CYTOMIC

Cytomic SIEMConnect Infrastructure Guide_

Legal notice.

Registered trademarks.

Contact information.

Corporate Headquarters:

Cytomic (Business Unit of Panda Security S.L.)

Calle Santiago de Compostela 12

Bilbao (Bizkaia) 48003 Spain.

https://www.pandasecurity.com/uk/about/contact/

Version: 3.00.00-02

Author: Cytomic

Date: 14/03/2024

About the Cytomic SIEMConnect Infraestructure guide

You can find the most recent version of this guide at:

https://info.cytomicmodel.com/resources/guides/SIEMConnect/en/SIEMCONNECT-Manual-EN.pdf

Downloading the Cytomic Importer software

To get the Cytomic Importer installation package, refer to the following URL: **https://www.pandasecurity.com/en-us/support/card?id=950031**

Cytomic SIEMConnect events guide

https://info.cytomicmodel.com/resources/guides/SIEMConnect/en/SIEMCONNECT-EventDescriptionGuide-EN.pdf

Cytomic EDR and Cytomic EDPR guides

https://info.cytomicmodel.com/resources/guides/EPDR/latest/en/EPDR-guide-EN.pdf

https://info.cytomicmodel.com/resources/guides/EDR/latest/en/EDR-guide-EN.pdf

Technical information about the modules and services compatible with Cytomic SIEMConnect

To access the Cytomic Insights User's Guide, go to the following URL:

https://info.cytomicmodel.com/resources/guides/Insights/en/INSIGHTS-guide-EN.pdf

CYTOMIC Nexus Guide

- You can find the most recent version of this guide at:

http://nexus-documents.cytomic.ai/AdvancedGuide/Nexus-Manual-EN.pdf

- For more information about a specific topic, please refer to the product's online help, available at:

https://documents.managedprotection.pandasecurity.com/Help/v77000/Partners/en-us/Content/index.htmhttps://nexus-documents.cytomic.ai/Help/v77000/Partners/en-us/Content/index.htm

Technical Support

Cytomic provides global support services aimed at responding to specific questions regarding the operation of the company's products. The technical support team also generates documentation covering technical aspects of our products. This documentation is available in the eKnowledge Base portal.

- To access specific information about the product, please go to the following URL:

https://www.pandasecurity.com/en/support/siemfeeder/

Survey on the Cytomic SIEMConnect Infraestructure guide

Rate this guide and send us suggestions and requests for future versions of our documentation:

https://es.surveymonkey.com/r/feedbackDWGuideEN

Table of contents

Chapter 1

Preface

This guide provides the information and procedures necessary for the implementation of the Cytomic SIEMConnect and Panda SIEMConnect for Partners services.

CHAPTER CONTENTS

Intended audience

This document is intended for:

- The technical staff responsible for managing the IT systems of companies that purchased the Cytomic Cytomic SIEMConnect service.
- The technical staff of the managed security service provider (MSSP) that purchased the Cytomic Panda SIEMConnect for Partners service.

Compatible security products

Cytomic SIEMConnect and Panda SIEMConnect for Partners require that one of the following products be installed on protected computers:

- Cytomic EDR
- Cytomic EDPR

The procedures and instructions in this guide apply equally to all of the aforementioned products. Also, the term "Cytomic EDR" is used generically to refer to all of them, as there is no difference among them with regard to the service.

Document structure

The information in this guide is divided into three sections, each of which is intended for different areas/ technical profiles within the IT department of a company or MSSP:

- **Architecture information** (chapters 2 and 3): intended for systems architects who need to have global visibility into the service to assess the impact of any changes made to the organization's IT

infrastructure and generate management and recovery procedures.

- **Service requirements information** (chapter 4): intended for system administrators who need to provision the resources needed for the service to work correctly.
- **Service deployment information** (chapters 5 and 6): intended for IT security specialists who configure the network access required to enable integration of the service into the company's or the MSSP's SIEM server.

Icons

The following icons are used in this guide:

Explanations and additional information, such as an alternate method for performing a certain task.

Suggestions and recommendations.

Important advice regarding the correct use of options available in Cytomic SIEMConnect or Panda SIEMConnect for Partners.

Refer to other chapters or sections in the guide for more information.

Chapter 2

Cytomic SIEMConnect architecture

Cytomic SIEMConnect is Cytomic's service for delivering the information and knowledge generated by the Cytomic EDR products to end customers' SIEM platforms. Cytomic SIEMConnect helps administrators uncover unknown threats, targeted attacks, and advanced malware (APT, Advanced Persistent Threats), providing deeper visibility into the activity of the processes run across organizations' IT structures.

*For more information about the equivalent solution to Cytomic SIEMConnect for security service providers, Panda SIEMConnect for Partners, refer to "**Panda SIEMConnect for Partners architecture**" on page 15.*

CHAPTER CONTENTS

Main objectives of the Cytomic SIEMConnect service

The principal objective of Cytomic SIEMConnect is to act as a link between the protection software installed on the computers on the network and the company's SIEM server, following the information flow below:

- Cytomic EDR's continuous monitoring sends the Cytomic cloud the telemetry generated by the applications run on a customer's systems.
- Cytomic SIEMConnect enriches this information with the security intelligence generated by Cytomic.
- Cytomic Importer recovers the enriched information from the Azure infrastructure assigned to the customer and sends it directly to the SIEM server or to one of the supported platforms (Kafka and

Syslog), for it to be leveraged later.

Enriching the monitored activity

Cytomic EDR monitors the actions executed by processes on computers. These actions are sent to the Cytomic cloud platform, where they are analyzed using machine learning techniques on a big data infrastructure in order to extract advanced security intelligence. This information enables Cytomic to classify each and every process run by customers, with 99.999% accuracy and near-zero false positive and false negative errors.

Cytomic SIEMConnect combines the information collected from the events monitored by Cytomic EDR and the security data generated, creating a single data flow compatible with the customer's SIEM server.

In order to take full advantage of Cytomic SIEMConnect, it is not necessary to make any changes to users' computer settings: the service operates within the Cytomic infrastructure, receiving data from every computer on the customer's IT systems. This data is normalized, enriched, and sent to the customer's designated SIEM for exploitation.

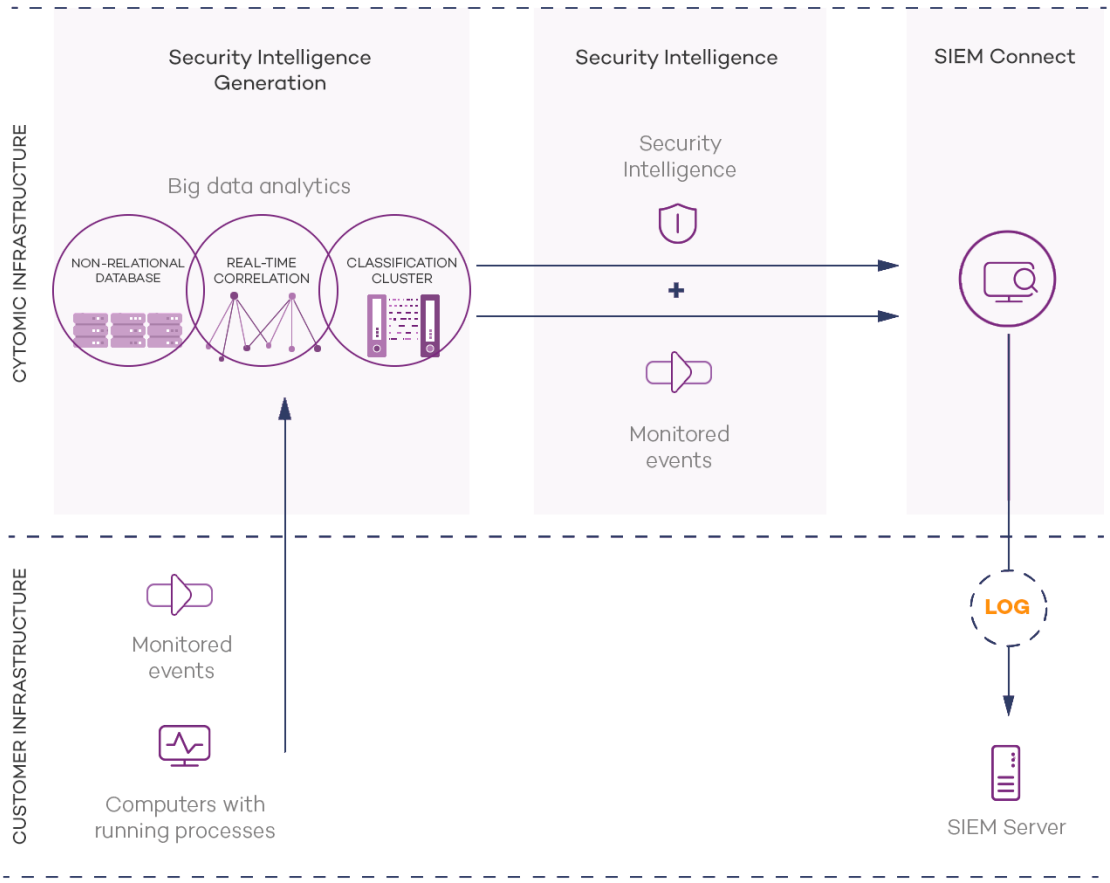

Figure 2.1: Information flow generated by Cytomic EDR and Cytomic SIEMConnect

Main benefits of the Cytomic SIEMConnect service

Cytomic SIEMConnect provides customers' SIEM solutions with information about the activities performed by the processes run on their networks. This information enables administrators to:

- **Obtain visual information about the malware detected on the network**, whether it was run or not, the infection vector, and the actions taken by processes. This information helps administrators make decisions with regard to defining remedial actions and adjusting security policies.
- **View the actions executed by each process**, whether goodware, malware, or temporarily unknown, in order to detect the suspicious activity of recent programs. Cytomic SIEMConnect compiles indicators that can be used to reach conclusions about their potential threat.
- **Monitor attempts to access confidential information**, preventing data leakage and theft. The service displays the Office files, databases, and other repositories of confidential information accessed by malware.
- **View network connections made by processes** in order to identify suspicious or potentially dangerous connections that could be used to steal data.
- **Locate all executed programs**, especially those with known vulnerabilities installed on users' computers, in order to design plans for updating software and adjust security policies.

General architecture

Cytomic SIEMConnect consists of the modules displayed in the following diagram:

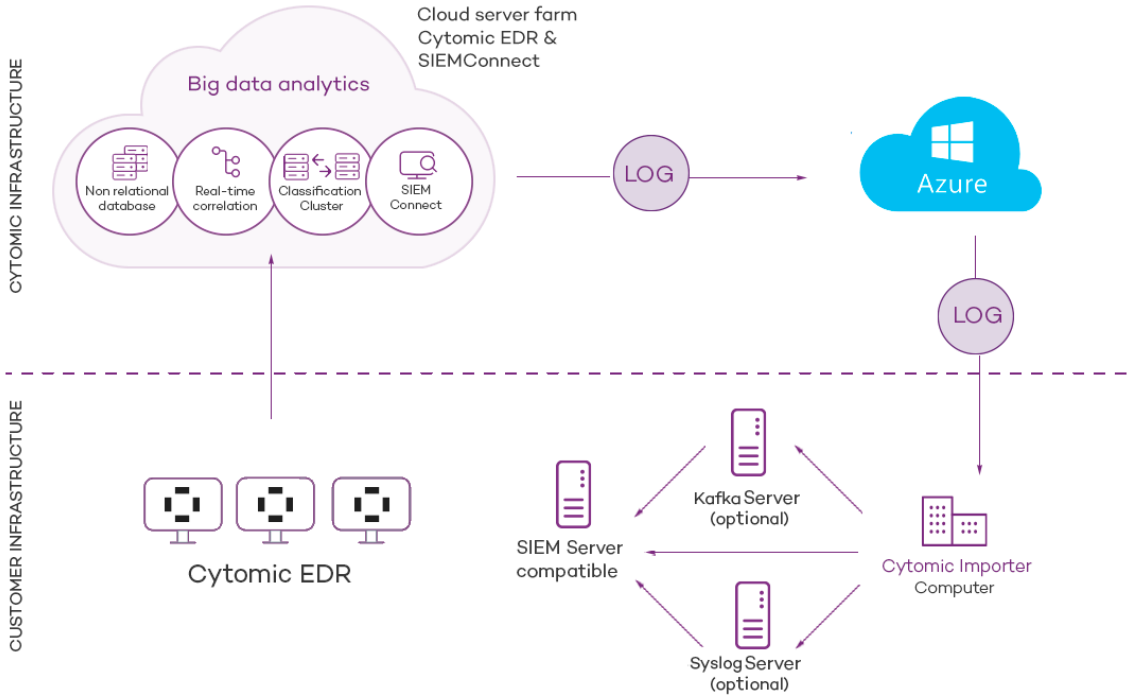

Figure 2.2: Logical diagram of the modules that make up Cytomic SIEMConnect and their relationships

The figure shows the following items:

- Computers on the customer's network protected by Cytomic EDR or Cytomic EDPR.
- Cytomic **cloud**: this stores information from the processes run and analyzes it to extract security intelligence.
- Cytomic SIEMConnect **service:** this collects events and security intelligence information and encapsulates it in the form of logs which are later sent to the customer.
- **Azure infrastructure**: this receives the logs from the Cytomic SIEMConnect service and stores them temporarily until they are downloaded by the customer.
- **Cytomic Importer computer**: a computer on the customer's network that runs the Cytomic Importer process and checks if there are new logs available on the Azure infrastructure to download and store.
- **Kafka server (optional)**: the computer on the customer's network that manages the queues of logs received by Cytomic Importer and sends them to the company's SIEM server.
- **Syslog server (optional)**: the computer on the customer's network that collects the logs received by Cytomic Importer and sends them to the company's SIEM server.
- **SIEM server**: installed on the customer's premises, this collects the data downloaded by the Cytomic Importer computer in order to generate dashboards that help detect suspicious processes that may pose a threat to corporate security.
- **Local and perimeter firewalls**: these protect inbound and outbound data traffic between the Cytomic Importer computer and the Azure infrastructure.

Benefits of the Azure infrastructure

Cytomic SIEMConnect asynchronously generates logs with the information collected from the protected computers, and stores them temporarily until they are downloaded and integrated by the customer into their SIEM system. To do that, Cytomic makes use of cloud-based services hosted on the Azure infrastructure. These services have the following characteristics:

- **Cloud-based storage**: high availability service, meaning the service is available 24 hours a day, from anywhere in the world. Each customer is assigned 80 GB and the information is stored for a maximum of 7 days.
- **Encrypted communications**: the information exchanged between the Cytomic Importer computer and the Azure platform is encrypted with the SSL cryptographic protocol.
- **Authenticated communications**: in order to manage the Cytomic Importerauthentication and authorization processes, two independent tokens are used to negotiate the shared key required to access the Cytomic SIEMConnectplatform and download the customer's network information. These tokens have different expiration times in order to ensure data confidentiality.

*For more information, refer to "**Appendix 2: Security Architecture**" on page **55**.*

- **Compressed communications**: data sent is compressed to reduce bandwidth usage on the

customer's premises.

- **Push-based data delivery mechanism**: to facilitate firewall configuration, the direction of the connections established between the Azure infrastructure and the customer's network is outbound from the customer's network. After the communication channel has been established, Azure sends all new logs available on the platform using push messages, instead of mechanisms such as pooling.

Information flow

The activity of the processes run on the customer's network is collected by Cytomic EDR through its continuous monitoring of the organization, and sent to the Cytomic cloud. There, it is enriched with security intelligence and placed in the Azure infrastructure, where it is temporarily stored until it is downloaded by the customer subscribed to the Cytomic SIEMConnect service.

The Cytomic Importer program, which is run on a server on the customer's network, downloads the generated logs and manages them in different ways depending on how it is configured:

- Stores the logs in a folder accessible to the organization's SIEM server, managing the volume of stored files so as not to exceed the limits set by the administrator.
- Sends the logs to a Kafka queue server from which they are collected by the SIEM server at the rate its resources allow.
- Sends the logs to a Syslog server from which they are collected by the SIEM server at the rate its resources allow.

The customer's SIEM server imports the logs and analyzes them periodically to incorporate the information into its repository and generate the relevant data in dashboards.

Chapter 3

Panda SIEMConnect for Partners architecture

Panda SIEMConnect for Partners is the Cytomic service that provides the MSSP's SIEM platform with all the information and knowledge generated by the Cytomic EDR products installed on customers' computers. Cytomic SIEMConnect helps MSSPs discover unknown threats, targeted attacks, and APTs (Advanced Persistent Threats) on customers' systems.

CHAPTER CONTENTS

Aims of the Panda SIEMConnect for Partners service

The primary objective of Panda SIEMConnect for Partners is to act as a link between the protection software installed on customers' computers and the MSSP's SIEM server, within the following general information flow:

- Cytomic EDR's continuous monitoring sends the Cytomic cloud the telemetry generated by the applications run on the MSSP's customers' systems.
- Panda SIEMConnect for Partners enriches this information with the security intelligence generated by Cytomic.
- Cytomic Importer recovers the enriched information from the Azure infrastructure assigned to the MSSP and sends it directly to the SIEM server or to one of the supported storage and queue management platforms (Kafka and Syslog), for it to be leveraged later.

Enriching the monitored activity

Cytomic EDR monitors the actions executed by processes on customers' computers. These actions are sent to the Cytomic cloud platform, where they are analyzed using machine learning techniques on a big data infrastructure in order to extract advanced security intelligence. This information enables Cytomic to classify each and every process run by customers, with 99.999% accuracy and near-zero false positive and false negative errors.

Panda SIEMConnect for Partners combines the information collected from the events monitored by Cytomic EDR and the security data generated, creating a single data flow compatible with the MSSP's SIEM server.

In order to take full advantage of Panda SIEMConnect for Partners, it is not necessary to make any changes to users' computer settings: the service operates within the Cytomic infrastructure, receiving data from every computer on the customer's IT systems. This data is normalized, enriched, and sent to the MSSP's designated SIEM server.

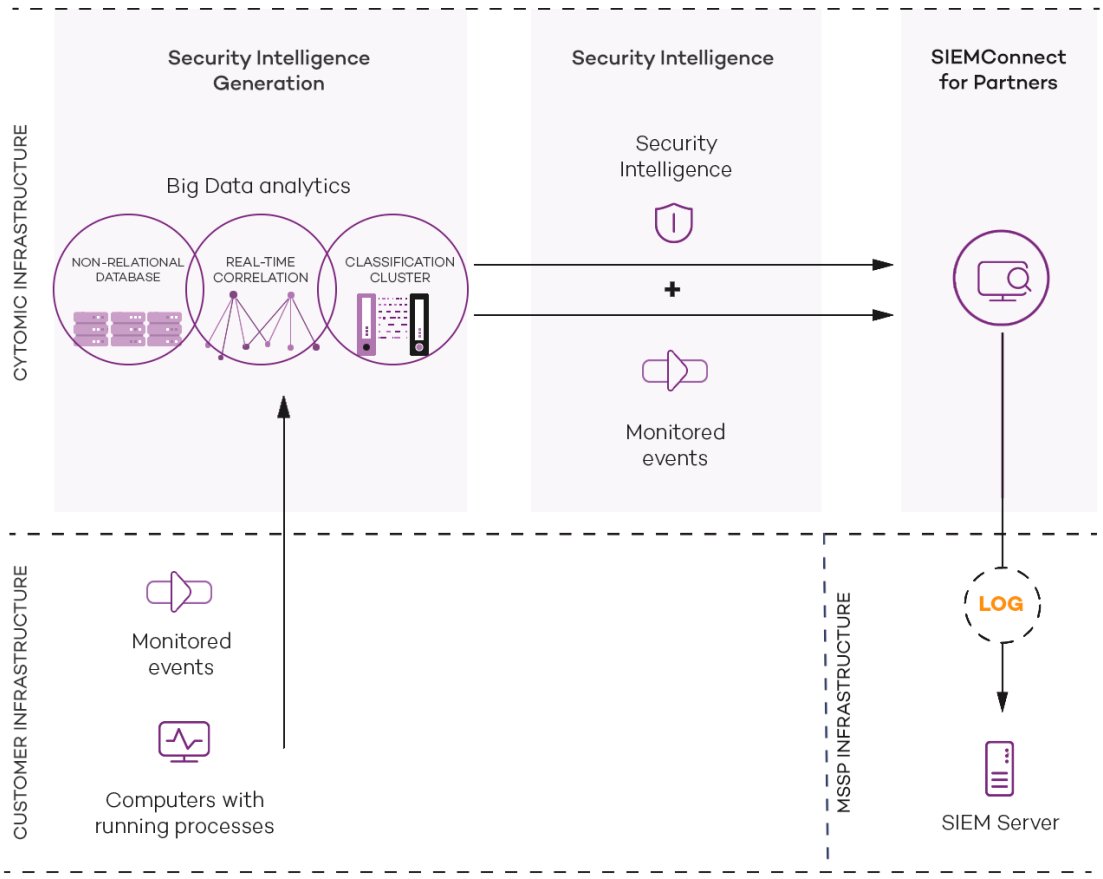

Figure 3.1: Information flow generated by Cytomic EDR and Cytomic SIEMConnect

Benefits of the Panda SIEMConnect for Partners service

Panda SIEMConnect for Partners delivers information about the activity of processes run on the MSSP's customers' systems. This information enables MSSPs to:

- **View details of the malware detected on customers' networks**, and see whether it was run or not, the infection vector, and the actions taken by processes. This information helps administrators make decisions with regard to defining remedial actions and adapting security policies.
- **View the actions executed by each process**, whether goodware, malware, or temporarily unknown, in order to detect the suspicious activity of recent programs. Panda SIEMConnect for Partners compiles indicators that can be used to reach conclusions about their potential threat.
- **View access by processes to confidential information**, preventing data leakage and theft. The service displays the Office files, databases, and other repositories of confidential information accessed by malware.
- **View network connections made by processes** in order to identify suspicious or potentially dangerous connections that could be used to steal data.
- **Locate all executed programs**, especially those with known vulnerabilities installed on users' computers, in order to design plans for updating software and adapting corporate security policies.
- **Apply centralized settings through CYTOMIC Nexus.** All MSSPs' customers' settings can be pushed out simultaneously.
- **Install the service simply and securely** as the telemetry download service only needs to be configured once, and new customers are added without having to deploy or install any additional components. Downloads are also secure thanks to secure TLS (Transport Layer Security) connections from the Cytomic cloud.
- **Keep storage costs down** by filtering events before they reach the MSSP's infrastructure, thereby minimizing the bandwidth and space consumed.

General architecture

The Panda SIEMConnect for Partners service consists of the modules displayed in figure 3.2:

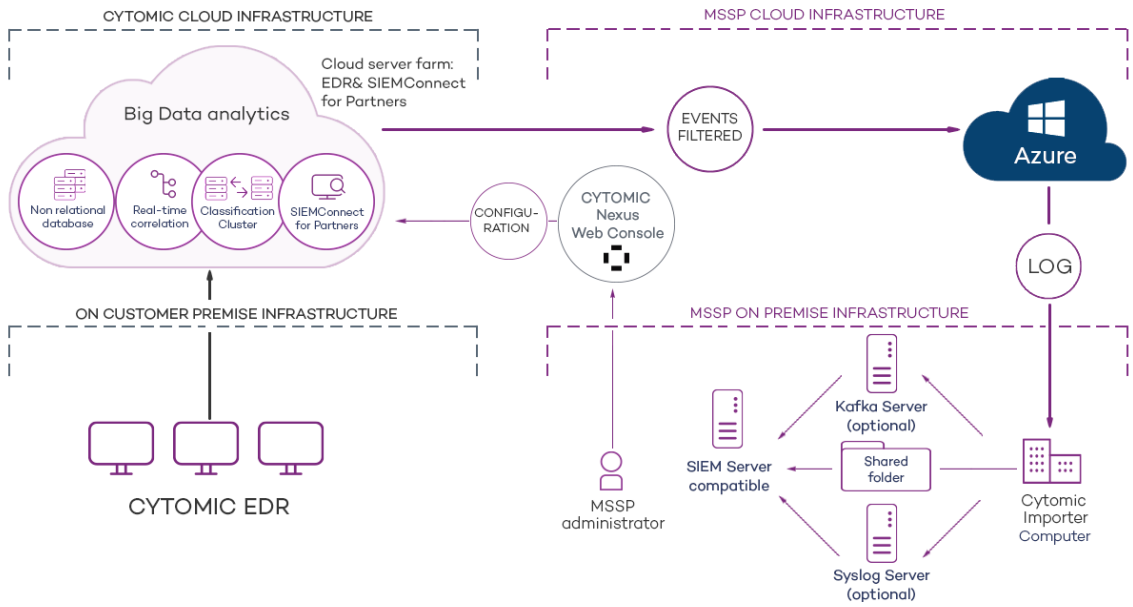

Figure 3.2: Logical diagram of the modules that make up Panda SIEMConnect for Partners and their relationships

The figure shows the following items:

- **Computers on the customer's network:** protected by Cytomic EDR or Cytomic EDPR.
- Cytomic **cloud**: this stores information from the processes run and analyzes it to extract security intelligence.
- **Panda SIEMConnect for Partners service**: this receives events and security intelligence information and encapsulates them in the form of logs which are later sent to the MSSP's Azure infrastructure.
- **MSSP's Azure infrastructure**: this receives the logs from the Panda SIEMConnect for Partners service on the cloud infrastructure assigned to the MSSP and stores them temporarily until downloaded from Cytomic Importer.
- **Cytomic Importer computer**: a computer on the MSSP network that runs the Cytomic Importer process and checks if there are new logs available on the Azure infrastructure to download and store.
- **Kafka server (optional)**: the computer on the MSSP's network that manages the queues of logs received by Cytomic Importer and sends them to the SIEM server.
- **Syslog server (optional)**: the computer on the MSSP's network that collects the logs received by Cytomic Importer and sends them to the SIEM server.
- **Share folderd (optional)**: a storage system on the MSSP's network where Cytomic Importer deposits logs in the absence of more advanced resources, such as a Syslog or Kafka server.
- **SIEM server**: installed on the MSSP's premises, this collects the data downloaded by the Cytomic Importer computer in order to generate dashboards that help detect suspicious processes that may

pose a threat to customers.

- **Local and perimeter firewalls**: these protect inbound and outbound data traffic between the Cytomic Importer computer and the Azure infrastructure.
- **CYTOMIC Nexus console**: this enables the MSSP to activate the Panda SIEMConnect for Partners service for customers and configure it in order to receive only selected events.

Benefits of the Azure infrastructure

Panda SIEMConnect for Partners generates logs asynchronously and stores them temporarily until they are collected and integrated in the SIEM server. To do that, Cytomic makes use of cloud-based services hosted on the Azure infrastructure. These services have the following characteristics:

- **Cloud-based storage**: high availability service, meaning the service is available 24 hours a day, from anywhere in the world. Each MSSP is assigned 80 GB and the information is stored for a maximum of 7 days.
- **Encrypted communications**: the information exchanged between the MSSP's Cytomic Importer computer and the Azure platform is encrypted with the SSL cryptographic protocol.
- **Authenticated communications**: in order to manage the Cytomic Importer authentication and authorization processes, two independent tokens are used to negotiate the shared key required to access the Panda SIEMConnect for Partners platform and recover the MSSP's customers' network information. These tokens have different expiration times in order to ensure data confidentiality.

*For more information, refer to "**Appendix 2: Security Architecture**" on page **55**.*

- **Compressed communications**: the data sent is compressed to reduce bandwidth usage on the MSSP's premises.
- **Push-based data delivery**: To aid configuration of firewalls, the direction of the connections established with the Azure infrastructure is outbound from the MSSP's network. After the communication channel has been established, Azure sends all new logs available on the platform using push messages, instead of mechanisms such as pooling.

General functions of the service provider

To recover the security information generated by customers, the security service provider has to interact with the items shown in figure **3.2**. To do this, it is necessary, in general, to follow the action list below:

- Check that the subscription to the Panda SIEMConnect for Partners service is active in CYTOMIC Nexus. Refer to **License management** in the **CYTOMIC Nexus Administration Guide**. If your subscription has expired or you are not a Panda SIEMConnect for Partners, customer, contact your assigned Cytomic sales representative.

- Check the connectivity of the items shown in figure 3.2, especially with respect to communication between Cytomic Importer and Azure. Refer to "**Firewall settings**" on page 23.
- Check the deployment and installation requirements. Refer to "**Deployment and integration requirements**" on page 21.
- Install Cytomic Importer on your IT infrastructure. Refer to "**Installing and configuring Cytomic Importer on Windows systems**" on page 29 or "**Installing and configuring Cytomic Importer on Linux systems**" on page 41.
- Create a Panda SIEMConnect for Partners setting profile in CYTOMIC Nexus detailing the groups of events to be sent to the Azure infrastructure. Refer to **SIEMConnect for Partners settings**in the **CYTOMIC Nexus Administration Guide**.
- Associate the newly created settings with the customers. Refer to **Assigning and sending settings** in the **CYTOMIC Nexus Administration Guide**. Depending on the **Enable real-time communication** option selected in the security product console for each customer, settings will either be assigned immediately or with a maximum delay of 10 minutes. Refer to "**Cytomic EDR and Cytomic EDPR guides**" for the administration guide of the product contracted by the customer.

After all steps have been completed, the MSSP's customers' computers start to send information which is stored on the Azure infrastructure until the MSSP's Cytomic Importer computer downloads it.

Chapter 4

Deployment and integration requirements

The requirements necessary to correctly implement the Cytomic SIEMConnect and Panda SIEMConnect for Partners services fall into three categories:

- Licenses and user information
- Deployment and operation
- Integration into the existing IT infrastructure

CHAPTER CONTENTS

Licenses and required information

Both Cytomic SIEMConnect and Panda SIEMConnect for Partners require the customer ID sent in the welcome email and the details of a user of the management console of the security product purchased by the organization and compatible with the service:

- Cytomic EDR (compatible with Cytomic SIEMConnect and Panda SIEMConnect for Partners)

- Cytomic EDPR (compatible with Cytomic SIEMConnect and Panda SIEMConnect for Partners)

Cytomic SIEMConnect

- For Cytomic EDR or Cytomic EDR, use the email address and password of a console user with the Full Control role.
- An email account managed by the administrator. This will be used to receive notifications about the service status.
- The customer ID included in the welcome email sent to the network administrator when the **Cytomic EDR** service was provisioned.

Panda SIEMConnect for Partners

- The email address and password of a Cytomic EDR or Cytomic EDR user in the MSSP with the Full Control role.
- The customer ID included in the welcome email sent to the MSSP technician when the **Cytomic EDR** service was provisioned.
- An email account managed by the administrator. This will be used to receive notifications about the service status.

Deployment and integration requirements

The requirements necessary to deploy and use Cytomic SIEMConnect and Panda SIEMConnect for Partners are described in the following sections:

- User computers protected by Cytomic EDR.
- Active Cytomic SIEMConnect or Panda SIEMConnect for Partners license.
- Computer with Cytomic Importer installed.
- Firewall settings.
- Proxy server settings.
- Enough bandwidth to receive data.

Cytomic Importer computer

This computer must meet the following minimum requirements:

- 1 GHz or faster processor.
- At least 512 MB of RAM.
- Enough free space to store the information received. The average space taken up on the storage

device is **1 MB per computer/hour**. The information is stored in uncompressed log files in LEEF format.

> *To change to CEF format in Cytomic SIEMConnect, send an email to siemconnect@cytomic.ai. To change the format of log files in Panda SIEMConnect for Partners, access CYTOMIC Nexus and modify the assigned settings.* Refer to **SIEMConnect for Partners settings** in the **CYTOMIC Nexus Administration Guide**.

- The Cytomic Importer program must be installed and configured correctly. For more information, refer to "**Installing and configuring Cytomic Importer on Windows systems**" on page **29**.
- The access information specified in section "**Licenses and required information**".
- An NTP server is required for computer time synchronization.

Firewall settings

In order for the Cytomic Importer computer to be able to download log files from the Azure infrastructure, all intermediate firewalls must allow network traffic with the following characteristics:

- Access to the URL **https://auth.pandasecurity.com**.
- Access to the URL **https://storage.accesscontrolmngr.pandasecurity.com**.
- Access to the URL **sb://pac100siemfeeder.servicebus.windows.net**.
- **Communication source**: Cytomic Importer computer.
- **Communication target**: Azure infrastructure.
- **Connection type**: outbound from the customer's network.
- **Layer 3 (transport) protocol**: TLS 1.2.
- **Layer 4 (application) protocol**: HTTPS (port 443), Amqp (ports 5671 and 5672), Amqp WebSockets (port 443).

Proxy server settings

If Cytomic Importer accesses the Cytomic cloud via a proxy server, the proxy server must have access via Web Sockets enabled. In that case, Cytomic Importer will use the Amqp WebSockets protocol instead of Amqp.

Bandwidth

Each user computer generates an average of 500 KB of compressed information in gzip format per hour.

The required bandwidth depends directly on the number of user computers monitored on the customer's network and the maximum allowable delay based on the organization's needs. Two thresholds can be established:

- **Minimum threshold**: minimum bandwidth required to receive all logs without losing files due to expiration of the log retention period. For more information, refer to "**Service availability**" on page **26**.

The log generation rate depends on multiple factors (computer activity, the computer's role within the organization, etc.). With a low bandwidth value, the service will leverage valley hours (non-working hours when most computers are turned off) to receive the log files generated during peak hours.

A low bandwidth value will lead to delays in receiving logs, preventing them from being received and processed in real time by the organization's SIEM server.

- **Maximum threshold**: the bandwidth required to download all log files as they are generated.

Data leverage requirements

To leverage the delivered data, install and configure a SIEM server compatible with any of the supported log formats.

Supported SIEM servers

The SIEM products compatible with the Cytomic SIEMConnect or Panda SIEMConnect for Partners service are those that support the Common Event Format (CEF) developed by ArcSight and the Log Event Extended Format (LEEF) developed by QRadar.

The data is received in one of the two formats: CEF or LEEF. The following is a partial list of SIEM servers compatible with the aforementioned formats:

- AlienVault Unified Security Management (USM)
- Fortinet (AccelOps) FortiSIEM
- Hewlett Packard Enterprise (HPE) ArcSight
- IBM's QRadar Security Intelligence Platform
- Intel Security's McAfee Enterprise Security Manager (ESM)
- LogRhythm
- SolarWinds Log & Event Manager (LEM)
- Splunk's Security Intelligence Platform

SIEM server settings

In order for the SIEM server to receive the data, you must set as source the system chosen for storing the data and correctly map the delivered events and fields. Available sources are:

- The folder where the Cytomic Importer computer stores the received logs.
- The Kafka queue server that collects the logs sent by Cytomic Importer.

- The Syslog server that collects the logs sent by Cytomic Importer.

> *For a full description of the delivered data, refer to the* **Cytomic SIEMConnect Event Description Guide**.

Characteristics of the downloaded log files

- Each log file has a maximum size of 256 KB in compressed format.
- Log files are stored decompressed in the configured folder respecting the maximum size defined in the settings.
- Each log file has a name in the format yyyymmdd-hhmm-(xxxxxx), where:
  - **yyyy**: year created.
  - **mm**: month created.
  - **dd**: day created.
  - **hh**: time created (hours).
  - **mm**: time created (minutes).
  - **–(xxxxxx)**: number assigned to the log file if more than one file is created within the same minute.

Sizing recommendations for the Cytomic Importer computer

Bandwidth sizing

- Calculate the bandwidth required based on the number of monitored computers on your network (500 KB per computer/hour).
- Use the value calculated in the previous point to set QoS rules on the router in your organization that connects the Cytomic Importer computer to the Internet. Monitor your bandwidth usage.at all times.
- Compare the date the log files were received on the Cytomic Importer computer to the date the events were generated to find out if there are delays in receiving data. The generation date of the log files is provided by the operating system. The generation date of each event is part of the log file's internal information schema. For a detailed description of all the fields included in log files, refer to the **Cytomic SIEMConnect Event Description Guide**.
- If the difference between the log receiving date and the event generation date increases gradually over time, check the received data flow.
- If the data flow is using all bandwidth reserved by the QoS rule, it means that Cytomic SIEMConnect is generating a number of log files too large for the bandwidth allocated to the Cytomic Importer computer. If, after 7 days (a full week to include periods of lower activity), the above mentioned

difference does not decrease, or the organization requires a shorter event receiving time, increase the bandwidth allocated to the service by the QoS rule.

- If the bandwidth allocated to the service is not completely used, but the difference between the log receiving date and the event generation date increases, there is a bottleneck in the Cytomic Importer computer hardware. Refer to "**Hardware sizing recommendations for the Cytomic Importer computer**".

Hardware sizing recommendations for the Cytomic Importer computer

If the difference between the log receiving date and the event generation date increases gradually over time, but the bandwidth allocated to the service is not completely used by the received data flow, it is very likely that there is a bottleneck in the Cytomic Importer computer hardware.

Because Cytomic Importer is a program that recovers messages from a queue-type structure, its CPU and RAM requirements are relatively low. In complex networks, with a large number of monitored computers, the main reason for slow download speeds is usually a bottleneck in the storage system of the computer running Cytomic Importer. Follow the advice below to determine the source of bottlenecks and resolve them. To see CPU and hard disk performance stats, you must start the Windows task manager with Cytomic Importer running in command-line or service mode.

- **High CPU usage with free cores**: Cytomic Importer is a single-threaded program, that is, it uses only one of the cores of the processor installed on the server. If the Windows task manager indicates that there is sustained CPU usage of more than 80% on one of the cores, run several instances of the program with different target folders. A conservative recommendation is to run one instance of Cytomic Importer for each core on the computer. For more information, refer to chapter "**Installing and configuring Cytomic Importer on Windows systems**" on page **29**
- **High CPU usage without free cores**: if the Windows task manager indicates that there is sustained CPU usage of more than 80% on all cores, it is advisable to install Cytomic Importer on an additional server or change the CPU to a more powerful one.
- **High bandwidth consumption in the storage system**: if the Windows task manager indicates high hard-disk usage, it is advisable to replace one or all of the components of the storage subsystem:
  - Replace mechanical disk drives with solid state drives (SSD).
  - Install a RAID–0 system or equivalent that allows data to be written to several drives at a time.
  - Replace the data bus interface with a more up-to-date version (SATA, eSATA, SAS, etc.).

Service availability

Cytomic SIEMConnect is available 24/7. Any service interruption is notified via email to the administrator account provided during the registration process.

To prevent data loss in the event of connectivity failure, unavailability of the customer's Cytomic Importer computer, or any other error, Cytomic retains the logs generated and not delivered to the customer for the following time periods:

- **Maximum number of days that logs are retained on the Azure platform:** 7 days.
- **Maximum amount of data retained on the Azure platform**: 80 GB for each customer.

Chapter 5

Installing and configuring Cytomic Importer on Windows systems

Cytomic Importer is the application responsible for downloading the events logged by Cytomic SIEMConnect and Panda SIEMConnect for Partners from the Azure infrastructure. These events are stored in log files which, depending on the configured settings, are decompressed by Cytomic Importer and placed in a local or remote folder, or sent to a compatible server (Kafka or Syslog).

CHAPTER CONTENTS

Installation requirements

Required information

For the information required by Cytomic Importer for its correct operation, refer to "**Licenses and required information**" on page **21**.

Operating system and required libraries

Make sure the computer that will run the Cytomic Importer program meets the following requirements:

It has .NET Framework 4.6.2 or higher installed: if an earlier version is installed, go to **https://dotnet.microsoft.com/en-us/download/dotnet-framework/net462** to download the appropriate version. Cytomic Importer is compatible with .NET Framework up to version 4.8.

- **Supported operating systems**: Windows 11, Windows 10, Windows 8.1, Windows 8, Windows 7 Service Pack 1, Windows Server 2022, Windows Server 2019, Windows Server 2016, Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2 SP1.

Required permissions

Cytomic Importer can be run as a command-line program or unattended as a Windows service.

- When run as a service, Cytomic Importer runs under the `local system` account and needs admin permissions to be correctly installed.
- When run as a command-line program, it doesn't require any specific permissions beyond access to the resources it may need; for example, write access to the folder configured to store the downloaded logs.

Firewall configuration

In order for the Cytomic Importer computer to be able to download log files from the Azure infrastructure, all intermediate firewalls must allow network traffic with the following characteristics:

- Access to the URL **https://auth.pandasecurity.com**.
- Access to the URL **https://storage.accesscontrolmngr.pandasecurity.com**.
- Access to the URL **sb://pac100siemfeeder.servicebus.windows.net**.
- **Communication source**: Cytomic Importer computer.
- **Communication target**: Azure infrastructure.
- **Connection type**: outbound from the customer's network.
- **Layer 3 (transport) protocol**: TLS 1.2.

- **Layer 4 (application) protocol**: HTTPS (port 443), Amqp (ports 5671 and 5672), Amqp WebSockets (port 443).

NTP server

To download the logs stored in the Azure infrastructure, an authentication and authorization process must be completed that involves generating a token. This token is issued with an expiration date to improve the security of the entire process, therefore, the clocks of both communication endpoints must be synchronized. For this reason, it is required that the computer that runs Cytomic Importer have the Windows Time Service (or an equivalent service) up and running in order to get the time from an NTP server. For more information, refer to **https://docs.microsoft.com/en-us/windows-server/networking/windows-time-service/accurate-time**.

Installation and configuration

For more information about the origin of the errors that may occur during the installation process, refer to "**Appendix 1: Troubleshooting**" on page **53**.

To install and configure Cytomic Importer, perform the steps below in the listed order:

1. Download and decompress the .ZIP file containing the installer: Refer to "**Download the installation package**".
2. Indicate the connection method supported by the IT infrastructure that will host the Cytomic Importer computer: direct connection or through a corporate proxy. Refer to "**Update the configuration.json file**".
3. Indicate the platform where your Cytomic security products reside. Refer to "**Configure the platform to use**".
4. Enter the credentials of the account used to access the service. Refer to "**Enter the access credentials**".
5. Configure the method to be used to send and store the received logs. Refer to "**Configure the log storage and send mode**".
6. Configure how Cytomic SIEMConnect will be run: as a service or from the command line. Refer to "**Configure the execution mode**".
7. Update the `configuration.json` file with the new installation settings. Refer to "**Update the configuration.json file**".

Download the installation package

Download the .ZIP package of the Windows version of Cytomic Importer from **https://www.cytomic.ai/en/support/id-950031** and decompress it to a folder on your computer. This package contains the following main files:

- `EventsFeederImporter.Host.exe`: downloads the log files containing the events that occurred on the customer's computers, and stores them on the computer hard disk or forwards them to

another computer depending on the configuration defined by the administrator. This program can be run as a process from the command line or as a service (refer to "**Configure the execution mode**").

- `EventsFeederImporter.ConfigAssistant.exe`: displays the configuration wizard containing the parameters required to configure Cytomic Importer.
- `Configuration.json`: this file contains the program settings. All personal data is stored obfuscated to prevent security leaks.

Cytomic Importer can be run multiple times simultaneously on the same computer. Each Cytomic Importer instance requires a separate configuration file. Refer to "**Configuring multiple instances**".

Configuration

This section describes the steps you must take to generate the configuration file required for a single Cytomic Importer instance to run in command-line or service mode and connect to the Azure infrastructure in order to download logs. All other scenarios are based on this configuration.

To configure Cytomic Importer, you must run the `EventsFeederImporter. ConfigAssistant.exe` program in command-line mode and answer "Yes" to the question **Do you want to change the configuration settings? [Yes / No]** A new configuration file that overrides the existing one is generated and the configuration wizard is launched.

To install Cytomic Importer as a service, you must run `EventsFeederImporter.ConfigAssistant.exe` *with administrator permissions. Right-click the file and select* ***Run as administrator***.

Configure the connection method

If the computer is behind a proxy server, answer **Y** to the question **Is Event Importer behind a proxy server? [Yes/No]** You will be prompted to enter the proxy server IP address, as well as the user name and password if the proxy server requires authentication.

Access through the configured proxy server is only used to connect to the Azure infrastructure assigned to the customer or MSSP. It is not used to connect to other resources such as the file server, the Kafka server, or the Syslog server.

Configure the platform to use

- Answer the question **Select your platform: [C]urrent, [L]egacy or [W]G Endpoint Security** with **C** (Current platform).

Enter the access credentials

- Enter the email address of the user account used to access the Cytomic EDR console.
- Enter the password. If the account has 2FA enabled, enter the 6-digit OTP code immediately after

the password, without spaces.

- Enter the Customer ID specified in the welcome email. After you enter it, Cytomic Importer generates a new access token it uses internally to subscribe to the service and download the generated log files.

To determine if the access account has 2FA enabled, go to the Cytomic EDR management console:

- Click **https://central.cytomic.ai**
- Click the account name in the upper-right corner of the page. A drop-down menu appears.
- Click **Set up my profile**. The **Cytomic Account** page opens**.** This page indicates whether 2FA is enabled or not.

For more information about how to enable 2FA, refer to **https://info.cytomic.ai/central/en/index.htm**

Configure the log storage and send mode

For more information about how to choose the method to be used to store and send the downloaded logs, refer to "**Configuring log storage and forwarding**".

Configure the execution mode

Cytomic Importer can be run as a service or in command-line mode. Answer **Y** or **N** to the question **Do you want to register Event importer as a Windows service? [Yes / No]**

- **Y:** this installs the program as a service. Cytomic Importer will install automatically as a service only if the user who started the installation process has admin permissions.

*Select the (**Y**) option only if you are going to install and run a single Cytomic Importer instance as a service on your computer. In all other cases, select (**N**). Refer to* "**Configuring multiple instances**".

- **N**: select this option to run one or multiple instances from the command line or to run multiple instances of the program as a service. Refer to "**Configuring multiple instances**".

Update the configuration.json file

After it finishes running, the configuration wizard updates, with the entered information, the `configuration.json` file located in the same folder. Then, Cytomic Importer starts downloading the logs stored in the Azure infrastructure.

The `configuration.json` file contains the following data:

- Information about the customer whose logs are downloaded.
- Information about the method selected to send and store the downloaded logs.
- Information about the execution mode (command line or service).

Configuring multiple instances

You must configure multiple instances of Cytomic Importer in the following cases:

- If the computer that runs the Cytomic Importer program shows symptoms indicating a lack of resources similar to those described in section "**Hardware sizing recommendations for the Cytomic Importer computer**" on page **26**, we recommend that you install one or more additional instances of the program and run them concurrently.
- If you require that a single computer with Cytomic Importer download logs from more than one customer simultaneously, but you are not using Panda SIEMConnect for Partners.

*To download logs files from multiple customers and centralize all downloads through a single Cytomic Importer instance, use Panda SIEMConnect for Partners. Refer to "**Panda SIEMConnect for Partners architecture**" on page **15**.*

Multiple instances in command-line mode

- Download the latest version of Cytomic Importer from **https://www.cytomic.ai/en/soporte/id-950031** and decompress it to a separate folder for each customer whose logs you want to download.
- Configure each application independently using the steps described in section "**Configuration**" to install it in command-line mode.
- Run each application independently.

Multiple instances in service mode

To run multiple instances of Cytomic Importer in service mode, you must first install it in command-line mode and then register it manually as a service. To do this, follow the steps below:

- This example uses folders `c:\users\customer1` and `c:\users\customer2`.
- Follow the steps in section "**Multiple instances in command-line mode**".
- Use `control + c` to stop the execution of each running instance.
- If you intend to run Cytomic Importer as a service, you must register it manually. To do this, you must give each instance a different name, using the following parameters: `servicename`, `description`, and `displayname`. Run the following commands as administrator:

```
PS C:\> cd c:\users\customer1
PS C:\users\customer1> EventsFeederImporter.Host.exe install
-servicename:ServiceCustomer1
-description: ServiceCustomer1
-displayname: ServiceCustomer1
PS C:\> cd c:\users\customer2
PS C:\users\customer2> EventsFeederImporter.Host.exe install
-servicename:ServiceCustomer2
-description: ServiceCustomer2
-displayname: ServiceCustomer2
```

- Start each Cytomic Importer instance.

```
PS C:\> cd c:\users\customer1
PS C:\users\customer1> EventsFeederImporter.Host.exe start
-servicename:ServiceCustomer1
PS C:\> cd c:\users\customer2
PS C:\users\customer2> EventsFeederImporter.Host.exe start
-servicename:ServiceCustomer2
```

Configuring log storage and forwarding

Cytomic Importer provides several methods to store and forward logs based on the network architecture, the available resources, and the volume of information received from the Azure infrastructure:

- Storing the logs in a local or remote folder.
- Sending the logs to a Kafka server.
- Sending the logs to a Syslog server.

The storage method is configured in the configuration wizard when the following question is shown: **Importer enables you to send received events simultaneously to various channels. Do you want to change the current channel settings? [Yes / No]**. Selecting **Y** deletes the existing storage and forwarding settings (if any) and generates new settings.

Storing logs in a local or remote folder

- First, create the folder where the log files will be stored. This can be a local folder on the computer running Cytomic Importer or a remote, shared drive/resource.
- If you intend to run multiple instances of Cytomic Importer, create a separate folder for each instance. Otherwise, some logs might be lost during collection and storage.
- Select **F** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server.**
- Enter the full folder path for each instance of Cytomic Importer.
- Enter the extension of the files the events received from Cytomic Importer will be dumped to.
- To finish configuring the storage method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]**.

Sending logs to a Kafka server

- Select **K** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server**.
- Enter the IP address or domain name of the Kafka server and the listening port separated by ":".
- Enter the name of the queue/topic that logs will be sent to on the Kafka server.

- Enter the communication protocol that Cytomic Importer will use to send logs to the Kafka server:
  - **[N]one**: press **N** to send logs in unencrypted format.
  - **[S]sl:** press **S** to send logs using SSL encryption.
  - **s[A]slssl:** press **A** to send logs using SASL/SSL encryption.
  - **saslplain[T]ext:** press **T** to send logs using SASL/PLAIN encryption.
- Depending on whether the communication protocol chosen encrypts data or not, you will have to indicate the path of the file that contains the certificate issued by the CA configured on the Kafka server.
- To finish configuring the delivery method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]**.

Sending logs to a Syslog server

- Select **S** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server**.
- Select the format configured on the Syslog server for the received logs: **RFC[5]424** or **RFC[3]164**.
- Enter the IP address or domain name of the Syslog server and the listening port separated by ":".
- Select the transport protocol configured on the Syslog server for the received logs: **[T]cp** or **[U]dp**.

To make sure all log files sent by Cytomic Importer are received on the Syslog server, we recommended configuring the use of the TCP transport protocol on both ends. Otherwise, in overload situations, the UDP protocol might inadvertently discard logs.

- Select the secure protocol to use to encrypt communications between the Syslog server and Cytomic Importer. **[N]one** or **Tls1.[2]**.
- Select the end-of-message marker configured on the Syslog server for the received logs: **[C]r**, **[L]f**, **c[R]lf**.

If the transport protocol chosen is UDP, no end-of-line marker is used.

If the transport protocol chosen is TCP or TLS, the Null end-of-line marker is always used.

- If the communication protocol chosen encrypts data, indicate the location of the certificate issued by the CA configured on the Syslog server.
  - **[F]ile**: the CA's certificate is in a separate file**.**
  - **[C]ert Store**: the CA's certificate is located in the local certificate store on the computer where Cytomic Importer is run; more specifically, in the Trusted People Certificates branch.
- To finish configuring the delivery method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]**.

Copying downloaded logs to multiple locations

Cytomic Importer allows logs to be downloaded to multiple locations simultaneously. Each log downloaded this way will be removed from the download queue if at least one of the target locations is correctly updated. Therefore, should errors occur during log collection, the different locations might end up with a different number of logs at the end of the process. To implement the ability to download logs to multiple locations, the 'channels' feature is used. A channel indicates the storage type used by Cytomic Importer and its settings.

To configure Cytomic Importer, follow the steps below:

- Install Cytomic Importer as described in section "**Configuration**".
- Stop Cytomic Importer as described in section "**Running and stopping the program**".
- Add a new channel to the existing collection in file `configuration.json`. The `channels` parameter syntax is the following:

```
"Channels": [{ channel 1 parameters} , {channel 2 parameters}, …]
```

- Each channel indicates the storage type used to store logs and its associated settings.

Below you can find and example of a Cytomic Importer configuration with two channels: the first channel leaves logs in the `Log1` folder, and the second channel leaves them in folder `Log2`:

```
"Channels": [{
            "Type": "LocalDisk",
            "Name": "LD1",
            "Configuration": {
                        "fullPath":
"D:\\\\SIEMFeeder\\\\EventFeederImporter 1.0.3 Pro\\\\Log1",
                        "filesSplitFormat": "1m",
                        "filesSizeLimitInBytes": 102400,
                        "directoryMaxSizeInMb": 1024
                    }
            }, {
            "Type": "LocalDisk",
            "Name": "LD2",
            "Configuration";{
            "fullPath:
"D:\\\\SIEMFeeder\\\\EventFeederImporter 1.0.3 Pro\\\\Log2",
                        "fileSplitFormat":"1m",
                        "filesSizeLimitInBytes": 102400,
                        "directoryMaxSizeInMb": 1024
                    }
            },]
```

Running and stopping the program

After the configuration wizard is completed, Cytomic Importer starts automatically. To stop or start the program later, follow the steps below:

In command-line mode

- To start Cytomic Importer, double-click the `EventsFeederImporter.Host.exe` file or run it from the command line.
- To stop Cytomic Importer, press `Control+C` in the command window.

In service mode

- The service is automatically configured to start when the operating system boots. To run Cytomic Importer after a manual stop, access the Services snap-in of the operating system's MMC console and find the `EventsFeederImporter` service (this is the default name used by the installation wizard provided the service was not manually registered). Right-click it and select Start.
- To stop Cytomic Importer, access the Services snap-in of the operating system's MMC console and find the `EventsFeederImporter` service. Right-click it and select Stop.

Changing the settings

To regenerate the Cytomic Importer settings, follow these steps:

- Stop the process if it is running. Refer to "**Running and stopping the program**".
- Run the `EventsFeederImporter.ConfigAssistant.exe` program again, either from the command-line window or by double-clicking on it.
- Answer **Y** to the question **Do you want to change the configuration settings?**
- Enter the information described in section "**Configure the platform to use**".
- Start Cytomic Importer. Refer to "**Running and stopping the program**".

Manually editing the Cytomic Importer settings

The `configuration.json` file is in the same folder as Cytomic Importer and contains its execution parameters.

The first time it is run, the program asks for certain authentication information that it later dumps into the configuration file. Once the installation and execution processes are completed, you'll be able to edit the configuration file in order to fine-tune its parameters.

Every time you modify the `configuration.json` *file, you must stop and start the Cytomic Importer process so that it applies the changes made to the file. Refer to "***Running and stopping the program***".*

The `configuration.json` file follows the JSON syntax.

Parameters related to log files containing events

The parameters that determine how Cytomic Importer behaves in order to generate the log files that contain the events logged by Cytomic SIEMConnect are:

- **fullPath**: absolute path to the folder the log files are downloaded to.
- **fileSizeLimitInBytes**: maximum allowed size of the log files.
- **directoryMaxSizeInMB**: maximum allowed size of the folder that stores the log files. When the maximum size is reached, 10 percent of the oldest files are deleted.
- **FileSplitFormat**: rotation interval of the log files. The file name contains the year(yyyy), month(MM), day (dd), hour(HH), and minute (mm) when the file was created.
  - "**1h**" or **empty**: yyyyMMdd-HH format. A file is generated every hour.
  - "**1m**": yyyyMMdd-HHmm format. A file is generated every minute.
  - "**5m**": yyyyMMdd-HHmm format. A file is generated every 5 minutes.
  - "**10m**": yyyyMMdd-HHmm format. A file is generated every 10 minutes.
  - "**15m**": yyyyMMdd-HHmm format. A file is generated every 15 minutes.
  - "**30m**": yyyyMMdd-HHmm format. A file is generated every 30 minutes.
- **Channels**: indicates the characteristics of the channel used to download the log files.
- **Type**: storage type used in the channel.
- **Name**: channel name.
- **Configuration**: channel settings (fullPath, fileSplitFormat, fileSizeLimitInBytes, directoryMaxSizeInMB)).

Parameters related to the execution log

All operations executed by Cytomic Importer are logged to text files stored in the `Log` folder inside the same folder that contains the program.

*Refer to "***Appendix 1: Troubleshooting***" on page* **53** *for a description of the errors Cytomic Importer can generate.*

The parameters that determine how Cytomic Importer behaves in order to generate the log files that log the actions it takes are:

- **LogsPath**: absolute or relative path and file name. The backlash character ("\") must be escaped. For example `".\\log\\log.txt"`.
- **LogFileSizeLimitKBytes**: rotates the log file when it reaches a certain size in Kbytes, adding the suffix "-SequenceNumber" to it. For example `"log-3.txt"`.
- **LogRetainedFileCountLimit**: indicates the maximum number of files that Cytomic Importer stores on the storage device. When this number is reached, Cytomic Importer deletes the oldest file.
- **Interval**: rotation interval of the log files:
  - **0**: No rotation. The suffix is null. The file name matches the name defined in the **LogsPath** parameter.
  - **1**: the file is rotated every year. The suffix for the name defined in **LogsPath** is LognameYear(YYYY). For example `"log2021.txt"`.
  - **2**: the file is rotated every month. The suffix for the name defined in **LogsPath** is LognameYearMonth(YYYYMM). For example "`log202107.txt`"
  - **3**: the file is rotated every day. The suffix for the name defined in **LogsPath** is LognameYearMonthDay(YYYYMMDD). For example "`log20210722.txt`"
  - **4**: the file is rotated every hour. The suffix for the name defined in **LogsPath** is LognameYearMonthDayHour(YYYYMMDDhh). For example "`log2021072210.txt`"
  - **5**: the file is rotated every minute. The suffix for the name defined in **LogsPath** is LognameYearMonthDayHourMinute(YYYYMMDDhhmm). For example "`log202107221055.txt`"

Chapter 6

Installing and configuring Cytomic Importer on Linux systems

Cytomic Importer is the application responsible for downloading the events logged by Cytomic SIEMConnect and Panda SIEMConnect for Partners from the Azure infrastructure. These events are stored in log files which, depending on the configured settings, are decompressed by Cytomic Importer and placed in a local or remote folder, or sent to a compatible server (Kafka or Syslog).

CHAPTER CONTENTS

Installation requirements

Required information

For the information required by Cytomic Importer for its correct operation, refer to "**Licenses and required information**" on page **21**.

Operating system and required libraries

The download package contains everything Cytomic Importer requires to work on these distributions:

- Ubuntu 18.04.4 LTS Desktop (64bits)
- Red Hat Enterprise Linux 7.2 Server (64bits)

Required permissions

Cytomic Importer can be run as a command-line program or unattended as a system daemon.

- In daemon mode, the program runs under a user account, but root permissions are required for the administrator to configure it.
- When run as a command-line program, it doesn't require any specific permissions beyond access to the resources it may need; for example, write access to the folder configured to store the downloaded logs.

Firewall configuration

In order for the Cytomic Importer computer to be able to download log files from the Azure infrastructure, all intermediate firewalls must allow network traffic with the following characteristics:

- Access to the URL **https://auth.pandasecurity.com**.
- Access to the URL **https://storage.accesscontrolmngr.pandasecurity.com**.
- Access to the URL **sb:// pac100siemfeeder.servicebus.windows.net**.
- **Communication source**: Cytomic Importer computer.
- **Communication target**: Azure infrastructure.
- **Connection type**: outbound from the customer's network.
- **Layer 3 (transport) protocol**: TLS 1.2.
- **Layer 4 (application) protocol**: HTTPS (port 443), Amqp (ports 5671 and 5672), Amqp WebSockets (port 443).

NTP server

To download the logs stored in the Azure infrastructure, an authentication and authorization process must be completed that involves generating a token. This token is issued with an expiration date to improve the security of the entire process, therefore, the clocks of both communication endpoints must be synchronized. For this reason, it is required that the computer that runs Cytomic Importer have the `ntpd` service (or equivalent) up and running in order to get the time from an NTP server. For more information, refer to **https://www.ntppool.org/en/use.html**.

Installation and configuration

For more information about the origin of the errors that may occur during the installation process, refer to "**Appendix 1: Troubleshooting**" on page **53**.

To install and configure Cytomic Importer, perform the steps below in the listed order:

1. Download and decompress the .GZ file containing the installer: Refer to "**Download the installation package**".
2. (Optional) If needed, modify the execution attribute of the files.
3. Indicate the connection method supported by the IT infrastructure that will host the Cytomic Importer computer: direct connection or through a corporate proxy. Refer to "**Configure the connection method**".
4. Indicate the platform where your Cytomic security products reside. Refer to "**Configure the platform to use**".
5. Enter the credentials of the account used to access the service. Refer to "**Enter the access credentials**".
6. Configure the method to be used to send and store the received logs containing the monitored events. Refer to "**Configure the log storage and send mode**".
7. Update the `configuration.json` file with the new installation settings. Refer to "**Update the configuration.json file**".
8. (optional) Configure Cytomic Importer to run as a daemon. Refer to "**Configuring Cytomic Importer as a daemon**".

Download the installation package

Download the .GZ package of the Linux version of Cytomic Importer from **https://www.cytomic.ai/en/support/id-950031** and decompress it to a folder on your computer. This package contains the following main files:

- `Configuration.json`: contains the program settings. All personal data is stored obfuscated to prevent security leaks.

Cytomic Importer can be run multiple times simultaneously on the same computer. Each Cytomic Importer instance requires a separate configuration file. Refer to "**Configuring multiple instances**".

Modify the execution attribute of the files

For a Linux system to run a program, the execute bit of the file must be turned on. Run the following commands from a command line:

```
$ sudo chmod a+x /#_SiemFeeder_SAMPLEFOLDER#/EventsFeederImporter.Multiplatform.
Host
$ sudo chmod a+x /#_SiemFeeder_SAMPLEFOLDER#/EventsFeederImporter.Multiplatform.
ConfigAssistant
```

The variable `/#_SiemFeeder_SAMPLEFOLDER#/` is the full path to the folder where the uncompressed package resides.

Configuration

This section describes the steps you must take to generate the configuration file required for a single Cytomic Importer instance to run in command-line mode and connect to the Azure infrastructure in order to download logs. All other scenarios are based on this configuration.

To configure Cytomic Importer, you must run the `EventsFeederImporter.Multiplatform.ConfigAssistant` program and answer "Yes" to the question **Do you want to change the configuration settings? [Yes/No]**). A new configuration file that overrides the existing one is generated and the configuration wizard is launched.

Configure the connection method

If the computer is behind a proxy server, answer **Y** to the question **Is Event Importer behind a proxy server? [Yes/No]** You will be prompted to enter the proxy server IP address, as well as the user name and password if the proxy server requires authentication.

The password must be a string of alphanumeric characters, spaces, and symbols, except for: ":", "/", "?", "#", "[", "]", "@", "!", "$", "&", "'", "(", ")","*", "+",",",";","=",","

Access through the configured proxy server is only used to connect to the Azure infrastructure assigned to the customer or MSSP. It is not used to connect to other resources such as the file server, the Kafka server, or the Syslog server.

Configure the platform to use

- Answer the question **Select your platform: [C]urrent, [L]egacy or [W]G Endpoint Security** with **C** (Current platform).

Enter the access credentials

- Enter the email address of the user account used to access the Cytomic EDR console.
- Enter the password. If the account has 2FA enabled, enter the 6-digit OTP code immediately after the password, without spaces.
- Enter the Customer ID specified in the welcome email. After you enter it, Cytomic Importer generates a new access token it uses internally to subscribe to the service and download the generated log files.

To determine if the access account has 2FA enabled, go to the Cytomic EDR management console:

- Click **https://central.cytomic.ai**
- Click the account name in the upper-right corner of the page. A drop-down menu appears.
- Click **Set up my profile**. The **Cytomic Account** page opens**.** This page indicates whether 2FA is enabled or not.

For more information about how to enable 2FA, refer to **https://info.cytomic.ai/central/en/index.htm**

Configure the log storage and send mode

For more information about how to choose the method to be used to store and send the downloaded logs, refer to "**Configuring log storage and forwarding**".

Update the configuration.json file

After it finishes running, the configuration wizard updates, with the entered information, the `configuration.json` file located in the same folder. Then, Cytomic Importer starts downloading the logs stored in the Azure infrastructure.

The `configuration.json` file contains the following data:

- Information about the customer whose logs are downloaded.
- Information about the method selected to send and store the downloaded logs.
- Information about the execution mode (command line or daemon).

Configuring Cytomic Importer as a daemon

Cytomic Importer can run automatically as a process in the background at system startup. In such case, it does not show any messages on screen:

- If you are configuring Cytomic Importer as described in section "**Configuration**", stop the process after the procedure is complete by following these steps:
  - Open a Command Prompt window and run the command `ps ax | grep`

"`EventsFeederImporter.Multiplatform.Host`" to get the process PID.

- Write down the process PID and run the command `kill -9 #PID#`
- Go back to the initial window and press `Control + C`.

- If Cytomic Importer was already configured and is running in command-line mode, press Control + C.
- Edit the `siemfeeder.service` file included in the .GZ package by replacing the line starting with `ExecStart` with the full path to the `EventsFeederImporter.Multiplatform.Host` program.
- Copy the `siemfeeder.service` file to the system directory of the Linux distribution used. The most common paths are:
  - `/lib/systemd/system`
  - `/usr/lib/systemd/system`
- Run the `sudo systemctl enable siemfeeder` command to add the script to the system startup sequence.

Configuring multiple instances

You must configure multiple instances of Cytomic Importer in the following cases:

- If the computer that runs the Cytomic Importer program shows symptoms indicating a lack of resources similar to those described in section "**Hardware sizing recommendations for the Cytomic Importer computer**" on page **26**, we recommend that you install one or more additional instances of the program and run them concurrently.
- If you require that a single computer with Cytomic Importer download logs from more than one customer simultaneously, but you are not using Panda SIEMConnect for Partners.

*To download logs files from multiple customers and centralize all downloads through a single Cytomic Importer instance, use Panda SIEMConnect for Partners. Refer to "**Panda SIEMConnect for Partners architecture**" on page **15**.*

Multiple instances in command-line mode

- Download the latest version of Cytomic Importer from **https://www.cytomic.ai/en/support/id-950031** and decompress it to a separate folder for each customer whose logs you want to download.
- Configure each application independently using the steps described in section "**Configuration**" to install it in command-line mode.
- Run each application independently.

Configuring log storage and forwarding

Cytomic Importer provides several methods to store and forward logs based on the network architecture, the available resources, and the volume of information received from the Azure infrastructure:

- Storing the logs in a local or remote folder.
- Sending the logs to a Kafka server.
- Sending the logs to a Syslog server.

The storage method is configured in the configuration wizard when the following question is shown: **Importer enables you to send received events simultaneously to various channels. Do you want to change the current channel settings? [Yes / No]**. Selecting **Y** deletes the existing storage and forwarding settings (if any) and generates new settings.

Storing logs in a local or remote folder

- First, create the folder where the log files will be stored. This can be a local folder on the computer running Cytomic Importer or a remote, shared drive/resource.
- If you intend to run multiple instances of Cytomic Importer, create a separate folder for each instance. Otherwise, some logs might be lost during collection and storage.
- Select **F** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server.**
- Enter the full folder path for each instance of Cytomic Importer.
- To finish configuring the storage method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]**.

Sending logs to a Kafka server

- Select **K** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server**.
- Enter the IP address or domain name of the Kafka server and the listening port separated by ":".
- Enter the name of the queue/topic that logs will be sent to on the Kafka server.
- Enter the communication protocol that Cytomic Importer will use to send logs to the Kafka server:
  - **[N]one**: press **N** to send logs in unencrypted format.
  - **[S]sl:** press **S** to send logs using SSL encryption.
  - **s[A]slssl:** press **A** to send logs using SASL/SSL encryption.
  - **saslplain[T]ext:** press **T** to send logs using SASL/PLAIN encryption.
- Depending on whether the communication protocol chosen encrypts data or not, you will have to indicate the path of the file that contains the certificate issued by the CA configured on the Kafka

server.

- To finish configuring the delivery method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]?**

Sending logs to a Syslog server

- Select **S** when the following message appears: **Select where you want to deliver received events: [F]ile on disk, [K]afka topic/queue, or [S]yslog server**.
- Select the format configured on the Syslog server for the received logs: **RFC[5]424** or **RFC[3]164**.
- Enter the IP address or domain name of the Syslog server and the listening port separated by ":".
- Select the transport protocol configured on the Syslog server for the received logs: **[T]cp** or **[U]dp**.

> *To make sure all log files sent by Cytomic Importer are received on the Syslog server, we recommended configuring the use of the TCP transport protocol on both ends. Otherwise, in overload situations, the UDP protocol might inadvertently discard logs.*

- Select the secure protocol to use to encrypt communications between the Syslog server and Cytomic Importer. **[N]one** or **Tls1.[2]**.
- Select the end-of-message marker configured on the Syslog server for the received logs: **[C]r, [L]f, c[R]lf.**

> *If the transport protocol chosen is UDP, no end-of-line marker is used.*
>
> *If the transport protocol chosen is TCP or TLS, the Null end-of-line marker is always used.*

- If the communication protocol chosen encrypts data, indicate the location of the certificate issued by the CA configured on the Syslog server.
- To finish configuring the delivery method, answer **No** to the question **Do you want to configure another delivery channel? [Yes / No]**.

Copying downloaded logs to multiple locations

Cytomic Importer allows logs to be downloaded to multiple locations simultaneously. Each log downloaded this way will be removed from the download queue if at least one of the target locations is correctly updated. Therefore, should errors occur during log collection, the different locations might end up with a different number of logs at the end of the process. To implement the ability to download logs to multiple locations, the 'channels' feature is used. A channel indicates the storage type used by Cytomic Importer and its settings.

To configure Cytomic Importer, follow the steps below:

- Install Cytomic Importer as described in section "**Configuration**".

- Stop Cytomic Importer as described in section "**Running and stopping the program**".
- Add a new channel to the existing collection in file `configuration.json`. The `channels` parameter syntax is the following:

```
"Channels": [{ channel 1 parameters} , {channel 2 parameters}, …]
```

- Each channel indicates the storage type used to store logs and its associated settings.

Below you can find and example of a Cytomic Importer configuration with two channels: the first channel leaves logs in the `Log1` folder, and the second channel leaves them in folder `Log2`:

```
"Channels": [{
            "Type": "LocalDisk",
            "Name": "LD1",
            "Configuration": {
                        "fullPath":
"D:\\\\SIEMFeeder\\\\EventFeederImporter 1.0.3 Pro\\\\Log1",
                        "filesSplitFormat": "1m",
                        "filesSizeLimitInBytes": 102400,
                        "directoryMaxSizeInMb": 1024
                    }
            }, {
            "Type": "LocalDisk",
            "Name": "LD2",
            "Configuration";{
            "fullPath:
"D:\\\\SIEMFeeder\\\\EventFeederImporter 1.0.3 Pro\\\\Log2",
                        "fileSplitFormat":"1m",
                        "filesSizeLimitInBytes": 102400,
                        "directoryMaxSizeInMb": 1024
                    }
            },]
```

Running and stopping the program

In command-line mode

- To start Cytomic Importer, double-click the `EventsFeederImporter.Multiplatform.Host` file or run it from the command line.
- If you are configuring Cytomic Importer as described in section "**Configuration**", to stop the process after the procedure follow these steps:
  - Open a Command Prompt window and run the command `ps ax | grep "EventsFeederImporter.Multiplatform.Host"` to get the process PID.
  - Write down the process PID and run the command `kill -9 #PID#`
  - Go back to the initial window and press `Control + C`.
- If Cytomic Importer was already configured and is running in command-line mode, press Control + C to stop it.

In daemon mode

- To start Cytomic Importer, run `sudo service siemfeeder start` from the command line.
- To stop Cytomic Importer, run `sudo service siemfeeder stop` from the command line.
- To get the running status of Cytomic Importer, run `systemctl status siemfeeder.service` from the command line.

Changing the settings

To modify the Cytomic Importer settings, follow these steps:

- Stop the process if it is running.
  - If Cytomic Importer is running in command-line mode, press Control + C.
  - If Cytomic Importer is running in daemon mode, run `sudo service siemfeeder stop` from the command line.
- Run the `EventsFeederImporter.Multiplatform.ConfigAssistant` program from the command line or by double-clicking on it.
- Answer **Y** to the question **Do you want to change the configuration settings? [Yes / No]**
- Complete the configuration wizard steps.
- Start Cytomic Importer.
  - If Cytomic Importer was running in command-line mode, double-click the `EventsFeederImporter.Multiplatform.Host` file or run it from the command line.
  - If Cytomic Importer was running in daemon mode, run `sudo service siemfeeder Start` from the command line.

Manually editing the Cytomic Importer settings

The `configuration.json` file is in the same folder as Cytomic Importer and contains its execution parameters.

The first time it is run, the program asks for certain authentication information that it later dumps into the configuration file. Once the installation and execution processes are completed, you'll be able to edit the configuration file in order to fine-tune its parameters.

Every time you modify the `configuration.json` *file, you must stop and start the Cytomic Importer process so that it applies the changes made to the file. Refer to "***Running and stopping the program***".*

The `configuration.json` file follows the JSON syntax.

Parameters related to log files containing events

The parameters that determine how Cytomic Importer behaves in order to generate the log files that contain the events logged by Cytomic SIEMConnect are:

- **fullPath**: absolute path to the folder the log files are downloaded to.
- **fileSizeLimitInBytes**: maximum allowed size of the log files.
- **directoryMaxSizeInMB**: maximum allowed size of the folder that stores the log files. When the maximum size is reached, 10 percent of the oldest files are deleted.
- **FileSplitFormat**: rotation interval of the log files. The file name contains the year(yyyy), month(MM), day (dd), hour(HH), and minute (mm) when the file was created.
  - "**1h**" or **empty**: yyyyMMdd-HH format. A file is generated every hour.
  - "**1m**": yyyyMMdd-HHmm format. A file is generated every minute.
  - "**5m**": yyyyMMdd-HHmm format. A file is generated every 5 minutes.
  - "**10m**": yyyyMMdd-HHmm format. A file is generated every 10 minutes.
  - "**15m**": yyyyMMdd-HHmm format. A file is generated every 15 minutes.
  - "**30m**": yyyyMMdd-HHmm format. A file is generated every 30 minutes.
- **Channels**: indicates the characteristics of the channel used to download the log files.
- **Type**: storage type used in the channel.
- **Name**: channel name.
- **Configuration**: channel settings (fullPath, fileSplitFormat, fileSizeLimitInBytes, directoryMaxSizeInMB)).

Parameters related to the execution log

All operations executed by Cytomic Importer are logged to text files stored in the `Log` folder inside the same folder that contains the program.

*Refer to "**Appendix 1: Troubleshooting**" on page **53** for a description of the errors Cytomic Importer can generate.*

The parameters that determine how Cytomic Importer behaves in order to generate the log files that log the actions it takes are:

- **LogsPath**: absolute or relative path and file name. The backlash character ("\") must be escaped. For example `".\\log\\log.txt"`.
- **LogFileSizeLimitKBytes**: rotates the log file when it reaches a certain size in Kbytes, adding the suffix "-SequenceNumber" to it. For example `"log-3.txt"`.
- **LogRetainedFileCountLimit**: indicates the maximum number of files that Cytomic Importer stores on the storage device. When this number is reached, Cytomic Importer deletes the oldest file.
- **Interval**: rotation interval of the log files:

- **0**: No rotation. The suffix is null. The file name matches the name defined in the **LogsPath** parameter.
- **1**: the file is rotated every year. The suffix for the name defined in **LogsPath** is LognameYear(YYYY). For example "`log2021.txt`".
- **2**: the file is rotated every month. The suffix for the name defined in **LogsPath** is LognameYearMonth(YYYYMM). For example "`log202107.txt`"
- **3**: the file is rotated every day. The suffix for the name defined in **LogsPath** is LognameYearMonthDay(YYYYMMDD). For example "`log20210722.txt`"
- **4**: the file is rotated every hour. The suffix for the name defined in **LogsPath** is LognameYearMonthDayHour(YYYYMMDDhh). For example "`log2021072210.txt`"
- **5**: the file is rotated every minute. The suffix for the name defined in **LogsPath** is LognameYearMonthDayHourMinute(YYYYMMDDhhmm). For example "`log202107221055.txt`"

Chapter 7

Appendix 1: Troubleshooting

Below we describe some of the most commonly encountered issues and their solution:

Symptom/error	Cause	Solution
Error initializing .NET Framework	Cytomic Importer cannot find .NET Framework 4.6.1 or later on the administrator's computer.	Make sure .NET Framework 4.6.1 is installed. Go to **https://www.microsoft.com/en-us/download/details.aspx?id=49981** to download it.
invalid_redirect_uri **unrecognized_client_id** **unsupported_scope**	Customer ID not recognized	Make sure the customer is correctly registered in the Cytomic SIEMConnect service. Make sure the email address used to log in to the Cytomic EDR management console is correctly entered in Cytomic Importer
unrecognized_client_secret **unsupported_grant_type** **invalid_grant**	Customer password not recognized	Make sure the Cytomic account used to access the Cytomic SIEMConnect service has the Full Control role assigned. Run Cytomic Importer and answer 'Y' to the question '**Do you want to change the configuration settings**?' Re-enter the password. Check that the computer on which Cytomic Importer runs has the time synchronized with NTP or a similar service. Check that Windows Time Service is running.
unauthorized_client **unsupported_response_type** **invalid_scope** **access_denied** **invalid_request**	The authentication information is correct, but there is a problem downloading data	Contact Cytomic Support department.

Table 7.1: Potential problems and solutions

Symptom/error	Cause	Solution
temporarily_unavailable **server_error**	The Cytomic SIEMConnect service is temporarily unavailable due to technical issues	Run Cytomic Importer again after some time. Check the email account used to register the service. If the error persists, an email will be sent to the administrator explaining the reasons for the service stop and the available options.

Table 7.1: Potential problems and solutions

Chapter 8

Appendix 2: Security Architecture

This chapter deals with the AAA-based (Authentication, Authorization, and Access) security architecture implemented in Cytomic SIEMConnect as well as the encryption of all communications between the Cytomic Importer software and all the other components that make up the solution.

CHAPTER CONTENTS

AAA security architecture overview

Security architecture: Components

Figure shows the components responsible for authenticating customers and granting them access to the platform resources required to download the log files that contain the information collected from the organization's IT network.

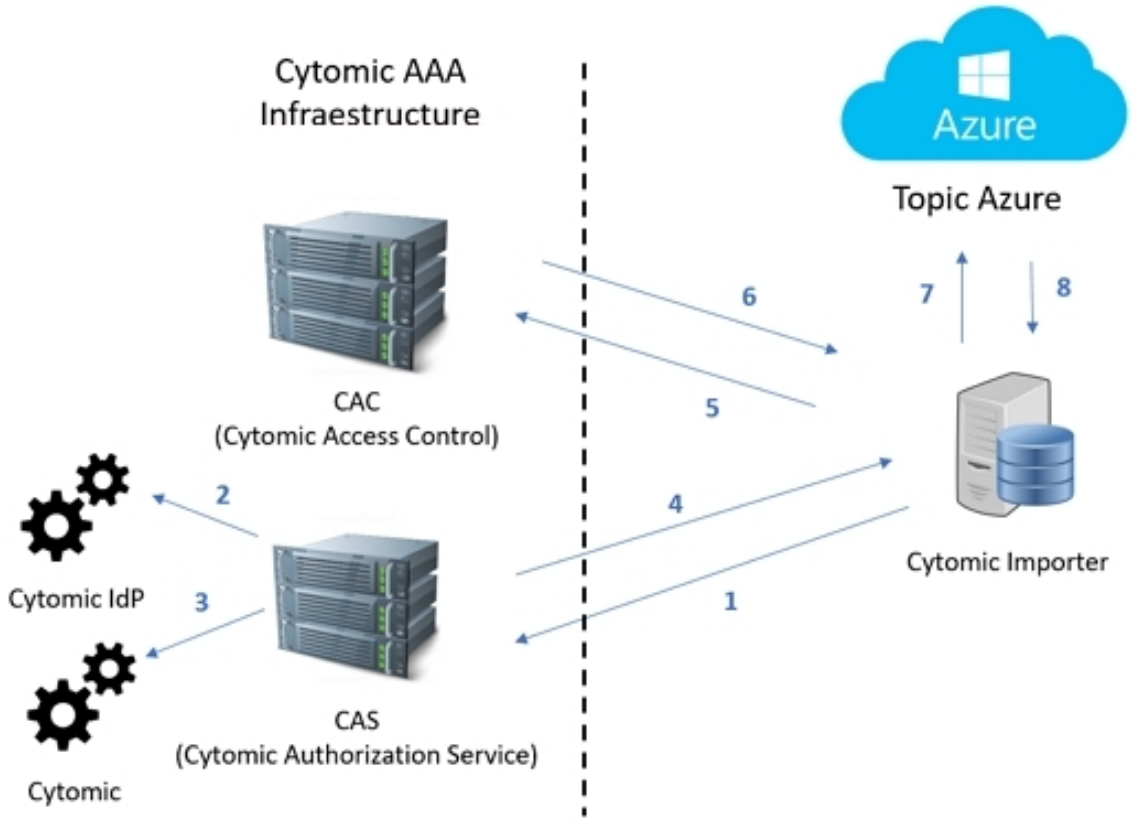

Figure 8.1: AAA security architecture overview

- **Cytomic Importer**: program provided by Cytomic and designed to collect the log files stored on the Azure platform.
- **Azure Topic**: a queue-type resource generated on the Azure platform. It stores the log files received from Panda Security with the information collected from the organization's IT network.
- **CAS (Cytomic Authorization Service)**: service that authenticates and authorizes access to the Azure topic. It receives, from Cytomic Importer, the credentials assigned to the customer when purchasing the service, and returns to it an access token and a refresh token.
- **CAC (Cytomic Access Control)**: service that enables Cytomic Importer to access the Azure topic provisioned to the customer. It receives the refresh token from Cytomic Importer and returns a shared access signature (SAS) key.
- **Cytomic IdP (Identity Provider)**: service that authenticates the sent credentials.

- **Cytomic**: service that authorizes access to Cytomic SIEMConnect.

Initial message exchange

To access the Cytomic SIEMConnect service securely, an initial message exchange must take place between the Cytomic Importer computer and Cytomic SIEMConnect. This exchange must take place successfully; otherwise, it won't be possible to access the information published in the Azure topic.

Below is a diagram showing the message flow established the first time that Cytomic Importer is run (numbered based on 8.1). This message flow must be established every time the user is removed from the system or is unassigned the Full Control role assigned via Cytomic.

1. **Cytomic Importer** sends the credentials (email address and password) assigned to the customer**.**
2. **Authentication Phase**: the CAS service connects to the Cytomic IdP service to validate the credentials.
3. **Authorization Phase**: the CAS service connects to the Cytomic service to check whether the customer has access to the Cytomic SIEMConnect service

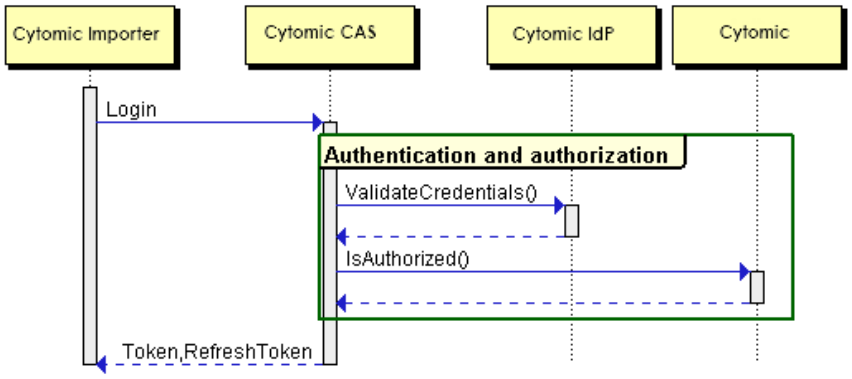

Figure 8.2: Steps 1 to 4 in the initial message exchange

4. The CAS service generates and delivers an access token and a refresh token to Cytomic Importer .
5. **Cytomic Importer** sends the refresh token to the PAC service.
6. **Access Phase**: the CAC service generates a shared access signature (SAS) key.
7. **Access to the topic**: Cytomic Importer accesses the assigned topic using the SAS key.
8. **Cytomic Importer** receives the logs from the subscribed topic.

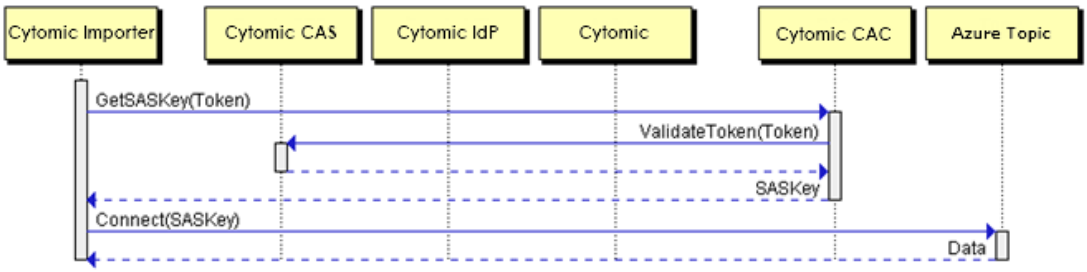

Figure 8.3: Steps 5 to 8 in the initial message exchange

Subsequent message exchange

Cytomic Importer uses the refresh token to obtain the SAS key. Both the token and the SAS key have an expiration date and are short lived for security reasons. As soon as the refresh token expires, Cytomic Importer will generate the following alternative message flow:

1. Cytomic Importer asks the PAS service for a new refresh token. To do that, it sends the access token that was assigned to it during the above-mentioned initial flow.
2. With the new refresh token, Cytomic Importer asks the PAC service for a new SAS key.
3. With the new SAS key, Cytomic Importer connects to the Azure topic and continues collecting log files.

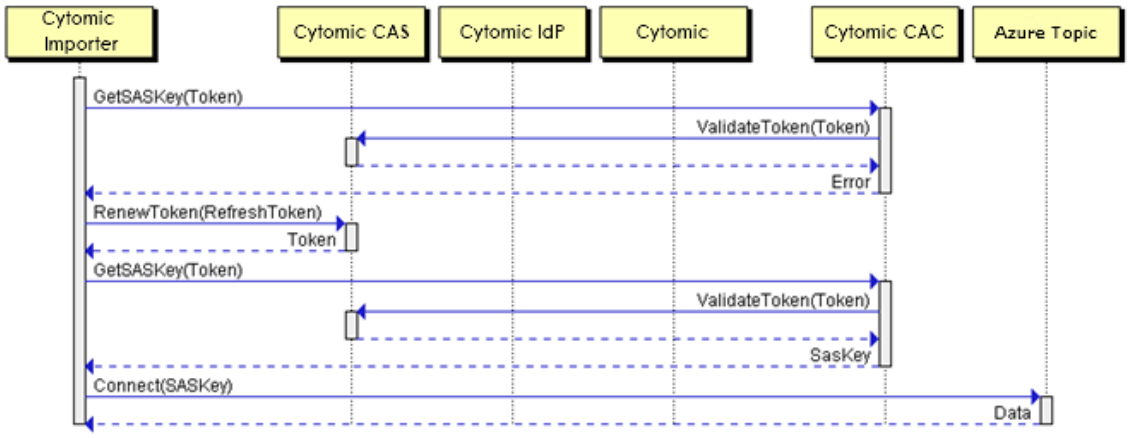

Figure 8.4: Message flow when the refresh token expires

Communication characteristics

AAA communication encryption

All communications for requesting and sending tokens are encrypted with HTTPS protocol SSL SHA256 – G3.

Lifetime of the tokens assigned by Cytomic SIEMConnect

- **CAS refresh token**: 14 days
- **CAS access token**: 20 minutes
- **SAS key**: 1 day

Cytomic Importer uses the refresh token to access the Azure topic. Once the refresh token expires, a new access token will be generated containing the account details entered in the Cytomic Importer program. In addition to this, a new refresh token will also be generated for Cytomic Importer to continue accessing the Azure topic.

Even if the account used when configuring the service is no longer available or doesn't have the Full Control role assigned to it, the customer will be able to continue accessing the service provided the refresh token has not expired (maximum lifetime: 14 days). If the refresh token expires, it won't be possible to generate a new refresh token and access will be denied.

Encrypted communications for downloading log files

All communications established for downloading log files are encrypted with the TLS/SSL and SASL protocols