



## Purpose

This document explains the settings in Nagios Fusion that can be tuned to improve performance.

## Target Audience

This document is intended for use by Nagios Administrators who need to understand what options are available for performance tuning.

## Summary

Nagios Fusion functions as a high-level dashboard for multiple instances of Nagios XI and Nagios Core. In larger deployments of Nagios XI and Nagios Core you may be required to tune certain performance settings to improve response times and overall load on the Fusion server. There are performance settings in Nagios Fusion that can be changed globally or per Fused server.

## Global Settings

The settings that apply globally can be adjusted by navigating to **Admin > System Configuration > System Settings > Data & Polling**.

The screenshot shows the Nagios Fusion interface. The top navigation bar includes 'Home', 'Views', 'Dashboards', 'Servers', 'Help', and 'Admin' (circled in blue). The left sidebar shows a menu with 'System Settings' circled in blue. The main content area is titled 'System Settings' and has two tabs: 'General' and 'Data & Polling' (highlighted in yellow). Under 'Data & Polling', there are two settings:

- Global Authentication Interval:** 300 Seconds. Define the frequency at which the fused servers are checked for authentication.
- Global Polling Interval:** 300 Seconds. Define the frequency at which the fused servers are polled for data.

## Default Interval Settings

- **Global Authentication Interval**
  - Define the frequency (seconds) at which the fused servers are checked for authentication
  - Authentication is a single check against the server to ensure both connectivity (information used for polling) and that the specified credentials are still correct
  - You want to keep this value relatively low
  - You can override this setting on a per-server basis ([Per Fused Server Settings](#))
- **Global Polling Interval**
  - Define the frequency (seconds) at which the fused servers are polled for data
  - Larger systems may be able to benefit from a larger number here
  - You can override this setting on a per-server basis ([Per Fused Server Settings](#))

## Polling Configuration

### Polling Configuration

<b>Polling Subsystem Memory Limit:</b>	<input type="text"/>	The maximum amount of memory that a poller may consume while running. ?
<b>Simultaneous Pollers:</b>	<input type="text" value="1"/> Pollers	The maximum number of pollers that Fusion can spawn to grab data. ?
<b>Live Data Timeout:</b>	<input type="text" value="45"/> Seconds	The length of time that Fusion waits for a chunk of data from a fused server before timing out. ?
<b>Polling Lock Max Age:</b>	<input type="text" value="1200"/> Seconds	The maximum amount of time that a poller can take while grabbing data from a fused server. ?
<b>Poll Record Count:</b>	<input type="text" value="1000"/> Records	The number of records to grab in a single API call to a fused server. ?

- **Polling Subsystem Memory Limit**
  - The maximum amount of memory that a poller may consume while running.
  - This is same value that would be accepted by `php.memory_limit`
  - Acceptable values can be:
    - `128M` = 128 Megabytes

- 1G = 1 Gigabyte
  - -1 = unlimited memory
  - If you're connecting to large systems, you may want to keep this value high
- **Simultaneous Pollers**
    - The maximum number of pollers that Fusion can spawn to grab data
    - 1 poller should be sufficient
    - If you're connecting to at least one large system, you'll want to increase this value so that the polling subsystem doesn't get choked waiting on a large batch of API calls
    - A decent formula in that case would be at least one poller for every fused server
  - **Live Data Timeout**
    - The length of time (seconds) that Fusion waits for a chunk of data from a fused server before timing out
    - This is the timeout value for each individual API call made to a fused server
    - This can be kept relatively low unless you are connected to a large Nagios Core install
  - **Polling Lock Max Age**
    - The maximum amount of time (seconds) that a poller can take while grabbing data from a fused server
    - If something goes wrong with a poller, and it isn't able to remove the lock on the fused server it's grabbing data from, then this is the amount of time that the subsystem will wait before it spawns a new poller
  - **Poll Record Count**
    - The number of records to grab in a single API call to a fused server
    - When fusion polls the data, it calls the API and grabs only a handful of records at a time
    - This is the value you set to determine the amount of records it grabs
    - If you don't have any fused servers that are large systems, then you can keep this number very high

## Database Retention Settings

## Database Retention Settings

Log Data Retention:  Days

The maximum amount of time to retain log data in the database. ?

- **Log Data Retention**

- The maximum amount of time (days) to retain log data in the database
- If you have debug data being written to the database, you'll want to keep this value as low as 1 or 2 days while that's enabled
- Having the log data in the database allows for a nice interface where you can view recent logs, but this is not meant as a replacement for the log file

## Per Fused Server Settings

Navigate to **Admin > Servers > Managed Fused Servers**. The settings that can be adjusted per Fused server can be defined when adding a new Fused server or by editing an existing Fused server.

## Polling Settings

## Polling Settings

Authentication Interval:  Seconds

Define the frequency at which this server is checked for authentication. ?

Polling Interval:  Seconds

Define the frequency at which this server is polled for data. ?

 Ignore XI host/service meta data ?*Enabling this option is not recommended unless you have an extremely large system.*

- **Authentication Interval**

- Define the frequency at which this server is checked for authentication
- Authentication is a single check against the server to ensure both connectivity (information used for polling) and that the specified credentials are still correct
- You want to keep this value relatively low

- **Polling Interval**

- Define the frequency at which this server is polled for data
- Larger Nagios XI and Core systems will benefit from a larger number here, due to the fact that it takes

longer to poll and process the data for these large systems.

- **Ignore XI host/service meta data**
  - Only used if you have an extremely large system
  - This only applies to XI servers that you're fusing to. Meta data refers to the following data:
    - Hosts and Services that are disabled
    - Host and Service acknowledgment information
    - Hosts and Services that are flapping
    - Hosts and Services that are in downtime
    - For a large system, this can allow you to set a much lower value for the `polling_interval`, as poll speed is reduced by up to 70%
    - This is not recommended for an XI system with less than 20,000 Host and Service checks

## Finishing Up

This completes the documentation on performance tuning in Nagios Fusion.

If you have additional questions or other support related questions, please visit us at our Nagios Support Forums:

<https://support.nagios.com/forum>

The Nagios Support Knowledgebase is also a great support resource:

<https://support.nagios.com/kb>