ManageEngine
EventLog Analyzer

Disaster Recovery and High Availability Configuration Guide

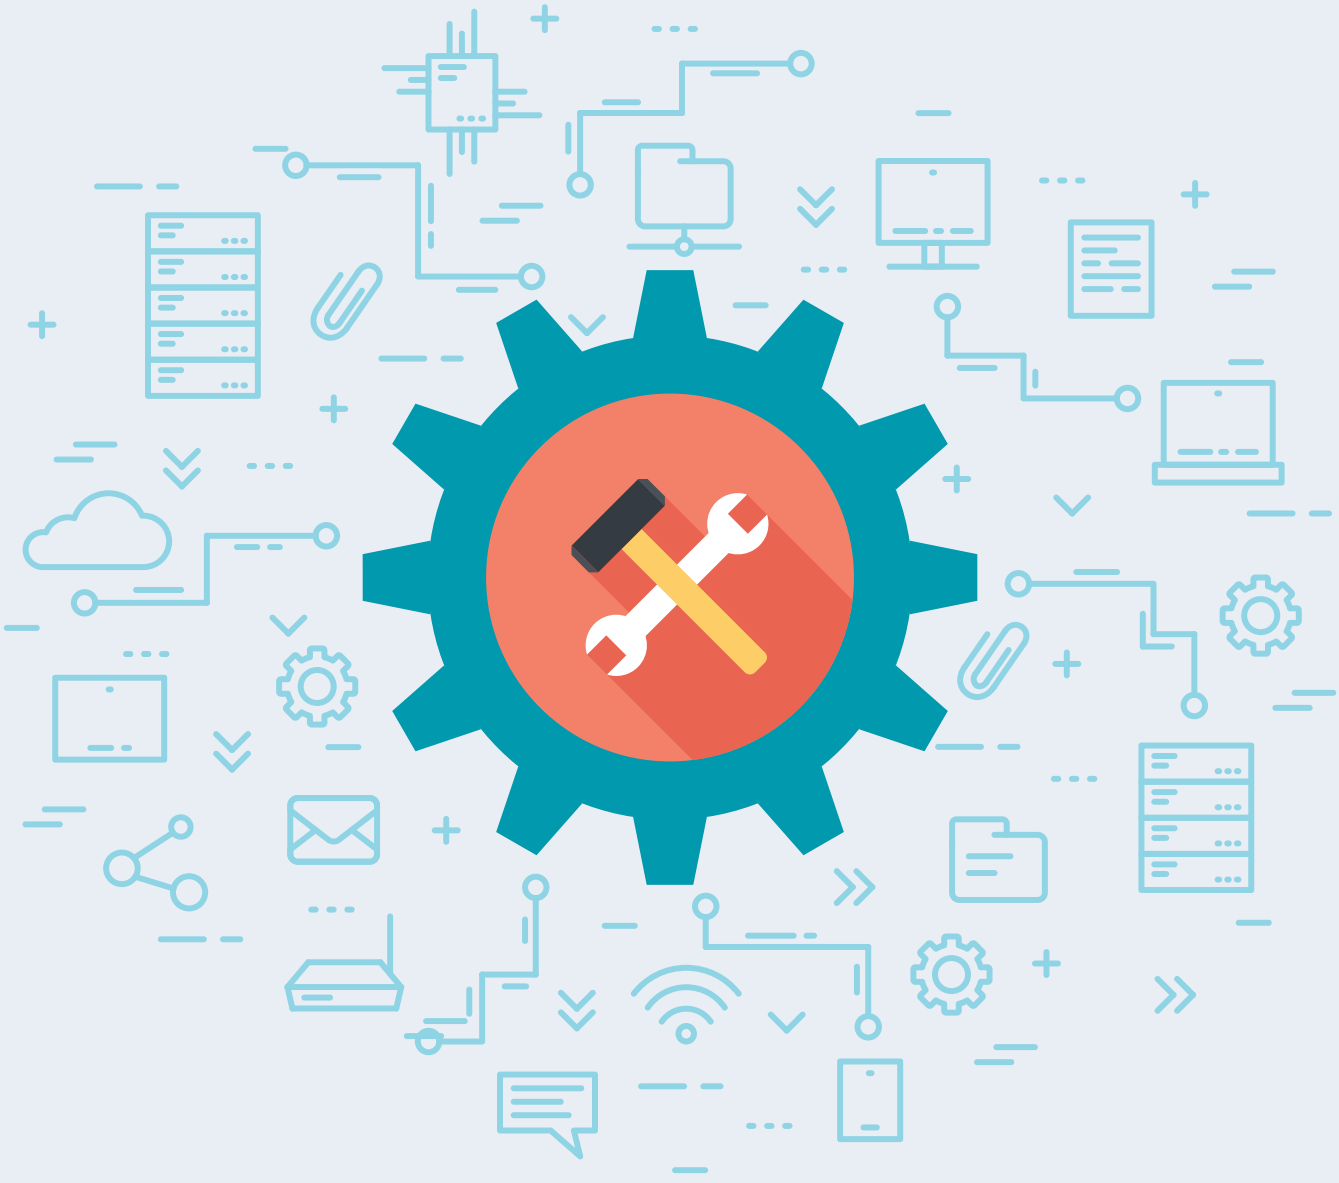

www.eventloganalyzer.com

Table of Contents

Purpose of the document

This document highlights the provisions for disaster recovery in EventLog Analyzer. It also illustrates the working and benefits of the high availability feature in the product.

Disaster Recovery for EventLog Analyzer

Log data from all possible log sources is collected and stored in the EventLog Analyzer server. This data is analyzed to detect anomalies and network security threats. Hence, the EventLog Analyzer server is a critical component from the perspective of an organization's network security. In the unlikely event of a major glitch in your environment which causes the EventLog Analyzer server to go down, log processing and analysis would come to a halt. This stoppage might turn out to be a gateway for security breaches. To avert such disasters, EventLog Analyzer has a backup mechanism.

As a disaster recovery measure, EventLog Analyzer offers the high availability feature. It allows for every EventLog Analyzer server to be configured with a standby server. This standby server would continuously monitor the primary server. In case the primary server fails, the standby server would immediately step in and start performing all the duties of the primary one without any lapse. Read more about the working of EventLog Analyzer's high availability module in the upcoming sections.

High availability in EventLog Analyzer

To configure high availability, the below mentioned procedure needs to performed on each installation of EventLog Analyzer.

Why it is necessary to ensure high availability of EventLog Analyzer?

Being a network security solution, EventLog Analyzer constantly monitors log data, looks for anomalies and attack patterns, validates threats, and helps in combating security attacks.

If the EventLog Analyzer server goes down, it would result in stoppage of log data collection and analysis. This could cause failure in identifying security incidents and in turn result in serious data breaches. Such breaches can cause not just huge financial losses and non-compliance penalties but also loss of credibility and reputation. Hence it's advisable to ensure high availability of EventLog Analyzer and thereby keep it running all the time.

Working of High Availability in EventLog Analyzer

EventLog Analyzer's high availability setup includes two separate installations. One of them acts as a primary server while the other acts as a standby server. Both the installations would point to the same database. And the archived log data and ES data will be available in the common network share.

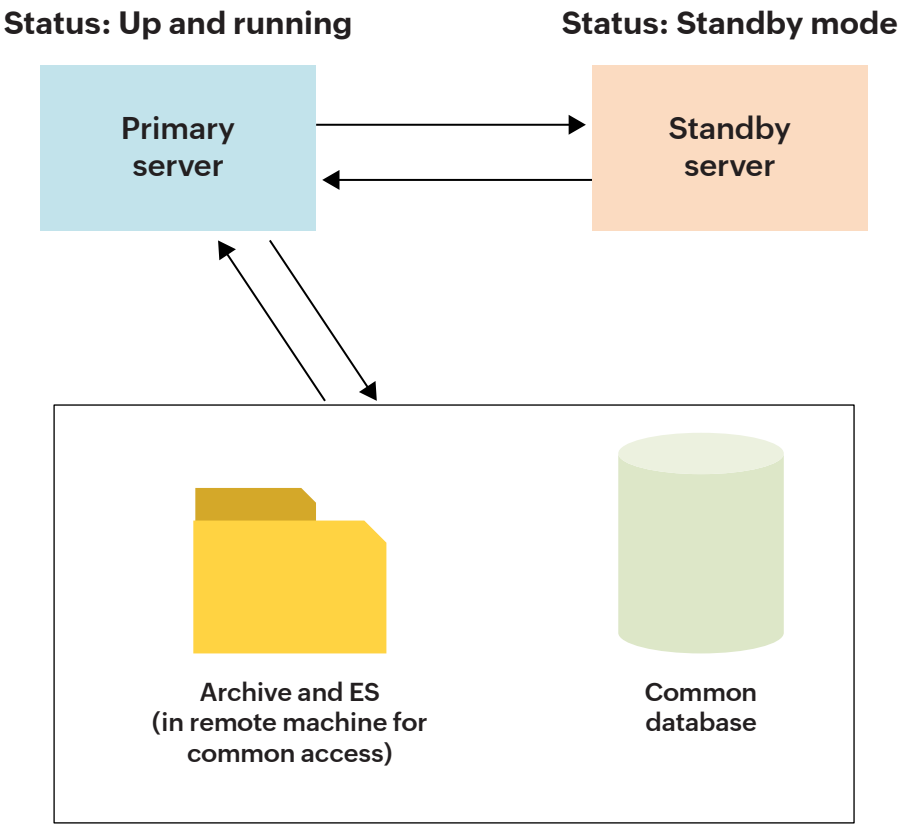

By default, the primary server will deliver all the required services. The standby server will also be started but it will remain in the standby mode. But it will continuously keep monitoring the primary server's status. Whenever the primary server fails, the standby server will kick in and take up the role of the primary server. It will start collecting the logs to prevent any data loss and continue to perform all the functions of the primary server until the actual primary server is brought back into service.

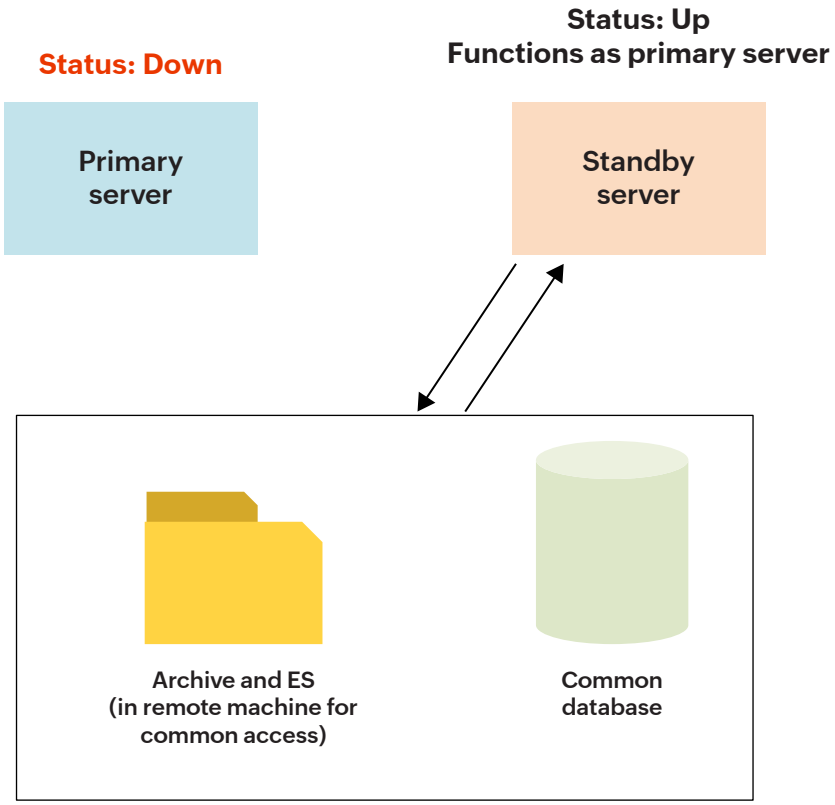

Prerequisites

Before configuring EventLog Analyzer for high availability, make sure you have two static IP addresses and one virtual IP address. All IP addresses should be in the same subnet. Both the primary and the standby servers should be of the same version.

Follow step 'a' if you are installing the product for the first time, step 'b' if you own the bundled edition, or step 'c' if you have the distributed edition of EventLog Analyzer.

a. If you are going to install the product for the **first time**, download EventLog Analyzer using this Standalone EventLog Analyzer link and then install it on two separate servers. It is recommended to use Standalone EventLog Analyzer for High Availability (HA) setup. After this, proceed to the section, "Steps to configure high availability."

Note:
Both the primary and standby servers should be on the same network, must run the same versions of the product, and should use the same port and protocol.

b. If you have already installed **EventLog Analyzer along with Log360**, please follow the following steps to convert the bundled EventLog Analyzer into Standalone.

i. To configure HA, EventLog Analyzer should use Standalone Elastic Search (ES).

ii. To check this, login to Log360 console and navigate to **Admin → Administration → Search Engine Management**

iii. If both EventLog Analyzer and Log360 are listed on the Search Engine Management tab, your installation is using standalone ES. If Log360 alone is listed, your installation is using Log360 common ES.

iv. In order to convert your installation to standalone ES, you will need to remove EventLog Analyzer from Log360.

To remove EventLog Analyzer from Log360: In Log360 console, go to **Admin → Administration → Log360 Integration → EventLog Analyzer and click Remove.**

Note:
The process might take some time to complete. Don't exit or refresh the page. After configuring high availability for EventLog Analyzer, you can integrate EventLog Analyzer back with Log360 using the virtual IP address (step 12 of the next section).

v. Post removal of EventLog Analyzer from Log360, proceed to the section "Steps to configure high availability."

c. If you are going to implement high availability for the **distributed edition of EventLog Analyzer,** perform the steps under the section "Steps to configure high availability" on every managed server, and add their virtual IP addresses to the admin server.

Note:
High availability cannot be configured in Linux installations.

Steps to configure high availability

Configuring high availability in EventLog Analyzer is simple. You can follow the following steps to do this:

1. After installation, change one of the EventLog Analyzer server's database to SQL by executing the **changeDBserver.bat** file located in <EventLog_Analyzer Home>\tools. In the dialog box that appears, enter the required details and save.

2. Now run the same changeDBserver.bat --config in the other server and point to the same database. Also, ensure that the first server is down while executing the changeDBserver.bat file on the second server.

3. Please note that both the primary and standby servers should have static IP addresses.
   To configure static IP address,
   a. Navigate to **Start > Control Panel > Network Sharing Center >Ethernet (Local Area Connection).**
   b. Select **Properties** menu.
   c. Now, uncheck **Internet Protocol Version 6 (TCP/IPv6).**
   d. Select **Internet Protocol Version 4 (TCP/IPv4)** and click on **Properties.**
   e. Select **Use the following IP address** radio button.
   f. Enter a static IP address and the subnet mask.
   g. Finally, click **OK** to save the configuration.

   The same steps mentioned above need to be followed in the standby server to configure static IP address

4. Now add the below entry in **wrapper.conf** file located in <EventLog_Analyzer Home>\server\conf.
   In the primary server, include the lines below
   a. wrapper.java.additional.x+1=-DremoteIp=<Standby Server IP>
   b. wrapper.java.additional.x+2=-DlocalIp=<Primary Server IP>
   c. wrapper.java.additional.x+3=-DvirtualIp=<Virtual IP>

   In the standby server, add the lines below:
   a. wrapper.java.additional.x+1=-DremoteIp=<Primary Server IP>
   b. wrapper.java.additional.x+2=-DlocalIp=<Standby Server IP>
   c. wrapper.java.additional.x+3=-DvirtualIp=<Virtual IP>
   d. wrapper.java.additional.x+4=-DSecondary=true

Note:
Both the primary and standby servers should be configured with the same virtual IP address

The value of x varies depending on the setup in your organization. To find the value of x that you need to enter,

- Navigate to \server\conf\wrapper.conf and search for "**wrapper.java.additional.**". Navigate to the last occurrence of the search result and note down the numerical value that is next to "wrapper.java.additonal.". It is your value for x.

- Add the commands for primary and secondary servers based on this value of x.

For example, let us consider the last occurrence of searching for "wrapper.java.additional." to be "wrapper.java.additional.36". In this scenario, your value for x is 36 and the lines you would need to add in the **primary server** would be:

- wrapper.java.additional.37=-DremoteIp=123.456.789.123
- wrapper.java.additional.38=-DlocalIp=123.456.789.124
- wrapper.java.additional.39=-DvirtualIp=123.456.789.125

The lines to be added in the **standby server** are:

- wrapper.java.additional.37=-DremoteIp=123.456.789.124
- wrapper.java.additional.38=-DlocalIp=123.456.789.123
- wrapper.java.additional.39=-DvirtualIp=123.456.789.125
- wrapper.java.additional.40=-DSecondary=true

Also ensure that,

- The virtual IP address is in the local network IP range and same subnet. Using this IP address, the high availability script will automatically add or remove the virtual IP during the product startup and shutdown.

- EventLog Analyzer processes are bound to the virtual IP. In case of syslog monitoring, the syslog devices should be configured to forward their log data to this virtual IP address.

Note:

- Please apply the required license for the primary server. For instructions, refer here.

- To update the license in the secondary server, kindly copy the files mentioned below from the primary server to the secondary server on the same location:
  1. <EventLog Analyzer Home>\lib\petinfo.dat
  2. <EventLog Analyzer Home>\lib\AdventNetLicense.xml

5. Now, in both the primary and standby servers, edit and update the interface name (*interfaceName field*) and virtual IP netmask (*VirtualIPNetMask field*) in the **StartHA.vbs** and **StopHA.vbs** files located in \tools directory. The value of the interfaceName field should be of the connection name found in your **Network Sharing Center**. The *VirtualIPNetMask* field should be filled with the subnet mask of the virtual IP.

6. Edit the path data in the elasticsearch.yml file in <EventLogAnalyzer_Home/ES/ Config> folder to install the product as a service. The value of the path.data field should be that of the common shared location, so that it can store logs of both primary and standby servers in ES data. Add the following property if it is not already present in the elasticsearch.yml file. If it is already present, just modify the property.

    a. node.max_local_storage_nodes: 2

7. Before starting EventLog Analyzer, ensure that it is installed as a service. If it is not installed as a service, execute the **service.bat** –I command from <EventLog Analyzer_Home>\bin directory to install the product as a service.

8. Start the primary server from the Windows Services console.

    Note:

    Please use only an administrator credential to start EventLog Analyzer service in both primary and standby servers.

9. Now in EventLog Analyzer, navigate to **Settings → Admin Settings → Data Storage → Archives → Settings** and change the Archive Location of archive log data to the common shared folder by providing its exact UNC path

10. If you want to change report export storage location to a common shared folder, follow the below steps.

    a. Navigate to **Settings → Admin Settings → General → Product Settings.**

    b. In the Product Configurations tab, specify the option in Reporting Mode Field as **Save to Location** or **Send Mail & Save to Location**. After selecting, click on the settings icon and provide the common shared folder location in the UNC path box.

    c. This will change the location of exported reports to the common shared folder.

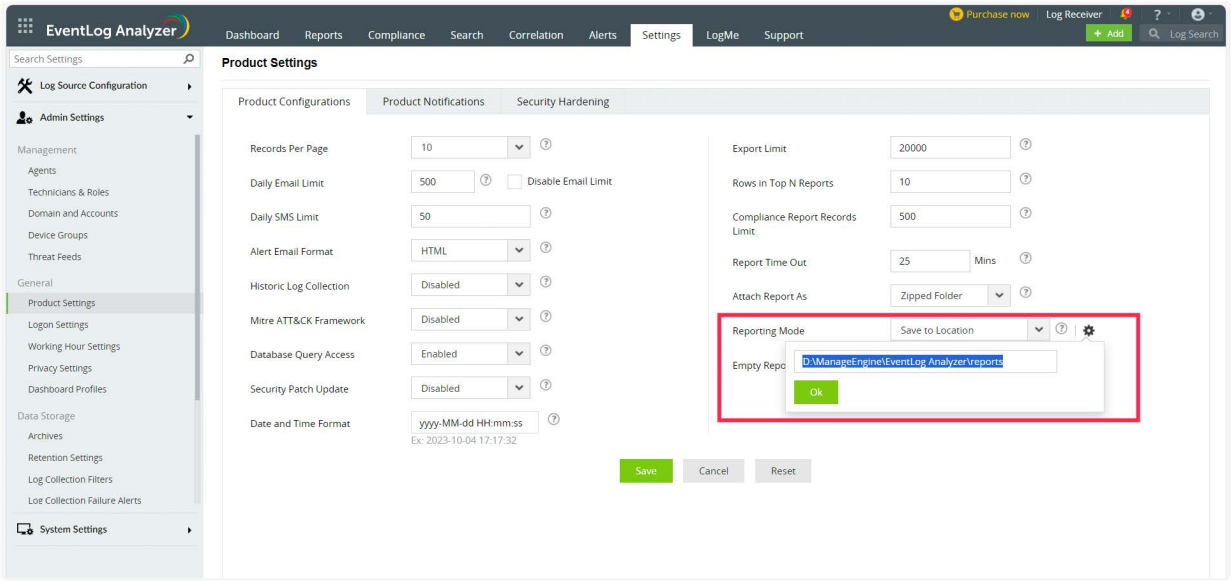

11. Email notification will be sent to the product users who have administrator privileges. To configure or change the email address of admin user, navigate to **Settings → Admin Settings → Management → Technicians and Roles**. This will display the product's technicians and their corresponding roles. Click on the edit icon for the admin user and you will be prompted with the 'Update technician details' dialog box, where you can edit the email address of the admin user.

12. After configuring high availability, follow these steps to integrate EventLog Analyzer with Log360.

    a. Go to **Admin → Administration → Log360 integration**. Click on modify button for EventLog Analyzer component.

    b. Enter the **virtual IP address** and the **port** number of the HA configured EventLog Analyzer server.

    c. Select the connection **Protocol** from the drop-down menu. Click **Integrate Now.**

Steps to activate standby server automatically

- Try to start the EventLog Analyzer service in the standby server while the primary server is up and running. The service startup will fail but this would trigger a process called **wscript.exe** that will start monitoring the primary server's availability.
- Once the primary server goes down, the standby server will automatically get initiated and also email notifications will be sent to administrators immediately.
- Troubleshoot the primary server when it goes down. Upon finishing troubleshooting, shutdown the standby server manually and then start the primary server.
- When the primary server is up and running, perform step 1 to initiate the script in the standby server.

Steps to upgrade EventLog Analyzer to the latest build

Prerequisites:

- Before proceeding with the upgrade process, please ensure that there is enough disk space on the EventLog Analyzer server.
- Please make a copy of the entire EventLog folder or take a snapshot of the server to have a backup, in case the upgrade fails.
- If you use Microsoft SQL database, we request you to take a snapshot of your database.
- Do not interrupt or cancel the upgrade process. If the upgrade fails, please contact support.
- It is important to start the product after each successful upgrade. Please follow the steps again to perform another upgrade.

Follow these steps to upgrade EventLog Analyzer

1. Before upgrading, make sure to stop both the primary and standby servers.
2. Apply the service pack on the primary server. Once the upgrade is completed, start the primary server. Ensure the primary server is working fine after the upgrade.

3. Stop the primary server and then apply the service pack on the standby server. Once the upgrade is completed, start the standby server. Make sure it's working fine and then stop the standby server.

4. Now both the servers will be upgraded to the latest version. You can start the primary server and then the standby server respectively.

Steps to restore high availability database

1. Primary server backup can be restored on primary server and vice verse.

2. If primary server backup needs to be restored in secondary server or vice-verse, move the following files under primary server conf folder to the secondary server conf folder or vice-versa,
   - database_params.conf
   - customer-config.conf
   - ela_key.key

Steps for migrating the high availability server

Prerequisites

1. Verify that both the primary and standby servers are reachable via hostnames or IP addresses.

2. Confirm that the existing Microsoft SQL Server instance is accessible from both the primary and standby servers.

3. Ensure that the primary and standby servers have sufficient permissions to access the archive and Elasticsearch locations.

4. Note that both the primary and standby servers must have static IP addresses.

Migrating the primary or standby server to a different server

1. Copy the entire folder where EventLog Analyzer is installed from the old machine to the new machine on which the server migration has to be done.

   Note: Ensure that EventLog Analyzer is not running on either the primary or standby server.

2. Update the following entries in the wrapper.conf file located in *<EventLog_Analyzer Home>\server\conf:*

i. On the primary server, edit the commands related to the configurations as mentioned below:

- Update the IP address of the current standby server in *wrapper.java.additional.x+1=-DremoteIp=<Standby Server IP address>.*
- Update the IP address of the current primary server in *wrapper.java.additional.x+2=-DlocalIp=<Primary Server IP address>.*
- Update the current virtual IP address in *wrapper.java.additional.x+3=-DvirtualIp=<Virtual IP address>.*

ii. On the standby server, edit the commands related to the configurations as mentioned below:

- Update the IP address of the current primary server in *wrapper.java.additional.x+1=-DremoteIp=<Primary Server IP address>.*
- Update the IP address of the current standby server in *wrapper.java.additional.x+2=-DlocalIp=<Standby Server IP address>.*
- Update the current virtual IP address in *wrapper.java.additional.x+3=-DvirtualIp=<Virtual IP address>.*
- Make sure that the command in *wrapper.java.additional.x+4=-DStandby=true is always true.*

3. After updating the commands, ensure that the high availability configuration steps outlined in the documentation are followed properly as mentioned

Migrating the primary or standby server to a different directory on the same machine

- **Case 1:** If Elasticsearch, the archive, and the database are kept in the default locations, there is no need to make any modifications to the primary and standby EventLog Analyzer servers.

- **Case 2:** If Elasticsearch, the archive, and the database are migrated to a different directory or server, please follow the steps given below, respectively:

Elasticsearch migration

1. Ensure that the Elasticsearch data is stored in a shared network location that is accessible by both the primary and standby servers.

2. Make sure to change the path.data and path.repo with new paths in the Elasticsearch.yml file in both instances: [Path : *<ELA_HOME>\ES\config\]*. The value of the path.data field should be that of the common shared location so that it can store the logs of both the primary and standby servers in Elasticsearch data.

3. Add the property `*node.max_local_storage_nodes:* 2` if it is not already present in the *Elasticsearch.yml file.*

Archive migration

Follow the steps mentioned here to migrate the archive data.

Database migration

Follow the steps mentioned here to migrate the database.

For any further clarifications and queries, contact **eventlog-support@manageengine.com**.