

EventLog Analyzer:

# BEST PRACTICES GUIDE

## Table of Contents

<b>System requirements</b>	2
Hardware Requirements	2
General Recommendations:	3
<b>Optimizing hard disk space</b>	5
Required hard disk space	5
Manage database size	5
Manage archive size	5
<b>Securing EventLog Analyzer</b>	6
Installation configuration	6
User configuration	6
SSL certification	6
<b>Database best practices</b>	6
Secure database	6
Optimize PostgreSQL database performance	6
Optimize MySQL database performance	7
Back up database	8
<b>Optimizing log search performance</b>	8
Heap allocation	8
Distributing indexing load	11
<b>Support best practices</b>	12
Create Support Information File	12

This guide details best practices which, if followed, ensure smooth operation and optimum performance of EventLog Analyzer.

## System requirements

### Hardware Requirements

Log management solutions are resource-intensive and selecting the right hardware plays a major role in ensuring optimal performance.

The following table denotes the suggested hardware requirements based on the type of flow.

	Low Flow	Normal Flow	High Flow
<b>Processor cores</b>	6	12	24
<b>RAM</b>	16 GB	32 GB	64 GB
<b>IOPS</b>	150	750	1500 *
<b>Disk space</b>	1.2 TB	3 TB *	4 TB *
<b>Network card capacity</b>	1 GB/s	1 GB/s	10 GB/s
<b>CPU Architecture</b>	64-bit	64-bit	64-bit

**Note:**

- The above-mentioned **values are approximate**. It is recommended to run a test environment similar to the production environment with the setup details mentioned in the above table. Based on the exact flow and data size, the system requirements can be fine-tuned.
- For higher IOPS, we can use RAID or SSD.

Use the following table to determine the type of flow for your instance.

Log type	Size (in Bytes)	Category	Log Units		
			Low Flow (EPS)	Normal Flow (EPS)	High Flow (EPS)
Windows	900	Windows	300	1500	3000
Linux, HP, pfSense, Juniper	150	Type 1 Syslogs	2000	10000	20000
Cisco, Sonicwall, Huawei, Netscreen, Meraki, H3C	300	Type 2 Syslogs	1500	6000	12000
Barracuda, Fortinet, Checkpoint	450	Type 3 Syslogs	1200	4000	7000
Palo Alto, Sophos, F5, Firepower, and other syslogs	600	Type 4 Syslogs	800	2500	5000

**Note:**

- A single-installation server can handle either a maximum of 3000 Windows logs or any of the high flow values mentioned for each log type in the above table.
- For log types which are not mentioned in the above table, choose the appropriate category based on the log size. For example, in the case of SQL Server logs when the byte size is 900 bytes, and EPS is 3000, it should be considered as High Flow.
- If the combined flow is higher than what a single node can handle, it is recommended to implement **distributed setup**.
- It is recommended to choose the next higher band if advanced threat analytics and a large number of correlation rules have been used.

## General Recommendations

### VM infrastructure

- Allocate 100 percent RAM/CPU to the virtual machine running EventLog Analyzer. Sharing memory /CPU with other virtual machines on the same host may result in RAM/CPU starvation and may negatively impact EventLog Analyzer's performance.
- Employ thick provisioning, as thin provisioning increases I/O latency. In case of VMware, Select Thick provisioned, eagerly zeroed as lazily zeroed is lower in performance.
- Enabling VM snapshots is not recommended as the host duplicates data in multiple blocks by increasing reads and writes, resulting in increased IO latency and degraded performance.

### CPU & RAM

- Server CPU utilization should always be maintained below 85% to ensure optimal performance.
- 50% of server RAM should be kept free for off-heap utilization of Elasticsearch for optimal performance.

### Disk

- Disk latency greatly affects the performance of EventLog Analyzer. Direct-attached storage (DAS) is recommended on par with the throughput of an SSD with near-zero latency and high throughput. An enterprise storage area network (SAN) can be faster than SSD.
- Currently only local and remote (NAS) drives are supported by EventLog Analyzer for storing live search index and archive data.

**Additional note:** Search indices require fast random access to the index files, which is not possible with blob storage-type data stores such as S3 and Azure Blob store.

## Web browsers

EventLog Analyzer has been tested to support the following browsers and versions with at least a 1024x768 display resolution:

- Microsoft Edge
- Firefox 4 and later
- Chrome 8 and later

## Databases

EventLog Analyzer can use the following databases as its back-end database.

### Bundled with the product

- PostgreSQL

### External databases

- Microsoft SQL 2012 & above

Please note the hardware requirements needed to configure the MS SQL database for EventLog Analyzer:

RAM	CPU	IOPS	Disk space
8GB	6	300-500	300-500 GB

## Operating systems

EventLog Analyzer can be installed in machines running the following operating systems and versions:

- Windows 7 & above, and Windows Server 2008 & above
- Linux: Red Hat 8.0 and above/all versions of RHEL, Mandrake/Mandriva, SUSE, Fedora, CentOS, Ubuntu, Debian

## Installation server

- SIEM solutions are resource-intensive. It is recommended to provide a dedicated server for their optimal performance.
- Eventlog Analyzer uses Elasticsearch. Elasticsearch process is expected to utilize off-heap memory for better performance. Off-heap memory is maintained by the operating system and will free up when necessary.

## Additional Elasticsearch Node Recommendations

Hardware	Minimum	Recommended
Base Speed	2.4 GHz	3 GHz
Core	12	16
RAM	64	64
Disk Space	1.2 TB	1.5 TB
IOPS	1500*	1500*

## Optimizing hard disk space

The two main contributing factors to hard disk space are the database and archive files. The database (or index) files contain the most recent log data which can be reported on and searched, while the archive files contain the older, historic log data. Archive files need to be loaded into the product first before they can be searched or reported on.

### Required hard disk space

The hard disk space required to store logs can be calculated by using the procedure detailed in the [performance optimization guide](#) in the EventLog Analyzer website

### Manage database size

Log data is stored in the database and is periodically compressed and stored among the archive files. The longer the retention period in the database, the greater is the hard disk space needed and lower the database performance. The default retention period is 32 days and is configurable (Settings > Admin settings > DB retention settings). Minimize this value to obtain optimum performance.

### Manage archive size

The archive files are retained for a specific period before being deleted permanently. As they can even be stored forever, the size of the archive folder could grow indefinitely. The archive retention period is forever and is configurable (Settings > Configuration settings > View archived files > Settings with Settings > Admin settings > View archived files > Settings). The archive folder size can also be managed by assigning a separate dedicated drive as the archive location, or manually transferring the contents to a tape drive or high capacity storage drive periodically.

## Securing EventLog Analyzer

### Installation configuration

The operating system user account used to install and run the product must be the same and must have permissions on all installed folders and subfolders. While it is not necessary for the root account to be used on a Linux system, on a Windows system, only the default administrator account must be used.

### User configuration

It is best to change the default passwords for the admin and guest user accounts in the EventLog Analyzer web client (Settings > Admin settings > Manage technician)

### SSL certification

EventLog Analyzer server-client communication can be secured using the SSL (Secure Sockets Layer) protocol. The SSL certification guide offers detailed steps on how to obtain SSL certification.

## Database best practices

### Secure database

For smooth and seamless installation, EventLog Analyzer makes use of the MySQL or PostgreSQL database default root/postgres user without password. It is recommended to assign a password to this account in order to further secure the database.

This is not needed in case of MS SQL, as a valid user account with credentials needs to be provided during installation itself.

### Optimize PostgreSQL database performance

To optimize performance of the PostgreSQL database:

- Stop EventLog Analyzer.
- Navigate to <EventLog Analyzer home>/pgsql/data/directory.
- Open the file postgres\_ext.txt.
- Replace the existing values of the parameters, with the values mentioned below.
- Save and restart EventLog Analyzer.

Parameter	Comment
shared_buffers=128 MB	Minimum requirement is 128 KB.
work_mem=12 MB	Minimum requirement is 64 KB.
maintenance_work_mem=100 MB	Minimum requirement is 1 MB.
checkpoint_segments=15	Logfile segments minimum 1 and 16 MB each.
checkpoint_timeout=11 minutes	Range: 30 seconds to 1 hour.
checkpoint_completion_target=0.9	checkpoint target duration is 0.0 - 1.0.
seq_page_cost=1.0	This parameter is measured in an arbitrary scale.
random_page_cost=2.0	This parameter is measured in same scale as above.
effective_cache_size=512MB	
synchronous_commit=off	

## Optimize MySQL database performance

To optimize performance of the MySQL database:

- Stop EventLog Analyzer.
- Navigate to <EventLog Analyzer home>/bin.
- Open the file startDB.bat (startDB.sh in case of a Linux machine).
- Replace the existing value of the parameter "--innodb\_buffer\_pool\_size", with a value suited to the RAM size of the machine, as given in the table below. For example, if the RAM size is 8 GB, the parameter should be "--innodb\_buffer\_pool\_size=3000M".
- Save and restart EventLog Analyzer.

RAM Size	Value
1 GB	Default value (no need to replace)
2 GB	1200M
3 GB	1500M
4 GB	1500M
8 GB	3000M
16 GB	3000M



## Back up database

It is recommended to back up the EventLog Analyzer database every fortnight, so that data is not lost in case of any disaster. The database files are located in the <EventLog Analyzer home>/mysql or <EventLog Analyzer home>/pgsql folder, as applicable to the build number. To back up the data, stop the EventLog Analyzer service, and take a copy of all files and folders in the location. This can be done manually or using any third party back up software. The procedure to back up MS SQL database data can be found in this link. It is also advisable to keep a backup of the archive files, found in <EventLog Analyzer>/archive. If restoring data from a backup, ensure that the build number of the product is the same as when the backup was taken.

## Optimizing log search performance

### 1. Give sufficient heap to the Elasticsearch

To ensure fair performance maintain the heap to data ratio of 1:60. This means that you can allocate approximately 1GB of memory (heap) for every 60GB of data in the Elasticsearch node (the maximum ratio). But for better performance, you can lower this ratio (i.e., 1:30 is better than 1:60) and increase speed.

**Note:** In the older Build 12320, heap to data ratio is 1:30.

Elasticsearch also uses file-system cache to provide faster searches. It is recommended to have enough free space on your RAM equivalent to that of the heap memory allocated for Elasticsearch. If this is not feasible, then ensure at least 30% of the server's RAM is free. OS will use this free RAM to cache the Elasticsearch's indices to provide better performance.

**Note:** Heap allocated to Elasticsearch shouldn't exceed 32GB.

#### Example:

Suppose we have **100GB of search data**,

then the heap size for Elasticsearch should be at least  $\rightarrow 100/30 \sim 4\text{GB}$

Insufficient heap is the underlying reason for several performance issues such as:

- Slow log processing / indexing performance
- Cached record
- Delayed search results
- Failed searches

### Find out the total size of data stored in Elasticsearch

Elasticsearch can run in either shared common (<ManageEngine>/elasticsearch/ES) or local instance (<EventlogAnalyzer>/ES)

### Steps to determine the Elasticsearch location (ES directory):

- When EventLog Analyzer is installed as a standalone application (i.e., running without Log360) the local ES will be in use, located in the <EventlogAnalyzer>\ES directory.
- If EventLog Analyzer has been installed along with Log360, the default Elasticsearch configuration (common ES) will be in use, located in <ManageEngine>\elasticsearch\ES directory.

### Steps to check the Elasticsearch data size:

1. Navigate to the <ES directory>\config.
2. Open the **elasticsearch.yml** file in the config folder.
3. Look for **path.data** setting in this file.

Navigate to the data folder specified in the **path.data** setting and check the size of the folder.

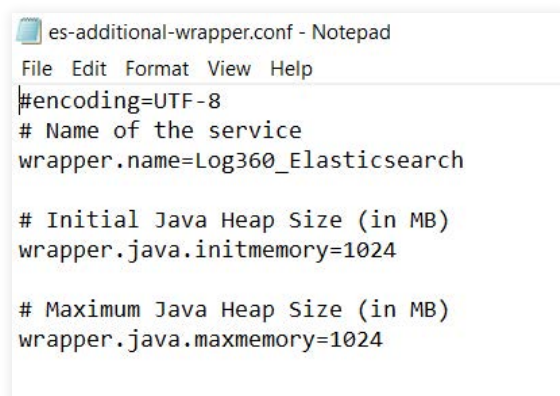
To ensure optimal performance, monitor and maintain your Elasticsearch data regularly and limit the size of single ES node between 1.5TB-1.9TB.

**Note:** In older Build 12320, ES can hold 800GB-1.2TB of data.

### Steps to adjust heap (memory):

1. Navigate to the ES directory depending on whether it is a standalone build or a bundled build (with Log360).
2. Navigate to /ES/config.
3. Open the configuration file → **es-additional-wrapper.conf** and view the heap size.

**Note:** Make sure the logged in user has permissions to write.



```

es-additional-wrapper.conf - Notepad
File Edit Format View Help
#encoding=UTF-8
# Name of the service
wrapper.name=Log360_Elasticsearch

# Initial Java Heap Size (in MB)
wrapper.java.initmemory=1024

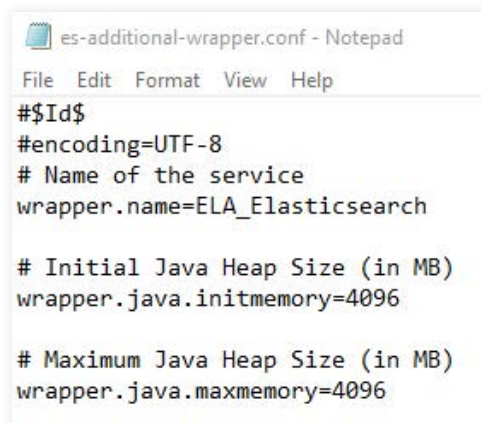
# Maximum Java Heap Size (in MB)
wrapper.java.maxmemory=1024

```

4. Heap size is written in MB.
  - a. wrapper.java.initmemory & wrapper.java.maxmemory both needs to be set to the same value. Here it's set to 1024, i.e., Elasticsearch's memory is set to (1024 MB/1024) = 1GB.
  - b. If it has to be increased to 25 GB, then we need to set both the values to 25 \* 1024 = 25600

### Steps to increase ES heap size:

1. Open the configuration file → **es-additional-wrapper.conf**.
2. Edit **wrapper.java.initmemory** and **wrapper.java.maxmemory** values to increase the heap size.
3. Make sure both the values of **wrapper.java.initmemory** and **wrapper.java.maxmemory** are the same. Otherwise the product will not start up properly.
4. If the product is running, stop Elasticsearch by going to **ES/bin** and run **stopES.bat** using the admin command prompt or just restart EventLog Analyzer. This will restart Elasticsearch with the new heap.
5. If **<ManageEngine>\elasticsearch\ES** heap was updated, you need to manually run the **stopES.bat** command from **<ManageEngine>\elasticsearch\ES\bin** before restarting EventLog Analyzer.



```

es-additional-wrapper.conf - Notepad
File Edit Format View Help
#$Id$
#encoding=UTF-8
# Name of the service
wrapper.name=ELA_Elasticsearch

# Initial Java Heap Size (in MB)
wrapper.java.initmemory=4096

# Maximum Java Heap Size (in MB)
wrapper.java.maxmemory=4096

```

### Note:

- In the event of **OutOfMemory** and **LowMemory** errors, the Elasticsearch heap will automatically expand up to one-third of the available RAM on the machine.
- It's important to note that increasing the heap size isn't always the solution to improve performance. Apart from heap, other factors like Disk and CPU may also cause performance problems. Ensure that the [System Requirements](#) are met.
- It is also important to monitor the memory usage regularly to ensure that the system is performing efficiently and to adjust the settings if necessary.
- Keep in mind that increasing the Elasticsearch heap size should be done with careful consideration of the available resources on your machine.

### Ensure that disk is not the bottleneck:

If the server is generating cached records (i.e., log processing is slow) or if the searches are slow, then you can:

- a. Use faster storage as mentioned in the [System Requirements](#) page.
- b. Check if the disk where the data is stored is not fragmented.
- c. In **Windows Resource Monitor**, you can check the **Disk** tab. If the **Disk Activity** shows the **Highest Active Time** to be always 100%, it indicates that the disk might have issues or is not fast enough.

## 2. Use additional search nodes to distribute search/indexing load for better performance

You can use the [Search Engine Management](#) feature present in Log360 (Log360 → Admin → SearchEngineManagement) to add additional Elasticsearch nodes to distribute search and indexing load using extra machines.

- a. If the data size is too big for a single node, it's better to add additional nodes to distribute the search/indexing load.
- b. If the search performance is not good enough, then add additional nodes.
- c. More nodes help in processing logs faster.

### Best practice for search:

- In Eventlog Analyzer, the **retention period is 32 days by default** (can be increased in **Settings → DB Settings** in UI). If it is updated to 90 days, then 32 days of data will be stored as live data which can be accessed fast. The data beyond that will be stored as cold data which needs to be unarchived and loaded to search engine. Hence, searching beyond live data will take more time than usual.
- While searching the data, both the heap (memory assigned to Elasticsearch) & off-heap (free RAM on the System) are used. Free RAM on the system allows Elasticsearch to read the indices faster. Hence, it is advised to keep at least the same amount of RAM free on the server equivalent to the heap provided to Elasticsearch for better performance. If this is not feasible then ensure at least 30% of the server's RAM is free. OS will use this free RAM to cache the Elasticsearch's indices to provide better performance.

It is necessary to have a disk with good sequential and random read speed because the search process involves iterating through a lot of files, which is an IO heavy operation. SSDs should be preferred as it reduces the I/O load and I/O waits and helps to exploit the full CPU power.

## Support best practices

### Create Support Information File (SIF)

When support is required, creating a Support Information File (SIF) to send to the support team (eventloganalyzer-support@manageengine.com) would be helpful and time saving. To create a SIF from the web client, go to the Support tab of the product. Click on 'Create Support Information File', wait 30-40 seconds, and click on the Support tab again. Click on download and send the downloaded SIF to the support team, or click 'Upload to FTP Server', provide the required details and submit. If the server or web client is not working, zip the files found at <EventLog Analyzer Home>/server/default/log and upload the zip file in this FTP link.

## About EventLog Analyzer

EventLog Analyzer is a comprehensive IT compliance and log management software for SIEM. It provides detailed insights into your machine logs in the form of reports to help mitigate threats in order to achieve complete network security. <https://blogs.manageengine.com/eventloganalyzer>

## About ManageEngine

ManageEngine delivers the real-time IT management tools that empower an IT team to meet an organization's need for real-time services and support. Worldwide, more than 60,000 established and emerging enterprises — including more than 60 percent of the Fortune 500 — rely on ManageEngine products to ensure the optimal performance of their critical IT infrastructure, including networks, servers, applications, desktops and more. ManageEngine is a division of Zoho Corp. with offices worldwide, including the United States, United Kingdom, India, Japan and China.

## Our Products

AD360 | Log360 | ADAudit Plus | Exchange Reporter Plus | DataSecurity Plus | SharePoint Manager Plus