

Ebook

Streamlining Network Monitoring and SIEM Integration with Cubro Solutions

Table of contents

Introduction.....	3
Network Monitoring vs. SIEM; What's the Difference?.....	4
What is Network Monitoring?	5
What is SIEM?	6
The Difference Between Network Monitoring and SIEM.....	7
Logs vs KPI.....	9
NetMon feeds SIEM	10
Disadvantages of SIEM Solution	11
Benefits of Cubro's integration with NetMon and SIEM tools.....	13
Optimising data feeds to SIEM	15
Importance of Filtering data to SIEM.....	20
Alleviating the burden on SIEM.....	22
Summary	28



Introduction



This ebook serves as a **comprehensive guide** to **understanding Network Monitoring** and **SIEM** (Security Information and Event Management) and the **differences between the two**. Network monitoring is the process of monitoring network traffic to identify issues, optimize network performance, and ensure network security. SIEM, on the other hand, is a security solution that combines security information management and security event management.

In this ebook, we will delve into the importance of network monitoring and SIEM in ensuring network security and maintaining network performance. We will also explore how Cubro Solutions offer a cost-effective and simplified approach to network monitoring and SIEM integration. Cubro Solutions provide powerful tools for network monitoring and SIEM integration, and can be easily customized to meet the specific needs of any organization.

Whether you are a network administrator, security analyst, or IT manager, this ebook will provide valuable insights into the world of network monitoring and SIEM integration, and how Cubro Solutions can help streamline your organization's security and network monitoring processes.

Network Monitoring vs. SIEM; What's the Difference?

Have you ever heard of Network Monitoring (NetMon) and Security Information and Event Management (SIEM)? These two tools are widely used, but what exactly do they do? Moreover, how do IT teams use NetMon and SIEM platforms to prevent cyber attacks and safeguard critical data for enterprises?



What is Network Monitoring?



Network Monitoring refers to a tool, software or hardware continuously observing the network and the data that flows through it. Depending on how the tool monitors the network, it may collect data directly as it moves across the network or gather data stored by a network node. After collecting the data, the tool interprets it and presents it through a dashboard on your device. Using a network monitoring tool, you can easily monitor your network's performance from one location.

The monitoring tool can identify any issues with the network, such as indications that a particular part of the network is about to fail. The tool will then send alerts to your IT team, enabling them to resolve the problem promptly. A network monitoring solution helps your enterprise to keep the network in check, ensuring that it operates optimally.

What is SIEM?



According to the leading analysts, Security Information and Event Management (SIEM) refers to collecting event data generated by various monitoring, assessment, detection, and response solutions deployed across different environments such as applications, networks, endpoints, and the cloud. SIEM capabilities typically include identifying threats through correlation, user and entity behaviour analytics (UEBA), and integration with security orchestration, automation, and response (SOAR) for response management. Additionally, standard SIEM integrations also include security reporting and continuously updated threat content via threat intelligence platform (TIP) functionality. While SIEM is primarily a cloud-based service, it may also support on-premises deployment.

The Difference Between Network Monitoring and SIEM

NetMon and SIEM can be deployed as a Software-as-a-Service (SaaS). However, the main difference between NetMon and SIEM is that SIEM follows the data across everything, including the network and endpoint devices. At the same time, NetMon solely focuses on data passing through the network. This does not necessarily imply that SIEM is superior to NetMon.

Using Network Monitoring, an enterprise concentrates all its resources on safeguarding the critical data flowing through the network. In today's business landscape, companies store vast amounts of information in the cloud, and data is transferred from endpoint to endpoint via the network. However, the drawback of relying solely on a NetMon tool is that it overlooks everything else, leaving the system vulnerable to external attacks.

On the other hand, with SIEM, the aim is to obtain a comprehensive overview of the entire system through a single pane. SIEM collects information from every endpoint, every network pass, every open application, and so forth.

However, the drawback of SIEM is that it can result in 'information overload', where the IT team has too much data to analyse and may miss critical details.





In conclusion, it would be beneficial for the IT team to use both tools to optimise the security of the enterprise's network. This way, the workload can be divided among team members, and the information can be better analysed and acted upon.



Network Monitoring (NetMon) solutions can deliver data to Security Information and Event Management (SIEM) solutions for analysis and correlation. NetMon solutions typically monitor network traffic and collect data about network activity, such as device connections, bandwidth utilisation, and protocol usage. This data can be used by SIEM solutions to provide a more comprehensive view of an enterprise's security posture by correlating network activity with other security events, such as endpoint activity or user behaviour. By combining data from multiple sources, SIEM solutions can help enterprises detect and respond to security threats more effectively.

Logs vs KPI

Logs and KPIs (Key Performance Indicators) are two types of data used for monitoring and analysis in various fields, including IT and business.

Logs refer to a record of events or transactions within a system, network, or application. They are often used for troubleshooting, debugging, auditing, and compliance purposes. Logs can contain a wealth of information, such as timestamps, source and destination IP addresses, protocols, actions performed, errors, and other metadata.

KPIs, on the other hand, are specific metrics or indicators used to evaluate the performance or effectiveness of a particular aspect of a system or process. They are often used to track progress toward specific goals or objectives and to identify areas for improvement. Examples of KPIs in IT may include network uptime, response times, server load, and user satisfaction.

While logs provide detailed information about system activity, KPIs offer a high-level view of performance that can help identify trends and patterns over time. Both logs and KPIs are useful for monitoring and managing IT systems, but they serve different purposes and are often used in conjunction with each other.

NetMon feeds SIEM

Today it is ubiquitous that a NetMon solution can deliver data to SIEM solutions for analysis and correlation. NetMon solutions typically monitor network traffic and collect data about network activity, such as device connections, bandwidth utilisation, and protocol usage. This data can be used by SIEM solutions to provide a more comprehensive view of an enterprise's security posture by correlating network activity with other security events, such as endpoint activity or user behaviour. By combining data from multiple sources, SIEM solutions can help enterprises detect and respond to security threats more effectively.

Disadvantages of SIEM Solution

While SIEM solutions offer several benefits for an enterprise, they also have certain drawbacks that should be considered. Here are some common disadvantages associated with SIEM solutions:

1.

High Cost:

SIEM solutions can be expensive to purchase, deploy, and maintain. Enterprises must invest in the hardware, software, and personnel required to implement and manage these solutions.

2.

Complex Implementation:

Implementing SIEM solutions can be complex and time-consuming. Enterprises must have a clear understanding of their security needs and objectives to ensure the SIEM solution is correctly configured.

3.

Resource-Intensive:

SIEM solutions can consume significant resources, such as storage, processing power, and network bandwidth, to collect and analyse data. Enterprises must ensure that they have sufficient resources to support these solutions.

Disadvantages of SIEM Solution

4.

Skill Requirements:

Effective SIEM implementation and management requires a skilled IT team with expertise in security, data analysis, and threat intelligence. Enterprises must ensure that they have the necessary resources and skill sets to deploy and manage these solutions effectively.

Overall, SIEM solutions can significantly benefit an enterprise's security posture. However, enterprises must carefully consider the potential drawbacks and weigh them against the benefits to determine if SIEM is the right solution for their needs.

5.

False Positives:

SIEM solutions may generate many false positives, which can be time-consuming for IT teams to investigate and remediate. This can result in alert fatigue and make it difficult to identify genuine security threats.

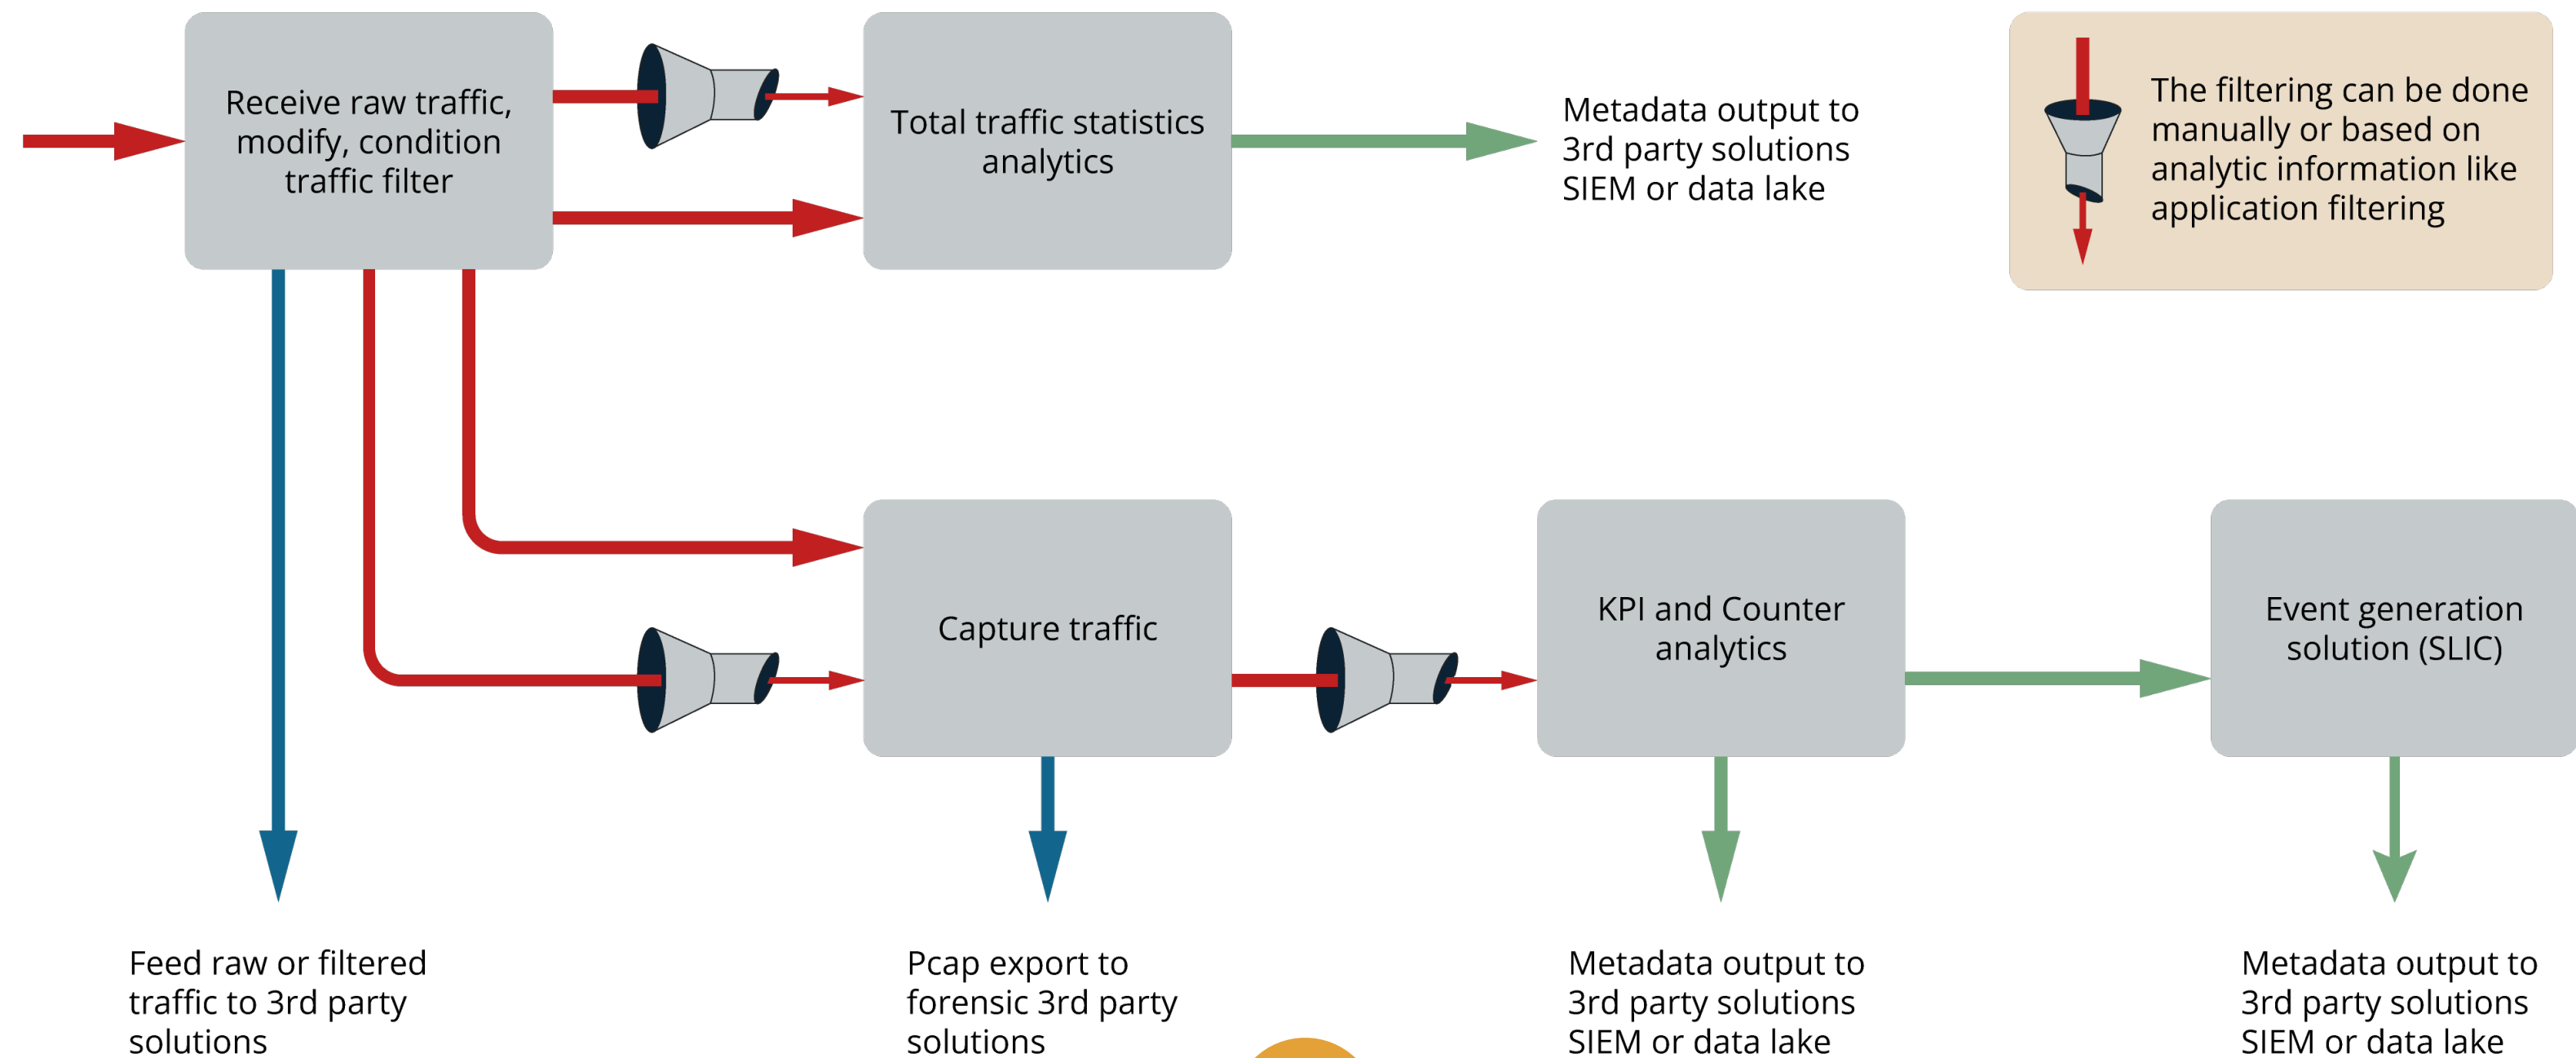
6.

Delay:

When working with RAW data, the delay becomes a significant concern when using SIEM. The initial delay occurs during the production of LOG files, which is not a high-priority task. Next, the SIEM client requires time to upload the LOG, followed by processing time by the SIEM. While this delay may not be significant for statistical purposes, it is critical for troubleshooting. Additionally, this delay makes it challenging to correlate LOG and metrics from multiple sources.

Benefits of Cubro's integration with NetMon and SIEM tools

Cubro solutions offer a cost-effective and simplified approach to network monitoring and SIEM integration. With our solutions, organisations can reduce the cost and complexity associated with monitoring and analysing network traffic.



1.

Cubro solutions can improve the efficiency of data feeds to SIEM by reducing the number of flows and eliminating irrelevant information.

2.

Filtering data is crucial for a SIEM, as irrelevant data can cause alert fatigue and false positives. Examples of irrelevant data include:

- Noisy WEB 2.0 traffic resulting from multiple flows and virtual endpoints within a short period
- Non-security-related data from widely-used applications like Netflix and YouTube. Analysing such traffic would result in unnecessary costs and time expenditures.
- Redundant data from complex networks. As networks become more complex, and with the prevalence of leaf-spine architecture, traffic is often sent to NetMon Solutions multiple times, leading to redundant analysis of the same data.

3.

To alleviate the burden on SIEM, it is best to avoid providing raw data like CDRs and instead offer KPIs that provide insights into both security and performance. Sharing only the essential events can further streamline the process.

Optimising data feeds to SIEM

Efficient data feeds to the SIEM can be achieved by reducing the number of flows and removing irrelevant information

As described in the previous slides, there are several disadvantages to performing flow-based computation, mainly due to a lack of resources (CPU, memory, storage). Aggregating metadata time window based solves most of these problems and allows a much higher computation throughput without losing essential metadata attributes.

Significant metadata cannot be extracted from single flows as they mainly contain technical data. The important information, such as device or user-specific data, is obtained from multiple flows. This information includes the type of traffic generated from a specific device, the extent of services used by a specific device, the volume of traffic generated by a user over time and the servers involved. The time window-based approach focuses on this critical data.

The key is collecting Metadata for uploads, downloads, and internal traffic, as well as DPI information from a device/user perspective over time. Extracting the essential information out of big data streams with the benefit of not wasting resources and storage for absolutely meaningless information (contained within single flows) is the crucial point.



The time window-based approach offers several essential data perspectives, such as the service perspective and client/device perspective.

Service perspective:

The service perspective provides insights into the amount of traffic generated for each service, usage frequency over time, and the volume of uploaded and downloaded traffic.

Client/Device perspective:



The client perspective enables analysis of network usage for each client or user, including upload and download activity, the services and locations used, and usage frequency over time.



Server perspective:

Time-window based

We create a bucket for each application for every client, if needed, and collect/count all packets within a specified time window (configurable). Once the window closes, we produce and enrich an XDR and transmit it. This approach reduces traffic on most far points, thereby avoiding workflow constraints along the way to the database.



Flow-based

One flow is generated for each 5-tuple connection, but many of these flows cannot be recognized as Amazon related because the external domain cannot be resolved. Once a flow is produced, it is sent to the Flow Cache. A Flow Cache can contain hundreds of thousands or even millions of entries, utilizing memory resources.

Once the flows expire, they're exported to the NetFlow Collector, which constantly analyses and archives them for future reference.

Flow based

The flow-based solution is concerned about sessions, including perpetually open sessions (TCP handshake).

Typically, a flow probe has a limitation in terms of the number of flows (FPS), not bandwidth

Irregularly terminated/established sessions are also problematic, as an irregularly terminated session stays open until the timeout and consumes unnecessary resources.

A session where the initial handshake is not seen for any reason will not be detected. This could be an issue for IoT devices, as they may communicate infrequently, and it could take days for a session to be detected again.

The resulting metadata stream size is configurable and typically ranges from 2% - 3% of the input traffic. See the table for performance figures.

Time-window based

We are not concerned about perpetually open sessions (TCP handshake) or irregularly terminated/established sessions.

The configurable time window offers the option to balance between performance constraints and the granularity of the output.

The resulting metadata stream size is configurable, typically ranging between 0.1% - 0.5% of the input traffic.

With the exponential growth in traffic, both solutions have their advantages and disadvantages. However, a time-window-based solution is much more efficient and helps to reduce Capex and Opex costs.

With the exponential growth in traffic, both solutions have their advantages and disadvantages. However, a time-window-based solution is much more efficient and helps to reduce Capex and Opex costs.

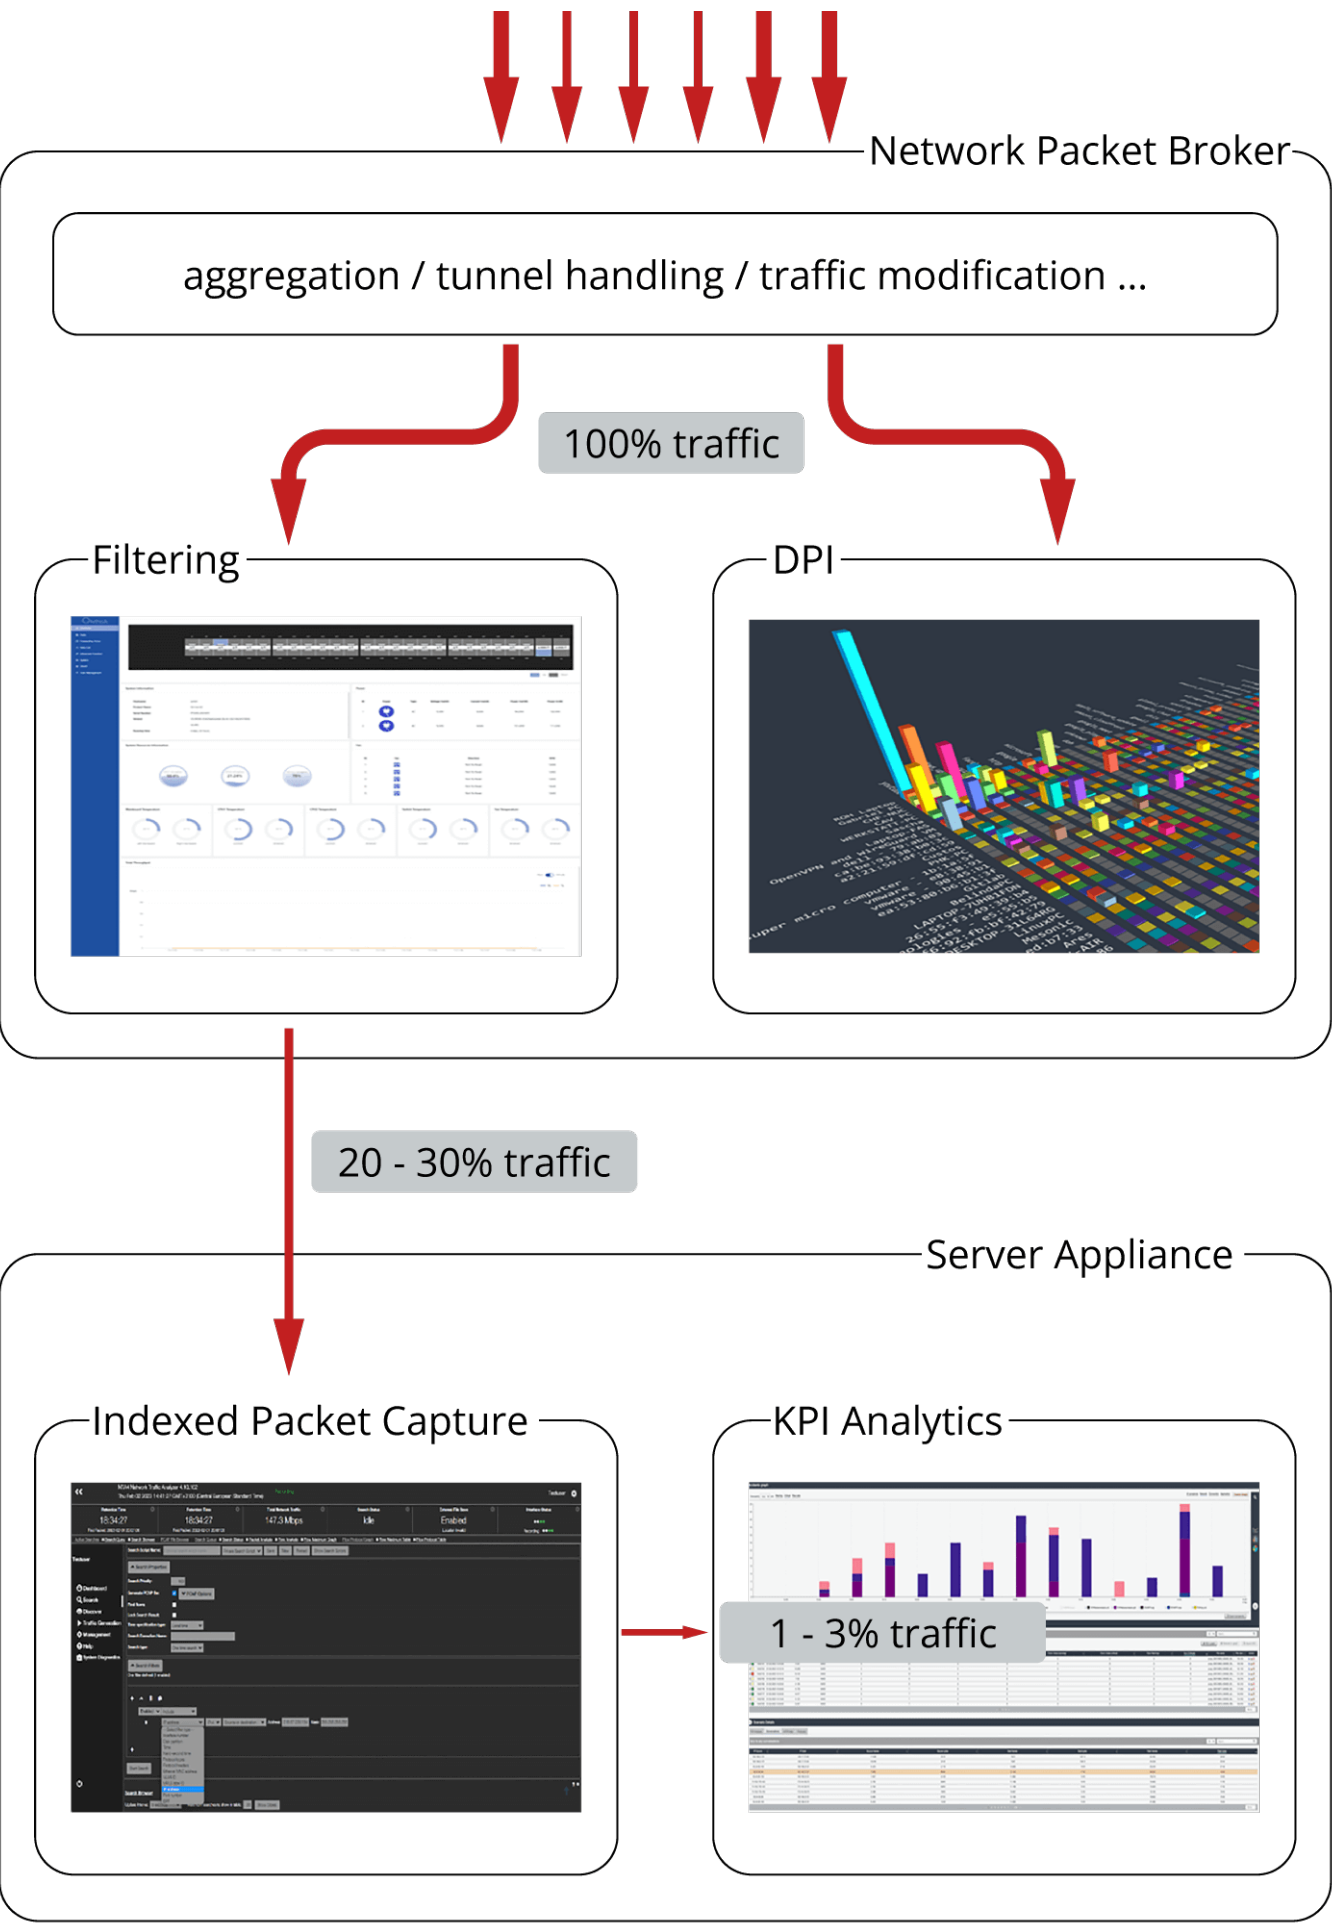
Estimated FLOW based Metadata retention time

Bandwidth in Gbit/s	1	5	10	30	60	90	Retention time in days
0,1	0,060	0,180	0,330	0,990	1,950	2,940	
0,5	0,180	0,810	1,620	4,860	9,720	14,040	
1	0,330	1,62	3,240	9,720	19,440	29,160	
5	1,62	8,10	16,20	48,60	97,20	145,80	
10	3,24	16,20	32,40	97,20	194,40	291,60	
50	16,20	81	162	486	972	1.458	
100	32,40	162	324	972	1.944	2.916	
500	162	810	1.620	4.860	9.720	14.580	
1000	324	1.620	3.240	9.720	19.440	29.160	in TB storage

Estimated Cubro Metadata retention time

Bandwidth in Gbit/s	1	5	10	30	60	90	Retention time in days
0,1	0,00	0,014	0,029	0,086	0,173	0,259	
0,5	0,01	0,072	0,144	0,432	0,864	1,296	
1	0,03	0,14	0,29	0,86	1,73	2,59	
5	0,29	1,44	2,88	8,64	17,28	25,92	
10	0,72	3,60	7,20	21,60	43,20	64,80	
50	1,44	7,20	14,40	43,20	86,40	129,60	
100	2,88	14,40	28,80	86,40	172,80	259,20	
500	14,40	72	144	432	864	1.296	
1000	28,80	144	288	864	1.728	2.592	in TB storage

Importance of Filtering data to SIEM



Not all data is relevant to a SIEM, and including irrelevant data can lead to alert fatigue and false positives. Here are some examples of data that may be irrelevant to a SIEM.

Raw traffic

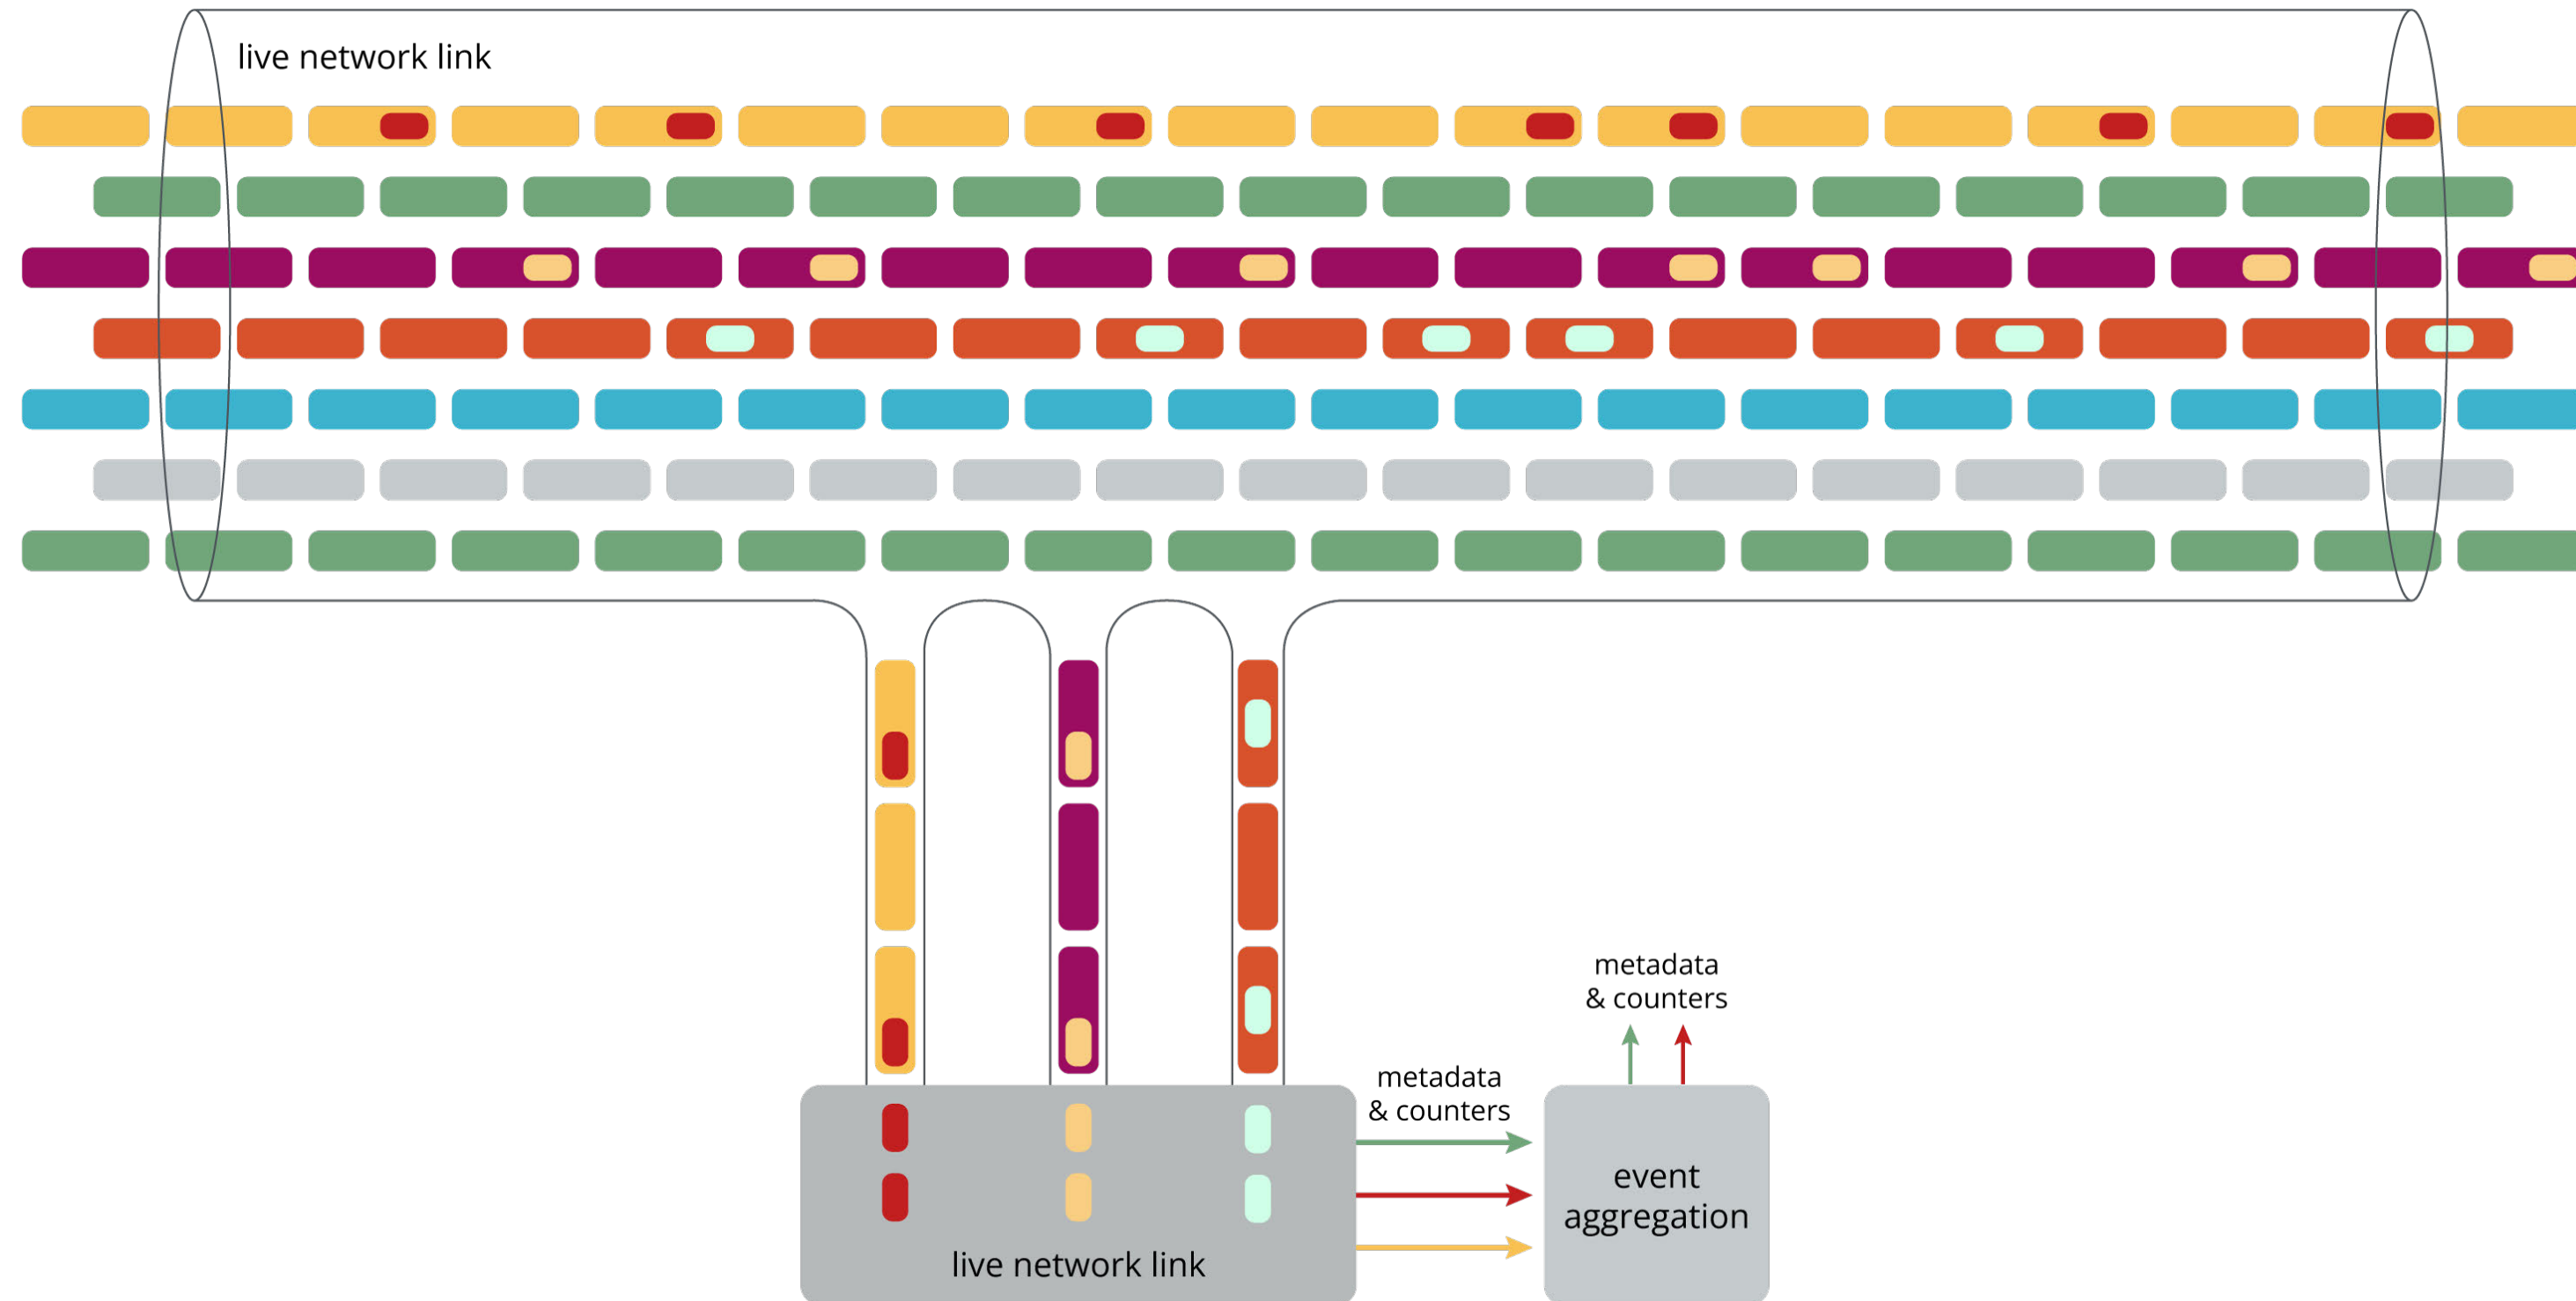
Bandwidth in Gbit/s	1.000	3.000	6.000	12.000	24.000	Retention time in hours
0,1	0,08	0,5	0,5	1.000	2.000	
0,5	0,25	0,75	1,5	3.000	6.000	
1.000	0,46	1375.000	2,75	5,5	11.000	
5.000	2,25	6,75	13,5	27.000	54.000	
10.000	4,50	13,5	27.000	54.000	108.000	
50.000	22,50	67,5	135.000	270.000	540.000	
100.000	45,00	135.000	270.000	540.000	1.080	In TB storage

After filtering, more than 70% of the traffic can be removed.

Bandwidth in Gbit/s	1.000	3.000	6.000	12.000	24.000	Retention time in hours
0,1	0,024	0,075	0,15	0,3	0,6	
0,5	0,075	0,225	0,45	0,9	1,8	
1.000	0,138	0,4125	0,825	1,65	3,3	
5.000	0,675	2,025	4,05	8,1	16,2	
10.000	1,35	4,05	8,1	16,2	32,4	
50.000	6,75	20,25	40,5	81	162	
100.000	13,5	40,5	81	162	324	In TB storage

Alleviating the burden on SIEM

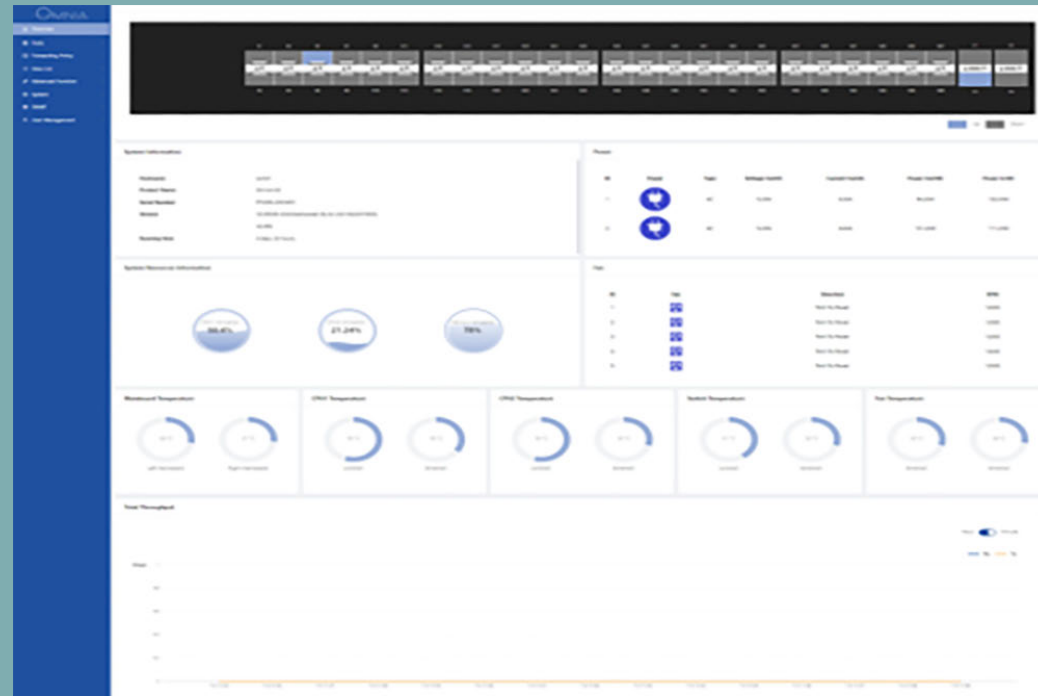
To alleviate the burden on SIEM, avoid offering raw data like CDRs and instead provide KPIs that offer insights into both security and performance. Ideally, sharing only the essential events would further streamline the process.



The first step in the process involves extracting the relevant traffic connections from live links, which can only be accomplished with a Network Packet Broker (NPB). A Cubro NPB can perform this task efficiently due to its hardware-based approach.

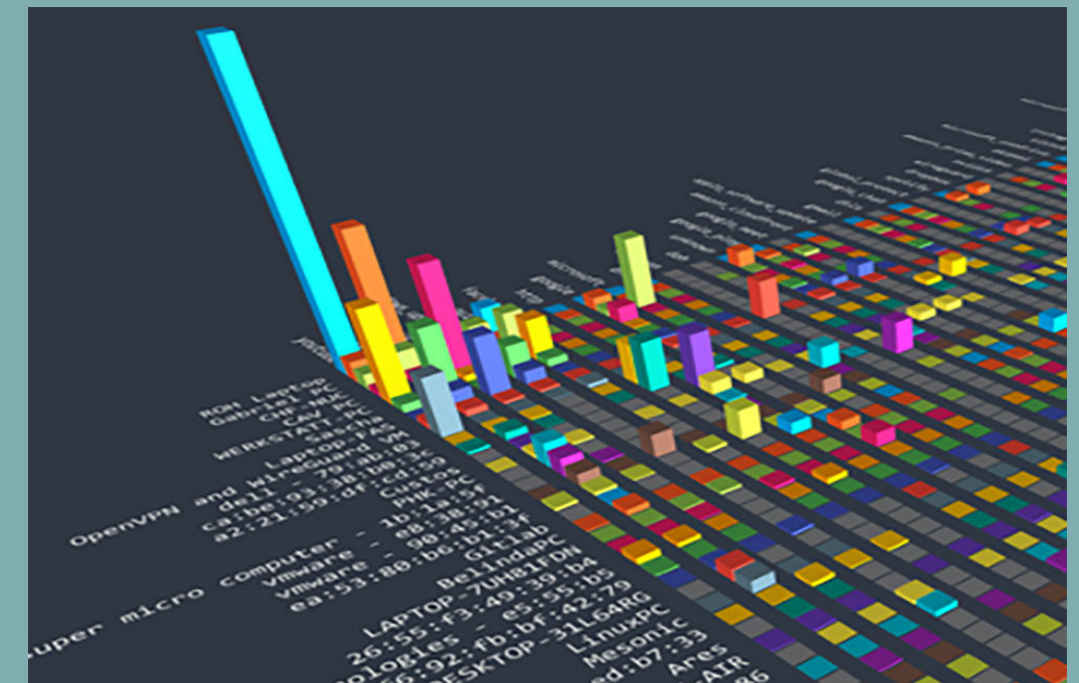
The second step involves using an analytic engine to extract only the relevant packets from these network connections. Thanks to the upfront filtering accomplished in the first step, the analytic engine can operate more efficiently.

The final step involves aggregating the metadata and counters, which in this case, represent a relatively small amount of data. This data is then correlated with events and incidents and fed into SIEM or any other database.



The initial stage in managing network traffic is to process it in hardware. Raw network traffic often requires processing to enable effective monitoring. This traffic can be in tunnels, distributed over multiple physical links, or fragmented. Conditioning the traffic to make it suitable for monitoring is a hardware task performed by a network packet broker.

Real-time statistics on the total traffic received are crucial for obtaining an overview of the current network situation. When using deep packet inspection (DPI), this information becomes even more valuable as it provides details on how filtering can be configured. Although this information can be sent to a SIEM as a CDR, doing so can result in an excessive amount of data that may overwhelm the SIEM and lead to increased costs.



The second function of the capture is that Cubro Analytics uses the capture to retrieve relevant information via API calls and then perform up to 30k possible metrics.

These metrics can be any kind of measurement, from any kind of delay calculator to all OSI layers or only pure counters. This measurement can be provided as a stream to any 3rd party tool; as a transport stream, we prefer Kafka, but also CSV is a possible option when the data is not too much.

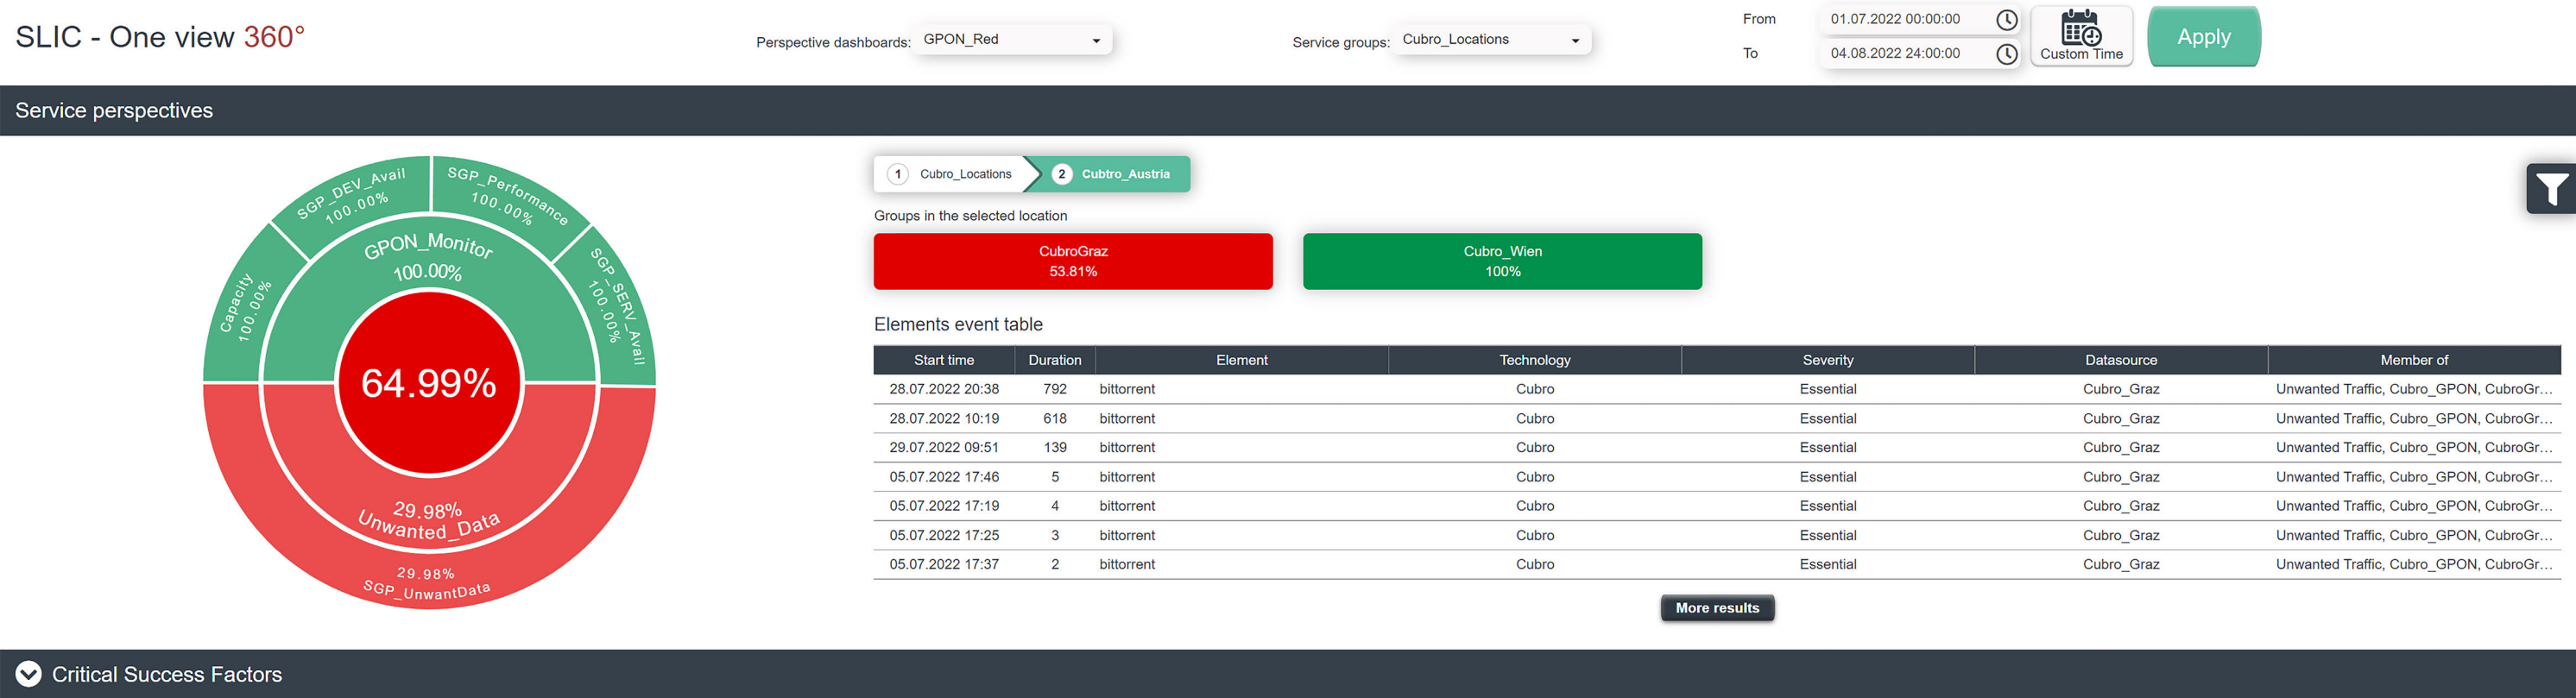
The advantage of this approach is that we offload the SIEM by the amount of data and processing; in this case, the information is available in near real-time.

Example: A media gateway event log can easily have a few megabytes per 15 minutes, but if only a few specific measures are needed, such as identifying unsuccessful calls, the Cubro KPI with these measures will only be a few kilobytes in size and require no further processing.



And finally, these KPIs can be converted as standalone or in combination with other KPIs in a single event.

This reduces the load on the SIEM close to zero.



Example: Unwanted or harmful traffic on a network

This is a common request, but it can be challenging to accomplish due to the high cost, particularly the cost of the data lake or SIEM, which makes the project difficult to implement. The Cubro solution can reduce the amount of information to a single event. It can indicate the presence of unwanted and harmful traffic during a specific time frame, along with a KPI showing the device that generated the traffic and the time of occurrence.

Summary

Cubro Solutions offer a cost-effective and simplified approach to network monitoring and SIEM integration. By optimising data feeds to SIEM and filtering out irrelevant data, organisations can reduce the cost and complexity associated with monitoring and analysing network traffic. While there are advantages and disadvantages to both flow-based and time-window-based solutions, the latter is more efficient and helps to reduce Capex and Opex costs. Overall, the importance of filtering data to SIEM cannot be overstated, as including irrelevant data can lead to alert fatigue and false positives. By sharing only essential events and offering KPIs that provide insights into both security and performance, organisations can streamline the process and make better use of their resources.

Some potential values that Cubro Solutions offer for cost-effective and simplified network monitoring and SIEM integration include:

1.

Scalability:

Cubro Solutions can be scaled up or down depending on the size of the network, making it an ideal choice for businesses of all sizes.

2.

Cost-effectiveness:

By providing a single platform for network monitoring and SIEM integration, Cubro Solutions can reduce costs associated with managing multiple tools.

3.

Ease of use:

Cubro Solutions are designed to be easy to use and implement, so businesses can quickly start monitoring their networks and integrating with SIEM solutions.

4.

Customization:

Cubro Solutions can be customized to meet specific business needs, allowing businesses to tailor their network monitoring and SIEM integration capabilities to their unique requirements.

5.

Advanced capabilities:

Cubro Solutions offer advanced features such as packet capture, analysis, and filtering, making it easier for businesses to identify and address network issues.

Thank
you

For more information, contact us at support@cubro.com.

Distributed by Telnet Networks Sales@telnetnetworks.ca
800-561-4019