



⚠ RayVentory Data Hub 12.3

12.3.4547.69 [Update 07]

Released on Nov 14, 2022

Resolved issues

- **DATA HUB AGENT** All third-party libraries are signed in this build. [RR-3257](#)
 - **DATA HUB AGENT** Fixed a possible DLL hijacking in case of non-standard Windows configuration of PATH environment variables. [RR-3257](#)
-

12.3.4543.67 [Update 06]

Released on Oct 11, 2022

Improvements

- Improved performance of ODBC queries, which - depending on the underlying driver - are now processed up to 10 times faster. [RR-3408](#)

Resolved issues

- Fixed various bugs in cloud connectors (Update 2022.2 Patch 2). [RR-3373](#)
 - Fixed invalid URL returned by the Azure API. [KOT-134](#)
 - Fixed bad request exception in SAP Qualtrics connector. [KOT-107](#)
 - Fixed bad support token format. [KOT-5](#)
 - Fixed Google Cloud BigQuery connector that was returning items marked for deletion. [KOT-143](#)
 - Improved stability and minor technical changes. [KOT-139](#) [KOT-121](#)
-

12.3.4540.41 [Update 05]

Released on Sep 12, 2022

New Cloud Connectors (Update 2022.2 Patch 1) [RR-3262](#)

In this version, the following new connectors have been added:



Workday



Huawei Cloud Stack



Nutanix Prism Central



OpenStack



Proxmox



Qualys



SAP Hybris Cloud for Customer (SAP C4C)

Security improvements

- Fixed a security issue with handling of user sessions, which - under rare circumstances - when using correct timing and certain unlikely preconditions could lead to an issue known as “session hijacking”. [RR-3315](#) [ZEN-20095](#)
- Fixed a security issue, where JWT sessions tokens could contain encrypted sensitive information. [RR-3315](#) [ZEN-20095](#)

⚠ Issue **RR-3315** means that an encrypted token, derived from user password, could be saved in application local storage. Successful use of this vulnerability requires physical access to the user machine and successful reverse engineering of RayVentory Data Hub code. This release removes the offending values from the security token for general Data Hub sign-in purposes.

Adjustments in this version mean, that the encrypted tokens are never written in typical Data Hub authentication flow. These are however still required and present when LDAP-based authentication is enabled (Windows builds only).

12.3.4537.38 [Update 04]

Released on [Jul 22, 2022](#)

This build introduces fixes and minor improvements for cloud/SaaS connector (version 2022.1 patch 5) and security improvements.

Resolved issues

- Resolved multiple issues which may lead to catalog recognition leaving some raw data unrecognized. [RR-3172](#)
- Fixed connector Microsoft 365 returning error `HTTP Exception 400 Bad Request` when fetching app usage data. [RR-3242](#)

Security improvements

- Fixed a security issue (CVSS 3.1 scoring = 4.3) where a user outside of the Data Administrator group could get access to the list of tasks via REST API. [RR-2862](#)
- Fixed a possible XSS injection in the history of changes view, where the input from the user was not properly escaped (CVSS 3.1 scoring = 4.6). [RR-3181](#)

- Fixed a possible XSS injection using the file storage (CVSS 3.1 scoring = 4.6). [RR-2860](#)
 - Removed detailed stack trace from exceptions in case finding a task by table name fails. (CVSS 3.1 scoring = 5.3). [RR-2859](#)
-

12.3.4530.37 [Update 03]

Released on [Jun 30, 2022](#)

This build introduces fixes and minor improvements for cloud/SaaS connector (version 2022.1 patch 4).

Resolved issues

- Airwatch logging now distinguishes Applications and ArpForDevices. [KOT-102](#)
 - Inventory endpoint for Google Cloud VM Manager has improved paging algorithm, which relies on tokens instead of offsets. [KOT-90](#)
 - Added a new parameter `directOnly` for Adobe Admin Console connector. This flag is used to return only those users who have a direct membership in the product profile. [KOT-99](#)
 - Fixed an issue where it was impossible to set private keys for Google Cloud connector. [KOT-103](#)
 - Internal documentation of various endpoints used by the DocuSign connector. [KOT-78](#)
-

12.3.4529.36 [Update 02]

Released on [May 25, 2022](#)

This build introduces fixes and minor improvements for cloud/SaaS connector (version 2022.1 patch 3).

Resolved issues

- Fixed forbidden dependencies in connector Big Query. [KOT-77](#)
 - Fixed an exception `NoSuchMethodException` in debug logging. [KOT-84](#)
 - Fixed `OAuth2TokenCache` bug in SAP Qualtrics. [KOT-82](#)
 - Fixed usage of deprecated API version for SOAP login in connector Salesforce. [KOT-73](#)
 - Fixed missing organizational unit hierarchy in connector AWS Organizations. [KOT-68](#)
 - Fixed missing `UserId` property for user targets in connector Intune. [KOT-72](#)
 - Added missing parameter `serverUrl` for connector 15five. [KOT-62](#)
 - Fixed missing `LocaleSidKey` in connector Salesforce. [KOT-50](#) [KOT-89](#)
-

12.3.4525.33 [Update 01]

Released on [Apr 28, 2022](#)

Feature Highlights

Active Directory / LDAP Authentication [RR-80](#) [RR-2732](#)

It is now possible to relay the complete configuration and group assignment to the Active Directory. In this hybrid mode, the administrator creates local groups in Data Hub and assigns them (via a full DN string) to Active Directory groups. The assignment of users to groups is then managed by the manipulation of AD group assignment.

A configurable setting dictates whether only selected users or all AD users should be able to log in.

New Cloud Connectors (Update 2021.5 Patch 4) RR-2968 RR-2087

The following connectors have been added in this update:

- Cisco Meraki COKCON-725
- Cisco Umbrella COKCON-717
- Typetalk COKCON-697
- JamaSoftware COKCON-684
- Stripe COKCON-680
- Falcon COKCON-663
- ActiveCampaign COKCON-586
- Google Cloud Compute Instances COKCON-383
- Cisco Webex COKCON-41
- Google Kubernetes COKCON-386
- Google Cloud Storage Download COKCON-385
- Google Cloud Database Instances COKCON-384
- Microsoft Intune Reports COKCON-730
- Dot4Cmdb COKCON-478
- VMware vSphere COKCON-704
- Vonage Unified Communications COKCON-723
- G2 COKCON-721
- RankWatch COKCON-655

Advanced Transformation Editing (JSON) RR-3057 RR-3007

It is now possible to edit transformation steps using a raw, JSON editor. This is also a bridge to a CLI usage of transformations (with help of ETL CLI), which makes simple copy-and-paste, import, and exporting even easier.

Ability to Download Log Files from the Task View RR-2933

Complete log files from agent execution can now be downloaded from the Task view. This feature has a configurable option, specifying how long the logs should be stored. Additionally, an excerpt of each log is stored permanently, even for tasks that have already been cleaned-up.

Other Improvements

- Logs for transformation tasks are now more detailed and easy to follow. For example, the executed steps are referenced by their names and each step contains extra information about tables used and produced, the number of input and output rows, etc. RR-2965 RR-2992
- Optimized the disk space consumption by the transformation process. Depending on the size and structure of the transformation, up to 80-90% of savings can be expected. RR-2891
- When migrating from 12.2 to 12.3 Update 1, the existing transformation definitions will now be automatically migrated to the new format. RR-2943
- It is now possible to change the temporary location for transformation files and configure some further diagnostic options. RR-2999 RR-2963
- Several internal changes in the docker images and default compose file for idiomatic usage of environment variables. RR-3004
- Improved performance of VMware AirWatch connector. COKCON-767

Resolved Issues

- Improved performance and reworked storing of agent logs from the Data Hub Agent. This issue caused visible UI lags and timeouts when accessing the task list and task history page in case of big agent logs with multiple repetitions. [RR-2933](#)
- Fixed an issue with improper schedule settings (daily vs hourly). [RR-2998](#)
- Fixed a UI issue where an infinite loop could be triggered by creating a step that references itself as a source. [RR-3013](#)
- Changed the default place where Azure Active Directory reports and tasks are imported. The correct place is the Directory Services folder. [RR-2934](#)
- Fixed various issues with the Transformation editor, where some combinations of inputs could produce invalid JSON content. [RR-2913](#)
[RR-2917](#) [RR-2924](#)
- Fixed an issue with bad performance of the task view in case of several runs with big agent log tails. [RR-2933](#)
- Fixed an issue, where default database log appender would write no logs in the Docker installation. [RR-3043](#)
- Fixed an issue, where saving a report in a Docker-hosted instance would show an error. [RR-3068](#)
- Fixed an issue, where a failed transformation (in server mode) was not properly reported to the backend. [RR-2920](#)
- Fixed an issue with improper values in the Output column in the Transformation editor. [RR-3088](#)
- Resolved several contradictory information in the User Guide. [RR-2944](#)
- Removed unnecessary API calls between UI and the backend. [RR-3004](#)
- Fixed not working domain fetching by the Veeva Vault connector. [COKCON-750](#)
- Fixed Azure SDK connector bypassing proxy with DNS requests. [COKCON-790](#)
- Fixed missing endpoint in the Veeva Vault connector. [COKCON-791](#)
- Fixed `java.nio.channels.UnresolvedAddressException` in Azure SDK connector [KOT2](#)
- Fixed handling of 404 response for the apps endpoints in VMware air watch. [KOT9](#)
- Fixed issue 405 Method not allowed in Veeva Vault. [KOT25](#)
- Fixed an `OutOfMemory` exception in Microsoft Azure Compute. [KOT27](#)
- Resolved security issues [CVE-2020-9488](#), [CVE-2018-10237](#), [CVE-2021-4104](#), [CVE-2018-1000844](#), [CVE-2018-1000850](#), [CVE-2021-43797](#) in various cloud connectors. [KOT13](#) [KOT14](#) [KOT15](#) [KOT16](#) [KOT17](#) [KOT18](#)

12.3 [RTM]

Released on [Feb 22, 2022](#)

Feature Highlights

New and Updated Cloud Connectors

In this version, many new cloud and SaaS connectors have been added.


- | | | |
|-------------------------|----------------------------|---|
| • 15Five | • EasySoftware EasyProject | • Microsoft 365 (renamed from Microsoft Office 365) |
| • Acquia Cloud Platform | • Envoy | • Nulab Backlog |
| • Acquire | • ExaVault | • okta |
| • Adobe Captivate Prime | • foreman | • Pipedrive |
| • Adobe Marketo | • Gainsight PX | • Quip |
| • Articulate Rise 360 | • Help Scout | • Salesforce SOQL |
| • Asana | • HubSpot | • SalesFusion Sugar CRM |
| • Atlassian Cloudadmin | • Jamf Pro | • Sophos Central |
| • Atlassian Opsgenie | • JetBrains Hub | • Tableau |
| • BambooHR | • Malwarebytes Nebula | • teamgantt |
| • Bloomfire | • Media Beacon | |

- Box
- Caspio
- ClickMeeting
- Cloud Academy
- Coupa
- Cvent
- Cyberark
- Drift
- Microsoft Azure Compute Inventory
- Microsoft Azure Consumption
- Microsoft Azure Log Analytics
- Microsoft Azure SQL
- Microsoft Graph Generic Query
- TeamViewer
- Tibco Scribe
- TravelPerk
- Typeform
- Twilio
- Vonage Contact Centers
- Workable
- Zeplin.io
- Zoho Books



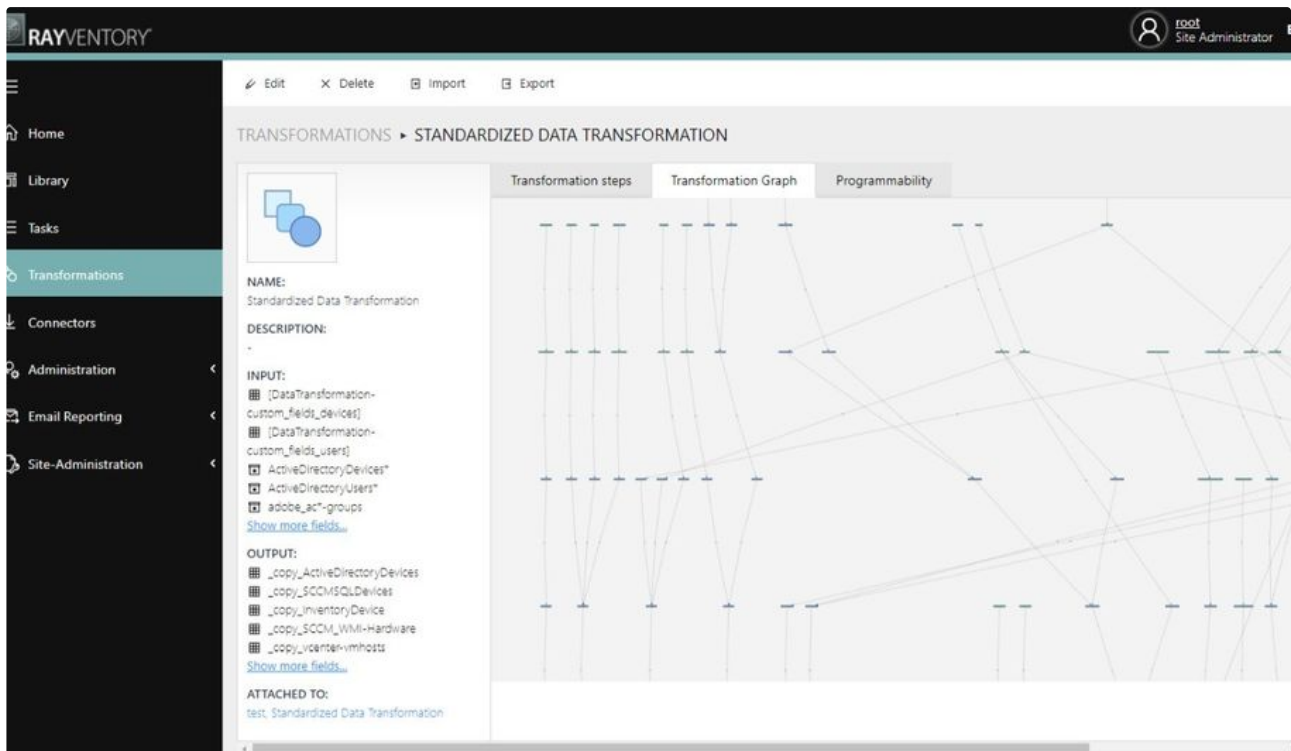
Integrated Transformation Capabilities

In this version, we added a new feature to handle data transformation. Whereas the previous releases of Data Hub had a way to achieve this task manually via JSON-based configuration, this version has a fully featured UI-driven editor where all data operations can be defined without coding. [RR-1144](#)



Sample ETL use case - transforming user information from different (usually only partially compatible) sources.

The new editor is available from the newly introduced **Transformation** tab. Each data transformation operates on a set of input tables and produces one or more output tables. In the basic editing view, the user defines atomic operations like mapping, joining, deduplicating, filtering, etc. The connections between them are determined automatically by chaining input and output of each step.



Graph viewer representing data flow and transformation steps of a complex process.

The built-in graph viewer helps to find out the data flow and to understand even the most complex transformation flows.

Custom functionality and advanced functions can be programmed as SQL macros and used directly by supporting action types.

Transformation supports the following basic operations:

- Mapping a table or multiple tables into a single table. Existing columns can be taken as-is, transformed, or discarded. New columns with auto-generated, fixed, or derived values can also be created.
- Filtering a table, using logical operators, conditions, and even SQL-snippets.
- Taking top X or bottom X rows from a table.
- Deduplicating a table based on some specific values and their combinations.
- Grouping rows in a table, aggregating columns, and counting elements.
- Joining two or more tables with custom join rules and wildcard support.
- Enriching a table by adding missing information from other tables.
- Creating a split table from aggregated values.

These types of operations can be combined together and chained, so that even the most complex data flow can be seen as a breakdown of smaller, simple operations.

- Default import comes with a preconfigured transformation layer which seamlessly joins various sources and creates data sources for integrated SAM, HAM, and IT Visibility reporting.

New Default Reports Structure

In this version, there is a new, better structured, and well organized default structure of standard reports and dashboards. Many of them have undergone significant changes, making sure that the data is presented in optimal way. Also their visuals have been updated. **RR-1920**

Once all new reports are imported, the following folder structure will be created:

- **Data collection**
All tasks and reports regarding raw, unprocessed data.
- **Data transformation**
A scope for tasks. Its main purpose is to prepare the data for the analysis and further processing. Reports and dashboards from this place can be used to control the quality and understand possible data gaps and other issues.
- **Data analysis**
Based on the transformed data, the purpose of this section is to show interactive, business critical data in an inspiring and appealing form. Various different use cases are covered here, from IT professionals and system administrators to managers and license owners.

i Your existing reports will not be migrated automatically. This change only affects newly imported reports.

Docker-Based Installation

As an alternative to the current MSI+IIS based installation, a Linux-based Docker image of the server component is available and can be pulled from the Docker Hub. **RR-2496**



More information about the installation of the Docker-based RayVentory Data Hub can be found in the following location:

<https://hub.docker.com/r/raynetgmbh/rayventory-datahub>

i Note: The Agent is still a Windows application available in MSI format. It is bundled in both Docker and MSI sources.

Catalog Software Recognition Now Supports Data from SCCM/ECM

The Catalog task has been extended with evidences from the System Center Configuration Manager. In addition to RayVentory, SCCM data is now enriched with key metrics such as end of life or release dates, business functionality, license classifications, and vulnerability information and can be viewed through dashboards. [RR-2471](#)

For this, the process of how the software normalization step works has been changed. In previous version, the Catalog connector performed direct queries on RayManageSofti (RMSi) / RayVentory database. In this release, it reuses the data from other, already executed tasks.

Software Information Now Has More Enriched Properties

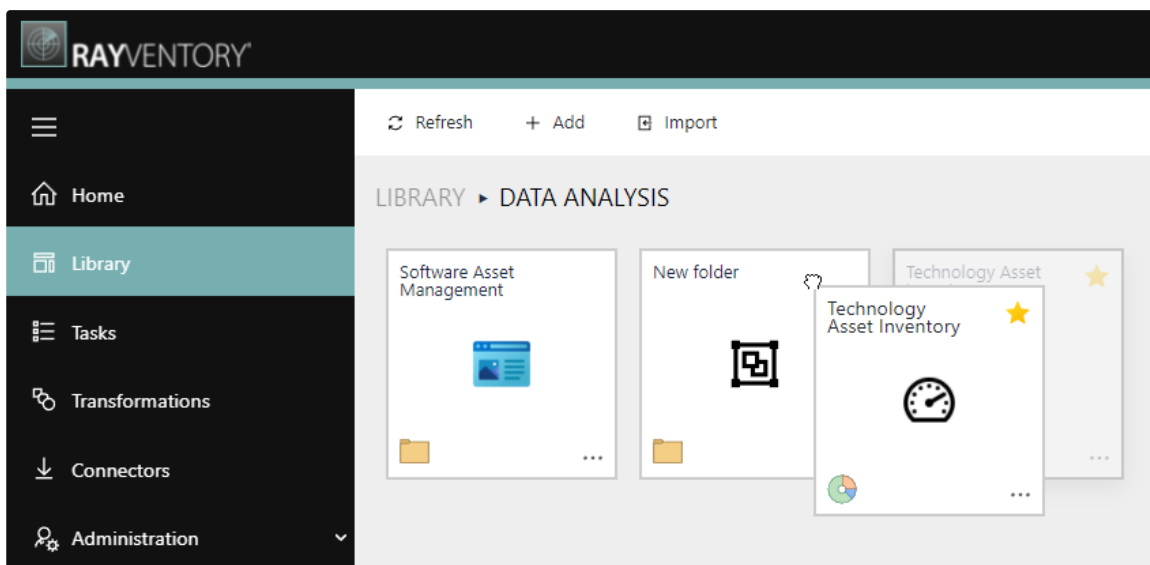
We extended the number of extra information pulled for each recognized software coming from RayVentory Catalog [RR-2462](#). The following information is now available:

- Product edition
- Extended support EOL date
- Software classification based on UNSPCS taxonomy (class / commodity)

Ignored fingerprints are now also written for troubleshooting and reporting purposes.

Reordering Tasks with Drag and Drop Functionality

Tasks can now be moved between folders by a simple usage of drag and drop technique. [RR-1131](#) [RR-161](#)



Dragging an element to reorder or move to a sub-container is now possible.

Drag-and-drop functionality can be used to change the position of each item or to move items to subfolders.

Getting Data from Catalog Is Now Much Faster and More Accurate

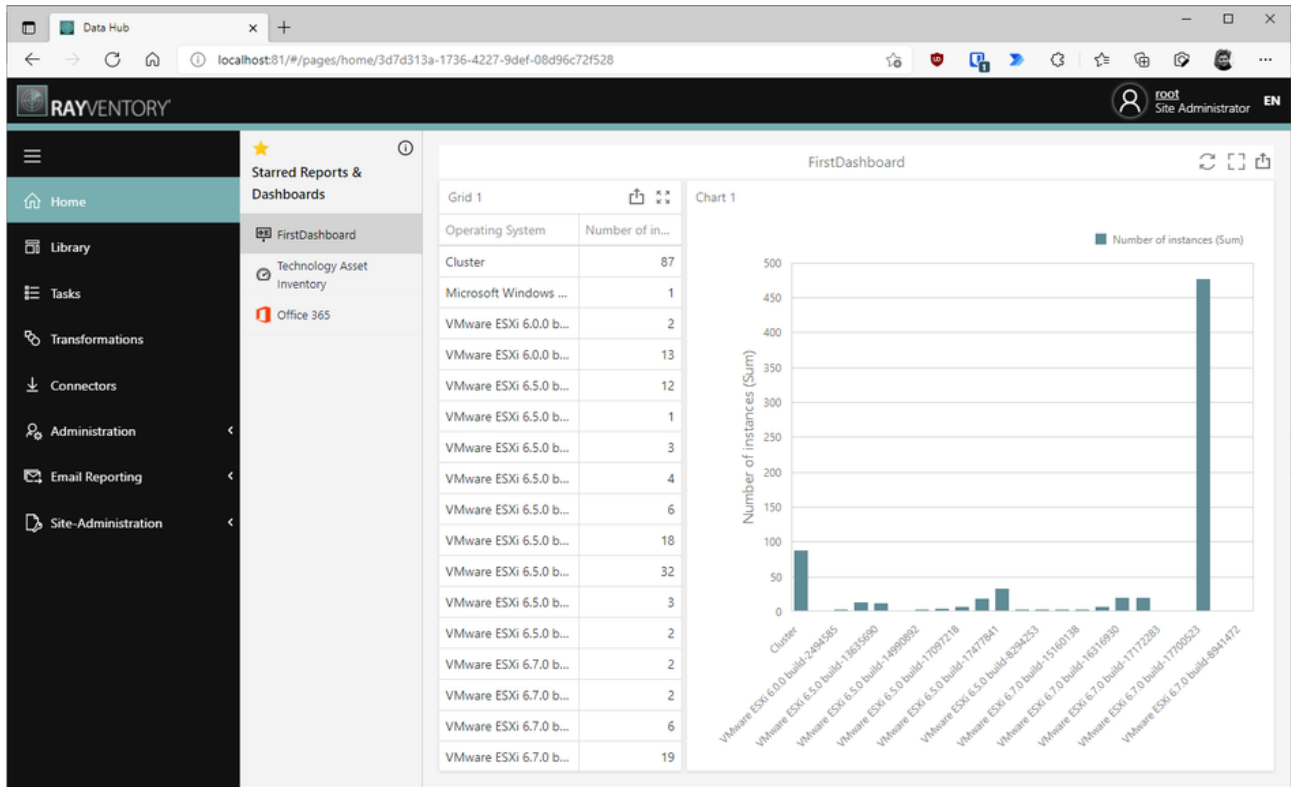
There are substantial changes in the catalog normalization algorithm, which - depending on the size of the data - bring huge performance improvements. The process runs much faster, shortening up lengthy hours-long operation into a matter of minutes. [RR-2471](#)

Additionally, it is now possible to add a maximum number of processed fingerprints for troubleshooting and quick runs. [RR-2381](#)

Thanks to many internal changes and optimizations, the standard recognition rate of installed software should also be higher in this new version. [CAT-1019](#) [CAT-1002](#) [RR-2463](#)

New Dashboard Screen

A new Home screen with the possibility to pin featured reports and dashboards and instantly access to them. [RR-2104](#)

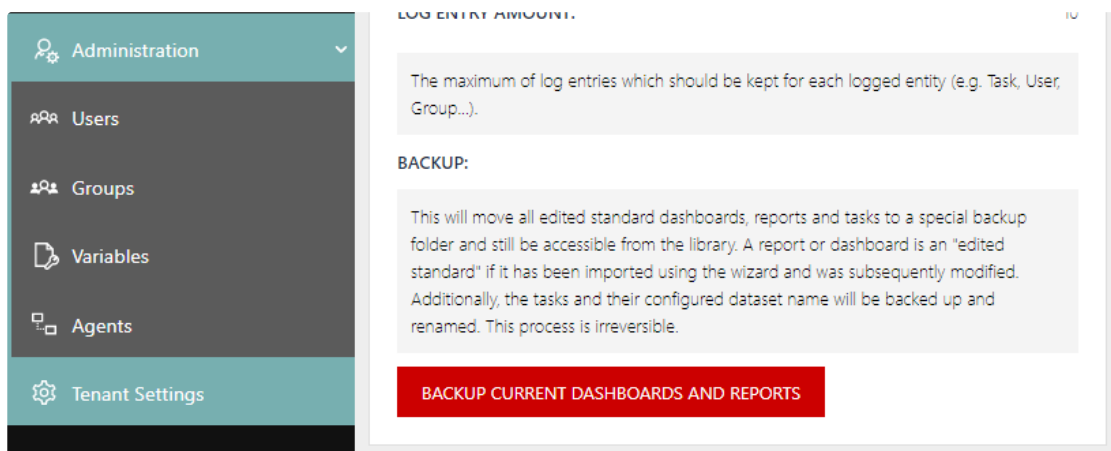


The dashboard with a list of favorite reports.

You can easily pin your favorite reports and dashboards and have them accessible from a single place directly in the Home screen. To start a dashboard or a report, simply add it to favorites list by pressing the star icon on its tile.

Easy Backup of Existing Reports, Dashboards, and Tasks

With a single click, the complete structure of current reports, dashboards, and their tasks can be moved to a separate folder. This feature can be used to quickly migrate from older versions of Data Hub, by using all new dashboards and being able to move forward step-by-step with current, already customized, and modified versions of previous reports and dashboards. [RR-2402](#)



Tenant-based setting to prepare a backup of existing reports.

Better Task Progress

The sidebar with information about each task received a substantial update, so that it is now more compact and provides only the information relevant for the current state (depending on whether a task is running, pending, or finished). [RR-1376](#)

Additionally, in case of a failure during the data collection, an excerpt of the log is now sent and shown in order to better understand the cause of issues. [RR-2411](#)

Task History Details

STATUS:
● Failed

TRIGGERED:
Nov 19, 2021, 3:31:57 PM by root

START DATE: Nov 19, 2021, 3:32:02 PM **END DATE:** Nov 19, 2021, 3:32:02 PM

ERROR:

```
RayVentory Catalog API key is missing.
```

AGENT LOG:

```
[19-11-2021 14:32:02] [DEBUG] : Starting RayVentory Catalog Collector...
```

An example of a failed task with extra log information and precise messages.

Russian Translation

The multi-language Data Hub Server UI is now also available in Russian (next to already being available in English and German). [RR-2840](#)
[RR-2725](#)

Performance Improvements

In this version, several areas were optimized for a better performance and stability.

- Server backend has been improved to handle large data and process it faster. [RR-1945](#)
- Improved performance of parallel import of data from agents. [RR-1843](#)

Security Improvements

i The following security improvements have already been released as hotfixes for affected installations.

- Fixed the security vulnerability [CVE-2021-44228](#). In this version, `log4j` is not used anymore and it was removed from the bundle. [RR-2696](#)
[RR-2704](#)
- Fixed an issue where debug logging could put out sensitive information in log files for SaaS connectors. [RR-2717](#)

Other Improvements

- Ability to create a new task directly from the Collectors page. [RR-1823](#)
- Extended configuration of proxy settings. [RR-2145](#)
- Added WMI-based capabilities to the SCCM collector. [RR-1616](#)
- Email notification now supports SMTP mail servers with NTLM authentication mechanism. [RR-2089](#)
- Changed SMTP port for TLS/SSL from 465 to 587 in email notification. [RR-2088](#)
- Improved resolving of variable placeholders containing special characters. [RR-2075](#)
- Enabled sorting of the "Last run" column of the Task-Table in Tasks-View. [RR-2103](#)
- Improved the detection of the installed `java.exe`. Specific registry entry search is now also done for *OpenJDK* and *Microsoft* JDK installations. [RR-2124](#)
- Increased the upper limit of lengths of variables to handle large strings. [RR-1944](#)
- New type of task – to the already existing Container and Data collection types, a new type Transformation has been added for data transforming purposes. [RR-2156](#)
- In the Import wizard it is possible to overwrite previously imported reports. [RR-1918](#)
- Added icons representing types (dashboard, report, or folder) in the Library screen. [RR-2605](#)
- It is now possible to clean the target existing tables without dropping them. [RR-2845](#)

Resolved Issues

- Fixed an issue with insufficient escaping of special characters in LDAP strings. [RR-2150](#)
- Fixed an issue where the data import could skip some entries depending on the maximum number of allowed tables. [RR-2172](#)
- Fixed an issue with a wrong default value of the *ActiveDirectoryGroups* task, which could lead to unnecessary extraction of nested groups. [RR-2153](#)
- Fixed an issue with duplicated group names in the Active Directory collector. [RR-2149](#)
- Fixed handling of decimal values with custom scale/precision in the ODBC Data Collector. [RR-2039](#)
- Fixed a problem with timeout of 30 seconds with long-running SQL queries. [RR-2063](#)
- Fixed an issue with an error caused by internal scheduler. [RR-1831](#)
- Fixed some log messages not being logged properly, regardless of the log level selection. [RR-2087](#)
- Fixed direct file import of CSV, XLS files, where some rows could be skipped. [RR-2095](#)
- Fixed an issue where SQL tasks might fail due to newline operators within string values. [RR-2132](#)
- Fixed an issue where task configurations could not be edited and saved if an input field was of the type array. [RR-2139](#)
- Fixed behavior of tooltips on IE11. [RR-1990](#)
- Fixed an issue where direct file import was not writing correct table names in case of multi-table import. [RR-1816](#)
- Fixed an issue with the setting "ignore duplicates" not working when performing a direct file import. [RR-1826](#)
- Fixed the message informing the user about the minimum database version in case of missed or incorrectly performed database migration. [RR-1828](#)
- Fixed an issue with import wizard, where referenced variables from the configuration tab were not shown. [RR-1844](#)
- Fixed importing of RPA files with sections. [RR-1854](#)
- Fixed an issue with direct file import, where ignoring of duplicates was not correctly respected. [RR-1859](#)
- Fixed incorrect parsing of attributes for custom AD queries. [RR-1898](#)
- Fixed an issue with wrong handling of integer values from auto-suggestion configuration fields in the task configuration. [RR-1911](#)
- Fixed an issue with improper conversion of string values to integer values, resulting in some connectors returning an invalid cast exception. [RR-1912](#)
- Improved logging, which should output less warnings in case of disabled SSL validation. [RR-1916](#)
- Fixed an issue with the Active Directory collector returning incomplete results for the "Users and group mapping" sub-task. [RR-1930](#)
- Fixed an issue with the length of some task configuration fields, which were not accepting more than a few hundred characters and rendered some longer strings truncated. [RR-1944](#)

- Fixed an issue with incorrect message in the Report if a required parameter was not set. [RR-2647](#)
- Fixed an issue with wrongly rendered line breaks in tooltips for cloud connectors. [RR-2687](#)
- Removed dependency on ASP.NET Core from Data Hub Agent. [RR-2396](#)
- Fixed a critical issue where only 10,000 rows could be returned by the agent when using target type SQL in the Task Agent Settings. [RR-2815](#)
- Fixed an issue with protected variables not being properly masked with a placeholder. [RR-2819](#)

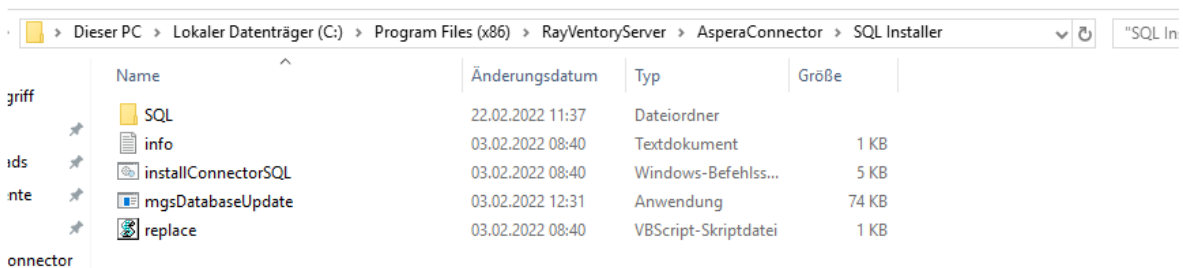
Breaking Changes and Known Issues

- RayVentry Data Hub 12.3 requires the following components to be installed:
 - Installation of the Aspera Connector [12.3.0.11566](#) is required. Previous versions of connectors are unable to deliver the required data.
 - JDK 11 or higher is now required for cloud connectors. Having previously supported JDK 8 will result in failing tasks.
- When upgrading from previous version of the Data Hub Server, on some machines the web application does not start and ends with an error `50x`. This can be fixed by repairing the MSI installation (right click the `.msi` file, and from the context menu select *Repair*). [RR-2826](#)
- Important changes for users migrating from version 12.2 (without Service Packs)
 - We have changed the way the report table name resolution works. The previously used separator (dot) is not used anymore in favor of a minus sign.
 - This means that the tables created by 12.2 extraction tasks followed the following naming convention `dataset.tablename`, whereas the new one puts out `dataset-tablename`. The migration script included in this build should be able to convert your existing reports to the new format, but the data will be only available once all tasks are restarted. You may also do the renaming on your own.

Installing the New Version of Aspera Connector

Several standard reports and tasks have a dependency on Aspera Connector. The connector must be installed on the server, that will be queried by Data Hub tasks.


The sources for the installation can be found after the installation of the RayVentry Server in the following folder (the path may be different if a custom installation was used): `C:\Program Files (x86)\RayVentryServer\AsperaConnector\SQL Installer` :



You can also access these files by unpacking the original MSI file, with the following command line: `msiexec /a <path-to-msi-file> TARGETDIR=<extraction-dir> /qb`.

To install Aspera Connector, execute the script `installConnectorSQL.cmd` as Administrator:

```
C:\Windows\system32\cmd.exe
SQL Server Settings for ManageSoft Database
-----
SQL Server      (localhost):
Instance        (SQLExpress):
Database Name   (ManageSoft):
Database Version (12.3):
```

 The default values (displayed in round brackets) are standard, and do not have to be re-entered if a change is not required. You can simply confirm them by pressing `Enter` to take them over.

Once the script has finished, the installation is complete.