

Intrusion detection with SOCs: Part 1

threat intelligence,
monitoring,
integration and
processes

David Crooks

UKRI STFC

EGI CSIRT/IRIS Security team

david.crooks@stfc.ac.uk




Introduction

- We need to match our monitoring capabilities and methodology to our circumstance
- Following our architecture: what are our threats, how do we defend ourselves?

Landscape

Landscape: the world has changed

- In the past, biggest risk for academic security
 - Relatively simple, untargeted attacks
 - Belief that research computing was major risk
- This is no longer the case
 - Determined, well-resourced attackers
 - **9-5 jobs** working on malware services
 - Phishing and identity theft are major risk
 - Research computing security can be **major asset**
- Big business: we are targets



© 2021 AO Kaspersky Lab. All Rights Reserved

kaspersky

Impact

- In our community (research and education) we are faced by determined attackers
- The impact of successful attacks can be **catastrophic**
- **Months** of site/facility downtime
- Major reputational and financial damage

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 (TLP)

RED	Not for disclosure, restricted to participants only.	Recipients may not share TLP:RED information with any parties outside of the specific exchange, meeting, or conversation in which it was originally disclosed. In the context of a meeting, for example, TLP:RED information is limited to those present at the meeting.
AMBER	Limited disclosure, restricted to participants' organisations.	Recipients may only share TLP:AMBER information with members of their own organisation, and with clients or customers who need to know the information to protect themselves or prevent further harm. Sources are at liberty to specify additional intended limits of the sharing: these must be adhered to.
GREEN	Limited disclosure, restricted to the community.	Recipients may share TLP:GREEN information with peers and partner organisations within their sector or community, but not via publicly accessible channels
WHITE	Disclosure is not limited	Subject to standard copyright rules, TLP:WHITE information may be distributed without restriction.

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

TLP Examples

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

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



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**
 - Under what 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

Malicious activities

Event ID: 10878
 Uuid: 5aec700c-0ebb-46b-...
 Org: CIRCL
 Owner org: CIRCL
 Contributors: alexandre.dulaunoy
 Tags: [empty]
 Date: 2018-05-04
 Threat Level: Low
 Analysis: Initial
 Distribution: All communities
 Info: Malicious activities
 Published: No
 #Attributes: 2
 Last change: 2018/05/04 02:38:11
 Extends: [empty]
 Extended by: [empty]
 Sightings: 0 (0)
 Activity: [empty]

Distribution graph [atomic event]

All
 Attributes
 Object attributes

Event Graph

Threat Level: Low
 Analysis: Initial
 Event Info: Ransomware found on a production server
 Extends event: 5ad8687b-0e10-4a8b-a157-46a5950d210f
 Matched event: Id: 10728, Analysis: Completed, Threat level: Low
 Tags: CIRCL.osint.feed, ip-white, malware_classification:malware-category-"Ransomware", coint.source-type-"blog-post", misp-galaxy:ransomware-"CS:GO Ransomware", misp-galaxy:ransomware-"MC Ransomware"
 Info: OSINT - Microsoft & CS:GO Ransomware Shine For Media Attention

estimative-language:confidence-in-analytic-judgment="high"
 High

View Dashboard

- Add Widget
- Import Config JSON
- Export Config JSON
- Save Dashboard Config
- List Dashboard Templates

Authentication Failure Data

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

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!

- Congratulations, you have shared your first event!
- You have been using tags, good job!
- Taxonomies have been used in your events.
- 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

MISP Demo

MISP

- We'll look more at this including in the workshop
- Training and documentation is also [available](#)

R&E threat intelligence + EGI CSIRT

- 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

- 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

Security Operations Centres

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

Security Operations Centres

- From a high level, however, a SOC is the combination of
 - Technology
 - People
 - Processes
- Developing the team that uses a SOC and develops good information from it

Security Operations Centres

- Developing the processes by which you disseminate and coordinate the alerting from the SOC
- Are **equally important to the tooling**

Teams and Processes

- Who maintains the SOC?
 - Next year?
- Where does the next tranche of hardware come from?
- Who analyses the alerts?
- ...

How to deploy a SOC

- First question: what is the scope?
 - Individual batch farm?
 - Single site organisation?
 - Multi-site organisation?
 - Country?
- What considerations might come into play?
 - Effectiveness of intelligence
 - Network logging

How to deploy a SOC

Example	Deployment
Individual batch farm	Not clear that intelligence will be most useful
Single-site organisation	Identify network choke points
Multi-site organisation	How do we ship data around?
Country	Can't use DPI for the backbone of a country

How to deploy a SOC

- 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!

Specific questions

- Where are the network choke points that are most relevant?


Example

- STFC is a multi-site organization and we are deploying a SOC against the RAL campus

Specific questions

Example


- STFC is a multi-site organization and we are deploying a SOC against the RAL campus



Specific questions

Example

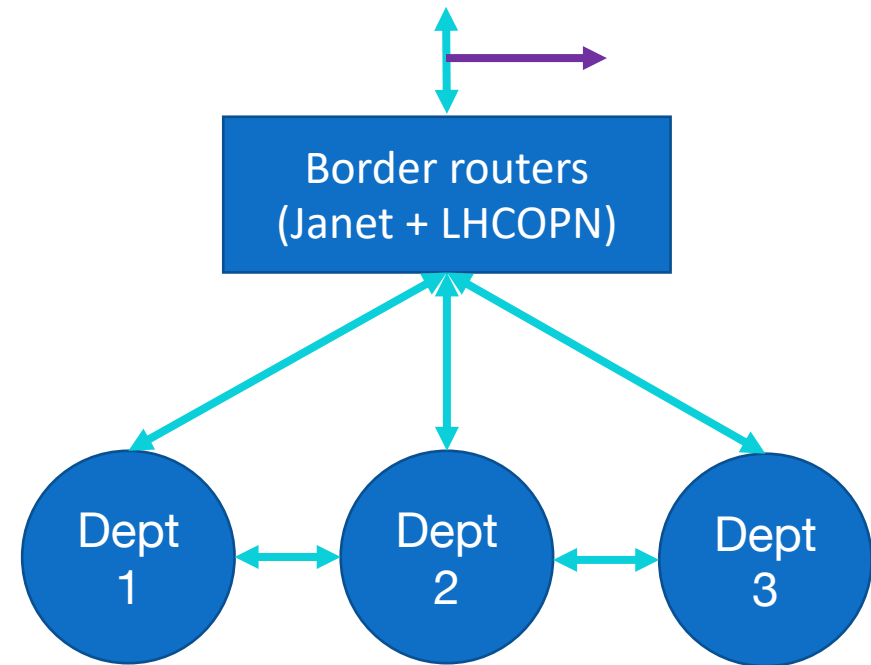
- Where do we put the network tap?



Specific questions

Example


- Where do we put the network tap?



Specific questions

Example

- North South traffic
 - Into and out of a site
- East West traffic
 - Traffic within a site



Specific questions

Example

- STFC Multi-site
 - What could our approach be?



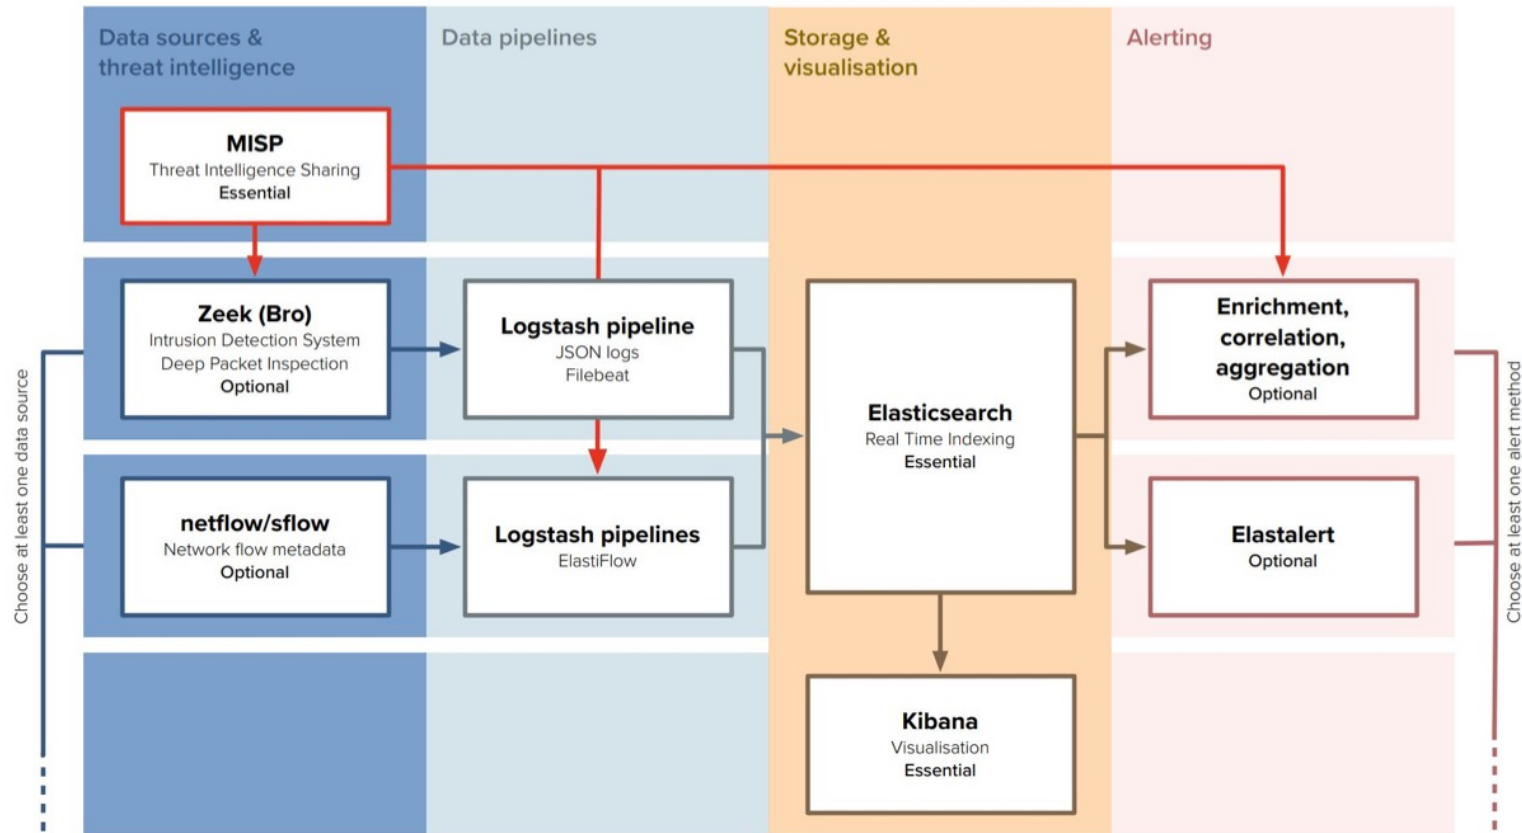
SOC Components

- Talked about some of the key components
 - Threat intelligence
 - Fine-grained network monitoring
- Let's look at an overall structural diagram

SOC Components

- **NOTE:** this is the reference design created by the SOC WG
 - Coordinated by WLCG but open to R&E
 - Not the only way of going forward
 - Contains the necessary core elements

SOC 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

Data pipelines

- Logstash works as part of the standard elastic stack
 - Starting point

BUT

- Is typically not performant enough at high load
 - Kafka, ...

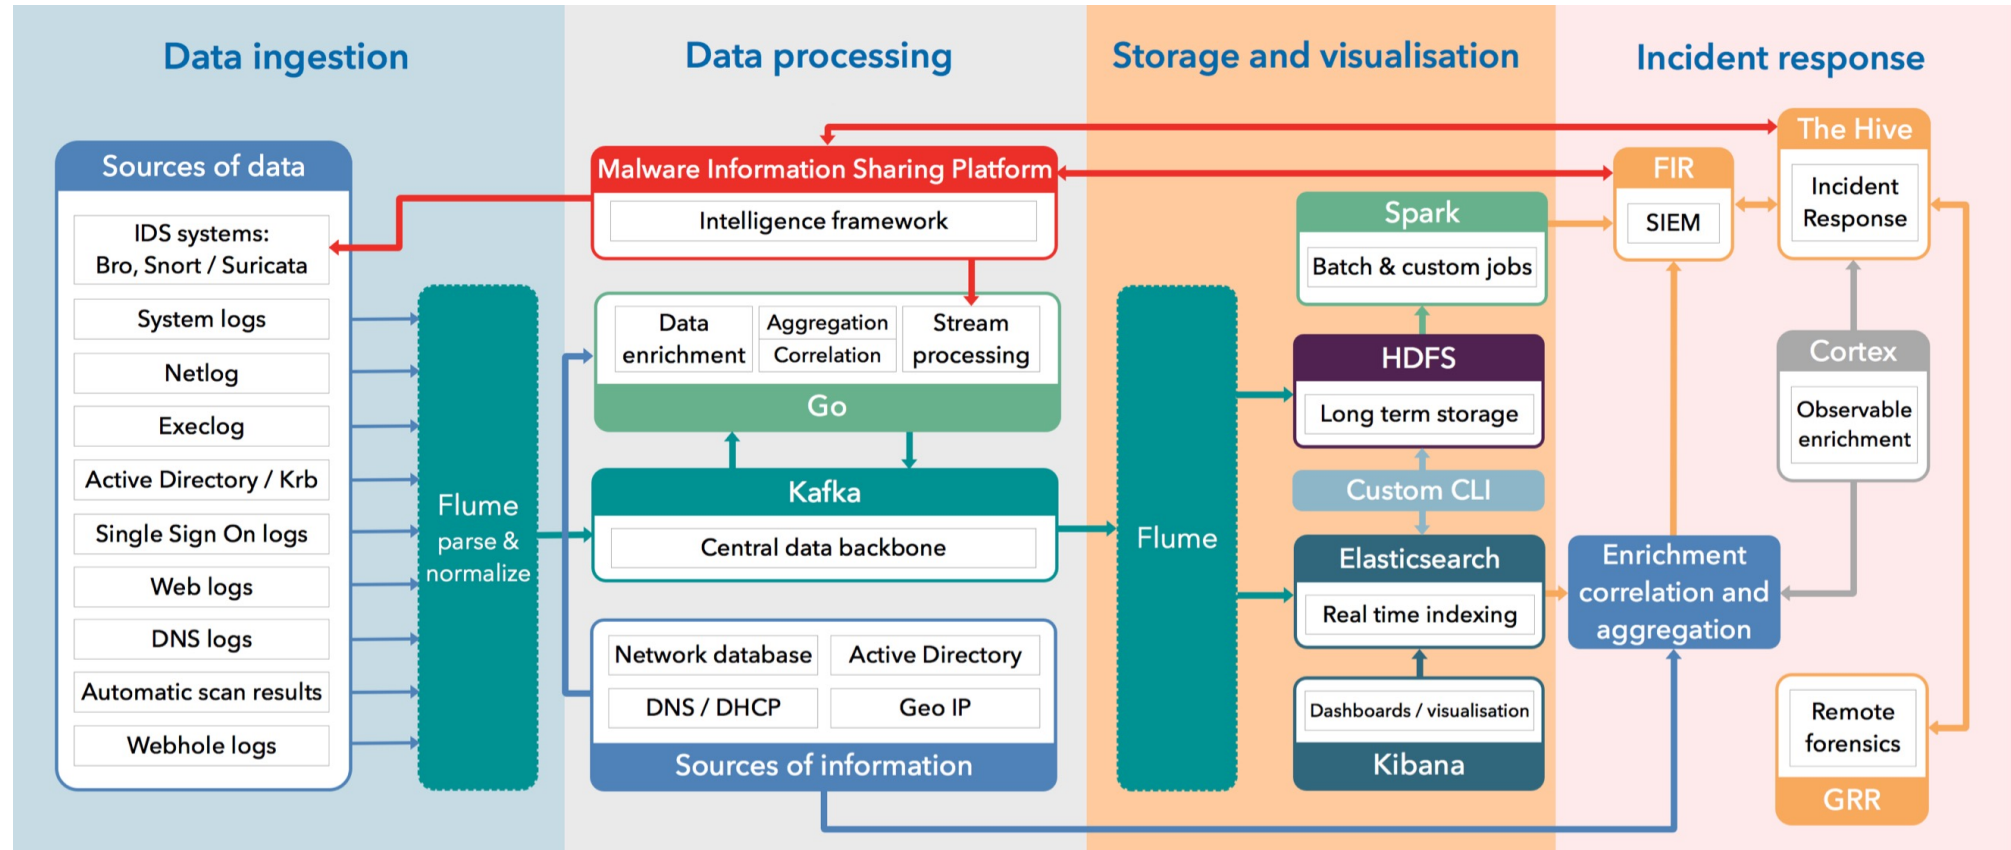
Storage and visualisation

- Elasticsearch + Kibana
 - Common, well understood components
- CERN has Elasticsearch service
- OpenSearch is a useful new distribution
 - Includes security plugins from the outset

Alerting

- Alerting directly from Zeek
(see this during the week)
- Alerting from elasticsearch
- Aggregation of information
into emails
 - In use in CERN SOC

CERN SOC



Conclusions

- The implication of our current security landscape is that we have to work together to defend our community
- We can do this by making sure our teams have
 - High quality shared intelligence
 - The tools to use this with

Conclusions

- There are existing methods for classifying how to share information
- We can use these in the sharing of threat intelligence
- We have considered some aspects of how to deploy a SOC