
Energy-Log-Server-2.x Documentation

Release latest

Oct 02, 2018

Contents

1	About	1
2	Introduction	3
3	Elasticsearch	5
4	Kibana	7
5	Logstash	9
6	ELK	11
7	Where does the data come from?	13
8	System services	15
9	First configuration steps	17
10	First login	19
11	Index selection	23
12	Time settings and refresh	25
13	Fields	27
14	Filtering and syntax building	29
15	Saving and deleting queries	31
16	Visualizations	35
17	Creating visualization	37
18	Vizualization types	39
19	Edit visualization and saving	41
20	Dashboards	43

21	Sharing	45
22	Reports	47
23	CSV Report	49
24	PDF Report	53
25	Scheduler Report (Schedule Export Dashboard)	55
26	Users, roles and settings	59
27	Creating a User (Create User)	61
28	User's modification and deletion, (User List)	63
29	Create, modify and delete a role (Create Role), (Role List)	65
30	Object access permissions (Objects permissions)	69
31	Default user and passwords	73
32	Settings	75
33	License (License Info)	77
34	Special accounts	79
35	Alert Module	81
36	Enabling the Alert Module	83
37	Creating Alerts	85
38	List of Alert rules	89
39	Alerts status	91
40	Type of the Alert module rules	93
41	Example of rules	95
42	Example of rules	97
42.1	Unix - Authentication Fail	97
42.2	Windows - Firewall disable or modify	97
43	Intelligence Module	99
44	Create AI Rule	101
45	The fixed part of the screen	103
46	Screen content for regressive algorithms	105
47	Screen content for the Trend algorithm	109
48	Screen content for the neural network (MLP) algorithm	113
49	AI Rules List	117

50 AI Learn	121
51 AI Learn Tasks	125
52 Scenarios of using algorithms implemented in the Intelligence module	127
53 Results of algorithms	129
54 Scheduler Module	131
55 Permission	133
56 Verification of Elasticsearch service	135
57 Verification of Logstash service	137
58 Verificatoin of Energy Log Server GUI service	139
59 Node roles	141
60 Naming convention	143
61 Config files	145
62 Example setup	147
63 Integration with AD	149
64 AD configuration	151
65 Configure SSL suport for AD authentication.	155
66 Role mapping.	163
67 Configuring Single Sign On (SSO)	165
67.1 Create an User Account for Elasticsearch auth plugin	165
67.2 Define Service Principal Name (SPN) and Create a Keytab file for it	166
67.3 Create a file named <i>krb5Login.conf</i> :	166
67.4 Append the following JVM arguments (on Elasticsearch node in <i>/etc/sysconfig/elasticsearch</i>):	167
67.5 Add the following additional (<i>sso.domain</i> , <i>service_principal_name</i> , <i>service_principal_name_password</i>) settings for <i>ldap</i> in <i>elasticsearch.yml</i> or <i>properties.yml</i> file wherever the <i>ldap</i> settings are configured:	167
67.6 To apply the changes restart Elasticsearch service	167
67.7 Enable SSO feature in <i>kibana.yml</i> file:	167
68 Client (Browser) Configuration	169
68.1 Internet Explorer configuration	169
68.2 Chrome configuration	172
68.3 Firefox configuration	172
69 Configure email delivery for sending PDF reports in Scheduler.	175
70 Configuration file for postfix mail client	177
71 Basic postfix configuration	181
72 Example of postfix configuration with SSL encryption enabled	183

73 Elasticsearch API	185
74 Elasticsearch Index API	187
75 Elasticsearch Index API - Adding Index	189
76 Elasticsearch Index API	191
77 Elasticsearch Index_API useful commands.	193
78 Elasticsearch Document API	195
79 Elasticsearch Document API - Create_Document	197
80 Elasticsearch Document API - Delete Document	199
81 Elasticsearch Document API - useful commands	201
82 Elasticsearch Cluster API	205
83 Elasticsearch Search API	207
84 Mapping, Fielddata and Templates	209
85 Logstash	211
86 Logstash - Input “beats”	213
87 Logstash - Input “network”	215
88 Logstash - Filter “beats syslog”	217
89 Logstash Filter “network”	221
90 Logstash - Filter “geoip”	225
91 Logstash - Output to Elasticsearch	227
92 Logstash plugging for “naemon beat”	229
93 Logstash plugging for “perflog”	231
94 CHANGELOG	233
95 2.1.17	235
95.1 Added	235
95.2 Changed	235
96 2.1.24	237
96.1 Added	237
96.2 Changed	237
97 2.1.26	239
97.1 Added	239
97.2 Changed	239
98 2.1.27	241
98.1 Added	241

99 2.1.28	243
99.1 Added	243
99.2 Changed	243
1002.1.29	245
100.1 Added	245
1012.1.30	247
101.1 Changed	247
1022.1.31	249
102.1 Changed	249
1032.1.32	251
103.1 Added	251
1042.1.33	253
104.1 Added	253
104.2 Changed	253
1052.1.34	255
105.1 Added	255
105.2 Changed	255
1062.1.35	257
106.1 Added	257
1072.1.37	259
107.1 Added	259
1082.2.1	261
108.1 Added	261
1092.2.2	263
109.1 Added	263
1102.2.3	265
110.1 Added	265
110.2 Changed	265
1112.3.0	267
111.1 Added	267
111.2 Changed	267
1122.3.3	269
112.1 Added	269
112.2 Changed	269
1132.3.4	271
113.1 Added	271
113.2 Changed	271
1142.3.5	273
114.1 Added	273
114.2 Changed	273

1152.3.6	275
115.1 Added	275
115.2 Changed	275
1162.3.7	277
116.1 Added	277
116.2 Changed	277
1172.3.8	279
117.1 Added	279
117.2 Changed	279
1182.3.9	281
118.1 Added	281
118.2 Changed	281
1192.4.0	283
119.1 Added	283
119.2 Changed	283
1202.4.1	285
120.1 Changed	285
1212.4.2	287
121.1 Added	287
121.2 Changed	287
1222.4.3	289
122.1 Added	289
1232.4.4	291
123.1 Added	291
123.2 Changed	291
1242.4.5	293
124.1 Added	293
124.2 Changed	293

CHAPTER 1

About



Energy Log Server User Guide

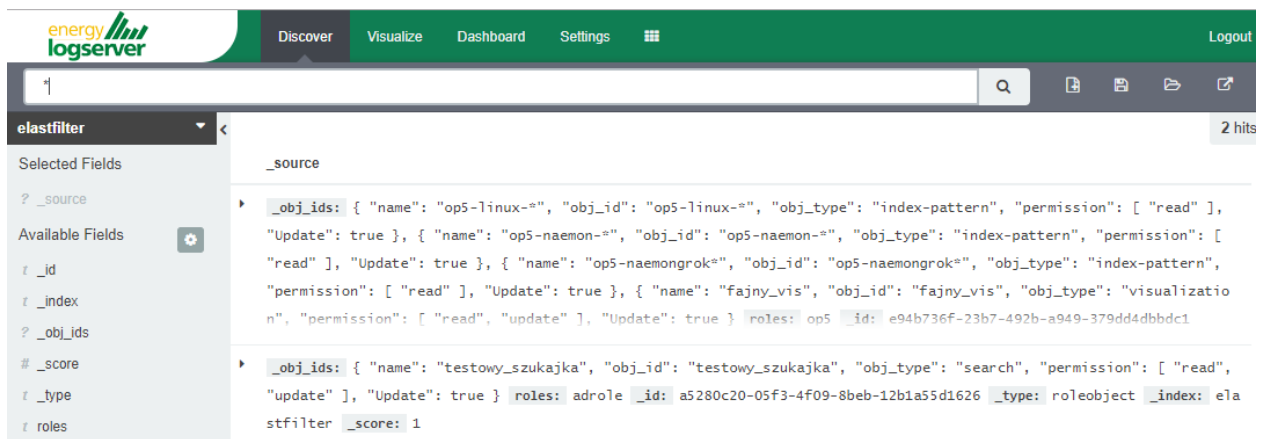
Software ver. 2.x

Document version. 0.8

CHAPTER 2

Introduction

Energy Log Server is innovation solution allowing for centralize IT systems events. It allows for an immediately review, analyze and reporting of system logs - the amount of data does not matter. Energy Log Server is a response to the huge demand for storage and analysis of the large amounts of data from IT systems. Energy Log Server is innovation solution that responds to the need of effectively processing large amounts of data coming from IT environments of today's organizations. Based on the open-source project Elasticsearch valued on the market, we have created an efficient solution with powerful data storage and searching capabilities. The System has been enriched of functionality that ensures the security of stored information, verification of users, data correlation and visualization, alerting and reporting.



Energy Log Server project was created to centralize events of all IT areas in the organization. We focused on creating a tool that functionality is most expected by IT departments. Because an effective licensing model has been applied, the solution can be implemented in the scope expected by the customer even with very large volume of data. At the same time, the innovation architecture allows for servicing a large portion of data, which cannot be dedicated to solution with limited scalability.

CHAPTER 3

Elasticsearch

Elasticsearch is a NoSQL database solution that is the heart of our system. Text information sent to the system, application and system logs are processed by Logstash filters and directed to Elasticsearch. This storage environment creates, based on the received data, their respective layout in a binary form, called a data index. The Index is kept on Elasticsearch nodes, implementing the appropriate assumptions from the configuration, such as:

- Replication index between nodes,
- Distribution index between nodes.

The Elasticsearch environment consists of nodes:

- Data node - responsible for storing documents in indexes,
- Master node - responsible for the supervisions of nodes,
- Client node - responsible for cooperation with the client.

Data, Master and Client elements are found even in the smallest Elasticsearch installations, therefore often the environment is referred to as a cluster, regardless of the number of nodes configured. Within the cluster, Elasticsearch decides which data portions are held on a specific node.

Index layout, their name, set of fields is arbitrary and depends on the form of system usage. It is common practice to put data of a similar nature to the same type of index that has a permanent first part of the name. The second part of the name often remains the date the index was created, which in practice means that the new index is created every day. This practice, however, is conventional and every index can have its own rotation convention, name convention, construction scheme and its own set of other features. As a result of passing document through the Logstash engine, each entry receives a data field, which allows to work with data in relation to time.

The Indexes are built with elementary part called shards. It is good practice to create Indexes with the number of shards that is the multiple of the Elasticsearch data nodes number.

CHAPTER 4

Kibana

Kibana lets you visualize your Elasticsearch data and navigate the Elastic Stack. Kibana gives you the freedom to select the way you give shape to your data. And you don't always have to know what you're looking for. Kibana core ships with the classics: histograms, line graphs, pie charts, sunbursts, and more. Plus, you can use Vega grammar to design your own visualizations. All leverage the full aggregation capabilities of Elasticsearch. Perform advanced time series analysis on your Elasticsearch data with our curated time series UIs. Describe queries, transformations, and visualizations with powerful, easy-to-learn expressions.

CHAPTER 5

Logstash

Logstash is an open source data collection engine with real-time pipelining capabilities. Logstash can dynamically unify data from disparate sources and normalize the data into destinations of your choice. Cleanse and democratize all your data for diverse advanced downstream analytics and visualization use cases.

While Logstash originally drove innovation in log collection, its capabilities extend well beyond that use case. Any type of event can be enriched and transformed with a broad array of input, filter, and output plugins, with many native codecs further simplifying the ingestion process. Logstash accelerates your insights by harnessing a greater volume and variety of data.

CHAPTER 6

ELK

“ELK” is the acronym for three open source projects: Elasticsearch, Logstash, and Kibana. Elasticsearch is a search and analytics engine. Logstash is a serverside data processing pipeline that ingests data from multiple sources simultaneously, transforms it, and then sends it to a “stash” like Elasticsearch. Kibana lets users visualize data with charts and graphs in Elasticsearch. The Elastic Stack is the next evolution of the ELK Stack.

Where does the data come from?

Energy Log Server is a solution allowing effective data processing from the IT environment that exists in the organization.

The Elasticsearch engine allows building a database in which large amounts of data are stored in ordered indexes. The Logstash module is responsible for load data into Indexes, whose function is to collect data on specific tcp/udp ports, filter them, normalize them and place them in the appropriate index. Additional plugins, that we can use in Logstash reinforce the work of the module, increase its efficiency, enabling the module to quickly interpret data and parse it.

Below is an example of several of the many available Logstash plugins:

exec - receive output of the shell function as an event;

imap - read email from IMAP servers;

jdbc - create events based on JDC data;

jms - create events from Jms broker;

Both Elasticsearch and Logstash are free Open-Source solutions.

More information about Elasticsearch module can be found at: <https://github.com/elastic/elasticsearch>

List of available Logstash plugins: <https://github.com/elastic/logstash-docs/tree/master/docs/plugins>

System services

For proper operation Energy Log Server requires starting the following system services:

- elasticsearch.service - we can run it with a command:

```
systemctl start elasticsearch.service
```

we can check its status with a command:

```
systemctl status elasticsearch.service
```

```
[root@energy-logserver ~]# systemctl status elasticsearch
● elasticsearch.service - Elasticsearch
   Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor preset: disabled)
   Active: active (running) since Wed 2018-08-29 10:11:43 CEST; 1 day 1h ago
     Docs: http://www.elastic.co
   Process: 13349 ExecStartPre=/usr/share/elasticsearch/bin/elasticsearch-systemd-pre-exec (code=exited, status=0/SUCCESS)
    Main PID: 13351 (java)
      CGroup: /system.slice/elasticsearch.service
              └─13351 /bin/java -Xms2g -Xmx2g -Djava.awt.headless=true -XX:+UseP...

Aug 30 11:41:06 energy-logserver elasticsearch[13351]: at org.jboss.netty.uti...
Aug 30 11:41:06 energy-logserver elasticsearch[13351]: at org.jboss.netty.uti...
Aug 30 11:41:06 energy-logserver elasticsearch[13351]: at java.util.concurren...
Aug 30 11:41:06 energy-logserver elasticsearch[13351]: at java.util.concurren...
Aug 30 11:41:06 energy-logserver elasticsearch[13351]: at java.lang.Thread.ru...
Aug 30 11:41:12 energy-logserver elasticsearch[13351]: [2018-08-30 11:41:12,5...
Aug 30 11:41:12 energy-logserver elasticsearch[13351]: [2018-08-30 11:41:12,5...
Aug 30 11:41:13 energy-logserver elasticsearch[13351]: [2018-08-30 11:41:13,8...
Aug 30 11:41:13 energy-logserver elasticsearch[13351]: [2018-08-30 11:41:13,8...
Aug 30 11:46:33 energy-logserver elasticsearch[13351]: [2018-08-30 11:46:33,8...
Hint: Some lines were ellipsized, use -l to show in full.
[root@energy-logserver ~]#
```

- kibana.service - we can run it with a command:

```
systemctl start kibana.service
```

we can check its status with a command:

```
systemctl status kibana.service
```

```
[root@energy-logserver ~]# systemctl status kibana
● kibana.service - Kibana
   Loaded: loaded (/usr/lib/systemd/system/kibana.service; enabled; vendor prese
t: disabled)
   Active: active (running) since Thu 2018-08-30 11:27:23 CEST; 43min ago
 Main PID: 1489 (node)
   CGroup: /system.slice/kibana.service
           └─1489 /opt/kibana/bin/./node/bin/node /opt/kibana/bin/./src/cli...

Aug 30 11:45:03 energy-logserver kibana[1489]: response statusCode: 200
Aug 30 11:45:03 energy-logserver kibana[1489]: {"type":"response","@timestam...C
Aug 30 11:45:04 energy-logserver kibana[1489]: {"type":"response","@timestam...D
Aug 30 11:45:04 energy-logserver kibana[1489]: {"type":"response","@timestam...t
Aug 30 11:45:04 energy-logserver kibana[1489]: {"type":"response","@timestam...k
Aug 30 11:45:04 energy-logserver kibana[1489]: {"type":"response","@timestam...D
Aug 30 11:45:04 energy-logserver kibana[1489]: {"type":"response","@timestam...p
Aug 30 11:54:43 energy-logserver kibana[1489]: {"type":"response","@timestam...}
Aug 30 11:54:43 energy-logserver kibana[1489]: {"type":"response","@timestam...6
Aug 30 11:54:44 energy-logserver kibana[1489]: {"type":"response","@timestam...
Hint: Some lines were ellipsized, use -l to show in full.
[root@energy-logserver ~]#
```

- logstash.service - we can run it with a command:

```
systemctl start logstash.service
```

we can check its status with a command:

```
systemctl status logstash.service
```

```
[root@energy-logserver ~]# systemctl status logstash
● logstash.service - logstash
   Loaded: loaded (/etc/systemd/system/logstash.service; disabled; vendor preset
: disabled)
   Active: active (running) since Thu 2018-08-30 12:14:03 CEST; 8s ago
 Main PID: 4163 (java)
   CGroup: /system.slice/logstash.service
           └─4163 /bin/java -XX:+UseParNewGC -XX:+UseConcMarkSweepGC -XX:CMSI...

Aug 30 12:14:03 energy-logserver systemd[1]: Started logstash.
Aug 30 12:14:03 energy-logserver systemd[1]: Starting logstash...
```

First configuration steps

To install and configure Energy Log Server on the CentOS Linux system you should:

- copy archive Energy Log Server tar.bz2 to the hosted server;
- extract archive Energy Log Server tar.bz2 contain application:

```
cd /root/  
tar xvfj archive.tar.bz2
```

- go to the application directory and run installation script as a root user:

```
cd /roo/insatll/  
./install.sh
```

During instalation you will be ask about following tasks:

- add firewall exeption on ports 22(ssh), 5044, 5514 (Logstash), 5601 (Kibana), 9200 (Elastisearch), 9300 (ES cross-JVM);
- installation of Java environment (Open-JDK), if you use your own Java environment - answer “N”;
- connect to the Energy Log Server CentOS repository, which provides Python libraries, and some fonts;
- installation of Logstash application;
- configuration of Logstash with custom Energy Log Server configuration;
- installation of mail components for Energy Log Server notification;
- installation of Kibana, the Energy Log Server GUI;
- installation of data-node of Elasticsearch;
- configuration of Elasticsearch as Data Node;
- configuration of Elasticsearch as Master Node.

Optionally you can:

- install and configure the filebeat agent;

- install and configure the winlogbeat agent;
- configure Energy Log Server perf_data to integrated with the Energy Log Server Monitor;
- configure naemonLogs to integrated with the Naemon;
- configure integration with Active Directory and SSO servers. You can find necessary information in [12-00-00-Integration_with_AD](#) and [13-00-00-Windows-SSO](#);
- install and configure monitoring with Marver:

```
cd /usr/share/elasticsearch
sudo bin/plugin install license
sudo bin/plugin install marvel-agent
systemctl restart elasticsearch
```

- enable predictive functionality in Intelligence module:

```
curl -XPOST 'http://localhost:9200/_aliases' -d '{
  "actions" : [
    { "add" : { "index" : "intelligence", "alias" : "predictive" } },
    { "add" : { "index" : "perfdata-linux", "alias" : "predictive" } }
  ]}'
```

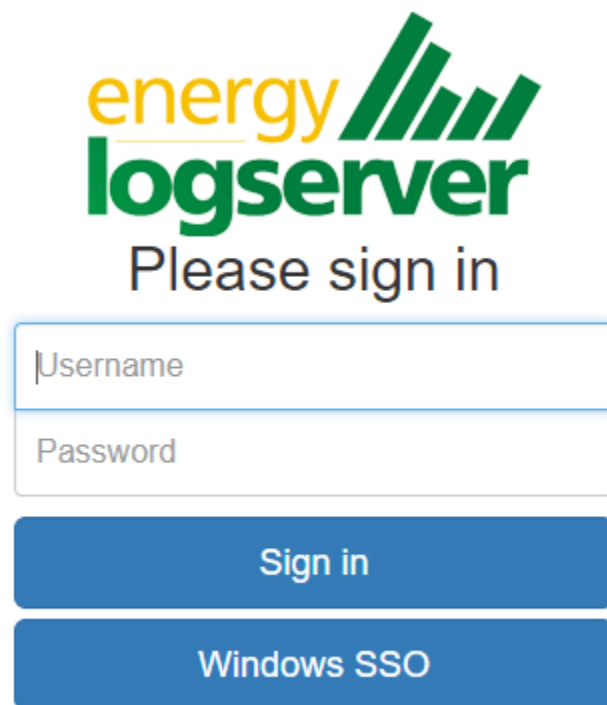
- generate writeback index for Alert service:

```
/opt/alert/bin/elastalert-create-index --config /opt/alert/config.yaml
```

CHAPTER 10

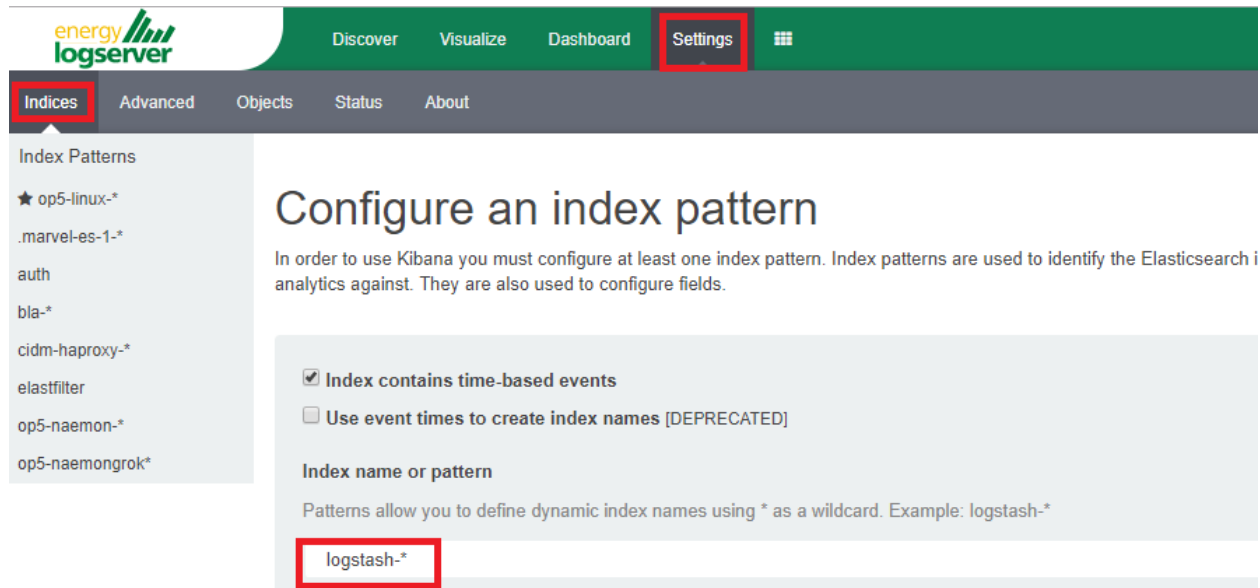
First login

If you log in to Energy Log Server for the first time, you must specify the Index to be searched. We have the option of entering the name of your index, indicate a specific index from a given day, or using the asterix (*) to indicate all of them matching a specific index pattern. Therefore, to start working with Energy Log Server application, we log in to it (by default the user: logserver/password:logserver).



The image shows the login interface for Energy Log Server. At the top is the logo, which consists of the word "energy" in orange, "logserver" in green, and a green bar chart icon to the right. Below the logo, the text "Please sign in" is displayed. There are two input fields: the first is labeled "Username" and the second is labeled "Password". Below these fields are two blue buttons: the top one is labeled "Sign in" and the bottom one is labeled "Windows SSO".


The application at the first login is set by default on the tab: **Settings\Indices**



The screenshot shows the Energy-Log-Server web interface. The top navigation bar has tabs for Discover, Visualize, Dashboard, and Settings (highlighted with a red box). The left sidebar has tabs for Indices (highlighted with a red box), Advanced, Objects, Status, and About. The main content area is titled 'Configure an index pattern'. It includes a list of index patterns on the left: op5-linux-*, .marvel-es-1-*, auth, bla-*, cidm-haproxy-*, elastfilter, op5-naemon-*, and op5-naemongrok*. The main content area has a section for 'Index contains time-based events' with two checkboxes: 'Index contains time-based events' (checked) and 'Use event times to create index names [DEPRECATED]' (unchecked). Below this is a section for 'Index name or pattern' with a text input field containing 'logstash-*' (highlighted with a red box). The text input field has a placeholder text: 'Patterns allow you to define dynamic index names using * as a wildcard. Example: logstash-*'.

In the place where application by default sets name of the Logstash-* pattern, enter the name of the index or index pattern (after confirming that the index or sets of indexes exists).

In additional, the field name should be given, after witch individual event (events) should be sorter. By default the *timestamp* is set, which is the time of occurrence of the event, but depending of the preferences. It may also be the time of the indexing or other selected based on the fields indicate on the event.



Configure an index pattern

In order to use Kibana you must configure at least one index pattern. Index configuration fields.

☒ **Index contains time-based events**

☐ **Use event times to create index names [DEPRECATED]**

Index name or pattern

Patterns allow you to define dynamic index names using * as a wildcard

☐ **Do not expand index pattern when searching (Not recommended)**

By default, searches against any time-based index pattern that contains selected time range.

Searching against the index pattern *logstash-** will actually query elastic

Time-field name ⓘ [refresh fields](#)

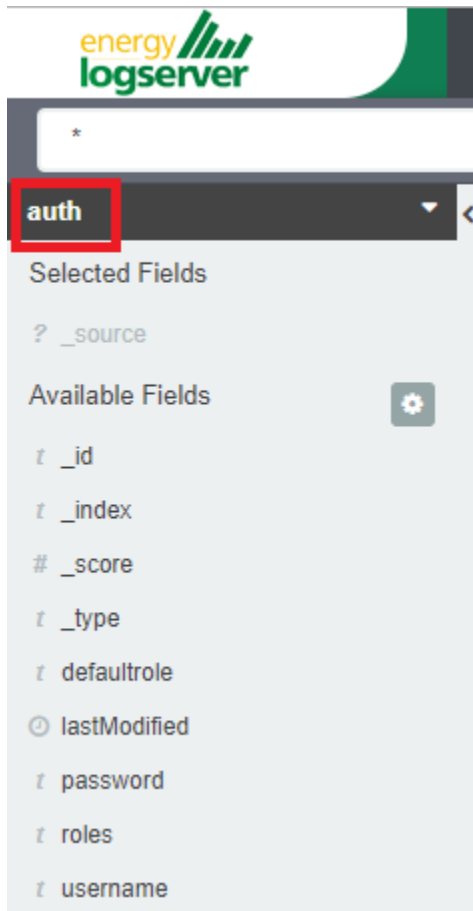
At any time, you can add more indexes or index patterns by going to the main tab select „Settings” and next select „Indices”.

CHAPTER 11

Index selection


After login into Energy Log Server you will going to „Discover” tab, where you can interactively explore your data. You have access to every document in every index that matches the selected index patterns.

If you want to change selected index, click on the black field with the name of the current object in the left panel. Clicking on the object from the expanded list of previously create index patterns, will change the searched index.



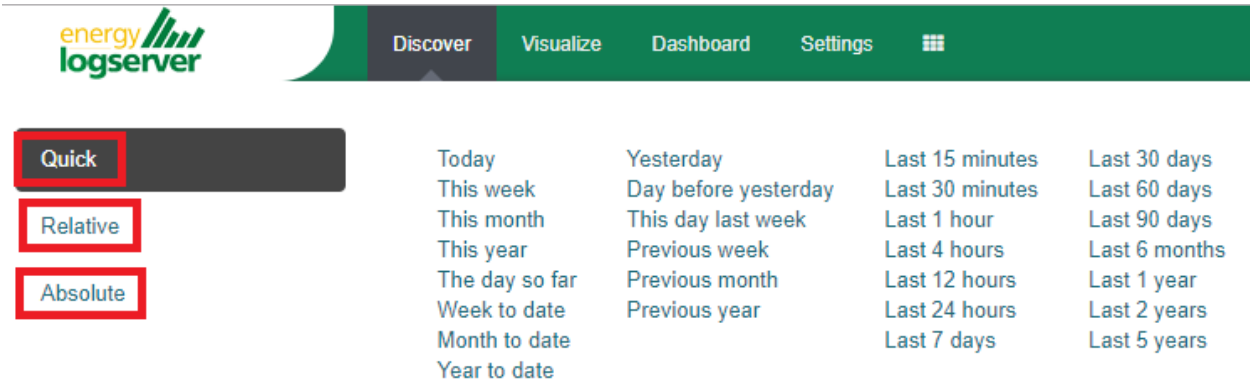
Time settings and refresh

In the upper right corner there is a section in which it defines the range of time that Energy Log Server will search in terms of conditions contained in the search bar. The default value is the last 15 minutes.






After clicking this selection, we can adjust the scope of search by selecting one of the three tabs in the drop-down window:

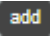
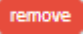
- Quick:** contain several predefined ranges that should be clicked.
- Relative:** in this windows specify the day from which Energy Log Server should search for data.
- Absolute:** using two calendars we define the time range for which the search results are to be returned.



CHAPTER 13

Fields

Energy Log Server in the body of searched events, it recognize fields that can be used to created more precision queries. The extracted fields are visible in the left panel. They are divided on three types: timestamp, marked on clock icon  `timestamp`; text, marked with the letter “t”  `params.level` and digital, marked witch hashtag  `_score`.

Pointing to them and clicking on icon , they are automatically transferred to the „Selected Fields” column and in the place of events a table with selected columns is created on regular basis. In the “Selected Fields” selection you can also delete specific fields from the table by clicking  on the selected element.

Selected Fields

☒ timestamp

☐ operation

☐ username

☐ _id

Available Fields



☐ _index

☐ #_score

☐ _type

☐ method

☐ params._

☐ params.allow_no_indices

☐ params.ignore_unavailable

☐ params.include_defaults

☐ params.level

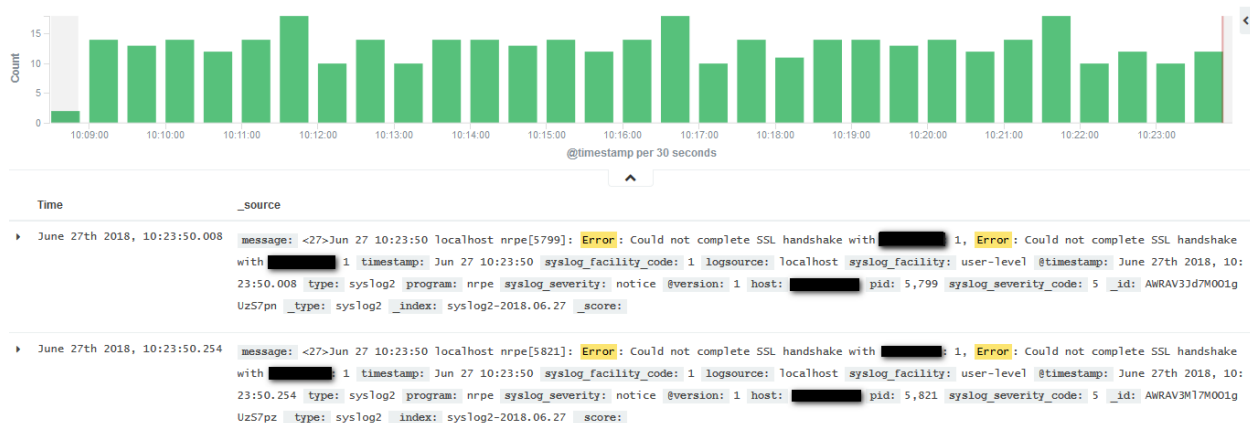
☐ params.preference

☐ params.timeout

☐ request

Filtering and syntax building

We use the query bar to search interesting events. For example, after entering the word „error”, all events that contain the word will be displayed, additionally highlighting them with an orange background.



Fields can be used in the similar way by defining conditions that interesting us. The syntax of such queries is:

```
<fields_name:<fields_value>
```

Example:

```
status:500
```

This query will display all events that contain the „status” fields with a value of 500.

The field value does not have to be a single, specific value. For digital fields we can specify range in the following scheme:

```
<fields_name:[<range_from TO <range_to]
```

Example:

```
status:[500 TO 599]
```

This query will return events with status fields that are in the range 500 to 599.

The search language used in Energy Log Server allows to you use logical operators „AND”, „OR” and „NOT”, which are key and necessary to build more complex queries.

- **AND** is used to combined expressions, e.g. „error AND „access denied”. If an event contain only one expression or the words error and denied but not the word access, then it will not be displayed.
- **OR** is used to search for the events that contain one OR other expression, e.g. „status:500” OR “denied”. This query will display events that contain word „denied” or status field value of 500. Energy Log Server uses this operator by default, so query „status:500” “denied” would return the same results.
- **NOT** is used to exclude the following expression e.g. „status:[500 TO 599] NOT status:505” will display all events that have a status fields, and the value of the field is between 500 and 599 but will eliminate from the result events whose status field value is exactly 505.
- **The above methods** can be combined with each other by building even more complex queries. Understanding how they work and joining it, is the basis for effective searching and full use of Energy Log Server.


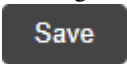
Example of query built from connected logical operations:


```
status:[500 TO 599] AND („access denied" OR error) NOT status:505
```

Returns in the results all events for which the value of status fields are in the range of 500 to 599, simultaneously contain the word „access denied” or „error”, omitting those events for which the status field value is 505.

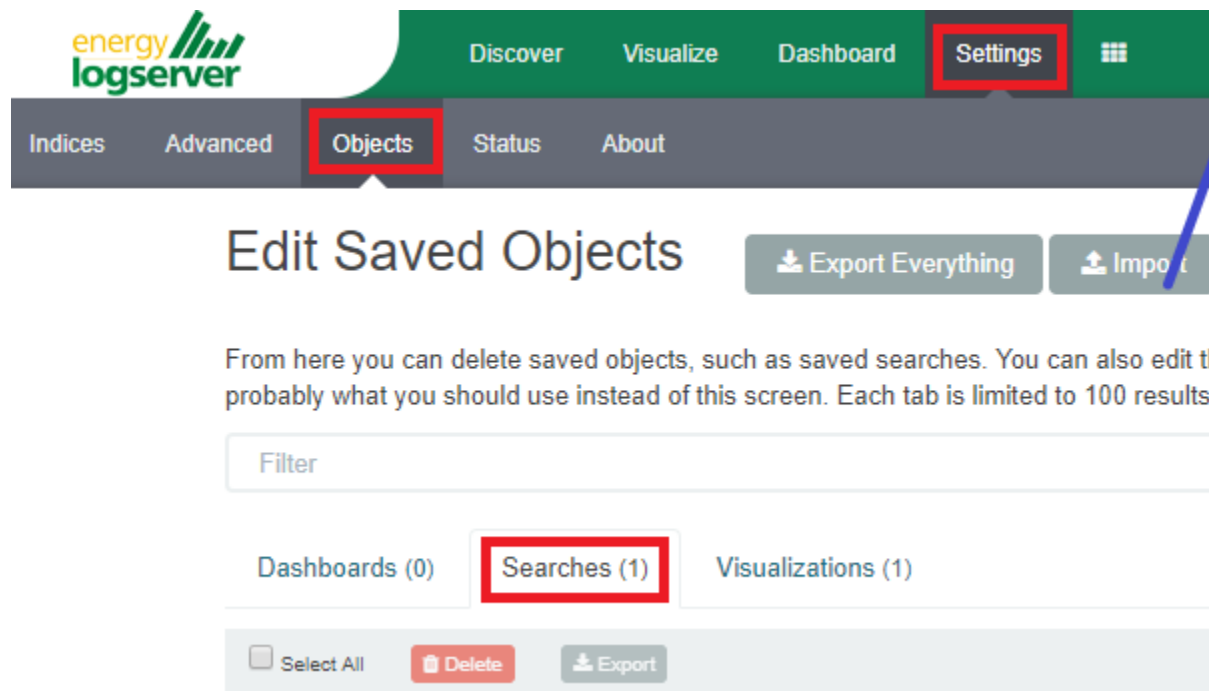
CHAPTER 15



Saving and deleting queries



Saving queries enables you to reload and use them in the future. To do this, click on the icon  to the right on the query bar. This will bring up a window in which we give the query a name and then click the button .

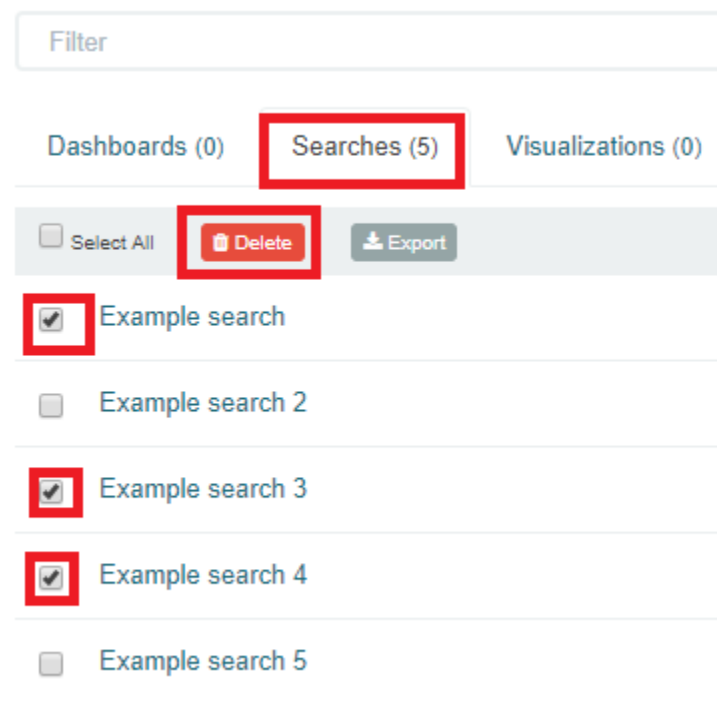




Saved queries can be opened by going to „Settings” from the main menu at the top of the page, then from the submenu select „Object” and finally the „Searches” tab.



To open a saved query in the search window, you can click on the icon  to the right of the query you are interested in. After clicking on the icon  or the name of the saved query, we will gain access to the advanced editing mode, so that we can change the query on a lower level. It is a powerful tool designed for advanced users, designed to modify the query and the way it is presented by Energy Log Server.

To delete a saved query, select it in the list (the icon  to the left of the query name), and then click . In this way, you can delete many saved queries simultaneously.



From this level, you can also export saved queries in the same way. To do this, you need to click on  and choose the save location. The file will be saved in .JSON format. If you then want to import such a file to Energy Log Server, click on button , at the top of the page and select the desired file.

CHAPTER 16










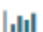
Visualizations

Visualize enables you to create visualizations of the data in your Energy Log Server indices. You can then build dashboards that display related visualizations. Visualizations are based on Energy Log Server queries. By using a series of Energy Log Server aggregations to extract and process your data, you can create charts that show you the trends, spikes, and dips.

Creating visualization

To create visualization, go to the „Visualize” tab from the main menu.

Create a new visualization

	Area chart	Great for stacked timelines in which the total of all series changes of unrelated data points as changes in a series.
	Data table	The data table provides a detailed breakdown, in tabular form, of the data in the chart. Clicking the grey bar at the bottom of the chart will toggle the data table on and off.
	Line chart	Often the best chart for high density time series. Grey bars can be misleading.
	Markdown widget	Useful for displaying explanations or instructions for the chart.
	Metric	One big number for all of your one big number needs.
	Pie chart	Pie charts are ideal for displaying the parts of some whole with no more than 7 slices per pie.
	Tag cloud	A tag cloud visualization is a visual representation of each tag is shown with font size or color.
	Tile map	Your source for geographic maps. Requires an elastic longitude coordinates.
	Timeseries	Create timeseries charts using the timeline expression and moving averages.
	Vertical bar chart	The goto chart for oh-so-many needs. Great for time series you need, you could do worse than to start here.

A new page will be appearing where you can create or load visualization. To load previously created and saved visualization, you must select it from the list.

Or, open a saved visualization

Visualization Filter	manage visualizations
	5 visualizations
Event Levels	
Number of Events	
Number of Events Over Time By Event Log	
Sources	
Top Event IDs	

In order to create a new visualization, you should choose the preferred method of data presentation. Next, specify whether the created visualization will be based on a new or previously saved query. If on new one, select the index whose visualization should concern. If visualization is created from a saved query, you just need to select the appropriate query from the list, or (if there are many saved searches) search for them by name.

CHAPTER 18


Vizualization types

Before the data visualization will be created, first you have to choose the presentation method from an existing list. Currently there are 7 types of visualization. Each of them serves different purposes. If you want to see only the current number of products sold, it is best to choose „Metric”, which presents one value.

69,592

Count

However, if we would like to see user activity trends on pages in different hour and days, a better choice will be „Area chart”, which displays a chart with time division.



The „Markdown widget” views is used to place text e.g. information about the dashboard, explanations and instruc-

tion on how to navigate. Markdown language was used to format the text (the most popular use is GitHub). More information and instruction can be found at this link:<https://help.github.com/categories/writing-on-github/>

Edit visualization and saving

Editing a saved visualization enables you to directly modify the object definition. You can change the object title, add a description, and modify the JSON that defines the object properties. After selecting the index and the method of data presentation, you can enter the editing mode. This will open a new window with empty visualization.


The screenshot shows the Kibana visualization editor interface. At the top, there is a search bar with an asterisk. Below it, the visualization title is "logstash-*". There are two tabs: "Data" and "Options", both highlighted with red boxes. To the right of the tabs are a green play button and a grey close button. Below the tabs, the "metrics" section is highlighted with a red box. It contains a "Y-Axis" dropdown menu set to "Count" and a "+ Add metrics" button. Below the metrics section, the "buckets" section is highlighted with a red box. It contains a "Select buckets type" dropdown menu with three options: "X-Axis", "Split Area", and "Split Chart". At the bottom of the interface is a "Cancel" button.

At the very top there is a bar of queries that can be edited throughout the creation of the visualization. It works in the same way as in the “Discover” tab, which means searching the raw data, but instead of the data being displayed, the visualization will be edited. The following example will be based on the „Area chart”. The visualization modification panel on the left is divided into two tabs: „Data” and „Options”.

In the „Data” tab, you can modify the elements responsible for which data and how should be presented. In this tab there are two sectors: “metrics”, in which we set what data should be displayed, and „buckets” in which we specify how they should be presented. In the „Options” tab, there are settings relating mainly to visual aesthetics. Each type of visualization has separate options. To create the first graph in the chart modification panel, in the „Data” tab we add X-Axis in the “buckets” sections. In „Aggregation” choose „Date Histogram”, in „Field” should automatically be located “timestamp”, and “interval”: “Auto” (if not, this is how we set it). Click on the icon on the panel. Now our first graph should show up.



Some of the options for „Area Chart” are:

Smooth Lines - is used to smooth the graph line.



- **Current time marker** – places a vertical line on the graph that determines the current time.
- **Set Y-Axis Extents** – allows you to set minimum and maximum values for the Y axis, which increases the readability of the graphs. This is useful, if we know that the data will never be less than (the minimum value), or to indicate the goals the company (maximum value).
- **Show Tooltip** – option for displaying the information window under the mouse cursor, after pointing to the point on the graph.



To save the visualization, click the icon  next to the query bar, give it a name and click the button . To load the visualization, go to the „Settings” tab from the main menu at the top of the page, then from the submenu select „Objects” and finally the “Visualizations” tab. From this place, we can also go into advanced editing mode, export and delete visualization.


CHAPTER 20




Dashboards

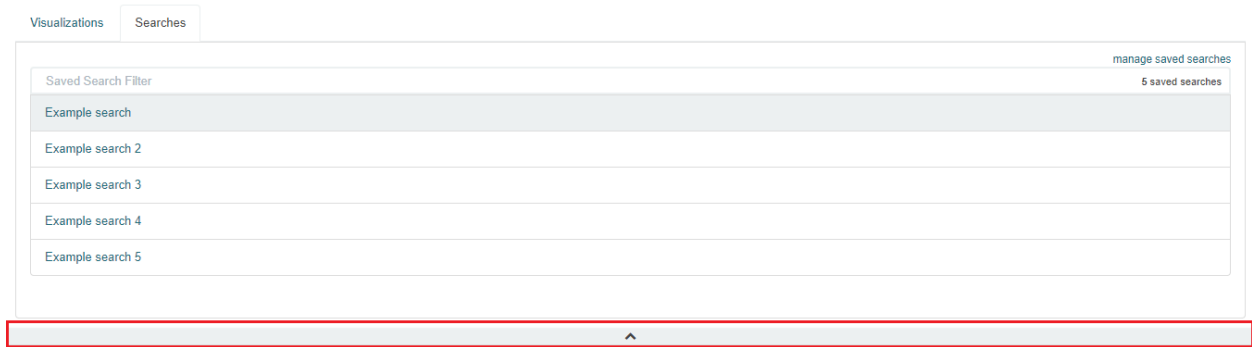
Dashboard is a collection of several visualizations or searches. Depending on how it is build and what visualization it contains, it can be designed for different teams e.g.:

- SOC - which is responsible for detecting failures or threats in the company;
- business - which thanks to the listings can determine the popularity of products and define the strategy of future sales and promotions;
- managers and directors - who may immediately have access to information about the performance units or branches.


To create a dashboard from previously saved visualization and queries, go to the „Dashboard” tab in the main menu bar. When you open it, a new page will appear.



Clicking on the icon  and selecting a saved query and / or visualization from the list will add them to the dashboard. If, there are a large number of saved objects, use the bar to search for them by name. To collapse the window for adding objects, click on the gray bar with icon  (the windows can be re-opened by clicking the icon  to the right of the query bar again).




Elements of the dashboard can be enlarged arbitrarily (by clicking on the right bottom corner of object and dragging the border) and moving (by clicking on the title bar of the object and moving it).

To save a dashboard, click on the icon  to the right of the query bar and give it a name. To load the dashboard go to the “Settings” tab in the main menu at the top of the page, then “Object” from the submenu and finally in the “Dashboard” tab. Here you can also go into advanced editing mode, export and delete dashboard.


CHAPTER 21

Sharing

The dashboard can be share with other Energy Log Server users as well as on any page - by placing a snippet of code. Provided that it cans retrieve information from Energy Log Server.

To do this, open the saved dashboard and click on the icon  to the right of the query bar. A window will appear with two gray bars. The content of the first one is used to provide the dashboard in the page code, and the second is a link that can be passed on to another user. There are two icons next to them, the first is to shorten the length of the

link, and second on copies to clipboard the contest of the given bar

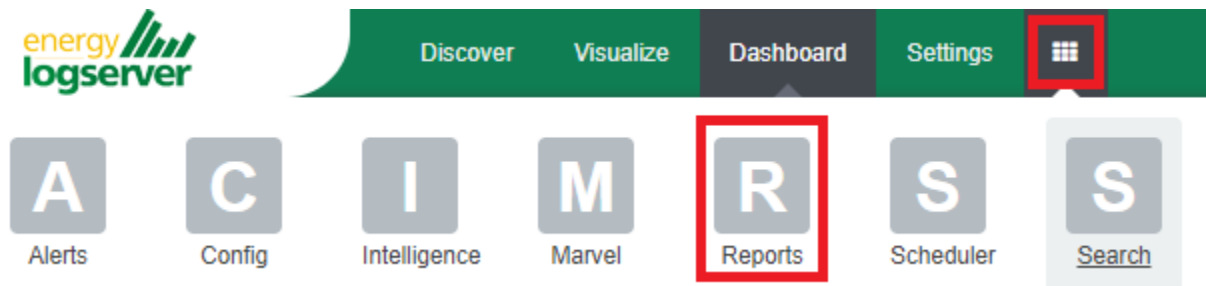


CHAPTER 22

Reports

Energy Log Server contains a module for creating reports that can be run cyclically and contain only interesting data, e.g. a weekly sales report.

To go to the reports windows, select to tiles icon from the main menu bar, and then go to the „Reports” icon (To go back, go to the „Search” icon).



CHAPTER 23

CSV Report

To export data to CSV Report click the Reports icon, you immediately go to the first tab - Export Task Management.

Export Task Mangement

Export Dashboard

Schedule Export Dashboard

Task Management

Start Time	Index path	Search query	Field to export	Status	Action
------------	------------	--------------	-----------------	--------	--------

☒ Index pattern

Index

☐ Index name

Search query

Search query

Time Criteria Field Name

Time Criteria Field

From date

2018-01-01

HH

00

mm

00

ss

00

To date

2018-01-01

HH

00

mm

00

ss

00

Field to export

☐ Include meta fields in export

Submit

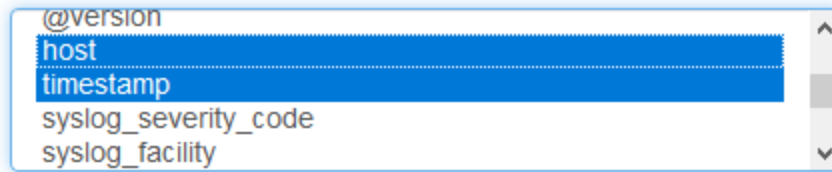
In this tab we have the opportunity to specify the source from which we want to do export. It can be an index pattern. After selecting it, we confirm the selection with the Submit button and a report is created at the moment. The symbol



can refresh the list of reports and see what its status is.


We can also create a report by pointing to a specific index from the drop-down list of indexes.

We can also check which fields are to be included in the report. The selection is confirmed by the Submit button.

Field to export☐ Include meta fields in export

When the process of generating the report (Status:Completed) is finished, we can download it (Download button) or delete (Delete button). The downloaded report in the form of *.csv file can be opened in the browser or saved to the disk.

[Export Task Management](#) [Export Dashboard](#) [Schedule Export Dashboard](#)

Task Management 

Start Time	Index path	Search query	Field to export	Status	Action
2018-08-30 11:32:58	op5-linux-2018.08.30	hostname:emPRD_songo-www	hostname servicestate servicestatety pe servicedescription	Complete	Download Delete

☒ Index pattern

Index

☐ Index name

Search query

Search query

Time Criteria Field Name

Time Criteria Field

From date

2018-01-01

HH

00 ▾

mm

00 ▾

ss

00 ▾

To date

2018-01-01

HH

00 ▾

mm

00 ▾

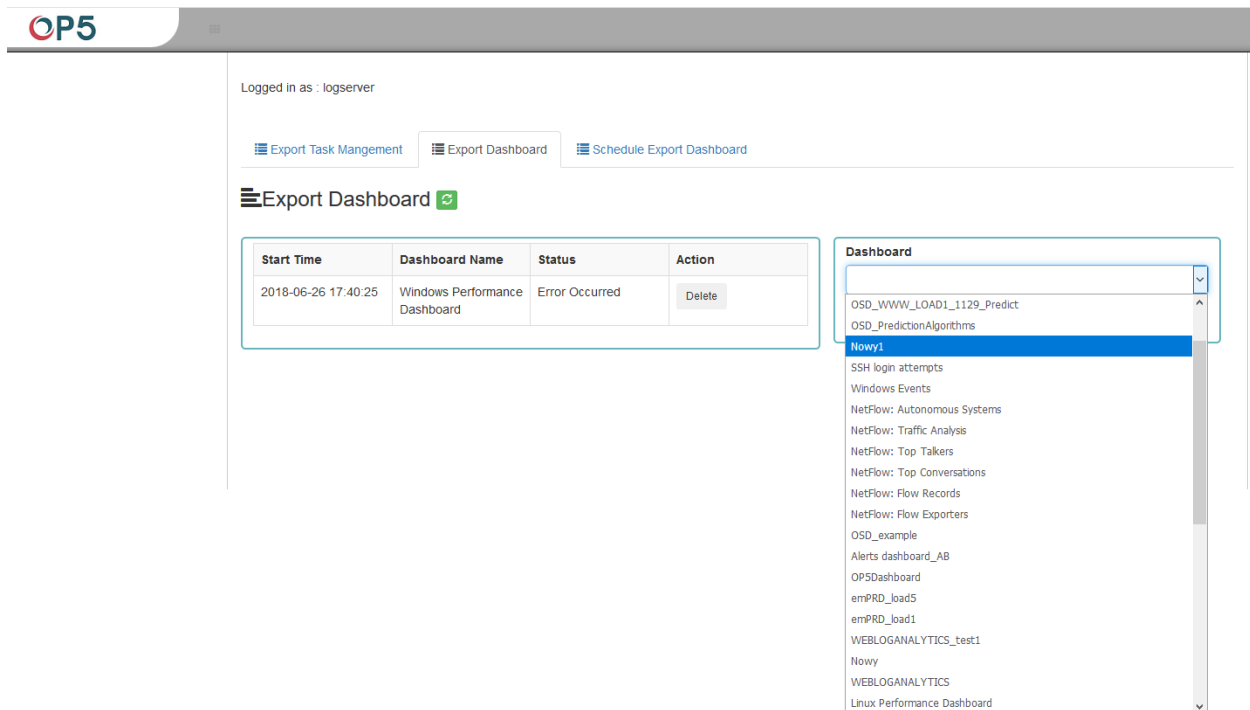
ss

00 ▾

In this tab, the downloaded data has a format that we can import into other systems for further analysis.

PDF Report

In the Export Dashboard tab we have the possibility to create graphic reports in PDF files. To create such a report, just from the drop-down list of previously created and saved Dashboards, indicate the one we are interested in, and then confirm the selection with the Submit button. A newly created export with the Processing status will appear on the list under Dashboard Name. When the processing is completed, the Status changes to Complete and it will be possible to download the report.



OP5

Logged in as : logserver

Export Task Mangement Export Dashboard Schedule Export Dashboard

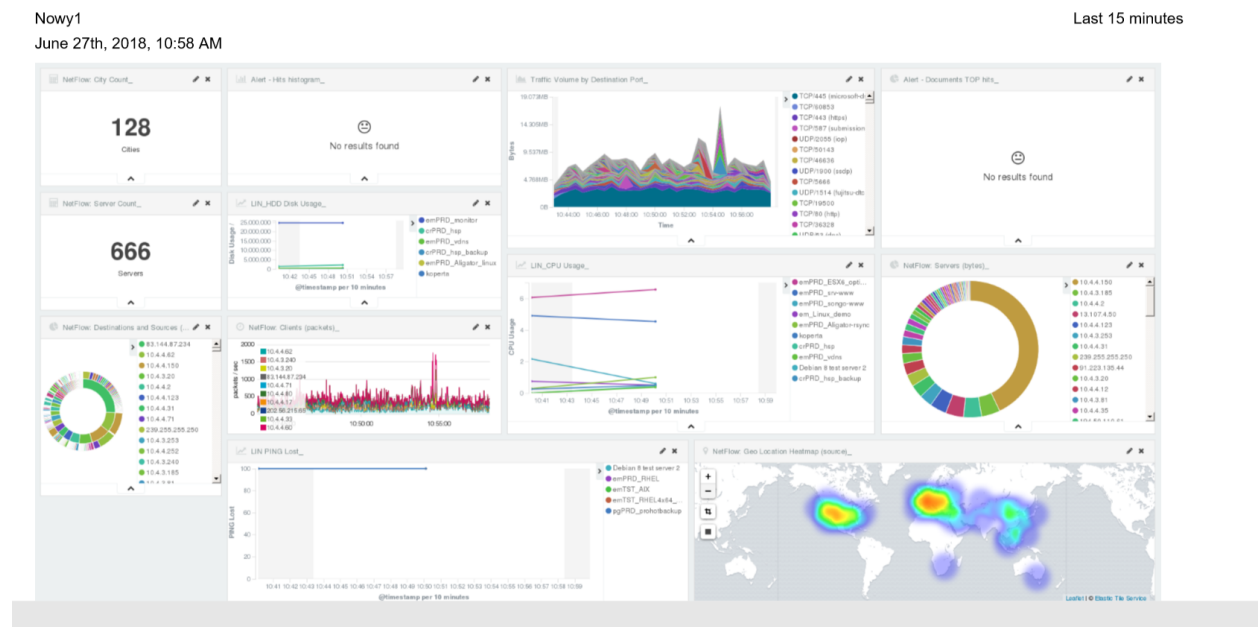
Export Dashboard

Start Time	Dashboard Name	Status	Action
2018-06-26 17:40:25	Windows Performance Dashboard	Error Occurred	Delete

Dashboard

- OSD_WWW_LOAD1_1129_Predict
- OSD_PredictionAlgorithms
- Nowy1
- SSH login attempts
- Windows Events
- NetFlow: Autonomous Systems
- NetFlow: Traffic Analysis
- NetFlow: Top Talkers
- NetFlow: Top Conversations
- NetFlow: Flow Records
- NetFlow: Flow Exporters
- OSD_example
- Alerts dashboard_AB
- OP5Dashboard
- emPRD_load5
- emPRD_load1
- WEBLOGANALYTICS_test1
- Nowy
- WEBLOGANALYTICS
- Linux Performance Dashboard


By clicking the Download button, the report is downloaded to the disk or we can open it in the PDF file browser. There is also to option of deleting the report with the Delete button.



Scheduler Report (Schedule Export Dashboard)

In the Report selection, we have the option of setting the Scheduler which from Dashboard template can generate a report at time intervals. To do this goes to the Schedule Export Dashboard tab.

Logged in as : logserver

 Export Task Mangement

 Export Dashboard

 Schedule Export Dashboard

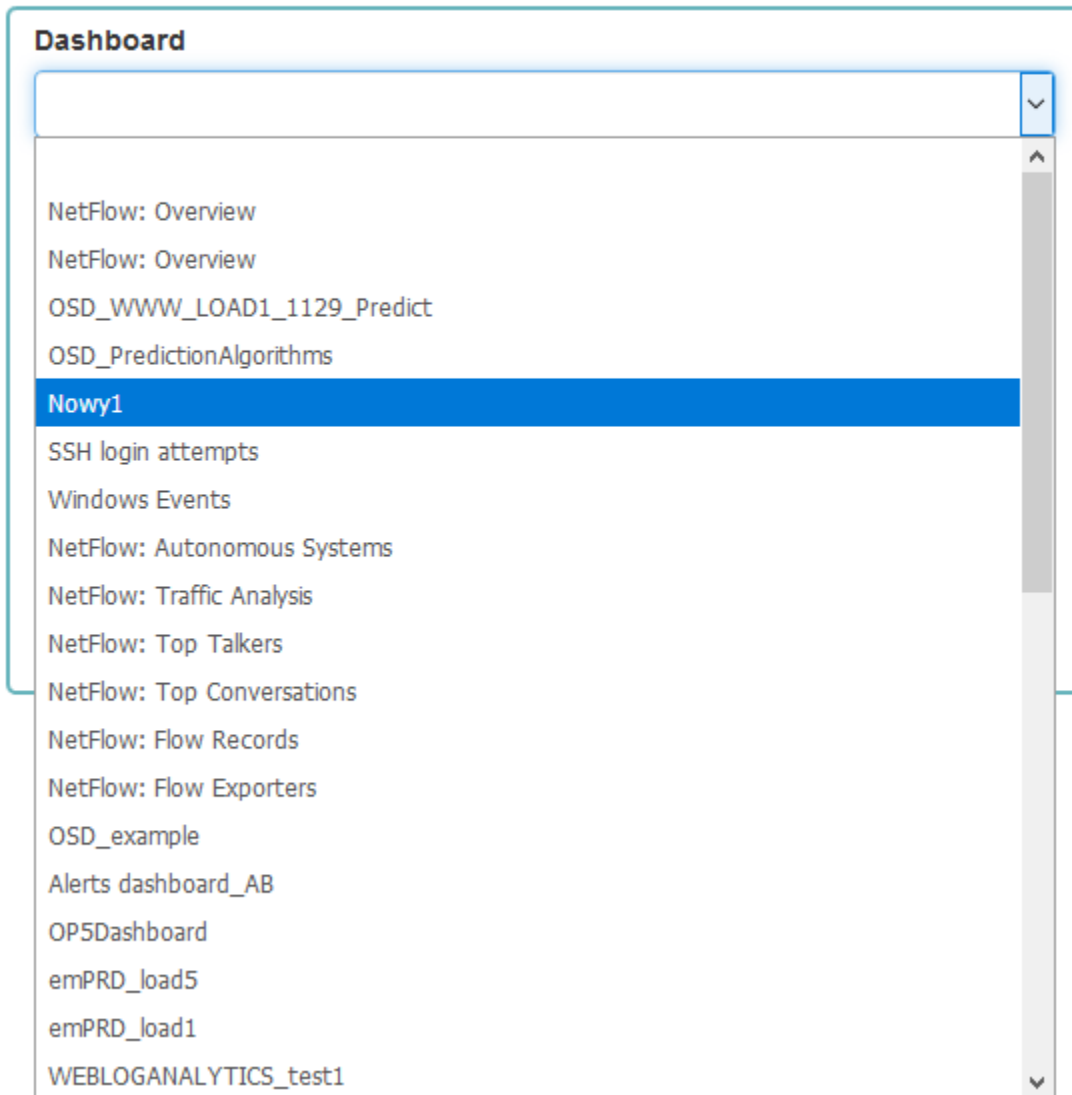
Scheduler for Dashboard

Dashboard
<input type="text"/>
Email Topic
<input type="text"/>
Emails
<input type="text"/>
Cron Schedule
<input type="text"/>
<input type="button" value="Submit"/>

Dashboard Name	Scheduled	Status	Action
Nowy1	Monthly	STOPPED	<input type="button" value="Cancel"/>

In this tab mark the saved Dashboard.

Scheduler for Dashboard



In the Email Topic field, enter the Message title, in the Email field enter the email address to which the report should be sent. From drop-down list choose at what frequency you want the report to be generated and sent. The action configured in this way is confirmed with the Submit button.

Scheduler for Dashboard

Dashboard

Nowy1

Email Topic

Dashboard Nowy1

Emails

emca@it.emca.pl

Cron Schedule

Daily

Weekly

Monthly

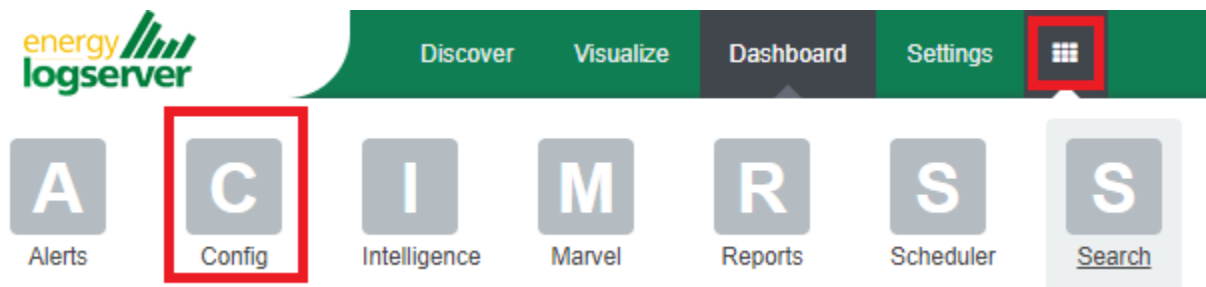
Cron Tab Format

The defined action goes to the list and will generate a report to the e-mail address, with the cycle we set, until we cannot cancel it with the Cancel button.

Dashboard Name	Scheduled	Status	Action
Nowy1	Daily	ENABLED	<div>Cancel</div>

Users, roles and settings

Energy Log Server allows you to manage users and permission for indexes and methods used by them. To go to the management window, select the tile icon from the main menu bar and then go to the „Config” icon (To go back, go to the „Search” icon).



A new window will appear with three main tabs: „User Management”, „Settings” and „License Info”.

From the „User Management” level we have access to the following possibilities: Creating a user in „Create User”, displaying users in „User List”, creating new roles in „Create roles” and displaying existing roles in „List Role”.

CHAPTER 27

Creating a User (Create User)

To create a new user click on the Config icon and you immediately enter the administration panel, where the first tab is to create a new user (**Create User**).

Logged in as : logserver



User Management



Settings



License Info

Create User

User List

Create Role

Role List

Objects permission



Create User

Username**Password****Roles**☒ admin
☐ adrole
☒ Audit only
☐ authsystem
☐ testrole**Default Role**

In the wizard that opens, we enter a unique username (Username field), password for the user (field Password) and assign a role (field Role). In this field we have the option of assigning more than one role. Until we select role in the Roles field, the Default Role field remains empty. When we mark several roles, these roles appear in the Default Role field. In this field we have the opportunity to indicate which role for a new user will be the default role with which the user will be associated in the first place when logging in. The default role field has one more important task - it binds all users with the field / role set in one group. When one of the users of this group create Visualization or Dashboard it will be available to other users from this role(group). Creating the account is confirmed with the Submit button.

User's modification and deletion, (User List)

Once we have created users, we can display their list. We do it in next tab (**User List**).

User List

Username	Roles	Default Role	Actions
abicki	monitoringrole,	monitoringrole	Delete Update
alert	admin,		Delete Update
audit	Audit only,		Delete Update
bla	import_test,	import_test	Delete Update

New Password (bla)

New Password

Re-enter New Password

Re-enter New Password

Roles

autnsystem
Audit only
test
adrole
import_test

Default Role

import_test

Submit

In this view, we get a list of user account with assigned roles and we have two buttons: Delete and Update. The first of these is ability to delete a user account. Under the Update button is a drop-down menu in which we can change the previous password to a new one (New password), change the password (Re-enter Ne Password), change the previously assigned roles (Roles), to other (we can take the role assigned earlier and give a new one, extend user permissions with new roles). The introduced changes are confirmed with the Submit button.

We can also see current user setting and clicking the Update button collapses the previously expanded menu.

Create, modify and delete a role (Create Role), (Role List)

In the Create Role tab we can define a new role with permissions that we assign to a pattern or several index patterns.

Logged in as : logserver



User Management



Settings



License Info

[Create User](#)[User List](#)[Create Role](#)[Role List](#)[Objects permission](#)

+ Create Role

Paths

Methods

get

post

delete

put

head

Roles

In example, we use the syslog2* index pattern. We give this name in the Paths field. We can provide one or more index patterns, their names should be separated by a comma. In the next Methods field, we select one or many methods that will be assigned to the role. Available methods:

- PUT - sends data to the server
- POST - sends a request to the server for a change
- DELETE - deletes the index / document
- GET - gets information about the index /document
- HEAD - is used to check if the index /document exists

In the role field, enter the unique name of the role. We confirm addition of a new role with the Submit button. To see if a new role has been added, go to the net Role List tab.

[Create User](#)
[User List](#)
[Create Role](#)
[Role List](#)
[Objects permission](#)

Role List

Paths	Methods	Roles	Actions
audit*,audit,	get,post,delete,put,head,	Audit only,	Delete Update
security,auth,_auth, .marvel-es-data*,.marvel-es-1*, audit,auditbeat-*,	get,post,delete,put,head,	admin,	Delete Update
		adrole,	Delete Update
.kibana*,	get,post,put,head,	authsystem,	Delete Update
beats-*,	get,post,put,head,	beat-role,	Delete Update
test_raporty_idx,	get,post,head,	import_test,	Delete Update
op5*,	get,post,delete,put,head,	monitoringrole,	Delete Update
op5*,	get,	readonlyop5,	Delete Update
syslog2*,	get,	search,	Delete Update
Paths ["search"] <input type="text" value="syslog2*"/>			
Methods <input type="text" value="get"/> <input type="text" value="post"/>			

As we can see, the new role has been added to the list. With the Delete button we have the option of deleting it, while under the Update button we have a drop-down menu thanks to which we can add or remove an index pattern and add or remove a method. When we want to confirm the changes, we choose the Submit button. Pressing the Update button again will close the menu.

Fresh installation of the application have sewn solid roles which granting user special rights:

- admin - this role gives unlimited permissions to administer / manage the application
- alert - a role for users who want to see the Alert module
- kibana - a role for users who want to see the application GUI
- Intelligence - a role for users who are to see the Intelligence module

Object access permissions (Objects permissions)

In the User Manager tab we can parameterize access to the newly created role as well as existing roles. In this tab we can indicate to which object in the application the role has access.

Example:

In the Role List tab we have a role called **sys2**, it refers to all index patterns beginning with syslog* and the methods get, post, delete, put and head are assigned.

[Create User](#) [User List](#) [Create Role](#) [Role List](#) [Objects permission](#)

☰

Role List

Paths	Methods	Roles	Actions
audit*,audit,	get,post,delete,put,head,	Audit only,	<div>Delete</div> <div>Update</div>
security,auth,_auth, .marvel-es-data*,.marvel-es-1*, audit,auditbeat-*,	get,post,delete,put,head,	admin,	<div>Delete</div> <div>Update</div>
		adrole,	<div>Delete</div> <div>Update</div>
.kibana*,	get,post,put,head,	authsystem,	<div>Delete</div> <div>Update</div>
beats-*,	get,post,put,head,	beat-role,	<div>Delete</div> <div>Update</div>
test_raporty_idx,	get,post,head,	import_test,	<div>Delete</div> <div>Update</div>
op5*,	get,post,delete,put,head,	monitoringrole,	<div>Delete</div> <div>Update</div>
op5*,	get,	readonlyop5,	<div>Delete</div> <div>Update</div>
audit,	get,post,delete,put,head,	auditrole	<div>Delete</div> <div>Update</div>
syslog*,	get,post,delete,put,head,	sys2,	<div>Delete</div> <div>Update</div>
op5*,	get,post,delete,put,head,	syslogrole,	<div>Delete</div> <div>Update</div>
winad*,	get,post,delete,put,head,	test,	<div>Delete</div> <div>Update</div>

When we go to the Object permission tab, we have the option to choose the sys2 role in the drop-down list choose a role:

[Create User](#) [User List](#) [Create Role](#) [Role List](#) [Objects permission](#)

▼

Objects permissions

Choose a role:

sys2

readonlyop5

sys2

admin

syslogrole

skozak

authsystem

Audit only

test

adrole

import_test

monitoringrole

tylk0ADS

beat-role

NetFlow: Destinations (packets)

NetFlow: Servers (flow records)

NetFlow: Destination Ports (flow records)

NetFlow: VLANs (flow records)

>

<

Save

<input type="checkbox"/>	Object name	Type	Read	Update	<input type="checkbox"/>
<input type="checkbox"/>	syslog2*	index-pattern	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
<input type="checkbox"/>	Windows Events	dashboard	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
<input type="checkbox"/>	op5-linux	index-pattern	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

After selecting, we can see that we already have access to the objects: two index patterns syslog2* and Energy Log

70

Chapter 30. Object access permissions (Objects permissions)

Server-* and on dashboard Windows Events. There are also appropriate read or updates permissions.

[Create User](#) [User List](#) [Create Role](#) [Role List](#) [Objects permission](#)

▼ Objects permissions

Choose a role:

sys2

Save

Find:

Write object name

NetFlow: ToS Count

NetFlow: Destinations and Ports (packets)

NetFlow: Destinations (flow records)

NetFlow: VLANs (packets)

NetFlow: Countries and Cities (packets)

NetFlow: Autonomous Systems (flow records)

NetFlow: Top Cities

NetFlow: Geo Location Heatmap (client)

NetFlow: Traffic Locality (flow records)

NetFlow: Source ASs (flow records)

NetFlow: Sources (bytes)

NetFlow: Cities (packets)

NetFlow: Sources and Ports (bytes)

NetFlow: Destinations (packets)

NetFlow: Servers (flow records)

NetFlow: Destination Ports (flow records)

NetFlow: VLANs (flow records)

Select object type:

All

Dashboard

Index pattern

Search

Visualization

>

<

<input type="checkbox"/>	Object name	Type	Read	Update	<input type="checkbox"/>
<input type="checkbox"/>	syslog2*	index-pattern	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
<input type="checkbox"/>	Windows Events	dashboard	<input checked="" type="checkbox"/>	<input type="checkbox"/>	
<input type="checkbox"/>	op5-linux	index-pattern	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	

From the list we have the opportunity to choose another object that we can add to the role. We have the ability to quickly find this object in the search engine (Find) and narrowing the object class in the drop-down field “Select object type”. The object type are associated with saved previously documents in the sections Dashboard, Index pattern,

Search and Visualization. By buttons and Save button to save the selection. we have the ability to add or remove or object,

71

Default user and passwords

The table below contains built-in user accounts and default passwords:

Address	User	Password	Description
↪ -----	Usage		-----
↪ -----	-----	-----	-----
↪ https://localhost:5601	logserver	logserver	A built-in *superuser*
↪ account			
	alert	alert	A built-in account for the
↪ Alert module			
	intelligence	intelligence	A built-in account for the
↪ Intelligence module	usage	for	authorizing communication with elasticsearch
↪ server, for put and get the data from indexes			
	scheduler	scheduler	A built-in account for the
↪ Scheduler module			

CHAPTER 32

Settings

The Settings tab is used to set the audit on different activates or events and consists of several fields:

Logged in as : logserver

[User Management](#) [Settings](#) [License Info](#)

Time Out in seconds (use 0 for longer time-out)

3600

Submit

Delete Audit Data (in days)

30

Submit

Delete Application Tokens (in days)

2

Submit

Delete All Tokens

☒ Login ☒ Logout ☒ Create User ☒ Delete User ☒ Update User ☒ Create Role ☒ Delete Role ☒ Update Role ☒ Export Start ☒ Export Delete ☒ Queries
☐ Content ☐ Bulk

Update Audit Setting

Delete Exported CSVs (in days)

30

Submit

Delete Exported PDFs (in days)

30

- Time Out in seconds field - this field defines the time after how many minutes the application will automatically log you off
- Delete Audit Data (in days) field - in this field we specify after what time the data from the audit should be deleted
- Delete Application Tokens (in days) - in this field we specify after what time the data from the audit should be deleted
- Next field are checkboxes in which we specify what kind of events are to be logged (saved) in the audit index. The events that can be monitored are: logging (Login), logging out (Logout), creating a user (Create User), deleting a user (Delete User), updating user (Update User), creating a role (Create Role), deleting a role (Delete Role), update of the role (Update Role), start of export (Export Start), delete of export (Export Delete), queries (Queries), result of the query (Content), if attempt was made to perform a series of operation (Bulk)
- Delete Exported CSVs (in days) field - in this field we specify after which time exported file with CSV extension have to be removed
- Delete Exported PDFs (in days) field - in this field we specify after which time exported file with PDF extension have to be removed
- To each field is assigned button thanks to which we can confirm the changes.

CHAPTER 33

License (License Info)

The License Information tab consists of several non-editable information fields.

Logged in as : logserver



User Management



Settings



License Info

Company : Foo Bar S.A.

Data nodes in cluster : 1

No of documents :

Indices : [*]

Issued on : 2018-06-08T10:24:27.490

Validity : 3 months

These fields contain information:

- Company field, who owns the license - in this case EMCA S.A.
- Data nodes in cluster field - how many nodes we can put in one cluster - in this case 100
- No of documents field - empty field

- Indices field - number of indexes, symbol[*] means that we can create any number of indices
- Issued on field - date of issue
- Validity field - validity, in this case for 360000 months

CHAPTER 34

Special accounts

At the first installation of the Energy Log Server application, apart from the administrative account (logserver), special applications are created in the application: alert, intelligence and scheduler.

Logged in as : logserver

User Management Settings License Info

Create User **User List** Create Role Role List Objects permission

User List

Username	Roles	Actions
alert	admin,	Delete Update
intelligence	admin,	Delete Update
logserver	admin,	Delete Update
scheduler	admin,	Delete Update

- **Alert Account** - this account is connected to the Alert Module which is designed to track events written to the index for the previously defined parameters. If these are met the information action is started (more on the action in the Alert section)
- **Intelligence Account** - with this account is related to the module of artificial intelligence which is designed to track events and learn the network based on previously defined rules artificial intelligence based on one of the available algorithms (more on operation in the Intelligence chapter)

- **Scheduler Account** - the scheduler module is associated with this account, which corresponds to, among others for generating reports

CHAPTER 35

Alert Module

Energy Log Server allows you to create alerts, i.e. monitoring queries. These are constant queries that run in the background and when the conditions specified in the alert are met, the specify action is taken.

Logged in as : logserver

Create alert rule

Alert rules List

Alerts Status

For example, if you want to know when more than 20 „status:500” responscode from on our homepage appear within an one hour, then we create an alert that check the number of occurrences of the „status:500” query for a specific index every 5 minutes. If the condition we are interested in is met, we send an action in the form of sending a message to our e-mail address. In the action, you can also set the launch of any script.

Enabling the Alert Module

To enabling the alert module you should:

- `generate writeback index for Alert service/opt/alert/bin/elastalert-create-index --config /opt/alert/config.yaml`

Configure an index pattern

In order to use Kibana you must configure at least one index pattern. Index patterns are used to identify the Elasticsearch index to run search and analysis on. Index patterns are also used to configure fields.

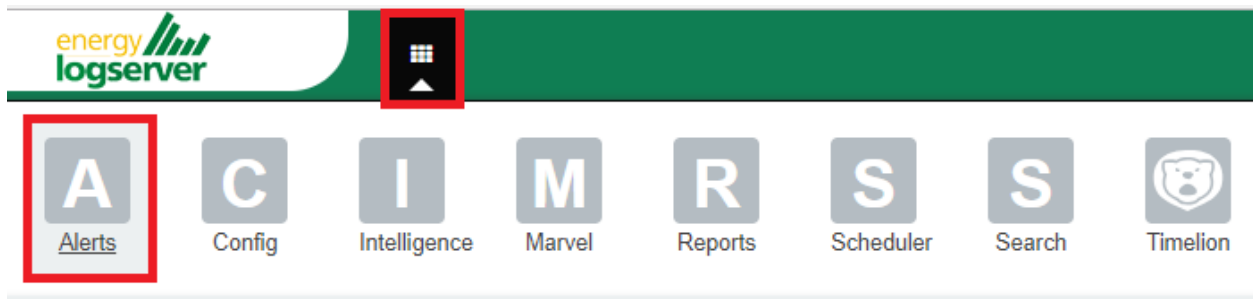
The screenshot shows the 'Configure an index pattern' dialog in Kibana. It has two tabs: 'Index' (selected) and 'Fields'. Under the 'Index' tab, there are two checkboxes: 'Index contains time-based events' (checked) and 'Use event times to create index names [DEPRECATED]' (unchecked). Below these is a section titled 'Index name or pattern' with a text input field containing 'alert*'. A note below the input says: 'Patterns allow you to define dynamic index names using * as a wildcard. Example: logstash-*'. Below the input field is another checkbox: 'Do not expand index pattern when searching (Not recommended)' (unchecked). A paragraph explains: 'By default, searches against any time-based index pattern that contains a wildcard will automatically be expanded to query only the indices that cover the currently selected time range. Searching against the index pattern logstash-* will actually query elasticsearch for the specific matching indices (e.g. logstash-2015.12.21) that fall within the selected time range.' At the bottom, there is a 'Time-field name' section with a dropdown menu showing '@timestamp' and a 'refresh fields' link. A green 'Create' button is at the bottom right.

- `configure the index pattern for alert*`
- `start Alert service systemctl start alert`

CHAPTER 37

Creating Alerts

To create the alert, select the tile icon from the main menu bar and then go to the „Alerts” icon (To go back, go to the „Search” icon).



We will display a page with tree tabs: Create new alerts in „Create alert rule”, manage alerts in „Alert rules List” and check alert status „Alert Status”.

In the alert creation windows we have an alert creation form:

Logged in as : logserver

[Create alert rule](#) [Alert rules List](#) [Alerts Status](#)

+ Create Alert

Name

Index pattern

Role

admin

adrole

alert

intelligence

Type

Description

Example

Example

Alert method

Any

Submit

- Name - the name of the alert, after which we will recognize and search for it.
- Index pattern - a pattern of indexes after which the alert will be searched.
- Role - the role of the user for whom an alert will be available
- Type - type of alert
- Description - description of the alert.
- Example - an example of using a given type of alert. Descriptive field
- Alert method - the action the alert will take if the conditions are met (sending an email message or executing a

command)

- Any - additional descriptive field.

List of Alert rules

The “Alert Rule List” tab contain complete list of previously created alert rules:

Logged in as : logserver

[Create alert rule](#) [Alert rules List](#) [Alerts Status](#)

Alert Rule List

Name	Index pattern	Type	Role	Alert method	Actions
test	*	any	["admin"]	email	Show Enable Delete Update

In this winsow, you can activate / deactivate, delete and update alerts by clicking on the selected icon with the given alert. [Enable](#) [Delete](#) [Update](#).



CHAPTER 39

Alerts status

In the “Alert status” tab, you can check the current alert status: if it activated, when it started and when it ended, how long it lasted, how many event sit found and how many times it worked.

Logged in as : logserver

[Create alert rule](#) [Alert rules List](#) [Alerts Status](#)

 Alerts List  Alert module status: RUNNING [Recovery Alert Dashboard](#)

Name	Start time	End time	Time taken	Hits	Matches
------	------------	----------	------------	------	---------

Also, on this tab, you can recover the alert dashboard, by clicking the “Recovery Alert Dashboard” button.

Type of the Alert module rules

The various RuleType classes, defined in Energy Log Server-Log-Aalytics. An instance is held in memory for each rule, passed all of the data returned by querying Elasticsearch with a given filter, and generates matches based on that data.

- **Any** - The any rule will match everything. Every hit that the query returns will generate an alert.
- **Blacklist** - The blacklist rule will check a certain field against a blacklist, and match if it is in the blacklist.
- **Whitelist** - Similar to blacklist, this rule will compare a certain field to a whitelist, and match if the list does not contain the term.
- **Change** - This rule will monitor a certain field and match if that field changes.
- **Frequency** - This rule matches when there are at least a certain number of events in a given time frame.
- **Spike** - This rule matches when the volume of events during a given time period is spike_height times larger or smaller than during the previous time period.
- **Flatline** - This rule matches when the total number of events is under a given threshold for a time period.
- **New Term** - This rule matches when a new value appears in a field that has never been seen before.
- **Cardinality** - This rule matches when the total number of unique values for a certain field within a time frame is higher or lower than a threshold.
- **Metric Aggregation** - This rule matches when the value of a metric within the calculation window is higher or lower than a threshold.
- **Percentage Match** - This rule matches when the percentage of document in the match bucket within a calculation window is higher or lower than a threshold.

CHAPTER 41

Example of rules

Example of rules

42.1 Unix - Authentication Fail

- index pattern:

syslog-*

- Type:

Frequency

- Alert Method:

Email

- Any:

```
num_events: 4
timeframe:
  minutes: 5

filter:
- query_string:
  query: "program: (ssh OR sshd OR su OR sudo) AND message: \"Failed password\"
↪ "
```

42.2 Windows - Firewall disable or modify

- index pattern:

beats-*

- Type:

Any

- Alert Method:

Email

- Any:

filter:

```
- query_string:  
  query: "event_id:(4947 OR 4948 OR 4946 OR 4949 OR 4954 OR 4956 OR 5025) "
```

CHAPTER 43

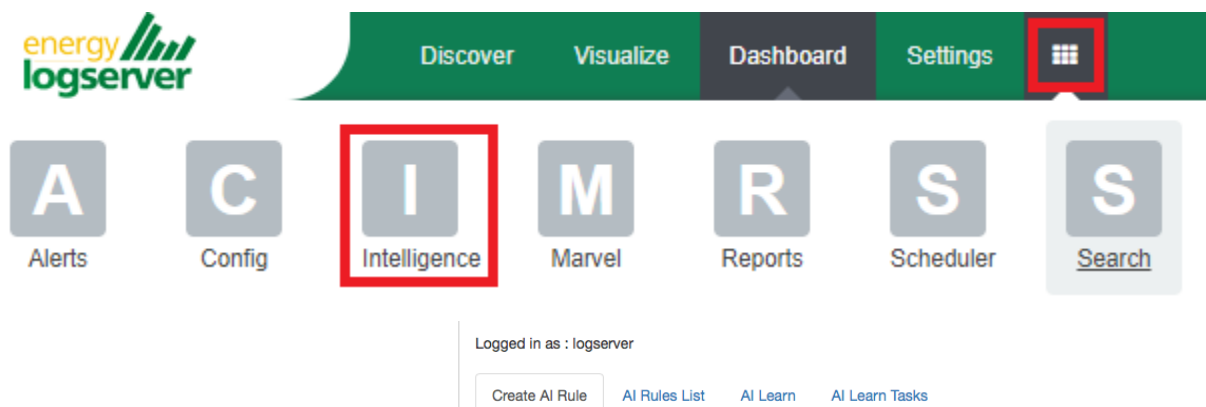
Intelligence Module

A dedicated artificial intelligence module has been built in the Energy Log Server system that allows prediction of parameter values relevant to the maintenance of infrastructure and IT systems. Such parameters include:

- use of disk resources,
- use of network resources,
- using the power of processors
- detection of known incorrect behavior of IT systems In the future, it is planned to launch algorithms that enable automatic detection of anomalies, e.g. non-standard network traffic, which may suggest hacking attempts.

(A detailed description of the implemented algorithms can be found in a separate document)

To access of the Intelligence module, click the tile icon from the main menu bar and then go to the „Intelligence” icon (To go back, click to the „Search” icon).



There are 4 screens available in the module:

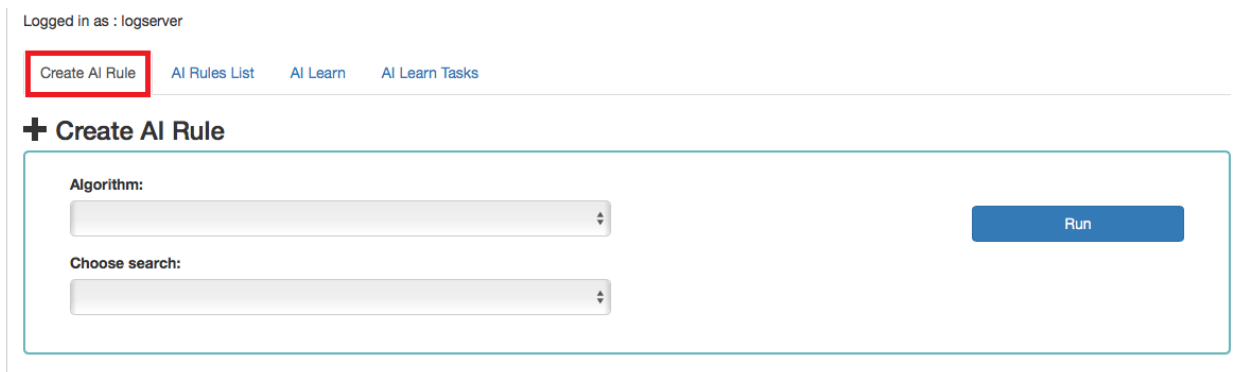
- **Create AI Rule** - the screen allows you to create artificial intelligence rules and run them in scheduler mode or immediately
- **AI Rules List** - the screen presents a list of created artificial intelligence rules with the option of editing, previewing and deleting them

- **AI Learn** - the screen allows to define the conditions for teaching the MLP neural network
- **AI Learn Tasks** - a screen on which the initiated and completed learning processes of neural networks with the ability to preview learning results are presented.

CHAPTER 44

Create AI Rule

To create the AI Rule, click on the tile icon from the main menu bar, go to the „Intelligence” icon and select “Create AI Rule” tab. The screen allows to defining the rules of artificial intelligence based on one of the available algorithms (a detailed description of the available algorithms is available in a separate document).



Logged in as : logserver

Create AI Rule AI Rules List AI Learn AI Learn Tasks

+ Create AI Rule

Algorithm:

Choose search:

Run

Description of the controls available on the fixed part of screen:

- **Algorithm** - the name of the algorithm that forms the basis of the artificial intelligence rule
- **Choose search** - search defined in the Energy Log Server system, which is used to select a set of data on which the artificial intelligence rule will operate
- **Run** - a button that allows running the defined AI rule or saving it to the scheduler and run as planned

The rest of the screen will depend on the chosen artificial intelligence algorithm.

The fixed part of the screen

The screenshot shows a web interface for creating an AI rule. At the top, it says 'Logged in as : logserver'. Below this is a navigation bar with four links: 'Create AI Rule' (highlighted), 'AI Rules List', 'AI Learn', and 'AI Learn Tasks'. The main section is titled '+ Create AI Rule'. It contains a form with two dropdown menus: 'Algorithm:' and 'Choose search:'. To the right of these dropdowns is a blue button labeled 'Run'.

Description of the controls available on the fixed part of screen:

Algorithm - the name of the algorithm that forms the basis of the artificial intelligence rule Choose search - search defined in the Energy Log Server system, which is used to select a set of data on which the artificial intelligence rule will operate Run - a button that allows running the defined AI rule or saving it to the scheduler and run as planned The rest of the screen will depend on the chosen artificial intelligence algorithm.

CHAPTER 46

Screen content for regressive algorithms

Algorithm:

Simple Moving Average

Choose search:

Uslugi_WWW_with_cols

AI Rule Name:

my_test_

Feature to analyse (from search):

perf_data./

Multiply by field (from search):

hostname

Multiply by values (from search):

emPRD_Aligator_linux
emPRD_Cyberoam_public_FC
emPRD_ESX6_optima64
emPRD_RHEL

Time frame:

Day

Value type:

Average

Max probes:

20

Max predictions:

30

Description of controls:

- **feature to analyze from search** - analyzed feature (dictated)
- **multiply by field** - enable multiplication of algorithms after unique values of the feature indicated here. Multiplication allows you to run the AI rule one for e.g. all servers. The value “none” in this field means no multiplication.
- **multiply by values** - if a trait is indicated in the „multiply by field”, then unique values of this trait will appear in this field. Multiplications will be made for the selected values. If at least one of value is not selected, the „Run” buttons will be inactive.
- **time frame** - feature aggregation method (1 minute, 5 minute, 15 minute, 30 minute, hourly, weekly, monthly, 6 months, 12 months)
- **max probes** - how many samples back will be taken into account for analysis. A single sample is an aggregated data according to the aggregation method.
- **value type** - which values to take into account when aggregating for a given time frame (e.g. maximum from time frame, minimum, average)
- **max predictions** - how many estimates we make for ahead (we take time frame)
- **data limit** - limits the amount of data downloaded from the source. It speeds up processing but reduces its quality
- **start date** - you can set a date earlier than the current date in order to verify how the selected algorithm would work on historical data
- **Scheduler** - a tag if the rule should be run according to the plan for the scheduler. If selected, additional fields will appear;

Scheduler: ☒

Prediction cycle
(crontab format):

Enable: ☐

- **Prediction cycle** - plan definition for the scheduler, i.e. the cycle in which the prediction rule is run (e.g. once a day, every hour, once a week). In the field, enter the command that complies with the cron standard. Enable – whether to immediately launch the scheduler plan or save only the definition
- **Role** - only users with the roles selected here and the administrator will be able to run the defend AI rules The selected „time frame” also affects the prediction period. If we choose “time frame = monthly”, we will be able to predict a one month ahead from the moment of prediction (according to the “prediction cycle” value)

CHAPTER 47

Screen content for the Trend algorithm

Algorithm:

Trend

Choose search:

Uslugi_WWW_with_cols

AI Rule Name:

rpa_trend

Feature to analyse (from search):

perf_data./

Time frame:

Day

Value type:

Average

Max probes:

10

Max predictions:

20

Data limit:

10000

Start date:

2018-03-01

Threshold:

-1

Scheduler:



Role:

admin
ALL_test

Description of controls:

- **feature to analyze from search** - analyzed feature (dictated)
- **multiply by field** - enable multiplication of algorithms after unique values of the feature indicated here. Multiplication allows you to run the AI rule one for e.g. all servers. The value “none” in this field means no multiplication.
- **multiply by values** - if a trait is indicated in the „multiply by field”, then unique values of this trait will appear in this field. Multiplications will be made for the selected values. If at least one of value is not selected, the „Run” buttons will be inactive.
- **time frame** - feature aggregation method (1 minute, 5 minute, 15 minute, 30 minute, hourly, weekly, monthly, 6 months, 12 months)
- **max probes** - how many samples back will be taken into account for analysis. A single sample is an aggregated data according to the aggregation method.
- **value type** - which values to take into account when aggregating for a given time frame (e.g. maximum from time frame, minimum, average)
- **max predictions** - how many estimates we make for ahead (we take time frame)
- **data limit** - limits the amount of data downloaded from the source. It speeds up processing but reduces its quality
- **start date** - you can set a date earlier than the current date in order to verify how the selected algorithm would work on historical data
- **Scheduler** - a tag if the rule should be run according to the plan for the scheduler. If selected, additional fields will appear;

Scheduler: ☒

Prediction cycle
(crontab format):

Enable: ☐

- **Prediction cycle** - plan definition for the scheduler, i.e. the cycle in which the prediction rule is run (e.g. once a day, every hour, once a week). In the field, enter the command that complies with the cron standard. Enable – whether to immediately launch the scheduler plan or save only the definition
- **Role** - only users with the roles selected here and the administrator will be able to run the defend AI rules The selected „time frame” also affects the prediction period. If we choose “time frame = monthly”, we will be able to predict a one month ahead from the moment of prediction (according to the “prediction cycle” value)
- **Threshold** - default values -1 (do not search). Specifies the algorithm what level of exceeding the value of the feature „feature to analyze from cheese” is to look for. The parameter currently used only by the “Trend” algorithm.

CHAPTER 48

Screen content for the neural network (MLP) algorithm

Algorithm:

Multi Layer Perceptron ANN

Name:

rpa_ann_2000_ANN_20180503_104024

Choose search:

Uslugi_WWW_with_cols

Accuracy: 0.6149193548387096

Weighted precision: 0.3781258129552549

Overall efficiency: 0.45834267049146893

Run

Attributes to analyse from search		Analysed weight	Attribute analyzed
perf_data./	perf_data./	-0.19525205216734406	perf_data.time
perf_data.free_memory	perf_data.free_mem	-0.07863953880113653	
perf_data.cpu_usage	perf_data.cpu_usa	-0.06251180295737524	
perf_data.mem_usage	perf_data.mem_us	0.05181616786061537	
perf_data.avgqu-sz	perf_data.avgqu-sz	-0.045473151254527465	
perf_data.load15	perf_data.load15	-0.02556274656942572	
perf_data.cpu_user	perf_data.cpu_user	-0.02232814630493624	
perf_data.load5	perf_data.load5	-0.020889999164069112	
perf_data.cpu_idle	perf_data.cpu_idle	0.019885681122719448	
perf_data.await	perf_data.await	0.01827435049755162	
perf_data.cpu_sys	perf_data.cpu_sys	-0.015911517530838776	
perf_data.load1	perf_data.load1	-0.012822584228478538	
perf_data.io_write	perf_data.io_write	0.01221505604864565	
perf_data.r	perf_data.r	-0.012244672087087087	
perf_data.cpu_iowait	perf_data.cpu_iow	-0.011977745509837864	
perf_data.pl		0.006104901588956799	

114

Chapter 48. Screen content for the neural network (MLP) algorithm

Descriptions of controls:

- **Name** - name of the learned neural network
- **Choose search** - search defined in Energy Log Server, which is used to select a set of data on which the rule of artificial intelligence will work
- **Below**, on the left, a list of attributes and their weights based on teaching ANN will be defined during the teaching. The user for each attribute will be able to indicate the field from the above mentioned search, which contain the values of the attribute and which will be analyzed in the algorithm. The presented list (for input and output attributes) will have a static and dynamic part. Static creation by presenting key with the highest weights. The key will be presented in the original form, i.e. `perf_data./` The second part is a DropDown type list that will serve as a key update according to the user's naming. On the right side, the attribute will be examined in a given rule / pattern. Here also the user must indicate a specific field from the search. In both cases, the input and output are narrowed based on the search fields indicated in Choose search.
- **Data limit** - limits the amount of data downloaded from the source. It speeds up the processing, but reduces its quality.
- **Scheduler** - a tag if the rule should be run according to the plan or the scheduler. If selected, additional fields will appear:

Scheduler: ☒

Prediction cycle
(crontab format):

Enable: ☐

- **Prediction cycle** - plan definition for the scheduler, i.e. the cycle in which the prediction rule is run (e.g. once a day, every hour, once a week). In the field, enter the command that complies with the *cron* standard
- **Enable** - whether to immediately launch the scheduler plan or save only the definition
- **Role** - only users with the roles selected here and the administrator will be able to run the defined AI rules

CHAPTER 49

AI Rules List

Logged in as : logserver

Create AI Rule

AI Rules List

AI Learn

AI Learn Tasks

AI Rules List

	Name	Search	Method	Actions
✓	int1	Uslugi_WWW_with_cols	Trend	<div>ShowDeleteUpdatePreview</div>
✗	k1	Uslugi_WWW_with_cols	Trend	<div>ShowDeleteUpdate</div>
✓	k2	Uslugi_WWW_with_cols	Trend	<div>ShowDeleteUpdatePreview</div>
✓	k3	Uslugi_WWW_with_cols	Trend	<div>ShowDeleteUpdatePreview</div>
✓	ko4	Uslugi_WWW_with_cols	Random Forest Regression Shift	<div>ShowDeleteUpdatePreview</div>
✓	ko5	Uslugi_WWW_with_cols	Trend	<div>ShowDeleteUpdatePreview</div>
✓	rpa_lrs_day_2	Linux_host_load	Linear Regression Shift Trend	<div>ShowDeleteUpdatePreview</div>
✓	rpa_lrst_day_100	Linux_host_load	Linear Regression Shift Trend	<div>ShowDeleteUpdatePreview</div>

Choose search:

Linux_host_load

Feature to analyse (from search):

perf_data.load5

Time frame:

Day

Value type:

Average

Max probes:

20

Max predictions:

10

Scheduler:

☐

Role:

adminALL_testauditdatabases

Column description:

- **Status:**
 - ☉ - the process is being processed (the pid of the process is in brackets)
 - ✓ - process completed correctly
 - * - the process ended with an error
- **Name** - the name of the rule
- **Search** - the search on which the rule was run
- **Method** - an algorithm used in the AI rule
- **Actions** - allowed actions:
 - **Show** - preview of the rule definition
 - **Enable/Disable** - rule activation /deactivation
 - **Delete** - deleting the rule
 - **Update** - update of the rule definition
 - **Preview** - preview of the prediction results (the action is available after the processing has been completed correctly).

CHAPTER 50

AI Learn

Logged in as : logserver

[Create AI Rule](#) [AI Rules List](#) [AI Learn](#) [AI Learn Tasks](#)

+ AI Learn

Choose search:

Uslugi_WWW_with_cols

Build (18)

Prefix name:

test_cache_ann_

Choose input cols (25):

- perf_data.size
- perf_data.slow_queries_rate
- perf_data.time
- perf_data.tps
- hostname
- hoststate
- @timestamp
- type
- perf_data.cpu_usage
- perf_data./

Choose output col:

perf_data.time

Time frame:

Minute

Output class category:

if((outputCol) < 10,(floor((outputCol)+1), Double{1

Timeframes Output shift:

0 1 minute

Output class count:

20

Value type:

Average

Split data to train&test:

0.8

Max iter (x100):

from: 1 to: 2

Max probes:

1000

Neurons:

	1st	2nd	3rd
from:	22	80	40
to:	30	80	40

Data limit:

10000000

Results: 18 / 18 [Refresh](#) ☐ Autorefresh

Internal name	Overall efficiency
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200
test_cache_ann_Multi_Layer_Perceptron_ANN_2018050-	0.4402956393200

Algorithm data:

Confusion matrix:

```
0.0 63.0
0.0 106.0
```

Accuracy = 0.6272189349112426

Labels rows count:

Save algorithm: test_cache_ann_Multi_Layer_Perceptron_

Description of controls:

- **Search** - a source of data for teaching the network
- **prefix name** - a prefix added to the id of the learned model that allows the user to recognize the model
- **Input cols** - list of fields that are analyzed / input features. Here, the column that will be selected in the output col should not be indicated. Only those columns that are related to processing should be selected. **
- **Output col** - result field, the recognition of which is learned by the network. **This field should exist in the learning and testing data, but in the production data is unnecessary and should not occur. This field cannot be on the list of selected fields in “input col”.**
- **Output class category** - here you can enter a condition in SQL format to limit the number of output categories e.g. `if((outputCol) < 10, (floor((outputCol))+1), Double(10))`. This condition limits the number of output categories to 10. **Such conditions are necessary for fields selected in “output col” that have continuous values. They must necessarily be divided into categories. In the Condition, use your own outputCol name instead of the field name from the index that points to the value of the “output col” attribute.**
- **Time frame** - a method of aggregation of features to improve their quality (e.g. 1 minute, 5 minutes, 15 minutes, 30 minutes, 1 hour, 1 daily).
- **Time frames output shift** - indicates how many time frame units to move the output category. This allows teaching the network with current attributes, but for categories for the future.
- **Value type** - which values to take into account when aggregating for a given time frame (e.g. maximum from time frame, minimum, average)
- **Output class count** - the expected number of result classes. **If during learning the network identifies more classes than the user entered, the process will be interrupted with an error, therefore it is better to set up more classes than less, but you have to keep in mind that this number affects the learning time.**
- **Neurons in first hidden layer (from, to)** - the number of neurons in the first hidden layer. Must have a value > 0. Jump every 1.
- **Neurons in second hidden layer (from, to)** - the number of neurons in second hidden layer. If = 0, then this layer is missing. Jump every 1.
- **Neurons in third hidden layer (from, to)** - the number of neurons in third hidden layer. If = 0 then this layer is missing. Jump every 1.
- **Max iter** (from, to) - maximum number of network teaching repetitions (the same data is used for learning many times in internal processes of the neural network). The slower it is. Jump every 100. The maximum value is 10, the default is 1.
- **Split data to train&test** - for example, the entered value of 0.8 means that the input data for the network will be divided in the ratio 0.8 to learning, 0.2 for the tests of the network learned.
- **Data limit** - limits the amount of data downloaded from the source. It speeds up the processing, but reduces its quality.
- **Max probes** - limits the number of samples taken to learn the network. Samples are already aggregated according to the selected “Time frame” parameter. It speed up teaching but reduces its quality.
- **Build** - a button to start teaching the network. The button contains the number of required teaching courses. You should be careful and avoid one-time learning for more than 1000 courses. It is better to divide them into several smaller ones. One pass after a full data load take about 1-3 minutes on a 4 core 2.4.GHz server. **The module has implemented the best practices related to the number of neurons in individual hidden layers. The values suggested by the system are optimal from the point of view of these practices, but the user can decide on these values himself.**

Under the parameters for learning the network there is an area in which teaching results will appear.

After pressing the “Refresh” button, the list of the resulting models will be refreshed.

Autorefresh - selecting the field automatically refreshes the list of learning results every 10s.

The following information will be available in the table on the left:

- **Internal name** - the model name given by the system, including the user - specified prefix
- **Overall efficiency** - the network adjustment indicator - allow to see at a glance whether it is worth dealing with the model. The greater the value, the better.

After clicking on the table row, detailed data collected during the learning of the given model will be displayed. This data will be visible in the box on the right.

The selected model can be saved under its own name using the “Save algorithm” button. This saved algorithm will be available in the “Choose AI Rule” list when creating the rule (see Create AI Rule).

AI Learn Tasks


The “AI Learn Task” tab shows the list of processes initiated teaching the ANN network with the possibility of managing processes.

Each user can see only the process they run. The user in the role of Intelligence sees all running processes.

Logged in as : logserver

[Create AI Rule](#)
[AI Rules List](#)
[AI Learn](#)
[AI Learn Tasks](#)

+AI Learn Tasks



Algorithm prefix	Progress	Processing time	Actions
ko2_	16 / 2	1272	Cancel Show
rpa_ann_3	0 / 2	0	Cancel Show Pause
rpa_ann_1_	0 / 2	0	Cancel Show Pause
rpa_ann_2_	0 / 2	0	Cancel Show

Description of controls:

- **Algorithm prefix** - this is the value set by the user on the AI Learn screen in the Prefix name field
- **Progress** - here is the number of algorithms generated / the number of all to be generated
- **Processing time** - duration of algorithm generation in seconds (or maybe minutes or hours)
- **Actions:**
 - **Cancel** - deletes the algorithm generation task (user require confirmation of operation)
 - **Pause / Release** - pause / resume algorithm generation process.

AI Learn tab contain the Show in the preview mode of the ANN hyperparameters After completing the learning activity or after the user has interrupted it, the “Delete” button appears in “Action” field. This button allows you to permanently

delete the learning results of a specific network.

Logged in as : logserver

[Create AI Rule](#) [AI Rules List](#) [AI Learn](#) **AI Learn Tasks**

+ AI Learn Tasks

Algorithm prefix	Progress	Processing time (s)	Actions
kk	0 / 4	0	Show Delete

Scenarios of using algorithms implemented in the Intelligence module

1. Teaching MLP networks and choosing the algorithm to use:
 - a. Go to the AI Learn tab,
 - b. We introduce the network teaching parameters,
 - c. Enter your own prefix for the names of the algorithms you have learned,
 - d. Press Build,
 - e. We observe the learned networks on the list (we can also stop the observation at any moment and go to other functions of the system. We will return to the learning results by going to the AI Learn Tasks tab and clicking the show action),
 - f. We choose the best model from our point of view and save it under our own name,
 - g. From this moment the algorithm is visible in the Create AI Rule tab.
2. Starting the MLP network algorithm:
 - a. Go to the Create AI Rule tab and create rules,
 - b. Select the previously saved model of the learned network,
 - c. Specify parameters visible on the screen (specific to MLP),
 - d. Press the Run button.
3. Starting regression algorithm:
 - a. Go to the Create AI Rule tab and create rules,
 - b. We choose AI Rule, e.g. Simple Moving Average, Linear Regression or Random Forest Regression, etc.,
 - c. Enter your own rule name (specific to regression),
 - d. Set the parameters of the rule (specific to regression),
 - e. Press the Run button.
4. Management of available rules:
 - a. Go to the AI Rules List tab,
 - b. A list of AI rules available for our role is displayed,
 - c. We can perform the actions available on the right for each rule.\

Results of algorithms

The results of the “AI algorithms” are saved to the index „intelligence” specially created for this purpose. The index with the prediction result. These following fields are available in the index (where xxx is the name of the attribute being analyzed):

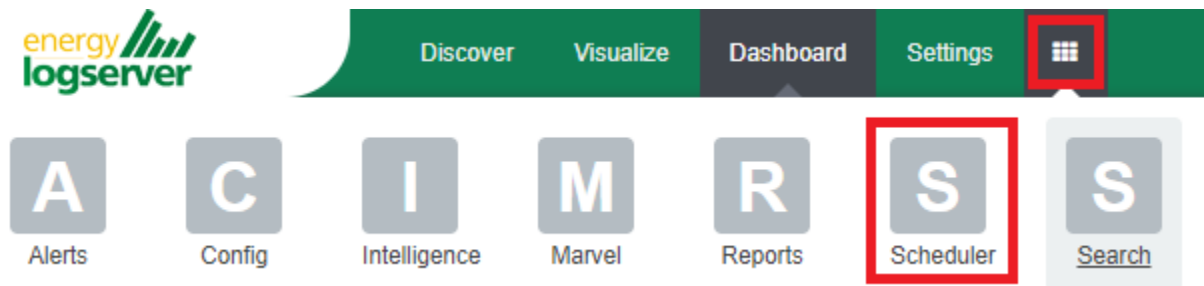
- **xxx_pre** - estimate value
- **xxx_cur** - current value at the moment of estimation
- **method_name** - name of the algorithm used
- **rmse** - average square error for the analysis in which _cur values were available. **The smaller the value, the better.**
- **rmse_normalized** - mean square error for the analysis in which _cur values were available, normalized with _pre values. **The smaller the value, the better.**
- **overall_efficiency** - efficiency of the model. **The greater the value, the better. A value less than 0 may indicate too little data to correctly calculate the indicator**
- **linear_function_a** - directional coefficient of the linear function $y = ax + b$. **Only for the Trend and Linear Regression Trend algorithm**
- **linear_function_b** - the intersection of the line with the Y axis for the linear function $y = ax + b$. **Only for the Trend and Linear Regression Trend algorithm.**

Visualization and signals related to the results of data analysis should be created from this index. The index should be available to users of the Intelligence module.

Scheduler Module

Energy Log Server has a built-in task schedule. In this module, we can define a command or a list of commands whose execution we instruct the application in the form of tasks. We can determine the time and frequency of tasks. Tasks can contain a simple syntax, but they can also be associated with modules, e.g. with Intelligence module.

To go to the Scheduler window, select the tile icon from the main menu bar and then go to the „Scheduler” icon (To go back, go to the „Search” icon)



The page with three tabs will be displayed: Creating new tasks in the „Create Scheduler Job”, managing tasks in the „Job List” and checking the status of tasks in „Jobs Status”

In the window for creating new tasks we have a form consisting of fields:

- **Name** - in which we enter the name of the task
- **Cron Pattern** - a field in which in cron notation we define the time and frequency of the task
- **Command** - we give the syntax of the command that will be executed in this task. These can be simple system commands, but also complex commands related to the Intelligence module. In the task management window, we can activate /deactivate, delete and update the task by clicking on the selected icon for a given task



In the task status windows you can check the current status of the task: if it activated, when it started and when it ended, how long it took. This window is not editable and indicates historical data.

Permission

Permission have been implemented in the following way:

- Only the user in the admin role can create / update rules.
- When creating rules, the roles that will be able to enables / disengage / view the rules will be indicated.

We assume that the Learn process works as an administrator.

We assume that the visibility of Search in AI Learn is preceded by receiving the search permission in the module object permission.

The role of “Intelligence” launches the appropriate tabs.

An ordinary user only sees his models. The administrator sees all models.

Verification of Elasticsearch service

To verify of Elasticsearch service you can use following command:

- Control of the Elastisearch system service via **systemd**: # `systemctl status elasticsearch` output:

```
`elasticsearch.service - Elasticsearch`
`Loaded: loaded (/usr/lib/systemd/system/elasticsearch.service; enabled; vendor_
→preset: disabled)`
`Active: active (running) since Fri 2017-09-01 10:36:52 CEST; 3h 7min ago`
`Docs: http://www.elastic.co`
`Main PID: 8362 (java)`
`CGroup: /system.slice/elasticsearch.service`
        └─8362 /bin/java -Xms256m -Xmx1g -Djava.awt.headless=true`
```

- Control of Elasticsearch instance via **tcp port**:

```
# curl -XGET '127.0.0.1:9200/'
```

output:

```
{
  "name" : "Henry Peter Gyrich",
  "cluster_name" : "elasticsearch",
  "version" : {
    "number" : "2.3.5",
    "build_hash" : "90f439ff60a3c0f497f91663701e64ccd01edbb4",
    "build_timestamp" : "2016-07-27T10:36:52Z",
    "build_snapshot" : false,
    "lucene_version" : "5.5.0"
  },
  "tagline" : "You Know, for Search"
}
```

- Control of Elasticsearch instance via **log file**:

```
# tail -f /var/log/elasticsearch/elasticsearch.log
```

- other control commands via **curl** application:

```
`curl -XGET "http://localhost:9200/_cat/health?v"`\n`curl -XGET "http://localhost:9200/_cat/nodes?v"`\n`curl -XGET "http://localhost:9200/_cat/indices?v"``
```

Verification of Logstash service

To verify of Logstash service you can use following command:

- control Logstash service via **systemd**:

```
# systemctl status logstash
```

output:

```
`logstash.service - logstash`
  `Loaded: loaded (/etc/systemd/system/logstash.service; enabled; vendor preset: _
  ↪disabled)`
  `Active: active (running) since Wed 2017-07-12 10:30:55 CEST; 1 months 23 days ago`
  `Main PID: 87818 (java)`
  `CGroup: /system.slice/logstash.service`
          └─87818 /usr/bin/java -XX:+UseParNewGC -XX:+UseConcMarkSweepGC`
```

- control Logstash service via **port tcp**:

```
# curl -XGET '127.0.0.1:9600'
```

output:

```
`{`
  ` "host": "skywalker",`
  ` "version": "4.5.3",`
  ` "http_address": "127.0.0.1:9600"`
`}`
```

- control Logstash service via **log file**:

```
# tail -f /var/log/logstash/logstash-plain.log
```

Debugging:

- dynamically update logging levels through the logging API (service restart not needed):

```
curl -XPUT 'localhost:9600/_node/logging?pretty' -H 'Content-Type:↵
↵application/json' -d'
{
  "logger.logstash.outputs.elasticsearch" : "DEBUG"
}
```

- permanent change of logging level (service need to be restarted):
- edit file `/etc/logstash/logstash.yml` and set the following parameter:

```
*`log.level: debug`*
```

- restart logstash service:

```
systemctl restart logstash
```

- checking correct syntax of configuration files:

```
/usr/share/logstash/bin/logstash -tf /etc/logstash/conf.d
```

- get information about load of the Logstash:

```
# curl -XGET '127.0.0.1:9600/_node/jvm?pretty=true'
```

output:

```
{
  "host" : "iosssmes1",
  "version" : "5.4.3",
  "http_address" : "100.127.111.14:9600",
  "id" : "defdf2f-2f5a-4c7e-ae03-e163e550bdb0",
  "name" : "iosssmes1",
  "jvm" : {
    "pid" : 3535,
    "version" : "1.8.0_131",
    "vm_name" : "OpenJDK 64-Bit Server VM",
    "vm_version" : "1.8.0_131",
    "vm_vendor" : "Oracle Corporation",
    "start_time_in_millis" : 1503057722599,
    "mem" : {
      "heap_init_in_bytes" : 268435456,
      "heap_max_in_bytes" : 1038876672,
      "non_heap_init_in_bytes" : 2555904,
      "non_heap_max_in_bytes" : 0
    },
    "gc_collectors" : [ "ParNew", "ConcurrentMarkSweep" ]
  }
}
```

Verificatoin of Energy Log Server GUI service

To verify of Energy Log Server GUI service you can use following command:

- control the Energy Log Server GUI service via **systemd**:

```
`# systemctl status kibana`
```

output:

```
`kibana.service - no description given`
`Loaded: loaded (/usr/lib/systemd/system/kibana.service; enabled; vendor preset:↵
↵disabled)`
`Active: active (running) since Fri 2017-09-01 10:39:46 CEST; 2 days ago`
`Main PID: 8527 (node)`
`CGroup: /system.slice/kibana.service`
        └─8527 /opt/kibana/bin/./node/bin/node /opt/kibana/bin/./src/cl`
```

- control the Energy Log Server GUI via **port tcp**:

```
`# curl -XGET '127.0.0.1:5601/'`
```

output:

```
`<script>var hashRoute = '/app/kibana';`
`var defaultRoute = '/app/kibana';`
`var hash = window.location.hash;`
`if (hash.length) {`
`  window.location = hashRoute + hash;`
`} else {`
`  window.location = defaultRoute;`
`}</script>`
```

- Control the Energy Log Server GUI via **log file**:

```
`# tail -f /var/log/messages`
```

Node roles

Every instance of Elasticsearch server is called a *node*. A collection of connected nodes is called a *cluster*. All nodes know about all the other nodes in the cluster and can forward client requests to the appropriate node. Besides that, each node serves one or more purpose:

- Master-eligible node - A node that has `node.master` set to `true` (default), which makes it eligible to be elected as the master node, which controls the cluster
- Data node - A node that has `node.data` set to `true` (default). Data nodes hold data and perform data related operations such as CRUD, search, and aggregations
- Client node - A client node has both `node.master` and `node.data` set to `false`. It can neither hold data nor become the master node. It behaves as a “smart router” and is used to forward cluster-level requests to the master node and data-related requests (such as search) to the appropriate data nodes
- Tribe node - A tribe node, configured via the `tribe.*` settings, is a special type of client node that can connect to multiple clusters and perform search and other operations across all connected clusters.

Naming convention

Elasticsearch require little configuration before going into work. The following settings must be considered before going to production:

- **path.data** and **path.logs** - default locations of these files are:

`/var/lib/elasticsearch` and `/var/log/elasticsearch`.

- **cluster.name** - A node can only join a cluster when it shares its `cluster.name` with all the other nodes in the cluster. The default name is “elasticsearch”, but you should change it to an appropriate name which describes the purpose of the cluster. You can do this in `/etc/elasticsearch/elasticsearch.yml` file.
- **node.name** - By default, Elasticsearch will use the first seven characters of the randomly generated UUID as the node id. Node id is persisted and does not change when a node restarts. It is worth configuring a more human readable name: `node.name: prod-data-2` in file `/etc/elasticsearch/elasticsearch.yml`
- **network.host** - parameter specifying network interfaces to which Elasticsearch can bind. Default is `network.host: ["_local_", "_site_"]`.
- **discovery** - Elasticsearch uses a custom discovery implementation called “Zen Discovery”. There are two important settings:
 - `discovery.zen.ping.unicast.hosts` - specify list of other nodes in the cluster that are likely to be live and contactable;
 - `discovery.zen.minimum_master_nodes` - to prevent data loss, you can configure this setting so that each master-eligible node knows the minimum number of master-eligible nodes that must be visible in order to form a cluster.
- **heap size** - By default, Elasticsearch tells the JVM to use a heap with a minimum (Xms) and maximum (Xmx) size of 1 GB. When moving to production, it is important to configure heap size to ensure that Elasticsearch has enough heap available

Config files

To configure the Elasticsearch cluster you must specify some parameters in the following configuration files on every node that will be connected to the cluster:

- `/etc/elasticsearch/elasticsearch.yml`:

```
- `cluster.name:name_of_the_cluster` - same for every node;  
- `node.name:name_of_the_node` - uniq for every node;  
- `node.master:true_or_false`  
- `node.data:true_or_false`  
- `network.host:["_local_", "_site_"]`  
- `discovery.zen.ping.multicast.enabled`  
- `discovery.zen.ping.unicast.hosts`  
- `index.number_of_shards`
```

```
- index.number_of_replicas
```

- `/etc/elasticsearch/logging.yml`:

```
logger:  action:  DEBUG - for easier debugging.
```


Example setup

Example of the Elasticsearch cluster configuration:

- file `/etc/elasticsearch/elasticsearch.yml`:

```
cluster.name: tm-lab node.name: "elk01" node.master: true *node.data: true network.host:
127.0.0.1,10.0.0.4 http.port: 9200 discovery.zen.ping.multicast.enabled: false discovery.zen.ping.unicast.hosts:
["10.0.0.4:9300","10.0.0.5:9300","10.0.0.6:9300"] index.number_of_shards: 3 index.number_of_replicas: 1
```

- to start the Elasticsearch cluster execute command:

```
# systemctl restart elasticsearch`
```

- to check status of the Elasticsearch cluster execute command:

- check of the Elasticsearch cluster nodes status via tcp port:

```
# curl -XGET '127.0.0.1:9200/_cat/nodes?v'
host          ip          heap.percent ram.percent load node.role
↪master name
10.0.0.4      10.0.0.4      18           91           0.00 -      -
↪elk01
10.0.0.5      10.0.0.5      66           91           0.00 d      *
↪elk02
10.0.0.6      10.0.0.6      43           86           0.65 d      m
↪elk03
10.0.0.7      10.0.0.7      45           77           0.26 d      m
↪elk04
```

- check status of the Elasticsearch cluster via log file:\

```
tail -f /var/log/elasticsearch/tm-lab.log (cluster.name)
```

Integration with AD

You can configure the Energy Log Server to communicate with Active Directory to authenticate users. To integrate with Active Directory, you configure an Active Directory realm and assign Active Directory users and groups to the Energy Log Server roles in the role mapping file.

To protect passwords, communications between the Energy Log Server and the LDAP server should be encrypted using SSL/TLS. Clients and nodes that connect via SSL/TLS to the LDAP server need to have the LDAP server's certificate or the server's root CA certificate installed in their keystore or truststore.

CHAPTER 64

AD configuration

The AD configuration should be done in the `/etc/elasticsearch/elasticsearch.yml` file.

Below is a list of settings to be made in the `elasticsearch.yml` file (the commented section in the file in order for the AD settings to start working, this fragment should be uncommented):

```
|**Directive**|**Description**|
|-----|-----|
| # LDAP|
| #ldaps:|
| # - name: \"example.com\"|/# domain that is configured|
| # host: \"127.0.0.1,127.0.0.2\"|/# list of server for this domain|
| # port: 389|/# optional, default 389 for|
|unencrypted session or 636 for encrypted sessions|
|/# ssl_enabled: false|/# optional, default true|
|/# ssl_trust_all_certs: true|/# optional, default false|
|/# ssl.keystore.file: \"path\"|/# path to the truststore store|
|/# ssl.keystore.password: \"path\"|/# password to the trusted|
|certificate store|
|/# bind_dn: [admin@example.com]|/# account name administrator|
|/# bind_password: \"password\"|/# password for the administrator|
|account|
|/# search_user_base_DN: \"OU=lab,DC=example,DC=com\"|/# search for the DN user|
|tree database|
|/# user_id_attribute: \"uid\"|/# search for a user attribute|
|optional, by default \"uid\"|
```

(continues on next page)

(continued from previous page)

```
|# search\_groups\_base\_DN: \"OU=lab,DC=example,DC=com\"|# group database search.↵
↵This is a catalog main, after which the groups will be sought.|
|# unique\_member\_attribute: \"uniqueMember\"      |# optional, default\
↵\"uniqueMember\"      |
|# connection\_pool\_size: 10                        |# optional, default 30      ↵
↵      |
|# connection\_timeout\_in\_sec: 10                  |# optional, default 1      ↵
↵      |
|# request\_timeout\_in\_sec: 10                    |# optional, default 1      ↵
↵      |
|# cache\_ttl\_in\_sec: 60                          |# optional, default 0 -↵
↵cache disabled      |
```

If we want to configure multiple domains, then in this configuration file we copy the # LDAP section below and configure it for the next domain.

Below is an example of how an entry for 2 domains should look like. (It is important to take the interpreter to read these values correctly).



```
# LDAP
ldaps:
- name: "example1.com" # DOMENA1
  host: "127.0.0.1,127.0.0.2"
  port: 389 # optional, default 389
  ssl_enabled: false # optional, default true
  ssl_trust_all_certs: true # optional, default false
  ssl.keystore.file: "path1"
  ssl.keystore.password: "path1"
  bind_dn: "admin@example1.com"
  bind_password: "password1"
  search_user_base_DN: "OU=lab,DC=example1,DC=com"
  user_id_attribute: "uid" # optional, default "uid"
  search_groups_base_DN: "OU=lab,DC=example1,DC=com"
  unique_member_attribute: "uniqueMember" # optional, default "uniqueMember"
  connection_pool_size: 10 # optional, default 30
  connection_timeout_in_sec: 10 # optional, default 1
  request_timeout_in_sec: 10 # optional, default 1
  cache_ttl_in_sec: 60 # optional, default 0 - cache disabled
- name: "example2.com" # DOMENA2
  host: "127.0.0.1,127.0.0.2"
  port: 389 # optional, default 389
  ssl_enabled: false # optional, default true
  ssl_trust_all_certs: true # optional, default false
  ssl.keystore.file: "path2"
  ssl.keystore.password: "path2"
  bind_dn: "admin@example2.com"
  bind_password: "password2"
  search_user_base_DN: "OU=lab,DC=example2,DC=com"
  user_id_attribute: "uid" # optional, default "uid"
  search_groups_base_DN: "OU=lab,DC=example2,DC=com"
  unique_member_attribute: "uniqueMember" # optional, default "uniqueMember"
  connection_pool_size: 10 # optional, default 30
  connection_timeout_in_sec: 10 # optional, default 1
  request_timeout_in_sec: 10 # optional, default 1
  cache_ttl_in_sec: 60 # optional, default 0 - cache disabled
```

After completing the LDAP section entry in the `elasticsearch.yml` file, save the changes and restart the service with the command:

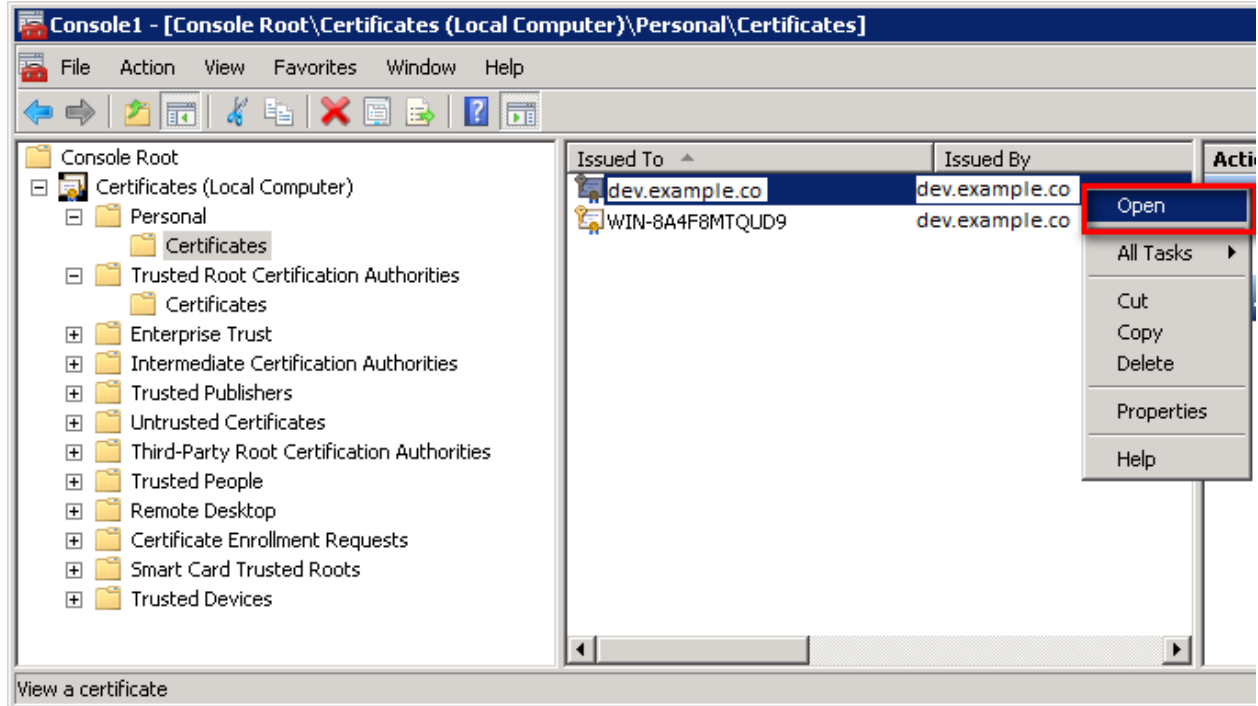
```
# systemctl restart elasticsearch
```


Configure SSL support for AD authentication.

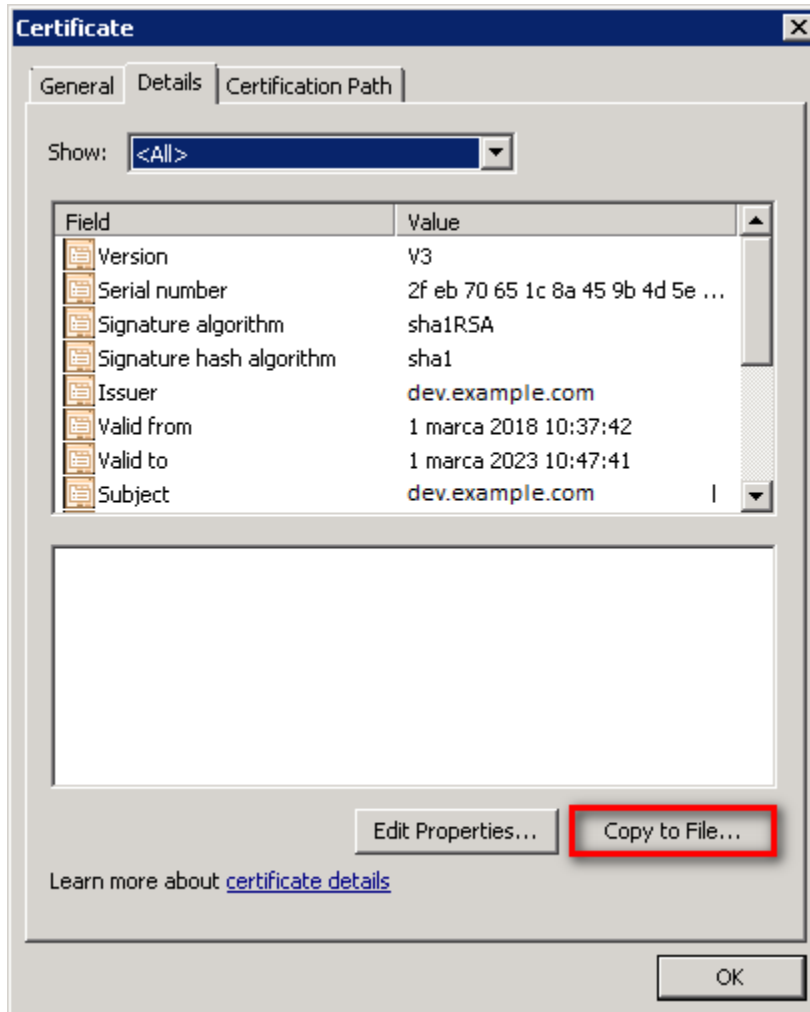
1. Open the certificate manager on the AD server.




1. Select the certificate and open it



1. Select the option of copying to a file in the Details tab



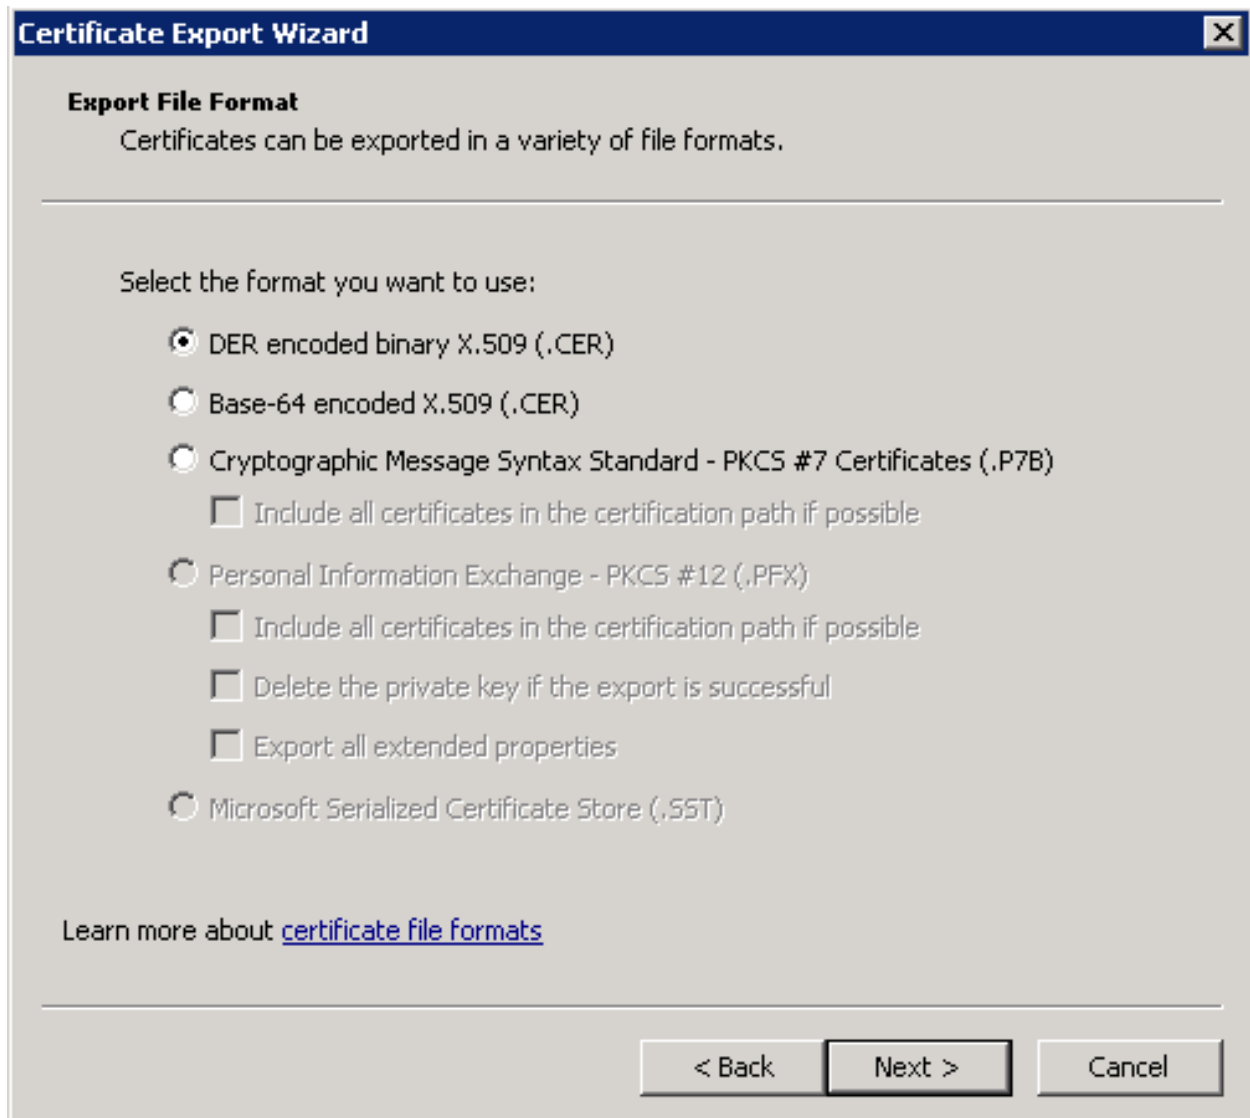
1. Click the Next button




1. Keep the setting as shown below and click Next



1. Keep the setting as shown below and click Next.



1. Give the name a certificate



After the certificate is exported, this certificate should be imported into a trusted certificate file that will be used by the Elasticsearch plugin.

To import a certificate into a trusted certificate file, a tool called „keytool.exe” is located in the JDK installation directory.

Use the following command to import a certificate file:

```
keytool -import -alias adding_certificate_keystore -file certificate.cer
-keystore certificatestore
```

The values for RED should be changed accordingly.

By doing this, he will ask you to set a password for the trusted certificate store. Remember this password, because it must be set in the configuration of the elasticsearch plugin. The following settings must be set in the `elasticsearch.yml` configuration for `SSL:ssl.keystore.file`: "<path to the trust certificate store>"

```
ssl.keystore.password: "< password to the trust certificate store>"
```


CHAPTER 66

Role mapping.

In the `/etc/elasticsearch/elasticsearch.yml` configuration file you can find a section for configuring role mapping:

```
# LDAP ROLE MAPPING FILE
# rolemapping.file.path: /etc/elasticsearch/role-mappings.yml
```

This variable points to the file `/etc/elasticsearch/role-mappings.yml`. Below is the sample content for this file:

admin:

- "CN=Admins,OU=lab,DC=dev,DC=it,DC=example,DC=com"

bank:

- "CN=security,OU=lab,DC=dev,DC=it,DC=example,DC=com"

The mapping mechanism works in this way: An AD user logs in to Energy Log Server. In the application, there is an admin role, which through the file `role-mapping.yml` binds to the name of the admin role to which the Admins container from AD is assigned. It is enough for the user from the AD account to log in to the application with the privileges that are assigned to the admin role in the Energy Log Server. At the same time, if it is the first login in the Energy Log Server, an account is created with an entry that informs the application administrator that it was created by logging in with AD.

Similar, the mechanism will work if we have a role with an arbitrary name created in Energy Log Server Logistics and connected to the name of the `role-mappings.yml` and existing in AD any container.

Below is a screenshot of the console on which are marked accounts that were created by users logging in from AD

Logged in as : logserver

User Management Settings License Info

Create User User List Create Role Role List Objects permission

User List

Username	Roles	Default Role	Actions
alert	admin,		Delete Update
user1@example.com	adrole,	adrole	Delete Update
audit	authsystem,alert,test_role_for_ad,	test_role_for_ad	Delete Update
intelligence	admin,		Delete Update
logserver	admin,		Delete Update
scheduler	admin,		Delete Update
user2@example.com	adrole,	adrole	Delete Update

If you map roles with from several domains, for example dev.examlloe1.com, dev.example2.com then in User List we will see which user from which domain with which role logged in Energy Log Server.

Configuring Single Sign On (SSO)

In order to configure SSO, the system should be accessible by domain name URL, not IP address nor localhost.

Ok : `https://loggui.com:5601/login`. **Wrong :** `https://localhost:5601/login`, `https://10.0.10.120:5601/login`

In order to enable SSO on your system follow below steps. The configuration is made for AD: `dev.example.com`, GUI URL: `loggui.com`

67.1 Create an User Account for Elasticsearch auth plugin

In this step, a Kerberos Principal representing Elasticsearch auth plugin is created on the Active Directory. The principal name would be `name@DEV.EXAMPLE.COM`, while the `DEV.EXAMPLE.COM` is the administrative name of the realm. In our case, the principal name will be `esauth@DEV.EXAMPLE.COM`.

Create User in AD. Set “Password never expires” and “Other encryption options” as shown below:

67.2 Define Service Principal Name (SPN) and Create a Keytab file for it

Use the following command to create the keytab file and SPN:

```
C:> ktpass -out c:\Users\Administrator\esauth.keytab -princ HTTP/loggui.com@DEV.EXAMPLE.COM
-mapUser esauth -mapOp set -pass 'Sprint$123' -crypto ALL -pType KRB5_NT_PRINCIPAL
```

Values highlighted in bold should be adjusted for your system. The **esauth.keytab** file should be placed on your elasticsearch node - preferably `/etc/elasticsearch/` with read permissions for elasticsearch user: `chmod 640 /etc/elasticsearch/esauth.keytab chown elasticsearch: /etc/elasticsearch/esauth.keytab`

67.3 Create a file named *krb5Login.conf*:

```
com.sun.security.jgss.initiate{
  com.sun.security.auth.module.Krb5LoginModule required
  principal="esauth@DEV.EXAMPLE.COM" useKeyTab=true
  keyTab=/etc/elasticsearch/esauth.keytab storeKey=true debug=true;
};
com.sun.security.jgss.krb5.accept {
  com.sun.security.auth.module.Krb5LoginModule required
  principal="esauth@DEV.EXAMPLE.COM" useKeyTab=true
  keyTab=/etc/elasticsearch/esauth.keytab storeKey=true debug=true;
};
```

Principal user and keyTab location should be changed as per the values created in the step 2. Make sure the domain is in UPPERCASE as shown above. The `krb5Login.conf` file should be placed on your elasticsearch node, for instance `/etc/elasticsearch/` with read permissions for elasticsearch user:

```
sudo chmod 640 /etc/elasticsearch/krb5Login.conf
sudo chown elasticsearch: /etc/elasticsearch/krb5Login.conf
```

67.4 Append the following JVM arguments (on Elasticsearch node in `/etc/sysconfig/elasticsearch`):

```
-Dsun.security.krb5.debug=true          -Djava.security.krb5.realm=DEV.EXAMPLE.COM -
Djava.security.krb5.kdc=AD_HOST_IP_ADDRESS -Djava.security.auth.login.config=/etc/elasticsearch/krb5Login.conf
-Djavax.security.auth.useSubjectCredsOnly=false
```

Change the appropriate values in the bold. This JVM arguments has to be set for Elasticsearch server.

67.5 Add the following additional (`sso.domain`, `service_principal_name`, `service_principal_name_password`) settings for ldap in `elasticsearch.yml` or `properties.yml` file wherever the ldap settings are configured:

```
sso.domain: "dev.example.com"
ldaps:
- name: "dev.example.com"
  host: "IP_address"
  port: 389                                # optional, default 389
  ssl_enabled: false                      # optional, default 
  ↪ true
    ssl_trust_all_certs: false            # optional, default 
  ↪ false
    bind_dn: "Administrator@dev.example.com" # optional, skip for 
  ↪ anonymous bind
    bind_password: "administrator_password" # 
  ↪ optional, skip for anonymous bind
    search_user_base_DN: "OU=lab,DC=dev,DC=it,DC=example,DC=com"
    user_id_attribute: "uid"              # optional, default "uid"
  ↪ "
    search_groups_base_DN: "OU=lab,DC=dev,DC=it,DC=example,DC=com"
    unique_member_attribute: "uniqueMember" # optional, default
  ↪ "uniqueMember"
    service_principal_name: "esauth@DEV.EXAMPLE.COM"
    service_principal_name_password : "Sprint$123"
```

Note: At this moment, SSO works for only single domain. So you have to mention for what domain SSO should work in the above property `sso.domain`

67.6 To apply the changes restart Elasticsearch service

```
sudo systemctl restart elasticsearch.service
```

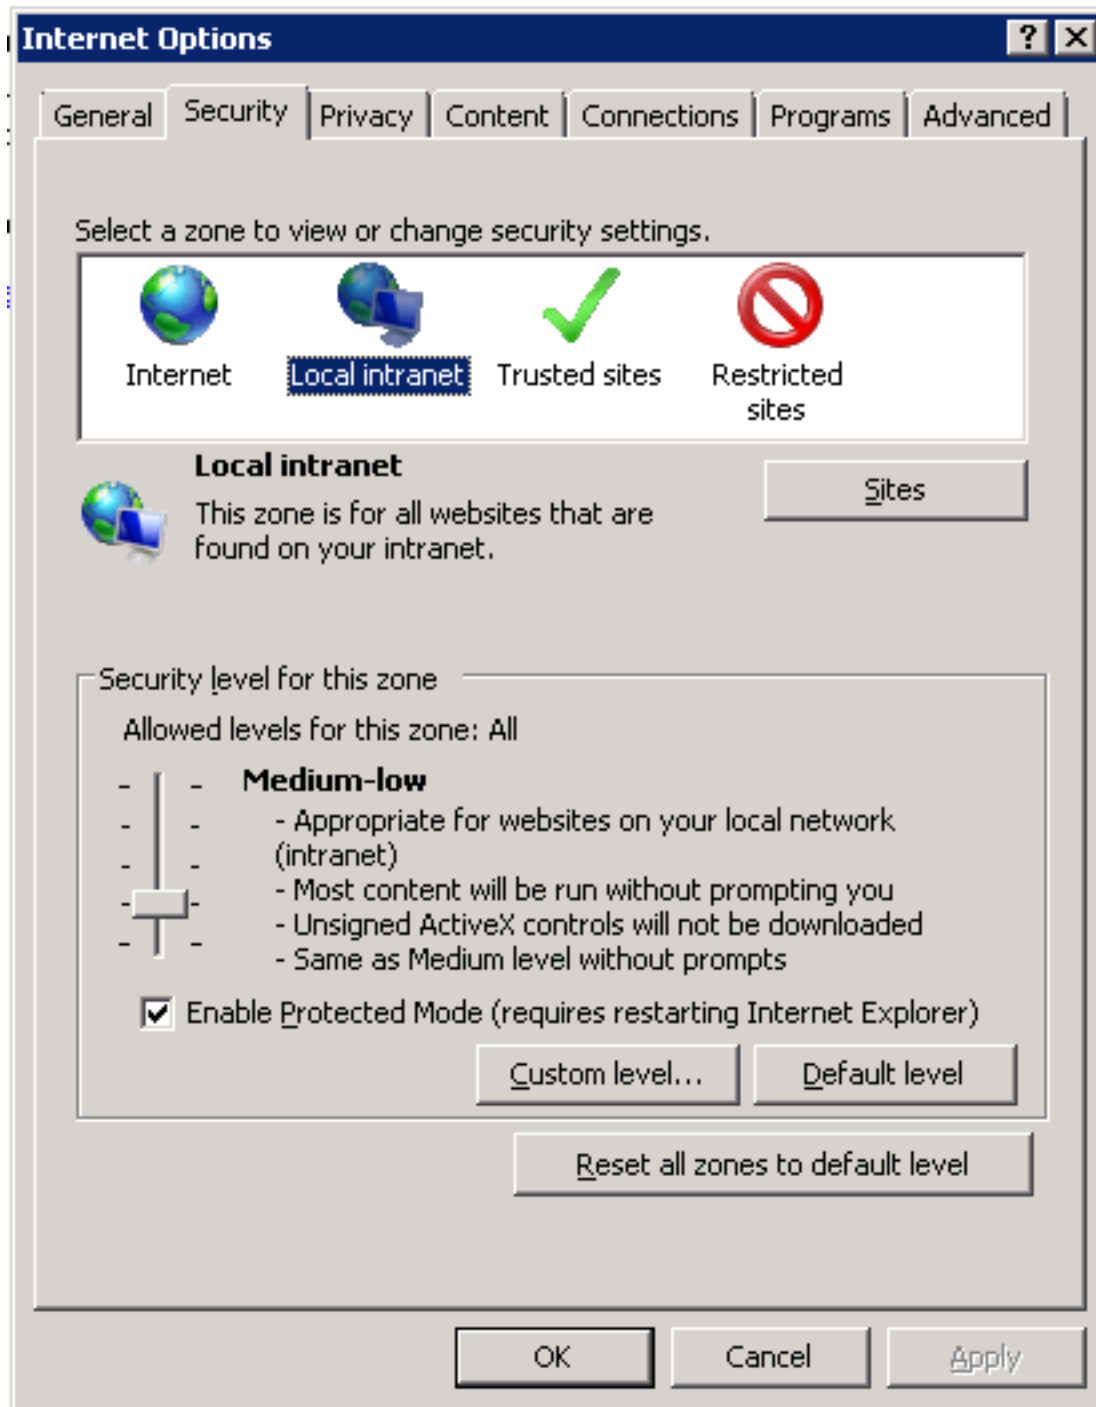
67.7 Enable SSO feature in `kibana.yml` file:

```
kibana.sso_enabled: true After that Kibana has to be restarted: sudo systemctl restart kibana.service
```


Client (Browser) Configuration

68.1 Internet Explorer configuration

1. Goto Internet Options from Tools menu and click on Security Tab:




1. Select Local intranet, click on Site -> Advanced -> Add the url:



After adding the site click close.

1. Click on custom level and select the option as shown below:

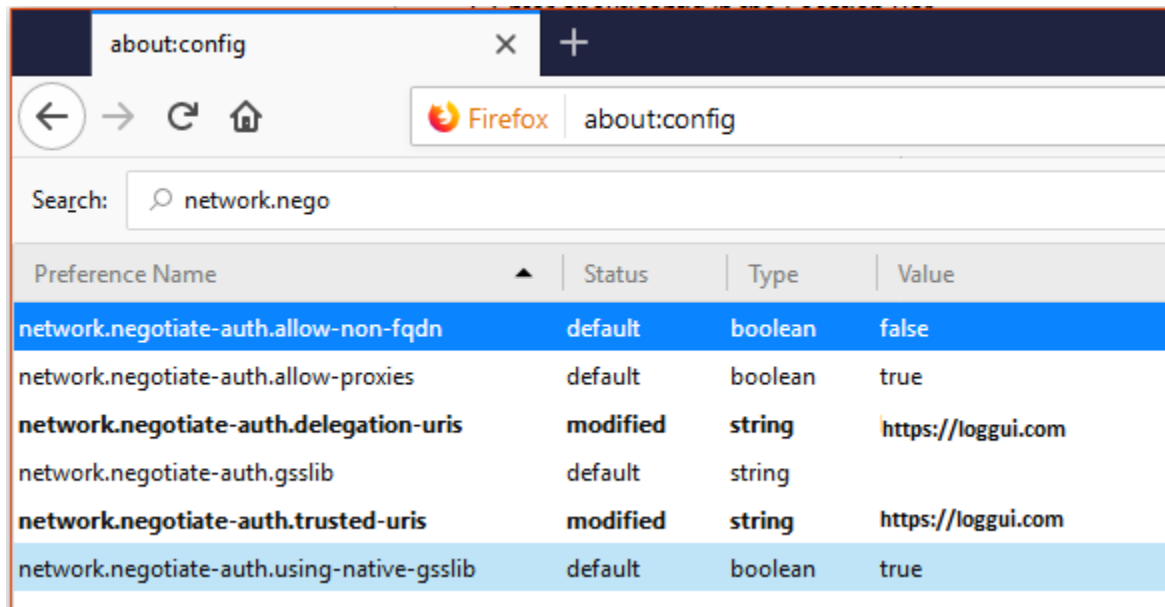


68.2 Chrome configuration

For Chrome, the settings are taken from IE browser.

68.3 Firefox configuration

Update the following config:



Configure email delivery for sending PDF reports in Scheduler.

The default e-mail client that installs with the Linux CentOS system, which is used by Energy Log Server to send reports (Section 5.3 of the [Reports](#) chapter), is *postfix*.

Configuration file for **postfix** mail client

The *postfix* configuration directory for CentOS is */etc/postfix*. It contains files:

main.cf - the main configuration file for the program specifying the basics parameters

Some of its directives:

Directive	**Description**
queue_directory	The postfix queue location.
command_directory	The location of Postfix commands.
daemon_directory	Location of Postfix daemons.
mail_owner	The owner of Postfix domain name of the server
myhostname	The fully qualified domain name of the server.
mydomain	Server domain
myorigin	Host or domain to be displayed as origin on email leaving the server.
inet_interfaces	Network interface to be used for incoming email.
mydestination	Domains from which the server accepts mail.
mynetworks	The IP address of trusted networks.
relayhost	Host or other mail server through which mail will be sent. This server will act as an outbound gateway.
alias_maps	Database of aliases used by the local delivery agent.
alias_database	Alias database generated by the new aliases.
command	
mail_spool_directory	The location where user boxes will be stored.

master.cf - defines the configuration settings for the master daemon and the way it should work with other agents to deliver mail. For each service installed in the master.cf file there are seven columns that define how the service should be used.

Column	Description

(continues on next page)

(continued from previous page)

service		The name of the service
type		The transport mechanism to be user.
private		Is the service only for user by Postfix.
unpriv		Can the service be run by ordinary users
chroot		Whether the service is to change the main directory (chroot) ↳
↳for the mail. Queue.		
wakeup		Wake up interval for the service.
maxproc		The maximum number of processes on which the service can be ↳
↳forked (to divide in branches)		
command + args		A command associated with the service plus any argument

access - can be used to control access based on e-mail address, host address, domain or network address.

Examples of entries in the file

Description	Example
-----	-----
To allow access for specific IP address:	192.168.122.20 OK
To allow access for a specific domain:	example.com OK
To deny access from the 192.168.3.0/24 network:	192.168.3 REJECT

After making changes to the access file, you must convert its contents to the access.db database with the postmap command:

```
# postmap /etc/postfix/access
# ll /etc/postfix/access*

-rw-r--r--. 1 root root 20876 Jan 26 2014 /etc/postfix/access
-rw-r--r--. 1 root root 12288 Feb 12 07:47 /etc/postfix/access.db
```

canonical - mapping incoming e-mails to local users.

Examples of entries in the file:

To forward emails to user1 to the [[user1@yahoo.com] mailbox:

```
user1 user1\@yahoo.com
```

To forward all emails for example.org to another example.com domain:

```
@example.org @example.com
```

After making changes to the canonical file, you must convert its contents to the canonical.db database with the postmap command:

```
# postmap /etc/postfix/canonical
# ll /etc/postfix/canonical*

-rw-r--r--. 1 root root 11681 2014-06-10 /etc/postfix/canonical
-rw-r--r--. 1 root root 12288 07-31 20:56 /etc/postfix/canonical.db
```

generic - mapping of outgoing e-mails to local users. The syntax is the same as a canonical file. After you make change to this file, you must also run the postmap command.

```
# postmap /etc/postfix/generic
# ll /etc/postfix/generic*
```

(continues on next page)

(continued from previous page)

```
-rw-r--r--. 1 root root 9904 2014-06-10 /etc/postfix/generic
-rw-r--r--. 1 root root 12288 07-31 21:15 /etc/postfix/generic.db
```

relocated – information about users who have been transferred. The syntax of the file is the same as canonical and generic files.

Assuming tha user1 was moved from example.com to example.net, you can forward all emails received on the old address to the new address:

Example of an entry in the file:

```
user1@example.com user1@example.net
```

After you make change to this file, you must also run the postmap command.

```
# postmap /etc/postfix/relocated
# ll /etc/postfix/relocated*

-rw-r--r--. 1 root root 6816 2014-06-10 /etc/postfix/relocated
-rw-r--r--. 1 root root 12288 07-31 21:26 /etc/postfix/relocated.d
```

transport – mapping between e-mail addresses and server through which these e-mails are to be sent (next hops) int the transport format: nexthop.

Example of an entry in the file:

```
user1@example.com smtp:host1.example.com
```

After you make changes to this file, you must also run the postmap command.

```
# postmap /etc/postfix/transport
[root@server1 postfix]# ll /etc/postfix/transport*

-rw-r--r--. 1 root root 12549 2014-06-10 /etc/postfix/transport
-rw-r--r--. 1 root root 12288 07-31 21:32 /etc/postfix/transport.db
```

virtual - user to redirect e-mails intended for a certain user to the account of another user or multiple users. It can also be used to implement the domain alias mechanism.

Examples of the entry in the file:

Redirecting email for user1, to root users and user3:

```
user1 root, user3
```

Redirecting email for user 1 in the example.com domain to the root user:

```
user1@example.com root
```

After you make change to this file, you must also run the postmap command:

```
# postmap /etc/postfix/virtual
# ll /etc/postfix/virtual

-rw-r--r--. 1 root root 12494 2014-06-10 /etc/postfix/virtual
-rw-r--r--. 1 root root 12288 07-31 21:58 /etc/postfix/virtual.db
```

Basic postfix configuration

Base configuration of *postfix* application you can make in `/etc/postfix/main.cfg` configuration file, which must complete with the following entry:

- section # *RECEIVING MAIL*

```
`inet\_interfaces = all`  
`inet\_protocols = ipv4`
```

- section # *INTERNET OR INTRANET*

```
relayhost = \[IP mail server\]:25 (port number)
```

I the next step you must complete the canonical file of *postfix*

At the end you should restart the *postfix*:

```
`systemctl restart postfix`
```

Example of postfix configuration with SSL encryption enabled

To configure email delivery with SSL encryption you need to make the following changes in the *postfix* configuration files:

- **/etc/postfix/main.cf** - file should contain the following entries in addition to standard (unchecked entries):

```
- mydestination = $myhostname, localhost.$mydomain, localhost
- myhostname = example.com
- relayhost = [smtp.example.com]:587
- smtp_sasl_auth_enable = yes
- smtp_sasl_password_maps = hash:/etc/postfix/sasl_passwd
- smtp_sasl_security_options = noanonymous
- smtp_tls_CAfile = /root/certs/cacert.cer
- smtp_use_tls = yes
- smtp_sasl_mechanism_filter = plain, login
- smtp_sasl_tls_security_options = noanonymous
- canonical_maps = hash:/etc/postfix/canonical
- smtp_generic_maps = hash:/etc/postfix/generic
- smtpd_recipient_restrictions = permit_sasl_authenticated
```

- **/etc/postfix/sasl/passwd** - file should define the data for authorized

```
*`[smtp.example.com\]:587 [[USER@example.com:PASS]] (mailto:USER@example.
↪com:PASS)`*
```

You need to give appropriate permissions:

```
*`chmod 400 /etc/postfix/sasl_passwd`*
```

and map configuration to database:

```
*`postmap /etc/postfix/sasl_passwd`*
```

next you need to generate a ca cert file:

```
*`cat /etc/ssl/certs/Example\_Server\_CA.pem | tee -a etc/postfix/cacert.pem`*
```

And finally, you need to restart postfix

```
*`/etc/init.d/postfix restart`*
```

Elasticsearch API

The Elasticsearch has a typical REST API and data is received in JSON format after the HTTP protocol. By default the tcp/9200 port is used to communicate with the Elasticsearch API. For purposes of examples, communication with the Elasticsearch API will be carried out using the *curl* application.

Program syntax:

```
*` # curl -X<metoda> -u login:password '<adres_ip_elasticsearch>:<port>'\n*` # curl -XGET -u login:password '127.0.0.1:9200'`*
```

Available methods:

- PUT - sends data to the server;
- POST - sends a request to the server for a change;
- DELETE - deletes the index / document;
- GET - gets information about the index /document;
- HEAD - is used to check if the index / document exists.

Avilable APIs by roles:

- Index API - manages indexes;
- Document API - manges documnets;
- Cluster API - manage the cluster;
- Search API - is used to search for data.

CHAPTER 74

Elasticsearch Index API

The indices APIs are used to manage individual indices, index settings, aliases, mappings, and index templates.

Elasticsearch Index API - Adding Index

Adding Index - automatic method:

```
*`# curl -XPUT -u login:password '127.0.0.1:9200/twitter/tweet/1?pretty=true' -d'{`*\`
  *`"user" : "elk01",`*\`
  *`"post_date" : "2017-09-05T10:00:00",`*\`
  *`"message" : "tests auto index generation"`*\`
  *`}'`*`
```

You should see the output:

```
*`{`*\`
*`"_index" : "twitter",`*\`
*`"_type" : "tweet",`*\`
*`"_id" : "1",`*\`
*`"_version" : 1,`*\`
*`"_shards" : {`*\`
  *`"total" : 2,`*\`
  *`"successful" : 1,`*\`
  *`"failed" : 0`*\`
*`},`*\`
*`"created" : true`*\`
*`}``*`
```

The parameter `action.auto_create_index` must be set on `true`.

Adding Index – manual method:

- settings the number of shards and replicas:

```
# curl -XPUT -u login:password '127.0.0.1:9200/twitter2?pretty=true'
-d'{"settings" : {"number_of_shards" : 1,"number_of_replicas" : 1}}`*`
```

You should see the output:

```
*`{*`*\n  *`"acknowledged" : true`*\n*`}``*
```

- command for manual index generation:

```
# curl -XPUT -u login:password '127.0.0.1:9200/twitter2/tweet/1?pretty=true' -d'{"user" : "elk01","post_date" : "2017-09-05T10:00:00", "message" : "tests manual index generation"}'
```

You should see the output:

```
*`{*`*\n  *`"_index" : "twitter2",`*\n  *`"_type" : "tweet",`*\n  *`"_id" : "1",`*\n  *`"_version" : 1,`*\n  *`"_shards" : {`*\n    *`"total" : 2,`*\n    *`"successful" : 1,`*\n    *`"failed" : 0`*\n  *`},`*\n  *`"created" : true`*\n*`}``*
```

Elasticsearch Index API

Delete Index - to delete *twitter* index you need use the following command:

```
# curl -XDELETE -u login:password '127.0.0.1:9200/twitter?pretty=true'
```

The delete index API can also be applied to more than one index, by either using a comma separated list, or on all indices by using `_all` or `*` as index:

```
# curl -XDELETE -u login:password '127.0.0.1:9200/twitter*?pretty=true'
```

To allowing to delete indices via wildcards set `action.destructive_requires_name` setting in the config to `false`.

Elasticsearch Index_API useful commands.

- get information about Replicas and Shards:

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_settings?pretty=true'#  
curl -XGET -u login:password '127.0.0.1:9200/twitter2/_settings?pretty=true'
```

- get information about mapping and alias in the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_mappings?pretty=true'#  
curl -XGET -u login:password '127.0.0.1:9200/twitter/_aliases?pretty=true'
```

- get all information about the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter?pretty=true'
```

- checking does the index exist:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter?pretty=true'
```

- close the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_close?pretty=true'
```

- open the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_open?pretty=true'
```

- get the status of all indexes:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/indices?v'
```

- get the status of one specific index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/indices/twitter?v'
```

- display how much memory is used by the indexes:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/indices?v&h=i,  
tm&s=tm:desc'
```

- display details of the shards:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/shards?v'
```

CHAPTER 78

Elasticsearch Document API

Elasticsearch Document API - Create_Document

- create a document with a specify ID:

```
# curl -XPUT -u login:password '127.0.0.1:9200/twitter/tweet/1?
pretty=true' -d '{"user" : "lab1", "post_date" : "2017-08-25T10:00:00",
"message" : "testuje Elasticsearch"}'
```

You should see the output:

```
*`{`*\
*`"_index" : "twitter",`*\
*`"_type" : "tweet",`*\
*`"_id" : "1",`*\
*`"_version" : 1,`*\
*`"_shards" : {`*\
*`"total" : 2,`*\
*`"successful" : 1,`*\
*`"failed" : 0`*\
*`},`*\
*`"created" : true`*\
*`}``*`
```

- creating a document with an automatically generated ID: (note: PUT-> POST):\

```
*`# curl -XPOST -u login:password '127.0.0.1:9200/twitter/tweet?pretty=true' -d'
↪{`*\
*`"user" : "lab1",`*\
*`"post_date" : "2017-08-25T10:10:00",`*\
*`"message" : "testuje automatyczne generowanie ID"`*\
*`}``*`
```

You should see the output:\

```
*`{`*\
*`"_index" : "twitter",`*\
*`"_type" : "tweet",`*\
```

(continues on next page)

(continued from previous page)

```
*`"_id" : "AV49sTlM8NzerkV9qJfh",`*\n*`"_version" : 1,`*\n*`"_shards" : {`*\n  *`"total" : 2,`*\n  *`"successful" : 1,`*\n  *`"failed" : 0`*\n*`},`*\n*`"created" : true`*\n*`}``*
```

Elasticsearch Document API - Delete Document

- delete a document by ID:\

```
# curl -XDELETE -u login:password '127.0.0.1:9200/twitter/tweet/1?pretty=true'# curl -XDELETE -u login:password '127.0.0.1:9200/twitter/tweet/AV49sTlM8NzerkV9qJfh?pretty=true'
```

- delete a document using a wildcard:\

```
# curl -XDELETE -u login:password '127.0.0.1:9200/twitter/tweet/1*?pretty=true'
```

(parametr: action.destructive_requires_name must be set to false)

Elasticsearch Document API - useful commands

- get information about the document:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/tweet/1?pretty=true'
```

You should see the output:\

```
*`{`*\
  *`"_index" : "twitter",`*\
  *`"_type" : "tweet",`*\
  *`"_id" : "1",`*\
  *`"_version" : 1,`*\
  *`"found" : true,`*\
  *`"_source" : {`*\
    *`"user" : "lab1",`*\
    *`"post_date" : "2017-08-25T10:00:00",`*\
    *`"message" : "testuje Elasticsearch"`*\
  *`}`*\
```

```
}
```

- get the source of the document:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/tweet/1/_source?pretty=true'
```

You should see the output:\

```
*`{`*\
  *`"user" : "lab1",`*\
  *`"post_date" : "2017-08-25T10:00:00",`*\
  *`"message" : "testuje Elasticsearch"`*\
```

```
}\\
```

- get information about all documents in the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter*/_search?
q=* &pretty=true'
```

You should see the output:

```
{
  "took" : 7,
  "timed_out" : false,
  "_shards" : {
    "total" : 10,
    "successful" : 10,
    "failed" : 0
  },
  "hits" : {
    "total" : 3,
    "max_score" : 1.0,
    "hits" : [ {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "AV49sTlM8NzerkV9qJfh",
      "_score" : 1.0,
      "_source" : {
        "user" : "lab1",
        "post_date" : "2017-08-25T10:10:00",
        "message" : "auto generated ID"
      }
    }, {
      "_index" : "twitter",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 1.0,
      "_source" : {
        "user" : "lab1",
        "post_date" : "2017-08-25T10:00:00",
        "message" : "Elasticsearch test"
      }
    }, {
      "_index" : "twitter2",
      "_type" : "tweet",
      "_id" : "1",
      "_score" : 1.0,
      "_source" : {
        "user" : "elk01",
        "post_date" : "2017-09-05T10:00:00",
        "message" : "manual index created test"
      }
    } ]
  }
}
```

- the sum of all documents in a specified index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/count/twitter?v'
```

You should see the output:\

```
*epoch          timestamp count*\
*1504281400      17:56:40      2*
```

- the sum of all document in Elasticsearch database:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/count?v'
```

You should see the output:\

<pre>*`epoch timestamp count`*\n*`1504281518 17:58:38 493658`*</pre>

Elasticsearch Cluster API

Example of use:

- information about the cluster state:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cluster/health?pretty=true'
```

- information about the role and load of nodes in the cluster:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/nodes?v'
```

- information about the available and used place on the cluster nodes:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/allocation?v'
```

- information which node is currently in the master role:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/master?v'
```

- information about currently performed operations by the cluster:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/pending_tasks?v'
```

- information on recoveries / transferred indices:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/recovery?v'
```

- information about shards in a cluster:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cat/shards?v'
```

- detailed information about the cluster:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_cluster/stats?human&pretty'
```

- detailed information about the nodes:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_nodes/stats?human&pretty'
```


Elasticsearch Search API

- searching for documents by the string:\

```
*`# curl -XPOST -u login:password '127.0.0.1:9200/twitter*/tweet/_search?
→pretty=true' -d '{`*\
  *`"query": {`*\
    *`"bool" : {`*\
      *`"must" : {`*\
        *`"query_string" : {`*\
          *`"query" : "test"`*\
        *`}`*\
      *`}`*\
    *`}`*\
  *`}`*\
*`}'`*`
```

- searching for document by the string and filtering:\

```
*`# curl -XPOST -u login:password '127.0.0.1:9200/twitter*/tweet/_search?
→pretty=true' -d '{`*\
  *`"query": {`*\
    *`"bool" : {`*\
      *`"must" : {`*\
        *`"query_string" : {`*\
          *`"query" : "testuje"`*\
        *`}`*\
      *`},`*\
      *`"filter" : {`*\
        *`"term" : { "user" : "lab1" }`*\
      *`}`*\
    *`}`*\
  *`}`*\
*`}'`*`
```

- simple search in a specific field (in this case user) uri query:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter*/_search?
q=user:lab1&pretty=true'
```

- simple search in a specific field:\

```
*`# curl -XPOST -u login:password '127.0.0.1:9200/twitter*/_search?pretty=true'
↪-d '{`*\
    *`"query" : {`*\
        *`"term" : { "user" : "lab1" }`*\
    *`}``*\
*`}``*`
```

Mapping, Fielddata and Templates

Mapping is a collection of fields along with a specific data type Fielddata is the field in which the data is stored (requires a specific type - string, float) Template is a template based on which fielddata will be created in a given index.

- Information on all set mappings:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_mapping?pretty=true'
```

- Information about all mappings set in the index:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_mapping/*?pretty=true'
```

- Information about the type of a specific field:\

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter/_mapping/field/message*?pretty=true'
```

- Information on all set templates:\

```
# curl -XGET -u login:password '127.0.0.1:9200/_template/*?pretty=true'
```

- Create - Mapping / Fielddata - It creates index twitter-float and the tweet message field sets to float:

```
*`# curl -XPUT -u login:password '127.0.0.1:9200/twitter-float?pretty=true' -d
→ '{`*\
  *`"mappings": {`*\
    *`"tweet": {`*\
      *`"properties": {`*\
        *`"message": {`*\
          *`"type": "float"`*\
        *`}`*\
      *`}`*\
    *`}`*\
  *`}`*\
*`}'`*
```

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter-float/_mapping/field/message?pretty=true'
```

- Create Template:\

```
# curl -XPUT -u login:password '127.0.0.1:9200/_template/
template_1' -d '{"template" : "twitter4", "order" : 0, "settings" :
{"number_of_shards" : 2}}'

# curl -XPOST -u login:password '127.0.0.1:9200/twitter4/tweet?pretty=true'
-d '{\
```

```
*`"user" : "lab1",`*\
*`"post_date" : "2017-08-25T10:10:00",`*\
*`"message" : "test of ID generation"`*\
*`}``*\
```

```
# curl -XGET -u login:password '127.0.0.1:9200/twitter4/_settings?pretty=true'
```

– Create Template2 - Sets the mapping template for all new indexes specifying that the tweet data, in the field called message, should be of the “string” type:\

```
*`# curl -XPUT -u login:password '127.0.0.1:9200/_template/template_2' -d`{'*\
*`"template" : "*",`*\
*`"mappings": {`*\
*`"tweet": {`*\
*`"properties": {`*\
*`"message": {`*\
*`"type": "string"`*\
*`}``*\
*`}``*\
*`}``*\
*`}``*\
*`}``*\
*`}``*\
*`}``*\
```

- Delete Mapping - Deleting a specific index mapping (no possibility to delete - you need to index):\

```
# curl -XDELETE -u login:password '127.0.0.1:9200/twitter2'
```

- Delete Template:\

```
# curl -XDELETE -u login:password '127.0.0.1:9200/_template/template_1?
pretty=true'
```

The Energy Log Server use Logstash service to dynamically unify data from disparate sources and normalize the data into destination of your choice. A Logstash pipeline has two required elements, *input* and *output*, and one optional element *filter*. The input plugins consume data from a source, the filter plugins modify the data as you specify, and the output plugins write the data to a destination. The default location of the Logstash plugin files is: `/etc/logstash/conf.d/`. This location contains the following Energy Log Server default plugins:

- `01-input-beats.conf`
- `01-input-syslog.conf`
- `020-filter-beats-syslog.conf`
- `020-filter-network.conf`
- `099-filter-geoip.conf`
- `100-output-elasticsearch.conf`
- `naemon_beat.example`
- `perflogs.example`

Logstash - Input “beats”

This plugin wait for reciving data from remote beats services. It use tcp /5044 port for communicaton:

```
input {  
  beats {  
    port => 5044  
  }  
}
```

Logstash - Input “network”

This plugin read events over a TCP or UDP socket assigns the appropriate tags:

```
input {  
  tcp {  
    port => 5514  
    type => "network"  
  
    tags => [ "LAN", "TCP" ]  
  }  
  
  udp {  
    port => 5514  
    type => "network"  
  
    tags => [ "LAN", "UDP" ]  
  }  
}
```


Logstash - Filter “beats syslog”

This filter processing an event data with syslog type:

```
filter {

  if [type] == "syslog" {
    grok {
      match => {
        "message" => [
          # auth: ssh/sudo/su

          "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
          ↳{SYSLOGHOST:[system][auth][hostname]} sshd(?:\[%{POSINT:[system][auth][pid]}\])?: %
          ↳{DATA:[system][auth][ssh][event]} %{DATA:[system][auth][ssh][method]} for (invalid_
          ↳user )?%{DATA:[system][auth][user]} from %{IPORHOST:[system][auth][ssh][ip]} port %
          ↳{NUMBER:[system][auth][ssh][port]} ssh2(: %
          ↳{GREEDYDATA:[system][auth][ssh][signature]})?"",

          "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
          ↳{SYSLOGHOST:[system][auth][hostname]} sshd(?:\[%{POSINT:[system][auth][pid]}\])?: %
          ↳{DATA:[system][auth][ssh][event]} user %{DATA:[system][auth][user]} from %
          ↳{IPORHOST:[system][auth][ssh][ip]}",

          "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
          ↳{SYSLOGHOST:[system][auth][hostname]} sshd(?:\[%{POSINT:[system][auth][pid]}\])?:
          ↳Did not receive identification string from %{IPORHOST:[system][auth][ssh][dropped_
          ↳ip]}",

          "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
          ↳{SYSLOGHOST:[system][auth][hostname]} sudo(?:\[%{POSINT:[system][auth][pid]}\])?:
          ↳\s*%{DATA:[system][auth][user]} : ( %{DATA:[system][auth][sudo][error]} ;)? TTY=%
          ↳{DATA:[system][auth][sudo][tty]} ; PWD=%{DATA:[system][auth][sudo][pwd]} ; USER=%
          ↳{DATA:[system][auth][sudo][user]} ; COMMAND=%
          ↳{GREEDYDATA:[system][auth][sudo][command]}",
```

(continues on next page)

(continued from previous page)

```

                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} %{DATA:[system][auth][program]}(?:\[%
→{POSINT:[system][auth][pid]}\})?: %{GREEDYMULTILINE:[system][auth][message]}",

                                # add/remove user or group
                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} groupadd(?:\[%{POSINT:[system][auth][pid]}\})?
→: new group: name=%{DATA:system.auth.groupadd.name}, GID=%{NUMBER:system.auth.
→groupadd.gid}",

                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} userdel(?:\[%{POSINT:[system][auth][pid]}\})?
→: removed group '%{DATA:[system][auth][groupdel][name]}' owned by '%
→{DATA:[system][auth][group][owner]}'",

                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} useradd(?:\[%{POSINT:[system][auth][pid]}\})?
→: new user: name=%{DATA:[system][auth][user][add][name]}, UID=%
→{NUMBER:[system][auth][user][add][uid]}, GID=%
→{NUMBER:[system][auth][user][add][gid]}, home=%
→{DATA:[system][auth][user][add][home]}, shell=%
→{DATA:[system][auth][user][add][shell]}$",

                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} userdel(?:\[%{POSINT:[system][auth][pid]}\})?
→: delete user '%{WORD:[system][auth][user][del][name]}'$",

                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{SYSLOGHOST:[system][auth][hostname]} usermod(?:\[%{POSINT:[system][auth][pid]}\})?
→: add '%{WORD:[system][auth][user][name]}' to group '%
→{WORD:[system][auth][user][memberof]}'",

                                # yum install/erase/update package
                                "%{SYSLOGTIMESTAMP:[system][auth][timestamp]} %"
→{DATA:[system][package][action]}: %{NOTSPACE:[system][package][name]}"
                                ]
                                }

                                pattern_definitions => {
                                    "GREEDYMULTILINE"=> "(.|\n)*"
                                }
                                }

                                date {
                                    match => [ "[system][auth][timestamp]",
                                        "MMM d HH:mm:ss",
                                        "MMM dd HH:mm:ss"
                                    ]
                                    target => "[system][auth][timestamp]"
                                }

                                mutate {
                                    convert => { "[system][auth][pid]" => "integer" }
                                    convert => { "[system][auth][groupadd][gid]" => "integer" }
                                    convert => { "[system][auth][user][add][uid]" => "integer" }
                                    convert => { "[system][auth][user][add][gid]" => "integer" }
                                }

```

(continues on next page)

(continued from previous page)

```
}  
}  
}
```


Logstash Filter “network”

This filter processing an event data with network type:

```
filter {
  if [type] == "network" {
    grok {
      named_captures_only => true
      match => {
        "message" => [

          # Cisco Firewall
          "%{SYSLOG5424PRI}%{NUMBER:log_sequence#}: %{SPACE}%{IPORHOST:device_ip}: (?
→:.)? %{CISCOTIMESTAMP:log_data} CET: %{CISCO_REASON:facility}-%{INT:severity_level}-
→%{CISCO_REASON:facility_mnemonic}: %{SPACE}%{GREEDYDATA:event_message}",

          # Cisco Routers
          "%{SYSLOG5424PRI}%{NUMBER:log_sequence#}: %{SPACE}%{IPORHOST:device_ip}: (?
→:.)? %{CISCOTIMESTAMP:log_data} CET: %{CISCO_REASON:facility}-%{INT:severity_level}-
→%{CISCO_REASON:facility_mnemonic}: %{SPACE}%{GREEDYDATA:event_message}",

          # Cisco Switches
          "%{SYSLOG5424PRI}%{NUMBER:log_sequence#}: %{SPACE}%{IPORHOST:device_ip}: (?
→:.)? %{CISCOTIMESTAMP:log_data} CET: %{CISCO_REASON:facility}-%{INT:severity_level}-
→%{CISCO_REASON:facility_mnemonic}: %{SPACE}%{GREEDYDATA:event_message}",
          "%{SYSLOG5424PRI}%{NUMBER:log_sequence#}: %{SPACE} (? : . ) ? %
→{CISCOTIMESTAMP:log_data} CET: %{CISCO_REASON:facility}-%{INT:severity_level}-%
→{CISCO_REASON:facility_mnemonic}: %{SPACE}%{GREEDYDATA:event_message}",

          # urzadzenia HP switches
          "%{SYSLOG5424PRI}%{SPACE}%{CISCOTIMESTAMP:log_data} %{IPORHOST:device_ip}
→%{CISCO_REASON:facility}: %{SPACE}%{GREEDYDATA:event_message}"
        ]
      }
    }
  }
}
```

(continues on next page)

(continued from previous page)

```

syslog_pri { }

if [severity_level] {

    translate {
        dictionary_path => "/etc/logstash/dictionaries/cisco_syslog_severity.yml"
        field => "severity_level"
        destination => "severity_level_descr"
    }

}

if [facility] {

    translate {
        dictionary_path => "/etc/logstash/dictionaries/cisco_syslog_facility.yml"
        field => "facility"
        destination => "facility_full_descr"
    }

}

# okreslamy locality dla adresow IP na listach ACL
if [event_message] =~ /\(d+\.d+\.d+\.d+\)/ {
    grok {
        match => {
            "event_message" => [
                "list %{NOTSPACE:[acl][name]} %{WORD:[acl][action]} %{WORD:[acl][proto]} %
↪{IP:[src][ip]}.*%{IP:[dst][ip]}",
                "list %{NOTSPACE:[acl][name]} %{WORD:[acl][action]} %{IP:[src][ip]}",
                "^list %{NOTSPACE:[acl][name]} %{WORD:[acl][action]} %{WORD:[acl][proto]}
↪%{IP:[src][ip]}.*%{IP:[dst][ip]}"
            ]
        }
    }
}

if [src][ip] {

    cidr {
        address => [ "%{[src][ip]}" ]
        network => [ "0.0.0.0/32", "10.0.0.0/8", "172.16.0.0/12", "192.168.0.0/16",
↪ "fc00::/7", "127.0.0.0/8", "::1/128", "169.254.0.0/16", "fe80::/10", "224.0.0.0/4",
↪ "ff00::/8", "255.255.255.255/32" ]
        add_field => { "[src][locality]" => "private" }
    }

    if ![src][locality] {
        mutate {
            add_field => { "[src][locality]" => "public" }
        }
    }
}

if [dst][ip] {

```

(continues on next page)

(continued from previous page)

```
cidr {
  address => [ "%{[dst][ip]}" ]
  network => [ "0.0.0.0/32", "10.0.0.0/8", "172.16.0.0/12", "192.168.0.0/16",
↪ "fc00::/7", "127.0.0.0/8", "::1/128",
    "169.254.0.0/16", "fe80::/10", "224.0.0.0/4", "ff00::/8", "255.255.255.
↪ 255/32" ]
  add_field => { "[dst][locality]" => "private" }
}

if ![dst][locality] {
  mutate {
    add_field => { "[dst][locality]" => "public" }
  }
}

}

# date format
date {
  match => [ "log_data",
    "MMM dd HH:mm:ss",
    "MMM dd HH:mm:ss",
    "MMM dd HH:mm:ss.SSS",
    "MMM dd HH:mm:ss.SSS",
    "ISO8601"
  ]
  target => "log_data"
}

}
}
```

Logstash - Filter “geoip”

This filter processing an events data with IP address and check localization:

```
filter {
  if [src][locality] == "public" {

    geoip {
      source => "[src][ip]"
      target => "[src][geoip]"
      database => "/etc/logstash/geoipdb/GeoLite2-City.mmdb"
      fields => [ "city_name", "country_name", "continent_code", "country_code2"
↪, "location" ]
      remove_field => [ "[src][geoip][ip]" ]
    }

    geoip {
      source => "[src][ip]"
      target => "[src][geoip]"
      database => "/etc/logstash/geoipdb/GeoLite2-ASN.mmdb"
      remove_field => [ "[src][geoip][ip]" ]
    }
  }

  if [dst][locality] == "public" {

    geoip {
      source => "[dst][ip]"
      target => "[dst][geoip]"
      database => "/etc/logstash/geoipdb/GeoLite2-City.mmdb"
      fields => [ "city_name", "country_name", "continent_code", "country_code2"
↪, "location" ]
      remove_field => [ "[dst][geoip][ip]" ]
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
    geoip {
      source => "[dst][ip]"
      target => "[dst][geoip]"
      database => "/etc/logstash/geoipdb/GeoLite2-ASN.mmdb"
      remove_field => [ "[dst][geoip][ip]" ]
    }
  }
}
```

Logstash - Output to Elasticsearch

This output plugin sends all data to the local Elasticsearch instance and create indexes:

```
output {  
  elasticsearch {  
    hosts => [ "127.0.0.1:9200" ]  
    #http_compression => true  
    #sniffing => true  
  
    index => "%{type}-%{+YYYY.MM.dd}"  
  
    user => "logserver"  
    password => "logserver"  
  }  
}
```

Logstash plugging for “naemon beat”

This Logstash plugin has example of complete configuration for integration with naemon application:

```
input {
  beats {
    port => FILEBEAT_PORT
    type => "naemon"
  }
}

filter {
  if [type] == "naemon" {
    grok {
      patterns_dir => [ "/etc/logstash/patterns" ]
      match => { "message" => "%{NAEMONLOGLINE}" }
      remove_field => [ "message" ]
    }
    date {
      match => [ "naemon_epoch", "UNIX" ]
      target => "@timestamp"
      remove_field => [ "naemon_epoch" ]
    }
  }
}

output {
  # Single index
  # if [type] == "naemon" {
  #   elasticsearch {
  #     hosts => ["ELASTICSEARCH_HOST:ES_PORT"]
  #     index => "naemon-%{+YYYY.MM.dd}"
  #   }
  # }

  # Separate indexes
```

(continues on next page)

(continued from previous page)

```
if [type] == "naemon" {
  if "_grokparsefailure" in [tags] {
    elasticsearch {
      hosts => ["ELASTICSEARCH_HOST:ES_PORT"]
      index => "naemongrokfailure"
    }
  }
  else {
    elasticsearch {
      hosts => ["ELASTICSEARCH_HOST:ES_PORT"]
      index => "naemon-%{+YYYY.MM.dd}"
    }
  }
}
```

Logstash plugging for “perflog”

This Logstash plugin has example of complete configuration for integration with perflog:

```
input {
  tcp {
    port => 6868
    host => "0.0.0.0"
    type => "perflogs"
  }
}

filter {
  if [type] == "perflogs" {
    grok {
      break_on_match => "true"
      match => {
        "message" => [
          "DATATYPE::%{WORD:datatype}\tTIMET::%{NUMBER:timestamp}\tHOSTNAME::%
→{DATA:hostname}\tSERVICEDESC::%{DATA:servicedescription}\tSERVICEPERFDATA::%
→{DATA:performance}\tSERVICECHECKCOMMAND::.*?HOSTSTATE::%{WORD:hoststate}
→\tHOSTSTATETYPE::.*?SERVICESTATE::%{WORD:servicestate}\tSERVICESTATETYPE::%
→{WORD:servicestatetype}",
          "DATATYPE::%{WORD:datatype}\tTIMET::%{NUMBER:timestamp}\tHOSTNAME::%
→{DATA:hostname}\tHOSTPERFDATA::%{DATA:performance}\tHOSTCHECKCOMMAND::.*?HOSTSTATE::
→%{WORD:hoststate}\tHOSTSTATETYPE::%{WORD:hoststatetype}"
        ]
      }
      remove_field => [ "message" ]
    }
    kv {
      source => "performance"
      field_split => "\t"
      remove_char_key => "\.\'"
      trim_key => " "
      target => "perf_data"
    }
  }
}
```

(continues on next page)

(continued from previous page)

```
    remove_field => [ "performance" ]
    allow_duplicate_values => "false"
    transform_key => "lowercase"
  }
  date {
    match => [ "timestamp", "UNIX" ]
    target => "@timestamp"
    remove_field => [ "timestamp" ]
  }
}

output {
  if [type] == "perflogs" {
    elasticsearch {
      hosts => ["127.0.0.1:9200"]
      index => "perflogs-%{+YYYY.MM.dd}"
    }
  }
}
```


CHAPTER 94

CHANGELOG

95.1 Added

- Export Dashboard to PDF

95.2 Changed

- bugfix dashboard generation using phantomjs

96.1 Added

- Report generation time added to the PDF

96.2 Changed

- bugfix export Dashboard relative paths
- bugfix bigger dashboards
- bugfix relative path bug fix
- bugfix Export to CSV

97.1 Added

- Kibana Reporting plugin

97.2 Changed

- bugfix for report generation date on PDF
- bugfix elasticsearch auth jar
- bugfix reports validation

98.1 Added

- Password Change function
- Deleting old CSV/PDF function

99.1 Added

- CSV Reports moved to reports tab

99.2 Changed

- bugfix ES Auth Plugin Jar
- bugfix Kibana
- bugfix for scheduler index

CHAPTER 100

2.1.29

100.1 Added

- Basic Auth function

CHAPTER 101

2.1.30

101.1 Changed

- bugfix From and To date on CSV CSV Task Export

102.1 Changed

- bugfix auth-local-cookie.js change url to css file to relative
- bugfix fter submit command on 'export dashboard' page, Dashboard field shines in red
- bugfix server should not allow GET (even if user manually change request) server side verification
- bugfix In Export Task Managment tab use Search query "*" i Time Criteria Field Name as "@timestamp" by default

103.1 Added

- Introduction of CSRF Tokens for each action. The CSRF Token can be generated once at login and appended to all GET / POST actions.

104.1 Added

- Changing your password require you to enter your current password.
- Password change form require you to enter a new password twice.
- Cookies new flags: expires, secure, domain.

104.2 Changed

- bugfix Password Change function

105.1 Added

- Auto create user elastalert by kibana with random password (like logserver user)
- Change Admin tab to Config
- Deleting a user or role should require confirmation that the user is sure about this action.

105.2 Changed

- bugfix: Marvel plugin problem with auth plugin enabled.
- bugfix: Public directory <https://127.0.0.0.1:5601/bundles/>

106.1 Added

- In audit.log should not appear _bulk (default off)
- In audit.log should not appear password update/create request.
- In audit.log query output on/off (default on)
- Userlist and Rolelist alphabetically
- Add username in Password change page info about role on change page

CHAPTER 107

2.1.37

107.1 Added

- Non logged user which clicks on the link to the dashboard is redirected to the login page, and after successfully logging user is redirected to Discovery section not to the dashboard from the link

CHAPTER 108

2.2.1

108.1 Added

- Introduction to Licensing module

CHAPTER 109

2.2.2

109.1 Added

- Licensing module

CHAPTER 110

2.2.3

110.1 Added

- Code Documentation

110.2 Changed

- bugfix: in User/Role Update

CHAPTER 111

2.3.0

111.1 Added

- Object Management Elastafilter (kibana)
- Active Directory AD Integration (elasticsearch-auth)

111.2 Changed

- bugfix: in Alerts module (kibana)

112.1 Added

- AD config moved to elasticsearch.yml (elasticsearch-auth)

112.2 Changed

- bugfix: Dashboards elastafilter (kibana)

113.1 Added

- In Config-> settings: Miliseconds to seconds change
- Ldap/AD auth as optional

113.2 Changed

- bugfix Userlist problem no roles selected after Update action
- bugfix elasticsearch logging

114.1 Added

- Timepicker on task export changed from 2015-05-19 to 2018-01-01*
- Object permission selector set/unset all
- Java jdk1.8.0_161 added to git
- Phantomjs added as kibana dependency
- Marvel added to git

114.2 Changed

- bugfix: If admin create role with /* pattern causes elasticsearch and kibana (no connection to elasticsearch) crash in loop. . . .
- bugfix: Default config settings with elastafilter problem with timeout 0/token expire etc
- bugfix: In Alert -> Flatline bad example
- bugfix: In Alert -> Alert writeback index based on the configuration file.

115.1 Added

- LDAP/AD SSL Support
- Title and Logos change to Logserver

115.2 Changed

- bugfix: no Rules directory for alert
- bugfix: Example LDAP settings change

116.1 Added

- Logserver/Energy tab changed to Search tab in Kibana
- Show license details in Kibana
- Added default system settings for services [sysctl,limits.d]
- Java dir changed => jdk-1.8.0_161.tar.bz2
- New default users: logserver, alert, scheduler, intelligence

116.2 Changed

- bugfix: Exception in thread “Timer-1”
- bugfix: Default elastalert password changed

117.1 Added

- Intelligence plugin
- Schedule plugin
- Added kibana/elasticsearch PayLoad settings to config files.

117.2 Changed

- bugfix: All shards failed in Discovery Tab => (elastafilter limits)
- bugfix: LDAP settings format

118.1 Added

- Intelligence plugin upgrade
- Schedule plugin upgrade
- Alerts upgrade
- Java upgrade to 1.8_171
- Added index-rotation.sh script for audit & alert index

118.2 Changed

- bugfix: All shards failed in Discovery Tab => (elastafilter limits)
- bugfix: Max clause limit => 10240

119.1 Added

- Preview & Status in Intelligence section
- Role-Mapping for AD
- Default Users Role
- Added scripts\ldap-alias-create.sh for AD User Alias
- Example Alert Dashboard

119.2 Changed

- bugfix: Intelligence Trend

120.1 Changed

- bugfixes: Intelligence Scheduler Login Filter
- bugfixes: elasticsearch-auth plugin

121.1 Added

- Timelion
- More detailed logs Kibana
- More detailed logs Intelligence
- PDF/CSV reports for AD users
- Alerts preview for Users
- New Alert SAVE button

121.2 Changed

- bugfix: audit index: @timestamp
- bugfix: Intelligence Trend Threshold fix
- bugfix: Intelligence (SPARK_LOCAL_IP)
- bugfix: Mangement typo in Reports
- bugfix: Objects import
- bugfix: Recovery Alert Dashboard fix
- bugfix: Password update/change disabled for AD users

122.1 Added

- Kibana role for GUI users
- Alert role for Alert Module users
- Intelligence role for A/I Module users
- Intelligence temporary directory
- Intelligence Preview update

123.1 Added

- SSO single sign on implementation
- Improved installation process
- OpenJDK Support
- Removed all unnecessary dependencies
- Removed python libraries (pip)
- Logstash configuration examples
- OVF changed to OVA

123.2 Changed

- bugfix: Conflict during system upgrade (yum update)
- bugfix: journald.conf conflict
- bugfix: Defaultrole AD users
- bugfix: start_date Intelligence module

CHAPTER 124

2.4.5

124.1 Added

- New Login Page with SSO support
- Alert module update

124.2 Changed

- bugfix: metric_aggregation, new term, percentage_match