untouched to select all data centers or storage vendors.

Sample

Mapped but not Masked Volumes										
Data Center	Tier	Storage	Volume	Capacity (GB)	Volume Type	LUN	Protocol Controller	Storage Ports	Storage Targets	
QDC-A	N/A	Sym-000190103253	1FF9	499.27	BCV, Meta	106		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2025	499.27	BCV, Meta	107		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2051	499.27	BCV, Meta	108		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2070	499.27	BCV, Meta	109		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			20A9	499.27	BCV, Meta	110		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			20D5	499.27	BCV, Meta	111		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2101	499.27	BCV, Meta	112		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			212D	499.27	BCV, Meta	113		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2159	499.27	BCV, Meta	114		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2185	499.27	BCV, Meta	115		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	
			2181	499.27	BCV, Meta	116		FA-03A Port: 0, FA-14A Port: 0	50:06:04:82:D5:2F:B5:42, 50:06:04:82:D5:2F:B5:4D	

Interpretation

Use this report to locate any unmasked volumes.

Report columns

The report includes these columns:

Data Center

Location of the storage array.

Tier

The storage tier (access priority) of the volume.

Storage

Name of the storage array.

Volume

Volume name.

Capacity (GB)

Sizes of the disks comprising the volume.

Volume Type

Types of disks comprising the volume, for example SSD for solid state or FC for Fibre Channel.

LUN

Logical Unit Number or host volume ID.

Protocol Controller

Protocol converter through which the volume is mapped to the storage port (applicable only for some types of storage arrays, such as EMC CLARiiON).

Storage Ports

Worldwide Port Name of the storage controller ports.

Storage Targets

Worldwide Node Name of the storage controller.

Storage Summary report

This report gives an overview of the storage devices in the environment. Storage devices are listed with their vendor, model, and microcode or firmware version, as well as a brief summary of their utilization.

Accessing this report

To display this report, select any of these methods:

- Public Folders > Reports > Storage Summary

Report generation options

You can select one or more data centers, storage vendors, storage families, storage models, and storage names. Take no action in the selection boxes to select all data.

Sample

Storage Summary															
		Show/Hide Filters Show/Hide Glossary													
		Data Centers: All Storage Families: All Storage Vendors: All Storage Names: All													
		SAN NAS 													
Data Center	Storage Name	Backend/Virtual	Raw Capacity GB	Unconfigured Raw Capacity GB	Storage Pool Capacity GB	Storage Pool Used Capacity GB	Storage Pool Unused Capacity GB	Allocated Volume Capacity GB	Unallocated Volume Capacity GB	Orphaned Volume Capacity GB	Protection Volume Capacity GB	Internal Volume Capacity GB	Internal Volume Used Capacity GB	Internal Volume Protection Capacity GB	Handled
CRG	EVA03C		48,615.04	25,047.54	18,854	18,854	0	18,854	0	0	0	0	0	0	0
HJA	opaparray01c,opaparray02c		79,181.41	2,751.8	57,295.13	9,972.57	47,322.56	12,460.19	0	0	0	21,628.5	8,613.63	0	
	opaparray01n,opaparray02n		158,362.81	5,503.6	114,590.27	47,652.34	66,937.93	56,693.92	0	146.22	0	51,521.15	42,676.96	0	

Report columns

The report includes these columns:

Data Center

Location of the storage array.

Storage Name

Name of storage array.

Backend/Virtual

Indication of virtualization. Identifies whether the storage acts as virtual storage (V), backend storage (B), or both virtual and backend storage (BV). If virtualization is not associated with the host, the column in the row for that host is displayed as empty.

Raw Capacity

The total physical capacity of all disks in the array.

Unconfigured Raw Capacity

Physical capacity of disks that are not configured into a storage pool, as well as spare and failed disks.

Storage Pool Capacity

The size of the storage pool that is available for storing user data.

Storage Pool Used Capacity

The amount of usable capacity in use on the storage pool. This includes the capacity consumed by internal volumes, volumes, Snapshot Reserve, and other storage efficiency technology overheads.

Storage Pool Unused Capacity

The usable capacity that might be available for storing additional data on the storage pool.

Allocated Volume Capacity

The capacity of volumes accessed by hosts, used for replication and volumes that are masked but no access path to them exists.

Unallocated Volume Capacity

The capacity of volumes that are not allocated. These volumes are not used for any purpose.

Orphaned Volume Capacity

The capacity of volumes that are masked but no access path to them exists.

Protection Volume Capacity

The capacity of volumes that are used for replication.

Internal Volume Capacity

The usable capacity of the internal volumes.

Internal Volume Used Capacity

The used capacity of the internal volumes.

Internal Volume Protection Capacity

The provisioned capacity of internal volumes that are used for replication.

Manufacturer

Name of manufacturer of the storage array.

Family

Product classification of the storage array (for example, CLARiiON or Symmetrix).

Model

Model number supplied by the manufacturer.

Microcode version

Version of firmware running on the storage array.

Serial Number

Serial number associated with the device.

VM Paths report

With this report, you can quickly identify the current virtual configuration.

Accessing this report

To display this report, select **Public Folders > Reports > VM Paths**.

Sample

VM Paths															
VM	VM ID	Host	Host ID	OS	State	Resource Name	Resource Technology	Resource Capacity (GB)	Datstore	Storage	VM Capacity (GB)	Datstore Capacity (GB)	Datstore Free Capacity (GB)	Datstore Utilization	Powered Off Since
QA-VM-0	vm-90	QA-Ext1	UX	running	QA-Share	NAS		0.99	QA-Datstore	STAP-Perform	1.15	97.66	19.53	20.00%	
QA-VM-1	vm-91	QA-Ext1	UX	resetting	QA-Share	NAS		0.99	QA-Datstore	STAP-Perform		97.66	19.53	20.00%	

Interpretation

You might want to use this report to see the following:

- Which virtual machine is running on which host
- Which hosts are accessing which shared volumes
- What the active access path is
- What comprises capacity allocation and usage

Report columns

The report includes these columns:

VM

Virtual machine name. Is blank when the row represents a SAN volume or a NAS share that is being accessed from a virtual host (ESX) but does not contain virtual machine data.

VM IP

IP address assigned to the virtual machine.

Host

Name of the host.

Host IP

IP address of the host.

OS

Guest operating system.

State

Shows the combined power state of the virtual hardware of the virtual machine and of the guest OS: Running, Powered Off, Powered On, Not Running, Standby, Shutting Down, Unknown, Suspended.

Resource Name

Name of the SAN volume or NAS logical disk where the data for this path resides. Might be blank if one of the following occurs:

- The SAN volume is accessible to the host but is not part of any database.
- OnCommand Insight cannot identify the storage containing the virtual machine.

Resource Technology

Populates with either SAN or NAS values that indicate what storage technology is used by the VM Path. Unlike the Client view, Resource Technology does not differentiate the SAN value between FC and iSCSI.

Resource Capacity (GB)

Size of the volume, in gigabytes.

Datastore

Name of the datastore that resides on the volume and is used by the virtual machine.

Storage

SAN or NAS storage array name. Might be blank if OnCommand Insight cannot identify the storage containing the virtual machine.

VM Capacity

Capacity allocated to the virtual machine on the datastore. This is the sum of all virtual disks stored on the datastore that are used by the virtual machine.

Datastore Capacity

Storage capacity that is allocated to the datastore (provided by the SDK).

Datastore Free Capacity

Datastore capacity that is not allocated to virtual machines (provided by the SDK).

Datastore Utilization

The ratio of used capacity divided by total capacity. Shows how much capacity is being used.

Powered Off Since

Time when the virtual machine was shut down, if it is not running.

Capacity and chargeback dashboards and reports

The capacity, chargeback, and trending data provides information to plan the overall storage allocation at a business level.

Senior managers, business unit managers, and application managers have different requirements for capacity report information. A wide variety of reports to satisfy these needs are provided in these report groups:

- **Capacity** is a large group of reports with folders for file system utilization and capacity information for internal volumes, ports, qtrees, storage, storage pools, and virtual machines.
- **Chargeback** provides capacity accountability, historical capacity accountability, and trending information.
- **Storage Efficiency** supplies the NetApp Storage Capacity Efficiency Rating report.

Some report folders, reports, and reporting tabs are hidden for use only by advanced users. See information about switching between simple and advanced reporting views.

Accessing capacity reports

In addition to the report dashboards, you can open individual capacity reports. To locate a report, search for it in the Public Folders tab of the **OnCommand Insight Reporting Connection**. You can also open a dashboard and select an individual report from the group of reports listed.

Steps

1. Open the **OnCommand Insight Reporting Connection**.
2. Click the **Public Folders** tab.
3. From the Public Folders tab, click **Packages**.

The report category names also include the version number of OnCommand Insight that is installed in your environment.

4. Open any of these folders to access the individual capacity reports:
 - File System Utilization
 - Internal Volume Capacity
 - Port Capacity
 - Qtree Capacity
 - Storage and Storage Pool Capacity
 - Storage Capacity
 - VM Capacity
 - Volume Capacity

Storage Capacity and Storage Efficiency report folders, reports, and reporting tabs are hidden for use only with advanced views. See information about switching between simple and advanced reporting views.

5. Generate reports using the Report authoring tools.

Storage Summary report

This report gives an overview of the storage devices in the environment. Storage devices are listed with their vendor, model, and microcode or firmware version, as well as a brief summary of their utilization.

Accessing this report

To display this report, select any of these methods:

- **Public Folders > Reports > Storage Summary**

Report generation options

You can select one or more data centers, storage vendors, storage families, storage models, and storage names. Take no action in the selection boxes to select all data.

Sample

Storage Summary															
		Show/Hide Filters Show/Hide Glossary Data Centers: All Storage Families: All Storage Vendors: All Storage Names: All													
		SAN NAS													
Data Center	Storage Name	Backend/Virtual	Raw Capacity GB	Unconfigured Raw Capacity GB	Storage Pool Capacity GB	Storage Pool Used Capacity GB	Storage Pool Unused Capacity GB	Allocated Volume Capacity GB	Unallocated Volume Capacity GB	Orphaned Volume Capacity GB	Protection Volume Capacity GB	Internal Volume Capacity GB	Internal Volume Used Capacity GB	Internal Volume Protection Capacity GB	Handled
ORG	EVA01C		48,615.04	25,047.54	18,854	18,854	0	18,854	0	0	0	0	0	0	
N/A	opaparray01c,opaparray02c		79,181.41	2,751.8	57,295.13	9,972.57	47,322.56	12,460.19	0	0	0	21,628.5	8,613.63	0	
	opaparray01n,opaparray02n		158,362.81	5,503.6	114,590.27	47,652.34	66,937.93	56,693.92	0	146.22	0	51,521.15	42,676.96	0	

Report columns

The report includes these columns:

Data Center

Location of the storage array.

Storage Name

Name of storage array.

Backend/Virtual

Indication of virtualization. Identifies whether the storage acts as virtual storage (V), backend storage (B), or both virtual and backend storage (BV). If virtualization is not associated with the host, the column in the row for that host is displayed as empty.

Raw Capacity

The total physical capacity of all disks in the array.

Unconfigured Raw Capacity

Physical capacity of disks that are not configured into a storage pool, as well as spare and failed disks.

Storage Pool Capacity

The size of the storage pool that is available for storing user data.

Storage Pool Used Capacity

The amount of usable capacity in use on the storage pool. This includes the capacity consumed by internal volumes, volumes, Snapshot Reserve, and other storage efficiency technology overheads.

Storage Pool Unused Capacity

The usable capacity that might be available for storing additional data on the storage pool.

Allocated Volume Capacity

The capacity of volumes accessed by hosts, used for replication and volumes that are masked but no access path to them exists.

Unallocated Volume Capacity

The capacity of volumes that are not allocated. These volumes are not used for any purpose.

Orphaned Volume Capacity

The capacity of volumes that are masked but no access path to them exists.

Protection Volume Capacity

The capacity of volumes that are used for replication.

Internal Volume Capacity

The usable capacity of the internal volumes.

Internal Volume Used Capacity

The used capacity of the internal volumes.

Internal Volume Protection Capacity

The provisioned capacity of internal volumes that are used for replication.

Manufacturer

Name of manufacturer of the storage array.

Family

Product classification of the storage array (for example, CLARiiON or Symmetrix).

Model

Model number supplied by the manufacturer.

Microcode version

Version of firmware running on the storage array.

Serial Number

Serial number associated with the device.

VM Capacity Summary report

The VM Capacity Summary report shows capacity information that includes provisioned capacity and commit ratio by data stores. This report provides details about the virtual machine, including its total capacity, provisioned capacity, and commit ratio.

Accessing this report

To display this report, select **Public Folders > Capacity > VM Capacity > Reports > VM Capacity Summary**.

Report generation options

For this report, you select the data store names to be represented.

Sample

VM Capacity Summary					
Datastore	Actual Capacity (GB)	VM Name	ESX Host	Provisioned Capacity (GB)	Commit Ratio
ds-0	97.66	N/A	NtapESX-1	0.00	0.00%
		vm-0	NtapESX-1	0.98	1.00%
		vm-1	NtapESX-1	1.95	2.00%
ds-0	97.66			2.93	3.00%
ds-1	97.66	N/A	NtapESX-1	0.00	0.00%
		vm-2	NtapESX-1	0.98	1.00%
		vm-3	NtapESX-1	1.95	2.00%
ds-1	97.66			2.93	3.00%
ds-10	97.66	N/A	NtapESX-5	0.00	0.00%
		vm-20	NtapESX-5	0.98	1.00%
		vm-21	NtapESX-5	1.95	2.00%
ds-10	97.66			2.93	3.00%
ds-11	97.66	N/A	NtapESX-5	0.00	0.00%
		vm-22	NtapESX-5	0.98	1.00%
		vm-23	NtapESX-5	1.95	2.00%
ds-11	97.66			2.93	3.00%
ds-12	97.66	N/A	NtapESX-5	0.00	0.00%
		vm-24	NtapESX-5	0.98	1.00%
		vm-25	NtapESX-5	1.95	2.00%
ds-12	97.66			2.93	3.00%
ds-13	97.66	N/A	NtapESX-6	0.00	0.00%
		vm-26	NtapESX-6	0.98	1.00%
		vm-27	NtapESX-6	1.95	2.00%
ds-13	97.66			2.93	3.00%
ds-14	97.66	N/A	NtapESX-6	0.00	0.00%
		vm-28	NtapESX-6	0.98	1.00%

Interpretation

These definitions and formulas are presented in this report:

Datastore

Name of the data store or volume.

Actual Capacity

Total amount of space available in the data store.

VM Name

Name of the virtual machine or machines.

ESX Host

Name of the Hypervisor host.

Provisioned Capacity

Amount of space vCenter reports as per guest VMs. Provisioned capacity on the data store on the whole.

Commit Ratio

Ratio of data store provisioned capacity/Total capacity of the data store.

Creating custom ad hoc reports

You can use the report authoring tools to create custom reports. You can also save reports to run on a regular schedule and email results to yourself and others.

The example in this section shows the following processes, which can be used for any of the OnCommand Insight data models:

- Identifying a question to be answered with a report
- Determining the data needed to support the results
- Selecting data elements for the report

This section shows you how to create reports using different report authoring tools: Query Studio and Workspace Advanced.

Report authoring tools

The OnCommand Insight enterprise reporting data models provide data elements and interactive relationships among data elements that yield business views of the data. Using the data elements and relationships, you can create reports using Query Studio, Report Studio, Workspace or Workspace Advanced, which are report generation tools.

You can use the following report authoring tools to create reports with OnCommand Insight data models:

- With Query Studio, you drag and drop data elements to create customized ad hoc reports that address your business needs.

Note: Although the OnCommand Insight documentation provides basic information about creating reports using Query Studio, for details about Query Studio format options, see the Cognos documentation.

- With Report Studio, you can create more complex reports with complex filters, prompts, and conditions.
You can create very complex reports using the Report Studio tool and OnCommand Insight data. Report Studio is a powerful tool, and its use is beyond the scope of this guide.
- With Workspace, you can create and assemble dashboards and create interactive reports.
- With Workspace Advanced, you can create rich, interactive reports with a tool that has all the functionality of Query Studio, Report Studio Express, and Analysis Studio.

Some authoring tools are hidden for use only with advanced views.

What you need to do before you design your report

Before you design your custom report, you need to complete some prerequisite tasks. If you do not complete these, reports could be inaccurate or incomplete.

For example, if you do not finish the identification process, your capacity reports will not be accurate. Or, if you do not finish setting annotations (such as tiers, business units, and data centers), your custom reports will not show reports accurately across your domain and might show "N/A" for some data points.

Before you design your reports, complete the following tasks:

- Configure all data sources. For details, see the *OnCommand Insight Installation and Administration Guide*.
- Identify all devices in your environment. For details, see the identification chapters in the *OnCommand Insight Inventory User Guide*.
- Enter as many annotations (such as tiers, data centers, and business units) on devices and resources in your environment as possible. It is beneficial to have annotations stable before generating reports, because OnCommand Insight Data Warehouse collects historical information.
- Configure OnCommand Insight Data Warehouse to accept the data from the OnCommand Insight server in the Extract, Transform, and Load (ETL) process.

Process of creating reports

The process of creating ad hoc reports involves several tasks.

- Plan the results of your report.
- Identify data to support your results.
- Select the data model (for example, Chargeback data model, Inventory data model, and so on) that contains the data.
- Select data elements for the report.
- Optionally format, sort, and filter report results.

How to plan the results of your custom report

Before you open the Query Studio or Business Insight Advanced report design tools, you might want to plan the results you want from the report. With these report authoring tools, you can create reports easily and might not need a great deal of planning; however, it is a good idea to get a sense from the report requestor about the report requirements.

- Identify the exact question you want to answer.
- Identify the data elements that you need to support the answer.
- Identify the relationships between data that you want to see in the answer. Do not include illogical relationships in your question, for example, “I want to see the ports that relate to capacity.”
- Identify any calculations needed on data.
- Determine what types of filtering are needed to limit the results.
- Identify how the report will be distributed. For example, should it be emailed on a set schedule, located on a dashboard, or included in a OnCommand Insight Reporting Connection My Folders area?
- Determine who will maintain the report. This might affect the complexity of the design.
- Create a mockup of the report.

Here are some examples of questions you can answer by creating your custom reports and the OnCommand Insight Data Warehouse model:

- How much capacity do I have left?
- What are the chargeback costs per business unit?
- What is the capacity by tier to ensure that business units are aligned at the proper tier of storage?

- How can I forecast power and cooling requirements? (Add customized metadata by adding annotations to resources.)

Tips for designing reports

Several tips might be helpful when you are designing reports.

- Determine whether you need to use current or historical data.
Most reports only need to report on the latest data available in the Data Warehouse.
- Data Warehouse provides historical information on capacity and performance, but not on inventory.
- Everybody sees all data; however, you might need to limit data to specific audiences.
To segment the information for different users, you can create reports and set access permissions on them.

Creating your first Query Studio report

The most important part of the of the report design is determining the question you want to answer with your report and selecting data to support that result.

This example creates a report that answers the following question: Which data center is using the most capacity in the top level (most expensive) tier?

This is the report that this example wants to achieve.

Business Unit	Tier	Provisioned Raw Capacity GB	Tier Cost
Customer Support	Gold	48.00	10
Dev	Gold	432.00	10
Dev	Gold-Fast	2,608.00	12
Engineering	Gold	80.00	10
Legal	Bronze	64.00	3
Legal	Gold	128.00	10
Legal	Gold-Fast	3,936.00	12
Marketing	Gold	144.00	10
Marketing	Gold-Fast	8,080.00	12
Marketing	Silver	336.00	7
N/A	Bronze	1,884.61	3
N/A	Gold	4,385.56	10
N/A	Gold-Fast	30,935.81	12
N/A	Silver	3,870.34	7
Operations	Gold-Fast	16.00	12
PM	Gold	48.00	10
Stores	Bronze	560.00	3
Stores	Gold	1,430.37	10
Stores	Gold-Fast	128.00	12
Trading	Bronze	240.00	3

To design this report, you use the Storage Capacity data model from the Advanced data mart including the following data elements:

- Business Unit (or if you want the entire Business Entity information, you could select Business Entity)
- Tier
- Physical Raw Capacity
- Tier Cost

You might want to enhance the report by doing the following:

- Group by business unit.
- Sort the data by capacity in descending order.
- Filter any business units that are not needed.

Accessing Query Studio

You must open the OnCommand Insight Reporting Portal and access Query Studio to create sample reports in all data models.

About this task

Use these steps to access Query Studio for any specific data model.

Steps

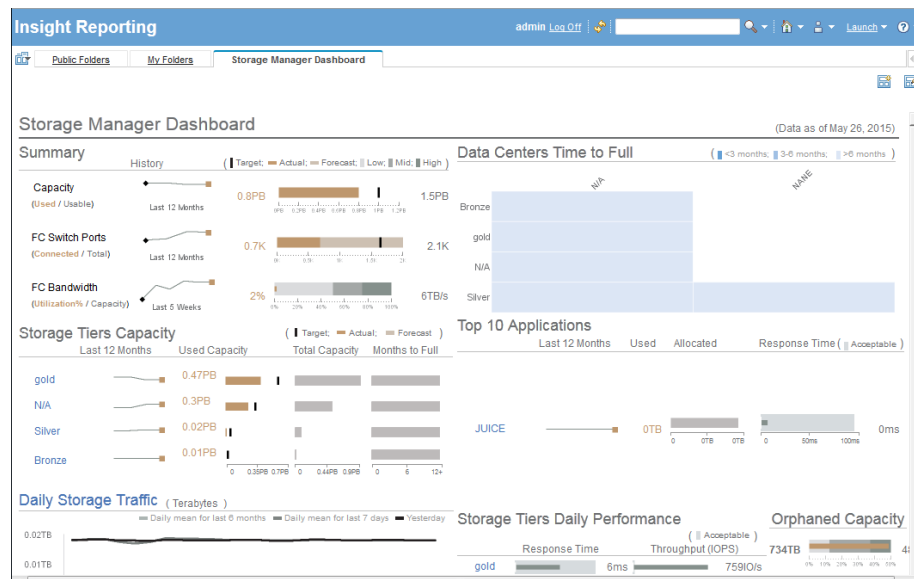
1. Open a web browser and enter the following URL:

http://<host-name>/reporting

2. Enter your user name and password, and then click **OK**.

This gives you access to the OnCommand Insight Reporting Portal. The Portal can be accessed in several ways. For other methods, see information about accessing the OnCommand Insight Reporting Portal.

3. Select **Insight Reporting Content** to open the OnCommand Insight Reporting Connection.



4. Click the link for the data model you want to access.
5. When the next view opens, from the Launch Menu on the upper right-hand corner, select **Query Studio**.

The Query Studio opens and provides access to the data model you selected.

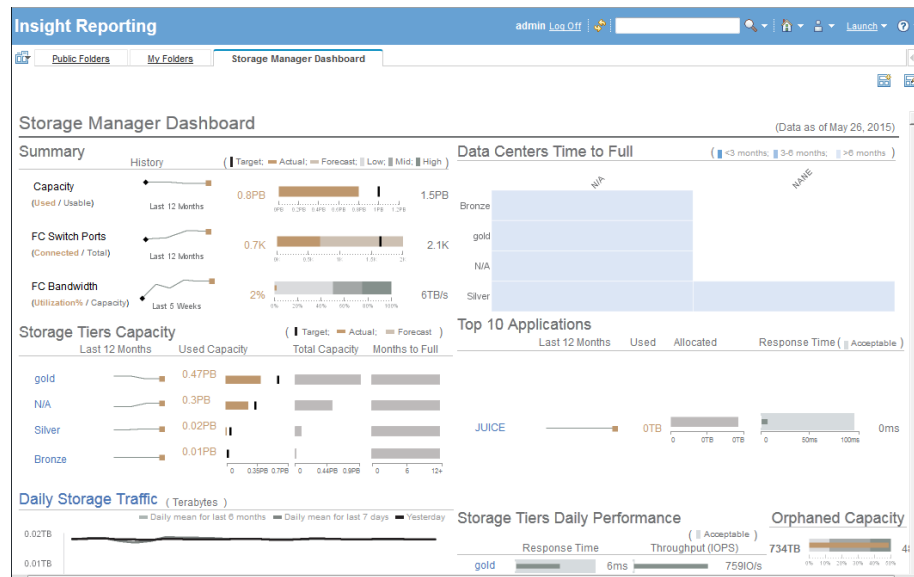
Selecting the data model and data elements for your report

The OnCommand Insight Reporting Portal provides several data models that focus the area of your report. Each data model appears as a separate folder tab.

Steps


1. From the OnCommand Insight Reporting Portal, select **Insight Reporting Content**.

The list of data models appears. If you do not have a license for the specific tabs, a warning icon appears on the tab. If you login as an administrator, you see folders instead of tabs.



2. Click on the link of the data model for which you are running the report. For this example, click the **Capacity** folder, then the **Storage Capacity** link.
3. From the Launch menu on the upper right corner, select a report authoring tool, such as **Query Studio**.

The report authoring tool opens and provides access to the data model you selected.



Filtering and sorting data

Filter data to limit the results and sort the data to show capacity in descending order.

Steps

1. Highlight the **Business Unit** column. Click the **Filter** icon and remove "N/A" from the list.

2. Highlight the **Provisioned Raw Capacity** column and click on the **Sort** icon. Arranging in descending order.

Pivoting the report

You can click on a column and create a crosstab report by pivoting the report. When you do this, the selected column values become the values across the top of the report.

Steps

1. To create the crosstab report, select the column that will become the values across the top.
2. Click the **Pivot** icon.

Here is a sample report before pivoting.

Business Unit	Tier	Provisioned Raw Capacity GB	Tier Cost
Customer Support	Gold	48.00	10
Dev	Gold	432.00	10
Dev	Gold-Fast	2,608.00	12
Engineering	Gold	80.00	10
Legal	Bronze	64.00	3
Legal	Gold	128.00	10
Legal	Gold-Fast	3,936.00	12
Marketing	Gold	144.00	10
Marketing	Gold-Fast	8,080.00	12
Marketing	Silver	336.00	7
N/A	Bronze	1,894.61	3
N/A	Gold	4,385.56	10
N/A	Gold-Fast	30,935.81	12
N/A	Silver	3,870.34	7
Operations	Gold-Fast	16.00	12
PM	Gold	48.00	10
Stores	Bronze	560.00	3
Stores	Gold	1,430.37	10
Stores	Gold-Fast	128.00	12
Trading	Bronze	240.00	3

Here is a sample of the same report after pivoting.

	Provisioned Raw Capacity GB	Customer Support	Dev	Engineering	Legal	Marketing	N/A	Operations	PM	Stores	Trading	Summary
Gold	10	48.00	432.00	80.00	128.00	144.00	4,385.56	16.00	48.00	1,430.37	192.00	6,887.93
Gold	Gold	48.00	432.00	80.00	128.00	144.00	4,385.56	16.00	48.00	1,430.37	192.00	6,887.93
Gold-Fast	12		2,608.00		3,936.00	8,080.00	30,935.81	16.00		128.00	8,736.00	54,439.81
Gold-Fast	Gold-Fast		2,608.00		3,936.00	8,080.00	30,935.81	16.00		128.00	8,736.00	54,439.81
Bronze	3				64.00		1,894.61			560.00	240.00	2,748.61
Bronze	Bronze				64.00		1,894.61			560.00	240.00	2,748.61
Silver	7						336.00	3,870.34				4,206.34
Silver	Silver						336.00	3,870.34				4,206.34
Summary		48.00	3,040.00	80.00	4,128.00	8,560.00	41,076.32	16.00	48.00	2,118.37	9,168.00	68,282.69

Saving your report

You should save your reports in a folder other than the My Folders folder. When reports are backed up and restored in OnCommand Insight, the My Folders content is not included in the backup.

Steps

1. From Query Studio, click on the **Save** icon.
2. Select a folder other than the My Folders folder.

3. Enter a name for your report.

Selecting another data model for more reporting

When you want to create reports using another data model, you can exit the current data model and select another.

Steps

1. From Query Studio, click the **Home** icon in the upper-right corner.
2. From the Launch menu in the upper-right corner, select **Query Studio**.

Returning to OnCommand Insight Data Warehouse

When you have finished using Query Studio, you can return to Data Warehouse easily.

Steps

1. From the upper right side of Query Studio, click on the **Return** arrow.
If you are creating a report, a message appears at the bottom asking if you want to save the report.
2. To save the report, click **Yes**.

Creating reports with Workspace Advanced

You can create custom reports using Workspace Advanced, a report authoring tool.

About this task

You can create the following sample reports:

- Report sample 1: Raw Capacity by Data Center (table)
- Report sample 2: Raw Capacity by Data Center (chart added to first example)

Creating a sample report: Raw Capacity by Data Center (table)

This example shows how to create a basic table report using OnCommand Insight Workspace Advanced.

About this task

This report shows the distribution of raw capacity by data center. The report is created using the Storage and Storage Pool Capacity data model.

Data Center	Raw Capacity (TB)
New York	14,116.39
Tokyo	4,804.66
Overall - Summary	18,921.05

Steps

1. Access the OnCommand Insight Reporting Connection using this URL: `https://<reporting-server>/reporting`

2. From the Launch menu (at the top right corner of the Storage Manager Dashboard), click **Workspace Advanced**.
3. From the Public Folders list, click **Capacity > Storage and Storage Pool Capacity**.
4. In the Workspace Advanced landing page, click **Create new**.
5. From the pre-defined report layouts New pop-up, choose **List** and click **OK**.
6. In the right pane, click **Source** tab.
7. In the tree, expand **Storage and Storage Pool Capacity**.
8. Expand **Data Center** to see a child Data Center.
9. Drag the child **Data Center** to the list.
10. Expand **Capacity (TB)** and drag **Raw Capacity (TB)** into the report area to the right of the Data Center column.
11. Click **Save** and save the report in the **Customer Reports** folder.

Creating a sample report: Raw Capacity by Data Center (chart)


This example shows you how to add a chart to a basic table report using OnCommand Insight Workspace Advanced.

About this task

This report shows the distribution of raw capacity by data center with a chart. The report is created using the Storage and Storage Pool Capacity data model.

Steps

1. Use the sample Raw Capacity by Data Center report as a base.
2. To create a page layout that includes both the List and Chart, select the Page Layout drop-down menu and choose the two-column layout.



The report page updates showing two columns.

3. In the right pane, click the **Toolbox** tab.
4. Select **Chart** and place it over the right column work area.
5. From the Insert Chart dialog box, select **Pie, Donut**, (use the pie chart selected by default), and then click **OK**.
6. In the right pane, select the **Source** tab.
7. In the tree, expand **Storage and Storage Pool Capacity > Physical Capacity > Capacity (TB) > Raw Capacity (TB)**.
8. Drag **Raw Capacity (TB)** and place it on the Default Measure section of the chart.

9. From the Physical Capacity folder, drag the child **Data Center** element and place it on the Series (pie-slices) section of the chart.
10. Click **Save** and save the report in the **Customer Reports** folder.

Importing and exporting report definitions in Report Studio

In OnCommand Insight Reporting, the Report Studio application provides an easy method for importing and exporting report definitions in an XML format.

This allows you to add new reports posted in NetApp forums to your OnCommand Insight Reporting installation. In addition, this feature is useful to back up report definitions and exchange reports with colleagues.

To create these reports, you must have a license for the OnCommand Insight Plan module.

You do not have to specify a package when importing XML. The package is specified as part of the XML.

Reporting data model

Your enterprise can benefit from the data that is discovered and stored in OnCommand Insight Data Warehouse. The OnCommand Insight Data Warehouse is a centralized repository that stores data from multiple information sources and transforms them into a common, multidimensional data model for efficient querying and analysis.

From this repository, you can generate custom reports such as chargeback, consumption analysis, and forecasting reports that answer questions such as the following:

- What inventory do I have?
- Where is my inventory?
- Who is using our assets?
- What is the chargeback for allocated storage for a business unit?
- How much headroom do I have on switch ports?
- How long until I need to acquire additional storage capacity?
- Are business units aligned along the proper storage tiers?
- How is storage allocation changing over a month, quarter, or year?

Using the data model provided with OnCommand Insight Reporting, you can use report authoring tools, such as Query Studio, Report Studio, Business Insight, or Business Insight Advanced, to design and schedule reports. You can access the report authoring tools from the OnCommand Insight Reporting portal.

Data model overview

OnCommand Insight provides several data models for use in report development. Each data model is an aggregation that summarizes data so that it can be queried and searched. For example, reports about capacity planning use the Capacity data model.

The OnCommand Insight enterprise reporting data models provide data elements and interactive relationships among data elements that yield business views of the data. Using the data elements and relationships, you can create reports using Query Studio, Report Studio, Business Insight or Business Insight Advanced, which are report generation tools.

OnCommand Insight also provides data marts that can be used to develop your own SQL queries. There is a distinction between these SQL query data marts and the data models used in reporting. The individual OnCommand Insight reporting data models use the underlying OnCommand Insight database schema provided in the data marts; however, the data models use additional tables and sometimes new elements in the tables. For instance, the data model includes a Monthly Capacity Fact table in the Storage Capacity data model that is based on the Capacity Fact table from the database schema and data mart. The data model filters out the values from the database schema table to show only month information.

Another example of a difference between the database schema used in data marts and the data model is in the Violation table and the Violation Type column. The data model translates programmatic-named values in the database to match the text displayed in the OnCommand Insight Client.

Related concepts

[Report authoring tools](#) on page 42
[Chargeback data model](#) on page 77
[File System Utilization data model](#) on page 54
[Internal Volume Capacity data model](#) on page 55
[Inventory data model](#) on page 81
[Performance data models](#) on page 86
[Qtree Capacity data model](#) on page 64
[Storage and Storage Pool Capacity data model](#) on page 67
[Storage Efficiency data model](#) on page 92
[VM Capacity data model](#) on page 70
[Volume Capacity data model](#) on page 75

OnCommand Insight data model descriptions

OnCommand Insight includes several data models from which you can select predefined reports or create your own custom report.

Capacity data model

Enables you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, and VM capacity. The Capacity data model is a container for several capacity data models:

Storage and Storage Pool Capacity data model

Enables you to answer questions about storage capacity resource planning, including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model, because it is a simpler, targeted data model.

File System Utilization data model

Enables you to answer questions about file system utilization. This data model provides visibility into capacity utilization by hosts on the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type.

Internal Volume Capacity data model

Enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time.

Port Capacity data model

Enables you to answer questions about switch port connectivity, port status, and speed over time. This data can be used when planning for purchases of new SAN switches.

Qtree Capacity data model

Enables you to trend qtree utilization (with data such as used versus allocated capacity) over time. You can also see the information by different dimensions, for example, by business entity, application, tier, and service level.

VM Capacity data model

Enables you to report your virtual machine environment and its capacity usage. This data model lets you report on changes in capacity usage over time for virtual machines and data stores. The data model also provides thin provisioning and virtual machine chargeback data.

Volume Capacity data model

Enables you to analyze all aspects of volumes in your environment and look at them by vendor, model, tier, service level, and data center. You can view capacity related to orphan volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technology (iSCSI or FC) and compare virtual to non-virtual volumes for array virtualization issues.

Chargeback data model

Enables you to answer questions about used capacity and allocated capacity on storage resources (volumes, internal volumes, and qtrees). This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities, and includes both current and historical data. Report data can be categorized by service level and storage tier.

You can use this data model to generate chargeback reports by finding the amount of capacity used by a business entity. This data model enables you to create unified reporting of multiple protocols (including NAS, SAN, FC, and iSCSI).

Inventory data model

Enables you to answer questions about inventory resources such as hosts, storage systems, switches, disks, tapes, qtrees, quotas, virtual machines and servers, and generic devices. The Inventory data model includes several submarts that enable you to see information about replications, FC paths, iSCSI paths, NFS paths, and violations. The Inventory data model does not include historical data.

Performance data model

Enables you to answer questions about performance for volumes, application volumes, internal volumes, switches, applications, VMs, ESX vs VM, hosts, and application nodes.

The Performance data model includes information that helps you determine appropriateness of tiers, storage misconfigurations for applications, and volume and internal volume last access times. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

Storage Efficiency data model

Enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount used or consumed (the physical measurement). For example, when thin provisioning is enabled, OnCommand Insight indicates how much capacity is actually taken from the device. You can also use this model to determine efficiency when deduplication is enabled.

Data model fact and dimension tables

Each data model includes both fact and dimension tables.

- Fact tables: Contain data that is measured, for example, quantity, raw and usable capacity. Contain foreign keys to dimension tables.
- Dimension tables: Contain descriptive information about facts, for example, data center and business units. A dimension is a structure, often composed of hierarchies, that categorizes data. Dimensional attributes help describe the dimensional values.

Using different or multiple dimension attributes (seen as columns in the reports), you construct reports that access data for each dimension described in the data model.

For explanations of all data elements used in creating reports, see the Data Glossary.

Colors used in data model elements

Colors on data model elements have different indications.

- Yellow assets: Represent measurements.
- Non-yellow assets: Represent attributes. These values do not aggregate.

Using multiple data models in one report

Typically, you use one data model per report. However, you can write a report that combines data from multiple data models.

To write a report that contains only one data model, choose the data model and use one of two report authoring tools, Query Studio or Report Studio.

To write a report that combines data from multiple data models, choose one of the data models to use as the base, then write SQL queries to access the data from the additional data marts. You can use the SQL Join feature to combine the data from the different queries into a single query that you can use to write the report.

For example, say you want the current capacity for each storage array and you want to capture custom annotations on the arrays. You could create the report using the Storage Capacity data model. You could use the elements from the Current Capacity and dimension tables and add a separate SQL query to access the annotations information in the Inventory data model. Finally, you could combine the data by linking the Inventory storage data to the Storage Dimension table using the storage name and the join criteria.

Capacity data models

Capacity data models enable you to answer questions about storage capacity, file system utilization, internal volume capacity, port capacity, qtree capacity, storage pools, and VM capacity. The Capacity data model is a container for several capacity data models.

- Storage and Storage Pool Capacity data model. If you are new to capacity reporting, you should start with this data model, because it is a simpler, targeted data model.
- File System Utilization data model
- Internal Volume Capacity data model
- Port Capacity data model
- Qtree Capacity data model
- Storage Capacity data model

- VM Capacity data model
- Volume Capacity data model

File System Utilization data model

The File System Utilization data model enables you to answer questions about file system utilization. This data model provides visibility into capacity utilization by hosts on the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type.

File systems make use of an underlying data storage device that offers access to an array of fixed-size blocks, for example, a local hard drive, SAN LUN, or SAN volume. For each file system, you can determine the type of file system, for example, NTFS or FAT32.

You can answer the following types of questions using the File System Utilization data model:

- What is the capacity (allocated and used) per FS?
- Where is the data kept and how is it accessed, for example, local, SAN, or NAS?
- What type of file systems (for example, NTFS, FAT, UFS, VxFS, and ISO 9660) do we have and how many of each?
- What are the trending statistics for the file system type and capacity?

FS Utilization data model details

The File System Utilization data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced: Provides all values and details available from the simple data mart, including access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[FS data](#) on page 54

[FS advanced data mart](#) on page 55

FS data

This data model provides visibility into capacity utilization by hosts on the file system level. Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type.

The fields and field descriptions are as follows:

Allocated Capacity

See "Capacity concepts" information.

File System Name

Mounting point or drive letter.

Hardware ID

Host-specific identifier.

Host name

Representative host accessing the capacity.

Host OS

Operating system running on the host, if available to OnCommand Insight.

Host IP

IP address assigned to the host device.

Location

Indicates where the data is kept and how it is accessed, for example, local, SAN, or NAS.

Manufacturer

Name of manufacturer that makes the device.

Model

Manufacturer's model number.

Type

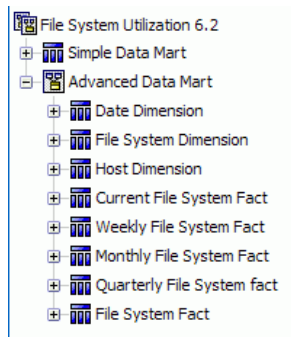
File system type, for example, NTFS, FAT, UFS, UDF, ZFS, ISO 9660, or VXFS.

Used Capacity

See "Capacity concepts" information.

FS advanced data mart

The FS advanced data mart contains various fact and dimension tables.



- **Date Dimension:** Includes dates and time periods from which you can select for the report. Each capacity fact is associated with the date when the Data Warehouse build ran that generated the fact.
- **File System Dimension:** Provides attributes (such as FS name, location, and type).
- **Host Dimension:** A representative host from the group of hosts using the capacity.
- **File System fact:** Includes the capacity data for the FS.

Related concepts

[Date Dimension](#) on page 141

[File System Dimension](#) on page 143

[Host Dimension](#) on page 144

Internal Volume Capacity data model

The Internal Volume Capacity data model enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time.

You can answer the following types of questions using the Internal Volume Capacity data model:

- Which internal volumes have a utilization higher than a predefined threshold?
- Which internal volumes are in danger of running out of capacity based on a trend?

- What is the used capacity versus the allocated capacity on our internal volumes?

Creating a sample report with the Internal Volume Capacity data model

This report shows internal volume capacity per data center and vendor allocation by internal volume, storage and tier.

About this task

The following illustration shows the report that you will create with these instructions, although your data will be different.

Internal Volume Capacity Allocation				
by Internal Volume, Storage and Tier				
Storage Name	Internal Volume Name	Tier Name	Allocated Capacity GB	Used Capacity GB
Chicago	NetApp-8:0:VOL0	Bronze	183.37	7.00
Chicago	NetApp-8:0:VOL1	Bronze	183.37	7.00
Chicago	NetApp-8:0:VOL2	Bronze	183.37	7.00
Chicago	NetApp-8:1:VOL0	Bronze	183.37	7.00
Chicago	NetApp-8:1:VOL1	Bronze	183.37	7.00
Chicago	NetApp-8:1:VOL2	Bronze	183.37	7.00
Chicago	NetApp-8:2:VOL0	Bronze	183.37	7.00
Chicago	NetApp-8:2:VOL1	Bronze	183.37	7.00
Chicago	NetApp-8:2:VOL2	Bronze	183.37	7.00
Dallas	NetApp-9:0:VOL0	Bronze	183.37	7.00
Dallas	NetApp-9:0:VOL1	Bronze	183.37	7.00
Dallas	NetApp-9:0:VOL2	Bronze	183.37	7.00
Dallas	NetApp-9:1:VOL0	Bronze	183.37	7.00
Dallas	NetApp-9:1:VOL1	Bronze	183.37	7.00
Dallas	NetApp-9:1:VOL2	Bronze	183.37	7.00
Dallas	NetApp-9:2:VOL0	Bronze	183.37	7.00
Dallas	NetApp-9:2:VOL1	Bronze	183.37	7.00
Dallas	NetApp-9:2:VOL2	Bronze	183.37	7.00
Detroit	NetApp-10:0:VOL0	Bronze	183.37	7.00
Detroit	NetApp-10:0:VOL1	Bronze	183.37	7.00

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. From a blank screen in Query Studio showing the Internal Volume Capacity data model, select **Advanced Data Mart**.
3. Select **Storage Dimension** > **Storage Name** and drag it into the work area.
4. Select **Internal Volume Dimension** > **Internal Volume Name** and drag it into the work area.
5. Select **Tier Dimension** > **Tier Name** and drag it into the work area.
6. Select **Internal Volume Capacity Fact** > **Allocated Capacity GB** and drag it into the work area.
7. Select **Internal Volume Capacity Fact** > **Used Capacity GB** and drag it into the work area.
8. Optionally, group the report by Storage Name by clicking the Storage name column and then the **Group** icon in the upper right.

Related tasks

[Accessing Query Studio](#) on page 45

Internal Volume Capacity data model details

The Internal Volume Capacity data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced: Provides access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[Internal Volume Capacity data](#) on page 57

[Internal Volume Capacity advanced data mart](#) on page 59

Internal Volume Capacity data

This data model enables you to answer questions about internal volume used capacity, allocated capacity, and capacity usage over time.

The fields and field descriptions are as follows:

Application

Applications such as your email, finance, development or human resources applications.

Application URL

URL of the application.

Internal Volume

The logical layer in storage devices that is not externally accessible as a SAN volume or NAS share. An internal volume is carved from a storage pool and exposed to hosts as a share or LUN.

Internal Volume URL

URL of the internal volume.

Allocated Capacity

The allocated size of the internal volume.

Business Unit

The business unit for the representative application of all applications accessing this internal volume.

Consumed Capacity

The amount of capacity taken from the underlying storage pool for this internal volume after applying storage efficiency technology overhead. Consumed capacity is less than or equal to used capacity, depending on the storage efficiency technology applied to this internal volume.

Data Unused Capacity

Capacity in the internal volume that is available for writing NAS data or allocating volumes.

Data Used Capacity

Capacity in the internal volume that is used by hosts (volumes or NAS shares).

Deduplication Ratio

A number between 0 and 1 that represents the deduplication compression ratio applicable to the capacity measured in this fact.

Note: The OnCommand Insight Client includes a Deduplication Savings column on the Internal Volumes detail view. This represents the rate of savings in the internal volume.

For example, 1/3 of capacity is saved by using deduplication. The Data Warehouse and reports show a Deduplication Ratio data element. This represents the rate of compression that occurs from using deduplication, complementary data related to Deduplication Savings. In the example where Deduplication Savings is 1/3, the Deduplication Ratio is 2/3.

Deduplication Savings

Rate of savings in the internal volume by using deduplication.

Nondeduplication Capacity

Capacity on the internal volume that does not have deduplication technology applied to it. Estimation of the consumed capacity assuming deduplication was turned off.

$(\text{ConsumedCapacity}) \div (\text{Dedupe Ratio})$

Number of Snapshots

A count of how many Snapshot copies were created on the internal volume.

Protection Type

Type of protection used on the storage, for example, Mirrored or Striped.

Raw Capacity

The physical disk capacity. This differs from usable capacity when technologies such as RAID are used, where some of the raw capacity is used for protection purposes.

Raw Allocated Capacity

The raw allocated size of the internal volume.

Raw Data Used Capacity

Raw capacity in the internal volume that is used by hosts (volumes or NAS shares).

Raw Consumed Capacity

The raw amount of capacity taken from the underlying storage pool for this internal volume.

Raw Data Unused Capacity

Raw capacity in the internal volume that is available for writing NAS data or allocating volumes.

Raw Total Clone Saved Capacity

Raw capacity on cloned internal volumes that is shared with the source of the clone.

Raw Snapshot Allocated Capacity

Raw capacity on the internal volume that is reserved for Snapshot copies.

Raw Snapshot Unused Capacity

Raw capacity remaining after some capacity that was reserved for Snapshot copies was not used. This excludes capacity that was used for protection purposes.

Raw Snapshot Used Capacity

Raw capacity that has been used for Snapshot copies.

Snapshot Allocated Capacity

Capacity on the internal volume that is reserved for Snapshot copies.

Snapshot Unused Capacity

Capacity remaining after some capacity that was reserved for Snapshot copies was not used. This excludes capacity that was used for protection purposes.

Snapshot Used Capacity

Capacity that has been used for Snapshot copies. Can exceed Snapshot Allocated Capacity in the case of Snapshot Overflow.

Service Level

Service level represents the agreement between a service consumer or application and the storage service provider, associated with an application. Service level differs from tier.

For example, you might have multi-tiered storage pools including 1 FC pool and one SSD pool. Your service level policy might require that Tier 1 should sustain 5,000 IOPS.

However, because this LUN uses only 100 IOPS, the service provider might decide to move this service to a lower tier that sustains only 500 IOPS. The service level is 5,000 or 500 IOPS, while the tier represents the quality of service, for example, gold or silver.

Service Level Cost

Cost / GB is the cost per gigabyte for a service level.

The service level cost can be the cost of providing the service or the actual price to service the consumer, while the tier cost is the average cost per GB for the specific tier.

Storage Access Type

Indication of the type of access to the storage, SAN or NAS.

Tier Name

Level of service (for example, gold, silver, or bronze).

Tier Cost

Unit of cost per gigabyte assigned to the tier level.

Time of Last Snapshot

The date and time when the most recent Snapshot copy was created.

Total Clone Saved Capacity

Capacity on cloned internal volumes that is shared with the source of the clone.

Used Capacity

Capacity in the internal volume used by hosts (volumes and NAS shares) including Snapshot Used Capacity. Used Capacity is no less than Consumed Capacity, depending on the storage efficiency technology applied to this internal volume.

Internal Volume Capacity advanced data mart

The Advanced data mart contains the full table data structure, which contains all current and historical data.

The Advanced data mart is organized into the following fact and dimension tables:

- Application Dimension
- Application Group Dimension
- Business Entity Hierarchy
- Current Internal Volume Capacity Fact: Includes only the current capacity values. This table is used most often, for writing reports that reflect the current state of the storage environment (as of the last Data Warehouse build).
- Internal Volume Capacity Fact
- Weekly Internal Volume Capacity Fact
- Monthly Internal Volume Capacity Fact
- Quarterly Internal Volume Capacity Fact
- Date Dimension: Includes dates and time periods from which you can select for the report. Each capacity fact is associated with the date when the Data Warehouse build ran that generated the fact.

- **Internal Volume Dimension:** Defines the internal volumes, which are the logical layers in storage devices that are not externally accessible as SAN volumes or NAS shares.
- **Service Level Dimension**
- **Storage Dimension:** Provides attributes (such as storage name, IP, model, manufacturer, family, and type). Also includes data center.
- **Storage Pool Dimension:** Defines the storage pools from which the capacity is taken.
- **Tier Dimension:** Defines the tier based on the tier annotations in OnCommand Insight.
- **Virtual Storage Dimension:** Includes the name, data center, IP, and other information about virtual front-end storage for back-end capacity in storage virtualization.

Related concepts

[Application Dimension](#) on page 135

[Application Group Dimension](#) on page 136

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Service Level Dimension](#) on page 146

[Storage Dimension](#) on page 146

[Storage Pool Dimension](#) on page 147

[Tier Dimension](#) on page 151

[Time period fact tables](#) on page 152

Port Capacity data model

The Port Capacity data model enables you to answer questions about switch port connectivity, port status, and speed over time. This data can be used when planning for purchases of new SAN switches.

The Port Capacity data model:

- Provides easy access to switch capacity information recorded in the OnCommand Insight Data Warehouse.
- Provides access to information about switches, switch ports and port connectivity.
- Enables the creation of detailed and accurate port capacity and connectivity reports based on various criteria, such as data center, port speed, switch vendor, and connectivity type.
- Maintains historical data to generate reports for past periods (last week, last month, last quarter, and so on)
- Allows the creation of capacity trend reports using historical data.

You can answer the following types of questions using the Port Capacity data model:

- What types of ports are needed by port speed or vendor?
- In what location is this capacity needed?
- What is the optimal time to purchase that capacity and make it available?
- How much more capacity exists on the switch ports?

Creating a sample report with the Port Capacity data model

The report generated in this example will answer the question: How do you get an active view of ports in your environment?

About this task

The following example shows the process, and it can be used for any of the tables described in the Port Capacity data model.

Ports per Data Center		
Data Center	Port Count	Port Speed
New York	170	1G
New York	32	4G
Sydney	176	1G
Tokyo	202	1G
Zurich	272	1G
Summary	852	

Steps

1. Access Query Studio.

For detailed steps, see the related links about accessing Query Studio.

2. In Query Studio, expand the **Simple Data Mart**.
3. Select the **Port Count** element and drag it to the work area.
4. Select the **Port Speed** element and drag it into the work area.
5. Select the **Data Center** element and drag that into the work area.

The result is a table of values listing all ports and their speeds according to the Data Center location.

6. Optionally, to filter out any missing values or "N/A" values, click on the **Data Center** column and click **Filter**.
7. To group the data, click on the **Data Center** column and click the **Group** icon.
8. Optionally, click on the **Pivot** icon to create a crosstab format for the report.

In this example, you can see how many active ports each data center has, which ones are running at 1G or 4G, and their totals. This information provides a clear and rapid view into resource allocation.

Related tasks

[Accessing Query Studio](#) on page 45

Port Capacity data model details

The Port Capacity data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.

- Advanced: Provides access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[Port Capacity data](#) on page 62

[Port Capacity advanced data mart](#) on page 63

[Time period fact tables](#) on page 152

Port Capacity data

This data model enables you to answer questions about switch port connectivity, port status, and speed over time.

The fields and field descriptions are as follows:

Connected to Firmware

Firmware running on the port to which the port identified by the Name field is physically connected.

Connected to Name

Device and port to which the port identified by the Name field is physically connected.

Connected to Type

Protocol to which the port identified by the Name field is physically connected.

Data Center

Location of data center, based on the switch's data center annotation in OnCommand Insight.

Free Port Count

Number of free ports on the switch.

GBIC Type

The type of Gigabit Interface Converter (GBIC) used for the port.

Generated

Real FC ports for which OnCommand Insight does not know their WWN. Generated ports for NPV switches or clustered Data ONTAP storage.

NPV

N-port virtualization (NPV) devices that OnCommand Insight detects in your environment. Applies to FC network only.

Port Connectivity Type

Type of connectivity.

Port Count

Number of ports on the switch.

Port Fabric

Fabric on which the port is configured to operate.

Port Speed

Speed capability of the port, in gigabytes.

Port Status

Current status of the port, taken from the device itself. Possible values are OK, NOT LICENSED, ERROR, INACTIVE and UNKNOWN.

Port Type

Typically used to identify the protocol used to communicate with the other device, taken from the device itself: F-port, E-port, and so on.

Port WWN

WWN (World Wide Name) to identify the port.

Switch Firmware

Version of the firmware running on the switch.

Switch Level

Core, Host, Storage or User-defined (Annotation level assigned to the switch).

Switch Level Cost

Port cost associated with each switch level, as set in the OnCommand Insight Client's OnCommand Insight Settings > Cost Accounting page.

Switch Manufacturer

Name of the manufacturer who makes the switch.

Switch Name

Switch name, as provided by the SAN manager at the switch console.

Port Capacity advanced data mart

The full table data structure is provided for the Port Capacity data model in the Advanced Data Mart, which contains all current and historical data.

The data mart includes the following dimensions and fact tables:

- **Connected To Dimension:** Information about the device the switch port is connected to (storage array / host / switch / tape / generic device).
- **Date Dimension:** Includes dates and time periods from which you can select for the report. Each port fact is associated with the date when the Data Warehouse build ran generated the fact.
- **Port Dimension:** Describes a switch port.
- **Switch Dimension:** Defines the switch owning the port.
- **Time period fact tables:** Provides historical data in your reports.
 - **Current Ports Capacity Fact:** Includes only the current port values. This table is used most often, for writing reports that reflect the current state of the storage environment (as of the last Data Warehouse build).
 - **Weekly Ports Capacity Fact**
 - **Monthly Ports Capacity Fact**
 - **Quarterly Ports Capacity Fact**
- **Ports Fact table:** Includes the full set of port data for every Data Warehouse build performed.

Related concepts

[*Connected to Dimension*](#) on page 141

[*Date Dimension*](#) on page 141

[*Port Dimension*](#) on page 145

[*Switch Dimension*](#) on page 148

[*Time period fact tables*](#) on page 152

Qtree Capacity data model

The Qtree Capacity data model enables you to track qtree utilization trends (with data such as used versus allocated capacity) over time. You can also organize the information by different dimensions, for example, by business entity, application, tier, and service level.

Quotas are typically set up by storage administrators on shares or qtrees. Space quotas that indicate the limits on the amount of usable space or the number of files that can be created on a NAS share or qtree. The following types of limits apply:

- **Hard limit:** The disk space limit imposed on this qtree by a quota. If this is exceeded, data cannot be written.
- **Soft limit:** The soft quota space limit that, if exceeded, issues warnings rather than rejects space requests.

OnCommand Insight displays the limit data:

- If a hard limit exists, but a soft limit does not, the soft limit receives the value of the hard limit, and vice versa.
- If both hard and soft limits are not specified, they use the value of the internal volume allocated capacity.

Qtree capacity facts are purged after a default time threshold. For details, see Data Warehouse processes.

You can answer the following types of questions using the Qtree Capacity data model:

- What is the used capacity for qtrees versus the limits set per application or business entity?
- What are the trends of our used and free capacity so that we can do capacity planning?
- Which business entities are using the most capacity?
- Which applications consume the most capacity?

Creating a sample report with the Qtree Capacity data model

This example creates a report that shows qtree information, its internal volume, and hard and soft limit capacities.

About this task

Qtree Capacities						
Business Unit	Storage	Qtree	Internal Volume	Used Capacity GB	Hard Limit Capacity GB	Soft Limit Capacity GB
Billing	NTAP2	ImageArchive	NTAP2:Vol6	350.00	1,400.00	700.00
Billing	NTAP2	ImageArchive	NTAP2:Vol8	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingPublications	NTAP1:Vol0	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingPublications	NTAP1:Vol2	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingPublications	NTAP1:Vol4	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingVideos	NTAP1:Vol0	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingVideos	NTAP1:Vol2	350.00	1,400.00	700.00
Marketing	NTAP1	MarketingVideos	NTAP1:Vol4	350.00	1,400.00	700.00

Steps

1. Access Query Studio.

For detailed steps, see the related links about accessing Query Studio.

2. From a blank screen in Query Studio showing the Qtree Capacity data model, expand **Advanced Data Mart**.

3. Expand **Business Entity Hierarchy**, select **Business Unit**, and drag it into the work area.
4. Expand **Storage Dimension**, select **Storage**, and drag it to the right of Business Unit in the work area.
5. Expand **Qtree Dimension**, select **Qtree**, and drag it into the work area.
6. Expand **Internal Volume Dimension**, select **Internal Volume**, and drag it into the work area.
7. Expand **Qtree Capacity Fact**, select **Used Capacity**, **Hard Limit Capacity**, and **Soft Limit Capacity** and drag them into the work area.
8. Optionally, group the report by Business Unit by clicking the Business Unit column and then the **Group** icon in the upper right.
9. Click the **Title** and enter the report title.

Related tasks

[Accessing Query Studio](#) on page 45

Qtree Capacity data model details

The Qtree Capacity data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced: Provides access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[Qtree Capacity data](#) on page 65

[Qtree Capacity advanced data mart](#) on page 66

Qtree Capacity data

This data model enables you to trend qtree utilization (with data such as used compared to allocated capacity) over time. You can also organize the information by different dimensions, for example, by business entity, application, tier, and service level.

The fields and field descriptions are as follows:

Application

Application annotation associated with this qtree.

Business Unit

Business unit annotation associated with this qtree.

Internal Volume

Name of the internal volume on which the qtree resides.

Missing Quota Limits

Indication of a quota limit that is not set. If limits have been set, "NO_LIMITS_MISSING" appears.

Protection Type

Type of protection used on the storage, for example, Mirrored or Striped.

Qtree

Name of qtree.

Service Level

Service level associated with the application.

Service Level Cost

Cost of providing the service or the actual price to service the consumer.

Storage

Physical storage where the capacity resides.

Storage Access Type

Indication of the type of access to the storage, SAN or NAS.

Storage Pool

Name of the storage pool on which the qtree resides.

Tier

Level of service (for example, Gold, Silver, Bronze).

Tier Cost

Unit of cost per gigabyte assigned to the tier level.

Virtual Storage

The front-end array in case this qtree is on virtual internal volume.

Hard Limit Capacity

Disk space limit imposed on this qtree by a quota. If this is exceeded, data cannot be written.

Soft Limit Capacity

Soft quota space limit that, if exceeded, issues warnings rather than rejects space requests.

Used Capacity

Capacity in the qtree including Snapshot Used Capacity. Used Capacity is no less than Consumed Capacity, depending on the storage efficiency technology applied to this resource.

See capacity concepts.

Qtree Capacity advanced data mart

The Advanced data mart contains the full table data structure, which contains all current and historical data.

The Advanced data mart is organized into the following fact and dimension tables:

- Qtree Capacity Fact table
- <Time> Qtree Capacity Fact tables
- Application Dimension
- Business Entity Hierarchy
- Date Dimension
- Internal Volume Dimension
- Qtree Dimension
- Service Level Dimension
- Storage Dimension
- Storage Pool Dimension
- Tier Dimension

- Virtual Storage Dimension

Qtree Capacity data model fact table

The Qtree Capacity Fact table shows information about the hard and soft limits on the qtree and includes historical information.

Time Qtree Capacity Fact table

Use the elements in these tables to obtain historical data in your reports. The information in these time period tables reflects a given time period only if the ETL runs in that period. For example, if OnCommand Insight stops running for the entire month of February, there won't be a date point for February.

The fact tables and table descriptions are as follows:

Current Qtree Capacity Fact

Qtree capacity fact data for only the last ETL run (the most recent date point).

Weekly Qtree Capacity Fact

Qtree capacity fact data for a single date point each week. Includes historical data. Contains the values from the Qtree capacity table, filtered to include only values associated with dates where the Is Week Representative equals 1.

Monthly Qtree Capacity Fact

Qtree capacity fact data for a single date point each month. Includes historical data. Contains the values from the Qtree capacity table, filtered to include only values associated with dates where the Is Month Representative equals 1.

Quarterly Qtree Capacity Fact

Qtree capacity fact data for a single date point each quarter. Includes historical data.

Q1 is Jan 1 - Mar 31

Q2 is Apr 1 - Jun 30

Q3 is Jul 1 - Sep 31

Q4 is Oct 1 - Dec 31

Storage and Storage Pool Capacity data model

The Storage and Storage Pool Capacity data model enables you to answer questions about storage capacity resource planning including storage and storage pools, and includes both physical and virtual storage pool data. This simple data model can help you answer questions related to capacity on the floor and capacity usage of storage pools by tier and data center over time.

If you are new to capacity reporting, you should start with this data model, because it is a simpler, targeted data model.

You can answer the following questions using the Storage and Storage Pool Capacity data model:

- What is the projected date for reaching the capacity threshold of 80% of my physical storage?
- What is the physical storage capacity on an array for a given tier?
- What is my storage capacity by manufacturer and family as well as by data center?
- What is the storage utilization trend on an array for all of the tiers?
- What are my top 10 storage systems with the highest utilization?
- What is the storage utilization trend the storage pools?
- How much capacity is already allocated?
- What capacity is available for allocation?

Creating a sample report with the Storage and Storage Pool Capacity data model

This report shows the total used and raw capacity per storage pool.

About this task

The instructions provide steps to create the following report:

Storage	Storage Pool	Raw Capacity (MB)	Capacity (MB)	Used Capacity (MB)
rtoeq1	bk:aggr_02	10,000.00	800.00	200.00
rtoeq2	bk07:aggr_03	200,000,000.00	100,000.00	2,000.00
rtpd05	br5:aggr_01	200,000.00	100,000.00	2,000.00
rtps d77	sr7:aggr_03	200,000,000.00	1,000,000.00	10,000.00
svr88	ttr3:aggr_04	200,000,040,000.00	1,000,100,000.00	520,000.00
Summary		200,400,250,000.00	1,001,300,800.00	534,200.00

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. Select **Packages > Storage and Storage Pool Capacity**.
3. From Query Studio, select **Physical Capacity**.
4. Select **Storage > Storage Pool** and insert it into the work area.
5. Select **Capacity (MB) > Capacity** and insert it into the work area.
6. Select **Capacity (MB) > Raw Capacity** and insert it into the work area.
7. Select **Capacity (MB) > Used Capacity** and insert it into the work area.
8. Select **Capacity (MB) > Used Raw Capacity** and insert it into the work area.
9. Optionally, group the report by Storage Pool by clicking the Storage Pool ID column and then the **Group** icon in the upper right.

Related tasks

[Accessing Query Studio](#) on page 45

Storage and Storage Pool Capacity data model submarts

The Storage and Storage Pool Capacity data model includes two submarts.

- All Capacity submart: Includes information about all storage pools, both physical and virtual. Use this to answer questions about storage and storage pools, including virtual storage pools. This submart excludes tier and data center to prevent double counting; however, this dimension itself does not prevent double counting.
- Physical Capacity submart: Includes information about only physical storage pools. Excludes virtual storage pools. Use this submart if you want to report on tiers and data centers. This submart itself does prevent double counting.


You can use only one submart per table or chart in a report. If you create a report that has multiple tables or charts, both submarts still cannot be used for a single table or chart in that report.

Projected capacity display in reports

You can show how much capacity will be used in a future by using the Projected Capacity data elements, which are included in both the Physical Capacity and All Capacity submart within the Storage and Storage Pool Capacity data model.

Projected capacity is an indication of how much capacity will be used at a future date, based on a linear regression algorithm that is applied to a set of recently observed capacity usage metrics.

The following example shows that if the capacity remains the same, you will run out of capacity within the next twelve months:



Future Capacity (dark green dots)

Future Capacity in the system assuming that there is no addition of storage systems. This corresponds to the capacity currently in the system.

Used Capacity (orange dots)

Used Capacity has been increasing.

Projected Used Capacity (light green dots)

According to the linear regression projection on the used capacity (orange dots), Projected Used Capacity will reach Projected Capacity as shown in the top right where the dark green and light green lines converge.

All Capacity submart details

You can use capacity data elements in the All Capacity submart in reports to show capacity utilization. The All Capacity submart includes information about all storage pools, both physical and virtual, but excludes tier and data center information.

Use this submart to answer these questions:

- What are the top 10 highly utilized storage pools in the last quarter?
- What is the storage pool capacity utilization over time?

Note: To report only on virtualized storage pools, you might want to filter using the Only Virtual Storage Pools filter.

Several items denoted with a hierarchical icon include hierarchies of data. If you place one of these items in your report, you will be able to drill down to its components.

Date data elements

Use the date data elements to write different reports for different time frames. This is useful, for example, when you want to view different time points in the past. The report using date data displays a line for each time that the ETL ran.

For details about dates, see the full description of date information in the cross-reference links.

Only Virtual Storage Pools filter

Use this filter to display only those storage pools that are virtual. When array virtualization is in use, the storage pool value reflects the storage pool from the front-end virtualizer.

Related concepts

[Date Dimension](#) on page 141

[Storage Dimension](#) on page 146

Physical Capacity submart details

You can use capacity data elements in the Physical Capacity submart in reports to show capacity utilization. This submart includes information about only physical storage pools and excludes virtual storage pools. Use this submart to report by tier and data center.

Use this submart to answer these questions:

- What is the raw capacity on the floor?
- What is the capacity utilization over time, across data centers, or across tiers?

Several items denoted with an icon are hierarchical. If you place one of these items in your report, you can drill down to its components.

Related concepts

[Date Dimension](#) on page 141

[Storage Dimension](#) on page 146

[Storage Pool Dimension](#) on page 147

[Tier Dimension](#) on page 151

Related references

[Data Center](#) on page 141

VM Capacity data model

The VM Capacity data model enables you to report your virtual machine environment and its capacity usage. This data model lets you report on changes in capacity usage over time for virtual machines and data stores. The data model also provides thin provisioning and virtual machine chargeback data.

The VM Capacity data model includes the following features:

- Provides access to VM storage capacity information recorded in the OnCommand Insight Data Warehouse.
- Provides access to information about individual VMs, such as data store usage, guest OS, and guest state.
- Enables the creation of detailed capacity accountability (chargeback) reports on the VM, application and business unit levels.
- Maintains historical data to show trending data (last week, last month, last quarter, and so on).

You can answer the following types of questions using the VM Capacity data model:

- How can I determine capacity chargeback based on capacity provisioned to VMs and data stores?
- What capacity is not used by VMs and which portion of unused is free, orphaned, or other?
- What do we need to purchase based on consumption trends?
- What are my storage efficiency savings achieved by using storage thin provisioning and deduplication technologies?

Creating a sample report with the VM Capacity data model

You can create a report that shows the VM OS, its power state, VM count, and provisioned capacity.

About this task

This process can be used for any of the tables described in the VM Capacity data model.

VM OS, Power State and Capacity

VM Count	Powered Off	Powered On	Suspended	Summary
UX	10	10	10	30
Windows		50		50
			0	0
Summary	10	60	10	80

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. From a blank screen in Query Studio showing the VM Capacity data model, select the **Advanced Data Mart**.
3. Select **VM Dimension > VM OS** and drag it into the work area.
4. Select **VM Dimension > VM Power State** and drag it to the right of the VM OS in the work area.
5. Select **VM Dimension > VM Count** and drag it to the right of the VM Power State in the work area.
6. To group the data by the power state, click on the VM Power State column in the work area and click the **Group** icon.
7. Optionally, click Power State column and click the **Pivot** icon to create a crosstab format for the report.

Related tasks

[Accessing Query Studio](#) on page 45

VM Capacity data model details

The VM Capacity data model contains a simple data mart and an advanced data mart.

- Simple data mart
- Advanced data mart

Capacities in the VM Capacity data model are taken from virtual disks (VMDKs). This means that the provisioned size of a VM using the VM Capacity data model is the size of its virtual disks. This is different from the provisioned capacity in the Virtual Machines view in the OnCommand Insight Client, which shows the provisioned size for the VM itself.

VM Capacity data

This data model enables you to report your virtual machine environment and its capacity usage. This data model lets you report on changes in capacity usage over time for virtual machines and data stores.

The fields and field descriptions are as follows:

Actual Capacity

The amount of data store capacity being used. This measurement ignores savings done in the backend storage using thin-provisioning and deduplication technologies.

See capacity concepts information in the cross references.

Application

Representative application.

Business Unit

The VM's business unit or the business unit for the application.

Capacity Type

See capacity concepts information in the cross references.

Data Store

Data store name. If capacity has been associated with a virtual machine, then that capacity has information about data stores to match. The virtual machine accesses the data store, which is defined on the ESX hosts, and the hosts access the capacity.

Host

Representative host accessing the capacity.

Host Data Center

The physical location of the host resource where the capacity resides.

Moid

Managed object identification

Provisioned Capacity

The amount of capacity that was allocated for a VM. This measurement includes VMDKs and other VM-related files.

See capacity concepts information in the cross references.

Service Level

Service level represents the agreement between a service consumer or application and the storage service provider, associated with an application. Service level differs from tier.

For example, you might have multi-tiered storage pools including 1 FC pool and one SSD pool. Your service level policy might require that Tier 1 should sustain 5,000 IOPS.

However, because this LUN uses only 100 IOPS, the service provider might decide to move this service to a lower tier that sustains only 500 IOPS. The service level is 5,000 or 500 IOPS, while the tier represents the quality of service, for example, gold or silver.

Service Level Cost

Cost / GB is the cost per gigabyte for a service level.

The service level cost can be the cost of providing the service or the actual price to service the consumer, while the tier cost is the average cost per GB for the specific tier.

Storage

Physical storage where the capacity resides.

Storage Access Type

Indication of the type of access to the storage, for example, SAN, NAS, or unknown.

Storage Data Center

Location of data center.

Tier

Level of service (for example, Gold, Silver, Bronze).

Tier Cost

Cost assigned to tier service per gigabyte.

Virtual Center IP

IP is an attribute of a capacity fact. It is the IP of the virtual center instance that reported this capacity. IP address of the Virtual Center for VM Capacity. Required because VM Capacity types may have the VM dimension empty for the following Capacity Types: data store Free, data store Overhead, and VM Volume Unused. See also VM Dimension in the Reporting Glossary.

VM

Name of the virtual machine.

Note: Capacities not associated with hosts, applications, tiers, or VMs use “N/A” to indicate null values.

VM IPs

IP of the VM.

VM OS

OS of the VM.

VM Power State

Indication of whether the VM power is on or off.

Related concepts

[VM Dimension](#) on page 152

VM Capacity advanced data mart

The full table data structure is provided for the VM Capacity data model in the Advanced Data Mart, which contains all current and historical data.

The Advanced data mart is organized into the following dimensions and attributes for capacity:

- Current VM Capacity Fact table: Includes only the current VM Capacity values. This table is used most often, for writing reports that reflect the current state of the storage environment (as of the last Data Warehouse build).
- Application Dimension: A representative application from the group of applications using the VM Capacity.
- Application Group Dimension: Defines the application group (set of applications) that is using the VM Capacity.
- Business Entity Hierarchy: Includes the business entity elements in hierarchical order (tenant, line of business, business unit, and project).
- Capacity Type is an attribute of a capacity fact. See [Capacity Types](#) on page 139.

- **Date dimension:** Includes dates and time periods from which you can select for the report. Each capacity fact is associated with the date when the Data Warehouse build ran that generated the fact.
- **Datastore dimension:** defined as NAS, SAN, or UNKNOWN when the data store is not connected a volume.
- **Host Dimension:** A representative host from the group of hosts using the VM Capacity.
- **Host Group Dimension:** Defines the host group (set of hosts) that is using the VM Capacity.
- **Internal Volume Dimension:** Defines the internal volumes on which the capacity resides.
- **Service Level Dimension:** Defines the storage service levels and costs.
- **Storage Dimension:** Defines the physical storage from which the VM Capacity is taken.
- **Tier Dimension:** Defines the tier based on the tier annotations in OnCommand Insight.
- **Virtual Storage Dimension:** Defines the front-end storage in configurations where capacity uses storage virtualization.
- **VM Dimension:** Defines the Virtual Machine using the capacity. The VM dimension table provides a current (live) snapshot of inventory. If capacity is associated with a VM, the table shows information about that VM.
- **VM Capacity Fact:** Includes the full set of VM Capacity data for every Data Warehouse build performed.
- **Time period fact tables:** Provides historical data in your reports.
 - Weekly VM Capacity Fact
 - Monthly VM Capacity Fact
 - Quarterly VM Capacity Fact

Related concepts

[Application Dimension](#) on page 135

[Application Group Dimension](#) on page 136

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Host Dimension](#) on page 144

[Host Group Dimension](#) on page 144

[Internal Volume Dimension](#) on page 144

[Service Level Dimension](#) on page 146

[Storage Dimension](#) on page 146

[Storage Pool Dimension](#) on page 147

[Storage Pool Dimension](#) on page 147

[Tier Dimension](#) on page 151

[Time period fact tables](#) on page 152

[VM Dimension](#) on page 152

Current VM Capacity Fact

The Current VM Capacity Fact table contains the values from the VM Capacity Fact table filtered to include only values associated with the Is Latest date (included in the Date Dimension table).

Use the VM Current Capacity Fact table to write reports that reflect only the latest data and not historical data. (To include historical data, use the VM Capacity Fact elements.)

VM Capacity Fact

This fact table includes all VM Capacity fact data and includes a date point for each ETL process. This table is useful for reporting on the data from a specific historical date, or providing more detailed trending reports.

You might use this fact table less frequently than other tables. For descriptions see Current VM Capacity Fact.

Related concepts

[Current VM Capacity Fact](#) on page 74

Volume Capacity data model

You can use this data model to analyze all aspects of volumes in your environment and look at them by vendor, model, tier, service level, and data center. You can view capacity related to orphan volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technology (iSCSI or FC) and compare virtual to non-virtual volumes for array virtualization issues.

You can answer several types of questions using the Volume Capacity data model:

- Which volumes have a utilization higher than a predefined threshold?
- What is the trend in my data center for orphan volume capacity?
- How much of my data center capacity is virtualized or thin provisioned?
- How much of my data center capacity must be reserved for replication?


Creating a sample report with the Volume Capacity data model

This example displays all orphaned volumes in a specific storage system.

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. From the blank screen in Query Studio showing the Volume Capacity data model, select **All Volumes**.
3. Select **Storage > Storage Name** and drag it into the work area.
4. Select **Volume > Volume Name** and drag it into the work area.
5. Select **Capacity (GB) > Orphaned Volume Capacity (GB)** and drag it into the work area.
6. Filter the report by orphaned volumes by selecting **Filters > Include Only Orphaned Volumes** and dragging it into the work area.
7. Optionally, filter the report by a specific storage system by right-clicking on the storage system column heading and selecting **Filter**.

Result

All Orphaned Volumes		
 Include Only Orphaned Volumes AND Storage Name: Sym-Perf, USP-2222111, USP-2222112		
Storage Name	Volume Name	Orphaned Volume Capacity (GB)
Sym-Perf	01:00	32.00
Sym-Perf	01:05	32.00
USP-2222112	00:0	16.00
USP-2222112	00:1	16.00
USP-2222112	00:2	16.00
USP-2222112	00:3	16.00
USP-2222112	00:4	16.00
USP-2222112	00:5	16.00
USP-2222112	00:6	16.00
USP-2222112	00:7	16.00
Summary		192.00

Related tasks

[Accessing Query Studio](#) on page 45

Volume Capacity data model details

This data model includes different submarts that enable you to include or exclude backend volumes.

This data model includes the following submarts:

- Host Visible Volumes: Excludes backend volumes.
- All Volumes: Includes backend volumes.

This data model includes the following data elements and dimensions:

- Capacity
- Data Center
- Host
- Service Level
- Storage
- Tier
- Additional Summaries

You can add rows or columns to your report that are summarized by the following data:

Volume count

The summary row displays the total number of volumes in the report.

Days since last access

The summary row displays the number of days since the resource was last accessed for capacity requests.

This data model includes the following filters:

- Include Only Accessed Volumes: Volumes that are accessed by hosts.
- Include Only Orphaned Volumes: Volumes that are masked but do not have active paths to them.

- Include Only Protection Volumes: Volumes that are used for replication.
- Include Only Unused Volumes: Volumes that are not accessed by hosts.
- Include Only Virtual Volumes: Volumes that are virtual.

Related concepts

[Host Dimension](#) on page 144

[Service Level Dimension](#) on page 146

[Storage Dimension](#) on page 146

[Tier Dimension](#) on page 151

Related references

[Data Center](#) on page 141

Chargeback data model

The Chargeback data model enables you to answer questions about used capacity and allocated capacity on storage resources (volumes, internal volumes, and qtrees). This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities and includes both current and historical data. Report data can be categorized by service level and storage tier.

Chargeback data is impacted by internal volumes:

- For storage without internal volumes, chargeback reports show chargeback by volumes.
- For storage with internal volumes:
 - If business entities are assigned to volumes, chargeback reports show chargeback by volumes.
 - If business entities are not assigned to volumes but assigned to qtrees, chargeback reports show chargeback by qtrees.
 - If business entities are not assigned to volumes and not assigned to qtrees, chargeback reports show the internal volume.
 - The decision whether to show chargeback by volume, qtree or internal volume is made per each internal volume, so it is possible for different internal volumes in the same storage pool to show chargeback at different level.

Capacity facts are purged after a default time interval. For details, see Data Warehouse processes.

Reports using the Chargeback data model might display different values than those reports using the Storage Capacity data model.

- For storage arrays that are not NetApp storage systems, the data from both data models is the same.
- For NetApp and Celerra storage systems, the Chargeback data model uses a single layer (of volumes, internal volumes, or qtrees) to base its charges, while the Storage Capacity data model uses multiple layers (of volumes and internal volumes) to base its charges.

You can answer the following types of questions using the Chargeback data model:

- What is the used capacity for qtrees versus the limits set per application or business entity (tenant, line of business, business unit, and project)?
- What are the trends of our used and free capacity so that we can do capacity planning?

- How much should we charge back each business unit for its used capacity (based on qtrees and visibility into internal values for NAS)?
- Which application, business entity, or host consumes the most capacity?

Creating a sample report with the Chargeback data model

This example creates a report that shows qtree information, its provisioned capacity, and tier.

About this task

Provisioned Capacity (TB)						
Business Unit	Application	Host	Storage	Protection Type	Resource Type	Tier
N/A	N/A	N/A	CX600-_34_APM00351029334	N/A	VOLUME	Bronze
N/A	N/A	N/A	Phoenix	N/A	INTERNAL_VOLUME	Bronze
Legal	Legal docs	ISL_Host	CX600-_34_APM00351029334	N/A	VOLUME	Bronze
R&D	N/A	N/A	Chicago	N/A	QTREE	Bronze
R&D	N/A	N/A	Dallas	N/A	QTREE	Bronze
R&D	N/A	N/A	Detroit	N/A	QTREE	Bronze
R&D	N/A	N/A	Jacksonville	N/A	QTREE	Bronze
R&D	N/A	N/A	Philadelphia	N/A	QTREE	Bronze
R&D	N/A	N/A	Phoenix	N/A	QTREE	Bronze
R&D	N/A	N/A	San Antonio	N/A	QTREE	Bronze
Stores	N/A	N/A	CX600-_27_APM00351029327	N/A	VOLUME	Bronze
Trading	OracleERP	AI_Host2	CX600-_27_APM00351029327	N/A	VOLUME	Bronze
Trading	OracleERP	AI_Host1	CX600-_27_APM00351029327	N/A	VOLUME	Bronze
Trading	OracleERP	AI_Host3	CX600-_27_APM00351029327	N/A	VOLUME	Bronze
Trading	OracleERP	AI_Host3	CX600-_27_APM00351029327	Volume Replica	VOLUME	Bronze
N/A	N/A	N/A	NTAP1	Internal Volume Replica	INTERNAL_VOLUME	Silver
N/A	N/A	N/A	NTAP2	Internal Volume Replica	INTERNAL_VOLUME	Silver
N/A	N/A	N/A	NetApp-17	N/A	INTERNAL_VOLUME	Silver
N/A	N/A	N/A	NetApp-18	N/A	INTERNAL_VOLUME	Silver
N/A	N/A	N/A	NetApp-19	N/A	INTERNAL_VOLUME	Silver

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. From a blank screen in Query Studio showing the Chargeback data model, select **Advanced Data mart**.
3. In the data model, expand **Storage Dimension** and select **Storage** and drag it in the work area.
4. Expand **Host Dimension** and select **Host** and drag it to the left of Storage in the work area.
5. Expand **Application Dimension** and select **Application** and drag it to the left of Host.
6. Expand **Business Entity Hierarchy** and select **Business Unit** and drag it to the left of Application.
7. Expand **Chargeback Fact** and select **Protection Type** and **Resource Type** and drag each to the right of Storage in the work area.
8. Expand **Tier Dimension** and select **Tier** and drag it to the right of Resource Type.
9. Optionally, group the report by Business Unit by clicking the **Business Unit** column and then the **Group** icon in the upper right.
10. Optionally, pivot the report by clicking on the **Tier** column and then the **Pivot** icon in the upper right.
11. Click on the **Title** and enter a report title.

Related tasks

[Accessing Query Studio](#) on page 45

Chargeback data model details

The Chargeback data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced: Provides access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[Chargeback data](#) on page 79

[Chargeback advanced data mart](#) on page 80

Chargeback data

This data model provides storage capacity chargeback and accountability information by hosts, application, and business entities and includes both current and historical data. Report data can be categorized by service level and storage tier.

The fields and field descriptions are as follows:

Application

Application annotation associated with this qtree.

Business Unit

Business unit annotation associated with this qtree.

Data Center

Location of data center, based on the switch's data center annotation in OnCommand Insight.

Host

Name of the host associated with the qtree.

Internal Volume

Name of the internal volume on which the qtree resides.

Mapped by VM

Indication whether this volume, qtree, or internal volume is used by a data store.

Protection Type

Type of protection used on the storage, for example, mirrored or striped.

Resource Name

Name of device or resource. In the Qtree data model, this is the qtree name. This can appear as "N/A" for volumes or show the qtree name.

Resource Type

Type of device or resource, for example, host or qtree. In the Qtree data model, this type shows as "QTREE."

Service Level

Service level represents the agreement between a service consumer or application and the storage service provider, associated with an application. Service level differs from tier.

For example, you might have multi-tiered storage pools including one FC pool and one SSD pool. Your service level policy might require that Tier 1 should sustain 5,000 IOPS.

However, because this LUN uses only 100 IOPS, the service provider might decide to move this service to a lower tier that sustains only 500 IOPS. The service level is 5,000 or 500 IOPS, while the tier represents the quality of service, for example, gold or silver.

Service Level Cost

Cost / GB is the cost per gigabyte for a service level.

The service level cost can be the cost of providing the service or the actual price to service the consumer, while the tier cost is the average cost per GB for the specific tier.

Qtree

Name of the qtree.

Storage

Physical storage where the capacity resides.

Storage Access Type

Indication of the type of access to the storage, for example, SAN, NAS, or unknown.

Storage Pool

Name of the storage pool on which the qtree resides.

Tier

Level of service (for example, gold, silver, or bronze).

Tier Cost

Unit of cost per gigabyte assigned to the tier level.

Provisioned Capacity

Capacity promised to storage initiators. Often used for chargeback. See "Capacity concepts" information.

Used Capacity

Capacity in the qtree including Snapshot Used Capacity. Used Capacity is no less than Consumed Capacity, depending on the storage efficiency technology applied to this qtree. See "Capacity concepts" information.

Chargeback advanced data mart

The Advanced data mart contains the full table data structure, which contains all current and historical associations from hosts to applications (for example, multiple applications for each host or hosts that are gone).

The Advanced data mart is organized into several fact and dimension tables.

Related concepts

[*Application Dimension*](#) on page 135

[*Application Group Dimension*](#) on page 136

[*Business Entity Hierarchy*](#) on page 137

[*Date Dimension*](#) on page 141

[*Host Dimension*](#) on page 144

[*Host Group Dimension*](#) on page 144

[*Internal Volume Dimension*](#) on page 144

[*Qtree Dimension*](#) on page 145

[*Service Level Dimension*](#) on page 146

[*Storage Dimension*](#) on page 146

[*Storage Pool Dimension*](#) on page 147

[*Tier Dimension*](#) on page 151

Time Chargeback Fact table

Use the elements in these tables to obtain historical data in your reports.

The information in these time period tables reflects a given time period only if the ETL runs in that period. For example, if OnCommand Insight stops running for the entire month of February, there won't be a date point for the February table.

The fact tables and descriptions are as follows:

Current Chargeback Fact

Chargeback fact data for only the last ETL run (the most recent date point).

Weekly Chargeback Fact

Chargeback fact data for a single date point each week. Includes historical data. Contains the values from the Chargeback table, filtered to include only values associated with dates where the Is Week Representative equals 1.

Monthly Chargeback Fact

Chargeback fact data for a single date point each month. Includes historical data. Contains the values from the Chargeback table, filtered to include only values associated with dates where the Is Month Representative equals 1.

Quarterly Chargeback Fact

Chargeback fact data for a single date point each quarter. Includes historical data.

Q1 is Jan 1 - Mar 31

Q2 is Apr 1 - Jun 30

Q3 is Jul 1 - Sep 31

Q4 is Oct 1 - Dec 31

Inventory data model

The Inventory data model lets you report on objects such as hosts, storage systems, switches, disks, tapes, qtrees, quotas, virtual machines, servers, and generic devices. It includes several submarts that enable you to see information about replications, FC paths, iSCSI paths, NFS paths, and violations. It does not include historical data.

The Inventory data model includes the following features:

- Provides easy access to most of the inventory information recorded in the OnCommand Insight Data Warehouse.
- Includes submarts that enable you to answer questions pertinent to an area of research.
- Reflects the state of inventory at the time of the last Data Warehouse build. The Inventory data model does not contain historical data.
- Provides access to annotations, including user-defined annotations, set on objects in the storage environment.

Using the Inventory data model, you can create reports that answer the following types of questions:

- What assets do I have and where are they?
- Who is using the assets?
- What types of devices do I have and what are components of those devices?
- How many hosts per OS do I have and how many ports exist on those hosts?
- What storage arrays per vendor exist in each data center?

- How many switches per vendor do I have in each data center?
- How many ports are not licensed?
- What vendor tapes are we using and how many ports exist on each tape?
- Are all the generic devices identified before we begin working on reports?
- What are the paths between hosts and storage volumes or tapes?
- What are the paths between generic devices and storage volumes or tapes?
- How many violations of each type do I have per data center?
- For each replicated volume, what are the source and target volumes?
- Do I have any firmware incompatibilities or port speed mismatches between Fibre Channel host HBAs and switches?

Creating a sample report with the Inventory data model

This example generates a report that answers the question: What types of violations occur per data center?

About this task

This report includes both iSCSI and FC violations.

Violations per data center

Violation Count	Boston	New York	Sydney	Tokyo	Zurich	Summary
Connection Count	24					24
Missing Security	16					16
Path Outage	8		3		46	57
Session Count	16					16
Unauthorized Sharing	8					8
Missing Redundancy		99	12	20	3	134
Missing Virtual Cluster Paths				1		1
Summary	72	99	15	21	49	256

The following example shows the process, and it can be used for any table in the Inventory data model.

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. In Query Studio, expand the **Inventory > Violation** submart.
3. Expand the **Violation > Simple** data mart.
4. Select the **Data Center** element and drag it to the work area.
5. Select the **Violation Type** element and drag it into the work area to the right of the Data Center column.
6. Select the **Violation Count** element and drag that into the work area to the right of Violation Type.

The result is a report listing all violations and their types and count per data center location.

7. Optionally, to filter out any missing values or “N/A” values, click on the **Data Center** column and click **Filter**. Exclude the "N/A" values.
8. To group the data, click on the **Data Center** column and click the **Group** icon.

Violations per data center		
Data Center	Violation Type	Violation Count
Boston	Connection Count	24
	Missing Security	16
	Path Outage	8
	Session Count	16
	Unauthorized Sharing	8
Boston		72
New York	Missing Redundancy	99
New York		99
Sydney	Missing Redundancy	12
	Path Outage	3
Sydney		15
Tokyo	Missing Redundancy	20
	Missing Virtual Cluster Paths	1
Tokyo		21
Zurich	Missing Redundancy	3
	Path Outage	46
Zurich		49
Summary		256

9. Optionally, click the **Pivot** icon to create a crosstab format for the report. The example in the introduction of this section shows the pivoted report.

Related tasks

[Accessing Query Studio](#) on page 45

Inventory data model details

In the Inventory data model, the data is presented in various submarts. You can drag and drop elements within one submart; you cannot combine elements from different submarts.

For example, the Inventory data model includes a Storage Asset submart and a Host submart. You can drag and drop elements from the Storage Asset submart or from the Host submart, but not from both submarts in the same report.

Each of the non-device specific submarts includes a Simple data mart.

Related concepts

[DR Replication submart](#) on page 84

[FC Path submart](#) on page 84

[Generic Device Assets submart](#) on page 84

[Host Assets submarts](#) on page 84

[Host Port Connectivity submart](#) on page 84

[iSCSI Path submart](#) on page 85

[NFS Path submart](#) on page 85

[Storage Assets submart](#) on page 85

[Switch Assets submart](#) on page 85

[Tape Assets submart](#) on page 85

[Violation submart](#) on page 86

Using submarts

Device submarts provide access to information about a particular type of device and the components within that device, for example, all the volumes, shares, and ports on a storage array.

For definitions of each inventory element, refer to the Inventory chapter in *OnCommand Insight Inventory User Guide*.

DR Replication submart

You can use the DR Replication submart to create reports on volume-to-volume replications. It includes information about hosts accessing the source volumes and about hosts accessing the target volumes. It also includes replicated count and volume capacity measurements.

Using this submart, you can answer the following questions:

- Which volumes are being replicated and on which hosts?
- For each replicated volume, what are the source and target volumes?

FC Path submart

You can use the FC Path submart to create reports on Fibre Channel paths between hosts and storage volumes or tapes and paths between generic devices and volumes or tapes.

Using this submart, you can answer the following questions:

- What is the path between hosts and volumes?
- Are there redundant paths for critical storage needs?
- What is the volume capacity related to this path?

Generic Device Assets submart

You can use the Generic Device Assets submart to create reports on generic devices and their ports.

Using this submart, you can answer the following questions:

- How many generic devices exist in the environment?
- Should we perform more identifications before we rely on incomplete report data?

Host Assets submarts

You can use the Host Assets submart to create reports on hosts and their adapters, ports, and annotations.

Using this submart, you can answer the following questions:

- How many hosts per OS do I have?
- How many ports exist on each host?

Host Port Connectivity submart

You can use the Host Port Connectivity submart to create reports on the connectivity between host and switch FC ports.

Using this submart, you can answer the following questions:

- Do I have firmware incompatibilities between Fibre Channel host HBAs and switches?
- Do I have port speed mismatches between Fibre Channel host HBAs and switches?

iSCSI Path submart

You can use the iSCSI Path submart to create reports on iSCSI paths between hosts and storage volumes.

Using this submart, you can answer the following questions:

- What is the SAN path for iSCSI?
- What are the session and connection counts on paths for critical storage needs?
- What is the volume capacity related to this path?

NFS Path submart

You can use the NFS Path submart to create reports on NFS paths to show services or end points of NAS paths to NFS Exports.

Using this submart, you can answer the following questions:

- What shares are exported from storage?
- Which initiator hosts are accessing those shares?

Storage Assets submart

You can use the Storage Assets submart to create reports on storage arrays and their volumes, storage pools, internal volumes, shares, disks, qtrees, quotas, controllers, and FC ports.

Using this submart, you can answer the following questions:

- How many storage arrays per OS do I have in each data center?
- What is the available and used capacity of the volumes, internal volumes, and storage pools?
- What storage arrays per vendor exist in each data center?
- What is the hard capacity quota set on qtrees per storage system?

Switch Assets submart

You can use the Switch Assets submart to create reports on switches and their fabrics, ports, and annotations.

Using this submart, you can answer the following questions:

- How many switches per vendor do I have?
- Are there any nonlicensed ports in each data center? If so, how many?

Tape Assets submart

Using the Tape Assets submart you can create reports on tapes and their controllers, ports and annotations.

Using this submart, you can answer the following questions:

- How many tapes per vendor do I have?
- How many ports are associated with each tape?

Violation submart

You can use the Violation submart to create reports on OnCommand Insight violations between hosts and storage volumes or tapes. Create reports on violations between generic devices and storage volumes or tapes.

Using this submart, you can answer the following questions:

- What are the types of violations of each type per data center?
- What is the volume capacity related to the violations per data center?

Performance data models

The Performance data models enable you to answer questions about performance for volumes, internal volumes, applications, switches, and hosts.

The Performance data models include information that helps you determine appropriateness of tiers, storage misconfigurations for applications, traffic data through the switches, and volume and internal volume last access times. This data model provides data such as response times, IOPs, throughput, number of writes pending, and accessed status.

Performance information is provided in the separate Performance data models.

Using the Performance data model, you can create reports that answer the following types of questions:

- What volumes or internal volumes have not been used or accessed during a specific period?
- Can we pinpoint any potential misconfiguration for storage for an application (unused)?
- What was the overall access behavior pattern for an application?
- Are tiered volumes assigned appropriately for a given application?
- Could we use less expensive storage for an application currently running without impact to application performance?
- What are the applications that are producing more accesses to currently configured storage?
- Is my host traffic through connected ports balanced?
- What are the most used switches based on port performance?
- What is the host trending throughput based on port performance?
- Which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch)?

Related concepts

[*Application Volume Hourly Performance data model*](#) on page 88

[*Host Volume Hourly Performance data model*](#) on page 89

[*Internal Volume Daily Performance and Internal Volume Hourly Performance data models*](#) on page 89

[*Switch Performance data model*](#) on page 90

[*Volume Daily Performance and Volume Hourly Performance data models*](#) on page 91

Creating a sample report with one of the Performance data models

You can access the data model and create a sample report. This example generates a report that answers the question: Are there any misalignments having a maximum IOPS of less than 100 or maximum throughput of less than 100?

About this task

You cannot drag elements from different data models into the same report.

Tier Underuse Misalignment Report			
Misalignments for tier "3.Silver", for having maximum IOPS less than 100 or maximum throughput less than 100			
Business Unit	Application	Max. Iops	Max. Throughput
AXA	AXA - Unknown	12	3
CSC	Corporate Shared Services	12	3
	EHUB	12	3
	OMEGA SRM	12	3
	SAP	12	3
	USD	12	3
DuPont	Capital Management	12	3
	CT Project	12	3
	LIMS	12	3
	NB_Media	12	3
	Print	12	3
ETS	Six Sigma	12	3
	Accession Number Generator	12	3
	AOL	12	3
	AP Studio Art	12	3
	APMISSPC	12	3
	APPLY	12	3
	Autosys	12	3
	Blackberry	12	3
	CASHEE	12	3

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. In OnCommand Insight Reporting, in the Public Folders, click the **Packages** folder link.
3. In the Packages folder, click the **Volume Daily Performance** link.
4. Expand the Data mart.
5. Expand the **Application Dimension**, select the **Application Name** element, and drag that into the work area.
6. Expand **Volume Daily Performance Fact**.
 - a. Select the **Max Total IOPS** element and insert it into the work area.
 - b. Select the **Max Total Throughput** element and insert it to the work area.

The result is a table listing all applications and their maximum total IOPS and total throughput. However, you want to limit your results to those with maximum IOPS of less than 100 or maximum throughput less than 100.
7. Right-click the **Max Total Throughput** column and select **Filter**. Enter the range of 0 to 100.
8. Right-click the **Max Total IOPS** column and select **Filter**. Enter the range of 0 to 100.
9. Expand the Business Entity Dimension, select the **Business Unit** element and drag it into the work area to the left of the Application Name.

10. Optionally, to filter out any missing values or “N/A” values, click on the **Application** column, right-click and select **Filter**. Select the applications you want to include.
11. To group the data, click on the **Business Unit** column and click the **Group** icon.

Related tasks

[Accessing Query Studio](#) on page 45

Performance data model details

Various models are included in the Performance data models.

All throughput measurements are in megabytes per second.

Related concepts

[Application Volume Hourly Performance data model](#) on page 88

[Host Volume Hourly Performance data model](#) on page 89

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Switch Performance data model](#) on page 90

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Combining related data models

You can drag and drop elements within one Performance data model; you cannot combine elements from different Performance data models.

For example, the Performance data model includes a Volumes Daily submart and an Internal Volumes Daily submart. You can drag and drop elements from the Volumes Daily submart or from the Internal Volumes Daily submart, but not from both submarts in the same report.

Application Volume Hourly Performance data model

This data model provides application performance data extracted from volume hourly data.

The data model contains the following fact and dimension tables:

- Application Volume Hourly Dimension
- Business Entity Hierarchy
- Date Dimension: Includes dates and time periods from which you can select for the report. Each fact is associated with the date when samples were taken into the system, not the time when Data Warehouse was processed.
- Service Level Dimension
- Tier Dimension
- Time Dimension
- Virtual Storage Dimension
- Time Facts: You can use Last Week, Last Month, and Last Quarter to filter different time periods.
- Application Volume Hourly Performance Fact: This table displays the response time, IOPS, throughput, and cache hit ratio information.

Related concepts

[Application Dimension](#) on page 136

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Service Level Dimension](#) on page 146

[Tier Dimension](#) on page 151

[Time Dimension](#) on page 152

Host Volume Hourly Performance data model

This data model provides host performance data, which was extracted from volume hourly data. There are Simple and Advanced data marts.

The data model contains the following fact and dimension tables.

- Business Entity Hierarchy
- Date Dimension: Includes dates and time periods from which you can select for the report. Each performance fact is associated with the date when the data was sampled.
- Host Dimension: A representative host from the group of hosts.
- Service Level Dimension
- Tier Dimension
- Time Dimension
- Host Performance Fact
- Time Facts

Related concepts

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Host Dimension](#) on page 144

[Service Level Dimension](#) on page 146

[Tier Dimension](#) on page 151

[Time Dimension](#) on page 152

Internal Volume Daily Performance and Internal Volume Hourly Performance data models

This data model provides internal volume performance data, which was extracted from either internal volume daily or hourly data.

The data model contains the following fact and dimension tables:

- Application Dimension
- Business Entity Hierarchy
- Date Dimension: Includes dates and time periods from which you can select for the report. Each performance fact is associated with the date when the data was sampled.
- Internal Volume Dimension: A representative internal volume from the group of internal volumes.
- Service Level Dimension
- Storage Dimension
- Storage Pool Dimension
- Tier Dimension

- Time facts: You can use Last Week, Last Month, and Last Quarter to filter different time periods.
- Internal Volume Daily/Hourly Performance Fact: From this, you can include IOPS, throughput, response time, and accessed information.

Related concepts

[Application Dimension](#) on page 135

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Service Level Dimension](#) on page 146

[Storage Dimension](#) on page 146

[Storage Pool Dimension](#) on page 147

[Tier Dimension](#) on page 151

Switch Performance data model

The Switch Performance data model enables you to answer questions about host traffic, under-utilized and over-utilized switches, and throughput.

You can answer the following types of questions using the Switch Performance data model:

- Is my host traffic through connected ports balanced?
- Show me switches/ports exhibiting a high number of errors.
- Show me most used switches based on port performance.
- Show me underutilized switches based on port performance.
- Show me host trending throughput based on port performance.
- Show me performance utilization for last X days for one specified host, storage system, tape, or switch.
- Show me which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch).
- Show me throughput for a specific business unit in our environment.

The data model includes the following fact and dimension tables:

- Application
- Business Entity Dimension
- Connected Device Dimension
- Data Center Dimension
- Fabric Dimension
- Host Dimension
- Storage Dimension
- Switch Port Dimension
- Tape Dimension
- Time Dimension
- Hourly Performance for Host

- Hourly Performance for Storage
- Hourly Performance for Switch Port
- Hourly Performance for Tape

Related concepts

[Application Dimension](#) on page 135
[Business Entity Hierarchy](#) on page 137
[Connected to Dimension](#) on page 141
[Fabric Dimension](#) on page 143
[Host Dimension](#) on page 144
[Storage Dimension](#) on page 146
[Switch Port Dimension](#) on page 149
[Switch Performance](#) on page 149
[Tape Dimension](#) on page 151
[Time Dimension](#) on page 152
[Connected Device dimension in data marts](#) on page 154
[Fabric Dimension](#) on page 143
[Port Dimension](#) on page 145
[Switch Dimension](#) on page 148
[Tape Dimension](#) on page 151

Related references

[Data Center](#) on page 141

Switch Performance fact tables

The Switch Performance data model includes switch performance fact tables for hosts, ports, storage systems, and tapes.

The following fact tables are included:

Hourly Performance for Host

Contains hourly rollups of port performance data aggregated per host. Each line represents one hour of port performance for ports connected to a specific host.

Hourly Performance for Switch Port

Contains hourly rollups of port performance per switch port. Each line represents one hour of port performance for one port within a switch.

Hourly Performance for Storage

Contains hourly rollups of port performance data aggregated per storage. Each line represents one hour of port performance for ports connected to a specific storage.

Hourly Performance for Tape

Contains hourly rollups of port performance data aggregated per tape. Each line represents one hour of port performance for ports connected to a specific tape.

Volume Daily Performance and Volume Hourly Performance data models

This data model provides daily and hourly aggregated data for volume performance.

The data model contains the following fact and dimension tables:

- Application Dimension

- Application Group Dimension
- Business Entity Hierarchy Dimension
- Date Dimension: Includes dates and time periods from which you can select for the report. Each performance fact is associated with the date when the data was sampled.
- Host Dimension
- Host Group Dimension
- Service Level Dimension
- Storage Dimension
- Storage Pool Dimension
- Tier Dimension
- Time Facts: You can use Last Week, Last Month, and Last Quarter to filter different time periods.
- Virtual Storage Dimension
- Volume Dimension
- Volume Daily or Hourly Performance Fact: Displays the response time, IOPS, throughput, and cache hit ratio information.

Related concepts

[Application Dimension](#) on page 135

[Application Group Dimension](#) on page 136

[Business Entity Hierarchy](#) on page 137

[Date Dimension](#) on page 141

[Host Dimension](#) on page 144

[Host Group Dimension](#) on page 144

[Service Level Dimension](#) on page 146

[Storage Dimension](#) on page 146

[Storage Pool Dimension](#) on page 147

[Tier Dimension](#) on page 151

[Volume Dimension](#) on page 153

Storage Efficiency data model

The Storage Efficiency data model enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount used or consumed (the physical measurement). For example, when thin provisioning is enabled, OnCommand Insight indicates how much capacity is actually taken from the device. Use this model also to determine efficiency when deduplication is enabled.

You can review storage by name, model, manufacturer, family, and data center. You can also review raw capacity, backend capacity, gain/loss, and potential gain/loss.

You can answer the following types of questions using the Storage Efficiency data model:

- What does the implementation of thin provisioning and deduplication technologies yield in storage efficiency savings?
- What are the storage savings across data centers?
- What is the historical trend of disk-based storage for the last three quarters?

- Based on actual utilization data, are we using the disk resources we have efficiently?
- Based on historical capacity trends, when do we need to purchase additional storage?
- What would be the capacity gain if we enabled technologies such as thin provisioning and deduplication?
- Am I at risk of running low on storage capacity now?

Storage Efficiency data model concepts

Different storage technologies can yield increases (gains) or decreases (losses) to capacity. This table describes some of these technologies and indicates whether the technology represents a gain or a loss to capacity.

Storage Technology	Gain or Loss	Notes
Disk failure protection	Loss	Represents the amount of capacity that is used to protect disk groups by RAID level. Example: RAID-4 where a single disk that is using every disk group for parity. Summing the size of the parity disk generates capacity loss.
Deduplication	Gain	Represents data compression or saved capacity in a compression ratio. Example: In a 2:1 compression ratio, for every 1 GB compressed capacity, there should be 2 GB of uncompressed capacity. This means that more space is needed to store the same data without deduplication.
Thin provisioning	Gain	Helps to eliminate over provisioning by consuming capacity only from the capacity that is used by the volume.
Clones	Gain	The clone shares capacity with its source volume and consumes only the capacity that is not mutual to itself and the source volume.

Storage efficiency calculation

Storage efficiency is calculated with a specific mathematical formula.

The efficiency of storage is calculated by taking the total storage raw capacity and adding the gainers, then subtracting the consumers. Finally, divide the total by the total storage raw capacity as follows:

$$\frac{\text{TotalArrayCapacity} + \text{GainedCapacity} - \text{LostCapacity}}{\text{TotalArrayCapacity}} \times 100$$

Creating a sample report with the Storage Efficiency data model

This example creates a sample report showing storage efficiency per data center and manufacturing vendor.

About this task

Storage Capacity Efficiency per Manufacturer and Data Center							
NOT dataCenter: NA							
dataCenter	manufacturer	model	name	technology	rawCapacityMB	potentialGainMB	potentialLossMB
Chicopee	HDS	HD59500V	HD59570V@10.12.104.200		846,544,798	0	0
		HD59570V	HD59570V@10.12.104.200				
	HDS				846,544,798	0	0
Chicopee					846,544,798	0	0
Data Center - Dev	HDS	R601	SBK_USP_VM@172.26.64.46				
		R601	USP_VM@172.26.64.46		373,475,830	0	0
	HDS				373,475,830	0	0
Data Center - Dev					373,475,830	0	0
Engineering Lab	NetAPP	FAS3170	engnetapp01,engnetapp02		677,763,072	0	0
	NetAPP				677,763,072	0	0
Engineering Lab					677,763,072	0	0
Harborside	HP	HSV110	eva_bcp		618,545,391	0	0
	HP				618,545,391	0	0
Harborside					618,545,391	0	0
London	HDS	HD59500V	HD59580V@10.16.17.50		382,476,900	0	0
		HD59580V	HD59580V@10.16.17.50				
		R601	USP_VM@10.16.17.185		285,016,626	0	0
	HDS				667,493,526	0	0

Steps

1. Access Query Studio.
For detailed steps, see the related links about accessing Query Studio.
2. From a blank screen in Query Studio showing the Storage Efficiency data model, select **Advanced Data Mart**.
3. Expand **Storage Dimension**.
4. Select and drag **Data Center**, **Manufacturer**, **Model**, and **Storage Name** into the work area.
5. Expand **Efficiency Fact**.
6. Select and drag **Storage Technology**, **Raw Capacity**, **Potential Gain**, and **Potential Loss** into the work area.
7. Optionally, group the report by data center by clicking the **Data Center** column and then the **Group** icon in the upper right.
8. Optionally, group the report next by Manufacturer.

Related tasks

[Accessing Query Studio](#) on page 45

Storage Efficiency data model details

The Storage Efficiency data model contains a simple data mart and an advanced data mart.

- Simple: Provides quick access to the most commonly used data elements. Includes only the last snapshot of Data Warehouse data; it does not include historical data.
- Advanced: Provides access to historical data values. All data elements in the Simple data mart are also included in the Advanced data mart.

Related concepts

[Storage Efficiency data](#) on page 95

[Storage Efficiency advanced data mart](#) on page 96

Storage Efficiency data

This data model enables you to track the storage efficiency score and potential over time. This data model stores measurements of not only the provisioned capacity, but also the amount used or consumed (the physical measurement).

The fields and field descriptions are as follows:

Backend Capacity

Total capacity of backend LUNs. Relevant for virtual arrays only.

Data Center

The physical location of the storage device where the capacity resides.

Gain/Loss in MB

Capacity gained or lost by using the technology. “0” appears when the technology is not specified.

See Storage Efficiency concepts in the cross references.

IP

IP address assigned to the storage array.

Manufacturer

Name of the manufacturer of the storage array.

Microcode Version

Version of the firmware running on the storage array.

Model

Manufacturer’s model number.

Potential Gain/Loss in MB

Potential capacity to gain or lose by using the technology to the full extent. “0” appears the when technology is not specified.

Raw Capacity

See capacity concepts in cross references.

Serial Number

Serial number for the storage array.

Storage Family

Product classification of the storage array (for example, CLARiiON or Symmetrix).

Storage Name

Name of the storage where the capacity resides.

Technology

Storage efficiency and protection technology such as deduplication, thin provisioning, or type of protection used on the storage, for example, mirrored or striped. Blank when representing the storage's raw and backend capacity (one row storage).

Related concepts

[Storage Efficiency data model concepts](#) on page 93

Storage Efficiency advanced data mart

The Advanced data mart contains the full table data structure, which contains all current and historical associations from hosts to applications (for example, multiple applications for each host, hosts that are gone).

The Advanced data mart is organized into the following fact and dimension tables:

- **Date Dimension:** Includes dates and time periods from which you can select for the report. Each capacity fact is associated with the date when the Data Warehouse build ran that generated the fact.
- **Efficiency fact:** Includes the capacity data.
- **Storage Dimension:** Provides attributes (such as storage name, IP, model, manufacturer, family, and type). Also includes data center.

Related concepts

[Date Dimension](#) on page 141

[Storage Dimension](#) on page 146

Efficiency Fact

This fact table includes all capacity efficiency fact data and includes a date point for each ETL process. You might use this fact table less frequently than other tables. This table is useful for reporting on the data from a specific historical date, or providing more detailed trending reports.

Time Efficiency Fact table

Use the elements in these tables to obtain historical data in your reports.

The information in these time period tables reflects a given time period only if the ETL runs in that period. For example, if OnCommand Insight stops running for the entire month of February, there won't be a date point for February.

The fact tables and descriptions are as follows:

Current Efficiency Fact

Efficiency fact data for only the last ETL run (the most recent date point).

Weekly Efficiency Fact

Efficiency fact data for a single date point each week. Includes historical data. Contains the values from the Efficiency table, filtered to include only values associated with dates where the Is Week Representative equals 1.

Monthly Efficiency Fact

Efficiency fact data for a single date point each month. Includes historical data. Contains the values from the Efficiency table, filtered to include only values associated with dates where the Is Month Representative equals 1.

Quarterly Efficiency Fact

Efficiency fact data for a single date point each quarter. Includes historical data.

Q1 is Jan 1 - Mar 31

Q2 is Apr 1 - Jun 30

Q3 is Jul 1 - Sep 31

Q4 is Oct 1 - Dec 31

Data marts for SQL queries

Data Warehouse provides data in its database schema that you can use in your custom SQL queries. The database includes individual data marts.

Related concepts

Inventory data mart for applications in OnCommand Insight (dwh_inventory) on page 97

Capacity data mart (dwh_capacity) on page 98

Capacity Efficiency data mart (dwh_capacity_efficiency) on page 98

File System Utilization data mart (dwh_fs_util) on page 98

Performance data mart (dwh_performance) on page 98

Port Capacity data mart for switches (dwh_ports) on page 98

Storage and Storage Pool Capacity data mart (storage_and_storage_pool_capacity_fact) on page 98

Volume Capacity data mart for SQL queries on page 133

Data marts in Data Warehouse

Data appears in various data marts in Data Warehouse. Some of these data marts require additional licensing.

- Inventory data mart
- Capacity data mart
- Capacity Efficiency data mart
- File System Utilization data mart
- Performance data mart
- Port Capacity data mart
- Volume Capacity data mart

Inventory data mart for applications in OnCommand Insight (dwh_inventory)

The Inventory data mart contains physical and logical data on the environment being monitored by OnCommand Insight. The Inventory data mart publishes the device configurations and relationships between devices.

The Inventory data mart is generated from OnCommand Insight Assure data (switch and storage data) that is acquired from the data sources.

The Inventory data mart consolidates data from multiple OnCommand Insight databases. The data mart contains current data only.

Use this data mart to do the following:

- Generate inventory reports
- Work with an API for retrieving details on the devices that are monitored by multiple OnCommand Insight servers
- Interface with CMDB systems

Capacity data mart (dwh_capacity)

The Capacity data mart is based on the Inventory data mart, but displays changes in capacity usage over time. The Capacity data mart includes fact and dimensional tables. Unlike the Inventory data mart, historical data is contained in the Capacity data mart.

Use this data mart to do the following:

- Generate chargeback reports by finding the amount of capacity being used by a business unit.
- Identify capacity planning trends.
- Monitor historical capacity usage trends related to tiers, data centers, applications, and business units.

Then, project anticipated use in forecast reports.

- Generate VM capacity chargeback and planning (included in the VM Capacity data mart, which is included in the Capacity data mart).

The data mart includes two fact tables: Capacity fact, which contains storage capacity information, and the VM Capacity fact table, which provides capacity insight into resources such as the virtual center, ESX hosts, data stores, and VMs.

Capacity Efficiency data mart (dwh_capacity_efficiency)

The goal of this data mart is to estimate the gain and loss (or the efficiency) of different technologies used in the storage.

For example, you can determine how capacity is saved by using storage efficiency technologies, such as thin provisioning, deduplication and RAID.

File System Utilization data mart (dwh_fs_util)

This data mart includes file system utilization data. This enables visibility into capacity utilization by hosts on the file system level.

Administrators can determine allocated and used capacity per file system, determine the type of file system, and identify trending statistics by file system type.

Performance data mart (dwh_performance)

The Performance data mart collects performance information for volumes, internal volumes, applications and hosts.

Volume and internal volume performance data is aggregated hourly and daily, while application and host performance data is aggregated hourly.

Port Capacity data mart for switches (dwh_ports)

This data mart provides historical data on data such as switch ports connectivity, status and speed.

This data can be used when planning for purchasing new SAN switches.

Storage and Storage Pool Capacity data mart (storage_and_storage_pool_capacity_fact)

You can create reports about storage capacity resource planning, including physical and virtual storage pool data for storage and storage pools.

Volume Capacity data mart (dwh_capacity)

Use this data mart (dwh_capacity. volume_capacity_fact table) to report on all aspects of volumes in your environment and look at them by vendor, model, tier, service level, and data center.

You can view capacity related to orphan volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technology (iSCSI or FC) and compare virtual to non-virtual volumes for array virtualization issues.

Accessing MySQL using the command-line interface

In addition to accessing Data Warehouse data elements through the report authoring tools (such as Query Studio and Report Studio), you can obtain access to Data Warehouse data elements directly by connecting as a MySQL user. You might want to connect as a MySQL user to use the data elements in your own applications.

About this task

There are many ways to connect. The following steps show one way.

When accessing MySQL, connect to the MySQL database on the machine where Data Warehouse is installed. The MySQL port is 3306 by default; however, you can change it during installation. The user name and password is dwhuser/netapp123.

Steps

1. On the machine where Data Warehouse is installed, open a command-line window.
2. Access the MySQL directory in the OnCommand Insight directory.
3. Type the following user name and password:

```
mysql -udwhuser -pnetapp123
```

The following is displayed, depending on where Data Warehouse is installed:

```
c:\Program Files\SANscreen\mysql\bin> mysql -udwhuser -pnetapp123
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.
Your MySQL connection id is 882
Server version: 5.1.28-rc-community MySQL Community Server (GPL)
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the buffer.
```

4. Show the Data Warehouse databases:

```
show databases;
```

The following is displayed:

```
mysql> show databases;
+-----+
| Database |
+-----+
| information_schema |
| dwh_capacity |
| dwh_capacity_efficiency |
| dwh_fs_util |
| dwh_inventory |
| dwh_performance |
| dwh_ports |
+-----+
```

Accessing the database schema

The entire Data Warehouse database schema is shown in OnCommand Insight Data Warehouse.

About this task

For each table, the DB Schema documentation describes all the columns. References between tables are documented in the “source” side of the relationship. For example, column `hostId` in `host_adapter` refers to `host` (using `id`).

Steps

1. From the Data Warehouse Main Menu, select **Documentation > Database Schema**.
2. In the Database pane, click one of the databases, for example, **dwh_inventory**.
3. For example, in the Tables pane, click the **annotation_value** table.

The `dwh_inventory.annotation_value` table appears.

Database Schema Databases dwh_inventory acq_acquisition_ur acq_data_source annotation_value application backend_lun backend_lun_to_vr backend_path business_entity cm_lun_requests cm_lun_requireme cm_port_requests cm_port_requireme	dwh_inventory.annotation_value			
	Column	Type	Nullable	Description
	<code>id</code>	<code>int(11)</code>	false	GUID for the annotation.
	<code>annotationType</code>	<code>varchar(255)</code>	false	System or user defined type such as Tier, Data center, etc.
	<code>valueIdentifier</code>	<code>varchar(255)</code>	false	Value of the annotation.
	<code>enumSequence</code>	<code>int(11)</code>	true	Sequence number determining the order of enumeration values. This is used primarily for display purposes.
	<code>enumCost</code>	<code>double</code>	true	Optional cost associated with the annotation. Applicable for Tier annotation.
	dwh_inventory.application			
	Column	Type	Nullable	Description
	<code>id</code>	<code>int(11)</code>	false	GUID for the application.
	<code>businessEntityId</code>	<code>int(11)</code>	true	GUID of the business entity References: <ul style="list-style-type: none"> • <code>id</code> in dwh_inventory.business_entity

Inventory data mart for SQL queries

The Inventory data mart (`dwh_inventory`) provides a single database schema that consolidates inventory data.

The data in the Inventory data mart is used for generating inventory reports and for building the capacity, performance, port capacity, and other data marts. The Inventory data mart contains physical and logical data on the environment being monitored by OnCommand Insight and includes device configurations and relationships between devices. The Inventory data mart consolidates data from multiple OnCommand Insight databases. The data mart contains current data only.

Use this data mart to do the following:

- To generate inventory reports
- As an API for retrieving details on the devices that are monitored by multiple OnCommand Insight servers
- As an interface for CMDB systems

Questions you can answer with the Inventory data mart

The Inventory data mart can be used to answer various questions.

- What assets do I have and where are they?
- Who is using the assets?
- What types of devices do I have and what are components of those devices?
- How many hosts per OS do I have and how many ports exist on those hosts?
- What storage arrays per vendor exist in each data center?
- How many switches per vendor do I have in each data center? How many ports are not licensed?
- What vendor tapes are we using and how many ports exist on each tape?
- Are all the generic devices identified before we begin working on reports?
- What are the paths between hosts and storage volumes or tapes? What are the paths between generic devices and storage volumes or tapes?
- How many violations of each type do I have per data center?
- For each replicated volume, what are the source and target volumes?
- Do I have any firmware incompatibilities or port speed mismatches between Fibre Channel host HBAs and switches?

Inventory data mart considerations

For each device (for example, host, storage, or switch) supported by DWH, FC ports are presented in device-specific tables providing strong references between ports, nodes and devices. Connectivity between FC ports is displayed through the port_connectivity table.

The Inventory data mart includes access to the following data:

- Applications and annotations
- Business entity information
- Capacity information
- Disk information
- File systems
- NAS
- Paths and violations
- Qtree information
- Quota information
- Storage devices
- SAN fabrics and zoning
- SAN connectivity and active ports
- Virtualization and replication
- VM

Applications and annotations

Annotations can be accessed using any one of several methods.

- **Annotation tables:** The `annotation_value` table contains all the annotations extracted from OnCommand Insight regardless of their type as defined in OnCommand Insight. The `object_to_annotation` table links annotation values to entities in the `dwh_inventory` database. Use the `objectType` column to determine to which table in `dwh_inventory` an annotation refers to. If the `objectType` is set to `Port`, the annotation might be referring to any of the `host_port`, `switch_port` or `storage_port` inventory tables.
- **Dynamic annotations:** Using the Data Warehouse Portal, you can choose to add annotations permanently as columns to inventory tables. After this is done, the annotation value is visible when selecting from the entity's inventory table. See information about importing annotations in the *OnCommand Insight Data Warehouse User Guide*.

Application data that is extracted from the following OnCommand Insight sources are merged into the Data Warehouse application table:

- `host_to_application` and `vm_to_application` tables
- Application and Business Entity tables and the annotations of virtual machines

Applications are linked to hosts and virtual machines using the `host_to_application` and `vm_to_application` relationship tables.

Because the Application and Business Unit annotations of virtual machines are merged into a single "application and business unit" entity, there might be some differences between OnCommand Insight and Data Warehouse. For example, if a VM is annotated with a single application and multiple business units, only one business unit is selected.

Business entity information

The hierarchical business entity includes the following components: tenant, line of business, business unit, and project. You do not have to specify a value for each component. Any value can be left out if not used.

For example, if a site does not have a line of business and projects, but wants to use tenant and business unit, you can leave the line of business and project fields blank when setting tenant and business unit.

Business entities can be associated with one or more applications. If a business entity is associated with the same object (Host for example) on two different servers, but contains two different values, Data Warehouse chooses one and logs the event.

Full Name

Complete business entity name.

Tenant

Primarily used by service providers to associate resources within a business entity to a customer.

LOB

A line of business within a company, for example "Storage Management" or "Database Replication."

Business Unit

Business unit for the application. If there are multiple business units associated with a host, then the business unit shown in the table is a representative. Name of the Business Unit that owns the application. A specific application has the same Business Unit on all hosts.

Note: The value will be N/A if the application is not associated with a Business Unit.

Project

A project to which you might want to further break down capacity chargeback.

Capacity information

The Inventory data mart includes physical and provisioned capacity along with usable and raw capacity. These capacity data elements can be applied to resources such as arrays, hosts, business units, applications, tiers, and data center.

- Physical capacity: Capacity actually present on the storage device.
- Provisioned capacity: Capacity promised to storage initiators. Often used for chargeback.
- Usable capacity: Capacity that is presented after RAID.
- Raw capacity: Capacity that is presented before RAID.

For details about capacity concepts, see the *OnCommand Insight Planning User Guide*.

Disk information

Use the disk table to display data such as speed, disk capacity, disk location, disk role, vendor, model, status, and serial number.

File systems

The file_system table contains hosts file systems and their utilization.

NAS

There are various tables related to supporting NAS information.

- volume
- storage_pool
- internal_volume
- storage
- quota
- nas_file_share: Includes internal volume ID, path, qtree, and FS status.
- nas_logical: Includes host, storage, and share IDs
- nas_share: Includes share ID, share name, and protocol used.
- nas_share_initiator: Includes share initiator and assigned permission.

iSCSI information

The Inventory data mart displays iSCSI logical path information.

Use the following tables to help display iSCSI information.

- iscsi_logical: Shows the iSCSI logical paths
- violation: Uses the technologyType column, which displays values such as FC, iSCSI, and HV. This helps identify violations by the type of technology. Also, this table stores iSCSI-specific violations, such as the iSCSI security violation, number of sessions violation, and number of connections violation.

Paths and violations

An access path represents a relationship between a particular application on a given server and its data on a storage device.

An access path is a fundamental notion in OnCommand Insight. It consists of a sequence of physically connected components, each of which is configured to allow data to flow along that path (assuming proper zoning and masking). Each path has various properties. For each path, OnCommand Insight ensures that these properties conform to a policy associated with the path, where that policy establishes the minimal requirements for the path's availability, performance, and security.

For details about access paths in the Client see the *OnCommand Insight Inventory User Guide*.

The following tables contain data related to the access paths in the storage environment:

- Logical path ("logical" table): Lists all the access paths in the storage network. If a host can access a certain volume in some storage, there is a logical path record with the IDs of these three. Additional fields include properties of the logical path are used to ensure that the path complies with the desired service level.
These properties include: minimum number of hops between the host and the storage and number of host and storage ports.
- Physical path (physical_path table)
- Violation (violation table): Lists all paths that are missing or do not comply to the policy. The triplet of host, storage and volume IDs that represent the access path are included. Also the violation type (for example, "Missing Path" or "Missing Redundancy") and the time the violation appeared are included, too. The violation table uses the technologyType column, which displays values such as FC, iSCSI, and HV. This helps identify violations by the type of technology. Also, this table stores iSCSI-specific violations, such as the iSCSI security violation, number of sessions violation, and number of connections violation.

You might find it helpful to use these tables also when working on identification and violation issues:

- generic_device
- storage
- host
- tape

Qtree information

The qtree table includes information about the qtree subdivision of OnCommand Insight internal volumes. You might want to look at the qtree details to view the security state, qtree status, or qtree type.

Quota information

The quota table provides data about space quotas that indicates the limits on the amount of usable space or the number of files that can be created on a NAS share or qtree.

Quotas are typically set up by storage administrators on the shares or qtrees. Multiple quotas can exist on a tree.

SAN connectivity and active ports

The SAN connectivity can be identified using multiple tables.

- Host (host table)

- HBA (host_adapter table)
- Switch (switch table)
- Storage (storage, storage_controller tables)
- Virtual switch (virtual_switch table)
- Tape (tape, tape_controller tables)
- Port connectivity (port_connectivity table)

The port tables for storage, switches, tapes, hosts, and generic devices in the Inventory database include an “active” column. This column indicates whether the port is active. The column appears in the following tables:

- switch_port
- storage_port
- tape_port
- host_port
- generic_device_port

SAN fabrics and zoning

Zoning information is included multiple tables.

- zone table: Lists all the zones of active zone sets (zone_capability in OnCommand Insight) and includes fields from the zone capabilities, such as the WWN of the fabric and the configuration name.
- zone_member table: Lists all the members of zones in the active zone set. For ease of usage, zoneName and fabricId fields are included, too.

Non-active zones and zone members are filtered out. Also, zone members are not merged. When two connectors are reporting on the same zone, the first reporting determines the zone_members.

Storage devices

Storage data is available in the Inventory data mart.

- Volume access: The volume_map and volume_mask tables present the mapping and masking configuration of each FC volume. Volume mapping in OnCommand Insight defines the LUN presented to FC initiators through a selected storage port. Volume masking limits access to the volume for selected initiators identifiable by HBA/port WWN. The protocolController column is used for matching mapping and masking settings of a volume.
- Virtual volume: A volume is flagged as “virtual” (virtual=1) if its capacity is mapped to volumes on other storage arrays. See virtual_to_backend_volume and see Virtualization and replication on page 175. The “actualCapacity” of virtual, front-end volumes should be set to 0.
- Volumes involved in replications (local or remote) are flagged as “source” and “target” to indicate the direction of the replication.

Virtualization and replication

Both SAN array virtualization and replications are modeled as a relationship between SAN volumes.

- Array Virtualization (virtual_to_backend_volume table): A many-to-many relationship mapping multiple front end (virtual) volumes to a pool of back end volumes from which the virtual-volumes are carved.

- Replication (dr_replica table): A one-to-many relationship between “replica source” and multiple “replica targets.” Replicas could be chained to form a tree-like structure in which some “replica target” volumes act as a “replica source” for other replications.

VM

Various tables are related to VM data, and map VMware virtual disks and data stores to SAN volumes or NAS shares.

These tables are as follows:

- Application (application table)
- Host (hv_server, host tables): Contains information about an ESX server is hosting the VMs. The host table contains all the hosts known to OnCommand Insight including non-ESX hosts. The hv_server table adds information on ESX hosts such as datacenter name.
- Host to NAS share (hv_host_to_nas_share table)
- NAS/SAN data store (hv_data_store, hv_virtual_disk tables): Supports both SAN and NAS data stores. SAN data stores are referred to as LUNs (hv_lun), while NAS data stores are referred to as NAS-Shares (hv_nas_share).
- LUN and path (hv_lun table): Represents the host’s view of SCSI LUNs. Paths from the LUNs to the storage are defined in the hv_path table. When multiple paths are available for a LUN, the path that is used is referred to as the active path. OnCommand Insight uses those tables to map LUNs to SAN volumes. This mapping is available through hv_data_store_to_volume and hv_virtual_disk_to_volume tables.
- Server (hv_server table)
- Paths (hv_path table)
- Internal volumes (internal_volume table)
- Volumes (volume table)
- Virtual machine to application host (host table)
- Virtual machine to disk (hv_virtual_machine_to_disk table)
- Virtual disk (hv_virtual_disk table): Contains data about virtual disks, which are stored inside data stores, represented by data storeId in hv_virtual_disk. The hv_virtual_machine_to_disk table maps virtual disks to the VMs that are using them.
- Virtual disk to volume (hv_virtual_disk_to_volume table)
- Data store and data store to volume (hv_data_store and hv_data_store_to_volume tables). Data store to internal volume (hv_data_store_to_internal_volume table)

The act of mapping a virtual disk to a SAN volume (so VM SCSI commands are delegated to the volume) is called “raw device mapping.” In the Inventory datamart, the lunId (LUN ID) of such virtual disks points to the mapped LUN (the lunId is null for non-RDM virtual disks).

The OnCommand Insight Inventory data mart includes the following information:

- Map of data stores and virtual disks (RDM only) to SAN volumes (hv_data_store_to_volume and hv_virtual_disk_to_volume tables).
- Definitions of a policy and reporting on policy violations (violation table) such as the following:
 - HV_ACTIVE_PATH_CONFLICT_VIOLATION: When ESX hosts use active/passive SAN arrays, each volume must be accessed using the same "storage processor" (SP). If this policy

is enabled and OnCommand Insight identifies the SAN array as active/passive, then a violation is issued for each volume accessed through different SPs simultaneously.

- **HV_MISSING_VIRTUAL_CLUSTER_PATH_VIOLATION:** All the hosts in a cluster (identifiable by VMware-HA or manually by the user) are required to access the same set of SAN volumes. When this policy is enabled, a violation is issued if a host in a cluster does not access a volume that is being accessed by other hosts in the cluster.
- **HV_INCONSISTENT_LUNS_VIOLATION:** If this policy is enabled, a violation is issued when multiple hosts access the same SAN volume using different LUN numbers.
- **HV_MISSING_VIRTUAL_CLUSTER_NAS_PATH_VIOLATION:** When this policy is enabled, a violation is issued if a host in a cluster does not access a NAS share that is being accessed by other hosts in the cluster.
- Association of VMs with applications and business units (vm_to_application and application tables) that could be used for capacity chargeback and planning.

Inventory data mart sample SQL queries

Various SQL queries can be executed in the Inventory data mart.

Inventory data mart SQL Query Example 1

The following query demonstrates selecting all the host ports and the connected switch ports. Notice that host ports connected directly to storage ports are listed having the switch name column empty.

```
SELECT
    h.name AS 'Host',
    hp.wwn AS 'Host port',
    pc.connectedWwn AS 'Connected port',
    s.Name AS 'Switch name'

FROM
    dwh_inventory.host h

JOIN
    dwh_inventory.host_port hp ON hp.hostId = h.id

LEFT JOIN
    dwh_inventory.port_connectivity pc ON pc.portId = hp.id

LEFT JOIN
    dwh_inventory.switch_port sp ON sp.id = pc.connectedId

LEFT JOIN
    dwh_inventory.switch s ON s.id = sp.switchId
```

The following shows the results of this query.

Host	Host port	Connected port	Switch name
Clearcase1	10:B0:00:00:00:00:36	10:B0:00:00:00:00:53	hbrcd1
Clearcase1	10:B0:00:00:00:00:37	10:B0:00:00:00:00:69	stbrcd1
Clearcase2	10:B0:00:00:00:00:38	10:B0:00:00:00:00:54	hbrcd1

Host	Host port	Connected port	Switch name
Clearcase2	10:B0:00:00:00:00:39	10:B0:00:00:00:00:00:70	stbrcd1
Filer1	10:B0:00:00:00:00:00:40	10:B0:00:00:00:00:00:55	hbrcd1
Filer1	10:B0:00:00:00:00:00:41	10:B0:00:00:00:00:00:71	stbrcd1
Filer2	10:B0:00:00:00:00:00:42	10:B0:00:00:00:00:00:56	hbrcd1
Filer2	10:B0:00:00:00:00:00:43	10:B0:00:00:00:00:00:72	stbrcd1
csNAS1	10:B0:00:00:00:00:00:44	10:B0:00:00:00:00:00:73	stbrcd1
csNAS1	10:B0:00:00:00:00:00:345	10:B0:00:00:00:00:00:85	hbrcd2

Inventory data mart SQL Query Example 2

The following query demonstrates using annotations and querying for all the hosts and their data center.

```
SELECT
    h.name 'Host',
    a.valueIdentifier 'Data center'
```

```
FROM
    dwh_inventory.host h
```

```
LEFT JOIN
    dwh_inventory.object_to_annotation o2a
    ON o2a.objectId = h.id
```

```
LEFT JOIN
    dwh_inventory.annotation_value a
    ON a.id = o2a.annotationValueId AND a.annotationType =
    'Data_Center'
```

The following shows the results of this query.

Host	Data center
Clearcase1	Tokyo
Clearcase2	Tokyo
Filer1	Tokyo
Filer2	Tokyo
csNAS1	Tokyo
csNAS2	Tokyo
csNAS3	Tokyo
lotus	Tokyo
Nas1Backup	Tokyo
Nas2Backup	Tokyo

Capacity data mart for SQL queries

The Capacity data mart lets you generate chargeback data, identify capacity planning trends, analyze consumption, and identify capacity used and not used. The Capacity data mart (`dwh_capacity`) is based on the inventory data mart, but displays changes in capacity usage over time.

The Capacity data mart includes information for hosts, volumes, internal volumes, qtrees, virtual machines, storage, and storage pools. It also includes fact and dimensional tables. Unlike the Inventory data mart, historical data is contained in the Capacity data mart.

Use this data mart to do the following:

- Generate chargeback data based on capacity provisioned to VMs and data stores. Obtain more granular capacity usage by using the host, application, storage, application and tier dimensions.
- Identify capacity planning trends.
- Generate VM capacity chargeback and planning reports (included in the VM Capacity table, which is included in the Capacity data mart).
- Identify capacity not used by VMs and classify the unused capacity as free, other or orphan. This is limited to provisioned capacity only.
- Identify used versus allocated capacity on internal volumes.
- Identify capacity trends on qtrees, look at missing quota limits, or determine whether qtrees surpass their hard or soft limits. For details about limits, see the [Qtree Capacity data model](#) on page 64.

Questions you can answer with the Capacity data mart

You can create reports answering various types of questions using this data model.

- How much capacity is there?
- How is capacity being consumed? Or not being consumed?
- How is capacity being consumed by applications and hosts in different tiers?
- What is the chargeback for allocated storage for a business unit?
- How long until I need to acquire additional storage capacity?
- Are business units aligned along the proper storage tiers?
- How is storage allocation changing over a given period of time?
- Which internal volumes have a utilization higher than a predefined threshold?
- Which internal volumes are in danger of running out of capacity based on a trend?
- What is the used capacity versus the allocated capacity on our internal volumes?
- What is the used capacity for qtrees versus the limits set per application or business entity?
- How much should we charge back each business entity for its used capacity (based on qtrees and visibility into internal volumes for NAS)?

Capacity data mart considerations

The Capacity data mart is based on the inventory data mart and provides fact and dimensional tables.

The capacity dimension tables are not shared with the inventory database, because the inventory has only the current data in its tables and the capacity data mart needs to maintain history by using a slowly changing dimension type (SCD) in its dimension tables. (For details about "SCD 2", see http://en.wikipedia.org/wiki/Slowly_changing_dimension#Type_2.)

Capacity concepts

The Storage Capacity data model and the Capacity data mart both use various concepts.

- Physical capacity: Capacity actually present on the storage device.
- Provisioned capacity: Capacity promised to storage initiators. Often used for chargeback.
- Usable capacity: Capacity that is presented after RAID.
- Raw capacity: Capacity that is presented before RAID.
- Backend capacity: Total capacity of backend LUNs. Relevant for virtual arrays only.

Note: Backend capacity is used in the Storage Capacity data model, but not the Capacity data mart.

For example, let's say that we have the following resources:

- A 100 GB volume allocated on a thin-provisioning storage pool
- A host that uses 50 GB on the volume
- A storage pool that is on top of a RAID 5 group with a 20% overhead ratio

The following table displays the example's capacity concepts and how they relate to each other:

Type of capacity	Usable Capacity	Raw Capacity
Physical	50 GB	60 GB
Provisioned	100 GB	120 GB

Capacity data mapped to the OnCommand Insight Client

The Data Warehouse capacity data maps to the data in the OnCommand Insight Client.

Resource	Data Warehouse	OnCommand Insight Client
Volume	Physical Capacity	Consumed Capacity
	Provisioned Capacity	Capacity
Internal Volume	Physical Capacity	Consumed Capacity
	Provisioned Capacity	Total Allocated Capacity
Storage Pool	Physical Capacity	Total Allocated Capacity
	Provisioned Capacity (Zero appears in the capacity_fact table to indicate N/A.)	N/A

OnCommand Insight Client data in Capacity data mart

Some data in the Capacity data mart needs to be defined in the OnCommand Insight Client. Use this table to determine what needs to be set in the Client to be able to see that information in Capacity data mart. "NA" indicates that the object does not have an association.

For example, to determine where the business unit dimension is coming from, look at the Business Unit column, look at the object row, and see if business unit information comes from an annotation on the object or from the business unit attribute on the application.

The following table lists from where in the OnCommand Insight Client the Capacity data mart dimensions are derived.

Object	Capacity Type	Host	Application	Business Unit	Tier
Storage	<ul style="list-style-type: none"> Unused Failed Spare 	N/A	N/A	N/A	Annotation on storage
Storage Pool	<ul style="list-style-type: none"> Storage Pool Other Storage Pool Unused Reserved 	N/A	N/A	N/A	Annotation on storage pool (Array virtualization - tier of representative backend volume)
Internal Volume	<ul style="list-style-type: none"> Data Files Internal Volume Reserved Snapshot Used Snapshot Unused 	N/A	N/A	Annotation on internal volume. (Replication - target inherits from source)	Annotation on internal volume (Array virtualization - tier of representative backend volume)
Volume	<ul style="list-style-type: none"> Accessed Masked 	Representative FC Host	Can be assigned directly on volume	Business unit of the representative application	Annotation on volume (Array virtualization - tier of representative backend volume)
Volume	Protection	Derived from primary volume	Derived from primary volume	Derived from primary volume	Annotation on volume (Array virtualization - tier of representative backend volume)

Object	Capacity Type	Host	Application	Business Unit	Tier
Volume	Volume Unused	N/A	N/A	N/A	Annotation on volume (Array virtualization - tier of representative backend volume)
Volume	Backend Volume Other	N/A	N/A	N/A	Annotation on volume
VM	VM Used	The host owning the VM	Annotation on VM	Annotation on VM or from host application	<ul style="list-style-type: none"> RDM - Annotation on volume Non-RDM - same as "data store Unused"
Data store	Data store Unused	Representative ESX Host	Representative ESX host application (application in host view)	Business unit of the representative application	<ul style="list-style-type: none"> SAN - annotation on representative volume NAS - annotation on representative internal volume
Data store	Data store Other	Representative ESX Host	Representative ESX host application (application in host view)	Business unit of the representative application	<ul style="list-style-type: none"> SAN - annotation on representative volume NAS - annotation on representative internal volume

Capacity data mart components

The Capacity data mart contains fact and dimension tables.

- Fact tables. A Capacity fact table records capacity facts. The Capacity data mart also includes a VM Capacity fact table and the Internal Capacity fact.
- Dimension tables that describe the criteria by which these capacity measurements are divided. Dimensions make the measurements in the fact table significant, because they answer the questions of what, when and where for the data.

Capacity data mart fact tables

The Capacity data mart includes these fact tables.

- Capacity fact table
- Internal Volume Capacity fact table

- Qtree Capacity fact table
- Chargeback fact table
- VM Capacity fact table
- Volume fact current table

Capacity Fact table

The Capacity Fact table holds capacity facts (measurements of the physical and provisioned usable and raw capacity) recorded in the data mart. This information is kept in the context of various dimension tables, which describe the criteria by which OnCommand Insight Data Warehouse divides these capacity measurements.

The Capacity Fact table includes a business unit and a business entity column that enables you to show capacity per business unit or business entity.

The Capacity data mart includes the Capacity Fact table, which holds historical data. This table includes these views:

- Capacity Fact version 5: Provides a view to the capacity_fact that makes is compatible with the fact table of Data Warehouse 5.1. This is included for backward compatibility purposes.
- Capacity Fact Current: Provides a view to only the current data that exists in capacity_fact. Capacity_fact holds historical data.

Internal Volume Capacity fact table

The Internal Volume Capacity fact table provides information about internal volume capacity.

Qtree Capacity fact table

This table provides information about qtree capacity.

The qtree capacity fact table has one row per qtree with used quota > 0, while the chargeback fact table has all the data necessary for chargeback in a qtree-based environment. The chargeback table also has volumes and internal volumes.

Chargeback fact table in data marts

This table provides information about the provisioned capacity and used capacity by service level and tiers for internal volumes, qtrees, and volumes. This information would be helpful in identifying costs by application or business entity (tenant, line of business, business unit, or project).

VM Capacity fact table

The VM Capacity fact table displays the breakdown of capacity provisioned to virtual machines and data stores. With this fact table you can generate reports that show information on Chargeback and Capacity not used by VM.

- Chargeback: Review capacity chargeback based on capacity provisioned to VMs and data stores. You can achieve more granular capacity usage by using dimensions such as host, application, storage, virtual storage, internal volume, VM, data store, and tier dimensions.
- Capacity not used by VM: Review unused capacity, datastore Unused, or datastore Other.

The VM Capacity fact table displays the following measurements as "usable MB":

- Actual capacity: The amount of data store capacity being used. This measurement does not count savings achieved in the backend storage using thin provisioning and deduplication.

- **Provisioned capacity:** The amount of capacity that was allocated for a VM. This measurement includes VMDKs and other VM-related files.

The VM Capacity fact table does not correlate provisioned to physical capacity, including the amount of physical capacity consumed by a VM and the amount of unused VM volumes or LUNs.

The VM Capacity fact table includes several fields, including the following:

- **datastore Type:** The following values could appear:
 - NAS when a data store is mapped to a share through the NFS/CIFS protocol.
 - SAN when the data store is mapped to volumes.
 - UNKNOWN when the data store is not connected to a storage.

Capacity Types related to VM capacity

Various capacity types are related to the VM Capacity fact table.

The Capacity types and descriptions are as follows:

VM Used

The amount of data store capacity used by VMDK and other VM-related files (such as swap and Snapshot copy files).

This is further categorized into:

Actual

- Thick VMDKs, the size of the VMDK
- Thin VMDK, the amount of capacity consumed from the data store
- RDM VMDK, the size of the VMDK
- Other files, the size of the thick provisioned files

Provisioned: The size of the VMDK or file.

Data store free

The amount of capacity not used in a VMFS data store:

- Actual: The data store's free capacity.
- Provisioned: Not Applicable. In NFS data stores, the data store free equals the amount of free-space in the internal volume that can be shared with other applications and data stores.

Data store other

The amount of data-store capacity that is used but isn't accounted for any VM. Applicable for VMFS data-stores only.

- Actual: The amount of used data store capacity that is not associated with any VM.
- Provisioned: Not applicable.

Volume fact current table

The volume_fact_current table maps capacity fact rows to the volumes that are rolled up into them. This table provides volume usage data for non-VM volumes only.

This table can be used to:

- Show the volumes included in given a certain amount of capacity presented in a report (for a specific storage, host, tier, and capacity type).

- Given one or more volumes, show where its capacity is used. Indicate storage (or backend storage for virtual volumes), hosts, applications, and its capacity type (for example, accessed, masked, or unused).

For capacity that is not virtualized (not used by VMware and not by storage virtualization), there is a row in the table for each volume. This row indicates this volume's capacity. For capacity that is used by storage virtualization, there are rows included in the table for the frontend volumes only. (Backend volumes do not appear in the table, as their capacity appears as part of the frontend volumes.) There might be more than one row in the table per frontend volume. For example, a frontend volume is mapped to backend volumes that reside on two different backend arrays or assigned different tiers. In that case, the following appears in the capacity columns:

- `physicalUsable` and `provisionedUsable`: The physical and provisioned usable capacity of the frontend volume.
- `physicalRaw` and `provisionedRaw`: The physical and provisioned raw capacity of the frontend volume (usually 0).

The `volume_fact_current` table is limited to mapping current fact rows only (see the latest column in the `date_dimension`).

This table does not support mapping of capacity used by virtual machines.

Capacity data mart dimensions and lists

By using different or multiple dimension attributes (seen in columns in the dimension tables), you can construct SQL queries that access data at a higher or lower level of detail for each dimension in the Capacity data mart.

The Capacity data mart includes the following dimensions:

- Application dimension (`application_dimension`)
- Application group dimension (`application_group_dimension`)
- Business entity dimension (`business_entity_dimension`)
- Capacity type list (`capacity_type`)
- Date dimension (`date_dimension`)
- Data store dimension (`datastore_dimension`)
- Host dimension (`host_dimension`)
- Host group dimension (`host_group_dimension`)
- Internal volume dimension (`internal_volume_dimension`)
- Qtree dimension (`qtree_dimension`)
- Service level dimension (`service_level_dimension`)
- Storage dimension (`storage_dimension`)
- Storage pool dimension (`storage_pool_dimension`)
- Tier dimension (`tier_dimension`)
- Virtual machine dimension (`vm_dimension`)
- Volume dimension (`volume_dimension`)

Business entity (Tenant, Line of Business, Business Unit, and Project) information is obtained as columns in the Capacity Fact table.

Related concepts

[Capacity Fact table](#) on page 113
[Application dimension in data marts](#) on page 154
[Application group dimension in data marts](#) on page 154
[Date dimension in data marts](#) on page 154
[Hosts and host group dimensions](#) on page 156
[Internal volume dimension in data marts](#) on page 157
[Qtree dimension in data marts](#) on page 157
[Service Level dimension in data marts](#) on page 157
[Storage dimension in data marts](#) on page 157
[Storage pool dimension in data marts](#) on page 157
[Tier dimension in data marts](#) on page 158
[VM dimension in data marts](#) on page 158

Capacity type list

This defines the different Capacity Types. Capacity types are mutually exclusive. This means that every single byte on the disk is counted only once using a single capacity type.

The capacity data mart supports the following capacity types:

Data store Other

Capacity at the data store level that is not specified by OnCommand Insight, for example, overhead attributed to different technologies.

Data store Unused

Capacity at the data store level that is not used by any virtual disk.

VM Used

Capacity of virtual disks that are used by virtual machines. (The host in this capacity type is the host running/owning the VM. The host group is always a group with a single host in it.)

Array virtualization and capacity data mart

Virtual volumes are treated differently than other volumes in the environment in several aspects.

- Storage dimension: Points to the backend storage array. Array virtualization is supported when the front-end array supports thin provisioning.
- Tier dimension: Points to the tier defined on the backend volumes.
- Virtual storage dimension: Points to the frontend storage array.

The following data is available:

- Mapping between the storage pool and the backend LUNs from which it is made. (The backend LUNs serve as disks for the storage pool.)
- Mapping between the backend LUN and the volume in the backend storage that it reflects.
- Mapping between the virtual volume, the backend LUN, and the volume in the backend storage that the backend LUN reflects.

The raw capacity calculation is based on the RAID level of the backend volumes assuming no overhead is added by the virtualizer. The raw capacity is calculated based on the multiplication of the usable-to-raw ratio of the storage pool in the frontend array and the usable-to-raw ratio of the backend volumes. For example, if the usable-to-raw ratio of the frontend storage pool is 0.5 and the

usable-to-raw ratio of the backend volumes is 0.75, the usable-to-raw ratio for storage objects on top of the virtualized storage pool is 0.75 times 0.5.

The implementation supports a configuration where a frontend volume is made of backend volumes that reside on different backend storage arrays and/or have different tiers. In that case, the capacity of the virtual volume is broken according to the ratio of capacity of the backend volumes so that a percentage of the frontend volume's capacity is associated with each applicable backend storage array or tier.

Capacity on backend volumes that is not included in the frontend volumes is classified with capacity type **Backend Volume Other** and has the backend storage array in the storage dimension, the frontend array in the virtual storage dimension, and the backend volumes tier in the tier dimension. Capacity might not be included in frontend volumes either because it is connected to the frontend storage (however, not all the capacity is carved into frontend volume) or because the virtualization process requires some overhead of capacity.

Capacity Efficiency data mart for SQL queries

The Capacity Efficiency data mart (`dwh_capacity_efficiency`) tracks the storage efficiency and potential gain or loss over time based on efficiency technologies. You can use the Capacity Efficiency data mart to measure the improvement achieved using efficiency technologies and determine which technologies yield the greatest improvements.

The Capacity Efficiency data mart estimates the gain and loss (or the efficiency) of different technologies used in the storage. For example, you can determine how capacity is saved by using storage efficiency technologies, such as thin provisioning, deduplication, and RAID.

Administrators look for technologies that improve the storage capacity consumption by increasing the ratio between what is allocated to the consumer versus how much raw capacity is needed to support the allocation. When administrators improve storage capacity efficiency, they improve this ratio.

Questions you can answer with the Capacity Efficiency data mart

You can answer various questions using the Capacity Efficiency data mart.

- What is our storage efficiency savings as a result of implementing thin provisioning and deduplication technologies?
- What are the storage savings across data centers?
- Based on historical capacity trends, when do we need to purchase additional storage?
- What would be the capacity gain if we enabled technologies such as thin provisioning and deduplication?
- Regarding storage capacity, am I at risk now?

Capacity producers and consumers

Different storage technologies can yield increases (gains) or decreases (losses) to capacity.

The following table describes some of these technologies and indicates whether the technology represents a gain or a loss to capacity.

Storage Technology	Gain or Loss	Notes
Disk failure protection	Loss	Represents the amount of capacity that is used to protect disk groups by RAID level. Example: RAID-4 where a single disk that is using every disk group for parity. Summing the size of the parity disk generates capacity loss.
Deduplication	Gain	Represents data compression or saved capacity in a compression ratio. Example: In a 2:1 compression ratio, for every 1 GB compressed capacity, there should be 2 GB of uncompressed capacity. This means that more space is needed to store the same data without deduplication.
Thin provisioning	Gain	Helps to eliminate over provisioning by consuming capacity only from the capacity that is used by the volume.
Clones	Gain	The clone shares capacity with its source volume and consumes only the capacity that is not mutual between itself and the source volume.

Storage efficiency calculation

Storage efficiency is calculated with a specific mathematical formula.

The efficiency of storage is calculated by taking the total storage raw capacity and adding the gainers, then subtracting the consumers. Finally, divide the total by the total storage raw capacity as follows:

$$\frac{\text{TotalArrayCapacity} + \text{GainedCapacity} - \text{LostCapacity}}{\text{TotalArrayCapacity}} \times 100$$

Capacity Efficiency data mart considerations

For each storage system, certain parameters are collected over time and stored in the fact and dimension tables.

Those parameters are:

- Raw Capacity: Storage raw capacity.
- Backend Capacity: Total capacity of backend LUNs. Relevant for virtual arrays only.
- For each technology:
 - Name: Name of the technology, such as deduplication or clone
 - Gain: The total capacity gained in this storage by using the specified technology, for example, 100 MB of redundant blocks saved by implementing deduplication.
 - Loss: The total capacity lost in this storage due to the specified technology, for example, 10 GB that was used as spare.
 - Potential gain: The total capacity potentially gained in this storage if the specified technology is used to the full extent, for example, by enabling deduplication in all volumes, a total of 300 MB can be saved.

- Potential loss: The total capacity potentially lost in this storage if the specified technology is used to the full extent.
- Each storage resource has one row for the storage itself:
 - Raw Capacity & Backend Capacity are specified.
 - Technology column is NULL.
 - Gain, Loss, Potential Gain and Potential Loss are set to 0.
- Plus a row for each relevant technology:
 - Raw Capacity & Backend Capacity are set to 0.
 - Technology column is specified.
 - Gain, Loss, Potential Gain and Potential Loss are set to the relevant values.

Follow these rules when you develop queries using the Capacity Efficiency data mart:

- If no technology is reported for a storage, only one row is included with the Raw Capacity & Backend Capacity columns.
- All capacity values are in MBs.
- To include additional columns or filter your query, use the storage and date dimensions.
- Different storage arrays might have different number of supported technologies and thus a different number of rows.
- No efficiency data is available in the inventory.

The data mart is based on the Inventory data mart and provides fact and dimensional tables. It contains:

- A Fact table in which capacity facts are recorded.
- Dimension tables that describe the criteria by which these capacity measurements are divided.

Capacity Efficiency data mart fact table

The Capacity Efficiency data mart includes the Efficiency Fact table, which provides information about backend capacity, raw capacity, technology used, and capacity gain or loss based on technologies used.

The gain or loss indicates the result of using storage efficiency and protection technologies such as deduplication, thin provisioning, or protection. Protection types used on the storage include mirrored or striped. The gain or loss data is null when it represents the storage's raw and backend capacity (one row storage).

Capacity efficiency data mart dimension tables

The Capacity Efficiency data mart includes date and storage dimension views.

- Date dimension
- Storage dimension, which includes information such as model, manufacturer, version, and data center

Related concepts

[Date dimension in data marts](#) on page 154

[Storage dimension in data marts](#) on page 157

Capacity Efficiency data mart sample SQL query

This query lists all the storage arrays and their efficiency score.

```
SELECT
    s.name,
    (SUM(f.rawCapacityMB)+SUM(f.gainMB)-SUM(f.lossMB))/
    SUM(f.rawCapacityMB)
    AS efficiency

FROM
    dwh_capacity_efficiency. efficiency_fact f,
    dwh_capacity_efficiency.date_dimension d,
    dwh_capacity_efficiency.storage_dimension s

WHERE
    d.latest = TRUE AND
    f.dateTk = d.tk AND
    f.storageTk = s.tk

GROUP BY s.id
ORDER BY s.name
```

File System Utilization data mart for SQL queries

File System Utilization data mart (dwh_fs_util) includes file system utilization data. This data mart enables visibility into capacity utilization by hosts on the file system level. Administrators can determine allocated and used capacity per file system, the type of file system, and identify trending statistics by file system type.

The File System Utilization data mart tracks capacity utilization of file systems in hosts over time.

Questions you can answer with the File System Utilization data mart

Using the File System Utilization data mart, you can create queries that can answer various questions.

- What is the size of the file system?
- Where is the data kept and how is it accessed, for example, local or SAN?
- What are the historical trends for the file system capacity? Then, based on this, what can we anticipate for future needs?

FS Utilization data mart considerations

The File System Utilization data mart contains information from fact and dimension tables.

FS Utilization data mart fact table

The FS Utilization Fact table describes file system utilization. It includes the size of the file system and the amount of file system capacity that is used.

FS Utilization data mart dimension tables

The File System Utilization data model includes several dimension tables.

- Date dimension
- File System dimension
- Host dimension

Related concepts

[Date dimension in data marts](#) on page 154

[File System Dimension in data marts](#) on page 156

[Host dimension in data marts](#) on page 156

FS Utilization data mart sample SQL query

This query lists the host, the file systems on this host, and the allocated and used capacity for each file system.

```
SELECT
    h.name AS HostName,
    fsd.name AS FileSystemName,
    fsd.location AS Location,
    fsd.type AS FileType,
    sizeMB AS AllocatedCapacity,
    UsedMb AS UsedCapacity

FROM
    dwh_fs_util.fs_util_fact fs,
    dwh_fs_util.file_system_dimension fsd,
    dwh_fs_util.date_dimension dd,
    dwh_fs_util.host_dimension h

WHERE
    dd.latest = 1 AND
    fs.dateTk = dd.tk AND
    fs.fsTk = fsd.tk AND
    h.tk=fs.hosttk

GROUP BY h.name
ORDER BY h.name
```

Performance data mart for SQL queries

The Performance data mart (dwh_performance) collects performance information for volumes, internal volumes, applications and hosts. Volume and internal volume performance data is aggregated hourly and daily, while application and host performance data is aggregated hourly.

The Performance data mart includes fact and dimensional tables. Unlike the Inventory data mart, historical data is contained in the Performance data mart.

Note: The Performance data mart functions only if you have the OnCommand Insight Perform license installed.

Questions you can answer with the Performance data mart

Using this data model, you can create reports that answer several types of questions.

- What volumes or internal volumes have not been used or accessed during a specific period?
- Can we pinpoint any potential misconfiguration for storage for an application (unused)?
- What was the overall access behavior pattern for an application?
- Are tiered volumes assigned appropriately for a given application?
- Could we use cheaper storage for an application currently running without impact to application performance?
- What are the applications that are producing more accesses to currently configured storage?

When you use the switch performance tables, you can obtain the following information:

- Is my host traffic through connected ports balanced?
- Which switches or ports are exhibiting a high number of errors?
- What are the most used switches based on port performance?
- What are the underutilized switches based on port performance?
- What is the host trending throughput based on port performance?
- What is the performance utilization for last X days for one specified host, storage system, tape, or switch?
- Which devices are producing traffic on a specific switch (for example, which devices are responsible for use of a highly utilized switch)?
- What is the throughput for a specific business unit in our environment?

When you use the disk performance tables, you can obtain the following information:

- What is the throughput for a specified storage pool based on disk performance data?
- What is the highest used storage pool?
- What is the average disk utilization for a specific storage?
- What is the trend of usage for a storage system or storage pool based on disk performance data?
- What is the disk usage trending for a specific storage pool?

Performance data mart considerations

The Performance data mart is based on the inventory data mart and provides fact and dimensional tables.

The performance dimension tables are not shared with the inventory database, because the inventory has only the current data in its tables and the Performance data mart needs to maintain history by using a slowly changing dimension type (SCD) in its dimension tables.

All throughput measurements are in megabytes per second.

The Performance fact tables include the following key data:

readResponseTime

The maximum read response time for a volume (internal volume, application volume, or host volume).

writeResponseTime

The maximum write response time for a volume (internal volume, application volume, or host volume) for the time specified. For applications and hosts, this is the maximum write response time for all volumes accessed by the application or host during the summary hour.

totalResponseTime

The maximum total response time for all volumes (internal volumes, application volumes, or host volumes) for the time specified.

totalResponseTimeMax

The maximum response time (peak) for a volume (internal volume, application volume, or host volume) for the time specified.

readThroughput

The average read throughput for a volume (internal volume, application volume, or host volume) for the time specified. For application and hosts, this displays the sum of the read throughputs of all accessed volumes during the hour.

writeThroughput

The average write throughput for a volume (internal volume, application volume, or host volume) for the time specified. For application and hosts, this displays the sum of the write throughputs of all accessed volumes during the hour.

totalThroughput

The readThroughput plus the writeThroughput.

totalThroughputMax

The maximum throughput (peak) for a volume (internal volume) throughout the day. Not applicable to application and host fact tables.

readIops

The average read IOPS for a volume (internal volume, application volume, or host volume) for the time specified. For application and hosts, this displays the sum of the read IOPS of all accessed volumes during the hour.

writeIops

The average write IOPS for a volume (internal volume, application volume, or host volume) for the time specified. For application and hosts, this displays the sum of the write IOPS of all accessed volumes during the hour.

totalIops

The readIops plus the writeIops.

totalIopsMax

The maximum total IOPS (peak) for a volume (internal volume) for the time specified. Not applicable to application and host fact tables.

readCacheHitRatio

The average read cache hit ratio for a volume (internal volume, application volume, or host volume) for the time specified.

writeCacheHitRatio

The average write cache hit ratio for a volume (internal volume, application volume, or host volume) for the time specified.

totalCacheHitRatio

The average total cache hit ratio as the sum of readCacheHitRatio with writeCacheHitRatio for all volumes (internal volume, application volume, or host volume) for the time specified.

writePending

The maximum number of pending write requests for a volume (internal volume, application volume, or host volume) for the time specified.

accessed

The accessed status (1 for accessed or 0 for not accessed) of a volume (internal volume) for the time specified. Not applicable to application and host fact tables.

sumOfAveragesVolumeThroughput

The sum of all averages of throughput for all application volumes and host volumes for the time specified.

maxOfMaxVolumeThroughput

The maximum throughput (peak) for all application volumes or host volumes for the time specified.

sumOfMaxVolumeThroughput

The sum of all maximum throughput for all application volumes or host volumes.

The Disk Performance fact tables include the following key data:

readThroughput

Disk read throughput average (MB/s).

writeThroughput

Disk write throughput average (MB/s).

totalThroughput

Disk total throughput average (READ+WRITE, MB/s).

totalThroughputMax

Maximum disk total throughput (MB/s).

readIops

Disk read I/O per second.

writeIops

Disk write I/O per second.

totalIops

Total disk I/O per second average (READ+WRITE).

totalIopsMax

Maximum disk I/O per second.

readUtilization

Disk read utilization (%).

writeUtilization

Disk write utilization (%).

totalUtilization

Total disk utilization (%).

totalUtilizationMax

Maximum disk utilization (%).

accessed

Indication whether the disk was accessed.

Performance data mart components

The Performance data mart contains fact and dimension tables.

- Fact tables. A fact table records performance facts.
- Dimension tables that describe the criteria by which these measurements are divided. Dimensions make the measurements in the fact table significant, because they answer the questions of what, when, and where for the data.

Performance data mart fact tables

The Performance data mart includes various fact tables.

Application Volume Hourly Performance fact table

Collects summaries of performance numbers for application by tier by hour based on volume performance data.

Disk Daily Performance fact table

Contains daily aggregated data for a specific disk in the system. Use this fact table to relate a disk to its owning storage system and storage pool.

Disk Hourly Performance fact table

Contains hourly aggregated data for a specific disk in the system. Use this fact table to relate a disk to its owning storage system and storage pool.

Host VM Daily Performance fact table

Summarizes performance information per VM host by day.

Host VM Hourly Performance fact table

Summarizes performance information per VM host by hour.

Internal Volume Daily Performance fact table

Summarizes performance information per internal volume and day. This fact table includes a dimension to virtual storage.

Internal Volume Hourly Performance fact table

Summarizes performance information per internal volume and hour. This fact table includes a dimension to virtual storage.

Storage Node Daily Performance fact table

Summarizes performance information per storage node and day.

Storage Node Hourly Performance fact table

Summarizes performance information per storage node and hour.

Switch Performance for Host Hourly Performance fact table

Contains hourly rollups of port performance data aggregated per host. Each line represents one hour of port performance for ports connected to a specific host.

Switch Performance for Port Hourly Performance fact table

Contains hourly rollups of port performance per switch port. Each line represents one hour of port performance for one port within a switch.

Switch Performance for Storage Hourly Performance fact table

Contains hourly rollups of port performance data aggregated per storage. Each line represents one hour of port performance for ports connected to a specific storage.

Switch Performance for Tape Hourly Performance fact table

Contains hourly rollups of port performance data aggregated per tape. Each line represents one hour of port performance for ports connected to a specific tape.

VM Daily Performance fact table

Summarizes performance information per VM by day.

VM Hourly Performance fact table

Summarizes performance information per VM by hour.

Volume Daily Performance fact table

Summarizes performance information per volume and day. This fact table includes a dimension to virtual storage.

Volume Hourly Performance fact table

Summarizes performance information per volume and hour. This fact table includes a dimension to virtual storage.

Performance data mart dimensions and lists

By using different or multiple dimension attributes (seen in columns in the dimension tables), you can construct SQL queries that access data at a higher or lower level of detail for each dimension in the Performance data mart.

The Performance data mart submarts include the following dimensions. To identify which of the submarts include each of the dimensions, see the schema diagrams in Data Warehouse.

- Application dimension (application_dimension)
- Application group dimension (application_group_dimension)
- Business Entity dimension (business_entity_dimension)
- Connected Device dimension (connected_device_dimension)
- Date dimension (date_dimension)
- Disk dimension (disk_dimension)
- Fabric dimension (fabric_dimension)
- Host dimension (host_dimension)
- Host group dimension (host_group_dimension)
- Internal volume dimension (internal_volume_dimension)
- Port dimension (port_dimension)
- Service level dimension (service_level_dimension)
- Storage dimension (storage_dimension)
- Storage pool dimension (storage_pool_dimension)
- Switch dimension (switch_dimension)
- Tape dimension (tape_dimension)
- Tier dimension (tier_dimension)
- Time dimension (time_dimension)
- Volume dimension (volume_dimension)

Related concepts

[Application dimension in data marts](#) on page 154

[Application group dimension in data marts](#) on page 154

[Business Entity dimension in data marts](#) on page 154
[Connected Device dimension in data marts](#) on page 154
[Date dimension in data marts](#) on page 154
[Fabric dimension in data marts](#) on page 156
[Hosts and host group dimensions](#) on page 156
[Internal volume dimension in data marts](#) on page 157
[Port dimension in data marts](#) on page 157
[Service Level dimension in data marts](#) on page 157
[Storage dimension in data marts](#) on page 157
[Storage pool dimension in data marts](#) on page 157
[Tape dimension in data marts](#) on page 158
[Tier dimension in data marts](#) on page 158
[Time dimension in data marts](#) on page 158
[Volume dimension in data marts](#) on page 158

Performance data mart sample SQL query

Various SQL queries can be executed using the Performance data mart.

Performance data mart SQL Query Example

The following query identifies which volumes were not accessed since beginning of the year. To help troubleshoot the nonuse, you can list the business units and applications that are supposed to be accessing the volume. This helps you identify potential application-to-volume misconfigurations.

```

SELECT
    GROUP_CONCAT(DISTINCT COALESCE(ad.businessUnit, 'N/A' ))
    AS businessUnit,
    GROUP_CONCAT(DISTINCT ad.name) AS applicationName,
  
```

```

sd.name AS storageName,
vd.name AS volumeName

FROM
    dwh_performance.volume_daily_performance_fact vpf

JOIN
    dwh_performance.volume_dimension vd

ON
    vd.tk = vpf.volumeTk

JOIN
    dwh_performance.storage_dimension sd

ON
    sd.tk = vpf.storageTk

JOIN
    dwh_performance.application_dimension ad

ON
    ad.tk = vpf.applicationTk

WHERE
    vpf.timestamp BETWEEN UNIX_TIMESTAMP('2011-01-01')*1000 AND
    UNIX_TIMESTAMP('2011-01-31')*1000

GROUP BY
    vd.id

HAVING
    MAX(vpf.accessed) = 0

```

That example results in the following:

businessUnit	applicationName	storageName	volumeName
Data Warehouse Unit	ONC	DAT1_12555	0B:13
Finance	DS	DAT1_12666	0B:02
Eng_Services	VMware ESX	DAT4_65555	B1:11
MEI_BU	BU_reporting	DAT4_76777	53:00
IT	Exchange	DAT10_48888	91:13
IT	Emp_DB	26000@192.168.1.444	11:06

Port Capacity data mart for SQL queries

The OnCommand Insight Port Capacity data mart (dwh_ports) enables access to an open database that facilitates generation of switch port capacity planning reports.

The Port Capacity data mart provides current and historical data on items such as switch ports connectivity, status, and speed. This data can be used when planning new switch purchases.

Note: The Port Capacity data mart is functional only if you have the OnCommand Insight Plan license installed.

Questions you can answer with the Port Capacity data mart

Using the Port Capacity data mart, you can answer several questions.

- How can I create a port consumption forecast that predicts resource (port) availability (according to data center, switch vendor and port speed)?
- Which ports are likely to run out of capacity, providing data speed, data center, vendor and number of Host and storage ports?
- What are the switch port capacity trends over time?
- What are the port speeds?
- What type of port capacity is needed and which organization is about to run out of a certain tier of port type or vendor?
- What is the optimal time to purchase capacity and make it available?

Port Capacity data mart considerations

The port capacity data mart is based on the inventory data mart and provides fact and dimensional tables.

The following are contained in this data mart:

- A Fact table in which port facts are recorded.
- Dimension tables that describe the criteria by which these port measurements are divided.

Port Capacity data mart fact table

The Port Capacity data mart includes the Ports Fact table, which holds historical data about the port and its connectivity. The Ports Fact table has the ports_fact_current view, which provides a view only to the current data that exists in the ports_fact table.

The Ports Fact table includes the following data:

- Fabric
- Port speed
- Connectivity type
- Port type, typically used to identify the protocol used to communicate with other devices, for example, F-port or E-port
- Port status, such as licensed or inactive.

Port Capacity data mart dimension tables

The Port Capacity data mart is divided into several dimensions.

- Date dimension: Defines the time (determined according to the schedule in the OnCommand Insight Data Warehouse Portal) from which the measurements are taken.
- Port dimension: Describes a switch port.
- Switch dimension: Defines the switch owning the port.

Related concepts

[Connected to Dimension](#) on page 141

[Date dimension in data marts](#) on page 154

[Port dimension in data marts](#) on page 157

[Switch dimension in data marts](#) on page 158

Port Capacity data mart sample SQL query

This query lists the switch port capacity by data center.

```
SELECT
pf.speed AS 'Speed',
sd.manufacturer AS Manufacturer,
sd.datacenter AS DataCenter,
pf.connectivitytype AS ConnectivityType,
COUNT(pf.connectivitytype) AS PortCapacity
```

```
FROM
dwh_ports.ports_fact pf,
dwh_ports.switch_dimension sd,
dwh_ports.date_dimension dd
```

```
WHERE
sd.Tk=pf.SwitchTk
AND pf.datetk=dd.tk
AND pf.connectivitytype <> 'free'
AND dd.latest=1
```

```
GROUP BY
sd.datacenter, sd.manufacturer,pf.speed
```

```
ORDER BY speed
```

This example results in the following:

Speed	Manufacturer	DataCenter	ConnectivityType	PortCapacity
1G	Brocade	N/A	Host	12
1G	Brocade	Data Center II	Host	5
2G	Brocade	N/A	Host	721
2G	McData	N/A	Switch	394
2G	McData	Data Center 1	Switch	197

Speed	Manufacturer	DataCenter	ConnectivityType	PortCapacity
2G	McData	Data Center II	Switch	102
2G	Brocade	Data Center II	Host	334
2G	Brocade	Data Center 1	Host	338
4G	McData	Data Center 1	Host	103
4G	McData	Data Center II	Host	17
4G	Brocade	Data Center II	Storage	17
4G	Brocade	N/A	Host	438
4G	Brocade	Data Center 1	Host	37
4G	McData	N/A	Storage	39

Storage and Storage Pool Capacity data mart for SQL queries

The Storage and Storage Pool Capacity data mart (storage_and_storage_pool_capacity_fact) enables access to an open database that facilitates generation of storage and storage pool capacity reports.

You can create reports about storage capacity resource planning, including physical and virtual storage pool data for storage and storage pools by using the Storage and Storage Pool Capacity data mart.

Questions you can answer with the Storage and Storage Pool Capacity data mart

You can answer various questions using the Storage and Storage Pool Capacity data mart.

- What is the projected date for reaching the capacity threshold of 80% of my physical storage?
- What is the physical storage capacity on an array for a given tier?
- What is the storage utilization trend on an array for all of the tiers?
- What are my top 10 storage systems with the highest utilization?
- What is the storage utilization trend for the storage pools?
- How much capacity is already allocated?
- What capacity is available for allocation?

Storage and Storage Pool Capacity data mart considerations

The Storage and Storage Pool Capacity data mart facilitates the generation of storage and storage pool capacity reports.

The capacity dimension tables are not shared with the inventory database, because the inventory database has only the current data in its tables and the capacity data mart needs to maintain history.

Note: This data mart does not include business entity information and so you cannot use the multiple tenancy filter in reports.

The data mart includes the following information:

- Dimensions
 - Storage
 - Storage Pool
 - Tier
 - Date
- Measures
 - backendCapacity: True when the storage is backend
 - virtualCapacity: True when the storage is virtual
 - capacityMB: Capacity allocated for a storage-pool as reported by the storage-array in MB
 - RawCapacityMB: capacityMB * rawToUsableRatio
 - usedCapacityMB: Capacity used in a storage-pool as reported by the storage-array in MB
 - usedRawCapacityMB: UsedCapacityMB * rawToUsableRatio
 - snapshotUsedCapacityMB: Capacity used for storage-pool level snapshots in MB
 - snapshotUsedRawCapacityMB: snapshotUsedCapacityMB * rawToUsableRatio

Related concepts

[Date dimension in data marts](#) on page 154

[Storage dimension in data marts](#) on page 157

[Storage pool dimension in data marts](#) on page 157

[Tier dimension in data marts](#) on page 158

Storage and Storage Pool Capacity data mart sample SQL query

Various SQL queries can be executed in the Storage and Storage Pool Capacity data mart.

Storage and Storage Pool Capacity data mart SQL Query Example 1

The following query retrieves capacity and raw capacity for all storage arrays.

```
SELECT
  sd.name AS 'Storage Name',
  SUM(spcf.capacityMB) AS 'Capacity MB',
  SUM(spcf.rawCapacityMB) AS 'RawCapacity MB'

FROM
  dwh_capacity.storage_and_storage_pool_capacity_fact spcf
  JOIN dwh_capacity.storage_dimension sd
    ON spcf.storageTk = sd.tk
  JOIN dwh_capacity.date_dimension d
    ON d.tk=spcf.dateTk
    AND d.latest = 1

GROUP BY sd.name
```

This example results in the following:

Storage Name	Capacity MB	RawCapacity MB
AMS2500_1	77999576	87999576
DS8300_5	636887630	1344733480
ml1	2040599168	4328142336

Storage and Storage Pool Capacity data mart SQL Query Example 2

The following query obtains all storage pools, capacity usage, and raw capacity.

```
SELECT
    spd.name AS 'Storage Pool Name',
    SUM(spcf.capacityMB) AS 'Capacity MB',
    SUM(spcf.rawCapacityMB) AS 'RawCapacity MB'

FROM
    dwh_capacity.storage_and_storage_pool_capacity_fact spcf
    JOIN dwh_capacity.storage_pool_dimension spd
      ON spcf.storagePoolTk = spd.tk
    JOIN dwh_capacity.date_dimension d
      ON d.tk=spcf.dateTk
      AND d.latest = 1

GROUP BY spd.name
```

This example results in the following:

Storage Pool Name	Capacity MB	RawCapacity MB
B001	10192896	11683098
B002	6432768	7385787
B003	1302016	1494613

Volume Capacity data mart for SQL queries

Use this data mart to report on all aspects of volumes in your environment and look at them by vendor, model, tier, service level, and data center. You can view capacity related to orphan volumes, unused volumes, and protection volumes (used for replication). You can also see different volume technology (iSCSI or FC) and compare virtual to non-virtual volumes for array virtualization issues.

Questions you can answer with the Volume Capacity data mart

You can answer various questions using the Volume Capacity data mart.

- Which volumes have a utilization higher than a predefined threshold?
- What is the trend for orphan volume capacity by my data center?
- How much of my data center capacity is virtualized or thin provisioned?
- How much of my data center capacity must be reserved for replication?

Volume Capacity data mart sample SQL query

Various SQL queries can be executed using the Volume Capacity data mart.

Masked Orphaned Volumes by Data Center

The following query retrieves storage name, masked orphaned volume name, and volume capacity and groups it by data center.

```
SELECT
    sd.datacenter AS DataCenter,
    sd.name AS Storagename,
    vd.name AS volumenname,
    cf.ProvisionedCapacityMB/1024 AS VolumeCapacity
FROM
    volume_capacity_fact cf,
    volume_dimension vd,
    storage_dimension sd
WHERE
    cf.storageTk=sd.tk AND
    cf.volumetk=vd.tk AND
    isOrphaned=1
```

This example results in the following:

Data Center	Storage	Volume Name	Volume Capacity (GB)
New York	USP-2222112	00:0	16.00
New York	USP-2222112	00:1	16.00
New York	USP-2222112	00:2	16.00
New York	USP-2222112	00:3	16.00
New York	USP-2222112	00:4	16.00
New York	USP-2222112	00:5	16.00
New York	USP-2222112	00:6	16.00
New York	USP-2222112	00:7	16.00
Sydney	Sym-Perf	01:00	32.00
Sydney	Sym-Perf	01:05	32.00

OnCommand Insight Reporting data model glossary

Use this glossary to understand terminology in OnCommand Insight data models and reports.

Data model dimension tables for Reporting Connection

OnCommand Insight includes many dimension tables that you can use in reports that you use or create with the Reporting Connection.

Application Dimension

This dimension shows application annotation, description, and priority, such as High, Medium, or Low.

- In the Chargeback data model, this table shows application information associated with the hosts that access the capacity. "N/A" is displayed for rows where the capacity has not been associated with any application.
- In the Qtree Capacity data model, this table shows application information associated with this qtree.

If there are multiple applications accessing the capacity, only one application is chosen as the representative of all of them, and that is the application shown in this table. This is done to avoid double-counting capacity when writing chargeback reports.

The fields and field descriptions are as follows:

Application Name

Name of the representative application.

Description

Description text about the application.

Priority

Priority for each application, listed in the same order as the applications are listed in the Application Name field.

Business Unit

Name of the Business Unit that owns the application. Business units are part of the business entity hierarchy. A specific application has the same Business Unit on all hosts. The value will be N/A if the application is not associated with a Business Unit.

Note: The Business Unit field only appears in the Application dimension of older data models. It was migrated to the Business Entity dimension in newer models. To obtain a more hierarchical view of business information, see the Business Entity dimension.

Related concepts

[Chargeback data model](#) on page 77

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity data model details](#) on page 65

[Switch Performance data model](#) on page 90

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Application Group Dimension

This dimension defines the set of applications that is using the capacity. In this dimension, all applications are associated with capacity. Use the table data elements to display all applications that are associated with capacity fact values, rather than just the representative application that is selected by the Data Warehouse.

The fields and field descriptions are as follows:

Application Names

Full list of applications accessing the capacity.

Cardinality

Number of applications in the application group accessing the capacity.

Representative Application

Name of the application representing the application group.

The application group shown in the table is representative. If there are multiple applications accessing the capacity, only one application group is chosen as the representative, and that is what is shown in this table. This is done to avoid double-counting capacity when writing chargeback reports.

Related concepts

[Chargeback data model](#) on page 77

[Internal Volume Capacity data model details](#) on page 57

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Application Dimension

This table shows application information associated with the hosts that access the capacity. “N/A” appears for rows where the capacity has not been associated with any application.

If there are multiple applications accessing the capacity, only one application is chosen as the representative of all of them, and that is what is shown in this table. This is done to avoid double-counting data.

The fields and field descriptions are as follows:

Application Name

Name of the representative application.

Business Unit

The business units that are using the capacity. Business units are part of the business entity hierarchy. Use this to capture previous versions of OnCommand Insight, which used business unit annotations without the business entity hierarchy.

Name of the Business Unit that owns the application. A specific application has the same Business Unit on all hosts.

Note: The value will be N/A if the application is not associated with a Business Unit.

Description

Descriptive text about the application.

Priority

Priority for each application, listed in the same order as the applications are listed in the Application Name field.

Business Entity Hierarchy

Use the Business Entity Hierarchy dimension if you want to show information in a hierarchy.

The hierarchical business entity includes the following components: tenant, business unit, and project. You do not have to specify a value for each component. Any value can be left out if not used.

Business entities can be associated with one or more applications. If a business entity is associated with the same object (Host, for example) on two different servers, but contains two different values, Data Warehouse chooses one and logs the event.

You can use the following fields in the business entity hierarchy:

Tenant

Primarily used by service providers to associate resources within a business entity to a customer.

Business Unit

Name of the business unit that owns the application. If there are multiple business units associated with a host, then the business unit shown in the table is a representative. A specific application has the same Business Unit on all hosts.

Note: The value will be N/A if the application is not associated with a business unit.

LOB

A line of business within a company, for example "Storage Management" or "Database Replication."

Project

A project to which you might want to further break down capacity chargeback.

Related concepts

[Application Volume Hourly Performance data model](#) on page 88

[Chargeback advanced data mart](#) on page 80

[Host Volume Hourly Performance data model](#) on page 89

[Internal Volume Capacity data model details](#) on page 57

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity advanced data mart](#) on page 66

[Switch Performance data model](#) on page 90

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Business Unit Dimension

All business units and applications that have ever been set in the enterprise environment are shown in this table.

The fields and field descriptions are as follows:

Business Unit

Business unit for the application. If there are multiple business units associated with a host, then the business unit shown in the table is a representative.

Business units are part of the business entity hierarchy. Use this to capture previous versions of OnCommand Insight, which used business unit annotations without the business entity hierarchy.

Application Name

Name of application in the business unit.

Capacity

You can identify how much physical storage is available, used, or unused and how much storage is projected to be available or used.

Capacity data elements

You can use any of the Capacity data elements in MB, GB, or TB.

Accessed Volume Capacity

Provisioned capacity of volumes that are accessed by hosts.

Capacity

The size of the storage or storage pool that is available for storing user data.

Orphaned Volume Capacity

The provisioned capacity of volumes that are masked but do not have active paths to them.

Protection Volume Capacity

the provisioned capacity of volumes that are used for replication.

Raw Capacity

The physical disk capacity of the storage or storage pool. Raw capacity is derived by the device manufacturer. This differs from usable capacity when technologies such as RAID-5 are used, where some of the raw capacity is used for protection purposes. Capacity that is presented before RAID.

Snapshot Used Capacity

The capacity at the internal volume and storage pool level that is used for Snapshot copies.

Snapshot Used Raw Capacity

The physical capacity at the internal volume and storage pool level that is used for Snapshot copies

Unconfigured Raw Capacity

The raw capacity that is unconfigured and might be available for storing additional data on the storage or storage pool.

Unused Capacity

The size of the storage or storage pool that is available for storing user data.

Unused Raw Capacity

The raw capacity that might be available for storing additional data on the storage or storage pool.

Unused Volume Capacity

The provisioned capacity in the volume that has not been accessed by hosts and is available for storing additional data.

Used Capacity

The capacity consumed in the device. The amount of usable capacity in use on the storage or storage pool. This includes the capacity consumed by internal volumes, volumes, Snapshot Reserve, and other storage efficiency technology overheads.

Used Raw Capacity

The amount of raw capacity in use on the storage or storage pool, including reserved capacity, Snapshot used capacity, Snapshot reserved capacity, and storage efficiency technology overhead.

Volume Capacity

The provisioned capacity of the volume.

Projected Capacity data elements

To obtain projected capacity statistics, OnCommand Insight uses a linear regression algorithm applied to the set of recently observed capacity.

You can use any of the Projected Capacity data elements in MB, GB, or TB.

Projected Capacity

An indication of how much capacity will be used at a future date, based on a linear regression algorithm that is applied to a set of recently observed capacity usage.

Projected Raw Capacity

The expected physical disk capacity of the storage or storage pool. Raw capacity is derived by the device manufacturer. This differs from usable capacity when technologies such as RAID-5 are used, where some of the raw capacity is used for protection purposes. Capacity that is presented before RAID.

Projected Used Capacity

The capacity expected to be consumed in the device. This includes the capacity consumed by internal volumes, volumes, Snapshot Reserve, and other storage efficiency technology overheads.

Projected Used Raw Capacity

The amount of raw capacity expected to be in use on the storage or storage pool, including reserved capacity, Snapshot used capacity, Snapshot reserved capacity, and storage efficiency technology overhead.

Related concepts

[Chargeback data model details](#) on page 79

[FS Utilization data model details](#) on page 54

[Internal Volume Capacity data model details](#) on page 57

[Storage and Storage Pool Capacity data model](#) on page 67

[Storage Efficiency data model details](#) on page 94

[VM Capacity data model details](#) on page 71

Capacity Types full description

OnCommand Insight describes capacity using capacity types. Capacity types are mutually exclusive. This means that each byte on the disk is counted only once using a single capacity type. The items in this table describe properties of the capacity types. For instance, “Used” will have a value of 1 for all capacity types that can be considered as used capacity.

The capacity types and descriptions are as follows:

Accessed

Capacity of volumes that are being accessed by hosts (supported only on FC paths).

Backend Volume Other

Capacity of backend volumes that is not used by virtual storages, for example, overhead attributed to different technologies.

Data Files

Capacity at the internal volume level that is used for files that are not LUNs. Used for NAS environments.

Data store Other

Capacity at the data store level that is not specified by OnCommand Insight, for example, overhead attributed to different technologies.

Data store Unused

Capacity at the data store level that is not used by any virtual disk.

Failed

Capacity of failed disks at the storage level.

Internal Volume Reserved

Capacity on the internal volume that is reserved on the storage pool, but is not used.

Masked

Capacity of volumes that are masked to hosts but not accessed by any hosts.

Protection

Capacity of volumes that are used as replicas of other volumes.

Reserved

Capacity at the storage pool level that is reserved for volumes.

Snapshot Unused

Capacity at the internal volume and storage pool level that reserved for Snapshot copies but is not used.

Snapshot Used

Capacity at the internal volume and storage pool level that is used for Snapshot copies.

Spare

Capacity of spare disks at the storage level.

Storage Pool Other

Capacity at the storage pool level that is not specified by OnCommand Insight, for example, overhead attributed to different technologies.

Storage Pool Unused

Capacity at the storage pool level that is not used.

Unused

Capacity at the storage level that is not used. Does not include failed or spare disks.

VM Used

Capacity of virtual disks that are used by virtual machines. (The host in this capacity type is the host running/owning the VM. The host group is always a group with a single host in it.)

Volume Unused

Capacity of volumes that are not masked and are not used as replicas.

Connected to Dimension

This table displays what device is connected to the switch port: storage array, host, switch, tape, or generic device.

The fields and field descriptions are as follows:

Connected to Name

Type of device the switch port is connected to.

Firmware

Firmware running on the port to which the port identified by the Name field is physically connected.

Manufacturer

Manufacturer of the device the switch port is connected to.

Model

Model of the device.

Type

Port identified by the name field.

Related concepts

[Port Capacity data model](#) on page 60

[Switch Performance data model](#) on page 90

Data Center

The Data Center element defines the physical location associated with the device.

Data Center data elements

Data Center (All)

If you select this data element, Data Center (All) information appears in the report indicating that you are at top of a hierarchy and that you can drill down into more detail and see specific data centers.

Data Center

Data center associated with this device.

Related concepts

[Chargeback data model details](#) on page 79

[FS Utilization data model details](#) on page 54

[Internal Volume Capacity data model details](#) on page 57

[Storage and Storage Pool Capacity data model](#) on page 67

[Storage Efficiency data model details](#) on page 94

[VM Capacity data model details](#) on page 71

Date Dimension

The Date dimension has an entry for every run of the ETL. Use the data elements to write different reports for different time frames. This is useful, for example, when you want to view different time points in the past.

Adding a Capacity Fact column gives you the time the ETL ran and the capacity at that time. The report also shows you a line for each time the ETL ran.

The Date dimension includes several date-related elements, including the following which warrant some explanation:

- **Day In <time period>:** For example, the Day in Week field indicates which day of the week the date is (Sunday is 1, Monday is 2, Tuesday is 3, and so on).
- **Is <time period> Representative:** The Is Day Representative flag is like the other representative flags, used to mark the first time the ETL runs on a given day, because the ETL can be scheduled to run multiple times in a single day (or may be run multiple times manually).
The representative flags mark the first date_dimension value for the given field. For example, the first date that the ETL was run for every month has an Is Month Representative value of 1 and all other dates for that month have an Is Month Representative value of 0. With this, you can easily pick out a single date per month, quarter, or week for generating trending reports and historical reports.
Usually the representative point is the first date in the given time span. For example, usually the Is Month Representative for February is Feb 1st. If the ETL does not run on the 1st, the point becomes the 2nd (or the first time the ETL is run in February).

Date (All)

If you select this data element, Date (All) information appears in the report indicating that you are at top of a hierarchy and that you can drill down into more detail and see data by different time periods.

Date Year

Month, day, and year in four-digit format.

Full Date

Month, day, year, and time of day, in alphanumeric format.

Day in Week

Day of the week. The first day of the week has the value of 1.

Day In Month

The day of the month. (The first day in the month is 1, the second day is 2, and the third day is 3, etc.).

Day In Quarter

The day of the quarter. (The first day in the quarter is 1, the second day is 2, and the third day is 3, etc.).

Day in Year

Day in the year. First day of the year has the value of 1.

Month

Month of the year. January has the value of 1.

Quarter

Quarter of the year. First quarter has the value of 1.

Year

Year in four-digit format.

Is Day Representative

Indicates whether this date value is the representative for the Day. This is useful if the ETL is scheduled to run more than once a day.

Note: "1" indicates the date is a representative for the value. "0" indicates it is not a representative date.

Is Week Representative

Indicates whether this date is the first date available for its week.

Is Month Representative

Indicates whether this date is the first date available for its month.

Is Quarter Representative

Indicates whether this date is the first date available for its quarter.

Is Latest Date

Indicates whether this date is the last date available in the table. Represents the data collected in the last run of the ETL.

Related concepts

[Application Volume Hourly Performance data model](#) on page 88

[Chargeback advanced data mart](#) on page 80

[FS advanced data mart](#) on page 55

[Host Volume Hourly Performance data model](#) on page 89

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Port Capacity advanced data mart](#) on page 63

[Qtree Capacity advanced data mart](#) on page 66

[Storage Efficiency advanced data mart](#) on page 96

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Fabric Dimension

Use the Fabric Dimension to identify the association of one or more FC switches and review their performance data. By using this dimension, you can retrieve fabric name and then look at WWN, status on whether fabric is VSAN, VSAN identifier, and status on whether the fabric has zone enforcement enabled.

The field and description is as follows:

Fabric Name

The fabric or VSAN on which the port is configured to operate.

Related concepts

[Switch Performance data model](#) on page 90

File System Dimension

This dimension table provides FS attributes (such as FS name, location, and type).

The fields and field descriptions are as follows:

Hardware ID

Host-specific identifier.

Location

Indicates where the data is kept and how it is accessed. For example: local, SAN, NAS.

Name

Mounting point or drive letter.

Type

File system type, for example, NTFS, FAT, UFS, UDF, ZFS, ISO 9660, or VXFS.

Host Dimension

The Host Dimension displays the host that is accessing the capacity. If multiple hosts are accessing the same volume (as in a cluster), one host is selected as the representative.

This dimension includes the following data elements:

Host Name

Name to identify the host.

IP

IP address assigned to the host device.

Model

Manufacturer's model number.

Manufacturer

Name of manufacturer who makes the device.

Host OS

Operating system running on the host, if available to OnCommand Insight.

Related concepts

[Chargeback advanced data mart](#) on page 80

[FS advanced data mart](#) on page 55

[Switch Performance data model](#) on page 90

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Host Group Dimension

This dimension provides data elements that represent a group of hosts. In this dimension all hosts are associated with capacity. This dimension enables you to access all hosts that are associated with capacity fact values, rather than just the representative host that is selected by the Data Warehouse.

This dimension includes the following data elements:

Cardinality

Number of hosts in the host group accessing the capacity.

Representative Host

Name of the application representing the host group.

Host Names

List of all hosts accessing the capacity.

Related concepts

[Chargeback advanced data mart](#) on page 80

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Internal Volume Dimension

Use the data elements in this table to identify the internal volumes on which the capacity resides. When array virtualization is in use, the internal volume dimension value reflects the internal volume from the front-end virtualizer.

This dimension includes the following data elements:

Internal Volume Name

Name of the internal volume.

Is Space Guaranteed

Vendor-specific indication of how the internal volume's space is allocated if thin provisioning is in use. For Data ONTAP, displays one of the following types:

- Volume: Data ONTAP pre-allocates space in the storage pool for the volume. The pre-allocated space cannot be allocated to any other volume in that storage pool.
- File: Data ONTAP pre-allocates the space in the volume so that any file in the volume with space reservation enabled can be completely rewritten, even if its blocks are pinned for a Snapshot copy.
- None: Data ONTAP reserves no extra space for the volume. Writes to the LUNs or files contained by that volume could fail if the containing storage pool does not have enough available space to accommodate those writes.

Is Thin Provisioned

Indicates whether thin provisioning is enabled.

Is Thin Provisioning Supported

Indicates whether the device supports thin provisioning.

Storage IP

IP associated with the storage.

Storage Pool Identifier

Identifier of the storage pool on which the internal volume resides.

Type

Type of internal volume, for example, FlexVol or FlexClone.

Related concepts

[Chargeback advanced data mart](#) on page 80

[Internal Volume Capacity advanced data mart](#) on page 59

[Qtree Capacity advanced data mart](#) on page 66

[VM Capacity advanced data mart](#) on page 73

Port Dimension

This table defines the WWN for a switch port.

The field and field description is as follows:

WWN

WWN (World Wide Name) to identify the port.

Related concepts

[Port Capacity data model](#) on page 60

[Switch Performance data model](#) on page 90

Qtree Dimension

This dimension provides information about the qtree itself, its security, qtree type, and status.

The fields and field descriptions are as follows:

Storage Identifier

Identifier of the storage array.

Type

Type of qtree. This column displays one of the following qtree types:

- Default: Qtree was created by the system as a logical construct on the internal volume.
- Explicit: Qtree was created by the user.

Related concepts

[Chargeback advanced data mart](#) on page 80

Service Level Dimension

This table provides data that defines service level, which is the agreement between a service consumer or application and the storage service provider.

The service level cost can be the cost of providing the service or the actual price to service the consumer, while the tier cost is the average cost per TB for the specific tier. Use the tier dimension to include storage data, but use the service level dimension to describe service, which could be a combination of tier, bandwidth, and more. For example, a customer might request that the level of storage service sustains 5,000 IOPS. However, because the customer uses only 100 IOPS, the service provider might decide to move this service to a lower tier that sustains only 500 IOPS.

Using the Chargeback data model, you can use the Service Level dimension to report on services provided to storage consumers together with storage tiers and their costs, typically used for chargeback purposes.

Using the VM Capacity data model, you can show how performance is maintaining at 1,500 IOPS and compare that to the agreed-upon service level agreement. This differs from tier, which is the physical aspect of the storage, for example, 15K RPM on the disks.

The fields and field descriptions of the Service Level Dimension table are as follows:

Sequence

Sequential number of the tier defining the order of the values for the level of service.

Cost

Unit of cost per gigabyte assigned to the service level.

Related concepts

[Chargeback advanced data mart](#) on page 80

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity advanced data mart](#) on page 66

[VM Capacity advanced data mart](#) on page 73

[Volume Capacity data model](#) on page 75

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Storage Dimension

This table describes the actual physical storage where the capacity resides. If you use the data elements from this table and data elements from the virtualized storage environment, the resulting information represents backend storage.

In the VM Capacity data model, when a data store is mapped to a virtualized storage, storage and tier dimensions are selected according to the backend array and tier.

IP

IP address assigned to the storage array.

Manufacturer

Name of the manufacturer of the storage array.

Microcode Version

Version of the firmware running on the storage array.

Model

Manufacturer's model number.

Serial Number

Serial number for the storage array.

Storage Name

Name of the storage array.

Storage Data Center

Location of data center.

Storage Family

Product classification of the storage array (for example, CLARiiON or Symmetrix).

Related concepts

[Chargeback advanced data mart](#) on page 80

[FS Utilization data model details](#) on page 54

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity advanced data mart](#) on page 66

[Storage and Storage Pool Capacity data model](#) on page 67

[Storage Efficiency advanced data mart](#) on page 96

[Switch Performance data model](#) on page 90

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Storage Pool Dimension

This table defines the capacity available in storage pools. When array virtualization is in use, the storage pool dimension value reflects the storage pool as seen from the front-end virtualizer.

This dimension includes the following data elements:

Storage Name

Name of the storage array associated with the storage pool.

Storage IP

IP address assigned to the storage array.

Storage Pool Name

Name of the storage pool, which is the physical layer of storage that consists of the disks within a RAID group that provide thin provisioning capabilities. An aggregation of storage from which space for volumes or shares can be created.

Redundancy

Level of mirroring defined for the device based on the storage technology, for example, RAID-DP, underlying the device. This is taken from the device itself.

Thin Provisioning Supported

Indicates whether the device supports thin provisioning.

Type

A vendor-specific description of the storage pool. One of the following values can appear:

- Aggregate: For NetApp storage systems.
- Backend Group: For array virtualization.
- Flash Pool aggregate: For a storage pool using Flash disks as cache.
- RAID group: A non-vendor specific type to indicate RAID technology.
- Reserved: For a storage pool that cannot be used for data.
- Thin Provisioning: For a storage pool using thin provisioning technology.

Uses Flash Pool aggregates

An indication that the selected storage pool is using SSD disks as cache (for example, Flash Pool aggregate technology). In this case, the SSD disks are used for cache and do not contribute to the usable size of the storage pool. A "1" indicates that Flash Pool aggregate technology is used; a "0" indicates that it is not.

Virtual

True indicates that the storage pool is virtual.

Related concepts

[Chargeback advanced data mart](#) on page 80

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity advanced data mart](#) on page 66

[VM Capacity advanced data mart](#) on page 73

Switch Dimension

This table defines the characteristics of the switch owning the port.

The fields and field descriptions are as follows:

Data Center

Location of Data Center, based on the switch's data center annotation in OnCommand Insight.

Firmware

Version of the firmware running on the switch.

IP

IP address of the switch.

Manufacturer

Name of the manufacturer who makes the switch.

Model

Manufacturer's model of the switch.

Serial Number

Serial number for the port.

Switch Level

Core, Host, Storage or User-defined (Annotation level assigned to the switch).

Switch Name

Switch name, as provided by the SAN manager at the switch console.

WWN

WWN (World Wide Name) to identify the switch.

Switch Port Dimension

This dimension provides information about the switch manufacturer and port speed, type, and WWN.

The fields and field descriptions are as follows:

Manufacturer

Name of the manufacturer who makes the switch.

Model

Manufacturer's model of the switch.

Switch Name

Switch name, as provided by the SAN manager at the switch console.

Port Name

Name of the port or, if no name is available, the WWN of the port.

Port WWN

WWN (World Wide Name) to identify the port

Port Speed

The speed capability of the port.

Port Connectivity Type

One of the following types of device to which the port is connected: Free, Generic, Host, ISL, Storage, Tape, or N/A.

Is ISL Port

Indication whether this port is a connection between two inter-switch links (ISLs) using the E-port.

Port Type

The port type, typically used to identify the protocol used to communicate with the other device, taken from the device itself. For example, E-Port, F-Port, G-Port, or TE-Port.

Port GBIC Type

The type of Gigabit Interface Converter (GBIC) used for the port.

Related concepts

[Port Capacity data model](#) on page 60

[Switch Performance data model](#) on page 90

Switch Performance

Use data elements in Switch Performance to view traffic distribution data across the different switches on the storage network.

The fields and field descriptions are as follows:

Rx Value (MB)

Amount of traffic received on this port.

Tx Value (MB)

Amount of traffic transmitted on this port.

Has Errors

Indication that there are errors in the port.

Has CRC Errors

Indication that there are CRC frame errors. CRC frame errors indicate bit errors somewhere in the data path and point to poor connections, bad cables, or links that are too long.

Sync Loss Count

Number of synchronization loss errors. If a Loss of Sync error occurs, the hardware cannot make sense of the traffic or lock onto it. All of the equipment might not be using the same data rate or the optics or physical connections might be of poor quality. The port must resync after each such error, which impacts system performance.

Signal Loss Count

Number of signal loss errors. If a Loss of Signal error occurs, there is no electrical connection and a physical problem exists.

Class 3 Discard Count

The count of FC Class 3 data transport discards.

Frame Too Short Count

The count of FC data transmission frames that are too short.

Frame Too Long Count

The count of FC data transmission frames that are too long.

BB Credit Error Count

FC uses buffer-to-buffer credits to control transmission flow. The credit value is decreased by one when a frame is sent and increased by one when a response is received. As the available credits for a given port approach zero, the error warns that the port will stop receiving transmissions when zero is reached and will not resume until the BB credits can be replenished.

Rx Usage (%)

Percentage of available bandwidth used for receiving transmissions.

Tx Usage (%)

Percentage of available bandwidth used for sending transmissions.

Rx Max Usage (%)

The percentage of maximum bandwidth used for receiving transmissions.

Tx Max Usage (%)

The percentage of maximum bandwidth used for sending transmissions.

FC Balance Index

A calculated standard deviation that measures how balanced the load is across connected switch ports.

Weighted FC Balance Index

A calculated standard deviation, weighted on port traffic, that measures how balanced the load is across connected switch ports. This value removes the outliers.

Ports Speed (Gbit/s)

The sum of the speeds of the switch ports connected to the device. For example, if three ports with a speed of 8 Gbit/s are connected to the device, the ports speed shows as 24 (3*8).

Port Count

Number of ports to which this device is connected. Each line shows performance for one device (storage system, host, and so on).

Tape Dimension

Use the Tape Dimension to identify the tape device for which switch port performance data is collected.

Related concepts

[Port Dimension](#) on page 145

[Switch Performance data model](#) on page 90

Tier Dimension

The Tier Dimension table is populated if tier annotations were associated with capacity in the OnCommand Insight Client either manually or using rules.

For details about tier annotations, see the *OnCommand Insight Planning User Guide*.

In the Tier Dimension table, when a data store is mapped to a virtualized storage, storage and tier dimensions are selected according to the backend array and tier.

The N/A value is for all capacity that has not yet been associated with a tier.

(All)

If you select this data element, Tier (All) information appears in the report indicating that you are at top of a hierarchy and that you can drill down into more detail and see Tier > Tier Name > Tier Cost.

Tier Cost

Unit of cost per gigabyte assigned to the tier level.

Tier Sequence

Sequential number of the tier defining the order of the values for the tier.

Tier Name

Name of the tier, for example, Gold, Silver, or Bronze.

Related concepts

[Chargeback advanced data mart](#) on page 80

[FS Utilization data model details](#) on page 54

[Internal Volume Capacity advanced data mart](#) on page 59

[Internal Volume Daily Performance and Internal Volume Hourly Performance data models](#) on page 89

[Qtree Capacity advanced data mart](#) on page 66

[Storage and Storage Pool Capacity data model](#) on page 67

[Storage Efficiency data model details](#) on page 94

[VM Capacity advanced data mart](#) on page 73

[Volume Daily Performance and Volume Hourly Performance data models](#) on page 91

Time Dimension

This dimension table displays hours, minutes, seconds, and microseconds. Each fact is associated with the time when data was sampled.

Time period fact tables

Use the elements in these tables to obtain historical data in your reports.

These tables are used for all data models.

The information in these time period tables reflects a given time period only if the ETL runs in that period. For example, if OnCommand Insight stops running for the entire month of February, there won't be a date point for the February table.

The fact tables and descriptions are as follows:

Current <data model> Fact

Data model fact data for only the last ETL run (the most recent date point).

Weekly <data model> Fact

Data model fact data for a single date point each week. Includes historical data. Contains the values from the specific data model Fact table, filtered to include only values associated with dates where the Is Week Representative equals 1.

Monthly <data model> Fact

Data model fact data for a single date point each month. Includes historical data. Contains the values from the specific data model Fact table, filtered to include only values associated with dates where the Is Month Representative equals 1.

Quarterly <data model> Fact

Specific data model fact data for a single date point each quarter. Includes historical data.

Q1 is Jan 1 - Mar 31

Q2 is Apr 1 - Jun 30

Q3 is Jul 1 - Sep 31

Q3 is Jul 1 - Sep 31

Related concepts

[Internal Volume Capacity advanced data mart](#) on page 59

[Port Capacity data model](#) on page 60

[Storage and Storage Pool Capacity data model](#) on page 67

[VM Capacity advanced data mart](#) on page 73

VM Dimension

The VM dimension table provides a current (live) snapshot of inventory.

If capacity is associated with a VM, the table shows information about that VM. However, sometimes capacity data will only be available if the VM is in the current OnCommand Insight environment (not historical).

A report using this table will always reflect current data only. If there is no historical data available, the row is left blank.

VM Name

Name of the virtual machine.

VM OS

Name of guest host.

DNS Name

DNS lookup requests using in IP identify validation.

IPs

IP address assigned to the host device.

VM Power State

Power state of the virtual machine: Powered On, Powered Off, Suspended.

VM Power State Change Date

Date and time when the VM Power State was changed to the current value.

Powered Off Since

Date and time when a system was last active. If the VM is running, this field is blank.

Guest OS State

Combined power state of the virtual hardware of the virtual machine and of the guest OS: Running, Powered Off, Powered Off, Not Running, Standby, Shutting Down, Unknown, Suspended.

Processors

Number of processors used by the virtual machine.

Memory

Memory used by the VM, in MB.

Virtual Center IP

IP address of the virtual machine.

Virtual Disks

A file that represents a disk of a VM. The VM presents the virtual disk to its operating system as a hard drive. In a VMware implementation, this could be the VMDK file name (relative to the data store).

VM Count

The number of VMs in the selected group.

Volume Dimension

This dimension provides information about the volume, whether thin provisioning is enabled, whether it is a meta volume, and what type of volume it is.

The fields and field descriptions are as follows:

Is Meta

Indicates whether this volume is a meta volume. A meta volume is a volume that serves as a head in a group of volume members.

Is Thin Provisioned

Indicates whether thin provisioning is enabled.

Is Virtual

Indicates whether this item is virtual storage.

Label

Label or description of the volume.

Storage Identifier

Physical storage on which the volume resides.

Type

Type of volume, for example, SFS or BCV.

Volume name

Name of volume.

Data mart dimension tables for SQL queries

OnCommand Insight includes many dimension tables that you can use in SQL queries.

Application dimension in data marts

This is a representative application from the group of applications using the capacity.

Application group dimension in data marts

This is a representative application from the group of applications using the capacity. This defines the application group (set of applications) that is using the capacity.

Applications and application groups behave in a similar way to hosts and host groups. Therefore, a similar solution is applied to applications.

Related concepts

[Hosts and host group dimensions](#) on page 156

Business Entity dimension in data marts

The `business_entity_dimension` is used to identify the tenant, line of business (LOB), business unit, and project.

Connected Device dimension in data marts

This table (`connected_device_dimension`) displays what device is connected to the switch port: storage array, host, switch, tape or generic device.

Date dimension in data marts

Each time the capacity job runs, a row is inserted into the date dimension table. The row represents the time of the last build time.

Use the data elements to write different reports for different time frames. This is useful, for example, when you want to view different time points in the past.

Adding a Capacity Fact data element gives you the time the ETL ran and the capacity at that time. The data also includes a line for each time the ETL ran.

The Date Dimension includes several date-related elements, including the following which warrant some explanation:

- **Day In <time period>**: For example, the Day in Week field indicates which day of the week the date is (Sunday is 1, Monday is 2, Tuesday is 3, etc.).
- **Rep<time_period>**: The Representative flag is like the other representative flags, used to mark the first time the ETL runs on a given day, since the ETL can be scheduled to run multiple times in a single day (or may be run multiple times manually).

The representative flags mark the first `date_dimension` value for the given field. For example, the first date that the ETL was run for every month has a Month Representative value of 1 and all other dates for that month have a Month Representative value of 0. With this, you can easily pick out a single date per month, quarter, or week for writing trending reports and historical reports.

Usually the representative point is the first date in for the given time span. For example, usually the Month Representative for February is Feb 1st. If the ETL does not run on the 1st, the point becomes the 2nd (or the first time the ETL is run in February).

The columns and column descriptions are as follows:

fullDate

Month, day, year, and time of day, in alphanumeric format.

dayInWeek

Day of the week. The first day of the week has the value of 1.

dayInMonth

Day in the month. The first day of the month has the value of 1.

dayInQuarter

Day in the quarter. The first day of the quarter has the value of 1.

dayInYear

Day in the year. First day of the year has the value of 1.

Month

Month of the year. January has the value of 1.

Quarter

Quarter of the year. First quarter has the value of 1.

Year

Year in four-digit format.

RepDay

Indicates whether this date value is the representative for the Day. This is useful if the ETL is scheduled to run more than once a day.

RepWeek

Indicates whether this date is the first date available for its week.

RepMonth

Indicates whether this date is the first date available for its month.

RepQuarter

Indicates whether this date is the first date available for its quarter.

Latest

Indicates whether this date is the last date available in the table. Represents the data collected in the last run of the ETL.

Data store dimension

This dimension provides capacity and performance characteristics of the ESX, storage and fabric.

Disk dimension in data marts

Use the disk_dimension to obtain information about disks for which performance is being captured. By using this table, you can retrieve disk name, speed, location, role (for example, data, parity, or

spare), vendor, model, type (for example, ATA, FATA, SATA, FC, or SSD), disk group, status (for example, normal or failed), and serial number.

Fabric dimension in data marts

Use the Fabric dimension (fabric_dimension) to identify the association of one or more FC switches and review their performance data. By using this dimension, you can retrieve fabric name and then look at WWN, status on whether fabric is VSAN, VSAN identifier, and status on whether the fabric has zone enforcement enabled.

File System Dimension in data marts

File system data is kept in the dwh_fs_util.file_system_dimension dimension and has several columns.

Those columns are as follows:

- **Name:** Mounting point or drive letter. Part of the natural key.
- **Location:** Indicates where the data is kept and how it is accessed. For example: local, SAN, or NAS. If the location is unknown, it is set to null.
- **HardwareId:** Host specific identifier. If unknown, it will be set to null.
- **Type:** File system type. For example: NTFS, FAT, UFS, UDF, ZFS, ISO 9660, or VXFS. If unknown, this is set to null. The measurements (size and used capacity) are kept in the dwh_fs_util.fs_util_fact table.

Files systems that do not correlate to a host (for example, when the host IP of the file system does not correlate to any identifier or IP of a host) are not loaded to DWH.

The Inventory database does not include a file system table.

Host dimension in data marts

The Host dimension represents the host accessing the capacity. HostTK refers to the representative host; however, this might not be the only host that is accessing the capacity.

Hosts and host group dimensions

This is a representative host from the group of hosts using the capacity. The Host dimension represents the group of hosts sharing some capacity.

This might lead to confusing results when grouping by host. Consider the following issues:

- The host group dimension joined with host group bridge defines which hosts belong to which groups.
- A host may belong to multiple groups.
- The groups are determined by looking at hosts that are accessing volumes.

The fact table includes the hostGroupTk dimension that refers to the host-group accessing the capacity. hostTK refers to the representative host; however, this is not the only host that is accessing the capacity.

OnCommand Insight chooses a representative from every group, and counts the capacity only for the representative. The application and business units are selected according to the ordering of their name. In this way, instead of associating a capacity with the host that accesses it, OnCommand Insight associates the capacity with the group of hosts that accesses it. Therefore, a specific capacity appears only once in the fact table at any given time point. The representative host is chosen for every group of hosts, using an alphabetical order with the host names.

Internal volume dimension in data marts

Use the internal_volume dimension to identify storage pool, type of internal volume, and whether the internal volume uses thin provisioning technologies.

Related concepts

[Capacity type list](#) on page 116

Port dimension in data marts

Use the Port dimension (port_dimension) to identify the switch port with which performance is associated. By using this dimension, you can retrieve port name, port WWN, port Gigabit Interface Converter (GBIC) type, port type (for example, F-Port or E-Port), port speed, and indication of whether switch port has generated WWN (for NPV switches).

Qtree dimension in data marts

Use the qtree dimension to identify qtree subdivision of OnCommand Insight internal volumes. You might want to look at the qtree details to view the security state, qtree status, or qtree type.

For descriptions of columns in this data mart, see the [Qtree Capacity data model](#) on page 64.

Service Level dimension in data marts

Use the service_level_dimension to report on services provided to storage consumers together with storage tiers and their costs, typically used for chargeback purposes. Use the tier dimension to include storage data, but use the service level dimension to describe service, which could be a combination of tier, bandwidth, and more.

Storage dimension in data marts

Use the storage dimension to identify the physical storage from which the capacity is taken. By using this table, you can retrieve storage array name, IP address, model, manufacturer, serial number, family, and firmware version running on the storage array.

Storage pool dimension in data marts

Use the storage_pool dimension to obtain information about storage pools. Only thin provisioning storage pools have capacity associated with them. Other storage pools have their capacity accounted for under their volumes and under the storage level.

By using this table, you can retrieve storage pool name, IP of the storage where the storage pool resides, type of storage pool (for example, RAID or aggregate), and whether thin provisioning technologies are used. A virtual column also indicates whether the resource is a frontend storage pool in array virtualization.

The capacity types applicable for storage pools are:

- Snapshot used: Capacity of the storage pool that is used for Snapshot copies.
- Snapshot unused: Capacity of the storage pool that is reserved for Snapshot copies, but unused.
- Unused: Available capacity on a thin provisioned storage pool.
- Reserved: Capacity on the storage pool that is reserved (such as in CLARiiON virtual provisioning).

- Other: Used capacity of the storage pool that is not a Snapshot copy and is not counted under the physical column in any of the objects (internal volumes or volumes) that reside on the storage pool.

Switch dimension in data marts

The switch_dimension table defines the characteristics of the switch owning the port. Use this dimension to retrieve switch firmware, IP, manufacturer, WWN, and switch level (core, host, storage, or user-defined)

Tape dimension in data marts

Use the Tape dimension table (tape_dimension) to tape name, IP, manufacturer, and serial number.

Tier dimension in data marts

The tier dimension defines the tier of the measured capacity-supports tier annotation. This refers to the tier of the LUN, internal volume, storage pool, or storage array. The cost column can be used to retrieve the cost associated with the tier.

Time dimension in data marts

Use the time_dimension to identify the hour, minute, second, or microsecond, if available, of an event. This dimension is used to record the time of DWH hourly performance summaries.

VM dimension in data marts

Use this dimension to define the Virtual Machine using the capacity. The VM dimension table provides a current (live) snapshot of inventory. If capacity is associated with a VM, the table shows information about that VM.

Volume dimension in data marts

Use the volume_dimension to identify the volume's name and type, whether the volume is a virtual volume, and if the volume uses thin provisioning technologies.

Copyright information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- accessing
 - Data Warehouse portal [12](#)
 - database schemas [100](#)
 - MySQL [99](#)
 - Reporting Portal [21](#)
- advanced data marts
 - Chargeback [80](#)
 - FS Utilization [55](#)
 - Internal Volume Capacity [59](#)
 - Port Capacity [63](#)
 - Qtree Capacity [66](#)
 - Storage Efficiency [96](#)
 - VM Capacity [73](#)
- annotations
 - business entities [137](#)
- Application dimension
 - definitions [136](#)
- Application Dimension
 - definitions [135](#)
- application groups [136](#)
- authoring tools
 - for reports [42](#)

B

- browsers
 - IE required for reports [7](#)
- business entities [137](#)
- Business Unit Dimension
 - definitions [137](#)

C

- capacity
 - concepts [110](#)
 - definitions [138](#)
 - projected [69](#)
 - report types [38](#)
 - reports [38](#)
 - virtual machines report [40](#)
- capacity data mart
 - tasks you can perform [98](#)
- Capacity data mart
 - fact tables [112](#)
 - overview [109](#)
- Capacity data mart dimensions
 - in SQL queries [115](#)
- capacity data model overview [53](#)
- capacity efficiency
 - views [119](#)
- Capacity Efficiency data mart
 - overview [117](#)
 - sample SQL queries [120](#)
- capacity information
 - types of [103](#)
- capacity types

- for SQL queries [116](#)
- chargeback
 - reports [38](#)
- Chargeback data model
 - advanced data marts [80](#)
 - data [79](#)
 - sample reports [78](#)
- Cognos
 - documentation locations [22](#)
- comments
 - how to send feedback about documentation [161](#)
- connected devices
 - in SQL queries [154](#)
- Connected to Dimension
 - definitions [141](#)
- customizing
 - report output format [24](#)

D

- dashboards
 - adding from User Preferences [27](#)
 - detailed data links [26](#)
 - viewing [28](#)
- data capacity fact maintenance
 - processes [15](#)
- data centers
 - data element definitions [141](#)
- data marts
 - Capacity [109](#)
 - Capacity Efficiency [117](#)
 - File System Utilization [120](#)
 - Inventory [100](#)
 - Performance [121, 122, 125](#)
 - Port Capacity [129](#)
 - Storage and Storage Pool Capacity [98](#)
 - Volume Capacity [99, 133, 134](#)
- data model
 - selecting [48](#)
- data models
 - Chargeback [77](#)
 - descriptions [11, 51, 53](#)
 - fact and dimension tables [53](#)
 - File System Utilization [54](#)
 - glossary [135](#)
 - Internal Volume Capacity [55](#)
 - introduction [50](#)
 - Inventory [81](#)
 - Performance [86](#)
 - Port Capacity [60](#)
 - Qtree Capacity [64](#)
 - Storage and Storage Pool Capacity [67](#)
 - Storage Efficiency [92](#)
 - Switch Performance [90](#)
 - using multiple [53](#)
 - VM Capacity [70](#)
 - Volume Capacity [75, 76](#)
- Data Warehouse
 - accessing data as MySQL user [99](#)

- data marts [97](#)
- Extract, Transform, Load (ETL)
 - process for Data Warehouse [14](#)
- Extract, Transform, Load (ETL) process [14](#)
- login and password [12](#)
- software components [13](#)
- Data Warehouse portal
 - accessing [12](#)
 - login and password [12](#)
- database schemas
 - data marts [100](#)
- dates
 - data element definitions [141](#)
- designing reports
 - tips for [44](#)
- devices
 - report [34, 39](#)
- dimension tables
 - capacity efficiency [119](#)
- dimensions
 - application group [136](#)
 - date [141](#)
 - fabric [143](#)
 - host [144](#)
 - host group [144](#)
 - internal volume [144](#)
 - list of [135](#)
 - port [145](#)
 - qtree [145](#)
 - service level [146](#)
 - storage pool [147](#)
 - switch [148](#)
 - switch port [149](#)
 - switch port performance [149](#)
 - tape [151](#)
 - tier [151](#)
 - VM [152](#)
 - volume [153](#)
- dimensions in data marts
 - connected device [154](#)
 - disk [155](#)
 - fabric [156](#)
 - list of [154](#)
 - port [157](#)
 - switch [158](#)
 - tape [158](#)
- disks
 - in SQL queries [155](#)
- documentation
 - how to receive automatic notification of changes to [161](#)
 - how to send feedback about [161](#)
 - IBM Cognos [22](#)
 - list of [18](#)
 - Reporting Connection Help [17](#)

E

- email notification
 - reports [25](#)
- emailing reports
 - setting up email notification [25](#)
- environment

- hosts [31](#)
 - storage devices report [34, 39](#)
- exporting report definitions [50](#)
- extracting data
 - from OnCommand Insight into Data Warehouse [14](#)
- extracting data from OnCommand Insight into Data Warehouse
 - data retention [16](#)
 - frequency [14](#)

F

- Fabric Dimension [143](#)
- fabrics
 - in SQL queries [156](#)
- fact and dimension tables
 - data models [53](#)
- feedback
 - how to send comments about documentation [161](#)
- File System Utilization data mart
 - overview [120](#)
 - sample SQL queries [121](#)
- File System Utilization data model
 - advanced data marts [55](#)
 - data [54](#)
- filtering
 - data for reports [23](#)

G

- generating
 - pre-defined reports [23](#)
- getting started
 - accessing the Data Warehouse portal [12](#)
- glossary [135](#)

H

- Help
 - accessing in Reporting Connection [17](#)
- historical data retention
 - in Data Warehouse [15](#)
- host groups [144](#)
- hosts
 - accessing capacity [144](#)
 - HBAs [31](#)
 - summary [32](#)

I

- importing
 - report definitions in Report Studio [50](#)
- information
 - how to send feedback about improving documentation [161](#)
- Insight
 - license types [8](#)
- Internal Volume Capacity data model
 - advanced data marts [59](#)
 - data [57](#)
 - sample reports [56](#)
- internal volumes [144](#)

- inventory
 - information categories [30](#)
- Inventory data mart
 - accessing the database schema [100](#)
 - overview [100](#)
 - sample SQL queries [107](#)
- Inventory data model
 - sample reports [82](#)
 - submarts [84](#)

L

- licenses
 - data in reports depends on [8](#)
 - Reporting roles [9](#)
 - types [8](#)

M

- mapping
 - volumes [33](#)
- models, data
 - descriptions [11, 51](#)

O

- OnCommand Insight
 - extracting data from [14](#)
 - inventory reports [30](#)
 - reporting [7](#)

P

- paths
 - virtual machine [36](#)
- Performance data mart for SQL queries
 - fact tables [125](#)
 - key data [122](#)
 - overview [121](#)
 - sample SQL queries [127](#)
- Performance data model
 - Application Volume Hourly [88](#)
- Performance data models
 - sample reports [87](#)
- pivoting reports [47](#)
- Plan
 - reporting user roles [9](#)
 - reports [38](#)
- Port Capacity data mart
 - overview [129](#)
 - sample SQL queries [130](#)
- Port Capacity data model
 - advanced data marts [63](#)
 - data [62](#)
 - sample reports [61](#)
- ports
 - fabric summary [30](#)
 - in SQL queries [157](#)
- processes
 - consolidation [13](#)
 - data capacity fact maintenance [15](#)
 - ETL [13](#)

- Jobs [13](#)
- requests to DWH [13](#)
- product documentation
 - list of [18](#)
- projected capacity in reports [69](#)

Q

- Qtree Capacity data model
 - advanced data marts [66](#)
 - data [65](#)
 - sample reports [64](#)
- qtrees [145](#)
- Query Studio
 - accessing [45](#)
 - creating custom reports [42](#)

R

- Reporting
 - user roles [9](#)
- Reporting Portal
 - accessing [21](#)
- reports
 - accessing [20](#)
 - accessing capacity [38](#)
 - authoring tools [42](#)
 - capacity [38](#)
 - common questions answered [7](#)
 - creating [43](#)
 - creating custom ad hoc [49](#)
 - creation process [43](#)
 - custom ad hoc [42, 48](#)
 - custom ad hoc, introduction to creating [48](#)
 - design tips [44](#)
 - exporting definitions [50](#)
 - filtering data before generating [23](#)
 - finding predefined [8](#)
 - formatting and scheduling [24](#)
 - glossary [135](#)
 - IBM documentation [22](#)
 - importing definitions [50](#)
 - inventory [30](#)
 - licenses control data [8](#)
 - locating pre-defined [21](#)
 - locations of predefined [8](#)
 - pivoting [47](#)
 - pre-defined [7](#)
 - saving [47](#)
 - setting up email notification [25](#)
 - types [20](#)
- roles
 - Reporting [9](#)

S

- sample reports
 - Chargeback data model [78](#)
 - Internal Volume Capacity data model [56](#)
 - Inventory data model [82](#)
 - Performance data models [87](#)
 - Port Capacity data model [61](#)

- Qtree Capacity data model [64](#)
 - Storage and Storage Pool Capacity data model [68](#)
 - Storage Efficiency data model [93](#)
 - VM Capacity data model [71](#)
 - Volume Capacity data model [75](#)
 - sample SQL queries
 - Capacity Efficiency data mart [120](#)
 - File System Utilization data mart [121](#)
 - Inventory data mart [107](#)
 - Performance data mart [127](#)
 - Port Capacity data mart [130](#)
 - Storage and Storage Pool Capacity data mart [132](#)
 - selecting output format
 - reports [7](#), [8](#), [20–25](#), [30](#), [38](#), [42–44](#), [47–50](#), [135](#)
 - service levels [146](#)
 - simple data marts
 - Chargeback [79](#)
 - FS Utilization [54](#)
 - Internal Volume Capacity [57](#)
 - Port Capacity [62](#)
 - Qtree Capacity [65](#)
 - Storage Efficiency [95](#), [96](#)
 - VM Capacity [72](#), [74](#)
 - SQL queries
 - capacity data mart dimensions [115](#)
 - capacity types [116](#)
 - dimension tables for [135](#), [154](#)
 - user access [99](#)
 - virtual volumes [116](#)
 - storage
 - data element definitions [146](#)
 - host utilization [32](#)
 - summary report [34](#), [39](#)
 - Storage and Storage Pool Capacity data mart
 - considerations [131](#)
 - overview [98](#)
 - questions you can answer using [131](#)
 - sample SQL queries [132](#)
 - Storage and Storage Pool Capacity data model
 - All Capacity submart [69](#)
 - capacity concepts [110](#)
 - Physical Capacity submart [70](#)
 - projected capacity [69](#)
 - sample reports [68](#)
 - submarts [68](#)
 - Storage capacity
 - producers and consumers [117](#)
 - Storage Capacity data model
 - capacity concepts [110](#)
 - Storage Efficiency data model
 - advanced data mart [96](#)
 - data [95](#), [96](#)
 - sample reports [93](#)
 - storage pools [147](#)
 - submarts [84](#)
 - suggestions
 - how to send feedback about documentation [161](#)
 - Switch Performance data model
 - fact tables [91](#)
 - switch port performance [149](#)
 - switch ports [149](#)
 - switches
 - data for custom reports [90](#)
 - fabric summary [30](#)
 - in SQL queries [158](#)
- ## T
- tapes
 - in SQL queries [158](#)
 - tenant annotations
 - line of business
 - project annotations [137](#)
 - terminology
 - Application Dimension [135](#), [136](#)
 - Business Unit Dimension [137](#)
 - capacity definitions [138](#)
 - capacity types [139](#)
 - Connected to Dimension [141](#)
 - data centers [141](#)
 - dates [141](#)
 - fabric [143](#)
 - Fabric Dimension [143](#)
 - port [145](#)
 - qtree [145](#)
 - storage [146](#)
 - switch [148](#)
 - time [152](#)
 - VM [152](#)
 - volume [153](#)
 - tiers [151](#)
 - time
 - data element definitions [152](#)
 - tips
 - for designing reports [44](#)
 - twitter
 - how to receive automatic notification of documentation changes [161](#)
- ## U
- users
 - Reporting roles [9](#)
- ## V
- videos
 - report formatting and scheduling [24](#)
 - virtual machines
 - capacity summary report [40](#)
 - dimension table [152](#)
 - path [36](#)
 - virtual volumens
 - in SQL queries [116](#)
 - VM Capacity data model
 - advanced data marts [73](#)
 - data [72](#), [74](#)
 - sample reports [71](#)
 - Volume Capacity data mart
 - overview [99](#), [133](#)
 - questions answered [133](#)
 - sample SQL query [134](#)
 - Volume Capacity data model
 - details [76](#)
 - overview [75](#)

- sample reports [75](#)
- Volume Daily Performance data model [91](#)
- Volume Hourly Performance data model [91](#)
- volumes
 - unmasked [33](#)

- IE required for reports [7](#)
- Workspace Advanced
 - creating custom reports [48, 49](#)
- Workspace Advanced tool
 - introduction to creating custom reports using [48](#)

W

- web browsers