

Proxedo API Security in Kubernetes: Administration Guide

Copyright © 2019 Balasys IT Ltd.. All rights reserved. This document is protected by copyright and is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of Balasys.

This documentation and the product it describes are considered protected by copyright according to the applicable laws.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). This product includes cryptographic software written by Eric Young ([eay@cryptsoft.com](mailto:eyay@cryptsoft.com))

Linux™ is a registered trademark of Linus Torvalds.

Windows™ 10 is registered trademarks of Microsoft Corporation.

The Balasys™ name and the Balasys™ logo are registered trademarks of Balasys IT Ltd.

The Proxedo™ name and the Proxedo™ logo are registered trademarks of Balasys IT Ltd.

AMD Ryzen™ and AMD EPYC™ are registered trademarks of Advanced Micro Devices, Inc.

Intel® Core™ and Intel® Xeon™ are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

All other product names mentioned herein are the trademarks of their respective owners.

DISCLAIMER

Balasys is not responsible for any third-party websites mentioned in this document. Balasys does not endorse and is not responsible or liable for any content, advertising, products, or other material on or available from such sites or resources. Balasys will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through any such sites or resources.

2026-07-16

Preface

Typographical conventions

Before you start using this guide, it is important to understand the terms and typographical conventions used in the documentation. For more information on specialized terms and abbreviations used in the documentation, see the [Glossary](#) at the end of this document.

The following text formatting principles and icons identify special information in the document.



Tips provide best practices and recommendations.



Notes provide additional information on a topic, and emphasize important facts and considerations.



Warnings mark situations where loss of data or misconfiguration of the device is possible if the instructions are not obeyed.

Command

Commands you have to execute.

Emphasis

Reference items, additional readings.

`/path/to/file`

File names.

Parameters

Parameter and attribute names.

Additional marks used specifically in the Web User Interface (UI):

Key	Description
*	The elements marked with * in the configuration reference tables are mandatory to be configured.
(Default)	For some of the configuration elements there are recommended default values, marked as (Default). In case the value is not defined during the configuration, the default value will be considered for the actual element.
+	By clicking this sign you can add the actual element to the list of configuration elements.

Contact and support information

This product is developed and maintained by Balasys IT Ltd..

Contact:

Balasys IT Ltd.
4 Alíz Street
H-1117 Budapest, Hungary
Tel: +36 1 646 4740
E-mail: <info@balasys.hu>
Web: <http://balasys.hu/>

Sales contact

You can directly contact us with sales-related topics at the e-mail address <sales@balasys.hu>, or leave us your contact information and we call you back.

Support contact

To access the Balasys Support System, sign up for an account at the Balasys Support System page. Online support is available 24 hours a day.

Balasys Support System is available only for registered users with a valid support package.

Support e-mail address: <support@balasys.hu>.

Training

Balasys IT Ltd. holds courses on using its products for new and experienced users. For dates, details, and application forms, visit the <https://www.balasys.hu/en/services#training> webpage.

1. Scope of this document

This document describes installation, configuration and operation of Proxedo API Security in Kubernetes. The purpose of this document is to present the designed approach for different installation scenarios, base system configuration, and the usage of the Web User Interface (UI). It also documents common use cases for operation and troubleshooting. The primary intended audience of this document are system engineers and system designers for configuring Proxedo API Security systems.

2. Introduction to Proxedo API Security

2.1. What is Proxedo API Security

The Proxedo API Security (PAS) is a security solution that protects [API](#) serving endpoints. It is positioned in the network flow between consumers of the APIs (clients) and backend solutions serving the API (servers) as a transparent [HTTP](#) proxy.

Proxedo API Security can:

- handle incoming Transport Layer Security v1 ([TLS](#)) connections from clients & outgoing TLS connections to servers separately and selectively
- verify that the communication conforms to HTTP specifications
- verify that the content of the messages conform to their specified content type
- verify that the content of messages conform to API specification(s) as described in schemas
- evaluate the level of risk with regards to the API call using the data collected from call data

- provide rule-based protection against a variety of web-based application layer attacks
- extract parts of the content of the messages and relay them to external data stores such as log servers, [SIEM](#) systems or other data warehouses

2.2. Where to start

Depending on what you need to do the following starting points are suggested:

- To understand what the product does and how, see [Overview of Proxedo API Security](#).
 - If you are familiar with API terminology jump right to [Architecture for Proxedo API Security](#).
- See [Installation of Proxedo API Security in Kubernetes environment](#) if you need to set up a new PAS.
- The [Operation of Proxedo API Security in Kubernetes environment](#) chapter is about how to manage a working system on the level of the operating system.
- [Configuration of Proxedo API Security on the Web User Interface](#) contains in-depth information about everything that can be configured with the help of the Web User Interface.
- If you are already familiar with the system and need to find a component that suits your needs consult the [Matcher types](#), [Comparators](#), [Extractor types](#) or [Insight Target](#).

3. Overview of Proxedo API Security

3.1. Main features

3.1.1. TLS

Transport Layer Security v1 (TLS) (successor of the now obsoleted Secure Socket Layer v3 (SSL)) is a widely used crypto protocol, guaranteeing data integrity and confidentiality in many PKI and e-commerce systems.

The TLS framework inspects TLS connections, and also any other connections embedded into the encrypted TLS channel. TLS connections initiated from the client are terminated on the Proxedo API Security, and two separate TLS connections are built: one between the client and the firewall, and one between the firewall and the server. If both connections match the configuration settings of PAS (for example, the certificates are valid, and only the allowed encryption algorithms are used), PAS inspects the protocol embedded into the secure channel as well. Note that the configuration settings can be different for the two connections, for example, it is possible to permit different protocol versions and encryption settings.

3.1.2. Enforcement

Proxedo API Security acts as an HTTP proxy and verifies that the traffic passing through conforms to HTTP's specifications. By using OpenAPI schemas, as defined in OpenAPI specifications (also known as Swagger), it also verifies that the traffic passing through conforms to the API endpoint's specification and can log or deny non-conforming traffic.

PAS also provides its own versatile filtering system to control passing traffic.

3.1.3. Fraud Detection

The Fraud Detection module of Proxedo API Security reduces the number of fraudulent transactions by harnessing device fingerprinting and enriching incoming data with alternate sources to provide the best accuracy and details about transactions.

3.1.4. Rule-based Enforcement

Besides its positive security model approach, Proxedo API Security also has a web application firewall module.

The *WAF Enforcer* protects against a variety of application layer attacks including credential theft, code injection, cross-site scripting (XSS), cookie poisoning, CSRF, SQL injection, DoS, ransomware, and more.

3.1.5. Insights

With Proxedo API Security it is possible to extract business-relevant information with extremely high resolution from the traffic and relay it to external data stores where further analysis can be implemented.

Thus, it is possible to feed Log Management solutions, Monitoring and SIEM systems, Data visualization tools with data extracted from the traffic, even to the level of specific fields deep inside API calls or URI parameters.

3.1.6. Security flow

The security flow binds most of PAS's features together. It allows flexible configuration for handling the traffic. *Multiple Enforcement*, *Filter* and *Insight plugins* can be mix-and-matched with control over error policies.

3.2. Main Concepts in Proxedo API Security

This chapter provides an overview of the Proxedo API Security solution, introduces its main concepts, and explains the relationship of the various components.

API Endpoint

Proxedo API Security protects API endpoints. An API endpoint is the serving part of the communication channel and is the collection of all functions of a service. It resides at a list of well-known top URIs under which all the functions are accessible. APIs have well-defined HTTP Endpoints for all exposed calls, resources etc., usually through providing a schema that describes all parameters of these URI paths, including possible HTTP response codes, the format and fields of the data structure in the request's and response's body.

Client

It is a consumer of API endpoints. It is the source of the requests.

Backend

The backend constitutes of one or more servers that serve the API endpoint. It receives the requests of the client and sends the responses.

HTTP message

It can be an HTTP request coming from the client or an HTTP response coming from the backend.

Call

An HTTP conversation constitutes of a request — response interchange of HTTP messages between the client and the backend. Whenever the direction is irrelevant in the context — it applies to both requests and responses — the message is named Call.

Listener

It is the part of PAS that listens to incoming traffic for given API Endpoints. It is bound to a network port. Clients address this port when accessing API Endpoints through the gateway.

TLS

Transport Layer Security is the cryptographic protocol that secures HTTPS communications. PAS can apply TLS encryption both when communicating with Clients and Backends. TLS encryption can also be used with *Syslog Insight Target* and *Elastic Insight Target*.

Security flow

It provides a collection of security rules that PAS applies to a Call. It is two series of *Plugins*: one for requests and one for responses.

Plugin

It is an element of the security flow that applies a specific security function. It has different types based on the role they do.

Deserializer

A *Plugin* responsible for decompressing and parsing the HTTP message's body to structured data. This ensures that a message is well-formed. The structured data will also be consumed by other *Plugins* that operate on the body of the message.

Serializer

A *Plugin* responsible for serializing and compressing the structured data to the format of the HTTP message's body.

Filter

A *Plugin* that rejects calls when they match defined rules.

Enforcer

A *Plugin* that validates calls against externally defined schemas.

Insight

A *Plugin* that extracts various data from the call and sends it to external systems (log servers, SIEMs, and other data analysis tools).

Brick

They are reusable components. They can be defined on their own and then shared by multiple other components.

Error policy

It is a brick that defines what happens if the *Plugin* has found an error. It decides if calls are rejected or merely logged, and defines the details of the HTTP error response sent to the client if a call is rejected.

Matcher

It is a brick that decides if the *Plugin* should be executed for a given call by checking various data in the HTTP message.

Selector

Selector is a brick that can extract a piece of information from a call. It is used by *Insight plugins*.

Insight Target

It is a brick that defines an external system to send extracted data to. It is used by *Insight plugins*.

3.3. Architecture for Proxedo API Security

Proxedo API Security is based on a micro-services architecture separated into three deployment units: *Management*, *Storage*, and *Core*. These deployment units (or infrastructure components) can be scaled or moved between hosts to accommodate different throughput and reliability requirements.

3.3.1. Management component

Responsible for handling the security component configuration of the *Core* component, while the data itself resides in the *Storage* component. Contains the following services:

Config API

Exposes a configuration API that can be used to manage the product:

- Editing the security component configuration

- Applying the security component configuration

Config WebUI

Provides a browser-based user interface to the configuration API.

3.3.2. Storage component

Stores and distributes different versions of the security component configuration to the *Core* component. Contains the following services:

Consul

Stores the different versions of the security component configuration, and monitors the status of PAS services.

Blob Store

Stores file resources that are part of the security component configuration.

Log Manager

Collects operational logs from other services.

3.3.3. Core component

The *Core* services are each responsible for a well-defined subset of handling traffic between the client and the backend. Contains the following services:

Transport Director

Manages the transport layer of API connections:

- Handles network connections from the client
- Handles network connections towards the backends
- Handles TLS on these connections
- Load balances between multiple backend servers
- Load balances between multiple *Flow Directors*
- Enforces HTTP protocol validity in calls

Flow Director

Responsible for the execution of the *Plugins* in the *Endpoints'* flow and for applying *Error Policies* as necessary.

Insight Director

Manages the connections to *Insight Targets*. Responsible for sending the data collected by *Insight plugins* to *Insight Target* systems.

Content Filtering Director

Provides content filtering capabilities for the *WAF Enforcer plugin*.

3.3.4. The configuration process



While the configuration most commonly takes place on the Web UI, the process works the same way through the configuration API.

1. When a user logs in to the *Web UI*, the currently running configuration is visible.
 - When logging in to the *Web UI* for the first time after a fresh install, the current configuration is empty. Only a few mandatory and default components are added, and some mandatory components must be added to the configuration for the first configuration to become valid.

- The running configuration is always stored in the *Storage* component.
- 2. The user can edit the configuration: add new components, delete existing components, and change fields on existing components.
 - The changes the user makes are only visible to the user, other users can only see the running configuration and their own changes.
 - The user's changes are always stored in the *Storage* component.
- 3. Individual components and the configuration as a whole are validated.
 - Partially configured components can be saved with missing fields, but they won't become valid until all mandatory fields are properly filled.
 - An invalid configuration is still saved, and can be fixed at a later time. Every user has their own set of changes.
- 4. When the configuration is valid, it can be applied to the running system.
 - When a user's configuration is applied, the changes are merged with the running configuration.
 - Applying the changes means reloading the *Core* services with the new configuration.
 - The new running configuration becomes visible to every user.

3.3.5. Connection handling example




Figure 1. PAS Architecture

1. Incoming connections are accepted by the *Transport Director*.
 - It handles TLS with the client if necessary.
 - Chooses the *Endpoint* based on the URL.
2. It hands over the connection to the *Flow Director*.
 - The *Flow Director* applies the *Endpoint* specific *Request Security Flow*.
3. If an *Insight plugin* needs to send data to an external *Insight Target* it sends the collected data to the *Insight Director*.

4. The *Insight Director* sends the data further to the *Insight Target* with the appropriate protocol.
5. If a *WAF Enforcer plugin* is present in the *Request Security Flow* it sends data to the *Content Filtering Director* and receives a verdict.
6. If a *Fraud Detector plugin* is present in the *Request Security Flow* it sends data to the external *Fraud API* and receives a score.
7. The *Flow Director* hands the connection back to the *Transport Director*.
8. The *Transport Director* then sends the data to the *Backend*.
 - It handles TLS with the backends if necessary.
 - It performs load balancing among *Backend* servers if necessary.

The same procedure is executed with the response coming from the *Backend*.

3.3.6. Understanding processing flow

The figure on Proxedo API Security architecture and the steps following that describe how client connection is handled. The following figure explains how calls are processed in more details:




Figure 2. PAS processing flow

1. As shown in the figure above, the incoming connection from the client is handled by the *Transport Director*, applying TLS if needed.
2. The *Transport Director* then chooses the *Endpoint* based on the URL in the request. First endpoint that has a matching URL is chosen.
3. The *Transport Director* hands over the connection to the *Flow Director*, indicating which *Endpoint* the connection belongs to.
4. The *Flow Director* then starts applying the request part of the *Security Flow* definition.
5. For each *Plugin* the *Flow Director*:
 - Checks if the *Plugin*'s matcher matches the request.
 - If so, it executes the *Plugin*, if not, it executes the next *Plugin*.
 - If the *Plugin* indicates success it executes the next *Plugin*.
 - If the *Plugin* indicates an error it applies the *Plugin*'s error policy. If the policy dictates to abort the connection:

- It fills error details and hands back the connection to the *Transport Director*, aborting the execution of the flow.
 - The *Transport Director* closes the connection, sending error details to the client if allowed by the policy.
6. Once the last *Plugin* has been executed the connection is handed back to the *Transport Director*.
 7. The *Transport Director* initiates the connection towards the *Backend*:
 - It handles load balancing if necessary.
 - It handles TLS if necessary.
 - It sends the request itself to the *Backend* server.
 8. The *Backend* server sends its response to the *Transport Director*.
 9. Once, the response has been received the *Transport Director* again hands over the connection to the *Flow Director*.
 10. The *Flow Director* then starts applying the response part of the *Security Flow* definition, executing the *Plugins* as above.
 11. Once, the last *Plugin* has been executed the connection is handed back to the *Transport Director*.
 12. Finally, the *Transport Director* sends the response to the client.

Usually, *Plugins* are organized in the following manner:

- A *Deserializer Plugin* extracts and processes the request body to understand the details.
- Filters are applied to filter unnecessary traffic.
- Enforcers are applied for detailed validation of calls.
- Insights are applied to collect data from the call.
- *Serializer Plugin* serializes and compresses the body.

Though the order of the plugins can be changed based on the needs, note the followings:

- When a *Plugin* needs access to the request body it requires *Deserialized* data. It is therefore strongly recommended that the first plugin is a *Decompressor* followed by a *Deserializer*.
- At the end of the flow it is strongly recommended to place a *Serializer* plugin followed by a *Compressor*.
- Generally *Insights* are applied after *Filters* and *Enforcers* so that they are not executed on possibly invalid calls.
- Anything that operates on the HTTP headers or the body of the message will be aware of the call direction: The same *Plugin* in the request and response flow will act on the request or response data.
- However, the *Flow Director* handles a request-response exchange together, so you can still use details from the request in *Plugins* of the response flow. The most notable example of this is using *URI* or *method* matchers in the response flow.
- *Plugins* in the request flow, however, cannot access details of the response flow (since they are not available yet).

It is also worth noting that *Insight Plugins* instantly hand over data to the *Insight Director*, and let the execution continue.

4. Installation of Proxedo API Security in Kubernetes environment

The forthcoming sections describe the installation of PAS in Kubernetes.



To manage Kubernetes (K8s) applications, [Helm](#), the package manager for Kubernetes is used. Packages are called *charts* in the *Helm* context.

4.1. Prerequisites for installing PAS

The followings are needed prior to the installation of PAS:

- a technical user for accessing Balasys' download site
- the Helm chart



Prior to the installation of the *Helm chart*, the *Helm chart* itself must be configured. For minimum configuration of the *Helm chart* see section [Minimum configuration settings for the Helm chart](#).

4.1.1. Cluster components necessary for PAS

To make use of some of the features, PAS shall be deployed in a cluster, with the following components installed:

- metrics server for auto-scaling
- *Persistent volume* for storing configuration in the management component



Persistent Volume Claim parameters can be set up to match a manually managed *Persistent volume*, so is *Storage Class* name.

- access for the target namespace to deploy PAS in

4.1.2. Tools necessary for the installation

To create the basic configuration for the installation, the following tools are necessary:

- openssl for storage certificate generation
- the htpasswd tool, which is part of the apache2-utils package on debian distributions, the httpd-tools package on Red Hat based distributions
- the helm command line tool to manage the package installation
- the kubectl command line tool to communicate with the Kubernetes cluster

4.1.3. Minimum configuration settings for the Helm chart

The *Helm* chart contains the following:

- configuration parameters to bootstrap PAS in K8s
- definitions of
 - pods
 - services
 - autoscaling configuration for the core component
 - a *Persistent Volume Claim* for the management



Ingress configuration for any component is not included.



HTTP and HTTPS management access is recommended to be configured using an Ingress (kubernetes object).



In order to be able to install the *Helm chart* the minimum configuration settings have to be completed. The following sections contain the details only for the necessary minimum configuration, however for checking further possible configuration options, see section [Base system configuration for PAS in Kubernetes](#).

The files detailed in the next sections need to be created and filled in prior to PAS installation.

4.1.3.1. Using values.yml file

1. Use the values.yml (values file) with the default and necessary values. Run the following command to output the configuration options:

```
helm show values /path/to/chart/proxedo-api-security-4.15.0.tgz
```

2. Create a local values.yml file with the preferred values to overwrite the default values if required. The values file with minimum configuration is as follows (with example values):

```
config:
  storage:
    consul:
      gossip_encryption_key: MhstT80sqle63WC7kn0ak+c7GfK7k50Y2n/4Qk/fSXs=
    blob_store:
      access_key: GK3f53e895f3ade9c5763725dc
      secret_key: 3786755aee1989655f81c97e537d6d569dd9d321225522774bfc77ffb2d91f82
      rpc_secret: e36b62f4d0c63ff2c3d9d29cfcf4b2a1ced3d7f71c3b4250555287c57466712f
      admin_token: SS+Kp50sDg4zL69ZivriOnY0PDWpMYtpmN+BmsNmxeY=
```

3. Generate these necessary secrets with the help of the following command. The values above are examples, they shall not be copied directly.

```
# config.consul.gossip_encryption_key
$ openssl rand -base64 32
gI97yg2Zcq4XL20ne8NBwH2e0PbzkMxjqMFdp8jQZac=

# config.blob_store.access_key
$ echo GK$(openssl rand -hex 12)
GK3f53e895f3ade9c5763725dc

# config.blob_store.secret_key
$ openssl rand -hex 32
3786755aee1989655f81c97e537d6d569dd9d321225522774bfc77ffb2d91f82

# config.blob_store.rpc_secret
$ openssl rand -hex 32
e36b62f4d0c63ff2c3d9d29cfcf4b2a1ced3d7f71c3b4250555287c57466712f
```

```
# config.blob_store.admin_token
$ openssl rand -base64 32
SS+Kp50sDg4zL69ZivriOnY0PDWpMYtpmN+BmsNmxeY=
```

4.1.3.2. Creating certificates for storage

For technical reasons, a TLS certificate is necessary for configuration storage purposes. Create the internal CAs and signed certificates either with a preferred method, or else the necessary files can be created with the following example commands as well.

1. Generate a CA key pair.



The `-days` parameter in the example commands define the validity period of the generated certificates in days. Change it, if it is required.



The certificate files generated here and used with the *Helm chart* are sensitive pieces of information, therefore handle those with attention.

```
openssl req -nodes -new -x509 -days +3650 -keyout storage-ca-key.pem -out storage-ca.pem
-subj "/CN=PAS Storage CA"
```

2. Generate a private server key and a Certificate Signing Request (CSR).

```
openssl req -nodes -new -keyout consul-0-key.pem -out consul-0.csr -days +3650 -subj
"/CN=storage.pas"
```

3. Sign the CSR using the CA.

```
openssl x509 -req -days +3650 -in consul-0.csr -CA storage-ca.pem -CAkey storage-ca-
key.pem -CAcreateserial -out consul-0.pem
```

With the help of the above examples, further files need to be generated. These files will need to be provided for the *Helm chart*:

- consul-0.csr
- consul-0-key.pem
- consul-0.pem
- storage-ca-key.pem
- storage-ca.pem

4.1.3.3. Creating management users' file

For logging into the management component, the `users.htpass` file is required. Run the following command to generate one, and provide the password.

```
htpasswd -B -c users.htpass username
```

4.2. Installing PAS in Kubernetes

The following sections and the example commands use the `proxecto-api-security` kubernetes namespace as an example, but it can be replaced with any other namespace name.



It is recommended to install PAS in a namespace separate from the backend application(s).

To create a new namespace, run the following command:

```
kubectl create namespace proxecto-api-security
```

4.2.1. Setting up docker registry connection

1. Log in to the PAS docker registry to access the docker images of PAS.
2. Create the `proxecto-api-security-registry-credentials` secret using the following command to enable kubernetes to access the docker images:

```
kubectl create --namespace proxecto-api-security \  
  secret docker-registry proxecto-api-security-registry-credentials \  
  --docker-server=docker.balysys.hu \  
  --docker-username=<<your username>> \  
  --docker-password="$(read -sp "Docker registry password: " DOCKER_PASSWORD; echo  
  $DOCKER_PASSWORD)"
```

4.2.2. Providing the necessary files for *Helm* installation

Provide the created files for the *Helm* install command, an example of which can be seen below (substitute your values):

```
helm upgrade --install proxecto-api-security --namespace=proxecto-api-security \  
  --values /path/to/config/files/values.yml \  
  --set-file mgmt_users=/path/to/config/files/users.htpass \  
  --set-file storage_ca_key=/path/to/config/files/storage-ca-key.pem \  
  --set-file storage_ca_cert=/path/to/config/files/storage-ca.pem \  
  --set-file storage_server_key=/path/to/config/files/consul-0-key.pem \  
  --set-file storage_server_cert=/path/to/config/files/consul-0.pem \  
  /path/to/chart/proxecto-api-security-4.15.0.tgz
```

4.3. Verifying the installation of PAS in Kubernetes

If everything is correct, the Helm command will present the following output:

```
NAME: proxecto-api-security  
LAST DEPLOYED: Mon May  2 13:51:46 2022  
NAMESPACE: proxecto-api-security  
STATUS: deployed  
REVISION: 1  
TEST SUITE: None
```

1. Run the `kubectl get pods --namespace=proxecto-api-security --selector=app=proxecto-api`

`-security` command to investigate the running pods. The output shall be similar to the following example:

NAME	READY	STATUS
RESTARTS AGE		
proxedo-api-security-blob-store-86ccc6d864-frc5k0 40s	1/1	Running
proxedo-api-security-config-api-76d587d6cd-wpw5d0 40s	1/1	Running
proxedo-api-security-consul-68c5c87f75-mvlct0 40s	1/1	Running
proxedo-api-security-content-filtering-director-55b859df9-sztwp0 40s	1/1	Running
proxedo-api-security-flow-director-5cddf58677-qxczd	0/1	
ContainerCreating 0 40s		
proxedo-api-security-frontend-676bfd8956-qrtm40 40s	1/1	Running
proxedo-api-security-insight-director-585cc5f86-j8rrz	0/1	
ContainerCreating 0 40s		
proxedo-api-security-transport-director-5bbdf58d7d-whzsq	0/1	
ContainerCreating 0 40s		

The core pod is missing the core configuration, therefore it will not enter the "Running" state until the first configuration is applied in the management.

2. Run the following command to access the management component for verification.

```
kubectl port-forward --namespace=proxedo-api-security service/proxedo-api-security-frontend 8080:80
```

3. Open the <http://127.0.0.1:8080/> in the browser.

5. Base system configuration for PAS in Kubernetes

This chapter explains configuration details for setting up a working PAS. Configuration settings are detailed here, which are based on the installation of the *Helm chart*.

The *Helm chart* carries Kubernetes manifest files for each component, and requires a set of parameters to be configured by the user for the installation.

The values.yml file

The configuration of PAS components is condensed into a `values.yml` file. The default version of this file can be printed by using the following command:

```
helm show values /path/to/chart/proxedo-api-security-4.15.0.tgz
```

To configure the necessary parameters and to overwrite the not suitable default values, save the output to a file, and keep only those parts that has to be overwritten. The modified file can be provided as `--values my-values.yml` to the Helm installation command.

There are two main sections of this file:

1. Infrastructure - This section defines the options necessary for kubernetes to deploy the components.
2. Configuration - This section defines the options for PAS itself. The main configuration of the storage and management components is defined in this file.

The format of this file must adhere to the [YAML 1.1 specification](#).

There are different sections in this configuration file, some of which, as for example, the 'config.mgmt.frontend' section, might not need specific configuration. However, the default values of these sections must be set by `{}`.

For information on how to provide the custom `values.yml` file, see section [Providing the necessary files for Helm installation](#). See configuration examples in [Appendix B](#).

5.1. Infrastructure configuration

In this infrastructure part of the configuration, many parameter fields are directly associated with the configuration attributes defined for the Kubernetes objects. For such parameters that have a Kubernetes equivalent, the Kubernetes parameter is referenced in the format that can directly be used with the `kubectl explain` command. This command provides the most specific documentation of each field. However, for using this command, access to a cluster is required.

In case it is not feasible to use the `kubectl explain` command, the referenced format can also be used to navigate to the correct object and field at the following site: [Kubernetes API](#).

The following tables describe the infrastructure parameters and their Kubernetes equivalent if that exists.

Table 1. Docker-related parameters

Parameter field	Default value	Description
infrastructure.docker.registry	docker.balysys.hu	The registry to download docker images from.
infrastructure.docker.pull_policy	IfNotPresent	This parameter has a Kubernetes equivalent in all pods: <code>pod.spec.containers</code> .
infrastructure.docker.image_tag		The image tag to use instead of the one corresponding to the current PAS version.

Table 2. Storage-related infrastructure parameters

Parameter field	Default value	Description
infrastructure.storage.volume_claim		This parameter has a Kubernetes equivalent: <code>PersistentVolumeClaim</code> .
infrastructure.storage.storage_class_name		This parameter has a Kubernetes equivalent: <code>PersistentVolumeClaim.spec.storageClassName</code> .
infrastructure.storage.access_modes	ReadWriteOnce	This parameter has a Kubernetes equivalent: <code>PersistentVolumeClaim.spec.accessModes</code> .
infrastructure.storage.requests		This parameter has a Kubernetes equivalent: <code>PersistentVolumeClaim.spec.resources.requests</code> .
infrastructure.storage.requests.storage	100Mi	This parameter has a Kubernetes equivalent: <code>PersistentVolumeClaim.spec.resources.requests.storage</code> .

Table 3. Blob-store infrastructure parameters

Parameter field	Default value	Description
Resources		

Parameter field	Default value	Description
infrastructure.storage.blob_store.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.storage.blob_store.resources.auto_fill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.storage.blob_store.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.storage.blob_store.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.storage.blob_store.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .
infrastructure.storage.blob_store.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .
infrastructure.storage.blob_store.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.storage.blob_store.resources.requests.cpu	350 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.storage.blob_store.resources.requests.memory	450 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.storage.blob_store.resources.requests.ephemeral_storage	50Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .

Table 4. Consul infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.storage.consul.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.storage.consul.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.storage.consul.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.storage.consul.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.storage.consul.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .

Parameter field	Default value	Description
infrastructure.storage.consul.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .
infrastructure.storage.consul.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.storage.consul.resources.requests.cpu	350 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.storage.consul.resources.requests.memory	450 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.storage.consul.resources.requests.ephemeral_storage	50 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .

Table 5. Transport Director infrastructure parameters

Parameter field	Default value	Description
Service		
infrastructure.core.transport_director.service		This parameter has a Kubernetes equivalent: <i>service</i> .
infrastructure.core.transport_director.service.type	ClusterIP	This parameter has a Kubernetes equivalent: <i>service.spec.type</i> .
infrastructure.core.transport_director.service.ports		This parameter has a Kubernetes equivalent: <i>service.spec.ports</i> . A port with a specific <i>target_port</i> value needs to be set up for each listener port in the PAS configuration on the management interface.
infrastructure.core.transport_director.service.ports.name	HTTP	This parameter has a Kubernetes equivalent: <i>service.spec.ports.name</i> .
infrastructure.core.transport_director.service.ports.port	80	This parameter has a Kubernetes equivalent: <i>service.spec.ports.port</i> .
infrastructure.core.transport_director.service.ports.protocol	TCP	This parameter has a Kubernetes equivalent: <i>service.spec.ports.protocol</i> .
infrastructure.core.transport_director.service.ports.target_port	49 000	This parameter has a Kubernetes equivalent: <i>service.spec.ports.targetPort</i> .
infrastructure.core.transport_director.service.ports.node_port		This parameter has a Kubernetes equivalent: <i>service.spec.ports.nodePort</i> .
Resources		
infrastructure.core.transport_director.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.core.transport_director.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.

Parameter field	Default value	Description
infrastructure.core.transport_director.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.core.transport_director.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.core.transport_director.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .
infrastructure.core.transport_director.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .
infrastructure.core.transport_director.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.core.transport_director.resources.requests.cpu	250 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.core.transport_director.resources.requests.memory	450 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.core.transport_director.resources.requests.ephemeral_storage	50 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .
Scaling		
infrastructure.core.transport_director.scaling		For scaling parameters, see the separate table on scaling, Parameters for Scaling - Transport Director, Flow Director, Insight Director .

Table 6. Flow Director infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.core.flow_director.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.core.flow_director.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.core.flow_director.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.core.flow_director.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.core.flow_director.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .

Parameter field	Default value	Description
infrastructure.core.flow_director.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .
infrastructure.core.flow_director.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.core.flow_director.resources.requests.cpu	250 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.core.flow_director.resources.requests.memory	600 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.core.flow_director.resources.requests.ephemeral_storage	200 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .
Scaling		
infrastructure.core.flow_director.scaling		For scaling parameters, see the separate table on scaling, Parameters for Scaling - Transport Director, Flow Director, Insight Director .

Table 7. Insight Director infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.core.insight_director.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.core.insight_director.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.core.insight_director.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.core.insight_director.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.core.insight_director.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .
infrastructure.core.insight_director.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.core.insight_director.resources.requests.cpu	120 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.core.insight_director.resources.requests.memory	350 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
Scaling		

Parameter field	Default value	Description
infrastructure.core.insight_director.scaling		For scaling parameters, see the separate table on scaling, Parameters for Scaling - Transport Director, Flow Director, Insight Director .

Table 8. Content Filtering Director infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.core.content_filtering_director.resources		This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources</code> .
infrastructure.core.content_filtering_director.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.core.content_filtering_director.resources.limits		This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.limits</code> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.core.content_filtering_director.resources.limits.cpu		This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.limits.cpu</code> .
infrastructure.core.content_filtering_director.resources.limits.memory	1 Gi	This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.limits.memory</code> .
infrastructure.core.content_filtering_director.resources.limits.ephemeral_storage	20 Gi	This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.limits.ephemeral-storage</code> .
infrastructure.core.content_filtering_director.resources.requests		This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.requests</code> .
infrastructure.core.content_filtering_director.resources.requests.cpu	250 m	This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.requests.cpu</code> .
infrastructure.core.content_filtering_director.resources.requests.memory	600 Mi	This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.requests.memory</code> .
infrastructure.core.content_filtering_director.resources.requests.ephemeral_storage	1 Gi	This parameter has a Kubernetes equivalent: <code>pod.spec.containers.resources.requests.ephemeral-storage</code> .

Table 9. Parameters for Scaling - Transport Director, Flow Director, Insight Director

Parameter field	Default value	Description
infrastructure.core.<transport/flow/insight>_director.scaling		This parameter has a Kubernetes equivalent: <code>HorizontalPodAutoscaler</code> .
infrastructure.core.<transport/flow/insight>_director.scaling.create_autoscaler	true	This parameter defines whether to create the <code>HorizontalPodAutoscaler</code> object with the forthcoming configuration options. If it is set to false, the HPA object to enable core autoscaling will need to be created manually.

Parameter field	Default value	Description
infrastructure.core.<transport/flow/insight>_director.scaling.min_replicas	1	This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.minReplicas</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.max_replicas	10	This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.maxReplicas</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.metrics</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.cpu		This parameter defines the CPU metric configuration.
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.cpu.average_utilization	80	This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.metrics.resource.target.averageUtilization</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.memory		This parameter defines the memory metric configuration.
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.memory.average_utilization	80	This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.metrics.resource.target.averageUtilization</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior</i> . If it is defined, either <i>scale_down</i> or <i>scale_up</i> parameter must be defined.
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown</i> . If it is defined, all included parameters need to be defined.
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.stabilization_window_seconds		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.stabilizationWindowSeconds</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.policies		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.policies</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.policies.type		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.policies.type</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.policies.value		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.policies.value</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.policies.period_seconds		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.policies.periodSeconds</i> .
infrastructure.core.<transport/flow/insight>_director.scaling.metrics.behavior.scale_down.select_policy		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleDown.selectPolicy</i> .

Parameter field	Default value	Description
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp</i> . If it is defined, all included parameters need to be defined.
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.stabilization_window_seconds		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.stabilizationWindowSeconds</i> .
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.policies		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.policies</i> .
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.policies.type		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.policies.type</i> .
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.policies.value		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.policies.value</i> .
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.policies.period_seconds		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.policies.periodSeconds</i> .
infrastructure.core.<transport/flow/insight>_director. scaling.metrics.behavior.scale_up.select_policy		This parameter has a Kubernetes equivalent: <i>horizontalpodautoscaler.spec.behavior.scaleUp.selectPolicy</i> .

Table 10. Config-api infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.mgmt.config_api.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.mgmt.config_api.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.mgmt.config_api.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.mgmt.config_api.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.mgmt.config_api.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .
infrastructure.mgmt.config_api.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .

Parameter field	Default value	Description
infrastructure.mgmt.config_api.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.mgmt.config_api.resources.requests.cpu	350 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.mgmt.config_api.resources.requests.memory	450 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.mgmt.config_api.resources.requests.ephemeral_storage	100 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .

Table 11. Frontend infrastructure parameters

Parameter field	Default value	Description
Resources		
infrastructure.mgmt.frontend.resources		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources</i> .
infrastructure.core.mgmt.frontend.resources.autofill_limits	false	When true and limits are not defined, limits will be the same as the requests. When false and limits are not defined, there are no limits. Setting limits overrides this setting.
infrastructure.core.mgmt.frontend.resources.limits		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits</i> . If this is defined, either a CPU limit value, a memory limit value, or both must be provided.
infrastructure.core.mgmt.frontend.resources.limits.cpu		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.cpu</i> .
infrastructure.core.mgmt.frontend.resources.limits.memory		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.memory</i> .
infrastructure.core.mgmt.frontend.resources.limits.ephemeral_storage		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.limits.ephemeral-storage</i> .
infrastructure.core.mgmt.frontend.resources.requests		This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests</i> .
infrastructure.core.mgmt.frontend.resources.requests.cpu	350 m	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.cpu</i> .
infrastructure.core.mgmt.frontend.resources.requests.memory	450 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.memory</i> .
infrastructure.core.mgmt.frontend.resources.requests.ephemeral_storage	70 Mi	This parameter has a Kubernetes equivalent: <i>pod.spec.containers.resources.requests.ephemeral-storage</i> .

5.2. PAS configuration in Kubernetes

5.2.1. Configuration options for the storage component

The `config.storage` section controls keys to be used between the management and storage components.

The configuration file has three main sections, namely **common**, **consul** and **blob-store**.

Table 12. Storage configuration **common** options

Key	Default	Description
<code>config.storage.common.short_product_name</code>	pas	This parameter must be set to 'pas'.
<code>config.storage.common.standalone_mode</code>	true	This parameter must be set to 'true'. It denotes whether the storage is run in standalone or in cluster mode. The non-standalone mode is not relevant in Kubernetes environment, therefore it is not supported.

Table 13. Storage configuration **consul** options

Key	Default	Description
<code>config.storage.consul.bind_cluster_addr</code>	127.0.0.1	The address to bind on as a cluster member. This will be used to communicate with other members. This is a required parameter.
<code>config.storage.consul.gossip_encryption_key</code>		The encryption key to use for the gossip protocol. It is a 32-byte shared key encoded into base64 format. Use <code>openssl rand -base64 32</code> to generate it. For more information on the keys produced as part of the configuration, see Using values.yml file . This is a required parameter.
<code>config.storage.consul.log_level</code>	INFO	The log level of consul. The possible values are: TRACE, DEBUG, INFO, WARN, ERR



The options with 'N/A' default value are such sections that cannot have exact values, only the values described afterwards in the table.

Table 14. Storage configuration **blob-store** options

Key	Default	Description
<code>config.storage.blob_store.access_key</code>		The access key used for connecting to the blob-store. It starts with <code>GK</code> , followed by 12 hex-encoded bytes. This is a required parameter.
<code>config.storage.blob_store.secret_key</code>		The secret key is a 32-byte hex-encoded random string, which can be generated with a command such as <code>openssl rand -hex 32</code> . This is a required parameter.
<code>config.storage.blob_store.rpc_secret</code>		The RPC secret is a 32-byte hex-encoded random string, which can be generated with a command such as <code>openssl rand -hex 32</code> . This is a required parameter.

Key	Default	Description
config.storage.blob_store.admin_token		The token for accessing the API administration endpoints. Any string can be used for this value. A random token generated with <code>openssl rand -base64 32</code> is recommended. This is a required parameter.



The options with 'N/A' default value are such sections that cannot have exact values, only the values described afterwards in the table.

For configuration examples, see section [Minimal configuration](#).

5.2.2. Configuration options for the management component

The `config.mgmt` section controls:

- Web service parameters
- Authentication

The configuration file has two main sections, namely **frontend** and **configapi**.

The default values for both **frontend** and **configapi** sections are automatically effective. If the attributes have to be configured with specific values, other than the default values, the `{ }` curly braces have to be deleted and the new values have to be added.

Table 15. Management configuration **frontend** options

Key	Default	Description
config.mgmt.frontend.server_name	—	The hostname the web server should serve the requests on. The default value means that the management interface will be served regardless of the provided hostname.
config.mgmt.frontend.cors_api	N/A	This section configures cross origin request sharing options for API access.
config.mgmt.frontend.allow_origin		The value of the Access-Control-Allow-Origin header. This is a required parameter in case of enabled CORS API.



The options with 'N/A' default value are such sections that cannot have exact values, only the values described afterwards in the table.

Table 16. Management configuration log level setting options - **configapi** section

Key	Default	Description
config.mgmt.configapi.log_level	INFO	The log level can be set to DEBUG, INFO, WARNING, ERROR, CRITICAL.

Table 17. Management configuration user session options - **configapi** section

Key	Default	Description
config.mgmt.configapi.session	N/A	This section configures the options for session lifetimes.
config.mgmt.configapi.session.session_validity	600	The allowed lifetime of a login session token in seconds. It determines the time period between group membership and user existence checks. This DOES NOT control the length of a user session.
config.mgmt.configapi.session.renew_validity	36000	The validity of the renew token. It determines for how long session tokens can be renewed. Therefore the maximum length of a user session is the sum of the two parameters.



The options with 'N/A' default value are such sections that cannot have exact values, only the values described afterwards in the table.

For further details on **configapi** section parameters related to LDAP authentication, see [Management configuration LDAP authentication options - configapi section](#).

For configuration examples on the management component, see section [Minimal configuration](#) and section [Management configuration with LDAP authentication](#).

5.2.2.1. Configuring authentication and local users in PAS

There are two methods available to configure authentication in PAS:

- *htpasswd* authentication
- Lightweight Directory Access Protocol (LDAP) authentication



It is required to provide the *htpass* file already for the *Helm chart* installation. See section [Providing the necessary files for Helm installation](#).

Using *htpasswd* for authentication and for the configuration of local users

By using *htpasswd* authentication, the administrator can define individual user credentials directly in the *htpasswd* file. This file is created and provided for the *Helm* installation command. As local users are stored in an *htpasswd* file, the standard *htpasswd* tool needs to be used.

It is not possible to configure user groups, or to define different access levels for the users with *htpasswd* authentication, yet it is possible to define as many user credentials as necessary one by one. The user credentials are encrypted in the configuration file using the bcrypt encryption method. If you want to add new users to the *htpasswd* file, run the `htpasswd users.htpass username` command and provide the password.

Example command and output

```
$ htpasswd -B users.htpass new-user
New password:
Re-type new password:
Adding password for user new-user
```

Consider the followings related to the command and the example output:

- the `htpasswd` file is created and provided for the `Helm` installation command
- `new-user` is the name of the new user

As a result, similar content is expected to appear in the referred file:

```
new-user:$2y$05$jsvtfYMP1HJZ1WCNGV6d.j4yWU5gJ4D97Vr6z8yK9A2wy80g1iD.
```

LDAP authentication

LDAP authentication is a more elaborate way to configure authentication for PAS. With LDAP authentication it is possible to define user groups and attach different levels of access to these users, however, PAS does not support different levels of authorization based on these attributes yet at the moment.



If LDAP authentication is used, only the administrator user - and no other user - can authenticate with the `htpasswd` file.

The following **configapi** parameters, which are part of the configuration file's **configapi** section, take part in LDAP authentication:

Table 18. Management configuration LDAP authentication options - **configapi** section

Key	Default	Description
<code>config.mgmt.configapi.ldap</code>	N/A	This section configures the options for LDAP authentication. LDAP authentication is disabled by default.
<code>config.mgmt.configapi.ldap.ldap_url</code>		The URL of the LDAP server. It must start with <code>ldap[s]://</code> . This is a required parameter in case of LDAP authentication.
<code>config.mgmt.configapi.ldap.bind_user</code>		The service user to use, for searching the LDAP server. If <code>use_ntlm</code> parameter is OFF, this must be the whole DN. If it is ON, it must be the Active Directory domain and the username concatenated by a backslash (eg. <code>AD_domain\administrator</code>). This is a required parameter in case of LDAP authentication.
<code>config.mgmt.configapi.ldap.bind_password</code>		The password of the service user. This is a required parameter in case of LDAP authentication.
<code>config.mgmt.configapi.ldap.use_ntlm</code>	OFF	Set this parameter to ON to use NTLM authentication. This is only available when the LDAP server is Microsoft Active Directory.
<code>config.mgmt.configapi.ldap.tls_version</code>	TLSv1_2	The TLS version for the LDAPS connection. It must be one of the following: SSLv23, TLS, TLS_CLIENT, TLS_SERVER, TLSv1, TLSv1_1, TLSv1_2, TLSv1_3.
<code>config.mgmt.configapi.ldap.validate_cert</code>	no	Set it to yes to validate certificates.

Key	Default	Description
config.mgmt.configapi.ldap.ca_certs_file	/opt/balasy/etc/ldap_ca_certs.pem	This file contains the certificate files of the certificate authorities. Provide the path and filename for the certificate file. The certificate file must be in PEM format. See a single CA file configuration example in Single CA file example. In case a self-signed certificate is used, the server certificate must also be included in this file. In case a chain of certificates is used, the certificate of each level must be included in this file, beginning with the certificate of the signer of the server certificate, followed by the signer of that certificate up to the root certificate. For example on a Certificate chain with multiple CA, see Example on certificate chain with multiple CAs. In case multiple chains of certificates are used, the chains must be concatenated in the same file. The first matching chain will be used for verification. The above details are based on the Python SSL library documentation, available at https://docs.python.org/3.12/library/ssl.html#certificates. Use the <code>--set-file mgmt_ldap_ca_certs_file=<path/to/file></code> command during helm installation to specify this file. Also uncomment the <code>ca_certs_file</code> parameter without changing its value.
config.mgmt.configapi.ldap.user_base_dn		The base DN under which users reside. This is a required parameter in case of LDAP authentication.
config.mgmt.configapi.ldap.username_attribute	sAMAccountName	The attribute that contains the name of the user.
config.mgmt.configapi.ldap.user_object_class	user	The object class of the users.
config.mgmt.configapi.ldap.memberOf_attribute	memberOf	The attribute that contains membership information (groups) on user objects.
config.mgmt.configapi.ldap.group_base_dn		The base DN under which groups reside. This is a required parameter in case of LDAP authentication.
config.mgmt.configapi.ldap.groupname_attribute	name	The attribute that contains the name of the group.
config.mgmt.configapi.ldap.member_attribute	member	The attribute that contains membership information (users) on group objects.

Key	Default	Description
config.mgmt.configapi.ldap.group_object_class	group	The object class for groups.
config.mgmt.configapi.ldap.allowed_groups		A list of group names (as contained by 'groupname_attribute') allowed to log in. This is a required parameter in case of LDAP authentication.

6. Configuration of Proxedo API Security on the Web User Interface

This chapter explains configuration details for setting up a working Proxedo API Security (PAS) with the help of the Web User Interface.

The Proxedo API Security Web User Interface (UI) is installed together with the installation of Proxedo API Security. The URL for Proxedo API Security Web UI and the necessary credentials are generated when the management component is first started. The password for the administrator can be set or generated during installation.

For information on how to set up more users, see section *Configuring authentication and local users in PAS*.

By using OpenAPI schemas, as defined in OpenAPI specifications (also known as Swagger), Proxedo API Security verifies that the traffic passing through conforms to the API endpoint's specification. An OpenAPI Swagger schema detailing the Configuration API is available at: `<frontend_url>/api/v1/openapi`. `<frontend_url>` here refers to the URL address of the user's Proxedo API Security Web User Interface.

6.1. Minimum configuration

It is possible to run PAS with a minimum, basic configuration. For a minimum configuration the following items need to be configured in the Web UI:

- [BRICKS / File](#)
 - Name
 - Type: License
 - File

For more details on the License *File*'s requirements, see [File types](#).
- [SYSTEM / License](#)
 - License File

For more details on the *License*'s parameters, see [License's configuration](#).
- [SERVICES / Backend](#)
 - Name
 - Servers

For more details on the *Backend*'s parameters, see [Backend's configuration](#).
- [SERVICES / Endpoint](#)
 - Name
 - URLs
 - Backend
 - Request
 - Response

For more details on the *Endpoint's* parameters, see [Security Flow](#) and [Endpoint's configuration](#).

- [SERVICES / Listener](#)

- Name

- Endpoints

For more details on the *Listener's* parameters, see [Listener's configuration](#).

This basic configuration can be further improved with the completion of more configuration units later. The minimum configuration can also be used to test the installation settings.

6.2. Login Page

The main component of the Login page is the login form where the user needs to provide the credentials in order to be authorized to use the Web UI of Proxedo API Security.

As part of the initial configuration of Proxedo API Security, the administrator defines the necessary credentials, which can now be used.




Figure 3. Login page for Proxedo API Security Web User Interface

For accessing the Web User Interface:

1. Enter the valid user credentials.
2. Click the **Log In** button.

After a successful login, the user has access to the Proxedo API Security Web UI.

6.3. Proxedo API Security Web User Interface main page

The configuration elements are organized into a logical order for easier usage.




Figure 4. Proxedo API Security Web User Interface main page

6.3.1. Navigation

The PAS Web UI has the following navigation areas:




Figure 5. Navigation areas in the Proxedo API Security Web User Interface

The navigation areas are described here in more details:

Left navigation area (1)

This navigation area (1) presents the navigation units available for configuration.

When opening up the Proxedo API Security Web UI, four main navigation units are available, that is, BRICKS, PLUGINS, SERVICES, and SYSTEM.

These four main navigation units can be opened for further sub-navigation units by clicking on either the navigation item itself or on the  arrow icon next to it. Alternatively, when the sub-navigation units are not in use, they can be hidden by clicking the arrow navigation icons next to the main navigation items, or similarly

by clicking on the navigation item itself.

Top navigation area (2)

This Top navigation area (2) presents the *Changes*, *Status* and *Configuration Backup* buttons in the top left corner. For more information on these services, see [Checking and finalizing changes in Proxedo API Security configuration](#), [System-wide status information](#) and [Backup and restore running or user configuration for Proxedo API Security](#). The top right corner presents the *Help* button and a *Profile* button that shows the current user's name. The *Logout* option is present under the *Profile* button.

Main configuration area (3)


This is the main configuration area of the Web UI. Any navigation unit selected in the Left navigation area (1) presents the configuration details in this Main configuration area (3). The configuration details can be edited in this area.

In case there are already configured parameters, those are displayed in a table in the Main configuration area (3).

In order to add more configuration details, select the *New* navigation button in the upper right corner.

The Main configuration area (3) provides the following navigation and activity options. Note that some of these activities are also available when the configuration parameters are presented in list view:

Table 19. Navigation and activity options in the Main configuration area (3)

Navigation option	Description
	By selecting the <i>New</i> navigation button on the active window of a component, a new component can be configured.
	By selecting the <i>Pen</i> navigation button next to a component, the Web UI navigates back to the configuration page of the selected element. The so far configured details can be changed or new configuration details can be added.
	By selecting the <i>Copy</i> navigation button next to a component, the Web UI copies all the information of that component into a new instance, which instance can be saved with a new name, inheriting the same, copied parameters.
	By selecting the <i>Bin</i> button next to a component, the configuration element can be deleted. If an element is selected for deletion, a <i>Warning appears</i> , requesting confirmation on the deletion of the element.
	This icon is visible at the right side of every drop-down list during configuration. By selecting this icon it is possible to unselect an item of the drop-down list and to clear the selection field from any data. Clearing the field from data with the help of this icon gains importance when an earlier selected drop-down list item, saved in our configuration, has to be cleared from the configuration data.
	By selecting the <i>Next page</i> button it is possible to navigate to the next page of the parameter keys listed.

6.3.2. Naming Configuration components in the Web UI

When configuring the Proxedo API Security Web UI, name the configuration components with the usage of the English alphabet and numerals. When the name is composed of more than one word, use underscore. It is not allowed to use spacing or any special characters though.

6.4. BRICKS - Configuration units

Bricks are reusable components. They do not provide a complete security function themselves, instead, they are used as building blocks elsewhere (hence the name). They can be used by *Plugins* (like *Selectors*), or utilized by

other bricks (like Extractors).

Certain bricks are so called *default* objects, which are in 'read-only' state and cannot be configured or modified. Such default objects are listed in the following table:

Table 20. Default objects - BRICKS

Default object name	Class
always	Matcher
never	Matcher
content_type_json	Matcher
content_type_json_pattern	Matcher
json_content	Matcher
content_type_xml_base	Matcher
content_type_xml_dtd	Matcher
content_type_xml_ext_parsed	Matcher
content_type_xml_pattern	Matcher
content_type_xml_text	Matcher
content_type_xml_text_ext_parsed	Matcher
xml_content	Matcher
error_policy	Error policy
enforcer_default	Error policy
insight_default	Error policy
client_address	Selector
client_port	Selector
server_address	Selector
server_port	Selector
error_policy	Selector
error_policy_action	Selector
error_policy_status_code	Selector
error_policy_silent	Selector
error_policy_message	Selector
plugin_name	Selector
plugin_verdict	Selector
plugin_error_message	Selector

These default objects are listed under the actual classes in the Web UI.

The *BRICKS* main page in the Web UI is as follows:




Figure 6. The BRICKS main page in the Web User Interface

1. Click on the *BRICKS* main configuration item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Click on the sub-navigation unit you would like to configure. The details of the sub-navigation menu open up in the Main configuration area.

6.4.1. Error Policy

Error Policies define how to proceed if a *Plugin* decides to have found an error. For example, when an *Enforcer plugin* decides that the call is invalid.


It is the error policy that enables the user to act differently in case the error appears in a request or a response.

Every Plugin has a default error policy, namely, the 'error_policy', except for the Enforcer and the Insight Plugins, which have their own default error policies already configured for usage, the enforcer_default and the insight_default error policies.

6.4.1.1. Configuring Error Policies

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *Error Policy*.

The configuration window that appears presents the default error policies, as listed in [Default objects - BRICKS](#) and the configuration values already set by the user:



Changes Status Configuration Backup

Help user

Error Policy

New

Name	Request action	Request code	Request silent	Response action	Response code	Response silent	Actions
enforcer_default	Abort	422	False	Abort	502	True	Edit Add Delete
error_policy	Abort	400	False	Abort	502	True	Edit Add Delete
insight_default	Log	400	False	Log	502	True	Edit Add Delete

< 1 >

Figure 7. Error policy's main page in the Web User Interface

- Click on the *New* navigation button to create an Error Policy.

Error Policies have default values for each of their fields. They form a strict security policy: all errors are fatal, and only errors made by the client are reported in detail.

- Configure the necessary parameters for the error policy based on the details provided in the table [Error policy configuration options](#).
- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The error policies configured here can be used in the *Plugin's* configuration, by referencing their name.

The following values can be configured for the *Error Policy Brick*:

Error Policy

Name : *
 Request ^
 Action :
 Silent :
 Code :
 Message :
 Response ^
 Action :
 Silent :
 Code :
 Message :
 Additional Error Headers :
 Validate Save Cancel

Enter Name
 Choose Action (Default: Abort)
 False Default True (Default: False)
 Enter or choose Code (Default: 400)
 Enter Message (Default: Request Error)
 Choose Action (Default: Abort)
 False Default True (Default: True)
 Enter or choose Code (Default: 500)
 Enter Message (Default: Response Error)
 Header Name Header Value +

Figure 8. Configuring error policies in the Web User Interface

Table 21. Error policy configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the error policy. This name of the error policy can be referenced from other parts of the configuration, that is, the error policy is reusable.
Request Action	The available values are: - Abort - Log	Abort	It defines what action shall take place if there is an error on the request side: - Abort: the request is denied if the <i>Plugin</i> fails. Use the other parameters to control the content of the error sent to the client. - Log: the invalid requests are allowed, but are logged.
Request Silent	True or False.	False	When turned on, the <i>Plugins</i> do not report on the denial of the invalid request. When turned off, the <i>Plugins</i> have the ability to report the error in detail in the body of the HTTP error request.
Request Code	The values are available from a drop-down list. If the elements of the drop-down list are selected, it will make the list of the actual request codes visible. The applicable request code can be selected.	422	The HTTP status code to be used when denying invalid requests.

Key	Values	Default value	Description
Request Message	The message can be provided in free text.	Request error	The HTTP response line when denying invalid requests.
Response Action	The available values are: • Abort • Log	Abort	It defines what action shall take place if there is an error on the request side: • Abort: the request is denied if the <i>Plugin</i> fails. Use the other parameters to control the content of the error sent to the client. • Log: the invalid requests are allowed, but are logged.
Response Silent	True or False.	True	When turned on, the <i>Plugins</i> do not report on the denial of the invalid response. When turned off, the <i>Plugins</i> have the ability to report the error in detail in the body of the HTTP error response.
Response Code	The values are available from a drop-down list. Note that the response codes are grouped, so that if the elements of the drop-down list are selected, further groups of response codes will be made visible in a tree structure. The applicable request code can be selected.	502	The HTTP status code to be used when denying invalid requests.
Response Message	The message can be provided in free text.	Response error	The HTTP response line when denying invalid requests.
Additional Error Headers	A list of additional header name-value pairs.	Empty list	A list of HTTP header key-value pairs to include in the response when denying invalid requests.

6.4.2. Matcher

Matchers decide if the Plugin should be executed for a given call by checking various data in the HTTP message. They provide an extremely versatile way of defining the circumstances that must be met for the *Plugin* to execute.

Matchers need four pieces of information:

- **Name:** The **Name** field can be defined in free text and it is not related to the extractor that will be used. This **Name** can be referenced in Plugins.
- **Type:** This parameter defines what part of the call needs to be checked.
- **Comparator:** The Comparator shows by what means the collected value of the call is compared with the provided pattern. (Some comparators also take flags or arguments.)
- **Expression:** A regular expression specifies a set of strings that match it. A complete explanation on how to write expressions is not in the scope of this document.

The matchers can be used in Plugin configurations' match option by referencing their name.



There are some named Matchers available without explicit configuration:

- **always** and **never** are instances of Always matcher and Never matcher.
- **json_content** that matches requests with the Content-Type headers representing JSON.

Also note that no other matchers can be defined with these names.

Matchers internally utilize Extractors to fetch the information from the call to compare with. The **Type** of the matcher resembles the name of the extractor that will be used.

All matchers have a default comparator that is applied implicitly.




If you want to use comparator parameters, the comparator name should be given even if the default comparator is used.

6.4.2.1. Configuring Matchers

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *Matcher*.

The configuration window that appears presents the default matchers, as listed in [Default objects - BRICKS](#) and the configuration values already set by the user:



Name	Type	Comparator	Expression	Actions
always	Always			
content_type_json	Content Type	Equals	application/json	
content_type_json_pattern	Content Type	Pattern	application/*[.*]json	
content_type_xml_base	Content Type	Equals	application/xml	
content_type_xml_dtd	Content Type	Equals	application/xml-dtd	
content_type_xml_ext_parsed	Content Type	Equals	application/xml-external-parsed-entity	
content_type_xml_pattern	Content Type	Pattern	application/*[.*]xml	
content_type_xml_text	Content Type	Equals	text/xml	
content_type_xml_text_ext_parsed	Content Type	Equals	text/xml-external-parsed-entity	
json_content	Any			
never	Never			
xml_content	Any			

Figure 9. Matchers' main page in the Web User Interface

3. Click on the *New* navigation button to create a Matcher.
4. Provide the name of the matcher.
5. Choose the type of the matcher from the drop-down list.
6. Configure the necessary parameters with the help of the below tables.
7. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.

8. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Matcher Brick*:




Figure 10. Configuring matchers in the Web User Interface

Table 22. Matcher configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.	It can be defined in free text.	The Name of the matcher which can be referenced in Plugins.
Type*	For the available values, see Matcher types .		The preferred matcher type has to be selected from the drop-down list.

Matcher types

Depending on the choice of the matcher type, some more required configuration fields might appear on this page. The following tables describe the matcher types in details and provide the necessary information for the additional configuration fields, required for setting the matcher types:

- [Matcher types and their settings - Simple matchers](#)
- [Matcher types and their settings - Compound matchers](#)
- [Matcher types and their settings - URI matchers](#)
- [Matcher types and their settings - SOAP matchers](#)

Table 23. Matcher types and their settings - Simple matchers

Matcher type	Description
Always	This matcher always matches.
Never	This matcher never matches. It can be used to turn off a <i>Plugin</i> .
Call Direction	Matches the direction of the message (request or response).
Backend Response Time	Matches the time spent between the sending the request towards the server and receiving the response from the server, in milliseconds. Only matches in a response flow.

Matcher type	Description
Method	Matches the HTTP method of the request. Note that the standard and the practice differs regarding upper and lower casing, set case sensitivity according to needs. When choosing the <i>Method</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
Header	It matches the value of an HTTP header. Some HTTP headers can be present more than once in a call. To accommodate this, matching is completed against the value of each occurrence of the header. Matching occurs if there is any match. For example, if the <i>Accept</i> header was repeated as follows: <pre>Accept: application/json Accept: application/xml</pre> Consequently, in this example above both <code>header.accept: application/json</code> and <code>header.accept: application/xml</code> would match. To match against the header named server the key will be <code>header.server</code>, possibly completed with comparator specification, like <code>header.server.regex.ignorecase</code>.  While the values are not, the HTTP header names are case insensitive, so you can write them all lowercase in the configuration key. The syntax of this matcher differs from the others because the name of the <i>Header</i> must be added.  While the values are not, the HTTP header names are case insensitive, so you can write them all lowercase in the configuration key.
Cookie	Matches the value of a key in the Cookie HTTP header. A Cookie header key can be present more than once in a call. To accommodate this, matching is completed against the value of each occurrence of the key. Matching occurs if there is any match.
Content Type	Matches the content type of the message. It is a more robust solution than using the <i>Header</i> matcher on the <i>Content-Type</i> header because that can contain parameters as well. When choosing the <i>Content type</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
Status	Matches the status code of the response.  See the default Status class comparator which allows convenient matching on HTTP status classes. When choosing the <i>Status</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

Matcher type	Description
Raw Content	Matches the raw bytes of the request or response. It requires an expression in the form of a hexadecimal string. For example, for matching a PNG image file, the expression shall be '89504e470d0a1a0a', which is equivalent to '89 50 4e 47 0d 0a 1a 0a', as whitespaces can also be used.
Text Content	Matches the request's or response's content as a decoded string.
Client Address	Matches the client's IP address (both IPv4 and IPv6). Use the <i>subnet</i> type comparator with that matcher type. The <i>subnet</i> comparator examines if the IP address of the Client is in the specified subnet. The format for the input of the subnet comparator is the CIDR notation for IPv4 (for example, 192.0.2.0/24) and canonical prefix notation for IPv6 (for example, 2001:db8::/32).
Client Port	Matches the client's port (TCP).
Server Address	Matches the server's IP address (both IPv4 and IPv6). Use the <i>subnet</i> type comparator with that matcher type. The <i>subnet</i> comparator examines if the IP address of the Server is in the specified subnet. The format for the input of the subnet comparator is the CIDR notation for IPv4 (for example, 192.0.2.0/24) and canonical prefix notation for IPv6 (for example, 2001:db8::/32).
Server Port	Matches the server's port (TCP).
Error Policy Action	Matches the <i>Request Action</i> or <i>Response Action</i> field of the <i>Error Policy</i> of the <i>Plugin</i> selected by the <i>plugin</i> field. If the <i>plugin</i> field contains the special value "Previous", the data will be extracted from the last evaluated plugin in the same HTTP request or response, if there is one. If the <i>plugin</i> field refers to a specific plugin instance, the data will be extracted from the last evaluation of the referred plugin instance. If the matcher is in the HTTP response, and the last evaluation is in the HTTP request, then the evaluation result from the HTTP request will be selected. Only plugins with a negative verdict will return data.
Error Policy Status Code	Works like the <i>Error Policy Action</i> matcher, but matches only the <i>Request Code</i> or <i>Response Code</i> field of the referred plugin's <i>Error Policy</i> .
Error Policy Silent	Works like the <i>Error Policy Action</i> matcher, but matches only the <i>Request Silent</i> or <i>Response Silent</i> field of the referred plugin's <i>Error Policy</i> .
Error Policy Message	Works like the <i>Error Policy Action</i> matcher, but matches only the <i>Request Message</i> or <i>Response Message</i> field of the referred plugin's <i>Error Policy</i> .
Plugin Name	Works like the <i>Error Policy Action</i> matcher, but matches the name of the referred plugin. Returns data regardless of verdict.
Plugin Verdict	Works like the <i>Error Policy Action</i> matcher, but matches the verdict of the referred plugin. Returns data regardless of verdict.
Plugin Error Message	Works like the <i>Error Policy Action</i> matcher, but matches the error message provided in the referred plugin's negative verdict, if there is one.

Matcher type	Description
XPath	Matches the data from the body of an XML call with the help of the XPath expression. XPath is a query language for XML. It is a very versatile tool for extracting the needed information from the body of the call, and organizing it according to needs. A complete explanation on how to write XPath expressions is not in the scope of this document. To learn more about it visit the main website. For more details on XPath configuration options, see XPath extractor configuration options.
JMESPath	Matches the data from the body of a JSON call with the help of the JMESPath expression. JMESPath is a query language for JSON. It is a very versatile tool for extracting the needed information from the body of the call, and for organizing it according to needs. A complete explanation on how to write JMESPath expressions is not in the scope of this document. To learn more about it visit the JMESPath website: • There is a tutorial. • There are examples. • There is also a formal specification. When choosing the <i>JMESPath</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.  The result of the JMESPath expression should be a string when using string comparators (<i>Equals</i>, <i>Starts with</i>, etc.), and number when using number comparators (<i>Min</i>, <i>Max</i>, <i>Range</i>). In case of boolean or complex types, convert to string in the JMESPath expression and use the string representation of the result. Example: instead of comparing the boolean result of <code>address != ''</code>, use <code>to_string(address != '')</code> with a string comparator.
Fraud Detector Score	Matches the score value provided by the <i>Fraud Detector</i> plugin.

Table 24. Matcher types and their settings - Compound matchers

Any	Any is a Compound matcher that matches if any of its sub-matchers matches. The sub-matcher can also be a compound matcher.
All	All is a Compound matcher that matches if all of its sub-matchers match. The sub-matcher can also be a compound matcher.
None	None is a Compound matcher that matches if none of its sub-matchers match. The sub-matcher can also be a compound matcher.
One	One is a Compound matcher that matches if exactly one of its sub-matchers matches. The sub-matcher can also be a compound matcher.

Table 25. Matcher types and their settings - URI matchers

Matcher type	Description
URI matchers	A range of matchers is available to match different parts of the URI. The structure of an URI looks as follows: <pre>scheme://[username[:password]@]host[:port][/path][?query][#fragment]</pre> That is, for example: <pre>https://john.doe:secret123@example.com:8443/some/resource?foo=bar&baz=qux#some-anchor</pre>  The fragment part is used by the client locally, and is never sent in the HTTP requests, therefore PAS cannot do anything with it. These matchers use the URI extractors. It has an extensive list of examples of what each extractor extracts from the URI.
URI	Matches against the whole request URI as received from the client. When choosing the <i>URI</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI netloc	Matches the network location in the URI. It includes: • username and password if present • host • port if present unless scheme default  If the port is the default port for the scheme - that is 80 and 443 for HTTP and HTTPS, respectively - the port will not be included even if explicitly sent by the client. Therefore if the client used http://example.com:80/path then the <i>netloc</i> would be http://example.com, not http://example.com:80. When choosing the <i>URI netloc</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

Matcher type	Description
URI origin	Matches the <i>origin</i> part of the URI. It includes: • scheme • host • port if present, unless the default port for the scheme is used  If the port is the default port for the scheme - that is 80 and 443 for HTTP and HTTPS, respectively - the port will not be included, even if explicitly sent by the client. Therefore if the client used <code>http://example.com:80/path</code>, then the <i>origin</i> would be <code>http://example.com</code>, not <code>http://example.com:80</code>. When choosing the <i>URI origin</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI scheme	Matches the scheme of request (http or https). Note that the scheme is case insensitive by definition, therefore the case will always be ignored. When choosing the <i>URI scheme</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI username	Matches the <i>username</i> in the request if present. When choosing the <i>URI username</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI password	Matches the <i>password</i> in the request if present. When choosing the <i>URI password</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI host	Matches the <i>host</i> in the request. When choosing the <i>URI host</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI port	Matches the <i>port</i> of the request. Note that this matches the default <i>port</i> — that is 80 and 443 for HTTP and HTTPS, respectively — even if it is not explicitly in the request. When choosing the <i>URI port</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

Matcher type	Description
URI path	Matches the <i>path</i> part of the URI. The path is normalized to allow more robust matching and cleaner reporting. This means that: • If the path is missing <code>/</code> it is extracted. • Repeating forward-slash (<code>/</code>) characters are replaced with a single one. • dot (<code>.</code>) and double-dot (<code>..</code>) path segments are resolved. Consequently, if the path present in the <i>URI</i> would be <code>//some/./nothere/./resource///./somewhere</code> the <i>path</i> would be <code>/some/resource/somewhere</code>. If you need to match the <i>path</i> exactly as received, use URI raw path matcher. When choosing the <i>URI path</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI raw path	Matches the <i>path</i> part of the URI, without the normalization of URI path matcher carried out.  If the <i>path</i> is missing, the match still runs against a single forward slash (<code>/</code>). It is recommended to use URI path matcher unless there is an explicit need for matching the raw path. One such example would be logging or filtering out badly formed requests. When choosing the <i>URI raw path</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI raw query	Matches the <i>query</i> part of the URI as a string. It is recommended to use URI query parameter matcher unless there is an explicit need for matching the raw string. An example on this might be if there is a match on <code>foo=barbar</code> or <code>tofoo=bar</code> as well, even though it was not intended. When choosing the <i>URI raw query</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
URI query parameter	Matches the value of a query parameter. It is also valid for URIs to include a query parameter more than once. That is, it could be <code>foo=bar&qux=quz&foo=baz</code>. To accommodate this, matching is done against the value of <i>each</i> occurrence of the parameter. Matching occurs if any value is matched. When choosing the <i>URI query parameter</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

Table 26. Matcher types and their settings - SOAP matchers

Matcher type	Description
SOAP Matchers	A range of matchers is available to match different parts of the SOAP message. These matchers extend the XPath matcher with predefined expressions. They use the SOAP extractors. It has an extensive list of examples of what each extractor extracts from the SOAP message. When choosing the <i>SOAP Matchers</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP version	Matches the SOAP message version. It identifies with the SOAP namespace. The possible values are: • soapv1_1 - the message version is SOAP v1.1 • soapv1_2 - the message version is SOAP v1.2 When choosing the <i>SOAP version</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP envelope	Matches the SOAP envelope. When choosing the <i>SOAP envelope</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP header	Matches the SOAP header. When choosing the <i>SOAP header</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP body	Matches the SOAP body. When choosing the <i>SOAP body</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP fault	Matches the SOAP fault. When choosing the <i>SOAP fault</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP fault code	Matches the SOAP fault 'code'. The expression depends on the SOAP version. • faultcode - it is the SOAP v1.1 node tag. • Code - it is the SOAP v1.2 node tag. When choosing the <i>SOAP fault code</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

Matcher type	Description
SOAP fault detail	Matches the SOAP fault 'detail'. The expression depends on the SOAP version. • Detail - it is the SOAP v1.1 node tag. • Detail - it is the SOAP v1.2 node tag. When choosing the <i>SOAP fault details</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP 1.1 fault faultstring	Matches the SOAP fault 'faultstring'. This matcher only works with SOAP version 1.1. When choosing the <i>SOAP 1.1 fault faultstring</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP 1.1 fault faultactor	Matches the SOAP fault 'faultactor'. This matcher only works with SOAP version 1.1. When choosing the <i>SOAP 1.1 fault faultactor</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP 1.2 fault reason	Matches the SOAP fault 'Reason'. This matcher only works with SOAP version 1.2. When choosing the <i>SOAP 1.2 fault reason</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP 1.2 fault node	Matches the SOAP fault 'Node'. This matcher only works with SOAP version 1.2. When choosing the <i>SOAP 1.2 fault node</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.
SOAP 1.2 fault role	Matches the SOAP fault 'Role'. This matcher only works with SOAP version 1.2. When choosing the <i>SOAP 1.2 fault role</i> matcher from the drop-down list, additional parameters appear. For more information on the configuration of these parameters, see Matcher types' additional configuration options.

For details on comparator types, see [Types of comparators](#).

Depending on the matcher type selected, the administrator might need to fill in further parameters. These parameters are described in the following table.

Table 27. *Matcher types' additional configuration options*

Key	Values	Default value	Description
Comparator			The matchers need the information on the Comparator, which shows by what means the collected value of the call is compared with the provided pattern.
Type*	The available comparator types can be checked from the drop-down list.	Equals	This configuration option has to be defined for the Comparator. For details on the comparator types, see Types of comparators .

Key	Values	Default value	Description
Ignorecase	True or False.	False	This configuration option has to be defined for the Comparator. It sets the IGNORECASE flag for the selected comparator type. For matcher types that work with numeric data type or with IP addresses, the 'Equals' and 'Not Equals' comparator types do not have ignorecase field.
Expression*			This configuration option has to be defined for the Comparator. A regular expression specifies a set of strings that match it.
JMESPath Expression*	A valid JMESPath expression in text.		A complete explanation on how to write JMESPath expressions is not in the scope of this document. To learn more about it visit the JMESPath website: • There is a tutorial. • There are examples. • There is also a formal specification.
Query Parameter			It is also valid for URIs to include a query parameter more than once. That is, it could be <i>foo=bar&qux=quz&foo=baz</i> . To accommodate this, matching is done against the value of <i>each</i> occurrence of the parameter. Matching occurs if any value is matched.
Header			Extracts the value of an HTTP header. It is valid for some HTTP headers to be present more than once in a call. In this case, all the values are extracted as a list. It provides the name of the header in the configuration.
Namespaces	A list of key and expression pairs, in text.		The namespaces to use during extraction.
XPath Expression	A valid XPath expression in text.		A complete explanation on how to write XPath expressions is not in the scope of this document. • There is a tutorial. • There are examples. • There is also a formal specification.
Multiline			Sets the Multiline flag for the <i>Regex</i> comparator.
Minimum*			Matches if the pattern is larger or equal to the value.
Maximum*			Matches if the pattern is smaller or equal to the value.

Key	Values	Default value	Description
Source Plugin	Reference to a <i>Fraud Detector Plugin</i> or "Last".	Last: In case there are more Fraud Detector plugins defined in the Security Flow, by using this default value, the selector will use the score value provided for the last run Fraud Detector plugin.	The Fraud Detector plugin to be used in case there are more than one defined.
Plugin	Reference to a <i>Plugin</i> or "Previous".	Previous: A special value that dynamically selects the plugin that was most recently evaluated in the security flow.	The plugin from which data will be extracted.

6.4.3. Selector

Selectors are responsible for collecting information from the call. They utilize [Extractor bricks](#) for this purpose.

Most extractors return simple string values. However, some (might) return dictionaries. For example, you can get all the HTTP headers, or all the URI query parameters.

They are used by [Insight plugins](#) and [Fraud Detector plugins](#).

6.4.3.1. Configuring Selectors

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *Selector*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 11. Selector main page in the Web User Interface

- Click on the *New* navigation button to create a Selector.
- Name the *Selector* key.
- Fill in any more desired parameters.
- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Selector Brick*:




Figure 12. Configuring Selector in the Web User Interface

Table 28. Selector configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name of the parameter can be referenced.
Type*	Choose the selector type from the drop-down list. For more details on the values, see Extractor types .		Extractors are used to extract data from the call. They are utilized by Selector (and Matcher as well). Extractors are included by their type in Selectors, and are used by a special syntax in matchers. For details, see Extractors and Extractor types .
Save As Key	Text with ASCII characters. Space, '=', '' and ']' are not allowed.	A special token, "<Selector Name>", which will use the Selector's name.	The key under which the results of a selector are saved in the <i>Insight</i> plugin's dictionary.

Depending on what value is selected for the *Type* parameter, additional parameters might appear for configuration. The following table provides details on these additional parameters.

Table 29. Additional Selector configuration options

Key	Values	Default value	Description
Save Under Key	True or False.	True	When set to False, returned dictionaries are merged into the <i>Insight</i> plugin's dictionary instead of being stored under a key. NOTE: this could lead to conflicting keys if multiple selectors would return the same key, in this case, keys could be overwritten.
Clear Text	True or False.	False	When turned on, whitespaces are stripped from the beginning and end of the result.
Namespaces	A list of key and expression pairs, in text.		The namespaces to use during extraction.
XPath Expression	A valid XPath expression in text.		A complete explanation on how to write XPath expressions is not in the scope of this document. • There is a tutorial. • There are examples. • There is also a formal specification.
JMESPath Expression*	A valid JMESPath expression in text.		A complete explanation on how to write JMESPath expressions is not in the scope of this document. To learn more about it visit the JMESPath website: • There is a tutorial. • There are examples. • There is also a formal specification.
Expression*			A regular expression specifies a set of strings that match it.

Key	Values	Default value	Description
Time Format	A valid time format string in text.	YYYY-MM-DDT HH:mm:ss.SSSS SSZZ (line breaks for display purposes only)	The time format to use, see: Timestamp format options .
Time Zone	A time zone specifier in text.	UTC	The name of the time zone, or the time zone offset. The time zone can be specified by using the name, for example, "Europe/Budapest", or as the time zone offset in +/-HH:MM format, for example, +01:00.
Source Plugin	Reference to a <i>Fraud Detector Plugin</i> or "Last".	Last: In case there are more Fraud Detector plugins defined in the Security Flow, by using this default value, the selector will use the score value provided for the last run Fraud Detector plugin.	The Fraud Detector plugin to be used in case there are more than one defined.
Plugin	Reference to a <i>Plugin</i> or "Previous".	Previous: A special value that dynamically selects the plugin that was most recently evaluated in the security flow.	The plugin from which data will be extracted.
Include Request Counter	True or False.	True	When turned on, the request counter is included in the Session ID. See [session-id] for details.

6.4.4. Insight Target

Insight Target bricks define where the data collected by the [Insight](#) will be sent to.

The **Insight Target** configuration tree contains named *Insight Targets* with their respective configuration.

See the [Insight Target configuration options](#) for the available *Insight Target* types and their configuration options.

6.4.4.1. Data flattening

To ensure compatibility with a wide range of *Insight Target* types, the results collected by the *Insight plugin* can be flattened. The path inside the complex data structure is encoded into the key for each value:

- The merged key describes the path to the value in the data structure as a string.
- The parts of the path will be separated by a separator character, forward slash by default ("/").
- Keys in nested dictionaries are added to the path by name.

- List items are added to the path by their index.

For example, take the following data structure:

```
{
  "a": 1,
  "b": 2,
  "c": [
    {
      "d": [2, 3, 4],
      "e": {
        "f": 5,
        "g": 6
      }
    }
  ]
}
```

This will be flattened to this:

```
{
  "a": 1,
  "b": 2,
  "c/0/d/0": 2,
  "c/0/d/1": 3,
  "c/0/d/2": 4,
  "c/0/e/f": 5,
  "c/0/e/g": 6
}
```

6.4.4.2. Configuring Insight Targets

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *Insight Target*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 13. Insight Target main page in the Web User Interface

- Click on the *New* navigation button to create an Insight Target.
- Provide the name for your *Insight Target* configuration.
- Select the *Insight Target* type.
- Continue with the Syslog, Elastic and Local log configurations with the help of the following tables: [Syslog Insight Target configuration parameters](#), [Elastic Insight Target configuration parameters](#) and [Local log Insight Target configuration parameters](#).
- Configure any more desired parameter details.
- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Insight Target Brick*:




Figure 14. Configuring Insight Target in the Web User Interface

Table 30. Insight Target configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the <i>Insight Target</i> . This name of the <i>Insight Target</i> can be referenced from other parts of the configuration.
Type*	The values can be selected from the drop-down list. The available values are: • Local log • Syslog • Elastic		• Local log: Logs the result of the insight in the local system log. For more details on configuration settings with Local log, see Local log Insight Target configuration parameters. • Syslog: Sends the insight to a remote syslog server using the IETF syslog protocol defined in RFC5424. For more details on configuration settings with syslog, see table Syslog Insight Target configuration parameters. • Elastic: Sends the insight to an <i>Elasticsearch</i> engine in JSON. Since <i>OpenObserve</i> acts as a drop-in replacement for <i>Elasticsearch</i>, this target can be configured for <i>OpenObserve</i> instances. For more details on configuration settings, see Elastic Insight Target configuration parameters.
Flatten	True or False.	True	Flatten the <i>Insight Target</i> message.
Flatten Separator		/	The separator in the flattened message.
Level		3	The log level for the logged message.
Tag	The value can be selected from a drop-down list.	info	The log tag for the logged message.

The following table presents the configuration parameters for the Local log *Insight Target* type:

Table 31. Local log *Insight Target* configuration parameters

Key	Values	Default value	Description
Flatten	True or False.	True	Flatten the <i>Insight Target</i> message.
Flatten separator		/	The separator in the flattened message. Only visible when Flatten is True.
Level		3	The log level for the logged message.
Tag		info	The log tag for the logged message.

The following table presents the configuration parameters for the syslog *Insight Target* type:

Table 32. Syslog *Insight Target* configuration parameters

Key	Values	Default value	Description
Flatten	True or False.	True	Flattens the <i>Insight Target</i> message. Only visible when Data Format is JSON.
Flatten Separator		/	The separator in the flattened message. Only visible when Flatten is True.
Remote Connection	• Host: Hostname or IP address as text. • Port: The available values are integers. • Protocol: The available values are: TCP and UDP. • IP Protocol: The available values are: 4 and 6, corresponding to IPv4 and IPv6. • Use TLS: True or False. • Syslog TLS*: Reference to a <i>TLS Brick</i> of type <i>Syslog TLS</i>.	• Protocol: TCP, Port: 601 (6514 if <i>Use TLS</i> is <i>True</i>.) • Protocol: UDP, Port: 514 • IP Protocol: 4 • Use TLS: False	• Host: The hostname or the IP address of the syslog server. • Port: Add the port number here to connect to the remote system. • Protocol: Add the transport protocol to send messages over. The available values are: TCP and UDP. • IP Protocol: The internet protocol version of the given driver. • Use TLS: It enables using TLS for the Syslog communication. • Syslog TLS*: It is mandatory if the <i>Use TLS</i> option is set to True.
Flush Lines			It specifies how many lines are flushed to a destination at a time. The <i>Insights Director</i> waits for this number of lines to accumulate and sends them off in a single batch. Increasing this number increases the throughput, as more messages are sent in a single batch, but also increases the message latency.
Data Format	The possible values are: SData, JSON.	SData	This is the data format of the insight.
Include Message	True or False.	False	Whether to include the Insight plugin's Message field in the JSON output.
Message Key	Free text.	message	The key where the Insight plugin's Message field will appear in the JSON output.
Second Fraction Digits	Integer between 0 and 6 inclusive	3	The number of digits representing the fractions of seconds in the Syslog timestamp.
Time Zone	See table Time zones for time zone names.	GMT	The name of the time zone (for example, "Europe/Budapest") or the time zone offset in +/-HH:MM format (for example, +01:00).
Report Config Load	True or False.	False	It reports the event of a configuration being loaded with a cryptographic hash of the loaded configuration. This informs the <i>Insight Target</i> about changes in the configuration.

Key	Values	Default value	Description
Mask Credit Card Numbers	True or False.	False	It masks the middle section of recognised credit card numbers in any fields of the log message. Recognised credit cards are from one of the following issuers: American Express, Discover Card, Mastercard, VISA.
Enable Heartbeat	True or False.	False	It enables sending heartbeat (-- MARK --) messages to the <i>Insight Target</i> .
Heartbeat	- Frequency: A number greater than or equal to 1. - Mode: The possible values are: 'idle' (heartbeat messages are only sent when there is no traffic towards the <i>Insight Target</i>) and 'periodical' (heartbeat messages are sent regardless of activity).	- Frequency: 30 - Mode: 'periodical'	- Frequency: The number of seconds between heartbeat messages. - Mode: The operation mode of the heartbeat functionality.

The following table presents the configuration parameters for the Elastic *Insight Target* type:

Table 33. Elastic *Insight Target* configuration parameters

Key	Values	Default value	Description
Flatten	True or False.	True	It flattens the <i>Insight Target</i> message.
Flatten Separator		/	The separator in the flattened message. Only visible when Flatten is True.
Remote Connection			Settings related to the remote connection.
Include Message	True or False.	False	Whether to include the Insight plugin's Message field in the JSON output.
Message Key	Free text.	message	The key where the Insight plugin's Message field will appear in the JSON output.
Username*			The username to authenticate with on the servers.
Password*			The password to authenticate with on the servers.
Servers*	There are two values to be configured: Host*: The hostname or IP address of the Elasticsearch instance. Port: The port on host to connect to. Defaults to 9200. (Add the values by clicking the '+' sign.)		The list of Elasticsearch servers. Messages will be load balanced between servers if multiple servers are given.

Key	Values	Default value	Description
Index*			The name of the index in the Elasticsearch instance.
Use TLS	True or False.	False	Enables using TLS in the connection towards the Elastic servers.
Elastic TLS*	Reference to a <i>TLS Brick</i> of type <i>Elastic TLS</i> .		The TLS configuration towards the Elastic servers. Mandatory if <i>Use TLS</i> is set to <i>True</i> .
Workers		4	The number of workers to use for communicating with the Elasticsearch servers. Should at least equal the number of servers.
Mask Credit Card Numbers	True or False.	False	It masks the middle section of recognised credit card numbers in any fields of the log message. Recognised credit cards are from one of the following issuers: American Express, Discover Card, Mastercard, VISA.

6.4.5. TLS

Transport Layer Security (TLS) is the cryptographic protocol that secures HTTPS communications. PAS can apply TLS encryption both when communicating with Clients and Backends. TLS encryption can also be used with *Syslog* and *Elastic Insight Targets*.

When HTTPS is used the *TLS* settings must be configured.

6.4.5.1. Configuring TLS Bricks

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *TLS*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:





Figure 15. TLS main page in the Web User Interface

3. Click on the *New* navigation button to create a TLS.

The following values can be configured for the *TLS Brick*:



TLS

Name : *

Type : *

Enable Certificate : ☐ ☒ ☐ (Default: False)

Options ^

Ciphers : (Default: ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384)

Minimum TLS Version : (Default: TLS v1.2)

Maximum TLS Version : (Default: TLS v1.3)

Timeout : (Default: 300)

Enable Session Cache : ☐ ☒ ☐ (Default: False)

Disable Ticket : ☐ ☒ ☐ (Default: True)

Enable Verification : ☐ ☒ ☐ (Default: False)

Figure 16. Configuring TLS in the Web User Interface

The configuration of the first two parameters determines the TLS type and from these two steps on, it is either a *Backend-facing TLS* configuration, a *Client-facing TLS* configuration, a *Syslog TLS* configuration or an *Elastic TLS* configuration.

6.4.5.1.1. Configuring Client-facing TLS Bricks

The following parameters need to be configured for *Client-facing TLS*:

TLS

Name : *	Enter Name	
Type : *	Client-facing TLS	
Certificate ^		
Certificate File : *	Choose File	
Key File : *	Choose File	
Options ^		
Ciphers :	Enter Ciphers	(Default: ECDHE-ECDSA-AES128-GCM-SHA256:ECDSA-AES128-GCM-SHA256:ECDSA-AES256-GCM-SHA384:ECDSA-AES256-GCM-SHA384:ECDSA-AES128-GCM-SHA256:ECDSA-AES256-GCM-SHA384:ECDSA-AES128-GCM-SHA256:ECDSA-AES256-GCM-SHA384)
Minimum TLS Version :	Choose Minimum TLS Version	(Default: TLS v1.2)
Maximum TLS Version :	Choose Maximum TLS Version	(Default: TLS v1.3)
Timeout :	Enter Timeout	(Default: 300)
Enable Session Cache :	False <input checked="" type="button" value="Default"/> True	(Default: False)
Disable Ticket :	False <input checked="" type="button" value="Default"/> True	(Default: True)
Cipher Server Preference :	False <input checked="" type="button" value="Default"/> True	(Default: True)
Disable Renegotiation :	False <input checked="" type="button" value="Default"/> True	(Default: True)
Diffie-Hellman Parameters File :	Choose File	
Prioritize ChaCha20-Poly1305 :	False <input checked="" type="button" value="Default"/> True	(Default: False)
Enable Verification :	False <input checked="" type="button" value="Default"/> True	(Default: False)
Validate Save Cancel		

Figure 17. Configuring Client-facing TLS in the Web User Interface, TLS options

TLS

Name : *	Enter Name	
Type : *	Client-facing TLS	
Certificate ^		
Certificate File : *	Choose File	
Key File : *	Choose File	
Options v		
Enable Verification :	False <input checked="" type="button" value="Default"/> True	(Default: False)
Client Verification ^		
Trusted Certs :	Choose File	
Required :	False <input checked="" type="button" value="Default"/> True	(Default: True)
Trust Level :	Choose Trust Level	(Default: Full)
Verify Depth :	Enter Depth	(Default: 4)
CA Bundle :	Choose File	
Verify CRL :	False <input checked="" type="button" value="Default"/> True	(Default: False)
Intermediate Revocation Check Type :	Choose Intermediate Revocation Check ...	(Default: Hard Fail)
Leaf Revocation Check Type :	Choose Leaf Revocation Check Type	(Default: Hard Fail)
Validate Save Cancel		

Figure 18. Configuring Client-facing TLS in the Web User Interface, Certificate options

1. Name the Client-facing TLS configuration.

2. Select the *Type* of the TLS, *Client-facing TLS* in this case, from the drop-down list to configure TLS.

For details on these parameters, see the following table:

Table 34. TLS configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name of the parameter can be referenced.
Type*	Choose the required value from the drop-down list.		Client-facing TLS, Backend-facing TLS, Syslog TLS and Elastic TLS configurations can be defined here.

3. Configure the mandatory parameters for *Client-facing TLS*, based on the information provided in Table [Client-facing TLS configuration](#).

Table 35. Client-facing TLS configuration

Key	Values	Default value	Description
Certificate			Configuration for the X.509 certificate used for TLS connections on the listener.
Certificate File*	Reference to a <i>File Brick</i> of type <i>Server Certificate</i> .		The certificate to be presented to clients.
Key File*	Reference to a <i>File Brick</i> of type <i>TLS Key</i> .		The private key corresponding to the certificate file.
Options			TLS protocol options used on the listener.
Ciphers		ECDHE-ECDSA-AES128-GCM-SHA256: ECDHE-RSA-AES128-GCM-SHA256: ECDHE-ECDSA-AES256-GCM-SHA384: ECDHE-RSA-AES256-GCM-SHA384: ECDHE-ECDSA-CHACHA20-POLY1305: ECDHE-RSA-CHACHA20-POLY1305: DHE-RSA-AES128-GCM-SHA256: DHE-RSA-AES256-GCM-SHA384	Specifies the allowed ciphers. Can be set to all, high, medium, low, or a string representation of the selected ciphers.

Key	Values	Default value	Description
Minimum TLS Version	Select one of the following options in the drop-down menu: • TLS v1.0 • TLS v1.1 • TLS v1.2 • TLS v1.3	TLS v1.2	The minimum version of TLS. Minimum TLS version must be less than or equal to the maximum TLS version.
Maximum TLS Version	Select one of the following options in the drop-down menu: • TLS v1.0 • TLS v1.1 • TLS v1.2 • TLS v1.3	TLS v1.3	The maximum version of TLS. Maximum TLS version must be greater than or equal to the minimum TLS version.
Timeout		300	It drops idle connection if the timeout value (in seconds) expires.
Enable Session Cache	True or False.	False	Store session information in the session cache. Set this option to 'On' to enable TLS session reuse.
Session Cache Size		20480	The number of sessions stored in the session cache for TLS session reuse.
Disable Ticket	True or False.	False	Session tickets are a method for TLS session reuse, described in RFC 5077. Set this option to 'On' to disable TLS session reuse using session tickets.
Cipher Server Preference	True or False.	True	Use server and not client preference order when determining which cipher suite, signature algorithm or elliptic curve to use for an incoming connection.
Disable Renegotiation	True or False.	True	Set this parameter <i>On</i> to disable client-initiated renegotiation.
Diffie-Hellman Parameters File	Reference to a <i>File Brick</i> of type <i>Diffie-Hellman Parameters</i> .		Contains the Diffie-Hellman parameters to be used by the TLS connection.
Prioritize ChaCha20-Poly1305	True or False.	False	Set this parameter <i>On</i> to prioritize using the ChaCha20-Poly1305 encryption.
Enable Verification	True or False.	False	It is an option for verifying client side X.509 certificates. By default no client verification takes place.
Client Verification			Client verification options

Key	Values	Default value	Description
Trusted Certs	Reference to a <i>File Brick</i> of type <i>Certificates</i> .		A directory where trusted IP addresses-certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
Required	True or False.	True	If it is set to True, PAS requires a certificate from the peer.
Trust Level	The values can be selected from the drop-down list. The available values are: • none • untrusted • full	full	It defines the trust level for certificate verification: • none: Accept even invalid certificates, for example self-signed certificates. • untrusted: Both trusted and untrusted certificates are accepted. • full: Only valid certificates signed by a trusted CA are accepted.
Verify Depth		4	The length of the longest accepted CA verification chain. PAS will automatically reject longer CA chains.
CA Bundle	Reference to a <i>File Brick</i> of type <i>CA Bundle</i> .		A bundle of trusted CA certificates and CRL files. CA certificates are loaded on-demand when PAS verifies the certificate of the peer.
Verify CRL	True or False.	False	If it is set to True, PAS checks the CRLs (Certificate Revocation Lists) associated with trusted CAs. CRLs will load automatically if PAS verifies the certificate of the peer.
Intermediate Revocation Check Type	The values can be selected from the drop-down list. The available values are: • none • soft_fail • hard_fail	hard_fail	The revocation check type for all certificates in the chain, except the Leaf Certificate: • none: Ignore the result certificate revocation status check • soft_fail: It fails if the check is successful and the certificate is revoked, it will pass otherwise • hard_fail: It passes only if the check is successful and the certificate is not revoked
Leaf Revocation Check Type	The values can be selected from the drop-down list. The available values are: • none • soft_fail • hard_fail	hard_fail	The revocation check types for the Leaf certificate are as follows: • none: With this option the result of the certificate revocation status check is ignored • soft_fail: It fails if the check is successful and the certificate is revoked, it passes otherwise • hard_fail: It passes only if the check is successful and the certificate is not revoked


- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the

component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.

5. Save the component configuration by clicking the *Save* button.

6.4.5.1.2. Configuring Backend-facing TLS Bricks

The following parameters need to be configured for *Backend-facing TLS*:




The screenshot shows the 'TLS' configuration page in the web user interface. The 'Type' is set to 'Backend-facing TLS'. The 'Enable Certificate' checkbox is checked. Under the 'Options' section, the following parameters are visible:

- Ciphers:** Enter Ciphers (Default: ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384)
- Minimum TLS Version:** Choose Minimum TLS Version (Default: TLS v1.2)
- Maximum TLS Version:** Choose Maximum TLS Version (Default: TLS v1.3)
- Timeout:** Enter Timeout (Default: 300)
- Enable Session Cache:** False (Default: False)
- Disable Ticket:** False (Default: True)
- Enable Verification:** False (Default: False)

Buttons at the bottom: Validate, Save, Cancel.

Figure 19. Configuring Backend-facing TLS in the Web User Interface, TLS options



The screenshot shows the 'TLS' configuration page in the web user interface, specifically the 'Certificate' section. The 'Type' is set to 'Backend-facing TLS'. The 'Enable Certificate' checkbox is checked. Under the 'Certificate' section, the following parameters are visible:

- Certificate File:** Choose File
- Key File:** Choose File

Under the 'Options' section, the following parameters are visible:

- Enable Verification:** False (Default: False)

Under the 'Backend Verification' section, the following parameters are visible:

- Trusted Certs:** Choose File
- Trust Level:** Choose Trust Level (Default: Full)
- Verify Depth:** Enter Depth (Default: 4)
- CA Bundle:** Choose File
- Verify CRL:** False (Default: False)
- Check Subject:** False (Default: False)
- Intermediate Revocation Check Type:** Choose Intermediate Revocation Check ... (Default: Soft Fail)
- Leaf Revocation Check Type:** Choose Leaf Revocation Check Type (Default: Soft Fail)

Buttons at the bottom: Validate, Save, Cancel.

Figure 20. Configuring Backend-facing TLS in the Web User Interface, Certificate options

1. Name the *Backend-facing TLS* configuration.
2. Select *Backend-facing TLS* from the drop-down list to configure *Backend-facing TLS*.

For details on these parameters, see the following table:

Table 36. TLS configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name of the parameter can be referenced.
Type*	Choose the required value from the drop-down list.		Client-facing TLS, Backend-facing TLS, Syslog TLS and Elastic TLS configurations can be defined here.

3. Configure the mandatory parameters for *Backend-facing TLS*, based on the information provided in Table [Backend-facing TLS configuration](#).

The configuration parameters are described in details in the following table:

Table 37. Backend-facing TLS configuration

Key	Values	Default value	Description
Enable Certificate	True or False.	False	It is an option for enabling backend side X.509 certificates. By default no backend verification takes place.
Certificate			Configuration for the X.509 certificate used for TLS connections on the listener.
Certificate File*	Reference to a <i>File Brick</i> of type <i>Client Certificate</i> .		The certificate to be presented to the backend.
Key File*	Reference to a <i>File Brick</i> of type <i>TLS Key</i> .		The private key corresponding to the certificate file.
Options			TLS protocol options used on the listener.

Key	Values	Default value	Description
Ciphers		ECDHE-ECDSA-AES128-GCM-SHA256: ECDHE-RSA-AES128-GCM-SHA256: ECDHE-ECDSA-AES256-GCM-SHA384: ECDHE-RSA-AES256-GCM-SHA384: ECDHE-ECDSA-CHACHA20-POLY1305: ECDHE-RSA-CHACHA20-POLY1305: DHE-RSA-AES128-GCM-SHA256: DHE-RSA-AES256-GCM-SHA384	Specifies the allowed ciphers. Can be set to all, high, medium, low, or a string representation of the selected ciphers.
Minimum TLS Version	Select one of the following options in the drop-down menu: • TLS v1.0 • TLS v1.1 • TLS v1.2 • TLS v1.3	TLS v1.2	The minimum version of TLS. Minimum TLS version must be less than or equal to the maximum TLS version.
Maximum TLS Version	Select one of the following options in the drop-down menu: • TLS v1.0 • TLS v1.1 • TLS v1.2 • TLS v1.3	TLS v1.3	The maximum version of TLS. Maximum TLS version must be greater than or equal to the minimum TLS version.
Timeout		300	It drops idle connection if the timeout value (in seconds) expires.
Enable Session Cache	True or False.	False	Store session information in the session cache. Set this option to 'On' to enable TLS session reuse.
Session Cache Size		20480	The number of sessions stored in the session cache for TLS session reuse.
Disable Ticket	True or False.	False	Session tickets are a method for TLS session reuse, described in RFC 5077. Set this option to 'On' to disable TLS session reuse using session tickets.

Key	Values	Default value	Description
Enable Verification	True or False.	False	It is an option for verifying Backend side X.509 certificates. By default no backend verification takes place.
Backend verification			Backend verification options
Trusted Certs	Reference to a <i>File Brick</i> of type <i>Certificates</i> .		A directory where trusted IP addresses-certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
Trust Level	The values can be selected from the drop-down list. The available values are: • none • untrusted • full	full	The trust level for certificate verification: • none: Accept even invalid certificates, for example self-signed certificates. • untrusted: Both trusted and untrusted certificates are accepted. • full: Only valid certificates signed by a trusted CA are accepted.
Verify Depth		4	It defines the length of the longest accepted CA verification chain. PAS will automatically reject longer CA chains.
CA Bundle	Reference to a <i>File Brick</i> of type <i>CA Bundle</i> .		A bundle of trusted CA certificates and CRL files. CA certificates are loaded on-demand when PAS verifies the certificate of the peer.
Verify CRL	True or False.	False	If it is set to True PAS checks the CRLs (Certificate Revocation Lists) associated with trusted CAs. CRLs will load automatically if PAS verifies the certificate of the peer.
Check Subject	True or False.	False	If it is set to, PAS compares the subject of the server-side certificate with application-layer information (for example, it checks whether the Subject matches the hostname in the URL).
Intermediate Revocation Check Type	The values can be selected from the drop-down list. The available values are: • none • soft_fail • hard_fail	soft_fail	The revocation check types for all certificates in the chain, except for the Leaf Certificate are as follows: • none: If this options is set, the certificate revocation status check results are ignored • soft_fail: If this option is set, the certificate revocation check fails, if the check is successful and the certificate is revoked. The check passes otherwise. • hard_fail: If this option is set, the check passes only if the check is successful, and the certificate is not revoked.

Key	Values	Default value	Description
Leaf Revocation Check Type	The values can be selected from the drop-down list. The available values are: • none • soft_fail • hard_fail	soft_fail	The revocation check type for the Leaf Certificate. • none: The result of the Certificate Revocation Status Check is ignored. • soft_fail: If this option is set, the certificate revocation check fails, if the check is successful and the certificate is revoked. The check passes otherwise. • hard_fail: If this option is set, the check passes only if the check is successful, and the certificate is not revoked.

4. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.

5. Save the component configuration by clicking the *Save* button.

6.4.5.1.3. Revocation checks for certificates

PAS tries to complete both CRL and OCSP-stapling checks for certificates.

The result for a certificate, according to the revocation check types is as follows:

Table 38. Certificate revocation checks

CRL check	OCSP stapling check	Soft fail result	Hard fail result
PASS	PASS	PASS	PASS
PASS	unsuccessful	PASS	PASS
unsuccessful	PASS	PASS	PASS
unsuccessful	unsuccessful	PASS	FAIL
PASS	FAIL	FAIL	FAIL
FAIL	PASS	FAIL	FAIL
unsuccessful	FAIL	FAIL	FAIL
FAIL	unsuccessful	FAIL	FAIL
FAIL	FAIL	FAIL	FAIL

6.4.5.1.4. Configuring Syslog TLS Bricks

The following parameters need to be configured for *Syslog TLS*:

TLS

Name : *
Type : *
Enable Client TLS Authentication :
Client TLS Authentication ^
Certificate File : *
Key File : *
Options ^
Ciphers :
Disable TLS v1.2 :
Disable TLS v1.3 :
ECDH Curve List :
Peer Verify :
Diffie-Hellman Parameters File :
Enable Verification :

Syslog TLS

False Default True

 (Default: False)

Choose File

Choose File

Enter Ciphers

 (Default: ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384)

False Default True

 (Default: False)

False Default True

 (Default: False)

Enter List

 +

Choose Peer Verify

 (Default: required-trusted)

Choose File

False Default True

 (Default: False)

Validate Save Cancel

Figure 21. Configuring Syslog TLS in the Web User Interface

1. Name the Syslog TLS configuration.
2. Select the *Type* of the TLS, *Syslog TLS* in this case, from the drop-down list to configure TLS.

For details on these parameters, see the following table:

Table 39. TLS configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name of the parameter can be referenced.
Type*	Choose the required value from the drop-down list.		Client-facing TLS, Backend-facing TLS, Syslog TLS and Elastic TLS configurations can be defined here.

3. Configure the mandatory parameters for *Syslog TLS*, based on the information provided in Table [Syslog TLS configuration](#).

Table 40. Syslog TLS configuration


Key	Values	Default value	Description
Enable Client TLS Authentication	True or False.	False	Option for enabling TLS authentication towards the server.
Client TLS Authentication			Configuration for the X.509 certificate used for TLS connections on the <i>Insight Target</i> .
Certificate File*	Reference to a <i>File Brick</i> of type <i>Client Certificate</i> .		The certificate to be used for the connection.

Key	Values	Default value	Description
Key File*	Reference to a <i>File Brick</i> of type <i>TLS Key</i> .		The private key corresponding to the certificate file. The private key file must not be encrypted.
Options			TLS protocol options used on the <i>Syslog Insight target</i> .
Ciphers	Colon-separated list of ciphers from the list supported by OpenSSL 3.0.2.	ECDH+AESGCM: DH+AESGCM:EC DH+AES256: DH+AES256:ECD H+AES128: DH+AES:!aNULL :!MD5: !DSS!aNULL: !MD5: !DSS	Specifies the allowed ciphers.
Disable TLS v1.2	True or False.	False	Disables the usage of TLSv1.2 in the connection.
Disable TLS v1.3	True or False.	False	Disables the usage of TLSv1.3 in the connection.
ECDH Curve List	Add the names of one or more ECDH curves. The possible values are the ones supported by OpenSSL 3.0.2.	empty list	A list of curves permitted in the connection when using Elliptic Curve Cryptography (ECC).
Peer Verify	Select one of the following options in the drop-down menu: optional-trusted, optional-untrusted, required-trusted, required-untrusted	required-trusted	Defines the verification method of the peer. The four possible values are a combination of two properties of validation: whether the peer is required to provide a certificate (required or optional prefix), and whether the certificate provided needs to be valid (trusted or untrusted suffix).
Diffie-Hellman Parameters File	Reference to a <i>File Brick</i> of type <i>Diffie-Hellman Parameters</i> .		Contains the Diffie-Hellman parameters to be used by the TLS connection.
Enable Verification	True or False.	False	Option for enabling the verification of server side X.509 certificates.
Server Verification*			Server verification options are mandatory if <i>Enable Verification</i> is set to <i>True</i> .
CA Bundle	Reference to a <i>File Brick</i> of type <i>CA Bundle</i> .		A bundle of trusted CA certificates and CRL files. CA certificates are loaded on-demand when PAS verifies the certificate of the peer.
Verify CRL	True or False.	False	Verifies that certificates used in the connection are not revoked by any CRLs in the CA directory.

- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

6.4.5.1.5. Configuring Elastic TLS Bricks

The following parameters need to be configured for *Elastic TLS*:



TLS

Name: *

Type: *

Enable Client TLS Authentication: ☐ False ☒ Default ☐ True (Default: False)

Client TLS Authentication ^

Certificate File: *

Key File: *

Options ^

Ciphers: (Default: ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-POLY1305:DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384)

TLS Version: (Default: TLS v1.2)

Peer Verify: ☐ False ☒ Default ☐ True (Default: True)

Enable Verification: ☐ False ☒ Default ☐ True (Default: False)

Figure 22. Configuring Elastic TLS in the Web User Interface

1. Name the Elastic TLS configuration.
2. Select the *Type* of the *TLS brick*, *Elastic TLS* in this case, from the drop-down list to configure the encryption used with the Elastic server.

For details on these parameters, see the following table:

Table 41. TLS configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name of the parameter can be referenced.
Type*	Choose the required value from the drop-down list.		Client-facing TLS, Backend-facing TLS, Syslog TLS and Elastic TLS configurations can be defined here.

3. Configure the mandatory parameters for *Elastic TLS*, based on the information provided in Table [Elastic TLS configuration](#).

Table 42. Elastic TLS configuration

Key	Values	Default value	Description
Enable Client TLS Authentication	True or False.	False	Option for enabling TLS authentication towards the server.
Client TLS Authentication			Configuration for the X.509 certificate used for TLS connections on the <i>Insight Target</i> .
Certificate File*	Reference to a <i>File Brick</i> of type <i>Client Certificate</i> .		The certificate to be used for the connection.

Key	Values	Default value	Description
Key File*	Reference to a <i>File Brick</i> of type <i>TLS Key</i> .		The private key corresponding to the certificate file. The private key file must not be encrypted.
Options			TLS protocol options used on the <i>Elastic Insight target</i> .
Ciphers	Colon-separated list of ciphers from the list supported by OpenSSL 3.0.2.	ECDH+AESGCM: DH+AESGCM:EC DH+AES256: DH+AES256:ECD H+AES128: DH+AES:!aNULL :!MD5: !DSS!aNULL: !MD5: !DSS	Specifies the allowed ciphers.
TLS Version	Select one of the following options in the drop-down menu: • TLS v1.0 • TLS v1.1 • TLS v1.2 • TLS v1.3	TLS v1.2	Defines the TLS version used in the connection.
Peer Verify	True or False.	True	Defines whether the peer is verified. If set to true, the peer is required to provide a certificate, and the certificate provided needs to be valid.
Enable Verification	True or False.	False	Option for enabling the verification of the X.509 certificate presented by the Elastic server.
CA Bundle	Reference to a <i>File Brick</i> of type <i>CA Bundle</i> .		A bundle of trusted CA certificates and CRL files. CA certificates are loaded on-demand when PAS verifies the certificate of the peer.

4. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.

5. Save the component configuration by clicking the *Save* button.

6.4.6. File

The *File* configuration element enables the administrator to upload files used by various plugins.

6.4.6.1. Configuring File Bricks

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *BRICKS*.
2. Select *File*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already

configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 23. File main page in the Web User Interface

- Click on the *New* navigation button to create a File Brick.
- Choose the type of the file brick from the drop-down list.
- Upload a file according to the selected type.
- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The following values can be configured for the *File Brick*:




Figure 24. Configuring File in the Web User Interface

Table 43. File configuration parameters

Key	Values	Default	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		It defines the file-related configuration.
Type*	The available values are: • Swagger • OpenAPI 3.0 • OpenAPI 3.1 • OpenAPI 3.2 • XSD • WSDL • CA Bundle • Certificates • Diffie-Hellman Parameters • TLS Key • Client Certificate • Server Certificate • License See table Requirements for specific file types for specific requirements for each type.		The type selected here defines by which <i>PLUGINS</i> it can be used. The file uploaded here with the <i>Type Swagger</i> , for example, can be used by <i>Swagger Plugins</i> .
File*	The required file can be uploaded here.		
Passphrase	String value, could be empty.		<i>Only available for TLS Key files.</i> The passphrase to access an encrypted private key. Leave empty if the private key is unencrypted.

Table 44. Requirements for specific file types

File type	Requirements
CA Bundle	1. The file must be a flat ZIP file with the CA certificates inside. 2. It can contain copies of the certificates named following the <code><hash>.0</code> format. The value of the <code><hash></code> part can be produced with the following command: <code>openssl x509 -noout -hash -in /path/to/cert/file</code>. These copies will be generated automatically after saving if they are not present already, and the original File brick will be overwritten. 3. It can contain CRL files, and it also can contain the copies of them following the <code><hash_of_the_related_ca_file>.r0</code> format. The hash can be produced as described above. These copies will be generated automatically after saving if they are not present already, and the original File brick will be overwritten.

File type	Requirements
Certificates	1. The file must be a flat ZIP file with the certificates inside. 2. The certificates must be named after IPv4 or IPv6 addresses.
Diffie-Hellman Parameters	1. Must be in PEM format. 2. Must be a parameters file, such as one generated by the <code>openssl dhparam</code> utility.
TLS Key	1. Must be in PEM format. 2. Must be a private key file. 3. Could be encrypted or unencrypted. If the file is encrypted, the passphrase must be provided in the Passphrase field.
Client Certificate	1. Must be in PEM format. 2. Must be a certificate file. 3. Must have a Common Name attribute, and have the CLIENT_AUTH ExtendedKeyUsage.
Server Certificate	1. Must be in PEM format. 2. Must be a certificate file. 3. Must have a Common Name attribute, and have the SERVER_AUTH ExtendedKeyUsage.
License	1. Must be a ZIP file with a single <code>pas</code> directory, with a single <code>license.txt</code> inside. 2. Must be a valid PAS license. 3. Expirations are not validated. Limits are validated in integrity checks based on the <i>License File</i> brick selected in the <i>License</i> system.
Swagger	The file must be a Swagger schema as described in the OpenAPI 2.0 specification .
OpenAPI 3.0	The file must be an OpenAPI 3.0 schema as described in the OpenAPI 3.0 specification .
OpenAPI 3.1	The file must be an OpenAPI 3.1 schema as described in the OpenAPI 3.1 specification .
OpenAPI 3.2	The file must be an OpenAPI 3.2 schema as described in the OpenAPI 3.2 specification .
XSD	The file must be an XML Schema Definition as described in XML Schema Part 1: Structures , XML Schema Part 2: Datatypes , XSD 1.1 Part 1: Structures and XSD 1.1 Part 2: Datatypes .
WSDL	The file must be a WSDL service descriptor as described in the Web Services Description Language 1.1 specification or in the Web Services Description Language 2.0 specification .

File editor

Files in certain *File* brick types are editable when configuring the *File* brick. A *File editor* is available for the following types:

- Swagger
- OpenAPI 3.0
- OpenAPI 3.1

The uploaded file can be opened and edited by clicking the *Edit* button. The contents of the file open inside a new window, with the *Edit* tab selected:

File Editor

Edit Overview

```

1 openapi: "3.0.0"
2 info:
3   version: 1.0.0
4   title: Swagger Petstore
5   license:
6     name: MIT
7 servers:
8   - url: http://petstore.swagger.io/v1
9 paths:
10  /pets:
11    get:
12      summary: List all pets
13      operationId: listPets
14      tags:
15        - pets
16      parameters:
17        - name: limit
18          in: query
19          description: How many items to return at one time (max 100)
20          required: false
21          schema:
22            type: integer
23            maximum: 100
24            format: int32
25      responses:
26        200:
27          description: A paged array of pets
28          headers:
29            x-next:
30              description: A link to the next page of responses
31              schema:
32                type: string
33          content:
34            application/json:
35              schema:
36                $ref: "#/components/schemas/Pets"

```

Save Close

The editor can be closed without saving any changes to the file with the *Close* button. The changes are saved and the editor is closed with the *Save* button.

The *Overview* tab shows errors if there are any, and the structure of the schema that the file describes:

File Editor

Edit Overview

pets ^

- GET /pets List all pets
- POST /pets Create a pet
- GET /pets/{petId} Info for a specific pet

Schemas ^

- Pet >
- Pets >
- Error >

Save Close

6.4.7. Authentication

Authentications contain settings for ways to authenticate incoming calls.



The *Authentication* brick only defines the authentication method used for the call, but not the scope. For example, if two *Endpoints* share an *Authentication* brick, a user who authenticates on one *Endpoint* will still have to authenticate separately on the other *Endpoint*.

Authentication is based on OpenID Connect (OIDC), an authentication protocol based on OAuth 2.0. It requires an external OpenID Identity Provider. See [the OpenID Foundation website](#) for details. See the [Authentication configuration options](#) for the available *Authentication* configuration options.

6.4.7.1. Configuring Authentications

1. Click on the *BRICKS* main navigation item in the Left navigation area. Alternatively you can also click on the

 sign to open up the sub-navigation items of *BRICKS*.

2. Select *Authentication*.

In the configuration window, you will either see empty parameter fields for the component you are configuring or a list of components with preconfigured parameters. These components may be system defaults or may have been configured by an administrator:




Figure 25. Authentication main page in the Web User Interface

3. Click on the *New* navigation button to create an Authentication.
4. Provide the name for your *Authentication* configuration.
5. Proceed with configuring OpenID Connect parameters. The list of available configuration options is described the following table: [OpenID Connect Authentication configuration parameters](#).
6. Set up any optional parameters if needed.
7. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
8. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Authentication Brick*:

Authentication

Name : *
 Client ID : *
 Client Secret : *
 Discovery URL : *
 Callback URL : *
 Use TLS :
 Username Claim :
 Discovery Timeout :
 Token Timeout :
 Cache Expiration Time :

Enter Name
 Enter ID
 Enter Secret
 Enter URL
 Enter URL
 False Default True (Default: False)
 Enter Claim (Default: email)
 Enter Timeout (Default: 10)
 Enter Timeout (Default: 10)
 Enter Time in seconds (Default: 3600)

Validate Save Cancel

Figure 26. Configuring Authentication in the Web User Interface

Table 45. Authentication configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the <i>Authentication</i> . This name of the <i>Authentication</i> can be referenced from other parts of the configuration.

The following table describes the configuration parameters for the OpenID Connect *Authentication* type:

Table 46. OpenID Connect Authentication configuration parameters

Key	Values	Default value	Description
Client ID*	String.		The ID used to register PAS at the identity provider.
Client Secret*	String.		A secret provided by the identity provider.
Discovery URL*	URL.		The identity provider's configuration endpoint, usually ending in <code>"/.well-known/openid-configuration"</code> .
Callback URL*	URL.		The Redirection Endpoint the identity provider should redirect to after successful authentication. NOTE: as this call contains authentication tokens, it is highly recommended to use an URL with HTTPS for this.
Use TLS	True or False.	False	Enables using TLS encryption when connecting to the identity provider.
CA Bundle*	Reference to a <i>File Brick</i> of type <i>CA Bundle</i> .		The CA bundle to use in validating TLS connections towards the provider servers. Mandatory if <i>Use TLS</i> is set to <i>True</i> .

Key	Values	Default value	Description
Username Claim	String.	email	The name of the claim in the ID token that contains the username used for authorization. In OpenID Connect, a claim is a piece of information asserted about a user, such as their email address, name, or unique identifier.
Discovery Timeout	Integer.	10	The maximum time, in seconds, to wait for a response from the configuration endpoint.
Token Timeout	Integer.	10	The maximum time, in seconds, to wait for a response from the token endpoint.
Cache Expiration Time	Integer.	3600	The duration, in seconds, for which authentication information is cached.

6.4.8. Common configuration elements for BRICKS

6.4.8.1. Extractors

Extractors are used to extract data from the call.

Extractors are not independent configuration components, but common configuration elements that are utilized by [Matchers](#) and [Selectors](#). In fact, when configuring matchers and selectors, it is extractors that are listed at their type fields. Extractors are configured and used as part of matchers and selectors. There are no named extractors.

Most extractors return simple string values. However, some (might) return dictionaries. For example, you can get all the HTTP headers, or all the URI query parameters.

See the [Extractor types](#) for more details on extractors and their configuration options.

The following table provides details on extractor types:

Table 47. Extractor types

Key	Description
Method	Extracts the HTTP method of the request. It does not require configuration.
Status	Extracts the status code of the response. It does not require configuration.
JMESPath	Extracts data from the body of a JSON call with the help of a JMESPath expression. JMESPath is a query language for JSON. It is a very versatile tool for extracting the needed information from the body of the call, and organizing it according to requirements. A complete explanation on how to write JMESPath expressions is not in the scope of this document. To learn more about it visit the JMESPath website: • There is a tutorial. • There are examples. • There is also a formal specification.
Header	Extracts the value of an HTTP header. It is valid for some HTTP headers to be present more than once in a call. In this case, all the values are extracted as a list. It provides the name of the header in the configuration.
Header Force List	A <i>Header</i> extractor that returns a list even if there is only a single extracted value.

Key	Description
Header First	A <i>Header</i> extractor that only returns the first extracted value even if there is a list of extracted values.
Headers	The <i>Headers</i> extractor returns all the headers from the call. The results are stored as a dictionary, therefore it is recommended to set 'Save Under Key' to False if you use this from a Selector. It is valid for some HTTP headers to be present more than once in a call. In such cases all the values are stored under the header's key as a list. It does not require configuration.
Fraud Detector Score	Extracts the score value provided by the <i>Fraud Detector</i> plugin.
URI	Extracts the whole request URI as received from the client. It does not require configuration.
URI Netloc	Extracts the <i>network location</i> in the URI. It does not require configuration. It includes: • username and password if present • host • port if present unless scheme default  If the port is the default port for the scheme - that is 80 and 443 for HTTP and HTTPS, respectively - the port will not be included even if explicitly sent by the client. Therefore if the client used http://example.com:80/path then the <i>netloc</i> would be http://example.com, not http://example.com:80.
URI Origin	Extracts the <i>origin</i> part of the URI. It does not require configuration. It includes: • scheme • host • port if present, unless the default port for the scheme is used  If the port is the default port for the scheme - that is 80 and 443 for HTTP and HTTPS, respectively - the port will not be included, even if explicitly sent by the client. Therefore if the client used http://example.com:80/path, then the <i>origin</i> would be http://example.com, not http://example.com:80.
URI Scheme	Extracts the <i>scheme</i> of the request (http or https). It does not require configuration.
URI Username	Extracts the <i>username</i> in the request if present. It does not require configuration.
URI Password	Extracts the <i>password</i> in the request if present. It does not require configuration.
URI Host	Extracts the host in the request. It does not require configuration.
URI Port	Extracts the port of the request, the default port — that is 80 and 443 for HTTP and HTTPS, respectively — even if it is not displayed explicitly in the request. It does not require configuration.

Key	Description
URI Path	Extracts the <i>path</i> part of the URI. It does not require configuration. The path is normalized to allow more robust matching and cleaner reporting. This means that: • If the path is missing <i>/</i> it is extracted. • Repeating forward-slash (<i>/</i>) characters are replaced with a single one. • dot (<i>.</i>) and double-dot (<i>..</i>) path segments are resolved. Consequently, if the path present in the <i>URI</i> was <i>//some/./nothere/./resource///./somewhere</i> the <i>path</i> would be <i>/some/resource/somewhere</i>. If you need to extract the <i>path</i> exactly as received, use URI raw path parameter.
URI Raw Path	Extracts the path part of the URI, without the normalization of URI path carried out. NOTE: If the <i>path</i> is missing a single forward slash (<i>"/</i>) is extracted. It does not require configuration.
URI Raw Query	Extracts the query part of the URI as a string. It does not require configuration.
URI Query	Extracts the query part of the URI. The results are stored as a dictionary, therefore it is recommended to set 'Save Under Key' to False if you use this from a Selector.
URI Query Param	Extracts the value of a query parameter. It is also valid for URIs to include a query parameter more than once. That is, it could be 'foo=bar&qux=quz&foo=baz'. In this case both values are extracted as a list. Provide the name of the parameter in the configuration.
URI Query Param Force List	An <i>URI Query Param</i> extractor that returns a list even if there is only a single extracted value.
URI Query Param First	An <i>URI Query Param</i> extractor that only returns the first extracted value even if there is a list of extracted values.
Client Address	Extracts the client's IP address.
Client Port	Extracts the client's port (TCP).
Server Address	Extracts the server's IP address.
Server Port	Extracts the server's port (TCP).
Parsed Content	Extracts the content. It does not require configuration.
Raw Content	Extracts the raw bytes of the request or response. It saves the results as a base64 encoded string.
Text Content	Extracts the request's or response's content as a decoded string.
Cookie	Extracts the values for a given key from the Cookie HTTP header. It is valid for multiple key-value pairs to be present in a Cookie header for the same key. In this case, all the values are extracted as a list. It requires the name of the Cookie key in the configuration.
Cookie Force List	A <i>Cookie</i> extractor that returns a list even if there is only a single extracted value.
Cookie First	A <i>Cookie</i> extractor that only returns the first extracted value even if there is a list of extracted values.

Key	Description
Cookies	The <i>Cookies</i> extractor returns all the key-value pairs from the Cookie header. The results are stored as a dictionary, therefore it is recommended to set 'Save Under Key' to False if you use this from a Selector. It is valid for multiple key-value pairs to be present in a Cookie header for the same key. In such cases, all the values are stored under the Cookie's key as a list. It does not require configuration.
Content Type	Extracts the content type from the HTTP header. It does not require configuration.
Content Type Charset	Extracts the charset from the content type HTTP header. It does not require configuration.
Call Direction	Extracts the call direction (request, response). It does not require configuration.
Session Id	Extracts the internal identifier of the HTTP session in keep-alive HTTP connections. Its 'Include request counter' option enables adding a request counter representing the number of requests in the session. See [session-id] for details.
Backend Response Time	Extracts the time spent between the sending the request towards the server and receiving the response from the server, in milliseconds. Only returns a value in a response flow.
Backend Name	Extracts the name of the <i>Backend Service</i> component handling the call.
Endpoint Name	Extracts the name of the <i>Endpoint Service</i> component handling the call.
Static	Extracts a string, integer, number, object, array, boolean as string from the configuration.
Timestamp	Extracts the current time. Also see the tables on Configuring timestamps and Timestamp format options .
Error Policy	Extracts <i>Request Action</i> or <i>Response Action</i> field of the <i>Error Policy</i> of the <i>Plugin</i> selected by the <i>plugin</i> field. If the <i>plugin</i> field contains the special value "Previous", the data will be extracted from the last evaluated plugin in the same HTTP request or response, if there is one. If the <i>plugin</i> field refers to a specific plugin instance, the data will be extracted from the last evaluation of the referred plugin instance. If the extractor is in the HTTP response, and the last evaluation is in the HTTP request, then the evaluation result from the HTTP request will be selected. Only plugins with a negative verdict will return data.
Error Policy Action	Works like the <i>Error Policy</i> extractor, but extracts only the <i>Request Action</i> or <i>Response Action</i> field of the referred plugin's <i>Error Policy</i> .
Error Policy Status Code	Works like the <i>Error Policy</i> extractor, but extracts only the <i>Request Code</i> or <i>Response Code</i> field of the referred plugin's <i>Error Policy</i> .
Error Policy Silent	Works like the <i>Error Policy</i> extractor, but extracts only the <i>Request Silent</i> or <i>Response Silent</i> field of the referred plugin's <i>Error Policy</i> .
Error Policy Message	Works like the <i>Error Policy</i> extractor, but extracts only the <i>Request Message</i> or <i>Response Message</i> field of the referred plugin's <i>Error Policy</i> .
Plugin Name	Works like the <i>Error Policy</i> extractor, but extracts the name of the referred plugin. Returns data regardless of verdict.
Plugin Verdict	Works like the <i>Error Policy</i> extractor, but extracts the verdict of the referred plugin. Returns data regardless of verdict.
Plugin Error Message	Works like the <i>Error Policy</i> extractor, but extracts the error message provided in the referred plugin's negative verdict, if there is one.

Key	Description
XPath	Extracts data from the body of an XML call with the help of a XPath expression. XPath is a query language for XML. It is a very versatile tool for extracting the needed information from the body of the call, and organizing it according to needs. A complete explanation on how to write XPath expressions is not in the scope of this document. To learn more about it visit the main website. Also see table XPath extractor configuration options. Provide the XPath expression in the configuration. Depending on the expression, the return value is a single node or a list of nodes. If you want a single value or a list independent from the expression, use <i>XPath First</i> or <i>XPath Force List</i>.
XPath Force List	Works like <i>XPath</i> but it returns a list even if there is only a single extracted value.
XPath First	Works like <i>XPath</i> but it only returns the first extracted value even if there is a list of extracted values.
SOAP Version	Extends the XPath extractor with predefined expressions. Extracts the SOAP message version. It identify with the SOAP namespace. Possible values: • soapv1_1 - the message version is SOAP v1.1 • soapv1_2 - the message version is SOAP v1.2
SOAP Envelope	Extends the XPath extractor with predefined expressions. Extracts the SOAP envelope.
SOAP Header	Extracts the SOAP header. Extends the XPath extractor with predefined expressions.
SOAP Body	Extracts the SOAP body. Extends the XPath extractor with predefined expressions.
SOAP Fault	Extracts the SOAP fault. Extends the XPath extractor with predefined expressions.
SOAP Fault Code	Extracts the SOAP fault 'code'. Extends the XPath extractor with predefined expressions. This extractor expression depends on the SOAP version. • faultcode - the SOAP v1.1 node tag • Code - the SOAP v1.2 node tag

Key	Description
SOAP Fault Detail	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'detail'. This matcher expression depends on the SOAP version. • Detail - the SOAP v1.1 node tag • Detail - the SOAP v1.2 node tag
SOAP 1.1 Fault Faultstring	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'faultstring'. This extractor only works with SOAP version 1.1.
SOAP 1.1 Fault Faultactor	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'faultactor'. This extractor only works with SOAP version 1.1.
SOAP 1.2 Fault Reason	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'Reason'. This extractor only works with SOAP version 1.2.
SOAP 1.2 Fault Node	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'Node'. This extractor only works with SOAP version 1.2.
SOAP 1.2 Fault Role	Extends the XPath extractor with predefined expressions. Extracts the SOAP fault 'Role'. This extractor only works with SOAP version 1.2.



You can still use **Save As Key** for extractors returning dictionaries. For example, you can save all the headers under the [headers](#) key and the URI query parameters under the [parameters](#) key.

Timestamp extractors can be configured as follows:

Table 48. Configuring timestamps

Name	Default	Description
Time Zone	'UTC'	Set the time zone. • An <i>str</i> describing a time zone, similar to 'US/Pacific', or 'Europe/Berlin'. See: Time zones • An <i>str</i> in ISO 8601 style, as in '+07:00'. • An <i>str</i>, one of the following: 'local', 'utc', 'UTC'.
Time Format	YYYY-MM-DDT HH:mm:ss.SSSSSZZ (line breaks for display purposes only)	Set the format. See: Timestamp format options
Use Request Start Time	True	If set to True, uses the request's start time. This value is set once for each call. If set to False, uses the time when the selector is processed during a session. This value can change every time the selector's value is queried during a call.

Table 49. Timestamp format options

Name	Token	Output
Year	YYYY YY	2000, 2001, 2002 ... 2012, 2013 00, 01, 02 ... 12, 13
Month	MMMM MMM MM M	January, February, March Jan, Feb, Mar 01, 02, 03 ... 11, 12 1, 2, 3 ... 11, 12
Day of Year	DDDD DDD	001, 002, 003 ... 364, 365 1, 2, 3 ... 364, 365
Day of Month	DD D Do	01, 02, 03 ... 30, 31 1, 2, 3 ... 30, 31 1st, 2nd, 3rd ... 30th, 31st
Day of Week	dddd ddd d	Monday, Tuesday, Wednesday Mon, Tue, Wed 1, 2, 3 ... 6, 7
Hour	HH H hh h	00, 01, 02 ... 23, 24 0, 1, 2 ... 23, 24 01, 02, 03 ... 11, 12 1, 2, 3 ... 11, 12
AM / PM	A a	AM, PM, am, pm am, pm
Minute	mm m	00, 01, 02 ... 58, 59 0, 1, 2 ... 58, 59
Second	ss s	00, 01, 02 ... 58, 59 0, 1, 2 ... 58, 59
Sub-second	S...	0, 02, 003, 000006, 123123123123 the result is truncated to microseconds, with half-to-even rounding
Time zone	ZZZ ZZ Z	Asia/Baku, Europe/Warsaw, GMT -07:00, -06:00 ... +06:00, +07:00, +08, Z -0700, -0600 ... +0600, +0700, +08, Z
Seconds Timestamp	X	1381685817, 1381685817.915482
ms or µs Timestamp	x	1569980330813, 1569980330813221

Table 50. XPath extractor configuration options

Key	Default	Description
xpath_expression		The expression to extract the node from the call to match against.
namespaces		Defines the XML namespaces.
clear_text	False	Whether to remove white spaces at the beginning and at the end of the string.

6.4.8.2. Comparators

Comparators are used for comparing the pattern with the result of the XPath expression.

Table 51. Types of comparators

Key	Description	Parameters
Equals	Matches if the parameter is exactly the same as the value matched. For matchers that work with numeric data type or with IP addresses it validates if the input is a valid number or IP address.	Ignorecase: Case differences (lower case, upper case) are ignored. When the present Value would match value . For matcher types that work with numeric data type or with IP addresses, the 'Equals' and 'Not Equals' comparator types do not have ignorecase field.
Not equals	Matches if the parameter is not exactly the same as the value matched. For matchers that work with numeric data type or with IP addresses it validates if the input is a valid number or IP address.	Ignorecase: Case differences are ignored. When the present Value would not match value . For matcher types that work with numeric data type or with IP addresses, the 'Equals' and 'Not Equals' comparator types do not have ignorecase field.
Starts with	Matches if the value starts exactly with the pattern.	Ignore case: Case differences are ignored. When the present Value would match value_given .
Ends with	Matches if the value ends exactly with the pattern.	Ignore case: Case differences are ignored. When the present Value would match given_value .
Contains	Matches if the exact pattern is found somewhere in the value.	Ignore case: Case differences are ignored. When the present Value would match some-value-given .
Pattern	The Pattern treats the input as Unix shell-style wildcards. There are special characters used in shell-style wildcards: • '*' Matches everything. • '?' Matches a single character. • [seq] Matches any character in seq.  For a literal match, wrap the meta-characters in brackets. For example, <code>[?]</code> matches a literal question-mark (?).	Ignore case: Case differences are ignored. When the present Value would match some-value-given .
Regex	<i>Regex</i> treats input as a regular expression for matching. Consult Python's regular expression documentation and their Regular Expression HOWTO .	• Ignore case: It sets the IGNORECASE flag for the regex. • Multiline: It sets the MULTILINE flag for the regex.
Minimum	Matches if the pattern is larger or equal to the value.	
Maximum	Matches if the pattern is smaller or equal to the value.	
Range	Matches if the value is between the limits in the pattern, including boundaries. The format of the pattern must be minimum..maximum.	

Key	Description	Parameters
Status class	Status class is a special comparator for conveniently matching HTTP status code classes in a <i>Status</i> matcher. It takes the name of the class and checks if the status code is in the given range as stated in Checking status code range .	
Subnet	The <i>subnet</i> comparator examines if an extracted IP address is in the specified subnet. The format for the input of the subnet comparator is the CIDR notation for IPv4 (for example, 192.0.2.0/24) and canonical prefix notation for IPv6 (for example, 2001:db8::/32).	
Error Policy Action	Matches the value in the <i>Request Action</i> or <i>Response Action</i> fields of an <i>Error Policy</i> brick.	
Plugin Verdict	Matches the verdict of a <i>Plugin</i> evaluation.	

Table 52. Checking status code range

Pattern	Status code range	Description
Info	1xx	Informational response
Success	2xx	Successful response
Redirect	3xx	Redirects
Client Error	4xx	Client Errors
Server Error	5xx	Server Errors

6.5. PLUGINS - Configuration units

A plugin is an element of the security flow that applies a specific security function. Plugins have different types based on the role they do:




Figure 27. The PLUGINS main page in the Web User Interface

Plugins are named, so that they can be referenced in other parts of the configuration.



This means that *Plugin* configurations are reusable.

Certain Plugins are so called *default* objects, which are in 'read-only' state and cannot be configured or modified. Such default objects are listed in the following table:

Table 53. Default objects - PLUGINS

Default object name	Key
default_json	Serializer
default_xml	Serializer
default_json	Deserializer
default_xml	Deserializer

6.5.1. Common Plugin parameters

Regardless of what plugins do, all plugins share some common parameters.

Table 54. Plugins' common parameters

Key	Values	Default value	Description
Matcher	Reference to a <i>Matcher Brick</i> .	Always: If the value is not defined, the plugin is always executed.	Optional parameter. It decides if the Plugin should be executed based on the call's details. If no matcher is configured the Plugin is always executed. For more details, see Matcher .

Key	Values	Default value	Description
Error Policy	Reference to an <i>Error Policy Brick</i> .		Optional parameter. It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. If no error policy is configured, the plugin type's default error policy is applied. For more details, see Error Policy .

Plugins are always named so that their names refer to a *Plugin* that represents a certain configuration. The names themselves are referenced from the [Security Flow](#).

6.5.2. Enforcer

An *Enforcer Plugin* validates calls against externally defined schemas.

The *Plugin* supports validation against OpenAPI (Swagger) schemas, XSD schemas, WSDL schema or WAF ruleset.

Understanding the format of these schemas is not in the scope of this document. Further information is available at:

- [The OpenAPI 2.0 format](#)
- [The OpenAPI 2.0 Specification](#)
- [The OpenAPI 3.0 format](#)
- [The OpenAPI 3.0 Specification](#)
- [The OpenAPI 3.1 Specification](#)
- [OpenAPI 3.2 specification](#)
- [XSD 1.1 Specification](#)
- [XSD Tutorial](#)
- [WSDL Tutorial](#)
- [WSDL 1.1 Specification](#)
- [WSDL 2.0 Specification](#)

The Enforcer Plugin uses its own default error policy, that is, the 'enforcer_default' error policy. The Plugin overrides the following fields of the [default error policy](#):

Table 55. Default Enforcer Error Policy

Policy Setting	Default
request_code	422
request_message	Request Error

Problems are considered errors that lead to the termination of the call. Problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.

6.5.2.1. Configuring Enforcer Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Enforcer* plugin.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 28. Enforcer Plugin's main page in the Web User Interface

3. Click on the *New* navigation button to create an Enforcer.
4. Name the *Enforcer* Plugin.
5. Choose the type of the *Enforcer* plugin.
6. Choose an *Error policy* from the drop-down list. The drop-down list will offer the error policy options configured under *BRICKS*.
7. Choose a *Matcher* from the drop-down list. The drop-down list will offer the matcher options configured under *BRICKS*.
8. Depending on the choice of the *Enforcer plugin* type selected earlier, different fields appear here for further configuration:
 - Swagger - Choose an uploaded Swagger file if the Enforcer type selected at the *Type* field was Swagger.
 - OpenAPI 3.0 - Choose an uploaded OpenAPI 3.0 file if the Enforcer type selected at the *Type* field was OpenAPI 3.0.
 - OpenAPI 3.1 - Choose an uploaded OpenAPI 3.1 file if the Enforcer type selected at the *Type* field was OpenAPI 3.1.
 - OpenAPI 3.2 - Choose an uploaded OpenAPI 3.2 file if the Enforcer type selected at the *Type* field was OpenAPI 3.2.
 - WSDL - Choose an uploaded WSDL file if the Enforcer type selected at the *Type* field was WSDL.
 - Operations - Fill in the *Operations* fields according to [XSD enforcer plugin configuration options for Operations](#) if the Enforcer type selected at the *Type* field was XSD.
 - Request Limit in Kilobytes - Fill in a number if you want to overwrite the default setting.
 - Harden Additional Properties Defaults - Choose the desired setting.
9. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
10. Save the component configuration by clicking the *Save* button.

6.5.2.2. Swagger

The Swagger enforcer *Plugin* validates against OpenAPI 2.0 schemas.

The *Plugin* needs the schema definition file of the API Endpoint. This file must be in JSON or YML format.

6.5.2.3. OpenAPI 3.0

The OpenAPI 3.0 enforcer *Plugin* validates against OpenAPI 3.0 schemas.

The *Plugin* needs the schema definition file of the API Endpoint. This file must be in JSON or YML format.

6.5.2.4. OpenAPI 3.1

The OpenAPI 3.1 enforcer *Plugin* validates against OpenAPI 3.1 schemas.

The *Plugin* needs the schema definition file of the API Endpoint. This file must be in JSON or YML format.

6.5.2.5. OpenAPI 3.2

The OpenAPI 3.2 enforcer *Plugin* validates against OpenAPI 3.2 schemas.

The *Plugin* needs the schema definition file of the API Endpoint. This file must be in JSON or YML format.

6.5.2.6. XSD

XSD enforcer *Plugin* validates against XSD schemas. Both XSD 1.0 and 1.1 are supported.



As XSD enforcer requires parsed XML content an xml deserializer plugin needs to be included before XSD enforcer.

In the XSD enforcer you can define operations. Each operation contains criteria for identifying the call, and path of an XSD schema. If the HTTP message meets all criteria, its content will be validated using the schema.

XSD enforcer schema must contain at least one operation.

6.5.2.7. WSDL

WSDL enforcer *Plugin* validates against WSDL 1.0-2.0 schemas.



As WSDL enforcer requires parsed XML content, an xml deserializer plugin needs to be included before WSDL enforcer.

The Enforcer Plugin uses its own default error policy, that is, the 'enforcer_default' error policy. The Plugin overrides the following fields of the [default error policy](#):

Table 56. Default Enforcer Error Policy

Policy Setting	Default
request_code	422
request_message	Request Error

Problems are considered errors that lead to the termination of the call. Problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.

The plugin needs the schema definition file. This file must be in XML format.



WSDL schema validates request and response as well. Make sure that wsdl enforcer included in request and response flow as well.



In simple cases — when the listener/endpoint is serving a single version of a single API endpoint — a matcher is usually not needed as the schemas define all known URLs in the API. If however multiple API endpoints are consolidated under a single endpoint definition, you can define multiple enforcers each matching on a sub-path by using an URI path matcher and putting them all in the Security Flow.

6.5.2.8. WAF

The Web Application Firewall (WAF) enforcer *Plugin* protects against known attacks. The list of known attacks is updated by the [WAF Ruleset Updater](#).

The following values can be configured for the *Enforcer Plugin*:




Figure 29. Configuring an enforcer plugin in the Web User Interface

Table 57. Enforcer Plugin's configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		This name identifies the Enforcer Plugin. The name of the plugin can be referenced from other parts of the configuration.

Key	Values	Default value	Description
Type*	Can be selected from the drop-down list. The available values are: • Swagger • OpenAPI 3.0 • OpenAPI 3.1 • OpenAPI 3.2 • XSD • WSDL • WAF		The type of the <i>Enforcer</i> plugin.
Matcher	Reference to a <i>Matcher Brick</i> .	Always: If the value is not defined the plugin is always executed.	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If omitted the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	enforcer_default	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. For details see Error Policy .
Swagger*/OpenAPI 3.0*/OpenAPI 3.1*/OpenAPI 3.2*/WSDL*/Operations*	Depending on which type of the component was selected above, the following values are available: • For Swagger, OpenAPI 3.0, OpenAPI 3.1, OpenAPI 3.2, WSDL, and XSD a reference to a <i>File Brick</i> of the appropriate type. • For XSD the configuration options for Operations can also be selected here. For details on parameters for Operations, see XSD enforcer plugin configuration options for Operations.		The Swagger enforcer <i>Plugin</i> validates against OpenAPI 2.0 schemas. The OpenAPI 3.0 enforcer <i>Plugin</i> validates against OpenAPI 3.0 schemas. The OpenAPI 3.1 enforcer <i>Plugin</i> validates against OpenAPI 3.1 schemas. The OpenAPI 3.2 enforcer <i>Plugin</i> validates against OpenAPI 3.2 schemas. WSDL enforcer <i>Plugin</i> validates against WSDL 1.0-2.0 schemas. XSD enforcer <i>Plugin</i> validates against XSD schemas.

Key	Values	Default value	Description
Harden Additional Properties Defaults	True or False.	False	Only available for OpenAPI 3.0, OpenAPI 3.1 and OpenAPI 3.2 enforcers. If set to True, the Enforcer will check calls as if the default value of <code>additionalProperties</code> would be False for <code>Schema Objects</code> , triggering the error policy if a non-specified property is present in the call, unless <code>additionalProperties=True</code> is explicitly set on the object. If set to False, the original behavior of OpenAPI where <code>additionalProperties</code> defaults to True is retained.
Request Limit in Kilobytes		100000	Only available for WAF enforcers. It defines the size limit for requests in kilobytes.

XSD has the following configuration options for the *Operations* parameters:

Table 58. XSD enforcer plugin configuration options for *Operations*

Key	Default	Description
URI Path	*	The pattern for <code>uri_path</code> .
Choose Method		The method of the HTTP message. The following values are available for <i>Method</i> : • get • head • post • put • delete • connect • options • trace • patch
Status		The status of the HTTP message.
Choose Call Direction		The direction of the message, which must be either request or response.
Choose File		The XSD schema.

6.5.3. Filter

Filter Plugins are lightweight alternatives of *Enforcer Plugins* for filtering unwanted traffic. They only consist of a matcher and an error policy. If the matcher matches, the error policy is applied. This way you can use matchers inline, instead of creating a whole schema-based *Enforcer Plugin* for the simple use cases.

6.5.3.1. Configuring Filter Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Filter* plugin.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 30. Filter Plugin's main page in the Web User Interface



Make sure that any component referenced in the configuration of this component, for example an Error policy or a Matcher selected from the drop-down lists, must remain part of the configuration later as well. Removing any of the referenced components might lead to invalid configuration.

3. Click on the *New* navigation button to create a Filter.
4. Add the name of the Filter Plugin.
5. Add the Body content for the error policy. (Optional)
6. Define the Content type.
7. Choose an error policy from the drop-down list. (Optional)
8. Choose a matcher from the drop-down list. (Optional)
9. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
10. Save the component configuration by clicking the *Save* button.

The Plugin does not override any of the [default error policy](#) options.


Problems are considered errors that lead to the termination of the call. Problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.



If you omit the matcher, the *Plugin* will always be executed. For *Filter plugins* this means aborting **all** calls.

The following values can be configured for the *Filter Plugin*:



Filter

Name : *

Matcher : (Default: always)

Error Policy : (Default: error_policy)

Body :

Content Type :

Figure 31. Configuring a filter plugin in the Web User Interface

Table 59. Filter Plugin's configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the Filter Plugin. This name of the plugin can be referenced from other parts of the configuration.
Matcher	Reference to a <i>Matcher Brick</i> .	Always: If the value is not defined the plugin is always executed.	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If omitted the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	error_policy	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. For details see Error Policy .
Body	Can be defined in free text.		The body of the message sent in case an error policy is applied.
Content Type			The content type of HTTP error request sent, if the filter stops the call. It can be referenced by its name.

6.5.4. Fraud Detector

The Fraud Detector Plugin, leveraging the data collected from the calls by the selectors, evaluates the level of risk with regards to the call. The risk calculated by the Fraud Detector plugin is translated to a score between *0.0* and *100.0*. The lower the score is, the more secure and trustworthy the actual call is. Consequently, the value *0.0* means that the call is perfectly secure, until the value *100.0* identifies a malicious act with the call.

6.5.4.1. Configuring Fraud Detector Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Fraud Detector* plugin.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 32. *Fraud Detector's* main page in the Web User Interface



Make sure that any component referenced in the configuration of this component, for example an Error policy or a Matcher selected from the drop-down lists, must remain part of the configuration later as well. Removing any of the referenced components might lead to invalid configuration.

3. Click on the *New* navigation button to create a Fraud Detector.
4. Add the name of the Fraud Detector.
5. Choose an error policy from the drop-down list. (Optional)
6. Choose a matcher from the drop-down list. (Optional)
7. Choose a *Selector* from the drop-down list. When it is selected click on the plus sign to add it to the configuration.
8. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
9. Save the component configuration by clicking the *Save* button.

See [Error Policy](#) to understand how they shall be applied here.

The following values can be configured for the *Fraud Detector Plugin*:

Fraud Detector

Name : *

Enter Name

Matcher :

Choose Matcher

(Default: always)

Error Policy :

Choose Error Policy

(Default: error_policy)

Selectors : *

Choose Selector

+

Validate

Save

Cancel

Figure 33. Configuring the Fraud Detector plugin in the Web User Interface

Table 60. Fraud Detector Plugin's configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the Fraud Detector. This name of the plugin can be referenced from other parts of the configuration.
Matcher	Reference to a <i>Matcher Brick</i> .	Always: If the value is not defined the plugin is always executed.	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If omitted the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	error_policy	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. For details see Error Policy .
Selectors*	A list of references to <i>Selector Bricks</i> .		A list of Selector that collect information from the call. Selectors can be configured as listed in Selector configuration for the Fraud Detector Plugin.  It is possible to add more data from the selectors to the Fraud Detector Plugin using custom fields, apart from the list in section Selector configuration for the Fraud Detector Plugin. In such cases contact the Balasys Support team.

6.5.5. Insight

It is a Plugin that extracts various data from the call and sends it to external systems (log servers, SIEMs, and other data analysis tools).

6.5.5.1. Configuring Insight Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Insight* plugin.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 34. Insight Plugin's main page in the Web User Interface

3. Click on the *New* navigation button to create an Insight.

The Plugin uses the default *Error policy* by default, that is, the 'insight_default'.

The Plugin overrides the following fields of the [default error policy](#):

Table 61. Default Insight Error Policy

Policy Setting	Default
request	log
response	log

Problems are considered errors that only need to be logged. If that is overridden then problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.

The *Plugin* collects the information from all the selectors and sends them to all the targets.

The collected information from all the selectors is arranged into a dictionary: a list of *key* — *value* pairs. The key can be configured in each selector. Certain selectors might return complex data structures, that are made up of other dictionaries and/or lists. To ensure compatibility with a wide range of *Insight Target* types, such results are flattened. The path inside the complex data structure is encoded into the key for each value. More details are available on this in [Data flattening](#).

4. Add the name of the Insight Plugin.
5. Choose an error policy from the drop-down list. (optional)
6. Choose a matcher from the drop-down list. (optional)
7. Add the message content for the error policy. (optional)
8. Choose a selector from the drop-down list.
9. Select the *Insight Target*.
10. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
11. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Insight Plugin*:




Figure 35. Configuring an insight plugin in the Web User Interface

Table 62. *Insight Plugin's configuration options*

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the insight. This name of the insight can be referenced from other parts of the configuration.

Key	Values	Default value	Description
Matcher	Reference to a <i>Matcher Brick</i> .	Always: If the value is not defined the plugin is always executed.	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If omitted the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	insight_default	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. For details see Error Policy .
Message	Can be defined in free text.	The name of the plugin.	The message part of the log message.
Selectors*	A list of references to <i>Selector Bricks</i> .		A list of Selectors that collect information from the call. It is possible to multiselect more than one selector in this list by clicking on them. The multiple selected elements can then be added to the configuration by clicking on the plus sign.
Targets*	A list of references to <i>Insight Target Bricks</i> .		A list of Insight Targets where the collected information will be sent to.

6.5.6. Serializer

The *Serializer Plugin* is responsible for serializing and compressing the structured data to the format of the HTTP message's body.

A typical Security Flow configuration starts with a *Deserializer* and finishes with a *Serializer*. If the Security Flow starts with a *Deserializer*, then the body will be automatically decompressed and then re-compressed at the end. This ensures that transferred HTTP bodies are syntactically correct and that they are reconstructed to avoid transferring potentially crafted content.

The Serializer Plugin understands the Content-Type HTTP header and can work with JSON and XML content.

6.5.6.1. Configuring Serializer Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Serializer*.

The configuration window that appears presents the default Serializers, as listed in [Default objects - PLUGINS](#) and the configuration values already set by the user:




Figure 36. The serializer main page in the Web User Interface

3. Click on the *New* navigation button to create a Serializer.

The Plugin does not override any of the [default error policy](#) options.

Problems are considered errors that lead to the termination of the call. Problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.

Continue configuring the serializer with the following steps:

4. Add the name of the serializer.
5. Select the type of the *Serializer*.
6. Choose an Error policy from the drop-down list.
7. Choose a Matcher from the drop-down list.
8. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
9. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Serializer Plugin*:

Serializer




Figure 37. Configuring a serializer in the Web User Interface

Table 63. Serializers' configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the serializer. This name of the serializer can be referenced from other parts of the configuration, that is, the Plugin is reusable.
Type*	The value can be selected from a drop-down list. The value can be: • JSON • XML		There are two types of predefined (de)serializer plugins.
Matcher	Reference to a <i>Matcher Brick</i> .	Depending on which 'Type' was selected for the <i>Serializer</i> , the default value can be: <code>json_content</code> or <code>xml_content</code> .	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If no matcher is configured the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	<code>error_policy</code>	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. If no error policy is configured, the plugin type's default error policy is applied. For details see Error Policy .

6.5.7. Deserializer

It is a Plugin responsible for parsing the HTTP message's body to structured data. This ensures that a message is well-formed. The structured data will also be consumed by other Plugins that operate on the body of the message.

A typical Security Flow configuration starts with a *Deserializer* and finishes with a *Serializer*. If the Security Flow starts with a *Deserializer*, then the body will be automatically decompressed and then re-compressed at the end. This ensures that transferred HTTP bodies are syntactically correct and that they are reconstructed to avoid transferring potentially crafted content.

Deserializers understand the *Transfer-Encoding* HTTP header and decompress data by using the *gzip*, *deflate* and *brrotli* algorithms.

6.5.7.1. Configuring Deserializer Plugins

1. Click on the *PLUGINS* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *PLUGINS*.
2. Select *Deserializer* plugin.

The configuration window that appears presents the default Deserializers, as listed in [Default objects - PLUGINS](#) and the configuration values already set by the user:




Figure 38. The deserializer's main page in the Web User Interface

2. Click on the *New* navigation button to create a Deserializer.

The Plugin does not override any of the [default error policy](#) options.

Problems are considered errors that lead to the termination of the call. Problems in the request are reported back to the client, while errors in the response are suppressed to avoid information leak.

See [Error Policy](#) to understand how defaults are applied.

3. Add the name of the deserializer.
4. Select the Type of the Deserializer.
5. Choose an Error policy from the drop-down list.
6. Choose a Matcher from the drop-down list.
7. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
8. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Deserializer Plugin*:




Figure 39. Configuring a deserializer in the Web User Interface

Table 64. Deserializers' configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the deserializer. This name of the deserializer can be referenced from other parts of the configuration.
Type*	The value can be selected from a drop-down list. The value can be: • JSON • XML		There are two types of predefined (de)serializer plugins.
Matcher	Reference to a <i>Matcher Brick</i> .	Depending on which 'Type' was selected for the <i>Deserializer</i> , the default value can be: <code>json_content</code> or <code>xml_content</code> .	It decides if the Plugin should be executed based on the call's details. For details see Matcher . If omitted the Plugin is always executed.
Error Policy	Reference to an <i>Error Policy Brick</i> .	<code>error_policy</code>	It defines a custom error policy to be applied if the Plugin reports an error. The settings of the Error policy here override the Security Flow's default error policy. For details see Error Policy .

Key	Values	Default value	Description
Charset Conflict	- drop: If this parameter is set to 'drop', the configuration instructs to drop the call in case there is conflict for the character set in the message's header. - log: If the value is set to 'log', the system will use either type of the character set defined and will log the error.	drop	This parameter needs to be configured in case the 'Type' of the Deserializer is set to XML. In XML messages, there might be a conflict in the definition of the character set. The XML and the HTTP headers might instruct to use different character sets. The conflicting information on the character set can be configured to be handled in two different ways, that is the call dropped, or the call maintained and the error logged, depending on the settings of this parameter.

6.6. SERVICES - Configuration units

Proxedo API Security is based on a micro-services architecture.




Figure 40. The SERVICES main page in the Web User Interface

6.6.1. Backend

Backends are a set of servers for a given API endpoint.

Their configuration is made up of two main parts:

- a list of servers: port pairs and how to route traffic to them
- TLS configuration for talking to the servers

6.6.1.1. Configuring Backends

1. Click on the *SERVICES* main navigation item in the Left navigation area. Alternatively you can also click on the

 sign to open up the sub-navigation items of *SERVICES*.

2. Select *Backend*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 41. The main page for Backend

3. Click on the *New* navigation button to create a Backend.
4. Name the *Backend* configuration.
5. Provide the values for the Servers parameter: *Host* and *Port*.
6. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
7. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Backend Service*:

Backend

Name : *

Servers : *

+

Load Balancing ^

Method :

▼

(Default: Failover)

Connection Timeout :

(Default: 30000)

Retry After :

(Default: 600000)

Use TLS :

(Default: False)

Figure 42. Configuring backend in the Web User Interface

Table 65. Backend configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the backend. This name of the backend can be referenced from other parts of the configuration.
Servers*	There are two values to be configured: • Host: The name or IP address of the host to connect to. • Port: The port on host to connect to. (You can add the values by clicking the '+' sign.)		The list of servers that serve API endpoint(s). See Backend servers' configuration for details.
Load Balancing / Method	One of the following methods can be used: • Failover: use the first server while available, then fail over to the next • Round Robin: use all servers in a round-robin fashion If the value is not configured the default value will be added.	Failover	Load balancing method to use.
Load Balancing / Connection Timeout	If the value is not configured the default value will be added.	30000	The connection timeout in milliseconds towards the backend server.
Load Balancing / Retry After	If the value is not configured the default value will be added.	600000	The time in milliseconds before connection towards the backend server is started again in case of a connection failure.

Key	Values	Default value	Description
Use TLS	True or False.	False	Enables using TLS in the connection towards the backend servers.
Backend-facing TLS*	Reference to a <i>TLS Brick</i> of type <i>Backend-facing TLS</i> .		The TLS configuration towards the backend servers. See Configuring Backend-facing TLS Bricks for details. Mandatory if <i>Use TLS</i> is set to <i>True</i> .

6.6.2. Endpoint

An endpoint holds together all the policies that apply to a certain API endpoint:

- List of URLs
- The default error policy for the endpoint
- The backend to which requests will be forwarded
- The authentication method
- The security flow that will be applied to the traffic

6.6.2.1. Security Flow

The Security Flow definition in an endpoint lists what happens to the traffic on a given endpoint.

To understand how requests flow through PAS, see [Understanding processing flow](#). The Security Flow starts when the Transport Director has already set up client connection and routed the request to the Flow Director. At this point the TLS and HTTP layers are already processed, but the content in the body of the request is available only in raw format and has not been parsed yet.

At this stage, the configuration security flow decides on what happens to the traffic by applying a list of *Plugins* one by one. *Plugin* is a collective name for Enforcers, Insights, Filters, etc. Once, all the *plugins* have processed the request, the control is handed back to the *Transport Director* which routes the request to a backend server, and comes back with the response after handling TLS and HTTP. At this point, the *Flow Director* applies another list of *Plugins* to response, and once done, it hands back the response to the *Transport Director* which in turn returns that to the client.

If at any point an error occurs, the error policy is applied — which might either mean to lead to logging the error or to terminating processing and returning an error indication to the client.

Plugins can override the endpoint's error policy.

Also note that different *Plugins* need different data. An Insight that applies a JMESPath query needs parsed JSON, while one that extracts value from an HTTP header field does not. Other *Plugins* provide these required values, like a JSON deserializer *Plugin*. It is important that the *Plugins* are configured in such an order that the required data is made available beforehand.

6.6.2.2. Configuring Endpoints

1. Click on the *SERVICES* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *SERVICES*.
2. Select *Endpoint*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 43. The main page for Endpoint

3. Click on the *New* navigation button to create an Endpoint.
4. Name the *Endpoint Service*.
5. Select the *Backend* parameter from the drop-down list. Backend servers are configured under the *SERVICES* main navigation item.
6. Complete a Security Flow from the configured (and the default) plugins. For more details, see [Security Flow](#).
 - Choose the *Request* plugin from the drop-down list. The Plugin options available from the drop-down list have been configured under the *PLUGINS* main navigation item.
 - Choose the *Response* plugin from the drop-down list. The Plugin options available from the drop-down list have been configured under the *PLUGINS* main navigation item.
7. Provide the URL to address the API endpoint.
8. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
9. Save the component configuration by clicking the *Save* button.



While ports must be unique, as only one listener can bind to a specific port, it is perfectly valid to route incoming traffic from multiple listeners to the same endpoint.

A typical security flow is configured with the *plugins* in the following order:

- a *Deserializer Plugin* that decompresses and parses the content of the request
- an *Enforcer Plugins* that ensure the call is valid
- *Insight Plugins* that extract important data from certain calls
- a *Serializer Plugin* that rebuilds the contents of the request



The *Plugin* configurations are reusable.

The following values can be configured for the *Endpoint Service*:

Endpoint

Name : *
 Traffic Mode :
 URLs : *
 Use Authentication :
 URL Rewrite Rule :
 SNI Rewrite Rule :
 Backend : *
Failure Policy ^
 Silent :
 Code :
Security Flow ^
 Request :
 Response :

Enter Name
 Choose Traffic Mode (Default: Enabled)
 Enter URLs +
 False Default True (Default: False)
 Enter Rule
 Enter Rule (Default: <Dynamic>)
 Choose Backend
 False Default True (Default: True)
 Enter or choose Code (Default: 500)
 Choose Plugin +
 Choose Plugin +

Validate Save Cancel

Figure 44. Configuring endpoint in the Web User Interface

Table 66. Endpoint configuration

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the endpoint. This name of the endpoint can be referenced from other parts of the configuration.
Traffic Mode	There are three options: • Enabled: the plugins will process the calls. This is the normal operation. • Disabled: calls will be automatically rejected with a silent HTTP 400 error code. This mode completely prevents traffic from reaching the backends. • Pass-through: allows calls to pass through the endpoint without interception or modification.	Enabled	This option allows the user to disable security flows on the endpoint for troubleshooting purposes, without altering the existing configuration. The pass-through option enables traversing traffic, while prevents any modification or interference.
URLs*			The URLs which the clients use to address the API endpoint.

Key	Values	Default value	Description
URL Rewrite Rule			The URL by which the backend servers understand incoming requests. When set, two transformations take place: • The original URL will be replaced by the matching URL configured for the <i>Endpoint</i>. • The <i>Host</i> header will be replaced by the host indicated in the URL rewrite rule.
SNI Rewrite Rule		<Dynamic>	It can be used to rewrite the Server Name Indication (SNI) field in a TLS handshake towards the backends. The <Dynamic> default value means that the SNI value used towards the backend will be the same as the value of the Host header, either coming from the client or defined in the URL Rewrite Rule.
Backend*	Reference to a <i>Backend Service</i> .		Backends are a set of servers for a given API endpoint. For more details, see Backend .
Use Authentication	True or False.	False	Enables using authentication for incoming calls. NOTE: for authenticated <i>Endpoints</i> , it is recommended to use a <i>Listener</i> with TLS.
Authentication*	Reference to an <i>Authentication Brick</i> .		The configuration of the selected authentication method. Mandatory if <i>Use Authentication</i> is set to <i>True</i> .
Failure Policy	Two values have to be configured: • Silent • Code	Silent: True; Code: 500	With the help of the Failure Policy, it can be configured whether the client shall receive notification or not, and whether the notification shall contain the code on the type of the failure. The values in details are as follows: • Silent: If the silent value is active, the Failure policy is not reported. If the silent value is inactive, the failure policy is reported towards the user. • Code: Code is an HTTP response code here, that can be set manually or from the provided drop-down list.

Key	Values	Default value	Description
Security Flow*	The security flow process requires the configuration of the following values, each containing a list of <i>Plugins</i>. • Request* • Response*		The values in details are as follows: • Request: The Transport Director sets up client connection and routes the request to the Flow Director. The Request has numerous values to be configured. For more details, see Security Flow. • Response: The Transport Director routes the request to a backend server, and comes back with the response after handling TLS and HTTP. For more details, see Security Flow. Note, that both for the Request and Response parameters, it is possible to multiselect more than one element in the list by clicking on them. The multiple selected elements can then be added to the configuration by clicking on the plus sign.

6.6.3. Listeners

Listeners are network endpoints where services are exposed to the network. They consist of:

- a listening port
- an optional client-side TLS configuration if HTTPS is used
- a list of endpoints that handle the traffic.

Since these are the entry points for client traffic it must be routed here on the network.

6.6.3.1. Configuring Listeners

1. Click on the *SERVICES* main navigation item in the Left navigation area. Alternatively you can also click on the  sign to open up the sub-navigation items of *SERVICES*.
2. Select *Listener*.

In the configuration window that appears, you can either see the empty parameter values that can be configured for the actual component or you can see already configured component(s) and their parameters. The already configured components with defined parameters can be default components available in the system by default, or can be components configured by the administrator:




Figure 45. Listener's main page in the Web User Interface

- Click on the *New* navigation button to create a Listener.

At least one listener must always be configured in the Proxedo API Security configuration.

- Name the *Listener Service*.
- Select the *Client TLS* parameter from the drop-down list. The client side TLS parameter values have to be defined previously under *BRICKS*.
- Select the *Endpoint* from the drop-down list. The endpoint values have to be defined previously under *SERVICES/Endpoint*.



All endpoints in the list must have the same backend and backend URL configured.

- Fill in the *Port* information. If it is not configured, the default value will be applied.



Ports must be unique, only one listener can bind to a specific port. It is however perfectly valid to route incoming traffic from multiple listeners to the same endpoint.

- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Listener Service*:




Figure 46. Configuring a listener in the Web User Interface

Table 67. Listeners' configuration options

Key	Values	Default value	Description
Name*	Free text. Alphanumeric, may contain underscores, may not start with a number.		The name identifying the listener. This name of the listener can be referenced from other parts of the configuration.
Port	Any port value can be defined. Note that the port value has to be within the range configured in the docker.	49000	The number of the port the listener binds to.
Endpoints*	A list of references to <i>Endpoint Services</i> .		The list of endpoint(s), as defined under Endpoint that serve traffic coming in on the listener.
Use TLS	True or False.	False	Enables using TLS in the connection towards the clients.
Client TLS*	Reference to a <i>TLS Brick</i> of type <i>Client TLS</i> .		The TLS configuration towards the clients. See TLS for details. Mandatory if <i>Use TLS</i> is set to <i>True</i> .

6.7. SYSTEM - Configuration units




Figure 47. The SYSTEM main page in the Web User Interface

6.7.1. Log

If at any point an error occurs during the Security Flow, the error policy is applied and logging takes place if configured so.

6.7.1.1. Configuring Log

1. {step_open_systems}
2. Select *Log*.
3. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
4. Save the component configuration by clicking the *Save* button.




Increasing the verbosity hugely increases the amount of logs generated, and will reduce performance.



The logs at the highest level of verbosity (9) might include sensitive information, such as passwords.

The following values can be configured for the *Log System*:



Changes Status Configuration Backup

Help user

Log : Default

Name : * default

Verbosity : Enter Verbosity (Default: 3)

Message Filter Expression : Enter Expression +

Validate Save Cancel

Figure 48. The main page for Logs

Table 68. Log configuration

Key	Values	Default value	Description
Name*	Log has a default name 'default', that cannot be changed.		The name identifying the log configuration.
Verbosity	The value can take number format.	3	The verbosity of logging. It must be between 1-9.
Message Filter Expression	A list of message filter expressions. A single message filter expression consists of a log category, a colon, and a number specifying the verbosity level of that given category. Categories match from left to right and wildcards can be used. For example: <code>http.*:5,core.info:3</code> . The last matching entry will be used as the verbosity of the given category. If no match is found the default verbosity specified with <i>verbosity</i> is used.	*.accounting:4,core.summary:4	Set verbosity mask on a per category basis. Each log message has an assigned multi-level category, where levels are separated by a dot.

6.7.2. Transport Director

The **Transport Director** manages the transport layer of API connections:

- handles network connections from the client
- handles network connections towards the backends
- handles TLS on these connections
- load-balances between multiple backend servers
- load-balances between multiple *Flow Directors*
- enforces HTTP protocol validity in calls

6.7.2.1. Configuring Transport Director

1. {step_open_systems}
2. Select *Transport Director*.
3. Click the *Validate* button to check if the defined parameters are of the correct type and all required

parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.

4. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Transport Director System*:




Figure 49. The main page for Transport Director

Table 69. Transport Director configuration

Key	Values	Default value	Description
Name*	<i>Transport Director</i> has a default name 'default', that cannot be changed.		The name identifying the Transport Director configuration. This name of the <i>Transport Director</i> can be referenced from other parts of the configuration.
Enable Core Dump	True or False.	False	Enables core dumps on failures.
Threads		0	Set the maximum number of threads that can be used in parallel. Note, that setting the value to zero means that the number of threads that can be used in parallel is unlimited.

6.7.3. Fraud Detector

The Fraud Detector, leveraging the data collected by the Fraud Detector plugin, establishes the actual connection with the Fraud API for an evaluation on the data of the calls.

Although the average response time of the Fraud API is half second, depending on the size and the complexity of the traffic to be investigated the response time might increase up to three seconds. Consequently, it is recommended to carefully identify the content selected for detection.

It is also recommended to consider that the API evaluates the maximum of 10 requests per second, therefore it is important to carefully define the matcher for the fraud detection, so that the load of requests is not unnecessarily high and the requests exceeding the value of 10 requests per second do not get dropped.

There are three recommended data types to be configured as selectors when configuring the Fraud Detector plugin, namely the IP address, the phone number and the e-mail address. For more details on how to configure Fraud Detector plugin, see [Fraud Detector Plugin's configuration options](#).

6.7.3.1. Configuring Fraud Detector

1. {step_open_systems}
2. Select *Fraud Detector*.

Continue with the steps if the Fraud Detector is required in active state:

3. Set the Fraud Detector system to active state. The Fraud Detector is set to 'inactive' state by default, as for the 'active' state license is required.
4. Define the API Endpoint destination.
5. Fill in the API key. The API Key is provided together with the license purchased for the Fraud Detector.
6. Add the value for the Connection Timeout parameter. The value has to be provided in seconds.
7. Provide the value for the Response Timeout parameter. The value has to be provided in seconds.
8. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
9. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *Fraud Detector System*:




Figure 50. The Fraud Detector's main page in the Web User Interface

Table 70. Fraud Detector's configuration options

Key	Values	Default value	Description
Name*	<i>Fraud Detector</i> has a default name 'default', that cannot be changed.	default	The name identifying the Fraud Detector. This name of the Fraud Detector can be referenced from other parts of the configuration.
Active	True or False.	False	If the license for the Fraud Detector is purchased, the system can be activated.

If the Fraud Detector system is set to active, the following further parameters are available:

Fraud Detector : Default




Figure 51. Configuring an active Fraud Detector in the Web User Interface

Table 71. The active Fraud Detector's configuration options

Key	Values	Default value	Description
Client Configuration			Configure the parameters of Fraud Detector.
API Endpoint		The default value is as follows: https://fraud-api.balasys.hu/api .	The API endpoint.
API Key*	The value for the API Key is provided by the purchase of the Fraud Detector license.		The API key is provided when the license for the Fraud Detector is purchased.
Connection Timeout	The value can be provided in seconds.	5	The time limit for establishing connection with the provided URL.
Response Timeout	The value can be provided in seconds.	10	The time limit for how long the PAS awaits the answer from the Fraud API after an established connection.

6.7.4. WAF Ruleset Updater

The Web Application Firewall (WAF) Ruleset Updater System is designed to automatically update the ruleset used for WAF enforcers, it is thereby critical in ensuring real-time protection against zero-day attacks by maintaining an up-to-date defense mechanism.

To activate this system, extra credentials will be necessary which can be obtained from the Balasys sales team.

6.7.4.1. Configuring WAF Ruleset Updater

1. {step_open_systems}

2. Select *WAF Ruleset Updater*.

Continue with the steps if the WAF Ruleset Updater is required in active state:

- Set the WAF Ruleset Updater system to active state. To activate the WAF Ruleset Updater a license is required. To acquire a license, contact our sales team at the e-mail address <sales@balasys.hu>.
- Fill in the API Username. The API Username is provided together with the license purchased for the WAF API.
- Fill in the API Password. The API Password is provided together with the license purchased for the WAF API.
- Add the value for the Poll Interval Seconds parameter. The value has to be provided in seconds.
- Add the value for the Connection Timeout Seconds parameter. The value has to be provided in seconds.
- Provide the value for the Response Timeout Seconds parameter. The value has to be provided in seconds.
- Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
- Save the component configuration by clicking the *Save* button.

The following values can be configured for the *WAF Ruleset Updater System*:





Figure 52. The WAF Ruleset Updater's main page in the Web User Interface

Table 72. WAF Ruleset Updater's configuration options

Key	Values	Default value	Description
Name*	<i>WAF Ruleset Updater</i> has a default name 'default', that cannot be changed.	default	The name identifying the WAF Ruleset Updater. This name of the WAF Ruleset Updater can be referenced from other parts of the configuration.
Active	True or False.	False	The system needs to be activated only if a <i>WAF Enforcer</i> is to be used.

If the WAF Ruleset Updater system is set to active, the following further parameters are available:

WAF Ruleset Updater : Default



Name : *

Active : ☐ False ☒ Default ☐ True (Default: False)

Updater Configuration ^

API Username : *

API Password : *

Poll Interval Seconds : (Default: 3600)

Connection Timeout Seconds : (Default: 5)

Response Timeout Seconds : (Default: 10)

Figure 53. Configuring an active WAF Ruleset Updater in the Web User Interface

Table 73. The active WAF Ruleset Updater's configuration options

Key	Values	Default value	Description
Updater Configuration			Configure the parameters of WAF Ruleset Updater.
API Username*	The username required to download and update the WAF enforcer's ruleset. The value for the API Username is provided with the purchase of the WAF license.		The API Username is provided when the license for the WAF API is purchased.
API Password*	The password required to download and update the WAF enforcer's ruleset. The value for the API Password is provided with the purchase of the WAF license.		The API Password is provided when the license for the WAF API is purchased.
Poll Interval Seconds	The value must be provided in seconds.	3600	The time between two ruleset updates.
Connection Timeout Seconds	The value must be provided in seconds.	5	The time limit for how long the PAS awaits the answer from the WAF API to establish the connection.
Response Timeout Seconds	The value must be provided in seconds.	10	The time limit for how long the PAS awaits the answer from the WAF API after an established connection.

6.7.5. License

The License System holds the License File brick currently in use.

6.7.5.1. Configuring License

1. {step_open_systems}

2. Select *License*.
3. Choose an uploaded License File brick from the drop-down list.
4. Click the *Validate* button to check if the defined parameters are of the correct type and all required parameters have been filled out for configuring the component. If the configuration is erroneous or incomplete, the Web UI provides a warning that the 'Component validation failed'. Also a warning with information on the missing or faulty elements appears at the problematic field. If the configuration of the component is valid, after clicking the *Validate* button, a 'Component validation successful' notification is shown.
5. Save the component configuration by clicking the *Save* button.

The following values can be configured for the *License System*:





Figure 54. The License's main page in the Web User Interface

Table 74. License's configuration options

Key	Values	Default value	Description
Name*	<i>License</i> has a default name 'default', that cannot be changed.	default	The name identifying the License. This name of the License can be referenced from other parts of the configuration.
License File*	A reference to a <i>File Brick</i> of the License type.		The License File that is to be used.

6.8. System-wide status information



Changes Status Configuration Backup

Help user

Status

License

Serial	Version	Backend Limit	Endpoint Limit	Date	Customer	Valid Not Before	Valid Not After	License Type
ae949306-9fba-c928-4ce4-66f24cf285d0	4.0	50	50	2023/08/25	Balasys IT Zrt.	2023/08/25	2026/08/25	!!!!!!This copy is an evaluation version!!!!!!

Figure 55. System-wide status information

6.8.1. License details

Details on the active license are presented in a table.

- **Serial:** The unique ID of the license.
- **Backend Limit:** The number of configured Backends that this license supports.
- **Endpoint Limit:** The number of configured Endpoints that this license supports.
- **Date:** The date this license was issued.
- **Customer:** The customer this license was issued to.
- **Valid Not Before:** The license is not valid before this date.
- **Valid Not After:** The license is not valid after this date.
- **License Type:** The type of this license.

6.9. Checking and finalizing changes in Proxedo API Security configuration

It is possible to list and check any changes made to the PAS configuration until the changes have not been applied with the *Apply Configuration* button.

Click on the *Changes* button in the Top-left navigation area to list the changes made to the configuration.

The following pieces of information are displayed:

- configuration integrity problems
- changes made to any of the configuration components




Figure 56. Checking changes made to the configuration

6.9.1. Configuration Integrity

For errors in configuration integrity, the following pieces of information are displayed in table format:

- **Type:** The type of the integrity problem, for example cycle detection.
- **Description:** Provides details on the nature of the integrity error.



Until the configuration integrity errors listed here are not corrected, the configuration cannot be applied.

For details on configuration integrity errors, see the examples in section [Integrity errors](#).

6.9.2. Configuration Changes

For changes on the configuration components, the following pieces of information are displayed in table format:

- **Type:** Type denotes the category (Brick, Plugin, Service, System) and the class (for example, Matcher, Filter, Log) of the configuration component, for example Brick/Matcher.
- **Name:** The name of the configuration component is displayed here, to which the actual change has been made.
- **Change:** The nature of the change made to the configuration component is provided here, that is, *Added*, *Edited*, *Deleted* or *Broken Reference* (no direct change, but it refers to another component that is now missing or invalid).
- **Validity:** This field informs the user on whether the configured component is valid or not, as follows:
 -  - Any instance marked with this sign is invalid.
 -  - Any instance marked with this sign is valid.



Click on the  sign to see more information on why the instance was found invalid.

Invalid configuration components can be corrected and revalidated by using the *Validate* button, available at each component's configuration page. For more information, see section *Component-level validation* in chapter [Applying and validating Proxedo API Security configuration](#).

- **Actions:** This field provides several actions that can be taken.
 - Difference - Show the difference between the edited and the running versions of the component.
 - Edit - Edit the configuration data for a component. If the edit button is disabled, it means that the instance has been deleted.
 - Discard - Undo any configuration changes to a component. If the discard button is disabled, no direct changes have been made to the actual component.

By selecting the *Discard All* button, it is possible to discard all changes made to the configuration. However, the changes that have been applied to the configuration already and the factory-supplied configuration elements cannot be discarded.

6.10. Applying and validating Proxedo API Security configuration

PAS configuration can be checked and validated on two levels:

- component-level validation
- validating the whole configuration

6.10.1. Component-level validation

Component-level validation takes place while configuring the actual elements of the configuration and by using the *Validate* button on the Web UI page of the specific component.

If the configuration of the component is erroneous or not adequate, the Web UI provides a warning that the *Component validation failed*. Also a warning with information on the missing details appears at the problematic field for the user.




Figure 57. Component validation failed

If the configuration of the component is satisfactory, after clicking the *Validate* button, the user receives the *Component Validation successful* notification. Click OK. For related errors see, section [Validation errors](#).




Figure 58. Component validation successful

6.10.2. Validating the whole configuration

Configuration integrity errors can be viewed on the *Changes page*, along with a summary of valid and invalid component changes. To make it available click the *Changes* button so that all the changes made to any component by the user will be visible. For related errors, see section [Validation errors](#).

6.10.3. Applying the whole configuration

The *Apply Configuration* button is available from the *Changes* page. To make it available click the *Changes* button so that all the changes made to any component by the user will be visible. In order to take the changes into effect, click the *Apply Configuration* button. The configuration can only be applied if all changes are valid. When applying the configuration by using the *Apply Configuration* button, the Web UI provides either of the following messages:

- The configuration is applied successfully. Click *OK*.




Figure 59. *Apply Configuration* result - successful

- The configuration failed.
 If applying the configuration failed, the Web UI also provides an additional pop-up window with the description of the problem. The problems can be as follows:
 - At least one of the services failed to start, the previous configuration settings have been restored.
 - Restoring the original configuration was not successful.



During the process of applying the configuration, no changes can be completed to the configuration. The process however shall not take more than 10 seconds.

6.10.4. Validation errors

In case the configuration could not be applied, the following result messages help the user to correct the configuration and achieve a valid configuration.

6.10.4.1. Component-related errors

These errors are the results of the validation of the actual components. By correcting these the user can achieve a functioning configuration.

6.10.4.1.1. Missing data for required fields

Each component has compulsory configuration fields that must be filled in. In case any of those fields are left empty, the Web UI provides a *Missing data for required field* notification when the component is validated, that is, the *Validate* button is used. Each compulsory field is highlighted with a * sign.

Example

The *Insight Target* component requires the *Host* field to be filled in, otherwise the component's configuration is not valid.

Error message: **Missing data for required field.**

Changes
Status
Configuration Backup

Component validation failed

Insight Target

Name : *
Type : *
Flatten :
Flatten Separator :
Remote Connection ^
Host : *
Port :
Protocol :
IP Protocol :
Use TLS :
Flush Lines :
Data Format :
Second Fraction Digits :
Time Zone :
Report Config Load :
Mask Credit Card Numbers :
Enable Heartbeat :

Insight Target example

invalid name.

Syslog

False Default True

(Default: True)

Enter Separator

(Default: /)

Enter Host

Missing data for required field.

Enter Port

Choose Protocol

(Default: TCP)

Choose IP Protocol

(Default: 4)

False Default True

(Default: False)

Enter Lines

(Default: 0)

Choose Data Format

(Default: SData)

Enter Digits

(Default: 3)

Enter Zone

(Default: GMT)

False Default True

(Default: False)

False Default True

(Default: False)

False Default True

(Default: False)

Validate
Save
Cancel

Figure 60. Missing required field - Insight Target

6.10.4.1.2. Missing reference

This error indicates that the component references a non-existing component.

Example

The user creates an Error Policy, *error_policy_1* which is referenced in a Filter. Following that, this specific Error Policy, *error_policy_1* is deleted from the configuration. This results in a missing reference in the Filter.

Error message: **Reference to a non-existing component: error_policy_1.**



To correct the missing reference, navigate to the Filter component. In order to clear the invalid

reference to the missing component, the  icon has to be selected on the right side of the Error Policy drop-down list. By clicking this icon, the configuration data is cleared from this selection.

6.10.4.1.3. Port conflict

This error indicates that two or more Listeners are configured to use the same port. This leads to a failed configuration.

Example

Two Listeners are configured to use the same port.

Error message: **listener_1 uses the same port as listener_2.**

6.10.4.2. Integrity errors

6.10.4.2.1. Cycle detection

Error message: **Reference cycle detected in configuration: brick/matcher/matcher_1→brick/matcher/matcher_2→brick/matcher/matcher_1.**

This error indicates that there is a cycle of references between component instances.

Example

If the compound matcher *matcher_1* is configured to reference the compound matcher *matcher_2* and the compound matcher *matcher_2* is also referencing the compound matcher *matcher_1*, there will be a cycle of references between these two matchers.


6.10.4.2.2. Required instance is missing

Error message: **At least one service/listener must be configured.**

This error indicates that a required instance is not configured.

6.10.4.2.3. Fraud Detector plugin configured with the Fraud Detector system in inactive state

Error message: **Fraud Detector Plug-in must not be part of Security Flows while the Fraud Detector System is disabled.**



The screenshot shows the Proxedo API Security configuration interface. On the left is a sidebar with navigation links: BRICKS, PLUGINS, SERVICES, and SYSTEM. The main area is titled 'Changes' and contains a 'Configuration Integrity' section with a table of errors. Below this is a 'Configuration Changes' table listing various components and their status.

Type	Description
INCOMPATIBLE COMPONENTS	Fraud Detector Plug-in must not be part of Security Flows while the Fraud Detector System is disabled. Offending instance: test_endpoint endpoint-> request-> test_fraud_detector plugin

Type	Name	Change	Validity	Actions
Plugins/Fraud Detector	test_fraud_detector	Added	⊙	🗑️ 🔗 🔄
Services/Backend	test_backend	Added	⊙	🗑️ 🔗 🔄
Services/Endpoint	test_endpoint	Added	⊙	🗑️ 🔗 🔄
Services/Listener	test_listener	Added	⊙	🗑️ 🔗 🔄
System/Fraud Detector	default	Added	⊙	🗑️ 🔗 🔄
System/License	default	Added	⊙	🗑️ 🔗 🔄
System/Log	default	Added	⊙	🗑️ 🔗 🔄
System/Transport Director	default	Added	⊙	🗑️ 🔗 🔄
System/WAF Ruleset Updater	default	Added	⊙	🗑️ 🔗 🔄

Figure 61. Fraud Detector plugin integrity error

This error indicates that there is a Fraud Detector plugin configured, however, the Fraud Detector system is not activated. In order to solve this integrity error, either the Fraud Detector plugin has to be removed from the configuration, or, in case the license for the Fraud Detector is purchased, the Fraud Detector system has to be activated and configured.

6.10.4.2.4. WAF Enforcer plugin configured with the WAF Ruleset Updater system in inactive state

Error message: **WAF Enforcer Plug-in cannot be cannot be configured if the WAF Ruleset Updater System is disabled.**

This error indicates that there is a WAF Enforcer plugin configured, however, the WAF Ruleset Updater system is not activated. In order to solve this integrity error, either the WAF Enforcer plugin has to be removed from the configuration, or, in case the license for the WAF Ruleset is purchased, the WAF Ruleset Updater system has to be activated and configured.

6.10.4.2.5. Insight Message field collision

Error message: **The message keys of some JSON formatted Syslog Insight Targets are in conflict with the Save As Key fields of some Selectors.**

This error indicates that there is an Insight that contains Selectors and Targets that have conflicting configurations. Insight plugins have a Message field that can identify the originating plugin in a log. Certain Insight Target bricks can be configured using the Include Message field to include this Message field in the log line even if the data format does not support a dedicated field for this (JSON for example). In this case, the field where the message should be keyed to can be configured using the Message Key field on the Insight Target. This key can conflict with the Save As Key field of Selectors in the same Insight, leading to lost data.

6.10.4.2.6. License limit exceeded

Error message: **The number of configured Backends (11) exceeds the limit allowed by the active license (10).**

This error indicates that the Backend or Endpoint limit of the selected license is exceeded. Either a different license should be purchased and configured in the License system, or fewer Backend or Endpoint services configured to stay below the limit.

Error message: **The active license's version (2.0) is not valid for this product version.**

This error indicates that the active license is not supported by the installed product version. This typically occurs when the license was issued for an older or newer product version than the currently running.

6.11. Backup and restore running or user configuration for Proxedo API Security

It is possible to backup and restore the Proxedo API Security configuration in the Web UI.




Figure 62. Backup and restore running or user configuration for Proxedo API Security

In order to export any configuration information from the system, complete the following steps:

1. Select the *Configuration Backup* button.
2. To export a configuration, select the type of the configuration to be exported at the *Export configuration* button. The following options can be selected from the drop-down menu:
 - Running: This export option downloads the configuration settings of the currently running configuration.
 - User: This export option downloads the default configuration settings of the system.

The configuration will be downloaded in .zip file format.

3. To import an existing configuration file, select the empty field beside *Import configuration*. Only .zip file formats can be uploaded.
4. Select the *Download* or the *Upload* buttons to finish the activity. The system will ask you to define the *Insight Target* or source destination for the activity. Note that only files in .zip format can be downloaded or uploaded.



In case of importing a configuration file, the system will notify the user that by importing a configuration file, the existing configuration will be overwritten: 'This operation overwrites user configuration. Are You sure?'

7. Operation of Proxedo API Security in Kubernetes environment

This section introduces different methods of inspecting a PAS service state. For inspecting a PAS service state, it is recommended to use selectors, as selectors utilize all the three labels that are added to most objects of the PAS installation.

The three labels are as follows:

- **app**: This label is present on each object with the value of proxedo-api-security.
- **component**: This label is present on all objects that can be associated with any of the three main components of PAS, such as :
 - *mgmt* for the management object
 - *core* for the core objects
 - *storage* for the storage objects
- **subcomponent**: This label is attached to all objects that are directly and exclusively associated with one subcomponent (services, deployments, pods, network policies, etc.). The value of this label is always the name of the subcomponent, for example, *flow-director*, *blob-store*, *config-api*, etc. Since objects are named, using the *proxedo-api-security-<subcomponent-name>* convention, using the *proxedo-api-security-flow-director* object name is most often equivalent to using the *subcomponent=flow-director* selector. Using the selector can be more advantageous, especially with pods, if there are multiple running instances. Since pod names are suffixed with dynamically changing hashes, using a specific pod name can be both inconvenient and sometimes too narrow.

These labels are useful for semantically narrowing down the focus of queries about kubernetes objects.

7.1. Querying objects

By using the `kubectl get` command, objects can be queried with basic information about them.

Run the `kubectl get pods --namespace=proxedo-api-security --selector=app=proxedo-api-security` command to get the list of pods related to PAS.

The output will be similar to the following example:

Example output for querying objects

NAME	READY	STATUS
RESTARTS AGE		
proxedo-api-security-blob-store-768f54bddd-fpd2v	1/1	Running
0 20m		

proxedo-api-security-config-api-5b8b845744-htswp	1/1	Running
0 20m		
proxedo-api-security-consul-65f4c78f-26bsg	1/1	Running
0 20m		
proxedo-api-security-content-filtering-director-55b859df9-sztpw	1/1	Running
0 20m		
proxedo-api-security-flow-director-7459896d6c-k9ttm	0/1	
ContainerCreating 0 20m		
proxedo-api-security-frontend-84798447c4-srvvj	1/1	Running
0 20m		
proxedo-api-security-insight-director-5756f4f4b4-jw4vv	0/1	
ContainerCreating 0 20m		
proxedo-api-security-transport-director-7d4f7fbdf-sh9kw	0/1	
ContainerCreating 0 20m		

In this example, the core components do not have configuration, as that is to be set on the Web UI, and for this reason they are not in *Running* state in the example.

To get PAS services, network policies, and so on, the relevant part of the command referring to 'pods' needs to be changed to the object type in question.

7.2. Inspecting objects

To get more detailed information about any specific kubernetes object, use the `kubectl describe` command. Selectors can also be used with this command, however it is recommended to use this command with a specific object name.

Based on the previous example where core pods were not in *Running* state, the `kubectl kubectl --namespace=proxedo-api-security describe pod proxedo-api-security-flow-director-7459896d6c-k9ttm` command can be used to find out the reason behind its malfunction.

The output will be similar to the following example:

Example output for inspecting objects

```

Name:          proxedo-api-security-flow-director-7459896d6c-k9ttm
Namespace:     mate
Priority:       0
Node:          api-kube-node-2/10.90.31.63
Start Time:    Mon, 04 Jul 2022 10:40:56 +0200
Labels:        app=proxedo-api-security
               component=core
               pod-template-hash=7459896d6c
               subcomponent=flow-director
Annotations:   <none>
Status:        Pending
IP:            <none>
IPs:           <none>
Controlled By: ReplicaSet/proxedo-api-security-flow-director-7459896d6c
Containers:
  flow-director:
    Container ID:
    Image:        docker.balasys.hu/api-security/flow-director:4.15.0
    Image ID:
    Ports:        1318/TCP, 8080/TCP
    Host Ports:   0/TCP, 0/TCP
    State:        Waiting
      Reason:     ContainerCreating
    Ready:        False
    Restart Count: 0

```

```

Requests:
  cpu:                250m
  ephemeral-storage:  200Mi
  memory:             550Mi
Readiness:            http-get http://:8000/health delay=0s timeout=2s period=10s
#success=1 #failure=1
Startup:              http-get http://:8000/health delay=0s timeout=2s period=1s
#success=1 #failure=30
Environment:
  CONTENT_FILTERING_DIRECTOR_HOSTNAME: proxedo-api-security-content-filtering-
director
  INSIGHT_DIRECTOR_HOSTNAME: proxedo-api-security-insight-director
  SERVICE_ADAPTOR_PORT:      8000
Mounts:
  /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-qbnnv (ro)
Conditions:
  Type                Status
  Initialized          True
  Ready                False
  ContainersReady      False
  PodScheduled         True
Volumes:
  kube-api-access-qbnnv:
    Type:              Projected (a volume that contains injected data from
multiple sources)
    TokenExpirationSeconds: 3607
    ConfigMapName:       kube-root-ca.crt
    ConfigMapOptional:    <nil>
    DownwardAPI:         true
QoS Class:             Burstable
Node-Selectors:         <none>
Tolerations:           node.kubernetes.io/not-ready:NoExecute op=Exists for 300s
                       node.kubernetes.io/unreachable:NoExecute op=Exists for 300s
Events:
  Type    Reason      Age    From          Message
  ----    -
  Normal  Scheduled   37m    default-scheduler  Successfully assigned
mate/proxedo-api-security-flow-director-7459896d6c-k9ttm to api-kube-node-2

```

In this example, the *Events* section of the output shows (among other details) that two necessary configuration objects do not exist, and therefore the pods cannot be started. It also describes the volumes, ports, environment variables and many more attributes that can be helpful for finding out the reason behind its malfunction.

7.3. Checking logs

Logs of PAS components are by default available through the `kubectl logs` command. An extract of the output of `kubectl logs --namespace=proxedo-api-security pods/proxedo-api-security-frontend-84798447c4-srvvj` command is displayed in the following example:

Example output for checking logs

```

2025-11-11T13:34:54 config-webui 10.90.101.13 - "POST /api/v1/auth/login HTTP/1.1" 200
1005 "http://api-kube-node-3.dev.balasys:30001/login" "Mozilla/5.0 (X11; Linux x86_64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /api/v1/ui-adaptor/menu HTTP/1.1"
200 1942 "http://api-kube-node-3.dev.balasys:30001/login" "Mozilla/5.0 (X11; Linux
x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /assets/outline/appstore.svg
HTTP/1.1" 200 574 "http://api-kube-node-3.dev.balasys:30001/" "Mozilla/5.0 (X11; Linux
x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"

```

```

2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /assets/outline/api.svg HTTP/1.1"
200 1134 "http://api-kube-node-3.dev.balasy:30001/" "Mozilla/5.0 (X11; Linux x86_64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET
/assets/images/proxecto_API_transparent.svg HTTP/1.1" 200 3975 "http://api-kube-node-
3.dev.balasy:30001/changes" "Mozilla/5.0 (X11; Linux x86_64) AppleWebKit/537.36 (KHTML,
like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /assets/outline/setting.svg
HTTP/1.1" 200 1873 "http://api-kube-node-3.dev.balasy:30001/" "Mozilla/5.0 (X11; Linux
x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /SourceSansPro-
SemiBold.43cc81b496222dc9ce3c.ttf HTTP/1.1" 200 268280 "http://api-kube-node-
3.dev.balasy:30001/styles.e68c8c26486c2eba6127.css" "Mozilla/5.0 (X11; Linux x86_64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /api/v1/ui-adaptor/config/changes
HTTP/1.1" 200 1969 "http://api-kube-node-3.dev.balasy:30001/changes" "Mozilla/5.0 (X11;
Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"
2025-11-11T13:34:54 config-webui 10.90.101.13 - "GET /assets/outline/rollback.svg
HTTP/1.1" 200 265 "http://api-kube-node-3.dev.balasy:30001/changes" "Mozilla/5.0 (X11;
Linux x86_64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/103.0.5060.53 Safari/537.36"

```

The `kubectl logs` command can also be used with *Selectors* and other object types like deployments or services. In this case, its scope is wider and can sometimes be more adequate.

7.3.1. Understanding logs

As multiple pieces of software run in each container, there are two layers of logs in each containers' output. The first field is always an ISO formatted date. Then the name of the process inside the container follows. The remaining fields are the output of the process itself. In the below example, we see logs from the `flow-director` container. It prints output for processes called `pre`, `event-handler`, `flow-director` and `service-adaptor`.

Container log output

```

2021-04-20T09:15:30 pre Container starts
2021-04-20T09:15:33 pre INFO:confgen: Generating configuration files;
service_name='twisted'
2021-04-20T09:15:34 event-handler INFO:SupervisordEventDispatcher:Dispatching event;
processname='pre', eventname='PROCESS_STATE_EXITED'
2021-04-20T09:15:34 event-handler INFO:SupervisordEventDispatcher:Process exited;
processname=pre, success=True
2021-04-20T09:15:34 event-handler INFO:SupervisordEventDispatcher:Starting main
processes.
2021-04-20T09:15:34 event-handler INFO:SupervisordEventDispatcher:Starting process;
process='flow-director'
[...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_builder.info(3) (nosession): Loaded plugin; [...]
2021-04-20T09:15:37 flow-director flow_set.info(3) (nosession): Start building flows
[...]
2021-04-20T09:15:39 event-handler INFO:SupervisordEventDispatcher:Starting process;
process='service-adaptor'
[...]
2021-05-07T14:23:55 service-adaptor INFO:HealthCheck:All services are healthy; [...]
2021-05-07T14:23:55 service-adaptor Request received; [...]

```


7.3.1.1. Flow Director and Transport Director logs

As from the API security perspective, the most important components are *Flow Director* and *Transport Director*, we discuss their logs more in detail. There are two important concepts related to these logs: categories and Session IDs.

- **Categories** help filtering logs based on their relevance. They are composed of a *component*, a *tag*, and a *severity*, for example: *http.info(3)*.
 - The **component** helps to identify the part of the solution. For the *Transport Director* this is usually *core* or *http*, for the *Flow Director* it is either *core*, or the *Plugin*'s type, such as *serializer* or *enforcer*.
 - The **tag** helps to define the type of the message. Usually one of *info*, *error*, *debug*, *policy* or *accounting*.
 - The **severity** defines how important the message is. It is a number between 1-9 where 1 is the highest.
- **Session ID** helps identifying log lines that belong to the same session. This is especially important as the calls travel between the *Transport Director* and the *Flow Director*.

It is usually in the form of `svc/default/<listener>:<transport-director-session>/default/http#<http-request-count>/flow:<flow-director-id>/ch:<flow-director-channel>/<endpoint_name>/<plugin_type>/<plugin_name>`, for example: `svc/default/httpbin:14/default/http#0/flow:1/ch:28/endpoint_test/enforcer/manualtest`.

Information that is not available at the time, will be missing from the Session ID. Generally, the part until `/flow:` belongs to the *Transport Director*. Consequently, the *Transport Director* will never see that part. The *Flow Director* however will fetch and include that information. Nevertheless, in early phases it might not be available, and the Session ID will start with *flow*.

Despite some parts not being always available, the ID is constructed in such a manner that grepping on any part will find other messages with extra information as well.

7.4. Troubleshooting containers

7.4.1. Inspecting files

For inspecting file content inside PAS containers, it is recommended to log into a specific container and use the shell for navigating in the file system and the available tools for reading them.

Example for inspecting container filesystem

```
kubectl --namespace proxedo-api-security exec --stdin --tty pods/proxedo-api-security-config-api-579b59fc8d-59hqh -- /bin/bash
root@proxedo-api-security-config-api-579b59fc8d-59hqh:/# ls -l /opt/balasys/etc/mgmt/
config.yml
users.htpass
```

7.4.2. Inspecting processes and network

Processes and network can best be troubleshot using the `kubectl debug` tool and the PAS debug image. This image holds basic process and network debugging tools and enables installing new ones.

Example for inspecting processes and network traffic of containers

```
kubectl --namespace proxedo-api-security debug --stdin --tty
--image=docker.balasys.hu/api-security/container-debugger:4.15.0 pods/proxedo-api-security-config-api-579b59fc8d-59hqh -- tcpdump -i any
Defaulting debug container name to debugger-wcrnx.
If you don't see a command prompt, try pressing enter.
[...]
```

7.5. Changing bootstrap configuration

Since bootstrap configuration is provided during *Helm* installation, the parameters used there can be changed in the provided files. Moreover, all the input files may be changed. As soon as the changes are made, they can be made effective by running the installation command, as displayed in [Providing the necessary files for *Helm* installation](#).

7.6. Backup and restore

7.6.1. Bootstrap configuration

As the whole bootstrap configuration is provided at the time of installation, the directory, in which the installation was carried out, needs to be saved, so that the installation procedure can be repeated.

7.6.2. Running configuration

To completely backup the running configuration, the storage component's *Persistent Volume* needs to be backed up. This can be done by directly backing up the *Persistent Volume* that is assigned to the `proxedo-api-security-storage` *Persistent Volume Claim*. This solution is specific to the Kubernetes Cluster and therefore it is the responsibility of the cluster administrator. In this case, the cluster administrator also needs to make sure that the restored *Persistent Volume* gets assigned to the new *Persistent Volume Claim* from the new PAS installation.

Another method for creating a backup of the running configuration is to use the backup mechanism available on the Web UI, see [Backup and restore running or user configuration for Proxedo API Security](#).

7.7. Factory reset

In case a factory reset is necessary, the simplest solution is to delete the namespace, PAS is installed in. If that is not feasible, an alternative is to explicitly delete Kubernetes objects related to PAS. To do so, two main steps are required:

1. Uninstall the PAS *Helm* chart using the `helm uninstall --namespace=proxedo-api-security proxedo-api-security` command. This will remove all kubernetes objects managed by the *Helm* charts, including the *Persistent Volume Claim* associated with the storage components.
2. Delete the core configuration objects. These objects are not managed by the *Helm* chart but by the management component. To complete this, run the following command:
 - `kubectl --namespace=proxedo-api-security delete secrets proxedo-api-security-registry-credentials`

Following these steps, PAS shall be installed from scratch. For more details, see [Installation of Proxedo API Security in Kubernetes environment](#).

Appendix A: Selector configuration for the Fraud Detector Plugin

The following fields can be defined in the *Save As Key* field when creating a new *Selector*. The saved *Selector* can be used by the *Fraud Detector* plugin.



The data type selected in the API for the actual selector option shall be the one listed in this table as *Type* for the actual selector. Currently, no data type conversion is possible for selectors.

Table 75. Selector configuration for the Fraud Detector Plugin

Values for Save As Key field	Data type	Description	Example
action_type	string	The type of the user action being scored. Any string can be valid.	update_content, verification or account_login_fail
client_address	string	The user's IP address at the time of the transaction. It shall include the full IPv4 or IPv6 address.	
transaction_id	string	A unique identifier for the transaction, as found in the system. If it is not specified, it is automatically generated.	98db9a56b2e3
affiliate_id	string	The user's unique affiliate identifier in the system.	
affiliate_name	string	The name of the affiliate for the registered user. Can be ASCII-encoded via a secure hash algorithm, such as MD5 or SHA-2.	jdoe345
order_memo	string	The description of the transaction found in the system.	
email	string	The full email address of the registered user.	
email_domain	string	The email address domain of the registered user.	
password_hash	string	The hash of the user's password in ASCII encoding (we recommend using HMAC-SHA256 or RSA-SHA256).	
user_fullname	string	The user's registered full name. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	John Doe
user_name	string	The user's registered username. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	jdoe325
user_id	string	The user's unique identifier. If the request was sent without a user_id value, a unique ID is automatically generated based on the user_name and/or the email fields, based on which is available. If none of these identifiers were included in the request, the user ID is generated randomly.	00ab11-as2233
user_created	integer	The date when the user first registered to the protected service, using the UNIX time format and UTC time zone, without milliseconds.	1446370717 (Sun, 01 Nov 2015 09:38:37 +0000)
user_category	string	The user's category.	VIP
user_account_status	string	The user's current account status.	login_blocked
user_bank_account	string	The user's bank account number for monetary transfer.	IBAN number
user_bank_name	string	The name of the user's bank account.	
user_balance	float	The user's current balance.	1010.25

Values for <i>Save As Key</i> field	Data type	Description	Example
user_verification_level	string	The user's verification level.	ID_verified
user_dob	date	The user's date of birth in the format of YYYY-MM-DD.	1983-01-01
user_country	string	The country code for the user's registered address. Uses the two-character ISO 3166-1 format.	US, DE
user_city	string	The complete name of the city associated with the user's registered address.	London, New York
user_region	string	The state or region code for the user's registered address. Uses the two-character ISO 3166-2 format.	NY, DE
user_zip	string	The zip/postal code of the user's registered address.	10005, PH1 1EU
user_street	string	The first line of the user's registered street address. Can be hashed in ASCII encoding as well.	MD5, SHA-2 family: 157 W 26th St
user_street2	string	The second line of the user's registered street address. Can be hashed in ASCII encoding as well.	MD5, SHA-2 family: Apt.432
session_id	string	The session ID is a custom, unique ID that links the user's device data with the transactions. It shall be based on the user's current browsing session, by tracking cookies for example. If JavaScript Agent v4 is used, the encrypted payload returned by the SDK (supported by JS Agent v4, iOS SDK 3.0.0, Android SDK 3.0.0) shall be sent in the session field, instead of the session_id.	
session	string	The base64 encoded session data returned by the SDKs.	
device_id	string	This field shall only be used if a device fingerprinting solution is used already. This is the ID that shall be linked to the transactions or in case rules are required to be built on those IDs.	
payment_mode	string	The method of payment used.	card, paypal, wire transfer, bitcoin
payment_provider	string	The name of the payment service provider related to the transaction.	skrill
card_fullname	string	The user's full name found on the card. Can be hashed in ASCII encoding as well.	MD5, SHA-2 family
card_bin	string	The first 4, 6 or 8 digits of the card number.	

Values for <i>Save As Key</i> field	Data type	Description	Example
card_hash	string	The hash of the credit card used by the user in ASCII encoding. We recommend using HMAC-SHA256 or RSA-SHA256 formats and strictly advise not to use MD5 hash format.	
card_expire	string	The card's expiration date.	2022-01
card_last	string	The last 4 digits of the card number. These help to identify the card.	
avs_result	string	The standard Address verification Service (AVS) codes sent by the credit card processor.	N, A
cvv_result	boolean	The Card Verification Value (CVV) result.	true, false
status_3d	string	The Card Verification Value (CVV) result.	true, false
sca_method	string	The result of the Strong Customer Authentication method.	2FA
phone_number	string	The user's registered phone number, including the country code. Cannot include spaces or hyphens, the + sign is optional. The maximum length is 19 characters.	36704316088
transaction_type	string	The transaction type of the actual business.	purchase, return
transaction_amount	float	The full transaction amount. As a decimal point use '.' (full stop).	539.99
transaction_currency	string	The currency used by the user, in ISO 4217 format. Crypto currencies are also supported.	EUR, BTC, USDT
shipping_country	string	A two-character ISO 3166-1 country code for the country associated with the user's shipping address.	US, DE
shipping_city	string	The full name of the city associated with the user's shipping address.	London, New York
shipping_region	string	The state or region code for the user's shipping address. Uses the two-character ISO 3166-2 format	NY, DE
shipping_zip	string	The zip/postal code of the user's shipping address.	10005, PH1 1EU
shipping_street	string	The first line of the user's shipping street address. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	157 W 26th St
shipping_street2	string	The second line of the user's shipping street address. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	Apt.432
shipping_phone	string	The phone number associated with the user's shipping address, including the country code. Cannot include spaces or hyphens, the + sign is optional. The maximum length is 19 characters.	36704316088

Values for <i>Save As Key</i> field	Data type	Description	Example
shipping_fullname	string	The user's registered full name. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	John Doe
shipping_method	string	The type of the shipping method used by the customer.	standard, UPS, FedEx
billing_country	string	The country code for the user's billing address. Uses the two-character ISO 3166-1 format.	US, DE
billing_city	string	The full name of the city associated with the user's billing address.	London, New York
billing_region	string	The state or region code for the user's billing address. Uses the two-character ISO 3166-2 format	NY, DE
billing_zip	string	The zip/postal code of the user's billing address.	10005, PH1 1EU
billing_street	string	The user's billing street address line 1. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	157 W 26th St
billing_street2	string	The user's billing street address line 2. Can be hashed in ASCII encoding as well (e.g. MD5, SHA-2 family).	Apt.432
billing_phone	string	The phone number associated with the user's billing address, including the country code. Cannot include spaces or hyphens, the + sign is optional. The maximum length is 19 characters.	36704316088
discount_code	string	The discount code that the user applied during the checkout.	
gift	boolean	The user can mark the order with true or false value, dependent on if it is a gift or not.	
gift_message	boolean	The user can mark the order with true or false value, dependent on if the order has a gift message or not.	
merchant_category	string	The category of the merchant.	digital_item_seller
merchant_id	string	The unique merchant identifier in case the orders are from different merchants.	ab01-cd23-4567
merchant_created_at	integer	The date the merchant was created, using the UNIX time format and UTC time zone.	1446370717 (Sun, 01 Nov 2015 09:38:37 +0000)
merchant_country	string	The country code for the merchant's address. Uses the two-character ISO 3166-1 format.	US, DE
receiver_fullname	string	The receiver's full name for monetary transfer.	IBAN number
details_url	string	The URL of the transaction in the management platform.	

Values for <i>Save As Key</i> field	Data type	Description	Example
regulation	string	The license or market name for gambling operator.	MGA
bonus_campaign_id	string	The bonus campaign's unique identifier.	bonus100a
brand_id	string	The brand's unique identifier.	brand123



The maximum length of all request parameters is 100 characters, except for the following: **500 characters for card_hash** 64 characters for the session_id (sent directly or within the session field) **19 characters for the phone_number** 15 characters for card_bin **4 characters for transaction_currency** 50 characters for discount_code and shipping_method **** 255 characters for transaction_id**

Appendix B: Time zones

Country Code	Time zone Name
AD	Europe/Andorra
AE	Asia/Dubai
AF	Asia/Kabul
AG	America/Antigua
AI	America/Anguilla
AL	Europe/Tirane
AM	Asia/Yerevan
AO	Africa/Luanda
AQ	Antarctica/McMurdo
AQ	Antarctica/Casey
AQ	Antarctica/Davis
AQ	Antarctica/DumontDUrville
AQ	Antarctica/Mawson
AQ	Antarctica/Palmer
AQ	Antarctica/Rothera
AQ	Antarctica/Syowa
AQ	Antarctica/Troll
AQ	Antarctica/Vostok
AR	America/Argentina/Buenos_Aires
AR	America/Argentina/Cordoba
AR	America/Argentina/Salta

Country Code	Time zone Name
AR	America/Argentina/Jujuy
AR	America/Argentina/Tucuman
AR	America/Argentina/Catamarca
AR	America/Argentina/La_Rioja
AR	America/Argentina/San_Juan
AR	America/Argentina/Mendoza
AR	America/Argentina/San_Luis
AR	America/Argentina/Rio_Gallegos
AR	America/Argentina/Ushuaia
AS	Pacific/Pago_Pago
AT	Europe/Vienna
AU	Australia/Lord_Howe
AU	Antarctica/Macquarie
AU	Australia/Hobart
AU	Australia/Currie
AU	Australia/Melbourne
AU	Australia/Sydney
AU	Australia/Broken_Hill
AU	Australia/Brisbane
AU	Australia/Lindeman
AU	Australia/Adelaide
AU	Australia/Darwin
AU	Australia/Perth
AU	Australia/Eucla
AW	America/Aruba
AX	Europe/Mariehamn
AZ	Asia/Baku
BA	Europe/Sarajevo
BB	America/Barbados
BD	Asia/Dhaka
BE	Europe/Brussels
BF	Africa/Ouagadougou
BG	Europe/Sofia
BH	Asia/Bahrain

Country Code	Time zone Name
BI	Africa/Bujumbura
BJ	Africa/Porto-Novo
BL	America/St_Barthelemy
BM	Atlantic/Bermuda
BN	Asia/Brunei
BO	America/La_Paz
BQ	America/Kralendijk
BR	America/Noronha
BR	America/Belem
BR	America/Fortaleza
BR	America/Recife
BR	America/Araguaina
BR	America/Maceio
BR	America/Bahia
BR	America/Sao_Paulo
BR	America/Campo_Grande
BR	America/Cuiaba
BR	America/Santarem
BR	America/Porto_Velho
BR	America/Boa_Vista
BR	America/Manaus
BR	America/Eirunepe
BR	America/Rio_Branco
BS	America/Nassau
BT	Asia/Thimphu
BW	Africa/Gaborone
BY	Europe/Minsk
BZ	America/Belize
CA	America/St_Johns
CA	America/Halifax
CA	America/Glace_Bay
CA	America/Moncton
CA	America/Goose_Bay
CA	America/Blanc-Sablon

Country Code	Time zone Name
CA	America/Toronto
CA	America/Nipigon
CA	America/Thunder_Bay
CA	America/Iqaluit
CA	America/Pangnirtung
CA	America/Atikokan
CA	America/Winnipeg
CA	America/Rainy_River
CA	America/Resolute
CA	America/Rankin_Inlet
CA	America/Regina
CA	America/Swift_Current
CA	America/Edmonton
CA	America/Cambridge_Bay
CA	America/Yellowknife
CA	America/Inuvik
CA	America/Creston
CA	America/Dawson_Creek
CA	America/Fort_Nelson
CA	America/Vancouver
CA	America/Whitehorse
CA	America/Dawson
CC	Indian/Cocos
CD	Africa/Kinshasa
CD	Africa/Lubumbashi
CF	Africa/Bangui
CG	Africa/Brazzaville
CH	Europe/Zurich
CI	Africa/Abidjan
CK	Pacific/Rarotonga
CL	America/Santiago
CL	America/Punta_Arenas
CL	Pacific/Easter
CM	Africa/Douala

Country Code	Time zone Name
CN	Asia/Shanghai
CN	Asia/Urumqi
CO	America/Bogota
CR	America/Costa_Rica
CU	America/Havana
CV	Atlantic/Cape_Verde
CW	America/Curacao
CX	Indian/Christmas
CY	Asia/Nicosia
CY	Asia/Famagusta
CZ	Europe/Prague
DE	Europe/Berlin
DE	Europe/Busingen
DJ	Africa/Djibouti
DK	Europe/Copenhagen
DM	America/Dominica
DO	America/Santo_Domingo
DZ	Africa/Algiers
EC	America/Guayaquil
EC	Pacific/Galapagos
EE	Europe/Tallinn
EG	Africa/Cairo
EH	Africa/El_Aaiun
ER	Africa/Asmara
ES	Europe/Madrid
ES	Africa/Ceuta
ES	Atlantic/Canary
ET	Africa/Addis_Ababa
FI	Europe/Helsinki
FJ	Pacific/Fiji
FK	Atlantic/Stanley
FM	Pacific/Chuuk
FM	Pacific/Pohnpei
FM	Pacific/Kosrae

Country Code	Time zone Name
FO	Atlantic/Faroe
FR	Europe/Paris
GA	Africa/Libreville
GB	Europe/London
GD	America/Grenada
GE	Asia/Tbilisi
GF	America/Cayenne
GG	Europe/Guernsey
GH	Africa/Accra
GI	Europe/Gibraltar
GL	America/Godthab
GL	America/Danmarkshavn
GL	America/Scoresbysund
GL	America/Thule
GM	Africa/Banjul
GN	Africa/Conakry
GP	America/Guadeloupe
GQ	Africa/Malabo
GR	Europe/Athens
GS	Atlantic/South_Georgia
GT	America/Guatemala
GU	Pacific/Guam
GW	Africa/Bissau
GY	America/Guyana
HK	Asia/Hong_Kong
HN	America/Tegucigalpa
HR	Europe/Zagreb
HT	America/Port-au-Prince
HU	Europe/Budapest
ID	Asia/Jakarta
ID	Asia/Pontianak
ID	Asia/Makassar
ID	Asia/Jayapura
IE	Europe/Dublin

Country Code	Time zone Name
IL	Asia/Jerusalem
IM	Europe/Isle_of_Man
IN	Asia/Kolkata
IO	Indian/Chagos
IQ	Asia/Baghdad
IR	Asia/Tehran
IS	Atlantic/Reykjavik
IT	Europe/Rome
JE	Europe/Jersey
JM	America/Jamaica
JO	Asia/Amman
JP	Asia/Tokyo
KE	Africa/Nairobi
KG	Asia/Bishkek
KH	Asia/Phnom_Penh
KI	Pacific/Tarawa
KI	Pacific/Enderbury
KI	Pacific/Kiritimati
KM	Indian/Comoro
KN	America/St_Kitts
KP	Asia/Pyongyang
KR	Asia/Seoul
KW	Asia/Kuwait
KY	America/Cayman
KZ	Asia/Almaty
KZ	Asia/Qyzylorda
KZ	Asia/Qostanay
KZ	Asia/Aqtobe
KZ	Asia/Aqtau
KZ	Asia/Atyrau
KZ	Asia/Oral
LA	Asia/Vientiane
LB	Asia/Beirut
LC	America/St_Lucia

Country Code	Time zone Name
LI	Europe/Vaduz
LK	Asia/Colombo
LR	Africa/Monrovia
LS	Africa/Maseru
LT	Europe/Vilnius
LU	Europe/Luxembourg
LV	Europe/Riga
LY	Africa/Tripoli
MA	Africa/Casablanca
MC	Europe/Monaco
MD	Europe/Chisinau
ME	Europe/Podgorica
MF	America/Marigot
MG	Indian/Antananarivo
MH	Pacific/Majuro
MH	Pacific/Kwajalein
MK	Europe/Skopje
ML	Africa/Bamako
MM	Asia/Yangon
MN	Asia/Ulaanbaatar
MN	Asia/Hovd
MN	Asia/Choibalsan
MO	Asia/Macau
MP	Pacific/Saipan
MQ	America/Martinique
MR	Africa/Nouakchott
MS	America/Montserrat
MT	Europe/Malta
MU	Indian/Mauritius
MV	Indian/Maldives
MW	Africa/Blantyre
MX	America/Mexico_City
MX	America/Cancun
MX	America/Merida

Country Code	Time zone Name
MX	America/Monterrey
MX	America/Matamoros
MX	America/Mazatlan
MX	America/Chihuahua
MX	America/Ojinaga
MX	America/Hermosillo
MX	America/Tijuana
MX	America/Bahia_Banderas
MY	Asia/Kuala_Lumpur
MY	Asia/Kuching
MZ	Africa/Maputo
NA	Africa/Windhoek
NC	Pacific/Noumea
NE	Africa/Niamey
NF	Pacific/Norfolk
NG	Africa/Lagos
NI	America/Managua
NL	Europe/Amsterdam
NO	Europe/Oslo
NP	Asia/Kathmandu
NR	Pacific/Nauru
NU	Pacific/Niue
NZ	Pacific/Auckland
NZ	Pacific/Chatham
OM	Asia/Muscat
PA	America/Panama
PE	America/Lima
PF	Pacific/Tahiti
PF	Pacific/Marquesas
PF	Pacific/Gambier
PG	Pacific/Port_Moresby
PG	Pacific/Bougainville
PH	Asia/Manila
PK	Asia/Karachi

Country Code	Time zone Name
PL	Europe/Warsaw
PM	America/Miquelon
PN	Pacific/Pitcairn
PR	America/Puerto_Rico
PS	Asia/Gaza
PS	Asia/Hebron
PT	Europe/Lisbon
PT	Atlantic/Madeira
PT	Atlantic/Azores
PW	Pacific/Palau
PY	America/Asuncion
QA	Asia/Qatar
RE	Indian/Reunion
RO	Europe/Bucharest
RS	Europe/Belgrade
RU	Europe/Kaliningrad
RU	Europe/Moscow
UA	Europe/Simferopol
RU	Europe/Kirov
RU	Europe/Astrakhan
RU	Europe/Volgograd
RU	Europe/Saratov
RU	Europe/Ulyanovsk
RU	Europe/Samara
RU	Asia/Yekaterinburg
RU	Asia/Omsk
RU	Asia/Novosibirsk
RU	Asia/Barnaul
RU	Asia/Tomsk
RU	Asia/Novokuznetsk
RU	Asia/Krasnoyarsk
RU	Asia/Irkutsk
RU	Asia/Chita
RU	Asia/Yakutsk

Country Code	Time zone Name
RU	Asia/Khandyga
RU	Asia/Vladivostok
RU	Asia/Ust-Nera
RU	Asia/Magadan
RU	Asia/Sakhalin
RU	Asia/Srednekolymsk
RU	Asia/Kamchatka
RU	Asia/Anadyr
RW	Africa/Kigali
SA	Asia/Riyadh
SB	Pacific/Guadalcanal
SC	Indian/Mahe
SD	Africa/Khartoum
SE	Europe/Stockholm
SG	Asia/Singapore
SH	Atlantic/St_Helena
SI	Europe/Ljubljana
SJ	Arctic/Longyearbyen
SK	Europe/Bratislava
SL	Africa/Freetown
SM	Europe/San_Marino
SN	Africa/Dakar
SO	Africa/Mogadishu
SR	America/Paramaribo
SS	Africa/Juba
ST	Africa/Sao_Tome
SV	America/El_Salvador
SX	America/Lower_Princes
SY	Asia/Damascus
SZ	Africa/Mbabane
TC	America/Grand_Turk
TD	Africa/Ndjamena
TF	Indian/Kerguelen
TG	Africa/Lome

Country Code	Time zone Name
TH	Asia/Bangkok
TJ	Asia/Dushanbe
TK	Pacific/Fakaofo
TL	Asia/Dili
TM	Asia/Ashgabat
TN	Africa/Tunis
TO	Pacific/Tongatapu
TR	Europe/Istanbul
TT	America/Port_of_Spain
TV	Pacific/Funafuti
TW	Asia/Taipei
TZ	Africa/Dar_es_Salaam
UA	Europe/Kiev
UA	Europe/Uzhgorod
UA	Europe/Zaporozhye
UG	Africa/Kampala
UM	Pacific/Midway
UM	Pacific/Wake
US	America/New_York
US	America/Detroit
US	America/Kentucky/Louisville
US	America/Kentucky/Monticello
US	America/Indiana/Indianapolis
US	America/Indiana/Vincennes
US	America/Indiana/Winamac
US	America/Indiana/Marengo
US	America/Indiana/Petersburg
US	America/Indiana/Vevay
US	America/Chicago
US	America/Indiana/Tell_City
US	America/Indiana/Knox
US	America/Menominee
US	America/North_Dakota/Center
US	America/North_Dakota/New_Salem

Country Code	Time zone Name
US	America/North_Dakota/Beulah
US	America/Denver
US	America/Boise
US	America/Phoenix
US	America/Los_Angeles
US	America/Anchorage
US	America/Juneau
US	America/Sitka
US	America/Metlakatla
US	America/Yakutat
US	America/Nome
US	America/Adak
US	Pacific/Honolulu
UY	America/Montevideo
UZ	Asia/Samarkand
UZ	Asia/Tashkent
VA	Europe/Vatican
VC	America/St_Vincent
VE	America/Caracas
VG	America/Tortola
VI	America/St_Thomas
VN	Asia/Ho_Chi_Minh
VU	Pacific/Efate
WF	Pacific/Wallis
WS	Pacific/Apia
YE	Asia/Aden
YT	Indian/Mayotte
ZA	Africa/Johannesburg
ZM	Africa/Lusaka
ZW	Africa/Harare

Appendix C: values.yml examples

C.1. Minimal configuration

The configuration example is set as follows:

- Default TLS settings are used for storage-storage configuration
- Certificates and encryption key are generated by openssl commands
- INFO log level is defined
- If the parameters for the management configuration are not defined, the default values will be used.

Example values.yml file

```
config:
  storage:
    consul:
      gossip_encryption_key: MhstT80sqle63WC7kn0ak+c7GfK7k50Y2n/4Qk/fSXs=

  blob_store:
    access_key: your_access_key
    secret_key: your_secret_key
    rpc_secret: your_rpc_secret
    admin_token: your_admin_token
```

C.2. Management configuration with LDAP authentication

The configuration examples are set as follows:

- LDAP authentication is configured without TLS.
- The authentication configuration was tested using Microsoft Active Directory.

Example values.yml with NTLM on

```
config:
  mgmt:
    configapi:
      ldap:
        ldap_url: ldap://ad.example.com
        use_ntlm: on
        bind_user: AD_domain\administrator # The name of the user follows the domain.
        bind_password: your_administrator_password
        user_base_dn: CN=Users,DC=example,DC=com
        group_base_dn: CN=Users,CN=Builtin,DC=example,DC=com
        allowed_groups:
          - Users

  storage:
    consul:
      gossip_encryption_key: MhstT80sqle63WC7kn0ak+c7GfK7k50Y2n/4Qk/fSXs=

  blob_store:
    access_key: your_access_key
    secret_key: your_secret_key
    rpc_secret: your_rpc_secret
    admin_token: your_admin_token
```

Example values.yml with NTLM off

```
config:
  mgmt:
```

```
configapi:
  ldap:
    ldap_url: ldap://ad.example.com
    use_ntlm: off
    bind_user: CN=administrator,CN=Users,DC=example,DC=com # This must be the DN of
the user
    bind_password: your_administrator_password
    user_base_dn: CN=Users,DC=example,DC=com
    group_base_dn: CN=Users,CN=Builtin,DC=example,DC=com
    allowed_groups:
      - Users

  storage:
    consul:
      gossip_encryption_key: MhstT80sqle63WC7kn0ak+c7GfK7k50Y2n/4Qk/fSXs=

  blob_store:
    access_key: your_access_key
    secret_key: your_secret_key
    rpc_secret: your_rpc_secret
    admin_token: your_admin_token
```

Appendix D: LDAP certificate examples

Single CA file example

```
-----BEGIN CERTIFICATE-----
... (the certificate for the CA)...
-----END CERTIFICATE-----
```

Example on certificate chain with multiple CAs

```
-----BEGIN CERTIFICATE-----
... (the certificate for the CA)...
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
... (the root certificate for the CA's issuer)...
-----END CERTIFICATE-----
```

Glossary

<i>API</i>	Application Programming Interface
<i>CA</i>	Certification Authority
<i>CRL</i>	Certificate Revocation List
<i>HTTP</i>	HyperText Transport Protocol
<i>HTTPS</i>	HyperText Transport Protocol Secure
<i>JSON</i>	JavaScript Object Notation
<i>LDAP</i>	Lightweight Directory Access Protocol
<i>MIB</i>	Management Information Base

<i>NTLM</i>	NT LAN Manager
<i>PEM</i>	Privacy Enhanced Mail
<i>SN1</i>	Server Name Indication
<i>SNMP</i>	Simple Network Management Protocol
<i>SOAP</i>	Simple Object Access Protocol
<i>SSL</i>	Secure Socket Layer
<i>SIEM</i>	Security Information and Event Management
<i>TLS</i>	Transport Layer Security
<i>URI</i>	Universal Resource Indicator
<i>URL</i>	Universal Resource Locator
<i>WSDL</i>	Web Service Definition Language
<i>XML</i>	Extensible Markup Language
<i>XSD</i>	XML Schema Definition