Nagios Log Server 2024 Technical Overview

High Level Overview

Nagios Log Server is an application that provides organizations a central location to send their machine generated logs that will be indexed and stored for later retrieval or querying and analysis in near real-time. Examples of machine generated logs are:

- Windows Eventlogs
- Linux syslogs
- Mail server logs
- Web server logs
- Application logs

Once log data has been indexed (indexing usually happens within 5 seconds from arrival) it can be easily analyzed using the graphical query and filtering tools on the dashboard. It also provides quick search functionality to search any log event item on Google, Bing, and Stack Overflow.

Alerts can be created based on the query used in the dashboard and:

- Send to Nagios XI or Nagios Core via NRDP
- Send as a SNMP Trap
- Send an Email
- Execute a custom script

Finally, the data that is sent to Nagios Log Server can be automatically archived to a shared network drive. The archived data can be restored and re-analyzed at any point in the future.

What that means in plain English is that it can be used to record any log events that are happening across all of the machines and network devices organization wide. Users of Nagios Log Server can access all of this data in a central location, searching it through the UI. Having all of the data in one location has the added benefit of being able to compare or correlate log data from multiple devices. The automated archiving of the log data will assist in maintaining compliance with certain standards that require log data to be stored for a specific amount of time.

An Example Application

- An obvious application where Log Server could be used would be as an advanced system to analyze the received logged events and send important items (e.g. Critical Errors) to Nagios Core or Nagios XI for alerting

Less Obvious Applications

- Developers can send debug logs to Nagios Log Server, and easily filter out the information that isn't important leaving just the key items of interest
- Organizations can utilize the graphical and analytic capabilities of Nagios Log Server to analyze web server logs, not only for errors, but to determine what are the most requested pages, what is the geo-location of their visitors, popular browsers and more
- With a small script, users could archive Nagios Check Results including performance data, and have the ability to setup custom dashboards visualizing the data however they wish (table, bar graph, pie chart)
- Nagios Log Server could be used to index and archive message from an IMAP mailbox, for security or historical reference
- Nagios Log Server can also receive SNMP Traps, again allowing all of the previous talked about functionality on the Traps received

The Benefits of Nagios Log Server Over Text Based Systems

Nagios Log Server allows all of your organizations machine generated data to be stored and indexed in one central location, allowing for queries to be performed on all of the log data at the same time providing the ability for correlative analysis.

This data can be presented to the user running the query in customized views including a table of results, bar charts, pie charts, line graphs, etc. Fields in the logs that are determined to be numeric can have calculations done when creating / using the graphing / table functionality to provide data like total, min, max, mean, etc.

Log Server Terminology

The following section will outline common term used in Nagios Log Server and their meaning.

- Instance - A single Nagios Log Server installation. All Instances become of a member of a Nagios Log Server Cluster
- Cluster - A collection of Nagios Log Server Instances. Each member of the Cluster stores a portion of the data in the Indexes as well as shared in the workload of performing indexing, searching, alerting and maintenance operation.
- Elasticsearch - Storage/Indexing engine used in Nagios Log Server
    - Lucene - Elastic search is built on top of Apache Lucene which is a full text indexing and search engine. Coincidentally, Lucene was first introduced in 1999 the same year Ethan Galstad created NetSaint. Lucene joined the Apache Project just after NetSaint was Renamed to Nagios.
- Logstash - Package used to receive and pre-process log messages before sending them to Elasticsearch for indexing. Logstash has dozens of possible Inputs and Filters that can be added through the Configure menu to enable additional capabilities.

  - Inputs - Additional Inputs can be added to allow Nagios Log Server to collect data from various places, like TCP/UDP ports, SNMP Traps, unix sockets, long running command pipes, etc.
  - Filters - Filters can be applied to messages before they are sent to Elasticsearch for indexing. These can contain items such as breaking apart messages into fields for easy searching, adding geo location information, resolving IP to DNS names, dropping messages you do not want indexed, etc.
- Kibana - A visualization package that was used as a base for the dashboards in Nagios Log Server. Nagios Log Server's version of Kibana has been highly modified to give users private dashboards, added alerting capabilities, save/save as functionality, and the ability to create and save queries and apply them to different dashboards.
  - Query - One or more queries can be added giving them different color representation for each items matching the query fields.
  - Filtering - Filtering limits the items in the result set to either contain or not contain certain elements. On a technical note, Filters are cached in Elasticsearch and usage of filters can dramatically speed up the response time.
  - Row - The dashboard is made up of one or more rows. Rows have a specific height and have a total span of 12.
  - Panel - Panels are added to rows and are sized according to the span specified and the height of the Row in which they reside. Panels can be dragged and dropped to be placed in different Rows. There are many different panel types allowing users to create combinations of graphs, table views, or even add a text/HTML block describing other elements on the page.
- Index - The primary unit of storage, an Elasticsearch index, akin to a database. Nagios Log Server created a new index for each days logs. Additionally there are several indexes used to store setting, user information and internal audit logs.

    - Shard - Each Index in broken up into multiple shards, which allows the distributing of data and workload across the cluster. By default each index in Nagios Log Server has 5 shards. Each shard is a single Lucene instance. Nagios Log Server knows exactly which Shard in the cluster contains each log message, so for direct retrieval it does not need to interact with and other Shards.
- Snapshot Repository - This must be a shared Network file system that is accessible by ALL Instances. The Repository will hold all archived Snapshots until they are deleted
- Snapshot - A backup of a particular index. Snapshots can be restored at future dates to if data need to be recovered.
- Cluster Hostname - Utilized primarily in the Help Sections, a cluster hostname can be specified and will populated Help section with that name in lieu of the IP/hostname of the directly access machine. This allows those using round robin DNS or load balancers to send logs to one hostname and have it distributed to any member of the cluster.