

Nagios Log Server 2024 Performance and Storage Walkthrough

Minimum Requirements for Nagios Log Server

The base VMware virtual machine running CentOS 7.x with Nagios Log Server installed has the following default performance statistics:

- 2 x CPU (single core)
- 2GB of RAM
- 100 GB of disk space

This is the minimum requirements that we have set to work consistently with a 500 MBs of log data a day. This should be used as a baseline since the numbers will easily scale. Here is the download page that contains a vSphere OVF templates, VMware Virtual Machines and a source installer:

<https://www.nagios.com/downloads/nagios-log-server/>

Many customers have reported that memory is one of the first resources that needs increasing on a Nagios Log Server instance, 8GB seems to be the common amount of RAM required.

Nagios Log Server Cluster

Nagios Log Server is a clustered application, it consists of one or more instances of Nagios Log Server. An instance is an installation of Nagios Log Server, it participates in the cluster and acts as a location for the received log data to reside. The log data is spread across the instances using the Elasticsearch database, a special database used by Nagios Log Server.

We will describe how Nagios Log Server operates using an example system cluster called nagioslsdev so we can quantify requirements throughout this document.

Instance	Total Disk Space	Memory	Processor	Primary Shard	Replica Shard
nagioslsdev01	1.5 TB	16 GB	2GHz / dual core	500 GB	500 GB
nagioslsdev02	1.5 TB	16 GB	2GHz / dual core	500 GB	500 GB
nagioslsdev03	1.5 TB	16 GB	2GHz / dual core	500 GB	500 GB
nagioslsdev04	1.5 TB	16 GB	2GHz / dual core	500 GB	500 GB

- Total Disk Space
 - The actual disk space on each Nagios Log Server Instance
 - Using Solid State Drives (SSD) are recommended, they significantly boost queries and indexing
- Memory
 - Relates to log processing power and speed of displaying, querying and filtering dashboards
 - This larger amount of memory will allow more sophisticated filters and queries as the entire database will need to be read to display the requested dashboard(s)
 - Memory is more important than processing power
- Processor
 - More CPUs cores will perform better than CPUs with strictly better clock-speeds
- Primary and Replica Shards
 - Each Index in a cluster has a primary and a replica shard, a new index is created every day
 - If you have 2 instances in a cluster then all the data is replicated

- When you have 3 or more instances the data is distributed over all the instances allowing for better redundancy and can improve query performance

Let's assume the following about our example cluster:

- The cluster is receiving 500 GB of data every day
- There are 4 Instances in our cluster each with a 1.5 TB storage capacity
- Each Storage index will have a 500 GB primary shard and a 500 GB replica shard per day
- 500 GB/day Primary shard and 500 GB/day Replica shards = 1 TB/day
- Over a 30 day period the total data is 30 TB (1 TB a day for each daily index)
- Each system will need extra space for the operating system and other locally run processes

Performance Balancing

In our example we have 4 instances that will share the overall data storage and because of this they will also share processing and memory power when running the queries on their local primary index shard. At this point each shard has been allocated, you can find out if the shards have been successfully allocated under the Admin > System > Cluster Status page.

When running a query for something like '*error' in the dashboard, Nagios Log Server will search through each primary shard from each Instance, which processes the request with its native resources. Once each cluster has been searched through, processed and collected the dashboards graphs and tables are drawn with the data query result.

What Happens When An Instance Goes Down?

There will be times where an instance in a cluster will lose connection or go down. Nagios Log Server is able to account for this by reallocating shards from the server into Instances that are still working. This allows for the Administrator to react to the downed machine without losing functionality.

The Replica shard is designed to replace a primary shard if one of the servers was to go down. For example let's say that nagioslsdev02 went down. The other instances that contain the replica shards for the instance that is currently down would use their replica shards to reallocate the missing data from the downed server and distribute them among the healthy servers. This way none of the data from nagioslsdev02 will be lost since the replica shard recreated the missing primary shards in the cluster and the cluster will operate as normal until the downed server can be fixed.

In Nagios Log Server no single instance is more important than another instance, it is the cluster functionality which provides this flexibility. There is always a "master" instance however if the master goes down then another one of the instances takes over instantly.

Archived Logs - Snapshots

When you are bringing in archived logs (snapshots) from another Elasticsearch, Logstash, Kibana (ELK) configuration you will not need to allocate extra space for these logs. Snapshots logs are going to be located on a file system with separate storage so they won't take any space on your current clusters drives. More information on this can be found in the following documentation:

[Backing Up And Restoring Nagios Log Server](#)

Bottlenecks

When you look at the nagioslsdev cluster you will observe that they all have the same amount of memory. If, for example, 3 of the instances in the cluster had 16 GB of memory and the 4th instance only had 4 GB of memory, there would be a significant delay for the smaller instance to finish a query over its data. The amount of data will be the same over the 4 instances, but the slower instance will finish its query request 4 times slower than the other 3 instances.

In this situation you could remove the lower memory instance and this would increase the performance of the cluster dramatically. This should be considered when creating or engineering your Nagios Log Server cluster.