Monitoring Nginx Server

Restricted Rights Legend

The information contained in this document is confidential and subject to change without notice. No part of this document may be reproduced or disclosed to others without the prior permission of eG Innovations Inc. eG Innovations Inc. makes no warranty of any kind with regard to the software and documentation, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose.

Trademarks

Microsoft Windows, Windows 2008, Windows 7, Windows 8, Windows 10, Windows 2012 and Windows 2016 are either registered trademarks or trademarks of Microsoft Corporation in United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright

©2016 eG Innovations Inc. All rights reserved.

Table of contents

MONITORING NGINX SERVER 1
1.1 Enabling the Nginx status page 2
1.2 The Web Server Layer 4
1.2.1 Nginx Status Test 4
CONCLUSION 7

Table of Figures

Figure 1.1: The layer model of the Nginx Server 1
Figure 1.2: Opening nginx.conf file 2
Figure 1.3: Adding code block for enabling the status page 3
Figure 1.4: Sample output massage 4
Figure 1.5: The tests mapped to the Web Server layer 4

1

Monitoring Nginx Server

Nginx (pronounced “engine-x”) is a free, open-source, high-performance HTTP server (aka web server) and reverse proxy, as well as an IMAP/POP3 proxy server. Therefore, NGINX provides a unique combination of web server, caching proxy and load balancing solution to any website that just wants to be consistently efficient. Because of its design and architecture NGINX has already enabled more performance, scalability, reliability and security to many organizations across the world. Today NGINX is one of the most popular open source web servers on the Internet.

NGINX optimizes the usage of the operating system and the hardware resources with its modular, event-driven, asynchronous, non-blocking architecture. Using event notifications and asynchronous handling of a variety of consequent actions associated with accepting, processing and managing network connections and content retrieval, NGINX provides hints to the operating system and gets timely feedback in regards to when expect an inbound or outbound traffic, when check disk operation, when refresh content and so on. Therefore administrators prefer to monitor the Nginx Server and obtain critical metrics from the server such as the Active connections, connections accepted, Total requests etc. To cater to the requirements of the administrators, eG Enterprise provides a specialized *Nginx Server* monitoring model. Figure 1.1 depicts the model used by the eG Enterprise suite to monitor an Nginx Server.

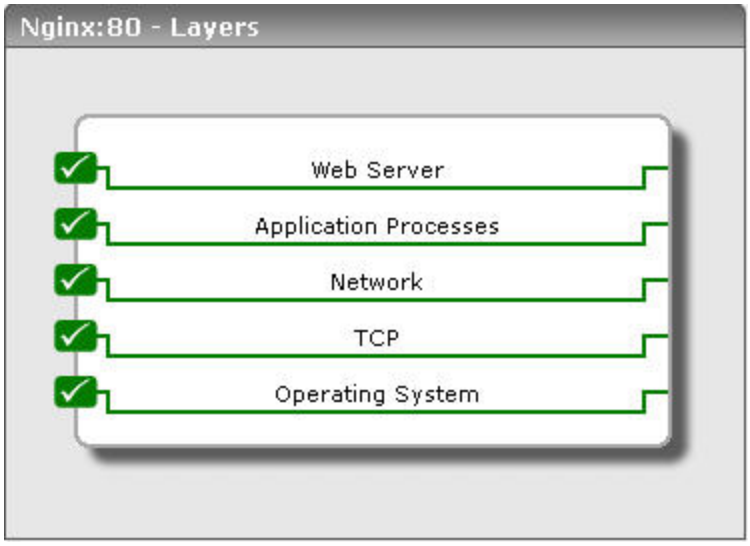

Figure 1.1: The layer model of the Nginx Server

The **Operating System**, **Network**, **TCP** and **Application Processes** layers of an Nginx Server model are similar to that of a Windows Generic server model. Since these tests have been dealt with in the *Monitoring Unix and Windows Servers* document, Section 1.1 focuses on the Web Server layer.

1.1 Enabling the Nginx status page

The Nginx status test fetches and reports the status of Nginx server from a Nginx status page. If the Nginx status test fails to report measures, first check whether the Nginx status page is enabled or not. If the status page is not enabled, then, you need to enable it manually. To enable the status page in the Windows system, follow the steps given below:

1. Go to the configuration folder and open *nginx.conf* file to edit.

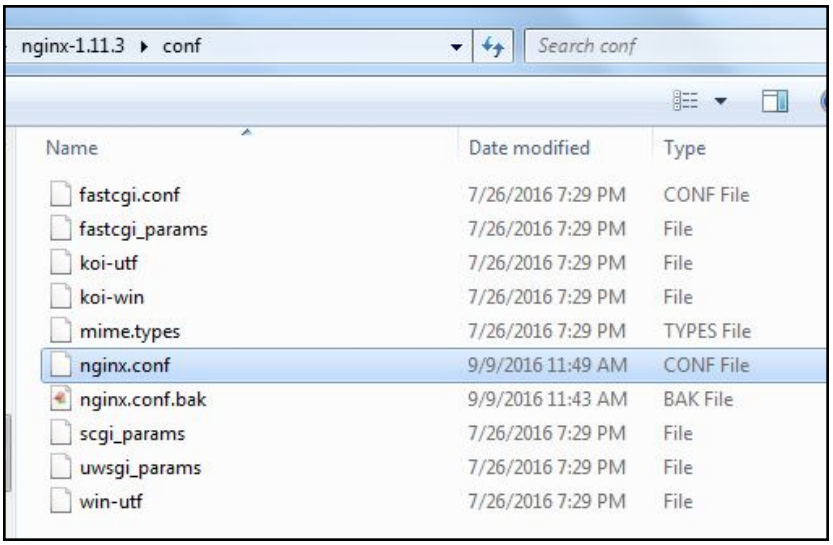

Figure 1.2: Opening nginx.conf file

2. Then, include a code block as shown in Figure 1.3, inside the server block of the *nginx.conf* file.

```
server {
    listen       80;
    server_name  localhost;

    #charset koi8-r;

    #access_log  logs/host.access.log  main;

    location / {
        root   html;
        index  index.html index.htm;
    }

location /nginx_status {
    # Turn on nginx stats
    stub_status on;
    # I do not need logs for stats
    access_log   off;

}

    #error_page  404              /404.html;

    # redirect server error pages to the static page /50x.html
    #
    error_page   500 502 503 504  /50x.html;
    location = /50x.html {
        root   html;
    }
```

Figure 1.3: Adding code block for enabling the status page

3. Once you added the code block, save the *nginx.conf* file and reload it again for reflecting the changes.
4. To check whether the status page is enabled or not, go to a browser and type the URL as "example.com/nignx_status"
5. If you see output message as shown in Figure 1.4, you can be assured that the nginx status page is enabled and the Nginx Status test will report measures.

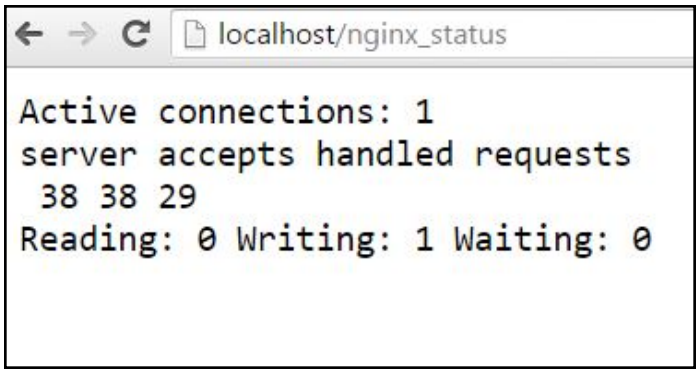

Figure 1.4: Sample output massage

1.2 The Web Server Layer

This layer tracks the health of the Nginx Server and the critical performance statistics of the Nginx Server. Since the HTTP test is discussed in detail in the Monitoring Web Servers document, let us now discuss the Nginx Status test in the following section.

Figure 1.5: The tests mapped to the Web Server layer

1.2.1 Nginx Status Test

This test reveals critical performance statistics pertaining to an Nginx server. This test, upon execution, accesses a specific URL on the Nginx server, which contains the required metrics.

Target of the test : An Nginx server

Agent deploying the test : An external agent

Outputs of the test : One set of results for the URL that is to be accessed

Configurable parameters for the test

1. **TEST PERIOD** - How often should the test be executed
2. **HOST** - The host for which the test is to be configured.
3. **PORT** – The port number at which the specified **HOST** listens to.
4. **URL** – In the **URL** text box, the URL to be accessed by this test for extracting the performance statistics of the Nginx server, will be displayed by default. The URL is: *http:///{Nginx web server host}:{Nginx web server port}/nginx_status*.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Active connections:	Indicates the number of connections that are currently active on this server.	Number	A high value is desired for this measure.
Connections accepted:	Indicates the number of connections that are currently accepted by this server.	Number	
Connections handled:	Indicates the number of connections that are currently handled by this server.	Number	
Total requests:	Indicates the number of requests that are currently handled by this server.	Number	
Reading request headers:	Indicates the number of request headers that are currently read by this server.	Number	
Writing requests:	Indicates the number of write responses provided to the client by this server after reading and processing the	Number	

Measurement	Description	Measurement Unit	Interpretation
	requests.		
Waiting requests:	Indicates the number of waiting connections i.e., keep alive connections (including reading and writing connections).	Number	
Requests per connection:	Indicates the average number of requests that are handled for each connection.	Number	The value of this measure is a ratio of the measures - Total requests: Connections handled.

2

Conclusion

This document has described in detail the monitoring paradigm used and the measurement capabilities of the eG Enterprise suite of products with respect to the **Nginx Server**. For details of how to administer and use the eG Enterprise suite of products, refer to the user manuals.

We will be adding new measurement capabilities into the future versions of the eG Enterprise suite. If you can identify new capabilities that you would like us to incorporate in the eG Enterprise suite of products, please contact support@eginnovations.com. We look forward to your support and cooperation. Any feedback regarding this manual or any other aspects of the eG Enterprise suite can be forwarded to feedback@eginnovations.com.