

ViPR SRM Advanced Reporting Student Guide

Education Services
January 2015

ViPR SRM Advanced Reporting



Copyright © 2015 EMC Corporation. All Rights Reserved.

Support Contact: Education Services

Welcome to EMC ViPR SRM Advanced Reporting.

Copyright © 1996, 2000, 2001, 2002, 2003, 2004, 2005, 2006, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015 EMC Corporation. All Rights Reserved. EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC2, EMC, Data Domain, RSA, EMC Centera, EMC ControlCenter, EMC LifeLine, EMC OnCourse, EMC Proven, EMC Snap, EMC SourceOne, EMC Storage Administrator, Acartus, Access Logix, AdvantEdge, AlphaStor, ApplicationXtender, ArchiveXtender, Atmos, Authentica, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, Avamar, Captiva, Catalog Solution, C-Clip, Celerra, Celerra Replicator, Centera, CenterStage, CentraStar, ClaimPack, ClaimsEditor, CLARiiON, ClientPak, Codebook Correlation Technology, Common Information Model, Configuration Intelligence, Configuresoft, Connectrix, CopyCross, CopyPoint, Dantz, DatabaseXtender, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Document Sciences, Documentum, elnput, E-Lab, EmailXaminer, EmailXtender, Enginuity, eRoom, Event Explorer, FarPoint, FirstPass, FLARE, FormWare, Geosynchrony, Global File Virtualization, Graphic Visualization, Greenplum, HighRoad, HomeBase, InfoMover, Infoscapes, Infra, InputAccel, InputAccel Express, Invista, Ionix, ISIS, Max Retriever, MediaStor, MirrorView, Navisphere, NetWorker, nLayers, OnAlert, OpenScale, PixTools, Powerlink, PowerPath, PowerSnap, QuickScan, Rainfinity, RepliCare, RepliStor, ResourcePak, Retrospect, RSA, the RSA logo, SafeLine, SAN Advisor, SAN Copy, SAN Manager, Smarts, SnapImage, SnapSure, SnapView, SRDF, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix, Symmetrix DMX, Symmetrix VMAX, TimeFinder, UltraFlex, UltraPoint, UltraScale, Unisphere, VMAX, Vblock, Viewlets, Virtual Matrix, Virtual Matrix Architecture, Virtual Provisioning, VisualSAN, VisualSRM, Voyence, VPLEX, VSAM-Assist, WebXtender, xPression, xPresso, YottaYotta, the EMC logo, and where information lives, are registered trademarks or trademarks of EMC Corporation in the United States and other countries.

All other trademarks used herein are the property of their respective owners.

© Copyright 2015 EMC Corporation. All rights reserved. Published in the USA.

Revision Date: January 2015

Revision Number: MR-1CN-SRMREP - 3.5SP1.1.1

Course Overview

Description

This course provides training for modifying reports and creating standard and enhanced reports within the EMC ViPR SRM solution. ViPR SRM provides enterprise level performance and SLA management information for storage environments. After completing this course, you will understand the ViPR SRM's reporting capabilities, architecture, and be able to modify existing reports or create unique reports within your storage management realm.

Audience

ViPR SRM Report administrators.

Objectives

Upon completion of this course, you should be able to:

- Describe ViPR SRM's report structure design and data manipulation.
- Use the tools and features of ViPR SRM in report design and creation.
- Create various unique reports based upon selected criteria employing the many report types provided in ViPR SRM, including, dashboard, inventory, advanced graph reports (time-ratio, heat map, Geo-Map), mixed-reports, etc.
- Demonstrate recognition of the advantages and features of the different report types by selecting appropriate report types to satisfy specified conditions.
- Demonstrate an understanding of the uses for formulas in report design and creation.
- Create enhanced reports and manage the content of the report tree.



This course covers altering, creating, and editing reports using ViPR SRM. It also includes coverage of all available report types and features like formulas used to transform the default data set.

Agenda

	Modules/Lessons	Labs
Day 1	- Module 0 – Introduction and Agenda - Module 1 – ViPR SRM Overview - ViPR SRM Background - Overview of portal tools and features - Examples demonstrating the use of reporting - Overview of Data Enrichment - Module 2 – Basics of Reporting - Creating reports and report design - Base organization of data - Browsing the data set	- Lab Demo – Group Management - Lab 1: Discovering Available Datasets
Day 2	- Module 3 – Reporting Key Concepts - Report Definition Key Concepts - Expansion - Module 4 – Base Report Types - Base Report Types: Tables	- Lab 2: Organizing a Report Structure - Lab 3: Expansion in Reporting - Lab 4: Using Table Report Types - Lab 5: Using Graph Report Types



Copyright © 2015 EMC Corporation. All Rights Reserved.

Course Overview

3

We begin our Reporting journey with a basic understanding of the portal tools and features. In some instances, much can be done from this point, without having to alter a report in report edition mode. A review of the Data Enrichment concept will be covered as well as how to use simple tools that will allow you to dig into the data set and discover its organization and richness.

The second day we will cover the report edition mode itself. You will learn how important the report structure is and also define the metric context and how the data flows in reports. You will also start to work on creating table reports by learning how they are built.

The third day will cover graph and mixed reports. An intensive tour of the Formulas feature will then be presented. Hands on labs will help you understand the infinite possibilities in using formulas.

Agenda

	Modules/Lessons	Labs
Day 3	- Module 4 (continued) - Base Report Types: Graph - Base Report Types: Mixed Reports - Module 5 - Formulas	- Lab 6: Using Mixed Report Types - Lab 7: Completing Reports - Lab 8: Integration of Formulas in Reports
Day 4	- Module 5 (continued) - Module 6 – More Report Types - Aggregated Report Types - TopN Report Types	- Lab 9: Practicing Aggregated Report Types - Lab 10: Using TopN Report Types
Day 5	- Module 6 (continued) - Advanced Report Types - Miscellaneous Report Types - Publishing Your Reports	- Lab 11: Advanced Report Types - Lab 12: ReportPack Creation



The fourth day will allow you to spend even more time with the application of the formula concepts as you start exploring some advanced report types. A strong focus on labs will allow you to take time to digest the possibilities in using such advanced report types.

The final day will allow you to wrap up the advanced report types topic and practice labs in order to see use cases of such report types. Finally, you will be able to wrap up your report work in an a concise package and publish in the public section of the reporting portal.

Module: ViPR SRM Overview

Upon completion of this module, you should be able to:

- Describe the function of ViPR SRM in the creation of the Software Defined Data Center (SDDC)
- Explain how reporting is an integral function of managing the SDDC
- Describe the tools and features available through the portal to run, modify, and create reports useful in managing the SDDC
- Discuss how data enrichment can be used to enhance collected data
- Define uses for and examples of enhanced reporting



This module focuses on a brief introduction to the Software Defined Data Center (SDDC) made possible with ViPR and how ViPR SRM can be used to monitor and report on the condition of the infrastructure that is used to implement the SDDC. Included is an introduction to the tools and features of ViPR SRM, including the use of data enrichment, as well as a discussion of enhanced reporting usage for managing the physical and virtual elements of the SDDC.

Module: ViPR SRM Overview

Lesson: ViPR SRM Background

This lesson covers the following topics:

- Virtualization and the software defined datacenter
- ViPR SRM essentials



Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

2

This lesson covers an overview of the software defined data center where virtualization is a core component. A focus is placed on software defined storage and the value ViPR brings to a virtual storage environment making heterogenous arrays work together in a simple, extensible, and open platform.

Infrastructure Management Challenges

DATA GROWTH

Unconstrained growth, stranded investments driving up fixed costs

OVERHEAD

Siloed management for physical, virtual, and cloud creating complexity


OPERATIONS

Service level, configuration, and fault management increasing variable costs



Rapid data growth combined with increased service level expectations introduces new management challenges. Continuously evolving IT environments frequently lead to orphaned or under utilized storage resources. Rapid adoption of virtualization and cloud technologies have reduced visibility into service dependencies, health, and performance. IT frequently lacks the ability to quickly understand who is using capacity, how much they are using, if the right tier of service level is being used, and when more resources will be required. The traditional silo forms of infrastructure management lead to higher capital and operational costs, often resulting in too much time being spent with lengthy problem resolution times.

Transformation of IT into the 3rd Platform



Source: IDC, 2013

Copyright © 2015 EMC Corporation. All Rights Reserved.

EMC²
ViPR SRM Overview

4

We are now in what leading industry analysts refer to as the 3rd platform of the IT Technology Evolution. A platform that must be architected for web-scale and capable of managing billions of users and millions of applications; designed for speed, agility, mobility and growth.

However, Enterprise customers are not starting with a blank canvas, and while they want all the benefits that the 3rd platform offers, they have invested millions if not billions of dollars into an infrastructure that they must continue to maintain and grow. So how do they Mind the Gap? How do they gain the value of the 3rd platform while leveraging their current IT infrastructure? It's not an easy task. It's a journey that the customer must embark upon with a trusted advisor like EMC.

Impact to Infrastructure Managers

Key challenges for managers of data center infrastructure:

Meeting SLAs



- Relationship mapping
- Configuration and compliance

Visualize

Controlling Costs



- Capacity planning
- Utilization

Analyze

Business Agility



- New apps and services
- Impact analysis

Optimize



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


ViPR SRM Overview

5

IT managers generally have three key concerns when managing IT infrastructure: meeting SLAs, controlling costs, and increasing business agility.

Meeting SLAs requires the ability to understand end-to-end relationships and service dependencies. It also requires insight into end-to-end performance trends and compliance with configuration best practices and interoperability guidelines to ensure the environment is always configured right to meet service levels. Controlling costs requires the ability to identify under-utilized resources and track capacity consumption trends to improve capacity planning. Enhancing business agility means rolling out new applications and resolving problems faster. When discussing the Storage Resource Management Suite, you can categorize these features into three general areas: Visualize, Analyze, and Optimize.

ViPR SRM



- **Visualize** applications to storage dependencies & performance
- **Analyze** configurations and capacity growth
- **Optimize** your environment to improve return on investment while meeting SLAs

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


ViPR SRM Overview

6

ViPR SRM enables you to visualize application to storage dependencies, analyze configurations and capacity growth, and optimize your environment to improve return on investment. It also offers an extensible framework to support investments in EMC ViPR, software defined storage.

Common Presentation Layer

Configurable Views Across the Data Center



Different personas in the datacenter require specific tools, data and metrics. The Storage Team needs to understand capacity utilization and trends, troubleshoot performance issues and ensure configuration compliance to assure service levels. With greater understanding and visibility they're better prepared to optimize their resources, improve agility, and ROI.

The Operations Team need end-to-end data center visualization to quickly identify root cause and act on issues while optimizing performance and utilization.

The Management Team needs to have access to executive level dashboards that synthesize the information they need to make better decisions for the business for controlling costs while still achieving SLAs. ViPR SRM, Service Assurance Suite, and ViPR, share a common user interface and reporting technology that enable IT organizations to consolidate data and create custom dashboards and reports that meet the unique needs of each user and role.

Comprehensive End-to-end Analysis and Reporting



The EMC ViPR SRM provides customers with comprehensive end-to-end analysis and reporting in heterogeneous environments for compute, network, and storage. The total solution begins with discovery of topology and extraction of configuration data to determine and map relationships for supported devices. Polling of devices in the topology is established to gather data to provide performance, availability, and capacity management. Comparison of results to policies and compliance rules can then be used to report on compliance to best practices, mandatory requirements, and SLAs. The reports may also include alerts when policy requirements are breached.

Centralize Management

A Simple Way to Manage Heterogeneous Storage

- ViPR M&R
 - ▶ Included with ViPR
 - ▶ Capacity and chargeback reporting for ViPR virtual resources
 - ▶ Discovers physical arrays and available capacity to support provisioning capacity to ViPR
- ViPR SRM
 - ▶ Enables management of the physical infrastructure under ViPR control
 - ▶ Tracks capacity on physical arrays to support purchasing processes
 - ▶ Enables SLA management for physical and ViPR virtual storage



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

9

ViPR and ViPR SRM share a common console that centralizes management across heterogeneous storage environments. With the ViPR SRM, you can optimize today's storage investments while putting in place an extensible environment to ease your transition to software defined storage.

ViPR includes a solution known as ViPR M&R designed to help storage administrators manage the virtual resources in their ViPR environment. It offers capacity reporting and chargeback for virtual storage resources presented to ViPR tenants. And, it discovers capacity available from physical arrays to support the provisioning of storage to ViPR.

ViPR SRM extends customers' investments with comprehensive tools for reporting on the physical infrastructure managed by ViPR. It helps storage teams ensure service levels in both traditional and software defined storage environments. ViPR SRM tracks capacity consumption across the storage environment including the arrays that present capacity to ViPR to improve capacity planning and purchasing processes. It also monitors and reports on performance trends in ViPR environments to identify when service levels are not being met and provides end-to-end views of the physical infrastructure to simplify troubleshooting processes. Finally, it tracks changes and monitors compliance with design best practices and the EMC Support Matrix to ensure your environment is always configured right to meet service levels.

Module: ViPR SRM Overview

Lesson Summary

During this lesson the following topic was covered:

- Virtualization and the software defined datacenter
- ViPR overview



This lesson covered the value of virtualization across the compute, network, and storage components of the datacenter. Discussion focused on the software defined datacenter and how ViPR is used as a software only solution to virtualize storage in a simple, extensible, and open platform. These values in the ViPR product are gained through abstracting the physical storage in virtual arrays, and pooling the storage in virtual pools and making that storage available to end users placed in ViPR projects.

Module: ViPR SRM Overview


Lesson: Overview of Portal Tools and Features

This lesson covers the following topics:

- Discussion of the ViPR SRM architecture
- Frontend Portal tools and features (Browse mode)




This lesson covers an overview of the ViPR SRM architecture plus a discussion of the user interface (Frontend portal) in the user's browse mode.



ViPR SRM consists of building blocks: Frontend, Backend, Database, and Data Collectors. These are the four main base components of the suite.

- **Data Collection:** Performs the discovery and regular data-collection of infrastructure components. Multiple data-collectors can run on a single host or be spread onto multiple hosts so long as available resources allow (mainly CPU, memory).
- **Database:** The data repository for the whole application. Actually, it is best to distinguish between physical DB and logical DB in this discussion. More precisely, MySQL runs a database instance (the physical, the DB engine) and can manage many "logical" databases. Each logical database must have one and only one APG-Backend driving updates to it. Typically there is only a single MySQL Database Engine on a given Backend host.
- **Backend:** The APG-Backend is responsible for receiving data from any data collector, performing validation, and processing before storing the data into the database.
- **Frontend:** The home of the Tomcat web server application running multiple web applications that allow users and administrators to use a convenient User Interface (UI) for report browsing, modification, and creating, as well as GUI-led administration tasks.

Please keep in mind that the above illustration does not ascribe to any specific attribution to any host, they are just components and are all always present within a ViPR SRM platform. It is possible to have each on a separate host, combine some into one, or, for that matter, the components can exist as a combination of virtual Apps and/or stand alone servers integrated into a solution suite.



The basic user interface will display when the user logs into the **{Frontend Hostname or IP}:58080/APG** user portal. On the screen, you will see some items that are organized based on use case reporting groups. Down the left side of the display is the report tree groupings, including the use cases.

At the top of the page are menu entries that group scheduled reports, stored reports, and possibly reports that have been marked as favorites.

The largest part of the portal contains the icons displayed in the report panel on the right. Here, you will find dashboards and reports that aggregate related data from more than one collector. You might see VNX and VMAX switches in the same arrays report, and Brocade and Cisco switches in the same switch report.

Remember, the use-case reports only aggregate data from SolutionPacks available in the ViPR SRM suite (see Module 2 for the complete list). Data from other SolutionPacks are only available in the Report Library which becomes visible in the report tree once the first SolutionPack has been installed.

Finally, is the My Reports area which holds any modified, personalized, or customized reports you create.

Note: this entry will not show upon initial installation.

Navigating The Report Panel



Dashboards – Views high level summaries for storage and networks



Explore – Investigates configuration items including configuration and performance data



Operations – Monitors alerts, compliance, service level agreements, and other situations to watch



Planning – Analyzes capacity trends and assists in plans for future growth



Report Library – Views reports by SolutionPack

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

14

- **Dashboards:** Allows the user to display groups of pre-defined storage related reports.
- **Explore:** Provides reports categorized on discovered configuration items including storage, storage connectivity, and hosts. Additionally, since there may be many different discovered objects, the user can use the search function to search by either device type or device name.
- **Operations:** Provides reports used to monitor and validate alerts, compliance, service level agreements, and other situations to watch.
- **Planning:** Use this report grouping to analyze capacity and growth trends to plan for expansion of the infrastructure.
- **Report Library:** You will find the SolutionPack reports in the Report Library. Here, each SolutionPack reports only on configuration items discovered by the related collector. So you will find VNX arrays and VMAX arrays in separate reports in the Report Library.

Navigating The Report Tree



- **Scheduled Reports** – Timed reports
- **Stored Reports** – “Snapshot” reports
- **Favorite Reports** – Frequently used reports
- **Report Library** – List of reports
- **My Reports** – Personalized and enhanced reports specific to the logged in user

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

VIPR SRM Overview

15

The navigation tree contains reports in a foldable structure which can be expanded for more detail.


Scheduled Reports: Contains all reports that have been scheduled, as well as some statistics about their generation. Scheduled reports will be generated at the given rate, and then sent by email once finished.

Stored Reports: Exist as "snapshot" reports, which are generated at a *certain point in time*. They are useful for large reports where generation time would be excessive. Stored reports can be set up as scheduled reports.

Favorite Reports: Reports that are used frequently may be saved as bookmarks in the Favorite Reports node of the report tree. If a particular report is used most frequently for a particular user, it may be marked for viewing upon login to the application under that user's account.

Report Library: Users and Administrators will see the complete list of installed SolutionPack reports listed in this area of the Report Tree. Clicking on any report will cause the details of the selected SolutionPack report to be displayed in the Report Panel.

My Reports: Enhanced reports can be created with a Report Wizard. The report output of the Report Wizard is saved in My Reports. Pre-defined reports may also be copied into My Reports for special modification. The My Report branch of the tree is specific to the logged in user. In order for others to use these reports they must be published and imported into the other user's report tree.




The Help System enables the end user to find information on how to use and edit reports, manage users, and customize the interface to meet their specific needs.

- **Browsing Reports:** For users who want to manipulate how reports are displayed.
- **Working with Reports in Edit mode:** For users and administrators who want to edit and create reports.
- **Managing ReportPacks:** For users and administrators who want to create, import, and export ReportPacks.
 - * **Note:** A ReportPack is a set of one or more reports usually packaged together as part of a SolutionPack
- **Managing Users:** For administrators who want to add users and set up user restrictions.
- **Administering the Interface:** For administrators who want to customize the interface for all users.
- **Customizing User Settings:** For users who want to set personal preferences.
- **Report Library:** For users and administrators who work with installed SolutionPacks.


The question mark in the white circle next to the Help option in the Menu bar allows the end user to view all of the installed licenses in the system. Each license will identify the type of license (whether it is a trial license or was granted by EMC) and the expiration date.

Toolbar: Administration



Assuming the user has administrative privileges, the **Administration** selection gives access to the administration portal. As you will see later, there are several options available to administrators which are very useful to the data storage administrator. Of particular interest to storage administrators are the **Management of Database Metrics** under the Modules group and List/Import ReportPack. We will discuss these options in detail later.

The Tools Bar Explained: Lower Level



The lower part of the tools bar is used to mostly interact with the current report pane. These tools deserve slides for themselves. Let's begin with a brief description of each:

Edit Mode Button: This allows the user to toggle between Browse Mode (to view reports) and Edit Mode (to customize reports). Toggling may be disabled by an administrator.

Simple Search: The search bar exists in two forms: simple and advanced search. The textbox is part of the simple search: just type in the object name you are looking for and let the search engine find it for you.

Advanced Search: The advanced search is under the drop down icon, just to the right of the textbox. It offers much more flexibility in search. This item will be described in greater detail later.

Display Tab: Used to change the look of a report. May be disabled by an administrator.

Exports Tab: Contains options to export a report in one of several formats to be viewed offline.

Tools Tab: Permits several actions to be taken on a report (such as printing, storing, emailing, and others). There are also related tools for managing reports including a wizard.

Menu Bar: Lower Utilities - Display


The screenshot shows the 'Menu Bar: Lower Utilities - Display' section of the EMC² VIPR SRM Overview interface. The top navigation bar includes 'Settings', 'Administration', and 'Help'. Below this, the 'DISPLAY' menu is highlighted, with a tooltip indicating it 'Controls how data is displayed in a report'. The 'DISPLAY' menu is open, showing various configuration options:

- Display:** 'two reports per line' (dropdown)
- Report:** 'mixed using defaults' (dropdown)
- Rendering:** 'static charts' (dropdown)
- Summary Statistics:** 'None' (dropdown)
- Time Selection:** '1 hour' (dropdown), 'average' (dropdown), 'last' (dropdown), '1 day' (dropdown)
- Actions:** 'View in Full Page', 'Revert to Default Report', 'Revert Layout Positions', and an 'Apply' button.

The interface also includes a search bar and an 'EDIT MODE' button. The bottom of the page features the EMC² logo, copyright information, and the page number 19.

The Display utility permits the end user to select how the data in the report will be displayed. Selections are made for the report type (charts and graphs), aggregation (summarization of group values) and report time base settings. Some report display parameters may be customized using the Display option in the menu bar. Which parameters may be modified and to what degree depend upon the selected report definition.

Display Tab



Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

20

The Display tab allows you to alter the view of the current report (to tweak it). Three sections can be identified under Display tab: Report Type, Time Selection, and Actions.

Actions: shows action once you are done with other sections.

View in Full Page: displays current report in full page, in a new window.

Revert to Default Report: brings the report back to its original, yet default, settings after you have made several changes (or tweaks).

Apply: applies the changes set from Report Type and/or Time Selection sections.

Time Selection: groups Time range selection options of our current report.

Display [...] values: used to change the value aggregation to be used; whether real-time values or aggregated values.

Using [...] aggregation: used to select the type of aggregation (average, max, etc.) to use (sometimes called “aggregation function”).

Over [...]: used to change time range selection of the report. Note that the Time Range Quick Switch arrows allow movement backward or forward in time. Once clicked, it then shows clickable calendar icons that allow you to select date-time start or date-time end to exact time range you’re looking for your report.

Padlock icon: toggles the lock to the modified Time Selection so default report time range selection is locked while browsing to any other reports and thus compare more easily.

Custom time range: defined by clicking the opening or closing double quote sign. Free form text box allows you to type in the time expression or using the calendar page icon in order to use the calendar pop up helper.


Report Type: groups options that, based on default type of report, will allow you to change options in your current report.

Display: used to change the number of reports per line. Especially useful for mixed reports.

Report: used to change on the fly the report type. This option depends on the default report’s family currently set.

Rendering: some graphics can be rendered either in static mode or in dynamic mode. In Dynamic charts, hovering mouse pointer over a graph line will show a tip tool containing pointed metric name with exact value and time. While dynamic charts look nicer and are easier to use, static charts may be quicker to load.

Summary Statistics: used to overlay useful information on your graph report type: statistics information, pin points, trending, etc.



ViPR SRM bases its time range selection using three framing options: “previous”, “current”, or “last”.

From Display menu and when using report editor, you have to choose among these options:

- **previous**: selects the entire and complete previous range, according to period boundaries.
- **last**: selects one complete period range back from present time.
- **current**: selects the current period range, according to period boundaries.

Let’s assume the current date and time is: 15th of May, 08:00 AM.

- The previous day (24 hours) boundaries become exactly: May 14th from 00:00AM to 11:59:59 inclusive.

This could be translated in selecting ‘yesterday’ time range selection.

- The last day (24 hours) boundaries become exactly: May 14th 08:00AM to May 15th 08:00AM inclusive.

This could be translated in selecting the very last 24 hours from the actual present time.

- The current day boundary becomes exactly: May 15th 00:00AM to 11:59:59PM inclusive (some projection in the future, could be useful for trending and forecasting report time range selection).


The same concept applies for weeks, months, years, quarters and so on.

This concept is used in report tweaking but also in report editing.

Aggregate Introduction

Definition

An aggregate is a data record summarizing many values with an effect of reducing the sampling rate. It compensates this effect by providing more statistical details.



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

22

Another important notion in ViPR SRM is the use of aggregates.

The Backend's responsibility is to compute and populate aggregates in real-time.

ViPR SRM keeps two types of values in its database: real-time values and aggregated values.


Real-time values are the value of metrics as they are collected, representing 1 data point in time. They are aligned with the Collector's polling cycle.

Aggregated values are the compound summary of several real-times values put together in order to improve report performance and save disk space. ViPR SRM's default defines these aggregates as:

- 1-hour: when polling cycle is 5 minutes, it takes 12 samples to generate one single aggregate record for a given metric.
- 1-day: when polling cycle is 5 minutes, it takes 288 samples to generate one single aggregate record for a given metric.
- 1-week: when polling cycle is 5 minutes, it takes 288×7 samples to generate one single aggregate record for a given metric.

Aggregated values are created and updated in real-time. The values are collected and received at the APG-Backend module. This real-time aggregation of values always assures up-to-date and precise values.


Menu Bar: Display Actions



The Actions block of the the Display menu allows for management of the report view.

- **View in Full Page** - Allows the report to be viewed in a new browser window. The browser window will open automatically when this option is selected.
- **Revert to Default Report** - If Time Selection changes are made, the report has been zoomed in, or the Report Type block has been modified, the changes can be quickly reversed by selecting this option.
- **Revert Layout Positions** - If a Report contains several graphs in a single window and the positions of the graphs on the window are changed for any reason, this will cause the entry **Revert Layout Positions** to appear in the Actions menu as shown above. If this option is selected, the graph positions will realign to their original defined positions.

Menu Bar: Lower Utilities - Exports



EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

24

Exports options allow reports to be saved as files in any of several different formats compatible with the report. For example: pdf, csv, xls, and others. Any format that is not compatible with the currently selected report will be shaded indicating it is unavailable for this report type.


Tools - Action on Reports



The Tools menu allows actions on the current report.

Show Report URL will pop up a message with a textbox containing the URL of the report, already highlighted and ready to copy in buffer. This could then be used to paste it in an email, IM service, or in a document.

Tools - Action on Reports



Prints the selected report to a locally available printer

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

26

Print this Report will reformat the display in a printer-friendly version of the current report, allowing you to directly print the report.

Tools - Action on Reports



Stores the selected report into the **Stored Reports** branch of the report tree

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

27

Store this Report allows you to take a snapshot of the current report, as it is shown. The current report layout along with its dataset will be copied (snapshot) and available under the **Stored Reports** section in Report Tree. From then on, if you want to return back to this report, all you have to do is to expand the Stored Reports entry and select the report. If an updated version of the report is required, you will have to process it again like you would from report's original node.


Tools - Action on Reports

The screenshot displays the EMC SRM user interface. At the top, a blue navigation bar contains the date 'Friday, February 14, 2014' and links for 'Settings', 'Administration', and 'Help'. Below this is a secondary bar with 'EDIT MODE', a search box, and tabs for 'DISPLAY', 'EXPORTS', and 'TOOLS'. The 'TOOLS' tab is active, showing a 'Tools' section with 'Real-Time Grapher' and 'Report Wizard', and an 'Actions' section with 'Show Report URL', 'Favorite this Report', 'Print this Report', 'Schedule this Report', 'Store this report', and 'E-mail me this Report'. A blue arrow points from the 'Favorite this Report' button to the 'Favorite Reports Management' dialog box. This dialog has a title bar and a description: 'Favorite reports act like bookmarks to some specific reports so that they are always quickly available. One of them can be set as "first report to appear after login"'. It features a 'New Favorite' button and a 'Configured Usable' section with a 'Link Name' field (containing 'Configured Usable') and a 'Use as Login Report' checkbox. At the bottom are 'Clear Login Report' and 'Save' buttons. On the left, a sidebar shows a tree view with 'Favorite Reports' selected. The footer includes the EMC logo, copyright text 'Copyright © 2015 EMC Corporation. All Rights Reserved.', and page information 'VIPR SRM Overview 28'.

Favorite this Report feature allows you to quickly go back to a report node you use more often, like a browser's bookmark. The bookmarked report is then available under "**Favorite Reports**" section of the Report Tree. Selecting an entry from the Report Tree will always bring you to the original targeted report.

This action is in contrast with previous "Store this Report" where a snapshot is completely set and kept aside. If an original report node gets deleted or moved, this Favorite Report entry will show as a Broken link, like any browser would do with a web page bookmark. A favorite report can be set as your default report page when logging in.

Tools - Action on Reports



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview


29

Schedule this Report is similar to the “Stored this report” feature but with more options and flexibility. When clicking this option, a new pop-up will show called the **New Scheduled Report**. This is where the user can set various parameters of a scheduled report through tabs.


Here is an overview of each tab when setting a Scheduled Report:

- **Scheduling:** This is used to set report name and set schedule definition and status.
- **Storage:** Select if report is mean to be stored under “Stored Report”. A confirmation setting about report generation can be sent out.
- **Email:** Formatting your output report for email.
- **Remote Transfer:** Sending you report to remote service (network drive, FTP, etc.)
- **Alert:** Transfer the report to Alerting Module adapter.

Tools - Action on Reports



E-mail me this Report allows you to send yourself a copy of the current report in PDF format. Please note this will work only if you set your e-mail address correctly under the Settings menu above.




The Report Wizard will permit Table, Graph, and TopN to be quickly generated depending on the the filter used to select metrics. The **View Type** filters used for metric selection are: APG Topology, Core Networking, Device Browser, or Devices by Location. In the case shown, APG Topology is used to select metrics for a server. The metrics chosen are Availability, CurrentCPUUtilization, and UsedMemory (not shown). When the final selections are made, the Run button and the Add Metrics button will highlight.

Report Wizard Output

The screenshot displays the 'Report Wizard' interface. At the top, there are buttons for 'Run', 'Save in "My Reports"', and 'Quit Wizard'. Below these, a section titled 'My Reports / Wizard' shows a line graph of 'UsedMemory' over time. A blue arrow points from the 'Collector 235' entry in the 'My Reports' list to the 'Specify a report name...' dialog box. The dialog box contains the text 'Please specify a name for the report.' and a text input field with 'Collector 235' entered. There are 'Cancel' and 'Ok' buttons at the bottom of the dialog. The background graph shows multiple data series for 'UsedMemory (bytes)' over a 24-hour period.

Clicking 'Save in "My Reports"' pops up the dialog box to name the report. Once named, it is saved automatically in the My Reports branch of the report tree.

Basic User Interface Report Elements



The basic user interface report element includes multiple reports for each SolutionPack. Hovering the cursor over a report entry may pop up a definition box for the specific report. The example shown here reveals a pop-up definition box for the Throughput by Switch(/s) report in the Brocade FC Switch SolutionPack. The icons in the upper right hand corner of the report, starting from the left, allow the legend to be hidden, a link to this element to be created in the personal view, make the element taller, wider, or moved to another area. To view this report individually, click on the report title.

Control Icons for Graphs and Tables

Graph Control Icons		Table Control Icons	
	Displays the graph line legend at bottom of the graph		Creates a link to the graph in a personal view
	Creates a link to the graph in a personal view		Makes the table shorter
	Makes the graph view taller		Narrows the table
	Makes the graph view wider		Makes the table taller
	Allows the graph to be dragged to a different place		Allows the table to be dragged to a different place

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

34

To view the graph and table control icons, place the cursor in the field of the graph or table and the icons will appear in the upper right-hand corner of the view. When individual graphs or tables are viewed, only the pin icon is available to link the viewing element to a personal view in the My Reports node of the Report Tree.

Report Drill Down: Saving a Personal View



Clicking on a report title in the Report Library creates a display for that specific report which may contain one or more graphs and/or tables. Each graph and table title has a hyperlink to a more detailed view of that particular graph or table. When individual graph or table views are being examined, only the pin control icon is available. Clicking on the pin icon allows for linking this report to a personal view in My Reports.

At the bottom of the graph, a legend displays all graph lines in rectangular boxes. The boxes (and the graph lines they represent) may be selected or deselected manually for view in the graph. They may also be selected as a group with the All, None, Auto, and Invert buttons. Graph customization allows the graph lines to be modified in color, curve type, width, and marker.


Report Drill Down: Graph Customization

All >> Report Library >> Brocade FC Switch >> Throughput by Switch(/s)

Brocade FC Switch / Throughput by Switch(/s)

March 2014, Tuesday 18 » Wednesday 19, 2:54 PM EDT

This report displays the Throughput by switches. Typically, this is the sum of throughput for all FC ports belonging to a switch.



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

36


At the bottom of the graph, a legend displays all graph lines in rectangular boxes. The boxes (and the graph lines they represent) may be selected or deselected manually for view in the graph. They may also be selected as a group with the All, None, Auto, and Invert buttons. Graph customization allows the graph lines to be modified in color, curve type, width, and marker. Any customization areas grayed-out are not available for this graph.

Report Drill Down: Zooming



Drilling down on a section of the graph for a greater level of detail is possible when clicking on an area of the time line and while holding the mouse button, drag the mouse cursor to the opposite corner of the in the highlighted space until the complete area of interest in time and amplitude is shaded. Once the cursor is released, the new graph will be created.

Graph Modifiers: CPU% Example



Graph area covers the last full day from the time the graph was created. Thresholds were set at 60% for a major state, and 80% for a critical state. Threshold settings may be changed in edit mode. Edit mode will be discussed in another module.

In the Display utility, the type of graph may be changed to make reading the data in the graph easier and understanding the data displayed more accurate to the conditions the end user wishes to examine in the infrastructure. Report rendering is set to either static or dynamic. By default all reports display dynamic data in real-time values. Static reports are generated on a schedule and display data at a point in time and do not change over time.

To show or hide the labels at the bottom for the graph lines, simply click on the label of choice and it will gray out. Once the label has been grayed out, click on the Apply button and the graph line for that label will no longer be present in the graph. To return it, click the label once again, and when it is highlighted click the Apply button. The graph customization button next to the Apply button may be used to change the color of the graph line associated with the particular label. Other modifications that may be made within Graph customization are: curve style, width of the graph line, and graph marker.

Table Report Modifiers

All >> Report Library >> EMC VNX >> Summary >> Array Summary

Summary / Array Summary

March 2014, Thursday 13 > Thursday 20, 3:26 PM EDT | Last 1 Week (cancel filtering)

2 elements found, displaying all elements.

System Type	Serial #	System Name	IPs
Block	APM00133851629	myVNXblock	10.126.66.45, 10.126.66.46
File	APM00133851629	myVNXblock	10.126.66.44

Report generated in 0.201s.

System Type

Please enter a value to filter on (press Enter to activate):

Filters pass through, not out



Row Ordering

Click to stop and remove filtering

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

39

The icon in a table report, which appears in every column as two triangles, one on top of the other, is used to order the rows in the table from that column. When clicked, the last row entry becomes first and all rows invert order from that point. If all column entries are the same, no reordering takes place. If all column entries are alpha characters, then the order simply reverses top to bottom. If the column entries are numeric, the rows are ordered by the numerical ordering in that chosen column.

Column filters permit rows to be filtered for specific content within a column. For example, the word File could be used to filter only rows in the System Type column which have the word File in them. Filtering filters through, not out for value selections. Once in filtering mode, a message in the window next to the path title Summary/Array Summary will state "cancel filtering" to return the table rows back to their original entries.

Wildcard entries may be used with the character (*) star to retrieve a large number of rows. For example: the Serial number column could be filtered with APM* to return all EMC data storage arrays using this designation in the serial number.

Module: Overview of the ViPR SRM User Interface

Lesson Summary

During this lesson the following topics were covered:

- Reports panel
- Navigation tree
- Menu bar



This lesson covered logging into SRM, managing reports, and accessing the utilities of the menu bar.

Module: Overview of the ViPR SRM User Interface

Lesson: Examples Demonstrating the Use of Reporting

This lesson covers the following topics:

- A discussion of use case examples for enhanced reporting



This lesson covers an overview of some examples to illustrate the purpose and need for enhanced report creation.

Use Case 1 – Chargeback by Group

Initial situation:

- The default report mixes data from any business unit and application together, thus making it harder for final chargeback or view per application consumers

Improvement:

- Re-structure of the report layout by integrating a drill down per business unit and application

Benefits:

- Allows a report of service level tagged storage consumption per application inside of business units

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

42

This use case example will demonstrate the re-purpose of a Chargeback by Group report.

Chargeback by Group report displays the Service Level Chargeable capacities roll-up for all Physical Hosts, Hypervisors, and VM guests per user-defined group.

The use case for report customization is about drilling down this report per Business Unit and Application prior to the Chargeback report. The work is then to simply re-structure the report layout by integrating a drill down definition.

The final benefit from this customization, is that the final report will list service levels consumption based on a selection of business units and applications.

Original source report: All>>Dashboards>>Storage>>Chargeback by Group

Use Case 1 – Chargeback by Group

Original report: All>>Dashboards>>Storage>>Chargeback by Group

Primary Used Chargeable

10 elements found, displaying all elements.

Group	Host Type	Platinum	Gold	Silver	Bronze	Fast VP	Virtual Storage	Other	Primary Used Chargeable	Primary Used
All Physical Hosts	Host	0.00 GB	49.21 GB	71.33 GB	29.26 GB	0.00 GB	0.00 GB	0.00 GB	149.80 GB	149.80 GB
Provider Tenant	Hypervisor	0.00 GB	721.44 GB	341.33 GB	240.50 GB	0.00 GB	0.00 GB	0.00 GB	1.27 TB	3.82 TB
Sales	VirtualMachine	0.00 GB	91.33 GB	275.93 GB	430.63 GB	0.00 GB	0.00 GB	51.67 GB	849.57 GB	849.57 GB
Human Resources	Host, VirtualMachine	0.00 GB	263.92 GB	71.33 GB	23.20 GB	0.00 GB	0.00 GB	0.00 GB	358.45 GB	358.45 GB
Manufacturing	Host, VirtualMachine	0.00 GB	0.00 GB	0.00 GB	610.77 GB	0.00 GB	0.00 GB	0.00 GB	610.77 GB	610.77 GB
All Passive Hosts	PassiveHost	0.00 GB	0.00 GB	0.00 GB	0.00 GB	0.00 GB	0.00 GB	0.00 GB	0.00 GB	0.00 GB

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

VIPI SRM Overview

43

The original Chargeback by Group report consists of four (4) sub reports: Primary Used Chargeable, Primary Presented Chargeable, Total Used Chargeable, and Total Presented Chargeable. For each of the defined groups, you have the storage service level amount associated.

Original source report: All>>Dashboards>>Storage>>Chargeback by Group

Use Case 1 – Chargeback by Group

Modified report/final result

UC1 - Chargeback by BU and Application

- Engineering]
 - vCloud
- Human Resources]
 - PeopleSoft
- T]
 - Asset Tracking
- Manufacturing]
 - Inventory Tracking
- Sales]
 - CRM
- Sales|Human Resources]
 - CRM

Application

Business Units

UC1 - Chargeback by BU and Application

- Engineering]
 - vCloud
 - Chargeback by Group**
 - Primary Used Chargeable
 - Primary Presented Chargeable
 - Total Used Chargeable
 - Total Presented Chargeable

Primary Used Chargeable

2 elements found, displaying all elements.

Group	Host Type	Platinum	Gold	Silver	Bronze	VF
All VM Guests	VirtualMachine	0.00 GB	2.11 TB	37.36 GB	44.00 MB	0.0
Engineering	VirtualMachine	0.00 GB	1.85 TB	37.36 GB	44.00 MB	0.0

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

44

With the modified version of the report, better granularity and more drill down are possible before reaching the Chargeback by Group report details.

Source of modified report: All>>ViPR SRM Reporting class>>Module 01 - ViPR SRM Overview>>UC1 - Chargeback by BU and Application

Use Case 2 – Storage Systems Report

Initial situation:

- For a customer delivering service across multiple sites, location details are missing from the standard Storage Systems report

Improvement:

- Restructuring the report adds location information and some useful details

Benefits:

- Array summary per location, along with numbers of array and total capacity size. Can be on a GeoMap report

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

45

This use case scenario calls for improvement of a Storage Systems report. The default report does not contain any details about the location (or sites) of the listed storage systems. This is missing for the customer who does business across multiple physical sites.

Fortunately, ViPR SRM M&R flexible reporting can fix this. Simply by restructuring the report by location, the customer can have summary and detailed information with the click of mouse. This improvement also prepares a GeoMap report where sites can be pinned on a geographical map (Google or OpenStreetMap). Notice that location property can also be replaced by datacenter name or site code.

Use Case 2 – Storage Systems Report

Original report: All>>Explore>>Storage>>Storage Systems

Storage System Type	Serial Number	Storage System Alias	Model	Vendor	Firmware	Total Raw (TB)	Last Collection Time
Block	888973588710	888973588710	VMAX40K	EMC Corporation	5876.268.174 date:02.05.2014 patch_date:02.05.2014	410.53	Tuesday, July 22, 2014 4:29:52 PM EDT
Block	5603141	XIV 5603141	114	IBM	11.4.1.a	323.18	Tuesday, July 22, 2014 4:15:09 PM EDT
Block	888973588727	888973588727	VMAX40K	EMC Corporation	5876.268.174 date:02.05.2014 patch_date:02.05.2014	233.07	Tuesday, July 22, 2014 4:29:56 PM EDT
Block	888973588777	888973588777	VMAX40K	EMC Corporation	5876.268.174 date:02.05.2014 patch_date:02.05.2014	131.45	Tuesday, July 22, 2014 4:11:14 PM EDT

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

46

The above shows what the initial Storage System report looks like before modification. This report lists useful information about storage systems that are under the management of a ViPR SRM installation. Notice that no information about the location of any of the storage systems (site and/or datacenter name) are presented.

Source or original report: All>>Explore>>Storage>>Storage Systems

Use Case 2 – Storage Systems Report

Modified report/final result

- UC2 - Location of Storage Systems
 - Storage Systems**
 - Hopkinton
 - 888973588727
 - 888973588777
 - APM88902001234
 - Montreal
 - Santa Clara

UC2 - Location of Storage Systems / Storage Systems

July 2014, Monday 21 » Tuesday 22, 4:00 PM EDT | Last 1 Day: average on real-time

Storage Systems		
3 elements found, displaying all elements.		
Site Name	Number of arrays	Total Raw Capacity
Hopkinton	3	364.53 TB
Montreal	3	436.97 TB
Santa Clara	2	72.50 TB

Hopkinton

3 elements found, displaying all elements.

Storage System Type	Serial Number	Storage System Alias	Model	Vendor	Firmware	Total Raw (TB)	Last Collection Time
Block	888973588727	888973588727	VMAX40K	EMC Corporation	5876.268.174 date:02.05.2014 patch_date:02.05.2014	233.07	Wednesday, July 30, 2014 7:41:55 PM EDT
					5876.268.174		Wednesday,

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

47

The final result includes an intermediate report that lists the sites along with a number of arrays and the total raw capacity of all the arrays. The total simply gives an idea about the size (importance) of the site and can be compared among other sites, at a glance. At the bottom of the report, the storage systems of a site are listed, along with the same level of details from the original report. Also, this intermediate report is interactive. This means selecting a site entry will update the original report content.

Notice that the success of such a report is possible only when the appropriate location information is tagged correctly. This is done by the Data Enrichment feature, under the responsibility of a ViPR SRM Administrator.

The new report is available on the demo, as follows: All>>ViPR SRM Reporting class>>Module 01 - ViPR SRM Overview>>UC2 - Location of Storage Systems>>Storage Systems

Use Case 3 – Chargeback Virtual Infrastructure

Initial situation:

- Customer defines a Chargeback model for shared virtual infrastructure. Standard library has none specifically defined

Improvement:

- New report definition, application of a Chargeback model in a report

Benefits:

- Customer report tailored with its own Chargeback model



Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

48

This use case demonstrates how a new report set can be created in response to a specific need. This time, the customer would like to apply a chargeback model for its tenants of the virtual infrastructure by applying the following model, on a monthly basis:

- vCPU is charged 25\$
- Memory is charged 10\$ per GB
- Disk Space Usage is charged 0.50\$ per GB (based on 95th percentile)

Use Case 3 – Chargeback Virtual Infrastructure

Tabular report proposal:

Tabular report / Chargeback VDC - Drill-down

July 2014, Monday 21 x Tuesday 22, 4:00 PM EDT | Last 1 Day: average on 1 hour

Chargeback VDC - Drill-down			
Drill-down Chargeback report - per Customer			
3 elements found, displaying all elements.			
Customer	VM Count	Customer unit total cost (\$)	
VMware	29	5,157.14	
VCE	15	3,835.76	
EMC	61	9,707.56	
Total VM	105	Total	18,700.46

EMC

Drill-down Chargeback report - per Business Unit

3 elements found, displaying all elements.

Business unit	VM Count	Business unit total cost (\$)
Human Resources	4	741.63
Sales	38	4,597.70
IT	19	4,368.22
Total		9,707.56

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

49

The above demonstrates the final result of the report in a tabular view format. The drill down is per customer name and business units (of the tenants). For each level, customers can have the total of the chargeback: per customer and a breakdown per business unit.

Use Case 3 – Chargeback Virtual Infrastructure

Chargeback VDC - details

This is a chargeback report of a VDC.

Charge rules are as follow (per VM):

- vCPU is charged 25\$/vCPU/month
- Memory is charged 10\$/GB/month
- Disk usage is 0.50\$/GB/month (95th percentile)

19 elements found, displaying 1 to 10.

<< < 1 2 > >> Show 10 ▼

Name	Customer	Business Unit	IP	Guest OS	State, Status	Nb vCPU	Cost of vCPU (\$)	Memory Size	Memory Cost (\$)	Disk Used space (95th)	Disk Used space Cost (\$)	VM Total Cost (\$)
mhmcb097 - SRM 3.0	EMC	IT	172.8.166.97	SUSE Linux Enterprise 11 (64-bit)	Powered On, Normal	4	100.00	40 GB	400.00	72.71	36.30	536.30
mhmcb040 - DB	EMC	IT	172.8.166.40	Microsoft Windows Server 2008 R2 (64-bit)	Powered On, Normal	4	100.00	12 GB	120.00	424.39	212.19	432.19
mhmcb098 - SRM3.5 - FE&BE	EMC	IT	172.8.166.98	SUSE Linux Enterprise 11 (64-bit)	Powered On, Normal	4	100.00	32 GB	320.00	0.02	0.01	420.01

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


VIPI SRM Overview

50

Drilling down to the detail, per business unit shows the values for each component of the chargeback model: cpu, memory, disk space used.

Use Case 3 – Chargeback Virtual Infrastructure

Advanced transformation sample:



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

51

The above demonstrates that ViPR SRM presents various report types that can be used to better visualize some important details or bring another view of the results. The above shows a derivation of the chargeback model report in a treemap. The treemap groups data per level (customer, business unit, VM) where cell size and color define visual parameters. In the above, cell size is about how much money a given VM generates (then group by business unit and then by customer). The cell color is about the footprint of the VM. This footprint is based on two important parameters: memory and cpu allocation. The footprint is calculated by a multiplication of memory size (in GB) and number of vCPU. The footprint value is then compared with a given (report configured) chart and assigned a color. The legend at the bottom explains the various possibilities defined. More details about a given VM is provided when hovering the mouse on a VM.

Conclusion about Use Cases

- Standard report library meets the average requirements
- Above use case customization relies on data enrichment
- Typical data enrichment involves (but not limited to):
 - ▶ Customer and/or Business Units
 - ▶ Application name
 - ▶ Location
 - ▶ Device grouping
 - ▶ Service Level
- Always work on a copy of the report from the standard library
- Report customization is endless



From the previous use cases, you can note that customization can be done on reports from the provided standard library of reports. This standard library is meant to meet the average requirements of customers in terms of reporting in ViPR SRM. This means not all use cases are fulfilled out of the box. When thinking about report customization, most of the requirements involve data enrichment (DE). This DE is a feature of ViPR SRM under the administration of a ViPR SRM Administrator (or its delegated). DE must be defined on the collection side of the ViPR SRM platform and before report customization can be started.

It is important to notice that customization of report from the standard library should never be done on original report. One of the main reasons is that upon an upgrade of ViPR SRM, standard report library can be overwritten. Thus any modification to an original report will be lost. Plus, many standard reports rely on other reports (from the standard library). A rule of thumb is to always work on a complete copy of the original report and never directly on the original ones.

Module: Overview of the ViPR SRM User Interface

Lesson Summary

During this lesson the following topics were covered:

- Use case examples for enhanced reporting



This lesson covered possible enhancement of some report from the standard library.

Module: Overview of the ViPR SRM User Interface

Lesson: Overview of Data Enrichment

This lesson covers the following topics:

- Data enrichment features



This lesson covers an overview of the data enrichment feature in the context of reporting.

Data-Enrichment – Discussion

Characteristics:

- Complements collection: machine vs human view
- Customization of reports (segregation, grouping, include/exclude)
- Used for Role-based reports
- Expand convention: tag properties based on convention
- Location information for map reports
- Customer and/or Business Unit information for Chargeback reports
- Used in a broad range of SolutionPacks (normalization)
- Group Management allows manage a defined set of Data Enrichment rules



Copyright © 2015 EMC Corporation. All Rights Reserved.


ViPR SRM Overview

55

Use cases where Data-Enrichment applies are broad:

- **Complements collection:** When raw data is extracted from the source device or through a tool, data can be difficult to understand. In the networking world for example, the model of devices is extracted from SNMP OID and is identified by a series of numbers in dot notation. Translation into real model numbers for human readability is better suited in reports. As another example, an Array is identified by its Serial Number... one could decide to complement this with a machine name - better for reports and human readability.
- **Customization:** One of the most wanted features in report management, customization allows the matching of generic collection with some personalized/custom properties for report enhancement purposes.
- **Role-based reports:** A service provider may want to filter out information that does not belong to a specific customer or business unit/organization.
- **Expands convention:** Tagging information is based on accessible information. For example, when devices have a naming convention, Data Enrichment can split, match, and tag complete values to expand the information and make it more valuable.
- **Tagging complete information for reports:** Used in geomap and chargeback reports, tagging information is available only from another tier source - typically from a centralized CMDB.
- **Pre-defined SolutionPack reports:** VMAX, VNX, Brocade, Cisco, VMware, etc... all have specific property corrections applied while using Data-Enrichment.
- **Group Management:** A management feature in ViPR SRM that allows you to manage a define set of data enrichment rules.

Use Case – VDC Inventory



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

56

When collecting from the VDC using VMware vCenter SolutionPack, the remote source does not provide details about which VMs belong to which customer and/or business unit. So it can be difficult to create a segregating report (requires a lot of hard coding). In order to simplify this task and to automate the report creation, “tagging” the VMs can leverage the data into more valuable capabilities.

The above report, customized with the help of Data Enrichment, is straightforward and provides the number of VMs per customer and then, through drill-down reports, on per business unit basis.

Source: modified version of: All>>ViPR SRM Reporting class>>Module 05 - Formulas>>Lab 5 - vCenter Inventory (part 2)

Use Case – VMAX Service Level

Service Level Capacity

6 elements found, displaying all elements.

Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used
Platinum	971.16	438.14	533.02	431.90	45.11
Silver	17,913.70	17,132.45	781.24	15,842.69	95.64
Bronze	147,312.51	61,509.07	85,803.44	60,351.75	41.75
FAST VP	91.22	91.22	0.00	91.22	100.00
Other	8.18	0.06	8.11	412.82	0.79
Gold	149,541.80	145,123.49	4,418.32	140,810.44	97.05
	315,838.56	224,294.44	91,544.13		

Gold

4 elements found, displaying all elements.

Array	Vendor	Array Model	Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used
888973588727	EMC Corporation	VMAX40K	Gold	3,890.42	3,882.85	7.56	3,847.08	99.81
888973588777	EMC Corporation	VMAX40K	Gold	6,523.36	5,919.24	604.12	6,033.48	90.74
888973588710	EMC Corporation	VMAX40K	Gold	139,050.94	135,281.40	3,769.55	130,890.99	97.29
888976588482	EMC Corporation	VMAX10K	Gold	77.08	40.00	37.08	38.46	51.90
				149,541.80	145,123.49	4,418.32		

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

VIPI SRM Overview

57

Service levels are seen in the chargeback reports, and there is also a Service Level Capacity report that catalogs all of the array storage by service level. The difference is that this report sums all LUNs from the array point of view. LUNs are counted even if they are not allocated to a host.

The table at the top is a master view. Click on a service level and the component arrays will display in the detail view below. You cannot drill further to identify the individual LUNs in this report, but the All LUNs report does show individual array LUNs by service level.

The VMAX Service Level report shown above reports capacity per service level.

Service Level is not known by the arrays. This information is not available on the array at collection time. Data-Enrichment helps in categorization of the objects of the array. Thus a report can be shown by Service Level. You can change the Service Level tag name by configuring Data-Enrichment.

Challenges

Raw collection sometimes cannot provide all proper/complete properties

Missing/incomplete information from source objects

Objects not described by non-technical attributes

Object attributes reside in external source

Raw attributes not human understandable (encoded)

Raw attributes may require translation: device status of "1" ...?!

Attributes inspection/reformat/workaround for tagging

Platform Admin should study the data available and know what they want to achieve

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

VIPR SRM Overview

58

The installation of the discovered SolutionPack objects are defined by the properties attached. This property set is taken from the available data from the source system. There are challenges that face an SRM platform administrator (and Storage administrator) in regards to discovered objects. As a platform administrator, your goal should be to achieve/fix the following:


- **Missing information from the source:** Some attributes can give you clues. Ex: Storage Service level for a given LUN or SG.
- **Objects do not have information:** Customer name, device usage, status, group ownership.
- **Some information only available externally:** Either from XLS, CSV, WebService, Database source, CMDB, etc.
- **Source provides information in an incomplete or cryptic format:** Need to rework before using in reports.
- **Attributes are not formatted/normalized correctly and need rework:** Can patch or workaround the source problem.

In any case, the platform administrator should know what data is available and envision the desired result. Use Data Enrichment to properly identify the scope of the targeted metrics, select the proper components, and configure correctly. Some testing may be required, along with some familiarity with the Data-Enrichment feature and reporting.

Data Enrichment

Definition

Data enrichment refers to the process used to enhance, refine, or improve raw data properties with the intent to make the data a more valuable asset.



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

VIPR SRM Overview

59

Data enrichment generally refers to the process used to enhance, refine, or improve raw data properties with the intent to make the data a more valuable asset.

1. We start with the raw data that is collectable.
2. Add additional elements to the data, through the use of data enrichment configurations.
3. Modify existing reports to add the enriched data to the report.
4. Resulting in reports that are more informative and targeted to the local environment.

Data-Enrichment – Features

In-line, real-time rule analysis
Support of multiple rule sets (serial/in series)
All SolutionPacks include a base rule set named global-enrichment
Ability to tag both TimeSeries and event-based data
Implemented by filtering <u>at the collection layer</u>
Update timer: re-read file, re-fetch from source
Flexible key value matching methods: string, regex, wildcard, or range
Flexible tagged value: re-use, overwrite, complement
Un-numbered matching key and property assignments

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


ViPR SRM Overview

60

Features of the Data-Enrichment:

- In-line components sitting on the collection layer.
- A single Data-Enrichment entity can process many rule sets for different purposes within the same collector.
- All SP includes a generic Data-Enrichment components used for Customization. It also contains a template rule set (Customer, Business Unit).
- Data-Enrichment has the ability to work with all Timeseries and Event-based data.
- All Data-Enrichment are implemented by the Property-Tagging-Filter component and sitting at the collection layer for any given collector.
- Rule sets can be updated by re-reading information source on time-based. It can also be updated real-time if using WebService accessor.
- Flexible rule key evaluation for an exact string match, regular expression, sql-like statement, or range value (only for numbers).
- A matching key, provides flexible value tagging for re-use of property value, overwrite of existing property value, complement value (add text string to existing).
- Un-numbered matching key and set properties. For example for a given device name, you can tag with 100 properties. Or the reverse, very restrictive (many key differentiation) will tag with one single property value. And both mix... also handles empty values.

Group and Service Level Property Tagging



- **Device Grouping:** Tag **hosts, hypervisors, or virtual machines** with groups based on name or pattern
- **Service Level:** Tag **LUNs** based on performance criteria: disk technology, disk size, disk speed, etc.
- **Business, Customer & Application Ownership:** per device name, assign Customer, business unit and application data

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

61

Chargeback groups and service levels are created through collector property tagging filters. When the collector retrieves information that matches a given criteria, it adds an additional group or service level property. Reports can then use the new properties to categorize the objects.

Groups are created by tagging hosts, hypervisors, or virtual machines with a “nodegrp” property. You can specify hosts by name or regular expression. The host-type collectors (Physical Host, VMware) include group tagging filters. In fact, the other collectors that enable the use-case reports also include group tagging filters, but their groups will not appear in the default chargeback reports. If you tag other objects like arrays and switches with groups, you will have to create a completely customized report to display your results.

Service levels are created by tagging LUNs with a “svclevel” property. The array-type collectors (VMAX and VNX) include service level tagging filters that characterize LUNs based on their disk technology, disk size, disk speed, RAID type, and whether they are FAST Cache enabled. A separate filter can also characterize FAST VP LUNs based on the policy used.

Business (unit), Customer and Application ownership in another category of data tagging based on device name.

User Defined Device Grouping

Welcome admin (logged as admin) C.M. - EMC M&R Monday, July 21, 2014

ONLINE UPDATE SOLUTIONPACK CENTER USER SETTINGS CONFIGURATION DOCUMENTATION

Your Company Logo

- Physical Overview
- Logical Overview
- Packages Management
- Licenses Management
- Audit Logs
- Discovery Center
- Central Configuration Repository
- SolutionPacks
- Data Enrichment
- Groups Management**
 - Business, Customer & Application own
 - Device Grouping**
 - Service Level by LUNs
- Scheduled Tasks

Device Groups

Create groups from devices.

Create Edit Delete Cut Copy Paste

All

- All Hypervisors**
- All Passive Hosts
- All Physical Hosts
- All VM Guests
- EMC

This table is only a preview. Use the search for more detailed results. Note that the search feature might take up to several minutes depending on how much data you have.

Show 10 entries Search: (Wildcard)

Device name	Device type	Device model
10.7.23.17	Hypervisor	PowerEdge R610
10.7.23.19	Hypervisor	PowerEdge R610
10.7.23.2	Hypervisor	PowerEdge R610
10.7.23.4	Hypervisor	PowerEdge R610
losar202.iss.emc.com	Hypervisor	PowerEdge R710
losar203.iss.emc.com	Hypervisor	PowerEdge R710

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. ViPR SRM Overview 62

Device Groups capability allows you to take discovered (including passively discovered) devices and place them and organize them into a collection. As an example, in configuring for VMware, one can select all hosts, hypervisors, virtual machines, and passive hosts, and name them. To make your own groups, select Create. Choose to define membership dynamically or from a static list. Dynamic selection allows the user to specify criteria for selection, while Static selection allows that user to specify an object from the list of discovered objects. For Dynamic, the user can create rules based on: device name, type, model, or any property value. For static the user can choose to add a discovered device to the group.

Once created, groups can be edited and deleted and participate in cut, copy, and paste operations.

User Defined Service Levels

- Physical Overview
- Logical Overview
- Packages Management
- Licenses Management
- Audit Logs
- Discovery Center
- Central Configuration Repository
- SolutionPacks
- Data Enrichment
- Groups Management
 - Business, Customer & Application owner
 - Device Grouping
 - Service Level by LUNs**
- Scheduled Tasks

[ONLINE UPDATE](#)
[SOLUTIONPACK CENTER](#)
[USER SETTINGS](#)
[CONFIGURATION](#)

Service Level

Service Level rules are applied in the order shown. The first matching service level is assigned.

[Edit](#)
[Reorder](#)
[View All Service Level Rules](#)

System Resource
Pool Contributor
FAST VP
Platinum
Gold
Silver
Bronze
Custom1
Custom2
Custom3
Other

This table is only a preview. Use the search for more detailed results. Note that the search may take several minutes depending on how much data you have.

Show entries Search: (Wildcard)

Device name	Device type	Device model	RAID Type	Disk Technology	Disk Size	Component type	Component name
888973588710	Array	VMAX40K	RAID-1	SAS	2795	LUN	0051
888973588710	Array	VMAX40K	RAID-1	SAS	2795	LUN	0052
888973588710	Array	VMAX40K	RAID-1	SAS	2795	LUN	0053
888973588710	Array	VMAX40K	RAID-1	SAS	2795	LUN	0054

Copyright © 2015 EMC Corporation. All Rights Reserved. VIPI SRM Overview 63

Service Level rule sets are meant to categorize LUN objects based on some of the disk characteristics that define them.

“System Resource” and “Pool Contributor” have greater precedence than those below. Each metric goes through this rule set and the first matching service level is the one that the LUN will be assigned to.

With the exception of System Resource, Pool Contributor, and Other, the service levels can be reordered. This allows the user to specify the order of service levels in which criteria are tested.

“View All Service Level Rules” allows the viewing of all service levels at the same time. This makes it easier to determine which service level applies to a LUN.

VMAX – Rules for Service Levels

Disk Technology Service Levels

Name	RAID Type	Disk Technology	Disk Size	Disk Speed
Platinum	*	Flash Drive, Enterprise Flash Drive	*	*
Gold	RAID-1	Fibre Channel	*	*
Silver	RAID-5	Fibre Channel	*	*
Bronze	*	SATA	*	*

Other Pre-defined Service Levels

Name	Criteria
System Resource	purpose=='System Resource'
Pool Contributor	purpose=='Pool Contributor'
FAST VP	FAST managed LUNs not tagged with a FAST service level
Other	Thick LUNs or Non-FAST thin LUNs not already tagged with a service level



The above describes managing VMAX service levels. There are four default service levels - Platinum, Gold, Silver, and Bronze. You can see the criteria used to determine the service levels in this table. The criteria depends entirely on Disk Technology, except for Fiber Channel LUNs, where the RAID Type is considered.

There are four other default service levels. System Resource and Pool Contributor service levels are tagged for LUNs that have a 'purpose' attribute of that name. The collector tags the LUN with the name during discovery. LUNs are tagged with the FAST VP service level if they are marked as FAST VP LUNs during collection, and they are not otherwise tagged with a FAST VP service level. In other words, all FAST VP LUNs are added to the generic FAST VP service level until an administrator configures more specific FAST VP filters.

Finally, the other service level is tagged for any LUNs not otherwise categorized.

User Defined Business, Customer, and Application Ownership

Groups Management

- Business, Customer & Application ownership
- Device Grouping
- Service Level by LUNs

Device Groups

Create groups from devices.

Create Edit Delete Cut Copy Paste

All Engineering Human Resources IT Manufacturing Sales

Tag properties

Tag values

Ed Device Group

Business Unit

customer

appli

owner

Engineering

VMware

vCloud

VCloud_Dev


Device type	Device name	Device description	Component type	Component name
Array	888973588727	5876.268.174 date:02.05.2014 patch_date:02.05.2014	Access	BM_Iglov037_1_36A6
Array	888973588727	5876.268.174 date:02.05.2014 patch_date:02.05.2014	Access	BM_Iglov037_2_36A7
Array	888973588727	5876.268.174 date:02.05.2014 patch_date:02.05.2014	Access	BM_Iglov038_1_7F60
Array	888973588727	5876.268.174 date:02.05.2014	Access	BM_Iglov038_2_7F61

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. VIPR SRM Overview 65

The third section of Groups Management that allows you to set the rules for four more properties: business unit, customer, application (appli), and owner. You can see in the above example that five rules are defined: Engineering, Human Resources, IT, Manufacturing, and Sales.

For any metric that meets a rule (filter rule set inside an entry), the tagged properties and values will be added to the property set.



If cases supported by Groups Management feature tool cannot meet your specific case of data enrichment, you can still use the individual data enrichment rule set of a specific SolutionPack. In that mode, as a ViPR SRM Administrator, you will have to register the module before administering the data enrichment rule set.

Upon initial installation of SRM, no PTF module is registered. You first need to register the module of an existing SolutionPack in order to edit the configuration, thus the rule set.

1. The **Register a new module** button is the indicated action to take for that step.

The new popup window presents selectable items from left to right:

- **Servers:** Select which host of your SRM platform you want to select Categories from.
- **Categories:** Select the category (the type) of collector. There are basically two types - Collecting and Events.
- **Modules:** In the list of modules, select the component to register. The format of the module name is:

Collector-Manager::[InstanceName]::[Property-Tagging-Filter Name]

InstanceName is the unique name of the collector. Most of the time it is related to a SolutionPack's name.

Property-Tagging-Filter Name is the unique component name, inside the collection layer. You can have multiple Property-Tagging-Filters inside of a single collector.

2. Some SolutionPacks are meant for DataEnrichment customization. For VMware Group management, you can use User_Defined_Group_Management. For even more flexibility, use the Property-Tagging-Filter with the name **"DataEnrichment"**, unless documented otherwise. If you use others, you might disorder the collection and reports may not show as expected.
3. Once the component (module) selected, click the "Register" button in order to add the selected to the list of manageable DataEnrichment components.

Demo – Group Management

This demo video provides details on how to perform management of the data enrichment feature.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

ViPR SRM Overview

67

This short video shows how to manage device grouping, Service Level (for LUN), Customer, Business Units, Application, and Owner tagging. This feature is key to Reporting in ViPR SRM as data must be tagged before being organized in custom grouping reports.

Module: Overview of the ViPR SRM User Interface

Lesson Summary

During this lesson the following topics were covered:

- The value and features associated with data enrichment



This lesson covered data enrichment feature and how it integrates in reporting.

Module Summary

Key points covered in this module:

- Described the function of ViPR SRM in the creation of the Software Defined Data Center (SDDC)
- Explained how reporting is an integral function of managing the SDDC
- Described the tools and features available through the portal to run, modify, and create reports useful in managing the SDDC
- Defined uses for and examples of enhanced reporting
- Discussed how data enrichment can be used to enhance collected data
- Described Group Management and Data-Enrichment feature tools



This module covered the function of ViPR SRM in the creation of the Software Defined Data Center (SDDC). We explained how reporting is an integral function of managing the SDDC and described the tools and features available through the portal to run, modify, and create reports useful in managing the SDDC. Also covered were examples of enhanced reporting, how data enrichment can be used to enhance collected data, and Group Management and Data-Enrichment feature tools.

This slide is intentionally left blank.



Module: Basics of Reporting

Upon completion of this module, you should be able to:

- Discuss the concepts important to good report design
- Describe how the reporting engine works
- Explain the base organization of data as it is used by ViPR SRM
- Demonstrate the functionality of SRM used to browse the data set

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

1

This module focuses on the basics of reporting in ViPR SRM with concentration on understanding how the data collected is represented to the report engine. Emphasis is given to understanding how to use the data set to modify or create reports unique to the customer's needs.

Module: Basics of Reporting

Lesson: Creating Reports and Report Design

This lesson covers the following topics:

- Expected outcomes and report design
- New reports or repurposing existing reports
- Arranging a draft report
- The iterative approach to report creation

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

2

This lesson covers creating reports, specifically the expected outcomes and how they affect report design. Discussions include deciding whether to repurpose existing reports by adding or removing content or creating a new report. The use of the Report Wizard to begin the creation of a report and the iterative process of report creation are also highlighted.

Expected Outcomes

- Where do expectations come from?
 - ▶ Business requirements
 - ▶ Business strategy data alignment
 - ▶ Ad-hoc on demand requirements
 - ▶ Introduction of a new component
 - ▶ Investigation of component services
 - ▶ Analytics
- What are some common expectations?
 - ▶ Reuse and tailor standard reports
 - ▶ Modify time scale
 - ▶ Adjusting thresholds
 - ▶ Employ benefits from data enrichment



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

3

This first lesson will take some time to help you design reports appropriately. The above is meant to provide you with some advice/indications for report design, it is not meant to be a single reference.

Let's first talk about the requirements for a report... where do expectations come from? Who's asking what?

Business requirements: The business you are in every day requires information from various sources. ViPR SRM can adapt in so many ways such as:

- purpose of the report
- structure of required reports (drill down)
- metrics used (scope)
- format (table, graph, or other)
- calculations of KPIs
- Data aggregation: level of details, granularity

Business strategy data alignment: Reporting with ViPR SRM has infinite possibilities. Creating a new report structure (base or through SolutionPack) is completely free to you. You may want to design your report structure with respect to your business strategy. As an example, for a Service Provider case, a separate report structure would be defined and made available to your customers. The report structure is a template and data applied to it depends on the customer logging in (and tagged dataset). This customer view will obviously be different from yours, as a global administrator. As an example, a customer report would list pools and their used and free capacity only for the service they subscribed. As a global service provider, you would add numbers of events, per severity. This simple case requires you to define two separate reports, but reuse one or the other.

Expected Outcomes – (cont.)

- Where do expectations come from?
 - ▶ Business requirements
 - ▶ Business strategy data alignment
 - ▶ Ad-hoc on demand requirements
 - ▶ Introduction of a new component
 - ▶ Investigation of component services
 - ▶ Analytics
- What are some common expectations?
 - ▶ Reuse and tailor standard reports
 - ▶ Modify time scale
 - ▶ Adjusting thresholds
 - ▶ Employ benefits from data enrichment



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

4

Ad-hoc on demand: a very specific report that needs follow up after some service. You might want to follow up on components behavior. The characteristic of such an expectation is that its scope is well known and reports should be fast to define.

Introduction of new component or service: studying the effect of a change in the IT infrastructure - closely looking at very specific performance of a component. The characteristic of such an expectation is that it might re-use some reports while focused on more specific component(s).

Investigation of components services: a close follow up on some critical components. The provided, OOTB reports (or from SolutionPack) are built for generic and global. You might want to have your own specific report, looking more closely to some chosen components - easier and faster to read/understand.

Analytics: more advanced reporting needs focusing on specific components (by criticality). Could be used for planning and engineering the IT infrastructure. It could also be justified by way of doing in the business you are in every day. You might want to use some advanced math and computation over components data in order to plan or to provide recommendation.

Because report definition in ViPR SRM is always a template (report context), you will want to re-use your work as much as possible. Over time, you will save time when you base your report on an already existing one. All you need to do is alter/tailor the initial report. Sometimes base reports use data covering the last 24 hours. Most users reading a report will notice and will ask for a longer term (horizon of the report). Although this is a common change that can be done by the user of the report, more permanent views can be derived from the report's definition. OOTB and provided SolutionPack reports are defined with comprehensive, yet generic or best practice values of threshold. It is normal for you to change these definition and set them according to your standards.

Data enrichment: OOTB and provided SP's are developed to take advantage of some data enrichment but cannot plan for all possibilities and specifics alignment with your business organization. In order to benefit from the work of data enrichment in such reports, you will have to revisit report definition in order to augment the level of details that re-define the break down.

Target Audience

Management

Accounting

Customer

Operations

Engineering

Usable Capacity by Service Level

5 elements found, displaying all elements.

Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used
Silver	1,363.28	93.90	1,269.38	325.60	6.89
Bronze	1,012.08	962.18	49.90	962.18	95.07
FAST VP	25.31	25.31		25.29	100.00
Other	0.00		0.00		
Gold	5,396.39	21.46	5,374.93	13.97	0.40
	7,797.06	1,102.85	6,694.21		

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

5

Who's the audience of your report? This crucial question will allow you to determine the level of details of your report. The above shows various audience's for reports. Here are some defined audience's, as an example:

- **Operations:** looks at current issues. It could be based on a customer problem.
- **Engineering:** looks for trends in capacity to plan for growth.
- **Management:** interested in global views, dashboards (no tiny details) and making sure SLAs are met. This will allow them to take action accordingly.
- **Accounting:** precision, reproducible, and exact are key factors for billing. If a customer/support wants to return back in time, data must be accessible, reports must be reproducible and exact.
- **Customers:** in a Service Provider mode, organizations may want to allow their customers to get their own service reports, instead of having the organization pushing the reports out (monthly, weekly basis, email format, etc). The same rules apply from a customer's point-of-view. A Service Provider needs to make sure the customer sees only their reports without disclosing the data of other customers.

ViPR SRM displays personalized, actionable, interactive information to the appropriate audience within the organization (and beyond).

ViPR SRM Reports

What is a report in ViPR SRM?

- Every report page is defined by a report template
- Snapshot in time
- Comprehensive format
- Selects a set of metrics and KPIs
- Static or dynamic



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

6

There are two modes of operation with ViPR SRM reporting: browse and edit modes.

-Browse mode views what has been defined in edit mode. Edit mode defines various parameters. The amount of parameters used to define a report, highly depends on: format of report, scope of metrics, and some specific parameters.

- Any report with ViPR SRM is dynamically generated when asked for. Either you browse (live mode) reports using the report navigation, or when you schedule the report. At this moment, a time range is determined and the metrics data are recovered from the database in order to let the report generator build the asked for report.

- A report must have certain format: the base formats are table and chart (graph). ViPR SRM supports many other types of reports. Your report goal should make use of the most effective format in order to be comprehensive as much as possible and self explanatory.

- It is important to scope your report: the list of metrics available for reports is always the full content of the database. You always have to first set a filter to limit the scope of metric to what is required for your report purpose. This is what we will later call "metric context".

- A report can be set to be static or dynamic: when static, we mean that the content does not change over time and does not depend on the component collection to be shown. This type is useful for report structure organization. Presented in list or items of a table, they are very handy. By default, when adding a new report in the navigation tree, the static format is selected. Dynamic reports depend on context. The content depends on components of the collection. Obviously an inventory-like report is dynamic as its content depends on the components available in the database.

In a simple format, there are report types that are meant to be a report of reports. Either this type of report is used describe a drill down structure (selecting and refining elements) or to aggregate multiple smaller report of specific purpose.

Using Existing Reports

- Existing reports may be used as a starting point:
 - ▶ Remove one or more columns from a report
 - ▶ Adjust thresholds
 - ▶ Modify column names or the labels of graph axes
 - ▶ Calculate an average, showing peak/minimum value or a total
 - ▶ Change order of appearance of elements
 - ▶ Re-ordering rows, child items
- Always work from a copy of an existing report template (never from the original)
- Use your personal report sandbox for development purposes



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

7

Report creation is a process, we'll touch base on this shortly. There are many benefits from re-purposing existing reports. The OOTB report, or those provided after the deployment of a SolutionPack, can be used as a starting point to create unique and personalized report that will better meet the expectations. Some of the advantages are listed here.

In any case, the existing report template should be copied and pasted into your personal report sandbox (ie: "My Reports") where it can be modified to meet specified requirements. Later you will be able to put this report definition into a global report pack in order to "publicize" your work.

Laying Out a Basic Report

1. Identify the goal of the report and target audience
2. Identify the report type that fits the goal
3. Scope the report (filtering by objects)
4. Define the report levels, generic to specific
5. Review and refine, add small details and description

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

8

Main goal of the report: When creating a report, make sure to identify the target audience. You might also want to see if this requirement comes from an already existing report - ask to have an example/sample of this source report.

Find the best format of the report: Would the information be presented in a table (inventory or list like) or would it be simply put in a chart (graph)? Sometime aggregation of data is required and is presented by type, but ViPR SRM supports this easily. It is really the responsibility of the report designer to align the goals with the proper format in order to give meaning to a given report.

The scope of the report: Metric scoping is key to reporting - either for direct use in report or for KPI calculation (2nd processing). But more important: do you have this kind of metric data in the dataset? Some investigation might be required. Also consider that the report engine has access to the full content of the database: timeseries and events. Mixing all those metrics from all SolutionPack does not make sense without some structure and organization. It is then important to scope the metrics selections by defining a filter at the report node definition. The report generator will then take care of the rest.

Report levels: As you would create a file directory structure on your local computer in order to retrieve files, you would do the same for report. The report navigation will show your report structure as you define it. Usually, the top level (close to root level) is used to summarize and list available components from the database. As you drill down, you refine the metric selection (thanks to compound filter feature of the tool), thus refining your report. This is often referred to from generic to specific. The generic form shows information of high level: global capacity per array down to finer details.

Review and Refine: When a goal is reached with the core report, review and add some level of details like properties, proper time range selection, report description, etc.

Laying Out a Basic Report – (cont.)

Report edition => editing a template:

1. Report node behavior (dynamic vs static)
2. Report format (table, graph, etc)
3. Report relative time range (report's horizon)
4. Report layout
5. Metric filter, KPI calculation
6. ... and other specifics parameters by context!

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

9

When ready to move on, keep in mind that when you use the report edition mode, you are working out a report template that will be in use at browse time. Many report aspects will be required:

- Is it a dynamic or a static report? Depending on the context and the level details, this report could be targeted. In reverse, it could report on all objects, no matter which device they belong to.
- Report format that speaks to your report audience. What level of details should be presented at each layer/level?
- Report time range: What should a given report cover? Last 6 hours, 24h, week, month, semester, year?
- Report layout: Multiple or single reports.
- Metric context: Uses filtering in order to select/limit the metric scope that will be used by the report engine in computing the report.
- By the type of report you select, you will have specific details to provide and options to select.

Report Creation

- Successful report creation:
 - ▶ Define the goal of your report
 - ▶ Scope your report in term of metrics
 - ▶ Define the best way to expose the result (report format)
 - ▶ Define the drill down
- The report engine processes in order:
 - ▶ Metrics selection by context (apply filter)
 - ▶ Retrieve the data set (time series / events)
 - ▶ Build up the report (table, graph, aggregation, etc)


EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

10

Before you start selecting report nodes and choosing report types, defining filter, etc, you should consider how the report engine works. The above exposes some key points in three simple steps. Refer to next page in order to view those steps in a chart.




What happens when you select a report from report list/inventory?

At the heart of the process, there's the Report Engine (RE). Based on report context (selected report), RE will first query database for applicable metric, based on filter expression (combined and based on context from the report structure). Once the list of metric is known, RE will fetch data values (timeseries) for each metric and depending on the report format asked for. It will then adjust based on some properties of the report before outputting the final report.

In some cases, a report is a mix of multiple smaller reports - the process is just repeated for each smaller report and a layout in a larger view.

The key point to understand is the metric selection based on properties is always done before applying the timeseries data. Thus it is important to set the metric context correctly when defining a report. Otherwise, RE will consume tons of data from database, will consume a rather large amount of processing time, and generate a questionable result.



A report is selected from the navigation tree. Even if a report shows static entries, these are considered a report. The above shows a static report named “Dashboards”. This report shows “Storage” report, which is a list of 4 reports. You must select a report from the list in order to see a report from the list. You can also choose to unfold the report list by clicking the small triangle icon next to the report’s name. In the above, up to “Enterprise Capacity Dashboard” has been unfolded.

The report layout shows many items such as:

- **Breadcrumb:** Follows the report list as you drill down.
- **[Parent/current] report names:** The last report name and the current report name are shown on this line. It helps focus on the report context. Multiple report levels are involved in producing a single report.
- **Report date and coverage:** the exact report date and time for the data selection that are covered by the current. Take note that under some circumstances, this could be overridden. Coverage describes the printed date/time range (1 day, 1 week, 1 month, 1 quarter, etc).
- **Aggregation level:** A report always selects metric data from an aggregation level. Four levels of aggregation are normally available. Like the above, if not shown, it means report engine makes use of real-time data aggregation.
- **Report processing time:** Is always printed at the bottom of any report. It helps determine how long it took to generate a given report. If a report takes long to generate, you might want to schedule or pre-generate the report (other features).
- **Report description:** This section is optional and set by the report designer. It is of good practice provide a detailed description of the current report and, perhaps, the limitations (included/not included). The report reader can then better understand the areas of the report.

Components of a Report

Report Library / EMC VMAX

July 2014, Tuesday 8 » Wednesday 9, 8:08 AM EDT | Last 1 Day: average on 1 hour

Pool Capacity Details

15 elements found, displaying all elements.

Show ALL ▾

Array ▾	Pool Name ▾	Pool State ▾	Used ▾	Free ▾
000195901097	SRM-Pool	Enabled	18.000 MB	8.925 TB
000192606327	DEFAULT_POOL	Enabled	0.000 GB	94.757 GB
000192606304	DEFAULT_POOL	Enabled	4.211 GB	90.545 GB
000192606327	BC_SATA_R6	Enabled	0.000 GB	50.520 GB

Metric properties

Metrics data
(summary of
timeseries)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

13

Let's start our study of the basics of reporting by looking at some components of a report. Above shows two examples of reports: Pool Capacity Details and Pools Capacity Summary. First is the time range where the report time range period covered. It is always important to relate this information with the presented results. In the above, the report covers the last 24 hours (1 day) from the date and time you clicked for the report. It also reveals use of the 1-hour average of data collection (time-based aggregation), often referred as the 1-hour aggregate timeseries.

In the first case, the organization of the information is presented in a table by pool name, per array, the status, and Used and Free capacity. These numbers are metrics from the Pool Name. Metrics is at the base of reporting.

The second case (IOPS) shows the behavior of IOPS of a given array over time. The report designer shows both metrics stacked one over the other (thus Read + Write Requests). In the case of a graph as above, limited details are available (in term of properties). Only metric names are available in the legend. The remaining details should tell by the selected report context.

Module: Basics of Reporting

Lesson Summary

During this lesson the following topics were covered:

- Expected outcomes and report design
- New reports or repurposing existing reports
- Arranging a draft report
- The iterative approach to report creation



This lesson covered the elements of basic reporting including metrics and properties, how they are used in both time series and event data collection and finally how data is aggregated to satisfy many different types of reporting needs.

Module: Basics of Reporting

Lesson: Base Organization of Data

This lesson covers the following topics:

- Metrics and properties
- Timeseries and event data
- Aggregation of data



This lesson covers base organization of data including discussion of metrics and properties, comparison of timeseries and event based data, and how data is aggregated.

Metric Defined

Definition

A metric is the smallest unit of data you can regularly poll from a given instance out of a given object, device, service, or system.

Examples:

- From an array, there are disks. A given collector is configured to retrieve data units from each single disk reported by the array's interface:
 - ▶ ReadOperations
 - ▶ WriteOperations
 - ▶ [...]
- From an array, there are LUNs. A given collector is configured to retrieve data units from each single LUNs found/reported by the array's interface:
 - ▶ ReadOperations
 - ▶ WriteOperations
 - ▶ [...]




The term “metric” is the base item of the reporting concept and is used in each single module of the ViPR SRM platform. So, what is a metric? It is the smallest data unit that is extracted from a given object of a targeted system, sometimes compared to Key Performance Indicator (KPI). The metric is more of a raw data item than KPI - it represents / identifies the raw value extracted from a given object of a targeted system. In ViPR SRM context, a KPI is the result calculated from a given amount of, yet aggregated, metrics. SRM collects and stores metrics that the Report Generator uses to generate reports. The values in the reports are KPIs, they are the final result that can be compared to evaluate, and make the context of the report understandable. You can also consider some raw metrics to be a KPI also.

From the above, we have two examples: one is from Disk context and the second from LUN context. They are using the same metric name set but that does not mean they are the same... the context become key in here. “Disk” and “LUN” are object types. You can have multiple object instances of each type (Disk number, LUN name).

The array's interface refers to the way to get the metrics from the array. As an example, with EMC SAN products, SRM gets VMAX metrics through Solution Enabler and Unisphere.

What is a Metric?



An Array is a physical entity that encompasses a number of objects of different types. Depicted above, we can identify: Ports, LUNs, Raid Groups (RG), Storage Pools (SP), and Disks (to just name a few). With SRM the common property name for such object types is defined as the property named “parttype”. So we have parttype=“Disk” or parttype=“LUN” would identify a piece of a metric’s context.

For each type of object, we can have one or many “instances”: there are 6 ports, there are 2 LUNs, 2 RGs, 2 SPs, 11 Disks. Each of those objects can be uniquely identified by a number or a location (Disk’s case) or a label. With SRM the common property used to designate an instance of an object is the “part” property name.

For each object we can have one or many metrics. As depicted above, the same metric name is used for many object instances of the same object type. With SRM, the common property name for a metric name is “name” (simply!).

In order to have a more complete picture, a property named “device” is added to each single metric of the above in order to better reference the group. This will be used to differentiate metric set of an array device name from another array device name. **Note:** the serial number is used as the device name in identifying the array.

With the same concept as the previous, a property named “devtype” is added to each single metric of the above in order to differentiate with the metric set from, as an example, a SAN Switch or from any other type of device in the infrastructure (host, application, vCenter, etc).

This way, you can have multiple arrays, types of arrays, mixed-environment of data source: hosts, VM hosts, SAN Switches, Routers, Firewalls, Network switches, etc.

Property Defined

Definition

A property is meta data that is associated with and provides context for a metric.

Examples continued:

- From the same array, where disk metrics would be ReadOperations, WriteOperations, etc., s. Collection layer assigns properties that describes the data units for each single disk, such as:
 - ▶ Disk number
 - ▶ Array's SN
 - ▶ Disk type
 - ▶ [...]
- From the previous array, there are also LUNs metrics ReadOperation, WriteOperations, etc. Collection layer assigns properties that describes the data units for each single LUN, such as:
 - ▶ LUN Name
 - ▶ UUID
 - ▶ Storage Pool
 - ▶ [...]



Any system could end up having several metrics (ie. per component type). Since the metric is the smallest data unit that is extracted from a given component of a targeted system, a metric must always be described by its context. This context is described by the properties associated with it. In order to be meaningful, **a metric must have properties**. That's the foundation of ViPR SRM for collection and, most important, for reporting.

Simply said, properties in report context will allow you to specify which metrics you are interested in working with for a given report. As an example, if you want to graph IOPS over time, you need to select (thus filter out) the pertinent metrics.

In these two examples shown above, the amount of properties assigned are more than shown. As an example, in VMAX context, a Disk is described by about 40 properties while LUN is described by more than 64 properties.

Timeseries Defined

Definition

A timeseries belongs to a metric. It is the time and value data a collector regularly retrieves from a target device/object. Timeseries data is always stored in database.

Examples:

From an array, there are disks. A given collector is configured to retrieve data units from each single disk reported by the array's interface, as such: Read and Write Requests. Data is stored as follow in the database:


Timestamp	WriteRequests (IOPS)	Timestamp	ReadRequests (IOPS)
7/21/14 16:02:30	12.14669991	7/21/14 16:02:30	2.673330069
7/21/14 16:07:30	48.22999954	7/21/14 16:07:30	6.083330154
7/21/14 16:12:30	9.033329964	7/21/14 16:12:30	1.763329983
7/21/14 16:17:30	8.483329773	7/21/14 16:17:30	2.043329954
7/21/14 16:22:30	8.350000381	7/21/14 16:22:30	1.46333003
7/21/14 16:27:30	9.583330154	7/21/14 16:27:30	1.653329968
7/21/14 16:32:30	8.676670074	7/21/14 16:32:30	1.436669946
7/21/14 16:37:30	11.7833004	7/21/14 16:37:30	2.20333004
7/21/14 16:42:30	8.996669769	7/21/14 16:42:30	1.379999995
[...]	[...]	[...]	[...]



Timeseries belongs to a metric. This is the time value data that a given collector (setup by a SolutionPack) recovers from a target device's object. The time and value data is processed by a backend before being stored in database and referenced by the metric.

Common Convention

One single raw data item generated by a collector:



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

20

The fundamental element of the data model is the raw value, which represents a floating point value, at a specific date and time for a specific metric. A raw value must be joined with meta-data about the metric. These meta-data are called **properties**.

As from the above “property card”, we can easily identify the origin of the raw value, in this case describing the number of ReadRequests: it contains the name of the device the value comes from along with the type of object (parttype), the object instance name (part), and along with many other properties. This is a simple sample of a collection. Let’s then think of a given SRM platform is collecting 1000’s of 1000’s of metrics from the infrastructure (a small deployment). When looking at the report, anyone should be able to relate back to the infrastructure.

Internally, SRM has a flat data model: it does not have any knowledge about the relationship between metrics. Only the metric’ properties can help in creating the organization of the data.

As an example: successively grouping by devtype, then by device, then by parttype, then by part would explore in a logical manner the data collection. At first, it may appear as a drawback but actually, it is a major advantage since there is no implicit hierarchy between data properties. It doesn’t make any assumptions like “a device contains ports which contains several indicators” and being limited to this schema. The suite can collect data from heterogeneous sources and still organize them in a convenient way in the Frontend, at report time. Furthermore, thanks to this, users can define their own hierarchical data model, according to their needs.


The format of such properties is specific: on the left-hand side, there are “property name” and on the right-hand side we have the “property value”. This forms a pair.

The “[...]” sign in the above means that the number of property name and value pairs are unlimited in number to describe the metric.

The “variableid” depicted is a text string that uniquely identifies a metric. It is generally composed of a selection of multiple property values.

The “group” property is a mandatory field set at collection configuration and is used by APG-Backend module. It is mainly used to give data a retention scheme. This detail is covered in an advanced course of SRM Suite.

Common Property Convention



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Basics of Reporting

21

From the previous metric concept, we can start to see that the model is unlimited and can be used to describe any object (type, size, flavor, etc) that has metrics. More, we can even better describe metrics by adding more properties... and even create your own data property name and definition.

This puts flexibility in the model. Although there is a lot of flexibility, SRM brings the "Common Convention" concept in order to set a base common property naming convention. The concept is to make sure any collection that is developed can re-use some of the generic report layouts provided in a base report installation. The common convention defines the following properties to be as strongly suggested in each collection. The most conventional are depicted in the above slide, along with their generic definition.

Expanding Common Convention



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

22

Depending on the SolutionPack applied and the type of component collected, the number and type of property changes. The base concept is the Common Convention. The concept is expanded, meaning the number of properties increases by context. When a SolutionPack is created, engineering defines properly the associated properties for any metric collected.

Because ViPR SRM uses a flexible flat data model, the number of properties will vary depending on the collection context (metrics). This is what gives ViPR SRM its unique flexibility to provide a unified view across many different technologies & vendors. A metric may have one or more properties attached than others metrics. Not all metrics require the same set of attached properties.

From the above, is an extension of the metric property set brought about by Storage application.


Note: Not all objects from Storage make use of the above extensions. Usage of the extensions is dependent upon the context of collection (component type, type of array, etc). This is just an example of expanding the common convention showing how flexible the model is.

Data Enrichment, another flexibility in ViPR SRM, allows you to further augment/complement the data collection. This is a feature of administration of the platform. Adding such information to metrics will allow for more flexibility (and customization) when managing/creating reports.

Values and Properties

What is collected?

Timestamp:1205848501, Value:64	Source=XMLCollector devtype=Array device=000194900149 parttype=LUN part=02A0 unit=Nb/s name=ReadRequests ...
Timestamp:1205848501, Value:65	Source=XMLCollector devtype=Array device=000194900149 parttype=LUN part=02A0 unit=Nb/s name=WriteRequests ...
Timestamp:1205848501, Value:66	Source=INCHARGE-AM-PM devtype=Host device=W4N-MTL-MAIL01 parttype=FileSystem part=C:\ unit=% name=UtilizationPct ...
Timestamp:1205848501, Value:64	Source=INCHARGE-AM-PM devtype=Host device=W4N-MTL-MAIL01 parttype=Memory part=Physical Memory unit=% name=UtilizationPct ...



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

23

Now let's see how this works in the implementation. The above figure shows a collection of four raw data metrics migrating from Collector toward the APG-Backend. Notice how each single metric is a packet record in itself and travels between the Collection toward the APG-Backend. A packet record contains the complete property set along with data value and timestamp. This allows APG-Backend to process data correctly, deciding if updates in the database are required or not (saving on costly data storage operations) and performing aggregation based on the information that join the value of the metric. It is then important to understand that when leaving the Collector, the data is "finalized", meaning it has its complete set of properties. It won't be possible to alter the content of the database or at the APG-Backend level.

Here are the details of a data packet record:

- A timestamp, that is, the moment in time when the value was collected.
 - ▶ Collection uses UNIX timestamp; or the amount of seconds since January 1st, 1970, at midnight GMT.
- A value, collected at a given timestamp.
 - ▶ Values can be displayed onto graphs, averaged, etc
- Several properties, used to better identify the metric.

Sample Metrics (4) Property Set

VMAX Disk metric	VNX Storage Pool	VMAX LUN metric (1)	VMAX LUN metric (2)																																																																																																																																																																																
Available properties	Available properties	Available properties	Available properties																																																																																																																																																																																
<table><tr><th>Name</th><th>Value</th></tr><tr><td>model</td><td>VMAX-1SE</td></tr><tr><td>datagr</td><td>SYMMETRIX-DISKS</td></tr><tr><td>memberid</td><td>7</td></tr><tr><td>dname</td><td>DISK_GROUP_001</td></tr><tr><td>memberof</td><td>DF-7A</td></tr><tr><td>diskcost</td><td>0.280</td></tr><tr><td>apgdb</td><td>APG-DB</td></tr><tr><td>partmode</td><td>Assigned</td></tr><tr><td>name</td><td>WriteRequests</td></tr><tr><td>apgrp</td><td>group</td></tr><tr><td>devtype</td><td>Array</td></tr><tr><td>firmware</td><td>5874.248.194</td></tr><tr><td>diskgrp</td><td>1</td></tr><tr><td>hyperid</td><td>49</td></tr><tr><td>lunname</td><td>06A9</td></tr><tr><td>raidgrp</td><td>53</td></tr><tr><td>vendor</td><td>EMC</td></tr><tr><td>ip</td><td>000194900149</td></tr><tr><td>unit</td><td>Nb/s</td></tr><tr><td>disktype</td><td>SATA</td></tr><tr><td>source</td><td>XmlCollector</td></tr><tr><td>parttype</td><td>Disk</td></tr><tr><td>partdesc</td><td>SATAHGST</td></tr><tr><td>device</td><td>000194900149</td></tr><tr><td>diskrpm</td><td>7200</td></tr><tr><td>part</td><td>07A:C0</td></tr><tr><td>partvrs</td><td>N02THJK</td></tr></table>	Name	Value	model	VMAX-1SE	datagr	SYMMETRIX-DISKS	memberid	7	dname	DISK_GROUP_001	memberof	DF-7A	diskcost	0.280	apgdb	APG-DB	partmode	Assigned	name	WriteRequests	apgrp	group	devtype	Array	firmware	5874.248.194	diskgrp	1	hyperid	49	lunname	06A9	raidgrp	53	vendor	EMC	ip	000194900149	unit	Nb/s	disktype	SATA	source	XmlCollector	parttype	Disk	partdesc	SATAHGST	device	000194900149	diskrpm	7200	part	07A:C0	partvrs	N02THJK	<table><tr><th>Name</th><th>Value</th></tr><tr><td>model</td><td>VNX7500</td></tr><tr><td>module</td><td>SAS</td></tr><tr><td>vols</td><td>18</td></tr><tr><td>serialnb</td><td>APM00121301071</td></tr><tr><td>datagr</td><td>VNXBlock-StoragePool</td></tr><tr><td>maxdisk</td><td>1000</td></tr><tr><td>apgdb</td><td>APG-DB</td></tr><tr><td>ip</td><td>Iglad039</td></tr><tr><td>unit</td><td>GB</td></tr><tr><td>revision</td><td>05.31.000.5.709</td></tr><tr><td>source</td><td>VNXBlock-Collector</td></tr><tr><td>parttype</td><td>Storage Pool</td></tr><tr><td>dgraid</td><td>RAD5</td></tr><tr><td>name</td><td>UserCapacity</td></tr><tr><td>device</td><td>vrnx1071</td></tr><tr><td>apgrp</td><td>group</td></tr><tr><td>maxhost</td><td>8192</td></tr><tr><td>devtype</td><td>Array</td></tr><tr><td>part</td><td>Cabinet Pool 1</td></tr></table>	Name	Value	model	VNX7500	module	SAS	vols	18	serialnb	APM00121301071	datagr	VNXBlock-StoragePool	maxdisk	1000	apgdb	APG-DB	ip	Iglad039	unit	GB	revision	05.31.000.5.709	source	VNXBlock-Collector	parttype	Storage Pool	dgraid	RAD5	name	UserCapacity	device	vrnx1071	apgrp	group	maxhost	8192	devtype	Array	part	Cabinet Pool 1	<table><tr><th>Name</th><th>Value</th></tr><tr><td>uid</td><td>60000970000194900149533030324130</td></tr><tr><td>sgname</td><td>Iglob110_SG</td></tr><tr><td>firmware</td><td>5874.248.194</td></tr><tr><td>model</td><td>VMAX-1SE</td></tr><tr><td>serialnb</td><td>N/A</td></tr><tr><td>datagr</td><td>SYMMETRIX-LUNS</td></tr><tr><td>vendor</td><td>EMC</td></tr><tr><td>pdname</td><td>Not Visible</td></tr><tr><td>memberof</td><td>07F08F</td></tr><tr><td>apgdb</td><td>APG-DB</td></tr><tr><td>ip</td><td>000194900149</td></tr><tr><td>unit</td><td>KB/s</td></tr><tr><td>source</td><td>XmlCollector</td></tr><tr><td>parttype</td><td>LUN</td></tr><tr><td>name</td><td>WriteThroughput</td></tr><tr><td>device</td><td>000194900149</td></tr><tr><td>apgrp</td><td>group</td></tr><tr><td>devtype</td><td>Array</td></tr><tr><td>part</td><td>02A0</td></tr></table>	Name	Value	uid	60000970000194900149533030324130	sgname	Iglob110_SG	firmware	5874.248.194	model	VMAX-1SE	serialnb	N/A	datagr	SYMMETRIX-LUNS	vendor	EMC	pdname	Not Visible	memberof	07F08F	apgdb	APG-DB	ip	000194900149	unit	KB/s	source	XmlCollector	parttype	LUN	name	WriteThroughput	device	000194900149	apgrp	group	devtype	Array	part	02A0	<table><tr><th>Name</th><th>Value</th></tr><tr><td>uid</td><td>60000970000194900149533030324130</td></tr><tr><td>sgname</td><td>Iglob110_SG</td></tr><tr><td>firmware</td><td>5874.248.194</td></tr><tr><td>model</td><td>VMAX-1SE</td></tr><tr><td>serialnb</td><td>N/A</td></tr><tr><td>datagr</td><td>SYMMETRIX-LUNS</td></tr><tr><td>vendor</td><td>EMC</td></tr><tr><td>pdname</td><td>Not Visible</td></tr><tr><td>memberof</td><td>07F08F</td></tr><tr><td>apgdb</td><td>APG-DB</td></tr><tr><td>ip</td><td>000194900149</td></tr><tr><td>unit</td><td>Nb/s</td></tr><tr><td>source</td><td>XmlCollector</td></tr><tr><td>parttype</td><td>LUN</td></tr><tr><td>name</td><td>ReadRequests</td></tr><tr><td>device</td><td>000194900149</td></tr><tr><td>apgrp</td><td>group</td></tr><tr><td>devtype</td><td>Array</td></tr><tr><td>part</td><td>02A0</td></tr></table>	Name	Value	uid	60000970000194900149533030324130	sgname	Iglob110_SG	firmware	5874.248.194	model	VMAX-1SE	serialnb	N/A	datagr	SYMMETRIX-LUNS	vendor	EMC	pdname	Not Visible	memberof	07F08F	apgdb	APG-DB	ip	000194900149	unit	Nb/s	source	XmlCollector	parttype	LUN	name	ReadRequests	device	000194900149	apgrp	group	devtype	Array	part	02A0
Name	Value																																																																																																																																																																																		
model	VMAX-1SE																																																																																																																																																																																		
datagr	SYMMETRIX-DISKS																																																																																																																																																																																		
memberid	7																																																																																																																																																																																		
dname	DISK_GROUP_001																																																																																																																																																																																		
memberof	DF-7A																																																																																																																																																																																		
diskcost	0.280																																																																																																																																																																																		
apgdb	APG-DB																																																																																																																																																																																		
partmode	Assigned																																																																																																																																																																																		
name	WriteRequests																																																																																																																																																																																		
apgrp	group																																																																																																																																																																																		
devtype	Array																																																																																																																																																																																		
firmware	5874.248.194																																																																																																																																																																																		
diskgrp	1																																																																																																																																																																																		
hyperid	49																																																																																																																																																																																		
lunname	06A9																																																																																																																																																																																		
raidgrp	53																																																																																																																																																																																		
vendor	EMC																																																																																																																																																																																		
ip	000194900149																																																																																																																																																																																		
unit	Nb/s																																																																																																																																																																																		
disktype	SATA																																																																																																																																																																																		
source	XmlCollector																																																																																																																																																																																		
parttype	Disk																																																																																																																																																																																		
partdesc	SATAHGST																																																																																																																																																																																		
device	000194900149																																																																																																																																																																																		
diskrpm	7200																																																																																																																																																																																		
part	07A:C0																																																																																																																																																																																		
partvrs	N02THJK																																																																																																																																																																																		
Name	Value																																																																																																																																																																																		
model	VNX7500																																																																																																																																																																																		
module	SAS																																																																																																																																																																																		
vols	18																																																																																																																																																																																		
serialnb	APM00121301071																																																																																																																																																																																		
datagr	VNXBlock-StoragePool																																																																																																																																																																																		
maxdisk	1000																																																																																																																																																																																		
apgdb	APG-DB																																																																																																																																																																																		
ip	Iglad039																																																																																																																																																																																		
unit	GB																																																																																																																																																																																		
revision	05.31.000.5.709																																																																																																																																																																																		
source	VNXBlock-Collector																																																																																																																																																																																		
parttype	Storage Pool																																																																																																																																																																																		
dgraid	RAD5																																																																																																																																																																																		
name	UserCapacity																																																																																																																																																																																		
device	vrnx1071																																																																																																																																																																																		
apgrp	group																																																																																																																																																																																		
maxhost	8192																																																																																																																																																																																		
devtype	Array																																																																																																																																																																																		
part	Cabinet Pool 1																																																																																																																																																																																		
Name	Value																																																																																																																																																																																		
uid	60000970000194900149533030324130																																																																																																																																																																																		
sgname	Iglob110_SG																																																																																																																																																																																		
firmware	5874.248.194																																																																																																																																																																																		
model	VMAX-1SE																																																																																																																																																																																		
serialnb	N/A																																																																																																																																																																																		
datagr	SYMMETRIX-LUNS																																																																																																																																																																																		
vendor	EMC																																																																																																																																																																																		
pdname	Not Visible																																																																																																																																																																																		
memberof	07F08F																																																																																																																																																																																		
apgdb	APG-DB																																																																																																																																																																																		
ip	000194900149																																																																																																																																																																																		
unit	KB/s																																																																																																																																																																																		
source	XmlCollector																																																																																																																																																																																		
parttype	LUN																																																																																																																																																																																		
name	WriteThroughput																																																																																																																																																																																		
device	000194900149																																																																																																																																																																																		
apgrp	group																																																																																																																																																																																		
devtype	Array																																																																																																																																																																																		
part	02A0																																																																																																																																																																																		
Name	Value																																																																																																																																																																																		
uid	60000970000194900149533030324130																																																																																																																																																																																		
sgname	Iglob110_SG																																																																																																																																																																																		
firmware	5874.248.194																																																																																																																																																																																		
model	VMAX-1SE																																																																																																																																																																																		
serialnb	N/A																																																																																																																																																																																		
datagr	SYMMETRIX-LUNS																																																																																																																																																																																		
vendor	EMC																																																																																																																																																																																		
pdname	Not Visible																																																																																																																																																																																		
memberof	07F08F																																																																																																																																																																																		
apgdb	APG-DB																																																																																																																																																																																		
ip	000194900149																																																																																																																																																																																		
unit	Nb/s																																																																																																																																																																																		
source	XmlCollector																																																																																																																																																																																		
parttype	LUN																																																																																																																																																																																		
name	ReadRequests																																																																																																																																																																																		
device	000194900149																																																																																																																																																																																		
apgrp	group																																																																																																																																																																																		
devtype	Array																																																																																																																																																																																		
part	02A0																																																																																																																																																																																		

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

24


Shown here is an example of what you would find in property database for a sample of four different metrics, randomly chosen. No matter how many properties there are, the minimum of the common convention is always present, namely: source, devtype, device, parttype, part, name, and unit.

Take note how they are various in nature and thus the number of properties for each metric varies. Simply by reading the “property card”, you can perfectly and precisely identify where this data comes from when referring to the whole infrastructure.

Defining the property for a given collection is already done for you when implementing a SolutionPack. You should not worry about creating or taking care of this step.

Note: the above samples are taken from (left to right):

- a VMAX Disk metric
- a VNX Storage Pool metric
- a VMAX LUN metric (1)
- a VMAX LUN metric (2 – same object but 2nd metric)



Since a report is defined as a template, you can choose to cover a period of time. You must at least define the default report time range. In order to better understand the terms used, the above depicts the time range selection. Of course, at run time, when viewing a report, a user (if allowed) can override the default selection in order to adjust the coverage.

ViPR SRM comes pre-configured with default time ranges, listed under the display menu. When defining a report, you can also choose the default time range. The other concept is to tell the report engine how to reference the selected time range: based on current present time (last) or based on an exact time boundary (previous or current). The above depicts their meaning. ViPR SRM bases its time range selector using three options: “previous”, “current”, or “last”.

From display menu, and when using report editor choose:

- **previous:** selects the entire and complete previous range, according to period’ boundaries.
- **last:** selects one complete period range back from present time.
- **current:** selects the current period range, according to period’ boundaries.

Let’s assume the current date and time is: 15th of May, 08:00 AM.

- The previous day (24 hours) boundaries become exactly: May 14th from 00:00AM to 11:59:59 inclusive.
This could be translated in selecting ‘yesterday’ time range selection.
- The last day (24 hours) boundaries become exactly: May 14th 08:00AM to May 15th 08:00AM inclusive.
This could be translated in selecting the very last 24 hours from the actual present time.
- The current day boundary becomes exactly: May 15th 00:00AM to 11:59:59PM inclusive (some projection in the future, could be useful for trending and forecasting report time range selection).
This could be translated in selecting complete present period definition (present day).


The same concept applies for weeks, months, years, quarter and so on.

This concept is used in report tweaking but also in Report Edition.

Aggregates

Definition

An aggregate is a data record summarizing many values with an effect of reducing the sampling rate. It compensates this effect by providing more statistical details.



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

26

Two types of values are stored in the database: real-time values (Raw) and (time based) aggregated values. Real-time values are the actual raw values collected by the collectors. Aggregated values are the compounding of several real-time values into one value. Aggregated values improve report performance and save disk space. Aggregated values are created and updated as real-time values are collected so that aggregated values are always current and precise. A default installation contains 1 hour, 1 day, and 1 week aggregates. Recall that this is Backend's responsibility is to compute and populates aggregates in real-time.



Real-time values are the raw value of metrics as they are collected, representing 1 data point in time. They are aligned with collector's polling cycle (more on time alignment later).

Each aggregate record holds an average, sum, and count of the samples in the aggregate as well as the maximum, minimum, and last values taken from the raw values of the aggregated samples.

Default aggregated values are:

- 1-hour: when polling cycle is 5 minutes, it takes approximately 12 samples to generate one single aggregate record for a given metric.
- 1-day: when polling cycle is 5 minutes, it takes approximately 288 samples to generate one single aggregate record for a given metric.
- 1-week: when polling cycle is 5 minutes, it takes approximately 288*7 samples to generate one single aggregate record for a given metric.

Using Aggregates



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Basics of Reporting

27

The main reason to use aggregated values is to improve and optimize report generation and calculation performance. Also, real-time values are most often used for a short period of time selection, typically less than 1-day (24 hours) as opposed to aggregated values which are more suitable for longer time range reporting and trending analysis.

For example, if you generate a yearly report, you probably do not need the >105,000 value dataset packed into your graph (assuming polling every 5 minutes over 365 days). A report using weekly averaged values is likely to be more effective. Do not forget that retention policy may result in raw data not being available for a period selection larger than the policy.

Also remember that the aggregates “functions” (average, min, max, sum, etc) are available only when selecting an aggregate (1-hour, 1-day or 1-week). They have no effect on real-time selection.




Using aggregates in reports is faster and lighter for the Report Generator. Especially when using large time range selections (week, month) where numbers of values in time-series could be substantial compared to real-time aggregates when computing over 500 LUNs down into one single chart.

In a graph, there will be a change in the degree of precision of the graph. When using or creating a report, always ask yourself the question about the relevant level of detail you want to achieve in the report.

The above examples are questionable. If the report goal is to show the general aspect of the behavior of a given metric over time in order to, let say, visually identify the peak hours, the hourly report (bottom one) is preferred. On another hand, if the goal the report is to identify the peak value, maybe going with a graph is not best, a table would be sufficient (and still, selection would make use of “max” value of aggregate over the period instead of real-time value). From the above, you can notice how the top graph packs the data in the graph frame, less interesting.

The benefits are at calculation time, stress and load on the Database and Portal Report Engine.



Another concept with ViPR SRM is to differentiate the nature of timeseries data from events data. Basically, timeseries is the most common. Any metric from data collection involves a series of values collected over time, regularly. Predictability of timeseries is key: collection kicks in and collects data every X number of seconds (ie: 300 for 5 minutes), yet regularly.

Events are the counterpart of data collection. In that case, no data collection process is involved. Generally speaking, events are generated from a given source system (array, switch, etc) and the information (time/date and event details) are forwarded to a ViPR SRM listener that writes the data into the database. Predictability of an event cannot be established. Events also have properties, as demonstrated above, when selecting the event on the graph. Events can also have a start and an end.

With ViPR SRM, both types of data can be manipulated at the report level. Some report types will allow you to merge both on the same report view, which speaks very well in visualizing events with metrics behavior over a period of time (as shown above).

Source: All>>ViPR SRM Reporting class>>Module 04>>Lab 10 - Overlay - VNX

Module: Basics of Reporting

Lesson Summary

During this lesson the following topics were covered:

- Identifying the metrics and properties of many component types
- Uses of time-series and event driven data in the report engine
- Aggregation of data into aggregated groupings based upon needs



This lesson covered the elements of basic reporting including metrics and properties, how they are used in both time series and event data collection and finally how data is aggregated to satisfy many different types of reporting needs.

Module: Basics of Reporting

Lesson: Browsing the Data Set


This lesson covers the following topics:

- Scoping the data set
- Advanced search engine
- Inspecting database metrics
- Report Wizard



This lesson covers browsing the data set including discussion of using the user interface to browse metrics and properties. Particular emphasis is given to using filtering, advanced search tool, management of database metrics, and report wizard to gain an understanding of how the data is represented and how it can be used in the your reports.

Using Filter Wizard



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

32

An important concept in report customization (and creation) is to select the proper metric scope/context. This is achieved by setting a filter. The filter selects the data metrics that will be included for report computation. A tool that helps in creating the filter expression is the Filter Wizard.

The filter wizard is extensively used throughout SRM and not just for report customization. The filter wizard can be defined as the GUI tool that allows you to work with a filtering expression. The blue shape over the above snapshot delimits the filter wizard working area. Here is a list of some of other UI tools that use the filter wizard throughout Watch4net UI:


- **Report Edition:** report filtering, metric selection in table columns, formulas.
- **Search Engine:** selecting a metric scope.
- **Alerting Engine:** selecting metric in analysis process.
- **Management of Database Metric:** selecting metric for deletion or for reference.

The filter wizard shows up as of logical boxes, each one containing one or several Watch4net properties and values expression, linked together by logical operators like AND and OR.

As you build your filter, the boxes assembly evolves to reveal your filter expression.

At the bottom of the graphical view of the filter, you have the equivalent in text, as you would type it in (aka filter expression).

Filter Wizard



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

33

The Initial state of a new filter expression is showing the expression as “everything”. This means the current context is selecting all metrics for report computation. In order to limit the scope of the report context, you need a filter expression. Filter works ONLY with metric properties.

To filter on a given property, hover the mouse on the box and left-click on the “Everything” box, a contextual menu will appear, and choose “Refine / using a wizard...”. This will launch the Refine Wizard of filter wizard.

Or you can click directly on other menu items on the contextual menu, as such:

- “with an expression...”: allows you to type in the filter’s expression manually
- “a device type...”: launches the Refine Wizard and set left-hand side of the expression with the “devtype” property
- “a device...”: launches the Refine Wizard and set left-hand side of the expression with the “device” property
- “a component type...”: launches the Refine Wizard and set left-hand side of the expression with the “parttype” property
- “a component...”: launches the Refine Wizard and set left-hand side of the expression with the “part” property
- “a metric type...”: launches the Refine Wizard and set left-hand side of the expression with the “name” property

Filter Wizard

Wizard automatically makes suggestions from property values available

Refine *

Please enter a value for property devtype

devtype is Array

is

is not

is like

is not like

is present

is not present

starts with

doesn't start with

ends with

doesn't end with

Click on "is" to pull down the list of operators

Refine devtype=="Array"

Please enter a value for property partype

partype is

Application

CifsShare

Controller

Disk

FileSystem

Host

LUN

Memory

Pool

Port

...

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Basics of Reporting 34

Once you have chosen the property name (left hand of the expression), you have to enter (or select) a property value. Again, just type in and the Wizard will provide you with some matching suggestions.

When using filter wizard, default operator is set to an exact match (ie "is" keyword) between property name and property value. Maybe this is not what you want.

To change the operator, click the "is" keyword between the property name and the property value input text box, a drop-down menu will show up, giving you list of available operators. You could then set proper operation you want about the relation between property and the value.

The must be operator one of:

- "is": an exact match of a property value
- "is not": negation of an exact match of a property value
- "is like": is a flexible match of a property value
- "is not like": a negation of the "is like"
- "is present": no matter what a property value is, match if value is present
- "is not present": negation of the "is present"
- "starts with": flexible match of a property value, constrain by the start of the text string
- "doesn't start with": negation of "starts with"
- "ends with": flexible match of a property value, constrain by the end of the text string
- "doesn't end with": negation of "ends with"
- "contains": flexible match of a property value, constrain by the middle part of the text string
- "do not contains": negation of "contains"

Filter Operands

Operation	Operand
AND	&
OR	(pipe)
NOT	!
Is like	=
Strictly is	==
Wildcard (many characters, including no character)	%
Wildcard (one character)	_ (underscore)
Priority, expression grouping	()



Example: `devtype=='Array' & (parttype=='Disk' | parttype=='LUN')`

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Basics of Reporting

35

Any filter expression created using the filter wizard end up in a textual expression that is similar to the SQL query language.

For the more advanced user, it sometimes faster to type in the filter expression using the filter wizard. The user would have to select the “Edit expression” from the filter wizard contextual menu. The above table shows the available operand characters to use for the desired operation.

Filter Wizard



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting 36

Once the first property of the filter expression entry is set, another property filter can be added by selecting the existing filter and using either the **AND** or **OR** operand.

Just click the first filtered property and the contextual menu will appear with the **AND** and **OR** logical operators added as shown here.

Filter by Device


The screenshot shows the EMC² management console interface. The top navigation bar includes 'Welcome (logged in as admin)', the date 'Wednesday, December 18, 2013', and links for 'Settings', 'Administration', 'Help', and a power icon. Below the navigation bar, there's a 'Your Company Logo' placeholder and a sidebar menu with options like 'Scheduled Reports', 'Stored Reports', 'Favorite Reports', 'Dashboards', and 'Explore'. The 'Explore' section is expanded, showing a list of device types: 'Storage', 'Storage Connectivity', 'IP Connectivity', 'Hosts', and 'Search for any device'. The 'Search for any device' option is selected, and a list of device IDs is shown below it. The main content area displays the 'Explore / Search for any device' report. It includes a search bar with 'Sun 25' entered, a 'DISPLAY' button, and 'EXPORTS' and 'TOOLS' options. The report title is 'Explore / Search for any device', and it shows the date 'December 2013, Tuesday 17 » Wednesday 18, 2:45 PM EST' and 'Last 1 Day'. Below the title, it says 'List every devices of every device type (including 'Datastore', 'PassiveHost', etc)'. A message indicates '190 elements found, displaying 1 to 25'. The table has two columns: 'Device Type' and 'Device Name'. A green callout box points to the 'Filter by devices or types' dropdown menu in the table header. The table lists various devices, including Arrays and Datastores.

Device Type	Device Name
Array	000194900084
Array	000195700399
Array	000195700407
Array	VNX-112-112
Array	VNX-54-60
Datastore	10.127.96.57-InternalDrive1
Datastore	10.127.96.57-InternalDrive2
Datastore	AppSync Shared Production
Datastore	AppSync_Shared_Recovery
Datastore	AppSyncLabinfra

Another way to filter is to use the **Search for any device** report in the **Explore** area, in browse mode. It lists every device you have discovered. Sub-components are not listed in this report. You can use the filters to search by device name or type. Once you find your device, just click on it to jump to the configuration item detail.

This method of searching is more focused on devices. A single device will appear only once in this output, while the previous search feature might show many reports that contain the device.

Advanced Search



One of the first applications of filtering is depicted with Advanced Search engine.

At first, let's say that Advanced Search is more flexible than Simple Search. Those operate at different levels.


Simple Search matches your search criteria throughout the report inventory and finds the best match.

The Advanced Search engine selects a metric scope described by the filter parameter and promptly generates a new report node described by the "Expansion" parameter (drilldown structure).

The key in using Advanced Search is the filter and the expansion. For now, just keep in mind that Expansion is just a way to group metrics by the given property. Each grouping property will create a drilldown level, down to the last item.

Usage of this feature is recommended for users who are familiar with reporting and the collection details (properties). The "Search Base" drop list has two options. "search from root" means to perform the search against the whole database content. The second option is "search from the selected node" which means performing the search from the current context of the current report node you are looking at.

For a general search, set a filter in order to limit the scope of metric selection, set the expansion expression and select "search from root".



There is a formula string displayed below the filter's logical boxes. This is actually the filter's expression, displayed in a textual manner.

When you click on the filter, the "Edit expression" panel shows up and allows you to edit this expression without using the Wizard. You can build the same filters as you would with the Wizard. Whether you opt for editing the expression by hand or using the Wizard, the filter will always be displayed in both manners at the same time.

Here are the rules for building a successful filter without the Wizard:

- Values should be entered between single quotation marks.
- Use parenthesis, but don't forget to close them.
- If unsure, first build your filter with the Wizard, then look at the expression.
- Look at the next page for a list of available operands.


Advanced Search Expansion

The screenshot displays the 'Advanced Search Expansion' dialog in the EMC2 interface. The 'Filter' section contains a rule: 'devtype is FabricSwitch' and 'parttype is Port or parttype is Interface'. The 'Expansion' section shows 'device part' with a book icon highlighted by a green circle. The 'Search Base' is set to 'search from root'. The 'Available properties' list includes 'Component name (part)' which is selected. The 'Description' pane shows 'device' with no description available. The dialog has 'Cancel' and 'Apply' buttons.

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Basics of Reporting 40

Unless you already know by which property names you want the search result to be presented, you might want to explore the available property list by clicking on the book icon. This will bring up a new pop-up allowing you to choose among the context available property list. It is also possible to set the order of appearance of the selected properties by performing some drag and drop.

Remember that the sequence of the properties in the “Expansion” box will affect their sequence in the search result.



This is an example of Advanced Search usage:

The current context of the search example is: “Which director frontend port of the array serial number '000194900854' are instrumented by the metric named IORate?” If such, list the port numbers and show a graph of the metric performance.

1. Based on the previous statement, the filter expression is set to select the scope of desired metric selection:

device=='000194900854' & parttype=='Port' & name=='IORate'

Which translates into more human reading to:

“device is 000194900854 “ AND “parttype is Port” AND (metric) “name is IORate”.

2. Then we want the resulting metric context to be grouped by port number and director (port, director property expansion). To achieve this, expansion is set with “part,director” (which describes the instance of the entity of parttype “Port”).
3. Once applied, the report pane shows a table with the list of distinct values for “part,director” context defined by the expansion expression.
4. When clicking on an entry of the table (drill-down), the report page changes to a graph showing the metric IORate (IOPS) over the default period (1 day).

Report Structure using Advanced Search

The screenshot displays the EMC ViPR SRM Advanced Search interface. At the top, there's a search bar with the text 'director part'. Below it, the 'All / Search Results' pane shows a list of 14 elements found, displaying all elements. The elements are listed in a table with a 'Name' column. The elements are: FA-10E, FA-10F, FA-10G, FA-10E, FA-7H, FA-7E, FA-7G, FA-9E, FA-7F, FA-9G, FA-7H, FA-7E, FA-7G, FA-9E, FA-7F, FA-9G. The right pane shows a detailed view of 'FabricSwitch / losam240' with 5 elements found, displaying all elements. The elements are listed in a table with a 'Name' column. The elements are: Sensor, Processor, Port, Interface, Memory. The bottom of the interface features the EMC logo and copyright information: Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Basics of Reporting 42.

From the previous example of Advanced Search usage, notice how the report structure showed up. Also notice the grouping: per port and per director (unified), on a single report level. You could have decided to have them structured differently: on two separate report levels. You can do this by using a space separator instead of the comma under.


Also notice that a chart (graph) of all applicable metrics are plotted because of advanced search. Notice how quick you can drill down data using this simple tool.

By extending that capability, combined with the knowledge about Common Convention, you can easily explore the assets of a given ViPR SRM installation. Simply use Advanced Search, with minimum filter (ie a few device name, or simply without filter) but with an expansion expression set to the common convention property list, as follows (IMPORTANT: use SPACES as property separator in Expansion box):

source devtype device parttype part name

Management of Database Metrics

Administration→Modules→Frontend→Management of Database Metrics



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

43

Another great tool for accessing inventory like reports and identifying the metrics associated with different devices and parts is the **Management of Database Metrics** module. However, this tool is only available to those users with administrative privileges. To use the module, enter the Administration view and select the module under the Frontend group.

Filter for Metrics Management

EMC²

Administration >> Modules >> Metric Management

Management of Database Metrics | Administration

This module currently allows metrics deletion. Please note that the changes will not be immediately visible. The Backend service and Tomcat service need some time to refresh the data. If the property store or the property flat table is used, the deleted metrics will still be visible in the reports until it is refreshed.

Metric Selection

Filter

1

devtype is Hypervisor and parttype is Memory and device is lglac019.lss.emc.com

devtype='Hypervisor' & parttype='Memory' & device='lglac019.lss.emc.com'

Maximum results 100

Show last timestamp for all results

2

Properties to show name device part

Query

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Basics of Reporting 44

Once selected, the Management of Database Metrics module displays a Metric Selection screen which has a filter section. The filter here works the same as filtering already discussed. The goal here is to use the filter to build a query that will result in a list of metrics.

1. For example, here we set the filter for:

devtype is Hypervisor and parttype is Memory and device is lglac019.lss.emc.com

This filter will result in a list of Memory metrics on the designated Hypervisor.

2. The **Properties to show** block allows you to define the columns to be displayed in the result of the query. In this example, the resulting report will show name (metric name), device (device instance name), and part (name of the part associated with the metric).

Query Result

Device = lgla019.lss.emc.com

- parttype = Memory (from filter)
- part – System (memory)
- name = Metrics associated with system memory

10 results.

	name	device	part
☐	CurrentUtilization	lgla019.lss.emc.com	System
☐	TotalMemory	lgla019.lss.emc.com	System
☐	SharedCommon	lgla019.lss.emc.com	System
☐	SwapIn	lgla019.lss.emc.com	System
☐	SwapOut	lgla019.lss.emc.com	System
☐	Granted	lgla019.lss.emc.com	System
☐	Active	lgla019.lss.emc.com	System
☐	Balloon	lgla019.lss.emc.com	System
☐	Shared	lgla019.lss.emc.com	System
☐	SwapUsed	lgla019.lss.emc.com	System

⏮ All | None | Invert

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

45

The result of the query appears. In this case, the result is the list of the 10 different values associated with System memory. At this point, you can use the list to further research all available properties for any of the metrics by selecting the desired metric. This can be very useful in learning the properties that are associated with any given metric to help you learn how to build better reports using all the properties and metrics available.

Available Properties Card

The screenshot displays the EMC² web interface. On the left is a navigation menu with categories like Administration, Modules, and Management of Database Metrics. The 'Management of Database Metrics' section is active, showing various tools. On the right, a table lists metrics: CurrentUtilization (checked), TotalMemory, and SharedCommon. An 'Available properties' dialog box is open, showing a list of properties and their values for the selected metric.

name	device
☒ CurrentUtilization	Iglac019.lss.emc.com
☐ TotalMemory	Iglac019.lss.emc.com
☐ SharedCommon	Iglac019.lss.emc.com

Name	Value
nbcore	4
parentt	ClusterComputeResource
model	PowerEdge 2950
vportgrp	Management Network
eraseme	
memberof	CSE-Hopkinton
deviceid	44454c4c-4300-104b-805a-b8c04f4e324a
apgdb	APG-VMWARE
maintmod	false
lpsubnet	255.255.255.0
nbcpu	2
devdesc	VMware ESX 5.1.0 build-1065491
vmotion	Yes
vcip	10.247.166.41
nodegrp	All Hypervisors

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

46

One of the benefits of this module is that for any selected metric, you can display the available properties associated with it. Simply click on the metric and the **Available properties** card will display. This is extremely useful, especially for those new to the data structures used in ViPR SRM and to those who need to design reports around the metric. Use of this module will, over time, strengthen your knowledge of the metrics and their associated properties available in the ViPR SRM data set.

Report Wizard



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Basics of Reporting

47

While in Browse Mode, the Tools menu provides a complete set of tools for those who want to make life easier. One of those tools is the **Report Wizard**.

The Report Wizard is a useful tool that allows you to quickly build a simple report.

Using Report Wizard Tool



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting 48

Here is an example of what the Report Wizard looks like.


The **View Type** drop-down menu at the top left allows you to choose the view to sort elements in the first column. It is either:

- APG Topology
- Core Networking
- Device Browser
- Devices by Location

In order to be able to run the report, you must select at least one metric.

At first, when you open the Report Wizard, only the left-most column is visible. Other columns appear as you select items in the current column. At the top of each column, there is a drop-down menu that allows you to change the elements displayed. There is also, at the bottom of each column, a magnifying glass icon, circled above, that applies a filter to the column.

Using Report Wizard Tool




EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting 49

Since we are using the device browser here, the first column becomes a list of device types. Selecting one or more device types causes the second box to open in the display. this column displays the known devices of the selected device type. Again, select one or more devices to continue.

Using Report Wizard Tool



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting 50

The next column to appear displays Component Types (parttype) for the selected devices. Again, you can select one or more components to build a report on to get to the next column.

Using Report Wizard Tool

The screenshot displays the 'Report Wizard' interface with a 'Device Browser' view type. It features four main columns for selection:

- Class:** A list of classes including 'Array' (selected), 'FabricSwitch', 'FileServer', and 'Host'.
- Hostname:** A list of hostnames including '000192606304', '000192606327' (selected), and 'myVNXsystem'.
- Component Types:** A list of component types including 'Access', 'AccessTolun', 'Control', 'Disk', 'LU', 'Member', 'Port' (selected), 'Service Level', and 'Storage Group'.
- Components:** A list of components including '0' (selected) and '1'.

Below the columns, there are two expandable sections: '0 Selected Metrics' and 'Optional Report Settings'. A yellow arrow points from the 'Port' component type in the third column to the '0' component in the fourth column.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Basics of Reporting 51

Here, the **Port** component type is selected then in the next column, both ports are selected. The port 0 is a component name for the selected Component Type. The components are actually known as **part** properties.

Available Metrics

For the chosen ports on the array, the available metrics are the last stage of the Report Wizards construction

- devtype=Array
- device= 000192606327
- parttype=port
- part= 0 (port#)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

52

Finally, a list of metrics for the selected components are displayed. Specifically, these are the metrics associated with the devicetype=Array, device=000192606327, parttype=port and part= 0.

Click on Add Metrics to set the metrics for the report. The metrics will be added in the bottom of the display as you will see on the next page. The upper portion of the display will reset to the Device Type entry ready to add another devtype/device to the report.

Using Report Wizard

Selected Metrics

Optional Report Settings

5 Selected Metrics

1 Represent all the currently selected metrics. Bear in mind that each displayed metric is related to an underlying set of selected components (that are not displayed here).

Availability AverageRequestSize IORate KBytesTransferred Throughput

Optional Report Settings

Display real-time average over last 1 day

Automatic report type selection

Table

Single Graph

Multi-Graph

Stacked Chart

Horizontal Bar Graph

A table report with one line per metric	Display all metrics in one graph	Each metric displayed in its own graph	Metrics are stacked one over the other	Each metric is a horizontal bar in the graph
---	----------------------------------	--	--	--

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting 53

While the selected metrics are displayed, the trash can icon next to each metric allows you to remove any metrics you may not desire.

Below the selected metrics, at the bottom of the window, the Optional Report Settings contains more options for the report. With these options, you can:

- select one of five report types or allow the wizard to do so automatically (default)
- manage Axis (even add a second axis as shown)
- manage scaling
- manage aggregation of selected metrics

A report must be run before it can be saved to **My Reports** and the save button will be shaded out until the report has been run. Once saved to My Reports, the report may be edited again using either the Edit Mode or the Report Wizard.

My Reports/Wizard Table

Report Wizard

Run Save in "My Reports" Quit Wizard

1

My Reports / Wizard

June 2014, Sunday 29 » Monday 30, 10:24 AM EDT | Last 1 Day

Device type: Array Device name: 000192606327

Component name: 0

5 elements found, displaying all elements.

Specify a report name...

Please specify a name for the report.

Report Name: Port table

2

Cancel Ok

Name	Value	Unit
Availability	100.00	%
AverageRequestSize	0.00	KB
Throughput	0.00	MB/s
IORate	2.33	KB/s
KBytesTransferred	0.00	KB/s

3

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Basics of Reporting 54

In the example above, the table report type is selected then run. Once the report has been run, the **Save in "My Reports"** button is enabled.

1. Click the **Save in "My Reports"** button.
2. Enter a report name and click **OK**.
3. Report will appear in the **My Reports** branch of the report tree.

Display Management Settings

Single Graph or Stacked Graph

Metric Management:

- Main Axis
- Scaling Mode
- List of Metrics
- Add a Secondary Axis

Metric management: handle each metric separately ▼

• Main Axis

Scaling mode: -None- ▼

Availability ⚙

AverageRequestSize ⚙

IORate ⚙

KBytesTransferred ⚙

Throughput ⚙

Add a Secondary Axis

• Secondary Axis

Scaling mode: -None- ▼

Drag and drop the metrics from the other axis...

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

55

In addition to the display and aggregation settings, both the single graph and the stacked graph allow settings adjusting the left and right axes of the graph. You can set metrics to display on either a single axis or you can define a secondary axis and determine which metrics will be represented by dragging and dropping the desired metric onto the secondary axis. The use of the axis settings and setting up a secondary axis will be covered in greater detail when we discuss the different report types later in this course.

Module: Basics of Reporting

Lesson Summary

During this lesson the following topics were covered:

- Use of filters and how they are employed in scoping the data set
- Extended use of filters and other options in refining output while using the advanced search engine
- Using the GUI management of database metrics to identify the properties and metrics for given device types
- Using the Report Wizard to explore and create simple report



This lesson covered the use of filtering to limit the scope of data analyzed by either search or advanced search operations. Emphasis is given to understanding how the same operations are used in the building of reports which we will soon explore. To aid in understanding how metrics and properties are represented for different devices, we also explored devices using the management of database metrics function of the user interface.

Lab 1: Discovering Available Datasets

This lab covers table report creation:

- Part 1 – Familiarize with Advanced Search Engine
- Part 2 – Using Management of Database Metrics
- Part 3 – Practice report creation using Report Wizard

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Basics of Reporting

57

This lab 1 session allows you to practice of basic features of the portal.

Part 1 – Familiarize with Advanced Search Engine tool

Part 2 – Use of Management of Database Metrics tool

Part 3 – Practice using Report Wizard (basic reporting)

Lab 1 Debrief: Discovering Available Datasets

This lab exercise allowed you to:

- Familiarize with Advanced Search Engine, Management of Data Metrics tools, and Report Wizard tools.
- Recognize the tools you have on hand to discover the collected data set on a given platform.
- Familiarize yourself with metric, properties, timeseries, and aggregates concepts.



In this lab, we became familiar with the Advanced Search Engine tool, use of Management of Database Metrics tools, and practiced using the Report Wizard (basic reporting).

Module Summary

Key points covered in this module:

- Discuss the concepts important to good report design
- Describe how the reporting engine works
- Explain the base organization of data as it is used by ViPR SRM
- Demonstrate the functionality of ViPR SRM used to browse the data set



This module covered the data that is collected and how it can be used to generate many reports that are necessary to understand the core underlying health of the data storage infrastructure in the data center. The primary purpose was to identify the data collected, how the human administrator of storage data centers can browse the data to gain familiarity with the structure and representation of data, how that data is represented in both time series and event-based collections, and finally how that data is aggregated and used by the report engine in the generation of reports as defined and required.

This slide is intentionally left blank.



Module: Reporting Key Concepts

Upon completion of this module, you should be able to:

- Recognize the difference between Edit and Browse modes
- Describe Filtering and Expansion concepts



This module focuses on the key concepts of reporting as they apply to the ViPR SRM data set. Emphasis is given metric context by the use of Filtering and Expansion concepts.

Module: Reporting Key Concepts

Lesson: Report Definition Key Concepts

This lesson covers the following topics:

- Introduction to the Edit Mode features
- Scoping the report using filter expressions
- Parent → Child node structure concepts
- Copying pre-existing report templates into My Reports
- Pinning reports vs copy and paste




This lesson discusses both the use of the filter wizard and editing filter expressions. Following filtering, is a discussion of report organization in the iterative approach and the parent/child node structure concept beginning with the static report structure. Finally, is a discussion of the use of the copy and paste feature to copy existing report templates into the My Reports tree branch as well as a quick review and comparison with the report pinning feature.

Browse/Edit Mode Toggle

Browse Mode

- Edit Mode displayed
- Search, advanced search available
- Display, Exports, and Tools available



Mode Toggle
Switch

Edit Mode

- Browse Mode displayed
- Only Exports menu available




EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

3


The first item of the menu bar is the mode toggle switch. Initially, it displays **Edit Mode**. This indicates that the current mode is the browse mode. It is the default mode upon entry into the user interface. Clicking on the selection causes a transition into the edit mode. Notice that the mode displayed changes to **Browse Mode** indicating that the current mode is the edit mode. Notice that while in edit mode, the only other menu item available on the menu bar is the **EXPORTS** group of features.



Welcome to ViPR SRM's Edit Mode, the rich interface where you can tweak existing reports or create your own.

It is immediately noticeable that the GUI looks somewhat different from the Browse Mode:

- Report Tree on the left does not look the same. It now displays report definitions and templates for all currently available reports including the **My Reports** branch which is where all your personally created or modified reports will be kept.
- Buttons appear at the top of the Report Tree panel that allow you to create, cut, copy, paste, paste as link, or delete the selected report.
- Report Pane is also different in many ways. Most significantly it now displays the settings tabs for the selected report.



While in the Edit mode, these buttons are located above the Report Tree under the logo graphic.


They allow you to edit the tree:

- Add a new report
- Move (cut) a report
- Copy a report
- Paste a report
- Paste a report as a link
- Delete a report

When pasting a report as a link, three choices are given:

- **U:** Paste according to the UID (is defined) Not available for some reports
- **A:** Paste as an absolute link
- **R:** Paste as a relative link

- **UID Links:** (if available from source report node) refers to a named report node, the best available selection.
- **Absolute Link:** The complete path is embedded into the link - good to use when the destination of the link will not move, but the link itself may move.
- **Relative Link:** A shorter path is embedded into the link. The link is broken if either the link or its destination moves.




The report tree lists all the currently available report definitions including those of all installed SolutionPacks. The node navigation tree allows you to create report definition structures or to copy existing reports into your **My Reports** folder for your personal usage. It could be compared to any operating system directory organization with parent and child folders, except that in this case, each entry in the Edit Mode tree is a report definition.

Note that the case-based folders that contain reports in the Browse Mode are repeated in Edit Mode. However, in the case of the edit mode, the reports under these folders are in definition form.

As noted earlier, you may copy any of the report definitions into your My Reports folder so that you may modify them as required. You should not modify any of the pre-defined reports. Instead, copy the desired report and then modify the copy.

The reason for this is two fold. One, you will always have the original definition of the report, and two, you will not over write your modifications when you upgrade the core of SolutionPack reports.

Report Pane



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

7

The report pane allows you to define several and various parameters of the report currently selected in the report tree (on left).

1. At the top of the report pane is the report path, sometimes called the “bread crumb” as it indicate where in the Report tree the selected report exists. It is useful for navigation and reference to report definitions.

2. The Edit Mode allows you to set parameters for you currently selected report node. From the main panel, you will notice up to six main tabs:

- **Filtering & Expansion:** Tab is first, you will use it to set a report node’s name and specify any filters and expansion (if required, depending on your report design and goal).
- **Report Configuration:** Used to set the Display mode of a Table, Graph, other type from the drop down list.
- **Report Details:** [Graph | Table | ...] This tab will change depending on the selection in the Report Configuration tab of the report type.
- **Formula:** Allows the use of formula content to be defined and selected for the currently selected report node.
- **UI Wiring:** An advanced option is outside the scope of this course.
- **Pre-Generation:** Reports that are automatically generated on a predetermined schedule.

Filtering & Expansion

The screenshot shows the 'Filtering & Expansion' tab for a report node named 'Infrastructure'. The window has a title bar 'Infrastructure | Edit Mode on ReportPack "Default ReportPack", template "Device Summary"'. Below the title bar are several tabs: 'Filtering & Expansion' (active), 'Report Configuration', 'Report Details: Table', 'Formula', 'UI Wiring', and 'Pre Generation'. The main area contains the following fields and controls:

- Name:** A text field containing 'Infrastructure'.
- Unique Identifier:** A text field containing 'W4N_DS_INFRA'.
- Filter:** A dropdown menu currently set to 'Everything'.
- Use the filter of this node for:** A dropdown menu set to 'expansion only'.
- Display Policy:** A dropdown menu set to 'always show the node'.
- Expansion:** A section with a collapsed arrow. It contains an information icon and the text: 'The specified properties will be used for the expansion of the current node'.
- At the bottom of the Expansion section are two buttons: 'Expand on every' and 'Add a Property'.
- At the bottom right of the Expansion section is a button: 'Switch to Input Field'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

8

Filtering and expansion

Filtering and expansion affects the report node properties. These properties are crucial to every node and report type because they determine the scope of the data retrieved from the database.

The name of the report node, the filter used to limit metric selection, and some advanced settings that you can adjust are on the **Filtering & Expansion** tab.

Filtering


The filter of a node is an expression used to select the metrics that make up the basic content of the report. Because the filter determines the report's data set, it is a vital part of the overall report configuration. To set the filter properly, you must know which properties are being collected by the collectors and what they mean. There are several ways to create a filter. If you know what properties and values you want to set, you can manually edit a filter. You can also use the **Filter Wizard** to help you define a filter.

Expansion


The expansion setting is central to how the report tree works. It enables the quick generation of report trees, expanding its branches as it transforms individual nodes into multiple nodes. Expansion can automatically create and delete nodes as they appear or disappear in the database.

More on Filtering and Expansion

Filtering



Expansion



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

9

A filter is an expression that limits the number of metrics displayed on a report. A filter is vital to configuring a report because it determines the report's data set. You can manually create a filter when you know which properties and values to include in a filter, or you can use the **Filter Wizard** to help you define the filter. In a previous module, we discussed the uses of filtering and expansion in relation to the advanced search capability of ViPR SRM. In the case of report filtering, the methods for generating filters is the same as that used in the advanced search.

Unlike the use of expansion in the advanced search to define the metrics to expand in the search, Expansion properties in report definitions dynamically create sub nodes. Sub nodes add levels to a report so you that can drill down. When a node is set to expand, the node transforms itself from a single node to as many nodes as there are matches of the specified property name in the data set. By using expansion, you can define a report that automatically creates and deletes nodes based on the properties found in the data set. For example, as devices are added to the network, they are automatically added to the report when their properties match the expansion criteria.

Throughout this course we will demonstrate the use of filtering and expansion in many ways to create reports.

Report Configuration

The screenshot shows the 'Report Configuration' tab in the 'Infrastructure' report editor. The title bar indicates 'Edit Mode on ReportPack "Default ReportPack"; template "Device Summary"'. The tab is part of a sequence including 'Filtering & Expansion', 'Report Configuration', 'Report Details: Table', 'Formula', 'UI Wiring', and 'Pre Generation'. The configuration settings are as follows:

- Display Mode:** default display
- Report Type:** standard table
- Report Rendering:** static charts by default
- Default Duration:** last, 1 week
- Sampling Period:** 1 day
- Sampling Type:** average
- Propagate time selection on drill-down:** unchecked

At the bottom, there are expandable sections for 'Description and Properties' and 'Advanced Settings'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Reporting Key Concepts

10

Use the **Report Configuration** tab to define the core settings for a report, including the report type, the time period of the report, and the aggregation used.

The settings on this tab deal with the report preferences that are used by default when the corresponding node is selected in Browse mode.

Report Types



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

11

As you can see, there are numerous report types available to choose from. Since there is a great variety of report types, we will discuss the features and uses of each type and demonstrate configuration of a large number of those in a later module.

Report Details

Table

The screenshot shows the 'Report Details: Table' configuration window for the 'Infrastructure' report. The window has a title bar 'Infrastructure | Edit Mode on ReportPack "Default ReportPack", template "Device Summary"'. Below the title bar are tabs: 'Filtering & Expansion', 'Report Configuration', 'Report Details: Table' (selected), 'Formula', 'UI Wiring', and 'Pre Generation'. Below the tabs are buttons: 'Sort', 'Paste (0)', '+ Attribute', '+ Property', and '+ Value'. The main area contains a list of items with expand/collapse icons and delete icons:

- Attribute: Technologies*
- Value: Devices Count*
- Value: Availability (%)*
- Value: Current Availability (%)*

Graph

The screenshot shows the 'Report Details: Graph' configuration window for the 'File Systems Utilization' report. The window has a title bar 'File Systems Utilization | Edit Mode on ReportPack "EMC Watch4net Health", template "EMC Watch4net Health"'. Below the title bar are tabs: 'Filtering & Expansion', 'Report Configuration', 'Report Details: Graph' (selected), 'Formula', 'UI Wiring', and 'Pre Generation'. Below the tabs are input fields for 'Main Y-Axis Boundaries' (between 0.0 and 109.0) and 'Main Y-Axis Scaling Mode' (none). Below these is a 'Thresholds' section with a warning icon and text: 'Thresholds must include a unique value and a name. Only the last definition is kept in case of duplicate values. Values are mandatory and must be in decimal format. Duplicate names may be used.'

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

12

Use the Report Details tabs to define the settings for a specific report type.

The Report Details tab content changes according to the type of report you select on the **Report Configuration** tab. Not all report types will use the Report Details tab. The following sets of report details are available by Report Type:

Report Details: Table

Report Details: Graph

Report Details: TopN Graph

Report Details: Map

Report Details: Heat Map

Report Details: Treemap


Report Details: External

Report Details: Topology

Report Details: Overlays

Report Details: Status

Formula



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

13

You can apply formulas to any report node of the tree to perform advanced computations over metric data. Although not stated explicitly, formula transforms metric data into KPI. All formulas have at least one result, and each result has additional options for displaying them in different contexts. Again, this adds great flexibility to reporting therefore, we will dedicate a later module to this topic along with several labs and demos to illustrate their use.

Pre-Generation

The screenshot shows the 'Pre-Generation' tab in the 'device' configuration window. The window title is 'device | Edit Mode on ReportPack "EMC WatchNet Health"; template "EMC WatchNet Health"'. The 'Pre-Generation' tab is active, and a 'New Pre-Generated Report' button is visible. The 'new report' section contains the following fields:

- Recipients:** 'Schedule this report for' dropdown set to 'All Users'.
- Details:**
 - Name:** Text input field.
 - Active:** Checkmark icon.
 - Schedule:** 'Every day' dropdown, 'at' dropdown, '00' dropdown, '00' dropdown.
 - Report Link:** Two radio button options: 'All / EMC WatchNet Health / WatchNet Servers Summary / device/' and 'All / Report Library / EMC WatchNet Health / WatchNet Servers Summary / device/'.
 - Buttons:** 'None' and 'All' buttons.
 - Send confirmation email:** Checkmark icon.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

14


Pre-generated reports are reports that are automatically generated on a predetermined schedule. Pre-generated reports can save users time when they are used to generate reports that take a long time to process. Only administrators can create pre-generated reports.

This feature must be differentiated from the Scheduled Report feature. Pre-Generation changes the behavior of report from being dynamically generated to statically generated for browse mode. When in browse mode, by default, every report gets computed at the moment you select it (thus the “dynamically generated” term). With pre-generated, the report will be presented with data from the last generation time and not with the current time (thus the “statically generated” term).

The use case for pre-generation is for report that takes long time to generate due to their nature (aggregation of thousands of objects, lots of KPI computation, etc.).

By presenting pre-generated report to users, the portal that runs the report gets less hit on resource consumption, and thus enhances the user experience by providing a faster response time in report generation time.

Understanding the Report Structure



If you start a Report design from scratch, here is what you might see from your report pane (Browse Mode). Out of the box, the My Reports branch does not display in browse mode because it has no leaf nodes (reports). The default setting is to not display any branches that have no leaves.

Switching to **Edit Mode** shows an empty structure for **My Reports** at the bottom of the display.

My Reports is a special report starting node given to you when the administrator creates your user account. It is your personal report sandbox. What you define under this node, will only be seen by you. Remember, that if you need to share your report design or want to **publish** it under a public node, you will have to first export the report design and bring it in as a ReportPack (as a new node or a template). We will discuss this action later in the course.

Hierarchical Report Structure

The image displays two screenshots of the ViPR SRM reporting interface, illustrating a hierarchical report structure.

Top Screenshot (BROWSE MODE): The interface shows a sidebar menu with categories like Scheduled Reports, Stored Reports, Favorite Reports, and My Reports. Under My Reports, the 'Parent Child Reports (list)' is expanded, and 'Report 001 - Parent' is selected. The main pane shows the breadcrumb 'All >> My Reports >> Parent Child Reports (list) >> Report 001 - Parent' and a list of child reports: Report 001-01 - Child, Report 001-02 - Child, and Report 001-03 - Child. The date range is 'July 2014, Wednesday 16 - Thursday 17, 10:32 AM EDT' and the last day's average is shown.

Bottom Screenshot (EDIT MODE): The interface shows the 'Report 001 - Parent' in edit mode. The sidebar menu is similar, but the main pane shows the breadcrumb 'All >> My Reports >> Parent Child Reports (list) >> Report 001 - Parent' and a list of child reports: Report 001-01 - Child, Report 001-02 - Child, Report 001-03 - Child, Sub Report 03-001 - Child, and Sub Report 03-002 - Child. The date range is 'July 2014, Wednesday 16 - Thursday 17, 10:32 AM EDT' and the last day's average is shown. The interface also includes a 'Filtering & Expansion' section with a 'Name' field set to 'Report 001 - Parent', a 'Filter' dropdown set to 'Everything', and a 'Use the filter of this node for' dropdown set to 'expansion only'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

16

Understanding the reporting structure is important in ViPR SRM.

Parent-Child relationships create a hierarchical report structure in which a parent node report is affected by the children nodes below it. Basically, there are two types of report entry: **Static or Dynamic**. The above all show examples of static reports. A static report entry will show based on its name in both Edit and Browse mode. Report structure is creating the report drill-down capability. As demonstrated above, selecting “Report 001 – Parent” from report tree shows the underlying child reports in report pane: Parent-has-a-view-over-the-child effect, or downstream view. For demonstration purposes, all reports above, no matter the entry in the structure is set to be of “list” type, a trim-down version of a table report. In the end, the above report structure is meant to demonstrate the feature of organization of reporting in ViPR SRM.

Report structures are similar to a computer’s directory organization. The difference is that each entry in the report structure is a report you can browse. Notice the result of selecting “Report 001 – Parent” in the report pane: child reports are listed. Each report level can be basically defined as a graph or a table (note that there are multiple types of graph, there are basically only two types of tables).

A very basic report structure is one that is completely static: each node is a static entry with some filtering applied (optional). In the example above, the node structure does not do anything, it is just for demonstration of the structural concepts.

The secret of a well-structured and designed report is to begin from the top level where the report viewers are presented with a global view report and as they drill down (clicking items in your report), they get closer to the details. This method brings filtering into play: from parent report node down to a child (or children) report node(s), You can use filtering rules that select from a range of settings ranging from a global set of metrics to a more precise, yet limited, set of metrics that will make up a more detailed report (in terms of specific details).

Understanding the Report Structure: Parallel Nodes

The screenshot displays the EMC ViPR SRM reporting interface. On the left, a sidebar shows a navigation tree with 'BROWSE MODE' highlighted. The main content area is titled 'My Reports / Parallel report nodes' and shows a list of four report nodes: Report 001, Report 002, Report 003, and Report 004. A 'BROWSE MODE' button is visible in the top left, and an 'EDIT MODE' button is in the top right. The interface also shows a date range from July 2014 to Thursday 17, 10:36 AM EDT, and a message indicating '4 elements found, displaying all elements.' On the right, a detailed view of the 'Parallel report nodes' structure is shown, including a 'Drilldown structure' with 'Parent 001' and 'Child 001' (which further branches into 'Sub Child 001' and 'Sub Sub Child 001'). The EMC logo is visible in the bottom right corner of the interface.

Shown above is another report structure with multiple child nodes under the parent **Parallel report nodes**. This report serves no useful purpose, its goal is simply to show report structure capabilities.


Shown here is an example of a drilldown report structure.

All those examples use a static report format making no difference between Browse mode and the Edit mode (construction mode).

Later you will add some dynamic report nodes, where you will define an entry in Edit mode that will generate multiple report nodes in Browse mode. This is an important feature in ViPR SRM reporting.

Remember that each report node may have a filter expression that refines the metrics used for the report. The important rule to keep in mind is that filters are carried down and compounded with an "AND" clause.

Metric “flows” in Report Structure



- Metrics “flow” from Root down to report levels
- Child inherits filter expression from parent
- Ex: report1 metric context limited to VMAX and Disk

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

18


When developing reports while using ViPR SRM, how is data selected and compared with the report structure?

The metric context flows from root report level down to child and leaf reports. As you define a filter in order to limit the metric context for your report, child reports are inheriting this context (defined by filtering expression).

From the above, all of the object metrics are available from a report placed on the root level. The metrics then flow into the different reports when filtering takes place in order to better circumscribe the metric context.

Copy and Paste a Report

1. Enter Edit Mode
2. Select the report to copy
3. Select the **Copy the Selected Report** Icon
4. Navigate to My Reports*
5. Select the Paste icon
6. Report appears in the My Reports tree
7. Revisit upstream reports and copy filter expression



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

19

The simplest way to get a report ready to edit is to copy an already existing report from the report tree. The process is quite easy.




1. Enter Edit Mode.
2. Select the report to copy.
3. Select the **Copy the Selected Report** Icon.
4. Navigate to My Reports.
5. Select the Paste icon.
6. Report appears in the My Reports tree.
7. Revisit all upstream reports and take a copy the filter expression that you must bring at the top of the copied report.

* Take note the any pasted report will be placed as a child report, right under the selected destination report.

Pinning a Report

In Browse Mode, hover on a report:

1. Select the push pin icon for the report
2. Click on **New Personal View**
3. Enter a new personal view name
4. Report will appear under My Reports branch of the tree



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

20

Using the pin icon in a report will create a report entry under “My Report” (your report sandbox) and copy the selected report definition into it.

Take note that this creates a link to the report, it does not bring a full copy of the report definition. You can change report name but not the linked one. The benefit of doing this way is that if the original report gets updated, this will be reflected in your custom view as well. This is providing that the report has not moved under a different level.

Pinned Reports are Linked to a Template

The top screenshot shows the 'My Reports' tree on the left with the path: **My Reports** > **Pinned Report** > **Enterprise Capacity Dashboard** > **Raw Capacity Usage**. The right pane shows the 'Raw Capacity Usage' report in 'Edit Mode'. A section titled 'Link to Another Node' contains a message: 'This node is linked to another node of the tree, and thus cannot be edited directly.' Below this, the 'Linked to' field shows 'Raw Capacity Usage' with a link icon, and the 'Hide if Broken?' field has a question mark icon.

The bottom screenshot shows the '<<Templates>>' tree on the left with the path: **<<Templates>>** > **Alerts Summary** > **<<Devices>>** > **<<Storage>>** > **WWN Linker** > **VPLEX WWN Linker** > **Host Linker** > **Enterprise Capacity Dashboard** > **Raw Capacity Usage**. The right pane shows the 'Raw Capacity Usage' report in 'Edit Mode on ReportPack "Default ReportPack", template "Templates"'. It features tabs for 'Filtering & Expansion', 'Report Configuration', 'Report Details: TopI Graph', and 'Formula'. Below these are 'UI Wiring' and 'Pre Generation' tabs. The 'Name' field is 'Raw Capacity Usage', the 'Unique Identifier' is 'W4N-ENTERPRISE-RAW-CAP', and the 'Filter' field shows an error: 'partype is not present'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts


21

As noted on the previous page, the definition for the pinned report is not entirely complete. Some of the definition will link to other reports and cannot be edited as shown here. The top two captures show the tree for the pinned report and the edit pane for the Raw Capacity Usage branch. Note that the node name for this branch is in italic indication that it is a link, not a definition.

The actual source definition is in the All>><<Templates>>><<Storage>>>Enterprise Capacity Dashboard>>Raw Capacity Usage branch of the report tree.

One thing to note, as mentioned previously, if this file is moved to a different branch of the tree the link in My Reports will be broken. It is not best practice to move any of the template files however as they may be linked to numerous reports.

Tools - Export Report (template)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

22

When in Edit Mode, it is possible to export a report tree structure.


You have two choices:

- **Report Template:** Exports the node and its sub-nodes into an xml file ready to be imported by another ViPR SRM M&R user. Linked report are not resolved, they are kept as is, pointing to referenced report.
- **Report Definition:** Exports the node and its sub-nodes into an xml file, like the “Report Template” feature, but includes all linked reports resolved. This may create a bigger report description file.

If you migrate report set within the same portal, use the Template option. If you migrate report externally to another ViPR SRM M&R Frontend, use the Definition option. You might need to edit the structure and re-build the links as this later option removes them and this inflates the file containing the report structure.

Doing some exports may also be used to start creation of a ReportPack. The later groups report templates in a convenient package and that can be managed by an Administrator. You can then have reports made available on the root report structure and better control users or group of user access the reports (edit and/or just browse (read only)). Refer to ReportPack administration for more details.

Import Report (template)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

23

Once you have Exported a report template (or definition), you might want to import it into either another user's **My Report** sandbox or under a ReportPack.

Above illustrates how to import a report definition under "My Report".

1. Select a report node under which you are interested and where you want the report definition to be imported.
2. From **Edit Mode**, select "Settings" from top right menu bar.
3. Select the **Custom Tree** tab.
4. Select your file using **Choose File** button.
5. Click save and the report definition will show up under the selected report node.

Module: Reporting Key Concepts

Lesson Summary

During this lesson the following topics were covered:

- Introduction to the Edit Mode features
- Scoping the report using filter expressions
- Parent → Child node structure concepts
- Copying pre-existing report templates into My Reports
- Pinning reports vs copy and paste



This lesson covered the use of the filter wizard and editing filter expressions. Following filtering, was a discussion of report organization in the iterative approach and the parent/child node structure concept beginning with the static report structure. Finally, was a discussion of the use of the copy and paste feature to copy existing report templates into the My Reports tree branch as well as a quick review and comparison with the report pinning feature.

Lab 2: Organizing A Report Structure

This lab covers table report creation

- Part 1 – Set the starting point
- Part 2 – Practice report structure
- Part 3 – Pin a report
- Part 4 – Copy and paste a report

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

25

This Lab 2 allows you to practice table report creation.

Part 1 – Initialize “My Report” section: per module, per lab, reports. Good practice in organizing “My Reports” section.

Part 2 – Familiarize with report structure creation: parent – child – sub child – sub sub child.

Part 3 – Pin a report ... notice the difference in edit mode

Part 4 – Copy and paste a report: list of HBAs, Usable Capacity by Pool report

Lab 2 Debrief: Organizing A Report Structure

This lab exercise allowed you to:

- Familiarize with reporting structure and their construction
- Discover the difference between “Pin” a report and “Copy and Paste” a report
- Practice report structure layout

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

26

This lab exercise allowed you to:

- Familiarize with reporting structure and their construction
- Discover the difference between Pin a report and Copy & Paste a report
- Practice report structure layout

Module: Reporting Key Concepts

Lesson: Expansion

This lesson covers the following topics:

- Expansion concept and its effects in reporting
- Expansion modes: simple vs complex
- Effect of group expansion



This lesson covers the concepts of expansion and its effects in reporting including group expansion. Also covered are simple versus complex expansion modes.

Expansion Concept

Definition

Expansion is a report engine concept that allows you to group metrics by distinct values of a given property name.

- Result of expansion = report nodes
- Expansion splits the metric context: 1 per report node
- Quick example: collector of 5 metrics for each of 4 devices:
 - ▶ Before expansion: metric context of 20 metrics
 - ▶ Expansion on device property = 4 report nodes, 5 metrics per context
 - ▶ Expansion on name property = 5 report nodes, 1 metric x 4 devices per context



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

28

Expansion concept in ViPR SRM is key. Basically it splits the metric context in groups. Each metric context has the characteristics of having the same property value (from the expanded property in use).

The output of expansion results in generating report nodes: 1 per property value.

Simple example: Say you have a collector that recovers 5 metrics from 4 devices. Before expansion, the metric context holds a total of 20 metrics. Let's also say that each single metric is tagged with (at very minimum) metric name and device name. This way looking a given metric, we know what this metric measures and from which device it belong to.

Performing expansion on device property generates 4 report nodes, 1 for each device property value found in the original metric context. The resulting metric context holds 5 metrics, this for each of the report nodes.

Performing expansion on [metric] name generates 5 report nodes, 1 for each [metric] name property value found in the original metric context. The resulting metric context holds 4 metrics, all of the same name, but each are linked to a different device.

Static vs. Dynamic Description

- Static Report:
 - ▶ Specific to an instance or group of instances
 - ▶ Filtered manually by desired instances
 - ▶ Does not use expansion feature
 - ▶ Stays the same between Edit and Browse modes
- Dynamic Report:
 - ▶ Metric context driven
 - ▶ Report for each instance is presented automatically
 - ▶ Makes use of expansion feature
 - ▶ Defined in Edit mode, transformed in Browse mode




There are two implicit types of report node in use with ViPR SRM: **Static and Dynamic reports**. This should not be confused with pre-generation feature, two separate subjects.

When laying down a report structure, you will make use of both type of reports. The difference is in the behavior at browse time. The **Static Report** does not change the report based on the metric context. No matter how many elements are selected from the database, this report stays the same format, one single report node is generated. Filter expression is required in order to select the metric context. This way you can select metric of a specific device, group of devices, type of objects, etc. You can build a complete static report structure but it may require constant maintenance as new elements are added to the collection. Best used in high-level of report structure, less use for drill down types of structure.

The **Dynamic Report** is driven by the metric context. The report name will be replaced by a series of reports, based on the context. This type of report makes use of the Expansion feature. This will automatically divide (or group) metric context by the expansion expression and then to generate a report node per group - very convenient for report structure that depends on the data collected. When defined, in edit mode, only a single entry is required. Then in browse mode, the report is replaced by the result of the expansion. Each report is then automatically filtered.

Expansion

- Dynamically create report nodes
- Each nodes adds a level to a report that allows drilling down
- Use expansion to automatically create nodes based on properties found in the data set



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

30


Expansion dynamically creates report nodes. Such report nodes add levels to a report so you can drill down. When a node is set to expand, the node transforms itself from a single node to as many nodes as there values from the specified property name in the data context. By using expansion, you can define a report that automatically creates nodes based on the properties found in the data set (of the data context). For example, as hosts are added to the infrastructure, they are automatically added to the report when their properties match the expansion criteria (ie: hostname).

In the example shown here, notice that the expansion node (Device Type node) becomes a list of device types. If any new device types are added in the future, this report would add the new device type automatically. It is critically important that you recognize where the expansion is performed.

Filtering to Control Expansion

Filtering can be used to control expansion:

- In this example, the parent node has been filtered for devtype is Array **or** devtype is Host



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts


31

As you can see, applying a filter to the parent node will limit the expansion result (ie, set the metric context) to only those device types that pass the filter. Actually, the filter can be applied to the parent node as shown or in the child node and the resulting output would be the same.

Taking Expansion to Another Level (device)

Add another level to expand on device names (device property):

1. Add a new node called device
2. Expand on the **device** property
3. Save



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts


32

By adding another dynamic node using the expansion on the device property, the result is that each already expanded device type is itself expanded to become a list of all devices of that type. Shown here, you can see that the **Array** device type has expanded to include a list of all devices that are of type Array. Notice the blue arrow next to Host. If it is clicked, a list of all hosts will be shown. As you should begin to see, it is possible to create a multi-layered report that allows you to drill down into the inventory to see the complete set of data about your infrastructure as desired.

Taking Expansion to Another Level (parttype)

What types of parts would you find on an array?

1. Add another node under device and name it parttype
2. Add the expansion property **parttype**
3. Save



Note:

parttype = object type
part = single object entity

The answer depends on the type of array

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

33

With this expansion, you can begin to see that as we drill down, we can uncover the differences that exist between elements in the infrastructure and do not necessarily need to know what those differences are. We can use reporting to uncover them for us.

As we have stated before, but worthy of mentioning again, expansion can be compared with a GROUP BY statement of an SQL query: selected metrics (after filter) are grouped together into smaller sets according to a given property name. The list of nodes created at browse time is the list of distinct instances of that property within the data set. Also, for each report node created there is a filter implicitly set to match `property_name='[property value]'`.

Advanced Search Revisited (1)

BROWSE MODE

My Reports

Adv Search like report

source → APG_HEALTH

devtype → RSC-LunMappingDetection

device → SRMRHEL1

device → SRMRHEL2

device → SRMSM2

device → SRMSM3

devtype → VMAX-Collector

device → Array

parttype → Access

parttype → AccessToLUN

parttype → Controller

parttype → Disk

part → DF-7A:C 0

part → DF-7A:C 2

name → Availability

name → Capacity

drill-down

My Reports / Adv Se

July 2014, Wednesday 16 » Thursday

• APG_HEALTH

• RSC-LunMappingDetection

• VMAX-Collector

• VMWareCollector

• VNXBlock-Collector

• VNXFile-Collector

List of sources

Report generated in 0.249s.


EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Reporting Key Concepts 34

Remember the drilldown report example you created from Advanced Search using Common Convention? The result of the Advanced Search engine was that the drilldown structure you created cannot be transformed into a report node. Here you have everything on hand to have this report created, using expansion and report hierarchical structure.

Recall of the used common convention for drill-down: source devtype device parttype part name. The resulting report nodes is as demonstrated above.

Advanced Search Revisited (2)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

35

How do you create this report? Simply create a parent-child structure, one per property of the common convention.

The above structure makes use of squared brackets to indicate the node is dynamic and uses expansion, this is left up to you, not mandatory and not universal convention.

Indicated for each report name, the expansion property is in use.

Using “Filtering & Expansion” tab, you can successively define the structure and set the expansion level, as indicated. Returning back to browse mode, you will see the list of available items from the database.

Access to Databases

Expansion type: Simple Complex

APG COMPLIANCE EVENTS VMWARE VNX WHATIF

Selected properties:

Available properties:

- ☐ acknowledged
- ☐ active
- ☐ category
- ☐ certainty
- ☐ closedat
- ☐ count
- ☐ device
- ☐ devtype
- ☐ duration

Description: Move the cursor over a property to display its description here.

Timeseries database

Event type databases

- Timeseries Properties
- Event data fields

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

36

With ViPR SRM, there are many datasources to select properties from. One must differentiate the timeseries from the event. Mainly you have:

- Timeseries are the most common and event if data is spread across multiple timeseries databases, Report Engine group them under the APG tab (above).
- Event databases are separated by application: Compliance, VMware, VNX and (generic) Events.
- WhatIf database is a separate database used to keep temporary data using the What If Scenario feature (available under some SolutionPacks).

The reason why there are more Events databases is mainly because of the available type of fields. Not all system types provides the same level of details for a given event. As an example, VMware events are different from compliance. With VMware, you can have task-related events and system-related events. As opposed to compliance that generates “Breaches”, or configuration issues notifications (with a level of severity). And those two are different from the generic ones, found under Events database. They mainly differentiate by their list of fields.

Depending on the type of SolutionPacks deployed on your ViPR SRM setup, there will be a variable number of events databases. Report Engines must be authorized to access the content of those databases and a ViPR SRM Administrator should set this up at deployment time. One of the ways to discover which events database is available to the Report Engine (so you can workout reports from them), is to look at the expansion property selector menu.

Limiting the Scope

The screenshot displays the EMC Reporting interface. On the left, a sidebar shows the 'BROWSE MODE' section with a tree view of reports. Under 'My Reports', the 'Adv Search like report' is highlighted. The main area shows 'My Reports / Adv Search like report' with a list of reports: 'VMAX-Collector', 'VNXBlock-Collector', and 'Watch4NetSnmCollector'. A green box labeled 'Filtering for Event context' points to the 'Filter' field in the 'Reboot & PoweredOn/Off' report's 'Filtering & Expansion' tab. The filter is set to 'devtype is Array or devtype is FabricSwitch'. The 'EDIT MODE' section shows the 'Reboot & PoweredOn/Off' report with the same filter applied. The footer includes the EMC logo, copyright information, and the module name 'Reporting Key Concepts'.

Like in Advanced Search, you can return to top level and set a filter to limit the scope of items presented in the defined report.

In the above, we limited the scope to Arrays and Switches. So right on the top, we set the following:

`devtype='Array' or devtype='FabricSwitch'`

The result in Browse mode shows the sources where such device types are collected from and you can drilldown and unfold to confirm you have the expected scope of metrics.

Note that if you are preparing a report that makes use of events data, you have the same concept. Above, "Reboot & Powered On/Off" report (edit mode) is meant to limit the event data selection to those that are related with the intended report.

Grouping with Expansion

The screenshot displays the EMC2 reporting interface. At the top, a green oval labeled 'BROWSE MODE' points to the 'My Reports / Array's port' section. This section shows a list of reports under 'My Reports', including 'Array's port and directors listing'. Below this, a yellow oval labeled 'EDIT MODE' points to the 'Expansion' configuration area. The 'Expansion' section includes a message: 'The specified properties will be used for the expansion of the current node'. It features an 'Expand on every' dropdown, an 'Add a Property' button, and a 'Switch to Input Field' button. Under 'Group 1', two properties are listed: 'Simple: director' and 'Simple: port', each with edit, delete, and sort icons. The bottom of the interface shows the EMC2 logo, copyright information (© 2015 EMC Corporation), and the module name (Module: Reporting Key Concepts) along with the page number (38).

With expansion, you are not limited to expanding on a single factor (property).

There are multiple cases where this feature is applicable, and not just for expansion. Some reports require usage grouping functions to show up correctly. Typically, treemap, heatmap, and stacked bars could make use of grouping expansion.

This example shows expansion on ports and directors of a VMAX array.

The grouping is done using the comma as a property separator, if you're using "Switch to Input Field" option and manually type in the properties.

If you use the property selector "Add a Property", you must select the checkbox "Selected properties must be grouped".

Grouping – Usable Capacity by Pool

BROWSE MODE

Pool	Array	Pool Type	Pool Used Trend	Usable (GB)	Pool Used (GB)	Used for Block (GB)	Used for File (GB)	Used for Virtual Storage (GB)	Pool Free (GB)	% Pool Used
BC_FC_RAID1	000192606327	Thin Pool	10.51 ↗	50.52	21.08	21.08	0.00	0.00	29.43	41.74
BC_SATA_R6	000192606304	Thin Pool	14.08 ↗	50.52	18.98	18.98	0.00	0.00	31.54	37.56
Pool 0	myVIXsystem	Storage Pool	342.67 ↗	1,071.84	342.67	126.67	0.00	0.00	729.17	31.97
BC_FC_RAID1	000192606304	Thin Pool	11.21 ↗	50.52	6.32	6.32	0.00	0.00	44.20	12.51

Part **Device** **Pooltype**

Enterprise Capacity Dashboard

- Raw Capacity Usage
- Usable Capacity
- Used Capacity by Purpose
- Usable Capacity Trend
- Usable Capacity by Service Level
- Usable Capacity by Pool
 - Pools Expansion**
- File Capacity Dashboard
- Virtual Storage Dashboard

Expansion

EDIT MODE

The specified properties are used for the expansion of the current node

Expand on every

Add a Property

Group 1

Simple: device

Simple: part

Simple: pooltype

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Reporting Key Concepts 39

One application of a report using property grouping in expansion is the Usable Capacity by Pool report, found under Dashboards (see source below).

The report context first filters for Block storage device and also for Storage Pool (both defined as parttype). Then the expansion used to generate the table report lists items uniquely per (1) device, (2) Pool name (part), and (3) pool type.

The report template shows that the report is configured to do expansion over three properties and the result shows one row per group, along with the numbers.

Later, when studying advanced report types, we will revisit how grouping applies. Some advanced report types require grouping in order to correctly show results.

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Usable Capacity by Pool

Expansion Effects

Expansion generates an **implicit filter**

Component types list	Component count	Filter
LUN	868	* & source=="VMAX-Collector" & devtype=="Array" & device=="000198700604" & parttype=="LUN"
StorageGroupToLUN	151	* & source=="VMAX-Collector" & devtype=="Array" & device=="000198700604" & parttype=="StorageGroupToLUN"
AccessToLUN	146	* & source=="VMAX-Collector" & devtype=="Array" & device=="000198700604" & parttype=="AccessToLUN"
MetalMember	128	* & source=="VMAX-Collector" & devtype=="Array" & device=="000198700604" & parttype=="MetalMember"
Disk	120	* & source=="VMAX-Collector" & devtype=="Array" & device=="000198700604" & parttype=="Disk"

As you drill down expanded reports, filter expression adds up

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Reporting Key Concepts 40

The effect of using expansion is that it implicitly creates a filter on the report node, based on the expanded property and its value (at browse time). This filter defines a metric context for the report. The global filter (shown above) is built from previous parent report nodes (compound filter expressions) down to the current report node.

The above example is report defined from a successive expansion (edit mode), and as follows:

per source -> per device type -> per device -> per parttype -> per part -> per name

The resulting report (browse mode) shows drill down up to per device level, showing a list of available parttypes. At this level, for each row, the global filter select metrics that will match:

source is "VMAX-Collector"

AND device type is "Array"


AND device is "000198700604"

AND {for each row} parttype is {row's parttype}

This behavior is important to recognize. In mode advanced reporting, the only way to change this filter is by using complex expansion (split type).

Note: the initial "*" means "from all metrics found in available databases", that equals to a filter that selects "Everything".

Complex Expansion – Split Type



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Reporting Key Concepts

41

Let's introduce this concept of "Complex Expansion" - one of the ViPR reporting capabilities. This advanced feature differs in extending the possibilities of (simple) expansion by manipulating property value. In one of complex expansion features, there is the split type. Split type is used to create metric cross connections, allowing to redefine the metric context, thus allowing you to move from a given data collection domain to another.

The split type of complex expansion is considered as a table join in data sets. Let's imagine that Data set A is from Storage Array collection and Data set B is from Host collection. At a certain point in the report context (because of filtering), you have selected metrics from Data set A and now you are required to join metrics from Data set B. In order to perform this operation, you'll require a common property. Some properties in any data set are designed for jointure. Complex expansion (split type) will do the operation for each metric in Data set A to match the common property value in another data set. The result will be a new metric context, common to both data set, that is the intersection.

Complex Expansion – Split Type



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

42

The concept is as follows: remember that any report has a metric context. This could be Array's LUNs, Disk, Pools, etc. or VMware Hypervisors, etc. or FabricSwitch Ports, etc. You now want to define a report structure that will drill down from a given metric context to another context. The pre-requisite is that, in between, there must be a common (or shared) property value that can act as a link.

Simply put, from a given metric context you want to join data from another domain of metrics. Typical example is to link Array's LUN metric context with Hosts Disk of VMware, thus allowing you to create an end-to-end report structure drill down.

There are multiple flavors of complex expansion where you can manipulate property value in order to recreate the metric context. The above describes the complex expansion using partsn property value matching from VMAX Collection Domain to VMware Collection Domain. Actually, what complex expansion does is to recreate a whole new metric context where each given property value on the source side match the same property value in a new context (local or global).

The starting point of the above is on the left hand side. From a given context (VMAX), we isolated a LUN number (0580), defined in a report structure (filter and expansion). From the property list of this LUN, we chose partsn and ask the report engine to find any matching partsn value counterpart throughout the whole database... we ended up with a metric (and many others!) from VMware context, especially from Disk collection. Report Engine can now generate the report based on this new context, thus showing the above hypervisor host in the list of related hosts.

Partsn property is a common property often used throughout various SolutionPacks in order to relate objects. Collection layer is responsible to mark this property value.

Complex Expansion – Simulation Demo

1. Source metric context:

```
devtype=='Array' & device=='888973588710' &  
parttype=='LUN' & part=='0580'
```

2. Look at **partsn** value

3. Filter of the new metric context:

```
partsn=='60000970888973588710533030353830'
```

4. Explore the new metric context content

Reference:

All>>ViPR SRM Reporting class>>Module 03>>Demo -
Cmplx expansion on partsn (array 710)

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

43

The Management of Database Metrics tool uses the above steps to view metric contexts.

Use “Properties to show”: source devtype device part name. Also click on a metric to view the property card.



The **Display Policy** dictates the rule of presentation for the current report node (at browse time). Some report nodes can be hidden and used only for some background computations (using formulas, as an example). The report will not show if expected data is not available from the current context. In between, there are multiple possibilities as follows:

- **Always show the node** is the default behavior: the report node always shows in the tree.
- **Hide if no variable is selected**: hide the node if the filter does not select any metric (if resulting context is empty).
- **Hide if the node is a leaf**: hide report node if it does not have any child report. This is especially useful with expanded nodes, as they may expand into an empty list (from an empty context!).
- **Hide if no variable is selected or if the node is a leaf**: is a combination of the previous two options.
- **Show this node in the report but not in the tree**: Report can be drilled-down from the report pane area but will not show in the report tree. This use case when expansion exposes a long list you don't want to show up in report tree.
- **Always hide this node**: never show the node. It is useful to temporary disable a report without removing it. This also disables formulas applied to this report node. This is convenient when developing reports and for testing. Set at the parent level, this option is useful when defining report templates that won't show at browse time, but will be referenced from other reports.
- **Hide except when it's the target of a link**: hide the report node, except when it is the target of a link.
- **Show this node only in the tree and not in the report**: forbid drilling-down to this node from a parent report.
- **Hide in browse mode but use it for formula results**: is useful when defining a node that should be used for building reports above it only, but not as a report itself (for instance, when chaining formulas).

Use this Filter for...

The screenshot shows the EMC Reporting interface for the report 'Expand on partwnn (Cisco fabrics)'. The 'Filtering & Expansion' tab is selected. The 'Name' field is 'Expand on partwnn (Cisco fabrics)'. The 'Filter' field contains 'datagrps starts with CISCO' and 'datagrps=CISCO%'. The 'Display Policy' dropdown is set to 'expansion only'. The 'Expansion' section shows 'Expand on every' with 'Simple: partwnn'.

“Use the filter of this node for” option:

For the cases where you have to use a filter and an expansion on the same report node, you should ask the report engine about how to use the filter: should it be used in order to just generate the expansion, should it be used for narrowing down the current metric context or for both? This option is intended to be used only when performing Filtering AND Expansion on a given node. In any other case, it does not have any effect.

- **Expansion only** applies the filter while expanding the node using the expansion. This means that nodes that could have been generated through this field, but that do not match the given filter, will not be created. However, the metric context will not be affected by the given filter.
- **Selection only** applies the filter after the node expansion. It means the expansion is performed without taking the filter into account and after this, each resulting node is individually passed through the filter. The child node context is affected by the filter expression.
- **Expansion and selection** is the option that combines both previous options. The filter of the node is used for the node generation as well as each resulting node for the metric context it selects.

Source: All>>Explore>>Storage Connectivity>>Fabrics>>Expand on partwnn (Cisco fabrics)

Module: Reporting Key Concepts

Lesson Summary

During this lesson the following topics were covered:

- Expansion concept and its effects in reporting
- Expansion modes: simple vs complex
- Effect of group expansion



This lesson covered expansion concepts and modes.

Lab 3: Expansion In Reporting

This lab covers table report creation:

- Part 1 – Drill down collection inventory (using advanced search)
- Part 2 – Inventory of arrays (drill down)
- Part 3 – Inventory of switches (drill down)
- Part 4 – Inventory of hosts: split VMs from physical

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

47

This Lab 3 allows you to practice table report creation.

Use Advanced Search

Part 1 – Create report like Advanced Search, drill down the common convention

Part 2 – Report 1: drill down inventory of arrays (**set filters** (like devtype, then array models), structure your report) devtype=array & (parttype=X | parttype=Y) | !parttype (add 1 at a time, observe)... put !parttype is a separate branch named (metric directly attached to array)

Part 3 – Report 2: drill down inventory of switches (**same as above**)

Part 4 – Report 3: drill down VMs vs physicals (host and hypervisor)

Lab 3 Debrief: Expansion In Reporting

This lab exercise allowed you to:

- Practice filtering and expansion concepts
- Understand the metric context by using filtering



This lab exercise allowed you to familiarize yourself with the expansion concept, practice filtering and expansion concepts, and understand the metric context by using filtering.

Module Summary

Key points covered in this module:

- Edit Mode versus Browse mode
- Importance of Filtering and Expansion concepts

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Reporting Key Concepts

49

This module covered the key concepts of reporting as they apply to the ViPR SRM data set. Emphasis is given metric context by the use of Filtering and Expansion concepts.

This slide is intentionally left blank.



Module: Base Report Types

Upon completion of this module, you should be able to:

- Describe the base report types
- Create and edit the different types of reports



This module focuses on using base report types of ViPR SRM.

Module: Base Report Types

Lesson: Base Report Types: Tables


This lesson covers the following topics:

- Standard table
- List type
- Interactive table



This lesson covers table report format.

Tables Report Types



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

3

When turning into Edit mode, under “Report Configuration” tab, then under “Report Types”, you will see a complete list of available report format. In this lesson, we will explore Tables reports.

The Report Types are grouped into sub groups:

- **List**: the simplest report type, allowing you to list child report entries in a convenient format (up to 2 columns).
- **Standard table**: most used format, fully customizable table report. You basically define columns, filter the metrics, state the cell content format, etc.
- **Metric-based table**: like standard table and just list metrics available on the context.
- **Interactive table**: on the top level, defined as a standard table, item selected in parent triggers update of the child report. Adds interactivity to reporting.

NOTE: Whenever you create a new report node, it inherits the list report type by default.

List Report Type

The screenshot displays the EMC2 report configuration interface. On the left, a sidebar shows a tree view of reports: 'My Reports' > 'Adv Search like report' > '[Per Source]' > '[Per devtype]' > '[Per device]' (highlighted in blue). Above this sidebar is a yellow 'EDIT MODE' button. The main area is titled '[Per device] | Edit Mode' and contains a 'Report Configuration' tab. Under this tab, various settings are visible: 'Display Mode' (default display), 'Report Type' (list), 'Report Rendering' (static charts by default), 'Default Duration' (last 1 day), 'Sampling Period' (real-time), and 'Sampling Type' (average). A 'Propagate time selection on drill-down' checkbox is also present. Below these settings are links for 'Description and Properties' and 'Advanced Settings'. On the right side of the main area, there is a green 'BROWSE MODE' button and a preview of the report output. The preview shows the title 'Array / 000192606304', the date range 'July 2014, Wednesday 16 » Thursday 17, 3', and a list of storage components: Access, AccessToLUN, Controller, Disk, LUN, MetaMember, Port, Service Level, Storage Group, Storage Pool, StorageGroupToLUN, Thick Lun, and Tier.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

4

List reports are defined by default when adding a new report node in Edit Mode. It is a simple report, the content of which is a list of all the underlying children report nodes.

Unlike most other reports, there are no further parameters (found under Report Details, which grays out anyway) to configure when defining a report node to be a list report.

The above is a re-use one of the report that we visited when building up the report based on Common Convention.

Report Configuration

EDIT MODE

Report Configuration

Display Mode: default display

Report Type: list

Report Rendering: static charts by default

Default Duration: last 1 day

Sampling Period: real-time

Sampling Type: average

Propagate time selection on drill-down: ☐

Description and Properties

Advanced Settings

BROWSE MODE

FabricSwitc...

July 2014, Wednesday 16 x Thursday 17, 3:24 PM EDT | Last 1 Day: average on real-time

This is a report description. You can use standard HTML tags (table, list, bold, italic, fonts, etc).

Device model	Device description
MDS 9216A	Cisco SAN-OS(tm) m9200. Software (m9200-ek9-mz). Version 3.3(4). RELEASE SOFTWARE (fc2) Copyright (c) 2002-2005 by Cisco Systems, Inc. Compiled 7/2/2009 1:00:00

- FCS
- InterConnect
- Interface
- Memory
- NameServer
- Port
- PowerSupply
- Processor
- Sensor

List report

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 5

While we're using "Report Configuration", let's take a moment to look at other parameters in this tab.

Report Rendering: allows you to state in a Chart type report what is to be presented in an interactive format or a static format. Interactive allows the user to hover the mouse on the chart and tooltips will show values.


Default Duration: where you set the report's global period coverage. You can use the default and later, in report details, specify some override. Mostly apply to Chart reports. This will tell the time horizon when the Report Engine should recover data.

Sampling Period: the report's default data aggregation level to use when recovering data. Remember the defaults are real-time (as data been collected), 1-hour, 1-day, or 1-week aggregation levels.

Sampling Type: this option selected when multiple metrics must be aggregated together: the default is to perform an average of the values.

Description and Properties: adds some description to the top of the report (free text where you can use standard HTML tags). Also adds a nicely formatted list of properties (above using "model" and "devdesc", device description). The Report Engine automatically generates the format.

Table Report – Base Concept



EMC VMAX / Arrays Summary

July 2014 Wednesday 16 » Thursday 17 3:37 PM EDT | Last 1 Day

3 elements found, displaying all elements.

Array	Model	Engenuity / Microcode	Cache Size	Disks Count	Hot Spares Count	Unconfigured Disks Count	Devices Count
000192606327	VMAX20K	5876.251.161 date:09.03.2013 patch_date:09.03.2013	28 GB	56	8	0	829
000195901097	VMAX10K	5876.268.174 date:02.05.2014 patch_date:02.05.2014	42 GB	58	6	0	263
000192606304	VMAX20K	5876.251.161 date:09.03.2013 patch_date:09.03.2013	28 GB	56	8	0	831

Rows

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 6

When you create a table report in ViPR SRM, you need to define what will be displayed in both the rows and the columns.

Rows : Typically, table report rows are populated from metrics found on child report nodes. Thus, you have to create a child report node and configure it. Usually, the child report node is set to be the expansion of a property. In the above case, “device” property is used to list 1 array per row.

Columns : Columns display values, properties, and attributes. The table report’s columns are defined in the Edit mode, under “Report Details:Table” tab.

Note : When you create a new report node, this new node is by default a list table report with one column displaying the sub node name. You have to change report type from “list” to “standard table”.

It is IMPORTANT TO UNDERSTAND that what populates a table is from the Child report node (ie: from an expansion or from a static list of reports).

Source report: Arrays Summary (Report Library>>EMC VMAX>>Arrays Summary)

Table Report Defined


The screenshot above shows the same report as the previous page, but this time in Edit Mode. What is not shown but must be obvious by now: under Report Configuration, “Standard table” is selected from Report type.

We can see that there are several columns defined in the main section (Report Details: Table). These are called “Attribute : Device”; “Property : IP” and so on. There are basically **three types** of columns available in defining a table reports:


- **Attribute:** Display attributes about the report node. These attributes can be very useful for troubleshooting.
- **Property:** Displays the column to display a property value (ie: device, source, part, etc.)
- **Value:** Display a metric value that can be calculated over a specified time period.

Reordering and Sorting a Column

Reordering



Sorting



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

8

You can change the order of the columns simply by drag and drop action:

- **Drag and Drop:** drag a column to its new location and drop it there.

It is easy to delete a column - simply click the trash bin icon in the column's header bar.

You can specify a sort of the rows, based on a per columns. Use "Sort" button on top or column definition and from the pop up and simply drag and drop columns you want to use for sorting.

You can also take a copy of a column you have already defined. Simply use the check option on left of the column's header and click "Copy" button on bottom (notice the value of the Paste button). Move around to the destination report and hit "Paste" button.

Column Type: Attribute

Attribute: Array

The lines of a standard table report are defined using the direct children of the selected node. An attribute column displays the information related to this node, such as its name or number of children.

* Column Name: Array

* Attribute: node name (selected)

Advanced Properties

Attribute

configures the attribute to recover from the node.

- node ID displays the node short ID, relative to its parent
- global ID displays the complete node ID, starting from the tree root
- node name displays the name of the node
- global name displays the complete name of the node, starting from the root of the tree
- expansion names array choose this only if a multi-expansion is applied to the node that will appear in the array table. In the field box type the index of this zero-based array, that contains each part of the multi-expansion.
- node filter displays the node short filter, relative to its parent
- global filter displays the complete node filter, starting from the tree root
- child count displays the number of children the node has
- other displays the value of the property you type in the field box.

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 9

To display an attribute, there are two sections in this type of column:

1. **The attribute definition:** Column Name and attribute to display from the drop-down menu.
2. **Advanced Properties:** Display Conditions allows to hide or show the column and under some specific cases (ie: if device name match a specific regular expression).

Listed here are some uses for attributes:

- **Node ID:** report engine's internal report node ID.
- **Global ID:** compound internal report node ID from root down to present report.
- **Node name:** the report's name: if using an expansion, it is the property value. If using static report, the report's name, as defined under "Filtering & Expansion" tab. This is the most often used item.
- **Global name:** the compound names (equals to the breadcrumbs, or the path of the report).
- **Expansion name array:** reserved for special usage.
- **Node filter:** the filter expression on the current report node.
- **Global filter:** the compound filter expression: used mainly for troubleshooting reports.
- **Child count:** returns the number of underlying report nodes (if expansion is set further down).

Hint: hit the Question mark next to the Attribute in order to get details about the use of the different attributes.

Column Type: Property

Property: Model*

A property column displays a property about the selected variables of the child. This can be any collected property, such as a part, for example http, https or ssh, or the device, which could be the name of the server.

* Column Name: Model

* Property: model

Selected Value: use the filter below

Filter to Apply: Everything

Advanced Properties

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

10

Display a property in a column (per row).

1. The column's name: shown as the column header of the table report.
2. The property definition: choosing the property to display from either a list or typing it directly.
3. Selected value: either you are using a property from a formula or you want to be choose a specific metric using a filter expression
4. Advanced Properties: Other advanced properties, like Display Condition.

Column Type: Value

Value: Cache Size*

A value column fetches a value from the database. You can specify the value with a filter, or as a formula result, such as the amount of your CPU utilization, and can express the value in the format you choose, such as percentage or bits per second.

* Column Name: ?

Selected Value: ?

Filter to Apply:

? ?

▶ Time Management

▶ Value settings

▶ Advanced Properties

▶ Summary settings

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

11

Displays a metric value. This column type is the most flexible type of column. There are five sections in this type of column:

1. **The selected value:** where you set the column's name (the header). Then which value (metric or from a formula) you want to use for this column (chosen from the report's context). Either "use the filter below" or select from a formula result.
2. **Time management:** selects the aggregation level, sampling period and range to the columns.
3. **Value format:** sets the value format of the column (how this value will show up: units, format, threshold, etc).
4. **Advanced Properties:** other settings applicable for a value format.
5. **Summary setting:** summarizes the column (total or average), at the bottom of the column of the report.

Each value type column defined in a table report can calculate a value over selected periods of time. You can show a report for minimum, maximum, and average value of a single metric (3 columns) and over two different time ranges (thus 6 columns values to define).

Value: Time Management

Time Management

The Time Management section controls the aggregation and time periods of the selected values. If more than one time period is defined, the values will be represented by several columns in Browse mode.

Sampling Period: and

Sampling Type:

Column Time Range(s):

Recover...:

Temporal Aggregation:

Select aggregation

Select time period

inherit report time range

all the values in the time range

average

same as sampling type

use a fixed aggregation

average

min

max

sum

last

count

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 12

With ViPR SRM, the first step in choosing how to present results in is time range selection.

Sampling Period: the already computed temporal aggregates in the database from which you choose to use. This option is especially relevant when you use a longer time range. For example, with a month or year long period, selecting the day or week aggregate instead of the real-time, dramatically increases the speed of the report generation. With longer time ranges, fine grain detail (with a multitude of dots appearing along the graph) is not necessary. There are two options: “Allow higher periods to be selected” and “Enforce the specified period” allows flexibility for user experience. In fact, ViPR SRM may select a higher aggregate when it fits the report time-range. If you want the specified aggregate, and want to prevent ViPR SRM from making a change, then this option has to be selected accordingly.

Sampling Type: Type of aggregate you want to use for the sampling period. No aggregation is permitted with the **real-time** sampling period.

Column Time Range(s): There are three options: use “inherit report time range”, “use fixed time range” or “generate a column on the”. The first one uses default’s report time range definition whereas the second one is used to specify a specific time range independently overwriting the default report time range. The third one is used to generate multiple column time periods (more details to follow).

Recover ... There are two options, “the last value only” or “all the values in the time range”. The first one retrieves the last value whereas the second option retrieves all the values in the report time range. The second option results in several values, so you must specify an aggregation to calculate a single value to display in the table report. You do this by using the Temporal Aggregation drop-down list that shows up.

Temporal Aggregation: This option specifies the aggregation to apply to the values of the whole time range. Although you force one or you keep report’s default.

Value Settings

The screenshot shows the 'Value settings' configuration panel. It includes a descriptive note at the top: 'Sets up the format for displaying values and the specific thresholds for the column. Values may not automatically display themselves the way you want, have too many decimal numbers, not be in the desired unit or visual alerts must be triggered.' Below this, there are three main sections: 'Use Rounding' with a checked checkbox and a field set to '2 decimals'; 'Scaling Mode' with a dropdown set to 'by unit'; and 'Value Display' with a dropdown set to 'display all values'. On the right side, there are two unit conversion settings, each with 'Unit Type' and 'Unit' dropdowns. The first conversion is from 'Byte' to 'GB', and the second is from 'Byte' to 'Auto'. A green box labeled 'Rounding a value' has an arrow pointing to the '2 decimals' field. Another green box labeled 'Scaling mode:' has an arrow pointing to the 'by unit' dropdown.

Value settings

Settings up the format for displaying values and the specific thresholds for the column. Values may not automatically display themselves the way you want, have too many decimal numbers, not be in the desired unit or visual alerts must be triggered.

Use Rounding ☒ 2 decimals ?

Scaling Mode by unit ?

Unit Type Byte TO Unit Type Byte

Unit GB Unit Auto

Value Display display all values ?

Scaling mode:

- Supports division and multiplication
- Automate data conversion based on metric's units

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

13

Value settings section allows you to configure how the value will show up for each cell in the column. Among the available options, you have:

- **Rounding:** sets the number of figures after decimal point.
- **Scaling:** converts numbers from a given unit format to a new one. Not all systems use the same base unit format (KB, MB, GB). Report Engine can streamline this detail for you.
- **Threshold Definition:** the value in the cell will be presented with either an colored icon or a background color, depending on the value formatter you choose. This section allows you to define such thresholds.
- **Value Display:** allows you include or exclude some rows based on the threshold definitions.
- **Value Formatter:** there are multiple way the value can be printed out in the cell of your report.

Value Settings: Thresholds in Report

Value: Buffer Credit Lost*

EDIT MODE

Value settings

Thresholds Definition

Add a new threshold

Normal	Normal	-Infinity
Major	Major	0.001
Critical	Critical	0.5

BROWSE MODE

Switch	Slot/Port	Type	Buffer Credit Lost	Throughput (/s)	Avg. Utilization (%)
losar180	72	E-Port	1,153.68412	3.9 MIB/s	0.3
losar180	76	E-Port	1,135.86815	4.5 MIB/s	0.4
losar180	79	E-Port	168.10291	13.9 MIB/s	1.0
losar180	77	E-Port	0.12826	32.6 MIB/s	2.9
losar180	40	F-Port	0.09146	90.0 B/s	0.0
losar180	58	F-Port	0.01924	543.9 KiB/s	0.1
losar180	59	F-Port	0.01747	543.8 KiB/s	0.1
losar180	12	F-Port	0.00757	5.6 MIB/s	0.3
losar180	37	F-Port	0.00033	232.0 KiB/s	0.1

Value Display

- display all values
- display all values
- display Normal Severity and above
- display Unknown Severity and above
- display Minor Severity and above
- display Major Severity and above
- display Critical Severity

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 14

For each defined value column, there may be critical multiple levels that could trigger **visual alerts**. Threshold definitions are embedded under the “Value Settings” definition. To enable this feature, you can simply set a value in one or more of the Level fields.

Thresholds Definition: Thresholds define Severity ranges. Ranges are defined by their lower boundary value, their upper boundary being defined by the next value in the Threshold sequence. Boundary values are inclusive if the Threshold Severity is greater than its neighbor. Only the last definition is kept in case of duplicate values. Duplicate names may be used. Values are mandatory and must be in decimal format, including – Infinity and Infinity.

Values Display: There are multiple options (depending on the number of thresholds you defined) that allows you to list only a subset that matches a given level. For instance, by changing to “display major severity and above”, the report displays only rows that fit the major severity and critical level conditions.

Source: Report Library>>Brocade FC Switch>>Inventory & Performance>>All FC Ports Statistics>>FC Ports Summary

Here is a source report to look at that combines the use of Thresholds and Value Display features:

All>>Operations>>Situations to Watch>>Performance: High Traffic Utilization (%)

Value Settings: Value Formatter

Value Formatter

You can select here how you want to display the values of this column.

- Status:** Displays the status icon for the value according to the threshold levels that are set.
- Value:** Displays the value in numeric form and the status icon.
- Value with background:** Displays the value

Mapping (key, value)

key	value
200	OK

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 15

Value Formatter allows you to customize how the value in the cell shows. This gives your report consumers more detail about the behavior of the value over time. If the value in return is a time/date result, it can be interpreted and printed correctly. Click the question icon in order to get the full description of each available formatters.

There are 12 types of value formatters:

- **Status:** Displays only the status icon for the value according to the threshold levels that are set.
- **Value:** Displays the value in numeric form and the status icon.
- **Value with Background:** Displays the value and changes the background according to the thresholds.
- **Value with Graph:** Displays the value and status icon, and is clickable so you can display an inline graph popup plotting the value over the selected time period.
- **Date and Time:** Displays the date and time when the value was collected, converted from the UNIX timestamp, along with the status icon.
- **Time:** Displays the time when the value was collected, converted from the UNIX timestamp, along with the status icon.
- **Duration:** Formats a number of seconds into a duration.
- **Sparkline:** Displays the value and status icon along with an in-line graph, or “sparkline”, plotting the value over the selected time period.
- **Status Sparkline:** Displays the value and status icon along with an inline status bar indicating the status value over the selected time period.
- **Performance Sparkline:** Displays the value and status icon along with an inline performance bar indicating the performance of this value relative to the other values on the page for this column, over the selected time period.
- **Bar Chart Sparkline:** Displays the value and status icon and a bar chart showing the values relative to the other values on the page for this column, over the time period. The color matches to the threshold of the bar for the matching period.
- **Value to String:** Allows you to convert some metric value to string, using a mapping set (shows when the formatter is selected).

Value: Advanced Properties

- Per column display condition
- Complement value will do 1-value:
 - ▶ Example: you have the Uptime
 - ▶ If you want the Downtime, use the Complement value
- Replace empty values with zero
- Exclude data when covered by maintenance period
- Filter out rows based on filter expression

Advanced Properties

Contains advanced settings that are only useful in some v

Display Condition: show

Complement value? ☐

Replace empty values by zero? ☐

Use per-line maintenance period? ☐

Allow filtering on this value? ☐

Default filter expression: `$v >= 0.00 || $v < 0.00 || empty($v)`

Allow filtering on this value?

This triggers the use of a JEXL boolean expression that the value must match in order to be displayed. \$v is the variable used to represent the numeric value.

In the following examples, rows are only shown when

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

16

Advanced Properties allow flexibility in how values show up.

Display conditions could be useful for designing a report that will serve as template.

Depending on the situation, you can choose to display or hide a column... or you can let the report engine decide, depending on the metric context (conditional).

In some cases, when a report gets generated, some metric value selections will lead to a null selection (no data for the selected time range). Instead of printing an empty cell, it might be better to print a zero value.

If you are using featured Outage Manager, you might want to exclude data values from calculation based on a maintenance period. This is a nice feature to have for SLA reports. A device marked a for maintenance period will report data but a Report Engine will exclude them at report time.

Value filtering allows the row to be printed based on the value and the expression. Click on the question mark icon to get the complete details on the usage of this feature.

Filtering on a Value vs. a Metric

Value: Raw Capacity

A value column fetches a value from the database. You can specify the value with a filter, or as a formula result, such as the amount of your CPU utilization, and can express the value in the format you choose, such as percentage or bits per second.

Column Name: Raw Capacity

Selected Value: use the filter below

Filter to Apply: name is RawCapacity and part is not present

name=="RawCapacity" & !part

Filtering based on properties

Filtering based on values

Advanced Properties

Contains advanced settings

Display Condition

Complement value?

Replace empty values by zero

Use per-line maintenance period

Allow filtering on this value?

Default filter expression: \$v >= 0.00 || \$v < 0.00 || empty(\$v)

Allow filtering on this value?

This triggers the use of a JEXL boolean expression that the value must match in order to be displayed. \$v is the variable used to represent the numeric value.

In the following examples, rows are only shown when:

- empty(\$v): the value is empty
- (\$v >= 0 or \$v < 0) and not empty(\$v): the value is (greater or equal to 0 or less than 0) and not empty
- \$v > 10 or \$v < 5: the value is greater than 10 or less than 5
- \$v != 0 and not empty(\$v): the value is not equal to 0 and not empty
- \$v == 0 && not empty(\$v): the value is equal to 0 and not empty

For more details, see the JEXL reference.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

17

In order for you to have better control over what's going to show up in your report, a filter expression allows you to specifically select a metric from the available report's context.

More than that, you can also specify that if the resulting value is not relevant, you can discard the row entry. This is done using the "Allow filtering on this value", under "Advanced Properties" section of value's column definition. As an example, raw capacity values lower than a 10GB are not relevant for your report, just filter them out. Hitting the question mark icon next to the value filtering option will give you more details on the expression lingo (JEXL).

Summary Settings

Summary settings

EDIT MODE

These settings allow you to add a summary row to this column.

Label

Generate column

sum
-none-
sum
average

of all visible rows

BROWSE MODE

Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used
Platinum	246.97	94.95	152.02	94.88	38.45
Silver	18,415.68	16,564.22	1,851.47	16,419.62	89.95
Bronze	128,815.58	42,421.30	86,394.28	42,521.72	32.93
FASTVP	0.00		0.00		
Other	750.60	64.01	686.59	1,747.37	8.53
Gold	135,063.56	100,886.81	34,176.75	100,900.74	74.70
	283,292.39	160,031.29	123,261.10		

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

18

Summary settings allows you to print an extra cell at the bottom of the column that will either show the sum of the values or an average value.

Please note that the summary value will add up values for the current view only, not from all of the table rows. This means that if you have a table result of 100 rows, only first 25 (1st page) numbers will add up, not the whole 100 rows. A Feature request has been made for the next version to support a summary that will include all of the rows.

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Usable Capacity by Service Level

Laying out Table Report

EDIT MODE

Storage

- Storage Systems
- More Reports
 - Enterprise Capacity**
 - device

Report Type: standard table

Expansion: The specified properties will be used for the expansion. Expand on every. Add a Property. Simple: device

BROWSE MODE

Enterprise Capacity

Array	Array Model	Firmware	Array Type	Total Raw (GB)	Configured Including RAID (GB)	Configured Usable (GB)	RAID Overhead (GB)	Unconfigured (GB)	Hot Spare (GB)
0001957009	VMAX40K	5876.268.1 date:02.05.2014 patch_date:02.05.2014	Symmetrix	134,605.16	89,986.43	50,533.84	39,452.59	29,262.80	15,360.19
000195700999	VMAX40K	5876.268.1 date:02.05.2014 patch_date:02.05.2014	Symmetrix	134,605.16	89,986.43	50,533.84	39,452.59	29,262.80	15,360.19

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 19

A table report depends on a report structure of at least two defined report nodes. Remember that the content (rows) of the report is given by the child report while the table report is defined as a parent report. In a word, the tables content is populated by the expansion of the child report. Most of the time, you will have an expansion set on the child report of a parent that is a table report. In short, a table report ALWAYS requires at least two report levels.

In the above example: Enterprise Capacity is defined as a standard table and holds the column definition. Child report (device) hold the expansion on the property named device. That means for the table report to show up one row per array (expansion). Basically, this is where the metric grouping takes place since each array exposes metrics show in column (exception for some columns value comes from some calculation (formulas)).

Metrics an array exposes [Column name => metric name]:

Total Raw (GB) => RawCapacity

Configured Including RAID (GB) => ConfiguredRawCapacity

Configured Usable (GB) => ConfiguredUsableCapacity

RAID Overhead (GB) => RAIDOverheadCapacity

Unconfigured (GB) => UnconfiguredCapacity

Hot Spare (GB) => HotSpareCapacity

Unusable (GB) => (using calculation)

Free (GB) => FreeCapacity

Used (GB) => UsedCapacity

Used For Block (GB) => BlockUsedCapacity

Used For File (GB) => FileUsedCapacity

Used For Virtual Storage (GB) => VirtualUsedCapacity

[so on...]

Hint: use Advanced Search to see for yourself (and more details using Metric Management). Filter out using the following expression: [datatype=='Block' & !parttype] then set expansion on device and name.

Source: All>>Explore>>Storage>>More Reports>>Enterprise Capacity

Interactive Table

Service Level Capacity						
6 elements found, displaying all elements.						
Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used	
Platinum	246.97	94.95	152.02	94.88	38.45	
Silver	18,415.68	16,564.22	1,851.47	16,419.62	89.95	
Bronze	128,815.58	42,421.30	86,394.28	42,521.72	32.93	
FAST VP	0.00		0.00			
Other	750.60	64.01	686.59	1,747.37	8.53	
Total	135,063.56	100,886.81	34,176.75	100,900.74	74.70	
	283,292.39	160,031.29	123,261.10			

Bronze									
4 elements found, displaying all elements.									
Array	Vendor	Array Model	Service Level	Usable (GB)	Used (GB)	Free (GB)	Used Trend	% Used	
000195700949	EMC Corporation	VMAX40K	Bronze	85,886.70	31,668.87	54,217.83	31,627.38	36.87	
000195700999	EMC Corporation	VMAX40K	Bronze	40,270.04	10,669.25	29,600.79	10,703.73	26.49	
AMS200@10.247.55.161	Hitachi	AMS200	Bronze	95.92	72.88	23.04	72.88	75.98	
AMS2100@10.247.55.163	Hitachi	AMS2100	Bronze	2,562.92	10.30	2,552.62	10.30	0.40	
				128,815.58	42,421.30	86,394.28			

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types


20

An Interactive table is a “master table” where details are expanded in a child report (table, chart, etc) when you select a row. The “master table” kicks in the generation of the underlying report (whatever this report is).

In the above example, Service Level Capacity (SL) exposes capacity globally, per SL. On the click of a row (Bronze), the child report (Bronze) is exposed and details the capacity, per array.

This report layout requires at least three report nodes. First is the master table, defined as an interactive table (no expansion). Second, the report expands so the master table lists entries. From the above, the second report is also a table and depends on the third level to show up row entries.

Source: All>>Explore>>Storage>>More Reports>>Service Level Capacity



The above demonstrates the configuration of an interactive table from the Service Level Capacity report (see source).

The important configuration items are indicated and as follows:

Level 1 (master table): this first report node is defined as an interactive report and hold the master table columns definitions (as a standard table).

Level 2 (child report): this second report node is defined as a standard table (it could have been other), and holds the expansion for the master table.

Level 3 (sub child report): this third report holds the expansion that allow the second report to populate the table.

In short: Level 1 show list of Service Levels. Once a SL is selected, it triggers an update of the second report (Service Level Capacity Details). The later list arrays and details.

Source: All>>Explore>>Storage>>More Reports>>Service Level Capacity

Multi-Column

Array SN	IORate [12 AM]	IORate [01 AM]	IORate [02 AM]	IORate [03 AM]	IORate [04 AM]	IORate [05 AM]	IORate [06 AM]	IORate [07 AM]	IORate [08 AM]	IORate [09 AM]
888973588727	5,157.32	5,119.99	5,117.21	5,128.08	5,067.41	5,278.70	5,219.65	5,249.28	5,523.12	5,523.12
888973588777	723.86	696.92	692.02	692.81	716.88	701.79	758.53	699.08	724.37	724.37
888973588710	3,892.08	3,827.32	3,607.00	3,675.14	3,590.53	4,010.75	4,022.91	4,150.06	4,098.62	4,098.62
888976588482	288.98	587.61	735.35	292.73	291.04	306.77	295.61	658.32	755.24	755.24

EDIT MODE

The Time Management section controls the aggregation and time range of the data. If more than one time period is defined, the values will be represented by several columns in Browse mode.

Sampling Period: and

Sampling Type:

Column Time Range(s): for each (max 25 columns)

Recover...:

Temporal Aggregation:

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

22

Add some extra to your data table presentation: multi-column and background set to threshold value. Multiple Columns Table Reports allow you to easily view the evolution of a metric in a table.

Thresholds are found under the Value settings section of column definition. Under same section, you will have to choose “Background” value formatter in order to generate the nice cell background matching the threshold value.

Source: All>>ViPR SRM Reporting class>>Module 04>>Lab Extra 1 - VMAX IORate multi-column

Module: Base Report Types

Lesson Summary

During this lesson the following topics were covered:

- Standard table
- List type
- Interactive table



This lesson covered base table report types: list, standard, interactive.

Lab 4: Using Table Report Types

This lab covers table report creation:

- Part 1 – Adding a column to existing report
- Part 2 – Inventory Like Report – Table Drill Down
- Part 3 – Inventory Like Report – Flat Table
- Part 4 – Enterprise Capacity report
- Part 5 – Array performance (Read/WriteReq)
- Part 6 – Table Report using Event Data
- Part 7 – Multi-column

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

24

This Lab 4 allows you to practice table report creation.

Part 1 – Add a column to an existing report (must take a copy before!!)

Part 2 - Take a copy of inventory labs, turn list report type to standard table. Add columns: property, values. Do for arrays and switches reports.

VMAX: TotalDisk TotalHotSpare TotalLun TotalMemory TotalUnconfiguredDisk

Part 3 – Inventory like report revisited with switches: flat report table

Part 4 – VMAX - Table Report Capacity Report (metrics from array)

Part 5 – VMAX - Table report array' global ReadRequest and WriteRequest: use min, max, avg over 3 defined periods: last 6 hours, last day, last week. Use value formatter.

Part 6 – Events list table report

Part 7 – Table report with Multi-columns

Lab 4 Debrief: Using Table Report Types

This lab exercise allowed you to:

- Familiarize yourself with table reports
- Re-enforce report layout and expansion
- Familiarize yourself with some features of a table report
- Understand events data selection and display in a table report



This lab exercise allowed you to practice table reports, re-enforce the understanding of report layout and expansion, become familiar with some features of table reports and understand events data selection and display in a table report.

Module: Base Report Types


Lesson: Base Report Types - Charts

This lesson covers the following topics:

- Simple chart (graph) report
- Bar chart



This lesson covers base graph report formats: simple chart and bar.



Graphs are reports that display the selected metrics (the report context) or values computed through formulas over time. Various graph report types are available - the first two are the most often used types. In rare cases you would use children-based (last three).

When clicking on a report defined as a graph, a Report Engine:

- Applies context filter and select the metrics from the database.
- Looks for report's time range definition and recovers, for each metric, timeseries values.
- Plots one single timeseries per metric.

From the above, as an example, if report context recovers 10,000 metrics (because of a bad filter expression), there will be 10,000 timeseries on a single chart.

An exception to the above is if you use formulas on the report context (and selected "show in graph" option). This will be part of our discussion when looking into formulas.

Report Details: Graph – Metric Customization

Filtering & Expansion | Report Configuration | **Report Details: Graph** | Formula

Main Y-Axis Boundaries between and

Main Y-Axis Scaling Mode

Metric Customization

Allows color and shape settings to be customized on this report for each metric.

Match metrics using a

Configuration

curve

width

marker

color

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: Base Report Types 28

Once you select a graph report type, you will notice the Report Details tab changing, allowing you to select options for the selected type of report.

Among the available options:

- **Y-axis boundaries limitation:** Either you specify an upper and lower boundary or leave empty. The report engine will make use when rendering the graph.
- **Y-axis scaling:** You can simply scale the axis that changes the values. Beware that this feature does not change the values printed when using the dynamic chart options.
- **Metric customization:** This section of Report Details will expand, allowing you to customize the look and feel of the plotted lines (colors, thickness, marker, shape, etc). The options can also allow you the use right or left Y-Axis (double axis), and specify which metrics go on which Y-Axis.

Report Details: Graph – Legend

Filtering & Expansion

Report Configuration

Report Details: Graph

Formula

Main Y-Axis Boundaries

between

0.0

and

Main Y-Axis Scaling Mode

none

Legend

Configure which properties are displayed in the graph's legend to better identify the plotted variables.

Properties in Legend

Edit Properties

Switch to Input Field

Simple: part

Simple: name (Metric name)

Legend Visibility

always visible

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

29

Once you select a graph report type, you will notice the Report Details tab changing, allowing you to select options for the selected type of report. The Legend options allows you to set the information that will make the bottom labels speak. You can add as many properties as you want in the label. Usually, most contains two to three items. Otherwise, the label gets a long print name. Legend Visibility allows you to gray out items, when applicable.

Report Details: Graph – Thresholds

Filtering & Expansion | Report Configuration | **Report Details: Graph** | Formula

Thresholds

Thresholds must include a Severity (or a name for Custom Severity) and a unique value. Only the last definition is kept in case of duplicate values. Values are mandatory and must be in decimal format. Duplicate Severity or names may be used.

Thresholds Definition

[Add a new threshold](#)

Custom	Low	0.0		
Normal	Normal	0.015		
Custom	Medium	0.05		
Minor	Minor	0.075		
Major	Major	0.1		
Critical	Critical	0.125		

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

30

Thresholds Definition allows you to define markers in your graph (dotted horizontal lines). Or depending on the report type context, it will paint the bars of a bar chart using the defined threshold levels.

Report Details: Graph – Settings

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.


Module: Base Report Types

31

Summary Statistics allows you to set various presentation info:

- **Pinpoint** and **Boxed** are embedded in the graph and show some information about Max, Min, Average, and 95th percentile values.
- **Trend** adds a trending curve over a displayed metric (1 single metric at a time).
- **Logarithmic** uses Y-Axis log-scale.
- **Composite** allows you to display small graphs for different time ranges.

Static vs. Dynamic Charts



Under Report Configuration, you have Report Rendering options. Static charts plots dots and lines of a simple chart and uses a free crosshair pointer that does not stick to any value on the chart. The tooltip of this option is free to move.

Dynamic chart plots the same form of chart, but the pointer sticks (or snap) to the polled values, showing the tooltip with the metric name, date and time, and the exact value of the sample.

You can also turn from dynamic to static using the Display menu. In edit mode, you can configure both modes.

Graph Report Tips

Element	Usage and effect
Bar Chart	Helpful in studying polling cycle and its regularity
Dynamic Rendering	Helpful in pinpointing values in time and for area zooming
Children Base report type	Useful to drill down metric context
Logarithmic Scale	Helps to better display metric values of large scale variation (10 ² factors)
Composite Graph	Helpful for at-a-glance short and long term analysis
Metric Customization	Helpful for color blindness
Spline	Smoothen the shape of the curve on a graph
Second Y-Axis	Useful for metrics of different magnitude (arbitrary large values vs limited ones (IOPS vs %))
Legend	Adds complement of information in aggregated graph reports

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

33

Some tips from the pros of reporting:

- **Bar chart** is a report type you can use to momentarily look at the regularity of the polling cycle of a given metric. Another option would be to pull out the values using Export tools. You can also use browse mode to switch between simple graph and bar graph.
- **Dynamic rendering** is a feature of Report Configuration, under Report Rendering. This option is also available in browse mode.
- **Children-base report type** is sometimes used for study of the available metric in the child context. Handy to roll down the list of metrics and seeing if data is polled and give a quick view of the behavior of the metric.
- **Logarithmic scale** is sometime useful for showing large variations of a single (or several) metrics. If values differs from a 100 or 1000 (or more) factors between low and high values, consider viewing the behavior of such metrics using log scale.
- **Composite graph** is a feature allowing you to view in a single report short and long term behavior of metrics.
- **Metric customization** allows you to change the format of the plotted lines: type, thickness, markers, and colors. Sometimes reports must be made available to color blind people.
- **Spline** is where by default the report engine simply links the datapoints. If you want to smoothen the point, the spline option is handy.
- **Second Y-Axis** is often used for correlation and behavior of two metrics with different magnitudes. Typically, you'll find the raw values for which are not bounded (or limited) and other metrics are expressed in term of a percentage. If both refers to same Y-Axis, percentage will be squeezed as a flat line in the bottom of the graph, not very relevant/useful. This option allows you to keep them separate and improve your report presentation.
- **Legend** is by adding properties, you can improve the reading of the graph legend. This is especially true when graphing aggregated timeseries data (single metric from multiple objects from a device perspective). As an example: IOPS per Pool from an Array perspective. Default legend won't show Pool name, you must add manually.



Wondering about the regularity of the data collection? Here's a hint about how to check the polling cycle consistency. You can use this hint while in browser mode (recommended) instead of building a specific report for this task. This will work with any simple chart report, because the bar chart is a member of the report family.

A collector with a normal collection rate (polling cycle) will gather a value for each metric every five minutes (for the cases where polling cycle is set so). The 1-hour aggregate data will then summarize 12 samples per hour and the 1-day aggregate 288 samples. Among the aggregate fields, you have "count" which represents the total number of samples for the given aggregate. A quick way to check the consistency of the collection is to use the bar chart with sampling period set to 1-hour and use the sampling type to "count". When looking at the chart, all bars should be aligned to 12 for the whole report period. It will align to 288 if using the 1-day aggregate. If this is not the case, you can say that either the collector missed some polls OR the polling cycle is set differently.

If you build a report, don't forget to be very specific in using this report (specific filter) OR make sure to expand/drill down the structure. Otherwise, all metrics or the database will be aggregated in a single chart... potentially blocking/clogging the portal.


From the above, in exhibit 1, you can notice that count values from the 727 array is consistent over time (12 sample/hour), for the selected metric (IORate).

From the above, in exhibit 2, you can notice that some there was a missed polled data between 1PM and 2PM. You can also notice the polling cycle is every 20 minutes (3 polls/hour), for the selected metric (RawCapacity).

Hint: Bar Chart

- Bar colors go along with defined threshold values
- Single metric only
- Define thresholds in Report Details tab

Storage Performance / Global Disk Throughput



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

35

Bar chart bar color can also be used with threshold values. The color of the bar quickly helps identify which timeslot/bar is in which defined threshold. The above is a modified version of Global Disk Throughput report. It filters on WriteThroughput metrics and reports and includes thresholds values defined. Notice that this will work only if the report exposes data from a single metric. It cannot be used with more than one.

Note: be careful with the definition of the threshold values as you must refer to original scale of values, before applying scaling mode (multiply, divide or by units).

Source: All>>Report Library>>VMware vCenter>>Inventory>>Datacenters>>CSE 176 Lab>>Hosts Clusters>>CSE Cluster 2>>Hosts>>mqqb053.lss.emc.com>>Storage Performance>>Global Disk Throughput

Lab 5: Using Graph Report Types

This lab covers table report creation:

- Part 1 – Array performance: IORate, HitPercent
- Part 2 – Switch performance: per port
- Part 3 – VM's CPU and memory usage

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

36

This lab 5 allows you to practice chart report creation.

Part 1 – performance array, plot global IORate and HitPercent. IORate is “unlimited” while “HitPercent” is in %.

Part 2 – performance SAN switch, plot CurrentUtilization, per port

Part 3 – performance VM: processor CurrentUtilization, Active memory, usedCapacity (datastore)

Lab 5 Debrief: Using Graph Report Types

This lab exercise allowed you to:

- Familiarize yourself with graph report types
- Familiarize yourself with parameters of graph reports



This lab exercise allowed you to become familiarized with graph report types and the parameters of the graph report.

Module: Base Report Types

Lesson Summary

During this lesson the following topics were covered:

- Simple chart (graph) report
- Bar chart



This lesson covered different types of charts.

Module: Base Report Types

Lesson: Base Report Types: Mixed Reports

This lesson covers the following topics:

- Mixed reports



This lesson covers mixed report formats.

Mixed Report Types

Parent report - enforces parent parameters

Enforces child's parameters

Mixes timeseries and events data


Mixed



mixed reports



mixed using defaults



overlay

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types


40

Mixed reports are often used in global reports, like Dashboards. The main feature is to merge in a single report layout, multiple child report (Graphs, tables, etc). A mixed report allows you to display all the reports from the sub nodes onto a single page. When you run a mixed report, ViPR SRM processes all children reports to bring their data onto the mixed report's page.

There are three types of Mixed-Reports:

- **Mixed Reports:** generates a report using temporal information **from the selected node** (the parent).
- **Mixed Defaults:** uses all the **default settings of each child** (including temporal settings and bounds). This option can be overridden when using "Propagate time selection on drill-down" option.
- **Overlay:** used for mixing timeseries with events onto a single report.

Mixed Reports



EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

41

There are two types of Mixed reports:

- **Mixed Reports:** enforces (override) the child time-range selection and aggregation to the mix-report level. No matter how the child report's time-range and aggregation are defined, the parent report is enforced.
- **Mixed using defaults:** makes use of each child's configured time-range and aggregation when passing to the mixed-report level. Thus parent reports are not used.



This type of report does not require any parameters other than those found under Report Configuration tab (Report Details tab is grayed out). Sometimes this type of report is referred to report “pass-through” as it builds its report based on child report(s).


As you might already notice, this type of report will require more than a single node definition. The parent report defined as a Mixed report does nothing by itself, it requires child report(s) (at least 1).

The above example demonstrates the layout of File Capacity Dashboard. There are four child reports layout in a “three reports per line” format (Display Mode).

Use the Display Mode setting to configure the number of reports per row (horizontal) displayed. Let’s imagine the report page is laid out as a grid. The “default display” option depends on user’s settings. Using Browse Mode, you will be able to use the report icons to stretch or shrink and re-organize the layout more conveniently. As an example, the first report in left upper-hand corner can be scaled on two columns, thus taking 2/3 of the first row.

Source: All>>Dashboards>>Storage>>File Capacity Dashboard

Mixed Report Example



This is an example from the previous layout demonstrated. The parent report is set as a Mix Report while each child report is defined as an “Horizontal bars” report. As you can see, each child is a drillable report. That means you can get closer details out of each of them.

Source: All>>Dashboards>>Storage>>File Capacity Dashboard

Overlay Report



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

44


Overlay reports are used to graphically display event based information on top of a time-series graph. Events that occur at discrete moments in time are represented using markers. Events that have duration are represented lines lines or areas.

Events are also displaying information (tooltips) when hovering the pointer on a given event.

The graph report can be any graph or aggregated report that shows all the results in a single graph. You should consider filtering the table (event data) in order to have it correspond with the data shown on the graph.

The overlay report can have multiple overlays, each defining a different set of events to display. These appear in the graph in user defined zones.

Source: All>>Report Library>>VMware vCenter>>Inventory>>Hosts>>List of Hosts>>lglba180.lss.emc.com>>Performance CPU & Memory>>CPU



The Overlay report node must be a direct parent of 1) a graph report and 2) a table report (events).

Tip: define the report structure without applying the Overlay report at first, instead, use Mixed Report. Then define both the Graph and Event table report. Run the report and see if it makes sense. Pay attention to the content of the table: are those the selected events expected? Do you need to refine (filter)?

In the Overlay report details tab, use (+) Add an Overlay button in order to define a zone where selected events will show and how they will show on the graph.

[This slide describes how to create report structure laying out an overlay report – usually this requires a table report and a graph (set as overlay). Find an already existing report – describe options in setting up such report]

Source: All>>VMware vCenter>><<TEMPLATES>>><<ESX>>>Hosts>>ESX Device Summary>>Performance CPU & Memory>>CPU

Define an Overlay

Major → line [21 : 40]

1 An overlay defines a zone in the graph on which events will be drawn.

Display Title: Major

Selection Filter: EventType is ScheduledTaskFailedEvent or EventType is VmCloneFailedEvent

Boundaries: from 21% to 40%

Number of Lines: 5

Main Color: #ffffff

Display Level: top

Display Mode: line

Event Limit: 99

▶ Tooltip Information

▶ Momentary Events

▶ Durable Events

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

46

An overlay defines a zone in the graph on which events will be drawn. When defining an overlay, you have many options to set.

First, among the general options you have:

- Name of the Overlay and its associated filter: The Selection Filter is the expression that will limit the scope of selected event base on available properties for events. There is a long list of available properties for event, depending on the available events databases (VMware, Generic, Compliance, VNX, ViPR).
- Boundaries: this the region of the Y-Axis where you want this overlay to be drawn.

For the remaining options, refer to the question mark icon for further details. Lot of flexibility and control over how an overlay can be displayed, thus covering many use cases of event handling.

Define an Overlay – (cont.)

▼ Tooltip Information

Select the properties that will be displayed in any event tooltip.

Tooltip Title: ?

Additional Properties:

Simple:

▼ Momentary Events

Momentary events have no duration, and thus the same start and end times. They will be represented by a single point on the graph.

Events Color: ? 

Point Size: ?

Marker Type: ?

▼ Durable Events

Durable events have a duration, and may even be infinite. They will be represented by a line on the graph.

Events Color: ? 

Line Width: ?

Line Style: ?

Start Marker: ?

End Marker: ?

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

47

When defining an overlay, there are three other sections:

- **Tooltip Information:** define what information to be included in tooltip bubble when hovering the mouse on a event (in browse mode). You can select any properties available from the event context.
- **Momentary Events:** if event type is defined as a single moment in time (as opposed to durable events), you can specify how this overlay will look on the graph (color, shape and size).
- **Durable Events:** if event type is defined as durable in time (with a start and an end indicators), you can specify how this overlay will look on the graph.

Hint: It's better to categorize/separate event types per overlay: critical ones defined for a certain area, color and shape, different from informational types.

Module: Base Report Types

Lesson Summary

During this lesson the following topics were covered:

- Mixed reports



This lesson covered mixed reports.

Lab 6: Using Mixed Report Types

This lab covers table report creation:

- Part 1 – Revisiting VM report (lab 5 part 3)
- Part 2 - Array and SAN switch performance (from labs 5p1 and 5p2 in mix mode)
- Part 3 – Overlay report (VNX)
- Part 4 – Overlay report (VMware Hypervisor)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

49

This lab 6 allows you to practice report creation and using mixed reports

Part 1 – Revisiting VM report (lab 7), adding mix and interactivity

Part 2 – Mixing reports from lab 5 and 6

Part 3 – VNX Overlay report

Part 4 – VMware hypervisor Overlay report

Lab 6 Debrief : Using Mixed Report Types

This lab exercise allowed you to:

- Familiarize yourself with mixed report structures
- Improve your reporting skills by practice
- Mix timeseries and events data in a single report (overlay)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

50

This lab exercise allowed you to become familiarized with mixed report structure, improve your reporting skills by practice, and mix timeseries and events data in a single report (overlay).

Lab 7: Completing Reports

This lab covers table report creation:

- VMware VM list per customer/business unit + count# (iteration 1)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Base Report Types

51

Lab 7

This lab allows you to practice report creation.

This is iteration 1: no formulas. Count# is directly from child, no VM at the customer report level.

List of customers (# of BU per cust)

List of business units (# of VMs per BU)

List of VMs

Iteration 2 (next module), with formulas, will report count of VMs per customer (1 formula on the list of BU report node)... later.

Lab 7 Debrief: Completing Reports

This lab exercise allowed you to:

- Become familiar with VMware VM list per customer/business unit + count# (iteration 1)



This lab exercise allowed you to become familiar with VMware VM list per customer/business unit + count# (iteration 1).

Module Summary

Key points covered in this module:

- Base report types: table and graph
- Mixed reports usage
- Detailed configuration of report types



This module covered: base report types: table and graph, mixed reports usage, and detailed configuration of report types.

This slide is intentionally left blank.



Module: Formulas

Upon completion of this module, you should be able to:

- Describe formula concepts
- Integrate formulas into reports
- Discuss temporal vs spatial aggregation
- Apply application of formulas in reporting

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

1

This module focuses on presentation, use case, and practice using formulas in ViPR SRM reporting.

What is a Formula?

Characteristic
A reporting feature of ViPR SRM
Defined in edit mode, used in browse mode
Transforms metrics in KPI
Works with Timeseries and Events
Computes and aggregates data
Conveniently available in a formula library
Generates various outputs
Makes use of metric context and 1 st lower node level
Has dedicated tab definition
Multiple on a single node, serializable

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

2


In the reporting context, formulas have been introduced to enrich the feature set and processing capabilities of ViPR SRM. Formula is introduced and parameterized in edit mode but is in action only at the proper time in browse mode. Usually a KPI is the result of metric manipulation from the application of a formula. From a given context, formulas can transform dataset in multiple ways even if the data is timeseries or events types.

A formula can use many inputs and/or many outputs (depend on the formula in use). The formula library gives you the choice of a rich set of applicability. On a given report node, formulas are used and defined in a separate, yet dedicated, tab in the UI of Edit Mode.

The result of a formula can be used in another formula. Results are defined under the same tab or from the child report, so you can create a chain of data processing, by chaining results from one formula to another formula. Thus you can create powerful transformation.

Outputs of a formula are always available to the parent report (case of a table report) or directly on the same node (case of a graph report).

Formula – Use Case

- Use a report showing:
 - ▶ **UsedCapacity**
 - Given:
 - ▶ **Capacity**
 - ▶ **FreeCapacity**
 - Compute:
 - ▶ **UsedCapacity = Capacity – FreeCapacity**
 - Applicable to each single object (using expansion)
 - Thousands of other applications: 95thPercentile, summing capacities of all storage groups, forecasting, deviation, count items, etc.
- 

This metric is not directly available

Common metrics for Disk, SG, SP object types.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

3

As an introduction, here's a simple Use Case.

Many Storage objects instruments the Capacity metric which measures the capacity value each single instance. As an example, each single Disk has a measured value of its Capacity. Also, many have FreeCapacity metric (ex: Disks and Pools).

Now, let's say that you want to report on used (consumed) capacity, for Chargeback purposes. You will need to compute something based on what you have available: Capacity and FreeCapacity.

This will be set at the level of each single object. That means in report editing mode, you define the formula at the same level where the expansion takes place. At browse time, the report engine will apply the formula for each single expanded value available. This can then be used in a report listing the different objects (result of an expansion).

There are many other applications for formulas. This page is just a simple Use Case. Discovering the richness of the formula library will tell you about the possibilities.

Formula Library

Formula group	Sample list of formulas
Mathematical	Addition, Multiplication, Division, Ratio, Scale, Power, Percentile, Rate, Min, Max, StandardDeviation, Sum, etc. Spatial and temporal aggregation
Availability and SLA	Availability from uptime, number of reboots, downtime duration, Health Factor, etc.
Utilization	Bandwidth, Busy Hours, Conditional Value Count, StorageUtilization, etc.
Thresholds	Sustained Threshold, DevicePercentage analysis, etc.
Forecasting	Smoothing, Trending, Linear Forecast, Polynomial Forecast, ScenarioBased Forecast, etc.
General Purpose (Util)	Nop, Child Count, Event Count, Convert to Constant

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

4

The ViPR SRM standard formula library contains close to 100 formulas. Those are grouped by type, as shown above.

Mathematical: usual math functions. Some formulas compute on timeseries value only (temporal). Some compute over many metrics of the same type (from different objects) - these are then called “spatial” formulas.

Availability and SLA: those formulas analyze the given metric behavior. It can also computer SLA based on some criteria.


Utilization: computes utilization base on basic metrics.

Thresholds: those formulas report analysis of given metric based on a threshold setting. Basically, they count the number of occurrence of a given threshold.

Forecast: those formulas transform the given metric timeseries to compute various forecasting functions.

Generic purposes: those are tagged as ‘util’ formulas. They are used as general purposes and various usage.

Temporal Aggregation



- a.k.a Temporal Aggregation: per specified metrics and timeslot

EMC²


Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

5

How do formula work with datasets? The new generated result exists only at browse time AND is not permanent - not stored anywhere, this is in the report engine's memory. Its lifetime is on the report node and for the parent report node only.

Formula Concept in Report Structure



Applying formulas: “from the leaves to the root” concept

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

6

Several formulas can be used at different levels to modify metrics values (or generate KPIs).


The result of an application of a formula on a given report node is available on a parent node... thus “from the leaves to the root” concept.

It is then possible to take the result of a formula on a sub-node and display it, for example, in a table report one level higher than the node where the formula is applied. Calculation can also make reuse of a calculation done on the same node (cascading formulas).

Overall, formulas are applied from the leaves to the root. Results of formulas can be used on the node where it's applied, and on the parent nodes. But not the sub-nodes (not downward).

Applying the Concept

Example : using metric Utilization (%) from any FC port of a Cisco SAN Fabric, we want to factor by 66%. Calculations will be $(CurrentUtilization * 20) / 30$



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

7

This example demonstrates the “from the leaves to the root” concept. You can perform all those operations on a single node.

That said, let's say you want to see the result of this formula: $(Availability * 20) / 30$. We can see that there are two formulas to apply: Multiplication and Division.

The first formula to apply is the multiplication. Then, the division formula grabs the multiplication result (from the leaves to the root concept) and computes the division calculation.


The report is designed as follows (under Report testing yard):

- **Cisco Fabric Switch – 66% CurrenUtilization:** is a List report, no expansion, filter expression selects CurrentUtilization metrics of any FC Port of any Cisco FabricSwitch.
- **List of SAN Fabric:** is an expansion over device, thus it lists Cisco Fabric.
- **FC port list:** is an expansion over part, thus it lists FC ports attached to a given Cisco Fabric.
- **Result of CurrentUtilization *20 /30:** a nop formula used simply to forward the division result to the parent.
- **Division Formula:** a division formula taking as one of its inputs the result from the multiplication formula underneath and a constant value of 30.
- **Multiplication Formula:** a multiplication formula taking original CurrentUtilization metric as one of its input and a constant value of 20.

Please note that all of these could have been put on a single report node: FC port list report.

Source: All>>ViPR SRM Reporting class>>Module 05>>Lab 1 - Cisco FabricSwitch - 66% CurrentUtilization

Formula – Use with Graph vs. Table Reports



Depending on the report type you use, think about where to put your formula. If you plan to show the result of a formula on a graph report, you must place the formula that computes the final result on the same report node. Plus, you will have to check the “show in graph” box option of the final formula. Doing this has a little drawback: any original metrics on the context is replaced by the formula timeseries for which the option is selected (show in graph). As a workaround, simply recall the metrics you want to join on the graph by using a pass-through formula: nop (or nop4). Nop does “no operation” on the value, thus simply copy the values from the input to the output.

Remember that a table report is populated from the child reports. Child reports can be static report nodes OR, most of the time, the result of an expansion. If you plan to calculate a KPI (result of a formula) to show up for each single row, then the final formula must be where the expansion is, thus the child report node (where the parent is a table report).

The above re-uses the same example but in a more compact way. All calculations of 66% of CurrentUtilization is performed on the leaf [FC port list] report node. This report node expands the list of FC ports. Thus, for each port, a formula calculation will be applied. This report is a graph type. Thus, the “show in graph” option is selected. Plus, in order to be able to compare with original value, a pass-through formula is included, filtering on CurrentUtilization.

The [list of SAN Fabric] report is a table, populated by the child report node, just where the formulas are set. Thus, by simply defining a value column, and selecting the result from a formula, the column will retrieve the computed value, one for each FC port.

Source: All>>ViPR SRM Reporting class>>Module 05>>Lab 2 - Final - Cisco FabricSwitch - 66% CurrentUtilization

Applying a Formula - General

[list of SAN Fabric] | Edit Mode

The screenshot shows the 'Formula' tab in the 'Edit Mode' of a report titled '[list of SAN Fabric]'. The interface includes a top navigation bar with tabs: 'Filtering & Expansion', 'Report Configuration', 'Report Details: Table', and 'Formula'. Below the navigation bar, there are buttons for 'Paste (0)' and 'Add a Formula'. A toolbar shows 'All / Invert / None' for selection and 'Copy' and 'Delete' for actions. The main area is titled 'Formulas' and shows a 'Selected formula: Multiplication'. On the left, a list of 'Available formulas' includes 'MetricAvailability (availability)', 'MetricCount (util)', 'Min (math.spatial)', 'MonthlyAggregation (calendar)', 'MosAndRFactor (utilization)', 'MulPercentage (math)', 'Multiplication (math)', 'NodeMaintenancePeriod (util)', and 'Nop (util)'. The 'Multiplication (math)' formula is selected. On the right, the 'Description' section for 'Multiplication' is displayed, showing a tip to move the cursor over a formula to display its description. The description states: 'This formula computes the product of the given metrics. The first parameter (optionally a combined one whose components will be summed) will be the multiplied by the other parameter (optionally a combined one whose components will be summed)'. Below the description, the 'Parameters' section lists: 'first operand - This is the first operand of the multiplication.' and 'second operand - This is the second operand of the multiplication.' The 'Settings' section is also visible but empty. The bottom of the interface features the EMC2 logo and copyright information: 'Copyright © 2015 EMC Corporation. All Rights Reserved.' and 'Module: Formulas 9'.

To apply a formula: In Edit Mode select the report that needs a formula and click the "Formula" tab. Click the "Add a Formula" button and then select a formula from the library. Set the parameters, fill in the formula settings, and set the result(s). Click Save to validate your formula.

When selecting a formula (or browsing through the library), you always have complete details about the formula: general description, use case, and three sections: Parameters, Settings, and Results.

Applying Formula – Sections

The diagram illustrates the structure of a formula configuration in the EMC² interface. It is divided into three main sections:

- Parameters:** This section is titled "2 Parameters to Enter". It includes fields for "Name", "first operand", and "second operand". A green box labeled "Parameters" points to this section, with the text "Formula's input parameters(s)".
- Settings:** This section is titled "1 Setting to be Set". It includes a field for "Name" and a "scale" field. A green box labeled "Settings" points to this section, with the text "Formula's settings".
- Results:** This section is titled "1 Result Returned". It includes a field for "Name" and a "CurUtil_times" field. A green box labeled "Results" points to this section, with the text "Formula's output parameters(s)".

On the right side, a blue box titled "Formula "" - Parameter ""sec" lists the available options for each section:

- Parameters:** Constant Value, Filter on this Node, Formula Result, Constant Value (highlighted), Property Value, Combined Parameters, Empty Parameter.
- Settings:** Various settings, depending on the selected formula.
- Results:** Named formula outputs.

At the bottom right, the EMC² logo is displayed. The footer contains the text "Copyright © 2015 EMC Corporation. All Rights Reserved." and "Module: Formulas 10".

From a chosen formula, you will always have three sections for parameters:

Parameters to Enter => Input parameters definition:

All parameters have a short description indicating what is expected as a value. They also have a definition, that can be edited so you can select the exact parameter you need. You have choice of six types of input parameters:

- **Filter on this node:** A basic Watch4net filter. For instance : name='Availability'
- **Child formula result:** Get the child formula result (From the leaves to the root concept) .
- **Constant value:** Constant float value, for instance: 3.0
- **Property value:** Watch4net property must be filled here.
- **Combined parameters:** Combined several input parameters. For instance a "filter on this node" parameter and a "constant value" parameter.
- **Empty parameter:** Leave empty the input parameter.


Settings to be Set => Various options available, depending the on the formula:

For each setting, we have a short description and an editable value field. When the expected value is not a number, existing possibilities are usually included in the description.

Result Returned => Output metric names:

The names of the results are editable. Each result has a short description. Each result also has options. The first option, "show in graph" will be detailed in the following slides. The other, "default result" is used for aggregated reports.

Formula Input Parameters



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

11

Each single input parameter of a formula allows you to select the type of data from 6 formats. The initial default type is Empty Parameter, which selects empty data set.

Filter on this node: allows you to use a filter for metric selection. Remember that report context apply.

Formula Result: allows you to choose an already existing formula result from the current or child report node. This allows you for formula cascading.

Constant Value: allows you to select a constant number. Thus this creates a timeseries which fills timeslots with the provided value.

Property Value: allows you to convert a property value that is a number and, like constant value, to create a constant timeseries.

Combined Parameters: allows you to merge declared metrics together (sum, average, etc).

Empty Parameter: when an input parameter is not required.

Formula - Example

The screenshot shows the 'math.Division' formula editor. It has a title bar 'math.Division' and a subtitle 'Results [%PoolUsed]'. The main text says: 'This formula computes the quotient of the given metrics. The first parameter (optionally a combined one whose components will be summed) will be the numerator, while the other parameter (optionally a combined one whose components will be summed) will be the denominator.'

There are three sections:

- 2 Parameters to Enter:** A table with columns 'Name', 'Description', and 'Definition'.

Name	Description	Definition
numerator	This is the value to use as numerator.	Result on this node: "PoolUsedCapacity"
denominator	This is the value to use as denominator.	Filter: "name=="PoolUsableCapacity" & !(parttype)"
- 1 Setting to be Set:** A table with columns 'Name', 'Description', and 'Value'.

Name	Description	Value
scale	This value will be used to scale the result. For instance, it may be 100 in order to get a percentage.	100
- 1 Result Returned:** A table with columns 'Name', 'Description', and 'Options'.

Name	Description	Options
%PoolUsed	The computed result, ie numerator / denominator.	☒ show in graphs ☐ default result

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

12

The above demonstrates the usage of a math.Division formula. The goal of this formula application is to calculate the percentage of used capacity of a pool. In the report context, this formula is defined where a single pool is selected, that is where the report node expands (where metric context is grouped by pool name). The numerator that makes up the division is selected from an already existing value from another formula, sitting on the same report node: PoolUsedCapacity. The denominator defines a metric filter where the expression is very specific: metric is PoolUsableCapacity with no parttype defined... this expression makes sure the chosen metric is the right one.

The division formula then proceeds to the calculation on each 5-minute timeslot.

In the second section, each timeslot of the timeseries from the above is then scaled by 100 (thus a percentage value).

In the last section, the KPI is set with a name (%PoolUsed), thus it can be used: either in a graph, a column of a table or re-used in another formula.

At this point, if the result is to be used in a graph of the current report node, you are required to specify by checking the option "show in graphs".

Source: All>>Explore>>Storage>>More Reports>>Enterprise Capacity>>device

Use Formula in a Table Report

The screenshot shows the 'Value: % Pool Used' configuration window. It includes a description: 'A value column fetches a value from the database. You can specify the value with a filter, or as a formula result, such as the amount of your CPU utilization, and can express the value in the format you choose, such as percentage or bits per second.' The 'Column Name' is '% Pool Used'. The 'Selected Value' dropdown is set to '%PoolUsed'. The 'Use time settings from' dropdown is set to 'this column'. The 'Time Management' section includes: 'Sampling Period' (1 hour), 'Sampling Type' (average), 'Column Time Range(s)' (inherit report time range), 'Recover...' (the last value only), and 'Time Threshold' (2 period(s)). A green box highlights the '% Pool Used' column name and the '%PoolUsed' selected value. A green arrow points from this box to a list of formula results on the right. The list includes: '%PoolSubscription', '%PoolUsed' (highlighted in blue), 'Capacity-less-Usable', 'PoolUsedCapacity', 'TotalCapacity', 'TotalSubscribed', 'TotalUsed', 'UnusableCapacity', and 'UsedCapacity'. A text box above the list says 'use the filter below use a formula result'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

13

From the previous example, if you are required to use the calculated value into a table report, you need to have a “value” column type and select the formula output name from the drop down “Selected Value” selector. When expansion the selection, you can see you have access to all children formula results.

Note: if the formula is defined on the same report node where the table report is defined, the value will not show up.

The “Time Management” section allows you to specify what transformation you want to perform from the resulting KPI. In the above, report designer just wanted to recover the last value from the output of the formula. It also tells that 1-hour aggregate using average value will be passed to the formula. Obviously, the output value of the formula is another 1-hour aggregate.

Source: All>>Explore>>Storage>>More Reports>>Enterprise Capacity

Spatial Aggregation

Default behavior of Report Engine when multiple entities of a metric exists on a report node (or un-filtered column of a table): **average**.

Spatial aggregation formulas are meant to compute aggregation over metrics from multiple/underlying report nodes.

Metrics from expansion		formula						ParentReportNode
	Timeslot	1	2	3	...	x	y	
	metricA of obj1	v11	v12	v13	...	v1x	v1y	Child1
	metricA of obj2	v21	v22	v23	...	v2x	v2y	Child2
	metricA of obj3	v31	v32	v33	...	v3x	v3y	Child3
	metricA of objZ	vZ1	vZ2	vZ3	...	vZx	vZy	ChildZ
	Sum(MetricA of all objects)	r1	r2	r3		rx	ry	

le: Math.Spatial.Sum= timeseries or timeslot sums of given metric

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

14

Formulas can be used to aggregate results of underlying objects, having the same metric name. Example: Each LUN has a Capacity metric. You want to add them all together. Each LUN is a different object. The above depicts the concept of what spatial aggregation is. It is important to understand that if you apply a temporal formula when you are in a context of multiple (yet “spatial”) objects, the formula will first perform an average aggregation per timeslot: sum of all values divided by the number of metrics. In some cases, this default behavior is not what you expect. In order to workaround this default behavior, choose the appropriate Spatial formula. This one will be able to process, per timeslot, the appropriate function to all of the values, without performing average beforehand.

Example of default behavior: LUN report

- Parent report is a table
 - Child is expansion over Storage Array names
 - Sub-child exposes list of LUNs (per Array)

Parent is defined to be a three column table report:


a) Array Name, b) filter over ‘ReadRequests’, c) filter over ‘WriteRequests’

We want to report data per Array, but database metric is single per object... we have to consolidate objects of a given array. Under columns ReadRequests and WriteRequests, exists multiple entities of each metrics (one per LUN). If no proper action is taken for report computation, default behavior of Report Engine is to average all values per timeslots. This might not be what you want!

In order to aggregate the metrics and to add them per timeslot:

If there is a SpatialAggregation formula applied at parent level, that sums all ReadRequest together, and another formula that sums all WriteRequests together, the result presented to the table will make more sense.

Formula – Spatial Aggregation



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

15

In order to demonstrate appropriate use of Spatial Aggregation, let's investigate a report from the Dashboard (global report).

The consumable report "Raw Capacity Usage" is defined as a type that requires a value from each child report so it is able to generate a Bar chart from this context. The different categories are described by children reports where each child report filters on a specific metric (ConfiguredUsableCapacity, HotSpareCapacity, etc). The "Arrays" report node is an expansion on each found array device of the context. The parent report is also responsible to aggregate (sum) the values for each found device (ie: array) using a `math.spatial.Sum` formula. A default value is generated for each child and the parent will aggregate them.

The above example is taken from a template, re-usable report, that will create the context from a parent node (of Raw Capacity Usage).

You can discover from the standard report library, that same concept applies for the drill down of Usable Capacity and by Purpose - also applies to Used Capacity Trend report.

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Raw Capacity Usage

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Usable Capacity

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Usable Capacity Trend

Spatial Aggregation in Action

Raw Capacity Usage							
Configured Usable							
Arrays	Timeslot	1	2	3	...	x	y
ConfiguredUsableCapacity of Array1	v11	v12	v13	...	v1x	v1y	
ConfiguredUsableCapacity of Array2	v21	v22	v23	...	v2x	v2y	
ConfiguredUsableCapacity of Array3	v31	v32	v33	...	v3x	v3y	
ConfiguredUsableCapacity of Array{Z}	vZ1	vZ2	vZ3	...	vZx	vZy	
Sum.Spatial :	r1	r2	r3		rx	ry	

Vab = Value of metric a on timeslot b
Rn = Result value of timeslot n

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

16


Above depicts an example of what the Report Engine is doing at browse time, for the first report node (Configured Usable). For each found Array (defined as an expansion), it performs a sum of the values for each timeslot of the metric context (which is filtered on only 1 metric, ConfiguredUsableCapacity in the above case).

The same is performed at each level of the whole report under “Raw Capacity Usage”, but for different metric (filtering selects a metric context for each).

The timeslots are per report’s default aggregation. If it’s Real-Time, it uses 300 seconds period, if it’s 1-hour, it makes use of 1-hour aggregation data, and so on.

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Raw Capacity Usage

Spatial Sum - Raw Capacity Usage



The resulting report. Each report entry filters for one specific metric, from all found arrays. Spatial.Sum formula aggregates and sums them up at the report definition. The bar report consumes the returned values from each report (Configured Usable, Hot Spare, etc). We will detail the construction of such report type in a separate topic. The green rectangle shows our studied example from previous pages.

Source: All>>Dashboards>>Storage>>Enterprise Capacity Dashboard>>Raw Capacity Usage

Frequently Used Formulas

Formula name
Nop, Nop4
Math.spatial.sum
Math.Addition/Substraction
Math.Division
Util.ChildCount
Util.distinctValueCountFormula
Math.standardDeviation
Math.temporal.percentile
forecasting.LinearForecast
util.FindMinMaxTimestamp
util.ConditionalTimeValueCount

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

18

Nop (and its variant, Nop4) used to forward results from a child node to the root level. The second use case for Nop is data re-alignment (our next topic). Nop4 is simply forwarding up to four values using a single formula definition, faster, easier to manage than a bunch of individual Nop formulas.

Math.spatial.sum aggregates data for a global report view: all arrays together, all VMs, all Hypervisor, etc.

Math.Addition or Math.Substraction computes complementary values: totalCapacity = Free+used, Used=Total-Free, etc

Math.Division is used for ratio and percentage computing. Other variations are: Math.Ratio, math.percentage, math.DifferentialRatio. If the goal is to scale a value, use math.scale or other formulas that have a settings.

Util.ChildCount allows you to compute the number of items resulting from an expansion. Typical use case is where you want to count the number of hosts and re-use that number is another formula calculation in a column of a table report. Please note that ChildCount is also a feature of an Attribute of a column of a table. You have more flexibility in using a formula than an attribute.

Util.distinctValueCountFormula is a new formula in 3.5 that is similar to Util.ChildCount and allows you to count the number of items. The difference is that for Util.ChildCount, you need to count from a report structure (result of an expansion is typical). Rather than doing this, distinctValueCountFormula directly counts the number of distinct values without having to generate the nodes (ie perform the expansion on the report structure). This leaves your report structure cleaner.

For the advanced statistical computing:

- **StandardDeviation:** used in statistical analysis, indication of the spread of the values around the average.
- **Percentile:** set with the 95 percentile, is shows the value where 95% of the samples are equal or lower than this value.
- **Forecasting:** used to forecast a metric value based on the selected report time range. You can have multiple results: expected value in the future, deviation (per second), etc. Calculation uses linear regression method.
- **FindMinMaxTimestamp:** recovers various values: Min/Max values and their respective date/time.
- **ConditionalTimeValueCount:** many formulas of this type. Base on conditions, you can count the number of occurrences where the value reaches a given value (configured in formula). Thus, you can count the number of events when a value exceeded a value. This formula is presented here just to demonstrate the richness of the formula library.

Metric Timeslots Misalignment

Timestamp mis-alignment / Original values

Tuesday, July 22, 2014, 1:00 PM » 4:00 PM, EDT | Last 3 Hours average on real-time



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

19

A caveat you can meet when designing an aggregated report is the misalignment of data. This will occur mostly when you have multiple SolutionPacks installed of the same type (VMAX, VNX, VMware, etc). Above capture shows data misalignment between two arrays. You can clearly notice the diamond data points are not aligned between blue and yellow timeseries. Notice the polling cycle: 20 minutes. The default timeslot of the report engine is five minutes (real-time data is in used).

For demonstration purposes, a single metric of each array has been selected (filtered): RawCapacity. This is also the case with all metrics from the array and for which the polling cycle is not aligned with the default of the report engine.

This has an impact when trying to aggregate object metrics (arrays) all together in an overall report. You must either re-align the data to the closest aggregation (there's a formula for that). Or use a higher aggregation level (ie: 1-hour would suffice for this case).

As an example, if you need to Spatial Sum RawCapacity metrics from various Arrays, you will have to either:

- Use a re-aggregation formula with a time base of 20 minutes (1200 seconds) then perform Spatial Sum formula.
- Use Spatial Sum formula with 1-hour aggregate (either at the report's definition or in a column of a table report).

Please note, in the above use case, simply passing the metric through Nop formula won't suffice: the poling cycle of the original data is different from the default base aggregation of the report engine (5 minutes).

Timeslot Misalignment Consideration

Issue:

- Aggregation results (spatial.sum) does not add up correctly

Causes:

- Two distinct SolutionPacks installation
- Collectors polling cycle not in sync
- Polling cycle different from Report Engine's default period

Solutions:

- Application of a formula (Nop, AggregationOfAggregation)
- Use of common time based denominator (aggregation)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

20

Each SolutionPack has its own associated collector manager (CM). Each CM is a distinct process and cannot all be in synch on a single polling cycle, thus they are skewed at time stamping. This wouldn't be much of a problem if they were aligned with the default time base of the report engine. The default of the later is 300 seconds, or 5 minutes. If a collector's polling cycle is, for example, 20 minutes, it is much higher than the default time base. Thus when aligning data at report time, the report engine will have some "holes" in timeslots and could not work correctly when going through formula calculation that requires aggregation of objects.

There are two valuable solutions to workaround this situation. First, the use of higher aggregation level that both original data and the report engine will align on (common denominator). As an example, from the above, use of 1-hour data aggregation level will make sure data is aligned across multiple objects, even if they are polled from out-of-synch collectors.

In some cases, you cannot afford the above solution. In that case, you'll have to define an aggregation level using a formula that will force the report engine to align with. This formula is `math.temporal.AggregationOfAggregation`. This formula lets you generate an aggregation at report time and set a new time base. Automatically, all object metrics and the report engine will align. You will also notice that if data need to be aligned before being spatial summed up, simple use of `Nop` formula will suffice.

Demo: Misaligned Data

This demo will help you deal with a misaligned dataset.

▶ Demo - Timestamp mis-alignment

Original values

Nop effect - Real-Time

Re-alignment using 1-hour aggregate (AVG)

Re-alignment using formula

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

21

This demo will show you:

- How to detect such misalignment of data.
- Expose some solutions that could fix the situation.

Be sure to check the demo server for a report named “Demo - Timestamp misalignment”.

Application: Chargeback Reports

Dashboards

- Storage
 - Enterprise Capacity Dashboard
 - File Capacity Dashboard
 - Virtual Storage Dashboard
 - Chargeback by Group**
 - Primary Used Chargeable
 - Primary Presented Chargeable
 - Total Used Chargeable
 - Total Presented Chargeable

Primary Used Chargeable

9 elements found, displaying all elements.

Group	Host Type	Platinum	Gold	Silver
All Physical Hosts	Host	0.00 GB	49.21 GB	131.33
Provider Tenant	Hypervisor	0.00 GB	721.44 GB	341.33
Sales	VirtualMachine	0.00 GB	221.32 GB	275.93
Human Resources	Host, VirtualMachine	0.00 GB	263.92 GB	71.33
Manufacturing	Host, VirtualMachine	0.00 GB	0.00 GB	0.00
IT	Host, VirtualMachine	0.00 GB	802.25 GB	1.35
All Hypervisors	Hypervisor	440.44 GB	44.67 TB	5.76
All VM Guests	VirtualMachine	228.83 GB	14.33 TB	3.94
Engineering	VirtualMachine	0.00 GB	1.85 TB	37.36

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

22


With ViPR SRM 3.5 M&R you have a standard Chargeback report by Service Level. The report can be viewed under the following source:

All>>Dashboards>>Storage>>Chargeback by Group >> Primary Used Chargeable

This report allows you to view which groups (tenants) consume which amounts of Service Level resources, thus you can charge them back. There are pre-defined Service Levels showing up in the report: Platinum, Gold, Silver, Bronze, Fast VP, Virtuals and Others. This report merges two type of information: Groups (rows) and Service Level (columns). Those must be defined at the collection layer before correct data shows up in the report: refer to ViPR SRM Administrator to perform this.

The use case is that you want to modify this report and have additional Service Levels (Custom[1,2,3]) showing up. This report makes use of many transformations and formulas. The good thing is that Custom Service Levels are defined but simply not considered in report. All you have to do is to enable them.

Application: Chargeback Reports – (cont.)




NEVER work directly on the OOTB (base/default) reports! Always take a local copy of the source reports then you can modify.

From the report copy, under Formula tab of “Group Expansion”, add an aggregation formula (Spatial.Sum) that sums up one of the Custom capacity KPIs. Name this result meaningfully. Note: we recommend you to take a copy of an existing Spatial.Sum formula and modify the input and output parameters.

In a second step, under “Primary Used Chargeable” report, add a new value column (or take a copy from an existing one) on the existing table. As the selected value, choose the result from the formula you just added. Name this column meaningfully.

You can also change the column headers to bet fit your definition of Service Level. Do not change Service Level at the source (Groups Management).



With ViPR SRM 3.5, you have access to three additional custom Service Level definitions that are not initially showing in the default report. This feature usually requires some elevated privileges that a ViPR SRM platform Administrator can provide to you.

You can change the definition using Group Management feature, found in Centralized-Management tool. Please refer to a ViPR SRM Administrator in order to either perform the change. You can create any device grouping you want, even change/rename the default ones.

Basically, when you manage the User Groups, you categorize storage based on criteria. You already have access to defined groups such as Platinum, Gold, Silver, and Bronze... and some labeled "Custom", numbered from 1 to 3. The later ones are the "define you own".

The limitation is that you can hardly define new ones and/or change names of such Service Level. In fact, those names are tightly referenced in the defined reports. So instead of changing the Service Level names, it is recommended to change the column title and/or add the custom ones.

Once you are done, you can now start (first) copying the reports, and then alter in order to enable the Custom Service Levels to show up in report.

Formulas – Writing Your Own

Modules

Frontend

- Formula Editor
- Management of Database Metrics
- Outage Editor
- Real Time Grapher
- Topology Management

```
File  Wizards
package com.watch4net.apg.v2.gui.formula.library.util;

import com.watch4net.apg.v2.gui.data.TimeValuePairArray;
import com.watch4net.apg.v2.gui.formula.StandardFormula;
import com.watch4net.apg.v2.gui.formula.StandardFormulaComputationContext;

/**
 * This formula does absolutely nothing, except passing the parameter va
 *
 * @author © 2005-2007 - Watch4Net Solutions
 */
public class Nop extends StandardFormula {
    /**
     * Creates a new {@link Nop} formula instance.
     */
    public Nop() {
        setDescription(
            "<h3>This formula forwards the given parameter to the next l<br>"
            "<p>This proves useful when we do not want a value to be agg<br>"
            "in the database. Thus, using expansion, we can get it at t<br>"
            "level and use it in a formula without any database aggrega<br>"
            "<i>(which is the default behavior on multiple selections)</i>"
        );
        setParametersNames(new String[] { "metric" });
        setParametersDescriptions(new String[] {

```

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

25

If you are familiar with Java programming language, you could write your own ViPR SRM M&R formula right from the GUI. You don't need any IDE tool.

Please note that formula development from EMC (ViPR SRM development team) is available with Professional Service engagement only. Thus, there is no official support with the use of this feature.

Lab 8: Integration of Formulas in Reports

This series of lab will allow you to:

- Practice application of formulas
- Discover metric set from various collection contexts
- Explore the formula library and various usage and possibilities
- Practice more report structure

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

26

This lab 8 allows you to practice many use cases of formulas in reporting.

Part 1: IOPS = ReadReq + WriteReq -> per VMAX Array.

Part 2: Brocade port utilization (%)

Part 3: Property Value conversion using formula (maxspeed Cisco)

Part 4: Add Custom group to OOTB Chargeback report

Part 5: Baseline using formulas

Part 6: VMware: from mod 4, lab add child count – intro to spatial aggr

Part 7: IOPS on VNX Block, per SP , requires Spatial Sum - simple

Part 8: Timeslots realignment

Part 9: Discovering Percentile, and other advanced formula application from Lab 1 (max,min,avg, trending, deviation, stddev, period below X, Timestamp of Max/Min

Part 10: SAN Switch port count (DistinctValueFormula) – advanced use of formula, special case.

Part 11: VMware Chargeback

Lab 8 Debrief: Integration of Formulas in Reports

This lab exercise allowed you to:

- Familiarize with formula concept
- Explore the reporting possibilities with the use of formulas
- Understand how to build more complex (and useful) reports
- Transform metrics into KPIs

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

27

This lab exercise allowed you to become familiar with formula concept, explore the reporting possibilities with the use of formula, understand how to build more complex (and useful) reports and transform metrics into KPIs.

Module Summary

Key points covered in this module:

- Applications of formulas in reporting
- Temporal and spatial aggregation
- Timeslot misalignment: impacts, causes, and solutions

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: Formulas

28

This module covered the formula concepts in ViPR SRM reporting.

Module: More Report Types

Upon completion of this module, you should be able to:

- Describe the different types of report
- Explain usage of other report types



This module focuses on exploring remaining report types that were not covered yet.

Module: More Report Types

Lesson: Aggregated Report Types

This lesson covers the following topics:

- Stacked charts
- Stacked bars
- Status reports
- Gauge reports






This lesson covers stacked charts and bars and status, and gauge reports.

Aggregated Report Types

Characteristics:

- At-a-glance total: cumulative effect
- One area per found child report
- Small areas automatically grouped together
- Legend lists active and grouped areas
- Status makes use of defined thresholds
- Gauge chart shows single value

Aggregated



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

3

The most detailed aggregated report is the Stacked chart. It displays an area for each child of a node, and the areas are stacked on top of each other in the resulting graph. To enhance clarity, small areas are grouped and aggregated in a section called Other, displayed in light gray. The legend lists the grayed variables so you can enable them if you want to. Each area of a stacked chart is the result of the aggregation set in the report options for the specified period, as many times as the period can fit into the time range.


A status report aggregates data for each child using the specified aggregation function over the selected period. Each child node is displayed as a large colored symbol with the aggregated value and the name of the node displayed. Color thresholds can be defined. Different symbols (squares, triangles, ...) can be configured.

A Gauge chart provides you with a snapshot of the health of an indicator. Gauge charts display the absolute aggregated value of a single indicator.

Heatmap and Treemap will be covered in a separate lesson.

Stacked Chart

Cisco MDS/Nexus / Overall Throughput



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

4

Stacked chart is a standard chart where metric timeseries are aligned and values stacked on top of one another. The total height then gives a visual idea of the total and it allows to relate the most significant values to each other, over time. Most significant numbers are stacked first and remaining ones are summarized (grouped) together, per timeslot and shown as a total. In the use case above, there is no “Other” category because the number of metrics is less than the minimum required as per report configuration.

This type of chart uses children report nodes to generate the stacked timeseries. It also proves useful when you want to summarize and relate a single metric type among multiple objects or device and where stacked value makes sense. It proves also useful when showing complimentary values that can be stacked. As an example: WriteRequests and ReadRequests of a storage object (Pool, LUN, Disk, Array, etc).


The above report aligns the data from each port of each FabricSwitch. Device Traffic (MiB/s) KPI values were computed from In and Out octets and aggregated on children report nodes. We can notice the trend: every hour there is a peak of global usage across all switches. This is due to some hourly scheduled backup updates.

Note that because of the nature of this type of report, you must ensure to select a sampling mode where data is time aligned. Otherwise, you might end up with false results (total not adding up correctly).

Source: All>>Report Library>>Cisco MDS/Nexus>>Overall Throughput

Stacked Chart – Structure

Stacked chart requires children definition:



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

5

The above is the construction of a stacked report: a list of report nodes can be used on the child level of the defined stacked chart report. This could be from a defined expansion or a list of report nodes returning a single value from each. If a stacked report is a leaf report, the report engine will create a chart based on the metric available from the context, one metric on the top of the other. For that later case, you must ensure you have properly filtered the context.

The example shown above demonstrates that the report no longer requires the “total bandwidth” KPI to be present on the graph, due to the nature of the report: one will see the sum effect of stacking base metrics. A lab will allow you to practice this use case, demonstrated above, on left.



In the structure of the above report, please note that filter selecting Read and Write metrics is set at the expansion level. One could simplify the report by moving the filter up to the parent and remove the last report node (“expansion node report”). This use case is fine because the report is a leaf.

Source: All>>ViPR SRM Reporting class>>Module 06>>Lab 1 - Stacked Chart/Bar>>VNX-Block Use LUNs Perf>>[serialnb,part]>>Bandwidth (MB/s) - Stacked>>[selected metrics]

Source: All>>Report Library>>EMC VNX>>Details>>Block Systems (All Reports)>>Performance>>User LUNs Performance>>APM88928219360, LOGICAL UNIT NUMBER 0>>Bandwidth (MB/s)

Stacked Bar Chart

- Vertical bars based on time
- Viewable total
- Relate metrics
- Data time stamps must align
- Bar color based on thresholds (single metric only)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

6


The stacked bar chart shows data with vertical bars instead of being linked together with a line. One bar is the width of a time slot, defined by the sampling type selection of the report. It is same definition as for stacked chart: children define the categories to stack or the metric context tells the metrics to stack.

Shown above is the relative level (or magnitude) of the data from two metrics: Read and Write Throughputs. At the same time, you have a good idea of the total (bandwidth).

If your report is shown a single metric (or KPI) and if you use threshold (horizontal dotted lines), bar color will adjust automatically.

A definition of this report requires, as member of aggregated type, children report node definition that the stacked bar chart will consume to generate the report.

Source: All>>Report Library>>VMware vCenter>>Datacenters Summary>>CSE 176 Lab>>Hosts Clusters>>CSE Cluster 2>>Hosts>>mqqb053.lss.emc.com>>Storage Performance>>Global Disk Throughput



Wondering about the regularity of the data collection? Here's a hint about how to check the polling cycle consistency. You can use this hint while in browser mode (recommended) instead of building a specific report for this task.

A collector with a normal collection rate (polling cycle) will gather a value for each metric every 5 minutes (for the cases where polling cycle is set so). The 1-hour aggregate data will then summarize 12 samples per hour and the 1-day aggregate 288 samples. Among the aggregate fields, you have "count" which represents the total number of samples for the given aggregate. A quick way to check the consistency of the collection is then to use bar chart with sampling period set to 1-hour and use the sampling type to "count". When looking at the chart, all bars should be aligned to 12 for the whole report period. It will align to 288 if using the 1-day aggregate. If this is not the case, you can say that either the collector missed some polls OR the polling cycle is set differently.

If you build a report, don't forget to be specific in using this report (specific filter) OR make sure to expand/drill down the structure. Otherwise, all metrics or the database will be aggregated in a single chart... potentially blocking/clogging the portal.

From the above, in exhibit 1, you can notice that stacking the count values for four different arrays give consistent over time (4 arrays times 12 sample/hour=48). Each count value of each array is consistent (always the same height).

From the above, in exhibit two, you can notice that some arrays are polling data 3 times per hour while others are polling for 4. Between 1PM and 2PM, some arrays missed a polling cycle so one stack is lower in total than the other.

Status Report Type

Characteristics:

- Quick Visual report
- Generally used for overview reports
- Button color from thresholds
- Single value result
- Tooltip provides details
- Clickable (drill-down)

Fabrics Overall Health / Switches

The current operational status of the switch.

The states are as follow:

- online(100%) means the switch is accessible by an external Fibre Channel port;
- offline(0%) means the switch is not accesible;
- testing(NaN) means the switch is in a built-in test mode and is not accessible by an external
- faulty(50%) means the switch is not operational.



mgmt207106
Value 100
Device description Connectrix ED-DCX8510-8B
Device model Brocade DCX 8510-8 Backbone
Device IP address 172.6.207.106
firmware v7.1.1c
FC Port count 640

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

8

The status report is a relatively simple and basic report type. It is of Aggregated type because as Stacked reports, it requires children reports in order to enumerates the elements of the report. This report aligns colored icon with an optional value printed on it and the name of the element on the bottom.

Status reports are meant to provide a quick visual report with color code (although optional). It is best used to visualize multiple elements or aspects of health of a good working environment. You can see this type of report in overview reports.

The chosen icon can be colored according to defined threshold values.

A single metric value can be computer and shown on the icon and checked against the thresholds. If more the one metric is selected in the context, all values are aggregated (depending on the configuration, default is to average).

A tooltip can be defined that will show up when mouse is over the icon. List of properties can be provided in report configuration.

Any icon is clickable, allowing user to drill down deeper in the report.

Source: All>>Brocade FC Switch>>Fabrics Overall Health>>Switches

Gauge Chart

Characteristics:

- Quick view of range value
- Single metric or KPI
- Graphic and text value
- Visual threshold zones

mgmt207106, FC port 2/9



--- Normal (Values less than 70) --- Major (70) --- Critical (90)

CurrentUtilization (%)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

9

A Gauge chart provides you with a snapshot of the health of an indicator. Gauge charts display the absolute aggregated value of a single indicator. Both text values and graphical values are shown - useful for comparison and to relate values in their range of possibilities.

Thresholds are shown in a visual mode: the report uses threshold settings to color code the range of values. Any range of metric value is acceptable, gauge range is from defined boundaries, set under Report Detail.

In the above, the metric CurrentUtilization has been filtered out and shown as an average over a period for a given port on a given SAN Switch. The value is expressed in percentages.

Module: More Report Types

Lesson Summary

During this lesson the following topics were covered:

Aggregated report types:

- Stacked charts
- Stacked bars
- Status reports
- Gauge reports



This lesson covered aggregated report types: stacked chart and bars and status reports.

Lab 9: Practicing Aggregated Report Types

This lab will allow you to practice Aggregated report types:

- Part 1 – Stacked Chart report
- Part 2 – Stacked Bar report
- Part 3 – Status report
- Part 4 – Gauge Chart report

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

11

This lab 9 time will allow you to practice with some of the presented Aggregated report types construction.

Part 1 – Stacked Chart report: VNX-Block Use LUNs performance: This report is based on an already existing report from the VNX report Library.

Part 2 – Stacked Bars report: Brocade SAN Switch status: provided a view of the current status of the environment.

Part 3 – Status report: this report is based on an existing report from Brocade report library.

Part 4 – Gauge chart report: Brocade SAN switch port utilization in an interactive report.

Lab 9 Debrief: Practicing Aggregated Report Types

This lab exercise allowed you to:

- Familiarize with stacked report types
- Familiarize with status and gauge report types
- Understand the use of other report types
- Explore the possibilities in reporting

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

12

This lab exercise allowed you to: become familiar with stacked report types, status and gauge report types, understand the use of report types, and explore the possibilities of reporting.

Module: More Report Types

Lesson: TopN Report Types

This lesson covers the following topics:

- Definition
- Horizontal bars
- Single horizontal bar
- Pie chart
- Time ratio bar



This lesson covers the grouping report types (TopN).

TopN Report Types

Characteristics:

- Data presented in categories
- Put categories in perspective with the others
- No time axis – timeless chart
- Single metric per chart
- Timeseries metric is aggregated
- Results presented in percentage
- Most significant categories are presented first
- Least significant categories are aggregated
- Involves at least two reports nodes (parent+child)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

14

TopN report types are meant to visually compare data of same type altogether, in a global way.

Two important datas are used to create a TopN report: categories and the value they represent.


The categories are compared and for this reason, there is no time axis showing, although, with ViPR SRM, the report always defines some time range. Only a single metric or KPI can be used to generate TopN report. This means you cannot compare the same TopN report, let's say, Capacity and Throughput together (or any kind)... that wouldn't make sense.

So what happens to the metric timeseries data that is recovered from the database? It is simply aggregated as defined by the report's definition (average, sum, min/max, last value, etc). Most often you will find that TopN reports makes use of a default value from a formula output... which is selected to return a single value from each category. Once generated, the report presents the data in percentage, easier to compare between categories. You can hover the mouse over a category (a bar, a slice, an area) in order to read the raw values used to generate this piece of the chart.

You will also notice the report engine will rank categories by their weight. The most significant shows first while the least significant are aggregated when the number of categories is greater than a defined value (configurable).

Finally, the TopN report consumes the child reports for two components of the report: categories and values. This is the reason why a TopN report defines at least a structure of two (or more) report entries. The report in TopN list that differs from the above is the time ratio bars report.

Bar Graph Definition



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

15


A bar chart or bar graph report is a chart with rectangular bars with lengths proportional to the values that they represent.

A bar graph report is a chart that uses horizontal bars to show comparisons among categories. The Y-axis of the chart shows the specific categories being compared, and the X-axis represents the discrete value.

When this type of report is used, it has a list of categories (defined by sub child nodes) and a default value (AVG/MAX/MIN/SUM from timeseries) from each child (from each category).

Source: All>>Report Library>>VMware vCenter>>Inventory>>Virtual Machines>>Distribution by Guest O.S

Bar Graph Modes



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

16

The bar graph report can also present bars clustered in groups of more than one (summary grouping mode), and can also show the bars divided into subparts to show a cumulative effect (stacked grouping mode). Stacked and summary grouping modes prove useful when you want to see information from the same representation: relative among categories and distribution (or weight) within a category.

Default behavior of is the tooltip: user simply hovers the mouse over a category (or sub-category) in order to to generate the report - no need to click or to drill down. When clicking on a category, this will bring you down to the child report (if defined and available).

Source: All>>ViPR SRM Reporting class>>Module 06>>Lab X - Bar Chart>>VDC VM Customer and Business Unit distribution

Bar Chart - Options

- Scaling: convert units (ie: from GB to TB)
- Metric customization: select color of categories
- Legend: select legend format
- Settings for TopN Graphs: number of categories, ordering, modes, colors of default
- Tooltip information: properties to add to tooltip pop up

Settings for TopN Graphs

Section order: descending ?

Section grouping mode: summary ?

Number of sections: 10 ?

Others section color: [color picker] ?

Expand Others section: ☒ ?

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

17

Some options are available under “Report Details: Horizontal Bars” tab:

Scaling: allows you to scale the values of each category. For example when you know you are aggregating values in GB and that would be more meaningful to show up in TB or to let Report Engine decide at run time. As usual you can simply multiply/divide the values.

Metric customization: allows you to control the color of categories.

Legend: allows you to set how a legend will appear. This also allows you to select the “Others” category.

Settings for TopN Graphs: the most relevant for this type of graph report, allows you select category ordering rule, number of sections, mode (stacked or summary) and define color of the default category.

Tooltip information: allows you to select additional information to show up when hovering the mouse on a category. You can add as many properties as you want.

Bar Graph – Grouping Mode

Settings for TopN Graphs

- Section order: descending
- Section grouping mode: **summary** (selected)
- Number of sections: 3
- Others section color: #FFFFFF
- Expand Others section: ☐

Expansion

The specified properties will be used for the expansion of the current node

Expand on every:

New Group

Simple: vendor

Group 1

Simple: model (Device model)

Typical report structure

- {parent report – bar chart}
- {expansion node report}

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: More Report Types 18

The Grouping mode of the bar graph report depends on the expansion expression.

From the report where you chose Bar Chart report type, you have to set the “Section grouping mode” from the “Report Details: TopN Graphs” tab: choose among summary or stacked. Then from child report’s “Filtering & Expansion” tab, select the two properties that will factorize (group) the metrics. Then move up one property that will represent the main grouping. You should see the UI layout showing “New Group” separate from “Group 1” and a yellow zone behind. Drop the property over the yellow zone. You should then see Group 1 and Group 2.

Group 1 is the first grouping, this will set the first level of category. Group 2 set for second level of category, thus creating the stack or the second distribution.

If you already know the property names, choose to use the Property Input Field button and then fill in the property surrounded by parentheses and separated by a comma. Example: (vendor),(model) defines main category as vendor and sub-category as model.


Note: There is a formula in use on the expanded node in order to present a value of the category that the parent report (the bar chart report itself) will consume.

Pie Chart

Similar to bar chart in definition and use:

Pools / Distribution by Target Protection

July 2014, Monday 21 » Tuesday 22, 4:00 PM EDT | Last 1 Day: average on real-time



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

19

Pie chart is a variation of the bar chart. It is defined as a report type and makes use of the same options. The difference is that there are no stacked options, only summary is in effect. Report structure is the same, parent report consumes categories and their respective values. Hovering the mouse pointer over a section will detail the value in a tooltip pop up.


Source: All>>Report Library>>EMC VMAX>>Inventory>>Pools>>Distribution by Target Protection

Single Horizontal Bar

- Similar to bar chart in definition and use
- Categories stacked horizontally

Pools Capacity / Pools Capacity Summary

July 2014, Monday 21 » Tuesday 22, 4:00 PM EDT | Last 1 Day: average on 1 hour



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

20

A Single Horizontal bar report uses the same data as the horizontal bar and pie chart reports but displays the proportion of each child with regards to all the other children on a single horizontal bar. Group expansion expression is not in effect for this type of report.

Source: All>>Report Library>>EMC VMAX>>Pools Capacity>>Pools Capacity Summary

Time Ratio Bar

Analysis over time, per time interval (threshold)


Lab X - Single Horizontal bar / ESX hosts CPU utilization

This report categorize VMware hosts CPU utilization over the period using 1-hour aggregate.

The defined intervals are:

- [0-30)%
- [30-45)%
- [45-55)%
- [55-70)%
- [70-90)%
- [90-100)%

CPU of this host was in the range of 45-55% for about 77% of the selected report time range.



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

21

EMC²

The time ratio bar report proves useful when you want to provide an analysis of the behavior of a metric over time, based on value intervals (defined as thresholds). This report does not require definition of categories. Instead, it requires definition of value intervals. Each value of the metric timeseries is categorized according to the defined value intervals. Value intervals are then compared among themselves and the horizontal line can be generated. Defined value intervals are showing in the legend, at the bottom of the report.

This report would answer a question like: over a certain period, how long (in % of time) will the metric value fall into the defined intervals? As depicted above, an Administrator is interested in monitoring the CPU usage of the hosts over a one week period. Quickly, he can determine the proportion of the period the CPU sustained in an interval. From the above, it shows that most of the hosts CPU are in a safe zone of operation. Only two hosts were shooting in the interval of between 55-70%, but that was for a relatively short period of time (19% and 8%, respectively). You can then play with some variations: real-time values over a short period of time, etc. Note: not all hosts are shown above, see the source report.

Hint: in order to help the report engine to optimize computation, make sure to use a proper time range and aggregation level conveniently. The above makes use of 1-hours aggregate over a whole period of 1 week. A bad example would be to use real-time value over a 1 week or a longer period. This has a direct impact on the compute time and can worsen if the number of hosts to analyze becomes too long.

Source: All>>ViPR SRM Reporting class>>Module 06>>Lab X - Single Horizontal bar>>ESX hosts CPU utilization

Module: More Report Types

Lesson Summary

During this lesson the following topics were covered:

TopN report types:

- Horizontal bars
- Single horizontal bar
- Pie chart
- Time ratio bar



This lesson covered aggregated report types: horizontal bars, pie charts, and time ratio.

Lab 10: Using TopN Report Types

This lab covers TopN types of reports:

Part 1 – Bar charts:

- A) FabricSwitch Inventory I
- B) FabricSwitch Inventory II
- C) VM guest OS distribution I
- D) VM guest OS distribution II
- E) VDC VM Customer/Bunit distribution
- F) Check data polling

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

23

The Labs including TopN report types allows you to practice various way of presenting the data.

Lab 10 – part 1 is about using bar chart report types:

a: FabricSwitch Inventory using simple expansion only

b: FabricSwitch Inventory using group expansion

c: VM Distribution by guess OS using simple expansion

d: VM Distribution by guess OS: an alternate view

e: VDC VM Customer and Business Unit distribution. This report re-use some of the VM Chargeback report.

f: This report is used to practice data collection polling among various collections and let you conclude. (Array context vs Switch Context)

Lab 10: Using TopN Report Types (cont.)

This lab covers TopN type of reports:

Part 2 – Pie charts:

- A) part 1 – FabricSwitch Inventory III
- B) part 2 – Storage Pool Distribution by Target Protection

Part 3 – Single Horizontal bar report:

- A) Part 1 – Datastore Capacity Distribution
- B) Part 2 – FabricSwitch Inventory IV

Part 4 – Timeratio report:

- A) Part 1 – ESX host CPU utilization
- B) Part 2 – IORate analysis (revisited)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

24

The Labs including the TopN report types allows you to practice various way of presenting the data.

Lab 10 part 2: using pie chart report type:

a: FabricSwitch Inventory using group expansion

b: Storage Pool Distribution by Target Protection. Re-use of an existing report from the VMAX Report Library.

Lab 10 part 3: Single Horizontal bar report types:

a: Datastore Capacity Distribution: re-use of an existing report from VMware report library.

b: FabricSwitch Inventory IV (revisited): cascaded expansion by vendor and model

Lab 10 part 4: Timeratio report types:

a: ESX host CPU utilization: an analysis of the CPU utilization over time

b: IORate analysis: re-purpose report from module 5.

Lab 10 Debrief: Using TopN Report Types

These lab exercises allowed you to:

- Familiarize with TopN report types
- Understand the use of one or the other of report types
- Explore the possibilities in reporting



These lab exercises allowed you to become familiar with TopN report types, understand the use of one or the other of report types, and Explore the possibilities in reporting.

Module: More Report Types

Lesson: Advanced Report Types

This lesson covers the following topics:

- Heat map
- Tree map
- Geo map




This lesson covers advanced report types: Heat, Tree, and Geo maps.

Advanced Report – Heat Map

Characteristics:

- Helps visualize aggregated data on two axes
- Suitable for analysis of a single metric/KPI
- Spatial and temporal analysis
- Spatial requires group expansion
- Thresholds are required
- Cells areas are of the same size

Windows VM Filesystem Heat Map (%Free)



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

27

A Heat map is a report type representing data in a matrix and using a cell color code, based on an aggregated value. The information you want to extract from such report is to quickly identify spots where to look at priority, ie: hot spots. That way, you will be able to plan your action on the area where it is the most required.

A single metric or KPI type can be used for a Heat map report, but it could be from multiple objects. You have two possibilities over expansion that will generate the two axes: use standard property expansion or use time expansion. The above example uses two expanded properties for both axes: host (device) on the x-axis and drive object (part where parttype='FileSystem') on the y-axis. Time expansion mode would allow analysis of the metric over timeslots, covering the report time range.


For each cell, the metric/KPI timeseries is aggregated in a single value (ie avg, min, max, etc) before being associated to a color, based on thresholds definition.

Hovering the mouse on a cell will pop up a tooltip helper giving the details of the current cell value (and other configuration properties, if set).

Source: All>>Report Library>>VMware vCenter>>Situations to Watch>>Heat Map of the VM Filesystems>>Windows VM Filesystem Heat Map (%Free)

Heat Map - Definition

- Used to visualize aggregated data along two axes (X,Y)
- Analysis of single metric per heat map report



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

28

Use heat maps to visualize aggregated data along two axes. Each dimension is the result of an expansion of a property. For each property couple, a value is associated and is represented using a color. The color depends on the threshold set under Report Detail: Graph tab. There are other settings that add some flexibility to the report: Tooltip Information, Order, Show grid and Use gradient colors.

Any white cell in the resulting matrix indicates missing value for current couple (X and Y).

Heat maps are similar to status report, but using two dimension definition instead of a single one.

Use multiple expansion with grouping in order to define the rows and columns of the heat map report type. For the following example shown above, (device),(part) expansion expression is used. Device will represent the columns and part the rows.

Heat map requires children nodes to display something, more precisely, a grouped expansion.

By default, the rows and columns will be sorted using the order of the children nodes. An option allows to change that sort. The smart sort will try to group values with same colors together (best effort).

Heat Map – Report Structure

Create the following:

- Parent:
 - ▶ Report type = Heat map
 - ▶ Filter expression
 - ▶ Sampling period, type, duration
 - ▶ Thresholds definition and parameters
- Child:
 - ▶ Group expansion
 - ▶ Formula or default metric context (from filter)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

29

Construction of a heat map report involves at least two report nodes. The parent report defines the heat map report type and some associated details under “Report Details: Heat Map” tab. In order to work correctly and to be valuable in a heat map report, a filter expression must select specific metric context (a metric that is common to multiple object and object types).

Remember, only one metric make sense with this report type. You should also specify what result you want to show in the cells. This is done by selecting the sampling period, type and the duration. By selecting, as an example, 1-hour, average over last 2 days, the computed value will be an average of the timeseries values over last 2 days using 1-hour aggregate data. For each cell, Report Engine will make the average calculation over $2 \times 24 = 48$ samples. If you have, let's say, 100 by 60 objects, that is a total of 288,000 calculations, which is still fair. If calculation would have been using Real-Time data (5' sampling) over a month, that would lead to over 51 millions of calculations... processing time (thus user waiting time) is a precious resource to save.

It is also on the Parent report node where you set the threshold definitions, that will color the cell area based on the metric value. Some parameters are also available like order of the elements, grid lines and gradient. Tooltips are helpful when user looking at the report want more information about a selected cell (just hover the mouse on a cell): metric value, and properties.

Child report node defines the expansion levels that will make the axes of the Heat map. The report type can be any; it will show up only when a user clicks on a cell of the heat map report.

Note that the chosen metric or KPIs elected to be used must be either filter on the metric context at the parent level or must be defined and presented as the default value from a formula, and defined at the child level.

Heat Map – Example

Heat Map of the VM Filesystems

- Unix VM Filesystem Heat Map (%Free)
- Windows VM Filesystem Heat Map (%Free)**

(vsaname),(vspname)

devtype=='VirtualMachine' & parttype=='FileSystem' & 'part=/'


Expansion

The specified properties will be used for the expansion of the current node

Group 1	
Simple:	device

Group 2	
Simple:	part

Report Type



heat map

Formula

math.Percentage → Results [%Free]

Axes Configuration

X Axis

Expansion Type: use report expansion

Y Axis

Expansion Type: use report expansion

Add a new threshold

Custom	Critical	10.0	
Custom	Warning	25.0	
Custom	Good	40.0	

Copyright © 2015 EMC Corporation. All Rights Reserved.
Module: More Report Types 30

The above is an example of construction of a heat map report.

On parent report node, there is a filter that set the metric context (you must go towards the root report level to retrieve the complete filter expression, not shown). Heat map report type is set at this level, along with time range period, aggregation level and aggregation function (avg, min, max, etc).

For this use case, the report designer chose to use a formula, where the result will be used for the creation of the Heat map report. Looking closer to the details of the formula, the result (%Free) is set to be the default value. This choice is an alternative to single metric selection.

“Report Details: The Heat Map” tab allows you to set parameters like the X and Y axis expansion to use but also some other options. Among those, you’ll find the “severity thresholds” definition that will fill the cell with the color based on the resulting metric/KPI value.

The Tooltip section will allow you to define which properties to show when the mouse is over a given cell, on the report.

Source: All>>Report Library>>VMware vCenter>>Situations to Watch>>Heat Map of the VM Filesystems>>Windows VM Filesystem Heat Map (%Free)

Advanced Report – Tree Map

Characteristics:

“Improved” Heat map report

Global view of infrastructure health


Cell size proportional to metric/KPI value

Cell color area based on metric/KPI value (threshold)

Spatial expansion only

Multi-level grouping (up to 3)

Improved Tooltips: properties and values



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

31


Tree map is similar to heat map but with more flexibility and information. It is meant to provide a global visual health report over the infrastructure. With the color code, you can quickly pinpoint areas where to focus attention.

The main differences with heat map are the grouping contexts and the cell areas that are proportional to the value compared with the total. The cell area color indicates the severity (based on threshold) of the value of the given object. The grouping context allows you to group objects of up to three levels.

Compared to heat map, only spatial expansion is possible, no temporal expansion.

The tooltip formatter allows you to complement the information when hovering over a cell of the report. Not only properties can be printed, but also some aggregated values, like it would show in a column of a standard table report.

Tree Map – 1 Level Example



The above is an example of a single level of tree map. This reports show all VMs that are running a hypervisor. Cell size is proportional to the number of CPU while color is according to the average CPU utilization of the report period.

Larger cell areas mean larger VM in terms of CPU. The above report is the “leaf” result of a drill down from clusters down to VMs on hypervisors (3 levels). The tooltip helps you better locate a selected VM, along with numbers.

Source: All>>Operations>>Workload Distribution>>Virtual Workload Distribution>>VMware vCenter>>CSE-Hopkinton>>Cpu Based Tree-map>>CSE-Cluster-CISCO>>mhmcb242.lss.emc.com

Tree Map - Definition

Base Tree map definition (1 level of expansion)

- Parent:
 - ▶ Report type = Tree map
 - ▶ (Filter expression)
 - ▶ Sampling period, type, duration
 - ▶ Number of levels = 1
 - ▶ Size: select metric (or formula) that will make the size
 - ▶ Color: select the metric (or formula) that will make the color
 - ▶ Thresholds definition (under value setting of Color)
- Child (level 1):
 - ▶ Expansion
 - ▶ (Formula)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

33

A basic definition of a Tree map report involves at least two report nodes in a parent / child structure. By basic, we mean use of one level of expansion, no groups definition (or “cluster”).


The parent holds the definition of the Tree map and the report standard definitions: sampling period (Real-time, 1-hour, 1-day, etc), type (average, min, max, etc) and the default duration (previous/last/current of day/week/month/etc).

The filter expression is not mandatory at this point, it could be set on another parent node. It is just a recall that it is good practice to set the right metric context for the report to show up correctly.

Then among the “Report Details: Tree map” tab, you will have to set the level of grouping (1,2, or 3). Select a metric that will be used for the size of the cell and another one that will be compared with a threshold definition and that will make the cell color.

The second required element is the child report node that will make the tree map cells. It is based on expansion of a property (device name, object type, cluster name, etc). This is also where you can have a formula that will return values for the above parent (defining size and/or color).

Tree Map – Multi-level Example




The above is an example of multi-level use with tree map report. The higher level aggregates the number from the lower, automatically (for assigned size metric). One can notice that the size of the VM is the starting point, but also the size of an hypervisor and in the end, the size of the cluster. Each can be compared in their respective nest: cluster, hypervisor, VM.

The above is showing average CPU load per VM, grouped by hypervisors and by cluster. You can, at a glance, have a view over the whole datacenter and spot the VM in problems.

Source: All>>Operations>>Workload Distribution>>Virtual Workload Distribution>>VMware vCenter>>CSE-Hopkinton

Tree Map – Cluster/Hypervisor/VM UC



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

35

Here we look at the construction of the VMware Tree map report. Let's first look at the layout of the report structure from the following source: All>>Operations>>Workload Distribution>>Virtual Workload Distribution. The VMware vCenter is set as an interactive table that makes use of the next underlying expansion set on "mixed report" node: DataCenter list (based on value of memberof property). This report lets you choose the datacenter you are interested in investigating using a Tree map report.

While the mix report is expanding on the DataCenter (memberof), it is forwarding you to two Tree map analysis's - CPU size and memory size. The construction is the same on both reports (Expand on memberof) - even the filter are the same. The difference is in parameterization of the Tree map, where one selects CPU and the second memory.

The first "Expand on memberof" (a) is tricky here because although, by its name, it lets you suppose there is an expansion on memberof, it actually hides the selected report entry from the previous level (mixed report, which you select a DataCenter from memberof property). In fact, this node hides the DataCenter name by the report name "Cpu-based Tree-map". This node is set to be the first level of Tree map, grouping data from three levels down, that is, per VM. This first Tree map report consumes the underlying expansion, by cluster (and not by DataCenter, as it would look like). The next report node (b), expands on cluster name and is set to be the second level of the Tree map report, grouping elements per hypervisor (the underlying expansion).

The next report node (c), expands on hypervisor by using complex expansion. This is because the metric context of Hypervisor is too restrictive and we want to recover VM metrics by making a map from hypervisor device name to hypervisor property, which any VM is tagged with. This report node is also set to be the Tree map level 3, which will consume the VM device expansion, just below. The next report node (d) expands on VM device name and is set as a standard report... actually it is an interactive table.

This organization of the report allows you to have a general overview at the top and lets you drill down from the DataCenter point of view, down to the individual VM.

Advanced Report – Geo Map

Characteristics:

- Pinpoint objects over a geographical map
- Provides a overview of how objects are spread geographically
- Pin can be color-coded based on metric tested against thresholds
- Pin size based on a metric value
- Tooltip enables to detail an object on the map
- Single level of expansion
- Report layout based on a table
- Meaningful only with a property holding location information
- Google map render will best match location
- Requires license and valid Google map account

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

36

Geomap reports are primarily used to pin objects on a geographical map. Different view sof objects can be used, but typically are device or datacenter related. Pin icon on the map can be presented in two flavors: size and color. Pin size will be relative to all of the pins. Pin color is tested against threshold values.


Tooltip shows more details about the pinned object.

The Geomap report makes use of a single level of expansion (as opposed to Treemap). The best way to build a geomap report is to start with a standard table report that can be automatically mapped in a geomap report. It is assumed that the location information is attached to the metric set prior to define a geomap report. Otherwise, the report will show nothing and default map will be presented.

If location property is not providing all the details about a precise location, Google map will try to locate on the map based on regional query (closest to where server is located in case of location name confusion (ex: same city name)).

It is also assumed that ViPR SRM Administrator configured correctly the frontend installation in registering a valid license (Geomap integration module). Default map rendered is Google map service, this one is *not* free for businesses.

Geo Map - Definition



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types 37

This report has the ability to display elements on a geographical map.

You **MUST** have in your property set a **valuable location information**. This information will be processed by Google Maps® engine.

There are five main attributes that can be set for a map report:

- **Location:** locates an element on the map.
- **Name:** used as the title in the popup window for the current element on the map.
- **Size:** sets the size of the icon on . The size of icon is based on a given metric value. For example: number of devices per site.
- **Color:** sets the color of the icon. This relies on threshold definition. For example: global average availability of devices per site.
- **Icon:** sets the icon aspect (circle, square, triangle, etc). This is a static selection given at report definition time.

The tooltip content is free and can hold (unlimited) values and/or properties information.

Geo Map Report – Construction

The screenshot illustrates the construction of a Geo Map Report in three steps:

- Define your report structure:** A green arrow points to the 'Geomap - Array location, size and count' section in the left sidebar, which contains '[Location]' and '[device]' fields.
- As a draft, define and test your report using a table type report:** A green arrow points to the 'Report Details: Table' tab, which is active. It shows a list of attributes: 'Attribute: Location*', 'Value: Total Raw Capacity*', 'Property: Customer*', 'Value: Device count*', and 'Value: Total Disk*'. Below this, the 'Map Item Configuration' section shows a table with columns 'Name', 'Location', 'Color', 'Size', 'Search', and 'Icon'. The 'Selected Column' is set to 'Use 'Customer' column'. The 'Displayed Information' section shows 'Attribute: Location'.
- Switch to map report type:** A green arrow points to the 'Report Details: Map' tab, which is selected. It shows the 'Map Item Configuration' section with the same table structure.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

38

Geographic Map reports can be seen as similar to a standard table. Thus it is a recommended practice to first build and test your report as a table and then switch to a map report and finally assign columns to the required map attributes (Name, Location, Color, and Size).

Geo Map Report – Location Override

Map Item Configuration

Map Item Configuration

Name Location Color Size Search Icon

Selected Column Use 'Location' column

Advanced Properties

Secondary Location Column Use 'Location' column

Location overrides

Santa Clara=Albany

Location overrides

The locations entered here will have precedence over the ones that are found in the selected data. The correct format is: name=location, one pair per line (i.e.: Watch4Net=4388 Saint-Denis, Montreal, Quebec, Canada, H2J 2L1).

Before:



After:



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

39

Advanced Properties on the Location tab of the Map Item Configuration view allows you to override locations that Google Maps could have difficulty with, but that are present in the database.

From the example above, relocate the “Santa Clara” office to “Albany”, just for demonstration purposes.

It can also be used to provide better precision on the location.

Module: More Report Types

Lesson Summary

During this lesson the following topics were covered:

Advanced report types:

- Heat map
- Tree map
- Geo map



This lesson covered advanced report types: Heat, Tree, and Geo maps.

Lab 11: Advanced Report Types

This lab covers Advanced type of reports:

- Part 1 – Heat map:
 - Heat map - Cisco ports usage
- Part 2 – Tree map:
 - a – Single level tree map
 - b – Double level tree map
 - c – Triple level tree map
 - d – VNX Array open slots
- Part 3 – Geo map report:
 - Geomap - Array location, size, and count

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

41

Lab 11 part 1,2, 3 will give you practice of advanced report types.

Lab 11 Debrief: Advanced Report Types

This lab exercise allowed you to:

- Familiarize yourself with advanced report types: heat, tree, and geo map
- Practice real life examples of using advanced report types

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

42

This lab exercise allowed you to familiarize yourself with Advanced report types: Heat, Tree, and Geo map, better understand the advanced report type and their use, and practice real life examples of using advanced report types.

Module: More Report Types

Lesson: Miscellaneous Report Types


This lesson covers the following topics:

- Icons and external report types
- Topology report



This lesson covers remaining report types: Icons, External, and Topology.

Icons Report Type



EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

44

Icons report type is mainly used for presentation of a list of reports. You could consider this report as a replacement of List report type. Icons report type simply lists the children reports (static or from an expansion) in 1, 2 or 3 columns format, selected from Display Mode. Along with the report name and an icon, the report description is shown, under the icon. Each child report is associated with an icon.

Source: All>>Explore>>Storage>>More Reports>>Enterprise Capacity

Icons Report Type - Definition

Explore **EDIT MODE**

- Storage
 - Storage Systems
 - More Reports**
 - Enterprise Capacity
 - File Capacity
 - Virtual Storage Capacity
 - Chargeback by Host
 - Service Level Capacity
 - Storage Pools
 - RAID/Parity Group
 - NAS Pools
 - More Storage Inventories
 - View by Service Level

More Reports | Edit Mode on ReportPack "Default Rep

Filtering & Expansion **Report Configuration**

Display Mode: three reports per line

Report Type: Icons

Filtering & Expansion

Advanced Parameters

Use filter on unmatched variables: No

Report Display Preference: Use Real Time Values

Node Icon: [Icon Selection]

Top N for events: [Input Field]

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: More Report Types 45

The parent report is defined as the Icons report type while each of the children are associated with an icon, along with the report name. For each child report you associate an icon by selecting the Node Icon parameter under Advanced Parameters of Filtering & Expansion tab. The parent also defines the layout, in 1, 2 or 3 columns. This is done by selecting Display Mode, under Report Configuration tab.

Source: All>>Explore>>Storage>>More Reports>>Enterprise Capacity

External Report Type

The screenshot displays the ViPR SRM user interface. On the left is a sidebar with a 'BROWSE MODE' button and a list of report types under 'Module 06 - More Report Types', including 'Lab 1 - Stacked Chart' through 'Lab 11 - GeoMap', and 'Extras', 'WIP', 'Explore', 'Operations', 'Planning', 'Report Library', and 'My Reports'. The 'Demo - External report' option is highlighted. The main content area shows a breadcrumb trail: 'All >> ViPR SRM Reporting class >> Module 06 - More Report Types >> Demo - External report'. Below this, the title 'Module 06 - More Report Types / Demo - External report' is displayed. The embedded content, highlighted by a green box and labeled 'iFramed URL content', is a webpage for 'EMC² ViPR SRM'. This webpage features the EMC² logo, navigation links for 'PRODUCTS & SOLUTIONS', 'SERVICES', 'SUPPORT', 'HOW TO BUY', and 'COMMUNITIES', and a section titled 'INFRASTRUCTURE MANAGEMENT' with the heading 'ViPR SRM'. It describes the software as 'SOFTWARE TO VISUALIZE, ANALYZE, OPTIMIZE STORAGE RESOURCES.' and includes a 'WATCH VIDEO' button and a link to an 'EMC Pulse Blog' article. The footer of the main interface shows 'Copyright © 2015 EMC Corporation. All Rights Reserved.', 'Module: More Report Types', and the page number '46'.


External report type is used to embed external and web site or web tool inside of ViPR SRM report page using iFrame. The page then becomes like a browser, allowing you to click and browse the content.

External Report Type - Definition

The screenshot displays the 'Demo - External report' configuration page in the EMC SRM Reporting interface. The left sidebar shows a tree view of report types, with 'Demo - External report' selected. The top navigation bar includes 'EDIT MODE' and icons for DISPLAY, EXPORTS, and TOOLS. The main content area has three tabs: 'Filtering & Expansion', 'Report Configuration', and 'Report Details: External'. The 'Report Details: External' tab is active, showing a text input field for the 'External report address' with the value 'http://www.emc.com/data-center-management/vipr-srr'. A 'Save' button is located below the input field. A blue footer bar contains the EMC logo and copyright information.

On the report node, to refer to an external page, simply select the External as the Report Type (Report Configuration tab). Next, under the Report “Details: External” tab, enter the URL of the site. If you are pointing to a web site where you can pass on some arguments, you can forge the URL accordingly using metric properties available on the current metric context of the report node. You can use **`${property}`** syntax to inject property values into the URL and **`@{attribute}`** syntax to inject report attribute values (node name, node ID, filter expression, child count, etc). Example: **`http://www.company.com/query?device=${device}`**.

Topology Report Type



Topology maps help visualize the connectivity relationships between hosts, fabrics, and arrays. Virtual machines are shown with their physical hypervisor hosts and data stores. You can find topology maps in the configuration item view of host or array objects in the use-case reports area, under the Explore menu option in the SRM Suite User Interface.

As an example, Topology map, in a single shot gives you an ability to traverse from the host, to its ports, to the Fabric(s) its connected to, all the way to the Storage Systems.

There is also more than just a picture: elements on the topology map are clickable and then brings you to available related reports. Thus it becomes a navigation tool. The report list will change according to the selected elements.

Map Types options allows to select the level of details:

- End to End Physical Connectivity view: allows you to see physical connectivity between objects.
- Masked Storage Systems view: built upon only Masking Views, useful when only interested in provisioned storage.
- Masked Storage Systems with Replicas view: built upon Masking Views and Replicas.

In addition to having multiple perspectives on topology maps, there are some overlays. For instance, if there are Alerts on a given device (Severe or High), an indicator would be shown on top of that device indicating as such. In order to not clutter the map though, this indicator will be shown only if there is a Severe or a High Alert on that device.

Tooltips are also providing some details upon hovering over the object icon: For Host, you would see some attributes of the host and a spark line for its CPU Utilization. For an Array, you would see the aggregated Port IOPs in addition to some vendor information.

Topology Report Type - Definition

EDIT MODE

Topology Map | Edit Mode on ReportPack "Default ReportPack", template "Explore"

Explore

- ▶ Storage
- ▶ Storage Connectivity
- ▶ **Hosts**
 - ▶ **Hosts List**

Topology Map

Device Summary
 - ▶ VMs

Topology Map

Device Summary

Filtering & Expansion
Report Configuration
✔ Report Details: Topology

1 The Simple and Advanced mode are based on the topology. The Drilldown mode is based on the subreports, while the Custom mode is completely based on an XML file.

Mode Custom ▼

1 This defines how the XML describing the report will be retrieved. In each case (file path or URL), dynamic values can be injected at runtime through the use of placeholders. In order to inject a property, use \$[property] (ex: http://www.company.com/query?mydevice=\$[device]). For topomap services that support it, "hop-count" and "map-type" can also be specified with their default values, using @{hop-count[default-value]} or @{map-type[default-value]} (ex: http://www.company.com/query?hopcount=@{hop-count:1}&maplayer=@{map-type}).

XML data Retrieval Remote XML through URL ▼

http://localhost:9876/srm/maps?device=\$[device]&dev

Rendering

Layout Hierarchical ▼ ?

Template

☐ APG Topology.xml (Database, Portal, Web Server, E

☐ Core Networking.xml (Switch, FileServer, Host, IMAP

☒ Storage Topology.xml (Switch, UnifiedStorageEmc, V

Copyright © 2015 EMC Corporation. All Rights Reserved.
Module: More Report Types 49

Topology map in ViPR SRM relies on an external component of the main Report Engine: Maps-Service and Topology-Service. Those service components are responsible to maintain the relationship database and to generate the report.

Based on this, Topology report type is simplified. The construction of a topology is shown above, from the Explore standard report library, under Hosts report.

If you look carefully, Topology Map report node is a child report of a Host List (or VMs [list]). This parent report expands on device. The metric context for the Topology Map is only for one device. The key details are passed to the Maps-Service thru an URL, that will generate the topology. The URL is defined in under "Report Details: Topology" tab. In the Rendering section, the template supported for storage is "Storage Topology.xml" file descriptor. Others are available for other application use.

Module: More Report Types

Lesson Summary

During this lesson the following topics were covered:

Miscellaneous report types:

- Icons and external report types
- Topology report



This lesson covered remaining report types: Icons, External, and Topology.

Module: More Report Types

Lesson: Publishing Your Reports

This lesson covers the following topics:

- ReportPack concepts
- Building your own ReportPack



This lesson covers ReportPack feature.

ReportPack

Characteristics:

- Composed of one or many report templates
- Is the place for public, sharable custom reports
- Is convenient and flexible for custom report repository
- Can be placed directly on the root report node tree
- Read and Edit rights can be applied (through Roles)
- Can be exported and backed-up
- Rule of thumb: never edit reports from Report Library
- Always work on a copy of any report from the Report Library



Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

52

The ReportPack allows you to place and centralize custom reports. This could become a place to share report development among a group of users. Combined with roles, you can define users who can read only reports from those who can alter template reports.

A ReportPack is made of one to many report templates. A report template could define a deep and complete report structure from all the richness of the reporting features available with SRM.

A ReportPack can be extracted from the report structure, backed-up and kept in a safe place of your own (saved, versioned, etc).

The main use case for using ReportPacks is for custom reporting from the Report Library. Never replace/change any base report from the Report Library. You never know when later updates can be applied to those reports, thus overwriting your customizations. Always take a copy of the base report back to your personal sandbox before making any modifications. Work out your report and then, when ready, place them in a Template of a ReportPack.

The report's structure you defined then allows other users to use the report... or modify the report(s). This is done thru a Role-based definition (ViPR SRM Administrator).

The templates of the defined ReportPack can be placed directly on the root node of the tree structure, or placed under an existing node ID, complementing the existing report node.

ReportPack Management

ReportPacks Management | Administration

Complete ReportPacks (in arp format) can be imported through the following form. The list below contains the ReportPacks that have been installed or created.

Select the corresponding file and then click on Upload for the ReportPack name, version, description, template and formulas to be processed, imported, and made available.

[Upload a ReportPack](#)

Click on any line to view details for a specific ReportPack.

[+ New ReportPack](#)

Name	Description	Number of Templates
Brocade Fabric Switch	This SolutionPack contains reports that are specifically designed for the Brocade Fabric Switch devices	1
Cisco MDS Nexus	SolutionPack for Cisco MDS/Nexus	1
Default ReportPack	The default ReportPack, included with any Watch4net install.	10
EMC VMAX	Watch4net Solutions ReportPack for EMC VMAX (formerly Symmetrix)	2
EMC VNX	SolutionPack for EMC VNX	3
EMC Watch4net Health	This SolutionPack enables you to monitor the performance and health of all the Watch4net components.	1
Physical Hosts		2
VMware vCenter	SolutionPack for VMware vCenter	4

Click on the ReportPack name to see the templates

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: More Report Types 53

Typically, a ReportPack is a set of SRM reports. These reports are named “Templates”. ReportPacks allow you to group many reports together.

From Administration Menu, you have access to an entry named “ReportPack” from the left-hand menu.

Upload a ReportPack: Complete ReportPacks (in arp format) can be imported through the form. Select the corresponding file and then click on’ Upload’ for the ReportPack name, version, description, template, and formulas to be processed and imported into W4N.

New ReportPack: You can create a ReportPack on your own.

ReportPack Management

ReportPacks

- New ReportPack
- Brocade Fabric Switch
- Cisco MDS Nexus
- Default ReportPack
- EMC VMAX
- EMC VNX
- EMC Watch4net Health**
- Physical Hosts
- VMware vCenter

Click on the EMC Watch4net Health Template to see the details

ReportPack Modification: EMC Watch4net Health | Administration

This section allows the creation and modification of ReportPacks.

ReportPack Information

This information is for administration purposes only. It will not be displayed in the user interface.

Name: EMC Watch4net Health

Version: 1.3

Description: This SolutionPack enables you to monitor the performance and health of all the Watch4net components.

[Delete](#) [Export](#) [Save](#)

This list contains all the templates that have been created for this ReportPack. [New Template](#)

Template Name	Description
EMC Watch4net Health	This SolutionPack enables you to monitor the performance and health of all the Watch4net components.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: More Report Types 54

The templates are SRM reports generally designed for a ReportPack. The template list contains all the templates that have been created for the ReportPack. In this form, you can either save, export, or delete the ReportPack. Templates can be added here. By clicking any template, you can access the template management where you can change the description. The list contains report Templates that compose the current ReportPack. Click any Template from the table in order to view details of a specific Template, to modify its description, export, or delete it.

ReportPack Management

Templates within a ReportPack can be placed directly on the main report structure or within it

ReportPack Modification: Default ReportPack | Administration

This section allows the creation and modification of ReportPacks.

ReportPack Information

Template Name	Description	
Dashboard	Enterprise wise reports	◆
Device Summary	Legacy - A summarized overview of all the devices on the network.	◆
Explore	Explore your IT infrastructure	◆
Global Reports	Legacy - This node contain a set of general reports.	◆
Operations	Situations to watch, Alerts, Compliance, etc.	◆
Planning	Capacity planning reports	◆
ReportPacks	This node helps you to hook ReportPacks.	◆
SolutionPacks	This node helps you to hook SolutionPacks.	◆
Templates	This report contains template to be reused in Global Reports & SolutionPacks reports	
Transactions and Applications	Legacy - This set of reports is designed to help monitoring the performance of transactions and applications	

-
- ▶ Scheduled Reports
 - ▶ Stored Reports
 - ▶ Favorite Reports
 - ▶ Dashboards
 - ▶ Explore
 - ▶ Operations
 - ▶ Planning
 - ▶ Report Library
 - ▶ Brocade FC Switch
 - ▶ Cisco MDS/Nexus
 - ▶ EMC VMAX
 - ▶ EMC VNX
 - ▶ EMC Watch4net Health
 - ▶ Physical Hosts
 - ▶ VMware vCenter

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

55

Once a template is part of a ReportPack, it can be positioned on the main root level or within an existing branch (provided by a UID). This way you can better organize the global report structure. If you don't select any "hook" option, the template will be placed directly under the root level of the global report structure. This setting is within a template.

Steps to Create a ReportPack


Step	Description
1	(Edit Mode) Export your reports to XML format -> Templates
2	Go to Administration Panel under the ReportPacks node
3	Create a new ReportPack
4	Create and import templates
5	Export the ReportPack



The basic steps to create a ReportPack are shown here.

ReportPack Management – Step 1

- From Edit Mode
- Highlight and export the report you need to compose your ReportPack
- Each report will become a **template** entry part of your ReportPack



In Edit mode, highlight your Report and select the Report Definition button from the Exports menu

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

57

1. From Edit Mode, select and Export your Reports. Each extracted export becomes a template, even if you select “Report Template” or “Report Definition”.

The difference between the two options is that if you know you are using Links inside a given report structure, “Report Template” will keep those references. On the other hand, “Report Definition” will resolve the links and replace with the real report definition. This will inflate the file size and you lose all of the advantages of reusability.

ReportPack Management – Demo – Step 2 and 3

The screenshot displays the 'ReportPacks Management' interface under the 'Administration' tab. On the left, a navigation menu lists 'Administration' (Portal, Users, Profiles, Roles, Centralized Management) and 'ReportPacks' (Modules). The main content area is titled 'ReportPacks Management | Administration' and includes an 'Upload a ReportPack' button. A dashed box contains instructions: 'Complete ReportPacks (in arp format) can be imported through the following form. The list below contains the ReportPacks that have been installed or created.' Below this, another instruction states: 'Select the corresponding file and then click on Upload for the ReportPack name, version, description, template and formulas to be processed, imported, and made available.'

The 'ReportPack Creation | Administration' form is shown, with a sub-header 'This section allows the creation and modification of ReportPacks.' A note indicates: 'This information is for administration purposes only. It will not be displayed in the user interface.' The form fields are: 'Name' (My Custom ReportPack), 'Version' (0.1), and 'Description' (This is an optional ReportPack description (but recommended). Version number is for you own purpose). A 'Save' button is at the bottom. A green box highlights the 'New ReportPack' button in the top right, and a green arrow points from the 'Description' field to a text box that says: 'Give your new ReportPack a name, version, and description (optional) ...and save!'.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved. Module: More Report Types 58

2. Go to the Administration Panel under the ReportPack menu.
3. Create a New ReportPack: fill-in the details (ReportPack name, Version, Description) and Save your ReportPack definition.

ReportPack Management – Demo – Step 4

ReportPack Modification: My Custom ReportPack | Administration

This list contains all the templates that have been created for this ReportPack.

Template Name	Description
No template found for this ReportPack.	

+ New Template

Report Template Creation | Administration

This section allows the creation and modification of Report Templates.

Template Information

Information other than the tree template is for administration purposes only. It will not be displayed in the user interface.

Template Name

Version

Description

Hook Nodes

Hook to Another Node ?

Tree Template

Choose File No file chosen ?

Fill in the details about the template (name, version, description)

Several options are available to position your report

Click Browse and retrieve your XML report template (previously exported)

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

59

4. Start to incorporate your template. Click new Template to define one. Fill in the (optional) details. Then import your template using “Choose File” (your template must be local to your browser). Select the Hook Nodes option to give you the ability to specify where the template will appear in the tree (could be multiple places). When completed, select Save.

ReportPack Management – Demo – Step 5

Keep your work in a safe place or under revision control

Export your work to an ARP package (SRM ReportPack format)

ARP = APG Report Pack

ReportPack Modification: My Custom ReportPack | Administration

This section allows the creation and modification of ReportPacks.

ReportPack Information

This information is for administration purposes only. It will not be displayed in the user interface.

Name: My Custom ReportPack

Version: 0.1

Description: This is an optional ReportPack description (but recommended). Version number is for your own purpose.

[Delete](#) [Export](#) [Save](#)

This list contains all the templates that have been created for this ReportPack.

[+ New Template](#)

Template Name	Description
Customer and Business Unit VM report	This report is about Customer and Business Unit VM report.

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

60

5. Saving your ReportPack to ARP (Default SRM Export Format) makes your ReportPack compatible for other SRM installations. This a good practice to keep a copy in safe place and under revision control system (CVS-like).

Custom ReportPack

The screenshot displays the EMC² interface. On the left, a navigation pane lists various report categories, with 'List of Virtual Machines' highlighted. A green arrow points from this item to a green-bordered box containing the text: 'ReportPack template now shows on main branch'. On the right, the 'Role Modification: Base role | Administration' window is open, specifically the 'Template Access' tab. It contains two dropdown menus: 'RP: My Custom ReportPack' and 'Template: Customer and Business Unit VM report', both currently set to '(No Access)'. A green arrow points from the second dropdown to another green-bordered box containing the text: 'ReportPack and template can now be access controlled through Template access of a role'. A descriptive note at the top of the role modification window explains the access behavior: 'This defines which kind of access this role will provide. Basically, it is a bottom to top behavior. It means that if the template one is defined, it will be the ReportPack one. And finally, as last resort, if global access will be used.'

Now you have a custom ReportPack, you can see that RP in the report tree. You can control who accesses the ReportPack and its templates by setting the roles accordingly.

Module: More Report Types

Lesson Summary

During this lesson the following topics were covered:

- ReportPack concepts
- Building your own ReportPack



This lesson covered ReportPack feature.

Lab 12: ReportPack Creation

This lab covers creation of a ReportPack:

- Create an exportable report structure
- Create a ReportPack
- Visualize the ReportPack on the top report tree of the demo server

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module: More Report Types

63

This Lab 12 will allow you to practice ReportPack creation and exportation.

Lab 12 Debrief: ReportPack Creation

This lab exercise allowed you to:

- Familiarize yourself with the ReportPack concept
- Package report presentations
- Publicize report structures



This lab exercise allowed you become familiar with the ReportPack concept, package a report presentation , and publicize a report structure.

Module Summary

Key points covered in this module:

- Report types definition and use cases:
 - ▶ Aggregated report types
 - ▶ TopN report types
 - ▶ Advanced report types
 - ▶ Miscellaneous report types
- ReportPack concept and usage



This module covered the use of various report type along with describing their use case context. Also covered the ReportPack concept.

Course Summary

Key points covered in this course:

- ViPR SRM report: design, structure, editing
- Report types: dashboard, inventory, advanced types
- Formula concept: definition and application
- Features of portal (user and administrator standpoint)
- Data inspection, selection, and manipulation
- Data Enrichment feature
- Report Packaging

EMC²

Copyright © 2015 EMC Corporation. All Rights Reserved.

Module #: Module Name

66

This course covered the subject of Reporting using ViPR SRM. This included all aspects from the user standpoint to report designer and some administration features.

Thank You!



Copyright © 2015 EMC Corporation. All Rights Reserved.