ASGARD Analysis Cockpit v4 Manual

Nextron Systems

Apr 02, 2024

CONTENTS

Welcome to Nextron System's Manual for the ASGARD Analysis Cockpit v4.

Note: If you are still using an older version of the Analysis Cockpit, please click here to see the older version of the documentation.

Analysis Cockpit is the central platform for analyzing THOR events and SIGMA matches generated by ASGARDs real time agents.

It can be used in an environment where scans results can be automatically collected from ASGARD Management Centers or environments in which THOR is executed by scripts or any other 3rd party solution.

In the following chapters we will describe how the Analysis Cockpit works, how to install the system, and how to use it. Additional information regarding troubleshooting, known issues, or general administrative tasks can also be found.

CHAPTER

ONE

REQUIREMENTS

In this chapter we will go over the requirements needed to get your Analysis Cockpit set up correctly. Please follow those steps carefully and don't skip anything, as you might encounter problems during or after the installation.

1.1 Introduction

Analysis Cockpit is the central platform for analyzing THOR events and SIGMA matches generated by ASGARDs real time agents.

It can be used in an environment where scans results can be automatically collected from ASGARD Management Centers or environments in which THOR is executed by scripts or any other 3^{rd} party solution.

It is available as a virtual appliance on VMWare and also as a dedicated hardware appliance.

While THOR can also be seen or used as hunting solution THOR is optimized to avoid false negatives – meaning optimized to not miss an indicator of compromise. On the other side this clearly leads to more anomalies and false positives being reported.

In a scenario where you scan your infrastructure frequently you would either be seeing the same anomalies again and again or you would need to create many rules to filter out these anomalies in order to save analysis time.

Analysis Cockpit is designed to facilitate this process and help you generate these rules automatically, so that you can set your baseline-filters after the first scan. After setting the first baseline it is now easy to focus on relevant Alerts and Warnings as only differences between the first and second scans are shown.

Analysis Cockpit comes with an integrated and highly configurable ticketing system that helps organizing your analysis workflow and – if required – interfaces to your existing ticketing system through a configurable connector. If ASGARD's Bifrost_2 Service is used to collect suspicious samples, the Analysis Cockpit can submit the samples to various Sandbox systems and include the report in the graphical frontend.

Furthermore, Analysis Cockpit comes with a rule-based alert forwarding and SIEM integration that makes it easy for your organization to react quickly on new incidents. For organizations or projects where a SIEM system is not available, Analysis Cockpit features a separate notification section to deal with alerts and notifications you would normally process in a SIEM system.

The following document describes requirements, the installation process and best practices to group, classify and dispatch events for further analysis.

1.2 Hardware Requirements

There are a few things to consider, before you start with the installation.

If you install on VMWare, the minimum requirements for the virtual machine are as follows:

- System memory: 16 GB
- Hard disk: 200 GB
- CPU cores: 2

The disk size of 200 GB is fine in scenarios where you import only Alerts and Warnings into the Cockpit, scan less than 1.000 systems on a weekly basis and want to keep the logs for less than one year. If you also import Notices and Info messages for these 1.000 servers, we recommend a disk size of at least 500 GB.

The Analysis Cockpit does not have any filters on which type of events will be imported into the database. This has to be controlled by chaning your scan parameters in your ASGARD Management Center. To change the log level, you can use the `--reduced` parameter for all of your scans.

Hint: `--reduced` - Reduced output mode - only warnings, alerts and errors will be printed.

For an Installation of up to 20.000 endpoints the following specifications are recommended:

- System memory: 32 GB
- Hard disk: 2 TB SSD
- CPU cores: 4

1.3 Network Requirements

The Analysis Cockpit and other systems which will have to communicate with each other, need the following ports opened within the network. For a detailed and up to date list of our update and licensing servers, please visit https://www.nextron-systems.com/resources/hosts/.

The Analysis Cockpit requires the following open ports (incoming).

1.3.1 From Management Workstation to Analysis Cockpit

Description	Ports
Administrative Web Interface	443/tcp
Command Line Access	22/tcp

1.3.2 From Analyst Workstation to Analysis Cockpit

Description	Ports
Administrative Web Interface	443/tcp

1.3.3 From ASGARD Management Center to Analysis Cockpit

Description	Ports
Syslog Forwarding	514/tcp, 514/udp
Asset Synchronization	7443/tcp

1.3.4 From Analysis Cockpit to SIEM (optional)

Description	Ports
Syslog Forwarding	514/tcp, 514/udp

1.3.5 From Analysis Cockpit to the Internet

The Analysis Cockpit is configured to retrieve updates from the following URLs:

- https://update-301.nextron-systems.com

A proxy system should be configured to allow access to these URLs without TLS/SSL interception (Analysis Cockpit uses client-side SSL certificates for authentication). It is possible to configure a proxy server, username and password during the setup process of the Analysis Cockpit platform. It only supports BASIC authentication, not NTLM Authentication.

1.3.6 From Analysis Cockpit to Sandbox Systems (optional)

Depending on the Sandbox system and your individual configuration.

Description	Ports
Sandbox (typically)	443/tcp, 8080/tcp

1.3.7 Time Synchronization

Analysis Cockpit tries to reach the public Debian time servers by default.

Server	Port
0.debian.pool.ntp.org	123/udp
1.debian.pool.ntp.org	123/udp
2.debian.pool.ntp.org	123/udp

The NTP server configuration can be changed in the settings.

1.3.8 DNS

Analysis Cockpit needs to be able to resolve internal and external IP addresses.

> **Warning:** Please make sure that you install your Analysis Cockpit with a `domain name` (see *Network Configuration*). If you do not set the domain name and install the ASGARD package, you will have problems connecting your ASGARD(s) to the Analysis Cockpit.
>
> All components you install should have a proper domain name configured to avoid issues further during the configuration.

1.3.9 Internet Access during Installation

The Analysis Cockpit installer requires Internet access during the setup. The installation process will fail if required packages cannot be loaded from https://update3.nextron-systems.com

SSL/TLS Interception

The installation and update processes do not accept an unknown but valid SSL/TLS certificate presented by an intercepting entity and therefore don't support SSL/TLS interception.

Since our products are usually used in possibly compromised environments, the integrity of our software and update packages has highest priority.

1.3.10 Architecture Overview

The following image shows an architecture overview with all products and their communication relationships.

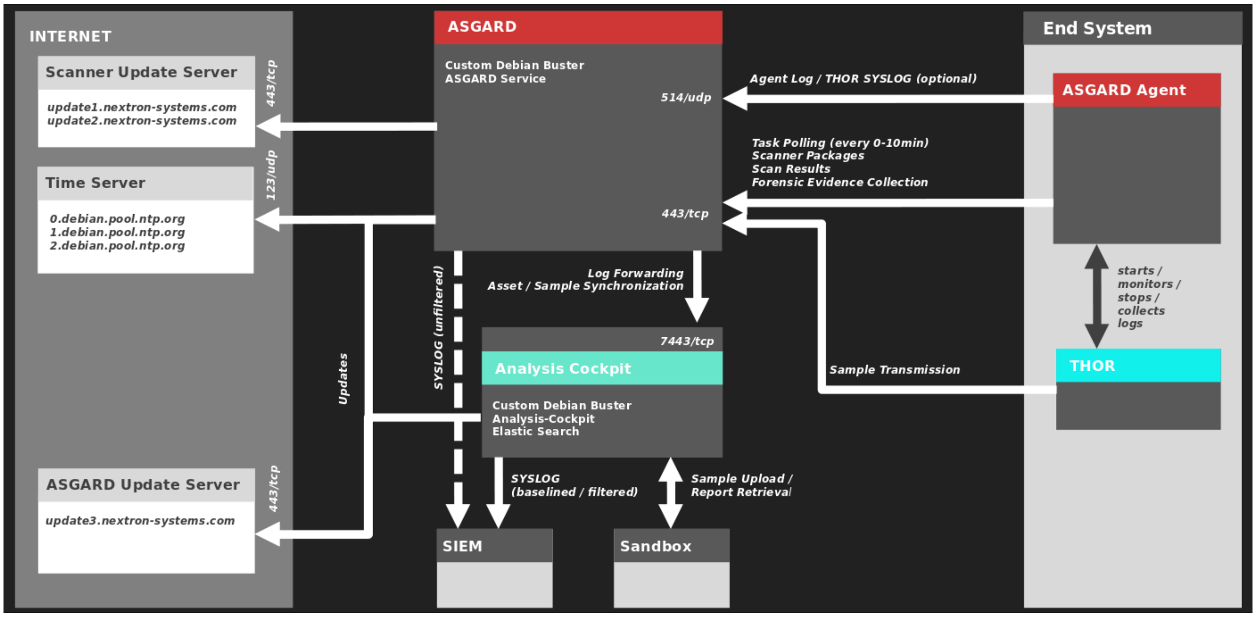

Fig. 1: Full Architecture

1.4 Verify the Downloaded ISO (Optional)

You can do a quick hash check to verify that the download was not corrupted. We recommend to verify the downloaded ISO's signature as this is the cryptographically sound method.

The hash and signature file are both part of the ZIP archive you download from our portal server.

1.4.1 Via Hash

Extract the ZIP and check the sha256 hash:

On Linux

```
user@unix:~/nextron-universal-installer$ sha256sum -c nextron-universal-installer.iso.
↪sha256
nextron-universal-installer.iso: OK
```

or in Windows command prompt

```
C:\temp\nextron-universal-installer>type nextron-universal-installer.iso.sha256
efccb4df0a95aa8e562d42707cb5409b866bd5ae8071c4f05eec6a10778f354b  nextron-universal-
↪installer.iso
C:\temp\nextron-universal-installer>certutil -hashfile nextron-universal-installer.iso␣
↪SHA256
SHA256 hash of nextron-universal-installer.iso:
efccb4df0a95aa8e562d42707cb5409b866bd5ae8071c4f05eec6a10778f354b
CertUtil: -hashfile command completed successfully.
```

or in Powershell

```
PS C:\temp\nextron-universal-installer>type .\nextron-universal-installer.iso.sha256
efccb4df0a95aa8e562d42707cb5409b866bd5ae8071c4f05eec6a10778f354b  nextron-universal-
↪installer.iso
PS C:\temp\nextron-universal-installer>Get-FileHash .\nextron-universal-installer.iso

Algorithm       Hash                                                                   ␣
↪Path
---------       ----                                                                    --
↪--
SHA256          EFCCB4DF0A95AA8E562D42707CB5409B866BD5AE8071C4F05EEC6A10778F354B       ␣
↪C:\Users\user\Desktop\asgard2-installer\nextron-universal-installer.iso
```

1.4.2 Via Signature (Recommended)

Extract the ZIP, download the public signature and verify the signed ISO:

On Linux

```
use@unix:~/temp$ wget https://www.nextron-systems.com/certs/codesign.pem
use@unix:~/temp$ openssl dgst -sha256 -verify codesign.pem -signature nextron-universal-
↪installer.iso.sig nextron-universal-installer.iso
Verified OK
```

or in Powershell

```
PS C:\temp\nextron-universal-installer>Invoke-WebRequest -Uri https://www.nextron-
↪systems.com/certs/codesign.pem -OutFile codesign.pem
PS C:\temp\nextron-universal-installer>"C:\Program Files\OpenSSL-Win64\bin\openssl.exe"␣
↪dgst -sha256 -verify codesign.pem -signature nextron-universal-installer.iso.sig␣
↪nextron-universal-installer.iso
Verified OK
```

Note: If `openssl` is not present on your system you can easily install it using winget: `winget install openssl`.

1.5 Other Optional Requirements

1.5.1 Usage of a Reverse Proxy

If you are planing to make the Analysis Cockpit available through a reverse proxy, see *Reverse Proxy to access the Analysis Cockpit*.

CHAPTER

TWO

SETUP GUIDE

In this chapter we will show and example installation with VMware ESXi and the provided ISO image to install the Analysis Cockpit. Please pay good attention to the setup during the Debian Installer, since this contains important steps which might break your installation!

2.1 Create a New ESX VM and Mount the ISO

Create a new VM with your virtualization software. In this case, we will use VMWare ESX managed through a VMWare VCenter.

The new VM must be configured with a Linux base system and Debian GNU/Linux 10 (64 bits) as target version. It is recommended to upload the Nextron Universal Installer ISO to an accessible datastore and mount the same to your newly created VM.

Please make sure to select a suitable v-switch or physical interface that reflects the IP address scheme you are planning to use for the new Analysis Cockpit. Only use one Hard Disk for the installation.

2.2 Navigate through the Installer

Start the installation confirming the only available option in the boot loader screen.

The installer then loads the additional components from the ISO image and lets you select a location and language.

Note: If DHCP is available, network parameters will be configured automatically. Without DHCP, ASGARD proceeds with the manual network configuration dialogue.

2.3 Network Configuration

The next step prompts for a hostname for the device. After entering a hostname and clicking `Continue`, it also prompts for the Domain Name. After this Information is submitted, the Installer tries to get network configurations from a DHCP-Server. If there is none to be found, it will prompt for a static IP-Configuration.

Enter the IP address that Analysis Cockpit should use and optimally directly add a netmask in CIDR notation. (see below) If you don't append the netmask, you'll be asked for a network mask in the following dialogue.

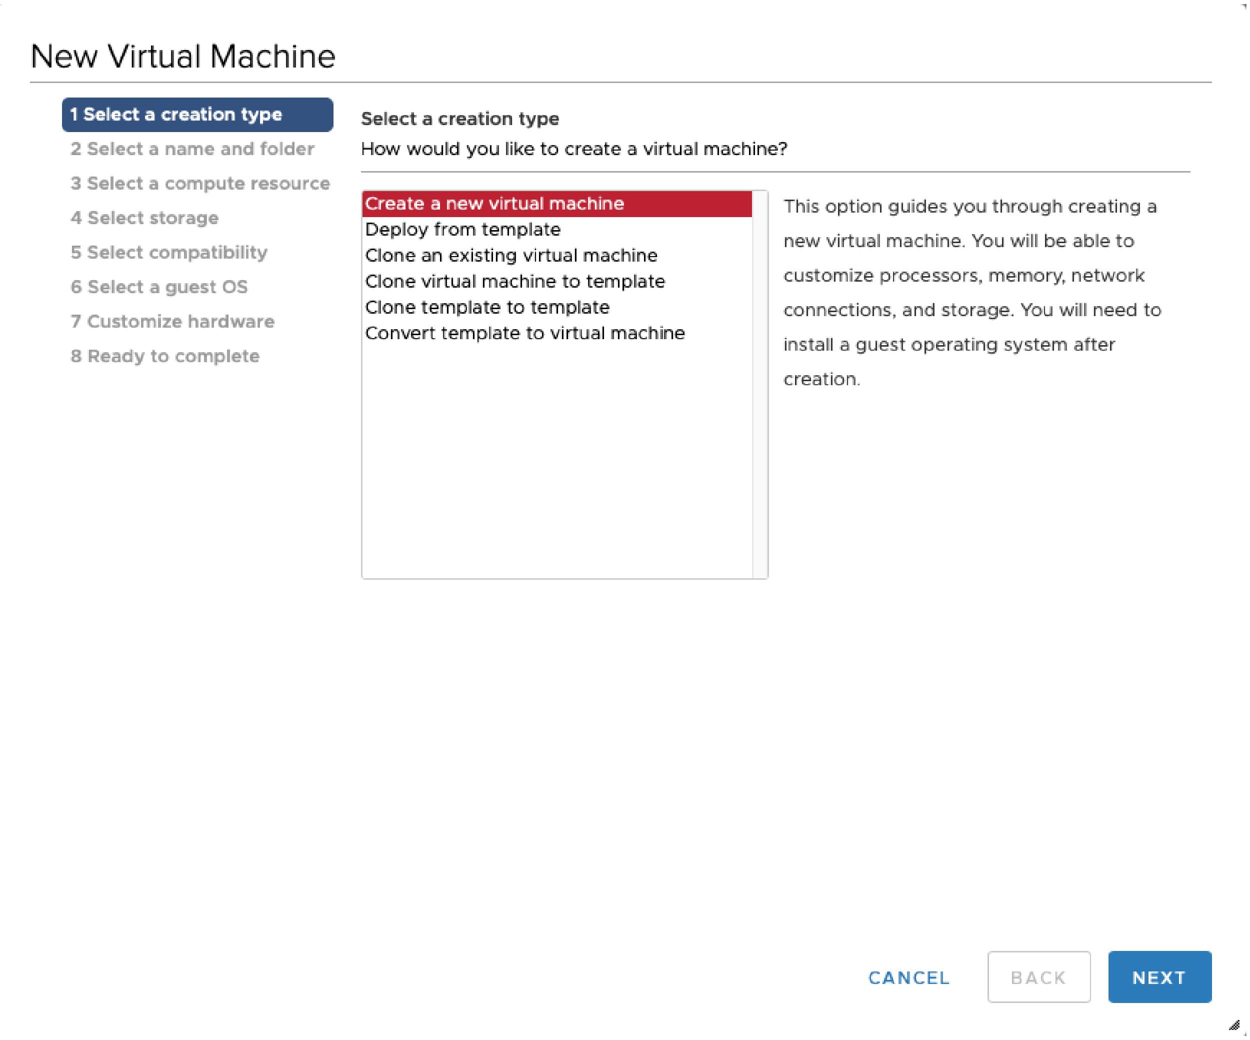

Select a creation type

How would you like to create a virtual machine?

- Create a new virtual machine
- Deploy from template
- Clone an existing virtual machine
- Clone virtual machine to template
- Clone template to template
- Convert template to virtual machine

This option guides you through creating a new virtual machine. You will be able to customize processors, memory, network connections, and storage. You will need to install a guest operating system after creation.

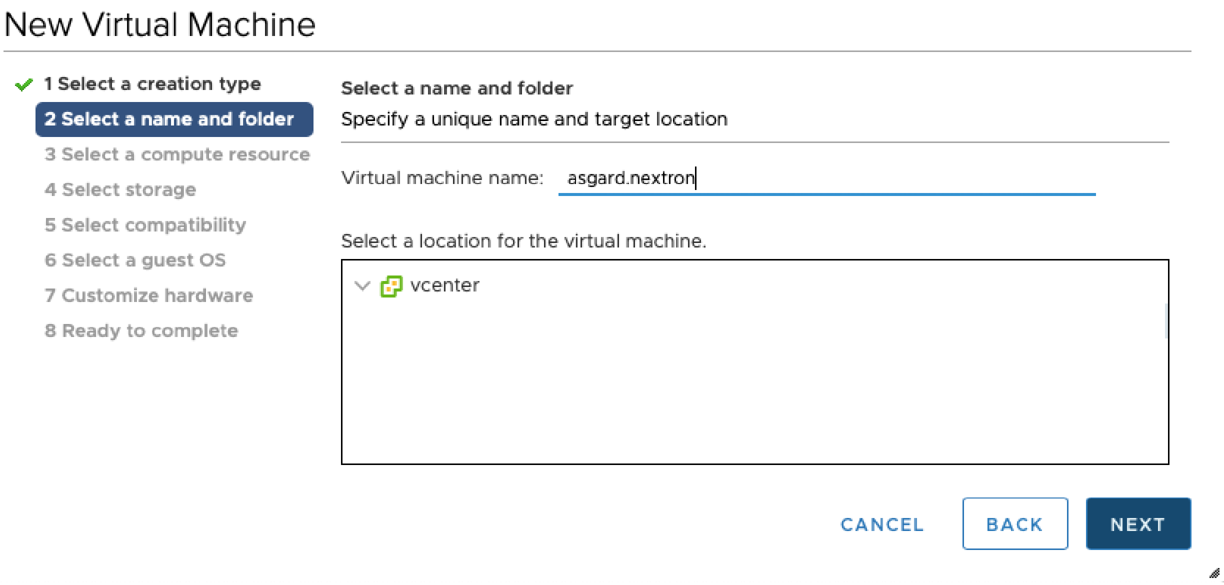

Select a name and folder

Specify a unique name and target location

Virtual machine name: asgard.nextron

Select a location for the virtual machine.

vcenter

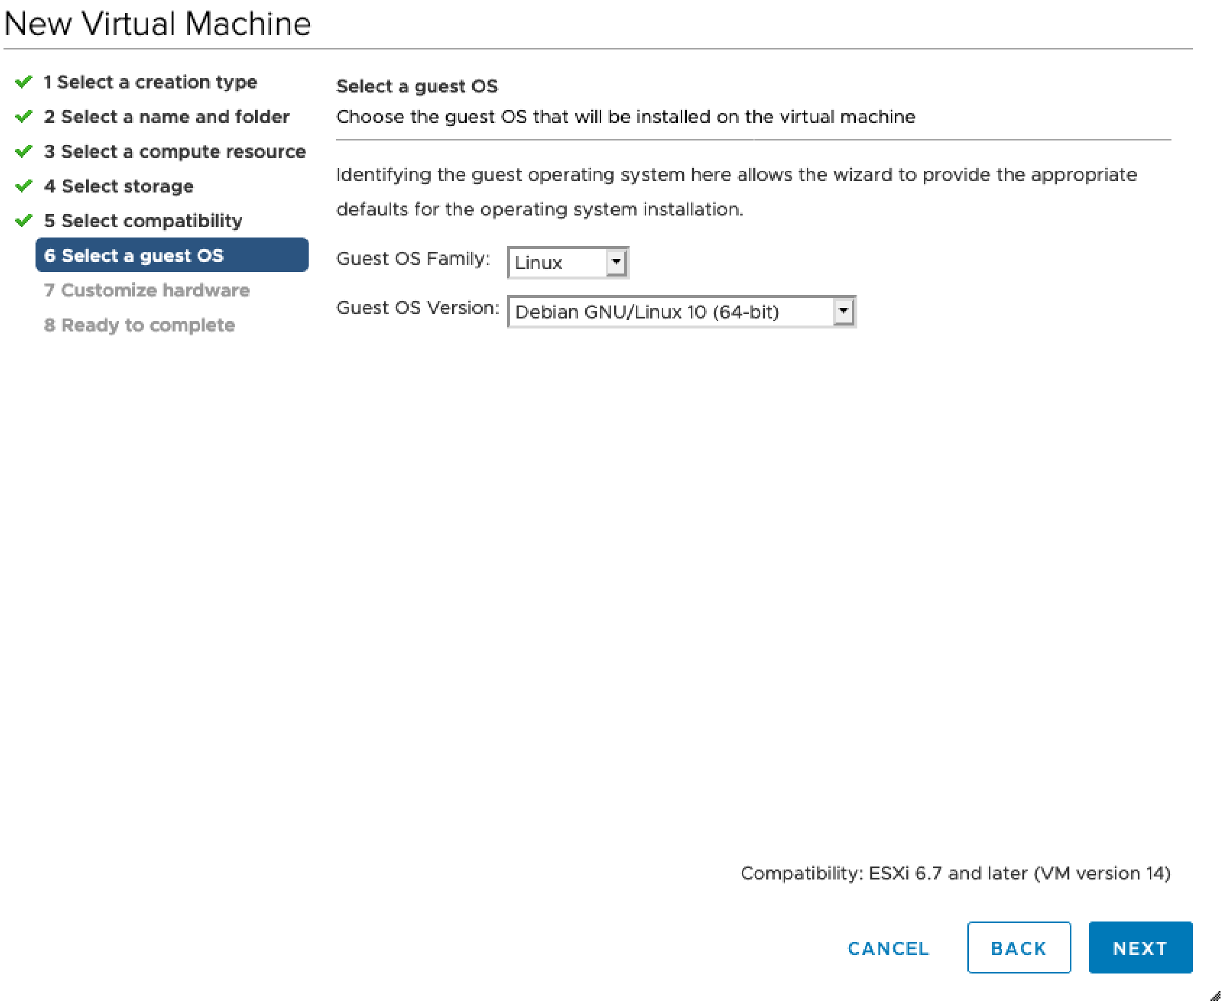

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select a guest OS

Choose the guest OS that will be installed on the virtual machine

Identifying the guest operating system here allows the wizard to provide the appropriate defaults for the operating system installation.

Guest OS Family: Linux

Guest OS Version: Debian GNU/Linux 10 (64-bit)

Compatibility: ESXi 6.7 and later (VM version 14)

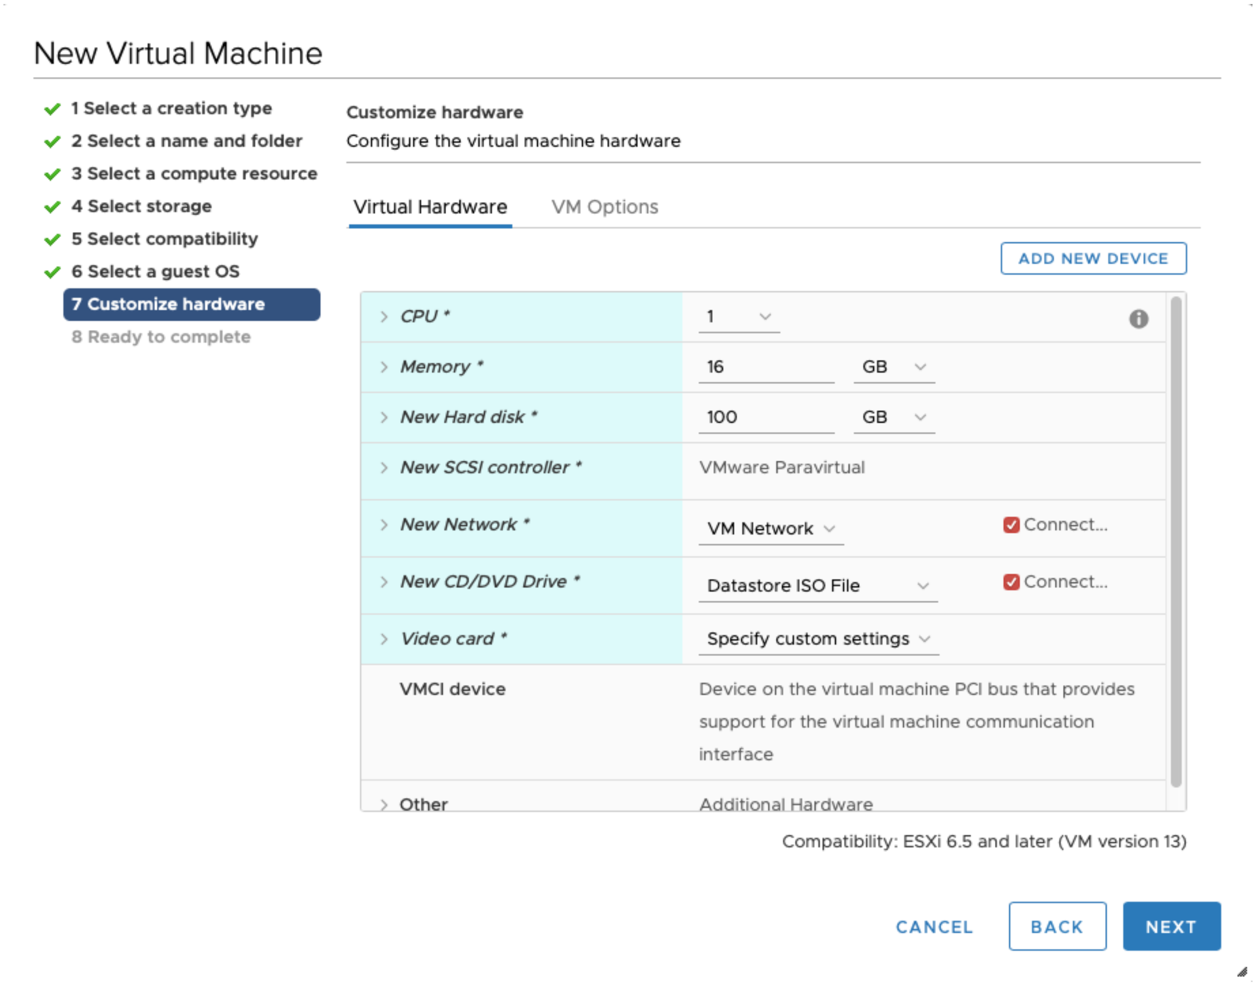

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- ✓ 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Customize hardware

Configure the virtual machine hardware

CPU *	1
Memory *	16 GB
New Hard disk *	100 GB
New SCSI controller *	VMware Paravirtual
New Network *	VM Network ☑ Connect...
New CD/DVD Drive *	Datastore ISO File ☑ Connect...
Video card *	Specify custom settings
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface
Other	Additional Hardware

Compatibility: ESXi 6.5 and later (VM version 13)

Fig. 1: Create a new virtual Machine

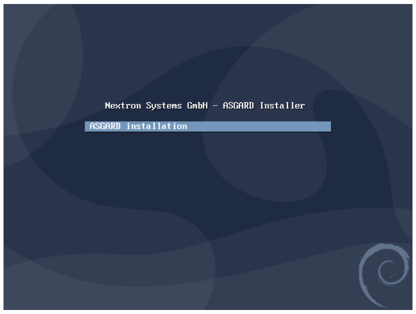

Nextron Systems GmbH - ASGARD Installer

ASGARD installation

Fig. 2: Starting the installation

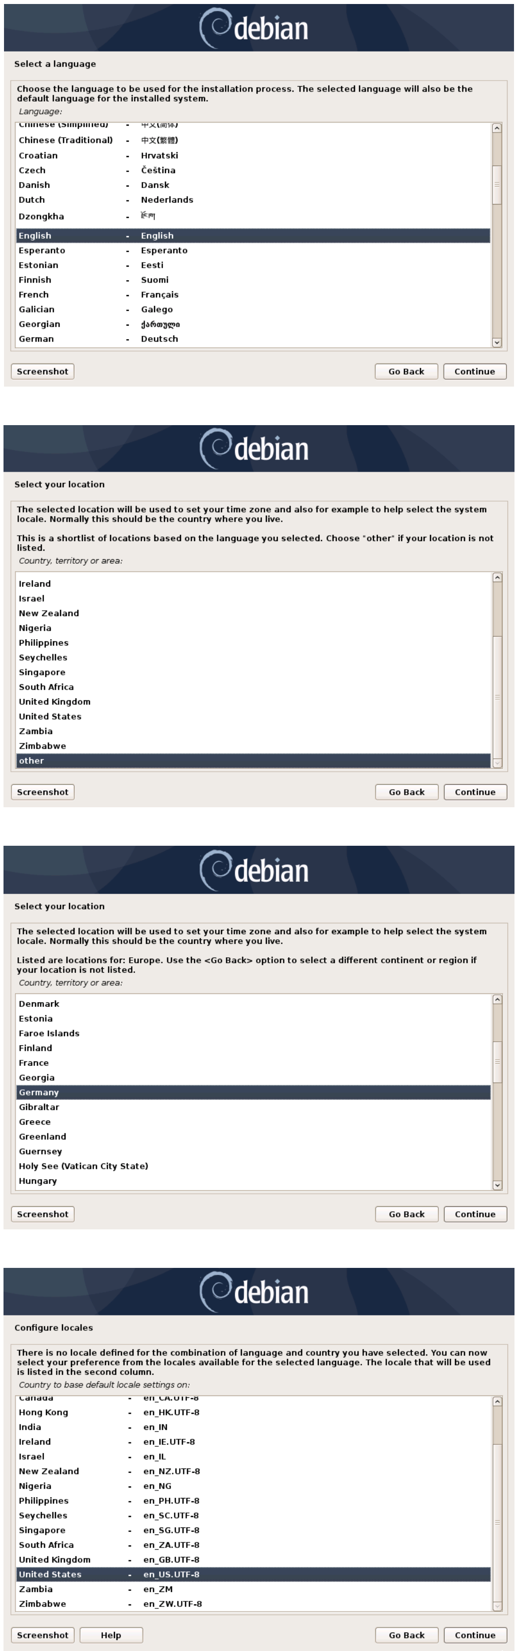

Select a language

Choose the language to be used for the installation process. The selected language will also be the default language for the installed system.

Language:

Chinese (Simplified)	-	中文(简体)
Chinese (Traditional)	-	中文(繁體)
Croatian	-	Hrvatski
Czech	-	Čeština
Danish	-	Dansk
Dutch	-	Nederlands
Dzongkha	-	རྫོང་ཁ
English	-	**English**
Esperanto	-	Esperanto
Estonian	-	Eesti
Finnish	-	Suomi
French	-	Français
Galician	-	Galego
Georgian	-	ქართული
German	-	Deutsch

Select your location

The selected location will be used to set your time zone and also for example to help select the system locale. Normally this should be the country where you live.

This is a shortlist of locations based on the language you selected. Choose "other" if your location is not listed.

Country, territory or area:

Ireland
Israel
New Zealand
Nigeria
Philippines
Seychelles
Singapore
South Africa
United Kingdom
United States
Zambia
Zimbabwe
other

Select your location

The selected location will be used to set your time zone and also for example to help select the system locale. Normally this should be the country where you live.

Listed are locations for: Europe. Use the <Go Back> option to select a different continent or region if your location is not listed.

Country, territory or area:

Denmark
Estonia
Faroe Islands
Finland
France
Georgia
Germany
Gibraltar
Greece
Greenland
Guernsey
Holy See (Vatican City State)
Hungary

Configure locales

There is no locale defined for the combination of language and country you have selected. You can now select your preference from the locales available for the selected language. The locale that will be used is listed in the second column.

Country to base default locale settings on:

Canada	-	en_CA.UTF-8
Hong Kong	-	en_HK.UTF-8
India	-	en_IN
Ireland	-	en_IE.UTF-8
Israel	-	en_IL
New Zealand	-	en_NZ.UTF-8
Nigeria	-	en_NG
Philippines	-	en_PH.UTF-8
Seychelles	-	en_SC.UTF-8
Singapore	-	en_SG.UTF-8
South Africa	-	en_ZA.UTF-8
United Kingdom	-	en_GB.UTF-8
United States	-	**en_US.UTF-8**
Zambia	-	en_ZM
Zimbabwe	-	en_ZW.UTF-8

Fig. 3: Choosing language, location and locales

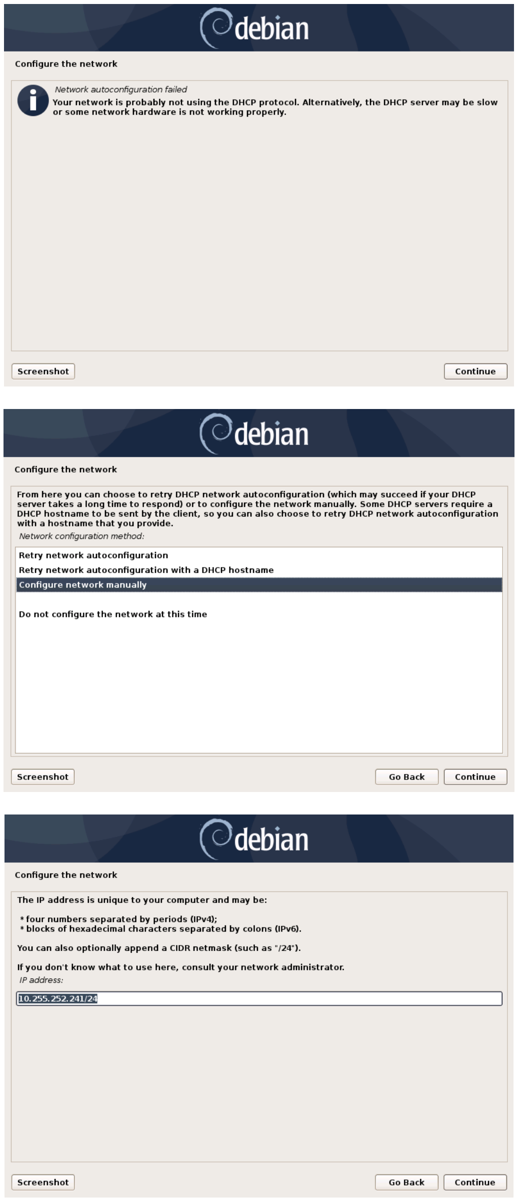

Configure the network

Network autoconfiguration failed

Your network is probably not using the DHCP protocol. Alternatively, the DHCP server may be slow or some network hardware is not working properly.

Configure the network

From here you can choose to retry DHCP network autoconfiguration (which may succeed if your DHCP server takes a long time to respond) or to configure the network manually. Some DHCP servers require a DHCP hostname to be sent by the client, so you can also choose to retry DHCP network autoconfiguration with a hostname that you provide.

Network configuration method:

- Retry network autoconfiguration
- Retry network autoconfiguration with a DHCP hostname
- Configure network manually
- Do not configure the network at this time

Configure the network

The IP address is unique to your computer and may be:

- four numbers separated by periods (IPv4);
- blocks of hexadecimal characters separated by colons (IPv6).

You can also optionally append a CIDR netmask (such as "/24").

If you don't know what to use here, consult your network administrator.

IP address:

10.255.252.241/24

Fig. 4: Network Configuration - IP addresses

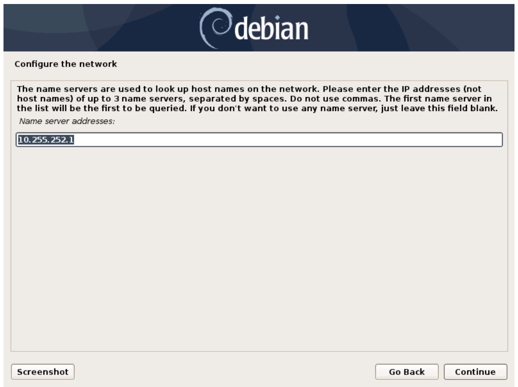

Configure the network

The name servers are used to look up host names on the network. Please enter the IP addresses (not host names) of up to 3 name servers, separated by spaces. Do not use commas. The first name server in the list will be the first to be queried. If you don't want to use any name server, just leave this field blank.

Name server addresses:

10.255.252.1

Fig. 5: Network Configuration – Enter the DNS server addresses

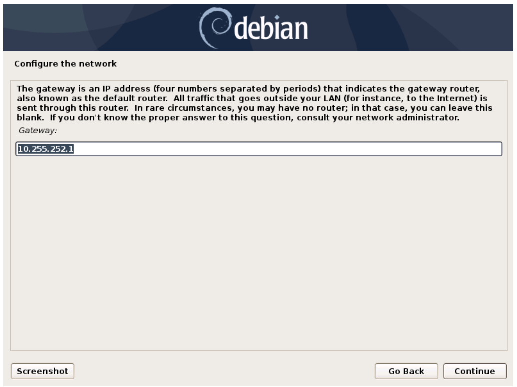

Configure the network

The gateway is an IP address (four numbers separated by periods) that indicates the gateway router, also known as the default router. All traffic that goes outside your LAN (for instance, to the Internet) is sent through this router. In rare circumstances, you may have no router; in that case, you can leave this blank. If you don't know the proper answer to this question, consult your network administrator.

Gateway: 10.255.252.1

Fig. 6: Network Configuration - Enter the Gateway

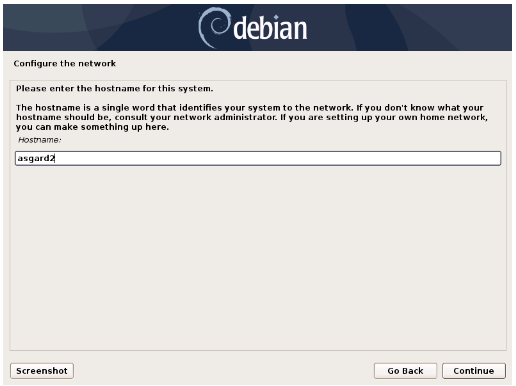

Configure the network

Please enter the hostname for this system.

The hostname is a single word that identifies your system to the network. If you don't know what your hostname should be, consult your network administrator. If you are setting up your own home network, you can make something up here.

Hostname: asgard2

Fig. 7: Network Configuration - Enter the Hostname

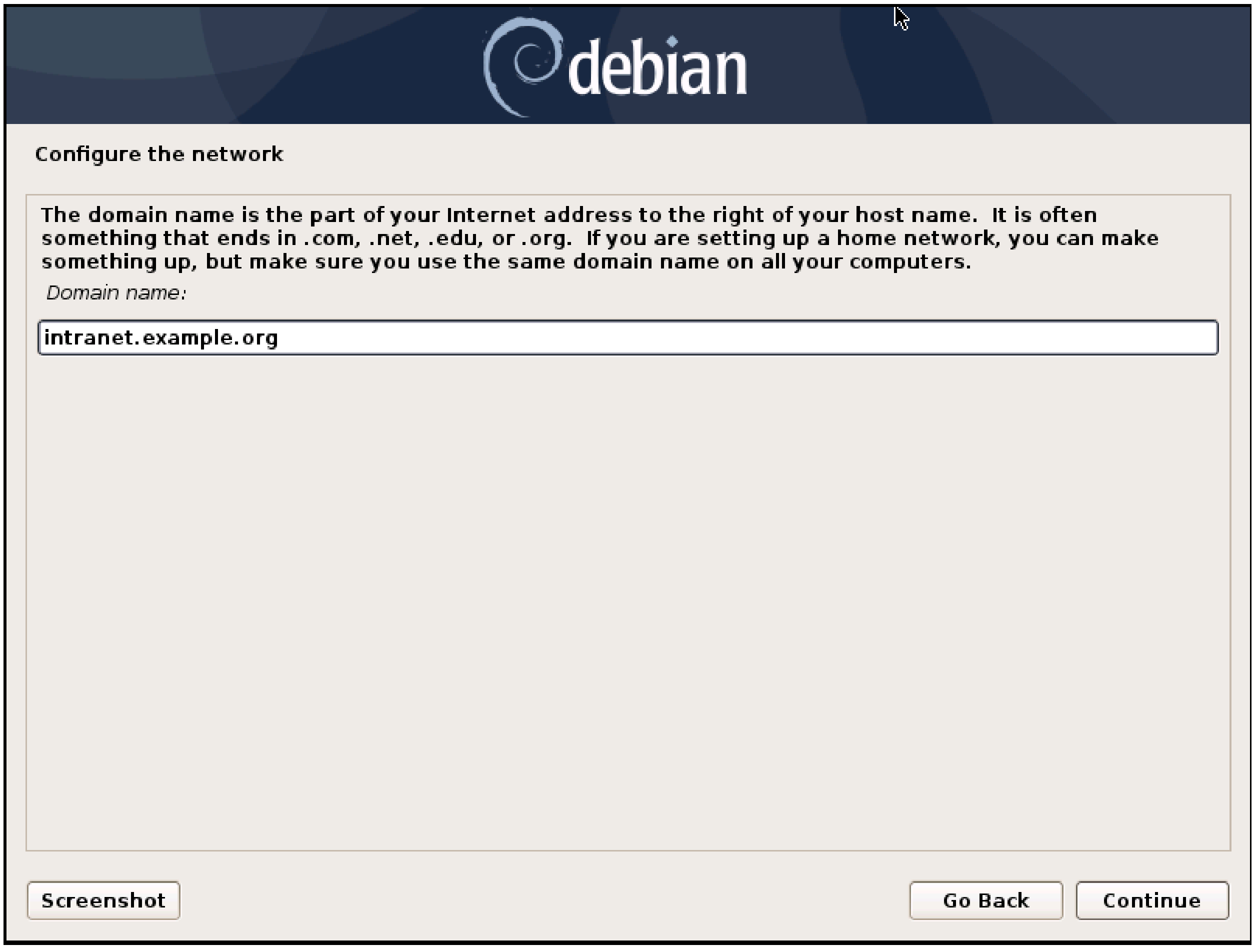

Configure the network

The domain name is the part of your Internet address to the right of your host name. It is often something that ends in .com, .net, .edu, or .org. If you are setting up a home network, you can make something up, but make sure you use the same domain name on all your computers.

Domain name:

intranet.example.org

Fig. 8: Network Configuration - Enter the Domain name

Danger: **Important:** Make sure that the combination of hostname and domain creates an FQDN that can be resolved from the ASGARD Management Center(s) you want to connect with your Analysis Cockpit. If you've configured a FQDN (hostname + domain) that cannot be resolved, your ASGARDs will encounter an error during connection.

This is especially important since your Analysis Cockpit will create some certificates during the installation, which will not contain an IP Address as its Subject Alternative Name (SAN), but only the FQDN! You will not be able to connect your ASGARD Management Center with your Analysis Cockpit via IP Address.

2.4 Choosing a password

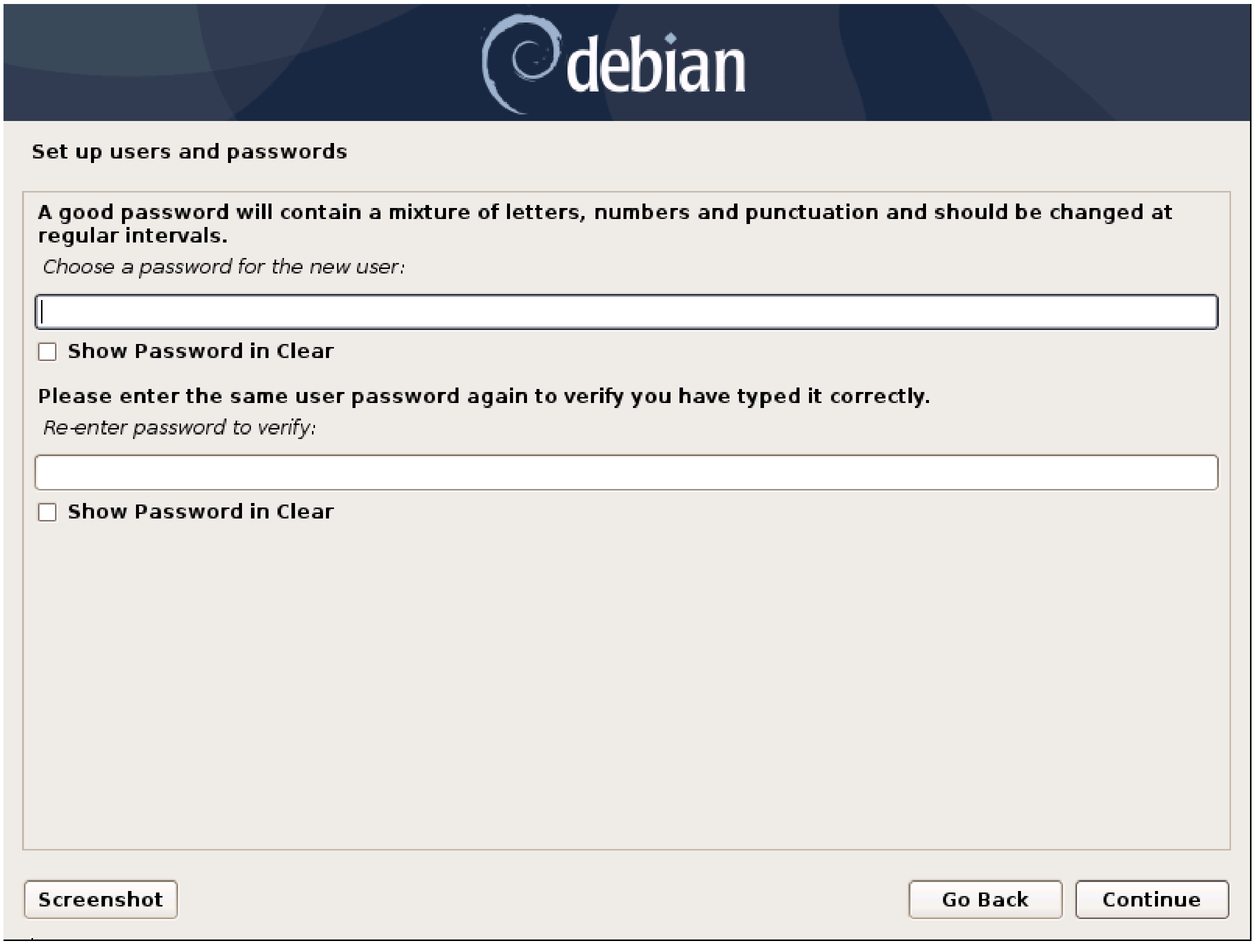

Set up users and passwords

A good password will contain a mixture of letters, numbers and punctuation and should be changed at regular intervals.

Choose a password for the new user:

□ Show Password in Clear

Please enter the same user password again to verify you have typed it correctly.

Re-enter password to verify:

□ Show Password in Clear

Fig. 9: Choosing a password for the nextron user

2.5 Partitioning of the Hard Disk

Warning: The Analysis Cockpit is intended to be installed with only one disk. Do not configure your server with multiple disks. The system won't configure additional disks. Make sure that your disk has the recommended size. See *Hardware Requirements* for more information.

Finally, confirm the settings, select “Yes” and click “Continue”.

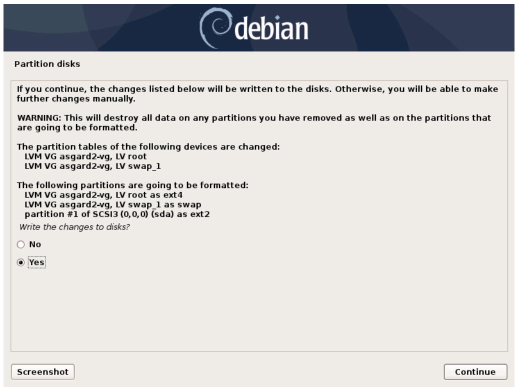

Partition disks

If you continue, the changes listed below will be written to the disks. Otherwise, you will be able to make further changes manually.

WARNING: This will destroy all data on any partitions you have removed as well as on the partitions that are going to be formatted.

The partition tables of the following devices are changed:
  LVM VG asgard2-vg, LV root
  LVM VG asgard2-vg, LV swap_1

The following partitions are going to be formatted:
  LVM VG asgard2-vg, LV root as ext4
  LVM VG asgard2-vg, LV swap_1 as swap
  partition #1 of SCSI3 (0,0,0) (sda) as ext2

Write the changes to disks?

○ No

● Yes

Fig. 10: Partition Disks – Write changes to disks

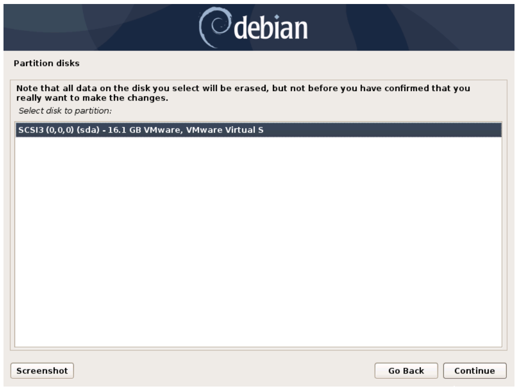

Partition disks

Note that all data on the disk you select will be erased, but not before you have confirmed that you really want to make the changes.

Select disk to partition:

SCSI3 (0,0,0) (sda) - 16.1 GB VMware, VMware Virtual S

Fig. 11: Partition disks – Select disk to partition

2.6 Proxy Configuration

If you are using a proxy to access the Internet, enter the proxy details in the next step. Please note, Internet connectivity is required for the next step – the installation of the ASGARD Analysis Cockpit service.

The proxy configuration supports unauthorized access and HTTP Auth, for example `http://our-proxy.local:8080` and `http://username:password@our-proxy.local:8000`

Hint: Your Installer will finish now the installation. After the installation is done, you will be prompted to log in to your server.

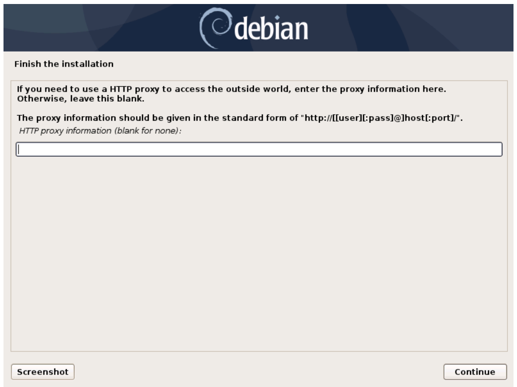

Finish the installation

If you need to use a HTTP proxy to access the outside world, enter the proxy information here. Otherwise, leave this blank.

The proxy information should be given in the standard form of "http://[[user][:pass]@]host[:port]/".

HTTP proxy information (blank for none):

Fig. 12: Proxy Configuration

2.7 Changing the IP-Address

The Analysis Cockpit's IP-Address can be changed in **/etc/network/interfaces**. The IP is configured with the `address` variable.

```
nextron@asgard-ac:~$ sudo vi /etc/network/interfaces
```

```
auto ens32
iface ens32 inet static
   address 172.16.2.7/24
   gateway 172.16.2.254
   dns-nameservers 172.16.20.20
```

You can now restart `networking.service` to apply the changes.

```
nextron@asgard-ac:~$ sudo systemctl restart networking.service
```

Important:

- The network interface might have a different name, so pay attention to the name (in this example `ens32`).
- If restarting the `networking.service` is throwing an error, you you can restart the server

The new IP can be applied with the command **sudo systemctl restart networking**

2.7.1 Verifying DNS Settings

To verify if ASGARD is using the correct DNS Server, you can inspect the file `/etc/resolv.conf`:

```
nextron@asgard-ac:~$ cat /etc/resolv.conf
search example.org
nameserver 172.16.200.2
```

If you see errors in this configuration, you can change it with the following command:

```
nextron@asgard-ac:~$ sudoedit /etc/resolv.conf
```

2.8 Install the Analysis Cockpit Services

The base installation is now complete. In the next step we'll install the Analysis Cockpit service.

Important:

- Internet connectivity is required for this step.
- Use an upper case `i` in the word `nextronInstaller`.

Use the VMWare console or SSH to the appliance using the user `nextron`.

To start the Analysis Cockpit installation run the following command:

```
nextron@asgard-ac:~$ sudo nextronInstaller -cockpit
```

After the installer has completed its operations successfully, the system is ready to be used.

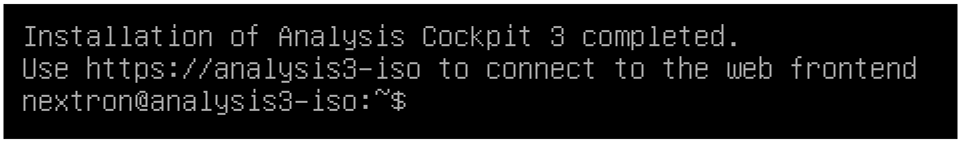

```
Installation of Analysis Cockpit 3 completed.
Use https://analysis3-iso to connect to the web frontend
nextron@analysis3-iso:~$
```

Fig. 13: Message upon successful completion

Note that the FQDN shown after `https://` has to be resolvable by the connected ASGARD Management Centers and users that try to access the Analysis Cockpit.

CHAPTER

THREE

ADMINISTRATION

This chapter assumes, that you have read the chapter *Basic Concepts*.

In order to configure the Analysis Cockpit for the first use, the following steps need to be done:

- License installation
- System update
- Set users and set user rights
- Define canned responses
- Decide about syslog forwarding
- Integrate your log source

These steps are described in detail in the following sections.

3.1 Initial Tasks

3.1.1 License Installation

Before you can use the cockpit, you must install a license. Navigate to `Licensing` section in the `Settings` menu, click the `Upload License` button, select your license file and click `Upload`. After verifying if your license is valid, you will be able to use your Analysis Cockpit.

3.1.2 System Update

All updates are applied from the Web GUI. Simply navigate to the `Updates` section in the `Settings` menu, review the release notes and click the update button. You can also check for new updates by clicking the `Check for Updates`.

Elasticsearch Cluster Update

If you are running an Elasticsearch Cluster with your Analysis Cockpit, we recommend to update the cluster members anytime you are installing an update on your Analysis Cockpit. Not only might an update for the Analysis Cockpit contain an update for Elasticsearch, but more importantly, system and security updates for the underlying debian system are also included.

To update your cluster members, run the following commands on each of them:

```
nextron@node-1:~$ sudo apt update
nextron@node-1:~$ sudo apt upgrade
```

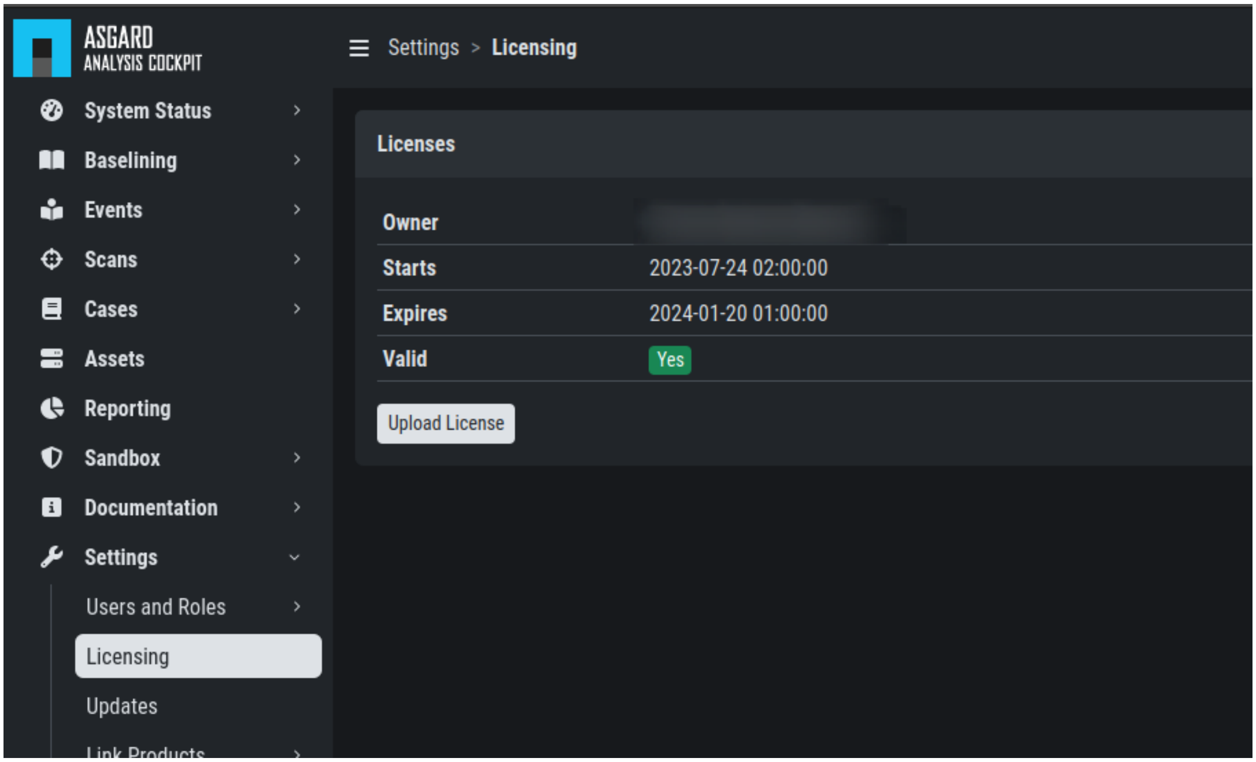

Licenses

Owner	[redacted]
Starts	2023-07-24 02:00:00
Expires	2024-01-20 01:00:00
Valid	Yes

Fig. 1: Licensing

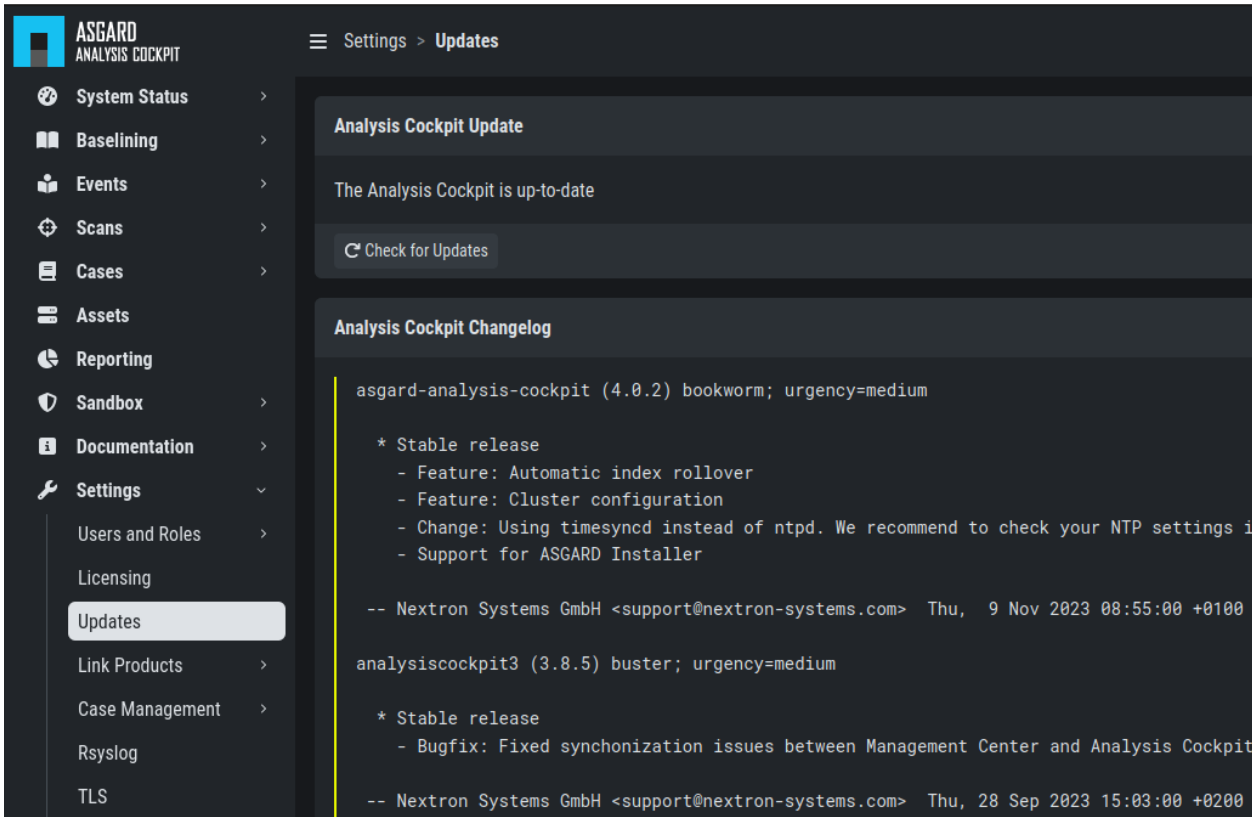

Analysis Cockpit Update

The Analysis Cockpit is up-to-date

Analysis Cockpit Changelog

```
asgard-analysis-cockpit (4.0.2) bookworm; urgency=medium

  * Stable release
    - Feature: Automatic index rollover
    - Feature: Cluster configuration
    - Change: Using timesyncd instead of ntpd. We recommend to check your NTP settings i
    - Support for ASGARD Installer

 -- Nextron Systems GmbH <support@nextron-systems.com>  Thu,  9 Nov 2023 08:55:00 +0100

analysiscockpit3 (3.8.5) buster; urgency=medium

  * Stable release
    - Bugfix: Fixed synchonization issues between Management Center and Analysis Cockpit

 -- Nextron Systems GmbH <support@nextron-systems.com>  Thu, 28 Sep 2023 15:03:00 +0200
```

Fig. 2: Updating the System

Note: Performing system updates is usually risk free. However, we still recommend that you create a backup/snapshot before updating your cluster nodes.

3.1.3 Set Users and User Rights

The chapter *Understanding Users, Roles, Rights and Case Status* already describes how to set up a 2-level analyst model for working with cases. The roles defined in that section are non-administrative roles, meaning they are only allowed to access cases based on the respective status of a ticket. The following permissions are related to the Analysis Cockpit as a whole.

Additionally, roles can have the following rights:

- Administrator
- Universal
- View Notifications
- Acknowledge Notifications
- Upload Events
- Delete Events
- Upload File(s) for Sandbox Analysis
- Download File(s) for Sandbox Analysis

Roles can be granted these privileges by choosing them in the `New Role` dialogue.

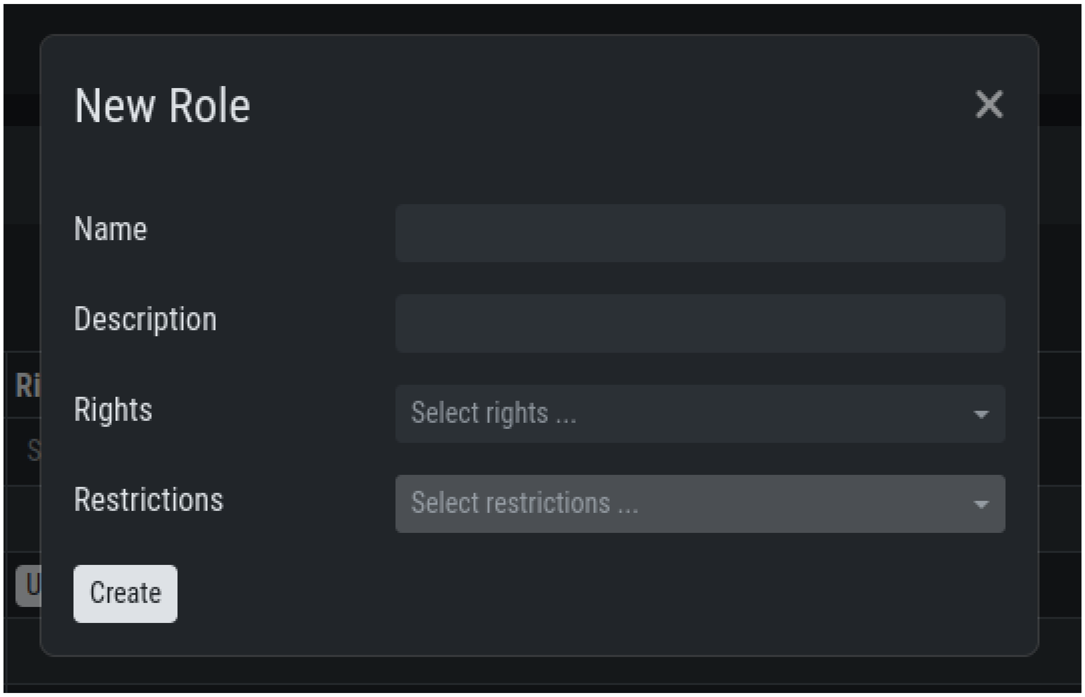

Fig. 3: New Role

3.2 Configure Canned Recommendations

Canned recommendations are predefined actions that can be used within a case. The recommendations are fully configurable and are aimed to facilitate choice making regarding the action that should be applied for a specific case. There is no need to set this up, but we suggest doing some planning and provide recommendations that are suitable for your organization. Some recommendations such as `Verify Legitimacy`, `Provide Sample File / Sample Directory`, `Run full Antivirus Scan` are already generated by default. You are free to use, modify or delete them. Recommendations can also be added by any user from within a case.

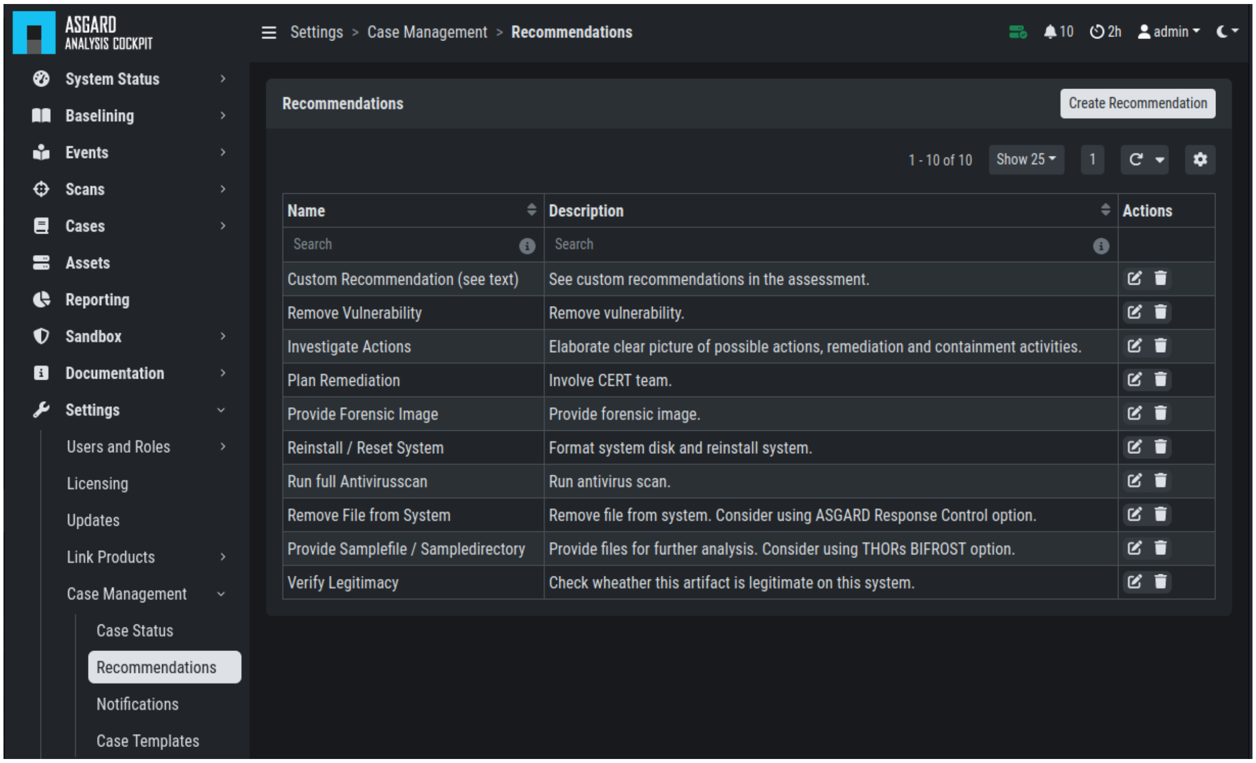

Name	Description	Actions
Search	Search	
Custom Recommendation (see text)	See custom recommendations in the assessment.	
Remove Vulnerability	Remove vulnerability.	
Investigate Actions	Elaborate clear picture of possible actions, remediation and containment activities.	
Plan Remediation	Involve CERT team.	
Provide Forensic Image	Provide forensic image.	
Reinstall / Reset System	Format system disk and reinstall system.	
Run full Antivirusscan	Run antivirus scan.	
Remove File from System	Remove file from system. Consider using ASGARD Response Control option.	
Provide Samplefile / Sampledirectory	Provide files for further analysis. Consider using THORs BIFROST option.	
Verify Legitimacy	Check wheather this artifact is legitimate on this system.	

Fig. 4: Case Management- Recommendations

3.3 Syslog Forwarding

The `Rsyslog` tab in the `Settings` menu allows forwarding of all incoming THOR events, along with all audit logs and all other Cockpit related logs.

Please note, that forwarding THOR Logs through syslog might lead to a certain loss of information as THOR events might exceed syslog length restrictions.

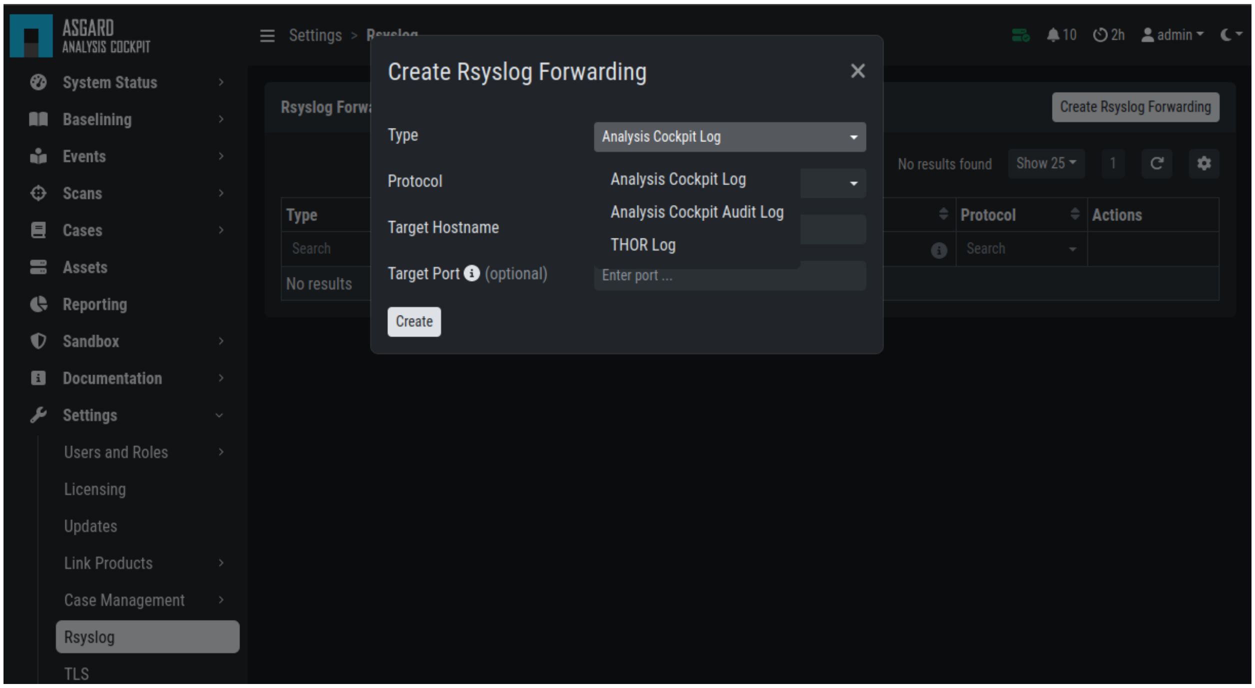

Create Rsyslog Forwarding

Type: Analysis Cockpit Log

Protocol

Target Hostname

Target Port (optional): Enter port ...

Fig. 5: Add Rsyslog Forwarding II

3.4 TLS Certificate Installation

Instead of using the pre-installed self-signed TLS Certificate, users should upload their own TLS Certificate for ASGARD. This will avoid browser warnings when navigating to your Analysis Cockpit's web interface.

In order to achieve the best possible compatibility with the most common browsers, we recommend using the system's FQDN in both fields `Common Name` AND `Hostnames`.

Navigate to the `TLS` section via the `Settings` menu. You can click `Generate CSR` to open the following modal.

Hint: Please note that generating a CSR on the command line is not supported.

The generated CSR can be used to generate a TLS Certificate. Subsequently, this TLS Certificate can be uploaded in the in the same section of your Analysis Cockpit.

3.5 Configure LDAP

The `LDAP` tab in the `Users and Roles` section lets you configure an LDAP server and define mappings between LDAP groups and roles within the Analysis Cockpit.

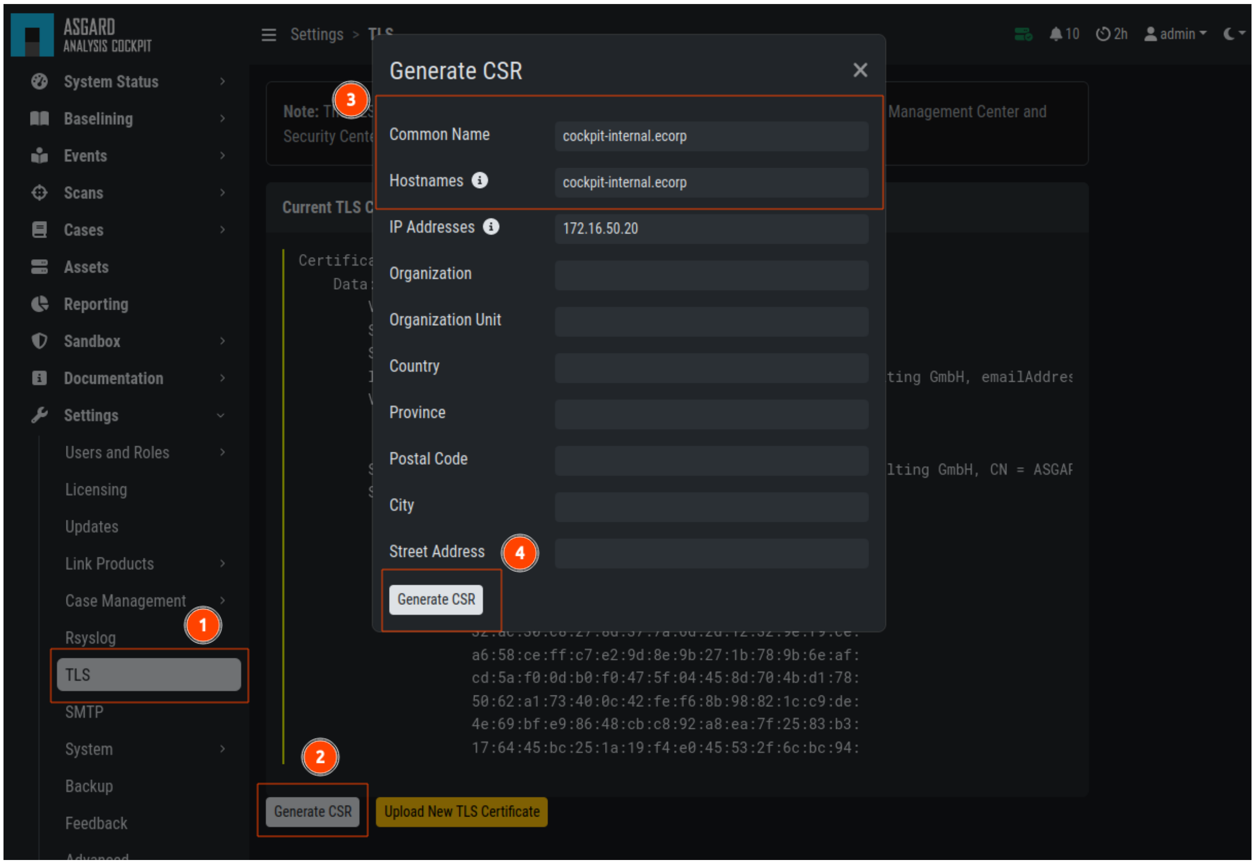

Generate CSR

Common Name	cockpit-internal.ecorp
Hostnames	cockpit-internal.ecorp
IP Addresses	172.16.50.20
Organization	
Organization Unit	
Country	
Province	
Postal Code	
City	
Street Address	

Generate CSR

```
a6:58:ce:ff:c7:e2:9d:8e:9b:27:1b:78:9b:6e:af:
cd:5a:f0:0d:b0:f0:47:5f:04:45:8d:70:4b:d1:78:
50:62:a1:73:40:0c:42:fe:f6:8b:98:82:1c:c9:de:
4e:69:bf:e9:86:48:cb:c8:92:a8:ea:7f:25:83:b3:
17:64:45:bc:25:1a:19:f4:e0:45:53:2f:6c:bc:94:
```

Fig. 6: Generate a Certificate Signing Request (CSR)

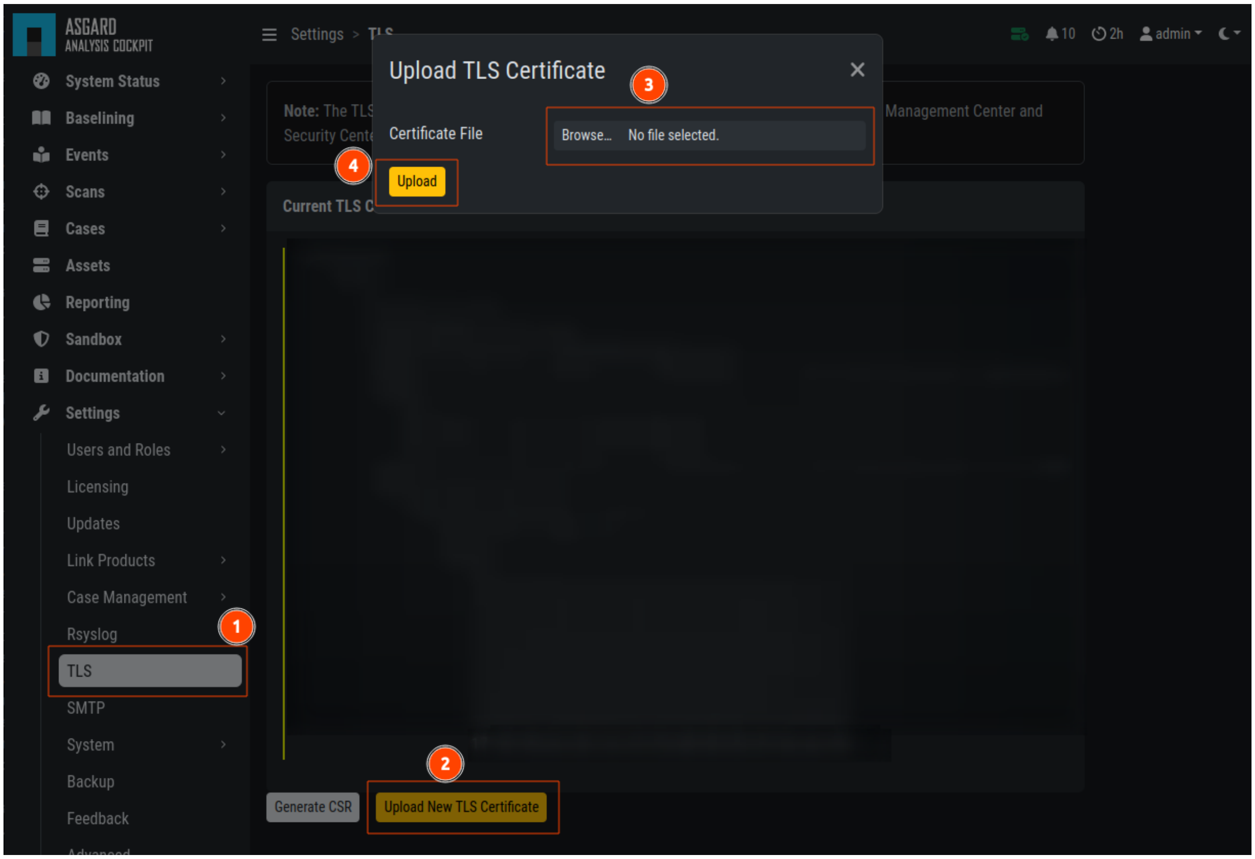

Note: The TLS ... Security Cente... Management Center and

Fig. 7: Upload a TLS Certificate

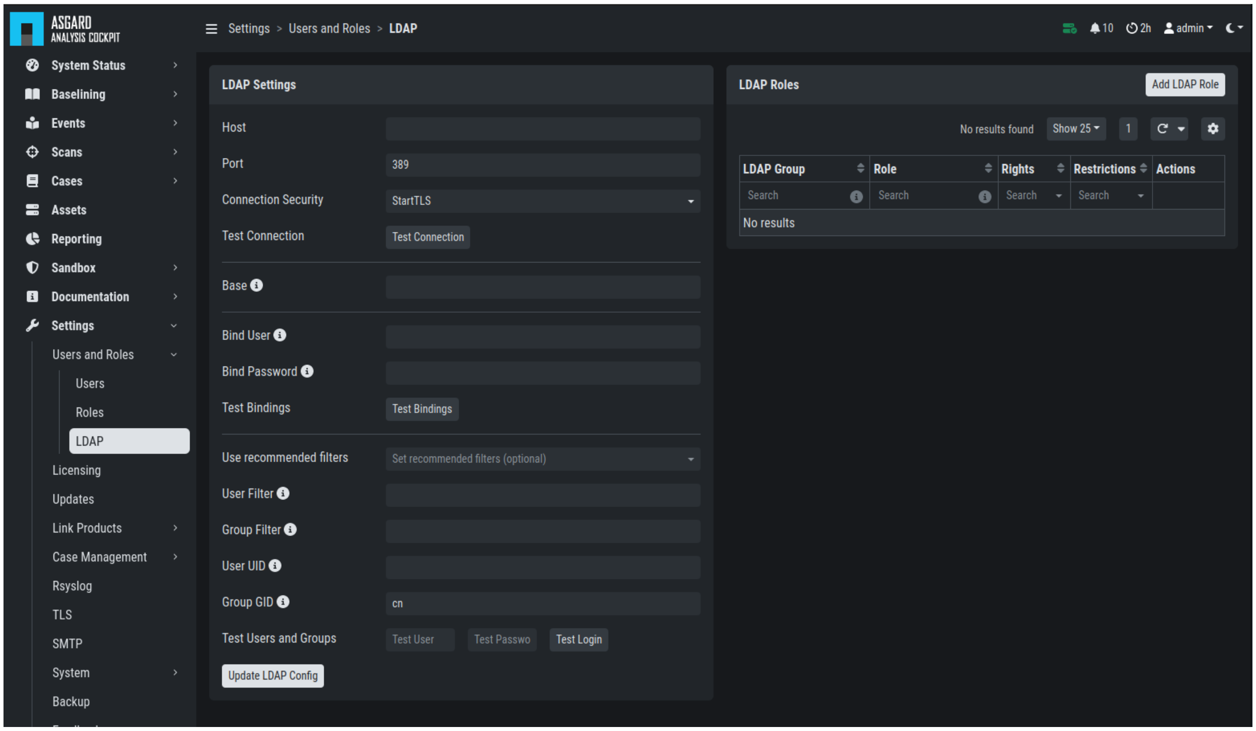

LDAP Settings

Host:
Port: 389
Connection Security: StartTLS
Base:
Bind User:
Bind Password:
Use recommended filters: Set recommended filters (optional)
User Filter:
Group Filter:
User UID:
Group GID: cn
Test Users and Groups

LDAP Group	Role	Rights	Restrictions	Actions
Search	Search	Search	Search	
No results				

Fig. 8: Configure LDAP

3.6 Configure Notifications

As described in *Cases and Log Processing*, the Analysis Cockpit is able to forward logs to a SIEM system in case this particular log line was added automatically to a case with the type "Incident".

The `Notifications` section in the `Case Management` settings allow you to define custom notifications for event assignments (Event Assignment Notifications). It is recommended to at least configure an Event Assignment Notification for events that get added to existing **Incident** cases.

Additionally, notifications can be defined for changes to cases (Case Change Notifications), so Level 2 analysts can get notified if a case gets added to their in-queue (e.g., Finished Level 1).

The notification itself can be a syslog message or an email. In order to use email for notifications you have to setup an email account in the `Mail Account` Tab. Additionally webhook support has been added to facilitate interfacing to services like Slack.

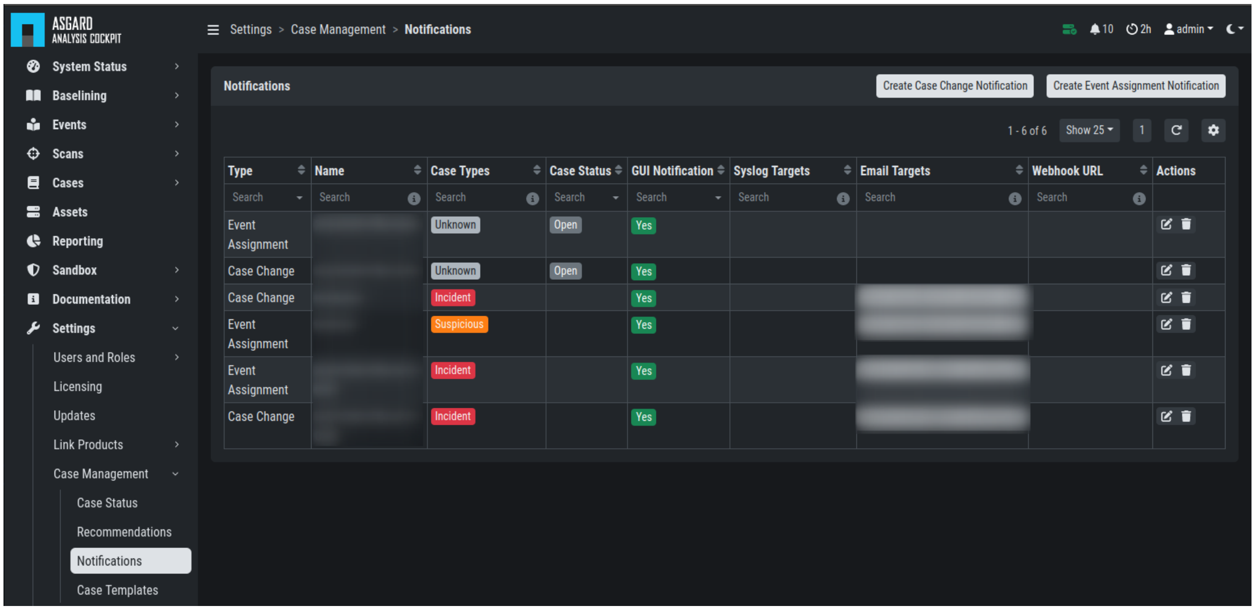

Notifications

Type	Name	Case Types	Case Status	GUI Notification	Syslog Targets	Email Targets	Webhook URL	Actions
Search	Search	Search	Search	Search	Search	Search	Search	
Event Assignment	[redacted]	Unknown	Open	Yes				
Case Change	[redacted]	Unknown	Open	Yes				
Case Change	[redacted]	Incident		Yes		[redacted]		
Event Assignment	[redacted]	Suspicious		Yes		[redacted]		
Event Assignment	[redacted]	Incident		Yes		[redacted]		
Case Change	[redacted]	Incident		Yes		[redacted]		

Fig. 9: Case Management - Notifications

Note: The Analysis Cockpit will collect all triggering events and send only one email every 15 minutes. Syslog and Webhooks are triggered in real time for every single event.

Additionally, you can see the notifications in the top right corner (bell icon) and inspect them. You will see all `Unread` notifications, which can be `Acknowledged` by selecting one or more notification and clicking `Acknowledge`. Only `Unread` notifications will show up in the top right status bar of the Cockpit.

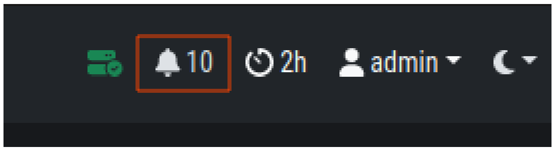

Fig. 10: UI Notification Bell

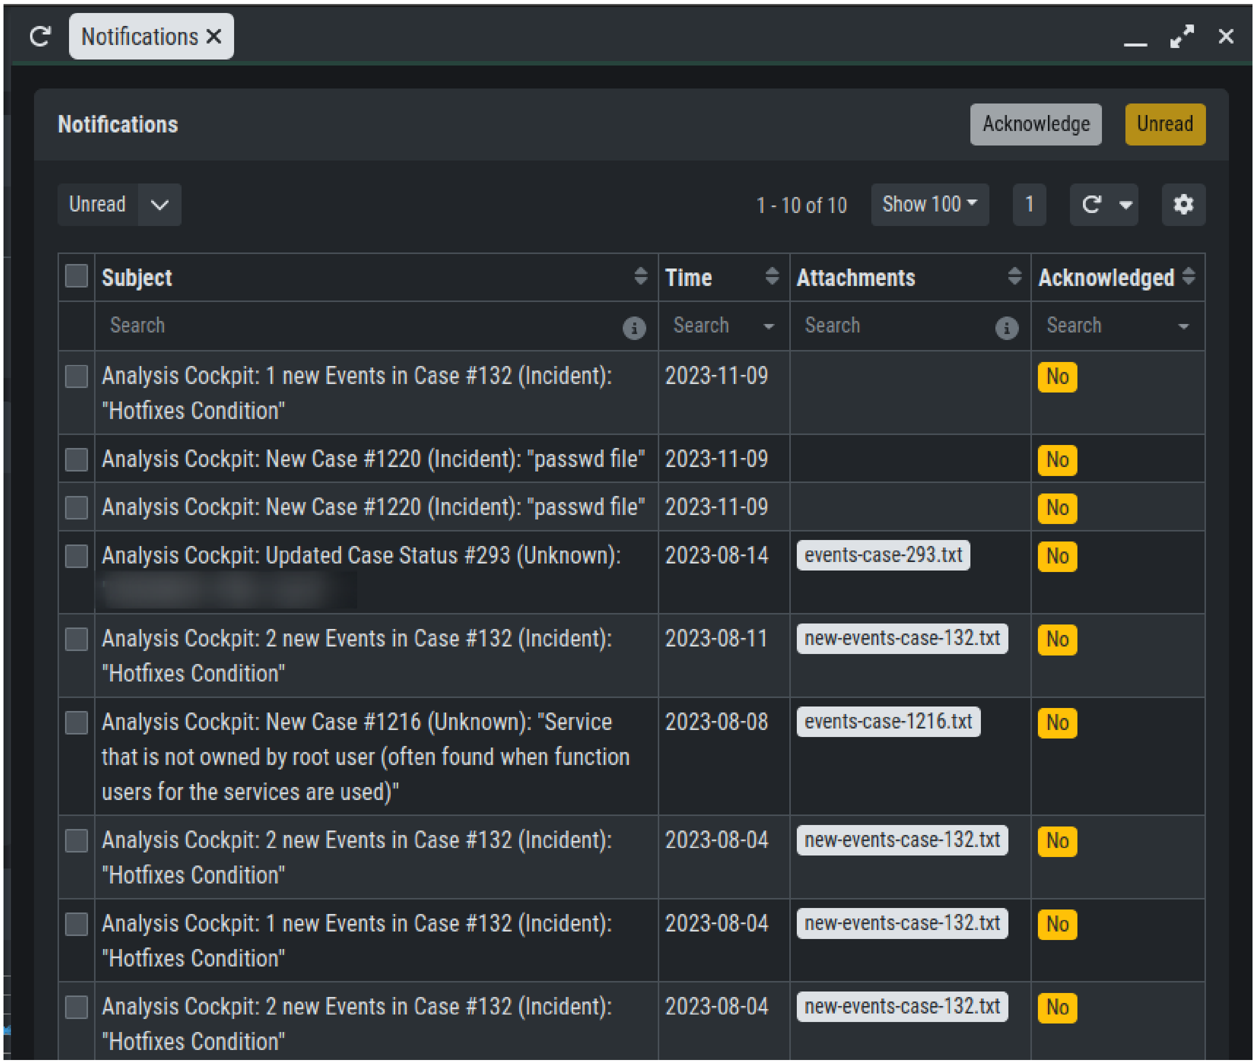

□	Subject	Time	Attachments	Acknowledged
	Search	Search	Search	Search
□	Analysis Cockpit: 1 new Events in Case #132 (Incident): "Hotfixes Condition"	2023-11-09		No
□	Analysis Cockpit: New Case #1220 (Incident): "passwd file"	2023-11-09		No
□	Analysis Cockpit: New Case #1220 (Incident): "passwd file"	2023-11-09		No
□	Analysis Cockpit: Updated Case Status #293 (Unknown): [redacted]	2023-08-14	events-case-293.txt	No
□	Analysis Cockpit: 2 new Events in Case #132 (Incident): "Hotfixes Condition"	2023-08-11	new-events-case-132.txt	No
□	Analysis Cockpit: New Case #1216 (Unknown): "Service that is not owned by root user (often found when function users for the services are used)"	2023-08-08	events-case-1216.txt	No
□	Analysis Cockpit: 2 new Events in Case #132 (Incident): "Hotfixes Condition"	2023-08-04	new-events-case-132.txt	No
□	Analysis Cockpit: 1 new Events in Case #132 (Incident): "Hotfixes Condition"	2023-08-04	new-events-case-132.txt	No
□	Analysis Cockpit: 2 new Events in Case #132 (Incident): "Hotfixes Condition"	2023-08-04	new-events-case-132.txt	No

Fig. 11: UI Notifications

3.6.1 Configure Event Assignment Notifications

To configure log notifications, click the `Add Event Assignment Notification` button in the `Notifications` section of the `Case Management` menu. This leads you to a form that allows you to set a name for your notification, the notification type (syslog, email, webhook or notification within the Analysis Cockpit) and the condition that will trigger your notification.

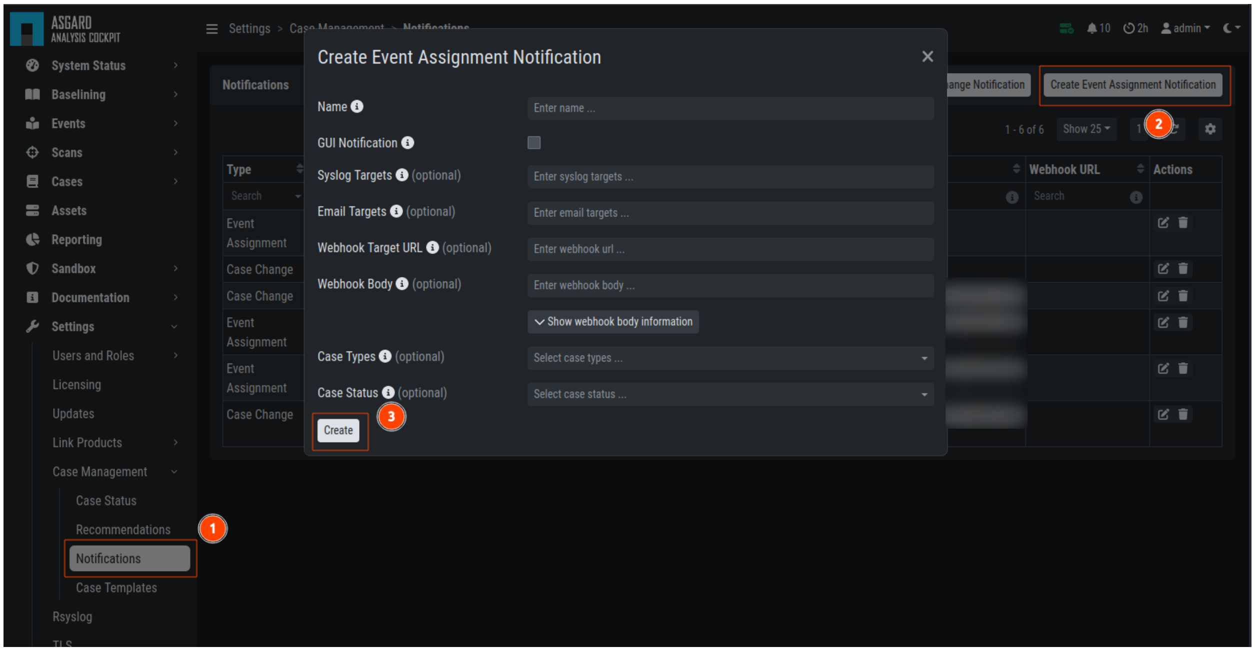

Fig. 12: Event Assignment Notification

3.6.2 Configure Case Change Notifications

To configure Case Change Notifications, click the `Add Case Change Notification` button in the `Notifications` section of the `Case Management` menu. This leads you to a form that allows setting a name for your notification, the notification type (syslog, email, webhook or notification within the Analysis Cockpit) and the condition that will trigger your notification.

3.7 Log File Import

3.7.1 Basic Concepts

In general, all logs show up in the Events section. Additionally, all Alerts and Warnings that are not matching a particular case will show up in the `Baselining` section. Notices and informational events will NOT show up in the Baselining Section as they match the predefined default cases for these events.

All logs are tagged with a specific scan id – regardless of how the log was integrated. This enables filtering down to all logs contained in a specific scan.

If ASGARD Management Center is connected and the events was generated as part of a group scan the event is also tagged with this particular group scan id. This allows for filtering down to all logs a particular group scan.

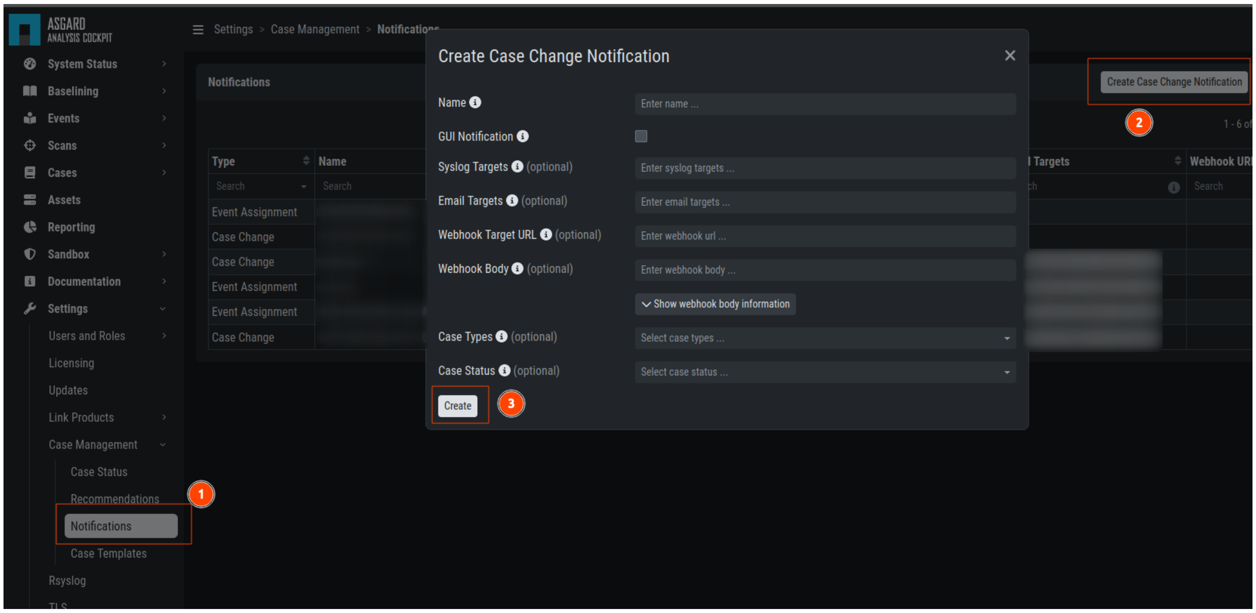

Type	Name
Search	Search
Event Assignment	
Case Change	
Case Change	
Event Assignment	
Event Assignment	
Case Change	

Create Case Change Notification

Name: Enter name ...

GUI Notification: □

Syslog Targets (optional): Enter syslog targets ...

Email Targets (optional): Enter email targets ...

Webhook Target URL (optional): Enter webhook url ...

Webhook Body (optional): Enter webhook body ...

Show webhook body information

Case Types (optional): Select case types ...

Case Status (optional): Select case status ...

Create

Fig. 13: Case Change Notification

Assets are identified through the asset ID that was issued by ASGARD Management Center during the setup of the ASGARD Agent. If this ID is not available to the Analysis Cockpit (e.g. log has been uploaded manually or sent through syslog) the hostname (NOT the FQDN) will be used instead.

3.7.2 Direct Integration with ASGARD Management Center

If the Analysis Cockpit is linked to one or more ASGARD Management Centers, all THOR logs get integrated automatically and will show up in the Baselining and/or the Events section. Aurora Events will also automatically show up.

To see how to connect an ASGARD Management Center with your Analysis Cockpit, follow the instructions in the chapter *Connect to ASGARD Management Center*.

You can retrieve old scans performed by ASGARD Management Center before you connected it to Analysis Cockpit using the `Request Events` button in the `Scans` section.

3.7.3 Syslog Input

Another way to import log data is by using SYSLOG messages.

The ANALYSIS COCKPIT listens on port 514/udp and 514/tcp for incoming log data and all logs will show up in the Baselining and/or the Events section.

Incoming syslog messages get assigned to single scan using the "ScanID" value that's unique in each scan.

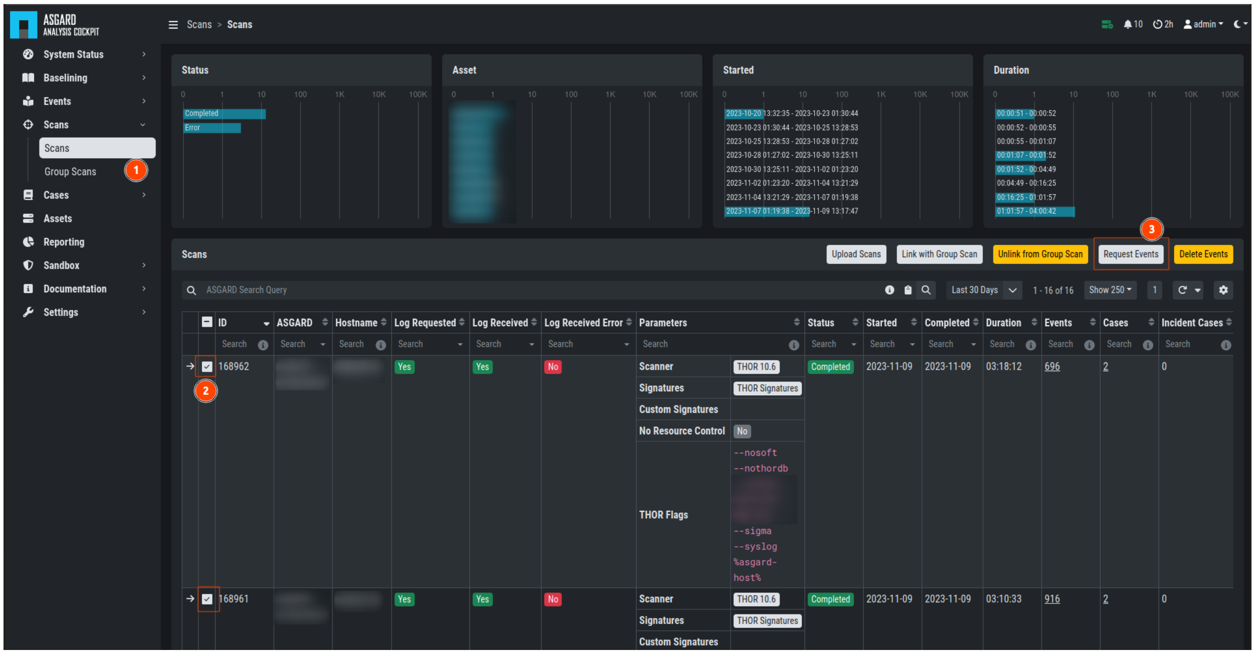

Started:
2023-10-20 13:32:35 - 2023-10-23 01:30:44
2023-10-23 01:30:44 - 2023-10-25 13:28:53
2023-10-25 13:28:53 - 2023-10-28 01:27:02
2023-10-28 01:27:02 - 2023-10-30 13:25:11
2023-10-30 13:25:11 - 2023-11-02 01:23:20
2023-11-02 01:23:20 - 2023-11-04 13:21:29
2023-11-04 13:21:29 - 2023-11-07 01:19:38
2023-11-07 01:19:38 - 2023-11-09 13:17:47

Duration:
00:00:51 - 00:00:52
00:00:52 - 00:00:55
00:00:55 - 00:01:07
00:01:07 - 00:01:52
00:01:52 - 00:04:49
00:04:49 - 00:16:25
00:16:25 - 01:01:57
01:01:57 - 04:00:42

	ID	ASGARD	Hostname	Log Requested	Log Received	Log Received Error	Parameters		Status	Started	Completed	Duration	Events	Cases	Incident Cases
☑	168962	[redacted]	[redacted]	Yes	Yes	No	Scanner	THOR 10.6	Completed	2023-11-09	2023-11-09	03:18:12	696	2	0
							Signatures	THOR Signatures							
							Custom Signatures								
							No Resource Control	No							
							THOR Flags	--nosoft --nothordb [redacted] --sigma --syslog %asgard-host%							
☑	168961	[redacted]	[redacted]	Yes	Yes	No	Scanner	THOR 10.6	Completed	2023-11-09	2023-11-09	03:10:33	916	2	0
							Signatures	THOR Signatures							
							Custom Signatures								

Fig. 14: Request Events from Scan

3.7.4 File Import Through Web-Based GUI

Alternatively, logs can be uploaded through the web-based interface by selecting the particular log file (must be the .txt format, html import is not supported) and clicking the `Upload Scans` button within the Scans section.

Note: You can upload one or more THOR scans in one or more text files. The Analysis Cockpit will automatically generate scans in the database, based on the scanned assets and the SCAN_IDs in the events. Only .txt, .log, .txt.gz and .log.gz files are supported.

After a successful upload, the scans should appear in the list below.

Important: If you can not see events in the `Events` or `Baselining` views, please make sure that you've selected the correct time frame as filter. Often times manually uploaded scans happened days or weeks before the upload. The log data gets indexed with the timestamp of their creation and not the import, and can therefore be hidden in the default view.

After the upload, you're able to link the recently uploaded scans with an existing or new group scan. You can also unlink scans from a group scan.

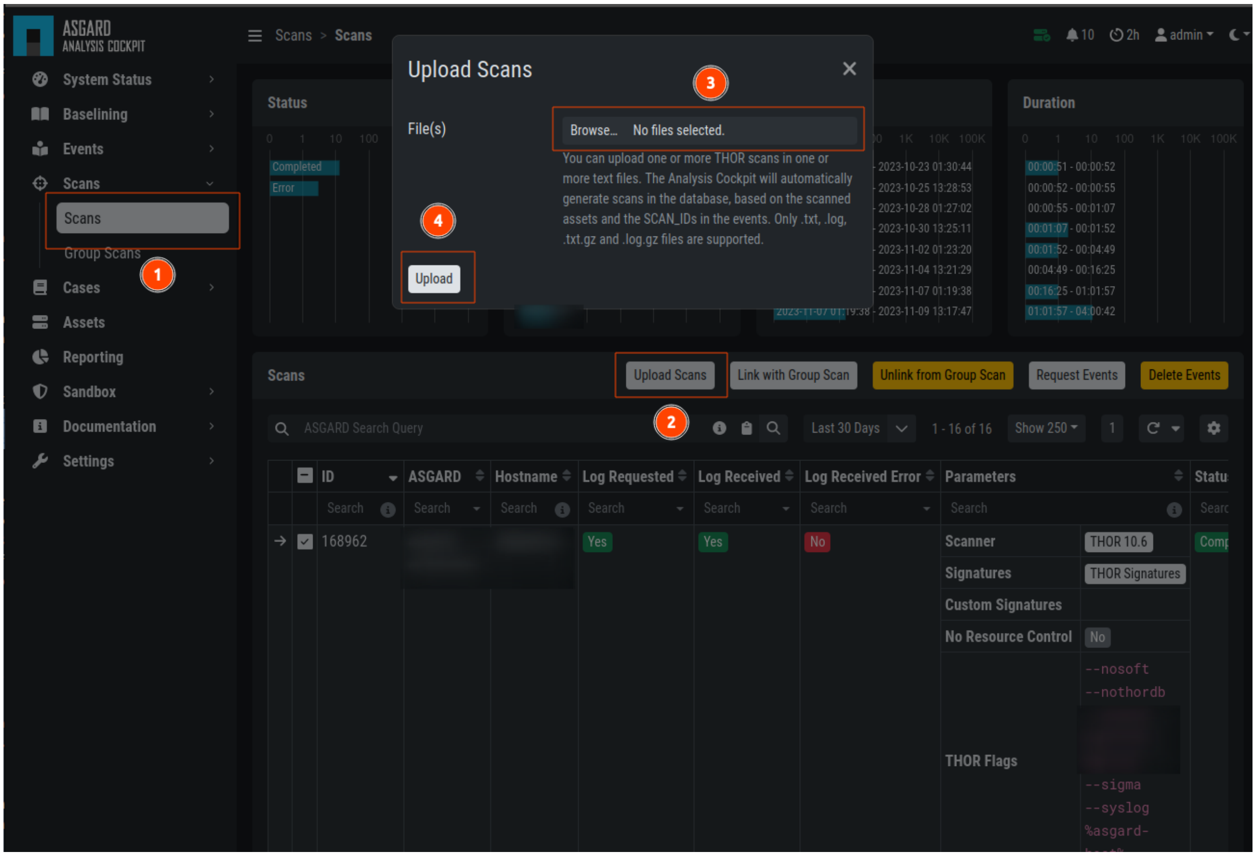

You can upload one or more THOR scans in one or more text files. The Analysis Cockpit will automatically generate scans in the database, based on the scanned assets and the SCAN_IDs in the events. Only .txt, .log, .txt.gz and .log.gz files are supported.

Fig. 15: Upload logs using the web-based interface

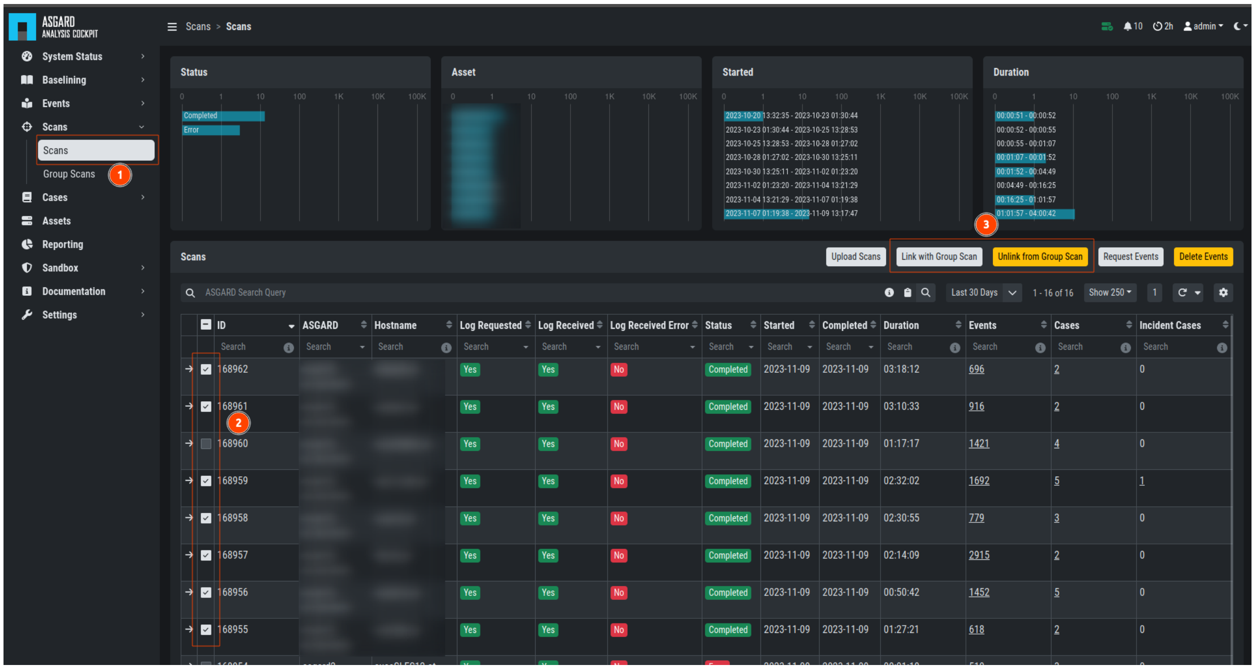

	ID	ASGARD	Hostname	Log Requested	Log Received	Log Received Error	Status	Started	Completed	Duration	Events	Cases	Incident Cases
☑	168962	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	03:18:12	696	2	0
☑	168961	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	03:10:33	916	2	0
□	168960	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	01:17:17	1421	4	0
☑	168959	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	02:32:02	1692	5	1
☑	168958	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	02:30:55	779	3	0
☑	168957	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	02:14:09	2915	2	0
☑	168956	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	00:50:42	1452	5	0
☑	168955	[redacted]	[redacted]	Yes	Yes	No	Completed	2023-11-09	2023-11-09	01:27:21	618	2	0

Fig. 16: Link/Unlink scans with an existing or new group scan

3.7.5 File Import Using the Command Line

This option can be helpful in an environment where you scan without ASGARD Management Center but want to automate analysis by dropping the log data into that import directory.

Log files can be imported by placing the files in the following directory:

`/var/lib/asgard-analysis-cockpit/events`

Make sure that user and group of these files is set to `cockpit`.

You can change the owner and group manually by using:

```
nextron@asgard-ac:~$ sudo chown cockpit:cockpit <file>
```

Successfully imported files get a new extension named `.ok`.

When the file is moved to that folder with the wrong permissions, Analysis Cockpit tries to handle these situations in the appropriate way. If the Analysis cockpit had read access but no rights to write/delete/rotate/rename the file, the file gets blacklisted in memory and will not be imported as long as the service doesn't get restarted. A restart of the service would cause the service to re-index the log data placed in that folder.

Important: We highly recommend not to directly copy (scp, rsync) files into that folder, but use a staging folder in which you set the right permissions and then copy the files to the import folder.

Copying files directly to that folder has many problematic side effects, e.g. files partly composed of binary zeros because the file transfer is still in progress.

3.8 Connect to ASGARD Management Center

In order to receive log data from ASGARD Management Center systems, add them in the corresponding section in the system settings.

`Settings` > `Link Products` > `Management Center`

After you have connected the two components, all assets along with additional information from ASGARD will show up in the `Assets` section of your Analysis Cockpit.

3.9 Asset View

In most cases working with the `Baselining` section and the `Cases` section can be seen as the best practice approach for setting baselines and dealing with alerts and warnings.

However, in some cases it makes sense to change perspective and rather go for a host centric approach. The Analysis Cockpit will calculate numbers of lines in different case types (Incident, Suspicious, Anomaly, etc.) on a per host basis for a given time frame. Along with information from ASGARD on last scan dates, labels, host availability etc. this gives an entirely different perspective.

By using the "Asset View" you can e.g., easily answer questions like:

- Which systems appear most often in “Incident” cases?
- Which systems haven't reported a single event for more than a month?
- Which Domain Controllers have not been scanned yet?

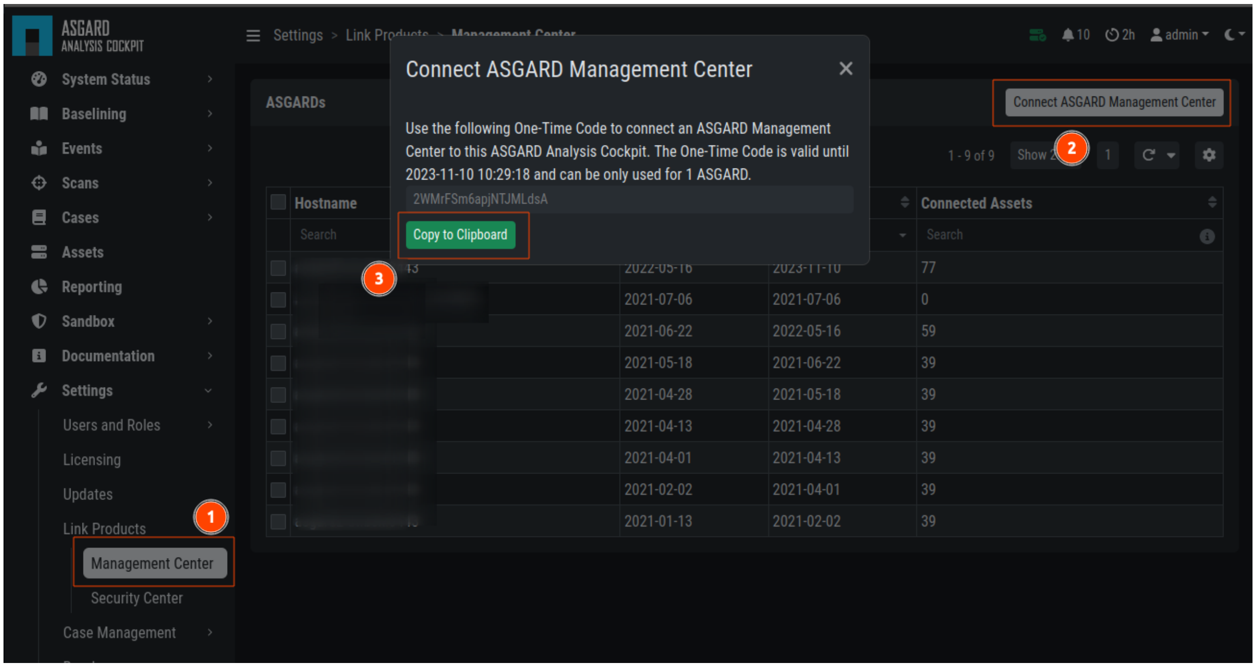

Connect ASGARD Management Center

Use the following One-Time Code to connect an ASGARD Management Center to this ASGARD Analysis Cockpit. The One-Time Code is valid until 2023-11-10 10:29:18 and can be only used for 1 ASGARD.

2WMrFSm6apjNTJMLdsA

Copy to Clipboard

Hostname			Connected Assets
Search			Search
[redacted]	2022-05-16	2023-11-10	77
[redacted]	2021-07-06	2021-07-06	0
[redacted]	2021-06-22	2022-05-16	59
[redacted]	2021-05-18	2021-06-22	39
[redacted]	2021-04-28	2021-05-18	39
[redacted]	2021-04-13	2021-04-28	39
[redacted]	2021-04-01	2021-04-13	39
[redacted]	2021-02-02	2021-04-01	39
[redacted]	2021-01-13	2021-02-02	39

Fig. 17: Link ASGARD Management Center

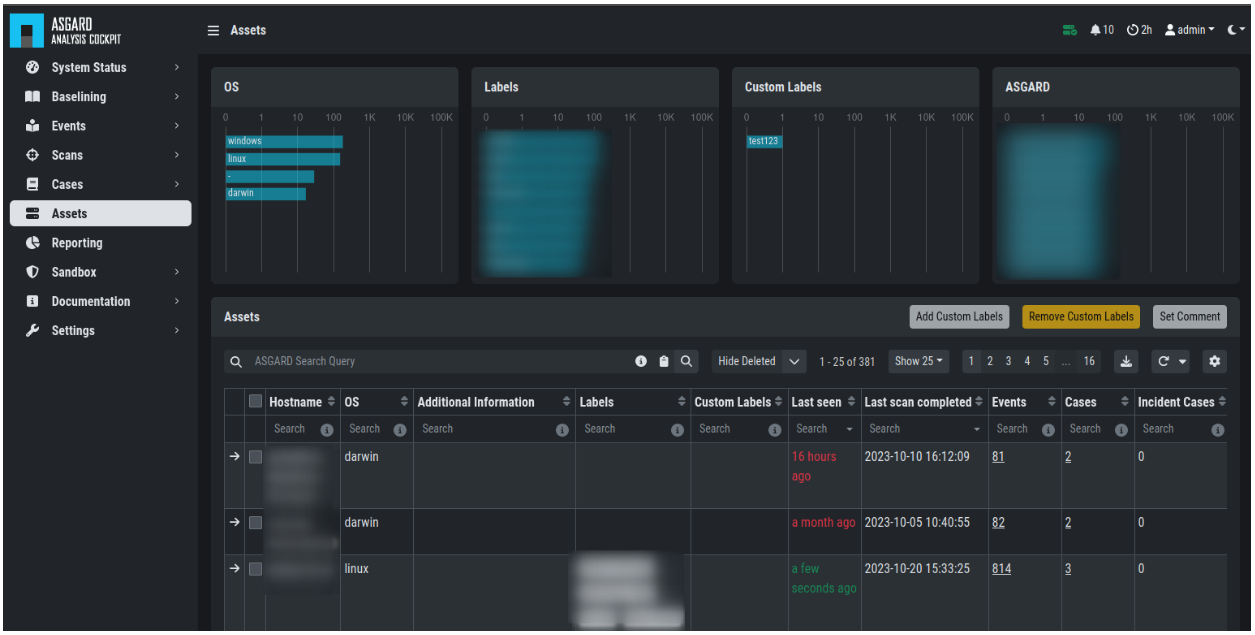

	Hostname	OS	Additional Information	Labels	Custom Labels	Last seen	Last scan completed	Events	Cases	Incident Cases
	Search	Search	Search	Search	Search	Search	Search	Search	Search	Search
→ □	[redacted]	darwin				16 hours ago	2023-10-10 16:12:09	81	2	0
→ □	[redacted]	darwin				a month ago	2023-10-05 10:40:55	82	2	0
→ □	[redacted]	linux		[redacted]		a few seconds ago	2023-10-20 15:33:25	814	3	0

Fig. 18: Asset View after a Successful Connection

- Which systems with IP addresses starting "192.168." appear in "Incident" cases?

In combination with the `ASGARD Query` and `Labels`, which are identical to your ASGARD, you can even narrow down the events by system group (e.g., Domain Controllers, or certain locations).

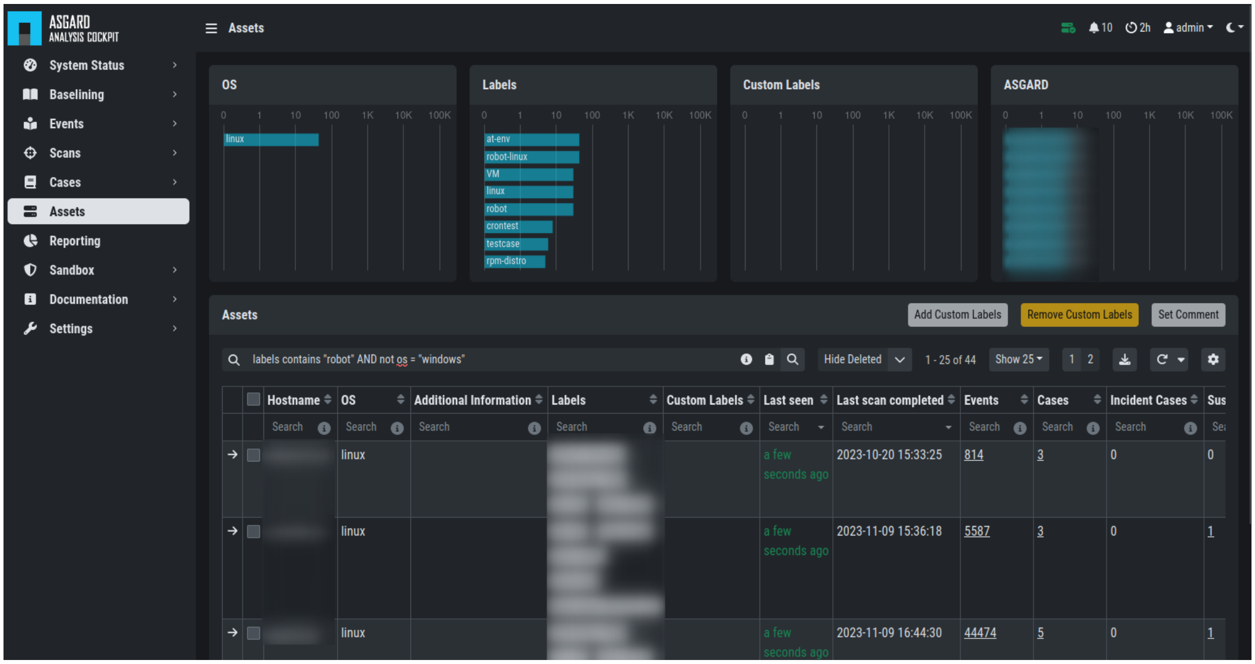

Assets

labels contains "robot" AND not os = "windows"

1 - 25 of 44

	□	Hostname	OS	Additional Information	Labels	Custom Labels	Last seen	Last scan completed	Events	Cases	Incident Cases	Sus
		Search	Search	Search	Search	Search	Search	Search	Search	Search	Search	Se
→	□	[redacted]	linux		[redacted]		a few seconds ago	2023-10-20 15:33:25	814	3	0	0
→	□	[redacted]	linux		[redacted]		a few seconds ago	2023-11-09 15:36:18	5587	3	0	1
→	□	[redacted]	linux		[redacted]		a few seconds ago	2023-11-09 16:44:30	44474	5	0	1

Fig. 19: Filtering within the Assets view

3.10 Sandbox Integration

You can configure your Analysis Cockpit to upload files to a local sandbox. Currently you can use CAPEv2 (recommended) or Cuckoo.

Additionally, you can look at the following `python` file and write your own connector, for a different sandbox, if you need to: `/usr/share/asgard-analysis-cockpit/sandbox/connector/capev2.py`

Note: This section only focuses on the integration of your Analysis Cockpit with an existing sandbox. We will not cover how to set up the sandbox.

3.10.1 Analysis Cockpit Sandbox Configuration

In the web view of your Analysis Cockpit, navigate to `Sandbox > Sandboxes`. Click `Add Sandbox` in the top right corner. Keep the `Name` short and add a proper `Description`.

If you wish to enable automatic scanning for uploaded files (Bifrost), you can do so by checking the `Automatic Mode`.

Once you click `Add` the page will display an API token. Copy this token, we will need it later.

Connect to your Analysis Cockpit via SSH and follow the steps below.

Change the user to the root user:

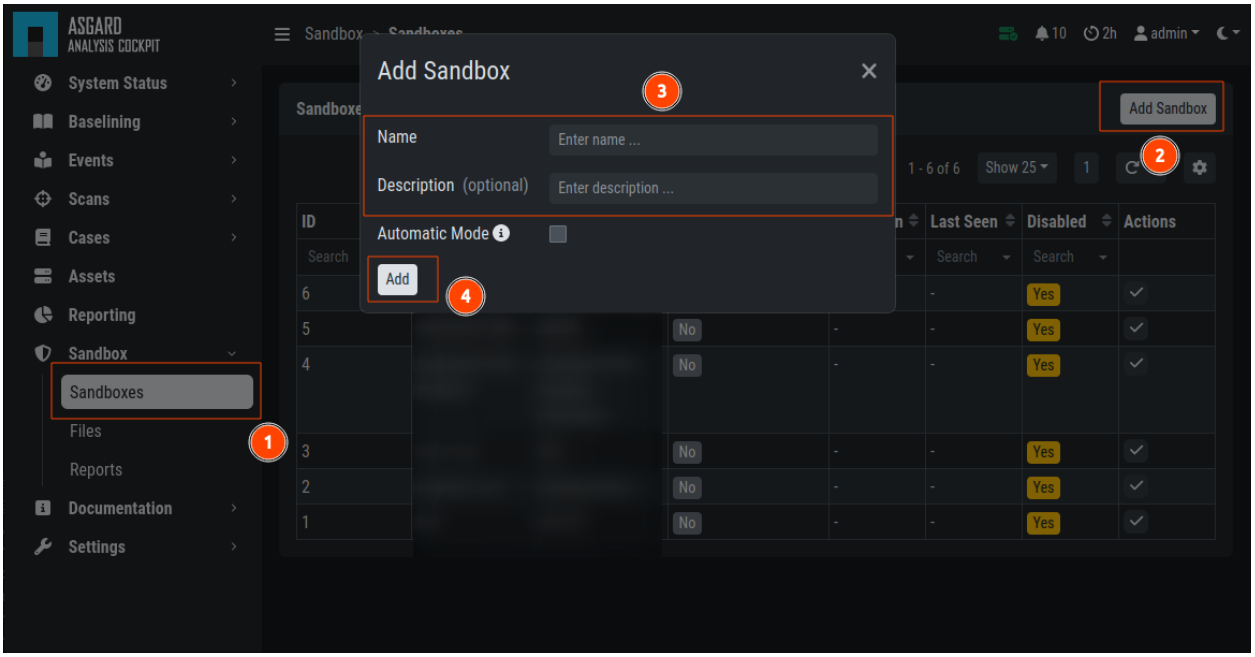

Add Sandbox

Name: Enter name ...

Description (optional): Enter description ...

Automatic Mode □

Add

Fig. 20: Adding a new Sandbox

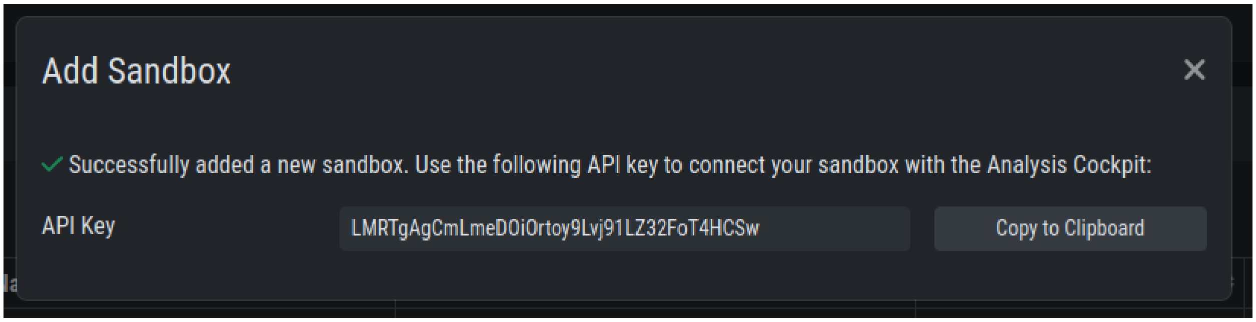

Add Sandbox

✓ Successfully added a new sandbox. Use the following API key to connect your sandbox with the Analysis Cockpit:

API Key: LMRTgAgCmLmeDOiOrtoy9Lvj91LZ32FoT4HCSw

Fig. 21: Sandbox API Token

```
nextron@cockpit:~$ sudo su -
[sudo] password for nextron:
root@cockpit:~#
```

We change into the configuration directory of the Analysis Cockpit:

```
root@cockpit:~# cd /usr/share/asgard-analysis-cockpit/sandbox/connector
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector#
```

Here you can find multiple files and folders. The `.py` and `.ini` files each represent the type of sandbox you want to integrate with. In this example, we will configure the CAPEv2 sandbox with our Analysis Cockpit.

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# ls -lA
total 40
drwxr-xr-x 2 root root 4096 16. Jan 11:20 analysiscockpit
-rw-r--r-- 1 root root  252 16. Jan 10:26 capev2.ini
-rwxr-xr-x 1 root root 9834 16. Jan 10:26 capev2.py
-rw-r--r-- 1 root root  277 16. Jan 10:26 cuckoo.ini
-rw-r--r-- 1 root root 9867 16. Jan 10:26 cuckoo.py
drwxr-xr-x 2 root root 4096 16. Jan 11:20 sandboxapi
```

Here we have two files which are of relevance to us:

- capev2.ini
  - This holds the configuration for both the sandbox and your Analysis Cockpit
- capev2.py
  - This has the systemd configuration to create the actual service on the system (we don't change anything in here)

Change the `capev2.ini` with a text editor. The important lines, which need to be changed accordingly to your environment, are marked:

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# nano capev2.ini
```

```
[DEFAULT]
debug = yes
tmp_directory = /usr/share/asgard-analysis-cockpit/sandbox/capev2

[capev2]
protocol = http
host = 192.168.0.50
port = 8000
token = <your CAPEv2 API token here>
verify = no
all = yes
html = yes

[analysis-cockpit]
host = localhost:443
apikey = <your API Key here>
verify = no
```

For lines 6-10, please fill the information accordingly. `host` is the IP/FQDN of your sandbox. `port` is the listening port of the web interface of your sandbox. `token` is the API token generated in the user management of your sandbox. `verify` is for verification of the TLS certificate (if you don't use TLS or don't want to verify the certificate, set this option to `no`).

For lines 16-17 you have to set the `apikey` of your Analysis Cockpit (see "Add Sandbox" step in the beginning of this section) and `verify`, which can be set to `no`; this will verify the TLS certificate.

Save your files after you made your changes.

Now you have to create a new directory and give the `analysiscockpit` user permission:

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# mkdir /usr/share/
↪asgard-analysis-cockpit/sandbox/capev2
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# chown -R␣
↪analysiscockpit: /usr/share/asgard-analysis-cockpit/sandbox/
```

We need to create a systemd service file in order to run the CAPEv2 connector on your Analysis Cockpit. Below you can find a predefined service file which we will use:

```
[Unit]
Description=CAPEv2 Sandbox Connector
After=network.target

[Service]
ExecStart=/usr/bin/python3 /usr/share/asgard-analysis-cockpit/sandbox/connector/capev2.py
Restart=on-failure
User=analysiscockpit
Group=analysiscockpit
SyslogIdentifier=capev2_connector

[Install]
WantedBy=multi-user.target
```

Now we run the following command and paste the content from the output earlier into it:

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# nano /lib/systemd/
↪system/capev2-connector.service
```

The file should now look like this:

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# cat /lib/systemd/
↪system/capev2-connector.service
[Unit]
Description=CAPEv2 Sandbox Connector
After=network.target

[Service]
ExecStart=/usr/bin/python3 /usr/share/asgard-analysis-cockpit/sandbox/connector/capev2.py
Restart=on-failure
User=analysiscockpit
Group=analysiscockpit
SyslogIdentifier=capev2_connector

[Install]
WantedBy=multi-user.target
```

(continues on next page)

(continued from previous page)

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector#
```

Now that the systemd service file is created, we need to activate it. Run the following command:

```
root@cockpit:/usr/share/asgard-analysis-cockpit/sandbox/connector# systemctl daemon-
↪reload && systemctl enable capev2-connector && systemctl start capev2-connector
Created symlink /etc/systemd/system/multi-user.target.wants/capev2-connector.service → /
↪lib/systemd/system/capev2-connector.service.
```

The connection to your sandbox should work now. You can see the `capev2.log` for debug output and troubleshooting:

```
root@cockpit:~# tail /usr/share/asgard-analysis-cockpit/sandbox/capev2.log
22-11-15 12:07:46 DEBUG: Starting new HTTPS connection (1): localhost:443
22-11-15 12:07:46 DEBUG: https://localhost:443 "GET /api/sandboxes/a/reports/pending?
↪limit=10&offset=0 HTTP/1.1" 200 13
22-11-15 12:07:46 DEBUG: no pending references found
22-11-15 12:08:46 DEBUG: Starting new HTTP connection (1): 192.168.0.50:8000
22-11-15 12:08:46 DEBUG: http://192.168.0.50:8000 "GET /apiv2/cuckoo/status/ HTTP/1.1"␣
↪200 289
22-11-15 12:08:46 DEBUG: Starting new HTTPS connection (1): localhost:443
22-11-15 12:08:46 DEBUG: https://localhost:443 "GET /api/sandboxes/a/get-sha256s-without-
↪report?limit=10 HTTP/1.1" 200 13
22-11-15 12:08:46 DEBUG: Starting new HTTPS connection (1): localhost:443
22-11-15 12:08:46 DEBUG: https://localhost:443 "GET /api/sandboxes/a/reports/pending?
↪limit=10&offset=0 HTTP/1.1" 200 13
22-11-15 12:08:46 DEBUG: no pending references found
root@cockpit:~#
```

3.10.2 Analysis Cockpit Sandbox Usage

Once your sandbox is set up and running, you can see the status of it in the sandbox view (Last Seen):

In the `Files` view you can see previously analyzed files or upload files for analysis by yourself:

Note: If you did not enable `auto mode` of your configured sandbox, you have to manually add the file for scanning in here. You can do this by pressing the `Scan file with sandbox` button to the right of your file.

After your file has been uploaded, you have to wait until your sandbox is finished with analyzing the file. Change to the `Reports` view to see the status of the files.

Once the file was analyzed and the reports are ready, you will see that the status of the file changed to `SUCCESS` and the buttons `REPORT`, `JSON` and `HTML` can be clicked. You can now download the report.

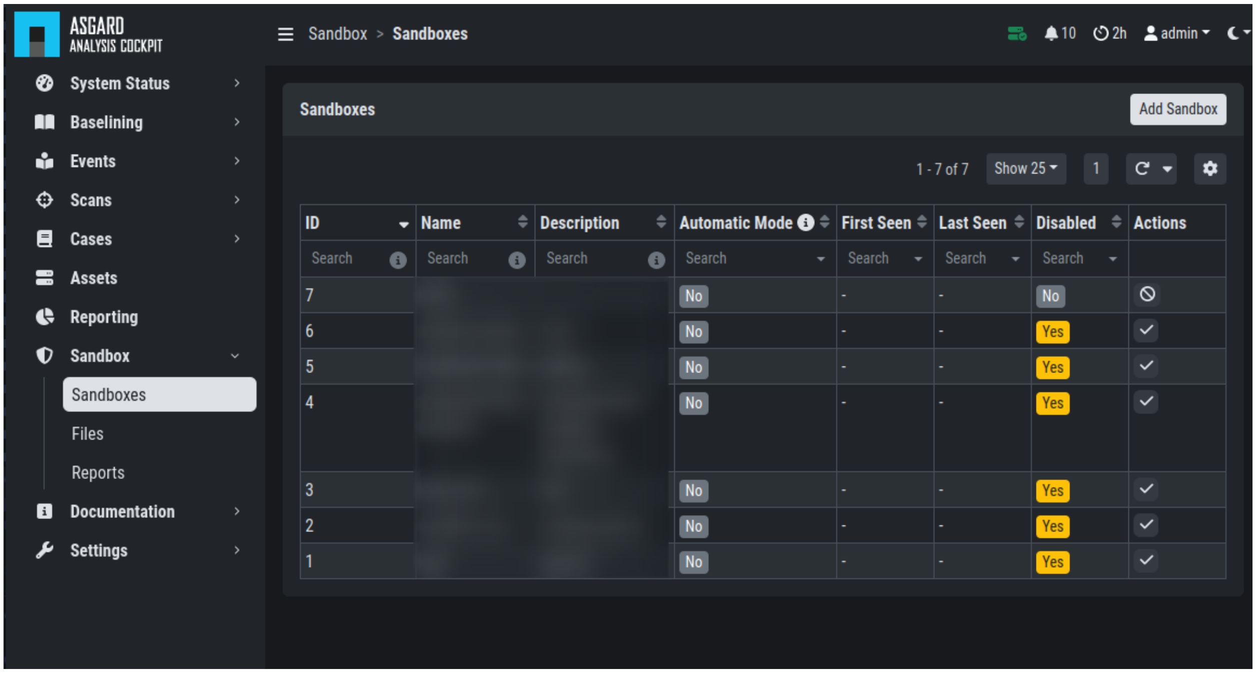

Sandboxes

ID	Name	Description	Automatic Mode	First Seen	Last Seen	Disabled	Actions
Search	Search	Search	Search	Search	Search	Search	
7	[redacted]	[redacted]	No	-	-	No	⊘
6	[redacted]	[redacted]	No	-	-	Yes	✓
5	[redacted]	[redacted]	No	-	-	Yes	✓
4	[redacted]	[redacted]	No	-	-	Yes	✓
3	[redacted]	[redacted]	No	-	-	Yes	✓
2	[redacted]	[redacted]	No	-	-	Yes	✓
1	[redacted]	[redacted]	No	-	-	Yes	✓

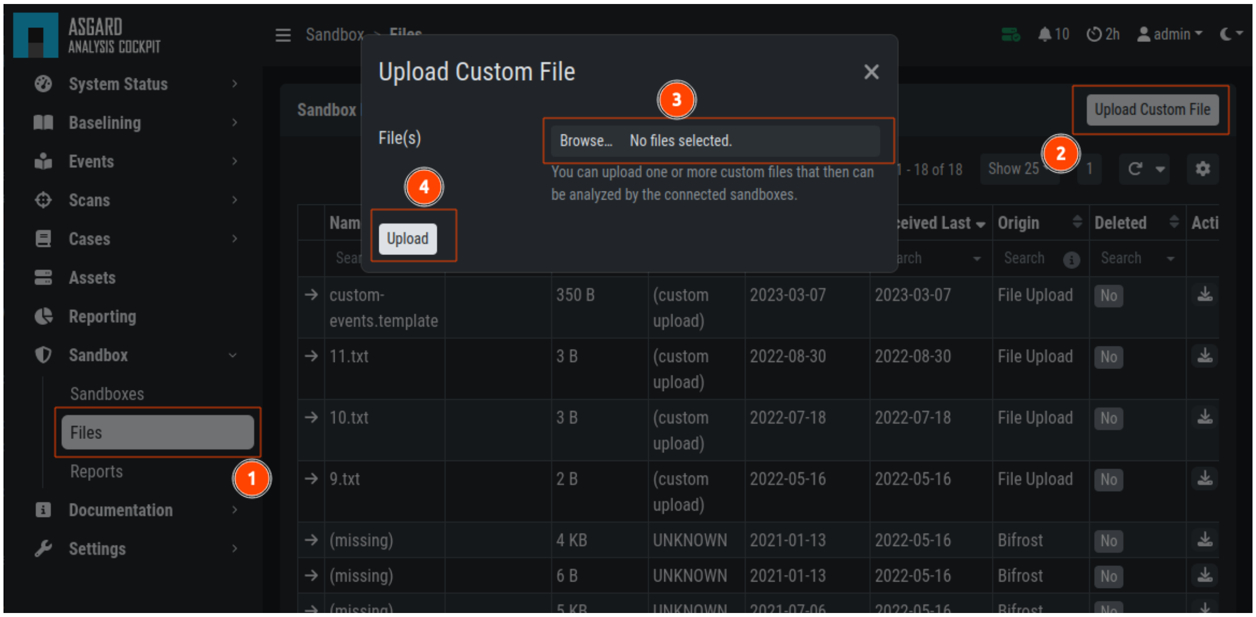

Upload Custom File

File(s): Browse... No files selected.

You can upload one or more custom files that then can be analyzed by the connected sandboxes.

Upload

	Nam				eived Last	Origin	Deleted	Acti
	Sear				arch	Search	Search	
→	custom-events.template	350 B	(custom upload)	2023-03-07	2023-03-07	File Upload	No	⤓
→	11.txt	3 B	(custom upload)	2022-08-30	2022-08-30	File Upload	No	⤓
→	10.txt	3 B	(custom upload)	2022-07-18	2022-07-18	File Upload	No	⤓
→	9.txt	2 B	(custom upload)	2022-05-16	2022-05-16	File Upload	No	⤓
→	(missing)	4 KB	UNKNOWN	2021-01-13	2022-05-16	Bifrost	No	⤓
→	(missing)	6 B	UNKNOWN	2021-01-13	2022-05-16	Bifrost	No	⤓
→	(missing)	5 KB	UNKNOWN	2021-07-06	2022-05-16	Bifrost	No	⤓

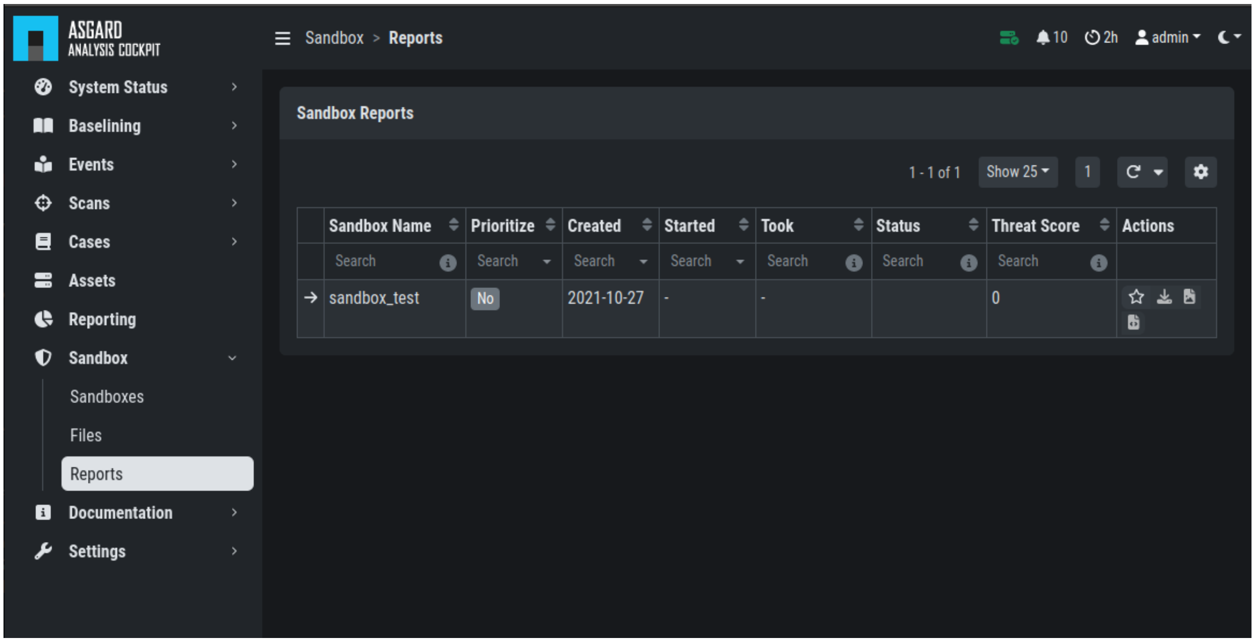

Sandbox Reports

	Sandbox Name	Prioritize	Created	Started	Took	Status	Threat Score	Actions
	Search	Search	Search	Search	Search	Search	Search	
→	sandbox_test	No	2021-10-27	-	-		0	

3.11 API

The API documentation has been integrated into the web interface. You can find it in the `Documentation` menu.

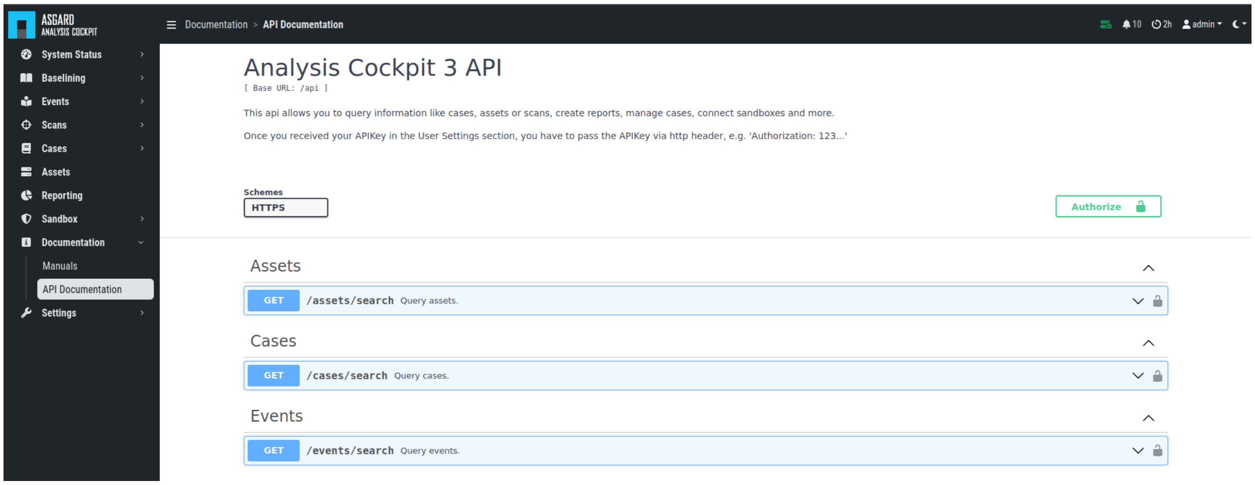

Analysis Cockpit 3 API

[Base URL: /api]

This api allows you to query information like cases, assets or scans, create reports, manage cases, connect sandboxes and more.

Once you received your APIKey in the User Settings section, you have to pass the APIKey via http header, e.g. 'Authorization: 123...'

Assets

GET `/assets/search` Query assets.

Cases

GET `/cases/search` Query cases.

Events

GET `/events/search` Query events.

Fig. 22: API Documentation

CHAPTER

FOUR

BASIC CONCEPTS

In this chapter we will go over the basic concepts of your Analysis Cockpit. We explain how Events and Cases are handled, and how to establish a streamlined workflow within your Analysis Cockpit.

4.1 Events

All events that have been stored in your Analysis Cockpit – regardless if they are assigned to a particular case or not – are displayed in the section `Events`. This section can be seen as your threat hunting pool. The section provides powerful filtering options. The Events Section is split into the different sources of your Events:

- THOR Events
- Aurora Events
- Log Watcher Events (deprecated)

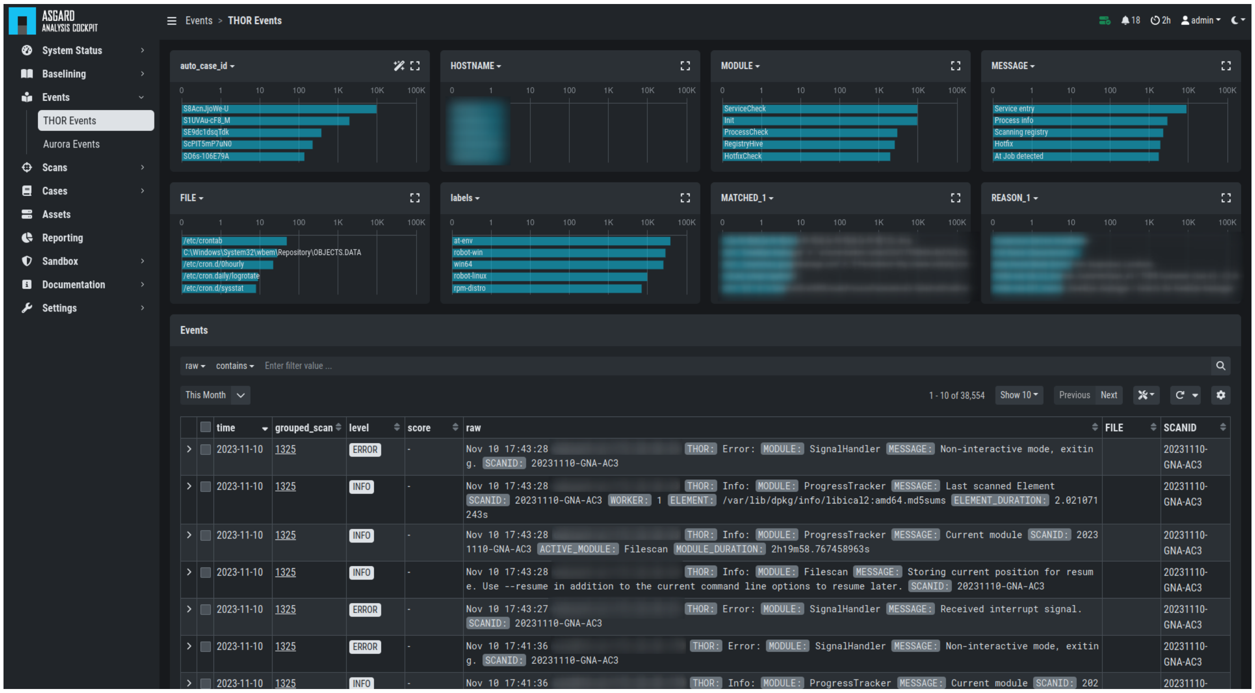

	time	grouped_scan	level	score	raw	FILE	SCANID
>	2023-11-10	1325	ERROR	-	Nov 10 17:43:28 [redacted] THOR: Error: MODULE: SignalHandler MESSAGE: Non-interactive mode, exiting. SCANID: 20231110-GNA-AC3		20231110-GNA-AC3
>	2023-11-10	1325	INFO	-	Nov 10 17:43:28 [redacted] THOR: Info: MODULE: ProgressTracker MESSAGE: Last scanned Element SCANID: 20231110-GNA-AC3 WORKER: 1 ELEMENT: /var/lib/dpkg/info/libical2:amd64.md5sums ELEMENT_DURATION: 2.021071243s		20231110-GNA-AC3
>	2023-11-10	1325	INFO	-	Nov 10 17:43:28 [redacted] THOR: Info: MODULE: ProgressTracker MESSAGE: Current module SCANID: 20231110-GNA-AC3 ACTIVE_MODULE: Filescan MODULE_DURATION: 2h19m58.767458963s		20231110-GNA-AC3
>	2023-11-10	1325	INFO	-	Nov 10 17:43:28 [redacted] THOR: Info: MODULE: Filescan MESSAGE: Storing current position for resume. Use --resume in addition to the current command line options to resume later. SCANID: 20231110-GNA-AC3		20231110-GNA-AC3
>	2023-11-10	1325	ERROR	-	Nov 10 17:43:27 [redacted] THOR: Error: MODULE: SignalHandler MESSAGE: Received interrupt signal. SCANID: 20231110-GNA-AC3		20231110-GNA-AC3
>	2023-11-10	1325	ERROR	-	Nov 10 17:41:36 [redacted] THOR: Error: MODULE: SignalHandler MESSAGE: Non-interactive mode, exiting. SCANID: 20231110-GNA-AC3		20231110-GNA-AC3
>	2023-11-10	1325	INFO	-	Nov 10 17:41:36 [redacted] THOR: Info: MODULE: ProgressTracker MESSAGE: Current module SCANID: 202		20231110-

Fig. 1: Events Section

4.2 Baselining

All events that have **not** been assigned to a particular case are displayed in the `Baselining` section of the Analysis Cockpit.

Again, the Baselining Section is split into the different sources of our events. Additionally, you can see the `Suggested Cases`, which will suggest cases based on predefined *Case Templates*.

- THOR Events
- Aurora Events
- Log Watcher Events (deprecated)
- Suggested Cases

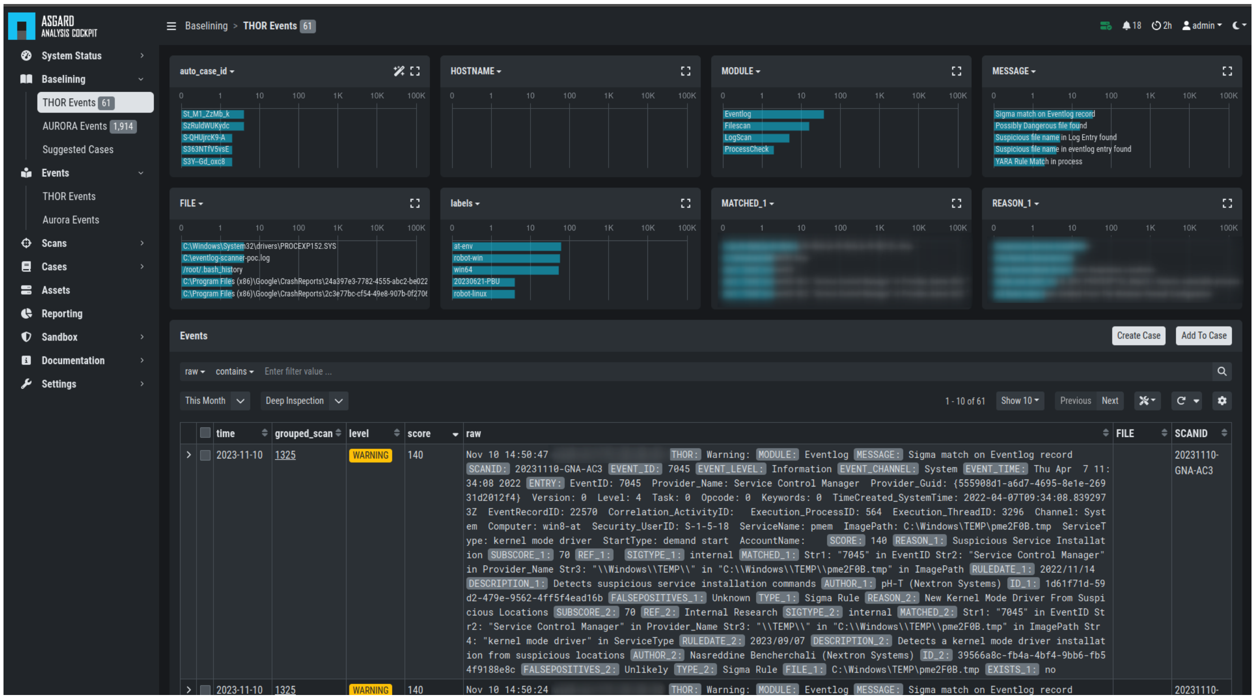

time	grouped_scan	level	score	raw	FILE	SCANID
2023-11-10	1325	WARNING	140	Nov 10 14:50:47 [redacted] THOR: Warning: MODULE: Eventlog MESSAGE: Sigma match on Eventlog record SCANID: 20231110-GNA-AC3 EVENT_ID: 7045 EVENT_LEVEL: Information EVENT_CHANNEL: System EVENT_TIME: Thu Apr 7 11:34:08 2022 ENTRY: EventID: 7045 Provider_Name: Service Control Manager Provider_Guid: {555908d1-a6d7-4695-8e1e-26931d2012f4} Version: 0 Level: 4 Task: 0 Opcode: 0 Keywords: 0 TimeCreated_SystemTime: 2022-04-07T09:34:08.8392973Z EventRecordID: 22570 Correlation_ActivityID: Execution_ProcessID: 564 Execution_ThreadID: 3296 Channel: System Computer: win8-at Security_UserID: S-1-5-18 ServiceName: pmem ImagePath: C:\Windows\TEMP\pme2F0B.tmp ServiceType: kernel mode driver StartType: demand start AccountName: SCORE: 140 REASON_1: Suspicious Service Installation SUBSCORE_1: 70 REF_1: SIGTYPE_1: internal MATCHED_1: Str1: "7045" in EventID Str2: "Service Control Manager" in Provider_Name Str3: "\\Windows\\TEMP\\" in "C:\\Windows\\TEMP\\pme2F0B.tmp" in ImagePath RULEDATE_1: 2022/11/14 DESCRIPTION_1: Detects suspicious service installation commands AUTHOR_1: pH-T (Nextron Systems) ID_1: 1d61f71d-59d2-479e-9562-4ff5f4ead16b FALSEPOSITIVES_1: Unknown TYPE_1: Sigma Rule REASON_2: New Kernel Mode Driver From Suspicious Locations SUBSCORE_2: 70 REF_2: Internal Research SIGTYPE_2: internal MATCHED_2: Str1: "7045" in EventID Str2: "Service Control Manager" in Provider_Name Str3: "\\TEMP\\" in "C:\\Windows\\TEMP\\pme2F0B.tmp" in ImagePath Str4: "kernel mode driver" in ServiceType RULEDATE_2: 2023/09/07 DESCRIPTION_2: Detects a kernel mode driver installation from suspicious locations AUTHOR_2: Nasreddine Bencherchali (Nextron Systems) ID_2: 39566a8c-fb4a-4bf4-9bb6-fb54f9188e8c FALSEPOSITIVES_2: Unlikely TYPE_2: Sigma Rule FILE_1: C:\Windows\TEMP\pme2F0B.tmp EXISTS_1: no		20231110-GNA-AC3
2023-11-10	1325	WARNING	140	Nov 10 14:50:24 [redacted] THOR: Warning: MODULE: Eventlog MESSAGE: Sigma match on Eventlog record		20231110-

Fig. 2: Baselining Section

Logs that represent the same type of anomaly or incident can be grouped together using the various filters and then be stored in a Case for further analysis. Grouping can be done manually by filtering and clicking `Create Case`, selecting individual Events and clicking `Create Case`, or automatically by simply clicking the `Advanced Tools` button and `Auto Baselining`. With `Auto Baselining`, the Cockpit automatically calculates groups of "similar" log lines.

Once stored in a case, the logs will disappear from the Baselining section.

The Analysis Cockpit can automatically check for events that can be added to existing cases. By clicking the `Optimize` button, the Analysis Cockpit will iterate through all unassigned events and check if there is a matching case.

Note: The optimization will iterate through all unassigned events and assign them to cases if a match were found. This may take a while.

In an ideal organization, the Baselining section should always be empty at the end of a day, as these logs represent suspicious elements that have not yet been looked at.

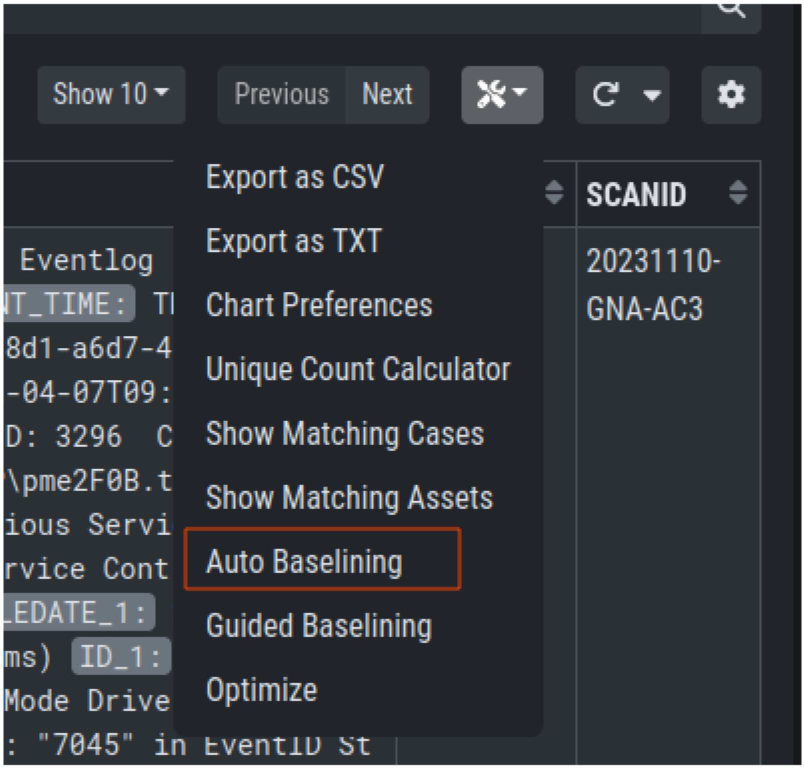

Fig. 3: Auto Baselining

4.2.1 Baselining Views

In the `Baselining` section there are two main views, the `Compromise Assessment Mode` and the `Deep Inspection Mode`. Additionally, you can find the `Custom Signatures Only Mode`, which will only show events found by custom signatures. This can be helpful if you scanned your environment with customer signatures, for example during or after an incident.

By default, the Analysis Cockpit Baselining Mode is set to `Compromise Assessment`.

4.2.2 Compromise Assessment Mode

The `Compromise Assessment Mode` is a new filter/representation of events created and reviewed by our security experts.

It includes our most successful detections. In this context "success" means, that the detection uncovered malicious activity in the wild and at the same time had a low anomaly and false positive rate. Additionally we also consider a detection to be successful that caused little or no false positives or anomalies.

The new view will combine and apply different techniques and filters to all the unclassified events in the `Baselining` section, providing a reduced set of logs which proved to be relevant from an analyst perspective.

This new "Compromise Assessment Mode" dramatically reduces your baselining effort. In our tests we noticed a decrease of events in the Baselining section of more than 90%. We believe that especially entities that follow our "Continuous Compromise Assessment" approach should switch into this new mode. We've also challenged the new mode with the post exploitation tools and techniques found in the context of HAFNIUM / Exchange exploitations in March 2021 and covered almost every aspect of the attacks in the new view.

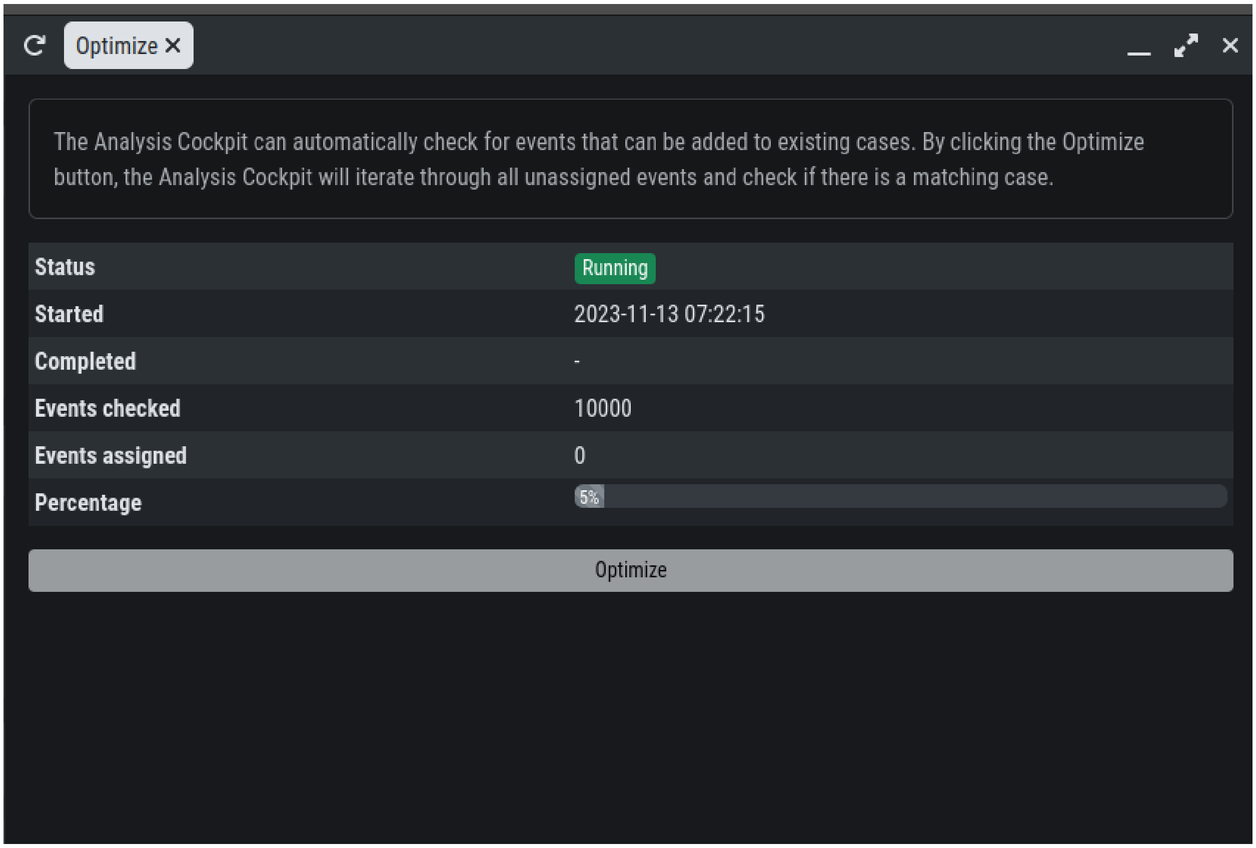

The Analysis Cockpit can automatically check for events that can be added to existing cases. By clicking the Optimize button, the Analysis Cockpit will iterate through all unassigned events and check if there is a matching case.

Status	Running
Started	2023-11-13 07:22:15
Completed	-
Events checked	10000
Events assigned	0
Percentage	5%

Fig. 4: Optimize Function

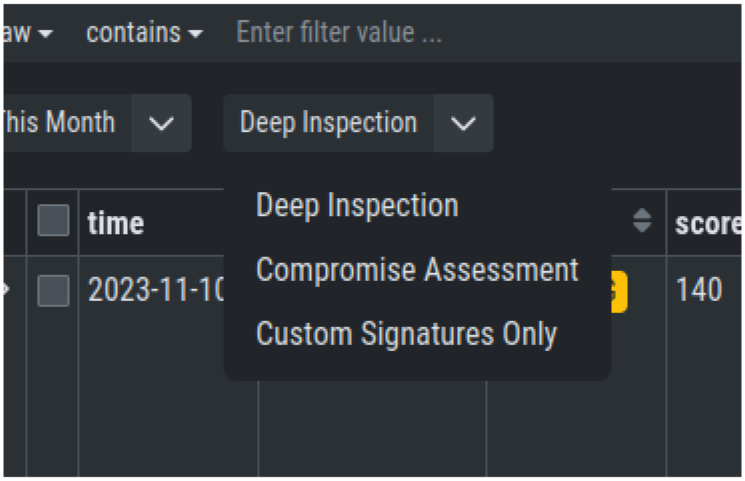

Fig. 5: Select your view

Note: In case of an Incident Response, the `Deep Inspection Mode` is always recommended, since nothing is filtered here.

4.2.3 Deep Inspection Mode

This view is basically how it used to be (the old default view). It shows all Alerts and Warnings unless they are already part of an existing case.

4.2.4 Custom Signatures Only Mode

The `Custom Signatures Only` view will only show you events, which:

- Are not part of a case
- Where found by a custom signature

This view can be helpful if you only want to see events found by one of your custom signatures during a THOR scan. This can be helpful if you want to see only those events and nothing else.

4.3 Cases and Log Processing

The Cases section gives a good overview regarding the existing cases and also provides various filtering options. Column visibility can be configured by clicking on the Columns button of this section.

The Cases Section is split into the different sources of your Cases:

- THOR Cases
- Aurora Cases
- Log Watcher Cases (deprecated)

Additionally, you can find more information regarding:

- Grouping Criteria
- Case Changes
- Security Center

When a case is created, the state will be "Open" and the type will be set to "Noteworthy" by default.

The following states can be set (by default):

- Open
- Level 1 Finished
- Level 1 Working
- Level 2 Working
- Closed

It is possible to configure custom states.

The following types can be set:

- Incident

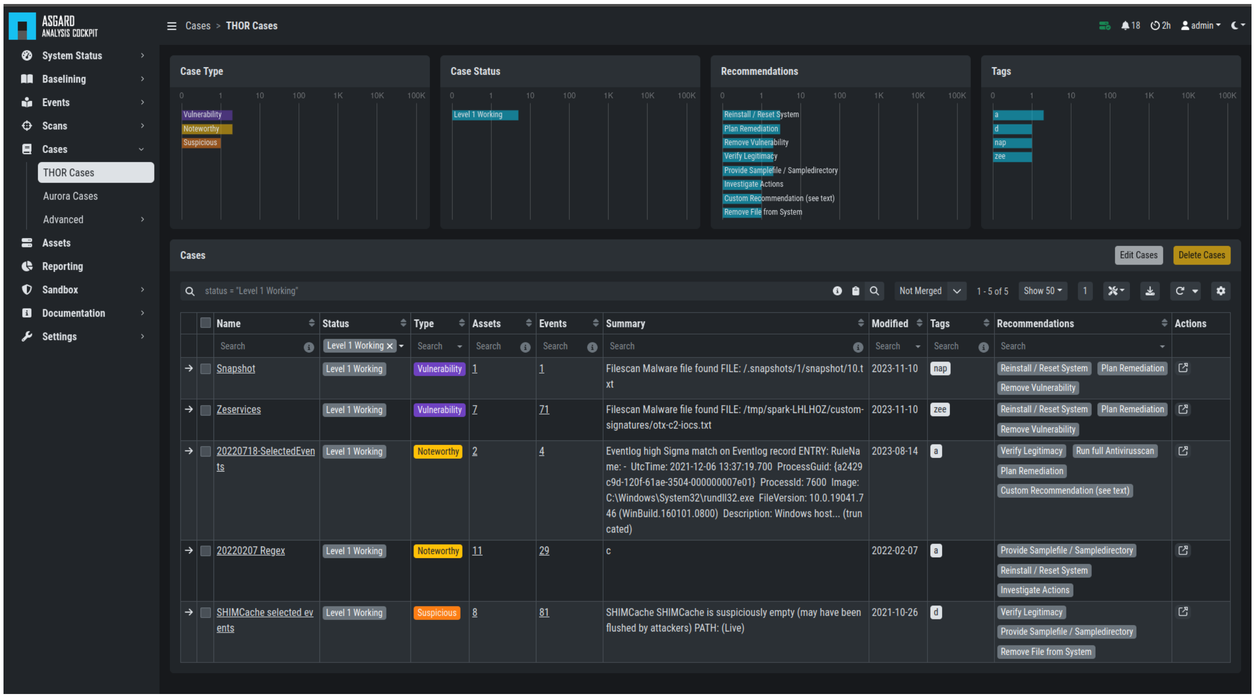

	Name	Status	Type	Assets	Events	Summary	Modified	Tags	Recommendations	Actions
	Search	Level 1 Working ×	Search	Search	Search	Search	Search	Search	Search	
→	Snapshot	Level 1 Working	Vulnerability	1	1	Filescan Malware file found FILE: /.snapshots/1/snapshot/10.txt	2023-11-10	nap	Reinstall / Reset System Plan Remediation Remove Vulnerability	
→	Zeservices	Level 1 Working	Vulnerability	7	71	Filescan Malware file found FILE: /tmp/spark-LHLHOZ/custom-signatures/otx-c2-iocs.txt	2023-11-10	zee	Reinstall / Reset System Plan Remediation Remove Vulnerability	
→	20220718-SelectedEvents	Level 1 Working	Noteworthy	2	4	Eventlog high Sigma match on Eventlog record ENTRY: RuleName: - UtcTime: 2021-12-06 13:37:19.700 ProcessGuid: {a2429c9d-120f-61ae-3504-000000007e01} ProcessId: 7600 Image: C:\Windows\System32\rundll32.exe FileVersion: 10.0.19041.746 (WinBuild.160101.0800) Description: Windows host... (truncated)	2023-08-14	a	Verify Legitimacy Run full Antivirusscan Plan Remediation Custom Recommendation (see text)	
→	20220207 Regex	Level 1 Working	Noteworthy	11	29	c	2022-02-07	a	Provide Samplefile / Sampledirectory Reinstall / Reset System Investigate Actions	
→	SHIMCache selected events	Level 1 Working	Suspicious	8	81	SHIMCache SHIMCache is suspiciously empty (may have been flushed by attackers) PATH: (Live)	2021-10-26	d	Verify Legitimacy Provide Samplefile / Sampledirectory Remove File from System	

Fig. 6: Cases Section

- Suspicious
- Vulnerability
- Noteworthy
- Unknown
- Legitimate Anomaly
- False Positive

See chapter *Glossary* for a detailed description of these terms.

Within a case, it is possible to add various information, write a summary, provide canned recommendations or add assessment information.

The log lines contained in the case can of course be analyzed in detail and changes to the case are tracked automatically.

The cockpit will automatically calculate rules (auto_case_id), that make sure, future incoming logs that are similar to the log lines in this particular case are automatically assigned to this case and **will not show up** in the `Baselining` section.

Important: `Automatically assign newly incoming events to this case` needs to be selected during case creation to automatically assign new events to an existing case.

In order to understand this better, let's assume you have decided a group of logs are legitimate anomalies. Then all future logs that are similar to these anomalies will automatically be added to this case and not show up in the Baselining section.

In case you have decided a group of log lines represent a security incident, the same thing will happen. Future log lines that represent a security incident will show up only in the case and not in the Baselining section.

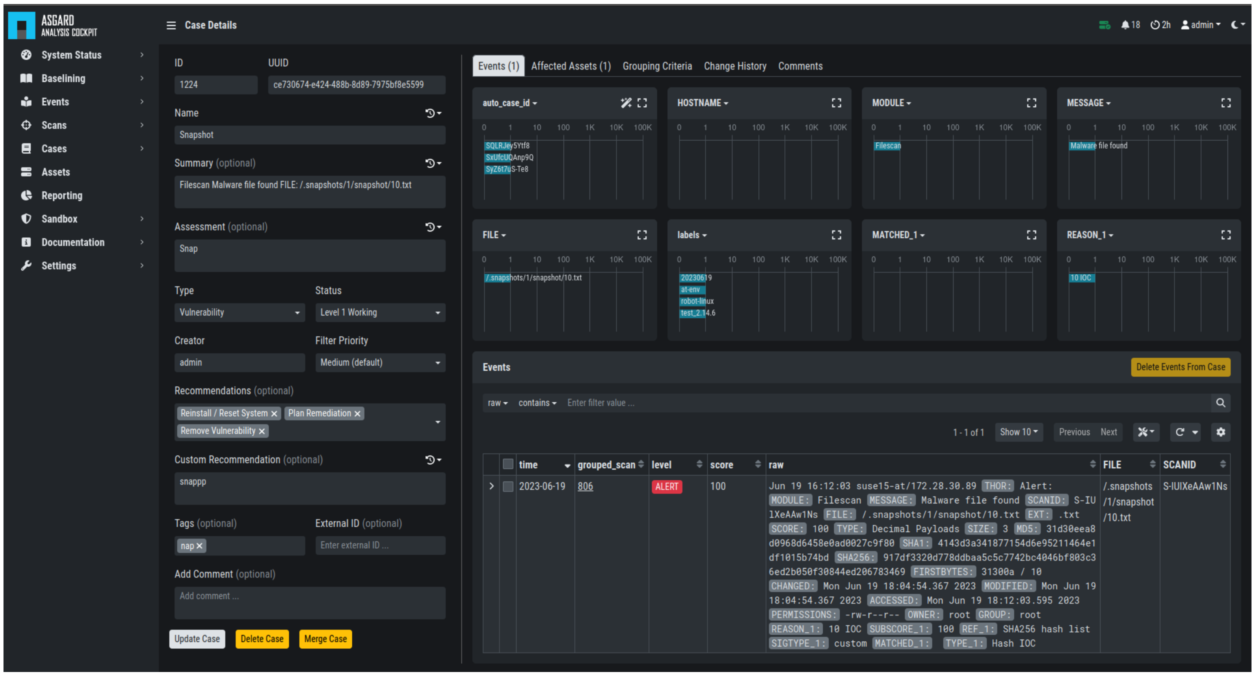

ID: 1224

UUID: ce730674-e424-488b-8d89-7975bf8e5599

Name: Snapshot

Summary (optional): Filescan Malware file found FILE: /.snapshots/1/snapshot/10.txt

Assessment (optional): Snap

Type: Vulnerability

Status: Level 1 Working

Creator: admin

Filter Priority: Medium (default)

Recommendations (optional): Reinstall / Reset System ✕ Plan Remediation ✕ Remove Vulnerability ✕

Custom Recommendation (optional): snappp

Tags (optional): nap ✕

External ID (optional): Enter external ID ...

Add Comment (optional): Add comment ...

Events

time	grouped_scan	level	score	raw	FILE	SCANID
2023-06-19	806	ALERT	100	Jun 19 16:12:03 suse15-at/172.28.30.89 THOR: Alert: MODULE: Filescan MESSAGE: Malware file found SCANID: S-IU1XeAAw1Ns FILE: /.snapshots/1/snapshot/10.txt EXT: .txt SCORE: 100 TYPE: Decimal Payloads SIZE: 3 MD5: 31d30eea8d0968d6458e0ad0027c9f80 SHA1: 4143d3a341877154d6e95211464e1df1015b74bd SHA256: 917df3320d778ddbaa5c5c7742bc4046bf803c36ed2b050f30844ed206783469 FIRSTBYTES: 31300a / 10 CHANGED: Mon Jun 19 18:04:54.367 2023 MODIFIED: Mon Jun 19 18:04:54.367 2023 ACCESSED: Mon Jun 19 18:12:03.595 2023 PERMISSIONS: -rw-r--r-- OWNER: root GROUP: root REASON_1: 10 IOC SUBSCORE_1: 100 REF_1: SHA256 hash list SIGTYPE_1: custom MATCHED_1: TYPE_1: Hash IOC	/.snapshots/1/snapshot/10.txt	S-lUlXeAAw1Ns

Fig. 7: Case Details

Most organizations want to be alerted in case of a security incident. The Cockpit can be configured to forward all logs that are automatically assigned to an incident case to the organizations' SIEM System via syslog. Organizations that prefer to handle THOR Events entirely within the Analysis Cockpit and not forward anything to a SIEM system may choose to configure a notification that shows up in the Cockpit's Notification Section.

The following picture shows the recommended log processing.

As one can see, an incoming log line only shows up in the `Baselining` section when it matches no existing case.

This behavior is highly configurable and can be changed in the `Settings` section of the Analysis Cockpit. One can even decide not to forward anything to a SIEM System or may decide to also forward suspicious elements in addition.

In other Words:

Cases represent the means of setting and maintaining the log baseline within the Cockpit. When you scan your infrastructure once, assign all logs to cases and then scan it for the second time, the `Baselining` section should be empty if nothing has changed. All incoming logs should be similar to the ones in the first scan and therefore be assigned to the respective cases and not show up in the `Baselining` section.

Working with cases is explained in detail in the sections below.

4.3.1 Case Templates

Case Templates can be used to suggest new cases in the `Suggested Cases` section. If there are no Suggested Cases in the view, no events match the Case Templates in your Analysis Cockpit.

To import new Case Templates, you need to create a `.yaml` file with the conditions first. This can be done by navigating to the `Cases` view and exporting your search results as Case Templates. You will be able to download a `.yaml` file from here, which can be used to import as a Case Template.

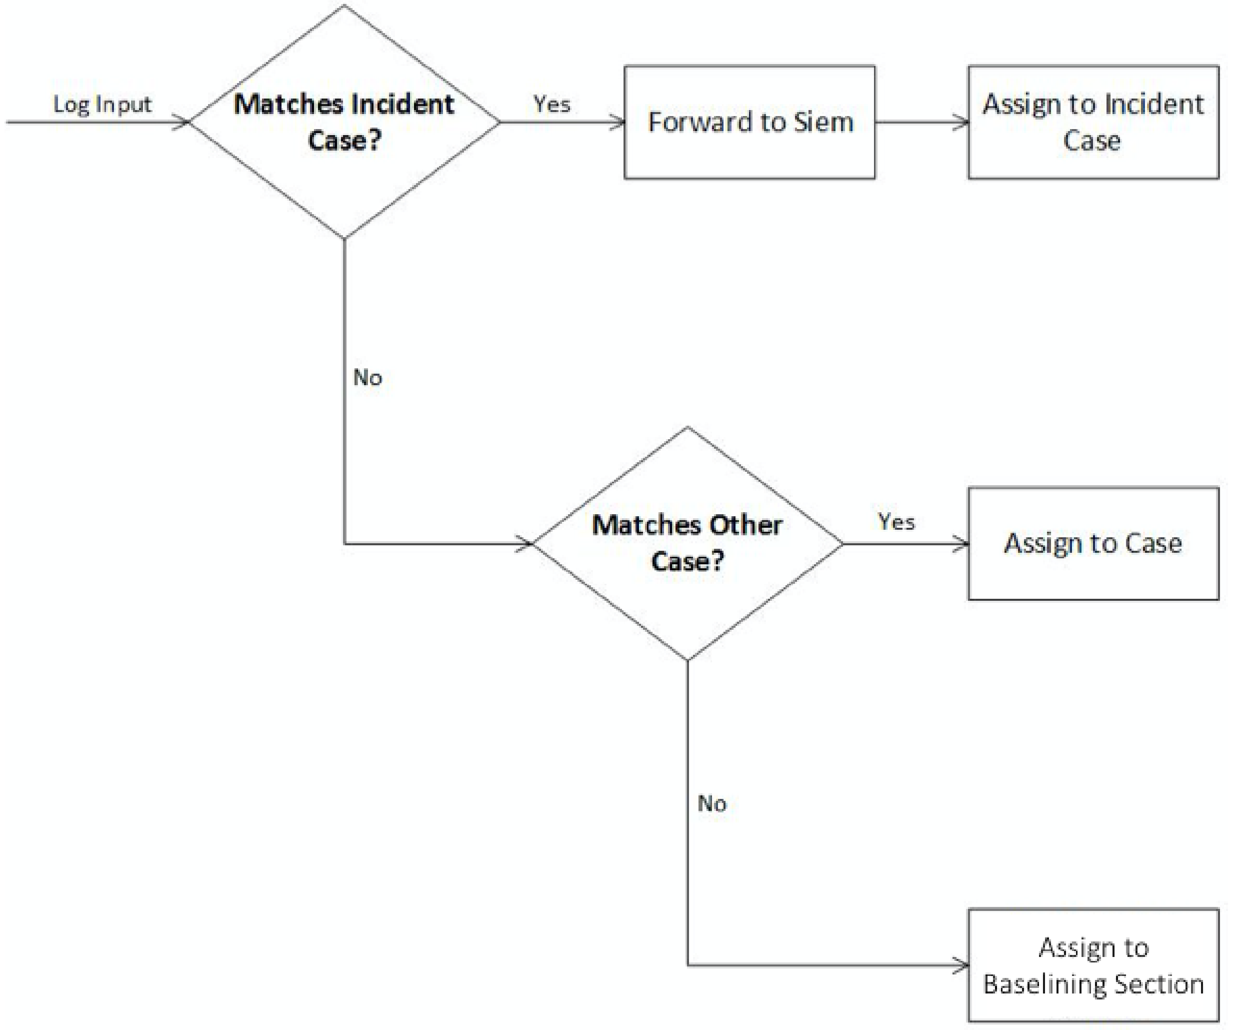

Fig. 8: Log Processing

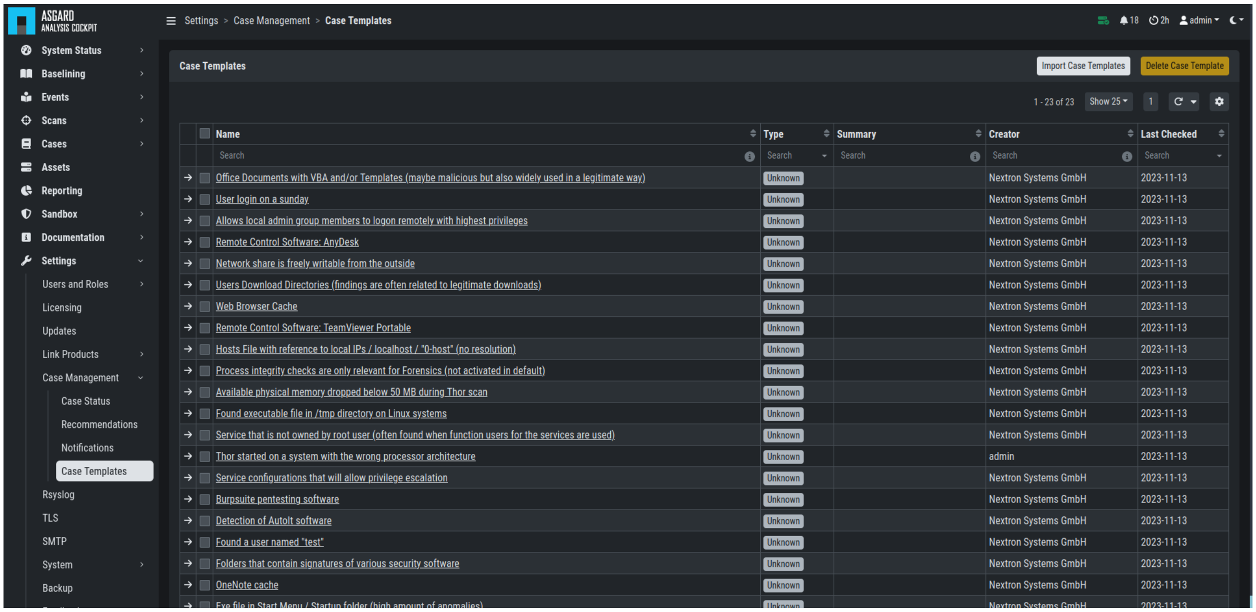

Case Templates

Name	Type	Summary	Creator	Last Checked
Office Documents with VBA and/or Templates (maybe malicious but also widely used in a legitimate way)	Unknown		Nextron Systems GmbH	2023-11-13
User login on a sunday	Unknown		Nextron Systems GmbH	2023-11-13
Allows local admin group members to logon remotely with highest privileges	Unknown		Nextron Systems GmbH	2023-11-13
Remote Control Software: AnyDesk	Unknown		Nextron Systems GmbH	2023-11-13
Network share is freely writable from the outside	Unknown		Nextron Systems GmbH	2023-11-13
Users Download Directories (findings are often related to legitimate downloads)	Unknown		Nextron Systems GmbH	2023-11-13
Web Browser Cache	Unknown		Nextron Systems GmbH	2023-11-13
Remote Control Software: TeamViewer Portable	Unknown		Nextron Systems GmbH	2023-11-13
Hosts File with reference to local IPs / localhost / "0-host" (no resolution)	Unknown		Nextron Systems GmbH	2023-11-13
Process integrity checks are only relevant for Forensics (not activated in default)	Unknown		Nextron Systems GmbH	2023-11-13
Available physical memory dropped below 50 MB during Thor scan	Unknown		Nextron Systems GmbH	2023-11-13
Found executable file in /tmp directory on Linux systems	Unknown		Nextron Systems GmbH	2023-11-13
Service that is not owned by root user (often found when function users for the services are used)	Unknown		Nextron Systems GmbH	2023-11-13
Thor started on a system with the wrong processor architecture	Unknown		admin	2023-11-13
Service configurations that will allow privilege escalation	Unknown		Nextron Systems GmbH	2023-11-13
Burpsuite pentesting software	Unknown		Nextron Systems GmbH	2023-11-13
Detection of AutoIt software	Unknown		Nextron Systems GmbH	2023-11-13
Found a user named "test"	Unknown		Nextron Systems GmbH	2023-11-13
Folders that contain signatures of various security software	Unknown		Nextron Systems GmbH	2023-11-13
OneNote cache	Unknown		Nextron Systems GmbH	2023-11-13
Exe file in Start Menu / Startup folder (high amount of anomalies)	Unknown		Nextron Systems GmbH	2023-11-13

Fig. 9: Case Templates Overview

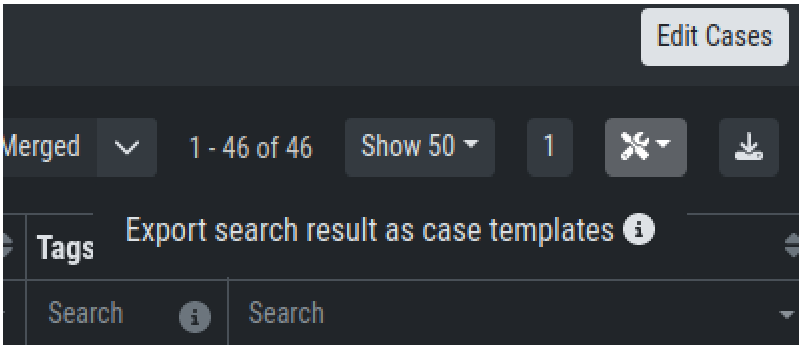

Fig. 10: Exporting Search Results as Case Templates

Listing 1: Exported Case Template

```yaml
uuid: 94565b82-45fc-47f8-82eb-e9c5352c37c2
name: Thor started on a system with the wrong processor architecture
summary: ""
type: 5
scanner: THOR
creator: admin
condition: "\"MODULE: Startup\" AND \"MESSAGE: 32 bit THOR was executed on 64 bit
  system. For improved results, use the 64 bit version of THOR.\"\r\n"
```

After you downloaded the Case Templates, you can import them in the `Case Templates` view.

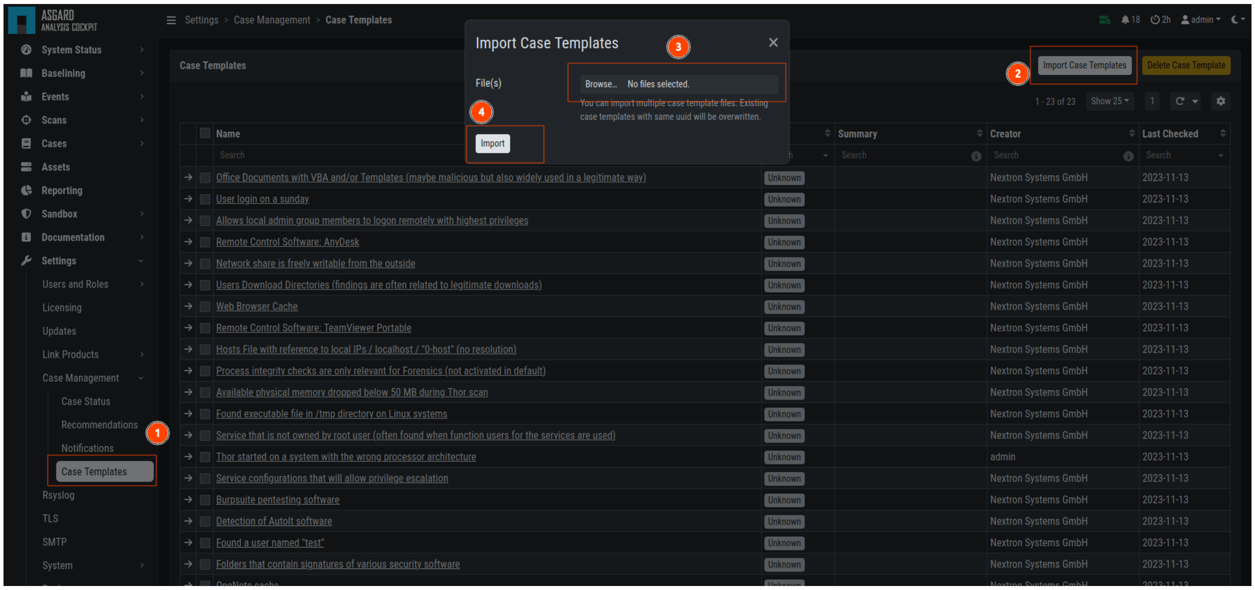

Import Case Templates

File(s): Browse... No files selected.

You can import multiple case template files. Existing case templates with same uuid will be overwritten.

Import

Name		Summary	Creator	Last Checked
Office Documents with VBA and/or Templates (maybe malicious but also widely used in a legitimate way)	Unknown		Nextron Systems GmbH	2023-11-13
User login on a sunday	Unknown		Nextron Systems GmbH	2023-11-13
Allows local admin group members to logon remotely with highest privileges	Unknown		Nextron Systems GmbH	2023-11-13
Remote Control Software: AnyDesk	Unknown		Nextron Systems GmbH	2023-11-13
Network share is freely writable from the outside	Unknown		Nextron Systems GmbH	2023-11-13
Users Download Directories (findings are often related to legitimate downloads)	Unknown		Nextron Systems GmbH	2023-11-13
Web Browser Cache	Unknown		Nextron Systems GmbH	2023-11-13
Remote Control Software: TeamViewer Portable	Unknown		Nextron Systems GmbH	2023-11-13
Hosts File with reference to local IPs / localhost / "0-host" (no resolution)	Unknown		Nextron Systems GmbH	2023-11-13
Process integrity checks are only relevant for Forensics (not activated in default)	Unknown		Nextron Systems GmbH	2023-11-13
Available physical memory dropped below 50 MB during Thor scan	Unknown		Nextron Systems GmbH	2023-11-13
Found executable file in /tmp directory on Linux systems	Unknown		Nextron Systems GmbH	2023-11-13
Service that is not owned by root user (often found when function users for the services are used)	Unknown		Nextron Systems GmbH	2023-11-13
Thor started on a system with the wrong processor architecture	Unknown		admin	2023-11-13
Service configurations that will allow privilege escalation	Unknown		Nextron Systems GmbH	2023-11-13
Burpsuite pentesting software	Unknown		Nextron Systems GmbH	2023-11-13
Detection of AutoIt software	Unknown		Nextron Systems GmbH	2023-11-13
Found a user named "test"	Unknown		Nextron Systems GmbH	2023-11-13
Folders that contain signatures of various security software	Unknown		Nextron Systems GmbH	2023-11-13
OneNote cache	Unknown		Nextron Systems GmbH	2023-11-13

Fig. 11: Import Case Template

You can now inspect the Case Template. You can find it by either looking for the name or filter by who created it. You can see that the conditions match the contents of your exported Case Template (`.yaml` file).

4.4 Understanding Users, Roles, Rights and Case Status

The rights and roles model within the cockpit is aimed to support large multinational organizations with different independent users working with the case management at the same time. An organization responsible for analyzing THOR logs might be split up in groups of analysts.

Within the cockpit, all users have the right to access the logs and create cases. Within the `Case Management` section, access rights are granted depending on the particular state the case is in.

In order to setup your rights management you must first decide about the states you want your cases to have, then assign rights for a particular state to a role and after that you add users to that particular role.

In order to understand this better, let's look at an example.

Let's assume we have an organization where a Level 1 analyst group located in Frankfurt is responsible for creating cases and providing an initial assessment for cases, while a Level 2 analyst group located in Hamburg is responsible

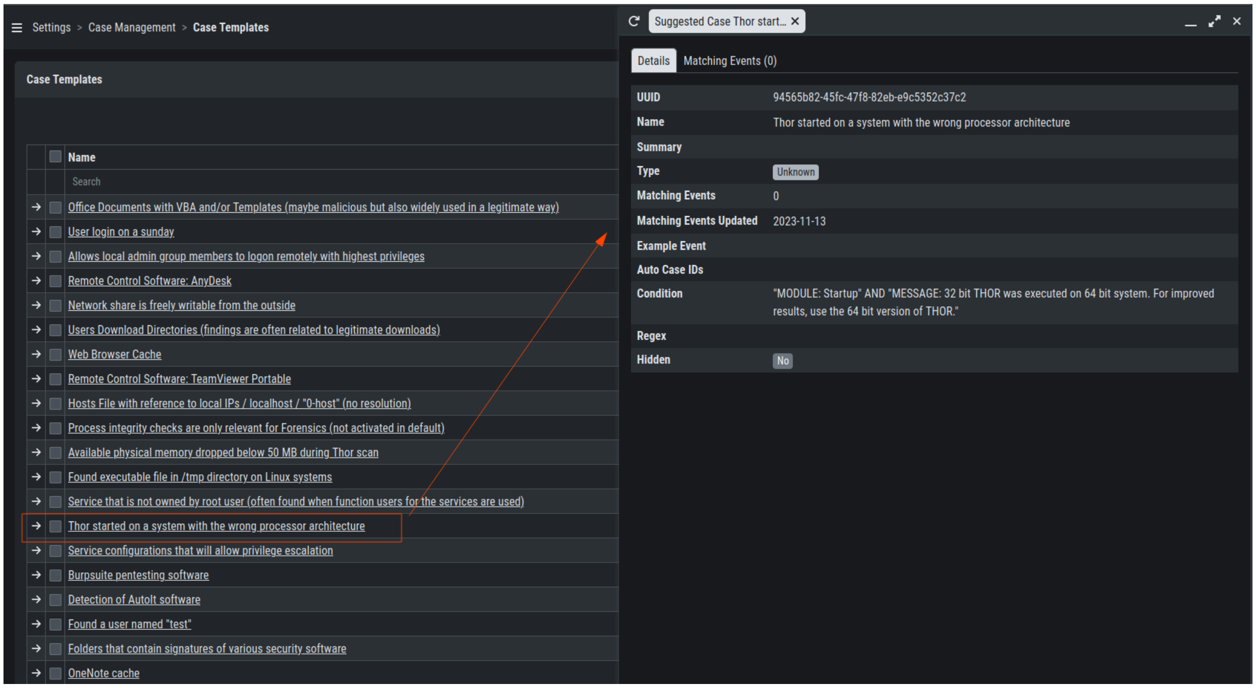
Settings > Case Management > Case Templates
Case Templates
Suggested Case Thor start... ×
Details
Matching Events (0)

Name
Search
Office Documents with VBA and/or Templates (maybe malicious but also widely used in a legitimate way)
User login on a sunday
Allows local admin group members to logon remotely with highest privileges
Remote Control Software: AnyDesk
Network share is freely writable from the outside
Users Download Directories (findings are often related to legitimate downloads)
Web Browser Cache
Remote Control Software: TeamViewer Portable
Hosts File with reference to local IPs / localhost / "0-host" (no resolution)
Process integrity checks are only relevant for Forensics (not activated in default)
Available physical memory dropped below 50 MB during Thor scan
Found executable file in /tmp directory on Linux systems
Service that is not owned by root user (often found when function users for the services are used)
Thor started on a system with the wrong processor architecture
Service configurations that will allow privilege escalation
Burpsuite pentesting software
Detection of AutoIt software
Found a user named "test"
Folders that contain signatures of various security software
OneNote cache

UUID	94565b82-45fc-47f8-82eb-e9c5352c37c2
Name	Thor started on a system with the wrong processor architecture
Summary	
Type	Unknown
Matching Events	0
Matching Events Updated	2023-11-13
Example Event	
Auto Case IDs	
Condition	"MODULE: Startup" AND "MESSAGE: 32 bit THOR was executed on 64 bit system. For improved results, use the 64 bit version of THOR."
Regex	
Hidden	No

Fig. 12: Inspect Imported Case Template

for reviewing, final decision and closing of cases. In order to support an efficient workflow, you would at least need the following states for your cases:

- Open (nobody is yet working on this case)
- Level 1 Working (Level 1 is working on this case)
- Level 1 Finished (Level 1 has finished and nobody is now working on this case)
- Level 2 Working (Level 2 is working on this case)
- Closed (Case closed)

A workflow could look like this:

For your convenience, we already did the setup for this example and ship all Analysis Cockpit with this workable template by default. You are free to use, modify or delete the corresponding rights, statuses and roles.

However, in order to explain the concepts and the setup of roles and statuses better we assume for a while, we had an empty cockpit with no roles and statuses pre-configured.

In order to set up our pre-configured example, we navigate to the `Settings` section and create the following roles:

Every role can have different rights. We will explain this in detail in the next section. Firstly, we create Level 1 Analyst and Level 2 Analyst without rights at all.

After that we define the following statuses:

In the lower table you can manage the access rights for every role and every Case Status. We can give the suitable rights to our generated roles by clicking the `New Rights for Case Status` button on the right.

For Level 1 Analyst we add the right to read and write all "Open" cases and change the case status to this status (set).

Additionally, we grant Level 1 Analyst the rights to read, write and set all cases for "Level 1 Working".

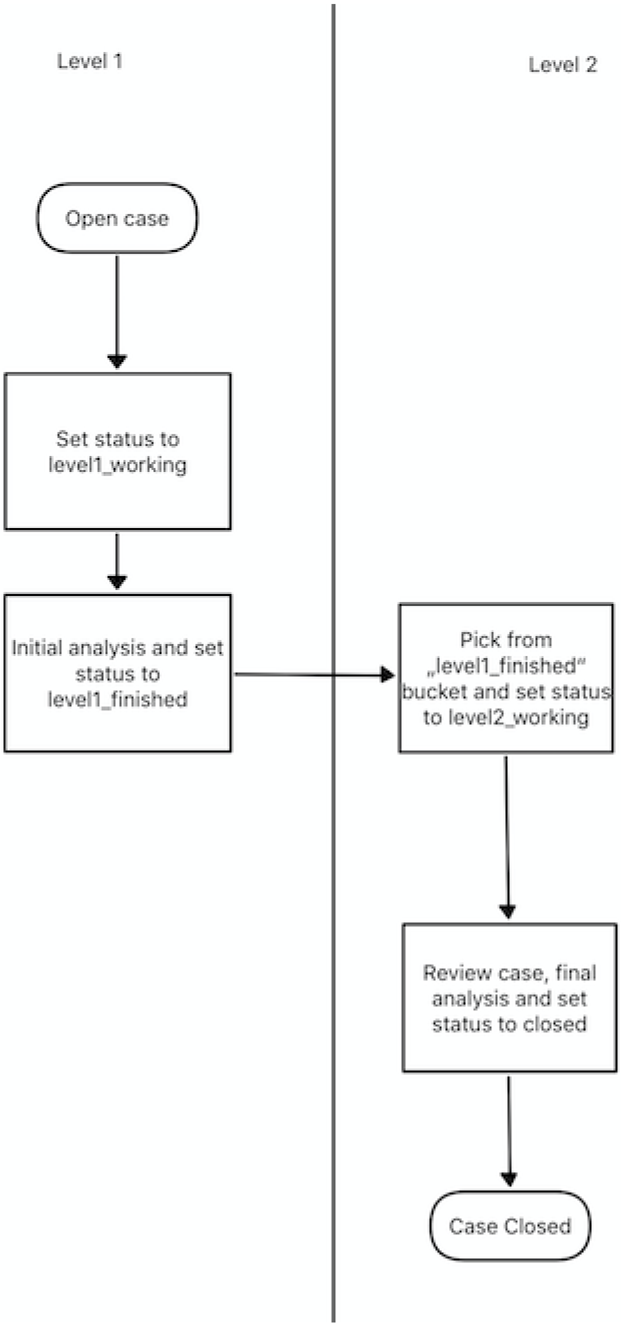
Level 1
Level 2
Open case
Set status to level1_working
Initial analysis and set status to level1_finished
Pick from „level1_finished“ bucket and set status to level2_working
Review case, final analysis and set status to closed
Case Closed

Fig. 13: Workflow open Cases

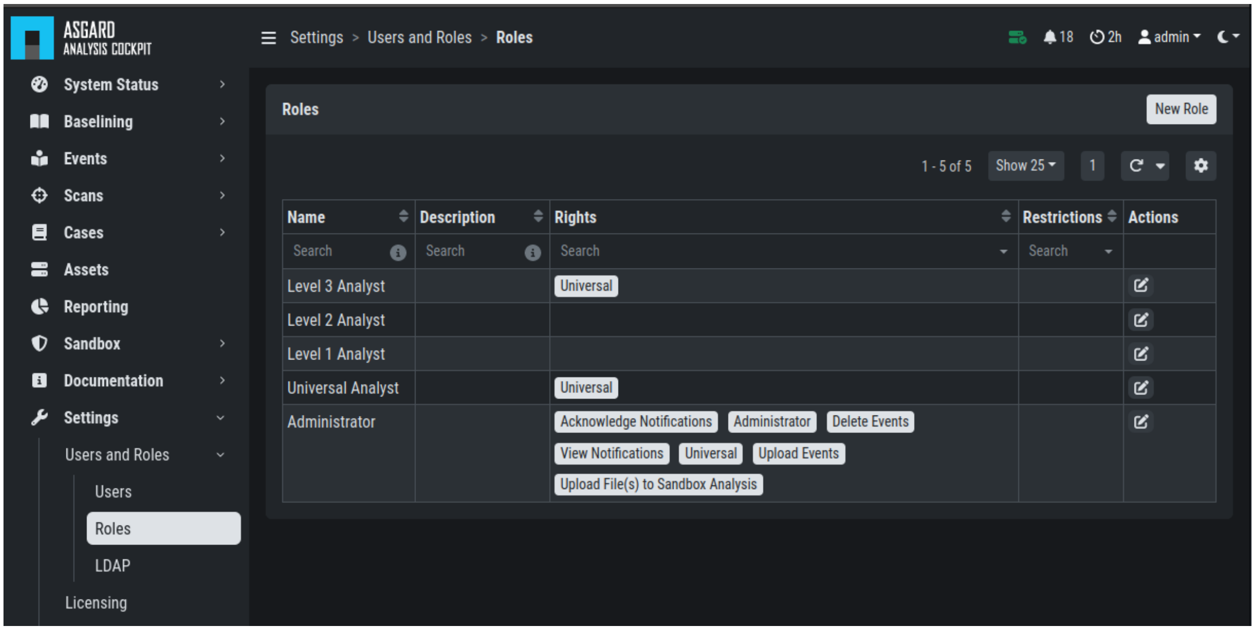

Name	Description	Rights	Restrictions	Actions
Level 3 Analyst		Universal		
Level 2 Analyst				
Level 1 Analyst				
Universal Analyst		Universal		
Administrator		Acknowledge Notifications Administrator Delete Events View Notifications Universal Upload Events Upload File(s) to Sandbox Analysis		

Fig. 14: Settings – adding additional roles

Finally, we grant the right to read and set cases for the status `Level1 Finished`. This allows Level 1 Analysts to set a particular case to "Level 1 Finished" and restricts them from modifying this case once they have passed it to this status.

For Level 2 we now add the rights to read and write cases for "Level 1 Finished" and the rights to read, write and set cases for "Level 2 Working". This allows Level 2 analysts to pick cases from the "Level 1 Finished" status and start working on them.

As we do not want Level 2 Analysts to reopen cases, that have already been closed we only grant them rights to read and set for the status "Closed".

Additionally, we give Level 2 Analyst the right to set the case status to "Open".

After that, the `Access rights for Case Status` section looks like this:

Of course, this is only an example. You may of course decide to give Level 2 full access to all cases, and it may also be a good means of training to grant Level 1 Analysts the right to see the "Level 2 Working" and "Closed" cases. You may also want Level 2 Analysts to reopen "Closed" cases or may restrict this right to an additional role. This just illustrates, that the system is highly configurable with an almost infinite number of statuses, roles and rights.

Finally, you simply add users and add them to their particular role.

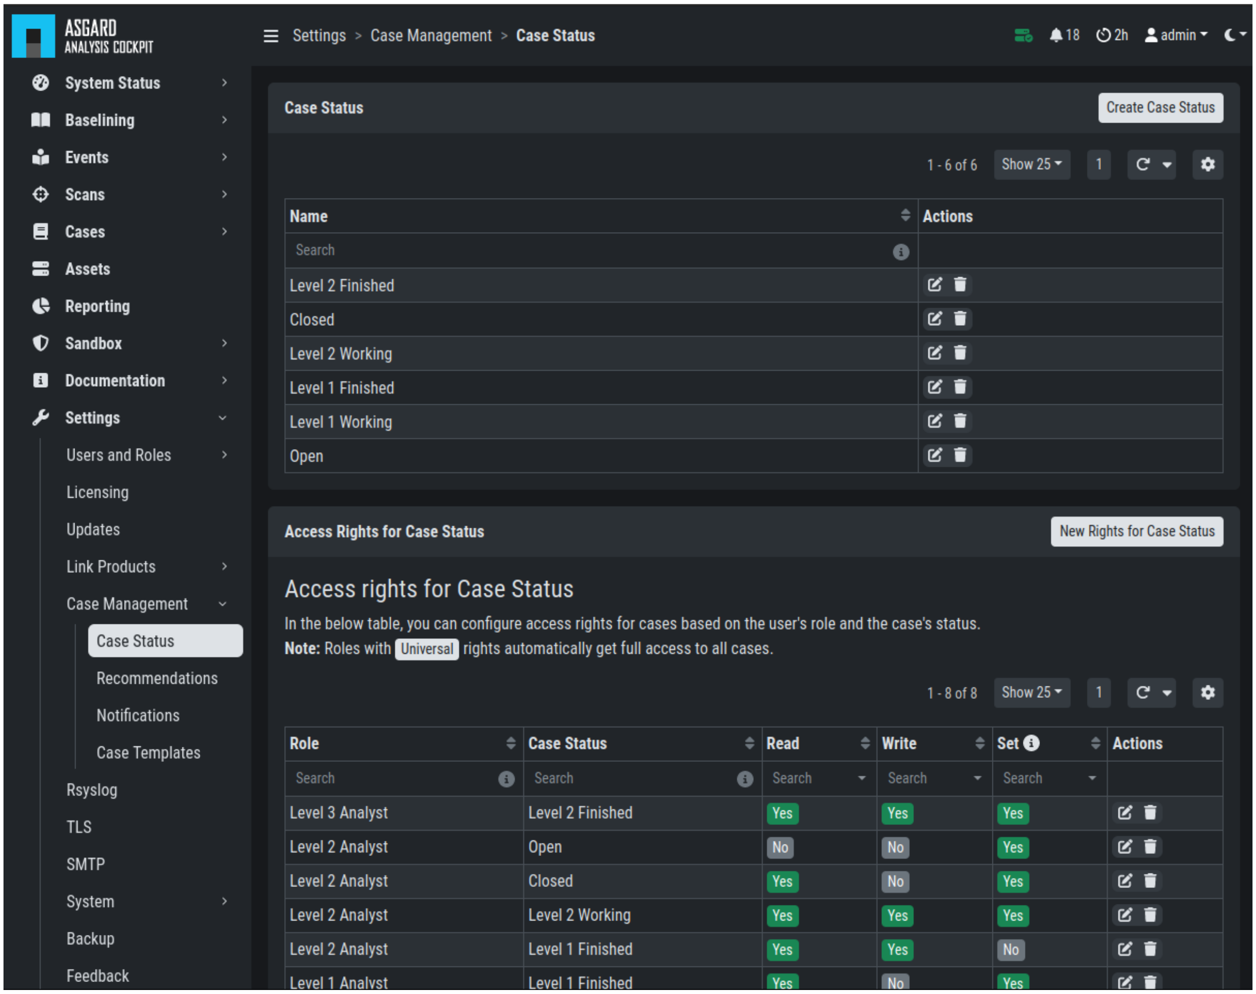

Case Status

Name	Actions
Level 2 Finished	
Closed	
Level 2 Working	
Level 1 Finished	
Level 1 Working	
Open	

Access Rights for Case Status

Access rights for Case Status

In the below table, you can configure access rights for cases based on the user's role and the case's status.

Note: Roles with Universal rights automatically get full access to all cases.

Role	Case Status	Read	Write	Set	Actions
Level 3 Analyst	Level 2 Finished	Yes	Yes	Yes	
Level 2 Analyst	Open	No	No	Yes	
Level 2 Analyst	Closed	Yes	No	Yes	
Level 2 Analyst	Level 2 Working	Yes	Yes	Yes	
Level 2 Analyst	Level 1 Finished	Yes	Yes	No	
Level 1 Analyst	Level 1 Finished	Yes	No	Yes	

Fig. 15: Settings – Case Status

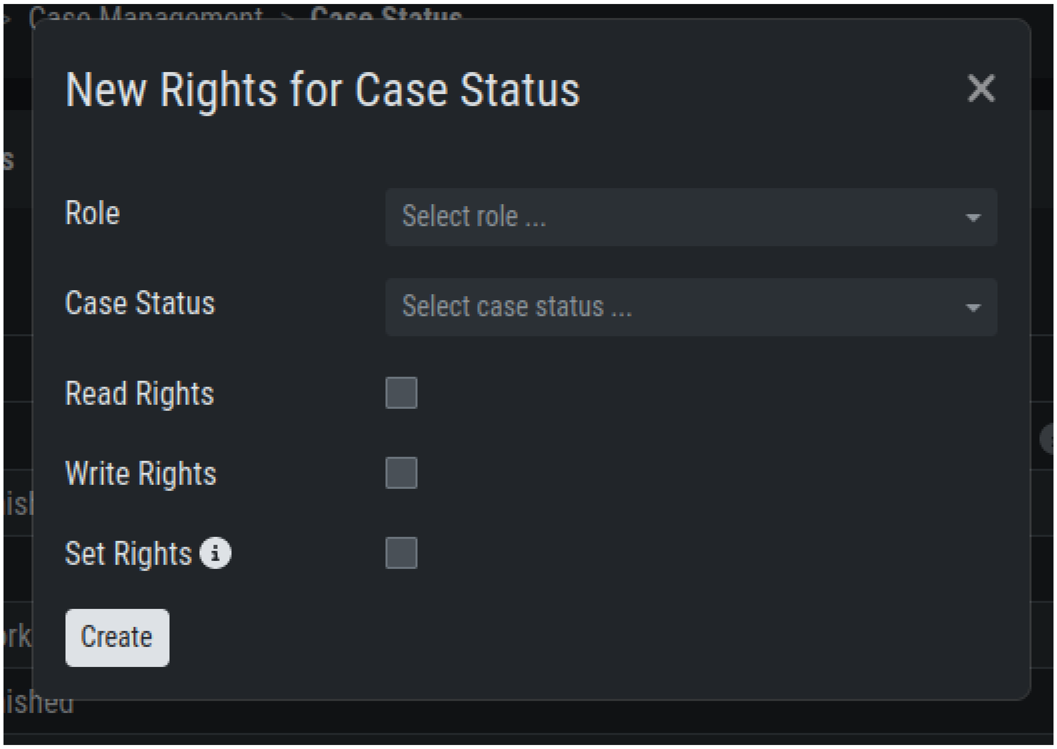

New Rights for Case Status

Role: Select role ...

Case Status: Select case status ...

Read Rights □

Write Rights □

Set Rights □

Fig. 16: Edit Rights – Read, Write, Set

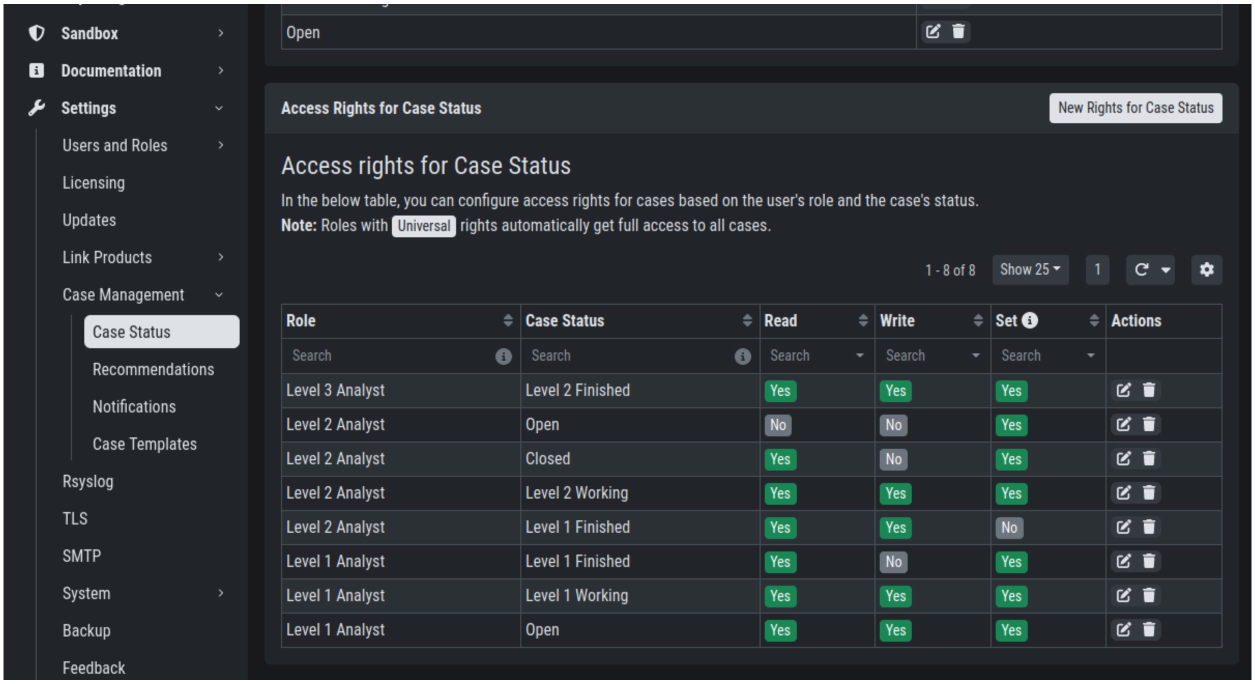

Access Rights for Case Status

Access rights for Case Status

In the below table, you can configure access rights for cases based on the user's role and the case's status.

Note: Roles with Universal rights automatically get full access to all cases.

Role	Case Status	Read	Write	Set	Actions
Search	Search	Search	Search	Search	
Level 3 Analyst	Level 2 Finished	Yes	Yes	Yes	
Level 2 Analyst	Open	No	No	Yes	
Level 2 Analyst	Closed	Yes	No	Yes	
Level 2 Analyst	Level 2 Working	Yes	Yes	Yes	
Level 2 Analyst	Level 1 Finished	Yes	Yes	No	
Level 1 Analyst	Level 1 Finished	Yes	No	Yes	
Level 1 Analyst	Level 1 Working	Yes	Yes	Yes	
Level 1 Analyst	Open	Yes	Yes	Yes	

Fig. 17: Settings – Access rights for Case Status

CHAPTER

FIVE

BASELINING

This section assumes, that you have read the chapter *Basic Concepts*.

All incoming logs, that do not match an existing case, will show up in the `Baselining` section. From here you can create cases and define your baseline, meaning every event showing in the baseline is in theory something unknown/new.

5.1 Customize Your View

By default, the Analysis Cockpit `Baselining` view ships with multiple bar charts and a table with the most relevant columns in order to help you find meaningful groups of logs. You can add additional bar charts by clicking on the `Advanced Tools` button and selecting `Chart Preferences`.

You can also modify which bar charts are shown by the name/field-name of the chart and choose the category you want to see. To get more details about a bar chart, you can click on square symbol in the heading of the bar chart.

Click the `Columns` button to manage which columns are shown.

Since the column preferences have an overvelming amount of fields you can choose from, we made looking for specific columns easier by integrating a search into the top right corner.

Hint: All views are personalized and changes will only affect your user.

5.2 Manual Case Creation

This section walks you through the manual case creation. This method gives you more flexibility in terms of conditions and details regarding the cases, but is more time consuming.

The results of the cases depending on specific settings you are setting during the case creation.

5.2.1 Case Creation Basics

Create a new case following these steps:

1. Select on which conditions the case should be built. `Search Result`

   will take your filters from the baselining view into consideration, and build a case with the condition of your search.

1. Inspect the Case Assignment and Conditions. Set Auto Assign if needed.

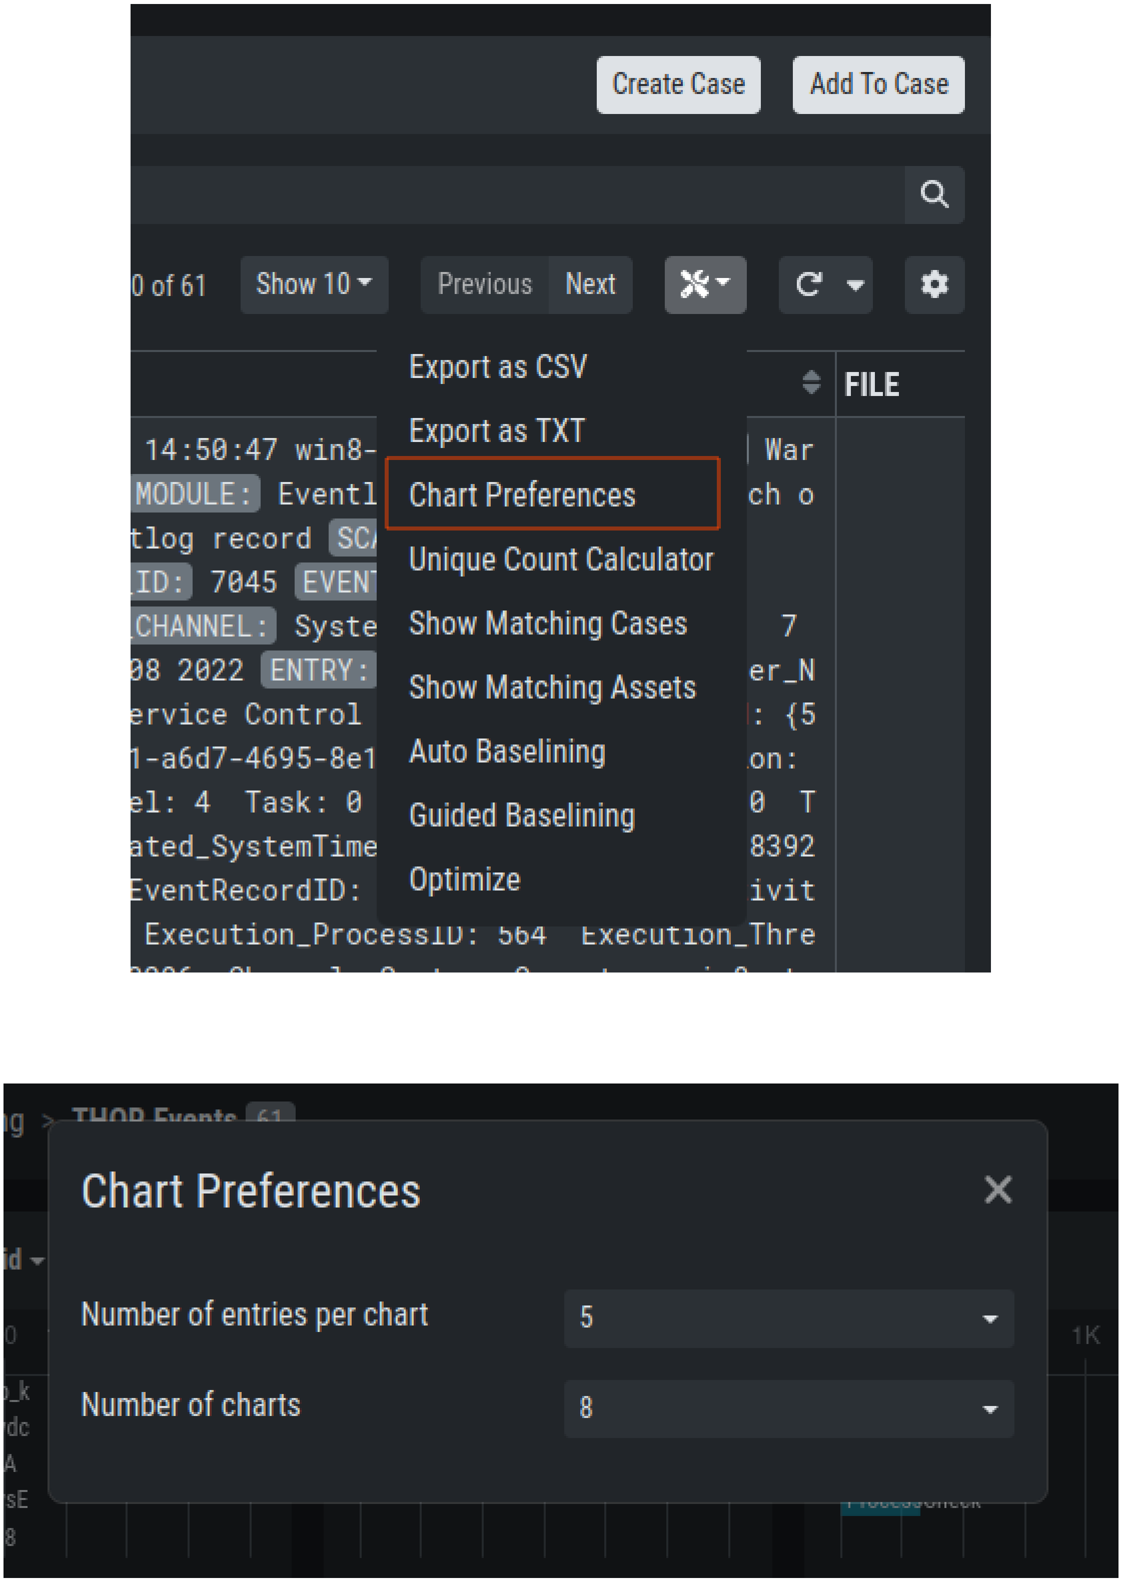

Chart Preferences

Number of entries per chart: 5

Number of charts: 8

Fig. 1: Chart Preferences

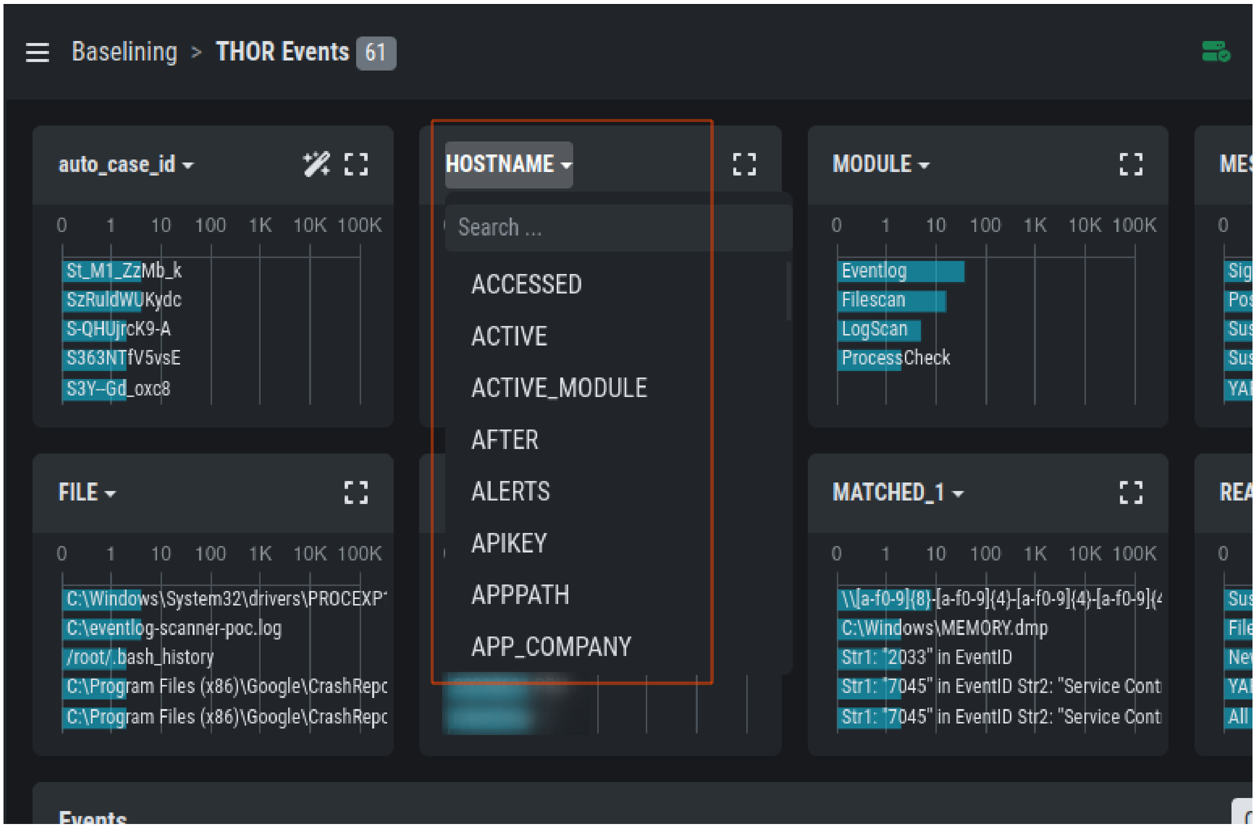

Fig. 2: Bar Chart Selector

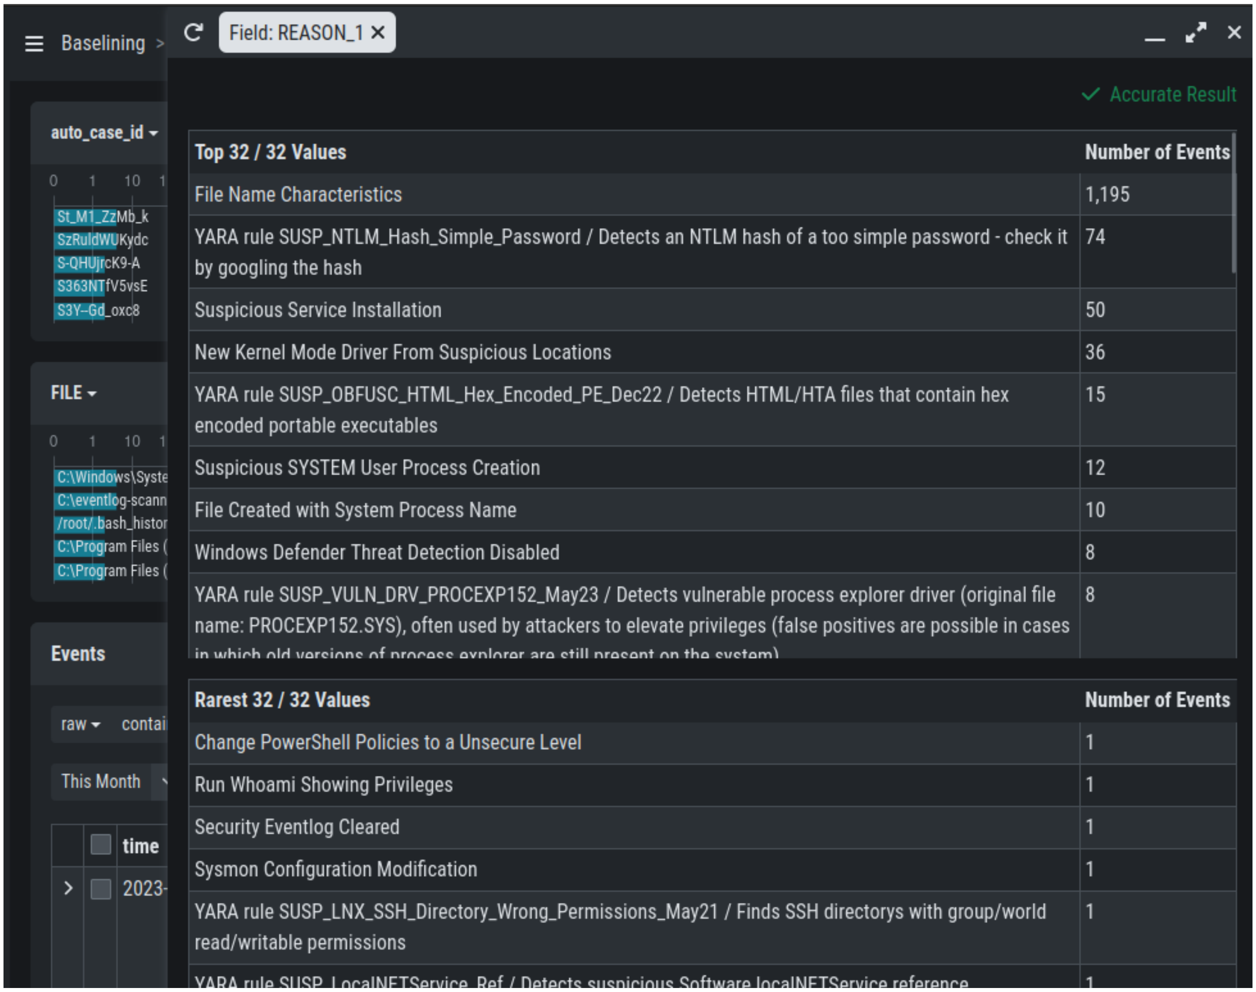

Top 32 / 32 Values	Number of Events
File Name Characteristics	1,195
YARA rule SUSP_NTLM_Hash_Simple_Password / Detects an NTLM hash of a too simple password - check it by googling the hash	74
Suspicious Service Installation	50
New Kernel Mode Driver From Suspicious Locations	36
YARA rule SUSP_OBFUSC_HTML_Hex_Encoded_PE_Dec22 / Detects HTML/HTA files that contain hex encoded portable executables	15
Suspicious SYSTEM User Process Creation	12
File Created with System Process Name	10
Windows Defender Threat Detection Disabled	8
YARA rule SUSP_VULN_DRV_PROCEXP152_May23 / Detects vulnerable process explorer driver (original file name: PROCEXP152.SYS), often used by attackers to elevate privileges (false positives are possible in cases in which old versions of process explorer are still present on the system)	8

Rarest 32 / 32 Values	Number of Events
Change PowerShell Policies to a Unsecure Level	1
Run Whoami Showing Privileges	1
Security Eventlog Cleared	1
Sysmon Configuration Modification	1
YARA rule SUSP_LNX_SSH_Directory_Wrong_Permissions_May21 / Finds SSH directorys with group/world read/writable permissions	1
YARA rule SUSP_LocalNETService_Ref / Detects suspicious Software localNETService reference	1

Fig. 3: Bar Chart Details

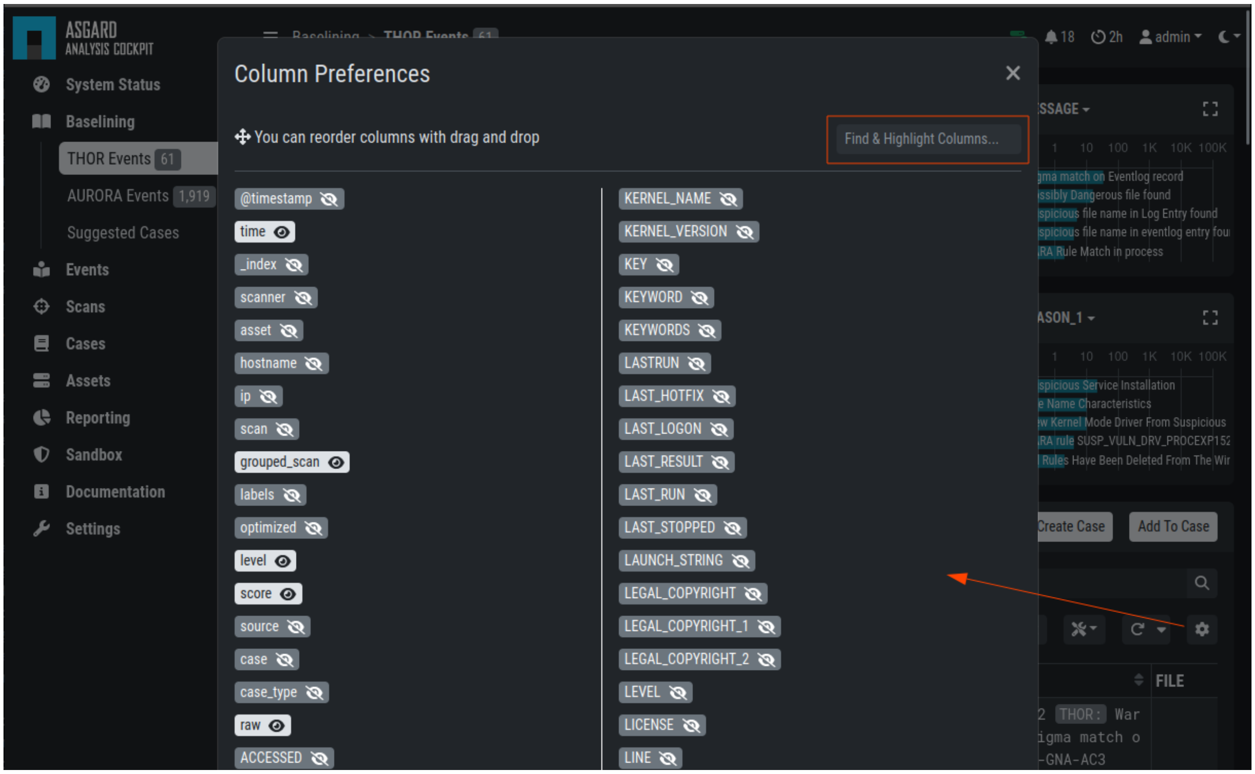

Column Preferences

You can reorder columns with drag and drop

Fig. 4: Column Preferences

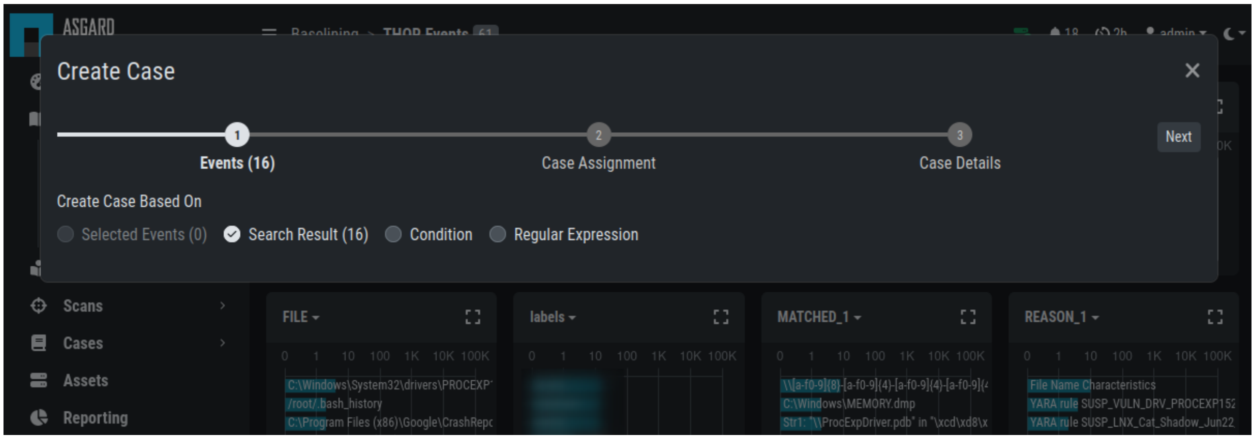

Create Case

1 Events (16) — 2 Case Assignment — 3 Case Details

Next

Create Case Based On

○ Selected Events (0) ● Search Result (16) ○ Condition ○ Regular Expression

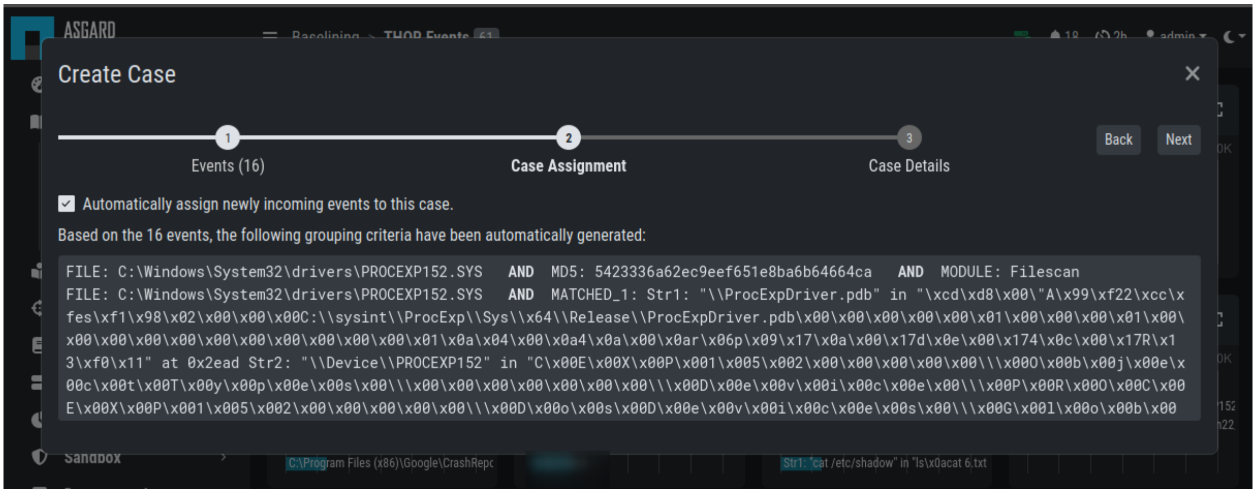

☑ Automatically assign newly incoming events to this case.

Based on the 16 events, the following grouping criteria have been automatically generated:

```
FILE: C:\Windows\System32\drivers\PROCEXP152.SYS   AND  MD5: 5423336a62ec9eef651e8ba6b64664ca   AND  MODULE: Filescan
FILE: C:\Windows\System32\drivers\PROCEXP152.SYS   AND  MATCHED_1: Str1: "\\ProcExpDriver.pdb" in "\xcd\xd8\x00\"A\x99\xf22\xcc\x
fes\xf1\x98\x02\x00\x00\x00C:\\sysint\\ProcExp\\Sys\\x64\\Release\\ProcExpDriver.pdb\x00\x00\x00\x00\x00\x01\x00\x00\x00\x01\x00\
x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x01\x0a\x04\x00\x0a4\x0a\x00\x0ar\x06p\x09\x17\x0a\x00\x17d\x0e\x00\x174\x0c\x00\x17R\x1
3\xf0\x11" at 0x2ead Str2: "\\Device\\PROCEXP152" in "C\x00E\x00X\x00P\x001\x005\x002\x00\x00\x00\x00\x00\\\x00O\x00b\x00j\x00e\x
00c\x00t\x00T\x00y\x00p\x00e\x00s\x00\\\x00\x00\x00\x00\x00\x00\x00\\\x00D\x00e\x00v\x00i\x00c\x00e\x00\\\x00P\x00R\x00O\x00C\x00
E\x00X\x00P\x001\x005\x002\x00\x00\x00\x00\x00\\\x00D\x00o\x00s\x00D\x00e\x00v\x00i\x00c\x00e\x00s\x00\\\x00G\x00l\x00o\x00b\x00
```

C:\Program Files (x86)\Google\CrashRepc
Str1: "cat /etc/shadow" in "ls\x0acat 6.txt

3. Set a case **name**, which serves as title - use keywords that make it easy for other analysts to find it based on a few terms (e.g. if a false positive was caused by matches in **savedsearch.conf**, use this filename in the title of your case)
4. Select a sample event for the **summary** field
5. Add your **assessment**
6. Choose one or more **recommendations**
7. Select a **case type** (see the *Case Types* for a detailed description of every case type)
8. Select a **case status** (usually used to mark it as 'work in progress' or to forward it to the next team)
9. Submit case by clicking the **Create Case** button

5.2.2 Select Log Messages for a Case

In order to create a meaningful case, you typically start with selecting logs or groups of logs that you want to be contained in the case. This can be done in various ways:

- by adding a custom filter in the search bar
- by clicking on one of the bars in the bar chart
- by clicking on the filter symbol in a field in a log line
- by using the Lucene Search Query

You can generate a filter condition using an expression in the search field, choosing a category, deciding whether the expression should be contained, equal etc. and clicking the search button. Clicking on one of the bars in the bar chart or on the filter symbol in a field in a log line will generate a filter condition, too.

Hint: Filters can be negated by clicking on the two arrows symbol or delete it by clicking on the cross symbol.

Using the built-in custom filters is the most common and easiest way to select groups of logs.

For those who prefer Lucene, an additional Lucene search bar can be activated and can even be combined with the built-in custom search.

In order to activate the Lucene Query search just click the `contains` button and choose `Lucene Query`.

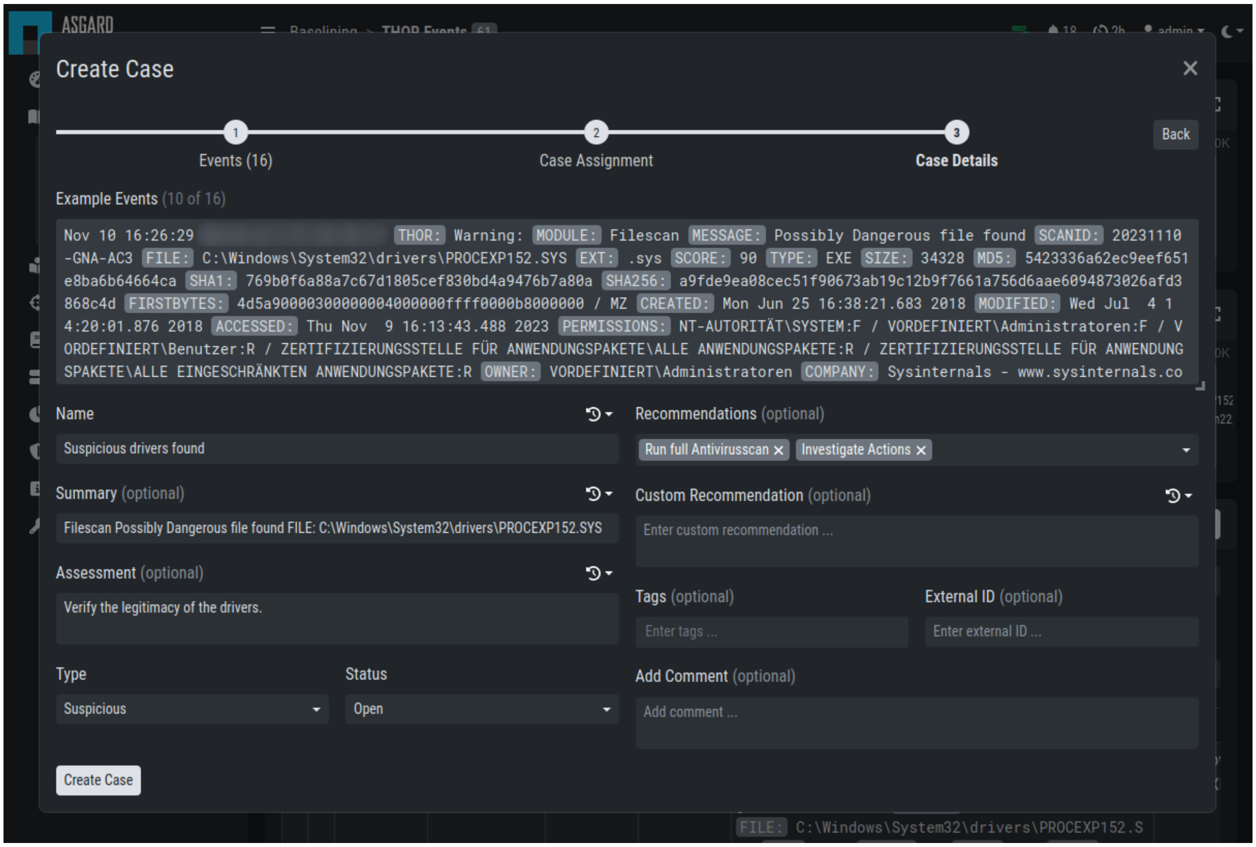

Example Events (10 of 16)

```
Nov 10 16:26:29 [redacted] THOR: Warning: MODULE: Filescan MESSAGE: Possibly Dangerous file found SCANID: 20231110
-GNA-AC3 FILE: C:\Windows\System32\drivers\PROCEXP152.SYS EXT: .sys SCORE: 90 TYPE: EXE SIZE: 34328 MD5: 5423336a62ec9eef651
e8ba6b64664ca SHA1: 769b0f6a88a7c67d1805cef830bd4a9476b7a80a SHA256: a9fde9ea08cec51f90673ab19c12b9f7661a756d6aae6094873026afd3
868c4d FIRSTBYTES: 4d5a90000300000004000000ffff0000b8000000 / MZ CREATED: Mon Jun 25 16:38:21.683 2018 MODIFIED: Wed Jul  4 1
4:20:01.876 2018 ACCESSED: Thu Nov  9 16:13:43.488 2023 PERMISSIONS: NT-AUTORITÄT\SYSTEM:F / VORDEFINIERT\Administratoren:F / V
ORDEFINIERT\Benutzer:R / ZERTIFIZIERUNGSSTELLE FÜR ANWENDUNGSPAKETE\ALLE ANWENDUNGSPAKETE:R / ZERTIFIZIERUNGSSTELLE FÜR ANWENDUNG
SPAKETE\ALLE EINGESCHRÄNKTEN ANWENDUNGSPAKETE:R OWNER: VORDEFINIERT\Administratoren COMPANY: Sysinternals - www.sysinternals.co
```

Name: Suspicious drivers found

Recommendations (optional): Run full Antivirusscan ✕ Investigate Actions ✕

Summary (optional): Filescan Possibly Dangerous file found FILE: C:\Windows\System32\drivers\PROCEXP152.SYS

Custom Recommendation (optional): Enter custom recommendation ...

Assessment (optional): Verify the legitimacy of the drivers.

Tags (optional): Enter tags ...

External ID (optional): Enter external ID ...

Type: Suspicious

Status: Open

Add Comment (optional): Add comment ...

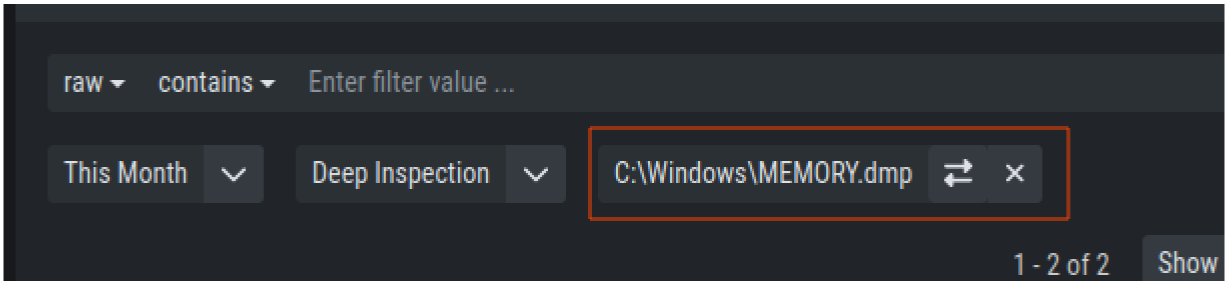

Fig. 5: Active Filters

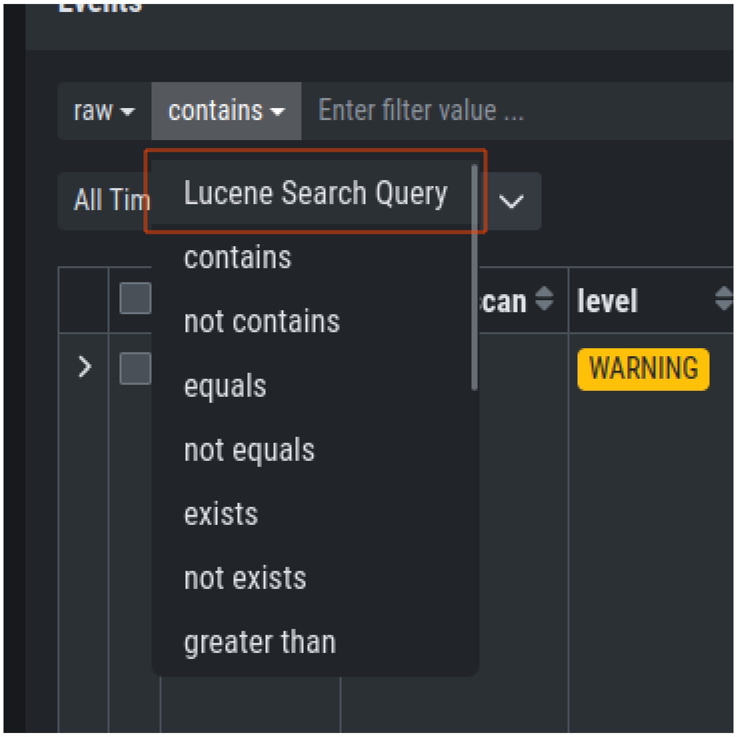

Fig. 6: Lucene Query

Note: You can Alt/Shift click items in the top field view to add them as a `NOT` filter to your search.

5.2.3 Case Creation from Search Results

This is the most relevant way to create a case. Create the filters, so that you only see the logs you want to be contained in your case. Then click the `Create Case` button, select `Search results` and add a name, that makes sense to you.

If you want future incoming logs with the same log lines automatically assigned to this case, you have to tick the checkbox `Automatically assign newly incoming events to this case.`.

You may add a summary to your case.

You may or may not add assessment, case type, recommendations or a comment. After closing you will find the log section empty, as it is still using your filter, but the matching log lines have been removed from this section and added to the case.

Simply remove the filter and the remaining log lines will show up.

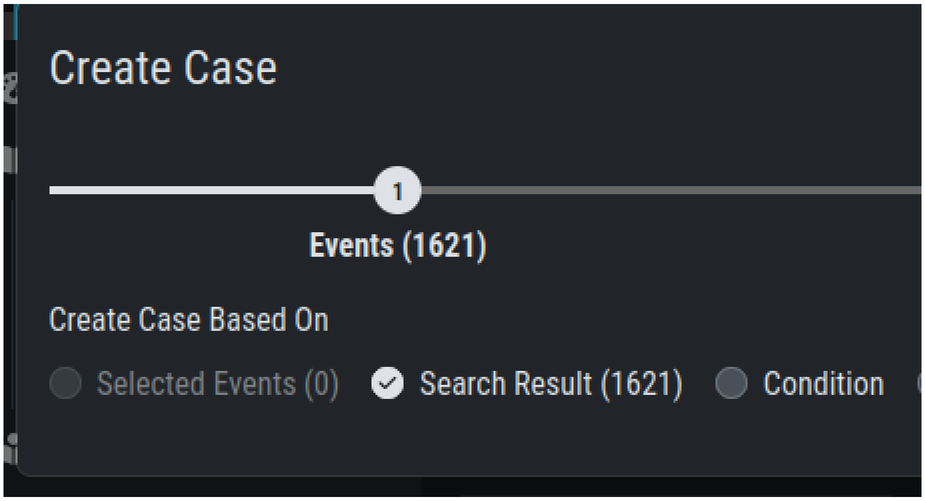

Create Case

1 Events (1621)

Create Case Based On

○ Selected Events (0) ☑ Search Result (1621) ○ Condition

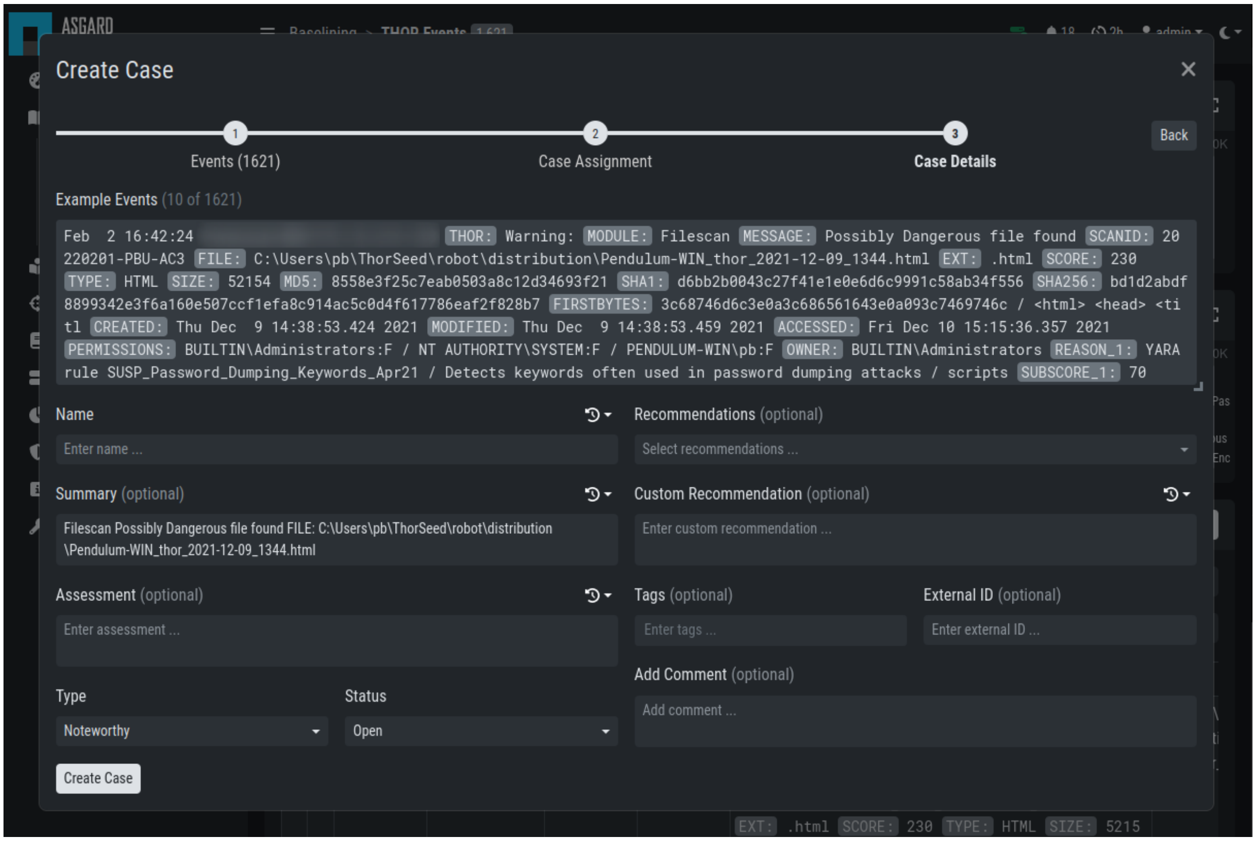

Create Case

Example Events (10 of 1621)

```
Feb  2 16:42:24 [redacted] THOR: Warning: MODULE: Filescan MESSAGE: Possibly Dangerous file found SCANID: 20
220201-PBU-AC3 FILE: C:\Users\pb\ThorSeed\robot\distribution\Pendulum-WIN_thor_2021-12-09_1344.html EXT: .html SCORE: 230
TYPE: HTML SIZE: 52154 MD5: 8558e3f25c7eab0503a8c12d34693f21 SHA1: d6bb2b0043c27f41e1e0e6d6c9991c58ab34f556 SHA256: bd1d2abdf
8899342e3f6a160e507ccf1efa8c914ac5c0d4f617786eaf2f828b7 FIRSTBYTES: 3c68746d6c3e0a3c686561643e0a093c7469746c / <html> <head> <ti
tl CREATED: Thu Dec  9 14:38:53.424 2021 MODIFIED: Thu Dec  9 14:38:53.459 2021 ACCESSED: Fri Dec 10 15:15:36.357 2021
PERMISSIONS: BUILTIN\Administrators:F / NT AUTHORITY\SYSTEM:F / PENDULUM-WIN\pb:F OWNER: BUILTIN\Administrators REASON_1: YARA
rule SUSP_Password_Dumping_Keywords_Apr21 / Detects keywords often used in password dumping attacks / scripts SUBSCORE_1: 70
```

Name: Enter name ...

Recommendations (optional): Select recommendations ...

Summary (optional): Filescan Possibly Dangerous file found FILE: C:\Users\pb\ThorSeed\robot\distribution \Pendulum-WIN_thor_2021-12-09_1344.html

Custom Recommendation (optional): Enter custom recommendation ...

Assessment (optional): Enter assessment ...

Tags (optional): Enter tags ...

External ID (optional): Enter external ID ...

Add Comment (optional): Add comment ...

Type: Noteworthy

Status: Open

Fig. 7: Baselining – Create Case

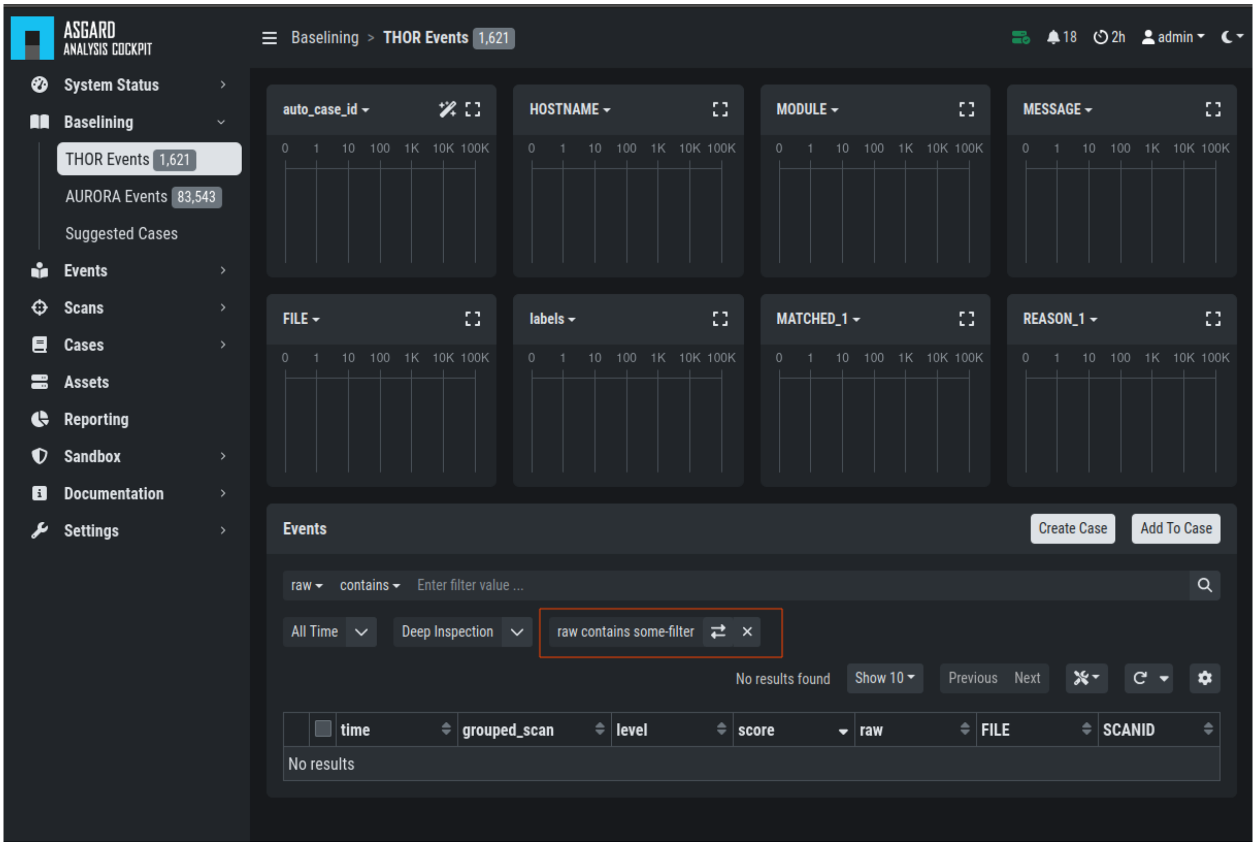

raw contains some-filter

No results found

time	grouped_scan	level	score	raw	FILE	SCANID
No results						

Fig. 8: Log Section empty

5.2.4 Case Creation from Selection

In order to create a case from a specific selection of logs simply use the checkboxes at the very left side of the table and click the `Create Case` button select `Selected events` and add a name, that makes sense to you.

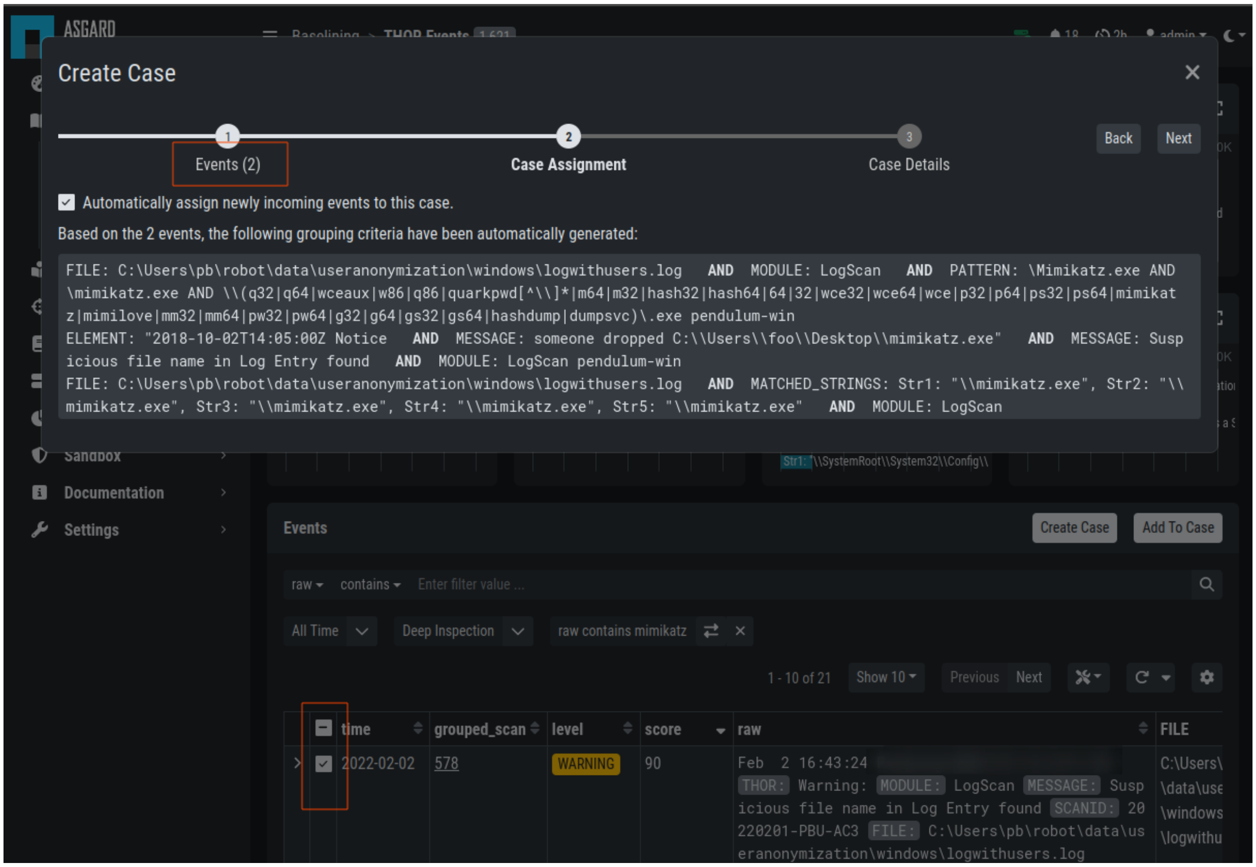

☑ Automatically assign newly incoming events to this case.

Based on the 2 events, the following grouping criteria have been automatically generated:

```
FILE: C:\Users\pb\robot\data\useranonymization\windows\logwithusers.log   AND  MODULE: LogScan   AND  PATTERN: \Mimikatz.exe AND
\mimikatz.exe AND \\(q32|q64|wceaux|w86|q86|quarkpwd[^\\]*|m64|m32|hash32|hash64|64|32|wce32|wce64|wce|p32|p64|ps32|ps64|mimikat
z|mimilove|mm32|mm64|pw32|pw64|g32|g64|gs32|gs64|hashdump|dumpsvc)\.exe pendulum-win
ELEMENT: "2018-10-02T14:05:00Z Notice   AND  MESSAGE: someone dropped C:\\Users\\foo\\Desktop\\mimikatz.exe"   AND  MESSAGE: Susp
icious file name in Log Entry found   AND  MODULE: LogScan pendulum-win
FILE: C:\Users\pb\robot\data\useranonymization\windows\logwithusers.log   AND  MATCHED_STRINGS: Str1: "\\mimikatz.exe", Str2: "\\
mimikatz.exe", Str3: "\\mimikatz.exe", Str4: "\\mimikatz.exe", Str5: "\\mimikatz.exe"   AND  MODULE: LogScan
```

Str1: \\SystemRoot\\System32\\Config\\

☐	time	grouped_scan	level	score	raw	FILE
☑	2022-02-02	578	WARNING	90	Feb 2 16:43:24 [redacted] THOR: Warning: MODULE: LogScan MESSAGE: Suspicious file name in Log Entry found SCANID: 20220201-PBU-AC3 FILE: C:\Users\pb\robot\data\useranonymization\windows\logwithusers.log	C:\Users\ \data\use \windows \logwithu

Fig. 9: Creating Cases from Selection

After closing you will find the selected logs have been removed from the logs section.

5.2.5 Case Creation Using a Custom Condition

To create a case with a condition, click the `Create Case` button and select `Condition`. Now you can build a condition by entering keywords in the field.

Keywords in the same field are combined by **OR**, you can negate them by clicking the `NOT` button or combine them with **AND** by clicking the `Add AND Condition` button. The filter bubbles you have generated before will be used as default. You are free to use, modify or delete them. Conditions only match on the `raw` field.

The `Test Condition / Regular Expression` button will calculate the numbers of hits and return some matching and some non-matching events as an example.

Again, you may or may not add auto assignment for future incoming log lines, summary, assessment, case type, recommendations or a comment. After closing you will find the selected logs have been removed from the logs section.

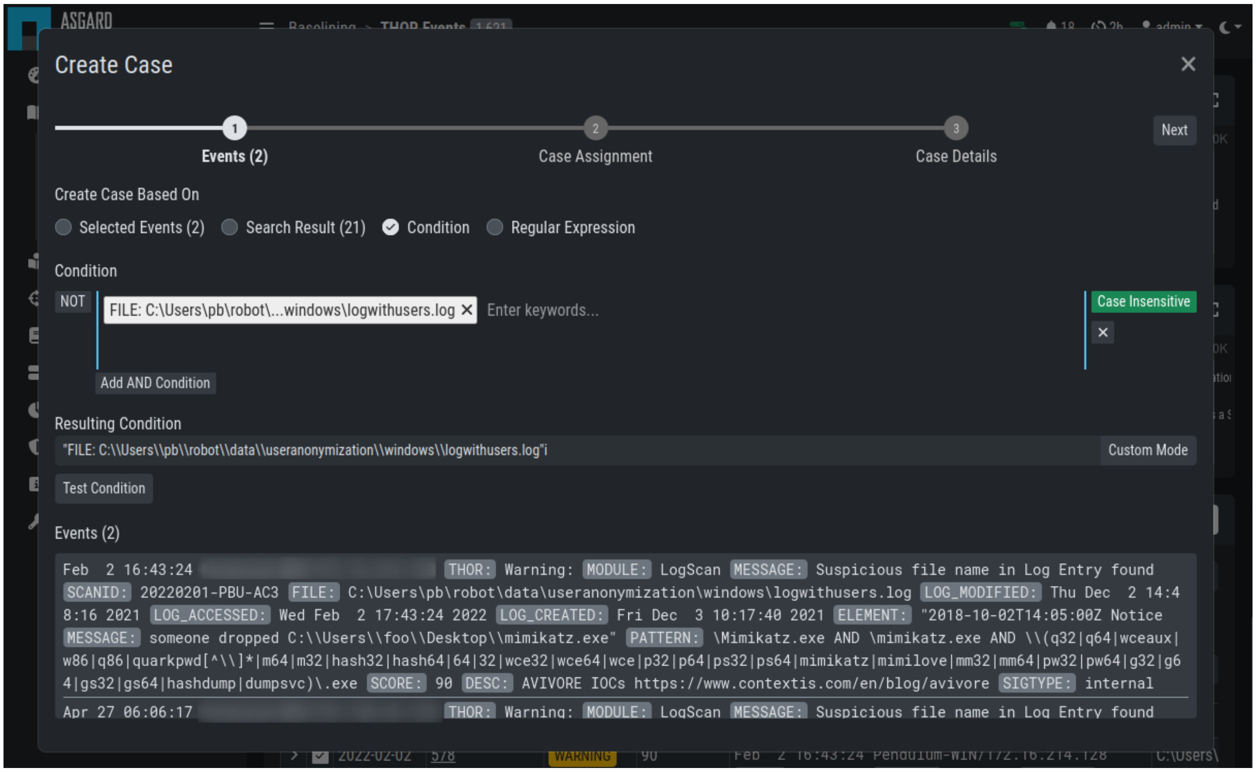

Create Case Based On

○ Selected Events (2) ○ Search Result (21) ● Condition ○ Regular Expression

Condition

NOT | FILE: C:\Users\pb\robot\...windows\logwithusers.log ✕ | Enter keywords...

Resulting Condition

"FILE: C:\\Users\\pb\\robot\\data\\useranonymization\\windows\\logwithusers.log"i

Events (2)

```
Feb  2 16:43:24 [redacted] THOR: Warning: MODULE: LogScan MESSAGE: Suspicious file name in Log Entry found
SCANID: 20220201-PBU-AC3 FILE: C:\Users\pb\robot\data\useranonymization\windows\logwithusers.log LOG_MODIFIED: Thu Dec  2 14:4
8:16 2021 LOG_ACCESSED: Wed Feb  2 17:43:24 2022 LOG_CREATED: Fri Dec  3 10:17:40 2021 ELEMENT: "2018-10-02T14:05:00Z Notice
MESSAGE: someone dropped C:\\Users\\foo\\Desktop\\mimikatz.exe" PATTERN: \Mimikatz.exe AND \mimikatz.exe AND \\(q32|q64|wceaux|
w86|q86|quarkpwd[^\\]*|m64|m32|hash32|hash64|64|32|wce32|wce64|wce|p32|p64|ps32|ps64|mimikatz|mimilove|mm32|mm64|pw32|pw64|g32|g6
4|gs32|gs64|hashdump|dumpsvc)\.exe SCORE: 90 DESC: AVIVORE IOCs https://www.contextis.com/en/blog/avivore SIGTYPE: internal
Apr 27 06:06:17 [redacted] THOR: Warning: MODULE: LogScan MESSAGE: Suspicious file name in Log Entry found
```

Fig. 10: Creating Cases through Condition

5.2.6 Case Creation Using a Regular Expressions

In order to create a case from a regular expression just click the `Create Case` button and select `Regular Expression`. This lets you write and test your regular expression.

The `Test Condition / Regular Expression` button will calculate the numbers of hits and return some matching and some non-matching events as an example.

Again, you may or may not add auto assignment for future incoming log lines, summary, assessment, case type, recommendations or a comment. After clicking the `Create Case` button, the matching lines will get removed from the log management view.

> **Warning:** It is recommended to use regular expressions only rarely and with caution. This feature can severely impact the performance of the system.

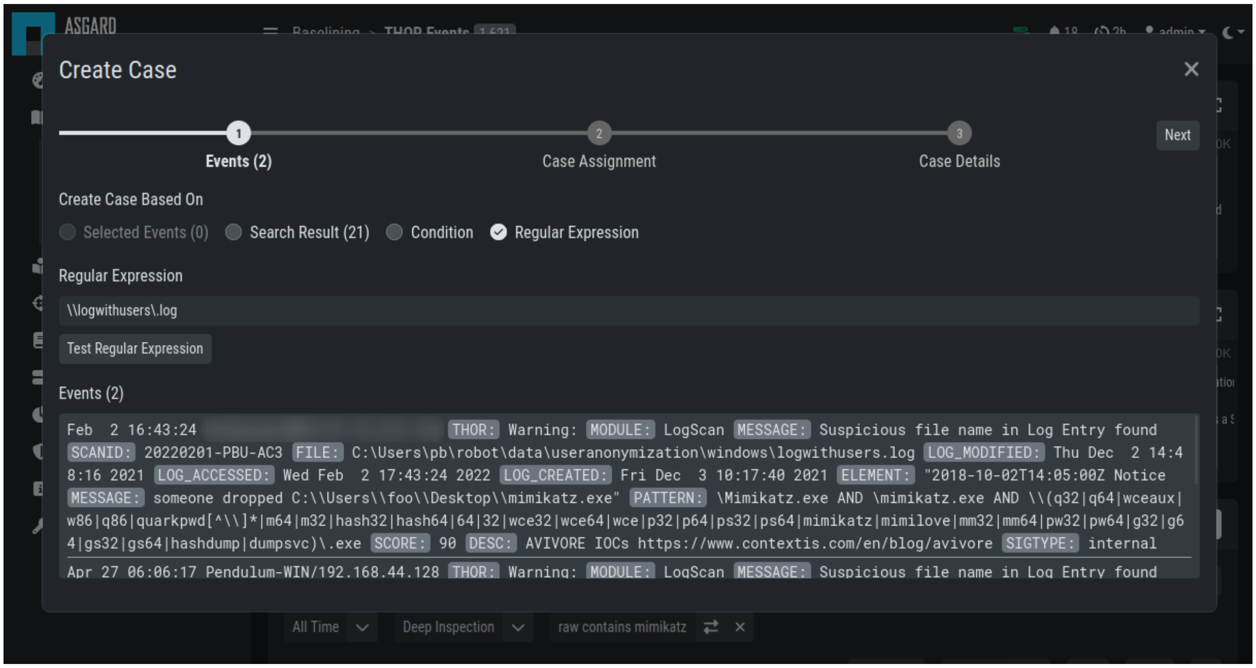

Create Case Based On

○ Selected Events (0) ○ Search Result (21) ○ Condition ● Regular Expression

Regular Expression: \\logwithusers\.log

Test Regular Expression

Events (2)

```
Feb  2 16:43:24 [redacted] THOR: Warning: MODULE: LogScan MESSAGE: Suspicious file name in Log Entry found
SCANID: 20220201-PBU-AC3 FILE: C:\Users\pb\robot\data\useranonymization\windows\logwithusers.log LOG_MODIFIED: Thu Dec  2 14:4
8:16 2021 LOG_ACCESSED: Wed Feb  2 17:43:24 2022 LOG_CREATED: Fri Dec  3 10:17:40 2021 ELEMENT: "2018-10-02T14:05:00Z Notice
MESSAGE: someone dropped C:\\Users\\foo\\Desktop\\mimikatz.exe" PATTERN: \Mimikatz.exe AND \mimikatz.exe AND \\(q32|q64|wceaux|
w86|q86|quarkpwd[^\\]*|m64|m32|hash32|hash64|64|32|wce32|wce64|wce|p32|p64|ps32|ps64|mimikatz|mimilove|mm32|mm64|pw32|pw64|g32|g6
4|gs32|gs64|hashdump|dumpsvc)\.exe SCORE: 90 DESC: AVIVORE IOCs https://www.contextis.com/en/blog/avivore SIGTYPE: internal
Apr 27 06:06:17 Pendulum-WIN/192.168.44.128 THOR: Warning: MODULE: LogScan MESSAGE: Suspicious file name in Log Entry found
```

Fig. 11: Creating Cases through Regular Expressions

5.3 Automated Case Creation

With Auto Baselining, the Cockpit will automatically generate cases for groups of logs that are similar, or in other words: Have the same auto_case_id.

After clicking the button `Automatically generate Cases` button in the `Auto Baselining` tab you will be prompted for a threshold. This means: Do only create a case when you find at least that many similar logs. In our example below the Cockpit will now generate cases for all groups of at least 2000 similar events.

After pressing the `Start` button, the Cockpit will start calculating and create cases. Depending on the data volume this may take a while and you will be presented a page that shows that Auto Cases is still running along with the current number of cases.

It is safe to leave this page, once the status in `Running`. It will continue in the background.

Important: The Analysis Cockpit generates auto_case_ids only for Alerts and Warnings. Don't use the Autocase feature for Notice and Info level events.

5.4 Add to Case

Sometimes you may want to add log lines to an already existing case because they represent the same security context. To do this you can just click the `Add to Case` button and select the suitable case. It is also possible to add an additional comment to this case for the addition.

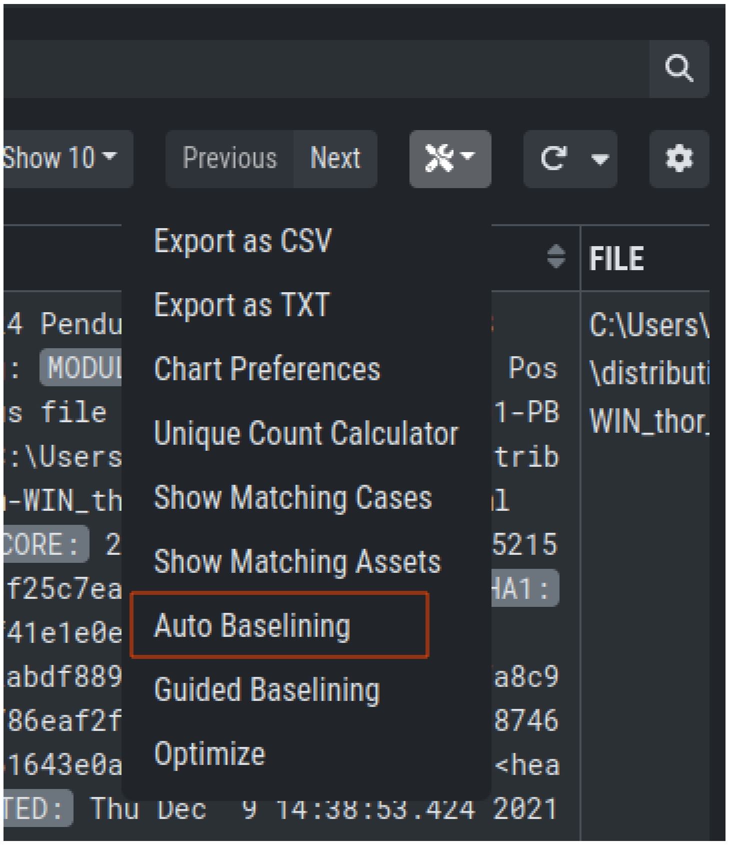

Fig. 12: Automatically create cases

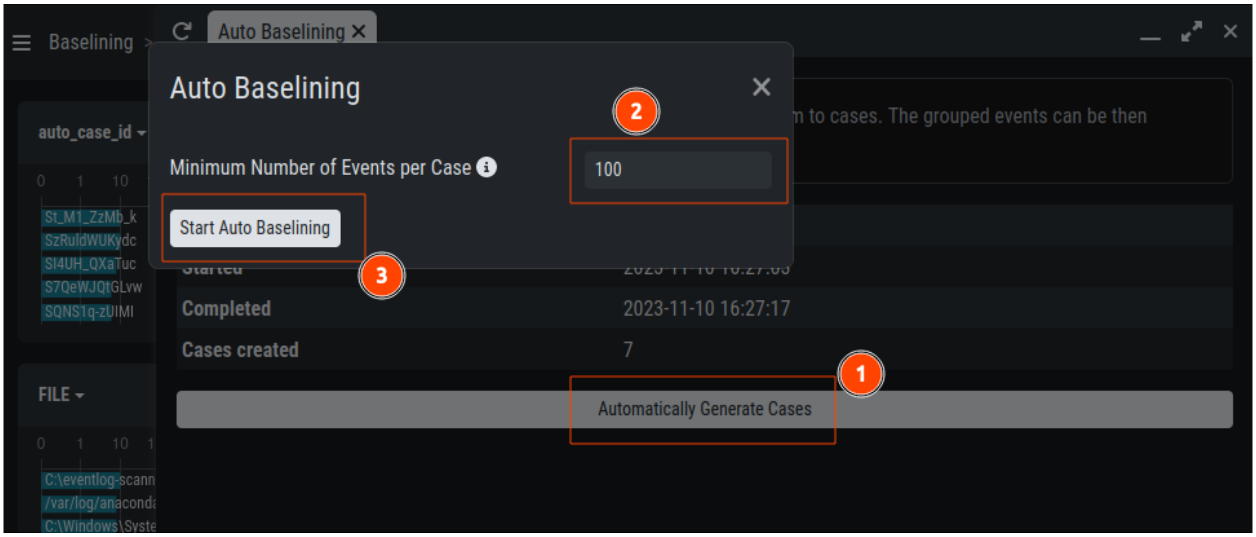

Minimum Number of Events per Case: 100

Completed 2023-11-10 16:27:17

Cases created 7

Fig. 13: Automatically create cases

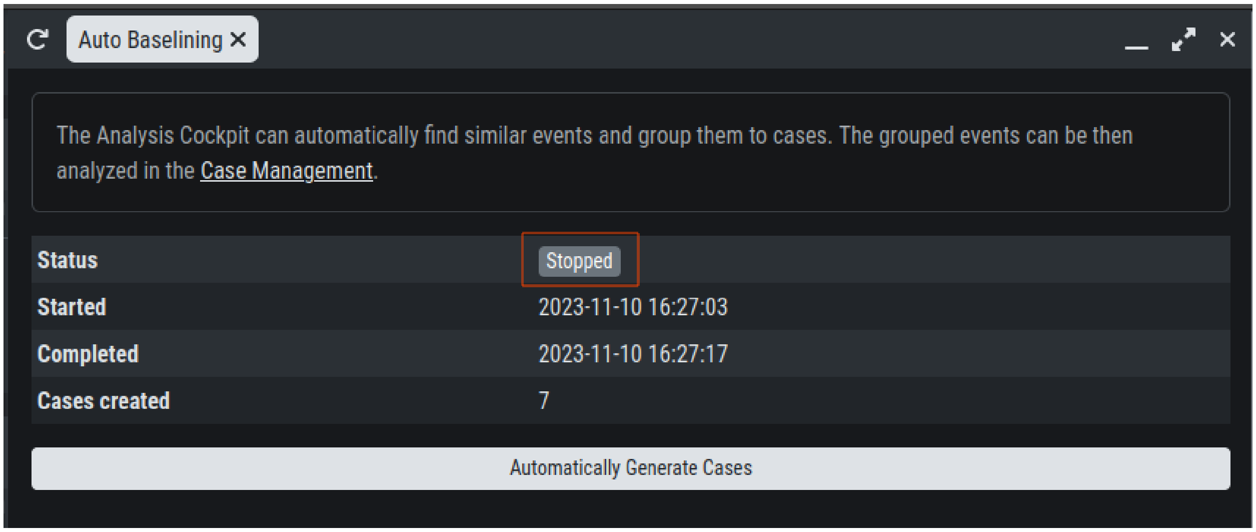

The Analysis Cockpit can automatically find similar events and group them to cases. The grouped events can be then analyzed in the Case Management.

Status	Stopped
Started	2023-11-10 16:27:03
Completed	2023-11-10 16:27:17
Cases created	7

Fig. 14: Auto Cases Status

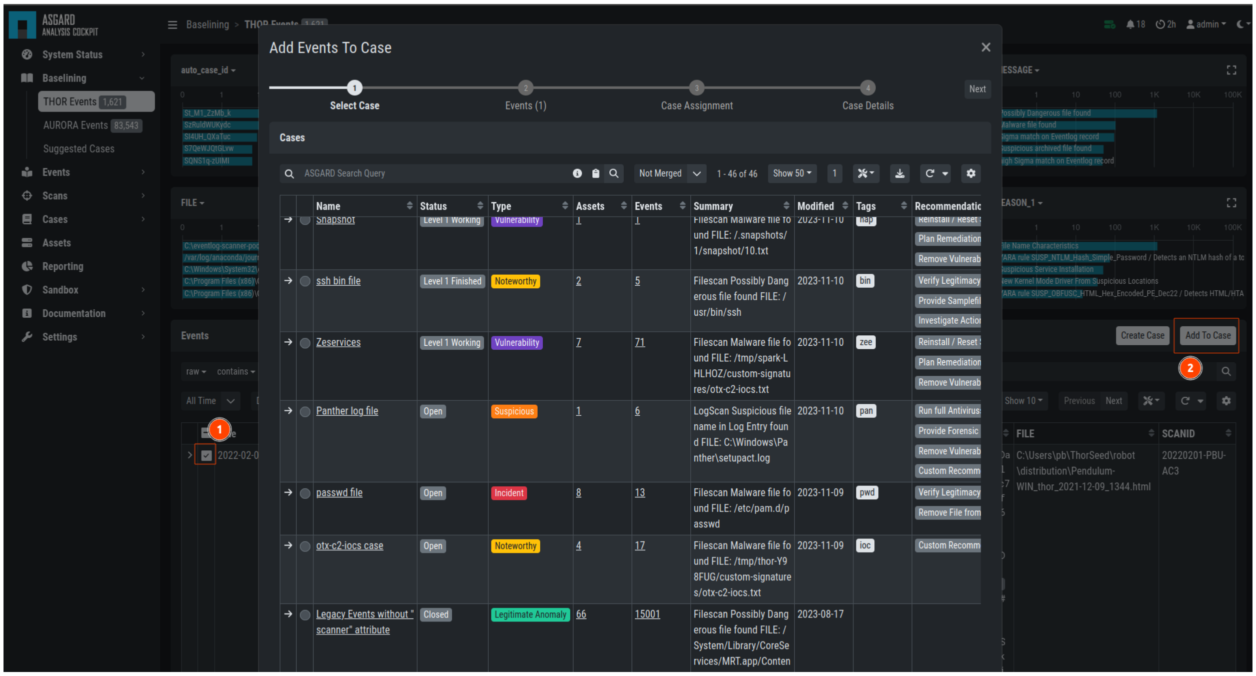

Add Events To Case

1 Select Case — 2 Events (1) — 3 Case Assignment — 4 Case Details

Cases

ASGARD Search Query | Not Merged | 1 - 46 of 46 | Show 50 | 1

Name	Status	Type	Assets	Events	Summary	Modified	Tags	Recommendatic
Snapshot	Level 1 Working	Vulnerability	1	1	Filescan Malware file found FILE: /.snapshots/1/snapshot/10.txt	2023-11-10	nap	Reinstall / Reset Plan Remediation Remove Vulnerab
ssh bin file	Level 1 Finished	Noteworthy	2	5	Filescan Possibly Dangerous file found FILE: /usr/bin/ssh	2023-11-10	bin	Verify Legitimacy Provide Samplefil Investigate Actior
Zeservices	Level 1 Working	Vulnerability	7	71	Filescan Malware file found FILE: /tmp/spark-LHLHOZ/custom-signatures/otx-c2-iocs.txt	2023-11-10	zee	Reinstall / Reset Plan Remediation Remove Vulnerab
Panther log file	Open	Suspicious	1	6	LogScan Suspicious file name in Log Entry found FILE: C:\Windows\Panther\setupact.log	2023-11-10	pan	Run full Antivirus Provide Forensic Remove Vulnerab Custom Recomm
passwd file	Open	Incident	8	13	Filescan Malware file found FILE: /etc/pam.d/passwd	2023-11-09	pwd	Verify Legitimacy Remove File from
otx-c2-iocs case	Open	Noteworthy	4	17	Filescan Malware file found FILE: /tmp/thor-Y98FUG/custom-signatures/otx-c2-iocs.txt	2023-11-09	ioc	Custom Recomm
Legacy Events without " scanner" attribute	Closed	Legitimate Anomaly	66	15001	Filescan Possibly Dangerous file found FILE: /System/Library/CoreServices/MRT.app/Conten	2023-08-17		

Fig. 15: Add to Case

5.5 Customizing the Detailed View of Log Lines

The detailed view for log lines opens by clicking on a log line. Within this view you can select some fields as favorite fields by clicking on the star symbol. They will always be shown at the top of this view. `MESSAGE`, `MODULE` and `hostname` are selected by default.

To search for all log lines with the same entry as this log line in a particular field, you can click the dropdown on the left hand side of the field.

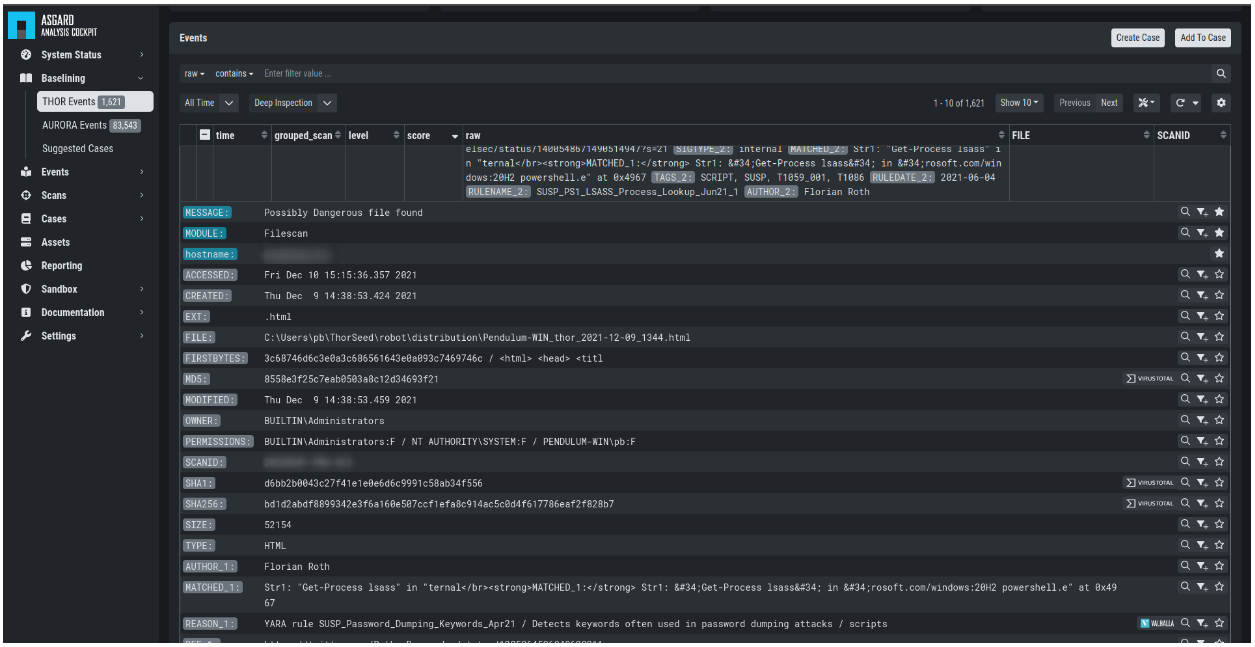

```
elsec/status/14005486/149051494/?s=21 SIGTYPE_2: internal MATCHED_2: Str1: "Get-Process lsass" i
n "ternal</br><strong>MATCHED_1:</strong> Str1: &#34;Get-Process lsass&#34; in &#34;rosoft.com/win
dows:20H2 powershell.e" at 0x4967 TAGS_2: SCRIPT, SUSP, T1059_001, T1086 RULEDATE_2: 2021-06-04
RULENAME_2: SUSP_PS1_LSASS_Process_Lookup_Jun21_1 AUTHOR_2: Florian Roth
MESSAGE:      Possibly Dangerous file found
MODULE:       Filescan
hostname:     [redacted]
ACCESSED:     Fri Dec 10 15:15:36.357 2021
CREATED:      Thu Dec  9 14:38:53.424 2021
EXT:          .html
FILE:         C:\Users\pb\ThorSeed\robot\distribution\Pendulum-WIN_thor_2021-12-09_1344.html
FIRSTBYTES:   3c68746d6c3e0a3c686561643e0a093c7469746c / <html> <head> <titl
MD5:          8558e3f25c7eab0503a8c12d34693f21
MODIFIED:     Thu Dec  9 14:38:53.459 2021
OWNER:        BUILTIN\Administrators
PERMISSIONS:  BUILTIN\Administrators:F / NT AUTHORITY\SYSTEM:F / PENDULUM-WIN\pb:F
SCANID:       [redacted]
SHA1:         d6bb2b0043c27f41e1e0e6d6c9991c58ab34f556
SHA256:       bd1d2abdf8899342e3f6a160e507ccf1efa8c914ac5c0d4f617786eaf2f828b7
SIZE:         52154
TYPE:         HTML
AUTHOR_1:     Florian Roth
MATCHED_1:    Str1: "Get-Process lsass" in "ternal</br><strong>MATCHED_1:</strong> Str1: &#34;Get-Process lsass&#34; in &#34;rosoft.com/windows:20H2 powershell.e" at 0x49
              67
REASON_1:     YARA rule SUSP_Password_Dumping_Keywords_Apr21 / Detects keywords often used in password dumping attacks / scripts
```

Fig. 16: Customizing the detailed view for log lines

Additionally, you can find a `VIRUSTOTAL` button in every hash field and a `VALHALLA` button in every reason field. By clicking `VIRUSTOTAL` the hash will be searched on Virustotal. By clicking `VALHALLA` you will get more information about the matching rule from valhalla.nextron-systems.com.

5.6 Usage of the Context Menu

You can use the context menu on any **value** in your logs to get an action menu. Within this menu, you can do different actions:

You can filter, search for similar events, or even create cases based on the value you right-clicked.

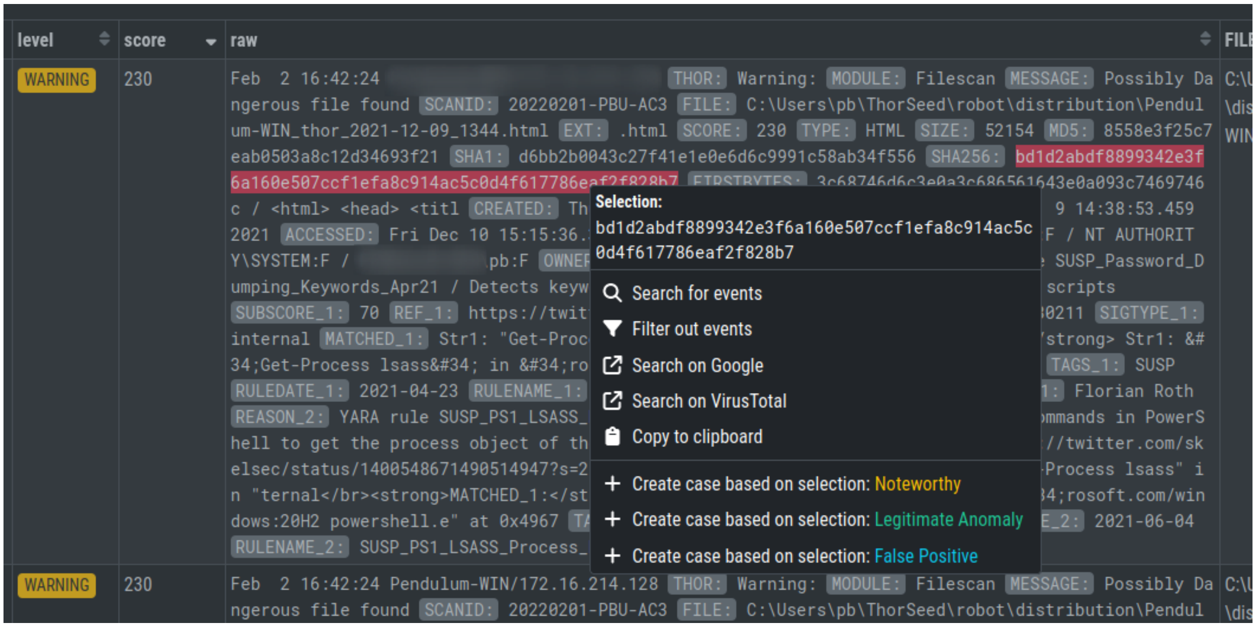

```
level     score   raw                                                                              FILE
WARNING   230     Feb  2 16:42:24 [redacted] THOR: Warning: MODULE: Filescan MESSAGE: Possibly Da  C:\U
                  ngerous file found SCANID: 20220201-PBU-AC3 FILE: C:\Users\pb\ThorSeed\robot\distribution\Pendul  \dis
                  um-WIN_thor_2021-12-09_1344.html EXT: .html SCORE: 230 TYPE: HTML SIZE: 52154 MD5: 8558e3f25c7  WIN
                  eab0503a8c12d34693f21 SHA1: d6bb2b0043c27f41e1e0e6d6c9991c58ab34f556 SHA256: bd1d2abdf8899342e3f
                  6a160e507ccf1efa8c914ac5c0d4f617786eaf2f828b7 FIRSTBYTES: 3c68746d6c3e0a3c686561643e0a093c7469746
                  c / <html> <head> <titl CREATED: Th                       9 14:38:53.459
                  2021 ACCESSED: Fri Dec 10 15:15:36.                       F / NT AUTHORIT
                  Y\SYSTEM:F / [redacted] pb:F OWNER                        SUSP_Password_D
                  umping_Keywords_Apr21 / Detects keyw                      scripts
                  SUBSCORE_1: 70 REF_1: https://twit                        0211 SIGTYPE_1:
                  internal MATCHED_1: Str1: "Get-Proc                       strong> Str1: &#
                  34;Get-Process lsass&#34; in &#34;ro                      TAGS_1: SUSP
                  RULEDATE_1: 2021-04-23 RULENAME_1:                        1: Florian Roth
                  REASON_2: YARA rule SUSP_PS1_LSASS_                       mmands in PowerS
                  hell to get the process object of th                      //twitter.com/sk
                  elsec/status/1400548671490514947?s=2                      Process lsass" i
                  n "ternal</br><strong>MATCHED_1:</st                      4;rosoft.com/win
                  dows:20H2 powershell.e" at 0x4967 TA                      E_2: 2021-06-04
                  RULENAME_2: SUSP_PS1_LSASS_Process_
WARNING   230     Feb  2 16:42:24 Pendulum-WIN/172.16.214.128 THOR: Warning: MODULE: Filescan MESSAGE: Possibly Da  C:\U
                  ngerous file found SCANID: 20220201-PBU-AC3 FILE: C:\Users\pb\ThorSeed\robot\distribution\Pendul  \dis
```

Selection:
bd1d2abdf8899342e3f6a160e507ccf1efa8c914ac5c
0d4f617786eaf2f828b7

Search for events
Filter out events
Search on Google
Search on VirusTotal
Copy to clipboard
Create case based on selection: Noteworthy
Create case based on selection: Legitimate Anomaly
Create case based on selection: False Positive

Fig. 17: Context Menu

CHAPTER

SIX

CASE MANAGEMENT BEST PRACTICES

This section assumes, that a 2-Level model as described in *Understanding Users, Roles, Rights and Case Status* is used.

The following actions will be explained:

- Opening a case
- Handing a case over to the next level
- Closing a case
- Reviewing the rules that add future logs to this case (Grouping Criteria)

6.1 Open a Case for Editing

The picture below shows the Case Management view with cases that have been created with the `Auto Case` feature. You can see those cases by filtering the `Name` column to only contain the `Auto Case` value and the `Status` column to only contain the `Open` value.

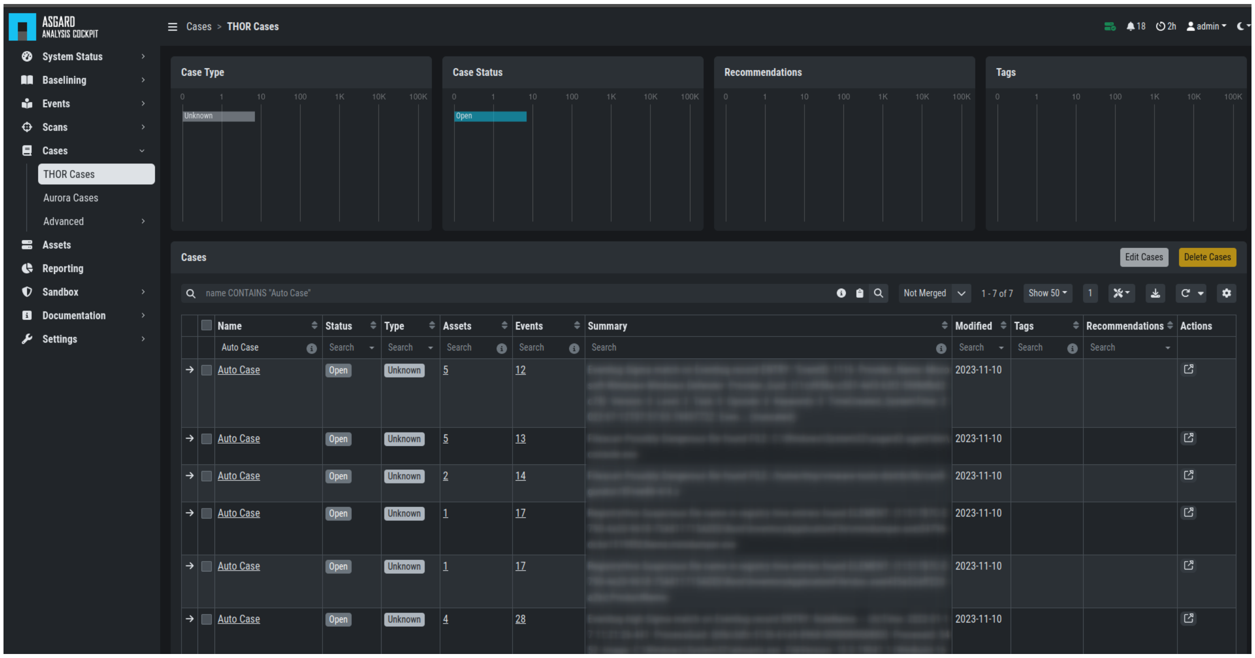

name CONTAINS "Auto Case"

Name	Status	Type	Assets	Events	Summary	Modified	Tags	Recommendations	Actions
Auto Case	Search	Search	Search	Search	Search	Search	Search	Search	
Auto Case	Open	Unknown	5	12	[redacted]	2023-11-10			
Auto Case	Open	Unknown	5	13	[redacted]	2023-11-10			
Auto Case	Open	Unknown	2	14	[redacted]	2023-11-10			
Auto Case	Open	Unknown	1	17	[redacted]	2023-11-10			
Auto Case	Open	Unknown	1	17	[redacted]	2023-11-10			
Auto Case	Open	Unknown	4	28	[redacted]	2023-11-10			

Fig. 1: Opening a Case for editing

In our example a Level 1 Analyst would now pick one of these open cases and set the Status to "Level 1 Working". To do this, they would open the case by clicking on the magnifier button and modify the status to `Level 1 Working` and then click `Update`.

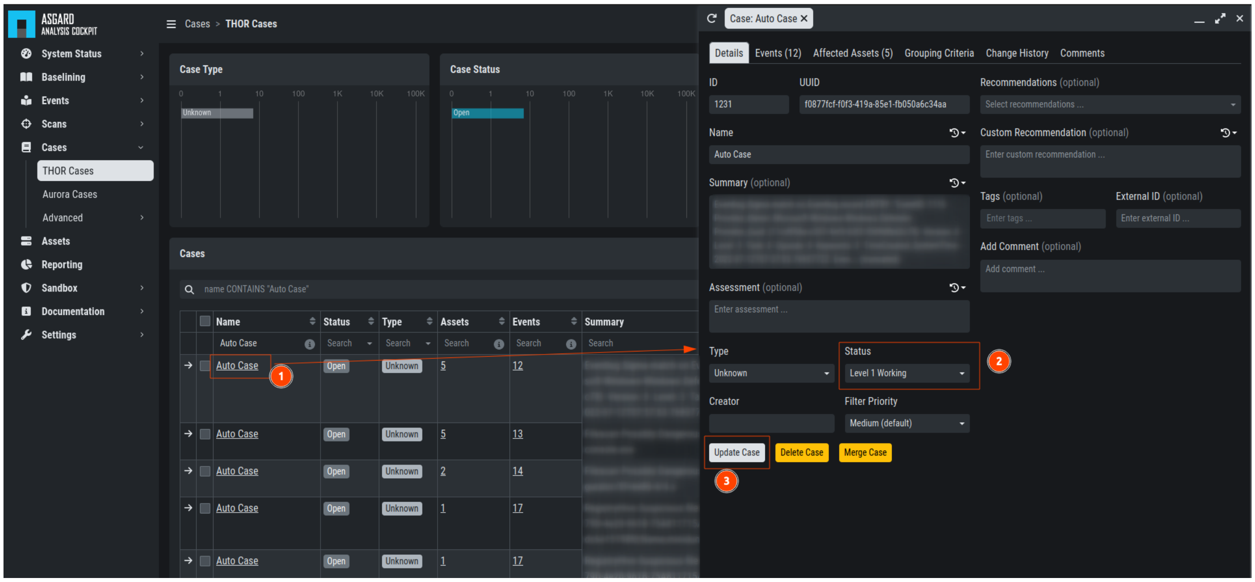

name CONTAINS "Auto Case"

Name	Status	Type	Assets	Events	Summary
Auto Case	Search	Search	Search	Search	Search
Auto Case	Open	Unknown	5	12	[redacted]
Auto Case	Open	Unknown	5	13	[redacted]
Auto Case	Open	Unknown	2	14	[redacted]
Auto Case	Open	Unknown	1	17	[redacted]
Auto Case	Open	Unknown	1	17	[redacted]

ID: 1231
UUID: f0877fcf-f0f3-419a-85e1-fb050a6c34aa
Recommendations (optional): Select recommendations ...
Name: Auto Case
Custom Recommendation (optional): Enter custom recommendation ...
Summary (optional): [redacted]
Tags (optional): Enter tags ...
External ID (optional): Enter external ID ...
Add Comment (optional): Add comment ...
Assessment (optional): Enter assessment ...
Type: Unknown
Status: Level 1 Working
Creator:
Filter Priority: Medium (default)

Fig. 2: Change Status

Now the logs within the case can be analyzed and results can be documented in the assessment field. Recommendations can be set from the canned recommendations list. Columns can be faded in and out and comments can be added.

6.2 Case Dispatching

Let's assume, our Level 1 Analyst concludes, that this is a "Legitimate Anomaly". They will now set the status to "Level 1 Finished" and update the case. After setting the case to "Level 1 Finished" the case becomes visible to the Level 2 Analyst.

6.3 Closing a Case

Let's assume, that a Level 2 Analyst now picks one of the cases in status "Level 1 Finished" and starts working on this case.

In this respect we assume, that something suspicious has been found, that needs further analysis by the system administration team. In most organizations this will be controlled through the organization's action request or ticketing system. So, we assume, that we will close the case in the Analysis Cockpit as it is progressed in another system. The status is changed to `closed` and the case gets updated.

Note: The Analysis Cockpit provides interfacing to action-request and external ticketing systems using the API.

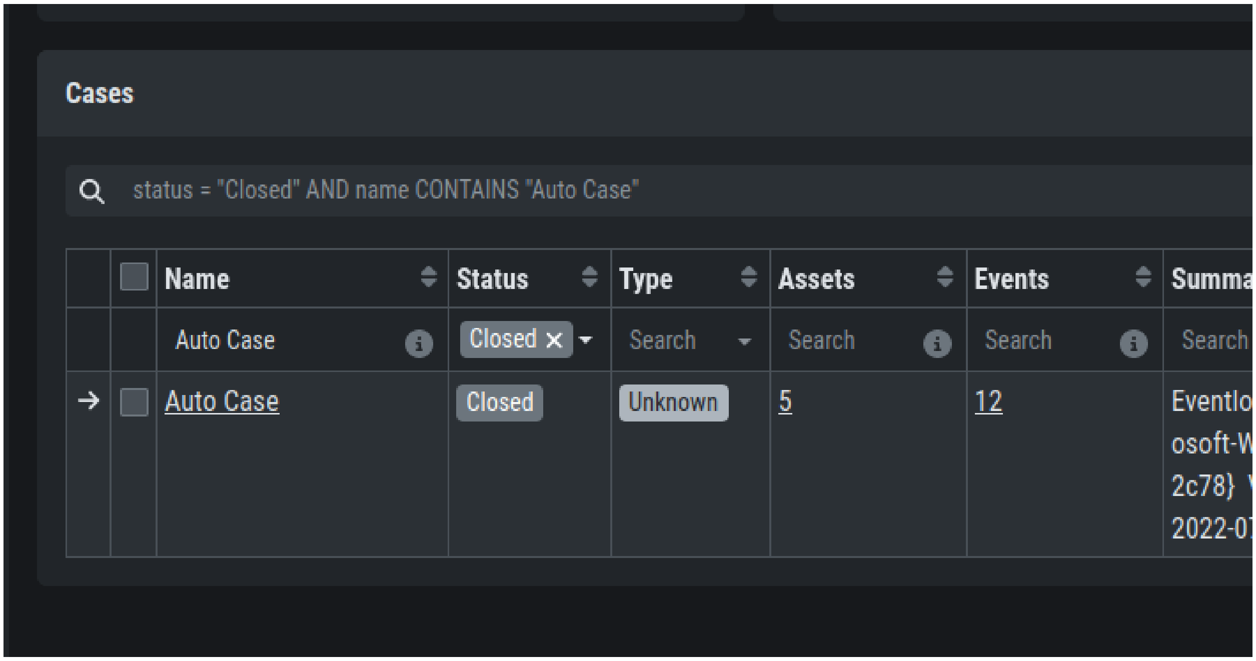

Cases

status = "Closed" AND name CONTAINS "Auto Case"

	□	Name	Status	Type	Assets	Events	Summa
		Auto Case	Closed ✕	Search	Search	Search	Search
→	□	Auto Case	Closed	Unknown	5	12	Eventlo osoft-W 2c78} 2022-0

Fig. 3: Closing a Case

6.4 Generate and Review auto_case_ids

These auto_case_ids can be reviewed in the `Grouping Criteria` section of the case.

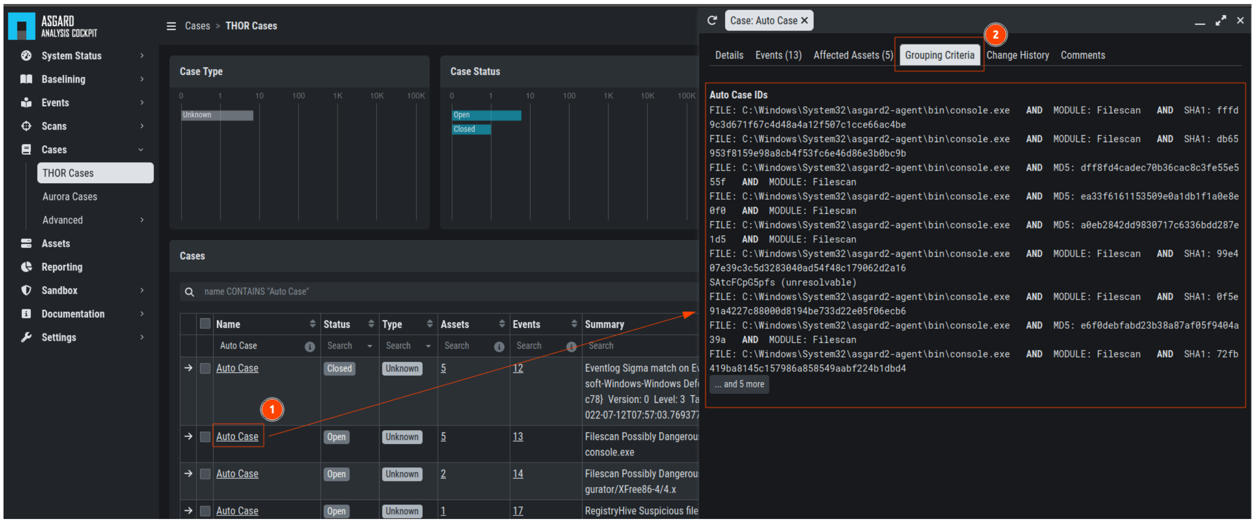

Cases

name CONTAINS "Auto Case"

	□	Name	Status	Type	Assets	Events	Summary
		Auto Case	Search	Search	Search	Search	Search
→	□	Auto Case	Closed	Unknown	5	12	Eventlog Sigma match on Ev soft-Windows-Windows Def c78} Version: 0 Level: 3 Ta 022-07-12T07:57:03.769377
→	□	Auto Case	Open	Unknown	5	13	Filescan Possibly Dangerou console.exe
→	□	Auto Case	Open	Unknown	2	14	Filescan Possibly Dangerou gurator/XFree86-4/4.x
→	□	Auto Case	Open	Unknown	1	17	RegistryHive Suspicious file

Auto Case IDs

```
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MODULE: Filescan   AND  SHA1: fffd
9c3d671f67c4d48a4a12f507c1cce66ac4be
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MODULE: Filescan   AND  SHA1: db65
953f8159e98a8cb4f53fc6e46d86e3b0bc9b
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MD5: dff8fd4cadec70b36cac8c3fe55e5
55f   AND  MODULE: Filescan
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MD5: ea33f6161153509e0a1db1f1a0e8e
0f0   AND  MODULE: Filescan
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MD5: a0eb2842dd9830717c6336bdd287e
1d5   AND  MODULE: Filescan
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MODULE: Filescan   AND  SHA1: 99e4
07e39c3c5d3283040ad54f48c179062d2a16
SAtcFCpG5pfs (unresolvable)
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MODULE: Filescan   AND  SHA1: 0f5e
91a4227c88000d8194be733d22e05f06ecb6
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MD5: e6f0debfabd23b38a87af05f9404a
39a   AND  MODULE: Filescan
FILE: C:\Windows\System32\asgard2-agent\bin\console.exe   AND  MODULE: Filescan   AND  SHA1: 72fb
419ba8145c157986a858549aabf224b1dbd4
```

... and 5 more

Fig. 4: Reviewing Grouping Criteria

In our example, three auto_case_ids were added that match all 1,000 log lines. In the future all incoming logs, that match one of the three "Detailed Reasons" will be added to this case directly and will not show up in the Log Management section.

6.4.1 Limitations

There are limitations to the visibility of grouping criteria. Grouping Criteria are only calculated for Alerts and Warnings. For all other types of logs (Notices, Info, Error) auto_case_ids are not calculated, so every log line gets its own highly specific filter that matches future occurrences of exactly the same log line but will not do any kind of generic matching. These highly specific filters are not displayed in the case for simplicities sake.

In rare cases the Analysis Cockpit will find it difficult to calculate auto_case_ids even for Alerts and Warnings. These logs will get tagged with `optimized\_template=false`. In this case, the behavior is like for Notices, Info and Error messages. Grouping Criteria will not show up as it will be one highly specific filter per log line.

6.5 More Information about Cases

The `Affected Assets` tab of a case shows assets that have contributed at least one log line to this case. In this example 5 assets are affected. All of them have the same operating system "windows".

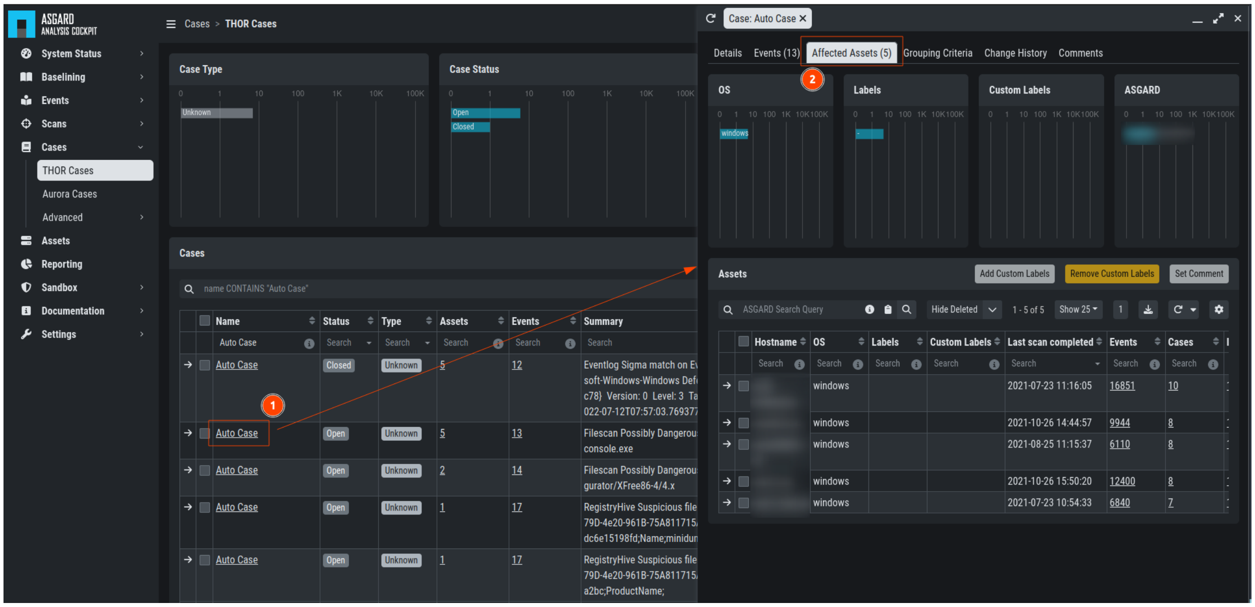

Cases

name CONTAINS "Auto Case"

Name	Status	Type	Assets	Events	Summary
Auto Case	Search	Search	Search	Search	Search
Auto Case	Closed	Unknown	5	12	Eventlog Sigma match on E soft-Windows-Windows Def c78} Version: 0 Level: 3 Ta 022-07-12T07:57:03.769377
Auto Case	Open	Unknown	5	13	Filescan Possibly Dangerou console.exe
Auto Case	Open	Unknown	2	14	Filescan Possibly Dangerou gurator/XFree86-4/4.x
Auto Case	Open	Unknown	1	17	RegistryHive Suspicious file 79D-4e20-961B-75A811715 dc6e15198fd;Name;minidun
Auto Case	Open	Unknown	1	17	RegistryHive Suspicious file 79D-4e20-961B-75A811715 a2bc;ProductName;

Assets

ASGARD Search Query

Hostname	OS	Labels	Custom Labels	Last scan completed	Events	Cases
Search	Search	Search	Search	Search	Search	Search
[redacted]	windows			2021-07-23 11:16:05	16851	10
[redacted]	windows			2021-10-26 14:44:57	9944	8
[redacted]	windows			2021-08-25 11:15:37	6110	8
[redacted]	windows			2021-10-26 15:50:20	12400	8
[redacted]	windows			2021-07-23 10:54:33	6840	7

Fig. 5: Case – Assets Tab

In the `Comments` tab you can add comments and attachments to this case. Attachments can be used to pass additional information to members of the analysis team (e.g. memory dump for further analysis).

The `Changes` tab shows information about changes to this case.

In other words: This is your case audit log.

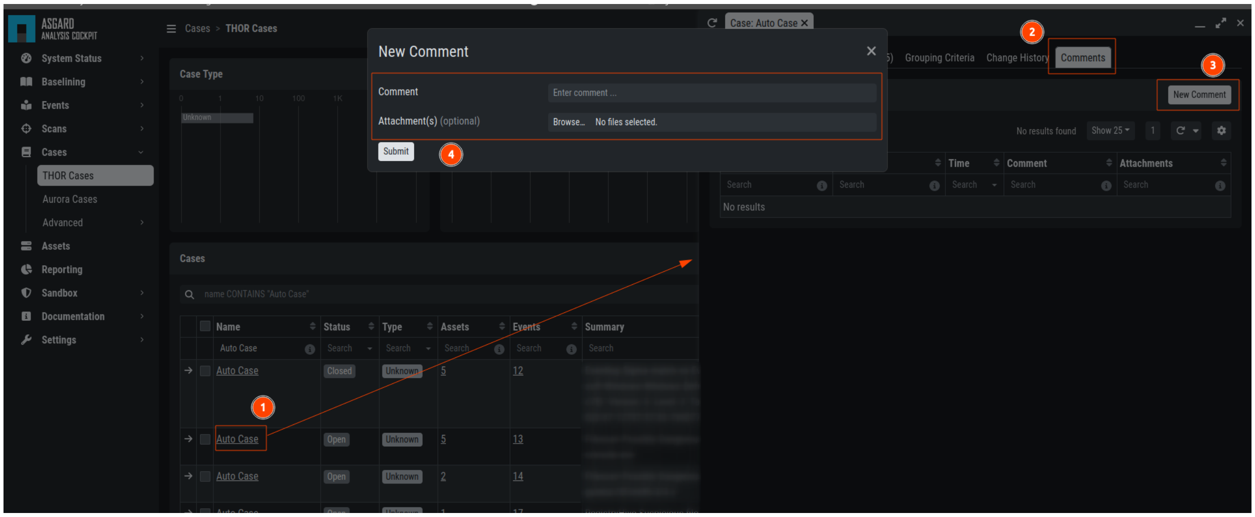

New Comment

Comment: Enter comment ...

Attachment(s) (optional): Browse... No files selected.

Submit

No results found

Time	Comment	Attachments
No results		

Fig. 6: Case – Comments

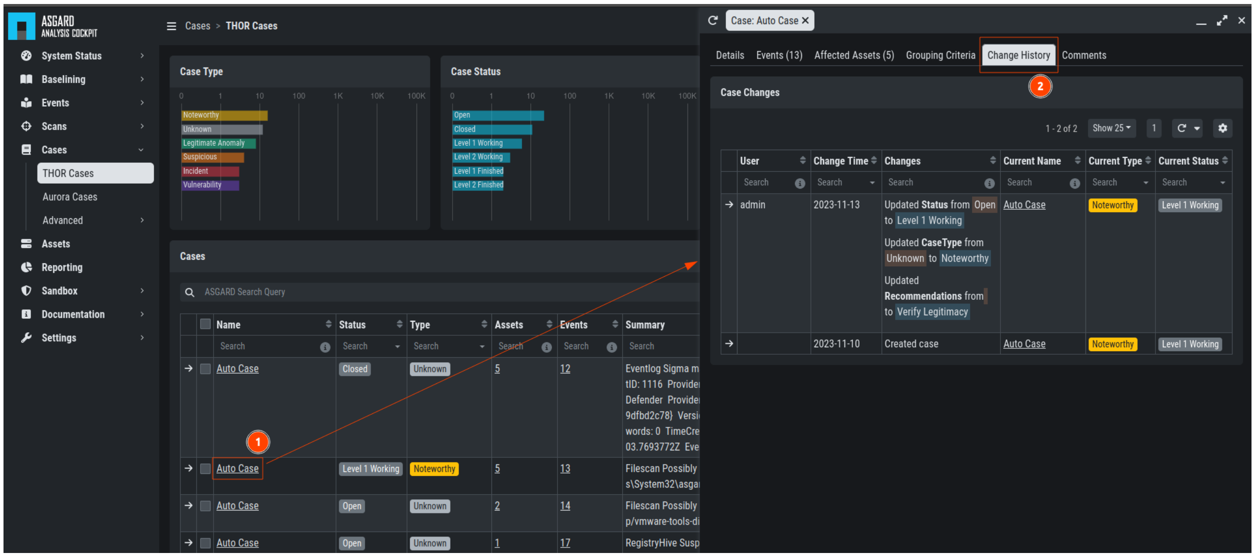

Case Changes

1 - 2 of 2 Show 25

	User	Change Time	Changes	Current Name	Current Type	Current Status
→	admin	2023-11-13	Updated **Status** from Open to Level 1 Working Updated **CaseType** from Unknown to Noteworthy Updated **Recommendations** from to Verify Legitimacy	Auto Case	Noteworthy	Level 1 Working
→		2023-11-10	Created case	Auto Case	Noteworthy	Level 1 Working

Fig. 7: Case – Changes tab

6.6 Bulk Edit / Bulk Delete

The Analysis Cockpit features a convenient way to make certain changes to groups of cases. Just select the case in the left column and click the `Edit Cases` or `Delete Cases` button.

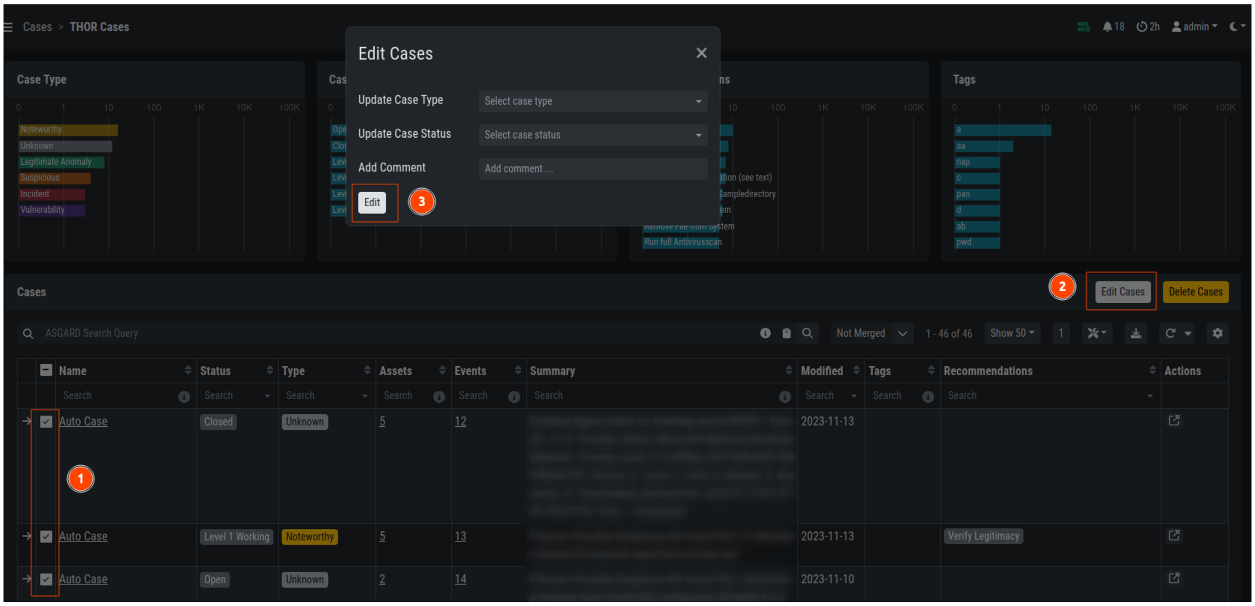

	Name	Status	Type	Assets	Events	Summary	Modified	Tags	Recommendations	Actions
	Search	Search	Search	Search	Search	Search	Search	Search	Search	
☑	Auto Case	Closed	Unknown	5	12	[redacted]	2023-11-13			
☑	Auto Case	Level 1 Working	Noteworthy	5	13	[redacted]	2023-11-13		Verify Legitimacy	
☑	Auto Case	Open	Unknown	2	14	[redacted]	2023-11-10			

Fig. 8: Bulk Edit

CHAPTER

SEVEN

MAINTENANCE

This chapter contains information on how you can regain disk space on your system or change settings which are not found in your Web UI.

7.1 Configuration Backup & Restore

The Analysis Cockpit comes with a backup and restore function for its configuration. The Configuration Backup contains the following data:

- Cases, Grouping Criteria, Recommendations, Case Changes, Case Comments
- Users, Roles, LDAP Roles, Role Rights
- User Configurations

To perform a backup, you can simply go to `Settings > Backup` and click `Create Configuration Backup`. To restore from an old backup, it is important to understand the implications of the restore. From the Backup page of the Analysis Cockpit:

> The restore procedure will install a previously generated configuration backup on this Analysis Cockpit. All data on this Analysis Cockpit will be deleted before. This can only be done on newly installed Analysis Cockpits and not on an Analysis Cockpit that is already in use. Do not use the restore to rollback to an earlier point in time, this will cause inconsistent data.

> **Warning:** Installing a configuration backup of an earlier Analysis Cockpit Release Version is not supported and may fail. The currently installed version is 3.5.6. The version of the configuration backup can be found in the file name. The backup's file name has the following pattern: `analysis_cockpit_%VERSION%_backup_%DATE%.sql.gz`

7.2 Regain Disk Space

If your disk is already at or close to 100% and AC no longer works properly, see section *Recover from a Full Disk*.

If your disk usage is growing too fast and free disk space is running out, you have several options:

1. Increase the size of your disk
2. Delete files that are not needed for operation
3. Delete files that are used by AC but are unneeded / dated

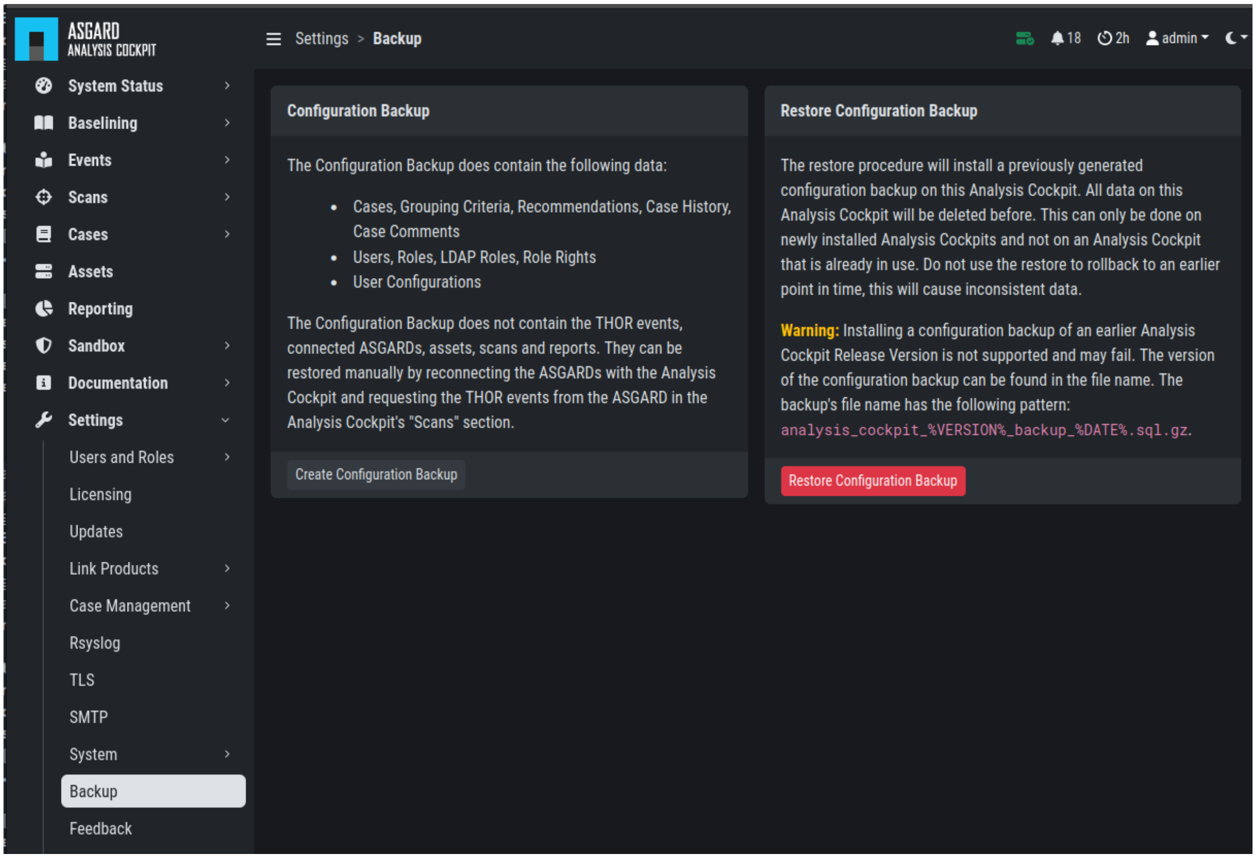

Configuration Backup

The Configuration Backup does contain the following data:

- Cases, Grouping Criteria, Recommendations, Case History, Case Comments
- Users, Roles, LDAP Roles, Role Rights
- User Configurations

The Configuration Backup does not contain the THOR events, connected ASGARDs, assets, scans and reports. They can be restored manually by reconnecting the ASGARDs with the Analysis Cockpit and requesting the THOR events from the ASGARD in the Analysis Cockpit's "Scans" section.

Create Configuration Backup

Restore Configuration Backup

The restore procedure will install a previously generated configuration backup on this Analysis Cockpit. All data on this Analysis Cockpit will be deleted before. This can only be done on newly installed Analysis Cockpits and not on an Analysis Cockpit that is already in use. Do not use the restore to rollback to an earlier point in time, this will cause inconsistent data.

Warning: Installing a configuration backup of an earlier Analysis Cockpit Release Version is not supported and may fail. The version of the configuration backup can be found in the file name. The backup's file name has the following pattern: `analysis_cockpit_%VERSION%_backup_%DATE%.sql.gz`.

Restore Configuration Backup

Fig. 1: Configuration Backup & Restore

7.2.1 Safe-to-Delete Files

The following files are safe to delete. They are not needed for AC to operate.

- `/var/lib/asgard-analysis-cockpit/log/*.gz`
- `/var/lib/asgard-analysis-cockpit/events/*.ok`

They are only kept on the system if needed for further processing. E.g. saving/sending the log files to another system or keeping the THOR scans (found in `events`) for backup reasons. If you do not need or plan to use those, they can be deleted. If you are unsure make a copy to another system before deleting them.

More details can be found in section *Recover from a Full Disk*.

7.2.2 Potentially Unneeded / Dated Files

This method is only advised as a last resort if increasing your disk space is not an option.

If your AC is running for a long time, there might be data ingested that you no longer need and therefore can be deleted to regain disk space. This includes:

- Scans
- Reports

Deleting Unneeded Scans

Warning: Deleting old scans deletes information ASGARD Analysis Cockpit uses.

As an example: If you delete a scan with which an asset was marked in an incident case, this connection is no longer made and the asset will be shown with 0 incident cases.

Therefore only delete scans you no longer need. This can be done under `Scans > Scans` by selecting the scans with check marks and clicking `Delete Events`.

You can filter events for deletion with the time range picker in the completed column and e.g. selecting only scans with 0 incident and 0 suspicious cases. (Add columns using the `Columns` button).

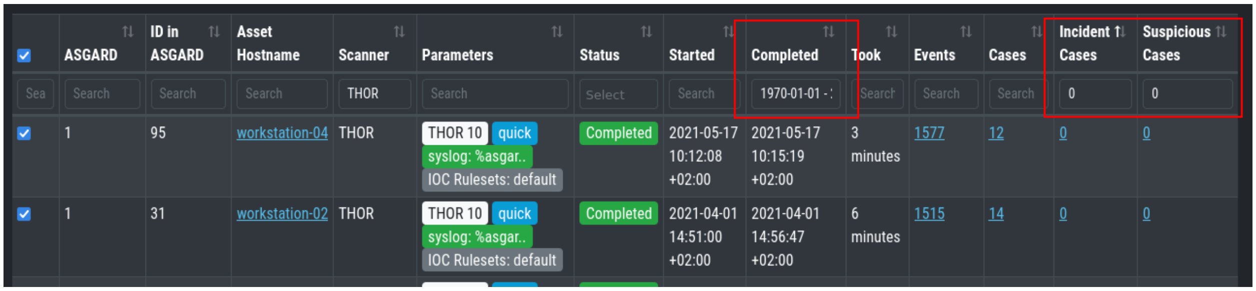

☑	ASGARD	ID in ASGARD	Asset Hostname	Scanner	Parameters	Status	Started	Completed	Took	Events	Cases	Incident Cases	Suspicious Cases
Sea	Search	Search	Search	THOR	Search	Select	Search	1970-01-01 -:	Search	Search	Search	0	0
☑	1	95	workstation-04	THOR	THOR 10 quick syslog: %asgar.. IOC Rulesets: default	Completed	2021-05-17 10:12:08 +02:00	2021-05-17 10:15:19 +02:00	3 minutes	1577	12	0	0
☑	1	31	workstation-02	THOR	THOR 10 quick syslog: %asgar.. IOC Rulesets: default	Completed	2021-04-01 14:51:00 +02:00	2021-04-01 14:56:47 +02:00	6 minutes	1515	14	0	0

Fig. 2: Possible Filter for Selecting Scans for Deletion

Another possibility is searching for assets which are no longer part of your infrastructure and deleting their scans.

Deleting Unneeded Reports

Old unneeded reports can be deleted via command line and are found at `/var/lib/asgard-analysis-cockpit/reports`.

Note: The reports are still listed in the UI after removal, but a download attempt will fail.

7.3 Increasing ElasticSearch's Heap Space

When installing your Analysis Cockpit, your Elasticsearch instance running on the same server will be initialized with a heap space of 2GB. This means, it will only use up to 2GB RAM to perform search queries. This might also cause unused RAM and could lead to issues in rare cases.

This issue is related to ElasticSearch, which stores your Analysis Cockpit's events. Elasticsearch calculates the required RAM for operations before executing them.

If you recently increased the RAM of your server, you need to change the below configuration to make use of it. If you did not increase your RAM, but initially set up your server with more RAM, you can also go ahead and change this.

To increase heap space for ElasticSearch, edit the following configuration file on your Analysis Cockpit:

```
nextron@cockpit:~$ sudoedit /etc/elasticsearch/jvm.options.d/10-cockpit.options
```

You should see the following default values:

```
-Xms2g
-Xmx2g
```

- Xms represents the initial size of total heap space
- Xmx represents the maximum size of total heap space

The `2g` part of the values indicates the heap space in gigabytes. We advise to use 50% of your system's memory for ElasticSearch. On a system with a maximum of 8 GB of RAM, this would be `4g`:

```
-Xms4g
-Xmx4g
```

After you saved your changes, restart the elasticsearch service (this could take a few seconds!):

```
nextron@cockpit:~$ sudo systemctl restart elasticsearch.service
```

Make sure the service is in `active (running)` state after you restarted it:

```
nextron@cockpit:~$ sudo systemctl status elasticsearch.service
```

CHAPTER

EIGHT

TYPICAL PITFALLS

This chapter contains typical pitfalls.

8.1 Certificate Validation Failed

If you receive the following error, SSL/TLS interception interrupted the installation process.

```
nextron@cockpit:~$ sudo nextronInstaller -cockpit
[sudo] password for nextron:
Ign:1 https://update3.nextron-systems.com analysis InRelease
Err:2 https://update3.nextron-systems.com analysis Release
Certificate verification failed: The certificate is NOT trusted. The certificate issuer␣
↪is unknown. Could not handshake: Error in the certificate verification. [IP: 192.168.3.
↪21 8080]
```

Since we do not support setups in which the connections to our update servers are intercepted (see chapter *SSL/TLS Interception*), the only way to resolve this problem is to deactivate SSL/TLS interception for our update servers.

8.2 Log File Import of Previous Years

The log file format of (old) THOR scan logs is the original SYSLOG format, which contains no year value in the timestamp of the message header.

You can modify the timestamp of old THOR logs by using the following script:

https://github.com/NextronSystems/nextron-helper-scripts/blob/master/asgard-analysis-cockpit/thor-timestamp-coverter.py

8.3 Disk Watermark

Elasticsearch has a disk watermark that it uses to determine if it should go into read-only mode. If the disk is too full, Elasticsearch will stop accepting new data. This watermark is set to 95% by default and will prevent data loss by stopping the system from writing to the disk.

The disk watermark of the Analysis Cockpit however is set to 90%. This means on a 1TB drive you need at least 100GB of free space or the Analysis Cockpit will put itself into read-only mode. We set this value lower than the default of Elasticsearch to give you more time to react before Elasticsearch goes into read-only mode, in which case you would need to reset the read-only mode manually on Elasticsearch.

The below message shows up in the Analysis Cockpit if the disk watermark is reached:

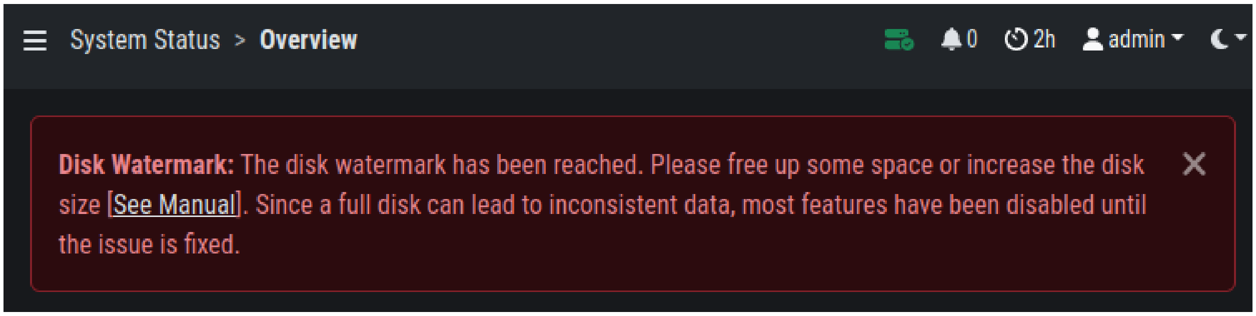

Disk Watermark: The disk watermark has been reached. Please free up some space or increase the disk size [See Manual]. Since a full disk can lead to inconsistent data, most features have been disabled until the issue is fixed.

Fig. 1: Disk Watermark

If you see this message, the Analysis Cockpit went into a read-only mode and you need to free up some disk space or increase the disk space by allocating more storage to the virtual machine.

To free up some disk space, you can follow the instructions in the next chapter (*Regain Disk Space*).

If your disk usage somehow got above 95%, Elasticsearch will go into read-only mode nontheless. If this happened, you need to reset the Elasticsearch state after you freed up some disk space. You can achieve this by running the following command:

```
nextron@cockpit:~$ curl -X PUT -s -u elastic:$(cat /etc/asgard-analysis-cockpit/elastic.
↪password) \
-H 'Content-Type: application/json' \
-d '{"index.blocks.read_only_allow_delete": null}' \
http://localhost:9200/_all/_settings
```

You should get the following output if the command was successful:

```
{"acknowledged":true}
```

Note: Please note that the password changes after the Analysis Cockpit was restarted, this is why we `cat` the password directly from the file.

8.4 Recover from a Full Disk

If your disk is full or near full, ASGARD Analysis Cockpit will not work properly. In order to resume its operation you need to make free space on the disk.

We suggest to save the files to another system beforehand, if you want to keep the information for future usage. ASGARD will not need the following files to function and they can be removed safely:

- `/var/lib/asgard-analysis-cockpit/log/*.gz`
- `/var/lib/asgard-analysis-cockpit/events/*.ok`

Especially the assignment log can grow big in production environments. If deleting the logs is not enough, deleting the already read-in events (ending on `.ok`) is the next best location to regain disk space. If there are too many files for a simple `rm *.ok`, you can use find to delete them:

```
nextron@cockpit:~$ sudo su -
[sudo] password for nextron:
root@cockpit:~# find /var/lib/asgard-analysis-cockpit/events -name "*.ok" -print0 |␣
↪xargs -0 -I'{}' rm '{}'
```

If Elasticsearch does not automatically work again after cleaning up some disk space, restart it under `Settings` > `System` > `Services` or with `sudo systemctl restart elasticsearch.service`. If this is not working either, you may need to disable Elasticsearch's read-only mode. See *Disk Watermark* for a how-to.

Deleting the files given above should be enough to resume operation. If the disk on your ASGARD Analysis Cockpit is full because of growing data over time, the disk space should be increased. If that is not an option you can delete old scans as described in section *Regain Disk Space*.

8.5 Debug Failed File Imports

Check for reported problems using this command:

```
nextron@cockpit:~$ sudo su -
[sudo] password for root:
nextron@cockpit:~$ find /var/lib/asgard-analysis-cockpit/events -name "\*.problem"
```

Make sure that you're able to see the imported log data and review the selected time range in the time range picker in whatever view you're reviewing the data. Be aware that the log data gets indexed with the creation timestamp of the log lines not the time of their import.

This means that if you're importing log data that is old, the default date range set in the date range picker may be too narrowly defined so that you're just unable to see the imported data.

8.6 Fixing a Broken Proxy Configuration

Sometimes during installation, proxy settings get mixed up or a typo in the proxy URL leads to a broken Internet connection.

It is not trivial to fix this situation, since the proxy settings collected during installation are changed in so many different locations on a Linux system for all the different services and command line tools.

8.6.1 Broken before Analysis Cockpit Installation

If you have set a wrong proxy before the package installation using the **sudo nextronInstaller -cockpit** command and the installer failed to fetch the required packages from our update servers, perform the following steps.

Fix the proxy string in the file `/etc/apt/apt.conf.d/00proxy`

```
nextron@cockpit:~$ sudoedit /etc/apt/apt.conf.d/00proxy
```

Then rerun the installer.

8.6.2 Broken after the Analysis Cockpit Installation

If your infrastructure has changed and you have to change the proxy server sometime later, edit the proxy settings in the Web GUI.

`Settings` > `System` > `Proxy`

CHAPTER

NINE

FAQS

This section has frequently asked questions and answers to them. You will find that the format of this section is split into a question as the introduction of each chapter and the explanation right after.

9.1 Disabling Assignment Logs

Q: My assignment Logs on the server are growing quickly, how can I turn them off

The assignment logs located at `/var/lib/asgard-analysis-cockpit/log/assignment.log` write warnings and errors for the `Optimize` function of the Cockpit.

If you have the feeling that the log is filling up too quickly, you can turn off those logs completely. It is advised to try and see what the problem is before turning off the log completely, as this might indicate an underlying issue.

Run the following command on your Analysis Cockpit (warning: this will restart your Analysis Cockpit. If you do not want to restart the Analysis Cockpit, you can run the second command at a later time):

```
nextron@cockpit:~$ echo "REPLACE INTO config VALUES ('write-assignment-log','false')" |␣
↪sudo mysql analysiscockpit
nextron@cockpit:~$ sudo systemctl restart asgard-analysis-cockpit.service
```

To turn back on the `assignment.log`, run the following command:

```
nextron@cockpit:~$ echo "REPLACE INTO config VALUES ('write-assignment-log','true')" |␣
↪sudo mysql analysiscockpit
nextron@cockpit:~$ sudo systemctl restart asgard-analysis-cockpit.service
```

9.2 No Events visible

Q: It seems that events are not visible or have been lost. What can I do to verify that they're still in the database?

If you think that some events are not visible or have been lost, you can do the following to verify that they still exist in the database.

First, check your date range picker.

Very often, analysts forget to set it to the right time frame and old events accidentally disappear from the view.

Secondly, make sure you're using the search in the `Events` section and not the `Baselining` section.

9.3 No new Events in Case

Q: I have created a case but it seems that no new incoming events are assigned to that existing case. How can I check what's wrong?

The first thing that you should check are the `auto_case_ids` of the events in that case (`Cases` > `Open Case` > `Events` > `auto_case_id` Panel).

If they are distributed as in the following screenshot, it seems that auto-casing doesn't work on this case.

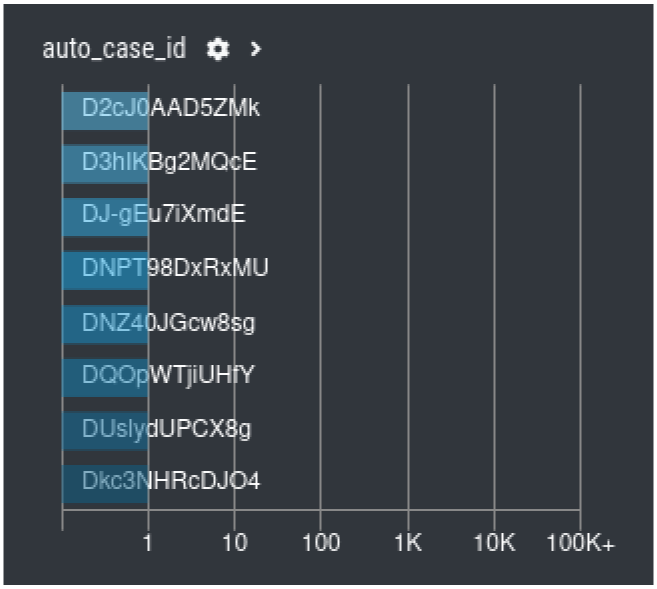

This case doesn't have groupable contents and uses only so-called "Dynamic Auto Case IDs", which are used whenever the Analysis cockpit was unable to find a suitable filter template to create usable filters for this type of events.

Also check the grouping criteria of that case:

`Cases` > `Open Case` > Tab `Grouping Criteria`

What are the conditions defined to assign new events to that case?

9.4 Location of Scan Logs

Q: Where are Scan Logs on the system located?

You can find the Scan Logs in `/var/lib/asgard-analysis-cockpit/events`. In this folder you will find three different naming schemes:

- **.txt.gz** - Logs which are not imported yet
- **.txt.gz.ok** - Logs which were imported successfully
- **.txt.gz.problem** - Logs which could not be imported correctly due to an error

If you need to manually investigate logs which failed during the import (.gz.problem), you can do so by copying the files to a different location (`/tmp` for example) and remove the suffix .problem. After that you can use `gunzip` to extract the log and inspect it. Most likely you will find that the file did not transfer correctly over to the Analysis Cockpit. This can be seen if you open the file and scroll to the very end. In this case the file will just end in the middle of a log line.

The Logs can be imported into the Cockpit via the `Scans` menu. Select the Asset which had a problem with the log transfer and click `Request Events`. This will transfer the Events from the corresponding ASGARD. You can also use the Fields **Log Requested**, **Log Received** and **Log Received Error** to filter and look for other failed log transmissions.

9.5 Default password for file downloads

Q: What is the password used to protect file downloads?

Artifacts uploaded to a case might be malware. To ensure the file is not automatically deleted by antivirus or executed by an unknowing user, we zip all files in the attachments and encrypt the ZIP file with a default password. The default password `infected` can be used to extract the file.

9.6 Disk Space filling up quickly

Q: My disk is getting full soon. What options do I have?

If your disk is already at or close to 100% and AC no longer works properly, see section *Recover from a Full Disk*.

In other cases check section *Regain Disk Space*.

9.7 Reverse Proxy to access the Analysis Cockpit

Q: I am using a Reverse Proxy to access the Analysis Cockpit. What do I have to take care of?

The Analysis Cockpit partially uses large URLs to communicate with its backend. Proxy server usually do not allow arbitrary large URLs.

In case of nginx the default header size is 8k (see http://nginx.org/en/docs/http/ngx_http_core_module.html#large_client_header_buffers). If you want to use the Analyst Cockpit behind a nginx reverse proxy, you need to increase the *large_client_header_buffer*. A size of 100k should be sufficient. Also the HTTP2 protocol has to be disabled.

A minimal example configuration for nginx looks as follows:

```
server {
   listen 443 ssl; # !! no http2 !!
   ssl_certificate /path/to/your/certificate.crt;
   ssl_certificate_key /path/to/your/private.key;
   location / {
      proxy_pass https://analysis-cockpit.your.org;
      proxy_set_header Host $http_host;
      proxy_set_header X-Real-IP $remote_addr;
      proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
   }
   large_client_header_buffers 4 100k; # increase maximal allowed URL length
}
```

9.8 Internet Explorer

Q: I am using Internet Explorer and the Analyst Cockpit seems to run into a timeout. What can I do?

Modern browsers (e.g. Firefox, Chrome, Edge, Safari) support large URLs. Internet Explorer does not. If you want to access the Analyst Cockpit and all its features, you need to switch your browser.

9.9 Admin Password reset

Q: I forgot my admin password and lost access to the WebUI. How do I reset the admin user password?

If you've lost the password of the local `admin` user (Web GUI) but still have access the system via SSH, you can reset it via command line using the following command.

```
nextron@cockpit:~$ sudo mysql analysiscockpit -e "UPDATE users SET password =
↪'7951GYqdAjLAoO1NaQu1ManJDIk' WHERE name = 'admin';"
```

This resets the password to `admin`. You should then change that password immediately.

9.10 Multi Factor Authentication reset

Q: How do I reset Multi Factor Authentication for a specific user

If you or another user lost their second factor (MFA) to log into the ASGARD Web UI, you can reset the users MFA Settings with the following command (in this example we assume that the user is called `john`):

```
nextron@cockpit:~$ sudo mysql analysiscockpit --execute "UPDATE users SET tfa_valid = 0␣
↪WHERE name = 'john';"
```

CHAPTER

TEN

KNOWN ISSUES

You can find a list of known issues in this section.

10.1 AAC#003: [WAR] could not create case

Introduced Version	Fixed Version
4.0.10	4.0.12

There is currently a bug in the Analysis Cockpit when creating cases, if the following conditions are given:

- Baseline view with a limited time frame (e.g. 30 days)
- Trying to create a `Condition Case`

When you try to create a case with the above criteria given, you will receive the following error:

```
[WAR] could not create case ERROR: can not create case with more than one source (query /
↪ condition / regex) set
```

10.1.1 AAC#003: Workaround

While we are working on the fix, you can do the following to avoid the above error:

- Change the time frame to `All time` in your baseline view

10.2 AAC#002: Scan stuck at Status "Unknown"

Introduced Version	Fixed Version
4.0.10	4.0.11

There is currently a bug in the Analysis Cockpit which prevents some Scans from being imported correctly.

This is caused by very big events (a single event bigger than 64 Kb), which will cause the parser to error. The Analysis Cockpit can never finish importing this Scan.

10.2.1 AAC#002: Fix

We are currently testing the fix, which will skip larger events and finish importing the scan logs.

You will also have the possibility to set the maximum size of a single log line in the advanced options, once the fix is released.

You will additionally see failed Log imports in the Dashboard of your Analysis Cockpit.

10.2.2 AAC#002: Check

You can check if one of your scan logs is effected if the following conditions are met:

You will see a scan which is in the Status `Unknown`

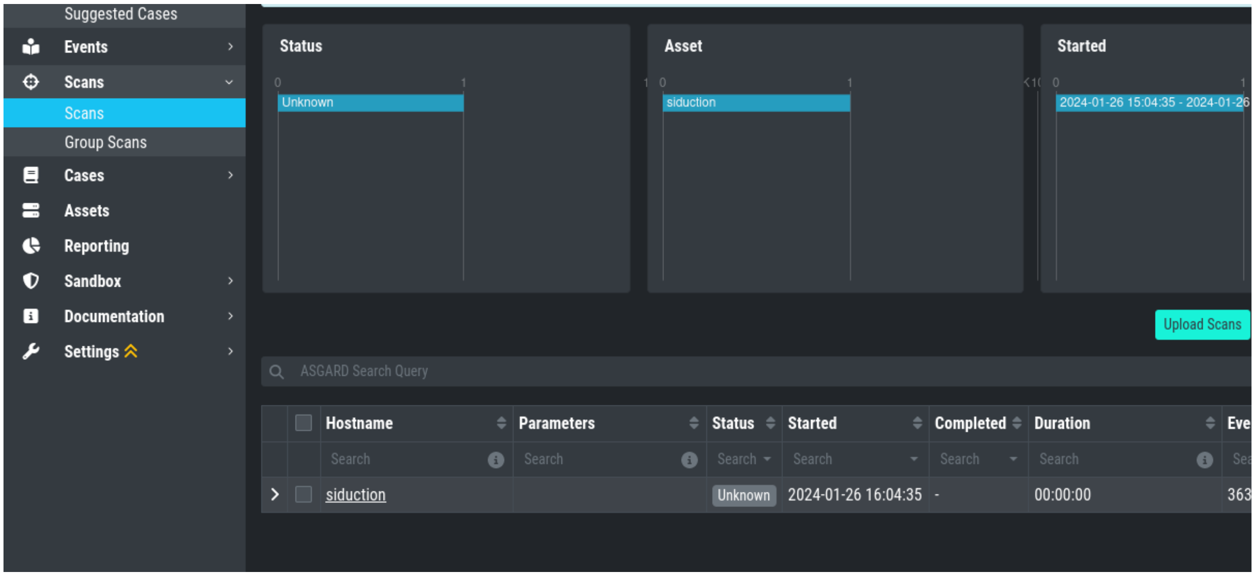

Hostname	Parameters	Status	Started	Completed	Duration	Eve
Search	Search	Search	Search	Search	Search	Se
siduction		Unknown	2024-01-26 16:04:35	-	00:00:00	363

When you connect to your Analysis Cockpit via SSH and enter a root session, you can execute the following command to see if the error occured on one or more log files:

```
root@analysis:# grep -R "ERROR: bufio.Scanner: token too long" /var/log/asgard-analysis-
↪cockpit
Jan 26 16:18:49 analysis analysiscockpit4[29459]: 2024-01-26T15:18:49Z [ERR] could not␣
↪read events from file PATH: /var/lib/asgard-analysis-cockpit/events/upload_siduction_
↪thor_2024-01-06.txt ERROR: bufio.Scanner: token too long
```

You should see from the above output which log had problems, which should also be reflected in the filename:

```
root@analysis:# ls /var/lib/asgard-analysis-cockpit/events
upload_siduction_thor_2024-01-06.txt.problem
```

The file has the `.problem` suffix, which indicates a problem during the import.

Once you installed the update you can re-import the failed scan logs. You can either upload them manually again, or rename the files from the output above (remove the `.problem` suffix).

10.3 AAC#001: Could not get table data: Data too large

This issue is related to ElasticSearch, which stores your Analysis Cockpit's events. Elasticsearch calculates the required RAM for operations before executing them.

The below error might occur on complex searches or aggregations (e.g. for the graphs in the baselining view). To fix this issue, you have to increase the RAM of your Analysis Cockpit and reconfigure ElasticSearch to actually use more RAM.

10.3.1 AAC#001: Fix

To actually fix the problem, you have to allocate more RAM to your Analysis Cockpit. You should be able to do this via your hypervisor.

To increase heap space for ElasticSearch, edit the following configuration file on your Analysis Cockpit:

```
nextron@cockpit:~$ sudoedit /etc/elasticsearch/jvm.options.d/10-cockpit.options
```

You should see the following default values:

```
-Xms2g
-Xmx2g
```

- Xms represents the initial size of total heap space
- Xmx represents the maximum size of total heap space

The `2g` part of the values indicates the heap space in gigabytes. We advise to use 50% of your system's memory for ElasticSearch. On a system with a maximum of 8 GB of RAM, this would be `4g`:

```
-Xms4g
-Xmx4g
```

After you saved your changes, restart the elasticsearch service (this could take a few seconds!):

```
nextron@cockpit:~$ sudo systemctl restart elasticsearch.service
```

Make sure the service is in `active (running)` state after you restarted it:

```
nextron@cockpit:~$ sudo systemctl status elasticsearch.service
```

CHAPTER

ELEVEN

UPGRADE FROM COCKPIT V3.10.1 TO COCKPIT V4.X

This Chapter contains instructions on how to upgrade your running Analysis Cockpit version 3.10.1 to the newest version 4.

Hint: If you want to stay on version 3 of the Analysis Cockpit and still install updates, you can run the following commands via SSH:

```
nextron@analysis:~$ sudo apt update
nextron@analysis:~$ sudo apt upgrade
```

There are two chapters, one for a **standalone installation** and one for an installation with an **Elasticsearch Cluster**. Depending on your environment, please follow **only one** section:

- Standalone Installation - *Standalone Upgrade*
- Cluster Installation - *Cluster Upgrade*

We developed an update program which helps you through the upgrade by automating as much as possible. You still have to upgrade your Elasticsearch cluster manually, due to version limitations with your master node.

Important: Your cockpit will be unavailable for an extended period of time during the upgrade proccess, usually between 30 and 60 minutes. Please make sure that no user is working with the Analysis Cockpit during the upgrade process.

11.1 Cluster Upgrade

This chapter guides you through the upgrade process of your Analysis Cockpit version 3.10.1 to version 4.x.

It is important to follow the steps carefully. We advise you to create a snapshot of all your Elasticsearch cluster nodes and the Analysis Cockpit itself before starting your upgrade.

Danger: Please do not update your Analysis Cockpit before you update your Elasticsearch Cluster Nodes. This can potentially break your envionment.

11.1.1 Preparation

To prepare for your upgrade, we compiled a list of tasks you should follow:

Task	Description
Snapshot of your Analysis Cockpit	For disaster recovery
Snapshot of your Elasticsearch Cluster Nodes	For disaster recovery
Analysis Cockpit running version 3.10.1	Prerequisite for the Major Upgrade
Newest `asgard-updater` is installed	This performs the update and has to be in the newest version
Cluster status is "green"	We don't want to upgrade a non functional cluster
Connection to our new update servers	New update server infrastructure
Stop your Analysis Cockpit - `optional`	To ensure no scans are being synchronized from your Management Center

For details regarding some of the above tasks, see the next section in this manual.

With the new version of your Analysis Cockpit, we also made changes to our update servers. Please make sure that all your components can reach the following servers:

Server	Port	Description
update3.nextron-systems.com	tcp/443	Old update server
update-301.nextron-systems.com	tcp/443	New update Server

The old update server is needed to fetch the updater and other prerequisites. The new update server is needed to upgrade your servers to Debian 12 and also to install any new packages, which are needed for your Analysis Cockpit v4.

You can find the corresponding IP-Addresses to the above FQDNs here: https://www.nextron-systems.com/resources/hosts/.

Analysis Cockpit running version 3.10.1

To check if your Analysis Cockpit is running on the correct version. You can navigate to `Settings` and `Updates`. The page should look like this:

Newest `asgard-updater` is installed

This step should be performed on your Analysis Cockpit and all your cluster nodes.

To check if a newer version of the `asgard-updater` is available, we have to run the following commands. If you get the highlighted output, you have already the newest version installed (the version from the output might be newer in your case):

```
nextron@analysis:~$ sudo apt update
nextron@analysis:~$ sudo apt install asgard-updater
Reading package lists... Done
Building dependency tree
Reading state information... Done
asgard-updater is already the newest version (1.0.17).
0 upgraded, 0 newly installed, 0 to remove and 18 not upgraded.
```

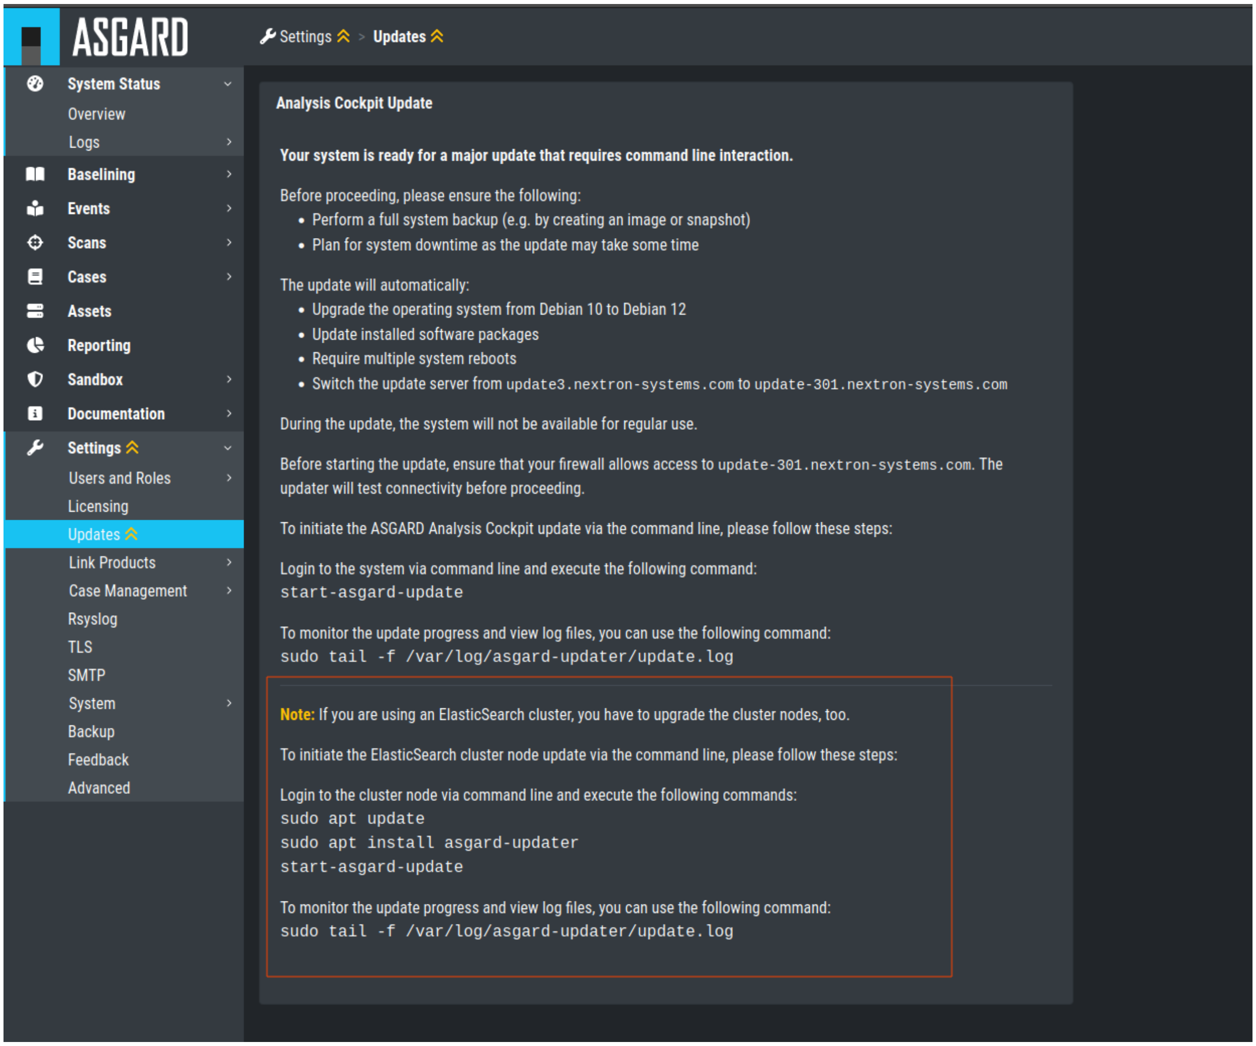

Analysis Cockpit Update

Your system is ready for a major update that requires command line interaction.

Before proceeding, please ensure the following:

- Perform a full system backup (e.g. by creating an image or snapshot)
- Plan for system downtime as the update may take some time

The update will automatically:

- Upgrade the operating system from Debian 10 to Debian 12
- Update installed software packages
- Require multiple system reboots
- Switch the update server from `update3.nextron-systems.com` to `update-301.nextron-systems.com`

During the update, the system will not be available for regular use.

Before starting the update, ensure that your firewall allows access to `update-301.nextron-systems.com`. The updater will test connectivity before proceeding.

To initiate the ASGARD Analysis Cockpit update via the command line, please follow these steps:

Login to the system via command line and execute the following command:

```
start-asgard-update
```

To monitor the update progress and view log files, you can use the following command:

```
sudo tail -f /var/log/asgard-updater/update.log
```

Note: If you are using an ElasticSearch cluster, you have to upgrade the cluster nodes, too.

To initiate the ElasticSearch cluster node update via the command line, please follow these steps:

Login to the cluster node via command line and execute the following commands:

```
sudo apt update
sudo apt install asgard-updater
start-asgard-update
```

To monitor the update progress and view log files, you can use the following command:

```
sudo tail -f /var/log/asgard-updater/update.log
```

Fig. 1: Update Section

You can now run the asgard-updater with the following command:

```
nextron@analysis:~$ start-asgard-update
```

Cluster status is "green"

You can see the status of your Elasticsearch Cluster with one of the following two methods:

Via the Web UI of your Analysis Cockpit:

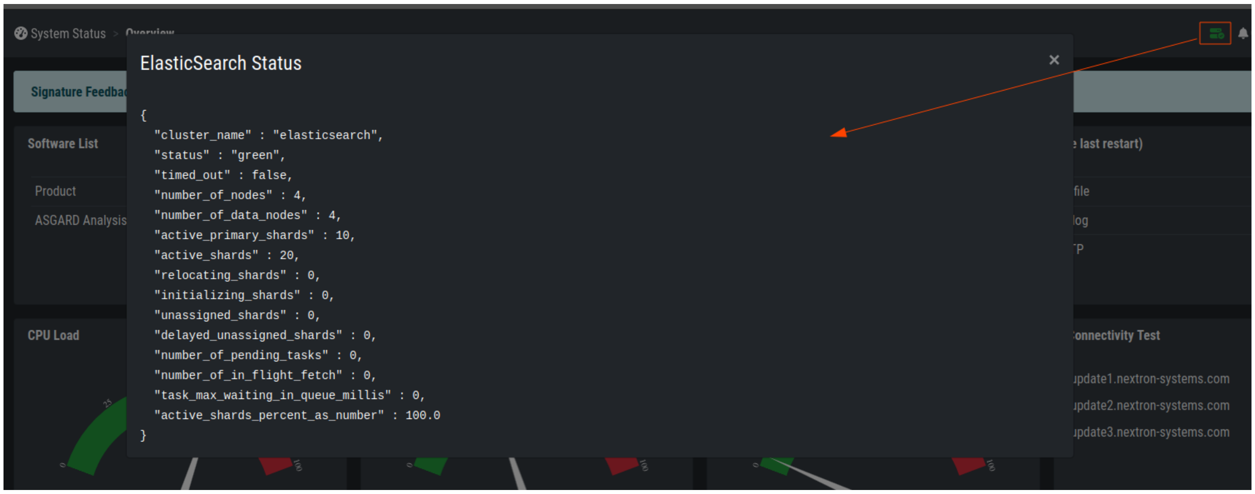

```
{
  "cluster_name" : "elasticsearch",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 4,
  "number_of_data_nodes" : 4,
  "active_primary_shards" : 10,
  "active_shards" : 20,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

Fig. 2: Elasticsearch Cluster Status

Or via SSH. To do this, connect to your Analysis Cockpit via SSH and run the following command:

```
nextron@analysis:~$ curl -s http://127.0.0.1:9200/_cluster/health | jq
{
  "cluster_name": "elasticsearch",
  "status": "green",
  "timed_out": false,
  "number_of_nodes": 4,
  "number_of_data_nodes": 4,
  "active_primary_shards": 8,
  "active_shards": 16,
  "relocating_shards": 0,
  "initializing_shards": 0,
  "unassigned_shards": 0,
  "delayed_unassigned_shards": 0,
  "number_of_pending_tasks": 0,
  "number_of_in_flight_fetch": 0,
  "task_max_waiting_in_queue_millis": 0,
  "active_shards_percent_as_number": 100
}
```

If you are unsure what your cluster nodes are, you can run the following command. Please note, the cluster marked as dim is your master node, or in our case the Analysis Cockpit.

```
nextron@analysis:~$ curl -s http://127.0.0.1:9200/_cat/nodes
172.28.30.53  23 61  0 0.03 0.10 0.04 di  - elastic-test-03
172.28.30.52  20 61  0 0.01 0.03 0.00 di  - elastic-test-02
172.28.30.225 68 97 20 1.17 1.48 1.60 dim * analysis
172.28.30.51  23 86  0 0.08 0.02 0.01 di  - elastic-test-01
```

For more information, run the following command:

```
nextron@analysis:~$ curl -s 'http://127.0.0.1:9200/_cat/nodes?format=json&filter_path=ip,
↪name' | jq
[
  {
    "ip": "172.28.123.53",
    "name": "elastic-test-03"
  },
  {
    "ip": "172.28.123.52",
    "name": "elastic-test-02"
  },
  {
    "ip": "172.28.123.225",
    "name": "analysis"
  },
  {
    "ip": "172.28.123.51",
    "name": "elastic-test-01"
  }
]
```

Stop your Analysis Cockpit - `optional`

You can optionally stop your Analysis Cockpit service to ensure no new scan logs are being synchronized from your Management Center. This will reduce the risk of losing new scan logs during the upgrade process.

To stop and disable the Analysis Cockpit service on your server, connect via SSH and run the following commands. Please keep in mind that stopping the service might take a while.

```
nextron@analysis:~$ sudo systemctl disable analysiscockpit3.service
Removed /etc/systemd/system/multi-user.target.wants/analysiscockpit3.service.
nextron@analysis:~$ sudo systemctl stop analysiscockpit3.service
```

11.1.2 Performing the upgrade

In this section we will perform the actual upgrade of the Analysis Cockpit and your cluster nodes. Please following the instructions carefully, and follow the sequence of updates according to this manual. Please do not continue if you don't have a backup/snapshot ready to restore your cluster in case of a disaster.

Cluster Node Upgrade

Hint: It is recommended that you update all your nodes at the same time. Do not update your Analysis Cockpit until all your notes are finished with the update.

If all the above tasks from the checklist are completed, you can start to upgrade your cluster nodes. Connect to your cluster nodes via SSH and run the following commands:

```
nextron@node-01:~$ sudo apt update
nextron@node-01:~$ sudo apt install asgard-updater
nextron@node-01:~$ start-asgard-update
```

This will install the asgard-updater, which will take care of the update task. The tool will upgrade your Elasticsearch version to the latest minor version available. After this, it will upgrade the OS from Debian 10 to Debian 12. Your system will restart many times during the update. If you have the feeling the upgrade is stuck at one point, you can run the following command and see the latest logs:

```
nextron@node-01:~$ sudo tail -f /var/log/asgard-updater/update.log
```

The update is finished if you are seeing the following lines:

```
nextron@node-01:~$ sudo tail -f /var/log/asgard-updater/update.log
2023-11-10T09:29:04.835115+01:00 elastic-test-01 asgard-updater[536]: Elasticsearch␣
↪service status: active
2023-11-10T09:29:04.835194+01:00 elastic-test-01 asgard-updater[536]: Upgrade finished.␣
↪Deactivating service...
2023-11-10T09:29:04.844839+01:00 elastic-test-01 asgard-updater[536]: Removed "/etc/
↪systemd/system/multi-user.target.wants/asgard-updater.service".
```

Please continue with the next step to finish the upgrade.

Analysis Cockpit Upgrade

You Elasticsearch Cluster will now be in a degraded ("red") state, since your Analysis Cockpit is still running on an older version of Elasticsearch. This is expected as long as you did not finish your Analysis Cockpit upgrade. You should see the cluster changing to a normal ("green") state throughout the upgrade of your Analysis Cockpit.

To finish your upgrade, connect to your Analysis Cockpit via SSH. We will run the following command on the command line to initiate the upgrade:

```
nextron@node-01:~$ sudo apt update
nextron@node-01:~$ sudo apt install asgard-updater
nextron@node-01:~$ start-asgard-update
```

The server running your Analysis Cockpit will now restart multiple times. It is important to not interrupt the upgrade process and let the server do all the tasks. You can however see if any errors occurred during the upgrade or just observe at what stage the upgrade is.

Run the following command to see the status of your upgrade:

```
nextron@analysis:~$ sudo tail -f /var/log/asgard-updater/update.log
```

There is a chance that you see the following log lines repeatedly in the output:

```
nextron@analysis:~$ sudo tail -f /var/log/asgard-updater/update.log
Nov 14 12:30:17 analysis asgard-updater[2403]: 2023-11-14T12:30:17+01:00 STARTING /usr/share/asgard-updater/bin/step05.sh
Nov 14 12:30:17 analysis asgard-updater[2403]: Checking for Elasticsearch Cluster Nodes...
Nov 14 12:30:17 analysis asgard-updater[2403]: Elasticsearch service status: active
Nov 14 12:30:17 analysis asgard-updater[2403]: Elasticsearch cluster is not healthy (status: red).
Nov 14 12:30:17 analysis asgard-updater[2403]: Elasticsearch cluster setup is enabled but no nodes are connected.
Nov 14 12:30:17 analysis asgard-updater[2403]: 2023-11-14T12:30:17+01:00 FINISHED /usr/share/asgard-updater/bin/step05.sh RC=54
```

If this is the case, your cluster nodes might not be fully online yet. The updater tries this check every minute to make sure the cluster is fully online and healthy, before continuing with the next steps. Even if it looks like the updater is stuck, you have to give it some time and wait for it to continue by itself.

The update is finished if you are seeing the following lines:

```
nextron@node-01:~$ sudo tail -f /var/log/asgard-updater/update.log
2023-11-10T09:29:04.835115+01:00 analysis asgard-updater[536]: Elasticsearch service status: active
2023-11-10T09:29:04.835194+01:00 analysis asgard-updater[536]: Upgrade finished. Deactivating service...
2023-11-10T09:29:04.844839+01:00 analysis asgard-updater[536]: Removed "/etc/systemd/system/multi-user.target.wants/asgard-updater.service".
```

Your cluster status should change back to a "green" status once all the updates of your Analysis Cockpit are installed. You can see the status in your Analysis Cockpit in the top right corner:

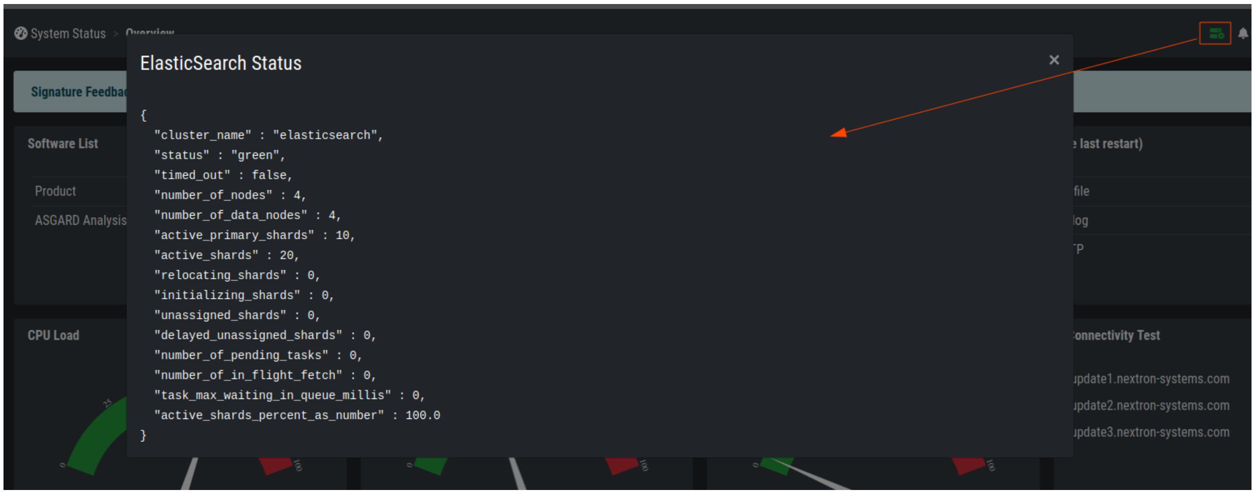

ElasticSearch Status

```
{
  "cluster_name" : "elasticsearch",
  "status" : "green",
  "timed_out" : false,
  "number_of_nodes" : 4,
  "number_of_data_nodes" : 4,
  "active_primary_shards" : 10,
  "active_shards" : 20,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 0,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

Fig. 3: Elasticsearch Cluster Status

Your upgrade is now finished, and you can use your Analysis Cockpit with the newest version.

11.2 Standalone Upgrade

This chapter guides you through the upgrade process of your Analysis Cockpit version 3.10.1 to version 4.x.

It is important to follow the steps carefully. We advise you to create a snapshot of the Analysis Cockpit itself before starting your upgrade.

11.2.1 Preparation

To prepare for your upgrade, we compiled a list of tasks you should follow:

Task	Description
Snapshot of your Analysis Cockpit	For disaster recovery
Analysis Cockpit running version 3.10.1	Prerequisite for the Major Upgrade
Connection to our new update servers	New update server infrastructure

For details regarding some of the above tasks, see the next section in this manual.

With the new version of your Analysis Cockpit, we also made changes to our update servers. Please make sure that all your components can reach the following servers:

Server	Port	Description
update3.nextron-systems.com	tcp/443	Old update server
update-301.nextron-systems.com	tcp/443	New update Server

The old update server is needed to fetch the updater and other prerequisites. The new update server is needed to upgrade your servers to Debian 12 and also to install any new packages, which are needed for your Analysis Cockpit v4.

You can find the corresponding IP-Addresses to the above FQDNs here: https://www.nextron-systems.com/resources/hosts/.

Analysis Cockpit running version 3.10.1

To check if your Analysis Cockpit is running on the correct version you can navigate to `Settings` and `Updates`. The page should looks like this:

11.2.2 Performing the upgrade

In this section we will perform the actual upgrade of the Analysis Cockpit.

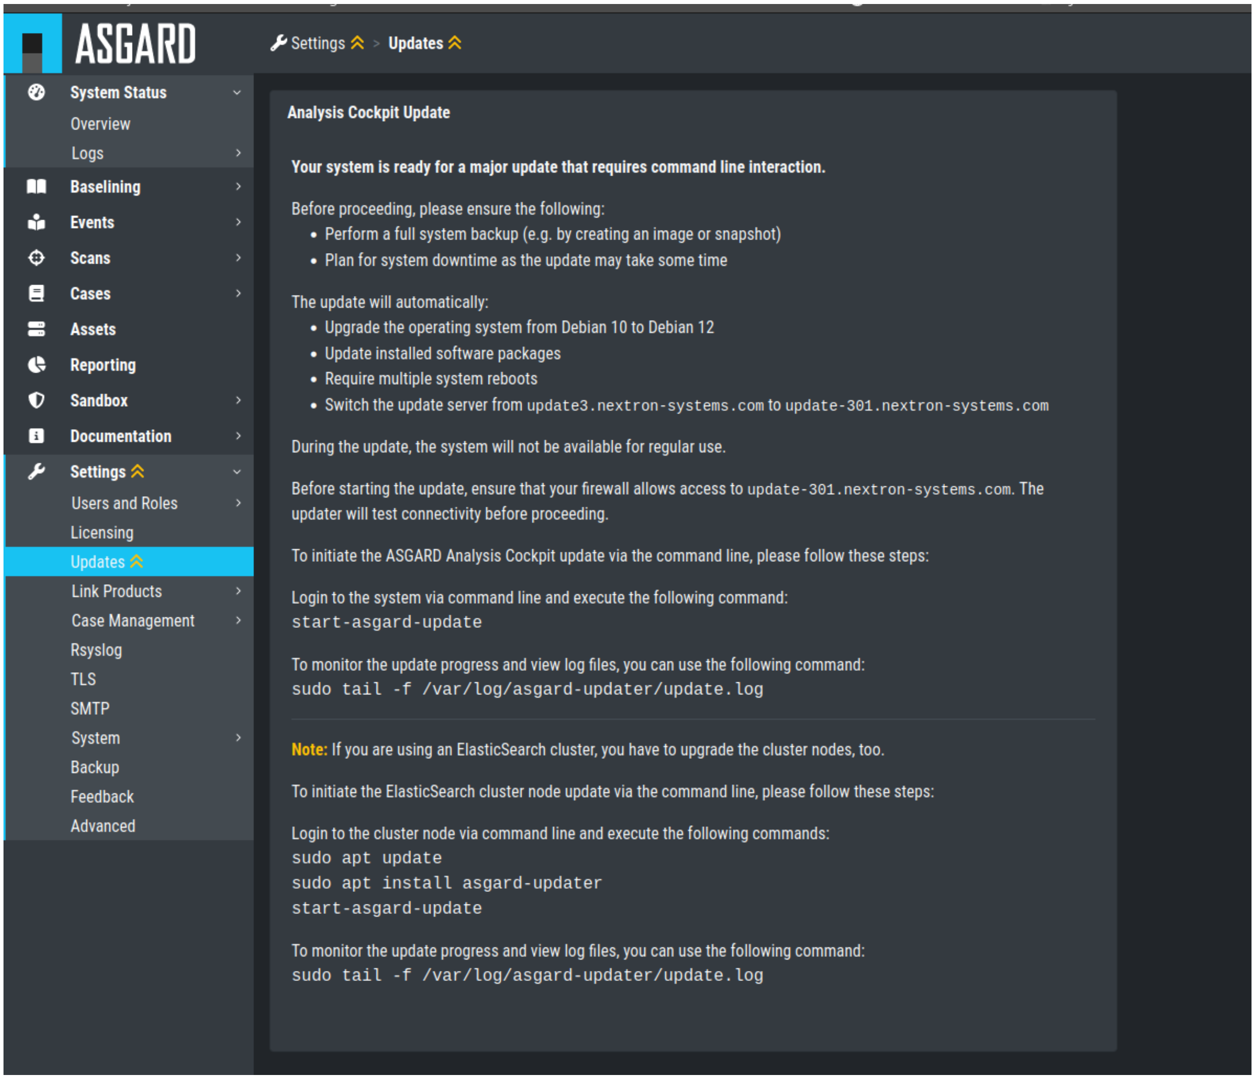

Analysis Cockpit Update

Your system is ready for a major update that requires command line interaction.

Before proceeding, please ensure the following:

- Perform a full system backup (e.g. by creating an image or snapshot)
- Plan for system downtime as the update may take some time

The update will automatically:

- Upgrade the operating system from Debian 10 to Debian 12
- Update installed software packages
- Require multiple system reboots
- Switch the update server from `update3.nextron-systems.com` to `update-301.nextron-systems.com`

During the update, the system will not be available for regular use.

Before starting the update, ensure that your firewall allows access to `update-301.nextron-systems.com`. The updater will test connectivity before proceeding.

To initiate the ASGARD Analysis Cockpit update via the command line, please follow these steps:

Login to the system via command line and execute the following command:

```
start-asgard-update
```

To monitor the update progress and view log files, you can use the following command:

```
sudo tail -f /var/log/asgard-updater/update.log
```

Note: If you are using an ElasticSearch cluster, you have to upgrade the cluster nodes, too.

To initiate the ElasticSearch cluster node update via the command line, please follow these steps:

Login to the cluster node via command line and execute the following commands:

```
sudo apt update
sudo apt install asgard-updater
start-asgard-update
```

To monitor the update progress and view log files, you can use the following command:

```
sudo tail -f /var/log/asgard-updater/update.log
```

Fig. 4: Update Section

Analysis Cockpit Upgrade

To start your upgrade, connect to your Analysis Cockpit via SSH. We will utilize `asgard-updater` to perform the upgrade. First we need to check if a newer version of the `asgard-updater` is available. If you get the highlighted output, you have already the newest version installed (the version might differ from the output here):

```
nextron@analysis:~$ sudo apt update
nextron@analysis:~$ sudo apt install asgard-updater
Reading package lists... Done
Building dependency tree
Reading state information... Done
asgard-updater is already the newest version (1.0.17).
0 upgraded, 0 newly installed, 0 to remove and 18 not upgraded.
```

You can now run the `asgard-updater` with the following command:

```
nextron@analysis:~$ start-asgard-update
```

The server running your Analysis Cockpit will now restart multiple times. It is important to not interrupt the upgrade process and let the server do all the tasks. You can, however, see if any errors occurred during the upgrade or just observe at what stage the upgrade is.

Run the following command to see the status of your upgrade:

```
nextron@analysis:~$ sudo tail -f /var/log/asgard-updater/update.log
```

The update is finished if you are seeing the following lines:

```
nextron@analysis:~$ sudo tail -f /var/log/asgard-updater/update.log
2023-11-10T09:29:04.835115+01:00 analysis asgard-updater[536]: Elasticsearch service␣
↪status: active
2023-11-10T09:29:04.835194+01:00 analysis asgard-updater[536]: Upgrade finished.␣
↪Deactivating service...
2023-11-10T09:29:04.844839+01:00 analysis asgard-updater[536]: Removed "/etc/systemd/
↪system/multi-user.target.wants/asgard-updater.service".
```

Your upgrade is now finished, and you can use your Analysis Cockpit with the newest version.

CHAPTER

TWELVE

GLOSSARY

This chapter contains explanations about wording used in the Analysis Cockpit and our ASGARD products.

12.1 Baselining

The `Baselining` section is meant to assign all unassigned events to cases, so that newly incoming events are immediately assigns to existing cases and only differences to the last scan have to be reviewed. (new malware, new signatures)

12.1.1 Auto Baselining

The Auto Baselining feature can be used to assign events to new and automatically generated cases. Auto Casing makes use of the so-called Auto Case ID, which is the same for all events of a certain type (see Glossary > Cases > Auto Case ID for details).

Auto Baselining is typically used to quickly reduce the remaining events in the Baselining section. It reduced the burden to manually group together events of a similar type.

Automatically cased events can then be reviewed in the `Cases` section.

Auto Baselining uses a threshold that defines the minimum number of events required in each of these automatically generated cases. A threshold of 10 instructs the process to create only cases for groups of at least 10 similar events. Obviously, we do not recommend using threshold of 1, but everything higher than 1 can be reasonable.

The best practice is to start the Auto Baselining process with a relatively high threshold (e.g., 10) and then subsequently perform iterations with lower thresholds.

12.1.2 Optimization

The Optimization is used to assign unassigned events to cases based on their filters.

Usually, an analyst selects events and creates a case with these events. During the case creation a set of filters gets generated to automatically assign newly incoming events to this case.

Often, characteristics of existing older events also align with the criteria described in the filters of the new case. However, they do not get assigned to that case, because they're already in the database and the case assignment only happens when new events arrive.

Optimization is used to assign unassigned events to existing cases.

12.2 Cases

12.2.1 Auto Case ID (formerly Group ID)

The Auto Case ID is an automatically generated ID for a group of events.

This group has been formed automatically by the use filters generated from predefined filter templates (see Glossary > Invisible > Filter Templates).

Auto Case IDs identify a groupable set of events.

The Auto Case IDs are used in case creation as the default method to assign new events to a case. You can review the Auto Case IDs used to assign new events to a case in the Tab `Grouping Criteria` of each case.

They are also used in "Auto Casing" and "Optimization". (see *Baselining*)

12.2.2 Dynamic Auto Case ID

The Dynamic Auto Case IDs are generated for events that don't have a corresponding filter template defined. This is often the case for very rare event types or events of a low level (Info, Notice).

You can think of them as a fallback in form of a less efficient grouping method.

They are usually created of all fields of an event except the ones that are highly specific, like the HOSTNAME field.

If your case uses many Dynamic Auto Case IDs, distinguishable by the leading uppercase "D", then automatic event assignment is nearly impossible. In these cases, you should rather use a string conditions to group events into that case and assign new incoming events automatically.

12.2.3 Filter Priority

It is possible that two or more cases include filters that would assign a new incoming event to them. This often happens when one case uses Auto Case ID (see Glossary section) and another case uses a string selection (String Condition).

The filter priority can be used to prioritize the filters of a case so that newly incoming events go to that prioritized case and not the one with the default priority.

You could also create fallback cases, e.g., for all events of the level "Notice" and intentionally reduce the priority of this case so that other cases that select events of that level always come first.

The full prioritization process looks like:

- Case with higher priority takes the event
- If both cases have the same priority, the case with the higher "Type" takes the event (Incident > Noteworthy)
- If both cases have the same priority and the same type, the case with the smaller case ID (older case) takes the events

12.2.4 External ID

This field is optional and can be used to refer to an ID that you use in a different system, like a ticket management system or an incident response platform.

12.2.5 Difference between Summary and Assessment

The field `Summary` is meant to include the elements (fields) of events that are used as characteristics to perform an assessment. You can think of it as the values that you as an analyst want to highlight for other analysts that review that specific case. It's often a special file name and location or a process name and YARA rule match on that process.

You can use the "Auto Summary" feature to get an auto-recommended content for this field.

The field `Assessment` is the one that requires the most effort. It contains the findings of the analyst's review.

12.2.6 Case Types

The following table describes the cases types taxonomy used in Analysis Cockpit.

Type	Description
Incident	Incident cases report a clear threat, indicated by a hard match and verified by research of an analyst. Analysts create incident cases to indicate the highest possible certainty and risk. Incident cases are also characterized by the fact that they do not need to be verified by someone else. They either indicate malware, a threat group or penetration testing activity and should trigger immediate response.
Suspicious	Suspicious cases are based on significant indicators that require a review by someone within the organization or more evidence to come to a final conclusion. Often, file samples or process memory dumps are required to verify/falsify a verdict. Cases of this type usually trigger evidence collection or review actions.
Noteworthy	Noteworthy cases are based on soft indicators or elements that should be reviewed whenever there is time to do that. They include all kinds of events that cannot be dismissed as false positives or anomalies but are likely uncritical. Noteworthy cases don't trigger an immediate response but should be reviewed whenever there is time to do that.
Vulnerability	Vulnerability cases contain detected software or configuration weaknesses that compromise system integrity. The reported vulnerabilities often include easy to exploit weaknesses that are frequently used by threat groups to execute code remotely, gain access or escalate privileges on affected systems. Cases classified as Vulnerability are typically integrated into a vulnerability management process as an additional input channel.
Legitimate Anomaly	Legitimate Anomaly cases contain events that are related to legitimate elements that are suspicious, but an ordinary finding in the context of the analyzed organization.The reason for an anomaly is not a malfunction of the scanner but a peculiarity within the analyzed environment. Legitimate Anomalies don't trigger any further activity.
False Positive	False Positive cases contain events that indicate suspicious or malicious activity, but the review revealed that it is actually legitimate software or other elements. The only reason for a false positive is a scanner malfunction or signatures that falsely report a threat (see section *Difference between False Positive and Legitimate Anomaly* for details). A false positive usually triggers a review by Nextron Systems and a signature adjustment.
Unknown	The default state of newly created cases.

12.2.7 Difference between False Positive and Legitimate Anomaly

We use "False Positive" and "Legitimate Anomaly" to distinguish between situations in which the scanner (THOR) made an error and situations in which a customer environment contains suspicious or malicious elements that are known.

E.g., a Winrar used by admins as `r.exe` in `C:\users\public` for software rollout purposes is not considered a "False Positive" but a "Legitimate Anomaly". It is a finding which doesn't have to be fixed in THOR's signature set but is simply a specific situation in the analyzed environment.

Matches that are clearly an error in THOR signatures should be classified as "False Positive".

Examples for "Legitimate Anomalies":

- Procdump.exe findings
- Suspicious RUN Key entries that use customer software
- Custom software that uses suspicious folders, e.g. `C:\Users\Public`, `%AppData%`
- Process memory match with a "ReflectiveLoader" YARA rule on a third party EDR agent process

Examples for "False Positives":

- YARA rule match on Bloomberg or SAP software
- Filename IOC match `w64.exe` on a Perl for Windows build tool
- YARA rule match with "Putty_Anomaly" on a legitimate and signed `putty.exe`

Another good example is one of the many anomaly signatures that triggers on an XORed MS-DOS Stub. A match with such a signature only qualifies as false positives when there is no XORed MS-DOS stub in that file and not when it turns out to be a legitimate file. The signature detects what it is designed to detect.

A signature with a rule named `MAL_Xrat_Mar21_1` that triggers on a legitimate and signed executable, however, is a false positive.

12.3 Invisible (Backend)

12.3.1 Filter Templates

The Analysis Cockpit uses so-called filter templates that describe which fields in which event types are specific enough to be used in a filter that can be used to automatically group events.

These groups can be identified by a common so-called "Auto Case ID" (formerly Group ID). See the respective entry in this Glossary.

The filter templates are static and predefined.

E.g., a typical filter template states that for events in the Module `Filescan`, the fields **FILE** and **SHA1** are sufficiently specific to group events based on equal values in these two fields.

CHAPTER

THIRTEEN

CHANGELOG

This chapter contains a list of all changes. Those changes are only related to the Analysis Cockpit version 4.

13.1 Analysis Cockpit v4

13.1.1 Analysis Cockpit 4.0.13

Release Date
Wed, 6 Mar 2024 07:32:00 +0200

Type	Description
Bugfix	Added second level of 'disk watermark' to prevent ElasticSearch from not working properly when disk space is low

13.1.2 Analysis Cockpit 4.0.12

Release Date
Tue, 20 Feb 2024 16:20:00 +0200

Type	Description
Bugfix	Fixed import of THOR logfiles with very long lines
Bugfix	Fixed error on creating new cases based on condition and time range
Bugfix	Fixed non-working ntp service restart after ntp configuration changes

This chapter contains all the changes of the ASGARD Analysis Cockpit.

13.1.3 Analysis Cockpit 4.0.10

Release Date
Tue, 16 Jan 2024 10:26:00 +0100

Type	Description
Feature	Automatic index rollover
Feature	Cluster configuration
Change	Using timesyncd instead of ntpd. We recommend to check your NTP settings in the UI after upgrade
Change	Improved sync performance between Management Center and Analysis Cockpit

CHAPTER FOURTEEN

INDEX

- genindex

INDEX