



oVirt Data Warehouse

02/11/11

Yaniv Dary
BI Software Engineer, Red Hat

oVirt Data Warehouse Overview

- API for historical configuration and statistical data .
- Used to create SQL reports using any tool.
- Currently provides information on host, vm, disks, storage domain and more.
- Uses ETL (**E**xtract **T**ransform **L**oad) service to receive data from the ovirt engine database.

Importance of Business Intelligence

- Right decisions on solid data.
- Know everything about your product.
- Preparing for the future.
- Summarizing and simplifying large data.
- Integral part of managing virtualization.

oVirt Engine DWH Views

- The engine API for the ETL.
- Creates a more user friendly interface of the engine tables.
- Gaps the differences between the engine to the history database tables.
- First stage in data transformation.

Talend Open Studio

- Powerful and versatile open source solution for data integration.
- Easy-to-use graphical development environment.
- Support for all types of data integration, data migration and data synchronization operations.
- Strong open source community.
- Based on the eclipse UI.

Talend Open Studio



The screenshot displays the Talend Open Studio interface. The top toolbar includes icons for navigation and actions, along with a zoom level of 50%. The main workspace shows a job design for 'Job HistoryETL 3.0'. The design is divided into PreJob, Main, and PostJob sections. The PreJob section contains components like tPrejob_1, tRowGenerator_1, tFileDelete_1, tFileInputProperties_1, tContextLoad_1, and tLogRow_2. The Main section features a complex flow with components such as tJavaFlex_1, tJavaFlex_2, tJavaFlex_3, tWaitForFile_1, tIterate_1, tRowGenerator_1, and tLogRow_1. The PostJob section includes tPostjob_1, tContextLoad_1, and tLogRow_1. The left sidebar shows the 'Repository' and 'Job Designs' sections, listing various job designs like 'AggregationToDaily 3.0', 'AggregationToHourly 3.0', 'ConfigurationSync 3.0', 'DailyTimeKeepingJob 3.0', 'DeleteTimeKeepingJob 3.0', 'HistoryDelete 3.0', 'HistoryETL 3.0', 'HostConfigCheck 3.0', 'HourlyTimeKeepingJob 3.0', 'MinimalVersionCheck 3.0', 'ParallelRun 3.0', 'SampleRunJobs 3.0', and 'SampleTimeKeepingJob 3.0'. The bottom status bar shows '0 errors, 0 warnings, 0 infos' and a table with columns 'Description' and 'Resource'. The table content is as follows:

Description	Resource
Errors (0 items)	
Warnings (0 items)	
Infos (0 items)	

TOS demo

DB Tables and API Views

- Sample data is collected at the end of every minute and is kept for up to 48 hours.
- Hourly level is aggregated every hour for the hour before last and is kept for 2 months.
- Daily level is aggregated every day for the day before last and is kept for 5 years.
- Upgrade is done via reentrant PL/pgSQL files.

Roadmap

- Support hibernate create database to enable cross database compatibility.
- Add fedora packaging.
- Add more columns of statistics.
- Add more sources other than ovirt-engine database.

oVirt

Q & A



oVirt Reports

02/11/11

Yaniv Dary
BI Software Engineer, Red Hat

JasperReports and JasperReports Server Overview



- JasperReports is an open source Java reporting tool that can produce reports and export them to PDF, HTML, Microsoft Excel, RTF, ODT, Comma-separated values and XML files.
- It generates the reports from an XML or .jasper file.
- JasperReports Server is a reporting server for JasperReports. Using it you can generate, organize, secure and deliver interactive reports.

oVirt Reports Overview

- oVirt reports package provides a suite of pre-configured reports that enable you to monitor the system.
- The reports module is based on JasperReports and JasperServer.

Reports Example




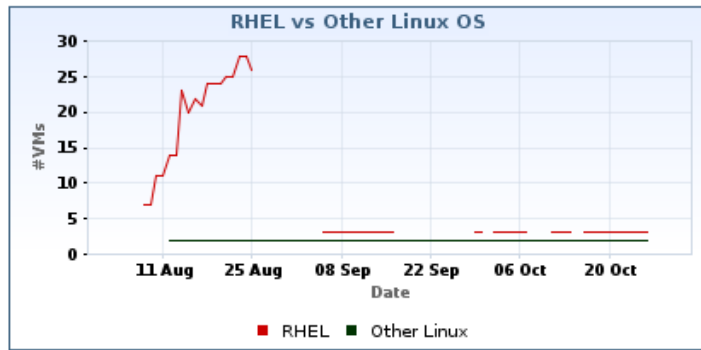


oVirt Reports

Oct 31, 2011

Active Virtual Machines by OS in Clusters of Data Center Default

Criteria: Datacenter: Default Date Range: 2011-08-01 - 2011-10-31 VM Type: All
 Cluster: All Period: Quarterly Show Deleted Virtual Machines: Yes



Input Controls

* Show Deleted Entities?: Yes

* Data Center: RHEVM-3

* Cluster: RHEVM-3

* VM Type: Server

* Period Range: Monthly

* Select Month: August 2011

* Start Date: 2011-08-01

* End Date: 2011-08-31

Apply OK Reset Cancel

oVirt Reports Topics

- Executive – High level reports that highly summarize data on resources and system entities.
- Inventory – Reports that list entities or display data about there status.
- Service Level – Report that are used to examine system health and service quality to users.
- Trends – Graphs that indicate resource usage for planning of resources.

Reports Tree



Folders

- root
 - Organizations
 - oVirt Reports
 - Reports
 - Executive**
 - Inventory
 - Service Level
 - Trend
 - Resources
 - temp
 - Themes
 - default
 - ovirt-reports-them

Repository

Sort By: **Name** | Modified Date

Run Edit Open Copy Cut Paste Delete

Active Virtual Machines by OS (BR18) /organizations/ovirtreports/Reports/Executive/active_vms_by_os_br18	The report contains comparative measurements number of running virtual machines and OS usage in for a selected cluster and a selected virtual machine's type within the requested period.	Report	October 18
Cluster Capacity Vs Usage (BR19) /organizations/ovirtreports/Reports/Executive/cluster_capacity_vs_usage_br19	This report contains charts displaying host's resources usage measurements (CPU core; physical Memory) and charts displaying virtual machine's resources usage measurements (virtual machine's total vCPU; Virtual Memory size) for a selected cluster.	Report	October 18
Host OS Break Down (BR22) /organizations/ovirtreports/Reports/Executive/host_os_break_down_BR22	This report contains a table and a chart displaying the number of hosts for each OS version for a selected cluster within a requested period.	Report	October 18
Summary of Host Usage Resources (BR17) /organizations/ovirtreports/Reports/Executive/summary_of_host_usage_resources_br17	The report contains a scattered chart of CPU and memory usage data within a requested period and for a selected cluster.	Report	October 18

IReport Demo

Roadmap

- Hibernate data source based reports.
- Adding more reports and reports types (Like events log reports).
- Adding localization to reports.
- Web admin integration.
- Integrated authentication with engine.

How To Contribute?

DWH

Git: `git://gerrit.ovirt.org/ovirt-dwh`

Wiki: http://ovirt.org/wiki/Ovirt_DWH

mailing list: engine-devel@ovirt.org

Reports

Git: `git://gerrit.ovirt.org/ovirt-reports`

Wiki: http://ovirt.org/wiki/Ovirt_Reports

mailing list: engine-devel@ovirt.org

oVirt

THANK YOU !

<http://www.ovirt.org>



oVirt Data Warehouse

02/11/11

Yaniv Dary
BI Software Engineer, Red Hat

oVirt Data Warehouse Overview



- API for historical configuration and statistical data .
- Used to create SQL reports using any tool.
- Currently provides information on host, vm, disks, storage domain and more.
- Uses ETL (**E**xtract **T**ransform **L**oad) service to receive data from the ovirt engine database.

Importance of Business Intelligence



- Right decisions on solid data.
- Know everything about your product.
- Preparing for the future.
- Summarizing and simplifying large data.
- Integral part of managing virtualization.

oVirt Engine DWH Views



- The engine API for the ETL.
- Creates a more user friendly interface of the engine tables.
- Gaps the differences between the engine to the history database tables.
- First stage in data transformation.

Talend Open Studio



- Powerful and versatile open source solution for data integration.
- Easy-to-use graphical development environment.
- Support for all types of data integration, data migration and data synchronization operations.
- Strong open source community.
- Based on the eclipse UI.

The screenshot displays the Talend Open Studio interface. The main workspace shows a job design for 'Job HistoryETL 3.0' with various components like 'tContextLoad', 'tFileDelete', and 'tFileInputProperties' connected in a flow. The left sidebar contains a 'Repository' view with 'Business Models' and 'Job Designs' listed. The bottom panel shows the 'Designer | Code' view with a 'Description' and 'Resource' table.

Description	Resource
Errors (0 items)	
Warnings (0 items)	
Infos (0 items)	

TOS demo

DB Tables and API Views

- Sample data is collected at the end of every minute and is kept for up to 48 hours.
- Hourly level is aggregated every hour for the hour before last and is kept for 2 months.
- Daily level is aggregated every day for the day before last and is kept for 5 years.
- Upgrade is done via reentrant PL/pgSQL files.

Roadmap



- Support hibernate create database to enable cross database compatibility.
- Add fedora packaging.
- Add more columns of statistics.
- Add more sources other than ovirt-engine database.

oVirt

Q & A



oVirt Reports

02/11/11

Yaniv Dary
BI Software Engineer, Red Hat

JasperReports and JasperReports Server Overview



- JasperReports is an open source Java reporting tool that can produce reports and export them to PDF, HTML, Microsoft Excel, RTF, ODT, Comma-separated values and XML files.
- It generates the reports from an XML or .jasper file.
- JasperReports Server is a reporting server for JasperReports. Using it you can generate, organize, secure and deliver interactive reports.

oVirt Reports Overview



- oVirt reports package provides a suite of pre-configured reports that enable you to monitor the system.
- The reports module is based on JasperReports and JasperServer.

Reports Example



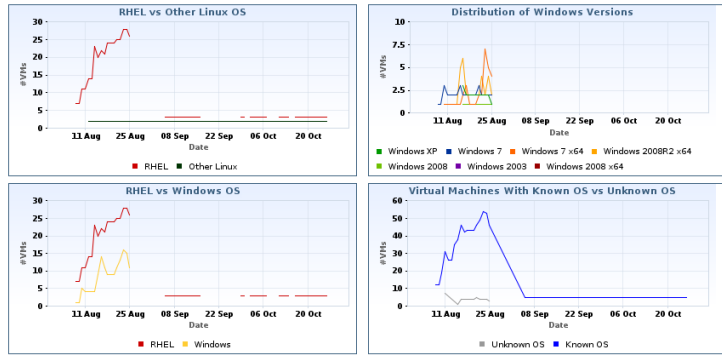
oVirt

oVirt Reports

Oct 31, 2011

Active Virtual Machines by OS in Clusters of Data Center Default

Criteria: **Datacenter:** Default **Date Range:** 2011-08-01 - 2011-10-31 **VM Type:** All
Cluster: All **Period:** Quarterly **Show Deleted Virtual Machines:** Yes



Page 1 of 1

Input Controls

* Show Deleted Entities?:

* Data Center:

* Cluster:

* VM Type:

* Period Range:

* Select Month:

* Start Date:

* End Date:

Apply OK Reset Cancel

oVirt Reports Topics



- Executive – High level reports that highly summarize data on resources and system entities.
- Inventory – Reports that list entities or display data about their status.
- Service Level – Reports that are used to examine system health and service quality to users.
- Trends – Graphs that indicate resource usage for planning of resources.

Reports Tree



oVirt Reports ovirt-admin | Log Out

View Manage

Folders

- root
 - Organizations
 - oVirt Reports
 - Reports
 - Executive**
 - Inventory
 - Service Level
 - Trend
 - Resources
 - temp
 - Themes
 - default
 - ovirt-reports-them

Repository Sort By: Name | Modified Date

Name	Description	Type	Modified Date
Active Virtual Machines by OS (BR18) /organizations/ovirtreports/Reports/Executive/active_vms_by_os_br18	The report contains comparative measurements number of running virtual machines and OS usage in for a selected cluster and a selected virtual machine's type within the requested period.	Report	October 18
Cluster Capacity Vs Usage (BR19) /organizations/ovirtreports/Reports/Executive/cluster_capacity_vs_usage_br19	This report contains charts displaying host's resources usage measurements (CPU core; physical Memory) and charts displaying virtual machine's resources usage measurements (virtual machine's total vCPU; Virtual Memory size) for a selected cluster.	Report	October 18
Host OS Break Down (BR22) /host_os_break_down_br22	This report contains a table and a chart displaying the number of hosts for each OS version for a selected cluster within a requested period.	Report	October 18
Summary of Host Usage Resources (BR17) /summary_of_host_usage_resources_br17	The report contains a scattered chart of CPU and memory usage data within a requested period and for a selected cluster.	Report	October 18

IReport Demo

Report creation demo.
Reports structure.
Pro edition demo
Themes and charts.

Roadmap

- Hibernate data source based reports.
- Adding more reports and reports types (Like events log reports).
- Adding localization to reports.
- Web admin integration.
- Integrated authentication with engine.



How To Contribute?

DWH

Git: <git://gerrit.ovirt.org/ovirt-dwh>

Wiki: http://ovirt.org/wiki/Ovirt_DWH

mailing list: engine-devel@ovirt.org

Reports

Git: <git://gerrit.ovirt.org/ovirt-reports>

Wiki: http://ovirt.org/wiki/Ovirt_Reports

mailing list: engine-devel@ovirt.org

oVirt

THANK YOU !

<http://www.ovirt.org>