

# Detection 2

## Threat Intelligence + Security Operations Centres

David Crooks

UKRI STFC

EGI CSIRT/IRIS Security team

[david.crooks@stfc.ac.uk](mailto:david.crooks@stfc.ac.uk)



# Introduction

- Building on yesterday, consider the development of Security Operations Centres
- Combination of people, processes and technology to augment our detection capabilities

# Recall the Landscape

- Recall from our discussion of the landscape
- Must work together
- Must share information about ongoing incidents

# The approach

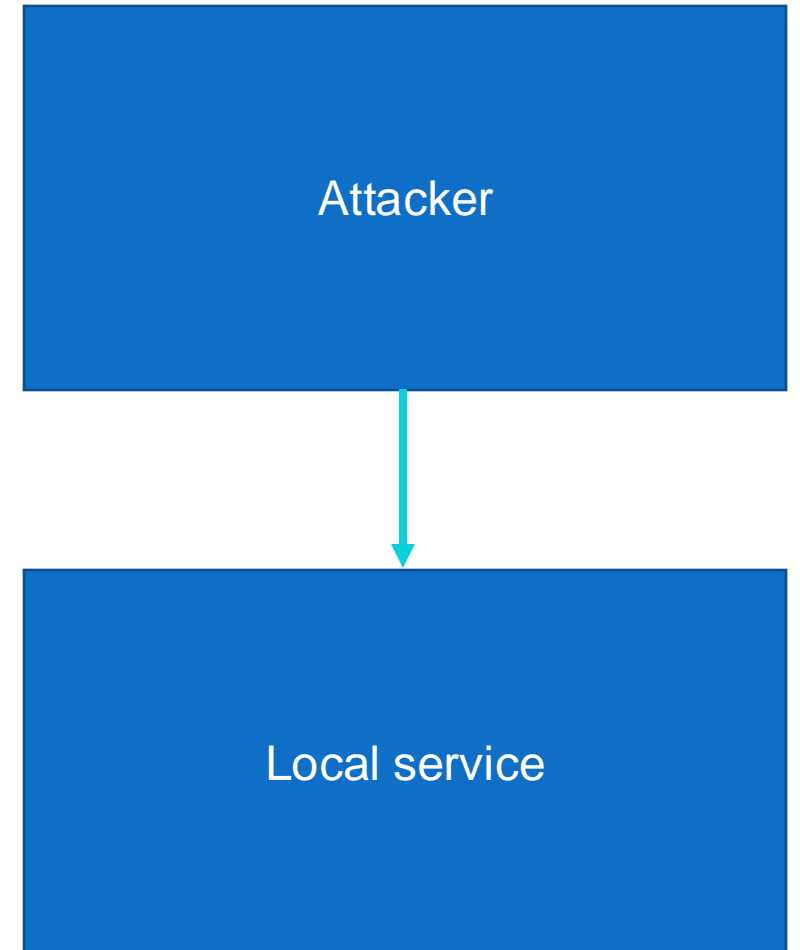
- During incident response, we generate useful Indicators of Compromise (IoCs)
- Give a fingerprint by which to identify malicious traffic and your or another site
- We **must** share this information

# Threat intelligence

- Threat intelligence is the collection of these IoCs in a way that can help identify an attack
- It **does not** include specific information about your facility or service

# Local vs attacker evidence

- Let's imagine that your Drupal CMS has been compromised via a recent unpatched vulnerability
- You're doing incident response and have a lot of information about the impact on your services
- You have some information on where the attacker came from and what actions they took on your network



# What information to share?

- The information that is useful to others are the **IoCs that identify the attacker**
- **Not** the impact on your service
- “The attacker’s IP was...”  
vs
- “My Drupal with all my group information was hacked and it’s a disaster!”

# Sharing threat intelligence

- Sharing information this way means you are giving others the most important information
- **Without** giving away sensitive information
  - not in the data protection sense here

# Type of IoCs

- Network
  - IP
  - Port
  - Timestamps
- Files
  - Checksums
- TTP information
  - Tactics Techniques Procedures

# Who to share with

- Build trust groups
- Share with others that are similar to you
  - What is useful to me?
  - What is useful to them?
- Make the information as useful as possible

# What makes good intelligence?

- Accuracy
- Timeliness
- Relevance
  
- Bulk lists of IPs are less useful than
  - I saw this set of indicators in active use today and these are developing
  - I saw evidence that X/Y/Z may be affected right now

# Traffic Light Protocol (TLP)

- TLP is a set of 4 designations
- Designed to indicate the conditions under which information can be shared
- And with which audience

# Traffic Light Protocol v2 (TLP)

<b>RED</b>	Not for disclosure, restricted to participants only.	For the eyes and ears of <i>individual</i> recipients only, no further disclosure. In the context of a meeting, for example, TLP:RED information is limited to those present at the meeting.
<b>AMBER</b>	Limited disclosure, restricted to participants' organisations.	Limited disclosure, recipients can only spread this on a need-to-know basis within their organization and its clients. Note that TLP:AMBER+STRICT restricts sharing to the organization only. <b>Sources are at liberty to specify additional intended limits of the sharing: these must be adhered to.</b>
<b>GREEN</b>	Limited disclosure, restricted to the community.	Limited disclosure, recipients can spread this within their community, but not via publicly accessible channels.
<b>CLEAR</b>	Disclosure is not limited	Recipients can spread this to the world, there is no limit on disclosure.

# TLP:AMBER

- For TLP:AMBER we can and **should** specify any specific restrictions
  - Only for security teams
  - Only for **this** security team, but all members of it

# TLP Examples

Example	Category
<b>Information about a vulnerability which impacts our community badly, but is not (yet) public knowledge</b>	
<b>I met my colleague and they had very timely information that would have an extremely high impact if it were to be generally available</b>	
<b>I have information that is timely and relevant about an ongoing incident that would be useful to my fellow incident responders</b>	
<b>I read about a critical vulnerability on The Register and \$GIANTPLATFORM is impacted!</b>	

# TLP Examples

Example	Category
<b>Information about a vulnerability which impacts our community badly, but is not (yet) public knowledge</b>	<b>TLP: GREEN</b>
<b>I met my colleague and they had very timely information that would have an extremely high impact if it were to be generally available</b>	<b>TLP: RED</b>
<b>I have information that is timely and relevant about an ongoing incident that would be useful to my fellow incident responders</b>	<b>TLP: AMBER</b>
<b>I read about a critical vulnerability on The Register and \$GIANTPLATFORM is impacted!</b>	<b>TLP: CLEAR</b>

# Data classification over time

- When determining which designation to use, what are the circumstances under which it will change?
  - We will tell you
  - After two weeks
  - ...
- Specificity is at the heart of all good communication

# Chatham House Rule

# Chatham House Rule



**When a meeting, or part thereof, is held under the Chatham House Rule, participants are free to use the information received, but neither the identity nor the affiliation of the speaker(s), nor that of any other participant, may be revealed.**

# Threat Intelligence technology

- OK, we now have
- **Intelligence**
  - That is timely and relevant
- And we know
- **Who we want to share with**
  - And under which restrictions

# Threat Intelligence technology

- How best to share this?
- Word of mouth
- Email
- ...
- Specific service

# MISP

- Previously **Malware Information Sharing Platform**
- Incredibly flexible threat intelligence sharing tool developed by CIRCL.LU
- Web application with API



<https://www.misp-project.org>

# MISP

The screenshot displays the MISP interface with several key components:

- Malicious activities:** A sidebar on the left showing event details for ID 10878, including metadata like Uuid, Org (CIRCL), and Threat Level (Low).
- Distribution graph:** A donut chart showing the distribution of the event across different sharing groups and communities.
- Event details:** A central panel showing event information such as Threat Level (Low), Analysis (Initial), and a description: "Ransomware found on a production server".
- Event graph:** A network diagram showing relationships between various entities like IP addresses and domains.
- Matched event:** A section showing related events, including one with ID 10728 and a classification of "Ransomware".
- Authentication Failure Data:** Two horizontal bar charts showing the number of failed authentication attempts for various users and IP addresses.
- Achievements of my organization:** A section with congratulatory messages and icons for sharing events, using tags, and using taxonomies.

Below the main interface, there are two dashboards:

### Authentication Failure Data

User	Count
admin	313
test	180
ks365908	146
kimsufi	141
user	131
postgres	123
ubuntu	109
oracle	81
git	72
deploy	69
ftpuser	68
nagios	60
mysql	49
support	39
111111	38
guest	38
testuser	36

### Authentication Failure Data

IP Address	Count
45.141.86.157	357
192.241.175.115	287
162.243.169.176	261
31.184.199.114	180
52.188.40.7	157
185.153.196.230	78
92.246.76.177	67
13.67.32.172	64
159.89.201.59	58
121.241.244.92	57
64.225.58.236	56
118.25.10.238	52
175.107.198.23	52
106.52.251.24	50
54.37.159.12	48
123.206.90.149	48
192.241.155.88	47

### Achievements of my organization

**Achievements Unlocked!**

- Event:** Congratulations, you have shared your first event!
- Tagging:** You have been using tags, good job!
- Taxonomy:** Taxonomies have been used in your events.
- Galaxies:** Galaxies have no secrets for you in this Threat Sharing universe.

**Next on your list:**

<https://www.misp-project.org>

# MISP

- Technical expression of trust
- Share information within a pre-defined set of sites / other MISP instances
  - Tags/comments/...

# MISP

- One of the most important tools we are using and will use
- Genuinely broad usage across gov/commerce/academia

# Research & Education threat intelligence

- R&E threat intelligence instance hosted by CERN
- Grew from activity for WLCG, available to the sector
- Either sync or use API

# R&E threat intelligence

- EGI CSIRT currently distributes IoCs via broadcasts to our sites
- Now working on incorporating threat intelligence sharing directly into our procedures
- Highly relevant intelligence on ongoing incidents to our scope

# Other sector intelligence sharing

- Jisc
  - UK NREN
- SAFER-TRUST
  - R&E trust group designed to support sharing intelligence across national and organisational boundaries

# Security Operations Centres (SOCs)

- We have a great source of intelligence: what now?
- We need to understand what is happening in our service/facility/network
  - Host/network logging
- Let's integrate these

# Security Operations Centres

- From a high level a SOC is the combination of
  - Technology
  - People
  - Processes
- **Note:** “SOC” is unfortunately an overloaded term
  - Take care in which aspect is being discussed

# UK NCSC guidance

- The NCSC has a useful overview of [building a SOC](#)
- “help organisations design a SOC and security monitoring capability proportionate to the threat they face, their resources and assets.”

# UK NCSC guidance

- The NCSC has a useful overview of [building a SOC](#)
- “help organisations design a SOC and security monitoring capability proportionate to the threat they face, their resources and assets.”

**Operating Model**

**Data Sources**

**Threat Intelligence**

**Incident  
Management**

# Operating Model

- Threat profile
  - What is the threat faced by your scope?
- Assets
  - What assets are you trying to protect?
- Governance + Resources

# Teams and Processes

- Does senior management support the SOC capability?
- Who maintains the SOC?
  - Next year?
- Where does the next tranche of hardware come from?
- Who analyses the alerts?

# SOC Pillars

- Informing the SOC
  - Threat intelligence
- Developing capability
  - Turn requirements into technical detection capability
- Detect and respond
- Supporting the SOC
  - Internal vs third-party support
  - Community efforts

# SOC Functions

- Threat Intelligence (TI)
- Threat hunting
- Analytics
- Engineering
- Incident Response
- Incident Management

# Data Sources and Incident Management

- Application Logs
  - Host Logs
  - **Network**
  - Cloud
- 
- How does the SOC fit into existing processes?
    - Does it operate 24x7?
    - Who responds to alerts?

# Security Operations Centres

- From a technology standpoint, a SOC is the combination of
  - threat intelligence
  - fine-grained logging information
  - storage and visualization
  - alerting

# How to deploy a SOC [engineering]

- Understand what scope you need to cover
- What outcome do you want?
- What logging capabilities do you already have?
- What staffing is available to you?
  
- **Start small enough to be useful**
  - MVP (minimum viable product)

# Considerations

- Important to identify a realistic starting point
- Your capabilities with the tools will grow with experience
- Want to make your processes effective rather than throw hardware at the problem
  - You do need some of that!

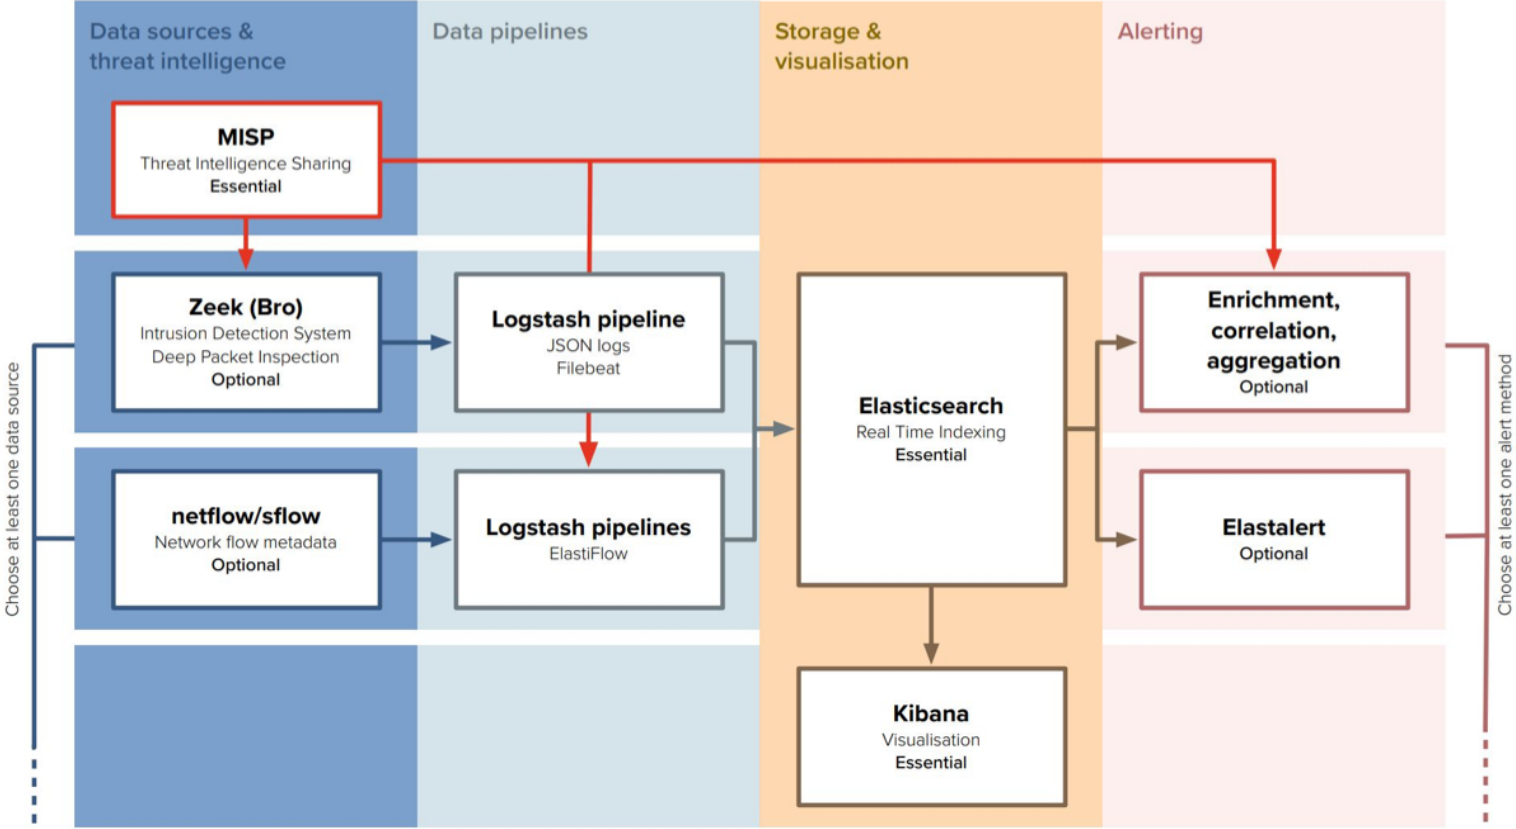
# SOC Engineering Components

- Talked about some of the key components
  - Threat intelligence
  - Fine-grained network monitoring
- Let's look at an overall structural diagram
- In the context of building a SOC process, we're now focusing on the engineering part
  - But don't forget the higher level considerations

# SOC Engineering Components

- **NOTE:** this is the reference design created by the SOC WG
  - Coordinated by WLCG but open to R&E
  - Contains necessary core elements
    - <https://wlcg-soc-wg-doc.web.cern.ch>

# SOC Engineering Components



# Data sources and threat intelligence

- Already discussed
  - MISP: threat intelligence
  - Zeek: network monitoring
  - Net/sflow: network monitoring
  - +host logs
- Start with one and grow from there

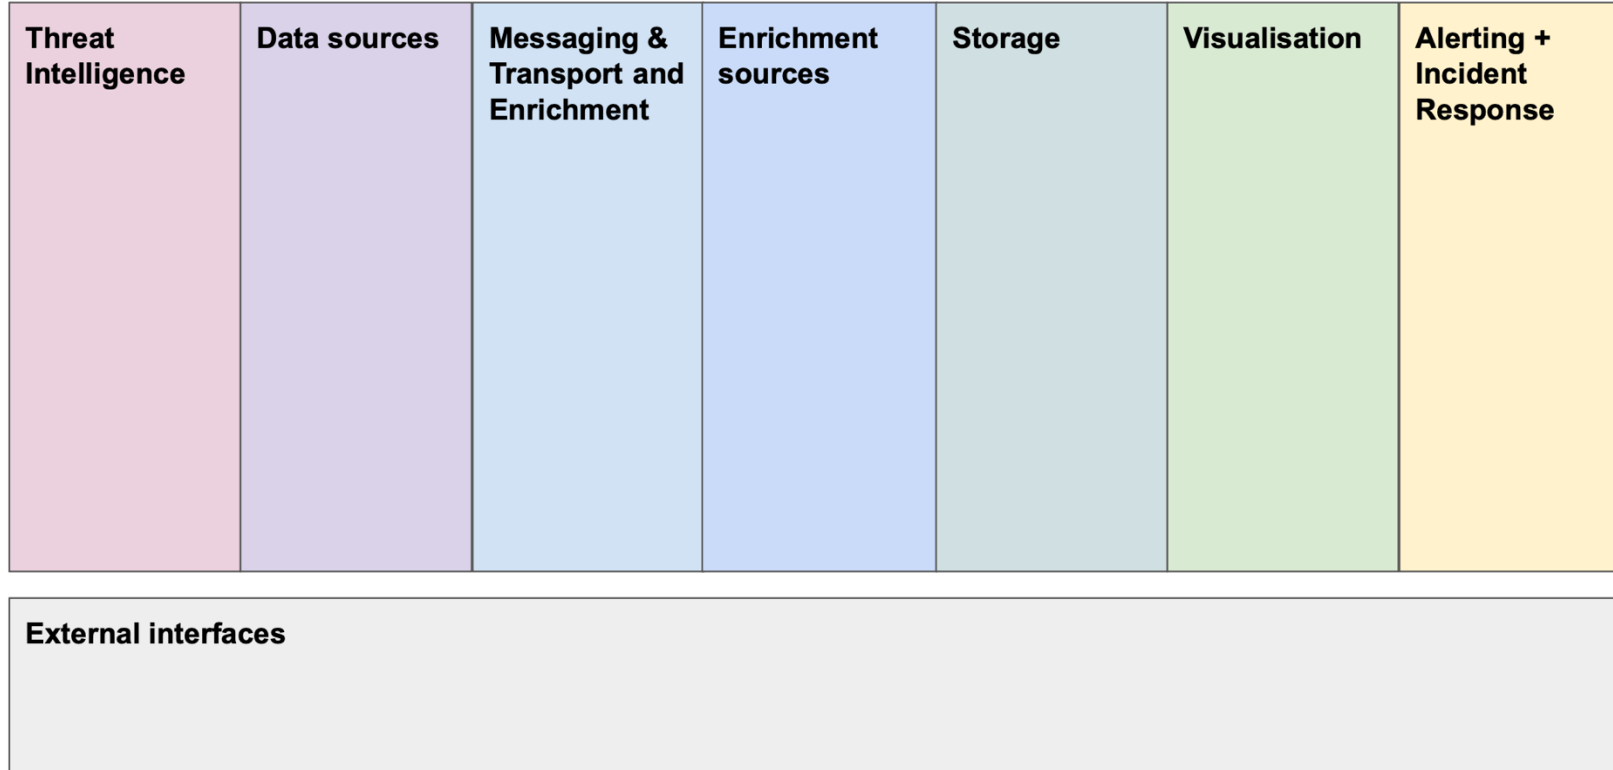
# Storage and visualisation

- OpenSearch + OpenSearch Dashboards
  - Common, well understood components
  - Open source distribution of Elasticsearch, Logstash and Kibana
  - Includes security plugins from the outset

# Alerting

- Alerting directly from Zeek
- Alerting from OpenSearch
- Aggregation of information into emails
  - In use in CERN SOC

# SOC Components

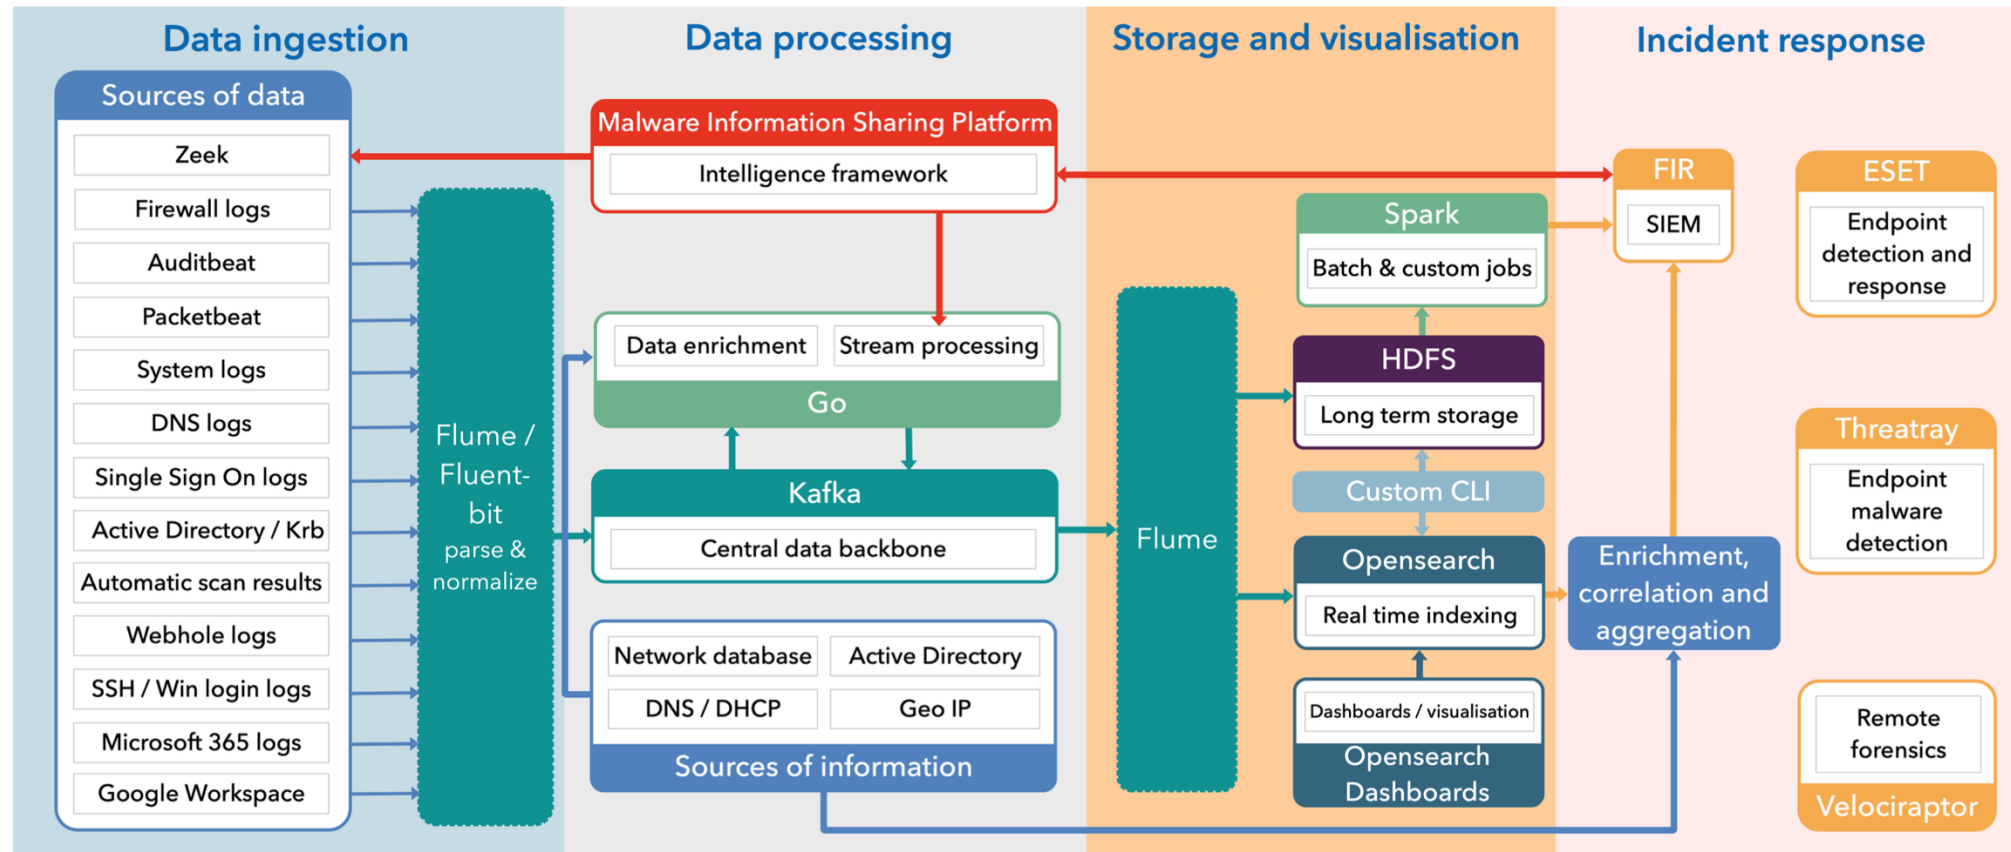


- Draft Reference Design v2


# External interfaces

- A key part of deploying SOC capabilities in our community is how they **interoperate**
- Thus, when considering options for different SOC elements, think about how they might – if they need to – interoperate with other facilities

# CERN SOC



# STFC SOC-NDR

<p><b>Threat Intelligence</b></p> <p>MISP</p> <ul style="list-style-type: none"> <li>- Specific UKRI instance</li> <li>- Pulling from CERN, SAFER-TRUST, Jisc</li> </ul> <p> JiscCTI / misp-docker</p>	<p><b>Data sources</b></p> <p>Zeek</p> <ul style="list-style-type: none"> <li>- 7.3</li> <li>- CERN-built RPMs</li> <li>- Rocky9</li> </ul>	<p><b>Messaging, Transport and Enrichment</b></p> <p>FluentBit</p> <ul style="list-style-type: none"> <li>- 3.0</li> <li>- Rocky9</li> <li>- Lua filters</li> </ul>	<p><b>Enrichment sources</b></p>	<p><b>Storage</b></p> <p>OpenSearch</p> <ul style="list-style-type: none"> <li>- 2.4*</li> <li>- Docker</li> <li>- Base OS is Rocky9</li> </ul>	<p><b>Visualisation</b></p> <p>OpenSearch Dashboards</p> <ul style="list-style-type: none"> <li>- 2.4*</li> <li>- Docker</li> <li>- Base OS is Rocky9</li> </ul> <p>* Graduate project in progress to work on upgrade paths</p>	<p><b>Alerting + Incident Response</b></p> <p><i>Zeek Alerting Framework</i></p> <ul style="list-style-type: none"> <li>- <i>ZeekJS</i></li> <li>- <i>Focus on webhooks/API calls</i></li> </ul>
---	---	---	----------------------------------	---	---	--

<p><b>External interfaces</b></p>
-----------------------------------

# Zeek specification

- When designing a zeek worker node, what are the main factors?
- Zeek works by:
  - splitting the traffic across cores
  - running a set of internal protocol analysers against each packet
  - running a set of scripts on top of these

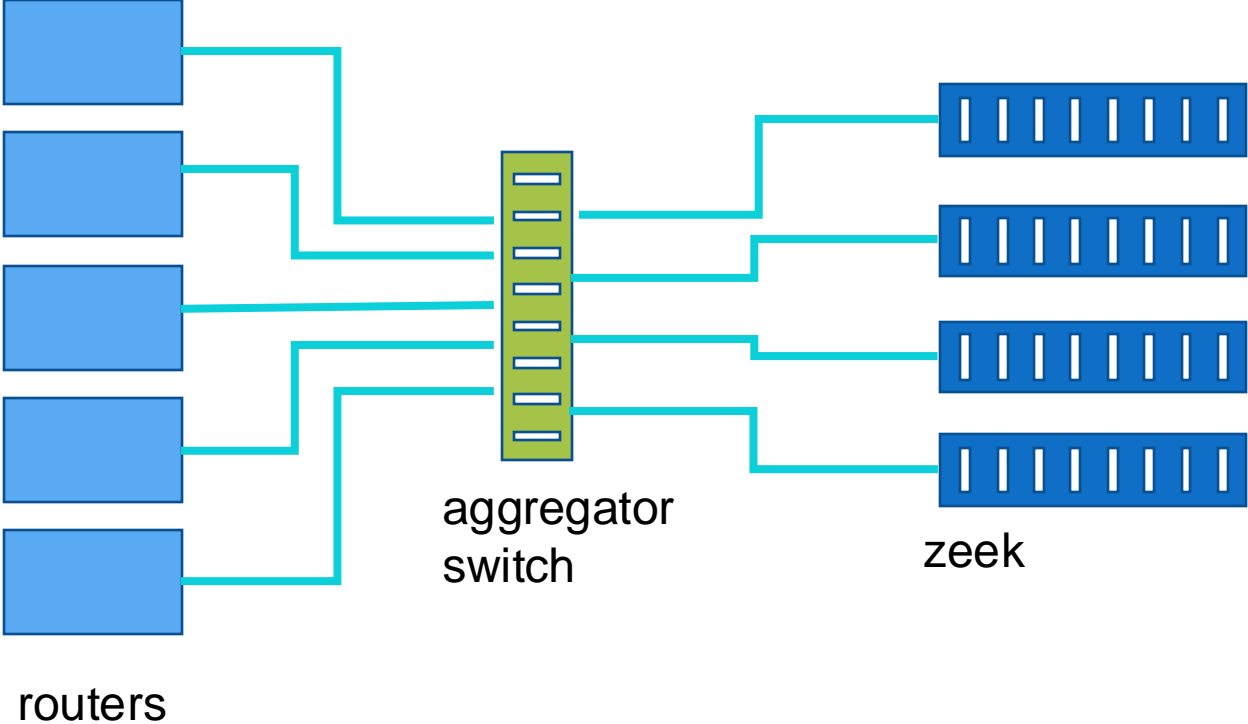
# Zeek specification + deployment

- Modern servers with ~hundreds of cores are particularly suitable for zeek processing
- A useful way of deploying the zeek farm is to consider the aggregation of network traffic

# STFC example

- 4 x 100 Gb/s links to Janet
  - General internet provided by Jisc (UK NREN)
  - In fact 8 links acting as primary/secondary
- 2 x 100 Gb/s LHCOPN links to CERN
- Need to make sure each zeek worker sees both halves of any one connection
- But can distribute these

# Traffic aggregation and load balancing



# Incident management / Alerting

- A key part of deploying a SOC is how to escalate information for operator/responder consideration
- SOC WG perspective is to leave this to local considerations
  - Different ticketing systems...

# Zeek + OpenSearch alerting

- Can alert from Zeek based on matches between MISP events and matching traffic
- Alert from OpenSearch based on outcomes of specific queries
- Key to put these alerts where they are most useful

# SOC deployment summary [1]

- We've spoken about the high level concept of a SOC
  - People
  - Processes
  - Technology

# SOC deployment summary

- We've spoken about some specifics around engineering
- Tomorrow in the workshop we will have a chance to look at some of these hands-on

# Conclusions: Detection

- In this two lecture block we've looked at
  - The basics of logging, and logging technologies
  - The importance of identifying the most useful logs to avoid "data as noise"
  - The vital role that central logging plays
  - The difference between flow based and deep packet inspection network monitoring

# Conclusions: Detection

- We've also discussed
  - The importance of sharing threat intelligence for our community
  - Tools to help share intelligence responsibly
  - The MISP platform
  - Components of a SOC
    - People, Process, Technology
- Last part of Detection track: containerised SOC deployment workshop