TIBCO LogLogic® Log Management Intelligence

TIBCO LogLogic® Enterprise Virtual Appliance

TIBCO LogLogic® Enterprise Virtual Appliance - Container Edition

User Guide

Version 6.4.0
February 2022

Contents

Overview

Log data can comprise up to 25 percent of all enterprise data. Log data also contains critical information that can improve security, compliance and availability. Until now most companies have relied on ineffective and inefficient homegrown solutions and manual processes to manage this data.

TIBCO LogLogic® provides the industry's first enterprise class, end-to-end log management solution. Using LogLogic® log management solutions, IT organizations can analyze and archive network log data for the purpose of compliance and legal protection, decision support for network security remediation, increased network performance, and improved availability.

LogLogic log management appliances help you to simplify, automate, and reduce the cost of log data aggregation and retention, eliminating the need for servers, tape libraries, and archival administrators. If the network grows, you can simply rack and stack additional appliances as needed.

Note: LogLogic LMI has been used to represent LogLogic LMI, LogLogic EVA, and LogLogic EVA - Container Edition. All processes and procedures applicable to LogLogic LMI are also applicable to LogLogic EVA and LogLogic EVA - Container Edition, unless explicitly stated otherwise.

Appliance Functions

There are two primary user types on a TIBCO LogLogic® appliance:

- Administrator - Configures and maintains the appliance itself, including managing log sources, user accounts, appliance configurations, running backups, and more.
- User - Monitors appliance operations, runs searches, manages alerts, and creates and runs reports based on collected data.

The appliance GUI provides access to administrator and user functions.

- Administrators can perform all functions on the appliance. *TIBCO LogLogic® Log*

Management Intelligence Administration describes administrator tasks and functions. The default `admin` user can perform certain tasks and functions exclusively. Such tasks are exclusively indicated in the relevant sections.

- Users can perform functions that have been assigned to them by the system administrator, and these might include a subset of the administrator task and functions. The *TIBCO LogLogic® Log Management Intelligence User Guide* describes user tasks and functions.

Reports, search, and alert functions can be opened by clicking their respective icons on the Home page or by clicking their buttons on the top menu on the Home page. Dashboard, management, and administration functions for the appliance can be accessed from the top menu on the Home page. For more information, see the *TIBCO LogLogic® Log Management Intelligence Administration*.

Online Help can be opened by clicking the **Help** icon on the top right corner. Brief video tutorials provide tips and guidance by example for many new LogLogic features. The tutorials are accessible from the Home page and from certain application pages.

LogLogic Appliance Home Page

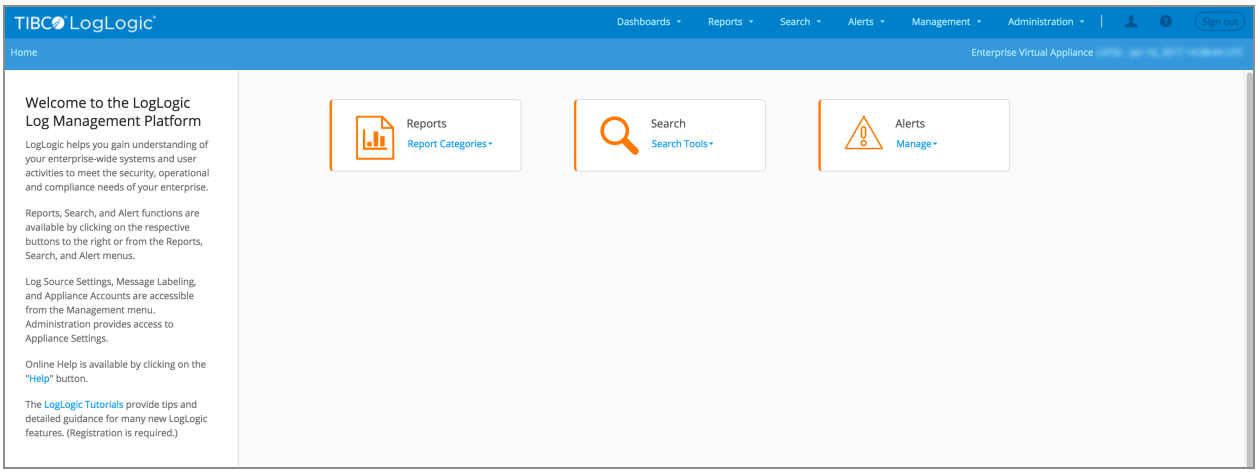

Note:

- The functions in the navigation menu vary depending on the appliance product family. For example, a TIBCO LogLogic® ST Appliance displays fewer options than a TIBCO LogLogic® LX Appliance because certain features are not available on LogLogic® ST Appliances. In addition, reports might show different entries, depending on the TIBCO LogLogic® Log Source Packages installed.
- For all text fields throughout the GUI, null is not a valid entry.

TIBCO LogLogic® Product Families

TIBCO LogLogic® appliances bring visibility of compliance activity metrics to CIOs and CSOs, and control over activities to the compliance team, permitting them to review the compliance timeliness and compliance posture mandated by Sarbanes-Oxley (SOX) and Payment Card Industry Data Security Standard (PCIDSS).

TIBCO LogLogic® appliances provide the highest log collection and analysis performance amongst all log management vendors. Log events are received and indexed in real-time. The TIBCO LogLogic® appliances have clearly stated metrics that cannot be matched.

Hardware Product Families

TIBCO LogLogic® offers the following hardware product families to provide better, faster, and smarter log management, database security, and regulatory compliance solutions to corporations:

- TIBCO LogLogic® LX Appliances are purpose-built appliances for real-time log data collection and analysis. These appliances slash response times to network security and utilization incidents, boost IT productivity, and reduce the corporate cost of security and performance event remediation.
- TIBCO LogLogic® MX Appliances perform real-time log data collection and analysis ideal for mid-size and large companies. These appliances slash response times to network security and utilization incidents, boost IT productivity, and are optimized to provide for log data needs in a non-enterprise environment.
- TIBCO LogLogic® ST Appliances automate the entire log data archival process,

minimizing administration costs while providing more secure log data capture and retention.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

Virtual Editions

TIBCO LogLogic® Enterprise Virtual Appliance provides an all-in-one software solution for log management.

TIBCO LogLogic® Enterprise Virtual Appliance and TIBCO LogLogic® Enterprise Virtual Appliance - Container Edition are virtual versions of the TIBCO LogLogic® Log Management Intelligence hardware appliance, and are distributed on the TIBCO eDelivery website.

For a comparison between the platforms and deployment instructions, see *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start.*

LogLogic® LX Appliance Product Family

Featuring a parallel processing architecture, the LogLogic LX Appliances centralize log data collection and retention by simultaneously processing raw log data and metalog data at high volume.

Distributed real-time reporting and targeted queries let administrators take immediate action on network issues from a centralized management console.

These appliances help enterprises harness the power of log data for a safer, more reliable network, while reducing corporate IT costs and providing a rapid return on investment.

Benefits

Appliances in the LogLogic LX Appliance product family offer the following benefits:

- Real-time reports, ad-hoc queries, and fast drill-downs to speed up identification, isolation, and repair of security and network incidents.
- Non-disruptive installation and plug-and-play operation: no changes to network configurations; no dependencies on other systems; no training required; available in

minutes.

- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® MX Appliance Product Family

The LogLogic MX Appliances centralize log data collection and retention by simultaneously processing raw log data and metalog data at any volume.

Designed specifically for mid-size and large companies, LogLogic MX Appliances provide the disk space and processing power required for most non-enterprise environments.

LogLogic MX Appliance features are focused on those needed to harness the power of log data for a safer, more reliable network, while reducing corporate IT costs and providing a rapid return on investment. LogLogic MX Appliances are designed for installations where data must be retained longer than LogLogic LX Appliances provide, but where managing other log appliances is not required.

Benefits

LogLogic MX Appliances offer the following benefits:

- Real-time reports, ad-hoc queries and fast drill-downs to speed up identification, isolation and repair of security and network incidents.
- Features and specifications targeted specifically to mid-size and large companies.
- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® ST Appliance Product Family

The LogLogic ST Appliances archive up to 10 years of log data while eliminating the need for servers, tape libraries, and archive administrators.

They are available in compact, rack-mountable systems with up to 8 terabytes of storage and interfaces to NAS devices. The LogLogic ST SAN (Storage Area Network) appliances offer virtually unlimited archive storage.

When used with LogLogic LX Appliances, LogLogic ST Appliances guarantee the complete and accurate transmission of network equipment logs from anywhere on the enterprise WAN or LAN. LogLogic ST Appliances feature an n-Tier architecture controlled by a management console that centralizes long-term log data archival while allowing for distributed log analysis and broader data accessibility.

Benefits

LogLogic ST Appliances offer the following benefits:

- High volume log data aggregation from centralized and remote log data sources.
- Long-term retention of unaltered, complete, raw log messages at a secure, central location to make archives unimpeachable.
- Distributed architecture of remote collection and central storage make log data collection and retention infinitely scalable.
- Self-maintaining, embedded database technology eliminates the need for database administration.

To view photographs of the appliance layout, see *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*.

LogLogic® EVA and LogLogic® EVA - Container Edition

LogLogic EVA is optimized for virtual and cloud platforms such as VMware server, Kernel-based Virtual Machine (KVM), Amazon Web Services (AWS), and Microsoft Azure; and provides an all-in-one software solution for log management.

LogLogic EVA - Container Edition is meant for container platforms such as Docker.

It helps you derive actionable information by capturing, indexing, and compressing log files and flow data.

Benefits

LogLogic EVA and LogLogic® EVA - Container Edition offer the following benefits:

- Provides all the alerting, searching, and reporting you need for both near-term activity and archived data.
- Real-time monitoring lets you alert as needed and monitor system behavior such as VPN session tracking, application distribution, port monitoring, hardware health, and security stance.
- Better IT operations, which leads to reduced time-to-resolution, simplified and improved security, improved IT efficiency, and best-practice implementations.
- Provides easy scalability.

For a comparison between the platforms and deployment instructions, see *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start.*

Scalable Infrastructure

The scalable LogLogic network infrastructure significantly accelerates response time to data center security and availability events, while providing complete log data archives for compliance and legal protection.

LogLogic appliances make log data in enterprise networks truly useful for the first time, improving corporate security, compliance, and network availability, while reducing IT costs and costly network downtime and improving corporations' return on IT investment.

Viewing Dashboards

LogLogic appliances let you monitor a large variety of data to observe the system's status and the widgets saved on your Dashboard.

- Viewing Multiple Systems Status (Management Station)
- Viewing Log Source Status

Viewing Multiple Systems Status (Management Station)

The Management Station System Status is the fastest way to view the condition and status of your appliances as traffic flows through your system.
You can use this information to provide for rapid reporting to the operations staff and acquire information about syslog messages at any particular time.

The appliance list can be sorted by ID, Model, or IP Address when the page opens. Administrators can set the default sort order.

The System Status information uses a proprietary technology for optimizing and then collecting security data for immediate use. Administrators can monitor the CPU usage when necessary to check on its congestion.

Procedure

1. Choose **Dashboards > Management Station** from the navigation menu.
2. View the following sections on the **Management Station** tab for information about an appliance's status. For detailed descriptions of each section, see the following table.
    - Message Statistics
    - Message Rate
    - New Alerts

- Message Counters

3. Click the **Refresh** button to view updated status information for the appliance.

Management Station Screen Elements

Element	Description
General information	
Software Version	Management Station appliance's software version.
Management Station sections	
Appliances	Lists the appliances in your Management Station cluster. To view the System Status for an Appliance, click its name. • A green square indicates the appliance is online. • A red square indicates the appliance is offline. • A blank square indicates the appliance entry is being updated.
Message Statistics	Displays the following message statistics: • Total, Processed, Dropped, Unapproved, and Skipped: Message processing information about each managed appliance. Click a number in these columns to change the displayed value to the nearest thousand, million, or billion value. Click the ID, Model, or IP columns to sort the appliances as required. • Message Rate/Sec: Message rate, per second, by time segments of 1, 5, and 15 minutes. Click on the message rate values to set the Message Rate graph to time scales of 4, 12, and 24 hours, respectively. • Time Skew: Time delta, in seconds, between the Management Station appliance and each remote appliance.

Element	Description
Message Rate Graph	Monitors the rate at which messages are collected. The Message Rate graph displays the current message rate by time segments of 1, 5, and 15 minutes. For example, 1 min – 100 msgs/sec. On LogLogic ST Appliances, to the right of the minutes is the number of messages per second (*xxx* msgs/sec) for the appliance. *xxx* does not reflect the number of messages that come in via the LogLogic TCP protocol. • The pink line represents the average number of messages per time segment. • The blue line represents the real-time incoming message rate for your appliance. • The red line appears when inbound traffic exceeds the preset threshold
New Alerts	The number of activated alerts, by hour and by priority (High, Medium, Low, All). Click an alert value to show the Aggregated Alert Log for LogLogic LX Appliance or LogLogic MX Appliance.
Message Counters	Statistics on each message category, stored in the syslog database. The count corresponds to a percentage related to the total number of messages received. This is useful in calculating data retention settings and maximum syslog message rates. The following is a list of message counters: • Total Received: Total number of incoming messages for all categories. • Processed: Total number of messages received and parsed into the database. • Skipped: Total number of messages ignored by the appliance when the log source entry in LogLogic® LMI

Element	Description
	exists but is disabled. • Unapproved: Messages received from a log source that is not in the Manage Devices table. These messages are discarded. The most recent 100 messages are accessible from the Data Sources screen. If auto-identify is on, all messages are auto-identified and no messages are unapproved. • Dropped: Total number of messages recognized but not processed due to network congestion.
	Updates the system status information for your appliance.

Viewing Log Source Status

The **Log Source Status** tab lets you view statistics for each source device.

Note: If during auto-discover a device has the same name as an existing device, a random number is appended to the device name.

Procedure

1. Choose **Dashboards > Log Source Status** from the navigation Menu.
2. View the following log status information for each source device:
   - Name
   - IP Address
   - Type
   - Last Received Time
   - Collector Domain
   - Total Message Count
   - Byte Rate/Sec

- Description

  For detailed descriptions of each item, see the Log Source Status Tab Elements table.

3. Click the **Refresh** button to update the view of your devices.
4. Optionally, click  to print all the items in the list.

Log Source Status Descriptions

The following table lists and describes the elements in the **Log Source Status** tab.

Log Source Status Tab Elements

Element	Description
	Saves the report in a CSV format. You can view the file in Excel as a spreadsheet. **Note:** The CSV file saves and displays a maximum of 10,000 lines. A generated report can contain more than this number.
	Displays the report in HTML format in a new window. You can save the HTML file to your local machine. **Note:** The HTML file saves and displays a maximum of 5,000 lines. A generated report can contain more than this number.
	Saves the report as a PDF file. You can save the PDF file to your local machine. Viewing the generated report as a PDF only works for Adobe Acrobat Reader version 6.0 and higher. **Note:** The PDF file saves and displays a maximum of 5,000 lines even though the generated report might contain more than this number.
	Click to print all the items in the list.
	Displays the first page or last page of detail for the device list.

Element	Description
	
•  •  • 	• Displays the previous page of detail for the device list. • Displays the next page of detail for the device list. • To display details for a specific page, type a page number and click **GO**. **Note:** For certain pages that display this option, you can only view a set number of rows. To set the number of rows to view, use the **Personal Preferences** tab.
Log Source Status section (all of the following columns are sortable)	
Name	Name of your source device. The format for this field is `<collector domain id>_<ip address>_ <device type>`, for example, `1_10.10.10.10._windows`.
IP Address	IP address for your source device.
Type	Type of source device.
Last Received Time	• For File-based devices, the time displayed shows the last event processed time • For Syslog-based devices, the time displayed shows when the last event was received
Collector Domain	This is the name used to identify each message sent from a specific device. This can either be the Collector Domain name added in the TIBCO LogLogic® Universal Collector or the name specified in the LogLogic LMI when the device was added.
Total Message Count	The following types of messages counts: • Total—Total number of messages processed for the specified device. • 1 Min—Total number of incoming messages during the previous one-minute period. • 5 Min—Total number of incoming messages during the previous five-

Element	Description
	minute period. • 15 Min—Total number of incoming messages during the previous 15 minute period.
1 Min (Byte Rate/Sec)	Byte rate per second for each device during the previous one-minute period.
Description	The description you defined for the Source Device in the **Management > Devices > Devices** tab and the **Management > Check Point Configuration > Interfaces** tab. If you selected the **Auto-identify Log Sources** option in the **Administration > System Settings > General** tab, the system displays that the source device is an auto-identified log source.
	Updates the view of your devices. If auto-identify is enabled and the appliance detects new devices, refresh displays them in this view.
Advanced Options	By default, all these options are displayed: • Name • IP Address - supports /prefix length <0-32> for IPv4 and / prefix length <0-128> for IPv6. The field supports the Classless Inter-Domain Routing (CIDR) notation for IPv4 and IPv6. Available options include: ◦ equals - only returns the pattern entered ◦ not equals - returns everything but the entered pattern ◦ in - several patterns can be entered separated by a comma; all matches are returned ◦ not in ◦ like - like behaves the same way as "in" ◦ not like **Note:** The use of asterisks (*) is no longer supported.

Element	Description
	• Type • Last Received Time • Collector Domain • Total • 1 Min • 5 Min • 15 Min • 1 Min (Byte Rate/Sec) • Description Use the list to view options in ascending or descending order.
Clear	Deletes all text in the **Advanced Options** text boxes.
Filter	Executes with the defined **Advanced Options** parameters.

Viewing Unapproved Messages

Use the **Unapproved Messages** tab to view information on up to 100 of the most recent real-time messages received from a recognized but unapproved source.
Unapproved messages are discarded.

Note: Messages from all file-based data are not listed here because they are not treated as real-time messages.

Procedure

1. Choose **Dashboards > Log Source Status** from the navigation menu.
2. Click the **Unapproved Messages** tab.
   This tab contains the following elements:

Element	Description
No	Number assigned to the message
Time	Time the message was received
IP Address	IP address of the appliance through which the message was received
Message	Text of the message

3. Click the Refresh button  to update the information.
4. (Optional) Click  to print all the messages in the list.

Viewing Recent Messages

Use the **Recent Messages** tab to view information on up to 100 of the most recently-received real-time messages.

Note: Messages from all file-based data are not listed here because they are not treated as real-time messages.

Procedure

1. Choose **Dashboards > Log Source Status** from the navigation menu.
2. Click the **Recent Messages** tab.
   This tab contains the following elements:

Element	Description
No	Number assigned to the message
Time	Time the message was received

Element	Description
IP Address	IP address of the appliance through which the message was received
Message	Text of the message

3. Click the Refresh button  to update the information.
4. (Optional) Click  to print all the messages in the list.

Search for Collected Log Messages

As the appliance collects log data from your log sources, you can search on those collected log messages.

In addition to running various simple and complex searches, you can define search filters and run reports.

By defining search filters in advance, you can include specific search criteria in an Index Search, a Regular Expression Search, and All Saved Searches without having to reenter the filtering criteria each time.

Note: To reload and open older, compressed log data for viewing on an appliance, download and view the files on your workstation.

Search Overview

LogLogic provides search and reporting tools for finding specific information in the collected log message content.

The tool you use varies depending on the task you want to perform.

Index Search

Search on indexed log source messages using a Boolean expression and see the results quickly. Use Index Search when a simple, fast search can provide the information you need to analyze failures or other anomalies.

Regular Expression (RegEx) Search

Search using a single regular expression or pre-defined search filter, either immediately or at a scheduled time.

Index Report

Generate a report based on indexed data using pre-defined Boolean search filters. Essentially, an Index Report is a compilation of multiple Index Searches run at once. You

can specify one or more pre-defined filters to use, and add additional criteria to those filters.

Search and Reporting Feature Comparison

Feature	Index Report	Index Search	RegEx Search
Multiple filters in search	Yes	No	No
Boolean Expressions	Yes	Yes	No
Regular Expressions	No	No	Yes
Graphical Results Available	Yes	Yes	No
Graphically view trends over time or log sources	No	Yes	No
Schedulable Search	Yes	Yes	Yes
Save customized search criteria for future use	Yes	Yes	Yes
View finished/past search results	No	No	Yes

Note: For a simple search to match a specific string, use Index Search. To search for strings that match more complex patterns, use RegEx Search.

Index Search

Use index search to perform targeted searches on log messages using keywords, Boolean expressions, and wildcards on the appliance or log sources.

Index Search lets you pinpoint problem areas on all log sources captured on the appliance and then view the search results quickly.

Due to the dynamic nature of LogLogic reporting, when paging between the last page of search results and other pages, additional messages matching the search criteria might have been received since the initiation of the original search. As such, you might see additional messages included on subsequent visits to the last search results page.

Index Search works on indexed logs making it faster than a search using regular expressions (RegEx search). The default criteria for the index search page is to search against all logs collected by the appliance (except the appliance's own logs) and for the last hour.

Search Expression Rules

Various rules apply when you enter a search expression.

- Use Boolean operators, such as AND, OR, or NOT for your search expression (but do not begin the expression with leading NOT)
- Use wildcard characters, such as an asterisk (*) or question mark (?) to match strings (but do not begin the expression with the wildcard)
- Do not use < or > as these are not valid characters
- Use parentheses to force an order of operations when the index search evaluates the search expression
- Enter up to 4096 characters for your search expression
- When using Index Search and Tag-Based search, the system does not support the use of search patterns shorter than 3 characters

Index Searches are case insensitive, so you do not have to use all uppercase letters when using Boolean operators, although it helps readability.

Some simple Index Search examples are listed in the following table.

Index Search Examples

Index Search Example	Rule
`tcp`	Use search expressions containing at least three characters.
`authenticate AND failed` `Tcp NOT Udp`	Use Boolean operators, such as AND, OR, or NOT.
`admin*` `10.*`	Use wildcard characters such as an asterisk (*) or a question mark (?) as shortcuts to match strings. **Note:** Wildcard character Index Search on IPv6 addresses work only if the asterisk or question mark is at the end of the address. The following examples demonstrate that the wildcards do not work if they are used anywhere else in the address: 2001:db8::ff00:42:83?? 2001:db8::ff00:*:8329 2001:db8::ff0?:42:8329 2001:db8::ff0*:42:8329 2001:db8::????:42:8329
`(tcp and udp) and service`	Use a delimiter such as parentheses to specify what gets evaluated first. In this example, tcp and udp are evaluated before the service keyword.

Running an Index Search

Index Search is available on all appliances.
The default criteria for the index search page is to search against all logs collected by the appliance (except the appliance's own logs) and for the last hour. You can search using these defaults or change them.

Procedure

1. Access the Index Search page from **Search > Index Search.**
2. Enter your search expression in the search text box and click the **Run** button.

   Note: Do not use < or > in your search expression, as these are not valid characters.

   If you want, you can adjust the search scope and rerun the search by selecting specific log sources and/or a different timeframe.

Selection of Specific Log Sources

Narrow the scope of the search to a group of log sources.

To perform a more targeted search, you can narrow the search scope to a group of log sources, such as all firewall interfaces, all routers, all General Syslog, Microsoft sources, other UNIX, or LogLogic appliances.

The default rule is set as `All Sources except LogLogic`. This includes all logs except LogLogic appliance logs. You can add any individual and/or group of non-LogLogic sources to this rule. However, if you specify any other log source, other than LogLogic source, the default rule is removed from the filter list (from the left pane) and the new log source is added. This enhancement applies to only system-defined groups and not the user-defined groups. For example, if you select a user-defined group that only includes LogLogic source, then the default rule is removed.

On the Management Station, you can select from one managed appliance or all appliances, or particular groups of appliances (for example, all LogLogic LX Appliances or all LogLogic ST Appliances) on which to run the search. The Choose Device page automatically populates the log sources included on all defined groups.

Note: When appliance selection is All, All LX/MX, or All ST, only system-defined groups (for example, All Cisco PIX) and user-defined global groups that reside on the management station are displayed.

Running a Targeted Index Search

Procedure

1. Click the **All Sources except LogLogic** button to open the Select Source(s) window.
2. Select log sources from the Add Log Sources pane. You can select sources by appliance, and filter by Name, Collector Domain, IP Address, Group or Type.
    - If you picked “Name”, enter a Source Name, a specific Device Name or a Name Mask. Wild cards are accepted in this field.
    - If you picked "Collector Domain", enter the name of the Collector Domain. This is the name used to identify each message sent from a specific device.
    - If you picked “IP Address”, enter a Source IP Address, a specific IP Address or an IP Address Mask. Wild cards are accepted in this field.
    - If you picked “Group”, enter a Group Name, or click the down arrow to the right of the text field and select “All” or one of the other Group names displayed in the drop-down box.
    - If you picked “Type”, enter a Source Type (a specific device type), or click the down arrow to the right of the text field and select “All” or one of the other Device Types displayed in the drop-down box

    Note: When adding a large number of devices, create a dynamic rule that contains all listed devices. To create a rule, first filter by Name or Type to retrieve the list of devices. Then click **<< Add filters as a rule**. This creates a dynamic rule containing all listed devices, on the right pane.
3. Enter a name for the dynamic rule in the pop-up window and click **OK**.
4. Click to select the sources you want in your report and then click **<< Add selected log sources** to add the selected devices and filters to the left-hand pane.
5. Click **Set**.

    The new Index Report search selection appears in the Sources row. The Index Search Sources field displays the newly added log sources.

    Note: If the number of devices selected is more than 60,000, the devices might not be loaded and the message `Loading` might be displayed.

Selecting Time Frame for an Index Search

Procedure

1. Click the calendar icon (to the right of Last Hour) to launch the **Date and Time Range Picker**.
2. Select a preset time interval by clicking the down arrow to the right of Last Hour, or pick a timeframe from the pop-up calendar. Click **Set**.
3. Click **Run**.
4. At the Search pop-up, select whether you want to retrieve all messages. Click **Yes**.

Result

After a few moments, the Index Search results are displayed.

The Search Results Tab

Index Search results are displayed on the **Search Results** tab.

The keywords you entered are highlighted in different colors. For example, when entering `login AND user` as your Boolean expression, the first keyword “login” are displayed in yellow and the second keyword “user” in turquoise.

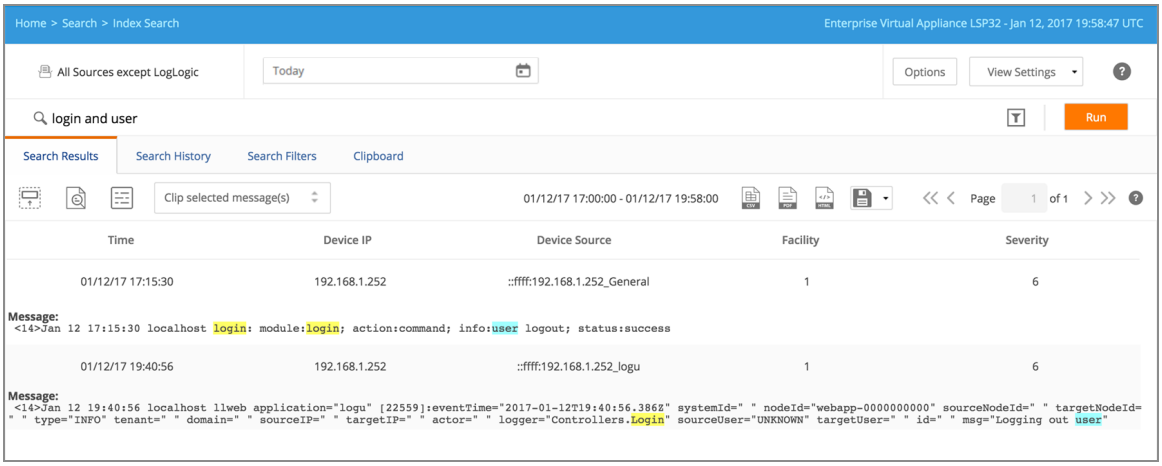

Different colors are used to highlight search keywords on the GUI after which it repeats the same color scheme.

The `llshred` utility destroys log event data. After running the utility, the log data is not deleted from the appliance. If the utility is run with `dryRun=true`, the log data is only processed. However, with `dryRun=false`, the log data is prefixed with the string `LLSHREDDER` and all characters are replaced with `X`. Running an index search after running the utility displays the shredded data in the search results.

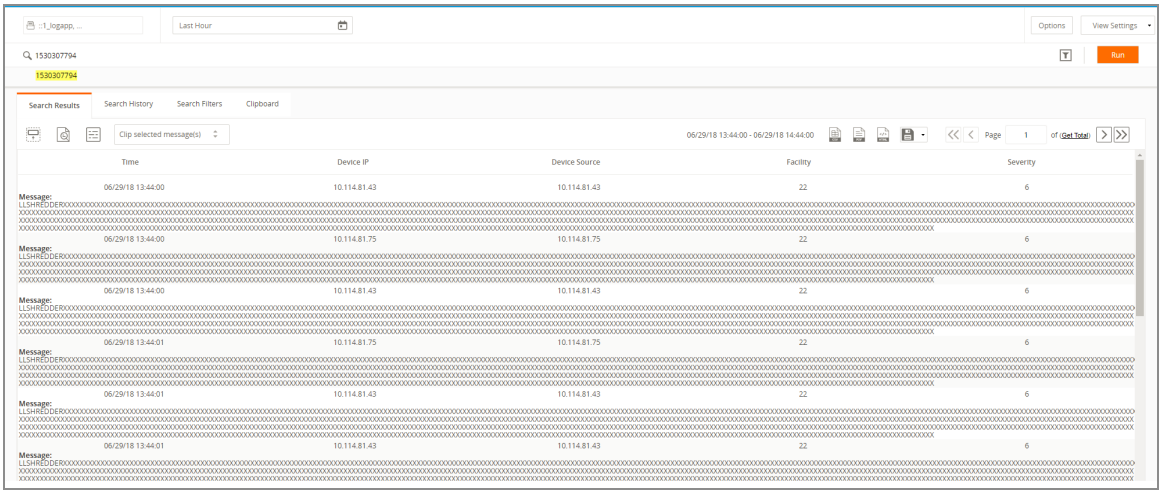

For more information about the utility, contact your administrator.

On the results tab, the Collector Domain is displayed in one of the following ways:

- For Collector Domains specified in a UC the following format; *<collectorDomainID>_<deviceIP>_<deviceType>* is displayed in the **Name** field. For example, a Windows machine with an IP address of 10.10.10.10 and collector domain windows is displayed as `1_10.10.10.10._windows`.
- For Collector Domains specified in LogLogic LMI (**Management > Devices > Add New**) the Collector Domain name is displayed in the **Collector Domain** field.

Viewing search results

You can view search results by using different view options.

From the top right of the **Index Search** page, select a suitable option from the **View** list:

Index Report Search-View options

Element	Description
Reset to Default	Reverts to default settings
Show Timeline	Select this check box to show timeline graph.
Hide Meta Header	Select this check box to hide the metadata header information.
View By	Select the option to view by Time or Device type.
Chart Type	Select the type. The options are Bar chart or Line chart.

Configuring Search Results Settings

Procedure

1. From the top right of the Index Search page, click the **Options** button. The Columns and Grouping window is displayed.
2. Optionally, enter a filter keyword in the **Keyword** field to narrow the displayed columns in your report.
3. Select the appropriate **Column Name** by selecting the check box to include or exclude that column from your report. You can change the column name by clicking the name. The column name field becomes an editable field allowing you to make the changes.

   Note: If you enter the same column name for two columns, the Index Search Results page displays the results for those two columns merged into one column.

4. Click  or  to move the selected column.
5. Choose the **Display** options.

Display Options

Element	Description
Raw	Select this option to display Index Search Results - both data in the columns, as well as the original raw message - in ascending order by time.
Grouped	Select this option to display Index Search Results - only the data in the columns without the original raw message - grouped by the selected column.
Group By	Choose the appropriate column to display group search results from the list. The default options are: • Time • Device IP • Device Source • Facility • Severity You can add more columns by creating custom tags using Log Labels.
Time Interval	This option is enabled when you select to Group By Time. The results are grouped based on the specified time interval. Select the Time Interval from the following options: • Every 5 Minutes • Every 30 Minutes • Every Hour • Every 3 Hours • Every 6 Hours • Every 12 Hours

Element	Description
	• Every Day • Every Week
Sum By	This optional setting allows you to add the numerical value of the selected column so that Search Results Summary displays the sum value of the grouped column instead of the count of message instances.
Aggregation Size	Select the option from the list. The results are sorted based on the selected option. The options are: • Top 1 • Top 5 • Top 50 • All If the search result fetches multiple rows that have identical log counts, the Aggregation Size element considers those rows as a single result group. Due to this, the Search Results tab might display more rows than the 'Top' option that was selected. For example, if there are seven result rows with log count as 4, 7, 4, 0, 91, 235, and 1029, then the 'Top 5' option returns six rows (4, 7, 4, 91, 235, and 1029), because two result rows that have identical log count (4) are considered as a single result group.

6. Click **Apply** to apply the new settings. The Index Search Results page displays the refined search results.

Search Results Options

The Search Results tab provides a toolbar with several options for managing Search results.

Element	Description
	Collapses and condenses the results display view.
	Allows you to view the selected message in relation to all others in your Index Search results. For details, see Viewing Index Search Results In Context
	Create a new log message pattern with the selected message. Highlight a message in the **Search Results** and click the **Create Message Pattern** button. The Message Pattern Editor is displayed, which can be used to select a particular message from a particular device and then create a pattern based on the parameters of that message for use in further searches.
Clip Selected message(s)	From the list use the default clipboard, a saved clipboard, or create a new clipboard to save results.
	Saves the search definition. You can choose to **Save** or **Save as** from the list to save your results. You can update your saved results using the **Save** option, see Search Results: Download Options.
<< < Page 1 of 22 > >>	Used for page navigation and for indicating the total number of pages of search results. This is particularly useful for large volumes of log messages as it lets you go through matched messages one page at a time. To page through the results, click the next arrow; to return to the previous page click the previous page arrow. You can also return to the first page or go to the last page by clicking on the first and last page arrows accordingly. The total number of results is automatically updated when you select the Show Timeline graphical view. If the time filter of the search query is a very large period of time, for example, 10 years, the number of results is huge, and displaying the results takes a long time. Furthermore, displaying the last page of results might take such a long time that the following error is displayed: `HTTP Error - 0. Please contact your admin!!!!` To avoid this, you can either reduce the time range of the classic Index Search, or add additional search terms to reduce the size of the result set, or do both.

Viewing Index Search Results In Context

When analyzing log events, you can select a particular message and see the log messages that immediately preceded or followed the message from your search results.

Note: The **In Context** tab appears only after the first time you click the icon in the search results toolbar.

Procedure

1. On the **Search Results** tab, select the message that you want to view and then select the icon.

   The **In Context** tab appears (next to the **Clipboard** tab) and the message you selected is immediately displayed in the **Search Results** tab.

2. By scrolling down on the page, the affected log message is highlighted in blue to show its relationship to the log messages that preceded this condition as well as those that occurred after this message.

3. Click the appropriate button to save the results. You can choose to save results in CSV, PDF, or HTML format.

Search Results: Download Options

You can download Index Search results in **CSV, PDF, or HTML** formats.

These buttons are located on the left side of the **Save** button. After the download is complete, the report in your chosen format is displayed.

Note: Viewing the query on the web page is faster than downloading and saving to a file. If the search results are large, downloading and saving as CSV format might take longer than as PDF or HTML.

Save Search Results

Output	Description
CSV	Use Microsoft Excel or another spreadsheet program to display Index Search results in a spreadsheet. By default, search results are written to `SearchExpressionHits.csv` and saved on the desktop.
PDF	Use Adobe Acrobat Reader to display the Index Search results. By default, search results are written to `report.pdf` and saved on the desktop. The first page includes a table of contents with links to the query used for the Index Search and the results table.
HTML	Opens a new tab in your Web browser and immediately displays HTML Index Search results as a LogLogic report. The HTML results include a table of contents with links to the query used for the Index Search and the results table. By default, the downloaded results are saved as `LogLogicReport.zip` in a temp folder on the local drive. You can use your own company logo on the report, see the General tab under System Settings.

Saving Search Results Report

You can save the results of the index or Regex search as a report.

Procedure

1. Click **Save As** option from the  icon list to save the report. You can update the saved report by using the **Save** option. The **Save As Report** window appears.
2. Enter the name and description of the report in the **Name** and **Description** fields respectively. The **Name** field is a mandatory field.

   > **Note:** Do not use any special characters in the **Description** field when saving the Index report.

3. Select the **Suite** option from the list.
4. Select the **Share?** check box if you want to share the report.
5. Select the desired print option. For Grouped Search, the options are:

- Print Summary Report
- Print Detailed Report
    - Order by group-by value: The printed report is ordered by the column selected as the group-by option.
    - Collapse Messages: The printed report is briefer, and does not include raw log data.

6. Click **Save**.

Trends

After running Index Searches, you can use the **View** menu to view search results graphically using the timeline option.

The trend output you see is based on your chosen time range and chosen devices referenced by the Index Search, and always includes only the messages and devices for that distribution.

The trend feature can be a powerful tool during your analysis of certain events and lets you see trends for certain activities by Time and Device.

Each option lets you view timeline data in either bar chart or line chart format. These charts show:

- the time or device on the x-axis
- the total number of messages on the y-axis

The procedure for viewing trends over time and by device is the same.

Viewing trends over time

Procedure

1. Click the **View** list and then select the **Show Timeline** check box.

   A timeline chart is displayed below the search text box. You can immediately see the distribution of messages over time and begin to get a sense of trends in the timeline

chart.

By hovering the mouse over an affected bar, you can get the total number of messages matching your search expression at that particular point in time.

View Menu – Viewing Trends by the Timeline Bar Chart

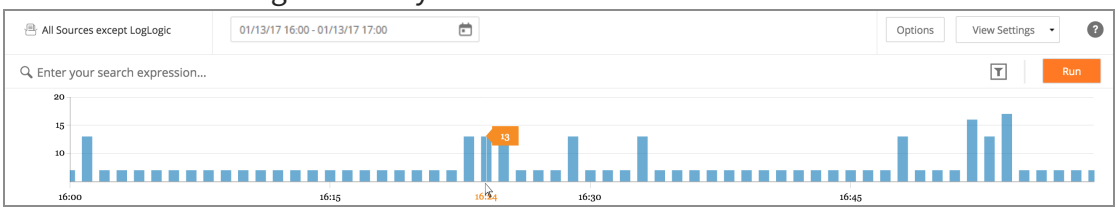

For example, in the following figure, you can see 13 log message instances at 04:24 in the evening. The scale on the x-axis shows the total number of messages while the y-axis shows the time distribution of those instances.

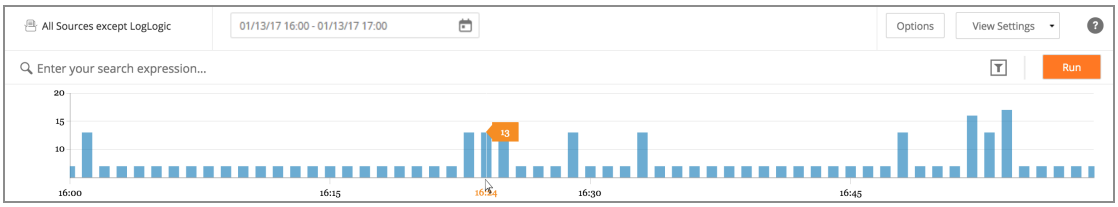

2. To zoom in on a particular area of interest, press and hold the left mouse button and drag over the area of interest.

   This refreshes the timeline view to show the zoom area in more detail.

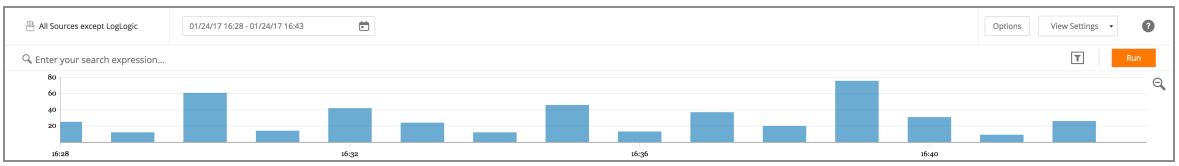

3. To return to the original view, click **Zoom Out**.
4. To view the same search in the line format, select **Chart Type > Line Chart** from the **View** menu.

   This displays the results in a line chart format. From this view, you can see spikes in the number of messages that match a keyword.

   Viewing Trends by the Timeline Line Chart

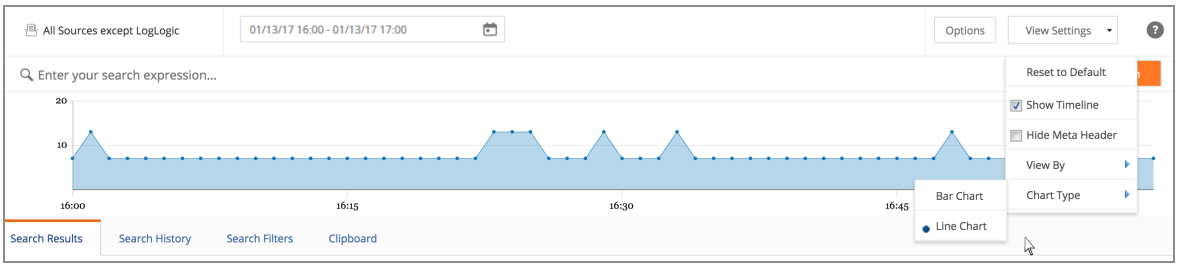

Result

Similarly, to view the same Index Search by log source, select **View By > Device** from the View menu.

The Search History Tab

Each time you run an Index Search, your search criteria are automatically saved on the **Search History** tab.

The **Search History** tab includes:

- Only those Index Searches with valid search criteria.
- User-specific Index Searches, which can be shared when saved as a search filter.
- Most recent searches on the top of the list

You can configure the search entries displayed (rows/page) on the **Search History** tab through the **Your LogApp Account** tab (see Viewing Your LogApp Account).

Saving an Index Search as a Filter

While search histories are user-specific, you can save an Index Search as a search filter.

You can use these saved search filters yourself or you can share these saved search filters with other users of the appliance.

Procedure

1. Click **Search History** to see the history of Index Searches.
2. Select the saved Index Search message and then click the button. The Save As

Filter dialog box is displayed.

3. Enter a name, description, and expression for the filter.

> **Note:** Do not use < or > in your search expression as these are not valid characters.

4. The filter name and description help you and other users to quickly understand the type of information that generates when running this Index Search.
5. If you want to share this filter with other users, click the **Shared with other users** check box.
6. Click **Add**.

   The Index Search is saved as a filter. You can use the filter in two places:

   - **Search > Index Search > Search Filters** tab
   - **Search > All Search Filters** tab

Running a Previously Saved Search Expression

Since your Index Searches are automatically saved for you on the **Search History** tab, you can browse through these previously saved sets of search criteria and run them again.

From the **Search History** tab, select the saved Index Search that you want to run and then click ▶.

The Search Filters Tab

The **Search Filters** tab lists all saved search filters created on the **Search History** tab.

The **Search Filters** tab includes the  button in the toolbar making it convenient to run a previously saved search filter.

The **Search Filters** tab organizes search filters by their name and displays the search expression used for the search filter in the **Expression** column.

> **Note:** All of your saved search filters show up on the **Search Filters** tab and on the **Index Report** tab.

Viewing or Using a Previously Saved Index Search Filter

Procedure

1. Select the filter from the table and then click  .

   This copies the search expression and enters it in the search expression text box.
2. Press **Enter** to run the search filter.

   This loads all the results of the search on the **Search Results** tab.

The Clipboard Tab

The Index Search Clipboard is an important tool for investigating and troubleshooting log events.

For example, during your analysis of a certain event, you might find an item of interest in one or more log messages. Once identified, you can create a Clipboard and copy and paste the affected log message(s) onto the Clipboard.

You can create several clipboards until you have found everything you need to help you with your analysis as you drill down on the details. After saving clipped messages to the clipboard, you can view them on the **Clipboard** tab and on the **Search Results** tab.

The **Clipboard** tab provides a toolbar with several options for using clipped messages. These options include:

	Adds a new clipboard
	Deletes one or more clipped messages
	Allows you to view or edit the clipped message

Adding a New Clipboard

You can add a clipboard from the **Search Results** page or the **Clipboard** tab.

Note: You can add up to 1,000 messages to a Clipboard. Each user is able to create up to 100 Clipboards.

The procedures are essentially the same for adding a new Clipboard. The next procedure shows how to add a Clipboard from the **Search Results** tab.

Procedure

1. On the Search Results tab, select messages to add to the clipboard from the search results.
2. To select more than one message to add to the Clipboard, hold the Shift key as you click on each message.
3. From the **Clip selected message(s)** list, select **New Clipboard**.

   The Add Clipboard dialog box opens.
4. Enter a name for the clipboard in the **Name** field.
5. If you enter an existing clipboard name, the messages are added to that existing clipboard.
6. Add a description for the clipped message in the **Annotate** field and click **Add**.

Result

The clipboard is added to the **Clipboard** tab and it is also available from the **Search Results** tab. You can go back and view or edit the clipped message(s) later on to allow for more analysis.

Viewing or Editing Clipped Messages

After saving clipped messages and annotating them, you can view or edit clipboards on the **Clipboard** tab.

Procedure

1. On the **Clipboard** tab, select the clipboard that you want to view or edit and click the edit icon .

   The Edit Clipboard dialog box appears. You can change the following:

- The **Name** of the clipped message
- The **Annotation** for the clipped message
- Remove one or more clipped log messages

2. Modify the **Name**, **Annotation**, or remove log messages and click **Update**.

Deleting Clipped Messages

You can manage the clipboard table by deleting unwanted clipped messages.

Procedure

1. On the **Clipboard** tab, select the Clipboard you want to delete and click the **Delete** button.
2. To delete more than one clipped message, hold down the shift key and select the messages you want to delete and then click the **Delete** button.

   The selected messages are deleted from the **Clipboard** tab.

Using the Tag Picker Interface

Using the Tag Picker Interface you can perform a tag-based search and access saved search terms, to quickly run an updated Index Report.

Procedure

1. Access the Index Search page by going to Search Index Search. Click the arrow  next to the text box labeled **Enter your search expression...** .

   The Tag Picker Interface opens.
2. Select an Event Type and left-click. The selected Event Type appears in the **Enter your search expression...** text box.
3. Add a Boolean operator (AND) to the search expression, and left-click a saved Field Tag. The selected Field Tag appears after the Boolean operator in the Search Expression text box.
4. Add a wild card (*) to recall all saved **Field Tags** with that name. Click **Run**.

> **Note:** You can specify special characters such as spaces, forward-slashes (/), so on, inside the quotes for **Field Tags**. For example: Identity: “John Smith”; Domain: “domain name / JOHN SMITH”.

5. Select **View** and display the Bar Chart for the search expression.
6. Compare with the previously saved Index Search results for this expression.

Regular Expression Search

Use the **RegEx Search Filter** tab to find specific types of data based on search expressions and time intervals you define.

RegEx Search provides more powerful search filter options than Index Search, though RegEx Search can take longer to process and is less interactive.

> **Note:** Working knowledge of regular expressions is a prerequisite.

Specifying Parameters for a New Search

Procedure

1. Select **Search > Regular Expression Search** from the navigation menu.
2. (Management Station only) Select the appliance (or **All Appliances**) on which to run the search.
3. Select the **Device Type**.
4. Select the **Source Device**, or all devices, connected to the appliance.

   To view Global groups created on this Management Station, you must select **Appliance > All Appliances**.

   Devices with Collector Domain are displayed in one of two ways:

   - For Collector Domains specified in a UC the following format: `<collector domainid>_<device IP>_<devicetype>` is displayed in the Name field. For

example, a Windows machine with an IP address of `10.10.10.10` and collector domain is displayed as `1_10.10.10.10_windows`.

- For Collector Domains specified in LogLogic LMI (**Management > Devices > Add New**) the Collector Domain name is displayed in the Collector Domain field.

5. Specify the **Time Interval** which to search for data passing through your appliance.
6. Define your **Search Filter**. Select one of the following options and specify the respective parameters.
   - Retrieve All—Use to retrieve all log files collected during a specified time interval regardless of the defined search expression parameters.
   - Pre-Defined—Select a pre-defined search expression (defined in/by search filters). All search filters you create appear in the list as a pre-defined search expression. If the selected filter includes multiple parameter fields, a text field for each parameter appears. The maximum length for each field is 25 characters.
   - Use Words—Use a specific word(s) as a search parameter.
   - Use Exact Phrase—Use an exact phrase as a search parameter.
   - Regular Expression—Use a regular expression as a search parameter.

     For more information about modifying or creating search expressions, see Search Filters.
7. Specify the **Time Interval** to search for data passing through your appliance.
8. Set a time for the search; do one of the following:
   - Select the **Schedule Search to Run Immediately** check box to start your search of archived data immediately.
   - Define a time to start the search of archived data. If the selected time is in the past, the search runs immediately. This search is useful if you know exactly which data source you want to search and do not need to search a time interval.
9. Enter a **Search Name** for the search.
10. Select the **Notify me when this search completes** check box to receive a notification that the search has completed.
11. To generate the report, click the **Run** button 

.

Note: Concurrent Regular Expression Searches apply only to the appliance models later than the 1000 series. You can select the number of concurrent searches to perform. The default is one, but you can choose to perform up to four searches concurrently. To specify more than four, you must edit the `capability.xml` file.

Generating a previously saved report

Procedure

1. Select **Search > Regular Expression Search** from the navigation menu.
2. In the **RegEx Search Filter** tab, select the report from the **Saved Custom Report** list.

   - To generate the report, click the **Run** button.
   - To export the report data to a file in CSV format, click the **Save as CSV** button.

Saving a Custom Report

After specifying the parameters for your report, you must save the report.

Procedure

1. Click  to expand the **Save Custom Report** section.
2. Type a name for your report and provide a brief description.
3. If you do not plan to share the report with other users logging in to the appliance, clear the **Share with Other Users** check box. By default, this check box is selected.
4. If packages are present on the appliance, the **Add Report to Package** list is visible letting you select a package in which to include this report.
5. Click the **Save Report** button to save your changes.

Using Distributed Regular Expression Search

Use Distributed RegEx Search to select all configured appliances to run a RegEx search and retrieve the merged results from the Remote Appliances and the Management Station.

Before you begin

- Add remote appliances — See the Creating a Management Station Cluster section in *TIBCO LogLogic® Log Management Intelligence Administration*.
- Your administrator must provide access to each of the remote appliances for you to have access to the data on the remote appliances. Access to appliances is provided via the Appliances tab of the User Edit page. For more information about user privileges, see the Managing Users section in *TIBCO LogLogic® Log Management Intelligence Administration*.

> **Note:** LogLogic LMI v5.4.2 or later must be installed on the Management Station and all Remote Appliances.

Procedure

1. Select **Search > Regular Expression Search** from the navigation menu.
2. For a Distributed RegEx Search you must select **All Appliances**.

> **Note:** The Distributed RegEx Search does not support Custom Reports on the Management Station.

3. Select the **Device Type** from the list of device types configured on the Management Station.

   If you select All, the Source Device field is disabled.
4. Select the **Source Device**.

   If All is selected, logs from both the Management Station and Remote Appliances are returned.

   Search results are based on the device name and are mostly returned from the Management Station. However, if the Management Station and Remote Appliances happen to have the same device name then the logs from both the Management Station and the Remote Appliance are returned.

5. Define your **Search Filter**. Select one of the following options and specify the respective parameters.

Option	Description
Retrieve All	Use to retrieve all log files collected during a specified time interval regardless of the defined search expression parameters.
Pre-Defined	Select a pre-defined search expression (defined in/by search filters). All search filters you create appear in the list as a pre-defined search expression. If the selected filter includes multiple parameter fields, a text field for each parameter appears. The maximum length for each field is 25 characters.
Use Words	Use a specific word(s) as a search parameter.
Use Exact Phrase	Use an exact phrase as a search parameter.
Regular Expression	Use a regular expression as a search parameter. For more information about modifying or creating search expressions, see Index Search.

6. Specify the **Time Interval** to search for data passing through your appliance.
7. Set a time for the search; do one of the following:
   - Select the **Schedule Search to Run Immediately** check box to start your search of archived data immediately.
   - Define a time to start the search of archived data. If the selected time is in the past, the search runs immediately. This search is useful if you know exactly which data source you want to search and do not need to search a time interval.
8. Select the **Notify me when this search completes** check box to receive a notification that the search has completed.
9. Enter a **Search Name** for the search. If a name is not entered in this field the results are displayed as distributed search <date<timestamp.

10. To generate the report, click the **Run** button.

> **Note:** Only the Management Station appliance can see the merged results from both the Management Station and Remote Appliances. A Remote Appliance can only see its own local results.

Distributed RegEx Search Results

To view a list of all the searches that are currently running, see the **Currently Running Searches** table in the **Pending Searches** tab.

For each running search, this table lists the search schedule, timespan, name, owner, Regular Expression, the approximate number of files processed, the total number to search, and the percentage completed.

For Distributed RegEx Searches two results are displayed on the Management Station search page. This is because two searches were run on the Management Station; one for the Management Station and one for the combined results from the Management Station and the selected Remote Appliances. The Remote Appliances see only their local results.

Finished Distributed RegEx Searches

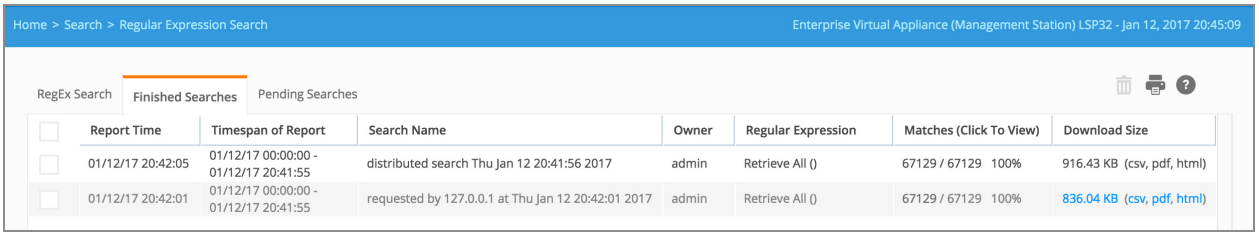

The Pending Searches Tab

From the **Pending Searches** tab, you can view running and pending searches.

The **Pending Searches** tab regularly refreshes to list all the pending, currently running, and finished RegEx and Distributed RegEx searches on the appliance.

To force a refresh, click the tab name.

List of Running Searches

To view a list of all the searches that are currently running, see the **Currently Running Searches** table on the **Pending Searches** tab.

For each running search, this table lists the following information:

- Search schedule
- Start time and timespan
- Search name followed by details of matches: *<number of matches>/<number of processed message>/<estimated number of messages>*
- Owner
- Regular Expression
- Approximate number of messages processed
- Total number to search
- Percentage completed

To suspend a running search, select its check box and click the **Stop** button. A suspended search stops processing; its partial results until that point appear in the **Finished Searches** tab.

Running and Pending RegEx Searches

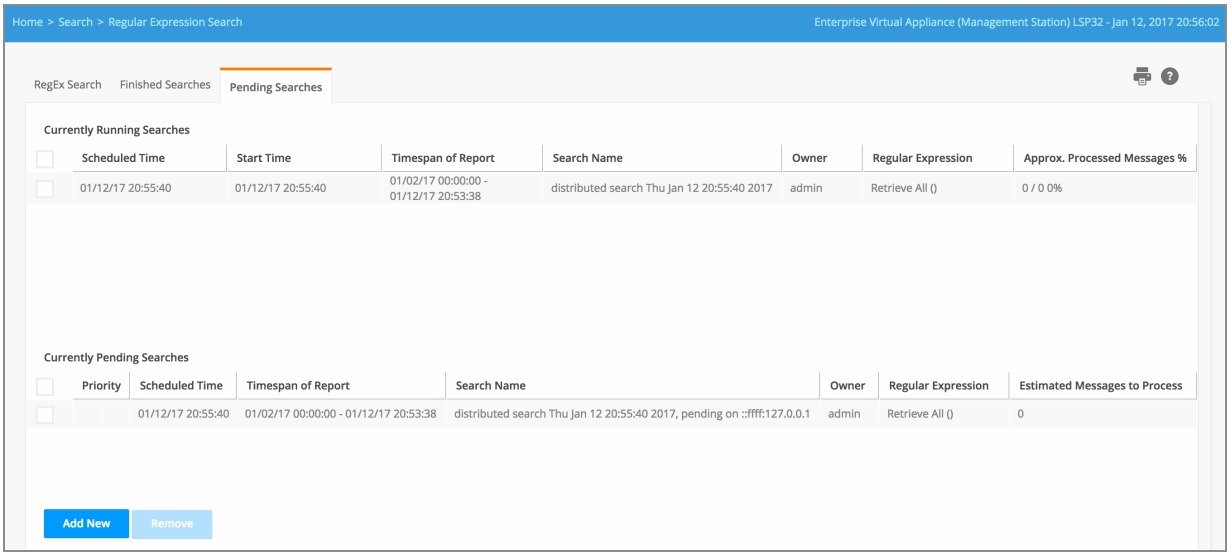

List of Pending Searches

To view a list of all the searches that are scheduled to run, see the **Currently Pending Searches** table in the **Pending Searches** tab.

For each pending search, this table lists:

- the priority for the search
- its schedule
- timespan
- name
- owner
- Regular Expression
- an estimate of the number of files to search

To remove a pending search from the queue, select its check box and click the **Remove** button. There is no confirmation prompt for removing a pending search.

To add a new RegEx search to the queue, click the **Add New** button. The **RegEx Search** tab appears.

RegEx Search Results

You can view pending, running, or finished searches in the **Finished Searches** or **Pending Searches** tabs under **Search > Regular Expression Search**.

To force a refresh of the tab and view the latest finished searches, click the tab name.

Finished Searches

To view the search results for any searches that have completed, click the **Finished Searches** tab.

Finished RegEx Searches

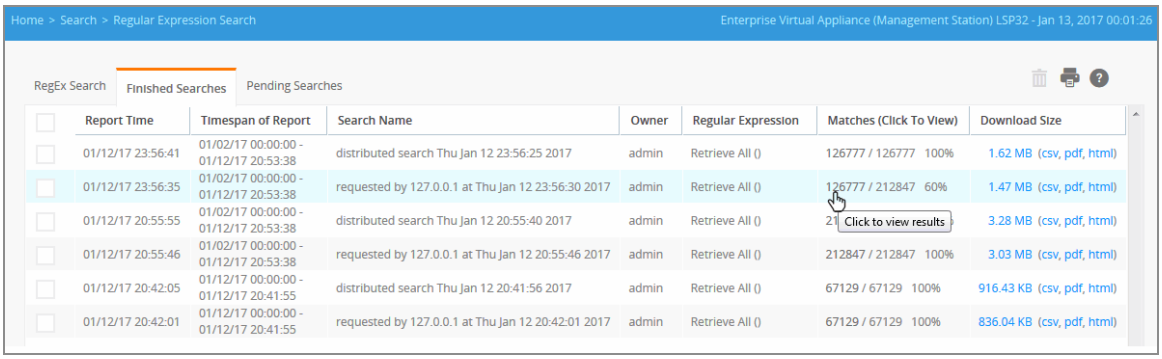

To view the search results for a particular search, click its number of matches. The **Matches** column indicates a ratio of the number of matches found to the total number of log messages during the selected time period. For example, 126777/212847 in the Matches column indicates that 212847 messages were logged during the time period in the **Timespan of Report** column, and 126777 of those messages matched the search expression.

To view or download the search results in HTML, PDF, or CSV, click the format extension in the **Download Size** column.

Clicking the download size number downloads a .txt file that is compressed to a .txt.gz file, and hence, the size of the downloaded file is less than the download size displayed. The file size also varies depending upon the file type. The results are in raw text format with one log per line in the file and with no metadata added by LogLogic LMI to the file contents.

The CSV format data is downloaded as a `.csv.gz` file and decompresses to `.csv`. When opened in Excel, it contains five fields of metadata for each raw message.

To delete a past search from the Appliance, select its check box and click the **Remove** button.

Search Filters

Search filters are user-created filters that include saved search patterns.

Such filters can be used in:

- Alerts
- Index Search

- RegEx Search
- Index Reports
- Message routing rules

You can also filter your results using the **Find** field. Enter the keywords in the **Find** field to view the filtered results based on your search keywords. You can filter results based on all columns.

Note: The **Find** field does not support the use of Japanese characters.

The All Search Filters page lists all search filters that:

- You created in the Add Search Filter page
- You created and saved from the Index **Search History** tab (see Saving an Index Search as a Filter)
- Are available to you, including shareable filters created or owned by other users

Note: Avoid using a regular expression when a non-regular expression alternative is available. Regular expressions are almost always less effective and more error-prone than non-regular expressions. For instance, instead of using the regular expression `r;^[^:]://.\.loglogic\.com/.*$` use `"r;url.domain=loglogic.com”r;`. You can also use a wild card symbol for searches. Using a wild card for regular expression searches indicates how many occurrences to match. For example, * means matching the preceding element zero or more times, whereas + means matching the element one or more times.

Adding a Search Filter

Add a search filter for complex pattern matching by using the Add Search Filter page.

Procedure

1. Select **Search > All Search Filters** from the navigation menu.
2. Click the **Add New** button.

3. Type a name for your new search filter.
4. **Sharing - Read Only** is the default setting for a new search filter; other users of this Appliance might see and use the new search filter. Set the radio button to **No** to prevent others from seeing and using the new search filter. Set the radio button to **Read Write** to allow others to see and modify the new search filter.
5. Type a brief description of the new search filter.

   This description helps you remember what the filter is for, and describes it to other users if you shared the filter.
6. Select a search filter option and enter the search filter criteria. See Search Filter Options.

   For this example, select the following option and a single filter criterion:

   a. Select the radio button **Use Exact Phrase**.

   b. Enter $username in the **Use Exact Phrase** text field.
7. Click the **Add** button.

> **Note:** When adding the very first Search Filter to the Appliance, you might see the following message immediately after clicking **Add**:
>
> ```
> There is no Search Filter defined in the system
> ```
>
> Refresh the appliance memory by clicking **Regular Expression Search** in the navigation menu; then click **Search Filters** in the menu, and your new Search Filter is displayed in the list.

Search Filter Options

You can use various types of search expressions when adding a search filter.

Search Filter Comparison

Filter type	Search criteria	Use predefined RegEx filters?	Where the filter is used
Use Words	A word, or two words with AND/OR	Yes	RegEx Search, Alerts

Filter type	Search criteria	Use predefined RegEx filters?	Where the filter is used
Use Exact Phrase	A phrase	Yes	RegEx Search, Alerts
Regular Expression	Regular expression	Yes	RegEx Search, Alerts
Boolean Expression	Keyword search using Boolean expressions	No	Index Search and Index Report

Note: Custom reports allow whichever filter types apply to the content of the custom report. For example, a custom report saved from an Index Search allows Boolean search filters. When creating a search filter to be used for Index Search or index report, ensure that you choose the Boolean expression as the filter type.

Use Words

Type a word as your search criteria. If you type more than one word, you can use the AND/OR list.

To specify any string of characters, use wildcards (*). For example, RADI*UDP would match the RADIUS opened UDP handle string.

Use Exact Phrase

Type a phrase as your search criteria. The appliance searches for strings including the phrase you specify.

To specify any string of characters, use wildcards (*). For example, RADI*UDP would match the RADIUS opened UDP handle string.

You can also define a parameter field using `$fieldname`. For example, `$username $zipcode $phone` displays text entry fields when you select the search filter in the RegEx Search tab. Field names with spaces in them display only the first word in the RegEx Search tab. For more information, see Additional Parameters to a Pre-Defined Regular Expression Search Filter.

Regular Expression

Type a regular expression as your search criteria; that is, a single character, a string of characters, or a string of numbers. A regular expression (RegEx) is a pattern that is matched against a subject string from left to right. Most characters stand for themselves in a pattern and match the corresponding characters in the subject.

The power of regular expressions comes from the ability to include alternatives and repetitions in the pattern. These are encoded in the pattern by use of metacharacters, which are interpreted in a special way (instead of standing for themselves).

Note: Avoid using a regular expression when a non-regular expression alternative is available. Regular expressions are almost always less effective and more error-prone than non-regular expressions. For instance, instead of using the regular expression `^[^:]*://.*\.loglogic\.com/.*$`" use `url.domain=loglogic.com`.

You can use a wildcard symbol (*) for searches. Using a wildcard for RegEx searches means the * matches the preceding element zero or more times.

Once you add a regular expression, the values you enter are stored as parameters in the database. To use this regular expression with alerts or RegEx Search, select the **Pre-Defined** option.

If you are creating a search filter for an alert, the search filter must be a regular expression.

Boolean Expression

Type a keyword that uses Boolean operators such as AND, OR, or NOT. For example:

```
“Portmapped translation built for gaddr” and NOT 155.363.777.53
```

Boolean expressions can search only indexed data. Indexing increases performance when searching unparsed data. It is most effective when used to find a rare occurrence of a string.

In addition to entering a keyword, you can also type:

- Numbers and words that are three or more characters
- Terms less than three characters, preceded by =. For example, for terms such as

`user=a` or `priority=7`, 'a' and '7' are indexed.

The Boolean expression should be no longer than 4096 characters in length.

Putting Your Logins Search Filter to Work

You must complete a few steps to start using your Logins search filter.

Procedure

1. Select **Regular Expression Search** from the navigation menu.
2. On the **RegEx Search Filter** tab that appears, select the **Pre-Defined** radio button.
3. In the **Pre-Defined** text field (Select Expression), click the list arrow, select **Logins search**, and click on the filter name. The filter form reloads and now displays Logins search in the **Pre-Defined** text field.

   > **Note:** If you specify the parameter `$username` in the **Use Exact Phrase** text field when you define your Logins search filter, the appliance opens a new text box next to username in which you can further define the type of user to search for.

4. Enter admin in the username text field to search for that class of user alone, or enter the wildcard * to search for logins from all users.
5. Select a **Start Time** to run your Logins search (immediately in this example).
6. Enter a name for your search in the **Search Name** text field.
7. Click the **Save Custom Report** menu expansion arrow and enter a Report Name and Report Description, and select whether to Share with Others.
8. Click Save Report.
9. Click **Run**.

   Report of Logins by username admin

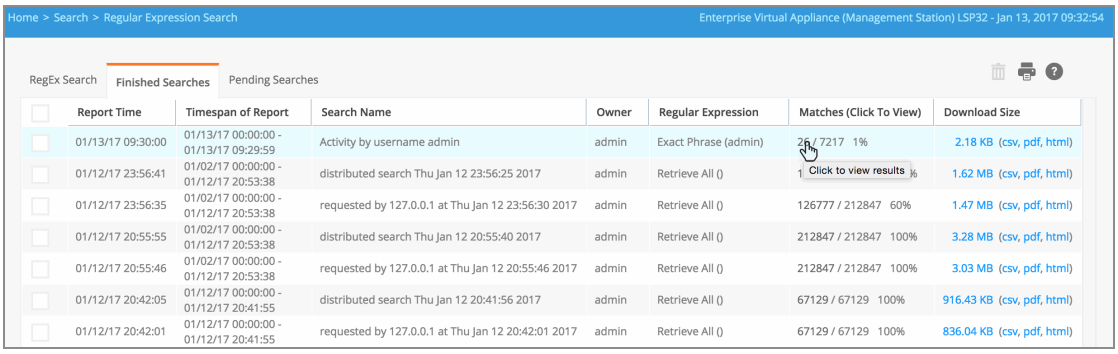

10. Click the number of matches to see the detailed report of the logins by username admin.

    Detailed Report of Logins by username admin

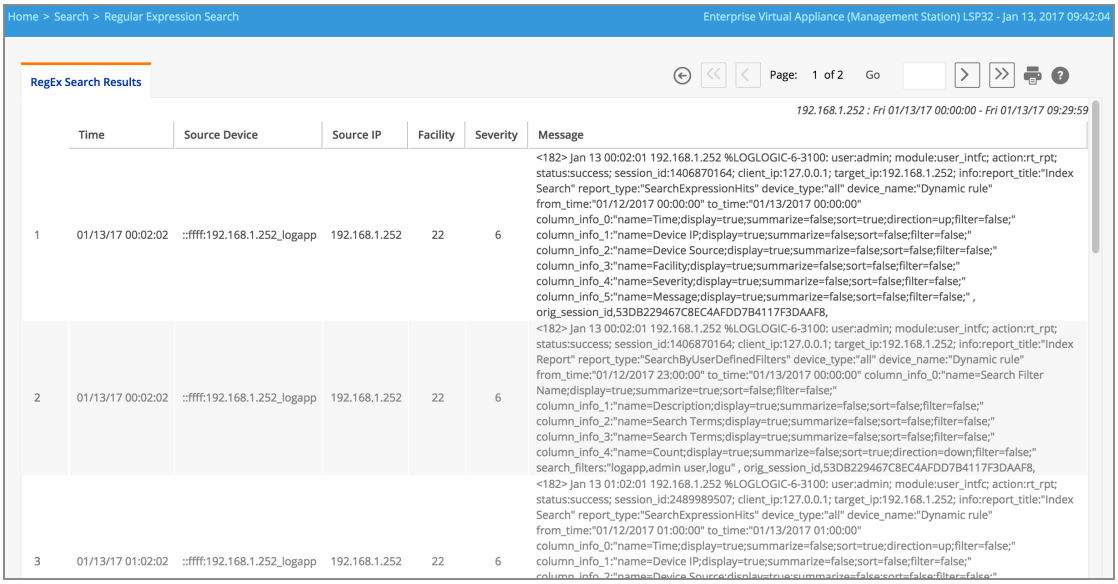

Additional Parameters to a Pre-Defined Regular Expression Search Filter

When creating a pre-defined search filter, you can define a parameter field by using the expression $fieldname.

The value you enter in the parameter replaces $field. In our example, we chose $username as our expression, and typed admin into the **User Name** field. This caused the regular expression search to return admin users wherever $username was specified.

The maximum length for each $field is 25 characters. Regular expressions can be up to 4096 characters in length.

This feature applies only to the Use Exact Phrase search filter and Regular Expression search.

Creating a Multi-Parameter Pre-Defined Regular Expression Search Filter

This example shows how to build on your single-parameter Logins search filter.

The example uses two additional parameters: `$zipcode` and `$phone`.

Procedure

1. Create a new pre-defined search filter exactly as the example Logins search filter we created earlier, except this time type $username $zipcode $phone in the **Use Exact Phrase** field.
2. Name your new search filter “Multi-parameter search” and click **Add**.

   Note: This time the new search filter appeared immediately after clicking **Add**, and both search filters are displayed in the list.

3. Select **Search > Regular Expression Search**, and select the **Pre-Defined** radio button; then select the pre-defined search filter that you just created (Multi-parameter search) from the list.

   The new form reloads, displaying each text field that corresponds to each new $field (search parameter) you define for this new search filter. The maximum length for each $field is 25 characters.
4. Click **Save Custom Report** at the bottom of the form, and enter a report name and description.
5. Click **Save Report**.
6. Type $username $zipcode $phone in the **Use Exact Phrase** field.

   In this example, we typed $username $zipcode $phone in the **Use Exact Phrase** field. The Appliance generated a text field in the search form for the part after the $. We

typed admin in the username field and used the wildcard * in the zipcode and phone fields to return the maximum number of user logins.

We elected to Save Custom Report, and named it Multi-parameter search, and we selected Schedule to run immediately for the Hourly Period: Last 24 Hours. See the results of our multi-parameter search filter query in the following figure.

The detailed Multi-parameter Search Report is revealed by clicking the number of matches returned by the search (see the arrow at the bottom of the top figure).

> **Note:** You can define this parameter for the **Use Exact Phrase** or **Regular Expression** fields from the Add or Modify page for any search filter.

7. Click the **Finished Searches** tab to see the results of the Parameter Search.

Modifying a Search Filter

In the second example earlier, we created a new search filter and added two more search parameters: $zipcode and $phone. As an alternative, we could have modified the first search filter we created, “Logins by username admin”. The following example demonstrates how to modify an existing search filter (assuming you no longer want to retain the original filter configuration).

Procedure

1. Select Search > Search Filters from the navigation menu.
2. Click on the name of the filter you want to change.
3. The **Modify Search Filter** tab appears with the same options as Adding a Search Filter.
4. Modify the search filter name, description, filter options and criteria, or sharing with other users as needed.
5. Now we think that IP address would be more valuable to us than zipcode and phone, so we elect to modify our multi-parameter search filter to suit our new needs.

> **Note:** We could also simply delete the filter and create a new one.

6. Click the Update button to modify the search filter.

7. Select **Regular Expression Search** from the navigation menu.
8. Click the **Pre-Defined** radio button on the **RegEx Search Filter** tab.
9. Select **Multi-parameter search** from the list in the **Select Expression** field (but do not enter search parameters until you complete Step 8).
10. Click the **Save Report** button at the bottom of the form and enter a new report name and description. Click **Save Report**.
11. Return to the search parameter text fields and enter your new parameters (username = admin, and ipaddress = wildcard *).
12. Click **Run**.
13. Click **Finished Searches** and then click the number of matches returned to see the results.

All Saved Searches

From the **Search > All Saved Searches** page, you can view a list of all saved searches for specific types of data based on search expressions and time intervals you have defined and saved in the past.

All the saved searches and types, such as Index Search, RegEx Search, Index Report, so on, which are stored in the system, are visible on this page.

Click the **Run** icon and regenerate the report with a different time range, or click the **Edit** icon and change the saved report parameters before rerunning the report. You can also filter the list of saved reports displayed by title by typing a keyword from the report title in the **Find** field and pressing **Enter**. The keyword or words are highlighted in the resulting list. To restore the full list of saved reports, clear the **Find** field and press **Enter** again.

You can also create reports from this page by clicking the down-arrow in the **Create Report** button and selecting from either Index Search or Regular Expression Search.

For more information on	See this topic
Index Search	Index Search
Regular Expression Search	Regular Expression Search

Using and Creating All Index Reports

Use the **All Index Reports** screen to view a list of all saved searches for specific types of data based on search expressions and time intervals you defined.
You can use these results to verify information found in your reports.

The results provide the number of hits for each selected search filter, which you can view in a table or a graphical chart. From the table, you can drill down to view the specific hits for a filter in detail similar to Index Search results.

Procedure

1. Go to **Search > All Index Reports**.
2. Click **Create Report** to open the **Properties** window.
3. Select log sources from the right-hand pane. You can select sources by Appliance, and filter returns by Name, IP Address, Group or Type.
    - If you picked “Name”, enter a Source Name, a specific Device Name or a Name Mask. Wild cards are accepted in this field.
    - If you picked "Collector Domain", enter the name of the Collector Domain. This is the name used to identify each message sent from a specific device.
    - If you picked “IP Address”, enter a Source IP Address, a specific IP Address or an IP Address Mask. Wild cards are accepted in this field.
    - If you picked “Group”, enter a Group Name, or click the down arrow to the right of the text field and select “All” or one of the other Group names displayed in the drop-down box.
    - If you picked “Type”, enter a Source Type (a specific device type), or click the down arrow to the right of the text field and select “All” or one of the other Device Types displayed in the drop-down box
4. Click **<<Add as a rule**, and enter a name in the text field of the dynamic rule pop-up.
5. Click **OK** to add the selected source and filters to the left-hand pane.
6. On the right-hand pane select a device name (or names) from the list by clicking its name.
7. Click **<<Add selected log sources** to add devices from the selected source to which you want to apply the filters when running the report.

8. Click **Columns and Filters** to select the columns for your report and choose filters for your results. Click in the field under the Value column and enter a term for the filter (such as login, id, and so on). Then click in the field under the Operator column and pick an operator from the drop-down.

   Click **Apply**. The selected operator and value move to the left-hand column.

9. Click **Index Report Search Selections** to select from the available expressions to be used in the report. If none are available, click **New Expression...** to add a new Boolean search expression for use in any Index Report.

10. In the **Add Search Expression...** popup that appears, enter Name, Description, Expression, and then click **Sharing** to define whether others can use or modify the new filter. Click **Save**.

    > **Note:** Do not use < or > in your search expression as these are not valid characters.

11. Place a checkmark next to the new search expression and click **<< Apply Selections** to add them to the left-hand pane for use in filtering your report. Then click **Save As > .**

12. Enter a name and description of the report in the pop-up. Select **Share with others** if desired. Click **Save & Close**. The new report appears in the list of all saved Index Reports.

13. Click in the Name field and enter a term to search for entries in the Saved Reports list. Hit **Enter**. Any terms found in the list of report titles are highlighted; all other reports not containing the search term is no longer shown in the list of Saved Reports. Clear the search term in the Name field and hit **Enter** to see all Saved Reports again.

14. Click the **Run** icon in the Actions column. The **Date and Time Range Picker** pops up, with Last Hour as the default setting. Click the down arrow next to Last Hour to reveal several other options (Last 2, 3, 6, 12, 18, or 24 Hours; Today; Yesterday). Select the timeframe from the **Date and Time Range Picker** and click **Run** again to execute the report.

Alerts

Alerts notify you of any unusual traffic on the network or detect anomalies on log sources or the LogLogic appliance itself.

You can create alerts specific to your monitoring needs, and use alerts that come pre-configured with TIBCO LogLogic® Compliance Suite or TIBCO LogLogic® Log Source Packages. You can also update existing alerts or remove them as needed. Similarly, you can define a new custom alert template and manage the existing custom alert templates. Using the template variables, you can define the alert email subject and alert message body for custom alerts.

You can import/export the custom alert templates and formats between appliances. For more details, refer to *TIBCO LogLogic® Log Management Intelligence Administration*.

For any alert, you can designate SNMP trap receivers, Syslog receivers, and Email recipients so people can receive notifications of alerts via email.

View and Handle Alerts

The **Show Triggered Alerts** page lists events triggered by rules defined for this appliance to monitor and report on.

From the **Show Triggered Alerts** page you can:

- View all alerts
- Filter shown alerts by alert category, priority, alert type, and keywords
- View all system alerts only, regardless of priority
- Change the alert category to Acknowledged
- Delete the alerts permanently
- (Management Station only) View alerts on a specific managed Appliance or on all managed Appliances

> **Note:** When the Data Privacy mode is enabled, these types of alerts are not displayed on the Show Triggered Alerts page: VPN Connection Alert, VPN Statistic Alert, VPN Message Alert, Pre-defined Search Filter Alert, Cisco PIX/ASA Messages Alert, and Network Policy Alert. For more information on Data Privacy mode, see the Managing System Settings section in *TIBCO LogLogic® Log Management Intelligence Administration*.

When an alert is triggered, Alert Viewer shows the alert category as **New**.

Filtering and Viewing Alerts

Procedure

1. Choose **Alerts > Show Triggered Alerts** from the home page.
2. Select the type of alerts to display from the **Show** list.
   - **All States** shows all alerts in all categories.
   - **New** or **Acknowledged Alerts** shows only alerts in the selected category.
3. Select the alert priority to view from the second list. The options are:
   - All Priorities
   - High
   - Medium
   - Low
   - All System Alerts

     To view all system alerts regardless of priority, select **All System Alerts**.
4. Select the type of alert from the third list. To view all types of alerts, select **All Types**.
5. (Management Station only) Select the appliance from which to view triggered alerts. To aggregate alerts from all managed Appliances into a single list, select **All**.
6. To filter using the keywords, enter the keywords in the **Find** field and press **Enter**. To search based on Priority and Type, select the respective lists. For the remaining columns, enter the keyword in the **Find** field to filter the list. The filtered results are displayed.

   The Show Triggered Alerts page displays the specified alerts with the following

details:

Alert Details

Element	Description
Time	The time when the alert triggered
Source IP	Source IP address contained in the syslog message If an alert is for multiple devices, **Device Group** is shown as the Source IP.
Priority	The priority of the alert The priority of an alert is specified in the **General** tab.
Type	The Log appliance alert type For a list of alert types, see the Preconfigured System Alerts and Types of Alerts tables.
Alert Destination	Email addresses, trap receivers, or the syslog receiver where notifications were sent when the alert triggered.

Paging Through Alerts

You can page through multiple results to your query in various ways.

- Use the navigation buttons  to go to the first, previous, next, or last page, respectively.
- Type the page number and click  to view the results on a specific page

Acknowledge, Print, or Remove Alerts

You can acknowledge, remove, or print alerts.

- To move alerts to the Acknowledged category, select their check boxes and click

.

- To delete selected alerts, select their check boxes and click .
- To delete all alerts permanently, regardless of priority, click Remove All.
- To print alerts, click .

Tip: Move an alert to the Acknowledged category once you have been notified of the alert. Remove an alert once the cause of the alert is corrected.

Manage Alert Templates

The Manage Alert Templates menu allows you to define a new alert template format and manage the custom alert templates.

Using the template variables, you can edit the alert message. The **Manage Alert Templates** page displays the following information:

Manage Alert Templates Details

Element	Description
Filter By Names	Filter using the template names. Enter the keywords and press the Enter key to view the filtered list.
Name	Name of the alert template.
Type	Type of the alert.
Template Type	Type of template.
Max Message Length	Indicates the maximum character length (including the alert email subject and the alert message) that is displayed.
Used By Alert(s)	Click the **List** link to view a list of alerts used by this template.

Adding a New Alert Template Format

You can define a new alert template format by using the **Add New Alert Format** option.

Procedure

1. Go to **Alerts > Manage Alert Templates**.
2. On the **Manage Alert Templates** page, click the **Add New** button.

   The Add New Alert Format window is displayed.
3. Define a template name in the **Name** field. This must be unique for each template.
4. From the **Alert Type** list, select a type of alert.

   > **Note:** For a LogLogic ST Appliance, only four alert types are available: Adaptive Baseline Alert, Message Volume Alert, Search Filter Alert, and System Alert.
5. Select the **Template Type** from the list. The options are: Email, Alert History, SNMP Trap, and Syslog. Once you select the template type, the default text for the selected type appears in the **Body** field.
6. Select a variable from the **Variables** list. Once you select a variable, the actual string for the selected variable appears in the **Variable Text** field.

   The valid variable string definitions are:

Variable Text	Description
`$ALERT_DESCRIPTION`	User-defined alert description.
`$ALERT_ID`	A number specific to the alert type. For example, 050300 for Message Volume Alert.
`$ALERT_LOG_SOURCES`	A list of log sources assigned to the alert.
`$ALERT_NAME`	User-defined alert name.

Variable Text	Description
$ALERT_TIME	The time when the alert was triggered.
$ALERT_TYPE	Type of Alerts. For example, Message Volume Alert.
$ALERT_URL	The URL that opens a page with alertable event details. Do not add any special characters after the $ALERT_URL.
$CUSTOM_EMAIL_SUBJECT	A portion of the email subject that is preconstructed based on the alert type. This field contains alert type-specific details. You cannot change this field.
$CUSTOM_STRING	A portion of the email body that is preconstructed based on the alert type. This field contains alert type-specific details. You cannot change this field.
$CUSTOM_SYSLOG_STRING	A portion of the alert syslog message that is preconstructed based on the alert type. This field contains alert type-specific details. You cannot change this field.
$FILTER	Text of a search-filter that matched as part of Search-filter alert.
$FILTER_NAME	A search-filter name. This filter is assigned to a Search-filter alert.
$HIGH_THRESHOLD	The high threshold value that was exceeded during alert monitoring.
$LOG	The log message that triggered the alert.
$LOG_SOURCES	The log sources that triggered the alert.
$LOG_SOURCE_IPS	IP addresses of log sources that triggered the alert.
$LOW_THRESHOLD	The low threshold value that was crossed during alert monitoring.

Variable Text	Description
$NUM_EVENTS	The number of alertable events that happened during the reset time. The reset time temporarily suppresses alerts.
$PRIORITY	The alert priority.
$RECIPIENT	Email, syslog, and SNMP to which the alert was sent.
$RESET_TIME	Alert reset time. Reset time temporarily suppresses alerts.
$SNMP_STRING	A portion of the alert SNMP message that is preconstructed based on the alert type. This field contains alert type-specific details. You cannot change this field.
$SRC_ APPLIANCE	The appliance that triggered the alert.
$TIME_SPAN	The time span value used in the alert definition.
$TYPE_SYSLOG	Alert type encoding as used in syslog alert message, such as “MESSAGE_VOLUME_ALERT”.

Note:

- The $$ variable is translated as $. For example, `$$ALERT_DESCRIPTION` is displayed on alert history as `$ALERT_DESCRIPTION`.
- If you define a number before the variable string, then only the specified number of characters is displayed in the alert message when the variable length is longer. For example, if you specify the variable string as `$10ALERT_DESCRIPTION`, then only the first 10 characters are displayed as the alert description. The remaining characters are truncated.
- Since some variables, such as `$LOW_THRESHOLD` and `$HIGH_THRESHOLD`, are not supported for a certain alert type, they might be displayed as empty or 0.
- When some alerts cannot distinguish log sources that have some messages or do not have any messages, such as Message Volume Alert and VPN Statistics Alert, they might list all assigned log sources in the `$LOG_SOURCES` variable.

7. The **Maximum Message Length** field displays the default maximum character length of the alert email subject and alert message that is displayed. You can update this value anytime. If the length of the alert email subject and the alert message is longer than the specified value, then the email subject is truncated.

   Note: When the selected **Template Type** is Email, the default maximum character length is 65503.

8. When you select the **Template Type** as Email, the **Subject** field appears with a default subject. Change the subject if required. The **Subject** field is mandatory but the **Body** field is optional.

   Note: The **Subject** or **Body** fields cannot include `<subject>`, `</subject>`, `<body>`, or `</body>` tags.

9. Add or change the default body of the selected template type in the **Body** field. You can select multiple variables. When adding, make sure you copy and paste the exact variable string (from the **Variable Text** field) in the **Body** field.

10. Click the **Add** button to save the new template format.

Result

The newly added template is displayed on the **Manage Alert Templates** page.

Viewing and Modifying an Alert Template

You can only view the default (system-defined) alert templates, but you cannot edit or delete them. However, you can update or delete the user-defined templates.

To viewing the default alert template format, go to **Alerts > Manage Alert Templates** and click the name of the alert template.

To modify a user-defined alert template format, perform the following steps.

Procedure

1. Go to **Alerts > Manage Alert Templates**.
2. On the **Manage Alert Templates** page, click the template name to update the format details.

   You can only update the user-defined alert templates.
3. Make the required changes and then click the **Update** button to save the changes.
4. To save the template format with a different name for later use, update the template name and click **Save As**.

Removing an Alert Template

You cannot delete the default alert templates.

However, you can delete the custom alert templates.

Procedure

1. Choose **Alerts > Manage Alert Templates** from the navigation menu.
2. Select the check box next to the template name and click the **Remove selected template(s)** button (that is located above the list on the top banner). You can delete

only the custom templates.

3. Click **Yes** on the confirmation window to delete the selected alert template. The confirmation window lists all associated alert rules for the selected template.

   Note: When you delete the selected template, all associated alert rules that are using this template use the default templates.

   The selected template is removed from the **Manage Alert Template** list.

Manage Alert Rules

From the **Alerts > Manage Alert Rules** page, you can define rules to detect unusual traffic on your network or detect appliance system anomalies.

You can add, modify, or remove alerts. You can configure alerts to generate SNMP events, a syslog receiver, or send an email notification when the alert rule is triggered. Each appliance includes a default set of alerts. You can modify these alerts and add to them as needed. You do not need to set up an SNMP or syslog server for the default alerts.

Note: If you have the Manage Alerts privileges, you can modify or delete alerts created by other users.

The **Manage Alert Rules** page displays the following information:

Field or Column	Description
Find	Filter using the keywords. Enter the keywords in the **Find** field and press **Enter**.
Name	Name of the alert.
Type	Type of the alert.
Priority	The defined priority of the alert.

Field or Column	Description
Enabled	Indicates whether the alert is active:  —You must assign a User and Alert Receiver for this alert.  —You must assign a Device for this alert.
Description	Description of the alert.

Types of Alerts

The following types of alerts are available:

Alert Type	Triggered when...
Adaptive Baseline Alert	The messages/second rate becomes more or less than the nominal rate for the traffic. **Note:** A baseline is established after 1 week from the alert activation time. After the baseline is established, the baseline is adjusted every 15 minutes. The new value is averaged in with a past baseline.
Cisco PIX/ASA Messages Alert	The messages per second rate for a specific PIX/ASA message code is greater or less than the specified rates.
Message Volume Alert	The messages/second rate is greater or less than the specified rates. If the user sets the “Zero Message Alert” check box, an alert is triggered only if zero messages are received within the timespan set. **Note:** Zero message alerts are supported only on local devices, and not on device groups spanning all LogLogic LMI appliances.

Alert Type	Triggered when...
Network Policy Alert	A network policy message is received with an Accept or Deny Policy Action. The appliance automatically pulls Check Point firewall rule bases via the Check Point Management Interface (CPMI), but you still must manually enter rules for a Network Policy Alert in the **Rules** tab. **Note:** The Rules tab is available for Network Policy Alerts and is accessible only after the new alert is initially saved.
Parsed Data Alert	The parsed data meets certain conditions specified for the alert. Parsed data alerts are different from other alert types; they are based on Pre-defined Search Filter alerts. See Creating Parsed Data Alerts.
Pre-defined Search Filter Alert	A text search filter matches message fields. This uses one of the search filters saved on the appliance: • Use Words • Use Exact Phrase • Regular Expression The pre-defined search filter is disabled if there are no search filters defined on the appliance. To create a pre-defined search filter, use **Search Filters** to add the filter. A search filter for an alert can contain words, phrase or a RegEx expression.
Ratio Based Alert	The specified message count is greater or less than a specified percentage of total messages. For example, “Login Success message count is fewer than 10% of total messages.” The appliance checks for any conditions that would

Alert Type	Triggered when...
	trigger a Ratio Based Alert every 60 seconds.
System Alert **Note:** The **Devices** tab is not available for system alerts.	An appliance system criterion is exceeded. For example, “Disk usage exceeds 80%”. By default, the priority of system alerts is high. You can change it to medium or low if needed. See also: Preconfigured System Alerts
VPN Connections Alert	A VPN connection is denied access and/or disconnected. The VPN Connections Alert is only applicable to Check Point VPN, Cisco VPN, Nortel VPN, and RADIUS Accounting device types.
VPN Messages Alert	Combinations of specific VPN message area, severity, and code. This alert is applicable to Cisco VPN devices.
VPN Statistics Alert	Recorded statistics on VPN or Radius messages match relative or absolute criteria. This alert is applicable to Check Point VPN, Cisco VPN, Nortel VPN, and RADIUS Accounting device types.

Note:

- For the LogLogic ST Appliance, an Adaptive Baseline Alert, a Message Volume Alert, and a Pre-defined Search Filter Alert can be created, along with a new System Alert.
- A LogLogic LX Appliance can create all types of Alerts.

Preconfigured System Alerts

System alerts notify you when system health and status criteria exceed the acceptable bounds.

All LogLogic LMI appliances include several system alerts that are preconfigured and enabled. By default, these alerts have the following settings:

- Email notifications are sent to the appliance admin user
- Priority is set to high
- Default reset time is 300 seconds except TCP Forward Falling Behind alert has a default reset time of 3600 seconds

All these alert settings can be customized as needed.

Alert	Description	Default High Threshold
System Alert - CPU/System temperature	The temperature of the appliance CPU has exceeded the specified High Threshold. **Note:** This alert is available in LogLogic EVA only if LogLogic EVA is used as a Management Station that manages physical Remote Appliances.	70 degrees Celsius
System Alert - Disk Usage	The usage of the specified drive on the appliance has exceeded the specified High Threshold.	80%
System Alert - Dropped Message	The number of messages dropped by the appliance has exceeded the specified High Threshold.	10 msg/sec
System Alert - Emergency Disk Usage	The disk storage is almost full, which indicates a state of emergency. An emergency disk usage alert is sent to warn the user 5% before the threshold is reached. When emergency purging begins, the alert is sent periodically until the disk usage level drops below the threshold value.	N/A
System Alert - Engine Status	An alert is triggered if the status of the selected engine changes to any of the statuses selected in the **Engine Status** section. The engine status values are Init, Running, Exited, Failsafe, or Disabled, of which Exited, Failsafe, and Disabled are selected by default. No alert is triggered for engines related to Advanced Search.	N/A

Alert	Description	Default High Threshold
System Alert - Fail Over	A failover has occurred on the appliance.	N/A
System Alert - Migration Complete	A data migration involving the appliance is successfully completed.	N/A
System Alert - Network Connection Speed	The speed of the network connection for the appliance has become lower than the specified Low Threshold.	10-Half
System Alert - Network Interface	A problem occurred with the appliance network interface.	N/A
System Alert - RAID Disk Failure	A failure occurred on an appliance RAID disk.	N/A
System Alert - Synchronization Failure	A failure occurred during log data synchronization on the appliance.	N/A

Preconfigured alerts for LogLogic EVA

For LogLogic EVA, only the following preconfigured system alerts are available:

- System Alert - Disk Usage
- System Alert - Dropped Message
- System Alert - Emergency Disk Usage
- System Alert - Engine Status
- System Alert - Fail Over
- System Alert - Migration Complete
- System Alert - Synchronization Failure

The System Alert - CPU/System temperature alert is available in LogLogic EVA only when it is used as a Management Station that manages **physical** Remote Appliances.

Adding a New Alert Rule

Adding an alert to the appliance involves selecting the type of alert, enabling the alert, specifying the log sources to monitor, and specifying alert recipients (SNMP traps, syslog receivers, and email user IDs).
Modifying an alert lets you change the same options as those for adding an alert.

> ⚠ **Warning:** When setting up an alert, do not pick search expressions that include variables. Doing so treats variables as having a literal meaning.

The **Devices**, **Alert Receivers**, and **Email Recipients** tabs list disabled log sources, receivers, or recipients marked as `(disabled)`.

Disabled entries are ignored during processing, but are listed on the following pages and are automatically present when enabled again:

- Log sources: **Management > Devices**
- Receivers: **Administration > Alert Receivers**
- Recipients: **Management > Users**

Procedure

1. Go to **Alerts > Manage Alert Rules**.
2. Click the **Add New** button.
3. In the **Type** tab, select an alert type. See Types of Alerts.

   After you select an alert type, the **General** tab for that alert type automatically appears; and the Devices, Alert Receivers, Email Recipients, and Templates tabs are enabled.
4. On the **General** tab, set up the alert.

   Options on the **General** tab vary depending on the alert type. The following table describes typical options:

Field	Description
Name	Alert name
Priority	Alert priority Default value: High
Alert Criteria	Alert criteria This field is available only for a System Alert.
Reset Time	The time in seconds after which the SNMP trap must be cleared
Enable	Click **Yes** to enable the alert. The alert is enabled after you click the **Add** button.
(Optional) SNMP OID	Enter a specific SNMP OID to further define the alert. For example, by defining this, your administrator or receiver knows that all alerts triggered with this SNMP OID originate from a specific device and alert.
Description	Alert description **Tip:** Enter a name and description unique enough to easily identify the alert in a large list.
Enable Schedule	Select the check box to specify the time period for scheduling the alerts. Select the appropriate **Time** and **Day** boxes to specify the schedule. The selected box turns blue. To remove a time slot, click the blue box.
Issue SNMP Trap Clear	Select the check box to clear the trap after the issue is resolved.

Field	Description
	You can clear the SNMP trap for system alerts where a critical condition is reported, such as disk usage alerts; but not for other system alerts that are issued only for information, such as data migration complete alert. For example, a disk usage alert might trigger when the disk usage crosses a threshold. After issuing this alert, if the disk usage later decreases to below the threshold, an SNMP clear trap is issued. The trap can only be sent via SNMP and to the same receiver that is configured for the alert. The trap contains a text message indicating the condition being cleared and the name of the alert. A record of the trap appears on the Show Triggered Alerts page as well as in the log file `sys.log`.

5. On the **Devices** tab, specify log sources for the alert.

   All the log sources on the appliance are listed in **Available Devices**. When you move a device to the **Selected Devices** section, the alerts you configure are activated for those devices. You can define different alerts for different devices.

   For available devices where the Collector Domain was specified in LogLogic Universal Collector, the following format is displayed:

   ```
   <collector domainid>_<device IP>_<devicetype>
   ```

   For example, a Windows machine with an IP address of 10.10.10.10 and collector domain is displayed as `1_10.10.10.10._windows`.

   Select the **Track all devices individually** check box to generate an independent alert messages for each selected device. The reset time tracks for the group as a whole and you can change alert properties using one alert for the device group.

> **Note:** When configuring an alert (except for System Alerts) on logs transferred using LogLogic TCP, the alert reporting can be slightly slower than in real-time. Because LogLogic TCP sends data in chunks that the appliance incrementally merges, an alert can appear anywhere between real-time and up to 5 minutes later. As a result, for example, Message Volume rates can be determined when averaging over a 5 minute or greater increment, but do not provide meaningful averages for smaller timespans. For Cisco PIX/ASA Messages alerts, the Timespan setting should be at least 60 seconds.

6. On the **Alert Receivers** tab, specify SNMP trap receivers and syslog receivers for the alert.

   You can define alerts for both SNMP traps and for syslog receivers and users; or only for SNMP traps. The **Alert Receivers** tab lists all the available traps and syslog for the appliance. You must configure SNMP traps, syslog receivers, and/or add specific traps. For more information about Alert Receivers, see *TIBCO LogLogic® Log Management Intelligence Administration*.

7. On the **Email Recipients** tab, specify people who should receive alerts via email.

   a. Select templates for each alert type from the list. The **Templates** tab displays all available templates for each alert type:

      - History
      - SNMP
      - Syslog
      - Email

      After you select the template, the format is displayed. To define or modify template formats, see Adding a New Alert Template Format.

   b. By default, the **Default** option for the Alert Email Template is selected to send the default email message. In this case, from the **Message Size** list, select Long or Short message forms.

   c. Select the **Enable View Alert Detail from Email** check box to provide additional alert detail in email. The email includes a link that you can click to open the Alert Notification page on the LogLogic LMI GUI.

Note:

- The size of email messages that include an alert is limited to 1024 bytes. Any additional alert text is truncated.
- If a LogLogic LMI session is open in a browser and then you click the link in an alert email received, you might have to log in again in a newer session. After logging in to the newer session, you are logged out of the earlier LogLogic LMI sessions to maintain access security.

You can define alerts for both users and SNMP traps; or only for users. **Available Users** lists all the users available for the appliance. For information about adding users, see *TIBCO LogLogic® Log Management Intelligence Administration*.

8. The **Rules** tab is enabled only for Network Policy Alerts. When adding a Network Policy alert, you must save the alert and then modify it to access the Rules tab. From the **Rules** tab, you can define the Accept (or Deny) Source and Destination IP Address Ranges, Port Ranges, and Protocols parameters for the alert. For example, define the firewall policy rules you want to monitor for this alert. A single alert can have a single rule or multiple rules. You must add an alert before defining rules. You can define up to 1000 rules for each alert. If you leave the fields blank and add the rule, you are still defining an alert. The appliance accepts all values if you leave the fields blank.
9. Click the **Add** button to add the new alert to the appliance.

Creating Parsed Data Alerts

Parsed Data alerts are created differently from other alert types.

There is no Parsed Data alert type to select in the interface; its creation is based on a Pre-defined Search Filter alert. The Filter specifies matching values that are extracted by the parser from the log messages.

To use Parsed Data alert, you need to know the name of the database table where parsed logs are stored along with the column names. You can find the exact column names using the **Management > Column Manager** page to create the search filter for this alert type. For more information, see the *Managing Column Manager* chapter in *TIBCO LogLogic® Log Management Intelligence Administration*. When specifying the matching values, data type should be considered for the relevant table columns. For example, IP addresses must be a numeric type, that is, a 32-bit integer and not the string representation such as 169.1.1.1.

Procedure

1. Create a Pre-defined Search Filter:
    a. Name the filter.
    b. For filter type, select **Use Exact Phrase**.
    c. For the DB table, specify _table=. (Only one _table= entry is allowed.)
    d. Specify columns and values to match as name-value pairs separated by commas. For example, this is a string-matching filter:

```
_table=Authentication,actionID=2,statusID=4
```

2. Create a Pre-defined Search Filter alert:
    a. Name the Search Filter alert with a prefix _parsed. For example, `_parsed_Login Failure`.
    b. Select the Pre-defined Search Filter you created for this alert.

    Usage notes:

    - Parsed data alerts apply only to messages from configured log sources.
    - Parsed data alerts apply only to the tables configured in the alert.
    - Parsed data alerts are not available on LogLogic ST Appliances.
    - Do not configure the same alert for both real-time and pulled data files. Instead, create separate alerts for each, with the same search expression.

Modifying or Removing An Alert

You can modify alert settings or remove alerts from the **Manage Alert Rules** page. The same tabs appear when you add an alert (see Adding a New Alert Rule).

To modify an alert:

1. Click the alert name in the **Name** column.
2. View the settings for the Alert Rule on the **General** tab, the **Alert Receivers** tab, the **Email Recipients** tab, and the **Templates** tab. Change the settings and click **Update** or **Cancel** to retain.

To remove an alert:

1. Select the check box next to the alert name.
2. Click the delete icon .
3. On the Remove Alerts page, you can confirm or cancel the removal.

Real-Time Reports

By using real-time reports, you can search and generate reports for monitoring various real-time activities.

The real-time activities are derived from the log data that is collected from your log sources. Each Real-Time report category contains multiple specific reports.

> **Warning:** Depending on LogLogic LSP packages and your selected log sources, you might see different types of reports, columns, and optional filters for each report.

Preparing a Real-Time Report

The real-time reports are a central component to the agile reporting in LogLogic. By using the real-time reports, you can quickly view detailed information about the collected log data, catered to your specific needs.

Real-time reports can take longer than Saved Reports because they run against all up-to-the-minute raw log data; and not against stored summarized log data. Real-time reports capture all hits in collected raw log data that meet the report criteria.

Sometimes, the message **Message: Unavailable** might be displayed in the report result. To view the accurate detailed messages, run an Advanced Search query by including the `sys_body` column.

> **Note:** When two devices have the same IP address but only one device has a Collector Domain ID, duplicate data (data combined from both domains) might be displayed.

To generate a real-time report, refer to the procedure and illustrations in Generating a Report: An Example - Denied Connections Report.

Selecting Sources and Search Filters

Procedure

1. In the navigation menu under **Reports**, select the category and type of report to generate.
2. Click **Create Report** to open the Properties window.
3. Under **Add Log Sources**, click the down arrow next to **Select** and pick a filter (Name, Collector Domain IP Address, Group or Type) to filter returns.
    - If you picked “Name”, enter a Source Name, a specific Device Name or a Name Mask. Wild cards are accepted in this field.
    - If you picked "Collector Domain", enter the name of the Collector Domain. This is the name used to identify each message sent from a specific device.
    - If you picked “IP Address”, enter a Source IP Address, a specific IP Address or an IP Address Mask. Wild cards are accepted in this field.
    - If you picked “Group”, enter a Group Name, or click the down arrow to the right of the text field and select “All” or one of the other Group names displayed in the drop-down box.
    - If you picked “Type”, enter a Source Type (a specific device type), or click the down arrow to the right of the text field and select “All” or one of the other Device Types displayed in the drop-down box

    > **Note:** When adding a large number of devices, create a dynamic rule that contains all listed devices. To create a rule, first filter by Name or Type to retrieve the list of devices. Then click **<< Add filters as a rule**. This creates a dynamic rule containing all listed devices, on the right pane.

4. If desired, add a second filter by clicking the **+ sign** and repeating Step 3 as often as you like.
5. To delete a filter, click the **- sign** to remove the last selection made (repeat if needed). Do not click **Cancel** unless you want to cancel your report.
6. Click **<< Add as a rule**, and enter a name in the text field of the dynamic rule pop-up.
7. Click **OK** to add the selected source and filters to the left-hand pane.

8. Select a device name (or names) by clicking its name.
9. Click **<< Add selected log sources** to add devices from the selected source to which you want to apply the filters when running the report.
10. Click **Run** to initiate a report of the selected source and devices with the filters you chose in Step 3.

Selecting Time Frame and Running a Report

Procedure

1. When you click **Run** in Step 10, the **Date and Time Range Picker** pops up, with Last Hour as the default setting. Click the down arrow next to Last Hour to reveal several other options (Last 2, 3, 6, 12, 18, or 24 Hours; Today; Yesterday).
2. To select a different date range, click the small calendar icon to the right of the current Date and Hour display and chose any month and day for the start of the report period. Move to the right and click the second small calendar icon to choose any month and day for the end of the report period.
3. Click **Run** again to execute the report.

Operations on the Reports

You can perform various operations on the report results page.

- To resize and move the columns to the positions you prefer, click on them and drag.
- To see detailed information for a particular source device, click the number of returns for the device in the **Count** column.
- To download reports in CSV, PDF, or HTML format, click the corresponding icon.

Modifying Report Settings and Time Frame

Procedure

1. Clicking the **Edit Settings** button opens up a Properties window again, this time allowing you to **Add Columns and Filters** if desired.
2. Enter your selections for **Add Columns and Filters** (if any) and click **Save As**.
3. Enter a name and description for the report in the pop-up window. Select **Share with others** if desired. Click **Save & Close**.
4. Click **Run Again** to execute your report with the new filtering criteria. The new report is displayed in the list of all Saved Reports (from **Reports > All Saved Reports**).
5. Click the date range (blue type at top left) to modify the timeframe for your report. The **Date and Time Range Picker** appears, with Last Hour as the default setting. Follow the steps listed in Selecting Time Frame and Running a Report.
6. From the list of Saved Reports (access **Reports > All Saved Reports**), click **Run** or **Edit** to modify the report settings of any Saved Report.
7. To search for a particular report or report series in the Saved Reports list, click in the **Find** field and enter a search term.
8. Press **Enter**. Any term found in the list of report titles are highlighted; all other reports not containing the search term is no longer shown in the list of Saved Reports. Clear the search term in the **Find** field and press **Enter** to see all Saved Reports again.
9. To add a schedule for a Saved Report, click the report **Name** and then click **Schedule selected**.

   The **Scheduling** window opens. You can define a Timeframe, Email Recipients (pre-defined system users), and Formatting options. Click the **Manage Recipients** button to update the appliance address book. Using this option, you can add new or modify recipient addresses that are non-defined system users (that are not defined under **Management > Users** page).
10. To delete a Saved Report from the list, click the report **Name** and then click **Remove selected**.

    A pop-up message asking you to **Confirm Deletion** is displayed.

Formats for Saving a Generated Report

You can save a generated report as a CSV, PDF, or HTML file by clicking the icons at the top of the report results.

When saving a generated report in any of the available formats, you can save and display a maximum of 5,000 lines, even if the number of lines in the generated report exceeds the defined limit. To change this limit, contact your administrator.

Format	Description
CSV	Downloads and saves the report data in a comma-separated .csv file, that can be viewed in spreadsheet applications such as Microsoft Excel. By default, reports are created in the CSV format. You can download the report by clicking the CSV icon.
PDF	Downloads and saves the report data in a PDF file, viewable in Acrobat format such as Adobe Acrobat Reader.
HTML	Open the report data in a new browser window or tab, from which you can also download the HTML file for archival.

Rerunning a Saved Report

If you have saved a report earlier, you can change some parameters and rerun it.

Procedure

1. To rerun a saved report, go to **Reports > All Saved Reports** and select a previously saved report.
2. To regenerate the report with a different time range, click the **Run** icon.
3. To change the saved report parameters before rerunning the report, click the **Edit** icon. All options are available, not just the ones originally selected.
4. To customize the new report, use new filters and wildcards.

> **Note:** Wildcard searches are supported for IP addresses and detailed messages.

Generating a Report: An Example - Denied Connections Report

This example shows how to generate a Network Activity report that displays denied connection activity related to the IP addresses you select.

These steps do not apply to the generation of the following reports on the appliance:

- Check Point Policies report, which lists current Check Point Firewall policy rules on log sources connected to your appliance.
- All Saved Reports, which lists previous search results, saved as reports, and selected to be shared with others at the time of generation.

Procedure

1. Select **Reports > Network Activity > Denied Connections** from the home page menu.
2. Click the **Create Report** button.
3. Select the log source connected to the appliance.
4. Select log sources from the list by clicking its name (or names). Click **Add selected log sources** to move them to the Log Sources list.
5. Click **Run** to run the report.
6. Specify the time interval to search for data passing through the appliance and click **Run**.
7. On the **Denied Connections** results page, adjust the order and position of columns.
8. At the top menu, select the CSV, PDF, or HTML icon to export the entire report to a file.
9. To choose another time to run the Denied Connections report, click the date range in the upper left section of the report.
10. Select the date and time and click **Run**.
11. Click the **Edit Settings** button to revise columns and filters in the report and **Run** the report again.

Rerunning and Editing settings of a previously saved report (Denied Connections)

Settings of a previously saved report can be edited and you can rerun the report.

Procedure

1. Select **Reports > Network Activity > Denied Connections** from the Home page.
2. To run the saved report, click ▶ and then click the **Run** button on the **Date and Time Range Picker** that pops up.
3. After the Denied Connections report opens, click the **Edit Settings** button.
4. Click **Properties** to open the Properties Dialog pane.
5. Enter your data and click **OK**.
6. To add a schedule for the Denied Connections report, click the **Scheduling** link.

   The **Add a Schedule** pane opens on the right side. You can define a Timeframe, Email Recipients (pre-defined system users), and Formatting options. Click the **Manage Recipients** button to update the appliance address book. Using this option, you can add new or modify recipient addresses that are non-defined system users (that are not defined under **Management > Users** page).
7. Click the **Add Schedule** button at the bottom of the Timeframe pane to confirm the schedule for the Denied Connections report.
8. Click **Save and Close** on the **Properties** window to save your entries.
9. View the saved schedule for the Denied Connections report.
10. To make further changes to the Denied Connections report, repeat Steps 1 — 9.

Available Operators

Multiple filter operators are available for each report.

These operators are listed in the Optional Filter Operators table.

Note: Some report columns display as empty when the actual value is either null or an empty string.

- If the value is null, you can filter using `--null--`.
- If the value is an empty string, you can filter using two single quotes `''`.

Optional Filter Operators

Operator	Description
=	Specifies an acceptable substitution for a word in a query.
!=	Specifies to not substitute a word in a query.
in	Displays data in the results that contains the specified word in a list.
not in	Excludes data in the results that contains the specified word in a list
like	Displays data that has a partial match to the value you type. For example, you can use this operator to type a partial IP address such as 10.2.3.*. This type of search returns all IP addresses which contain these numbers.
not like	Excludes data that contains a partial match to the value you type.
contain	Displays data that matches the alphanumeric string you type. For example, you can use this operator to type a string such as 'Accessed URL' for any detailed message. This type of search returns all detailed messages which contain, start with, or end with the value 'Accessed URL'.
not contain	Excludes data that matches the alphanumeric string you type.
start with	Displays data that begins with the alphanumeric value you type. For example, you can use this operator to type a string such as 'Accessed

Operator	Description
	URL' for any detailed message. This type of search returns all detailed messages which contain, start with, or end with the value 'Accessed URL'.
not start with	Excludes data that begins with the alphanumeric value you type.
end with	Displays data that ends with the alphanumeric value you type. For example, you can use this operator to type a string such as 'Accessed URL' for any detailed message. This type of search returns all detailed messages which contain, start with, or end with the value 'Accessed URL'.
not end with	Excludes data that ends with the alphanumeric value you type.
regexp	Displays data in the results only that contains the regular expression you define.
not regexp	Displays data in the results only that does not contain the regular expression you define.
>	Displays only data in the results that is higher than a threshold number.
<	Displays only data in the results that is less than a threshold number.
between	Displays data that is between (inclusive) the numeric values you type.

Access Control Reports

To search for and generate reports on the number of times a selected log source executes an authentication rule, use **Access Control** reports.

The submenu that appears when you click **Reports > Access Control**, lists the reports that are available for each log source.

Accessing the Access Control Reports

Choose **Reports > Access Control > report-name** from the navigation submenu, where *report-name* is any one of the Access Control reports.

Access Control Reports

Report	Definition	More information
Permission Modification	Use the Permission Modification screen to search for and create a report on changes made to user permissions on selected log sources during a specified time interval.	Permission Modification Reports
User Access	Use the User Access screen to search for and generate a report on user activities in accessing resources (for example, service, file, directory, application) on selected log sources during a specified time interval.	User Access Reports
User Authentication	Use the User Authentication screen to search for and generate a report on who has authenticated on selected log sources during a specified time interval.	User Authentication Reports
User Created/Deleted	Use the User Created/Deleted screen to search for and generate a report on what users have created or deleted during a specified time interval.	User Created/Deleted Reports
User Last Activity	Use the User Last Activity screen to search for and generate a report on activity of users during a specified time interval.	User Last Activity Reports
Windows Events	Use the Windows Events screen to search for and generate a report on data about all log events from the Microsoft Windows operating systems. For example, the captured log events include application, security, and system events.	Windows Events Reports

Preparing a Real-time Report includes the common options that you specify for all Real-

Time Reports.

Optional columns and filters can be sorted in ascending or descending order. Choose sort order using the list. Optional filter operators are different for each Access Control report, and are explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Permission Modification Reports

To search for and generate a report on activities related to modification of user permissions (for example, adding or deleting permissions) on selected log sources during a specified time interval, use the Permission Modification Real-Time Report.

Menu path: **Reports > Access Control > Permission Modification**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional columns and filters can be sorted in ascending or descending order. Choose sort order using the list. The optional filters are listed in the following table.

For information on saving the generated report, see Formats for Saving a Generated Report.

Permission Modification Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
User	User who is making the inquiry
Action	Action taken
Status	Status of the connection
Source IP	IP address of the source host device
Source Domain	Domain of the source host device

Option	Description
Target User	User for whom the inquiry is being made
Target IP	IP address of the accessed Appliance
Target Domain	Domain of the accessed Appliance
Type	Type of connection
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event
Application Type	The type of application that generated the event
Count	Number of connections

User Access Reports

To search for and generate a report on user activities in accessing resources (for example, service, file, directory, application) on selected log sources during a specified time interval, use the User Access Real-Time Report.

Menu path: **Reports > Access Control > User Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional columns and filters can be sorted in ascending or descending order. Choose sort order using the list. The optional filters are listed in the following table.

Option	Description
Source Device	Description of the device that sent these log messages
User	User who is making the inquiry
Source IP	IP address of the source host device
Source Domain	Domain of the source host device
Target User	User for whom the inquiry is being made
Target IP	IP address of the accessed Appliance
Target Domain	Domain of the accessed Appliance
Group	The name of the Policy group
Action	Action taken
Status	Status of the connection
Type	Type of connection
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event
Application Type	The type of application that generated the event
Count	Number of connections

User Authentication Reports

To search for and generate a report on who has authenticated on selected log sources during a specified time interval, use the User Authentication Real-Time Report.

Menu path: **Reports > Access Control > User Authentication**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display only the Source Device, User, Source IP, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

User Authentication Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
User	User who is making the inquiry
Source IP	IP address of the source host device
Source Domain	Domain of the source host device
Target User	User for whom the inquiry is made
Group	The name of the Policy group
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event

Option	Description
Application Type	The type of application that generated the event
Status	Status of the connection
Type	Type of connection
Disconnect Reason	Reason the connection was terminated
Count	Number of connections

User Created/Deleted Reports

To search for and generate a report on what users have been created or deleted on selected log sources during a specified time interval, use the Users Created/Deleted Real-Time Report.

Menu path: **Reports > Access Control > User Created/Deleted**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display only the Source Device, User, Source IP, Target User, Target IP, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

User Created/Deleted Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
User	User who is making the inquiry
Source IP	IP address of the source host device

Option	Description
Target User	User for whom the inquiry is being made
Target IP	IP address of the accessed Appliance
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event
Application Type	The type of application that generated the event
Action	Action taken
Action Details	Details of the action
Status	Status of use
Count	Number of connections

User Last Activity Reports

To search for and generate a report on the most recent activity of all users on selected log sources during a specified time interval, use the User Last Activity report.

Menu path: **Reports > Access Control > User Last Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional columns and filters can be sorted in ascending or descending order. Choose sort order using the list. The optional filters are listed in the following table.

Option	Description
Source Device	Description of the device that sent these log messages
Time	Time of connection
Connection ID	ID number for the connection
User	User who is making the inquiry
Source IP	IP address of the source host device
Target User	User for whom the inquiry is being made
Target IP	IP address of the accessed Appliance
Action	Action taken
Action Details	Details of the action
Status	Status of the activity
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event
Application Type	The type of application that generated the event
Access Details	Details of access

Windows Events Reports

To search for and generate a report on data on all Windows Event IDs, the number of events for each ID, and a description of each ID for selected log sources running the Microsoft Windows operating systems, use the Windows Events Real-Time Report.

For example, the captured log events include application, security, and system events.

Menu path: **Reports > Access Control > Windows Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display only the Source Device, Event ID, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Windows Events Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Event ID	Numeric ID corresponding to the source device
User	User ID on the source device
Source Domain	Domain name of the source device
Target User	User ID of the destination device
Target Domain	Domain name of the destination device
Originating Host	The original hostname where the event was originally created
Subsystem	The subsystem of the host
Originating IP	The original source IP address where the event was originally created
Event Name	Name of the event
Application Type	The type of application that generated the event
Action	Action taken
Status	Status of use

Option	Description
Type	Content type of the object as seen in the HTTP reply header
Count	Number of Windows events for the source device

Database Activity Reports

To search for and generate reports on various events occurring on database server log sources, use the **Database Activity** reports.

Accessing the Database Activity Reports

Choose **Reports > Database Activity > report-name** from the navigation menu, where *report-name* is any one of the following reports:

Database Activity Reports

Report	Description	More information
All Database Events	Use the All Database Events screen to search for and generate a report on the event types that are occurring.	All Database Events Reports
Database Access	Use the Database Access screen to search for and generate a report on all database server connections including user access and failed user access attempts.	Database Access Reports
Database Data Access	Use the Database Data Access screen to search for and generate a report on user access and changes to your data for a specified time period.	Database Data Access Reports
Database Privilege Modifications	Use the Database Privilege Modifications screen to search for and generate a report on database privilege changes, such as user reconfiguration and privilege manipulation.	Database Privilege Modifications Reports

Report	Description	More information
Database System Modifications	Use the Database System Modifications screen to search for and generate a report on system database changes such as drops and table drops.	Database System Modifications Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Database Activity report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

All Database Events Reports

To search for and generate a report on the event types that are occurring on specified database server log sources during a specified time interval, use the All Database Events Real-Time Report.

Menu path: **Reports > Database Activity > All Database Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are selected: Source Device, Database, Event Type ID, Event Type Name, Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

All Database Events Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log

Option	Description
	messages
Database	Database name on which the action occurred
DB User	User name of the database user whose actions were audited
Sys Priv	System privileges granted or revoked
Database Object Name	Name of the object affected by the action
Status	Status or return code of the action completion (numeric value)
Severity	Severity level of the event
OS User	Operating system login user name of the user whose actions were audited
Event Type ID	Database vendor audit code for the action type
Event Type Name	Type of database event such as DROP_TABLE, SQL_UPDATE, or CREATE_TABLE (names vary by vendor)
Object Priv	Object privileges granted or revoked on the database object
Count	Number of log entries returned with the given parameters

Database Access Report

To search for and generate a report on all database server connections, including user access and failed user access attempts, on specified database server log sources during a specified time interval, use the Database Access Real-Time Report.

Menu path: **Reports > Database Activity > Database Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are selected: Source Device, Database, Event Type ID, Event Type Name, Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Database Access Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent log data
Database	Database name on which the action occurred
DB User	User name of the database user whose actions were audited
Sys Priv	System privileges granted or revoked
Database Object Name	Name of the object affected by the action
Status	Status or return code of the action completion (numeric value)
Severity	Severity level of the event
OS User	Operating system login user name of the user whose actions were audited
Event Type ID	Database vendor audit code for the action type
Access Type	The action or method used to access any database object
Object Priv	Object privileges granted or revoked on the database object

Option	Description
Count	Number of log entries returned with the given parameters

Database Data Access Report

To search for and generate a report on user access and changes to your data on specified database server log sources during a specified time interval, use the Database Data Access Real-Time Report.

Menu path: **Reports > Database Activity > Database Data Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are selected: Source Device, Database, Event Type ID, Event Type Name, Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Database Data Access Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent log data
Database	Database name on which the action occurred
DB User	User name of the database user whose actions were audited
Sys Priv	System privileges granted or revoked
Database Object Name	Name of the object affected by the action

Option	Description
Status	Status or return code of the action completion (numeric value)
Severity	Severity level of the event
OS User	Operating system login user name of the user whose actions were audited
Event Type ID	Database vendor audit code for the action type
Access Type	The action or method used to access any database object
Object Priv	Object privileges granted or revoked on the database object
Count	Number of log entries returned with the given parameters

Database Privilege Modifications Report

To search for and generate a report on database privilege changes, such as user re-configuration and privilege manipulation, on specified database server log sources during a specified time interval, use the Database Privilege Modifications Real-Time Report.

Menu path: **Reports > Database Activity > Database Privilege Modifications**

In addition to setting the common report options in Real-Time Reports, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are selected: Source Device, Database, DB User, Modification Type, Object Priv, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Database Privilege Modifications Report - Optional Filter Operators

Advanced Option	Description
Source Device	Description of the device that sent log data
Database	Database name on which the action occurred
DB User	User name of the database user whose actions were audited
Sys Priv	System privileges granted or revoked
Database Object Name	Name of the object affected by the action
Status	Status or return code of the action completion (numeric value)
Severity	Severity level of the event
OS User	Operating system login user name of the user whose actions were audited
Event Type ID	Database vendor audit code for the action type
Modification Type	Modification action of a user, profile, or role privilege
Object Priv	Object privileges granted or revoked on the database object
Count	Number of log entries returned with the given parameters

Database System Modifications Report

To search for and generate a report on system database changes such as drops and table drops, use the Database System Modifications Real-Time Report.

Menu path: **Reports > Database Activity > Database System Modifications**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are selected: Source Device, Database, DB User, Database Object Name, Access/Modification Type, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Database System Modifications Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent log data
Database	Database name on which the action occurred
DB User	User name of the database user whose actions were audited
Sys Priv	System privileges granted or revoked
Database Object Name	Name of the object affected by the action
Status	Status or return code of the action completion (numeric value)
Severity	Severity level of the event
OS User	Operating system login user name of the user whose actions were audited
Event Type ID	Database vendor audit code for the action type
Access/Modification Type	Modification action of a user, profile, or role privilege

Option	Description
Object Priv	Object privileges granted or revoked on the database object
Count	Number of log entries returned with the given parameters

IBM i5/OS Activity Reports

To search for and generate reports on various events occurring on your IBM i5/OS log sources, use IBM i5/OS Activity reports.

Accessing the IBM i5/OS Activity Reports

Choose **Reports > IBM i5/OS Activity > report-name** from the navigation menu, where *report-name* is any one of the following reports:

IBM i5/OS Activity Reports

Report	Description	More information
All Log Entry Types	Use the IBM i5/OS Activity All Log Entry Types screen to search for and generate a report on all recorded entry types.	All Log Entry Types Reports
System Object Access	Use the IBM i5/OS Activity System Object Access screen to search for and generate a report on all failed access attempts throughout the system.	System Object Access Reports
User Access by Connection	Use the IBM i5/OS Activity User Access by Connection screen to search for and generate a report on all system access and system access attempts by user.	User Access By Connection Reports
User Actions	Use the IBM i5/OS Activity User Actions screen to search for	User Actions

Report	Description	More information
	and generate a report on all user actions performed and attempted.	Reports
User Jobs	Use the IBM i5/OS Activity User Jobs screen to search for and generate a report on all jobs that users are running.	User Jobs Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each IBM i5/OS Activity report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

All Log Entry Types Reports

To search for and generate a report on all recorded entry types, use the All Log Entry Types Real-Time Report.

Menu path: **Reports > IBM i5/OS Activity > All Log Entry Types**

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filter Operators

Option	Field	Description
Source Device	devIP	IP address of the device that sent log data

Option	Field	Description
Journal Type	jrnEntryType	Two-character Audit Journal record (entry) type
Journal Description	jrnTypeDesc	Description of the journal entry type
Journal Job	jobName	Name of the job that caused the entry to be created
Journal User	jrnUserName	Profile name of the user associated with Journal Job
Journal Number	jrnJobNbr	Job number of the Journal Job
Journal Program	jrnPgm	Name of the program that created the entry
Journal Library	jrnPgmLib	Program library
Journal System Name	jrnSyName	Name of the system where the journal resides
Journal Remote Port	jrnRmtPort	Remote port of the system associated with the journal entry
Journal Remote Address	jrnRmtIPAdr	Network address of the system associated with this entry
Action	action	An action associated with the entry type
Action Description	actionDesc	Description of the action
Attribute Name	attribute	Name of an attribute that

Option	Field	Description
		was the target of the action
Attribute Description	attributeDesc	Description of the attribute (if available)
Destination Server	destServer	Name of a remote workstation or server in a network event
DLO Folder	DLOFolder	Name of the Document Library Object folder
DLO User	DLOUser	Name of the Document Library Object owner or user creating or accessing the DLO
Entry Type	entryType	Type of event or entry in the journal type (can be considered a subtype of the journal type)
Entry Description	entryDesc	Description of the entry
Job Name	jobName	Name of the Journal Job or the job that was the target of the action described in the entry
Job Number	jobNumber	Number of the Journal Number or the job that was the target of the action described in the entry
Job User	jobUser	The Journal User of profile name of the user associated with the job that was the

Option	Field	Description
		target of the action described in the entry
Local IP Address	lclIPadr	Local IP address of the system involved in the network event
Object Library	lib	Library of the object that was acted on
Object Name	obj	Name of the object that was acted on
Object Type	objType	Type of object that was acted on
Remote IP Address	rmtIPadr	Remote IP address of the system involved in the network event
Source Server	srcServer	Name of a workstation or server where the audited event occurred, or that was the source system in a network event
Status	status	Status code
Status Description	statusDesc	Description of the status code (if available)
User ID/Profile	user	A user ID (UID) or user profile involved in the recorded event; typically the originator or target of the event
Journal Code	details	Provides event details.

Option	Field	Description
Count	(computed by the appliance)	A count of action attempts, entries, or other count information; dependent on Journal and Entry type

System Object Access Reports

To search for and generate a report on all failed access attempts throughout the system, use the System Object Access Real-Time Report.

Menu path: **Reports > IBM i5/OS Activity > System Object Access**

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filter Operators

Option	Field	Description
Source Device	devIP	IP address of the device that sent log data
Journal Type	jrnEntryType	Two-character Audit Journal record (entry) type
Journal Description	jrnTypeDesc	Description of the journal entry type
Journal Job	jobName	Name of the job that caused the entry to be created
Journal User	jrnUserName	Profile name of the user

Option	Field	Description
		associated with Journal Job
Journal Number	jrnJobNbr	Job number of the Journal Job
Journal Program	jrnPgm	Name of the program that created the entry
Journal Library	jrnPgmLib	Program library
Journal System Name	jrnSyName	Name of the system where the journal resides
Journal Remote Port	jrnRmtPort	Remote port of the system associated with the journal entry
Journal Remote Address	jrnRmtIPAdr	Network address of the system associated with this entry
Action	action	An action associated with the entry type
Action Description	actionDesc	Description of the action
Attribute Name	attribute	Name of an attribute that was the target of the action
Attribute Description	attributeDesc	Description of the attribute (if available)
Destination Server	destServer	Name of a remote workstation or server in a network event
DLO Folder	DLOFolder	Name of the Document

Option	Field	Description
		Library Object folder
DLO User	DLOUser	Name of the Document Library Object owner or user creating or accessing the DLO
Entry Type	entryType	Type of event or entry in the journal type (can be considered a subtype of the journal type)
Entry Description	entryDesc	Description of the entry
Job Name	jobName	Name of the Journal Job or the job that was the target of the action described in the entry
Job Number	jobNumber	Number of the Journal Number or the job that was the target of the action described in the entry
Job User	jobUser	The Journal User of profile name of the user associated with the job that was the target of the action described in the entry
Local IP Address	lclIPadr	Local IP address of the system involved in the network event
Object Library	lib	Library of the object that was acted on

Option	Field	Description
Object Name	obj	Name of the object that was acted on
Object Type	objType	Type of object that was acted on
Remote IP Address	rmtIPadr	Remote IP address of the system involved in the network event
Source Server	srcServer	Name of a workstation or server where the audited event occurred, or that was the source system in a network event
Status	status	Status code
Status Description	statusDesc	Description of the status code (if available)
User ID/Profile	user	A user ID (UID) or user profile involved in the recorded event; typically the originator or target of the event
Journal Code	details	Provides event details.
Count	(computed by the appliance)	A count of action attempts, entries, or other count information; dependent on Journal and Entry type

User Access by Connection Reports

To search for and generate a report on all system access and system access attempts by users, use the User Access By Connection Real-Time Report.

Menu path: **Reports > IBM i5/OS Activity > User Access By Connection**

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filter Operators

Option	Field	Description
Source Device	devIP	IP address of the device that sent log data
Journal Type	jrnEntryType	Two-character Audit Journal record (entry) type
Journal Description	jrnTypeDesc	Description of the journal entry type
Journal Job	jobName	Name of the job that caused the entry to be created
Journal User	jrnUserName	Profile name of the user associated with Journal Job
Journal Number	jrnJobNbr	Job number of the Journal Job
Journal Program	jrnPgm	Name of the program that created the entry
Journal Library	jrnPgmLib	Program library
Journal System Name	jrnSyName	Name of the system where the journal resides

Option	Field	Description
Journal Remote Port	jrnRmtPort	Remote port of the system associated with the journal entry
Journal Remote Address	jrnRmtIPAdr	Network address of the system associated with this entry
Action	action	An action associated with the entry type
Action Description	actionDesc	Description of the action
Attribute Name	attribute	Name of an attribute that was the target of the action
Attribute Description	attributeDesc	Description of the attribute (if available)
Destination Server	destServer	Name of a remote workstation or server in a network event
DLO Folder	DLOFolder	Name of the Document Library Object folder
DLO User	DLOUser	Name of the Document Library Object owner or user creating or accessing the DLO
Entry Type	entryType	Type of event or entry in the journal type (can be considered a subtype of the journal type)
Entry Description	entryDesc	Description of the entry

Option	Field	Description
Job Name	jobName	Name of the Journal Job or the job that was the target of the action described in the entry
Job Number	jobNumber	Number of the Journal Number or the job that was the target of the action described in the entry
Job User	jobUser	The Journal User of profile name of the user associated with the job that was the target of the action described in the entry
Local IP Address	lclIPadr	Local IP address of the system involved in the network event
Object Library	lib	Library of the object that was acted on
Object Name	obj	Name of the object that was acted on
Object Type	objType	Type of object that was acted on
Remote IP Address	rmtIPadr	Remote IP address of the system involved in the network event
Source Server	srcServer	Name of a workstation or server where the audited event occurred, or that was the source system in a

Option	Field	Description
		network event
Status	status	Status code
Status Description	statusDesc	Description of the status code (if available)
User ID/Profile	user	A user ID (UID) or user profile involved in the recorded event; typically the originator or target of the event
Journal Code	details	Provides event details.
Count	(computed by the appliance)	A count of action attempts, entries, or other count information; dependent on Journal and Entry type

User Actions Reports

To search for and generate a report on all user actions performed and attempted, use the User Actions Real-Time Report.

Menu path: **Reports > IBM i5/OS Activity > User Actions**

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filter Operators

Option	Field	Description
Source Device	devIP	IP address of the device that sent log data
Journal Type	jrnEntryType	Two-character Audit Journal record (entry) type
Journal Description	jrnTypeDesc	Description of the journal entry type
Journal Job	jobName	Name of the job that caused the entry to be created
Journal User	jrnUserName	Profile name of the user associated with Journal Job
Journal Number	jrnJobNbr	Job number of the Journal Job
Journal Program	jrnPgm	Name of the program that created the entry
Journal Library	jrnPgmLib	Program library
Journal System Name	jrnSyName	Name of the system where the journal resides
Journal Remote Port	jrnRmtPort	Remote port of the system associated with the journal entry
Journal Remote Address	jrnRmtIPAdr	Network address of the system associated with this entry
Action	action	An action associated with the entry type

Option	Field	Description
Action Description	actionDesc	Description of the action
Attribute Name	attribute	Name of an attribute that was the target of the action
Attribute Description	attributeDesc	Description of the attribute (if available)
Destination Server	destServer	Name of a remote workstation or server in a network event
DLO Folder	DLOFolder	Name of the Document Library Object folder
DLO User	DLOUser	Name of the Document Library Object owner or user creating or accessing the DLO
Entry Type	entryType	Type of event or entry in the journal type (can be considered a subtype of the journal type)
Entry Description	entryDesc	Description of the entry
Job Name	jobName	Name of the Journal Job or the job that was the target of the action described in the entry
Job Number	jobNumber	Number of the Journal Number or the job that was the target of the action described in the entry

Option	Field	Description
Job User	jobUser	The Journal User of profile name of the user associated with the job that was the target of the action described in the entry
Local IP Address	lclIPadr	Local IP address of the system involved in the network event
Object Library	lib	Library of the object that was acted on
Object Name	obj	Name of the object that was acted on
Object Type	objType	Type of object that was acted on
Remote IP Address	rmtIPadr	Remote IP address of the system involved in the network event
Source Server	srcServer	Name of a workstation or server where the audited event occurred, or that was the source system in a network event
Status	status	Status code
Status Description	statusDesc	Description of the status code (if available)
User ID/Profile	user	A user ID (UID) or user profile involved in the recorded event; typically the originator

Option	Field	Description
		or target of the event
Journal Code	details	Provides event details.
Count	(computed by the appliance)	A count of action attempts, entries, or other count information; dependent on Journal and Entry type

User Jobs Reports

To search for and generate a report on all jobs that users are running, use the User Jobs Real-Time Report.

Menu path: **Reports > IBM i5/OS Activity > User Jobs**

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filter Operators

Option	Field	Description
Source Device	devIP	IP address of the device that sent log data
Journal Type	jrnEntryType	Two-character Audit Journal record (entry) type
Journal Description	jrnTypeDesc	Description of the journal entry type

Option	Field	Description
Journal Job	jobName	Name of the job that caused the entry to be created
Journal User	jrnUserName	Profile name of the user associated with Journal Job
Journal Number	jrnJobNbr	Job number of the Journal Job
Journal Program	jrnPgm	Name of the program that created the entry
Journal Library	jrnPgmLib	Program library
Journal System Name	jrnSyName	Name of the system where the journal resides
Journal Remote Port	jrnRmtPort	Remote port of the system associated with the journal entry
Journal Remote Address	jrnRmtIPAdr	Network address of the system associated with this entry
Action	action	An action associated with the entry type
Action Description	actionDesc	Description of the action
Attribute Name	attribute	Name of an attribute that was the target of the action
Attribute Description	attributeDesc	Description of the attribute (if available)
Destination Server	destServer	Name of a remote

Option	Field	Description
		workstation or server in a network event
DLO Folder	DLOFolder	Name of the Document Library Object folder
DLO User	DLOUser	Name of the Document Library Object owner or user creating or accessing the DLO
Entry Type	entryType	Type of event or entry in the journal type (can be considered a subtype of the journal type)
Entry Description	entryDesc	Description of the entry
Job Name	jobName	Name of the Journal Job or the job that was the target of the action described in the entry
Job Number	jobNumber	Number of the Journal Number or the job that was the target of the action described in the entry
Job User	jobUser	The Journal User of profile name of the user associated with the job that was the target of the action described in the entry
Local IP Address	lclIPadr	Local IP address of the system involved in the network event

Option	Field	Description
Object Library	lib	Library of the object that was acted on
Object Name	obj	Name of the object that was acted on
Object Type	objType	Type of object that was acted on
Remote IP Address	rmtIPadr	Remote IP address of the system involved in the network event
Source Server	srcServer	Name of a workstation or server where the audited event occurred, or that was the source system in a network event
Status	status	Status code
Status Description	statusDesc	Description of the status code (if available)
User ID/Profile	user	A user ID (UID) or user profile involved in the recorded event; typically the originator or target of the event
Journal Code	details	Provides event details.
Count	(computed by the appliance)	A count of action attempts, entries, or other count information; dependent on Journal and Entry type

Threat Management Reports

To search for and generate reports on information about threat management, use the Threat Management reports.

Choose **Reports > Threat Management** from the navigation menu.

Threat Management Reports

Report	Description	More information
IDS/IPS Activity	Use the IDS/IPS Activity screen to search for and generate a report on all attack activities from Intrusion Detection/Prevention Systems (IDS/IPS).	IDS/IPS Activity Reports
Threat Activity	Use the Threat Activity screen to search for and generate a report on threats detected, eliminated, quarantined, and detected but unable to be mitigated.	Threat Activity Reports
Configuration Activity	Use the Configuration Activity screen to search for and generate a report on the following data; signature file installed, software update, configuration loaded.	Configuration Activity Reports
Scan Activity	Use the Scan Activity screen to search for and generate a report on the following data; scan delayed, scan aborted.	Scan Activity Reports
Security Summary	Use the Security Summary screen to search for and generate a report on summarized user and computer activity alongside other product's security interactions.	Security Summary Reports
DB IPS Activity	Use the DB IPS Activity screen to search for and generate a report on data (such as username, client/server IP addresses, so on) for various database intrusion prevention events.	DB IPS Activity Reports
HIPS Activity	Use the HIPS Activity screen to search for and generate a report on alerts from IPS/IDS signatures, DDOS attacks and port scan occurrences. See for more details.	HIPS Activity Reports

Preparing a Real-time Report includes the common options that you specify for Real-Time Reports.

For information on saving the generated report, see Formats for Saving a Generated Report.

IDS/IPS Activity Reports

To search for and generate a report on all attack activities from IDS/IPS systems, use the IDS/IPS Activity Real-Time Report.

Menu path: **Reports > Threat Management > IDS/IPS Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Log Source IP, Source IP, Destination IP, Destination Port, Signature, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

IDS/IPS Activity Report - Optional Filter Operators

Option	Description
Log Source IP	IP address of the device that sent these log messages
Source IP	IP address from which the attack originated
Source Port	Port from which the attack originated
Destination IP	IP address that was targeted
Destination Port	Port that was targeted
Action	Response of the intrusion prevention system (IPS) when it detects an attack reported by the IDS/IPS

Option	Description
	Note: If you do not have an IPS associated with your IDS/IPS, you might not see any results if using this filter.
Signature ID	Rule or numeric ID for the event **Note:** The Signature ID from the vendor might be more consistent than the Signature.
Protocol	Protocol of the destination device
Signature	Identifier from IDS/IPS for an event
Sensor	Device that sends events to a collector analysis system
Sensor IP	IP address of the device that detected the event
Classification	Type of attack
Priority	Priority level of the attack
Count	Number of attacks.

Threat Activity Reports

To search for and generate a report on all threats detected, eliminated, quarantined, and detected but unable to be mitigated, use the Threat Activity Real-Time Report.

Menu path: **Reports > Threat Management > Threat Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Event Name, Category , User Name, Target User, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Threat Activity Report - Optional Filter Operators

Option	Description
Source Device	IP address of the device that sent these log messages
Event ID	Numeric ID corresponding to the source device
Event Type	Type of event
Category	The category of the event
Event Response	Response to the event
Status ID	The ID of the status
Severity ID	The severity ID
Severity Name	The name of the severity code associated with the event
User Name	Name of the user who is making the inquiry
Target User	User for whom the inquiry is being made
Target Group	Group for who the inquiry is being made
Threat Name	Name of the threat
Source IP	IP address from which the attack originated
Destination IP	IP address that was targeted
Destination Host	Host that was targeted
Analyzer Name	Name of the analyzer

Option	Description
Analyzer Version	The version of the analyzer
Data Version	The version of the data associated with the event
Action	An action associated with the entry type
Status	Status of the connection
Count	Number of attacks.

Configuration Activity Reports

To search for and generate a report on all data such as; signature file installed, software update, and configuration loaded, use the Configuration Activity Real-Time Report.

Menu path: **Reports > Threat Management > Configuration Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Event Name, Category , User Name, Target User Name, Action, Status, and Count:

For information on saving the generated report, see Formats for Saving a Generated Report.

Configuration Activity Report - Optional Filter Operators

Option	Description
Source Device	Source device that sent these log messages
Event Name	Name of the event
Event Type	Type of event
Category	The category of the event

Option	Description
Severity ID	The severity ID
Severity Name	The name of the severity code associated with the event
User Name	Name of the user who is making the inquiry
Target User Name	User for whom the inquiry is being made
Threat Type	The type of threat associated with the event
Source IP	IP address from which the attack originated
Destination IP	IP address that was targeted
Analyzer Name	Name of the analyzer
Analyzer Version	The version of the analyzer
Data Version	The version of the data associated with the event
Action	An action associated with the entry type
Status	Status of the connection
Count	Number of attacks.

Scan Activity Reports

To search for and generate a report on all scan delayed or scan aborted data, use the Scan Activity Real-Time Report.

Menu path: **Reports > Threat Management > Scan Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Event Name, Category , User Name, Target User Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Scan Activity Report - Optional Filter Operators

Option	Description
Source Device	Source device that sent these log messages
Event Name	Name of the event
Event Type	Type of event
Category	The category of the event
Event Response	
Severity ID	The severity ID
Severity Name	The name of the severity code associated with the event
User Name	Name of the user who is making the inquiry
Target User Name	User for whom the inquiry is being made
Target Domain	Domain of the accessed Appliance
Target Group	Group for whom the inquiry is being made
Threat Name	The name of the threat

Option	Description
Threat Type	The type of threat associated with the event
Source IP	IP address from which the attack originated
Destination IP	IP address that was targeted
Destination Port	Port that was targeted
Analyzer Name	Name of the analyzer
Analyzer Version	The version of the analyzer
Action	An action associated with the entry type
Status	Status of the connection
Count	Number of attacks.

Security Summary Reports

To search for and generate a report on all summarized user and computer activity alongside other product's security interactions, use the Security Summary Real-Time Report.

Menu path: **Reports > Threat Management > Security Summary**

For this report, you can select to view various options in the generated report for your appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Source IP, Destination IP, User, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Security Summary Report - Optional Filter Operators

Option	Description
Source Device	Source device that sent these log messages
Source IP	IP address from which the attack originated
Destination IP	IP address that was targeted
Source Port	Port from which the attack originated
Destination Port	Port that was targeted
User	User who is making the inquiry
Source Host	Host from which the attack originated
Destination Host	Host that was targeted
Type	Type of connection
Event	Type of event
Action	An action associated with the entry type
Status	Status of the connection
Count	Number of attacks.

DB IPS Activity Reports

To search for and generate a report on all data (such as username, client/server IP addresses, so on) for various database intrusion prevention events, use the DB IPS Activity Real-Time Report.

Menu path: **Reports > Threat Management > DB IPS Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Client IP, Database User, Database IP, SQL Command, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

DB IPS Activity Report - Optional Filter Operators

Option	Description
Source Device	Source device that sent these log messages
Session ID	ID of the session
Client IP	IP address of the client
Client Hostname	Host name of the client
End User IP	IP address of the end user
Database User	Name of the database user
Database IP	IP address of the database
Database Hostname	Host name of the database
Database Name	Name of the database on which the action occurred
Schema	
Service Name	The name of the service
Database Type	The type of database

Option	Description
Database Port	The database port
SQL Command	
Object name	The name of the object
Source Program	
Count	Number of attacks.

HIPS Activity Reports

To search for and generate a report on all alerts from IPS/IDS signatures, DDOS attacks and port scan occurrences, use the HIPS Activity Real-Time Report.

Menu path: **Reports > Threat Management > HIPS Activity**

For this report, you can select to view various options in the generated report for your Appliance. Optional filter operators can be sorted in Ascending or Descending order. Choose sort order using the list. The default is to display only Source Device, Event Name, Target User, Threat Type, Source IP, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

HIPS Activity Report - Optional Filter Operators

Option	Description
Source Device	Source device that sent these log messages
Event ID	the ID of the event
Event Name	Name of the event

Option	Description
Event Type	The type of event
Event Response	
Severity Name	Name of the severity
Target User	User for whom the inquiry was made
Threat Type	The type of threat
Source IP	IP address from which the attack originated
Host IP	Host from which the attack originated
Destination IP	IP address that was targeted
Destination Host	Host that was targeted
Analyzer Name	Name of the analyzer
Analyzer Version	The version of the analyzer
Object Name	Name of the object affected
Destination Port	Port that was targeted
Target Process Name	Name of the target process
Count	Number of attacks.

Mail Activity Reports

To search for and generate reports on information about mail-related activities on mail server log sources, use Mail Activity reports.

The Report Information tab that appears when you click on **Reports > Mail Activity** lists which reports are available for each log source.

Choose **Reports > Mail Activity > report-name** from the navigation menu, where *report-name* is any one of the following reports:

Mail Activity Reports

Report	Description	More information
Exchange 2000/03 SMTP	Use the Exchange 2000/03 SMTP screen to search for and generate a report on all Exchange 2000/03 SMTP events recorded by your mail servers.	Exchange 2000/03 SMTP Reports
Exchange 2000/03 Activity	Use the Exchange 2000/03 Activity screen to search for and generate a report on all mail server activity for your Microsoft Exchange servers.	Exchange 2000/03 Activity Reports
Exchange 2000/03 Delay	Use the Exchange 2000/03 Delay screen to search for and generate a report on all delays in mail activity for your Microsoft Exchange servers.	Exchange 2000/03 Delay Reports
Exchange 2000/03 Size	Use the Exchange 2000/03 Size screen to search for and generate a report on mail size for all your Microsoft Exchange server mail activity.	Exchange 2000/03 Size Reports
Server Activity	Use the Server Activity screen to search for and generate a report on server activity.	Server Activity Reports
Exchange 2007/10 Activity	Use the Exchange 2007/10 Activity screen to search for and generate a report on all mail server activity for your Microsoft Exchange servers.	Exchange 2007/10 Activity Reports

Report	Description	More information
Exchange 2007/10 Mail Size	Use the Exchange 2007/10 Mail Size screen to search for and generate a report on mail size for all your Microsoft Exchange server mail activity.	Exchange 2007/10 Mail Size Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Mail Activity report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2000/03 SMTP Reports

To search for and generate a report on all mail server activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2000/03 Activity Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2000/03 SMTP**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, all options are shown except the Source User, Source Host, Domain Name, and Time Taken (ms).

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2000/03 SMTP Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages

Option	Description
Source User	User of the source device
Source IP	IP address of the source device
Source Host	Host name of the source device
Domain Name	Domain name of the source device
Destination IP	IP address of the destination device
Destination Port	Port of the destination device
Method	Request method to obtain an object; for example, GET
URL Query	URL requested
Status	SMTP result codes
Size	Number of bytes transferred
Time Taken (ms)	Time to complete the event
Count	Number of cache views

Exchange 2000/03 Activity Reports

To search for and generate a report on all delays in mail activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2000/03 Delay Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2000/03 Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Recipient Domain, Status, and Count are shown:

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2000/03 Activity Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device
Message ID	Numeric identifier of the message
Sender	Email address of the sender
Sender Domain	Domain name of the sender's email
Recipient	Email address of the recipient
Recipient Domain	Domain name of the recipient's email
Status	Exchange status
Count	Number of emails

Exchange 2000/03 Delay Reports

To search for and generate a report on all mail server activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2000/03 Activity Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2000/03 Delay**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Recipient Domain, Average Delay, Max Delay, and Count are shown.

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2000/03 Delay Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device
Message ID	Numeric identifier of the message
Sender	Email address of the sender
Sender Domain	Domain name of the sender's email
Recipient	Email address of the recipient
Recipient Domain	Domain name of the recipient's email
Average Delay	Average delay of each message
Max Delay	Maximum delay of each message
Count	Number of emails

Exchange 2000/03 Size Reports

To search for and generate a report on mail size for all server mail activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2000/03 Size Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2000/03 Size**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Sender, Total Size (Bytes), Max Size (Bytes), Count, and Actual Count are shown.

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2000/03 Size Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device
Message ID	Numeric identifier of the message
Sender	Email address of the sender
Sender Domain	Domain name of the sender's email
Recipient	Email address of the recipient
Recipient Domain	Domain name of the recipient's email
Total Size (Bytes)	Total number of bytes transferred
Max Size (Bytes)	Maximum number of bytes transferred
Count	Number of emails
Actual Count	

Server Activity Reports

To search for and generate a report on server activity, use the Server Activity Real-Time Report.

Menu path: **Reports > Mail Activity > Server Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Source IP, Source Port Destination IP, Destination Port , and Messages are shown.

For information on saving the generated report, see Formats for Saving a Generated Report.

Server Activity Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device
Source IP	IP address of the source host device
Source Port	Port of the source host device
Destination IP	IP address that was targeted
Destination Port	Port that was targeted
Messages	Number of log messages received representing this connection

Exchange 2007/10 Activity Reports

To search for and generate a report on all delays in mail activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2007/10 Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2007/10 Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Sender, Recipient, and Count are shown.

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2007/10 Activity Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device

Option	Description
Sender	Email address of the sender
Recipient	Email address of the recipient
Source	
Count	Number of emails

Exchange 2007/10 Mail Size Reports

To search for and generate a report on mail size for all server mail activity for selected Microsoft Exchange servers during a specified time interval, use the Exchange 2007/10 Mail Size Real-Time Report.

Menu path: **Reports > Mail Activity > Exchange 2007/10 Mail Size**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. By default, the Source Device, Sender, Total Size (Bytes), Max Size (Bytes), and Count are shown.

For information on saving the generated report, see Formats for Saving a Generated Report.

Exchange 2007/10 Mail Size Report - Optional Filter Operators

Option	Description
Source Device	Name of the Microsoft Exchange device
Sender	Email address of the sender
Total Size (Bytes)	Total number of bytes transferred
Max Size (Bytes)	Maximum number of bytes transferred
Count	Number of emails

Network Activity Reports

To search for and generate reports on information about connections on log sources, use Network Activity reports.

Choose **Reports > Network Activity > report-name** from the navigation menu, where *report-name* is any one of the following:

Network Activity Reports

Report	Description	More information
Accepted Connections	Use the Accepted Connections screen to search for and generate a report on IP connections that were accepted by a log source.	Accepted Connections Reports
Active FW Connections	Use the Active FW Connections screen to search for and generate a report on current active sessions from the selected firewall log sources.	Active FW Connections Reports
Active VPN Connections	Use the Active VPN Connections screen to search for and generate a report on current active sessions through Check Point Interface, Cisco VPN 3000, Nortel Connectivity, and RADIUS Acct Client log sources.	Active VPN Connections Reports
Application Distribution	Use the Application Distribution screen to search for and generate a report on information about messages, grouped by application ports, that were accepted by a device.	Application Distribution Reports
Denied Connections	Use the Denied Connections screen to search for and generate a report on connections denied by the selected firewall log sources.	Denied Connections Reports
FTP Connections	Use the FTP Connections screen to search for and generate a report on syslog messages related to FTP traffic through the selected firewall log sources.	FTP Connections Reports

Report	Description	More information
VPN Access	Use the VPN Access screen to search for and generate a report on the number of VPN connections that the log source either completed or denied.	VPN Access Reports
VPN Sessions	Use the VPN Sessions screen to search for and generate a report on data about separate invocations of sessions on log sources during a specified time interval.	VPN Sessions Reports
VPN Top Lists	Use the VPN Top Lists screen to search for and generate a report on the top users and IP addresses and statistics.	VPN Top Lists Reports
Web Cache Activity	Use the Web Cache Activity screen to search for and generate a report on locally stored web information served during a specified time interval.	Web Cache Activity Reports
Web Surfing Activity	Use the Web Surfing Activity screen to search for and generate a report on web information served during a specified time interval.	Web Surfing Activity Reports
DHCP Activity	Use the DHCP Activity screen to search for and generate a report on events related to all DHCP activity.	DHCP Activity Reports
DHCP Granted/Renewed Activity	Use the DHCP Granted/Renewed Activity screen to search for and generate a report on events related to DHCP requests that were granted or renewed.	DHCP Granted/Renewed Activity Reports
DHCP Denied Activity	Use the DHCP Denied Activity screen to search for and generate a report on events related to DHCP requests that were denied.	DHCP Denied Activity Reports
NAT64 Activity	Use the NAT64 Activity screen to search for and	NAT64 Activity

Report	Description	More information
	generate a report on each binding when sessions are built and destroyed.	Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Network Activity report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Accepted Connections Reports

To search for and generate a report on IP connections that were accepted by selected firewall log sources during a specified time interval, use the Accepted Connections Real-Time Report.

Note:
- Accepted Connections data is summarized in 10 minutes and 1 hour. If the report time interval is less than 2 hours, the time range is cut to 10 minutes, and if it is more than 2 hours, it is cut to 1 hour.
- To view the detail report, you must enable the **Administration > System Settings > General tab > Enable Accept Detail** option. This might require additional time and storage in downloading this report.

Menu path: **Reports > Network Activity > Accepted Connections**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the following options:

Note: Column headings differ for PIX and non-PIX devices.

Accepted Connections Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Translated IP	IP address as translated by the device*
Source IP	IP address of the source host (non-PIX devices only)
Destination IP	IP address of the destination host device (non-PIX devices only)
Port	Port number (service) of the destination host
Protocol	Protocol of the destination host
Description	Description of the port (service)
Messages	Number of log messages received representing this connection
In Bytes	Number of incoming bytes (Check Point Interface, Cisco PIX, and Juniper Firewall only)
Out Bytes	Number of outgoing bytes (Check Point Interface, Cisco PIX, and Juniper Firewall only)
Action	Accept or encrypt - Identifies if the connection was accepted or accepted with encryption (Check Point Interface only)

Note: * Under certain conditions Network Address Translation (NAT) addresses can show up as 0.0.0.0 in real time reports such as Accepted Connections Reports. This is not a bug since System Alert messages of a certain type (e.g., FWSM-4-106100 in Cisco Catalyst 6500 Series Switches) do not have a translated (mapped) address present in the logs. Therefore, zero is correct because there is no relevant IP address in the parsed logs for FWSM-4-106100.

For information on saving the generated report, see Formats for Saving a Generated Report.

Active FW Connections Reports

To search for and generate a report on current active sessions through selected Cisco PIX Firewall log sources, use the Active FW Connections Real-Time Report.

The Active Firewall Connection report is generated by monitoring the start and end messages of a particular connection in progress. Connections that have generated a start message but have not yet generated an end message are assumed to be active for a period of time before being timed-out.

Menu path: **Reports > Network Activity > Active FW Connections**

In Active FC Connections reports, you must specify the log source:

Element	Description
IP Address	IP address for the log source
Port	Port number for the log source
Protocol	Protocol type (from the list)

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the options.

For information on saving the generated report, see Formats for Saving a Generated Report.

Note: The generated list displays in real-time. As a result, the last page of connections might be closed/no longer active by the time you scroll to the last page. This results in no data displaying in the last page of the report

Active FW Connections Report - Optional Filter Operators

Option	Description
Create Time	Time the session began
Connection	ID in the log message assigned to the unique connection
Protocol	IP Protocol (TCP, UDP, so on) of the connection
Translated IP/Port	Public (NAT'ed) IP address of the source host (IP address only)
Source IP/Port	IP address of the internal host device (IP address only)
Destination IP/Port	IP address of the external host device (IP address only)
Direction	Inbound or Outbound connection attempt

Active VPN Connections Reports

To search for and generate a report on current active sessions through selected VPN and RADIUS log sources, use the Active VPN Connections Real-Time Report.

Menu path: **Reports > Network Activity > Active VPN Connections**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the following options.

For information on saving the generated report, see Formats for Saving a Generated Report.

Note: The generated list displays in real-time. As a result, the last page of connections might be closed/no longer active by the time you scroll to the last page. This results in no data displaying in the last page of the report.

Active VPN Connections Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Connections	Number of log messages received representing connections

Application Distribution Reports

To search for and generate a report that summarizes accepted traffic by application ports through selected firewall log sources during a specified time interval, use the Application Distribution Real-Time Report.

Note:

1. The Application Distribution data is summarized in 10 minutes and 1 hour. If the report time interval is less than 2 hours, the time range is cut to 10 minutes, and if it is more than 2 hours, it is cut to 1 hour.
2. To view the detail report, you must enable the **Administration > System Settings > General tab > Enable Accept Detail** option. This might require additional time and storage in downloading this report.

Menu path: **Reports > Network Activity > Application Distribution**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the following options.

For information on saving the generated report, see Formats for Saving a Generated Report.

Application Distribution Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Port	Port number (service) of the connection
Protocol	IP protocol (TCP, UDP, so on.) of the connection
Description	Description of the port (service)
Messages	Number of log messages received representing this connection
Src -> Dest Bytes	Number of outbound bytes sent (not for Nortel VPN)
Bar Graph	Percentage of total outbound bytes represented as a bar graph
Percentage	Number of outbound bytes represented as a percentage
Dst -> Src Bytes	Number of inbound bytes received (not for Nortel VPN)

Denied Connections Reports

To search for and generate a report on denied connections by selected firewall log sources during a specified time interval, use the Denied Connections Real-Time Report.

Menu path: **Reports > Network Activity > Denied Connections**

In addition to setting the common report options in Preparing a Real-time Report, you can select:

- The type of information the appliance aggregates for the generated report
- Various optional filter operators in the generated report for your appliance

Denied Connections Report - Summary Methods

Method	Description
Src IP/Any--> Any/Port	Aggregates records from a specific Source IP and any port going to any destination IP and a specific destination port. The system derives the Source IP and destination port from your Device Type and Source Device selections.
Src IP/Any --> Dest IP/Port	Aggregates records from a specific Source IP and any port going to a specific Destination IP and specific Destination port. The system derives the Source IP and Destination IP from your Device Type and Source Device selections.
Denied by Port	Aggregates records from the port numbers only

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the following optional filter operators.

For more information on saving the generated report, see Formats for Saving a Generated Report.

Denied Connections Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Attempts*	Number of times log messages denied the connection
Src IP	IP address of the source host device
Src Port	Port number of the source host device
Dest IP	IP address of the destination host device
Dest Port	Port number of the destination host device
Protocol	IP protocol (TCP, UDP, so on.) of the connection
Description	Description of the destination port (service)

Option	Description
Access Group	(Cisco PIX/ASA only) Lists any group of which you are a member
Rules	(Check Point Interface only) Condition set on the firewall to complete the security policy; identifies what is allowed and not allowed through a specific interface.
Policy ID	Unique policy identifier of the device on the firewall (Juniper Firewall only)
Direction	(Check Point Interface, Cisco PIX/ASA/FWSM, Juniper Firewall, and Nortel Connectivity only) Inbound or Outbound connection attempt. Direction is stored as a number internally, for INBOUND use 1, for OUTBOUND use 2, and for INTERNAL use 3.

Note: “Attempts” for Cisco router by “src IP/any” are larger than the number shown in the Denied Connections Report because IP packets are measured in this instance, instead of the actual number of messages sent.

FTP Connections Reports

To search for and generate a report on all syslog messages related to FTP traffic through the selected firewall device during a specified time interval, use the FTP Connections Real-Time Report.

Menu path: **Reports > Network Activity > FTP Connections**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the options.

For information on saving the generated report, see Formats for Saving a Generated Report.

FTP Connections Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Source Device IP	IP address of the source device that sent these log messages
From	IP address of the source device
To	IP address of the destination device
Count	Number of times syslog messages related to FTP traffic were generated

VPN Access Reports

To search for and generate reports on the VPN connections that the selected log sources either completed or denied during a specified time interval, use the VPN Access Real-Time Report.

Menu path: **Reports > Network Activity > VPN Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the options.

For information on saving the generated report, see Formats for Saving a Generated Report.

VPN Access Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages

Option	Description
Public IP	Public IP address originating the VPN connection
Group	VPN group of which the source device is a part
User	VPN user ID
Target User	VPN user ID of the originating VPN connection
Connections	Number of log messages received representing connections
Denies	Number of denied connection messages received
Avg Duration	Average duration of each connection
Byte Count	Number of bytes transferred during the session
Avg Bandwidth (Bytes/Sec)	Average bandwidth used for each connection

Appliances cannot receive disconnected messages. A VPN session is recorded permanently in the database table authentication after it is disconnected, prior to that the session is considered active. A Check Point VPN session is considered disconnected when a new connection attempt is made by the same user from the same IP address.

VPN Sessions Reports

To search for and generate a report on data about VPN sessions (including initiation and conclusion times) on selected log sources during a specified time interval, use the VPN Sessions Real-Time Report.

Menu path: **Reports > Network Activity > VPN Sessions**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display only the Source Device, User, Avg Duration, Avg Bytes, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

VPN Sessions Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
User	User ID
Target User	User ID on the device with which the source device attempted to connect
Source IP	IP address of the device that sent these log messages
Target IP	IP address of the device with which the source device attempted to connect
Avg Duration	Average duration of each connection
Avg Bytes	Average number of bytes
Count	Number of VPN sessions

Appliances cannot receive disconnected messages. A VPN session is recorded permanently in the database table authentication after it is disconnected, prior to that the session is considered active. A Check Point VPN session is considered disconnected when a new connection attempt is made by the same user from the same IP address.

VPN Top Lists Reports

To search for and generate a report on the top users, IP addresses, and other statistics, use the VPN Top Lists Real-Time Report.

Menu path: **Reports > Network Activity > VPN Top Lists**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Choose the Method from the list. The options are: Top Disconnect Reasons, By IP Address, and By User. Depending on the method selection, the default column options changes. Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. The default is to display all the options for Top Disconnect Reasons

For information on saving the generated report, see Formats for Saving a Generated Report.

VPN Top Lists Report - Types

Report Type	Description
Source Device	The description of the source device
Connections	Number of connections to the source device
Disconnect Reason	Reason for disconnection

Warning: If you run a report for the Top Disconnect Reasons, the “unknown” that displays in the Disconnect Reasons column, represents the disconnect reasons reported by RADIUS. If you have not properly plugged in your RADIUS server, all reasons display as “unknown”. Click a Connections number or Source Device to drill-down and view the Disconnect Details column. This column displays the VPN syslog messages associated with the disconnect reason.

Web Cache Activity Reports

To search for and generate a report on all URLs accessed through proxy or cache servers on specified log sources during a specified time interval, use the Web Cache Activity Real-Time Report.

Menu path: **Reports > Network Activity > Web Cache Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source IP, Destination IP, Status, Size, Filter Category, Filter Result, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Web Cache Activity Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent these log messages
Source User	User of the source device
Source IP	IP address of the source device
Source Host	Host name of the source device
Domain Name	Domain name of the source device
Destination IP	IP address of the destination device
Destination Port	Port of the destination device
Peer IP	IP address of the peer device
Peer Host	Host name of the peer device
Peer Status	A code that explains how the request was handled; for example, by forwarding it to a peer or returning the request to the source
Method	Request method to obtain an object; for example, GET
URL	URL requested

Option	Description
Cache Code	Information on the result of the transaction: the kind of request, how it was satisfied, or in what way it failed
Status	HTTP result codes
Type	Content type of the object as seen in the HTTP reply header
Size	Number of bytes transferred
Filter Category	The category of the filter
Filter Result	The results after using the filter
Count	Number of cache views

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

Web Surfing Activity Reports

To search for and generate a report on all URLs accessed via firewalls or web servers on selected log sources during a specified time interval, use the Web Surfing Activity Real-Time Report.

Menu path: **Reports > Network Activity > Web Surfing Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device IP, Source IP, Destination IP, Status, Size, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report

Web Surfing Activity Report - Optional Filter Operators

Option	Description
Source Device IP	IP address of the device that sent these log messages
Source User	User ID of the source device
Source IP	IP address of the device originating the connection
Source Host	Host name of the source device
Domain Name	Domain name of the source device
Destination IP	IP address of the destination device
Destination Port	Port of the destination device
Method	Request method to obtain an object; for example, GET
URL	URL requested
Status	HTTP result codes
Type	Content type of the object as seen in the HTTP reply header
Size	Number of bytes transferred
User Agent	
Referred By	
Count	Number of syslog messages received for this connection and status code

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

DHCP Activity Report

To search for and generate a report on events related to all DHCP activity, use the DHCP Activity Real-Time Report.

Menu path: **Reports > Network Activity > DHCP Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, MAC Address, Client Name, Lease Address, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
MAC Address	MAC IP address
Client Name	Name of the client
Lease Address	
Action	Action taken
Status	Status of the activity
Count	Number of connections

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

DHCP Granted/Renewed Activity Reports

To search for and generate a report on events related to DHCP requests that were granted or renewed, use the DHCP Granted/Renewed Activity Real-Time Report.

Menu path: **Reports > Network Activity > DHCP Granted/Renewed Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, MAC Address, Client Name, Lease Address, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
MAC Address	MAC IP address
Client Name	Name of the client
Lease Address	
Action	Action taken
Status	Status of the activity
Count	Number of connections

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

DHCP Denied Activity Report

To search for and generate a report on events related to DHCP requests that were denied, use the DHCP Denied Activity Real-Time Report.

Menu path: **Reports > Network Activity > DHCP Denied Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, MAC Address, Client Name, Lease Address, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
MAC Address	MAC IP address
Client Name	Name of the client
Lease Address	
Action	Action taken
Status	Status of the activity
Count	Number of connections

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

NAT64 Activity Reports

To search for and generate a report on each binding when sessions are built and destroyed, use the NAT64 Activity Real-Time Report.

Menu path: **Reports > Network Activity > NAT64 Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Time, Translated IPv6, Original IPv4, Original IPv6 Port, Original IPv4 Port, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Time	Time of connection
Translated IPv6	The translated IPv6 address
Original IPv4	The original IPv4 address
Original IPv6 port	The port of the original IPv6
Original IPv4 port	The port for the original IPv4
Count	Number of connections

When you drill down on the report results, there is no default sort-by selection. The drill-down results are generally in order by time. If you specify a sort-by selection for this report's drill-down, performance in generating the drill-down results is slower.

Operational Reports

To search for and generate reports on information about syslog messages on log sources, use Event Logs reports.

The Report Information tab that appears when you click on **Reports > Operational** lists which reports are available for each log source.

Choose **Reports > Operational > *report-name*** from the navigation menu, where *report-name* is any one of following reports:

Operational Reports

Report	Description	More information
All Unparsed Events	Use the All Unparsed Events screen to search for and generate a report on unparsed syslog messages for selected devices.	All Unparsed Events Reports
Firewall Statistics	Use the Firewall Statistics screen to search for and generate a report summarizing firewall syslog messages classified as security messages.	Firewall Statistics Reports
Total Message Count	Use the Total Message Count screen to search for and generate a report summarizing firewall or Nortel VPN device syslog messages classified as system messages.	Total Message Count Reports
Security Events	Use the Security Events screen to search for and generate a report on firewall syslog messages classified as security messages.	Security Events Reports
System Events	Use the System Events screen to search for and generate a report on firewall or Nortel VPN device syslog messages classified as system messages.	System Events Reports

Report	Description	More information
VPN Events	Use the VPN Events screen to search for and generate a report on the number of Cisco VPN syslog messages that appear with the type called “System Type”.	VPN Events Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

All Unparsed Events Reports

To search for and generate a report on syslog messages that are not parsed into the Security, System, or VPN Events reports, or into any other report table (for example, Authentication) for selected log sources during a specified time interval, use the All Unparsed Events Real-Time Report.

Menu path: **Reports > Operational > All Unparsed Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report. Optional filter operators are not visible if you select the Boolean Search in the Search Filter criteria.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are all selected:

For information on saving the generated report, see Formats for Saving a Generated Report.

All Unparsed Events Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent the log messages

Option	Description
Source Device IP	IP address of the source device that sent the log messages
Facility	Syslog facility associated with the message
Severity	Severity code associated with the message
Count	Number of times syslog messages were generated

Firewall Statistics Reports

To search for and generate a summary report of event types and messages per firewall, for selected log sources during a specified time interval, use the All Unparsed Events Real-Time Report.

Menu path: **Reports > Operational > Firewall Statistics**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report. Optional filter operators are not visible if you select the Boolean Search in the Search Filter criteria.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are all selected.

For information on saving the generated report, see Formats for Saving a Generated Report.

Firewall Statistics Report - Optional Filter Operators

Option	Description
Source Device	Description of the device that sent the log messages
System Messages	The number of system messages
Security Messages	The number of security messages

Option	Description
Accepted Messages	The number of accepted messages
Denied Messages	The number of denied messages
Total Messages	The total number of messages

Total Message Count Reports

To search for and generate a summary report of log messages for selected log sources at a specified time interval, use the Total Message Count Report.

Menu path: **Reports > Operational > Total Message Count**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report. Optional filter operators are not visible if you select the Boolean Search in the Search Filter criteria.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are all selected.

For information on saving the generated report, see Formats for Saving a Generated Report.

Total Message Count Report - Optional Filter Operators

Option	Description
Time	Time the syslog message was generated
Source Device	Description of the device that sent the log messages
Messages	The total number of messages

Security Events Reports

To search for and generate a report on firewall syslog messages classified as security messages for selected log sources during a specified time interval, use the Security Events Real-Time Report.

Menu path: **Reports > Operational > Security Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. By default, the following options are all selected.

For information on saving the generated report, see Formats for Saving a Generated Report.

Security Events Report - Optional Filter Operators

Option	Description
Source Device	Description of the device originating the connection
Source Device IP	IP address of the source device
Message Code	Code number of the security message
Message Code Description	Description of the security message (Cisco PIX only)
Module	Juniper Netscreen module name, that is, system (Juniper Firewall only)

Option	Description	
Severity	List of severity codes:	
	0	Emergency: system is unusable
	1	Alert: action must be taken immediately
	2	Critical: critical conditions
	3	Error: error conditions
	4	Warning: warning conditions
	5	Notice: normal but significant condition
	6	Informational: informational messages
	7	Debug: debug-level messages
	(Juniper Firewall only)	
Count	Number of syslog messages classified as security messages generated	

System Events Reports

To search for and generate a report on firewall or Nortel VPN device syslog messages classified as system messages for selected log sources during a specified time interval, use the System Events Real-Time Report.

Menu path: **Reports > Operational > System Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

Optional filter operators can be sorted in ascending or descending order. Choose sort order using the list. Optional filter operators are not visible if you select Boolean Search in the Search Filter criteria. By default, the following options are all selected.

For information on saving the generated report, see Formats for Saving a Generated Report.

System Events Report - Optional Filter Operators

Option	Description
Source Device	Description of the device originating the connection
Source Device IP	IP address of the source device
Message Code	Code number of the security message
Message Code Description	Description of the security message (Cisco PIX only)
Module	Juniper Netscreen module name, that is, system (Juniper Firewall only)

Option	Description	
Severity	List of severity codes:	
	0	Emergency: system is unusable
	1	Alert: action must be taken immediately
	2	Critical: critical conditions
	3	Error: error conditions
	4	Warning: warning conditions
	5	Notice: normal but significant condition
	6	Informational: informational messages
	7	Debug: debug-level messages
	(Juniper Firewall only)	
Count	Number of syslog messages classified as security messages generated	

VPN Events Reports

To search for and generate a report on Cisco VPN, Check Point VPN, Nortel VPN, or RADIUS syslog messages of the System Message type for selected log sources during a specified time interval, use the VPN Events Real-Time Report.

Menu path: **Reports > Operational > VPN Events**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

By default, the following options are all selected.

For information on saving the generated report, see Formats for Saving a Generated Report.

VPN Events Report - Optional Filter Operators

Option	Description
Time	Time the syslog message was generated
Source Device	IP address of the device originating the connection
Group	VPN group name
User	VPN user ID
Public IP	Public IP address originating the VPN connection
Severity	Severity Code associated with the message
Code	Code number of the system message
Area	Name of the defined VPN area
Detail Message	Text of the syslog message

Appliances cannot receive disconnected messages. A VPN session is recorded permanently in the database table authentication after it is disconnected, prior to that the session is considered active. A Check Point VPN session is considered disconnected when a new connection attempt is made by the same user from the same IP address.

Policy Reports

To search for and generate reports on information about policies that were exercised on a log source, use Policy reports.

The Report Information tab that appears when you click on **Reports > Policy Reports** lists which reports are available for each log source.

Choose **Reports > Policy Reports > *report-name*** from the navigation menu, where *report-name* is one of:

Policy Reports

Report	Reports Provide	More information
Check Point Policies	The Check Point Policies report lists current Check Point Firewall policy rules on log sources connected to your Appliance.	Check Point Policies Reports
Network Policies	Use the Network Policies screen to search for and generate a report on the number of times a particular network policy has been exercised by a selected firewall device.	Network Policies Reports
Rules/Policies	Use the Rules/Policies screen to search for and generate a report on enforcement of a particular rule or policy by a selected firewall device.	Rules/Policies Reports
ECM Policy	Use the ECM Policy screen to search for and generate a report on data leak protection events captured by the log source device.	ECM Policy Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Check Point Policies Reports

To search for and generate a report listing current Check Point Firewall policy rules on log sources connected to your appliance, use the Check Point Policy Real-Time Report.

Menu path: **Reports > Policy Reports > Check Point Policy**

For information on saving the generated report, see Formats for Saving a Generated Report.

Check Point Policy Screen Elements

Element	Description
LEA Server	LEA servers connected to your system.
Package	Security package that Check Point organizes for policy rules. For example, you can install one package on a firewall, but you can define several packages at the same time.
Rule Index	Rule numbers (represents Check Point indices) the CPMI process retrieves. You can view Check Point policy rules only if you configured your LEA server to use auto discovery (CPMI). **Note:** Rule 0 is not assigned by Check Point. It is assigned by LogLogic as a default for parsed messages that do not automatically have a rule number assigned by Check Point.
Rule	Description of the rule.

Network Policies Reports

To search for and generate a report on the number of times a particular network policy has been exercised by selected firewall log sources during a specified time interval, use the Network Policies Real-Time Report.

Menu path: **Reports > Policy Reports > Network Policies**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Log Source IP, Source IP, Destination IP, Destination Port, Signature, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Network Policies Report - Optional Filter Operators

Option	Description
Log Source IP	IP address of the device that sent these log messages
Source IP	IP address of the device that exercised the policy
Source Port	Port of source device
Destination IP	IP address of the destination device
Destination Port	Port of the destination device
Protocol	Protocol of the destination device
Signature	Identifier of the policy
Classification	Classification of the policy
Priority	Priority of the policy
Count	Number of times a policy was exercised

Rules/Policies Reports

To search for and generate a report on information about enforcement of a particular rule or policy by selected firewall devices during a specified time interval, use the Rules/Policies Real-Time Report.

Menu path: **Reports > Policy Reports > Rules/Policies**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display all the following options:

For information on saving the generated report, see Formats for Saving a Generated Report.

Rules/Policies Report - Optional Filter Operators

Option	Description
Interface	Name (or IP address) of the network interface that enforced the policy
Rule	Rule number that was enforced (Check Point Interface only)
Policy	Policy number that was enforced
Type	Type of rule/policy that was enforced
Messages	Number of messages received representing this policy
Bar Graph	Number of messages received expressed as a bar graph
Percentage	Number of messages received expressed as a percentage
Package	Security policy package (Check Point Interface only)
Rule Description	Displays Rule Details: Source, Destination, Service Description and Rule Actions: Permit, Deny, so on. (Check Point Interface only)

ECM Policy Reports

To search for and generate a report on data leak protection events captured by the log source device use the ECM Policy Real-Time Report.

Menu path: **Reports > Policy Reports > ECM Policy**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source Device IP, Performer Name, Parent Name, Event, Event Name, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

ECM Policy Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Source Device IP	IP address of the device that exercised the policy
Performer Name	Name of the performer
Parent Name	Name of the parent
Object Name	Name of the object that was acted on
Event	The type of event
Event Name	Name of the event
Source Name	Name of the source host device
Count	Number of attacks

Enterprise Content Management Reports

To search for and generate reports on information about enterprise content management, use Enterprise Content Management reports.

The Report Information tab that appears when you click on **Reports > Enterprise Content Management** lists which reports are available for each log source.

Choose **Reports > Enterprise Content Management > *report-name*** from the navigation menu, where *report-name* is one of:

Policy Reports

Report	Reports Provide	More information
ECM Activity	Use the ECM Activity screen to generate a report for ECM activity.	ECM Activity Reports
Content Management	Use the Content Management screen to generate a report containing logs of events which correspond to some action done on the contents of the site.	Content Management Reports
Security Settings	Use the Security Settings screen to generate a report containing logs of all the events related to creation, deletion, modification of user/group/roles.	Security Settings Reports
Expiration and Disposition	Use the Expiration and Disposition screen to generate a report containing logs of all events related to object expiration and disposition approvals.	Expiration and Disposition Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

ECM Activity Reports

To search for and generate a report on ECM activity use the ECM Activity Real-Time Report.

Menu path: **Reports > Enterprise Content Management > ECM Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source Device IP, Performer Name, Parent Name, Event, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

ECM Activity Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Source Device IP	IP address of the device that exercised the policy
Performer Name	Name of the performer
Parent Name	Name of the parent
Object Name	Name of the object that was acted on
Event	The type of event
Event Name	Name of the event
Source Name	Name of the source host device
Source IP	IP address of the source host
Destination IP	IP address that was targeted
Source Port	Port from which the attack originated
Destination Port	Port that was targeted
Protocol	Protocol of the destination device
Count	Number of attacks

For information on saving the generated report, see Formats for Saving a Generated Report.

Content Management Reports

To search for and generate a report containing logs of events which correspond to some action done on the contents of the site use the Content Management Real-Time Report.

Menu path: **Reports > Enterprise Content Management > Content Management**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source Device IP, Performer Name, Parent Name, Object Type, Event, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Content Management Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Source Device IP	IP address of the device that exercised the policy
Performer Name	Name of the performer
Parent Name	Name of the parent
Object Type	Type of object that was acted on
Object Name	Name of the object that was acted on
Event	The type of event
Event Name	Name of the event
Source Name	Name of the source host device
Count	Number of attacks

Security Settings Reports

To search for and generate a report containing logs of all the events related to creation, deletion, modification of user/group/roles use the Security Settings Real-Time Report.

Menu path: **Reports > Enterprise Content Management > Security Settings**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source Device IP, Performer Name, Parent Name, Event, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Security Settings Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Source Device IP	IP address of the device that exercised the policy
Performer Name	Name of the performer
Parent Name	Name of the parent
Object Name	Name of the object that was acted on
Event	The type of event
Event Name	Name of the event
Source Name	Name of the source host device
Count	Number of attacks

For information on saving the generated report, see Formats for Saving a Generated Report.

Expiration and Disposition Reports

To search for and generate a report containing logs of all events related to object expiration and disposition approvals use the Expiration and Disposition Real-Time Report.

Menu path: **Reports > Enterprise Content Management > Expiration and Disposition**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Source Device IP, Performer Name, Parent Name, Object Name, Event, Event Name, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Expiration and Disposition Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Source Device IP	IP address of the device that exercised the policy
Performer Name	Name of the performer
Parent Name	Name of the parent
Object Name	Name of the object that was acted on
Event	The type of event
Event Name	Name of the event
Source Name	Name of the source host device
Count	Number of attacks

HP NonStop Audit

To search for and generate reports on information about HP NonStop systems and generate Audit and EMS log data , use HP NonStop Audit reports.

The **Report Information** tab that appears when you click on **Reports > HP NonStop Audit** lists the reports that are available for each log source.

Choose **Reports > HP NonStop Audit > *report-name*** from the navigation menu, where *report-name* is one of:

HP NonStop Audit Reports

Report	Reports Provide	More information
Configuration Changes	Use the Configuration Changes screen to generate a report for all configuration changed done on an HP NonStop server during a specified time.	Configuration Changes Reports
Failed and Successful Logins	Use the Failed and Successful Logins screen to generate a report for all successful and failed logins on an HP NonStop Audit server.	Failed and Successful Logins Reports
Object Changes	Use the Object Changes screen to generate a report for all objects that are accessed on an HP NonStop Audit server.	Object Changes Reports
HP NonStop Audit Activity	Use the HP NonStop Audit Activity screen to generate a report for all audit activities on an HP NonStop Audit server.	HP NonStop Audit Activity Reports
User Actions	Use the User Actions screen to generate a report for all user actions done on an HP NonStop Audit server.	User Actions Reports
Object Access	Use the Object Access screen to generate a report for a list of all objects created, deleted, or modified on an HP NonStop Audit server.	Object Access Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Configuration Changes Reports

To search for and generate a report on all configuration changes done on an HP NonStop server during a specified time use the Configuration Changes Real-Time Report.

Menu path: **Reports > HP NonStop Audit > Configuration Changes**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Configuration Changes Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User Name	Username of the creator
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported

Option	Description
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken
Status	Status of the connection
Count	Number of attacks

Failed and Successful Logins Reports

To search for and generate a report for all successful and failed logins on an HP NonStop Audit server use the Failed and Successful Logins Real-Time Report.

Menu path: **Reports > HP NonStop Audit > Failed and Successful Logins**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Failed and Successful Logins Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User	Username of the creator

Option	Description
Name	
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken
Status	Status of the connection
Count	Number of attacks

Object Changes Reports

To search for and generate a report for all objects that are accessed on an HP NonStop Audit server use the Object Changes Real-Time Report.

Menu path: **Reports > HP NonStop Audit > Object Changes**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Object Changes Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User Name	Username of the creator
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken
Status	Status of the connection
Count	Number of attacks

HP NonStop Audit Activity Reports

To search for and generate a report for all audit activities on an HP NonStop Audit server use the HP NonStop Audit Activity Real-Time Report.

Menu path: **Reports > HP NonStop Audit > HP NonStop Audit Activity**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

HP NonStop Audit Activity Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User Name	Username of the creator
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken
Status	Status of the connection
Count	Number of attacks

User Actions Reports

To search for and generate a report for all user actions done on an HP NonStop Audit server use the User Actions Real-Time Report.

Menu path: **Reports > HP NonStop Audit > User Actions**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count:

For information on saving the generated report, see Formats for Saving a Generated Report.

User Actions Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User Name	Username of the creator
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken
Status	Status of the connection
Count	Number of attacks

Object Access Report

To search for and generate a report for a list of all objects created, deleted, or modified on an HP NonStop Audit server use the Object Access Real-Time Report.

Menu path: **Reports > HP NonStop Audit > Object Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User Name, Creator User Name, Target User, Event Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Object Access Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User Name	Name of the user making the inquiry
Creator User Name	Username of the creator
Target User	User for whom the inquiry is being made
User Group	Name of the group
Reported Time	Time the event was reported
Process Name	Name of the process
Event Name	Name of the event
Object Type	Type of object that was acted on
Action	Action taken

Option	Description
Status	Status of the connection
Count	Number of attacks

IBM z/OS Activity

To search for and generate reports on information about IBM z/OS system generated operational and audit logs in binary format, use IBM z/OS Activity reports.

The Report Information tab that appears when you click on **Reports > IBM z/OS Activity Reports** lists which reports are available for each log source.

Choose **Reports > IBM z/OS Activity Reports > *report-name*** from the navigation menu, where *report-name* is one of:

IBM z/Activity Reports

Report	Reports Provide	More information
Resource Access	Use the Resource Access screen to generate a report for resource access on z/OS.	Resource Access Reports
Security Modifications	Use the Security Modification screen to generate a report for security modification on z/OS.	Security Modifications Reports

Report	Reports Provide	More information
System Access/Configuration	Use the System Access/ Configuration screen to generate a report for access and configuration on z/OS.	System Access/Configuration Reports
Unix System Services	Use the Unix System Services screen to generate a report for Unix system services on z/OS.	Unix System Services Reports
Login/Logout	Use the Login/Logout screen to generate a report for login and logout activities on z/OS.	Login/Logout Reports
Violation	Use the Violation screen to generate a report for violation activities on z/OS.	Violation Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Resource Access Reports

To search for and generate a report on resource access on z/OS use the Resource Access Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > Resource Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Resource Access Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event

Option	Description
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Target Object Name	Name of the object that was acted on
Target Object Type	Type of target object that was acted on
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

Security Modifications Reports

To search for and generate a report for security modification activities on z/OS use the Security Modifications Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > Security Modifications**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Event Type, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Security Modifications Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Target Object Name	Name of the object that was acted on
Target Object Type	Type of target object that was acted on
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

System Access/Configuration Reports

To search for and generate a report for access and configuration activities on z/OS use the System Access/Configuration Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > System Access/Configuration**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Event Type, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

System Access/Configuration Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

Unix System Services Reports

To search for and generate a report for Unix system services on z/OS use the Unix System Services Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > Unix System Services**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Event Type, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Unix System Services Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Target Object Name	Name of the object that was acted on

Option	Description
Target Object Type	Type of target object that was acted on
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

Login/Logout Reports

To search for and generate a report for login and logout activities on z/OS use the Login/Logout Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > Login/Logout**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Event Type, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Login/Logout Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type

Option	Description
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event
Target User	User for whom inquiry is being made
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

Violation Reports

To search for and generate a report for violation activities on z/OS use the Violation Real-Time Report.

Menu path: **Reports > IBM z/OS Activity > Violation**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Record Type Description, Event Type, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Violation Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Record Type ID	The ID of the record type
Record Type Description	Description of the record type
SubType Description	Description of the sub type
Event Type	Type of event in the journal type
Logon ID/User ID	A user ID or login ID involved in the recorded event
Target User	User for whom inquiry is being made
Job Name	Name of the journal job or the job that was the target of the action described in the entry
Target Object Name	Name of the object that was acted on
Target Object Type	Type of target object that was acted on
Violation Occurred	
Action	Action taken
Status	Status of the connection
Count	A count of action attempts, entries, or other count information dependent on journal and entry type.

Storage Systems Activity Reports

To search for and generate reports on information about file and directory access, use Storage Systems Activity reports.

The Report Information tab that appears when you click on **Reports > Storage Systems Activity Reports** lists which reports are available for each log source.

Menu path: **Reports > Storage Systems Activity Reports > *report-name***

Storage Systems Activity Reports

Report	Reports Provide	More Information
Filer Access	Use the Filer Access screen to generate a report for individual file and directory access events such as user, timestamp, result, on z/OS.	Filer Access Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filer Access Reports

To search for and generate a report for individual file and directory access events use the Filer Access Real-Time Report.

Menu path: **Reports > Storage Systems Activity > Filer Access**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, User, Filer IP, Filer Name, Action, Status, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Filer Access Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
User	User who is making the inquiry
Source IP	IP address of the source host device
Target User	User for whom inquiry is being made
Filer IP	IP address of the filer
Filer Name	name of the filer
Action	Action taken
Status	Status of the connection
Count	Number of connections

Flow Activity Reports

To search for and generate reports on information about application usage, user browsing and top users, use Flow Activity reports.

The Report Information tab that appears when you click on **Reports > Flow Activity Reports** lists which reports are available for each log source.

Choose **Reports > Flow Activity Reports > *report-name*** from the navigation menu, where *report-name* is one of:

Flow Activity Reports

Report	Reports Provide	More information
Application Usage	Use the Application Usage screen to generate a report for application usage seen across all traffic.	Application Usage Reports
User Browsing Statistics	Use the User Browsing Statistics screen to generate a report for site destination statistics by user.	User Browsing Statistics Reports
Top Users	Use the Top Users screen to generate a report for top traffic users.	Top Users Reports

Preparing a Real-time Report includes the common options that you specify for all Real-Time Reports.

Optional filter operators are different for each Event Logs report, and explained in their respective sections.

For information on saving the generated report, see Formats for Saving a Generated Report.

Application Usage Reports

To search for and generate a report for application usage seen across all traffic use the Application Usage Real-Time Report.

Menu path: **Reports > Flow Activity > Application Usage**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Time, Category, Application Name, Bar Graph, Percentage, Total Traffic, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Application Usage Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Time	Time of connection
Category	The type of category
Application Name	Name of the application
Bar Graph	Percentage of total bytes represented as a bar graph
Percentage	Number of bytes represented as a percentage
Total Traffic	Total amount of traffic
Count	Number of connections

User Browsing Reports

To search for and generate a report for site destination statistics by user use the User Browsing Statistics Real-Time Report.

Menu path: **Reports > Flow Activity > User Browsing Statistics**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Time , User IP, Destination Name, and Number of times Accessed.

For information on saving the generated report, see Formats for Saving a Generated Report.

User Browsing Statistics Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Time	Time of connection
User IP	IP address of the user making the inquiry
Destination Address	IP address that was targeted
Number of times Accesses	The number of times accessed

Top Users Reports

To search for and generate a report for top traffic users use the Top Users Real-Time Report.

Menu path: **Reports > Flow Activity > Top Users**

In addition to setting the common report options in Preparing a Real-time Report, you can select optional filter operators in the generated report.

You can select to view various options in ascending or descending order. Choose the sort order by using the list. The default is to display the Source Device, Time , User IP, Bar Graph, Percentage, Total Traffic, and Count.

For information on saving the generated report, see Formats for Saving a Generated Report.

Top Users Report - Optional Filter Operators

Option	Description
Source Device	Device that sent these log messages
Time	Time of connection

Option	Description
Category	The type of category
User IP	IP address of the user making the inquiry
Bar Graph	Percentage of total bytes represented as a bar graph
Percentage	Number of bytes represented as a percentage
Total Traffic	Total amount of traffic
Count	Number of connections

All Saved Reports

The All Saved Reports screen displays a list of all saved reports for specific types of data based on search expressions and time intervals you have defined and saved in the past.

All saved searches and types, such as Index Search, RegEx Search, Index Report, and so on, that are stored in the system are visible on this page.

You can click the **Run** icon and regenerate the report with a different time range, or click the **Edit** icon and change the saved report parameters before rerunning the report. All options are available, not just the ones originally selected. You can customize the new report using new filters and wildcards. You can also filter the list of saved reports displayed by title or by typing a key word from the report title in the **Find** field and pressing the Enter key. The key word or words are highlighted in the resulting list. To restore the full list of saved reports, clear the **Find** field and press the Enter key again.

For more information on saving the generated report, see Formats for Saving a Generated Report.

User Preferences

The admin icon on the home page allows you to set values for your **Account Information, System Preferences**, and to **Change Password**.

Viewing Your LogApp Account

Procedure

1. Click the user icon on the home page.
2. Review and accept or change the default settings as explained in the Account Options table.

Account Options

Element	Description
Account Information	
User Login	The login name of the current user.
Email Address	The email address of the current user. This can be reset by the system administrator or user.
System Preferences	
Rows per Page	The number of rows that display in each report page. Can be set from 10 to 1000 rows by user.
Page Refresh Rate	The page refresh rate in seconds. Can be set from 30 to 600 seconds by user.
Session	Session Timeout can be set from 5 to 300 minutes by user. The default is

Element	Description
Timeout	300 minutes (5 hours).
Enable Multiline View	Checking this check box enables display of multiple lines in PDF and HTML reports.
Login Landing Page	The page that appears immediately after logging into the LogLogic LMI appliance. You can change this at any time. For instructions, see Changing Login Landing Page.

3. Click **Save**.

Changing the Login Landing Page

The login landing page appears immediately after logging in to the LogLogic LMI appliance. The LogLogic Overview page is the default landing page. To change the landing page, perform the following steps.

Procedure

1. Click the user icon at the top of the home page.
2. From the **Login landing page** list, select a page from the following options:
   - LogLogic Overview
   - Triggered Alerts
   - Index Search
   - All Saved Reports
   - All Saved Searches
   - Advanced Dashboard
   - Advanced Search
   - Monitoring Console
3. Click **Save**.

Result

The next time you log in to the appliance, the new home page that you selected in this step is displayed.

Changing the Account Password

You can change your password at any time.

Procedure

1. Click the user icon on the home page.
2. Click the **Change Password** button.
3. The Change Password dialog box appears. It displays date of last password update.
4. In the **Current Password** field, enter your current password.
5. In the **New Password** field, enter your new password as per the password requirements specified on the window.

   By default, the password must be at least six characters, containing at least one non-alphabetical character, and cannot be the same as the user ID. Spaces at the beginning and end of the password are removed before being stored in the system. The password requirements might differ if password control has been enabled in LogLogic LMI.
6. In the **Confirm New Password** field, enter your new password again for verification.

Advanced Features

This section describes advanced features in LogLogic LMI.

By default, these features are disabled. A user with administrator privileges can enable the Advanced Features, or set up a High Availability pair for Advanced Features, or both. For more information on enabling the features or setting up High Availability, see *TIBCO LogLogic® Log Management Intelligence Administration*.

- Advanced Search
- Advanced Dashboards
- Bloks: Filter Bloks, Correlation Bloks, and Time Bloks
- Data Models
- Enrichment Lists
- Exporting and Importing Configurations

  Only an administrator with CLI access using root can use this feature. For more information on exporting and importing configurations, see the "Import or Export Entities Between Appliances" section in *TIBCO LogLogic® Log Management Intelligence Administration*.
- Monitoring Console
- Monthly index
- Queries
  - Search queries
  - Scheduled queries
  - Tail queries
  - Distributed Advanced Search queries
- REST API support for Advanced Search
- Rule Management:
  - Triggers

- Aggregation rules
- Distributed aggregation rules

Important Considerations

- You cannot access advanced features on a standby node in a high availability setup. However, you can access them from the public IP address or the IP address of the active node.

  Before configuring HA, you must disable Advanced Features on both active and standby appliances.
- Use caution when enabling advanced features on LX4025, ST4025, MX4025, ST2025-SAN, or LX1025R1 models, because the memory requirements of these features when in use might cause performance issues. Also, continuous use of Advanced Features on these models can cause the appliance to run out of memory and lead to engine restarts or failure.
- On an appliance where the memory is less than the minimum required 32 GB, you cannot enable Advanced Features from the GUI. You must enable it from the CLI, by running the system logu enable command.
- For information about the behavior of Advanced Aggregation and Monitoring Console features after upgrading to LogLogic LMI 6.4.0, see the Upgrade Considerations section in *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*.

Advanced Search

Using Advanced Search, you can easily interact with your data. You can run simple and complex searches, save search elements and time ranges in the form of Bloks, and retrieve results to analyze failures or other anomalies.

The simple search mechanism retrieves all events that match the search terms. Advanced searches retrieve results by a "pipeline" concept, where expressions are separated by pipes ("|"). LogLogic LMI uses Event Query Language (EQL), which is an intuitive and efficient search query language that enables you to search large data and view results in seconds. A Structured Query Language (SQL) dialect is also supported.

For more information about how to form a search query and sample queries, see Search Syntax Reference.

You must specify the time in the **Search** or **Time** field. The **Search** and **Time** fields can be combined (by using AND) or used individually as described:

- If you define the time in either the **Search** or **Time** field, the results are retrieved for the specified time period.
- If you define the time in both, the **Search** field and the **Time** field, the results are retrieved for the time period that is common to both the fields.

Note: All dates and time values are defined in the local time zone of the installed system. They are not based on the browser's time zone.

By navigating to **Management > Advanced Features > Queries** and then to the Search page, you can view search queries that are currently running or are completed. From this page, you can select and delete any query, if required. Deleting the query from this page closes the search tab for that query from the Advanced Search page.

Note: Queries used in Advanced Dashboards are not listed on this page.

For complex queries, you can create different types of Bloks that can be reused in future searches. Bloks are query fragments that can be easily referenced from queries. For detailed steps about how to build and use Bloks, see Bloks.

For sample search examples, see Search Examples.

On the Advanced Search page, you can click  to open multiple search tabs. You can run multiple searches using different search elements on the same data to analyze any anomalies.

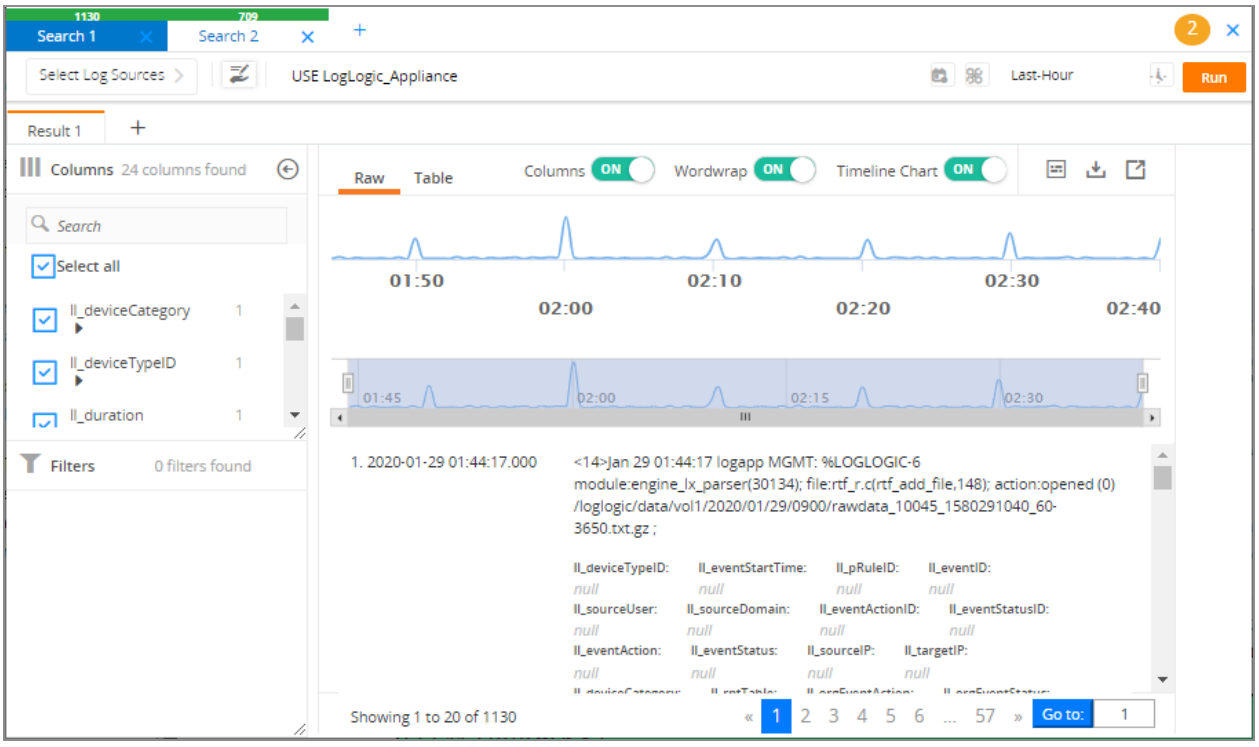

The Search Field

You can enter queries in any of the supported languages (SQL or EQL), retrieving data from data models, with filters of any kind such as LIKE, regular expressions, comparison operators, math, functions, and so on. You can use single or multiple terms.

- To start an EQL statement, enter USE.
- To start an SQL statement, enter SELECT.

You can also search data based on Bloks. For details on how to add a new Blok or use the existing Bloks, see Bloks.

For example, enter the following query in the **Search** field to retrieve events from the system data model within the last hour:

```
use system | sys_eventTime in -1h:NOW
```

The `system` data model refers to all the data in the system.

Note: Copying a query from a Rich Text Format (RTF) application (such as Microsoft Word) to LogLogic LMI might interfere with query processing. For example, when you copy the query, extraneous characters might be added, or straight quotation marks (") might be replaced with curly quotation marks (”), which are not part of a correct query string. Therefore, when copying from a rich-text source, review the search query syntax and correct any errors before proceeding.

When used in an Advanced Search query or an expression, the name must be enclosed in square brackets ([ ]) in any of the following scenarios:

Name	Examples
starts with a number	[123] [12model] [1col_0]
contains a hyphen	[data-a] [-col1] [aBc-]
contains all numbers	[1234]
contains a space	[abc model]

To view system notifications, if any, click the 2 icon in the upper-right corner of the page.

To close all search tabs at once, click the Close icon (**X**).

To view search results, click the Run button Run.

Additional assistance

You can use the following components to help you form a query quickly:

Log source picker

Instead of using data model names in the Advanced Search query, you can search by specific log sources. Click **Select Log Sources**, and then select multiple log sources from the log source picker. A query is automatically generated and displayed in the search field. For more information, see Log Source Picker.

Content Assist

You can use suggestions from a content assistant to help you create the search query. As you start typing, the Content Assist feature shows contextual matches and completions for each keyword into the **Search** field. For more information, see Content Assist.

Content Assist

The Content Assist feature shows type-ahead or contextual matches and completions for each keyword as you type in the **Search** field.

These contextual matches are retrieved from your data. You can get assistance for language syntax, column names, data model names, recent search history, and Blok names. You can disable Content Assist by clicking the **Disable Content Assist** button.

As you start typing in the **Search** field, the Content Assist panel is displayed:

- **Suggestions** help you build your search query by suggesting the next matching term.
- **Models** let you choose data models to be used in your query. Data models are the equivalent of tables in SQL.
- **History** displays all recent search entries that you can choose from to run a query.

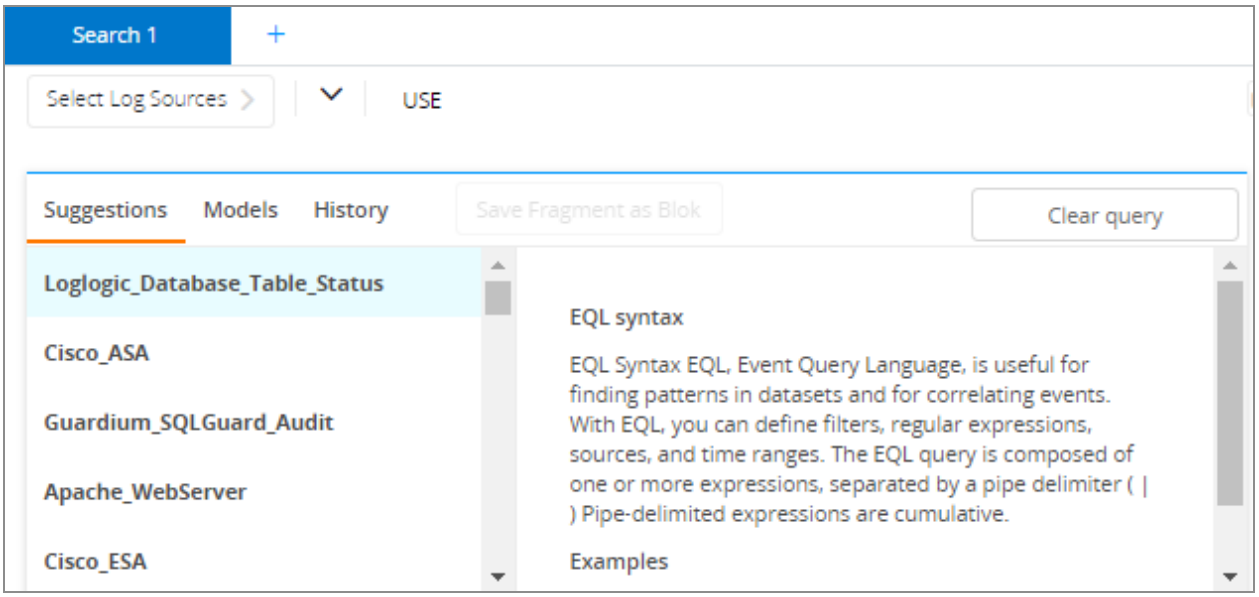

Click the term to select and add it in the **Search** field.

Log Source Picker

Instead of using data model names in the Advanced Search query, you can select log sources from the log source picker, and a query including the selected sources is generated for you. You can also use filters to create a dynamic rule to generate the search query.

The generated query includes the `system` data model and a filter for the selected log sources. For example:

```
USE system | sys_device IN ('::ffff:198.51.100.2_
otherUnix','::ffff:198.51.100.3_otherUnix')
```

```
USE system | DeviceInGroup("All Other Unix") OR sys_device = '::1_
logapp'
```

```
USE system | sys_concentratorId IN ('198.51.100.2', '198.51.100.30')
```

Limitations

The log source picker in Advanced Search has the following limitations:

- Even if you have access permissions to Remote Appliances in a Management Station setup, you cannot specify a Remote Appliance in the log source picker and then select the log sources created on that appliance.
- In the generated query, only the system columns are displayed in the results and in the Columns panel, and you can filter the results by the system columns. To make other data model columns available for filtering or parsing, you must replace the `system` data model with the appropriate one in the search query. For example:

```
USE Other_UNIX | sys_device IN ('::ffff:127.0.0.1_
otherUnix','::ffff:10.128.132.92_otherUnix')
```

For steps about using the log source picker and creating a dynamic rule, see Selecting Log Sources.

Selecting Log Sources

If you use the log source picker to select log sources, an Advanced Search query including the selected log sources is automatically generated for you. Then you can specify the time range in the Time field and run the query.

Procedure

1. Click the **Select Log Sources** button to open the log source picker.
2. From the List of Log Sources pane, select the check box of the required log sources.

   You can filter the list of log sources by the following parameters. After selecting a filter, type the value or select from the drop-down list to filter further:

Filter field	Additional information field
Device Name	Enter the device name or select from the list
Group Name	Enter the device group name or select from the list
Type	Enter a source type (a specific device type) or select from the list. In a Management Station setup, you can select Remote Appliance as the type

Filter field	Additional information field
Collector Domain	Enter the name of the collector domain.
Description	Enter a description of the log source.
IP Address	Enter the specific IP address of the log source.

You can apply multiple filters by clicking the + button. Click **Reset filters** to clear all filters.

You can also use the **Search** field to filter the list of devices and then select from the filtered list.

3. To add the selected devices and filters to the Selected Log Sources pane, select the log sources and then click **Add Selected Log Sources**.
4. (Optional) To add a large number of devices, you can create a dynamic rule that contains all listed devices. You can create multiple rules, if required.
    a. Use a filter to retrieve the list of devices.
    b. Click **Add Filters as a Rule**.
    c. Enter a name for the dynamic rule in the dialog box and click **Create**.

    A dynamic rule containing the listed devices is created and displayed in the **Selected log sources** pane.

    Note: The dynamic rules created in a Search tab can be used only in the same Search tab and until the Search tab is active.

5. Review the list of log sources in the **Selected log sources** pane. Repeat the steps, if required, to add log sources by filters, or by selecting their check boxes, or by creating more dynamic rules.

    Note: You can remove selected sources or dynamic rules by clicking **Delete source**.

6. Click **Set**.

Result

A search query that includes the selected log sources is automatically generated and displayed in the search field.

If you selected remote appliances as log sources, then:

- Selecting one appliance adds `sys_concentratorId = 'IP'` in the search query; where *IP* is its IP address.
- Selecting multiple appliances adds `sys_concentratorId IN ('IP1', 'IP2')` in the search query; where *IP1* and *IP2* are their IP addresses.
- Using the Type filter to add all appliances adds `sys_concentratorId IN ('ALL')` in the search query; where `ALL` indicates that all remote appliances and the management station are included.

What to do next

Specify the time range in the Time field and then click the Run button 

to run the query.

The Time Field

In the Time field, you can enter absolute and relative time ranges.

You can also search based on Bloks.

From the **Search** tab, enter the time period in the **Time** field and click the Run button 

. For details on how to add a new time Blok or use the existing Blok, see Time Bloks.

> **Note:** All dates and times are defined in the local time zone of the appliance rather than being based on the client system time zone.

For example, enter -5h to retrieve all events that occur in the last 5 hours.

Note: The **Time** field must be empty when using infrastructure search queries.

An example of invalid infrastructure search query is:

```
use LogLogic_Config_Bloks | sys_eventTime in -5d
```

Search Results

After running a search query, you can view search results in the **Result** tab.

You can visualize results using **Timeline Charts** or **Data** panel. After running a query, if you retrieve lots of results, you can group the results without having to issue a new query, and then drill-down into the information. You can see both aggregated counts as well as create visualization elements to better isolate trends and issues. You can include multiple filters to narrow your results. Create a filter in the context of an event, and view results based on a specific filter.

After clicking **Run**, a progress bar is displayed above the search tab name showing the progress of the query. Based on your data, it might take a few minutes to retrieve results into all panels. By default, results are returned in ascending order. After the query is run, the number of results is displayed above the tab name. Twenty results are displayed per page. You can jump to other pages using the pagination controls at the bottom of the panel.

If you use a GROUP_BY clause in the query, you can save the query as an aggregation rule by clicking the icon  .

Note: By default, a maximum of 100,000 results are displayed in the Result tab. To increase the limit, use the LIMIT clause in your query. For details, see LIMIT Statement.

Querying a large data set using Advanced Search might display an error or an exception if the result contains more than a few million records.

- LogLogic_Monitor_Cpu
- LogLogic_Monitor_Cpu_Load
- LogLogic_Monitor_Diskspace

- LogLogic_Monitor_Memory
- LogLogic_Monitor_Node_Memory

Click the Add icon ⊞ to add multiple result tabs to view the same data in different forms. When results are grouped together, a new **Result** tab is displayed showing the grouped results for the selected value.

> **Tip:** If you are using multiple search tabs, closing a tab that is no longer required frees the memory being used for displaying search results.

The **Result** tab is divided into the following panels:

- Data display data in raw format and normalized tabular format on the **Raw data** and **Table** tabs within the panel.
- Columns provide all available columns and their associated values based on each search query. You can turn the Columns view off by using the switch on the top of the panel.
- Timeline Charts display the distribution of events in time using a line chart at the top of the panel. The Timeline Charts view is on by default. You can turn the view off by using the slider at the top of the panel.
- Filters displays the filters you have used in the results. You can filter the search results based on time range or column values. You can also edit the filter values.

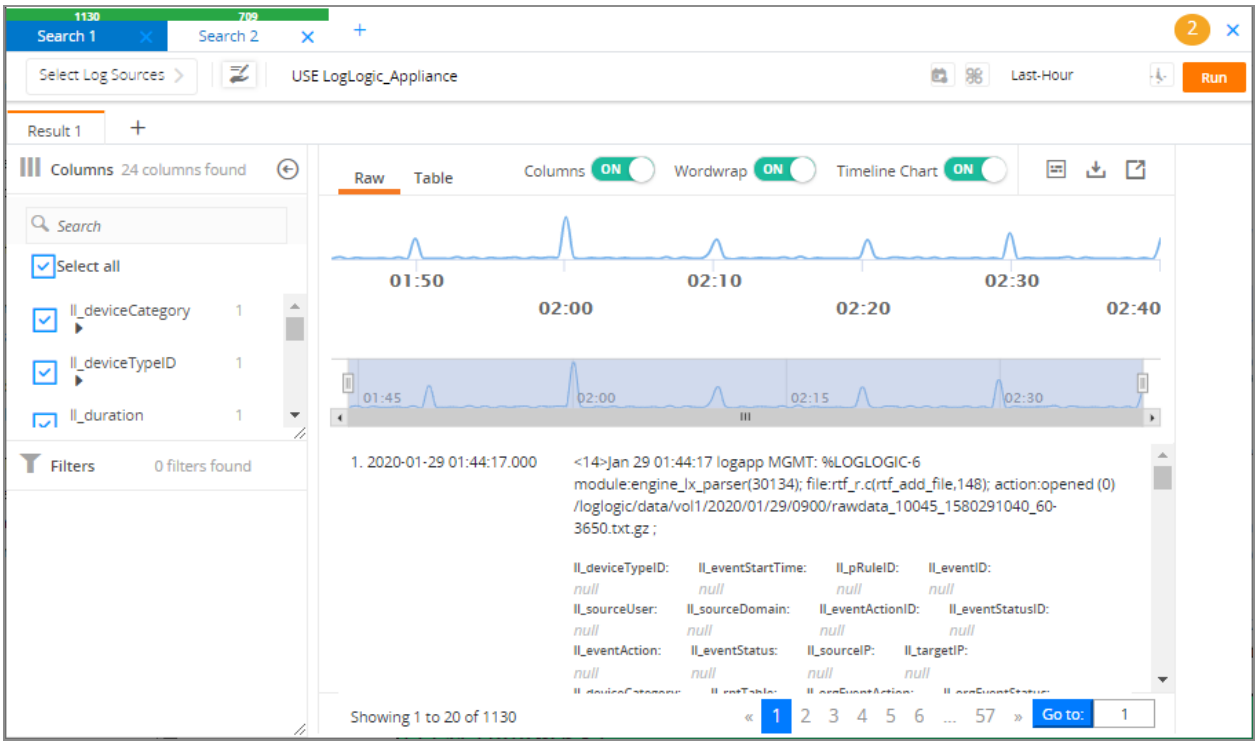

To page through the results, click the next arrow; to return to the previous page click the previous page arrow. You can also return to the first page or go to the last page by clicking on the first page and last page arrows, respectively. You can also jump to a page by entering the page number in the **Go to** field.

Timeline Charts

A chart is a visual representation of your data against time. By using elements such as lines (in a line chart), a chart displays a series of numeric data in a graphical format.

You can add multiple result tabs to view the same data in different formats. The timeline chart displays the event distribution for a specified time period. You can use different options to view chart details, zoom in and out of the chart, and show or hide the chart panel.

Note: Timeline charts are not available in GROUP BY queries. In case of queries that do not require a time filter (such as infrastructure queries), timeline charts are off by default. If turned on, the message `No data to display` is displayed.

From the Timeline Charts panel, you can perform the following tasks:

- **Showing or hiding the Timeline Charts panel**

  Click the slider icon 

 located in the upper-right corner of the results to show or hide the panel.

- **Zooming in or zooming out of timeline charts**

  You can zoom in or zoom out of a particular area of the timeline chart using the time-range picker.

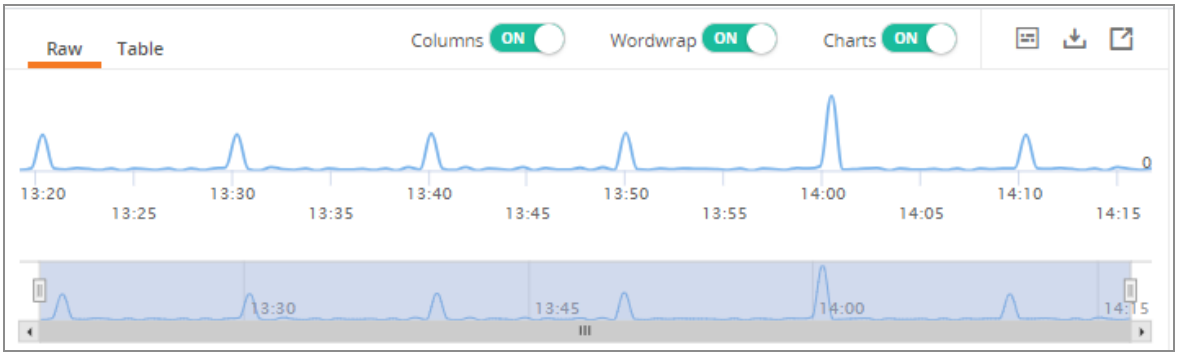

Grab the handles on the X-axis time-range picker, it turns into a slider. Drag the slider across the X-axis to define the time range that you want to zoom in. The timeline chart is updated for the selected time.

The following timeline chart displays the zoomed in data for a specified time range and the **Data** panel shows the filtered results for the corresponding time range.

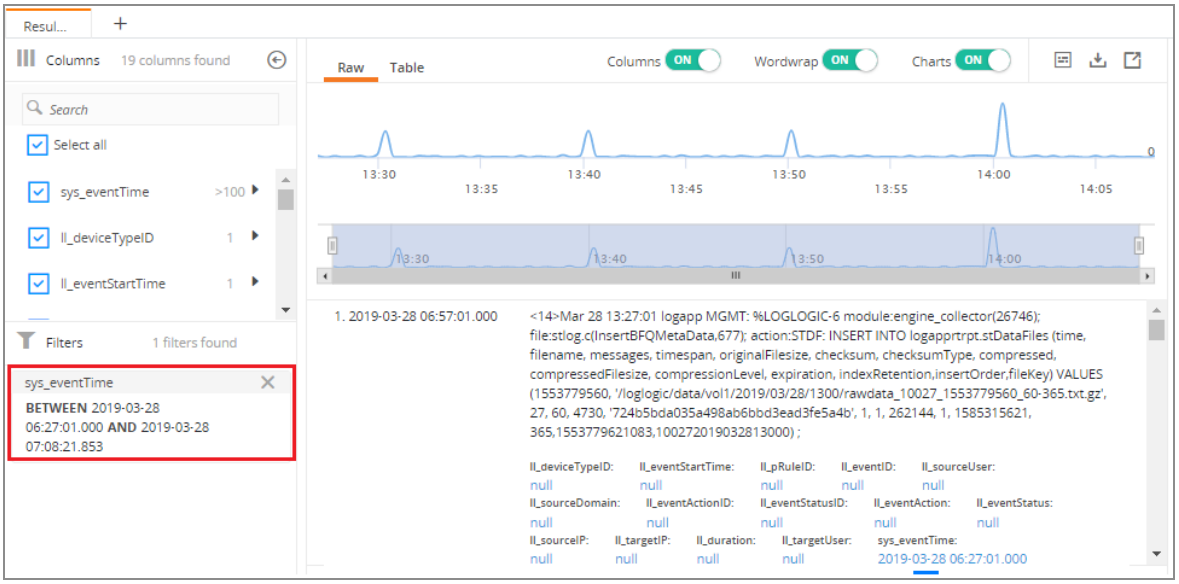

You can expand and collapse the time range by dragging the borders of the selected time range to the desired location. Once you define the time range, position the mouse inside the selected time range and drag the slider to define the new time range. Similarly, you can define a specific time by clicking on the timeline chart. The time range can be adjusted at any time.

> **Note:** As you adjust the time range on the timeline chart, the Columns and the Data panels are adjusted automatically for the selected time range.

- **Viewing timeline chart details**

  Hover your mouse over any part of the timeline chart to view the time and the number of messages.

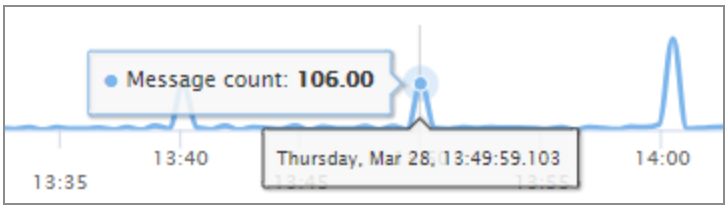

> **Note:** When you use Correlation Bloks in advanced search, hovering your mouse over any part of the chart displays the number of correlation events instead of the message count.

- **Filtering results based on the time range**

  You can fine-tune your search results based on the time range. Click the event count (the line that represents the number of events) on the timeline chart or define the time range by zooming in on the timeline chart to view results in the Data panel. A new filter is added for the defined time and the filtered results are displayed on the Data panel.

Columns

Based on your search query, all available columns are displayed in the Columns panel.

You can group together your results based on any column and the value associated with that column. Similarly, filtering helps you fine-tune your search results when analyzing large data sets.

System columns are columns with event metadata that are present in all data models by default. For a list of system columns, see Types of Columns. Additional columns are available depending on the data models involved in the query.

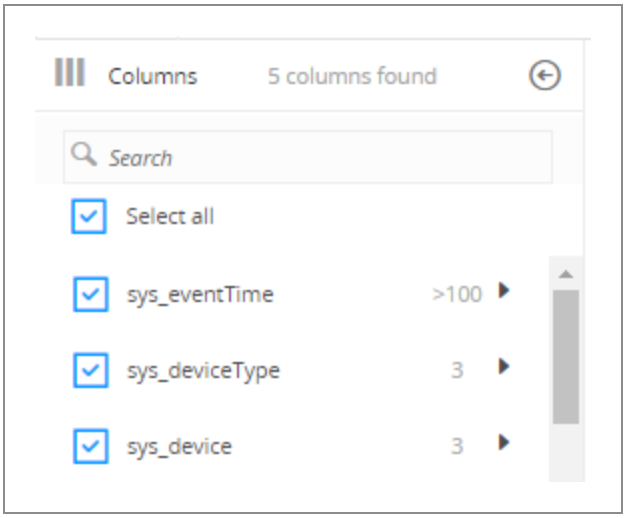

From the **Columns** panel, you can perform the following tasks:

- **Showing or hiding the Columns panel**

  Click the hide icon located on the right corner to hide the **Columns** panel. Click to show the **Columns** panel.

- **Finding columns**

  You can quickly find the desired column by typing the column name in the **Search** field. As you start typing a column name in the **Search** field, all possible columns that start with the letters that are typed get displayed in the pane. The **Columns** panel is refreshed based on the selection.

- **Showing or hiding columns from the Data panel**

  Select the check box to show the column in the **Data** panel. Clear the check box to hide the column from the **Data** panel. Click **Select all** to select or clear all columns.

  The add columns icon ✓ located on the left side of the column name defines that the column is displayed in the **Data** panel. The **Data** panel gets updated immediately based on your selection.

- **Viewing column value details**

  Click the column value and then select **Show values** to view the details of the selected value. The window displays a maximum of 100 distinct values for the

selected column. The *Percent* column is calculated using the maximum 100 distinct values. When the distinct values for a column exceed 100, the *Percent* column is not displayed. If you filter on a particular column value, then the percent value on the top shows the percentage of occurrence of this particular column value in the entire result set.

The following illustration displays values for the column `sys_eventTime`.

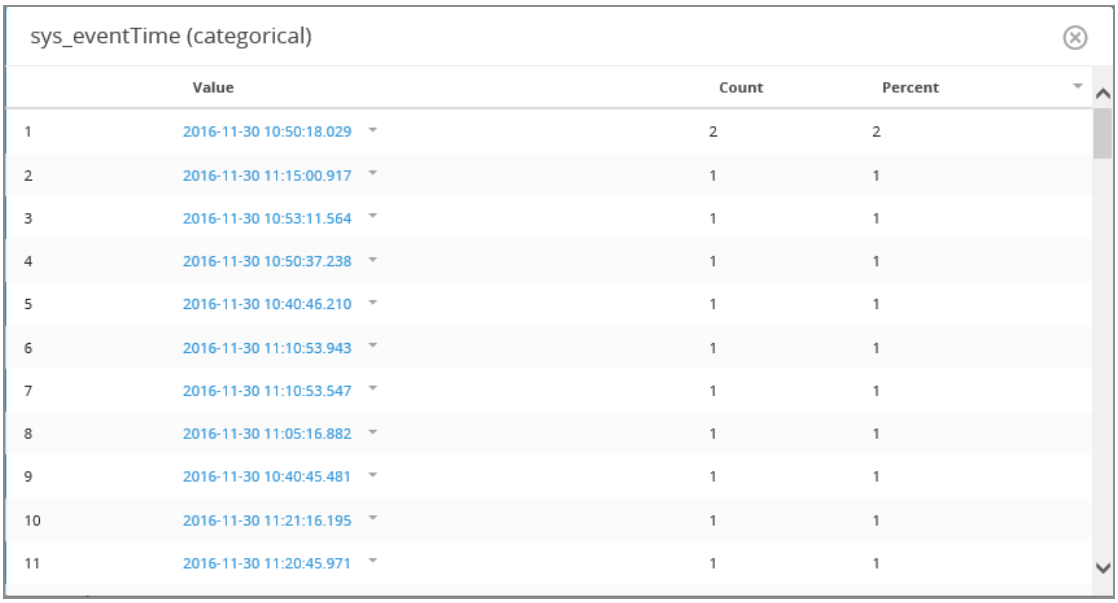

- **Filtering results based on the column value**

  Click the **Value** link and select **Include this filter** to filter results based on that value. If you select the **Exclude this filter** option, the results are displayed without the defined value. You can add multiple filters. Select **Remove this filter** to remove the selected filter from the results. The blue icon represents included values and red icon represents excluded values from filtering data on the **Data** panel.

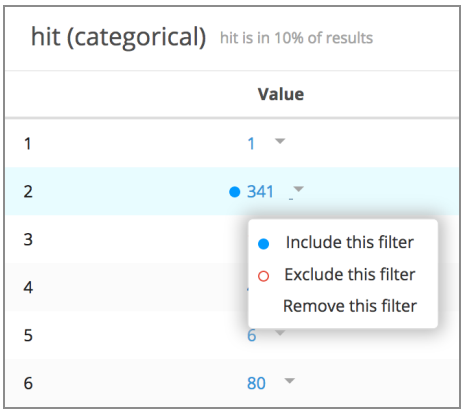

The following illustration displays filtered results based on the value filter Other UNIX included for a column `sys_deviceType`.

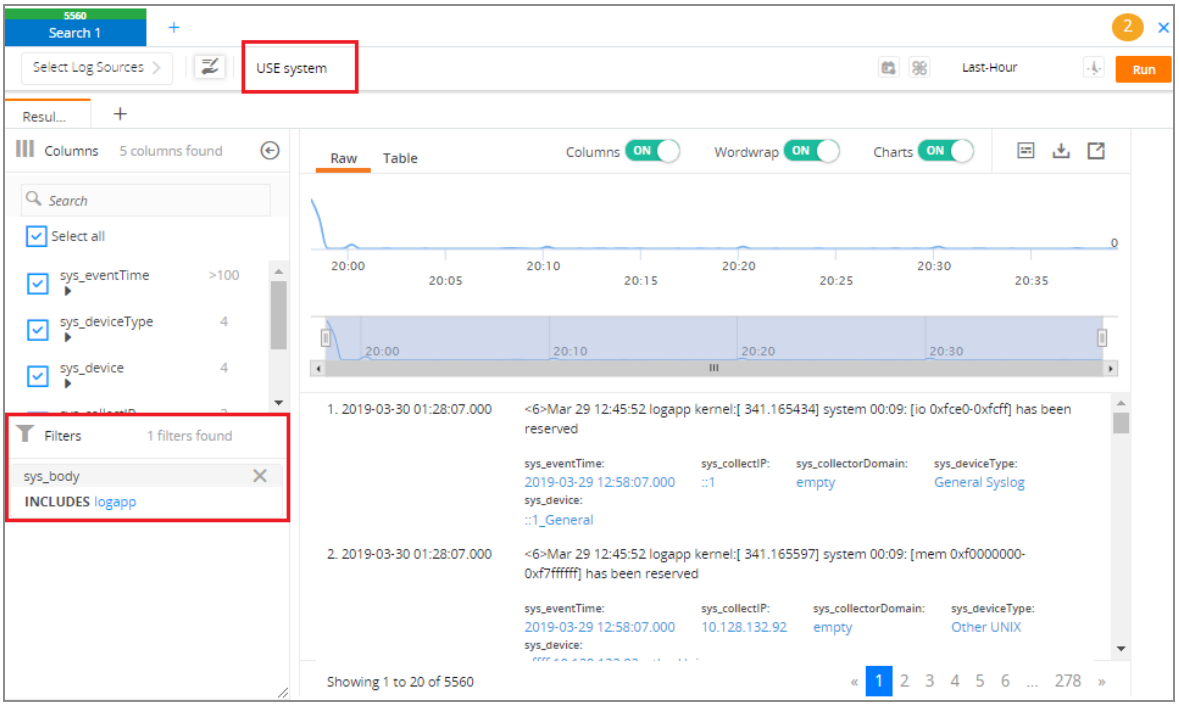

- **Editing value filters to refine results**

  Based on your selection, a new filter is added in the **Data** panel and the refined results are displayed based on the filter. Click inside the value filter box to edit the value. Click the check mark to update the value changes. The **Data** panel results are refreshed based on the updated filters.

Note: When updating the time value, enter it in the `YYYY:MM:DD HH:mm:ss` format.

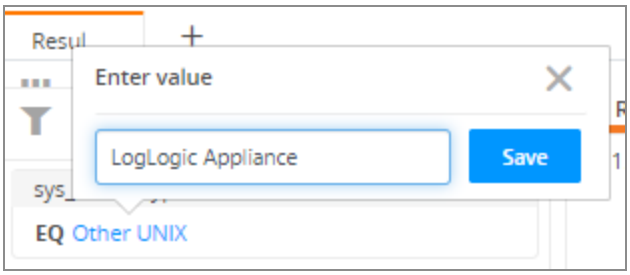

- **Grouping by values**

  Click the column value and then select **Group by** to view grouped results. A new **Result** tab opens showing the results that are grouped by the column. The number of groups is displayed against the column name in the Columns pane. However, for time-based columns, the number of unique values is displayed instead of the number of groups.

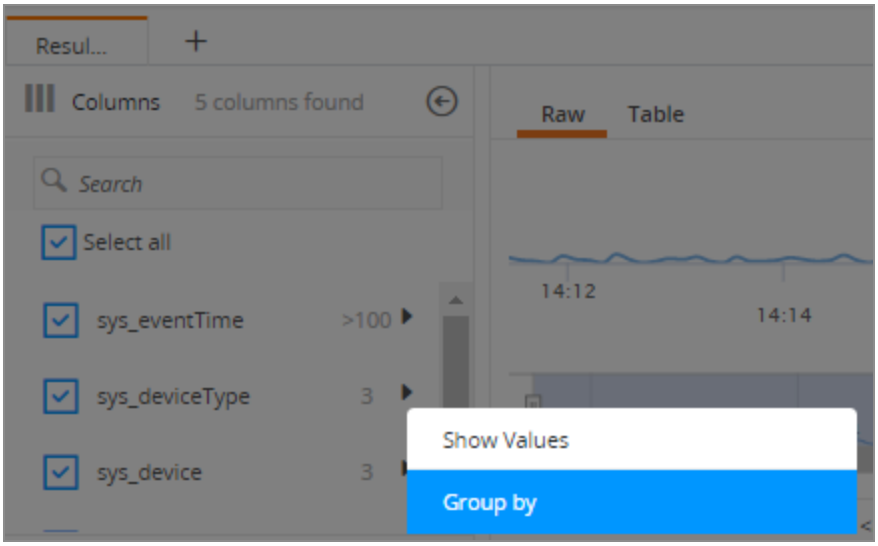

  The following illustration displays the results grouped by the `activity` column.

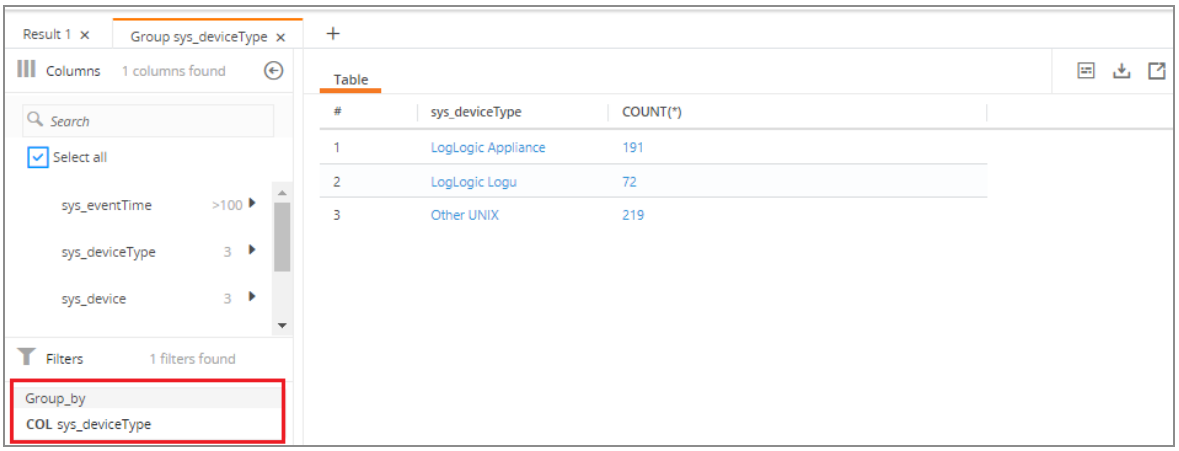

You can group by different time ranges. Click the timestamp value, and select the **Group Dates by** option. From the list, select the option to group your results by different time periods. A new **Result** tab opens showing the results that are grouped by different time units.

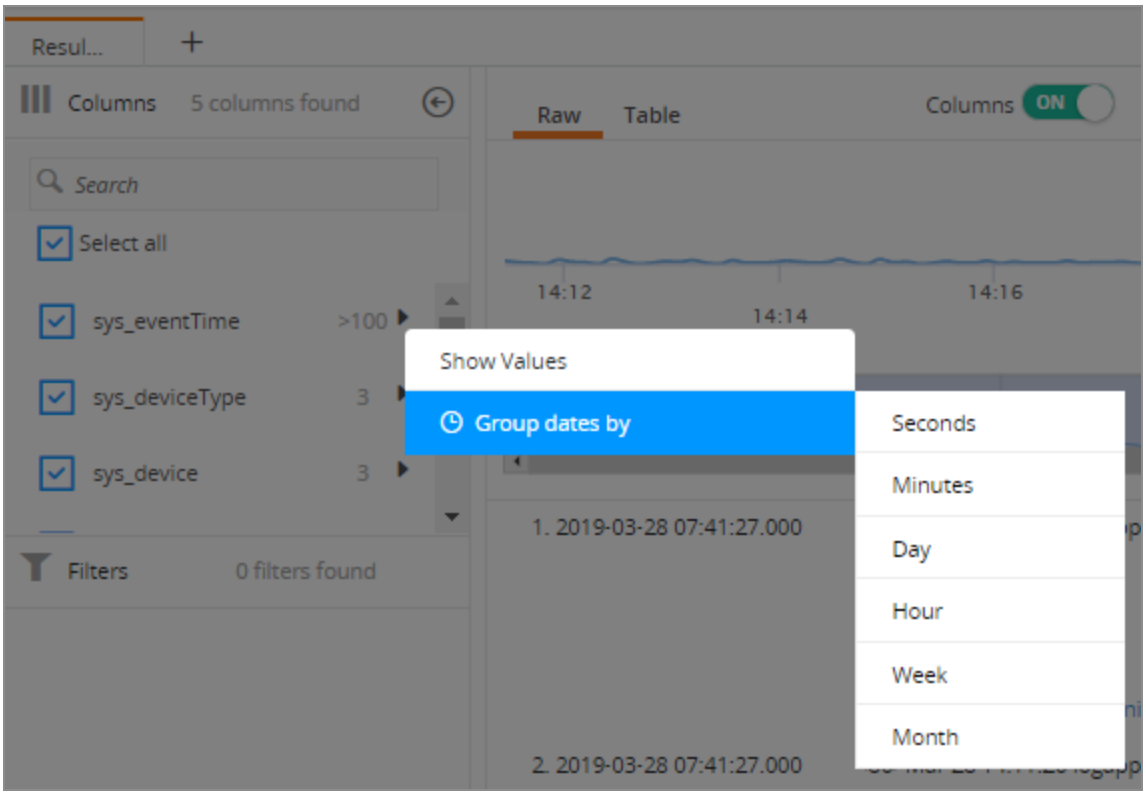

- **Add aggregation**

  You can aggregate columns that have Integers and Long values. Click the column

value and select **Add aggregation**. Define how to group values in the aggregation column. The options are: SUM, MIN, MAX, AVG. A new column is added in the **Data** panel.

> **Note:** The aggregation menu is only available after a GROUP BY operation, and only for numeric columns. However, if a GROUP BY operation is already used in the query, the results cannot be aggregated further.

The following illustration displays a new aggregation column (AVG) added in the Data panel.

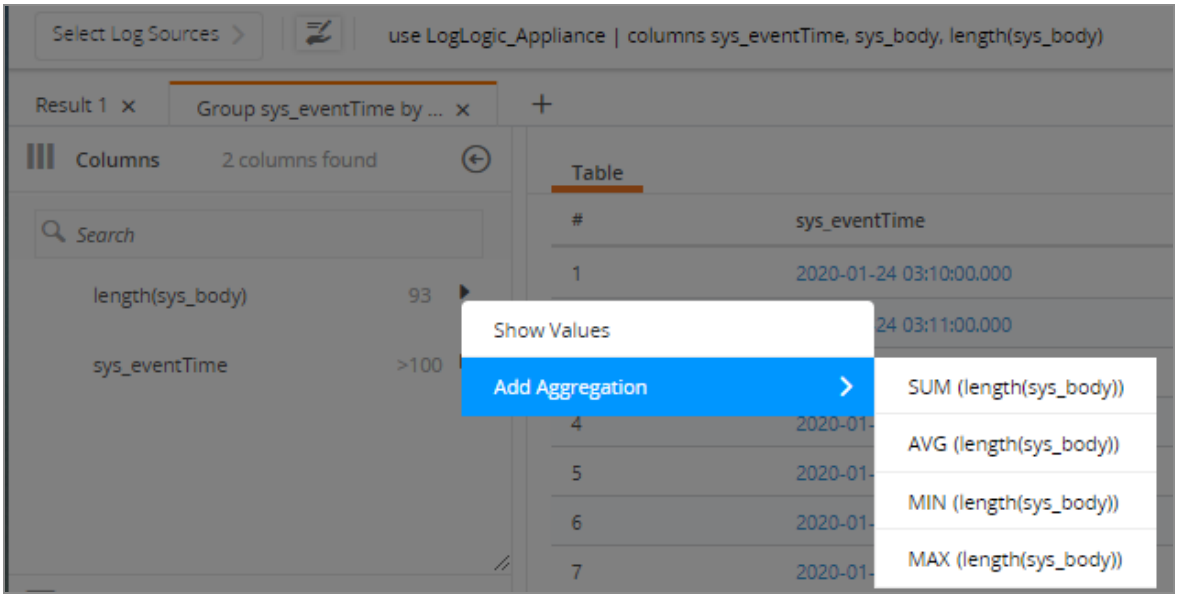

Example

1. Run the following query:

```
use LogLogic_Appliance | columns sys_eventTime, sys_body, length
(sys_body)
```

2. In the left pane, click the column `sys_eventTime` and group by minutes.
3. Click the length column, and choose an aggregation type from the **Add aggregation** sub-menu, for example, **AVG(length(sys_body))**.

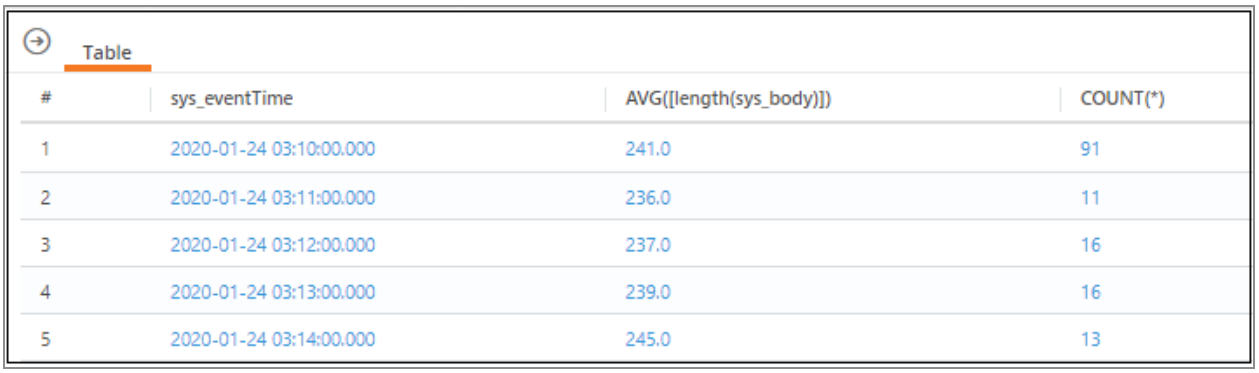

Table

#	sys_eventTime	AVG([length(sys_body)])	COUNT(*)
1	2020-01-24 03:10:00.000	241.0	91
2	2020-01-24 03:11:00.000	236.0	11
3	2020-01-24 03:12:00.000	237.0	16
4	2020-01-24 03:13:00.000	239.0	16
5	2020-01-24 03:14:00.000	245.0	13

Types of Columns

There are two types of columns: system columns and parsed columns.

The system columns are available by default in all data models. System columns contain event metadata such as event body (sys_body), event time (sys_eventTime), or the device name that produced the event (sys_device).

Parsed columns are specific to data models. They are defined in the data model and their values are parsed from the body of the event.

The following list describes all system-generated columns.

Name	Type	Description
sys_body	String	The text of the event
sys_collectIP	InetAddress	The IP from where the event originated. This supports both IPv4 and IPv6.
sys_collectTime	Long	The time when the event was ingested Currently unused.
sys_collectorDomain	String	Name of the collector domain for this event
sys_ collectorDomainId	long	ID of the collector domain for this event
sys_concentratorId	String	IP address in IPv4 format, of the LogLogic LMI

Name	Type	Description
		appliance or group of appliances on which a distributed Advanced Search query is run.
sys_device	String	Name of the device for this event
sys_deviceType	String	Name of the device type for this event
sys_eventKey	String	A unique key that identifies the event in the LogLogic storage
sys_eventTime	Timestamp	The UTC time of the event in Epoch milliseconds. For syslog data, sys_eventTime is the time the event was collected. For file log data, sys_eventTime is the original event time.
sys_filename	String	The file name for event collected from a file
sys_sourceSubType	String	Sub-classification of the source type Currently unused.
sys_sourceType	Integer	ID of the device type

Note: Some system columns are not returned by default in queries that retrieve all columns, such as a SELECT * query, as they are not generally needed in regular queries or they are currently unused. To obtain their values, their name must be explicitly specified in the SELECT or COLUMNS statement. Those columns are:

- sys_collectTime
- sys_collectorDomainId
- sys_concentratorId
- sys_domain
- sys_eventKey
- sys_filename
- sys_sourceSubType
- sys_sourceType

Data

Based on your search query, the retrieved data is displayed in a normalized tabular format, and each event is summarized per row.

You can view data in the raw data format or table format.

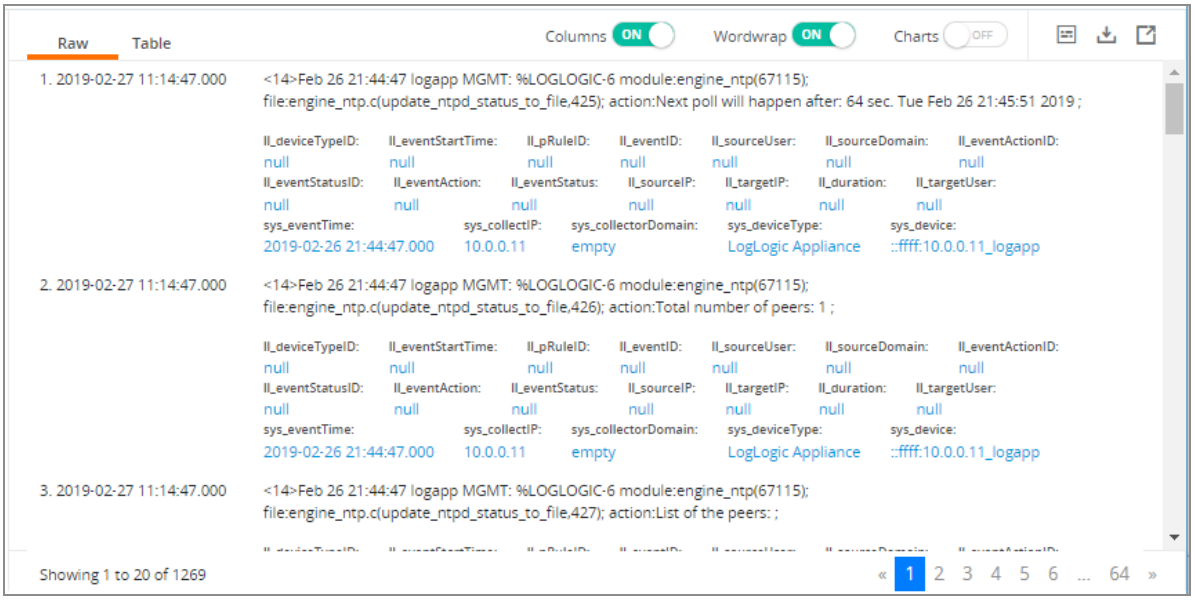

From the **Data** panel, you can perform the following tasks:

- **Viewing event count**

  The total number of retrieved events is displayed on the bottom-left side.

- **Filtering search results**

  You can create a filter using the column value and event body text to fine-tune your search results.

  Click  to show or hide filters from the **Data** panel.

- **Filtering data**

  Click the column value and select **Include this Filter** to filter the data based on the value. If you select **Exclude this Filter**, the results exclude the specified value.

  > **Note:** To filter by any text in the body of the log events (`sys_body` column), turn on the **Messages** view, select the required text, and right-click the selected text.

  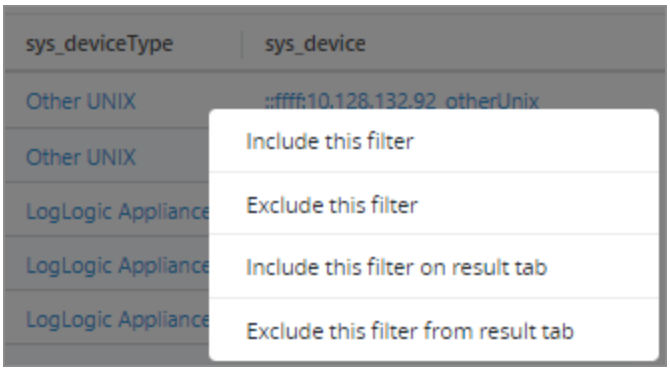

  The Data panel displays results immediately based on the defined filters. You can add multiple filters to fine-tune your search results. You can update the existing filter value. Click the value to open the **Enter value** field. Update the value in the field and click  . The results are refreshed immediately based on the new filter.

  The following image displays the raw data showing filtered results for `sys_body CONTAINS 'logapp'`:

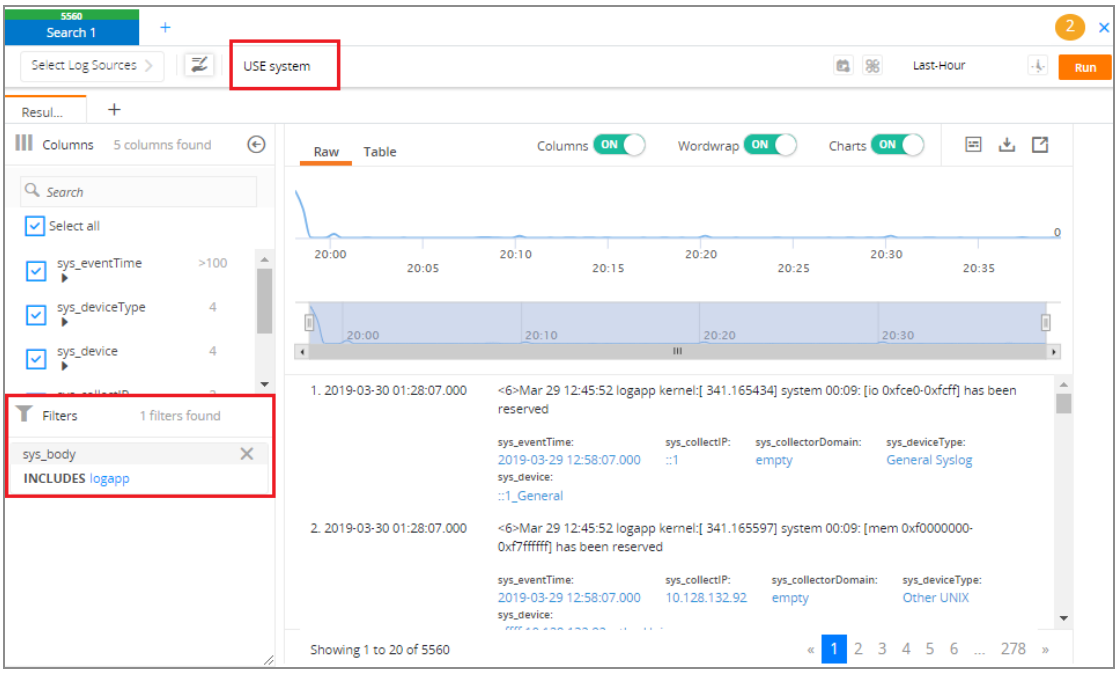

The following image displays the table showing filtered results for `threadid='8158'`.

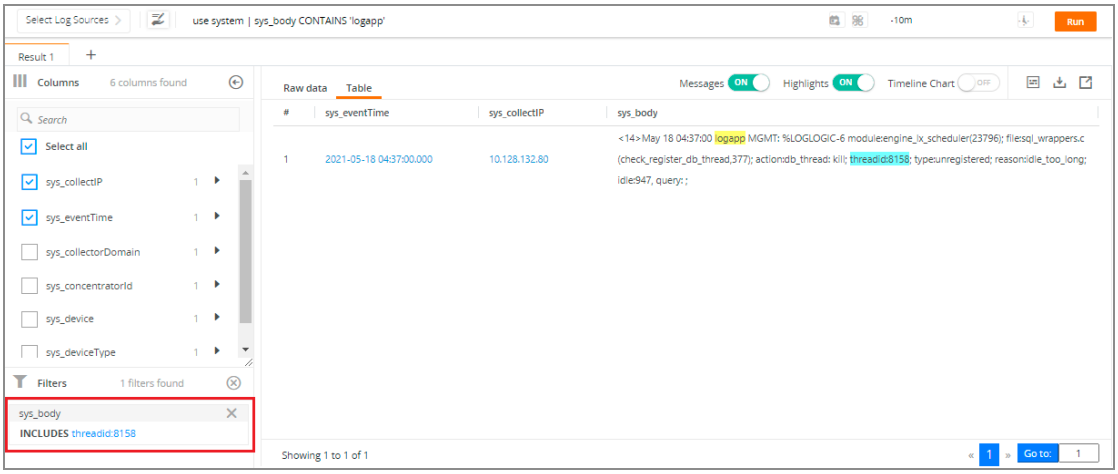

- Click  to show or hide filters from the **Data** panel.
- Click the column value and select **Include this filter** on the **Result** tab to filter the data based on the value in a new **Result** tab. If you select **Exclude this filter from Result** tab, a new **Result** tab displays results excluding the specified value.
- You can filter based on the event body. In the Table mode, ensure that the

Messages option is set to **on**. Then drag the mouse to select the event body and select **Include this filter** to filter your results based on the event body filter. The selected keyword is highlighted in the results. If you select **Exclude this filter**, then the results exclude the specified event body.

The following image shows results based on the event body `threadid='8158'`.

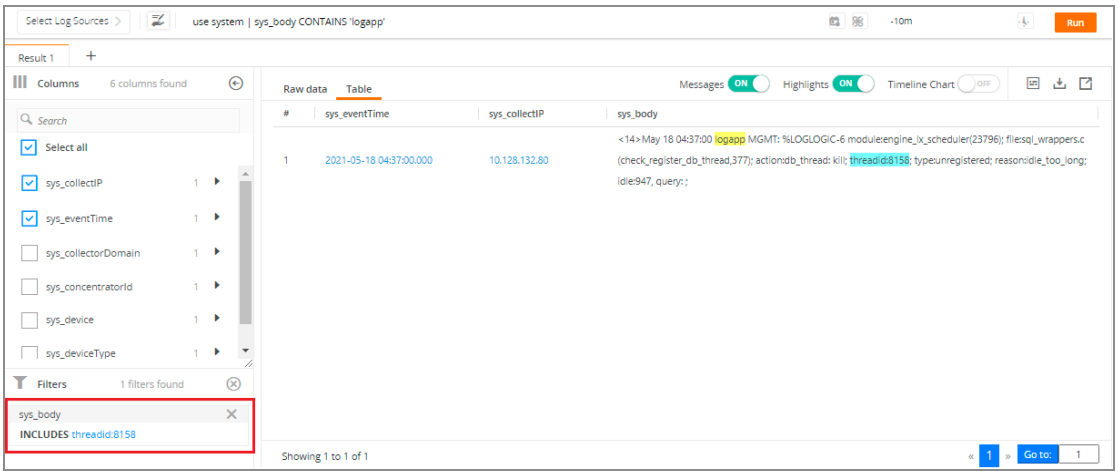

- **Highlighting keywords**

  By default, the Highlights option is set to **on** for queries that include CONTAINS or LIKE statements. Click the Highlights on or off link to highlight keywords or remove highlighting from the keywords. This option is not visible for queries that do not include CONTAINS or LIKE statements.

  Highlighting in the Raw data mode

  In the following image, the keyword `logapp` is highlighted when the search query is:

```
USE system | (sys_body CONTAINS 'logapp ')
```

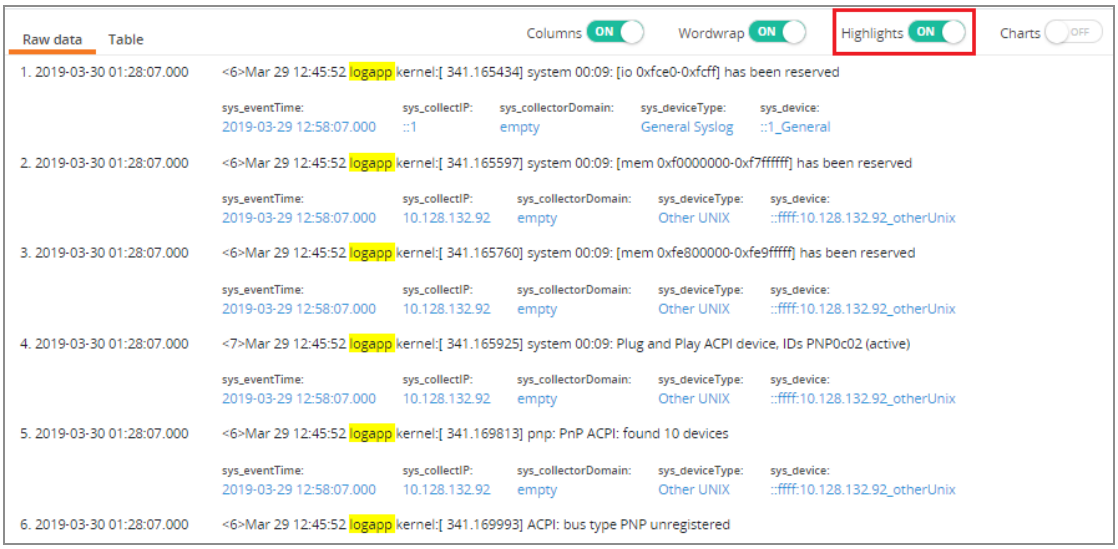

Highlighting in the Table mode

In the following image, the keyword `logapp` is highlighted when the search query is:

```
use system | sys_body CONTAINS 'logapp'
```

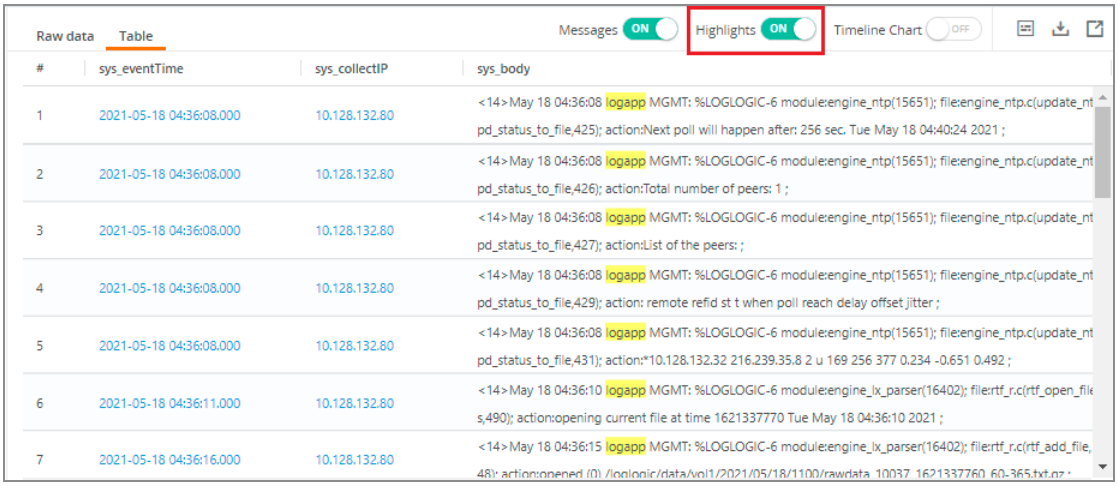

- **Sorting columns**

  You can sort on any column, including group-by count(*) column, group-by aggregation-columns, and other columns. Click the column value and then select **Sort Ascending** to sort columns in order. Click the column value and then select **Sort Descending** to sort columns in descending order.

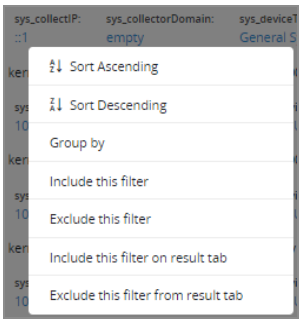

- **Grouping by values**

  Click the column value and select **Group by** to view grouped results. A new **Result** tab opens showing grouped results for the selected value. The number of groups is displayed against the column name in the **Columns** pane. However, for time-based columns, the number of unique values is displayed instead of the number of groups.

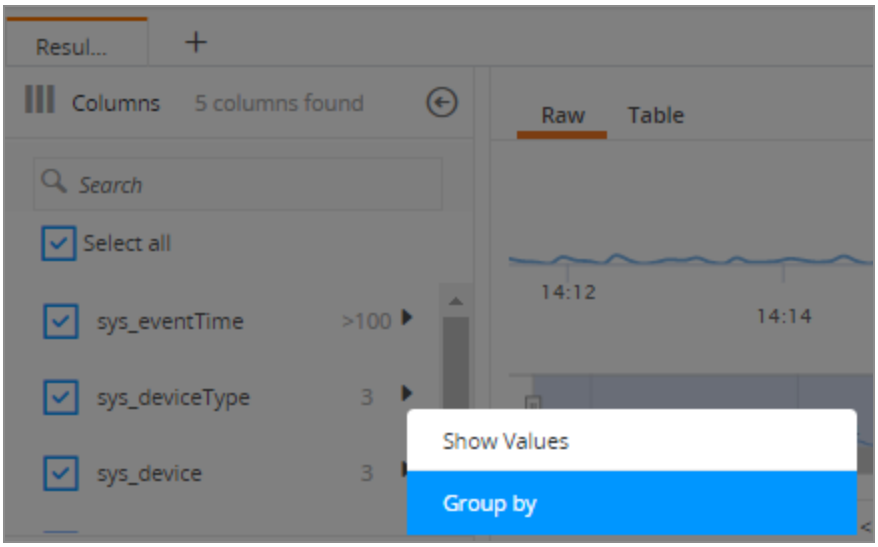

  To group by different time range options, click the time value, select **Group Dates by** option, and then select the period to group your results by different time periods. The **Table** panel is refreshed showing the results that are grouped by the defined time period. When grouped by `sys_eventTime`, the results are sorted in ascending order.

- **Hiding columns**

  Click the column header and then select **Hide** to hide the selected column from the **Data** panel.

- **Adding a new data model**

You can add a new data model from the **Data** panel. Click located on the upper-right corner of the **Data** panel to add a new data model. All events that are displayed in the **Results** tab are copied in the **Create Data Model** panel. For instructions on how to add a new data model, see Creating a Data Model in Graphical Mode.

> **Note:** If a search query contains a single data model, then the defined source filter is copied. If there are multiple data models defined in the query, the **Create source filter** panel does not display any value.

You can edit custom data models from the **Data** panel. Click located on the upper-right corner of the **Data** panel to edit the data model. All events that are displayed in the **Results** tab are copied in the **Create Data Model** panel. For instructions on how to update data models, see Editing Data Models.

> **Note:** The edit icon is only visible when search results are retrieved using custom data models. You cannot edit the *system* data models and the LogLogic LMI built-in data models.

- **Creating filtered query as a new search query**

  After adding filters on your results, click the icon, located on the upper-right corner of the **Data** panel, to create a new search query in a new **Search** tab for the same conditions.

  In the following image, a filter condition `sys_body INCLUDES logapp` is added on the **Data** panel in the **Search 1** tab.

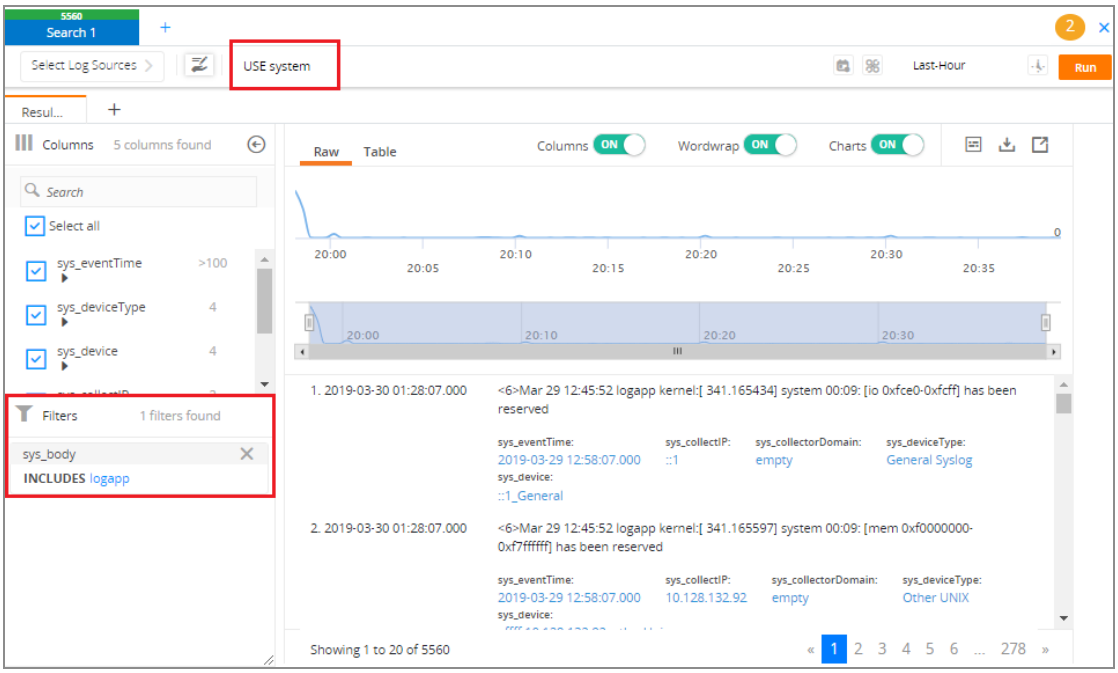

Now if you click the icon in the **Search 1** tab, a new tab Search 2 opens, showing the conditions in the **Search** field.

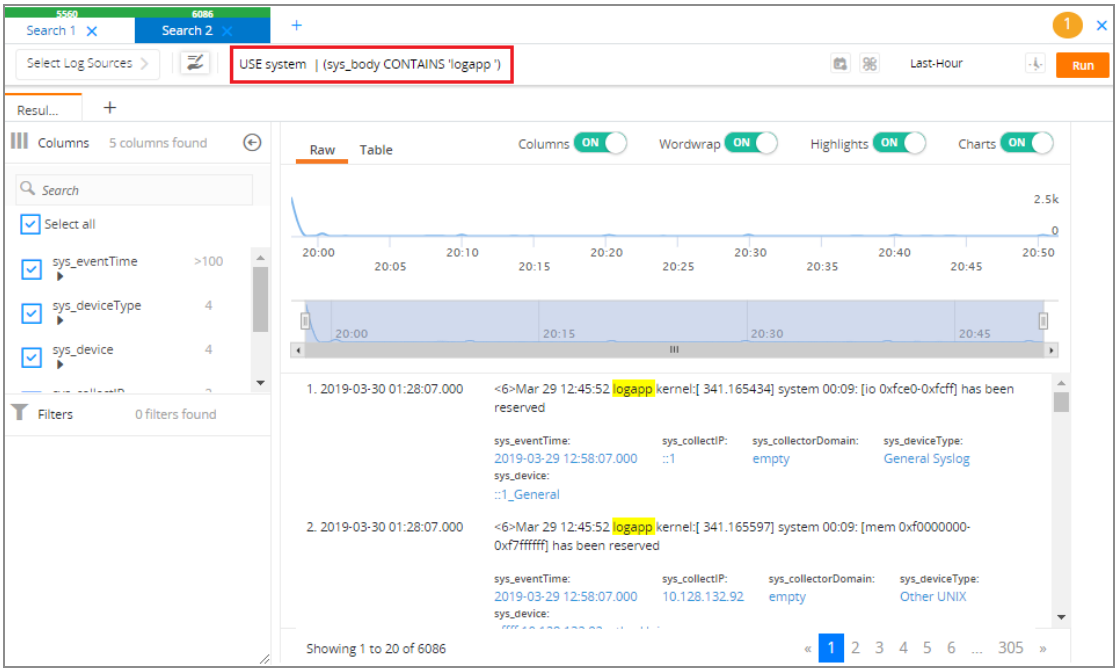

- **Downloading search results**

  You can download the search results as a CSV or PDF file. In the upper-right corner of

the **Data** panel, click the download icon and then specify your option in the Download Results dialog box.

You can also select how much data you want to download:

- **All data:** All pages, all columns
- **All data - Selected columns:** All pages, selected columns, applied filters
- **Current view:** Current page, selected columns, applied filters

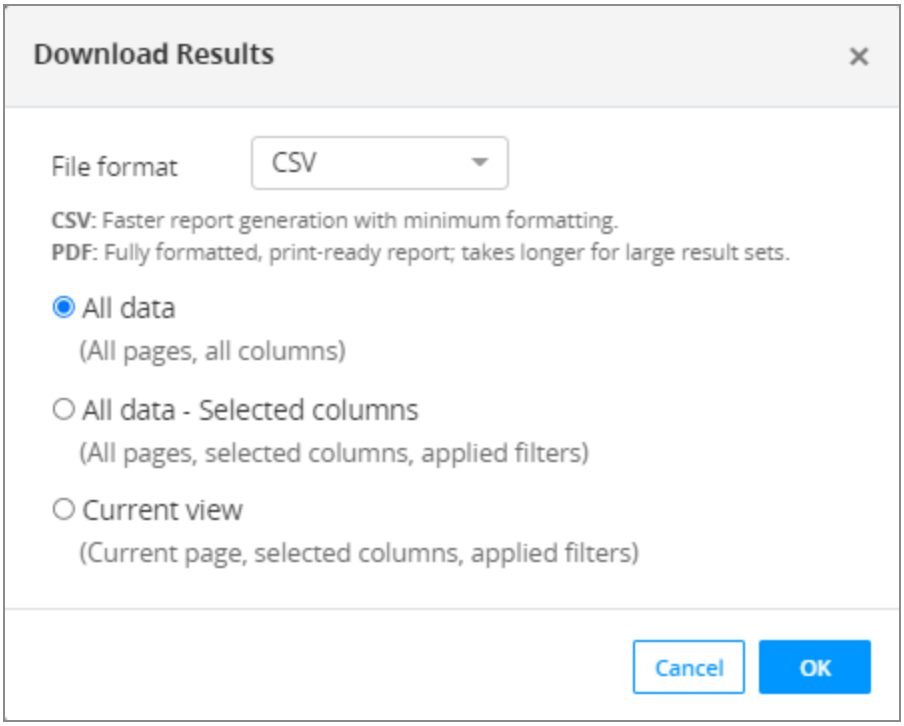

If you select CSV as the file format, the report is generated and downloaded to your computer. Whereas if you select PDF as the file format, the report generation starts in the background. After the PDF is ready, it is saved on the LogLogic LMI appliance and a notification appears in the notification area of the GUI.

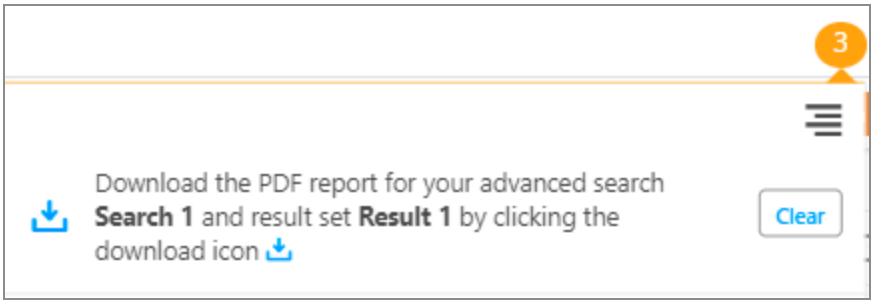

After clicking the download icon in the notification area, the PDF file is saved to your computer.

For PDF download:

- You can initiate a PDF report for different result sets simultaneously.
- While a PDF report is being generated for one result set, you cannot initiate another PDF for the same result set. You must wait till the current report generation is complete.
- If you clear the notification from the notifications panel or close the search tab, then the PDF report is deleted from the appliance and is no longer available for download.
- If you close the search tab for which a PDF report is being generated, the PDF generation is terminated.

For more information about data formats, see the following topics.

- Raw Data Format
- Table Format

Raw Data Format

Based on your search query, the results are displayed in **Raw data** format. Each event is summarized per row.

The column value options are displayed in the following illustration:

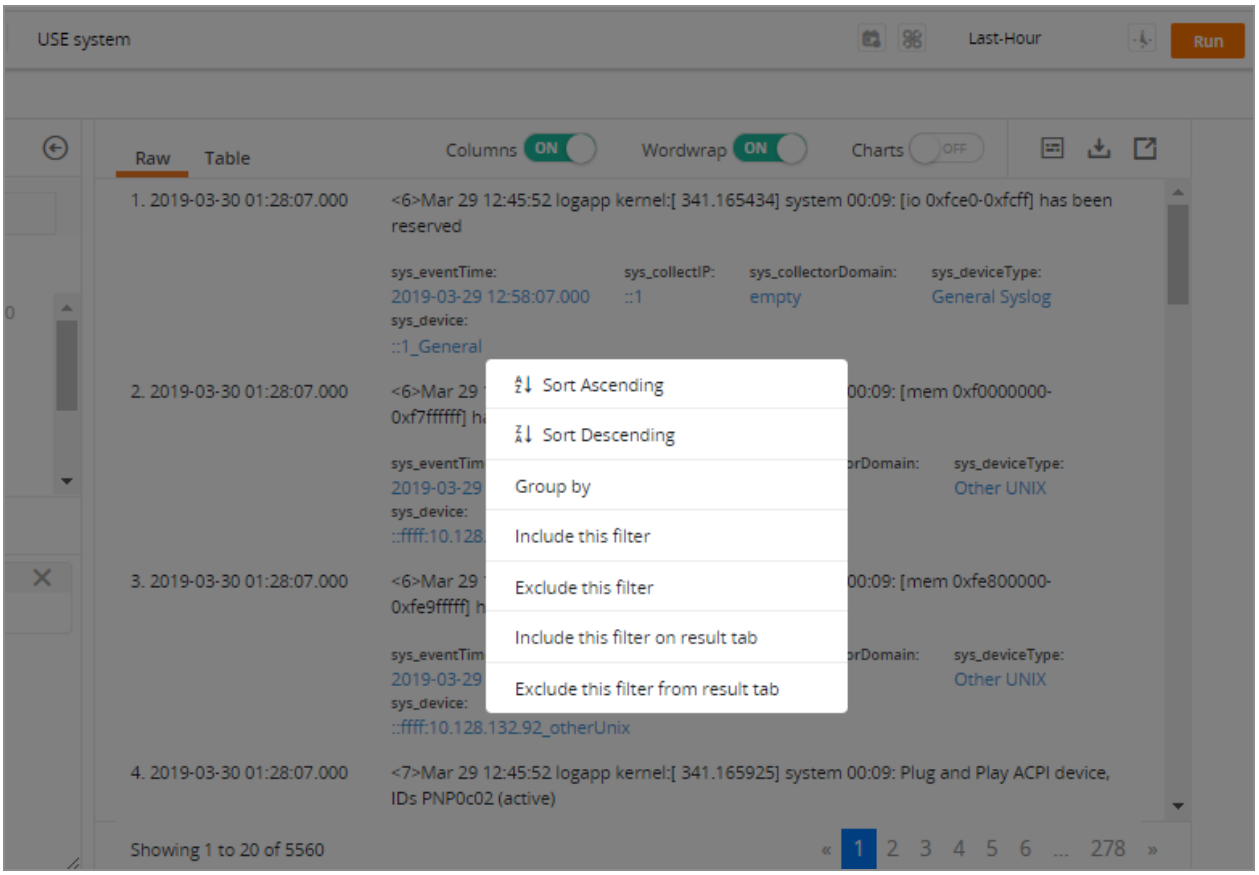

The same result set can be viewed in the Table Format.

Using the raw data format, you can perform the following tasks:

- **Showing or hiding columns from the Raw data**

  Click the Columns button to **on** or **off** to show the selected columns below the event, or to hide columns to view events in the raw format.

- **Wrapping long events**

  Click the Wordwrap button to **on** or **off** to indicate if long event should break at normal word break points or to display long events.

Table Format

Based on your search query, the results are displayed in normalized table format. Each event is summarized per row.

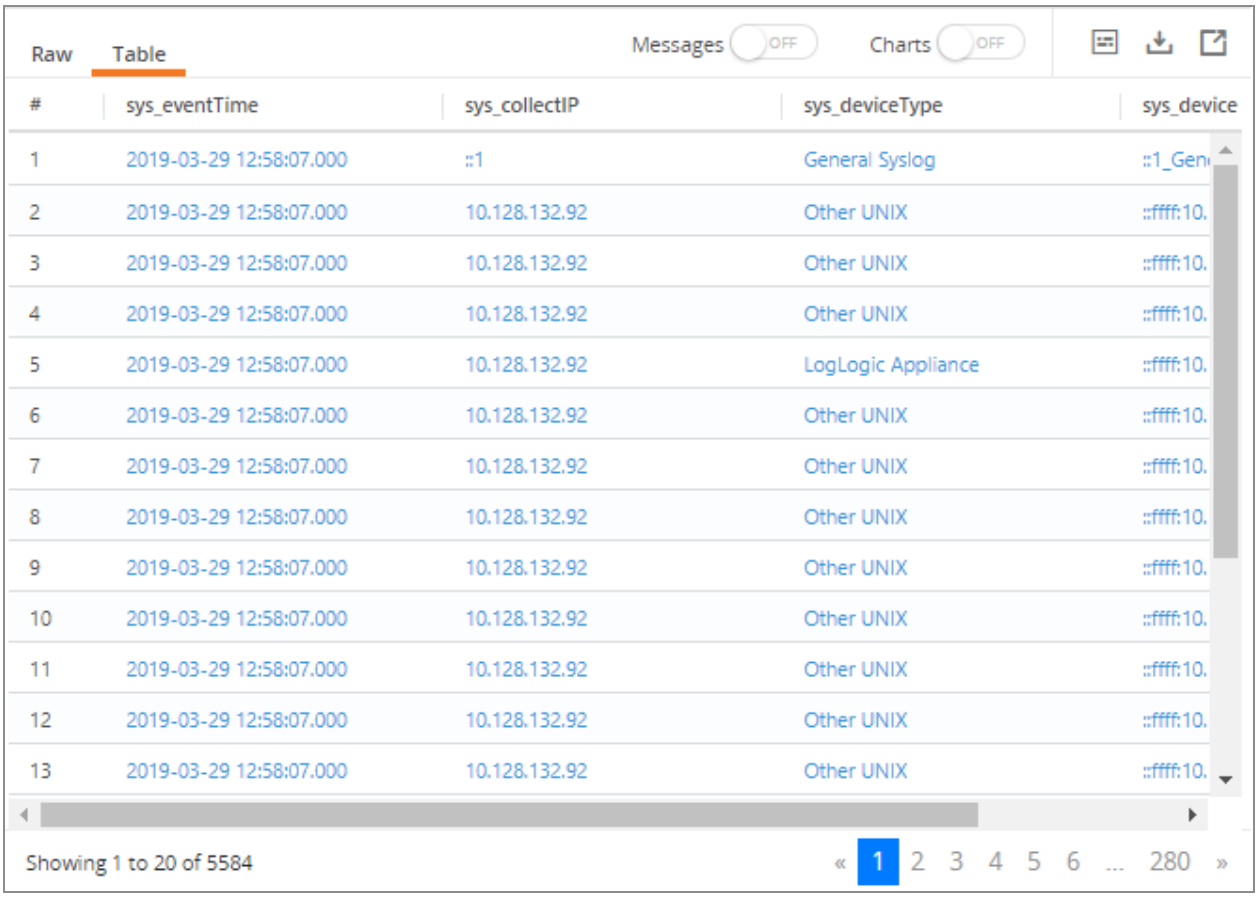

The same result set can be viewed in the Raw Data Format.

Using the table format, you can perform the following tasks:

- **Showing or hiding event body**

  Click Messages **on** or **off** to show or hide the event body in the `sys_body` column.

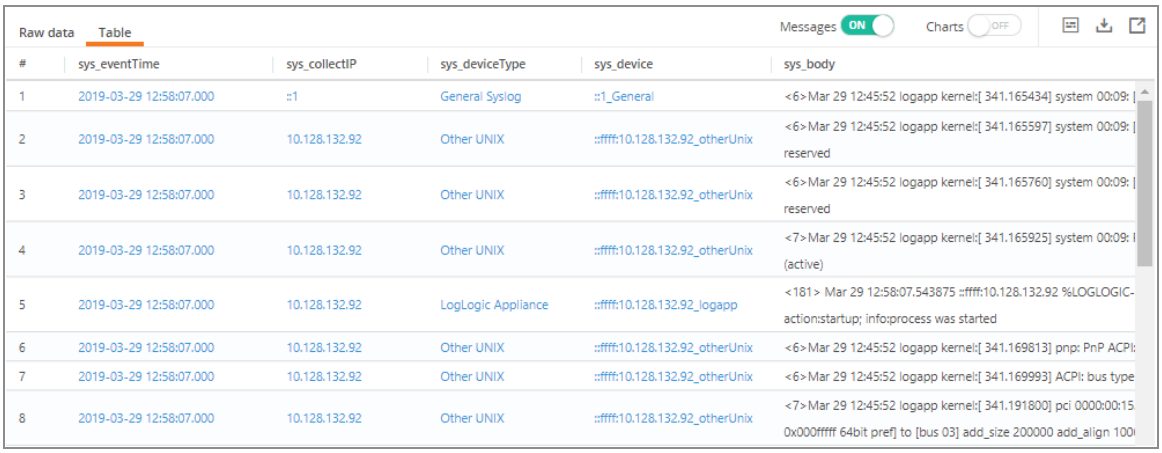

Monthly Index

When enabled, the monthly index feature increases the performance of the searches spanning monthly time ranges.

The monthly index adds a month-based index on top of the existing hourly indexes and enables searches to quickly locate the search terms within a given month.

Note:

- Only the advanced search uses monthly indexing.
- Monthly indexing can be enabled only if the Advanced Search features are enabled. To enable the advanced features and the monthly index feature, contact your administrator.
- To troubleshoot performance issues, you can disable monthly indexes in a query by specifying useMonthlyIndex=false after the OPTIONS keyword. For example:

```
use system | OPTIONS useMonthlyIndex=false | 'logapp' |
group by hours(sys_eventTime)
```

Artificial Intelligence Queries

LogLogic LMI includes artificial intelligence and machine learning capabilities to classify information from unknown log sources. This is achieved by using the TensorFlow trained model to automatically classify logs of access or audit types. By using an additional classifier in the advanced search query, you can view the additional information in the search results.

You can use the training model classification in advanced search queries, correlation, aggregation rules, and Bloks.

Before using the training model, ensure that Advanced Features are enabled on the appliance. To ensure that the training model is set up on the appliance, contact your LogLogic LMI administrator.

Example - Advanced Search

To search logs of a particular log source using the TensorFlow training model, use the following advanced search query:

```
Use system | options useClassifier='ll_tax_v1'
```

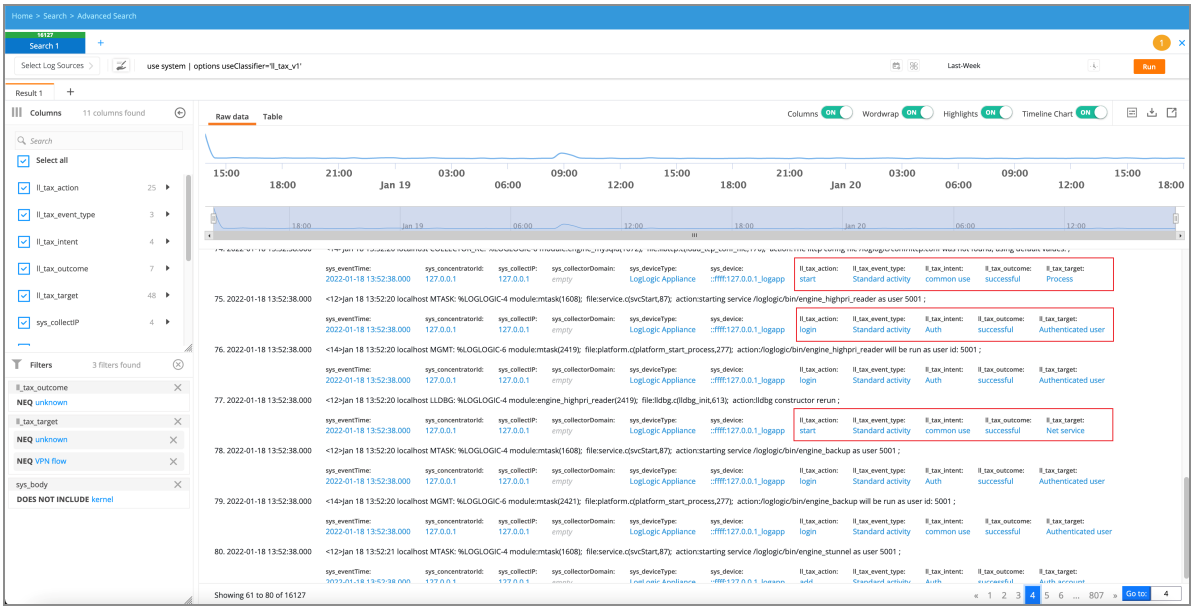

For more information about the classifier options and the result columns, see OPTIONS Statement.

Enrichment Lists

Running searches is often a very static experience. Users search for key words or phrases that they know in order to return specific results. An enrichment list is used to map a lookup key to an enriched value.

If the data you want to search is more dynamic and changes often, it might be difficult to create a query that would collect the right information. In such cases, you can use an enrichment list and reference the list in any query or in Predefined EQL Functions to achieve accurate results. From the **Advanced Search** tab, you can use an enrichment list in your search query. LogLogic LMI provides some built-in enrichment lists, for example, `mapRuleAction`. You can refer to the built-in lists as a guideline to create your own.

Use the following syntax to run a search query with an enrichment list:

```
use <nameOfDataModel> | $<NameofEnrichmentList>(<lookup_key> [,<default_
```

value>])='<*comparison_value*>' where:

- <*nameOfDataModel*> is the name of the data model to be queried
- <*NameofEnrichmentList*> is the name of the enrichment list
- <*lookup_key*> is the key to be looked up. If it is a column name, the value of the column for each row returned is used to look up the enriched value. It can also be a constant, in which case it always returns the same result. If there is no mapping for the lookup key, the enriched value is NULL.
- *default_value* is an optional parameter which, when provided, is the default return value when there is no mapping for the lookup key.
- <*comparison_value*> is the value that you want compared with the enriched value produced by the enrichment list.

For example, consider the search query:

```
use LogLogic_Appliance |$ipBlackList(sys_collectIP)='blacklisted'
```

The Enrichment List `ipBlackList` contains some value mapped to the key name `blacklisted`. This value is searched in the column `sys_collectIP` of the data model `LogLogic_Appliance`. We want the search result to display all log events that match the comparison value `blacklisted`.

Note: The behavior of the operators = and != is similar to that in regular SQL. If one side of the comparison has the value NULL, regardless of which operator is used, the result of the comparison is NULL, which is evaluated as FALSE in a WHERE clause. This results in both = and != returning FALSE if there is no mapping for <*lookup_key*>, which might not be what you expect, especially for !=. A better solution is to use the !== operator instead of !=. This returns TRUE if one side of the comparison is NULL, which is a more expected result. For more information about the operators == and !==, see FILTER Statement.

See also:

- Creating an Enrichment List
- Editing an Enrichment List
- Using enrichment lists in distributed Advanced Search

Note: This feature is available only if the Advanced Features option is enabled by using the **Advanced Features** option on the Administration > System Settings > General Settings page.

Viewing and Searching Enrichment Lists

From the **Management > Advanced Features > Enrichment Lists** page, you can view, find, or filter enrichment lists in the following ways:

Task	Steps
View enrichment lists based on filters	You can use filters to easily find enrichment lists. Click the **View** list to select the required filter. The following options are available: • All • Created by me • Created by system • Imported
Find enrichment lists	You can quickly find the desired enrichment lists by typing the enrichment lists name in the **Find** field. As you start typing the enrichment lists name in the **Find** field, the Enrichment Lists page is automatically refreshed showing your selection.
Sort enrichment lists	You can sort any column in ascending or descending order. Click on the column name or click the arrow (that is displayed on the right side of the column name when you click in the column) to sort the column.
Show or hide columns	You can show or hide columns, except the mandatory column, from the table. Click to view all available columns in the table. Select the check box to show the column. Clear the check box to hide the column from the table. The Enrichment Lists page is updated immediately.

Managing Enrichment Lists

You can manage enrichment lists from the **Management > Advanced Features > Enrichment Lists** page:

> **Note:** The duplicate, delete, and move icons are enabled after you select one or more enrichment lists.

Task	Steps
View and search enrichment lists	See Viewing and Searching Enrichment Lists
Create or edit an enrichment list	See: • Creating an Enrichment List • Editing an Enrichment List
Delete	Select the enrichment list and click the delete icon. After you delete an enrichment list, it cannot be recovered. **Note:** You cannot delete built-in and system-created enrichment lists.
Duplicate	Select the enrichment list and click the duplicate icon.
Move	Select the enrichment list and click the move icon. Select the group to which you want to move it. **Note:** You cannot move built-in and system-created enrichment lists.
Rename	Select the enrichment list and click **Edit** in the Details panel. Then rename the enrichment list and click **Save**.

Task	Steps
	Note: • You cannot rename an enrichment list using the rename icon . • You cannot rename built-in and system-created enrichment lists.

Enrichment List Groups

The **Management > Advanced Features > Enrichment Lists** page displays lists in groups. The **All Enrichment Lists** group displays all enrichment lists.

For information about groups, types of groups, and managing groups, see the Groups section.

Creating an Enrichment List

You can create an enrichment list from **Management > Advanced Features > Enrichment Lists**.

LogLogic LMI provides a few built-in Enrichment Lists. You can create your own, if required.

1. Go to **Management > Advanced Features > Enrichment Lists**.
2. Click **Create New Enrichment List**.
3. In the **Add Enrichment Lists** dialog box, specify the following fields for the enrichment list:
    - Parent Group: Select a parent group where you want to save the list.

      You can create a new group or select the **User** group, or select any user-created group

      Default parent group: When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.
    - Name: The name can include letters, numbers, or underscore (_).

- Description
- Value Type
- Mappings

4. Click **Save**.

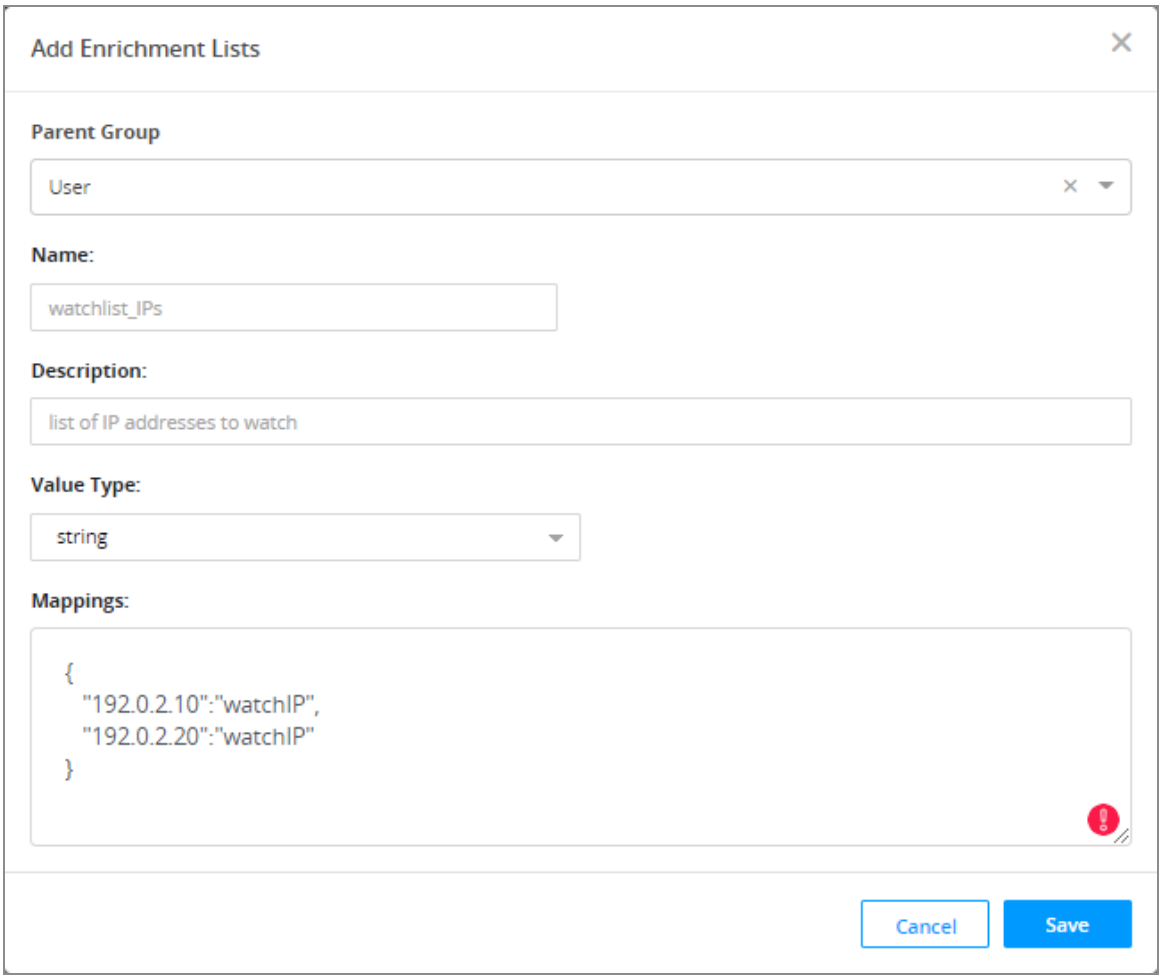

Result

The Enrichment List is displayed on the page.

If the mappings section contains the key default and there is no other mapping for the key used during lookup, the mapped value is returned as a default value. To override this default value at query time, specify the optional `<default_value>` parameter as a second parameter.

Note: Clicking any Enrichment List name in the list displays its information in a separate Details pane.

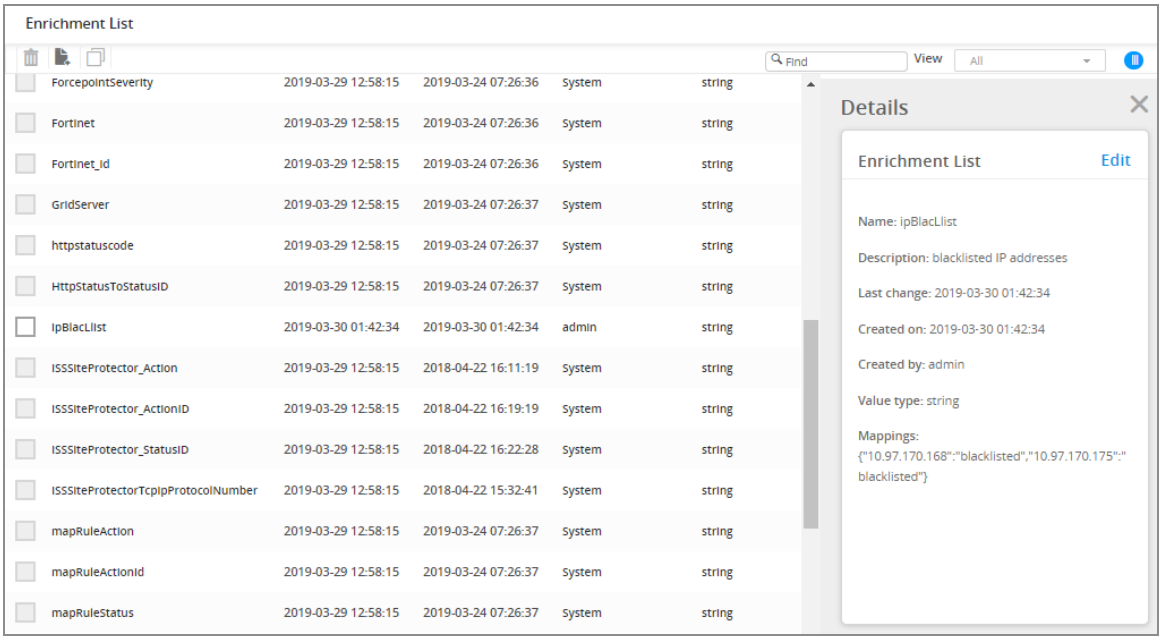

Editing an Enrichment List

You can edit an existing enrichment list from **Management > Advanced Features > Enrichment Lists**.
To edit an enrichment list, perform the following procedure. To delete an enrichment list, select the list and click the delete icon .

Procedure

1. Navigate to **Management > Advanced Features > Enrichment Lists**.
2. Click the name of the enrichment list and modify the required fields.

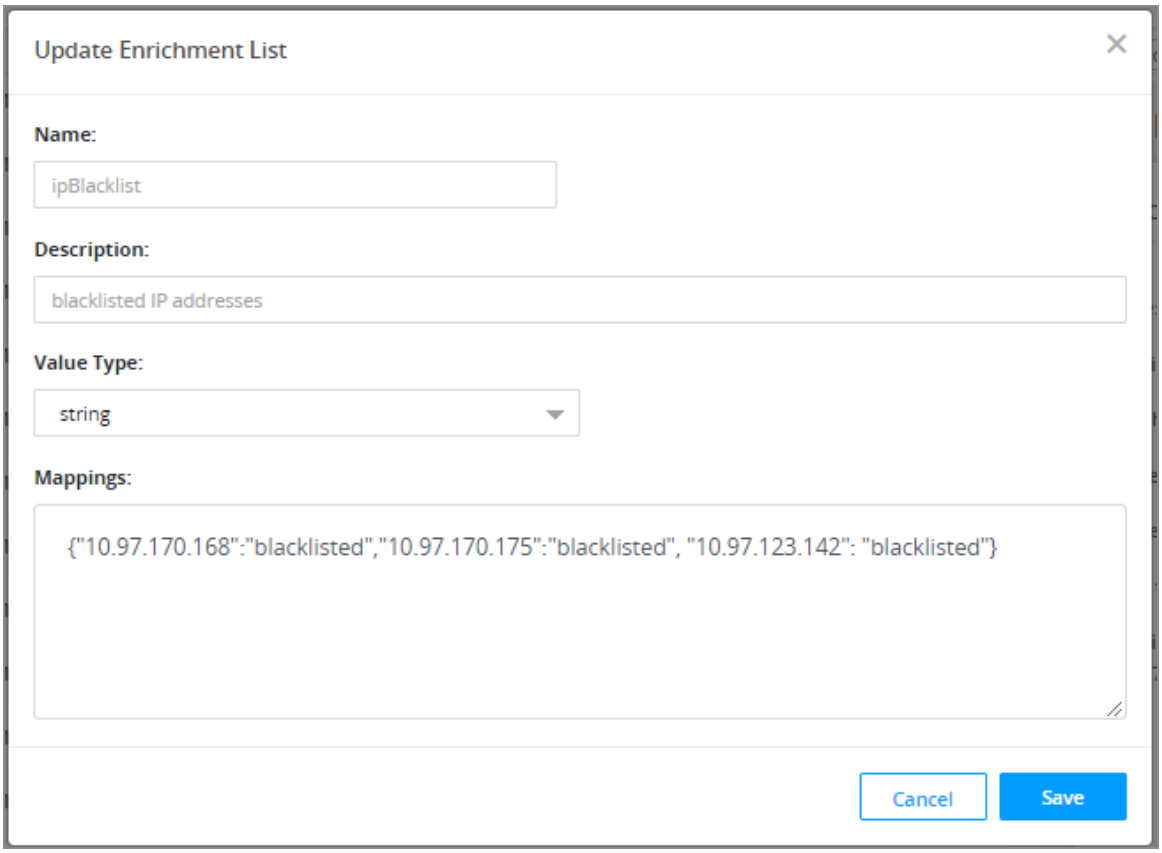

Note: When the mappings are stored, any white space characters are removed. Therefore, you might see the mappings in one line.

3. Click **Save**.

Infrastructure Queries

Infrastructure queries retrieve statistical information about LogLogic LMI data, such as its configuration or the amount of data ingested into LogLogic LMI. They work in the same way as other queries, except where indicated.

Infrastructure queries are not necessarily related to log events and do not contain an event time-stamp column like other data models. Thus, a time value need not be specified in infrastructure queries. If you use the Time field rather than embedding the time span explicitly in the query, then you must delete the value from the time Blok field to successfully execute an infrastructure query.

Fetching the Ingested Data

Query

```
USE LogLogic_System_Ingest_And_Index_Stats
```

Description

The data ingest count is taken when the files are collected by the appliance, whereas the index count is taken when those files (or, in the case of large pulled files, parts of those files) are actually indexed.

If a large amount of data is ingested in 1 hour, most of it might be indexed in the subsequent hour, resulting in a higher index count than the ingest count for that hour.

In a relatively quiet system, the data ingest count is updated only periodically, whereas the index count is updated every time a file is indexed. Therefore, the index count might be updated before the data ingest count. As a result, the index count might be more than the data ingest count for the most recent hour.

Therefore, for up-to-date values, check the count shortly after the end of any particular hour.

Query result

The following fields are returned in the search results:

Field	Description
lls_time	Time period to be queried
lls_ingestBytes	Number of data bytes ingested during the specified time period
lls_indexBytes	Number of bytes indexed during the specified time period
lls_indexSizeDelta	Change in the size of the physical disk space consumed by the index
lls_indexMsgCount	Number of messages indexed
lls_ingestMsgCount	Number of messages ingested
lls_unIndexBytes	Number of bytes that are not indexed yet

Examples

1. `USE LogLogic_System_Ingest_And_Index_Stats`

   Returns: the results in one-hour buckets

2. `USE LogLogic_System_Ingest_And_Index_Stats | COLUMNS lls_time, lls_ingestBytes / 1024.0 / 1024.0 / 1024.0 AS IngestGB, lls_indexBytes / 1024.0 / 1024.0 /1024.0 AS IndexGB`

   Returns: the results in GB

3. `USE LogLogic_System_Ingest_And_Index_Stats | COLUMNS DAYS(lls_time) as myTime, SUM(lls_ingestBytes ) / 1024.0 / 1024.0 / 1024.0 AS IngestGBPerDay, SUM(lls_indexBytes ) / 1024.0 / 1024.0 / 1024.0 AS IndexGBPerDay| GROUP BY DAYS(lls_time)`

   Returns: the result in one-day buckets in GB

Fetching the Advanced Application Pack Schema

The Advanced Application Pack (AAP) Schema query displays the metadata on the Advanced Search page.

Query

```
use LogLogic_AapSchema
```

Description

By using the `LogLogic_AapSchema` infrastructure query, you can view the common schema used by the built-in data models in LogLogic LMI, and the common schema used in creating custom data models. By using the common schema, you can leverage the normalization done in the LogLogic LMI models and optimize searches and reporting.

Query result

When you run this query, the fields from the AAP schema file are displayed in the search results area, for example:

- ll_eventID
- ll_deviceCategory
- ll_sourceUser
- ll_targetUser
- ll_sourceIP
- ll_targetIP

Correlation Alert SLA Status

The correlation alert SLA status query displays the status of correlation alerts.

Query

```
use LogLogic_AdvancedAlerts_SLA
```

Description

This query displays the status of alerts that are currently triggered in the system.

Query result

The following fields are returned in the search results:

Field	Data type	Description
lls_alertTime	TIMESTAMP	Date and time when the alert is created
lls_slaExpiration	DURATION	Time to expiration, displayed as Duration data type
lls_ruleName	STRING	Rule name
lls_state	STRING	Alert state
lls_ackTime	TIMESTAMP	Time when the alert is acknowledged
lls_timeToRespond	DURATION	Time to respond to or acknowledge the alert Displayed as Duration data type.
lls_category	STRING	Alert category
lls_severity	STRING	Alert severity
lls_comment	STRING	Comments you provide when acknowledging an alert

Aggregation Rule Metrics

To view the progress of aggregation computation, run the Aggregation Rule Metrics query in Advanced Search.

Query

```
USE Loglogic_Aggregation_Rules_Metrics
```

Description

When a new aggregation rule is created, it starts computing the metrics for the given GROUP BY expressions. It aggregates data for all events received in the system after rule start time and the data for the past events specified in the **Compute aggregation for** field. To view the progress of the aggregation computation, you can run this infrastructure query in the Advanced Search.

Query Result

The following fields are returned in the search results:

Field	Description
aggregationRuleName	The name of the aggregation rule
totalEventCount	The number of events processed by the aggregation rule
successCount	The number of valid events computed by the aggregation rule query
errorCount	The number of events that failed to validate by the aggregation rule query
savedRecordCount	The number of records stored in the database for the aggregation rule
dataSizeInMB	The size of saved records in megabytes (MB)
replayProgress	The number of events processed out of the total events for historical computation (the aggregation computed before rule creation) Applicable only if the aggregation start time is specified in the **Compute aggregation for** field.
replayPercentage	Percentage progress of the historical computation
isMerged	Indicates (Yes or No) if the historical computation results are

Field	Description
	merged in the database with the ongoing, real-time computation results

Queries

You can view the following types of queries in the system:

Query	Description
Scheduled Queries	A list of EQL and SQL queries that are scheduled to be executed at the frequency that you have set.
Search Queries	A list of Advanced Search queries that are currently running in the system or have completed.
Tail Queries	Queries that run on near-real-time data after the data is indexed.
Distributed Advanced Search	Advanced Search queries that are run on a Management Station and distributed to Remote Appliances.

Scheduled Queries

You can schedule search queries to run at a selected time and frequency.

You can create a schedule to run the Advanced Search queries at the required time intervals and have the reports sent as email attachments.

Note: Only SQL and EQL queries are supported.

You can view a list of scheduled queries from any of the following locations:

- By clicking **Search > Advanced Search >**  **> Show Scheduled Queries**

> **Note:** This option is disabled if there is no text in the search field.

- By navigating to the **Management > Advanced Features > Queries > Scheduled** page

From the Scheduled Queries page, you can:

- View a list of scheduled queries.
- Click a query to open the Details pane that displays more information about the query.
- Create, update, or delete a scheduled query.
- Configure multiple schedules to execute the query and save its reports.
- Enable or disable query schedules:
  - To enable or disable selected schedules, click the toggle button next to the schedule name.
  - To disable all schedules, click the toggle button next to the title **Schedules**.
- Download the search results or send the results as email attachments. For more information, see Configuring Query Schedules.

Adding Scheduled Queries

You can add a query and then configure schedules to execute the query. You can configure multiple schedules per query.

> **Note:** Only SQL and EQL queries are supported.

Procedure

1. Navigate to any of the following locations:
   - On the **Advanced Search** page, type your query and then click **> Save as Scheduled Query**. The Add Scheduled Query dialog box opens, in which your query is already available in the Query field.
   - Go to **Management > Advanced Features > Queries > Scheduled** tab and click . The Add Scheduled Query dialog box includes the following sections:

- Query - Displays the name, description, and query text.
- Schedules - displays a list of schedules for the query. For information on how to add or modify schedules, see Configuring Query Schedules.

2. Type the name, description, and query.

   The following options are available:

   - Click **Validate** to check whether the query syntax is correct.
   - Click **Format** to format the query keywords.

   When saving an infrastructure query as a scheduled query, if you include a time filter in the query, the query is saved successfully. However, after the report is saved in the system, an error is displayed in the **Query Scheduled > View Executions** section.

3. Click **Save**.

Result

The query is added to the list of queries on the Scheduled Queries page.

Configuring Query Schedules

You can schedule the execution of your queries and set the frequency of execution.

Procedure

1. Open the Scheduled Queries page in one of the following ways:
   - Click **Management > Advanced Features > Queries > Scheduled**.
   - On the Advanced Search page, click  **> Show as Scheduled Queries**.
2. On the Scheduled Queries page, click a query to open its Details pane.
3. On the Details pane, click **Edit**. On the Edit Scheduled Query page, configure the query schedule.
4. To add a schedule, click the plus icon **+** or the **Add Schedule** button. The default schedule name is displayed, but you can change it, if required.
5. In the **Time Frame** section, enter the following information.

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td>Run</td><td>Select one of the following frequency types:
<table>
<tr><th>Frequency Type</th><th>Description</th></tr>
<tr><td>Once every few hours</td><td>Select the hourly interval from the Hourly Interval drop-down list.</td></tr>
<tr><td>Once a day</td><td>Select the time of the day from the At drop-down list.</td></tr>
<tr><td>Once every few days</td><td>a. Select the number of days from the Interval drop-down list.
b. Select the time of the day from the At drop-down list.</td></tr>
<tr><td>Once a week</td><td>a. Select the time of the day from the At drop-down list.
b. Select the day of the week from the Every drop-down list.</td></tr>
<tr><td>Once a month</td><td>a. Select the time of the day from the At drop-down list.
b. Select the day of the month from the Every drop-down list. You can also select Last day of the month.</td></tr>
</table></td></tr>
<tr><td>For the Period</td><td>Type the period on which the query should be executed. The default value is -12h. For valid time ranges, see Time Range Expressions.</td></tr>
</table>

6. To run the query immediately, click **Run Now**.

> **Note:** The **Run Now** button is enabled only when you are editing a schedule that has already been saved earlier.

The previous and next execution time is displayed in the **Previous Execution** and

Next Execution fields.

7. In the **Number of reports to be saved in the system** field, enter the number of execution results that you would like to be stored in the system.

   To view the stored reports, click **View Executions**. From the Executions page you can delete or download selected reports as a CSV file.

   > **Note:** The **View Executions** link is visible only when you are editing a schedule that has already been saved earlier. The link is not available when adding a new schedule.

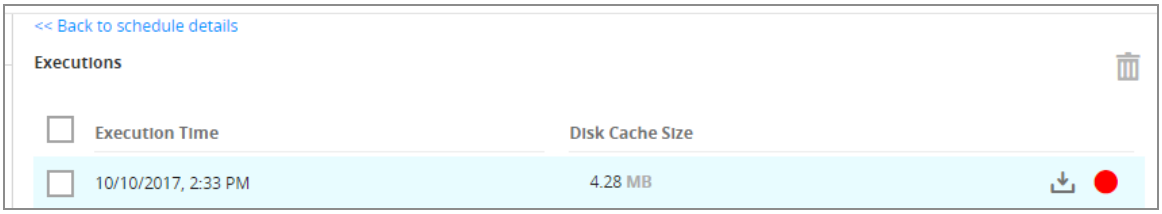

8. (Optional) **Email Recipients** section: If you want the report to be sent by email, enter the name and email ID of the email recipients and click the check mark to save. You can send the email attachments as a PDF, HTML, or CSV file.

   - If a user with the name and email ID is found in the database, the login name of the user is filled automatically in the **Login Name** field.
   - The file name of the attachment includes the schedule name followed by the date and time when the email is sent. The file name format is `<scheduleName><date>T<time>.<fileExtension>`. For example, `CiscoASA_Schedule_1_2019-12-23T01:03:22.226017.csv`.
   - The attachments are compressed before being attached to the email. If the attachment size is more than the maximum permitted value, the email does not include the attachment. You must download the attachment from the Executions page. To have the attachment size limit increased, contact your administrator.
   - If the number of columns you selected or the content within the columns is too large, then the PDF report might be distorted. In such a case, reduce the number of selected columns. Otherwise, select the CSV or HTML format as the report type.

Tip: If there is any error while executing the query or sending emails, an error icon is displayed. Hover over the error icon to read the error messages.

If you do not receive an email, then check if the mail has been redirected to the Spam folder.

Search Queries

The **Management > Advanced Features > Queries > Search** page displays the queries that are currently running in the system or have completed.

The following query information is displayed as a table:

- Query ID
- Query
- Total events
- Progress percentage

To delete queries, select the queries by clicking the corresponding check boxes, and then click .

Even if the execution of a query has completed, the query remains in the system cache until it is deleted. A query can be deleted in any of the following ways:

- On the GUI:
  - By closing the Advanced Search tab for the query
  - By deleting the query from the **Queries > Search** page
- From the REST API:
  - By using the `deleteQuery()` API
  - If the time to live value of the query elapses.

    The `timeToLive` parameter is applicable only to queries created with REST API. It indicates the number of seconds of time after which the query is automatically deleted. The default value is 120 seconds.

Tail Queries

Tail queries run on near-real-time data after the data is indexed.

The query results display new incoming events that match the query criteria. Because of the nature of these queries, they never end; you must cancel or delete them manually.

You can query real-time data from the Advanced Search or the Data Grid widget, by using any of the following methods:

- Use the `TAIL` keyword in the query
- Select **Real Time** from the time filter list (only on the Advanced Search tab)

Features of TAIL queries

- Real-time search results are appended to the results list. Navigate to the last page to view the latest results.
- Similar to other search queries, you can click **Pause** or **Resume** at the top of the results page to control the streaming of logs.
- Real-time streaming stops if you switch to another page or log out. However, it automatically resumes when you return to the search page or log in again.
- You can use `LIMIT` and `BUFFER` statements in a tail query:
    - `LIMIT` specifies the maximum number of results that must be displayed. The results are displayed in three pages. After retrieving more results, the same number of initial results are removed from the first page. Thus, if the limit is 50, only the latest 50 results fetched are displayed.
    - `BUFFER` specifies the maximum streaming time to wait before returning results for the Tail query. You can use `BUFFER` with the `TAIL` keyword; not with the **Real Time** option from the time filter.

Example of TAIL query

The example query that includes the `TAIL` keyword:

```
USE system |TAIL BUFFER 500 ms | LIMIT 50
```

In the following example, the time filter is set to **Real Time**:

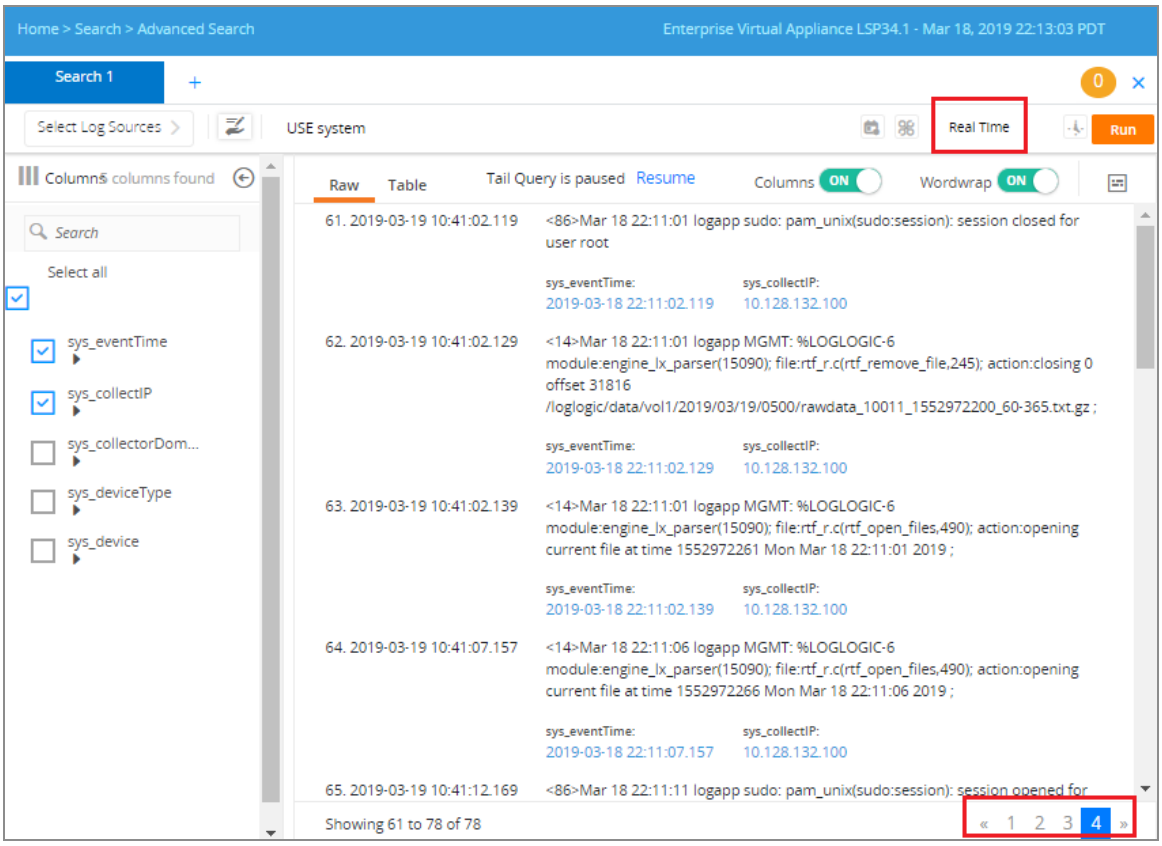

Limitations

Tail queries have the following restrictions:

- The results are always sorted by time.
- Filter options and the timeline chart view are not available.
- The options to download the results and to open search in a new tab are not available.
- The Tail keyword cannot be used in:
    - Infrastructure queries
    - Distributed Advanced Search
    - GROUP BY statement (aggregation)
    - Distributed aggregation functions
    - SQL queries

Instead, use the `where` clause to achieve the same results. For example:

- EQL query: `use LogLogic_Appliance | TAIL`
- SQL query: `select * from LogLogic_Appliance WHERE sys_collectTime > NOW`

Distributed Advanced Search

Running a query on one appliance returns results only from that appliance. Distributed Advanced Search queries can be run on a Management Station and distributed to multiple Remote Appliances in the setup. The results from all appliances are consolidated by the Management Station.

To run a distributed Advanced Search query on multiple appliances, specify a list of appliances by using the `sys_concentratorId` column or the `DeviceInGroup` query function within the query that you run on the Management Station.

Important: The query is sent to only those appliances to which you have access permissions. Furthermore, on the Remote Appliances, the query is run only on those log sources to which you have access permissions. If you run a query on the Management Station and omit the appliance list, the query is run only on the Management Station and not on the Remote Appliances.

Requirements to run distributed Advanced Search queries

You can use distributed queries in Advanced Search if all of the following requirements are met:

- The Management Station setup includes the required Remote Appliances. For information about how to set up a Management Station, see "Manage Appliances with Management Station" in the *TIBCO LogLogic® Log Management Intelligence Administration*.
- You have permission to access the Remote Appliances and devices on which the search query is to be run.
- Advanced Features are enabled on the Management Station and Remote Appliances.

- The data node port (9621) on each Remote Appliance is accessible by the Management Station.

Using the sys_concentratorId column

You can specify the value of `sys_concentratorId` using =, ==, or the IN construct. Valid values for `sys_concentratorId` can be any of the following:

- Explicit appliance IP addresses, where the IP addresses must exist in the Management Station list on the **Management > Management Station** page.
- Appliance grouping shortcuts, similar to those used in the Log Source Picker in Advanced Search.
- Enrichment lists

Examples with IP addresses and shortcuts used in the `sys_concentratorId` column

Type of value	Value of `sys_concentratorId`	The query is sent to...
Appliance IP address	sys_concentratorId = "127.0.0.1"	Only to the Management Station, which is always specified as 127.0.0.1. This is the default value if `sys_concentratorId` is not specified.
	sys_concentratorId IN ("127.0.0.1, *<IP1>*")	The Management Station itself and the specified Remote Appliance having IP address IP1.
	sys_concentratorId=" *<IP2>*"	Only to the specified Remote Appliance having IP address IP2.
Appliance grouping shortcuts	sys_concentratorId IN "ALL"	All members of the Management Station setup.
	sys_concentratorId IN "ALL_LX"	The types of appliances specified in the shortcut.
	sys_concentratorId IN "ALL_ST"	

Type of value	Value of sys_concentratorId	The query is sent to...
	sys_concentratorId IN "ALL_MX"	
	sys_concentratorId IN "ALL_LX_MX"	
	sys_concentratorId IN "ALL LX MX"	
	sys_concentratorId IN "ALL LX/MX"	

Using Enrichment Lists in the sys_concentratorId column

Enrichment lists can also be used in distributed Advanced Search queries to specify sets of appliances to which the query must be sent.

For example, you can create an enrichment list called `MyApplianceSets` with mappings such as:

```
{ "set1" : "127.0.0.1, <IP1>, <IP2>",
"set2" : "<IP3>" }
```

If you run a query like this:

```
USE system | sys_concentratorId = $MyApplianceSets('set1') | sys_body
CONTAINS 'joe'
```

the query is sent to the IP addresses in `set1`.

The following query is an example of using an enrichment list when using a device group in a distributed Advanced Search:

```
USE system | DeviceInGroup('MyLogLogicAppliances', $MyApplianceSets
('set1') ) | sys_body CONTAINS 'joe'
```

Note: The IP addresses in the enrichment list must match those in the Management Station setup.

Examples with enrichment lists used in the `sys_concentratorId` column

Type of value	Value of sys_concentratorId	The query is sent to...
Enrichment list	sys_ concentratorId=$MyLogLogicAppliances ("set1")	All appliances in set1 from the Enrichment List `MyLogLogicAppliances`
Enrichment list and IN	sys_concentratorId IN ($MyLogLogicAppliances("set2"))	All appliances in set2 from the Enrichment List `MyLogLogicAppliances`
Combination	sys_concentratorId IN ("*<IP_1>*", "ALL_ LX", $MyLogLogicAppliances("set1"))	The following appliances: • The appliance with IP address IP_1 • All LogLogic LX Appliances • All appliances in set1 from the Enrichment List `MyLogLogicAppliances`.

Using the DeviceInGroup function

You can specify a list of local or global device groups by using the `DeviceInGroup` function:

```
DeviceInGroup("<Group Name>", "<IP list or appliance grouping shortcuts>")
```

where:

- *<Group Name>* is the name of the device group
- *<IP list or appliance grouping shortcuts>* specifies the list of appliance concentrator IP addresses on which the device group is to be evaluated

Examples of the `DeviceInGroup` function

Value of DeviceInGroup	The query is sent to...
DeviceInGroup ("MyLogLogicAppliances", "ALL")	All log sources in the device group `MyLogLogicAppliances` on all appliances on which you have access permission.
DeviceInGroup ("GlobalGroup", "ALL_LX")	All log sources in the device group `GlobalGroup` on all LogLogic LX Appliances which have group members and on which you have access permission.
DeviceInGroup ("LocalGroup", "*<IP1>*")	All log sources in the device group `LocalGroup` on the appliance at IP address IP1.
DeviceInGroup ("LocalGroup", "*<IP2>*, *<IP3>*")	All log sources in the device group `LocalGroup` on the appliances at IP address IP2 and IP3.
DeviceInGroup ("GlobalGroup", $MyLogLogicAppliances ("set1"))	All log sources in the global group `GlobalGroup` on the appliances specified in the set1 entry in the Enrichment list `MyLogLogicAppliances`. The query is sent to only those appliances with members in `GlobalGroup` and on which you have permission.

Local device groups

You can also run the query on local device groups that use the same name on multiple appliances. For example, if the local device group `MyLocalGroup` exists on two appliances, Appliance1 (having IP address IP1) and Appliance2 (having IP address IP2), and you have permissions on `MyLocalGroup` on both appliances, then using the following function in a distributed Advanced Search query on the Management Station:

```
DeviceInGroup("MyLocalGroup", "<IP1>, <IP2>")
```

causes the query to be sent to both appliances Appliance1 and Appliance2, for their corresponding local group `MyLocalGroup`. Search results from all devices belonging to the `MyLocalGroup` device group on both appliances are sent back to the Management Station.

Global device groups

If you create a global device group `GlobalGroup` on the Management Station that includes log sources from Appliance1 and Appliance2, you can use the global group name in the `DeviceInGroup` function:

```
DeviceInGroup("GlobalGroup", "ALL")
```

This query is sent to Appliance1 and Appliance2.

Global Dynamic Device Groups

In regular expression search, you cannot run distributed queries on dynamic device groups because dynamic device groups cannot have a global scope.

However, for distributed Advanced Search queries, if you create **local** dynamic device groups that have the same group name on the appliances involved, and run a distributed Advanced Search query for those local dynamic device groups, you can achieve the effect of a global search result on dynamic device groups.

For more information about device groups, see "Device Group Management" in *TIBCO LogLogic® Log Management Intelligence Administration Guide*.

Limitations and errors

Limitations

You cannot use the Tail functionality in distributed Advanced Search.

Errors

An error is displayed if:

- The specified appliance shortcut is invalid.
- An appliance specified in the distributed query is not a member of the Management Station setup.

Examples of distributed Advanced Search queries

These are a few examples of complete queries that you can use in distributed Advanced Search:

Using `sys_concentratorId`:

- USE system | sys_concentratorId="ALL" | sys_eventTime in -10m
- USE system | sys_concentratorId IN ("127.0.0.1") | sys_eventTime in -10m
- USE system | sys_concentratorId IN ($MyApplianceGroups("remotes")) | sys_eventTime in -10m

Using `DeviceInGroup` function:

- USE system | DeviceInGroup("<group name>", "ALL") | sys_eventTime in -1y
- USE system | DeviceInGroup("<group name>", "RemoteAppliance1") | sys_eventTime in -1y

Using `DeviceInGroup` in the projection:

- USE system | sys_deviceType = "General Syslog" | COLUMNS sys_collectIP, DeviceInGroup("LocalGroup1") as dig | sys_concentratorId = "ALL" | sys_eventTime in -1y
- USE system | COLUMNS sys_deviceType, sys_device, sys_eventTime, DeviceInGroup("All Other UNIX") as isOtherUnix | sys_eventTime in -1h

Bloks

To analyze your data faster, you can create different types of Bloks in LogLogic LMI to help you accelerate your search process.

A Blok is a contextual element or filter that fits with other elements to form a search query. Bloks are reusable elements of a query. You can combine many types of Bloks together to create complex queries. Build and save different Bloks that can be used in future searches rather than searching every time by manually typing the same filter.

LogLogic LMI supports the following types of Bloks:

- Filter Bloks: contain filter statements, aggregation rules
- Correlation Bloks: contain correlation rules
- Time Bloks: contain absolute and relative time ranges

You can add one or more filters in a Blok. If you realize that you need to add another filter to the existing Blok, you can add more filters or build another Blok.

You can add new Bloks and modify existing Bloks from the Search tab. Similarly, you can manage all types of Bloks in a central location by clicking the **Management > Advanced Features > Bloks** menu.

When entering a Blok name in the search query field, start with the prefix defined for each type of Blok. Content assist can help you by showing all possible values for that type of Blok. For time Bloks, you can select the Blok in the time range field. The following is a list of prefixes that you can use in the search query field:

- *filter.Blok name*
- *correlation.Blok name*

For detailed information on how to create a Blok, see Creating a Blok.

Blok Groups

The **Management > Advanced Features > Bloks** page displays Bloks in groups. The **All Bloks** group displays all Bloks.

For information about groups, types of groups, and managing groups, see the Groups section.

Filter Bloks

You can create filter Bloks that contain one or more filters.

Each filter comprises one or more terms. A filter Blok supports valid EQL or SQL statements.

You can have one or more filters in a Blok. If you realize that you need to add another filter to the existing Blok, you can add more filters or build another Blok. Multiple Bloks of different types can be used in a single search query. For detailed information about valid filters, see FILTER Statement.

Example

Create a filter Blok and use it in a search query:

Create and save a filter Blok that includes `sys_deviceType='Other UNIX' AND sys_body like '%security%'`. Now when you run a query using this Blok, only events with `Other UNIX` and `security` are retrieved.

Use this filter Blok and add another element or filter to it, for example, type `sys_deviceType='Cisco ASA'` to the same query to create a more complex query. For example, `filter.`*`Blok name`*` AND sys_deviceType='Cisco ASA'`. Now when you run a query using this Blok, events with `Other UNIX`, `security`, and `Cisco ASA` are retrieved.

Correlation Bloks

For your forensic needs, you can search data using Correlation Bloks.

Correlation Bloks are created using Event Correlation Language Reference (ECL). You create a correlation Blok and use the Bloks on real-time data to set up triggers. The triggers are, in turn, configured to send alerts on the real-time data.

You can also use correlation Bloks in Advanced Search to search historical data and analyze the patterns in the data. When entering a Blok name in the **Search** field, start with the prefix *correlation.* for any existing correlation Blok. Content assist can help you by showing all possible values for that type of Blok. The correlation search results are displayed every time the rule's conditions are met. For more information, see Using Correlation Bloks in Advanced Search.

You cannot combine a correlation Blok with other Blok types in a single query. Only one correlation Blok can be used at a time in a query. In a correlation Blok query if there are more than one million events for the defined time duration, only the first one million events are processed for better performance. In such cases, it is a best practice to reduce the time duration to retrieve accurate results.

Note:

- The Blok name cannot contain a period (.).
- On the Advanced Search page, you cannot filter search results for Correlation Blok by clicking the timeline chart or using the chart time slider. This is because LogLogic LMI does not support adding filters and running subqueries on the search results for a Correlation Blok.

Using Correlation Bloks in Advanced Search

You can use correlation Bloks in Advanced Search to search historical data and analyze the patterns in the data.

Procedure

1. Click next to the **Search** field, click **Choose Blok**, and then select the correlation Blok from the list.
2. Enter the time period in the **Time** field and click **Run**.

   > **Note:** For valid time expressions, see Time Range Expressions. If you enter an invalid value in the **Time** field, the range 1970-2070 is used.

   The correlation results display all events that contributed to the triggering of the correlation rule. Based on the correlation rule, the columns (correlation events and event groups) are extracted in a table format. Each row helps you analyze the associated values of the columns and event groups. If you refresh the search tab by clicking **Search > Advanced Search** or refresh the browser, the search tab closes.

   The following illustration displays the defined correlation rule in the **Search** field and retrieved events in the **Timeline Charts**, **Columns**, and **Data** panels. When you use Correlation Bloks in advanced search, hovering your mouse over any part of the chart displays the number of correlation events instead of the message count.

   Correlation Rule and Timeline Charts

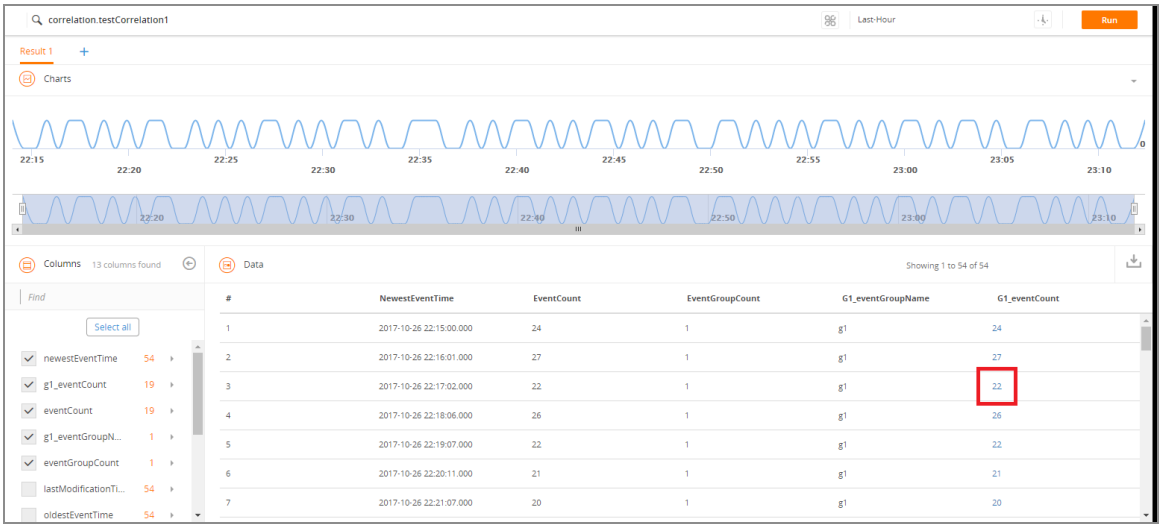

3. Click the event count link to view the event details on a new **Search** tab.

> **Note:** The event count link is available only when the count is less than 1024.

Click the event count link (22 in the illustration), and the new search tab opens with the auto-generated EQL query in the **Search** field for the events associated with that event count. The **Timeline Charts**, **Columns**, and **Data** panels display the results associated for that event count as shown in the following illustration.

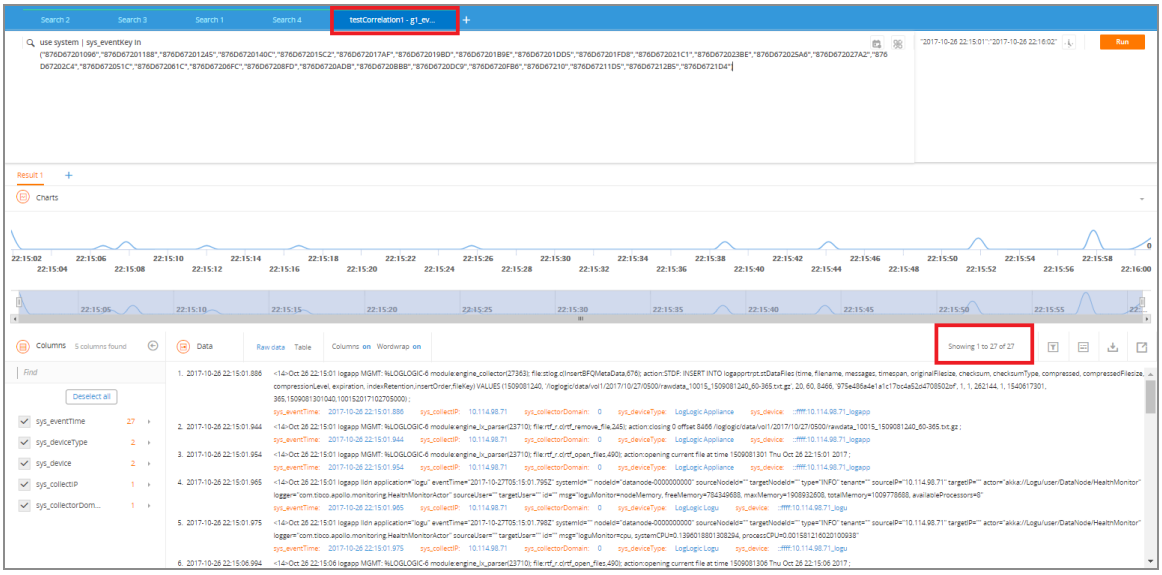

Time Bloks

Analyzing events based on a certain time range can help correlate results and find the root cause faster.

You can narrow your search results to a specific time range using the Time Blok. You can use the preset time Blok or create your custom time Blok that you can use any time.

Each time Blok is translated in a statement before executing the query. When entering the time Blok name in the **Search** field, start with the prefix *time* for any existing time Blok. You can use Content Assist to see all possible values for that type of Blok. For detailed information on how to create a time Blok, see Adding a Time Blok.

By default, the time range is set to last hour. You can define the absolute or relative time. For valid time ranges, see Time Range Expressions.

Viewing All Time Bloks

The default or user-defined time Bloks can be easily used to quickly search your data. The default time Bloks have preset time ranges. You cannot modify or delete the default time Bloks. However, you can update or delete user-defined time Bloks.

Procedure

1. From the Search page, click  located next to the **Time** field, and select **Choose Blok**.
2. In the **Find** field, type the Blok name to quickly find the desired time Blok.
3. Select the **Blok name** from the list of Bloks. The **Description** and **Source statement** fields are auto-filled based on the selected Blok.
4. Click **Save** to add the Blok in the **Time** field. The selected time Blok is displayed in the **Time** field.
5. Click 

to view results for the defined time range.

Adding a Time Blok

If you usually search for events that are in the specific time range, you can create a custom time Blok for that time range and save it for later use.

Procedure

1. From the Search page, click located next to the **Time** field and click **Select a date range** to open a window.
2. Specify the date and time in the **From** and **To** fields. Time must be in **Hours** and **Minutes** and click **OK**. The selected date and time range is displayed in the **Time** field. Alternatively, type in the time expression in the **Time** field. Content Assist shows you typeahead or contextual matches and completions for each keyword as you type it into the search field. To define a valid time statement, see Time Range Expressions.
3. To save a new time Blok, click next to the **Time** field and select **Save as Blok**. Alternatively, to add a new Blok, select **New Blok**.
4. In the Add new Blok window, enter the information in the following fields:

Field	Description
Name	Enter the name of the Blok. It must be a unique name that consists of a single word. This is a mandatory field.
Description	Enter the description of the Blok.
Source Statement	The statement of the source (time expression).

5. Click **Save** to save the new time Blok. The new time Blok is added in the Choose **Blok list**.

Modifying Time Bloks

You can modify the custom time Bloks at any time.
You cannot modify default time Bloks.

Procedure

1. From the **Search** tab, update the time range expression in the **Time** field. For detailed information about valid time statements, see Time Range Expressions.
2. To save a new time Blok or update the existing Blok, click | next to the **Time** field and select **Save as Blok**.
3. Update the information. For information about each field, see Adding a Time Blok.
4. Click **Save** to save the new time Blok. The new time Blok is added in the Choose **Blok list**.

Viewing and Searching Bloks

You can view and search Bloks, or use Bloks to search data.

You cannot modify or delete the system Bloks. However, you can update or delete any custom Bloks. See Managing Bloks.

Viewing Bloks

From the **Management > Advanced Features > Bloks** page, you can view, find, or filter Bloks in the following ways:

Task	Steps
View Bloks based on filters	You can use filters to easily find Bloks. Click the **View** list to select the required filter. The following options are available: • All • Created by me • Created by system • Imported
Find Bloks	You can quickly find the desired Blok by typing the Blok name in the **Find** field. As you start typing the Blok name in the **Find** field, the Bloks page is automatically refreshed showing your selection.
Sort Bloks	You can sort any column in ascending or descending order. Click on the column name or click the arrow (that is displayed on the right side of the column name when you click in the column) to sort the column.
Show or hide columns	You can show or hide columns, except the mandatory column, from the table. Click to view all available columns in the table. Select the check box to show the column. Clear the check box to hide the column from the table. The Bloks page is updated immediately.

Searching Data by Using Bloks

From the **Search > Advanced Search** page, you can use the system or existing Bloks to quickly search your data.

Procedure

1. From the **Search > Advanced Search** page, click located next to the **Search** field, and select **Choose Blok**.

2. Select the type of Blok from the list. The options are All, Filter, Correlation, and Time Bloks.
3. In the **Find** field, type the Blok name to quickly find the desired Blok.
4. Select the **Blok name** from the list of Bloks. The **Description** and **Source > statement** fields are auto-filled based on the selected Blok.
5. Click **OK** to add the Blok in the **Search** field. If you select a time Blok, it is displayed in the **Time** field.
6. Click 

to **view results** for the defined Blok.

Managing Bloks

A Blok is a contextual element or filter that fits with other elements to form a search query. Build and save different Bloks that can be used in future searches rather than searching every time with the same filter.

The system Bloks have preset values. You cannot modify or delete the system Bloks.

You can manage all types of Bloks from the **Management > Advanced Features > Bloks** page.

Note: The duplicate, delete, and move icons are enabled after you select one or more Bloks.

Task	Steps
View and search Bloks	See Viewing and Searching Bloks.
Create and edit Bloks	• Creating a Blok • Editing a Blok **Note:** You cannot edit system-generated Bloks.

Task	Steps
Delete	Select the Blok and click the delete icon . **Note:** • You cannot delete built-in and system-created Bloks. • You can delete a user-defined Blok at any time. After the Blok is deleted, search queries are not affected; but you cannot run a new query with a deleted Blok. Queries in the **Search > History** tab that use the deleted Blok cannot be run again.
Duplicate	Select the Blok and click the duplicate icon .
Move	Select the Blok and click the move icon . Select the group to which you want to move it. **Note:** You cannot move built-in and system-created Bloks.
Rename	Select the Blok and click **Edit** in the Details panel. Then rename the Blok and click **Save**. **Note:** You cannot rename a Blok using the icon.

Creating a Blok

If you usually search for events that provide you with specific information such as user name or severity, you can create a custom Blok for that criteria and save it for later use.

Procedure

1. From the Search page, click the **Choose Blok** icon located next to the **Search** field, and select **New Blok**.

   Alternatively, go to **Management > Advanced Features > Bloks** and click **Create New Blok**.

2. In the Add new Blok dialog box, provide the following information:

a. Parent Group: Select a parent group where you want to save the Blok.

   You can create a new group or select the **User** group, or select any user-created group

   Default parent group: When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.

b. Select the **Blok type** from the list.

c. Name: It must be a unique name that consists of a single word with no special characters.

   The Blok name cannot include a period (.).The name can include letters, numbers, hyphen, or underscore (_).

d. Description

e. Enter the statement of the source in the **Source statement** field. Make sure to enter a valid syntax. Filter and Time Bloks support SQL, EQL, and ECL syntax. For syntax information, see:

   - EQL Search Syntax Reference
   - ECL Search Syntax Reference

3. Click **Save**.

Result

The new Blok is added in the **Choose Blok** list and is displayed in the **Search** field. It is also displayed on the All Bloks page and in the parent group that you selected.

Editing a Blok

You can modify the user-defined Bloks at any time.

> **Note:** You cannot edit system-generated Bloks.

Procedure

1. From the Search page, update the statement in the **Search** field. Content assist

shows you contextual matches and completions for each keyword as you type into the **Search** field. For syntax information, see Search Syntax Reference.

2. Click the **Choose Blok** icon located next to the **Search** field and select **Save as Blok**.
3. Update the information. For information about each field, see Creating a Blok.
4. Click **Save** to save the Blok as a new Blok.

Result
The new Blok is added in the **Choose Blok** list and is displayed in the **Search** field.

Rules

From **Management > Advanced Features > Rules**, an administrator can add, edit, or delete triggers and aggregation rules.

- Triggers: can be created after defining a correlation Blok. Triggers describe what action should be taken once a correlation Blok is triggered.
- Aggregation rules: can optimize the performance of aggregation search queries.

Managing Triggers

Triggers describe what action should be taken when a correlation Blok is triggered. If several triggers are associated with the same correlation Blok, all of them are triggered.

Adding a Trigger

You can add triggers from the **Management > Advanced Features > Rules > Triggers** tab.

1. Enter a name and description for the trigger. The name can include letters, numbers, or underscore (_).
2. Enable or disable the trigger.

   You can enable or disable triggers at any time, but they must be synchronized in

order to be activated.

3. Specify a trigger group, severity, category, and correlation Blok. The built-in group `System` is provided, but you can also create your own trigger groups. However, you cannot delete a trigger group from the GUI.
4. Set the maximum number of alerts and the frequency (hour, minute, day).
5. (Optional) Set up alert notifications to be sent out when the trigger is activated. Notifications can be in the form of email, syslog, or SNMP and you can set multiple notifications for a single alert.

   For SNMP notifications, if you select the SNMP version as v3, the **Snmp Community** field is not included in the log data that is displayed in the search results.

For more information about alerts, see Advanced Alerts. For information about managing triggers and alerts by using REST API, see REST API for Aggregation Rules.

Synchronizing Triggers

Any updates to the trigger settings requires synchronization for the changes to take effect. Click the **Sync triggers** icon  on the **Triggers** tab and select the trigger groups to be synchronized. The synchronization process deploys all enabled and disabled triggers in the selected groups to the correlation node. If you create a new trigger, you must synchronize the trigger group containing the new trigger by clicking the **Sync triggers** icon  .

Note: Even if you select a single trigger and click the **Sync triggers** icon  , the trigger group for that trigger is selected for synchronization and all triggers contained in that trigger group are synchronized. You can select multiple trigger groups to be synchronized.

Aggregation Rules

You can define aggregation rules to optimize the performance of aggregation (GROUP BY) queries.

After you create and enable aggregation rules, the system precomputes the aggregations as and when log events arrive in the system. The aggregations are used to compute search

query results. As time progresses, the aggregations are precomputed at real time, providing query results much faster than the queries that were not optimized. Such optimized aggregation queries can be vital for creating responsive dashboards.

Note: To use aggregation rules, Advanced Features and Advanced Aggregation must be enabled. Contact your administrator to enable the features.

An aggregation rule is defined as a regular EQL or SQL query that contains a GROUP BY statement and aggregated projections. For details, see GROUP BY Statement.

For example, consider the data model for BW process information. This data model has columns such as `Domain`, `AppSpace`, `Application`, `Process`, `Activity`, `exectionTime`, `success`, and `failure`.

Consider the aggregation query as:

```
USE BWProcesses | GROUP BY Domain, AppSpace,  Application, Process
COLUMNS avg(executionTime), max(failure), min(success)
```

For this search query, the aggregation rule maintains the aggregate metrics `avg(executionTime)`, `max(failure)`, `min(success)` for each combination of the `GROUP BY` fields: `Domain`, `AppSpace`, `Application`, and `Process`. The computation happens on every relevant log event ingested in the system at real time.

When you use an aggregation query for which there is a matching aggregation rule, the query is a candidate for optimization. If pre-computed metrics are already present, the results are returned without any query-time computation.

The time filters in the aggregation query must be within the optimized scope of the rule, that is, between the time specified in the **Compute aggregation for** field and the time of the query.

For example, suppose the rule was created at 10 am on Jan 1, with a retention period of 2 weeks. After 2 days, on Jan 3, if you query for aggregated data which matches the rule but with a time filter condition of -2w, the result is still unoptimized because 2 weeks from the time of search results in time range starting from Dec 19, but the aggregation rule was not created on Dec 19. However, if you specify the **Compute aggregation for** value as -2w while creating the aggregation rule, then the aggregation is computed from Dec 19 onwards.

Aggregation rules also support grouping by time aggregates. The supported time aggregations are:

Time function	Groups by
`years(sys_eventTime, [multiplier])`	Years
`months(sys_eventTime, [multiplier])`	Months
`weeks(sys_eventTime, [multiplier])`	Weeks
`days(sys_eventTime, [multiplier])`	Days
`hours(sys_eventTime, [multiplier])`	Hours
`minutes(sys_eventTime, [multiplier])`	Minutes
`seconds(sys_eventTime, [multiplier])`	Seconds

In these functions, the `multiplier` parameter is optional. For more information, see: Time functions.

For example:

```
USE BWProcesses | GROUP BY Domain, AppSpace,  Application, Process, days
(sys_eventTime), hours(sys_eventTime), minutes(sys_eventTime, 10)
COLUMNS avg(executionTime), max(failure), min(success)
```

This rule computes the metrics `avg(executionTime)`, `max(failure)`, `min(success)` for each combination of `Domain`, `AppSpace`, `Application`, and `Process` across every 10 minutes, every hour, and every day.

Such queries retrieve the time series data for trend analysis. For example, `avg (executionTime)`, `max(failure)`, `min(success)` for `AppSpace='appSpace1' AND Process='Process1' AND Domain='domain1'` aggregated for each hour for the day='Monday'.

This query can be used in the Advanced Dashboard to create a time series chart showing the trend, for example, average execution time for a process across hours of the day.

Tip: After an aggregation rule is created, a filter Blok is automatically created in the system. If there are multiple time clauses in the GROUP BY query, multiple filter Bloks are created - one for each time clause. To save on query response time, you can use the filter Blok in Advanced Dashboards instead of typing the entire query.

An aggregation query can be equal to or subset of another rule query if all of the following conditions are met:

- The USE or FROM clauses are equal
- The GROUP BY non-time aggregates are exactly equal
- The time aggregates are equal or subset
- The projection aggregates of the two queries are equal or subset

Examples

The following examples illustrate how a search query is considered to match an existing aggregation rule.

Consider the aggregation rule as:

```
USE LogLogic_Monitor_Cpu | GROUP BY ll_nodeId, weeks(sys_eventTime),
days(sys_eventTime), hours(sys_eventTime, 5)
COLUMNS max(ll_systemCPU), max(ll_processCPU), avg(ll_systemCPU), avg
(ll_processCPU)
```

Now consider the following search queries:

Query1

```
USE LogLogic_Monitor_Cpu | GROUP BY ll_nodeId, weeks(sys_eventTime)
COLUMNS max(ll_systemCPU),avg(ll_processCPU)
```

Query1 matches the aggregation rule because all conditions match:

- The USE clause in Query1 is the same as that of the aggregation rule.

- The GROUP BY non-time aggregate `(ll_nodeId)` in Query1 is exactly the same as that in the aggregation rule.
- The time aggregate column `weeks(sys_eventTime)` is a subset of that in the aggregation rule. So, the GROUP BY clause of Query1 is a subset of the GROUP BY clause of the aggregation rule.
- The projection aggregates in Query1, `max(ll_systemCPU), avg(ll_processCPU)`, are a subset of that in the aggregation rule.

Query2

```
USE LogLogic_Monitor_Cpu, Loglogic_Appliance | GROUP BY ll_nodeId,
weeks(sys_eventTime) COLUMNS max(ll_systemCPU), avg(ll_processCPU)
```

Query2 does **not** match the aggregation rule because:

The USE clause in Query2 is not same as that of the aggregation rule; it has an extra data model `Loglogic_Appliance`.

Query3

```
USE LogLogic_Monitor_Cpu | GROUP BY ll_nodeId, ll_pRuleID, weeks(sys_
eventTime) COLUMNS max(ll_systemCPU), avg(ll_processCPU)
```

Query3 does **not** match the aggregation rule because:

The GROUP BY non-time aggregate in Query3 is not exactly the same as that in the aggregation rule; it has an extra column `ll_pRuleID`.

Query4

```
USE LogLogic_Monitor_Cpu | GROUP BY ll_nodeId, minutes(sys_eventTime)
COLUMNS max(ll_systemCPU), avg(ll_processCPU)
```

Query4 does **not** match the aggregation rule because:

The time GROUP BY in Query4 is not equal to or a subset of that in the aggregation rule; it has a different column `minutes(sys_eventTime)`.

Query5

```
USE LogLogic_Monitor_Cpu | GROUP BY ll_nodeId, minutes(sys_eventTime)
COLUMNS max(ll_systemCPU), sum(ll_processCPU)
```

Query5 does **not** match the aggregation rule because:

The projection columns in Query5 are not equal to or a subset of that in the aggregation rule; it has a different column `sum(ll_processCPU)`.

Limitations

You cannot create precomputed aggregation sets for infrastructure queries.

Aggregation Rule Groups

The **Management > Advanced Features > Rules > Aggregation** page displays rules in groups. The **All Rules** group displays all aggregation rules.

For information about groups, types of groups, and managing groups, see the Groups section.

Viewing and Searching Aggregation Rules

From the **Management > Advanced Features > Rules > Aggregation** page, you can view, find, or filter rules in the following ways:

Task	Steps
View rules based on filters	You can use filters to easily find rules in the system. Click the **View** list to select the required filter. The following options are available: • All • Created by me • Created by system • Imported

Task	Steps
Find rules	You can quickly find the desired rule by typing the rule name in the **Find** field. As you start typing a rule name in the **Find** field, the Aggregation page is automatically refreshed showing your selection.
Sort rules	You can sort any column in ascending or descending order on the Aggregation page. Click the column name or click the arrow (that is displayed on the right side of the column name when you click in that column) to sort the column.
Show or hide columns	You can show or hide columns, except the mandatory column, from the table. Click the column picker icon to view all available columns in the table. Select the check box to show the column. Clear the check box to hide the column from the table. The Aggregation page is updated immediately

Managing Aggregation Rules

You can view all aggregation rules, add a new rule, edit existing rules, or delete rules from the system.

You can manage rules from the **Management > Advanced Features > Rules > Aggregation** page.

Note: The duplicate , delete , and move icons are enabled after you select one or more rules.

Task	Steps
View and search rules	See Viewing and Searching Aggregation Rules.
Create a new rule	See Creating an Aggregation Rule.
Edit a rule	Clicking a rule displays the Details panel, in which more information about the query is displayed. You can click **Edit** to modify the information. See Editing an

Task	Steps
	Aggregation Rule. **Note:** You cannot edit built-in and system-created rules. However, you can duplicate the built-in or system-created rule, save it in the **User** group or a user-created group, and then edit the information in the duplicated rule.
Duplicate	Select the rule and click the duplicate icon  . Modify the query and then save the new rule to the **User** group or a user-created group.
Delete	Select the rule and click the delete icon  . After you delete a rule, it cannot be recovered. When an aggregation rule is deleted from the system, the aggregated data for that rule is deleted and the query is not optimized. Similarly, the corresponding filter Blok is also deleted from the system.
Move	Select the rule and click the move icon  . Select the group to which you want to move the rule.
Rename	You can rename groups but not rules.
Enable or disable	Aggregation rules can be disabled. After a rule is disabled, it stops computing aggregations and all the existing aggregated data that was precomputed for that rule is deleted. The optimization starts only after the rule is enabled. To enable or disable a rule, click **Enabled**.

Creating an Aggregation Rule

When a new aggregation rule is created, it starts computing the metrics for the given GROUP BY expressions. It aggregates data for the events starting from the aggregation start period up to the rule creation time. When an aggregation query matches an aggregation rule, the query results are fetched from the collected aggregation data. After an aggregation rule is created, a filter Blok is automatically created in the system. If there are multiple time clauses in the GROUP BY query, multiple filter Bloks are created - one for each time clause.

Attention:

- When creating aggregations for small periods, for example using GROUP BY minutes

or seconds, use the data type as `long` for the columns in your custom data model.

- When using your own data model to create aggregation rules, if you edit any column types of the data model after creating an aggregation rule, then for the changes to take effect you must either create a new aggregation rule again, or edit and save the existing aggregation rule.

> **Note:** Editing an aggregation rule causes all existing data to be purged, and the rule starts computing the aggregation using the updated query.

Procedure

1. Go to **Management > Advanced Features > Rules** and click the **Aggregation** tab.
2. Click **Create New Rule**.
3. In the **Add Aggregation Rule** dialog box, provide the following information:
    a. Parent Group: Select a parent group where you want to save the rule.

    You can create a new group or select the **User** group, or select any user-created group

    Default parent group: When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.

    b. Name: The name can include letters, numbers, hyphen, or underscore (_). The rule name cannot start with a number.

    c. (Optional) Description of the rule.

    d. Query: Ensure that you enter a valid syntax of a search query. An EQL or SQL query that contains GROUP BY statement and aggregated projections are supported. For details, see GROUP BY Statement. Note the following exceptions:

    - Time functions in the GROUP BY statement can only have `sys_eventTime` as an argument. Any other timestamp column is not supported.
    - Restrictions on GROUP BY statement:
        - GROUP BY statement can include both column names or expressions. However, only time functions are supported, for example, `days(sys_eventTime)`, `weeks(sys_eventTime)`. Generic

expressions such as `length(Process)/10` are not supported.

- Special characters in column names should be escaped, for example, `[Process Name]` or `[% Memory]`

- Restrictions on the projection clause (COLUMNS):
  - The projection items must contain aggregate functions such as `sum()`, `avg()`.
  - Projection items can be a complex expression, for example, `concatenate(max(cpu), '_vs_', min(cpu))`, but it should contain aggregation functions. An aggregation function can contain only simple functions as arguments. For example, `max(Memory)` is supported but `max(sqrt(Memory))` is not.

Click **Validate** to verify the query statement.

e. Select or clear the appropriate **Aggregation time** check boxes to add or remove time functions in GROUP BY clause of the query.

When you select this option, the time aggregation you selected is inserted in the GROUP BY clause of the rule query. The default time aggregation functions are:

- weeks(sys_eventTime)
- days(sys_eventTime)
- hours(sys_eventTime)
- minutes(sys_eventTime)

f. In the **Retention period** field, enter the retention time for which the computed aggregation values remain stored in the aggregated data.

You can search based on the same aggregation functions until the specified time has passed. By default, it is set to -1w (1 week).

For example, if the retention period is -2w (2 weeks), then pre-computed results remain in the system for 2 weeks since the time of rule creation.

g. In the **Compute aggregation for** field, specify the time when the aggregation computing must start. A one-time calculation of aggregation happens on the data that was already collected prior to the creation of the aggregation rule. To

start computing as soon as the rule is created, leave the field empty. To compute aggregation for the data that exists prior to rule creation, specify the time relative to rule creation time. For example, -1d, -1w, and so on.

> **Note:**
> - The **Compute aggregation for** period cannot be more than the retention period.
> - To modify the **Compute aggregation for** period, first modify the retention period and then the **Compute aggregation for** period.

The data model `Loglogic_Aggregation_Rules_Metrics` provides a disk usage statistics of all aggregation rules in the system. You can run the infrastructure query Aggregation Rule Metrics and view the results.

h. In the **Maximum Aggregated Data Storage Size in MB** field, enter the maximum storage size of aggregated data. The minimum value is 1 MB, the maximum value is 2147483647, and the default is 1024 MB if the field is empty or 0. You can specify the limit based on the storage availability in your setup.

   Data is purged daily, based on the retention period and the maximum storage values. Purging is triggered 24 hours after rule creation time. When one of these values is exceeded, the oldest 24 hours data is purged to match the values that you specified in the aggregation rule. For example, with the retention period as -1w and maximum storage size as 1024 MB, suppose that the aggregation data on the first day is 1000 MB and on the second day, is 2024 MB. On the second day, the oldest 24 hours data (in this example, 1000 MB data of the first day) is purged to match the storage size limit of 1024 MB.

i. To activate the aggregation rule, set **Enabled** to **ON**.

   The optimization starts only after the rule is enabled. The disabled rule does not compute real-time events.

4. Click **Save**.

Result

The newly added rule is displayed in the All Rules group and in the parent group that you selected.

Saving an Advanced Search Query as an Aggregation Rule

If the query is not already included in any saved aggregation rule, then the Optimize icon is displayed.

Procedure

1. On the Advanced Search page, click the optimize icon  to open the **Add Aggregation Rule** dialog box.

   The query that you typed on the Advanced Search page is automatically filled in the **Query** field.

2. Enter the other fields in the **Add Aggregation Rule** dialog box and then click **Save**.

   The aggregation rule is saved in the system.

Editing an Aggregation Rule

When the existing aggregation rule is updated in the system, the aggregated data is reset and computation starts again from the updated time, and the corresponding filter Blok is also updated automatically.

Important Considerations

- If you update any of the fields, the aggregation start time is updated on the Aggregation page.
- If you update the description or re-enable a disabled rule, the precomputed aggregation data before the updated rule start time is not purged from the system until the retention time elapses.
- If you update any other field, the existing precomputed data for the given aggregation rule is purged and the data is recomputed starting from the time specified in the **Compute aggregation for** field, or from the current time if this field is empty.

Procedure

1. Go to **Management > Advanced Features > Rules** and click the **Aggregation** tab.
2. On the Aggregation page, click the rule name that you want to update.

The **Details** panel opens on the right side of the page.

3. Click **Edit** and make the necessary updates.

   For more information about fields, see the steps in the Adding an Aggregation Rule section.

4. Click **Save** to save the updated information.

Result
The updated rule is displayed on the Aggregation page.

Viewing the Statistics of Aggregation Rule Computations

When a new aggregation rule is created, it starts computing the metrics for the given GROUP BY expressions. It aggregates data for the events starting from the aggregation start period up to the rule creation time. To view the progress of the aggregation computation, you can run the following infrastructure query in the Advanced Search:

```
USE Loglogic_Aggregation_Rules_Metrics
```

For more information about the infrastructure query, see Aggregation Rule Metrics.

Distributed Aggregation Rules

Starting from LogLogic LMI 6.3.0, you can create an aggregation rule on a Management Station and distribute the rule to multiple Remote Appliances in the setup.

Using the distributed aggregation rule, you can run an optimized search query on the Management Station, which in turn runs on the specified Remote Appliances. The results from the specified Remote Appliances are consolidated by the Management Station.

Requirements to use distributed aggregation rules

You can use distributed aggregation rules in Advanced Search if all of the following requirements are met:

- The Management Station setup includes the required Remote Appliances. For

information about how to set up a Management Station, see "Manage Appliances with Management Station" in the *TIBCO LogLogic® Log Management Intelligence Administration*.

- You have permission to access the Remote Appliances and devices on which the search query is to be run.
- Advanced Features are enabled on the Management Station and Remote Appliances.
- Advanced Aggregation is enabled on the Management Station and Remote Appliances.

Important Considerations

- After the aggregation rule is created and distributed, the rule can be modified or deleted only from the Management Station GUI. The rule cannot be modified or deleted from the GUI of Remote Appliances.
- After the aggregation rule is created and distributed, if you change the Management Station setup (add or remove Remote Appliances), the rule is redistributed to only those appliances to which the rule creator user has access.

Creating distributed aggregation rules

You create an aggregation rule on the Management Station and specify the list of appliances to which the rule is to be distributed. The rule is then distributed to the specified appliances. You specify the appliances using the appropriate values in `sys_concentratorId` column or the `DeviceInGroup` function.

Modifying distributed aggregation rules

If you modify a distributed aggregation rule on the Management Station, the rule is also updated on the appliances where the rule had been distributed at rule creation time. You can modify other fields and parameters in the rule except the list of appliances. If you want to modify the list of appliances in the distributed aggregation rule, you must delete the existing rule and create a new one for redistribution.

Deleting distributed aggregation rules

If you delete a distributed aggregation rule on the Management Station, the rule is also deleted on the Remote Appliances. However, if the appliances registered in the Management Station are modified after rule creation time, then the rule is deleted from only those Remote Appliances that are members of the Management Station at deletion time. When a rule is deleted, all aggregated data created by that rule is also deleted.

Searching using the distributed aggregation rule

If you run an optimized Advanced Search query on a Management Station, the query is run on all appliances specified in the aggregation rule, and the results are consolidated by the Management Station. The query is run on only those appliances that are accessible at run time, and search results from only those appliances are consolidated by the Management Station.

You can run the optimized distributed query by specifying the exact rule query or by using a filter Blok created for the rule.

Example using a query:

```
USE LogLogic_Appliance | GROUP BY ll_pRuleID | sys_concentratorId =
'ALL'
```

Example using a filter Blok: If you save the query as a filter Blok `LogLogicApplianceBlok`, you can search using the filter Blok:

```
filter.LogLogicApplianceBlok
```

For more information about distributed Advanced Search and how to specify the list of appliances, see Distributed Advanced Search.

Limitations and errors

Limitations

- You cannot use the Tail functionality in distributed aggregation rules.
- You cannot use Enrichment Lists as a value in the `sys_concentratorId` parameter

and in the `DeviceInGroup` function in distributed aggregation rules.

Errors

An error is displayed if:

- The specified appliance shortcut is invalid.
- An appliance specified in the distributed query is not a member of the Management Station setup.

Advanced Alerts

Alerts are generated when real-time events match a correlation Blok that has an active trigger linked to it. For information on how to define triggers, see Managing Triggers.

Alerts can be distributed by email to a pre-defined list of people. To receive email notifications, the SMTP connection must be configured. To configure the SMTP connection, contact your administrator.

The default retention period for all generated alerts is 90 days.

Note: Advanced alerts work only if the Advanced Features are enabled. To enable Advanced Features, contact your administrator.

Viewing and Searching Advanced Alerts

From the **Alerts > Advanced Alerts** page, you can view, find, or filter alerts in the following ways:

Task	Steps
View alert details	See Viewing Alert Details.
View alerts based on	You can use filters to easily find alerts. Click the **View** list to select the required filter. The following options are available:

Task	Steps
filters	• All - all alerts in the system • Acknowledged - alerts that have been acknowledged • High Severity - alerts with high severity • Unacknowledged - alerts that have not been acknowledged
Find alerts	You can quickly find the desired alert by typing the alert name in the **Find** field. As you start typing the alert name in the **Filter** field, the Alerts page is automatically refreshed showing your selection.
Sort alerts	You can sort any column in ascending or descending order. To sort by a column, click on the column name or the arrow next to the column name.
Show or hide columns	You can show or hide columns, except the mandatory column, from the table. Click the column picker icon to view all available columns in the table. Select the check box to show the column. Clear the check box to hide the column from the table. The Alerts page is updated immediately.

The following table describes the **Alerts** information displayed on the Alerts page:

Column	Description
Severity	The severity of the trigger. The options are: • Info • Low • Medium • High **Note:** An admin (a user with administrator privileges) can configure severity options. The options might differ if they have been configured.
SLA Expiration	The Service Level Agreement (SLA) expiration time is the time by which an operator is expected to acknowledge the alert. When the

Column	Description
	SLA time expires, it displays the time in negative hours or days in this column field.
Status	The icon indicates the alert status:  expired  acknowledged  unacknowledged
Acknowledged	A check mark  indicates that the alert is acknowledged. Otherwise this field is blank.
Name	The trigger name associated with the alert.
Trigger Group	The group to which the trigger belongs.
Description	The description of the alert.
Category	The category of the trigger. The options are: • Attack on third party • Authorized Activity • Authorized security testing • Emergency changes • False positive • Known error • LogLogic Event • Network Noise • Security Alert • Suspicious Activity • Unauthorized Activity

Column	Description
	• Unknown **Note:** An admin (a user with administrator privileges) can configure the category options. The options might differ if they have been configured.
Elapsed time	The time since the alert was created
Last updated	The time when the alert was last updated

Limitation

If you use a correlation Blok in an Advanced Search and the result includes 'null' values, the triggered alerts does not include the items including 'null' values.

Managing Advanced Alerts

From the **Alerts > Advanced Alerts** page, you can view all triggered alerts, and acknowledge or filter them. From the Alerts page, you can perform the following tasks:

Task	Steps
Acknowledge alerts	Acknowledging an alert indicates that you have recognized the alert. Once you acknowledge the alert, your user name gets associated with that alert. For instructions on how to acknowledge alerts, see Acknowledging Alerts.
View and search alerts	See Viewing and Searching Advanced Alerts.
View alert details	See Viewing Alert Details.
Auto-refresh the list of alerts	Click the down arrow next to the refresh button to set the refresh

Task	Steps
	interval in seconds. Enter the time in seconds. The **Alerts** table is refreshed as per the defined time interval. By default, it is refreshed every 30 seconds. Clicking the **Pause** button halts refreshing, and the **Pause** button changes to **Resume**. Clicking **Resume** resumes refreshing the list of alerts.
Delete alerts	Select alerts from the list and delete them by clicking **Delete**

Limitation

If you use a correlation Blok in an Advanced Search and the result includes 'null' values, the triggered alerts does not include the items including 'null' values.

Acknowledging Alerts

Acknowledging an alert indicates that you have received and recognized the alert. Once you acknowledge the alert, your name is associated with that alert.

Procedure

1. Navigate to **Alerts > Advanced Alerts**.
2. Select the check box for the alerts you want to acknowledge. To select all alerts, select the check box located to the left of the column headings.
3. Click **Acknowledge** to acknowledge the selected alerts.
4. Enter the following information:

Field	Description
Severity	Alert severity
Category	Alert category
URL	Web address of an external web page such as a customer relationship

Field	Description
	management system or a defect tracking system. This enables you to track actions triggered by the alert.
Comments (Optional)	Comments about the alert

5. Click **Acknowledge** to acknowledge alerts.

Result

The **Alerts** table shows a  icon in the **Status** column and a checkmark  in the **Acknowledged** column for the acknowledged alerts.

Viewing Alert Details

You can view details of any generated alert.

Procedure

1. Navigate to **Alerts > Advanced Alerts**.
2. From the Alerts page, click the alert name to view its details.

   In the Details window, you can view alert details, history, associated correlation rule, and event group details. An event group count is displayed with a clickable number. For more information, see Viewing Event Group Details.

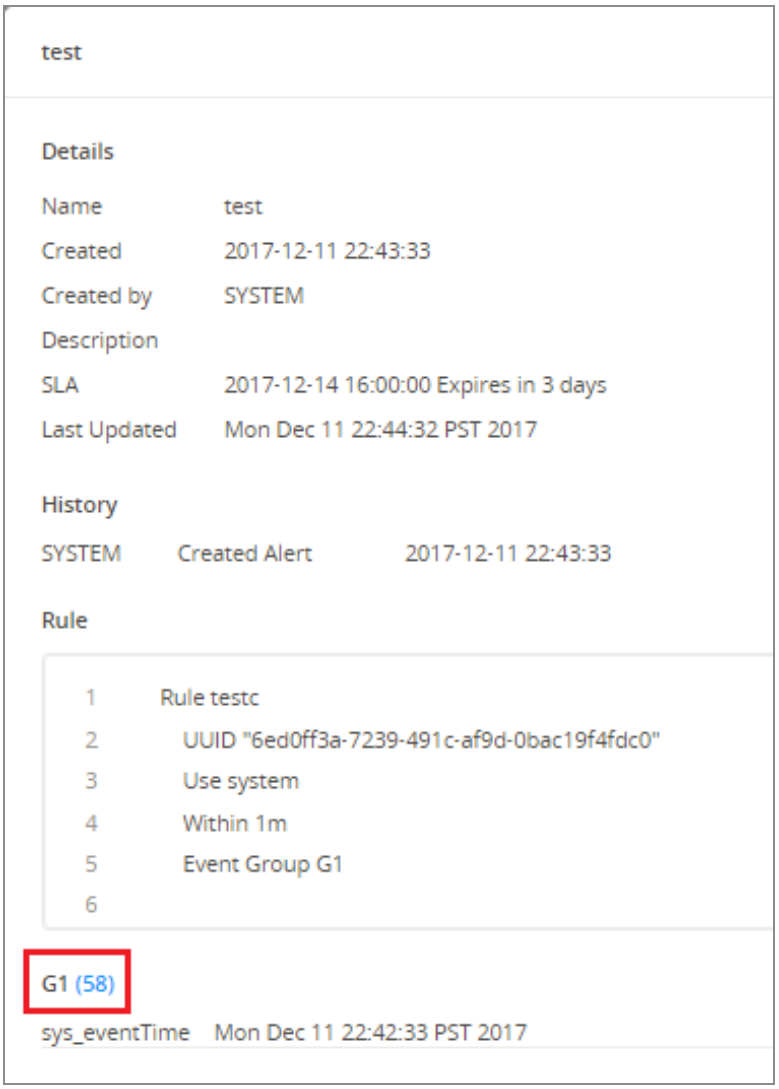

3. To acknowledge the alert, click **Acknowledge**. For details, see Acknowledging Alerts.

Viewing Event Group Details

Each event group describes the criteria that must combine events to be grouped together as part of the correlation rule. This is equivalent to a single search query defined in EQL.

Procedure

1. Navigate to **Alerts > Advanced Alerts**.
2. From the Alerts page, click the alert name to view its details.

In the Details window, you can view alert details, history, associated correlation rule, and event group details.

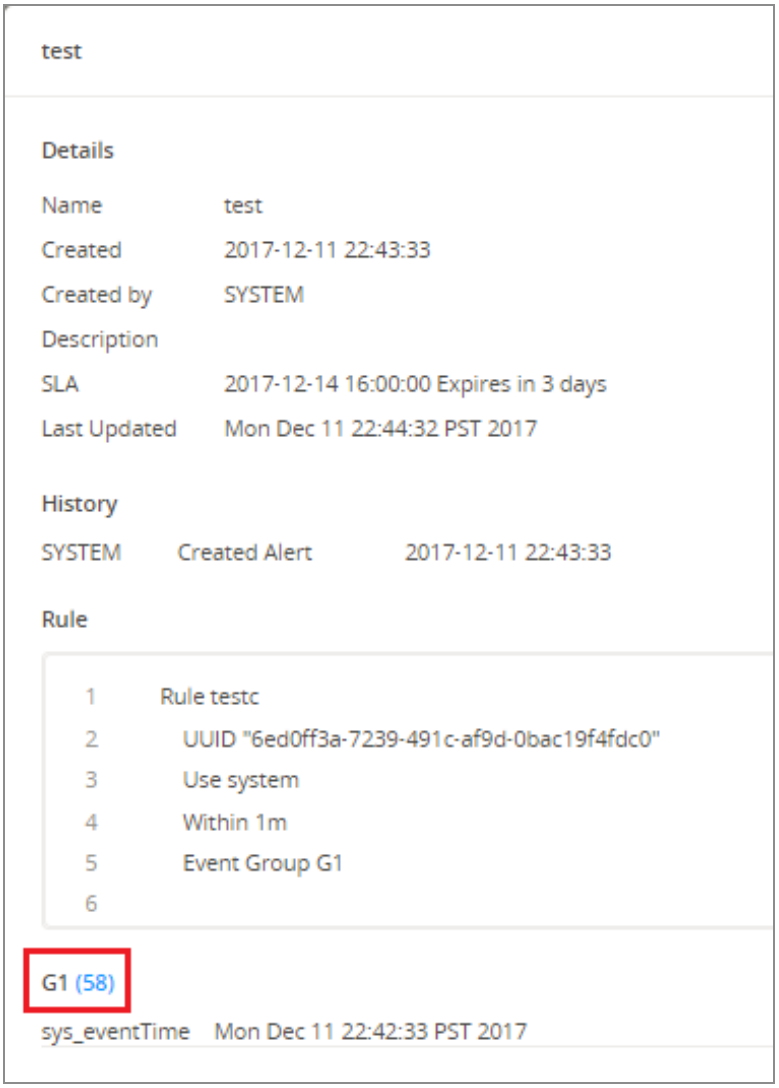

3. To view the associated event count query, click the event group **count** link, for example, **(58)** as shown in the example.

> **Note:** The event count link is available only when the count is less than 1024.

A new search tab is added showing the event count query in the **Search** field. The **Result** tab displays the retrieved results in the **Timeline Charts**, **Columns**, and **Data**

panels.

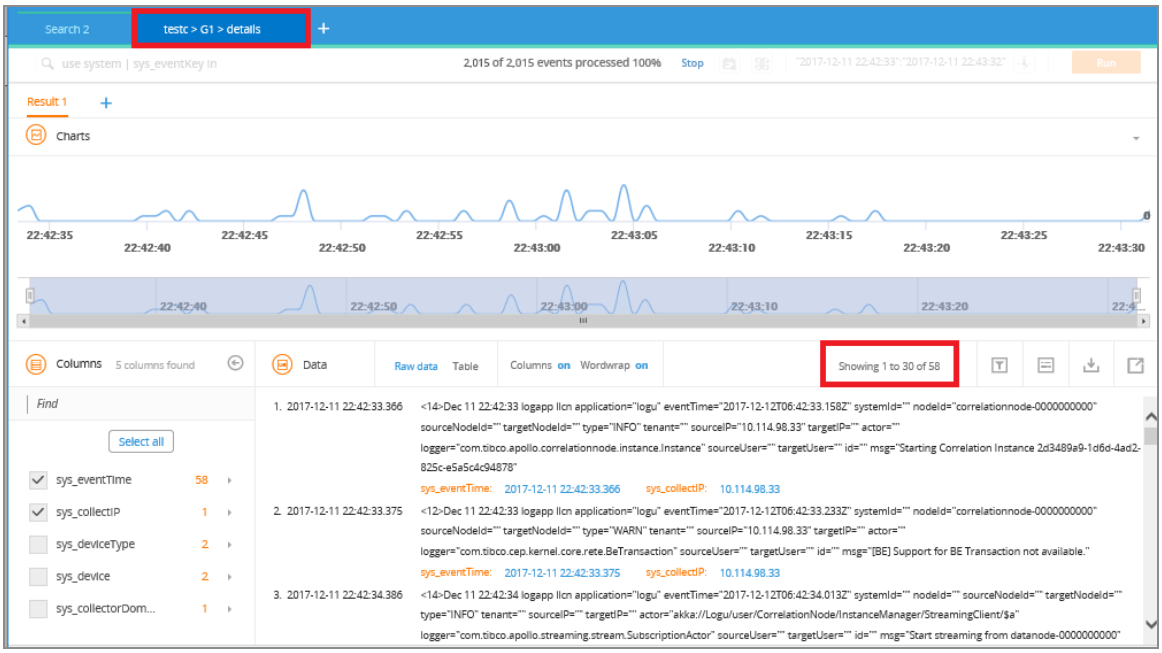

Advanced Dashboards

An Advanced Dashboard contains a collection of data widgets that provide a graphical representation such as a chart, a bar, or a count.

The use of dashboards is endless. For example, as an IT administrator, you can focus on all machines in your enterprise by creating a widget in a dashboard. Dashboards can be built as per your specifications. You can add multiple widgets in a dashboard.

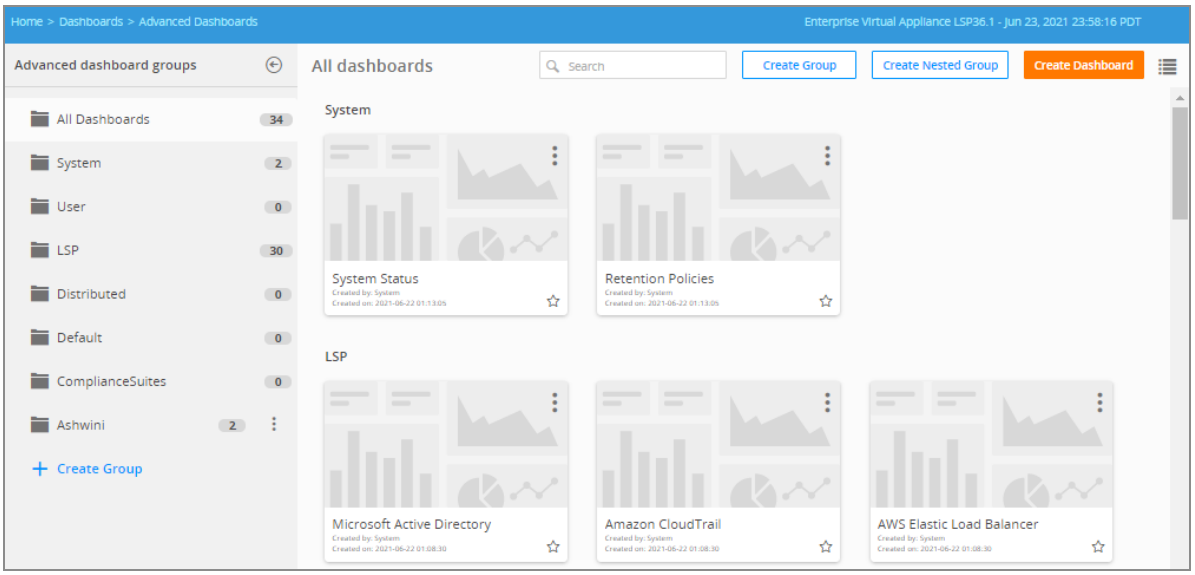

Dashboard Groups

The **Dashboards > Advanced Dashboards** page displays dashboards in groups. The **All Dashboards** group displays all dashboards. The System built-in group includes the following built-in dashboards:

- System Status
- Retention Policies

For information about groups, types of groups, and managing groups, see the Groups section.

Managing Advanced Dashboards

You can view all dashboards, add a new dashboard, copy or move an existing dashboard, or delete any dashboard in the system.

The Advanced Dashboards page displays all dashboards existing in the system. From the Advanced Dashboard page, you can perform the following tasks:

Note: You cannot move or delete built-in dashboards.

Task	In grid view...	In list view...
Create a new dashboard	1. Click **Create Dashboard**. 2. Enter a name for the dashboard. 3. Parent Group: Select a parent group where you want to save the dashboard. You can create a new group or select the **User** group, or select any user-created group **Default parent group:** When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.	
• Move • Delete	**To move or delete a single dashboard:** Click the context menu on the top-right corner and then click the icon for the required action. **To move or delete multiple dashboards:** Select the dashboards by clicking the top-left corner of the dashboard card. From the pane at the bottom of the screen, click the icon for the required action. **Note:** You cannot move or delete built-in and system-created dashboards.	Select one or more dashboards and then click the move or delete icons at the top of the list.
• Mark as favorite • Duplicate	**To duplicate a single dashboard or mark as favorite:** Click the context menu on the top-right corner and then click the icon for the required action.	Select one or more dashboards and then click the favorite or duplicate icons at the top of the list.

Task	In grid view...	In list view...
	To duplicate multiple dashboards or mark as favorite: Select the dashboards by clicking the top-left corner of the dashboard card. From the pane at the bottom of the screen, click the icon for the required action.	
View dashboards based on filters	Not applicable	You can use filters to easily find dashboards. Click the **View** list to select the required filter. The following options are available: • All - all dashboards in the system • Favorite - dashboards marked as favorite
Find dashboards	You can quickly find the desired dashboard by typing the dashboard name in the Search field. As you start typing a dashboard name in the **Search** field, the dashboard page is automatically refreshed showing your selection.	Search: Same as in grid view. Filter: Select a filter from the **View** drop-down list. The filter is not available in the grid view.
Sort dashboards	Not applicable	You can sort the list of dashboard groups by any column. Click the column name or click the arrow (that is displayed on the right side of the column name when you click in that column) to sort by the column.
Show or hide columns	Not applicable	You can select the columns to be used in the list. Click  to view all available columns in the table.

Task	In grid view...	In list view...
		Select the check box to show the column. Clear the check box to hide the column from the table. The Dashboard page is updated immediately.

System Group of Dashboards

The built-in System group includes the following dashboards for quick reference:

- System Status dashboard
- Retention Policies dashboard

System Status

The System Status dashboard provides an overview of the current status of the system. You can view this dashboard on the **Dashboards > Advanced Dashboards > Advanced System Status** page.

This dashboard displays the following widgets:

- Widgets displaying availability of disk space and archive storage:
    - Free Disk Space
    - Used Disk Space
    - Total Disk Space
    - Tables: stDataFile, indexFiles, metaIndexFiles - Size in MB
    - Remote Archive Storage
    - Estimated time to reach maximum disk usage
- Line chart widgets depicting CPU usage:
    - CPU Usage - 1 Minute
    - CPU Usage - 10 Minutes

    - CPU Usage - 15 Minutes

  On these charts, the X-axis displays the time and the Y-axis displays the CPU load.

- Line chart widgets depicting the message rate:
    - Message Rate - 1 Minute
    - Message Rate - 5 Minutes
    - Message Rate - 15 Minutes
    - Message Rate - 30 Days

  On these charts, the X-axis displays the time and the Y-axis displays the message rate.

- Bar widgets depicting the ingested and indexed data over the last 24 hours:
    - Ingested and Indexed Message Count - Last 24 hours
    - Ingested and Indexed Bytes - Last 24 hours
    - Real Time - Message Count - Last 24 hours
    - File Pull - Message Count - Last 24 hours
    - Ingested vs Indexed Messages in MB - Last 24 hours
    - Ingested Messages in MB - Last 24 hours
    - Indexed Messages in MB - Last 24 hours
    - Indexed vs Unindexed Messages in MB - Last 24 hours

  These bar widgets use two different colors for the message count or message bytes. On this chart, the X-axis displays time and the Y-axis displays the message count or message bytes.

- Widgets displaying status of rules:
    - Message Routing Rules
    - Forwarder Rules
    - File Transfer Rules
- Widgets displaying infrastructure information:
    - IO Statistics
    - Status of LogLogic LMI engines

Retention Policies

To view this dashboard, go to **Dashboards > Advanced Dashboards > Retention Policies**. The Retention Policies dashboard includes the following widgets:

- Number widgets showing default policy for raw message retention and default index retention in days:
    - Default Period of Raw Retention
    - Default Period of Index Retention
- Current Values of Data Retention Rules: An inverted bar chart that displays the number of retention in days for each rule. The X-axis displays the retention value in days and the Y-axis displays rule names.

Widgets

Widgets provide a graphical view of the data corresponding to a query. An advanced dashboard can include different types of widgets.

Some widgets are built-in and available on the built-in dashboards. See System Group of Dashboards. You can create more widgets for your use.

Managing Widgets

On an advanced dashboard page, you can perform the following widget tasks:

Task	Description
Add a new widget	Click a widget type from the Widgets pane to add a new widget. For instructions, see Adding Widgets to an Advanced Dashboard.
Configure or edit a widget	Click to update the configuration. You can configure, edit, delete, move, or duplicate the widget by clicking the appropriate icon. 

Task	Description
View a widget in full-screen mode	Click to view the widget in full-screen mode. Click to exit full-screen mode. In full-screen mode, you can open the View data panel to view the raw data displayed on the widget.
View chart details	Hover your mouse over a certain area of the chart to view the details.
View value details	Hover your mouse over the widget to view the value at that point. Clicking the value opens the search results of that value on an Advanced Search tab. However, if the widget query includes aliases (AS statement), then the Advanced Search tab displays search results for the entire query.
Resize a widget	Grab any corner of the widget and resize as required.

Widget Types

- Line Widget
- Bar Widget
- Bubble Chart Widget
- Pie Widget
- Number Widget
- Gauge Widget
- Stacked Column Widget
- Combined Widget
- Treemap Widget
- Heat Map Widget
- DataGrid Widget
- Range Bar
- Geomap Widget

Line Widget

This widget is used to show the distribution of the total count of one selected column over its distinct values.

Use the following information to configure the widget:

- Fetch data from source
- Line widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Line widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Line widget configuration	
X-axis data	Select the column name to define the X-axis.
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data	Select the two columns to define the Y-axis of the widget.

Field	Description
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Categorize by	Define the column name by which the Y-axis data is combined into a series.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
use LogLogic_Monitor_Memory | COLUMNS sys_eventTime, (ll_memTotal-ll_
memFree) as memUsed, ll_memTotal as memTotal
```

the X-axis is `sys_eventTime`, and the Y-axis is `memUsed`, `memTotal`.

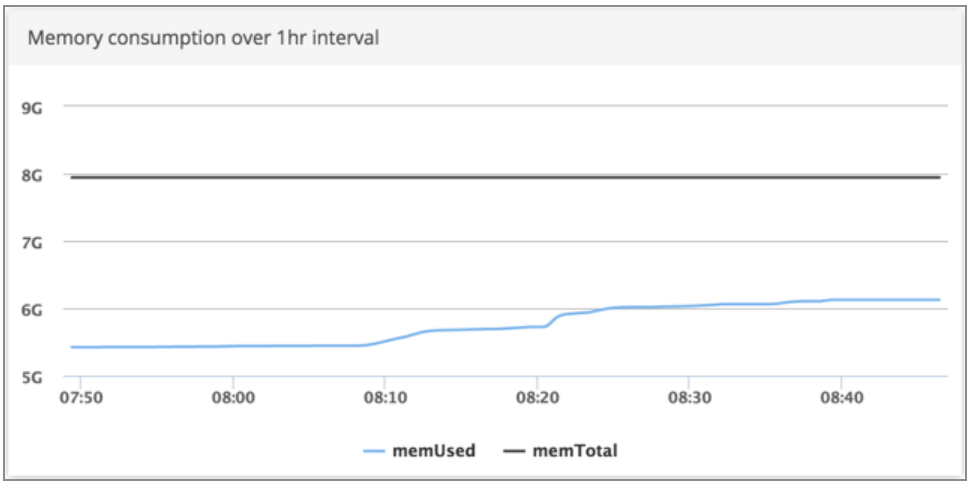

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Bar Widget

This widget is used to show the distribution of the total count of one selected column over its distinct values.

Use the following information to configure the widget:

- Fetch data from source
- Bar widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field

Field	Description
	options in the **Bar widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Bar widget configuration	
X-axis data	Select the column name to define the X-axis.
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data	Select the two columns to define the Y-axis of the widget.
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Show legends	Select the check box to display legends on the chart.
Show inverted	Select the check box to invert the X-axis and Y-axis values.
Categorize by	Define the column name by which the Y-axis data is combined into a series.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data.

Field	Description
	Default: • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
use LogLogic_Monitor_Memory | COLUMNS sys_eventTime, (ll_memTotal-ll_
memFree) as memUsed, ll_memTotal as memTotal
```

the X-axis is `sys_eventTime`, and the Y-axis is `memUsed`, `memTotal`.

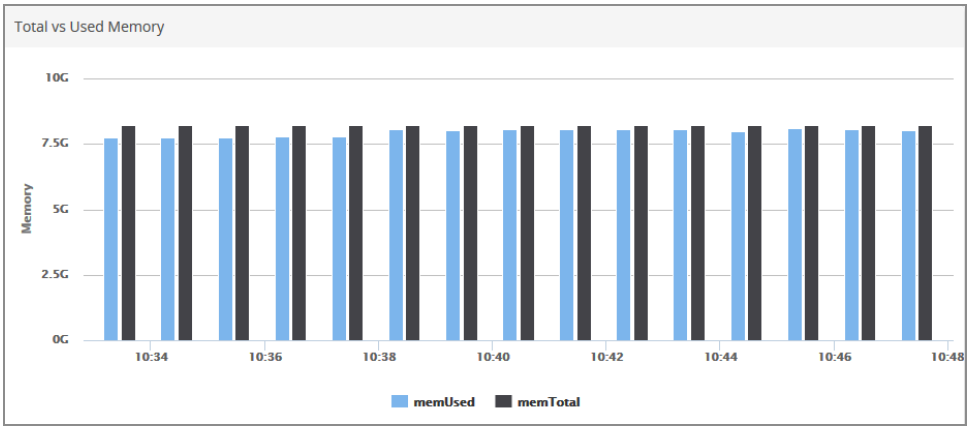

Related Topics

- Widgets

- Adding Widgets to an Advanced Dashboard

Bubble Widget

This widget displays data as a bubble chart or scatter plot. In a bubble chart, three dimensions of data are displayed - data from two columns is displayed on the x- and y-axes, and data from the third column is displayed through a bubble. In a scatter plot, two dimensions of data are displayed, and a third dimension can be displayed by coloring the points with different colors.

Use the following information to configure the widget:

- Fetch data from source
- Bubble widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Bubble widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Bubble widget configuration	
Bubble type	• To display a bubble chart, select Bubble.

Field	Description
	• To display a scatter plot, select Scatter.
X-axis data	Select the column name to define the X-axis.
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data	Select the two columns to define the Y-axis of the widget.
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Bubble weight	Select the column to be displayed as bubbles. The size of the bubble is proportionate to the column value.
Buffer size	Enter the number of rows to be displayed in the widget after refreshing the widget. For example, if the buffer size is 100 and the query returns 150 results, the latest 100 results are displayed. **Default:** 1000 rows
Use bubble color	Select the color of the bubbles or scatter points on the widget. Click the color value to change the color.
Use color threshold	Define the threshold range for the colors on the widget.
Use color value	Define the column name by selecting the column. You can use an Enrichment List or EQL conditional functions such as `IIF` in the query to return specific color values.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced

Field	Description
	Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example of bubble chart

For the search query:

```
USE General_Syslog | GROUP BY sys_collectIP , sys_eventTime
```

the X-axis is `sys_collectIP`, Y-axis is `sys_eventTime`, and Bubble value is `count(*)`.

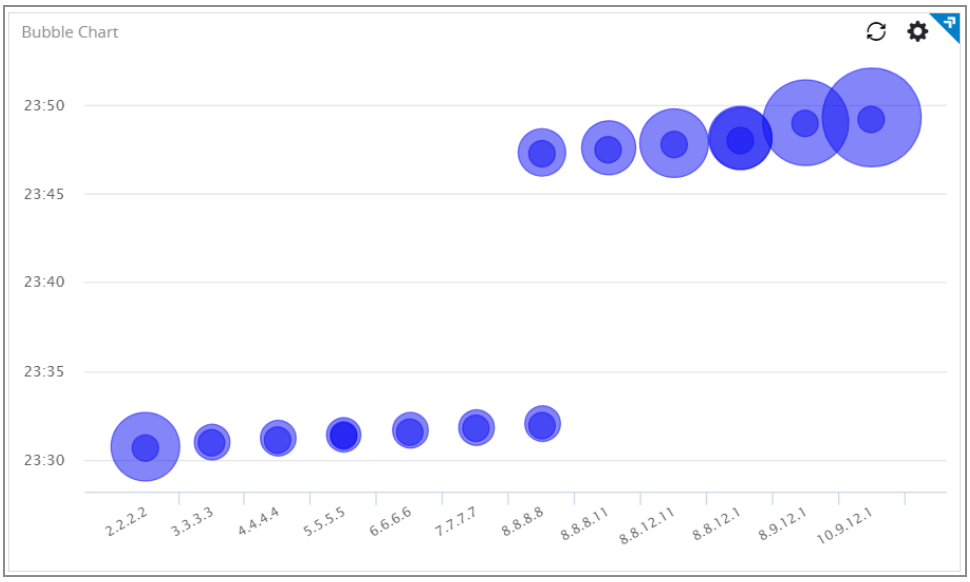

Example of scatter plot

For the search query:

```
USE General_Syslog | GROUP BY sys_collectIP, sys_deviceType
```

the X-axis is `sys_collectIP` and Y-axis is `count(*)`.

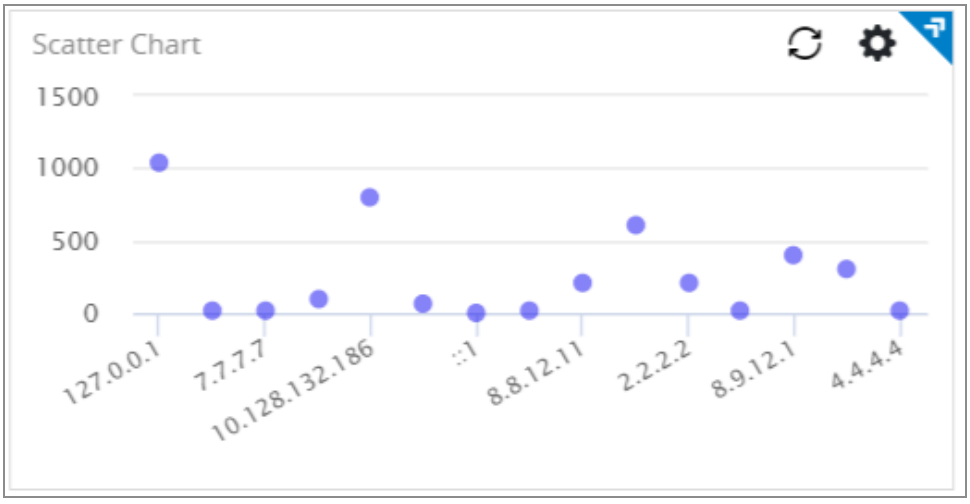

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Pie Widget

This widget uses one column at a time.

Each pie slice represents a distinct column value. The Pie widget data varies based on the selected column. Values that are not displayed in the specified the number of slices are grouped together into the 'Others' slice.

Use the following information to configure the widget:

- Fetch data from source
- Pie widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Pie widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Pie widget configuration	
Slice name	Select the column name to define the slice of the pie. If the column name is already defined in the search query, the Slice name column is automatically filled. Otherwise, as you start typing in the field, the available matching column names are displayed and you can select

Field	Description
	the appropriate one.
Slice value	Enter the slice value of the pie.
Show up to	Enter the number of slices to be displayed on the pie. If the slices are more than this number, they are combined and shown as 'Others'. For example, if number of slices is 5 and the value of **Show up to** is 3, then the third, fourth, and fifth slices are combined and shown as 'Others'.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

If the difference between the values of two pie slices is 1000 or more, the slice with the lesser value might not be displayed on the pie chart. To see the smaller slice, hide the larger slice by clicking it.

Example

For the search query:

```
use LogLogic_Appliance | GROUP BY ll_eventAction | COLUMNS ll_
eventAction AS EventAction, count(*) AS EventCount | (ll_eventAction IS
NOT NULL)
```

the Slice Name is `EventAction`, and the Slice Value is `EventCount`.

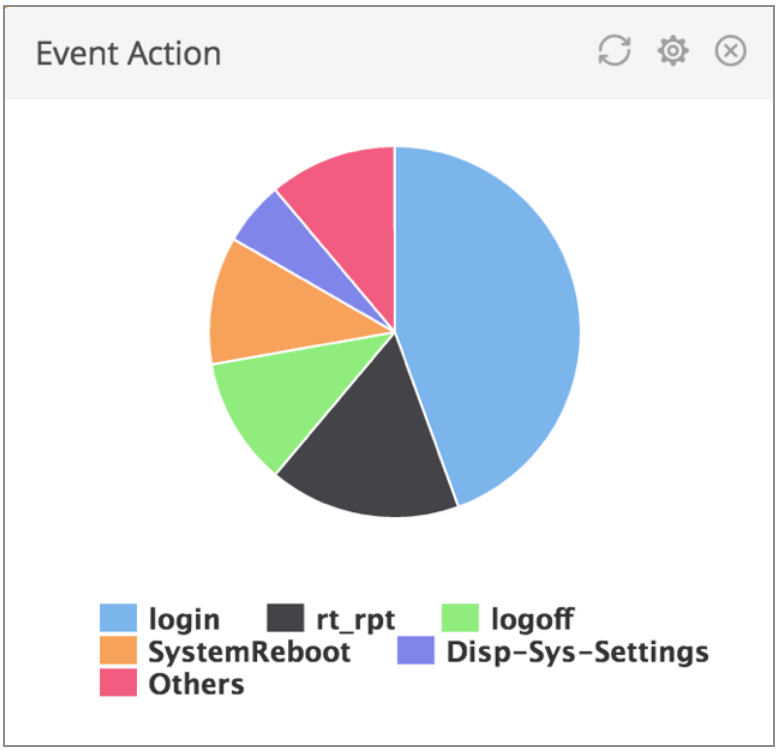

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Number Widget

A numerical value widget displays an important metric for single glance analysis.

Use the following information to configure the widget:

- Fetch data from source

- Number widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Number widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Number widget configuration	
Show value of	Select the column name from the list.
Unit label	Select the appropriate option or enter the desired unit.
Result label	Enter a label for the result. This label is displayed below the number on the widget.
Range	Specify the range values to define the size of the number widget and the range of each threshold slider.
Threshold	After selecting the **Threshold** check box, you can specify the threshold value. The color on the chart depends on the threshold values. • When the number is less than or equal to the minimum

Field	Description
	threshold value, the color changes to green. • When the number is greater than the minimum threshold value but less than or equal the maximum threshold value, the color changes to yellow. • When the number is higher than the maximum threshold value, the color changes to red.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
use LogLogic_Appliance | COLUMNS count(*)
```

the threshold value is set to 10000.

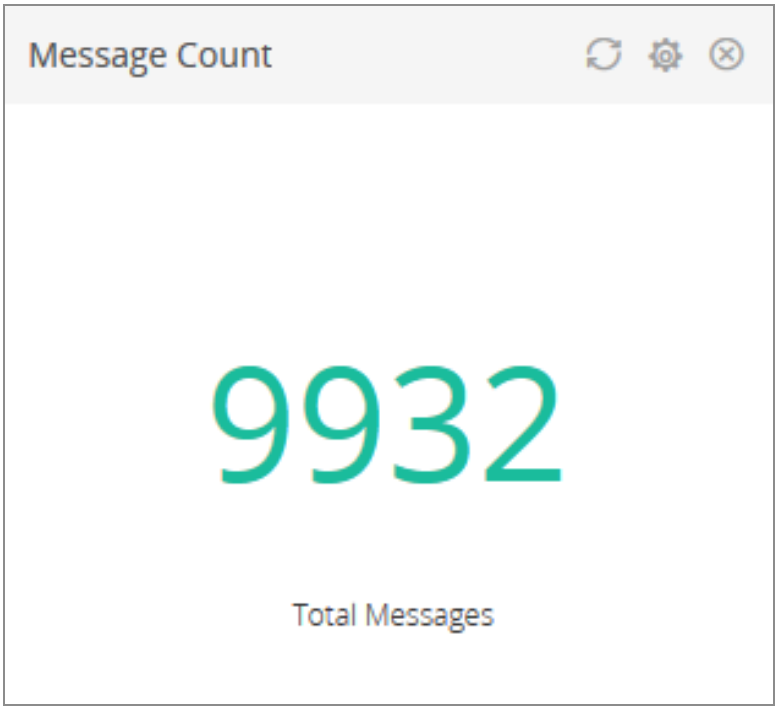

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Gauge Widget

Gauge chart displays information in the form of a dial. The values of selected columns can be displayed using a needle, as concentric circles, or like a speedometer.

All charts are configured in the same way, but displayed differently depending upon the type of gauge selected.

Use the following information to configure the widget:

- Fetch data from source
- Gauge widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Gauge widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Gauge widget configuration	
Gauge type	Select the type of gauge: • Needle • Solid • Wheel **Default:** Needle
Show value of	Select the columns to be displayed. You can select up to three columns of numeric type. • For Needle and Solid types, if you select more than one column, multiple charts are displayed, one per selected column.

Field	Description
	• For the Wheel type, one chart is displayed, and the selected columns are displayed as concentric circles.
Unit label	Type the desired unit.
Range	Specify the range values to define the size of the gauge chart and the range of each threshold slider.
Threshold	After selecting the **Threshold** check box, you can specify the threshold value. The color on the chart depends on the threshold values. • When the number is less than or equal to the minimum threshold value, the color changes to green. • When the number is greater than the minimum threshold value but less than or equal the maximum threshold value, the color changes to yellow. • When the number is higher than the maximum threshold value, the color changes to red.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds.

Field	Description
	This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
USE Other_UNIX | ll_eventActionID = 62 and ll_eventStatusID = 4 | GROUP
BY ll_eventActionID , ll_eventStatusID
```

the Show value is ll_eventActionID , ll_eventStatusID and Count(*).

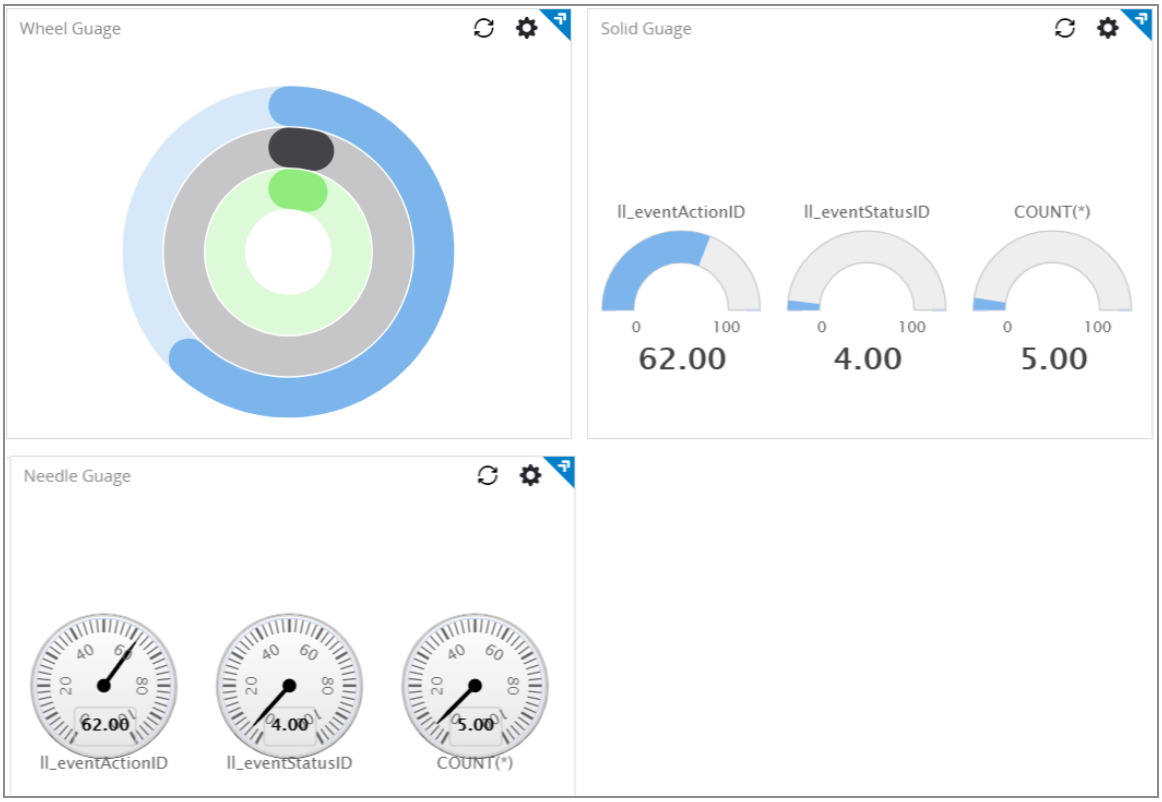

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Stacked Column Widget

This widget is used to show the distribution of the total count of one selected column over its distinct values.

Use the following information to configure the widget:

- Fetch data from source
- Stacked column widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Stacked column widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Stacked column widget configuration	
X-axis data	Select the column name to define the X-axis.

Field	Description
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data	Select the two columns to define the Y-axis of the widget.
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Categorize by	Define the column name by which the Y-axis data is combined into a series.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
use LogLogic_Logu | GROUP BY ll_node, ll_loglevel | COLUMNS ll_node, ll_
loglevel, count(*)
```

the X-axis is `ll_loglevel`, the Y-axis is `count (*)`, and the Categorize by is `ll_node`.

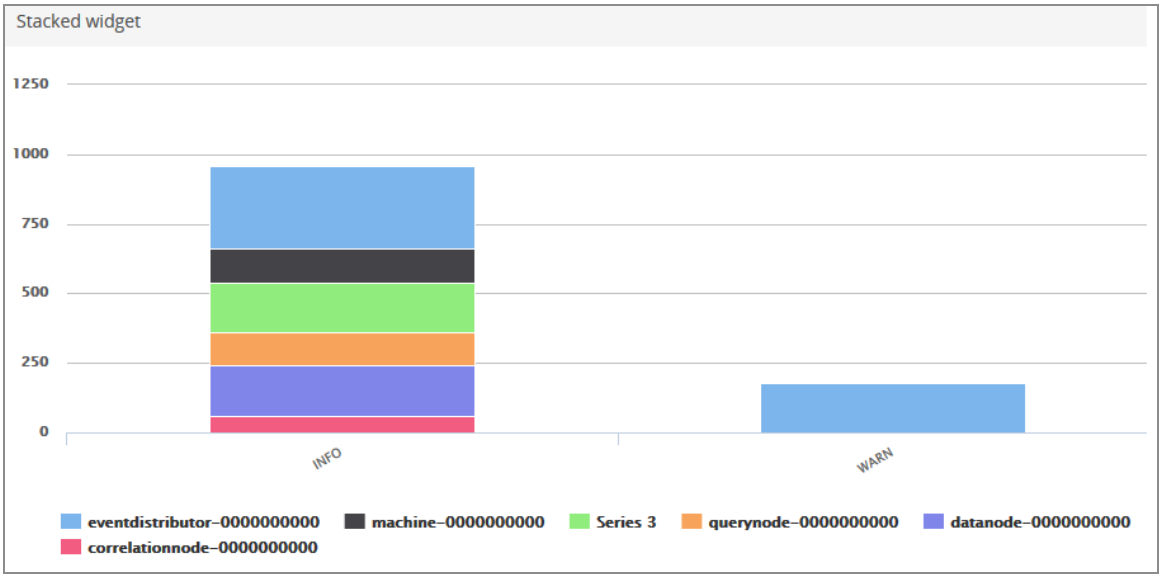

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Combined Widget

This widget is used to show the distribution of the total count of a selected column over its distinct values.

Use the following information to configure the widget:

- Fetch data from source
- Combined widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Combined widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Combined widget configuration	
X-axis data	Select the column name to define the X-axis.
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data	Select the two columns to define the Y-axis of the widget.
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Show Average	Select the check box to show the average in the line format.
Show Total	Select the check box to show the total in the pie format.
Categorize by	Define the column name by which the Y-axis data is combined into a series.
Widget description	Enter a short description for the widget. The description is displayed

Field	Description
	on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

The Combined Widget displays the pie, bar graph, and line graph for the results of the query. The pie and bar graphs display the values by the selected X-axis. The line graph joins the average values in each category. To calculate the average values, the sum of the Y-axis values is divided by the total number of items on the category.

Hover your mouse over the widget to view the value at that point. Clicking the value opens the search results of that value on an Advanced Search tab. However, if you click on the average line in the widget, the Advanced Search tab displays search results for the entire query.

To view the line and bar graphs properly, you can drag the pie to any area of the widget.

Example

For the search query:

```
USE LogLogic_Logu GROUP BY ll_node | COLUMNS ll_node, count(*)
```

the X-axis is `ll_node`, the Y-axis is `count (*)`, and Categorize by is `ll_node`.

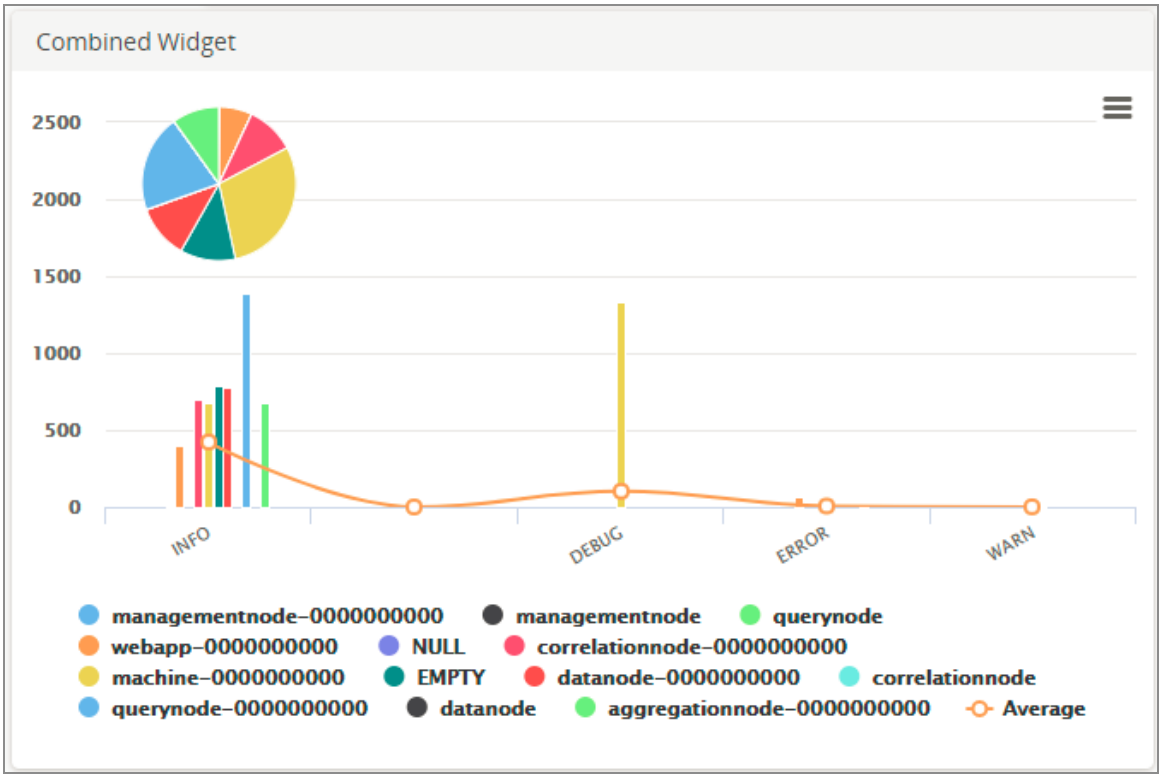

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Treemap Widget

This widget is used to visualize various thresholds in the form of a colored treemap.

The treemap widget is useful to view thresholds such as CPU, memory, indexer, and disk thresholds. You can visualize if some factors in the system have issues or are normal.

Use the following information to configure the widget:

- Fetch data from source
- Treemap widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Treemap widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Treemap widget configuration	
Tile Name	Select the column name. If the column names are already defined in the search query, the **Tile Name** column is automatically filled. Otherwise, as you start typing in the field, the available matching column names are displayed.
Tile Value	Select the column name by which the treemap tile is to be represented. The value of this column is used for the size of the tile. If the column names are already defined in the search query, the **Tile Value** column is automatically filled. Otherwise, as you start typing in the field, the available matching column names are displayed.
Categorize by	Define the column name by which the Y-axis data is combined into a

Field	Description
	series.
Use Color Axis	In the **Min Color** and **Max Color** fields, specify the range of minimum and maximum values of the color to be represented on the axis. To get the right color spread for the tile values, you must adjust the color axis.
Use Color Value	Define the column name by selecting the column. The color represented by the **Use Color Value** field is used to color the tiles on the chart. To return specific color values, you can use Enrichment List or EQL Conditional functions such as IIF in the query. If this field is specified, the **Use Color Axis** field is ignored.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to

Field	Description
	refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Note: Clicking the widget opens Advanced Search with the same query that you used for the widget.

Examples

For the search query:

```
use LogLogic_Appliance | GROUP BY ll_eventStatus, ll_sourceIP | COLUMNS
ll_eventStatus, ll_sourceIP, count(*) as count | ll_eventStatus is NOT
NULL | (ll_eventStatus != '')
```

the Tile Name is `ll_eventStatus`, and the Tile Value is `count(*)`. The treemap widget using the Color Axis value:

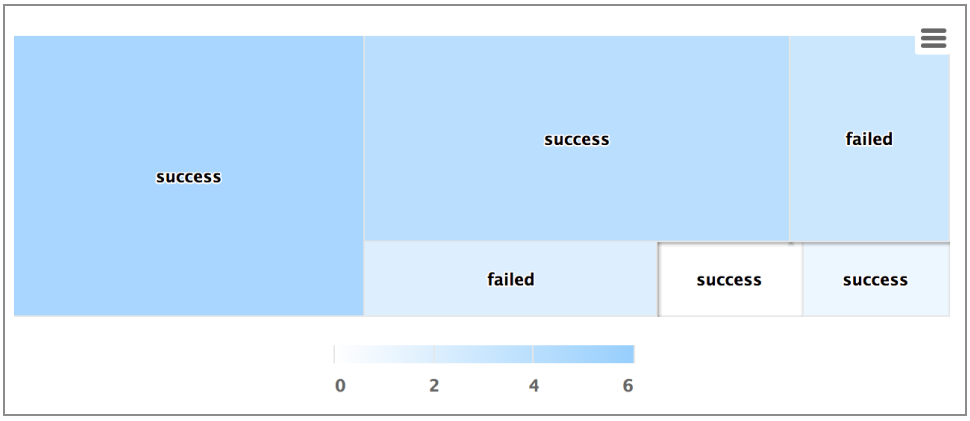

The widget using Color Axis with Categorize By `ll_sourceIP`:

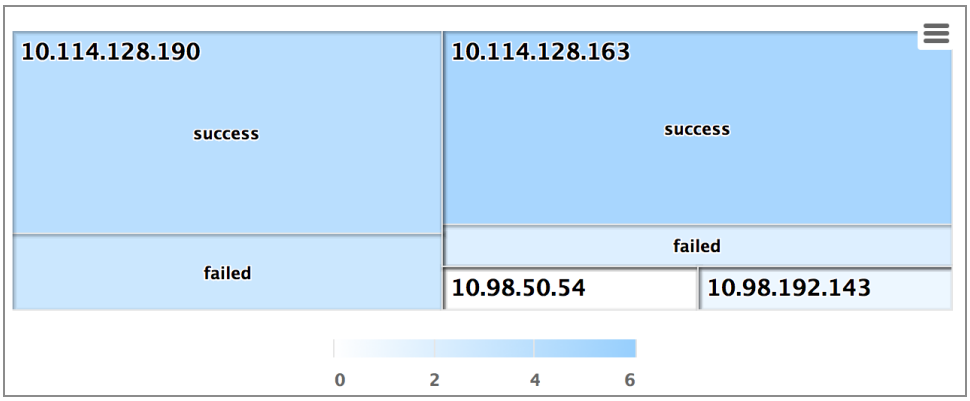

For the search query

```
use LogLogic_Appliance | GROUP BY
ll_eventStatus, ll_sourceIP | COLUMNS ll_eventStatus, ll_sourceIP,
count(*) as count, IIF(ll_eventStatus ='failed', 'red', 'green') AS
color | ll_eventStatus is NOT NULL | (ll_eventStatus != '')
```

using the color value column as color (from the query) and Categorize By `ll_sourceIP`:

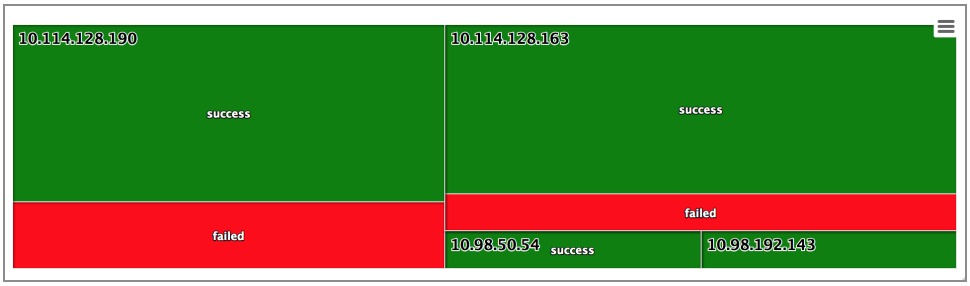

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Heat Map Widget

This widget is used to visualize various thresholds in the form of a colored heat map.

The Heat Map widget is useful to visualize data in the form of a heat map.

Use the following information to configure the widget:

- Fetch data from source

- Heatmap widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Heatmap widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Heatmap widget configuration	
X-axis data	Select the column name to define the X-axis.
Y-axis data	Select the two columns to define the Y-axis of the widget.
Tile Value	Select the column name by which the treemap tile is to be represented. The value of this column is used for the size of the tile. If the column names are already defined in the search query, the **Tile Value** column is automatically filled. Otherwise, as you start typing in the field, the available matching column names are displayed.
Use Color Axis	In the **Min Color** and **Max Color** fields, specify the range of minimum and maximum values of the color to be represented on the axis. To get the right color spread for the tile values, you must adjust the

Field	Description
	color axis.
Use Color Threshold	Define the threshold range for the colors on the heat map. • If **Tile Value** is below the threshold range, the tile color is green. • If **Tile Value** is above the threshold range, the tile color is red. • If **Tile Value** is in between the threshold range, the tile color is orange.
Use Color Value	Define the column name by selecting the column. The color represented by the **Use Color Value** field is used to color the tiles on the chart. To return specific color values, you can use Enrichment List or EQL Conditional functions such as IIF in the query. If this field is specified, the **Use Color Axis** field is ignored.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled.

Field	Description
	Default: OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Examples

For the search query:

```
use LogLogic_Appliance | GROUP BY ll_eventStatus, ll_sourceIP | COLUMNS
ll_eventStatus, ll_sourceIP, count(*) as count | ll_eventStatus is NOT
NULL | (ll_eventStatus != '')
```

the X-axis is ll_eventStatus and the Y-axis is ll_sourceIP

The following are examples of a Heat widget:

- Using default values of color axis

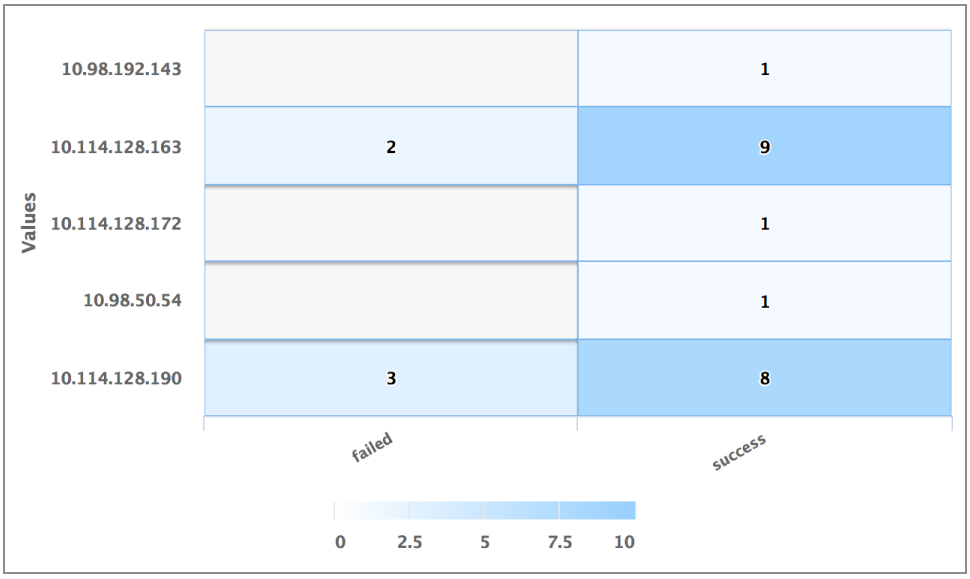

- Using threshold values 4 and 8

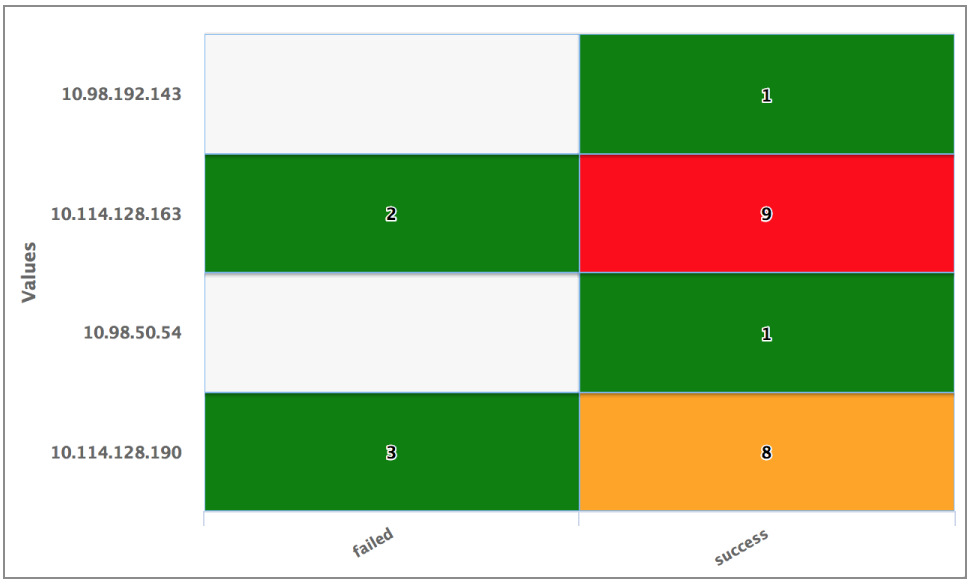

For the search query:

```
use LogLogic_Appliance | GROUP BY ll_eventStatus, ll_sourceIP | COLUMNS
ll_eventStatus, ll_sourceIP, count(*), IIF(ll_eventStatus ='failed',
'red', 'green') AS color | ll_eventStatus is NOT NULL | (ll_eventStatus
!= '')
```

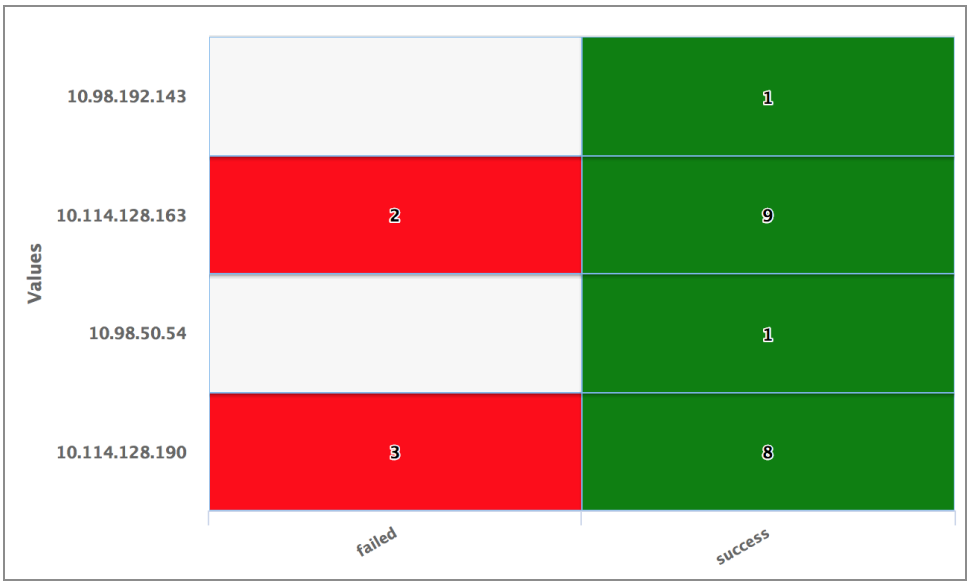

Related Topics

- Widgets

- Adding Widgets to an Advanced Dashboard

DataGrid Widget

This widget is used to visualize query results in the form of a table.

The DataGrid widget is useful to view results in the form of a table. TAIL query results are also displayed in the widget.

Use the following information to configure the widget:

- Fetch data from source
- DataGrid widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **DataGrid widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
DataGrid widget configuration	
Column(s)	Column value of the selected field is used to plot the widget.

Field	Description
Options	Select the required options to display in the widget result: • Show Row Numbers • Show Table Header • Hide Column Separator • Enable Column Selection • Enable Context Menu • Enable CSV Export • Enable Excel Export Even if you select the Enable CSV Export or Enable Excel Export options, the relevant menu options are available in the Filter menu only if the Enable Context Menu option is also selected.
Filter menu	This menu is available only if Enable Context Menu is selected in the **Options** field. 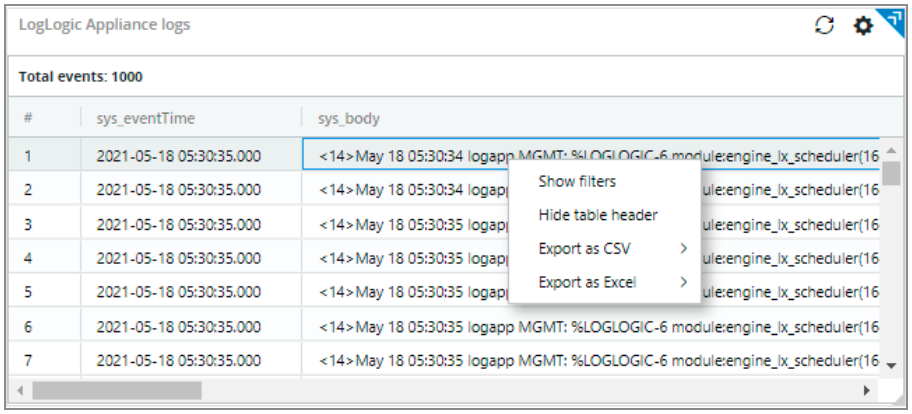 If Enable Context Menu is selected in the **Options** field, then the following options are available after right-clicking the widget: • Show/Hide filters • Show/Hide table header (if the Show Table Header option is also selected in the **Options** field) • Export as CSV - all data or selected data (if the Enable CSV Export option is also selected in the **Options** field)

Field	Description
	• Export as Excel - all data or selected data (if the Enable Excel Export option is also selected in the **Options** field)
Buffer size	Enter the number of rows to be displayed in the DataGrid after refreshing the widget. For example, if the buffer size is 100 and the query returns 150 results, the latest 100 results are displayed. **Default:**1000 rows
Font size	Select the font size of the text that is displayed on the DataGrid widget. Available options: • Comfortable: 12 px • Cozy: 11 px • Compact: 10 px **Default:** Comfortable (12 px)
Maximum number of decimal places	Enter the maximum number of decimal places to display for the columns that display float type of data. You can specify a value between 0 and 10. **Default:** 2
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and

Field	Description
	earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example

For the search query:

```
use LogLogic_Appliance
```

- The columns are `sys_eventTime` and `sys_body`.
- The total number of events is displayed if the **Options** field includes Show Table Header.
- The filter menu is available in the results if the **Options** field includes Enable Context Menu.

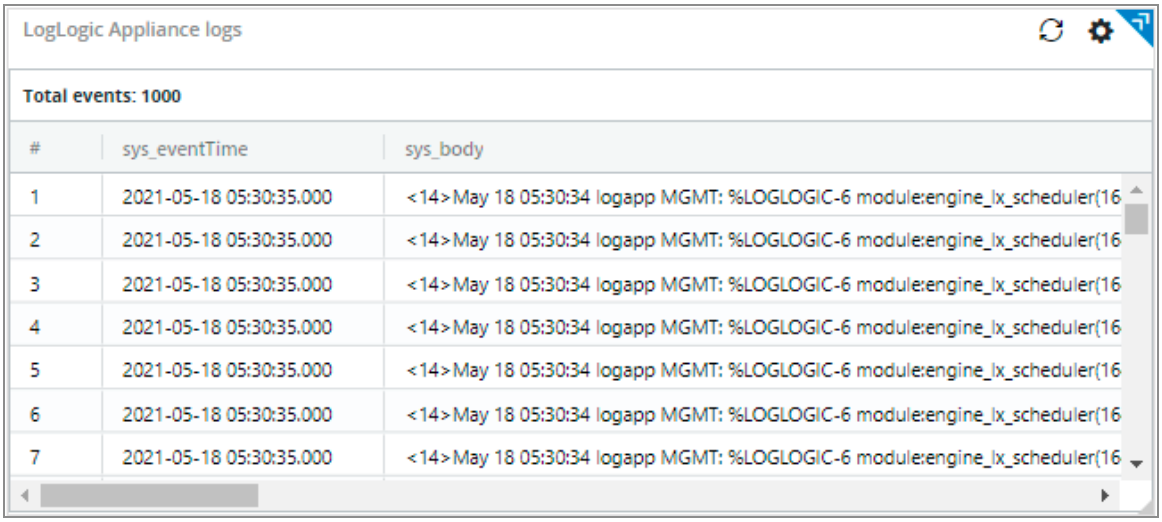

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Range Bar

This widget is used to show a range of values as a bar graph. Unlike a bar graph, the bars in a range bar need not start at zero. Each bar is rendered as a range of specified values.

Use the following information to configure the widget:

- Fetch data from source
- Range bar widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Range bar widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Range bar widget configuration	

Field	Description
X-axis data	Select the column name to define the X-axis.
X-axis label	Define the label name for the X-axis that is displayed on the widget.
Y-axis data (ranges)	Select the two columns to define the Y-axis of the widget.
Y-axis label	Define the label name for the Y-axis that is displayed on the widget.
Show legends	Select the check box to display legends on the chart.
Show Column Labels	Turn on the toggle to show or turn off to hide the data point on the range. **Default:** ON
Show inverted	Select the check box to invert the X-axis and Y-axis values.
Categorize by	Define the column name by which the Y-axis data is combined into a series.
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds.

Field	Description
	This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example Search Query

```
use LogLogic_System_CPU_Load_1_Min | GROUP BY hours(lls_time, 1) |
COLUMNS hours(lls_time, 1) AS Time, MIN(lls_loadAvg) AS MinLoad, MAX
(lls_loadAvg) AS MaxLoad, COUNT(lls_loadAvg) | lls_time IN -1d
```

In this example, `lls_time` values are plotted on the X-axis, and the range between `MinLoad` and `MaxLoad` on the Y-axis.

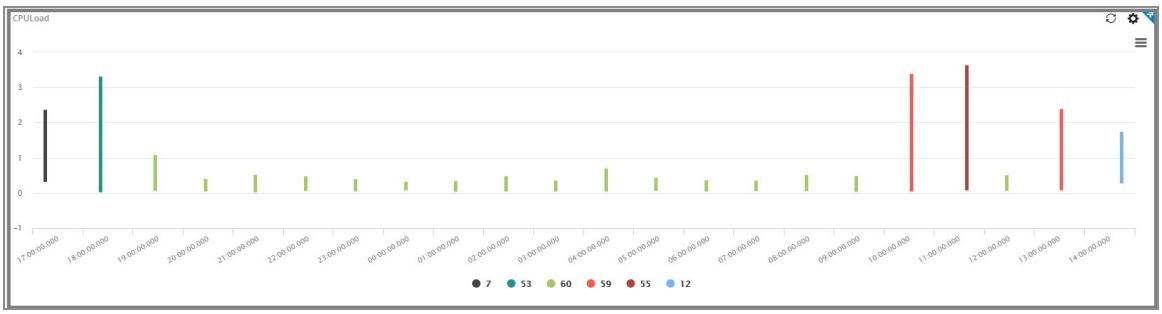

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Geomap Widget

At times, geographical distribution of data is as important as the data itself.

The Geomap widget gives you a unified view of your data visualization and its geographical distribution. For example, you can plot VPN connection logs and the IP addresses from

which they originate. The widget displays the IP addresses as points or a bubble chart on the geographical map.

Geomap widget uses the IP address and location information from your data model to plot the data on a map. You can select the geographical map as a particular region or country, or the world map. You can plot IP addresses as points or as a bubble for a region on the map.

Note: To use this widget, the MaxMind database must be available on your appliance. Contact your administrator.

Use the following information to configure the widget:

- Fetch data from source
- Geomap widget configuration

Field	Description
Fetch data from source	
Query	Enter a search query. The query must include the `geoiplookup()` function. For more information, see Miscellaneous Functions. • To start an EQL statement, enter USE. • To start an SQL statement, enter SELECT. You can search based on filter and time Bloks as well. After you enter the search query, the columns from the query are used as field options in the **Geomap widget configuration** section. For more information about EQL search syntax, see Event Query Language Reference.
Date & Time	You can enter absolute and relative time ranges. For example, enter -5h as a relative time range to display results for events that occurred in the past 5 hours. For more information and examples, see Time Range Expressions.
Geomap widget configuration	

Field	Description
Source IP	Select the data model column in which IP addresses are stored and that has a data type as INET_ADDR.
Location	Select the column with the location of the area. You must use the `geoiplookup()` function and specify the location parameter along with the **Source IP** field so that the `geoiplookup()` function calculates the latitude and longitude of the area: `geoiplookup(<columnName>, 'location')`
Location tooltip	Select the column with the geographical information of the required area. You must use the `geoiplookup()` function, such as country, city, postal, and so on, depending on the configuration of the geographical database available on your appliance. • The value of selected information is displayed as the tooltip. • If the field is empty, then the IP address is displayed as the tooltip of the point on the map. For example, if you use the `country` option, the country name is displayed as a tooltip of the point on the map.
Bubble weight	Select the column that can be used to indicate a number of IP addresses for the selected location. If multiple IP addresses belong to that region on the map, then a bubble is displayed instead of a dot. The size of the bubble is proportionate to the column value.
Add series (+)	Click + to add another row of the Source IP, Location, Location tooltip, and Bubble weight fields. You can add multiple series to the chart. For example, if a data model includes multiple columns that return an INET_ADDR type.
Region border color	Select the color of the region border. **Default:** #A0A0A0
Map background	Select the background color of the region.

Field	Description
color	**Default:** #E9E9E9
Map	Select a map type from the available list. **Default:** World continents
Widget description	Enter a short description for the widget. The description is displayed on the Advanced Dashboard when you hover over the widget.
Auto load	Turn on the toggle to automatically load widget data on the Advanced Dashboard as soon as you save the widget or when you navigate to the dashboard. Disabling the **Auto load** option also disables the **Auto refresh** option. However, you can manually refresh the widget on the Advanced Dashboard to load its data. **Default:** • Disabled for widgets created in LogLogic LMI 6.3.1 and later • Enabled for the widgets created in LogLogic LMI 6.3.0 and earlier
Auto refresh	Turn on the toggle to refresh the widget every few seconds. This setting is enabled only if the **Auto Load** option is enabled. **Default:** OFF
Refresh widget every	If **Auto refresh** is set to ON, then enter a time interval in seconds to refresh the widget. Refresh action starts after the data is completely retrieved and displayed.

Example Search Query

```
use system | COLUMNS sys_collectIP, geoiplookup(sys_
collectIP,"location"), geoiplookup(sys_collectIP,"country"), count(sys_
collectIP) | (sys_collectIP != '127.0.0.1' AND sys_collectIP !=
'198.51.100.0' AND sys_collectIP != '::1') | group by sys_collectIP
```

In this example, IP addresses of locations are marked as separate points and country

names are displayed as the tooltip. The IP address 198.51.100.0 is an example. Replace it with the IP address of your appliance.

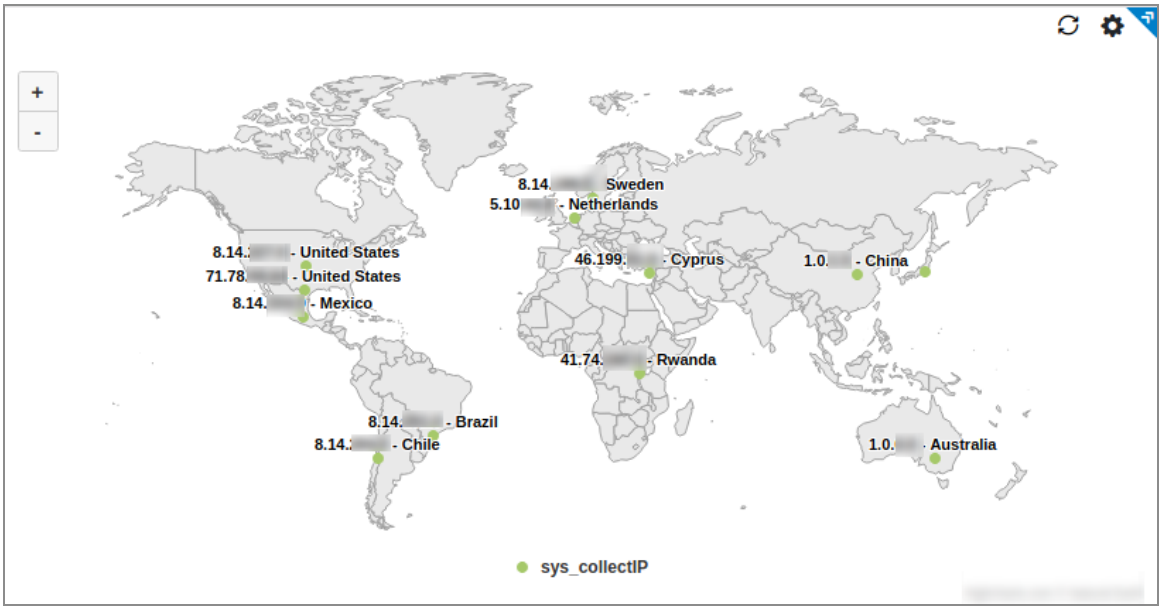

Note: Political boundaries or place names contained in the maps do not necessarily reflect TIBCO's view of any particular disputed border or place name and TIBCO does not warrant the accuracy of such boundaries or place names.

Related Topics

- Widgets
- Adding Widgets to an Advanced Dashboard

Adding Widgets to an Advanced Dashboard

You can create a new dashboard with multiple widgets based on your requirements. For best performance, limit the number of widgets to 20 per dashboard.

Tip: After an aggregation rule is created, a filter Blok is automatically created in the system. If there are multiple time clauses in the GROUP BY query, multiple filter Bloks are created - one for each time clause. To save on query response time, you can use the filter Blok in Advanced Dashboards instead of typing the entire query.

Procedure

1. Click **Dashboards > Advanced Dashboards** to display the **All Dashboards** page.
2. (Optional) Create a dashboard or a dashboard group to add the widget.
   - To create a new dashboard, click **Create Dashboard**.
   - To create a new dashboard group, click **Create Group** or **Create Nested Group**.
3. Click the dashboard to open the dashboard view.

   Note: You can add widgets with identical names in one dashboard.
4. From the **Widgets** pane, click the type of widget that you want to add.
5. To configure the widget, click the **Configure** link or click the Settings icon on the upper-right corner of the widget.
   a. To edit the widget name, click *Untitled widget* and enter the name of the widget in the field.
   b. Edit the fields as required. The configuration fields are different for each type of widget. See the Widgets section.
6. Click **Save** to save the widget.

Result

The widget is added and the retrieved results are displayed on the dashboard. While the data is loading, you can click the **Stop** button and change the widget configuration, and click **Start** to reload the data.

Note: By using columns with large content, the content might become unreadable on a widget, for example, if you use `sys_body` on the X-axis while using a Bar Chart widget.

Data Models

LogLogic LMI parses log data into structured formats to enhance search and analysis. Based on the log source type, you can define how to parse your data and which columns to extract.

From the **Management > Advanced Features > Data Models** page, you can view all data models available on your appliance.

Using data models, LogLogic LMI parses log data in a structured format to enhance search and analysis. Based on the log source type, you can define parsing rules within the data models to decide how to parse your data and which columns to extract.

The data models in LogLogic LMI can be broadly classified into the following categories:

- Advanced Data Models, which use parser types such as syslog, Regex, JSON, and so on.
- GP Parser-Based Data Models, which are data models that use a grouped-pattern parser (GP parser). GP parsers are especially efficient in handling complex Regex parsing rules that work on heterogeneous free-text logs.

Functions of Data Models

Using data models you can:

- Define parsing rules that extract columns from your data.
- Define a schema for an event.
- Name and specify the data types for extracted columns.

Parsing Rules

A data model can be associated with multiple parsing rules. Sometimes within the same source, some logs are completely different to others, and it is not practical, or even possible, to match them all with a single rule. You need a different way of parsing for each kind of log, and you can do that by defining several rules, each targeting one type of log.

If a data model has more than one parsing rule defined, then the extracted column set is the union of the column sets of all parsing rules and the additional system-defined columns. For example, create a data model and define a parsing rule, Rule1, to extract four defined columns and Rule2, to extract eight different defined columns. Now, when you run a search query on this data model, the 12 columns are displayed.

Parsing rules are applied top to bottom in the order they are defined in a data model. For example, if Rule1 matches some of your data then it is used to extract column values. If

Rule1 fails to match with your data, then only Rule2 is applied, and so on. You can change the order of parsing rules.

For an overview of the parsers in advanced data models, see Types of Parsers in Advanced Data Models. In GP parser-based data models, there is only one type of parser. See GP Parser-Based Data Models

For an overview of each parser, see Types of Parsers in Advanced Data Models.

Advanced Data Models

Advanced data models use different types of parsers to parse the logs from different log sources.

You can create a data model that defines which log source to use for parsing, based on the data relevance. For multiple log sources, the order of precedence can be defined in a specified query. The system columns are event metadata. All system columns are displayed with the prefix `sys_` and all columns from built-in parsers are displayed with the prefix `ll_` in the **Columns** panel.

LogLogic LMI also provides built-in data models. For the list of built-in data models, see the Supported Log Sources list in *TIBCO LogLogic® Log Source Packages Installation and Upgrade*, which is available on the TIBCO eDelivery website or TIBCO Support website after logging in.

Downstream Parsing

By configuring downstream parsing, you can chain a parsing rule of an advanced data model to another advanced or GP parser-based data model to further process the columns retrieved from the main parsing rule. This method of using two levels of parsing one after another helps to parse log messages having mixed formats. Consider the following example.

Example

In Windows Snare logs, the tab character is escaped with a backslash (\) and hence is a two-character sequence (\t). Consider the following sample log message received in LogLogic LMI:

```
<13>Apr 15 00:54:19 10.199.187.140
MSWinEventLog\t0\tSecurity\t406243509\tSat Dec 25 17:16:57
2004\t4699\tMicrosoft-Windows-Security-Auditing\tSYSTEM\tUser\tSuccess
Audit\tX78UNT2AJIC1Y\tObject Access\t\tA scheduled task was deleted.
Subject:    Security ID:    S-1-5-21-2798475463-3993569027-3406240830-
49243    Account Name:    QPAGjyT74D    Account Domain:    BJ    Logon
ID:    0x6b441b23    Task Information:    Task Name:     \dEmG_
Jclbu7ZtMGKUmLe3vArEqF_\erCw\TACZZvZ4mbn\M36kGL-2mYD-
oI\cefE4AJTOA0rGj1V7G0LLcTmen_    Task Content:     <Task version="1.2"
xmlns="http://schemas.microsoft.com/windows/2004/02/mit/task">
<RegistrationInfo> <Date>2017-04-21T17:54:26.989698</Date>
<Author>BJQPAGjyT74D</Author> </RegistrationInfo> <Triggers/>
<Principals> <Principal id="Author"> <RunLevel>LeastPrivilege</RunLevel>
<UserId>BJQPAGjyT74D</UserId> <LogonType>S4U</LogonType> </Principal>
</Principals> <Settings>
<MultipleInstancesPolicy>IgnoreNew</MultipleInstancesPolicy>
<DisallowStartIfOnBatteries>true</DisallowStartIfOnBatteries>
<StopIfGoingOnBatteries>true</StopIfGoingOnBatteries>
<AllowHardTerminate>true</AllowHardTerminate>
<StartWhenAvailable>false</StartWhenAvailable>
<RunOnlyIfNetworkAvailable>false</RunOnlyIfNetworkAvailable>
<IdleSettings> <StopOnIdleEnd>true</StopOnIdleEnd>
<RestartOnIdle>false</RestartOnIdle> </IdleSettings>
<AllowStartOnDemand>true</AllowStartOnDemand> <Enabled>true</Enabled>
<Hidden>false</Hidden> <RunOnlyIfIdle>false</RunOnlyIfIdle>
<WakeToRun>false</WakeToRun>
<ExecutionTimeLimit>P3D</ExecutionTimeLimit> <Priority>7</Priority>
</Settings> <Actions Context="Author"> <Exec>
<Command>CWindows\System32\svchost.exe</Command> </Exec>
</Actions></Task>\t4054120554
```

Such log messages, which include a two-character tab, are not parsed in LogLogic LMI. To accurately parse such logs, you can create a new data model, create a parsing rule with parser type as Regex, and add the following pattern in the Regex pattern field:

```
^(.*)
```

Next, select the downstream data model as `Microsoft_Windows` and provide an expression to substitute the two-character `\t` with a single-character `\t`:

- Data model: `Microsoft_Windows`
- Expression: `substitute($1,"\\t", "\t")`

The main parsing rule captures the entire log message in the `$1` variable. The `$1` variable is used as the parameter of the function performing the character substitution. Its output is used as the input of the downstream data model.

Downstream Parser Matrix

The following matrix depicts the combination of parsers that can work (indicated by OK), and the combinations that have limitations in downstream parsing. The row indicates the main parser type and the columns in that row indicate the downstream parser type.

For an overview of each parser, see Types of Parsers in Advanced Data Models.

Downstream parser type →	**Syslog**	**KVP**	**JSON**	**XML**	**Column**	**Regex**	**CEF**	**GP**
Main parser type ↓								
Syslog	N/A	OK	OK	OK	OK	OK	OK	OK
Key-value parser (KVP)	OK	N/A	Limited	Limited	Limited	OK	Limited	OK
JSON	OK	OK	N/A	OK	OK	OK	OK	OK
XML	OK	OK	OK	N/A	OK	OK	OK	OK
Columnar	OK	OK	OK	OK	OK	OK	OK	OK
Regex parser	OK	OK	OK	OK	OK	OK	OK	OK
CEF	OK	OK	OK	OK	OK	OK	N/A	OK

Limitations of Downstream Parsing

- You can configure the downstream parsing only for one parsing rule and to only one advanced data model. Only that parsing rule must be enabled and all other parsing rules must be disabled.
- If the main parser type is KVP, then there are limitations to configure JSON, XML, columnar, or CEF parsers as the downstream parser type:
    - When using JSON, XML, or columnar as the downstream parser type: the key separator must not be comma (,) and the JSON statement must be enclosed in single quotes (').
    - When using CEF as the downstream parser type: the key separator must not be vertical bar (|) and the JSON statement must be enclosed in single quotes (').
- When using GP as the downstream parser type, you must be extremely careful to provide an accurate and precise expression. The result of that expression must be a string that meets the downstream parser requirements.

> **Tip:** You can copy the exact expression from the GP parser-based data model and then paste it in the **Expression** field.

Types of Parsers in Advanced Data Models

The following types of parsers are available in advanced data models:

- Key-value Parser
- JSON Parser
- XML Parser
- Columnar Parser
- Regex Parser
- CEF Parser
- Syslog Parser

Key-value Parser

This parser uses simple key-value pair parsing rules to extract keys and values. The parser

recognizes patterns such as `k1=v1, k2=v2, k3=v3`. You can use key-value pair separators, for example, space, comma (,), or semi-colon (;), and key and value separators, for example, equal sign (=) or colon (:). Separators can be either one or more characters that have to be matched exactly or they can be regular expressions.

When referring to a value in a column expression, it is referred to as `$<key name>`. So, for a key with the name 'user' the value is referred to as `$user`.

Regular expressions can also be used to parse data from the beginning and ending of the event. This can be useful when parsing events that either start with or end with data that is not in the key-value pair format. If these regular expressions contain named groups, then those groups are extracted and can be used to populate columns.

You can specify the name of the last key in the data. Any data after that last key is treated as the value of that last key. This can be useful in situations where the last value in the data contains characters that might be interpreted as separators.

JSON Parser

This parser parses JSON logs and accepts valid JSON as input. The parser recognizes key-value pairs, key-object pairs, and arrays.

When referring to a key in a column expression, the column name is referred as `$<key name>`. When referring to array elements, the column name is referred as `$<key name>_<index_of_array>`, where *<index_of_array>* starts with 0. When referring to objects, the column name is referred as `$<key name>_<object name>`.

For example, for the following XML elements:

```
{"key1":"value1","key2":"value2"}
```

the column names would be:

- key1
- key2

You might need to change the default column names to ensure that only the supported characters are used. You can use underscore (_) and alphanumeric characters. However, a column name cannot begin with a number.

XML Parser

This parser parses XML logs, and accepts syntactically valid XML as input. The parser recognizes element nodes, text nodes, and attribute nodes.

When referring to a key in a column expression, the column name is referred as `$<key name>`. When referring to sibling elements, the column name is referred as `$<key name>_<index_of_element>`, where *<index_of_element>* starts with 1. When referring to XML elements, the column name is referred as `$<key name>_<element_name>`.

For example, for the following XML elements:

```
<Root>
  <child> This is child1</child>
  <child> This is child2</child>
</Root>
```

the column names would be:

- Root_child_1
- Root_child_2

You might need to change the default column names to ensure that only the supported characters are used. You can use underscore (_) and alphanumeric characters. However, a column name cannot begin with a number.

Columnar Parser

The data is extracted into different columns. This parser operates on data that is separated by a character or a sequence of characters, for example, comma or tab. There is no keyvalue; just the value. The data from different log sources extract different columns depending on keys identified in the data. When referring to a column in a column expression, it is referred to as `$<column number>`. So the first column is referred to as `$1`, the second column is `$2`, and so on.

You can also use regular expressions to parse data from the beginning and end of an event, for example, when parsing events that either start with or end with data that is not in columnar format. If the regular expressions contain named groups, those groups are extracted and are used to fill values in the columns.

Regex Parser

Regular expressions (Regex) are a sequence of characters that form a search pattern, mainly for use in pattern matching with strings or string matching. LogLogic LMI can use regular expressions for extracting columns from matched events.

Note: Working knowledge of regular expressions is a prerequisite.

Each character in a regular expression is either a meta character with its special meaning, or a regular character with its literal meaning. Together, they can be used to identify textual material of a given pattern, or process a number of instances of it that can vary from a precise equality to a very general similarity of the pattern.

LogLogic LMI supports the regular expression meta characters, based on Java regular expressions. For details, see Supported Regular Expression Characters.

Columns are extracted using either the capturing group pattern (simple parenthesis), the named capturing group pattern (?<name>), or a combination of both. When referring to a column in a column expression, when using named capturing groups the column name is that specified by the group name, preceded by “$”. When using unnamed capturing groups, the name is “$” followed by the group index. So the first unnamed group column is referred as `$1`, the second as `$2`, and so on, while a group named “user” is referred as `$user`. When using a combination of named and unnamed capturing groups, the named capturing group columns must be referred to by their given names rather than by "$" followed by their index.

CEF Parser

HP ArcSight Common Event Format (CEF) is an open log management standard. CEF defines a syntax that comprises a standard header and a variable extension, formatted as key-value pairs. Based on the ArcSight Extension Dictionary, the CEF header columns Version, Device Vendor, Device Product, Device Version, Signature ID, Name, and Severity are extracted into columns with their names, and expressions set to $cefVersion, $cefDeviceVendor, $cefDeviceProduct, $cefDeviceVersion, $cefSignatureID, $cefName, and $cefSeverity respectively.

The name of a column for an extension listed in the *ArcSight Extension Dictionary* is the full name of the extension. The name of a column for an extension that is not listed in the *ArcSight Extension Dictionary* is the key name as it is displayed in the data preceded with “$”.

The expressions of the non-timestamp extension columns are the CEF Key Names as defined in the *ArcSight Extension Dictionary*. The expressions of the timestamp extension columns are of the form ToTimestamp(<$CEF Key Name>, <proposed format>) where <proposed format> is a suggestion for the correct format to use when parsing the data.

Some extensions in the *ArcSight Extension Dictionary* have names that start with the asterisk (*). Since LogLogic LMI does not allow column names to start with asterisk (*), an asterisk (*) is omitted from the column name. For example, the *sourceProcessId extension is extracted into a column named *sourceProcessId*.

When the event was written, the pipe (|), equal sign (=), and backslash (\) characters might have been escaped by inserting a backslash (\) in front of them. The CEF parser removes the backslash (\) character, returning the data to its original form. For example, if the value of the Name header in the event is "detected a \| in message", the value of the cefName column is "detected a | in message".

Syslog Parser

Data conforming to the Syslog standard defined in RFC-5424 (https://tools.ietf.org/html/rfc5424) can be parsed using the Syslog Parser.

Note: The older, obsolete format described in RFC-3164 is not supported.

All the header fields defined in the format are extracted as is the Message component. If the log data contains Structured Data elements, those are extracted as well with the names of the resulting columns being composed of <element-name>.<key name> as shown in the following example:

```
<165>1 2003-10-11T22:14:15.003Z mymachine.example.com evntslog - ID47
[exampleSDID@32473 iut="3" eventSource="Application" eventID="1011"] An
application event log entry
```

The following columns are extracted:

- facility = local4;
- severity = notice;
- version = 1;
- timestamp = 2003-10-11 15:14:15 (if LogLogic LMI is running in the PDT time zone);
- hostname = mymachine.example.com;
- appname = evntslog;
- procid = <null>;
- msgid = ID47;

- exampleSDID@32473.iut = 3;
- exampleSDID@32473.eventSource = Application;
- exampleSDID@32473.eventID = 1011;
- msg = An application event log entry

Managing Data Models

You can manage data models from the **Management > Advanced Features > Data Models** page:

> **Note:** The duplicate icon, delete icon, and move icon are enabled after you select one or more data models.

Task	Steps
View and search data models	See Viewing and Searching Data Models.
Create a new data model	You can add data models using two modes. You can switch between the modes at any time. All information associated with a data model is preserved when you switch from graphical to raw mode. • Creating a Data Model in Graphical Mode: This is the default mode. A wizard helps you add data models and the associated rules. • Creating a Data Model in Raw Mode: This is for advanced users who understand the JSON syntax. Use JSON syntax to add a data model and associated rules.
Edit a data model	See Editing Data Models.
Duplicate	Select the data model and click the duplicate icon. **Note:** You cannot duplicate data models in the **System** group.

Task	Steps
Delete	Select the data model and click the delete icon . After you delete a model, it cannot be recovered. **Note:** You cannot delete built-in and system-created data models.
Move	Select the data model and click the move icon . Select the group to which you want to move it. **Note:** You cannot move built-in and system-created data models.
Rename	Select the data model and click **Edit** in the Details panel. Then rename the data model and click **Save**. **Note:** • You cannot rename built-in and system-created data models. • You cannot rename a data model using the rename icon .
Enable or disable	Data models can be enabled or disabled at any time. Select the data model and click **Enabled**. **Note:** By default, the built-in and system-created data models are enabled.

Data Model Groups

The **Management > Advanced Features > Data Models** page displays data models in groups. The **All Data Models** group displays all data models.

For information about groups, types of groups, and managing groups, see the Groups section.

Viewing and Searching Data Models

From the **Management > Advanced Features > Data Models** page, you can view, find, or filter data models in the following ways.

Task	Description
View models based on filters	You can use filters to easily find models. Click the **View** list to select the required filter. The following options are available: • All • Created by me • Created by system • Imported
Find data models	You can quickly find the desired model by typing the model name in the **Find** field. As you start typing a model name in the **Find** field, the Data Models page is automatically refreshed showing your selection.
Sort data models	You can sort any column in ascending or descending order on the Data Models page. Click the column name or click the arrow (that is displayed on the right side of the column name when you click in that column) to sort the column.
Show or hide columns	You can show or hide columns, except the mandatory column, from the table. Click  to view all available columns in the table. Select the check box to show the column. Clear the check box to hide the column from the table. The Data Models page is updated immediately.
Searching data models	All enabled data models can be searched using the source filter on the **Search > Advanced Search** page.

For more information and other operations on data models, see Managing Data Models.

Creating a Data Model in Graphical Mode

From the **Management > Advanced Features > Data Models** page, you can create a data model and enable it. Then you can use the data model to analyze results in the normalized format.

In graphical mode, a wizard helps you to create a data model in the following steps:

1. Define a source filter

2. Add sample events and parsing rules
3. Manage columns and data types

Procedure

1. Go to **Management > Advanced Features > Data Models**.
2. Click **Create New Data Model**. By default, the graphical mode opens.

   For instructions on how to add in raw mode, see Creating a Data Model in Raw Mode.
3. On the **Add data model** page, provide the following information:
   a. By default, the data model is enabled. Click the slider to **OFF** to disable the model.
   b. Parent Group: Select a parent group where you want to save the data model.

      You can create a new group or select the **User** group, or select any user-created group

      Default parent group: When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.
   c. Name: Enter the name of the data model.

      The name can include letters, numbers, or underscore (_).
   d. (Optional) Enter the description in the **Description** field.
4. Add a new source filter. For instructions, see Defining a Source Filter.
5. Add sample events and define a parsing rule. For instructions, see Adding a Parsing Rule in an Advanced Data Model
6. Manage your custom columns. For instructions, see Managing Columns and Data Types.
7. Click **Save**.

Result

The new data model is displayed on the **All Data Models** page and in the parent group that you selected.

Defining a Source Filter

You can add a new source filter that is assigned to the data model. The source filters bind multiple data models to a log source.

Procedure

1. In the **Source filter** field, enter the source filter statement that is assigned for this data model. Source filters can only be used on one or more system columns. All filter statements in the FILTER Statement section are supported. However, when running a full text search, the filter statement must be specified explicitly.

   In the following example, `165` is the device type ID that is retrieved from LogLogic LMI and the log message contains *<searchstring>*:

   ```
   sys_sourceType=165 AND sys_body CONTAINS '<searchstring>'
   ```

   Note: If you specify multiple data models, the first model whose filter matches with the event is used to parse that event, extracting all columns specified by that model.

2. Click **Validate** to validate the filter statement.
3. To add a new parsing rule, click **2. Add sample events and parsing rules** or click the right-arrow icon › located on the right side of the page, or, to add only the source filter and save the data model, click **Save**.

What to do next

- To add or edit a parsing rule in an advanced data model, see Adding a Parsing Rule in an Advanced Data Model.
- To edit a GP parsing rule, see Editing GP Parser Rules.

Adding a Parsing Rule in an Advanced Data Model

You can add one or more parsing rules that define how to parse log events.

Note: You cannot add a new GP parsing rule in a data model. You can duplicate an existing GP parsing rule, and make the required changes. For details, see Editing GP Parser Rules.

Procedure

1. Paste the sample log data in the **Sample events** panel.

   This data can be helpful in defining the parsing rule based on the log source. After saving the data model, the sample data is always available when editing the same data model or associated parsing rules.

   Note: You can paste a maximum of 100 KB sample data.

2. In the **Parsing rules** panel, click **Add new rule** to add a new parsing rule.

   You can add multiple rules for the same data model.

3. In the **Name** field, enter the name of the rule.

   The name must contain an alphanumeric character. It can also contain an underscore (_) and hyphen (-).

4. To enable the parsing rule, ensure that the slider is set to **ON**. To disable the rule, click the slider to **OFF**.

5. In the **Filter** field, enter the filter that is assigned to the parsing rule. All regular expression patterns are supported.

   Note: If you do not define the filter, all events are matched with this rule. Parsing rules that are listed after such parsing rule are ignored.

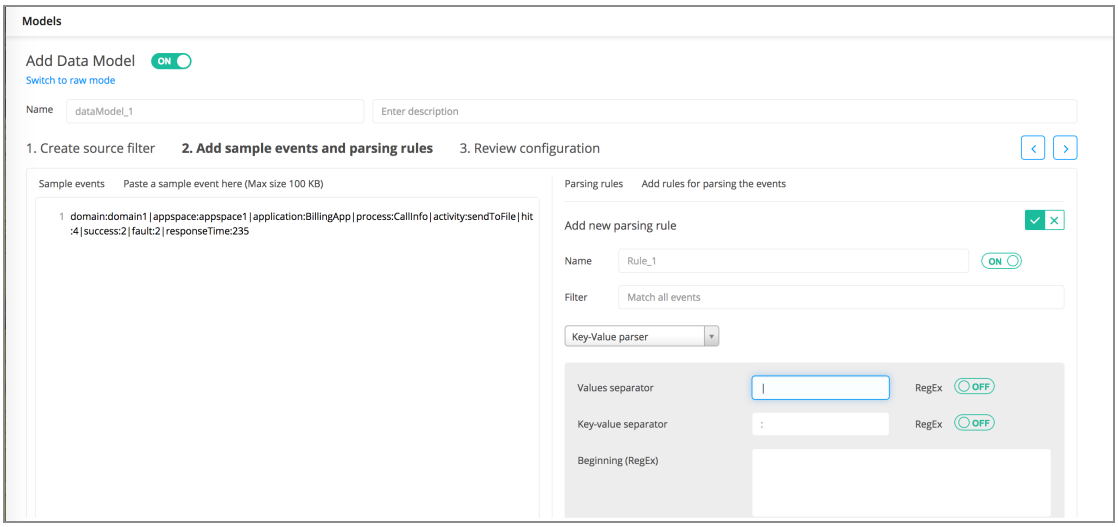

6. From the **Choose parser** list, select the type of parser you want to use. Depending upon the selected parser, you must provide additional information in various fields: see Parser Field Reference - Advanced Data Models.
7. To configure downstream parsing for this parsing rule to another data model, enable the **Manage downstream parsing** field.

> **Note:** By default, this field is disabled. You can configure the downstream parsing only for one parsing rule and to only one advanced data model. Only that parsing rule must be enabled and all other parsing rules must be disabled.

Field	Description
Data model	Select the data model whose parsing rules you want to apply immediately after the parent rule.
Expression	Select the expression to apply to the columns generated by the parent parsing rule.

For more information about downstream parsing, see Downstream Parsing.

8. To extract columns based on the parser type, click **Auto generate columns**. All custom columns are extracted in the **Manage columns for this rule** panel.

a. (Optional) By default, the **Store rule identifier on a new column** field is disabled. To store the rule identifier in the **ll_parsingRuleName** column, enable this field. Doing so displays the rule identifier in Advanced Search results. Then you can filter by the rule identifier column, and check which log messages matched this parsing rule. From the search results you can easily detect if the parsing rule works precisely or needs changes.

b. You can add, edit, or delete custom columns. To add a column, click  . To edit any values, click inside the **Column** and **Expression** fields. To delete a custom column, click the Delete  icon for that column.

- **Column** field: The name of the column that is displayed in the results. Click in the row to add or update any column name. The content assist shows contextual matches of the existing custom column names and you can select the required one.

  > **Note:** Two columns cannot have the same name. Column names are not case sensitive. When defining column names, follow the guidelines described in the COLUMNS Statement section.

- **Expression** field: Define how to map the values extracted by the parser into the defined columns. You can use arithmetic operators and conversion functions when defining an expression. The conversion functions are typically used when you need to define new columns where the expressions for new columns can use conversion functions to convert between data types and combine them using various operators. For more information, see the following topics:
  - For details about the arithmetic operators, see FILTER Statement.
  - For conversion functions, see Predefined EQL Functions.
  - The type of expression depends on the parser type: see Parser Field Reference - Advanced Data Models.

9. Click the **Refresh** icon. The **Parser preview** panel displays all extracted columns and their data types that are matched by the corresponding parsing rule.

   If you have configured downstream parsing, then the **Parser preview** panel displays the following information:

- Null as the values - if downstream parsing failed.
- Values as per the parsed fields - if downstream parsing is successful.

For easy readability, each event that matches with the corresponding rule is identified in the same color. To change the supported data type of a custom column, click in the corresponding **Type** field and select the appropriate data type from the list.

Note: The parser preview is displayed only if log data exists in the **Sample events** panel and if at least one parsing rule is enabled.

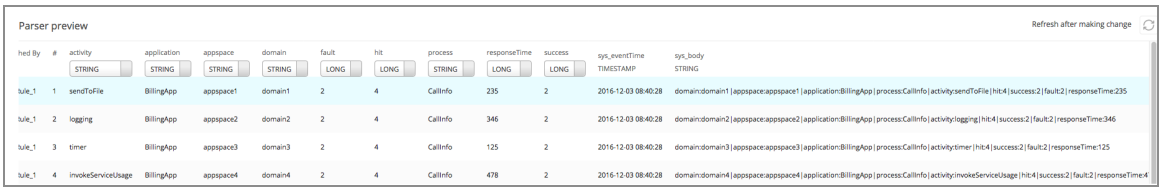

10. To save the new parsing rule, click the Save icon.

    The **Parsing rules** panel displays the newly added rule.

11. Click **3. Review configuration** or click › located on the right side of the page.

What to do next

Manage columns, review the configuration, and save the data model. For more information, see Managing Columns and Data Types.

Parser Field Reference - Advanced Data Models

While adding a parsing rule, you must provide different information depending upon the parser you have selected.

- Key-value parser
- JSON parser
- XML parser
- Columnar parser
- Regex parser
- CEP parser

Key-Value Parser

Field	Description
Values separator	Enter the delimiter that you want to use to separate key-value pairs. You can add only one separator at a time. The delimiters are case sensitive. For example, user=bob,vm=windows where user=bob is one pair and vm=windows is another pair separated with delimiter comma (,). The delimiter can be a single character, a string that has to be matched exactly, or a Java regular expression. RegEx: Select **ON** to use as a Java regular expression or **OFF** to use as a literal string.
Key-value separator	Enter the delimiter that you want to use to separate keys from their values. The delimiters are case sensitive. For example, user=bob where user is a key and bob is a value separated with delimiter equal sign (=). The delimiter can be a single character, a string that has to be matched exactly, or a Java regular expression. RegEx: Select **ON** to use as a Java regular expression or **OFF** to use as a literal string.
Beginning (RegEx)	If you want some initial characters in each line to be ignored, enter a regular expression for it. If a segment at the beginning of the line matches this regular expression, it is ignored. For example, if a line starts with Login and then followed by keyvalue pairs, then if you enter Login in this field, the first word Login is ignored when extracting columns. Named groups in the regular expression are extracted as columns. **Note:** For sending logs through UDP, when you create a new data model, type `.?.?.?` in the Beginning (RegEx) field so that LogLogic LMI can parse the logs correctly.
Ending	To ignore some characters at the end of each line, enter a

Field	Description
(RegEx)	regular expression for those characters. If a segment at the end of the line matches this regular expression, then it is ignored. Named groups in the regular expression are extracted as columns.
Predefined Columns	Used to define a fixed list of columns to be parsed. If predefined columns are specified: • The Key-value parser parses only the specified columns from logs. • The value in the **Last key** field is ignored. • The values in the **Values separator** and **Key-value separator** fields are considered as string literals. The Regex option is not supported. This field is useful when the column names are more than one word and separator is a space. For example, for the log: `Account Name:acc1, Account Domain:loglogic, Caller Computer Name:dell` specify "Account Name", "Account Domain", and "Caller Computer Name" in the **Predefined Columns** field to have the columns and their values extracted correctly.
Last key	Enter a key name. Whenever that key is found in a line, the parser stops searching for more key-value pairs in that line and the value for that key is the remaining content of the line. For example, if the line ends: Severity="high",EventSubClass="1",ObjectID="389576426" then if you specify Severity as the last key, then the value for severity is: `",EventSubClass="1",ObjectID="389576426".`

Field	Description
	Note: To specify a <space>, enter \s (backslash followed by s). For a <tab>, enter \t (backslash followed by t).
Expression	The expression uses a key name preceded with “$” to extract the value for the column. For example, $user is the value of the key "user" in the log line or empty if the key is not present.

Back to Adding a Parsing Rule in an Advanced Data Model

JSON Parser

Field	Description
Beginning (RegEx)	Beginning (RegEx): If you want some initial characters in each line to be ignored, enter a regular expression for it. If a segment at the beginning of the line matches this regular expression, it is ignored. For example, if a line starts with Login and is followed by array or objects, then if you enter Login in this field, the first word Login is ignored when extracting columns. Named groups in the regular expression are extracted as columns. **Note:** For sending logs through UDP, when you create a new data model, type `.?.?.?` in the Beginning (RegEx) field so that LogLogic LMI can parse the logs correctly.
Ending (RegEx)	To ignore some characters at the end of each line, enter a regular expression for those characters. If a segment at the end of the line matches this regular expression, then it is ignored. Named groups in the regular expression are extracted as columns.
Root element	The starting node of the JSON. In case of structured JSON, the root element specifies the starting column in the JSON structure. If the field is empty, all the keys within the JSON are mapped to individual columns. For example:

<table>
<tr><th>Field</th><th>Description</th></tr>
<tr><td></td><td>
<table>
<tr><th>Log</th><th>Root element</th><th>Meaning</th></tr>
<tr><td>{"key1":"value1","key2":["value2","value3"]}</td><td>key2[0]</td><td>Start parsing from 0th element of key2 array</td></tr>
<tr><td>{"key1":"value1","key2":{"key3":"value3"}}</td><td>key2.key3</td><td>Start parsing from inner structure key3</td></tr>
</table>
</td></tr>
<tr><td>New Line Delimiter</td><td>This field is required when the parser operates on multiline logs, especially those arriving from TIBCO LogLogic® Universal Collector. It is recommended to assign the same delimiter that is set in LogLogic® Universal Collector so that the parser removes the delimiters and the logs are parsed successfully.</td></tr>
</table>

Back to Adding a Parsing Rule in an Advanced Data Model

XML Parser

Field	Description
Beginning (RegEx)	Beginning (RegEx): If you want some initial characters in each line to be ignored, enter a regular expression for it. If a segment at the beginning of the line matches this regular expression, it is ignored. For example, if a line starts with Login and then followed by array or objects, then if you enter Login in this field, the first word 'Login' is ignored when extracting columns. Named groups in the regular expression are extracted as columns. **Note:** For sending logs through UDP, when you create a new data model, type `.?.?.?` in the Beginning (RegEx) field so that LogLogic LMI can parse the logs correctly.
Ending (RegEx)	To ignore some characters at the end of each line, enter a regular expression for those characters. If a segment at the end of the line matches this regular expression, then it is ignored. Named groups in the regular expression are extracted as columns.

Field	Description
Root path	If you need to extract a portion of the XML log, you can provide the starting point from a specific hierarchy within the XML log. If you leave the field empty or provide "/", the parser parses the entire XML log. To provide a starting point, use "/" as the separator between elements. For example, `/files/fileInfo_1/location`. If the XML log contains sibling elements with same name, you can address them using "[index]". For example, `/files/fileInfo[1] or /files/fileInfo[2]`. If an XML element has attributes, the attribute name is also separated by underscore (_) in the column name. For example, if the XML log is:

```
<files>
  <fileInfo sizeUnit = "kb">
     <fullName>/vaibhav/data.txt</fullName>
     <fileName>vaibhav.txt</fileName>
     <location>/vaibhav</location>
  </fileInfo>
  <fileInfo sizeUnit = "kb">
     <fullName>/shane/data.txt</fullName>
     <fileName>shane.txt</fileName>
     <location>/shane</location>
  </fileInfo>
</files>
```

Root path	Meaning
(empty)	Start parsing the entire XML from the root element
/files/fileInfo [1]/location	Start parsing from the `location` element of the first `fileInfo` element.
/files/fileInfo [2]	Start parsing from the second element `fileInfo`

New Line Delimiter	This field is required when the parser operates on multiline logs, especially those arriving from TIBCO LogLogic® Universal Collector. It is recommended to assign the same delimiter that is set in LogLogic® Universal Collector so that the parser removes the delimiters and the logs are parsed successfully.

Back to Adding a Parsing Rule in an Advanced Data Model

Columnar Parser

Field	Description
Separator	Enter the delimiter that you want to use as a column separator. The separator can be a string of one or more characters, or a Java regular expression. The delimiters are case sensitive. For example, `bob,windows` where comma (,) is a character used to separate two columns.
RegEx	Use this option to define how the separator should be interpreted. Select ON to use as a Java regular expression or OFF to use as a literal string.
Escape character	Define a character that is actually used to escape the character used as a column delimiter. The delimiters are case sensitive. For example, if you use a comma as a column separator and your column value has a comma in it, then that value has to be escaped so that a parser does not think that the instance of the comma is the start of a new column.
Beginning (RegEx)	If you want some initial characters in each line to be ignored, enter a regular expression for it. If a segment at the beginning of the line matches this regular expression, it is ignored. For example, if a line starts with Login and then followed by columnar data, then if you enter Login in this field, the first word Login is ignored when extracting columns. Named groups in the regular expression are extracted as columns. **Note:** For sending logs through UDP, when you create a new data model, type `.?.?.?` in the Beginning (RegEx) field so that LogLogic LMI can parse the logs correctly.
Ending (RegEx)	To ignore some characters at the end of each line, enter a regular expression for those characters. If a segment at the end of the line matches this regular expression, then it is ignored. Named groups in the regular expression are extracted as columns.
Max columns	Enter the maximum number of columns to be extracted. If more columns than maxColumns are found, then the content of the additional columns is included in the last column. For example, if the separator is <space> and the maxColumns value is 3 for a message like “a b c d”, then there are 3 columns with values “a”, “b” and “c <space> d”.

Field	Description
Trim values	If defined **ON**, then the extra (white) space that is generated at the beginning and end of the column is removed. If defined **OFF**, the extra space is not removed.
Expression	The expression uses the $<n> identifier where n is the column number for the value of column n. For example, $2 is the value of the column "2".

Back to Adding a Parsing Rule in an Advanced Data Model

Regex Parser

Field	Description
Regex pattern	Make sure to enter a valid PCRE regular expression that contains the groups (named or unnamed) to extracted into column values from the log event. Also, it is good practice to use one or more sample events to validate your regular expression and make sure that the correct values are extracted from the event. For a list of supported regular expression meta characters, based on Java regular expressions, see Supported Regular Expression Characters. For example, `(?<Sequence>\d+).*(?<ACL>\%\w+ \-\d\-\w+)\:\s(?<Name>\w+)\s(?<Version>\w+) \s(?<Status>\w+)\ s(?<Protocol>\w+)\s(?<SourceIP>\d{1,3}\.\d{1,3}\.\d{ 1,3}\. \d{1,3}).*(?<DestinationIP>\d{1,3}\.\d{1,3} \.\d{1,3}\.\d{1,3}).*` This extracts 8 fields: Sequence, ACL, Name, Version, Status, Protocol, SourceIP, and DestinationIP.
Expression	The columns are extracted using the capturing group pattern the named capturing group pattern or a combination of both. If you select the parser and the column list is empty, the parser tries to guess columns from the sample data.

Back to Adding a Parsing Rule in an Advanced Data Model

CEP Parser

Field	Description

Expression	Based on the ArcSight Extension Dictionary, the CEF header columns are extracted and the remaining data is formatted as key-value pairs. For example,
	`Sep 19 08:26:10 host CEF:0|Security|threatmanager|1.0|100|worm successfully stopped|10| src=10.0.0.1 dst=2.1.2.2 spt=1232`
	This extracts these columns and their values as follows:
	`$cefVersion=0, $cefDeviceVendor=Security, $cefDeviceProduct=threatmanager, $cefDeviceVersion=1.0, $cefSignatureID=100, $cefName=worm successfully stopped, $cefSeverity=10, $sourceAddress=10.0.0.1, $destinationAddress=2.1.2.2, $sourcePort=1232`

Back to Adding a Parsing Rule in an Advanced Data Model

Managing Columns and Data Types

From the **Review configuration** page, you can update columns and data types for the associated data model.
You can also review column statistics for each defined parsing rule.

Procedure

1. In the **Columns** panel, all system and custom columns are displayed. To add or remove columns, select the check box next to the column name. You can update any column name and type.
    - Name: The name of the column that is displayed in the results. To add or remove a column, click the check box next to the column name. To update the name of a user-defined column, click the column name.
        - The name can include letters, numbers, or underscore (_).

- When used in an Advanced Search query or an expression, the name must be enclosed in square brackets ([ ]) in any of the following scenarios:

Name	Examples
starts with a number	[123] [12model] [1col_0]
contains a hyphen	[data-a] [-col1] [aBc-]
contains all numbers	[1234]
contains a space	[abc model]

- A column name cannot include a period (.).
- Two columns cannot have the same name.
- You cannot use reserved keywords as the column name. For a list of reserved keywords, see Reserved Keywords.

- Type: The data type of the column. Click in the column to add or update the supported data types. Select the data type from the list.
- Parser rules: The rule name that includes the defined column.

2. Select the **Show system columns** check box to show all system columns.

   By default, some system columns are selected. If the check box is not selected, only the user defined columns and some default system columns are displayed. For a list of system columns, see Types of Columns.

3. After modifying column list, click ↻ to refresh the **Parser preview** panel to view all extracted columns and their data types for the defined parsing rule.

   - If you have configured downstream parsing, then the columns extracted from the parent rule and the downstream parser are displayed.

- To change the supported data type for custom columns, click in the **Type** field, and select the data type from the list.

The **Match statistics** panel displays overall information about events that are matched by specified parsing rules. It displays how many rules are enabled, how many columns are extracted by the rule, and how many events are matched with each rule.

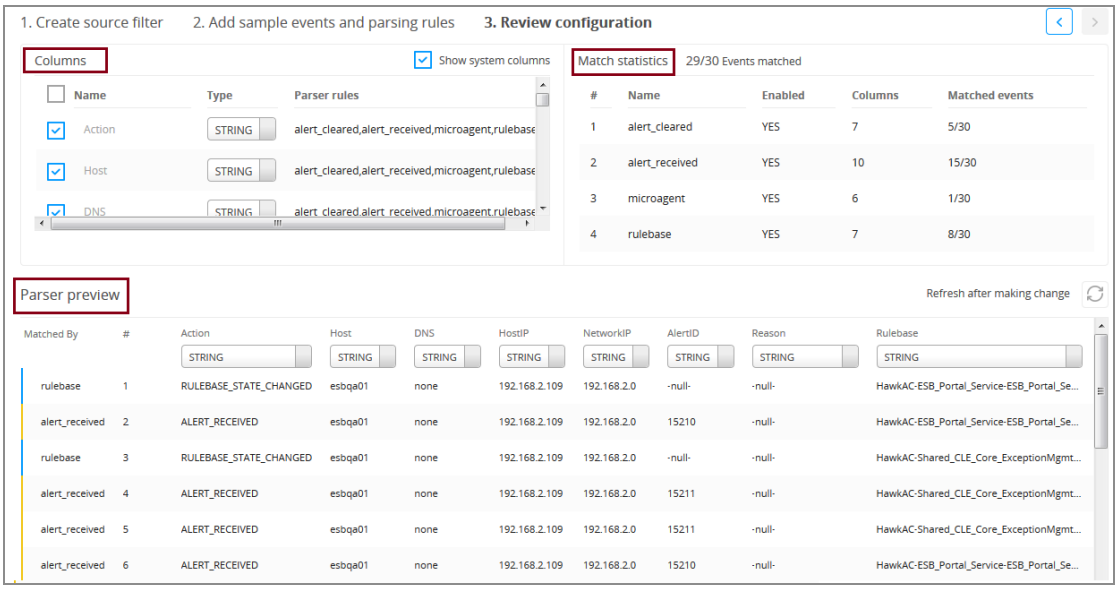

4. Click **Save** to save the data model.

What to do next

You can use the data model in advanced search (**Search > Advanced Search** menu).

Creating a Data Model in Raw Mode

You can add a new data model that can be activated to analyze results in the normalized format. All enabled models can be searched using the source filter from the Search tab.

Before you begin

This option is for advanced users who understand JSON syntax to create a new parsing rule. If otherwise, use the graphical mode to create new data model. For details, see Creating a Data Model in Graphical Mode.

Procedure

1. Go to **Management > Advanced Features > Data Models**.
2. On the **Data Models** page, click **Create a New Data Model**.
3. Click **Switch to raw mode** to add a new model in raw mode.
4. In the **Sample events** panel, paste the sample events to analyze data in normalized format. This data can be helpful in defining the parsing rule based on the log source. Once you add the data model, the sample data is always available when editing the same data model or associated parsing rules.

   Note: You can paste maximum of 100 KB sample data.

5. In the **Raw configuration mode** panel, enter the parsing rule. Ensure to define source filter, parsing rule, and parser properties in a valid JSON syntax. The content of your JSON file might vary depending on the fields selected and configured for your parsing rule.

   Note: For an example of the JSON file and syntax, you can open an existing data model and view its parsing rule in the raw mode.

6. Click **Validate** to ensure that the rule syntax is valid. Click **Format** to format the JSON.
7. Click to refresh the **Parser preview** panel to view all extracted columns and their data types that are matched by the defined parsing rule. Click in the **Type** field to change the supported data types and select the data type from the list.

   Note: This option is available only when the data is pasted in the **Sample events** panel and at least one parsing rule is enabled.

8. Click **Save** to add a new data model. The Data Models page displays the newly added model.

Editing Data Models

You can edit existing data models or save the same model as a new one.

Note: You cannot edit the *system* data model and built-in data models. For the list of built-in data models, see the Supported Log Sources list in *TIBCO LogLogic® Log Source Packages Installation and Upgrade*, which is available on the TIBCO eDelivery website or TIBCO Support website after logging in.

Procedure

1. Go to **Management > Advanced Features > Data Models**.
2. On the Data Models page, click the model name that you want to update.

   The **Details** panel opens on the right side of the page.
3. Click the **Edit** link to edit the data model.

   For detailed information, see Creating a Data Model in Graphical Mode.
4. Click ↻ to refresh the **Parser preview** panel to view all extracted columns and their data types that are matched by the corresponding parsing rule. Each event that matches with the corresponding rule are identified in the same color for easy readability.
5. Perform one of the following:
   - To save the updated information, click **Save**.
   - To save the data model as a new one, click **Save As**. Enter the new data model name in the **Name** field and click **OK**.

   The page is updated immediately.

Managing Parsing Rules

You can manage parsing rules within a data mode when you edit the data model.

Task	Steps
Define the sequence of parsing rules	See Defining the Sequence of Parsing Rules.
Create a new parsing rule	See Adding a Parsing Rule in an Advanced Data Model.

Task	Steps
Edit a parsing rule	See Editing Parsing Rules in Advanced Data Models.
Duplicate	Select the parsing rule and click the duplicate icon .
Delete	Select the parsing rule and click the delete icon . **Note:** You cannot delete parsing rules that are defined for the *system* data model and LogLogic LMI built-in data models.
Reorder	You can a parsing rule (move up or down), to change the sequence of the rules within the data model. **Note:** You cannot move a parsing rule from one data model to another.
Rename	Double-click the parsing rule, rename the rule, and click the **Save** icon .
Enable or disable	Select the parsing rule and click **Enabled**.

Defining the Sequence of Parsing Rules

When multiple parsing rules are defined for a single data model, you can set the rule sequence.
Columns are extracted per the first rule definition that matches the event and subsequent rules are not applied. For example, if Rule1 matches some of the data, it is used to extract column values. If Rule1 fails to match with your data, Rule2 is applied, and so on.
Therefore, sometimes you might want to change the sequence of the parsing rules for the data model.

Procedure

1. In the **Parsing rules** panel, hover over the rule row near the drag icon and the cursor turns into a hand, which you can use to drag the row up or down to change the sequence.

2. To refresh the **Parser preview** panel to view all extracted columns and their data types that are matched by the corresponding parsing rule, click the Refresh icon . Each event that matches with the corresponding rule is identified in the same color for easy readability.
3. To save the data model, click **Save**.

Editing Parsing Rules in Advanced Data Models

You can add or modify parsing rules within the data models that use one of the advanced data model parsers. See Types of Parsers in Advanced Data Models.

Procedure

1. In the **Parsing rules** panel, double-click the rule that you want to update.
2. In the **Edit parsing rule** panel, update the rule information. For details about each field, see Adding a Parsing Rule in an Advanced Data Model.
3. Click the **Refresh** icon to refresh the **Parser preview** panel to view all extracted columns and their data types that are matched by the corresponding parsing rule. Each event that matches with the corresponding rule is identified in the same color for easy readability.

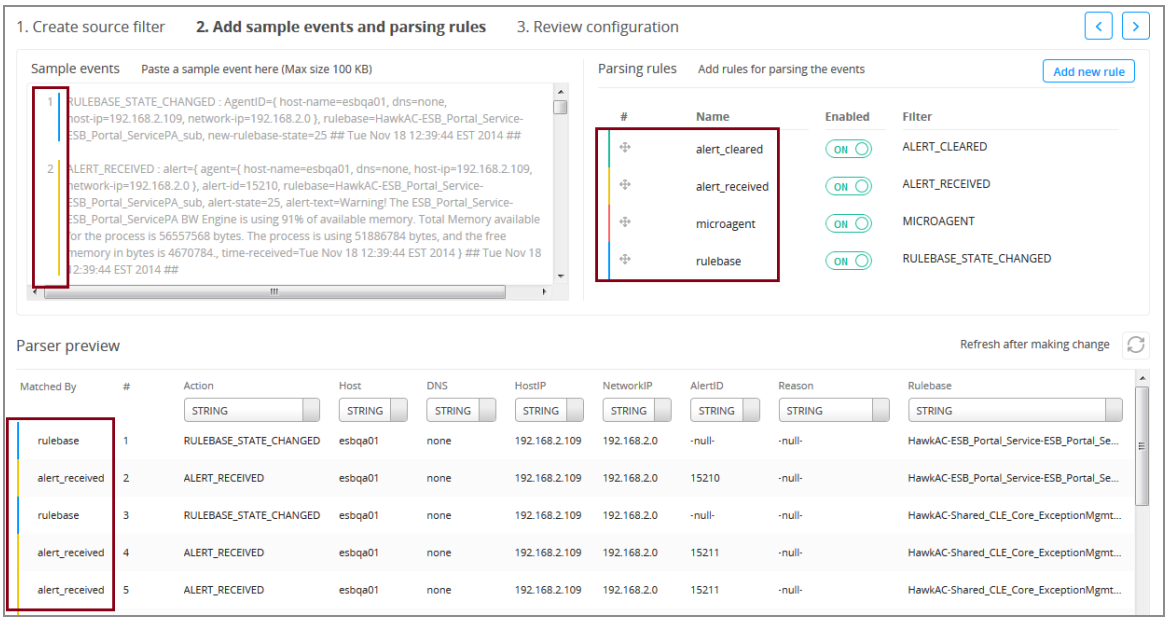

4. Click the **Save** icon to save the updated information. The Parsing rules panel is

updated immediately.

GP Parser-Based Data Models

The data models in LogLogic LMI can be broadly classified into the following categories:

- Advanced Data Models, which use parser types such as syslog, Regex, JSON, and so on.
- GP Parser-Based Data Models, which are data models that use a grouped-pattern parser (GP parser). GP parsers are especially efficient in handling complex Regex parsing rules that work on heterogeneous free-text logs.

GP parsers are very different in structure and syntax from the advanced data model parsers. The parsing rules in a GP parser use Regex patterns, and it is recommended that you have preliminary knowledge of Regex patterns. GP parsing rules have been designed to handle a large set of such complex rules. You cannot create a new data model or edit a built-in data model that uses the GP parser. However, you can duplicate a built-in data model, and then edit its parsing rules as required.

Usage: GP parser-based data models

Sometimes, if the format or pattern of the logs received in LogLogic LMI are slightly different than expected, then the logs are not parsed with the parsing rules defined in the built-in data model. Before LogLogic LMI version 6.4.0, you could not identify which parsing rules failed to match, or which rules were the closest match for such logs, or to edit the parsing rules expressions to accommodate the parsing of the changed log format.

Starting version 6.4.0, you can:

- View a graphical representation of the entire node tree of parsing rules.
- Identify or filter out rules sharing a given node.
- Validate a sample log event against the built-in parsing rules and inspect the graphical view to identify at which node the parsing failed, or the closest matching rule that you might need to tweak.
- Edit the rule expression as required, and save the modified rule or save it as a new rule. See Editing GP Parser-Based Data Models.

Note:

- To edit parsing rules of advanced data models, see Editing Parsing Rules in Advanced Data Models.
- For a list of parsers in advanced data models, see Types of Parsers in Advanced Data Models.

List and Graphical Views of GP Parsing Rules

GP parsing rules are displayed as a list on the **Rule list** tab and as a node tree on the **Node tree** tab.

Rule list tab

The **Rule list** tab displays a list of rules as a table. Each row in the table displays the rule number, its rule ID, and its Regex pattern. After clicking the **Edit** icon , the page displays a table that includes the rule ID and pattern; followed by basic rules and the event descriptor fields. To reorder (move up or down), edit, duplicate, or delete a rule, click the corresponding icon in the rule row.

Node tree tab

The **Node tree** tab displays the list of rules as a tree. Initially, in the node tree, the first node is green and all others are gray. After parsing a sample log that you provide in the **Sample event** field, the colors of the nodes change - depending on whether the node was visited; or whether the node pattern matched or failed.

The state of each node is displayed in a tooltip box after hovering over the node and also represented by the following colors:

- Grey: The node has not been visited during the parsing process
- Green: The node has been visited and matched
- Red: The node has been visited but pattern matching failed

Note:
- The colors red and green reflect the node parsing status when a sample log has been provided.
- In the collapsed state, the colors of the nodes are darker than in the expanded state. For example, dark green in collapsed state; but light green in expanded state.

In the node tree area, you can perform the following actions:

- To view node information: If you hover on a node, the node ID, pattern, and rule pattern of that node are displayed as a tooltip box.
- To zoom in or out of the tree: Move the mouse roller or use the mouse trackpad in the tree area.
- To move the tree: Click the background and drag the tree area (keep the mouse button pressed while dragging).
- To filter rules by node: Click the tooltip box of a node. The view switches to the **Rule list** tab; displaying a filtered list of rules. To clear the filters and view the entire rule list, click the **Refresh** icon  .
- To collapse or expand a node: Click the node. This action triggers an update of the tree layout and changes the view of the node tree. Expanding a node zooms in on the tree; collapsing a node zooms out and might display the entire tree in compact form.

Note: In the collapsed state, the colors of the nodes are darker than in the expanded state. For example, dark green in collapsed state; but light green in expanded state.

Editing GP Parser-Based Data Models

Starting from LogLogic LMI 6.4.0, you can edit the parsing rules of data models that use the GP parser. Because these are built-in data models, you cannot edit them directly. You must duplicate a built-in data model, save it in the **User** group or any other user-created group, and then edit the parsing rules of your new data model.

Similarly, you cannot create a new GP parsing rule. You can edit an existing rule, and then save it either as the same rule or as a new rule (the new rule is automatically assigned a new rule ID). By default, the GP parsing rules are enabled, and cannot be disabled.

> ⚠ **Caution:** The editing option is intended for advanced users who understand the complex Regex patterns. For assistance in editing GP parser rules, contact TIBCO Support.

Editing parsing rules involves the following steps:

1. Defining a source filter.

   This procedure is the same as that for advanced data models. See Defining a Source Filter.
2. Editing GP Parser Rules
3. Reviewing the Data Model Configuration

Editing GP Parser Rules

To edit a GP parser rule, you must add a sample event and then perform the following steps.

Procedure

1. In the **Sample events** box, paste a sample event.

   This data can help in defining the parsing rule based on the log source. After saving the data model, the sample data is always available when editing the same data model or associated parsing rules.

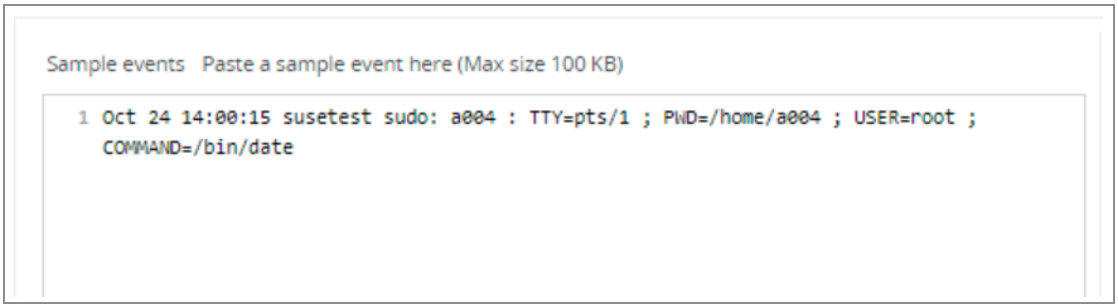

Note:

- You can paste a maximum of 100 KB of sample data.
- The tree view shows the node status only for the first event in the **Sample events** box.

2. Click the **Refresh** icon . The **Parser preview** panel displays all extracted columns and their data types that are matched by the corresponding parsing rule.

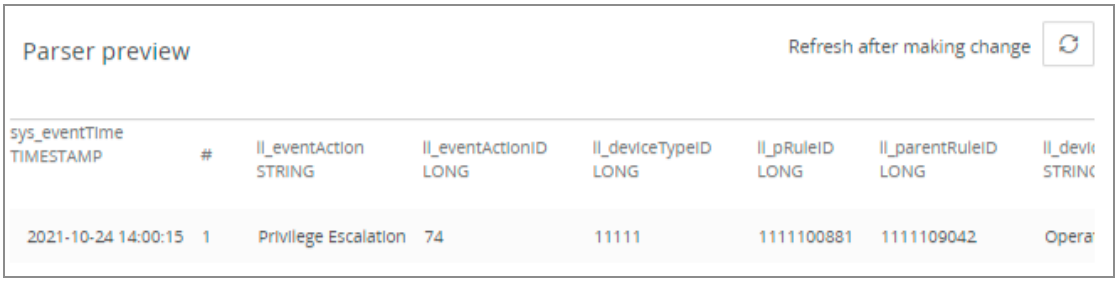

3. Click the **Node tree** tab to see the tree view. The node tree displays nodes with different colors:
   - Grey: The node has not been visited during the parsing process
   - Green: The node has been visited and matched
   - Red: The node has been visited but pattern matching failed

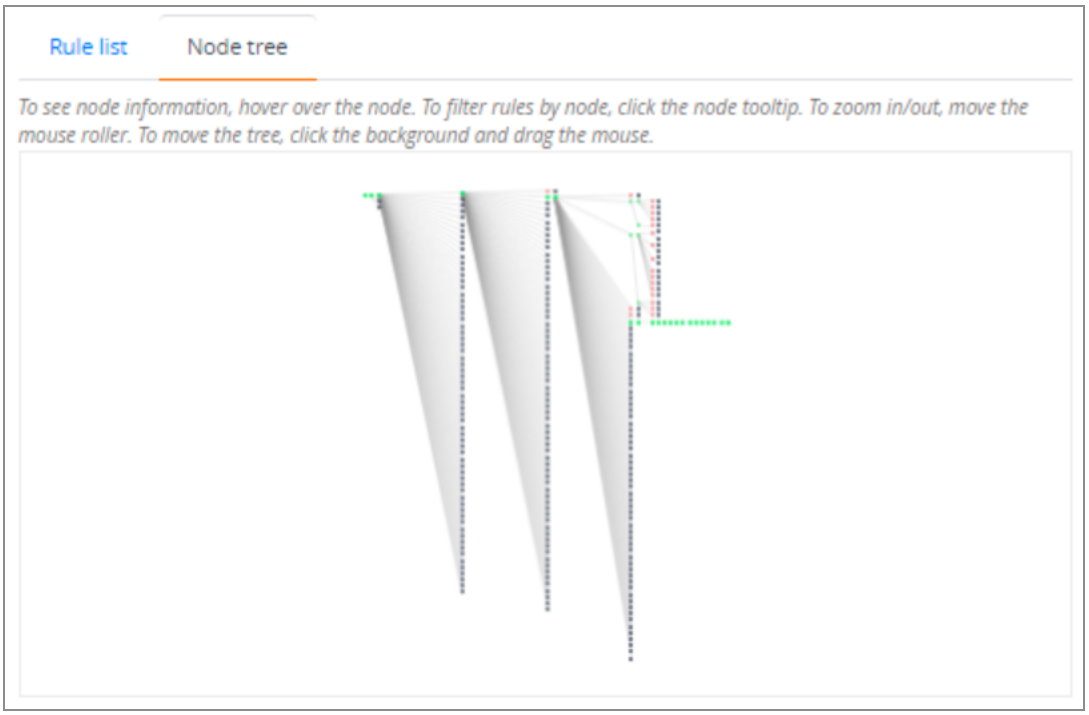

Note:

- The colors red and green reflect the node parsing status when a sample log has been provided.
- In the collapsed state, the colors of the nodes are darker than in the expanded state. For example, dark green in collapsed state; but light green in expanded state.

4. To zoom in or out of the view, move the mouse roller till you have a better view of the matching node. To view the node information of a node, hover over the node.

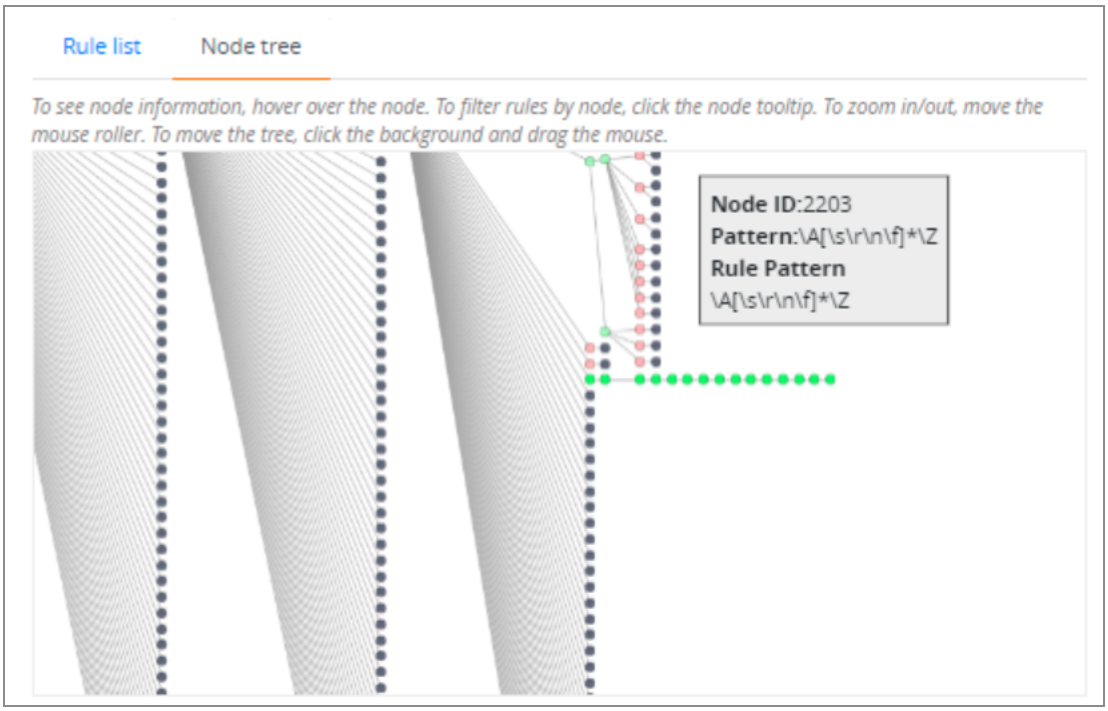

5. After clicking the tooltip box, the view switches to the **Rule list** tab, and only those parsing rules that share the node you clicked are displayed. Scroll to the rule you want to edit and click its **Edit or duplicate** icon .

6. In the **Edit or duplicate rule** section, you can edit any fields of the parsing rule except the **Rule ID** field. For detailed information about the syntax, description, and format rules, see the Field Reference section.

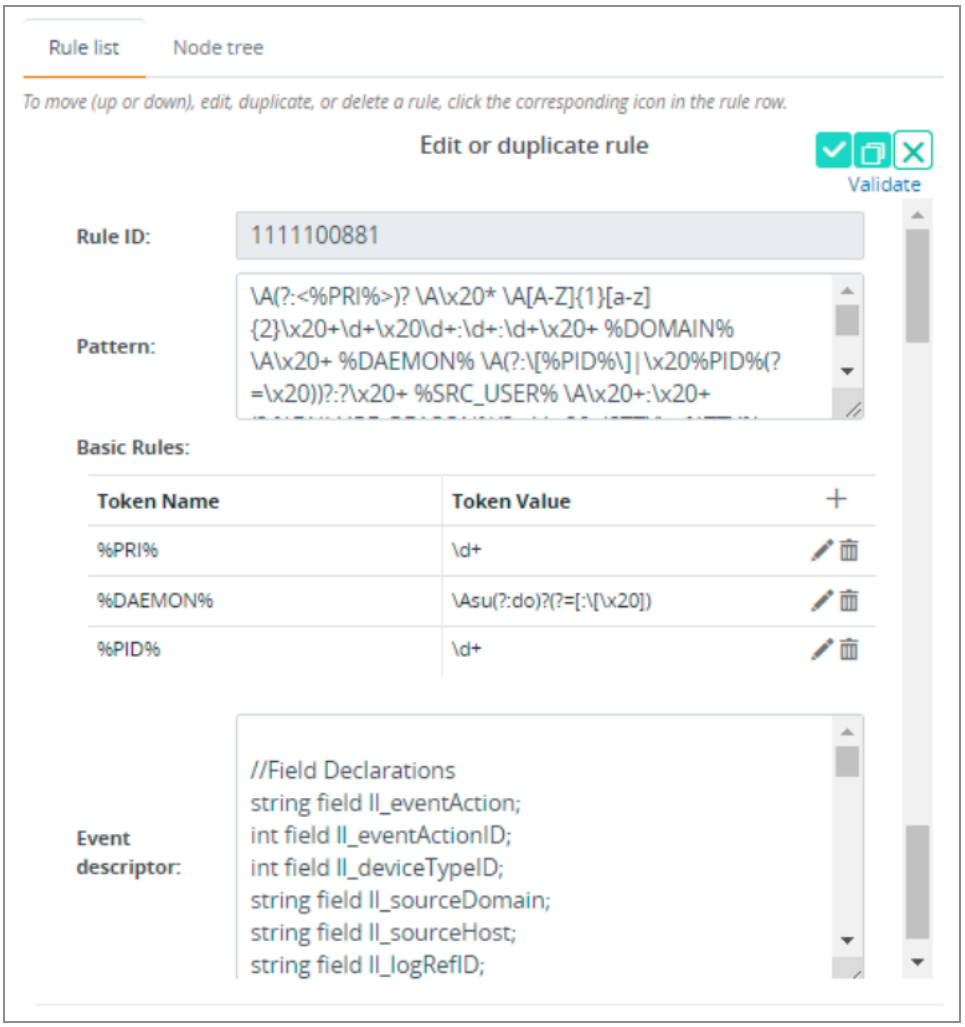

7. After modifying the required fields, click **Validate**. If validation of the rule is successful, a message is displayed at the top of the section. If not successful, an error message is displayed, providing the reason for failure.

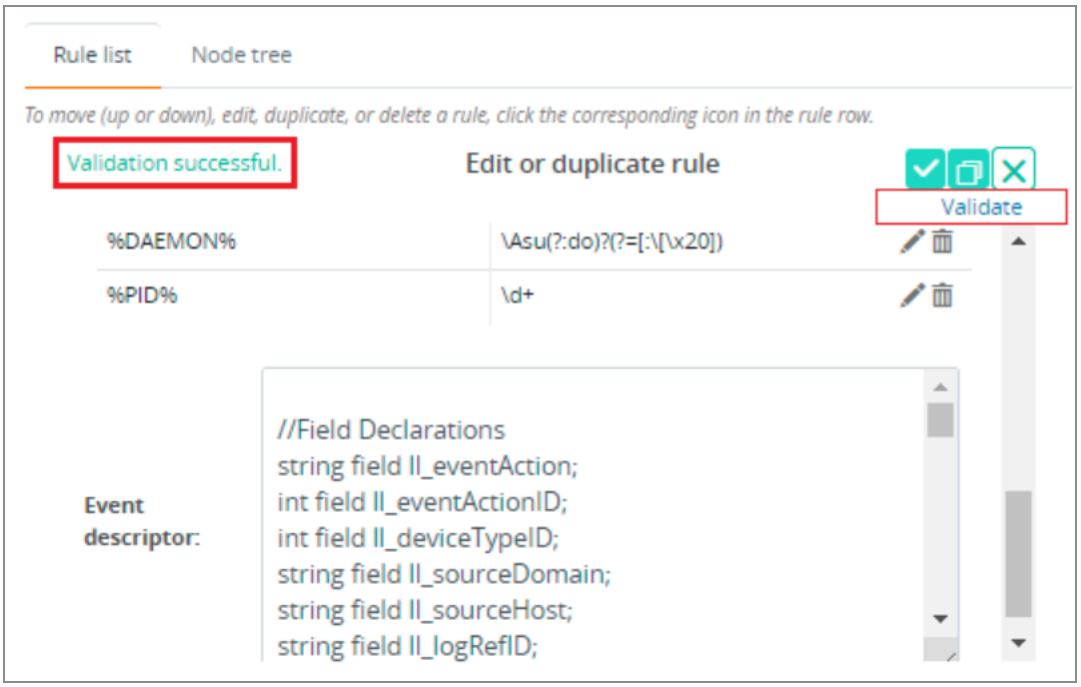

8. You can either save the changes in the same rule or save the rule as a new rule (duplicate the rule):
   - To save the changes in the same rule, click the **Save** icon .
   - To save a copy of the rule, click the **Save as** icon .
   - To cancel the changes, click the **Cancel** icon .

> **Note:**
> - The **Save** and **Save as** icons are enabled only after making changes in at least one field (except the **Rule ID** field) and if validation is successful.
> - Saving the parsing rule also saves the data model.

What to do next

Review the configuration and save the changes to the data model. See Reviewing the Data Model Configuration.

Reviewing the Data Model Configuration

From the **Review configuration** page, you can update columns and data types for the associated data model. You can also review column statistics for each defined parsing rule.

Procedure

1. In the **Columns** panel, all system and custom columns are displayed.

Field	Description
Name	The name of the column that is displayed in the results.
Type	The data type of the column.
Parser rules	The rule name that includes the defined column.

2. Select the **Show system** columns check box to show all system columns. By default, some system columns are selected. If the check box is not selected, only the user-defined columns and some default system columns are displayed. For a list of system columns, see Types of Columns.
3. The **Match statistics** panel displays overall information about events that are matched by specified parsing rules. It displays how many rules are enabled, how many columns are extracted by the rule, and how many events are matched with each rule.

What to do next

You can use the data model in advanced search (from the **Search > Advanced Search** menu).

Field Reference

This section explains the fields in a GP parser. You can edit any fields of the parsing rule except the **Rule ID** field. For information about how to edit these fields, see the Editing GP Parser Rules section.

A GP parsing rule is a sequence of pattern-matching nodes that parse a log. Each rule is made of pattern-matching nodes and event descriptors. Each node includes a set of basic

rules. The basic rules contain token definitions, which are the patterns to be used to match and capture substrings from the log. Such substrings are then assigned to the specified database columns. The definition of which substrings are stored in which database columns is stored in the **Event Descriptor** field. If a log matches the entire node sequence, an event is generated as defined in the **Event descriptor** field.

The following table describes the fields in a GP parsing rule:

Field Name	Description	Format
Rule ID	The unique identifier of the rule in a data model. It is displayed in Advanced Search results. The rule ID is an auto-generated number and cannot be modified.	Numeric, integer. Format : [0-9]+
Pattern	The Pattern field contains the rule definition.	It is a sequence of Regex or token nodes (%SOMETHING%). Note: The nodes are separated by a space character (' '). Therefore, to use a space character in the pattern, you must use \x20 or \s in the pattern.
Basic Rules	The **Basic Rules** field contains the list of token names and their corresponding patterns. The token names can be reused in the **Event Descriptor** field to extract substrings. These substrings can be used as function variables or field values. See the Event Descriptor Language Reference section.	%('A' .. 'Z' | 'a' .. 'z' | '_' | '0' .. '9')+%
Event descriptor	This field includes a set of columns with the values extracted from the log after the parsing rule completes parsing the log. The columns and their values are displayed in the **Parser preview** pane. This is a very large section, and includes other subsections.	You must follow the language syntax as described in the Event Descriptor Language Reference section.

Event Descriptor Language Reference

The event descriptor language is used in the parser rule file. Similar to other programming languages, it includes declarations, statement types, array constructs, and so on.

The following table describes the terms used in this section.

Term	Used in	Format	Example
Token	Used in the **Basic Rules** and **Event Descriptor** fields. A token is a pattern-capturing group. You can use a token to reference the strings in the **Event Descriptor** field.	%('A' .. 'Z' | 'a' .. 'z' | '_' | '0' .. '9')+%	ll_eventID := %EVENTID%
Static string	Used in the **Event Descriptor** field for constants	Any character sequence between double quotes. (escape character: '\')	"Some string", "with a \""
Field or temporary variable	Used in the **Event Descriptor** field	LETTER (LETTER | DIGIT)* Where LETTER includes upper case letters ('A' .. 'Z'), lower case letters ('a' .. 'z'), or underscore ('_')	ll_eventID
Function call	Used in the **Event Descriptor** field	FUNCTNAME(expressions) Where the format of FUNCTNAME is: LETTER (LETTER | DIGIT)*	concat2(“prefix_”, %EVENTID%)
Expressions	Used in the **Event Descriptor** field for passing arguments to function calls.	It can be a one-to-many static string, token, function call, or temporary field. When specifying a list of	• "string" • %TOKEN% • join("a", "b")

Term	Used in	Format	Example
		expressions, the separator is comma (,).	

Basic Operations

This section describes the syntax and format of the following operations:

- Declare a field.
- Declare a temporary variable.
- Assign a value to a field or variable.
- Use conditional statements. Iterate over an array.
- Generate multiple events for a log.

Operation	Syntax	Example
Declare a field	`TYPE field IDENT;` • TYPE is the type of the field: int, string, inetAddr, timestamp • IDENT is the field name. It can be either a function or an enrichment list (user-defined mapping function). Format: LETTER (LETTER | DIGIT)*	`int field ll_eventID;`
Declare a temporary variable	`temp IDENT`	• A simple temp variable: temp variable;

Operation	Syntax	Example
	`<ARRAY_MARKER>;` • IDENT is the name of the temp variable. • ARRAY_MARKER [] indicates the variable as an array. Optional.	• An array variable: temp array_var[];
Assign a value to a field or temporary variable	`IDENT := expression;`	• ll_eventID := %EVENTID%; • ll_eventStatusID := 1; • ll_eventAction := "some string value"; • ll_eventActionID := actionidentification (%EVENTID%);
Conditional statement	`if ( expression ) { statements } ( else ( if_statement | ({ statements } ) ) )?`	• if (strstr("success", %OUTCOME%)){ ll_eventStatusID = 1; } • if (strstr("success", %OUTCOME%)){ ll_eventStatusID = 1; } else{ ll_eventStatusID = 2}; • if (strstr("success", %OUTCOME%)){ ll_eventStatusID = 1; } else if (strstr("failure", %OUTCOME%)){ ll_eventStatusID = 2; } else{ ll_eventStatusID = 0;}

Operation	Syntax	Example
'Foreach' statement	`Foreach my_var in myTvar { statements }`	`Foreach my_var in myTvar { //create output data : set fields; insert; }`
'insert' statement Used in combination with `foreach`, it can be used to create multiple output data records for a single log.	insert;	`Foreach my_var in myTvar { //create output data : set fields; insert; }`

Functions

The following table describes a few functions, their syntax, and usage information.

Function name	Description
addnum (*expression1, expression2*)	Returns the sum of the expressions.
concat2 (*expression1, expression2*)	Concatenates the two expressions and returns the resultant string.
ipAddress (expression)	Converts the expression into an `inetAddr` value. The expression can be a 32-bit number or a string.
join(*string, expressions*)	Concatenates multiple expressions with the value of *string*.
logAppId()	Returns the app ID of the GP parser rule. It is an internal identifier and cannot be modified.

Function name	Description
logId()	Returns the log identifier.
logRuleid()	Returns the ID of the matching rule.
logsource()	Returns the log source name.
logtime()	Returns the log collection time.
mapRuleAction (*expression*)	Calls the enrichment list function for a data model, using the expression for the specific rule ID: `<Data model name>_mapRuleAction(Rule_ID)` Returns the corresponding value from the enrichment list.
mapRuleStatus (*expression*)	Calls the enrichment list function for a data model, using the expression for the specific rule ID: `<Data model name>_mapRuleStatus(RuleID_expression)` Returns the corresponding value from the enrichment list.
match(*<regex>, expression*)	Returns true if the *<regex>* expression matches *expression*. To use the capturing groups from the *<regex>* expression for further matching, you can use the `matchN()` function. The capturing groups from this function are automatically used as the parameter in the matchN() function.
matchList(*<some_regex>, expression*)	Finds all matches of the *<some_regex>* expression within the *expression* string and returns an array of strings. Use this function when you need demultiplexing.
matchN (*group_id*)	Returns a capturing group from the last match() function call. You can then use the capturing group value in other functions or assign the capturing group value to a variable.

Function name	Description

splitChar (*expression_split, expression*)	Splits the *expression* using the pattern in *<expression_split>* and returns an array of strings. Use this function when you need demultiplexing. It creates multiple, structured data output records for a single log.
	Use the following pattern:
	`foreach element in T_elementList {`
	`  //assign fields values.... ;`
	`  Insert;`
	`  }`

splitRegex(*"some_regex", expression*)	Similar to `splitChar()`, but with a Regex as the split condition.
strequal (*expression_string1, expression_string2*)	Returns true if the expression strings are identical.
strptime(*time_format, expression*)	Convert the *expression* string into a timestamp of the specified time format. For more information about how to specify the time format, see Time format specifiers.
strstr(*<source_expression>, search_expression*)	Returns true if the source expression string contains the search expression string.
upcase(*expression*)	Return uppercase string of the *expression* string.

Time format specifiers

Use the following letters to specify the time format parameter in the `strptime()` function.

Parameter	Description
'A', 'a'	The day of the week, using the locale's weekday names. Specify either the

Parameter	Description
	abbreviated or full name.
'B', 'b'	The month, using the locale's month names. Specify either the abbreviated or full name.
'c', 'C'	Not supported
'd'	The day of the month [1-31]. Leading zeros are permitted, but not required.
'D'	%D is the same as %m/%d/%y.
'e'	%e is the same as %d.
'h'	%h is the same as %b.
'H'	The hour (24-hour clock) [0-23]. Leading zeros are permitted; but not required.
'I'	The hour (12-hour clock) [1-12]. Leading zeros are permitted; but not required.
'j'	The day of the year [1366] Leading zeros are permitted; but not required.
'k'	%k is the same as %H.
'l'	%l is the same as %I.
'm'	The month number [1-12]. Leading zeros are permitted; but not required.
'M'	The minute [0-59]. Leading zeros are permitted; but not required.
'n'	Any white space.
'p'	The locale's equivalent of A.M. or P.M. It is an indicator for the 12-hour clock.
'r'	%r is the time as %I:%M:%S %p.

Parameter	Description
'R'	%R is the time as %H:%M.
'S'	The seconds [0-61]. Leading zeros are permitted; but not required. You can also specify milliseconds (optional).
't'	%t is any white space.
'T'	%T is the time as %H:%M:%S.
'U'	The week number of the year as a decimal number [0-53]. Leading zeros are permitted but not required. Sunday is considered as the first day of the week.
'w'	The weekday as a decimal number [0-6], with 0 representing Sunday. Leading zeros are permitted; but not required.
'W'	The week number of the year as a decimal number [0-53]. Leading zeros are permitted; but not required. Monday is considered the first day of the week.
'x'	Not supported.
'y', 'Y'	The year including the century. For example, 1998.
'%'	%% is replaced by %.

Monitoring Console

The Hawk Console is available in LogLogic LMI as the Monitoring Console. It provides a central view of all the distributed components interacting within the Hawk® environment.

It is easier to manage multiple domains from LogLogic LMI than configuring and controlling the domains individually. From the Monitoring Console, you can:

- Configure and manage domains of TIBCO Operational Intelligence Hawk® 6.2.1 HF2 or later
- Monitor and manage distributed applications and operating systems
- Manage alerts generated by Hawk agents

- Take action in response to predefined conditions

By default, the Monitoring Console is disabled in LogLogic LMI. After a user with administrator access enables it, you can access it from **Monitoring > Console**.

Domains

Monitoring Console provides a ready-to-use Hawk domain named `lmi_domain`, to which you can connect external Hawk agents. By default, `lmi_domain` is registered to the Monitoring Console. You can unregister it by clicking the Unregister icon on the domain card. To register it again, follow the procedure described in Configuring a Hawk domain.

You can also register other external domains and external Hawk agents can connect to the domains.

Related Topics

For more information, see the following documentation:

Topic	Reference
TIBCO Hawk Console	TIBCO Hawk Console User's Guide
TIBCO Operational Intelligence Hawk®	TIBCO Operational Intelligence Hawk® documentation
Enabling the Monitoring Console	Contact your LogLogic LMI administrator.
Configuring a Hawk domain within LogLogic LMI	Contact your TIBCO Operational Intelligence Hawk® administrator.
Securing the communication between the Monitoring Console and other Hawk components	Contact your TIBCO Operational Intelligence Hawk®administrator.

Configuring a Hawk domain

You can register and configure multiple Hawk domains within LogLogic LMI. The Monitoring Console connects to Hawk domains using TCP.

In addition to the default `lmi_domain`, you can configure and register other external domains in LogLogic LMI and external Hawk agents can connect to the domains.

If you have deployed LogLogic EVA on a Docker host, you must use the `proxy` domain type to connect to existing Hawk domains.

Domain type	Description
regular	An external Hawk domain (created outside of LogLogic LMI), to be connected over TCP. Provide a domain name and its TCP transport details.
proxy	An external Hawk domain (created outside of the LogLogic LMI), to be connected using the proxy method. Provide a domain name, URL, and login credentials.
LMI Domain	An internal domain, provided as a built-in, ready-to-use domain. External agents can connect securely to this domain using the `Self/Cluster Manager Host:Port` field.
New Domain	An internal domain, similar to the built-in LMI Domain External agents can connect securely to this domain using the `Self/Cluster Manager Host:Port` field.

Procedure

1. On the Monitoring Console, in the **Domain** section click the plus (+) icon.
2. In the **Configure Domain** dialog box, provide the following information:

   For details about these parameters, see TIBCO Hawk® Console User's Guide.

Field	Description	Default Value
Domain type	Choose one of the following domain types: • proxy • regular • New Domain • LMI Domain **Note:** LMI Domain is the default domain. If you select LMI Domain, all other fields are automatically set to their default values. LMI Domain is available in the list only if it is currently unregistered.	regular
Domain Name	Type a name for the domain	Automatically set to lmi_domain (if you select LMI Domain as the domain type)
Transport	The transport protocol **Note:** Not applicable to a proxy domain.	TCP
• Self Host: Port • Cluster Manager Host: Port	• **Self Host: Port**: Enter the IP address and port of the LogLogic LMI appliance • **Cluster Manager Host: Port**: Enter the IP address and port of the cluster manager	None

Field	Description	Default Value
	Note: These two fields are available when you select the domain type as regular.	
Self/Cluster Manager Host: Port	Enter the IP address and port of the appliance. The port number must be unique and open for bidirectional communication.	• If you select `LMI Domain` as the domain type, the value is set to `<appliance_IP_address>:9688` and the field is not editable. • If you select `New domain` as the domain type, you must enter the appliance IP address and port.
	If you select `LMI Domain` or `New domain` as the domain type, the two fields **Self Host: Port** and **Cluster Manager Host: Port** are combined as the **Self/Cluster Manager Host: Port** field.	

3. To configure a domain that uses SSL-based TCP transport, select the **Additional transport options** check box and provide the values for the following fields:

Field	Description
Key store	Absolute path of the key store that contains the Monitoring Console certificate and key to be loaded while communicating with the Hawk domain. You must provide a custom keystore that has a password-protected key. For example, the default keystore of LogLogic LMI is: `/loglogic/tomcat/conf/keystore`
Key store	Password to access the key store

Field	Description
password	
Key password	Password to access the private key
Trust store	Absolute path to the trust store to be used to validate the Hawk component certificates while communicating with the Hawk domain. For example, the default trust store of LogLogic LMI is: `/loglogic/tomcat/conf/truststore`
Trust store password	Password to access the trust store
SSL protocol (optional)	TLSv1.3 and TLSv1.2 protocols are supported.
SSL Enabled Algorithms (optional)	Comma-separated list of algorithms. Default value: `TLS_RSA_WITH_AES_128_CBC_SHA`
Security Policy	Choose one of the following: • None • Trusted • Trusted with Domains (only for Windows domains) • Custom

4. Click **Configure**.

Result

The domain is configured and displayed in the Domains section.

Note: Whenever the connection of the appliance with the domain is lost, the appliance tries to reconnect. However, if the appliance is not able to reconnect, a reconnect button appears on the Monitoring Console. You can click the button to manually reconnect to the domain.

What to do next

After configuring a Hawk domain, you can connect Hawk agents to it.

Note: When configuring Hawk agents, ensure that you use the domain name, IP address, and port number of the domain in LogLogic LMI.

For information about how to connect Hawk agents to the Hawk domain, contact your Hawk administrator or see TIBCO Operational Intelligence Hawk® documentation.

Groups

You can organize your data into groups. A collapsible pane on the left displays the dashboard groups in a folder-like structure. Groups are available in the following Advanced Features:

- Advanced aggregation
- Advanced dashboards
- Bloks
- Data models
- Enrichment lists

For more information about these features, see Advanced Features.

Shared and Nested Groups

To further organize artifacts, you can create nested groups. When you create a group, by default the **Shared** option is enabled. Setting a group as shared makes it available in other Advanced Features and you can store artifacts relevant to that functionality in the shared group. A group with its **Shared** option disabled is not available in other Advanced Features.

- If a parent group is shared, then that group, along with the nested groups that are shared, are available in other Advanced Features. The nested groups with **Shared** disabled are not available in other Advanced Features.
- If a parent group is not shared, then that parent group and all of its nested groups are not shared.
- The **Shared** option of a group cannot be changed later.
- Shared groups are indicated by the icon  and groups that are not shared are indicated by the icon  .

Note: By enabling the **Shared** option, only the groups are available in other Advanced Features; the artifacts stored in the groups are not available in other features. For example, the User group in the Data Models feature stores data models, but in the Bloks feature stores Bloks.

Built-in Groups

The following built-in groups are available:

Note: User-created artifacts cannot be stored in built-in groups.

Group Name	Contents of the group	Shared? Yes/ No
System	Artifacts created by the system	Yes
User	• New artifacts created by users (non-system users) in LogLogic LMI 6.4.0 • Artifacts created by users in previous LogLogic LMI versions and that are not associated with any group	Yes
LSP	Artifacts related to LogLogic LSP	Yes
Distributed	Distributed aggregation rules that are created	Yes

Group Name	Contents of the group	Shared? Yes/ No
	on a Management Station. On a Remote Appliance, you can see the aggregation rules created on the Management Station in the Distributed group. **Note:** You cannot create a new rule or group in the Distributed group.	
ComplianceSuites	Artifacts related to TIBCO LogLogic® Compliance Suite	Yes
Rules > Aggregation > All Rules	All aggregation rules	No
Advanced dashboards > All Dashboards	All dashboards	No
Bloks > All Bloks	All Bloks	No
Data models > All Data Models	All data models	No
Enrichment lists > All Enrichment Lists	All enrichment lists	No

After Upgrading to LogLogic LMI 6.4.0

In versions 6.3.x, groups were available only in the Advanced Dashboards. In other Advanced Features (Bloks, Data Models, Aggregation Rules, and Enrichment Lists), the corresponding artifacts were displayed on the main page of the feature.

After upgrading to LogLogic LMI 6.4.0, the artifacts from versions 6.3.x are stored in different groups, depending on whether the artifacts are built-in or how they are related to LogLogic LSP.

For details about these changes, see the "Upgrade Considerations" section in *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*.

Managing Groups

You can create, delete, or rename a group from any of the Advanced Features where the group is available.

Task	Steps
View all groups	The following types of views are available: • In the left pane, as a folder-like structure. The left pane can be collapsed by clicking  or expanded by clicking ☰. • On the main page, as a list view. The path of the group is also mentioned in the list. For dashboard groups, a grid view is also available. You can switch to the grid view by clicking the grid icon ⊞ or to the list view by clicking the list icon ☰.
Create a group	1. To create a parent group: click **Create Group**. To create a nested group: click **Create Nested Group**. 2. (Optional) By default, the group is shared across other Advanced Features. If you do not want it to be shared, then disable the **Shared** option. 3. Type a group name. 4. (Only for nested groups) Select the parent group from the list. You can create a new group or select the **User** group, or select any user-created group **Default parent group:** When creating a nested group within any 'All' group (for example, All Rules, All Bloks, and so on), the **User** group is the default group. Otherwise, the current parent group is selected as the default group.

Task	Steps
	5. Click **Save**.
• Delete • Rename	• Delete: Select the group and click the delete icon  . • Rename: Select the group and click the rename icon  . Enter the new group name. **Note:** You cannot delete or rename built-in groups.
• Move • Duplicate	**Note:** You cannot move or duplicate groups.

Universal Lossless Data Protocol Library

The Universal Lossless Data Protocol (ULDP) library provides a means for users to connect to a LogLogic LMI appliance and send logs to it.

Overview

The ULDP protocol is designed to send log data in a reliable and secure manner from their collection point to a LogLogic LMI appliance. TIBCO LogLogic® Universal Collector uses this protocol to communicate with a LogLogic LMI appliance.

Using the ULDP library, you can:

- Send data in raw mode.
- Send encrypted and compressed data.
- Tag files that are sent to LogLogic LMI to allow properly identifying the log source.
- Send log files in such a way that the time stamps are parsed.

The ULDP library is provided in the `ULDPClient` directory of the LogLogic LMI supplemental package. For more information, see the `README.txt` file in the same directory.

How Messages are Transmitted

When messages are sent, they are placed in a queue, waiting for their acknowledgements to arrive. If at some point the size of the queue is more than the maximum value (MaxQueueSize), a mandatory request for acknowledgment is sent and no further messages are sent until the acknowledgement is received, thereby effectively blocking the `sendMessage` call.

The sequence of calls is:

1. Create a `UldpConnectionSettings` object and set its properties. At minimum, the destination host should be set.
2. Create a `UldpSender` object, passing the `UldpConnectionSettings` as a parameter.
3. Call the `connect()` method on the `UldpSender` object.
4. For each log message element to be sent, perform the following steps:

a. Create a `UldpLogMessage` object.

b. Call the `sendMessage()` method of the `UldpSender` object, passing it the `UldpLogMessage` Object.

c. Optionally, call the `flush()` method if nothing is ready to be sent.

5. Call the `close()` method of the `UldpSender` object when desired.

   If the `close()` method raises the exception `PendingAckMessageException`, which means some messages have not been acknowledged by the LogLogic LMI appliance, and are possibly not received properly. Upon the next call to connect, those events is sent again, possibly ending up with duplicates, but avoiding any message loss.

Note: This library is not thread-safe. If multiple threads need to send messages through `sendMessage()` calls, the calls should be synchronized, for example by using:

```
synchronized(sender) {
    sender.sendMessage(...)
}
```

Connection Parameters

The connection properties are hosted in the `UldpConnectionSettings` object. They follow the typical get/set method naming of Java Beans.

Name	Type	Default Value	Description
Host	String	localhost	Part of the constructor
Port	Integer	5515 if secure; 5516 if	Part of the constructor

Name	Type	Default Value	Description
		not secure	
MaxQueueSize	Integer	51200 (50kB)	The size of the queue of non-acknowledged message to retain. The value 0 indicates not retaining messages after they are sent.
UseCompression	Boolean	false	If true, messages are compressed. Better compression is achieved by sending event in batches, by calling `flush()` after the end of the batch.
UseTls	Boolean	false	If true, initiates a TLS connection.
UseEncryption	Boolean	true	Implies that `UseAuthentication` is true. A ciphering algorithm should be used.
KeystorePath	String	N/A	Sets the keystore to be used to find an anchor of trust and a client certificate, if needed.
KeystorePassword	String	N/A	The password to be used to access the keystore and the private keys within it.
TlsProtocolName	String	TLSv1.2	The version of the TLS protocol to use.
CipherName	String	N/A	Name of the stack of crypto protocols to be used: • TLS_RSA_WITH_AES_128_CBC_SHA: If encryption is on • SSL_RSA_WITH_NULL_SHA: If encryption is off

Name	Type	Default Value	Description
			Note: Although encryption is off, a cipher suite is used for authentication. This might cause problems with Java 1.8.0_u201 or later, because null cipher suites are disabled by default.
NoServerAuthentication	Boolean	false	Accept any server certificate. Recommended only for testing purposes, such as when there is no threat of a man-in-the-middle (MITM) attack.
DomainName	String	N/A	Sets the domain name
ignoreHostnameValidation	Boolean	true	When true, the certificate subject is not validated against the host name or IP address used for connection.
addAcceptedCertificateFingerprints	String	N/A	Lists the fingerprints of the accepted certificates, separated by ','. An ASCII label identifying the hash function followed by a colon are concatenated at the beginning of the fingerprint. **Important:** Implementations must support SHA-1 as the hash algorithm and use the ASCII label "sha-1" to identify the SHA-1 algorithm. The length of an SHA-1 hash is 20 bytes and the length of the corresponding fingerprint string is 65 characters. An example certificate fingerprint is: `sha-1:E1:2D:53:2B:7C:6B:8A:29:A2:76:C8:64:36:0B:08:4B:7A:F1:9E:9D`

ULDP Log Types

ULDP defines several log types, each suited for distinct collection technologies and products.

Irrespective of its type, each log message contains at least:

- A timestamp
- A source address (InetAddress)

ULDP log types are:

Syslog Message

This is implemented by the `UldpSyslogMessage` class.

The only specific property is the log message itself.

For the products supported by LogLogic LMI and collected through Syslog, the source type is automatically identified, similar to when logs are transmitted using Syslog to the LogLogic LMI appliance.

Realtime LogFile Message

This is implemented by the `UldpFileTailMessage` class. This type is used for sending files line by line, as they are written into the file. Each line should have its own message.

Two properties must be defined: the content of the line itself and an `appName` (application name). LogLogic LMI uses the `appName` field to identify the source type. LogLogic LMI recognizes the following values for this field:

Value of `appName`	Description
AMXAdmin	TIBCO ActiveMatrix® Administrator
Business Works	TIBCO BusinessWorks™
HawkAgent	TIBCO Operational Intelligence Hawk®
TIBCO	TIBCO Generic
TIBCO ADMIN	TIBCO Administrator™

Value of `appName`	Description
TIBCO AMX BPM	TIBCO ActiveMatrix® BPM
TIBCO APIX	TIBCO® API Exchange
TIBCO ActiveSpaces	TIBCO ActiveSpaces®
TIBCO BE	TIBCO BusinessEvents® Server
TIBCO EMS	TIBCO Enterprise Message Service™ Server
TIBCO SILVER FABRIC	TIBCO Silver® Fabric

Even for unknown products, using this parameter enables easily retrieving different logs belonging to the same application because this value is part of the message that LogLogic LMI ingests.

FileChunk Message

A sequence of file chunk messages to send the content of a file in raw form. The `eof` attribute of the last file chunk must be set to true. A file identifier must be provided, which reflects both the location of the file and its content. If several overlapping chunks with the same `fileIdentifier` are sent, the last ones are ignored. If a file at the same location has different content than before, the file identifier should be changed. One way to achieve this is to create a compound file identifier with `<file path>:<checksum>`. Using this type of message is the only way to send logs in which the dates are extracted from the log message and are not assumed to be the current date, as is the case with other ULDP messages. This mechanism is similar to the file-pull mechanism in LogLogic LMI. The maximum size of a file chunk is 50 KB.

The following table lists the file format of supported file devices, when files are to be collected via LogLogic Universal Collector or via the ULDP library when sent over ULDP.

File format of supported file devices

FormatType	Name
0	Cisco ACS Failed Attempts

FormatType	Name
1	Cisco ACS Passed Authentication
2	Cisco ACS RADIUS Accounting
3	Cisco ACS TACACS+ Accounting
4	Cisco ACS Administration Audit
6	Microsoft IAS
7	Microsoft ISA Web (W3C)
9	Generic W3C
10	Others
11	W3C (NetCache)
12	W3C (BlueCoat)
13	Squid Native
14	MS Exchange 2003 Tracking Log
15	MS Exchange 2000 Tracking Log
16	MS Exchange 2003 SMTP (W3C)
17	MS Exchange 2000 SMTP (W3C)
18	Oracle Audit Log
19	Oracle DB Log
21	Oracle Listener Log

Sample Tools

Some tools are provided with LogLogic LMI as examples of how to use the ULDP library.

The sample tools do not intend to be production-ready pieces of software, but they demonstrate how to create the various message types and send them through a ULDP connection.

The source files are located in the `ULDPClient/examples` directory in the supplemental package of LogLogic LMI. The only dependency to run the Java classes is the ULDP-client library. For further information on how to use these commands and a list of options, see the source code.

The following sample tools can be used to send log files via a ULDP connection:

SendSyslogFileUldp

This tool sends the content of log files, optionally parsing the lines to extract the date and origin of the log, if they are formatted according to the syslog standard.

TailUldp

This tool sends the content of files line by line as they are appended at the end, similar to a UNIX `tail` command.

SendFileUldp

This tool sends entire files in a mechanism compatible with the file pulling features of LogLogic LMI.

Setting Up a Secure Connection With ULDP

To set up proper encryption with ULDP, you need to provide the API with a keystore at startup. The keystore contains the certificate and private key for authenticating with LogLogic LMI, and also an anchor of trust or certification authority (CA) to authenticate the server, unless you have chosen the `noServerAuthentication` option.

You can use the PKI of your choice to generate the keystore and its certificate. For testing purposes, the following procedure creates a minimal PKI. For actual implementation, refer to the documentation of the certificates that you use.

Procedure

1. Create a keystore that contains a new CA, for example: `CA.ks`.

```
keytool -genkeypair -alias CA -keyalg RSA -sigalg SHA256withRSA -
storetype JKS -keystore CA.ks
```

2. Export the root certificate into a file.

```
keytool -exportcert -alias CA -keystore CA.ks -rfc -file CA.cert
```

3. Create a keystore for the ULDP client API.

```
keytool -genkeypair -alias ULDP_client -keyalg RSA -sigalg
SHA256withRSA -storetype JKS -keystore client.ks -dname "CN=client IP
address"
```

> **Note:** The value of `dname` can be arbitrary and no validation is performed in LogLogic LMI. It is best practice to use an IP address that identifies the client.

4. Import the CA certificate into the client keystore.

```
keytool -importcert -alias CA -keystore client.ks -file CA.cert
```

5. Create a CSR from the client keystore.

```
keytool -certreq -rfc -alias ULDP_client -keystore client.ks -file
ULDP_client.csr
```

6. Sign the CSR with the CA.

```
keytool -gencert -infile ULDP_client.csr -alias CA -keystore CA.ks
-validity 365 -outfile ULDP_client.cert -rfc
```

   You can choose the validity arbitrarily. Ensure that the date on the host system is correct in order to produce a valid certificate.

7. Import the certificate in the client keystore.

```
keytool -importcert -file ULDP_client.cert -keystore client.ks -
alias ULDP_client
```

The client keystore for the ULDP library is now ready. Repeat steps 3-7 to create additional keystores for other clients, if needed.

8. Generate a CSR from the LogLogic LMI certificate.

```
system secureuldp create csr
```

Then copy and paste the contents of the certificate signing request into a file on the local machine, for example, `lmi.csr`.

9. Sign the LogLogic LMI CSR.

```
keytool -gencert -alias CA -infile lmi.csr -outfile lmi.cert -
validity 365 -keystore CA.ks -rfc
```

10. Import the CA root certificate into LogLogic LMI.

```
system secureuldp install rootCA
```

Then copy and paste the contents of `CA.cert` into the terminal.

11. Import the LogLogic LMI certificate.

```
system secureuldp install certificate
```

12. When calling the ULDP API, use the following methods on the `UldpConnectionSettings` object:

```
setUseAuthentication();
.setKeystorePath("location of client.ks");
.setKeystorePassword("keystore password");
```

13. On the LogLogic LMI GUI, navigate to **Home > Administration > System Settings**, and:

    a. Select **Enable Secure ULDP** to enable secure ULDP on the appliance.

    b. Click **Update**.

Syslog Host Field Character Sets

A list of the acceptable characters in an ASCII syslog header.

Character Description	Example	
Alpha chars, upper or lower case	A-Z and a-z	
Numbers	0-9	
Punctuation	at sign	@
	underscore	_
	period	.
	backslash	\
	colon	:
	asterisk	*
	brackets	[ ]
	parentheses	( )
	plus sign	+
	minus	-
space		
tab		

Exceptions

Some exceptions occur for ASCII syslog headers.

The following exceptions are noted for ASCII syslog headers:

- Some Unix/Linux syslog messages have a path in the process name. That is taken care of by looking for a leading backslash (/) and any number of the following characters:
    - Alpha characters, upper or lower case
        - A-Z
        - a-z
    - The numbers 0-9
    - Punctuation including:
        - underscore _
        - period .
        - dash -
- The use of space and tab depends on the log source. Some log sources have spaces at the point right before the log source target string is found. Others have only a tab. Specifically:
    - Windows messages require a space before the target string
    - Cisco VPN3000 requires a tab

Supported Regular Expression Characters

Advanced Search and data models in LogLogic LMI support the following regular expression meta characters, based on Java regular expressions:

Characters	Description
\a	Matches ASCII character code 0x07.
\d	Matches character in the set "0123456789".
\D	Matches any byte not in the set "0123456789".
\e	The escape character. Matches ASCII character code 0x1b.
\f	The form-feed character. Matches ASCII character code 0x0c.
\n	The new line (line feed) character. Matches ASCII character code 0x0a.
\r	The carriage return character. Matches ASCII character code 0x0d.
\s	A white space. Matches white space - \t \n 0x0b \f or \r.
\S	A non-white space. Matches any byte not in \s.
\t	The tab character. Matches any byte not in 0x09.
\w	A word character. Matches any ASCII character in the set underscore, digits, or upper or lower case letter.
\W	A non-word character. Matches any bytes not in \w.
\xHH	Matches a byte specified by the hex code HH. There must be exactly two characters after the \x.
\Q	Starts a quoted region. All meta characters lose their meaning until \E. A \\

Characters	Description
	can be used to put a backlash into the region.
\anything else	Matches the next character.
\k<name>	Refers to previous named capture.
[]	Specifies a character class - match anything inside the brackets. A leading ^ negates the sense of the class - match anything not inside the brackets. Negated character classes are computed from the set of code in the range 0-127 - in other words no bytes with the high bit set. Within a character class the following backslash characters mean the same thing as outside the character class: \a, \d, \D, \e, \f, \n, \r, \s, \S, \t, \w, \W, and \xHH.
{num} or {num:num}	Specifies a repetition count for the previous regular expression. Num must be less than 16. {num} is equivalent to {0:num}.
.	Matches any byte: 0x00 - 0xFF.
+	Specifies that the previous regular expression is repeated 1 or more times.
*	Specifies that the previous regular expression is repeated zero or more times.
( ) (?:)	Specifies capturing or non-capturing groups.
(?<name>)	Specifies capturing named groups.
|	Specifies alternation.
?	Specifies that the previous regular expression is repeated zero or one time.
anything else	Any other character matches itself.

Supported Data Types

This section provides a list of data types used in LogLogic LMI data models.

Standard data typesThe following standard data types are used in LogLogic LMI data models.

Data type	Description
Float	Stores double-precision floating-point numbers with up to 17 significant digits
String	Stores text, spaces, and numbers
Int	Stores integers
Long	Stores large integers
Boolean	Stores true or false
Timestamp	Stores both date and time
Inet_addr	Stores IP addresses in binary format

Data types created in LogLogic LMIThe following data types have been created for use in LogLogic LMI data models.

Data Type	Description	Examples
Duration	Displays time in a human-readable format on the GUI for time values that are stored in milliseconds (ms) as a Long data type.	• 1Y2M3W4D5h6m7s8ms • 1y4d5m 0 values are omitted while displaying time in the Duration format.

Search Syntax Reference

LogLogic Advanced Search query language is intuitive and efficient, you can search large data and view results in seconds.

The search query supports the following types of languages:

- Event Query Language (EQL)
- Structured Query Language (SQL) dialect.
- Event Correlation Language (ECL)

Both EQL and SQL are equally capable for searching, but the syntaxes are different in some cases. For example, simply providing a string in EQL is understood as a full text search, but it gives a syntax error in SQL. So the translation is not always literal. EQL is easy to use, however, SQL is more familiar and it is to write SQL using existing SQL tools.

Using EQL, you can define filters, regular expressions, sources, time ranges. ECL can be used to find patterns in a given set of data and for correlation purposes.

Event Query Language Reference

The search query supports the query languages: EQL and the LogLogic LMI SQL dialect.

The EQL query is composed of different parts separated by pipe (|) character. The pipe delimiter is used to separate the expression and each subsequent expression. Each pipe-delimited expression further processes search results from the preceding expression. For more structured queries, a subset of SQL is supported that is mainly focused on the SELECT statement. Both languages are supported in LogLogic LMI, but you cannot use a mix of both languages in the same query. All that is available in EQL can be achieved via SQL and vice versa except the following differences:

- EQL supports the full text search statement, but SQL does not support this statement. For details, see FILTER Statement.
- Multiple EQL filter expression statements, separated by a pipe, get automatically combined using the AND operator into a single filter expression. SQL does not

support this feature.

The EQL and SQL language rules are based on a Backus-Naur Form (BNF)-like notation as follows:

```
<symbol> ::= <expression> ;
```

where:

- Non-terminal symbols in syntax rules have angle brackets (< >). For example, in the rule <expression> ::= <expression> "+" <integer>; the <expression> is a non-terminal symbol and the rule specifies that as an expression is the addition of any number of integers.
- Terminal symbols are shown in double quotes (" "). For example, the "+" in the previous example.
- As an additional shortcut notation to BNF, optional symbols (that can occur zero or one times) are followed by a question mark (?). For example, in the rule <colNameForSort> ::= <colname> (ASC| DESC)?; a column name used for sort is a column name optionally followed by the keywords ASC or DESC.
- Optional symbols that can occur zero or any number of times are followed by an asterisk (*). For example, in the rule <itemList> ::= <item> ("," <item>)*; an itemList can contain one or more comma-separated items.
- Multiple symbols are grouped together using parenthesis ( ) when some common operation is applied, for example, the selection of one member of the group, or to indicate that the entire group can be repeated zero or more times. An example is shown in the previous bullet item.
- Words that are all capitalized represent keywords (special terminal symbols). For example, the keywords ASC and DESC in the column name for sort described in the previous example.

All parts of the query are optional, but overall the syntax is:

```
<EQL_statement> ::= <statement> ("|" <statement> )* ;
<statement> ::= <useStatement> | <filterStatement> | <groupByStatement>
|
<columnsStatement> | <sortStatement> | <limitStatement>;
```

String literals and identifiers (including keyspace, column family names, and data model names) are case sensitive but all EQL keywords are not case sensitive. For example, 'USE Windows' and 'use Windows' are treated in the same way.

String literal can be quoted with single (') or double (") quotation marks. The quotation marks (single or double) inside the string literal has to be prefixed with backslash (\ ) character. The \ character change to be prefixed with another backslash (\\ ). For example, "Mike's car" or 'Mike\'s car'.

A special syntax for time range can be used. For details, see Time Range Expressions.

Note: In this syntax reference topic, EQL keywords are mentioned in upper case letters for easy readability.

Examples

Expression	Definition
`sys_sourceType = 65536 and sys_eventTime in -5d | columns sys_eventTime, sys_collectIP, ll_eventStatus`	Events from source type '65536' in last 5 days, display result as a table with columns sys_eventTime, sys_collectIP, and ll_eventStatus
`USE Microsoft_Windows | ll_eventAction ='A user account was enabled.' | sys_eventTime IN -1h`	Using the data model Microsoft Windows, display results of all events where a user account was enabled during the past hour.

Common Search Commands

The search commands that LogLogic EQL uses.

command	Definition
USE	Defines the data models, which include the parsing configuration. For details, see USE Statement.

command	Definition
COLUMNS	Defines which columns should appear in search results. For details, see COLUMNS Statement.
GROUP BY	Groups search results based on specified columns. For details, see GROUP BY Statement.
SORT BY	Sorts search results based on the expression. For details, see Time Range Expressions.
LIMIT	Limits the size of search results to be displayed. For details, see LIMIT Statement.
FILTER	For detailed information about filters, see FILTER Statement.

USE Statement

The USE statement defines which data models to query.

A data model is a way to view a set of events, including columns parsed off the event body. The data model defines which events to parse, how to parse them, and what columns to extract in order to execute this query.

The USE statement is an optional parameter, but it is a good practice to improve performance by reducing the set of event sources and set of parsers used.

```
<useStatement> ::= "USE" <identifier> ( "," <identifier> )* ;
```

The USE statement consists of the USE keyword followed by one or many data model names separated by commas. An <identifier> is a letter followed by any sequence of letters, digits, or an underscore (_).

Note: If you do not specify any data model in the **Search** field, the results are retrieved in the following order:

1. All enabled built-in data model configurations
2. All enabled data models that are not LogLogic LMI-specific but have source filters defined
3. The system data model

The user-defined data models without the source filter are not included in the search query. For the list of built-in data models, see the Supported Log Sources list in the *TIBCO LogLogic® Log Source PackagesInstallation and Upgrade*, which is available on the TIBCO eDelivery website or TIBCO Support website after logging in. For more information about data models, see Data Models.

Certain data model expressions refer to a source of infrastructure data. This is defined by the corresponding data model itself and is typically defined by the name. The currently defined infrastructure data models are: `LogLogic_Config_Bloks` and `LogLogic_Config_Models` that represents the set of currently-defined Bloks and Data Model records respectively. For example, `use LogLogic_Config_Bloks | COLUMNS name, origin, created, type, description, value`

Do not use infrastructure queries within regular search queries. Infrastructure data model expression and an event data model expression are not allowed in the same query. An example of invalid mixed query is: `use LogLogic_Config_Bloks, system`.

Examples

Data Model Expression	Definition
use Windows	The result displays all events from Windows sources.
use Windows, Cisco	The result displays all events from Windows and Cisco log sources.

FILTER Statement

A filter is an expression that specifies the conditions that events must satisfy to be returned by this query.

The filter criteria can be in form of free text search of the entire body or value of a particular preparsed or parsed column.

Tip: The *system* (event metadata) columns are indexed so searching is faster on the *system* columns.

The list of available columns is determined by list of event sources. In case the list of event sources are not available, the system does the best to extract those columns using heuristics algorithms. For queries, the filter should contain a time condition, otherwise the default is used.

Note: When defining column names in a search query, follow the guidelines described in the COLUMNS Statement section.

A filter statement is any expression that evaluates to a result of type Boolean. Any event that does not satisfy this condition is eliminated from the results. An event that satisfies the condition if it returns true when the actual event values are substituted for any variable references.

The following table explains the types of filter statements that can be used.

Types of FILTER statements

Operator	Description
AND	Narrows your search results by only returning those events where each one of the AND conditions evaluates to true. For example, use AND to return results containing all specified keywords. When AND is used, the results contain all specified keywords and do not contain entries with just one of the specified keywords.
OR	Expands your search results by returning events where either of the OR conditions evaluates to true.

Operator	Description
	For example, use OR to return results containing any and all specified keywords. OR is ideal when you have common synonyms for a keyword. To narrow results as much as possible, combine OR statements with AND statements.
Full text search	Full text search on the body of each event can be performed by simply providing the phrase that needs to be enclosed in double quotes. For example, use system | "authentication failed" to retrieve all events that contain the phrase. The EQL full text search (specifically on sys_body) is exactly the same as the CONTAINS statement on the sys_body (so "use system | 'Bob'" is exactly the same as "select * from system where sys_body CONTAINS 'Bob'").
Equals (=), Not equals (<>), (!=), Natural equal to (==), Natural not equal to (!==) Lower than (<), Lower or equal (<=), Greater than (>), Greater or equal (>=)	A comparison condition compares two expressions using the operator specified in the comparison, which might be one of seven possible comparison operators with well-known meanings, and two more with special meanings. The comparison condition evaluates to true only if the comparison condition is satisfied. This can be used to narrow search results. These operators are case sensitive. For example, `col1 > col2/100`. Most of the operators behave as expected in SQL. If one or both values being compared is NULL, the result is NULL, otherwise the comparison is performed. However, the special operators == (natural equal to) and !== (natural not equal to) always return either true or false. If either value being compared is NULL, "natural equal to" returns false. If both values being compared are NULL, "natural equal to" returns true. If neither value being compared is NULL, the comparison is performed as expected. "Natural not equal to" returns the negation of "natural equal to". The operators == and !== can also be specified in functional form as NaturalEqualTo(val1, val2) and NaturalNotEqualTo(val1, val2). **Note:** If the field type is string, the comparison operators less than (<) and greater than (>) compare the data lexicographically (as strings) even if the data is numerical. For example, if the field type is string, ’21’ is considered less than ’3’.

Operator	Description
Plus (+), Minus (-), Multiply (asterisk (*)), Divide (forward slash (/), String concatenation (|||)	The arithmetic (+, -, *, /) and string concatenation (|||) operators can be used to create parts of other conditions. For example, `"column1 + column2 < 5"` or `"col3 * 4 - 1000 > col5"` The order of evaluation of the operators in an expression is according to the following precedence rules, from highest to lowest, with the highest precedence implying earlier evaluation: • Functions • Multiplication and division: both have equal priority and the evaluation order is from left to right • Addition and subtraction: both have equal priority and the evaluation order is also from left to right • String concatenation • Comparators (>, < and so on) For example, if you have an expression of the form "col1 > col2 + col3*col4", then col3*col4 is evaluated first, and then the result is added to col2. The col1 is then compared against the final result to see if it is greater. You can use a floating point number with the divide operator (/), to obtain a floating point number as the result. For example, `<Number> / 1024.0`.
Function	A set of predefined functions. For details, see Predefined EQL Functions. They can be used in filter, column expressions, or as part of Data Model expressions. **Note:** The parameters of the functions can be expressions themselves and is evaluated before the function is called. For example, `"ToInt(col1 + col2)"` adds the contents of the columns of the event named `col1` and `col2`, and pass the result to the `ToInt` function and the result of the function is used.
BETWEEN	Narrows your search results by only selecting those events where the left

Operator	Description
	hand side expression evaluates to a value that is between the two right hand side target expressions. Supports Timestamps, Long, and Integers. For time range syntax details, see Time Range Expressions.
IN, NOT IN	Narrows your search results. This is case sensitive. Checks whether the value matches any one of the values in a set or not. For example: • `"eventID IN ('id1', 'id2', 'id3')"` • `ll_eventID NOT IN ("6279","4749")` Supports all data types. For time range syntax details, see Time Range Expressions.
IS NULL, IS NOT NULL	Narrows your search results by accepting or rejecting the event based on whether the evaluated expression is null or not null. An expression most frequently becomes null if a column named in the expression has no value for the current event. Supports all data types.
LIKE, NOT LIKE	Expands your search results. Returns true if it matches the supplied pattern. This is case sensitive. The following rules are used to interpret the supplied string. • The character percent (%) is the wildcard character (matches zero or more characters). • The character underscore (_) means that it matches exactly one character. • The backslash character (\) is used to escape itself and the two characters mentioned earlier, if a literal search for any is desired.

Operator	Description
	Note: Since string literals in EQL and SQL require backslashes (\) to be escaped, note that additionally escaping for the LIKE statement doubles the escaping requirement. The simple rule to follow is to construct the match string using the rules stated earlier, then simply double up each backslash. The following examples show the actual syntax (not the escaping needed for Java): • `col1 LIKE "a_b"` - produces a match for "acb", "adb" and so on • `col1 LIKE "a\\_b"` - produces a match for "a_b" but not "acb". Note the double backslashes. • `col1 LIKE "a\\\\_b"` - produces a match for "a\cb" and "a\db" • `col1 LIKE "a%b"` - produces a match for "ab", "acb", "accb" and so on • `col1 LIKE "a\\%b"` - produces a match for "a %b" but not "acb"
CONTAINS, NOT CONTAINS	Expands your search results. Returns true when at least a part of the string matches the supplied pattern. This is not case sensitive. The sys_body column is special, because the supplied pattern is used to do a full text search on the event body. For all other columns, the following rules are used to interpret the supplied string. • The character asterisk (*) is the wildcard character (matches zero or more characters). • The character question mark (?) means that it matches exactly one character. • The backslash character (\) is used to escape itself and the two characters mentioned earlier, if a literal search for any is desired. The CONTAINS statement for columns starting with sys_ uses a full text search.

Operator	Description
	Note: Since string literals in EQL and SQL require backslashes (\) to be escaped, note that additionally escaping for the CONTAINS statement doubles the escaping requirement. The simple rule to follow is to construct the match string using the rules mentioned earlier, then simply double up each backslash. The following examples show the actual syntax (not the escaping needed for Java): • `col1 CONTAINS "a?b"` - produces a match for "ccc acb jjj", "adb" and so on • `col1 CONTAINS "a\\?b"` - produces a match for "a?b" but not "acb". Note the double backslashes. • `col1 CONTAINS "a\\\\?b"` - produces a match for "a\cb" and "a\db" • `col1 CONTAINS "a*b"` - produces a match for "ab", "acb", "accb" and so on • `col1 CONTAINS "a\\*b"` - produces a match for "a*b" but not "acb"
REGEXP, NOT REGEXP	Narrows your search results. By default, this is case sensitive but can be changed in the regular expression using the embedded flag (?i). Returns true if it matches the supplied pattern. The pattern syntax uses POSIX syntax. Since string literals in EQL/SQL require backslashes (\) to be escaped, note that all the backslashes inside a regular expression pattern must be doubled up, similar to the LIKE statement. Examples: • `col1 REGEXP "[a-z]b"` - produces a match for "ab", "cb" but not "Ab" or "_b" • `col1 REGEXP "\\w*"` - produces a match for a word, for example "this" or "that", but not "this and that"
DISTINCT	Fetches only the distinct values from data. The DISTINCT statement can be used in all advanced search queries.

Operator	Description

	If your search query includes a DISTINCT clause and the data includes multiple records with NULL values, all NULL values are included as distinct values in the search result. For example:
	`USE TIBCO_Hawk_Agent | COLUMNS DISTINCT ll_ collectHostName, ll_sensorName`
	In this example, `ll_collectHostName` and `ll_sensorName` can have NULL values.
	As a workaround, add the IS NOT NULL clause for all DISTINCT projections in the query. However, all results with NULL values are excluded from the search result.
	`USE TIBCO_Hawk_Agent | COLUMNS DISTINCT ll_ collectHostName, ll_sensorName | ll_collectHostName IS NOT NULL AND ll_sensorName IS NOT NULL`

Examples

Filter Expression	Definition
`"Authentication" and sys_eventTime in -1y`	The result displays all events that contain Authentication from the last one year.
`use sample | ll_sourceUser = 'SiteSvrAdmin' | sys_eventTime in '2014-02-02'`	The result displays all events that contain column 'll_sourceUser' and value is 'SiteSvrAdmin' on the 2 February 2014.

Predefined EQL Functions

A list of functions that are available in the EQL.

- Conversion functions
- String Functions
- Comparison Functions
- Math Functions
- Conditional Functions
- Time functions
- Miscellaneous Functions

The conversion functions are typically used when adding a new data model, or when you need to define new columns, where the expressions for new columns can use conversion functions to convert between data types and combine them using various operators. For instructions on how to add a new data model, see Creating a Data Model in Graphical Mode.

Conversion functions

Conversion functions

Function Name	Arguments	Returns
ToTimestamp	• (expression, formatString) • (expression, formatString, timezone) • (expression, formatString, timezone, defaultValue)	The expression, which should evaluate to a string, is interpreted as a time according to the supplied `formatString`. If the conversion fails, null is returned, unless a default string is provided, which is interpreted as a time and returned. Example: `ToTimestamp( logFileStringTimestampField, "dd, MM, yyyy HH:mm:ss", "America/ Los_ Angeles", "01, 01, 1970 00:00:00")` **Note:** If timezone is omitted or is empty, the system default timezone is used. If `formatString` does not contain a year, then when the function is being evaluated

Function Name	Arguments	Returns
		in the context of processing an event, the year from the event time (sys_eventTime) is used. If this results in a timestamp that is later than the event time, the prior year is used.
ToIP	• (expression) • (expression, defaultValue)	Converts the expression to an IP address (Java InetAddress). If the conversion fails, returns null; but if a default string is provided, it is interpreted as an IP address and returned. Example: `ToIP(ipAddressField, "10.0.0.1")`
ToTimestampString	• (expression, formatString) • (expression, formatString, timezone) • (expression, formatString, timezone, defaultValue)	Same as ToTimestamp, except that the expression is converted to string to get a printable timestamp. Example: `ToTimestampString(timestamp , "dd, MM, yyyy HH:mm:ss", "America/Los_Angeles", "01, 01, 1970 00:00:00")` **Note**: If timezone is omitted or is empty, the system default timezone is used.
ToInt	• (expression) • (expression, defaultValue)	Returns the integer value of the expression; or the default value if not convertible. Example: `ToInt("1348") or ToInt (numberField, 0)`
ToLong	• (expression) • (expression, defaultValue)	The obvious conversion to Long with default value taken if not convertible. Example: `ToLong("1348") or ToLong (numberField, 0)`
ToString	• (expression)	Returns the string string value of the

Function Name	Arguments	Returns
	• (expression, defaultValue)	expression; or the default value if not convertible. Example: `ToString(124.5) or ToString (numberField, "null")`
ToFloat	• (expression) • (expression, defaultValue)	Returns the Float value of the expression; or the default value if not convertible. Example: `ToFloat("1348.2") or ToLong (numberField, 0.0)` **Note**: LogLogic LMI uses double precision (that is 64 bits) when storing floating point numbers.
ToBool	• (expression) • (expression, defaultValue)	Returns the Boolean value of the expression; or the default value if not convertible. Example: `ToBool("FALSE") or ToBool (col1, FALSE)`
ExtractJson	• (expression, extraction path) • (expression, extraction path, default value)	The expression, which is a JSON string is parsed. A field is extracted from the expression using the extraction path. If either the expression or the path are invalid, an optional default value is returned. Example: `ExtractJson("{"cat": {"color":" blue"}}", "cat.color", "burlesque")` returns a string "blue" which is a JSON value of color, which is a JSON value of cat.
ExtractKvp	• (expression, extraction path) • (expression, extraction path,	The expression, which is a nested KVP string is parsed. A field is extracted from the expression using the extraction path. If either the expression or the path are invalid, an optional default value is returned.

Function Name	Arguments	Returns
	nested KVP delimiters /default "{}"/) • (expression, extraction path, nested KVP, deliiter / default ","/) • (expression, extraction path, nested KVP, delimiter, separator /default "="/) • (expression, extraction path, nested KVP, delimiter, separator, escape character / default "\\"/) • (expression, extraction path, nested KVP, delimiter, separator, escape character, default value)	Example: `ExtractKvp(" alert={ agent={ hostname=esbqa01, dns=none}}" , "alert.agent.dns")` returns a string "none". or `ExtractKvp("(abc^def|asd^aaa)" , "asd", "()", "|", "^")` returns "aaa".

String Functions

The smart list functions are usually used in filter expressions and data model.

Function Name	Arguments	Returns

lookup	(string 1, string 2, [string 3])	Returns the value associated with `string2` in the Enrichment list named `string1`.
		
		Examples:
		
		`lookup("list1", "key1")`
		`lookup ("list1", "key1", "default_return_value"`
		`$list1("key1")`
		`$list1("key1", "default_return_value")`

length	(expression)	Returns the length of the string value of the evaluated expression. For example, if an expression is an integer, it is converted to a string first. Example: `length("abc")` is 3, `length (3145)` is 4 (after converting the integer 3145 to the string "3145")
TransformString	• (stringToTransform, regularExpression, template) • (stringToTransform, regularExpression, template, defaultValue)	The function tries to match the `stringToTransform` string with the regular expression, and then returns the template with references to groups in the regular expression substituted with the actual values. To refer to groups, use $1, $2, so on, to refer to numbered groups, and $<name> to refer to named groups. If the string does not match, or is there any other error, the default value is returned (or NULL if not specified). Example: `TransformString("myName=loglogic", "myName=(\\S*)", "the name is $1")`

Function Name	Arguments	Returns
		returns: "the name is loglogic".
lower	(string)	Returns the lower case of the string
upper	(string)	Returns the upper case of the string
trim	(string)	Returns the trimmed string (without leading and trailing spaces).
substitute	(string1, string2, string3)	Substitutes `string2` by `string3` in `string1`.
left	(string, Int) <int>	Returns the left characters of string.
right	(string, Int) <int>	Returns the right characters of string.
• mid • substr • substring	(string, Int 1, Int 2)	Returns the characters from string starting at offset `int1` for a length of `int2`.
position	(string 1, string 2)	Returns the index of the first occurrence of `string2` within `string1`; -1 if no occurrence is found.
concatenate	(string 1, string 2, …)	Returns the concatenation of all strings passed as arguments.
find	(string1, string2, index)	Returns the position (from the beginning of `string1`) of the first occurrence of `string2` within `string1`; and -1 if no occurrence is found. Searching starts beyond the position specified by `index`. Example: `find("Here Here", "e", 3 )`

Function Name	Arguments	Returns
		returns 6
findlast	(string1, string2, index)	Returns the position (from the beginning of string1) of the last occurrence of `string2` within `string1`; and -1 if no occurrence is found. Searching starts beyond the position specified by `index`. Example: `findlast("Here Here", "e" )` returns 8
findnth	(string1, string2, n, index)	Returns the position (from the beginning of `string1`) of the 'n'th occurrence of `string2` within `string1`; and -1 if no occurrence is found. Searching starts beyond the position specified by `index`. Example: `findnth("Here Here", "r", 3, 2 )` returns -1
• substituteall • substitutefirst	(source, regex_pattern, replacement_string)	• Returns the new string after replacing the pattern in the source string with the replacement string. • Returns the original string if the pattern is invalid or not found. Examples: ◦ `substituteall('(4)john (3)hot(6)doggie(3)com', '(\d+)', '!' )`

Function Name	Arguments	Returns
		returns `'!john!hot!doggie!com'` ◦ `substitutefirst( '(4)john(3)hot(6)doggie(3)com', '(\d+)', '!' )` returns `'!john(3)hot(6)doggie(3)com'`
split	(input_string, 'separator', position_num)	Returns the string at the specified position, extracted from the original string specified in `input_string`; and an empty string if no occurrence is found. • `separator` can be any regex pattern • `position_num` starts at 0 • The function can be used within EQL, SQL, and ECL queries for advanced search queries, aggregation rule queries, correlation Blok definitions, and column parsing rules of a data model. However, the function **cannot** be used in a GROUP BY clause of aggregation rules.

For examples of these functions in SQL and EQL queries, see Search Examples.

Comparison Functions

Function Name	Arguments	Returns
NaturalEqualTo	(arg1, arg2)	• True if arg1 equals arg2

Function Name	Arguments	Returns
		• False if arg1 is not equal to arg2 • False if only one of the arguments is NULL • True if both arguments are NULL
NaturalNotEqualTo	(arg1, arg2)	• True if arg1 is not equal to arg2 • False if arg1 equals arg2 • True if only one of the arguments is NULL • False if both arguments are NULL

Math Functions

Function Name	Arguments	Returns
Power	base, exponent	Returns the value of base raised to the power of exponent. Any null argument returns null.
Round	• (numeric) • (numeric, [integer])	Returns a FLOAT value of the numeric argument rounded to the number of decimal places specified by *integer*. The argument *integer* is optional, and the default value is zero if not provided. The numeric value is rounded mid-way and away from zero. A null argument returns null. If the second argument is positive, it represents the number of decimal places to the right of the decimal point, whereas if it is negative, it represents the number of places to the left of the decimal point which become a zero.

<table>
<tr><th>Function Name</th><th>Arguments</th><th>Returns</th></tr>
<tr><td></td><td></td><td>For example:<table><tr><th>Function</th><th>Returns</th></tr><tr><td>ROUND(12.345)</td><td>12.0</td></tr><tr><td>ROUND(12.345, 2)</td><td>12.35</td></tr><tr><td>ROUND(123.45, -1)</td><td>120.0</td></tr><tr><td>ROUND(34567, -2)</td><td>34600</td></tr><tr><td>ROUND(-3456.5)</td><td>-3457.0</td></tr></table></td></tr>
<tr><td>Sqrt</td><td>numeric</td><td>Returns the square root of the argument or null if the argument is null.</td></tr>
<tr><td>sum</td><td>(<columnName>)
(distinct <columnName>)</td><td>Returns the sum of all values in the column specified by <columnName>. If <code>distinct</code> is specified, returns the sum of the distinct values in the column.

For example, if the column 'contentLength' contains the values 2,3,4,5,5,6:<table><tr><th>Function</th><th>Returns</th></tr><tr><td>sum (contentLength)</td><td>25</td></tr><tr><td>sum (distinct contentLength)</td><td>20</td></tr></table></td></tr>
</table>

Conditional Functions

Function Name	Arguments	Returns
IIF	Condition, then, else	Returns the value of the 'then' clause if the condition is true, otherwise the value of the 'else' clause. Example: `IIF(true, “a”, “b”)` returns “a” if true and “b” if false

Time functions

The following time functions are available:

- seconds (timestamp, [multiplier])
- minutes (timestamp, [multiplier])
- hours (timestamp, [multiplier])
- days (timestamp, [multiplier])
- weeks (timestamp, [multiplier])
- months (timestamp, [multiplier])
- years (timestamp, [multiplier])

Each function returns the value of the specified `timestamp` parameter truncated to the corresponding time unit (seconds, minutes, hours, and so on).

If the optional parameter `multiplier` is specified, then the function creates a time bucket of the specified units and with the precision of the multiplier, and returns the truncated timestamp at the start of the time interval specified by `multiplier`.

Example

Consider the following example:

```
seconds(sys_eventTime, 10)
```

This function creates a time bucket of 10 seconds, and returns the truncated timestamp out of the specified `sys_eventTime` at the start of the 10-second interval. That is, if the

value of sys_eventTime is "2020-06-26 10:57:24", then it truncates the value to "2020-06-26 10:57:20".

Miscellaneous Functions

Function Name	Arguments	Returns
`geoiplookup` By using this function within SQL and EQL queries, you can search logs that originated from a particular geographical area such as location, country, city, postal code, and so on. You can use the function in Advanced Search and Advanced Dashboards.	• (IPaddress) • (IPaddress, field_option) In the `IPaddress` parameter, you can specify an IP address or a column that stores an IP address such as Inet_Address. The `field_option` parameter can be one of the following values: • location • continent • country • city • postal • subdivision • domain • connectiontype • asn	Returns the geographical information of a specified IPv4 or IPv6 address. Returns the country name if the `field_option` parameter is not specified. **Note:** To use this function, the MaxMind database must be available on your appliance. If any errors are displayed, contact your administrator.
`matchcidr` You can use the function within SQL, EQL, and ECL	• (IP_string_format, IP_address_to_Match)	• Returns those IP addresses from the `IP_address_to_Match` parameter, which belong

Function Name	Arguments	Returns
queries, and in turn, in other functionality that make use of these queries.	• (IP_string_format, IP_address_to_Match, key) • (CIDR_expression, extractIP(input_column_name or custom_input_value)) In the `IP_string_format` parameter, you can specify a CIDR address, single IPv4 or IPv6 address, a comma-separated list of CIDR IP addresses or a range, a hyphen-separated range of IP addresses, or an enrichment list. In the `IP_address_to_Match` parameter, specify the column name from the logs, which are to be matched against the `IP_string_format` parameter. The data type of the column must be INET_ADDR. The `key` parameter specifies a key name in the enrichment list. The `key` parameter is mandatory when using this function in correlation Bloks. You must provide either a key	to the IP addresses specified in the `IP_string_format` list. • If you use `extractIPs` within `matchcidr`, then `extractIPs` extracts all IP addresses from either a column containing a list of IP addresses or a list of IP addresses. If any of the IP addresses matches with the `CIDR_expression`, then `matchcidr` returns true; otherwise false.

Function Name	Arguments	Returns
	name or an empty string ("").	
`extractIPs` You can use the function within SQL, EQL, and ECL queries, and in turn, in the `matchcidr` function.	• for extractIPs: (list_of_IP_addresses) • within matchcidr: (input_column) or (custom_input_value) The input list of IP addresses can contain IPv4, IPv6, or IPv6 compressed addresses.	Returns a list of IP addresses as a comma-separated string; or an empty string if no IP addresses are found within the input string.

For examples of `geoiplookup` and `matchcidr` functions in SQL and EQL queries, see Search Examples.

Time Range Expressions

The time range for IN operator understands both relative time and absolute time. Absolute time is the same as in BETWEEN operator.

Relative time is defined as: `<sign><number><unit>`, for example: `-5d` means 5 days ago.

Note: All dates and times are defined in the local time zone of the machine where the system is installed and it is not based on the browser time zone.

The following time units are available:

- s - second
- m - minute
- h - hour
- d - day

- w - week
- M - month
- q - quarter (3 months)
- y - year

The supported timestamp formats are:

- Any day of the week; for example, MON, TUE, WED, THU, FRI, SAT, SUN
- NOW specifies up to the current time
- Today specifies as the end of the day (23:59:59)
- yyyy-MM-dd HH:mm:ss, {d yyyy-MM-dd HH:mm:ss}, {t yyyy-MM-dd HH:mm:ss}, or {ts yyyy-
- MM-dd HH:mm:ss}
- MM/dd/yyyy HH:mm:ss
- BETWEEN and IN support dates (yyyy-MM-dd or MM/dd/yyyy). The interpretation depends on whether it is used as beginning or end of time period. When used in beginning it is equivalent to yyyy-MM-dd 00:00:00; when used at the end - yyyy-MM-dd 23:59:59.

Examples

Time Range Expression	Definition
-23h	Last 23 hours. For example, if the current time is 2014-10-20 08:00:00, -23h is the period from 2014-10-19 09:00:00, which is exactly the last 23 hours.
-1d	Last 1 day including today. **Note:** Starts from the midnight of the **previous** date till the current time. Therefore, the period is always more than 24 hours. For example, if the current time is 2014-10-20 08:00:00, -1d is the

Time Range Expression	Definition
	period from 2014-10-19 00:00:00, which is 32 hours.
-1M	Last month.
"2014-10-20"	From 2014-10-20 00:00:00 and 2014-10-20 23:59:59.
"2014-10-20":"2014-10-25"	From 2014-10-20 00:00:00 until 2014-10-25 23:59:59.
"2014-10-20 14:00:00":"2014-10-25 20:00:10"	From 2014-10-20 14:00:00 until 2014-10-25 20:00:10.
"2014-10-20 14:00:00":NOW	From 2014-10-20 14:00:00 until now (the time the query was issued).
MON:NOW	From beginning of last Monday till the current time.

COLUMNS Statement

COLUMNS is used to define which columns should appear in the results and how they should be computed.

```
<columnsStatement> ::= "COLUMNS" <columnsList> | <aggregationList> ;
<columnList> ::= <columnExpression> ( "," <columnExpression> )* ;
<aggregationList> ::= <aggregationExpression> ( ","
<aggregationExpression> )* ;
```

A COLUMNS statement can be a column based expression or an aggregate expression. A column based expression is any expression supporting mathematical and logical operators, functions, and other operators. For details, see FILTER Statement. An aggregate expression is a similar expression that contains an aggregationFunction. If all columns use aggregation functions, the result contains only one row with results of the aggregation. For details, see GROUP BY Statement.

When defining a column name in a search query, it must be enclosed in square brackets ([]) in the following scenarios:

- If a column name is also an EQL or SQL keyword, for example, `"use MyEventSourceConfiguration |[IN] >5"`.
- If a column name has a space, for example, `"use Hawk_getProcess | COLUMNS Status, [Virtual KBytes] | sys_eventTime in -10y"`.
- If a column name contains non-alphabetic or non-digit characters such as dash (-), for example, `"[ab]"`, to distinguish it from the subtraction expression "a-b".

The following data types for columns are supported:

- String
- Integer
- Long
- Double
- Boolean
- Timestamp
- IP address

Examples

Columns Expression	Definition
`columns sys_eventTime, sys_collectIP, sys_body`	The result is a table with three columns: `sys_eventTime, sys_collectIP, sys_body`. The columns could be one of the pre-parsed columns such as `sys_eventTime`, `sys_body`, or columns from configured parsers. See USE Statement.
`columns count(ll_sourceUser)`	The result has one column with one row with count of all events that has ll_sourceUser column with no empty value.
`columns ToInt(ll_eventActionID)+2 as action, sys_body`	The result is a table with two columns, the first column called 'action' with the value of converting ll_eventActionID to int and adding 2 to it, and the second column is sys_body.

Columns Expression	Definition
`columns max(length (sys_body)) – min(length(sys_ body))`	The result is a table with a column containing the difference in length between the longest and shortest events.

GROUP BY Statement

Grouping can be used to group values by one or more expressions involving columns.

Grouping requires a list of grouping expressions and the list of aggregation columns.

A group by expression can be a column name or an expression involving multiple columns, and an optional list of aggregation functions after the COLUMNS keyword. All the group by expressions and the aggregates listed after the COLUMNS keyword are displayed by the query.

```
<groupByStatment> ::= "GROUP BY" <columnExpression> ( ","
<columnExpression> )* )?
(COLUMNS <aggregationFunction> ("," <aggregationFunction> )* )?;
```

The following aggregation functions are supported:

- COUNT(*): Count all the rows.
- COUNT(columnName): Count all the rows in which the value of the column is not null.
- COUNT(DISTINCT columnName): Count all distinct values from the column.
- SUM(column): Sum of all values from the column. Supports numerical types (Integer, Long, Double).
- AVG(column): Provide average value for the column. Supports numerical types (Integer, Long, Double).
- MIN(column): Smallest value of the column. Supports all data types that can be ordered (Integer, Long, Double, Timestamp, String).
- MAX(column): Largest value of the column. Supports all data types that can be ordered (Integer, Long, Double, Timestamp, String).

- DURATION(timestamp): Returns the difference (in milliseconds) between the latest and the earliest time. Supports Timestamp only.
- Time functions: Groups events by time. Supports time functions (Seconds, Minutes, Hours, Days, Weeks, Months, Years).

Examples

Grouping Expression	Definition
`group by ll_sourceUser columns count(*)`	The result has two columns, the ll_sourceUser and count of users per distinct value.
`group by ll_sourceUser columns count(ll_sourceUser), min(sys_eventTime), max(sys_eventTime)`	The result has 4 columns ll_sourceUser, number of users for each distinct value of source user, minimum value of sys_eventTime and maximum value of sys_eventTime.
`group by ll_sourceUser columns Duration(sys_eventTime)`	The result has 2 columns, the source user and the duration.
`group by ToLong(sys_eventTime)/1000 COLUMN ToLong(sys_eventTime)/1000, AVG(LENGTH(sys_body)), COUNT(*)`	The result has three columns, ToLong(sys_eventTime)/1000, AVG (LENGTH(sys_body)), and COUNT(*). Grouping is done using the value of the expression in the first column, which results in events being grouped by the second at which they occurred. The next column shows the average length of the events every second. The last column shows the number of events that occurred every second.

SORT BY Statement

SORT BY causes the result rows to be sorted according to the specified expressions.

By default, results are sorted in ascending order.
`<sortByStatement> ::= "SORT BY" <expression> ( "," <expression> )* ;`

A SORT BY expression can be the name of a column.

If two rows are equal according to the leftmost expression, they are compared according to the next expression and so on. If they are equal according to all specified expressions, they are returned in an implementation-dependent order.

Note: Sorting by an expression is not supported directly, but you can do it if you put the expression in the projection (COLUMNS statement) and assign it a column name with the AS statement

The following functions are supported:

- ASC: Sort results in ascending order. This is the default order.
- DESC: Sort results in descending order.

Examples

Sorting Expression	Definition
`sort by sys_ eventTime ASC`	The result is sorted by time in ascending order.
`sort by ll_ sourceUser, sys_ eventTime DESC`	The result is sorted by ll_sourceUser in ascending order (default), in case ll_sourceUser is the same, sort by sys_eventTime in descending order.

LIMIT Statement

LIMIT indicates the maximum number of results that should be returned by the query.

```
<limitStatement> ::= "LIMIT" <number> ;
```

If you do not specify a LIMIT clause in the query, the default limit of 100,000 is used.

Example

```
USE system | LIMIT 100
```

Expression	Definition
`limit 100`	Limits the result set to top 100 rows.

Query Optimization for Better Performance

Besides narrowing the time span for the query, the best way to improve performance is to leverage the index, using the CONTAINS operator.

For example:

`sys_body CONTAINS 'string'`
quickly finds all the events that contain the token 'string' by using the index.

Note: The index only stores full words called tokens, and ignores characters such as punctuation signs, spaces, and so on.

Even if your query is based on other columns or operators, you can accelerate it if you know some tokens that appear verbatim in the events you are looking for, and add them to your query with the CONTAINS operator.

For example, the following query works as is:
```
USE Microsoft_Windows | ll_actionID = 4291
```

However, since we know that the token 4291 appears in the events we are looking for, we can get faster results by typing:
```
USE Microsoft_Windows | ll_actionID = 4291 AND sys_body CONTAINS '4291'
```

Text Search

Searching by text is an important feature when dealing with logs.

The LogLogic Query Language offers several operators to perform text matches:

	LIKE	REGEXP	CONTAINS
Matching level	Character	Character	Token
Syntax	Wildcards: _ %	Java Regular Expressions	Wildcards: ? *
Expression matches	Full string	Part of the string	Part of the string
Indexed	No	No	Yes
Case-sensitive	Yes	Per syntax	No

LIKE is the classical SQL operator. It matches the full string (so leading and trailing wildcards should be added if trying to match only a fragment). It has a granularity of character, that is, we can match character by character. The supported wildcards are _ for one character and % for many characters. It is not indexed, so it is not particularly fast.

REGEXP allows matching regular expressions. It searches a match within the string. It has character granularity. The syntax for the regular expression language is the same that provided by the Java language. It is not indexed.

CONTAINS searches within the index, with a token granularity. The index stores tokens, or full words, so we must search for the full words, or add wildcards. The wildcards allowed are ? for one character and * for multiple characters, and matching part of the string is enough. This operator takes advantage of the index, and hence CONTAINS speeds up queries.

See Types of FILTER statements for details on the syntax of this operators and available string functions that can be useful to manipulate text.

Search Examples

Some search examples for reference.

SQL Expression	EQL Expression	Definition
`select sys_eventTime,sys_body from LogLogic_Appliance where sys_eventTime between '2016-02-02' and '2016-02-03'`	`use LogLogic_Appliance |columns sys_eventTime, sys_body | sys_eventTime between '2016-02-02' and '2016-02-03'`	Displays results from the LogLogic_Appliance data model where the records have the timestamp between '2014-02-02' and '2014-02-03'.
`select * from LogLogic_Appliance where sys_body like '%Authentication%' and sys_eventTime between '2016-02-02' and '2016-02-03'`	`use LogLogic_Appliance | sys_body contains "Authentication" | sys_eventTime between '2016-02-02' and '2016-02-03'`	Displays results from the LogLogic_Appliance data model with "Authentication" in the event body.
`select * from LogLogic_Appliance where sys_body like '%logon%' and sys_eventTime between '2016-02-02' and '2016-02-03' limit 10`	`use LogLogic_Appliance | sys_body like '%logon%' | limit 10 | sys_eventTime between '2016-02-02' and '2016-02-03'`	Demonstrates an example of a 'like' statement to display a limit of 10 results.
`select * from LogLogic_Appliance where sys_eventKey REGEXP '[az0-9|]*' and sys_eventTime in -10y limit 10`	`use LogLogic_Appliance | sys_eventKey REGEXP '[a-z0-9|]*' | sys_eventTime in -10y | limit 10`	Demonstrates an example of the REGEXP expression matching.
`select * from LogLogic_Appliance where sys_eventTime between '2016-02-02' and '2016-02-03' order by sys_eventTime DESC`	`use LogLogic_Appliance | sys_eventTime between '2016-02-02' and '2016-02-03' | sort by sys_eventTime DESC`	Displays events sorted by time for records having timestamp for the specified dates in descending order.
`select * from LogLogic_Appliance where sys_eventTime between '2016-02-02' and '2016-02-03' order by sys_eventTime DESC limit 100`	`use LogLogic_Appliance | sys_eventTime between '2016-02-02' and '2016-02-03' | sort by sys_eventTime DESC | limit 100`	Displays top 100 results for records sorted by time in descending order.

SQL Expression	EQL Expression	Definition
`select sys_eventTime, sys_body from LogLogic_Appliance where sys_eventTime between '2016-02-02 14:34:34' and '2016-02-03 12:00:00' ORDER BY sys_eventTime DESC LIMIT 100`	`use LogLogic_Appliance | sys_eventTime between '2016-02-02 14:34:34' and '2016-02-03 12:00:00' | sort by sys_eventTime DESC | LIMIT 100`	Display sorted first page of results for events ordered by time in descending order.
`select ll_sourceUser, count(*) from LogLogic_Appliance where sys_eventTime between '2016-02-02' and '2016-02-03' group by ll_sourceUser`	`use LogLogic_Appliance | group by ll_sourceUser columns ll_sourceUser, count(*) | sys_eventTime between '2016-02-02' and '2016-02-03'`	Displays grouped results based on the source users.
`select ll_sourceUser, max(sys_eventTime), min (sys_eventTime), count (*) from LogLogic_Appliance where sys_eventTime between '2016-02-02' and '2016-02-03' group by ll_sourceUser`	`use LogLogic_Appliance | group by ll_sourceUser columns max(sys_eventTime), min(sys_eventTime), count(*) | sys_eventTime between '2016-02-02' and '2016-02-03'`	Displays the count of rows for distinct source users and its corresponding maximum timestamp and minimum timestamp.
`select ll_sourceUser, (max(ToLong(sys_eventTime ))- min(ToLong (sys_eventTime) ))/1000 as seconds from LogLogic_Appliance where sys_eventTime IN -10y group by ll_sourceUser`	`use LogLogic_Appliance | sys_eventTime in -10y | group by ll_sourceUser COLUMNS ll_sourceUser, (max(ToLong(sys_eventTime) )- min(ToLong (sys_eventTime)) )/1000 as seconds`	Demonstrates the use of a complex expression in the COLUMNS statement. For each user, calculate the difference in time between the earliest and latest events. The time values are first converted to LONG (milliseconds), then subtracted, and finally divided by 1000 to convert the milliseconds to seconds.

Examples for geoiplookup()

SQL Expression	EQL Expression	Definition
`select geoiplookup(sys_collectIP,"location"), sys_collectIP from Microsoft_Windows where sys_collectIP like '128%'`	`use Microsoft_Windows_Windows | sys_collectIP like '128%' |COLUMNS sys_collectIP, geoiplookup (sys_collectIP,"location")`	Returns location information such as latitude, longitude, time zone corresponding to the IPv4 address specified in the `sys_collectIP` parameter.
`select geoiplookup(sys_collectIP,"country"), sys_collectIP from Microsoft_Windows where sys_collectIP like '128%'`	`use Microsoft_Windows | sys_collectIP like '128%' |COLUMNS sys_collectIP, geoiplookup(sys_collectIP,"country")`	Returns the country corresponding to the IPv4 address specified in the `sys_collectIP` parameter.
`select geoiplookup(sys_collectIP,"country"), sys_collectIP from Microsoft_Windows where sys_collectIP like '2001%'`	`use Microsoft_Windows | sys_collectIP like '2001%' | COLUMNS sys_collectIP, geoiplookup (sys_collectIP,"country")`	Returns the country corresponding to the IPv6 address specified in the `sys_collectIP` parameter.

Examples for matchcidr() and extractIPs()

SQL Expression	EQL Expression	Definition
`select * from system where matchcidr ("198.51.100.101/32", sys_collectIP)`	If sys_collectIP="198.51.100.101", 203.0.113.101", then: `use system | matchcidr("198.51.100.101/32", sys_collectIP)`	Returns the data that includes `sys_collectIP` as 198.51.100.101, because the IP address is within the CIDR range of "198.51.100.101/32".

SQL Expression	EQL Expression	Definition
`select * from system where matchcidr ("198.168.0.0-198.168.255.255", sys_collectIP)`	If sys_collectIP="198.51.100.101",203.0.113.101", then: `use system | matchcidr("203.0.0.0-203.0.255.255",sys_collectIP)`	Returns the data that includes `sys_collectIP` as 203.0.113.101, because the IP address is within the range of "203.0.0.0-203.0.255.255".
• `select * from system where matchcidr ("IPList",sys_collectIP)` • `select * from system where matchcidr ("IPList",sys_collectIP, "IP1")`	For an enrichment list named `IPList`: `{"198.51.100.101/32":"IP1","192.0.2.101/32":"IP2","203.0.113.0-203.0.113.255":"IP3","2001:DB8:4860::8888/32":"IP4"}` and if `sys_collectIP`="198.51.100.101",203.0.113.101", then: • `use system | matchcidr("IPList",sys_collectIP)` • `use system | matchcidr("IPList",sys_collectIP,"IP1")`	Returns the data that includes `sys_collectIP` as 198.51.100.101, because this IP address is within the range 198.51.100.101/32 and matches the key IP1.
`select minutes (sys_eventTime), count(*) from system where sys_collectIP="198.51.100.0" AND extractIPs(sys_body) !="" AND matchCidr( "192.0.2.0/24" , extractIPs(sys_body) ) GROUP BY minutes(sys_eventTime)`	`USE system | sys_collectIP="198.51.100.0" AND extractIPs(sys_body) !="" AND matchCidr( "192.0.2.0/24" , extractIPs(sys_body) ) | GROUP BY minutes(sys_eventTime)`	Returns the output of the query aggregated per minute

Examples for split() and extractIPs()

SQL Expression	EQL Expression	Definition
`select split(sys_body,' ', 2) from system`	`USE system | COLUMNS split(sys_body, ' ', 2)`	**Parsing rule in a data model:** Returns the third element in the list, after splitting sys_body column using space as the separator.
`select * from system where sys_collectIP="192.0.2.0" and extractIPs(sys_body) !="" and split(extractIPs (sys_body),',',1) group by minutes(sys_eventTime)`	`USE system | sys_ collectIP="192.0.2.0" AND extractIPs(sys_body) !="" AND split(extractIPs(sys_ body),',',1)!="" | GROUP BY minutes(sys_eventTime)`	**Aggregation rule query:** Returns the output of the query aggregated per minute for results where `extractIPs` returns the list with more than one values. Otherwise `split` returns false.
`select split(sys_ collectIP, '\\.',3) from system`	`USE system | COLUMNS split(sys_collectIP, '\\.', 3)`	Returns the last numbers of the IP address provided in `sys_collectIP` and separated by dot (the regex '`\\.`' is for the dot character)
• `select split(sys_body, ' ', 4) from system` • `select extractIPs(sys_ body) from system where extractIPs(sys_ body)!="" and split (extractIPs(sys_ body),',',1)!=""`	• `USE system | COLUMNS split(sys_ body, ' ', 4)` • `USE system | extractIPs(sys_body) != "" AND split (extractIPs(sys_ body),',',1)!="" | COLUMNS extractIPs (sys_body)`	**Advanced search queries:** • Using space as a separator, returns the fifth element from the value of `sys_body` column. • Returns a list of extracted IPs from `sys_body` with more than one IP

SQL Expression	EQL Expression	Definition
		addresses.
`select * from system where sys_collectIP="192.0.2.0" and extarctIPs(sys_body)!="" and split (extractIPs(sys_body),',',1)!="" group by minutes(sys_eventTime)`	`USE system | sys_collectIP="192.0.2.0" AND extractIPs(sys_body) !="" AND split(extractIPs(sys_body),',',1)!="" | GROUP BY minutes(sys_eventTime)`	**Aggregation rule query:** Returns the output of the query aggregated per minute for results where `extractIPs()` returns a list with more than one values. Otherwise split() returns `false`.

OPTIONS Statement

You can use the training model classification in advanced search queries, correlation, aggregation rules, and Bloks, by adding the following syntax in the query:

```
OPTIONS useClassifier=<modelName>
```

The search result displays the following additional columns:

- `ll_tax_action`: Action performed by the logged-in user.
- `ll_tax_event_type`: Type of event, such as config activity or standard activity.
- `ll_tax_intent`: Intention of the action, such as authentication or common use.
- `ll_tax_outcome`: Outcome of the user's action, such as successful or failed.
- `ll_tax_target`: The parameter affected by the user's action, such as Application context or Auth policy.

The triplet `ll_tax_action`, `ll_tax_event_type`, and `ll_tax_intent` together indicate the user's action.

Event Correlation Language Reference

LogLogic Event Correlation Language (ECL) is effective in finding patterns in a given set of logs.

ECL is able to describe searches that are a bit complex for the regular EQL, especially when there is a need to join several types of events. Rules described in ECL can be used for advanced forensics searches and also for real-time alerting.

Rule Structure

A rule describes a pattern to look for within a given time window.

It contains a list of event group definitions (at least one), and the correlation criteria that are used to join those event groups (if there is more than one event group). A rule can also be valid for only a given period of time.

All mandatory parameters are explained here. The optional parameters are in square brackets [ ].

```
Valid From yyyy-MM-dd hh:mm:ss To yyyy-MM-dd hh:mm:ss ) ]
[ <identifier environment> ]
USE <source identifier> (, <source identifier>)*
Within <integer> [ d |h | m | s ][Fixed | Sliding ]
<event group 1>
<event group 2>
…
[Correlation
<correlation criteria 1>
<correlation criteria 2>
… ]
[Autofill]
(Set <expression> AS <identifier>)*
[Inject Correlation Event]
[ LIMIT <integer> CORRELATION EVENTS ]
```

Note: Each ruleset can have multiple rules. Each rule name must be unique in a defined ruleset.

Parameter	Description
Rule <identifier environment>	The rule name defined using an identifier and the environment. For details, see Identifier Environment.
USE	The list of log sources used by the rule. Multiple log sources must be separated by comma (,).
Within	The time period is defined as an integer in days, hours, minutes, or seconds.
Event Group	Each event group describes the criteria that must combine events to be group together as part of the rule. This is equivalent to a single search in EQL. For details, see Event Group Structure.
Correlation <correlation criteria>	The correlation criteria describes the joins and other constraints that various event groups must meet to trigger a rule. For details, see Correlation Criteria.
LIMIT	Limit on number of correlation events is only effective for "replay" instances when INJECT CORRELATION EVENT is not set. The default limit is 10,000.

Identifier Environment

An identifier environment is used to specify the data models, when those parts are not present in a key identifier.

The identifier environment is composed of:

- [Default data model <identifier>]

The identifier environment follows a hierarchical structure when resolving a missing part in an identifier. The order is as follows:

- Event Group Environment
- Correlation Rule Environment
- Ruleset Environment
- Root Environment (defined outside ECL itself)

- For Correlation REST API: this is the environment parameter.
- For Web application: this is related to the currently logged in user.

Limitations

- The only possible value for tenant is: tenant1
- The only possible value for domain is: shared
- The only possible value for source is: correlation

Simple Identifier

The simple identifier must be defined using letters, numbers, underscore (_), and dollar sign ($) with or without single quotes (' '). If single quotes (' ') are not used, use square brackets ([ ]), or back quotes (` `).

For example: `('a'..'z'|'A'..'Z') ('a'..'z'|'A'..'Z'|'0'..'9'|'_'|'$')*`

Key Identifier

An identifier is composed of four parts separated by dots. An identifier part follows the syntax of the simple identifier. The identifier parts are:

- Tenant name
- Domain name
- Source configuration name
- Field name (or column name)

The key identifier can be defined as follows:

```
[[[<tenant identifier>.]<domain identifier>.]<source config
identifier>.]<field
identifier>
```

The `<field identifier>` is mandatory part. If the other identifier parts are not defined, they are automatically extracted from the identifier environment.

Limitations:

- The only possible value for tenant is: tenant1
- The only possible value for domain is: shared
- The only possible value for source is: correlation

Event Group

An event group describes the criteria events should meet to be part of a rule. Event groups can be of the following types:

Type	Description
Required	The rule cannot be triggered if no event matches this event group. This is the default type.
Excluded	The rule is NOT triggered if an event matches this event group.
Optional	If any events match this event group and if other criteria are met, they are part of the triggering rule.

An event group can have the following parts:

- Conditions on the number of events
- A filtering clause
- A grouping clause
- A set of having clause
- Upper limits on the number of groups and events that can be created while this rule is run. This is a safeguard against a memory overflow.

An event group is defined as follows:

```
Event Group <identifier> [ Is ( Required | Optional | Excluded ) ] [
With Delayed Evaluation ]
[ At Least <integer> Events ]
[ At Most <integer> Events ]
[ <identifier environment> ]
```

```
[ Where <expression> ]
[ With The Same <expression> [ As <identifier> ]
( , <expression> [ As <identifier> ] )* ]
( Having <having clause> )*
[ Limits <integer> Groups And <integer> Events ]
```

When the `At Least` parameter is defined, it requires at least an integer more than 0. If it is omitted, this implies at least 1.

If the `Where` clause is defined, it must match the expressions. It is evaluated as a Boolean. For details, see Expressions.

Default Limits are 10000 groups and 100000 events.

Expressions

Expressions can be used to express how to compute a value in situations such as:

- in a condition
- in a grouping definition
- in field assignment

```
[ ( + | - ) ] <double>
[ ( + | - ) ] <long>
"<String>"
{ ( d | t | ts ) yyyy-MM-dd hh:mm:ss }
True
False
Null
<IPv4 address>
<IPv6 address>
<key identifier>
$<identifier>(<expression>)
( <expression> )
<expression> * <expression>
<expression> / <expression>
<expression> % <expression>
<expression> + <expression>
<expression> - <expression>
<expression> Is [ Not ] Null
Exists <expression>
<expression> [ Not ] Like <expression>
```

```
<expression> [ Not ] [ Any | All ] Contains <expression>
<expression> [ Not ] [ Any | All ] Regexp <expression>
<expression> [ Any | All ] = <expression>
<expression> [ Any | All ] != <expression>
<expression> [ Any | All ] > <expression>
<expression> [ Any | All ] >= <expression>
<expression> [ Any | All ] <= <expression>
<expression> [ Any | All ] < <expression>
<expression> [ Any | All ] <> <expression>
<expression> [ Any | All ] In ( <expression>, expression, … )
<expression> In <expression>/<expression>
<expression> [ Any | All ] Between <expression> And <expression>
Case <expression>
( When <expression> Then <expression ) +
[ Else <expression> ]
<function name> ( [ <expression> ] , [ <expression> ] , … )
<aggregation function>
```

The following operators are supported:

- Equals (=)
- Not equals (!=), (<>)
- Lower than (<)
- Lower or equal (<=)
- Greater than (>)
- Greater or equal (>=)
- In:
    - <list of expressions>: Checks if value matches any one of the values in a set. Supports all data types.
    - <network>/<net length>: Checks whether an IP address matches a network, defined as a network IP address and a network bitmask length.
- Between <expression> And <expression>: Supports Timestamps, Long, Integers and Float
- AND, OR

  Examples:

```
( sys_eventType = “1234”) and ( sys_body like “%login failed%”)
( sys_bodySize > 30) and (sys_bodySize < 20)
( ll_eventID is not null) and ( ll_eventID > -1 )
```

- Predefined ECL Functions
- Aggregation Functions
- Identifier Environment

Predefined ECL Functions

Functions are used to compute a value as output from parameters as input. Some functions are predefined in the language. You can also call a static Java function provided by the user. The predefined functions that are available in ECL are listed in the following tables.

- String functions
- List functions
- Conditional functions
- Smart List functions
- Conversion functions
- Miscellaneous functions

String Functions

Function Name	Arguments	Returns
• len • char_length • character_ length • length	(String)	Length of string 1.

Function Name	Arguments	Returns
lower	(String)	Lower case of string 1.
upper	(String)	Upper case of string 1.
trim	(String)	Trimmed string 1 (without leading and trailing spaces).
substitute	(String 1, String 2, String 3)	Substitute string 2 by string 3 in string 1.
left	(String, Int)	<int> left characters of string 1.
right	(String, Int)	<int> right characters of string 1 .
• mid • substr • substring	(String, Int 1, Int 2)	Characters from string1 starting at offset <int1> for a length of <int2>.
• find • position	(String 1, String 2)	Index of the first occurrence of string2 within string1, -1 if no occurrence is found.
concatenate	(String 1, String 2, …)	Concatenation of all strings passed as arguments.
TransformString	(stringToTransform, regularExpression, template) or (stringToTransform, regularExpression, template, defaultValue)	It tries to match the stringToTransform with the regular expression, and then returns the template with references to groups in the regular expression substituted with the actual values. To refer to groups, use $1, $2, etc to refer to numbered groups, and $<name> to refer to named groups. If the string does not match, or is there any other error, the default value is returned (or NULL if not specified).

List Functions

Function Name	Arguments	Returns
size	List	Size of the list

Conditional Functions

Function Name	Arguments	Returns
IIF	Condition, then, else	Returns the 'Then' value if condition is true, otherwise it should return the 'Else' value. For example: `IIF(true, “a”, “b”) returns “a” IIF (false, ”a”,”b”) returns “b”`

Smart List Functions

Smart List functions

Function Name	Arguments	Returns
lookup	(String 1, String 2)	The value associated with String2 in the smart list named String1.
isInList	(String 1, String 2)	True if the value String2 is defined in smart list named String1.

Conversion Functions

Function Name	Arguments	Returns
ToTimestamp	(expression, formatString) or (expression, formatString, timezone) or (expression, formatString, timezone, defaultValue)	The expression, which should evaluate to a string, is interpreted as a time according to the supplied `formatString`. If the conversion fails, null is returned, unless a default string is provided, which is interpreted as a time and returned. **Note:** If `timezone` is omitted or is empty, the system default `timezone` is used.
ToIP	(expression_ or (expression, defaultValue)	Convert the expression to an IP address (Java InetAddress). If the conversion fails, null is returned, unless a default string is provided, which is interpreted as an IP address and returned.
ToTimestampString	(expression, formatString) or (expression, formatString, timezone) or (expression, formatString, timezone, defaultValue)	Same as ToTimestamp, except the conversion is in the opposite direction

Function Name	Arguments	Returns
		to get a printable timestamp. **Note:** If timezone is omitted or is empty, the system default time zone is used.
ToInt	(expression) or (expression, defaultValue)	The obvious conversion to integer with default value taken if not convertible.
ToLong	(expression) or (expression, defaultValue)	The obvious conversion to Long with default value taken if not convertible.
ToString	(expression) or (expression, defaultValue)	The obvious conversion to String with default value taken if not convertible. **Note:** Using the ToString() function in a correlation Blok might result in inconsistent time format. Therefore, use the ToTimestamp() function instead.

Function Name	Arguments	Returns
ToFloat	(expression) or (expression, defaultValue)	The obvious conversion to Float with default value taken if not convertible. **Note:** LogLogic uses double precision (that is 64 bits) when storing floating point numbers.
ToBool	(expression) or (expression, defaultValue)	The obvious conversion to Boolean with default value taken if not convertible.
ToDouble	(expression) or (expression, defaultValue)	The obvious conversion to Double with default value taken if not convertible.
ExtractJson	(expression, extraction path) or (expression, extraction path, default value)	The expression, which is a JSON string is parsed. A field is extracted from the expression using the extraction path. If either the extpression or the path are invalid, an optional default value is returned.

Function Name	Arguments	Returns
ExtractKvp	(expression, extraction path) or (expression, extraction path, nested KVP delimiters /default "{}"/) or (expression, extraction path, nested KVP, delimiter / default ","/) or (expression, extraction path, nested KVP, delimiter, separator /default "="/) or (expression, extraction path, nested KVP, delimiter, separator, escape character / default "\\"/) or (expression, extraction path, nested KVP, delimiter, separator, escape character, default value)	The expression, which is a nested KVP string is parsed. A field is extracted from the expression using the extraction path. If either the expression or the path are invalid, an optional default value is returned.

Miscellaneous Functions

Function Name	Arguments	Returns
`matchcidr` You can use the function within SQL, EQL, and ECL queries, and in turn, in other functionality that make use of these queries.	`(IP_string_format, IP_address_to_Match, key)` In the `IP_string_format` parameter, you can specify a CIDR address, single IPv4 or IPv6 address, a comma-separated list of CIDR IP addresses or a range, a hyphen-separated range of IP addresses, or an enrichment list. In the `IP_address_to_Match` field, specify the column name from the logs, which are to be matched against the `IP_string_format` parameter. The data type of the column must be INET_ADDR. The `key` parameter specifies a key name in the enrichment list. The `key` parameter is mandatory when using this function in correlation bloks. You must provide a either key name or an empty string ("").	Returns those IP addresses from the `IP_address_to_Match` parameter, which belong to the IP addresses specified in the `IP_string_format` list.

Having Clause

The `Having` clause adds additional constraints on the events that have passed the filter and are grouped by the rule.

```
At (Least | Most) <integer>
Distinct <expression>
As <identifier>
Limit <integer>
Count Of <expression> Being <expression> (Greater | Less) Than <integer>
Percentage Of <expression> Being <expression> (Greater | Less) Than
<integer> %<condition>
```

Note: The `Having` clause expression must contain at least one aggregation function.

The supported parameters are:

Parameter	Description
Count Of	Counts the number of time two expressions are equals and check that this value is greater or less than a boundary.
Percentage Of	Counts the number of time two expressions are equals and make a ratio of this count versus the number of events in the group, then check whether the value is less or more than a value expressed as percent.

The `Having` clause can also be an expression using aggregation functions and resolving to a Boolean.

Aggregation Functions

Expressions used in the `Having` clause must contain at least one aggregation function. The `(*)` option applies the function to any event with no additional constraints. The `All` option applies the function to all values that are not null. The `Distinct` option applies the function only once per distinct values.

```
Count ( * )
Count ( [ Distinct | All ] <expression> Limit <integer> )
Sum ( [ Distinct | All ] <expression> Limit <integer> )
Avg ( [ Distinct | All ] <expression> Limit <integer> )
Max ( [ Distinct | All ] <expression> Limit <integer> )
Min ( [ Distinct | All ] <expression> Limit <integer> )
Var ( [ Distinct | All ] <expression> Limit <integer> )
Stdev ( [ Distinct | All ] <expression> Limit <integer> )
```

Function	Definition
Count()	Count of values
Sum	The total value
Avg	The average value
Max	The maximum value
Min	The minimum value
Var	The variance
Stdev	The standard deviation function

Having Clause

The `Having` clause adds additional constraints on the events that have passed the filter and are grouped by the rule.

```
At (Least | Most) <integer>
Distinct <expression>
As <identifier>
Limit <integer>
Count Of <expression> Being <expression> (Greater | Less) Than <integer>
Percentage Of <expression> Being <expression> (Greater | Less) Than
<integer> %<condition>
```

Note: The `Having` clause expression must contain at least one aggregation function.

The supported parameters are:

Parameter	Description
Count Of	Counts the number of time two expressions are equals and check that this value is greater or less than a boundary.
Percentage Of	Counts the number of time two expressions are equals and make a ratio of this count versus the number of events in the group, then check whether the value is less or more than a value expressed as percent.

The `Having` clause can also be an expression using aggregation functions and resolving to a Boolean.

Correlation Criteria

Correlation criteria can be of the following types:

- A join condition describing which fields should be equals in two event groups
- A sequencing constraint that describes the relative order in which two event groups should occurs
- An expression criteria that describes a condition among fields of different event groups

```
<event_group_identifier1> -> <field_identifier1> == <event_group_
identifier2> -> <field_identifier2> <event_group_identifier1>
(Begins | Ends) [At Least <integer> [ d | h | m | s ]] [Up To
<integer> [ d |h | m | s ]] (Before | After) <event_group_
identifier2> (Begins|Ends)
```

This is an expression criteria that is used to describe a condition between fields that belongs to different event groups.

```
<expression which uses syntax eventGroupIdentifier -> fieldIdentifier
for keys>
```

For example, `group1->sum_bytes >= group2->sum_bytes`

Note: The fields referenced in a join must be grouping fields for their respective event groups.

Correlation Blok (ECL) Examples

A correlation Blok can be used in Advanced Search on historic data for forensic analysis. It can also be used in a rule, which triggers in real time to generate alerts and notifications. Notifications can be in the form of email, syslog, or SNMP.

Blok	Definition
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [My EVENTS]`	Example 1: This Blok triggers a new alert at the first event and accumulate all events during 30 minutes time period.
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [My EVENTS] HAVING AT LEAST 1 DISTINCT [ll_ sourceDomain] HAVING AT LEAST 1 DISTINCT [ll_ deviceTypeID]`	Example 2: This Blok does the same as Blok Example 1, but the alerts generated then give information about the number of distinct `ll_sourceDomain` or `ll_deviceTypeID` and their values.

Blok	Definition
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [My EVENTS] WHERE [ll_ deviceTypeID] ="17" HAVING AT LEAST 2 DISTINCT [ll_ sourceIP]`	Example 3: This Blok filters events that have `ll_deviceTypeID` equal to "17", and at least 2 distinct values of `ll_sourceIP`
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [suspiciousSource s] AT LEAST 100 EVENTS WHERE [ll_ deviceTypeID] ="17" WITH THE SAME [ll_ sourceIP] HAVING AT LEAST 1 DISTINCT [ll_ eventStatus]`	Example 4: This Blok looks for at least 100 events with the same criteria as the previous one, coming from the same `ll_sourceIP` and giving information about the number of distinct `ll_ eventStatus` and their value.
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [suspiciousUsers] AT LEAST 100 EVENTS WHERE [ll_ deviceTypeID] ="17"`	Example 5: This Blok filters the event the same way as the previous one, and is looking for 100 events from the same `ll_sourceUser` that have at least 10 distinct values of `ll_sourceIP` and at most 1 distinct value of `ll_eventStatus`.

Blok	Definition
`WITH THE SAME [ll_ sourceUser] HAVING AT MOST 1 DISTINCT [ll_ eventStatus] HAVING AT LEAST 10 DISTINCT [ll_ sourceIP]`	
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [success] AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] = "success" WITH THE SAME [ll_ sourceUser],[ll_ sourceIP] EVENT GROUP [failed] AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] ="failure" WITH THE SAME [ll_ sourceUser],[ll_ sourceIP] CORRELATION success->[ll_ sourceIP]== failed->[ll_ sourceIP] success->[ll_`	Example 6: This Blok looks at two groups of events happening within 30 minutes. The first event group is success audit from the same `ll_sourceIP`/`ll_sourceUser` and the second group is failed status grouped the same way. The Blok is triggered if the fields grouped on both event groups are same.

Blok	Definition
sourceUser]== failed->[ll_ sourceUser]	
USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [success] is excluded AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] ="success" WITH THE SAME [ll_ sourceUser],[ll_ sourceIP] EVENT GROUP [failed] AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] = "failure" WITH THE SAME [ll_ sourceUser],[ll_ sourceIP] CORRELATION success->[ll_ sourceIP]== failed->[ll_ sourceIP] success->[ll_ sourceUser]== failed->[ll_ sourceUser]	Example 7: Same as the previous Blok, but this time the Blok is triggered if there are only failed events within 30m for the same `ll_sourceIP` / `ll_sourceUser`.

Blok	Definition
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [users] WHERE [ll_ eventStatus] ="failure" OR [ll_ eventStatus]="succ ess" WITH THE SAME [ll_ sourceUser] HAVING AT LEAST 2 DISTINCT [ll_ eventStatus]`	Example 8: This Blok looks for users that have events with `ll_eventStatus` equal to either failed or success.
`USE LogLogic_ Appliance WITHIN 30m EVENT GROUP [success] is excluded AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] ="Success" WITH THE SAME [ll_ sourceUser],[ll_ sourceIP] EVENT GROUP [failed] AT LEAST 1 EVENTS WHERE [ll_ deviceTypeID] ="17" AND [ll_ eventStatus] = "Failure" WITH THE SAME [ll_`	Example 9: Same as the previous Blok with an additional constraint that there are twice as many failed than success events.

Blok	Definition
`sourceUser],[ll_ sourceIP] CORRELATION success->[ll_ sourceIP]== failed->[ll_ sourceIP] success- >[ll_sourceUser]== failed->[ll_ sourceUser]`	
`use system WITHIN 20m EVENT GROUP eg1 WHERE matchcidr ("IPList",sys_ collectIP,"IP1")`	Example 10: For an enrichment list named `IPList`: `{"192.168.56.101/32":"IP1","191.163.56.101/32":"IP2","1 98.168.0.0- 198.168.255.255":"IP3","2001:4860:4860::8888/32":"IP4"}` and for `sys_collectIP=`"192.168.56.101",198.168.56.101", then:

REST API Reference

You can use the Representational State Transfer (REST) API to develop your own client application.

LogLogic LMI provides REST API that a client application can use to invoke services using simple HTTP methods. Starting from LogLogic LMI 6.1.1, queries via REST API are encrypted.

You can access LogLogic LMI API online documentation using the following URLs:

Category of services	Services provided	URL for the REST API
Query API	• Queries • Sub queries • Quick queries	`https://<hostName>:9681/docs` For more information, see REST API for Advanced Search.
Aggregation node API	Aggregation rules	`https://<hostName>:9685/docs`
Correlation API	• Correlation instances • Alerts • Triggers • Metrics	`https://<hostName>:9682/docs`
Monitoring console API	• Default • Authentication • Domain • Alert • Agent	`https://<hostName>:9687/HawkConsole/v1/docs` For more information, see TIBCO Operational Intelligence Hawk® Concepts Guide.

Category of services	Services provided	URL for the REST API
	• Invoke and Subscribe • Rulebase • Schedule • Microagent	

Note: The URLs include the default ports. If you change the default ports, update the URL accordingly.

To demonstrate the usage of REST API, some sample scripts are provided in the `REST/Samples/` directory in the supplemental package. For more information, see the `samples.txt` file in the supplemental package.

Authentication and Access

LogLogic LMI users who have Admin or Search Archived Data permissions can access REST API. If a user without these permissions tries to access REST API, authentication fails.

REST Request Format

REST API requests must be submitted in a specific format.

The format of a LogLogic REST API request is:

```
<METHOD> <baseurl>/<basePath>?<query_parameters>
```

where:

- `<METHOD>` is the HTTP method to be used on the resource (GET, POST, PUT, or DELETE)
- `<baseurl>` is the REST API Endpoint (baseurl).

- <basePath> is the part of the path that identifies the required LogLogic resource. It consists of:
    - a fixed part - for example, /api/v2/query when starting a query.
    - followed by (if required) path parameters - for example, id when starting a query
- <query_parameters> identifies any parameters to be passed as part of the request. Unless otherwise specified for a specific parameter, all query parameters are always optional. Multiple query parameters should be separated by ampersand (&) characters.

For example, the GET query request shows how to start a business service called GET query:

```
GET https://<hostName>:9681/api/v2/query/{id}
```

where:

- https://<hostName>:9681/ is the LogLogic REST API endpoint.
- api/v2 is the fixed part of the <basePath>.
- query is the path parameter in the <basePath>.
- {id} is the parameter of the path in the <basePath>.

REST API Endpoint (baseurl)

A specific endpoint must be used when submitting REST API requests.

The endpoint on which all LogLogic REST resources are exposed is:

```
<protocol>://<host>:<port>/api/<API_version_number>
```

where:

- <protocol> is the protocol used to communicate with the LogLogic runtime HTTP.
- <host> is the network name or IP address of the LogLogic runtime.
- <port> is the port used by the LogLogic runtime for incoming HTTP connections,

exposed by the HTTP connector shared resource.

- `api` is the context root used by the LogLogic REST API resource hierarchy.
- `<API_version_number>` is the API version number.

> **Note:** The LogLogic REST API endpoint is referenced throughout this document as `<baseurl>`.

Response Status Codes

All REST API requests return a response status code.

The main HTTP status codes that might be returned by LogLogic are:

HTTP status code	Description
200	Request completed successfully.
204	Request completed successfully but no content available to return.
400	Bad request/ Invalid query.
401	Authentication failure, invalid access credentials.
403	Insufficient permission.
404	<Component> id not found.
406	Not acceptable.
500	Unspecified internal server error.

REST API for Advanced Search

You can use REST API to run an Advanced Search query.

Querying LogLogic LMI using the REST API includes the following steps:

1. Create a query and obtain its ID.
2. Get status, results, or details from the query (in multiple invocations).
3. Delete the query after all the data has been obtained.

You can omit deleting the query if you create the query with a `timeToLive` parameter, in which case the query is automatically deleted after some time of inactivity.

> **Note:** TLSv1.2 and TLS v1.3 are supported for REST API.

Query Creation

This is a single synchronous call where the results are returned immediately after specifying the query:

```
POST <baseurl>/api/v2/query
```

The result can be any of the following items:

- Success: The query returns information about the query, including the ID number of the query and the schema of the results. When the query creation succeeds, although the call returns immediately, LogLogic starts the process of generating results in the background.
- Failure: The result is a failure if the query is invalid, for example, if there are syntax errors. In such case, the error details are returned.

Two more parameters can be specified in addition to the query itself:

- cached: Using this parameter, the results are cached temporarily. This allows you to retrieve any window of results from the total results, effectively allowing you to scroll the results up and down. It also allows you to run sub-queries, which is allowed only if the query is cached.

  > **Note:** Cached queries have a performance penalty and should be only used if needed.

- timeToLive: The time of inactivity in seconds, after which the query is automatically

deleted. The default is 0, which indicates never auto-deleting the query.

The query returns:

- The query ID
- The schema: An array of column descriptors, where each column descriptor contains the name and the type of the column. Similar to a header in a table, the column descriptor allows you to identify the values in each column. You use the column descriptor to interpret the results, as each row in the results is represented by a raw array.

API to Retrieve the Results

After creating the query and obtaining its ID, you can check its status, obtain details about the query, and retrieve results:

- `GET <baseurl>/api/v2/query/{id}/status`
- `GET <baseurl>/api/v2/query/{id}/details`
- `GET <baseurl>/api/v2/query/{id}/results`

The API to retrieve results takes the following parameters:

- id: The id of the query to get results from
- offset (optional): The offset of the results to be retrieved. Only available for cached queries. For non-cached queries, you cannot choose the offset; the results must be retrieved sequentially using multiple calls to this API.
- size (optional): The number of rows to return in this call. The API might return fewer or no rows, depending on the number of rows available.
- longPollTimeout (optional): The time in milliseconds to wait before returning results, if the result rows are not available. Only available for non-cached queries.

This API returns:

- rows: The array of rows. In most cases, it is a subset of the result rows. However, it could be all the results if there are only a few. Whether you need to fetch more results or not is indicated by the **hasMore** parameter. To fetch more rows, you must call this method again.
- offset: The offset of the subset of the rows returned.

- errorsOrWarnings: The details of error or warnings, if any.
- hasMore: A Boolean value. It is true if more rows are available to fetch, and false if there are no more rows to fetch.

Typically this API is repeatedly invoked to fetch the rows, until hasMore is false. After all the results have been fetched, you can delete the query.

API to Delete a Query

After using a query, you must delete it from the system; otherwise it continues using valuable resources. Deleting a query requires only the ID of the query:

```
DELETE <baseurl>/api/v2/query/{id}
```

Alternatively, you can set the `timeToLive` parameter while creating the query, so that the query is deleted automatically after the specified time of inactivity.

Sub-Queries

If a query returns too many rows, you can further refine the results by creating a sub-query. Similar to a filter, a sub-query adds modification to the original query, for example, sorting or grouping.

```
POST <baseurl>/api/v2/query/{id}/subquery/
```

This is especially useful for user interfaces that allow exploratory querying, where the user is not sure of what exactly to search.

Note: Use sub-queries only when needed, as they are supported only for cached queries and have a performance penalty compared to regular queries.

Creating a sub-query requires only the modification parameter. This parameter is an EQL fragment including the operations that must be applied. EQL and SQL both are supported.

After a sub-query is created, the APIs to retrieve results or status, delete, are equivalent to the APIs for the regular queries. The parameters and results are identical.

Examples of modifications:

Query	Description
`SORT BY ll_sourceUser`	Sort by the column ll_sourceUser
`GROUP BY ll_eventAction | COLUMNS ll_eventAction, COUNT(*)`	Get the count of events per ll_eventAction
`ll_device = 'MyDevice'`	Get only the events for device MyDevice

REST API for Aggregation Rules

The following table lists the REST API to perform the functions related to aggregation rules:

Resource API	Description
PUT /configuration/aggregationRule	Create an aggregation rule.
GET /configuration/aggregationRule	Retrieve all aggregation rules
POST /configuration/aggregationRule {ruleId}	Update an aggregation rule.
DELETE /configuration/aggregationRule {ruleId}	Delete an aggregation rule.
GET /configuration/aggregationRule {ruleId}	Retrieve an aggregation rule having the specified rule ID.
PUT /configuration/aggregationRule/validate/query	Validate an aggregation rule query.

REST API for Triggers and Trigger Groups

Using the correlation APIs, you can create a new trigger and test it on incoming data.

You can perform the following operations using REST API:

- Create a trigger and trigger group after creating a correlation Blok using ECL:

  ```
  PUT https://<host>:9682/api/v1/trigger-groups
  ```

  To create a trigger, use the following URL:

  ```
  PUT https://<host>:9682/api/v1/triggers
  ```

- Deploy or synchronize the trigger group:

  ```
  POST https://<host>:9682/api/v1/triggers/deploy
  ```

- Check the status of the deployed trigger group:

  ```
  GET https://<host>:9682/api/v1/trigger-groups
  ```

List of REST API services

REST API services to manage triggers

Resource API	Description
GET triggers	Get a list of triggers.
PUT triggers	Create a new trigger.
DELETE trigger/{id}	Delete a trigger from the node.
POST trigger/{id}	Update a trigger.
GET trigger/{id}	Get a trigger.

REST API services to manage trigger groups

Resource API	Description
GET trigger-groups	Get a list of trigger groups.
PUT trigger-groups	Create a new trigger group.
GET trigger-group/{id}	Get a trigger group.
POST trigger-group/{id}	Update a trigger group.
DELETE trigger-group/{id}	Delete a trigger group.
POST triggers/deploy	Deploy triggers.

Related topics

For information about how to manage triggers by using the GUI, see Managing Triggers.

REST API for Correlation Rules

Correlation instances are required to run correlation rules. A correlation instance includes an ECL ruleset, which is effective as soon as the correlation instance is active.

There are two types of instances, replay and real-time.

Replay Instances

Replay instances work on historical data. When you run an Advanced Search with a correlation Blok from the GUI, a corresponding replay instance is automatically created and executed.

You can perform the following operations on replay instances using REST API:

- Create a replay instance using ECL:

```
POST http://<host>:9682/api/v1/instances
```

Ensure that you specify `"instanceType": "replay"`. Note down the instance ID returned in the response body.

- Check the status of the instance to ensure it has finished processing the historical data:

```
GET https://<host>:9682/api/v1/instance/<instanceID>
```

`"status": "completed"` indicates that the instance has processed all the historical data in the system.

- Retrieve the correlation events to ensure the rule is operating as expected.

```
GET
https://<host>:9682/api/v1/instance/<instanceID>/correlationevents?
type=MESSAGES
```

- If the output is not as expected, update the instance with new ECL and view the new results.

```
POST https://<host>:9682/api/v1/instance/<instanceID>
```

After updating the instance, the modified rule is applied to the data and a new set of correlation events is generated.

- Delete the instance to free all the resources being used by the instance:

```
DELETE https://<host>:9682/api/v1/instance/<instanceID>
```

Real-Time Instances

Real-time instances work on real time events. When you create triggers and trigger groups from the GUI and synchronize trigger groups, a real-time instance is automatically created for each trigger group. The real-time instance includes all information associated with that trigger, such as correlation Blok and notification settings.

Real-time instances process only the new data ingested in the appliance. Therefore, after creating a real-time instance, you must wait until new data arrives before any alerts can be generated.

To create a real-time instance using REST API, use the same resource URL as that for creating replay instances; but specify `"instanceType": "realtime"`. REST API services for other operations are the same as those for replay instances.

Real-time instances generate alerts. To retrieve and acknowledge alerts using the REST API, see REST API for Advanced Alerts.

List of REST API services

Resource API	Description
GET activeInstances	Get a list of active instance states that currently exist in the node.
GET instances	Get a list of instances that currently exist in the node.
POST instances	Create a new instance and run it on the specified data.
GET instance/{id}	Get the summary of an instance.
POST instance/{id}	Update the state of an instance.
DELETE instance/{id}	Delete an instance from the node.
GET instance/{id}/correlationevents	Get correlation events for an instance. This supports only replay instances.
GET instance/{id}/correlationevent/{eventId}/group/{name}/eventKeys	Get a list of event references and sources. This supports only replay instances.
GET instance/{id}/result/{resultSetId}/columns	Get lists of correlation result set columns. This supports only for replay instances.
PUT validate/correlation/rule	Validate a correlation rule.

Limitations

You can use correlation instances from the REST API, but not from the GUI, because of the following reasons:

- A correlation instance created by using the REST API service cannot be viewed or edited from the GUI.
- A real-time correlation instance updated by using the REST API service is not synchronized with its corresponding triggers and trigger group created from the GUI. Similarly, a replay instance updated by using the REST API service is not synchronized with its corresponding search query created from the GUI.

REST API for Advanced Alerts

Using the correlation APIs, you can retrieve and acknowledge alerts. Real-time instances generate alerts; replay instances do not.

You can perform the following operations on alerts using REST API:

- Retrieve and view alerts:

```
GET
https://<host>:9682/api/v1/instance/<instanceID>/alerts?alertNameFi
lter=LoginFailureAttempt
```

- Retrieve alert details:

```
GET
https://<host>:9682/api/v1/instance/<instanceID>/alert/<alertID>
```

- Acknowledge alerts:

```
POST https://<host>:9682/api/v1/alerts
```

List of REST API services

Resource API	Description
GET instance/{id}/alerts	Get a list of alerts from an instance.
POST alerts	Acknowledge alerts.
GET instance/{id}/alert/{alertID}	Get the detailed summary of an alert.
GET instances/alert/severities	Get a list of alert severities.
GET instances/alert/categories	Get a list of alert categories.
POST instances/alert/fields	Get a list of alert fields for the rule and source environment. The alert fields are different depending on the rule and source environment.
GET instance/{id}/alert/{alertID}/group/{groupName}/ eventKeys	Get a list of event references and sources. This supports only realtime instances.

Related topics

For information about how to work with Advanced Alerts by using the GUI, see Advanced Alerts.

Reserved Keywords

A list of reserved keywords that cannot be used in user-defined fields.

The following keywords cannot be used as column names of data models.

Reserved keywords

MONTH	DAY	YEAR	HOUR
MIN	SEC	LENGTH	SUBSTR
SUM	AVG	USE	SUBSTRING
MID	LOWER	UPPER	TRIM
SUBSTITUTE	LEFT	RIGHT	FIND
CONCATENATE	POSITION	SECONDS	MINUTES
HOURS	DAYS	WEEKS	MONTHS
YEARS	IIF	TOTIMESTAMP	TOIP
TOSTRING	TOINT	TOLONG	TOFLOAT
TOBOOL	EXTRACTJSON	EXTRACTKVP	TOTIMESTAMPSTRING
INT	LONG	BOOLEAN	DATE
TIME	STRING	HEX	CASE
EXPRESSION			

Reserved keywords from the H2 database

ALL	CHECK	CONSTRAINT	CROSS

CURRENT_DATE	CURRENT_TIME	CURRENT_TIMESTAMP	DISTINCT
EXCEPT	EXISTS	FALSE	FETCH
FOR	FOREIGN	FROM	FULL
GROUP	HAVING	INNER	INTERSECT
IS	JOIN	LIKE	NULL
LIMIT	MINUS	NATURAL	NOT
OFFSET	ON	ORDER	PRIMARY
ROWNUM	SELECT	SYSDATE	SYSTIME
SYSTIMESTAMP	TODAY	TRUE	UNION
UNIQUE	WHERE	WITH	

TIBCO Documentation and Support Services

For information about this product, you can read the documentation, contact TIBCO Support, and join TIBCO Community.

How to Access TIBCO Documentation

Documentation for TIBCO products is available on the TIBCO Product Documentation website, mainly in HTML and PDF formats.

The TIBCO Product Documentation website is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the TIBCO LogLogic® Log Management Intelligence Product Documentation page.

- *TIBCO LogLogic® Log Management Intelligence Release Notes*
- *TIBCO LogLogic® Log Management Intelligence Administration*
- *TIBCO LogLogic® Log Management Intelligence Configuration and Upgrade*
- *TIBCO LogLogic® Log Management Intelligence Hardware Installation Guide*
- *TIBCO LogLogic® Log Management Intelligence Log Source Report Mapping*
- *TIBCO LogLogic® Log Management Intelligence Security Guidelines*
- *TIBCO LogLogic® Log Management Intelligence SSD Hardware Field Installation*
- *TIBCO LogLogic® Log Management Intelligence Syslog Alert Message Format Quick Reference*
- *TIBCO LogLogic® Log Management Intelligence User Guide*
- *TIBCO LogLogic® Log Management Intelligence XML Import/Export Entities Reference*
- *TIBCO LogLogic® Enterprise Virtual Appliance Quick Start*

Other TIBCO Product Documentation

The following documents for TIBCO LogLogic® Log Source Packages are available on the TIBCO eDelivery website or TIBCO Support website after logging in.

- *TIBCO LogLogic® Log Source Packages Release Notes*
- *TIBCO LogLogic® Log Source Packages Installation and Upgrade*
- *TIBCO LogLogic® Log Source Packages Log Configuration Guides*
- *TIBCO LogLogic® Log Source Packages Log Collector Guides*

How to Contact TIBCO Support

Get an overview of TIBCO Support. You can contact TIBCO Support in the following ways:

- For accessing the Support Knowledge Base and getting personalized content about products you are interested in, visit the TIBCO Support website.
- For creating a Support case, you must have a valid maintenance or support contract with TIBCO. You also need a user name and password to log in to TIBCO Support website. If you do not have a user name, you can request one by clicking **Register** on the website.

How to Join TIBCO Community

TIBCO Community is the official channel for TIBCO customers, partners, and employee subject matter experts to share and access their collective experience. TIBCO Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from TIBCO products. In addition, users can submit and vote on feature requests from within the TIBCO Ideas Portal. For a free registration, go to TIBCO Community.

Legal and Third-Party Notices

SOME TIBCO SOFTWARE EMBEDS OR BUNDLES OTHER TIBCO SOFTWARE. USE OF SUCH EMBEDDED OR BUNDLED TIBCO SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED TIBCO SOFTWARE. THE EMBEDDED OR BUNDLED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER TIBCO SOFTWARE OR FOR ANY OTHER PURPOSE.

USE OF TIBCO SOFTWARE AND THIS DOCUMENT IS SUBJECT TO THE TERMS AND CONDITIONS OF A LICENSE AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED SOFTWARE LICENSE AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER LICENSE AGREEMENT WHICH IS DISPLAYED DURING DOWNLOAD OR INSTALLATION OF THE SOFTWARE (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH SOFTWARE LICENSE AGREEMENT OR CLICKWRAP END USER LICENSE AGREEMENT, THE LICENSE(S) LOCATED IN THE “LICENSE” FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

TIBCO, the TIBCO logo, the TIBCO O logo, and LogLogic are either registered trademarks or trademarks of TIBCO Software Inc. in the United States and/or other countries.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only.

This software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the readme file for the availability of this software version on a specific operating system platform.

THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. TIBCO SOFTWARE INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS DOCUMENT AT ANY TIME.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "READ ME" FILES.

This and other products of TIBCO Software Inc. may be covered by registered patents. Please refer to TIBCO's Virtual Patent Marking document (https://www.tibco.com/patents) for details.