Monitoring NTP Server

eG Innovations Product Documentation

www.eginnovations.com

Total Performance Visibility

Table of Contents

Table of Figures

Chapter 1: Introduction

A time server is a server computer that reads the actual time from a reference clock and distributes this information to its clients using a computer network. The protocol most widely-used by time servers for distributing and synchronizing time over the Internet is the Network Time Protocol (NTP). The term NTP applies to both the protocol and the client/server programs that run on computers. The programs are compiled by the user as an NTP client, NTP server, or both. In basic terms, the NTP client initiates a time request exchange with the NTP server. As a result of this exchange, the client is able to calculate the link delay and its local offset, and adjust its local clock to match the clock at the server's computer.

On the other hand, if for any reason, the client is unable to contact the NTP server, time synchronization will not occur, resulting in serious failures - for instance, scheduled tasks may not run on time on the client, SSL certificate validity checks may go awry, domain controllers may not be able to authenticate the Windows clients, etc.

To avoid such ill effects, administrators must periodically check whether the NTP server is accessible to clients, check the responsiveness of the server to client requests, and if possible, even determine how different the client's time is from the server's time. This way, if a sudden loss of communication occurs between the client and the NTP server or if the time difference between the client and server is abnormally high, administrators can promptly detect the same and rapidly initiate remedial measures. This is where eG Enterprise helps administrators to achieve their duty in a smooth way.

Chapter 2: How to Monitor the NTP Server Using eG Enterprise?

eG Enterprise monitors the NTP Server using an **eG external agent** be deployed on any remote host – for example, an NTP client - in the environment. This external agent will run tests on the NTP server non-intrusively to check the network availability and accessibility of the NTP server, report how long the server takes to respond to client requests, and also measure the time difference between the client and the server.

2.1 Managing the NTP Server

The eG Enterprise cannot automatically discover the NTP Server. This implies that you need to manually add the component for monitoring. Remember that the eG Enterprise automatically manages the components that are added manually. To manage a NTP Server component, do the following:

1. Log into the eG administrative interface.
2. Follow the Components -> Add/Modify menu sequence in the **Infrastructure** tile of the **Admin** menu.
3. In the **COMPONENT** page that appears next, select *NTP Server* as the **Component type**. Then, click the **Add New Component** button. This will invoke Figure 2.1.

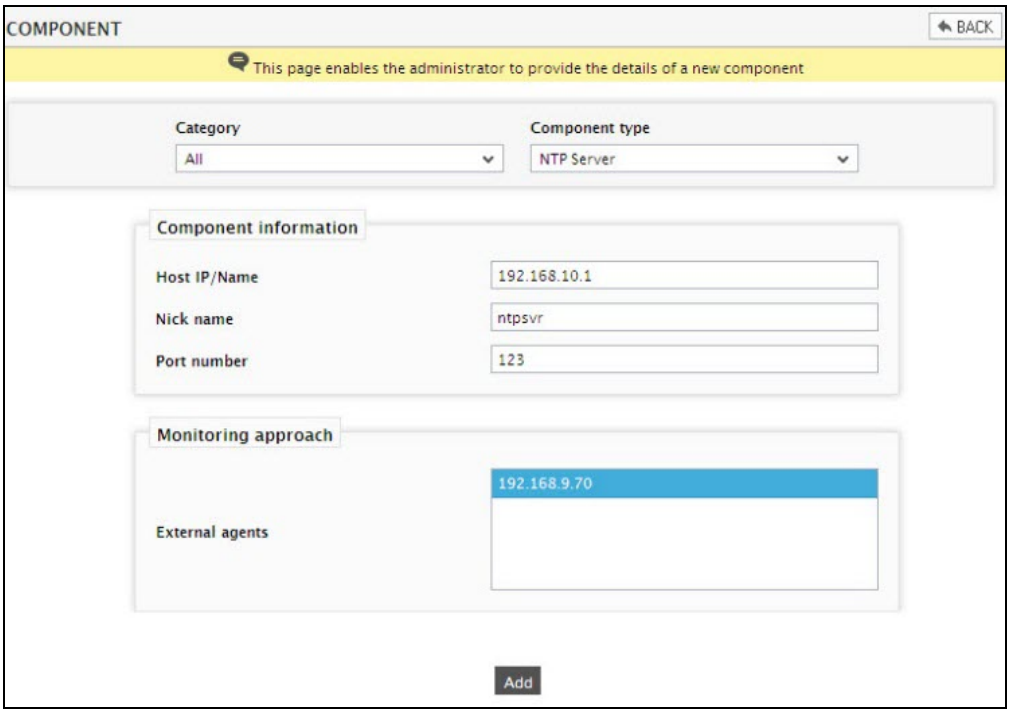

Figure 2.1: Adding a NTP Server

4. Specify the **Host IP** and the **Nick name** for the NTP Server in Figure 2.1. Then click the **Add** button to register the changes.
5. Finally, signout of the eG administrative interface.

Chapter 3: Monitoring the NTP Server

eG Enterprise offers a specialized NTP Server monitoring model to monitor the availability and overall health of the NTP server.

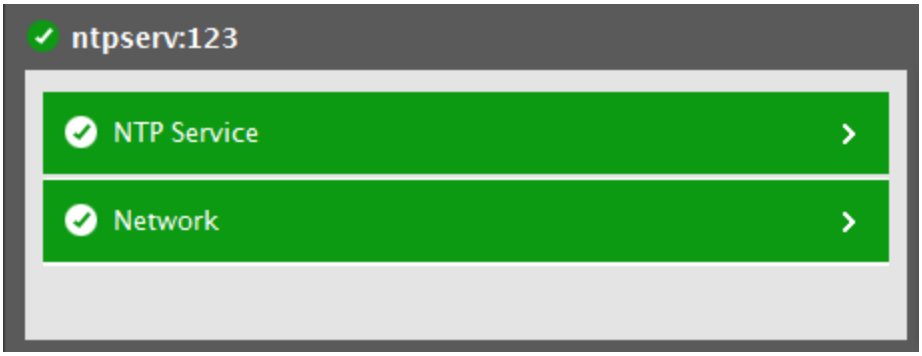

Figure 3.1: The layer model of the NTP Server

The **Network** layer of this model is mapped to a **Network** test that pings the NTP server at configured intervals to evaluate the health of the network connection between the client and the server. If the network link to the server is of a poor quality and may potentially break, this test will proactively alert administrators to it. Since this test has already been discussed at length in the *Monitoring Unix and Windows Servers* document, let us proceed to the NTP Service layer.

3.1 The NTP Service Layer

The NTP Service layer is associated with the **NTP Time Check** test that periodically checks time synchronization between an NTP client and server. The below section provides you great details on configuring the parameters of the **NTP Time Check** test and the metrics that the test reports.

3.1.1 NTP Time Check Test

The absence of time synchronization between an NTP client and server can have serious repercussions on the performance and operations of the client - for instance, scheduled tasks such as virus scans or backup routines may not run on time on the client, SSL certificate validity checks may go awry, domain controllers may not be able to authenticate Windows clients, etc. If these adversities are to be avoided, administrators should be proactively alerted to a potential non-sync between the client's time and the server's time and should also receive a 'heads-up' on the probable reasons for the same. This is where the **NTP Time Check** test helps! This test periodically checks the accessibility and responsiveness of the NTP server from an external location, and also indicates

how different the client's time is from the server's time. In the process, the test not only points to a time non-sync, but also reveals the probable reasons for the same - is it because the NTP server is down? Is it because the NTP server is slow in processing client requests? Or is it because the gap between the server's time and the client's time is very high?

Target of the test : An NTP server

Agent deploying the test : An external agent

Outputs of the test : One set of results for the target Server Node being monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port number at which the specified Host listens to.
Report Clock Offset	By default, this test reports the time difference between the NTP client and the server (the Clock offset value measure) and also indicates whether the client's time is ahead or behind the server's (the Client time relative to server time measure). This is why, the Report Clock Offset flag is set to **Yes** by default. However, the measures mentioned above are of significance only to an NTP client, which has to sync time with the monitored NTP server – say, a member server of a Windows domain that needs to sync time with its domain controller. On the contrary, for a host that does not seek to sync time with the NTP server, these two measures are meaningless! Such a situation may arise, if, owing to security constraints, an administrator prefers to deploy the external agent (that executes this test) on some remote host that need not sync time with the NTP server that is being monitored. Under such circumstances, the administrator may just want the test to report whether the NTP server is up and running or not, and if running, how responsive it is to requests. In this case, its best to turn off the report clock offset flag by setting it to **No**, so that the Clock offset value measure and the Client time relative to server time measure are no longer reported by the test.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	Indicates whether/not the NTP server is available.	Percent	If this measure reports the value 100, it indicates that the NTP server is accessible. The value 0 on the other hand indicates that the

Measurement	Description	Measurement Unit	Interpretation
			NTP server cannot be connected to.
Roundtrip delay	Indicates the length of time it takes for a signal to be sent plus the length of time it takes for an acknowledgment of that signal to be received. This time delay therefore consists of the transmission times between the two points of a signal.	Secs	To synchronize its clock with a remote server, the client must compute the round-trip delay time and the offset. The round-trip delay δ is computed as: $\delta = (t_3 - t_0) - (t_2 - t_1)$ where **t0** is the client's timestamp of the request packet transmission, 100 **t1** is the server's timestamp of the request packet reception, 150 **t2** is the server's timestamp of the response packet transmission and 160 **t3** is the client's timestamp of the response packet reception. 120 The shorter and more symmetric the round-trip time, the more accurate the estimate of the current time.
Clock offset of client	Indicates the number of seconds the client must add to its time to synchronize with the server's time.	Secs	The offset θ is given by $\theta = \frac{(t_1 - t_0) + (t_2 - t_3)}{2}$ A positive value indicates the server clock is higher. A negative value indicates the client clock is higher. Normally, if the client offset exceeds NTP's default panic threshold of 1000 secs, NTP exits with a message to the system log. You can however, configure NTP to allow the time to be set to any value without restriction; but, this can happen only once. If the panic threshold is exceeded after that, NTP will exit

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>with a message to the system log.

You can use the detailed diagnosis of this measure to know the client's time stamp, the server's time stamp, and the offset.</td></tr>
<tr><td>Client time relative to server time</td><td>Indicates whether the client is behind / ahead of the server in terms of time.</td><td></td><td>If the client's clock is running faster than the server's, this measure will report the value Ahead. If the client's clock is running slower than the server's, this measure will report the value Behind.

The numeric values that correspond to these measure values are as follows:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Ahead</td><td>1</td></tr><tr><td>Behind</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed above to indicate whether client is ahead or behind the server. In the graph of this measure however, the same represented using the numeric equivalents.</td></tr>
</table>

3.2 The NTP Service Layer

The NTP Service layer is associated with the **NTP Time Check** test that periodically checks time synchronization between an NTP client and server. The below section provides you great details on configuring the parameters of the **NTP Time Check** test and the metrics that the test reports.

3.2.1 NTP Time Check Test

The absence of time synchronization between an NTP client and server can have serious repercussions on the performance and operations of the client - for instance, scheduled tasks such as virus scans or backup routines may not run on time on the client, SSL certificate validity checks may go awry, domain controllers may not be able to authenticate Windows clients, etc. If these adversities are to be avoided, administrators should be proactively alerted to a potential non-sync

between the client's time and the server's time and should also receive a 'heads-up' on the probable reasons for the same. This is where the **NTP Time Check** test helps! This test periodically checks the accessibility and responsiveness of the NTP server from an external location, and also indicates how different the client's time is from the server's time. In the process, the test not only points to a time non-sync, but also reveals the probable reasons for the same - is it because the NTP server is down? Is it because the NTP server is slow in processing client requests? Or is it because the gap between the server's time and the client's time is very high?

Target of the test : An NTP server

Agent deploying the test : An external agent

Outputs of the test : One set of results for the target Server Node being monitored.

Configurable parameters for the test

Parameter	Description
Test period	How often should the test be executed.
Host	The host for which the test is to be configured.
Port	The port number at which the specified Host listens to.
Report Clock Offset	By default, this test reports the time difference between the NTP client and the server (the Clock offset value measure) and also indicates whether the client's time is ahead or behind the server's (the Client time relative to server time measure). This is why, the Report Clock Offset flag is set to **Yes** by default. However, the measures mentioned above are of significance only to an NTP client, which has to sync time with the monitored NTP server – say, a member server of a Windows domain that needs to sync time with its domain controller. On the contrary, for a host that does not seek to sync time with the NTP server, these two measures are meaningless! Such a situation may arise, if, owing to security constraints, an administrator prefers to deploy the external agent (that executes this test) on some remote host that need not sync time with the NTP server that is being monitored. Under such circumstances, the administrator may just want the test to report whether the NTP server is up and running or not, and if running, how responsive it is to requests. In this case, its best to turn off the report clock offset flag by setting it to **No**, so that the Clock offset value measure and the Client time relative to server time measure are no longer reported by the test.

Measurements made by the test

Measurement	Description	Measurement Unit	Interpretation
Availability	Indicates whether/not the NTP server is available.	Percent	If this measure reports the value 100, it indicates that the NTP server is accessible. The value 0 on the other hand indicates that the NTP server cannot be connected to.
Roundtrip delay	Indicates the length of time it takes for a signal to be sent plus the length of time it takes for an acknowledgment of that signal to be received. This time delay therefore consists of the transmission times between the two points of a signal.	Secs	To synchronize its clock with a remote server, the client must compute the round-trip delay time and the offset. The round-trip delay δ is computed as: $\delta = (t_3 - t_0) - (t_2 - t_1)$ where **t0** is the client's timestamp of the request packet transmission, 100 **t1** is the server's timestamp of the request packet reception, 150 **t2** is the server's timestamp of the response packet transmission and 160 **t3** is the client's timestamp of the response packet reception. 120 The shorter and more symmetric the round-trip time, the more accurate the estimate of the current time.
Clock offset of client	Indicates the number of seconds the client must add to its time to synchronize with the server's time.	Secs	The offset θ is given by $\theta = \frac{(t_1 - t_0) + (t_2 - t_3)}{2}$ A positive value indicates the server clock is higher. A negative value indicates the client clock is higher. Normally, if the client offset exceeds NTP's default panic threshold of 1000 secs, NTP

<table>
<tr><th>Measurement</th><th>Description</th><th>Measurement Unit</th><th>Interpretation</th></tr>
<tr><td></td><td></td><td></td><td>exits with a message to the system log. You can however, configure NTP to allow the time to be set to any value without restriction; but, this can happen only once. If the panic threshold is exceeded after that, NTP will exit with a message to the system log.

You can use the detailed diagnosis of this measure to know the client's time stamp, the server's time stamp, and the offset.</td></tr>
<tr><td>Client time relative to server time</td><td>Indicates whether the client is behind / ahead of the server in terms of time.</td><td></td><td>If the client's clock is running faster than the server's, this measure will report the value Ahead. If the client's clock is running slower than the server's, this measure will report the value Behind.

The numeric values that correspond to these measure values are as follows:

<table><tr><th>Measure Value</th><th>Numeric Value</th></tr><tr><td>Ahead</td><td>1</td></tr><tr><td>Behind</td><td>0</td></tr></table>
Note:

By default, this measure reports the Measure Values listed above to indicate whether client is ahead or behind the server. In the graph of this measure however, the same represented using the numeric equivalents.</td></tr>
</table>

About eG Innovations

eG Innovations provides intelligent performance management solutions that automate and dramatically accelerate the discovery, diagnosis, and resolution of IT performance issues in on-premises, cloud and hybrid environments. Where traditional monitoring tools often fail to provide insight into the performance drivers of business services and user experience, eG Innovations provides total performance visibility across every layer and every tier of the IT infrastructure that supports the business service chain. From desktops to applications, from servers to network and storage, from virtualization to cloud, eG Innovations helps companies proactively discover, instantly diagnose, and rapidly resolve even the most challenging performance and user experience issues.

eG Innovations is dedicated to helping businesses across the globe transform IT service delivery into a competitive advantage and a center for productivity, growth and profit. Many of the world's largest businesses use eG Enterprise to enhance IT service performance, increase operational efficiency, ensure IT effectiveness and deliver on the ROI promise of transformational IT investments across physical, virtual and cloud environments.

To learn more visit www.eginnovations.com.

Contact Us

For support queries, email support@eginnovations.com.

To contact eG Innovations sales team, email sales@eginnovations.com.