Retrospective Version 6.1

User Manual

Table of Contents

1 INTRODUCTION TO THE MANUAL

This manual describes the features and capabilities of Retrospective. It contains detailed descriptions of configuration tasks as well as troubleshooting procedures and end to end examples.

This document targets various IT specialists ranging from business analysts, software developers, testers to system administrators. Regardless of the function, Retrospective brings an immediate and constant benefit to everyone's daily work.

1.1 Conventions and Symbols

Element	Description
Italic	Used for denoting graphical user interface items, e.g. ***Preferences*** window
`Courier`	Used for file names and file system paths, e.g. `D:\jboss\server\default\log`
Boldface	Used for commands and exemplary values, e.g. Profile_20120314_131215
[square brackets]	Used for keyboard shortcuts and application buttons, e.g. [Ctrl] + [X]
Note This is an example text for additional information	Used for supplementary information on given topic.

2 INTRODUCTION TO RETROSPECTIVE

Retrospective is a desktop application for convenient and effective searching in local and distributed log files as well as log data from container platforms such as Docker and Kubernetes. Such data sources are accessed on local and/or remote computers and the log data of different format is extracted, combined and displayed in one place (the result table). Advanced features such as profile definition, data source monitoring, sorting/filtering result entries, highlighting, bookmarking, exporting etc. enable fast and robust log data exploration. Retrospective assists you in early error detection by monitoring application servers' logs, searching for exceptions, etc. The highly optimized search engine together with the aforementioned features ensures that much precious time is saved. Being able to access all log data in one place makes log data processing incredibly efficient, thus ensuring quick and effective reactions to problems, and therefore gain customers' trust and loyalty.

Unlike other expensive and complex solutions with monolithic event databases, Retrospective is easy to set up and offers immediate access to real time data in newly created log files and containers. There is no caching, indexing, or unnecessary data transfer involved, thus it does not violate any security compliances and offers rapid, just-in-time search results.

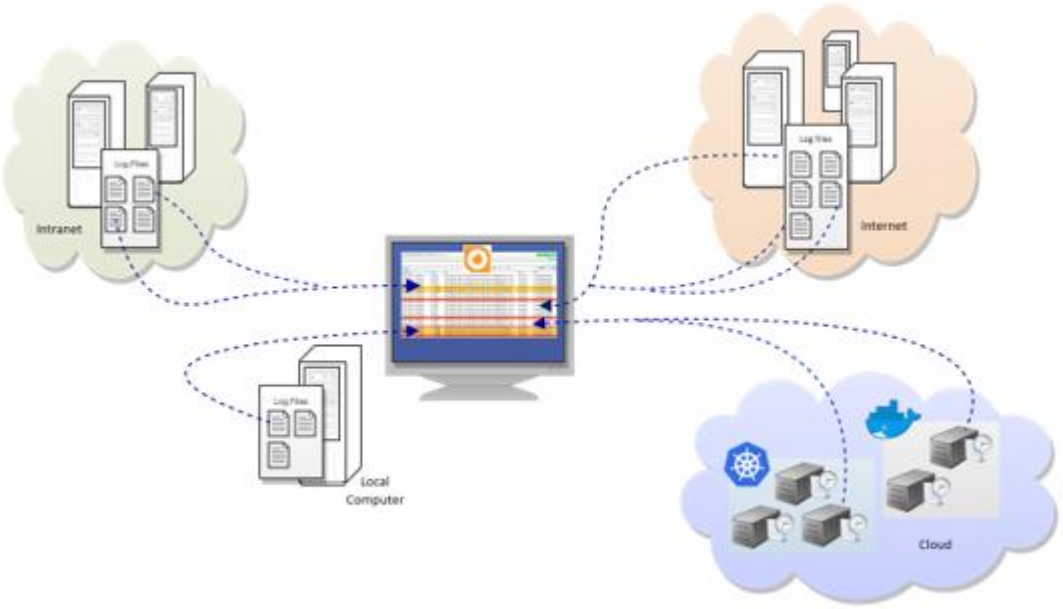

2.1 Getting Started

2.1.1 Upgrade from previous version

When a previous version of Retrospective was already installed on your computer and you launch a new program release for the first time, you will see dialog that lets you import the configuration from a previous release. In case the previous version was 5.7.0, result snapshots are also imported and after that, the last Retrospective session is completely restored.

For further details, please consult chapter 3.5.1 Retrospective Version Upgrade

2.1.2 First time use

When you launch Retrospective the first time, you'll be presented with the window shown below.

The search tab, located on the right, provides the interface to the Retrospective core functionality. It is needed to start searching data in one or several log files available locally or on remote computers accessible through SSH and/or to start searching log data of Docker or Kubernetes containers. Simply press the "Select Data Source(s)" button to enable you to choose the desired files/containers, then enter the search text and **Start Searching**. If you leave the search text field empty, Retrospective reads all log data from the selected data sources. If you enter a search text, the result will only be composed of log entries that contain this text. In case you want to define more complex search criteria, press the **"Advanced"** link that appears below the start button. In the advanced mode, you can combine different search criteria for matching the time range and the content of the desired log entries.

If you're interested in data written to the selected log files - or produced by the selected containers - in the future, Retrospective can monitor them and present matching data as soon as it appears. Simply switch to monitor mode by pressing the down arrow on the "Start Search" button, define your search criteria and start **monitoring**.

The **Welcome** banner, located on the left, leads you to the three main constituents of the program. These enable you to use a wide range of the most important Retrospective functionalities in a professional way right from the beginning.

Choice	Description
Start searching or monitoring	Follow this procedure if you want to search log data in one or several files chosen on the fly. You'll be presented with a dialog where you can select log files located on a local drive or on a remote server.
Manage connections to remote computers	This feature lets you define persistent SSH connection definitions for a set of remote servers within the host manager. Such servers (hosts) are then available to be used for ad-hoc search and monitoring but can also be used in different profiles.
Manage profiles and data sources	This option lets you define persistent profiles within the profile manager. Profiles are used to group files from one or several local drives, files from remote servers or containers, enabling you to perform search and monitoring on a complex set of data sources and have the retrieved result visualized in a larger context.

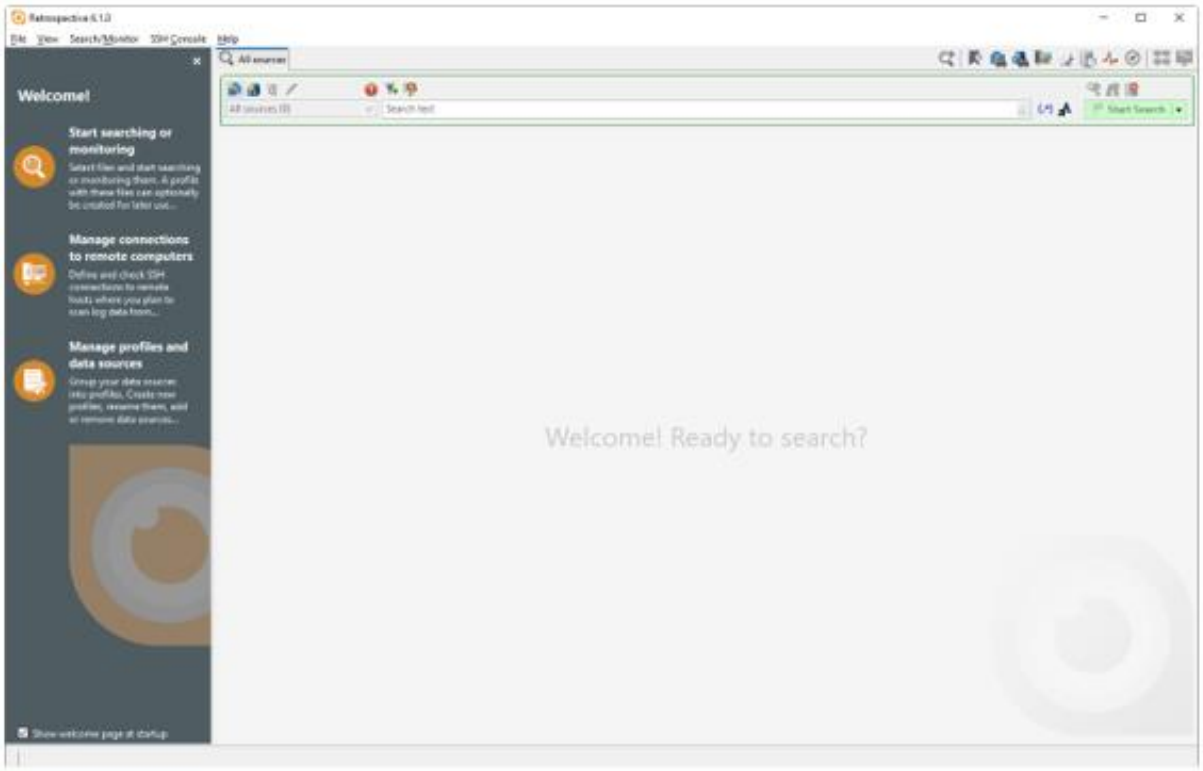

2.2 Basic Use Cases

2.3

This chapter describes a few common Retrospective use cases. It provides a rough overview on how the program can assist you in your daily work that's related to log data analysis. Thanks to the high flexibility of the program, the different steps can be rearranged and combined in comprehensive ways at any moment.

2.3.1 Ad-hoc Monitoring with Highlighting

Software developers typically run their programs from within their favorite integrated development environment (IDE). During development, programs usually write log data to a file on the local file system and thereby provide useful information about the processed data and the internal processing. Retrospective lets you monitor such log files and presents new data in real-time.

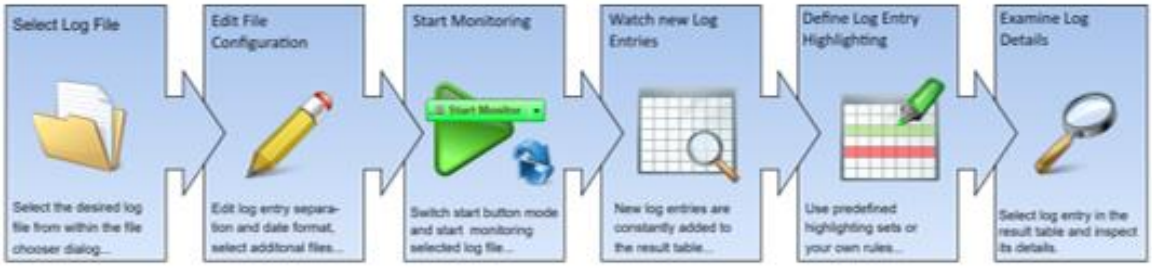

1. Open a new search tab and press the folder button that appears next to the profile selection drop-down list. If you have not defined any profile yet, press the "Open Data Source(s)" button instead. From within the file chooser dialog, select the file you'd like to monitor and press the OK button. Attention should be paid to the checkbox labeled "Autofind

configuration" that appears left of the OK button. If it's selected, Retrospective automatically analyzes the selected log file in order to determine its encoding, the log entry separation (delimiter) and the date/time format.

2. The chosen file now appears in the profile selection drop-down list. Pressing the edit button right of it opens a dialog where you can further configure the choice. You can for example manually define or correct the file encoding, the log entry separation (delimiter), date/time format and custom columns. You may also add other files or entire folders to your choice.
3. Once the data source(s) have been correctly defined, switch to monitor mode by activating the arrow on the start button (see 5.1.2 Defining Search Criteria). Now press the "Start Monitor" button.
4. Every time a new log entry is written to one of the chosen files, it immediately appears in the result table inside Retrospective.
5. In order to focus your attention on critical issues, individual log entry rows have a background color representing their log level (i.e. red for ERROR). You can change this default behavior at any time by selecting a predefined highlighting set or by defining your own rules, colors and fonts. Thus, you get immediate stylish feedback when data of interest appears.
6. If you want to see details of a multi-line log entry, simply select the related row in the result table.

2.3.2 Contextual Search

Software testers generally use program log data to describe the cause of a program failure in detail and thereby provide precious information to software developers that have to fix the problem. Testers often first try to find an error log entry that correlates with the particular program failure. From there, they fetch all log data that was generated during a specified time span. The fetched log data would then be attached to the relevant issue in the company bug tracking system.

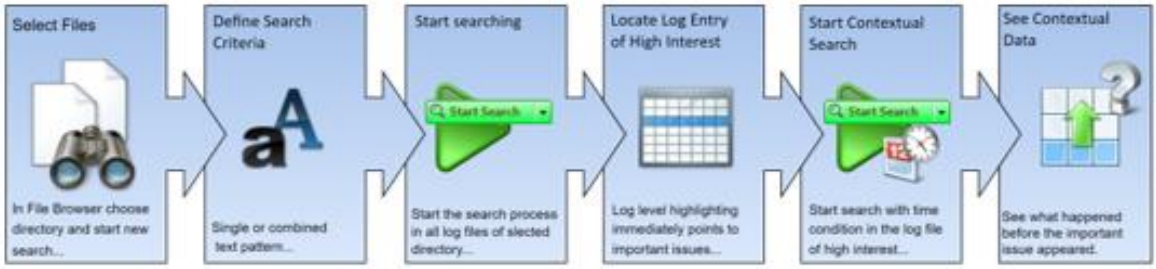

1. Open the File Browser tab and navigate to the relevant directory to which the program under test writes its log data. Choose the log file you're interested in and select the menu item "Start Searching" in the context menu that appears when you press the right mouse button. You may also choose several log files by selecting the check box that appears in front of each files' names. All files from the selected directory can be chosen by switching to "Filter Matching Selection" in the drop-down box that appears on top of the files section. The choice of available files can be limited by entering a text pattern in the file filter field (i.e. *.log). From there, you need to press the button, located top right, to start searching (create a new search tab).
2. On the search tab, enter a text pattern to be used for finding the log data you're interested in. Should you want to combine several search criteria, press the "Advanced" link that appears below the "Start Search" button. This enables you to add other search criteria (button) and have them joined with an AND/OR operator (see 5.1.2 Defining Search Criteria).
3. Once all your search criteria are defined, press the "Start Search" button.

4. Retrospective displays individual log entries that match your search criteria immediately upon discovery. Within the result table, log entries of level (severity) ERROR and WARNING are displayed with a red and orange background color to attract the user's attention to important issues.
5. If a log entry especially attracts your interest, move the mouse pointer over it and press the right cursor. From the pop-up context menu, you may select the menu item *Start new search in selected path > Start one minute beforehand*. This creates a new search tab where you start a new contextual search by pressing its "Start Search" button.
6. Retrospective now presents log data from the same file written within the minute preceding the log entry of special interest.

2.3.3 Profile Based File Search with Filtering/Sorting

Operations support team members usually scan well-defined sets of log files on different servers in order to proactively find problems reported by the many programs of which an enterprise application is built. The fetched log data can be filtered and sorted locally, and non-interesting data removed from the result set. The final result set with condensed and useful information about production issues can be saved to a file and shared with colleagues to be used as input for further investigations.

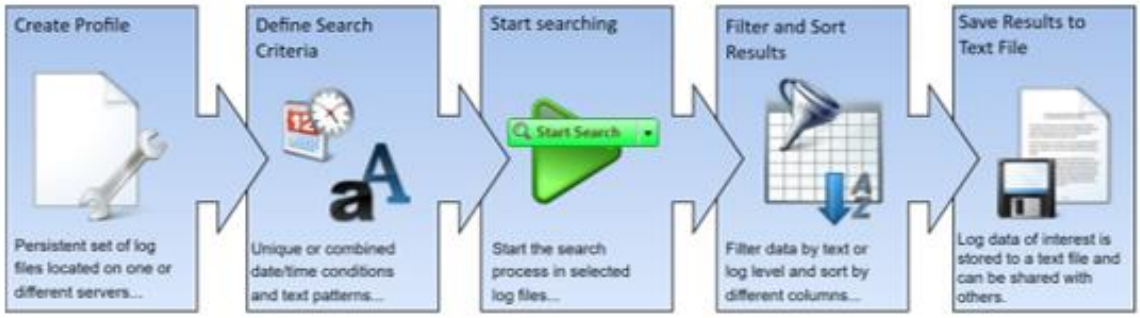

1. Create a profile within the Profile Manager and add as many file data sources as appropriate (see 3 Configuration and Setup). Data sources can be individual files on the local computer or on different remote servers. Data sources may also be entire directories or a bunch of files with a name that matches a user defined text pattern.
2. On the search tab, press the "Advanced" link that appears below the start button. Switch the search criteria type "Text" to "Date/Time" and enter the desired value(s). Now press the "Advanced" link that appears below the start button. This enables you to add other search criteria (button) and have them joined with an AND/OR operator. You could, for example, add a text pattern to be used for finding the log data of your interest written during a specified time period (see 5.1.2 Defining Search Criteria).
3. Once all your search criteria are defined, press the "Start Search" button.
4. Retrospective displays individual log entries that match your search criteria immediately upon discovery. Within the result table, make use of the many local sort and filter options, they help with re-arranging the result set and finding data of immediate interest. Local filters, for example, limit the number of result entries and allow you to avoid running the search once again with redefined search criteria (see 4.2.1 Result Entries Toolbar). The number of result entries can also be reduced by repeatedly selecting and deleting bulks of irrelevant log entries from the result table.
5. Press the "save result data to text file" button located on top of the result table. The arranged set of significant log data will now be saved to a text file for further analysis.

2.3.4 Profile Based Container Search with Formatting

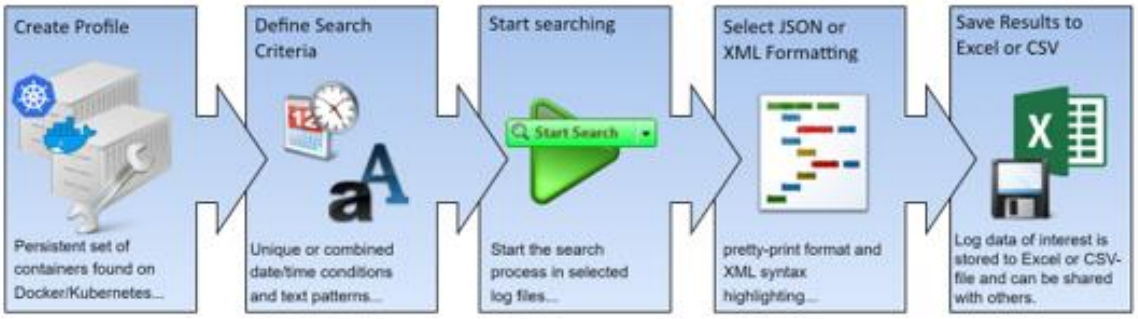

1. Create a profile within the Profile Manager and add as many Docker or Kubernetes container data sources as appropriate (see 3 Configuration and Setup). Data sources may also be entire directories (Images, Pods etc.) or a bunch of containers with a name that matches a user defined text pattern.
2. On the search tab, press the "Advanced" link that appears below the start button. Switch the search criteria type "Text" to "Date/Time" and enter the desired value(s). Now press the "Advanced" link that appears below the start button. This enables you to add other search criteria (button) and have them joined with an AND/OR operator. You could, for example, add a text pattern to be used for finding the log data of your interest written during a specified time period (see 5.1.2 Defining Search Criteria).
3. Once all your search criteria are defined, press the "Start Search" button.
4. Retrospective displays individual log entries that match your search criteria immediately upon discovery. Select a row to have its content displayed in a detail view that appears underneath the result table. Press the "Select formatting" button on top of the result table and choose JSON or XML depending on the expected log entry content. You'll notice that such structures are nicely formatted and that XML syntax is highlighted.
5. Press the "save to CSV file" button located on top of the result table. The log data will now be saved to an Excel file, formatting and syntax highlighting is carried over if desired.

2.3.5 Ad-Hoc Container Monitoring with Custom Columns

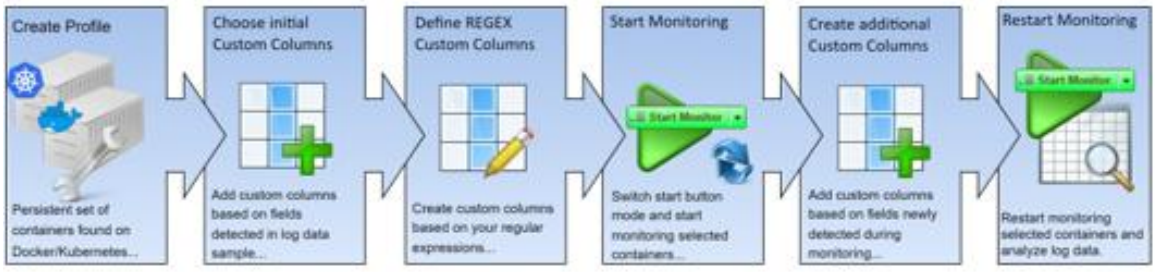

1. Create a profile within the Profile Manager and add as many Docker or Kubernetes container data sources as appropriate (see 3 Configuration and Setup). Data sources may also be entire directories (Images, Pods etc.) or a bunch of containers with a name that matches a user defined text pattern.
2. Press the "Add field to new column" button located in the bottom right section of the Profile Manager. This displays the "Add Field" dialog that presents different groups of items to be selected. The group "Detected in Data" presents all ready-to-use name/value pattern fields detected in data samples retrieved from the selected containers (see 3.4.6 Define Columns).

3. Optionally define additional custom columns based on expected fields that shall be extracted by freely editable regular expressions. The "Add Field" can assist you by proposing well known patterns such as name/value pairs, date/time, currency etc.
4. Once the data source(s) have been correctly defined, switch to monitor mode by activating the arrow on the start button (see 5.1.2 Defining Search Criteria). Now press the "Start Monitor" button.
5. Retrospective extracts the user defined fields and displays them in the appropriate custom columns. Optionally it detects new name/value pattern fields as monitoring keeps going on. Switch back to the Profile Manager and create additional custom columns if desired.
6. Press the "Start Monitor" button again and observe how Retrospective extracts the user defined fields and displays them in the appropriate custom columns. Use sorting and filtering on standard and custom columns and analyze your log data.

2.4 GUI Overview

Retrospective features a tabbed interface for smart and convenient discovery and analysis of log data from files and containers together with easy data sources management.

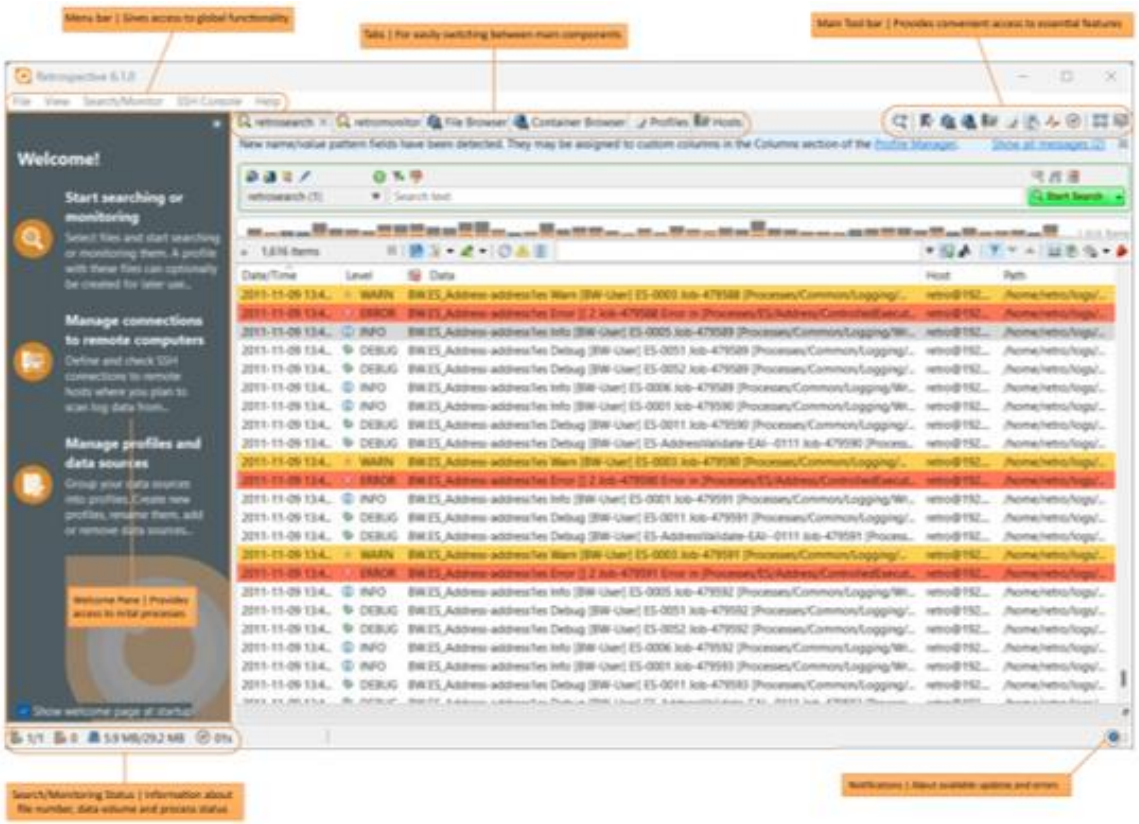

The content of the screen depends on the views currently open. User interface elements which always remain visible are...

- The menu that provides access to the preferences dialog, the different main views and most other features of Retrospective.
- The main toolbar that provides fast access to essential Retrospective features.

2.4.1 Menu Bar

The menu bar is composed of sub menus and menu items as described below.

Submenu/ Menu Item	Description
File	
Preferences	Displays the ***Preferences*** window containing application properties grouped in different topics (see 4 Preferences)
Export Hosts and Profiles	Exports the defined hosts and profiles to a local file for later use or to be shared with someone else. Prior to exporting the configuration to the file, a dialog is shown where you have to choose whether the passwords for accessing remote computers over SSH should be kept or if they should be cleared (purged). 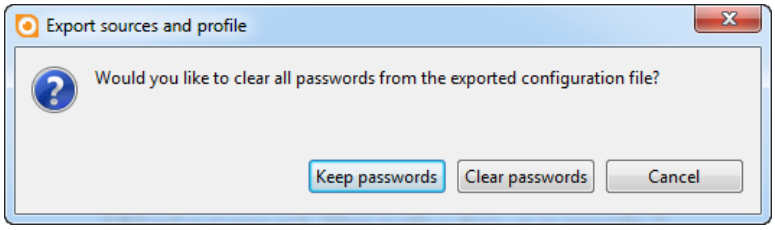 Choose "Clear passwords" if the file with the exported configuration should be shared with other people and if you want to make sure they can access individual remote computers only after they themselves have entered the required passwords.
Import Hosts and Profiles	Imports hosts and profiles configuration from a file that was previously exported either by you or by somebody else. Please be aware that this will completely overwrite your existing hosts and profiles configuration. Depending on how the file was exported, passwords for accessing remote computers over SSH may also not be available and will have to be re-entered in the host manager.
Export bookmarks	Exports bookmarked search criteria sets. This feature is of use only if the bookmarks are re-imported to a program session that has the exact same set of hosts and profiles configured.
Import bookmarks	Imports search criteria sets. This feature is of use only if the bookmarks have been exported from a program session that had the exact same set of hosts and profiles configured.
Import Configuration from previous Version	Opens the ***Configuration importer*** dialog aimed for importing configuration data from a previous compatible Retrospective release that was installed on your computer (see 3.5 Import from Previous Version).

Submenu/ Menu Item	Description
 Save View	Saves the current state of the application including opened tabs along with search criteria (see 10.6 Save/Reload/Manage Application State).
 Load View	Loads previously saved and named application states (see 10.6 Save/Reload/Manage Application State).
 Manage Views	Enables renaming or deletion of previously saved application states (see 10.6 Save/Reload/Manage Application State).
Exit	Closes the program.
View	
 Start new Search	Creates a new **Search** view ready to be configured for searching or monitoring log files (see 5 Searching and Monitoring).
 Bookmarks	Opens the ***Bookmarks*** view containing bookmarked search criteria sets.
 File Browser	Opens the **File Browser** that provides easy access to files and directories and lets you start new search or monitoring actions based on your selection (see 5.5 File Browser).
 Container Browser	Opens the **Container Browser** that provides easy access to containers located in Docker or Kubernetes subsystems and lets you start new search or monitoring actions based on your selection (see 5.6 Container Browser).
 Host Manager	Opens the ***Hosts*** view that lets you define SSH connection parameters to remote computer (see 3.3 Host Manager).
 Profile Manager	Opens the ***Profiles*** view where you define named sets of data sources located locally and on remote computers (see 3.4 Profile Manager).
 Result Snapshots	Opens the **Result Snapshots** view that lists all result data snapshots the user wants to keep for later use (see 5.9 Result Snapshots).
 Status	Opens the ***Status*** view displaying applications status information (see 5.12 Status View).
 History	Opens the ***History*** view containing recently executed search or monitoring actions (see 5.11 History View).
 Explode all Tabs	Explodes all tabs by creating an individual window for each of them (see 10.6 Save/Reload/Manage Application State).
 Implode all Windows	Closes additional windows and adds all their tabs to the window where this button was activated (see 10.6 Save/Reload/Manage Application State).
Search/Monitor	

Submenu/ Menu Item	Description
Choose new Files	Opens a dialog that lets you select files from a local drive or a remote computer. Subsequent search and monitoring actions will process (analyze the content) of all selected files.
Choose new Containers	Opens a dialog that lets you select containers from Docker or Kubernetes subsystems. Subsequent search and monitoring actions will process (analyze the streamed log data) of all selected containers.
Compose Data Sources	Opens a dialog that lets you compose a new set of data sources (optionally stored to a new permanent profile) from the ones that are defined in the existing profiles.
Configure current Data Sources	This lets you change the configuration of the current selected data sources. You may change their encoding, date/time pattern etc. but also add other data sources from one or many other computers, or remove existing ones.
Advanced Mode	Toggles between advanced and simple search definition mode.
Clone Tab	Creates a copy of the current search tab with the exact same search criteria.
Bookmark Tab	Stores the definition of search criteria as a user named bookmark for later user (see 5.10 Bookmarks View).
SSH Console	
New…	Opens a SSH Console for the selected host or all hosts within a folder. Please note that hosts must have been defined within the Host Manager prior to being available for opening a new SSH Console that connects to them.
Help	
User Manuel	Opens this user manual in the associated viewer.
About Retrospective	Opens information page about Retrospective.
Keyboard Shortcuts	Opens the ***Retrospective Keyboard Shortcuts*** window.
Release notes	Displays a list of changes implemented in Retrospective.
Program activation	Opens the ***Program Activation*** where you can enter the activation key and permanently activate Retrospective (see 3.2 Program Registration).
Update Software	Checks if there is a new software version available.
Welcome Page	Opens the ***Welcome*** banner that is located on the left.

2.4.2 Main Toolbar

The toolbar enables access to certain features such as opening new search tab, changing application appearance and accessing data source configuration, history and bookmark views.

The table below lists all available toolbar icons.

Icon	Description
 Start new Search	Creates a new **Search** view ready to be configured for searching or monitoring log files (see 5 Searching and Monitoring).
 Bookmarks	Opens the ***Bookmarks*** view containing bookmarked search criteria sets.
 File Browser	Opens the **File Browser** that provides easy access to files and directories and lets you start new search or monitoring actions based on your selection (see 5.5 File Browser)
 Container Browser	Opens the **Container Browser** that provides easy access to containers located in Docker or Kubernetes subsystems and lets you start new search or monitoring actions based on your selection (see 5.6 Container Browser).
 Host Manager	Opens the ***Hosts*** view that lets you define SSH connection parameters to remote computers (see 3.3 Host Manager).
 Profile Manager	Opens the ***Profiles*** view where you define named sets of data sources located locally and on remote computers (see 3.4 Profile Manager).
 Result Snapshots	Opens the **Result Snapshots** view that lists all result data snapshots the user wants to keep for later use (see 5.9 Result Snapshots).
 Status	Opens the ***Status*** view displaying applications status information (see 5.12 Status View).
 History	Opens the ***History*** view containing recently executed search or monitoring actions (see 5.11 History View).
 Explode all Tabs	Explodes all tabs by creating an individual window for each of them (see 10.6 Save/Reload/Manage Application State).
 Implode all Windows	Closes additional windows and adds all their tabs to the window where this button was activated (see 10.6 Save/Reload/Manage Application State).

2.5 Hosts

Hosts are remote servers accessed through SSH.

SSH Host definitions are required for searching and monitoring files stored on servers accessible via intranet or internet. They are used in cases when related data sources are defined on the fly (ad-hoc) or bundled within a profile. Hosts can be grouped into folders.

The same SSH Host definitions are also used for establishing a connection when you're working with the inbuilt SSH terminal.

2.6 Data Sources

Data sources refer to files that contain log data or containers that produce log data.

- **File** data sources are located on the local file system or on remote SSH hosts (servers). A file data source can point to a standard file, an archive file or it may use a naming pattern that matches multiple files of the same or different types each.
- **Container** data sources live inside the Docker platform or in the Kubernetes cluster – the two subsystems of container technologies supported by Retrospective. A container data source can point to an individual container or it may use a naming pattern that matches multiple containers of the same or different types each. Retrospective directly accesses log data from containers that is streamed out to stdout.

Data source definitions provide Retrospective with the following information.

- The files/containers that need to be included in the search or monitoring process.
- The file/data encoding (e.g. UTF-8 or latin1), which has to be known in order to properly read the file content or the log stream.
- How to cut the result data into individual entries (i.e. new line character or a date/time pattern) in order to present them in separate result table rows each.
- The Date pattern to be used for converting the log entry timestamp into a Date/Time instance that can be used for proper sorting and filtering.

All these aspects have to be dealt with when defining a data source. The user can either allow Retrospective to automatically discover the relevant information by using the Autofind procedure (see 3.4.3.1 Autofind) or he may provide the information manually using a wizard dialog.

2.7 Profiles

Profiles enable convenient grouping of data sources, allowing to precisely define which log files should be processed. They also contain information on how the data is to be presented in the result table.

Profiles are user defined persistent sets of data sources of the following types:

- Log file definitions from the local file system and/or from one or many remote servers.
- Container definitions from Docker or Kubernetes subsystems.

Search and monitoring actions performed on a specific profile will process (find matching entries) all contained files/containers and present data in the result table within standard and custom columns depending on your configuration. You can work with Retrospective with data sources that are created on the fly (ad-hoc) and not bundled in a profile. Sooner or later however you'll get to a point where you have to re-execute a search on the exact same set of data sources as in the past. That's when you realize that having a profile would be a good solution.

Retrospective features an easy to use Profile Manager which enables creating user profiles and adding data sources as well as defining result table columns quickly and conveniently. Log files stored locally can be dragged and dropped onto the Profile Manager tab, which depending on where the file will be dropped, will result in creating a new profile or adding a data source definition to a given profile.

2.8 Custom Columns

Log data found during the search and monitoring process is displayed in the result table within a predefined set of columns (Date/Time, Level, Data, Host and Path). The Data column contains the full, non-structured content of individual result entries. In order to have the most attracting cutouts of your log data displayed in individual new columns, you should define **Custom Columns**. Besides giving the column a name, you need to tell Retrospective, how to collect the data that shall appear in its table cells, you do this by defining **Fields**. In most cases there exist

no strict rules on how to produce log data, hence the semantically same fields can appear in different places and with distinct formats all over your log file(s). Therefore, Retrospective allows that more than one Field can be assigned to the same Custom Column. Retrospective currently supports the following types of Fields:

- Name/Value Pattern Fields
- Custom Regex Fields
- Split Fields (i.e. Character Separated Fields)
- Fixed Position Fields

Some of these field types will be described in the following sections.

Custom Columns are best defined for individual profiles within the Profile Manager. If you work with log files selected on the fly, you can define Custom Columns within the "Configure Data Source(s)" dialog for the current search tab. Keep in mind however, that in the latter case, these definitions are gone once your search tab is closed.

2.8.1 Name/Value Pattern Fields and Regex Fields

A common practice in programming is to log contextual information using name/value patterns. Unfortunately, too often many different patterns are used for logging the same data. Thus, it is hard to automatically extract the most meaningful values and present them in a custom column. Fortunately, Retrospective is aware of the programmer's habits and knows about commonly used name/value patterns they use. Thus, it assists you in detecting them during **Autofind** and even during the **search/monitoring** process and proposes them as Name/Value Pattern Fields to be assigned to Custom Columns if desired.

In fact, Name/Value Pattern Field is a special subtype of more general Regex Field. Regex Field is simply a field whose value is extracted by matching a defined regex expression. Name/Value Pattern is also defined by a regex expression but only a fragment of the matched expression – the one that represents the value – is extracted as the Field value. Retrospective provides some predefined Name/Value Pattern Fields as well as Regex Fields (e.g. amount with currency, e-mail address, IP address, date/time). As mentioned before, Name/Value Pattern Fields can be detected automatically by Retrospective, while Regex Fields must be defined manually. Thankfully the predefined Regex Fields makes this task easier in typical scenarios. For example, if you want to have an IP address in a given Custom Column, simply select a IPv4 Address predefined Regex Field and you are good to go.

In the figure below you can see, how different Name/Value Pattern Fields can be used to extract the same semantic information – some person ID number.

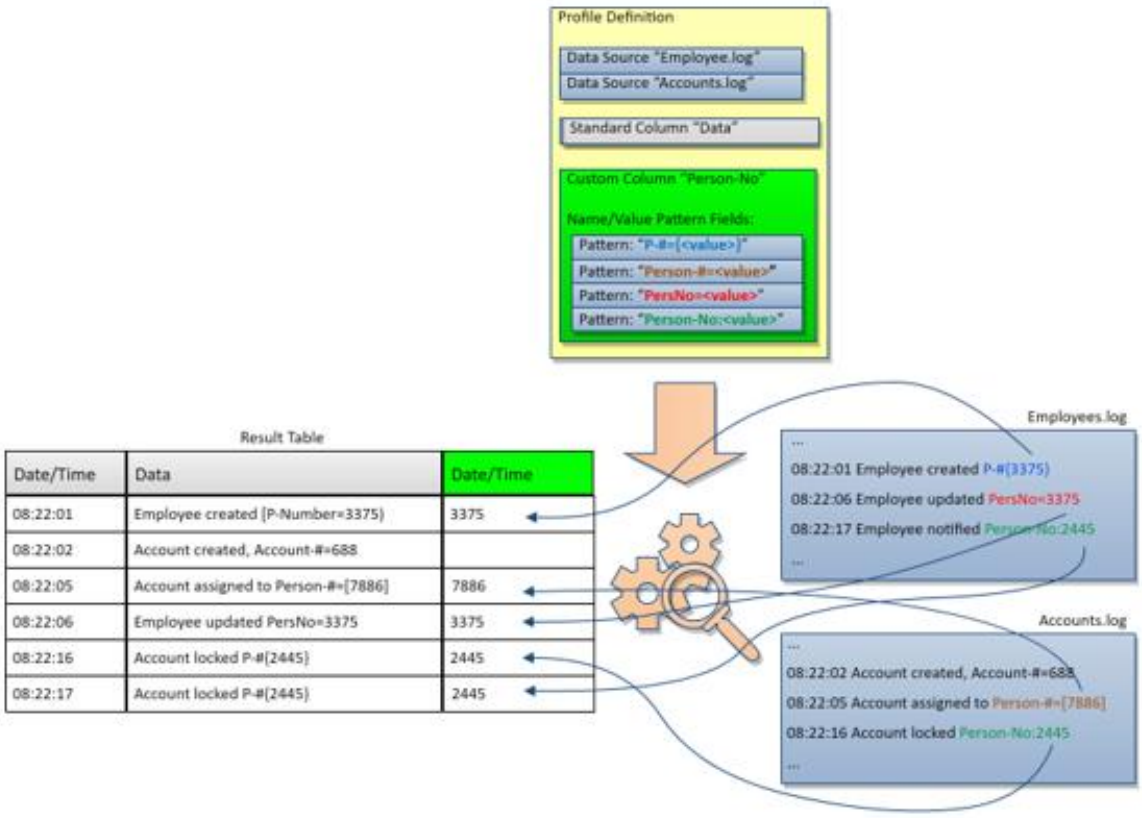

2.8.2 Character Separated Fields

Sometimes you'll face situations where your log files contain tabular data (known as CSV file format). When individual log entries of such files are split by a certain character (i.e. a comma), the same number of fields is obtained. All fields at a given position but from different log entries have the same semantics. The field at position one may contain the customer type, the field at position two its ID, the field at position three an action code, and so on. It's up to you to find out how such log files are structured. Once you're aware of it, simply define your Custom Columns and assign a desired Character Separated Field to each of them by specifying the split character and the field position. The figure below shows how different character separated fields can be used to extract some important information from given log entry. As can be noticed, it would be hard to extract this information with the use of Name/Value Pattern Field or Regex Field.

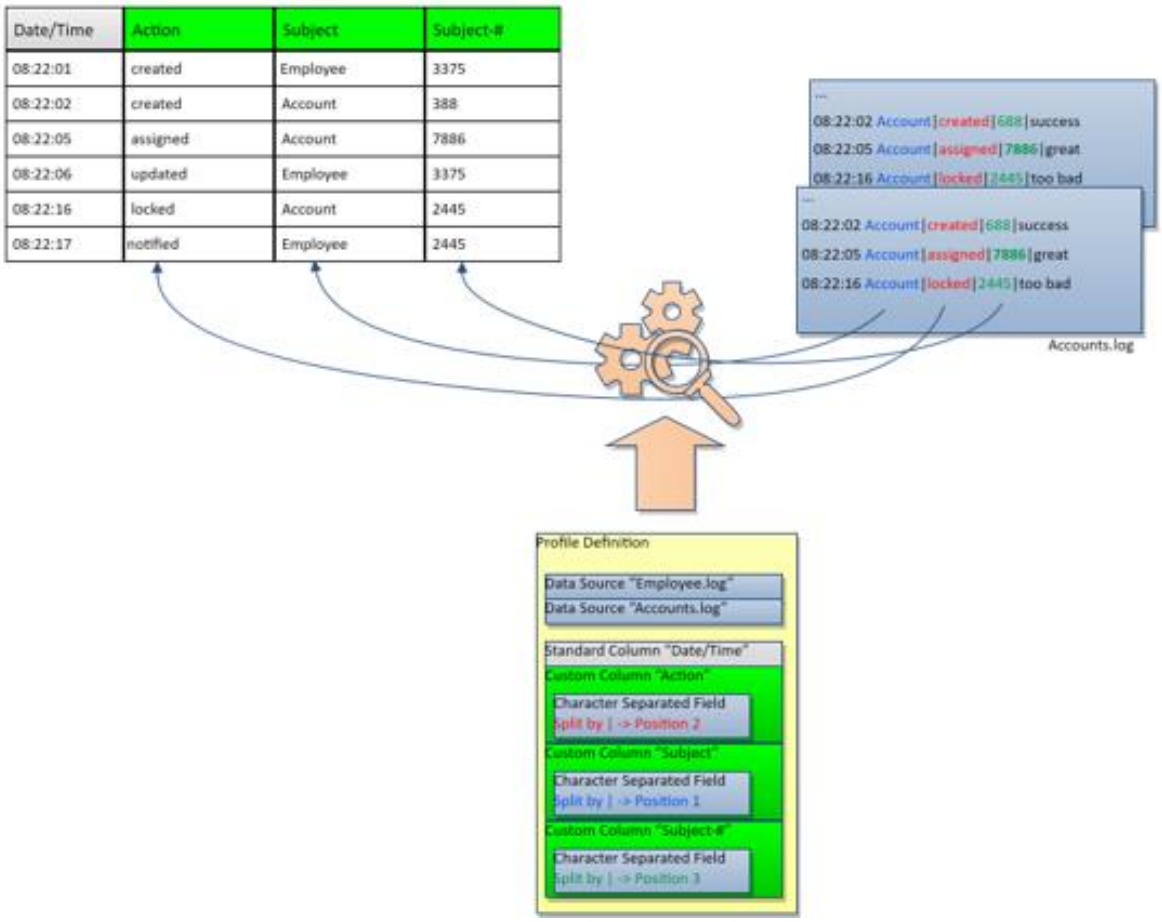

2.8.3 Field Detection

Field detection is about identifying and collecting potentially interesting Name/Value Pattern Fields present in log entries and to propose them to the user upon request to be assigned to custom columns. Currently Regex Fields and Character Separated Fields are not covered by the detection. Name/Value Pattern Field detection occurs in the following cases.

- Always During Autofind.
- During Search/Monitoring if the  toggle button next to the start button is enabled (in the Field Detection preferences page you can select/unselect the "Detection during search/monitoring" option to control whether detection during Search/Monitoring should be by default enabled or disabled).

Field detection tries to match the following name/value patterns.

Name/Value Pattern	Description
name=value	May appear at start/end of log entry or must be non-alphanumerically delimited. Name and value must not contain spaces.
name:value	May appear at start/end of log entry or must be non-alphanumerically delimited. Name and value must not contain spaces.
name='value'	Name must not contain spaces, but value may contain spaces.
name=”value”	Name must not contain spaces, but value may contain spaces.
name=(value)	Name must not contain spaces, but value may contain spaces.
name=[value]	Name must not contain spaces, but value may contain spaces.
name={value}	Name must not contain spaces, but value may contain spaces.
name=<value>	Name must not contain spaces, but value may contain spaces.
(name=value)	Name and value may contain spaces.
[name=value]	Name and value may contain spaces.
{name=value}	Name and value may contain spaces.
“name”=”value” (JSON)	Value in JSON may be a number, boolean, null or a quoted string.

2.8.4 Field Extraction

Field extraction is about matching individual field expressions during the search and monitoring process and to serve the extracted field values for being presented in corresponding result table custom columns.

- Field extraction is always performed during search and monitoring if the selected profile contains at least one visible custom column.
- Fields from individual custom columns are processed top-down for each result entry. As soon as a value is found, remaining fields are ignored.

It’s worth noting that a given Custom Column can have assigned several Fields of different types. For example, if you have two different log entry flavors in which in one entry a person number can be extracted using a Name/Value Pattern Field and in the second entry a person number has to be extracted using the Character Separated Field then Retrospective can easily support such situation. You just firstly assign the Name/Value Pattern Field to the Column as the first Field and then assign the Character Separated Field as the second Field. If the Name/Value Pattern Field will not match, then Retrospectives tries to extract the person number using the Character Separated Field.

3 CONFIGURATION AND SETUP

3.1 Changing Configuration Location

Retrospective uses a directory **.Retrospective** located in the user home to store configuration files inside release-dependent folders. In the following cases, this location

- Some users experienced long startup times and problems with the Retrospective disk storage (H2 database) because their user home is located on a network drive.
- When searching/monitoring large amounts of data, the space available in the user home may not be sufficient to hold all fetched data in the disk storage (H2 database).

To change the location of the configuration files to a preferred location, you need to proceed as follows:

- Shut down Retrospective.
- Create a local folder in your preferred location. This folder must be readable and writable for the user that runs Retrospective.
- Go to your Retrospective installation and locate the file "retrospective.ini".
  - Windows: <retrospective-installation-dir>\retrospective.ini
  - Linux: <retrospective-installation-dir>/retrospective.ini
  - Mac OS X: <retrospective-app-dir>/Contents/MacOS/retrospective.ini
- **Create a backup** of the file retrospective.ini
- Edit this file as follows:
  1. **Change** the property **-Dosgi.configuration.area** from

     `-Dosgi.configuration.area=@user.home/.Retrospective/6.0.0/.config`

     to

     `-Dosgi.configuration.area=C:/<some-path>/<your-customhome-folder>/6.0.0/.config`

  2. **Change** the property **-Dosgi.instance.area** from

     `-Dosgi.instance.area=@user.home/.Retrospective/6.0.0/.data`

     to

     `-Dosgi.instance.area=C:/<some-path>/<your-customhome-folder>/6.0.0/.data`

Note

There is a strict convention that you have to follow when configuring osgi.configuration.area and osgi.instance.area. The required pattern is as follows:

osgi.instance.area: `*/<your-custom-home-folder>/<retrospective-version>/.data`

osgi.configuration.area: `*/<your-custom-home-folder>/<retrospective-version>/.config`

Both folders, the .data and the .config, must reside in the same parent folder. See below for example configurations.

Example configuration for Windows:

….

```
-Dosgi.configuration.area=C:/Documents/customhome/6.0.0/.config
-Dosgi.instance.area=C:/Documents/customhome/6.0.0/.data
```

Example configuration for Linux:

```
…
-Dosgi.configuration.area=/home/dev/customhome/6.0.0/.config
-Dosgi.instance.area=/home/dev/customhome/6.0.0/.data
```

Example configuration for Mac OS X:

```
…
-Dosgi.configuration.area=/Users/dev/customhome/6.0.0/.config
-Dosgi.instance.area=/Users/dev/customhome/6.0.0/.data
```

Once you have applied these changes, you can save the file and restart Retrospective. In case you already created hosts, profiles and bookmark definitions you can either

- import these definitions from the ".Retrospective/6.0.0" folder found in your user home via the File→ Import Hosts and Profiles and File → Import Bookmarks
- copy the complete folder ".Retrospective/6.0.0" to <your-customhome-folder>

3.2 Program Registration

When installing Retrospective for the first time, you can use the UNREGISTERED program without any functional restrictions. Although such evaluation (trial) mode does not expire, a Retrospective license must be bought for continued use.

Please purchase a license online at:

http://www.retrospective.centeractive.com/content/retrospective-buy or request a quote from orders@centeractive.com in case you're interested in a multi-user license.

To register your Retrospective program, please proceed as follows:

1. Start the Retrospective program.
2. Select the menu item Help > Program Registration.
3. Within the program registration dialog, enter the license key.
4. Select the "Register Online" radio box.
5. Press the "Register Now" button.
6. In case your firewall blocks access to our registration server, please define a proxy server to get around this limitation and perform step 5 again.
   → The proxy server can be defined on the "Proxy Settings" page within the preferences dialog (menu item File > Preferences).

In case online registration didn't succeed; you may request registration by e-mail as follows:

7. Keep the registration dialog from previous steps open and make sure the license key is entered.
8. Select the "Request Registration by E-Mail" radio box.
9. Press the "Generate E-Mail" button. This should generate a registration request e-mail and open it in your preferred mail client.
   → In case no e-mail client is installed on your computer or Retrospective should be unable to locate or open it, the registration request data is stored at $USER_HOME/.Retrospective/6.0.0/registration_request.txt.
10. Send the registration request e-mail to support@centeractive.com.
11. Within a certain period, the centeractive support team will reply with the registration data.
12. Paste the registration data to the large text area and press "Register".

3.3 Host Manager

The ***Host Manager*** view lets you define and test SSH connections for accessing remote servers. This view can be opened by selecting View → Host Manager, clicking the icon or by using the [Ctrl] + [M] on Windows or [⌘] + [Z] on Mac key combination.

SSH Host definitions are required for searching and monitoring files stored on servers accessible via intranet or internet. They are used in cases when related data sources are defined on the fly (ad-hoc) or bundled within a profile. Hosts can be grouped into folders.

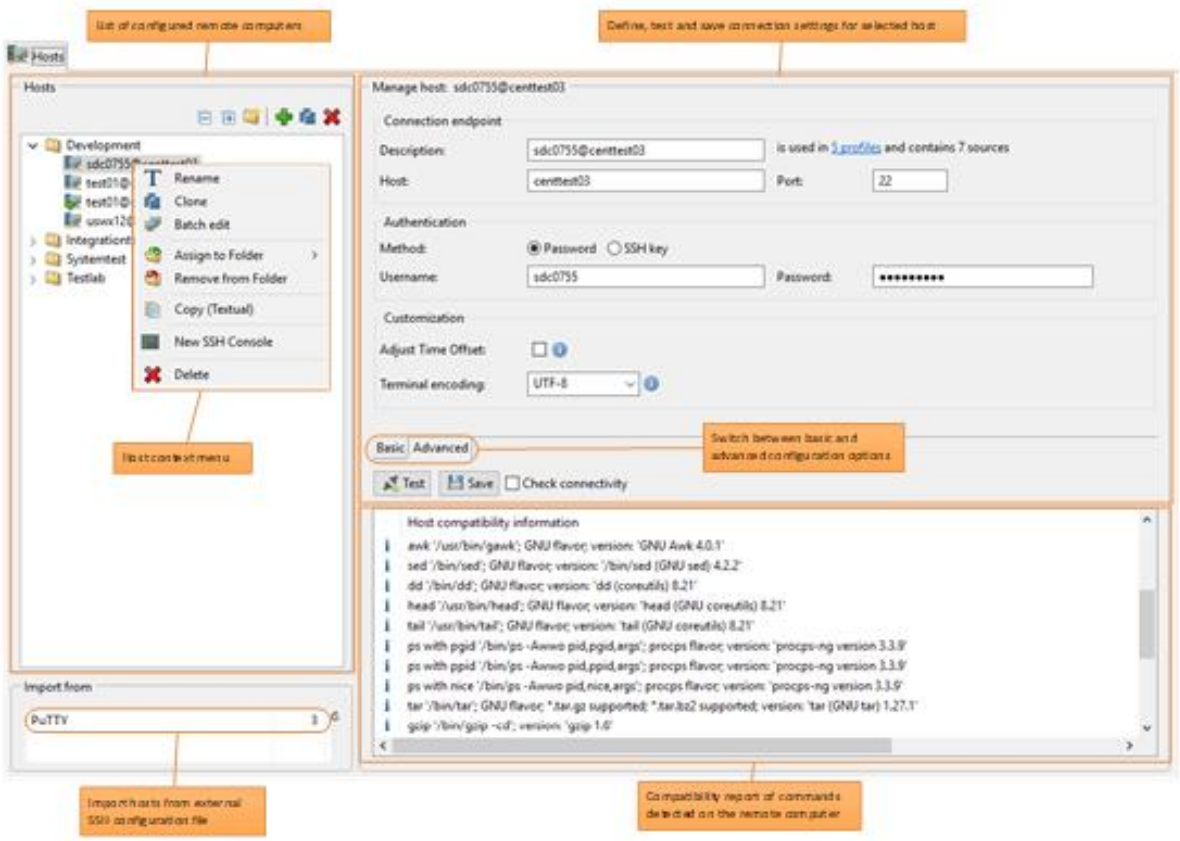

3.3.1 Create new Host

1. Click the [Add] button to create a new host definition. Optionally enter/choose the name of the folder, the host shall be part of.
2. Provide the computer name or the IP address together with the port number (22 is the default port for SSH).
3. Select authentication method and provide necessary authentication information.
4. To check, you can click the [Test] button to make sure that given host can be reached.
5. Select ***Check on save*** to have the connection tested before storing host definition.
6. Click the [Save] button to proceed with adding host definition.

3.3.2 Delete Host

1. Select desired host definition.

2. Click the [Delete] button or right-click given host definition and select ***delete*** to delete selected host definition.

Note

If data sources from a host are contained in data profiles, they will be automatically removed from there when you delete the host.

3.3.3 Edit Host

1. Select desired host definition to display connection details
2. Change connection parameters as desired and click the [Save] button to store changes.

3.3.4 Host Context Menu

Select a host from the list and choose the desired function in the context menu that pops up when you press the right mouse button.

Menu Item	Description
Rename	Changes the list entry to edit mode and lets you change the host description directly in the list. Once you've finished editing the host description, press [Enter] or select a different host from the list. Alternatively, you can also double click the desired host list entry to enter the edit mode.
Clone	Creates a copy of the selected host definition and adds it to the list. The description of the new entry is automatically extended with a postfix "(n)" since it is not allowed to have two hosts each with an identical description. Feel free to rename the auto-generated description and to adapt the connection parameters, then press the [Save] button to confirm your changes.

Batch Edit…

Retrospective lets you change the configuration parameters of multiple hosts in a single shot. The corresponding view appears if at least two hosts are selected. Press and hold the [Ctrl] button and click desired hosts or activate this menu and thereby automatically select the hosts from the context of the selected one (i.e. all hosts belonging to the same folder).

1. Activate the [Delete selected] button to delete the definition of all selected hosts.
2. Click the [Change] button to edit the SSH connection for the selected hosts
    - Enter the new username
    - Provide the new password
    - Press the [Change] button to store it.
3. Click the [Change] button to edit the SSH key for the selected hosts
    - Provide the new SSH key (and the key phrase if required)
    - Press the [Change] button to store the changes.

Assign to Folder

Assigns all selected hosts to the folder of your choice. If the desired folder does not yet exist, choose from the sub-menu and thus create a new one that will have the selected hosts added to it.

Hosts can also freely be dragged and dropped to the folder of your choice.

Note

After assigning hosts from a folder to a different one, you may end up having empty folders that are represented by a grayed-out folder icon.

Empty folders are not persisted and are lost when the program exits.

Remove from Folder

Removes the selected hosts from their folder and places them on the global area.

Hosts can also be removed from their folder by dragging and dropping them on the global area.

Note

After removing hosts from their folder, you may end up having empty folders that are represented by a grayed-out folder icon.

Empty folders are not persisted and are lost when the program exits.

Copy (Textual)	Copies the host definition in textual format to the clipboard. This can for example be pasted to a text editor for documentation purpose.
New SSH Console	Creates a new SSH Console and establishes a connection to the selected host using the configured authentication method (password- or key-based) and credentials.
Delete	Removes the selected host definition from the list

3.3.5 Connection Endpoint (Basic)

Here you need to provide the computer name or IP address together with the port number (22 is the default port for SSH).

The description is generated automatically from the entered data (format "user@host:port"), feel free to edit it for adapting it to your preferences.

Next to the description you'll notice a short text that indicates in how many profiles the host is currently used. A mouse-click on the link contained in this text opens a dialog where all these profiles and related data sources are listed.

3.3.6 Authentication (Basic)

Here you define the password- or key-based authentication method together with your credentials to be used for establishing new SSH connections to the selected host.

3.3.7 Customization (Basic)

3.3.7.1 Time Zone Fallback

When Retrospective collects log entries, it automatically converts their timestamps to the local time using the time zone information found in the log entry timestamps. This makes it easy to correlate log entries retrieved from servers lying in different time zones. In many cases however, applications write log entries without time zone information.

For further information, please consult chapter 8

Log Time Synchronization.

3.3.7.2 Adjust Time Offset

If you select the "Adjust time offset" checkbox, Retrospective will automatically adjust log entry dates coming from the remote host and the time search criteria passed to it. The goal of adjustment is to make the remote host time match the time on the local computer.

For further information, please consult chapter 8

Log Time Synchronization.

3.3.7.3 Terminal Encoding

This option defines the encoding used by the inbuilt SSH Console to:

- translate characters from the SSH console into bytes sent to the remote terminal
- translate bytes sent by the remote terminal into characters in the SSH console

You can either select one of the proposals from the drop-down list or enter an encoding of your choice in the text field. Selecting the default value "best-guess" means that Retrospective detects the remote terminal encoding automatically and then uses it in the SSH console. Remote terminal encoding is detected by looking into LC_ALL, LANG, LC_CTYPE and LC_MESSAGES environment variables on the remote host.

WARNING: In order for the SSH Console to work correctly, the selected/entered encoding has to be valid on your local Java Virtual Machine and it has to resemble the remote terminal encoding set in aforementioned environment variables. Therefore, please take caution when changing encoding from the default "best-guess". The "best-guess" setting should be fine and does not need to be changed in most typical situations.

3.3.8 Connection Options (Advanced)

3.3.8.1 Maximum Connections

By changing the maximum host connections, which currently defaults to 15, you can modify the number of simultaneous SSH connections that can be established to a host for a particular host configuration. This is useful in situations where you want to search or monitor log files while making sure that Retrospective consumes a limited amount of resources on the target host, which includes limited CPU consumption. CPU consumption limitation works exceptionally well if the host is a strong multicore machine. It can be assumed that one connection can maximally consume about 120% of a single core capacity during an intensive search operation. Therefore, if you have 8 cores, then limiting the connection number to 5 ensures that Retrospective will never consume approximately more than 5 * 120% = 600% of a single core, i.e. no more than 6 out of the 8 available cores. However, please be aware, that the less simultaneous connections, the worse performance Retrospective can offer. If you are concerned with Retrospective taking CPU from other host processes, then the execution priority modification (see section 3.3.9.2 Execution Priority) will allow you to prevent it in a less radical manner and without limiting the number of simultaneous connections.

Note that Retrospective does not open all connections specified in the discussed option right from the beginning of interacting with the host. Instead it increases the number of connections based on the amount of work it has to do and based on the host behavior. The contemporary number of connections is influenced by a **dynamic connection limit** that is always less or equal to the maximum connection limit. If there are no exceptions while opening new connections, Retrospective will gradually increase the **dynamic connection limit** allowing new connections to be opened. However, if some problems with connections opening start to appear, then the **dynamic connection limit** will be decreased. Here are some more detailed rules related to **dynamic connection limit** behavior:

- at the beginning the limit equals 2;
- increasing the limit is considered only when there is a need to open at least one connection more than the limit (for example the limit is 2 and you've just started to search data sources covering 3 different files);
- the limit is increased by one when connection to a given host was successfully opened at least 10 times in a row;

- the limit can be temporarily increased to prevent starvation (see below for explanation of starvation situation) of some activity, even if not enough subsequent connections were successfully opened. However, in such a case the limit is decreased as soon as any connection is released;
- the limit is decreased by 3 (or even more), when there is an error during connection's opening;
- the limit is never increased above maximum host connections and always at least one connection is allowed (the limit is always greater or equal to 1).

When the value is changed, there is no immediate effect to the connections that might already be open for this host configuration (because you may already be running a search on this host). Instead, a new connection pool is opened where the number of connections is restricted according to your settings. After saving the host definition, the new connection pool starts to be used and the connections from the old pool gradually expire. Actually, whether in a new pool or not, Retrospective ensures that a connection to given host is closed if not used for some assumed period. Currently this period is 5 minutes and cannot be changed in preferences (some connections needed for host compatibility checking - see section 3.3.11 Host Compatibility Information - are closed even sooner).

It has to be taken into account that the lower the maximum connection number, the higher the probability of encountering a starvation situation. A starvation situation happens when all connections to a given host are consumed. For example, if maximum connections is set to 1 and a long-term search on a big file is started, then any other activities such as searching or monitoring will idly wait until the search causing the starvation finishes. You will be warned periodically about starving activities (you can define the interval in seconds for these warnings, see section 4 Preferences). However, these warnings will appear only when the **dynamic connection limit** reaches the maximum connections specified in the Host Manager. Fortunately, even when starvation happens, you can still perform the light short-term activities such as host compatibility checking (see Section 3.3.11 Host Compatibility Information), file browsing (see section 5.5 File Browser), adding data source (see section 3.4.1 Create New Profile – point 4), configuring data source (see section 3.4.1 Create New Profile – point 9) and performing Autofind (see section 3.4.3.1 Autofind). This is possible because Retrospective allows opening just ONE more connection beyond the limit of the maximum host connections option. Therefore, when checking the currently opened connections on SSH host expect that there can be at most one more connection then the number specified in the Host Manager. Of course, since this is only a ONE additional connection, when experiencing starvation, you will not be able to perform Autofind and e.g. file browsing simultaneously. File browsing will have to wait until Autofind finishes. Rest assured that this ONE additional connection is never used for searching or monitoring, so it will naturally expire after light activities are finished. Also, be aware that searching is more aggressive in terms of connection acquisition then monitoring. Therefore, when you have searching competing with monitoring for some restricted number of connections, simply assume that searching always wins and monitoring starves.

Note

The maximum host connection configuration is independent from the jump server configuration (see 3.4.8.2 Jump Server). The restriction is related only to the target host and not to the jump server present in the connection chain for this host.

Note

The maximum connections setting applies only to the host configuration that you are currently editing. Assuming that you have another host configuration for the same host name, which differs in some other setting (for example different

credentials) then both configurations are independent in terms of connection pooling, i.e. both configurations have their own connection pool.

3.3.8.2 Jump Server

A jump server or jump host is a special-purpose computer on a network, typically used to manage devices in a separate security zone. If your environment needs to connect through a jump server to the target host in order to search/monitor log files, Retrospective provides the option of defining a jump server for a host. You can even chain more than two hosts together, i.e. you can access your target host through more than one jump server.

To define a jump server, switch to the Host Manager "Advanced" tab, then simply select a host from the drop-down list labeled "Jump server" in the "Connection options" section. Note: If you do not select a jump server, your host will be accessed directly, thus the default selection says "Direct Access".

The selectable hosts in this list are basically all the hosts you have defined through the host manager tab excluding the current host you are editing. Once you select a jump server, the "Connection chain" appears at the bottom of the "Connection options" section. It displays the chain of connected hosts, i.e. you are able to see the route that Retrospective takes to your host.

A green arrow on the host icon  means that this host is accessed through a jump host or a proxy server. This modified icon appears in the "Connection chain" as well as in the host list on the left side of the Host Manager.

The following screenshot shows the main elements related to Jump Server configuration described above:

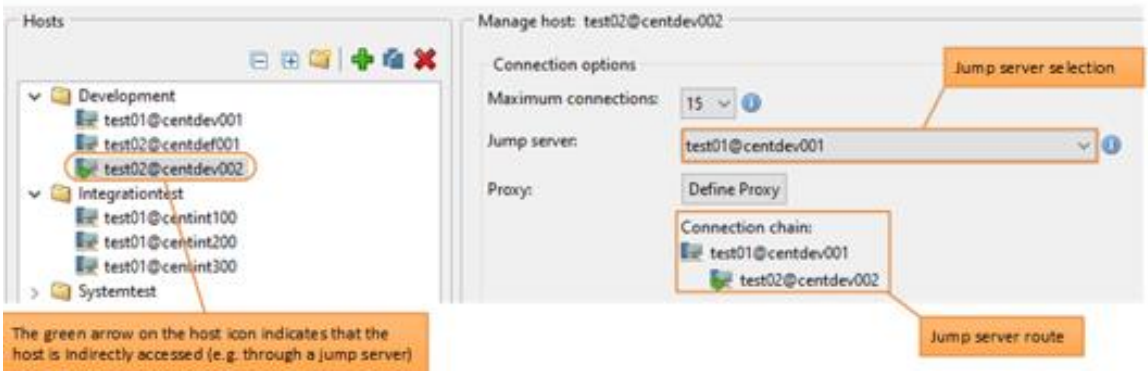

Note

Retrospective does not allow you to create jump server loops. Should your configuration contain a loop, an error is shown in the host manager form and you are not allowed to save your configuration.

Note

If a host is configured to be accessed through a jump server, then its IP address is evaluated by the jump server. Therefore if, for example, you specify the IP of a host in local network (e.g. 192.168.x.y), then the host in the local network accessible for the jump server will be accessed (not the host in the local network accessible for the host running Retrospective).

3.3.8.3 Proxy Server

A proxy server, also known as a "proxy", is a computer that acts as a gateway between a local network (e.g. all the computers at one company or in one building) and a larger-scale network such as the internet. Proxy servers provide increased security as they help preventing an attacker from invading a private network. If your target host needs to be accessed through a proxy server in order to search/monitor log files, Retrospective lets you define an **HTTP** or **SOCKS** proxy server for individual hosts. You can even chain a proxy server with a jump host, i.e. you can access your target host through a proxy server and a jump host.

To define a proxy server for a particular host, switch to the Host Manager "Advanced" tab. When you press the "Define Proxy" button inside the "Connection options" section, a dialog is shown that lets you configure the proxy server.

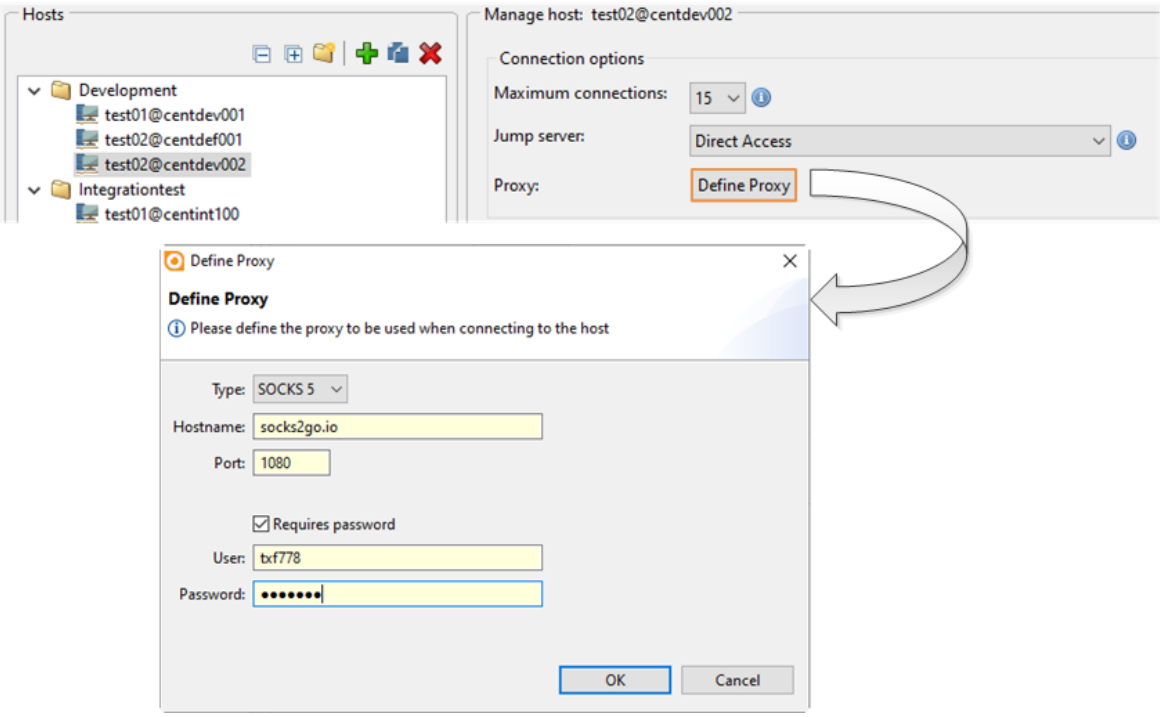

Once a proxy server is defined for a target host, the "Connection chain" appears at the bottom of the "Connection options" section. It displays the chain of devices that are involved when Retrospective establishes a SSH connection to the target hosts.

A green arrow on the host icon means that this host is accessed through a proxy server or a jump host. This modified icon appears in the "Connection chain" as well as in the host list on the left side of the Host Manager.

The following screenshot shows the main elements related to Jump Server configuration described above:

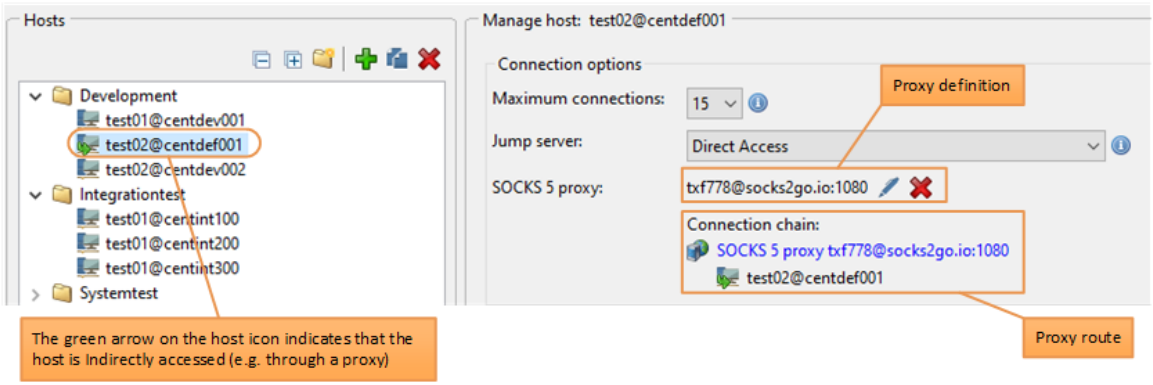
Hosts
Development
test01@centdev001
test02@centdef001
test02@centdev002
Integrationtest
test01@centint100
test01@centint200
test01@centint300
Systemtest
The green arrow on the host icon indicates that the host is indirectly accessed (e.g. through a proxy)
Manage host: test02@centdef001
Connection options
Maximum connections:
15
Jump server:
Direct Access
SOCKS 5 proxy:
txf778@socks2go.io:1080
Proxy definition
Connection chain:
SOCKS 5 proxy txf778@socks2go.io:1080
test02@centdef001
Proxy route

3.3.9 Target command modification (Advanced)

Retrospective allows the user to operate on remote hosts by executing shell commands through SSH. The "Target command modification" section provides the following options to modify the executed commands:

- Changing the user identity when executing shell commands on the host, see section 3.3.9.1 Identity Change.
- Decreasing the execution priority / CPU time consumption of executed shell commands on the host, see section 3.3.9.2 Execution Priority.

Note

These configuration options require certain tools to be available (for example the nice tool) or certain configurations on the target host. It is recommended that you always use the "Test" button to verify if your host configuration matches the setup of your target host or that you have the checkbox "Check connectivity" selected when saving a host configuration with target command modification settings.

3.3.9.1 Identity Change

If the log files that you want to search/monitor require the security privileges of a different user (normally the superuser – root) than the one you have configured in the Authentication section of the host manager tab, you can configure Retrospective to change your identity on a host and execute all shell commands on behalf of the configured target user. Retrospective offers the following approaches for the user identity change:

- sudo: Retrospective uses sudo to execute shell scripts as the target user
- sudo su: Retrospective uses sudo to become super user and then uses su to execute all shell scripts as the target user.

The configuration you create in Retrospective should correspond to the configuration in the /etc/sudoers file on the target host. The /etc/sudoers file offers a high number of configuration options, including some of the following: not requiring a password for certain commands, requiring a password for users or groups etc. Detailed description of /etc/sudoers syntax and configuration options is out of the scope of this manual, however the configuration elements required by Retrospective are discussed further below.

To change your identity on a host, proceed as follows:

1. Select an approach from the drop-down list labelled "Identity change" in the "Target command modification" section of the Host Manager. Selecting an approach makes all options related to configuring an identity change visible.
2. Enter the target user for the identity change. By default it is set to root, but you can enter another user. Note that the target user is mandatory so this text box cannot be left empty.
3. If (and only if) the user defined in the Authentication section of the Host Manager tab requires a password to execute sudo (as defined in /etc/sudoers), you have to mark the checkbox "Requires password".
4. Once you mark the "Requires password" checkbox (as stated in the previous step), Retrospective has to provide a password when executing sudo on the host. Depending on your authentication method configured in the Authentication section of the Host Manager tab, the following scenarios are possible:

- If you've chosen password authentication in the Authentication section, Retrospective has to provide the password of the user defined in the Authentication section when executing sudo. As this password is already available from the Authentication section you do not have to enter it again. Thus, the password text box in the "Identity change" section shows the hint "password reused of user <your-user>".
- If you've chosen "SSH key" in the Authentication section, you are required to enter the password of the user defined in the Authentication section by yourself. Thus, the password text box in the "User identity change" section appears red (required) and shows the hint "enter password of user <your-user>".

5. If you mark the checkbox "Initial login", the shell environment will be replaced with the shell environment of the target user. This means that for example login-specific resource files such as .profile or .login will be read by the shell.
6. The "Target command" in the "Target command modification" section shows the resulting shell command based on your configuration settings.

The following screenshot shows the input elements as described above:

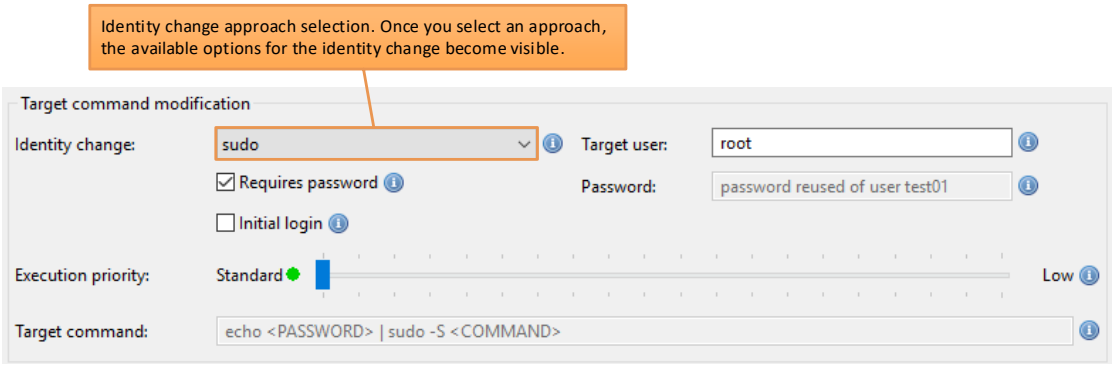

Note

Identity change settings have no effect if configured for a host that is used only as a jump server.

The basic syntax of `/etc/sudoers` file (note that the `sudoers` file can be also placed somewhere else in the filesystem other than the `/etc` directory) is as follows:

```
f1        f2    f3       f4
user      ALL=(ALL)      ALL
@group    ALL=(ALL)      ALL
```

The meaning of each field is as follows:

- *f1* – user (or a group) that is allowed to use sudo. This field has to contain the user provided in the Authentication section of Host Manager (or a group to which this user belongs).
- *f2* – list of hostnames. It is rarely used and is not that relevant for Retrospective. For explanation please see `sudoers` manual.
- *f3* – list of target users. Simply user (or a group) from field *f1* is allowed to change identity to these users. Depending on the selected approach for the user identity change, this field has to contain the following:
  - sudo: Target user provided in the Identity change section of Host Manager (or '`ALL`' that matches all users);
  - sudo su: `root` or '`ALL`' that matches all users.

- *f4* – list of commands that can be executed by the user (or a group) from field *f1* after changing identity. Depending on the selected identity change approach, this field has to contain the following:
  - sudo: '`ALL`' that matches all possible commands (Retrospective executes many typical *NIX commands on the target host and defining all of them cannot be easily done);
  - sudo su: Assuming that su tool is placed in `/bin/su` and `TARGET_USER` is the user provided in the Identity change section of Host Manager, then depending on the selection of Initial login option, the field has to contain the following:
    - Initial login NOT selected: `/bin/su TARGET_USER *` (please note that if the user provided in the Identity change section is root, then the field should contain: `/bin/su *`);
    - Initial login selected: `/bin/su - TARGET_USER *` (please note that if the user provided in the Identity change section is root, then the field should contain: `/bin/su - *`).

Note

When Requires password option is not selected in the Identity change section of Host Manager, then the content of *f4* field has to be prepended with the "NOPASSWD:" token.

In order to make the `/etc/sudoers` configuration required by Retrospective easier to comprehend, several examples are presented below. Therein we assume **retrouser** to be the username that was entered in the Authentication section.

Example 1:

Retrospective:	Identity change: **sudo**	Target user: **root**	Requires password: yes	Initial login: either
`/etc/sudoers:`	`retrouser`	`ALL=(ALL)`	`ALL`	

Example 2:

Retrospective:	Identity change: **sudo su**	Target user: **root**	Requires password: yes	Initial login: yes
`/etc/sudoers:`	`retrouser`	`ALL=(ALL)`	`/bin/su - *`	

Example 3:

Retrospective:	Identity change: **sudo**	Target user: **jboss**	Requires password: yes	Initial login: either
`/etc/sudoers:`	`retrouser`	`ALL=(jboss)`	`ALL`	

Example 4:

Retrospective:	Identity change: **sudo su**	Target user: **jboss**	Requires password: yes	Initial login: yes
`/etc/sudoers:`	`retrouser`	`ALL=(ALL)`	`/bin/su - jboss *`	

Example 5:

Retrospective:	Identity change: **sudo**	Target user: **jboss**	Requires password: no	Initial login: either
`/etc/sudoers:`	`retrouser`	`ALL=(jboss)`	`NOPASSWD:ALL`	

Example 6:

Retrospective:	Identity change: **sudo su**	Target user: **jboss**	Requires password: no	Initial login: yes
`/etc/sudoers:`	`retrouser`	`ALL=(ALL)`	`NOPASSWD:/bin/su - jboss *`	

Example 7:

Retrospective:	Identity change: **sudo su**	Target user: **root**	Requires password: no	Initial login: yes
`/etc/sudoers:`	`retrouser`	`ALL=(ALL)`	`NOPASSWD:/bin/su - *`	

Note

In all the examples above, `retrouser` present in field *f1* can be replaced with a group to which `retrouser` belongs.

Typically all users that are allowed to execute sudo are added to a `sudo` group.

3.3.9.2 Execution Priority

Execution priority configuration allows the user to lower the priority / CPU time consumption of all operations performed by Retrospective on a particular host. For example, this is useful when running a search on a production environment in which intensive usage of CPU could introduce an unwanted impact (e.g. application running in the environment slows down).

When moving the slider labeled "Execution priority" in the "Target command modification" section of the host manager tab to the right (towards "low") you decrease the priority of all Retrospective operations on this host.

Retrospective translates the execution priority setting into a POSITIVE nice value used for all shell commands it executes on this host.

The nice is a *NIX tool, used to modify the scheduling priority of a particular process. Note that Retrospective only allows you to configure an execution priority which results in more "niceness", i.e. you can only LOWER the scheduling priority of Retrospective operations on a host (lowering scheduling priority means increasing the "niceness" that is also referred to as "nice level").

In order to modify the execution priority, the nice tool has to be installed on the target host. If it is not installed, Retrospective will display an error message when you click "Test" on the host manager tab (the error will also be displayed during any other operation on the host, e.g. searching or monitoring some log files). In case of such an error message, you can either install the nice tool or simply disable modification of execution priority by moving the slider maximally to the left (then the nice tool will not be used at all – it will disappear entirely from the "Target command" text box).

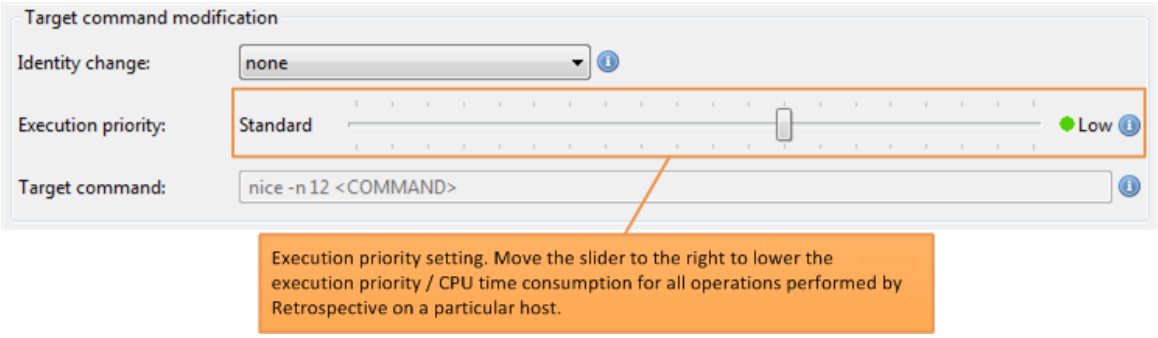

3.3.10 Import Host Configurations

Retrospective can read host definitions from `user/.ssh/config` or PuTTY configuration files stored in the `user/.ssh/` folder.

1. Select desired hosts configuration file in the ***Import from*** area, to display the list of host definitions.
2. Double-click desired host definition or right-click it and select ***Import*** from the context menu.
3. Edit connection parameters as needed and click the [Save] button to store new host definition.

Note

When importing hosts, PuTTY private keys (*.ppk) are ignored. Retrospective currently only support OpenSSH keys. There exist number of HOWTOs on the internet that explain how to convert a PuTTY private key into an OpenSSH key (i.e. https://www.simplified.guide/putty/convert-ppk-to-ssh-key).

3.3.11 Host Compatibility Information

The SSH connection to a remote server can be checked in two ways: when the user presses the "Test" button or when the user activates the "Save" button while the "Check connectivity" check box is selected. When the SSH connection can successfully be established, Retrospective performs some compatibility checks on the remote server and provides a report in a table that appears below the "Test" and "Save" buttons.

Retrospective operates on remote servers with the use of several flagship tools available at *NIX system. Information about host compatibility is related to availability of these tools and their functionality.

1) Why does Retrospective check the host compatibility?

In order to ensure that the set of available tools is sufficient for using Retrospective on the host and to notify the user about any minor (warnings) or major (errors) issues related to the tool set.

2) What exactly is checked and what is reported?

Checking the host compatibility covers the following tests:

- checking the operating system (currently the following are supported: Linux, FreeBSD, MacOS, Solaris, AIX, HP-UX)
- identification of the locale and encoding (performed by analyzing LC_ALL, LANG, LC_CTYPE, LC_MESSAGES environment variables in exactly this order)
- identification of the effective username and the effective user groups
- identification of the home directory
- availability of essential tools needed for basic scripting: command, echo, printf, grep, awk, uname, diff, ps
- availability and functionality of tools needed for searching: grep, fgrep, awk, sed
- availability and functionality of tools needed for monitoring: dd, head, tail
- availability and functionality of tools needed for uncompressing: tar, gzip, zip, bzip2
- availability and functionality of tools needed for target command modification (see section 3.4.9 Target command modification): sudo, su, sftp-server, nice (also identification of the effective nice level)

Note

When LC_ALL, LANG, LC_CTYPE and LC_MESSAGES environment variables are not set or cannot be parsed, then 'en_US' is assumed as the default locale and 'UTF-8' as the default encoding.

The report contains four types of messages:

- information about the detected operating system
- information, one for each available tool, which contains: tool name, path (sometimes arguments), flavor (typically GNU or OTHER) and version
- warning which communicates that a given tool is generally usable but is not fully functional or has performance issues
- error which communicates that Retrospective cannot be used because given tool is not available or has critical functional issues.

3) What should the user do if he sees warnings or errors?

In case of error, the user should install a tool which is missing for proper Retrospective operation. For example, if error notifies that awk tool is not available. Then it has to be installed to use Retrospective on given host.

In case of a warning, the user should install a tool, which is advised for optimal (in terms of performance or functionality) Retrospective operation. For example, if the gzip tool is missing, then user should install it to operate on *.gz files. Another possible warning example is notification that only FreeBSD fgrep is available, then the user should install GNU fgrep for improved searching performance.

4) What problems may arise if warnings and errors are ignored?

If an error is ignored, then Retrospective cannot be used on given host.

If a warning is ignored, then the user may experience limited functionality (e.g. cannot search through compressed files) or deteriorated performance (e.g. searching through large log sets takes significant amount of time).

3.4 Profile Manager

The ***Profiles*** view lets you configure data profiles and columns. This view can be opened by selecting View → Profiles, clicking the icon or by using the [Ctrl] + [P] keys combination.

Profiles are user defined persistent sets of data sources (log file and container definitions) from the local file system, from one or many remote servers and from Docker or Kubernetes subsystems. Search and monitoring actions performed on a specific profile will process (analyze the content) of all contained files and containers and present data in the result table using the columns you specified. You can work with Retrospective with data sources that are created on the fly (ad-hoc) and not bundled in a profile. Sooner or later however you'll get to a point where you have to re-execute a search on the exact same set of data sources as in the past. That's when you realize that having predefined profiles is definitely a good solution.

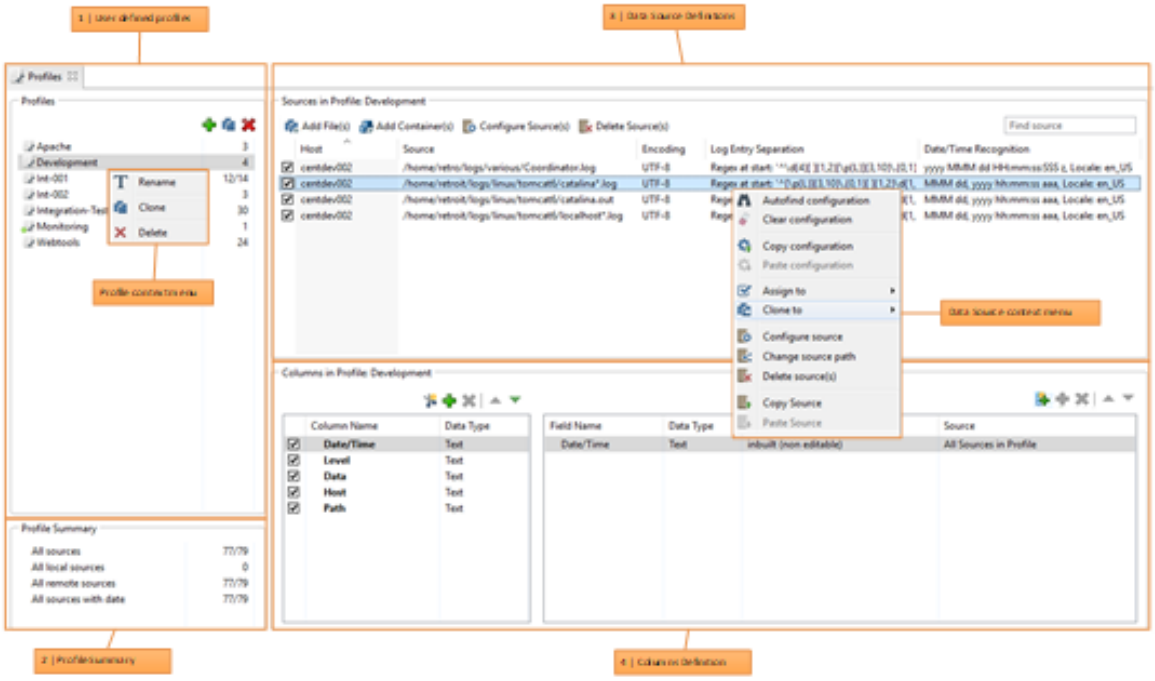

The ***Profiles*** view is divided into three sections:

1. ***Profiles***, containing the list of the user created profiles
2. ***Profile Summary***, containing data source summaries fetched from all existing profiles. Such groups let you view all data sources from any profile, and all local or remote data sources etc.
3. ***Sources in Profile***, containing the list of the data sources included in the currently selected profile. This section lets you add or remove data sources and enables you to individually configure them by specifying the encoding, date/time pattern etc.
4. ***Columns in Profile***, containing the columns that will be used for representing the search and monitoring data in the result table.

3.4.1 Create new Profile

1. Click the icon to open the ***Profiles*** tab.
2. Click the [Add] button to create a new data profile.
3. Change the profile's name if desired.
4. Click the [Add Files(s)] button in the ***Profiles*** detail section to add log files definitions (alternatively press the [Add Container(s)] button if you're interested in adding container definitions.
5. Select ***local drive*** in order to add files stored on the local file system or choose a remote SSH host. Alternatively press the [Add host...] button to define a new remote host.
6. Provide connection details to the remote machine and click the [OK] button to establish connection with remote host.
7. Browse the file system to find the folder containing the desired log files.

Note

User can also manually edit the directory path appearing in the text box. When editing the text box, wildcard characters ? and * can be used on any path elements. In order to understand how wildcards in the directory path are evaluated to directories see section 3.4.3.5 Deep Wildcards. When editing the text box, the ~ character can also be used to refer to the user home directory.

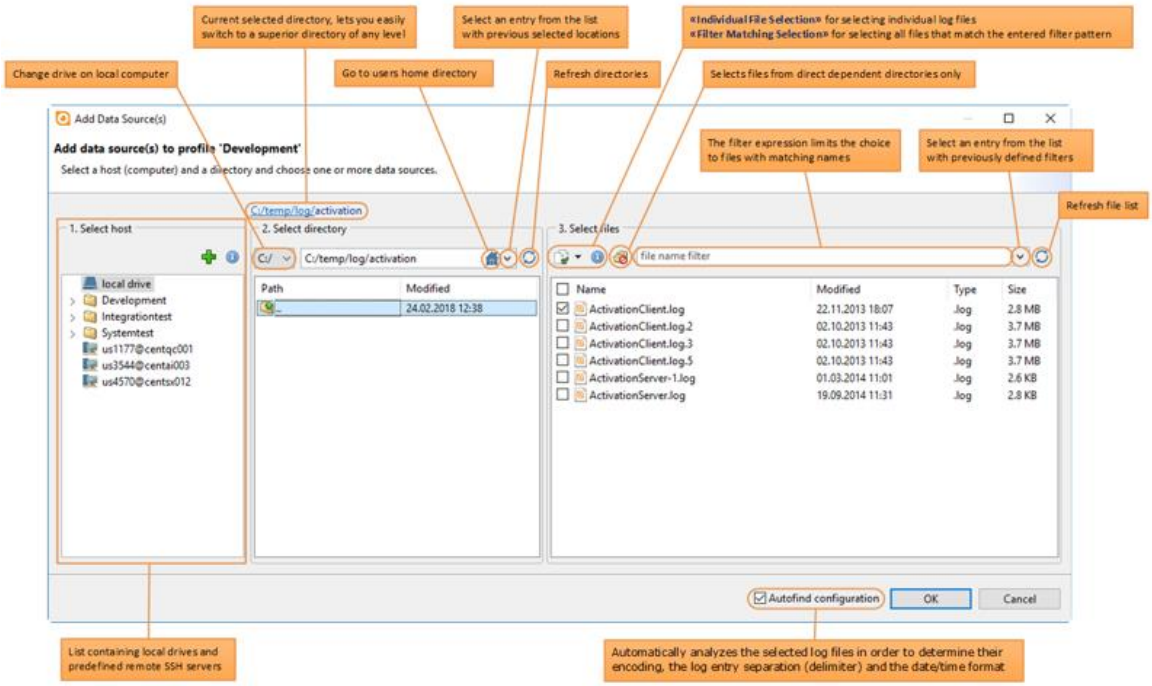

8. Select desired log files or containers and click the [OK] button.

Individual File Selection

The "Individual File Selection" mode by default lets you select individual log files from the current directory. You can reduce the number of listed files by entering a filter expression. Files will appear only if their name matches the file filter and a wild card is used. The file filter can contain the wildcards ? and * (i.e. *.log).

If the sub-directory button is enabled, Retrospective considers files from direct dependent directories only. The optional sub-directory filter lets you constrain the sub-directories to show the contained files. The sub-directory filter can contain the wildcards ? and *.


Filter Matching Selection

The "Filter Matching Selection" mode by default selects all files from the current directory. The choice of files can be limited by entering a filter expression. The group will only contain the files if their name matches the filter expression. The filter expression can contain the wildcards ? and *.

- An empty file filter or * selects all files from the current directory
- *.log selects all files with a name that ends with .log
- etc.

If the sub-directory button is enabled, Retrospective considers files from direct dependent directories only. The optional sub-directory filter lets you limit the number of sub-directories to be included in the group. The sub-directory filter can contain the wildcards ? and *.

9. Click the [Configure source...] button, or right-click desired log file and select ***Configure source*** from the context menu, to change log file encoding or define the log entries separation.

Note

Sources stored on the local file system can be dragged and dropped from ***Windows Explorer*** into the ***Profiles*** tab. Dropping a source file on the ***Profiles*** area will result in adding the given log file to the current profile. Dropping the file on the ***User Profiles*** section will result in the creation of a new profile with given source file defined.

3.4.2 Profile Context Menu

Select a profile from the list and select the desired function in the context menu that pops up when you press the right mouse button.

Menu Item	Description
Rename	Changes the list entry to edit mode and lets you change the profile name directly in the list. Once you've finished editing the profile name, press [Enter] or select a different profile from the list. Alternatively to enter the edit mode you can also double click the desired profile list entry. **Note** Renaming a profile used in a bookmarked search definition will render that bookmark unusable.
Clone	Creates a copy of the selected profile and adds it to the list. The name of the new entry is automatically extended with a postfix "Clone" since it is not possible to have two profiles with identical names. Feel free to rename the auto-generated profile name. You'll most probably also want to change the assigned data sources to make the profile different from the original one.

Menu Item	Description
Delete	Removes the selected profile from the list

3.4.3 Define Data Sources

A profile consists of one or many data sources. Individual data source can point to a local or remote (connected through SSH) computer or a container subsystem (Docker/Kubernetes) by a path definition representing either a single data source (file or container) or a group of data sources (e.g. '`Program*.log`'). It is mentioned in Section 3.4.1 Create new Profile, where the relevant selection options are discussed: "Item Selection" is pointing to one or several individual data sources whereas "Filter Matching Selection" includes all data sources from a directory or its sub-directories or a subset of them that match the entered filter.

Retrospective is most commonly used for analyzing log files and when it comes to log files, there are three aspects which need consideration. Firstly, each log file can have specific encoding (e.g. UTF-8 or latin1), which has to be known in order to properly read its contents. Secondly, log files typically contain specific strings or single control characters (i.e. line breaks) which allow Retrospective to differentiate one log entry from another and display search results in the way which is most convenient for the user. Thirdly, very often log entries contain date information, which can be appropriately interpreted and filtered by Retrospective. All three of these aspects have to be covered when defining a data source. The user can either allow Retrospective to automatically discover the relevant information by using the procedure presented in the first subsection or provide the information manually as discussed in the second subsection.

It is worth mentioning that although Retrospective is mostly used for processing textual log files, it is also capable of handling binary files. However, the Autofind procedure is not suited for binary content and Retrospective always tries to finally represent the content as characters (in accordance with the configured encoding). Therefore, with the exception of very simple file lookups, the processing of binary files is discouraged.

Another feature worth mentioning is a **support for compressed files**. When a compressed file is specified in the data source configuration, then during processing it is uncompressed on the fly and appropriately handled. For now the following compressed files are supported by Retrospective: *.zip, *.gz, *.bz2, *.tar, *.tar.gz, *.tgz, *.tar.bz2, *.tbz2. However, in order to ensure support for compressed files on remote hosts, appropriate tools have to be available on these hosts (for details please see section 3.3.11 Host Compatibility Information).

3.4.3.1 Autofind

Autofind is performed by Retrospective on one or several log files or container log streams in order to determine their encoding, log entry separation (delimiter) and date/time format. This information is needed for subsequent search and monitoring processes to identify individual log entries and correctly represent them in the result table. Autofind also tries to detect and collect key/value pattern fields that may be used for defining profile specific custom columns.

Note

When new files/containers or an unnamed set of data sources (Add Data Source(s) dialog) are added to a profile, Retrospective runs the Autofind procedure unless the "Autofind configuration" checkbox is deselected.

Retrospective however always performs Autofind when you drag and drop a log file.

1) Why is Autofind needed?

This step is needed to relieve the user of having to manually configure each data source. A data source can be an individual log file/container (specified in File/Container Browser by Item Selection) or a filter (specified in File/Container Browser by "Filter Matching Selection" e.g. '*.log'), which is a text pattern that matches the names of one or several files/containers within one or multiple (see section 3.4.3.5 Deep Wildcards) directories.

2) Why can Autofind not be performed directly during search or monitoring?

The Autofind procedure is time consuming and the user may want to change the discovered values prior to start searching or monitoring. Therefore, we assumed that performing Autofind during search or monitoring would, in most cases, be inconvenient for the user.

3) What data source related information is detected during Autofind?

Autofind needs to detect the following information:

- File encoding
- Log entry separation (delimiter) which includes:
  - i) Regex pattern of the delimiter.
  - ii) Its location in the log entry (either entry's start or end).
- Date/Time format which includes:
  - i) Date locale.
  - ii) Date format string compliant with the Java specification found at https://docs.oracle.com/javase/8/docs/api/java/text/SimpleDateFormat.html.
  - iii) Regex pattern automatically generated from date format string.
- Fields:
  - i) Fields matching predefined name/value patterns are detected and remembered in order to be proposed for custom column definition.

Note

Detection of file encoding is performed by analysis of the initial file bytes and matching them to the most probable encoding. When the initial file bytes contain only ASCII characters then the UTF-8 encoding is assumed. However, when in such situation, Autofind is performed on a remote ssh host, in which ISO-8859-1 was detected as the default encoding (see section 3.3.11 Host Compatibility Information), then Autofind assumes ISO-8859-1 encoding for such a file.

Note

Detection of the date locale is not easy. When Autofind is invoked:

- for a local data source, then initially, the default Java locale of the machine on which Retrospective is launched is checked.
- for a remote data source, then initially, the default locale of the remote server (specified by LC_ALL, LANG, LC_CTYPE or LC_MESSAGES environment variables) is checked.

4) What can go wrong with Autofind in general?

- i) There could be some connection issues when accessing the remote server. In such cases, an explicit message about the problem is displayed. Then, the user has to solve the problem and perform the Autofind procedure again.
- ii) Whenever a configuration problem occurs during Autofind (connection issues are not considered to be configuration problems), the user is notified and instructed to see

the data source configuration (presented in Section 3.5 Profile Manager) for additional information. If there are any warning or error icons, the user can hover over them to get the information. For example, if a file is not accessible (some permission problems) or it is not available at all (has been removed). Another example could be that the detected regex pattern of the delimiter is too long. Additional information about other possible warnings and errors is provided in point 5).

iii) Even though many precautions are taken on the implementation level, the Autofind could make a wrong assumption or simply make a mistake during detection. Then, the user has to make a manual correction of the values that are proposed by Retrospective.

iv) Several different issues have to be taken into account when a data source is specified by "Filter Matching Selection". This is covered in the subsequent point.

5) What can go wrong when Autofind is performed on a data source specified by "Filter Matching Selection"?

When a data source is specified by "Filter Matching Selection" then, most of the time, it refers not to a single file but to a set of files. In such situations, separate files could have different encodings, log entry separations or date/time formats.

In the case of "Filter Matching Selection", the data source can refer to a high number of files and Autofind cannot be performed on all of them because it would take too much time. The maximum number of files analyzed during Autofind is controlled by an option placed on the Data Sources page within the preferences dialog. If only a subset of available files is to be analyzed, the program attempts to parse different file types (each having largely distinct names) and thereby increases the chance of discovering differences in encodings, log entry separations and date/time formats.

Note

When data sources are specified in the "Item Selection" mode (each file/container selected individually) and the Autofind is performed for the whole profile, then the combined number of files/containers analyzed during Autofind is not limited. The preference option – Maximum number of files to be analyzed – limits the number of analyzed files/containers in the context of single data source (Filter Matching Selection) but not in the context of the whole profile.

When Autofind senses differences in one of the detected information (encoding, log entry separation or date/time formats), then the value present in the highest number of analyzed files is retained as the result of the Autofind procedure. Then, depending on the type of information, the following applies:

- ***Encoding***: The data source is marked with a warning icon whose tooltip provides information about files where encoding was different from the retained one;
- ***Log entry separation***: The data source is marked with an error icon whose tooltip provides information about files where the log entry separation was different from the retained one;
- ***Date/Time format***: The data source is marked with a warning icon whose tooltip provides information about files where the date/time format was different from the retained one;

On each mentioned tooltip, the user can hover over "files analyzed" to discover the files set that was analyzed by the Autofind procedure.
When a data source is marked with either warning or error, then the user should change the

data source configuration (e.g. refine the Filter) to only include files which have identical encoding, log entry separation and date/time format. This is especially critical in case of difference in log entry separation, because then, when searching or monitoring is started, Retrospective may not be able to split individual log entries from certain files and could use up all its available memory.

Note

Detection of file encoding in case of "Filter Matching Selection" tries to find the encoding that is compatible with all data source files/containers. Therefore, a following situation can occur. You have 5 files and when invoking autofind in File Selection separately for each of them, you get encoding X (e.g UTF-8) for the first 3 files and Y (e.g. ISO-8859-1) for the last two files. However, when you invoke autofind in "Filter Matching Selection" for all 5 files you will see that Y is chosen for all five files without any warnings (see above). It is possible because encoding detection not only considers single matching encoding but rather a list of encodings that can be potentially used for reading given file. Therefore, it may be possible that the encoding Y can be used for reading all five files although its matching probability is less than the one of encoding X for the first 3 files. Autofind in "Filter Matching Selection" tries to compromise to achieve the best possible encoding configuration for given set of files or containers.

6) When is manual configuration/correction required?

Manual configuration is required when the user does not want to rely on Autofind and in cases described in the previous points:

- When Autofind makes a mistake during detection;
- When dates in the log file entries contain not only milliseconds but also microseconds or even nanoseconds. Is such cases only the milliseconds will be included in the date format while micro and nanoseconds will be simply considered to be a part of the entry content. For example in case of the following log file content: "`04.05.2016 04:15:29,739345678 ERROR User authentication failure`"

  Autofind will detect the following date/time format: `dd.MM.yyyy HH:mm:ss,SSS` – no micro or nano seconds. The problem is that when the searching or monitoring is started in this file the entry content appearing in the Data column of the Result Table will be as follows: "`345678 ERROR User authentication failure`".

  Thankfully, Retrospective allows you to manually configure date/time format in a way that prevents micro/nano seconds from appearing in the Data column of the Result Table. You need to use a special "`i`" token in the date/time format to mark each character that should be ignored. So in the provided example you could manually set the following date/time format: `dd.MM.yyyy HH:mm:ss,SSSiiiiii` . Thanks to that, the Data column of the Result Table will not contain micro/nano seconds anymore. Please also be aware that, since Retrospective supports only a millisecond date/time precision, micro/nano seconds will also NOT be displayed in the Date column of the Result Table. This column will contain the date/time formatted in accordance with the Result Options -> Date/Time Format preference configuration. Please note that this preference configuration does not support the "`i`" token which can appear only in the Data Source definition. So in order to lookup micro/nano seconds that are present in given log entry visible in the Result Table, you need to click the entry and examine the Details Pane (5.3 Result Details Pane). The Pane contains the original date string of the entry in an unchanged date format, therefore micro/nano seconds will be available there.

- In the case of "Filter Matching Selection", throughout the analyzed file set, there are differences in encoding, log entry separation or date/time format. Retrospective indicates such configuration mismatches by displaying warning and error icons for individual fields inside the data source table. If the mouse pointer hovers over such an icon, you get a detailed description of the problem and should be able to correct it.

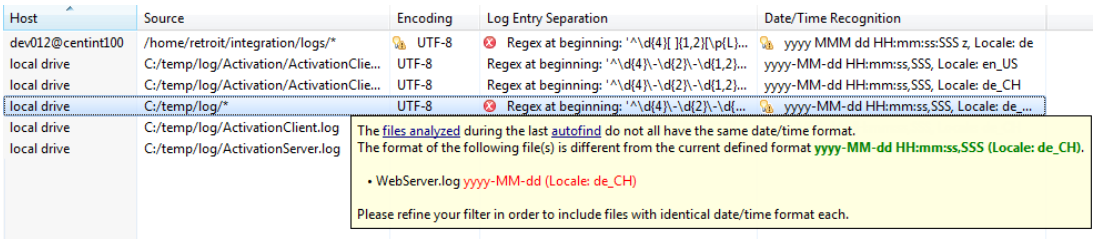

Running a search or monitoring session based on data sources that are marked with such a warning should be avoided. It generally leads into unexpected results.

As long as a data source definition is marked as erroneous, it must not be used for searching or monitoring log data. As explained in a previous section, Retrospective may for example not be able to split individual log entries from certain files and could use up all its available memory.

3.4.3.2 Manual Encoding and Pattern Definition

The file encoding, log entry separation through pattern recognition and date/time format can be manually configured for each data source (individual log file, directory or filter pattern). This is needed if Autofind is not desired or if it did not produce the expected result.

1. Select the desired data profile within the **Profile Manager** or in the data source(s) dialog.
2. Select desired data source and click the [Configure source...] button.
3. Define log entry separation. Select ***begins with*** option to define a regular expression determining the start of each log entry (e.g. a date) or ***ends with*** to define a regular expression determining the end of each log entry (e.g. new line character).
4. Click the [Next] button to proceed to date/time entry recognition.
5. Select locale, provide date format and regular expression.
6. Click the [Finish] button to end the configuration process. Please note that the [Finish] button is activated only if Retrospective does not find an obvious mismatch in the manually specified configuration.

Note

In the Configure source dialog a preview of the initial part (in case of one data source file) or initial parts (in case of multiple data source files) are displayed so the user can easily analyze them and manually set appropriate configuration. However, when a data source contains large number of files then, in order to make it consistent with the Autofind procedure, the preview shows initial parts of no more than the number of files defined in "Maximum number of files to be analyzed" preference in "Data Sources" preference page. In fact, the preview displays initial parts of exactly the same file set as the one analyzed by Autofind for given data source. Thanks to that it can be expected that clicking "auto-find" buttons available in the dialog will yield the same result as invoking Autofind procedure from the data source context menu (however in case of "Filter Matching Selection" there can be some differences, see the note below).

Note

For "Filter Matching Selection" data sources clicking "auto-find" buttons in Configure source dialog can sometimes yield different results in respect to "Log entry separation" and "Date/time format" than invoking Autofind procedure from the data source context menu. In particular when there is some mismatch i.e. one file has different "Log entry separation" (or "Date/time format") than other files, then clicking "auto-find" button will result in setting "Log entry separation" to the default, i.e. ending with newline (in case of "Date/time format", the format will be empty). In Autofind procedure, the result can be different and all mismatches are reported as proper error or warning icons (as described in Step 5 of Section 3.4.3.1 Autofind).

3.4.3.3 Disabling Data Sources

Individual data sources belonging to a profile can be disabled by deselecting the check boxes that appear in front of the table rows. Disabled data sources are shown with gray text color as shown below.

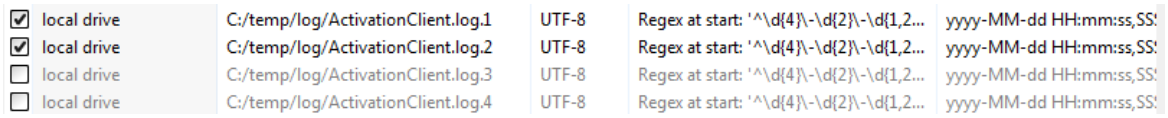

Disabled data sources are ignored during the log data search and monitoring process. If you want them to be included in processing, reselect the appropriate check boxes.

3.4.3.4 Data Source Context Menu

If you press the right mouse button while one or several data sources are selected, a comprehensive context menu pops up and lets you choose the desired contextual function. You can for example clone data source configurations to different hosts and thereby create a complex profile that includes identical data sources from many different remote computers.

Menu Item	Description
Autofind configuration	Starts the Autofind process for the selected data sources (see 3.4.3.1 Autofind)
Clear configuration	If you decide that you have to reconfigure the log file settings, you can clear existing configuration.
Copy configuration Paste configuration	If you have a configuration defined for one of the source files, you can easily apply the exact same settings to other source files. To do so proceed as follows: 1. Right-click the log data source with splitting strategy already defined and select **Copy configuration** from the context menu. 2. Right-click the log data source where you want to have the same configuration applied and select **Paste configuration** from the context menu.

Menu Item	Description
Assign to…	Assigns a different remote host to the selected data sources. Please note that this menu item is available only if the host of the selected data sources is a remote computer.
Clone to…	Clones the selected data sources but changes their hosts to the remote computer selected by the user. Please note that this menu item is available only if the host of the selected data sources is a remote computer.
Configure source	Opens a dialog where you can manually configure the selected data sources within a wizard-style dialog (see 3.4.4 Configure Source(s)).
Change source path	Opens a dialog that lets you change the path of the selected data source. Attention should be paid to the checkbox labeled "Autofind configuration" that appears left of the OK button. If it's selected, Retrospective automatically analyzes the selected data source (i.e. a file) in order to determine its encoding, the log entry separation (delimiter) and the date/time format. If you deselect the "Autofind configuration" checkbox, Retrospective will only change the path of the selected data source but keep the remaining configuration. Therefore, deselect the checkbox only if the data source was renamed but otherwise has remained unchanged.
Delete source(s)	Deletes the selected entries from the data source table
Copy Source	Copies the selected entries from the data source table to the clipboard. This is useful for pasting the identical data source(s) to the corresponding table of another profile. You may also paste it to the data source table of the same profile and adapt it to your needs.
Paste Source	Pastes the data sources from the clipboard (see "Copy Sources" explained above) to the table in which the context menu item is triggered.

3.4.3.5 Deep Wildcards

Retrospective supports **data source paths** and **directory paths** with so-called deep wildcards. **Data source path** contains deep wildcards if the wildcard characters (* and ?) appear in the level of directory. For example, the following path: "/path/*/*.log" contains deep wildcards and the following one: "/path/directory/*.log" does not. In order to explain how deep wildcards in **data source paths** are evaluated to **files** and how deep wildcards in **directory paths** are evaluated to **directories**, see the following example which assumes the following directory structure:

- /path/
  - dir1/
    - dir1file1.log
    - dir1file2.out

- dir2/
  - dir2file1.log
  - dir2file2.out
- dirOther/
  - dirDeep/
    - fileDeep
    - fileDeep2
  - fileOther.out
- file1
- file2

Data source paths are evaluated to files in the following manner:

- /path/*/*.log -> dir1file1.log, dir2file1.log
- /path/*/* -> dir1file1.log, dir1file2.out, dir2file1.log, dir2file2.out, fileOther.out
- /path/dir?/*.out -> dir1file1.out, dir2file1.out
- /path/*/*/* -> fileDeep, fileDeep2 (**not supported**)
- /path/*/*/fileDeep -> fileDeep (**not supported**)
- /path/* -> file1, file2 (**not a deep wildcard**)

Deep wildcards in **data source paths** are supported only on the level of the last directory in the path.

Data source paths with deep wildcards can be specified in Add Data Source(s) dialog or File Browser with the use of "Filter Matching Selection" combined with the "sub-directory button" (see section 3.4.1 Create New Profile).

Directory paths are evaluated to directories in the following manner:

- /path/*/ -> dir1, dir2, dirOther
- /path/*/*/ -> dirDeep
- /path/dir?/ -> dir1, dir2
- /path/dir?/*/ -> (no directory is matching the wildcards)
- /*/*/*/ -> dirDeep

There is no limitation for deep wildcards present in **directory paths**.

Directory paths with deep wildcards can be specified in Add Data Source(s) dialog or File Browser by editing the text box present in the "Select directory" panel (see section 3.4.1 Create New Profile).

Note

When using wildcards in the directory selection text box, the data source selection becomes disabled. Once you select a directory from the result list or remove the wildcard, the data source selection is enabled again.

3.4.4 Configure Sources

A wizard-style dialog appears when you press the 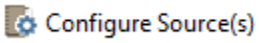

button at the top of the sources table or when you select the corresponding item from the sources context menu. Within the different pages, you can manually adjust the configuration of the selected data sources that was initially detected by **autofind**. Retrospective assists you in these tasks as follows:

1. It loads **sample data** from the selected data sources and displays it at the top of each page. Using the slider right to the data box, you can reduce or increase the amount of such sample data.
2. It displays a **preview** that instantly reflects the definitions you make on the page.

Note

For standard cases, powerful **autofind** is capable of detecting all settings needed to perfectly process your log data, hence you most often don't need to manually adjust the configuration as described in this section.

3.4.4.1 Log Entry Separation (Step 1 of 3)

On this page, you define the encoding of the data source and the pattern that is used for splitting individual log entries. In case the data source is a container, you can also define if Docker or Kubernetes shall add timestamps to its delivered log entries. This can be changed by clicking the toggle button at the top of the page ( or ).

Encoding: Determines how the content of the data source is encoded. In case the loaded sample data or the preview contain strange signs, you probably need to choose a different encoding.

Terminal Output: This checkbox only appears in case the data source is a container (Docker or Kubernetes). It lets you define whether associated log data is issued by terminal output. Such data contains special characters that control the cursor location, color, and other options on the terminal. Select this checkbox if you want to have these special characters be filtered out to produce human readable text. Please note that the detection and removal of special characters may affect performance.

New log entry: Tells Retrospective if the split pattern appears at start or at the end of log entries.

The **split pattern** determines how individual log entries shall be separated. It has the format of a regular expression (regex) and can be defined in different ways.

1. By manually entering a regex.
2. By marking  (dragging cursor along text) a chunk of text from the sample data and let Retrospective generate the regex out of it.
3. By pressing the 'auto-find' button and let Retrospective propose a regex assumed to be suitable.

Examples:

- If your log entries end with new line, you can select **'ends with'** and enter the expression **'\n'**.
- If your log entries start with date, you can select **'starts with'** and enter a date pattern such as **'20\d{2}-\d{2}-\d{2}'**.

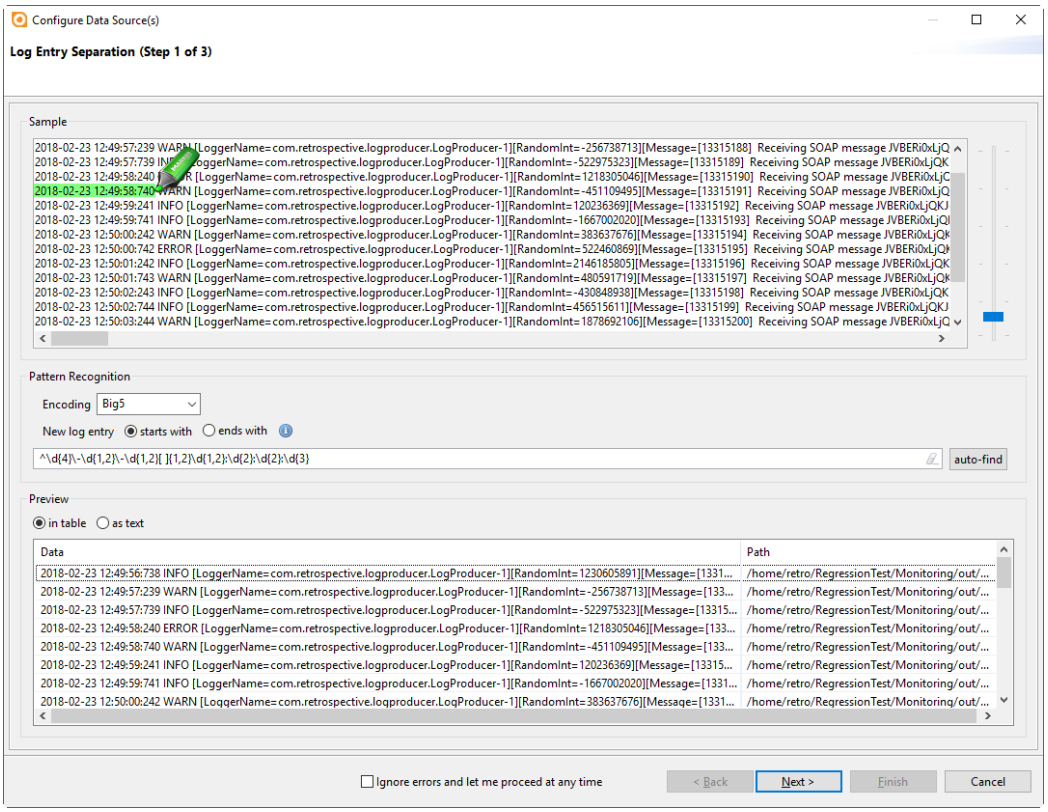

3.4.4.2 Date/Time Recognition (Step 2 of 3)

Locale: This dropdown field defines the expected format and language (e.g. important for month names) of the log entry timestamps.

Date Format: The format of the date/time section (timestamp) to be expected in the log entries is defined on this page. The date/time format is normally automatically recognized during **autofind** but sometimes it needs to be re-adjusted by the user. This can be done in different ways.

- By manually entering the format. Please keep in mind that the format must comply with the Java specification found at https://docs.oracle.com/en/java/javase/17/docs/api/java.base/java/text/SimpleDateFormat.html.
- By marking (dragging cursor along text) the date/time section in the sample data and let Retrospective generate its format. Be careful to select the log entry timestamp but not any date/time otherwise contained in the payload of the log entry.
- By pressing the 'auto-find' button and let Retrospective propose a format assumed to be suitable.

The **regular expression** (regex) is automatically generated from the date/time format and cannot be directly edited by the user

Time zone fallback: This checkbox only appears in case the data source is a container (Docker or Kubernetes). It determines whether the container time zone shall be used as a fallback for converting log entry timestamps that don't contain a time zone themselves. This option has no effect when the container has the same time zone as the local computer. This option should not be enabled if on the previous page you chose to let Docker or Kubernetes add timestamps to delivered log entries of the container.

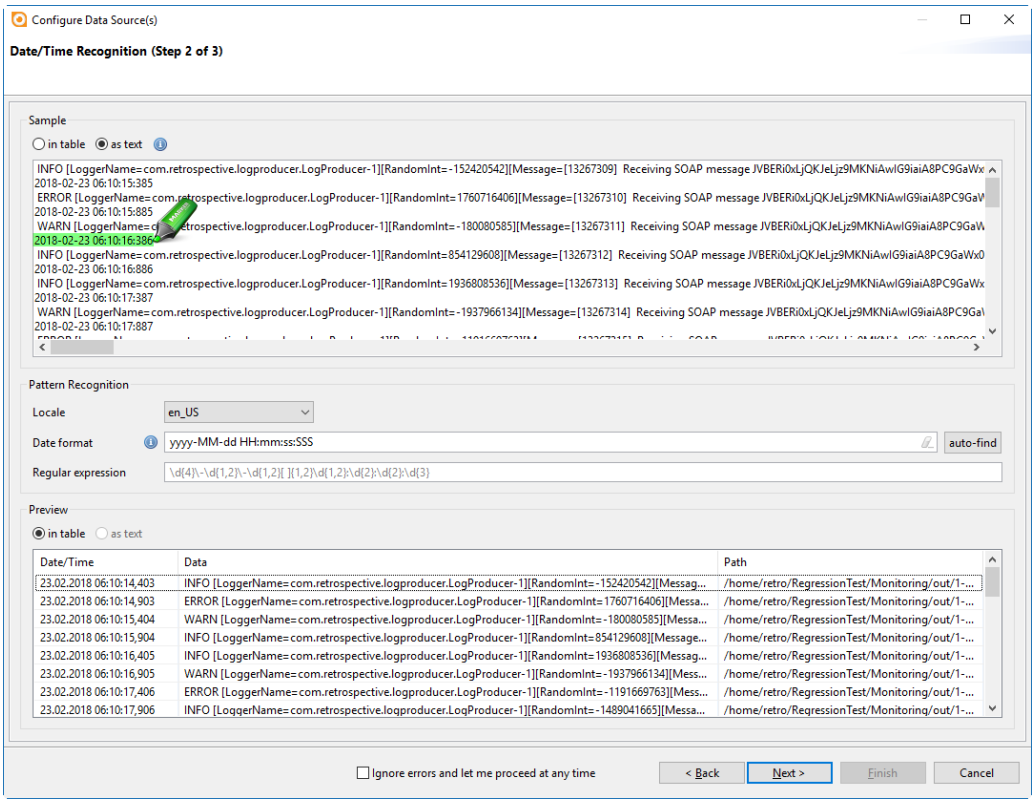
Configure Data Source(s)
Date/Time Recognition (Step 2 of 3)
Sample
in table
as text
Pattern Recognition
Locale
en_US
Date format
yyyy-MM-dd HH:mm:ss:SSS
auto-find
Regular expression
Preview
in table
as text
Date/Time
Data
Path
Ignore errors and let me proceed at any time
< Back
Next >
Finish
Cancel

3.4.4.3 Field Extraction (Step 3 of 3)

This page lets you define custom columns furnished by field values or other text sections extracted from your log data.

Retrospective highlights text in the sample data that matches predefined name/value pattern fields. Such fields are proposed within a pop-up dialog to be chosen when you press one of the "add field…" buttons. Changes are immediately reflected in the preview page. Otherwise this page offers pretty much the same functionality as the columns section within the Profile Manager (see 3.4.6 Define Columns).

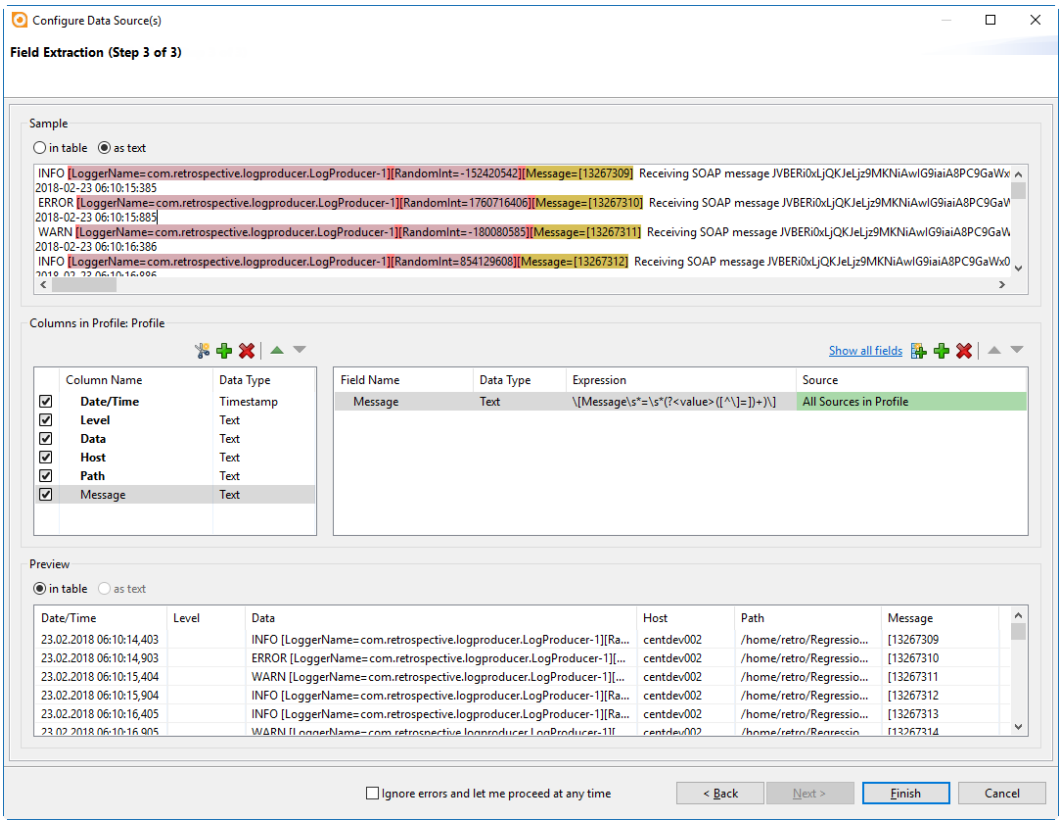

3.4.5 Delete Profile and Data Sources

1. Click the icon to open the ***Profiles*** tab.
2. To delete a data source from the profile containing this particular data source definition, simply select the profile in the ***Profiles*** section, select the data source in the ***Sources in Profile*** section and click the [Delete] button.

   Note

   Deleting a data source is only possible from within the profile which contains this particular data source. Selecting ***All Sources*** option and deleting a particular data source definition is not possible.

3. To delete a data profile, simply select desired profile and click the [Delete] button or right-click given profile and select ***Delete*** from the context menu.

3.4.6 Define Columns

By default, log data found during search and monitoring is displayed in the result table within a predefined set of columns (Date/Time, Level, Data, Host and Path). We call them **standard columns**. The visibility and position of these columns can be changed for individual profiles in the columns definition panel. It's not possible however to add, delete or change standard columns otherwise.

Custom columns on the other hand are entirely managed by the user. They can be added, moved, changed and deleted at will in order to obtain perfectly configured result tables.

3.4.6.1 Columns Table

When you create a new profile, the columns table initially contains all standard columns. Their names appear in bold letters to distinguish them from custom columns that can optionally be defined for individual profiles. The checkbox in front of the column indicates if it shall be visible (be shown) in the result table that is displayed when search or monitoring is performed.

Menu Item	Description
Add split field columns	Adds a variable number of columns each based on a split field. When pressing this button, you'll be prompted for the separator and the number of values you want to have extracted from individual result entries of the desired data source(s). Fields can be split by following types of separators: • Character • Space • Tab • String • Regular Expression 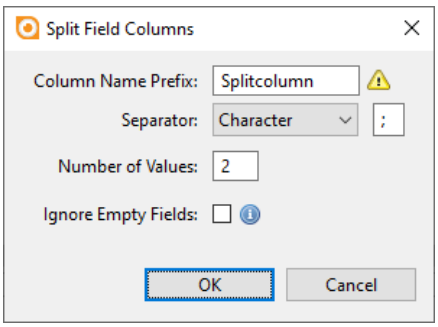 Choose a column name prefix that does not contain any space if you plan to use the columns in RQL queries. When the option "Ignore Empty Fields" is selected, consecutive separators are combined into a single separator. This is helpful in cases where individual fields are, for example, separated by one or many spaces. Retrospective creates as many split columns as needed to extract the entered number of values. Each of these split columns gets an initial split field assigned. The field in the first column starts at split position 1. The separator and the split position of individual fields can be changed at any time within the field table.

Menu Item	Description
Add	Adds a new custom column defined as non-visible as long as no fields are assigned to it.
Delete	Removes the selected custom column from the list. Please note that standard columns cannot be deleted but their visibility can be changed.
Move up	Moves the selected column up by one position
Move down	Moves the selected column down by one position

Standard columns have a single field assigned each. These are non-editable fields and the assignment cannot be altered. Custom column field assignments on the other hand are free of restrictions. Please note however that custom columns are of no use if they do not have at least one field assigned.

3.4.6.2 Fields Table

The fields table appears right to the columns table and shows the fields assigned to the selected column. A field specifies a well delimited chunk of data and the ways to extract them from different log entries.

Depending on the selection in the Source column, a field has a different scope.

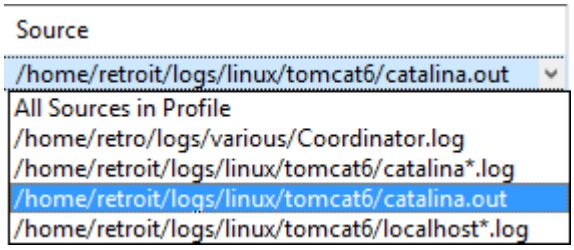

1. It is extracted from result entries of a specific data source defined in the same profile.
2. It is extracted from result entries of all data sources defined in the same profile ("All Sources in Profile").

Menu Item	Description
Add field to new column	Adds a new field and assigns it to a newly created custom column that has the same initial name as the field itself. Through a pop-up menu or a drag and drop operation, the added field can later be assigned to a different existing or new custom column. When this button is pressed, a dialog opens that contains a choice of fields of many possible types (Name/Value Pattern Fields, Regex Fields, Split Fields, Fixed Position Fields – see 2.7 Custom Columns), grouped by different categories. Name/Value Pattern Fields are present in the **Name/Value Pair** category and Regex Fields are present in all other categories. Some of the predefined Regex Fields cover common use cases (e.g. amount with currency, e-mail address, IP address and date/time).

Menu Item	Description

Within the **Other** category, you'll find the following items:

Regular Expression Field	This entry lets you define a custom regular expression for extracting the desired values out of your log data.
Split Field	This item lets you define a split field at a specific position using one of the following types of separators (see 2.7.2 Character Separated Fields): • Character • Space • Tab • String • Regular Expression
Fixed Position Field	This item lets you define a field with a fixed start position and fixed end position (or length).

In case that a previously performed **Autofind** has detected potentially interesting Name/Value Pattern Fields, they appear within the top located "Detected in data" category. Fields detected during search and monitoring with the same profile are also made available within this category. On top of the dialog you'll find a filter field that makes it easy to find the field of your choice.

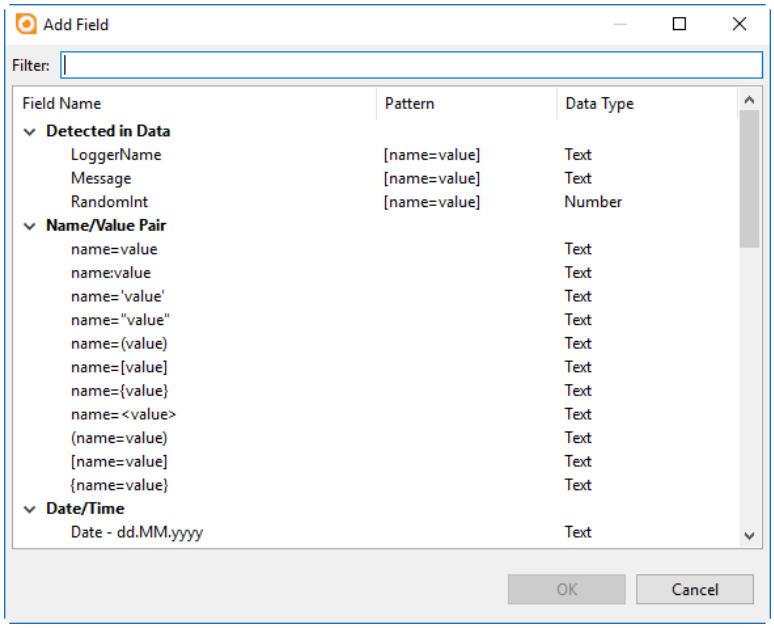

Menu Item	Description
 Add field to selected column	Adds a new field to the current selected custom column. Through a pop-up menu or a drag and drop operation, the added field can later be assigned to a different existing or a new custom column. Otherwise the button behaves the same as "Add field to new column" button described above.
 Delete	Removes the selected field from the custom column. Please note that fields from standard columns cannot be deleted.
 Move up	Moves the selected field up by one position
 Move down	Moves the selected field down by one position

3.5 Import from Previous Version

3.5.1 Retrospective Version Upgrade

When you install a new (higher) version of Retrospective, you can safely overwrite the current installation with the new one. No need to manually export and restore configuration and data files beforehand. These files are stored for each release separately in specific folders within $USER_HOME/.Retrospective.

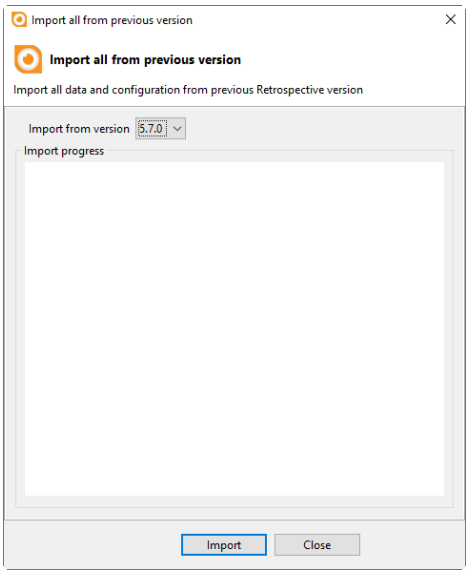

After the installation of the new program, Retrospective lets you import the configuration and data files from a previous compatible version of the program.

When the new program is launched, a dialog pops up and lets you import the configuration and data files through a click on the [Import] button.

Depending on the current and previously installed versions, Retrospective may however only allow the import of configuration data. This is the case when the versions of the underlying database are not compatible between the current and previous Retrospective release.

3.5.2 Through Menu

The configuration data from a previous compatible Retrospective release can be imported at any time as mentioned below. No collected log data (snapshots) will be imported from previous releases.

1. Select File → Import Configuration from previous Version to open the ***Configuration importer*** dialog.
2. Select previous Retrospective version from the ***Import from version*** drop-down list.
3. Click the [Import] button to proceed with the configuration import.

Note

Importing configuration data will overwrite the current settings.

4 PREFERENCES

This chapter describes the Retrospective preferences window that shows different options grouped by individual topics. Here it's possible to customize the behavior and look of the Retrospective program to a certain extent. The ***Preferences*** window is opened by selecting the menu item Select File → Preferences.

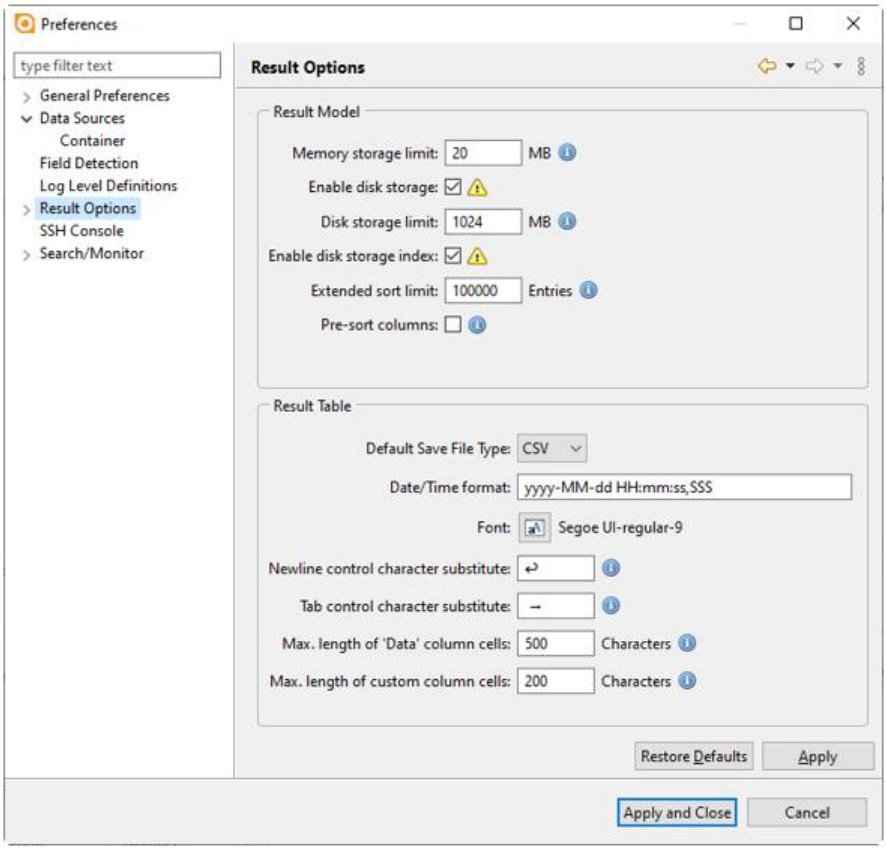

4.1 General Preferences

Option	Description
Show welcome page	Defines whether the welcome page located on the left is to be displayed every time when the Retrospective program is started.
Restore view from last session	Indicates if the window(s) and tab(s) should be restored to the state they had when the last program session was terminated.
Automatically restore result data	Automatically restores previously collected result data when restoring view from last session (see option above). • This only applies for result data collected with enabled disk storage (see 'Result Options' page). • When this option is disabled, you're prompted whether you want to restore result data or not.
Make update check	Tells the program whether it should check for a new available release of the Retrospective each time the program gets started. Please be aware that for the update check Retrospective contacts the program activation server over the internet. If you do not have direct

Option	Description
	access to the activation server, you'll have to define a proxy server (see following topic "Proxy Settings").
Derive search tab name from profile	Instructs Retrospective to automatically derive search tab names from the profile or the ad-hoc data sources they're based on. **Note** Once you manually change the name of a tab, Retrospective will consider that name as granted and no longer automatically change its name. Leave this checkbox unchecked if search tab names shall be of format 'Tab n' where n is any free number starting with 1.
Undo "Don't show again" in dialogs	Shows the dialogs again - in due time - where the option "Don't show again" has been previously selected.
Show prompt at application exit	Determines if the exit confirmation dialog is displayed before the program window is closed.

4.2 General Preferences > Proxy Settings

Option	Description
Enable proxy settings	Enables/disables the proxy server used for program update checks and for the program activation. A proxy server needs to be defined if your organization does not allow direct access to our activation server.
Proxy requires authentication	Authentication details for establishing the connection to the proxy server.

4.3 Data Sources

Option	Description
Ignore hidden files	Determines whether hidden files shall be ignored. • If this option is selected, hidden files are ignored, they won't show up in the UI (i.e. File Browser) and are ignored during search and monitoring. • If this option is unselected, hidden files and directories will be treated the same as any other file or directory.
Maximum number of files to be analyzed	Maximum number of files (or containers log streams) to be analyzed during Autofind procedure (see 3.4.3.1 Autofind) which automatically detects the file encoding and establishes a strategy for log entry separation and date/time recognition. This option is used for "Filter Matching Selection" mode only. • A lower value can be used if you're sure that all files in your selection have the same encoding and if their log entries are each of the same format. This shortens data source configuration time.

Option	Description
	• A higher value increases the chance to discover data sources of different encoding or with different internal data format. This will help in defining individual data sources for different groups of files. This option also limits the maximum number of files whose initial part is displayed for manual analysis in the Configure source dialog opened for "Filter Matching Selection" data source.
Alternate decoding character set	Character set to be used to decode the ZIP file entry names in case decoding with the default UTF-8 character set is not successful.

4.4 Data Sources > Container

Option	Description
Docker CLI	Instructs Retrospective where to find the **docker** command line interface (CLI) – or docker.exe on Windows - required to interact with the Docker subsystem. Docker is a tool designed to make it easier to create, deploy, and run applications by using containers. In order to discover Docker components and to fetch log data, Docker CLI must be found locally. Please add its directory to the 'PATH' ('Path' on Windows) environment variable or explicitly specify its location in the corresponding text field on this page.
Kubernetes CLI	Instructs Retrospective where to find the **kubectl** command line interface (CLI) – or kubectl.exe on Windows – needed to interact with the Kubernetes subsystem. Kubernetes is a system for automating deployment, scaling and management of containerized applications. In order to discover Kubernetes components and to fetch log data, kubectl CLI must be found locally. Please add its directory to the 'PATH' ('Path' on Windows) environment variable or explicitly specify its location in the corresponding text field on this page.
Label Selectors: Include in directories history	Determines if label selectors shall be included in the history that provides quick access to previous selected directories.
Log Entry Timestamps	*When a container is started or stopped, Docker usually produces log data of different format than the application it contains. Therefore, adding a timestamp to each delivered log entry may be helpful to allow Retrospective to accurately split and sort the log entries.* *Note that the added timestamp is produced by the Docker daemon. It may differ from the date/time logged by the application running in the Docker container. This is due to the fact that containers have their own time zone that is often different from that of the Docker daemon. For most popular Docker images, this is UTC.*

Option	Description
Let Docker add timestamps	If this checkbox is selected, Docker is instructed to add a timestamp to each delivered log entry. This is done by invoking the `docker logs` command with the option `--timestamps`. This is a default option that can be changed for individual data sources in the 'Add Data Source(s)' and the 'Configure Data Source(s)' dialog. A toggle button appears on them and lets you enable or disable Docker timestamps for the selected or configured data sources.
Docker time zone fallback	Determines if the container time zone shall be used as a fallback for converting log entry timestamps that don't contain a time zone themselves. This option has no effect when the container has the same time zone as the local computer. Note that selecting this option and option 'Let Docker add timestamps' at the same time may lead to an unexpected result. This is a default option that can be changed for individual data sources in the 'Configure Data Source(s)' dialog within the 'Date/Time Recognition' page.
Let Kubernetes add timestamps	If this checkbox is selected, Kubernetes is instructed to add a timestamp to each delivered log entry. This is done by invoking the `kubectl logs` command with the option `--timestamps`. This is a default option that can be changed for individual data sources in the 'Add Data Source(s)' and the 'Configure Data Source(s)' dialog. A toggle button appears on them and lets you enable or disable Kubernetes timestamps for the selected or configured data sources.
Kubernetes time zone fallback	Determines if the container time zone shall be used as a fallback for converting log entry timestamps that don't contain a time zone themselves. This option has no effect when the container has the same time zone as the local computer. Note that selecting this option and option 'Let Kubernetes add timestamps' at the same time may lead to an unexpected result. This is a default option that can be changed for individual data sources in the 'Configure Data Source(s)' dialog within the 'Date/Time Recognition' page.

Note

These are default options, so changing any of them will leave existing container data source configurations unchanged.

4.5 Field Detection

Option	Description
Spaces around separator	These options define the maximum number of spaces that may appear before and after the separator of name/value pattern fields. Should the settings made in this section also be used by the 'Name/Value Pair' Highlighting Set, apply the changes and switch to page 'Result Options > Highlighting Sets', then press the 'Restore Defaults' button. **Note** Restoring 'Highlighting Sets' defaults will overwrite your manual changes. Therefore, make sure to first export customized Highlighting Sets, then import them again once you've pressed the 'Restore Defaults' button.

Detection during search/monitoring	The detection of new name/value pattern fields is always performed during Autofind. This option determines if such fields shall also be detected during search and monitoring. Detected fields may then be assigned to profile specific custom columns.
	
	This is a default option that can be changed at any time on individual search/monitoring tabs by enabling or disabling the toggle button that appears above the start button.
	
	WARNING: Enabling field detection may heavily slow down the search and monitoring process. Depending on the data volume to be processed, it may even deteriorate the UI responsiveness. You may remedy the latter issue through the following measures.
	• Local Files: Decrease the option value "Max. threads for local file search/monitoring" on the same preferences page.
	• Containers: Decrease the following option values on the same preferences page.
	  ○ "Max. threads for search/monitoring in Docker"
	  ○ "Max. threads for search/monitoring in Kubernetes"
	
	Remote files accessed through SSH: Decrease the option value "Maximum connections" on the SSH host in the advance tab within the Host Manager.

4.6 Log Level Definitions

Option	Description
	This page lets you define individual log level definitions. Every such definition specifies how individual log entries shall be interpreted by Retrospective in order to find/extract the appropriate log level. It also contains information on how log entries with an assigned log level shall be displayed in the result table.

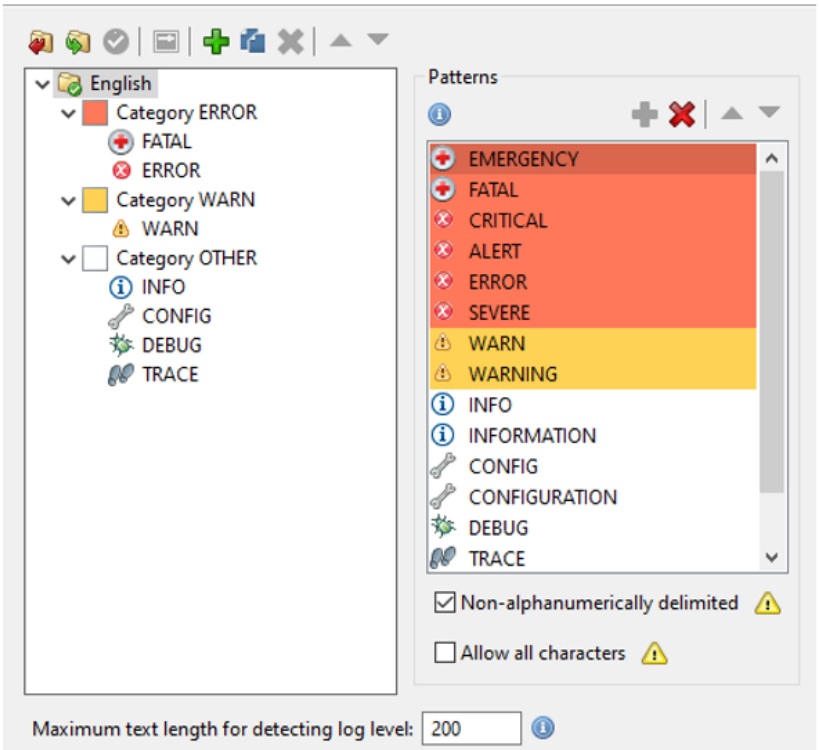

Log level definitions are composed of the hierarchy levels Configuration > Category > Level > Pattern.

Configuration

Root node of the tree-like structure.

- The configuration marked as default is currently in use by Retrospective. All other configurations are ready to be used upon user request.
- The selected log level configuration can be exported to a file for data backup or for being shared with other team members that would import them into their instance of Retrospective within this same preferences dialog.

Every log level configuration contains the following three categories.

- ERROR
- WARN
- OTHER

Option	Description
Category	The category groups types and defines the background color of log entries shown in the result table. They are also used to group log levels in the History Chart (horizontal chart that appears on top of the result table).
Level	Defines the log level name, icon and priority that are used to represent and filter corresponding log entries in the result table. A log level can be identified through one or several patterns. Log levels each have a different **priority**, which depends on their overall position. The priority is also referred to as **severity**. • You can change the priority of a log level by moving it up or down by one or several positions. • The level's icon can be changed at any time by pressing the button and choosing the preferred icon from the appearing dialog. **Note** The position of log levels matters when you locally filter result entries by the log level using an RQL query. If you use a Level column term together with a bounding operator (<, <=, >, >=) such as `Level>=fine`, Retrospective identifies matching result entries by checking the priority of their log level.
Pattern	**Standard Patterns** A case-insensitive word - a pattern - must be present in a log entry in order to match a specific log level. A valid pattern contains digits or uppercase letters only, it must not contain blanks or other special characters. **Note** The patterns 'INFO' and 'info', for example, are identical and must not appear together within the same log level configuration. During search and monitoring, collected log entries are analyzed in order to find a matching log level pattern and to identify the corresponding log level. **Enhanced Options** These options are valid for all patterns defined within the selected log level configuration. ☐ Non-alphanumerically delimited Deselect this checkbox if during log level detection, a text matching a given pattern shall be accepted regardless the type of surrounding characters. WARNING: Disabling this option can decrease the performance of log level detection and thus slow down the search process. ☐ Allow all characters Allows patterns that contain all kind of characters and even spaces.

Option	Description
	WARNING: Enabling this option can decrease the performance of log level detection and thus slow down the search process.
Maximum text length for detecting log level	Maximum text length to be analyzed per log entry for finding its log level. In order to determine the log level of individual log entries, Retrospective will analyze the text from the beginning up to the defined number of characters. • A small number of characters guarantees a low log entry processing time but prevents detection of log levels located beyond the analyzed text section. • A high number of characters ensures that log levels are detected even if they appear far from the text beginning. This may however have a negative impact on the log entry processing time in cases of high throughput.

4.7 Result Options

Option	Description
Memory storage limit	This option defines the in-memory storage limit in MB for individual search/monitor tabs. The value instructs Retrospective how much data (raw result data in MB) is allowed to be kept in memory for individual result tables. If this limit is exceeded, the program removes the oldest entries from the in-memory storage. Depending on the option "Enable disk storage", removed entries are lost or they remain available in disk storage and are reloaded into the memory when required (i.e. when scrolling to a different section of the result table). Every character of log data stored internally requires 2 bytes of memory and additional memory is allocated for indexing the data. Further items created during data processing also require additional memory to be allocated. Therefore, the memory effectively allocated for each search/memory tab may easily exceed twice the size specified by the `Memory storage limit` option. **Note** Retrospective by default uses up to 1GB of memory (maximum memory heap size). When several search tabs are opened, it is advised to set the `Memory storage limit` parameter to a lower value. In exceptional cases where the user still needs to use more memory per tab, the `-Xmx` attribute in `<Retrospective Installation Directory>/retrospective.ini` may temporarily be increased. Changing this parameter requires restarting Retrospective to take effect. Be aware that defining a high `Memory storage limit` may lead to out-of-memory errors, especially when working with multiple search/monitor tabs opened at the same time.
Enable disk storage	Tells the program if by default it should store result data on the hard drive when the in-memory storage limit is exceeded. • If this check box is selected, you can safely set above described `Memory storage limit` option to its default value of 20 MB. This will prevent the program from allocating a huge amount of memory. • If this box remains unchecked, Retrospective simply deletes the oldest entries from the result table when the in-memory storage limit is exceeded. Please note that a change of this value applies only to new search/monitoring sessions i.e. it has no influence on ongoing search/monitoring sessions.
Disk storage limit	Defines how much hard drive space can be used for storing search results. When this limit is exceeded, the oldest entries are deleted from the data store.
Enable disk storage index	If checked, results are indexed immediately after the search is finished. Please note that if you deactivate this option, sorting and local filtering will be applied only to data in memory and not to the entire result data.

Option	Description
Extended sort limit	Maximum number of result entries that may be present in order to have extended (smart) sorting applied (see 5.2.6 Extended (smart) Sorting). Extended sorting considers several visible and hidden columns beside the main sort column, and thus produces a result of higher accuracy. • if your result table contains MORE log entries than this value, then extended sort is NOT applied • if your result table contains LESS EQUAL number of log entries than this value, then extended sorting is applied. **Note** A high value may significantly increase the time required for sorting, especially when working with a high amount of data.
Pre-sort columns	If checked, visible columns (except the Data column) are pre-sorted for later use straight after the search/monitoring process is completed. This increases the result initialization time but ensures that changes in sorting of visible columns are processed more quickly. In case you expect huge volumes of result data, it is advisable to leave this box unchecked.
Default Save File Type	The default selection of the drop-down button that lets you save collected data to a file of a chosen type. This drop-down appears on the toolbar of the result table. The default file type is applied on every new search tab.
Date/Time format	Format to be used for representing Date/Time cells in the result table. The format must be compliant with the Java specification found at http://docs.oracle.com/javase/8/docs/api/java/text/SimpleDateFormat.html.
Font	Font to be used for representing the data in the result table.
Newline control character substitute	String to be shown in place of newline control characters (line break) inside 'Data' and custom column table cells. This option is useful to see more of your log entries inside the single line table row in case they contain newline control characters.
Tab control character substitute	String to be shown in place of tab control characters inside 'Data' and custom column table cells.
Max. length of 'Data' column cells	Maximum number of characters to be contained in individual cells of the 'Data' column. If the length of the data exceeds this number, it will be abbreviated using ellipses (...). Choose a low value, but large enough to see as much data as possible in the result table. **Note** If this value is too large, it can have a slight negative impact on memory consumption and performance in case your log entries are huge.

Option	Description
Max. length of custom column cells	Maximum number of characters to be contained in individual cells of custom columns. If the length of the cell content exceeds this number, it will be abbreviated using ellipses (...). Choose a low value, but large enough to see as much data as possible in the result table. **Note** If this value is too large, it can have a slight negative impact on memory consumption and performance in case your log entries are huge.

4.8 Result Options > Formatting

Option	Description
Word wrap	Indicates if long lines in the result detail view should be wrapped or not. This option can be changed for individual search tabs.
Tab size	Tab width measured in characters for data displayed in the result detail view.
Data Formatting	Defines the default formatting of the data shown in the result detail view. The default data formatting can be changed for individual search tabs.
JSON Indentation	Indentation for formatted JSON data.
XML Highlight syntax	Enables/disables XML syntax highlighting in the result detail view.
XML Pretty print	Enables/disables the formatting of XML data in the result detail view.
Replace XML entities	For replacing Retrospective XML character entities such as `'<'` and `'>'` with `'<'` and `'>'` respectively. **Note** Replacing such entities in a text that is embedded in an XML document can disrupt the document's well-formedness, thus making pretty print and syntax highlighting impossible.
XML indentation	Indentation for formatted XML data.

4.9 Result Options > Highlighting Sets

Option	Description
	This page lets you define highlighting sets that can be applied on any log entry result table in order to see terms or sections of interest highlighted by different colors. Each highlighting set (collection) contains one or many highlighting matchers.

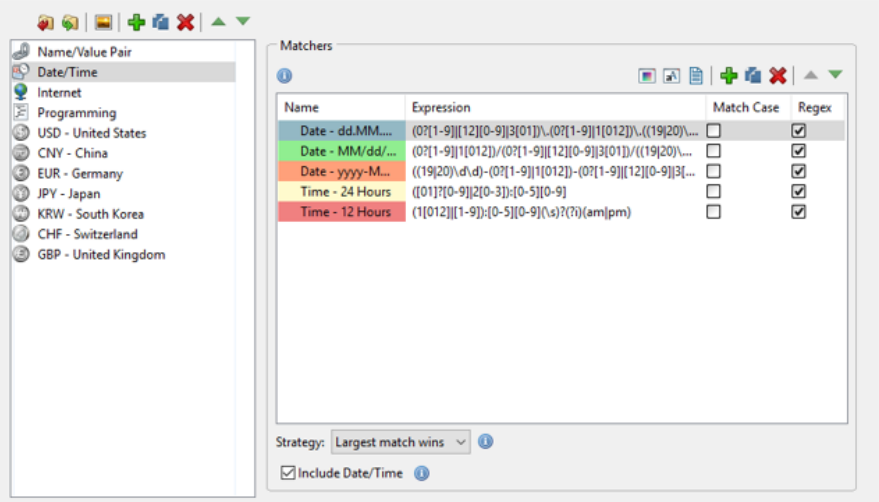

Retrospective offers a few predefined (default) highlighting sets that illustrate how single highlighting matchers may be defined. Therein contained matchers are samples and not optimized for correctness nor for performance in specific cases. It's up to you to adapt them or create new ones that perfectly fit your needs.

Highlighting Set

A collection of distinct highlighting matchers ready to be chosen on the fly for highlighting terms and sections of interest within the log entry result table.

- Single or multiple selected highlighting sets can be exported to a file for data backup or for being shared with other team members that would import them into their instance of Retrospective within this same preferences dialog.
- The highlighting set icon is useful for quickly identifying the set and the type of matchers it contains. The icon can be changed at any time by pressing the button and choosing the preferred icon from the appearing dialog.
- The position of the selected highlighting set can be changed through the move-up (▲) and move-down (▼) buttons. Available highlighting sets will be presented in the defined order within the menu that lets you choose the highlighting of the current log entry result table.

Matcher

A matcher defines an expression (text or regular expression) for finding a phrase of interest and the color that phrase shall be highlighted with.

Option	Description
	• When adding a new matcher, a default background color is applied. You can select a different color at any time by pressing the  button. • The highlighting font and text color can be adjusted within a dedicated dialog when you press the  button. • Matchers may be decorated with a description. A corresponding text field appears when the  button is pressed. The description is shown within the matcher tooltip that appears when the mouse pointer hovers over the table cells. • Matchers within the same highlighting set are evaluated according to the selected strategy (see below). The winning matcher within an individual log entry determines the highlighting of the entire result entry table row. Within the result entry detail view however, text portions are highlighted by specific colors taking into account all defined matchers (not only the winning matcher chosen according to the strategy). The position of the selected matcher within given highlighting set can be changed through the move-up  and move-down  buttons. WARNING: For displaying the result entry table, depending on the selected strategy, all matchers within the selected highlighting set may be evaluated for every result entry until a match is found. Depending on the matcher expressions, this may drastically slow down the Retrospective search and monitoring process. Therefore, you should avoid overstuffing highlighting sets with matchers that are not absolutely needed in the current context. It's good practice to create multiple different highlighting sets with a few matchers for a specific purpose and to switch between them on the fly. **This principle is also valid for most default highlighting sets** shipped with Retrospective. Remove unnecessary matchers from them or split them into distinct highlighting sets. **Text Matcher** When the "Regex" checkbox is **NOT** selected, the matcher expression represents simple text. In such cases you may use one or several **wildcard characters** to locate a text part to be highlighted without knowing the exact phrase you're looking for. Wildcards are special characters that represent unknown characters in a text. They're handy for locating similar but not identical sections. You can use them anywhere within the matcher expression. • The **asterisk (*)** matches any number of characters. • The **question mark (?)** matches a single character at a specific position. Since wildcard characters are interpreted in a special manner, you have to quote them with backslash (\) if you want to match them explicitly. • If you want to match **asterisk (*)**, you have to type ***** • If you want to match **question mark (?)**, you have to type **\?** • If you want to match **backslash (\)**, you have to type **** (please note, that using a single backslash that is not preceding another backslash, an asterisk or a question mark result in a syntax error)

Option	Description
	Regex Matcher When the "Regex" checkbox is selected, the matcher expression represents a regular expression. WARNING: be careful when defining regular expressions since wrongly defined regex can be very expensive (time consuming) and thus slow down the entire Retrospective search and monitoring process. It's preferable to create short and fast performing regex matchers for specific cases grouped in different highlighting sets than creating fancy, generic and slow regex matchers. *Please consult http://docs.oracle.com/javase/8/docs/api/java/util/regex/Pattern.html and check the summary of **regular expression** constructs or learn more about regular expressions by reading one of the many publically available reference pages.*
Strategy	Defines how different matchers of the highlighting set are evaluated and applied to individual result table rows. The following strategies are available. **First match wins** • Matchers are evaluated top-down, the first match within an individual result entry determines the highlighting of the table row. **Largest match wins** • The matcher with the largest cumulated size of conforming text sections determines the highlighting of the table row. **Note** The highlighting within the result entry detail view does not depend on the strategy. Within this component, search criteria, custom columns, local text filters and matchers are all simultaneously highlighted and overlapping sections are merged.
Include Date/Time	Determines if highlighting shall include the Date/Time part of table result entries. Please note that this option is valid for all matchers defined within the same highlighting set.

4.10 Result Options > Koia Data Analysis

Option	Description
	By pressing the  button located far right on the result table toolbar, result data can be exported to a CouchDB database and be further analyzed in a **Koia web-application**. To be able to do so, you need to install and configure CouchDB on your computer or a computer in a trusted environment according to the instruction found at https://github.com/centeractive/koia/blob/master/README.md#installing .
CouchDB	Defines how Retrospective can access your CouchDB.

Protocol	Hypertext transfer protocol to be used (HTTP or HTTPS).
	Note
	HTTPS communication between Retrospective and CouchDB requires the following settings:
	• CouchDB must be set up to work with SSL/TLS as described in section "HTTPS (SSL/TLS) Options" of the CouchDB documentation.
	• The user certificate must be imported into the Java 'cacerts' truststore. This can be done for example with the Java Keytool as follows:
	`$ cd $RETROSPECTIVE_HOME/jre/lib/security`
	`$ ../../bin/keytool -importcert -trustcacerts -file $COUCHDB_HOME/cert/couchdb.pem -keystore cacerts -alias "couchdb"`

Host	The host where CouchDB is installed. In the beginning, we recommend installing CouchDB on your local computer and to use the default value "localhost" or "127.0.0.1". Once this is working well, a more sophisticated setup may be chosen.
Port	Port to be used to access CouchDB over HTTP (Default value is "5984").
User	The name of the admin user you defined when performing the CouchDB "Single-Node-Setup".
Password	The password of the admin user you defined when performing the CouchDB "Single-Node-Setup".
Koia URL	The URL of the Koia web application. The official Koia URL is https://www.koia.io but you may also choose to install Koia on a server within a trusted environment. **Note** When using the official site https://www.koia.io or any other site accessed through HTTPS, your CoachDB **must** be installed on the local computer. Otherwise, your browser will block the access to CouchDB and complain about an insecure request made from a page that was loaded over HTTPS.

4.11 Result Options > Local Filter

Option	Description
Text Filter Type	**Retrospective Query Langauge (RQL)** The entered filter text is interpreted as an RQL query. **Classic Contains Filter** The entered filter text is used as is to match result entries. The text may contain spaces and the wildcards * and ?. **Note** If this option is changed, it becomes valid for new search tabs only.
Filter apply delay	Time in milliseconds that must be observed after the user stops typing before the entered filter is automatically applied. **Note** If this option is changed, it becomes valid for new search tabs only.
Immediately apply recovered filter	If this checkbox is checked, filters selected from the history drop down are immediately applied. If this checkbox is unchecked, such filters are executed only once the user presses the <ENTER> key.
Hide filtered-out entries	This option is used to display or hide results entries that match the local filter criteria. If this box is checked, only entries (rows) that match the local filter criteria are shown in the result table. Non-matching rows are no longer visible in the table. If the box is unchecked, all rows are still shown in the table and those that match the search criteria are highlighted (green background color). **Note** This is a default setting that can be changed at any time on individual search tabs.

4.12 SSH Console

Option	Description
SSH Console	
Line Buffer Size	Number of lines to be kept in the SSH Console's internal buffer.
Cursor Color	Let's you select the cursor color within a color picker dialog.
Show Blinking Cursor	Defines whether the cursor should be shown in blinking mode or not.
Enable Keyboard Shortcuts	The following keyboard shortcuts can be enabled: ***Windows/Linux*** / ***Mac OS*** Ctrl + A / ⌘ + A — Select all Ctrl + C / ⌘ + C — Copy selected text to the clipboard Ctrl + V / ⌘ + V — Paste text from the clipboard to the console Ctrl + W / ⌘ + W — Close tab
Quick Copy/Paste Mode	Quick copy/paste mode lets you: • Copy selected text to the clipboard through a left mouse button click on it. • Paste text from the clipboard to the console through a left mouse click at unselected position. Without quick copy/paste mode enabled, you're offered the same functionality through a context menu that appears when the right mouse button is activated.
Background Color	Let's you select the SSH Console background color within a color picker dialog.
Text Color	Let's you select the SSH Console text color within a color picker dialog.
Font	Let's you select the SSH Console text font within a font picker dialog.

4.13 Search/Monitor

Option	Description
Show time search criteria by default	Controls if a time search criteria is automatically added to the search criteria panel when a new search tab is opened.
Date/Time format	Format to be used for representing date/time fields in the search criteria panel. The format must be compliant with the Java specification found at http://docs.oracle.com/javase/8/docs/api/java/text/SimpleDateFormat.html.
Consider file modification date	Determines if the file modification date is to be considered during a search that is based on date/time criteria. If this option is enabled, Retrospective excludes files from the search process if the modification date is earlier then the specified time search criteria. For example, if the criteria specifies range from year 2013 to year 2014 and the file modification date is in year 2012 then the file is ignored since it cannot possibly contain matching entries.
Enable time range optimization	Informs Retrospective whether it should perform initial date checks and, in case of remote files, a narrow-down procedure, before searching log files (see 5.8.2 Search Optimization).
File start offset	Amount of existing data in bytes to be delivered when file monitoring is started. This value corresponds to the -c option of the Linux tail command. Enter 0 if monitoring should start without delivering any existing data.
Container start offset	Number of existing log lines to be delivered when container monitoring is started. This value corresponds to the -n option of the Linux tail command. Enter 0 if monitoring should start without delivering any existing log lines.
Enable dynamic discovery	Determines if Retrospective shall fetch data from files and containers that did not exist when the monitoring process was started. • If selected, new appearing files and containers are dynamically included in the ongoing monitoring process. • If not selected, new appearing files and containers are ignored. This is a default option that can be changed at any time on individual search/monitoring tabs by enabling or disabling the toggle button that appears above the start button.
and allow Autofind	Determines if **Autofind** (see 3.4.3.1 Autofind) shall be run for dynamically discovered files and containers in case the defined entry separation (newline control character, date/time etc.) cannot be found in the initial part of the log data.

Option	Description
	Note In case you had to manually adjust the log entry separation of a data source, you most probably don't want to allow **Autofind** to be run on dynamically discovered files and containers. Otherwise the manually changed configuration would simply be ignored. This is a default option that can be changed at any time on individual search/monitoring tabs by enabling or disabling the toggle button that appears above the start button.
Show Warnings for Compressed Files	Specifies if the program should show a warning for every compressed file that cannot be monitored. Please be aware that compressed files cannot be monitored at all. If this check box is selected, a warning is shown for every compressed file that is omitted by the monitoring process.

4.14 Search/Monitor > Log Entry Separation

Option	Description
Max size	The maximal size of accumulated unsplit content. If this size is exceeded, a warning is displayed and the search or monitoring process is aborted. This may occur if individual log entries are really large, for instance if they carry JSON or XML payload. • Then you need to increase this option and restart the search or monitoring process. • Alternatively you may have to adjust the log entry separation.
Max. size at beginning when monitoring	The maximal size of accumulated unsplit content at the beginning of a data source when monitoring. If this size is exceeded, a warning is displayed and the monitoring process is aborted. This may occur if the log entry separation was misconfigured. • Then you need to adjust the log entry separation and restart the monitoring process. • In case your log entries are large, you may need to increase this option.

4.15 Search/Monitor > System Resources

Option	Description
Max. threads for local file search/ monitoring	Maximum number of simultaneous threads that can be used for performing search and monitoring files on the local computer. When the value is changed, then it has no immediate effect on the thread number. Simply a new thread pool is created and the maximum thread number in this pool is appropriately restricted. The new pool starts to be used just after the preference option is saved. Then, threads from the old pool gradually expire. WARNING: the lower the maximum thread number the higher the probability of encountering a starvation situation. Starvation situation happens when all threads are consumed. For example, if maximum threads is set to 1 and a long-term search on a big file is started, then any other activities such as searching or monitoring will idly wait until the search causing the starvation finishes. The user will be periodically warned about starving activities.
ax. threads for search/ monitoring in Docker	Maximum number of simultaneous threads that can be used for performing search and monitoring containers from the Docker subsystem. When the value is changed, then it has no immediate effect on the thread number. Simply a new thread pool is created and the maximum thread number in this pool is appropriately restricted. The new pool starts to be used just after the preference option is saved. Then, threads from the old pool gradually expire. WARNING: the lower the maximum thread number the higher the probability of encountering a starvation situation. Starvation situation happens when all threads are consumed. For example, if maximum threads is set to 1 and a long-term search on a big file is started, then any other activities such as searching or monitoring will idly wait until the search causing the starvation finishes. The user will be periodically warned about starving activities.
Max. threads for search/ monitoring in Kubernetes	Maximum number of simultaneous threads that can be used for performing search and monitoring containers from the Kubernetes subsystem. When the value is changed, then it has no immediate effect on the thread number. Simply a new thread pool is created and the maximum thread number in this pool is appropriately restricted. The new pool starts to be used just after the preference option is saved. Then, threads from the old pool gradually expire. WARNING: the lower the maximum thread number the higher the probability of encountering a starvation situation. Starvation situation happens when all threads are consumed. For example, if maximum threads is set to 1 and a long-term search on a big file is started, then any other activities such as searching or monitoring will idly wait until the search causing the starvation finishes. The user will be periodically warned about starving activities.

Option	Description
Starvation warning interval	Interval in seconds between warnings that inform the user about a resource starvation situation. • A thread starvation occurs when all threads from the respective pool are consumed (in use). A connection starvation situation happens when all permitted connections to a given host are in use.

5 SEARCHING AND MONITORING

This chapter covers search and monitoring related topics, which is basically the main reason for which Retrospective was built. It's all about selecting data sources (log files), defining search criteria, performing search and monitoring actions. Finally, it explains how to deal with the result data.

The ***Search*** tab lets you select data sources (log files), define search criteria (see 5.1.2 Defining Search Criteria), analyze and tailor search results. This view can be opened by clicking the button or by using the [Ctrl] + [T] keys combination.

Note

Retrospective lets you work with a **maximum** of **10** search/monitoring tabs at the same time.

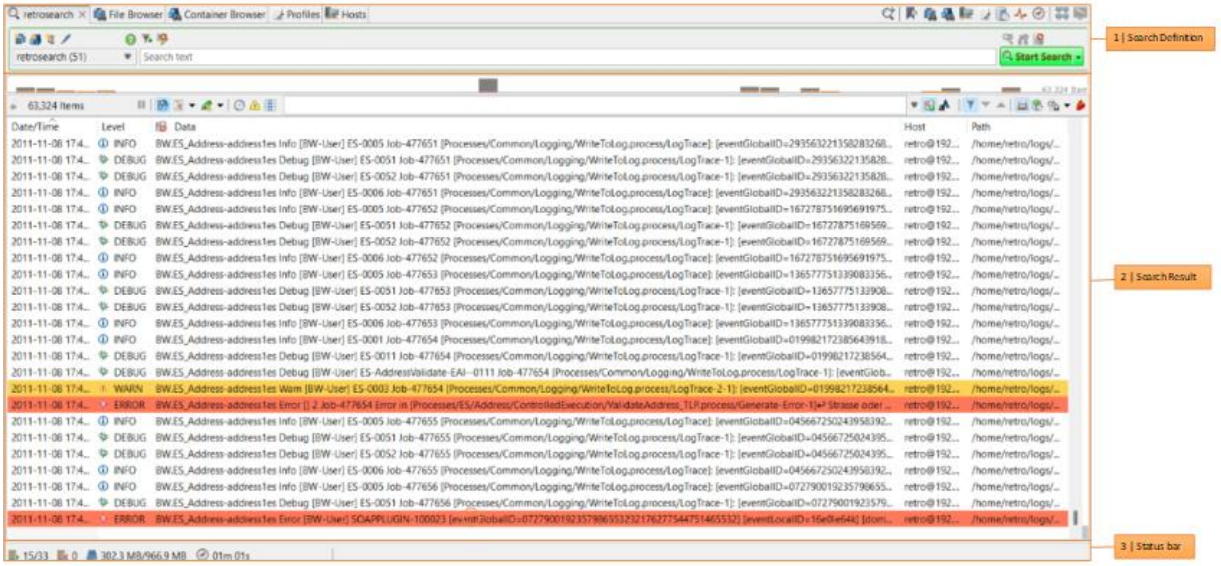

1 The ***Search Definition*** area contains fields that let you specify the data sources and the search/monitoring options for precise search definition.

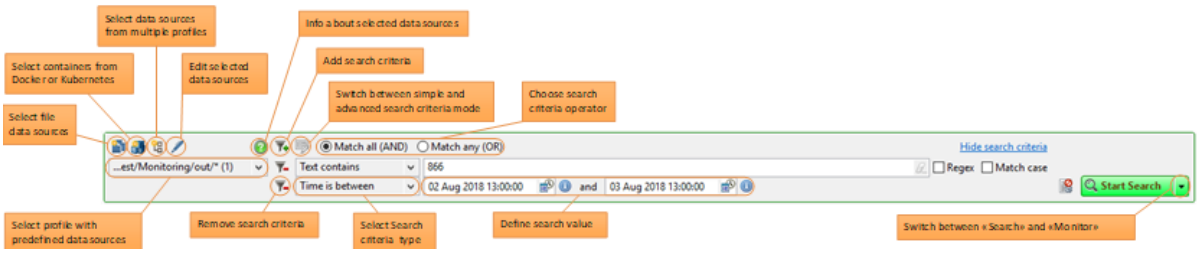

2 The ***Search results*** area contains the result entries and offers functionality for saving, highlighting and filtering.

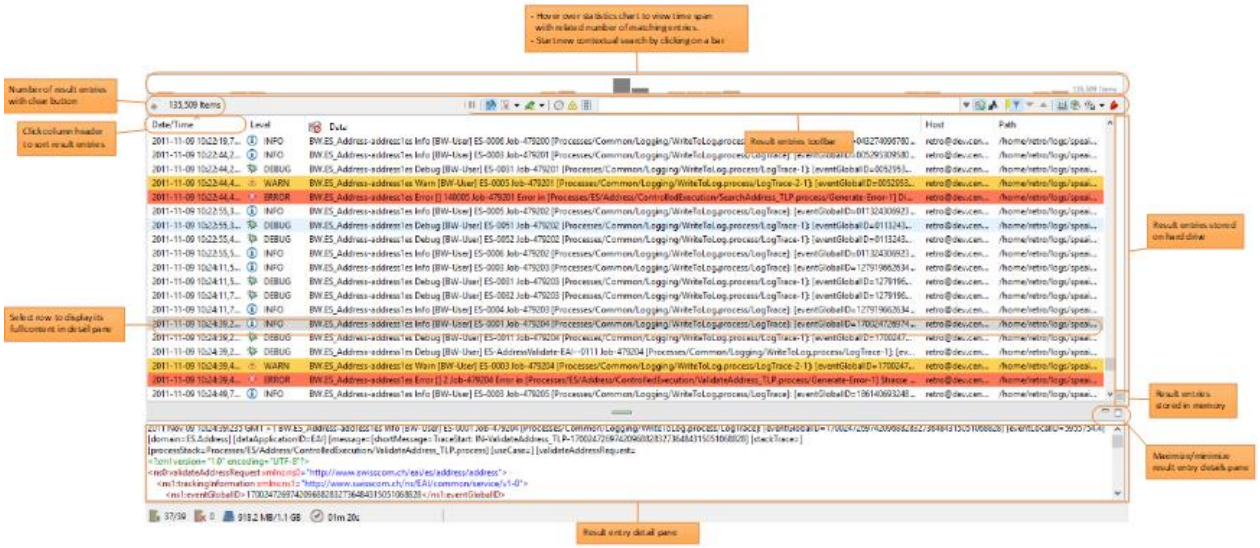

3 The ***Status bar*** area contains information on current search, running background processes and various notifications.

5.1 Search Definition

5.1.1 Search Definition Toolbar

The left located drop-down list on the search definition panel contains all available data profiles – and on-the-fly choices of data units – on which you may base your search or monitoring session. Below listed buttons help you to define the search for your needs.

Icon/Control	Description
	Opens a dialog from where you can choose **file** data sources you want to search log data from. Optionally your choice can be saved to a profile that will be available in future sessions as well.
	Opens a dialog from where you can choose **container** data sources you want to search log data from. Optionally your choice can be saved to a profile that will be available in future sessions as well.
	Open a dialog that lets you choose data sources from various existing profiles. These can either be used on the fly or be saved as a new profile that will be available in future sessions as well.
	This button lets you edit the current selected data sources. • If the selected entry in the drop-down list is a permanent profile, Retrospective switches to the profile manager and selects that same profile. That's where you can adapt the profile by adding new data units, modify existing ones or remove the ones that are not used anymore. • If the selected entry in the drop-down list is not a permanent profile, a dialog appears that shows all current chosen data units. Here you can add new data units, modify existing ones or remove obsolete ones.

Icon/Control	Description
	Placing the cursor on this icon shows a tool-tip with information about the selected permanent profile or the on-the-fly choice of data sources.
	Adds a new search criteria.
	Switches between the **simple** and **advanced** search criteria mode. In the default **simple** mode, the search definition panel contains a text field that may be left empty or may contain a case insensitive text. Log events must contain this text in order to be added to the search result.
Match all (AND) Match any (OR)	Search criteria operator that specifies how individual search criteria shall be combined. • AND instructs Retrospective to retain log entries matching all defined search criteria. • OR instructs Retrospective to retain log entries matching any of the defined search criteria.
	Determines if Retrospective shall fetch data from files and containers that did not exist when the monitoring process was started.  new appearing files and containers are dynamically included in the ongoing monitoring process.  new appearing files and containers are ignored.
	Determines if **Autofind** (see 3.4.3.1 Autofind) shall be run for dynamically discovered files and containers in case the defined entry separation (newline control character, date/time etc.) cannot be found in the initial part of the log data.  **Note** In case you had to manually adjust the log entry separation of a data source, you most probably don't want to allow **Autofind** to be run on dynamically discovered files and containers. Otherwise the manually changed configuration would simply be ignored.

	Determines if new name/value pattern fields shall be detected during search and monitoring. Detected fields may then be assigned to profile specific custom columns.
	WARNING: Enabling field detection may heavily slow down the search and monitoring process. Depending on the data volume to be processed, it may even deteriorate the UI responsiveness. You may remedy the latter issue through the following measures.
	• Local Files: Decrease the option value "Max. threads for local file search/monitoring" on the same preferences page.
	• Containers: Decrease the following option values on the same preferences page.
	  ○ "Max. threads for search/monitoring in Docker"
	  ○ "Max. threads for search/monitoring in Kubernetes"
	• Remote files accessed through SSH: Decrease the option value "Maximum connections" on the SSH host in the advance tab within the Host Manager.

5.1.2 Defining Search Criteria

Search criteria are parameters that specify which log data the search or monitoring action should retrieve and present in the result table. Search criteria help you reduce the amount of data, in order to obtain results of immediate interest. Retrospective uses text and time search criteria that may have some additional options (e.g. case-sensitivity in text criteria).

Depending on the general settings (see 4 Preferences), a ***Search*** tab initially contains one or two search criteria but Retrospective lets you define additional search criteria and thereby increase the precision of the search definition.

By default, Retrospective displays a search text field together with the following toggle buttons.

(.*) Press this button if the entered search text shall be interpreted as a regular expression (Regex)

Press this button if entered search text is aimed to be case-sensitive

If you want to define more complex search criteria, press the button "switch to advance search criteria mode" located above the search text field. You'll notice that a drop-down list appears left of your search text field.

The drop-down list lets you switch to a different search criteria type such as "Text starts with", "Time is between" or "Offset hours". The checkboxes let you further constrain a text search criteria and thereby instruct the search engine whether it has to be interpreted as a regular expression or if it is case sensitive.

Individual search criteria can be composed as follows.

1. Click the icon to add a search criteria. To remove a search criteria, click the icon next to the given search criteria.
2. Select logical operator for combining individual search criteria.

   Matches all (AND): Consider log entries matching all defined search criteria.

   Matches any (OR): Consider log entries matching any of the defined search criteria.
3. Select the search criteria type of your choice (e.g. "Text not contains").
4. Optionally select additional options such as "regular expression" or "case sensitive" through the appropriate toggle buttons.
5. Enter the search text, regular expression, time range or time offset.
6. Start searching by pressing the Start Search button.

5.1.2.1 Text Search Criteria

The drop-down list shown in the Advanced mode lets you choose from the **text** search criteria listed in the table below. The search text field should either contain a search string (given phrase) or a regular expression, depending on whether the "Regex" toggle button is enabled or not. Optionally you can also enable the "case-sensitive" toggle button that instructs the search engine to mind case sensitivity.

Type	Description
Text contains	Retains log entries that contain given phrase or match the entered regular expression.
Text starts with	Retains log entries in which the given phrase appears at the beginning.
Text ends with	Retains log entries in which the given phrase appears at the end.
Text not contains	Retains log entries that do not contain given phrase nor match the entered regular expression.

When the "Regex" checkbox is **NOT** selected, you can use one or several **wildcard characters** to locate a specific log entry without knowing the exact phrase to be searched.

Wildcards are special characters that can represent unknown characters in a text and are handy for locating log entries with similar, but not identical, data. You can use them anywhere in a text search criteria string.

- The **asterisk (*)** matches any number of characters.
- The **question mark (?)** matches a single character at a specific position. When used alone, it matches non-empty content.

Since wildcard characters are interpreted in a special manner, you have to quote them with backslash (\) if you want to search for them explicitly.

- If you want to search for **asterisk (*)**, you have to type *****.
- If you want to search for **question mark (?)**, you have to type **\?**.
- If you want to search for **backslash (\)**, you have to type **** (please note, that using a single backslash that is not preceding another backslash, an asterisk or a question mark results in a syntax error of the **text** criteria).

*In case of Regex-based **text** criteria consult https://docs.oracle.com/en/java/javase/17/docs/api/java.base/java/util/regex/Pattern.html and check the summary of **regular expression** constructs or learn more about regular expressions by reading one of the many publicly available reference pages.*

5.1.2.2 Time Search Criteria

Note

In search mode, Retrospective accepts multiple time search criteria per tab only when they are combined with **"Matches any (OR)"** operator.

When **"Matches all (AND)"** operator is used only one time search criteria per tab is allowed.

No time search criteria is allowed for monitoring.

Note

Time search criteria ensure that log entries newer than **now** (the moment in which search button is pressed) are not retained (however, in the case of **is between** and **is before absolute time** criteria, the user has to make sure that the upper time boundary is not newer than **now**). This, to some extent, prevents a race condition occurring when constantly changing files or container log streams are searched. However, when text search criteria are used, then there is no way to prevent such a race condition and it cannot be predicted which final parts of the changing files or container log streams will be searched by Retrospective and which will not.

The drop-down list shown in the Advanced mode lets you choose the following **absolute time** search criteria types.

Type	Description
Time is between	Retains log entries that were created between the INCLUDING lower time and the EXCLUDING upper time boundaries.
Time is before	Retains log entries that were created before the EXCLUDING upper time boundary.
Time is after	Retains log entries that were created between the INCLUDING lower time boundary and **now** (i.e. the moment in which search button was pressed). In the case of these criteria, the upper time boundary is enforced to prevent the race condition mentioned in the note above. In consequence, the lower time boundary must not be in the future. When the search is started, the EXCLUDING upper time boundary (now) appears behind the defined search criteria and thus informs you about the time range used by the core processing engine.

Note

The date/time format can be changed globally on the “Search/Monitor” preferences page.

When editing the individual date/time fields, the following user interaction through convenient functionality is possible.

Field Selection via Mouse

- Select any field with a direct mouse click. Clicking the center mouse button will cycle through the fields

Field Selection via Keyboard

- The left and right arrow keys move one field at a time.
- The Tab key will move to the next field (and from the last field to the next control)
- If the user is entering the value of the field via the keyboard and the maximum number of digits for that field has been typed, selection is automatically moved to the next field.

Field Editing via Mouse

- Scrolling with the mouse wheel will increment and decrement the field

Field Editing via Keyboard

- The Up and Down arrow keys will increment and decrement the field value
- The "+" and "-" keys will increment and decrement the field value
- Enter the numerical value directly with the keyboard

Full Selection & Copying (Windows/Linux)

- Select the full content of the field by pressing [Ctrl]+A while it has the focus.
- Copy the selected field content to the clipboard by pressing [Ctrl]+C
- Press [Ctrl]+V to paste the previously copied content from the clipboard to another field.

When working on Mac, use [⌘] instead of [Ctrl].

Choose date/time from dialog

- Click the  icon choose a date/time from a dialog.

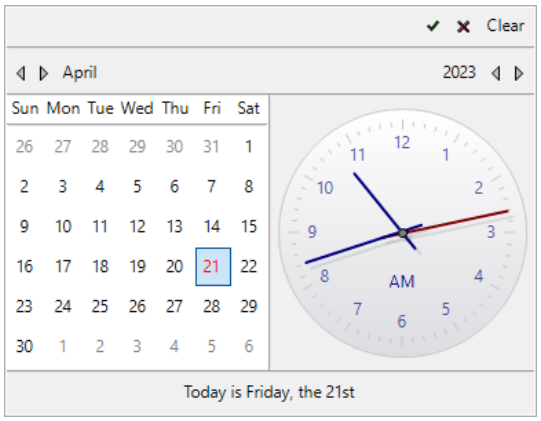

The drop-down list shown in the Advanced mode lets you choose the following **relative time** search criteria types (Please note that the date/time format used for the different examples is *'yyyy.MM.dd HH:mm:ss'*).

Type	Description
Offset days	Retains log entries that were created between the entered number of days ago and the current moment (now). *Example:* *When a 2 days offset is entered and the start button is pressed at '2015.08.04 15:34:27', log entries created between '2015.08.02 15:34:27' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Offset hours	Retains log entries that were created between the entered number of hours ago and the current moment (now). *Example:* *When a 5 hours offset is entered and the start button is pressed at '2015.08.04 15:34:27', log entries created between '2015.08.04 10:34:27' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Offset minutes	Retains log entries that were created between the entered number of minutes ago and the current moment (now). *Example:* *When a 10 minutes offset is entered and the start button is pressed at '2015.08.04 15:34:27', log entries created between '2015.08.04 15:24:27' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Offset seconds	Retains log entries that were created between the entered number of seconds ago and the current moment (now). *Example:* *When a 30 seconds offset is entered and the start button is pressed at '2015.08.04 15:34:27', log entries created between '2015.08.04 15:33:57' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Current day	Retains log entries that were created within the current day. *Example:* *If the start button is pressed at '2015.08.04 15:34:27', log entries created between '2015.08.04 00:00:00' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Current hour	Retains log entries that were created within the current hour. *Example:* *If the start button is pressed at '2015.08.04 15:34:27', all log entries created between '2015.08.04 15:00:00' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*
Current minute	Retains log entries that were created within the current minute. *Example:* *If the start button is pressed at '2015.08.04 15:34:27', all log entries created between '2015.08.04 15:34:00' (INCLUDING) and '2015.08.04 15:34:27' (EXCLUDING) are retained.*

Type	Description

Note

When the search is started, the INCLUDING lower time boundary and EXCLUDING upper time boundary (now) appear behind the defined search criteria and thus inform you about the time range used by the core processing engine.

5.2 Result Visualization (Result Table)

Search and monitoring results are shown in the result table where distinct information such as date/time, log level, data etc. appears in individual columns. To hide a column from the table or show it again, simply right click on a column header and select/deselect the desired column name from the pop-up menu. In the Retrospective properties on the page "***Result Options > Visible Columns"*** you can define the columns to be shown by default.

If a row from the table gets selected, its content appears in a **detail panel** just underneath the table. When several rows are selected, this detail panel contains the data of the top most selected row.

By default search entries are chronologically sorted, the most recent entries appearing at the bottom of the table. The default sorting can be changed by the user at any time through a mouse click on the desired column header. Here you can sort your entries by any column – except the data column – and also invert sorting by pressing the same column header again.

Monitoring entries, on the other hand, are added to the bottom of the table as they're detected regardless of their date/time. While monitoring goes on, individual columns cannot be sorted by the user. Once the monitoring process is stopped, the entries can be sorted by individual columns - except the data column.

5.2.1 Result Table Toolbar

Retrospective provides a toolbar on top of the log result table that makes logs exploration more convenient.

Icon / Control	Description
Clear results	Removes all current entries from the result table in monitoring mode.
Pause automatic scrolling	Pauses the automatic scrolling of the result table to the last discovered and displayed log entry. When automatic scrolling is paused, the icon appears on the button. Pressing the button again re-enables automatic scrolling.
Word wrap	Enables/Disables line wrapping at word boundaries within the result detail view. The default state for this word wrapping can be defined within the preferences dialog page Result Options > Formatting.
Data Formatting	This control lets you define the data formatting to be applied within the result detail view. The icon appearing on the control reflects the currently selected formatting option. The initial formatting for new search/monitoring tabs is adopted from the preferences page Result Options > Formatting.

Icon / Control	Description
	The following options are available from a pop-up menu that is displayed if you press the control arrow:
	JSON — Formats JSON data by using the corresponding indentation defined in the preferences dialog.
	XML — Formats XML data by using the corresponding options defined in the preferences dialog.
	Disable — Disables data formatting.
	Edit Formatting Options — Opens the preferences dialog for viewing/editing result data formatting options.
Define highlighting	This multi-function control lets you define the highlighting applied to individual result entry table rows and to different text parts within the result detail view. The icon shows the current selected highlighting option. By default, result entry table rows are highlighted by their log level. This fact is shown by the initially displayed icon. 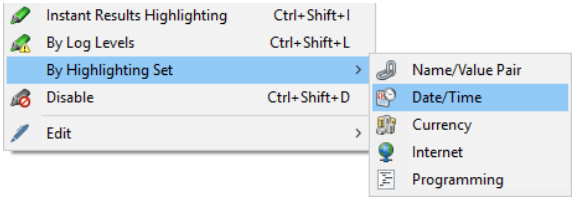 You can choose different highlighting options by pressing the button or by selecting the desired item from the pop-up menu that appears when you click the arrow located on the right side of the button. The highlighting menu and some individual options are also directly accessible through keyboard shortcuts. The following highlighting options are available:
	Instant Results Highlighting — Triggers instant highlighting where you define the color and the expression of a single matcher aimed to be used once (definition is not persisted) within the following pop-up dialog. This dialog also appears when you press the button/icon. 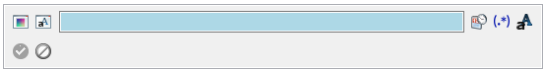
	By Log Level — Instructs Retrospective to highlight result entry table rows by their log level category **ERROR** or **WARN**. Log levels definitions can be viewed and edited within the preferences dialog that is also directly accessible through the "Edit" item of the pop-up menu shown above.

Icon / Control	Description	
	By Highlighting Set	Presents a sub-menu for selecting one of the predefined highlighting sets that contain matchers for locating specific terms or text sections and highlight them with specific colors. Highlighting sets can be viewed and defined within the preferences dialog in the Result Options section. The page is also directly accessible through the "Edit" item of above shown pop-up menu.
	Disable	Disables highlighting.
	Edit	Shows a sub-menu that lets you open the preference dialog for viewing/editing log level definitions or highlighting sets.

Filter results by date/time

Locally filters the result entries by the values that appears in the Date/Time column. Matching log entries are referred to as being filtered. Non-matching entries are hidden by default (filtered-out).

Note

Due to time zone conversion, adjusted time offsets for remote SSH hosts and to user defined formatting, the value displayed in the Date/Time column is usually different from the original timestamp found in the log entry.

If you need to filter log entries by the original date/time pattern, please refer to the *local text filter* further down.

Press the button and you'll see a pop-up box that lets you define a local date/time filter or remove the existing one.

The button icon is shown together with a green bullet when a local date/time filter is defined.

Note

The local date/time filter can be defined with the time(s) of the current selected result table entry. Simply press the right mouse button and select the menu "Set Local Time Filter with" and menu item of your choice.

Icon / Control	Description
Filter results by log level	Locally filters the result entries by their log level. Matching log entries are referred to as being filtered. Non-matching entries are hidden by default (filtered-out). 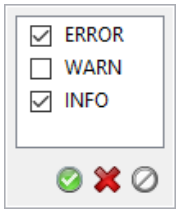 Press the button and you'll see a pop-up box that shows all distinct normalized log levels present in the result. Deselect the log levels you want to have filtered-out from the result table. The button icon is shown together with a green bullet when a local log level filter is defined.
Filter results by custom columns	Locally filters the result entries by their custom columns. Press the button and you'll see a pop-up box that contains a filter row for every visible custom column present in the current profile.  The button icon is shown together with a green bullet ( ) when at least one custom column filter is defined. Log entries are filtered if all defined custom column filters, together with the entered local filter text, match their content. Filtering results by custom columns works exactly the same as local filtering described below but the filter phrase is not matched against the Date/Time and Data columns but rather against the value of given custom column. Wildcard characters described in the local filtering can be used the same way in the custom columns filtering. It's worth noting that in case of custom column filtering, user can enable "Not empty" option through the toggle button . This filter shows results entries in which given custom column has a non-empty value. Since the introduction of RQL, custom column filters can directly be expressed as part of a query entered in the filter field described further down. The use of the custom column filter dialog however offers the following few subtle advantages or differences. 1. Individual user-defined column filters can be marked as case-sensitive. With RQL, on the other hand, the entire query is either case-sensitive or case-insensitive. 2. Each custom column filter maintains a dedicated history dropdown list. 3. The entered filter text matches if it is contained in the value of the custom column. Therefore, in the result details pane, only the matching part of the custom column field value is highlighted. Custom filters defined in this dialog may be combined with custom column filters defined as part of an RQL query.

Icon / Control	Description
	Locally filters the result entries by the entered phrase or RQL, depending on the specified local text filter type. For further information, please consult chapter 5.2.3. Local Filter Field. **Note** Due to time zone conversion, adjusted time offset for remote SSH hosts and to user defined formatting, the value displayed in the Date/Time column is usually different from the original timestamp pattern found in the log entry. The local text filter by default does not consider the original timestamp. If you need to filter log entries by the Date/Time column, please refer to the *Filter results by date/time* control described above. In case you really want to filter log entries by their original timestamp, press the toggle button explained further down. Matching log entries are referred to as being **filtered**. Non-matching entries are hidden by default **(filtered-out)**. The local filter control keeps track of the previously used filter texts. You can easily get them back by selecting an entry from the list that appears when you press the arrow button located right to the filter text field. Even with precisely defined search criteria, log search can return thousands of entries. Instead of refining the search criteria with additional parameters and retrieving the matching log entries from the source computers again, you can tell Retrospective to filter search results locally and display only those entries that contain given text or log levels.
Include original timestamp	Determines whether the original timestamps shall also be considered when locally filtering log entries. By default, timestamps are excluded - not taken into account - when the local filter is applied. Indicates that original timestamps are currently excluded. Press the button to have them included. Indicates that original timestamps are currently included. Press the button to have them excluded. You may include original timestamps in case you are trying to find a number or a word (e.g. the name of a month) not only in the log payload but also in the timestamp. For explicitly filtering log entries by date/time, use the *Filter results by date/time* control explained further up.
Match case	Determines whether the local filter phrase or RQL query described above should be case-sensitive. By default, the local filter phrase or RQL query is case-insensitive.

Icon / Control	Description
	Indicates that the local filter phrase or RQL query is currently considered to be case-insensitive. Press the button to switch to case-sensitive.
	Indicates that the local filter phrase or RQL query is currently considered to be case-insensitive. Press the button to switch to case-insensitive.
Show filtered-out results	Displays or hides log entries that are filtered-out by the local text filter and/or the log level filter. If the button is de-activated, all log entries are shown but filtered ones are highlighted with a green background color.
Next filtered result	Navigates to the next filtered row in the table in the case where all log entries are shown.
Previous filtered result	Navigates to the previous filtered row in the table in the case where all log entries are shown.
Enable disk storage index	Creates an index that enables including of hard drive stored results in local filtering and sorting. Note that this option becomes visible once Retrospective finishes searching/monitoring. *Note that this button is visible only if the option "Enable disk storage" in the "Result Options" preferences page is enabled.*
Create result data snapshot	Creates a snapshot (immutable copy) of the current result data. Result data snapshots appear in the "Result Snapshots" tab, from where you may restore a search tab from a snapshot at any time. Restored search tabs will be backed by another copy of the snapshot. That way we ensure that a snapshot remains unchanged after its creation.
Save result data to File	The drop-down button lets you save the data of desired columns from the result table to files of different types. The pre-selected save/file type can be changed in the preferences dialog within the "Result Options" page. **Note** When "Enable disk storage index" ( ) is not selected, SORTING and LOCAL FILTERING will be preserved but only the portion of data currently present in the local memory will be saved to the file.
Save to text file	Saves the result table data to one or several text files. The unchanged original content of individual entries is written to the text file without custom columns nor additional information such as host and path.
Save to CSV file	Saves the result table data to one or several comma-separated (CSV) files. CSV files contain tabular data in plain text where each line represents a data record, each of them consists of one or several fields, separated by the character(s) of your choice.

Icon / Control	Description
	Saving the data to CSV files only makes sense if the content of all fields within the result table fits on a single line.
 Save to Excel file	Saves the result table data to one or several Excel files. The following formats are available. • Excel Workbook, Decorated (*.xlsx) • Excel Workbook (*.xlsx) • Excel 97-2003 Workbook, Decorated (*.xls) • Excel 97-2003 Workbook (*.xls) In the resulting Excel spreadsheets, the background color of individual result entry rows is always derived from their log level.
 Save to JSON file	Saves the result table data to one or several JSON files. JSON (JavaScript Object Notation) is formatted and self-describing text for storing and interchanging data.
 Analyze result data in Koia	Writes the result data to a database (CouchD) in order to be further analyzed in a Koia web-application. This functionality becomes available only after you've performed the following steps. 1. Install CouchDB by following the instructions at https://github.com/centeractive/koia/tree/master/README.md#installing. 2. Configure the access to your CouchDB and the Koia web-application in the preferences page "Result Options/Koia Data Analysis" (menu item File > Preferences).

5.2.2 Result Table Headers

Except the "Data" column, individual columns can be ascending and descending sorted upon one or a repeated mouse click on the corresponding column header. Sorting can be active for a single column only. Therefore, as soon as you click on a column header, the sorting directives of previously sorted fields is not considered any more.

Through a mouse click on the "Data" column header, you change the content of the displayed values. The following icons located left on the column header, indicates what data is currently shown.

 The original timestamps found in the log entries are currently not shown. Click on the column header if you want to exclude the original timestamps from the Data column.

 The original timestamps found in the log entries are currently shown. Click on the column header if you don't want to include the original timestamps in the Data column.

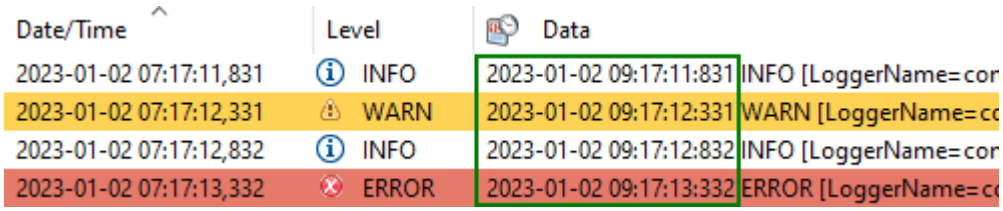

5.2.3 Local Filter Field

There are two types of local filter fields.

1. The Classic Contains Filter field
2. The Retrospective Query Language (RQL) field

Retrospective uses the same type of filter field for the result tables on all search tabs. The "Retrospective Query Language (RQL) field" is used by default. If you want to switch globally to the "Classic Contains Filter Field", open the Preferences Dialog and change the value of the "Text Filter Type" combo box on the page "Result Options > Local Filter".

5.2.3.1 Classic Contains Filter Field

The Classic Contains Filter field locally filters the result entries by the entered phrase. This filter type was the standard local text filter in Retrospective prior to version 6. Since the introduction or RQL, we call it classic.

The filter phrase is applied to the raw data of the log entries, which is what you see in the Result Detail Pane when you select individual log entries.

- The **asterisk (*)** matches any number of characters.
- The **question mark (?)** matches a single character at a specific position. When used alone, it matches non-empty values.
- If you want to search for one of these characters explicitly, they need to be escaped by a slash (`"\*"` or `"\?"`).

An arrow button appears right to the filter field. Press it and you'll see a drop down list that lets you select one of the last applied search terms again.

The main feature of the classic filter is still that you can just paste a block of text into the filter field, even with spaces between words, and you don't need to put double quotes around it to define a complete search term.

5.2.3.2 RQL Field

This field lets you define local filter criteria using the Retrospective Query Language (RQL). For detailed information about RQL, please consult chapter 6 Retrospective Query Language (RQL).

By default, the entered terms are applied to the Data column. You may however also define filter criteria for other columns. The supported operators are "=", "<", "<=", ">" and ">=". Type "=", "<" or ">" in the text field and you'll see a drop-down list with available column names that may be used in the RQL query.

Note that the drop-down list only appears if all following condition are met:

1. The operator "=", "<" or ">" is placed at the beginning of the query or it directly follows a space or one of the non-escaped characters listed below:
   - opening bracket "("
   - minus "-"
   - plus "+"
   - negation "!"
2. The operator is preceded by a space or it appears at the end of the query.
3. The operator is not located inside a phrase (a quoted term).

The following rules apply to the column list shown in the drop-down:

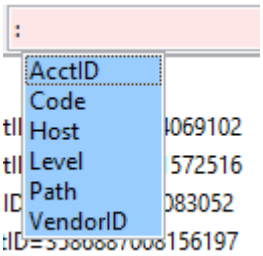

- The column names are sorted alphabetically. As soon as the drop-down appears, type the first letter of the desired column and it will be selected. You may also use the arrow-down and arrow-up keys to navigate to the desired column. Then type ENTER to have the column name inserted in the RQL query in front of the entered operator. At this point, the operator may be amended if for example you want to use "<=" or ">=".
- The list contains visible columns only (standard columns and custom columns.
- The list does not contain columns that contain a space in their name because it is not possible to define valid RQL terms with such columns.
- The list does not contain the "Data" column since this is the default column where no explicit column specifier is expected.
- The list does not contain the "Date/Time" column. Use the button from the toolbar to define a time filter.

An arrow button appears right to the filter field. Press it and you'll see a drop down list that lets you select one of the last applied RQL queries again. Note that each query is preceded by a sign ("✓" or "✗") that indicates if the query is valid in the current context.

5.2.4 Log Level Column

By default, the log level appears in its own column in a normalized form (e.g. "ERROR") regardless if it was logged in another language (e.g. "GRAVE" in French). When highlighting is enabled (button), Retrospective presents "ERROR" entries with a red background and "WARN" entries with an orange background. A log level specific icon appears in each row as long as a log level can be detected.

The result table can be filtered by log levels. Simply click on the button located on the result toolbar and you'll see a pop-up box that shows all distinct normalized log levels present in the result data. Select/deselect the log levels you want to have included/excluded in the result table. The icon show that a log level filter is currently active.

5.2.5 Sorting

The result table rows can be sorted by the values of individual columns through a simple mouse click on the corresponding column header. A click on an already sorted column inverses the sorting order (ascending <-> descending). The predefined **Data** column contains the non-structured content of the log entries and cannot be sorted.

Note

When you sort the log data by the **Date/Time** column, be aware that the date (year, month or day) of certain log entries needs to be "guessed" by Retrospective and may not correspond to the real log entry creation date.

- This is the case if an application omits the date but logs only the time.
- Another reason could be that you manually configured a data source and defined a date format that retains only the time.

In such cases, Retrospective deducts the date as follows:

- If the data source is a file, the date will be taken from the file modification date.
- If the data source is a container, the date will be taken from the log entry reception date.

5.2.6 Extended (smart) Sorting

Extended Sorting considers several visible and hidden columns beside the main sort column, and thus produces a result of higher accuracy. This corresponds to an SQL query ordered by several columns. Extended Sorting is applied only when both of the following conditions are met, otherwise it is automatically disabled.

- Sorting of one of the standard columns **Date/Time**, **Level**, **Host** or **Path** is requested.
- The number of result table entries doesn't exceed the value of the preference option **Extended sort limit.**

If Extended Sorting is disabled, Retrospective sorts result table entries ONLY by the selected column. If you click on the Level column header, you will have entries sorted by Level but all entries with the same level e.g. "INFO", will appear in an order depending on their insertion to the database. Therefore, when log data was collected from multiple hosts, you will probably not see them in chronologically order because the Date/Time won't be considered for sorting.

If Extended Sorting is applied, sorting is not only performed on the primary column but also on additional columns (secondary columns) according to the following table:

Primary Sort Column	Secondary Sort Columns
Date/Time	Path -> OrderOfArrivalToRetrospective
Level	Date/Time -> Path -> OrderOfArrivalToRetrospective
Host	Date/Time -> OrderOfArrivalToRetrospective
Path	Date/Time -> OrderOfArrivalToRetrospective

As you may imagine, when Extended Sorting is applied, sorting can take more time than otherwise. It however depends on the size of your entries and their date/time distribution. The following performance figures have been gathered during one of our performance testing

sessions. In this session we had 10 millions of log entries of rather small size, similar structure and regular time distribution.

- **With Extended Sorting**: Sorting by Date/Time took: **11 m 30s** (effective sorting by Date/Time and then Path and then OrderOfArrivalToRetrospective)
- **No Extended Sorting**: Sorting by Date/Time took: **4 m 41s** (effective sorting by Date/Time only).

The default value of the preference option **Extended sort limit** is 100000 log entries. This means that Extended Sorting is not applied when your result table contains more than 100000 entries in your result table. If your result table contains more than 100000 log entries and you really need Extended Sorting, then you can try to increase Extended sort limit and observe how it impacts the performance of sorting. If sorting with Extended Sorting does not take a lot more of time, then your settings are fine. Otherwise, you probably want to decrease Extended sort limit and live with the imperfect sorting. In case you only sort by Date/Time (most common cases), not having Extended Sorting applied, is typically not much of a problem.

5.2.7 Result Table Context Menu

The result table context menu appears upon a right-click on a log entry row. It is composed of the sub menus and menu items described below.

Submenu/Menu Item	Description
Copy	Copies the selected result entries to the clipboard.
Copy path	Copies distinct paths from the selected result entries to the clipboard.
Add Time Search Criteria with	***Sets the local date/time filter on the search tab.***
'Time is after' from selection	Sets a local filter of type 'Time is after' derived from the selected result entry. This menu item is enabled only if a single result entry is selected.
'Time is before' from selection	Sets a local filter of type 'Time is before' derived from the selected result entry. This menu item is enabled only if a single result entry is selected.
'Time is between' from selection	Sets a local filter of type 'Time is between' derived from the selected result entries (the ones with the oldest and the newest date/time). This menu item is enabled only if a multiple result entries are selected.
Add Time Search Criteria with	***Adds a new time search criteria to the search tab.***
'Time is after' from selection	Adds a new 'Time is after' search criteria derived from the selected result entry. This menu item is enabled only if a single result entry is selected.
'Time is before' from selection	Adds a new 'Time is before' search criteria derived from the selected result entry. This menu item is enabled only if a single result entry is selected.
'Time is between' from selection	Adds a new 'Time is between' search criteria derived from the selected result entries (the ones with the oldest and the newest date/time). This menu item is enabled only if a multiple result entries are selected.
Replace Time Search Criteria with	***Replaces all time search criteria on the search tab.***
'Time is after' from selection	Replaces all time search criteria with a new 'Time is after' search criteria derived from the selected result entry. This menu item is enabled only if a single result entry is selected.
'Time is before' from selection	Replaces all time search criteria with a new 'Time is before' search criteria derived from the selected result entry. This menu item is enabled only if a single result entry is selected.

Submenu/Menu Item	Description
'Time is between' from selection	Replaces all time search criteria with a new 'Time is between' search criteria derived from the selected result entries (the ones with the oldest and the newest date/time). This menu item is enabled only if a multiple result entries are selected.
Search with same profile	***Opens a new search tab that is based on the current profile, the current search criteria and a new time search criteria derived from the selected entries.***
<Menu items>	For details about dependent menu items please consult section 5.4.2 Based on Result Entries.
Search all data in selected path(s)	Opens a new search tab that is ready for searching log entries in the path of the selected entry or entries.
Search in selected path(s)	***Opens a new search tab that is based on the paths of the selected entries, the current search criteria and a new time search criteria derived from the selected entries.***
<Menu items>	For details about dependent menu items please consult section 5.4.2 Based on Result Entries.
Remove Selected	Removes the selected result entries from the table. This function can also be triggered by simply pressing the delete key.
Remove Filtered-Out	Removes the entries that don't match the local text and log level filters. **Note** This option is applied only to locally filtered results sets.

5.3 Result Details Pane

The details of the selected row from the result table are shown in the result details pane located at the bottom of the table. Any text that matches a text search criteria is highlighted with a different yellow color each. Custom column fields are highlighted in blue.

If you define a local text filter, a date/time or custom column filter, such text will also be highlighted in green, thus enabling you to quickly find the information you're looking for.

In case the local text filter is an RQL query, excluding terms are highlighted in red. If a matching term is directly or indirectly contained in an excluding group, it is marked with a dashed red border. This helps you understand, why certain log entries were excluded from the result. To also see non-matching log entries and be able to select them, you'll first have to toggle the button in the result table toolbar.

If your log data contains XML content, Retrospective can present it in a formatted (pretty-printed) manner and through this make it more easily readable. Optional XML syntax highlighting even improves readability. XML formatting can be enabled within the result table toolbar. Related settings are configured in the Preferences dialog within the "Result Options/Formatting" page.

If your log data contains JSON content, Retrospective formats (pretty-prints) it if you enable the option in the result table toolbar. Related settings can be adapted in the Preferences dialog within the "Result Options/Formatting" page.

When you press the right mouse button while the mouse pointer is located at the result details pane, the following pop-up menu offers useful functionality. Note that most of these menu items are enabled only if some text is selected.

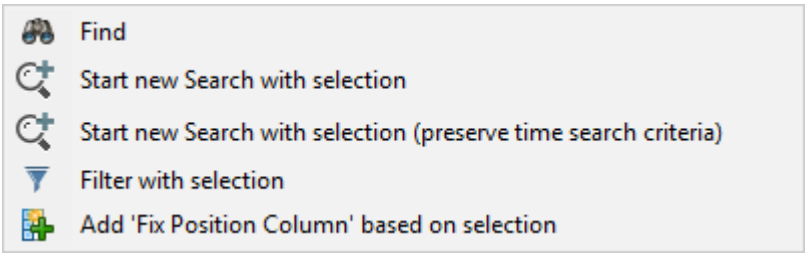

For a convenient result entry analysis, Retrospective lets you maximize the result details pane by clicking on the icon in the right corner of the entry details pane. For a better results set overview, minimize the result details pane by clicking on the icon.

5.4 Context-Diving

5.4.1 Based on Statistics Chart

The statistics chart located on top of the result table provides a quick overview on how the detected log entries are distributed over the search period. When you place the mouse pointer on individual bars, a tooltip appears with information about the corresponding time slice and the number and type of entries that were found in there.

If you click on a bar, a new search tab is created that contains the same search criteria as the original search tab. If the original search tab also contained a time search criteria, it will be ignored. A new time search criteria, corresponding to the time span of the activated bar, will be added to the new search tab instead.

Note

The time span of individual bars from the statistics chart include milliseconds but the format of the time search criteria of the target tab may be less precise. Therefore the result of a search within the newly created tab may produce different results.

5.4.2 Based on Result Entries

Individual entries in the result table often attract your interest but because they were obtained with restrictive search criteria, you're actually missing contextual information. If for example you search in log files using the text search criteria "Error", you'll only obtain log entries that contain this exact term. Retrospective supports you in such cases and enables you to easily retrieve contextual information as follows.

1. Within the same profile.
2. In the log file where one specific log entry was found.
3. In the log files where a set of selected log entries were found.

To obtain the contextual data, select the row (or rows) of interest in the result table and press the right mouse button. Within the appearing pop-up menu choose the desired menu item. This opens a **new search tab** where a search can be performed immediately or after further adjusting the search criteria.

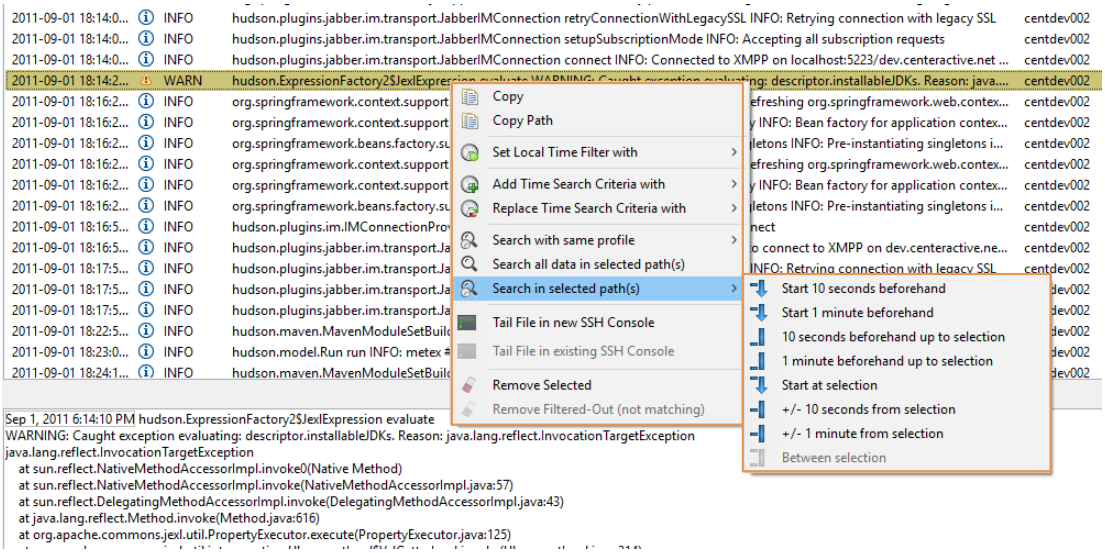

Depending on the chosen sub-menu and menu item, the new search tab will be based on the same data profile as the current tab or on an ad-hoc profile that contains files solely where the selected log entries were found in. The individual menu items have the following purposes.

Icon	Menu Item	Description
	Search all data in selected path(s)	Retrieves all data from the files in which the selected log entries were found.
	Start 10 seconds beforehand	Retrieves data starting 10 seconds before the selected log entry.
	Start 1 minute beforehand	Retrieves data starting 1 minute before the selected log entry.
	10 seconds beforehand up to selection	Retrieves data written during the 10 seconds that preceded the selected entry up to the entry itself.
	1 minute beforehand up to selection	Retrieves data written during the minute that preceded the selected entry up to the entry itself.
	Start at selection	Retrieves data written at the same time and after the selected entry.
	+/- 10 seconds from selection	Retrieves data written during ten seconds before and ten seconds after the selected entry.
	+/- 1 minute from selection	Retrieves data written during one minute before and one minute after the selected entry.
	between selection	Retrieves data written between the time ranges of the selected entries.

5.5 File Browser

The **File Browser** is a handy "entry point" for defining a new search/monitoring activity based on files. This view can be opened by selecting View → File Browser, clicking the icon or by using the [Ctrl] + [I] keys combination.

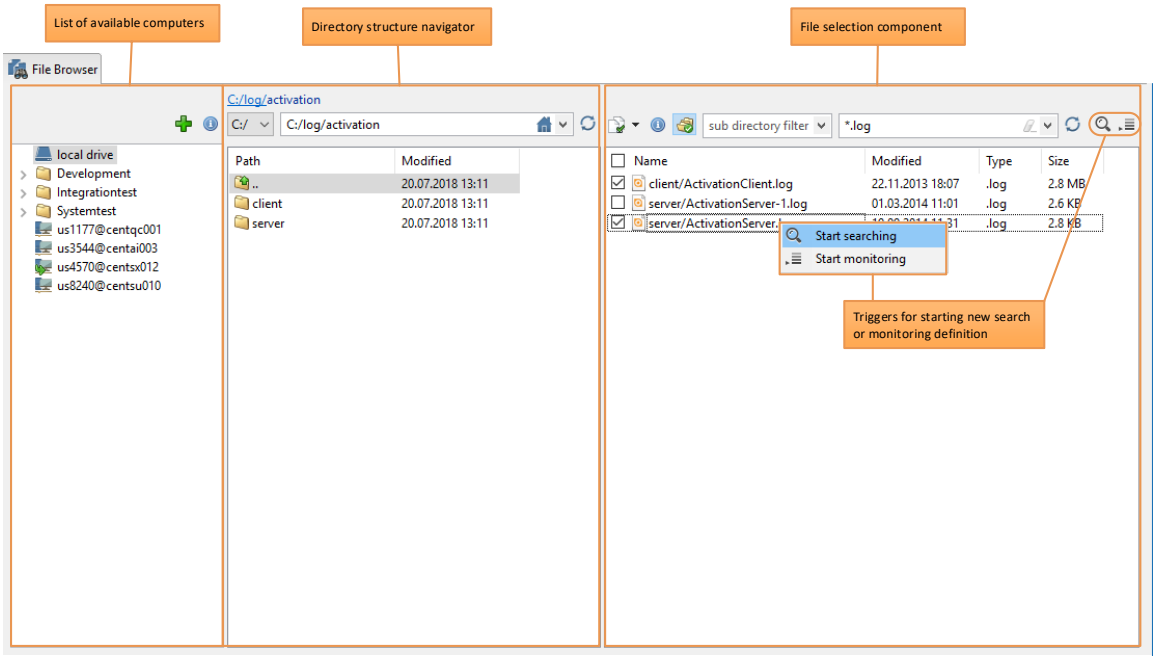

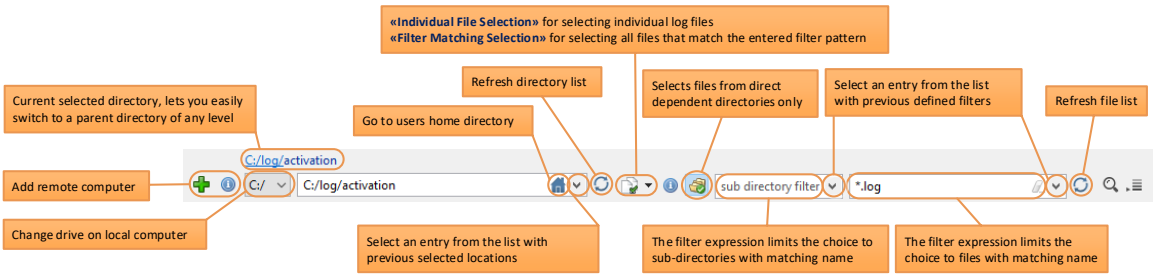

The **File Browser** lets you navigate through the file system on the local computer or on remote servers in order to locate a directory that contains log files you're interested in. Here you can create new search tabs based on the selected log files. This can be achieved in a few steps that are performed in a left-to-right sequence.

1. Choose "local drive" or a remote server in the list of hosts located on the left. The remote servers that appear in the list have typically been defined within the Host Manager view (see below) beforehand. New hosts may however also be defined directly from within the File Browser by pressing the button. Upon selection of a remote server, Retrospective immediately attempts to establish an SSH connection. If the connection is established, the content of the user's home directory is displayed in the directory structure component in the middle of the view.
2. Navigate to the folder that contains the log file(s) you want to search through or monitor. An individual folder opens if you double-click on it. The folder with a double dot, located at the top, lets you navigate to the parent folder.
3. In the file selection component, you can right-click an individual file and select "Start searching" or "Start monitoring" from the context menu that appears when you press

the right mouse button. You may also choose multiple log files by selecting the check box that appears in front of their names. Multiple files can also be chosen by switching to "Filter Matching Selection" ( ) in the drop-down box located at the top. When you press the or button, a new search tab is created and lets you search or monitor the selected log files.

Individual File Selection

The "Individual File Selection" mode by default lets you select individual log files from the current directory. You can reduce the number of listed files by entering a filter expression. Files will appear only if their name matches the file filter. The file filter can contain the wildcards ? and * (i.e. *.log).

If the sub-directory button is enabled, Retrospective only considers files from direct dependent directories. The optional sub-directory filter lets you constrain the sub-directories to show contained files. The sub-directory filter can contain the wildcards ? and *.

Filter Matching Selection

The "Filter Matching Selection" mode by default selects all files from the current directory. The choice of files can be limited by entering a filter expression. The group will only contain the files if their name matches the filter expression. The filter expression can contain the wildcards ? and *.

- An empty filter or '*' selects all files from the current directory
- '*.log' selects all files with a name that ends with '.log'
- 'billing*' selects all files with a name that starts with 'billing'
- etc.

If the sub-directory button is enabled, Retrospective considers files from direct dependent directories only. The optional sub-directory filter lets you limit the number of sub-directories to be included in the group. The sub-directory filter can contain the wildcards ? and *.

After performing other activities within Retrospective you can come back to the file browser at any time later and find the context from where you last triggered your search or monitoring action.

Note

Both in File Browser and in Add Data Source dialog (see 3.4.1 Create New Profile), when a remote host from the list of available computers is clicked, the home directory is the initial folder opened in the Directory structure navigator. The home directory detected on given host is displayed in Host Compatibility Information available in Host Manager (see 3.3 Host Manager).

5.6 Container Browser

The **Container Browser** is a handy "entry point" for defining a new search/monitoring activity analyzing the log data of containers located in Docker/Kubernetes subsystems. This view can be opened by selecting View → Container Browser, clicking the icon or by using the [Ctrl] + [G] keys combination.

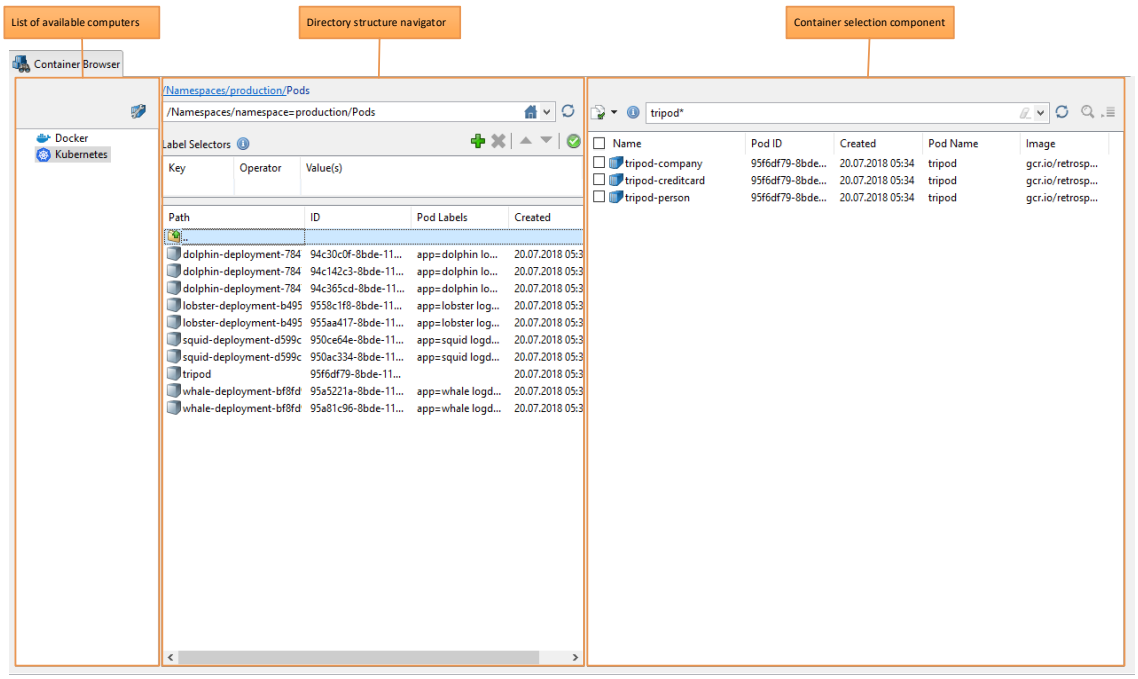

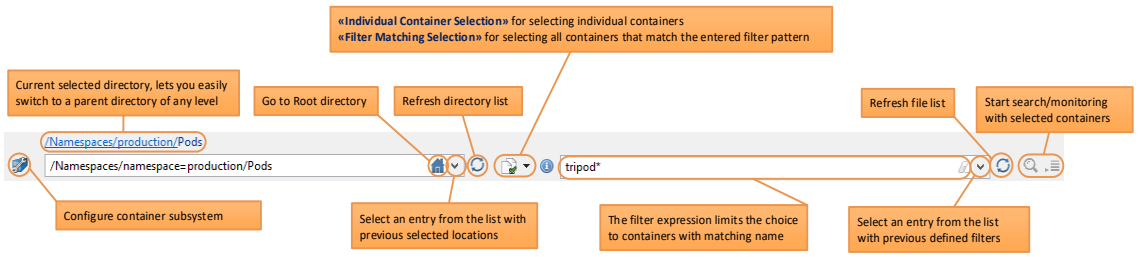

The **Container Browser** lets you navigate through the container subsystems of Docker and Kubernetes accessed through locally installed command line interfaces (CLI) `docker` and **`kubectl`** in order to locate containers you're interested in. Here you can create new search tabs based on the selected containers. This can be achieved in a few steps that are performed in a left-to-right sequence.

1. Choose the container subsystem: Docker or Kubernetes within the left located list. Upon selection of the container subsystem, Retrospective immediately attempts to fetch information about its internal structure and displays it in the directory structure navigator. If the CLI cannot be found, please press the button and properly configure the container subsystem of your choice.
2. Navigate to the folder that contains the containers you want to search through or monitor. An individual folder opens if you double-click on it. The folder with a double dot, located at the top, lets you navigate to the parent folder.
3. In the container selection component, you can right-click an individual container and select "Start searching" or "Start monitoring" from the context menu that appears when you press the right mouse button. You may also choose multiple containers by

selecting the check box that appears in front of their names. Multiple containers can also be chosen by switching to "Filter Matching Selection" ( ) in the drop-down box located at the top. When you press the or button, a new search tab is created and lets you search or monitor the selected containers.

Individual Container Selection	The "Individual Container Selection" mode by default lets you select individual containers from the current context. You can reduce the number of listed containers by entering a filter expression. Containers will appear only if their name matches the filter. The filter can contain the wildcards ? and * (i.e. *.log).
Filter Matching Selection	The "Filter Matching Selection" mode by default selects all files from the current context. The choice of containers can be limited by entering a filter expression. The group will only contain the containers if their name matches the filter expression. The filter expression can contain the wildcards ? and *. • An empty file filter or '*' selects all containers from the current context • '*validator' selects all containers with a name that ends with 'validator' • 'billing*' selects all containers with a name that starts with 'billing' • etc.

After performing other activities within Retrospective, you can come back to the file browser at any time and find the context from where you last triggered your search or monitoring action.

5.7 Typical User Interaction

Regular searching enables you to find desired entries in log files of your choice. Monitoring on the other hand continuously analyzes new entries written to the files of your choice and updates search results whenever the entry matches the search criteria (if any).

This procedure describes the typical steps a user may perform when he wants to search log files for specific content.

1. Click the icon or use [Ctrl]+[T] keys combination to open a new ***Search*** view.
2. In the search definition panel, select the data sources (log files) your search or monitoring action is to be based on.
    a. Select a profile with predefined data sources from the drop-down list located on the left.
    b. Alternatively, if you have not yet defined any profile, or if you want to perform ad-hoc searching or monitoring, choose your source files from the dialog that appears when you press the button.
    c. Click the icon (visible in advanced mode only) to open a dialog that lets you individually choose data sources from various existing profiles.
3. Enter the text that is to be used as search criteria. Only log entries that match the search criteria will be added to the result table. If you want to define more complex search criteria than simply a text, press the Advanced link located below the start button.
4. Choose **`Search`** for regular searching or **`Monitor`** for monitoring.

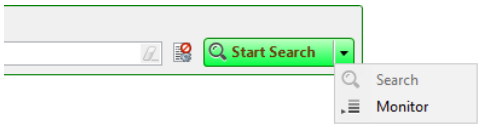

Note

When monitoring data sources defined with "Filter Matching Selection" (e.g. using a wild card '*.log'), Retrospective not only monitors the files/containers present when processing starts but also the ones that are created while monitoring is ongoing. Also support for data sources defined with "Item Selection" is provided. In this case, when the data source path does not exist when processing starts, monitoring waits until it appears.

If new files/containers created during an ongoing monitoring session should be ignored, you need to change the "Enable Dynamic Discovery" option in the preferences dialog. It can be found within the "Tail (Monitor)" box on the "Search/Monitor" page.

Define search criteria (see 5.1.2 Defining Search Criteria) and start searching by clicking the 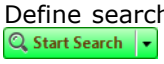

button. In simple mode with a single text search criteria, the search or monitoring process can also be started by pressing the [Enter] key while the cursor is placed in the text field.

5. Review and filter search results.

Note

Should the search result set exceed the memory limit, searching/monitoring does not stop, but the oldest rows are removed from the result table by default. If you enable disk storage on the "**Result Options**" preference page, the oldest rows are written to the local file system instead and remain available for local analysis.

5.8 Core Processing Engine

This section provides some insights into the core processing engine. Firstly, it describes what exactly happens "under the hood", when either searching or monitoring is performed by the user. Secondly, it discusses the details of search optimization, which can save the user a lot of time when searching in log files and containers in some typical use cases.

5.8.1 Inside View

The inside view of the core processing engine is provided by means of the answers to the following six questions:

1) **What happens when the user presses the start button until he sees log entries in the result table?**

In the most general case, searching and monitoring is performed on a Retrospective profile. A profile consists of one or many data sources. Individual data source can point to a local or remote (connected through SSH) computer by a path definition representing either a single file/container, a directory or a filter (e.g. '`Program*.log`'). When the start button is pressed by the user, the first step executed is the evaluation of all data sources in the profile. Data sources are evaluated in parallel by several work threads. The goal of the evaluation is to reveal the list of files/containers related to the data source. It comes to performing either file, directory or container listing and, if needed, applying the filter defined for the data source. As soon as the evaluation of individual data sources finishes, for each file/container, a separate thread takes over the task of either searching or monitoring (tailing) it.

Depending on the action mode (searching, monitoring), the following then applies for each file/container:

Searching:

- If the file is located on a remote computer, a specially crafted SSH command is executed there.
- If the file is compressed, an appropriate uncompressing mechanism is placed at the beginning of the reading pipeline.
- If the file is remote, then splitting data to log entries, extracting date/time from the entry (if time search criteria is present) and filtering defined in the search criteria is performed on the server side. As a result, only the data which matches the search criteria is transferred to Retrospective.
- Gradually, chunks of the file or container log stream are delivered to Retrospective and bytes are converted to characters in accordance with the encoding defined for the data source.
- Each character chunk is then processed by a separate chunk processor thread. If the file is remote, then the chunk processor simply transforms chunks to log entries so they can immediately be displayed in GUI. If the file is local or it is not a file but rather a container, then the following is applied:
    - The chunk processor splits chunks to log entries in accordance with the log entry separation (delimiter) and, if possible, extracts the date/time from the entry in accordance with the date/time format.
    - The chunk processor filters out log entries in accordance with the search criteria defined by the user (this step is omitted if filtering was already applied on the server side).
- Log entries are passed to the GUI layer and displayed to the user.

Graphical representation of the above searching description is presented in the figure below.

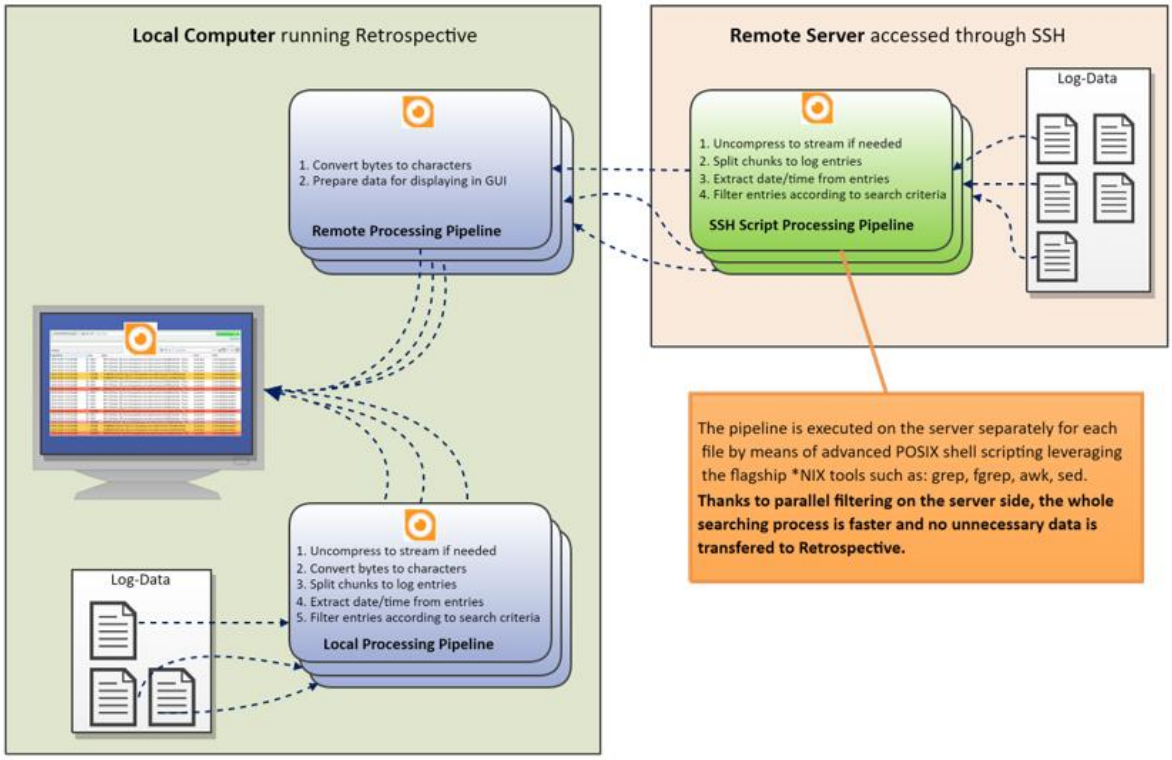

Monitoring initial processing:

- The monitoring task is repeatedly executed until the user cancels the monitoring action (by clicking the [Stop Monitor] button in the GUI).
- The goal of each monitoring task execution is to fetch a set amount of bytes from the end (tail) of the file or container log stream.
- If the file is remote, then during each task execution a specially crafted SSH command is executed on the remote host.
- If the file is compressed, it is ignored as it does not make much sense to monitor a file which is already compressed.
- During the first monitoring task execution, an initial amount of bytes, defined by the File start offset parameter within the preferences (or initial amount of lines defined by the Container start offset parameter) is retrieved from the end of the file or container log stream (the Fail start offset parameter corresponds to the -c parameter of the Linux tail command while Container start offset corresponds to the –n tail parameter).
- During each of the subsequent executions of monitoring task, only the new bytes which appeared in the file or container log stream (if it happens) are fetched.

Note

If the file size has decreased, then the whole file is fetched (it is assumed that the file was filled with new content)

Monitoring data processing (similar to searching)

- During each execution of a monitoring task, if new bytes have appeared in the file or container log stream, gradually, related chunks are delivered to Retrospective and bytes are converted to characters in accordance with the encoding defined for the data source
- Each character chunk is then processed by a separate chunk processor thread

- The chunk processor splits chunks to log entries in accordance with the log entry separation (delimiter) and, if possible, extracts the date/time from the entry in accordance with the date/time format
- The chunk processor filters out log entries in accordance with the search criteria defined by the user (even if file is remote, then filtering is still applied only locally - different from searching)
- Log entries are passed to the graphical interface and are displayed to the user

Graphical representation of the above monitoring description, covering both initial and data processing, is presented in the figure below.

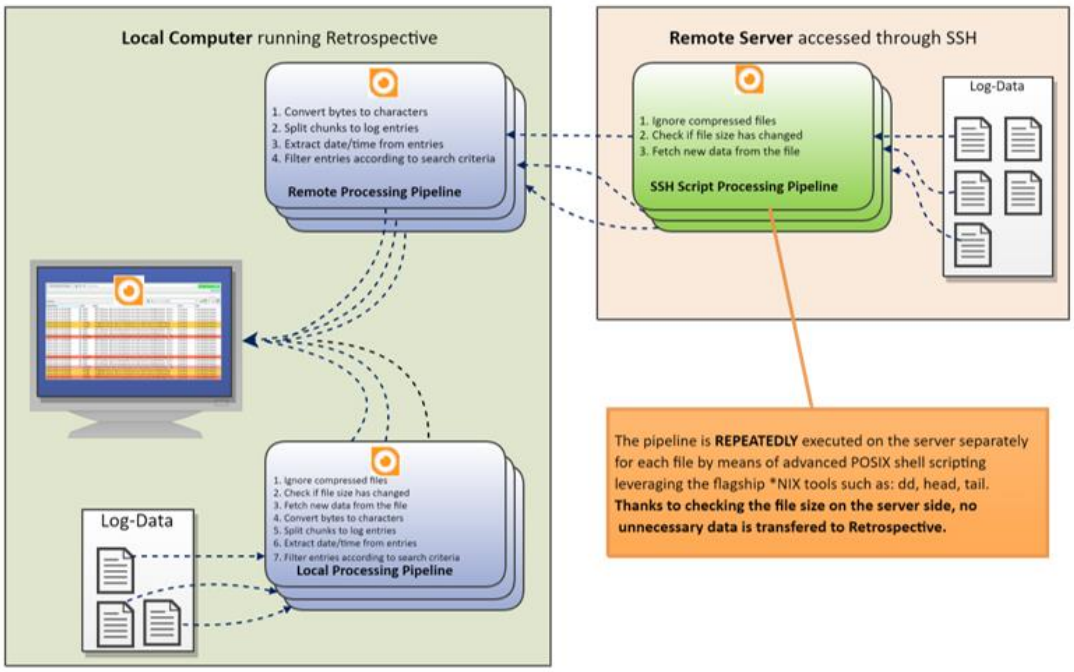

2) What are the differences between local and remote processing?

- In remote processing, evaluation of each data source is performed in a separate SSH command (most of the time a new TCP connection has to be established)
- Remote processing is implemented in shell, awk and sed scripts executed as SSH commands which are entirely performed on the server side, while local processing is implemented in Java and is performed inside Retrospective JVM instance
- In remote processing, during searching or monitoring, for each file a separate SSH command is executed on the server (a separate TCP connection is used). However, thanks to the use of thread pooling and connection pooling, there is no risk of creating too many connections

3) What's going on behind the scenes?

Retrospective ensures that searching or monitoring files on remote servers does not modify anything on the remote file system. Moreover, there is no need to install any agent software on the server. It is sufficient that the server is available through SSH. Retrospective simply exploits the flagship tools which are available on the remote server. More details about detection of these tools are provided in the Host Compatibility Information section (see 3.4.11 Host Compatibility Information). Thanks to detected tools, Retrospective moves the whole processing of remote files to the server side. This provides the following important benefits:

- Resources (CPU, memory) of the node on which Retrospective is launched are not consumed.

- When a file is searched with given search criteria, then only the data matching the criteria is actually transferred through the network to Retrospective.
- The above is also true for date/time filtering. In this case, a sophisticated polymorphic script executed as SSH commands is able to interpret and correctly filter dates in all possible formats and locales.
- In many cases, Retrospective scripts used for searching are faster than regular *NIX tools (awk, grep) thanks to the usage of special optimizations.
- For profiles with many data sources and many files, thanks to a parallel file processing, Retrospective could be compared to multiple instances of grep and awk tools executed simultaneously. This results in a significant performance boost.
- Relating to date/time filtering, search optimization is used (see "Search/Monitor" page on the preferences dialog). By looking at the beginning and end of each file, Retrospective ignores files which simply cannot match the specified time criteria. Yet again, this results in another performance boost.
- By assuming an adaptive approach, Retrospective exploits the resources of the remote server in an optimal manner. If servers respond quickly, then more simultaneous search/monitor SSH commands are allowed. If servers respond slowly, then the amount of simultaneous SSH commands is reduced. In the end, we get the results only as quickly as it is possible.

4) How is Date/Time interpreted on the servers?

The Retrospective assumption to not modify anything on the remote file system, makes the interpretation of date/time on server a real challenge. Please note that the date can have many different formats, locales and be expressed in many different time zones. Let's analyze the following examples:

- May 17 AD 12:20:39.227 GMT +1
- Mai 17 AD 12:20:39.227 CEST
- 2012-05-08T19:50:37.560+02:00

Clearly, it is not easy to have a unified way of interpreting the above date strings. Not only is there a need for understating what a given number represents (is 05 a day a month or a second?), but there is also a need for understanding different languages ('May' in English and 'Mai' in German) and different time zones representations (GMT +1, CEST, PST, +02:00).

Nevertheless, thanks to the introduction of the advanced polymorphic scripting approach, all date strings are entirely parsed on the server side. Retrospective analyzes the date format in remote files (see 3.4.3.1 Autofind) and then, when searching is started a date-format-driven process of search script generation is executed. The resulting SSH command is not only tailored for the date format present in the given file but is also highly performant.

5) What's the difference between searching and tailing (monitoring)?

The detailed differences in steps performed during searching and monitoring were provided in the scope of point 1. These differences can be summarized in the following way:

- Searching is executed once and it finishes, when all filtering of all files is finished, while monitoring is a continuous process implemented by repeated executions of monitoring tasks (one task for each file or container). Monitoring finishes when the user cancels it by clicking Stop Monitor button in the GUI.
- In the case of remote processing: during searching, log filtering is performed on the server side, while during monitoring, the filtering is performed locally.
- As stated in point 1, there are many similarities between implementation of searching and tailing which ensure that both features are highly performant and reliable.

6) What warnings or errors can be reported by Retrospective and how can the user correct the problem?

First of all, there could be some low-level issues such as: problems in connecting to the host; problems in accessing the files (e.g. insufficient permissions to read a file). In such cases, the user has to solve the problems before again executing searching or monitoring.

Secondly, there could be some problems with wrong date/time format. If the date strings in files belonging to given data source are not compliant with the date/time format configured in the data source, then warnings related to date parsing problems can appear. In such cases, the user should correct the data source date/time format, or decouple given data source covering multiple files (e.g. directory) to separate data sources (e.g. with the use of filter) which then will have different date/time formats configured. Problems with differences in date/time formats between files are preliminarily detected during the Autofind procedure (see 3.4.3.1 Autofind).

Thirdly, some problems could be caused by wrong encoding. If the user searches a file, of which the encoding is different from the one configured in the data source, then `SearchOffsetException` can be thrown. In such cases, the user should correct the data source encoding, or decouple given data source covering multiple files (e.g. directory) to separate data sources (e.g. with the use of filter) which then will have different encodings configured. Problems with differences in encodings between files are preliminarily detected during the Autofind procedure (see 3.4.3.1 Autofind).

5.8.2 Search Optimization

Retrospective has a search optimization feature that can be enabled or disabled on the Search/Monitor page within the preferences dialog. If this option is enabled and log data is searched with a time search criteria, for each file greater than 1000 kB (if file is local) or 500 kB (if file is on a remote server), Retrospective will analyze a small amount of log data (64 kb) from the beginning and the end of a file. The analyzed information lets the program find out whether individual files contain data that is located within the time period specified in the time search criteria. All files with log entry dates outside of this period are excluded from the search process. Additionally, when the file is on a remote server, Retrospective narrows down the search to the section of the file that matches the time search criteria. The rest of the file is excluded from the search process.

Note

Search optimization is currently not supported in the containerized subsystems of Docker and Kubernetes but this may change in future releases.

In the case of compressed non-tar files, such as .gz or .bz2 and compressed tar files containing only one compressed file:

- when the file is local, only the data from the beginning of the file is analyzed, because uncompressing the whole file would be too time consuming;
- when the file is remote, it is regularly processed, but when file entries stop matching the time search criteria, then searching is aborted (effectively this is very similar to the approach for local files).

In the case of tar archives, such as .tar, .tar.gz, .tar.bz2, .tgz, .tbz2, search optimization is not performed when an archive contains more than one compressed file because then assumptions about overall date ordering of compressed files cannot be made.

1) What does search optimization do?

The goal of search optimization is to avoid searching files (and, in case of remote files, file parts) whose content cannot possibly match the specified search criteria. Search optimization is used only when:

- the file is larger than 1000 kB (if file is local) or 500 kB (if file is on a remote server)
- the defined search criteria consists only of a **single** time criteria (multiple time criteria are supported only in cases of local searching)
- the defined search criteria consists of a time criteria and some other criteria joined by AND logical operator

When search optimization is enabled, then for each file on which it can be applied, the following is performed:

a) The sanitization (see yellow box below) is applied on the first 64 kB of the file.
b) If sanitization has failed, then the file is not skipped (optimization is not applied).
c) If sanitization is successful, then the date from the first log entry is confronted with the specified time search criteria. If the date is after the time frame of the criteria, then the file is skipped (optimization is applied).
d) The sanitization (see yellow box below) is applied on the last 64 kB of the file.
e) If sanitization has failed, then the file is not skipped (optimization is not applied).
f) If sanitization is successful, then the date from the last log entry is confronted with the specified time search criteria. If the date is before the time frame of the criteria, then the file is skipped (optimization is applied).
g) In case of compressed files (such as .gz or .bz2):
    i) If file is local only steps a), b) and c) are performed because uncompressing the whole file would be too time consuming.
    ii) If file is remote, it is regularly processed, but when file entries stop matching the time search criteria, then searching is aborted (effectively this is very similar to the approach for local files). Please note that in cases of remote searching, aborting is performed even if there are multiple time search criteria which prevent execution of the rest of the optimization logic.
h) If the file is remote and uncompressed and it was not skipped in the previous steps, then a narrow-down procedure is performed. The procedure uses effective divide and conquer approach to narrow down searching to a part of the file which could possibly match the time search criteria. The rest of the file is excluded from the search process.

Search optimization **Sanitization Procedure**:

1) Data from the file beginning or the end is split to log entries in accordance with the log entry separation (delimiter) and the date from each entry is extracted in accordance with the date/time format.
2) If the date could not be extracted from more than 2 log entries, then sanitization fails.
3) If the collected data contains less than 5 log entries from which the date was successfully extracted, then sanitization fails.
4) If dates extracted from log entries are not continuous (given log entry has an earlier date than the date present in the previous entry), then sanitization fails.
5) Otherwise, sanitization is successful.

2) Why and when should search optimization be enabled/disabled?

The optimization is enabled by default, because it provides a significant performance boost in most of the typical cases. For example, when the user has a directory with log files from some server (e.g. Apache HTTP or Tomcat) from the past six months and he or she wants to search for an occurrence of a specific error in a three days' time frame, then search optimization will make the searching process significantly shorter (only the file which can contain the dates of the three days' time frame will be searched).

However, the user has to be aware that, the search optimization is performed at some cost. I.e., fetching the data from the beginning and the end of a file and analyzing it introduces some overhead. It is not significant but is still present. Therefore, if for example the user has a case of 200 small files (just above 1000 kB) containing data from the past week and he or she wants to search the files in the context of the last six days, then search optimization is not advised. The majority of files will not be skipped by the optimization (the user is searching six of the seven last days) and the optimization overhead will be relatively high when searching through small files.

In the case of remote files, the optimization includes the narrow-down procedure, which can give significant performance boost when there are big files. For example, if the user has a 300 MB file and he specifies a date criteria covering 1 MB of the file, then the narrow-down procedure is able to decrease the duration of searching process from tens of seconds to several seconds. It is especially visible, when the file part matching the date criteria is quite late in the file and searching process cannot be aborted earlier (please see step g) ii) in point 1) above).

3) What can go wrong during search optimization and what's the remedy?

Typical problems that can occur during search optimization are related to the sanitization procedure (see yellow box above). They are as follows:

- The date format of data source is not appropriate, e.g. the locale is wrong, which does not permit extracting the date from more than 2 log entries (2nd sanitization step). This results in displaying relevant warnings. To fix the problem, the user should modify the data format accordingly.
- Log entry separation (delimiter) is not appropriate. It prevents extracting at least 5 log entries from the fetched data (3rd sanitization step). This results in the relevant warning being displayed. To fix the problem, the user should modify the log entry separation accordingly.
- Dates in given the file are not continuous which causes sanitization to fail (step 4). If there is some error in the log file which results in non-continuous dates, then the user may consider fixing this manually.

5.8.3 Dynamic Monitoring Discovery

Dynamic Tail is a feature available since the 3.3.0 release. By default, Dynamic Tail is enabled and it can be disabled by changing the "Enable Dynamic Discovery" option in the preferences dialog. The option can be found within the "Tail (Monitor)" box on the "Search/Monitor" page. If Dynamic Tail is enabled ("Enable Dynamic Discovery" option is selected), then Retrospective dynamically includes new appearing files/containers in the monitoring process. If Dynamic Tail is disabled, Retrospective ignores any new files/containers. Dynamic Tail is mostly used for data source specified by "Filter Matching Selection", because only these data sources can be evaluated to different file/container sets (e.g. a new file is added to the data source directory or a new container matching the data source filter is started). However, Dynamic Tail also supports Item Selection. When the path pointed by Item Selection does not exist, Dynamic Tail waits until it appears and then proceeds with the most appropriate monitoring strategy. Whether Dynamic Tail is enabled or not, monitoring is always dependent on the "File start offset" / "Container start offset" parameter (Amount of existing data to be delivered when monitoring is started. This value corresponds to the -c option of the Linux tail command for files and to –n tail option for containers) available in the preferences.

If Dynamic Tail is **disabled**, then monitoring of data source specified by Item Selection and "Filter Matching Selection" is performed in the following way:

1. If data source path does not exist or does not contain any files/containers, stop monitoring. Otherwise proceed to step 2.
2. Evaluate data source path to a file/container set
3. Spawn separate TailTask for each file/container and then:
    a. In the first TailTask execution fetch desired number of bytes/lines (see options "File start offset" and "Container start offset" on the Search/Monitoring preference page) from the file's end
    b. Repeat infinitely: Check if file or container log stream was enlarged and, in such cases, fetch new data which has appeared in the file or container log stream

If Dynamic Tail is **enabled**, then monitoring of data source specified by Individual File Selection is performed in the following way:

1. If data source path does not exist, then wait until it appears. When it appears, proceed to step 2.
2. Evaluate data source path. If data source path is a file, proceed to step 3. If it is a directory, proceed to step 2 of Dynamic Tail monitoring for data source specified by "Filter Matching Selection" and pass there all files present in the directory.
3. Spawn TailTask for the file evaluated from the data source path and then:
    a. In the first TailTask execution fetch "File start offset" bytes / "Container start offset" lines from the file's end
    b. Repeat infinitely: Check if file was enlarged and, in such cases, fetch new data which has appeared in the file

If Dynamic Tail is **enabled**, then monitoring of data source specified by "Filter Matching Selection" or by Individual Container Selection is performed in the following way:

1. Evaluate data source path to a file/container set
2. For each file/container fetch "File start offset" bytes / "Container start offset" lines from the end
3. Repeat infinitely: Evaluate data source path to a file/container set. Check for changes and perform the following:
    a. If a new file/container appeared, spawn separate TailTask which fetches its whole content (whole log stream in case of container)
    b. If a file was renamed, ignore it as its content was previously monitored
    c. If a file/container was enlarged, spawn separate TailTask which fetches new data that has appeared in the file

Note

Dynamic Tail is perfectly capable of handling log rotation. If Retrospective monitors directory in which log rotation is performed, it is properly detected as file renaming and no additional data, besides the actual new log content, is fetched.

Dynamic Tail is a lot more efficient for monitoring "Filter Matching Selection" data sources because it spawns TailTasks on demand and does not unnecessarily consume resources of local machine and remote servers.

Note

When Dynamic Tail is **enabled**, then in steps 3a and 3c of monitoring data source specified by "Filter Matching Selection" (or by Individual Container Selection), before TailTask is spawned, the following file / container log stream verification procedure is performed.

Fetch initial file or container log stream part and then check if log entry separation is present there.

a. If it is not present, then optionally run Autofind (see 3.4.3.1 Autofind) to detect correct encoding, log entry separation and date format. This can be switched on  or off  by the user through a toggle button that appears above the start button on individual search/monitoring tabs.
b. If it is present and the source has some date format configured, then split the initial file or container log stream part to log entries and check if date can be parsed in each entry. If date cannot be parsed in at least one entry, display a warning.

5.9 Result Snapshots

The **Result Snapshots** view lists all result data snapshots the user wants to keep for later use. This view can be opened by selecting View → Result Snapshots, clicking the icon or by using the [Ctrl] + [N] keys combination.

The view offers basic functionality through the top located toolbar and by a context menu that appears when the right mouse button is activated. The search overlay enables you to find desired snapshots. The search text field accepts wildcards ('*' and '?') but ignores regular expressions.

A **snapshot** is an immutable copy of collected result data. It can be created within the search/monitoring tab through the button, located top of the result table. In order to be able to create a snapshot, the result data must have been collected (fetched) while the disk storage option was enabled (see Result Options page within the preferences dialog).

When opening a snapshot through the button, its data is copied and a search tab is created, based on the duplicated data. Thus, we guarantee that a snapshot remains unchanged throughout its lifetime.

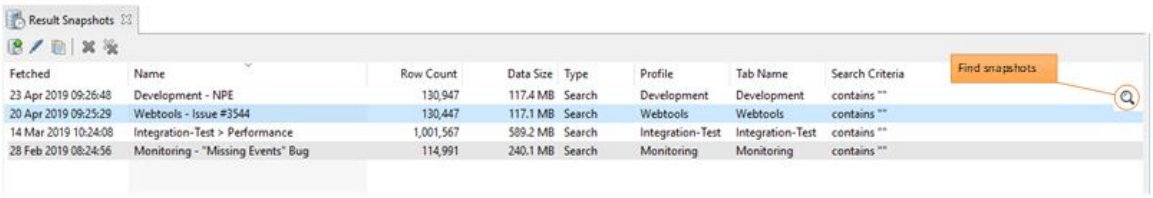

It's important to assiduously manage your snapshots because they may occupy a huge amount of disk space when created and preserved carelessly. Try to create snapshots only if there's a considerable chance you'll need them in the future without being able to collect the same data at given time again. Make sure to remove outdated snapshots.

5.10 Bookmarks View

The ***Bookmarks*** view allows you to access previously bookmarked search criteria. This view can be opened by selecting View → Bookmarks, clicking the icon or by using the [Ctrl] + [B] keys combination.

The search overlay enables you to find desired bookmarks. The search text field accepts wildcards ('*' and '?') but ignores regular expressions.

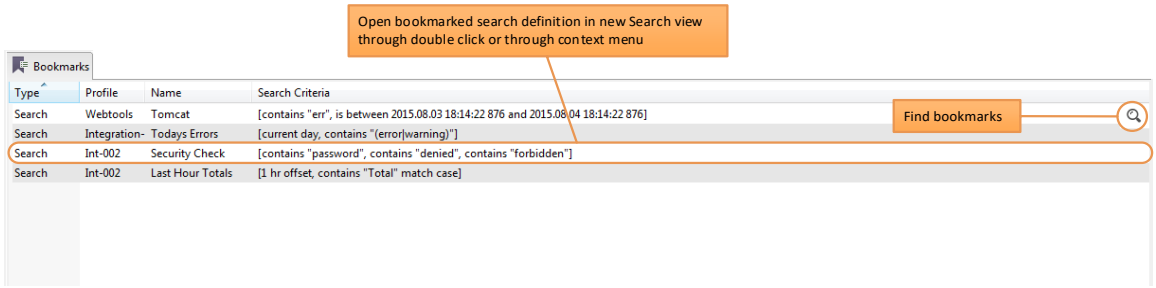

5.11 History View

The ***History*** view allows viewing the search criteria of the previously triggered search actions and opening a search tab with exactly the same search criteria. This view can be opened by selecting View → History, clicking the icon or by using the [Ctrl] + [H] keys combination.

The search overlay enables you to find desired history entries. The search text field accepts wildcards ('*' and '?') but ignores regular expressions.

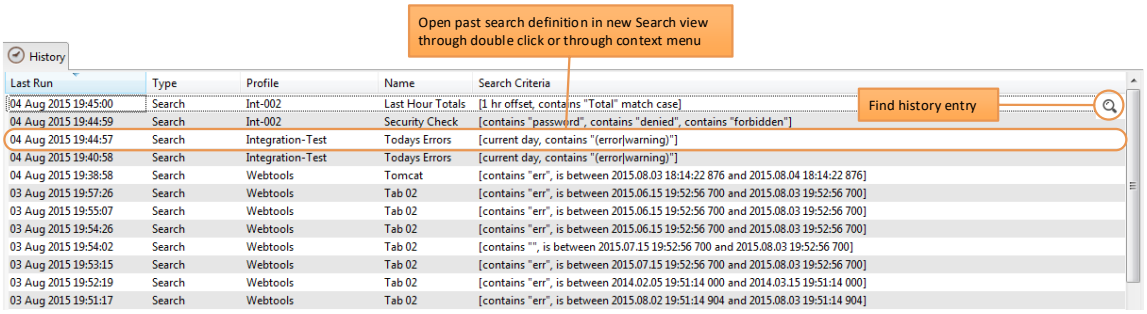

5.12 Status View

The ***Status*** view allows viewing of the application status information stored in Retrospective log file. This view can be opened by selecting View → Status, clicking the icon or by using the [Ctrl] + [S] keys combination.

The search overlay enables you to find desired status entries. The search text field accepts wildcards ('*' and '?') but ignores regular expressions.

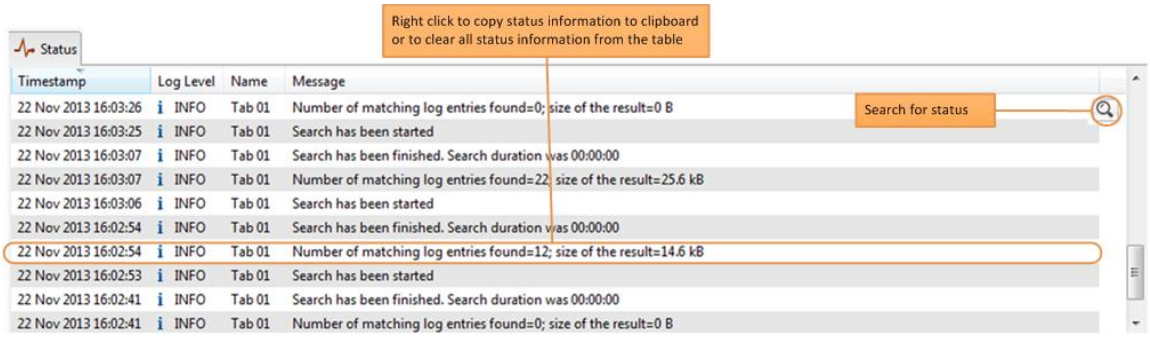

6 RETROSPECTIVE QUERY LANGUAGE (RQL)

The collected data can be filtered with the feature rich Retrospective Query Language (RQL). A query consists of one or multiple terms and groups. Terms and groups can be combined with Boolean operators to form a more complex query.

6.1 Terms

There are two types of terms: single terms and phrases. Phrases are surrounded by double quotation marks.

6.2 Columns

By default, the entered terms and phrases are used to search for content found in the "Data" column. You can however filter any other column by typing the column name followed by an operator (=, <, <=, >, >=) and then the term you are looking for.

Operator	Description
=	The **equal** operator tests if the column value and the filter term are equal. The term may contain wildcards, even if the column is of data type "Number".
<	The **less than** operator tests if the column value is less than the filter term. The term must not contain wildcards.
<=	The **less than or equal** operator tests if the column value is less than or equal to the filter term. The term must not contain wildcards.
>	The **greater than** operator tests if the column value is greater than the filter term. The term must not contain wildcards.
>=	The **greater than or equal** operator tests if the column value is greater than or equal to the filter term. The term must not contain wildcards.

Note

When using one of the bounding operators (<, <=, >, >=), the filter result may be different depending on whether the data type of the column is "Number" or "Text".

When the **data type** is "**Text**", the column values and the filter term are compared lexicographically, same as non-number text. If the first digits of the value and the filter term match, the second digits are compared and so on. Therefore, 10 is considered less than 2, 11 is considered greater than 1 but also greater than 100.

When the **data type** is "**Number**", the filter term will be interpreted as number even if it is quoted (enclosed in double-quotes).

Examples:

`Path=*trace*`	Filters log entries that contain "trace" in their path.
`Code=?`	Filters log entries where the column "Code" contains exactly one character.
`Name=*?*`	Filters log entries where the column "Name" is not empty.
`-ClientID=x`	Filters log entries where the value of the column "ClientID" is not equal to "x".
`!URL=https*`	Filters log entries where the value of the column "URL" does not start with "https".
`Amount<=10`	Filters log entries where the value of the column "Amount" is less or equal to 10.
`NOT Amount>10`	Filters log entries where the value of the column "Amount" is not greater than 10. This expression produces the same result as `Amount<=10`.

Note

If a column name contains one or several spaces, it cannot be used for defining an RQL term. Therefore try to choose column names that don't contain spaces.

6.2.1 Log Level Column

When the RQL query contains a Level column term together with a bounding operator (`<`, `<=`, `>`, `>=`), Retrospective identifies matching result entries by checking the **priority** of their log level. The priority is also referred to as **severity**.

Examples:

`Level<=info`	Filters log entries that contain a log level with a priority that is less or equal to the priority of the log level INFO, provided that the default log level configuration is active. In this example, "info" matches the "INFO" pattern associated with the identically named log level INFO.
`Level>=warning`	Filters log entries that contain a log level with a priority that is greater or equal to the priority of the log level WARN, provided that the default log level configuration is active. In this example, "warning" matches the "WARNING" pattern associated with the log level WARN.

You can change the priority of log levels by changing their position inside the "Log Level Definitions" preferences page.

6.3 Wildcards

RQL supports single and multiple character wildcards within single terms and phrases. Wildcards are special characters that represent unknown characters in a text. They're handy for locating similar but not identical sections.

Wildcard Character	Description
?	Matches a single character at a specific position. The single character wildcard search looks for terms that match with the single character replaced. For example, to search for "text" or "test", you can use the term: `te?t`
*	Matches any number of characters. A multiple character wildcard search looks for 0 or more characters. For example, to search for "test", "tests" or "tester", you can use the term: `test*`

Since wildcard characters are interpreted in a special manner, you have to escape them with a backslash (\) if you want to match them explicitly. If for example you want to match asterisk "*", you have to type `"\*"`.

If you want to explicitly match wildcards inside a phrase (a term surrounded by double quotation marks), you'll have to escape them with two backslashes (\\).

6.4 Boolean Operators

Boolean operators allow terms to be combined through logic operators. Retrospective supports AND, OR and NOT. Boolean operators must all be uppercase in order to taken into account.

Operator	Alternatives	Description

OR	||	The **OR** operator is the **default operator.** If there's no Boolean operator between two terms, the OR operator is used. The OR operator links two terms and finds a matching log entry if it matches either of the terms.
		To search for log entries that contain either "out of date" or just "outdated", use one of the following queries:
		`"out of date" OR outdated`
		`"out of date" || outdated`
		Since OR is the default operator, the same can be written as follows:
		`"out of date" outdated`

AND	&& / +	The AND operator matches log entries that contain both.
		To search for log entries that contain "HTTP" and " 500 ", use one of the queries below:
		`http AND " 500 "`
		`http && " 500 "`
		`http +" 500 "`

NOT	! / -	The NOT operator excludes log entries that contain the term after NOT.
		To search for log entries that contain "HTTP" but not " 200 ", use one of the queries below:
		`http NOT " 200 "`
		`http !" 200 "`
		`http -" 200 "`
		To search for log entries that contain "purchase" or do not contain "action", use one of the following queries:
		`purchase (NOT action)`
		`purchase (!action)`
		`purchase (-action)`

6.5 Grouping

You can use parentheses to group clauses to form sub queries. This can be very useful if you want to control the Boolean logic for a query. Entire groups may be excluded by placing "NOT", "!" or "-" in front of it.

To search for log entries that contain either of the terms "out of date" or "outdated", together with the term "message", use the following query:

```
("out of date" OR outdated) AND message
```

Grouping clauses can also be used for columns. If you want to search for log entries where the column "Code" contains either of the values "A", "M" or "X", your query could look as follows:

```
Path=(A OR M OR X)
```

or even simpler...

```
Path=(A M X)
```

6.6 Escaping Special Characters

Special characters that are part of the RQL syntax need to be escaped if you want to explicitly use them for filtering log content. RQL currently uses the following special characters:

```
+ - && || ! ( ) " * ? \ = < >
```

To escape these character, use the backslash "\" in front of the character. If you enclose the search term in double quotation marks, no escaping is needed. The exceptions to this rule are the double quotation mark itself, the wildcard characters "*" and "?" but also the backslash "\".

- For example to search for phrase that contains double quotation marks, use something like this:

  ```
  "Retrospective \"Release 6.0\""
  ```

- To explicitly match wildcard characters enclosed in double quotation marks, escape them by a backslash "\" as follows:

  ```
  "five\* solution"
  ```

  ```
  "what happened here\?"
  ```

- To explicitly filter log content that contains a backslash, your search term could be written as follows:

  ```
  "1\\2"
  ```

7 LOG TIME SYNCHRONIZATION

7.1 The challenge

A software system often consists of many individual programs that work together and perform complex tasks by exchanging messages. The programs may run on distinct servers and within containerized subsystem such as Docker and Kubernetes. The servers and containers involved can be in different locations and even in different time zones.

In the ideal case, all participating programs write log data with a timestamp that also contains the time zone. This allows Retrospective to automatically convert the timestamps to local date/time and present them in the Date/Time column of the result table. No further time synchronization would be needed for viewing and further analyze the log data, you should obtain a chronological correct picture from what happened in the whole system.

Unfortunately, log data is often written with a timestamp that does not include the time zone. When collecting log data from distant sources, individual entries cannot be easily correlated. Therefore, Retrospective needed to provide a mechanism to automatically synchronize log data if requested by the user.

7.2 SSH Hosts

Within the Host Manager, the following options let you control the time synchronization for individual SSH hosts:

- Time zone fallback
- Adjust time offset

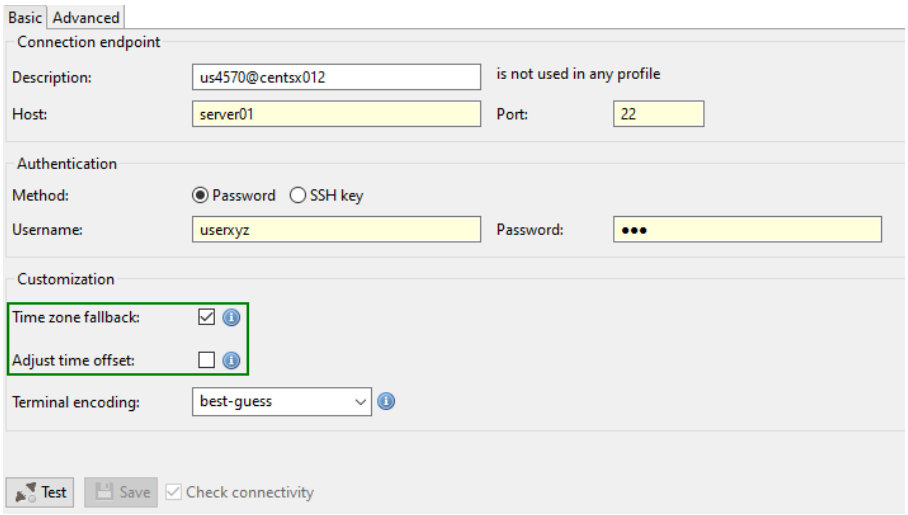

7.2.1 Time Zone Fallback

When Retrospective collects log entries, it automatically converts their timestamps to the local time using the time zone information found in the log entry timestamps. This makes it easy to correlate log entries retrieved from servers lying in different time zones. In many cases however, applications write log entries without time zone information.

The "Time zone fallback" option determines if the time zone of the host shall be used as a fallback for converting log entry timestamps that don't contain a time zone themselves. If this checkbox is selected, data sources will be treated as follows in case no time zone is contained in their "Date/Time Format":

- The time offset between the local computer and the remote host is computed using their time zones.
- All time search criteria passed to the remote host will be shifted by the computed time difference.
- The timestamps of retrieved log entries will be converted to the local time using the computed time difference.

The content of the result table Date/Time column will show the log entry timestamps converted to the local time. The tooltip of that column will however also report the original date/time in the format it was written to the log file.

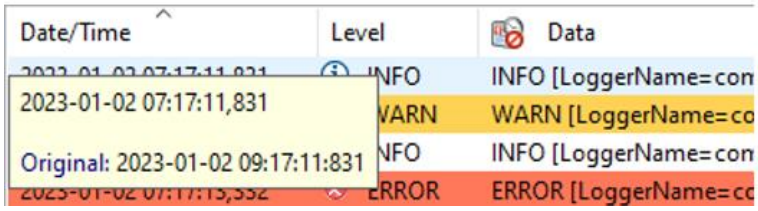

The original date/time can also be displayed permanently inside the result table Data column by clicking on the calendar button that appears in the header of that column.

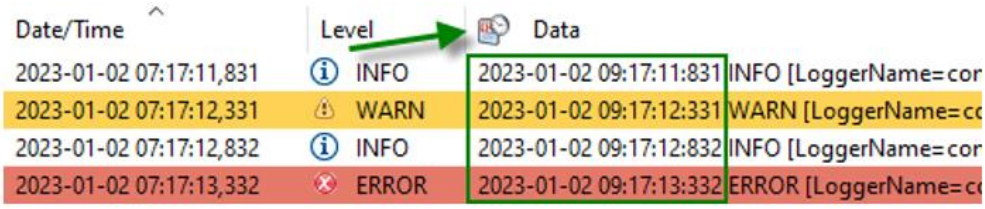

Note

This option has no effect when the remote host has the same time zone as the local computer.

7.2.2 Adjust Time Offset

If you select the "Adjust time offset" checkbox, Retrospective will automatically adjust log entry dates coming from the remote host and the time search criteria passed to it. The goal of adjustment is to make the remote host time match the time on the local computer.

If for example the time on the host is 638 milliseconds in the future compared to the local computer, the timestamps of all log entries found on that host will be decreased by 638 milliseconds and all time search criteria passed to it will be increased by 638 milliseconds. Due to this, you may consider the time between the remote host and the local computer to be synchronized.

The content of the result table Date/Time column will show the adjusted log entry timestamps. The tooltip of that column will however also report the original date/time in the format it was written to the log file.

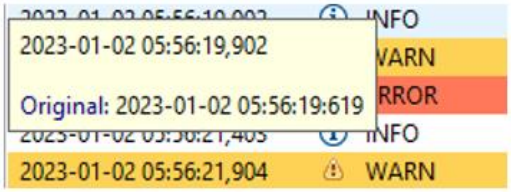

The original date/time can also be displayed permanently inside the result table Data column by clicking on the calendar button that appears in the header of that column.

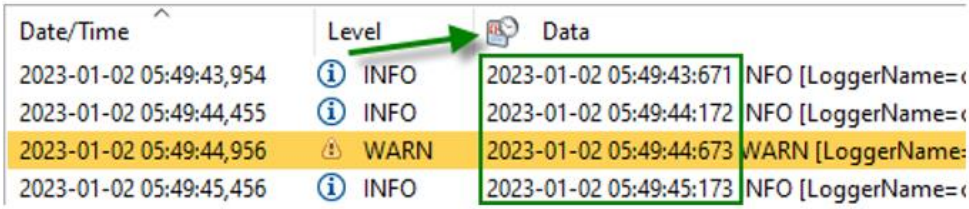

7.3 Containers

Within the preferences page "Data Sources > Container", the following options let you control the time synchronization for containerized sub-systems.

- Docker time zone fallback
- Kubernetes time zone fallback

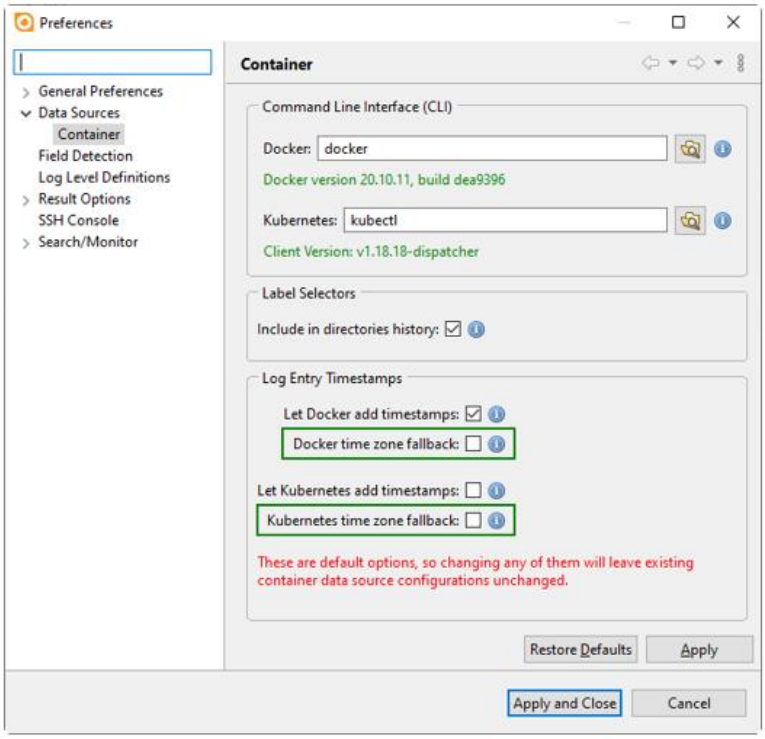

They both determine if the container time zone shall be used as a fallback for converting log entry timestamps that don't contain a time zone themselves. This options however have no effect when the container has the same time zone as the local computer.

These are default options that are considered when choosing new Docker or Kubernetes data sources. The option may be changed for individual data sources in the 'Configure Data Source(s)' dialog within the 'Date/Time Recognition' page.

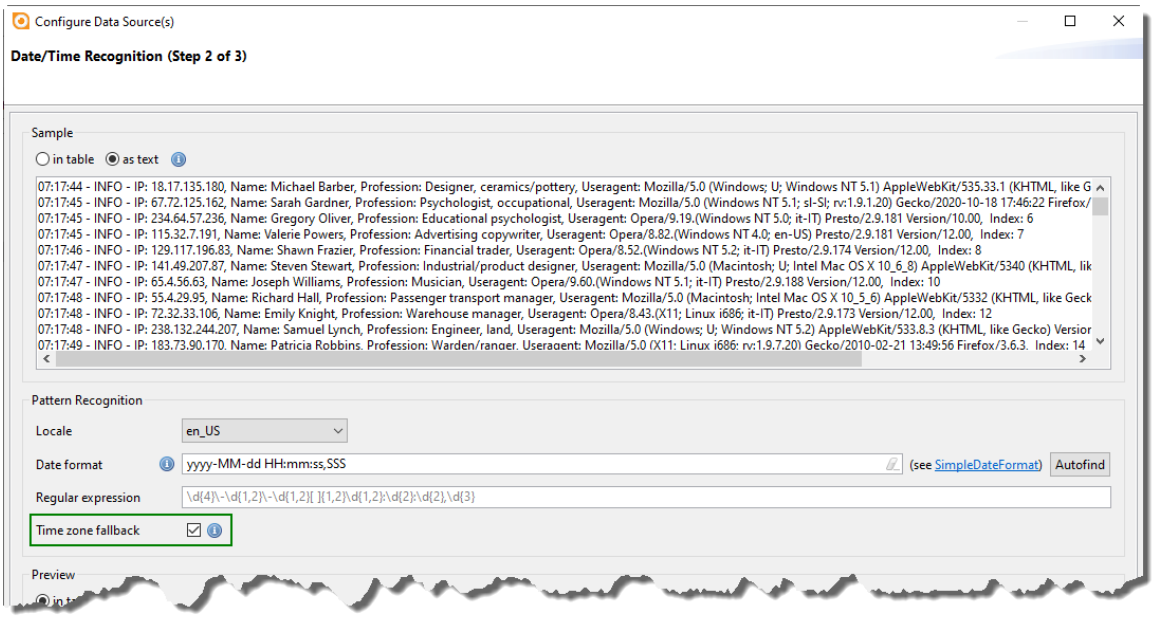

8 SSH CONSOLE

Retrospective contains a fully featured SSH client, named SSH Console, that uses the secure shell protocol to connect to remote hosts that have previously been defined in the Host Manager. The SSH Console primary emulates control sequences from **xterm** but supports also a subset of **vt100**.

8.1 Console creation

Individual SSH Consoles tabs can be created/started in different ways:

- Through the main menu "SSH Console/New/<host name>"
- Through the context menu item "New SSH Console" that appears when you right click on a host within the Host Manager.
- Through the context menu item "Tail File in new SSH Console" that appears when you right click a result entry from a remote computer within the result table.

When a new SSH Console is started, the SSH connection is established straight away using the credentials configured for the specific host. There's no need to type the username and password again.

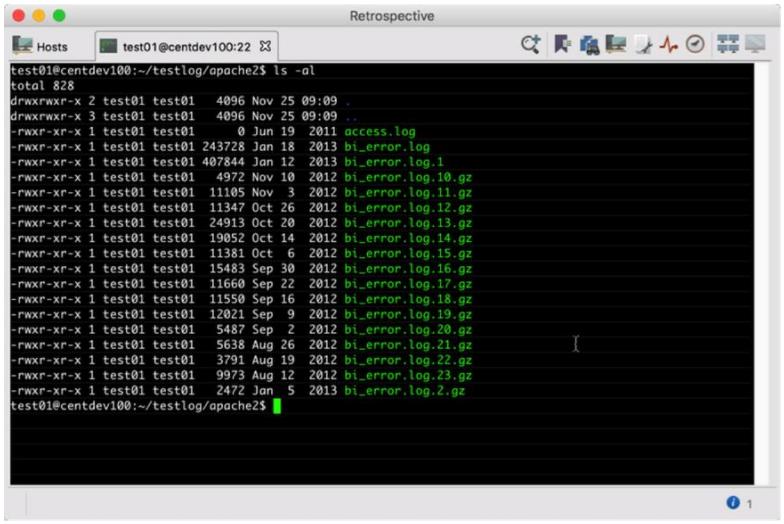

8.2 Console customization

The general behavior and look of the SSH Console, including colors, font and cursor, can be customized within the preferences dialog on the "SSH Console" page. The terminal encoding is defined individually for each host within the Host Manager (see section 3.3.7.3 Terminal Encoding). However, you usually do not have to even think about it because the default "best-guess" behavior automatically detects the encoding of the remote terminal by analyzing some of the LC_* environment variables.

Tabs of SSH Console can be exploded/imploded as any other tabs in Retrospective which allows you to adjust them to the needs of given program that is executed in the SSH Console. For example, when editing a big configuration file with **nano**, it is good to have a lot of vertical space. Exploding of SSH Console tabs is also very convenient when you need to make some action on a remote host that calls for log monitoring. For example, you need to deploy a new version of your application, so in the first tab, you start monitoring of the application server log file and in the second tab you open the SSH Console and make the deployment. Thanks to exploding, both tabs can be placed side by side, so any log changes are instantly visibly after deployment.

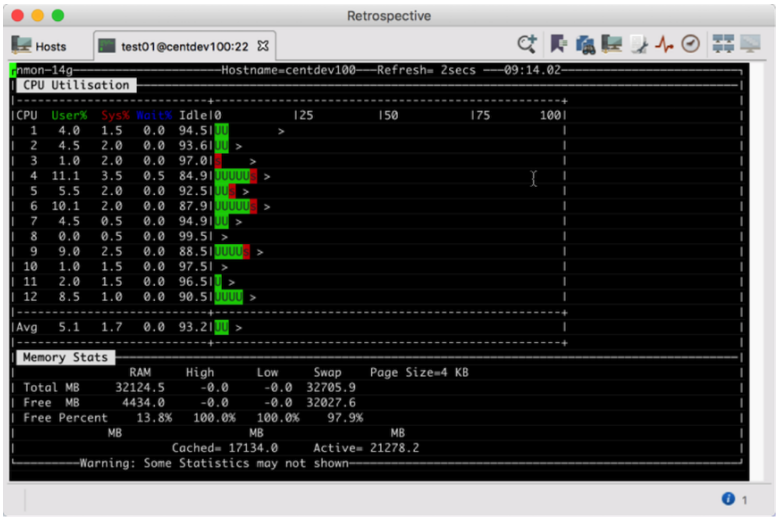

8.3 Console handy features

SSH Console has several more features that are worth mentioning.

Firstly, it supports Alternate Screen Buffer similarly to other feature-rich SSH clients such as PuTTY or iterm2 on MacOS. Thanks to this, the screen contents are saved between opening of typical console programs such as **nano**, **vi**, **nmon** or **midnight commander**. It is very handy when you list the contents of a directory, then open some file and after closing it you'd like to recall the directory contents.

It is very common to copy some text in and out of a console, thus other features provided by the SSH Console are: (i) quick copy/paste mode that allows you to copy/paste text with the use of left/right mouse buttons; (ii) keyboard shortcuts that allow you to copy/paste with the well-known shortcuts: ctrl+c/ctrl+v (also the combination ctrl+a, selecting all terminal contents, is supported). Both features can be enabled in the "SSH Console" page of the Retrospective preferences.

It is also definitely valuable that SSH Console is capable of supporting Alternate Character Set which ensures that line drawing characters in programs such as **nmon** or **midnight commander** are properly displayed even when single byte encoding (e.g. **latin1**, **latin2**) is used in the terminal.

8.4 Target command modification inside Console

It has to be highlighted that SSH Console does not support Target command modification (see 3.3.9 Target command modification). The shell started in SSH Console simply uses the default execution priority and does not perform any identity change. If you want to work in a shell with different execution priority or changed identity, you can spawn a new shell with the use of the following commands:

- execution priority change: `nice -n <adjustment> $SHELL` ; e.g. `nice -n 7 $SHELL`
- identity change: `sudo -u <username> -i` (or just `sudo -i` if sudoing to root) ; e.g. `sudo -u foobar -i`

To change both execution priority and identity, simply change the identity first, and then perform the execution priority change.

9 MANAGING WORKSPACE

Retrospective has many features which are intended to make it more convenient in day to day use. The following chapter contains descriptions of various useful procedures related to Retrospective workspace.

9.1 Rename Tabs

Retrospective provides the possibility of changing default search tab names, so they are more meaningful. Tab names are included in history entries and in status information, thus it is better to rename a tab before performing traceable actions. When you create a bookmark from your favorite search definition, Retrospective proposes the tab name as default bookmark name but gives you the choice to change it.

This procedure describes how to rename search tabs.

1. Double-click or right-click desired search tab and select ***Rename Tab*** from the context menu.
2. Provide new name and press the [Enter] key.

9.2 Change Tab Position

The position of individual tabs can be changed at any time. This feature is especially useful prior to saving the application state as a view. Retrospective guarantees that all tabs are placed at the same position when the view is reloaded again (see 6.6 Save/Reload/Manage Application State).

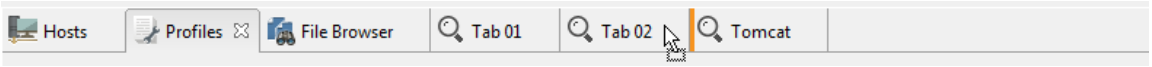

1. Press the left mouse button while the mouse pointer is on a tab.
2. Keep the mouse button pressed and drag the tab to the new position.
3. When an orange bar is shown at the desired new position, releasing the mouse button drops the tab to the new position.

9.3 Disjoin Tabs

One of the advantages of the tabbed user interface is the ability to freely disjoin any currently opened tab and have it opened in a separate window.

- To do that simply drag and drop desired tab outside the Retrospective window.

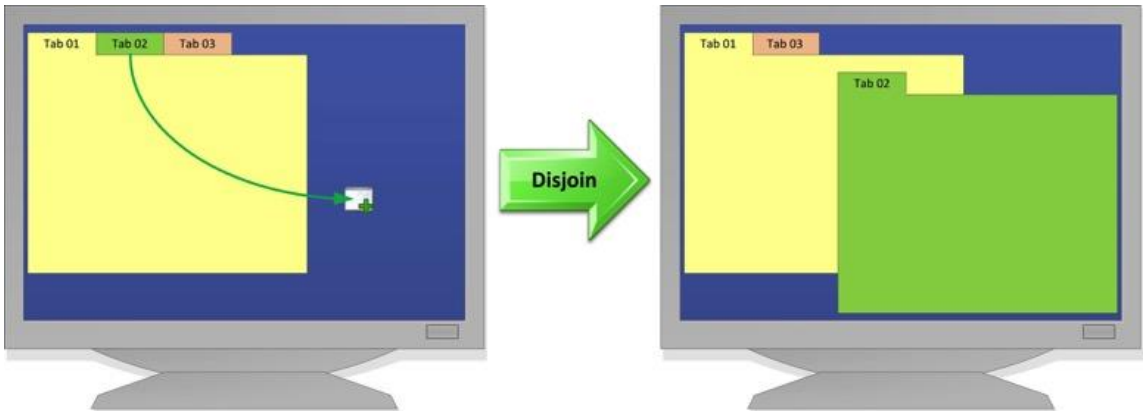

9.4 Explode Tabs

Retrospective allows disjoining and opening currently opened tabs as separate windows with a single click. To do that simply click the icon, select View → Explode all Tabs or use the [Ctrl] + [E] keys combination and every tab will be replaced with a separate window and the screen will be divided into as many equal parts as the number of the tabs that were opened.

Note

This feature works best when using high resolution monitors or multiple screens setups.

Retrospective follows the strategy described below when windows are "exploded":

1. A minimum window size of 800 x 500 (width x height) pixel is observed.
2. The tabs are equally distributed between the available monitors; no distinction is made depending on their screen size.
3. For each monitor the following rules are observed:
    a. As long as the minimal size can be observed for all windows, they are laid out side by side in order to occupy the entire screen. Depending on the number of tabs and on the screen format and size, the windows will be placed one beside the other, one on top of the other, or all within a grid.
    b. If the minimal size cannot be observed for all windows, they are laid out as cascading windows.

9.5 Implode Windows

Retrospective also allows imploding windows (previously exploded tabs).

- To implode, simply click the icon, select View → Implode all Windows or use the [Ctrl] + [Shift] + [E] keys combination and all currently opened Retrospective windows will emerge as tabs in a single Retrospective window.

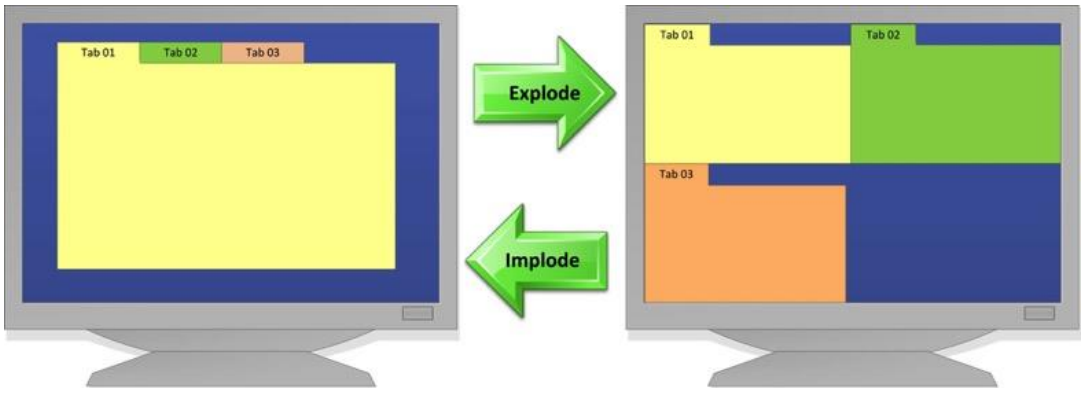

9.6 Save/Reload/Manage Application State

Retrospective can **automatically store** the application state when the program is terminated and reload once it is re-launched. This feature can be enabled/disabled on the General Preferences page within the preferences dialog. Simply select/deselect the checkbox "Restore view from last session at startup" to tell Retrospective how to respond. If the feature is enabled,

the number and position of the window(s) and tabs are restored to the state they had when the last program session was terminated. This includes restoring the search criteria defined for individual search tabs together with previously collected data.

Retrospective also lets you **manually save** the state of the application at any time.

1. To save the current state of the application simply select File → Save View.
2. Provide the view name.
3. Click the [OK] button to finalize.

While working with Retrospective, you can always reload the state of the application that was previously saved as a view. Be aware however that in such a case all current windows and tabs are closed and entirely replaced by the new application state.

1. To load the previously saved application state simply select File → Load View
2. Then select the menu item that corresponds to the name of the view

Manually saved application states can be renamed or deleted.

1. To rename/delete manually saved application states simply select File → Manage Views
2. To delete one or multiple saved views, select them in the list and click the [Delete] button.
3. To rename a saved view you can select the view and edit the name in the list.
4. To confirm your changes, click the [OK] button. To discard your changes, click the [Cancel] button.

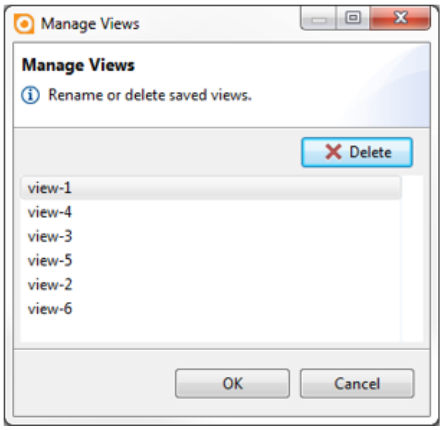

10 TROUBLESHOOTING AND BEST PRACTICES

This chapter covers potential problems with configuration and usage of Retrospective as well as guidelines how to use it effectively.

10.1 Troubleshooting

This table contains various problems which may occur when using Retrospective.

Case	Solution
Cannot access remote server	• Make sure that you have correct host IP address, SSH port number and user login credentials. • Make sure that the remote host you are trying to connect to has the SSH service enabled. • Check that the connection is not blocked by the firewall.
Data source not found	Make sure that the data source on the remote host has not been deleted or renamed.
Search does not return any results	Check and refine the search criteria.
Inconsistent filter information.	Make sure that there are no two text filters with the same option selected (`Starts with`/ `Ends with`) but different phrases.
Large data limit reached	Add additional search criteria to limit the search results.
Cannot load Sources	Make sure that the data sources defined in profile have not been deleted.
Cannot load View	Data storing user profile has been corrupted or deleted.
Not all profiles in the loaded view exist	Some data sources have been removed from original location.
Drag 'n Drop of sources onto search/tail tab was not successful	Repeat the operation or try adding data sources in an alternative way (see 3.4 Profile Manager).
Search or monitoring attempt on empty profile	Add data sources to given profile.
Cannot save results	Problem with the I/O system. Try repeating the action again.

10.2Best practices

This chapter covers best practices.

10.2.1 Bookmarking

Retrospective provides a bookmarking option allowing storing the exact definition of search criteria. If you frequently search for given phrases, define a bookmark which will allow you to execute certain searches more conveniently (see 5.10 Bookmarks View).

10.2.2 Pinning tabs

Retrospective allows pinning individual tabs. Pinned tabs cannot be closed unless they are un-pinned or moved to a different position. For pinning a tab, simply right-click its label and select ***Pin this tab*** from the context menu.

10.2.3 Application state saving

If you have to follow a given number of data sources on daily basis, saving the given state of the application (opened ***Search*** tabs) will save time whenever Retrospective is started.

10.2.4 Display options

Retrospective features imploding and exploding functionality (see 10.4 Explode Tabs and 10.5 Implode Windows) allowing the effective use of high resolution and multiple screens setups. If you have such a setup you can always implode application tabs and have them displayed as separate windows so you can monitor multiple search results without having to switch between the tabs.

10.2.5 Drag and Drop Data Sources

Retrospective supports drag and drop actions on data sources. It is possible to drag and drop a log file from the local file system onto the profiles tab. Depending on where you drop the selected files and/or directories, one of the following results is obtained.

- When dropped onto the data sources panel, selected files and directories are added to the selected profile (see illustration below).
- When dropped onto a profile name, selected files and directories are added to that profile.
- When dropped on an empty space within the profiles list, a new profile is created and the selected files and directories are added to it.

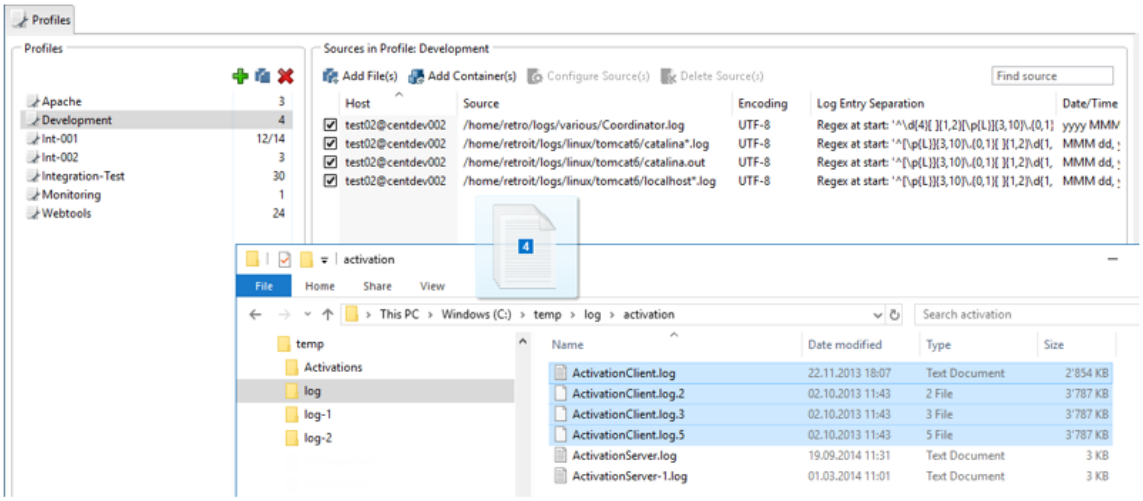

It is also possible to drag and drop existing data source definition from one profile to another.

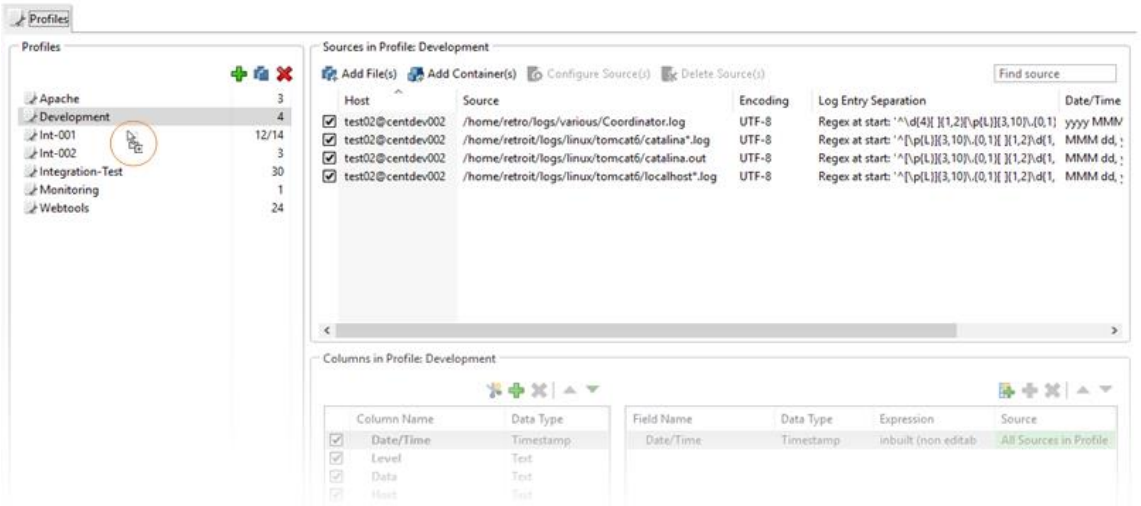

10.2.6 Keyboard shortcuts

The following Retrospective-specific keyboard shortcuts enable you to quickly access certain features.

Windows/Linux	Mac OS	Description
[Ctrl] + [W]	[⌘] + [W]	Closes current tab.
[Ctrl] + [T]	[⌘] + [T]	Opens new Search/Monitor tab.
[Ctrl] + [O]	[⌘] + [O]	Hides/Shows individual search criteria fields.
[Ctrl] + [Arrow Right]	[⌘] + [Arrow Right]	Navigates to next tab.
[Ctrl] + [Arrow Left]	[⌘] + [Arrow Left]	Navigates to previous tab.

Windows/Linux	Mac OS	Description
[Ctrl] + [N]	[⌘] + [N]	Opens the ***Result Snapshots*** view (see 5.9 Result Snapshots).
[Ctrl] + [B]	[⌘] + [B]	Opens the ***Bookmarks*** view (see 5.10 Bookmarks View).
[Ctrl] + [Y]	[⌘] + [Y]	Opens the ***History*** view (see 5.11 History View).
[Ctrl] + [S]	[⌘] + [S]	Opens the ***Status Information*** view (see 5.12 Status View).
[Ctrl] + [P]	[⌘] + [P]	Opens the **Profile Manager** (see 3.4 Profile Manager).
[Ctrl] + [I]	[⌘] + [I]	Opens the **File Browser** view (see 5.5 File Browser).
[Ctrl] + [G]	[⌘] + [G]	Opens the **Container Browser** view (see 5.6 Container Browser)
[Ctrl] + [E]	[⌘] + [E]	Explodes Retrospective tabs (see 10.4 Explode Tabs)
[Ctrl] + [Shift] + [E]	[⌘] + [Shift] + [E]	Implodes Retrospective windows (see 10.5 Implode Windows)
[F1]	[F1]	Opens Retrospective help pages.
[F2]	[F2]	Toggles full screen mode.
[Ctrl] + [R]	[⌘] + [R]	Toggles result entry detail view.
[Ctrl] + [Z]	[⌘] + [Z]	Opens the ***Host Manager*** (see 3.3 Host Manager).
[Ctrl] + [K]	[⌘] + [K]	Shows keyboard shortcuts reference.
[Ctrl] + [Enter]	[⌘] + [Enter]	Start Searching / Monitoring
[Ctrl] + [Shift] + [P]	[⌘] + [Shift] + [P]	Open Preferences Page
Formatting and Highlighting		
[Ctrl] + [Shift] + [F]	[⌘] + [Shift] + [F]	Show formatting menu
[Ctrl] + [Shift] + [H]	[⌘] + [Shift] + [H]	Show highlighting menu
[Ctrl] + [Shift] + [I]	[⌘] + [Shift] + [I]	Instant results highlighting
[Ctrl] + [Shift] + [L]	[⌘] + [Shift] + [L]	Highlighting results by log levels
[Ctrl] + [Shift] + [D]	[⌘] + [Shift] + [D]	Disable results highlighting

11 SSH SUPPORT

Retrospective uses the pure Java implementation of SSH2 of http://www.jcraft.com/jsch/ and https://github.com/hierynomus/sshj/. Supported are the following elements.

11.1 Key exchange

diffie-hellman-group-exchange-sha1
diffie-hellman-group1-sha1
diffie-hellman-group14-sha1
diffie-hellman-group-exchange-sha256
ecdh-sha2-nistp256
ecdh-sha2-nistp384
ecdh-sha2-nistp521

11.2 Cipher

blowfish-cbc
3des-cbc
aes128-cbc
aes192-cbc
aes256-cbc
aes128-ctr
aes192-ctr
aes256-ctr
3des-ctr
arcfour
arcfour128
arcfour256

11.3 MAC Message Authentication Code

hmac-md5
hmac-sha1
hmac-md5-96
hmac-sha1-96

11.4 Key type (Signatures)

Provided by http://www.jcraft.com/jsch/

ssh-rsa
ssh-dss

ecdsa-sha2-nistp256
ecdsa-sha2-nistp384
ecdsa-sha2-nistp521

Provided by https://github.com/hierynomus/sshj/

ssh-ed25519

12 KNOWN ISSUES

- Some users experienced long startup times and problems with the Retrospective storage because their user home is located on a network drive. To change the location of the configuration files to the preferred local location, please follow the instructions in chapter 3.1 Changing Configuration Location.
- Retrospective may terminate unexpectedly without any user notification because the underlying Java runtime environment has not sufficient memory to continue. In the Retrospective home directory a log file of format hs_err_pid*.log may be found in such cases. Please send this file to our support team at support@centeractive.com for further analysis in order to get assistance according to our EULA.
- Retrospective may have problems with file names and directory names which contain wildcard characters (* and ?).
- Due to limitations of Java, currently Retrospective does not support Windows special directory junctions e.g. "Documents and Settings", "Application Data". Such links are displayed as non-readable directories. Additionally, Retrospective can potentially have problems with directory junctions created by the user.
- In order to increase performance of browsing mapped network drives, Retrospective does not resolve Windows links (*.lnk files). Therefore, each such link is displayed as file even if it is a directory. If a user wants to enter a directory (on a mapped network drive) pointed by a link, then the full link path should be placed manually in Directories Panel textbox and then the Enter key should be pressed.
- Retrospective might fail to start if installed in a directory whose path contains certain invalid characters such as %#<>"!. To solve this issue, you have to install Retrospective in a directory whose path does not contain invalid characters.
- Log entry separation appearing several times in a single line is currently not supported for the remote hosts. For example, if you have a line with fields separated by colons and want to treat these fields as separate log entries, then this will only be possible locally, not remotely. We decided to not support such rare cases in order to squeeze more performance out of the searching SSH Script Processing Pipeline.
- In the SSH Console, the Bash keyboard shortcut ALT+F that jumps the cursor forward by one word is not working on Windows and Linux because the ALT+F shortcut is reserved for accessing the File item in the Retrospective Menu. However, ALT+F shortcut works fine on Mac OS. The complimentary shortcut: ALT+B that jumps the cursor backward by one word is working fine on all supported platforms: Windows, Linux and Mac OS.
- On Ubuntu, to enable the icons of context menus, it's necessary to ensure that the GSettings value of the key org.gnome.desktop.interface menus-have-icons is set to true.

13 GLOSSARY AND ABBREVIATIONS

A

Autofind

Autofind is performed by Retrospective in one or several log files in order to determine their encoding, log entry separation (delimiter) and date/time format.

B

Bookmark

Set of search criteria that make it easy to get back to your favorite search or monitoring definition

D

Data Source

One or several local or remote log files. A data source can be an individual log file or a directory that contains one or many log files. A data source can also be expressed as a filter (i.e. '*.log'), which is a text pattern that matches the names of one or several files within a same directory.

H

Host

A computer on which log files are stored

I

IP (Internet Protocol)

Communication protocol used for relaying datagrams based on unique addresses.

L

Log Entry

A log entry is a single record disclosing details from one or more events and incidents. A log entry is sometimes referred to as an event log, event record, log message, log record, or audit record.

P

Port

Port is an application or process specific software construct used for communication purposes.

Profile/Data Profile

User defined set of data sources.

R

Regex

A regular expression is a well-defined set of characters allowing certain text strings to be matched.

RQL

RQL stands for Retrospective Query Language that can be used for filtering collected log entries. A query consists of one or multiple terms and groups. Terms and groups can be combined with Boolean operators (AND, OR, NOT) to form a more complex query.

S

Split Strategy

Definition of how the log files should be split into separate entries.

SSH (Secure Shell)

Networking protocol which ensures secure data connection

T

Tail

Real time processing of data sources which are being constantly updated. Tail is deducted from the Linux tail command that outputs the last part of files. In Retrospective documentation, tailing is mostly referred as log file monitoring.