Fedcloud Security

Towards comprehensive Cloud SOC capabilities

David Crooks

EGI Conference 2023, Poznan, Poland

Contents

1. Elements of cybersecurity

   Identification; Prevention; Detection; Response

2. Laying the groundwork

   Identification + Prevention

3. What monitoring capabilities do we need?

   Detection

4. How do we respond?

   Response

Elements of Cybersecurity

Elements of Cybersecurity

- Identification
- Prevention
- Detection
- Response + Recovery

- May find these familiar from various frameworks (NIST/etc), useful way of breaking down the work to do

Identification

- What do we have; how do we structure our processes?

- Asset management
- Risk management
- Governance

Prevention

- What safeguards can we put in place to protect our systems from attack?

- Security controls
  - Both technical (eg firewalls) and management (policy)
- Architecture
- Training

Detection

- What monitoring and telemetry do we have access to to detect suspicious traffic in our environment?

- Network monitoring/IDS
- Central logging
- Threat Intelligence
- Security Operations Centres

Response + Recovery

- How do we respond in the case of an incident?
- How do we recover from an incident?

- Security team/CSIRT structures
- Incident response procedures
- Exercises

- Recovery procedures

  including final incident reports
- Communications
- Continuous improvement

  At the core of cybersecurity

Laying the groundwork

Identification

- *Discussion*

- What asset management systems are in use in the FedCloud?
- What risk management systems are in use in the FedCloud?

Prevention

- *Discussion*

- How do you manage patching/replacing unsupported components?

Detection

What monitoring capabilities do we need?

- Detection of suspicious activity requires instrumenting our environments
  - Logs

    central security and audit logs; user activity
  - Hosts

    Hypervisors/control plane vs cloud VMs
  - Network

    Where do we need to monitor?
  - DNS

    Monitoring without deep packet inspection

Detailed network monitoring

- The composed application layers in a cloud environment make fine-grained network monitoring particularly relevant
- Even with encrypted traffic, gives detailed picture of environment
- Coupled with community threat intelligence, allows real time correlation of the state of your estate with current threats

Threat Intelligence

- EGI CSIRT is now updating its procedures to make threat intelligence available throughout the investigation of incidents
- This will be available to anyone in our community via REST API with appropriate auth key
- This can then be integrated into scripts or a full-sized Security Operations Centre

Passive monitoring

- May typically see very high data throughput which is not consistent with the use of a perimeter firewall in all cases
  - And; we do not wish to alert the attacker that we are watching their progress
- Passive tapping of traffic, by optical taps, span ports or packet brokers, allows a clean copy of the traffic to be analysed without disrupting the connections or alerting the attackers
- We can then use tools like Zeek (comprehensive network logging) and Suricata (signature-based detection) to analyse the traffic
  - Focus here on Zeek

Summary so far

- So: we have
  - Threat intel via a REST API
  - A clean copy of the network traffic via network tap
  - A method of analysing this traffic

    Focusing on the Zeek network IDS

- How do we proceed for Cloud sites?

Location of network taps

- Let’s consider where best to locate our network tap
- Tapping North-South traffic gives optimal view of traffic to external internet
- If your cloud facility has its own direct internet links, this would be the place to start
  - *Discussion: is there anyone for whom this is true?*

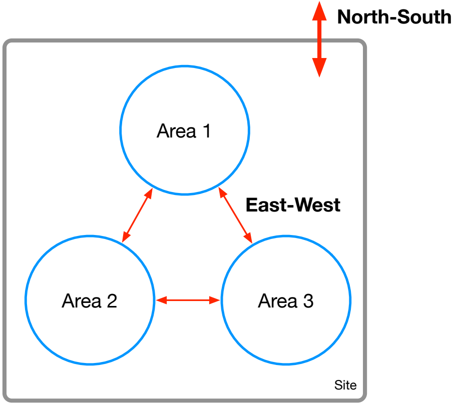

Internal/East-West traffic

- How should we approach internal traffic for a cloud facility deployed within a larger network?
- If the cloud is located within a particular network segment, tapping the uplink from that segment would be a good possibility
- What if there is no appropriate higher level tap point?

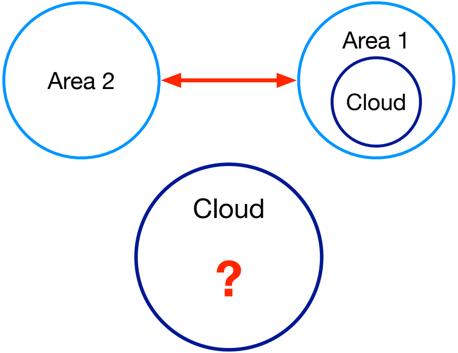

CloudSOC concept

- Zeek scales very well with bandwidth
  - Monitor a single 1Gb/s link with a RPi
- Can certainly run Zeek within a VM/container
  - Although tuning would certainly be required

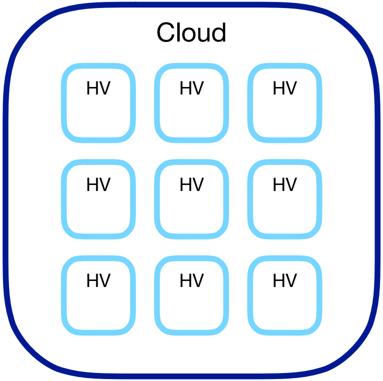

CloudSOC concept: HV Zeek

- Zeek scales very well with bandwidth
  - Monitor a single 1Gb/s link with a RPi
- Can certainly run Zeek within a VM/container
  - Although tuning would certainly be required
- A feasible technical option is to run Zeek on every Hypervisor

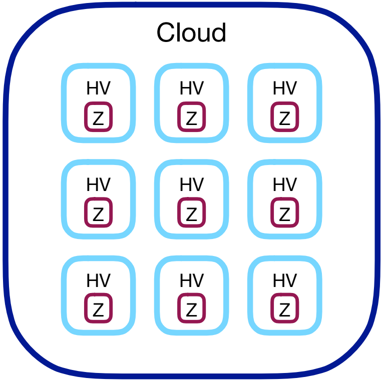

CloudSOC concept: Orchestration

- *How would we orchestrate hundreds of zeek sensors?*

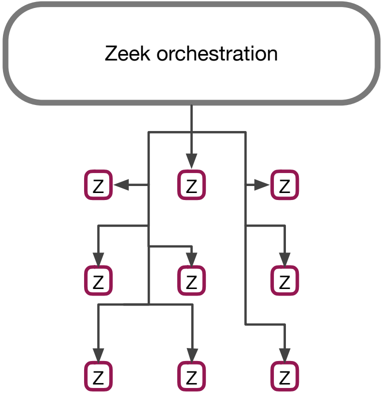

CloudSOC concept: Aggregation

- *How would we aggregate the data from hundreds of zeek sensors?*

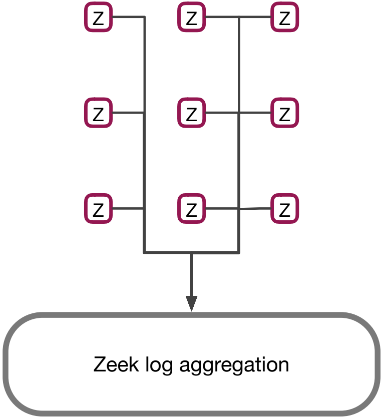

(passive) DNS

- Full scale network monitoring using Zeek is incredibly informative, but requires careful engineering
- In parallel, useful to consider the resolution of malicious domains that can be tracked using DNS results
- A new software component developed in the CERN Computer Security Team

- **pDNSSOC**

{p,D}DNS disambiguation

- P(rotective)DNS: A recursive DNS resolver that prevents malicious domains from being resolved
- P(ower)DNS: A supplier of open-source and commercial DNS software
- p(assive)DNS: A specific, anonymous subset of historical DNS resolution data including only: Domain name; Record type; Record value; Time stamp

{p,D}DNS disambiguation

- P(rotective)DNS: A recursive DNS resolver that prevents malicious domains from being resolved
- P(ower)DNS: A supplier of open-source and commercial DNS software
- **p(assive)DNS: A specific, anonymous subset of historical DNS resolution data including only: Domain name; Record type; Record value; Time stamp**

pDNSSOC

- PDNSSOC is a software component that correlates DNS logs with threat intel from MISP as an “80%” SOC
  - provides a turn-key solution to detect and respond to security incidents

- Allows for flexible deployment configurations
- Allows for rapid horizontal scaling across many sites

- Requires source of (passive) DNS resolution data

pDNSSOC architecture

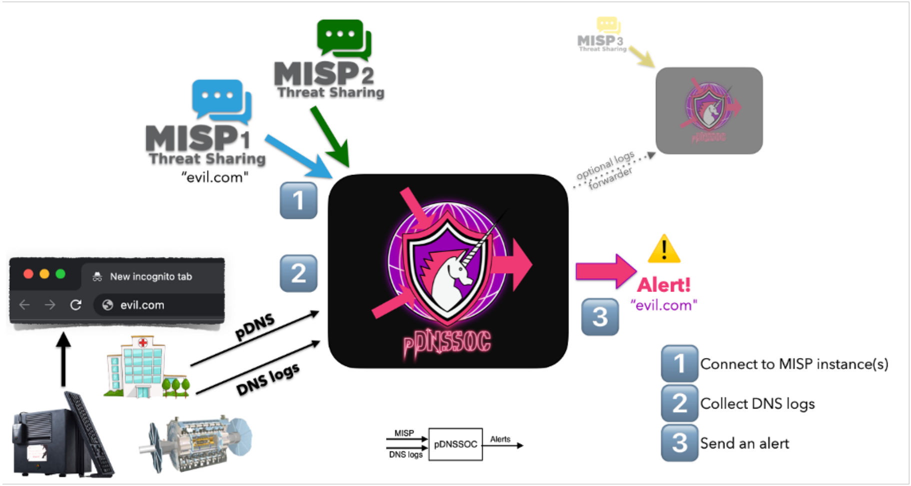

1 Connect to MISP instance(s)

2 Collect DNS logs

3 Send an alert

pDNSSOC Deployment: Federation

- An organisation forwards pDNS data using a pDNSSOC forwarder
- Can detect the intrusion at different levels while respecting TLP

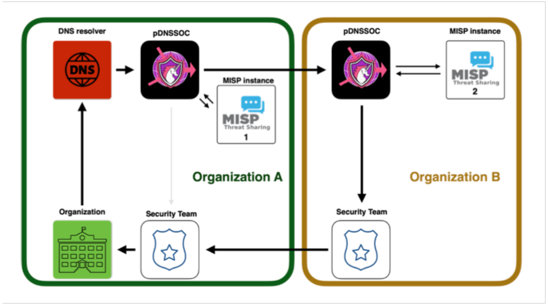

pDNSSOC Deployment: Collaboration

- An organisation forwards DNS/pDNS logs
- They cannot block the requests but gets alerts

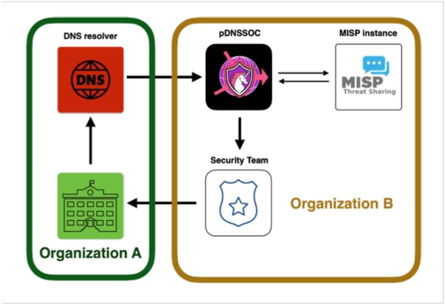

pDNSSOC Deployment: Responsive

- The organisation uses a remote DNS resolver.
- At this remote organisation, DNS + RPZ is used to block malicious domains from being resolved
  - pDNSSOC is used to generate alerts

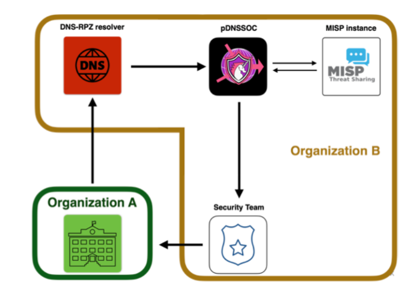

Response

Responding to incidents

- Important to understand how to respond in case of an incident

- Identify local security team(s)
- EGI SEC01

Discussion

Current capabilities

- *For cloud sites that already have network monitoring capabilities*

- Where in the network are you monitoring?
- What tools are you using?

Network topologies

- *For cloud sites that don't already have network monitoring capabilities*

- What does your network topology look like?
- Do you have a useful tap point outside your cloud?
- Looking at the internal cloud network topology, is trying to define internal virtual tap points worthwhile?

Next steps

- Deploying Zeek on every hypervisor, while feasible, is a considerable technical challenge
- pDNSSOC development is at the stage of looking for sites to test deployment: would cloud sites be interested?

SOC Working Group

- SOC Working Group focuses on building reference designs for Security Operations Centres
- Co-chaired by David C and Liviu Vâlsan at CERN

- Active Keybase community; to participate create an account at keybase.io and talk to David

SOC Hackathon

- 5 day technical + strategic meeting to work on technology used in Security Operations Centre deployment
  - Including the topics covered here today
- Could provide focus point to consider deployment of CloudSOC as well as pDNSSOC topologies
- https://indico.cern.ch/event/1268239/

Questions?