

NeMo / DDoS Application

User manual

May 2020

Document Information			
Non-disclosure notice	GEANT-Project / European NRENs		
Authors	Christian Keil (CKE) Klaus-Peter Kossakowski (KPK) Marcus Weseloh (MWE) Antonio Liu (LIU) Michel Gerdes (MGE) Andrej Zieger (AZI)		
Filename	NeMo-User_Manual-20200513.pdf		
Last editing	Wednesday 13 th May, 2020		
Page count	112		
Document status			
	Status	Date	Editor
Documentation to software version 4.0.3	V1.0	16.07.2013	CKE
Supplement alarm-email and report	V1.1	17.07.2013	CKE
Cleared inaccuracies	V1.1.1	19.07.2013	CKE
Update to software version 5.0.0	V1.2	15.04.2015	LIU
Update to software version 5.5.4	V1.3	30.09.2015	LIU
Minor corrections and additions to the most recent version	V1.3.1	21.12.2015	LIU
Update to software version 6.0.13	V1.4	26.04.2017	MGE
Update to software version 7.0.1	V1.5	08.11.2017	MGE
Translation to English, Update to software version 7.4.5, Adaption for international version	V1.6	14.08.2019	MGE
Cleanup for external release	V1.7	13.05.2020	AZI

Contents

1	Introduction and Objectives	7
1.1	Classification and objective of this document within the context of NeMo	7
1.2	Target group of the user manual	7
1.3	Structure of the user manual	7
1.4	Changes in the user manual	7
2	Requirements for accessing NeMo	9
2.1	Accessing the email interface	9
2.2	Accessing the web interface	9
3	NeMo email interface	10
3.1	Identification, status and classification of alarm messages	10
3.1.1	State and state changes of an alarm	11
3.1.2	Severity of an alarm	11
3.1.3	Categorization of alarms	12
3.1.4	Aggregation of alarms	12
3.2	Subject line of emails	12
3.3	Structure of an alarm message	13
3.3.1	Trigger of alarm messages	16
3.3.2	Temporal development of alarms	16
3.3.3	Changes of alarm messages over time	17
3.3.4	Meta alarm additions	18
4	NeMo web interface	20
4.1	General structure of the web interface	20
4.2	Overview of the most important alarm messages	22
4.3	Viewing all alarms	24
4.3.1	Manually creating an alarm	24
4.3.2	Filtering the list of alarms	25
4.3.3	Sorting the table	27
4.3.4	Visualization of closed alarms	27
4.4	Viewing all objects	28
4.5	Mitigations	31
4.6	Topology map of the Network	33
4.6.1	Adjustment of the timeframe to be visualized	34
4.6.2	Changing the appearance of the map	36
4.6.3	Adjustments of the visualization	37
4.6.4	Sparklines for selected objects	39
4.7	Visual Explorer	41
4.7.1	Adjustments of the visualization and bookmarks	42
4.7.2	Manipulation of the plots	43
4.7.3	Adding plots	49
4.8	Sparklines	53
4.8.1	Constructing a sparkline	53
4.8.2	Time frame selection	54
4.8.3	Controls for the selection of objects and indicators	54

4.8.4	Displaying sparklines	56
4.9	List of all detectors	57
4.9.1	Detailed view of a detector	57
4.10	Settings	59
4.10.1	Changing the indicator color	59
4.11	System information	61
4.11.1	Load	62
4.11.2	Status of necessary processes	62
4.11.3	Timeliness of the Networkmodel	62
4.11.4	Most recent topology update	62
4.11.5	Data size	63
4.11.6	State of the server	63
4.12	Detailed view of an object	64
4.12.1	Category and classification in topology	64
4.12.2	Traffic Analysis	65
4.12.3	Configuration of detectors	85
4.13	Details of an Alarm	86
4.13.1	Alarm summary	86
4.13.2	List of associated events	92
4.13.3	Analysis view of an alarm	95
4.13.4	Tracing back an alarm	100
4.13.5	Mitigation of an alarm	102
4.14	Detailed view of a Mitigation	103
4.14.1	Mitigation Details	103
4.14.2	Create a mitigation	107
4.14.3	Selecting an authorization	107
4.14.4	Target Filter	108
4.14.5	Statistics for an active or formerly active mitigation	109
5	References and Glossary	111
5.1	Glossary	112

List of Figures

1	Status transitions of alarm messages	11
2	Schematic structure of the alarm message from Listing 1	14
3	Development of an alarm based on time windows	16
4	One page of the web application	20
5	Content area if no alarms match the selected filters (part view)	22
6	View of all non-suppressed alarms	24
7	Form to create an alarm manually	25
8	Visualization of a suppressed alarm in the alarm list	26
9	Viewing all objects	28
10	Overview on mitigations (part view)	31
11	Topology of the Network	33
12	Zooming in shows grey pins indicating the proper location of routers	34
13	Time frame controls	35
14	Controls for moving and zooming the extract of the map	36
15	Control for selecting the map and objects of the Network	36
16	Visualization of the topology map using OpenStreetMap data	37
17	Controls for adjusting the objects in the topology map	38
18	Visualization of SNMP indicators in the topology map	38
19	Visualization of object details in the topology map	40
20	Visual Explorer	41
21	Aggregation, marker and bookmarks of the Visual Explorer	43
22	Stats Plot in the Visual Explorer	44
23	Selection of an object for a plot	44
24	Display of time and indicator values at mouse pointer	45
25	Selection of a time frame in a stats plot	45
26	Hiding and showing the options panel	46
27	The Add Metric link is shown once the mouse pointer is on the options menu	46
28	Adding colorful metrics	47
29	Control settings for plots	48
30	Comparison of enabling (top) and disabling (bottom) a true zero point	49
31	Sparkline view (part view)	53
32	Constructing a sparkline	54
33	Comparison of sparklines with different resolution	54
34	Sparklines in table based layout	56
35	A listing of all activated detectors	57
36	Detailed view of a detector	58
37	Detailed view of a detector without standard configuration (part view)	58
38	The unfolded menu item for the Preferences menu	59
39	Changing the coloring scheme for indicators (part view)	59
40	Page of system information	61
41	Detailed view of an object (part view)	64
42	Analysis view of an object	66
43	Start time of a NetFlow analysis	66
44	Start time is before the specified time period	67
45	Selection of a router for a NetFlow analysis	68

46	Available types of filters using the visual mode	69
47	A simple NetFlow filter in visual mode	70
48	A NetFlow filter using Groups of filters	71
49	A more complex NetFlow filter expression	71
50	Plot of a NetFlow filter	72
51	History of a filter expression	72
52	Evaluation of NetFlow data	73
53	Evaluation of top 10 protocols in the flow data	74
54	Time period of analysis is highlighted in the graph	75
55	Analyzing the top ten hosts (most communication)	76
56	Evaluation of NetFlow data using Parallel Coordinates	77
57	Coloring on linear scale	78
58	Coloring on logarithmic scale	79
59	Intensity of the lines with value .5	79
60	Eye-catching traffic pattern in parallel coordinates	80
61	Marking an interesting traffic patterns using value ranges	80
62	An isolated traffic pattern in parallel coordinates	81
63	Visualization of quartiles based on packets in parallel coordinates	81
64	Visualization of raw NetFlow data	82
65	Aggregation of NetFlow raw data using the source IP address	83
66	A selection of sparklines for the selected object	84
67	A Visual Explorer for an object (part view)	84
68	Configuration view of a detector	85
69	Summary of an alarm (part view)	86
70	Dialog for “Send to TTS”	87
71	List of alarm for manual merge	88
72	Creating a new Meta Alert	88
73	Adding an alarm to a Meta Alert	89
74	Details of an alarm	91
75	Trend of an indicator for an alarm	92
76	List of events associated with an alarm (part view)	92
77	Analysis view of an alarm	96
78	Controls to set time period are also links	96
79	Controls when live view is activated	97
80	Dialog to create a report	97
81	Dialog to print the report	98
82	Automatically generated NetFlow filter for an alarm	99
83	Selecting a new target for the analysis	99
84	Controls to determine the paths of potential attack traffic for an alarm	100
85	Finished trace back of a potential attack (part view)	101
86	The detailed view of a mitigation (the register Statistics is not active and therefore not accessible since the mitigation has not been deployed successfully yet) .	103
87	The detailed view of a mitigation for which the Statistics register is active . . .	104
88	Dialog to select the type of authorization	107
89	An analysis view for traffic that still passes the mitigation	109
90	Statistics register	110

1 Introduction and Objectives

Any operation of computer networks requires specific evaluations of network traffic data and adaption of appropriate measures to mitigate attacks on the infrastructure.

1.1 Classification and objective of this document within the context of NeMo

The NeMo DDoS application provides users at corresponding NOC an email interface and a web interface incorporating extensive tools for network analysis. These interfaces are used to disseminate alarm notifications and provide access to information about network traffic patterns detected as potential DDoS attacks. Specific technical information may be made available.

This user manual supports users by familiarizing with both interfaces and their functions, and providing a guide to look up the meaning of certain identifiers or categories.

1.2 Target group of the user manual

This manual is intended for users and decision makers who use any of the interfaces of the DDoS application of NeMo. It provides the level of knowledge necessary for effective use, but does not explain the technical concepts that are used, for example, for detection of network traffic patterns.

Reading the user manual requires a certain degree of familiarity with DDoS attacks and relevant characteristics, as listed in [1]. However, there is no need to be familiar with the selected interfaces and data formats chosen for the implementation as well as the algorithms, e.g. for detection of DDoS attacks. These are listed in [2] and [3].

1.3 Structure of the user manual

This user manual is structured into the following chapters:

Section 2 Requirements for access to NeMo The requirements for access to NeMo and its DDoS application must be fulfilled. These are described in this section.

Section 3 NeMo alert messages via email Various messages from the DDoS application are sent to users by email. This section explains the formats, the terms and their values used in these emails.

Section 4: NeMo Web interface While time-critical messages are sent to users by email, the same information is provided at the web interface. In addition, various lists and further functions are provided for analysis, further investigation and mitigation. The options of the web interface are explained together with the terms used and their possible values.

1.4 Changes in the user manual

The following changes and additions have been added in this version of the user manual:

- Added sections for functions which were not entirely documented.
- Added a description for [Target Filter](#) for mitigations.

- Removed DFN specifics for international edition

2 Requirements for accessing NeMo

In this section the requirements are described that must be fulfilled to be granted access to either email interface or web interface of NeMo.

2.1 Accessing the email interface

All notifications and messages from the NeMo application will be send to the list nemo@example.org. All email addresses being subscribed to this list will automatically receive those emails.

Users who want to subscribe to this list shall contact [Registration Contact <email>](#).

As there is no encryption enabled for this list, no further requirements must be fulfilled.

2.2 Accessing the web interface

The web interface is available at <https://nemo.example.org>.

To gain access to the web interface a valid user certificate must be installed in the web browser. The certificate must be signed by one of the following Certificate Authorities:

- To be configured
- [Name of CA](#)
Subject: [issuer DN of the CA](#)

A further access restriction—based on user certificates—can be implemented upon request, e.g. to grant only certain individuals access to a specific function.

The security of the transmission (confidentiality and integrity) is guaranteed likewise by a server certificate issued by the configured PKI.

3 NeMo email interface

NeMo alert messages, mostly referred to as “alarm” or “alert”, are automatically distributed to all users subscribed to the mailing list (as stated in Section 2.1). Each message contains detailed information about the circumstances leading to the event, as exemplified in Listing 1.

```
Subject: [NeMo] #22241 NEW INFO - Infrastructure: R1 - "High SYN/ACK ratio:
        36k ACK Packets/s, 550 SYN Packets/s"
From: nemo@example.org
To: nemo@example.org
Date: Tue, 24 May 2011 05:01:49 +0000 (UTC)

Observed high values of the ratio of SYN packets to ACK packets with
        36k ACK Packets/s, 550 SYN Packets/s on router R1.

Opened with severity Info.

Alert ID: 22241
Status: Open
Severity: Info
Tags: Infrastructure

Start Time: 2011-05-24 05:01:36
End Time: ongoing
Duration: 00:00:00

First Event Seen: 2011-05-24 04:56:00
Last Event Seen: 2011-05-24 05:00:00
Event Count: 4

Trigger: High SYN / ACK Ratio (ID 9)

Alert Description:
High ratio of SYN packets to ACK packets.

Affected Objects:
      Type      Name      Event Count
-----
Router    R1          4

Further Details:
https://nemo.example.org/nemo/alerts/details/22241/

All times expressed in UTC.
```

Listing 1: Example of an alarm notification via email

The following sections discuss the information contained in alarm messages, for which types of events these messages are created and how to interpret data within a message. We will start with explaining categories and terms followed by explaining the structure and each line.

3.1 Identification, status and classification of alarm messages

As soon as a new event is detected, that justifies an alarm message, the system assigns an [alert ID](#). This identifier is a consecutive number, allowing the identification of individual messages

and is therefore used for all subsequent emails generated with reference to the detected alarm situation. The identifier is also used consequently throughout the web interface.

3.1.1 State and state changes of an alarm

If a new alarm message is generated (Transition “[new](#)”), the status value is set to “[open](#)”. Its status may be changed depending on further development of circumstances, as shown in Figure 1.




Figure 1: Status transitions of alarm messages

Whenever new information is detected by NeMo that changes the rating of a detected alarm situation (see below in Section 3.1.2), the rating is updated accordingly (Transition “[upgrade](#)”). The status remains “[open](#)”. If no additional information is received for a certain period of time, the alarm situation will be stopped (Transition “[end](#)”) and its monitoring ceases. The status is set to “[closed](#)”. In rare cases, especially if the alarm situation is interrupted for technical reasons, e.g. after restarting the alarm server, the status may be set to “[aborted](#)”.

3.1.2 Severity of an alarm

The status does not indicate the severity of an incident, therefore a further classification is necessary. The severity may be any of these levels:

Informative (Info) These reports may indicate possible attacks. The alarms have just been triggered and the events have not yet affected network traffic negatively.

Warning These reports inform about anomalies that already have a negative impact on network traffic and should be investigated more closely.

Critical These reports indicate anomalies which are highly likely the result of an attack on the monitored network or an attack on a system that has been connected to the network.

All current alarm situations are re-assessed regularly, especially when new information is available or depending on the duration of a certain state; hence, the severity value may increase over time, whereas an automated decrease of the severity value is neither intended nor implemented.

3.1.3 Categorization of alarms

Each alarm may be categorized by assigning multiple alarm categories to allow further differentiation of alarms. These are also called “tags” within NeMo. Categories must be configured by the administration staff when setting up and configuring NeMo.

Tags are then associated automatically depending on the alarm procedure that triggered the alarm (To be configured by Administration staff as well). These tags are used consistently within the email and web interfaces. Tags are listed only if any were assigned.

3.1.4 Aggregation of alarms

Simple alarms summarize observations on a single network object, whereas “meta alarms” aggregate simple alarms and allow to combine various observations across several objects and conduct anticipated analysis tasks while raising the alarm. Meta alarms are categorized and identified as simple alarms and are assigned the same set of status. Meta alarms are automatically assigned to all tags of any child alarm, whereas additional categories may also be assigned to the meta alarm. The relationship model between alarms and meta alarms is **n to m**:

- each meta alarm includes at least one simple alarm and
- each simple alarm may belong to several meta alarms.

When a simple alarm is assigned to a meta alarm, notifications for the simple alarm are suppressed thereafter. The simple alarm is still monitored by NeMo, all state changes are applied accordingly and the simple alarm may be assigned to other meta alarms. To reduce the number of notifications sent by email, only state changes of meta alarms are sent, therefore no notifications are sent out for state changes of simple alarms while they are assigned to meta alarms.

In this document the term **alarm** refers to both simple alarms and meta alarms, as long as a restriction is not obvious from the context. Where the distinction is important, the respective full terms “simple alarm” or “meta alarm” shall be used.

3.2 Subject line of emails

The subject of an alarm email summarizes essential information, allowing users to filter and classify incoming emails without detailed analysis of the email body. An example is a good way to study the structure of the subject:

```
[NeMo] #22241 NEW INFO - Infrastructure: R1 - "High SYN/ACK ratio:36k ACK  
Packets/s, 550 SYN Packets/s"
```

1. Each subject line starts with “[NeMo]” to clarify the affiliation of the email¹.
2. This is followed by a hash and the alarm message’s alert ID, in this case “#22241”.
3. Next is a short version of the status change:

NEW for new alarms

UPD for an update and

¹This may also be configured by the administrative staff

END for closed or aborted alarms.

4. The severity is listed in a short form as “**INFO**”, “**WARN**” or “**CRIT**”, respectively.
5. If the alarm has been tagged or belongs to one or more categories this is indicated next, e.g. “**Infrastructure**” is this example. If no tag has been assigned, this is omitted in the subject line.
6. The name of the affected objects in the Network, in this example “**R1**”.
7. Finally, a short description in quotation marks. The description in the example “**High SYN/ACK ratio: 36k ACK Packets/s, 550 SYN Packets/s**” means, that a high ratio of packets with SYN flag set to packets with ACK flag set was observed on the previously listed router “**R1**”. In addition, the last observed values of the indicators triggering the alarm are listed.

3.3 Structure of an alarm message

A report is basically structured in eight segments, whereas the email’s header data is to be considered a ninth segment, see Figure 2. Only reports about meta-alarms contain another segment which lists the aggregated simple alarms. We will discuss this after explaining the structure and content of simple alarms’ messages and each segment.

1. Subject: [NeMo] #22241 NEW INFO - Infrastructure: R1 - "High SYN/ACK ratio: 36k ACK Packets/s, 550 SYN Packets/s"
 From: nemo@example.org
 To: nemo@example.org
 Date: Tue, 24 May 2011 05:01:49 +0000 (UTC)
2. Observed high values of the ratio of SYN packets to ACK packets with 36k ACK Packets/s, 550 SYN Packets/s on router R1.
 Opened with severity Info.
3. Alert ID: 22241
 Status: Open
 Severity: InfoTags: Infrastructure
4. Start Time: 2011-05-24 05:01:36
 End Time: ongoing
 Duration: 00:00:00
5. First Event Seen: 2011-05-24 04:56:00
 Last Event Seen: 2011-05-24 05:00:00
 Event count: 4
6. Trigger: High SYN / ACK ratio (ID 9)
7. Alert Description:
 High ratio of SYN packets to ACK packets.
8. Affected Objects:

Type	Name	Event Count
Router	R1	4
9. Further Details:
<https://nemo.example.org/nemo/alerts/details/22241/>
 All times expressed in UTC.

Figure 2: Schematic structure of the alarm message from Listing 1

Segment 1 In addition to the subject line already explained in Section 3.2 Segment 1 contains the sender address and the time stamps when the report was generated.

Segment 2 Segment 2 contains a long description as reason for the alarm status change. The new state and severity are listed as well. There are four cases to distinguish:

1. So, for a freshly generated alarm, Segment 2 might read as follows:

```
Observed high values of the ratio of SYN packets to ACK
packets with 36k ACK Packets/s, 550 SYN Packets/s on
router R1.

Opened with severity Info.
```

2. An update is marked with a long description similar to this:

```
Observed more high values of the ratio of SYN packets to ACK
packets on router R1.
```

```
Upgraded to severity Warning.
```

3. And when an alarm is closed, a long description similar to the following content appears.

```
The ratio of SYN packets to ACK packets on router R1 returned
to normal values.
```

```
Alert closed.
```

4. When the alarm server is restarted, the following long description appears, indicating that the alarm is aborted.

```
Alert was aborted due to server restart.
```

Segment 3 The information specified in the subject line about the identifying number of the alarm, state, severity and tags are repeated here. Tags are listed in a separate line with a “**Tags:** ” prefix. This line is omitted if no tags have been assigned for this alarm.

Segment 4 Information on event times that are relevant for the alarm are listed in this segment. For a new alarm, as in the example above, there is only the start time of the alarm. As long as the alarm is not closed or canceled, the end time is “**ongoing**”. An end time is only set when the status changes to “**Closed**” or “**Aborted**”. The duration is recalculated for each generated email for an alarm and shows how long an alarm situation has already been tracked. The format “**X days, HH:MM:SS**” is used. The number of days is omitted if the alarm was opened within the last 24 hours.

Segment 5 Information on the time span and number of events assigned to this alarm are comprised in this segment. As an alarm is opened when a specific number of events have been registered, the time stamp of the first detected event does not necessarily coincide with the timestamp of opening the alarm. The segment lists the timestamp of the first and last event, followed by the number of associated individual events.

Segment 6 The following segment displays information about the alarm. “**Trigger**” refers to the evaluation of events within the NeMo system sparking the alarm. The name of the trigger is followed by its unique ID. Alarms are generated for individual objects and are not aggregated across objects. See Section 3.3.1 for extensive details on triggers.

Segment 7 This segment contains a short “**Alert Description**”, which describes the entire alarm situation

Segment 8 In this segment the affected objects (i.e. routers and lines) are listed. In addition to type and name, the number of events that can be assigned to the respective object is specified, since the events combined in an alarm do not necessarily have to refer to only one object.

Segment 9 The final segment includes a URL that points to the same information at the web interface which also serves as an entry gate for further analysis.

3.3.1 Trigger of alarm messages

The triggers can be configured via the Admin menu as explained in the Admin manual.

3.3.2 Temporal development of alarms

The temporal development of simple alarms is based on time windows. An alarm creation or status change are conducted if a configured number of events is observed in a time window that is also configured, for example if four events matching the condition are registered for an object within five minutes. The procedure is shown in Figure 3 which is explained afterwards.

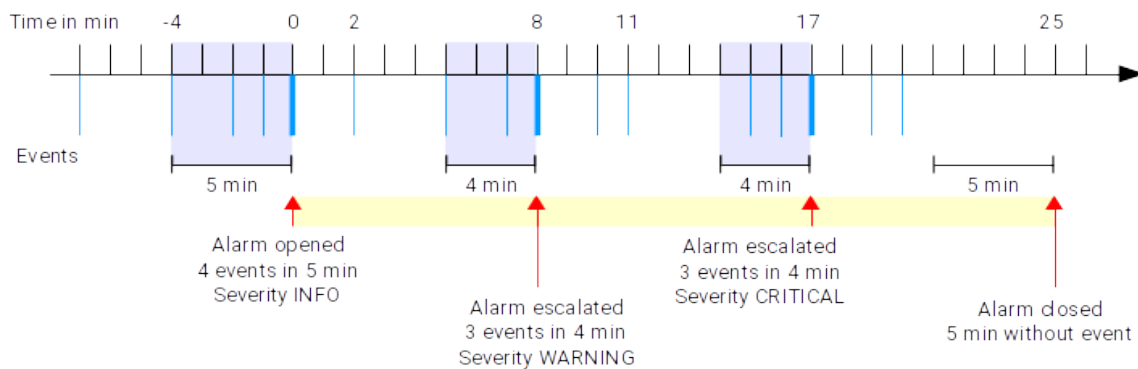


Figure 3: Development of an alarm based on time windows

- The horizontal black arrow represents the time, the division lines on the upper side are minutes, interesting times are marked by the minute indication relative to the time of opening the alarm (Minute 0).
- Below the arrow in blue are the times of single events.
- The time windows that lead to a status transition are highlighted in light blue, the event that ultimately triggered the transition is indicated by a bolder line.
- The status transitions themselves are marked by red arrows—these always coincide with an event and the end of the selected time window.
- The life of the alarm is indicated by a yellow bar at the level of the red arrows.
- For example, the individual events in Minutes -7, 2, 10 and 11 that are not highlighted in blue do not fall within a time window with event activity that justifies increasing the alarm's severity.
- In this case, two increases in severity occur as soon as three events are observed within four minutes:

1. In Minute 11 three events in the Minutes 8 to 11 were registered, but one event cannot be involved in two status transitions. Therefore, three additional events must be observed.
 2. This is given in Minute 17. Here, the three events are registered in three minutes, so the condition for the status transition is fulfilled.
- In Minute 25, a period of 5 minutes has elapsed without a single event, so the alarm is closed. Alternatively, an alarm can also be closed if the activity, i.e. the number of events observed, is lower than a configured threshold value within a configured time period. In this way, an alarm can also be closed if, for example, a single event is registered in an observation time period.

3.3.3 Changes of alarm messages over time

Listing 2 shows how the report may change over the life span of an alarm². Changes are marked in bold text and colored in red. The development of this alarm corresponds to that of the example alarm in Figure 3.

```
Subject: [NeMo] #22241 UPD WARN - Infrastructure: R1 - "Still high SYN/ACK
ratio"
Date: Tue, 24 May 2011 05:06:51 +0000 (UTC)

Observed more high values of the ratio of SYN packets to ACK packets on router R1.

Upgraded to severity Warning.

Severity: Warning

Duration: 00:05:15

Last Event Seen: 2011-05-24 05:06:00
Event count: 8

Affected Objects:
  Type           Name           Event Count
-----
Router          R1              8
```

Listing 2: Changes to an alarm message at the “Info” to “Warning” transition

More events (four new events to the original four) have been evaluated. Three of these new events were observed within a four minute window. This fulfills the condition for a status change of the trigger, so that the severity increases from “Info” to “Warning”.

In this case, the next transition is configured with the same condition—three events within four minutes—so that after further events were detected and processed and 17 minutes after the alarm was opened, the state switches to “Critical”. The short description is the same as the previous one, but this is only due to the identical configuration of the status transitions.

```
Subject: [NeMo] #22241 UPD CRIT - Infrastructure: R1 - "Still high
SYN/ACK "ratio"
Date: Tue, 24 May 2011 05:18:31 +0000 (UTC)
```

²Only the relevant lines are displayed.

Observed **critically** high values of the ratio of SYN packets to ACK packets on router R1.

Upgraded to severity **Critical**.

Severity: **Critical**

Duration: **00:16:55**

Last Event Seen: **2011-05-24 05:18:00**

Event count: **13**

Affected Objects:

Type	Name	Event Count
Router	R1	13

Listing 3: Changes to an alarm message at the “Warning” to “Critical” transition

After five minutes have passed since the last event being assigned to this alarm, the alarm is automatically closed and corresponding email notifications are sent.

Subject: [NeMo] #22241 **END CRIT - Infrastructure: R1 - "Normal SYN/ACK ratio"**

Date: Tue, 24 May 2011 **05:26:05 +0000 (UTC)**

The ratio of SYN packets to ACK packets on router R1 returned to normal values.

Alert closed.

Status: **Closed**

End Time: **2010-05-24 05:25:38**

Duration: **00:24:02**

Listing 4: Changes to an alarm message at the transition from “Open” to “Closed”

3.3.4 Meta alarm additions

Messages for meta alarms differ in the structure of simple alarm messages only by an additional segment following the list of affected objects. A list of aggregated alarms with their alert IDs and a short description is inserted, allowing a quick assessment of the overall situation. In the example in Figure 5, the system assesses the alarm as a UDP scan of a target host, since many small UDP packets are sent to many different target ports.

Type	Name	Event Count
Router	R1	36

Child Alerts:

Alert ID	Description
25294	Few to Many Ports.
25295	Small packets .
25296	High UDP packet counts.

Further Details:

Listing 5: List of aggregated simple alarms of a meta-alarm

4 NeMo web interface

The web interface offers extended information on alarm messages already sent by email and allows browsing all events. Individual events that led to the alarm can also be investigated. The web interface allows direct observation of current situations without having to refer to emails.

Accordingly, there are two points of entry:

The email interface provides a link to the alarm message in the web interface.

The web interface provides detailed information and is accessible at: <https://nemo.example.org>

4.1 General structure of the web interface

The web pages of the DDoS application are organized according to the following pattern. From top to bottom there are three areas:

- **navigation** (top, dark and light blue background),
- **content** (middle, white background) and
- **status information** (bottom, dark blue background).

The screenshot displays the DFN-NEMO web application interface. At the top, a navigation bar contains tabs for 'Summary', 'Alerts', 'Objects', 'Mitigations', 'Topology Map', 'Visual Explorer', 'Sparklines', 'Detectors', and 'Preferences'. The 'Summary' tab is active. Below the navigation bar, the page title is 'Summary of Open Alerts'. A filter section shows 'Tags' with buttons for 'All Severities (1)', 'Critical (1)', 'Warning (0)', and 'Info (0)'. A message indicates 'See all alerts opened during the last hour, the last 24 hours, or last 7 days.' Below this, a table lists the alerts. The table has columns: Alert ID, Workflow Status, Severity, Duration, Start Time, Event Count, Tags, Description, and Details. One alert is shown with ID 236477, status 'New', severity 'Critical', duration '13 h, 38 min (ongoing)', start time '03:35, today', and event count 43884. The description is 'Missing SNMP updates from one or several routers.' and the details are 'Missing heartbeats...'. The footer shows 'Logged in as: [user]', 'System Information', and 'DFN-NEMO Version 7.4.5'.

Figure 4: One page of the web application

The first line of the navigation area contains links to the various views of the application. The view currently displayed in the content area is indicated by bold text and a small light blue triangle underneath the view name (See “Summary” in the upper left corner in Figure 4). From left to right, the following views are available:

Summary The start page lists all current alarms. See Section 4.2 for more details.

Alerts The alerts page allows extensive access to all alarms and provides search and filtering options. See Section 4.3

Objects This page provides direct access to all objects modelled in NeMo (e.g. router, lines and networks). Individual objects may be isolated using search and filtering. See Section 4.4

Mitigations Active and inactive defensive measures are listed in this view. See Section 4.5.

Topology Map A graphic visualisation of routers and core lines of the Network, supplemented by the traffic volume on these devices. See Section 4.6

Visual Explorer This pages allow parallel and detailed visualisation of traffic volume on multiple objects using freely configurable time windows. This is especially helpful for visual analysis of traffic flows in combination with the visualization of traffic models. See Section 4.7.

Sparklines This link opens a flexibly configurable view of sparklines on multiple objects. Sparklines are visualisations of traffic flows, simplifying detection of parallel development on multiple objects. See Section 4.8 for details.

Detectors This pages shows all active detection methods (detectors), that can be selected for an object. For each detector type and number of associated objects are listed. Further details about individual detectors can be accessed as well. See Section 4.9 for details.

Settings Various settings can be configured using this page. See Section 4.10 for details.

On the right hand side of these menu items are an indication whether the instance is in managing (“(active)”) or reading mode (“(passive)”), and NeMo’s logo.

The second line of the navigation area displays descriptive information about the data being displayed for the current view in the content area (also called “extended navigation area”). It is also possible to select the time zone for the data being displayed. The three options are UTC³, CET⁴ and CEST⁵.

If appropriate, status information is displayed at the bottom of the page. This comprises currently of the username of the authenticated user in the left corner and the version of the software in the right corner. In the center a page [system information](#) is linked to. This view provides core data to assess the status of the application and its servers. This view is explained in detail in Section 4.11.

³Coordinated Universal Time or Greenwich Main Time (GMT). This is CET -1 hour.

⁴Central European Time, GMT +1 hour

⁵Central European Summer Time, GMT +2 hours

4.2 Overview of the most important alarm messages

This overview, accessible via the [Summary](#) link in the navigation area and displayed in Figure 4, lists current alarms and alarm analyses. Alarms can be filtered by alarm category (tag) and severity.

- The default is to display all open non-suppressed⁶ alarms, regardless of tags or severity.
- Filtering of alarms can be enabled with the checkboxes at the top of the content area. There is a checkbox for every tag. If tags are selected only those alarms are displayed that are associated with one of the selected tags. If no tags were assigned to alarms, the list is empty and filtering by tags is not available (see Figure 4).
- To filter by severity the tabs “[Critical](#)”, “[Warning](#)” and “[Info](#)” are used. The number of corresponding open alarms is already displayed in the tab’s header.

There are also three links “[the last hour](#)”, “[the last 24 hours](#)” and “[the last week](#)”, that will show all alarms—even the ones that are suppressed or have been closed—matching the selected filters for the respective time window. See Section 4.12 for more details.

Below, all alarms matching the selected filters are displayed. If open alarms match the filters their count is displayed on top and bottom. If the list contains more than 20 alarms, pagination techniques are used, distributing alarms to 20 per page. In this case and to switch between pages, navigation elements are enabled. If no alarms match the filters the message “[There are currently no open alerts](#)” is displayed instead of alarms (See Figure 5).

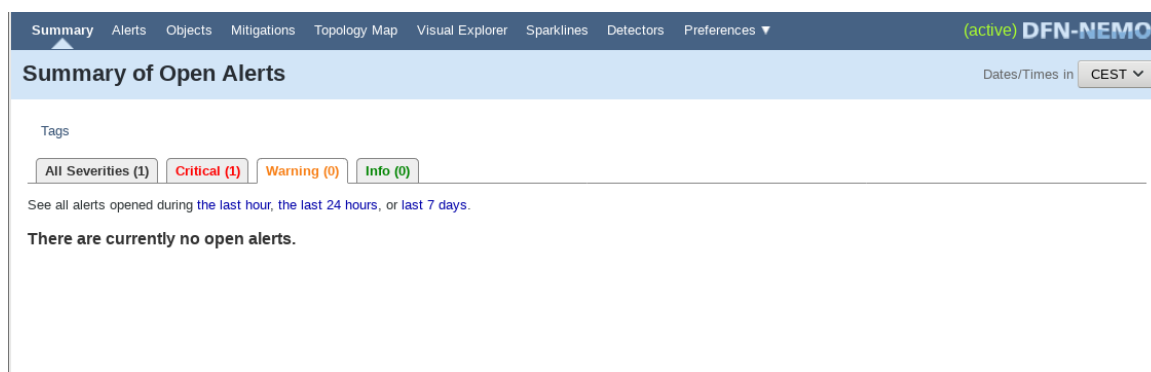


Figure 5: Content area if no alarms match the selected filters (part view)

If the set of alarm messages is not filtered by severity, the list is firstly sorted by severity in descending order and secondly by starting time in descending order: At the top of the list are the most severe and within each severity level the most recently opened alarms are listed atop. If all open alarms are associated with the same severity the list is sorted by starting time in descending order. Alarms with identical starting time are sorted by their alert ID in descending order.

Each entry in the alarm list starts with a “play symbol”⁷ identifying alarms running in various views of the web interface. For each alarm the following data is displayed:

⁶See the glossary in Section 5.1.

⁷A triangle pointing to the right within a circle. See Figure 4

Alert-ID An identifier is allocated by the system while creating a new alarm. The identifier is a consecutive number providing unambiguous identification of every single notification. The identifier is used consistently in all generated emails and within the web interface. The Alert ID links to the alert's detailed view (see Section 4.13).

Workflow status Possible values for the workflow status of an alarm are: “[New](#)”, “[Seen](#)” or “[Analyzed](#)”. Every alarm starts at [new](#). As soon as any user⁸ accesses the detailed view of an alarm (see Section 4.13) the state changes to [seen](#). When an alarm analysis for an alarm has been created, the state changes to [analyzed](#).

Severity The severity of an alarm, see Section 3.1.2.

Duration This field prints the duration of the alarm, by the minute. As the overview shows only ongoing alarms, this corresponds to the time span since opening the alarm and is therefore annotated with “([ongoing](#))”.

Start time Indicates the time when the alarm was opened, see Segment 5 in Section 3.3.

Event count The number of events associated with the alarm.

Tags Categories assigned to this alarm.

Description The description is identical to the “[alert description](#)” in Segment 5 of Section 3.3.

Details A sparkline visualises the timely course of network flow data responsible for the alarm.

Except for the alert ID there are no links. Therefore, these lists do not allow further investigation. Searching and filtering can be executed on these alarms at the overview of all alarms (see Section 4.3).

Underneath of this list the user generated analyses are displayed—if there are any to display—with the following data:

View Each entry in the list starts with a link “[View](#)” that navigates the user to the view of the analysis.

Alert ID An ID give to the analysis by the NeMo software. This is also a link to the detailed view of the analysis, see Section 4.13.3.

Object The name of the object on which the alarm was detected.

Comment A comment by the creator of the analysis. Should describe the alarm in more detail.

Created by The name of the creator of the analysis.

Created on The time when the analysis was created.

⁸The state is recorded once for all users.

4.3 Viewing all alarms

Following the link [Alerts](#) in the navigation area, all alarms are listed in this view. The visualization is similar to the overview, explained in Section 4.2. Only non-suppressed alarms are displayed, unless the user selects a subset of alarms to be displayed in which case suppressed alarms belonging to the selected subset are displayed as well. Results are limited to 20 per page and are therefore splitted over several pages if required. Navigation controls are then used to switch between pages.

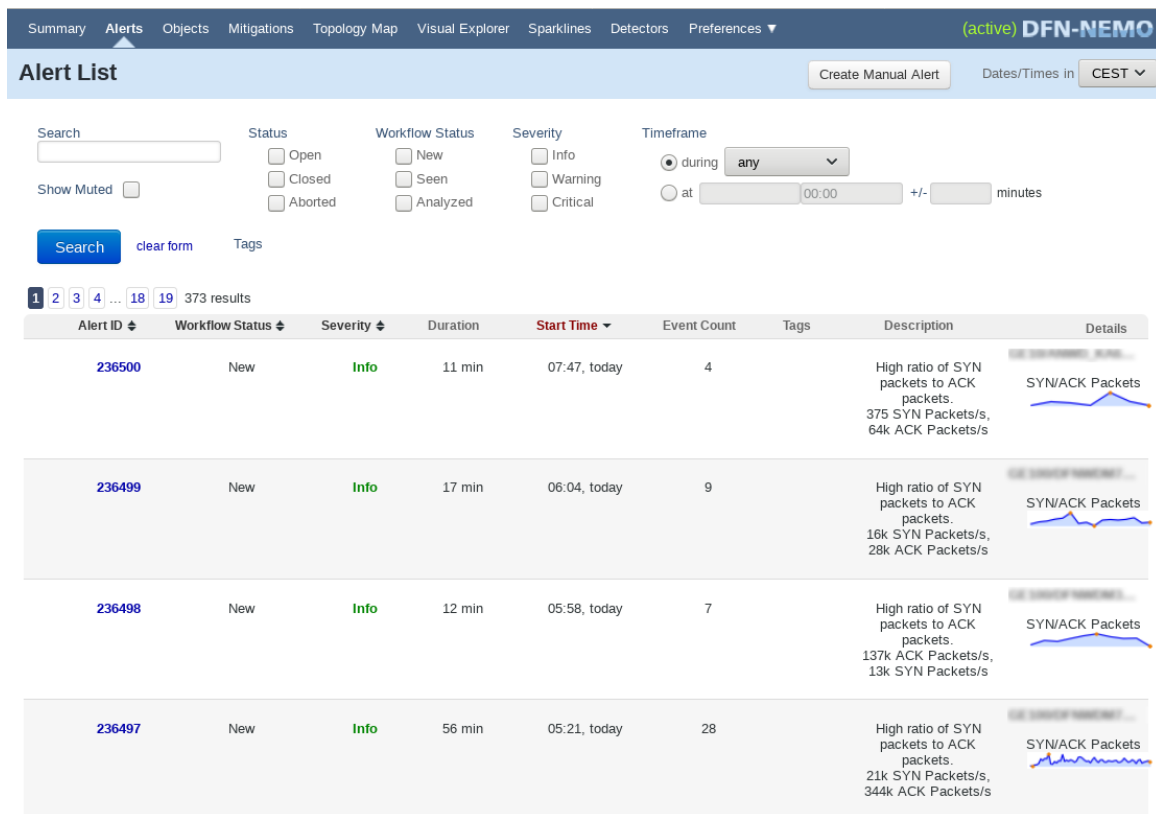


Figure 6: View of all non-suppressed alarms

4.3.1 Manually creating an alarm

In the upper right corner of the extended navigation area a button labeled “[Create manual alert](#)” is to be used to create a new alarm manually. This will open a form in which various data for the new alarm are to be entered or selected (See Figure 7). These are [severity](#), [start time](#), [end time](#), [description](#), [metrics](#)⁹ and [associated objects](#). The list of affected objects for selection can be searched and filtered using the search field in its upper right corner. The button “[Create alert](#)” will then create an alarm using the data provided in the form. The system will acknowledge this by displaying a notification “[Manual alert XXXX created](#)”.

⁹Metrics are measures to quantitatively assess the network traffic and behavior. See also [https://en.wikipedia.org/wiki/Metric_\(unit\)](https://en.wikipedia.org/wiki/Metric_(unit)).

Create Manual Alert Dates/Times in **CEST**

Severity: **Info**

Start time: 2019-04-23 09:19:30

End time:
Leave empty for ongoing alert

Description:

Metrics:
 Flows
 Traffic
 Packets
 TCP Packets
 ICMP Packets
 UDP Packets
 Incoming Packets
 Incoming Traffic
 Incoming TCP Packets
 Incoming UDP Packets
 Incoming ICMP Packets
 Outgoing Packets
 Outgoing Traffic
 Outgoing TCP Packets
 Outgoing UDP Packets
 Outgoing ICMP Packets
 Traffic/Packets
 Please choose at least one metric

Affected Objects

Show **10** entries

Search:

	Name	Type	Category
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	
<input type="checkbox"/>	...	Autonomoussystem	

Showing 1 to 10 of 1,690 entries

[Previous](#) [Next](#)

[Create Alert](#)

Logged in as: [User] System Information DFN-NEMO Version 7.4.5

Figure 7: Form to create an alarm manually

4.3.2 Filtering the list of alarms

Above the table control elements for filtering and searching alarms are located. The searching functionality can be controlled by entering terms to search for in the textfield. The filtering is configured by selecting checkboxes and values in the drop down boxes. The filters and search terms will be applied when the button “[Search](#)” is clicked or when the [return](#) key is pressed while the cursor is in a text field. “[Clear form](#)” can be used to reset the filter configuration that can be applied with another click on [Search](#).

The search is executed in data fields for description, alert id and the name of the alert. This can help to find certain alarms or alarms that match a certain situation (e.g. an anomaly in UDP protocol traffic) or all alarms on a specific object. The search term can be at any position

in the data fields to be searched, searching for “UDP” results in a list containing all alarms where the UDP protocol is mentioned in the description, whereas searching for “WDM4711” will list all alarms on the connection “GE10/DFNWDM4711_SOM_WHR”.

The checkboxes enable filtering by alarm status, workflow status or severity based on the options in Table 1. Each filter object can handle multiple filter values, enabling to search i.e. for [open](#) and [closed](#) alarms with severity [info](#) or [warning](#). If no selection is made for a filter object, filtering by this filter object is disabled, resulting in any objects being displayed regardless of their value in the respective field.

Filter object	Filter setting	Result
State	Open	Open alarms only
	Closed	closed alarms only
	Aborted	aborted alarms only
Workflow status	New	new alarms only
	Seen	seen alarms only
	Analyzed	analyzed alarms only
Severity	Info	Alarms with severity info only
	Warning	Alarms with severity warning only
	Critical	Alarms with severity critical only

Table 1: Filter options

The list may also be filtered by the start time of alarms using a drop down menu or explicit indication, see Table 2.

Filter object	Filter setting	Result
during	any	All start times
	last hour	Start time within last hour
	last 24 hours	Start time within last 24 hours
	last week	Start time within the last week
	last 30 days	Start time within the last 30 days
at	Start time +/- tolerance	Start time within -tolerance and +tolerance in minutes

Table 2: Filter options for start time

When the page is loaded, the default is to show only non-suppressed alarms to increase information concentration. If required, suppressed alarms can also be included in the list, by simply checking the Checkbox labeled “[Show muted](#)” placed above the search field. The suppressed alarms will then be integrated in the list, but marked by using darkgray textcolor and the label “([muted](#))” in their description. See Figure 8 for an example.

290903	New	Info	11 min	07:49, 2015-09-14	5	(muted) Low packet rate.	GE10/DFNWDM4711_SOM_WHR
--------	-----	------	--------	-------------------	---	--------------------------	-------------------------

Figure 8: Visualization of a suppressed alarm in the alarm list

Filtering by [tags](#) is similar to the filtering on the overview page (see Section 4.2), but filtering will only be applied, when the [Search](#) button was clicked. The form can be reset using the link

labeled “[clear form](#)” but will likewise only be applied to the results when the [Search](#) button is clicked afterwards.

4.3.3 Sorting the table

The list of alarms can be sorted using the arrow symbols in the columns’ header cells. The default sort order is descending by the start time, so the most recent alarm will be displayed in the first line. Clicking on a heading will sort the table in ascending order of the data fields in that column, e.g. from small values to large, for severity from [info](#) to [critical](#). The next click on the same header will then reverse the sort order. The currently selected criteria for sorting is highlighted using red text color in the header cell and by displaying only one arrow symbol. Therefore, when the default sort order is used, the [start time](#) header cell only contains a single arrow symbol pointing south.

4.3.4 Visualization of closed alarms

The visualization of closed alarms differs slightly from the listing of alarms. Closed alarms do not have the “play” symbol. The duration of a closed alarm is obviously not labeled ([ongoing](#)). The duration of an alarm is the time span between opening and closing the alarm. An alarm is closed, if no new events are detected within a defined time span that is associated with the alarm and depends on the detection model. Obviously the end time is not identical with the time of the last event associated with an alarm, due to the time out interval required to assess the end of an alarm situation.

4.4 Viewing all objects

The view on all objects is linked by the [Objects](#) item in the navigation area (Figure 9). It allows direct access to routers, lines, subnets and autonomous systems (AS¹⁰) modeled in NeMo. This view is similar to the [list of alarms](#) and its filter and search techniques (See Section 4.3.2).



Figure 9: Viewing all objects

Above the table control elements for search and filtering as well as navigation are located. After selecting filter and search options those will be applied when the [Go](#) button is clicked or

¹⁰An AS (autonomous system) is a number of IP ranges or IP subnets that are managed as a single entity.

when the [Enter](#) key is pressed while the cursor is placed in a search field.

The search is conducted on the name of the object, whereas the full name of the object does not have to be specified. The result set contains all objects for which the search string is a substring of the name of the particular object. This search is case insensitive.

Filtering of objects can be conducted on the type and category of objects as well as its state with the options as specified in Table 3.

Filter object	Filter setting	result
object type (category)	all	all objects
	Routers (all)	all routers
	Routers (core)	core routers
	Routers (other)	all routers except the core routers
	Lines (all)	all lines
	Lines (core-net)	all lines of the core net which are not connected to a core router
	Lines (core)	all lines connected to a core router
	Lines (user)	all lines connecting users
	Lines (peering)	all lines connecting the Network to other networks
	Lines (L3 VPN)	all lines that part of a L3 VPN
	Networks	all networks configured
	Autonomous Systems	all autonomous systems
	Router Groups	predefined groups of routers
Object state	only active	Only active objects
	only inactive	Inactive objects only
	all	All objects

Table 3: Filter options

The result set may be further reduced by using the predefined filters labeled [subcategories](#). Subcategories are for example be used to assign objects to various L3 VPNs or to mark objects with only a limited and sporadic amount of network traffic to be excluded from modeling typical network behaviour. All subcategory filter settings can be removed using the [clear selection](#) link.

Results are limited to 20 entries per page. Navigation controls allow to switch between pages. The order of entries may be changed using the triangle symbols provided in the heading of each column, whereas the default sorting is by name in ascending order. The following data fields are provided for each object:

name The name of the object. For lines and routers this is identical to the name used throughout the . For subnets under monitoring this is identical to their CIDR¹¹ notation. The name is also a link to the detailed view of the object (See Section 4.12).

type The type of the object, with possible values being “[Routers](#)”, “[Line](#)”, “[Net](#)” and “[Autonomous System](#)”.

¹¹CIDR is an abbreviation for Classless Inter-Domain Routing. It combines an IP address and a netmask as a combination of host and network IP ranges. E.g. 192.168.0.1/16 means the IP address 192.168.0.1 in the subnet 192.168.0.0. See also https://en.wikipedia.org/wiki/Classless_Inter-Domain_Routing.

category The category within the object type. The combination of type and category is identical to the one used for filtering and is described in Table 3.

subcategory Sub categories defined for this object are listed here. Subcategories are used within NeMo to illustrate specialities and similarities of objects in a flexible way and use them for automated processing. The following sub categories usually exist:

low traffic used to mark objects with a limited amount of network traffic. A limited amount of network traffic prevents to determine a characteristic network behaviour, which is required for machine learning processes implemented in NeMo for detection. Therefore, if an object is marked with this, such detection methods will not be applied to it.

further categories are used to mark lines belonging to L3 VPNs

active/inactive This field indicates whether the object is activated within NeMo and the DDoS application processes data for this object or whether the object has been deactivated.

last update Indicates the time when NeMo had received and processed data for this object. A relative time scale to the current time is used.

created on The time when the object was created within NeMo.

Flows This field provides a qualitative overview of the flows detected for this object within the last 24 hours. Values are averaged for 30 minute time frames and displayed using a sparkline¹².

¹²You will find more on sparklines in Section 4.8.

4.5 Mitigations

A table in the content area of the webpage lists currently active mitigations, as displayed in Figure 10.

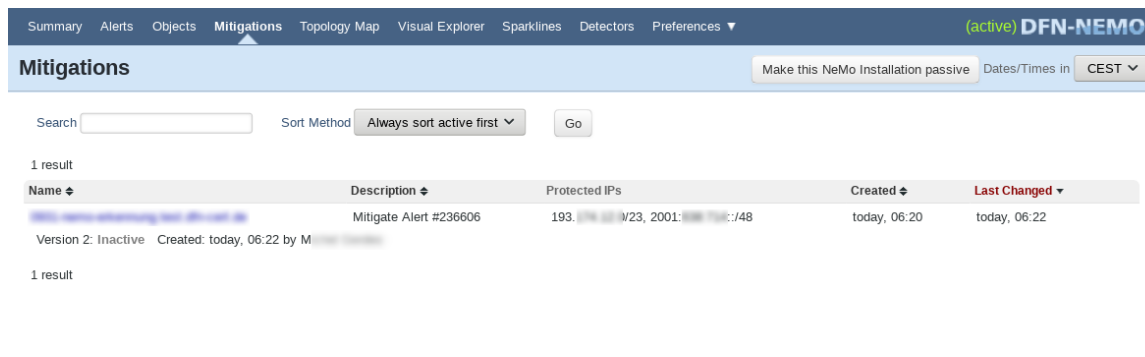


Figure 10: Overview on mitigations (part view)

The navigation area is extended by a button labeled “[Make this NeMo Installation passive](#)” or “[Make this NeMo Installation active](#)”. The active NeMo Installation is the one managing the mitigations. Switching between active and passive mode requires a dedicated authorization and should only be executed when no mitigations are active, otherwise the active mitigations will go into an state which is hardly to control. The user should be fully aware of what he or she is doing, so the system will ask for a confirmation before actually switching the mode.

The search field can be used to search on all fields and properties of mitigations using one or multiple search strings. Only those results will be listed that contain all search strings (separated by spaces). The search is case insensitive.

Results can be sorted by two different approaches:

Always sort active first This is the default, so active mitigations are listed atop inactive ones.

Standard sort order Mitigations may be sorted by their name, alternatively.

The following data fields are displayed for each mitigation. The sort order of the mitigations may be changed using the triangles in columns’ heading.

name A scheme is used for naming mitigations: <[serial number](#)>-<[name of detector](#)>. The name is also a link to the detailed view on the mitigation (see Section 4.14).

description A freetext description created by the editing person. This usually references the alarm.

created on The time of creating the mitigation.

last change The most recent time of editing the mitigation.

Some more data is displayed for each mitigation that cannot be used to sort mitigations:

protected ips The IP ranges (IPv4 as well as IPv6) to be protected by this mitigation.

Version A serial number of revision

Status An indication of the deployment of this mitigation. Possible states are explained in Section 4.14.

Name of last editor The name of the person creating the most recent revision of this mitigation. This is prefixed by the term “[Created:](#)” and the time of creation of this revision.

4.6 Topology map of the Network

The topology map can be accessed using the “[Topology Map](#)” link in the navigation area. It displays an overview on data about network flows and alarm situations in the Network that is available in NeMo of a certain period of time. Figure 11 shows the default page when opening the topology map.

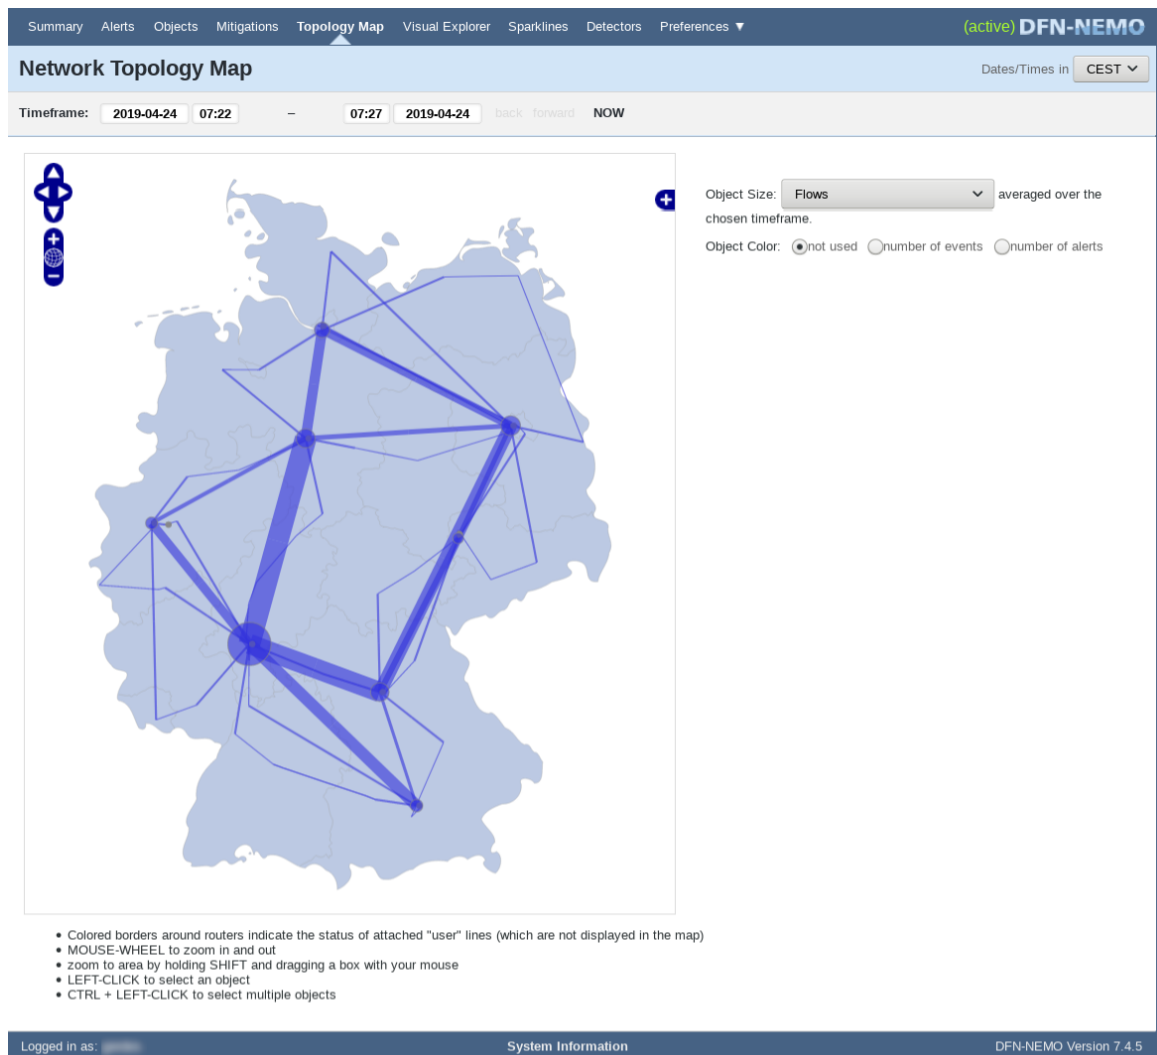


Figure 11: Topology of the Network

Below the navigation area and title of the page controls are embedded to select the time frame used for the map. The default is to show the last five minutes.

On the left hand side a map of Area is displayed in which locations of the routers of the Network are displayed by dark blue dots. The center of each dot matches with the GPS coordinates taken from the . If multiple routers are located close to each other, pins are used to illustrate each router’s location properly (See Figure 12 for an example). Core net lines are displayed by bold blue lines connecting the dots.

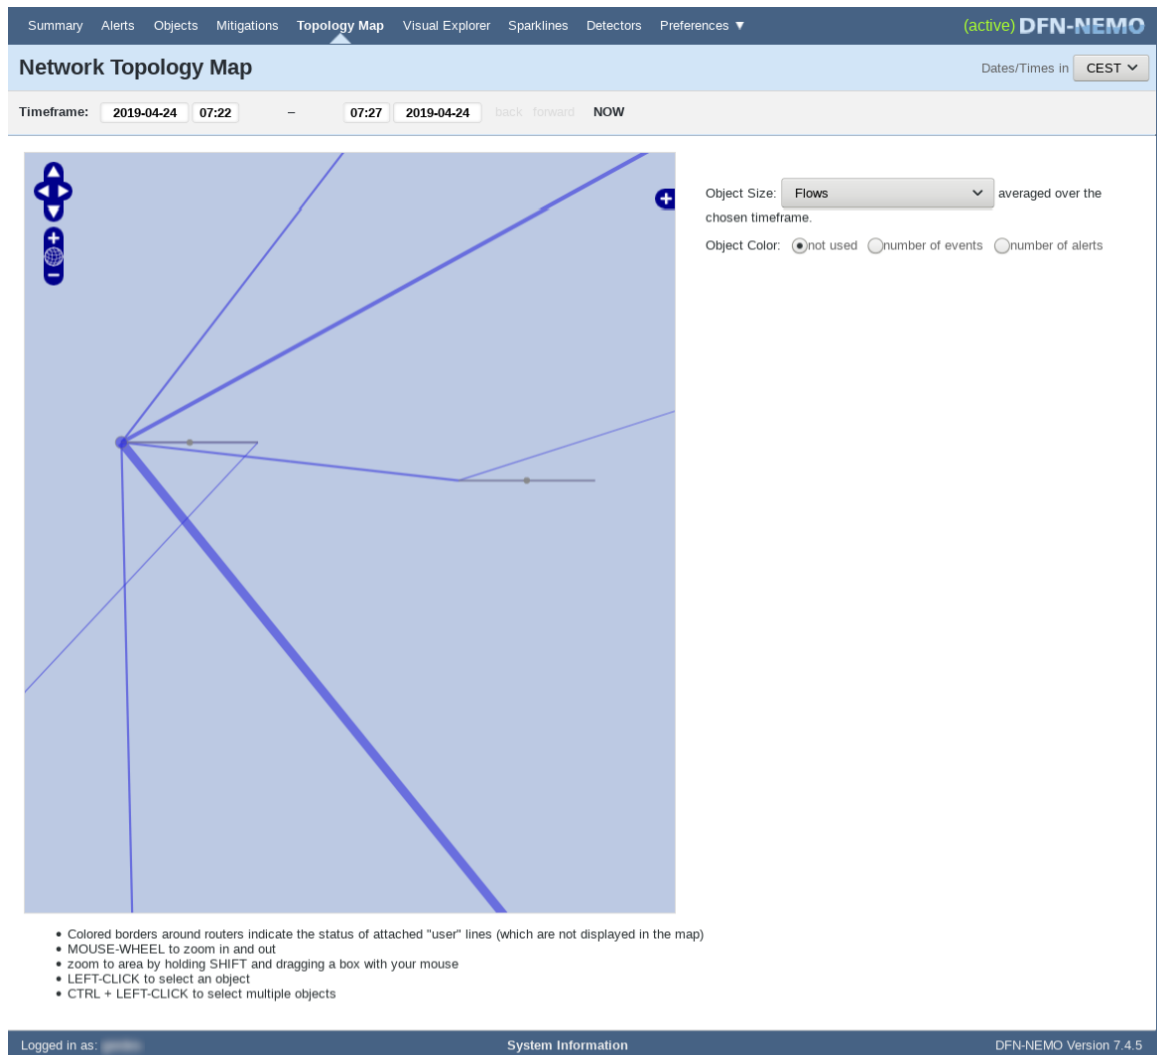


Figure 12: Zooming in shows grey pins indicating the proper location of routers

On the right hand side controls are located to change the visualisation of the map. These focus on the different net flow patterns and alarm situations.


4.6.1 Adjustment of the timeframe to be visualized

Controls to set the timeframe define it by specifying the start and end time. This is illustrated in Figure 13. Both start time and end time are to be entered using “YYYY-mm-dd” for the date and “HH:MM” for the time on that day, whereas “YYYY” means the year, “mm” the months of the year, “dd” the day of the month, “HH” the hour of the day and “MM” the minutes of the hour. Start and end time are separated by a dash. The start time is entered at the left hand side and the end time is to be entered on the right hand side of the dash. For the end time, the sort order of input fields for date and time are reversed.

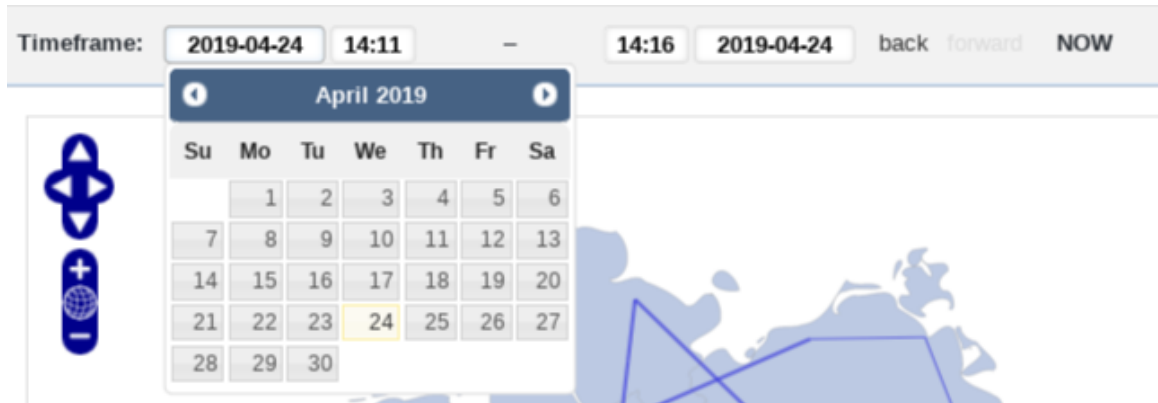
On the right hand side of these controls three links may be displayed. The link labeled “[NOW](#)” moves the timeframe to the actual time but keeps the length of the time frame (interval). If the interval is ten minutes long and the link is clicked the timeframe will be set from ten minutes ago until the time the link was clicked. If this link is clicked, the other two links are activated (see Figure 13b). These additional controls labeled “[back](#)” and “[forward](#)” allow to navigate within the time history, jumping back or forward by this interval.



(a) Controls to set the timeframe




(b) Activated link to navigate within the time history



(c) Entering the date using the calendar widget



(d) Controls to change the time frame using the mouse pointer



(e) Controls for a symmetric shift of the timeframe using the mouse pointer

Figure 13: Time frame controls

When specifying start and end time a calendar widget will pop up upon clicking the date field. This widget may help selecting the intended date. See Figure 13c.

Similar widgets will be displayed when the mouse pointer is on the time field (see Figure 13d) or when it is on the dash icon (see Figure 13e). When using the time widget the time can be increased using the buttons starting with a “+” or decreased using the buttons starting with a “-”, respectively. The predefined intervals are seven days (“7d”), one day (“1d”) or one hour (“1h”). The circle in the center can be used as a slider for setting the start time and end time. By clicking and pressing the main mouse button while moving the mouse pointer the start time will be increased when moving to the right and decreased when moving to the left.

When the time range widget (click on the dash) is used, the time frame is shrunk or enlarged depending on the button clicked, while the central circle is used to

zoom in shrinking the time frame by moving the mouse pointer left while pressing the mouse button

zoom out enlarging the timeframe by moving the mouse pointer to the right while pressing the mouse button.

4.6.2 Changing the appearance of the map

The map view shows a coarse map of Area, in which the locations of the routers in the Network are indicated as dots. The diameter corresponds to the number of flows detected for and processed by NeMo for the selected timeframe. Core net lines connecting routers are visualized by direct lines between two dots. The boldness of the lines corresponds to the number of flows on this particular line over the specified timerange. When the mouse pointer is over a Network component the name as defined by the is displayed.

The presentation can be adapted using either

graphical controls In the upper left corner of the map controls are located (see left hand side of Figure 14). The top triangles will move the extract to top, right, down or left, whereas each triangle points in the direction in which the extract will be moved. The plus or minus icons located below these are used to zoom in or zoom out, respectively.



Figure 14: Controls for moving and zooming the extract of the map

mouse The visible cutout can also be manipulated using the mouse. By clicking and moving the pressed mouse button, the extract of the map will move in the same direction. When the mouse wheel is used, the map is zoomed in or out at the location where the mouse pointer is at that moment. The third option to adjust the cutout is by pressing the [Shift](#) button on the keyboard and then clicking and moving the mouse pointer while keeping the button of the mouse pressed. This will create a rectangle that will fill the map's cutout once the pointer is released.

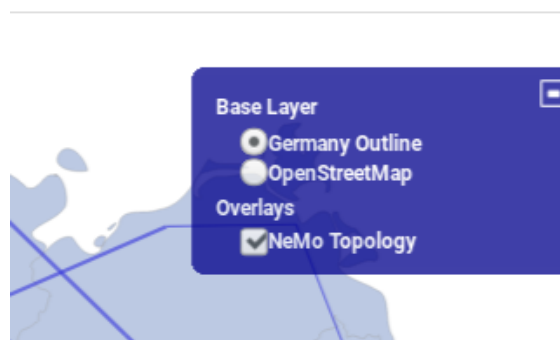


Figure 15: Control for selecting the map and objects of the Network

In the upper right corner of the extract an “+” symbol is located that will swing out more controls once clicked (see Figure 15). These controls allow on the one hand to switch between a contour map and a tourist map based on data of the OpenStreetMap project¹³ (see Figure 16) and on the other hand active and deactivate objects of the NeMo topology, including routers and lines, in this visualization.

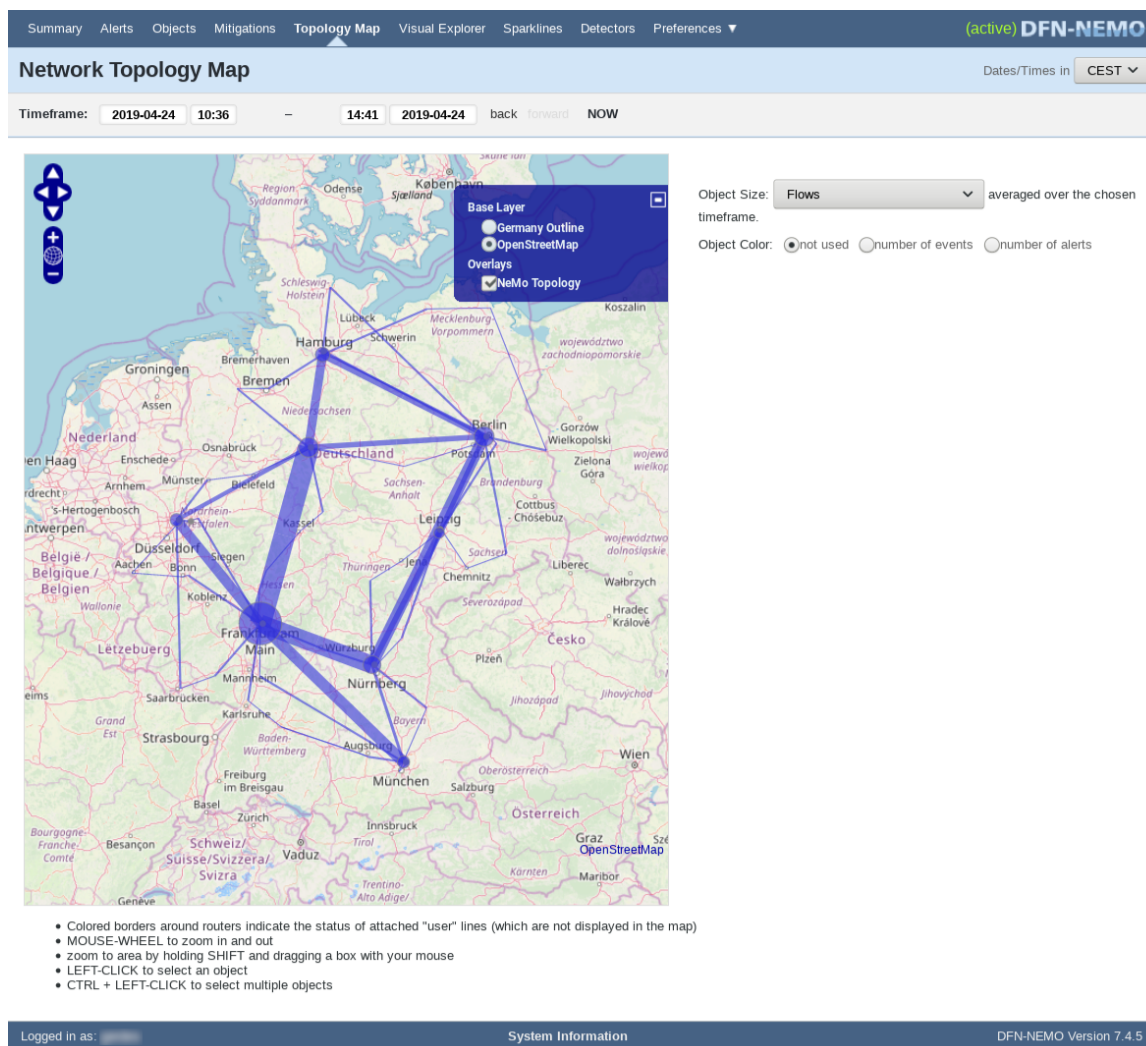


Figure 16: Visualization of the topology map using OpenStreetMap data

4.6.3 Adjustments of the visualization

On the right hand side, next to the map various controls allow to adjust the information being displayed in the map (see Figure 17), which will be explained in this section.

¹³OpenStreetMap is a project in which—similarly to processes at Wikipedia—free and open tourist maps are created by volunteers. See <http://www.openstreetmap.org>




Figure 17: Controls for adjusting the objects in the topology map

4.6.3.1 Object size The drop down menu defines the size of the Network objects within the map. The default is to set the size in relation to the number of registered flows at this object during the time frame. Other indicators that can be used for the size of objects on the map include the number of registered UDP packets or in relation¹⁴ to the object category (see Section 4.4). All objects may also be displayed with the same size. If SNMP counter for packets or bytes are used, the direction of the network traffic in relation to the network traffic on this line is visualized by splitting the line at the same aspect ratio as the traffic ratio (see Figure 18). If both parts are of the same length, a similar amount of traffic is send from each router over this line. If the split is at .66, so Router A sends twice the amount of network traffic over this line as Router B, if the split is closer to Router B.



Figure 18: Visualization of SNMP indicators in the topology map, the router in sends roughly three times as much traffic over the direct line as the router in

4.6.3.2 Object color Below the selection of the indicator the coloring scheme can be defined. The default is to color objects according to no scheme (“not used”), all objects are colored blue¹⁵. If either indicator “number of events” or indicator “number of alarms” is used, the colors as defined in the row labeled “Color legend” are used to color the object as defined by the caption for each value range within the colored boxes.

The following coloring schemes are available:

Based on number of events If coloring based on events is enabled, objects without events are colored green, objects with one to five events are colored in orange and objects with more than five events are colored in red.

¹⁴For routers the radius is scaled appropriately, for lines the boldness of the line is scaled appropriately.

¹⁵Except when the SNMP indicators are used, a complementary color (yellow) is used to color the other part of the line.

Based on the number of alarms If coloring based on alarms is enabled, objects without alarms are colored green, objects with one or two alarms are colored orange and objects with more than two alarms are colored red.

For reasons of clarity only core net lines are visualized in the map. Therefore, events or alarms on other lines (i.e. peerings or user lines) are visualized by a colored border of the router: e.g. a router with a red dot and an orange border means that there were more than five events (or more than two alarms, respectively) on this router detected and one to five events (or one or two alarms, respectively) on non core net lines connected on this router. Settings as defined in Figure 17 result in a visualization as in Figure 19.

If coloring is enabled, traffic direction visualization for SNMP values cannot be enabled as well.

4.6.4 Sparklines for selected objects

Objects can be selected¹⁶ and sparklines of the traffic patterns is displayed for these objects on the right hand side south of the coloring controls. Selected objects are colored purple. See Figure 19 for an illustration.

The sparklines can be adjusted using the controls at the top, that will appear when at least one object is selected. The controls allow to choose an indicator and the time frame for the sparklines. The default is to display the number of flows for two hours. The end time of the sparkline always matches the end time of the maps' end time (see Section 4.6.1), even when the interval of the sparklines is longer than the interval of the map. Next is a link labeled "[View selected objects in Visual Explorer](#)" allowing a more fundamental analysis of the selected objects. Clicking that link will open a new window showing the [Visual Explorer](#) that will be discussed in Section 4.7. The time frame of the map, the selected objects and the indicator for the sparklines are preselected in the [Visual Explorer](#) when the link is clicked. The link is followed by sparklines of the selected objects. For each object the number of events and the number of alarms are shown as well, if there were any. If the mouse pointer is located on a sparkline the actual value at that point appears in a pop up box. The object will be highlighted in the map as well by increasing the density of the coloring (color changes to dark violet) and showing its name.

¹⁶Holding the CTRL key to select more than one object or to remove an object from a selection, clicking in the empty space without holding the CTRL key will clear the selection.

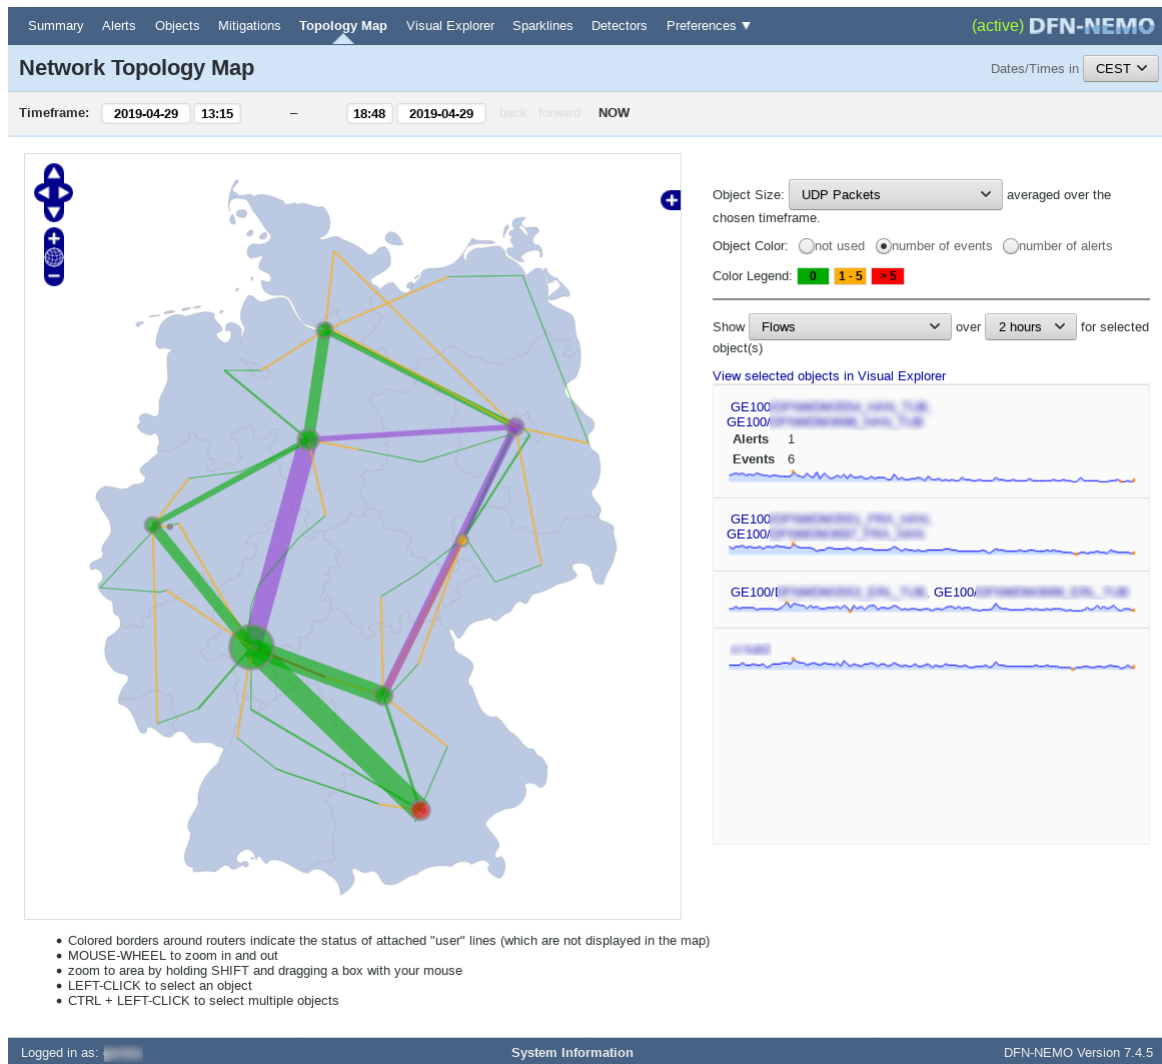


Figure 19: Visualization of object details in the topology map

4.7 Visual Explorer

The [Visual Explorer](#), accessible by a link of the same title in the navigation area or other links named alike within the NeMo application, is used for a graphical analysis of network traffic. For this purpose, indicator trends and traffic models of multiple objects are plotted next to each other, as shown in Figure 20.

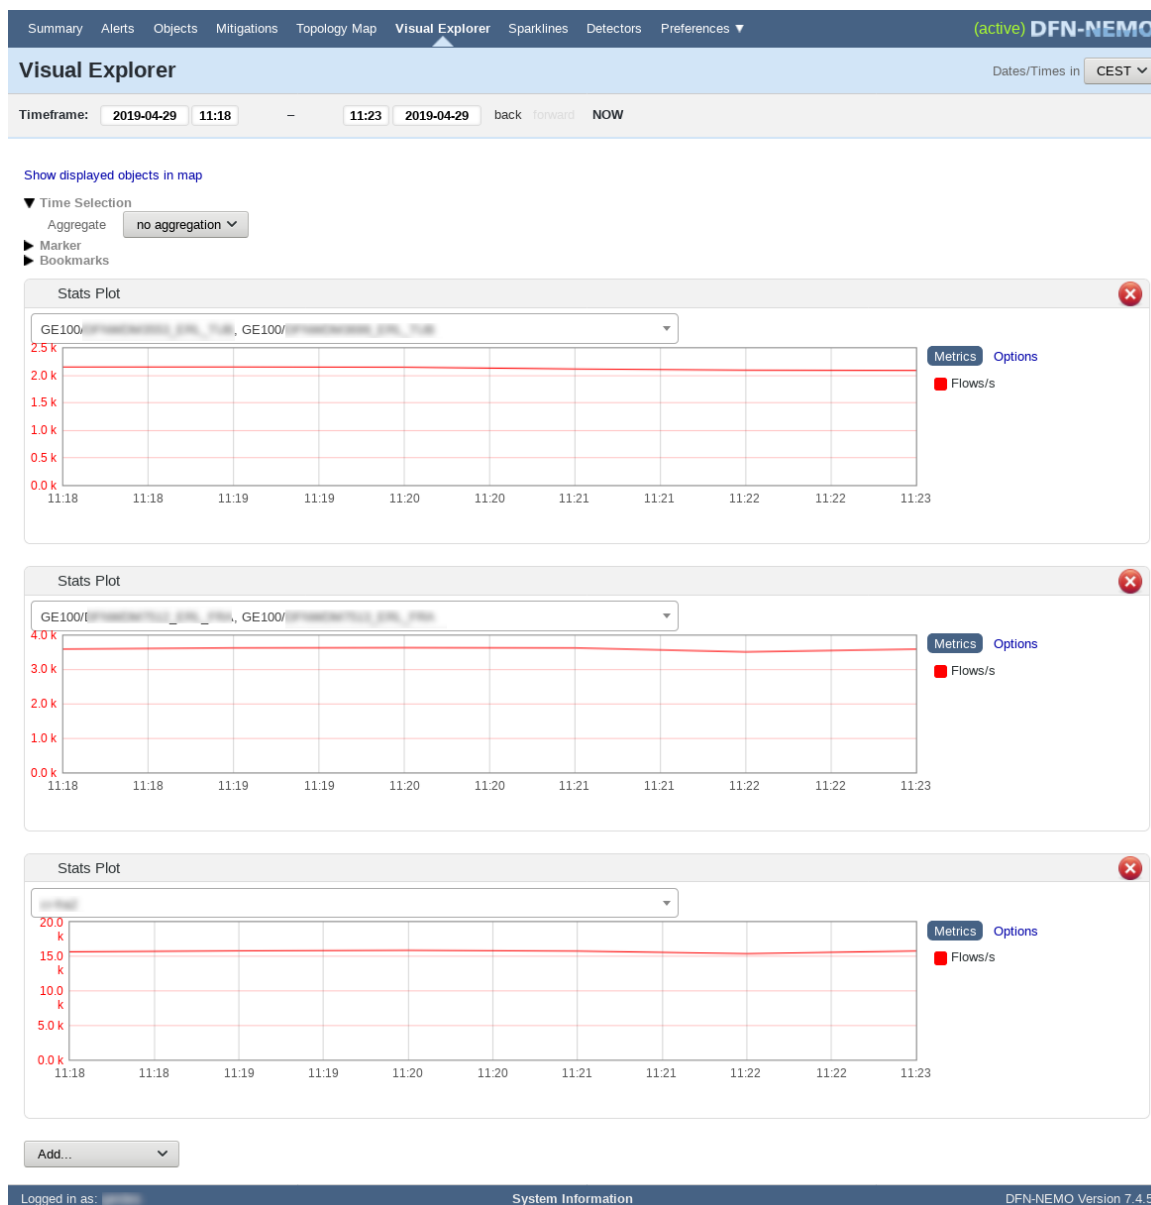


Figure 20: Visual Explorer

Below controls to define a timeframe for the analysis (similar to the controls for the topology map in Section 4.6), controls to adjust the visualization and to create bookmarks are placed. Bookmarks comprise settings and selected objects to an URL to being shared or stored and accessed at any later time. The link “[Show displayed objects in map](#)” will open the topology map with exactly the objects selected for plots in the Visual Explorer.

The control section is then followed by the visualization area with [indicator plots](#) (“[Stats Plots](#)”) and [model plots](#) (“[Detector Model Plots](#)”). Each plot consists of a box with a heading

in a grey colored background. The box contains the concise graph with indicator or model trend, respectively, and controls to adjust the visualization of the graph.

If the [Visual Explorer](#) is launched via the menu item in the navigation area one indicator plot is created showing the trend of the number of flows per second on the first object in the object database based on alphabetic sorting. If the Visual Explorer was launched by clicking on a link in the topology map, this plot is omitted, but a plot is created for the object and indicator as selected in the topology map. More plots can be created using the drop down menu at the bottom of the page.

4.7.1 Adjustments of the visualization and bookmarks

The visualization can be adjusted using the following settings:

The color associated with the indicators The color scheme can be changed using the submenu item “[Change Metric Colors](#)” within the “[Preferences](#)” menu item, accessible via the navigation area. A description of the configuration can be found in Section 4.10.

Aggregation of values, Markers and Bookmarks There are three ways for aggregation of values, each can be (un-)folded by clicking the triangle¹⁷.

Time Selection This section enables smoothing the graph by aggregating sequential values and hereby reducing the data set for long timeframes to detect important trends within the graphs. The default is not to aggregate (“[no aggregation](#)”), all graphs are rendered using all available data, whereas the resolution of this data depends on the type of data (e.g. one value per minute for NetFlow data, one value for five minutes for SNMP data). All other options in the drop down menu will aggregate sequential values within a predefined interval and this average value will be used for the graph.

Figure 21 shows an aggregation by five minute intervals. Therefore, NetFlow data that has been recorded every minute will be aggregated and the average value is shown in the graph. E.g. the graph will show the value for the values 12:00 until 12:04 at 12:00 and for values 12:05 until 12:09 at 12:05. The graph is then rendered by connecting these with direct lines.

Marker The section “[Marker](#)” allows to highlight parts of the graphs. Such markers are identified with a [name](#), a start time and (optionally) an end time. When the marker is created by clicking on “[Add Marker](#)” the background of the graph is colored in violet for the time period set for the marker. If no end time is specified only the start time is highlighted. If marker overlap they cannot be distinguished in the graph.

Below the form to create a new marker, the markers that have already been created are listed, where the red cross can be clicked on to remove that particular marker. There are two markers defined in Figure 21, one for a time frame and one for a single value.

¹⁷If the triangle points to the right, the section is folded, if the triangle points south the section is unfolded.

Bookmarks This section allows to create an URL comprising of the displayed graph and optionally the time frame and list of markers. These URLs can be shared by e.g. email allowing the recipient—if he or she has the required access rights—to access the same information and visualization the sender (creator) had at time of creating the bookmark. The two checkboxes “[include time](#)” and “[include markers](#)” include the displayed time frame and markers, respectively, in the creation of the URL. The URL/bookmark can be created using the button labeled “[Generate Bookmark URL](#)” and will thereafter be displayed in the textfield underneath.

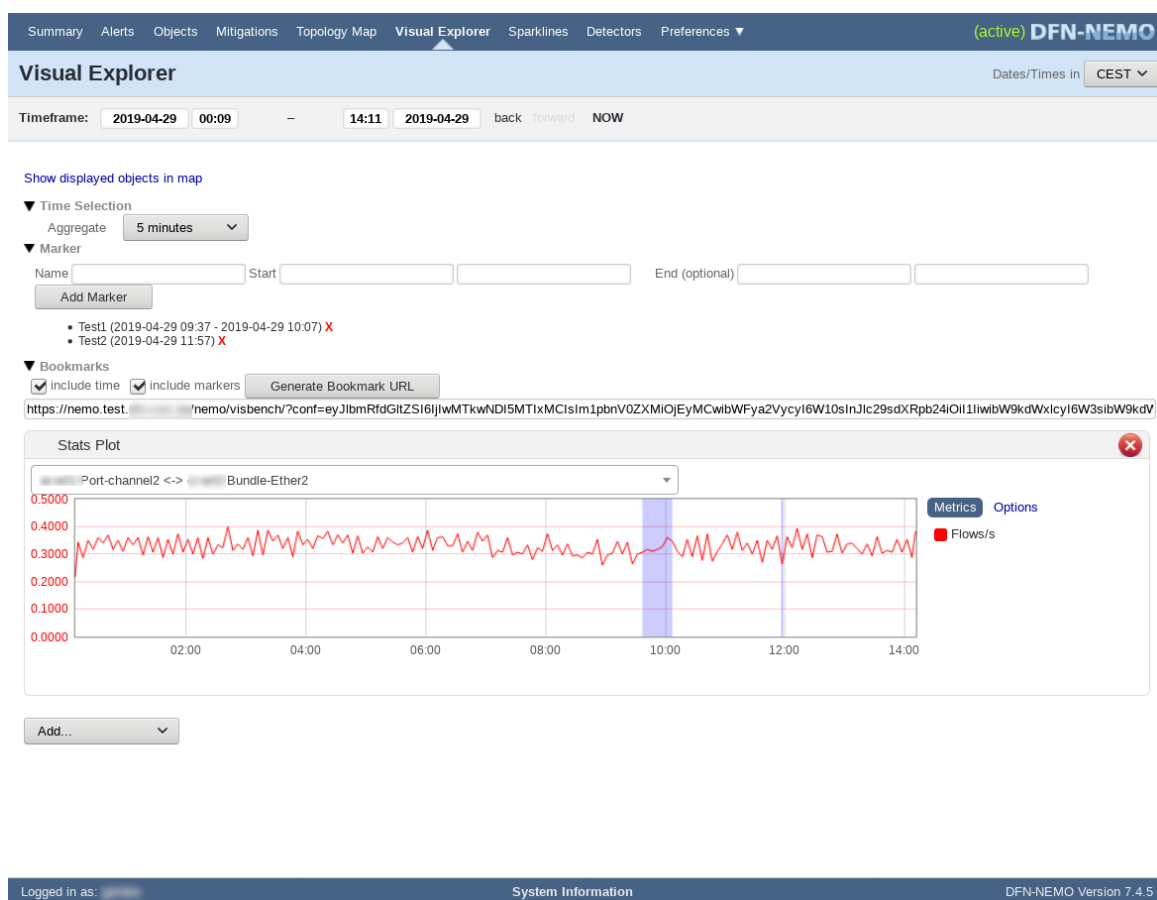


Figure 21: Aggregation, marker and bookmarks of the Visual Explorer

4.7.2 Manipulation of the plots

In the display area of the [Visual Explorer](#) indicator and model plots for various objects can be added. Each plot has the same structure, which is indicated in Figure 22:

Title Each plot has a title with a grey background. The title also shows the type of the plot:

Stats Plots used for indicator plots showing the trend of various indicators of objects within the Network.

Detector Model Plot used for model plots showing the normal behavior of the traffic at various objects. The normal behavior was automatically generated using machine learning.

On the right hand side of the title is a red dot with a white “X” within. A click on this icon will remove the plot from the display area.




Figure 22: Stats Plot in the Visual Explorer: The graph contains the indicators “UDP packets per second” and “Bytes per second divided by packets per second”. The number of UDP packets is colored in violet, as is its value range from “0.0K” to “8.0K”. The average packet size is displayed in green/yellow, its value range in green from “0.0” to “1000.0”.

Visualization The visualization for stats plots vary slightly from the one for detector model plots as the detector model plots allow to select the normal behavior to be investigated (e.g. the trend of the number of UDP packets per second on an object). Both types share a control item to select the object to display, consisting of a text field and a button labeled with an arrow pointing south. When the button is clicked, a list of all objects to be selected within the DDoS application appears. The list is grouped by the type of the object. Objects may also be selected by specifying parts of their name in the text field. The list is then reduced to objects containing all specified words (separated by spaces). Figure 23 shows this for input “cert”. An object can be selected by clicking with the mouse or by navigating the list of objects using the cursor keys and confirming the selection by hitting the enter key.

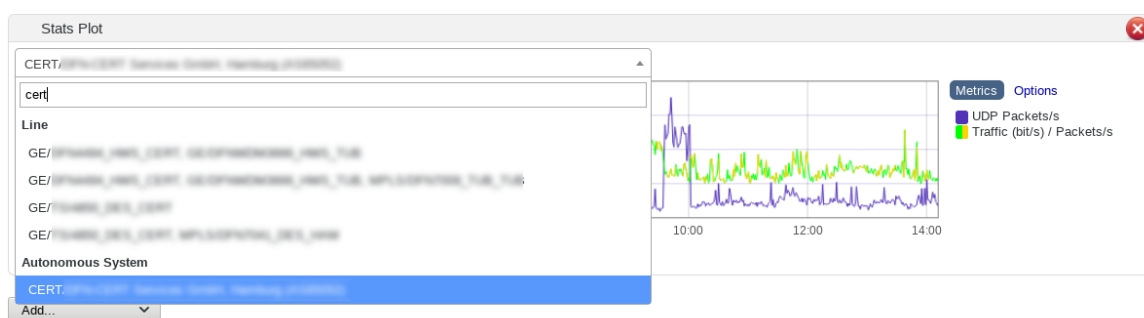


Figure 23: Selection of an object for a plot

Below the object selection the graph is displayed, showing the chronological trend of selected metrics for the time frame selected. The horizontal and vertical axis show times and ranges of values, respectively. Small values are listed with four positions after the decimal point, large values are shown only with a single position after the decimal point. Additionally, an appropriate suffix for the units is selected (“K”

for “Kilo”, 10^3 ; “M” for “Mega”, 10^6 ; “G” for “Giga”, 10^9). If several metrics are displayed, multiple vertical value ranges are displayed as well. Metrics and value ranges are colored using the same color. If a multiple colored indicator is selected, its first color is used for the value range coloring.

If the mouse pointer is within the graph (as seen in Figure 24 a vertical line is shown following the position of the mouse pointer and indicating the position within the graph. The exact time is displayed at the horizontal axis. If the mouse pointer is located on a point used for aggregation for any of the graphs, the values for each graph at this time are displayed.




Figure 24: Display of time and indicator values at mouse pointer

4.7.2.1 Changing the time frame on display Any plot allows to adjust the displayed time frame by clicking and moving the mouse pointer. This is especially helpful when details of a plot are to be investigated. Clicking and holding the main mouse key within the plot and followed by movements of the mouse pointer, the part of the plot that is marked by this is colored blue, as seen in Figure 25. When the left mouse key is let got, the selected time frame is transferred to the controls and subsequently displayed in all plots.





Figure 25: Selection of a time frame in a stats plot


A similar effect can be achieved using the mouse wheel. While the “ALT” key is pressed and the mouse wheel is spinned upwards, the plot is zoomed in at the location of the mouse pointer. Conversely, the plot is zoomed out, when the mouse wheel is turned downwards while pressing the “ALT” key allowing a broader view. Keep in mind that the key and wheel direction might be different depending on the operating system and web browser.

4.7.2.2 The options menu On the right hand side—next to the plot—control options are placed to manipulate the plot. There are two registers that allow to change the appearance of the graph: Firstly, there is a “Metrics” tab that enables the user to add and remove metrics to and from the graph. Secondly, there is the “Options” register controlling some options of the visualization. The default is to show the Metrics register.

4.7.2.2.1 Hiding the options menu First of all, the options panel can be hidden when clicking on the right border of the graph, as seen in Figure 26a and Figure 26b. When the mouse pointer is placed on the border, the border's size will be enhanced and a slim grey button will be displayed. When the options panel is hidden, the visualization of the plot will use the freed space.



(a) The grey button to hide the options panel



(b) The grey button to show the options panel

Figure 26: Hiding and showing the options panel

4.7.2.2.2 Adding metrics to the plot One option is to add metrics to the plot or remove metrics. A link labeled “[Add Metric](#)” will pop up while the mouse pointer is located on the options menu and the register is switched to “[Metrics](#)”, see Figure 27.






Figure 27: The [Add Metric](#) link is shown once the mouse pointer is on the options menu

Clicking the link will show a drop down menu as seen in Figure 28a, that lists all available metrics with its currently configured coloring scheme¹⁸. If any entry in the list is clicked, a graph for this metric is added to the plot.

¹⁸See Section 4.10.1 for configuration of the coloring scheme.



(a) Adding metrics to the graph



(b) Colored graphs if relative

Figure 28: Adding colorful metrics

Metric values used for this visualization are based on NetFlow data and SNMP data, using each router's individual sampling rate and then converted to average values per second.

There is one more deviation between [Stats plots](#) and [Detector model plots](#) which comes into effect when selecting an additional metric. If the mouse pointer is located on a [Stats plot](#) (as in Figure 28a), another link labeled “X/Y” pops up. If this link (“X/Y”) is clicked for two entries subsequently (e.g. first clicking on “X/Y” at “[Traffic](#)” and then on “X/Y” at “[Packets](#)”) the relation of these two metrics is added as a graph. The graph is colored one part after another in each of both metric's color, as visualized in Figure 28b for the relation of [Bytes](#) to [Packets](#). Some frequently used relations (average size of packets, number of source to number of destination ports, number of source to destination IP addresses, number of SYN packets to number of ACK packets) are already predefined in the list of indicators.

There is also a legend that lists all visualized metrics (or relation of metrics) and each graph's coloring. For relations the left part of the coloring scheme is the color for the numerator, the right part is colored in the same color as the denominator, the entry for the name lists first the name of the metric selected as numerator, followed by a slash and then the name of the metric selected as denominator.

4.7.2.2.3 Changing the appearance of the plot If the link labeled “Options” is clicked, the aforementioned legend disappears and controls to change the appearance of the plot pop up, as displayed in Figure 29.





Figure 29: Control settings for plots

These options are available with the following intentions:


Plot height This setting specifies the height of the graphic in pixels. The default value is 200 pixel.

Multiple Axes This setting controls whether multiple axes are displayed for the graphs, which is especially valuable if graphs are in value areas that have huge differences. If two graphs use the same scale, the one with smaller values will not show any significant value for interpretation, as the larger values of the other graph dominate the scale. A common axis is helpful if several values are to be compared directly, i.e. with [Detector model plots](#).

Y-Axis scale There is a checkbox labeled “[true zero point](#)”. If ticked, all axis will start with value **zero** instead of somewhere close to the smallest value within the graph. This helps a lot when comparing graphs, but variations for each graph are much better visible when the checkbox is unticked. See Figure 30 for a comparison.



(a) True zero point enabled



(b) True zero point disabled

Figure 30: Comparison of enabling (top) and disabling (bottom) a true zero point

4.7.3 Adding plots

At the bottom of the content area south of all plots a drop down menu is located which is to be used to add plots to the [Visual Explorer](#). Depending on the choice additional [Stats Plots](#) and [Detector Model Plots](#) are added subsequently.

We will now describe both plot types and explain their various options.

4.7.3.1 Detector Model Plots First a detector must be selected, before metrics for selected objects can be added. Detectors are explained in detail in Section 4.9. Detectors are part of the configuration of NeMo. [Detector Model Plots](#) require the detector to create a model ([Holt-Winters](#) or [Static deviation](#), see Section 4.13.2). There are multiple options for each of the detector models (see Tables 4 and 5), while metrics can also be displayed in relation to each other.

Metric	Description
Corridors	Defines a corridor for accepted values. Values detected outside of the corridor range raise events
Variance	Machine learned variance of detected values
Smoothed/s	Machine learned average values of detected values
Seasonsmoothed	Machine learned seasonal (weekly) variation of detected values
Predictions/s	Prediction of future values
Absolute values/s	The actual detected value

Table 4: Options for Holt-Winters detectors

Metric	Description
Corridors (I,W,C)	Defines corridors for accepted values. Values detected outside of the corridor range raise events of criticalities INFO (I) , WARNING (W) or CRITICAL (C)
Variance	Machine learned variance of detected values
Smoothed	Machine learned average values of detected values
Traffic / Packets	The average size of packets, as the relation of traffic volume to number of packets

Table 5: Options for static deviation detectors

4.7.3.2 Stats Plots A stats plot visualizes trends of one or multiple indicators. Indicators are explained in Tables 6 to 10.

Metric	Description
Flows/s	The number of flows per second.
Traffic (bit/s)	The transferred traffic in bit/s.
Packets/s	The number of packets per second.
TCP Packets/s	The number of packets for TCP protocol per second.
ICMP Packets/s	The number of packets for ICMP protocol per second.
UDP Packets/s	The number of packets for UDP protocol per second.
Incoming Packets/s	The number of incoming packets per second.
Incoming Traffic (bit/s)	The transferred incoming traffic in bit/s.
Incoming TCP Traffic (bit/s)	The transferred incoming traffic for TCP protocol in bit/s.
Incoming UDP Traffic (bit/s)	The transferred incoming traffic for UDP protocol in bit/s.
Incoming ICMP Traffic (bit/s)	The transferred incoming traffic for ICMP protocol in bit/s.
Outgoing Packets/s	The number of outgoing packets per second.
Outgoing Traffic (bit/s)	The transferred outgoing traffic in bit/s.
Outgoing TCP Traffic (bit/s)	The transferred outgoing traffic for TCP protocol in bit/s.
Outgoing UDP Traffic (bit/s)	The transferred outgoing traffic for UDP protocol in bit/s.
Outgoing ICMP Traffic (bit/s)	The transferred outgoing traffic for ICMP protocol in bit/s.
Traffic/Packets	The relation of Traffic to the number of Packets, $\frac{Traffic}{Number\ of\ Packets}$

Table 6: Options for Group Traffic for Stats Plots

Metric	Description
Src IPs/s	The number of different source IPs per second.
Dst IPs/s	The number of different destination IPs per second.
Src/Dst IPs	The relation of source to destination IPs, $\frac{Src\ IPs}{Dst\ IPs}$

Table 7: Options for Group IPs for Stats Plots

Metric	Description
Src Ports/s	The number of source ports per second.
Dst Ports/s	The number of destination ports per second.
Src/Dst Ports	The relation of source to destination ports, $\frac{Src\ ports}{Dst\ ports}$.

Table 8: Options for Group Ports for Stats Plots

Metric	Description
ACK Packets/s	The number of TCP packets with ACK flag enabled per second.
SYN Packets/s	The number of TCP packets with SYN flag enabled per second.
RST Packets/s	The number of TCP packets with RST flag enabled per second.
FIN Packets/s	The number of TCP packets with FIN flag enabled per second.
SYN/ACK Packets	The relation of the number of TCP packets with SYN flag enabled to the number of TCP packets with ACK flag enabled, $\frac{Number\ of\ TCP\ packets\ with\ SYN\ flag\ enabled}{Number\ of\ TCP\ packets\ with\ ACK\ flag\ enabled}$.

Table 9: Options for Group Flags for Stats Plots

Metric	Description
SNMP Traffic (bit/s)	The traffic volume in bits per second using SNMP data as source.
SNMP In Traffic (bit/s)	The incoming traffic volume in bits per second using SNMP data as source.
SNMP Out Traffic (bit/s)	The outgoing traffic volume in bits per second using SNMP data as source.
SNMP Packets/s	The number of packets per second using SNMP data as source.
SNMP In Packets/s	The number of incoming packets per second using SNMP data as source.
SNMP Out Packets/s	The number of outgoing packets per second using SNMP data as source.
SNMP In Errors/s	The number of incoming erroneous packets per second using SNMP data as source.
SNMP Out Errors/s	The number of outgoing erroneous packets per second using SNMP data as source.

Table 10: Options for Group SNMP for Stats Plots

4.8 Sparklines

The sparkline view allows to visually compare an indicator on multiple objects to detect network traffic trends either shared between these or detectable on a subset of objects only. Contrary to the [Visual Explorer](#) the visualization is reduced in detail allowing a large number of objects to be displayed, to enhance comparability and to highlight parallel trends. Figure 31 shows an example of a sparkline view.

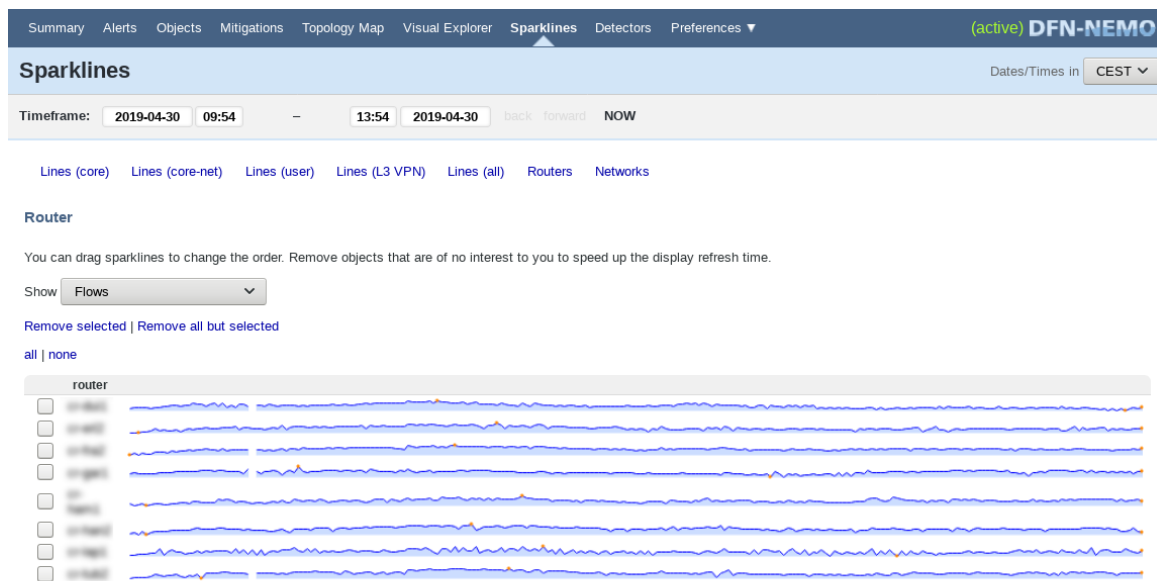


Figure 31: Sparkline view (part view)

The view comprises of three sections: the top section to select a time frame, see Section 4.8.2, a section with controls for selecting a combination of objects and object categories and to select indicators to be displayed, see Section 4.8.3 and the bottom section displaying the selected sparklines, see Section 4.8.4. But first, the construction of a sparkline is described in Section 4.8.1.

4.8.1 Constructing a sparkline

Each sparkline visualizes the trend of a metric that may change over time over a predefined time frame in a qualitative way. The objective of this visualization is to show the trend of a metric without consuming much space. The principle of construction is displayed in Figure 32.




Figure 32: Constructing a sparkline

To identify the trend of an indicator at one object over a defined period of time the times when maximum and minimum values are reached are determined. In this case¹⁹ the minimal value with 7.9 K flows at 7:38 and the maximum value with 17.9 K flows at 10:53. This area, defined by the time period and the extremal values, and colored in blue in Figure 32, is then scaled to the height of a text line and presented as qualitative trend. The web application highlights extremal values and the most recent values with red dots for orientation purposes.




Figure 33: Comparison of sparklines with different resolution

To highlight the differences between data in various resolutions, a different visualization for sparklines with fewer than one data value per minute is used. While the construction principle is similar, the sparkline is not rendered as a continuous line, rather as a sequence of not connected impulses as defined by the pattern of data collection. E.g. for SNMP data collected in five minute intervals a sparkline as in Figure 33 is rendered.

4.8.2 Time frame selection

The time frame can be configured as explained in Section 4.7.2.1.

4.8.3 Controls for the selection of objects and indicators

In the sparkline view indicator trends are visualized on similar objects. Which type these objects are of can be determined by the links at the top of the page. These include object types and object categories (see Section 4.4): . South of these links the currently visualized object type is noted. The default is to show the .

Next are controls for the selection of an indicator via a drop down menu. The default is to show the number of flows. All indicators are described in Table 11.

¹⁹In Figure 32 this is the number of flows on the router xr-des1 for four hours.

Indicator	Description
Flows	The number of flows per second
Traffic	The amount of network traffic in bit per second
Packets	The number of packets per second
TCP Packets	The number of TCP packets per second
ICMP Packets	The number of ICMP packets per second
UDP Packets	The number of UDP packets per second
Incoming Packets	The number of incoming packets per second
Incoming Traffic	The amount of incoming traffic in bit per second
Incoming TCP Packets	The number of incoming TCP packets per second
Incoming UDP Packets	The number of incoming UDP packets per second
Incoming ICMP Packets	The number of incoming ICMP packets per second
Outgoing Packets	The number of outgoing packets per second
Outgoing Traffic	The amount of outgoing traffic in bit per second
Outgoing TCP Packets	The number of outgoing TCP packets per second
Outgoing UDP Packets	The number of outgoing UDP packets per second
Outgoing ICMP Packets	The number of outgoing ICMP packets per second
Traffic/Packets	The average size of packets by dividing the amount of traffic in bit by the number of packets
Src IPs	The number of different source IP addresses per second
Dst IPs	The number of different destination IP addresses per second
Src/Dst IPs	The relation of source to destination IP addresses
Src Ports	The number of different source ports per second
Dst Ports	The number of different destination ports per second
Src/Dst Ports	The relation of source ports to destination ports
ACK Packets	The number of TCP packets with ACK flag set
SYN Packets	The number of TCP packets with SYN flag set
RST Packets	The number of TCP packets with RST flag set
FIN Packets	The number of TCP packets with FIN flag set
SYN/ACK Packets	The relation of TCP packets with SYN flag set to TCP packets with ACK flag set
SNMP In Traffic	The amount of incoming traffic in bit per second using SNMP data
SNMP Out Traffic	The amount of outgoing traffic in bit per second using SNMP data
SNMP In Packets	The number of incoming packets per second using SNMP data
SNMP Out Packets	The number of outgoing packets per second using SNMP data
SNMP In Errors	The number of incoming erroneous packets per second using SNMP data
SNMP Out Errors	The number of outgoing erroneous packets per second using SNMP data
Routing Processor	
Switching / LC Processor(s)	

Table 11: Indicators available for sparkline creation

As soon as an update of the view is triggered (e.g. changing the time period, object type,

object category or indicator) the sparklines and graphics are updated. Especially for lines this update process might take a while, therefore it is recommended to reduce the number of objects by excluding irrelevant objects or categories from the listing, as this might reduce the amount of time required to render the graphics considerably.

4.8.4 Displaying sparklines

The selected sparklines are displayed in a table based layout. For each object its name and a sparkline are shown as configured by the aforementioned controls. Figure 34 shows this. The name of the object is a link to a [Visual Explorer](#) that will show the combination of object, indicator and time period. While graphics are updated, the text “Loading...” is displayed instead of each graphic.




Figure 34: Sparklines in table based layout

To reduce the number of sparklines in this view, the user can remove sparklines by ticking checkboxes at the start of each entry line and delete these selections. There are two links at the top of the listing to enhance the selection of sparklines “all” and “none” by ticking checkboxes of all objects or of no object, respectively. There are also two options to remove objects from the list:

Remove selected which will remove the objects for which the checkbox was ticked

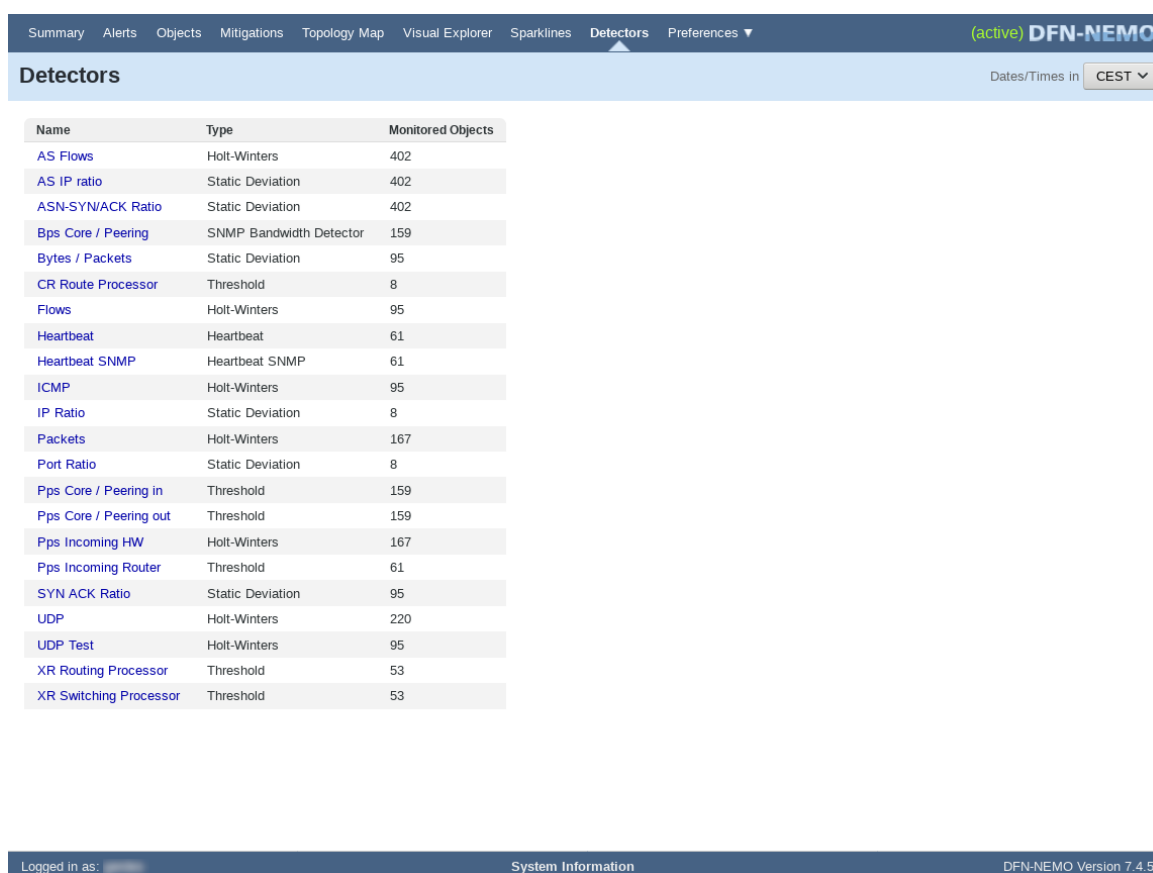
Remove all but selected which will remove all other objects (except those for which the checkbox was ticked).

Any will result in an improved loading of graphics.

To improve comparability of sparklines—or the traffic pattern they visualize—the order within the listing can be changed. The mouse pointer is positioned on the line to be moved and while pressing the mouse button the mouse pointer is moved to the position where the sparkline shall be displayed.

4.9 List of all detectors

A list of all currently supported detection methods (detectors) is accessible via the link “[Detectors](#)” in the navigation area. All activated detectors are directly accessible. The detectors are listed in alphabetic order and displayed with name, type and number of monitored objects for each detector (see Figure 35). The name is a link to more detailed information about each detector.



The screenshot shows the DFN-NEMO interface with the 'Detectors' tab selected. The page title is 'Detectors' and it includes a 'Dates/Times in' dropdown set to 'CEST'. Below the header is a table listing 25 detectors. Each row contains the detector name (a clickable link), its type, and the number of monitored objects. The detectors are listed in alphabetical order.

Name	Type	Monitored Objects
AS Flows	Holt-Winters	402
AS IP ratio	Static Deviation	402
ASN-SYNACK Ratio	Static Deviation	402
Bps Core / Peering	SNMP Bandwidth Detector	159
Bytes / Packets	Static Deviation	95
CR Route Processor	Threshold	8
Flows	Holt-Winters	95
Heartbeat	Heartbeat	61
Heartbeat SNMP	Heartbeat SNMP	61
ICMP	Holt-Winters	95
IP Ratio	Static Deviation	8
Packets	Holt-Winters	167
Port Ratio	Static Deviation	8
Pps Core / Peering in	Threshold	159
Pps Core / Peering out	Threshold	159
Pps Incoming HW	Holt-Winters	167
Pps Incoming Router	Threshold	61
SYN ACK Ratio	Static Deviation	95
UDP	Holt-Winters	220
UDP Test	Holt-Winters	95
XR Routing Processor	Threshold	53
XR Switching Processor	Threshold	53

At the bottom of the interface, there is a status bar showing 'Logged in as: [user]', 'System Information', and 'DFN-NEMO Version 7.4.5'.


Figure 35: A listing of all activated detectors

4.9.1 Detailed view of a detector

If a name is clicked in the listing of all detectors the detailed view for the selected detector is shown, see Figure 36.



At the top the name of the detector and its type²⁰ is displayed. On the left hand side of the next section the default configurations for this detector model are listed. On the right hand side two lists of objects are shown, for which this detector model is configured. First, all objects are listed for which an alternative configuration was selected, second, all objects are listed for which a default configuration was set up. For each list the number of objects is presented in the corresponding heading. If no default configuration was created for this detector, the left hand side is empty as well as the list at the top of the right hand side, see Figure 37. The object name in the listing is a link to this object's detector configuration, which will be discussed as part of the detailed object view in Section 4.12.



²⁰The various types of detectors are explained in Section 4.13.2.

4.10 Settings

The user interface of NeMo’s web interface can partly be configured by the user. These settings are accessible via the “[Preferences](#)” menu item in the [navigation](#) area. The unfolded menu has an item named “[Change Metric Colors](#)”—see Figure 38—allowing to configure the coloring scheme used for various indicators throughout the web interface.




Figure 38: The unfolded menu item for the Preferences menu

4.10.1 Changing the indicator color

Throughout the web application each indicator is colored identical. Using the aforementioned “[Change Metric Colors](#)” link a listing of indicators as in Figure 39 is shown.




Figure 39: Changing the coloring scheme for indicators (part view)

For each indicator an input field is located showing the currently selected color specification

as input text and the color as background color. The coloring specifications are taken from the RGB color space²¹ and are noted as “#RRGGBB”. The letters represent a two-place hexadecimal representation of the red (RR), green (GG) and blue (BB) part of the chosen color. Each part is between 0 (hexadecimal 00) and 255 (hexadecimal FF). The base colors red, green and blue are represented by the color codes #FF0000, #00FF00 and #0000FF. Darker colors have lower values like #800000 for a darker red. Mixed colors can be created by mixing the various colors, i.e. magenta with #FF00FF. Black is #000000, white #FFFFFF. If a valid color code is presented, the input field changes its background color to the color specified.

At the bottom of the page two buttons are placed to store the coloring scheme or reset all values to default values.

²¹Within the RGB color space colors are specified by their red, green and blue parts at additive color mixing: https://en.wikipedia.org/wiki/RGB_color_space

4.11 System information

To assess the state of the detection system and NeMo's servers important information is listed on the page "[System Information](#)" which is accessible via a link at the center of the [status information area](#) at the bottom of each page.

At the top of this page (see also Figure 40) various information about the operativeness of detection and analysis functions of this service is listed.

[Summary](#)
[Alerts](#)
[Objects](#)
[Mitigations](#)
[Topology Map](#)
[Visual Explorer](#)
[Sparklines](#)
[Detectors](#)
[Preferences ▼](#)

(active) DFN-NEMO

System Information Dates/Times in CEST ▼

Erkennung

Load 1.15, 0.72, 0.67

- ✓ condensed running
- ✓ eventd running
- ✓ mail system running
- ✓ last updated object: today, 10:39
- ✓ last detector run: today, 10:39
- ✓ all eventd threads running

Analyse

Load 0.26, 0.13, 0.10

- ✓ pure-ftpd running
- ✓ database running
- ✓ last data per FTP: today, 10:40
- ✓ last db access: today, 10:40

Update Timestamps

Last GIS data 2019-05-02 04:00
Last topology update run 2019-05-02 09:15

Last Topology Update Status

Core router OK
Significant router changes OK
Significant line changes OK
Significant AS changes OK

Database Table Stats

(Based on database statistics, not exact counts)

Line	48,420,300
Router	357,065
Net	44,634
Autonomous system	17,622,800
Router group	0
Filter	178,467

Analyse System

Memory

	total	used	free	shared	buffers	cached
Mem:	64328	63958	370	2125	645	56280
-/+ buffers/cache:		7032	57296			
Swap:	32764	0	32764			

Disk Space

Filesystem	Size	Used	Avail	Use%	Mounted on
devtmpfs	32G	8.0K	32G	1%	/dev
tmpfs	32G	16K	32G	1%	/dev/shm
tmpfs	32G	1.3M	32G	1%	/run
tmpfs	32G	0	32G	0%	/sys/fs/cgroup
/dev/sda3	50G	3.9G	43G	9%	/
/dev/sda4	5.0T	1.8T	3.2T	36%	/services/data
tmpfs	6.3G	0	6.3G	0%	/run/user/26

Number of pure-ftpd Processes

11

11 Database Processes

postgres	1585	0.0	0.0	45148	4920	?	Ss	Jan08	3:42	/usr/lib/systemd/systemd --user
postgres	1587	0.0	0.0	83544	1528	?	S	Jan08	0:00	(sd-pam)
postgres	3358	0.0	0.1	2254480	77316	?	S	Jan08	0:16	/usr/lib/postgresql96/bin/postgres -D /services/data/postgres/data
postgres	3448	0.0	0.0	68968	3160	?	Ss	Jan08	0:00	postgres: logger process
postgres	3449	0.1	3.2	2256816	2150892	?	Ss	Jan08	248:20	postgres: startup process recovering 00000001000004E9000000B2
postgres	3480	0.1	3.2	2258388	2137368	?	Ss	Jan08	164:20	postgres: checkpointing process
postgres	3481	0.0	3.2	2255024	2125340	?	Ss	Jan08	6:54	postgres: writer process
postgres	3579	0.0	0.0	70016	4000	?	Ss	Jan08	46:02	postgres: stats collector process
postgres	3609	0.0	0.0	2256728	30792	?	Ss	08:40	0:00	postgres: nemo nemo ::1(35322) idle
postgres	4979	0.0	0.0	2268040	5476	?	Ss	Jan08	159:53	postgres: wal receiver process streaming 4E9/B2CF94A0

Logged in as:

System Information

DFN-NEMO Version 7.4.5

Figure 40: Page of system information

4.11.1 Load

The average load is stated for both systems (detection and analysis). The load is a measure for the computational work to be executed by the computer system. The average load is the average system load over a period of time. The three numbers indicate the load for the last one, five and fifteen minute periods²².

4.11.2 Status of necessary processes

The status of required processes is listed for both systems. Also the last successful access to data is given. If everything is fine, a green check mark will also be listed for each process. If an error was detected, a red cross will be positioned left to the particular service's name. The following items are listed:

condensed running / not running This states whether the data processing for NetFlow and SNMP data is working.

last updated object The time for the last update of any line by the data preprocessing.

Eventd running / error in eventd States whether detection and alarming services are running or not.

Last detector run / all eventd threads running Time of the last run of the detector, and whether all processing steps of the detection are working. If an error is detected, the component in which the error was detected is also specified (detectors, alarm triggers or notification service via email).

Mail system running / has errors If errors at sending notifications via the email service are detected, the number of errors is specified here.

Pure-ftpd running / not running Informs whether the FTP component works and NetFlow and SNMP data are received.

Last data per FTP Time of the last successful transmission of data received at the FTP service.

Database running / not running States whether the database server is running.

Last db access States the last access to the database by the analysis component.

4.11.3 Timeliness of the Networkmodel

Next, two timestamps are given, that inform about the timeliness of the Networkmodel used in the application.

Last imported export Time when the data set used for the modelling was created.

Last topology update Time of the last update of the topology.

4.11.4 Most recent topology update

A list of data about the state of the most recent topology update.

²²For more information visit [https://en.wikipedia.org/wiki/Load_\(computing\)](https://en.wikipedia.org/wiki/Load_(computing)).

4.11.5 Data size

A statement of data size used by the services.

4.11.6 State of the server

Some information about the state of the server, like RAM usage, hard disk usage and various other indicators regarding the state of the system like number of FTP processes and a list of data base server processes.

4.12 Detailed view of an object

The detailed view of an object allows an enhanced investigation of traffic trends at this particular object. The detailed view is accessible from the objects overview, which is described in Section 4.4, when clicked on an object's name opening its detailed view. See Figure 41 for an example of a detailed view.

The extended navigation area contains the type of the object and its name as well as the time frame used for further displays²³.

The content area shows three registers, which will be described in detail in following subsections:

Details This is the default view and shows the categories specified for this object as well as an optional comment. There is also an overview of this objects' classification in the Network topology, all this is explained in Section 4.12.1.

Analysis This register allows a detailed analysis of the traffic detected on this object, see Section 4.12.2.

Detector Configuration This register shows the detectors configured for this object and also allows to configure these, see Section 4.12.3.

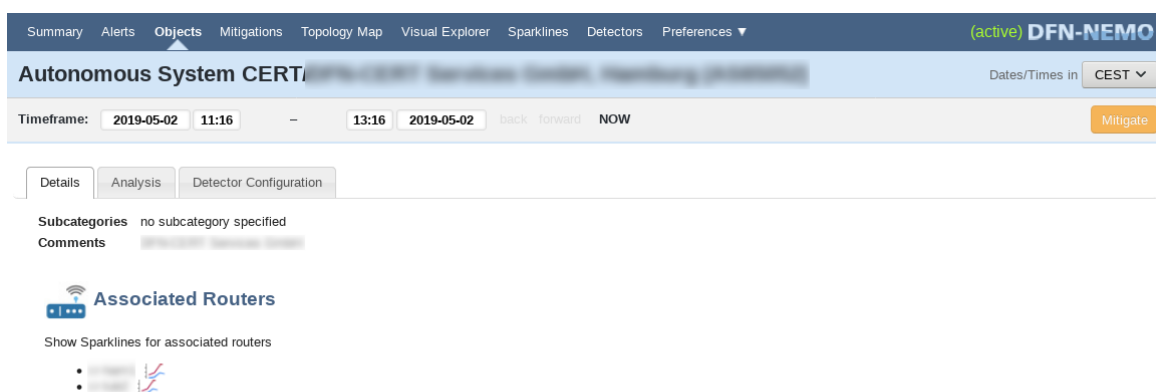


Figure 41: Detailed view of an object (part view)

4.12.1 Category and classification in topology

The register [Details](#) lists first the category and subcategories associated with this object. An optional comment is also listed, which may be specified with administrative privileges. If the object is not categorized or commented, these details are omitted. If no subcategories are specified for this object, the text “no subcategory specified” is printed.

Next is a section that lists connections in the topology to other objects associated with this object. This list differs depending on the type of the object:

Router For routers this consists of a list of connected lines and a list of connected networks and autonomous systems, which are monitored at this router.

Lines and Networks The detailed view for lines and networks shows a list of the routers this line connects to or those routers on which the network is monitored, respectively.

²³The time frame selected here is not used on the [Details](#) register, but on the other registers.

Autonomous System This lists the routers on which traffic data for this AS is collected.

If no objects are associated with this object, the text “[No associated object type](#)” is shown. Otherwise the list is titled with “[Show sparklines for associated object type](#)” which is a link to a sparkline view (see Section 4.8) of all objects associated with this object. Each entry in this list consists of the name of the object (linked to the detailed view of this object) and an icon that links to a Visual Explorer for this particular object.

4.12.2 Traffic Analysis

The traffic analysis view subdivides into three registers (see Figure 42, each will be discussed in detail in the following sections:

Target Details The default view offers several functions to analyse NetFlow data for this particular object, see Section 4.12.2.1.

Target Sparklines On this register sparklines of all basic indicators and commonly deduced indicators over the time period, see Section 4.12.2.2.

Visual Explorer This register shows a Visual Explorer of this particular object, see Section 4.12.2.3.

4.12.2.1 Target Details: Analysis of NetFlow data The register for analysis of NetFlow data divides into three section, as exemplary shown in Figure 42:

1. The first section shows a graph, that is very similar to a single indicator graph of the Visual Explorer of this object, see Section 4.12.2.1.1.
2. The second section allows to select routers as starting points for the analysis and to specify a NetFlow filter expression that will be used in the third section for analysis, see Section 4.12.2.1.2 and Section 4.12.2.1.3.
3. The third section presents multiple views on the data preselected in the second section, see Section 4.12.2.1.4.

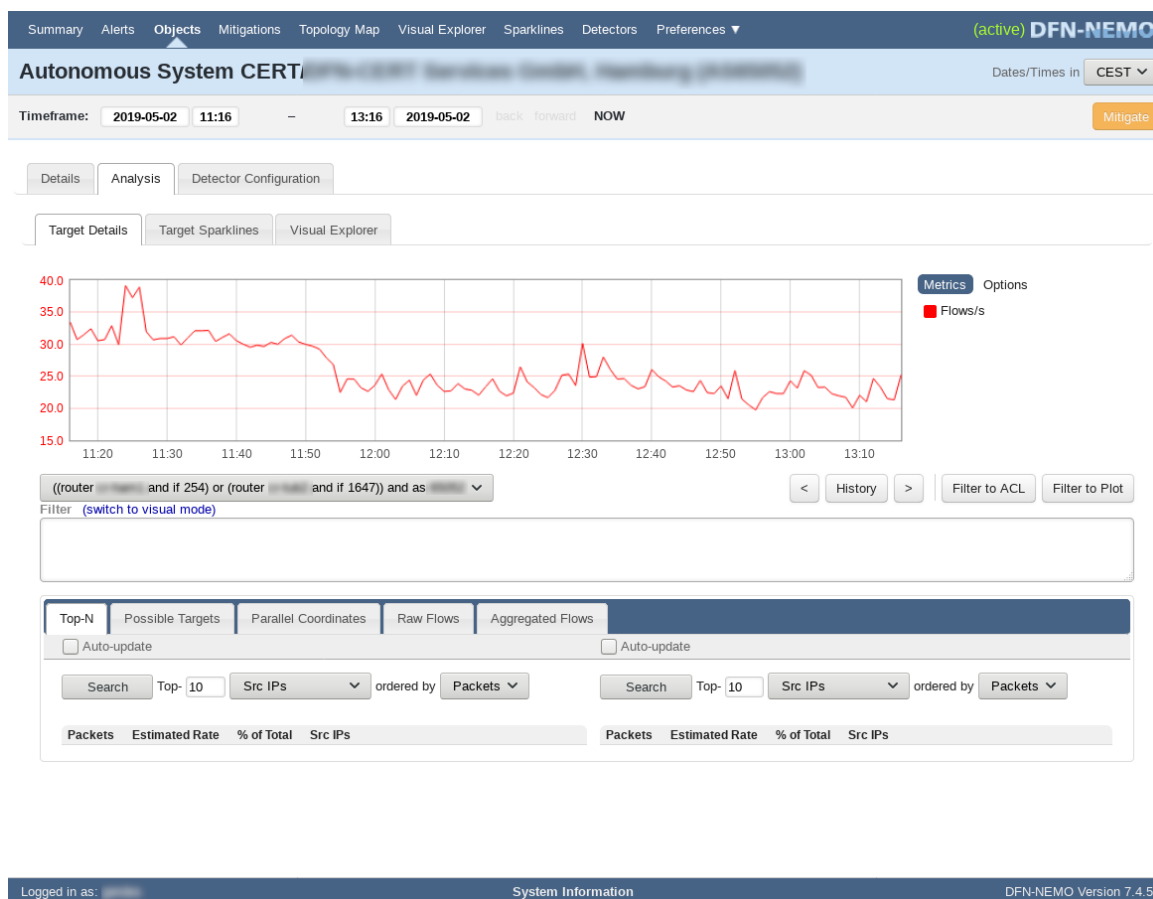


Figure 42: Analysis view of an object

The default is that all NetFlow data is used for analysis that was collected in the time period, specified in the extended navigation area. Contrary to a graph in the Visual Explorer the graph in the first section of this NetFlow analysis allows to specify a start time for analyses of the following sections.

4.12.2.1.1 Visual Explorer If the mouse pointer is within the graph's area the extended pointer is shown, marking the time and value of the indicator (if the mouse pointer is on the graph's line). If clicked, this time is used as the start time and marked with a bold red vertical bar as shown in Figure 43. Only those flows will be considered for analysis that occurred after this start time and before the specified end time. The bar can be removed by clicking the **X** at its top.




Figure 43: Start time of a NetFlow analysis

If the start time configured in the graph is not within the specified time period, this will be indicated by a directed mark at the border of the graph, see Figure 44. In this case the analysis will still consider all data between the start time configured in the graph until the end time specified in the navigation area.



Figure 44: Start time is before the specified time period

4.12.2.1.2 Selection of routers for analysis Above the filter definition in the second section and below the graph is a drop down menu located that can be used to filter the NetFlow data based on connected routers, see Figure 45. This selection defines NetFlow data used for further analysis.

As NetFlow data is collected at routers the NetFlow data is associated with the router. No NetFlow data is collected on lines, networks or for autonomous systems, so any analysis for these objects must be conducted on the NetFlow data of (all) to this object connected routers or at which this network or AS is monitored.

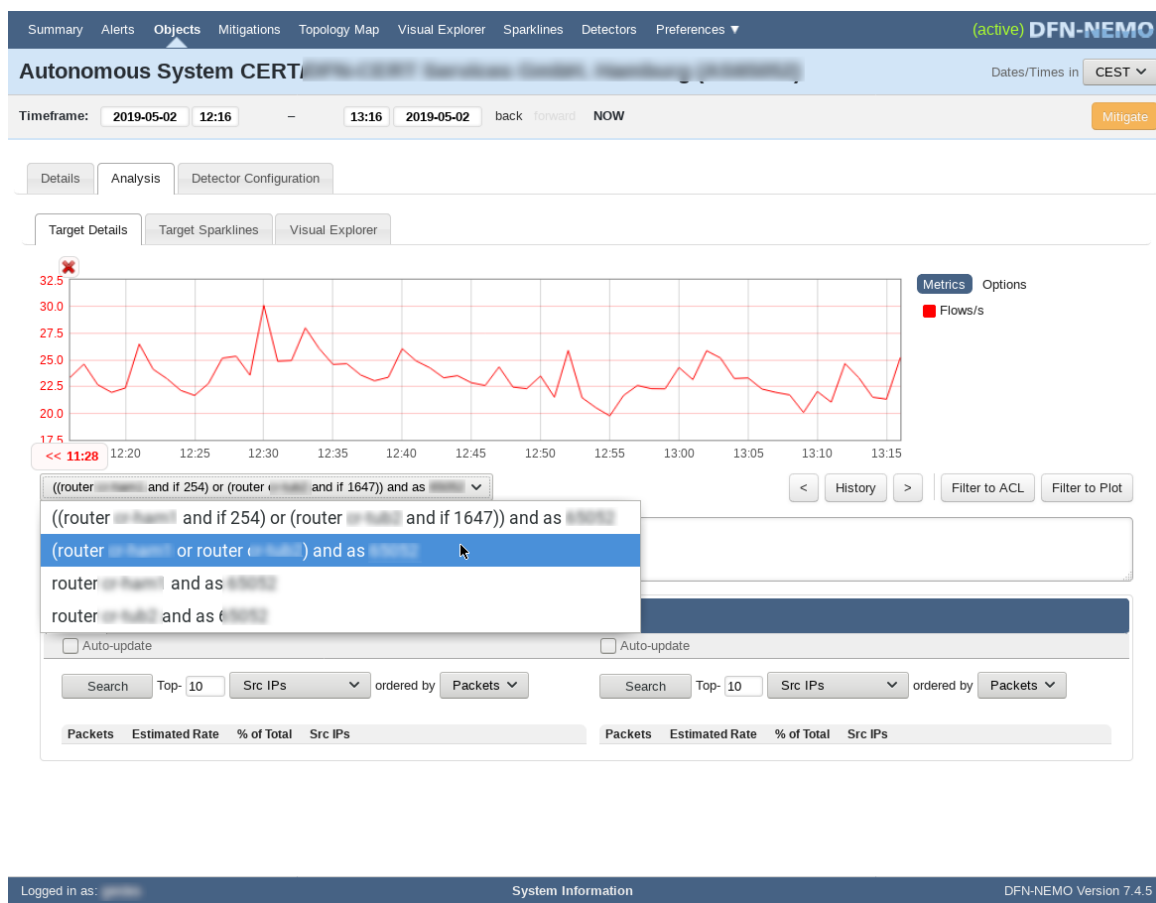


Figure 45: Selection of a router for a NetFlow analysis

The drop down menu lists for lines the routers this line is connected with, if only one is modeled in NeMo only that is listed. For analysis of networks or autonomous systems all routers are listed, that collected NetFlow data for the particular network, which requires a configuration for monitoring. Next to analyzing NetFlow data for each router on its own, a combination of all routers can be used as well.

4.12.2.1.3 Definition of filters The NetFlow filter in the second section can be used in two ways:

Text mode (default) The filter expression is specified in the text field using nfdump filter syntax²⁴.

Visual mode Alternatively the filter expression can be assembled using graphical input based on various modules.

The link to switch between the modes is above the text field and right to the word “Filter”, see Figure 42. Any click will switch between both modes, while transferring all input to the new mode. This includes checking it for correctness and conduct simple rewritings (i.e. deleting empty groups), yet the effect of the rewritten filter is the same.

²⁴Nfdump is a set of tools for the processing of NetFlow data. It was developed by Peter Haag at the Swiss Research Network SWITCH. <http://nfdump.sourceforge.net>.

In the visual mode filters consist of conditions or groups of conditions, whereas the latter contain conditions or groups of conditions. Conditions are added to the filter expression using the link labeled “[Add Condition](#)”, groups using the link labeled “[Add Group](#)”, respectively.

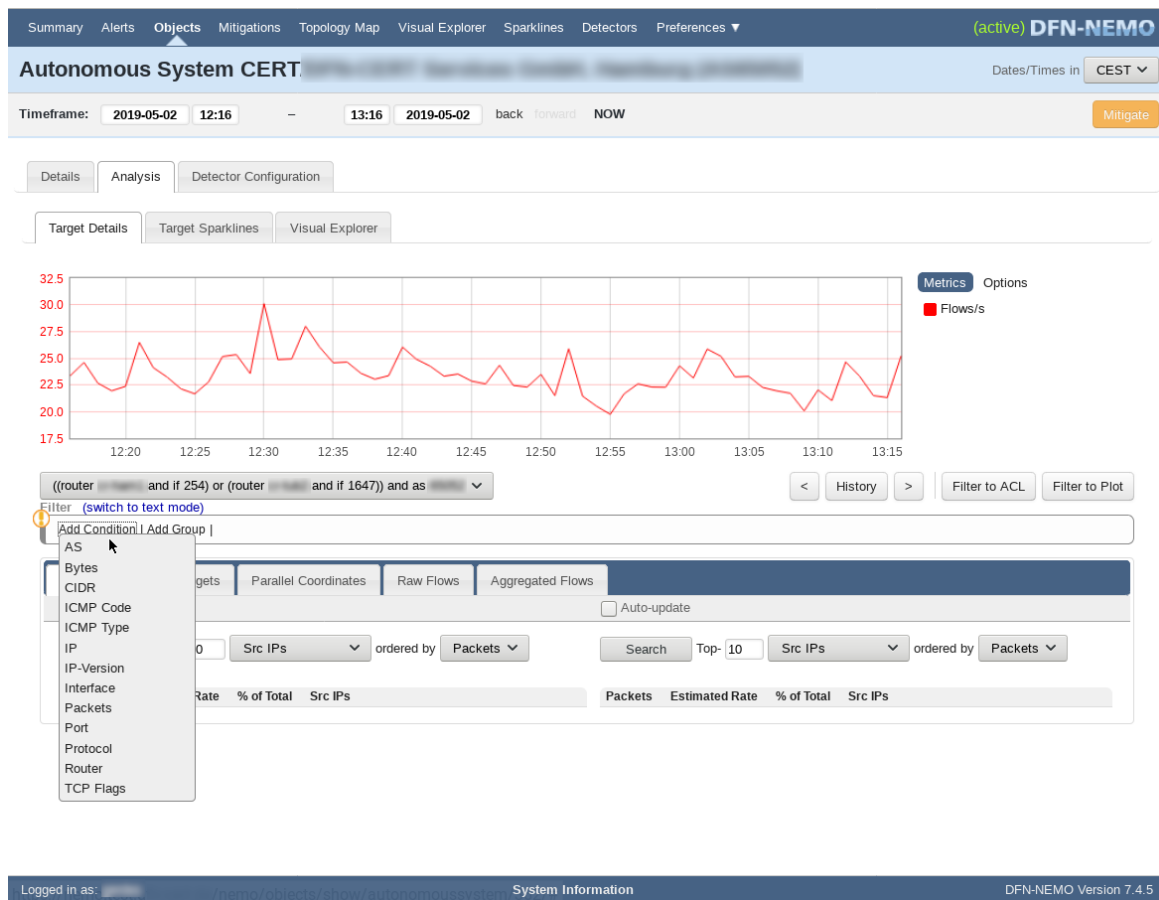


Figure 46: Available types of filters using the visual mode

When the “[Add Condition](#)” is clicked, a drop down menu is shown as seen in Figure 46. This menu lists all available properties of NetFlow data useable for filtering:

AS / IP / Port / Interface The AS filter accepts flows based on AS numbers in the NetFlow data. It is configured using two drop down fields and a text field:

- The first field is used to select for which AS number the condition will be enforced. “AS” (default) is used for either source or destination AS, whereas “Src AS” and “Dst AS” are fulfilled when the source AS or destination AS contain the AS number, respectively.
- The second field is used to set the comparison between the condition and expression: possible values are equal (“=”), larger (“>”), larger or equal (“>=”), smaller (“<”) or smaller or equal (“<=”).
- The text field is used to specify the value that shall be compared with.

This type of filter is used for any properties that appear twice in any NetFlow data set and that can be compared using the aforementioned way. Along with AS numbers these are IP addresses and port numbers (for both source and destination) and numbers of interfaces at the router (incoming and outgoing interface).

Bytes / Packets This filter condition accepts flows depending on the data transferred, measured in Bytes or Packets. The drop down field allows to specify the comparison with the input of the following text field. These comparisons are: equal (“=”), larger (“>”), larger or equal (“>=”), smaller (“<”) or smaller or equal (“<=”). It is not possible to use measure units like “k”, “M” or “10⁶”, the value has to be specified entirely.

CIDR This filter accepts flows if its IP addresses match the network range defined by the CIDR notation in the text field. As with the AS filter the first drop down field specifies whether the filter is executed on either source IP (“Src IP”) or destination IP (“Dst IP”) address fields or in any (“IP”, default).

ICMP Code / Type This filter will only accept flows that include ICMP traffic with the specified code or type.

IP version This filter accepts only flows of a particular IP version. “IPv4” will accept only traffic of IP version 4, “IPv6” similarly will accept only traffic of IP version 6.

Protocol The protocol filter accepts flows based on the protocols used within the flow. The protocols (e.g. TCP, UDP, ICMP) are selected using a drop down field. The number in brackets is the protocol’s number as specified by IANA²⁵.

Router This filter allows to filter flows based on routers connected with the object.

TCP Flags This filter accepts flows including TCP traffic, if the flags selected by ticking the appropriate checkboxes match those in the packets. Therefore, selecting “SYN” and “ACK” will match only flows which contain packets with SYN flag and packets with ACK flag.

A simple example for a visual NetFlow filter is shown in Figure 47. The filter comprises of two conditions linked by a logical operator.



Figure 47: A simple NetFlow filter in visual mode

Adding a second condition will also create a drop down field to select a logical operator to link the two conditions. The filter will accept flows only if any (“OR”) or both (“AND”) conditions are satisfied.

The second (right hand side) condition in Figure 47 shows two additional controls to specify a filter expression. If the mouse pointer is positioned on a filter on the left hand border two icons appear. The red cross in the upper left corner can be used to remove this condition from the filter expression, a click on the yellow exclamation mark inverts this condition, with the effect that only those flows will be accepted, that do **not** satisfy the condition. This is equivalent to a the logical “NOT” operator and visualized as such (see Figure 47). The entire filter expression can be inverted like this as well, the icons for the entire filter expression appear as soon as the mouse pointer is within the filter area but not within any condition area. The “NOT” operator

²⁵More on these here: <https://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>.

can be removed, by clicking on the yellow exclamation mark again. The filter in Figure 47 will therefore accept all flows that incorporate TCP traffic but do not use port 80.

If a filter expression requires a combination of “AND” and “OR” operators the controls do not suffice. For a filter expression only a single operator can be specified and all filter conditions within will be logically linked using this operator. To use both operators in one filter expression “Groups” must be used. A group is visualized using a new block in a new line within the filter expression area. A group is basically a filter expression itself and has the same links as discussed above to add conditions and groups but a different logical operator can be used for the conditions within the group. Figure 48 shows an example that cannot be expressed using the controls provided without using the group feature. The filter to be specified is: TCP connections using either port 80 or port 443.

Figure 48: A NetFlow filter using Groups of filters

Using the tools provided complex filter expressions can be specified, see Figure 49. This filter accepts flows with TCP traffic from the 10.0.0.0/8 subnet to destinations other than 10.1.2.3 using port 80 or port 443 or if the flow contains ICMP traffic of type 3—destination not reachable. The nfdump form would be:

```
(proto 6 and src net 10.0.0.0/8 and not dst ip = 10.1.2.3 and (port =
  80 or port = 443))
or (proto 1 and icmp-type 3)
```

Figure 49: A more complex NetFlow filter expression

NetFlow filters are used throughout the DDoS application to characterize a phenomenon indicated by indicators. To validate whether a filter expression matches the traffic, indicators can be determined for filtered traffic and added to the indicator plot. A link labeled “[Filter to Plot](#)” above the upper right hand corner of the filter expression area does exactly this. While the filtered indicators are determined, a notification is placed in the graph “[Loading plot data, click to cancel...](#)” which will cancel the calculation when clicked. As all raw data for the specified time period have to be processed, the processing might take a while, but can be

shortened by shortening the time period. If indicators are determined, their filtered variant is shown as an additional individual graph, see Figure 50.



Figure 50: Plot of a NetFlow filter

In case of Figure 50 the filter expression based on protocol, ip address and port characterizes the phenomenon at 10:55 to a large extent. The filter accepts 4 M UDP packets per second matching the increased traffic (from .5 M to 4.5 M UDP packets per second). Outside of the increased traffic the filter does not match any traffic, as the first increase of the packet rate at 10:20 is not part of the filter.

If the filter characterizes the event with sufficient precision, it can be translated to an [Access Control List \(ACL\)](#) for Cisco routers using the button labeled “[Filter to ACL](#)” placed on top of the right hand upper corner of the filter expression area. A text field will pop up that shows a translated notation of the filter for UDP packets from source IP address 192.0.2.10 to destination port 22 to the following ACL:

```
ip4 access-list router72
10 deny udp 192.0.2.10 0.0.0.0 0.0.0.0 255.255.255.255 eq 22
```

The button labeled “[History](#)” as well as the navigation buttons “<” and “>” can be used to see the history of specifying the filter expression. Every input is saved as a revision and can be retraced. The button [History](#) will show an overview as seen in Figure 51. Each entry can be used as a filter expression using the link “→”], whereas “<” jumps back one revision and “>” jumps forward one revision.

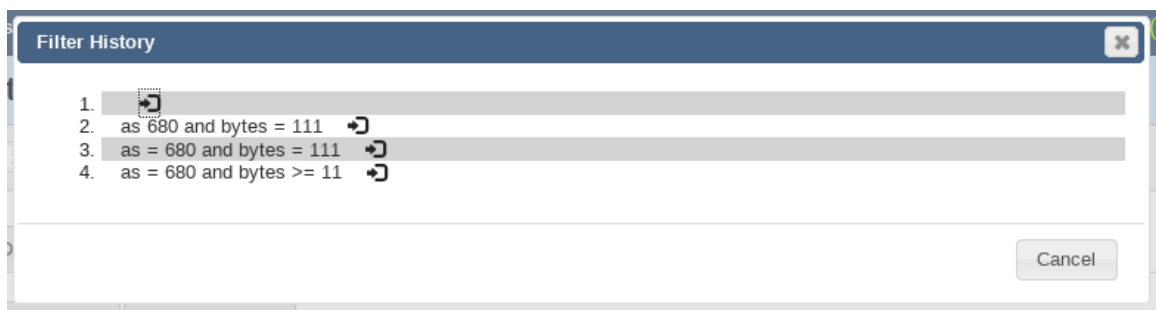


Figure 51: History of a filter expression

4.12.2.1.4 Evaluations The third section of the NetFlow view of an object (see Figure 52) allows various evaluations of filtered NetFlow data.




Figure 52: Evaluation of NetFlow data

The predefined evaluations are spread across five registers:

Top-N This tab shows the **N** most frequent values of various properties of flow data in the filtered flow data, like **N** most frequent source ports or destination IP addresses. **N** can be specified by the user, the default being **N** = 10, see Section 4.12.2.1.4.

Possible Targets In this tab the prevailing communication entities are listed. A communication entity is a combination of IP address and port, therefore, for each flow there is a source and a destination entity, see Section 4.12.2.1.4.

Parallel Coordinates This register allows a graphical analysis of flow data indicating various pattern and allowing a further filtering, see Section 4.12.2.1.4.

Raw Flows This is a textual representation of the flow data, see Section 4.12.2.1.4.

Aggregated Flows This register lists a textual view of flows aggregated by various properties, see Section 4.12.2.1.4.

All views of NetFlow data within the DDoS application show the values as collected at the routers. The sampling conducted by the routers is **not** included.

All searches within the NetFlow analysis require processing a huge amount of raw data. The processing is executed asynchronously in the background to enable the user using the object view without having to wait for the end of the processing. So further analysis can be started on another register while waiting for results of the first analysis. Results are presented in the register. Any processing of data for an analysis can be stopped using the button now labeled “**Cancel**”.

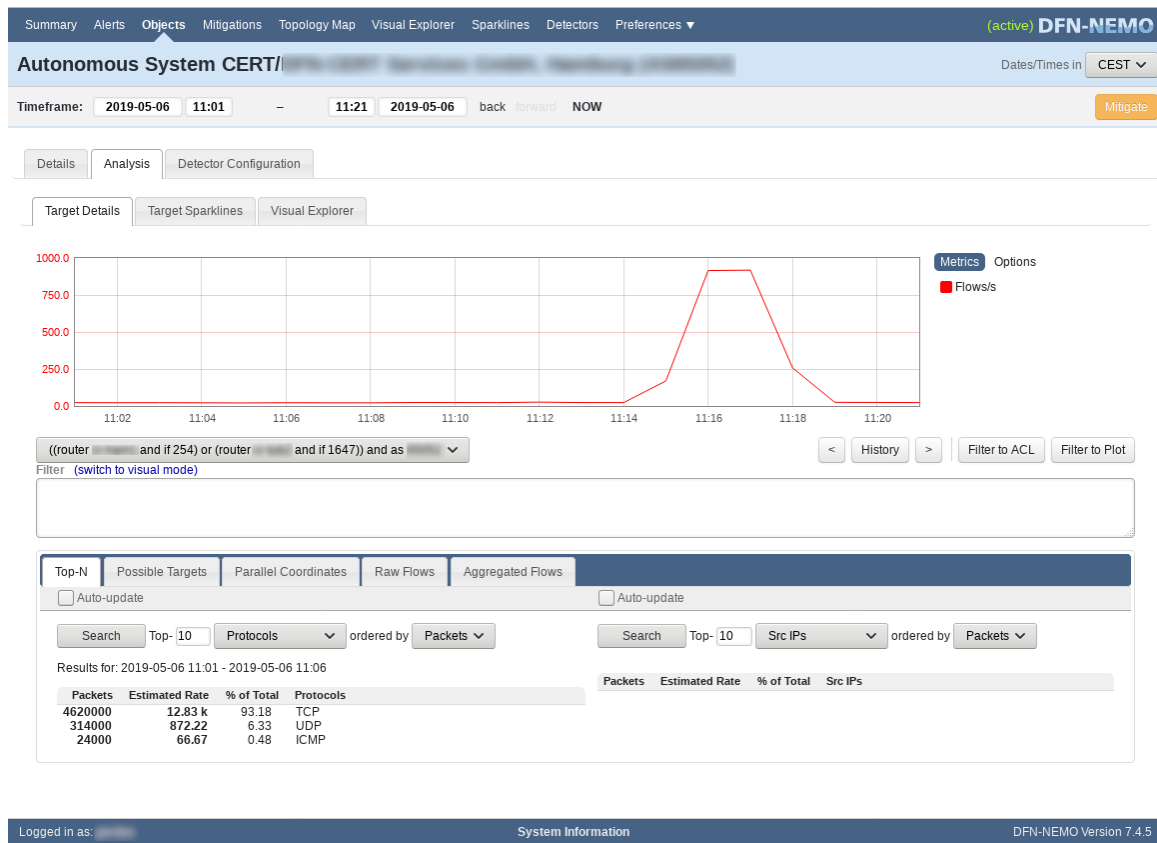


Figure 53: Evaluation of top 10 protocols in the flow data

Top-N analysis The [Top-N](#) register (see Figure 53) shows two [Top-N](#) evaluations next to one another, having similar features. Both evaluations show the [N](#) most frequent values of a NetFlow data property. All registers start with an “[Auto-update](#)” checkbox. If ticked, evaluations are updated if any filter is updated, (e.g. moving the start time of the evaluation will trigger an update of the evaluation results). If the checkbox is not ticked, the results must be updated manually.

Parameters for the [Top-N](#) lists are specified on top of the results list:

Top- This text field specifies the maximum number of displayed results. The default value is 10.

DropDown field This drop down field enables the user to select the NetFlow property to be evaluated. The list of properties is: Src Ports, Dst Ports, Src IPs, Dst IPs, Src AS, Dst AS, Bytes, Packets, Protocols, TCP flags, Input Interface, Output Interface

ordered by This drop down field specifies the count strategy. Depending on the selection the most frequent values are determined by packets, flows or bytes.

Search This buttons enables a manual evaluation of NetFlow data and must be clicked unless [Auto Update](#) is enabled as discussed above.

The Top-N results are titled with a line indicating the time period being analyzed. If the mouse pointer is placed on this line, the time period is also highlighted in the

graph, see Figure 54. In any case, only the first five minutes of the time period are analyzed for the Top-N evaluation. If raw data is insufficient, a notification will pop up.

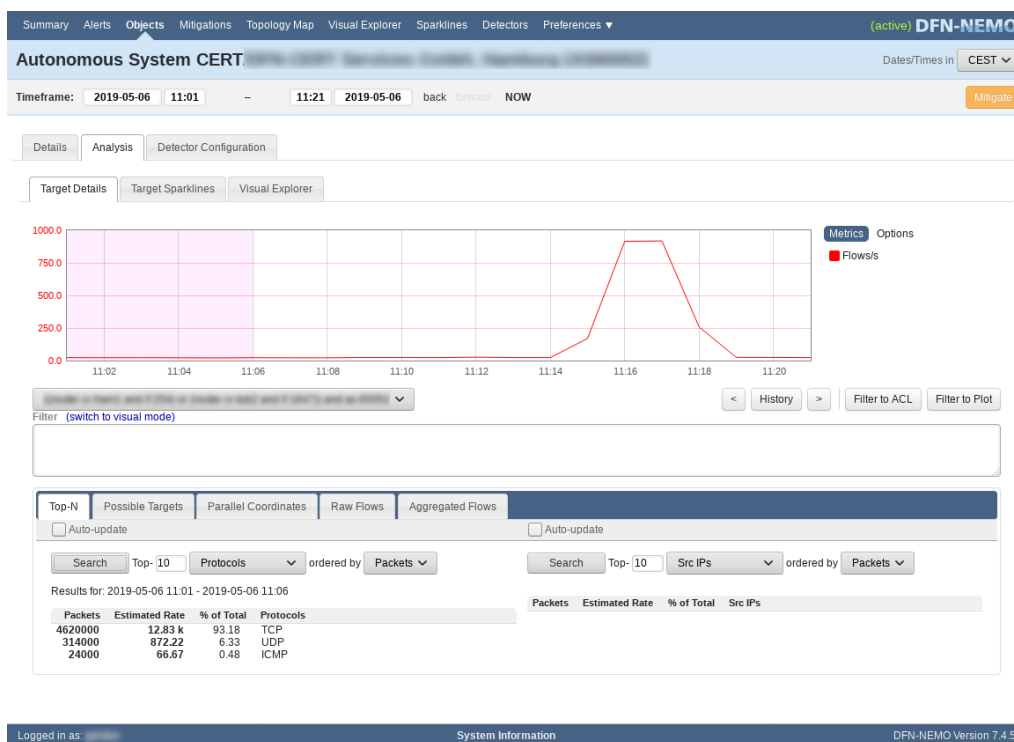


Figure 54: Time period of analysis is highlighted in the graph

The example in Figure 53 shows three protocols being used in the analyzed data, whereas there are 4,620,000 TCP packets being 93.18% of the total packets in the time frame. The numbers are the very same numbers of packets in the sampled data, so to get the total packet count the packet count must be multiplied with the inverted sampling rate to result in a number close to the actual packet rate.

Possible Targets Figure 55 shows the most communicating hosts within the sampled NetFlow data. The results can also be updated automatically if the [Auto-update](#) checkbox is ticked. The analysis can be configured by choosing the number of results and by selecting the way to count, the analysis will be conducted if [Auto-update](#) is ticked or when the button labeled [Search](#) is clicked.

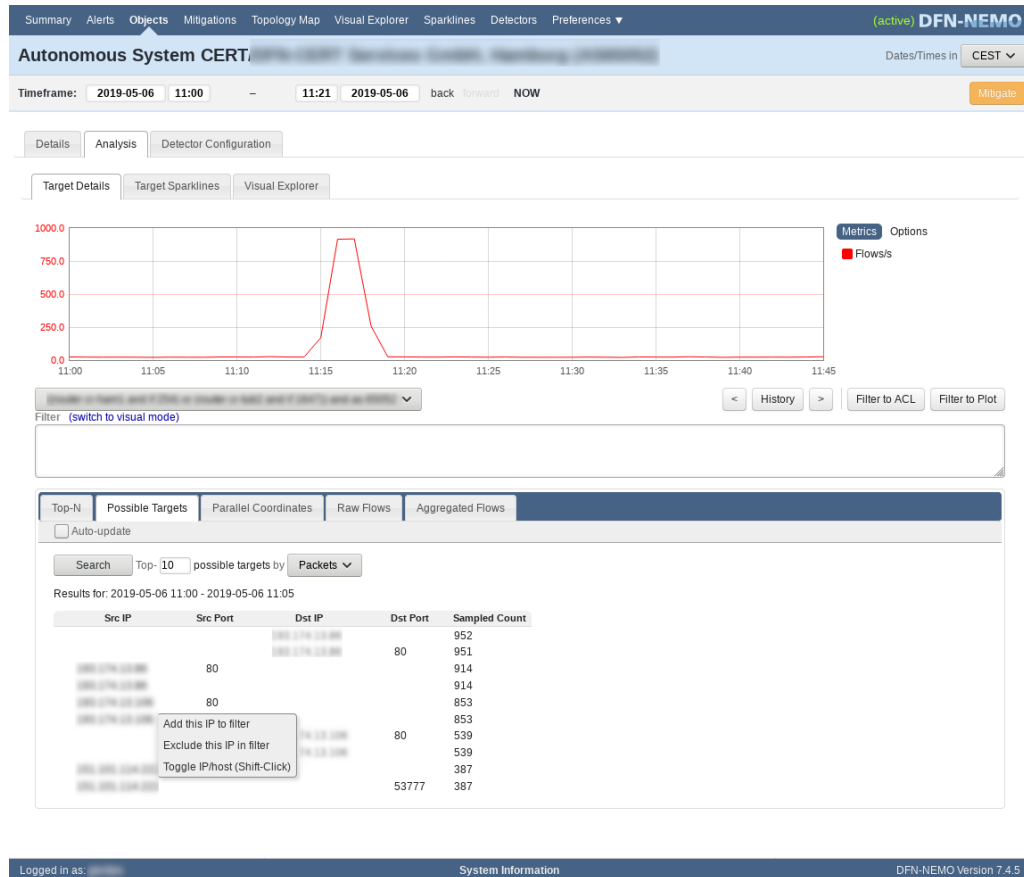


Figure 55: Analyzing the top ten hosts (most communication)

Results are titled with a line indicating the time period being searched. The results table consists of combinations of IP address and port for either source or destination as well as the number of their occurrence. Empty cells connotate multiple values. So the first result means that the host with the most connections is xxx.yyy.zzz.86 as a destination with 952 sampled packets, of which are 951 targeted to Port 80 (second most frequent result).

On the right hand side of the source IP address in line 6 is a pop up menu that will be shown once an icon is clicked which appears when the mouse pointer is on the IP address. This menu offers these features:

Add this IP to filter This will add the IP address to the filter in the second section using a logical “and” operator, resulting in reducing the results of the filter to those flows that contain this IP address as either source or destination.

Exclude this IP in filter This will add a condition to the filter in the second section with the intention of “and not this IP address”. Any analysis is only conducted for flows that fulfil the specified filter and not this IP address.

Toggle IP/host (Shift-Click) This will trigger a “reverse DNS lookup”²⁶ to associate a hostname to this IP address. If the lookup is successful,

²⁶A reverse DNS lookup will use the DNS system to associate a hostname to an IP address, see also https://en.wikipedia.org/wiki/Reverse_DNS_lookup.

the IP address will be replaced by the hostname, if not, the IP address will flash in red for a short time. This may also be triggered by clicking the IP address while pressing the Shift key.

Parallel Coordinates The next evaluation is using [Parallel Coordinates](#). In this view, flows are visualized enabling the user to detect characteristic traffic pattern. There are two visualizations, one for each of the IP protocol versions (see Figure 56 for an IPv4 example).

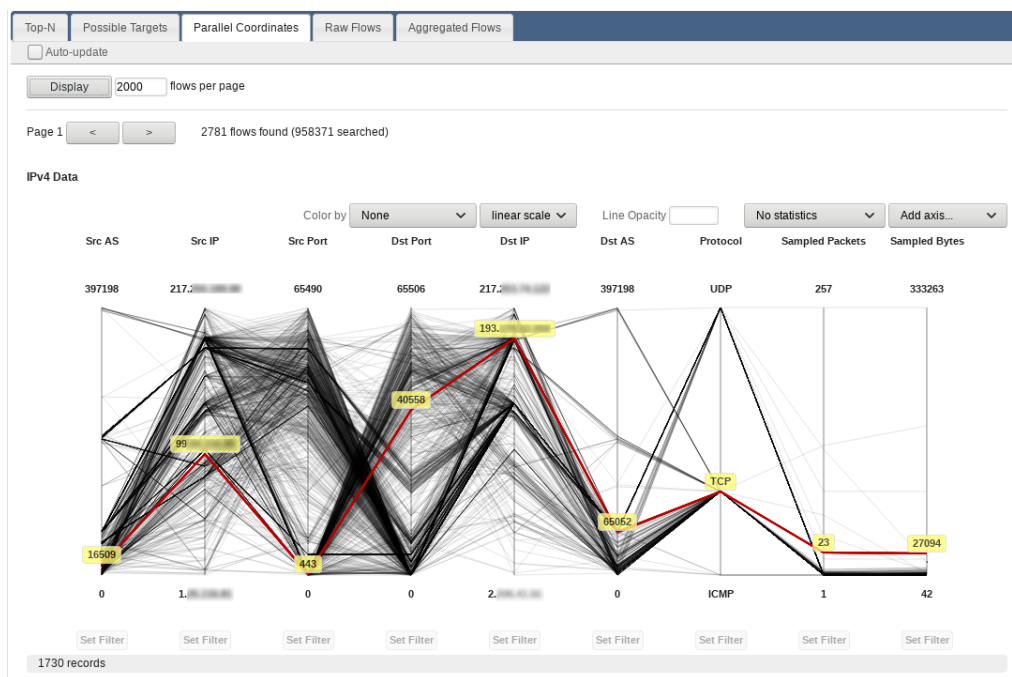


Figure 56: Evaluation of NetFlow data using [Parallel Coordinates](#)

The register starts with a statement on the number of flows used for each page of the results. If the number is too small, traffic patterns might not be detectable, if the number is too high, patterns will be invisible because of the multiplicity of flows. The preset value of 2,000 seems to be a good default.

A click on “[Display](#)” starts the search for matching NetFlow data. The next line “[x flows found \(y searched\)](#)” states the progress of the analysis. It contains the number [x](#) of flows matching the filter criteria while [y](#) flows have been evaluated. The analysis is conducted incrementally until the number of required flows for the current page is reached. If the statement says “[2000 flows found](#)” this does not mean that only 2,000 flows match the filters, but that only 2,000 have been identified in the flow set so far. If a different page is loaded the search is continued until the required number of flows is reached or no more matching flows are found. The current page number is indicated at the top. If the required number of flows (e.g. 2,000 as the default) is not found within a given time period, a notification appears.

If the search is cancelled or finishes successfully for the current results page the matching flows are visualized in plots for both IPv4 and IPv6. The plots comprise several parallel vertical axes which show various properties of the flows. Figure 56 shows axes for AS number, IP address and port for source and destination, protocol and number of flows and bytes, other properties are average packet size, duration

and start time of a flow and the input and output interface. Further axes can be added using the “Add axis...” drop down field on top of the plot. If the mouse pointer is located on the title of an axis, a red cross appears, which can be used to remove axis from the plot. The order of the axes can be changed by dragging and dropping the axis’ title. Directly above each axis the minimum and maximum values are indicated, the values inbetween are distributed equally along the axis. Clicking on the maximum or minimum value will switch the direction of the axis.

A single flow is visualized as a black line connecting the flow’s actual values of the properties. If the mouse pointer is on one flow, this line will be highlighted in red and its values shown at each axis (see Figure 56).

The visualization of the plot can be adjusted by two parameters:

Color by The color of the lines can be changed, selecting any of the axes. The number of flows is used as a scale²⁷ and differentiated between linear and logarithmic scale for coloring. Compare the images in Figures 57, 58 and 59.

Line Opacity The intensity or opacity of the lines can be specified with values between 0 and 1. The decimal point “.” is the delimiter. The default value is 0.1 without being shown in that field (see also Figure 59).

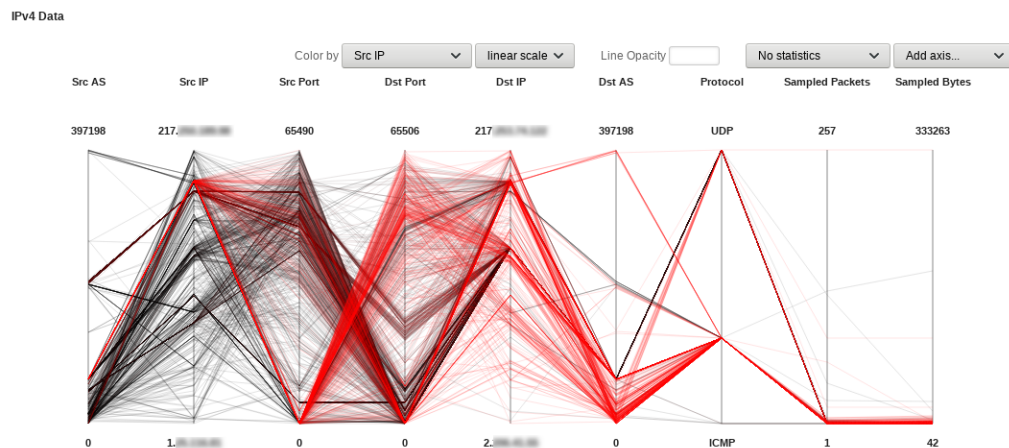


Figure 57: Coloring on linear scale

²⁷The highest number of flows is colored using the color with the most intense red value. For smaller flow values the color is changed using a color with less red.

- For linear scales the color progress is equally for color and value.
- For a logarithmic scale the same color is used for smaller values than at the linear scale.

On a logarithmic scale smaller values are colored more likely in red than on a linear scale with at least one large value.

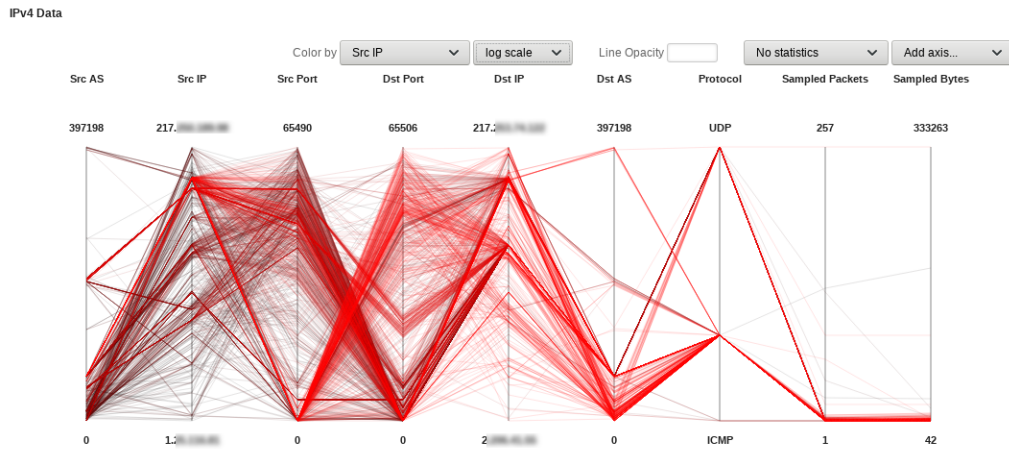


Figure 58: Coloring on logarithmic scale

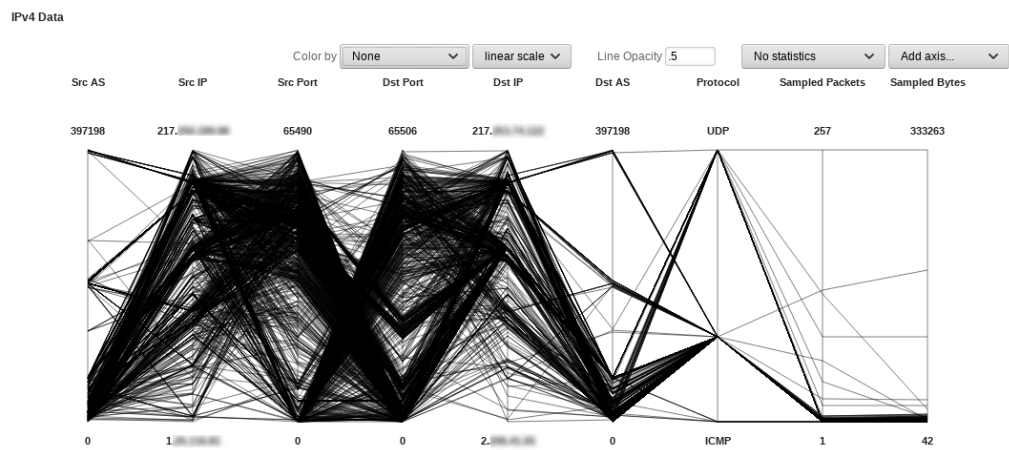


Figure 59: Intensity of the lines with value .5

Traffic patterns can be detected by simultaneous visualization. The visualization on the protocol axis shows that TCP is widely used. Many lines cross the axis at this point, therefore the corresponding flows contain TCP connections. The connection between source and destination port shall also be explained: Source and destination ports connect in a crossed manner, high source ports connect to low destination ports and low source ports connect with high destination ports. This is the typical scheme of connections in the internet in which clients using high source ports connect to services on low ports (smaller than 1024). As flows are unidirectional, the services' reply is considered an entire new flow.

Figure 60 shows a typical example for eye-catching traffic patterns in parallel coordinates. The plot contains axes for IP address and port for both source and destination. Eye-catching is the strong focus of the lines at one point of the scale for the destination IP address²⁸. Then the lines spread the entire scale for the destination port before focussing nearly on a single point on the source port scale.

²⁸This analysis is from right hand side to left hand side




Figure 60: Eye-catching traffic pattern in parallel coordinates




Figure 61: Marking an interesting traffic patterns using value ranges

To analyze this pattern clicking and dragging can be used to mark interesting sections on one or multiple axes. Figure 61 shows this for an interesting section of the source IP address.

The marked range is visualized by a red thick swelling on the axis. The flows with the property value in this range are marked blue while all other lines are colored in light grey. Underneath the plot the number of marked flows is stated: “**2000 records, 1425 brushed**”. With 1,425 lines nearly three out of four visualized flows match the marked pattern, making it characteristic. Underneath the axes on which a marker was placed the “**Set Filter**” button is enabled. Clicking the button will add the marked value range as an additional filter expression to the filter defined in the second section. If this process—mark and extend the filter—is repeatedly executed, a plot similar to Figure 62 is created. This shows one source using one port to send packets to all ports of the destination.




Figure 62: An isolated traffic pattern in parallel coordinates

Another analysis using parallel coordinates is based on showing “[Quartiles](#)”. Using a drop down menu on the right hand side quartiles can be calculated and visualized for all axes. The calculation of quartiles can be conducted on flows, packets or bytes. The quartils are visualized using transparent colored bars. The red line indicates the median, while the top and lower bar ends mark the top and lower quartile, respectively. Each bar visualizes 50% of all values, Figure 63 shows quartils of an analysis.

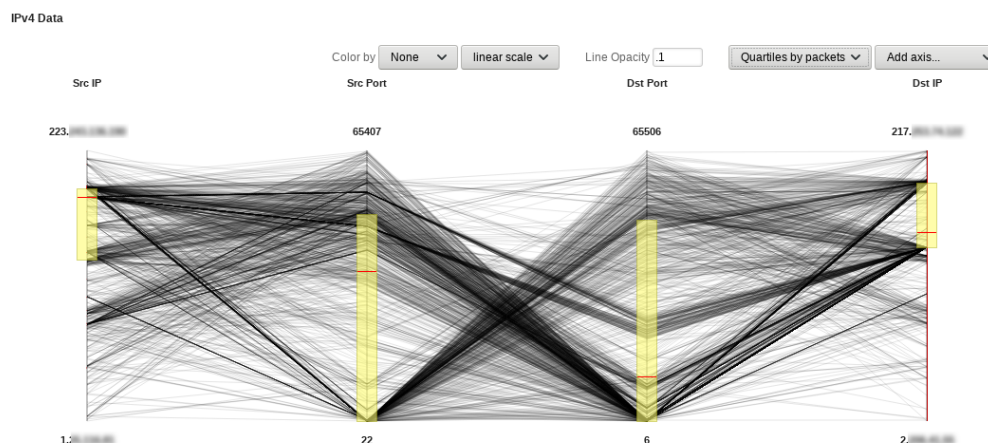


Figure 63: Visualization of quartiles based on packets in parallel coordinates

Raw Flows The fourth register is visualized in Figure 64. It shows the raw NetFlow data of each individual flow. It starts with a text field collecting the number of flows to visualize per page. The visualized data is paginated using mutiple pages if the number of flows exceeds the number of flows per page. The analysis is either started using “[Display](#)” or by ticking “[Auto-update](#)”.

No.	Router	Time	Src IP	Src Port	Src AS	Dst IP	Dst Port	Dst AS	Protocol	Flags
0		2019-05-06 10:59 - 10:59	193.50.135.100	35492	65052	193.50.135.100	443	14061	TCP	A
1		2019-05-06 10:59 - 10:59	193.50.135.100	9001	12876	193.50.135.100	50240	65052	TCP	A
2		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	51769	198570	TCP	AS
3		2019-05-06 10:59 - 10:59	193.50.135.100	10293	65052	193.50.135.100	443	2906	TCP	A
4		2019-05-06 10:59 - 10:59	193.50.135.100	443	2906	193.50.135.100	10293	65052	TCP	A
5		2019-05-06 10:59 - 10:59	193.50.135.100	63253	3320	193.50.135.100	80	65052	TCP	A
6		2019-05-06 10:59 - 10:59	193.50.135.100	55500	206485	193.50.135.100	23013	65052	TCP	S
7		2019-05-06 10:59 - 10:59	193.50.135.100	47713	137	193.50.135.100	3000	65052	TCP	AF
8		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	51772	198570	TCP	A
9		2019-05-06 10:59 - 10:59	193.50.135.100	56191	3320	193.50.135.100	80	65052	TCP	A
10		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	34556	65018	TCP	A
11		2019-05-06 10:59 - 10:59	193.50.135.100	28224	49234	193.50.135.100	443	65052	TCP	A
12		2019-05-06 10:59 - 10:59	193.50.135.100	51880	3320	193.50.135.100	80	65052	TCP	AP
13		2019-05-06 10:59 - 10:59	193.50.135.100	53843	65052	193.50.135.100	53	394353	TCP	A
14		2019-05-06 10:59 - 10:59	193.50.135.100	53	397194	193.50.135.100	33000	65052	TCP	AF
15		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	52380	65415	TCP	AS
16		2019-05-06 10:59 - 10:59	193.50.135.100	49480	3320	193.50.135.100	80	65052	TCP	A
17		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	51416	8255	TCP	AP
18		2019-05-06 10:59 - 10:59	193.50.135.100	80	65052	193.50.135.100	12711	65051	TCP	AP
19		2019-05-06 10:59 - 10:59	193.50.135.100	443	13335	193.50.135.100	46044	65052	TCP	A

Figure 64: Visualization of raw NetFlow data

Each line represents a flow and contains the following dataset: consecutive flow number, router data, start and end time of the flow, IP addresses, ports and AS association for both source and destination, protocol, TCP flags and ICMP type and code, number of transmitted packets and bytes, protocol version of the Internet Protocol, input and output interface on the router.

If the mouse pointer is on an interface number, a pop up will show the name of the connected line and the name of the interface. The visualization of the interface can be toggled using the link “[Toggle Iface Display](#)”. Values are: Number of the interface, name of the interface at the router and name of the connected line at the router. On the ICMP type and code field, a textual representation will be shown.

The default sort order of the flows is the sort order of the flow number, which is equal to the sort order of the raw data. The sorting can be changed using the triangles in each column’s header.

Aggregated Flows The last register visualizes aggregated flows, flows that share equal values for one or more properties. Hence, an overview of flows is created that have e.g. the same source IP address or the same destination port. Figure 65 exemplifies this for the aggregation on source IP addresses. The selection of parameters used for aggregation is placed at the top by checkboxes, one for each parameter/property. All flows will be aggregated that have equal values in all selected parameters. The visualization is paginated using the sort order as specified in the drop down fields below the selection checkboxes. Sorting can be conducted using flows, bytes or packets and in ascending or descending mode.

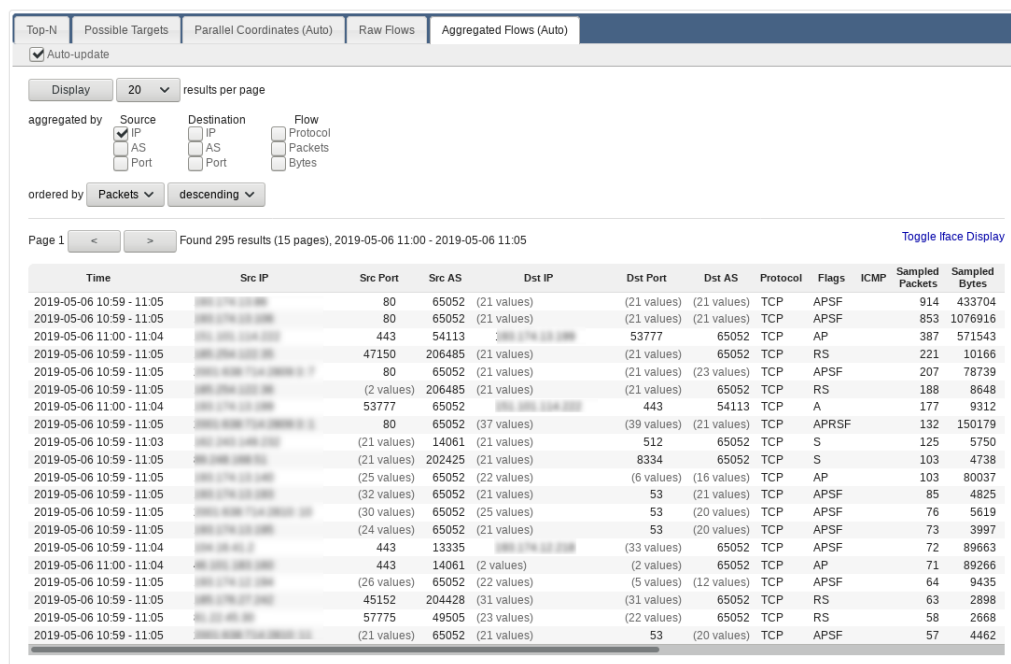


Figure 65: Aggregation of NetFlow raw data using the source IP address

In the example in Figure 65 all flows that originate at the same IP address are summed up into one entry. Therefore in the various columns the number of different values are listed as “(x values)”. Clicking this entry will expand the list of all individual entries. Clicking “(less)” will reduce the list.

Columns are aggregated as follows:

Time This will show the entire time range in which flows of this aggregated flow were detected. The start time equals the start time of the first flow, the end time equals the end of the last flow.

Flags This will indicate all TCP flags set in any of the aggregated flows' packets.

ICMP Each line may contain TCP and ICMP flows, therefore, ICMP type and code are separated from TCP flags in the aggregated view.

Packets and Bytes This column will hold the sum of packets and bytes transmitted in the aggregated flows in total.

4.12.2.2 Target Sparklines This section shows sparklines for all indicators for this particular object only, see Figure 66. See Section 4.8 for details on sparklines.

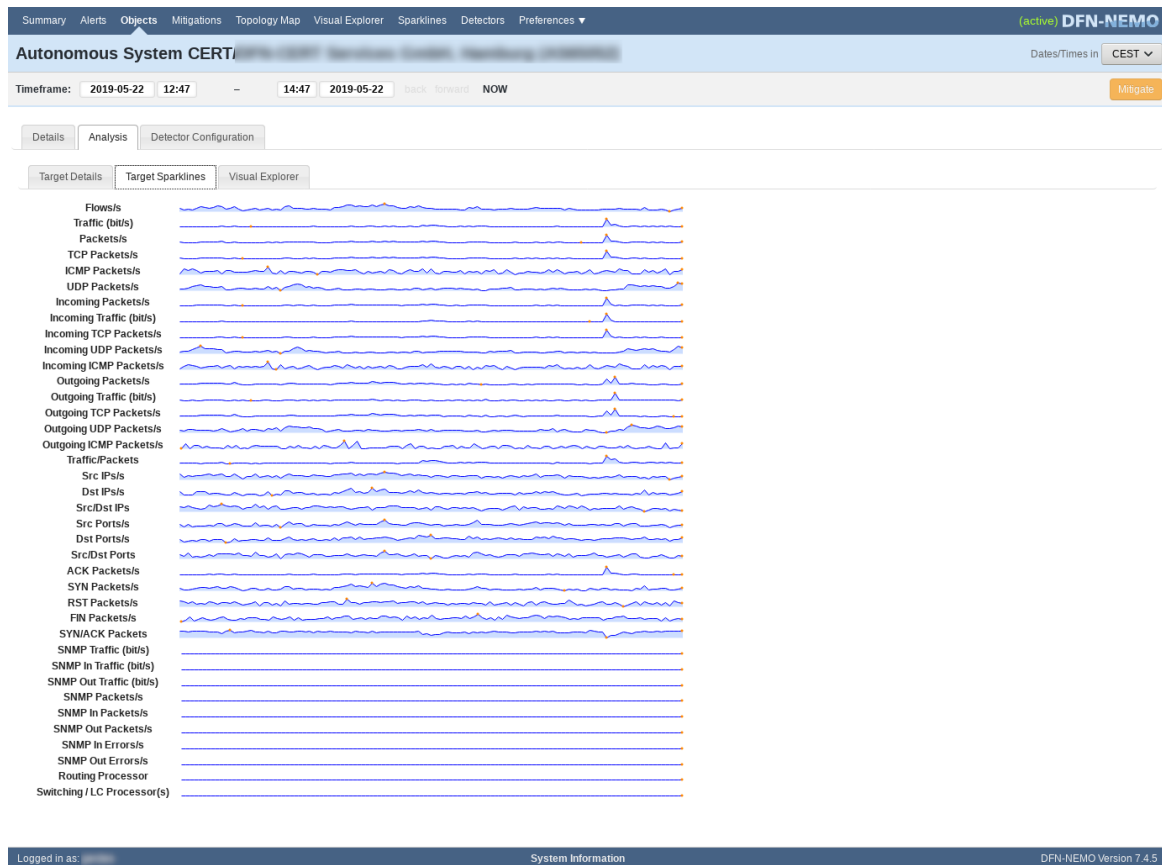


Figure 66: A selection of sparklines for the selected object

4.12.2.3 Visual Explorer The [Visual Explorer](#) provided on this register has the same options and functions as explained in Section 4.7, but traffic data is considered for this evaluation only if flows originate from or destined to this particular object. See Figure 67.

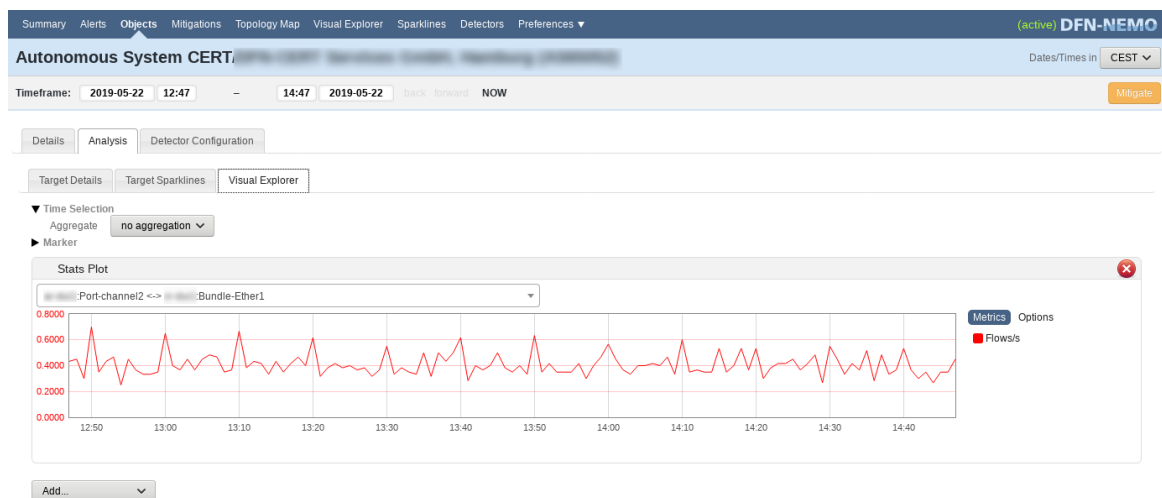


Figure 67: A Visual Explorer for an object (part view)

4.12.3 Configuration of detectors

The configuration view in Figure 68 shows the configuration of detectors for a single object. All detectors configured to monitor the object are sorted alphabetically with further configuration information.

The available configuration depends on the type of the detector. Below the name further information for each detector is listed in four columns:

Parameter name of the parameter

Default Value default configuration value

Object Value currently configured value, if this value is empty, the default value is active.

Help Text a clarifying text for this parameter

The column “**Object Value**” provides editable fields in which any value can be specified for each parameter (e.g. new numbered values for thresholds). A changed configuration can be stored using the “**Save**” button. A configuration is checked for plausibility before being deployed. Next to the “**Save**” button is a “**Remove**” button that removes and object value and restores the default values.

If a detector cannot be configured by the current user, the text “**No configuration to display**” is shown instead.

Summary Alerts **Objects** Mitigations Topology Map Visual Explorer Sparklines Detectors Preferences ▼ (active) DFN-NEMO

Router ar [object]

Dates/Times in CEST

Timeframe: 2019-05-06 14:31 – 16:31 2019-05-06 back forward NOW Mitigate

Details Analysis **Detector Configuration**

Heartbeat (Heartbeat)
No configuration to display

Heartbeat SNMP (Heartbeat SNMP)
No configuration to display

Pps Incoming Router (Threshold)

Parameter Name	Default Value	Object Value	Help Text
Field	incoming packets (SNMP) per minute	not editable	
Denominator	-	not editable	
Critical Below	0	None	
Warn Below	0	None	
Info Below	0	None	
Info Above	0	None	
Warn Above	0	None	
Critical Above	600,000,000	None	The number of packets per second
Ignore Field	-	None	
Ignore Below	-	None	

UDP (Holt-Winters) [-> Model plot]
No configuration to display

XR Routing Processor (Threshold)
No configuration to display

XR Switching Processor (Threshold)
No configuration to display

Logged in as: [user] System Information DFN-NEMO Version 7.4.5

Figure 68: Configuration view of a detector

4.13 Details of an Alarm

Any alarm can be analyzed using the web interface. This will yield information on the type and scope of the traffic responsible for the alarm, the communicating parties and the path through the Network used by the packets associated with the alarm.

4.13.1 Alarm summary

The alarm summary is the point of entry to the analysis of an alarm message. The summary as printed in Figure 69 is linked to at various places throughout the web application and also in each email sent by the application (see Section 3).

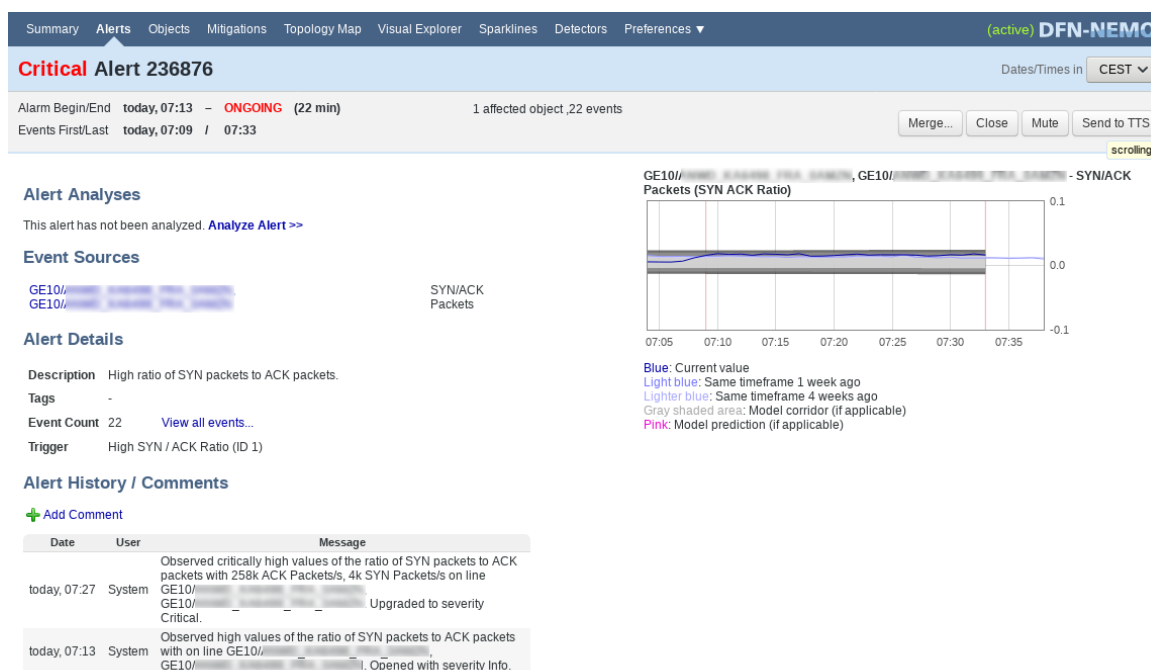


Figure 69: Summary of an alarm (part view)

4.13.1.1 Navigation area The summary page starts with an extended navigation area, that will be visible throughout the analysis of this alarm. Next to a title containing ID and severity of this alarm important data is stated, that has also been sent out in the email notification. This data includes

start and end time of the alarm and its duration The presentation format of the start and end time is “**Start date, start time – end date, end time (duration)**”.

the date and time of the first and last events associated with this alarm This data is formatted “**Date earliest event, time earliest event / Date most recent event, time most recent event**”.

If start date and end date are the same day, the end date is omitted. If any date is today or the day before instead of the date the text “**today**” or “**yesterday**” is printed. If the alarm is still ongoing, the end time is replaced by the red colored text “**ONGOING**”. The duration in this case is the time, that has passed since opening the alarm.

In the center, the number of affected objects in the Network as well as the number of events associated with this alarm is given.

On the right hand side, action buttons are placed. The number is three or four, depending on the state of the alarm:

Merge... Using this button alarms can manually be merged to Meta alerts. See Section 4.13.1.2 for an extensive description of the process.

Close This button is only visible when the alarm is still opened (ongoing). After clicking the button a confirmation is required before the current alarm is closed. The traffic responsible for opening the alert will not be ignored afterwards. If the traffic trend is ongoing or reappearing a new alarm will be created.

Mute A click on this button will oppress notifications to be send out by this alarm. No further notification emails are send and the alarm will not be listed in the alarm overview (see Section 4.3) by default anymore. If the current alarm is already muted, the button will be labeled “**Unmute**” implementing the reverse feature.

Send to TTS This button can be used to notify the NOC’s Trouble-Ticket-System of this alarm by email. The dialog is shown in Figure 70 and requires to specify the ID of a ticket in the TTS. A comment may be specified additionally.

A screenshot of a web dialog box titled "Sent Alert 236876 to TTS". The dialog has a dark blue header bar with a close button (X) on the right. Below the header, there is a label "TTS ID" followed by a text input field containing the placeholder text "Please enter a TTS ID". Below that is a label "Comment" followed by a larger text area. At the bottom right of the dialog, there are two buttons: "Send to TTS" and "Cancel".

Figure 70: Dialog for “Send to TTS”

Scrolling This button is placed slightly under the other buttons and can be used to control the scrolling behavior of the navigation area. A click on the label toggles the setting between different behaviors:

scrolling This is the default. The navigation area scrolls as well as the alarm summary northwards when the page scrolls southwards.

fixed An additional line will be added below the extended navigation area. When scrolling long pages, the navigation area sticks to its position at the top and the information and functions within are accessible to the user at all times.

4.13.1.2 Merging of Alarms Figure 71 shows the pop up listing an overview of alerts already registered with the system. This overview lists for each alarm its ID, the severity of the alarm, duration and start time of the alarm as well as its description. There are also navigation items as each page is limited to 15 results. The list may also be filtered using the search field in

the upper right corner (the search field acts as the one described in Section 4.3). This window can be closed using the button labeled “Close” or the “X” in the upper right corner.

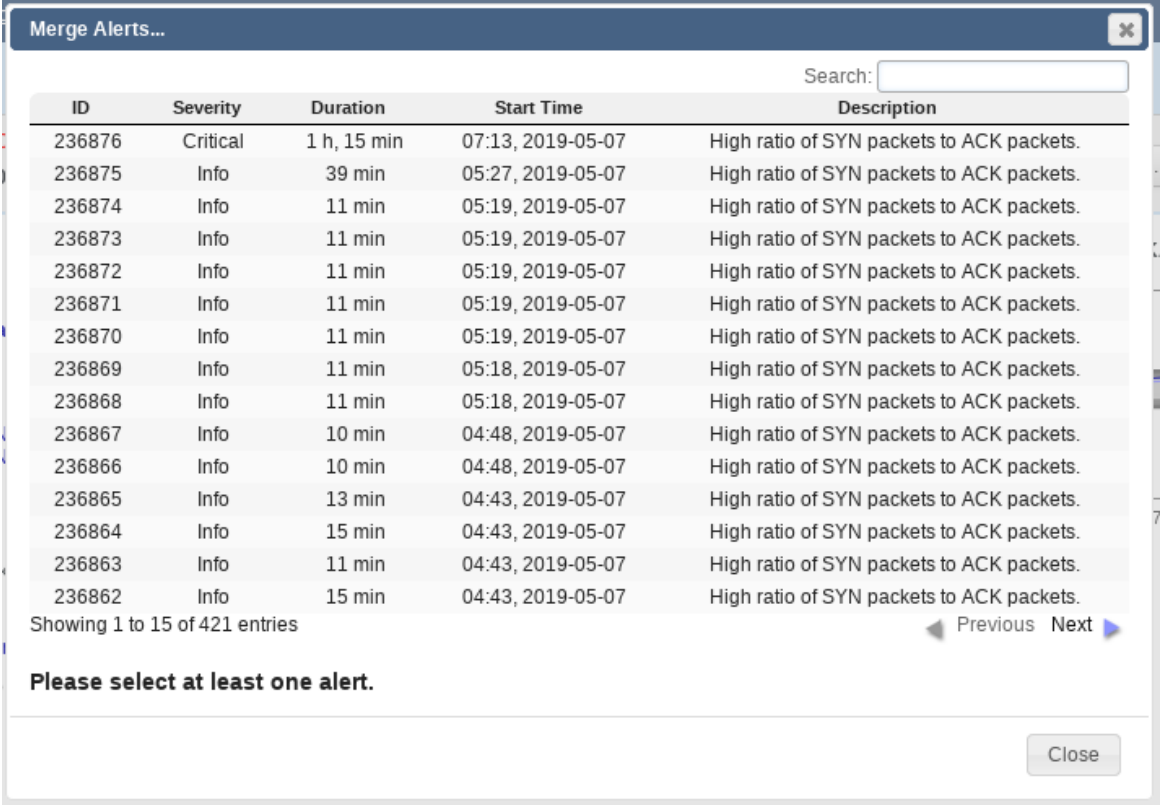


Figure 71: List of alarm for manual merge

Alarms can be selected to be merged by clicking their list entry. The list entry will be highlighted using a dark grey background. A selected alarm can also be unselected by clicking the list entry one more time.

- If the alarms are simple alarms, a dialog (see Figure 72) asks to select a type for the

new Meta Alert:

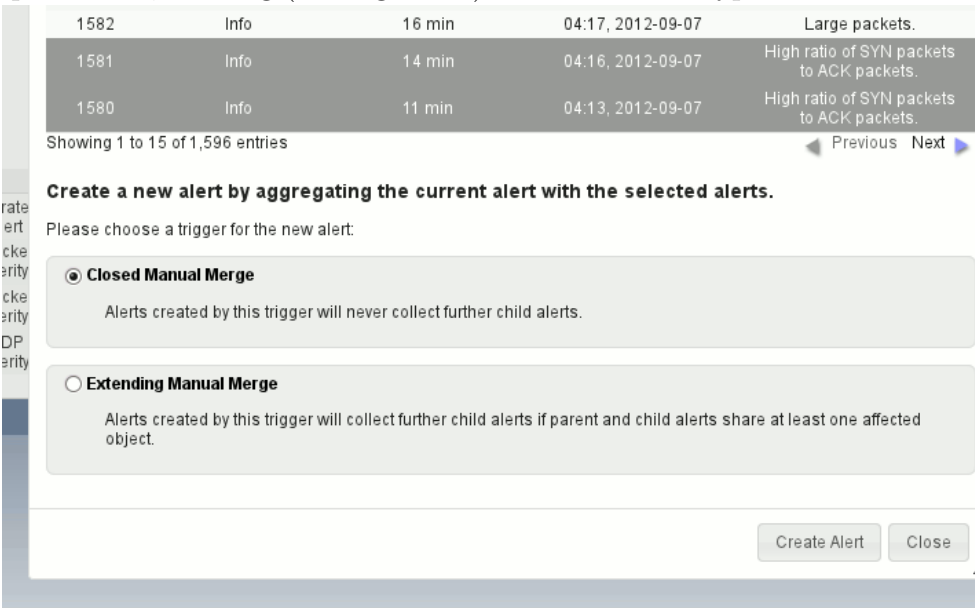


Figure 72: Creating a new Meta Alert

Closed Manual Merge This type creates a new Meta Alert that can only

be extended manually by additional alerts.

Extending Manual Merge This type creates a Meta Alert that is automatically extended by all alarms that are registered on any object already associated with this Meta Alert. An exception to this are heartbeat alerts²⁹ which are processed separately and which are therefore never associated to any manually created Meta Alert.

If the selection is final the new Meta Alert can be created using the button labeled “[Create Alert](#)” in the bottom right corner.

- If exactly one Meta Alert is selected from the list, the actual alarm can be added to this Meta Alert using a dialog as depicted in Figure 73. If the actual alarm is a Meta Alert itself, the note “[Add the current alert to the selected alert](#)” is replaced by the note “[Add the child alerts of the current alert to the selected alert](#)”, and by confirming clicking the button labeled “[Add Alert\(s\)](#)” all child alarms of the current Meta Alert are added to the selected Meta Alert.

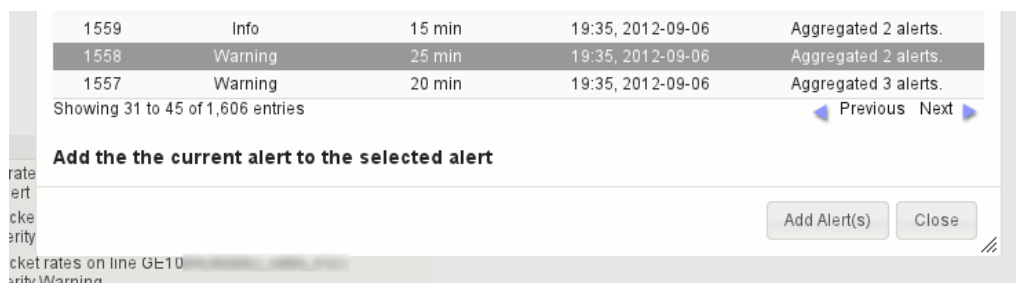


Figure 73: Adding an alarm to a Meta Alert

4.13.1.3 Summary area The summary area of alarm details is structured in two columns. The left hand column contains the information already send out via email about type and scope of the alarm, and provides links leading to detailed information. In the right hand column one or more plots are shown presenting an overview over the chronological trend of the indicators relevant for the alarm message.

4.13.1.3.1 Left column The left column consists of four to five sections which in turn will all be described:

Alert Analysis The first item in this list is “[Alert Analyses](#)” which lists all already available analyses for this alarm. An analysis comprises a describing text and all settings as provided in the analysis view as described in Section 4.13.3. If no analyses are present the text “[This alert has not been analyzed.](#)” next to the link “[Analyze Alert »](#)”, which directs the user to the analysis view for this alarm. If analyses are available, they are listed as shown in Figure 74a using creation time, creator, a description and a link to its analysis view. The link “[Start new Analysis](#)” may be used to start a new analysis of this alarm.

Event Sources The next section labeled “[Event Sources](#)” lists the objects for which the events were registered and associated with this alarm. The listing includes the name

²⁹A heartbeat alert is triggered when no NetFlow data is received for a router within a predefined time period. See also Section 4.13.2

of the object which is linked to its detailed view and the type or types of events registered on this object (e.g. “emphSYN/ACK” for a unusual relation between TCP packets with SYN flag and TCP packets with ACK flag), see also Figure 74b.

Description The next section provides the description of the alarm situation and the list of tags associated to the alarm. Then the total number of associated events is stated next to a link labeled “[View all events...](#)” that directs the user to an overview of individual events (see Section 4.13.2) and the name of the triggering detector.

Alert History The last section labeled “[Alert History / Comments](#)” shows important times and events during the development of the alarm and its editing by the users of the application, see Figure 74d. For each entry the time, the editor and a description are listed. The table lists the time of creating the alarm, the time of the adjustment of the severity and the time of closing the alarm. These entries list “System” as the editor, the description equals the one sent out via email as reason for the change of this alarm’s state. Any analysis is also noted with time, editor and the comment “[Analysis saved](#)”. Alarms and analyses can be commented using the link on top labeled “[Add Comment](#)”. This opens a dialog requesting a comment (see Figure 74e). If the comment is added, it appears in the history with time and editor³⁰ (see Figure 74f).

Child Alerts/Parent Alerts For Meta Alerts another section is shown, that lists the child alerts for this Meta Alert. Each entry consists of the following data (see Figure 74g):

- A red cross that removes—after confirmation—this child alert’s association with the Meta Alert
- The Alarm ID of the child alert which is also a link to the child alert’s summary
- The workflow status of this child alarm as described in Section 3.1.1
- The severity of this child alarm
- The duration and start time of this alarm
- The number of events associated with this child alarm
- The tags assigned to this child alarm.

Child alarms show a listing of Meta Alerts, this alarm is associated with (see Figure 74h).

³⁰The name of the editor are those provided in the user’s Client Certificate for authentication (see Section 2.2).

Alert Analyses

Start new Analysis

	Created	User	Comment
View Analysis	today, 12:43	[redacted]	Test

(a) Listing of available analyses

Alert Details

Description	High ratio of SYN packets to ACK packets.		
Tags	-		
Event Count	47	View all events...	
Trigger	High SYN / ACK Ratio (ID 1)		

(c) Providing details of the alarm

Add Comment for Critical Alert 236876

Please enter your comment:

(e) Dialog to add a comment

Child Alerts

	Alert ID	Workflow Status	Severity	Duration	Start Time	Event Count	Tags
X	236876	Analyzed	Critical	1 h, 15 min	07:13, 2019-05-07	47	
	(muted) Test (High ratio of SYN packets to ACK packets.)						
X	236889	Seen	Info	8 min (ongoing)	14:22, 2019-05-07	5	
	(muted) High ratio of SYN packets to ACK packets.						

(g) List of associated child alerts

Event Sources

GE10	[redacted]	SYN/ACK Packets
GE10	[redacted]	

(b) Listing of objects contributing events to this alarm

Alert History / Comments

[+ Add Comment](#)

Date	User	Message
today, 12:43	[redacted]	Analysis saved.
today, 08:28	System	The ratio of SYN packets to ACK packets on line GE10/[redacted] returned to normal values. Alert closed.
today, 07:27	System	Observed critically high values of the ratio of SYN packets to ACK packets with 258k ACK Packets/s, 4k SYN Packets/s on line GE10/[redacted]. Upgraded to severity Critical.
today, 07:13	System	Observed high values of the ratio of SYN packets to ACK packets with on line GE10/[redacted]. Opened with severity Info.

(d) View of the alert development

Alert History / Comments

[+ Add Comment](#)

Date	User	Message
today, 13:15	[redacted]	Test comment 123 This is really important!
today, 12:43	[redacted]	Analysis saved.
today, 08:28	System	The ratio of SYN packets to ACK packets on line GE10/[redacted] returned to normal values. Alert closed.
today, 07:27	System	Observed critically high values of the ratio of SYN packets to ACK packets with 258k ACK Packets/s, 4k SYN Packets/s on line GE10/[redacted]. Upgraded to severity Critical.
today, 07:13	System	Observed high values of the ratio of SYN packets to ACK packets with on line GE10/[redacted]. Opened with severity Info.

(f) View of the alert development including a comment

Parent Alerts

This alert is part of the following parent alerts:

Alert ID	Workflow Status	Severity	Duration	Start Time	Event Count	Tags
236890	Seen	Critical	7 h, 18 min (ongoing)	07:13, 2019-05-07	52	
Manually aggregated 2 alerts						

(h) List of associated Meta Alerts

Figure 74: Details of an alarm

4.13.1.3.2 Right column The right hand column comprises of plots for each of the indicators relevant for generating the alarm. A graph is created for each tuple of indicator and object on which events were registered for this alarm. Each graph is titled with the name of the object, the indicator and the detector responsible for raising this alarm. The format is “**Object name** — **Indicator (detector)**”. Each graph shows the trend of the indicator for a time range starting just before the first event associated with the alarm occurred until just after the last event associated with the alarm. The times of earliest and most recent events are indicated by

red vertical bars. To assess the trend of the indicator, in dark grey the indicator's trend one week ago for the same time period and in light grey for the same time period four weeks ago are also visualized to ease comparisons of the traffic trend. See Figure 75.




Figure 75: Trend of an indicator for an alarm

4.13.2 List of associated events

The list of events associated with an alarm lists time and type of the event further information regarding the registered traffic pattern and the assessment by the detector. The list can be accessed using the link labeled “[View all events...](#)” in the section “[Alert Details](#)” of an alarm summary.

Summary Alerts Objects Mitigations Topology Map Visual Explorer Sparklines Detectors Preferences ▼									
Critical Alert 236876 (muted)									
Toggle exact/humanized values Dates/Times in CEST ▼									
Alarm Begin/End today, 07:13 - 08:28 (1 h, 15 min) 1 affected object, 47 events									
Events First/Last today, 07:09 / 08:18 Merge... Unmute Send to TTS scrolling									
All displayed values are shown as amount per minute.									
2 Pages « 1 2 »									
Event ID	Severity	Timestamp	Event Type	Detector	Source	Threshold Value	Trigger Value	Absolute Value	Description
116852286	Info	2019-05-07 08:18	staticdeviation_above	SYN ACK Ratio		0.01	0.01	0.02	Line GE10/ triggered Static Deviation with observation 0.02, relative distance 0.01 is greater than the event threshold 0.01.
116851900	Info	2019-05-07 08:13	staticdeviation_above	SYN ACK Ratio		0.01	0.01	0.02	Line GE10/ triggered Static Deviation with observation 0.02, relative distance 0.01 is greater than the event threshold 0.01.
116851823	Info	2019-05-07 08:12	staticdeviation_above	SYN ACK Ratio		0.01	0.01	0.02	Line GE10/ triggered Static Deviation with observation 0.02, relative distance 0.01 is greater than the event threshold 0.01.
116851449	Info	2019-05-07 08:07	staticdeviation_above	SYN ACK Ratio		0.01	0.01	0.02	

Figure 76: List of events associated with an alarm (part view)

The web view in Figure 76 lists the events resulting in an alarm. The navigation area shows the relevant times of generation and closure of the alarm.

The content area contains the list of events. The events that actually triggered the alarm are also included. The list is strictly ordered by the time of the occurrence of single events and cannot be sorted. If the list contains more than 25 entries, pagination will spread the entries over multiple pages. The total number of events as well as navigation controls to switch between pages are located above and below the table.

For each event the following data is listed:

Event ID This is a consecutive number to ease differentiation of individual events.

Severity Every single event is assigned a severity level by the detecting detector. The levels “Info”, “Warning” and “Critical” are similar to the severity of alarms.

Timestamp The time when the event was detected.

Event Type The detecting algorithm is listed as it might be relevant for the analysis. There are several different types of events, the details are described in [2].

heartbeat These events indicate that no NetFlow data was received for the particular router.

holtwinters__below, holtwinters__above, staticdeviation__below, staticdeviation__above These events occur if the monitored traffic indicator differs significantly from its predicted value. The detection method differs depending on the conduct of the monitored indicator:

staticdeviation static conduct of the indicator

holtwinters periodic conduct of the indicator

The suffix of the event’s type indicates whether the prediction was exceeded or undercut.

threshold__below, threshold__above These events indicate that a predefined fixed threshold was exceeded or undercut by the monitored traffic indicator. Up to four thresholds can be defined for each object, resulting at exceeding or undercutting in events with severity levels “Warning” or “Critical”.

Detector The name of the detector. If the entry is “deleted” the detecting detector has been deleted in the meantime by changes to the configuration.

Source This is the unambiguous name of the router or line defined by the or the descriptor of the network used within the NeMo system on which this event was detected. Hence, source is the location of the detection and should not be mistaken for source or destination of a potential attack. This is also a link to the particular object.

Threshold Value This is the threshold, that was exceeded resulting in the event. All detectors facilitate thresholds to create events. The particular semantics depend on the type of event:

heartbeat The threshold is the time in minutes that has to have past since the last import of NetFlow data before creating an event.

threshold__above, threshold__below The defined threshold that has been exceeded leading to this event.

staticdeviation, holtwinters These detectors use a prediction model, based on the relative aberration of the monitored value to its predicted value. A threshold of .01 requires a relative aberration of 1%.

Trigger Value This is the actual value leading to the event. The interpretation depends on the type of event:

heartbeat The time since the last update.

threshold The actual value of the indicator collected at that minute.

staticdeviation, holtwinters Events that base on predictions the actual aberration of the prediction is printed.

Absolute Value The trigger value does not have to match the absolute value of the monitored indicator. This is especially relevant for detectors comprising predictions and compare these with the actual value. To assess the level of the traffic, the absolute value of the monitored indicator is specified.

A combination of the last three values could look like this: An event is created, when the monitored indicator deviates by .1 from the prediction. For a current value of 136 and a prediction of 100, the values are:

Threshold Value = .1 The configured threshold value. An event is created if the actual value deviates by 10% from its prediction.

Trigger Value = .36 The actual deviation of the monitored indicator from its prediction is

$$\frac{AbsoluteValue - Prediction}{Prediction} = \frac{136 - 100}{100} = \frac{36}{100} = .36$$

Absolute Value = 136 The actual value of the monitored indicator.

Description This field comprises of a free text description of the event. It may contain further information, characteristic for the particular type of event. The following types of description may appear:

```
No data has been received for Router <object_name> in the last
<minutes> minutes. The last update was at <YYYY-mm-dd
HH:MM:SS>
```

This description of a “[heartbeat](#)” event indicates that for the named router no NetFlow data was stored in the system. The name of the router matches the name specified in the . The most recent import of NetFlow data had this time stamp.

```
<object_type> <object_name> triggered Static Deviation with
observation <humanized_value>, relative distance
<humanized_distance> is greater than the event threshold
<humanized_threshold>
```

This description was created by a “[staticdeviation](#)” event. In addition to the name of the object it contains:

humanized_value The actual value of the monitored indicator in rounded format.

humanized_distance The relative aberration between the actual value and its prediction in rounded format.

humanized_threshold The threshold which violation triggered the event, displayed in rounded format.

```
<object_type> <object_name> triggered Holt-Winters, predicted
<humanized_prediction>, got <humanized_value>, and
 $\frac{\text{value} - \text{prediction}}{\text{prediction}} = \text{<humanized\_distance>}$  > <humanized\_threshold>
```

This description of a “[holtwinters](#)” event contains the following additional information:

humanized_prediction The prediction of the actual value as calculated by the detector, in rounded format.

humanized_value The actual value of the monitored indicator, in rounded format.

humanized_distance The relative aberration between the actual value and the prediction, in rounded format.

humanized_threshold The threshold value whose exceeding triggered the event, in rounded format.

In the list of events, as depicted in Figure 76, various values are shown. These values can be really large. To enhance readability, delimiters are used.

To enhance readability even further the exactness of values can be turned off. By clicking the link labeled “[Toggle exact/humanized values](#)” next to selecting the time zone, the display format can be changed. The “[humanized](#)” display shortens the values to two positions after the decimal point and facilitates the usual abbreviations K = Kilo, M = Mega and G = Giga. This will change the display of the absolute value “3,604,001.00” to “3.60 M”. This “rounded” format is then used throughout the event descriptions for all values.

If not stated otherwise, all visible average values are converted to time slices of one minute.

4.13.3 Analysis view of an alarm

Starting a new analysis or viewing an existing analysis leads to the view as shown in Figure 77. This view integrates and groups various views to support a consistent workflow for analyzing alarms.

The analysis view visualizes the network traffic over a selected time period, this view is similar to the network analysis in Section 4.12.2. The navigation area is extended by controls to select a time period and buttons to create reports and analyses.

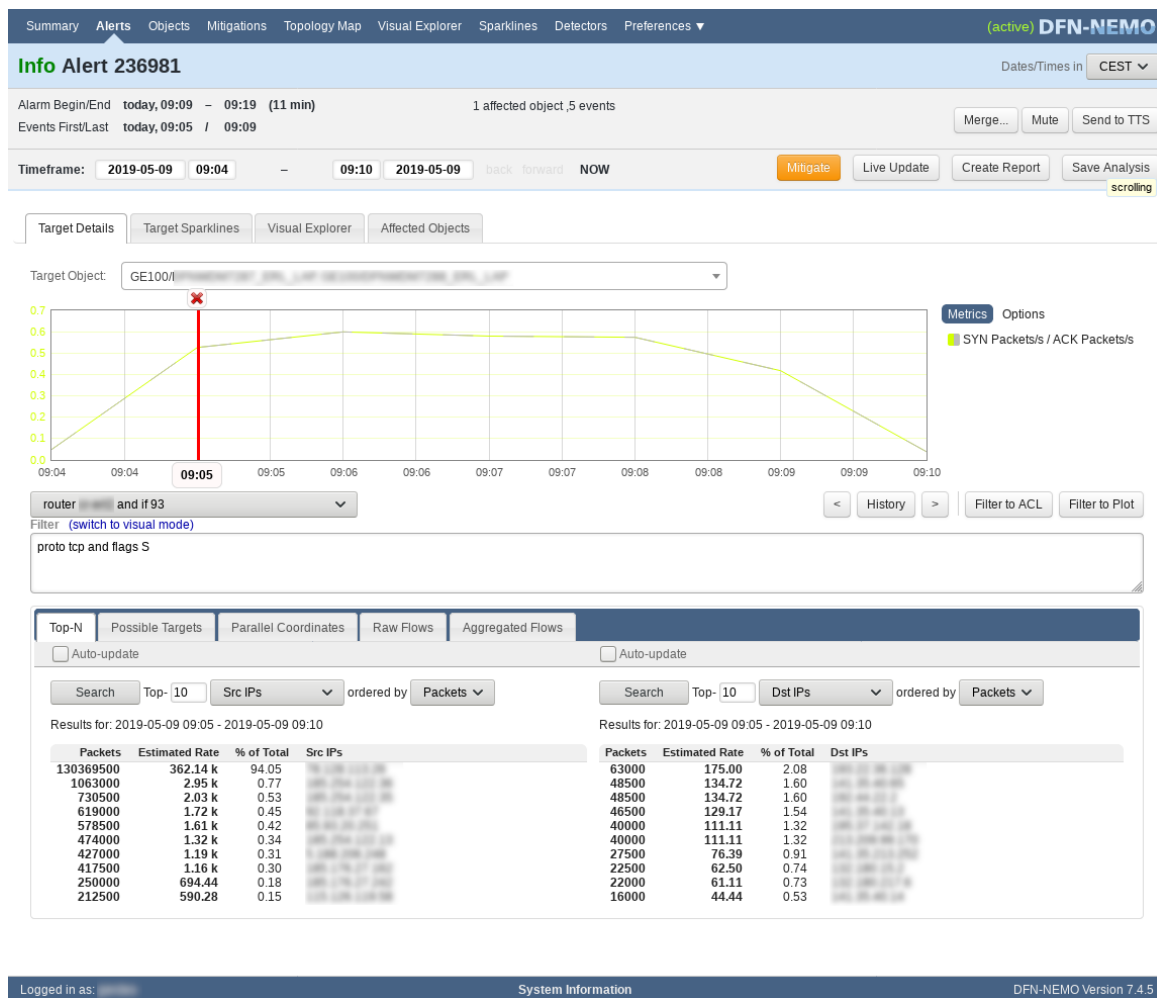


Figure 77: Analysis view of an alarm

The default value for the time period when starting an analysis is shortly before the start time of the alarm until its end. The visualization of begin and end of the alarm as well as times for the first and last event in the navigation area are also links that change the visualized time period accordingly, see Figure 78.




Figure 78: Controls to set time period are also links

4.13.3.1 Live Update Next, the analysis view can switch to a live update, integrating new data as soon as it has been processed. Clicking the button labeled “Live Update” fixates the end of the visualized time period to the actual time, and updates the graph once a minute. Figure 79 shows deactivated controls while live view is enabled. Clicking the “Live Update” button again will disable the live view and prevents the graph from being automatically updated

once a minute. The time period displayed last will be displayed until the user chooses a different time period using the (once again activated) controls.




Figure 79: Controls when live view is activated

As the live view cannot be used as grounds for creating reports and analyses of alarms, these buttons are deactivated as well while live view is enabled. Instead, a checkbox labeled “[Update cursor to latest timestamp](#)” is added, which—if ticked—will update the NetFlow cursor to the most recent time slice when the graph has been updated. This will produce a view on alarm traffic that is regularly updated.

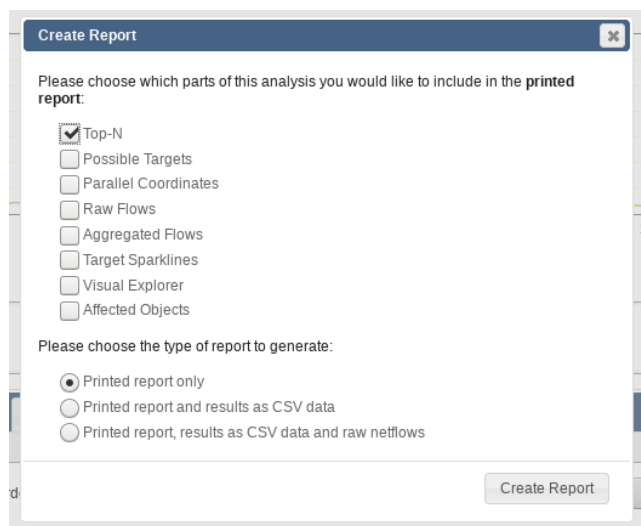


Figure 80: Dialog to create a report

4.13.3.2 Creating reports All settings selected in the report dialog will be saved and restored for this analysis when opening the dialog again.

Exported Data The button labeled “[Create Report](#)” will open a dialog as shown in Figure 80. A variety of analysis results is available to be used in the report. The default selection is limited to results of the “Top-N” evaluation, other results can be added by ticking their respective checkbox.

Export format The default is to only create a report.

Another option is to also publish the results in CSV³¹ format, to be used in spreadsheets or other applications.

³¹“Comma Separated Value” a data format which orders data in lines, separating them using a predefined separator. Therefore, multiple data connections can be stored in a single file, while each connection is stored in a single line `time, source ip, source t, destination IP, destination port`. See also https://en.wikipedia.org/wiki/Comma-separated_values.

Another option is to also include NetFlow raw data in nfdump format.

Exporting the report Clicking the button labeled “[Create Report](#)” will show the report with a dialog on top which will not be part of the printed report, see Figure 81.




Figure 81: Dialog to print the report

The export can be created by printing the page using the web browser’s printing option, hitting the enter key or by clicking the link labeled “[print this page](#)”, both resulting in opening the web browser’s printing dialog to print the page to paper³².

If exports (CSV or nfdump) were selected, links to files are listed (see Figure 81b, but this may take a short while, see Figure 81a). The CSV data is presented in a ZIP³³ file, which includes not only the CSV data but also an XML³⁴ file describing the parameters of the analysis. These include alarm ID, start and end time, analysis target object, position of the NetFlow cursor and applied NetFlow filters.

The dialog can be closed clicking either the “X” in the upper right corner of the dialog or the button labeled “[Close](#)”.

Saving the Analysis All settings made for the analysis, all selected NetFlow filter, visualized objects and configured time periods are part of the analysis and can be saved using the button labeled “[Save Analysis](#)”. Clicking the button opens a dialog in which a comment for this analysis may be specified. An ID for the NOC’s TTS may also be specified, see Section 4.13.1.1. An entry in the list of analyses is created. With this, existing analyses can be extended by further information and stored as a new analysis. Clicking “[Save Analysis](#)” will always create a new revision of the analysis.

4.13.3.3 Target of the Analysis An analysis has a defined target of the analysis at all times: an object, whose NetFlow data are the basis for the analysis. This target is the object on which the alarm was created when a new analysis is started. But during the analysis the target may change. The object is likely to be an object on which the detected attack traffic is especially good to identify, easing the separation of normal and abnormal traffic using NetFlow

³²An export to PDF is likewise possible if “print to file” is selected in the web browser’s print dialog. See your web browser’s help for instructions.

³³A data format for compressed data, reducing the size of the data files and works as a container for multiple files in one file, see also [https://en.wikipedia.org/wiki/Zip_\(file_format\)](https://en.wikipedia.org/wiki/Zip_(file_format)).

³⁴“Extensible Markup Language”, a data format for hierarchically structured data, see <https://en.wikipedia.org/wiki/XML>.

filter. This also eases the overall analysis. Mostly these are objects close³⁵ to the source or destination of an attack.

4.13.3.4 Analysis view The analysis view consists of four views, all accessible using the registers at the top of the content area, see also Figure 77.

Target Details This is the default view and shows the details of the current analysis target. It is similar to the detailed view of this object (see Section 4.12. The graph of indicator trends shows relevant indicators of the alarm when a new analysis is started. The relevant indicators are those, for which the alarm reported an abnormality. The start time for the NetFlow analysis will be the time of the first event associated with the alarm.

For an alarm, reporting a high rate of TCP SYN packets to TCP ACK packets the graph shows the particular metric.



Figure 82: Automatically generated NetFlow filter for an alarm

Based on the relevant indicators a NetFlow filter is generated automatically, see Figure 82.

Target Sparklines This sparkline view shows the traffic overview of the target of the analysis and is similar to the sparkline view explained for the object view in Section 4.12.2

Visual Explorer This register shows a Visual Explorer (see Section 4.7), comprising a plot for the object on which the alarm was generated and—if the object is a line or network—a plot for a router—ideally a core router. The plots show the relevant indicators for the alarm in analysis.

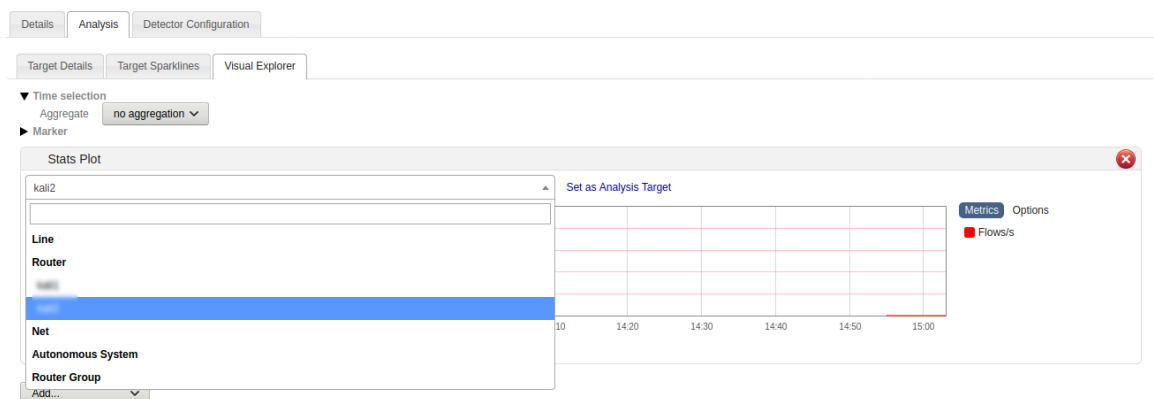


Figure 83: Selecting a new target for the analysis

³⁵Close means close in a topology point of view, the number of edges or routers in between is small, in best cases 1 and 0, respectively.

This Visual Explorer allows to select a new target for the analysis. Therefore a link labeled “[Set as Analysis Target](#)” is assigned to each plot. The link points to the view [Target Details](#) using the new target, see Figure 83.

This Visual Explorer does not create bookmarks, all settings will be stored when the analysis is saved.

Affected Objects This register allows to retrace potential attack traffic throughout the topology of the Network. NetFlow data is used to identify entry and exit routers as well as any intermediate routers. Its usage is explained in the following Section 4.13.4.

4.13.4 Tracing back an alarm

To determine the objects that were used by potential attack traffic NeMo provides the view as depicted in Figure 84.



Figure 84: Controls to determine the paths of potential attack traffic for an alarm

Tracing back traffic using raw NetFlow data is extremely costly in terms of CPU and is therefore limited to five minutes starting with the time of the start time defined on the target details page. The traffic to be searched will be defined by a NetFlow filter expression, while a filter specified in the detailed view can be used clicking the link labeled “[Use Target Filter](#)”. The filter can be adjusted, i.e. to remove TCP flags which are not collected on every router throughout the Network. Enclosing TCP flags in the filter might lead to inconsistent results.

The filter expression is then followed by selecting a “[Start Router](#)”. More routers may be included in the search, which is necessary in certain cases for a full trace back, especially when the packets use paths in the Network which are not collected in the raw data. To reduce the duration of the search time the number of selected additional routers should be small.

Further controls specify the type (“[Trace by](#)”):

Packets high packet rate

Bytes high data rate

Flows high number of connections

The search is executed when “[Find Affected Objects](#)” is clicked, but before the filter is checked for unsuited or erroneous conditions. Notifications will be shown if required. If the filter expression is suited the trace back is conducted that will take several seconds depending on the number of selected routers and depending on the filter expression. The results are a list of objects and a statement about the traffic’s data amount on these routers.

First the data amount is calculated for the [Start router](#), measuring as selected in packets, bytes or flows. If traffic matching the filter is identified, the router is added to the result list including the detected data amount. Of this data amount a threshold at 10% of this value is determined which is then used as a lower boundary, meaning that monitored traffic must have a certain data rate to be accepted as relevant traffic for this alarm.

All outgoing interfaces on this [Start router](#) are then checked for matching traffic with an amount larger than the threshold. If this is the case for an interface, the line connected at this interface and the router at the other end of the line are added to the results listing—unless the router is not modeled in NeMo. The router will then also be added to the search list of routers. All routers having been searched are ignored.

This search continues until no interface does show any relevant traffic amount of matching traffic³⁶. The results' listing is then shown next to a topology map of the Network in the content area, see Figure 85.

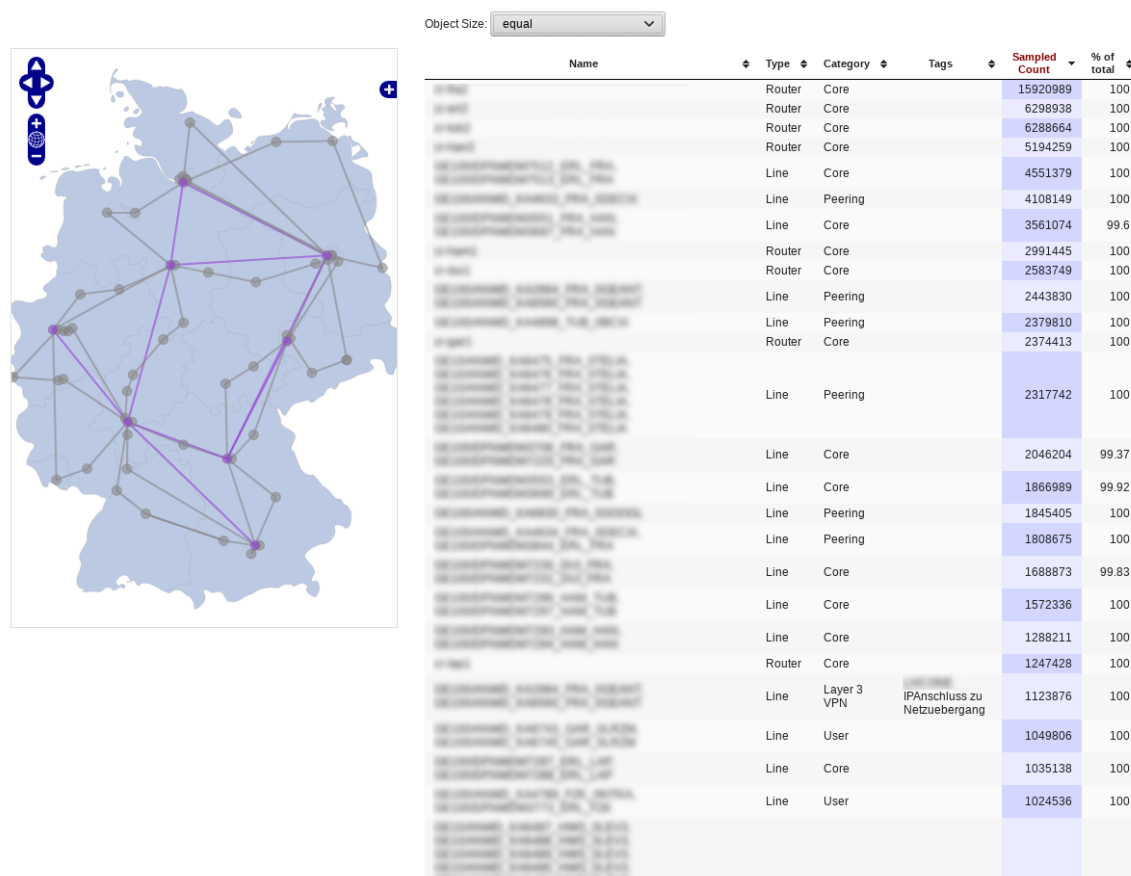


Figure 85: Finished trace back of a potential attack (part view)

On the left hand side is the topology map (see also Section 4.6) in which all objects are marked that had seen a relevant amount of matching traffic. On the right hand side a table shows each object including the traffic volume and the share of this traffic of the overall traffic at this object.

The list is sorted by descending traffic amount. The sort order can be changed using the

³⁶The algorithm is more complex than explained here, since routers in the Network do only register a part of the traffic in the NetFlow data they provide.

columns' titles. The values for traffic volume are taken from the NetFlow data and therefore do not respect the sampling rate implemented at the router.

4.13.5 Mitigation of an alarm

To create countermeasures against this traffic pattern mitigations can be created right from the analysis view. The orange colored button labeled “Mitigate” in the upper right corner (see Figure 77) will start the process that is described in detail in Section 4.14.2. The following data set is transferred and used to preliminary fill the form:

- source and destination IP range,
- protocol type,
- source and destination ports.

4.14 Detailed view of a Mitigation

The detailed view of a mitigation lists the name of the mitigation and its revision number in the extended navigation area, the naming scheme is explained in Section 4.5. The revision number indicates the number of stored edits to this mitigation.

There are three sections in the content area, the side bar on the right hand side and three registers “[Mitigation Details](#)”, “[Target Details](#)” and “[Statistics](#)”, which are visible in Figure 86. The [Target Details](#) register and the [Statistics](#) register will be discussed in Section 4.14.4 and Section 4.14.5, respectively.

4.14.1 Mitigation Details

The screenshot displays the 'Mitigation 0932-nemo-erkennung.' (Version 0) page in the NeMo application. The top navigation bar includes tabs for Summary, Alerts, Objects, Mitigations (selected), Topology Map, Visual Explorer, Sparklines, Detectors, and Preferences. The main content area is divided into three sections: Mitigation Details, Target Filter, and Statistics. The Mitigation Details section shows the description 'Mitigate Alert #236987', a red box for Protected Ranges, and a table for Rules. The Target Filter section shows the status 'Inactive' and 'nada-stu2: Inactive'. The Statistics section shows the status 'Not available' and a 'Request Authorization' button. The right sidebar contains sections for Mitigated Alert, Mitigation Target, and Comments / History.

Name	Direction	Source CIDRs	Dest. CIDRs	Protocol	Source Ports	Dest. Ports	Action	bps Limit
+ Add empty rule or Choose rule template...								

Countermeasures

- ☐ IP Header Validation
- ☐ TCP Flags Validation
- ☐ TCP Handshake Validator
- ☐ GeoIP Filter
- ☐ PCAP Filter

Save Save & Start Mitigation

Status: Inactive
nada-stu2: Inactive

Authorization
Status: Not available
Request Authorization

Mitigated Alert
Info Alert #236987
High ratio of SYN packets to ACK packets.

Mitigation Target
Router cr-lap1
Target Filter:
proto tcp and flags S

Comments / History
2019-05-09, 16:32
Mitigation created for alert #236987 by [user]
Add comment

Figure 86: The detailed view of a mitigation (the register [Statistics](#) is not active and therefore not accessible since the mitigation has not been deployed successfully yet)

The screenshot displays the 'Mitigation Details' page for 'Mitigation 0931-nemo-erkennung.test (Version 2)'. The top navigation bar includes links for Summary, Alerts, Objects, Mitigations (active), Topology Map, Visual Explorer, Sparklines, Detectors, and Preferences. The page title is 'Mitigation 0931-nemo-erkennung.test (Version 2)' with a 'Dates/Times In' dropdown set to 'CEST'. Below the title, there are tabs for 'Mitigation Details', 'Target Filter', and 'Statistics'. The 'Mitigation Details' tab is active, showing a description 'Mitigate Alert #236606' and 'Protected Ranges' '193.0/23, 2001::/48'. The 'Rules' section shows a table with one rule named 'test', direction 'incoming', protocol 'UDP', and action 'Drop'. The 'Countermeasures' section lists several options like IP Header Validation, TCP Flags Validation, etc. The right sidebar shows 'Status: Failed (Version 2)', 'Authorization' status 'valid', 'Mitigated Alert' 'Info Alert #236606', and 'Mitigation Target' 'as 65052'. The bottom status bar shows 'Logged in as: [user]', 'System Information', and 'DFN-NEMO Version 7.4.5'.

Figure 87: The detailed view of a mitigation for which the [Statistics](#) register is active

4.14.1.1 Mitigation Details This register lists the following information:

Description A description of the mitigation, usually referencing the alarm.

Protected Ranges The IP address ranges that are protected by this mitigation, containing IPv4 and IPv6 network ranges. The encircled (i) shows the net range for which the current user is authorized to create mitigations.

Rules The filtering rules that are active in this mitigation. A rule comprises the following data:

Name A freetext description of the rule. It is only used if an error occurs.

Direction Incoming The incoming network traffic of the protected ranges is filtered

Outgoing The outgoing network traffic of the protected ranges is filtered

Source CIDR The source IP network range for which this filter rule shall be applied, the default is to set this to the protected ranges if the direction of the rule is set to outgoing.

Destination CIDR The destination IP network range for which this filter rule shall be applied. The default is to set this to protected ranges if the direction of the rule is set to incoming.

Protocol The type of protocol for which this rule is to be applied. The selection contains: **TCP**, **UDP** and **ANY** (all/both).

Source Ports The ports used to send packets.

Destination Ports The ports to which packets shall be delivered. This enables the rule to filter specific traffic.

Action Actions can be specified that shall be executed for all IP packets for which the rule matches:

Pass The packets are delivered and not filtered / dropped. This is especially relevant if complex rule sets are created, that also match legitimate traffic that shall be delivered.

Drop Matching packets are dropped by the mitigation components and filtered from the network traffic.

Rate-Limit A limit for traffic bandwidth can be specified for matching traffic (An additional text field is shown).

Countermeasure A variety of countermeasures can be implemented. Each can be explained further by clicking the gray encircled “i”.

IP Header Validation Validates the conformity of the IP header to IP header specification

TCP Flags Validation Validates the correctness of TCP flags

TCP Handshake Validation This checks, whether the client can actually perform a TCP handshake before the connection is established. This is used to prevent SYN flooding.

GeoIP Filter Network traffic originating from selected countries based on GeoIP databases is dropped.

PCAP Filter This allows the specification of complex filter rules based on the “[TCPdump packet filter syntax](#)”³⁷.

There are two buttons placed underneath the specification of countermeasures. These are used to “[Save](#)” and “[Start](#)” the mitigation. A protected address range³⁸ and a filter rule are

³⁷See more information: <http://www.tcpdump.org/manpages/pcap-filter.7.html>.

³⁸A protected address range is accepted only if the user is authorized to create mitigations for this address range.

required to save the mitigation. If this is the case, the “[Save](#)” button is enabled. Starting the mitigation requires an authorization and is explained in Section 4.14.2. The “[Start](#)” button is enabled once an authorization is granted and selected.

4.14.1.2 Side bar The side bar lists various data about the mitigation.

State The current state of the deployment of the mitigation on the routers (mitigating components) and hence state of the mitigation. Possible states are:

Inactive The mitigation has not been deployed yet. Most common issue is an incomplete authorization.

Failed This state is reached if errors or problems occur while deploying the mitigation or if the mitigating components cannot be accessed.

Stopped This state is reached if a mitigation is stopped manually and all filter rules were removed successfully from all routers.

Active All rules were successfully deployed on all components. The authorization has not been revoked yet.

Partially active The filter rules were deployed not to all components or not all filter rules were installed on all components successfully.

Unknown This error state will be shown if NeMo is unable to assess the current state unambiguously.

There are temporary states:

Starting This state is reached if an authorization is active and any user activates the mitigation. This starts the deployment process.

Stopping This state is temporary if a mitigation was stopped in the system and NeMo is currently removing the filter rules from the components.

State of filter rules on components Figure 86 lists `nada-stu2`. Possible states are [Inactive](#), [active](#), [failed](#), [stopped](#), [starting](#) and [stopping](#) which all have the same meaning as for the state of the mitigation.

Authorization The state of the authorization of this mitigation. Possible values are:

not available If neither an existing one was selected or a new one requested.

selected If an authorization was selected, the state of the authorization is reused. This is explained in Section 4.14.3.

Mitigated Alarm The alarm for whose analysis the mitigation was started. This consists of a direct link to the alarm’s detailed view and its severity (as determined by the system).

Mitigation Target The network traffic that is to be influenced by the mitigation consisting of an object and a filter describing the network traffic (type of protocol, TCP flags, ports, abnormalities, ...). The detailed view is linked to.

History A log of changes to the mitigation in reversed chronological order (most recent entries are on top). Comments may be added by users to this log.

4.14.2 Create a mitigation

A mitigation is created from an analyzed alarm. Few fields of the detailed view are preliminary filled by NeMo, others must be edited by the user. The [Protected Range](#) is only then specified if the object is a network and if the IP range of the [Protected Range](#) is a subset of the net range for which the user is authorized to create mitigations. The following rule is applied: Either the entire IP range is used or no data is specified.

Before a mitigation can be started an authorization is required. If an authorization was granted the mitigation can be adjusted or stopped without further authorization. A mitigation is required if the net range of the protected range is enlarged.

4.14.3 Selecting an authorization

Before a mitigation can be started an authorization is required. The authorization is granted by a certain group of people specified by the organization to which this IP range belongs.

There are two types of authorizations, which will be explained in the following paragraphs. The dialog to select an authorization is shown in Figure 88.

Name	Granted on	Valid until
<input type="checkbox"/> Infrastrukturschutz_ALL	::0, 0.0.0.0/0	always valid

Select existing Authorization

Request new Authorization

Organisation:

IP-Ranges:

Request Authorization

Figure 88: Dialog to select the type of authorization

4.14.3.1 Infrastructure authorization To protect the infrastructure of the Network its NOC is authorized to start mitigations even without further authorization of affected organizations. This authorization is always valid, therefore the mitigation can be deployed at once.

Selecting this type of authorization is regulated and restricted by job instructions.

4.14.3.2 Organization bound authorization These authorizations require a temporally limited authorization for DDoS mitigations by the affected organizations. Already valid authorizations can be reused or new authorizations can be requested.

An authorization code is send to mobile devices specified by the organization. The code must be confirmed by any person authorized to do so. Two commands are implemented: **Start** and **Stop**. Only the first reply for each command is processed, whereas a **Stop** command cannot be replaced by a **Start** command. A **Start** command grants the authorization and allows NeMo to deploy the mitigation rules. The authorization can be reused until it is either revoked by a issuing a **Stop** command by any authorized person or if the life span of the authorization has ended.

An authorization has one of four states:

New An authorization request has not been send to the mobile devices.

Pending This state is temporary while the request was send but no command has been received from an authorized phone number.

Valid The request was confirmed by at least one authorized person but no person has stopped the authorization yet.

Invalid The authorization was stopped by any authorized person or the authorization was temporary and expired or an unrecoverable error occurred³⁹.

Organizationally bound authorizations are only temporary valid. The time span is seven working days⁴⁰. If the authorization expired a new authorization must be requested or selected if the mitigation shall be edited.

If a **Stop** command is received for any authorization or the authorization expires, this does not effect active (deployed) mitigations. An authorization is required only when the mitigation is changed.

4.14.4 Target Filter

This view—depicted in Figure 89—enables the user to conduct various analyses on the traffic that still passes the countermeasures in order to improve the scope of the mitigation. The options are similar to those explained in Section 4.12.2.

³⁹An error can occur if the SMS Gateway is unreachable or reports an error.

⁴⁰The working days are state specific.

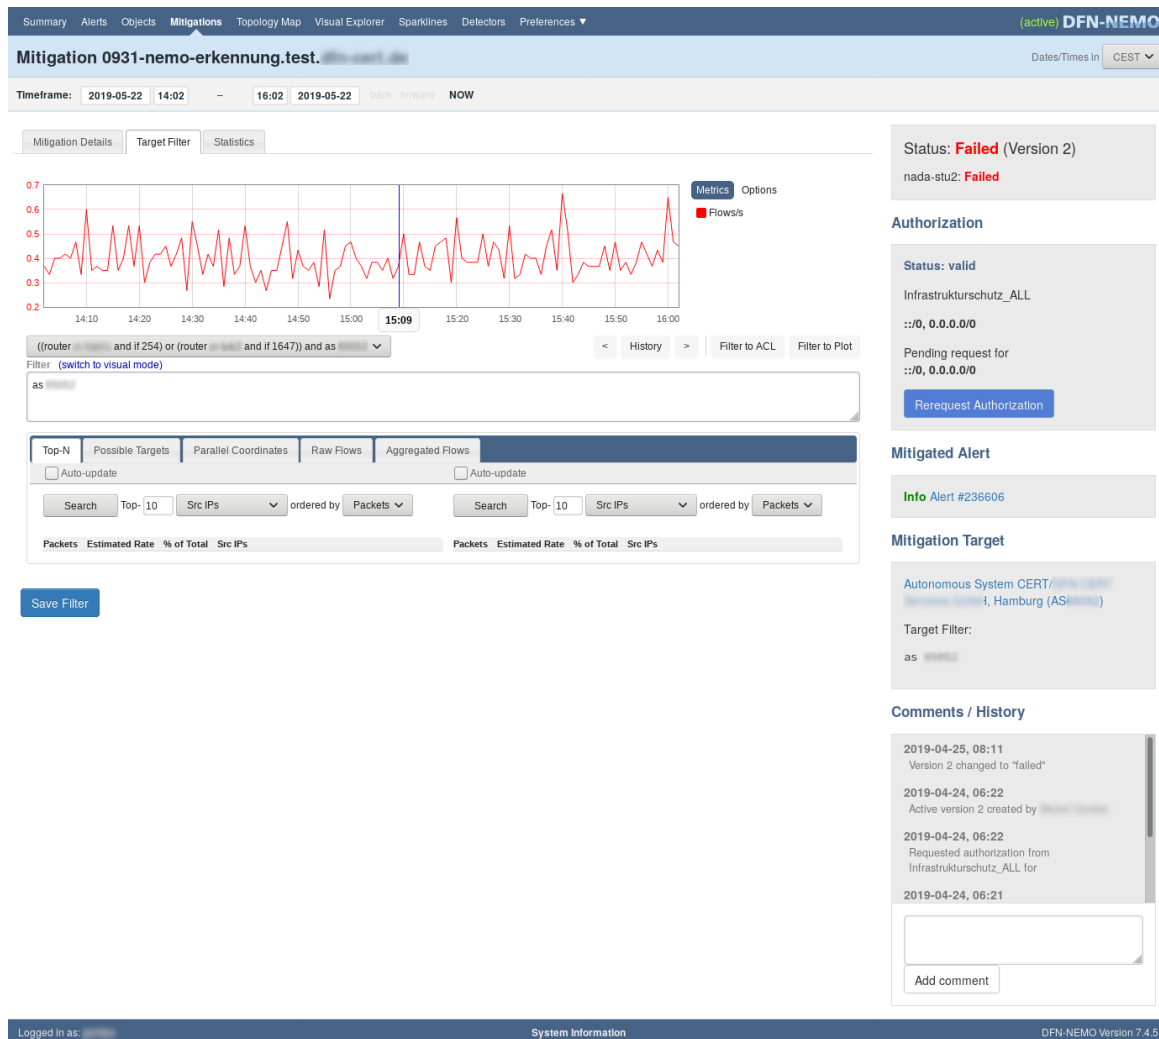


Figure 89: An analysis view for traffic that still passes the mitigation

4.14.5 Statistics for an active or formerly active mitigation

As soon as a mitigation is successfully deployed on any component, the [Statistics](#) register is activated. Its structure is divided in two parts, as depicted in Figure 90:

Attack Traffic All network traffic described by the [Mitigation target](#) is visualized in a timeline graph.

Rules For each filter rule the matching traffic is visualized in a timeline with relative vertical scale.

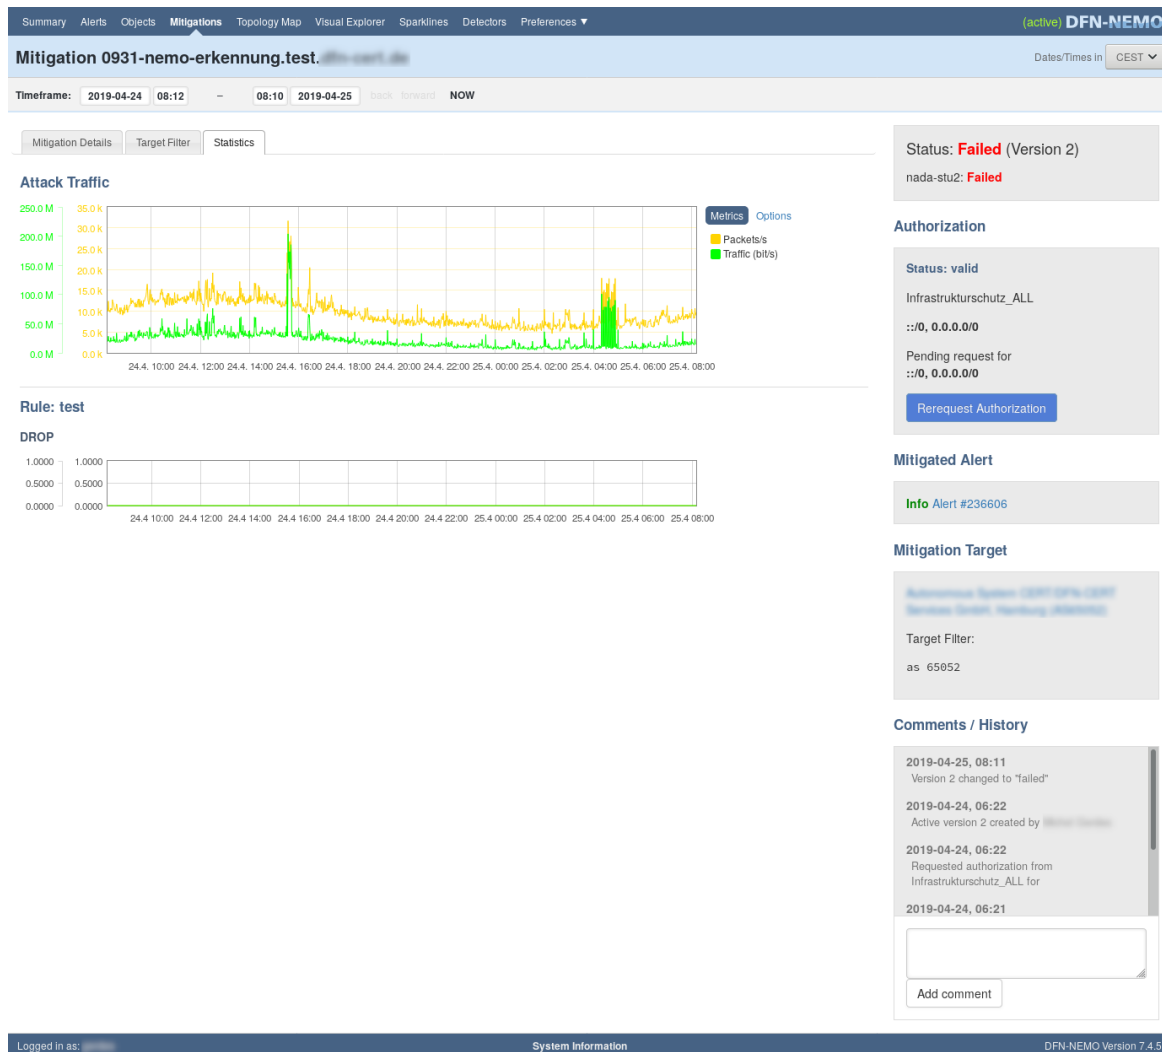


Figure 90: Statistics register

The default is to visualize two metrics in the **Attack Traffic** plot, each having its own vertical scale: “Packets per second” and “Transferred data rate per second”. Facilitating the link “Metrics” further metrics can be added to the graph. The “Options” link changes the appearance of the plot. Both options are identical to those described in Sections 4.7.2.2.2 and 4.7.2.2.3.

5 References and Glossary

References

- [1] BSI Bonn. Erkennung und abwehr von ddos-angriffen im ivbb - einschätzung und bewertung aktueller bedrohungen, 2007.
- [2] DFN-CERT Hamburg. Dfn-nemo / ddos: Erkennungsverfahren und algorithmen, 2010.
- [3] DFN-CERT Hamburg. Dfn-nemo / ddos: Schnittstellen und formate, 2010.

5.1 Glossary

Alarming procedure Procedure which identifies alarm situations and generates alarm messages by correlation of individual observations on the objects of the Network.

Simple Alarm An alarm that summarizes events identified by the detectors of the DDoS application, enabling the tracking of specific traffic situations.

Child alarm Also suppressed alarm, an alarm associated with a meta-alarm.

Indicator An indicator for traffic in Network recorded in the DDoS application. The indicators are derived directly from or derived from the NetFlow and SNMP data exported from the routers (basic indicator).

Meta-Alarm An alarm that aggregates simple alarms to reduce the number of messages and keep information density high. A meta-alarm contains one or more simple alarms, a simple alarm belongs to none, one or more meta-alarms.

Mitigation Mitigation of an anomaly is the creation and introduction of measures to reduce and minimize the effects of the anomaly.

Sparkline A qualitative representation of the course of a variable that changes over time and is designed to give an impression of the behaviour of the variable in terms of trends, periodic behaviour or parallelism with other variables, for example, with a small space requirement. A sparkline does not allow statements about the absolute values of a variable.

Suppressed alarm See child alarm.