Creating Content Packs in vRealize Log Insight 4.5

TE CHNI CA L W HI TE P APE R

vmware®

Table of Contents

Introduction

Content packs are read-only plug-ins to vRealize™ Log Insight™ that provide pre-defined knowledge about specific types of events such as log messages. The purpose of a content pack is to provide knowledge about a specific set of events in a format that is easily understandable by administrators, engineers, monitoring teams, and executives. A content pack should answer questions like, *"Is the product/application healthy?"* In addition, a content pack should create a greater understanding of how a product/application works.

A content pack comprises of information that can be saved from either the Dashboards or Interactive Analytics pages in vRealize Log Insight. This includes:

- Queries
- Fields
- Aggregations
- Alerts
- Dashboards
- Agent Groups
- Setup Instructions

By default, the current version of vRealize Log Insight ships with the vSphere and General content packs. Other content packs can be imported as required. In addition, any vRealize Log Insight user can create a content pack for private or public consumption.

Intended Audience

This paper provides information about each piece of information that can be saved in a content pack, as well as best practices for content pack creation. The information provided is specificay tailored to content pack authors using vRealize Log Insight 2.5 and newer. (Note: Log Insight 2.5, 3.0, 3.3, 3.6, 4.0 and 4.3 content packs are all compatible.)

Getting Started

Before creating a content pack, it is important to understand some concepts regarding the content pack workflow. The tips in this section will make creating and maintaining content packs easier.

Instance

Content packs are read-only plug-ins to vRealize Log Insight, which means imported content packs, cannot be edited. The easiest way to edit a content pack is to modify the saved definitions on the instance of vRealize Log Insight that was used to initially create the content pack. The original instance should be backed up to prevent data loss or corruption. If the instance used to create the content pack is lost and no backup exists, the content pack must be recreated on a new instance. Although certain components of a content pack can be cloned into a custom dashboard, also known as user space, doing so is not a recommended way to edit a content pack and might result in a content pack that is dependent on a separate content pack.

Alternatively, you can import a content pack into My Content (user space) and edit the content pack. However, if you have other widgets (dashboards, alerts, extracted fields and queries) from before you do the import, ensure you save and remove them before you import to avoid mixing of original content with the imported content.

vmware®

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Ideally you should create a new user via Administration \ Access Control to edit a content pack and import the content pack into My Content (user space) to edit a content pack so as to avoid any mixing of original content with the imported content,

User

Content packs are created in part from the content saved under Custom Dashboards, or more specifically either My Dashboards or Shared Dashboards on the Dashboards page. When exporting a content pack, everything within the selected custom dashboard is exported. For this reason, it is recommended that every individual content pack be authored by a separate user entity in vRealize Log Insight. For information on creating users in vRealize Log Insight, please refer to the vRealize Log Insight documentation.

Events

It is essential to collect relevant events before attempting to create a content pack, to ensure that the content pack covers all relevant events for a product/application. A common way to collect relevant events is with the assistance of quality assurance (QA) and/or support teams, because these teams usually have access to, and knowledge about, common events. Attempting to generate events while creating a content pack is time consuming and will likely result in missing important events. If QA and support teams are unable to supply events, simulated events could be used instead, assuming that product/application events are known and/or documented.

Once appropriate logs have been collected, they must be ingested into vRealize Log Insight. In the current version of vRealize Log Insight, it is possible to ingest events from the command line using the vRealize Log Insight Importer. In short, any file, directory, tarball, or ZIP file can be ingested by copying the events to the vRealize Log Insight virtual appliance and running:

For MS-windows you can use the tool as follows: (the tool is also available for the different flavors of Linux and can be used in a similar manner)

C:\my_logs>loginsight-importer.exe –manifest myLogsManifest.txt –source myLogs.tar –server 10.123.345.567 –debug_level 2 –logdir c:\my_logs Where:

- myLogsManifest.txt – is a Manifest file explaining to the importer how you would like to import your logs, this allows you to parse your logs if you so wish before they are ingested into vRealize Log Insight

  e.g. Contents of the myLogsManifest.txt –

  ```
  [filelog|data_logs] directory=D:\Logs
  include=*.txt
  parser=mysyslog
  tags={"product":"HP"}

  [parser|mysyslog]
  base_parser=syslog
  message_decoder=syslog_message_decoder
  debug=yes

  [parser|syslog_message_decoder
  ] base_parser=kvp fields=*
  ```

- myLogs.tar – This is your static logs in zip or tar format

- 10.123.345.567 – IP of your vRealize Log Insight server instance or vRealize Log Insight VIP
- 2 – debug level as 2 tells the importer to log details of the static import which can be essential for debugging purposes.
- c:\my_logs – The folder where the debug level log messages from the importer will be written For a detailed list of command line options for the importer tool refer to product documentation.

Authors

The authors of a content pack should possess the following competencies:

- Experience using VMware vRealize Log Insight.
- Real-world operating knowledge of the product/application.
- Understanding of and ability to generate optimized regular expressions.
- Experience with using logs to debug multiple problems with the product/application.
- Support background, with exposure to a variety of problems.
- System administrator background with previous syslog experience.

Queries

vRealize Log Insight allows queries to retrieve and summarize events. Queries can be created and saved from the Interactive Analytics page. A query comprises one or more of the following:

QUERY ELEMENT	DESCRIPTION
Keywords	Complete, or full-text, alphanumeric and/or hyphen characters
Globs	Asterisk (one or more) and/or question mark (exactly one) symbol used to match some quantity of keywords.
Regular expressions	Sophisticated string pattern matching, based on Java regular expressions.
Field operations	Keyword, regular expression, and pattern matches applied to extracted fields.
Aggregations	Functions that are applied to one or more subgroups of the results.

vRealize Log Insight supports the following types of queries:

QUERY TYPE	DESCRIPTION
Message	A query formed of keywords, regular expressions and/or field operations.
Regular expression or field	A query formed of keywords and/or regular expressions.
Aggregation	A query formed of a function, one or more groupings, and any number of fields.

Custom alerts can be defined in vRealize Log Insight and are triggered from scheduled queries of any type.

Saving Queries

Queries can be saved using one or more of the following methods:

METHOD	DESCRIPTION
Add to Dashboard	Saves the last-run query without time range as a chart, query in a query list, or field table widget in a dashboard group on the Dashboards page.
Save Current Query	Saves the last-run query with a time-specific time range as a loadable query on the Interactive Analytics page. Queries that are saved using Save Current Query that are exported as part of a content pack do not include any time range.

Figure 1. Note the Add to Dashboard link just below the bar on the Interactive Analytics page.

Figure 2. The Save Current Query… link under the menu drop-down navigation on the Interactive Analytics page.

Adding Query Notes

The notes section is very important and should be populated for every query. Information can be added as text, a link to documentation, a knowledge base article, or a forum. Information provided should answer the following questions:

- Why is this widget important?
- What is a “good” and a “bad” value?
- Where can more information be obtained?

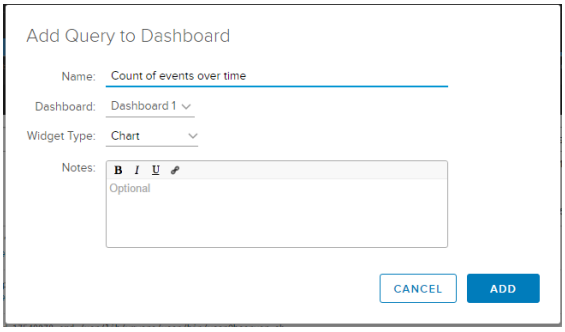

Figure 3. Add to Dashboard dialog box with notes section.

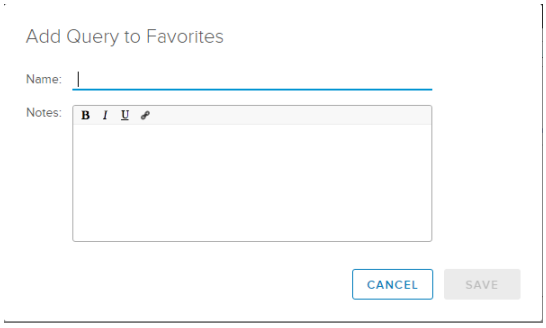

Figure 4. Save Current Query dialog box with notes section

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Message Queries

Message queries can be created using one or more of the following methods:

CREATING MESSAGE QUERIES	DESCRIPTION
Search bar	The search bar is one way to refine the results that are returned, given the existing events in a vRealize Log Insight instance. Although a constraint can be used instead of the search bar, it is often easier to understand a query that leverages the search bar over an equivalent constraint. As such, best practice is to use the search bar whenever possible, instead of an equivalent constraint.
Filters	A filter allows querying using a regular expression, a field, logical OR and AND operations, or a combination of search bar and constraint queries.

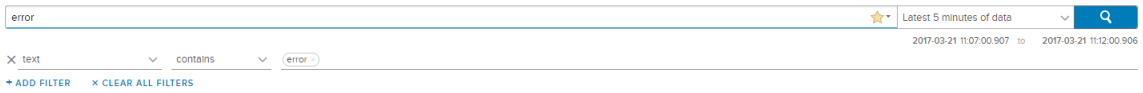

Figure 5. An example of the search bar with a keyword and a constraint with an equivalent query. Using the search bar is preferential.

Figure 6. An example of the search bar with a keyword, a constraint with a regular expression, and a constraint with a field operation. In order for the query to return a result, all three items need to return a match.

Creating Queries – Best Practices

Although query building is beyond the scope of this document, there are several important things to know about the search bar and constraints when creating content packs. In general, the following best practices apply:

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

- When constructing a query, use keywords whenever possible. When keywords are not sufficient, use globs and when globs are not sufficient, use regular expressions. Keyword queries are the least resource-intensive query type. Globs are a simplified version of regular expression and are the next least resource-intensive type of query. Regular expressions are the most resource-intensive query type and adversely affect query performance.
- Avoid regular expressions whenever possible. If a query can be written without regular expressions, it should be. This is primarily because, from a resource perspective, regular expressions are the most intensive query type. Leverage globs instead of regular expressions when keywords are not sufficient.
- Provide as many keywords as possible. When using regular expressions or fields, be sure to include as many keywords as possible. Keywords should be outside any regular expressions, including a logical OR such as (*this*|*that*). Regular expressions use a lot of resources. Keyword queries are the least resource-intensive query type and vRealize Log Insight is optimized to implement keyword queries before regular expressions, to minimize regular expression overhead.

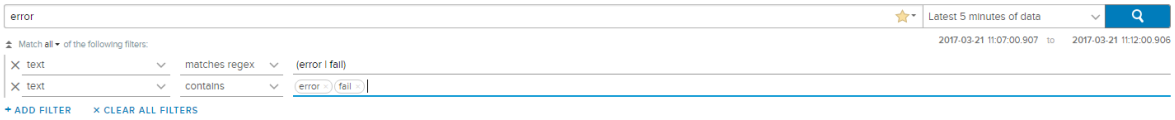

Figure 7. An example of two different ways to construct the same query. The first constraint is a regular expression. The second is a keyword, comma separated, logical OR match. The second constraint is always preferred over the first.

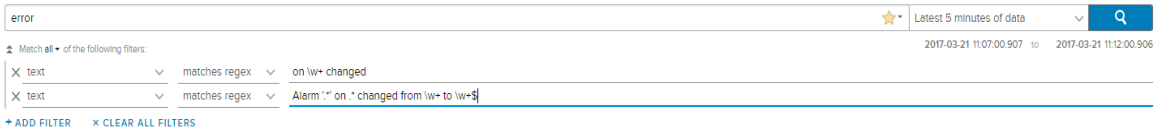

Figure 8 . An example of two different ways to query for the same field.The first constraint is generic and contains only two keywords; second constraint is specific and has five keywords. The second constraint is always preferred over the first.

Field Queries

Fields are a powerful way to add structure to unstructured events and allow for the manipulation of both the textual and visual representation of data. Fields are one of the most important items in a content pack because they can be used in multiple ways.

FIELDS	DESCRIPTION
Aggregations	Allowing for functions and groupings to be applied to fields.

Filters	Allowing for operations to be performed against fields. Any part of a log message that might be applicable to a query or aggregation should be extracted. Fields are a type of regular expression query and are especially useful for complex pattern matching, so a user does not need to know, remember, or learn complicated regular expressions.
Regex before value	This field should include as many keywords as possible. If the field is empty or only contains special characters, the Regex before value must include keywords.
Regex after value	This field should include as many keywords as possible. If this field is empty or only contains special characters, the Regex after value must include keywords.
Name	Only use alphanumeric characters. Ensure that all characters are lower case and use underscores instead of spaces as this makes fields easier to view. Important: Names for content pack fields and user fields can be the same, although content pack fields will have a namespace in parenthesis to the right of the field name. It is recommended to prefix content pack fields with an abbreviation (for example, vmw_) to avoid confusion.

FIELD TYPE	DESCRIPTION
Static	Static fields such as timestamp, host name, source, and appname are extracted at ingestion by Log Insight and their field definitions are nonmodifiable. Also any fields parsed out of the logs from the Log Insight agent via the parser or tags are also static fields in Log Insight with nonmodifiable field definition. We strongly recommend the use of static fields as additional context whenever possible which means we highly recommend the use of parsers for field extractions.
Extracted Fields	Any part of a log message that might be applicable to a query or aggregation can be dynamically extracted from the data by providing a regular expression. Extracted Fields are a type of regular expression query and are especially useful for complex pattern matching, so a user does not need to know, remember, or learn complicated regular expressions. However, if the regex definition of the field is not optimized for performance it can considerably slow down query performance. For e.g. a field definition with a pre-context as " **error.* name = ' "**and a post-context as **' ;** causes the query performance to be extremely slow, however changing the field definition to have pre-context as **"name = ' "** and post-context as **' ,** with additional context keyword as **error** doubles query performance.
Smart Fields	Machine learning analyzes events and discovers fields that similar log messages contain. The default name of a smart field is of the format smart field - *type number* [*event_type*]. You can rename a smart field and delete a smart field but you cannot modify its definition and is treated like a static field thereon. We strongly recommend the use of static fields as additional context whenever possible.

Additional Context

Beginning in vRealize Log Insight 2.5, you can add **Additional context** to refine your search and improve query performance. In the Additional context field, you can:

- Add keywords
- Add a filter on a static field with an operator and value

It is highly recommended to add additional context and more specifically a filter on a static field to every extracted field.

Multiple VIPs

Beginning in vRealize Log Insight 3.3, the integrated load balancer now allows for multiple VIPs to be configured with zero or more tags. This makes it possible to tag ingested log messages for devices that

cannot leverage the Log Insight agent and offers a query performance boost for content packs with limited keywords or varied log formats. These tags can be added as additional context on an extracted field to make the fields and hence the logs context specific.

If it is not possible to add keywords to an extracted field and it is not possible to add an existing static field to an extracted field, then the multiple VIP + tags feature must be used.

For e.g. - `2013-12-26T15:18:10.000-08:00</ent1:dateGenerated><ent1:lockFlag/><ent1:notes/></ent1:key><ent1:key xmlns:ent1="http://www.vmware.com/it/mw/entitlement/EntitlementManagement" id="1005259510"><ent1:value>J169M-65101-K89T8-AYWA2-88635</ent1:value><ent1:type><ent1:code>CLOUD</ent1:code></ent1:type><ent1:status><ent1:code>ACTIVE</ent1:code></ent1:status><ent1:quantity>200</ent1:quantity><ent1:dateGenerated>`

Creating Queries – Best Practices

In addition to the various components that comprise a field, several best practices must be considered.

- Only create fields for regular expression patterns. If a field can be queried using keyword queries, use keyword queries instead of a pre-defined field. Fields are intended to add structure to unstructured data and to provide a way to query specific parts of an event.
- Only create fields for regular expression patterns that return a fraction of the total events. Fields that match most events and/or return a very large number of results are not good candidates for field extraction because the regular expression will need to be applied to a large volume of events, resulting in a resource-intensive operation.
- When using filters in queries, do not use the match “any” operator unless one or more keywords are defined in the search bar.
- When using the text filter with multiple different values, one or more keywords should be defined in the search bar.
- Understanding of what "any" means vs "all": "any" means that each filter is a SEPARATE query -- so when multiple filters are used with ‘any’ operator it is actually multiple queries. In general more the queries, the slower the results. Think of "any" as "or" and "all" as "and" operators.
- Matching AQ (aggregated query) to MQ (message query) is not required for reasons mentioned above when the match “any” operator is used.

Examples

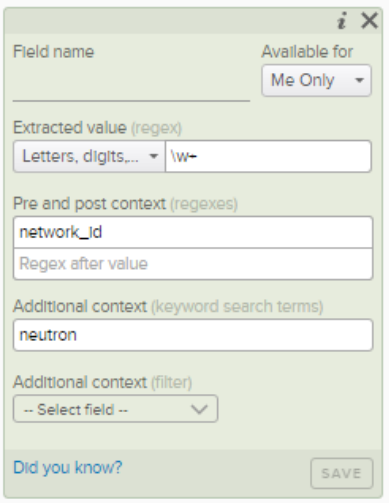

Figure 9. An example of an extracted field definition with multiple keywords

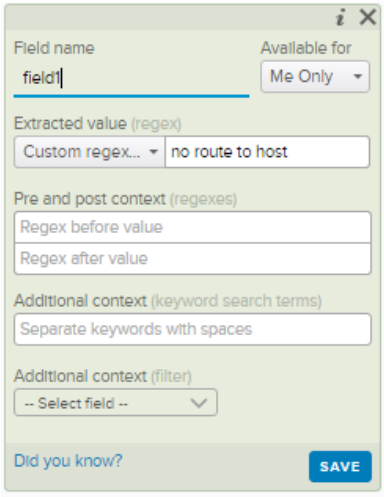

Figure 10.An example of a keyword field. Since this query can be constructed without a regular expression, it is not a good candidate for field extraction.

Figure 11. The recommended way to query for keyword matches. Information entered into the search bar or a constraint can also be saved for future usage by clicking the menu drop-down next to the Search button and selecting Save Current Query.

Temporary Fields

It is common for queries to contain one or more fields. For saved queries, it is important to note that the field definition used when a query is saved is always maintained. This means that, if a query is saved with a field and that field is later modified, the query will be modified when you update the field definition. In fact, if the field is used in other widgets such as dashboard chart or alert queries, those queries are also updated. Field modifications include:

- The *value* of the field is changed.
- The *regex before value* and/or the *regex after value* of the field are changed.
- The *name* of the field is changed.
- The *additional context* of the field is changed.
- The field is deleted.

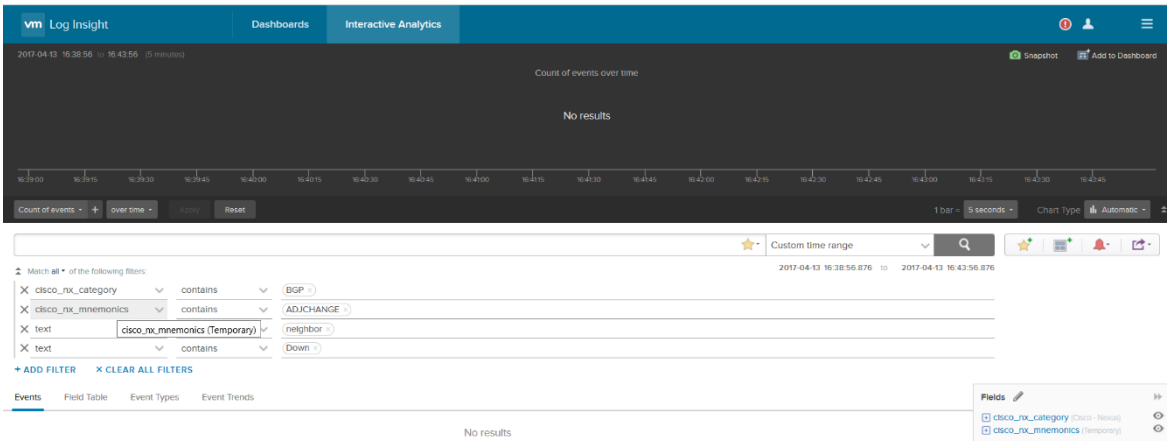

Figure 12. An example of running a query with a temporary field. Notice that the overview chart is grouped by cisco_asa_acl_hit_cnt and the cisco_asa_ace_hex field listed under the Fields section with (Temporary). This means the field does not exist in the vRealize Log Insight instance, but does exist as part of a chart widget or saved query.

Temporary fields are visible when running a saved query in the Interactive Analytics page, because the namespace *(Temporary)* appears next to the field name in the Fields section. Important: Saving, deleting, or modifying the field results in any use of the temporary field being removed from the query.

Ensure that content pack queries do not contain temporary fields. If a temporary field is found, recreate the saved query and delete the old saved query to remove the temporary field. To remove a temporary field from a chart widget:

1. Go to the widget on the Dashboards page.
2. Select the Edit in Interactive Analytics gear button within the widget.
3. Modify the field(s) used.
4. Select the Save followed by Return to Dashboard button on the Interactive Analytics page.

Aggregation Queries

vRealize Log Insight allows visual manipulation of events through the use of aggregation queries. An aggregation query is made up of two distinct attributes:

- Functions
- Groupings

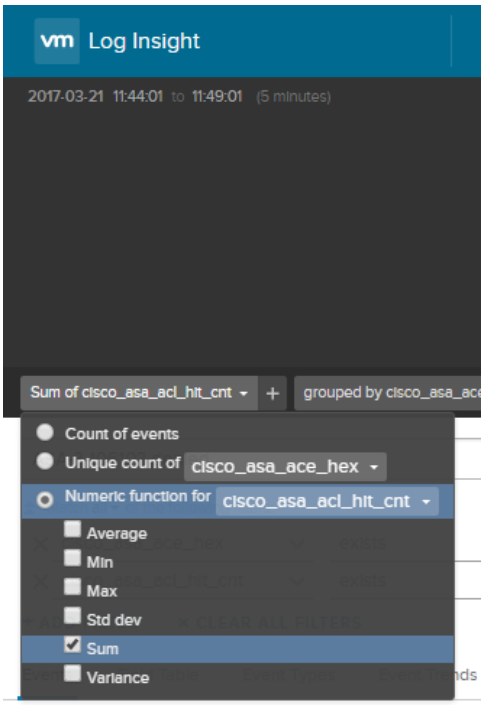

Figure 1 Functions

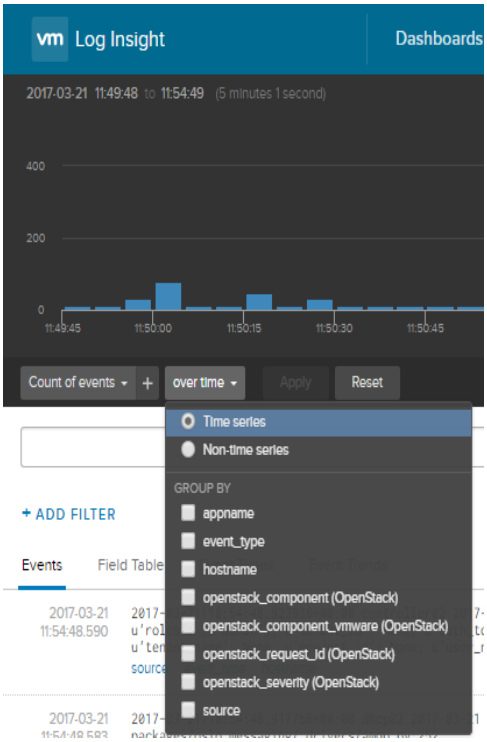

Figure 2 Groupings

In content packs, groupings are the most important consideration, but both functions and groupings will be addressed as they impact how charts are displayed.

An aggregation query requires one function and at least one grouping.

Charts

vRealize Log Insights displays analytics in different types of charts:

- Bar Charts
- Line Charts
- Stacked Charts
- Multi-colored Charts
- Table Charts
- Other Charts

Bar Charts

By default, the Interactive Analytics page of vRealize Log Insight displays a count of events over time in the overview chart. If the count function is used in conjunction with the time series grouping, a bar chart is created.

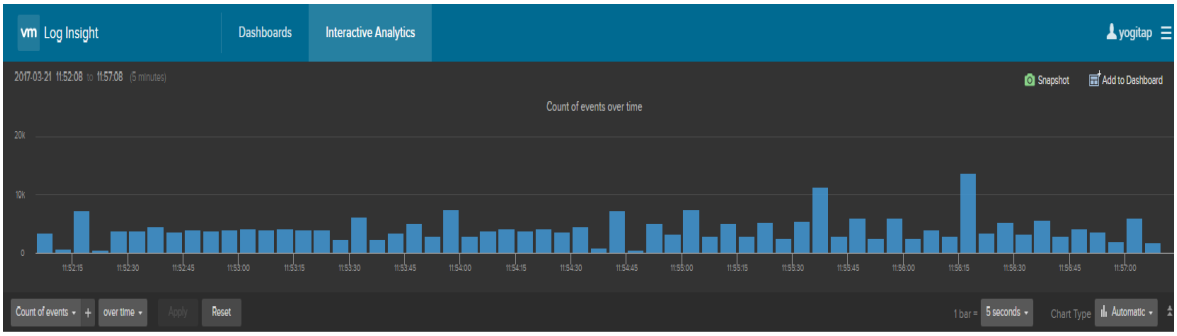

Figure 13. An example of a bar chart using count of events over time.

If the count function is used in conjunction with a single field grouping instead of time series, a bar chart is created with quantities listed from greatest to least.

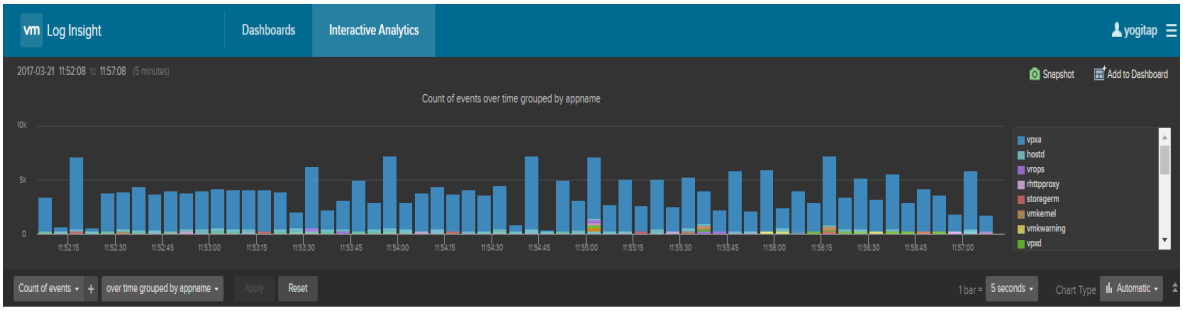

Figure 14. An example of a bar chart using count of events grouped by a field.

Line Charts

All functions, except the count function, are mathematical and require a field against which to apply the equation. When performing a mathematical function on a field and grouping by time series, a line chart is created.

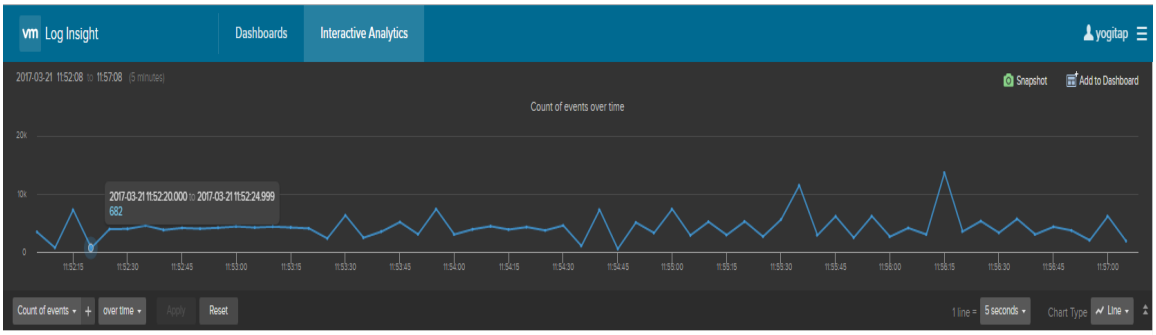

Figure 15. An example of a line chart using average of a field over time.

Stacked Charts

By default, the overview chart on the Interactive Analytics page of vRealize Log Insight is a count of events over time. If one field is added to the time series grouping, a stacked chart is created.

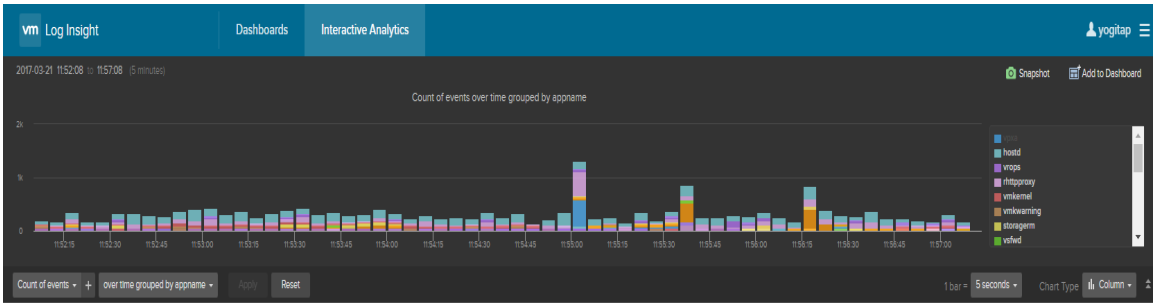

Figure 16. An example of a stacked bar chart using count of events over time with a field.

If grouping by time series, and a field and any function other than count is used, a stacked line chart is created.

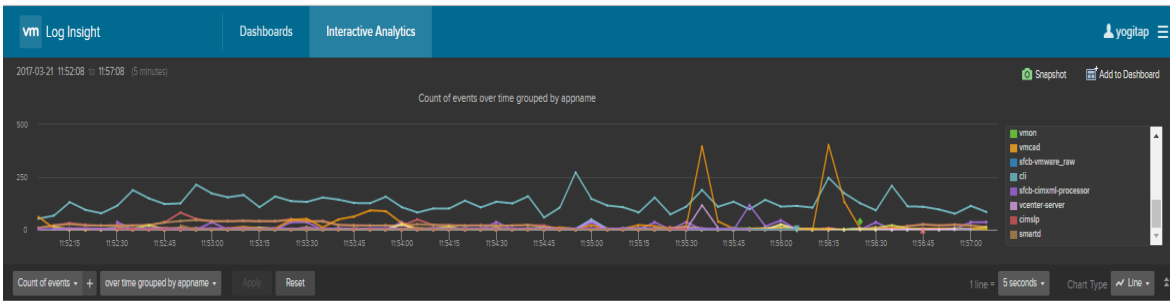

Figure 17. An example of a stacked line chart using average of a field over time grouped by a field.

Stacked charts are powerful when attempting to find anomalies for an object. Consideration needs to be given to the number of objects that could be returned. In general, the following best practices apply:

If the number of objects per bar returned will be less than ten, stacked charts are encouraged.

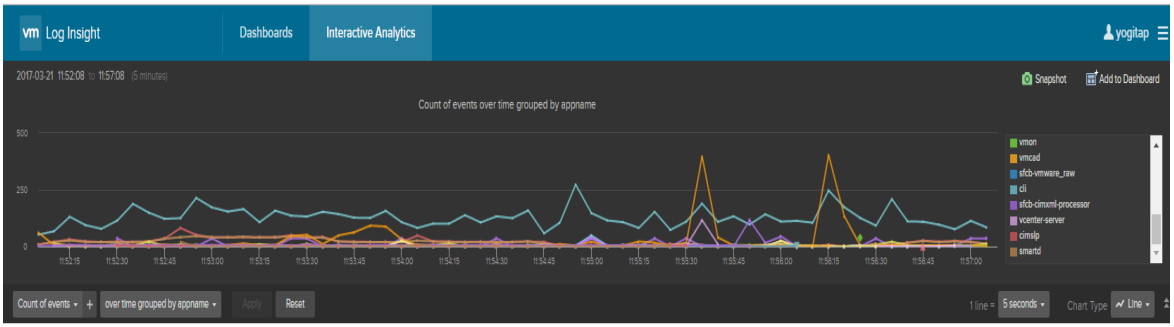

Figure 18. An example of a stacked line chart with a small number of objects. The chart is easy to read and understand.

If the number of objects returned per bar is or could be 10-20, stacked charts are good, but consideration must be taken when visually representing the chart in a content pack.

If the number of objects returned per bar is or could be greater than 20, stacked charts are discouraged.

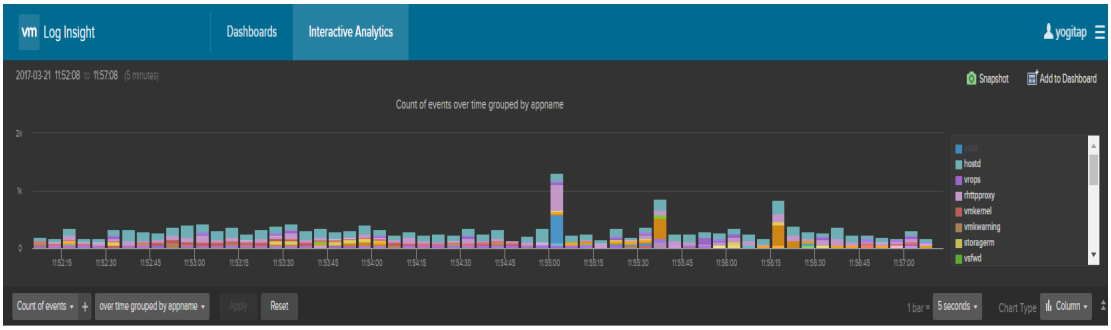

Figure 19. An example of a stacked bar chart with a large number of objects. The chart is hard to read and understand.

The recommendations above are made because a greater number of objects mean more resources are necessary to parse and display information. In addition, distinguishing between objects can become challenging when a large number of objects are returned.

Multi-Colored Charts

If a grouping is created using more than one field and time series, a multi-colored chart is created. The chart consists of two colors that interchange. Each interchange represents a new time range. Multi-colored charts can be hard to interpret so consider the value of such a chart before including it in a content pack.

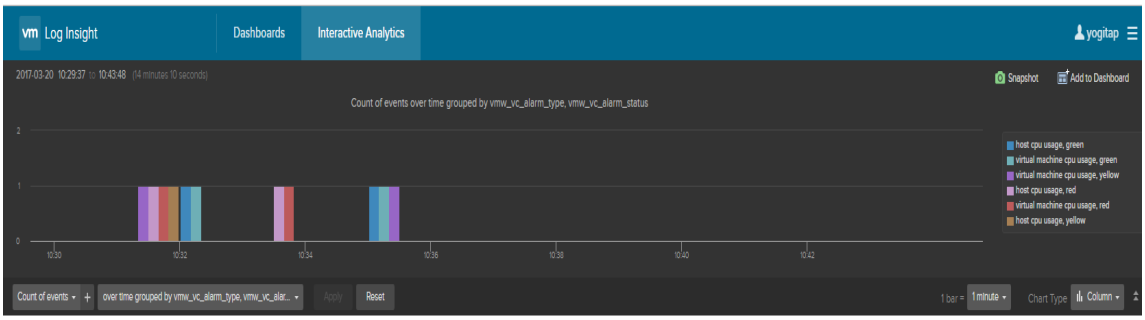

Figure 20. An example of a multi-colored bar chart using count of events over time, grouped by two fields.

When grouping by multiple fields, consider removing the time series for a more easily understood bar chart.

Figure 21. An example of a multi-field grouping bar chart using count of events, grouped by two fields.

If multiple fields are important over a time range, multiple charts could be created for each field individually over the time range. The charts could then be displayed in the same column of a dashboard group in a content pack.

Figure 22.An example of two similar charts stacked.Notice how one red alarm in blue matches mostly pink sources.

Table Charts

By default, the overview chart on the Interactive Analytics page of vRealize Log Insight is a bar chart. If one field is added to the time series grouping, you now have the option of viewing the chart as a table, you could create a table chart without any grouping but the data displayed may not be as meaningful.

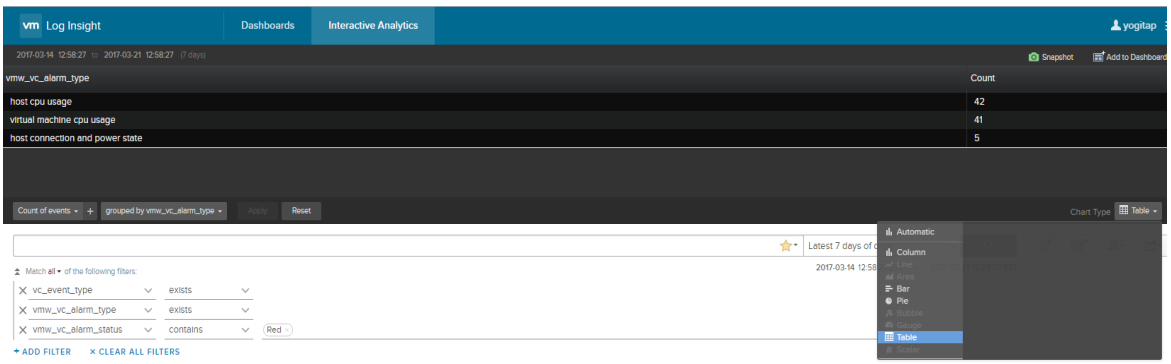

Figure 23. An example of a table chart using count of events over time with a field.

Other Charts

Several other chart types are available, including overlay, pie and bubble charts. To use these charts, a specific query type is required. If the option for these charts is available, you already have the correct query. If the option for these charts is not available, hover over the chart name you want to use. A pop-up message describes the type of query required for the chart type.

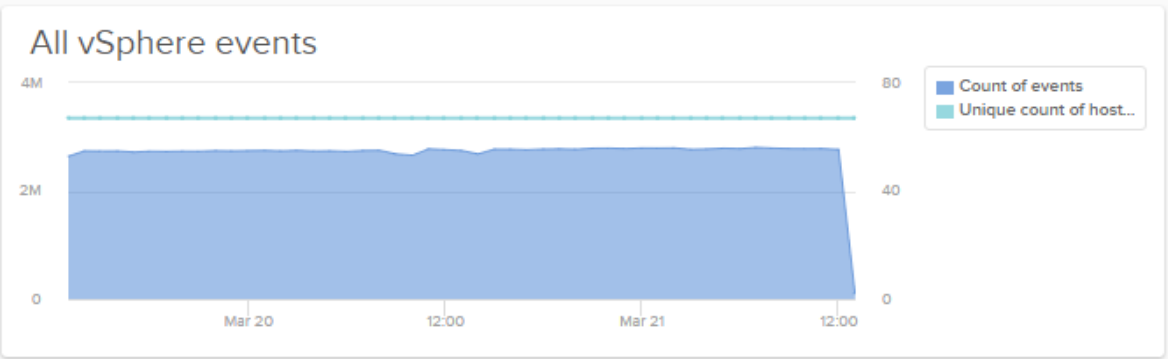

Figure 24. An example of overlay charts. Unique Count of hosts over Count of all vSphere events.

Message Queries

When constructing an aggregation query, the message query should only return results that are relevant to the aggregation query. This makes analyzing easier and ensures only relevant fields are shown.

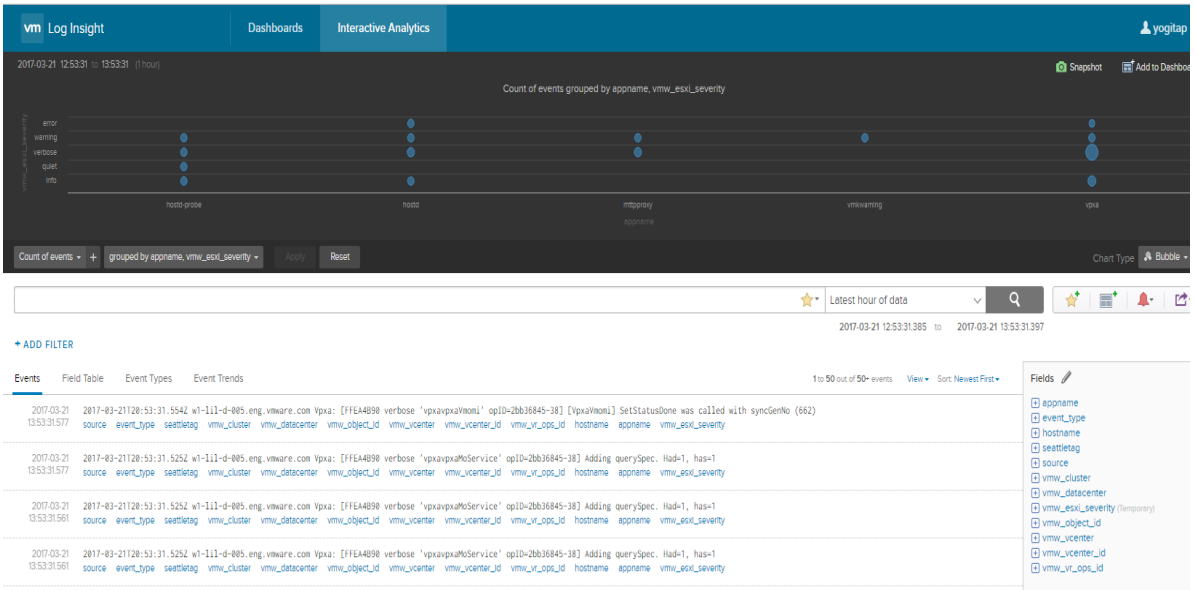

Figure 25. An example of an aggregation query without a message query. This is not recommended.

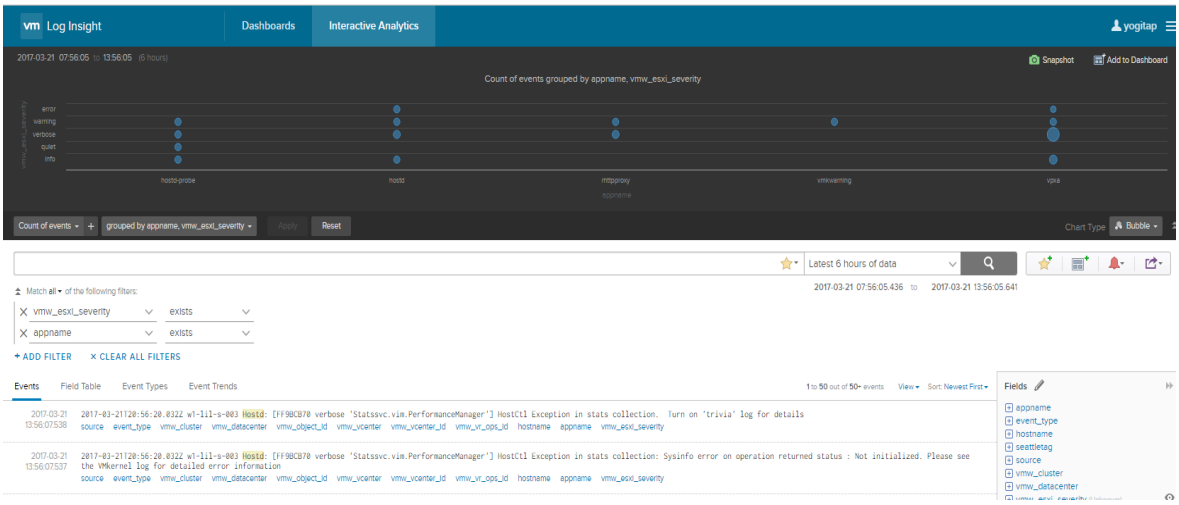

Figure 26. An example of an aggregation query with a message query. This is recommended. Notice the addition of filters for fields in the aggregation query with exists operator.

Alerts

Alerts provide a way to trigger a reaction when a certain type of event is seen. By default, vRealize Log Insight supports three different types of alerts:

- Email
- Webhook
- vRealize Operations Manager

Alerts can only be saved in user space and as such, all content pack alerts are disabled by default. If an enabled alert is created and then exported as part of a content pack, the alert is disabled in the content pack. This means that email, webhook and/or vRealize Operations Manager settings are not contained and cannot be added to a content pack.

Thresholds

It is important to understand how thresholds work to ensure that, if enabled, a content pack alert does not unintentionally spam a user. When considering a threshold, there are two things to keep in mind:

- How frequently to trigger the alert: vRealize Log Insight comes with pre-defined trigger frequencies. **Important**: Alerts only trigger once for a specific threshold window.
- How often to check if an alert state has occurred: An alert is triggered by a query. Alerts, such as queries, are not real-time in the current version. For each threshold window, a pre-determined query frequency has been allocated. Changing the threshold changes the query time.
- Alert can also be raised every time a new event type is seen but this can be noisy.

For alerts defined in a content pack, the “On any match” threshold should not be used.

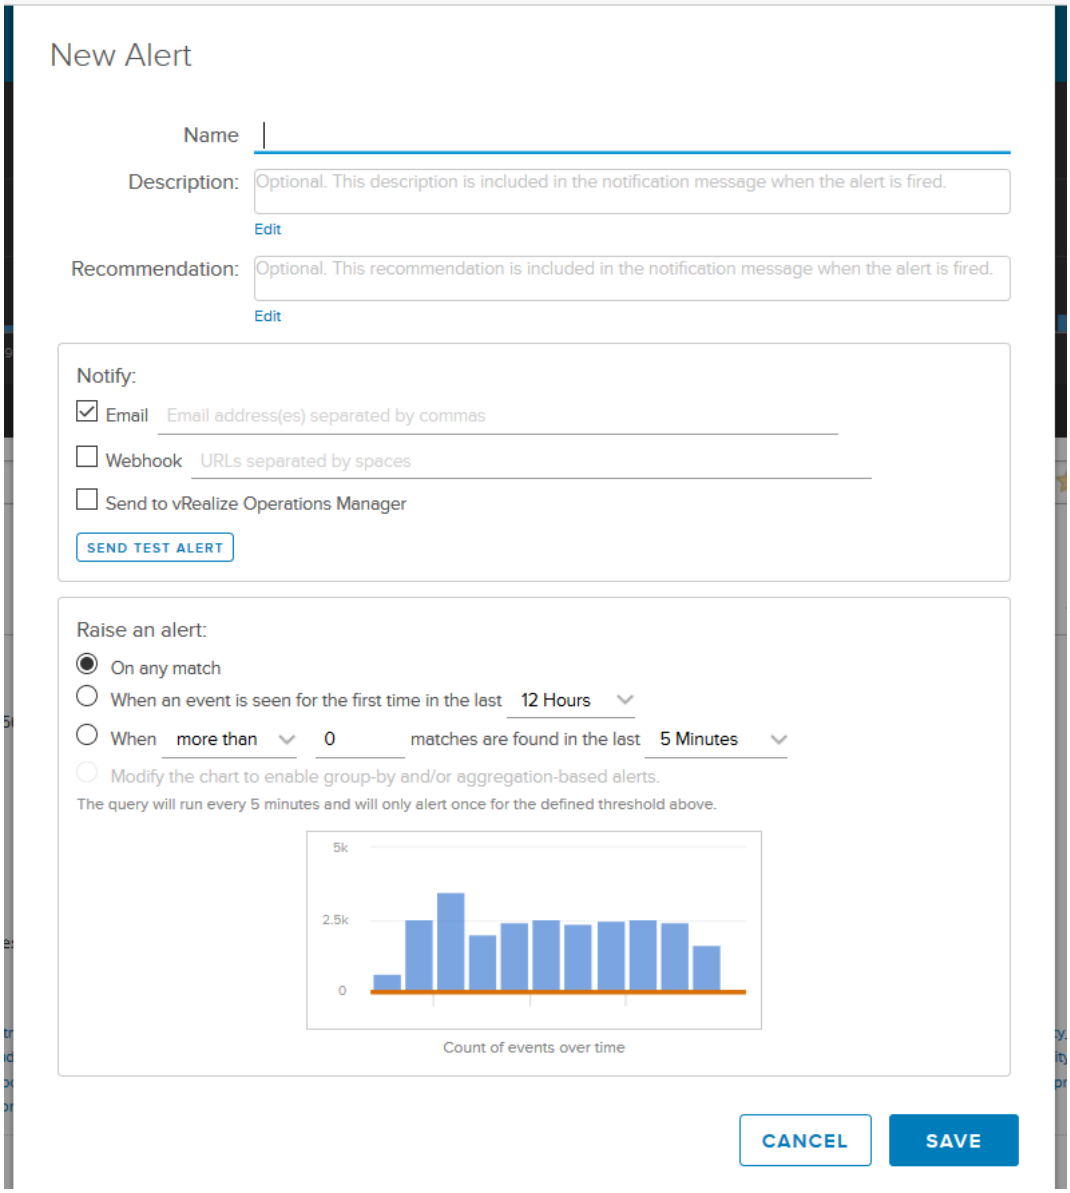

Figure 27. An example of an alert. The threshold has been set to trigger when a type of vCenter Server event for a hostname is seen in the last hour. The query runs every 10 minutes and if the alert triggers, it will not run again for one hour.

Dashboards

A content pack comprises one or more dashboard pages known as dashboard groups.

Dashboard Groups

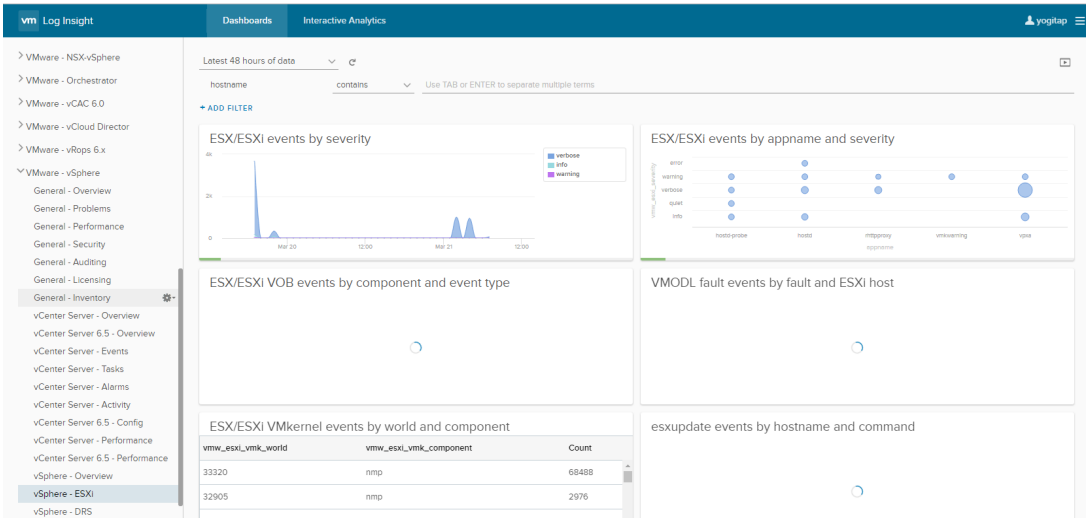

Figure 28. The vSphere content pack. In the left navigation bar, below the name of the content pack, are the dashboard groups.

Dashboard Group – Best Practices

When creating dashboard groups, the following best practices apply:

- Content packs commonly contain a minimum of three dashboard groups. The best practice is to start with an overview dashboard group to provide high-level information about the events for a specific product or application. In addition to the overview dashboard group, dashboard groups should be created on the basis of logical groupings of events. The logical groupings are product-specific, component-specific or application-specific, but some common approaches are performance, faults, and auditing. It is also common to create dashboard groups per component, such as disk and controller. With the component approach, it is important to note that it is only effective if queries can be constructed to return results from specific components. If this is not possible, the logical approach is recommended.
- When naming dashboard groups, make the title generic and avoid adding product-specific or application-specific names, unless they are being used in a component specific fashion. For example, in the vSphere content pack, the dashboard groups are called *ESX/ESXi hosts* and *SCSi/iSCSI and NFS* instead of *VMware ESX/ESXi hosts* and *VMware SCSI/iSCSI and NFS*.
- Dashboard widget names should start with Capital letter but rest of the letters should be lower case, not all CAPS.
- A dashboard group should contain a minimum of three, and a maximum of eight, dashboard widgets. With fewer than three dashboard widgets, the volume of knowledge that can be attained by the dashboard group is minimal. In addition, having many dashboard groups with only a limited number of dashboard widgets requires a user to switch between pages and does not provide information in a coherent way. Conversely, greater than eight dashboard widgets per dashboard group can result in the following:
- Too much information: A user might not know where to begin, or what is most important.
- Resource intensive: Each widget is a query that must be run against the system.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

- When nearing or exceeding eight dashboard widgets in a dashboard group, separate information and create multiple dashboard groups. If a dashboard widget is applicable to one or more dashboard groups, it is recommended to create the widget in each applicable dashboard group.

Dashboard Widgets

There are two different types of dashboard widgets in VRealize Log Insight:

- Chart: contains a visual representation of events with a link to a saved query.
- Query: contains title links to saved queries.

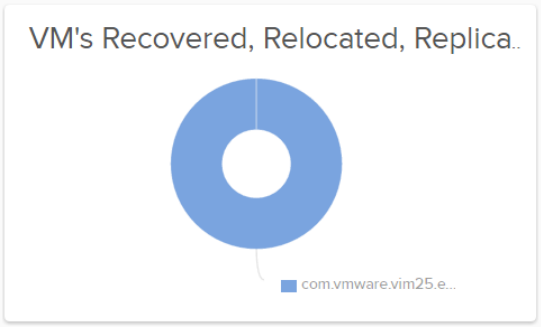

Figure 29. An example of a chart widget.

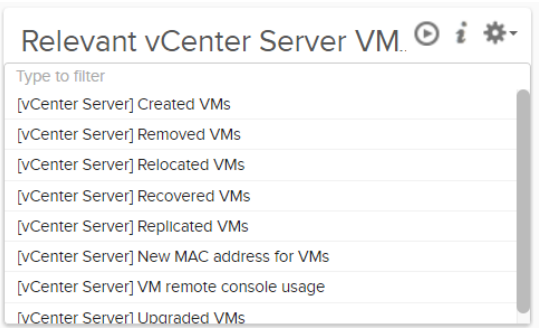

Figure 30. An example of a query widget.

Chart

A dashboard chart widget contains a visual representation of events. A chart can either be represented as a bar or line chart (or bubble, pie & area chart) and can be displayed in a stacked fashion. The following best practices apply:

Charts can contain a lot of information so avoid having more than two chart widgets per row. In some rare cases, three chart widgets can be used effectively, but more than three is strongly discouraged. When determining whether chart widgets are readable or not, use the minimum resolution supported by vRealize Log Insight (1280 x 800) because it cannot be assumed that users of the product will have a higher resolution.

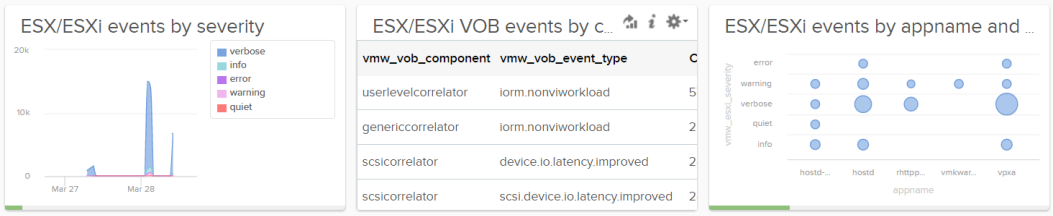

Figure 31. An example of three chart widgets in the same row with additional content, these widgets may become hard to read..

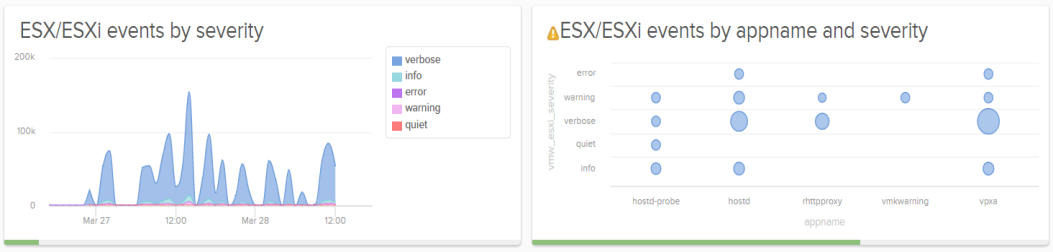

Figure 32. An example of two chart widgets per row. Event with additional content, these charts should be readable.

If any row, other than the last row, has a single chart widget, it is recommended to make that widget full width.

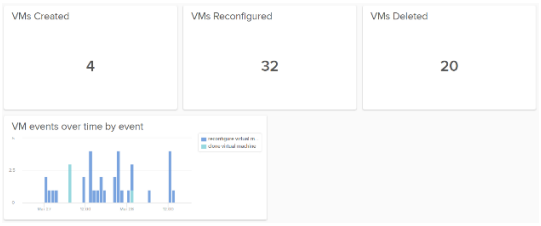

Figure 33. An example with a half-width chart on the top row. This is not recommended.

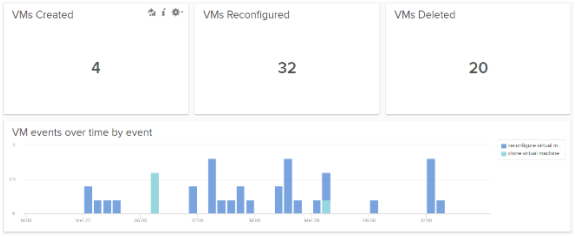

Figure 34. An example of a full-width chart on the top row. This is recommended.

When naming a chart widget, use a descriptive title and avoid cryptic field names. For example, an extracted field is called *vmw_error_message*. Instead of calling a chart *Count of vmw_error_message*, call it *Count of error messages.*

Multi-function charts can be used when you need to compare the same set of logs using different aggregation functions. For example, Max of events over time grouped by source & Avg of events over time grouped by source OR you want to compare Count of events over time and Unique count of hostname over time.

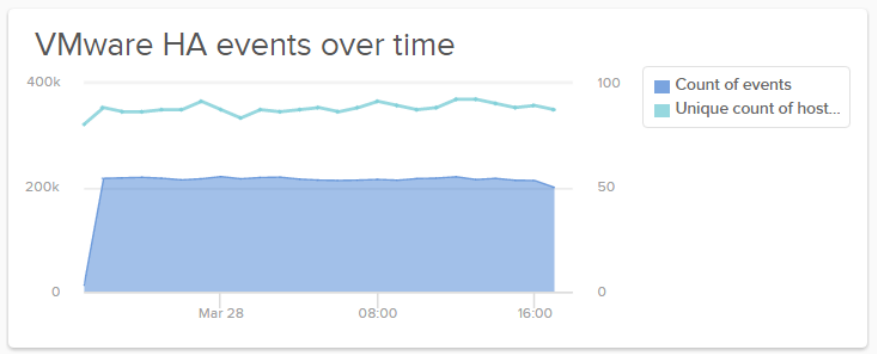

Similar charts can be saved and stacked in the same column of a dashboard group for visual comparison. Examples of such charts include:

Average X of events over time + Maximum X of events over time. Given the different functions used, it is possible that the Y-axis of the charts will not be the same scale.

Count of events over time grouped by X + Count of events over time grouped by Y.

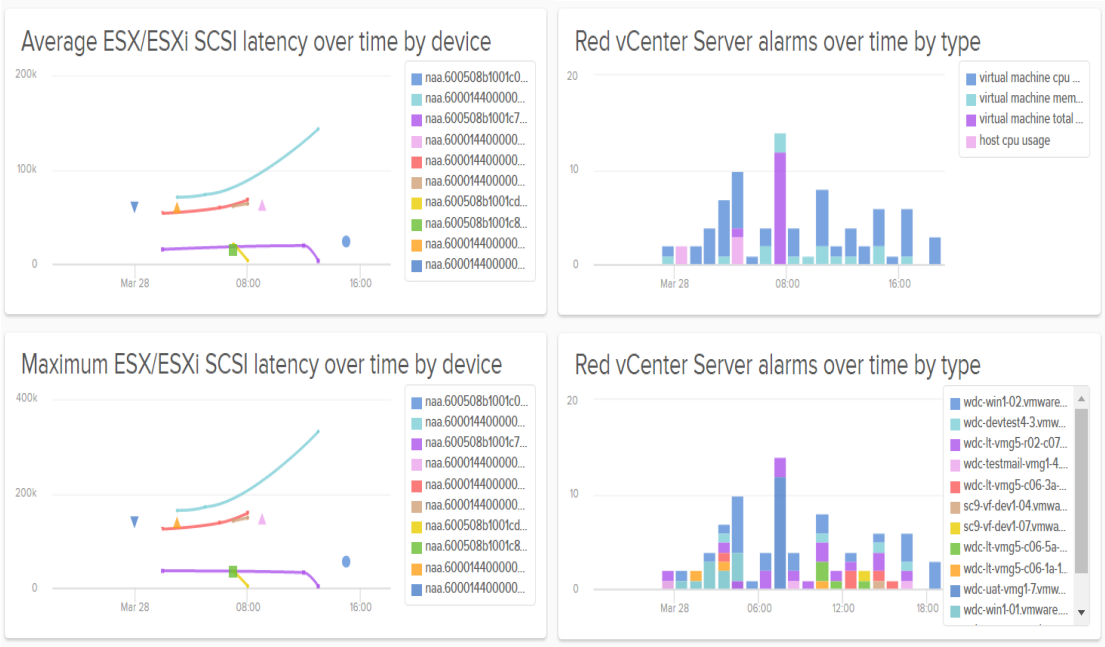

Figure 35. An example of two similar charts using different functions stacked. Notice that the scale of the charts does not match.

Figure 36. An example of two similar charts using different groupings stacked. For this type of query, the scale of the charts match.

Query

A dashboard query widget contains a title that is a link to a pre-defined query. Query widgets are often used when a chart widget does not provide a significant value, but a query does.

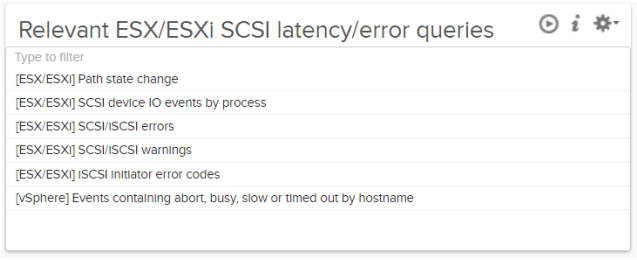

Figure 37. An example of a query widget.

Widgets

You can rename, resize, clone, and edit widgets.

Rename Widgets

To rename a widget, select the name of the widget. When naming a chart widget, use a descriptive title and avoid cryptic field names.

Resize Widgets

To resize a widget, hover over the right edge of a widget's contents.

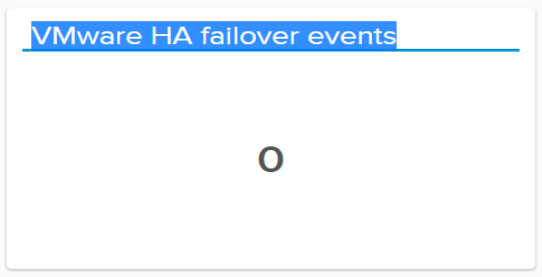

Figure 38. Renaming a widget.

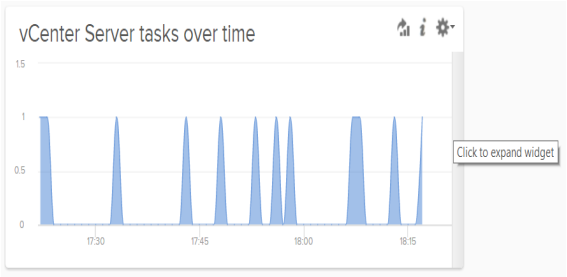

Figure 39. Resizing a widget.

Move Widgets

Within a dashboard group: To move a widget within a dashboard group, select between the title and the action buttons and drag to the new location. Important: It is not possible to create a new row between two existing rows.

Figure 40. Moving a widget within a dashboard group.

Figure 41. Attempting to add a new row in a dashboard group.

Instead, move the widget to the left-most position of the row below the row desired and move all widgets that follow the new widget down.

Figure 42. Moving a widget to create a new row in a dashboard group.

– Between dashboard groups: To move a widget between dashboard groups, click the gear action button and select *Move to Dashboard*.

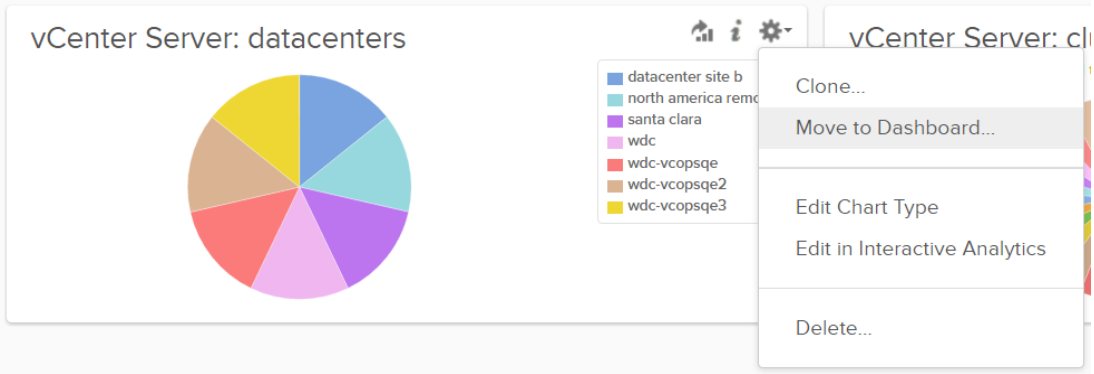

Figure 43. Moving a widget to a new dashboard group.

Clone Widgets

To clone a widget, click the gear action button and select *Clone*.

Important: When cloning a chart widget, any fields that the chart relies on are not cloned. Instead, cloned chart widget fields are defined by the cloned source. For this reason, cloned widgets should not be used in content packs because they might cause content packs to be dependent on other content packs.

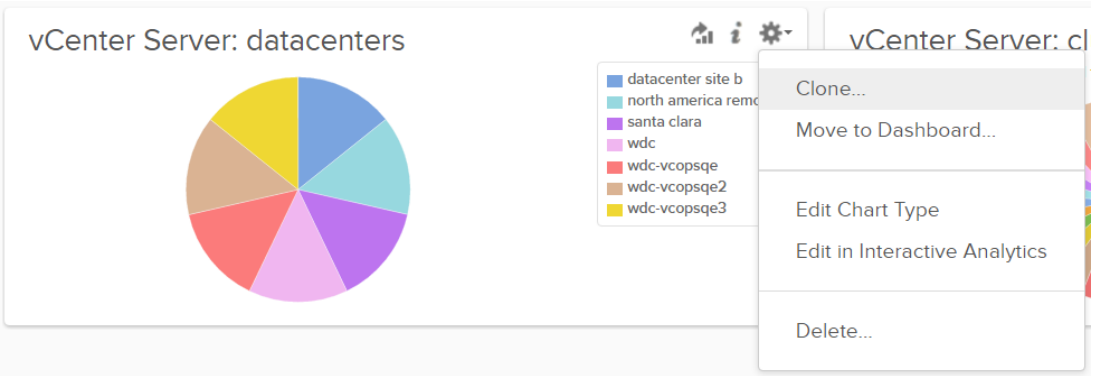

Figure 44. Cloning a widget.

Edit Widgets

To edit the notes section of a widget, click the ***i*** button and select ***Edit***.

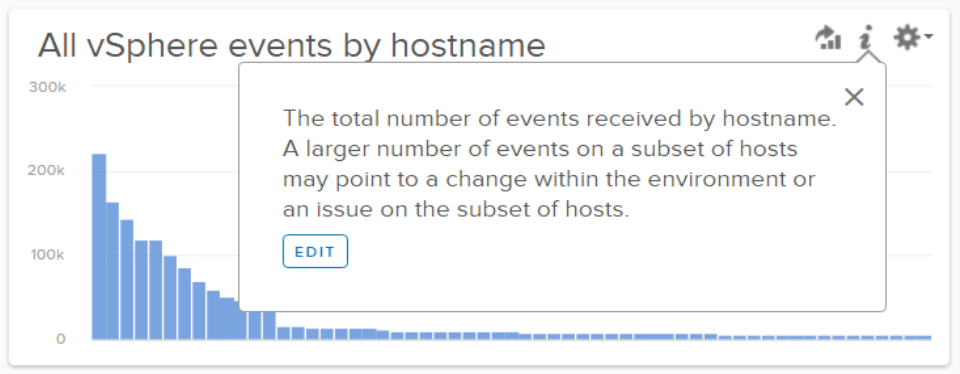

Figure 45. Editing the information field of a widget.

The notes section is very important and should be populated for every dashboard widget. Information that can be added can be text, a link to documentation, a knowledge base article, or a forum. Information provided should answer the following questions:

- Why is this widget important?
- What is a “good” and a “bad” value?
- Where can more information be obtained?

Edit Chart Type:

To change the chart type of a chart widget, click the gear icon of the widget and change the chart type or properties of the current chart type.

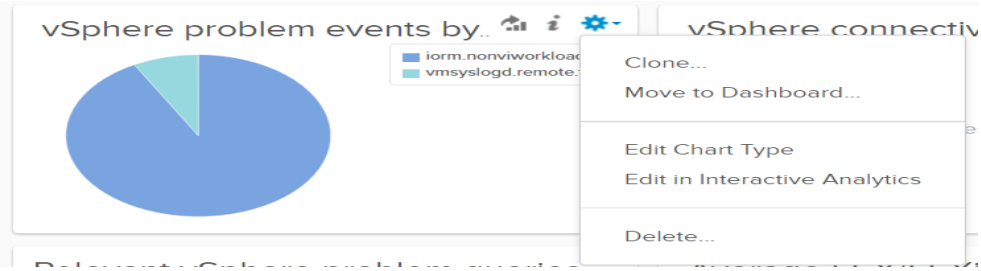

Edit Chart in Interactive Analytics: To change the underlying query of a widget, click the gear icon of the widget and select the menu to modify the widget’s query in interactive analytics.

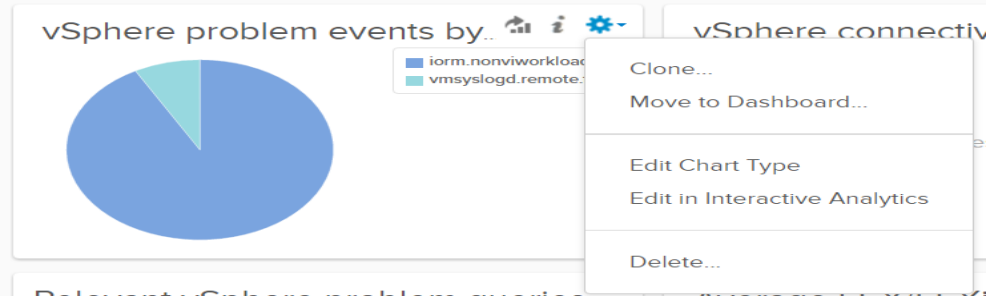

The underlying query for a widget can be modified. In order to change the underlying query, use the Edit Chart in Interactive Analytics menu. For chart widgets, the directions are:

1. Go to the widget on the Dashboards page.
2. Select Edit in Interactive Analytics within the widget from the gear menu.
3. Modify the query as required.
4. Click the Save button on the Interactive Analytics page.
5. Click the Return to Dashboard menu from the bottom of the Interactive Analytics page.

Agents, Agent Groups and Agent Configuration

A Log Insight agent supports sending syslog to third party destinations and using Log Insight's ingestion API to any remote destination. For systems that do not support the Log Insight agent you can also use agents like the rsyslog or syslog-ng agent to send events to vRealize Log Insight, for more information refer the **VMware vRealize Log Insight Documentation Center.** The Log Insight agent offers many benefits over third party syslog agent including the ability to keep track of where it left off to ensure the most recent logs are collected since the rotation as well as all rotated files, hence using the Log Insight agent is recommended.

The vRealize Log Insight server allows users with administration rights to configure agents from within the application UI. Agent configuration is picked up by each agent configured to send logs to the vRealize Log Insight server.

However, it may be necessary to push different configurations to different agents. This is where Agent Groups help the user. For example: To collect logs from an IIS server the IIS content pack uses logs in W3C format. For information regarding how to collect logs from an IIS server, refer to the setup instructions in the IIS content pack.

If IIS logs are stored in a log file directory, for example: C:\inetpub\logs\LogFiles\W3SVC1\ and your agent configuration looks similar to the following:

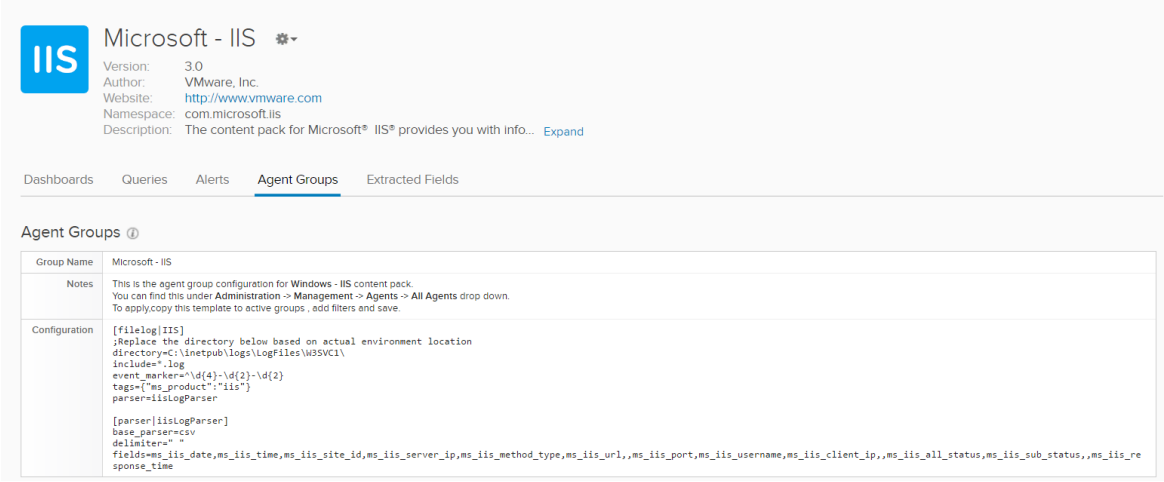

Users can actually eliminate the need for extracted fields when the logs are as structured as they are in the case of IIS.

To create an agent group:

1. From the Administration \ Agents UI, Select the drop down that says 'All Agents'
2. Select the option 'New Group'
3. Give it a name and enter notes about the alert including what it means, how to resolve it, and where to get more information.
4. Select 'New Group'. (Note: Your new group is created at this point but not saved)

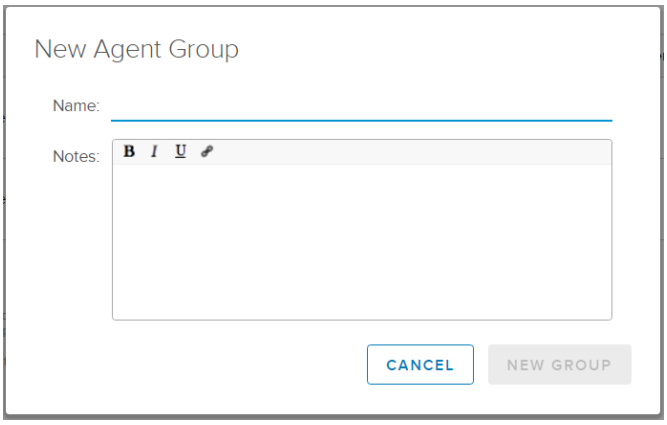

A filter is added to the UI, where you should select at least one filter.

Available filters are for IP address, hostname, version and OS.

For example, Select, OS, filter ‘starts with’ and value as ‘Microsoft’, to configure all your Microsoft windows agents. Note: all agents are listed until a filter is selected and you click Refresh.

Your Microsoft windows agents if any will be listed below, when the filter is selected.

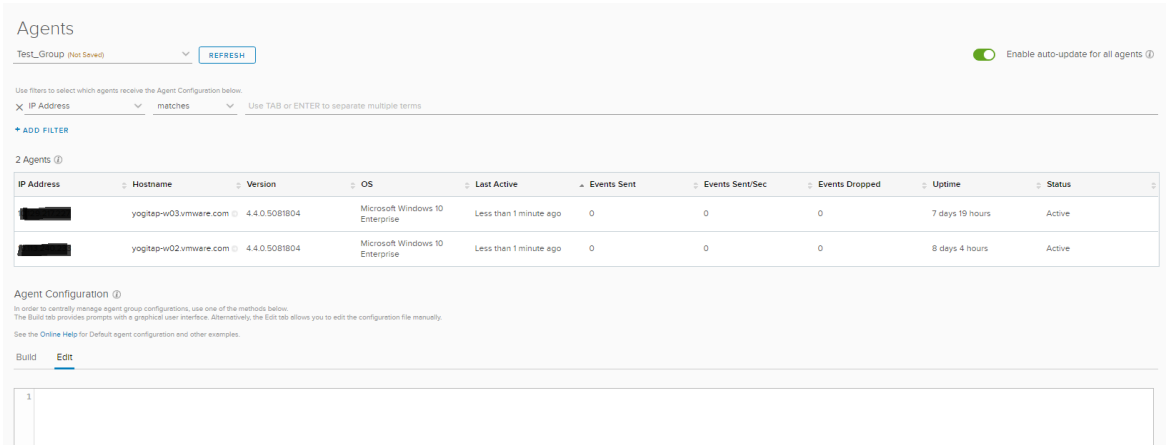

5. Type your agent configuration for your windows agents in the Agent Configuration box below.
6. Select button ‘Save New Group’.

The agent configuration will be picked up only the windows agents. If any windows agents are installed after the agent configuration is set, then the windows agents will automatically pick up this agent configuration.

After the agent configuration is pushed out to the agent, the `liagent-effective.ini` looks similar to:

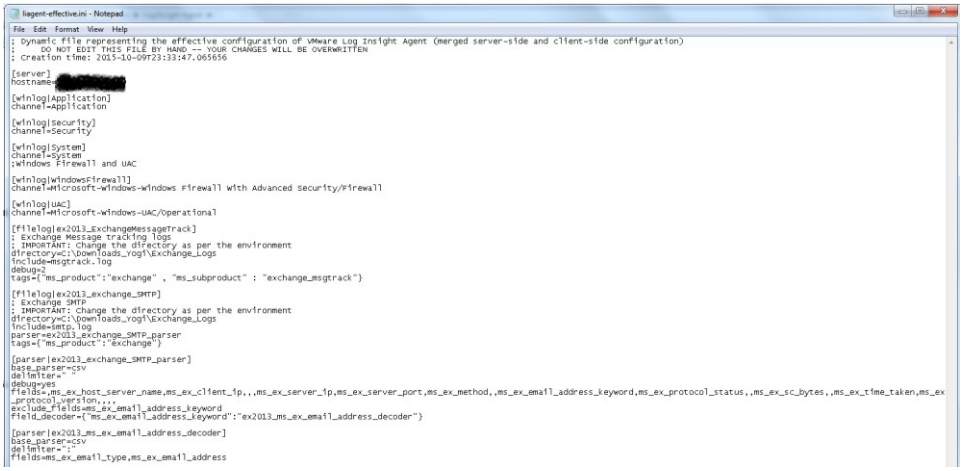

Note: You will see some agent groups listed automatically, these come from the installed content packs (for example, the vSphere content pack) or other installed content packs that are using agent groups, as in the picture above.

- Content pack agent groups are read-only groups.
- Agent group filter values are not exported with the content pack.

To apply an agent group from a content pack, 1.

Copy the required template to active groups,

2. Add filters and save.

If you have an existing filelog or winlog section with the same name as a section in the agent group, the content pack agent group will not merge with the client side configuration as the content pack agent group has a unique namespace associated with it.

Setup instructions for the content pack should give details about how to use agent groups configuration when using the content pack.

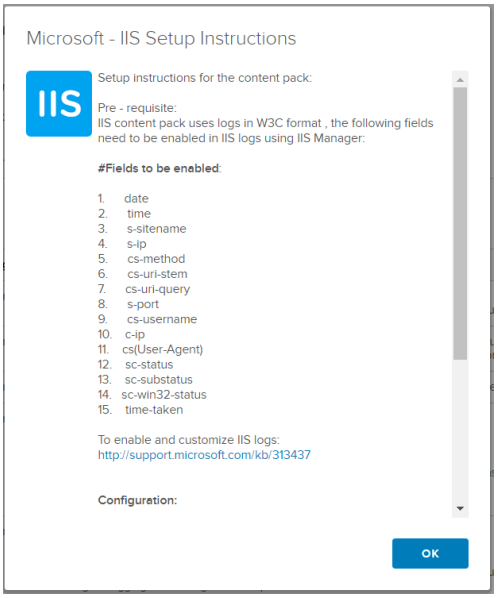

For information regarding editing and deleting agent groups refer to VRealize Log Insight documentation. For detailed parser information about setting up your agent configuration as per the structure of your logs, refer to the VMware vRealize Log Insight Documentation Center.

Suggested Template for Setup Instructions:

The <name of your content pack> content pack can forward logs to Log Insight in different ways and requires one or more of the below configurations.

Option 1:

Log Insight Agent Configuration: (preferred method)

The <name of your content pack> content pack requires the use of the Log Insight agent with the cfapi protocol (default) and the included agent group configuration. To apply the agent group configuration:

- Go to the *Administration* > *Agents* page (requires super admin privileges)
- Select the *All Agents* drop-down at the top of the window and select the *Copy Template* button to the right of the <*name of your agent group*> agent group (use the version as per your environment when there are multiple agent groups)
- Add the desired filters to restrict which agent(s) receive the configuration
- Select the *Refresh* button at the top of the page
- Check the agent group configuration for comments requiring user intervention such as names of directory locations and adjust as necessary
- Select the *Save Configuration* button at the bottom of the page.

Option 2:

Configuring <your product>:

- To set up event forwarding from <your device>, first log in to the device.
- In the 'Remote Syslog Server' field, enter the IP or FQDN of the server that logs will be forwarded to. This server may be your **IP address of the VMware Log Insight integrated load balancer** or Log Insight server, or a different server acting as an intermediary between your device and the Log Insight Server such as the Log Insight agent.

e.g. Navigate to Administration -> Management -> Remote Syslog of your product and enter the **FQDN or virtual IP address of the VMware Log Insight integrated load balancer to the list of syslog destination targets**.

Option 3:

Steps:

Beginning in vRealize Log Insight 3.3, the integrated load balancer now allows multiple VIPs to be configured with tags.

- Log into VMware Log Insight server as a **super admin user**. Go to Administration -> Cluster to add a new Virtual IP Address (VIP) e.g. *10.20.30.40* with the tag ***product=<your product>***
- On each device that is intended to be monitored by VMware Log Insight, add the FQDN or virtual IP address of the VMware Log Insight integrated load balancer to the list of syslog destination targets. For instance: *10.20.30.40*

NOTE: If it is not possible to add keywords to an extracted field and it is not possible to add an existing static field that uniquely identifies your product logs to an extracted field then the multiple VIP + tags feature must be used.

Content Packs

With an understanding of what comprises a content pack and the best practices when performing each operation, it is now time to view, export, import, edit, and publish the content.

View

To view saved content:

Navigate to the *Content Packs* section by clicking the gear icon in the navigation bar and selecting Content Packs.

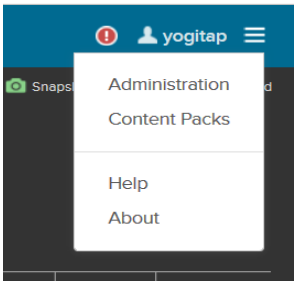

Figure 46. Content Packs menu option. Important: The Administration option will only be visible to Admin users.

Select where the content was saved. For content pack authors, content is saved under *Custom Content* and, if following the best practices described in the Getting Started section of this document, saved content will appear under *My Content*.

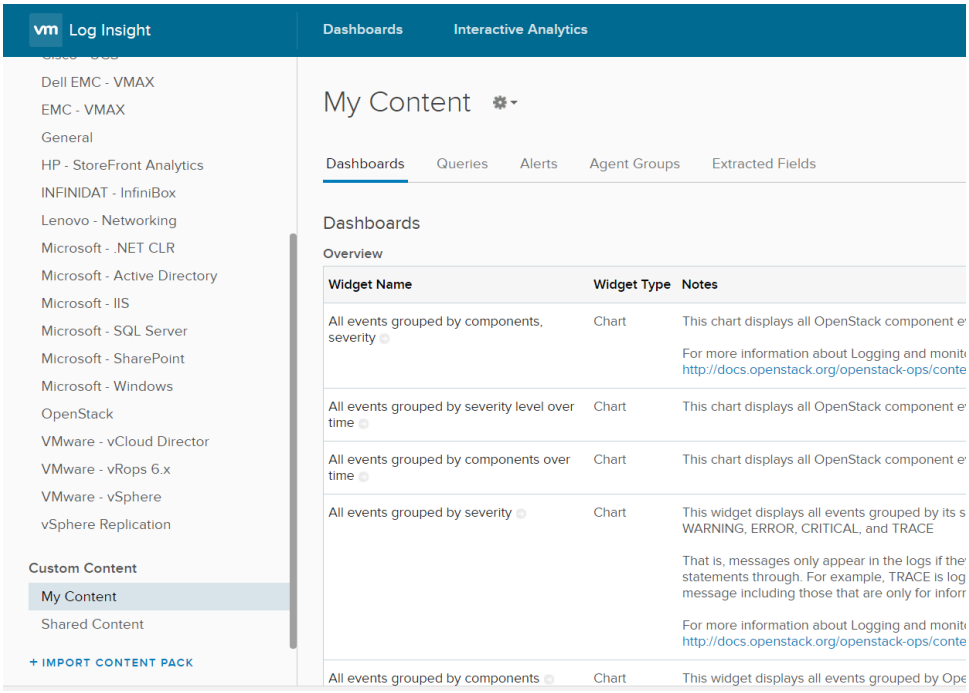

Figure 47. An example of saved content.

In general, a content pack should have:

- Three or more dashboards (dashboard groups)
- Three or more queries (chart/table widgets) per dashboard (nine or more in total)
- Five or more alerts
- Twenty or more extracted fields (if using agent groups together with parsers and tags, you can have a very effective content pack with very few or no extracted fields, especially in case of structured logs)

Export

Private Export

To export all information saved in a dashboard for private use:

1. Select the content, then click the gear icon to the right of the dashboard name.
2. Select Export.
3. Give the content pack a name. The recommended format is: <Company> – <Product> v<Version> (For example, VMware – vSphere v1.0). Ideally, the content pack name should be less than 30 characters to prevent word wrapping.
4. (Optional) Give the content pack a namespace of the format com.<company name>.<product name>. After a content pack is published DO NOT change the namespace because a user who is upgrading content packs will get a new copy of the content pack instead of an in-place upgrade.
5. Give the content pack a version number in the format MAJOR.MINOR.REVISION

- MAJOR - many changes to the content pack, for example one or more new dashboards
- MINOR - fixed a bug, changed a widget type, maybe added one or two widgets
- REVISION - for content pack authors, when preparing a new version to send to VMware with the revision set (starting from 1). Every time feedback is given and another revision needs to be validated,

the revision number is incremented. When the content pack is published officially, it does NOT include a revision number.

6. (Optional) Give the content pack an author name. (The company publishing the content pack, for example, VMware Inc.)
7. (Optional) Give the content pack a Website URL preferably the product website for which the content pack is for.
8. (Optional but recommended) Give the content pack a description that provides details about what the content pack is about and what kind of information the widgets display. We would like for you to include a link to your End User License Agreement so that our users are aware that the content is not distributable. Please include the link to your EULA at the end of your description. For example: “Please review our End User License Agreement before installation.”
9. (Optional but recommended for all content packs) Give the content pack Setup Instructions that provides details about what the user needs to do to effectively use the content pack. Setup Instructions should clearly outline the exact steps required in vRealize Log Insight and the product for which your content pack is. For example:
10. (Optional but recommended for versions greater than 1.0) Give the content pack Upgrade Instructions that provides details about any additional instructions the content pack user needs to be able to use all the features in the upgraded version of the content pack. For e.g any additional setup required for newer features to work or changes to earlier setup done to send logs to vRealize Log Insight. If there are no additional instructions required then its good practice to put a sentence that says something like. “There are no additional instructions required with this upgrade. This upgrade includes changes to:

    Alert ABC definition has changed, field definition of field vmw_pqr has changed, Dashboard Widget

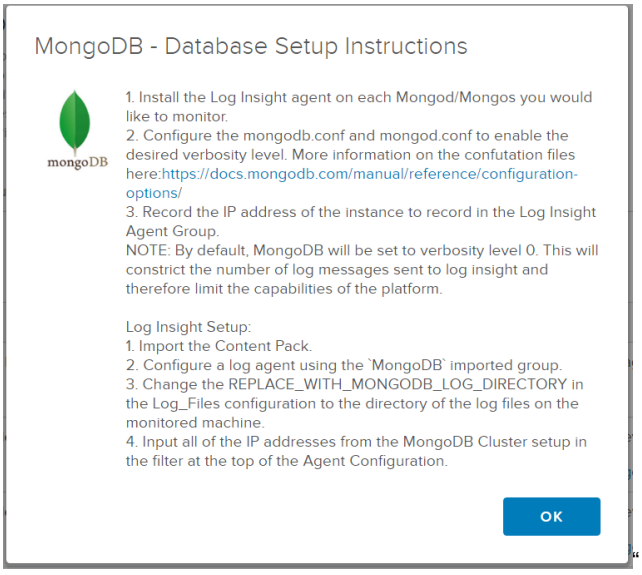

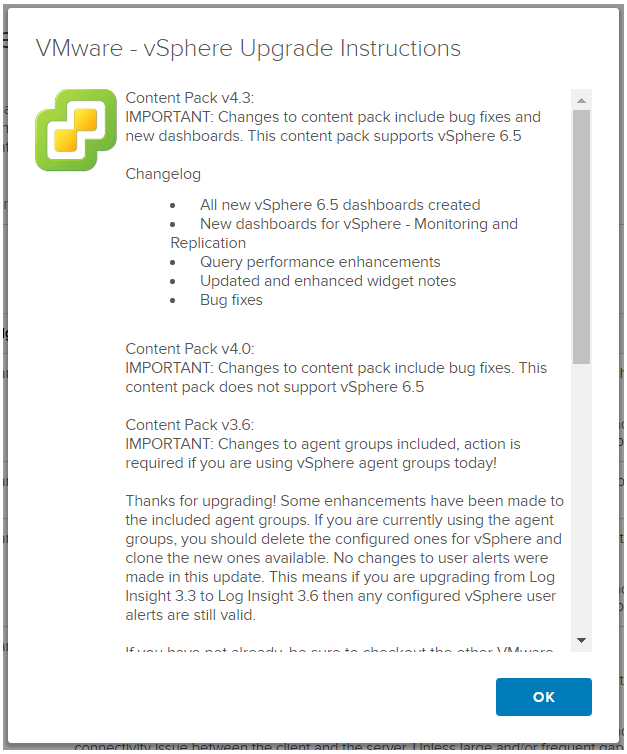

11. (Optional) Give the content pack an icon of size 144 x 144, with a PNG or JPG file type if possible, to help identify your content pack in the marketplace. If using icons that are governed by copyright laws, please ensure you have permission in writing to use the icon before you publish the content pack.
12. (Optional but Recommended) Give the content pack; a list of detailed setup instructions for the end user to be able to setup and use the content pack once published.
13. Select Export.

Once complete, a file ending with a VLCP extension, which stands for vRealize Log Insight Content Pack is downloaded.

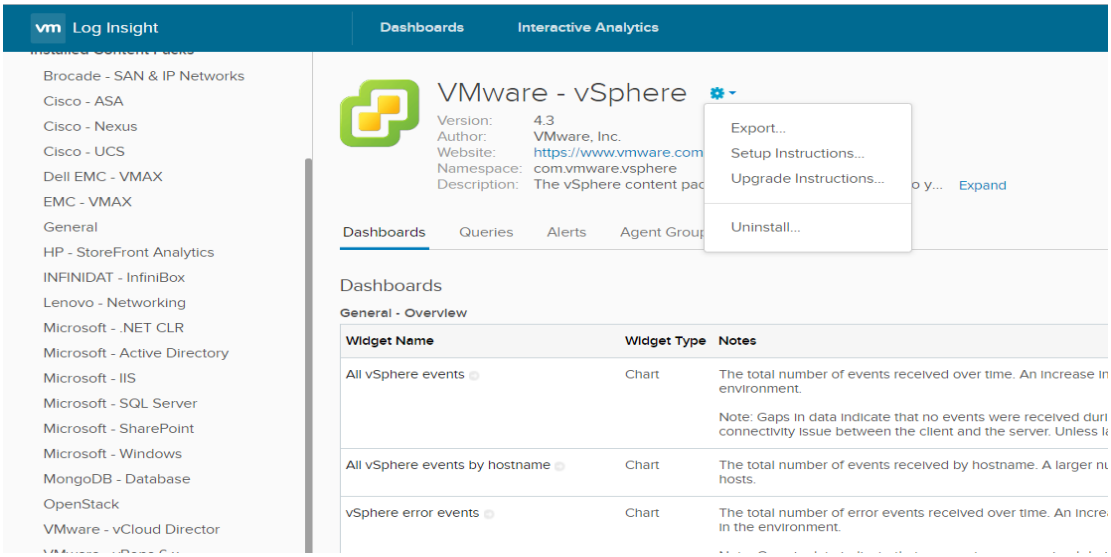

Figure 48. How to export a content pack.

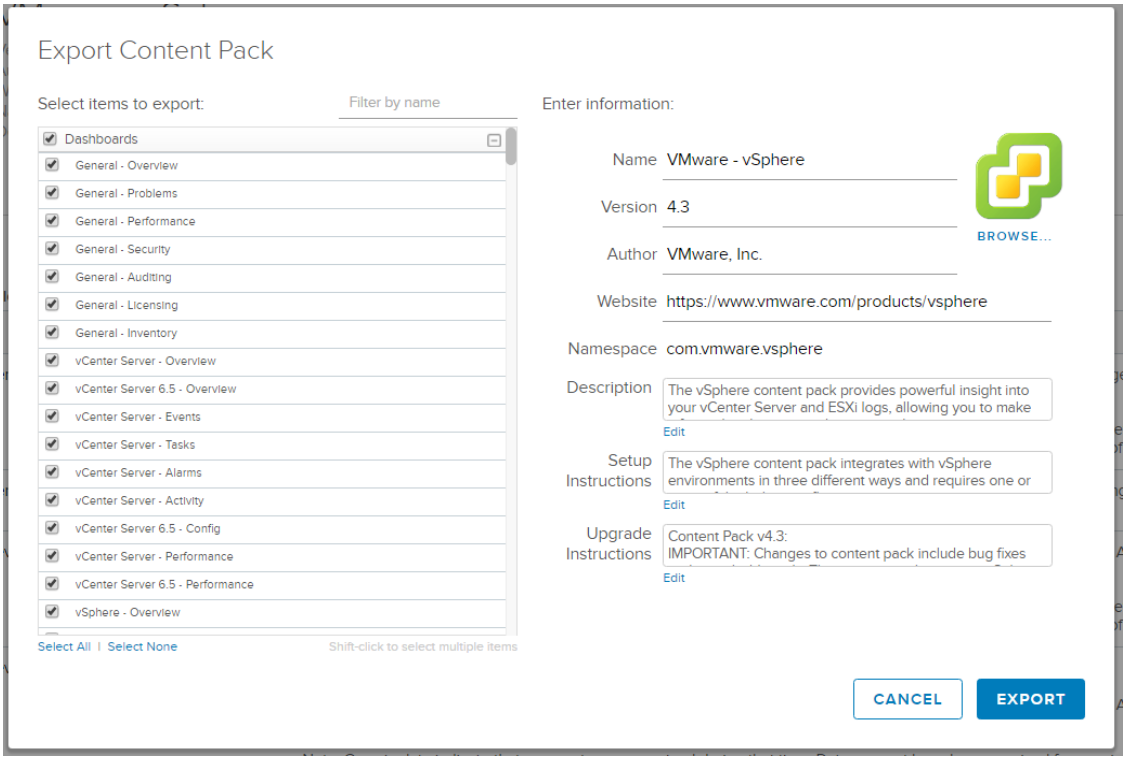

Figure 49. Export content pack dialog box.

Marketplace Import

To import a content pack from the in-product marketplace:

1. Click the *Marketplace tab in the Content Pack UI* link at the top of the left navigation bar.

2. Click *the content pack you are interested in*… select the check box for - By clicking "Install", I agree to the terms of any License Agreement included above.
3. Click *Install*
4. Content pack is installed to Installed Content Packs.

Import a content pack

To import a content pack:

1. Click the *Import Content Pack* link at the bottom of the left navigation bar.
2. Click *Browse*… to specify the location of the VLCP file.
3. Click *Import*.

You can install the imported content pack as an Installed Content Pack, where you can use the content pack but cannot edit the content pack.

OR, You can import the content pack into user content (My Content), where you can edit the content pack and queries, alerts and fields in the content pack.

Figure 50. Import button.

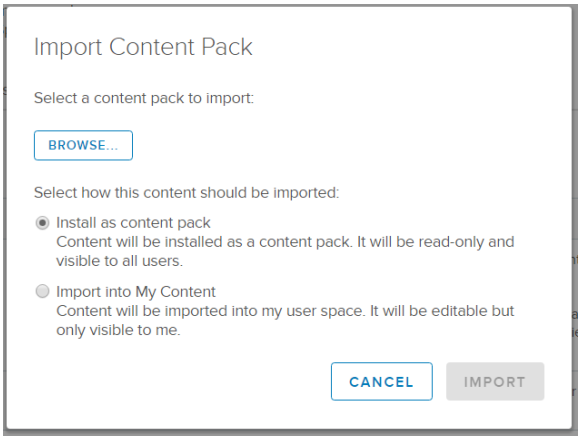

Figure 51. Importing a content pack dialog box.

When importing a content pack, warning and error events can occur. These include:

Duplicate Name: A Duplicate Name means that another content pack is installed in the system that has the same unique identifier. In this case, the options are to either choose *Overwrite* to replace the existing content pack or *Cancel* to keep the existing content pack.

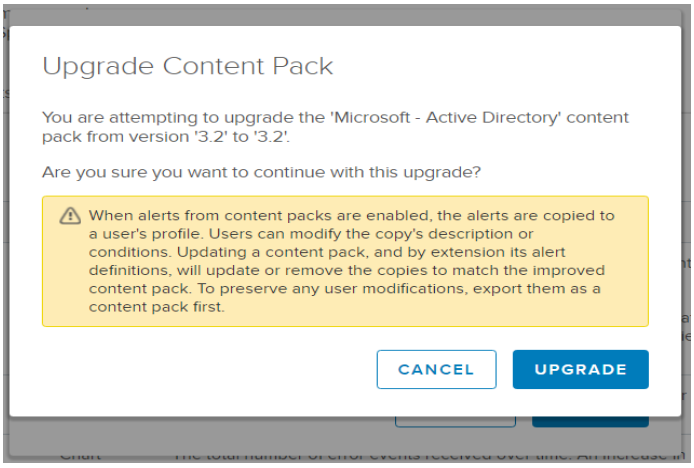

Figure 52. Duplicate content pack warning dialog box.

Invalid Format: Invalid Format means that the VLCP file was manually edited and contains syntax errors. The syntax errors must be fixed before the content pack can be imported. As VLCP files should not be manually edited, there is no easy way to locate and fix syntax errors.

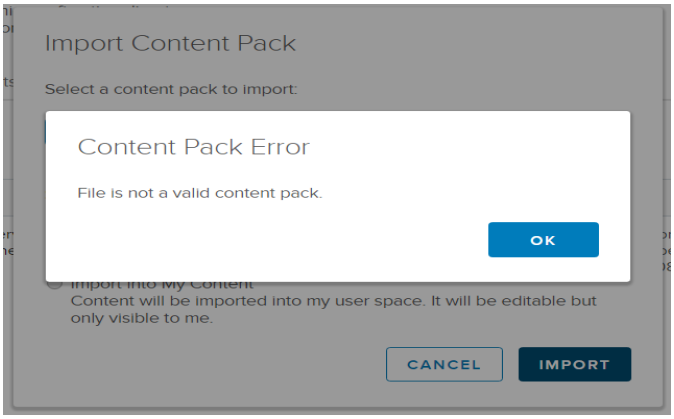

Figure 53. Content pack error dialog box

Edit

As imported content packs that are imported to Installed content packs are read-only, content packs should be edited from the instance of vRealize Log Insight on which they were created. It is possible to import content packs into user content, also known as user space, and to modify its contents. Take care that the user space does not contain any widgets, alerts, fields or queries from another content pack or the user will receive mixed content, resulting in confusion. The recommendation is to modify content packs on the instance of vRealize Log Insight used to create the original content pack or create a new user and import the content pack into the user space of this user to modify the content pack. The original vRealize Log Insight instance should be properly backed up.

Testing Your Content Pack

- Be sure to test your content pack on a Log Insight instance that has a very large number of (10+ million preferably) events ingested.
- Test the performance of your queries, dashboard, alerts for content correctness and time to complete the queries.
- If you followed the best practices described in this document the performance of the content pack should be good, however if you observe the performance in terms of time and correctness to be inefficient revisit the extracted field and query definitions to ensure you have followed the best practices and make necessary adjustments.

Publish

After a content pack has been created the content pack has to be reviewed and approved by the vRealize Log Insight CORD team via DCPN, post approval it can be published to the in-product vRealize Log Insight marketplace and on the VMware Solution Exchange. The requirements for content pack publishing are as follows:

- Content pack: A VLCP file.
- Events: Appropriate events that are necessary to validate content pack.
- Documentation: Information about how to configure the product/application to forward logs to vRealize Log Insight. Some release notes and upgrade instructions if it is an update to a published content pack. For more information, see the Resources section below.
- (Optional) Demo/Story: Example of how the content pack brings value (for example, YouTube video).

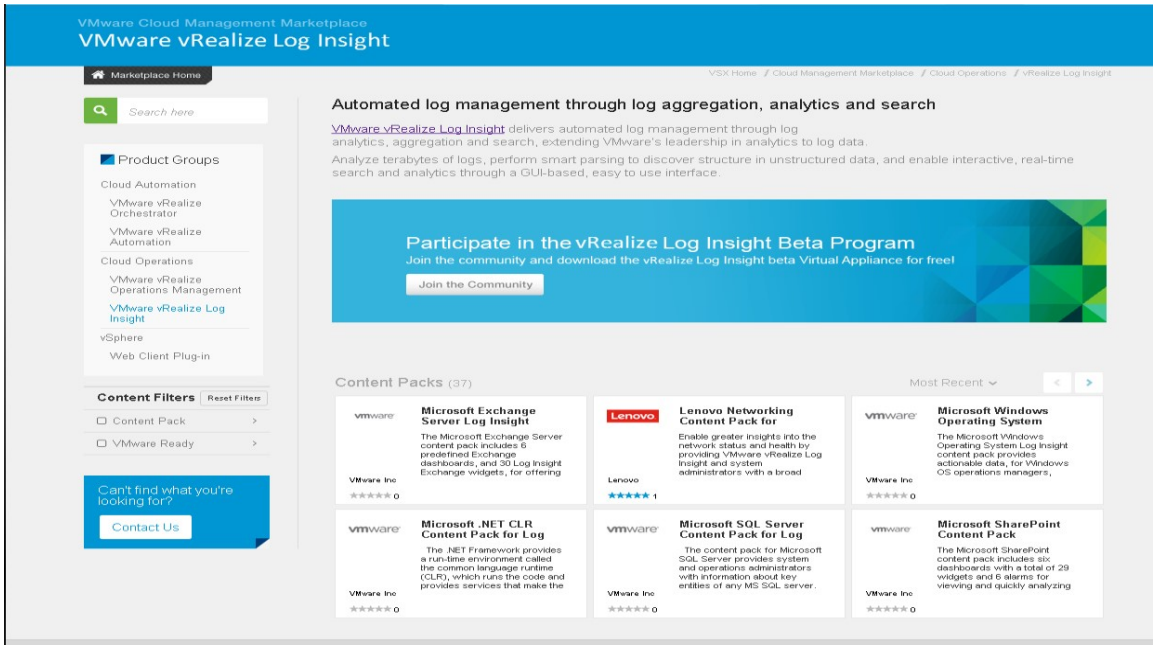

Figure 55. VMware Solution Exchange.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

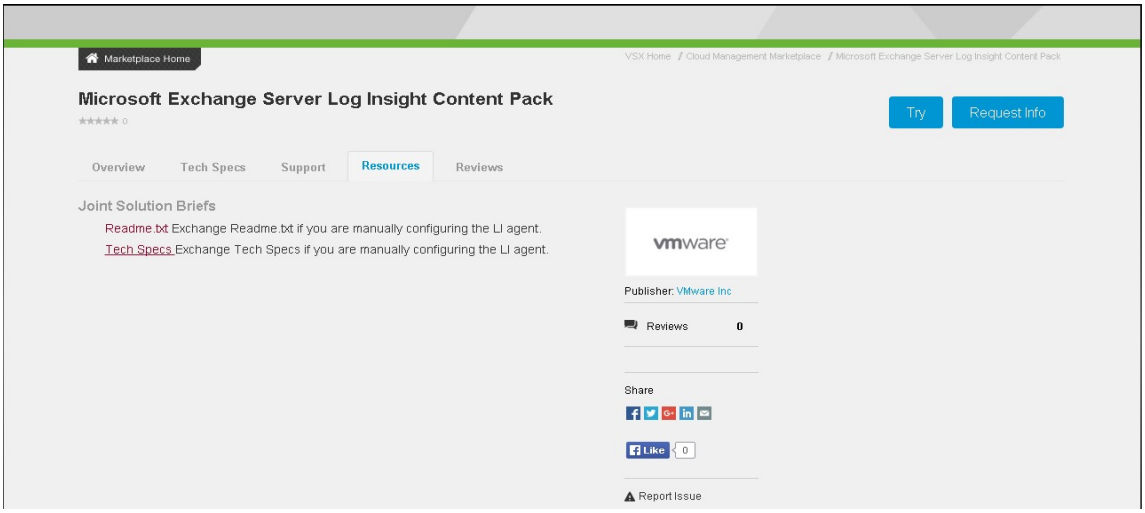

Figure 56. Additional resources for the content pack.

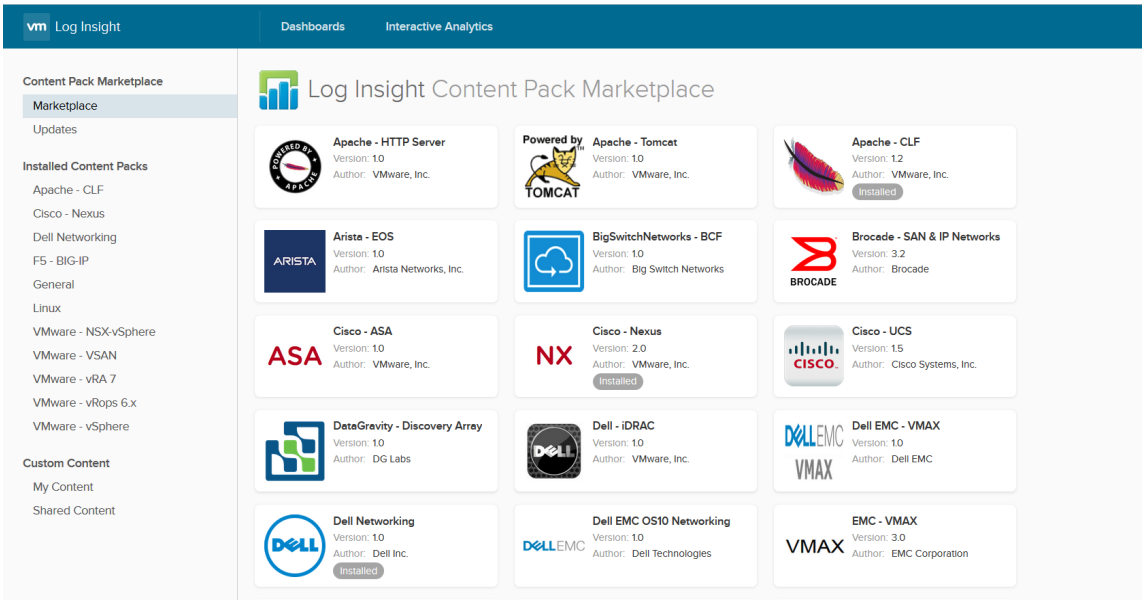

Figure 57. vRealize Log Insight In-product Marketplace.

Conclusions

Content packs are a powerful way to extend the knowledge contained within vRealize Log Insight. When creating a content pack, several best practices must be considered as outlined below.

Getting Started

Instance

- The instance of vRealize Log Insight used to create a content pack must be backed up. If the instance used to create a content pack becomes unusable, the content pack must be recreated on a different instance so that it can be modified.
- Do not attempt to edit a content pack from an instance of vRealize Log Insight other than the one that created the content pack, unless the intention is to recreate the content pack.
- Beginning in vRealize Log Insight 3.3, the integrated load balancer now allows multiple VIPs to be configured with tags. This makes it possible to tag ingested log messages for devices that cannot leverage the Log Insight agent and offers a query performance boost for content packs with limited keywords.

User

- Use a separate content pack author user on vRealize Log Insight for each content pack that is created.

Queries

Message Queries

- Use keyword queries whenever possible.
- If keyword queries are not sufficient, use globs.
- Use regular expressions only if keywords and globs are not sufficient. When using regular expressions, provide as many keywords as possible.
- Make queries as specific as possible. Content pack queries should only match events applicable to the product/application for which the content pack was designed.
- If you only need the exists operator on fields then you do not need to extract the field -- and extracting the field in this case actually causes performance issues.

Field Extraction

- Minimize the number of regular expressions that are used, whenever possible.
- Verify that a regular expression value will match every applicable log message.
- Provide as much pre-keyword and/or post-keyword context as possible. *Filters:*
- When using filters, do not use the match “any” operator unless one or more keywords are defined in the search bar.
- When using the text filter with multiple different values, one or more keywords should be defined in the search bar.

- Understanding of what "any" means vs "all": "any" means that each filter is a SEPARATE query -- so when multiple filters are used with ‘any’ operator it is actually multiple queries. In general more the queries, the slower the results. Think of "any" as "or" and "all" as "and" operators.
- Matching AQ (aggregated query) to MQ (message query) is not required for reasons mentioned above when the match “any” operator is used.

When naming a field:

- Use the following naming standard: <prefix>_<field>_<name> Use underscores, not spaces.
- Use all lowercase letters.
- <prefix> = something applicable to the content pack.
- Use keywords in additional context of field to help performance of field in queries.
- Use additional context filters on fields if possible to help field performance in queries.
- Use a VIP tag on fields as additional context for logs with limited keywords.
- Test to validate that an extracted field is working as expected.

Aggregation Queries

- When grouping by time series, do not add more than one field.
- Do not group by time series and one field if the number of unique fields is or could be more than 20.
- When grouping by more than one field and time series, ensure that the time series adds value.
- If the time series is important for more than one field, consider creating individual charts per field and per time series, and save charts in the same column of a dashboard group.
- When constructing aggregation queries, ensure that message queries return equivalent results. *Alerts Queries*

- Create alerts primarily for critical events.
- Limit alerts using thresholds. In general, a user should not receive more than six alerts per hour.
- Any saved alerts are disabled after they have been exported as part of a content pack. Email, Webhook and/or vRealize Operations Manager definitions are not included in a content pack.
- Be sure to enter descriptive information about an alert so a user will understand why it is important and who should be notified of the alert.

Dashboards

Dashboard Groups

- Consider starting with an overview dashboard group.
- Create dashboard groups based on a specific type of message (for example, overview, performance, and so on.), not based on a specific type of component (for example, compute, network, storage).
- It is recommended to duplicate the same dashboard widget in multiple dashboards groups if the dashboard widget is applicable in each dashboard group.
- Target at least three dashboard groups in a content pack.
- Dashboard groups and dashboard widgets cannot be reordered, except with user content.

Dashboard Widgets

- Target at least three dashboard widgets per dashboard group.
- Do not put more than three dashboard widgets in the same row.
- Do not put more than eight dashboard widgets in a dashboard group.
- When displaying similar information in different formats, ensure each format brings value.
- Stack related dashboards together for easier viewing.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

- Give the dashboard widgets descriptive names. Do not use field names in widget titles.
- Include notes for every dashboard widget. Ensure that the notes answer questions such as, *"Why is the widget important?"* and *"Where can additional information be found?"*
- Changing the definition of a field does not require that all dashboard widgets created with the previous field definition be re-created to take advantage of the new field definition. Saving the field definition should reflect the change in all occurrences of the field.
- The query definition of a dashboard widget can be modified, using the Edit chart in Interactive Analytics menu.

Content Packs

- A content pack should contain a minimum of three dashboards, nine total widgets, five alerts, and 20 fields.
- When exporting a content pack use the naming format: <*Company*> – <*Product*> *v*<*Version*>. Ideally, the content pack name should be less than 30 characters to prevent word wrapping.
- When exporting a content pack for publishing, export with a namespace & version number.
- When exporting with a namespace, use the namespace format: <*Ext*>.<*Domain*>.<*Product*>.
- When exporting a content pack always export with a detailed description of the product which the content pack addresses and how the content pack helps monitor the product.
- When exporting a content pack always add setup instructions to the content pack that will help the user forward logs to the vRealize Log Insight instance to be able to use the content pack to its potential.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Creating Content Packs in vRealize Log Insight 4.5

CHEAT SHEET TO CREATING CONTENT PACKS

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Introduction

Content packs are read-only plug-ins to vRealize™ Log Insight™ that provide pre-defined knowledge about specific types of events such as log messages. The purpose of a content pack is to provide knowledge about a specific set of events in a format that is easily understandable by administrators, engineers, monitoring teams, and executives. A content pack should answer questions like, *"Is the product/application healthy?"* In addition, a content pack should create a greater understanding of how a product/application works.

A content pack comprises information that can be saved from either the Dashboards or Interactive Analytics pages in vRealize Log Insight. This includes:

- Queries
- Fields
- Aggregations
- Alerts
- Dashboards
- Agent Groups
- Setup Instructions

By default, the current version of vRealize Log Insight ships with the vSphere and General content packs. Other content packs can be imported as required. In addition, any vRealize Log Insight user can create a content pack for private or public consumption.

Intended Audience

This paper provides information about each piece of information that can be saved in a content pack as well as best practices for content pack creation. The information provided is specifically tailored to give content pack authors using vRealize Log Insight 2.5 or later a quick start to developing content packs. (Note: Log Insight 2.5, 3.0, 3.3 and 3.6 content packs are all compatible.)

Step-by-step guide

Create a content pack as per these guidelines - this is a recommendation based on our experiences.

Prerequisites

Authors should have logs pertaining to the content pack to correctly extract meaningful content from the logs.

Import the logs into vRealize Log Insight using LI Agents, Event forwarders. In the current version of vRealize Log Insight, it is possible to ingest events from the command line using the vRealize Log Insight Importer. In short, any file, directory, tarball, or ZIP file can be ingested by copying the events to the vRealize Log Insight virtual appliance by running the importer.

For MS-windows you can use the Importer tool as follows: (tool is also available for the different flavors of linux)

```
C:\my_logs>loginsight-importer.exe –manifest myLogsManifest.txt –source myLogs.tar –server
10.123.345.567 –debug_level 2 –logdir c:\my_logs
```

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Developing Content for a Content Pack

Extracted Fields Components

Fields are one of the most important items in a content pack as they can be used in multiple ways for Aggregations (Allow for functions and groupings to be applied to fields) and Filters (Allow for operations to be performed against fields). From the Interactive Analytics UI in vRealize Log Insight start by creating extracted fields for events. The best practices for extracted fields are:

FIELD	DESCRIPTION
Regex before value	This field should include as many keywords as possible. If this field is empty or only contains special characters, then the Regex after value and/or the Additional Context must include keywords
Name	Only use alphanumeric characters. Ensure all characters are lower case and use underscores instead of spaces as this makes fields easier to view. It is recommended to prefix content pack fields with an abbreviation (for example, vmw_) to avoid confusion
Regex after value	This field should include as many keywords as possible. If this field is empty or only contains special characters, then the Regex before value and/or Additional Context must include keywords
Additional context (keywords)	In vRealize Log Insight 2.5 and newer you can also add keywords to a field by selecting Additional Context to further narrow down your search and improve field performance in a query. Use of Additional Context is recommended for all extracted fields
Additional context (static fields)	In vRealize Log Insight 2.5 and newer you can also add a filter on a static field with operator and a value by selecting Additional Context (filter) to further narrow down your search and improve field performance in a query. Use of Additional Context is recommended for all extracted fields.
Extracted Field Component Notes	In vRealize Log Insight 2.5 and newer releases, notes can and should be added to extracted fields to provide a more in-depth definition of the field, potential values for the field and additional information. It is recommended that all extracted fields contain notes.

- Only create fields for regular expression patterns. If a field can be queried using a single keyword/phrase query, then keyword/phrase queries should be used instead of a pre-defined field. Fields are meant to add structure to unstructured data as well as provide a way to query over specific parts of an event.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

For example, \(\d+\) is correct regex but the query returns all log events that contain numbers in parenthesis. The correct use would be to use this regex with additional context (keyword) such as “vcenter event for vm”, to make it specific. You could also make it specific by using the regex with additional context (filter) on static fields such as source, hostname.

- In an extracted field definition (all components), you should escape literal characters (for example, '.' should be '\.') so they are not treated as regular expressions.
- Extracted field definitions should contain as many keywords as possible. A minimum of one keyword and a recommended minimum of three keywords should be used. A keyword is defined as all sequential alphanumeric, hyphen and/or underscore characters. If keywords cannot be added to an extracted field because they do not exist in the event then either the VRealize Log Insight agent must be used with tags to overcome this limitation or the extracted field will need to be removed for performance reasons. Note that keywords defined within regular expression syntax DO NOT count as keywords.
- For example, Correct – “vcenter” , Incorrect – “vcenter-“ , Correct – “vcenter*” , Correct – “Microsoft Corporation” , Correct – “test.example.com”
- Field names should be lower case, without spaces, and follow the format: <product>[_<component>]_<fieldname>for example, vmw_latency_errors.
- Extracted fields should not be created for known syslog fields including timestamp, hostname and appname, source.
- Good candidates for extracted fields include key/value pairs and delimiter separators, but any useful information including usernames, sources, destinations, protocols, components, severity and status should also be extracted.
- vRealize Log Insight 3.3’s integrated load balancer now allows for multiple VIPs to be configured with zero or more tags. This makes it possible to tag ingested log messages for devices that cannot leverage the Log Insight agent and offers a query performance boost for content packs for logs with limited keywords.

Queries

- Every dashboard should contain a query list with at least three queries defined.
- Ensure queries are specific to events the content pack is designed for.
- Ensure the message query and the aggregation query return the same results, when editing a widget or alert query (TIP: you can check this by running the query in Interactive Analytics and adding filters with exists operator for fields in your chart query but not in your message query if necessary)
- Viz: A message query is a query made up of keywords, regular expressions and/or field operations.
- An aggregation query is a query made up of a function, one or more groupings, and any number of fields.
- If an extracted field has keywords defined in regular expression syntax then add a filter to all queries that use the extracted field for <field> contains <all keywords in regular expression syntax>
- When using filters in queries, do not use the match “any” operator unless one or more keywords are defined in the search bar.
- When using the text filter with multiple different values, one or more keywords should be defined in the search bar.
- Understanding of what "any" means vs "all": "any" means that each filter is a SEPARATE query -- so when multiple filters are used with ‘any’ operator it is actually multiple queries. In general more the queries, the slower the results. Think of "any" as "or" and "all" as "and" operators.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

- Matching AQ (aggregated query) to MQ (message query) is not required for reasons mentioned above when the match “any” operator is used.
- Notes for a widget should explain what the widget is trying to depict and also provide a pointer / web link to an article or website giving more details about the contents of the widget. Ideally, most widgets notes section should contain a link to more information.

Alerts

- Ensure the query is not run more than 6 times an hour
- Ensure alert queries are specific to events the content pack is designed for.
- Ensure the message query and the aggregation query return the same results, when editing a widget or alert query (TIP: you can check this by running the query in Interactive Analytics and adding filters with exists operator for fields in your chart query but not in your message query if necessary)
- Notes for an alert should explain what the widget is trying to depict and also provide a pointer / web link to an article or website giving more details about the contents of the widget. Ideally, most widgets notes section should contain a link to more information.
- Ensure the title of each alert is prefixed with something specific to the content pack. for example, IIS: Server too busy or MS-SQL: Hardware/Media Failure.
- If possible, separate critical alerts (alerts that should typically be enabled) from non-critical alerts (alerts that may or may not be applicable with some prefix such as “*** CRITICAL ***”).

Dashboards

- Extracted field names should not be used in the title of a widget as they could be cryptic and noninformational. for example, vSphere error events by cluster_guid, should be vSphere error events by cluster
- Every dashboard widget should contain detailed notes that go beyond the information provided in the title of the widget. Notes for a widget should explain what the widget is trying to depict and also provide a pointer / web link to an article or website giving more details about the contents of the widget. Ideally, most widgets notes section should contain a link to more information.
- Each dashboard should have a minimum of 3 widgets and not more than 6-8.
- Select the widget to run in Interactive Analytics and check the following points:
- Ensure dashboard widget queries are specific to events the content pack is designed for.
- Ensure the message query and the aggregation query return the same results, when editing a dashboard widget (TIP: you can check this by running the query in Interactive Analytics and adding filters with exists operator for fields in your chart query but not in your message query if necessary)
- Viz: A message query is a query made up of keywords, regular expressions and/or field operations.
- An aggregation query is a query made up of a function, one or more groupings, and any number of fields.
- If an extracted field has keywords defined in regular expression syntax then add a filter to all queries that use the extracted field for <field> contains <all keywords in regular expression syntax>
- Notes for a dashboard widget should explain what the widget is trying to depict and also provide a pointer / web link to an article or website giving more details about the contents of the widget. Ideally, most widgets notes section should contain a link to more information.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

- Each dashboard should have at least one dashboard filter (typically any field used in more than half the widget in the dashboard should be included as a field; for example, hostname though severity and component are good if available/ applicable).
- At least half the widgets in a dashboard should return results for a given dashboard filter.
- Dashboard widget names should start with Capital letter but rest of the letters should be lower case, not all CAPS.

Other General Considerations

- Log Insight was not built to service multiple products in a content pack -- the design is to have a single content pack per product.
- Log Insight does not support having separate content packs for the same underlying logs; due to the way things like extracted fields are handled (i.e. duplicates).
- Content packs should not be duplicating content in part or full in multiple content packs; Unfortunately extracted fields are duplicated affecting performance and creating maintenance woes.
- Content Packs should contain information about the author, website, description, setup instructions and an icon. The icon should be high resolution and should be 144 x 144. Use of non-square logos is not recommended. The icon file size should not be in the double digit and compressed icons are preferred.
- We would like for you to include a link to your End User License Agreement so that our users are aware that the content is not distributable. Please include the link to your EULA at the end of your description. For example: “Please review our End User License Agreement before installation.”
- If you need to modify your content pack for updates in your log content and if your content pack was developed on a version older than the one you will be using it on; import the older content pack in the vRealize Log Insight version where the newer logs are being ingested; complete the development of the content pack, and export it out with the new additions. NOTE: This new version of the content pack using newer features of the vRealize Log Insight may not be compatible on an older version of vRealize Log Insight.
- Information contained in the content pack must be specific to events for which the content pack covers ONLY. You can ensure this by using tags in VIPs, agents and parsers at time of event ingestion or by adding additional context keywords and filters to widgets in your content pack for example, if agent ingests event with tag “Microsoft: IIS”, you can filter on the field ‘Microsoft’ with a value of ‘IIS’ to get IIS events only in your IIS queries. Spelling/Grammatical errors as well as formatting issues reflect on the content pack author. The content pack needs to provide value and the deeper the value the better.

vmware®
VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Feature List and availability in vRealize Log Insight

The following is a listing of features that can be leveraged by different Log Insight versions. Note that Log Insight 2.5, 3.0, 3.3, 3.6, 4.0, 4.3 and 4.4 content packs are all compatible, but features available in newer versions will not work on older versions and simply be ignored.

FEATURE	LOG INSIGHT VERSION
Additional context (keywords/fields)	vRealize Log Insight v2.5 and newer
Setup Instructions for a content pack	vRealize Log Insight v3.0 and newer
Chart type - Table	vRealize Log Insight v3.3 and newer
Multi-VIP + Tags	vRealize Log Insight v3.3 and newer
Chart type – Gauge chart	vRealize Log Insight v4.0 and newer
Trend lines in Charts – Line & Area charts	vRealize Log Insight v4.3 and newer
Alert Recommendation	vRealize Log Insight v4.5 and newer
FEATURE	**LOG INSIGHT VERSION**
Importer Tool	vRealize Log Insight v3.3 and newer
Upgrade Instructions	vRealize Log Insight v3.6 and newer

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Adding Value to Your Content Pack

Do not have content pack with queries that group everything by hostname or severity as you could end up with too many distinct values or too many non-specific hosts if your LI instance is collecting logs from multiple sources.

For security reasons, the events containing key/value pairs should be extracted as fields. There should be a field for IP, info and so on and dashboards should be created showing information from all of these fields. You can use datasets to restrict access to sensitive data fields. See the vSphere content pack security dashboard for examples.

Add different types of charts to dashboards not necessarily, only time based charts that show the reasons for extracting a field and how it can help answer a problem solving question and add these details to the notes of your dashboard widget.

Add recommendations to alerts that will notify the receiver of the alert to take corrective action if necessary. Corrective actions can include but are not limited to tuning alert query or frequency; fix an environment issue.

Testing Your Content Pack

- Be sure to test your content pack on a Log Insight instance that has a very large number of (10+ million preferably) events ingested.
- Test the performance of your queries, dashboard, alerts for content correctness and time to complete the queries.
- If you followed the best practices described in this document the performance of the content pack should be good, however if you observe the performance in terms of time and correctness to be inefficient revisit the extracted field and query definitions to ensure you have followed the best practices and make necessary adjustments.

Publishing Your Content Pack

Done with your content pack! Now you can Publish and upload the content pack to the VMware Solution Exchange website.

Once your content pack has been reviewed and approved by the vRealize Log Insight CORD team via DCPN or Email, you are ready to publish your content pack.

1. Go to http://solutionexchange.vmware.com – If you already have a username and password, click the *Log In Now* link in the top right corner of the page.
2. Enter your username and password and click the *Log In Now* button. – If you do not have a username and password, click the *Register Now* button and select *Register Now* under the Partner Registration Request section. Fill out the required information within the Partner Registration Request and submit. You will receive a notification email if your login request is approved.

3. Once logged in you will be able to update or add new listing by clicking on Administration. This will take you to the Administration menu where you will be able to make changes to your personal account information, your company profile information, and add/edit a solution on your VSX profile.
4. Choose *Manage Solutions* from the Administration menu. – Find the solution that you are looking to edit, choose *Edit* under Actions and proceed to edit your solution. – If your solution is not listed, click *Add Solution* to begin a new listing. Make sure to use the *Save Draft* button frequently, just to make sure that you do not lose any of your work. Once you have completed your solution listing, click *Submit For Approval*. This will send your solution listing to the VSX Alliance Team to be reviewed for approval. You will receive an email regarding the approval of your solution, or information about why your solution may have been declined, along with information about what changes need to be made in order to get it approved.

The information required for filling out the web form to add your listing, and submitting your content pack are listed in the table below. Preparing this information in advance reduces the time it takes to submit the content pack.

REQUIRED INFORMATION	DESCRIPTION
Content Pack Name	VMware prefers a very descriptive name. This will be uploaded and posted in the banner. Examples are the following: • OpenStack Content Pack • NSX for vSphere Content Pack • Microsoft SharePoint Content Pack
Version	1.0 if it is your first submittal, and something greater than 1.0 for enhancements or bug fixes, and so on.
Your Company Logo	This will be uploaded and posted in the banner
Short Description	Short 2-3 sentence descriptions. This will be posted in the banner on VMware Solutions Exchange
Feature Description	3 bullet points describing the features and functions
Long Description	A lengthier 2-3-paragraph description along with a EULA description for users to understand that the content is not distributable.
VLCP File Name	The actual VLCP file which you will upload.

Compatible Technologies	Selection within the form regarding a long list of technologies your Content Pack is compatible with
Documentation	Any supportive documentation. These documents will be uploaded
REQUIRED INFORMATION	**DESCRIPTION**
	to the technical resource area
vRealize Log Insight Compatibility	Compatibility information with vRealize Log Insight, for example: vRealize Log Insight version 2.5, XYX Content Pack version 1.0
Support Contacts	Support information, either a URL, or a phone number.

Resources

Additional information about vRealize Log Insight and vRealize Log Insight content packs can be found using the links below.

VMware vRealize Log Insight documentation:

http://www.vmware.com/support/pubs/log-insight-pubs.html

VMware vRealize Log Insight communities:

http://communities.vmware.com/community/vmtn/vcenter/vcenter-log-insight

VMware vRealize Log Insight marketplace:

https://solutionexchange.vmware.com/store/loginsight

VMware vRealize Log Insight ideas:

http://loginsight.vmware.com

Acknowledgments

The author wishes to thank the following individuals for their technical review of this white paper: Jon Herlocker, Steve Flanders, and Yogita Patil.

About the Authors

Steve Flanders is the Log Insight CTO at VMware. He has an extensive background in designing and implementing cloud solutions with focuses on ensuring scalability and promoting a cloud vision.

Steve has helped architect a number of cloud offerings including VMware's Cloud Foundry, ATT's Synaptic Storage as a Service, and EMC's Atmos Online. Steve is the author of SFlanders.net, a technology-centric weblog focusing on a variety of topics including cloud computing, virtualization, and system administration.

Follow Steve's blog at http://sflanders.net

Follow Steve on Twitter: @smflanders

Yogita Patil is a CORD Engineer at VMware. She has an extensive background in troubleshooting and implementing cloud solutions. Yogita has helped test vRealize Log Insight product from its early days and helped shaped the usability of the product. Yogita is the author of a few blog posts including various content packs for the vRealize Log Insight product.

Follow Yogita on Twitter @nitinyogi6

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com