

IBM Operations Analytics - Log Analysis
Version 1.3.2

User's Guide



IBM Operations Analytics - Log Analysis
Version 1.3.2

User's Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 23.

Edition notice

This edition applies to IBM Operations Analytics - Log Analysis and to all subsequent releases and modifications until otherwise indicated in new editions.

References in content to IBM products, software, programs, services or associated technologies do not imply that they will be available in all countries in which IBM operates. Content, including any plans contained in content, may change at any time at IBM's sole discretion, based on market opportunities or other factors, and is not intended to be a commitment to future content, including product or feature availability, in any way. Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only. Please refer to the developerWorks terms of use for more information.

© Copyright IBM Corporation 2015.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this publication 1

Audience	1
Publications	1
Accessing terminology online.	1
Accessibility	1
Tivoli technical training.	1
Providing feedback	2
Conventions used in this publication	2
Typeface conventions	2

Searching and visualizing data 3

Search UI overview	3
Side bar icons	4
Search UI reference	4
Changing the search time zone	8
Searching data.	8
Query syntax.	10

Search results timeline.	12
List and Grid views	13
Refining search results.	14
Saving a search	15
Saved searches	16
Visualizing data	16
Creating charts and graphs	16
Percentile statistical functions	17
Dashboards	18
Search dashboards	20
Custom Search Dashboards	21

Notices 23

Trademarks	25
Terms and conditions for product documentation.	25
IBM Online Privacy Statement	26
Trademarks	26

About this publication

This guide contains information about how to use IBM® Operations Analytics - Log Analysis.

Audience

This publication is for users of the IBM Operations Analytics - Log Analysis product.

Publications

This section provides information about the IBM Operations Analytics - Log Analysis publications. It describes how to access and order publications.

Accessing terminology online

The IBM Terminology Web site consolidates the terminology from IBM product libraries in one convenient location. You can access the Terminology Web site at the following Web address:

<http://www.ibm.com/software/globalization/terminology>.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. In this release, the IBM Operations Analytics - Log Analysis user interface does not meet all accessibility requirements.

Accessibility features

This information center, and its related publications, are accessibility-enabled. To meet this requirement the user documentation in this information center is provided in HTML and PDF format and descriptive text is provided for all documentation images.

Related accessibility information

You can view the publications for IBM Operations Analytics - Log Analysis in Adobe Portable Document Format (PDF) using the Adobe Reader.

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see the IBM Human Ability and Accessibility Center. The IBM Human Ability and Accessibility Center is at the following web address: <http://www.ibm.com/able> (opens in a new browser window or tab)

Tivoli technical training

For Tivoli® technical training information, refer to the following IBM Tivoli Education Web site at <http://www.ibm.com/software/tivoli/education>.

Providing feedback

We appreciate your comments and ask you to submit your feedback to the IBM Operations Analytics - Log Analysis community.

Conventions used in this publication

This publication uses several conventions for special terms and actions, operating system-dependent commands and paths, and margin graphics.

Typeface conventions

This publication uses the following typeface conventions:

Bold

- Lowercase commands and mixed case commands that are otherwise difficult to distinguish from surrounding text
- Interface controls (check boxes, push buttons, radio buttons, spin buttons, fields, folders, icons, list boxes, items inside list boxes, multicolumn lists, containers, menu choices, menu names, tabs, property sheets), labels (such as **Tip:**, and **Operating system considerations:**)
- Keywords and parameters in text

Italic

- Citations (examples: titles of publications, diskettes, and CDs)
- Words defined in text (example: a nonswitched line is called a *point-to-point line*)
- Emphasis of words and letters (words as words example: "Use the word *that* to introduce a restrictive clause."; letters as letters example: "The LUN address must start with the letter *L*.")
- New terms in text (except in a definition list): a *view* is a frame in a workspace that contains data.
- Variables and values you must provide: ... where *myname* represents....

Monospace

- Examples and code examples
- File names, programming keywords, and other elements that are difficult to distinguish from surrounding text
- Message text and prompts addressed to the user
- Text that the user must type
- Values for arguments or command options

Searching and visualizing data

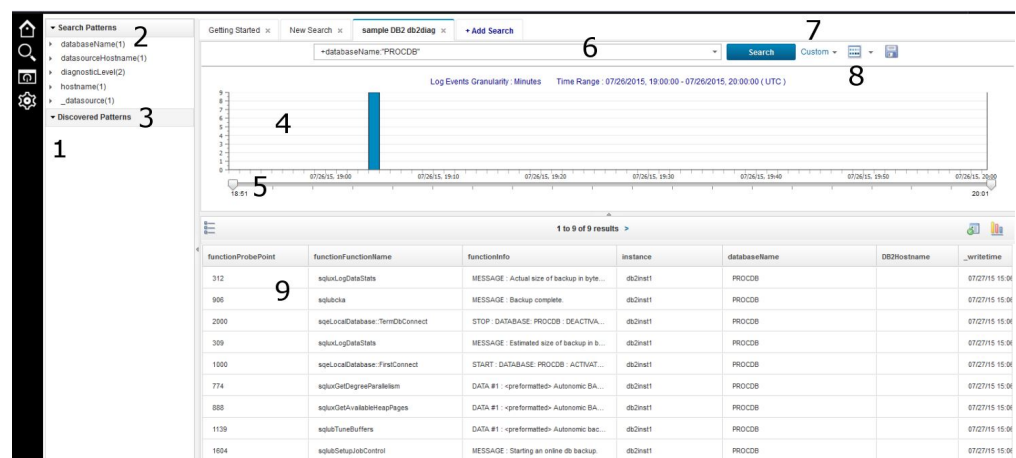
This section outlines how to use the IBM Operations Analytics - Log Analysis Search workspace to search your indexed data and to display this data in charts and dashboards.

To find the root cause of a problem experienced by users such as slowness or a failure, you can search through data such as log files, traces, configuration information, and utilization data. This type of search is iterative because the results for one search might lead to a set of other searches. An example of iterative search is finding the connection timeout in the error logs, which could lead you find the connection pool utilization details.

Search UI overview

Use this topic to help you to get started with the Search UI.

The following screen shot shows the capabilities of the Search UI:



1. Sidebar

Use the UI icons on the side bar to open the Getting Started UI, a saved search, a search dashboard or the Administrative Settings UI.

2. Search Patterns pane

The **Search Patterns** pane lists the fields that are found in your search. To filter the search for a field, click on the field and click **Search**.

3. Discovered Patterns pane

The **Discovered Patterns** pane lists fields and values. To display discovered patterns, you need to configure the source types in the data source.

4. Timeline pane

The **Timeline** pane displays the current search results filtered by time. To drill down to a specific time, click on a bar in the graph.

5. Timeline slider

Use the time line slider icon to narrow and broaden the time period.

6. Search box

Enter search queries in the **Search** field. When you click on a field in the Search or Discovered Patterns pane, the query is displayed in this field.

7 Time filter list

Use the time filter list to filter search results for a specified time range.

8. Data source filter

Use the data source filter icon to filter search results for a specific data source.

9. Table view / Grid view






Use the **List View / Grid View** icon to switch between both view. Use the list view to identify search results quickly. Use the table view to display search results in a tabular format.

Side bar icons

Use the side bar to quickly navigate the user interface.

The following table explains the available icons.

Table 1. . Side bar icons

Icon	Name	Description
	Getting Started icon	Use guided demonstrations and find links to useful information.
	Saved Searches icon	Run saved and example searches.
	Search Dashboards icon	Run custom and sample search dashboards to view charts based on search results.
	Administrative Settings icon	Create and administrate the data model, users, and roles. Display server statistics.
	Manage Alerts icon	 Create and administrate alerts. This feature is only available in the Standard Edition.

Search UI reference

Use this reference topic to help you to navigate the Search user interface (UI).

Buttons and fields on the Search UI

Table 2. Buttons and fields on the Search UI








Button or Field	Name	Description
	Search button and field	Search for a keyword or enter a search query.

Table 2. Buttons and fields on the Search UI (continued)

Button or Field	Name	Description
	Data Sources icon	Filter the search for specific data sources or groups of data sources.
Last 15 Minutes ▾	Time filter icon	Specify a relative or custom time filter.
	Save Quick Search icon	Save the search query and results for later use. To view the saved searches, click the Saved Searches icon on the side bar.
	Columns to be displayed icon	Filter the columns that are displayed in the Grid and Table views.
	Plot chart icon	Select the columns that you want to graph and click the Plot Chart icon.
	List View icon	To open the List View, click this icon.
	Grid View icon	To open the Grid View, click this icon.

Time line graph

Table 3. Time line graph


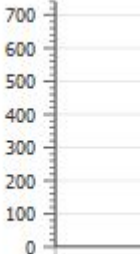
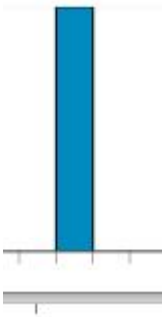






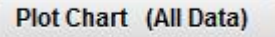
Field, icon or button		Description
	Time line slider icon	Filter the time that is displayed in the timeline graph.
	Y-axis	Shows the number of log events for each bar.
	Bar	Shows the number of log events for the specific time. To view more details, click a bar to drill down to more details about that time range.

Table 3. Time line graph (continued)

Field, icon or button		Description
	Time zone button	To change the time zone, click the time zone button.
	Log Events Granularity	Displays the granularity of the log records that are displayed. The level depends on the time that you have filtered for. You can drill down from years to months to days to hours to minutes and seconds.
	Time Range	Describes the time range that is displayed in the timeline.

Plot Chart editor

Table 4. Buttons and fields on the Plot Chart editor

Button or Field	Name	Description
 Generate Count	Generate Count check box	To generate a count of the selected columns, ensure that this check box is selected.
	Selected Columns	A list of the columns that you selected in the Grid view.
	Plot Chart (Current Page Data)	To create a chart of the data on the current page, click the Plot Chart (Current Page Data) button.
	Plot Chart (All Data)	To create a chart of all the data in the results, click the Plot Chart (All Data) button.

Render Chart editor

Table 5. Render chart editor





Button or Field	Name	Description
	Clear All	To clear the graph and close the window, click Clear All .
	Create New Dashboard	To create a new dashboard based on the data in the graph, click the Create New Dashboard icon.
	Add Charts to Existing Dashboard	To add the chart data to a dashboard, click the Add Charts to Existing Dashboard icon.
	Hide Portlet icon	To hide the graph, click the Hide Portlet icon.

Table 5. Render chart editor (continued)





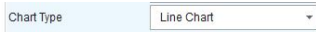


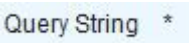


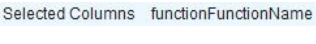
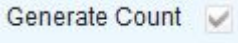
Button or Field	Name	Description
	Settings icon	To change the type of chart or the values that are displayed on the axes, click the Settings icon.
	Close Portlet icon	To close the graph and delete the chart, click the Close Portlet icon.

Chart Settings editor

Table 6. Chart settings editor

Button or Field	Name	Description
	Render	To create a chart, click Render .
Visualization tab		
	Title	Enter a name for the chart.
	Chart Type	Select the kind of chart that you want to use.
	Parameters: x-axis	Select the value that you want to display on the x-axis.
	Parameters: y-axis	Select the value that you want to display on the y-axis.
Query tab		
	Query String	The query string that is used by the search that generated the results.
	Time Filters	Select a time filter that for the chart.
	Datasource Filters	Filter the chart data for specific data sources.
	Selected Columns	The columns that the graph is based on.
	Generate Count	Indicates whether a count was generated when the chart was plotted.

Time zone dialog box

Table 7. Time zone dialog box



Button or field	Name	Description
	Time zone	Select the city or region for the time zone that you want to use.

Table 7. Time zone dialog box (continued)

Button or field	Name	Description
	Default time zone selection check box	Click this check box to use the specified timezone for all searches.

Changing the search time zone

By default, Log Analysis converts all times to the browser time zone.

The browser time may not match the time displayed in Log Analysis if there are issues with Daylight Savings Time.

For more information about this issue, see the *Search time zone does not match browser* topic in the Troubleshooting guide.

If you do not want to use the default time zone, you can change it. To change the time zone:

1. Click the time zone.
2. Select the city or region that represents the time zone that you want to use. For example, if you are in Ireland and you want to set the time zone as Greenwich Mean Time (GMT) you can select **Europe/Dublin (Greenwich Mean Time)**.
3. If you want to use this time zone in subsequent searches, click the **Use this time zone as the default for all searches** check box. This step is optional.
When you select a default time zone, you need to use the default time zone that is most commonly used by Log Analysis users because this helps Log Analysis to process log records more quickly and efficiently.
4. To save your changes, click **OK**


Searching data

You can search ingested data such as log files for keywords. Search results are displayed in a timeline and a table format.

Before you begin

Before you can search, you must first define a data source and ensure that the log file is configured for consumption or you can load the sample data.

Procedure

1. From the Search workspace , click the **New Search** or **Add Search** tab to open a new search table. Enter the search query.
2. Optional: You can filter data source by name, description, host name, log path, or tags or enter * to do a wildcard search. To limit the extent of the search to an individual data sources and any descendant data sources, select a leaf node from the Data Sources tree ().
3. In the **Time Filter** pane, click the **Time Filter** list (**Last 15 Minutes** ▾) and select the time period for which you want to search. Select **Custom** to specify a start time and date, and an end time and date for your search.

4. In the **Search** field, type the string for which you want to search in the log files. To view distribution information for all logs, in the Search field, type the wildcard character (*).

To search for a partial string, type an asterisk (*) at the start and end of your search string. For example, to search for strings that contain the phrase hostname, type *hostname*.

To narrow your search based on a service topology, type the service topology component on which you want to base your search, followed by a colon (:), followed by your search string. For example, service:DayTrader.

5. Click **Search**. The first search you perform after the IBM Operations Analytics - Log Analysis processes have been restarted might take longer to complete than subsequent searches.

The user interface refreshes every 10 seconds. The updated results are displayed in the **progress bar**.

Maximum search results: The search returns a maximum of 1000 records by default. This limit applies only to raw searches and not facet queries. This limit can be configured in `unitysetup.properties` file property:

`MAX_SEARCH_RESULTS=1000`. **Do not** to use a high value for the

`MAX_SEARCH_RESULTS` parameter. When a large number of results are returned, it degrades search performance.

Results

A graph displaying the distribution of matching events in the log is displayed. Log records containing a match for your search term are also displayed in Table view.

When you search for a specific term, the term is highlighted within the individual log records to facilitate faster analysis. If you search for a partial term, each term that contains the search phrase is highlighted. Fields that contain only tagged values, in other words values that are contained within angled brackets (<>), are not highlighted. If a field contains values that are tagged and values that are not tagged, the tagged terms are removed and the remaining terms are highlighted as appropriate.

If your search spans data that is stored in the archive, IBM Operations Analytics - Log Analysis displays the initial results while it retrieves the rest of the data. You can interact with the initial search results while IBM Operations Analytics - Log Analysis generates the search results. The progress bar displays the search progress.

To display the latest results during the search, click **We have more results for you**. To stop the search, close the tab. To start another search while you are waiting for the first search to complete, click the **Add Search** tab.

What to do next

If you want to load data that contains tags and want to keep the tagging, you can disable highlighting. To disable highlighting:

1. Open the `unitysetup.properties` file.
2. Locate the `ENABLE_KEYWORD_HIGHLIGHTING` property and set it to false.
3. Save the file.
4. To restart IBM Operations Analytics - Log Analysis run the following command from the <HOME>/IBM/LogAnalysis/utilities directory:

```
./unity.sh -restart
```

Related concepts:

Loading and streaming data

Before you can perform a search on log or other data, you must first load the data into IBM Operations Analytics - Log Analysis. When the file is loaded the data is indexed and is then available to be searched.

Data Source creation

You create data sources to ingest data from a specific source.

“Query syntax”

This section describes the combination of words, keywords, and symbols that you can use when searching for phrases using IBM Operations Analytics - Log Analysis.

Query syntax

This section describes the combination of words, keywords, and symbols that you can use when searching for phrases using IBM Operations Analytics - Log Analysis.

The query syntax is based on the Indexing Engine query syntax. For more information, see:

<https://wiki.apache.org/solr/SolrQuerySyntax>

Indexing Engines use a number of different query parser plug-ins. Log Analysis supports the Lucene query parser plug-in. For more information about the Lucene query syntax, see:

http://lucene.apache.org/core/5_1_0/queryparser/org/apache/lucene/queryparser/classic/package-summary.html

Standard keywords and operators

This topic lists the keywords and operators that you can use when searching in IBM Operations Analytics - Log Analysis.

Note: The operators such as AND and OR, which are part of this query syntax, are case sensitive. You need to use capitals for these operators.

- OR** This is the default operator. Either term or expression must be matched in the results. A variation to this keyword is `or`. For example, to search for a specific severity or message classifier, enter `severity:M OR msgclassifier:"WLTC0032W"`.
- +** To get AND like functions, use the `+` operator. You must add `+` as a prefix to these queries. For example, to search for a specific severity and message classifier, enter `+severity:W +msgclassifier:"WLTC0032W"`.
- AND** As an alternative to the `+` operator, you can use the AND operator. For example, to search for a specific severity and message classifier, enter `severity:W AND msgclassifier:"WLTC0032W"`.
- ""** Enables you to group individual terms into phrases that are searched for as a unit. For example, “document clustering”.
- ()** Enables you to group expressions to guarantee precedence. For example, document AND (cluster OR clustering).
- *** Wildcard operator that can be replaced in the returned value with a number of characters. This can be either passed as an operator to the

sources or expanded when the `meta.wildcard-expand` option is turned on. For example, `test*` might return `test`, `tests` or `tester`. You can also use the wildcard in the middle of the search term. For example, `t*est`.

Note: You cannot use this wildcard as the first character in a search. For example, you cannot use `*test`.

- ? Wild character operator that can be replaced in the returned value with a single character. This can be either passed as an operator to the sources or expanded when the `meta.wildcard-expand` option is turned on. For example, `te?t` might return `text` or `test`.

Note: You cannot use this wildcard as the first character in a search. For example, you cannot use `?test`.

- + Must operator. Forces the use of a keyword. For example `WAS +and DB2` searches for strings that contain the keyword `and`.

field: Enables you to restrict your query to a specific field. For example, `ID:123A` or `msgclassifier:"WLTC0032W"`. These operators are activated for every field defined in your syntax.

By default, the search engine supports the `title` field. When you are creating a search collection, you can extract any number of contents, for each document, and relate these contents to searchable fields. This is specified in the form of the source associated with each collection.

- NOT** The specified term or expression must not be matched in the search results. Variations to this keyword are `!` and `-`. For example, to search for log records that contain `WAS ID` but that do not contain `DB2 ID`, enter `"WAS ID" NOT "DB2 ID"`.

Note: You cannot use this operator for a single term.

Additional keywords and operators

This topic lists additional keywords that are more specific to the search and indexing operations performed by the search engine.

Range searches

To search for records in a range, use a range query. Range queries can include the terms in the range or they can exclude them. To include the query range terms, use brackets, for example:

```
[<search term> TO <search term>]
```

To exclude the query range terms, use braces, for example:

```
{<search term> TO <search term>}
```

Results are returned in lexicographical order.

For example, to search for all the log records modified on or between two dates, enter:

```
mod_date:[20020101 TO 20030101]
```

The search returns all the log records that have been modified in 2003, that is all the log records where the `mod_date` field is within the specified range.

You can also use range queries to search for fields that are not dates. For example, to search for all the log records that contain an ID between A to D but that do not include A or D, enter:

```
title:{A TO D}
```

DateMath queries

To help you to implement more efficient filter queries for dates, you can use DateMath queries.

For example, here are 4 possible DateMath queries:

- `timestamp:[* TO NOW]`
- `timestamp:[1976-03-06T23:59:59.999Z TO *]`
- `timestamp:[1995-12-31T23:59:59.999Z TO 2007-03-06T00:00:00Z]`
- `timestamp:[NOW-1YEAR/DAY TO NOW/DAY+1DAY]`

For more information, see the *DateMathParser* topic in the Lucene documentation at:

http://lucene.apache.org/solr/5_1_0/solr-core/org/apache/solr/util/DateMathParser.html

Escaping special characters

Regular expression or regex queries are supported by the query syntax.

Log Analysis supports escaping for the following special characters:

```
+ - && || ! ( ) { } [ ] ^ " ~ * ? : \ /
```

To escape a special character, use a back slash (\) before the special characters. For example, to search for the query (1+1):2, enter:

```
\\(1\\+1\\):2
```

To find multiple terms, use brackets. For example, to search for moat and boat, enter:

```
/[mb]oat/
```

Example query: Search for a keyword in a specified range

This example search query searches for a keyword in a specified time range. You can use queries like this one to search for a keyword in a specified range.

You want to search for all the instances of the error code 6543 in the SUMMARY field with a response time less than 5 seconds. For example:

```
fieldName:SUMMARY
```

```
fieldContents: "Transaction 12345 has failed with response time of 10  
seconds and error code of 6789."
```

You enter the following query. It specifies the summary field and a query for the time range:

```
"query" : "SUMMARY:/.*\s{[6-9]|\d\d+)\ssecond.*6789\."/>
```

Search results timeline

The search results timeline displays a graph showing the distribution of log events over a time period.

You can use the timeline slider to view the logs for a specific duration. You can zoom in and out to change the range of the data displayed. If there are a large number of dates in the log file, the timeline might display them as ### rather than displaying the dates. Use the timeline scroller to zoom in and display the appropriate date information.

The Timeline does not display data at the seconds level for data ingested using IBM Operations Analytics - Log Analysis Version 1.1. A message is displayed indicating that the data was indexed using a previous version of IBM Operations Analytics - Log Analysis. For this data, the Timeline cannot display information for multiple events that occur in time periods of less than one minute.

List and Grid views

Log records are displayed in both in a grid view and a list view. The default view is the List view. Log records are displayed in the grid view can be ordered by column for easy analysis. This view can be customized and used to display information in a range of ways:

Sorting in Grid view

You can also sort the information in the table columns by clicking on the column header. Not all columns are sortable. The Index configuration determines the fields that can be sorted.

The `_datasource` field is an internal field and cannot be sorted or used for charting purposes. If you want to sort your data by data source or to create a chart, create a field in the Index configuration for this purpose. This field can be used for sorting and in charts.

The order in which the fields are displayed is governed by the associated index configuration and source type. You can use the Index Configuration editor to add new fields or adjust the order of existing fields.

For more information about editing index configuration, see the *Editing an index configuration* topic in the *Administering and Installing* Guide.

Toggle views

Click **List View** or **Grid View** button to toggle between views. In each of these views, the button displayed allows you to toggle to the alternative view.

Customizing the columns displayed

To configure Grid view to display only the columns that you require, click the **Select Columns** icon on the Grid view toolbar, remove the columns that you do not want to display, and click **OK**.

Display a chart of your results

You can display the distribution of values for a number of the columns as a chart. To display the chart, select the column and click the **Plot Column** icon on the Grid view toolbar. The distinct values used to plot the chart can be viewed as hover help for each chart sector. The Plot feature is available for some values only. Where available, the **Plot Column** button is active when the columns are selected.

If the chart title contains a loading icon, the chart is loading data from the archive. The chart is automatically updated when all the searches are complete. If you log out before the search is completed, the search stops.

Running a Custom Search Dashboard from the Grid view toolbar

If you have configured a shortcut to a Custom Search Dashboard, click the icon on the toolbar to launch the Custom Search Dashboard.

Using a Custom Search Dashboard to display selected data from a column or cells

If you have created the required Custom Search Dashboard, you can select and display the contents of a column or individual cells. To display the data, select a column or individual cell in Grid view and then launch the application. If you select a column, only data from the currently displayed page is displayed by the application.

Related concepts:

“Custom Search Dashboards” on page 21

Custom Search Dashboards allow you to create and execute custom scripts and display the output of those scripts in a dashboard.

Refining search results

You can refine the search results.

You can narrow your search, by adding extra criteria in the search field. For example, the string `severity : E` returns log lines that contain errors. Alternatively, you can perform a free text search for a value in a column. All of the log lines that contain that text are returned. If more than 100 log lines are returned, click the arrows to view more log lines.

Note: If a host file contains the character sequence `::1` next to the host name, `::1` might be displayed as the value in the `sourceip` column.

You can also refine your search in these ways:

Search Patterns

To refine your search, use the values in the Search Patterns pane. For each new search, the list of fields with which you can filter your search is updated and listed in the Search Patterns pane. The number of occurrences of each value that is found is displayed with each unique keyword added as a child node. Click a keyword to add it to the **Search** field.

The keyword is added in the format `field:"value"`. You can add multiple keywords to refine your search. If you want to run an OR query, type the word OR between each added keyword search string. When you add all of the search criteria, click **Search** to return log lines that contain the values that you specified.

Discovered Patterns

When you search a data source that has been configured with a Source Type that uses the Generic annotator, the results of the search are listed in the Discovered Patterns pane.

For each new search, the list of fields with which you can filter your search is updated and listed. The counts in the Discovered Patterns pane indicate the number of records that contain a specific key or key-value pair. A key-value pair might occur multiple times in a record, but the total reflects the number of records in which the key-value pair occurs. The count of the value of nodes in a key-value pair tree might exceed the key count when multiple values occur for the same key in a single record.

Click a keyword to add it to the **Search** field. The keyword is added in the format `field:"value"`. You can add multiple keywords to refine your search. If you want to run an OR query, type the word OR between each added keyword search string. When you add all of the search criteria, click **Search** to return log lines that contain the values that you specified.

Data Source filtering

Refine your search by selecting a **Data Sources** leaf node. When you select a leaf node in the **Data Sources** tree, your search is refined to search only that data source and any descendant data sources. The **Data Sources** tree is defined by selecting a service topology node when you configure your data source. For more information, see *Editing groups in the service topology JSON file*.

Time Filters

Use the **Time Filters** list to refine your search based on a selected time period. Select a value from the list to limit the search period. The time period chosen limits the search time period based on the log entries. For example, choosing **Last Hour** limits the search to the final 60 minutes of log file entries.

Selecting a timeline value

Click a value in timeline to refine your search based on that value. Log events can be visualized up to second-level granularity.

Selecting a time zone

To change the time zone that is used in one or all of your searches, click the time zone button.

Related tasks:

Editing groups in the service topology JSON file

If your services are provided using a web application that is deployed across a range of servers such as web servers, application servers, and database servers, you might want to define the scope and structure of your application or service. To represent your application or service hierarchy, you can create a group or edit the default groups in the IBM Operations Analytics - Log Analysis service topology JSON file.

Saving a search

After you search for a keyword or series of keywords, you can save your search so that you can run it again at a later time. The searches you save are added to the Quick Searches pane.

About this task

Any directories that you created to organize your Quick Searches cannot be deleted. The directory structure is maintained.

Procedure

To save a search:

1. In the Search workspace, click the **Save Quick Search** icon. The Save Quick Search dialog box is displayed.
2. Enter a value in the **Name** and **Tag** fields. Adding a tag allows you to contain similar searches within a folder.
3. (Optional) Specify a time range as an absolute or relative time. The default option is relative time.
4. Click **OK**. The search is saved to the Save Quick Search pane.

What to do next

To use a saved search pattern, browse to the saved search in the Quick Searches pane and double-click the search pattern that you want to launch. You can also edit and delete the search from the right-click menu.

Saved searches

To display a list of saved searches, click the **Saved Searches** icon.

The following saved searches are available by default after you install the sample data:

sample WAS System Out

Example search that displays results from WebSphere® Application Server.

sample DB2 db2diag

Example search that displays results from DB2®.

sample MQ amqerr

Example search that displays results from IBM MQ.

sample Oracle alert

Example search based on alerts for the Oracle sample data.

sample App transaction log

Example search based on the sample application's transaction log.

sample Omnibus events

Example search based on the sample Omnibus events.

sample Windows OS events

Example search based on the sample Windows OS events.

Visualizing data

You can create charts and graphs that help users to process information quickly and efficiently.

Creating charts and graphs



After you select one or more columns in the grid view, you can use the **Plot Column** button to create charts to display the results.

About this task

To create a chart that is based on a specific field, you can create a search query and plot a chart based on the results. If your query returns blank fields, the chart might not display. To ensure that this does not happen, use queries that return data for the specific field. For example, to search for the severity field, enter `severity:*T0*`.

Procedure

1. In the Search workspace, select **Grid view**.
2. Select one or more columns.
3. Click the **Plot Column** icon in the Grid view toolbar. The **Plot Chart** UI is displayed.
4. To display counts for the selected columns, select the **Generate Counts** check box.

5. If the columns that you selected contain dates or numeric values, you can use the **Granularity** field to specify the granularity. You can use this setting only for columns that contain dates or numeric values and can be filtered.
6. If the columns that you selected contain numeric values, you can also apply statistical functions on the values.
 -  If you are using the Entry Edition, you can use the **sum**, **min**, **max**, **avg**, and **count** functions.
 -  If you use the Standard Edition, you can use the **missing**, **sumOfSquares**, **stddev**, and **percentile** functions.
7. To plot the chart on 100 or less of the records that you selected, click **Plot Chart (Current Page Data)**.
8. To plot the chart on all the indexed data, click **Plot Chart (All Data)**. If one or more of the fields in the selected columns is not filterable, the charts are only plotted if the total number of records is less than 1000. To change this setting, you must modify the `MAX_DATA_FACETS_IN_CHART` property in the `unitysetup.properties` file.

Results

The graph is rendered. To change the graph type, use the **Edit** icon.

If the chart title contains a loading icon, the chart is loading data from the archive. The chart is automatically updated when all the searches are complete. If you log out before the search is completed, the search stops.

Percentile statistical functions

You can use the statistical function to return values for specified percentiles.

You can use the Search REST API or the user interface to create percentile statistical functions.

For example, to query the maximum and minimum results for the 50th, 95th, and 99th percentiles with the Search REST API, you enter `"stats": ["min", "max", "percentile,50,95,99"]`. The facet response is:

```
{
  "min": 10,
  "max": 1000,
  "percentile": {
    "50": 100,
    "95": 200,
    "99": 225
  }
}
```

Percentile queries are not calculated incrementally, unlike the other queries that are used in Log Analysis. This fact means that the query needs to run over the entire time range before it can return any results. Log Analysis limits the number of asynchronous windows that can run simultaneously for this function. This property is set in the `MAX_NON_INCREMENTAL_WINDOWS` property in the `unitysetup.properties`. The default value is 2.

For example, if you specify a percentile query based on a time range from August 1 2015 to August 10 2015 and `MAX_NON_INCREMENTAL_WINDOWS=5` and `COLLECTION_ASYNC_WINDOW=1d`, only the most recent 5 days of data that is returned by the query are considered for percentile evaluation.

Dashboards

You can use dashboards to collate multiple charts, which are created during problem diagnosis, on a single user interface (UI).

For example, imagine an organization uses IBM Operations Analytics - Log Analysis to monitor all the server logs that it generates. The system administrator wants to be able to view the most critical errors, the highest severity errors, and the total number of errors on a single UI. To facilitate this scenario, you create a dashboard that is called System Admin and you add the charts that show the required information to it.

The data that is displayed on the dashboards is based on charts. For more information, see the *Charts* topic under **Custom Search Dashboard > Steps to create a Custom Search Dashboard > Application files** in the *Extending IBM Operations Analytics - Log Analysis* section.

Sample dashboards

Sample dashboards are included as part of the sample content for the following Custom Search Dashboard samples:

- Sample_EventInsightpack_v1.0.0.0
- Sample_AppTransInsightpack_v1.0.0.0
- Sample_weblogInsightpack_v1.0.0.0
- WASInsightPack_v1.1.0.3

Creating dashboards

You can create a dashboard to visualize data from multiple sources on a single UI.

About this task

This procedure describes how to create a dashboard and chart. You can also add a chart to an existing dashboard. To add a chart to an existing dashboard, click **Add Charts to an Existing Dashboard**. Select a dashboard from the list.

Use the drill-down feature to search the data for records that correspond to the area of the chart that you select. When you drill down on a field in a chart, a new search is created that is based on the field that you selected.

The drill-down feature is only supported for dashboards that are created in the UI. The drill-down feature is not supported in some cases where the underlying chart query does not support it. For example, if the query searches two date fields.

Procedure

1. Open an existing search UI or click **Add New Search** to create a new search UI.
2. Switch to the Grid View and plot the charts that you would like to include in the dashboard. You cannot use more than eight charts in a single dashboard. For more information, see “Creating charts and graphs” on page 16.
3. To plot the charts that you would like to include in the dashboard, select the columns that you are interested in and click the **Plot column** icon. Click **Plot Chart (All Data)**.

If you want to use the drill-down feature, you must click **Plot Chart (All Data)**. You cannot use the **Plot Chart (Current Page Data)** button. The drill-down function is not supported for this option.

You can also run a number of statistical operations on the selected columns. Select one of the following functions from the **Summary Function** drop-down list.

min The minimum values in a field.

max The maximum values in a field.

sum The sum of the values in a field.

avg The average of the values in a field.

count The count of the values in a field.

missing
The number of records for which the value for a field is missing.

sumOfSquares
Sum of the squares of the values in a field.

stddev
The standard deviation of the values in a field.

4. To create the dashboard, click the **Create New Dashboard** button. Enter a name and a tag. The tags are used to define groupings.
5. Save the dashboard.

What to do next

After the dashboard is created, a Custom Search Dashboard is automatically generated that represents the dashboard. You can view the Custom Search Dashboard and dashboard on the Search UI under **Search Dashboard > Dashboards**.

To view the visualization and query settings for the charts that you added to a custom dashboard, click **Search Dashboard > Dashboards > Dashboard name**. Select the required chart. Click the **Settings** icon. There are two tabs, the **Query** and the **Visualization** tabs.

Information about the query that provides the information that is visualized in the chart is displayed on the **Query** tab. This query is saved when the chart is created. To change the time filter from the default relative setting to match the absolute time that is used by the Search UI, open the **Query** tab and select the setting from the list. To view the effect of changing this setting without changing the dashboard, click **Render**.

The chart type and parameters are detailed on the **Visualization** tab. These settings are the settings that you made when you created the chart. For more information about these settings, see the *Application files* and *Charts* topics in the *Custom Search Dashboards* section of the *Extending IBM Operations Analytics - Log Analysis* guide.

To ensure that your dashboard displays current information, you can use the auto-refresh feature to regularly refresh a dashboard at scheduled intervals. For more information on configuring automatic dashboard refreshes, see the *Configuring automatic refreshes for new dashboards* topic in the *Administering IBM Operations Analytics - Log Analysis* guide.

Deleting search dashboards

If you no longer need a dashboard, you can delete it.

About this task

There are two types of dashboard that is displayed in the **Search Dashboards** list on the UI, dynamic dashboards and Custom Search Dashboards. Dynamic dashboards do not contain any custom logic and the type value that is defined in the associated JSON file is `DynamicDashboard`. Custom Search Dashboards are customized dashboards that are installed as part of an Insight Pack. They do contain custom logic.

The delete functions differently for each type. When you delete a dynamic dashboard, the dashboard and all related data are deleted. When you delete a Custom Search Dashboard, the Custom Search Dashboard file extension is changed to `.DELETED` for deletion by the IBM Operations Analytics - Log Analysis administrator.

Procedure

1. Open the Search UI.
2. Open the **Search Dashboards** list and select the dashboard that you want to delete.
3. Right-click the dashboard and click **Delete**. Confirm that you want to delete it when prompted.

Results

If the dashboard is a dynamic dashboard, the dashboard and associated data is deleted. If the dashboard is a Custom Search Dashboard, the Custom Search Dashboard file extension is changed to `.DELETED`. You can contact the IBM Operations Analytics - Log Analysis administrator and ask them to delete the Custom Search Dashboard if appropriate.

Search dashboards

To display a list of search dashboards, click the **Search Dashboards** icon on the side bar.

The **Dashboards** group contains the following search dashboards:

sample-events-hotspots

Displays example hot spot reports for the sample events.

WAS Errors and Warnings Dashboard

Displays example reports for errors and warnings that are generated by the sampleWebSphere Application Server application.

Sample-Web-App

Displays example reports for errors and warnings that are generated by the sample web application.

The **DB2AppInsightPack** group contains the following search dashboards:

DB2 Information Links

Displays useful information links for more information about DB2.

DB2 Troubleshooting

Displays example reports for DB2.

The **ExpertAdvice** group contains the following search dashboard:

IBMSupportPortal-ExpertAdvice

Displays search results based on your searches.

The **WASAppInsightPack** group contains the following search dashboards:

WAS Information Links

Displays useful information links for more information about WebSphere Application Server.

WAS Errors and Warnings

Displays example reports based on errors and warnings in WebSphere Application Server.

The **WindowsOSEventsInsightPack** group contains the following search dashboard:

Windows Events Log Dashboard

Displays example reports that are based in event data from the Windows operating system.

Custom Search Dashboards

Custom Search Dashboards allow you to create and execute custom scripts and display the output of those scripts in a dashboard.

To generate a Custom Search Dashboard, you create a JSON application file in the <HOME>/IBM/LogAnalysis/AppFramework/Apps directory. You can create sub directories in this directory to organize your applications.

After you have created your Custom Search Dashboard, the Custom Search Dashboards pane, in the Search workspace, displays the list of Custom Search Dashboards in the folder structure that you have specified.

To run a Custom Search Dashboard, locate it in the Custom Search Dashboards, right-click and click **Execute**.

For information about developing and customizing Custom Search Dashboards, see the *Extending* section of the documentation on the IBM Knowledge Center.

Expert advice

Expert advice is a Custom Search Dashboard that provides links to contextually relevant information to allow you to quickly resolve problems. Using the Expert advice Custom Search Dashboard, you can select any column or cells in Grid view and launch a search of the IBM support portal (IBMSupportPortal-ExpertAdvice.app). The Custom Search Dashboard searches for matches to unique terms contained in the column that you have selected. This Custom Search Dashboard can be launched from the Custom Search Dashboards panel in the left navigation pane of the Search workspace.

To increase the likelihood of success, the Custom Search Dashboard removes data that is specific to your environment for each search term. For example, a log message that contained the search string unable to access jarfile /myMachine/foo/bar/foobar.jar is not likely to return a specific match as the server path is likely to be specific to a user. This is abbreviated to unable to access jarfile to ensure better search results. The criteria used to exclude data can be configured. For more information about configuring the Expert advice

Custom Search Dashboard, see the *Configuring IBM Operations Analytics - Log Analysis* section of the IBM Knowledge Center.

To launch the Expert advice Custom Search Dashboard for the data returned by a search, select a column or cell of data in Grid view, right-click and click **Execute** to launch the IBMSupportPortal-ExpertAdvice.app Custom Search Dashboard.

Notices

This information was developed for products and services that are offered in the USA.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
United States of America*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

© Copyright IBM Corp. 2015. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

Privacy Policy Considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session and persistent cookies that collect each user's user name and password for purposes of session management, authentication, enhanced user usability, and single sign-on configuration. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> in the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.



Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other product and service names might be trademarks of IBM or other companies.



Product Number:

Printed in USA